
PeopleTools 8.60: Search Technology

March 2023

PeopleTools 8.60: Search Technology
Copyright © 1988, 2023, 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, software documentation, data (as defined in the Federal Acquisition Regulation), 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," "commercial computer software documentation," or "limited rights data" 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®, Java, and MySQL 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.

Documentation Accessibility

For information about Oracle's commitment to accessibility, visit the Oracle Accessibility Program website at <https://docs.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 <https://docs.oracle.com/pls/topic/lookup?ctx=acc&id=info> or visit <https://docs.oracle.com/pls/topic/lookup?ctx=acc&id=trs> if you are hearing impaired.

Contents

- Preface: Preface.....ix**
 - Understanding the PeopleSoft Online Help and PeopleBooks..... ix
 - Hosted PeopleSoft Online Help..... ix
 - Locally Installed Help..... ix
 - Downloadable PeopleBook PDF Files.....ix
 - Common Help Documentation.....ix
 - Field and Control Definitions..... x
 - Typographical Conventions..... x
 - ISO Country and Currency Codes..... x
 - Region and Industry Identifiers..... xi
 - Translations and Embedded Help..... xi
 - Using and Managing the PeopleSoft Online Help..... xii
 - PeopleTools Related Links..... xii
 - Contact Us..... xii
 - Follow Us..... xii
- Chapter 1: Getting Started with PeopleSoft Search Technology..... 15**
 - PeopleSoft Search Technology Overview..... 15
 - Implementing PeopleSoft Search Using Elasticsearch..... 15
- Chapter 2: Understanding PeopleSoft Search Framework Integration with Elasticsearch..... 17**
 - PeopleSoft Search Framework Features..... 17
 - Search Framework Definitions..... 18
 - Search Documents..... 18
 - Elasticsearch Concepts and Terminology..... 19
 - Elasticsearch Support Based on PeopleSoft Application Version..... 20
 - PeopleSoft Search Framework Supported Features in Elasticsearch..... 20
 - PeopleSoft Search Framework Architecture..... 23
 - Implementation Process Flow..... 28
 - PeopleSoft Search Features..... 30
 - Global Search..... 30
 - Search Pages..... 30
- Chapter 3: Working with Elasticsearch Clusters..... 33**
 - Understanding Clusters in Elasticsearch..... 33
 - Setting Up a Cluster During the Elasticsearch Installation..... 34
 - Naming the Cluster..... 34
 - Minimum Number of Master Nodes..... 34
 - List of Host Names where Elasticsearch will be Installed or is Installed..... 35
 - High-level Steps to Set Up a Cluster..... 35
 - Upgrading Existing Cluster Using ELK DPK..... 36
 - Adding a New Node to an Existing Cluster..... 36
 - Verifying the Cluster and its Nodes..... 37
 - Using Elasticsearch Interact Page..... 37
 - Setting the Number of Replicas..... 38
 - Starting and Stopping a Cluster..... 38
 - Starting and Stopping a Node..... 39
 - Removing a Node from a Cluster..... 40
 - Managing Large Indexes..... 41

Chapter 4: Defining Search Definition Queries.....	43
Understanding Search Definition Queries.....	43
Understanding the Search Designer Activity Guide.....	44
Defining Search Definition Queries with PeopleSoft Query.....	44
Adding a “Last Updated” Field to Records.....	44
Specifying a List of Fields to Index.....	45
Creating a Drilling URL.....	45
Creating a Prompt for the Last Modified Field.....	47
Defining Prompt Criteria.....	48
Defining Search Definition Queries with Connected Query.....	49
Defining a Deletion Query.....	50
Working with Images.....	50
Testing Your Search Definition Query.....	51
Chapter 5: Creating Query and Connected Query Search Definitions.....	53
Creating Search Definitions.....	53
Specifying General Settings.....	53
Adding Text and Bind Variables for Title and Summary.....	56
Mapping Search Attributes.....	58
Working with Attachment Properties.....	62
Understanding Attachment Processing in Elasticsearch.....	62
Specifying Attachments in a Search Definition.....	63
Specifying Attachment URL Properties.....	64
Specifying Static URL ID Attachment Properties.....	65
Specifying Dynamic URL ID Attachment Properties.....	67
Specifying Inline Attachment Properties.....	68
Managing Search Definition Security.....	68
Setting Source Level Security.....	69
Setting Document Level Security.....	70
Working With Advanced Settings.....	72
Mapping Components to Search Definitions.....	77
Viewing Search Attributes.....	78
Chapter 6: Creating File Source Search Definitions.....	81
Understanding File Source Search Definitions.....	81
Specifying File Source General Settings.....	81
Specifying Document Types.....	84
Chapter 7: Creating Web Source Search Definitions.....	89
Understanding Web Source Search Definitions.....	89
Specifying Web Source General Settings.....	89
Specifying Document Types for Web Source.....	91
Specifying Proxy Parameters in Elasticsearch.....	95
Chapter 8: Creating Activity Guide Search Definitions.....	97
Understanding Activity Guide Search Definitions.....	97
Specifying Activity Guide General Settings.....	97
Specifying Security Settings.....	102
Setting Source Level Security.....	103
Setting Document Level Security.....	104
Specifying Advanced Properties.....	106
Chapter 9: Creating Search Categories.....	109
Understanding Search Categories.....	109
Specifying General Search Category Settings.....	109
Selecting Advanced Search Field Settings.....	111

Selecting Facet Settings.....	112
Creating Associated Facets.....	114
Defining Advanced Settings for String Facets.....	115
Defining Advanced Settings for Number Facets.....	117
Defining Advanced Settings for Date Facets.....	120
Identifying Custom Search Pages.....	125
Working With Display Fields.....	126
Chapter 10: Administering PeopleSoft Search Framework.....	129
Understanding PeopleSoft Search Framework Administration.....	129
Understanding the Search Admin Activity Guide.....	129
Setting Up Local Nodes.....	129
Specifying Integration Broker Gateway.....	130
Setting Up Target Locations.....	130
Working With Search Instances.....	130
Creating Search Instances.....	130
Modifying Search Instances.....	138
Deleting Search Instances.....	139
Sharing a Search Instance Among Multiple PeopleSoft Application Systems.....	139
Administering Elasticsearch Using the Elasticsearch Interact Page.....	139
Setting Namespace Aliases.....	143
Administering Search Instances.....	145
Assigning Users to a Search Instance for Validating Indexes.....	146
Creating a Search Instance in Automated Configuration Manager.....	147
Enabling Global Search.....	148
Managing General Search Options.....	148
Administering Search Definitions and Search Categories.....	157
Understanding Search Definition Administration.....	157
Working with Search Definition Deployment.....	157
Working with Search Categories.....	165
Working with Search Indexes.....	167
Understanding the Index Build Process.....	167
Building Search Indexes.....	169
Viewing Search Index Build Process Details.....	171
Viewing the Exception Details in Attachment Processing.....	173
Partitioning Application Data in the Search Index Build Process.....	174
Using Direct Transfer.....	176
Understanding Direct Transfer.....	176
Configuration.....	177
System Considerations.....	178
Error Handling.....	179
Using Direct Transfer with SSL Enabled Elasticsearch.....	181
Real-time Indexing Using Direct Transfer Technology.....	182
Verifying Domain Status.....	183
Verifying Integration Broker Queue Status.....	183
Managing Search Context.....	183
Understanding Search Contexts.....	183
Defining Search Contexts.....	184
Viewing Search Contexts.....	184
Viewing Node in Network.....	186
Importing Remote Search Groups.....	186
Creating the Attachment URL ID List.....	188

Implementing Report Repository Search.....	189
Working with Search Framework Definitions During Upgrades.....	190
Boosting Score of Search Results.....	190
Testing Score Boosting.....	192
Chapter 11: Administering Real Time Indexing.....	193
Understanding Real Time Indexing.....	193
Understanding Set Based Processing in Real Time Indexing.....	193
Configuring Real Time Indexing.....	195
Prerequisites for Real Time Indexing.....	195
Steps to Enable Real Time Indexing for a Search Definition.....	197
Using the Configure Real Time Indexing Page.....	197
Turning Off Real Time Indexing During Batch Processing.....	203
Using an Application Engine Maintenance Program.....	205
Excluding Search Definitions from Real Time Indexing.....	206
Removing a Trigger.....	206
Viewing Error Details in the Real Time Indexing Process.....	207
Migrating Real Time Indexing Data Using ADS Definition.....	208
Troubleshooting.....	209
Chapter 12: Working with PeopleSoft Search Framework Utilities.....	211
Using the Search Test Page.....	211
Understanding the Search Test Page.....	211
Testing Search Categories.....	213
Testing Custom Search Attributes.....	213
Testing Filter Settings.....	213
Testing Facet Requests.....	214
Testing Grouping and Sorting Options.....	215
Testing Additional Parameters.....	215
Downloading Search Data.....	215
Running Diagnostics.....	216
Viewing Event Logs.....	218
Chapter 13: Working with PeopleSoft Search Framework Security Features.....	221
Understanding Search Framework Security.....	221
Applying PeopleSoft Permissions.....	221
Working with Authentication and Authorization.....	222
Authentication and Authorization in Elasticsearch.....	223
Configuring SSL between PeopleSoft and Elasticsearch.....	224
Configuring SSL for Elasticsearch.....	224
Configuring SSL for Kibana.....	226
Configuring SSL for Logstash.....	228
Setting Up Role-Based Search Group Access.....	229
Chapter 14: Managing Configurable Search.....	231
Understanding Configurable Search.....	231
Search Pages Delivered for Configurable Search.....	231
Considerations for Using Configurable Search.....	232
Guidelines on Using Configurable Search for Classic Components.....	233
Guidelines on Using Configurable Search for Fluid Components.....	233
Creating Custom Search Pages.....	235
Working with Configurable Search.....	235
Specifying System-wide Configurable Search Properties.....	235
Managing Configurable Search Configurations.....	237
Configuring a Search Page.....	238

Configuring an Add a New Value Page.....	247
Upgrade Considerations.....	251
Migrating Search Page Configuration Data Using Application Data Sets.....	252
Chapter 15: Working with PeopleSoft Search.....	255
Understanding PeopleSoft Search.....	255
Working with Search Pages.....	256
Understanding Search Pages.....	256
Enabling Keyword Search.....	258
Developing for Search Pages.....	259
Working with Keyword Search Modes.....	262
Working with Global Search.....	263
Understanding Global Search.....	264
Working with Search Groups and Search Contexts.....	264
Working With the Portal Registry Search.....	265
Working with Suggestions Search.....	267
Testing Suggestions Search.....	270
Working with Synonyms.....	271
Adding and Copying Synonyms.....	272
Clearing Synonyms from a Synonyms File.....	274
Disabling Synonym Search for an Index.....	275
Migrating Synonyms Data Using ADS Definition.....	276
Working with Search Operators in PeopleSoft Search.....	277
Working with the Search Results.....	280
Working With Facets.....	281
Viewing Search Results With Grid Format and List Format.....	285
Working with Related Actions.....	286
Working with Search Results in Fluid User Interface.....	287
Working with Search Results in Global Search.....	287
Working with Real-time Component Search Results.....	292
Working with Keyword Component Search Results.....	295
Working with Search Results on a Touch Friendly Device.....	296
Masking of Data in Search Results.....	298
Masking of Data in Search Results for Global Search.....	299
Chapter 16: Monitoring PeopleSoft Search Framework and Elasticsearch Using Kibana.....	301
Understanding Monitoring of Search Framework and Elasticsearch.....	301
Kibana Terminology.....	301
Understanding Dashboards and Visualizations.....	302
Common Elements on a Visualization.....	302
Setting Up Kibana in PeopleSoft Search Framework.....	303
Installing Kibana.....	304
Connecting Kibana to Elasticsearch.....	304
Starting the Kibana Service.....	305
Configuring the Monitoring Server.....	305
Monitoring Elasticsearch System Metrics and Indexing Metrics.....	308
Using the Monitoring Visualizations Page.....	309
Delivered Dashboards and Visualizations.....	310
Using the Indexing Metrics Dashboard and Visualizations.....	311
Using the Indexing Summary Dashboard and Visualizations.....	313
Using the System Metrics Dashboard and Visualization.....	314
Creating a Monitoring Visualization.....	315
Chapter 17: Working with Kibana Dashboards for PeopleSoft Application Data.....	317

Understanding Application Data and Kibana Dashboards.....	317
Accessing Kibana From PeopleSoft.....	318
Accessing Kibana Using the Kibana Menu Under Reporting Tools.....	319
Managing a Delivered Dashboard.....	320
Deploying a Dashboard.....	320
Assigning Roles for a Dashboard.....	321
Working with Field Alias.....	322
Understanding Field Alias.....	322
Understanding the Impact of Field Alias on Existing Dashboards.....	322
Recreating Existing Visualizations.....	323
Masking of Data in Kibana.....	323
Creating a Visualization for Application Data.....	324
Creating a Visualization.....	324
Importing a Dashboard.....	326
Configuring a Dashboard as a Tile or Related Information.....	327
Specifying User Privileges.....	332
Using a Dashboard.....	334
Using a Dashboard that is Configured as a Tile.....	334
Using a Dashboard that is Configured as Related Information of a Component.....	335
Copying Kibana Dashboards Using ADS Definitions.....	336
ADS Definitions for Elasticsearch.....	337
Records Used in ADS Definitions.....	337
Chapter 18: Configuring High Availability in PeopleSoft Search Framework.....	339
Understanding High Availability Feature in PeopleSoft Search Framework.....	339
Fail-over Mechanism.....	339
Implementing Fail-over Mechanism.....	339
Implementing Fail-over Mechanism Using Multiple Nodes Feature of Elasticsearch.....	340
Implementing Fail-over Mechanism Using PeopleSoft Search Framework's Multiple Search Instance Feature.....	340
Chapter 19: Integrating External Data with PeopleSoft.....	343
Understanding the Integration of External Data with PeopleSoft.....	343
Setting Up the Integration of External Data with PeopleSoft.....	343
Creating a Search Definition for External Data.....	344
Deploying an External Data Search Definition to Elasticsearch.....	346
Specifying a Data Source for External Data Search Definition.....	347
Building an Index for an External Data Search Definition.....	356
Chapter 20: Troubleshooting.....	359
Creating Users and Assigning Roles in Elasticsearch.....	359
Managing Elasticsearch Services.....	360
Managing an Elasticsearch Service on Linux.....	360
Troubleshooting Elasticsearch Configuration.....	362
Performing Administrative Tasks.....	362
Handling Common Errors.....	363
Managing Kibana Services.....	369
Managing Kibana Service on Linux.....	369
Troubleshooting Kibana Configuration.....	370
Kibana Menu is not Displayed Under Reporting Tools.....	370
Errors on Kibana Visualization Tiles.....	370
Certificate Error When Launching Kibana Visualization Tiles in the Chrome Browser.....	372
Unable to Save a Visualization.....	372

Preface

Understanding the PeopleSoft Online Help and PeopleBooks

The PeopleSoft Online Help is a website that enables you to view all help content for PeopleSoft applications and PeopleTools. The help provides standard navigation and full-text searching, as well as context-sensitive online help for PeopleSoft users.

Hosted PeopleSoft Online Help

You can access the hosted PeopleSoft Online Help on the [Oracle Help Center](#). The hosted PeopleSoft Online Help is updated on a regular schedule, ensuring that you have access to the most current documentation. This reduces the need to view separate documentation posts for application maintenance on My Oracle Support. The hosted PeopleSoft Online Help is available in English only.

To configure the context-sensitive help for your PeopleSoft applications to use the Oracle Help Center, see [Configuring Context-Sensitive Help Using the Hosted Online Help Website](#).

Locally Installed Help

If you're setting up an on-premise PeopleSoft environment, and your organization has firewall restrictions that prevent you from using the hosted PeopleSoft Online Help, you can install the online help locally. See [Configuring Context-Sensitive Help Using a Locally Installed Online Help Website](#).

Downloadable PeopleBook PDF Files

You can access downloadable PDF versions of the help content in the traditional PeopleBook format on the [Oracle Help Center](#). The content in the PeopleBook PDFs is the same as the content in the PeopleSoft Online Help, but it has a different structure and it does not include the interactive navigation features that are available in the online help.

Common Help Documentation

Common help documentation contains information that applies to multiple applications. The two main types of common help are:

- Application Fundamentals
- Using PeopleSoft Applications

Most product families provide a set of application fundamentals help topics that discuss essential information about the setup and design of your system. This information applies to many or all applications in the PeopleSoft product family. Whether you are implementing a single application, some combination of applications within the product family, or the entire product family, you should be familiar with the contents of the appropriate application fundamentals help. They provide the starting points for fundamental implementation tasks.

In addition, the *PeopleTools: Applications User's Guide* introduces you to the various elements of the PeopleSoft Pure Internet Architecture. It also explains how to use the navigational hierarchy, components, and pages to perform basic functions as you navigate through the system. While your application or implementation may differ, the topics in this user's guide provide general information about using PeopleSoft applications.

Field and Control Definitions

PeopleSoft documentation includes definitions for most fields and controls that appear on application pages. These definitions describe how to use a field or control, where populated values come from, the effects of selecting certain values, and so on. If a field or control is not defined, then it either requires no additional explanation or is documented in a common elements section earlier in the documentation. For example, the Date field rarely requires additional explanation and may not be defined in the documentation for some pages.

Typographical Conventions

The following table describes the typographical conventions that are used in the online help.

<i>Typographical Convention</i>	<i>Description</i>
Key+Key	Indicates a key combination action. For example, a plus sign (+) between keys means that you must hold down the first key while you press the second key. For Alt+W , hold down the Alt key while you press the W key.
... (ellipses)	Indicate that the preceding item or series can be repeated any number of times in PeopleCode syntax.
{ } (curly braces)	Indicate a choice between two options in PeopleCode syntax. Options are separated by a pipe ().
[] (square brackets)	Indicate optional items in PeopleCode syntax.
& (ampersand)	When placed before a parameter in PeopleCode syntax, an ampersand indicates that the parameter is an already instantiated object. Ampersands also precede all PeopleCode variables.
⇒	This continuation character has been inserted at the end of a line of code that has been wrapped at the page margin. The code should be viewed or entered as a single, continuous line of code without the continuation character.

ISO Country and Currency Codes

PeopleSoft Online Help topics use International Organization for Standardization (ISO) country and currency codes to identify country-specific information and monetary amounts.

ISO country codes may appear as country identifiers, and ISO currency codes may appear as currency identifiers in your PeopleSoft documentation. Reference to an ISO country code in your documentation does not imply that your application includes every ISO country code. The following example is a country-specific heading: "(FRA) Hiring an Employee."

The PeopleSoft Currency Code table (CURRENCY_CD_TBL) contains sample currency code data. The Currency Code table is based on ISO Standard 4217, "Codes for the representation of currencies," and also relies on ISO country codes in the Country table (COUNTRY_TBL). The navigation to the pages where you maintain currency code and country information depends on which PeopleSoft applications you are using. To access the pages for maintaining the Currency Code and Country tables, consult the online help for your applications for more information.

Region and Industry Identifiers

Information that applies only to a specific region or industry is preceded by a standard identifier in parentheses. This identifier typically appears at the beginning of a section heading, but it may also appear at the beginning of a note or other text.

Example of a region-specific heading: "(Latin America) Setting Up Depreciation"

Region Identifiers

Regions are identified by the region name. The following region identifiers may appear in the PeopleSoft Online Help:

- Asia Pacific
- Europe
- Latin America
- North America

Industry Identifiers

Industries are identified by the industry name or by an abbreviation for that industry. The following industry identifiers may appear in the PeopleSoft Online Help:

- USF (U.S. Federal)
- E&G (Education and Government)

Translations and Embedded Help

PeopleSoft 9.2 software applications include translated embedded help. With the 9.2 release, PeopleSoft aligns with the other Oracle applications by focusing our translation efforts on embedded help. We are not planning to translate our traditional online help and PeopleBooks documentation. Instead we offer very direct translated help at crucial spots within our application through our embedded help widgets. Additionally, we have a one-to-one mapping of application and help translations, meaning that the software and embedded help translation footprint is identical—something we were never able to accomplish in the past.

Using and Managing the PeopleSoft Online Help

Select About This Help in the left navigation panel on any page in the PeopleSoft Online Help to see information on the following topics:

- Using the PeopleSoft Online Help.
- Managing hosted Online Help.
- Managing locally installed PeopleSoft Online Help.

PeopleTools Related Links

[PeopleTools 8.60 Home Page](#)

[PeopleSoft Search and Kibana Analytics Home Page](#)

"PeopleTools Product/Feature PeopleBook Index" (Getting Started with PeopleTools)

[PeopleSoft Online Help](#)

[PeopleSoft Information Portal](#)

[PeopleSoft Spotlight Series](#)

[PeopleSoft Training and Certification | Oracle University](#)

[My Oracle Support](#)

[Oracle Help Center](#)

Contact Us

Send your suggestions to psft-infodev_us@oracle.com.

Please include the applications update image or PeopleTools release that you're using.

Follow Us

<i>Icon</i>	<i>Link</i>
	<u>YouTube</u>

Icon	Link
	Twitter@PeopleSoft_Info.
	PeopleSoft Blogs
	LinkedIn

Chapter 1

Getting Started with PeopleSoft Search Technology

PeopleSoft Search Technology Overview

The PeopleSoft Search Framework provides a standard, declarative method for creating, deploying, and maintaining search indexes for all of your PeopleSoft applications. Beginning with PeopleTools 8.56, the PeopleSoft Search Framework uses Elasticsearch as a search engine. Elasticsearch is an open-source search engine based on Apache Lucene™. PeopleTools is integrated with Elasticsearch version 7.10, Kibana version 7.10, and Logstash version 7.10.

The Elasticsearch search engine complements the PeopleSoft Search Framework search functionality with the following advantages:

- Seamless deployment using Elasticsearch DPK and ACM plug-ins.
- Ability to deploy highly available indexes, with fail over that is scalable.
- Excellent crawl and search performance.
- Efficient resource utilization on the search engine host making high volume crawling and searching practical.
- Excellent data recovery management.
- Efficient data collection engine and analytic and visualization platform.

Implementing PeopleSoft Search Using Elasticsearch

The PeopleSoft Search Framework involves the proper configuration and use of the following systems, tools, and technology. To administer or develop applications using the PeopleSoft Search Framework, a working knowledge of these systems is recommended.

Required Item	Documentation
PeopleTools	<i>PeopleTools Installation for <your platform></i> For PeopleTools 8.59 hardware and software requirements, see My Oracle Support, Certifications.

Required Item	Documentation
Elasticsearch	See <i>PeopleSoft Deployment Packages for Elasticsearch Installation on My Oracle Support (Doc ID 2205540.2)</i> . See Elasticsearch Reference [7.10] (www.elastic.co).
Kibana	See <i>PeopleSoft Deployment Packages for Elasticsearch Installation on My Oracle Support (Doc ID 2205540.2)</i> . See Kibana Guide [7.10] (www.elastic.co).
PeopleSoft Integration Broker	See Integration Broker Administration.
PeopleSoft Query and Connected Query	See Query.
PeopleCode	See "Understanding the PeopleSoft Search Framework Classes" (PeopleCode API Reference).

Understanding PeopleSoft Search Framework Integration with Elasticsearch

PeopleSoft Search Framework Features

The PeopleSoft Search Framework enables application developers and implementation teams to create search artifacts in a consistent, declarative manner and to deploy and maintain search indexes using one standard interface, regardless of the PeopleSoft application.

The PeopleSoft Search Framework consists of PeopleSoft components (pages and records provided by PeopleTools), which provide a centralized interface for configuring PeopleSoft integration with the search engine, creating search artifacts such as search definitions, search categories, and building and maintaining search indexes.

Some of the key features include:

- Search Administration Activity Guide provides a familiar PeopleSoft browser interface enabling you to: configure search engine connectivity, deploy search definitions, schedule index creation and maintenance, and run incremental index updates from one location.
- Search Designer Activity Guide provides a familiar PeopleSoft browser interface enabling you to: create search definitions, create search categories, define title and summary result display.
- Flexible security: indexes can be created with source-level security, document-level security, or no security. PeopleSoft Search Framework allows for fully authenticated and authorized search queries.
- Search Query API enabling rich display options, such as filtering.
- Visualizations of system metrics and indexing metrics using Kibana.

The PeopleSoft Search Framework enables you to generate search indexes using these source types:

- PeopleSoft Query and Connected Query: this option provides a familiar and intuitive way to declare the fields relevant for end user searches. With PeopleSoft Query you define your search meta data attributes. You use the Search Framework Designer to map query fields to meta data attributes, save the search definition to the database, and create search categories.
- Web source: this option enables you to index content deployed on a website that you want to make available for end user searches.
- File source: this option enables you to index files residing in your file system, such as reports.

After designing search definitions, the Search Framework Administration activity guide enables you to establish integration with Elasticsearch (by way of Integration Broker), deploy the search definitions and search categories, and build indexes.

Search Framework Definitions

The Search Framework is based on these PeopleSoft definitions:

Definition	Description
Search Definition	<p>Created in the Search Framework designer activity guide, a search definition maps the PeopleSoft Query fields, web source, or file source, to searchable attributes in the search engine.</p> <p>The search definition also enables you to attach security attributes to restrict access to the search results. The search definition provides the information required by the framework to enable the system to create search results (search documents).</p>
Search Category	<p>Also created in the Search Framework designer activity guide, search categories enable you to group search definitions logically. A search definition must belong to at least one search category. End users run searches against search categories, not individual search definitions. It is a requirement that at least one search category exists with exactly the same name as the search definition.</p>

Search Documents

Search documents describe the format of search results.

The main elements of a PeopleSoft search document are:

- URL
- Title
- Summary

In some cases search documents are non-structured documents, such as a Microsoft Word document or the text in a website. In a PeopleSoft application, the majority of information is structured, (for example, a Purchase Order). PeopleSoft information resides in a relational database where the document attributes constituting the search document are well known such as Employee Name, Customer Name, Product ID, and so on. While most of the PeopleSoft information can be displayed in the form of structured documents, the system also stores unstructured data in the form of attachments.

When a user runs a search based on a set of known attributes, the search returns “hits” in the form of search documents or search results. The user then analyzes the content of the search result to determine relevancy and uses the associated URL to navigate to the desired application page.

Elasticsearch Concepts and Terminology

The Elasticsearch search engine uses the following concepts and terminology.

Descriptions of the concepts and terminology listed in the following table are from the Elasticsearch website (www.elastic.co). Refer to Elasticsearch Reference [6.1], Basic Concepts.

<i>Elasticsearch Terminology</i>	<i>Description</i>
Cluster	<p>“A cluster is a collection of one or more nodes (servers) that together holds your entire data and provides federated indexing and search capabilities across all nodes. A cluster is identified by a unique name which by default is ‘elasticsearch’. This name is important because a node can only be part of a cluster if the node is set up to join the cluster by its name.”</p> <p>“Make sure that you don’t reuse the same cluster names in different environments, otherwise you might end up with nodes joining the wrong cluster. For instance you could use logging-dev, logging-stage, and logging-prod for the development, staging, and production clusters.”</p> <p>In a PeopleSoft implementation, you can specify a cluster name when you install Elasticsearch, and if you need to change the name, you need to edit the <code>elasticsearch.yml</code> configuration file.</p>
Node	<p>“A node is a single server that is part of your cluster, stores your data, and participates in the cluster’s indexing and search capabilities. Just like a cluster, a node is identified by a name. You can define any node name you want if you do not want the default. This name is important for administration purposes where you want to identify which servers in your network correspond to which nodes in your Elasticsearch cluster.”</p>
Index	<p>“An index is a collection of documents that have somewhat similar characteristics. In a single cluster, you can define as many indexes as you want.”</p> <p>An index is an equivalent of a relational database.</p> <p>In a PeopleSoft implementation, by default, each search definitions/search categories is deployed as an individual index.</p>
Type	<p>Type is the Elasticsearch meta object where the mapping for an index is stored.</p> <p>In a PeopleSoft implementation, each search definition corresponds to a type in Elasticsearch.</p>

<i>Elasticsearch Terminology</i>	<i>Description</i>
Alias	<p>Alias is a reference to an Elasticsearch index. An alias can be mapped to more than one index.</p> <p>In a PeopleSoft implementation, a search category is mapped as an Alias on the Elasticsearch server.</p>
Document	<p>“A document is a basic unit of information that can be indexed. This document is expressed in JavaScript Object Notation (JSON) format.”</p> <p>Connected query returns parent and child rows. In a PeopleSoft implementation, each row returned from the main query corresponds to a document in Elasticsearch and child information is attached to the main query and is sent as one document.</p>
Shards and Replicas	<p>“Elasticsearch provides the ability to subdivide your index into multiple pieces called shards. When you create an index, you can simply define the number of shards that you want. Each shard is in itself a fully-functional and independent ‘index’ that can be hosted on any node in the cluster.</p> <p>Elasticsearch allows you to make one or more copies of your index’s shards into what are called replica shards, or replicas for short.</p> <p>After the index is created, you may change the number of replicas dynamically anytime but you cannot change the number of shards after-the-fact.”</p>

Elasticsearch Support Based on PeopleSoft Application Version

In order to utilize Elasticsearch for your PeopleSoft application, you will need to apply updates and fixes from your application’s most current update image.

See *PeopleSoft Applications: Elasticsearch Updates and Fixes on My Oracle Support (Doc ID: 2181988.1)* for more information on which update images have the required fixes for your application.

PeopleSoft Search Framework Supported Features in Elasticsearch

This table lists the key functionality of Elasticsearch from the perspective of the use and implementation within a PeopleSoft application system.

<i>Item</i>	<i>Elasticsearch Functionality</i>
High availability	The recommended approach to high availability is by deploying an Elasticsearch cluster with at least 3 nodes on machines that are not on the same physical server and by following the configuration guidelines.

Item	Elasticsearch Functionality
Fail-over mechanism	Elasticsearch search engine is distributed by nature, that is, it knows how to manage multiple nodes. A cluster can contain one or more nodes and one of the nodes act as a master node. Elasticsearch automatically distributes the indexed data to the nodes based on its constructs of primary shards and replica shards, so indexed data is available even when a node fails.
Phonetic	<p>Elasticsearch supports phonetic search using a plug-in that is installed when you install Elasticsearch in a PeopleSoft implementation.</p> <hr/> <p>Note: For custom search pages, PeopleSoft Applications can use PeopleSoft APIs for phonetic queries in Elasticsearch.</p> <hr/>
Wild card search	Elasticsearch supports using a wild card as a leading character in a search text. For example, *racle.
<p>Stop Word</p> <p>Stop Words are words that are commonly used (such as articles, prepositions) and do not contain important significance to be used in a search text.</p>	If a word is in the Elasticsearch’s stop word list, it would be ignored.
Attachment processing	In Elasticsearch, PeopleSoft Search Framework downloads any attachment specified in a search definition, and directly pushes the attachment to the Elasticsearch search engine using the libcurl library bypassing Integration Gateway. Elasticsearch uses the ingest- attachment plug-in to extract the attachment contents and indexes the attachment data. This avoids multiple log-in attempts during crawling.


Item	Elasticsearch Functionality
<p>Stemming</p> <p>Stemming is the process of reducing inflected word to its root or stem. For example, talking, talked, talks will reduce to the root — talk.</p>	<p>Elasticsearch supports limited stemming support, for example, talk/talked/talks. Synonyms are supported by using a file to store the required synonyms as per your business needs.</p> <p>Elasticsearch supports stemming search for the following languages:</p> <ul style="list-style-type: none"> • Arabic • Czech • Danish • Dutch • English • Finnish • French • German • Greek • Hungarian • Italian • Japanese • Korean • Norwegian • Portuguese • Romanian • Russian • Simplified Chinese • Spanish • Swedish • Thai • Turkish • Traditional Chinese
<p>Relevancy Scoring</p>	<p>Relevancy scoring is performed by Elasticsearch, which uses an internal scoring algorithm. For more information, refer to Elasticsearch Reference [7.10].</p>
<p>Custom Scoring</p>	<p>In Elasticsearch, custom scoring is supported through Search Framework APIs.</p>

Item	Elasticsearch Functionality
Synonyms	<p>Elasticsearch uses a file (synonyms_grammar.txt) to store synonyms or any abbreviations that you want to include. If synonyms provided are not adequate for your installation, you can modify this file. The file is located at config\analysis\synonym_grammar.txt.</p> <p>You can enter synonyms in two ways:</p> <pre>"i-pod, i pod => ipod"</pre> <p>OR</p> <pre>"universe, cosmos"</pre>

PeopleSoft Search Framework Architecture

The server topology uses the essential elements of the PeopleSoft Internet Architecture, with the addition of Elasticsearch server and Kibana server.

The following illustration depicts the PeopleSoft server architecture connected to Elasticsearch (using Integration Broker and DirectTransfer), Kibana, and Logstash.




The web server and application server (the PeopleSoft Internet Architecture) provide the interface for the design, administration, and end user access. The PeopleSoft database stores the PeopleSoft queries that define the search data as well as the search definition meta data. PeopleSoft uses the Process Scheduler and Integration Broker’s publishing and subscription functionality to push application data to the Elasticsearch server. In Full Direct Transfer mode, PeopleSoft Search Framework pushes the documents with attachments and documents without attachments to Elasticsearch and bypasses Integration Broker. Elasticsearch stores the deployed search definitions and performs the typical search engine tasks, such as building indexes, crawling for updates, maintaining indexes, and servicing end-user search requests. An Elasticsearch cluster can have multiple nodes that assist in providing high availability and fail-over mechanism. Within the cluster, Elasticsearch internally maintains indexes among the nodes. Kibana is an analytics and visualizations platform, which enables users to create visualizations based on application

indexes in Elasticsearch. It is also a monitoring tool that provides real-time information on the health of the system and Elasticsearch metrics.

The following diagram depicts the relationships between all of the separate elements that are involved with the PeopleSoft Search Framework.

The following illustration depicts PeopleTools, Elasticsearch, Kibana, and Logstash elements interaction to enable the PeopleSoft Search Framework features.



PeopleTools

The following table describes the PeopleTools elements.

<i>PeopleTools Element</i>	<i>Description</i>
PeopleSoft Internet Architecture	Enables access to all aspects of the PeopleSoft Search Framework, including development, administrative, and end user interfaces.
Search Framework Designer	Enables application developers and implementation teams to perform design-time tasks, such as map query fields to search metadata, define search attributes, create search definitions, and create search categories.
Search Metadata	The metadata defined with both PeopleSoft Query and the Search Designer to describe the data that end users will run searches against.
PeopleSoft Query Connected Query	Provides familiar interface for creating SQL queries that declare exactly the data against which end users will run searches, and takes advantage of Query security.
Process Scheduler Integration Broker Publishing and Subscription	Facilitates the transfer of application data from PeopleSoft to the Elasticsearch server.
Search Framework Administration	Provides the interface for system administrators to perform tasks, such as creating a search instance, deploying search objects, scheduling search index builds, and administering Kibana.
SFTP/HTTP Repository Server	Facilitates storing of attachments, which are associated with search definitions.
Search Framework User Interface	Provides the PeopleSoft application end user with the appropriate prompts and fields by which they can submit a search request and review results.
Search Query API	A PeopleCode API that enables application developers to form a valid search request from user input.
PeopleSoft Integration Broker	Facilitates the integration between PeopleTools elements and the search engine (Elasticsearch).
DirectTransfer	Facilitates the transfer of application data to the search engine (Elasticsearch) directly; bypassing the Integration Broker.

Elasticsearch

The following table describes the Elasticsearch elements.

<i>Elasticsearch Element</i>	<i>Description</i>
Indexing Process	<p>The indexing process indexes document search attributes such as Title, Summary, search fields (fields to index), and so on. It can also index document-level security attributes if document-level security is being used.</p> <p>Elasticsearch uses the Ingest attachment plug-in to process attachments, both from the database and from the repository (FTP, SFTP, HTTP).</p>
Security Plug-in	<p>This module gets invoked from the search engine for a user who initiates a search from an application. The search security framework is responsible for authentication, search user validation, and authorization of the search query request.</p> <p>The search engine does a callback with the call back user credentials and passes the search user ID to the PeopleSoft Security Service (Web service) which authorizes the users and gets security attributes for that user. Search results are filtered based on this security attribute. The security plug-in (orcl-security-plugin) is built on top of the Elasticsearch search engine.</p>
Administrative REST APIs	<p>The administrative tasks include various processes, such as, deploying, undeploying, report synchronization, and updating.</p> <p>The http methods of GET, PUT, POST, and DELETE are used while sending the REST API requests to Elasticsearch.</p>
AddSynonyms Plug-in	<p>The AddSynonyms plug-in updates the synonyms file. The Copy to Elasticsearch action in PIA calls the AddSynonyms plug-in in Elasticsearch to update the synonyms file.</p>
Search Engine	<p>PeopleSoft Search Framework works with Elasticsearch search engine. Elasticsearch is an open-source search engine, which is based on Apache Lucene™.</p>

Kibana

The following table describes the Kibana elements.

<i>Kibana Element</i>	<i>Description</i>
Analytics and Visualizations platform	Kibana is an analytics and visualizations platform that uses indexes in Elasticsearch to display data in dashboards and visualizations.
Security Plug-in	The security plug-in module gets invoked from Kibana for a user who attempts to view a Kibana dashboard or visualization. Kibana does a callback with user credentials to Elasticsearch, which in turn does a callback with the call back user credentials and passes the search user ID to the PeopleSoft Security Service (Web service) which authorizes the users and gets security attributes for that user. Data in visualizations and dashboards are filtered based on this security attribute. The search security framework is responsible for authentication, search user validation, and authorization of the request.

Logstash

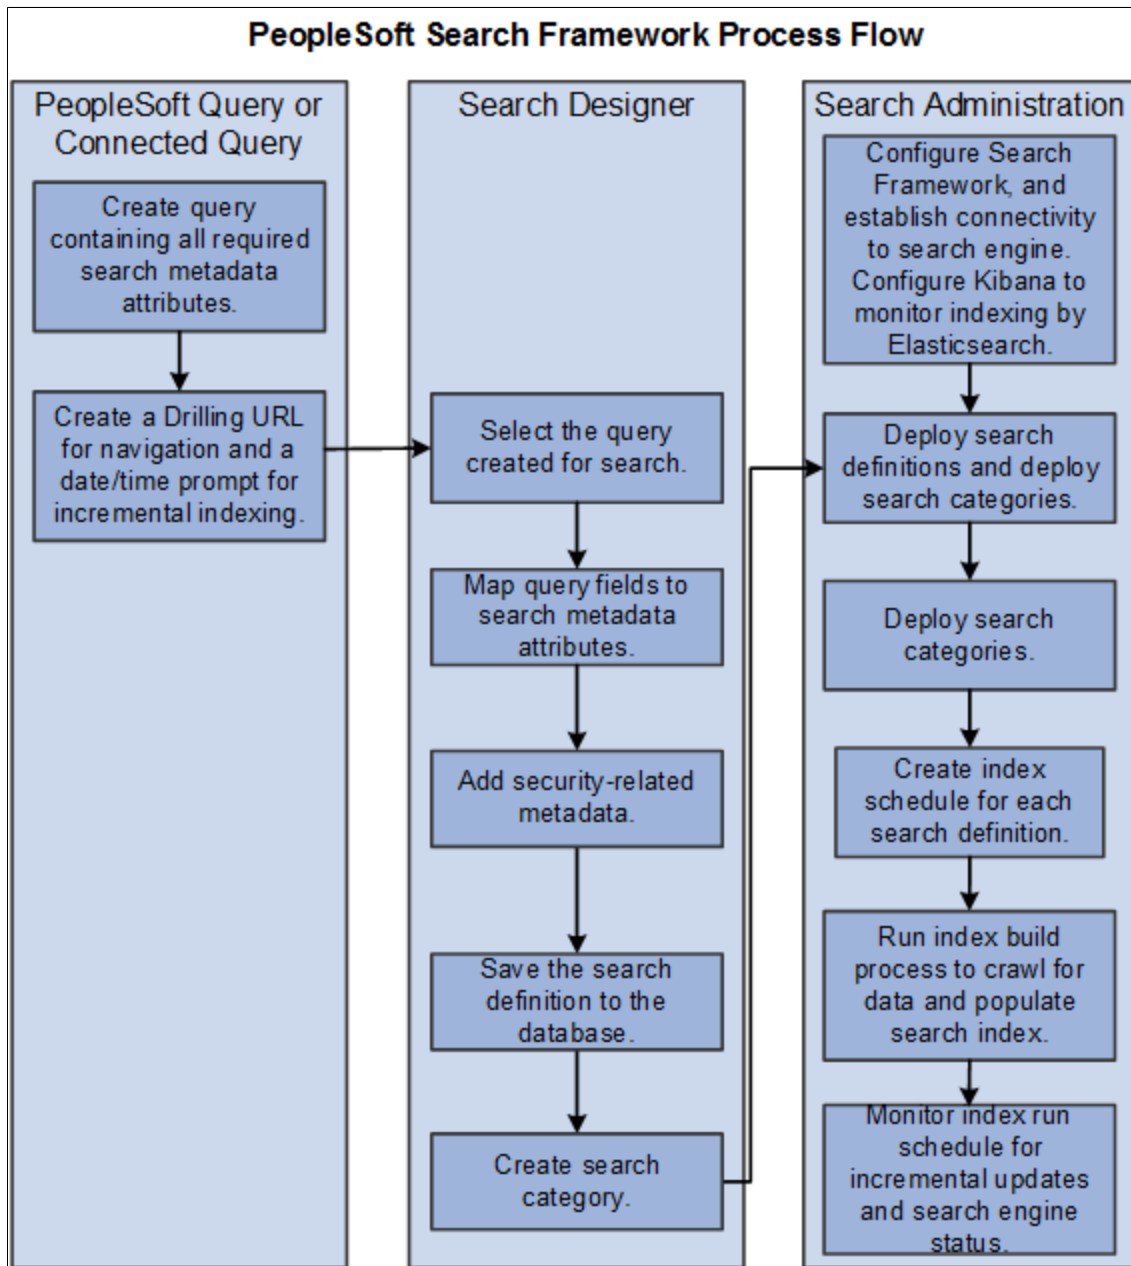
The following table describes the Logstash elements.

<i>Kibana Element</i>	<i>Description</i>
External Data Integration	External Data Integration enables to integrate external data with PeopleSoft. It uses Logstash to ingest the data from external sources to Elasticsearch.
PeopleSoft Health Center	PeopleSoft Health Center monitors the health of your PeopleSoft applications. It uses Logstash to collect data.

Implementation Process Flow

The following diagram illustrates the general process flow when implementing the PeopleSoft Search Framework and deploying search definitions. You first identify the business data you want to expose to text searches and create queries using PeopleSoft Query. PeopleSoft Query selects the appropriate data from your transaction tables. You map the query fields to the search metadata attributes and map the search definition to a search category in the Search Designer. In the Search Administration activity guide, you deploy the search definitions and categories, schedule index builds, and schedule index crawling so that the index can be updated as needed to reflect the current business data.

The following illustration depicts the Search Framework implementation, beginning with PeopleSoft Query and Connected Query, moving to the Search Designer for creating search definitions and categories, then moving to Search Administration for deploying search definitions and defining index builds.



Note: For web source and file source search definitions, PeopleSoft Query is not used. For these source types, you only need to point to the location of the web source or file source in the respective search definition.

PeopleSoft Search Features

Once Search Framework is configured and search definitions have been deployed and tested, these PeopleSoft Search features are available for your end users:

- Global Search.
- Search Pages.

Global Search

With Global Search enabled, the Global Search bar displays in the application header. The Global Search bar provides a drop-down list for selecting a specific search category against which to run a search.

See [Working with Global Search](#) and [Working with Search Results in Fluid User Interface](#).

See [Working with Suggestions Search](#).


This example illustrates the Global Search bar.



Search Pages

If you've enabled the Search Pages features for Search Framework by mapping a component to a search definition, the Keyword Search tab appears amongst the search pages the end user encounters when navigating to a component. The Keyword Search page enables users to execute a deeper, more free-form search to access application data. For example, the Find an Existing Value tab is limited to level 0 data, where the Keyword Search tab lets the users enter custom keywords and the system will search indexed data from levels 0-3.

This example illustrates the fields and controls on the Keyword Search page. You can find definitions for the fields and controls later on this page.



Note: The Find an Existing Value search page is equivalent to the component search pages in previous versions of PeopleSoft applications.

See [Working with Search Pages](#) and [Working with Search Results in Fluid User Interface](#).

Working with Elasticsearch Clusters

Understanding Clusters in Elasticsearch

Elasticsearch runs on Java Virtual Machines. Each JVM instance running Elasticsearch can be considered as an Elasticsearch node. To provide redundancy and scaling, Elasticsearch supports the concept of a cluster. Multiple nodes running on one or more hosts (physical or virtual) can be grouped into a cluster, using a unique "cluster name." Further, the index and source data managed by each node can be configured to suit load patterns and redundancy needs. Searches performed on a cluster will return integrated results from each of the nodes.

Different types of nodes can be present in a cluster. One node in the cluster is elected to be the master node, which is responsible for lightweight cluster-wide actions such as creating or deleting an index, tracking which nodes are part of the cluster, and deciding which shards to allocate to which nodes. Any master-eligible node (all nodes by default are master-eligible nodes) may be elected to become the master node by the master election process. Elasticsearch transparently manages the process of gathering the response from node or nodes holding the data.

Note: In the PeopleSoft implementation of an Elasticsearch cluster, the cluster contains only the master-data node, that is, the node can act both as a master and as a data node.

Depending on the purpose of your multi-node cluster, you may set up nodes on different machines or on a single machine. In a production environment, you would set up nodes on different machines when you want to achieve high availability. In a non-production environment, when you want to test the setup and multi-node cluster, you may set up multiple nodes on the same machine, but note that you may not be able to test the high-availability feature.

Elasticsearch provides the ability to subdivide your index into multiple pieces called shards. When you create an index, you can simply define the number of shards that you want. Each shard is in itself a fully-functional and independent 'index' that can be hosted on any node in the cluster. Elasticsearch allows you to make one or more copies of your index's shards into what are called replica shards, or replicas for short.

This topic discusses:

- Setting up a cluster when you perform a fresh deployment of Elasticsearch using the PeopleSoft delivered Deployment Packages (DPK).
- Adding new nodes to an existing Elasticsearch cluster.
- Verifying the cluster setup.
- Setting the number of replicas.
- Starting and stopping a cluster.
- Removing a node from a cluster.

- Managing large indexes.

Setting Up a Cluster During the Elasticsearch Installation

When you perform a fresh install of Elasticsearch using the PeopleSoft Deployment Packages (DPK), the installation process enables you to set up an Elasticsearch cluster with a single node or with multiple nodes.

For the step-by-step description of the installation process for your operating system, see PeopleSoft Deployment Packages for Elasticsearch Installation (Doc ID 2205540.2) on My Oracle Support.

During the installation process, you are prompted to specify various configuration parameters. If you are manually installing Elasticsearch, you are required to edit the `elasticsearch.yml` file where you can specify values for the parameters pertaining to a cluster. The configuration parameters that are specific to a cluster are described here:

- A name for the cluster.
- Minimum number of master nodes.
- List of host names where Elasticsearch will be installed or is installed.

Naming the Cluster

The setup script prompts you to enter a name for the Elasticsearch cluster.

If you are performing a manual installation of Elasticsearch, you need to enter a value for the **cluster.name** parameter in the `elasticsearch.yml` file.

A cluster is identified by a unique name. This name is important because a node can only be part of a cluster if the node is set up to join the cluster by its name. Make sure that you don't reuse the same cluster names in different environments, otherwise you might end up with nodes joining the wrong cluster. For instance you could use `logging-dev`, `logging-stage`, and `logging-prod` for the development, staging, and production clusters.

Minimum Number of Master Nodes

The setup script prompts you to enter the minimum number of master nodes.

If you are performing a manual installation of Elasticsearch, you need to enter a value for the **discovery.zen.minimum_master_nodes** parameter in the `elasticsearch.yml` file.

Oracle PeopleSoft recommends a cluster with three nodes for high availability and to prevent a split-brain scenario. Then, depending on the hardware, memory availability, and search performance, you may add nodes to the existing cluster.

This setting must be set to a quorum of your master-eligible nodes. It is recommended to avoid having only two master-eligible nodes, since a quorum of two is two. Therefore, a loss of either master-eligible node will result in an inoperable cluster. To prevent the split-brain scenario, PeopleSoft recommends an odd number of nodes in the cluster, so configure the majority of nodes ($\text{total number of master-eligible nodes} / 2 + 1$).

For more information on split-brain scenario, refer to the Elasticsearch online documentation.

List of Host Names where Elasticsearch will be Installed or is Installed

The setup script prompts you to enter host names of nodes present in the cluster.

If you are performing a manual installation of Elasticsearch, you need to enter a value for the **discovery.seed_hosts** parameter in the `elasticsearch.yml` file.

A node name is a symbolic name for identifying the node. Host name is the IP or DNS of the machine where Elasticsearch is installed. Host names(or IP/DNS) are required for letting each Elasticsearch server where it can ping and find other Elasticsearch servers during booting up.

Enter the host name for any nodes that are already members of a cluster. Enclose one or more host names in square brackets, with the host name or IP address in double quotes.

For example:

- For one host, you would enter: `["host1.example.com"]`
- To list two or more hosts, use commas: `["host1.example.com", "192.0.2.1"]`

High-level Steps to Set Up a Cluster

This section provides the high-level steps to set up a cluster when you are installing Elasticsearch. For steps 1 and 2, you may use the *PeopleSoft Deployment Packages for Elasticsearch Installation (Doc ID 2205540.2)* on My Oracle Support for detailed instructions.

1. Obtain the Elasticsearch DPK for your operating system.
2. Run the script to install Elasticsearch. The script creates a cluster with one or multiple nodes, which you need to specify.

While installing Elasticsearch ensure that you provide values for the three configuration parameters that are essential to set up a cluster, which are discussed in an earlier section.

- Cluster name.
 - Minimum number of master nodes.
 - List of host names where Elasticsearch is installed.
3. Verify the cluster setup by executing the following command in a browser:

```
http(s)://host:port/_cluster/stats?
```

Where, host refers to the Elasticsearch host.

Enter the Elasticsearch user credentials (esadmin) when prompted for login information.

The cluster and node information can be found in the `nodes/count/total` section of the response.

Other methods to verify the cluster and node setup, which involves using the PeopleSoft Search Framework and PeopleSoft Health Center, are discussed in a later section of this topic. See [Verifying the Cluster and its Nodes](#).

To complete the search functionality setup, you will need to follow the post-installation instructions in the *PeopleSoft Deployment Packages for Elasticsearch Installation (Doc ID 2205540.2)* on My Oracle Support.

Also, refer to [Understanding PeopleSoft Search Framework Administration](#).

Upgrading Existing Cluster Using ELK DPK

PeopleTools 8.59 is integrated with Elasticsearch version 7.10 and Kibana version 7.10. Search Framework enables you to upgrade from the previous version of Elasticsearch 6.1.2 or 7.0 and Kibana 6.1.2 or 7.0 without requiring you to perform a full index build.

You use the ELK DPK to upgrade from 6.1.2 or 7.0 to 7.10. You may choose to run the ELK DPK in interactive mode or in silent mode. The instructions to upgrade are described in *PeopleSoft Deployment Packages for Elasticsearch Installation on My Oracle Support (Doc ID 2205540.2)*.

Adding a New Node to an Existing Cluster

As determined by your business requirements for high availability and disaster recovery, you will need to determine the number of nodes appropriate for your installation. In addition to high availability and disaster recovery, system CPU and memory limitations may require addition of nodes to meet performance requirements. This section discusses how to add a node to your existing cluster.

To add a new node to an existing Elasticsearch cluster, you should complete the following steps:

1. Install Elasticsearch on a separate machine using the instructions provided in the PeopleSoft Deployment Packages for Elasticsearch Installation (Doc ID 2205540.2) on My Oracle Support for your operating system.
 - For the cluster name prompt, ensure that you enter the name of the existing cluster to which you want to add this node.
 - For the minimum number of master nodes prompt, before you enter a value consider the parameter description that is discussed in an earlier section.

See [Minimum Number of Master Nodes](#).

- For the host names of existing nodes, ensure that you enter the host names of the existing nodes so that the new node will discover the other nodes in the cluster.

Complete the Elasticsearch installation by following the steps described in PeopleSoft Deployment Packages for Elasticsearch Installation. After completing the installation, verify the cluster and node setup.

- Open the `elasticsearch.yml` configuration file of the existing nodes in the cluster and edit the following parameters.

Field or Control	Description
<code>discovery.zen.minimum_master_nodes</code>	Enter the value that you provided in Step 1.
<code>discovery.seed_hosts</code>	Add the host of the new node to the existing list of hosts. Elasticsearch nodes will find each other via unicast. Ensure that you enter the address in the correct format: ["host1", "host2"]. See List of Host Names where Elasticsearch will be Installed or is Installed .

- The `elasticsearch.yml` files of the existing nodes is updated with the host of the new node so that, in the future, whenever any of the nodes is restarted, the restarted node will be able to discover other nodes in the cluster mentioned in the `discovery.zen.ping.unicast.hosts` parameter.

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

- On PeopleSoft Search Framework, add the details of the new node to the Elasticsearch search instance on the Search Instance Properties page.

See [Working With Search Instances](#).

Verifying the Cluster and its Nodes

After setting up an Elasticsearch cluster, you can verify whether a cluster and its nodes are created by using the Elasticsearch Interact page.

Using Elasticsearch Interact Page

The Elasticsearch Interact page (available with PeopleTools 8.55.15) enables you to retrieve information from the Elasticsearch server regarding the statistics of a cluster and node among other information.

On the Elasticsearch Interact page, in the **Service Type** drop-down list, select *Cluster* and in the **Cluster API** drop-down list, select `_cluster/stats?pretty=true`.

See [Administering Elasticsearch Using the Elasticsearch Interact Page](#).

Alternatively, you can verify the cluster setup by executing the following command in a browser:

```
http(s)://host:port/_cluster/stats?
```

Where, host refers to the Elasticsearch host.

Enter the Elasticsearch user credentials (esadmin) when prompted for login information.

The cluster and node information can be found in the `nodes/count/total` section of the response.

Description of Color Codes for Elasticsearch Cluster

The status of the Elasticsearch cluster is indicated by the following color codes:

- Green - Indicates that the Elasticsearch cluster is fully operational. That is, all primary and replica shards are allocated.
- Yellow - Indicates a warning. All primary shards are allocated, but at least one replica is missing. Data is not missing, so search results will still be complete. However, high availability is compromised to some degree. If more shards disappear, data maybe lost.
- Red - Indicates that at least one primary shard (and all of its replicas) is missing. That is, data is missing and as a result searches will return partial results and indexing of the missing shard will return an exception.

Setting the Number of Replicas

In a multi-node cluster, you must ensure that the replica values are set correctly. You use the Search Options page to set the replica values.

On the Search Options page, when you specify a replica value, the replica value is applicable to a search instance and to the security index. Consider the following before you specify a replica value.

See [Managing General Search Options](#).

- Ensure that the replica value is set to 1 or greater than 1 so that in the event of a node failure data loss does not happen. The maximum number of nodes that can be down without loss of data will be equal to the number of replicas (replica count should be less than or equal to N-1), where N refers to the number of nodes in a cluster. Ensure that you set the replica value in all the nodes of a cluster.

Warning! Oracle PeopleSoft recommends that you do not set the replica value to 0.

- You may set the replica value to greater than 1 if you have more nodes available. For example, if you have 4 nodes in a cluster, and you want to provide high availability even if two nodes fail, then you need to set the replica value to 2.

If the replica value for an index is greater than N-1 (where N is the number of nodes in the cluster), you may find that the status of the cluster as yellow indicating a warning that at least one replica is missing.

For a description of color codes for cluster status, see [Verifying the Cluster and its Nodes](#).

Starting and Stopping a Cluster

Starting a cluster would mean starting all the nodes of a cluster. Similarly, stopping a cluster would mean stopping all the nodes of a cluster.

You start or stop a node by starting or stopping the Elasticsearch service on each node.

Important! Before you stop a node, you should ensure that no indexing requests or administration-related tasks are being made on the cluster. If you stop a node during indexing, the cluster meta data might get corrupted and the cluster could become non-operational (red color code). To ensure that no instances of PTSF_GENFEED are running, check the Process Monitor. If all processes are completed, you may stop all the nodes in a cluster and make the required modifications. After completing the modifications, you may start all the nodes of the cluster.

Starting and Stopping a Node

You start or stop a node by starting or stopping the Elasticsearch service on the node. Use the following steps to start or stop a node on Windows and on Linux.

Starting and Stopping the Elasticsearch Service on Windows

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

To start or stop an Elasticsearch service:

1. Open a command prompt, and change directory to ES_HOME/bin.
2. To see the usage for the service command:

```
service
service.bat install|remove|start|stop
```

- To start the Elasticsearch service:

```
service.bat start
```

- To stop the Elasticsearch service:

```
service.bat stop
```

Starting and Stopping the Elasticsearch Service on Linux

To start or stop the Elasticsearch service:

- In a terminal window, change directory to ES_HOME/bin.
- To start the Elasticsearch process:

```
nohup ./elasticsearch &
```

- To stop the Elasticsearch process:

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

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

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

```
kill <pid>
```

Removing a Node from a Cluster

To safely remove a node from a cluster, perform the following steps:

1. Before you remove a node, ensure that each index has at least one replica. You can use the Elasticsearch Interact page to verify the replica count by choosing `_cat/indices?v` option in the Cluster API drop-down list.

Alternatively, you can use the Elasticsearch Cluster Dashboard to view the replica information in PeopleSoft Health Center.

2. Identify the IP address of the Elasticsearch node that needs to be removed from the cluster.

Open a command prompt window and execute the following command where X.X.X.X stands for the IP address of the node that needs to be removed and also replace the host:port with the IP or port of any of the nodes in the cluster.

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

3. When the command is executed, Elasticsearch tries to move the existing shards out of the node that will be removed and moves it to other nodes in the cluster.
4. In a command prompt window, execute the following command to ensure that the relocation is complete and to validate that the `relocating_shards` attribute shows value as 0 (zero). Alternatively, you may use the `_cluster/health` option from the Cluster API drop-down list on the Elasticsearch Interact page.

```
curl -XGET http://<username:password@host:port>/_cluster/health?pretty
```

For example, `$>curl -XGET http://esadmin:esadmin@example.com:9200/_cluster/health?pretty`.

```
{
  "cluster_name" : "ES_CLUSTER1",
  "status" : "green",
  "timed_out" : false,
  "number_of_nodes" : 3,
  "number_of_data_nodes" : 3,
  "active_primary_shards" : 630,
  "active_shards" : 635,
  "relocating_shards" : 0,
  "initializing_shards" : 0,
  "unassigned_shards" : 0,
  "delayed_unassigned_shards" : 0,
  "number_of_pending_tasks" : 0,
  "number_of_in_flight_fetch" : 0,
  "task_max_waiting_in_queue_millis" : 0,
  "active_shards_percent_as_number" : 100.0
}
```

5. Shutdown the node that was identified for removal.
6. Set the exclusion rules to empty so that, in the future, the removed node can be added back to the cluster and can be used to hold the shards.

```
curl -XPUT "http://username:password@host:port/_cluster/settings" -d
```



```
'{
  "transient" :{
    "cluster.routing.allocation.exclude._ip" : ""
  }
}'
```

7. On the Search Instance Properties page, delete the row corresponding to the removed node.

See [Working With Search Instances](#).

Managing Large Indexes

To improve overall performance of the search functionality, Oracle recommends that you specify a value for shards on the Search Options page (located on the Search Framework Administration menu).

Elasticsearch enables you to distribute the indexed data into partitions, and this distribution of indexed data into partitions is performed by Elasticsearch. It does not affect the indexing process in PeopleSoft Search Framework. The default value is 5 and it can be overridden.

Chapter 4

Defining Search Definition Queries

Understanding Search Definition Queries

To define the PeopleSoft application data that end users can run their searches against, you use these standard PeopleSoft query design tools:

- PeopleSoft Query
- Connected Query

Using these tools provides a familiar interface for PeopleSoft implementation teams and a standard means of defining searchable data amongst all PeopleSoft applications. You will also need to make sure that the data you want to expose to Query is authorized by the Query Security Manager.

Note: If you are creating a search definition of the source type File Source or Web Source, you do not create a query or connected query to define the information to be indexed.

The data returned by the query/connected query serves as the source data that the search engine crawls to create the search index. If you want a particular field available for end user searches, you want to make sure that field is included in the list of fields of your query. Likewise, if you do not think a particular field is appropriate or useful to expose for end user searches, then you can make sure that field is not included in the list of fields for your query. This enables you to declare the scope of the search index and manage its size as well.

You do not need to create new queries solely for creating search definitions. You can re-purpose existing queries in PeopleSoft Query or connect multiple existing queries using Connected Query. Except for a handful of requirements for the Search Framework, queries used for creating search definitions are created exactly as you would any other query.

Queries need the following items defined for use with the Search Framework:

- List of fields to index from authorized records.
- "Last Updated" record field.
- Drilling URL.
- Prompt against the "Last Updated" field.
- Criteria for the "Last Updated" field.

This section describes the requirements and guidelines to follow when defining queries for use with the Search Framework. The PeopleSoft Query and Connected Query documentation is not duplicated within this guide. It is assumed that you have a working knowledge of those products.

Related Links

"PeopleSoft Query Overview" (Query)

"Creating New Queries" (Query)

Understanding the Search Designer Activity Guide

This topic provides overview information about the Search Designer Activity Guide.

The Search Designer Activity Guide is a dashboard in the Search Framework where you can:

- Create search definitions
- Create search categories
- View search attributes

Other page elements that appear on the page are discussed elsewhere in this book.

Defining Search Definition Queries with PeopleSoft Query

This section discusses how to define search definition queries with PeopleSoft query.

Adding a “Last Updated” Field to Records

The application data to be indexed *must* exist in a record containing a column that tracks the update history of the data. This column is referred to as the "Last Updated" field. The Last Updated field is a datetime field that captures when the set of data to be indexed has been changed. Keeping track of the data updates is *critical* for enabling incremental indexing. Rather than recreating the entire index each time the index needs to be updated, incremental indexing enables the system to gather only the information that has changed since the last time the index generation process has run.

The system keeps track of when the index generation process last ran, and compares that time to the Last Updated field value in the underlying record structure. Based on the comparison between those time values, the system can isolate only the data that has undergone a change since the last index generation process run. Once the system creates the initial full index, only those rows that have been updated or added since the last index process run will be collected and added to the existing index. Using incremental index updates improves performance and decreases system overhead.

If the query being used for the search definition only runs against a single record, then that record must contain a datetime field to capture data update date and time values. If the data exists within a hierarchy of tables (grand parent, parent, and child, for example), only one of the records within the hierarchy requires the existence of the datetime field.

Important! The Last Updated field *must* be of the type datetime. An example of this field is the LASTUPDTM field, which can be found in many delivered PeopleSoft applications. Whether you intend to implement an incremental indexing system or not, it is still a requirement to have a "last updated" field within the record.

Note: Many PeopleSoft application tables come with a LASTUPDTM field in place, especially those for which the application has provided search definitions. For any custom tables or tables that do not already track date and time updates, you need to ensure the field exists in the record or record hierarchy.

Note: For any search definitions delivered with your PeopleSoft applications, the underlying records will be configured to include the required datetime field as well as the program logic to ensure that the value of the datetime field is collected. For any custom applications, you will need to add the datetime field *manually*, alter the underlying SQL table, and include program logic to ensure the value of the datetime field is collected and updated accurately. For example, using SavePreChange PeopleCode you can test IsComponentChanged and if so, then update the LASTUPDTM field accordingly.

Specifying a List of Fields to Index

The query defines which fields will comprise the index for a search definition. When working in the query, use the Fields tab to determine if your list is complete. Keep in mind that fields in the index can be used for different purposes. While some fields are the ideal field against which end users would intuitively search (Customer Number, Order Date, and so on), others can be included for different reasons.

For example, some fields are used as metadata to help describe the data contained in the row (resume, invoice, sales order, and so on). Other fields are useful for security in restricting the viewing of the data only to users that have access to a certain type or level of data.

Because the underlying records require the existence of a datetime field to track the "last updated" value, the query must also contain the corresponding query field in the fields list.

Note: When using a query for component keyword searches, all search keys and alternate search keys for a component must be indexed. Select additional fields to be indexed as required by your business processes. Also keep in mind that some fields are codes and might need to be translated to their description for searching.

Important! When saving the query, make sure it is of type *Public*.

Creating a Drilling URL

The drilling URL defines the URL, the target, for the search result. The drilling URL enables the end user to click the link in the search result to display the appropriate PeopleSoft application page, with the appropriate data populated in the page.

Note: Drilling URL must be unique because it is the document ID in the search engine. If it is not unique, data loss in the index may happen.

Note: Elasticsearch has a restriction of 512 bytes for the maximum length of a document ID.

To create a drilling URL for Search Framework:

1. In PeopleSoft Query Manager, select the Expressions tab.
2. Click Add Expression.

3. On the Edit Expressions Properties dialog box, select Drilling URL from the Expression Type drop-down list.
4. Select the appropriate option for the type of drilling URL(s) you are creating.
 - Component URL: For creating the URL to components.
 - Attachment URL: For creating the URL to attachments.
 - Image URL: For creating the URL to images.

Note: For image URLs, also select the Query tab, select the Properties link, and select Image Hyperlink.

5. On the Select a Component page, provide the navigation to the appropriate page.
6. Click Search Keys to define which key(s) needs to be passed to launch the page to display the unique information automatically for a particular search result.

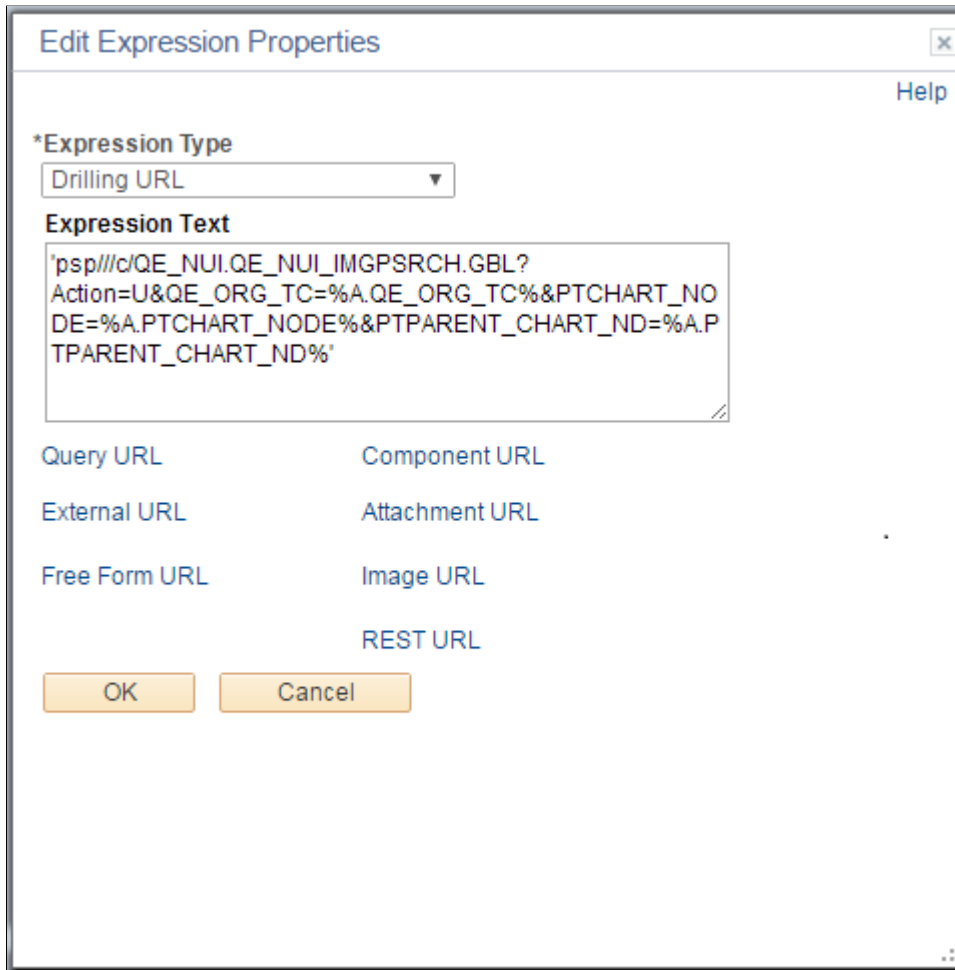
For example, Field Name = CUSTOMER ID and Key Value = A.CUSTOMER_ID.

Note: If you do not specify the correct search keys, the system will launch the page you have specified, but the user would still need to add the keys manually to view the information associated with a search result, which is not the desired behavior.

7. Click OK.
8. Click the Use as Field button to associate the Drilling URL as a field in the Query.

Note: Once you have clicked the Use as Field button, you must edit the 'created' field in the Fields table and give it an existing field name in your PeopleSoft database. Query will assign a temporary field name, but this should be changed for easier reference.

This example illustrates the fields and controls on the Edit Expression Properties dialog box. You can find definitions for the fields and controls later on this page.




When your expression is complete, click the Use as Field link in the Drilling URL grid on the Expressions tab. You should name the field with a valid field name, such as *DRILL_URL*. The field name should be valid in the PeopleSoft database.

See [Working with Attachment Properties](#) and [Working with Images](#).

Creating a Prompt for the Last Modified Field

You need to create a prompt for your query so that the system can use the last update date and time for enabling incremental index updates. This is required by the Search Framework even if you do not intend to implement incremental index updates. The Application Engine program performing incremental index updates requires the valid prompt data. Use the prompt tab in Query Manager to create your prompt. On the Edit Prompt Properties dialog box, select the "last updated" from the Field Name drop-down list.

This example illustrates the fields and controls on the Edit Prompt Properties dialog box. You can find definitions for the fields and controls later on this page.



Defining Prompt Criteria

So that the prompt identifies the correct rows for incremental updates, you must also add query criteria based on the “last updated” field. Use the Criteria tab in Query Manager, and click Add Criteria. Set the criteria as described in the following table.

<i>Item</i>	<i>Value Selected</i>
Expression 1	<ul style="list-style-type: none"> Field. Your “last updated” field.
Condition Type	<ul style="list-style-type: none"> <i>not less than</i>(recommended) greater than

<i>Item</i>	<i>Value Selected</i>
Expression 2	<ul style="list-style-type: none"> • Prompt • Your prompt

For example:

This example illustrates the fields and controls on the Edit Criteria Properties dialog box. You can find definitions for the fields and controls later on this page.

Edit Criteria Properties

Choose Expression 1 Type

Field

Expression

Expression 1

Choose Record and Field

A.ROW_LASTMANT_DTTM - Last Mod

*Condition Type greater than

Choose Expression 2 Type

Field

Expression

Constant

Prompt

Subquery

Expression 2

Define Prompt

Prompt New Prompt Edit Prompt

OK
Cancel

Note: As with all queries, query security must be provided in order to access query information. See Query for more information.

Defining Search Definition Queries with Connected Query

You can also use Connected Query to develop your queries for Search Framework. Connected Query enables you to connect multiple queries together in a parent-child relationship, where the child queries filter results for the parent queries. Using Connected Query lets you connect numerous simpler queries, rather than writing one more complicated query. The smaller more modular queries can be reused in different queries.

When using Connected Query, only the top-level (or root) parent query needs to have the requirements described in this topic. The child queries do not require these elements. For example, only the parent query requires a Drilling URL defined.

For more information on the requirements for the parent query, refer to [Defining Search Definition Queries with PeopleSoft Query](#).

Defining a Deletion Query

In many cases, you may want to consider also writing a deletion, or pruning, query to keep the index and the transactional tables in sync with regard to rows that have been deleted from the transactional tables.

For example, assume a row exists in the transactional table for Big Company, and that row has been included in the search query criteria and indexed by the Search Framework crawler. Now, suppose that row gets deleted from the transactional table because Big Company went out of business and no longer exists. Because the row no longer exists in the transactional table, you do not want the previously indexed data to appear in a user's search results. In a search definition, you associate a deletion query for an index on the Advanced tab in the Define Query to Delete SBO section. There you specify the query name and the Drilling URL field for the deletion query.

For more information on the Drilling URL, see [Working With Advanced Settings](#).

The Delete query generates a feed containing only the records that need to be deleted from the index or those documents/records that are marked for deletion from the index. One method of capturing the deleted rows is to create an audit record on the transactional table in which a record of all deleted rows gets inserted into the audit table. The delete query would then capture the rows in the audit table and remove from the index the document entries matching those rows. The delete query also needs a datetime prompt which returns rows deleted after that datetime value.

Working with Images

The Search Framework supports the use of images in search results. In some cases, a particular set of search documents would be easier to sort by the end user with the help of images appearing just to the left of the search results. For example, this might be useful in procurement catalogs showing items that can be ordered, or perhaps in an employee directory showing pictures of employees.

To include images, you need to ensure your underlying query is referring to the image properly. Similar to creating a drilling URL to a component or attachment, you create a drilling URL pointing to the image field. When creating the drilling URL, on the Edit Expression Properties dialog box, select Image URL. Then you supply the query name and image field name, making sure to set the expression to Use as Field.

When creating the search definition, if an image URL has been associated with the query, the Image URL field appears in the Search Result Field Mapping group box, which you need to set, similar to setting the URL Link.

Related Links

[Defining Search Definition Queries with PeopleSoft Query](#)
[Specifying General Settings](#)

Testing Your Search Definition Query

After you have defined your query in Query Manager with all the correct fields included along with the Search Framework requirements, you can test the query using the Run tab. This runs the query outside of the Search Framework, making sure the query is valid prior to incorporating it into a search definition.

This enables you to determine a variety of items, including:

- The correct rows are being returned by your query.
- The Drilling URL displays the appropriate page and data when clicked.
- You can open any associated attachments (if you can't open it from the query test results, then neither can the search engine nor the PeopleSoft Search Framework).

To test your query with the Run tab:

1. In Query Manager, with your query open, click the Run tab.
2. At the prompt, enter a valid value to run against the Last Modified field.

For example:

01/01/1900 0:0

3. View the results and click the Drilling URL and/or attachment links.

This example illustrates the fields and controls on the Query Manager Run tab.

Unit	Ledger	Account	Total Amt	Amount	Transaction Amt	Base Curr	Drilling URL
1	US001	LOCAL	621100	50000.000	50000.000	50000.000	USD http://slc11ohq.us.oracle.com:8000/pspt55doc11x/CUSTOMER/EMPL/q/?ICAction=ICQryNameURL=PUBLIC.ADB_MTD
2	US001	LOCAL	621100	50000.000	50000.000	50000.000	USD http://slc11ohq.us.oracle.com:8000/pspt55doc11x/CUSTOMER/EMPL/q/?ICAction=ICQryNameURL=PUBLIC.ADB_MTD
3	US001	LOCAL	621100	50000.000	50000.000	50000.000	USD http://slc11ohq.us.oracle.com:8000/pspt55doc11x/CUSTOMER/EMPL/q/?ICAction=ICQryNameURL=PUBLIC.ADB_MTD
4	US001	LOCAL	621100	100000.000	100000.000	100000.000	USD http://slc11ohq.us.oracle.com:8000/pspt55doc11x/CUSTOMER/EMPL/q/?ICAction=ICQryNameURL=PUBLIC.ADB_MTD
5	US001	LOCAL	621100	100000.000	100000.000	100000.000	USD http://slc11ohq.us.oracle.com:8000/pspt55doc11x/CUSTOMER/EMPL/q/?ICAction=ICQryNameURL=PUBLIC.ADB_MTD

4. Confirm that you arrive at the desired page with the appropriate data loaded as expected.

Chapter 5

Creating Query and Connected Query Search Definitions

Creating Search Definitions

To create a search definition:

1. Select PeopleTools, Search Framework, Search Designer Activity Guide, Define Search Definitions.
2. Click Add a New Value.
3. On the Add New Search Definition page, enter a name for the search definition in the Search Definition field.
4. Select the appropriate value from the Source Type drop-down list.
 - *Query\Connected Query*: requires an existing PeopleSoft query\connected query that defines the scope of the application data to be indexed.
 - *File Source*: enables you to index files stored within an accessible file system.
 - *Web Source*: enables you to index content within the structure of a website.
 - *Activity Guide*: enables you to index activity guides.
 - *External Data Source*: enables you to index external data for integrating with PeopleSoft.
5. Click Add.
6. Complete the required settings for the search definition source type, and save the search definition.

Specifying General Settings

Access the General page, by selecting PeopleTools, Search Framework, Search Designer Activity Guide, and from the left panel choose Search Definition.

Alternatively, select PeopleTools, Search Framework, Designer, Define Search Definitions.

This example illustrates the fields and controls on the General Settings page. You can find definitions for the fields and controls later on this page.

General

Map Search Attributes

Attachment Properties

Security

Advanced

Component Mapping

General Settings

Search Definition: PTPORTALREGISTRY
 Description: Navigator
 Object Owner ID:

Source Information

Source Type: Connected Query

*Source Name: View

Search Result Field Mapping

*Last Modified Date Time: Q

*URL Link: Q

*Image URL: Q

Title

%QUERYFIELD:PTPORTALREGISTRY:A.PORTAL_LABEL%

Title

Summary

%QUERYFIELD:PTPORTALREGISTRY:A.DESCR254%

Summary

Field or Control	Description
Search Definition	Displays the search definition's name as specified when adding a new value. <hr/> Note: At least one search definition must use the same name as the search category to which it belongs.
Description	Provide any additional information to distinguish the search definition.

Field or Control	Description
Source Type	Displays the type of query used to define the searchable data. Options are: Query or Connected Query.
Source Name	Displays the name of the base Query or Connected Query.
View	<p>Click to view the underlying query or connected query in either Query Manager or Connected Query Manager, respectively. This helps you to better understand the data being indexed by the search definition. When viewing the query or connected query from the PeopleSoft Search Framework, keep these items in mind:</p> <ul style="list-style-type: none"> • To view the query or connected query, you must have the appropriate security permissions in place. • You cannot make changes to the query or connected query when accessing it from the PeopleSoft Search Framework. The View button is intended for informational purposes only.
Last Modified Date Time	Select the field specified in the query that determines the underlying record's update date and time. For example, LASTUPDDTTM.
URL Link	<p>Select the query field containing the drilling URL defined for the query.</p> <hr/> <p>Note: Ensure that the drilling URL is unique.</p> <hr/>
Image URL	Select the Drilling URL pointing to the image field, as defined in the drilling URL properties for the query. This field appears only if an image URL has been associated with the query.
Title	<p>Define the title of the search result document. This is the bold, first line of the search result.</p> <p>Both Title and Summary should be granulated enough to be unique (that is, no duplicates among different documents).</p> <p>You can add text and bind variables to the title by clicking the *Title button. You cannot add text to the edit box directly.</p> <hr/> <p>Note: The title must be unique, that is, no duplicates among different documents.</p> <hr/> <p>See Adding Text and Bind Variables for Title and Summary</p>

Field or Control	Description
Summary	<p>Define the summary, or body, of the search result text. This is the text that appears below the title in a search result document.</p> <p>You can add text and bind variables to the summary by clicking the *Summary button. You cannot add text to the edit box directly.</p> <hr/> <p>Note: The summary must be unique, that is, no duplicates among different documents.</p> <hr/> <p>See Adding Text and Bind Variables for Title and Summary</p>
Save As	<p>Click to clone the search definition. In the subsequent page, you need to enter a name for the cloned search definition.</p> <p>When you clone a search definition, all the metadata associated with the existing search definition is saved in the cloned search definition.</p> <hr/> <p>Note: The Save As button is not available when you create or define a new search definition.</p> <hr/>

Adding Text and Bind Variables for Title and Summary

In the Search Result Title and Search Result Summary pages, you can insert custom text and bind variables, or a combination of the text and bind variables to accurately express the content of the search result document. You can insert multiple bind variables and bind variable types within the title and summary text, as needed.

Note: In a multi-language situation, only the base language structure for Title/Summary is respected. You cannot change the query fields to be displayed for different languages.

This example illustrates the fields and controls on the Search Result Title page. You can find definitions for the fields and controls later on this page.

Search Result Title ×

Query Fields

Query Field

Message Catalog

Message Set Number

Use Explain Text

Message Number

System Variables

System Variables

```
%QUERYFIELD:PTPORTALREGISTRY:A.PORTAL_LABEL%
```

<i>Field or Control</i>	<i>Description</i>
Query Field	Use the Query Fields drop-down list to select the desired fields from the underlying query of the search definition.
Message Set Number	Specify the Message Set Number to identify the message set to use.
Message Number	Specify the Message Number to identify the specific message to display.
Use Explain Text	By default, the message catalog text (that is, message text of the message catalog entry) is displayed. To use the message catalog explanation (that is, the description of the message catalog entry), select the Use Explain Text check box.
System Variables	Use the System Variables drop-down list to select the desired variable to display.

Field or Control	Description
Add	Select Add to form the expression.

The Search Result Summary page contains similar fields.

Mapping Search Attributes

Access the Map Search Attributes page, by selecting PeopleTools, Search Framework, Search Designer Activity Guide, Search Definition and selecting the Map Search Attributes tab.

The Map Search Attributes page displays all of the query fields of the query or connected query associated with the search definition.

This example illustrates the fields and controls on the Search Attribute Mapping page. You can find definitions for the fields and controls later on this page.

General Map Search Attributes Attachment Properties Security Advanced Component Mapping								
Search Attribute Mapping								
Search Definition		PTPORTALREGISTRY						
Description		Navigator						
Fields Included in the Index								Personalize Find First
Query Name	Query Field Name	Field to Index	Attribute Name	Attribute Display Name	Is Faceted	Is Hierarchy	Hierarchy Path	
1	PTPORTALREGISTRY	A.PORTAL_NAME	<input checked="" type="checkbox"/>	PORTAL_NAME	Portal Name	<input type="checkbox"/>	<input type="checkbox"/>	
2	PTPORTALREGISTRY	A.PORTAL_OBJNAME	<input type="checkbox"/>			<input type="checkbox"/>	<input type="checkbox"/>	
3	PTPORTALREGISTRY	A.PORTAL_LABEL	<input checked="" type="checkbox"/>	PORTAL_LABEL	Portal Label	<input type="checkbox"/>	<input type="checkbox"/>	
4	PTPORTALREGISTRY	A.DESCR254	<input checked="" type="checkbox"/>	DESCR254	Long Description	<input type="checkbox"/>	<input type="checkbox"/>	
5	PTPORTALREGISTRY	A.OPRID	<input type="checkbox"/>			<input type="checkbox"/>	<input type="checkbox"/>	
6	PTPORTALREGISTRY	A.PORTAL_EFFDT	<input type="checkbox"/>			<input type="checkbox"/>	<input type="checkbox"/>	
7	PTPORTALREGISTRY	A.PORTAL_CREATION_DT	<input type="checkbox"/>			<input type="checkbox"/>	<input type="checkbox"/>	
8	PTPORTALREGISTRY	A.PORTAL_EXPIRE_DT	<input type="checkbox"/>			<input type="checkbox"/>	<input type="checkbox"/>	
9	PTPORTALREGISTRY	A.PORTAL_EXPIRES	<input type="checkbox"/>			<input type="checkbox"/>	<input type="checkbox"/>	
10	PTPORTALREGISTRY	A.LASTUPDDTTM	<input type="checkbox"/>			<input type="checkbox"/>	<input type="checkbox"/>	
11	PTPORTALREGISTRY	A.PORTAL_CREF_URLT	<input type="checkbox"/>			<input type="checkbox"/>	<input type="checkbox"/>	
12	PTPORTALREGISTRY	A.PORTAL_URLXT	<input type="checkbox"/>			<input type="checkbox"/>	<input type="checkbox"/>	
13	PTPORTALREGISTRY	A.PORTAL_BASE_URI	<input type="checkbox"/>			<input type="checkbox"/>	<input type="checkbox"/>	
14	PTPORTALREGISTRY	A.PORTAL_CNTPRV_NAM	<input type="checkbox"/>			<input type="checkbox"/>	<input type="checkbox"/>	
15	PTPORTALREGISTRY	A.NODE_ITEMNAME	<input type="checkbox"/>			<input type="checkbox"/>	<input type="checkbox"/>	
16	PTPORTALREGISTRY	A.PT_URLTEXT	<input type="checkbox"/>			<input type="checkbox"/>	<input type="checkbox"/>	
17	PTPORTALREGISTRY	A.PORTAL_NODENAME	<input type="checkbox"/>			<input type="checkbox"/>	<input type="checkbox"/>	
18	PTPORTALREGISTRY	A.PORTAL_URLTEXT	<input type="checkbox"/>			<input type="checkbox"/>	<input type="checkbox"/>	
19	PTPORTALREGISTRY	PORTAL_DRILLURL	<input type="checkbox"/>			<input type="checkbox"/>	<input type="checkbox"/>	

Field or Control	Description
Field to Index	<p>Select the fields that you want to be indexed by the search engine. These fields would be those that you intend the end users would include in search queries intuitively. Fields that you do not select are not indexed. The remaining columns in the grid become enabled only after selecting the Field to Index check box for a particular row.</p> <p>When selecting fields also consider fields that may help to group results or help distinguish security access.</p> <p>Consider the following when you select fields for indexing:</p> <ul style="list-style-type: none"> • Number fields which are marked for indexing can be searched using the advanced search available on a component search, but cannot be searched using the global search. • For date and number fields that have multiple values, the condition between the values in the same document is always an OR condition. For example, if ST_DT>2010/01/01 and ST_DT<2000/01/01 and if ST_DT is a field with multiple values, then the effective condition will be ST_DT>2010/01/01 or ST_DT<2000/01/01, that is, documents which meet any of the conditions will be returned. <hr/> <p>Note: This does not impact single value fields.</p>
Attribute Name	<p>Displays the pre-defined attribute name that identifies the query field in the search engine.</p> <hr/> <p>Note: The attribute name must be unique in a search definition.</p> <hr/> <p>An Attribute Name is formed by concatenating the field name with its field label, unless the label is the default. If the field label is the default, just the field name is used for the attribute name. For example, if the field is VENDOR and the Display Name chosen is the default, then the attribute name is VENDOR. If the Display Name chosen is SUPP1 and it is not the default, then the attribute name is VENDOR_SUPP1.</p> <p>Due to this, there might be conflict between two fields which end up having same attribute name, and in these cases it is recommended to create a new label to make the attribute unique.</p>

Field or Control	Description
Attribute Display Name	<p>Determines how the attribute will display to the end user in the search results and in Kibana. Kibana uses the attribute display name as an alias when displaying the fields of a search definition for a visualization. Therefore, the attribute display name must be unique in a search definition.</p> <hr/> <p>Note: The display name comes from the field labels associated with the Query Field Name. It cannot be entered manually. If the wording of the display name is not appropriate, you need to add a new field label. If you select a field that is not the default field label, the Attribute Name will change to reflect your choice. If you change the field labels for existing search definitions, you must perform an update operation for the search definition on the Deploy Search Definition page.</p> <hr/> <p>Note: When creating search definitions be careful not to use the same Attribute Display Names that are used by search record key, alt keys, and list box items, unless their attribute name also matches the search record field name. Using the same display names affects how (if) fields appear to the end user on the Advanced search page.</p>
Is Faceted	<p>Select to make this field a facet field, meaning that this field can be used to categorize and narrow down search results based on its value.</p> <p>Defining a field as a facet requires some consideration. Facets cannot be blank, so if a field is not required, then the query needs to be structured so that there are not any blank values. You can use defaults on the record, or build COALESCE statements on a view to populate a field with a default (such as 'None', "NA" or "Blank") if it contains a blank value.</p> <p>If a facet does not have a value, then the query is dynamically modified to return <i>No Value</i> as the title of the facet. This processing is performed only on character attributes; not on long and XLAT attributes.</p> <hr/> <p>Note: Facets based on XLAT fields should be changed to use XLAT Long/Short in the query.</p> <hr/> <p>Note: There is a data size limit of 2000 characters for faceted fields. The system truncates anything beyond 2000 characters.</p> <hr/> <p>Important! Values with data types of datetime cannot be used as facets.</p>

Field or Control	Description
Is Hierarchy	<p>(Appears only for Connected Query.)</p> <p>This parameter is used to define a “soft hierarchy” which is programmed in the query.</p> <p>Soft hierarchy or data driven hierarchy means the data in a search attribute can be separated by a slash (/) to specify hierarchy in the data, for example, USA/California/Pleasanton. Also, the search attribute must be selected as a facet.</p> <p>This is different from the "hard" hierarchy discussed in Hierarchy Path.</p>
Hierarchy Path	<p>(Appears only for Connected Query.) Displays the hierarchy path for a Connected Query, indicating where in the hierarchy of connected queries a particular query field resides. When defining a hierarchy, you start with the highest level (most general) field at the top and sequentially list more granular fields, with the most granular being at the bottom.</p> <hr/> <p>Note: This is mainly used for hierarchical facets. The label in the user interface is determined by the field on which the hierarchy is defined.</p> <hr/> <p>This is considered a "hard" (predefined) hierarchy.</p>

Example: Define Hierarchy Path

This is an example of defining a hierarchy path.

This example illustrates the fields and controls on the Define Hierarchy Path page. You can find definitions for the fields and controls later on this page.

Define Hierarchy Path [Close]

[Help](#)

Query Name: QE_PERS_DATA

Source Field: A.QE_BIRTHCOUNTRY

Hierarchy Separator: /

Hierarchy			Personalize	Find	First	1-3 of 3	Last
Sequence number	*Related Query	Query Field					
1	QE_PERS_DATA	A.QE_BIRTHCOUNTRY	[+]	[-]			
2	QE_PERS_DATA	A.QE_BIRTHSTATE	[+]	[-]			
3	QE_PERS_DATA	A.QE_BIRTHPLACE	[+]	[-]			

OK Cancel

This definition is then displayed in the Hierarchy Path column of the Fields Included in the Index grid.

This example illustrates an example of hierarchy path.

```
A.QE_BIRTHCOUNTRY/A.QE_BIRTHSTATE/A.QE_BIRTHPLACE
```

Working with Attachment Properties

This section provides an overview of attachment processing and the attachment properties supported by PeopleSoft Search Framework.

Understanding Attachment Processing in Elasticsearch


PeopleSoft Search Framework downloads any attachment specified in a search definition, and pushes the encoded attachment data in the form of JSON document to the search engine. PeopleSoft Search Framework uses cURL handlers to transfer attachment data directly to the search engine bypassing the Integration Gateway. Search Framework uses libcurl library to communicate with the search engine. The transfer of data happens asynchronously, and each cURL handler can transfer data in parallel. You can specify the number of handlers on the Search Options page. The Schedule Search Index page displays an error if any data transfer does not proceed successfully. For more information, see [Managing General Search Options](#) and [Working with Search Indexes](#).

The search engine uses the ingest-attachment plug-in to extract the attachment contents. The attachment data is then indexed in the search engine.

Currently, the PeopleSoft Search Framework supports these file attachment storage location options:

- PeopleSoft database
- FTP
- SFTP
- HTTP

The following graphic depicts PeopleSoft downloading attachment from database and attachment repository, and then pushing encoded data to Elasticsearch. Elasticsearch uses Mapper attachment plug-in to extract attachment contents for indexing.




Specifying Attachments in a Search Definition

Access the Attachments page by selecting PeopleTools, Search Framework, Search Designer Activity Guide, Search Definition and selecting the Attachment Properties tab.

You use the Attachments page to select options related to the attachments that will be indexed as part of your search definition.

This example illustrates the fields and controls on the Attachment Properties page. You can find definitions for the fields and controls later on this page.



Field or Control	Description
Attachments	<p>The Attachments grid contains attachments you specify to be indexed as part of your search definition. These are the types of attachments supported:</p> <ul style="list-style-type: none"> • Attachment URL • Static URLID • Dynamic URLID <p>The options for each attachment type are described later in this topic where each attachment type is discussed.</p>
Edit	Click to edit the properties of a previously entered attachment.
Add Attachment	Click to add a new attachment to the Attachments grid.

Important! If the attachment URL properties or the URL of the attachment storage location changes, the PeopleSoft Search Framework recognizes the changes, and during the next index build or index update, the system will update the deployed search definitions to reflect the updated attachment properties and URL information.

Related Links

[Creating the Attachment URL ID List](#)

Specifying Attachment URL Properties

An Attachment URL type of attachment refers to attachments that are linked to the search definition by way of an Attachment URL type of drilling URL, specified on the Expression tab of PeopleSoft Query.

When using the Attachment URL type of attachment, the system uses the call-back credentials, defined for the search engine instance, to access the location of the attachment files. For example, if your attachments are stored in an FTP server, the FTP server needs to have the search engine call-back credentials configured as a valid user.

Note: The drilling URLs are different depending on the attachment repository type (database, FTP, and so on), but in all cases the system uses call-back user credentials to access the files.

Access the properties for this type of attachment by clicking the Add Attachment or Edit button on the Attachments page.

This example illustrates the fields and controls on the Define Attachment Properties: Attachment URL. You can find definitions for the fields and controls later on this page.

Define Attachment properties

Attachment Type

Attachment URL
 URL Id
 Inline Attachment

Query Name

Attachment URL

Attachment URL

<i>Field or Control</i>	<i>Description</i>
Attachment URL	Select if the attachment is linked to the search definition through a drilling URL in the underlying query as specified on the Expression tab in PeopleSoft Query.
Query Name	Select the query in which the Attachment URL drilling URL is defined.
Attachment URL	Select the query field in which the drilling URL query expression exists.

Specifying Static URL ID Attachment Properties

A Static URL ID type of attachment refers to an attachment that is associated with a URL ID that is static, or unlikely to change. You select the URL ID from the URL list defined by the Attachment URL ID List (PeopleTools, Search Framework, Search Administrator, Attachment URL ID List). For a Static URL ID type, the URL is generated from a URL ID (static) and a constant file reference path. This makes the URL always constant, and the file name changes based on the value received from the query field.

This example illustrates the fields and controls on the Define Attachment Properties: Static URL. You can find definitions for the fields and controls later on this page.

Define Attachment properties

Attachment Type

Attachment URL
 URL Id
 Inline Attachment

Query Name

URL Id

URL Id Type

URL Identifier

File Reference Path

File Name Query Field

<i>Field or Control</i>	<i>Description</i>
URL ID	Select if your attachment is associated with a URL ID, not an Attachment URL expression defined in PeopleSoft Query.
Query Name	Select the query name in which this attachment is referenced.
URL ID Type	If the URL ID is constant and unlikely to change, select <i>Static</i> .
URL Identifier	Select the URL Identifier from the drop-down list. This list is populated by the list defined by the Attachment URL ID List (PeopleTools, Search Framework, Search Admin Activity Guide, Attachment URL ID List). These URLs are a subset of the URLs defined in the PeopleSoft database, and this subset are those URLs identified to be used with the PeopleSoft Search Framework.
File Reference Path	(Optional) Enter a file reference path for mapping a dynamic file path.
File Name Query Field	Select the query field that will contain the name of the attached file.

Specifying Dynamic URL ID Attachment Properties

With a dynamic attachment type, all the components of a URL are generated from the value in query field. With a dynamic attachment type, much of the information related to the attachment is unknown prior to the transaction occurring.

For example, with a static attachment type, it is known beforehand what the URLID will be, so it can be specified when the attachment properties are set. With a dynamic attachment type, the values of the attachment properties are determined at the transaction time, and the property values, such as URLID, file reference path, and attachment file name, are stored in the row of data associated with the attachment. The attachment property values can vary between rows.

This example illustrates the fields and controls on the Define Attachment Properties: Dynamic URL. You can find definitions for the fields and controls later on this page.

Define Attachment properties

Attachment Type

Attachment URL
 URL Id
 Inline Attachment

Query Name

URL Id

URL Id Type

*URLID Query Field

File Path Query Field

File Name Query Field

Field or Control	Description
Query Name	Select the query name in which this attachment is referenced.
URL Id Type	If the URL ID is not constant and likely to change, select <i>Dynamic</i> .
URLID Query Field	Select the field in the query that will hold the URL ID of the attachment storage location.
File Path Query Field	Select the field in the query that will hold the URL ID of the attachment file reference path.

<i>Field or Control</i>	<i>Description</i>
File Name Query Field	Select the query field that will contain the name of the attached file.

Specifying Inline Attachment Properties

Important! The Inline Attachment option is deprecated, and this option should not be used. However, the Inline Attachment option is retained for backward compatibility. PeopleSoft Search Framework supports the Inline Attachment option for Elasticsearch in the event that you use a SES-based search definition that contains attachments specified as inline attachments.

Managing Search Definition Security

Access the Security page by selecting PeopleTools, Search Framework, Search Designer Activity Guide, Search Definition and selecting the Security tab.

This example illustrates the fields and controls on the Security page. You can find definitions for the fields and controls later on this page.



The Security page enables you to restrict access to search results generated by a search definition. Depending on the sensitivity of the search results, you can set these degrees of security:

<i>Field or Control</i>	<i>Description</i>
No Security	Select to define no security restriction for a search definition's search results. Anyone with access to the application can view the search results for a search definition set to No Security. That is, the search results are public to all users.
Source Level Security	Select to allow or restrict access to the entire search definition as per the specified user or role. That is, only specified users and roles are able to view search results for that search definition.

Field or Control	Description
Document Level Security	Select to restrict access to specific search results (documents) generated by a search definition. That is, with document level security, users can view search results generated by that search definition, but only the documents to which they have access. Note: This is generally referred to as row-level security in PeopleSoft applications.

Setting Source Level Security

Access the source-level security settings by selecting the Source Level Security radio button.

This example illustrates the fields and controls on the Security page - Source Level Security options. You can find definitions for the fields and controls later on this page.

The screenshot shows the 'Security Settings' page for search definition 'ITEMDEFN'. The 'Security Option' section has 'Source Level Security' selected. The 'Source Level' section contains a 'User/Role Mapping' table with the following data:

*Type	Name	*Privilege
Role	QE_CECLRK_ROLE	Allow
User	QETSTR	Allow

Field or Control	Description
Type	Select <i>Role</i> or <i>User</i> depending on the scope of your intended access restriction. Selecting <i>Role</i> restricts access to a specific PeopleSoft role. Selecting <i>User</i> restricts access to a specific PeopleSoft user.
Name	Select the user or role name.
Privilege	Define the access privilege or restriction. <ul style="list-style-type: none"> Allow. The specified role or user is allowed to view search results for this search definition. Deny. The specified role or user is not allowed to view search results for this search definition.

Note: Source-level security applies to every document within that search definition for the search engine.

Setting Document Level Security

Access the document-level security settings by selecting the Document Level Security radio button. Document-level security can also be thought of as attribute-based security.

With document-level security, one or more PeopleSoft Query columns act as the security attribute. The security attribute is cached on the search engine for each document. Only those users who have the correct security attribute will have the document returned. You can then specify an application class (AppClass) that returns a list of security values for the current user. The Appclass is called when the user submits the search request. When the user submits the search request, the search engine compiles a list of values returned from the application class associated with that specific user to build the security filter.

This example illustrates the fields and controls on the Security page: Document Level Security option. You can find definitions for the fields and controls later on this page.

Document Level

<i>Field or Control</i>	<i>Description</i>
Query Name	Select the name of the query or connected query containing the fields you want to use to restrict access.

Field or Control	Description
Source Field	<p>Select the field which will identify security values which will determine access.</p> <p>The source field(s) selected becomes the security attribute having the specified privilege. At indexing time, when the crawler inserts application data into the index, the values populating the selected source field(s) will carry the specified privilege.</p>
Privilege	<p>Define the access privilege or restriction.</p> <ul style="list-style-type: none"> • Allow. The value specified in the Source Field is allowed access to the data for this row in the search results. • Deny. The value specified in the Source field cannot see the data for this row in the search results. <hr/> <p>Note: The privilege of Deny is useful in situations where there are too many values for the security attribute if Allow were selected. For example, rather than enabling access to nine out of ten field values, it is more efficient to deny access only to the one you want to restrict.</p> <hr/> <p>Note: If multiple attributes appear in the grid the system effectively inserts an AND clause between the items in the grid.</p> <hr/>

Document Filter App Class

The application class specified in the Document Filter App Class section creates a list of values for the specific user performing the query. The application class enables you to define and run additional filters and logic against the application data contained in the indexed source fields.

As needed, PeopleSoft applications will provide filtering App Classes for delivered search definitions. For any custom search definitions, or additional filtering requirements, you will need to create or modify the filtering App Classes.

Field or Control	Description
Package Name	Select the name of the appropriate App Package.
Path	Select the path pointing to the App Class.
Class ID	Select the class ID for the App Class.

For example, assume you want to compile a list of valid setIDs to which the user may have access. You define an application package that would contain a method called `evaluateAttrValues`. This passes the search definition name (`sboName`), the name of the security attribute field that was identified in the search definition, and the user who will access the data. The application method would then build a list of valid values for that user and return it to search engine for comparison against the data to see if the attribute value on the data matches.

```
method evaluateAttrValues
  /+ &sboName as String, +/
  /+ &secAttr as String, +/
  /+ &srchUser as String +/
  /+ Returns Array of String +/
  /+ Extends/implements PTSF_SECURITY:SearchAuthnQueryFilter.evaluateAttrValues +/
  Local array of string &secValues;
  Local string &Role, &userPref, &csFullAccess, &csAdminAccess, &OnBehalfOf,
  &docOwner, &BU_Security, &Security_Type, &SID_Security, &PermList;
  Local SQL &sqlRoles, &sqlUserPrefs, &sqlDocOwners;

  &secValues = CreateArrayRept("", 0);

  /*&BU_Security, &Security_Type, &SID_Security*/
  SQLExec("Select SETID SECURITY, SECURITY_TYPE
  from PS_INSTALLATION_FS ", &SID_Security,
  &Security_Type);
  If &SID_Security = "N" Then
    &secValues.Push("A:ALL");
  Else
    Evaluate &Security_Type
    When "N"
      &secValues.Push("A:ALL");
    When "O"
      &secValues.Push("U:" | &srchUser);
    When "C"
      SQLExec("Select OPRCLASS from PSOPRDEFN where OPRID = :1",
      &srchUser, &PermList);
      &secValues.Push("P:" | &PermList);
    End-Evaluate;
  End-If;
End-method;
```

Note: Concatenating multiple attribute values using a separator, you can achieve an OR clause between attributes. For example, in the this sample "A:ALL" and "U:" | &srchUser are two different attributes merged into a single attribute to achieve the OR clause.

Note: To achieve improved performance, the security attribute should be chosen in such a way that no more than 50 values are returned per user per attribute.

Working With Advanced Settings

Access the Advanced page by selecting PeopleTools, Search Framework, Search Designer Activity Guide, Search Definition and selecting the Advanced tab.

This example illustrates the fields and controls on the Advanced Properties page. You can find definitions for the fields and controls later on this page.

General
Map Search Attributes
Attachment Properties
Security
Advanced
Component Mapping

Advanced Properties

Search Definition ITEMDEFN

Description Item Definition

Define Query to Delete SBO

Query Name

URL Link Field

Pre Processing AE Library

AE Library

AE Section

Post Processing AE Library

AE Library

AE Section

Encode Search Attributes

	Search Attribute	Attribute Display Name		
1	<input style="width: 90%;" type="text"/>	<input style="width: 90%;" type="text"/>	+	-

Full Indexing Criteria

Index Only Last Days Index Start Date 01/01/1900

Alternate search keywords

Define Query to Delete SBO

Field or Control	Description
Query Name	<p>Select the query that you've defined to remove orphaned search documents. This is your deletion query.</p> <p>For example, for rows of data that have been deleted from the database, you would want to ensure that those rows of data (or search result documents) no longer appear in the search index.</p> <hr/> <p>Note: Even though this query is defined the same way as the source type query is defined, the difference is that the result of the query is used for deleting the already indexed documents from the search engine. The Drilling URL field acts as the primary identifier to locate the documents to be deleted. This means the drilling URL must be identical to the drilling URL built for the original query.</p> <hr/> <p>Delete SBO Query will not be executed during the first index build of a search definition. It only becomes active after the initial index is populated.</p> <p>See Defining a Deletion Query.</p>
URL Link Field	Select the drilling URL field for the query selected.

Pre Processing AE Library

Before an index build is run, data can be manipulated to perform additional tasks. For example, it may be required to search through subqueries to determine if one or more subquery might have changed and to update the last updated datetime value on the parent so that the change is recognized. Another use might be to calculate or summarize a value to be indexed. Another common case for pre processing occurs when a query is created based on staging tables. In this case the pre processing Application Engine program can be used to populate the staging tables. This is useful when the data cannot be retrieved using a simple query (such as hierarchal or computed data).

You can define an Application Engine routine to perform these tasks.

Note: Pre and post processing adds to the overall indexing time, and it is recommended not to use this in scenarios where it can be avoided.

Note: Application Engine programs used for pre or post processing should be defined as type Library.

Field or Control	Description
AE Library	To call additional processing logic from Application Engine, specify the correct AE Library (Application Engine program) to run.
AE Section	Specify the Application Engine program section to run.

Post Processing AE Library

After an index is built, certain clean-up functions might need to be performed, such as removing the records in your delete query, or cleaning up staging tables. Use the post processing Application Engine routine to perform these functions.

Field or Control	Description
AE Library	To call additional processing logic from Application Engine, specify the correct AE Library (Application Engine program) to run.
AE Section	Specify the Application Engine program section to run.

Encode Search Attributes

Use the encode option in place of setting up custom lexers for applications that have indexed code-based fields for search attributes. Encoding these values enables the system to process the codes when generating the feed for the search index without requiring staging tables or a pre-processor Application Engine program. Encoding keeps words intact for indexing and searching that otherwise might be split during the parsing process of search index generation.

For example, a code-based search field, expects the end user to enter a code, known to the application users, that may not be interpreted as a single value by the search engine. In which case, the search engine may index the code as multiple values, losing the meaning of the code. For example, perhaps there is a product version represented by the code "M-2000". If the search engine indexes this code as two separate values, "M" and "2000" rather than "M-2000," the code loses its meaning and it no longer serves as a useful search attribute. In this case, if added to the Encode Search Attributes grid, PeopleSoft Search Framework encodes the "M-2000" value (to Mx2000 for example) so that it is indexed as a single value. While the value may be encoded for the search engine, the end user enters the code when searching as normal, requiring no knowledge of how the code is represented for the search engine indexing purposes.

If encoding search attributes, keep these items in mind:

- Encoded attributes are not exposed to the end user. The system uses encoded values only for programmatic operation.
- The search engine has a token limit of 64 bytes, so fields with more than 30 characters cannot be encoded. In a hierarchical value, use a word separator character to ensure the end value does not exceed 64 bytes.

- For these values, the system encodes the search string using EncodeSearchCode PeopleCode before passing the filters to the search API.
- For indexes containing encoded values, there can be slight overhead related to index size and query performance.

Note: Do NOT use encoded search attributes unless you are prepared to customize the underlying code. Encoded search attributes should never be displayed on a page, therefore, be careful what attributes you encode.

Full Indexing Criteria

<i>Field or Control</i>	<i>Description</i>
Index Only Last <> Days	<p>Enter the number of days for which you want to build the index. The number of days entered are treated as days prior to the current date.</p> <p>The default value is 0.</p>
Index Start Date	<p>The Index Start Date is calculated based on the number of days you enter in the Index Only Last <> Days field. The Index Start Date displays the date after deducting the number of days entered in the Index Only Last <> Days fields from the current date. The index will be built from the date that is displayed in the Index Start Date field.</p> <p>When the value of the Index Only Last <> Days field is 0, which is the default value, the Index Start Date is set to 01/01/1900.</p>

Alternate Search Keywords

<i>Field or Control</i>	<i>Description</i>
Alternate search keywords	<p>Enter any alternate search keywords that users are likely to submit when running search requests. For example, rather than submitting the more official "purchase order" phrase, a user may be more likely to search on "PO" instead.</p> <p>This field is translatable. If supporting multiple languages, the keywords need to be translated.</p>

Mapping Components to Search Definitions

Access the Component Mapping page by selecting PeopleTools, Search Framework, Search Designer Activity Guide, Search Definition and selecting the Component Mapping tab. The Component Mapping page enables you to map a component to a specific search definition to enable the search engine integration with Search Pages.

Mapping a component to a search definition effectively assigns component level security to the search definition. For this reason, if the Search Definition is assigned to a component, document level security is turned on by default.

Note: Multiple components can be mapped to the same search definition. When a component is mapped, the security type is changed to document level security, and a user having access to any one of the components will be able to search inside the search index. The security type cannot be changed, but you can add more security attributes to extend the security.

Important! Search definitions mapped to a component must use document level security.

This example illustrates the fields and controls on the Component Mapping page. You can find definitions for the fields and controls later on this page.

*Market	*Component Name	Description	Search Criteria
1 Global	MERCHANTCATTBL		

Field or Control	Description
Market	Select the Market for the component.
Component Name	Select the component.

Field or Control	Description
Search Criteria	<p>Enter any additional custom search criteria to be added automatically to the search engine search query during a Search Pages keyword search.</p> <p>The search criteria should take the form <code>ATTRIBUTENAME:VALUE</code> and can use keywords and brackets. For example:</p> <p><i>PTSF_SBO_NAME:EP_AP_VENDOR & (-STATUS:PND)</i></p> <hr/> <p>Note: The system appends additional search criteria as is to the search query. The Search Criteria field is "free text." The Search Framework performs no validation or parsing during design time or run time.</p>

Viewing Search Attributes

Access the Search Attributes page by selecting PeopleTools, Search Framework, Search Designer Activity Guide, View Search Attributes and selecting the Display Fields tab. Search attributes are used for creating aliases of the actual search query fields from PeopleSoft Query or Connected Query in the Oracle RSS feeds.

Only these three data types are supported:

- String
- Number
- Date

PeopleTools maps the PeopleSoft data types to the search engine data types. The Search Attributes page displays a search attributes mapped search engine data type.

This example illustrates the fields and controls on the Search Attributes page. You can find definitions for the fields and controls later on this page.

Search Attributes				
Search Attribute	Field Name	Attribute Display Name	Data Type	Related Search Definitions
1 ACCOMPLISHMENT	ACCOMPLISHMENT	Accomplishment	String	Related Search Definitions
2 ACCOUNT	ACCOUNT	Account	String	Related Search Definitions
3 ACCOUNTING_DT	ACCOUNTING_DT	Accounting Date	Date	Related Search Definitions
4 ACCOUNTING_PERIOD	ACCOUNTING_PERIOD	Accounting Period	Number	Related Search Definitions
5 ACCT_PERIOD_CHAR	ACCT_PERIOD_CHAR	Accounting Period	String	Related Search Definitions
6 ACQUISITION_DT_SES_ACQ_DT	ACQUISITION_DT	Asset Acquisition Date	Date	Related Search Definitions
7 ACQUISITION_DT_YR_H	ACQUISITION_DT_YR	Asset Acquisition Date	String	Related Search Definitions

The search attributes page is read-only and provides a quick view of all attributes currently authorized for sending to the search engine. Every time a field is mapped in a Search Definition, the system adds an entry to this grid for the field. Because a field can exist on more than one search definition, it is not generally removed from this grid.

Search attributes are retained separately from a search definition and this enables:

- **Reusability.** You can reuse the same search attributes amongst multiple search definitions.
- **Increased usability.** Search attributes enable you to modify the attribute name or query field name from the end user, who will see the attribute display name. For example, adding the display value of "Employee ID" is more readable and understandable than displaying the field name EMPLID.
- **Quick Access.** You can view and sort all search attributes in the system and check their data types without having to open the search engine or Application Designer. You can also view the search definitions associated with the search attribute by clicking the Related Search Definitions link.

Chapter 6

Creating File Source Search Definitions

Understanding File Source Search Definitions

File source definitions enable you to specify the location where external files reside in an accessible file system that you want to make available to the search engine crawling mechanism so that end users can search the content of these files to support PeopleSoft applications. Examples of external file types that can be indexed include Microsoft Word and Excel documents, plain text documents, HTML files, and so on.

Specifying File Source General Settings

Access General Settings page by selecting PeopleTools, Search Framework, Search Designer Activity Guide, Search Definition and selecting a Source Type of *File Source*.

Use the General Settings page to specify the location of the files to be indexed as well as the crawler settings.

This example illustrates the fields and controls on the File Source: General Settings page. You can find definitions for the fields and controls later on this page.

General
Document Types

General Settings

Search Definition QE_FILE_PIASRCH

*Description File type for piarsrch testing

Source Type File

Object Owner ID PeopleTools

Starting URLs (file://localhost) Personalize | Find | | | First 1-2 of 2 Last

Starting URL			
1	file://localhost/ds1/products/files/	...	+ -
2	file://localhost/mount/ptqa/ptqa_unix/PS_APP_HOME/PT854/PIASRCH_FILES/	...	+ -

Crawler Settings

Crawler Timeout (Seconds) 30 Default Language English

Max Document Size (MB) 10 Enable Language Detection

File URL Prefix file://localhost/ds1/products/

Display URL Prefix https://webhost/clients/

Field or Control	Description
Description	Add a brief description to help identify the purpose of the search definition.
Source Type	Displays the type of search definition, such as Query, Web, File, and so on.

Field or Control	Description
Starting URLs	<p>In the Starting URLs grid enter the location(s) in your file system where the files reside that you want to expose to the Search Framework crawling process. For each different location, add a new row to the grid.</p> <hr/> <p>Note: The starting URL is not case sensitive.</p> <hr/> <p>Note: For the search engine crawler to access the URL, the file starting URL must be fully qualified, as in <code>file://localhost/</code>.</p> <hr/> <p>On UNIX the starting URL format is:</p> <ul style="list-style-type: none"> To crawl local files use <code>file://localhost/<directory structure></code>. For example: <i>file://localhost/recruitment/resume/</i> To crawl from mounted file systems use <code>file://localhost//<mounted_dir_path></code>. For example: <i>file://localhost//dfs/recruitment/resume/</i> <p>On Windows the starting URL format is:</p> <ul style="list-style-type: none"> To crawl local files use <code>file://localhost/<Directory_Path></code>. For example: <i>file://localhost/D:/recruitment/resume/</i> To crawl from a mapped drive use <code>file://localhost//<machinename>/<shared_folder_path>/</code>. For example: <i>file://localhost//RTDC78067TLSBLD/recruitment/resume/</i> <p>A search engine can crawl files on directories located on the server where the search engine is installed or network file paths accessible by the search engine.</p> <p>When the search engine crawls files from a network drive, then the Oracle process/service should be started as a user who has access to the network drive, which you can accomplish by modifying the logon account of OracleServiceSID and OracleSIDTNSListener services to match the domain administrator and restart both services.</p> <p>In the Elasticsearch implementation, when the search engine crawls files from a network drive, then the Oracle process/service should be started as a user who has access to the network drive; the user must have at least read access to the network drive.</p>
Crawler Timeout	Indicates the maximum allowed time to retrieve a file for crawling.
Max Document Size	The maximum document size in megabytes that the system will crawl. Larger documents are not crawled.

Field or Control	Description
Enable Language Detection	<p>By Enabling Language Detection, the search engine automatically identifies the language of the document content and assigns the language code automatically.</p> <p>If the search engine crawler cannot determine a perfect match, it finds a best match from the trained set of languages and assigns. Otherwise the default language in the crawler configurations will be assigned.</p> <hr/> <p>Note: In the Elasticsearch implementation, automatic language detection is supported for all languages that are supported by PeopleSoft.</p> <hr/>
File URL Prefix	<p>The part of the access URL the system will not display in the search results due to security reasons.</p> <p>This is an optional feature where there is a need to hide the actual URL used for indexing. If a File URL Prefix is specified it is mandatory to have the Display URL Prefix set.</p>
Display URL Prefix	<p>The URL the system displays instead of the actual URL. For example, if the file URL is:</p> <p><i>file://localhost/home/operation/doc/file.doc</i></p> <p>and you want the display URL to appear as:</p> <p><i>https://webhost/client/doc/file.doc</i></p> <p>then specify the File URL Prefix as:</p> <p><i>file://localhost/home/operation</i></p> <p>and the Display URL Prefix as:</p> <p><i>https://webhost/client</i></p> <p>If you select Display URL Prefix, make sure that the files are reachable using the specified URL. The search engine crawler replaces the URL string specified for the File URL Prefix with the Display URL Prefix.</p> <hr/> <p>Note: When you select Display URL Prefix, you must consider the following: A File URL Prefix must be a fully resolved path; it should not be a symbolic link.</p> <hr/>

Specifying Document Types

Access the Document Types page by selecting PeopleTools, Search Framework, Search Designer Activity Guide, Search Definition and selecting a Source Type of *File Source*. Then click the Document Types tab.

Use the Document Types page to specify the document types, expressed as MIME types, of the documents to be crawled.

This example illustrates the fields and controls on the File Source: Document Types page. You can find definitions for the fields and controls later on this page.

The screenshot shows a web interface for configuring document types. At the top, there are two tabs: 'General' and 'Document Types', with 'Document Types' selected. Below the tabs, the title 'Document Types' is displayed. The search definition is 'QE_FILE_PIASRCH' and the description is 'File type for piarsrch testing'. A section titled 'Document types to crawl' contains four radio buttons: 'Include default types', 'Include all types', 'Add document types', and 'Exclude document types', with 'Exclude document types' selected. Below this is a table with two rows of document types: 'text/plain' and 'application/msword'. Each row has a search icon, a plus button, and a minus button. At the bottom, there are two empty boxes for 'Inclusion Rules' and 'Exclusion Rules' under the heading 'URL Boundary Rules'.

Document Types to Crawl

Field or Control	Description
Include default types	<p>Crawl only on the default document types, as defined by the search engine.</p> <p>The search engine default document types for crawling are:</p> <ul style="list-style-type: none"> • PDF • HTML • TXT (plain text) • Microsoft Word • Microsoft Excel • Microsoft PowerPoint <p>If no other document types are added to the Select grid, then the search engine considers only the default document types.</p>
Include all types	<p>Crawl all document types supported by the search engine.</p> <p>To see this list of supported document types, expressed as MIME types, select either the Add these types or Exclude these types radio button and click the lookup button for the Document Type column.</p>
Add these types	<p>In addition to the default document types, the system also crawls any document type added to the Document Type grid, with Add these types selected.</p>
Exclude these types	<p>Excludes specific document types added to the Document Type grid.</p> <p>Assume the majority of the document types supported by Elasticsearch crawler apply to your configuration, except for a small number of document types. In this case you can specifically include those document types in the Document Type grid. When the search engine crawls the file location, the search engine crawls all document types on the supported MIME list, except for those document types included in the Document Type grid, with Exclude these types selected.</p>

URL Boundary Rules

Field or Control	Description
Inclusion Rules	<p>Specify an inclusion rule that a URL must contain. For example:</p> <p><i>*es_xml*</i></p> <p>In this case, search engine crawls all documents with <i>es_xml</i> in the name.</p> <p>Specify an inclusion rule that URL must start with. For example:</p> <p><i>file://localhost/ds1/product/ES_ADD/es_doc*</i></p> <p>In this case, the search engine crawls all files starting with <i>file://localhost/ds1/product/ES_ADD/es_doc</i>.</p>
Exclusion Rules	<p>Specify an exclusion rule that a URL can't contain. For example:</p> <p><i>*.xml</i></p> <p>In this case, the search engine does not crawl anything with a <i>.xml</i> extension.</p>

URL boundary rules limit the crawling scope. When you add boundary rules, the crawler is restricted to URLs that match only the rules you specify. Inclusion and Exclusion rules can be formed to filter documents with patterns of *begins with*, *ends with*, *contains* or regular expressions. Rules with regular expression should start with the character *R*.

Rule	Description
Begins With: <i>file://localhost/example*</i>	In this case, the search engine considers URLs starting with <i>file://localhost/example</i> .
Ends With: <i>*.doc</i>	In this case, the search engine considers URLs ending with <i>.doc</i> .
Contains: <i>*contacts*</i>	In this case, the search engine considers URLs containing string <i>contacts</i> .
Regular Expression: <i>R.*es_html_lv[1-9].html</i>	In this case, the search engine considers URLs ending with numbers varying from 1 to 9 with file names ending with <i>es_html_lv</i> .

When working with these rules, keep in mind:

- Exclusion rules always override inclusion rules.
- Multiple inclusion and exclusion rules can be separated by a space or in a new line.
- Use an asterisk to represent a wildcard.
- Inclusion and exclusion rules are case-insensitive.

Creating Web Source Search Definitions

Understanding Web Source Search Definitions

Web source definitions enable you to make the content of internal or external websites available for search engine crawling and inclusion in the Search Framework end user searches. For example, if there is a website on your company's intranet that describes specific business processes for your end users, they can instigate a search against this web content from within the PeopleSoft application.

Specifying Web Source General Settings

Access the General Settings page by selecting PeopleTools, Search Framework, Search Designer Activity Guide, Search Definition and selecting a Source Type of *Web Source*.

Use the General Settings page to specify the location of the files to be indexed as well as the crawler settings.

This example illustrates the fields and controls on the Web Source — General Settings page. You can find definitions for the fields and controls later on this page.

The screenshot displays the 'General Settings' page for a search definition. At the top, there are tabs for 'General' and 'Document Types'. The main heading is 'General Settings'. Below this, the following fields are visible: 'Search Definition' with the value 'QE_WEB_SD', '*Description' with the value 'Web SD', 'Source Type' set to 'Web', and 'Object Owner ID' set to 'PeopleTools'. A section titled 'Starting URLs (http://)' contains a table with one row: '1 | http://www.unicode.org/standard/WhatsUnicode.html'. Below the table are 'Crawler Settings' with input fields for 'Crawler Timeout (Seconds)' (30), 'Max Document Size (MB)' (10), and 'Crawl Depth' (0). There are also radio button options for 'Honor Robots Exclusion' and 'Index Dynamic Pages', both of which are selected to 'Yes'.

Field or Control	Description
Description	Add a brief description to help identify the purpose of the search definition.

Field or Control	Description
Source Type	Displays the type of search definition, such as Query, Web, File, and so on.
Starting URLs	<p>Contains the URL of the web address. The search engine uses the URL as an entry point for starting to crawl a website.</p> <hr/> <p>Important! Only HTTP URLs are supported. The starting URL's mentioned should be accessible without any user credentials. The search engine crawler will ignore web sites requiring login.</p> <hr/>
Index All	This will index all the URLs which are allowed to access by the search engine crawler. This will not limit crawling in to a specific domain or host. As the number of URLs to index increases, time required to complete indexing also increases.
Stay in Host	<p>This will limit the indexing only to the specified host. For example, if you are indexing www.oracle.com and you select this option, you can index documents on www.oracle.com, but not on www.1.oracle.com.</p> <hr/> <p>Important! If neither option is selected, then the system switches to <i>Stay in Domain</i> mode. In this mode, indexing will be limited to a single domain. For example if you are indexing www.oracle.com it will consider all URLs with in this domain, including www.1.oracle.com, but URLs from a different domain, such as www.yahoo.com, would not be indexed.</p> <hr/>
Crawler Timeout	Indicates the maximum allowed time to retrieve a file for crawling.
Crawl Depth	<p>The number of nested links the crawler follows, with the initial URL, or home page, residing at a <i>depth of 0</i>.</p> <p>With a crawling depth of 1, the crawler also fetches any document linked to from the starting URL.</p> <p>With the crawling depth set to 2, the crawler fetches any document linked to from the starting URL (depth of 0), and also fetches any document linked to from the depth of 1, and so on.</p> <p>By adding a value for Crawl Depth, the system uses that value to enforce the crawling limit. If you enter no value, leaving the Crawl Depth blank, the system considers the crawling depth to be unlimited. As you increase the crawl depth, the content to be indexed can increase exponentially, which results in longer crawling durations.</p>
Max Document Size	The maximum document size in megabytes that the system will crawl. Larger documents are not crawled.

Field or Control	Description
Honor Robots Exclusion	<p>Robot exclusion policies are set at web server and the web page level. The Honor Robots Exclusion setting controls whether the search engine recognizes or ignores the robot exclusion settings.</p> <ul style="list-style-type: none"> • <i>Yes.</i> The crawler traverses the pages based on the access policy specified in the web server robots.txt file. The crawler also respects the page-level robot exclusion specified in HTML meta tags. • <i>No.</i> The crawler ignores any specified robot policy defined on the web server.
Index Dynamic Pages	<p>Controls whether search engine crawls and indexes dynamic pages. Typically, database applications serve dynamic pages, and the pages have a URL containing a question mark (?). The search engine considers URLs containing question marks dynamic pages.</p> <ul style="list-style-type: none"> • <i>Yes.</i> The search engine crawls dynamic pages. • <i>No.</i> The search engine does not crawl dynamic pages.

Specifying Document Types for Web Source

Access the Document Types page by selecting PeopleTools, Search Framework, Search Designer Activity Guide, Search Definition and selecting a Source Type of *Web Source*. Then click the Document Types tab.

Use the Document Types page to specify the document types of the documents to be crawled. Document types are expressed as MIME types.

This example illustrates the fields and controls on the Web Source — Document Types page. You can find definitions for the fields and controls later on this page.

General Document Types

Document Types

Search Definition QE_WEB_SD
Description Web SD

Document types to crawl

Include default types Include all types Add document types Exclude document types

Select	Document Type
1	

Personalize | Find | First 1 of 1 Last

URL Boundary Rules

Inclusion Rules

Exclusion Rules

.

Document Types to Crawl

Field or Control	Description
Include default types	<p>Crawl only on the default document types, as defined by the search engine.</p> <p>The search engine default document types for crawling are:</p> <ul style="list-style-type: none"> • PDF • HTML • TXT (plain text) • Microsoft Word • Microsoft Excel • Microsoft PowerPoint <hr/> <p>Note: When you select this option, notice that you cannot add or remove items from the Document Type grid.</p> <hr/> <p>Note: Exclusion rules take precedence. For example, if you enter *.doc in the Exclusion Rules edit box, none of the Microsoft Word documents will be processed.</p> <hr/>
Include all types	<p>Crawl all document types supported by the search engine.</p> <p>To see this list of supported document types, expressed as MIME types, select either the Add these types or Exclude these types radio button and click the lookup button for the Document Type column.</p> <hr/> <p>Note: When you select this option, notice that you cannot add or remove items from the Document Type grid.</p> <hr/>
Add document types	<p>In addition to the mandatory document types, which are HTML and TXT files, the system also crawls any document type added to the Document Type grid.</p>

Field or Control	Description
Exclude document types	<p>Excludes specific document types added to the Document Type grid.</p> <p>Assume the majority of the document types supported by the search engine crawler apply to your configuration, except for a small number of document types. In this case you can specifically include those document types in the Document Type grid. When the search engine crawls the file location, the search engine crawls all document types on the supported MIME list, except for those document types included in the Document Type grid.</p>

URL Boundary Rules

Field or Control	Description
Inclusion Rules	<p>Specify an inclusion rule that a URL must contain. For example:</p> <p><i>www.*.example.com</i></p> <p>In this case, the search engine crawls all content within <i>www.*.example.com</i>.</p>
Exclusion Rules	<p>Specify an exclusion rule that a URL can't contain. For example:</p> <p><i>www.*.uk.example.com</i></p> <p>In this case, the search engine does not crawl anything within <i>www.*.uk.example.com</i>.</p>

URL boundary rules limit the crawling scope. When you add boundary rules, the crawler is restricted to URLs that match only the rules you specify. Inclusion and Exclusion rules can be formed to filter documents with patterns of *begins with*, *ends with*, *contains* or regular expressions. Rules with regular expression should start with the character *R*.

Rule	Description
Begins With: <i>http://www.uk.example.com*</i>	In this case, the search engine considers URLs starting with <i>www.uk.example.com</i> .
Ends With: <i>*.xml</i>	In this case, the search engine considers URLs ending with <i>.xml</i> .

Rule	Description
Contains: <i>*contacts*</i>	In this case, the search engine considers URLs containing string <i>contacts</i> .
Regular Expression: <i>R^http://www.example.com/code.*./verson[1-9].html\$</i>	In this case, the search engine considers URLs from example.com with sites starting with string <i>code</i> , and has versions numbered from <i>1 to 9</i> and ends as <i>.html</i> .

When working with these rules, keep in mind:

- Exclusion rules always override inclusion rules.
- Multiple inclusion and exclusion rules can be separated by a space or in a new line.
- Use an asterisk to represent a wildcard.
- Inclusion and exclusion rules are case-insensitive.

Specifying Proxy Parameters in Elasticsearch

In Elasticsearch, if you want to crawl Web sites outside of your network, you need to specify the proxy parameters in the `elasticsearch.yml` configuration file.

To specify the proxy parameters:

1. Open the `elasticsearch.yml` configuration file, which is available at `<ES_HOME>/config`.
2. Locate the Web Crawler section, and enter your network proxy values for the following parameters:
 - `orcl.proxy.host` - Enter the network proxy host.
 - `orcl.proxy.port` - Enter the network proxy port.
3. Save the `elasticsearch.yml` file.

Creating Activity Guide Search Definitions

Understanding Activity Guide Search Definitions

Activity Guide search definition enables you to index the activity guide instances data in Elasticsearch that can be used for analytics in Kibana. You can also use an application table as the data source in order to provide supplemental information in the index.

Currently, the primary purpose of an activity guide search definition and index is for analytics in Kibana, that is, these indexes are not designed to be used in global search.

After creating a search definition and indexing it, you can create visualizations in Kibana. For information on creating visualizations and managing Kibana visualization and dashboards, see [Understanding Application Data and Kibana Dashboards](#).

Specifying Activity Guide General Settings

Use the General Settings page to specify the activity guide and application table (if required) data sources to be indexed for analytics. At this point, you should plan the data you want to build into a visualization, and in accordance with that plan include the required activity guide templates and application tables for the visualization. Adequate planning of a visualization at this stage ensures that you have the required data indexed and available in Kibana at the time of creating a visualization.

Access the General Settings page by selecting **PeopleTools >Search Framework >Search Designer Activity Guide** and from the left panel, choose Search Definition.

This example illustrates the fields and controls on the General Settings page. You can find definitions for the fields and controls later on this page.

Search Designer Activity guide

General Security Advanced

Search Definition: JOB_DATA_1
 *Description: Job Data
 Source Type: Activity Guide
 *Title: PTFAL_LABEL
 *Last Updated DateTime: LAST_UPDATE_DTT
 Object Owner ID: [dropdown]

Data Source

Activity Guide Templates

Activity Guide ID	Title
1 JOBDATA001	Job Data

Application Records

Record (Table) Name	Record Description
1 BD_JOBCODES	Job Codes

Search Attributes

Record (Table) Name	Field Name	Field to Index	Field Label ID	Attribute Display Name	Attribute Name
1 PTFAL_CONTEXT	ACTION	<input type="checkbox"/>	[input]		
2 PTFAL_CONTEXT	ACTION_CD	<input type="checkbox"/>	[input]		
3 PTFAL_CONTEXT	ACTION_REASON	<input type="checkbox"/>	[input]		
4 PTFAL_CONTEXT	BUSINESS_UNIT	<input type="checkbox"/>	[input]		
5 PTFAL_CONTEXT	EFFDT	<input checked="" type="checkbox"/>	EFFDT	Effective Date	EFFDT
6 PTFAL_CONTEXT	EFFDT_PAST	<input type="checkbox"/>	[input]		
7 PTFAL_CONTEXT	EFFSEQ	<input checked="" type="checkbox"/>	EFFSEQ	Effective Sequence	EFFSEQ
8 PTFAL_CONTEXT	EFFSEQ_PAST	<input type="checkbox"/>	[input]		
9 PTFAL_CONTEXT	EMPLID	<input checked="" type="checkbox"/>	EMPLID	Empl ID	EMPLID
10 PTFAL_CONTEXT	EMPL_RCD	<input checked="" type="checkbox"/>	EMPL_RCD#	Empl Record	EMPL_RCD,EMPL_RCD#
11 PTFAL_CONTEXT	HCR_ACTN_BEGRN	<input type="checkbox"/>	[input]		
12 PTFAL_CONTEXT	HCR_JOB_WIP_STATUS	<input type="checkbox"/>	[input]		
13 PTFAL_CONTEXT	MILITARY_SERVICE	<input type="checkbox"/>	[input]		
14 PTFAL_CONTEXT	NAME_DISPLAY	<input type="checkbox"/>	[input]		
15 PTFAL_CONTEXT	OPRID	<input checked="" type="checkbox"/>	USERID	User ID	OPRID,USERID
16 PTFAL_CONTEXT	PER_DETAILS	<input type="checkbox"/>	[input]		
17 PTFAL_CONTEXT	PER_ORG	<input type="checkbox"/>	[input]		
18 PTFAL_CONTEXT	POI_TYPE	<input type="checkbox"/>	[input]		
19 PTFAL_CONTEXT	REQ_REGION	<input type="checkbox"/>	[input]		
20 PTFAL_CONTEXT	REQUEST_INDCYR	<input type="checkbox"/>	[input]		
21 PTFAL_CONTEXT	TITLE_DETAILS	<input type="checkbox"/>	[input]		
22 PTFAL_LIST	CREATED_DTTM	<input checked="" type="checkbox"/>	CREATED_DTTM	Created Datetime	CREATED_DTTM
23 PTFAL_LIST	LAST_UPDATE_DTTM	<input type="checkbox"/>	[input]		
24 PTFAL_LIST	PTFAL_DUE_DATE	<input type="checkbox"/>	[input]		
25 PTFAL_LIST	PTFAL_LIST_ID	<input checked="" type="checkbox"/>	PTFAL_LIST_ID	Activity Guide ID	PTFAL_LIST_ID

Data Source Mapping

Application Record Mapping

Source Record	Source Field	Target Record	Target Field
1 BD_JOBCODES	JOBCODE	PTFAL_CONTEXT	PER_DETAILS

Effective Dated Processing

Source Record	Processing Type	Constant Date
1 BD_JOBCODES	Constant Date	[calendar icon]

Save

General | Security | Advanced

Field or Control	Description
Search Definition	Displays the search definition's name as specified when adding a new value.
Description	Provide any additional information to distinguish the search definition.

Field or Control	Description
Source Type	Displays the source type as specified when adding a new value.
Title	Currently, the default title is PTAI_LABEL, which is the label of a step. The same label is used in the summary text, so the title and summary text is the same.
Last Updated Date Time	The last updated date and time information is from the item table. However, If any updates are made on the list table fields that are selected in the search attributes grid, the updated date and time from the list table is propagated to a new field (PTSF_LAST_UPDDTTM) in the item table. This field is used for both incremental indexing and real time indexing updates on the list table.

Data Source

In addition to the activity guide templates, you can choose an application table to provide supplemental information. To help the indexing process fetch appropriate row or rows of data from application table, additional configuration is needed to map application table key fields with activity guide context table fields, which is specified in the Data Source Mapping grid. The mapping is used by the indexing process to fetch appropriate rows of data from application table and include it as part of the search document.

Field or Control	Description
Activity Guide ID	<p>Select the activity guide template whose instance data needs to be indexed. When you select an activity guide template, the Search Attributes grid is populated with its distinct context data fields and instance property fields, which can be selected for indexing.</p> <p>You may include multiple templates if the guided process contains more than one template. Ensure that you include functionally relevant templates when grouping multiple templates for a search definition.</p>

Field or Control	Description
Record (Table) Name	<p>Search Framework enables you to include an application table in the search definition to provide supplemental information that is not part of activity guide core tables. Therefore, if you specify an application table without specifying an activity guide template, it would result in an invalid search definition.</p> <p>Only application tables with key values are listed in order to facilitate correct data source mapping.</p> <p>Select the required application table. When you select an application table, the Search Attributes grid is populated with its fields, so you can choose the fields to be indexed.</p>

Search Attributes

The grid is populated with fields based on the data sources in the Data Source grid.

Field or Control	Description
Field to Index	<p>Select the fields that you want to be indexed by the search engine. These fields would be those that you intend to include in analytics. Fields that you do not select are not indexed.</p> <p>Fields that are considered necessary for any activity guide based visualization are selected, by default and cannot be edited. Oracle recommends that you do not change these mandatory fields.</p> <p>The remaining columns in the grid become enabled only after selecting the Field to Index check box for a particular row.</p>
Field Label ID	<p>This field is auto populated, however you can choose another field label in order to maintain a unique value in the Attribute Name field.</p> <p>If you select a field that is not the default field label, the Attribute Name will change to reflect your choice.</p>

Field or Control	Description
Attribute Display Name	<p>Determines how the attribute will display in Kibana when creating a visualization.</p> <hr/> <p>Note: The attribute display name must be unique in a search definition.</p> <hr/> <p>Note: The display name comes from the field labels associated with the Field Name. It cannot be entered manually. If the wording of the display name is not appropriate, you need to add a new field label.</p> <hr/> <p>On Save, a validation is built into the system to check whether duplicate attribute display names exist in the search definition. If you, a message prompts you to choose a different attribute display name.</p>
Attribute Name	<p>Displays the pre-defined attribute name that identifies the field in the search engine.</p> <hr/> <p>Note: The attribute name must be unique in a search definition.</p> <hr/> <p>An Attribute Name is formed by concatenating the field name with its field label, unless the label is the default. If the field label is the default, just the field name is used for the attribute name. For example, if the field is VENDOR and the Display Name chosen is the default, then the attribute name is VENDOR. If the Display Name chosen is SUPP1 and it is not the default, then the attribute name is VENDOR_SUPP1.</p> <p>Due to this, there might be conflict between two fields which end up having same attribute name, and in these cases it is recommended to create a new label to make the attribute unique.</p>

Data Source Mapping

The Data Source Mapping grid appears only when you specify an application table as a data source.

Use the Data Source Mapping grid to map application table key fields with activity guide context table fields.

All the application table key fields should be mapped to corresponding fields from PTAI_CONTEXT table. That is, all the application table key fields should be part of context fields of an activity guide template, which ensures fetching of correct rows of data to be included in the index. If an application table contains effective date, sequence or status fields (EFFDT, EFF_SEQ and EFF_STATUS fields) as one of the key fields, Search Framework ensures filtering of appropriate rows, applying the effective

dated processing logic for the respective application table, during indexing (full, incremental, and real time).

Note: The Effective Dated Processing grid appears only when you specify an application table which has an effective dated field.

Field or Control	Description
Source Record	Select the required application table.
Source Field	Select the required key field.
Target Record	Select the required activity guide context table.
Target Field	Select the required field from the activity guide context table.
Processing Type	Select the processing type for the application table. Options include: <ul style="list-style-type: none"> • Constant Date • Current Effective Dt • Instance Creation Dt
Constant Date	Select a date for the Constant Date process type. Note: This field is activated only when Constant Date is selected as the Processing Type .

Specifying Security Settings

Access the Security Settings page by selecting **PeopleTools >Search Framework >Search Designer Activity Guide** and from the left panel choosing Search Definition. Select the Security tab.

This example illustrates the fields and controls on the Security Settings page. You can find definitions for the fields and controls later on this page.

General	Security	Advanced
Security Settings		
Search Definition	ONBOARDING	
Description	onboarding	
Security Option		
<input checked="" type="radio"/> No Security	<input type="radio"/> Source Level Security	<input type="radio"/> Document Level Security

The Security page enables you to restrict access to data generated by a search definition. Depending on the sensitivity of the data, you can set these degrees of security:

Field or Control	Description
No Security	Select to define no security restriction for a search definition's data. Anyone with access to the application can view the data for a search definition set to No Security. That is, the data is public to all users.
Source Level Security	Select to allow or restrict access to the entire search definition as per the specified user or role. That is, only specified users and roles are able to view data generated by that search definition.
Document Level Security	Select to restrict access to specific data generated by a search definition. That is, with document level security, users can view data generated by that search definition to which they have access. Note: This is generally referred to as row-level security in PeopleSoft applications.

Setting Source Level Security

Access the source-level security settings by selecting the Source Level Security radio button.

This example illustrates the fields and controls on the Security Settings page - Source Level Security option. You can find definitions for the fields and controls later on this page.

<i>Field or Control</i>	<i>Description</i>
Type	Select Role or User depending on the scope of your intended access restriction. <ul style="list-style-type: none"> • Selecting Role restricts access to a specific PeopleSoft role. • Selecting User restricts access to a specific PeopleSoft user.
Name	Select the user or role name.
Privilege	Define the access privilege or restriction. <ul style="list-style-type: none"> • Allow - The specified role or user is allowed to view data for this search definition. • Deny - The specified role or user is not allowed to view data for this search definition.

Note: Source-level security applies to every document within that search definition.

Setting Document Level Security

Access the document-level security settings by selecting the Document Level Security radio button. Document-level security can also be thought of as attribute-based security.

This example illustrates the fields and controls on the Security Settings page - Document Level Security option. You can find definitions for the fields and controls later on this page.

General
Security
Advanced

Security Settings

Search Definition: ONBOARDING
Description: onboarding

Security Option

No Security
 Source Level Security
 Document Level Security

Document Level

Field Mapping

⌵
🔍
1-1 of 1
⏪ ⏩

Record	Field Name	Attribute Name	Privilege		
1	<input style="width: 90%;" type="text"/>	<input style="width: 90%;" type="text"/>	Allow	+	-

Document Filter App Class

Package Name:

Path:

Class ID:

Field or Control	Description
Record	Select the name of the record (table) containing the fields you want to use to restrict access.
Field Name	Select the field that will identify security values and determine access. The selected field becomes the security attribute having the specified privilege. At indexing time, when the crawler inserts application data into the index, the values populating the selected field will carry the specified privilege.

Field or Control	Description
Privilege	<p>Define the access privilege or restriction.</p> <ul style="list-style-type: none"> • Allow - The value specified in the Field Name field is allowed access to the data. • Deny - The value specified in the Field Name field cannot see the data. <hr/> <p>Note: The privilege of Deny is useful in situations where there are too many values for the security attribute if Allow were selected. For example, rather than enabling access to nine out of ten field values, it is more efficient to deny access only to the one you want to restrict.</p> <hr/> <p>Note: If multiple attributes appear in the grid the system effectively inserts an AND clause between the items in the grid.</p> <hr/>
Package Name	<p>The application class enables you to define and run additional filters and logic against the application data contained in the indexed source fields. The application class needs to be implemented to return the list of allowable values for the configured security attributes in the Field Mapping grid.</p> <p>Select the appropriate application package.</p>
Path	Select the path pointing to the application class.
Class ID	Select the appropriate application class.

Specifying Advanced Properties

Access the Advanced Properties page by selecting **PeopleTools >Search Framework >Search Designer Activity Guide** and from the left panel choosing Search Definition. Select the Advanced tab.

This example illustrates the fields and controls on the Advanced Properties page. You can find definitions for the fields and controls later on this page.

General	Security	Advanced
---------	----------	-----------------

Advanced Properties

Search Definition ONBOARDING

Description onboarding

Full Indexing Criteria

Index Only Last	<input type="text" value="0"/>	Days	Index Start Date	01/01/1900
-----------------	--------------------------------	------	------------------	------------

<i>Field or Control</i>	<i>Description</i>
Index Only Last <> Days	Enter the number of days for which you want to build the index. The number of days entered are treated as days prior to the current date. The default value is 0.
Index Start Date	The Index Start Date is calculated based on the number of days you enter in the Index Only Last <> Days field. The Index Start Date displays the date after deducting the number of days entered in the Index Only Last <> Days fields from the current date. The index will be built from the date that is displayed in the Index Start Date field. When the value of the Index Only Last <> Days field is 0, which is the default value, the Index Start Date is set to 01/01/1900.

Chapter 9

Creating Search Categories

Understanding Search Categories

Search Categories are essential for the Search Framework. Search categories:

- Organize search definitions into manageable, logical groups of data.
- Are required for searches. Search queries run only against search categories, not search definitions. A search definition must be a member of a search category before it can be searched.
- Enables users to focus searches to help refine results and restrict searches to a particular set of search indexes. For example, users can run searches against categories like Customers, Purchase Orders, or Products, rather than running a search against all indexes in the search engine.
- Allow users to search across several different groups of data simultaneously. For example if you were looking for a particular Item ID, with a Search category containing multiple Search Definitions, you can search across purchase order, sales orders and inventory at one time.
- Provide improved search performance and results by limiting the number of indexes searched.

Note: Search queries cannot be run directly against search definitions. Search queries can only be run against search categories.

Specifying General Search Category Settings

Access the General page by selecting PeopleTools, Search Framework, Search Designer Activity Guide, Search Category.

This example illustrates the fields and controls on the General page. You can find definitions for the fields and controls later on this page.

Field or Control	Description
<p>Search Category Name</p>	<p>Displays the search category name.</p> <hr/> <p>Important! Though any number of search definitions can be mapped to a search category, for each and every search definition, a search category of the <i>same</i> name must exist, and the search definition must belong to that search category. You can create search categories of a different name that are not associated with a single search definition. A search definition may belong to more than one search category, but it <i>must</i> belong to one with the same name.</p> <hr/> <p>For example, search definition XYZ must belong to search category XYZ, but search definition XYZ can also belong to search category ABC. Search category ABC does <i>not</i> require a corresponding search definition of the same name, and it can contain multiple search definitions. Search categories with more than one search definition need to be manually deployed.</p>
<p>Description</p>	<p>Add any additional information to distinguish search categories.</p> <hr/> <p>Note: If this search category becomes exposed through Global Search, then the description you enter may appear in the Global Search Bar. The description must be appropriate for end user viewing. Make the description as informative, concise, and intuitive as possible.</p> <hr/> <p>Note: The description is translatable.</p>

Field or Control	Description
Duplicates	Note: This option is not implemented with Elasticsearch. It is retained for backward compatibility.
Search Group	Select to enable this search category to be available for context searching within Global Search.
Search Definition	Add all search definitions that belong to this search category.
Save As	<p>Click to clone the search category. In the subsequent page, you need to enter a name for the cloned search category.</p> <p>When you clone a search category, all the meta-data associated with the existing search category is saved in the cloned search category.</p> <p>Note: The Save As button is not available when you create or define a new search category.</p>

Selecting Advanced Search Field Settings

Access the Advanced Search Fields page by selecting PeopleTools, Search Framework, Search Designer, Search Category and selecting the Advanced Search Fields tab.

The Advanced Search Fields page enables you to view and modify attributes to show on the advanced search page during end user searches to add more criteria to the basic search mode.

When working with the advanced search page attributes, keep these items in mind:

- This configured list can be programmatically retrieved from the SearchCategory App Class in the PT_SEARCH Application Package using method GetConfiguredFilterAttributes.
- All the attributes associated with this category can be retrieved using the same App Class with the method GetAllAttributes.
- Attributes not appearing on the Advanced Search Fields page can be retrieved using GetNonConfiguredFilterAttributes. This method will provide a list of attributes defined in the related search definitions which have not been configured here.

This example illustrates the fields and controls on the Advanced Search Fields page. You can find definitions for the fields and controls later on this page.

General | **Advanced Search Fields** | Facets | Custom Search Page | Display Fields

Advanced Search Fields

Search Category Name: PTPORTALREGISTRY
Description: Navigator

Auto detect common fields


Advanced Search Fields		Personalize Find [?] [grid]		First	1-4 of 4	Last
Sequence	*Search Attribute	Description				
1	PORTAL_LABEL	Portal Label				
2	DESCR254	Long Description				
3	PORTAL_PRNTLABEL_PORTAL_FL	Folder				
4	PORTAL_SOURCE	Source Application				

Field or Control	Description
Auto detect common fields	<p>When selected, this flag will gather all fields which are common to all search definitions listed on the General page. This will allow for searching across all search definitions in the category.</p> <p>Select to avoid manually listing the common fields from the joined records. (default)</p> <p>Deselect to delete some or to add additional attributes which are not common to all search definitions.</p> <hr/> <p>Note: When you select the Auto Detect flag, you must save the search category before the common fields will be properly displayed in the grid.</p>
Sequence	Control the sequence in which the search attributes appear in the advanced search interface.

Selecting Facet Settings

Access the Facets page by selecting PeopleTools, Search Framework, Search Designer Activity Guide, Search Category and selecting the Facets tab.

This example illustrates the fields and controls on the Facets page. You can find definitions for the fields and controls later on this page.



Field or Control	Description
Auto detect facets	<p>Select to detect common facets for all the search definitions belonging to the search category. When you save the search category, any facets that are common to all search definitions are displayed. This saves you from needing to discover which facets can be used to filter across all search definitions.</p> <p>Deselect this option if you wish to remove some facets or add ones that are not common to all search definitions (not recommended as this may distort facet counts).</p>
Sequence	<p>Control the sequence in which the faceted fields display for narrowing search results in an intuitive fashion for users.</p> <p>For example, perhaps you might want to ensure that these faceted fields display in this order: Country, State, County, City.</p>
Description	<p>Displays the value of the Attribute Display Name column in the grid on the Map Search Attributes page for the search definition.</p>
Is Associated	<p>Select if you intend to create a user-specific association with the facet.</p> <hr/> <p>Note: Creating associated facets applies only to certain situations and expects additional requirements to be in place.</p> <hr/> <p>For more information, refer to the topic below, “Creating Associated Facets.”</p>

Field or Control	Description
Facet Association	<p>Select how the value to which the facet should be associated will be determined. Options are:</p> <ul style="list-style-type: none"> • <i>Application Class</i>: When selected, the Associated Value App Class group box appears below the Facets grid on the Facets page, enabling you to specify the application package that will determine the facet association. • <i>Employee ID</i>: Associates the facet value with the employee ID value (%EmployeeID). • <i>User ID</i>: Associates the facet value with the currently signed on user (%UserID).
Advanced Settings	<p>Select to add further sorting and categorization of the facets depending on the data type of the facet. For example, you can create sort orders for string facets and ranges for number and date facets. See the following topics for more details on the options for each facet data type.</p> <p>By default, facet values are sorted by count, meaning that the system displays the facet value containing the most search results first, or at the top of the list. While this may be desirable in many situations, in some situations it may provide a better user experience if the facets are sorted based on criteria other than just the number of search results.</p> <p>For example, if the facet is a date field, it may make more sense to sort the facet values by the actual year, in a sequential order, (2013, 2012, 2011, and so on) as opposed to the number of search results with in each year. Likewise, in another situation, it may be more intuitive to sort a faceted Customer field alphabetically, regardless of search result count.</p> <hr/> <p>Important! Setting facet ranges applies only to date and number data types. Values with data types of datetime cannot be used as facets.</p> <hr/> <p>See Defining Advanced Settings for String Facets.</p> <p>See Defining Advanced Settings for Number Facets.</p> <p>See Defining Advanced Settings for Date Facets.</p>

Keep in mind that this configured list of facets can be programmatically retrieved using the GetFacetFilters method in the SearchCategory App Class in the PT_SEARCH Application Package. The result, which is an array of PT_SEARCH:FacetFilter, can be passed to the PT_SEARCH:SearchQuery class to get facet nodes in the search results.

Creating Associated Facets

An associated facet is a facet that is associated with a specific field value, such as an employee ID or the current user ID. This enables that associated facet value to be filtered automatically by the facet association value. These facets appear in the facet pane within a separate facet node entitled “My

Association” for example (default). All other facets in the facet pane behave normally, and the My Association facet node appears only if search results associated to the current user exist for the current search request.

For example, consider users accessing a component where job postings are created and managed. A hiring manager may want to have a facet category containing only search results associated with her, which might contain only those job postings she created. Likewise, a recruiter may want to have a facet category containing results associated only with him, such as the job openings for which he is the recruiter.

If you intend to implement associated facets for custom search categories, it is recommended that you examine the associated facets in your delivered PeopleSoft application search definitions and search categories to see examples. The following list provides general recommendations:

- You need to ensure that the facet association value is available in your table structure, your query or connected query, and that the value is indexed. Or, you can reference an application class, which will determine the association value programmatically.
- Only hierarchical search attributes can be associated facets. Define a hierarchy-based faceted attribute in the search definition, where the first level acts as the key for the facet association, for example user ID or employee ID, and the second level being the association, such as *A.OPRID/HRS_ASSOCIATION_CD*. It is recommended that at the association level, you reference a translate value (XLAT), such as hiring manager, recruiter, and so on.
- In cases, where the first level of the facet association is not user ID or employee ID, which PeopleTools populates, the PeopleSoft application delivers an application class to determine this value. If you are creating custom search definitions and categories, you will need to create the application class.

Defining Advanced Settings for String Facets

This example illustrates the fields and controls on the Advanced Facet Settings page (for string data types). You can find definitions for the fields and controls later on this page.

Advanced Facet Settings

Facet Name: PORTAL_PRNTLABEL_PORTAL_FLDR Data Type: String

Allow Multi Select
 Expand Tree Facet

Override default settings

Minimum Document Count: 0 Disable Facet Count

Maximum Facet Children: 0

Sort Order: Personalize | Find | View All | [Grid Icon] | First | 1 of 1 | Last

Level	Sort as	Sort order
1	[Dropdown]	[Dropdown] [+] [-]

OK Cancel

Field or Control	Description
Allow Multi Select	<p>By default, all facets enable you to choose multiple facet values on the Global Search results page and the component search results page to filter and narrow down the search results.</p> <p>Select the Allow Multi Select check box to enable multiple facet-values selection in hierarchical facets. If you deselect this option, single select applies to the current facet.</p> <hr/> <p>Note: Selecting multiple facet values is supported only on fluid search pages; not on classic or classic plus search pages.</p> <hr/> <p>See Working with Search Results in Global Search.</p>
Expand Tree Facet	<p>The Search Options page contains the Expand Tree Facet option to specify the setting at the system level. The default value is collapsed, that is, the Expand Tree Facet check box is not selected.</p> <p>Select the Expand Tree Facet option if you want the tree facets to be expanded. If you select this option, you are overriding the setting specified on the Search Options page, and the override setting applies to the current facet only.</p> <p>You can use this option to override the setting specified on the Search Options page.</p> <p>See Managing General Search Options.</p>
Override default settings	<p>Select to override the values entered on the Search Options page (PeopleTools, Search Framework, Search Admin Activity Guide, under Administration, select Search Options). When selected the disabled settings on the page become enabled so that you can override the equivalent values entered on the Search Options page. The override values apply to the current facet only. If not selected, the values entered on the Search Options page apply.</p>
Minimum Document Count	<p>Set the minimum number of search results (documents) required to display a facet. Default is <i>1</i>.</p>
Maximum Facet Children	<p>Set the maximum number of children to be returned for each facet node. Default is <i>100</i>.</p> <p>Options to enter are:</p> <ul style="list-style-type: none"> • An integer greater than 0 to limit the number of children displayed. • <i>0</i> to indicate no children should be displayed. • <i>-1</i> to return all children (no limit).

Field or Control	Description
Disable Facet Count	<p>Enables controlling the display of the result count (document count) for each facet node. For example, if 15 search results met a facet node's criteria, (15) appears next to the facet label.</p> <p>Enter <i>Y</i> for Yes or <i>N</i> for No.</p>
Level	<p>In the case of a nested facet, you can add additional rows to the sort order grid, and set different sort options at each level. The Level column indicates the level to which the sort options apply (1, 2, 3, and so on).</p> <p>For multi-select facets, level 1 sort option is used for all levels.</p>
Sort as	<p>Select one of these options to specify how the facet should be treated:</p> <ul style="list-style-type: none"> • <i>Count</i>. While the default sort order is to sort facet values by search result count in a descending order, the Count option enables you to apply an ascending order, for example. • <i>Number</i>. Apply the number data type to the facet value. • <i>String</i>. Apply the string, or character, data type to the facet value.
Sort order	<p>Specify the sort order that best suits your selected Sort as option and how users expect to see the values displayed, ascending or descending.</p>

Defining Advanced Settings for Number Facets

When the search attribute is a number data type, the Advanced Facet Settings page enables you to configure the sort order and ranges.

This example illustrates the fields and controls on the Advanced Facet Settings page (for date data types). You can find definitions for the fields and controls later on this page.

The screenshot shows the 'Advanced Facet Settings' window for a facet named 'QE_JOB_OPENING_ID' with a 'Number' data type. The 'Allow Multi Select' checkbox is checked, and 'Expand Tree Facet' is unchecked. Under 'Override default settings', both 'Minimum Document Count' and 'Maximum Facet Children' are set to 0. The 'Sort Order' section shows a single level with a 'Predefined' sort order. The 'Number Range' table is as follows:

Sequence	Node Name	Begin Operator	Begin Facet Value	End Operator	End Facet Value	Message Set Number	Message Number
1	Under 500000	Not Applicable		Less Than Equals	500000.00		
2	500000 - 600000	Greater Than	500000.00	Less Than Equals	600000.00		
3	600000 - 700000	Greater Than	600000.00	Less Than Equals	700000.00		
4	Above 900000	Greater Than	900000.00	Not Applicable			
5	Advanced Settings are in effect.	Greater Than	700000.00	Less Than Equals	800000.00	141	64

Field or Control	Description
Allow Multi Select	<p>By default, all facets enable you to choose multiple facet values on the Global Search results page and the component search results page to filter and narrow down the search results.</p> <p>Select the Allow Multi Select check box to enable multiple facet-values selection in hierarchical facets. If you deselect this option, single select applies to the current facet.</p> <hr/> <p>Note: Selecting multiple facet values is supported only on fluid search pages; not on classic or classic plus search pages.</p> <hr/> <p>See Working with Search Results in Global Search.</p>
Expand Tree Facet	<p>The Search Options page contains the Expand Tree Facet option to specify the setting at the system level. The default value is collapsed, that is, the Expand Tree Facet check box is not selected.</p> <p>You can use this option to override the setting specified on the Search Options page.</p> <p>See Managing General Search Options.</p>
Override default settings	<p>Select to override the values entered on the Search Options page (PeopleTools, Search Framework, Search Admin Activity Guide, under Administration, select Search Options). When selected the disabled settings on the page become enabled so that you can override the equivalent values entered on the Search Options page. The override values apply to the current facet only. If not selected, the values entered on the Search Options page apply.</p>

Field or Control	Description
Minimum Document Count	<p>Set the minimum number of search results (documents) required to display a facet. Default is <i>1</i>.</p> <p>This property can be overridden by selecting Advanced Settings on the Facets page (PeopleTools, Search Framework, Search Designer Activity Guide, Search Category, Facets, Advanced Settings).</p>
Maximum Facet Children	<p>Set the maximum number of children to be returned for each facet node. Default is <i>100</i>.</p> <p>Options to enter are:</p> <ul style="list-style-type: none"> • An integer greater than 0 to limit the number of children displayed. • <i>0</i> to indicate no children should be displayed. • <i>-1</i> to return all children (no limit).
Disable Facet Count	<p>Enables controlling the display of the result count (document count) for each facet node. For example, if 15 search results met a facet node's criteria, (15) appears next to the facet label.</p> <p>Enter <i>Y</i> for Yes or <i>N</i> for No.</p>

Sort Order

Displays an additional option than the string sort order, *Predefined*, which displays the facet node, based on the sequence number in which the ranges are defined.

Number Range

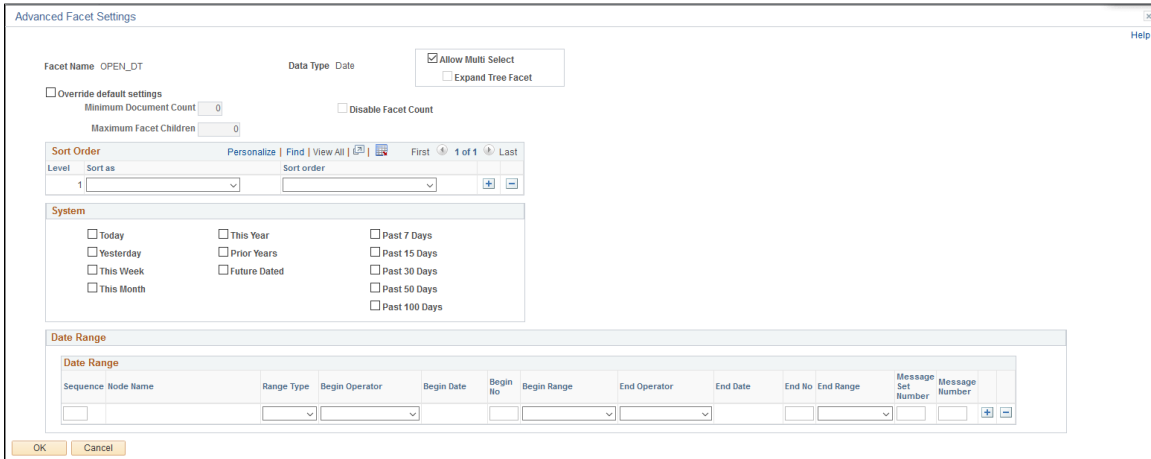
Field or Control	Description
Sequence	Order in which the ranges are defined.
Node Name	<p>Display name of the facet range. This is a display only field.</p> <p>The facet node is derived from the begin and end values specified for the ranges or by using the Message Catalog. Message Catalog has the precedence</p>

Field or Control	Description
Begin Operator	<p>The first operator for the range sets the beginning of the range. To set the beginning of the range, you have these options to operate against the beginning facet value specified:</p> <ul style="list-style-type: none"> • Greater Than • Greater Than Equals: greater than or equal to. • Not Applicable: no starting boundary.
Begin Facet Value	Set the beginning value for a facet range.
End Operator	<p>The second operator for the range, sets the end of the range. To set the end of the range, you have these options to operate against the end facet value specified:</p> <ul style="list-style-type: none"> • Less Than • Less Than Equals: less than or equal to. • Not Applicable: no ending boundary.
End Facet Value	The end value for the facet range.
Message Set Number	Message catalog set number to provide a custom facet node name.
Message Number	Message number to provide a custom facet node name.

Defining Advanced Settings for Date Facets

When the search attribute is a date data type, the Advanced Facet Settings page enables you to configure the sort order and ranges.

This example illustrates the fields and controls on the Advanced Facet Settings page (for date data types). You can find definitions for the fields and controls later on this page.



Field or Control	Description
<p>Allow Multi Select</p>	<p>By default, all facets enable you to choose multiple facet values on the Global Search results page and the component search results page to filter and narrow down the search results.</p> <p>Select the Allow Multi Select check box to enable multiple facet-values selection in hierarchical facets. If you deselect this option, single select applies to the current facet.</p> <hr/> <p>Note: Selecting multiple facet values is supported only on fluid search pages; not on classic or classic plus search pages.</p> <hr/> <p>See Working with Search Results in Global Search.</p>
<p>Expand Tree Facet</p>	<p>The Search Options page contains the Expand Tree Facet option to specify the setting at the system level. The default value is collapsed, that is, the Expand Tree Facet check box is not selected.</p> <p>You can use this option to override the setting specified on the Search Options page.</p> <p>See Managing General Search Options.</p>
<p>Override default settings</p>	<p>Select to override the values entered on the Search Options page (PeopleTools, Search Framework, Search Admin Activity Guide, under Administration, select Search Options). When selected the disabled settings on the page become enabled so that you can override the equivalent values entered on the Search Options page. The override values apply to the current facet only. If not selected, the values entered on the Search Options page apply.</p>

Field or Control	Description
Minimum Document Count	<p>Set the minimum number of search results (documents) required to display a facet. Default is <i>1</i>.</p> <p>This property can be overridden by selecting Advanced Settings on the Facets page (PeopleTools, Search Framework, Search Designer Activity Guide, Search Category, Facets, Advanced Settings).</p>
Maximum Facet Children	<p>Set the maximum number of children to be returned for each facet node. Default is <i>100</i>.</p> <p>Options to enter are:</p> <ul style="list-style-type: none"> • An integer greater than 0 to limit the number of children displayed. • <i>0</i> to indicate no children should be displayed. • <i>-1</i> to return all children (no limit).
Disable Facet Count	<p>Enables controlling the display of the result count (document count) for each facet node. For example, if 15 search results met a facet node's criteria, (15) appears next to the facet label.</p> <p>Enter <i>Y</i> for Yes or <i>N</i> for No.</p>

Sort Order

Displays an additional option than the string sort order, *Predefined*, which displays the facet node, based on the sequence number in which the ranges are defined.

System

This group box provides predefined set of date facet nodes.

Field or Control	Description
Today	Displays results with a date for the current day only.
Yesterday	Displays results with a date of the previous day only.
This Week	Displays results with a date for the current week only.
This Month	Displays results with a date for the current calendar month only.
This Year	Displays results with a date for the current calendar year only.

Field or Control	Description
Prior Years	Displays results with a date prior to the current year.
Future Dated	Displays results which are applicable beginning with tomorrow onward.
Past <i>N</i> Days	Displays results which are applicable for the past <i>N</i> days (current date - <i>N</i>), where <i>N</i> is a predefined number of days (7, 15, 30, 50, 100).

Date Range

In addition to the predefined system ranges, you can create additional date ranges using the Date Range grid. This allows you to create ranges based on static dates or relative ranges based on the current date.

Field or Control	Description
Sequence	Defines the order in which the facets appear.
Node Name	<p>A read-only field displaying how the facet node will appear to the end user. The value depends on whether the range is fixed or relative.</p> <p>If the range is fixed, then the system derives the facet node from the begin and end values specified for the ranges or by using the message catalog, with the message catalog taking precedence.</p> <p>If the range is relative, then it is required to provide a message catalog entry for the facet node name.</p>
Range Type	<p>Enables you to create ranges on fixed dates or dates relative to today. Options are:</p> <ul style="list-style-type: none"> • <i>Fixed</i>. Select to define static beginning and end dates. • <i>Relative</i>. Select to enable the Begin Range drop-down list to select dates relative to the current date. <hr/> <p>Note: Custom facet ranges do not support the facet nodes for rolling dates. For example, you cannot predefine a facet for each quarter on a financial year and have it dynamically update every year.</p> <hr/>

Field or Control	Description
Begin Operator	<p>The first operator for the range sets the beginning of the range. To set the beginning of the range, you have these options to operate against the beginning facet value specified:</p> <ul style="list-style-type: none"> • Greater Than • Greater Than Equals: greater than or equal to. • Not Applicable: no starting boundary.
Begin Date	Available only for fixed ranges. Set the beginning date for the facet range.
Begin No	This is the number of days/months/years to be set as the beginning of the facet range as per the selection from the Begin Range drop-down list.
Begin Range	<p>Set the span of the dates to be included in the range (past or future). Select in the form of:</p> <ul style="list-style-type: none"> • days/months/years before. • days/months/years from now.
End Operator	<p>The second operator for the range, sets the end of the range. To set the end of the range, you have these options to operate against the end facet value specified:</p> <ul style="list-style-type: none"> • Less Than • Less Than Equals: less than or equal to. • Not Applicable: no ending boundary.
End Date	Available only for fixed ranges. Set the end date for the facet range.
End No	This is the number of days/months/years to be set as the end of the facet range as per the selection from the End Range drop-down list.

Field or Control	Description
End Range	Set the span of the dates to be included in the range (past or future). Select in the form of: <ul style="list-style-type: none"> • days/months/years before. • days/months/years from now.
Message Set Number	Message catalog set number to provide a custom facet node name.
Message Number	Message number to provide a custom facet node name.

Identifying Custom Search Pages

Oracle provides default search pages to use with most application purposes. If you need to do something more than this, then you must create custom search pages for custom applications or to provide additional search features.

To indicate for the system to use these custom pages when doing a Global Search, then you should specify that search page, as well as any custom search result or advanced search pages.

Note: This functionality is only available for Global Search, not keyword component-level search.

Access the Custom Search Page page by selecting PeopleTools, Search Framework, Search Designer Activity Guide, Search Category and selecting the Custom Search Page tab.

This example illustrates the fields and controls on the Custom Search Page. You can find definitions for the fields and controls later on this page.

General
Advanced Search Fields
Facets
Custom Search Page
Display Fields

Custom Search Page Settings

Search Category Name PTPORTALREGISTRY

Description Navigator

Custom search page

Menu Name

Component Name

Market Global

Search results page

Advanced search page

<i>Field or Control</i>	<i>Description</i>
Menu Name	Select the menu to which the search page belongs.
Component Name	Select the search page's component name.
Market	Select the appropriate market.
Search results page	Select the search results page to use.
Advanced search page	Select the advanced search page to use.

Working With Display Fields

Access the Display Fields page by selecting PeopleTools, Search Framework, Search Designer Activity Guide, Search Category and selecting the Display Fields tab.

This example illustrates the fields and controls on the Display Fields page. You can find definitions for the fields and controls later on this page.

General
Advanced Search Fields
Facets
Custom Search Page
Display Fields

Display Fields

Search Category Name PTPORTALREGISTRY

Description Navigator

Display Fields
Personalize | Find | |
First ◀ 1-3 of 3 ▶ Last

Sequence	*	Search Attribute	Description		
1	1	PORTAL_LABEL	Portal Label		+ -
2	2	DESCR254	Long Description		+ -
3	3	PORTAL_SOURCE	Source Application		+ -

Field or Control	Description
Sequence	Use the Sequence edit box to specify the order in which the display fields appear.
Search Attribute	Select the search attribute from the drop-down list.

Administering PeopleSoft Search Framework

Understanding PeopleSoft Search Framework Administration

The system administration related to the Search Framework architecture can include tasks involving the search engine, Integration Broker, application server domains, and so on. PeopleTools provides a Search Framework activity guide that enables you to complete the administration tasks from the Search Framework administration activity guide, specifically under the menu, PeopleTools, Search Framework, Search Admin Activity Guide.

Note: If a task on the activity guide is completed in the relevant system, you may ignore the task in the activity guide.

Understanding the Search Admin Activity Guide

This topic provides overview information about the Search Admin Activity Guide.

The Search Admin Activity Guide is a dashboard in the Search Framework where you can deploy search definitions and manage the search definitions on the search engine. The tasks you can perform using the Search Admin Activity Guide include:

- Configuring search engine connectivity
- Deploying search definitions and categories
- Scheduling index creation and maintenance
- Running index build process
- Monitoring index run schedule
- Testing an index

Other page elements that appear on the page are discussed elsewhere in this book.

Setting Up Local Nodes

To set up local nodes, access the Nodes page by selecting PeopleTools, Search Framework, Search Admin Activity Guide, Configuration, Setup Local Nodes.

Important! Ensure that you define URLs in the Content URI Text and Portal URI Text fields for each local node that is in use.

For information on setting up local nodes, see "Configuring Nodes" (Integration Broker Administration) and "Understanding Managing Integration Gateways" (Integration Broker Administration).

Specifying Integration Broker Gateway

To specify Integration Broker gateway, access the Gateways page by selecting PeopleTools, Search Framework, Search Admin Activity Guide, Configuration, IB Gateway.

See "Administering Integration Gateways" (Integration Broker Administration) for more information.

Setting Up Target Locations

To set up target location, access the Target Locations page by selecting PeopleTools, Search Framework, Search Admin Activity Guide, Configuration, Setup Target Location.

See "Using the Target Locations Page to Set Target Locations for Services" (Integration Broker Administration) for more information.

Working With Search Instances

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

The search instance definition specifies connectivity and other administration settings required for:

- running queries against the search server.
- administrative tasks.

Creating Search Instances

To define a search instance, access the Search Instance Properties page by selecting PeopleTools, Search Framework, Search Admin Activity Guide, Configuration, Search Instance.

You can also define a search instance by using the SEARCH_TEMPLATE configuration template in Automated Configuration Manager (ACM). The steps to run the template are described later in this topic, see [Creating a Search Instance in Automated Configuration Manager](#).

This example illustrates the fields and controls on the Search Instance Properties page. You can find definitions for the fields and controls later on this page.

Search Instance Properties

Search Instance

Search Instance	PTSF_DEFAULT	
Search Provider	Elasticsearch	Search Options Config

Search Instance Properties | | [View All](#)

*Host Name	<input type="text" value="example.com"/>	<input type="button" value="Ping"/>	<input type="button" value="+"/>	<input type="button" value="-"/>
*Port	<input type="text" value="9200"/>			
*SSL Option	<input type="text" value="Disable"/>			
*User Name	<input type="text" value="esadmin"/>	<input type="button" value="Test Login"/>		
*Password	<input type="password"/>			
*Proxy Name	<input type="text" value="people"/>	<input type="button" value="Proxy Login"/>		
*Proxy Pwd	<input type="password"/>			

[Elasticsearch Interact](#)

Kibana

Host Name	<input type="text" value="example.com"/>
Port	<input type="text" value="5601"/>
SSL Option	<input type="text" value="Disable"/>

Logstash

Host Name	<input type="text" value="example.com"/>	<input type="button" value="Ping"/>
Logstash Port	<input type="text"/>	
*SSL Option	<input type="text" value="Enable"/>	

Call Back Properties [?](#)

URL	<input type="text" value="http://example.com:8000/PSIGW/RESTListeningConnector/T5901JDC"/>
User Name	<input type="text" value="PSAPPS"/>
*Password	<input type="password"/>
*Confirm Password	<input type="password"/>

[Update deployed definitions](#)

[Set Namespace Aliases](#)

Search Instance

To enable connectivity between your PeopleSoft system and the search engine, you need to provide these search instance values.

Field or Control	Description
Search Instance	<p>Enter a name for the search instance.</p> <hr/> <p>Note: PeopleSoft Search Framework enables you to create multiple search instances for the purpose of fail-over mechanism.</p> <hr/> <p>See Understanding High Availability Feature in PeopleSoft Search Framework and Implementing Fail-over Mechanism for more information.</p> <p>When you create multiple search instances, you can use the Search Instance Administration page to specify priority for each search instance. See Administering Search Instances.</p>
Search Provider	Elasticsearch is the default search provider.
Search Options Config	See Managing General Search Options .

Search Instance Properties

This section contains details for the user name and password used to connect to the search engine. These credentials are part of the authentication mechanism from PeopleSoft to the search engine. During an incoming search request, the search engine matches these credentials with values stored in the configuration file to authenticate the incoming search request.

Note: Use this section to add multiple nodes for an Elasticsearch search instance by clicking the scroll bar plus button. Configuring multiple nodes enables you to use the fail-over mechanism when one of the nodes may go down for any reason. In a multi-node Elasticsearch cluster, always ensure that the replica value is set to at least 1. For details, see [Managing General Search Options](#). For more information, see [Implementing Fail-over Mechanism](#).

Oracle PeopleSoft recommends setting up an Elasticsearch cluster with three nodes on machines that are not on the same physical server.

For information on creating additional Elasticsearch nodes, see [Adding a New Node to an Existing Cluster](#).

Field or Control	Description
Host Name	<p>Enter the server name of the host where the search engine is running, including the domain. The host name value can be a DNS name or an IP address.</p> <p>For example:</p> <p><i>example.com</i></p>

Field or Control	Description
Port	<p>Enter the port on which the search engine listens for requests.</p> <hr/> <p>Note: The default port for Elasticsearch is 9200. If a load balancer or reverse proxy is used between Elasticsearch and PeopleSoft, the Elasticsearch port number should be the same across all nodes in the cluster.</p> <hr/>
SSL Option	<p>Select one of the options for SSL:</p> <ul style="list-style-type: none"> • Select DISABLE in the following scenarios: <ul style="list-style-type: none"> • If SSL is not implemented on the search engine. • If http is implemented on both PeopleSoft and search engine. • Select ENABLE in the following scenarios: <ul style="list-style-type: none"> • If SSL is implemented on both PeopleSoft and the search engine. • If https is implemented on PeopleSoft and SSL is implemented on the search engine. <p>See Configuring SSL between PeopleSoft and Elasticsearch.</p>
Ping	<p>Click to make sure your host name and port are correct and that the search engine is available. A positive result displays a success message.</p> <p><i>Ping Test Result: Success. Elasticsearch Version 2.3.2 (262, 613)</i></p>
User Name	Enter the user name to connect to the search engine.
Password	Enter the password associated with the user name.
Test Login	<p>Click to test the credentials you have specified against the search engine. A positive result displays a success message.</p> <p><i>Login Success. (262, 615)</i></p>

Field or Control	Description
Proxy User	<p>Enter a trusted entity.</p> <p>This enables the PeopleSoft system to log into the search engine using a proxy identity to run a query. This proxy account is a Trusted Entity. You may need to contact your administrator to obtain Trusted Entity credentials or to have those credentials created for the Search Framework query service.</p> <hr/> <p>Important! Even though the Search Framework is authenticated by the search engine as the Trusted Entity, the search engine uses the identity of the currently signed on PeopleSoft user for query authorization.</p> <hr/> <p>Note: In Elasticsearch, Proxy User acts as an authentication user for every search request.</p> <hr/>
Proxy Pwd	Enter the password associated with the trusted entity.
Proxy Login	<p>Click to test the credentials you have specified against the search engine. A positive result displays the success message.</p> <p><i>Login Success (262,615)</i></p>
+ (plus sign)	<p>Click the + (plus) sign in the grid to enter details of another node if you plan to implement fail-over mechanism.</p> <hr/> <p>Note: Oracle PeopleSoft recommends a three-node cluster initially. Then, depending on the hardware, memory availability, and search performance, you may add nodes to the existing cluster.</p> <hr/> <p>See Understanding Clusters in Elasticsearch.</p>
Elasticsearch Interact	<p>Click this link to open the Elasticsearch Interact page that enables you to choose the cluster and index level options on which you want to obtain information from the Elasticsearch search engine.</p> <p>See Administering Elasticsearch Using the Elasticsearch Interact Page.</p>

Kibana

This section contains details of the server where Kibana is installed.

Field or Control	Description
Host Name	<p>Enter the server name of the host where Kibana is running, including the domain. The host name value can be a DNS name or an IP address.</p> <p>For example, <i>example.com</i></p> <hr/> <p>Note: Kibana can be installed on the same server as Elasticsearch or on a separate server. Oracle recommends that you install Kibana on one of the Elasticsearch nodes in the cluster to ensure better connectivity between Elasticsearch and Kibana.</p> <hr/>
Port	<p>Enter the port on which Kibana listens for requests.</p> <hr/> <p>Note: The default port for Kibana is 5601.</p> <hr/>
SSL Option	<p>Select one of the options for SSL:</p> <ul style="list-style-type: none"> • Select DISABLE in the following scenarios: <ul style="list-style-type: none"> • If SSL is not implemented on Kibana. • If http is implemented on both PeopleSoft and Kibana. • Select ENABLE in the following scenarios: <ul style="list-style-type: none"> • If SSL is implemented on both PeopleSoft and Kibana. • If https is implemented on PeopleSoft and SSL is implemented on Kibana. <p>For more information, see Configuring SSL for Kibana.</p>

Logstash

This section contains details of the server where Logstash is installed.

Field or Control	Description
Host Name	<p>Enter the server name of the host where Logstash is running, including the domain. The host name value can be a DNS name or an IP address.</p> <p>For example, <i>example.com</i>.</p>
Logstash Port	<p>Enter the port on which Logstash listens for requests.</p> <hr/> <p>Note: The default port for Logstash is 9800.</p> <hr/>
SSL Option	<p>Select one of the options for SSL:</p> <ul style="list-style-type: none"> • Select DISABLE in the following scenarios: <ul style="list-style-type: none"> • If SSL is not implemented on Logstash. • If http is implemented on both PeopleSoft and Logstash. • Select ENABLE in the following scenarios: <ul style="list-style-type: none"> • If SSL is implemented on both PeopleSoft and Logstash. • If https is implemented on PeopleSoft and SSL is implemented on Logstash.
Ping	<p>Click to make sure your host name and port are correct and that the search engine is available. A positive result displays a success message.</p>

Call Back Properties

The callback properties are used to perform a callback to PeopleSoft to fill in data such as search attribute values, so you need to provide the URL and password for this access.

Note: In Elasticsearch, PeopleSoft Search Framework uses the REST ListeningConnector for callback.

Field or Control	Description
URL	<p>Enter the URL for the PeopleSoft system listening connector, using the following syntax.</p> <p><i>http://<server>:<port>/PSIGW/RESTListeningConnector/<node></i></p> <hr/> <p>Note: If the callback URL is https and if SSL is not configured on Elasticsearch, then the PeopleSoft root CA needs to be installed in the Elasticsearch truststore and the truststore path and password should be specified in the <i>elasticsearch.yml</i> configuration file.</p> <p>If Elasticsearch is behind a load balancer and the setting is to terminate SSL at load balancer, the <i>orclssl.callback</i> property should be set to false in the <i>elasticsearch.yml</i> configuration file.</p> <hr/>
User Name	<p>Enter the PeopleSoft user name granted the permission list PTPT3300. This user must have the Search Server and Search Query Administrator roles.</p> <hr/> <p>Note: This user name must exist as an active user profile on the PeopleSoft system listed in the URL specified. This is generally a system user as opposed to an interactive user. In addition, if you have attachments on an FTP server that will be accessed by Elasticsearch, this user must also be authorized to access the FTP site.</p> <hr/>
Password\Confirm Password	<p>Enter and confirm the password associated with the PeopleSoft user name.</p>

Field or Control	Description
Validate	<p>Click to validate the URL, service operations, and credentials you have entered. A positive result displays the following message on the Validation Report page:</p> <p><i>All validations are successful.</i></p> <p>The following service-operation validations are performed to confirm:</p> <ul style="list-style-type: none"> • Service operations related to PTSF are active. • Handlers corresponding to service operation are active. • Routers corresponding to service operation are active. • Service operation alias name matches with router alias name. • Service-operation security is set.
Update deployed definitions	<p>Click this link to invoke a bulk update for all definitions currently deployed to the search engine.</p> <p>If you have changed any of the callback properties (URL, user name, or password), all currently deployed search definitions need to be updated with the new callback values to continue successful interaction with the Search Framework.</p>
Set Namespace Aliases	<p>Click this link to set namespace aliases when you clone an environment.</p> <p>See Setting Namespace Aliases for more information.</p>
Save	<p>When you save a search instance, Elasticsearch creates an orcl_es_auth index that stores the callback service URL for each PeopleSoft database. When a user accesses Kibana in PeopleSoft, the user is authenticated using the callback service request to the orcl_es_auth index of a specific database.</p>

Modifying Search Instances

You can modify any of the search instance values if the information changes for the specified server. That is, if the server receives a new IP address or a new DNS name, if the port for the search services changes, or if any credentials change, then the values can be updated as needed.

Important! You may not reuse a search instance entry for an entirely new instance of a search server if the existing search instance contains deployed search definitions and search categories. If you reuse an existing search instance entry that contains deployed search definitions and search categories, it may cause deployment errors.

Deleting Search Instances

Deleting a saved search instance is not supported.

Sharing a Search Instance Among Multiple PeopleSoft Application Systems

If you have multiple PeopleSoft applications, you can share a single search instance among the PeopleSoft systems. When sharing a search instance, keep these items in mind:

- Single sign-in needs to be implemented among all of the PeopleSoft systems sharing the search instance.
- To enable menu search in PeopleSoft Interaction Hub, you need to run the Remote Search Group Configuration plug-in in Automated Configuration Manager.

PeopleTools delivers the Remote Search Group Configuration plug-in, which is grouped under the Cluster_Template template. See Automated Configuration Management for information on running a template.

When you import remote search groups, you should consider the following items:

- Source and target PeopleSoft environments have to be mapped to the same search instance. For example, if HCM is mapped to the ES1 search instance and FSCM is mapped to the ES2 search instance, and if Interaction Hub is mapped to the ES1 search instance, then importing search groups from HCM to Interaction Hub is supported, however importing search groups from FSCM to Interaction Hub is not supported.
- Ensure that the Interaction Hub environment and the PeopleSoft application environments participating in Interaction Hub are on PeopleTools 8.55.11 release or later. However, it is not necessary that all the environments are on the same PeopleTools version, that is, while Interaction Hub is on the PeopleTools 8.55.11 release, HCM can be on 8.55.12 and FSCM can be on 8.55.13.
- While you do not need to synchronize user profiles among multiple PeopleSoft systems, if the same user ID exists on multiple systems, it must be associated with the same, individual user. That is, a user ID must be unique for all of the systems sharing the search instance, not just a single PeopleSoft application.

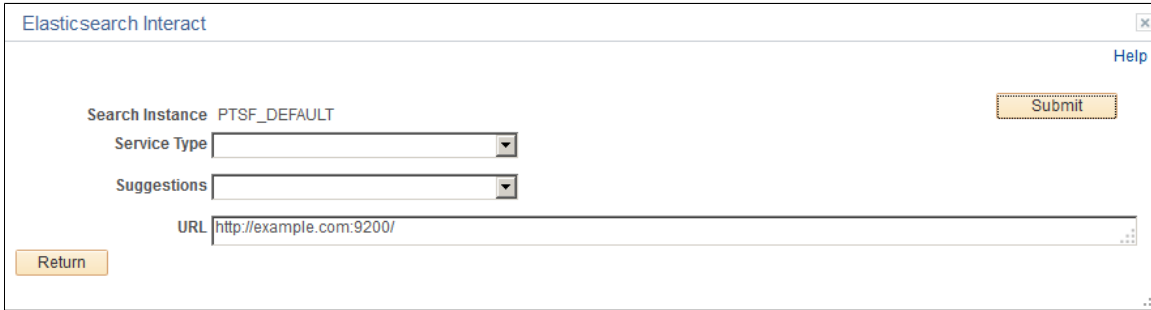
Administering Elasticsearch Using the Elasticsearch Interact Page

Elasticsearch provides several REST APIs that can be used to obtain details of a cluster or index in the PeopleSoft implementation. In the absence of a default user interface to administer Elasticsearch, PeopleSoft provides a user interface (called Elasticsearch Interact) that enables you to choose the cluster and index level options on which you want to obtain information from the Elasticsearch search engine. Based on the selected options, PeopleSoft Search Framework generates a URL, which is submitted to Elasticsearch, and the response is displayed on the Elasticsearch Interact page. The Elasticsearch Interact

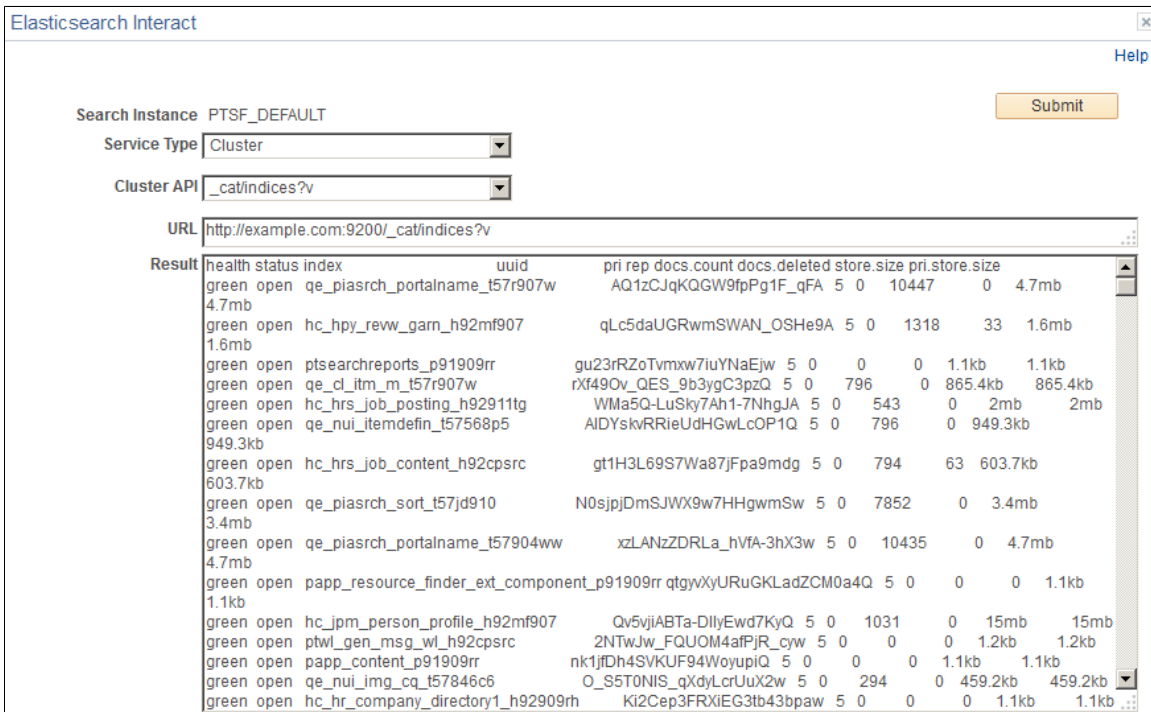
page is accessible from the Search Instance Properties page, and it provides administration requests only for those cluster and index validations that are exposed by PeopleSoft.

Access the Elasticsearch Interact page (PeopleTools, Search Framework, Search Admin Activity Guide, Configuration, Search Instance) by selecting an Elasticsearch search instance, and then selecting the Elasticsearch Interact link on the Search Instance Properties page.

This example illustrates the fields and controls on the Elasticsearch Interact page (initial). You can find definitions for the fields and controls later on this page.



This example illustrates the fields and controls on the Elasticsearch Interact page after selecting an option from the Service Type drop-down list. You can find definitions for the fields and controls later on this page.



Field or Control	Description
Search Instance	The search instance name.

Field or Control	Description
Service Type	<p>Displays a list of service type groups.</p> <p>The following service type groups are available:</p> <ul style="list-style-type: none">• Cluster - Use the Cluster service type to retrieve information related to the Elasticsearch cluster.• Count - Use the Count service type to retrieve the number of indexed documents related to search definitions.• Mapping - Use the Mapping service type to retrieve the attribute mapping information related to search definitions.• Search - Use the Search service type to retrieve index-level information related to search categories.

Field or Control	Description
Suggestions	<p>The Suggestions field displays on the initial Elasticsearch Interact page and it changes based on the service type you select.</p> <p>The Suggestions field changes to the following:</p> <ul style="list-style-type: none"> • When you select Cluster as service type, the Suggestions field changes to Cluster API, which provides cluster related options: <ul style="list-style-type: none"> • Aliases - provides information on currently configured aliases, which includes indices, filter, and routing information. • Allocation - provides a summary of the number of shards allocated to each data node and the disk space used. • Count - provides the document count of the entire cluster or individual indices. The document count includes only live documents, that is, the document count does not include deleted documents that have not yet been cleaned up by the merge process. • Health - provides the general health of a cluster. The information is similar to the cluster information when you select the <code>/_cluster/health</code> option. • Indices - provides a cross-section of each index, and it spans nodes. • Nodes - provides a topology of the cluster, which includes the Elasticsearch version, JVM version, heap, memory, load, uptime of a node, etc. • Plug ins - provides a list of running plug-ins in each node. • Shards - provides a detailed view of shards in a node. The detailed information includes whether its a primary or replica, the number of documents, the size of the shard in bytes, etc. • Cluster health - provides the general health of a cluster. It uses the color codes of green, yellow, and red. <p>For a description of the color codes used for cluster health in Elasticsearch, see Verifying the Cluster and its Nodes.</p> • Cluster statistics - provides cluster-wide statistics. The statistics include basic index metric (such as shard numbers, store size, memory usage) and information about the current nodes that form the cluster (such as number of nodes, roles, operating system, JVM versions, memory usage, CPU, and installed plugins).

Field or Control	Description
	<ul style="list-style-type: none"> Nodes statistics - provides statistics of all nodes in a cluster. By default, all statistics of a node are displayed. When you select Count or Mapping as service type, the Suggestions field changes to Search Definition, which provides a list of deployed search definitions in the search instance. When you select Search as service type, the Suggestions field changes to Search Category, which provides a list of deployed search categories in the search instance.
Method	<p>The Method field is displayed when the service type is Search. Choose the type of HTTP request that needs to be issued to the Elasticsearch server.</p> <p>Valid options are:</p> <ul style="list-style-type: none"> GET POST
URL	<p>The URL is automatically formed based on the selected options. For example, if you select Cluster as service type and select cat/indices as the cluster API, the URL is automatically formed:</p> <p><code>http://example.com:9200/_cat/indices?v</code></p> <hr/> <p>Note: The URL cannot be modified. To modify the URL, you need to change the options you have selected.</p> <hr/>
Body	<p>The Body field appears when you select POST as a method. This field must contain a JSON request.</p> <p>The contents of the Body field are passed along with the request to Elasticsearch.</p>
Result	<p>The Result area displays the response from Elasticsearch for each request. PeopleSoft Search Framework displays the response from Elasticsearch as is.</p>

Setting Namespace Aliases

Namespace aliasing enables you to re-use an index or indices built on an environment (original environment) on another environment (which is referred to as “new environment” for purpose of differentiation), thus avoiding the necessity to re-index the same search definitions on the new environment. The namespace aliasing feature is handy when you need to clone environments or when the middle-tier changes.

When you need to clone an environment, for example, you clone a test FSCM environment (FSCMTEST), where all the search definitions in FSCMTEST are already indexed to a search server, as a staging FSCM environment (FSCMSTAGING), you can re-use the index or indices of FSCMTEST from the FSCMSTAGING environment that is connected to the same search server by setting up namespace aliasing.

When the Application server or Web server changes, but is connected to the same database, you can use namespace aliasing to map the new environment to the original environment and thus re-use the index or indices of the original environment.

The namespace aliasing feature enables you to map the database name, node name, and node URI of the new environment to the original environment where the indices are built. The new environment re-uses the indices on the original environment, so users can perform search in the new environment without having to wait for indices to be built. Also, you can perform incremental indexing to update the indices from the new environment.

You can use the Namespace Alias Settings page to set namespace aliasing or you can use the PTSFConfigureSearch plug-in in Automated Configuration Manager to set namespace aliasing.

Access the Namespace Alias Settings page by selecting **PeopleTools >Search Framework >Search Admin Activity Guide >Configuration >Search Instance**. Then click the Set Namespace Aliases link.

Alternatively, access the Namespace Alias Settings page by selecting **PeopleTools >Search Framework >Administration >Custom Namespace Settings**.

This example illustrates the fields and controls on the Namespace Alias Settings page. You can find definitions for the fields and controls later on this page.

Field or Control	Description
Database Name	The name of the database in the new environment. You may enter the same database name as in the original environment.
Database Alias Name	The name of the database in the original environment.

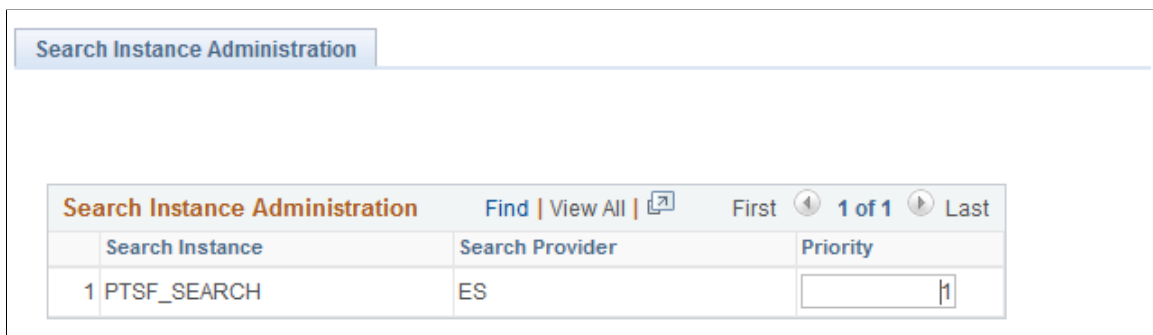
Field or Control	Description
Node Name	The name of the node in the new environment. You may enter the same node name as in the original environment.
Node Alias Name	The name of the node in the original environment.
Alias Node URI	The node URI in the original environment. When you enter the URI, follow this syntax: <i>http or https://<hostname>:<port>/psc/<sitename>/</i> Optionally, you can copy the URI entered for the node in the Content URI Text edit box on the Node Definitions page (Portal tab).
Clear Data	Clears the data entered on the page and also removes it from the database.

Administering Search Instances

To administer search instances, access the Search Instance Administration page by selecting **PeopleTools >Search Framework >Administration >Prioritize Search Instances**.

Search administrators can use the Search Instance Administration page to specify the priority for each search instance.

This example illustrates the fields and controls on the Search Instance Administration page. You can find definitions for the fields and controls later on this page.



Field or Control	Description
Search Instance	Lists the search instances created in PeopleSoft Search Framework.
Search Provider	Lists the search provider associated with a search instance.

Field or Control	Description
Priority	<p>Allows Search administrators to set the sequence in which the available search instances are to be used when any of the search instances fail. The lowest number is 1 and it has the highest priority.</p> <p>The search instance with the priority set at the lowest number has the highest priority, and it will be set as the default search instance.</p> <p>When you create a search instance on the Search Instance Properties page, the new search instance is assigned a sequence number. You can use this page to specify the sequence that you want.</p>

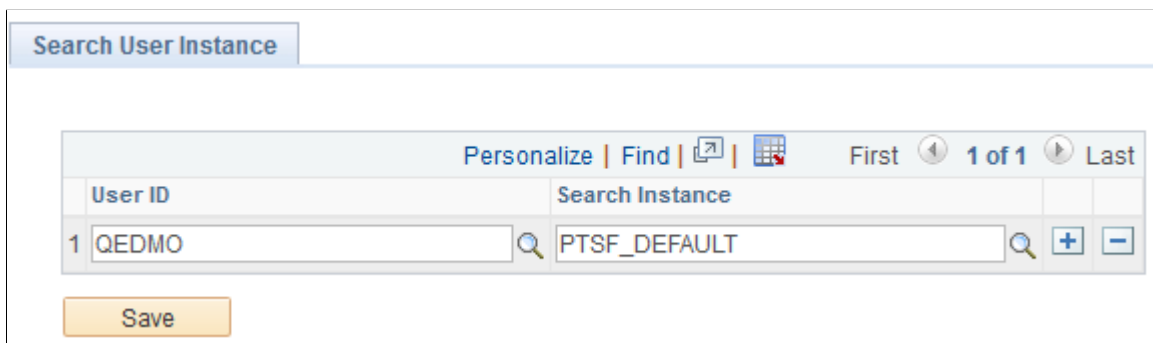
Assigning Users to a Search Instance for Validating Indexes

When you deploy an index on a new search instance, you may want to validate or test the index before you roll out the index in a production environment. Search administrators can use the Search User Instance page to assign users to a search instance on which the deployed index needs to be validated or tested.

Note: When you assign a user to a search instance using the Search User Instance page, you override the priority setting specified on the Search Instance Administration page for the particular user.

To assign users to a search instance where a deployed index needs to be validated or tested, access the Search User Instance page by selecting **PeopleTools >Search Framework >Administration >Search Instance by User**.

This example illustrates the fields and controls on the Search User Instance page. You can find definitions for the fields and controls later on this page.



Field or Control	Description
User ID	Select a user to whom you want to assign the task of validating or testing an index that is deployed.

Field or Control	Description
Search Instance	Select a search instance where you have deployed an index, which needs to be validated or tested before deploying the index to a production environment.

Note: You would use this page to test new Elasticsearch instances as they come online. It does not need to be specified for all users.

Creating a Search Instance in Automated Configuration Manager

The Automated Configuration Manager (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 SEARCH_TEMPLATE in ACM to configure an Elasticsearch search instance.

Using ACM allows you to automate the creating of search instance and deploying/indexing of search definitions. Though, you can create a search instance and deploy/index search definitions in a single execution of the template, PeopleSoft recommends to create the search instance first and then deploy/index the required search definitions.

Complete these steps to create a search instance and deploy/index search definitions:

1. Select PeopleTools, Automated Config Manager, ACM Templates, Template Definition.
2. Search for and open SEARCH_TEMPLATE.
3. On the Configuration Template Definition page, verify that the Configure Search Instance plugin is selected, and then select the Properties icon.
4. Specify the properties for the plug-in on the Configure Search Instance page.

If Kibana is installed, specify the properties associated with Kibana.

For descriptions of each of the properties, see "PTSFConfigureSrchInstance" (Automated Configuration Management).

5. Ensure that the Configure Search Instance plug-in is selected and the Deploy Search Definition plug-in is *not* selected.
6. Save the template and click Execute to begin the program run.

If the program run is successful, the system displays a success message, that is, a search instance is created. After successfully creating a search instance, proceed with the steps of deploying/indexing the required search definitions.

To verify that the search instance has been configured for Elasticsearch, select PeopleTools, Search Framework, Search Admin Activity Guide, Configuration, Search Instance.

7. On the Configuration Template Definition page, verify that Deploy Search Definition is selected and click the Properties icon.
8. Specify the properties for the plug-in on the Deploy Search Definition page and click OK.

For description of each of the properties, see "PTSFAdministerSearch" (Automated Configuration Management).

9. Ensure that the Deploy Search Definition plug-in is selected and the Configure Search Instance plug-in is *not* selected.
10. Save the template and click Execute to begin the program run.

If the program run is successful, the system displays a success message.

To verify the required search definitions are deployed and indexed, select PeopleTools, Search Framework, Search Admin Activity Guide, Administration, Schedule Search Index.

Enabling Global Search

Global Search is enabled by default based on three conditions:

- At least one search index must be deployed and crawled on the system and this search index is set as a Search Group.
- A user accessing global search must have the required permissions for the deployed search index.
- A user accessing global search must have access to the Fluid Global Search component.

When Global Search is enabled, the Global Search icon appears in the portal header.

Note: If a user does not have access to any search groups by way of their permissions lists, they will not see the Global Search icon.

Related Links

[Working with Global Search](#)

Managing General Search Options

Use the Search Options page to manage global settings for your search configuration.

Warning! These options apply to *all* deployed search indexes. Changing options after deployment may require index updates or redeploy.

Select PeopleTools, Search Framework, Search Admin Activity Guide, Administration, Search Options.

Alternatively, select PeopleTools, Search Framework, Administration, Search Instance Options or select the Search Options Config link on the Search Instance Properties page.

This example illustrates the fields and controls on the Search Options page. You can find definitions for the fields and controls later on this page.

Search Options

Search Instance

Default Values

Number of Shards	<input style="width: 50px;" type="text" value="5"/>
Number of Replicas	<input style="width: 50px;" type="text" value="0"/>
Minimum Document Count	<input style="width: 50px;" type="text" value="1"/>
Maximum Facet Children	<input style="width: 50px;" type="text" value="100"/>
Disable Facet Count	<input style="width: 50px;" type="text" value="N"/>
Min Match Percent	<input style="width: 50px;" type="text" value="100"/>
Cache Interval (Min)	<input style="width: 50px;" type="text" value="120"/>
Index Segment Size (MB)	<input style="width: 50px;" type="text" value="10"/>
Log ES requests for User ID	<input style="width: 50px;" type="text"/>
Log ES metrics for User ID	<input style="width: 50px;" type="text"/>
Enable Attachment Trace	<input style="width: 50px;" type="text" value="N"/>
Number of Indexing handlers	<input style="width: 50px;" type="text" value="5"/>
Max Sub Queue Size	<input style="width: 50px;" type="text" value="20"/>
Max Attachment Error Count	<input style="width: 50px;" type="text" value="100"/>
Full DirectTransfer	<input style="width: 50px;" type="text" value="Y"/>
Maximum Tree Facet Nodes	<input style="width: 50px;" type="text" value="1000"/>
Expand Tree Facet	<input style="width: 50px;" type="text" value="N"/>
Number of Days for Recent Ranking	<input style="width: 50px;" type="text" value="30"/>
Maximum Results in Suggestion Search	<input style="width: 50px;" type="text" value="50"/>
Recent Documents Retention (days)	<input style="width: 50px;" type="text" value="365"/>
Index Refresh Interval (sec)	<input style="width: 50px;" type="text" value="10"/>
Heartbeat Interval (sec)	<input style="width: 50px;" type="text" value="2"/>
Real Time Indexing Set Size	<input style="width: 50px;" type="text" value="300"/>
Real Time Indexing Threshold	<input style="width: 50px;" type="text" value="10000"/>

<i>Field or Control</i>	<i>Description</i>
Search Instance	Enables you to select the search instance for which you want to specify the global settings.

Field or Control	Description
Number of Shards	<p>Elasticsearch enables you to distribute the indexed data into partitions. Default value is 5.</p> <p>The value entered here is applicable to all search definition deployed from this PeopleSoft machine to Elasticsearch server unless it is overridden.</p> <p>You can override the value by selecting Index Settings on the Deploy Search Definition page (PeopleTools, Search Framework, Search Admin Activity Guide, Administration, Deploy/Delete Object).</p> <hr/> <p>Note: The value of this parameter cannot be changed when an index is in deployed status.</p> <hr/> <p>See Elasticsearch Concepts and Terminology.</p>
Number of Replicas	<p>A replica is a copy of the indexed data in a primary shard. Elasticsearch uses replicas for recovery management in the event of a failure in the primary shards.</p> <p>Set the number of copies you want to maintain of the primary shards.</p> <hr/> <p>Note: In a multi-node Elasticsearch cluster, always ensure that the replica value is set to at least 1 so that in the event of a node failure data loss does not happen. The maximum number of nodes that can be down without loss of data will be equal to the number of replicas (where replica count is less than <i>N</i>, that is, the number of nodes in a cluster.)</p> <hr/> <p>The value entered here is applicable to a search instance and to the security index (orcl_es_acl).</p> <p>See Setting the Number of Replicas.</p> <hr/> <p>Note: After the index is created, you may change the number of replicas dynamically anytime.</p> <hr/> <p>You can override the value by selecting Index Settings on the Deploy Search Definition page (PeopleTools, Search Framework, Search Admin Activity Guide, Administration, Deploy/Delete Object).</p> <p>See Elasticsearch Concepts and Terminology.</p>
Minimum Document Count	<p>Set the minimum number of search results (documents) required to display a facet. Default is 1.</p> <p>This property can be overridden by selecting Advanced Settings on the Facets page (PeopleTools, Search Framework, Search Designer Activity Guide, Search Category, Facets, Advanced Settings).</p>

Field or Control	Description
Maximum Facet Children	<p>Set the maximum number of children to be returned for each facet node. Default is <i>100</i>.</p> <p>Options to enter are:</p> <ul style="list-style-type: none"> • An integer greater than 0 to limit the number of children displayed. • <i>0</i> to indicate no children should be displayed. • <i>-1</i> to return all children (no limit). <p>This property can be overridden by selecting Advanced Settings on the Facets page (PeopleTools, Search Framework, Search Designer Activity Guide, Search Category, Facets, Advanced Settings).</p>
Disable Facet Count	<p>Enables controlling the display of the result count (document count) for each facet node. For example, if 15 search results met a facet node's criteria, (15) appears next to the facet label.</p> <p>Enter <i>Y</i> for Yes or <i>N</i> for No.</p> <ul style="list-style-type: none"> • <i>N</i>: Returns the document count. Default is <i>N</i>. • <i>Y</i>: Shows no document count.
Min Match Percent	<p>Enables controlling the display of search results based on the percentage of search keywords found in a document. The default value is 100.</p> <p>A value of 100 will display a document in search results only when all the search keywords are found in the document.</p> <p>A value of 50 will display a document in search results if at least 50% of the search keywords are found in the document.</p>
Cache Interval (Min)	<p>Set the duration of time (in minutes) for storing the security cache of a user. The default value is 120 minutes.</p> <p>For the security values, the system registers the time it is cached in the Elasticsearch index. If the duration of search time and cached time exceeds the cache interval or if the values are not available in the cache, the security values are obtained by a callback to PeopleSoft.</p> <p>The <i>acl.cache.delete.interval</i> property in the <i>elasticsearch.yml</i> configuration file runs a delete thread every 24 hours to clean the stale cache that has exceeded the cache interval. The default value is 24 hours. You may modify the value to run the delete thread at a frequency that you want.</p>

Field or Control	Description
Index Segment Size (MB)	<p>Set the size, in megabytes, by which the system segments (or chunks) the indexing data feed. Default is 10 megabytes.</p> <p>When the index segment size goes higher, the consumption of available resources such as JVM size, CPU, etc. also goes high. In a parallel processing enabled environment (that is, an environment with more than one sub handler), the segment size should be adjusted to the optimum level based on the available resources.</p>
Log ES requests for User ID	<p>Enter a user ID for which you want to log Elasticsearch requests and responses. The Elasticsearch requests and responses are written to the trace log file.</p> <hr/> <p>Note: Only one user ID can be entered at a time.</p> <hr/>
Log ES metrics for User ID	<p>Enter a user ID for which you want to log query performance metrics, that is, the time taken by Elasticsearch to execute a query. The Elasticsearch query performance metrics are written to the trace log file as part of the Elasticsearch response.</p> <hr/> <p>Note: Only one user ID can be entered at a time.</p> <hr/>
Enable Attachment Trace	<p>Enables controlling the detailed logging of each task on the trace file. You should consider setting this property to <i>Y</i> if you are planning to use trace files to troubleshoot any errors associated with the transfer of search data from PeopleSoft to the Elasticsearch search engine.</p> <p>Enter <i>Y</i> (yes) if you want to enable detailed logging of each task on the trace file.</p> <p>Enter <i>N</i> (no) to disable logging of each task on the trace file.</p>

Field or Control	Description
<p>Number of Indexing handlers</p>	<p>This property is associated with the direct transfer of search data from PeopleSoft to the search engine, which bypasses the Integration Gateway. The default value is 20.</p> <hr/> <p>Note: When Full Direct Transfer is enabled, this property is applicable for search documents with attachments and without attachments, but when Full Direct Transfer is not enabled, this property is applicable only for search documents with attachments.</p> <hr/> <p>Set the number of handlers that will be used to transfer data asynchronously to search engine. Each handler can transfer data to the search engine in parallel. For example, if you set the value of handlers to 20, it means that during indexing a maximum of 20 handlers are created. The handler value should be less than the bulk thread queue size on Elasticsearch (default value of bulk thread queue size is 50).</p> <hr/> <p>Note: The number of handlers should be set based on the size and volume of the attachments and the heap size that is set for Elasticsearch. You need to remember that a large number of parallel threads can cause congestion in Elasticsearch queue, which directly reduces the heap availability. If you are getting a number of Elasticsearch rejections when indexing, reduce the number of handlers to 5.</p> <hr/> <p>See Configuration.</p>
<p>Max Sub Queue Size</p>	<p>This property is associated with Integration Gateway, so this property is applicable only when the Full Direct Transfer property is set to N, that is, the search documents without attachments are transmitted through the Integration Gateway.</p> <p>Maximum subscription queue size indicates the maximum number of IB subscription queues that are running in parallel for indexing a given search definition. For example, if your environment has 10 subscription handlers in IB, and if you want to dedicate only three subscription handlers for search indexing purpose, you should enter the value as 3 for this property.</p> <hr/> <p>Note: The value entered for this property depends on the actual number of subscription handlers services activated in application server. Before you enter a value for this property, you should consider the actual number of subscription handlers services activated in application server. If you enter a value greater than the actual number of subscription handlers services, the system will use all the available subscription handlers services even though the Sub Queues would be partitioned based on the value entered for this property.</p>

Field or Control	Description
Max Attachment Error Count	<p>This property is associated with the direct transfer of search data from PeopleSoft to the search engine, which bypasses the Integration Gateway</p> <p>The maximum number of exceptions that should be displayed on the Attachment Transfer Exception Details page.</p> <p>The default value is 100.</p> <p>When the number of exceptions exceeds the maximum number, the PTSF_GENFEED Application engine program is aborted, and you must re-submit the Application Engine program.</p> <p>See Viewing Search Index Build Process Details.</p> <p>See Configuration.</p>
Full Direct Transfer	<p>The default value is <i>Y</i>.</p> <p>When the Full Direct Transfer setting is set to <i>Y</i> (yes), search documents with or without attachments are directly transmitted to the Elasticsearch search engine.</p> <p>If you set the Full Direct Transfer setting to <i>N</i> (no), then search documents without attachments are transmitted through the Integration Gateway, but the search documents with attachments are directly transmitted to the Elasticsearch search engine.</p> <p>See Configuration.</p>
Maximum Tree Facet Nodes	<p>Set the maximum number of tree facet nodes to be returned for each facet. Default is <i>1000</i>.</p> <p>For the Search Test page, this property can be overridden by selecting Additional Details in the Facet Request group box. (PeopleTools, Search Framework, Search Admin Activity Guide, Search Test Page, Facet Request. Additional Details)</p>
Expand Tree Facet	<p>This property specifies whether the tree facets should be expanded or collapsed. The default value is <i>N</i>, that is, collapsed.</p> <p>By setting the default behavior to collapsed, the tree facet displays a higher number of level 1 nodes.</p>

Field or Control	Description
<p>Number of Days for Recent Ranking</p>	<p>The Number of Days for Recent Ranking property is associated with the suggestions that are provided on the search box.</p> <p>This property specifies the number of days the pages visited by a user should be considered for ranking the pages as recently visited.</p> <p>The default value is 30, that is, the pages visited by a user during the last 30 days are considered for ranking. This is a dynamic property, that is, you can change its value without having to redeploy the index.</p> <p>For more information on ranking of pages, see Working with Suggestions Search.</p>
<p>Maximum Results in Suggestion Search</p>	<p>The Maximum Results in Suggestion Search property is associated with the suggestions that are provided on the search box.</p> <p>This property specifies the maximum number of suggestions that should be displayed, which includes the recently visited pages and suggested pages.</p> <p>The default value is 50 and it is also the maximum supported value, that is, you cannot increase the value. However, you may reduce the number of suggestions displayed. This is a dynamic property, that is, you can reduce the value without having to redeploy the index.</p> <p>For more information on suggested menu items, see Working with Suggestions Search.</p>
<p>Recent Documents Retention (days)</p>	<p>The Recent Documents Retention property is associated with the suggestions that are provided on the search box, which includes both recent menu suggestions and recent content keyword suggestions provided the corresponding indexes are enabled for real-time indexing.</p> <p>This property specifies the number of days the recent documents should be retained.</p> <p>The default value is 365.</p> <hr/> <p>Warning! The indexes may tend to grow large considering that the indexes hold the recently accessed menu and content keyword data for a year.</p> <hr/> <p>This is a dynamic property, that is, you can change its value without having to redeploy the index.</p>

Field or Control	Description
<p>Index Refresh Interval (sec)</p>	<p>Set the duration of time (in seconds) after which indexes in Elasticsearch are refreshed or saved. That is, after every specified number of seconds, the indexes are refreshed or saved. Elasticsearch aggregates all the documents received for indexing up to the specified interval and then saves to the index.</p> <p>This setting is to improve the performance of the indexing in Elasticsearch. The lowest value possible is 1 second.</p> <p>The default value is 10 seconds.</p> <p>If you set the interval duration to a higher value, it improves the indexing speed, but if you set the interval duration to a lower value, the indexed data is closer to being near real time. However, a lower value also means that Elasticsearch server is busy performing this task.</p> <hr/> <p>Warning! The Index Refresh Interval property is applicable to all deployed search definitions irrespective of whether real time indexing is enabled or not in a search definition.</p>
<p>Heartbeat Interval (sec)</p>	<p>The Heartbeat Interval property is associated with real time indexing.</p> <p>Set the duration of time (in seconds) for the real time indexing server process to sleep when no rows are submitted for processing. The server process looks for any new rows to index after the specified interval. You should set this property in accordance with the frequency of online transactions.</p> <p>The default value is 2 seconds.</p> <hr/> <p>Note: If you modify the value of this property, you should restart the Process Scheduler domain.</p>
<p>Real Time Indexing Set Size</p>	<p>Specify the maximum number of transactions in one set for real time processing.</p> <p>The default value is 300. That is, a set can have a maximum of 300 transactions.</p> <p>Any increase in this value will adversely affect the runtime execution of real time indexing.</p> <p>See Understanding Set Based Processing in Real Time Indexing.</p> <hr/> <p>Note: If you modify the value of this property, you should restart the Process Scheduler domain.</p>

<i>Field or Control</i>	<i>Description</i>
Real Time Indexing Threshold	<p>Specify the highest count of rows or transactions in the primary staging table.</p> <p>The Real Time Indexing Threshold property is used in the real time indexing maintenance AE program (PTRTI_TRUNC) to de-fragment the staging table when the count of rows exceeds the threshold value.</p> <p>See Using an Application Engine Maintenance Program.</p> <hr/> <p>Note: If you modify the value of this property, you should restart the Process Scheduler domain.</p> <hr/>

Administering Search Definitions and Search Categories

This section contains an overview of search definitions and search categories.

Understanding Search Definition Administration

You create search definitions using the Search Framework Designer Activity Guide after the data is identified using PeopleSoft Query and Connected Query. You use the Search Framework Administration Activity Guide to deploy the search definitions and manage the search definitions on the search engine.

Before end users can run searches against the search indexes, the search definitions need to be deployed to the search engine so that the search engine can create the structure of the search index based on the search definition, crawl the defined search criteria, and populate the index with the results of the search definition query. When your search definitions are deployed on the search engine, they become a source.

Once the search definitions are deployed to the search engine, you use the Search Framework administration Activity Guide to manage them by sorting them in search categories, updating them, undeploying them, or deleting them as needed.

Working with Search Definition Deployment

To manage search definitions, access the Search Definitions page by selecting PeopleTools, Search Framework, Search Admin Activity Guide, Administration, Deploy/Delete Object.

This example illustrates the fields and controls on the Deploy Search Definition page. You can find definitions for the fields and controls later on this page.

Definition	Description	Deploy Status	Crawl Status	Run Date/Time	User/Run Control ID	Index Settings
1 PTPORTALREGISTRY	Navigator	Deployed	Successful			Index Settings
2 PTSEARCHREPORTS	Reports	Deployed	Successful	09/03/2016 9:47:09AM	QEDMO / PTSEARCHREPORTS	Index Settings
3 PTWL_GEN_MSG_WL	Worklist Note	Deployed	Successful			Index Settings

Field or Control	Description
Search Instance	Use to display only those search definitions associated with a particular search instance.
Search Category Name	Use to display only those search definitions associated with a particular search category.
Filter by name	Use to enter full or partial definition names to reduce the list of search definitions appearing in the grid. Note: The values are case-sensitive.

Field or Control	Description
Deploy Search Definition (grid)	<p>Displays all existing search definitions defined in the Search Framework. Columns in the grid are:</p> <ul style="list-style-type: none"> • Definition. Displays the name of the search definitions you've created. Deployable search definitions are having a data source created in Query or Connected Query, as well as being defined in the Search Framework Designer activity guide. • Description. Displays your search definition description. • Deploy Status. Displays the state of search definition deployment: A search Definition is either Deployed or Undeployed. • Crawl Status. Displays the status of the index build process for a search definition (Success, Fail, Not Crawled). • Run Date\Time. Displays the date and time of the most recent index build process run. • User\Run Control ID. Displays the user ID and the run control ID used to run the index build process. • Update Status. After running the Report Sync Issues audit, the results appear in this column indicating any required updates. • Audit Status. After clicking Audit for a selected search definition, the Audit Status column appears displaying the result (<i>Success, Fail</i>). • Audit Report. If the Audit Status is <i>Fail</i>, the Audit Report column appears providing a link to the audit report for review.

Field or Control	Description
Index Settings	<p>Click to display the Advanced Index Settings dialog box where you can override some of the settings specified on the Search Options page.</p> <p>If you click the Index Settings link for a deployed search definition, you can change only the number of replicas.</p> <p>If you click the Index Settings link for a search definition that is not deployed, you can update the following settings:</p> <ul style="list-style-type: none"> • Override Settings - Select this option if you want to override the settings specified on the Search Options page. • Number of Shards - Specify a new number of shards, if desired. • Number of Replicas - Specify a new number of replicas, if desired. <hr/> <p>Note: In a multi-node Elasticsearch cluster, always ensure that the replica value is set to at least 1. For details, see Managing General Search Options.</p> <hr/> <ul style="list-style-type: none"> • Enable Synonyms - Select to enable synonym search for the index. This option is enabled by default. For details, see Disabling Synonym Search for an Index. <hr/> <p>Note: Before you change any of the values here, Oracle PeopleSoft recommends that you familiarize yourself with the functioning of the Elasticsearch cluster in your implementation.</p> <hr/>
Select All	<p>Click to select all the rows within the grid.</p> <hr/> <p>Note: If all rows in the grid have been selected, the button title changes to Deselect All.</p> <hr/>
Report Sync Issues	<p>Runs an audit checking routine which reports where properties for a search definition differ between PeopleSoft and the search engine. If you have updated a deployed search definition, it will no longer be synchronized with the representation on the search engine. Results of this check appear in the Update Status column (which, in the case of differences between the corresponding definitions becomes the View Report column). This applies only to deployed definitions.</p>

Field or Control	Description
Audit	<p>Click to run a series of audits on the selected search definitions to ensure they have all requirements in place.</p> <p>After running the audit process, these additional columns appear in the Deploy Search Definition grid: Audit Status and Audit Report. The Audit Status column will indicate <i>Success</i> or <i>Fail</i>. If the result is <i>Fail</i>, then the Audit Report column appears, displaying an Audit Report link, which you can click to review the audit results and correct the reported issues.</p> <p>See the section below: “Running the Search Definition Audit Process.”</p>
Deploy	<p>Deploy the PeopleSoft search definitions selected to the search engine, where it becomes a searchable data source.</p> <hr/> <p>Note: You can deploy multiple search definitions at a time by selecting the corresponding check box in the first column of the grid and then choosing Deploy.</p> <hr/> <p>Note: When you deploy a search definition, the search category of the same name is automatically deployed.</p> <hr/>
Undeploy	<p>Click to remove the representation of the selected search definitions from the search engine.</p> <hr/> <p>Note: You can undeploy multiple search definitions at a time by selecting the corresponding check box in the first column of the grid and then choosing undeploy.</p> <hr/> <p>Note: When you undeploy a search definition, the search category of the same name is automatically undeployed.</p> <hr/> <p>Note: When you undeploy a search definition, the index is not searchable, that is, a search against that index will not display any search results.</p> <hr/> <p>Warning! Oracle PeopleSoft does not recommend undeploying large indexes because re-indexing them requires additional time. However, for your business purposes, if you need to undeploy large indexes, it is recommended to undeploy them during the non-peak hours for performance reasons.</p> <p>The status of the search definition displays undeployed only after the actual deletion is completed. Undeploying a search definition mapped to a common index may take additional time.</p> <hr/>

Field or Control	Description
Update	Click to update the representation of the search definition on the search engine. If you have modified a deployed search definition, the PeopleSoft version and the version on the search engine will no longer match. Definitions that need to be updated will appear in the synchronization report. Clicking Update synchronizes the search engine representation of the definition to match the PeopleSoft version.
Delete	<p>Deletes the search definition from the Search Framework. This removes the search definition from the search definition list and from the search categories to which it is mapped. The underlying query and connected query are unaffected.</p> <hr/> <p>Note: When you delete a search definition, the indexed data for that search definition is no longer available, so searches against that indexed data will not display any results.</p> <hr/> <p>Important! Oracle PeopleSoft recommends that you do <i>not</i> delete delivered search definitions.</p> <hr/> <p>Warning! Delete is not reversible! Do not use this function unless you are certain it is what you want.</p> <hr/>

Running the Report Sync Issues Process

By selecting an index (or several indexes) and selecting the Report Sync Issues button, Search Framework will add a column to the Deploy/Delete Object grid to report differences found between Search Framework and the default search engine.

Depending on the status of your search definitions in both the PeopleSoft system and the search engine, the results of the synchronization report will vary. This table describes some likely scenarios and the recommended action.

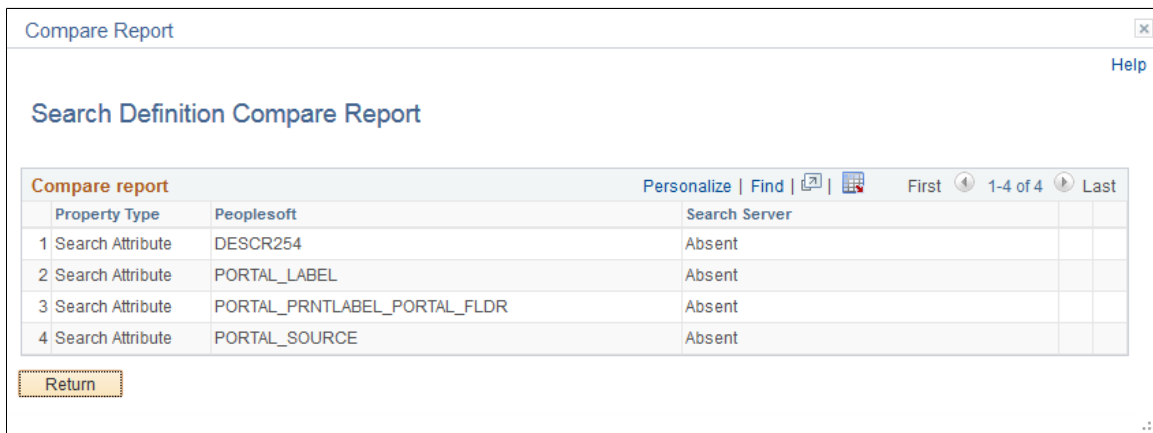
Update Status	Recommended Action
No update required for the definition.	No action is required. The definitions are identical between the two systems.
Update required definition out of sync.	Click the View Report link in the View Report column, and examine the reported differences. Resolve the differences as needed between the two systems, or click Update.
Definition unavailable in server.	Click the Reset Definition for Deploy link in the View Report column. This removes the previous definition on Elasticsearch or any remaining settings related to it, making it available to be deployed again.

Update Status	Recommended Action
Undeployed definition. Choose a deployed definition.	Ignore the report if the definition is purposely undeployed, or deploy the definition.

Note: You need to take the appropriate action to fix the compare issue. This may include, but is not limited to using the Update button, manually deploying the category or undeploying the search definition and reloading it.

When you select View Report, you may see a screen similar to this:

This example illustrates the fields and controls on the Search Definition Compare Report page. You can find definitions for the fields and controls later on this page.



Field or Control	Description
Property Type	Displays properties that have found to differ between the two systems. The values appearing in this column can be numerous, representing all the possible settings for a search definition, such as user ID, category, call back URL, and so on.
PeopleSoft	Displays the value defined for the search definition within the PeopleSoft system.
Search Server	Displays the value stored for the property on the search server.
OK	Click OK to leave the report and return to the previous page.

Running the Search Definition Audit Process

After selecting a search definition and clicking Audit, the system runs a series of audits on the search definition. The types of audits include checking to make sure:

- The underlying query or connected query exists.
- The user has permission to the query and the referenced records.
- The indexed fields exist in the query or connected query.
- The security App Class exists (if referenced in search definition).
- The pre processing AE library exists (if referenced in search definition).
- The post processing AE library exists (if referenced in search definition).
- The delete query exists (if referenced in search definition) and the user has permission to open the delete query.
- A search category with the same name as the search definition exists.
- The search definition belongs to at least one search category.



If the audit process finds no issues, the Audit Status column indicates Success. If the audit process does find issues, the system displays the Audit Report column containing a link to the audit report.

This example illustrates the audit related columns that appear after clicking Audit for a selected search definition.

Audit Status	Audit Report
Fail	Audit Report
Success	

If the audit discovers issues, the audit status gets set to *Fail*, and the **Audit Report** column appears. To review the issues on the Search Definition Audit Report page, click the **Audit Report** link.

This example illustrates the fields and controls on the Search Definition Audit Report page.

Search Definition Audit Report			
Audit Report		Personalize Find  	
		First	Last
	Audit Type	Audit Comment	
1	Search category check	Search category does not exist with the same name as search definition	+ -

Updating Definitions

In the event of updates being required, clicking Update makes the version of the search definition on the search server match the properties for that definition stored within the PeopleSoft system. After the process runs, you should see in the Update Status column the following message:

Updated definition <definition name>

Working with Search Categories

Search categories enable you to group search definitions within logical, manageable groups. You define search categories within the Search Framework Designer Activity Guide.

To manage search categories, access the Deploy Search Categories page by selecting PeopleTools, Search Framework, Search Admin Activity Guide, Administration, Deploy/Delete Object, Deploy Search Category.

This example illustrates the fields and controls on the Deploy Search Category page. You can find definitions for the fields and controls later on this page.

Select	Definition	Description	Deploy Related Definition	Deploy Status
1 <input type="checkbox"/>	PTPORTALREGISTRY	Navigator	Deploy Related Definition	Auto Deployed
2 <input type="checkbox"/>	PTSEARCHREPORTS	Reports	Deploy Related Definition	Auto Deployed
3 <input type="checkbox"/>	PTWL_GEN_MSG_WL	Worklist Note	Deploy Related Definition	Auto Deployed
4 <input type="checkbox"/>	QE_ALL_CAT_CQ_SD1	QE All cat CQ SD1	Deploy Related Definition	Undeployed

Field or Control	Description
Search Instance	Use to display only those search categories associated with a particular search instance.
Filter by name	Use to enter full or partial definition names to reduce the list of search categories appearing in the grid.

Field or Control	Description
Deploy Search Category (grid)	<p>Displays all existing search categories defined in the Search Framework. Columns in the grid are:</p> <ul style="list-style-type: none"> • Deployable Definition Name. Displays the name of the search categories you've created. Deployable search categories are those having been defined in the Search Framework Designer activity guide. • Description. Displays your search category description. • Deploy Related Search Definition. Click to display the Deploy Search Definition page with the grid populated only with search definitions within the selected search category. • Deploy Status. Displays these states of search object deployment: Deployed, Undeployed or Auto Deployed. Search categories will be automatically deployed (Auto Deployed) if a search definition with the same name is deployed. You need to deploy search categories covering multiple search definitions manually. That is, if the search category is undeployed, it must be manually deployed from this page. • Update Status. If you run a Report Sync Issues process from the Deploy Search Definition page, the process populates this column too if there are any differences between search categories on the two systems (PeopleSoft and the search engine).
Select All	<p>Click to select all the rows within the grid.</p> <hr/> <p>Note: If all rows in the grid have been selected, the button label changes to Deselect All.</p> <hr/>
Report Sync Issues	<p>Runs an audit checking routine which reports where properties for a search category differ between PeopleSoft and search engine. If you have updated a deployed search category it will no longer be synchronized with the representation on the search engine. Results of this check appear in the Update Status column (which, in the case of differences between the corresponding categories becomes the View Report column). This applies only to deployed categories.</p>
Deploy	<p>Deploy the PeopleSoft search category to the search engine.</p> <hr/> <p>Note: You can deploy multiple search categories at a time by selecting the corresponding check box in the first column of the grid and then choosing Deploy.</p> <hr/> <p>Important! You will receive a warning message when all search definitions contained within a category are not deployed, but as long as one search definition is deployed, you can deploy the search category.</p> <hr/>

Field or Control	Description
Undeploy	<p>Click to remove the search category from the search engine.</p> <p>You cannot undeploy an auto-deployed search category. You must undeploy the search definition of the same name and the auto-deployed category will be undeployed.</p> <hr/> <p>Note: The search category will be removed from the search engine, but it will still exist within the PeopleSoft system.</p> <hr/> <p>Note: This does not undeploy the search definitions associated with a search category.</p> <hr/>
Update	<p>Click to update the representation of the search category on the search engine. If you have modified a deployed search category, the PeopleSoft version and the version on the search engine will no longer match. Categories that need to be updated will appear in the synchronization report. Clicking Update synchronizes the search engine representation of the category to match the PeopleSoft version.</p>
Delete	<p>Deletes the search category from the Search Framework.</p> <hr/> <p>Note: This does not delete the search definitions associated with a search category.</p> <hr/> <p>Warning! This is not reversible! Oracle PeopleSoft does <i>not</i> recommend deleting delivered search categories.</p> <hr/>

Working with Search Indexes

This topic provides an overview of index building process and discusses how to build an index.

Understanding the Index Build Process

In Elasticsearch, an index is a logical namespace that maps to one or more primary shards and can have zero or more replica shards.

Before end users can submit search requests against the Search Framework deployed objects, the search indexes must first be built on the search engine. Prior to the index being built, a deployed search definition is an empty shell, containing no searchable data. A search index needs to be built for search definitions.

An Application Engine program, PTSF_GENFEED, builds the search index data and pushes it to Integration Broker, which makes it available to Elasticsearch.

When an attachment is specified in a search definition, PeopleSoft Search Framework transfers the attachment data directly to the search engine using cURL and libcurl libraries and does not use the Integration Broker.

Creating a search index with the Search Framework involves the following technologies:

- Search Framework
- Application Engine
- Query/Connected Query Manager
- Process Scheduler
- Integration Broker
- cURL and libcurl Libraries

Once you invoke a search index build from the Build Search Index page, the system automatically completes these general steps:

1. The Schedule Search Index page initiates the PTSF_GENFEED Application Engine program.
2. The Pre Processing Application Engine program defined for the search definition runs.
3. PTSF_GENFEED Application Engine program runs the query (PeopleSoft Query or Connected Query) associated with the search definition.

Multiple PTSF_GENFEED Application Engine programs will run for a Connected Query in cases where users specify the number of partitions. See [Partitioning Application Data in the Search Index Build Process](#) for more information.

4. The query output becomes a data source for PeopleSoft Search Framework.

If a query field is masked, Search Framework treats the masked field in two ways:

- If a query field is masked, but bypassed for Search Administrator, then all data is indexed.
- If a query field is masked, but not bypassed, then Search Framework removes the field from indexing.

Note: If you want search query masking to take effect for data that has already been indexed, you should do a full indexing.

See "Using Query Administration" (Query) for more information on bypassing a masked field.

5. The PeopleSoft Search Framework converts the query output to the JSON format and pushes the data to Elasticsearch and the Delete query (for incremental indexing) defined for the search definition runs.

During this step, the following steps are also executed:

- If it is full indexing, a delete request is sent to Elasticsearch to clear any indexed data that is present.

- If Full Direct Transfer is not enabled (segments of documents without attachments), on clearing indexed data, the data to be indexed is pushed to an Ordered IB Queue (PTSF_ES_SEND_Q) partitioned with segments. IB Queue processes the data asynchronously based on the partition.
- If Full Direct Transfer is enabled, on clearing indexed data, the data is transferred to Elasticsearch directly using cURL and libcurl libraries.
- The Delete Query is invoked when incremental indexing is run after a full index is built. Stale data is removed from the index as selected by Delete Query defined on the search definition.

Note: Each segment pushes the data in JSON format to Elasticsearch synchronously (that is, waits for response from Elasticsearch) and the response is read and acknowledged in PeopleSoft through exception handling.

6. The Post Processing Application Engine program defined for the search definition runs.

Related Links

[Monitoring Elasticsearch System Metrics and Indexing Metrics](#)

Building Search Indexes

A run control ID is bound to both, a search definition and a search instance.

Access the Build Search Index page. (Select PeopleTools, Search Framework, Search Admin Activity Guide, Administration, Schedule Search Index.)

This example illustrates the fields and controls on the Build Search Index page. You can find definitions for the fields and controls later on this page.

Build Search Index

Run Control ID ITEM_DEFN Report Manager Process Monitor Run

*Search Instance PTSF_DEFAULT

*Search Definition PTSEARCHREPORTS

Index By Date Range

Language Option

Base Language
 All Languages

Indexing type

Full index
 Incremental Index

Full Indexing Criteria

Number of days to index 0 Index Start Date 01/01/1900

To build a Search Framework search index:

1. Select PeopleTools, Search Framework, Search Admin Activity Guide, Schedule Search Index.

2. Enter a run control ID.
3. On the Build Search Index page, select the appropriate options.

Field or Control	Description
Search Instance	Select the search instance in which you want to build the search index.
Search Definition	<p>Select the search definition for which you are building the search index.</p> <hr/> <p>Note: Once the Run Control is executed, the page becomes read-only. To change any parameters of the page, you need to create a new run control.</p> <hr/>
Index By Date Range	<p>Select the Index By Date Range check box if you want to index data based on date range. The date range option enables you to partition the search data and run parallel indexing processes.</p> <p>For more information on the date range option, see Partitioning Application Data in the Search Index Build Process.</p>
Language Option	<p>Select one of the following:</p> <ul style="list-style-type: none"> • <i>All Languages</i>. Builds an index for each language enabled on the database. • <i>Base Language</i>. Builds an index only for the base language defined for the database. <p>Your selection depends on the languages enabled for your database and the languages through which you anticipate end users using to perform searches.</p>

Field or Control	Description
Indexing Type	<p>Select one of the following:</p> <ul style="list-style-type: none"> • <i>Full index.</i> Crawls all transactional data specified by the query criteria and rebuilds the entire index. This option requires the most time. This option must be selected the first time an index is built. • <i>Incremental index.</i> Updates the existing index and adds only documents associated with rows that have been added or updated since the last index build or update. The system determines the required updates based on a comparison between the timestamp of the index and the “last updated” field for the data row. <p>Items are also removed from the index as designated by the delete query on the search definition.</p> <hr/> <p>Note: You must create one run control for incremental indexing, scheduled to run very frequently, and create another run control for full index rebuilding set to run less frequently. For example, incremental indexing might run daily, where a full index rebuild may be set to run every six months.</p> <hr/>
Full indexing Criteria	<p>Applies only to full index builds. Displays the span of time in which the system will retrieve and index application data for this search definition. The date span is defined on the Advanced Settings tab for the search definition. This enables you to limit index processing for data that could potentially generate very large indexes, if needed.</p> <p>For more information on this option see, Working With Advanced Settings.</p> <p>If you plan to use the partitioning option through the Date Range grid, you must note that the indexing date span is overridden by the date span entered in the Date Range grid. For more information, see Partitioning Application Data in the Search Index Build Process.</p>

4. Select Run.
5. Use Process Monitor to verify program completion and success.

Note: Once the Run Control is executed, the page becomes read-only. To change any parameters of the page, you need to create a new run control.

Viewing Search Index Build Process Details

After the PTSF_GENFEED program begins to run, you can view the details which display on the Build Search Index page for that run control ID.

Note: The Details section shows the results of the most *recent* feed generation for this Search Definition. It may not be the same run control ID as the one you selected. If the run control ID differs from the one you selected, it will be highlighted.

This example illustrates the fields and controls on the Previous schedule details. You can find definitions for the fields and controls later on this page.

Previous schedule details

Previous schedule status **Success**

Last Successful Index Built On **09/09/2016 2:29:31AM**

AE Status **Schedule success**

Details

Process Instance **627** [Message Log](#)

Run Control ID **PTSEARCHREPORTS**

User ID **QEDMO**

Indexing type **Incremental Index**

Language Option **Base Language**


<i>Field or Control</i>	<i>Description</i>
Previous schedule status	<p>Displays the status of the most recently executed index build process.</p> <p>When errors occur while processing attachments, an Error link is displayed. Click the Error link to view the exceptions on the Attachment Transfer Exception Details page.</p> <p>See Viewing the Exception Details in Attachment Processing.</p>
Last Successful Index Built On	<p>Displays the date when the index was built successfully. If an index was never built, it displays <i>Never</i>.</p>
AE Status	<p>Indicates whether the most recently executed PTSF_GENFEED Application Program ran to completion.</p> <p>If the index build fails, this field displays the step where the process failed.</p>
Resume schedule	<p>Appears in the case where the index build failed. If the index build fails, you can use the Rebuild index option to regenerate the index.</p> <p>If the build fails, see Handling Common Errors to determine the cause of the failure to build an index.</p>
Process Instance	<p>Indicates whether the system successfully created the output that is required to populate the index.</p>

Field or Control	Description
Run Control ID	Displays the run control ID used to generate the index.
User ID	Displays the User ID who ran the run control, which may be a different user ID than the user who created the run control.
Indexing Type	Displays the index type (Full Index or Incremental Index).
Language Option	Displays the language options selected for the index build process (All Languages or Base Language).

Viewing the Exception Details in Attachment Processing

In the Previous schedule details section, select the Error link to view the details of exceptions with respect to attachment processing.

This example illustrates the fields and controls on the Attachment Transfer Exception Details page. You can find definitions for the fields and controls later on this page.



Field or Control	Description
Process Instance	The process instance ID of the PTSF_GENFEED Application Engine program that is causing the exception.
Request	The request sent by the PTSF_GENFEED Application Engine program to Elasticsearch.
Response	The response received by the PTSF_GENFEED Application Engine program from Elasticsearch.

Field or Control	Description
Rerun	<p>Use the Rerun button to resubmit the selected transactions to Elasticsearch. In the case where the maximum exceptions limit is not reached, when you select the Rerun button, the transactions are directly submitted to Elasticsearch without re-execution of the PTSF_GENFEED Application Engine program.</p> <p>In the case where the maximum exceptions limit is reached, the Rerun button acts as a re-submit of the PTSF_GENFEED Application Engine program.</p>

Related Links

[Working with Attachment Properties](#)

Partitioning Application Data in the Search Index Build Process

Indexing of search definitions that have reasonably large data volume can pose performance issues. To overcome such performance issues, PeopleSoft Search Framework enables you to partition the application data using the Index by Date Range option. After the date range and number of partitions are specified, the Partition Manager partitions the date range into equal date spans based on the number of partitions. The search index build process creates multiple PTSF_GENFEED Application Engine programs to run the query on the partitioned data. The search index build process also creates multiple run control IDs based on the number of partitions that enables you to re-run any of the run control IDs associated with a partition. The partitioning of application data is performed based on the date range and the number of partitions that are specified.

You can specify the number of partitions, but you cannot control the date span within a date range that is allotted for each partition. The Partition Manager allots the date span for the number of partitions specified by you. However, the Partition Manager does not consider data concentration to come up with the ranges. Hence, in case of large data volumes, the data distribution may not be equal for each partition range, resulting in some processes taking huge loads. Search Framework allows you to adjust these date ranges as needed to create an evenly spread data volume to index.

Partitioning of application data works best when the volume of data is evenly spread over the specified date range. The purpose of partitioning is to divide the data volume into chunks and to run a connected query on the partitioned data, so when you plan to use the partitioning option, you must consider whether the search documents for a specific search definition are spread evenly over a specific date range. That is, if within a date range, for example, 01/01/2017 to 03/31/2017, a huge volume of search data is concentrated in the month of March and the other two months have few attachments, and you specify 3 partitions, then the partitioning may not be very effective because only one partition may get the bulk of the search documents, so you may still face performance issues.

The naming convention followed for run control IDs that are generated for each partition is as follows:
 _<USERID>_RUNCNTL_<TIME>_<SEQUENCE NUMBER>

For example, _QEDMO_RUNCNTL_020022_2

Note: Partitioning is available for both full and incremental indexing. However, Oracle PeopleSoft recommends that you use partitioning only with full indexing.

This example illustrates the fields and controls on the Build Search Index page after enabling the Index by Date Range option. You can find definitions for the fields and controls later on this page.

Build Search Index

Run Control ID: PTPORTAL_DTRANGE [Report Manager](#) [Process Monitor](#)

*Search Instance

*Search Definition

Index By Date Range

Language Option

Base Language
 All Languages

Date Range

Start Date Number Of Partitions

End Date

	Start Date	End Date
1	<input type="text" value="01/01/2000"/> <input type="button" value="Calendar"/>	<input type="text" value="07/01/2000"/> <input type="button" value="Calendar"/>
2	<input type="text" value="07/02/2000"/> <input type="button" value="Calendar"/>	<input type="text" value="12/31/2001"/> <input type="button" value="Calendar"/>
3	<input type="text" value="01/01/2001"/> <input type="button" value="Calendar"/>	<input type="text" value="07/01/2001"/> <input type="button" value="Calendar"/>
4	<input type="text" value="07/02/2001"/> <input type="button" value="Calendar"/>	<input type="text" value="12/31/2002"/> <input type="button" value="Calendar"/>

<i>Field or Control</i>	<i>Description</i>
Start Date	Specify the start date of the date range.
End Date	Specify the end date of the date range.
Number of Partitions	Enter the number of partitions you want the search index build process to create.
Apply	When you select the Apply button, Partition Manager creates partitions based on the specified number. You can then modify the date ranges. Note: Save the settings before you select the Run button.

After you use the partitioning option and run the indexing process, the Build Search Index page displays the following information.

Note: Do not use Date Range partitioning on any index that uses preprocessing unless the preprocessing has been modified to also use the date ranges. If you use Date Range partitioning on an index that uses preprocessing, it may cause unintended results.

This example illustrates the fields and controls on the Date Range grid. You can find definitions for the fields and controls later on this page.

Date Range

Start Date End Date Partition Count

Process Instance	Start Date	End Date	Run Status	Message Log	Rerun
1	371 01/01/2000	08/31/2005	Success	Message Log	<input type="button" value="Rerun"/>
2	372 09/01/2005	05/02/2011	Success	Message Log	<input type="button" value="Rerun"/>
3	373 05/03/2011	12/31/2016	Success	Message Log	<input type="button" value="Rerun"/>

Field or Control	Description
Process Instance	Search index build process creates multiple PTSF_GENFEED Application Engine programs for the specified number of partitions. Displays the process instance for each partition.
Start Date	Displays the start date (within the specified date range) for each partition.
End Date	Displays the end date (within the specified date range) for each partition.
Run Status	Displays the status of each process instance.
Message Log	If the run status displays error status, use the Message Log link to view the error message.
Rerun	Use the Rerun button to run the partitioned process again.

Using Direct Transfer

Direct Transfer is a new PeopleSoft technology used to transmit search indexing data directly from PeopleSoft batch servers to Elasticsearch. This technology sends the data directly without having the overhead of passing the data through the Integration Broker Gateway and Web server.

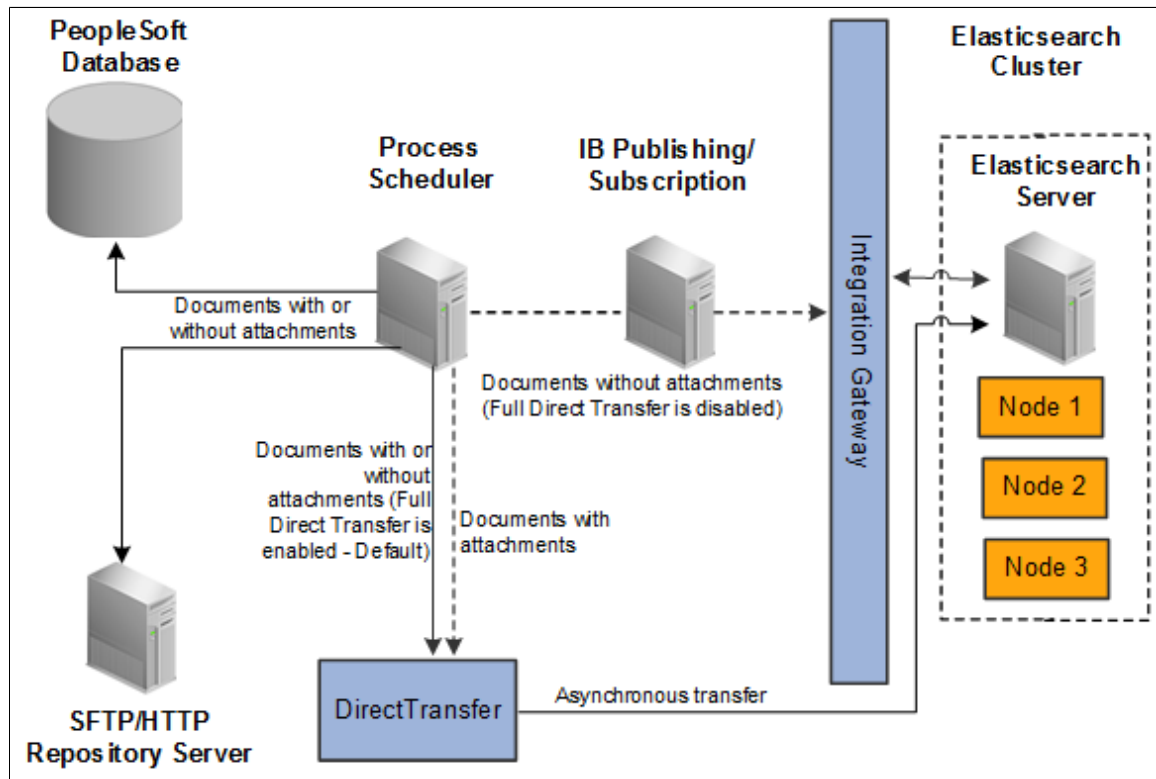
Understanding Direct Transfer

PeopleSoft Search Framework implements Direct Transfer technology to transfer search data (search documents with attachments and search documents without attachments) directly to Elasticsearch bypassing the Integration Gateway. This feature is called Full Direct Transfer, which is the default option in PeopleSoft Search Framework. However, PeopleSoft enables you to choose whether you want to

use the Full Direct Transfer option or not. If you choose not to use the Full Direct Transfer option, then search documents without attachments will be transmitted through the Integration Gateway but, the search documents with attachments will be transmitted directly to the Elasticsearch search engine.

Direct Transfer technology supports search definitions that are based on Connected Query and Query.

The following graphic depicts PeopleSoft downloading search documents from database, attachment repository, and HTTP repository, and then pushing encoded data to Elasticsearch using Direct Transfer. Elasticsearch uses Mapper attachment plug-in to parse attachment contents for indexing.



To use Direct Transfer, you need to:

- Complete the configurations required for Direct Transfer.
- Review the system considerations.
- Review error handling.

Configuration

Direct Transfer requires the following configurations:

- Specify the number of indexing handlers on the Search Options page.
- Specify the maximum attachment error count on the Search Options page.
- Specify whether you want to use the Full Direct Transfer option.

Specifying Number of Indexing Handlers

You use the Search Options page to enter the number of indexing handlers.

The number of indexing handlers determine how many parallel data transfers are done from PeopleSoft to Elasticsearch. The default value is 20, which means that during indexing a maximum of 20 handlers are created. This number should be less than the bulk thread queue size on Elasticsearch (default value of bulk thread queue size is 50).

For search definitions where an average attachment size is 100 KB and the bulk thread queue size on Elasticsearch is 50 (default value), then setting the value for Indexing Handlers to 20 would be optimal in most cases, subject to system considerations.

For indexing search definitions containing large attachments (for example, 10 MB or greater than 10 MB), then the value of Indexing Handlers can be reduced to 10.

The optimal value of Indexing Handlers is dependent on the data volume, pattern and system considerations on PeopleSoft and Elasticsearch.

Note: The Number of Indexing Handlers option is used for search documents with attachments and also for search documents without attachments (in the case when Full Direct Transfer is enabled). When Full Direct Transfer is not enabled, this property is applicable only for search documents with attachments.

Specifying Max Attachment Error Count

You use the Search Options page to enter a value for Max Attachment Error Count.

Max Attachment Error Count is used to determine the error tolerance of the indexing program. It is used to specify the maximum error transactions permitted during indexing. If during indexing, the number of errors exceeds the specified value, the indexing process will exit after completing the process of sending data that is already available in memory.

See [Error Handling](#).

Specifying Full Direct Transfer

You use the Search Options page to select Full Direct Transfer. The default value is *Y*.

If you set Full Direct Transfer to *Y* (yes), then search documents with or without attachments are directly transmitted to the Elasticsearch search engine.

If you set Full Direct Transfer to *N* (no), then search documents without attachments are transmitted through the Integration Gateway while the search documents with attachments are directly transmitted to the Elasticsearch search engine.

Related Links

[Managing General Search Options](#)

System Considerations

Review the following system considerations for Direct Transfer:

- Memory usage on PeopleSoft

- Memory usage on Elasticsearch

Memory Usage on PeopleSoft

The amount of memory used by Direct Transfer for runtime data storage is dependent on the segment size and number of attachment handlers. For example, for the default segment size of 10MB and 20 attachment handlers, Direct Transfer would utilize an additional 200 MB on the PeopleSoft Batch Server during data send.

Memory Usage on Elasticsearch

When Direct Transfer sends data to Elasticsearch, review the main memory considerations on Elasticsearch. This memory is allocated on the Elasticsearch JVM, and should be configured using the `ES_HEAP_SIZE` environment variable, which you set during the installation of Elasticsearch.

- The incoming data is stored in bulk thread queues on Elasticsearch during the ingestion process. The amount of memory used for this purpose is based on the bulk thread queue size.
- The amount of parallel ingestions which can happen on Elasticsearch is based on the cores on the system. The number of parallel bulk ingestions is equal to the number of cores.
- During ingestion, for documents containing attachments, Elasticsearch uses a large amount of memory for document parsing. An example is parse a large attachment of 10 MB size, the memory required for parsing it is around 100 MB.

Examples

1. To index a search definition where the average attachment size is 100 KB and Elasticsearch server is a 2 core system and memory available for Elasticsearch JVM is greater than or equal to 8 GB, then an Attachment Handler value of 20 should suit in most circumstances.
2. To index a search definition where average attachment size is 1 MB and Elasticsearch server is a 4 core system and memory available for Elasticsearch JVM is greater than or equal to 16 GB, then an Attachment Handler value of 10 should suit in most circumstances.
3. To index a search definition where average attachment size is 100 MB and Elasticsearch server is a 4 core system and memory available for Elasticsearch JVM is greater than or equal to 16 GB, then an Attachment Handler value of 5 should suit in most circumstances. The `http.max_content_length` setting on the `elasticsearch.yml` configuration file should also be increased in this scenario. The default value is 100 MB, but for large attachments you may set the value to a higher value, for example, `http.max_content_length=512mb`

Error Handling

To log error messages and to handle errors in Direct Transfer, PeopleSoft Search Framework provides the following:

- On the Search Options page, Search Framework provides the Enable Attachment Trace option to enable the detailed logging of each task on the trace file, which logs messages for search documents with attachments and search without attachments. For Direct Transfer, you should set this property to `Y` if you want to troubleshoot any errors associated with the transfer of search data from PeopleSoft to Elasticsearch.

For more information, see [Managing General Search Options](#).

- On the Search Options page, the Max Attachment Error Count option is used to determine the error tolerance of the indexing program. For more information, see [Configuration](#).
- On the Build Search Index page, in the Previous schedule details section, an Error link is displayed if an error occurs while indexing (PeopleTools, Search Framework, Search Admin Activity Guide, under Administration, Schedule Search Index).
 - Full Direct Transfer is enabled — During indexing, there could be possibilities that certain data transactions may not be successful either due to temporary system situations (for example, Elasticsearch is performing some internal maintenance operations) or due to data issues. In such cases, Direct Transfer stores the request and response of the transaction that goes into error in an error table, and proceeds with further indexing.

After the indexing process is complete, errors, if any, can be viewed on the Build Index page (Previous schedule details section).


- Full Direct Transfer is disabled — During the indexing of search documents with attachments (which are transmitted directly to Elasticsearch), if any error occurs, an Error link is displayed on the Build Search Index page.

During the indexing of search documents without attachments (which are transmitted through the Integration Gateway), if any error occurs in IB transactions, an Error link may not be displayed on the Build Search Index page. You need to check the Process Monitor whether the PTSF_GENFEED Application Engine program displays **No Success** status. If the status of the PTSF_GENFEED Application Engine program is **No Success**, administrators must review the IB monitor logs for details on the error.

- Based on the type of error, an administrator can use the **Rerun** option to resubmit the transactions if the issue was related to temporary environment problems, or can review the data and correct the data at the source, or can determine whether the data that is causing the error can be omitted for indexing. If the data needs to be corrected at source, then the indexing process should be executed again to get the correct data indexed.

Click the Error link to display the Attachment Transfer Exception Details page.

This example illustrates the fields and controls on the Attachment Transfer Exception Details page. You can find definitions for the fields and controls later on this page.



Field or Control	Description
Process Instance	The process instance ID of the PTSF_GENFEED Application Engine program that is causing the exception.
Request	The request sent by the PTSF_GENFEED Application Engine program to Elasticsearch.
Response	The response received by the PTSF_GENFEED Application Engine program from Elasticsearch.
Rerun	<p>Use the Rerun button to resubmit the selected error transactions to Elasticsearch. The progress of the transaction can be monitored using the Process Monitor page on PIA. When a transaction is successful, it is deleted from the error pages.</p> <p>In the case where the maximum exceptions limit is not reached, when you select the Rerun button:</p> <ul style="list-style-type: none"> • The transaction is directly submitted to Elasticsearch through the PTSF_GENFEED Application Engine program, but without executing the Connected Query program. • The request messages stored in error tables will be sent directly to Elasticsearch. <p>In the case where the maximum exceptions limit is reached, the Rerun button acts as a re-submit of the PTSF_GENFEED Application Engine program.</p> <p>After re-running the PTSF_GENFEED Application Engine program, if the transaction is not successful, the error tables are updated with the latest error message.</p>

Using Direct Transfer with SSL Enabled Elasticsearch

For Direct Transfer to work with SSL enabled Elasticsearch, the Certification Authority (CA) certificate should be available in the Digital Certificate repository of PeopleSoft. This is to ensure that when the Direct Transfer process is run, the root certificates are obtained from the repository to create a CA certificate bundle and use it for secured communication between the Process Scheduler server and Elasticsearch.

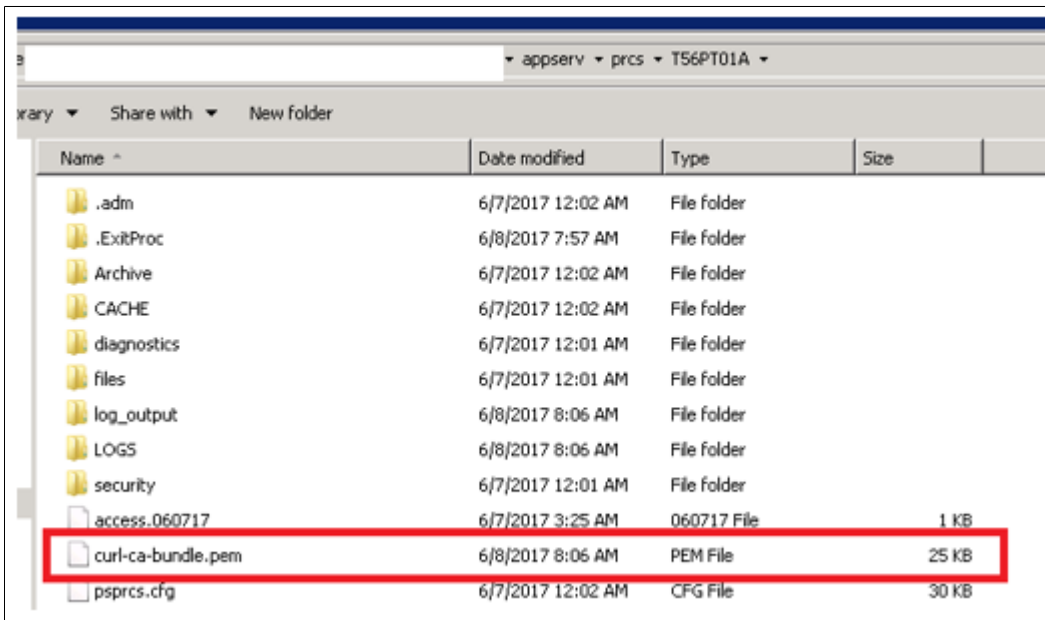
To make the CA certificate available in the Digital Certificate repository of PeopleSoft, add the CA certificate of Elasticsearch to **PeopleTools >Security >Security Objects >Manage Digital Certificates**.

This example illustrates the Digital Certificates page.

Type	*Alias	*Issuer Alias	Valid to	Links		
Root CA	GTE CyberTrust Global Root	GTE CyberTrust Global Root		Detail	+	-
Root CA	GTE CyberTrust Root	GTE CyberTrust Root		Detail	+	-
Root CA	KeyWitness Root	KeyWitness Root		Detail	+	-
Root CA	Novell RNO207225	Novell RNO207225	06/24/21 1:58:37PM	Detail	+	-
Root CA	Novell eDir	Novell eDir	04/12/19 1:23:29PM	Detail	+	-
Root CA	PeopleTools	PeopleTools		Detail	+	-

When the process starts, the certificate bundle in PEM format is available in the Process Scheduler server folder.

This example illustrates the certificate bundle in PEM format.



Real-time Indexing Using Direct Transfer Technology

Search Framework uses Direct Transfer technology during the indexing process. Search Framework enables applications to perform real-time indexing using the delivered methods of Direct Transfer.

See "PTDirectTransferObject Class Methods" (PeopleCode API Reference).

A sample code that performs real-time indexing using Direct Transfer:

```
Local JsonBuilder &jb = CreateJsonBuilder();
&jb.StartObject("");
&jb.AddProperty("FIRST_NAME", "MyFirstName");
&jb.AddProperty("LAST_NAME", "MyLastName");
&jb.EndObject("");

Local string &pdoc = &jb.ToString();
```

```
Local PTDirectTransferObj &PTDTobj = CreateDirecttransferObject();
&PTDTobj.SetIndex("myprofile");
REM &PTDTobj.SetIndex("vchr_chg_history_e92cmhk2");
&PTDTobj.SetDocId("1");
&PTDTobj.SetMethod("POST");
&PTDTobj.SetBody(&pdoc);

Local string &response = &PTDTobj.SyncTransfer();
```

The response from SyncTransfer can be parsed and acted upon accordingly.

Verifying Domain Status

To verify and set domain status, access the Domain Status page by selecting **PeopleTools >Search Framework >Search Admin Activity Guide >Administration >Domain Status**.

For more information, see "Working with the Domain Status Page" (Integration Broker Administration).

Verifying Integration Broker Queue Status

To verify the status of the PTSF_ES_SEND_Q queue, access the Asynchronous Services page by selecting **PeopleTools >Search Framework >Search Admin Activity Guide >Administration >Asynchronous Services**.

See [Handling Common Errors](#) for a description of the queue status.

Managing Search Context

This section provides an overview of managing search context.

Understanding Search Contexts

Use search contexts to define which search groups appear in the search group drop-down Global Search Bar, depending on the context of the user. The context determines what appears in the drop-down relative to where the user is located.

A search group is a search category that you enable for the purpose of setting up search contexts. Currently, search groups can belong to the Home Page search context.

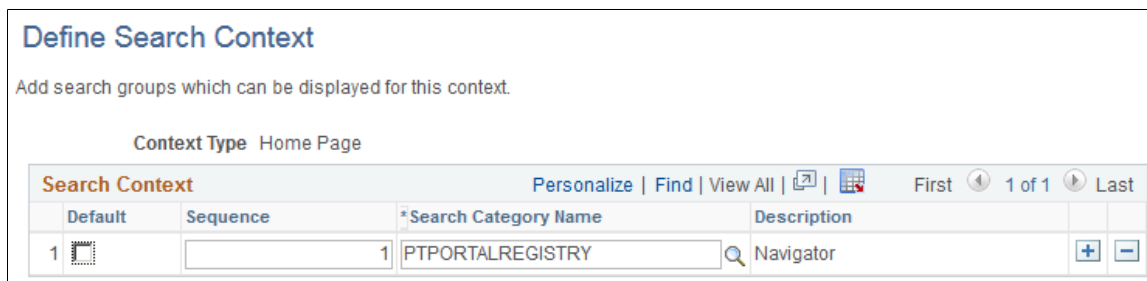
Note: From PeopleTools 8.55, only the Home Page search context is available. That is, the Portal Node and Work Center search contexts are *not* available.

For the Home Page search context, a default search group can be selected. It is not required to have a default search group for the search context.

Defining Search Contexts

Access the Define Search Context page by selecting PeopleTools, Search Framework, Search Admin Activity Guide, Administration, Define Search Context and select the Home Page search context.

This example illustrates the fields and controls on the Define Search Context page. You can find definitions for the fields and controls later on this page.



Use the plus and minus buttons to add and remove the search categories to and from the search context.

Field or Control	Description
Default	<p>Select the category that should be the default.</p> <p>Only PTPORTALREGISTRY related (local and remote search groups) search categories can be chosen as default. All other search categories are disabled.</p> <p>Only one default search category can be selected for a search context.</p> <hr/> <p>Note: If no categories are marked as Default, then Menu is the default. If Menu is not available to be chosen, then All is the default.</p> <hr/> <p>Note: Global Search allows searching by All allowable search categories. 'All' is an option that is automatically added to Global Search if in use. This means that users can search across all groups to which they have access. The 'All' category may <i>not</i> be removed.</p>
Sequence	<p>Specify a numerical sequence to define the order in which the search categories will appear in the Global Search Bar drop-down list.</p>
Search Category Name	<p>Select the search category to include for this search context.</p>

Viewing Search Contexts

To view your defined search contexts, access the View Search Contexts page by selecting PeopleTools, Search Framework, Search Admin Activity Guide, Administration, View Search Contexts.

This example illustrates the fields and controls on the View Search Context page. You can find definitions for the fields and controls later on this page.

View Search Context

List of all the search contexts defined for this system

Filter Contexts

Context Type

Node Name\Workcenter ID

Search Contexts

Context Type Home Page [Edit Context](#)

Default	Sequence	Search Category Name	Description
<input type="checkbox"/>	1	PTPORTALREGISTRY	Navigator

Filter Contexts

Use the Filter Context controls to customize and narrow the contexts displayed.

Field or Control	Description
Context Type	Select the context type you wish to view. Home Page is the only available context type.
Node Name\Workcenter ID	Enter a specific node name.
Clear	Click to clear the previous filter criteria.

Search Contexts

The Search Contexts section displays the search contexts that meet the criteria specified in the Filter Contexts section.

Field or Control	Description
Edit Context	Click to modify the search context definition. The link opens the Define Search Context page with the selected search context open.

Viewing Node in Network

Use the Node Network page to view and configure details about PeopleSoft nodes in the integration network.

For more information, see "Viewing Node Network Status" (Integration Broker Administration).

Importing Remote Search Groups

You can share search groups between databases to enable seamless navigation between systems. However, the nodes and the search groups need to be known to all systems involved. The imported search groups are treated as local search groups, and need to be modified as such. For example, imported groups need to be included in the search contexts to become available.

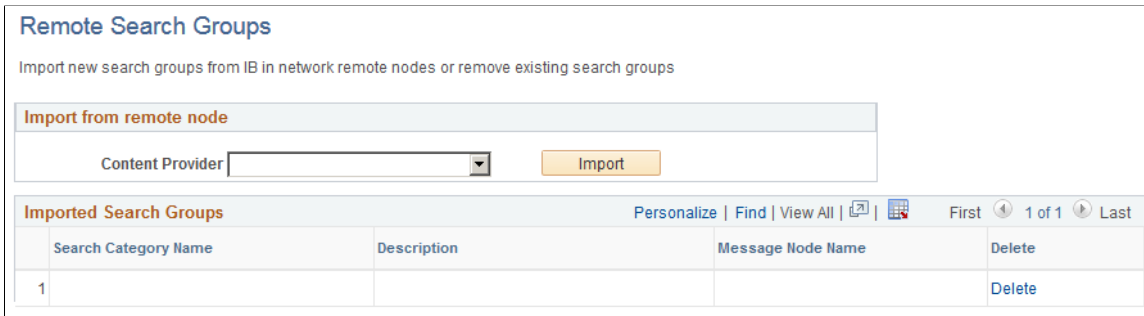
When you import search groups, you import them from a remote node, as defined in your Integration Broker configuration. When importing search groups, keep the following items in mind:

- Both the local and remote systems need to point to the same search instance. That is, both systems need to have the same search instance defined on the Search Instance Properties page.
- Both systems need to have unique Default Local Node names.
- Both systems need to have a unique database name.
- The remote node needs to be in the IB Network of the local system.
- The remote system needs to trust the local system. That is, single sign-in must be set up between the two systems, and they must share common PeopleSoft user profiles. If you need users to be able to search both indexes, the users must exist and be authorized in both systems.
- The remote node Portal URI Text and Content URI Text values need to be specified on the Portal tab of the node definition.

To import search groups from other systems, access the Remote Search Groups page by selecting PeopleTools, Search Framework, Search Admin Activity Guide, Administration, Remote Search Groups.

Note: Importing remote search groups retrieves all the search groups defined in that content system along with all related search categories, associated search definitions, and the list of faceted attributes. Be careful not to import all search groups unless you are actually using them. Turning on all search groups degrades performance.

This example illustrates the fields and controls on the Remote Search Groups page. You can find definitions for the fields and controls later on this page.



Field or Control	Description
Content Provider	Select the remote node defined in your system from which you want to import search groups. Content providers must be defined as nodes known to your system using Integration Broker.
Import	Click to invoke the import process. This will import all search groups from the selected node.
Delete	Click to remove any unnecessary or obsolete search groups from your system.

When working with remote search groups, consider these items:

- All the imported search groups need to be registered in the appropriate search context to appear in the Global Search Bar.
- When a user signs on to the PeopleSoft application, the host system sends requests to get the remote search category details. The category details received from the remote system are cached to avoid repeated calls to the remote system. If a content system goes down or if there have been any changes in the content system, it will be reflected in the host system only after a user signs on for a new session, not during the current session.
- For the remote search categories, the security of what needs to be displayed is controlled by the content system. If the user does not have access to the groups on the content system, the user will not be able access them in the Global Search Bar in the host system. If the content system is down, or the user does not exist in the content system, or the user does not have access to category details, the system logs an error into application server log and does not display the categories to the user.

Related Links

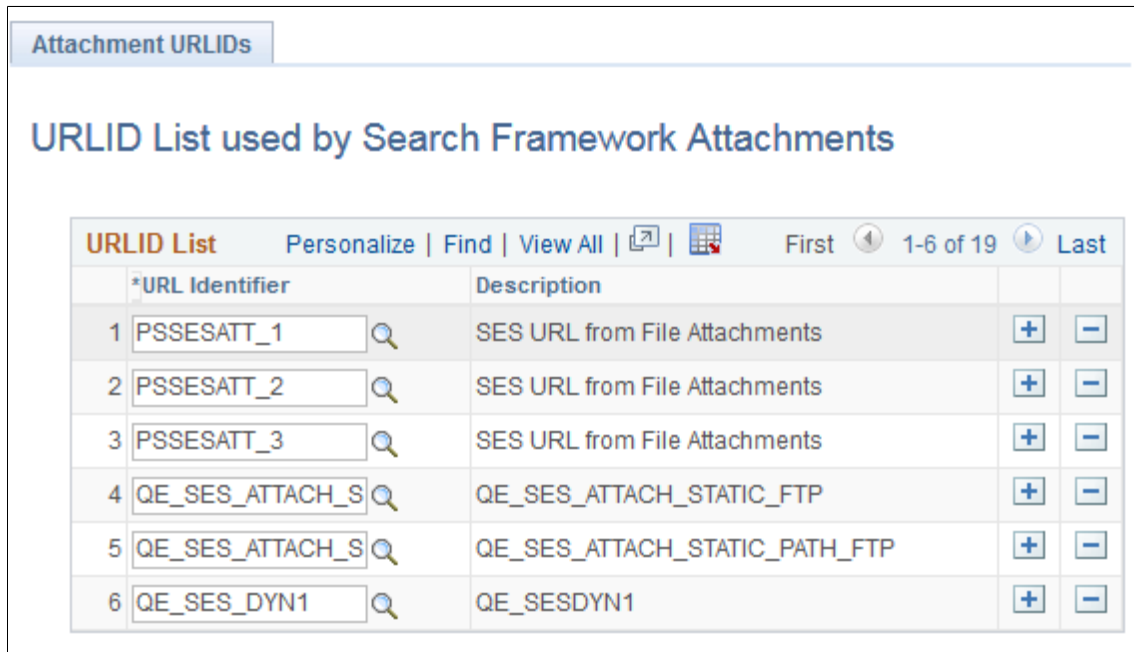
- "Understanding Nodes" (Integration Broker Administration)
- "Setting Portal Node Characteristics" (Portal Technology)

Creating the Attachment URL ID List

Access the Attachment URL ID List page by selecting PeopleTools, Search Framework, Search Admin Activity Guide, Administration, Attachment URL ID List.

Use the Attachment URL ID List page to define a list of URL IDs to be used specifically for attachments within the PeopleSoft Search Framework. URL definitions are created using the URL Maintenance page (PeopleTools, Utilities, Administration, Maintain URLs), and using the Attachment URL ID List page you identify a subset of these URL definitions to be recognized by the PeopleSoft Search Framework.

This example illustrates the fields and controls on the Attachment URL IDs page. You can find definitions for the fields and controls later on this page.



Field or Control	Description
URL ID List	This grid contains the subset of previously defined URL definitions that you want the PeopleSoft Search Framework to recognize as attachment storage locations.
URL Identifier	From the drop-down list, select the URL Identifier to add to the URL ID list. The drop-down list prompts against the list of URLs defined on the URL Maintenance page.

Note: If you need to use multiple user credentials for a single FTP server, create separate URLID's for each user. The URL definition allows only one user credential for each URLID. Password credentials must be encrypted on the URL definition.

Related Links

"URL Maintenance" (System and Server Administration)

Implementing Report Repository Search

The Search Framework enables you to index the contents (files and generated reports) of your report repository, enabling end users to find reports quickly using free text search.

PeopleTools delivers the Reports (PTSEARCHREPORTS) search definition and search category. To index your report repository, deploy PTSEARCHREPORTS search definition and search category and run the schedule index process (PTSF_GENFEED) for that search definition. The underlying connected query for the search definition gathers the required data from the Process Scheduler tables for report indexing and access.

Note: PeopleSoft delivers the PTSEARCHREPORTS index with all report types undefined. To define which reports to be included in the Report Repository Search, navigate to PeopleTools, Process Scheduler, Process Scheduler Settings and choose the Search Configuration tab.

End users can search the Reports index from:

- Global Search
- Component Search page
- Search Reports page (Reporting Tools > Search Reports)

The following example provides a sample search using the Search Reports page. Descriptions of the controls follow the example.

The screenshot shows the 'Search Reports' interface. At the top, there is a search input field containing the text 'reports' and a search icon. Below the search bar, the results are displayed in a table-like format. On the left side, there are filter options under the heading 'Filter by'. On the right side, there are search results under the heading 'Search Results'. The search results include a link to 'abcd', along with its last updated date and file name. The filter options include 'Created By' (VP1 (3)), 'Process Type' (Application Engine (1), SQR Report (1), nVision-Report (1)), 'Folder Name' (GENERAL (3)), and 'Created On' (2014 (3)).

Filter by	Search Results
Created By VP1 (3)	abcd Last Updated Date:2014-05-23 00:57:51 File Name: XRFMENU-10545.pdf
Process Type Application Engine (1) SQR Report (1) nVision-Report (1)	Detailed Financial Statement Last Updated Date:2014-05-22 10:50:20 File Name: psnvs_5908.nvt
Folder Name GENERAL (3)	Cross Reference Window Listing Last Updated Date:2014-05-20 04:55:41 File Name: XRFWIN_10528.PDF
Created On 2014 (3)	

The default facets by which users can filter the reports are:

- Created By: Filters results by the user ID submitting the report.
- Process Type: Filters results by process type (SQR, Application Engine, and so on).
- Folder Name: Filters results by the folder name where the report resides in the report repository.

- Created On: Filters results by date.

Related Links

[Administering Search Definitions and Search Categories](#)

[Working with Search Indexes](#)

"Using Search Reports" (Process Scheduler)

Working with Search Framework Definitions During Upgrades

These Search Framework definitions are managed objects, meaning they can be upgraded, added to projects, compared, copied, and so on:

- search attributes
- search definitions
- search categories
- search contexts

Boosting Score of Search Results

When end users search for content, the listing of search results is based on the score of a match. A match with the highest score is displayed at the top of the search-results list. PeopleSoft Search Framework allows application designers to programmatically boost the score of a result so that the most relevant results are displayed.

Note: Score boosting cannot be configured on Global and Component searches; only custom searches can be configured with score boosting.

PeopleSoft Search Framework uses the SearchFilterGenerator API to generate custom boosting queries where you provide the boost factor for a filter. The search engine supports score boosting on an individual filter level or on a group of filters and it supports boosting on any filters supported by the SearchFilterGenerator API.

Note: PeopleSoft Search Framework restricts the use of multiple score functions such as ADD, MIN, MAX, etc. within a single query. Only one operation is allowed in a single request.

The following code sample generates a custom score based on a filter on product for values of television, TV, monitor, or display with a boost factor.

```
If &fg.MatchAnyReturnTrue() Then
    &fg.ContainsWord("PRODUCT" , "Television");
    &fg.BoostFilterScore(0.03);
    &fg.ContainsWord("PRODUCT" , "TV");
    &fg.BoostFilterScore(0.03);
    &fg.ContainsWord("PRODUCT" , " Monitor");
    &fg.BoostFilterScore(0.02);
    &fg.ContainsWord("PRODUCT" , "Display");
    &fg.EndMatchAny();
    &fg.setBoostModeSum();
```

```
End-If;
```

For code examples of custom score boosting, see "Search Filter Generator Examples" (PeopleCode API Reference).

Note: Faceted navigation is supported in Elasticsearch.

Sample Custom Queries for Score Boosting

Applying a single filter:

```
If &fg.MatchAllReturnTrue() Then
    &fg.ContainsWord("DEPTNAME", "support");
    &fg.BoostFilterScore(0.02);
    &fg.EndMatchAll();
End-If;
```

Joining filters with logical operators:

```
If &fg.MatchAllReturnTrue() Then
    &fg.ContainsWord("DEPTNAME", "support");
    &fg.BoostFilterScore(0.02);
    &fg.ContainsWord("QE_STATUS", "open");
    &fg.BoostFilterScore(0.03);
    &fg.EndMatchAll();
End-If;
```

Applying individual scores in a function such as ADD, AVG, etc. to get overall score on the document based on filter matches:

```
If &fg.MatchAnyReturnTrue() Then
    &fg.ContainsWord("DEPTNAME", "support");
    &fg.BoostFilterScore(0.02);
    &fg.ContainsWord("QE_STATUS", "open");
    &fg.BoostFilterScore(0.03);
    &fg.EndMatchAny();
    &fg.setBoostModeSum();
End-If;
```

Using wild cards (use * and ? for multi character and single character match respectively):

```
If &fg.MatchAllReturnTrue() Then
    &fg.ContainsWord("MESSAGE_TEXT", "l*g");
    &fg.BoostFilterScore(0.01);
    &fg.ContainsWord("MESSAGE_TEXT", "b??ght");
    &fg.BoostFilterScore(0.02);
    &fg.EndMatchAll();
End-If;
```

Using stemming (add '\$' in front of the word to consider it for stemming):

```
If &fg.MatchAllReturnTrue() Then
    &fg.ContainsWord("MESSAGE_TEXT", "$wanted");
    &fg.BoostFilterScore(0.03);
    &fg.EndMatchAll();
End-If;
```

Using score boosting without specifying a field:

```
If &fg.MatchAnyReturnTrue() Then
    &fg.ContainsWord("", "long");
    &fg.BoostFilterScore(0.02);
    &fg.EndMatchAny();
End-If;
```

Using Fuzzy search (use ~ at the end of the word to treat it as fuzzy search request):

```
If &fg.MatchAnyReturnTrue() Then
    &fg.ContainsWord("MESSAGE_TEXT", "speaks~");
    &fg.BoostFilterScore(0.02);
    &fg.EndMatchAny();
End-If;
```

Related Links

"SearchFilterGenerator Class Methods" (PeopleCode API Reference)

Testing Score Boosting

The Search Test Page allows you to check the score boosting feature.

See [Using the Search Test Page](#).

Administering Real Time Indexing

Understanding Real Time Indexing

PeopleSoft Search Framework uses data that is indexed in Elasticsearch for search and analytics, so it is recommended that the data is indexed in real time to provide current search and analytics. This would eliminate stale data in Elasticsearch indexes and ensure that discrepancies do not exist between the data in the PeopleSoft database and the indexed data in Elasticsearch.

Any search definition with the source type of query or connected query can be configured for real time indexing. Refer to the PeopleSoft Application Fundamentals documentation *<for your product line>* for a list of delivered search definitions for which real time indexing can be configured.

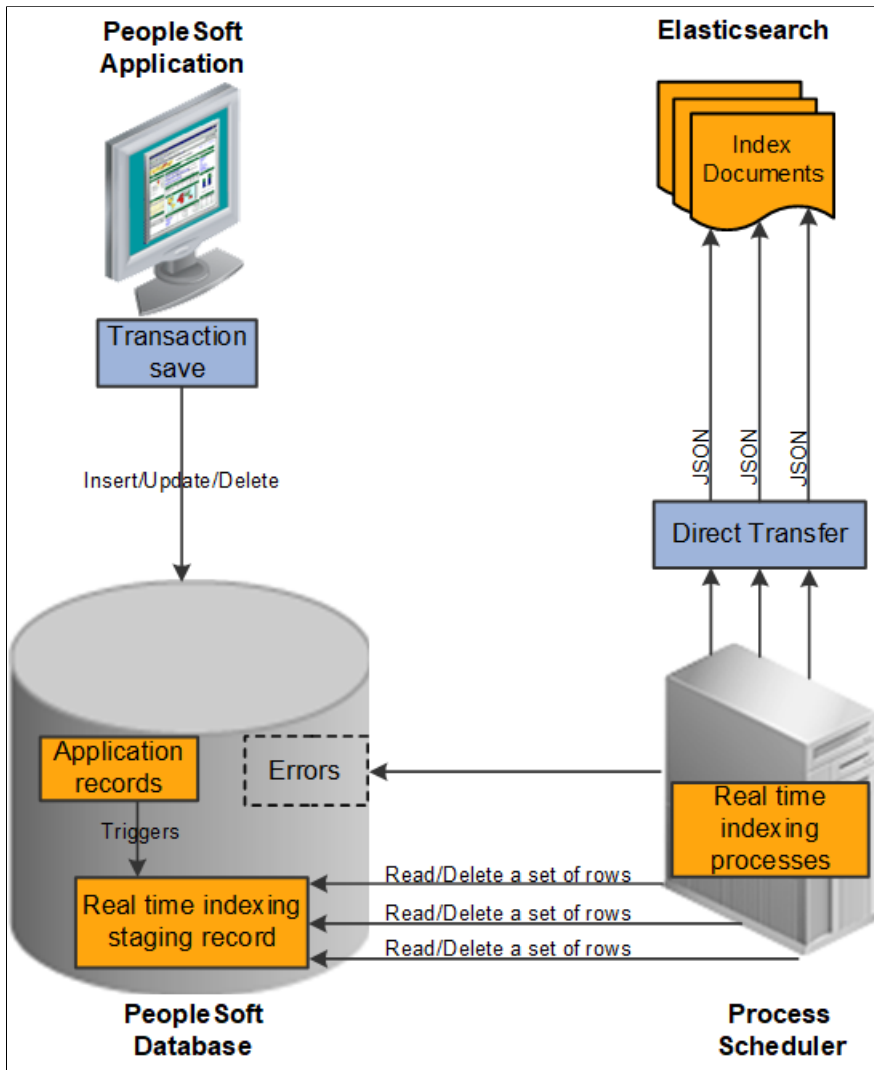
Understanding Set Based Processing in Real Time Indexing

Set based processing refers to processing groups, or sets of rows, at one time rather than processing each row individually.

The real time indexing process ensures that data is updated on the search server as soon as an application transaction is saved. Based on the volume of data under processing in the real time indexing queue, the transaction update may appear to be real time or near real time.

Search Framework implements set based processing in real time indexing as shown in the following diagram.

The following diagram illustrates the process flow of the real time indexing process.



Real time indexing uses database triggers as the initializing point for the communication to the search server. As the data is inserted, updated, or deleted, the database trigger associated with the application record inserts a row in the real time indexing staging table. The staging table acts as an interim data holder by storing the keys of the transaction.

Note: When batch processing is enabled for a search definition, updated data is not pushed to the real time indexing staging table for the search definition. For more information, see [Turning Off Real Time Indexing During Batch Processing](#).

A dedicated Process Scheduler process for real time indexing polls the staging table at regular intervals for any data to process. You can configure the number of processes in the Process Scheduler configuration properties file. By default, only one process is initiated. You may scale up the number of processes by changing the maximum number of instances. However, the current design is limited to enable real time indexing on one domain only. Enabling real time indexing on more than one domain causes the same dataset to be picked by both processes resulting in duplicate processing.

Real time indexing processes the transactions stored in the staging table as a set. The size of the set is a configurable parameter in the Search Options page, which can be specified as per the available resources

in the Process Scheduler server. A set can contain just one transaction or it can contain the maximum number of transactions as specified in the Real Time Indexing Set Size property based on the available data at any point in time.

Real time indexing again processes each set for a specific search definition and begins the data retrieval process for each search definition. After the data is collected and formatted to a JSON structure, it is transferred to the search server using Direct Transfer.

Note: When an application batch program updates a large volume of data in an insert or update action, you can expect delay in updating the data to the search server due to the relatively large volume of data as compared to the online transactions with low volume. Therefore, PeopleSoft recommends that you use the real time indexing batch switch for search definitions with batch programs having large volume indexing.

In most of the search definitions enabled for real time indexing, you may not need to schedule incremental indexing. However, indexes that have data with effective date changes may require the incremental indexing schedules for the actual data synchronization to take place. A change on the effective dated rows does not initiate a trigger, therefore such search definitions should be indexed using the existing indexing methods. While real time indexing ensures data synchronization based on transaction updates, incremental indexing updates the indexed data (from prior indexing schedules) to the current date.

Note: You should periodically run the AE maintenance program (PTRTI_TRUNC) of real time indexing to ensure that the real time indexing staging tables are de-fragmented after running large volume batch update for better performance. For a description of the AE maintenance program, see [Using an Application Engine Maintenance Program](#).

Configuring Real Time Indexing

Any search definition with the source type of query or connected query can be configured for real time indexing.

Search Framework supports real time indexing through SQL triggers and uses set based processing.

Important! Oracle does *not* deliver search definitions enabled for real time indexing automatically. Therefore, you should enable search definitions for real time indexing based on your business requirements. In addition, some search definitions need to be certified before being turned on. See the Application Fundamentals Online Help <for your product> for more information.

Oracle delivers two search definitions, PTSF_RECENTMENU and PTSF_RECENT_KW, for providing suggestions in search that are configured for real time indexing, but you have to enable real time indexing for these two search definitions.

See [Working with Suggestions Search](#).

Prerequisites for Real Time Indexing

Before you can begin configuring real time indexing for a search definition, you should configure the following elements or verify that these elements are set correctly:

- Set the Real Time Indexing property in Quick Configure Menu.

- Configure PSRTISRV server process.
- Create an application class for any preprocessing or post processing.
Application class may be provided by PeopleSoft for any delivered indexes.

Note: Real time indexing can be enabled on search definitions with trigger records up to 16 key fields.

Setting the Real Time Indexing Parameter in PSADMIN Quick Configure Menu

You should set the Real Time Indexing parameter on the Process Scheduler domain.

To set the Real Time Indexing parameter:

1. Run PSADMIN and choose the Process Scheduler domain.
2. Select 4) Configure this domain.
To configure the domain, the domain needs to be shut down.
3. Enter y to the question: Do you want to continue (y/n).
4. Toggle 6) Real Time Indexing to set it to Yes.
5. Select 7) Load config as shown.
6. Restart the domain.

Configuring the PSRTISRV Server Process

The real time indexing server process, PSRTISRV, starts when the Process Scheduler domain is booted up. The number of RTI server processes can be configured in the Process Scheduler section of the configuration file (psprcs.cfg). The default is set to 1 instance.

See "PSRTISRV Section" (Process Scheduler)

You should configure the following RTI server process parameters on the Search Options page (select **PeopleTools >Search Framework >Search Admin Activity Guide**, from the left panel choose Administration, and then choose Search Options):

- Index Refresh Interval
- Heartbeat Interval
- Real Time Indexing Set Size
- Real Time Indexing Threshold

For description of these parameters, see [Managing General Search Options](#).

Creating an Application Class

For full or incremental indexing of a search definition, any preprocessing is performed by an Application Engine program. In the real time indexing process, you should define an application class to perform preprocessing and post processing. The application class should extend the PTRTIPKG:PTRTI application class. The PTRTI class contains the Process and PostProcess methods, and you should

extend these methods with your application class to perform preprocessing and post processing tasks respectively.

See "PTRTI Class" (PeopleCode API Reference).

Steps to Enable Real Time Indexing for a Search Definition

To ensure that the real time indexing process runs smoothly, PeopleSoft suggests the following sequence of steps.

1. On the Configure Real Time Indexing page,
 - a. Select the Enable Real Time Indexing check box.
 - b. Verify that the metadata is populated correctly and ensure that you specify the required trigger records.
 - c. Select the Reload link if a search definition is updated after configuring it for real time indexing. You should use the Reload link to import the updates to real time indexing.

Note: The Reload link will appear on the page only if a configuration is already saved and stored in the database.

- d. Save the search definition.

Restart the Process Scheduler domain if you performed the Reload action.
 - e. Create the trigger.
2. Run incremental index for the search index for data synchronization. Perform this step to index records that were added before real time indexing was initialized.

The selected search definition is now configured for real time indexing. When data is inserted, modified, or deleted in the database for the search index, a SQL trigger runs to initiate real time indexing of data in Elasticsearch.

Note: Search definitions with queries using set operators like UNION, INTERSECTION etc. are not supported for real time indexing. Real time indexing requires equivalent queries without set operators.

Using the Configure Real Time Indexing Page

Use the Configure Real Time Indexing page to enable real time indexing for a specific search definition and to specify any additional information required for real time indexing of the search definition.

Access the Configure Real Time Indexing page by selecting **PeopleTools >Search Framework >Real Time Indexing >Configure Real Time Indexing**.

The search page displays search definitions with the source type as query and connected query from the default search instance. The search page also provides a comprehensive view of the deploy status retrieved from the Deploy Search Definition page, and the real time indexing configuration status of the search definitions, enabling you to take the necessary action on a search definition.

Select a search definition to display the real time indexing information. The search definition metadata is imported including security and attachment configurations.

This example illustrates the fields and controls on the Configure Real Time Indexing page. You can find definitions for the fields and controls later on this page.

Field or Control	Description
Enable Real Time Indexing	Select the check box to enable real time indexing for the search definition.

<i>Field or Control</i>	<i>Description</i>
Configure Batch	<p>Select the check box to configure batch programs, with a specific Process Type and Process Name, from an available list of batch programs.</p> <p>See Turning Off Real Time Indexing During Batch Processing.</p>
Trigger SQL	<p>After reviewing or modifying the search definition, you should save the search index. The Trigger SQL link appears only when you select Save on the page. Then, you should create a trigger. A trigger record is generated at runtime. Contact your database administrator to implement triggers on the database.</p> <hr/> <p>Note: PeopleSoft RTI trigger names have the prefix SFA and suffix TR. For example, SFA_PT_PORTAL_MRU_TR.</p> <hr/> <p>Any of the following ways can be used to create a trigger:</p> <ul style="list-style-type: none"> • Select the Trigger SQL link to download the .sql files for the supported database platforms. Run the downloaded trigger SQL in a SQL tool to create triggers. • Open the trigger record in Application Designer, and build the record using the Create Trigger option. • Run the Datamover command: CREATE_TRIGGER <recordname>, where recordname is the trigger record. <p>For more information on trigger records, see the section on Trigger Records later in this topic.</p>
Reload	<p>The Reload link appears only when you select Save on the page.</p> <p>If a search definition is updated after configuring it for real time indexing, you should use the Reload link to import the updates to real time indexing.</p> <hr/> <p>Note: After completing the reload action, ensure that you save the real time indexing configuration, and then restart the Process Scheduler domain.</p>

Data Source

<i>Field or Control</i>	<i>Description</i>
Query/Connected Query	<p>The system automatically selects Query/Connected Query as data source for those search definitions that are based on query or connected query.</p>

Field or Control	Description
Application Class	<p>The system automatically selects Query/Connected Query and Application Class as data sources for those search definitions that contain preprocessing or post processing Application Engine program.</p> <p>When both Query/Connected Query and Application Class are auto populated as data sources, the real time indexing process runs the application class method first and then runs SQL query. If query uses a temporary table, the application class method should insert the data to the temporary table.</p> <p>You can choose to use only the application class as the data source rather than Query/Connected Query. In this case, the application class methods perform the necessary preprocessing or post processing and also set the values for all search attributes of each search document to index. The real time indexing process then creates the JSON document and uses the Direct Transfer technology to index the data.</p> <p>See "PTRTI Class" (PeopleCode API Reference).</p>
Activity Guide	Reserved for future use.

Data Generator Application Class

The Data Generator Application Class section is displayed when application class is selected as the data source.

Field or Control	Description
Package Name	<p>Specifies the root package where the application class is created.</p> <p>You should ensure that the extended application class exists in the database.</p>
Path	Species the path, represented by a colon (:).
Class ID	Choose the relevant application class.

Field or Control	Description
Process	<p>Set the Process option to Yes if you want to perform any validation or preprocessing and to defer or discard processing.</p> <p>When you set this option to Yes, you should extend the Process method of the delivered application package PTRTIPKG:PTRTI.</p> <p>"PTRTI Class Methods" (PeopleCode API Reference)</p>
Post Process	<p>Set the Post Process option to Yes if you want to perform any validation or post processing.</p> <p>When you set this option to Yes, you should extend the PostProcess method of the delivered application package PTRTIPKG:PTRTI.</p> <p>"PTRTI Class Methods" (PeopleCode API Reference)</p>

General

Search definition metadata is imported and displayed in the Title, Summary, URL, and Last Update Date Time fields.

You cannot edit these fields.

Trigger Records

A trigger record is necessary for real time indexing. Therefore, you must ensure that the Record field contains a trigger record. When a search definition is based on Query or Connected Query, the system checks the query to try and determine the best trigger record and pre-populates the value for you.

The keys of the trigger record (transaction record) should exist in the primary record of the parent query. If trigger records are pre populated, you should retain these trigger records. You should add a trigger record only if no trigger record is given.

Optionally, you can add more trigger records but doing so will affect the performance. Also, keeping triggers on child records may cause undesired data insertions to the real time indexing staging table, which may increase the processing time.

Field or Control	Description
Record	<p>The Record field is pre populated based on the parent record used in Query/Connected Query of the search definition where the Last Modified Date Time is configured. If it is a view, the first record used in that view is auto selected as transaction record to create a trigger. However, you can change it to any other record in that view.</p> <p>The real time indexing process recognizes a transaction only when the trigger record inserts a row for the transaction. Hence, the trigger record should be decided based on this condition. Adding multiple trigger records can cause the same data to be processed multiple times to Elasticsearch.</p>
Temp Table	<p>Select this check box if the search definition's Last Modified Date Time field is using a temporary table.</p> <hr/> <p>Note: Temp Table and Transaction Record or Record fields are displayed for search definitions that have data source as Application Class.</p> <hr/>
Transaction Record	<p>Transaction Record field is used only when the Record field contains a view or a temporary table.</p>

Attachments

If a search definition contains attachments, the Include Attachments check box is displayed.

Field or Control	Description
Include Attachments	<p>Use this check box to either include attachments or exclude attachments from real time indexing.</p> <p>If a search definition contains attachments, it is ideal to include attachments in real time indexing. However, if you observe that attachments are causing delay in real time indexing, you can choose to exclude attachments. In this case, you can use incremental indexing because in incremental indexing attachments are also indexed.</p> <p>If an index does not contain attachments, the Include Attachments check box is not displayed.</p>

Search Attributes

The Search Attributes section displays metadata imported from the search definition, and it cannot be edited here.

Application developers can use this mapping from the search definition page to know what attributes should be loaded in the application class.

Turning Off Real Time Indexing During Batch Processing

While processing high volume batch updates for search definitions, real time indexing can cause a sudden increase in the number of entries to process in the real time indexing staging table. This can, in turn, cause the database table fragmentation to increase and eventually slow down the performance. Therefore, for large volume batch updates, it is recommended that real time indexing be temporarily suspended and changes to the search index be updated through incremental indexing.

Note: When real time indexing is off, the staging table will not receive any data for both online and batch programs.

After the batch programs are completed, incremental indexing for a search definition will take place at the scheduled time. When incremental indexing starts, real time indexing is automatically enabled for the search definition.

Enabling Batch on Real Time Indexing

Use the Configure Batch option on the Configure Real Time Indexing page to enable a batch program.

When you select the Configure Batch option for a search definition, you can access a list of batch programs, delivered by Applications for the search definition, that can cause high volume updates to the trigger records. You can choose additional batch programs from the drop-down, but these must be coded to use the batch switch. Oracle does not recommend adding batches other than the ones delivered for the search definition.

Select batch programs by process type and process name to turn off real time indexing for them.

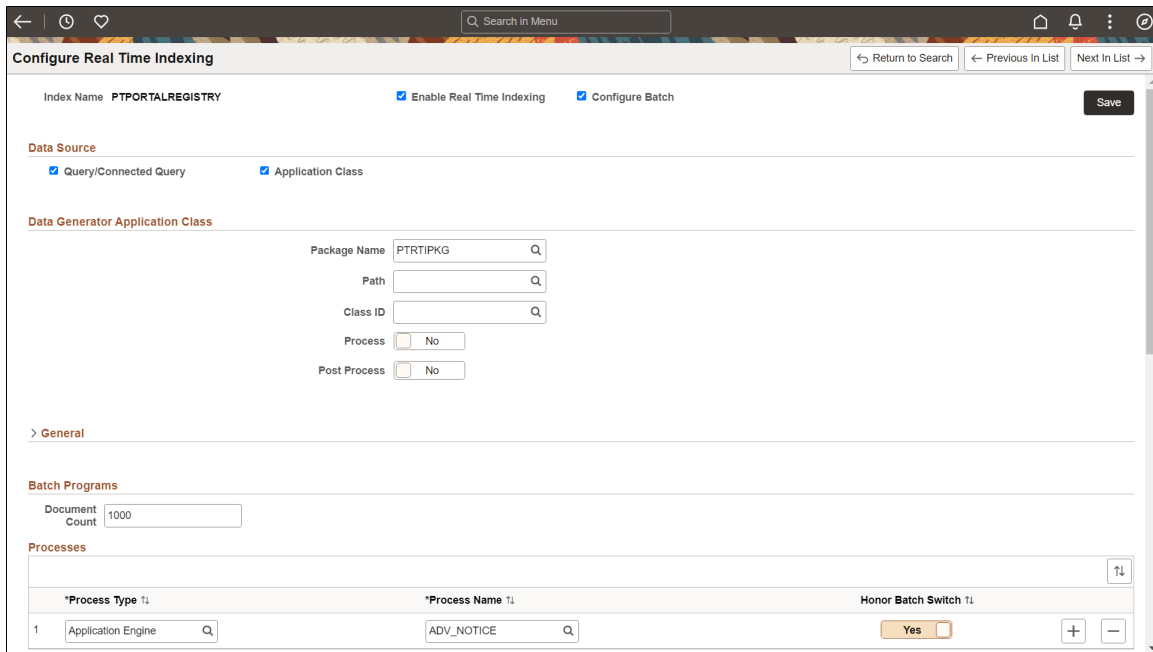
Oracle delivers functions that application developers can use to programmatically turn off real time indexing and set the document count value. However, these functions can only be used with Application Engine programs. See "TurnOffRTI" (PeopleCode Language Reference).

Programs using SQL or COBOL need to perform a direct SQL execution to turn off real time indexing.

Following is an example to turn off real time indexing in COBOL and SQR:

```
MERGE INTO PS_PTSF_RTI_STATUS d
USING(SELECT :1 PTSF_INDEX_NAME from dual)
ON(d.PTSF_INDEX_NAME = s.PTSF_INDEX_NAME)
WHEN MATCHED THEN UPDATE SET d.PTRTISTATUS = 0
WHEN NOT MATCHED THEN INSERT VALUES(:1, 0));
```

This example illustrates the fields and controls on the Configure Real Time Indexing page with the Configure Batch option selected:



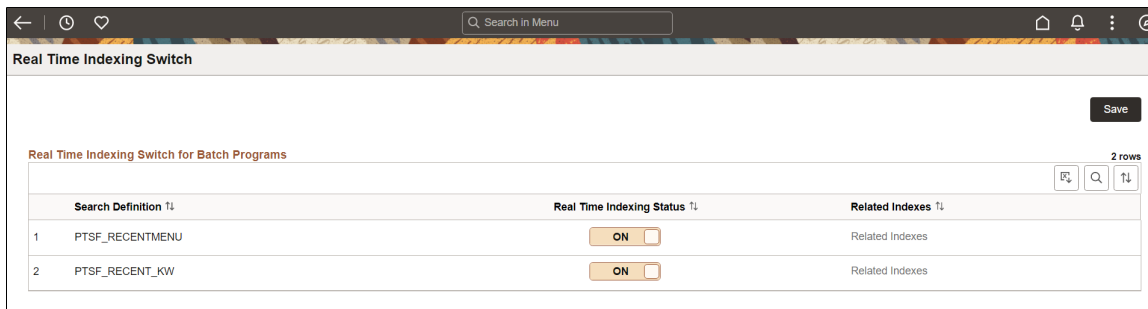
Field or Control	Description
Document Count	<p>The document count is a threshold beyond which the real time indexing process should be disabled during batch runs. Large batch jobs with records beyond this count should have real time indexing disabled.</p> <p>The default is set to 0. 0 means real time indexing is not turned off.</p> <hr/> <p>Note: This count will be used by application developers as a deciding factor to invoke the TurnOffRTI function. Applications will look for a number greater than 0 to compare to the record count for batch jobs to determine whether or not to disable real time indexing.</p>
Process Type	Select the process type.
Process Name	Select the process name.

Field or Control	Description
Honor Batch Switch	<p>Set the Honor Batch Switch option to Yes if you want the Real Time Indexing Switch to honor the process (batch program). When the Honor Batch Switch option is set to No, the request to turn off real time indexing for the batch program will be ignored.</p> <hr/> <p>Note: Applications will deliver the configuration with the Honor Batch Switch turned off. You must turn it to Yes if you have large batch runs and you must also set the Document Count.</p>

Using the Real Time Indexing Switch Page

The Real Time indexing Switch page displays the current real time indexing status of the listed search definitions. Administrators can use this page to check the real time indexing status and turn it on or off only in scenarios where real time indexing was not automatically turned on or off. Oracle recommends that administrators exercise caution before turning real time indexing on or off on this page.

This example illustrates the fields and controls on the Real Time Indexing Switch page:



Related Links

"TurnOffRTI" (PeopleCode Language Reference)

Using an Application Engine Maintenance Program

Real time indexing uses a staging table to store the keys of the application transactions to be indexed in Elasticsearch. The transaction data is removed from the staging table as soon as the transaction is indexed in Elasticsearch. Based on the number of search definitions enabled for real time indexing and the volume of transactions (online or batch), frequent data inserts and deletes take place in the staging table. This can eventually cause table fragmentation and over a period of time the performance of real time indexing may slow down due to the fragmentation as compared to the initial performance. Periodic maintenance of the staging table will restore the performance of the real time indexing process.

PeopleSoft provides an AE program, named PTRTI_TRUNC, which performs the maintenance activities on the staging table based on the value you set for the Real Time Indexing Threshold property on the Search Options page. The program obtains a lock on the staging table, backs up the data, and then truncates the table. After a successful truncate action, the backed-up data is copied to the staging table.

The successful run of the program depends on the activities on the staging table, that is, while locking and truncating the table, data insert should *not* take place on the table. If data insert takes place on the table after the lock on the table is obtained, the table is automatically unlocked to avoid any data loss. Therefore, the AE maintenance program should be run when no transactions are being inserted, deleted, or modified on the table.

For example, an administrator can plan to run the AE maintenance program after the successful completion of the application batch update. An administrator can schedule the AE program as a recurring job at regular intervals and the job, on each iteration, performs the maintenance activity based on the availability of the table lock. However, you should note that the AE program only performs the table truncate activity if the table crosses the Real Time Indexing Threshold value specified on the Search options page. The real time indexing process keeps track of the maximum count of the staging table rows and if the row count crosses the threshold value, the AE maintenance program is run. If the row count does not cross the threshold value, the AE program exits from the current run without proceeding for lock and truncate actions.

Related Links

[Managing General Search Options](#)

Excluding Search Definitions from Real Time Indexing

Use the Excluded Search Definitions page to hide a search definition in the Real Time Indexing Configuration page.

Note: Applications will deliver indexes to be excluded if the index is normally updated externally and thus is not eligible for real time indexing.

This example illustrates the fields and controls on the Excluded Search Definitions page:



Removing a Trigger

You may need to remove a trigger if you plan to remove real time indexing trigger from a specific transaction record. To remove an existing trigger, you should drop the trigger. However, if you drop a trigger, real time indexing can no longer update the associated search index.

Contact your database administrator to drop a trigger.

Complete these steps to drop a trigger:

1. Open an SQL tool.

2. Run the command: `DROP TRIGGER <trigger_name>;`

where `<trigger_name>` is the trigger that you want to drop.

For example, `DROP TRIGGER SFA_PT_PORTAL_MRU_TR;`

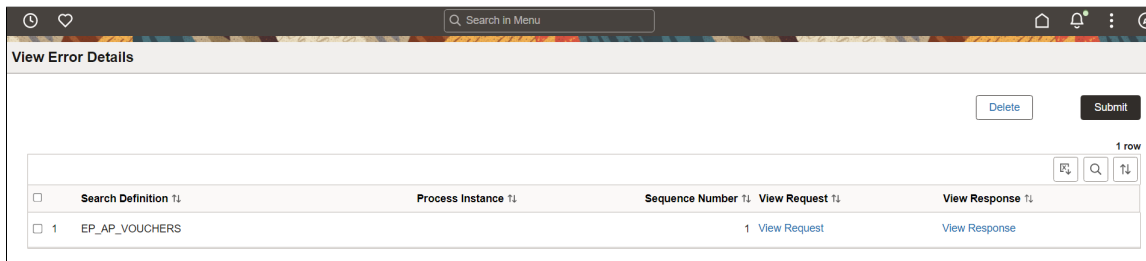
PeopleSoft RTI trigger names have the prefix SFA and suffix TR.

Viewing Error Details in the Real Time Indexing Process

Use the View Error Details page to check if any of the indexes configured for real time indexing encounter an error during the real time indexing process and to take required action.

Access the View Error Details page by selecting **PeopleTools >Search Framework >Real Time Indexing >View Error Details**.

This example illustrates the fields and controls on the View Error Details page. You can find definitions for the fields and controls later on this page.



Field or Control	Description
Index Name	Specifies the search index in which an error has occurred.
Request	Select the View Request link to view the real time indexing request information. Selecting the View Request link downloads a .json file.
Response	Select the View Response link to view the response from the Elasticsearch server. You can use the response information to identify the error and take an action to rectify the error. Selecting the View Response link downloads a .json file.
Submit	After reviewing the response, you may choose to run the real time indexing process again. Choose the search indexes that you want to run the real time indexing process again, and then select the Submit button. The Submit action recreates the .json file before sending to Elasticsearch.

Migrating Real Time Indexing Data Using ADS Definition

Oracle delivers an ADS definition to copy or migrate real time indexing metadata from one environment to another.

For information on using ADS, see "Understanding ADS Projects" (Lifecycle Management Guide).

ADS Definition for Real Time Indexing

ADS definition for real time indexing —PTSF_REALTIMEINDEXING.

PTSF_INDEX_GROUP

<i>Group</i>	<i>Details</i>
Group Name	Group for PTSF_INDEX
Group Description	Merge group for Search Index RTI metada
Group Member	<ul style="list-style-type: none"> • PTSF_INDEX.PTSF_ATTACHMENT • PTSF_INDEX.PTSF_RTI
Group Properties	Merge configurations

Tables for PTSF_REALTIMEINDEXING

<i>Table</i>	<i>Description</i>
PTSF_INDEX	Stores search definition details, such as connected query, query, application class, and so on.
PTSF_INDEXATT	Stores attachment information.
PTSF_INDEXATT_S	Stores SQL for attachments.
PTSF_INDEXCQRY	Stores parent child mapping information.
PTSF_INDEXDSEC	Stores document security.
PTSF_INDEXQRY	Stores queries in the index.

Table	Description
PTSF_INDEXRCFLD	Stores attribute information.
PTSF_RELLANGQRY	Stores related language information.
PTSF_INDEXTRREC	Stores trigger records for real time indexing.
PTSF_INDEXTRFLD	Stores trigger non-key fields for real time indexing.
PTSF_INDEX_BODY	Stores summary information.
PTSF_INDEX_DQ	Stores delete query information.
PTSF_INDEX_LUD	Stores last updated datetime details.
PTSF_INDEX_URL	Stores drilling URL information.

Troubleshooting

To troubleshoot any errors that you may encounter while real time indexing, you should ensure that you set the logging parameter, the Trace PC parameter, and the Trace SQL parameter.

Real Time Indexing Logs

In the Process Scheduler configuration file (psprcs.cfg), locate the PSRTISRV section, and enter the parameter for logging and set it to 5, for example:

```
LogFence = 5
```

Note: LogFence in the PSRTISRV section is applicable to the PSRTISRV server process only; not applicable to other servers.

Trace PC

In the Application Server configuration file (psappsrv.cfg), locate Trace PC and set the value to 2048, for example:

```
Trace PC = 2048
```

Trace SQL

In the Application Server configuration file (psappsrv.cfg), locate Trace SQL and set the value to 1, for example:

```
Trace SQL = 1
```

Chapter 12

Working with PeopleSoft Search Framework Utilities

Using the Search Test Page

This section contains an overview of search test page.

Understanding the Search Test Page

The search test page enables you to test an index independent of any existing PeopleSoft application PeopleCode or additional APIs. Using the Search Test page, you can isolate various search features and view the behavior, without signing on to the PeopleSoft application or the search engine. Also, by observing the underlying PeopleCode, the Search Test page can be a useful tool when studying the use and behavior of the query API. If you do not enter any criteria on the Search Test page, and click Search, the system will return all documents indexed for a search instance in the search engine. Using the criteria on this page, you can test how you anticipate search results to display for your application.

Access the Search Test page by selecting PeopleTools, Search Framework, Search Admin Activity Guide, Testing and Utilities, Search Test Page.

This example illustrates the fields and controls on the Search Test page. You can find definitions for the fields and controls later on this page.

The screenshot displays the Search Administration Activity Guide interface. At the top, there is a breadcrumb trail: "Search Administration Activity guide". Below this, there are two tabs: "Search" (active) and "Results". The main search area includes a "Search Instance" dropdown menu set to "PTSF_DEFAULT". Below this is a "Search Text" input field, a "Start Index" input field set to "1", and a "Docs Requested" input field set to "10". There is a "Language" dropdown menu set to "English" and a "Refresh Security Cache" button. A "Search" button is located to the right of the "Search Text" field. Below the search area, there is a "Search Category to Search in" section with a "Name" input field set to "PTPORTALREGISTRY" and a "Navigator" button. To the right of this section are buttons for "Get Facets", "Search Fields", and "Display Fields", along with a "Find" button and pagination controls showing "1 of 1". At the bottom, there are several expandable sections: "Custom Search Attributes", "Filter Settings", "Facet Request", "Grouping /Sorting Options", and "Additional Parameters".

Term	Definition
Search Instance	Choose a search instance from which you want to fetch search results.
Search Text	Enter the free text search request string to test.
Search	Select to initiate the search request against the desired search instance.
Start Index	Enables you control which result document displays first. For example, if you set it to <i>10</i> , the system displays the 10th search result at the top of the search result list.
Docs Requested	Enables you to control the number of search results displayed on the Results tab.
Language	Select the language for which you want to test the search index. Select blank to run against all languages.
Refresh Security Cache	<p>In Elasticsearch, when you select the Refresh Security Cache button, all the security values for the particular user are deleted from the Elasticsearch server. Then during the next search on a particular index, security values will be fetched from PeopleSoft for that index alone.</p> <p>When complete, you should see a message similar to:</p> <p><i>Security cache cleared successfully (262,1300)</i></p>
Duplicates	<p>Enables you to control how duplicates are displayed:</p> <ul style="list-style-type: none"> • <i>Remove</i>. Duplicates are not displayed. • <i>Mark</i>. Duplicates are displayed, but they are marked to indicate the duplicates. <hr/> <p>Important! To emulate the current default PeopleSoft Search Framework behavior, clear both <i>Remove</i> and <i>Mark</i> in the Duplicates group box.</p> <hr/>
Show image in result	If you have a drilling URL pointing to an image field, and you want to display the image in the test results, select this option.
Get Facets	Select the Get Facets button to automatically load the Facet Requests group box with all valid facets for the selected search category.
Search Fields	Select the Search Fields button to automatically load the Filter Settings group box with all valid filters for the selected search category.

Term	Definition
Display Fields	Select the Display Fields button to automatically load the Custom Search Attributes group box.

Testing Search Categories

The Search Category to Search in group box enables you to narrow the search only to a specific search category or set of search categories. Add search categories to the grid as needed to increase the scope of the search.

Testing Custom Search Attributes

The Custom Search Attributes group box enables you to identify specific search attributes and display the actual values for a search document in the search results. That is, in addition to the URL and summary in the search results, the results will also display the values for the fields you add to the Custom Search Attributes grid. The search attributes must be indexed fields in the search definition.

This group box can be automatically loaded by selecting the Display Fields button.

Testing Filter Settings

The Filter Settings group box enables you to search specific values of your attributes using the available operators.

This group box can be automatically loaded with all valid filters for the selected search category by selecting the Search Fields button.

If you have multiple attributes in the grid, use *Match Any* to return results that match any of the filters, and use *Match All* to display results that meet the criteria for all of the filters.

You can also enter a score for an attribute in order to boost the score of the attribute, which works in conjunction with the score operation option.

The following score operations can be used:

- Average - the averaged score is used.
- Maximum - the maximum score is used.
- Minimum - the minimum score is used.
- Sum - the scores are added.

Note: The operators displayed for testing filter settings on the Search Test page represent all possible operators that can be utilized by way of the delivered PeopleCode API. However, depending on the application's usage of the PeopleCode API and the data type of the field itself, end users will see varying display of the possible operators when defining search criteria. For example, *Equal To* applies only to number data types, while *Contains Phrase* applies only to character data types, and the appearance of the *Between* operator depends on the application's usage of the delivered PeopleCode. Refer to the PeopleSoft Search Framework PeopleCode documentation for more information.

Related Links

[Boosting Score of Search Results](#)

Testing Facet Requests

The Facet Request group box enables you to test and isolate specific facets. For example, you can find those that contain no search results in certain scenarios, or compare facet counts to attribute counts.

This group box can be automatically loaded with all valid facets for the selected search category by selecting the Get Facets button.

Field or Control	Description
Minimum Doc Count	Enables you to control the facets that display based on their count. For example, you can display facets that contain zero results, or only those which contain at least 10.
Maximum Facet Children	This applies to all facet requests. This determines the maximum number of facet nodes (facet values) returned by the API. The default is -1 where in all the children will be displayed.
Maximum Tree Facet Nodes	Enables you to control the nodes that display for hierarchical facets. The default value is 1000, which is configured on the Search Options page. You can override the value configured on the Search Options page by entering a new value.
Facet Name	Select the name of the facet you want to test.
Facet Path	This applies to all facet requests. The results will be filtered by the facet values specified in the path. In the case of hierarchical facets, the path can have the hierarchy information. For example, United States/CA/Pleasanton
Additional Details	Enables you to select or deselect the Allow Multi Select check box that controls the selecting of multiple facet values on the Results page.

Field or Control	Description
Return Count	Select to display the number of documents the facet contains.
Sort By	You can sort facets alphabetically or by count in ascending or descending order.

Testing Grouping and Sorting Options

The grouping and sorting options are not implemented in current release. These options are reserved for future use.

Testing Additional Parameters

The additional parameters options are not implemented in current release. These options are reserved for future use.

Downloading Search Data

If you need to share search definition information with Global Software Support (GSS) related to search definitions you are troubleshooting, use the Asynchronous Details page.

To download search data:

1. PeopleTools, Integration Broker, Service Operations Monitor, Monitoring, Monitor Asynchronous Services.
2. Select the Subscription Contracts tab.
3. In the Queue Name field, select PTSF_ES_SEND_Q.

PeopleSoft Search Framework uses this queue to communicate with the search engine server.

If you are troubleshooting the attachments, you need to select the PTSF_ES_ATTACH_SEND_Q queue.

4. Filter the search results using Status and Date to locate the desired sub-queue.
5. Click the Details link.
6. On the Asynchronous Details page, click the Download XML link.
7. Save the xml file to the location you want.
8. Open the file using a text editor program.

This example illustrates the Source XML Data.

```

<?xml version="1.0"?>
<data psnonxml="Yes">
  <![CDATA[--SEARCH INSTANCE START | PTSF_DEFAULT | SEARCH
INSTANCE END--{ "index" : { "_index" :
"orcl_es_defaultindex_t56805r2", "_type" :
"ptportalregistry_t56805r2", "_id" : "http://
\slc09kec.us.oracle.com:8000\psp\t56805r2x\PORTAL\QEDMO\c
\DEFINE_GENERAL_OPTIONS.RUN_FIN0004.GBL" } }
{"ORCL_ES_LINK" : "http://\slc09kec.us.oracle.com:8000\psp
\t56805r2x\PORTAL\QEDMO\c
\DEFINE_GENERAL_OPTIONS.RUN_FIN0004.GBL", "ORCL_ES_TITLE" :
"Account Types", "ORCL_ES_DESCRIPTION" : "<BR>Application:
QE_LOCAL", "ORCL_ES_LAST_MODIFIED_DATE" : "2016-08-
17T17:50:59.000Z", "PTSF_SBO_NAME" : "T56805R2
PTPORTALREGISTRY", "ORCL_ES_IMAGE_URL" : "http://
\slc09kec.us.oracle.com:8000\psc\t56805r2x\EMPLOYEE
\T56805R2\q\?
ICAAction=ICQryImageUrl=PUBLIC.PTPORTALTILES&IMGFIELD=CONTIMG&IMG
RECORD=PTSTILEIMG_VW&PORTAL_NAME=PORTAL&PORTAL_OBJNAME=RUN_FIN000
4&PORTAL_REFTYPE=C", "DESCR254" : [ "", null ], "PORTAL_LABEL" : [
"Account Types", null ], "PORTAL_NAME" : [ "PORTAL", null
], "PORTAL_PRNTLABEL_PORTAL_FLDR" : [ "Define Business
Rules", null ], "PORTAL_SOURCE" : [ "Application: QE_LOCAL", null
], "PTSF_NAV_MODE" : [ "1", null ], "PORTAL_SECTYPE_ORCL_ES_ACL" :
[ "P:ALLPNLS", "P:QEALLPGS", "S:Admin", null
], "ORCL_ES_LANGUAGE_CD" : "ENG"}
{ "index" : { "_index" : "orcl_es_defaultindex_t56805r2",
_type" : "ptportalregistry_t56805r2", "_id" : "http://
\slc09kec.us.oracle.com:8000\psp\t56805r2x\PORTAL\QEDMO\c
\DEFINE_GENERAL_OPTIONS.RUN_FIN0005.GBL" } }
{"ORCL_ES_LINK" : "http://\slc09kec.us.oracle.com:8000\psp
\t56805r2x\PORTAL\QEDMO\c
\DEFINE_GENERAL_OPTIONS.RUN_FIN0005.GBL", "ORCL_ES_TITLE" :
"Units of Measure", "ORCL_ES_DESCRIPTION" : "<BR>Application:
QE_LOCAL", "ORCL_ES_LAST_MODIFIED_DATE" : "2016-08-
17T17:54:59.000Z", "PTSF_SBO_NAME" : "T56805R2
PTPORTALREGISTRY", "ORCL_ES_IMAGE_URL" : "http://
\slc09kec.us.oracle.com:8000\psc\t56805r2x\EMPLOYEE
\T56805R2\q\?

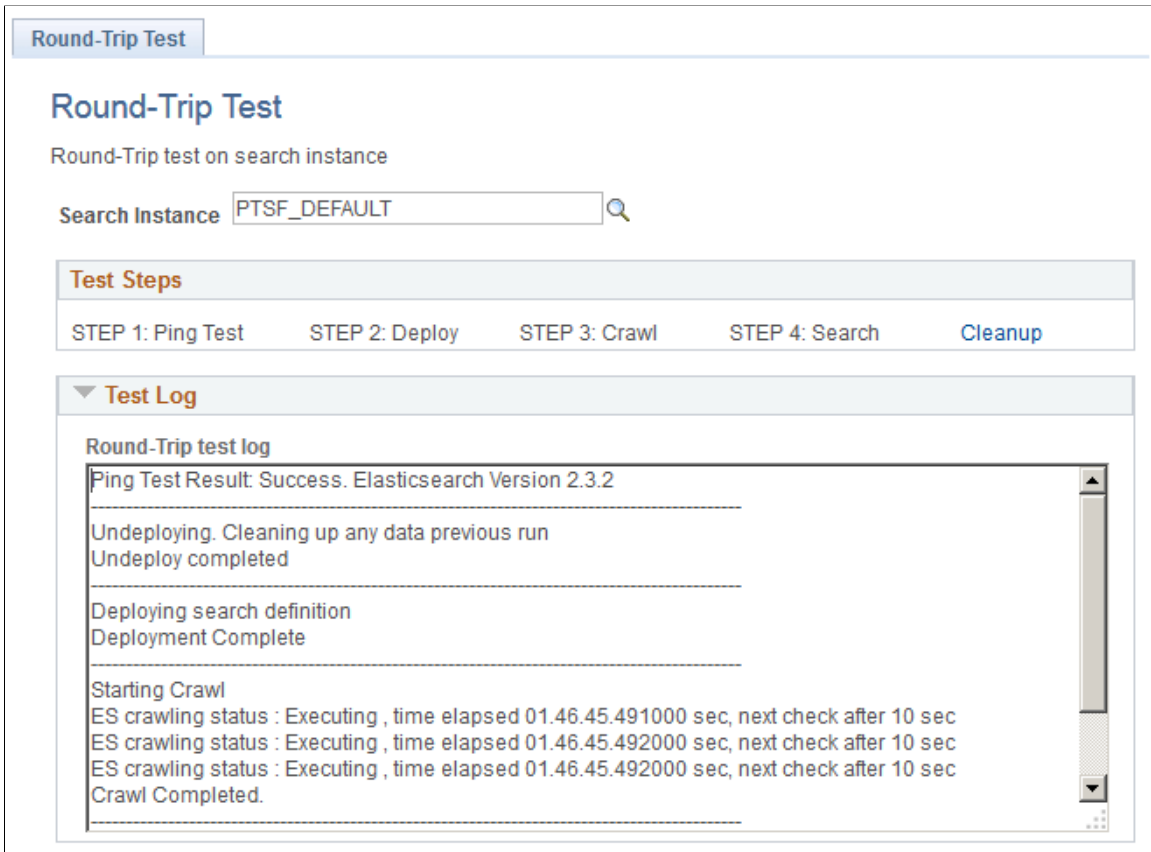
```

Running Diagnostics

Access the Round-Trip Test page by selecting PeopleTools, Search Framework, Search Admin Activity Guide, under Testing and Utilities, then select Round-Trip Test.

The round-trip test performs a set of diagnostic steps using a test search definition, delivered with PeopleTools named RNTRPTST. It provides end-to-end diagnostic tests to verify and troubleshoot the PeopleSoft and search engine integration.

This example illustrates the fields and controls on the Round-Trip Test page. You can find definitions for the fields and controls later on this page.



Field or Control	Description
Search Instance	Select the appropriate search instance defined for this PeopleSoft system.
Test Steps	<p>The Test Steps grid contains a series of links that are ordered in a typical troubleshooting sequence. Only one link is active at a time, and each link becomes active only after the previous step has completed successfully.</p> <ol style="list-style-type: none"> 1. Ping Test: Runs a ping test against the search engine server instance defined for your system. 2. Deploy: Deploys the RNTRPTST search definition and associated category to the search engine. 3. Crawl: Runs a test crawl using the RNTRPTST definition. 4. Search: Runs a sample search based on the index populated by the RNTRPTST search definition. 5. Cleanup: Click to remove the deployment of the RNTRPTST search definition and run any related cleanup tasks.

Field or Control	Description
Test Log	Indicates whether the system successfully created the feed output, required to populate the index.

Viewing Event Logs

View event logs on the Events Log page by selecting PeopleTools, Search Framework, Search Admin Activity Guide, Testing and Utilities, Events Log.

This example illustrates the fields and controls on the Event Logs page. You can find definitions for the fields and controls later on this page.

Field or Control	Description
Event Name	Select the name of the event by which to filter the logs.
Event Severity	Narrow the event display focus by selecting one of these severity levels: <ul style="list-style-type: none"> All: Displays all types of event logs (Error, Message, Warning). Error: Displays only error messages. Message: Displays the typical status messages the system records after each event. Warning: Displays only warning messages.
View Events	Click to display the events meeting your criteria in the Events grid based on Time Period settings. For example, after selecting a time period, click View Events to display only the logs that fall within that period.

<i>Field or Control</i>	<i>Description</i>
Purge Logs	After selecting a time period, click Purge Logs to remove from the system the logs from that time period. Click View Events to refresh the events list. This can help to keep the stored log information at a manageable size.
Time Period	Specify the specific time period for which you seek event log information.

Working with PeopleSoft Search Framework Security Features

Understanding Search Framework Security

By default, Search Framework and the related search engine possess security features designed to protect all of the data and processes within the system. PeopleSoft adds two plug-ins as security measures required for authenticating and authorizing users in the search system. These security plug-ins are bundled with the installation of the search engine. This section describes security topics that pertain to the integration between Search Framework and the search engine.

For installing Elasticsearch, see *PeopleSoft Deployment Packages for Elasticsearch Installation (PeopleSoft PeopleTools 8.56) on My Oracle Support (Doc ID: 2205540.2)*.

Related Links

[Working with Authentication and Authorization](#)

Applying PeopleSoft Permissions

The implementation, maintenance, and use of the Search Framework involve these user types:

Role	Description	Delivered PeopleTools Permission List
Search Administrator	Responsible for managing the deployment of search definitions and search categories, scheduling index builds, monitoring indexes, and establishing connectivity between Search Framework and the search engine. Note: This user needs to have access to all the queries (records) on which the search definition is based on in order to schedule the index generation.	PTPT3100
Search Developer	Responsible for creating search queries, search definitions and search categories.	PTPT3200

Role	Description	Delivered PeopleTools Permission List
Search Server	<p>Search engine search instance requiring access to the Search Framework service operations.</p> <hr/> <p>Note: This is the call back ID configured in the search instance page.</p> <hr/> <p>The Search Server role is required to be able to download attachments.</p>	PTPT3300
End User	<p>Runs search queries while using PeopleSoft applications, using Global Search or Search Pages.</p> <p>See the <i>Application Fundamentals PeopleBook</i> <for your product line> for the relevant roles needed for Global Search.</p>	<p>None specific to Search Framework. Restrictions to search results can be implemented by user profile or role.</p> <p>PTPT1000 is required for the PTPORTALREGISTRY search definition.</p>

Note: Search Administrator, Search Developer, and Search Server are roles delivered by PeopleTools.

Working with Authentication and Authorization

Search Framework handles various security related tasks, including:

- Authenticating users (development, administration, and end users).
- Authenticating systems requesting access to service operations.
- Authorizing end user search requests.

When managing search requests with Search Framework, it is important to distinguish between *authentication* and *authorization*.

Authentication determines if a user is a legitimate user, who can access the system. Authentication is configured using PeopleSoft user profiles, roles, and permission lists.

Authorization determines the access level for an authenticated user. Once a user is authenticated, the system invokes the authorization rules. You define authorization (access controls) per search definition on the Security tab. For some instances, applying No Security is a valid option. However, for other situations, you need to apply stricter control over what users can and cannot see. You can restrict access by the source (search definition) or by the document (search result).

Source-level security applies to all the documents in the data source. Setting source-level security is useful when you want to prevent global visibility of data source content. When defining a source-level security you specify the users and roles that can view the search results for that search definition. When a user searches the associated index, the system verifies the user's access level prior to displaying any search results.


Document-level security restricts access to specific search results. The document-level authorization uses security attributes. The attributes are defined using PeopleSoft Query or Connected Query during design time and are used to evaluate access during runtime. While defining a search definition, some of the fields chosen for the query may not be used necessarily for searching, but mainly as security attributes. For example, Department ID or Business Unit are examples of attributes that users may not necessarily search on, but their values can be used in the authorization process to evaluate if a user can view search results for a specific Department ID or Business Unit.

If applying any level of security, you should define an application class for every search definition. The application class is responsible for fetching a list of runtime values based on the security attributes. You associate the application class with a search definition on the Security tab of the search definition.

Authentication and Authorization in Elasticsearch

The following diagram illustrates the elements involved with authorization and authentication and the interaction between them.

This example illustrates how PeopleTools and Elasticsearch interact to authenticate users and authorize user access to search results.



Step	Description
1	During the index build, the crawler collects the defined security attributes and access restriction options associated with search definitions. The crawler applies those security attributes to the data sources (indexes).
2	An application end user runs a search query.

Step	Description
3	Elasticsearch security framework using the <i>orcl-security-plugin</i> invokes the PeopleSoft security service for user authentication.
4	Based on the authenticated user, Elasticsearch performs either a source-level or a document-level authorization using the <i>orcl-security-plugin</i> . For document level authorization PeopleSoft invokes the Application class defined for the Data Source. This would fetch a list of values for the security attribute which then Elasticsearch would use for filtering the search results.

Configuring SSL between PeopleSoft and Elasticsearch

You can configure SSL connections between your PeopleSoft system and Elasticsearch. When you have SSL configured, you then need to set the SSL Option field on the Search Instance Properties page to ENABLE.

To set up SSL on the PeopleSoft system, use the instructions provided in the System and Server Administration guide. See "Implementing WebLogic SSL Keys and Certificates" (System and Server Administration).

After you configure SSL in Elasticsearch, add the CA root certificate of Elasticsearch to the PeopleSoft keystore so that it will be trusted.

1. Copy the Elasticsearch CA root certificate (for example, trustedCA.cer) to the Webserver folder.
2. Use **pskeymanager -import** command to import the certificate to the PeopleSoft truststore.

Configuring SSL for Elasticsearch

To set up SSL in Elasticsearch, complete these steps in all the nodes of Elasticsearch cluster:

Note: Keystore and Truststore files should be placed under the ES_HOME/config directory.

1. Import the trusted root from Certification Authority (CA) and save it locally, for example, /home/certs/cacert.cer.
2. Set up truststore and import trusted root certificate from CA to the truststore using JAVA_HOME/bin keytool.

```
keytool -importcert -keystore ES_HOME/config/keystore/mytruststore.jks -file / =>
home/certs/cacert.cer -alias my_ca
```

When prompted, provide a password. This should be later encrypted and placed in the elasticsearch.yml configuration file.

In steps 7, 8, and 9, when prompted whether you trust this certificate, enter *Yes*.

3. Set up keystore and private key.

```
keytool -genkey -alias node_alias -keystore ES_HOME/config/keystore/mykeystore⇒
.jks -keyalg RSA -keysize 2048 -validity 90
```

When prompted, provide a password. This should be later encrypted and placed in the `elasticsearch.yml` configuration file.

Enter details for the questions prompted on the screen.

Enter key password for `<node_alias>` when prompted on the screen.

4. Create Certificate Signing request (CSR).

```
keytool -certreq -alias node_alias -keystore ES_HOME/config/keystore/mykeystor⇒
e.jks -file /home/certs/mycsr.csr -keyalg rsa
```

This step creates a CSR.

5. Use the CSR to get a certificate from CA. Download it and save it locally, for example, `/home/certs/signed.cer`.

6. Import the certificate to the Keystore. Use the same alias as the private key and certificate request.

```
keytool -importcert -keystore ES_HOME/config/keystore/mykeystore.jks -file /ho⇒
me/certs/signed.cer -alias node_alias
```

7. In the `elasticsearch.yml` configuration file, add these properties to set up SSL:

Field or Control	Description
<code>orclssl.http.ssl</code>	true
<code>orclssl.transport.ssl</code>	Set the property to <i>true</i> if you want to enable SSL in transport layer for node to node encryption. Enter <i>false</i> if you want to disable SSL in transport layer.
<code>orclssl.keystore</code>	<path to keystore>
<code>orclssl.keystore_password</code>	<keystore password>
<code>orclssl.truststore</code>	<path to truststore>
<code>orclssl.truststore_password</code>	<truststore password>
<code>orclssl.callback</code>	To set this property, refer to the Call Back Properties section in Creating Search Instances .

- Use `elasticsearchuser` script to encrypt the keystore and truststore password.

```
ES_HOME/bin> elasticsearchuser encrypt [password]
```

- Use the encrypted password for the `orclssl.keystore_password` and `orclssl.truststore_password` in the `elasticsearch.yml` configuration file.

Note: As both HTTP and HTTPS use the same channel, the port for HTTPS is the same as the one configured in `http.port` in the `elasticsearch.yml` configuration file. After SSL is enabled, the port configured for HTTP becomes HTTPS and HTTP will be disabled.

Configuring SSL for Kibana

You can configure SSL for Kibana. When you configure SSL for Kibana, you need to set the SSL Option field (in the Kibana section) on the Search Instance Properties page to ENABLE.

To set up SSL for Kibana, complete these steps:

- Configuration required in `kibana.yml`:

```
server.ssl.enabled: true
server.ssl.certificate: <certificate issued by server> (for example, signed.cer)
server.ssl.key: <keystore> (for example, mykeystore.key)
```

If Elasticsearch is SSL enabled, you need to complete the following additional configuration:

```
elasticsearch.ssl.certificateAuthorities: (for example, "C:\ES\elasticsearch7.0\plugins\orcl-security-plugin\config\properties\cacert.cer")
```

Note: Keystore and certificate should be placed in the `bin` (default) folder or `config` folder (specify the path in `kibana.yml`).

- Import the trusted root from CA and save it locally (for example, `D:\ca\cacert.cer`).
- Import root certificate from CA to the keystore using `keytool`:

```
JAVA_HOME\bin>keytool -importcert -keystore KIBANA_HOME/bin/mykeystore.jks -file D:\ca\cacert.cer -alias my_ca
```

- Setting up keystore and private key:

```
keytool -genkey -alias alias1 -keystore KIBANA_HOME/bin/mykeystore.jks -keyalg RSA -keysize 2048 -validity 712
```

- Create certificate signing request:

```
keytool -certreq -alias alias1 -keystore KIBANA_HOME/bin/mykeystore.jks -file D:\ca\mycsr.csr -keyalg rsa
```

This command creates a CSR as follows (this CSR is only for sample reference):

```
-----BEGIN NEW CERTIFICATE REQUEST-----
```

```

MIIC6zCCAdMCAQAwdjELMAkGA1UEBhMCSU4xEjAQBgNVBAsTCUthcm5hdGFrYTESMBAGA1UEBxMJ Q⇒
mFuZ2Fsb3JlMQ8wDQYDVQQKEwZPcmFjbgUxEzARBgNVBAsTC1Blb3BsZXNvZnQxGTAXBgNVBAMT EF⇒
JvaGluaSBQYWxsaXlhbGkwggEiMA0GCSqGSIb3DQEBAQUAA4IBDwAwggEKAoIBAQCCT1Tr4o6Yi v/d⇒
kwqhiCJ4d6ddVhyGTMzBHSQB6tv10GvNWjmJMqXWepAsu6gOgoECY0HwGC4L1bQh7hiwyT3ub ckTY⇒
nWULNgf1WulUIpyU9Z3Aj1BV3uhZWWJPnTC1JRyXIdvOMocpIM3YdQF1ZY3eOMY3Y1KT3ZMO GnDQz⇒
pSLFuXcVyAWbg32LKE9SW2zVIM8ueb6h1szv3U3KhJB7dI5inpoXg7cpnZxzjUK303HhB5l im003a⇒
PslhLb9rt9KjhAI4nxrpM9FEoAZI9V1NVpqfIksdBLVRnCcZGbnqH5n2nW3on5OFoNoRUI mTBc6Vs⇒
wPxHDz+bBAgcE0U8ETY39AgMBAAGgMDAuBqkqhkiG9w0BCQ4xITAFMB0GA1UdDgQWBbTJQ2nG1pF2k⇒
e5Z8HBrfvnBsgUjWTANBqkqhkiG9w0BAQsFAAOCAQEAEmcwNcQnAPTxfpHHbUIpsY+/ NhmfaltnhD⇒
q6AAcOL/rBmymgafdqDlIGWJ7tYJ/1zCjkFx9zeIh6RKhBHjzpuF8uM7Q1JJVB7CKI GP5UxJEs4bs⇒
gOqjSs8m71zGEW1D9gopYhGQJmcJr929NeR9k4cMMLpSbpsox2CPjstiepzIVKJEq ppqG3G+PXUZM⇒
W04Va/SCGHMndcLsd6XV3I5UWNCMNuJFEyBF4KVH6EJf59/vt3L2tQaV2X7JydcH 81laFxBPwCJe2⇒
GCB7Yo9Pi90lkDFIk2cAHry576Gf/T7GwHBg4/4T/vCRV1aZtKaRYJDmUW+5ae 4Nmzp3YwcTrdZg⇒
==
-----END NEW CERTIFICATE REQUEST-----

```

6. Use the CSR to get a certificate from CA. Download it and save it locally (for example, D:\ca\signed.cer).

7. Import the certificate to the keystore:

```

keytool -importcert -keystore KIBANA_HOME/bin/mykeystore.jks -file D:\ca\signe⇒
d.cer -alias alias1

```

8. Convert the keystore file to pem format because Kibana does not support any other format.

```

keytool -importkeystore -srckeystore KIBANA_HOME/bin/mykeystore.jks -destkeyst⇒
ore KIBANA_HOME/bin/mykeystore.p12 -srcstoretype jks -deststoretype pkcs12

```

Move the intermediate mykeystore.p12 file to Linux and execute the below commands:

```

openssl pkcs12 -in mykeystore.p12 -out mykeystore.pem

```

Or directly use this command:

```

keytool -list -rfc -keystore "mykeystore.jks" | sed -e "/-BEGIN [A-Z]*-*/,-*⇒
END [A-Z]*-*/!d" >> "myKeystore.pem"

```

9. Convert the pem file to .key format:

```
openssl rsa -in mykeystore.pem -out mykeystore.key
```
10. Update the kibana.yml file with the SSL settings.

Configuring SSL for Logstash

When you plan to configure SSL for Logstash, you may consider the following scenarios:

Scenario A: If you are planning to configure SSL during the installation of Logstash, follow these sequence of steps:

1. Complete the installation of Logstash.
2. Configure SSL for Logstash. Follow the steps described in the Configuring SSL section.
3. Start the Logstash service.
4. Set up external data integration with PeopleSoft.

See [Setting Up the Integration of External Data with PeopleSoft](#).

Scenario B: If you have installed Logstash and completed setting up external data integration with PeopleSoft, follow these sequence of steps:

1. Navigate to the Manage External Search Index page (**PeopleTools >Search Framework >Deploy External Source Defns >Manage External Search Index**) and select Stop to abort the indexing process.

See [Building an Index for an External Data Search Definition](#).

2. Delete the external search index if not required.
3. Stop the Logstash service.
4. Configure SSL for Logstash. Follow the steps described in the Configuring SSL section.
5. Start the Logstash service.
6. Set up external data integration with PeopleSoft.

See [Setting Up the Integration of External Data with PeopleSoft](#).

Configuring SSL

When you configure SSL for Logstash, you need to set the SSL Option field (in the Logstash section) on the Search Instance Properties page to Enable.

To set up SSL for Logstash, complete these steps:

1. Encrypt the keystore password using the following steps:

Windows:

```
set LOGSTASH_HOME= <Install_Dir>\pt\Logstash7.10.0
set JAVA_HOME= <JAVA_HOME>
LOGSTASH_HOME\pt\bin>PSLSCipher.bat <keystore password>
```

Linux:

```
export LOGSTASH_HOME= <Install_Dir>\pt\Logstash7.10.0
export JAVA_HOME= <JAVA_HOME>
```

```
LOGSTASH_HOME\pt\bin>PSLSCipher.sh <keystore password>
```

2. Edit the following parameters in the `http` section of the `orcl_psft_main.config` file, which is located at `<LOGSTASH_HOME>/orcl_psft_ext_config/orcl_psft_main.config`. Copy the encrypted password from step 1 and include in the `keystore password` parameter.

```
input {
  http {
    host => "xxxxxxx" # default: 0.0.0.0
    port => yyyy # default: 9800
    ssl => "true"
    keystore => "<Path to Keystore>/mykeystore.jks"
    keystore_password => "<Encrypted Keystore password>"
  }
}
```

3. Edit the `cacert` parameter in the `output_type_elasticsearch` section of the `psftextdata.template` file, which is located at `<LOGSTASH_HOME>/orcl_psft_ext_config/orcl_psft_main.config`.

For the `cacert` parameter, enter the path to the certificate file. For example, if path to the certificate file is `C:\dir1\dir2\RootCA.cer`, the `cacert` parameter value should be `cacert => "C:\dir1\dir2\RootCA.cer"` as shown in the following sample:

```
<output_type_elasticsearch>
  elasticsearch {
    hosts => ["<output.elasticsearch.hosts>"]
    index => "<output.elasticsearch.index>"
    user => "esadmin"
    password => "<Encrypted Password>"
    cacert => "C:\dir1\dir2\RootCA.cer"
  }
</output_type_elasticsearch>
```

4. Start the Logstash server using the following batch scripts:

Windows:

```
start_psftext_logstash.bat
```

Linux:

```
start_psftext_logstash.sh
```

After setting up SSL for Logstash, you can configure the external data setup, and then begin the indexing process.

Setting Up Role-Based Search Group Access

Global Search displays search groups in the Global Search Bar based on the user's role. If a search group assigned to a context belongs to the permissions for a role to which that user belongs, then the search group will appear in the search group drop-down list for that user. You configure search group access by selecting PeopleTools, Security, Permissions & Roles, Permission Lists, Search Groups.

Note: PeopleSoft Applications delivers pre-defined roles for each search group they define. Normally these roles will begin with the word 'Search'. See the *Application Fundamentals PeopleBook <for your product line>* for a complete list of all pre-defined roles.

This example illustrates the fields and controls on the PeopleTools Permission Lists: Search Groups page. You can find definitions for the fields and controls later on this page.

*Search Group Name	Description		
PTPORTALREGISTRY	Navigator	+	-
PTSEARCHREPORTS	Reports	+	-

Use the Search Groups grid to add search groups to which you want to add access. Search groups are those search categories that have the Search Group check box selected on the General tab of the search category definition.

Note: This does not define access to the data, but only gives a mechanism to control what search groups are seen in the Global Search Bar for any given user. Component-level security must still be granted to the user in order for them to click a result and be taken into the data.

Related Links

[Working with Global Search](#)

Managing Configurable Search

Understanding Configurable Search

Configurable search enables you to configure search pages for classic and fluid components based on your business requirements. It also enables you to maintain the same search experience across the PeopleSoft system. The PeopleTools delivered configurable search provides areas for configuration as described in this topic.

Configurable search enables you to manage search pages with minimal effort and reduces the dependency on Application Designer. Administrators can configure search pages for components to use either a search record (standard search) or a search definition (keyword search).

Search Pages Delivered for Configurable Search

PeopleTools delivers search pages for fluid and classic components, which are dynamically presented as the search page at runtime based on certain conditions that are described later in this topic.

PeopleTools delivers the following search pages:

- PTS_CFGSRCH_FL search page for fluid components.
- PTS_CFGSRCH_CL search page for classic components.

Using the Manage Configurable Search page, you can configure the dynamic presentation of the delivered search pages to suit your business requirements.

Example of configurable search for a classic component.

The screenshot shows a search interface titled "Find an Existing Value" with a button "Add a New Value" in the top right. Below the title is a section "Search Criteria" with a dropdown arrow and a note: "Enter any information you have and click Search. Leave fields blank for a list of all values." There are two search history sections: "Recent Searches" and "Saved Searches", each with a dropdown menu and an edit icon. The main search area includes a "Search by:" dropdown set to "Permission List", followed by the operator "contains" and an empty input field. Below this is a "Show more options" dropdown, which is expanded to show a "Case Sensitive" checkbox. At the bottom are "Search" and "Clear" buttons.

Related Links

[Specifying System-wide Configurable Search Properties](#)

[Working with Configurable Search](#)

Considerations for Using Configurable Search

Consider the following when you use configurable search in classic or fluid components:

- Supported primary action.

Configurable search does not support a combination of new, search, and keyword search as primary action for a single component. Configurable search supports the following combinations for primary action:

- New only.
- Search only.
- Keyword Search only.
- New and Search.
- New and Keyword Search.

For descriptions of these primary actions, see [Working with Search Pages](#).

If your business requirements demand both search and keyword search as primary action, you may consider registering two content references for the single component. In this case, you can configure one content reference for standard search and configure the other content reference for keyword search.

Component Properties page - primary actions and search/lookup types.

Component Properties

General
Use
Internet
Fluid
Style

Search Page

Primary Action New Search Keyword Search

Default Search Action: Update/Display ▼

Default Search/Lookup Type: Basic Advanced Keyword Only

Allow Action Mode Selection

Link To Add Page	Message Set/Nbr:	124	62
Link To Find Existing Value Page	Message Set/Nbr:	124	63
Link To Realtime Search Page	Message Set/Nbr:	124	63
Link To Keyword Search Page	Message Set/Nbr:	124	433
Instructional Text	Message Set/Nbr:	124	50

- Supported search or lookup types.

Configurable search supports only Basic and Advanced as the search/lookup types for the search and keyword search primary actions. Keyword Only search/lookup type is *not* supported. For example, if the Keyword Only search/lookup type is selected for a component, and if the component is enabled for configurable search, the Basic search is displayed at runtime; the Keyword Only option is not honored.

For descriptions of search/lookup types or search modes, see [Working with Search Pages](#).

- In configurable search, the number of search results is specified on the PeopleTools Options page. The default value for the number of search results is 300, however you may modify this value on the PeopleTools Options page. Configurable search uses this value to limit the number of search results across classic and fluid components and standard and keyword search.

See [Specifying System-wide Configurable Search Properties](#).

- In configurable search, related actions are *not* displayed in search results, by default. However, you can display the related actions in search results by setting the **Show Related Actions in Search Results** property to *Yes*.
- Configurable search supports facets on keyword search pages.
- Configurable search supports data masking for sensitive data and personally identifiable information (PII).

Guidelines on Using Configurable Search for Classic Components

Follow these guidelines when you plan to configure search for classic components:

- By default, the delivered (new) classic search page is displayed instead of the old standard component search page without any additional configuration on your part.
- If search is enabled for a component (that is, a valid search record is configured for the component and a search key exists), which search page gets displayed is based on these conditions:
 - If a custom search page (where page type is set to Search Page on the Use tab of the Page Properties dialog box in Application Designer) has been added to the component definition, the custom search page is displayed.
 - If configurable search has been configured for the content reference, then the delivered (new) classic search page is displayed reflecting that configuration. If configurable search has not been configured for the content reference, then the delivered (new) classic search page is displayed in its default state.
- If a component is not search enabled (that is, the component search record has no search keys and the Update/Display check box on the Use tab of the Component Properties dialog box is not selected), the system directly loads the transaction page, which is similar to the functionality prior to PeopleTools 8.60.

Guidelines on Using Configurable Search for Fluid Components

Follow these guidelines when you plan to configure search for fluid components:

- By default, search pages on fluid components are not replaced with the delivered (new) fluid search page. That is, for fluid components, you need to enable configurable search. You can enable configurable search either through Application Designer (using the Enable Configurable Search check box on the Component Properties dialog box) or through browser (using the Manage Configurable Search page). Oracle recommends that you enable configurable search for fluid components through browser because it's easier and cost effective.

For information on using the Enable Configurable Search check box, see "Using the Enable Configurable Search Check Box for Fluid Components" section later in this topic.

For detailed steps on configuring a search page, see [Working with Configurable Search](#).

- By default, new fluid components use the delivered (new) fluid search page.
- If search is enabled for a component (that is, a valid search record is configured for the component and a search key exists), the system displays a different search page based on these conditions:
 - If a custom search page (of type Search Page) has been added to the component definition, the custom search page is displayed.
 - If configurable search has been configured for the content reference, then the delivered (new) fluid search page is displayed reflecting that configuration. If configurable search is merely enabled on the component definition but there is no content reference configuration, then the delivered (new) fluid search page is displayed in its default state.
 - If configurable search is not enabled and a pivot grid configuration exists for the component, the pivot grid search page is displayed.
- If a component is not search enabled (that is, the component search record has no search keys and the Update/Display check box on the Use tab of the Component Properties dialog box is not selected), the transaction page is loaded directly, which is similar to the functionality prior to PeopleTools 8.60.

Using the Enable Configurable Search Check Box for Fluid Components

Application developers should note that to enable or disable configurable search for fluid components, you can use the **Enable Configurable Search** check box on the Fluid tab of the Component Properties dialog box in Application Designer. If you choose this route to enable configurable search, you have to deliver the component through a project. However, you should be aware that you do not have to use this route of delivering a component for configurable search because you can use the Manage Configurable Search page in browser to enable configurable search.

You may plan to deliver a component for some business requirements, and in that case you can enable configurable search on the component.

Note: If you plan to select the **Enable Configurable Search** check box on a delivered component, you should be aware that it might have future upgrade impact.

See "Setting Properties on the Fluid Tab" (Fluid User Interface Developer's Guide).

Considerations for enabling configurable search using the Enable Configurable Search check box:

- On existing components, if the **Search Page Type** field is set to None, the **Enable Configurable Search** check box is unavailable.

- If you set **Search Page Type** to Standard, the **Enable Configurable Search** check box is available for use.

In PeopleTools 8.60, only Standard search page type is supported for configurable search.

- On new components, the **Enable Configurable Search** check box is selected by default provided the Search Page Type is set to Standard.

Creating Custom Search Pages

If you plan to create custom search pages, Oracle recommends that you clone the delivered search pages and then modify them.

When you create a custom search page, ensure that you specify the page type as **Search Page** on the Use tab of the Page Properties dialog box in Application Designer. You should insert the custom search page in a component and redeliver the component to make the custom search page available to your end-users who can interact with it.

You can create a custom search page and apply it for a single component or you can apply it system-wide, which can be specified on the PeopleSoft Options page.

See [Specifying System-wide Configurable Search Properties](#).

Working with Configurable Search

The implementation of configurable search requires you to specify system-wide search properties on the PeopleTools Options page and specify component specific search properties on the Manage Configurable Search page.

Specifying System-wide Configurable Search Properties

On the PeopleTools Options page, use the Configurable Search Properties group box to specify system-wide configurable search properties.

Access the PeopleTools Options page by selecting **PeopleTools >Utilities >Administration >PeopleTools Options**.

This example illustrates the properties for configurable search on the PeopleTools Options page. Descriptions of the properties appear after the example.

Configurable Search Properties

Classic Search Page <input style="width: 100%;" type="text"/>	Default rows in search results <input style="width: 100%; text-align: center;" type="text" value="300"/>
Fluid Search Page <input style="width: 100%;" type="text"/>	Max Saved Searches <input style="width: 100%;" type="text"/>
Max Recent Searches <input style="width: 100%;" type="text"/>	

Recent/Saved Search Format

Label value For Example: Description: Manager

Label operator value For Example: Description (begins with): Manager

Field or Control	Description
Classic Search Page	<p>The search page that you specify here acts as the system-wide search page for classic components.</p> <p>Clone the delivered classic search page, PTS_CFGSRCH_CL, to create a custom search page. However, you may create your own search page without cloning the delivered search page.</p>
Fluid Search Page	<p>The search page that you specify here acts as the system-wide search page for fluid components.</p> <p>Clone the delivered fluid search page, PTS_CFGSRCH_FL, to create a custom search page. However, you may create your own search page without cloning the delivered search page.</p>
Max Recent Searches	<p>Specify the maximum number of recent searches that can be stored for a component.</p> <p>You can choose between 0 and 999.</p> <p>The default value is 5.</p> <hr/> <p>Note: The value entered here is overridden when you specify a value in the maximum recent searches field for a specific content reference on the Search Configuration page.</p> <hr/>
Default rows in search results	<p>Specify the number of search items that should be displayed to an end user.</p> <p>This value is applicable to classic and fluid components and to standard and keyword search types.</p> <p>The default value is 300.</p>

Field or Control	Description
Max Saved Searches	<p>Specify the maximum number of searches that can be stored for a component.</p> <p>You can choose between 0 and 999.</p> <p>The default value is 5.</p> <hr/> <p>Note: The value entered here is overridden when you specify a value in the maximum saved searches field for a specific content reference on the Search Configuration page.</p> <hr/>
Recent/Saved Search Format grid	<p>Options:</p> <ul style="list-style-type: none"> • Label value • Label operator value

Managing Configurable Search Configurations

Use the Manage Configurable Search page to list configured content references, to add a new configuration, and to review, edit, or delete an existing configuration.

Access the Manage Configurable Search page by selecting **PeopleTools >Utilities >Manage Configurable Search**.

Note: Manage Configurable Search page requires the PTPORTALREGISTRY search definition, so you must ensure that the search definition is deployed and indexed.

This example illustrates the Manage Configurable Search page. Descriptions of the fields appear after the example.

New Content References

The New Content References section enables you to search for content references that you want to configure for configurable search. When you select a content reference, the Search Configuration page displays the details of the selected content reference if the following conditions are met:

- Add, or Update/Display, or both action modes are enabled on the Use tab of the Component Properties dialog box in Application Designer.

- Search record contains search keys.

See "Setting Component Properties" (Application Designer Developer's Guide).

If any of these conditions are not met, the system does not display the Search Configuration page, but displays the following error message: Search and Add modes are not enabled for component, cannot configure the content reference.

Configured Searches

The Configured Searches section lists the content references that have been configured for configurable search. You may edit the configuration or delete the configuration. The list of configured content references could be quite long, so you can use the filter option to identify the content reference you require.

Configuring a Search Page

Use the Search Configuration page to configure a search page and specify how the configured search page should appear for an end user to interact with when accessing the component.

Access the Search Configuration page by selecting **PeopleTools >Utilities >Manage Configurable Search** and then selecting the content reference for which you want to configure a search page.

This example illustrates the Search Configuration page. Descriptions of the fields appear after the example.

Search Configuration

Preview Save

Portal Name: EMPLOYEE *Status: Active

Configurations < 1 of 1 > + -

Search Add New

*Effective Date: 08/16/2022

Main Options

*Search Type: Standard Search

*Search Record: PSOPRDEFN_SRCH

Search as default view: Yes

Search Page Heading:

Search Page Instructions:

Advanced Options

*Configured Search Presentation: Basic and Advanced
For Configured Searches, Basic is presented as "Fewer Options" and Advanced is presented as "More Options"

*Default Search Presentation: Basic

*Columnar Display: Single Column

Maximum Number of Recent Searches: Maximum: 999

Recent Search Format: Field label with value Field label with value and operator

Maximum Number of Saved Searches: Maximum: 999

Saved Search Format: User supplied name and field label with value User supplied name, field label with value and operator

Maximum Number of Search results:

Execute Auto Search: No Yes
Auto search shows pre-search results at launch

Collapse Search Criteria: No Yes
Once set, search criteria will be collapsed only when valid result is fetched

Show Related Actions in Search Results: No Yes

Small Form Factor Options

*Default Search Presentation: Basic
Basic search on small form factor allows you to see one field and search by one field

Image Options

Include Image in Search Results: Yes

Source Record: QE_EMPL_PHOTO

Source Field Name: QE_EMPLOYEE_PHOTO

Search Criteria and Result Options

Load Fields 2 rows

Search Criteria	Result Fields	Show All								
Field Name	Field Label	Modify Field Label	Required	Hidden	Display Only	Default Value	Include in Criteria	Criteria Sequence	Field Appearance in	
1 OPRID	User ID	<input type="checkbox"/> No <input checked="" type="checkbox"/> Yes	<input checked="" type="checkbox"/> Yes <input type="checkbox"/> No	<input type="checkbox"/> No <input checked="" type="checkbox"/> Yes	<input type="checkbox"/> No <input checked="" type="checkbox"/> Yes		<input checked="" type="checkbox"/> Yes <input type="checkbox"/> No	1	Both Basic & Advanced	
2 OPRDEFDESC	Description	<input type="checkbox"/> No <input checked="" type="checkbox"/> Yes	<input type="checkbox"/> No <input checked="" type="checkbox"/> Yes	<input type="checkbox"/> No <input checked="" type="checkbox"/> Yes	<input type="checkbox"/> No <input checked="" type="checkbox"/> Yes		<input checked="" type="checkbox"/> Yes <input type="checkbox"/> No	2	Advanced Search	

Configure Links 1 row

PeopleTools > Security > Permissions and Roles > Permission Lists	Order of Appearance	Open in New Window	*Link Label	Delete
PeopleTools > Security > Permissions and Roles > Permission Lists	1	<input checked="" type="checkbox"/> Yes <input type="checkbox"/> No	Permission Lists	<input type="button" value="Delete"/>

Add Link

Field or Control	Description
Status	<p>If you select Active, the search configuration is applied to the content reference and the configured search page is displayed based on conditions, which are described in the guidelines topic. See Guidelines on Using Configurable Search for Classic Components and Guidelines on Using Configurable Search for Fluid Components.</p> <p>If you select Inactive, the search configuration is not applied to the content reference and the system displays the search page specified in the component.</p>

Field or Control	Description
Preview	<p>Use the Preview button to view the configured search page. It opens a new window to display the newly configured search page.</p> <p>In the case of multiple effective-dated configurations for a content reference, selecting the Preview button displays the current effective-dated configuration.</p>
Effective Date	<p>PeopleSoft's effective-dating logic enables you to configure an effective-dated configuration for a future date. You can configure multiple effective-dated configurations for a content reference and use the effective-dating logic to enable the configuration on a future date. For example, you could configure two configurations for a content reference. While one configuration is based on keyword search and is set for a future date, the other configuration can be based on standard search and set for current effective date. Or you could set any of the configurable options and specify a future date for the configuration.</p> <p>Use the plus sign (+) to create a new search configuration for the content reference to which you can apply effective-dating logic.</p> <p>See "Using Effective Dates" (Applications User's Guide).</p>
Search or Add New	<p>The Search Configuration page displays the Search or Add New, or both tabs based on the configuration of the selected component. That is, if the Update/Display action mode is selected on the Component Properties page, the Search tab appears and if the Add action mode is selected, the Add New tab appears.</p> <p>While the Search tab displays the Search Configuration page, the Add New tab displays the Add Configuration page. Search is the default tab.</p>

Main Options

Field or Control	Description
Search Type	<p>The system displays the search type configured for the component.</p> <p>The following options are available:</p> <ul style="list-style-type: none"> • Keyword Search: Uses a search definition to perform a search. • Standard Search: Uses a search record to perform a search. <p>If you select a content reference that has a component mapped to a search definition and real time indexing is enabled, then the default search type is Keyword Search. However, you can change the search type to Standard Search and specify a search record for it.</p>
Search Record or Search Definition	<p>The Search Record or Search Definition field appears based on the search type.</p> <p>The Search Record field displays the search record specified on the Component Properties page. However, you may specify a different search record for the content reference. In this case, the search record specified on the Search Configuration page is used when a user accesses this content reference provided the status of this configuration is set to active. If the status is set to inactive, the search record specified on the Component Properties page is used.</p> <p>The Search Definition field displays the search definition mapped to the component. Similarly, if a search definition mapped to a component is enabled for real time indexing, this field displays the search definition, which is enabled for real time indexing. If a search definition is not mapped to a component, the field is blank.</p> <p>You may specify a different search definition for the content reference, but you should ensure that the search definition is mapped to the component. If you select an unmapped search definition, the following error message displays: Search definition is not mapped to component.</p> <p>See Mapping Components to Search Definitions.</p>

<i>Field or Control</i>	<i>Description</i>
Search as default view	<p>Use this option to specify whether search should be the default view. If you set Yes, search will be the default view and Add as the default view is automatically set to No.</p> <p>When a component is configured for Search and Add action modes, you can set only one action mode as the default view.</p>
Search Page Heading	<p>Specify a heading that will appear at the top of the search page. The heading should not exceed 100 characters.</p> <p>If you do not specify a page heading, the page displays the default heading, Find an Existing Value.</p>
Search Page Instructions	<p>Enter instructions for an end-user, for example information on how to complete the process.</p>

Advanced Options

<i>Field or Control</i>	<i>Description</i>
Configured Search Presentation	<p>Valid values for search lookup type:</p> <ul style="list-style-type: none"> • Advanced • Basic • Basic and Advanced <p>Default search lookup type is Basic and Advanced.</p> <p>For descriptions of search lookup types, see "Understanding Search Pages" (Applications User's Guide).</p>
Default Search Presentation	<p>This field appears when you select the Basic and Advanced option in the Configured Search Presentation field.</p> <p>Valid values for default search presentation:</p> <ul style="list-style-type: none"> • Advanced • Basic <p>See "Understanding Search Pages" (Applications User's Guide).</p>

Field or Control	Description
Columnar Display	Specify whether search criteria should be displayed as a single or double column.
Maximum Number of Recent Searches	<p>This field displays the value specified on the PeopleTools Options page.</p> <p>The default value is 5.</p> <p>However, you may specify the number of recent searches that should appear in the Recent Search Display area. Valid values are 0 to 999.</p> <p>The value entered here is applicable to the selected component and overrides the system-wide configurable search property.</p>
Recent Search Format	<p>Select a display option for the recent search:</p> <ul style="list-style-type: none"> Field label with value - Displays the name of the field and the value entered in the field. Field label with value and operator - Displays the name of the field along with the value entered in the field and the search operator.
Maximum Number of Saved Searches	<p>This field displays the value specified on the PeopleTools Options page.</p> <p>The default value is 5.</p> <p>However, you may specify the number of saved searches that should appear in the Saved Searches area. Valid values are 0 to 999.</p> <p>The value entered here is applicable to the selected component and overrides the system-wide configurable search property.</p>
Saved Search Format	<p>Select a display option for a saved search:</p> <ul style="list-style-type: none"> User supplied name and field label with value - Displays the name of the saved search along with field name and value entered in the field. User supplied name, field label with value and operator - Displays the name of the saved search and field name and the value entered in the field along with the search operator.

Field or Control	Description
Maximum Number of Search Results	<p>This field displays the value specified on the PeopleTools Options page.</p> <p>The recommended values for keyword search and standard search are 50 and 300, respectively.</p> <p>The default is 300.</p> <p>Specify the maximum number of search results to display. The value entered here is applicable to the selected component and overrides the system-wide configurable search property.</p>
Execute Auto Search	<p>Select Yes to display pre-search results based on the most recent search criteria.</p> <p>For initial auto search results, default search criteria is used. Subsequently, the most recent search criteria is used.</p>
Collapse Search Criteria	<p>Select Yes to collapse the search criteria fields. However, search criteria is collapsed after search results are fetched.</p>
Show Related Actions in Search Results	<p>By default, related actions are not displayed in search results even if related actions are configured for the component. To display related actions, you should ensure that related actions are configured for the component and set this field to Yes.</p>

Small Form Factor Options

Field or Control	Description
Default Search Presentation	<p>Specify the default lookup search for small form factor devices.</p>

Image Options

Field or Control	Description
Include Image in Search Results	<p>Specify whether to display an image with the search results.</p> <p>When you select Yes, two fields are displayed to provide the record and field information.</p>
Source Record	<p>The source record keys should be keys in the search record.</p>

Field or Control	Description
Source Field Name	Select a field of the record.

Search Criteria and Result Options

Select the Load Fields button to display the fields of the configured search record (standard search) or attributes of the configured search definition (keyword search).

The Search Criteria tab enables you to configure the appearance of fields or attributes and the sequence in which they should appear on the search page.

Field or Control	Description
Field Name or Attribute Name	Lists the fields of the configured search record for standard search or the attributes of the configured search definition for keyword search.
Field Label or Attribute Label	Lists the labels of the fields or attributes.
Modify Field Label or Modify Attribute Label	Set this option to Yes, which enables you to edit the field or attribute label.
Required	<p>Specify whether the field or attribute should be marked as a required field on the search page.</p> <p>The Required button for a corresponding record field is <i>not</i> enabled if the Required check box (on the Edits tab) and the Search Edit check box (on the Use tab) are selected in the Record Field Properties page in Application Designer.</p> <p>However, the Required button is enabled if the Required check box or the Search Edit check box is <i>not</i> selected in the Record Field Properties page.</p> <p>If you select the Required button, the Include in Criteria button is automatically selected if it was not previously selected. If you deselect the Required button, the Include in Criteria button is not automatically deselected.</p>
Hidden	<p>Specify whether the field or attribute should be hidden on the search page.</p> <p>You can set a field or attribute as hidden if the field or attributes has a default value associated with it.</p>

Field or Control	Description
Display Only	Specify whether the field or attribute should be display only. You can set a field or attribute as display only if the field or attribute has a default value associated with it.
Default Value	Lists the default value for a field or attribute if configured.
Include in Criteria	Specify whether the field or attribute should be used as search criteria on the search page. The Include in Criteria button corresponding to a record field is <i>not</i> enabled if the Required check box (on the Edits tab) and the Search Edit check box (on the Use tab) are selected in the Record Field Properties page in Application Designer. However, the Include in Criteria button is enabled if the Required or the Search Edit check box is <i>not</i> selected in the Record Field Properties page. If you deselect the Include in Criteria button, the Required button is automatically deselected if it was previously selected.
Criteria Sequence	Specify the sequence in which the fields or attributes should be displayed on the search page. A field or attribute with the lowest numeric value is displayed at the top.
Field Appearance in or Attribute Appearance in	Specify the search presentation for each field or attribute. For standard search, the default search presentation is Basic and Advanced.

The Result Fields tab enables you to configure the appearance of fields or attributes and the sequence in which they should appear on the search results page.

Field or Control	Description
Include in Results	Specify whether the field or attribute should be displayed in search results.
Include in Small FF	Specify whether the field or attribute should be displayed in search results on small form factor devices.

<i>Field or Control</i>	<i>Description</i>
Results Sequence	Specify the order in which the fields or attributes should appear on search results. The field or attribute with the lowest numeric value appears from left to right.

The Show All tab displays a combined list of configurable fields from the Search Criteria and Result Fields tabs.

Configure Links

Select the Add Link button to launch the Configure a link dialog where you can specify a content reference that will appear on the search page. This functionality enables you to provide links to other components that are relevant to the current component or the next task a user needs to perform.

<i>Field or Control</i>	<i>Description</i>
Order of Appearance	Specify the sequence in which the configured links should appear when you configure multiple links on the search page.
Open in New Window	Specify whether the configured content reference should appear in a new window. If you select No, the content reference opens in the current window.
Link Label	Enter the label for the content reference.

After completing the configuration, save the changes. Use the Preview button to preview the new page, which opens a new window to display the newly configured page.

Configuring an Add a New Value Page

Use the Add Configuration page to configure a page to add a new value to the PeopleSoft table and specify how the page should appear for an end user.

Click Add New on the Search Configuration page.

This example illustrates the Add Configuration page. Descriptions of the fields appear after the example.

Add Configuration
Save

Portal Name EMPLOYEE
Status Active
User Profiles

Configurations
Search Add New

Effective Date 09/09/2022

▼ Main Options

* Search Record

Add as Default View No

Add Page Heading

Add Page Instructions

▼ Configure Criteria
2 rows

Load Fields

Field Name	Field Label	Modify Field Label	Required	Hidden	Display Only	Included in Criteria	Sequence	Default Value
1 OPRID	User ID	<input type="checkbox"/> No	<input checked="" type="checkbox"/> Yes	<input type="checkbox"/> No	<input type="checkbox"/> No	<input checked="" type="checkbox"/> Yes	<input type="text" value="1"/>	
2 OPRDEFNDESC	Description	<input type="checkbox"/> No	<input type="checkbox"/> No	<input checked="" type="checkbox"/> No	<input checked="" type="checkbox"/> No	<input type="checkbox"/> No		

▼ Add button override

Package Name

Path

Class Name

Method

▼ Configure Links
1 row

PeopleTools > Security > Permissions and Roles > Permission Lists	Order of Appearance	Open in New Window	*Link Label	Delete
PeopleTools > Security > Permissions and Roles > Permission Lists	<input type="text" value="1"/>	<input checked="" type="checkbox"/> Yes	Permission Lists	

Add Link

<i>Field or Control</i>	<i>Description</i>
Status	The status specified on the Search Configuration page is extended to the Add Configuration page. Therefore, you cannot set or edit the status on this page.

Main Options

<i>Field or Control</i>	<i>Description</i>
Search Record	<p>The Search Record field displays the search record specified on the Component Properties page. However, you may specify a different search record for the content reference. In this case, the search record specified on the Search Configuration page is used when a user accesses this content reference provided the status of this configuration is set to active. If the status is set to inactive, the search record specified on the Component Properties page is used.</p> <p>If the search record does not contain a key or keys, the Load Fields button is disabled. Choosing a search record that contains a key or keys enables the Load Fields button.</p>

For descriptions of the fields in the Configure Criteria and Configure Links areas, which are similar to the Search Configuration page, see [Configuring a Search Page](#).

Configure Criteria

The Configure Criteria group box enables you to configure the appearance of fields or attributes and the sequence in which they should appear on the Add a New Value page.

Field or Control	Description
Field Name	Lists the fields of the configured search record for standard search or the attributes of the configured search definition for keyword search.
Field Label	Displays the labels of the fields or attributes.
Modify Field Label	Set this option to Yes, which enables you to edit the field or attribute label.
Required	<p>Specify whether the field or attribute should be marked as a required field on the Add a New Value page.</p> <p>The Required button for a corresponding record field is <i>not</i> enabled if the Required check box (on the Edits tab) and the Search Key are selected in the Record Field Properties page in Application Designer.</p> <p>However, the Required button is enabled if the Required check box or the Search Key is <i>not</i> selected in the Record Field Properties page.</p> <p>If you select the Required button, the Include in Criteria button is automatically selected if it was not previously selected. If you deselect the Required button, the Include in Criteria button is not automatically deselected.</p>
Hidden	<p>Specify whether the field or attribute should be hidden on the Add a New Value page.</p> <p>You can set a field or attribute as hidden only if the field or attribute has a default value associated with it.</p>
Display Only	<p>Specify whether the field or attribute should be display only.</p> <p>You can set a field or attribute as display only if the field or attribute has a default value associated with it.</p>

Field or Control	Description
Include in Criteria	<p>Specify whether the field or attribute should be used as search criteria on the Add a New Value page.</p> <p>The Include in Criteria button corresponding to a record field is <i>not</i> enabled if the Required check box (on the Edits tab) and the Search Key check box (on the Use tab) are selected in the Record Field Properties page in Application Designer.</p> <p>However, the Include in Criteria button is enabled if the Required check box or the Search Key check box is <i>not</i> selected in the Record Field Properties page.</p> <p>If you deselect the Include in Criteria button, the Required button is automatically deselected if it was previously selected.</p>
Sequence	Specify the sequence in which the fields or attributes should be displayed on the Add a New Value page.
Default Value	Lists the default value for a field or attribute if configured.

Add Button Override

Use the Add button override area to override the behavior of the Add button. Specify an application class or method for the target component where you want to redirect the user when they select the Add button.

Field or Control	Description
Package Name	Specify the root package where the application class is created.
Path	Specify the path, represented by a colon (:).
Class Name	Choose the relevant application class.
Method	Specify the method to be invoked.

Configure Links

Select the Add Link button to launch the Configure a link dialog where you can specify a content reference that will appear on the Add a New Value page. This functionality enables you to provide links to other components that are relevant to the current component or the next task a user needs to perform.

Field or Control	Description
Order of Appearance	Specify the sequence in which the configured links should appear when you configure multiple links on the search page.
Open in New Window	Specify whether the configured content reference should appear in a new window. If you select No, the content reference opens in the current window.
Link Label	Enter the label for the content reference.

After completing the configuration, save the changes. Use the Preview button to preview the new page, which opens a new window to display the newly configured page.

Upgrade Considerations

After upgrading to PeopleTools 8.60, you should be aware of the changes on search pages that will occur and follow the guidelines to apply configurable search to components.

- On classic components:
 - The existing PeopleTools-delivered component search pages are dynamically replaced with the new classic search page delivered in PeopleTools 8.60.
 - The existing custom search pages with the Standard search page type are not replaced.

For guidelines on using configurable search in classic components, see [Guidelines on Using Configurable Search for Classic Components](#).

- On fluid components:
 - By default, on existing fluid components, the search pages are not replaced with the new fluid search page delivered in PeopleTools 8.60. You should explicitly enable configurable search for fluid components.
 - The existing custom search pages are not replaced.
 - By default, new fluid components use the new fluid search page delivered in PeopleTools 8.60.
 - On a component definition (Fluid tab), you can use the Enable Configurable Search check box to enable or disable configurable search, see [Guidelines on Using Configurable Search for Fluid Components](#), Using the Enable Configurable Search Check Box for Fluid Components.

For guidelines on using configurable search in fluid components, see [Guidelines on Using Configurable Search for Fluid Components](#).

- PeopleTools 8.60 supports components that use the HCM delivered configurable search. If you want to use the PeopleTools delivered configurable search for a component (or content reference), you should deactivate the HCM delivered configurable search for that component:

1. Set the status of search configuration for the component or content reference in the HCM application to inactive.
2. Navigate to **PeopleTools >Portal >Event Mapping >Event Mapping Configuration**, select the content reference (component), and then remove the event mapping for SearchInit PeopleCode.

When you perform step 1, step 2 may not be necessary. However, it is recommended that administrators ensure that SearchInit PeopleCode for that component is removed by performing step 2.

Migrating Search Page Configuration Data Using Application Data Sets

Oracle delivers an Application Data Sets (ADS) definition, PTS_CFGSRCH, to copy or migrate configurable search configuration data from one database to another.

For information on using ADS, see "Understanding ADS Projects" (Lifecycle Management Guide).

The following table lists the ADS definitions:

ADS Definition	Description
PS_PTS_SRCH_ASGN	You can query this table to check whether configuration exists for a particular CREF or not; returns value 'A' if active configuration exists, otherwise returns 'I'.
PS_PTS_SRCH_DEFN	Common table that stores some of the information about search and add configurations. You can query this table to get information about configured search record/definition, add search record, page title and instructions, maximum number of recent and saved searches, display view options etc.
PS_PTS_SRCH_DTL_S PS_PTS_SRCHDTL_S_L	Stores search criteria fields for standard search.
PS_PTS_ADD_DTL PS_PTS_ADD_DTL_L	Stores information about add page fields and their configuration.
PS_PTS_SRCH_DTL_E PS_PTS_SRCHDTL_E_L	Stores search criteria fields for keyword search.

ADS Definition	Description
PS_PTS_CFG_SRCH	Stores information about configured portal object name and associated CREF label.
PS_PTS_LINK_DTL_A PS_PTS_LINK_DTL_S PS_PTS_LINKDTL_A_L PS_PTS_LINKDTL_A_S	Stores links configured for particular CREF on the Add or Search page.
PS_PTS_SRCH_DEF_L	Related language table for page title and page instructions.

Working with PeopleSoft Search

Understanding PeopleSoft Search

PeopleSoft Search refers to the set of features that PeopleSoft application end users access to submit search requests against a search engine. The PeopleSoft Search interfaces depend on:

- the Search Framework being completely configured.
- the PeopleSoft search definitions and categories deployed to the search engine instance.
- the search indexes are populated.

These are the PeopleSoft Search features:

Feature	Description
Global Search	Provides users a Global Search bar appearing in the header of the PeopleSoft browser session where users can select relevant search categories against which to run searches.
Search Pages	<p>When a user accesses a component, they use Search Pages to enter search keys to isolate the desired rows of data. Search Pages come in these formats: Find an Existing Value Search (or Standard Search) and Keyword Search.</p> <p>Find an Existing Value Search refers to the traditional component search, where users enter level 0 search keys running directly against the database on the Find an Existing Value tab, which appears regardless if Search Framework is configured.</p> <p>Keyword Search involves the use of the Keyword tab, which enables more free-form searches against a search engine index.</p> <p>In fluid user interface, keyword search and master-detail search is available, which can be configured as tiles to launch search in a component. While keyword search is against a search engine index, the master-detail search is configured in Pivot Grid.</p>

PeopleSoft Search Framework features that are key to enhancing user experience are:

Feature	Description
Facets	Facets filter and narrow search results by ordering search results by categories, which are the fields you have selected to be facets in the search definition. These are similar to Google's left-hand, navigational search facets, for example.
Search Groups	Search groups enable you to specify search categories to become available for display in the Global Search Bar. (Applies only to Global Search).
Search Contexts	Search contexts enable you to configure when a search group is available in the Global Search Bar. PeopleSoft Search Framework supports only the Homepage context. (Applies only to Global Search).

Working with Search Pages

This section provides an overview of working with search pages.

Understanding Search Pages

Configurable search enables you to configure the presentation of the delivered search pages for classic and fluid components. The search pages for PeopleSoft application components (classic and fluid) provide these options:

- Find an Existing Value (Standard search).
- Add a New Value.
- Keyword Search.

Note: Configurable search does not support a combination of standard search and keyword search for a component. For a list of supported combinations, see [Considerations for Using Configurable Search](#).

The PeopleTools delivered classic search page is displayed for classic components, by default. However, for fluid components, you need to specifically enable the delivered fluid search page for individual components or system-wide.

See [Understanding Configurable Search](#).

This example illustrates the fields and controls on the Classic PeopleSoft Application Component Search. You can find definitions for the fields and controls later on this page.

Departments

Find an Existing Value

▼ **Search Criteria**

Enter any information you have and click Search. Leave fields blank for a list of all values.

⊕ Add a New Value

🕒 Recent Searches

Choose from recent searches ▼ ✎

🔖 Saved Searches


Choose from saved searches ▼ ✎

Search by: Set ID ▼ = SHARE

▼ Show more options

Case Sensitive Include History Correct History

Search
Clear



Nothing yet
Your search results will appear here

While the Find an Existing Value page enables end users to find the rows of data they need to view or modify, it is limited to running searches only across level 0 fields. The Find an Existing Value search runs against the live data tables, the component search record, not a search engine index.

Using the Keyword search page, end users can perform a deeper, more free-form search against the data. The Keyword search runs across the indexed values stored on the search engine server, which contains fields across all scroll levels.

Note: The Keyword Search page only appears if a search index has been associated with the component in a search definition. Because the keyword search runs against the indexed data, depending on the frequency of index builds, the indexed data may not entirely match the Find an Existing Value search results. You can ensure that the indexed data is near real time by enabling real-time indexing for the mapped search definition.

For information on real-time indexing, see [Configuring Real Time Indexing](#).

This example illustrates the fields and controls on the Keyword search page. You can find definitions for the fields and controls later on this page.

Job Data

Find an Existing Value

▼ **Search Criteria**


Note: Search will return results last updated 5 hours ago.
Enter any information you have and click Search. Leave fields blank for a list of all values.

Keywords

Search

i

▼ Show more options



Nothing yet
Your search results will appear here

The search pages that display after you've enabled Keyword search for a component are:

- **Keyword Search:** This page enables the user to run free-text searches using the Keywords edit box. Implementation teams can index any attributes that they would like the end user to be able to search.
- **Add a New Value:** This is the same interface as the Add a New Value page used in previous PeopleSoft releases, which enables you to add new rows to PeopleSoft tables.

Enabling Keyword Search

To enable Keyword Search display when users access a component, you use the Component Mapping page in the Search Definition component for the Search Framework Designer. The Component Mapping page enables you to map a component to a specific search definition to enable search engine integration with Search Pages.

Note: Search definitions mapped to a component use document level security. This is enforced by the Search Designer activity guide.

For components that have been added to the grid on the Component Mapping page for a search definition, the search page for that component displays the Search Framework search options. The keyword Search Framework feature enables the user to search deeper (level 0-3) by running a free-formed search against crawled and indexed values in the search engine. The keyword search is powerful because it allows you to search across all fields in an index rather than just inside a specific set of fields defined in the search record.

When the search definition is deployed, the system updates the component meta data to indicate the search category to which the component is mapped. When the system renders the page, it displays the keyword search page for that component.

Note: Multiple components can be mapped to the same search definition.

To enable search pages for a component:

1. Select PeopleTools, Search Framework, Search Designer Activity Guide, Search Definition, and open the desired search definition.
2. Select the Component Mapping tab.
3. Add the desired component to the grid.
4. Update the security restrictions and access appropriately.
5. Save the search definition.
6. Deploy or update the search definition to the search engine.
7. Navigate to the component and verify that you see the Keyword search page.

Related Links

[Mapping Components to Search Definitions](#)

[Setting Document Level Security](#)

Developing for Search Pages

This section describes these items to consider while implementing Search Pages.

- Setting component properties for search pages.
- Working with Find an Existing Value and Keyword Search behavior differences.
- Working with Keyword Search data currency.
- Requirements for Keyword Search.

Setting Component Properties for Classic Search Pages

You configure properties for your search pages in the Component Properties dialog box in PeopleSoft Application Designer. On the Internet tab, in the Search Page group box there are numerous settings related to your application search pages. These settings apply specifically to the Keyword search page:

Field or Control	Description
Primary Action (Keyword Search)	Enables you to set the primary action of the component to display the Keyword Search page.
Default Search\Lookup Type (Keyword Only)	Enables you to set the default look up type to be used on the Keyword Search page.
Link to Realtime Search Page	Specifies the message set and number that stores the text displayed for the Find an Existing Value link in the Add a New Value tab and Keyword Search tab <i>if</i> Keyword Search is enabled.

Field or Control	Description
Link to Keyword Search Page	Specifies the message set and number that stores the text displayed on the Add a New Value tab and Find an Existing Value tab <i>if</i> Keyword Search is enabled.

Note: The Link to Keyword Search Page value only applies if Keyword Search is enabled for that component.

The Link to Realtime Search Page value applies to all components that use the PeopleTools search page, even if Keyword Search is *not* enabled.

Note: The term, *Realtime*, is used to differentiate between the currency of the data between the Find and Existing Value search and the Keyword search. While the Find an Existing Value search displays search results that are always a current reflection of the data in the database, the Keyword Search is only as current as the last index build or incremental index update. However, you can ensure that the indexed data is near real time by enabling real-time indexing for the component.

When implementing the Keyword search feature, make sure to adjust these settings accordingly. The general concepts surrounding these settings, along with other component properties, are discussed in the PeopleSoft Application Designer documentation.

See "Setting Internet Properties" (Application Designer Developer's Guide).

Setting Component Properties for Fluid Search Pages

Similar to classic search pages, you configure properties for your Fluid search pages in the Component Properties dialog box in PeopleSoft Application Designer. On the Fluid tab, in the Search Type field, select either None, Standard, or Master/Detail.

See "Implementing Search for Fluid Components" (Fluid User Interface Developer's Guide).

Working with Find an Existing Value and Keyword Search Behavior

Find an Existing Value or Standard search and Keyword search types have the same, general look and feel, however, these differences and features should be noted:

- With the Find an Existing Value search, the end user selects values based on the underlying record search keys and clicks Search. Results are rendered in a standard grid format, as seen in all previous PeopleSoft versions.
- With the Keyword search, the end user enters values into the Keywords edit box and clicks Search. The end user can also enter a combination of keywords in addition to using the provided search keys. The system runs this search against the search engine index. Results can be displayed in either a grid or a list format.
- The Keyword search will show facets, while the Find an Existing Value search cannot.
- Search operators are different between the Find and Existing Value search and Keyword search. For example, the typical Keyword search operators include *contains*, *=*, *not =*, where as the Find an Existing Value search operators include more, such as *begins with*, *contains phrase*, *=*, *not =*, *greater than*, *less than*, *between*, *in*, and so on.

Note: When you use the contains phrase operator, enter a complete word or a set of words, for example desktop or desktop computer. You can also use wild card characters; the wild card characters can be used at the beginning of a word, within a word or at the end of a word. Examples of valid search text containing wild cards: *test, t*st, test*.

- Keyword search fields do not contain prompts or drop-down lists. Search text needs to be added in a free-form manner, using wild cards as needed.

Note: In the Find an Existing Value tab you specify a wild card using "%", while in Keyword search you specify a "*".

Working with Keyword Search Data Currency

The Keyword search can display only search results that are included in the index. Search Framework provides the capability to enable real-time indexing for search definitions.

In search definitions that are *not* enabled for real-time indexing, the index is only as current as the last incremental index update or the last full build date and time. As such, the Keyword search results will not always contain the most current reflection of the transactional data.

For example, assume that an organization has the incremental index update process schedule to run once every 3 days. If a salesperson adds a new customer to the system just after the incremental update occurred, that customer information will not appear in the Keyword Search results until the next incremental index update process has completed. Likewise, if a customer has been deleted from the system just after an incremental index update, the search results for the Keyword Search will continue to include that customer information until the next incremental index update process has completed.

The Keyword search page (real-time indexing is *not* enabled) displays a message indicating how fresh the search results are. For example:

Note: Search will return results last updated 5 hours ago.

If you enable real-time indexing for a search definition, the data is indexed every time an update occurs on the search definition providing you near real-time data.

The Keyword search page (real-time indexing is enabled) does not display any message to indicate how fresh the search results are because the search results are near real time.

See [Understanding Real Time Indexing](#).

Requirements for Keyword Search

The requirements for keyword search slightly differ depending on whether you are using Application Designer or Manage Configurable Search page to set keyword search for a component.

When you use Application Designer to set keyword search for a component, ensure that the following settings are met:

- All the search keys shown on the search page must be indexed attributes.
- Only the fields marked as List Box items on the component search record will be shown.

Note: This is true when showing the results in the grid format for the keyword search. If you switch to list format, you can display additional fields that are indexed but not marked as List Box or search keys.

- All List Box item fields must be indexed attributes.

Note: Grid-based view of the keyword search results is based on the List Box item configuration on the associated search record.

When you use the Manage Configurable Search page to set keyword search for a component, ensure that the following settings are met:

- All the requirements listed for keyword search through Application Designer.
- Fields selected for search criteria and results must be indexed attributes.

See [Working with Configurable Search](#).

Working with Keyword Search Modes

The Keyword search page has these modes:

- Keyword-only
- Basic
- Advanced

These modes can be selected in the Component Properties dialog box.

Working with Keyword-Only Mode

The keyword-only mode is not supported. If keyword-only mode is selected for a component and if configurable search is enabled for the component, then keyword-only mode is not honored. Instead, basic mode is displayed.

See [Considerations for Using Configurable Search](#).

Working with Basic Mode

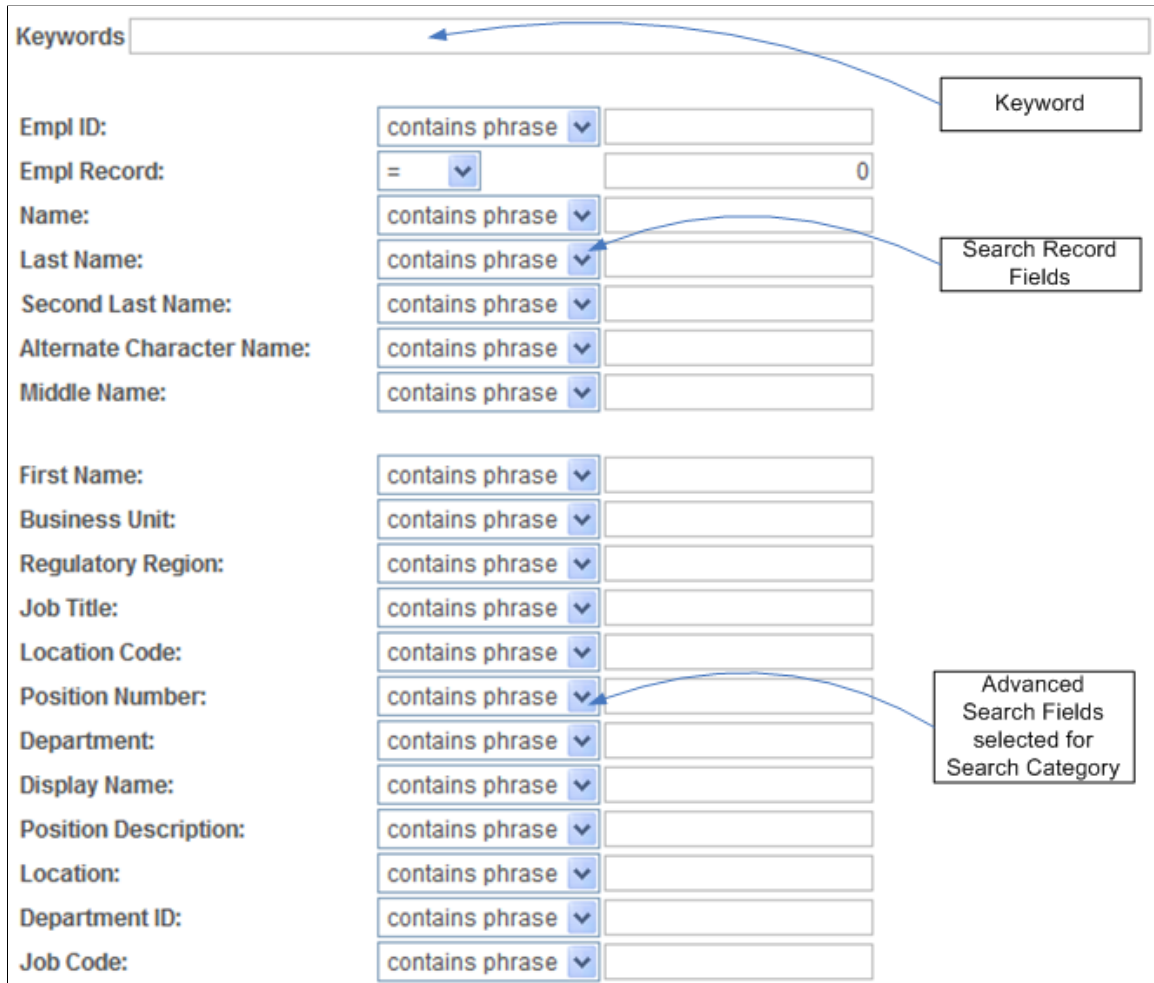
In basic mode, the Keywords edit box displays with an additional 'Search by' option with a drop-down list box to select a search record field to use for as search criteria.

Working with Advanced Mode

The advanced mode displays all search criteria:

- Keywords.
- Search record fields.
- Search fields defined for the Search Framework search category on the Advanced Search Fields tab.

The following illustration depicts the controls on the advanced search page and their association.



Note: The system does not display in the Advanced Search Fields list any search fields that are also displayed in the search record search field list.

Note: All search record fields (Key, Alt Key, and List box fields) must be part of the index. If any of the fields have translate values, those translate fields must also be part of the search index.

Note: The search criteria that appears depends on the application’s usage of the search filter PeopleCode and the type of data. For example, contains phrase is applicable to a character field, but not to a number field.

Working with Global Search

This section provides an overview of working with global search.

Understanding Global Search

Global Search provides a way for a user to search across all search indexes or a specific group of search indexes. End users submit a Global Search using the Global Search Bar, which is available in the Portal header throughout the user's session irrespective of the content the user is accessing in the target frame. The Global Search feature allows the user to search and drill down to a specific row of data or transaction from the search results, all without navigating through a menu structure to the component search page and entering search criteria there. In an Interaction Hub environment, Global Search can search across indexes from multiple content systems.

Note: From PeopleTools 8.55, the fluid global search page is used irrespective of the default theme and style the application is using. Global search does not support persistent search.

Global Search is enabled by default based on three conditions:

- At least one search index must be deployed and crawled on the system.
- A user accessing global search must have the required permissions for the deployed search index.
- A user accessing global search must have access to the Fluid Global Search component.


On homepages, the Global Search bar is presented in its expanded form:

This example illustrates the expanded form of the Global Search bar.



On transaction pages and dashboards, global search is presented in its muted form:

This example illustrates the muted form of the Global Search bar.



The Global Search Bar consists of:

- drop-down list for selecting a specific search group, containing the desired indexes.
- free text search edit box.

Related Links

[Working with Search Pages](#)

Working with Search Groups and Search Contexts

When implementing Global Search, it is highly recommended to employ search contexts and search groups. These constructs work together to refine the user interaction with the Global Search Bar drop-down list.

A search group is a type of search category that is exposed to Global Search to help manage search definitions deployed to the search engine. Search groups provide a way to show relevant categories based

on the context the user is in on the target frame. A search definition can belong to multiple search groups. You manage search groups in the Search Framework Administration Activity Guide. In an Interaction Hub installation, for example, search groups can be formed to include search definitions deployed across content systems.

Using search contexts, you can configure search groups such that based on the context in the target frame, only a subset of all possible search groups are available to the user in the Global Search Bar. For example, if the end user is accessing content from a PeopleSoft Financial application, then only search groups relevant to that application will be available. Likewise if a user is accessing content from a PeopleSoft Human Capital Management application, then only search groups related to that application will be available.

At any given time, end users have these types of search groups available:

- **All:** This search group will always be available and is a group of search categories that span across all the indexes that have been deployed and made searchable for a PeopleSoft application.
- **Component Specific:** This type of search group will match the component the user is currently viewing in the target frame. For example, if a user is viewing an expense report, then the Expense Report component would be the search group that is available, if the component is indexed, and it does not apply already to the Home Page context.

Note: Only Home Page search context can be used. Template level and Node level search contexts are no longer valid. If a search context is not defined for Home Page, then only the All search group is displayed. In an Interaction Hub installation, Portal's Home Page search context is used always.

Global Search shows search groups in the Global Search Bar based on a user's role. To set up permissions to search groups, you use PeopleTools Security permission lists. Setting up role-based search group access is discussed in the Security topics in this document.

Related Links

[Setting Up Role-Based Search Group Access](#)

[Managing Search Context](#)

[Specifying General Search Category Settings](#)

Working With the Portal Registry Search

This section contains an overview and describes how to set up the portal registry search.

Understanding the Portal Registry Search

PeopleTools provides the search definition, PTPORTALREGISTRY, which enables you to index the content references registered in the portal menu registry. If you deploy this search definition and build the index (just as you would any search definition), end users can search for content references and menu items using the free-text edit boxes and navigate to the underlying pages using the search results, rather than navigating the menu structure manually. This can provide an extremely efficient means of accessing application content.

Note: The portal registry search category is named Menu.

Note: The PTPORTALREGISTRY search definition provides a similar option as the portal registry search that has been available with the Verity search engine for previous application releases.

If you run the search from the menu pagelet or drop-down navigation Search Menu edit box, the search request only goes against the current portal. For example, if the user is currently logged into the EMPLOYEE portal, the search will be specific to content references defined within the EMPLOYEE portal only.

If you run the search from the Global Search Bar, the search request goes against *all* the portals defined in the local portal registry.

The PTPORTALREGISTRY search definition is based on the delivered PTPORTALREGISTRY query, which verifies the "Valid From" and "Valid To" dates for content references.

Note: For more accurate search results, it is recommended that you schedule running your index updates based on how often the content references in your system change effective dates and expire dates.

Setting Up the Portal Registry Search

To enable Portal Navigator Search:

1. Deploy the PTPORTALREGISTRY search definition and category.
2. Build the index.

There are some additional options that can be added as well:

- To add "Auto suggest," add permission using PeopleTools, Security to this web library: WEBLIB_PORTAL.PORTAL_SEARCH_PB.FieldFormula.IScript_SESPortalQry. This enables the system for anticipating the users desired search based on the partial input.
- To add the Portal Navigator index to Global Search, use PeopleTools, Security, Permissions & Roles, Permission Lists, Search Groups, and add the search category, PTPORTALREGISTRY, to the Search Groups grid.
- To enable the Navigator category to display in the Global Search drop-down list, select PeopleTools, Search Framework, Search Admin Activity Guide, View Search Contexts. Click the Edit Context link and add PTPORTALREGISTRY to the context type Homepage.
- To combine separate search definitions into a single search group, you use the PeopleSoft delivered Administer Remote Search Group ACM plug-in. You can combine search definitions from the same content provider or across content providers. For example, you can combine PTPORTALREGISTRY indexes from multiple content providers into a single Menu global search group.

When you run the Administer Remote Search Groups plug-in, the plug-in:

- Imports the search group from the source.
- Adds the search group to the search context.
- Merges the groups in the search engine server.

For information on using ACM plug-ins, see "Running the PSRUNACM Script" (Automated Configuration Management).

- The facet folder used for CREFs under Fluid Structure and Content, Fluid Pages is the folder that appears after the Fluid Pages folder. For example, in this folder structure (Fluid Structure and Content, Fluid Pages, FOLDER A, CREF A), the facet folder for CREF A will be FOLDER A.
- Every search result displays an icon, which is associated with the CREF. If you want to change the associated icon, consider the following:
 - PTNAVSRCH_DEF_ICN is the default icon used by the portal registry search (Navigator). If you want to replace this icon, you need to replace PTNAVSRCH_DEF_ICN with your own icon. The icon is uploaded in Application Designer.
 - If a CREF has a Tile image associated with it, it will use that Tile image, but only if the image is in the SVG format. If a Tile image is not associated with a CREF or if the CREF uses a non-SVG image type, then the system uses the default icon (PTNAVSRCH_DEF_ICN) for the search result.
- To exclude a CREF from the Navigator search index, ensure that you select the Hide from portal navigation option on the Content Ref Administration page.

Note: A specific CREF must be marked as hidden from navigation to be excluded from PTPORTALREGISTRY. A CREF that belongs to a portal folder that is marked hidden does not mean that the CREF is hidden.

Maintaining the Portal Registry Search

In order to ensure that the PTPORTALREGISTRY index provides accurate search results of content references registered in the portal menu registry, you must remove deleted content references from the index.

When you build an index for the PTPORTALREGISTRY search definition using the incremental index, the indexing process removes expired content references from the index. Currently, the PeopleSoft system does not track the deleted content references in the portal menu registry that should also be removed from the index. In this scenario, PeopleSoft recommends:

1. Ensure that content references that you want to delete are expired by setting the “Valid To” date to a past date on the Structure and Content page.
2. Run the incremental index for the PTPORTALREGISTRY search definition, which will delete the expired content references from the index.
3. Delete the expired content references from the portal menu registry.

Related Links

"Administering Content References" (Portal Technology)

Working with Suggestions Search

Search Framework provides an extremely efficient means to access content references and application content, without navigating the menu structure manually, using the global search bar that is provided in the banner. When you click in the search box or enter a keyword, Search Framework automatically lists suggestions relevant to a user based on several factors, which are described later in this topic.

Menu (PTPORTALREGISTRY search definition) is the default search category. Suggestions are supported for any search category configured for global search.

Note: In the search box, selecting the search icon displays the global search user interface where you can use advanced search features and facets.

See [Working with Global Search](#) for information on using global search.

Note: If Elasticsearch server is down, the search box displays the following message: No results to display.

General Guidelines Regarding Suggestions Search

Consider the following guidelines that control the display of suggestions:

- Three properties control the display of suggestions, and you specify these properties on the Search Options page.
 - Maximum Results in Suggestion Search
 - Number of Days for Recent Ranking
 - Recent Documents Retention (days)

For description of these parameters, see [Managing General Search Options](#).

- Ranking of suggestions takes into account the last accessed time of a component and the number of times the page is visited. If a page is visited frequently, the score of the page increases thereby pushing the page to the top of the suggestions list.
- PeopleSoft stores the data on recent navigation and keyword searches and the updated data is indexed in Elasticsearch using real time indexing, but you should enable these search definitions for real time indexing. Elasticsearch uses the PTSF_RECENTMENU and PTSF_RECENT_KW search definitions.

See [Configuring Real Time Indexing](#).

- Search Framework supports suggestions in a clustered environment.

Suggestions for Menu Search Category

The Menu category is the default category.

The suggestions for Menu category are automatically displayed in the following order:

- The first set of suggestions is based on your recent page activity. If your recent page activity does not fill up the configured number of results in the list, then the next set of suggestions is used.
- The second set of suggestions is based on page activity of other users, but only those pages that you have access to are displayed.

If the second set of suggestions do not fill up the remaining list, then the third set of suggestions is used.

- The third set of suggestions is the top results of a wildcard search on the PTPORTALREGISTRY search index that match the keyword. If a keyword is provided in the search box, the results get

updated to reflect the keyword or keywords you've entered and the suggestions list gets updated based on the prioritization rules.

Suggestions for Content Search

Choose a search category.

Suggestions for any indexed category are automatically displayed in the following order:

- The first set of suggestions is based on your recent page activity in the selected search category. If your recent page activity does not fill up the configured number of results in the list, then the next set of suggestions is used.
- The second set of suggestions is the top results from a wildcard search on the selected search category that match the keyword.

Note: When you select All as a category, a wildcard search is performed on all the indexes except the Menu index, and the top results that match the keyword are displayed.

Embedding Search Box in Component Pages

In component pages, the search box can be embedded that can be used for searching on specific component index. You need to add the search user interface subpage in the component page and call the `SetSearchbarContext` method in the Activate event of the component page.

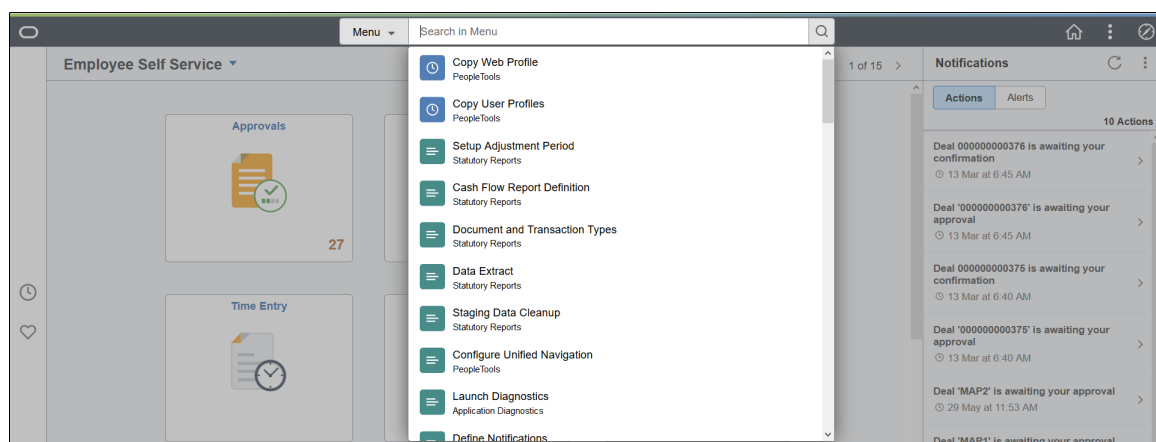
See "Embedding the Suggestions Search Box" (Fluid User Interface Developer's Guide).

For information on the application class, see "SearchbarContext Class" (PeopleCode API Reference).

Example of Suggestions

The following example uses the Menu category and shows the list of search suggestions.

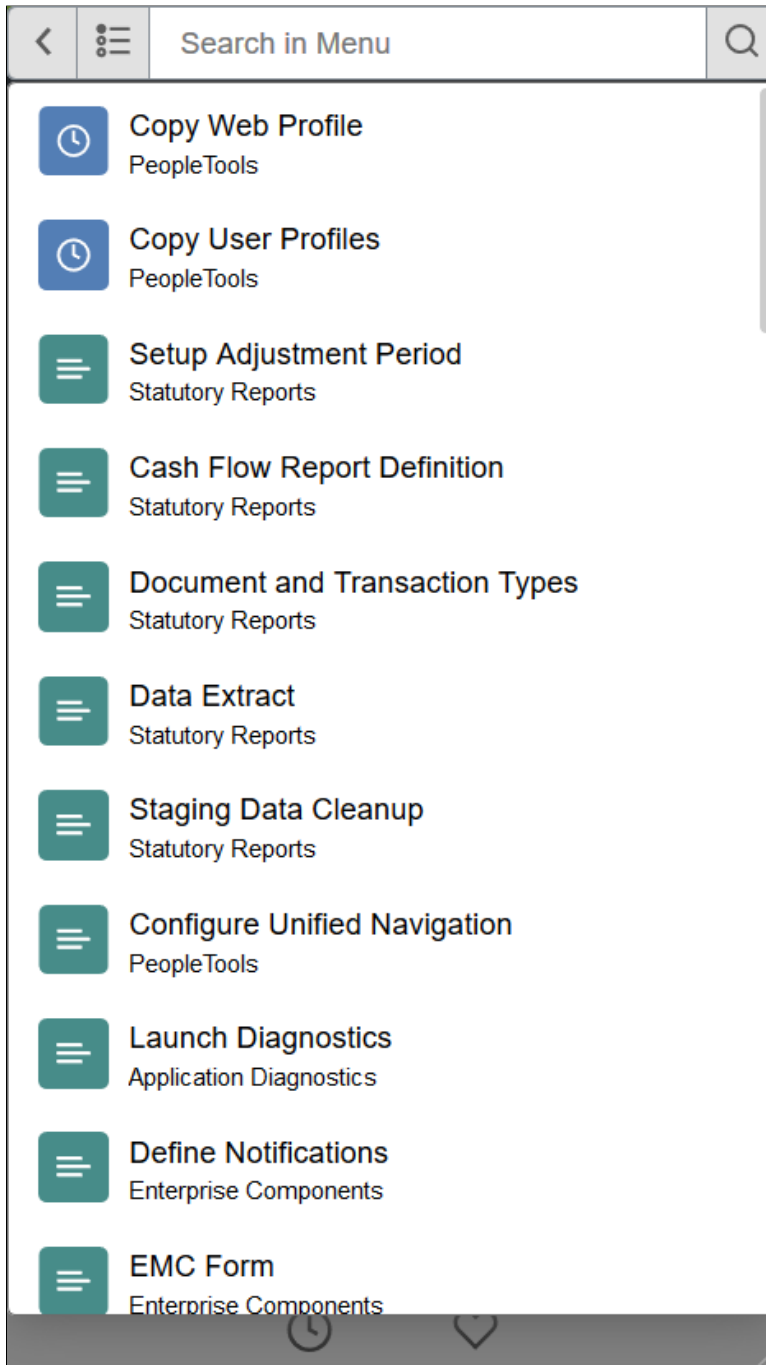
This example illustrates display of suggestions.



Example of Suggestions Search on Small Form Factor Device

On small form factor devices, the search box is displayed in the entire screen. You select categories by using the list icon in the search bar. The default category is Menu.

This example illustrates Suggestions Search on small form factor device.




Testing Suggestions Search

Search Framework provides the Suggestion Search Test page to simulate the suggestions list functionality that is available on the homepage and in transaction pages.

Access the Suggestion Search Test page by selecting **PeopleTools >Search Framework >Utilities >Suggestion Search Test**.

This example illustrates the fields and controls on the Suggestion Search Test page. You can find definitions for the fields and controls later on this page.



<i>Field or Control</i>	<i>Description</i>
Search Mode Selection	<p>You can use this page to simulate several use cases, such as:</p> <ul style="list-style-type: none"> • Header Like Search - Simulates the search available on the homepage using the search box. • Only Comp/Gen URL - Search only PeopleSoft components or generic URL. • Search Hidden Cref - Search that includes hidden content references. • Transaction Search - Simulates the search available on the transaction pages.

Working with Synonyms

Use of synonyms can help users find values without knowing the exact term they are looking for. When the system detects the term used is part of a synonym list, it will return values that also appear as a synonym, thus improving the probability that the result the user wishes is available. However, it should be noted, that users will see more search results though the entered keyword may not be present.

In prior versions, PeopleSoft enabled you to use synonyms during the search process through a configuration file in Elasticsearch. PeopleSoft delivered a configuration file (synonyms_grammar.txt) that enabled you to add synonyms based on your business requirements. PeopleSoft continues to support the

use of the synonyms_grammar.txt file. You should enter synonyms in this file and manually upload this file to Elasticsearch. The file is located at config\analysis\synonym_grammar.txt.

In addition to this functionality, PeopleSoft provides the Synonyms page to manage synonyms. Users can plan to start using the automated manner of maintaining their synonyms through PIA. Users can copy the contents of the synonyms_grammar.txt file and update the synonyms_custom file in PIA.

Synonyms are analyzed during the search process. PeopleSoft uses the synonym token filter of Elasticsearch to analyze the synonyms.

For information on synonym token filter, see Elasticsearch Reference [7.10], which is available on www.elastic.co.

Note: The following documentation is applicable to the use of synonyms managed through the Synonyms page.

Note: By default, synonym search is enabled.

Considerations for Using Synonyms

You should consider the following before you configure synonyms on the Synonyms page:

- PeopleSoft supports translation of synonyms that are delivered by PeopleSoft. Therefore, synonym files for related languages are suffixed with language codes. For example, synonyms_psft_eng.txt.
- Synonym file names are static and should not be changed because the exact file name is defined in the mapping of the index deployment requires .
- Synonym files are common to all indexes for which synonym search is enabled.
- In a scenario where an Elasticsearch cluster is used for multiple PeopleSoft applications, synonym entries are merged to the same file, but the synonyms are separated by owner ID.
- In an Elasticsearch cluster, synonyms are common for all databases for a specific owner ID.

Adding and Copying Synonyms

Use the Synonyms page to add synonyms and to copy the synonyms data to Elasticsearch.

Access the Synonyms page by selecting **PeopleTools >Search Framework >Administration >Synonyms**.

This example illustrates the fields and controls on the Synonyms page. You can find definitions for the fields and controls later on this page.

*Synonyms	*Object Owner ID	Synonym Enabled?	
email, e-mail -> email	PeopleTools	<input checked="" type="checkbox"/>	+ -
ipod, i-pod, i pod	PeopleTools	<input checked="" type="checkbox"/>	+ -

Buttons: Save, Copy to Elasticsearch, Download Synonym file

Field or Control	Description
Synonyms File Name	<p>PeopleSoft delivers two files for synonyms. The default file is synonyms_psft.</p> <ul style="list-style-type: none"> synonyms_custom - This file is for users to add synonyms based on their business requirements. Oracle recommends that users use this file for entering their lists of synonyms to ensure that their synonyms are not overwritten with the PeopleSoft delivered list of synonyms in the synonyms_psft file. synonyms_psft - PeopleSoft applications use this file to deliver its list of synonyms. <hr/> <p>Note: These file names are static and should not be changed.</p> <hr/> <p>Note: Synonym files are common to all indexes for which synonym search is enabled.</p> <hr/>
Synonyms	<p>Enter synonyms as comma separated values. The synonyms entered here are copied to the configuration file as is.</p> <p>An entry in this format: email, e-mail => email, means that when an end user searches for email or e-mail, they will get search results that contain email; not e-mail.</p> <p>An entry in this format: ipod, i-pod, I pod, means that when an end user searches using any of these keywords, they will get search results that contain all these keywords.</p> <hr/> <p>Note: PeopleSoft supports the solr synonym format of Elasticsearch.</p> <hr/> <p>For examples on using the solr format, see the Synonyms section in Elasticsearch Reference [7.10], which is available at www.elastic.co.</p>
Object Owner ID	<p>Select the owner of the synonyms. The owner ID helps determine the team that created or modified the synonyms and also helps to segregate the synonyms in the configuration file.</p>
Synonym Enabled	<p>The Synonym Enabled check box determines whether the synonyms on the row should be added to the selected file. If you do not select the check box, the synonyms are not added to the selected synonym file when you save the file.</p>
Save	<p>Saves the synonyms to the database.</p>

Field or Control	Description
Copy to Elasticsearch	<p>Copies the selected file immediately to the Elasticsearch analysis folder (ES_HOME/config/analysis/). You do <i>not</i> have to restart the Elasticsearch server. When the copy action is issued, the indexes that are enabled for synonyms are temporarily closed, which might affect indexing if it is in progress.</p> <hr/> <p>Important! The copy action should not be performed when indexing is in progress.</p> <hr/> <p>In order for both synonym files to be available in the Elasticsearch analysis folder, you must perform the Copy to Elasticsearch action for both synonyms files.</p> <p>During the copy action, only those blocks relevant to the owner ID are updated. In this manner, when you copy from another environment, it does not overwrite the synonyms as long as the owner IDs are different. For example, in a scenario where an Elasticsearch cluster is used for multiple PeopleSoft applications and the synonym file contains both FSCM and HCM entries, and you're updating synonyms from the FSCM application, the copy action updates only those entries of the FSCM owner IDs. So, if the file has entries of HCM owner IDs, those blocks are not overwritten.</p> <p>The copy action must be performed for all the nodes in a cluster. Before you perform this action, you should ensure that the nodes are configured on the Search Instance page.</p> <p>See Working With Search Instances.</p>
Download Synonym file	<p>Download a copy of the selected synonym file to your local system.</p> <hr/> <p>Note: Oracle does not recommend that you update the downloaded file manually and copy it to the analysis folder directly. If you do so, the data in the manually edited file can be overwritten when you click the Copy to Elasticsearch button.</p> <hr/> <p>If you prefer to manually edit a synonyms file, you can use the synonyms_grammar.txt file because this file is not controlled through PIA.</p>

Clearing Synonyms from a Synonyms File

If you need to clear or delete the synonyms entered in a synonym file, complete the following steps:

1. Access the Synonyms page, and select a synonym file.

Note: The synonyms_custom file is for users to add their list of synonyms. The synonyms_psft file is used by PeopleSoft applications to deliver its list of synonyms.

2. Clear the Synonym Enabled check boxes.

3. Select the Save button.
4. Select the Copy to Elasticsearch button.

Disabling Synonym Search for an Index

By default, synonym search is enabled. Synonym search is associated with an index.

Note: If you choose not to enable synonym search for particular indexes, then during an upgrade your choice is honored, that is, these indexes are not automatically enabled.

You use the Advanced Index Settings page to disable synonym search for an index.

Complete these steps to disable synonym search for an index:

1. Navigate to **PeopleTools >Search Framework >Search Admin Activity Guide** and from Administration in the left panel, choose Deploy/Delete Object.
2. On the Deploy Search Definition page, select the Index Settings link for the desired index.
3. On the Advanced Index Settings page, clear the Enable Synonyms check box, and select OK to save your setting.

This example illustrates the Advanced Index Settings page where you disable synonym search.

Advanced Index Settings ×

[Help](#)

Definition Name PTPORTALREGISTRY

Override Settings

Index Name

Number of Shards

Number of Replicas

Enable Synonyms:

OK

Cancel

Migrating Synonyms Data Using ADS Definition

Oracle delivers an ADS definition to copy or migrate synonym metadata from one environment to another.

For information on using ADS, see "Understanding ADS Projects" (Lifecycle Management Guide).

ADS Definition for Synonyms

ADS definition for synonyms —PTSF_SYNONYMS.

PTSF_SYNONYM_GROUP

Details of PTSF_SYNONYM_GROUP:

Group	Details
Group Name	PTSF_SYNONYM_GROUP
Group Description	Group for PTSF_SYNONYM
Group Members	PTSFSYNONYMS.PTSF_SYN_ENABLED
Group Properties	Merge configurations.

Tables for PTSF_SYNONYMS

Table	Description
PTSFSYNONYMS	Stores synonyms data.
PTSFSYNONYMLANG	Stores translated synonyms data.

Working with Search Operators in PeopleSoft Search

In the Global Search bar and the Keyword Search box, PeopleSoft Search Framework supports the use of search operators along with keywords to specify the type of search to be performed, such as a phonetic search, stemming search, fuzzy search, and so on.

Symbol	Usage	Description
*	document*	<p>Use * (asterisk) as a wildcard to search and match on one or more characters.</p> <p>You can use the * at the beginning of a keyword, at the end of a keyword, or in between a keyword.</p> <p>Examples:</p> <ul style="list-style-type: none"> • document* — returns search results with words beginning with document, such as documents, documentary, documentation. • *ment — returns search results with words ending in ment, such as department, management. <hr/> <p>Note: PeopleSoft Search Framework requires a minimum number of characters to perform a search. The normal minimum is at least four characters plus the wildcard character.</p> <hr/> <p>Note: In an Elasticsearch-based search, PeopleSoft Search Framework supports using * (asterisk) as a leading character.</p> <hr/> <p>Note: You can also use % (percent sign) as a wildcard. In PeopleSoft search, % is treated as a wildcard except when it appears within quotes.</p>
?	Orac?e	<p>Use ? (question mark) as a wildcard to search and match on one character.</p> <p>You can use multiple ? within a keyword, for example, Or?c?e.</p> <p>Orac?e returns search results with words such as Oracle.</p>

Symbol	Usage	Description
“ “	“year-end report”	Use “ “ (double quotes) around words that make up a phrase that must be matched exactly. However, because stemming is applied on the search text, search results containing different forms of the word or phrase will be displayed. For example, if your search text is "project manager," all forms of project manag (Elasticsearch uses shorter and generic form of a word) will be displayed. Therefore, your search results could include project management, project manager, project managers, project managing, etc. You may use facets to obtain the desired search results.
&	document & report	Use & (ampersand) to specify that all words must appear in the results.
	“year-end report” statement	Use (pipe) to search on any of the words or phrase.
!	!john	Use ! (exclamation mark) as a prefix to indicate that phonetic search is to be performed. Phonetic search returns results that have a similar sound. !john returns search results containing names, such as John, Jon, Juan.
\$	\$install	Use the \$ character as a prefix to indicate that stemming search is to be performed. Stemming search returns search results that match the root of a keyword. \$install returns search results containing words, such as install, installing, installed. Note: Elasticsearch supports dictionary and algorithmic stemming, however in the current Elasticsearch integration with PeopleSoft Search Framework, only algorithmic stemming is supported.

Symbol	Usage	Description
~	maintainance~	Use the ~ character as a suffix to indicate that fuzzy search is to be performed. Fuzzy search returns search results based on a likely relevance to the keyword even though the keyword and results spellings may not match. maintainance~ returns search results containing words such as maintenance.

Working with the Search Results

This topic discusses search results on a classic user interface.

After running a search using Keyword Search, the system presents the search results in an intuitive display, enabling interactive filtering and further navigation.

This example illustrates the fields and controls on the Search Results. You can find definitions for the fields and controls later on this page.

Find an Existing Value

Search Criteria
Note: Search will return results last updated 8 hours ago.
 Enter any information you have and click Search. Leave fields blank for a list of all values.

Keywords Search ?

[Show more options](#)

Search Results
50 of 66 results - Keywords "program manager"
[Hide Filter Panel](#)

Business Unit

- TQABU (25)
- LEBU (23)
- HXBUI (16)
- GBIBU (1)
- KR001 (1)

Department

- Foreign Exchange Programs Div (15)
- Controller (12)
- Business Services (10)
- Office of Analysis - TComm (8)

Include History Correct History

Employee Photograph	Empl ID	Empl Record	Name	First Name	Last Name	Second Last Name	Alternate Character Name	Middle Name	Related Actions
	LE0007	0	Hsui,Lewiston Michael	Lewiston	Hsui	(blank)	(blank)	Michael	⊙ >
	LE0003	0	Lyons,Margaret Kelly	Margaret	Lyons	(blank)	(blank)	Kelly	⊙ >
	LE0015	0	Zenith,Abigail Leigh	Abigail	Zenith	(blank)	(blank)	Leigh	⊙ >
	LE0002	0	Keyes,Sharon MIDDLE NAME	Sharon	Keyes	(blank)	(blank)	MIDDLE NAME	⊙ >
	LE0011	0	Thompson,Candace Lynne	Candace	Thompson	(blank)	(blank)	Lynne	⊙ >


Field or Control	Description
Facets	<p>This is the facet pane, which is the area to the left of the search results. Using the facets defined for the current search definition, the end user can filter the search results and drill down closer to the desired information.</p> <p>If searching the All category, you initially see the various search categories represented in the search results.</p> <p>By drilling into the category, you view the facets associated with attributes in the search index.</p>
Title	<p>The title enables you to identify the search result and navigate to the associated component. Select or tap the row to display the page you use to view or modify the underlying data for the search result.</p> <p>The title is defined using the Title edit box on the General tab of the Search Framework, Search Designer Activity Guide, Search Definition component.</p>
Summary	<p>Displays a general overview of the target data.</p> <p>The summary is defined using the Summary edit box on the General tab of the Search Framework, Search Designer Activity Guide, Search Definition component.</p>
Related Actions Indicator	<p>If the search definition is mapped to a component, and if there are any defined related actions associated with that component, the related actions indicator appears to the right of the search result if you've selected the Show Related Actions in Search Results field on the Search Configuration page. Clicking the related actions indicator displays a popup containing the defined related actions. See Configuring a Search Page.</p>

Working With Facets

A facet is an attribute that may be part of one or more search definitions. It provides an alternate representation of the list of values for a given attribute. Facets are attributes used to filter and narrow down a set of search results.

If a facet contains sensitive or PII data, it will not be displayed. For more information on masking of sensitive data, refer to [Masking of Data in Search Results](#).

This example illustrates the fields and controls on the Filter by: Facets. You can find definitions for the fields and controls later on this page.



Facets:

- Show search results across each index.
- When clicked show only results in that index related to that facet value.
- Can further narrow results using sub-facets in that index.

For example, assume the Customer component has a faceted attribute of Customer Level with values of L1, L2, and L3. By selecting the L1 facet link, L2 and L3 values will be removed from the search results list, leaving only the L1 for viewing.

For component Keyword search, the system determines facets associated with the search category to which the component is associated and renders them accordingly.

For Global Search, because the search results might span across multiple indexes that are not related to each other, each search definition within a specific search category should have common attributes marked as facets. These attributes should be generic so search results can be classified under one of these attributes. If non-common attributes are marked as facets in the search category, results not having that attribute will not be displayed.

Note: If an attribute is marked as a facet, and if for a given search document that facet attribute does not have a value, then the query is dynamically modified to return *No Value* as the title of the facet.

Facets display five values by default. If a particular facet has results for more than five facets, *More...* appears at the bottom of the list.

This example illustrates the More... control used to reveal more facet values.

▼ **Department**

- Foreign Exchange Programs Div (15)
- Controller (12)
- Business Services (10)
- Office of Analysis - TComm (8)
- Office of Analysis - EComm (5)
- Plannin & Comm.Relations - E (4)
- Plannin & Comm.Relations -T (4)
- Walnut Creek Branch (3)
- Program Planning & Analysis (2)
- CPSP-Program Plan and Analysis (1)

[More](#)

Click *More...* to expand the list.

Note: For any facet, the maximum number of values displayed is restricted to 100 for performance reasons.

This example illustrates the Less... control used to reduce the facet values displayed.

Once the entire list of facet values is expanded, click *Less...* to collapse the list, showing only five again.


▼ **Department**

- Foreign Exchange Programs Div (15)
- Controller (12)
- Business Services (10)
- Office of Analysis - TComm (8)
- Office of Analysis - EComm (5)
- Plannin & Comm.Relations - E (4)
- Plannin & Comm.Relations -T (4)
- Walnut Creek Branch (3)
- Program Planning & Analysis (2)
- CPSP-Program Plan and Analysis (1)
- Foreign Exchange Programs Div (15)
- Controller (12)

[Less](#)

As you drill into a set of facets, your path is expressed in the form of facet breadcrumbs.

This example illustrates the fields and controls on the Facet filter breadcrumbs.




You can deselect the check box for a filter to remove its effect on the search results, or to reverse the direction of your filtering and expand the results.

Click Clear all filters to remove all applied filters, returning the search results to the state prior to faceted filtering.

Viewing Search Results With Grid Format and List Format

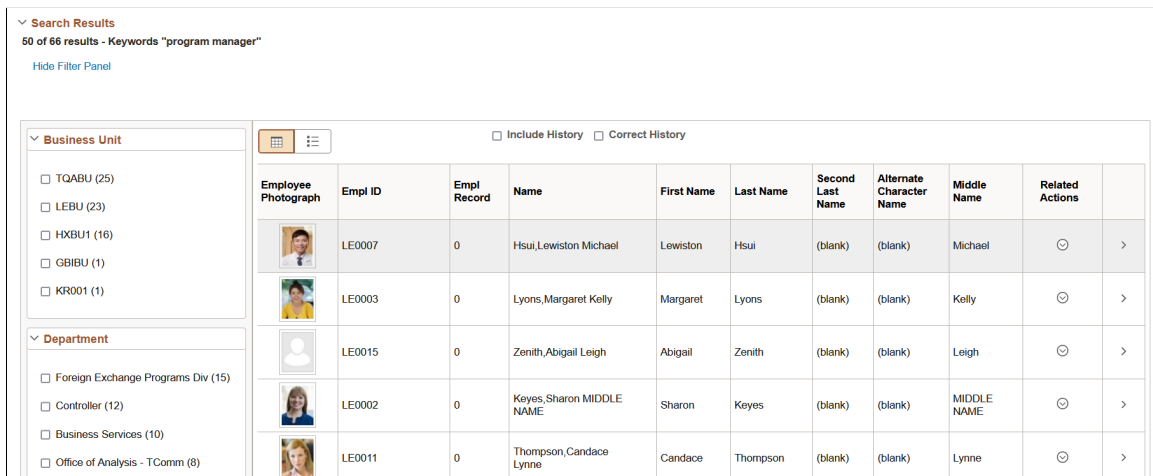
Use the grid and list buttons to toggle between the grid format and the list format for viewing search results.

Field or Control	Description
	Enables you to switch the view of the search results between list format and grid format. The icon representing the current view is greyed out.

Working with the Grid Format

The following example displays the grid format:

This example illustrates the fields and controls on the Search Results in grid format. You can find definitions for the fields and controls later on this page.



To navigate to the page loading the desired data, select the row.

Working with the List Format

The following example displays the list format:

This example illustrates the fields and controls on the Search Results in list format. You can find definitions for the fields and controls later on this page.

Search Results
66 results - Keywords "program manager"
Hide Filter Panel

Business Unit

- TQABU (25)
- LEBU (23)
- HXBU1 (16)
- GBIBU (1)
- KR001 (1)

Department

- Foreign Exchange Programs Div (15)
- Controller (12)
- Business Services (10)
- Office of Analysis - TComm (8)

Include History Correct History

	Job Data: Abigail Zenith Empl ID: LE0015 Empl Rcd: 0 Business Unit: LEBU Department: Plannin & Comm.Relations - E Job Title: Manager - Programs Planning Location: National Office Full/ Part Time: Full-Time
	Job Data: Margaret Lyons Empl ID: LE0003 Empl Rcd: 0 Business Unit: LEBU Department: Office of Analysis - TComm Job Title: Manager - Program Analysis Location: National Office Full/ Part Time: Full-Time
	Job Data: Candace Thompson Empl ID: LE0011 Empl Rcd: 0 Business Unit: LEBU Department: Office of Analysis - EComm Job Title: Manager - Program Analysis Location: National Office Full/ Part Time: Full-Time
	Job Data: Lewiston Hsui Empl ID: LE0007 Empl Rcd: 0 Business Unit: LEBU Department: Plannin & Comm.Relations - T Job Title: Manager - Programs Planning Location: National Office Full/ Part Time: Full-Time

To navigate to the page loading the desired data, select the row.

If the list view contains sensitive or PII data, list view is not displayed. For more information on masking of sensitive data, refer to [Masking of Data in Search Results](#).

Note: The default search result display format for Global Search is the list format.

Working with Related Actions

For components mapped to search definitions, end users can navigate to related actions defined for that component from search results without accessing the main result document. The related actions indicator appears if you've selected the **Show Related Actions in Search Results** field on the Search Configuration page.

Note: The related actions indicator appears only for component-level related *actions*. It does not appear for page-level related actions or for related *content*. Also, end users will only view related actions to which they have access. If there are none for that search result, the related actions indicator does not display.

This example illustrates the fields and controls on the Related actions indicator. You can find definitions for the fields and controls later on this page.

	LE0007	0	Hsui, Lewiston Michael	Lewiston	Hsui	(blank)	(blank)	Michael		>
--	--------	---	------------------------	----------	------	---------	---------	---------	--	---

By clicking on the related actions indicator, a related actions popup appears, showing the possible related actions for that user, according to security access.

After clicking a related action, the system displays the target related action as configured by the application (as in, in a modal window, a new browser window, and so on). The system displays the related action interface with the information associated with the search result displayed.

Related Links

[Configuring a Search Page](#)

Working with Search Results in Fluid User Interface

This section discusses the search results in fluid user interface and the search results on small form factor devices.

Working with Search Results in Global Search

End users use the Global Search bar available on the header to launch global search where they can enter their search criteria. Global Search provides a way for a user to search across all search indexes or a specific group of search indexes.

This example illustrates the Global Search bar.





After running a search using the Global Search bar, the system presents the search results in an intuitive display, enabling interactive filtering and further navigation.

This example illustrates the fields and controls on the Global Search results page. You can find definitions for the fields and controls later on this page.

Search Results	
<p>New Search</p> <p>Category: Company Directory</p> <p>Keywords: Joseph</p> <p>More Options</p> <p>Search Clear</p> <p>Country</p> <p><input type="checkbox"/> United States (5)</p> <p><input type="checkbox"/> Canada (3)</p> <p>State</p> <p><input type="checkbox"/> New Jersey (2)</p> <p><input type="checkbox"/> Quebec (2)</p> <p><input type="checkbox"/> California (1)</p> <p><input type="checkbox"/> New York (1)</p> <p><input type="checkbox"/> Oklahoma (1)</p> <p><input type="checkbox"/> Ontario (1)</p> <p>City</p> <p><input type="checkbox"/> Montreal (2)</p> <p><input type="checkbox"/> Teaneck (2)</p> <p><input type="checkbox"/> New York (1)</p> <p><input type="checkbox"/> Oklahoma City (1)</p>	<p>View Search Results</p> <p>8 results for keyword: "Joseph"</p> <p>8 rows</p> <p>Directory: Anna Joseph Empl ID: HCLA14 Department: Lab Facility Title: Nurse Location: New Jersey Operations Telephone: 879-8944 Email: gh@we.mmb</p> <p>Directory: Joseph Clarins Empl ID: KCD007 Department: Human Resources Title: Manager-HR Help Desk Location: Canadian Headquarters Telephone: 416-360-9239 Email:</p> <p>Directory: Joseph Millier Empl ID: KC0030 Department: Benefits Title: Specialist-Benefits Location: Montreal Office Telephone: 416-360-0384 Email:</p> <p>Directory: Joseph Sanders Empl ID: K0G011 Department: Payroll Title: Analyst-HR/Payroll Location: Corporation Headquarters Telephone: Email:</p> <p>Directory: Joseph Millier Empl ID: KC0030 Department: Finance and Administration Title: Clerk-Payroll Location: Montreal Office Telephone: 416-360-0384 Email:</p> <p>Directory: Anna Joseph Empl ID: HCLA14 Department: Lab Facility Title: Medical Assistant Location: New Jersey Operations Telephone: 879-8944 Email: gh@we.mmb</p> <p>Directory: Alex Joseph Empl ID: AJOSEPH Department: HRMS Division Title: HR-System Test Job Code Location: HR-System Test Location Telephone: Email:</p> <p>Directory: Anthony Bracco Empl ID: KUZ017 Department: Customer Service Title: Director-Customer Services Location: Oklahoma Telephone: 516/324-9838 Email: HCMGENUser1@mailtest.com</p>

The search results page has two panels: the left panel to display search options and facets and the right panel to display the search results.

Field or Control	Description
New Search	<p>The new search area enables you to initiate a new search by selecting a search category and entering keywords.</p> <hr/> <p>Note: A blank search is not supported. That is, you must enter valid search criteria.</p> <hr/>
Facets	<p>Facets are attributes used to filter and narrow down a set of search results. You can select a single facet value or you can select multiple facet values. For more information on facets, see the section on Using Facets later in this topic.</p>
	<p>Use the Hide Filters button to collapse the left panel so that the right panel occupies the entire real estate of the page.</p> <p>Toggle the button to expand the left panel.</p>
Number of returned rows	<p>The total number of rows that match a keyword in the selected search category is displayed at the top of the search page.</p> <p>The Global Search results page shows only the first 50 rows (default) of search results for performance reasons. You can configure the search page to display more rows, that is, to intentionally change the default number of rows, navigate to the PeopleTools Options page.</p> <p>If the total number of rows is greater than 50 rows, you can use the facets in the left panel to filter the desired information.</p>
	<p>If any bulk related actions are defined for a component, the bulk related actions indicator appears at the top of the search results page.</p>
Images	<p>Images can be configured to appear with the search results. If any of the search results is configured to display an image, the title and summary of a search result is indented.</p> <p>See Working with Images.</p>
Related Actions Indicator	<p>If any related actions are defined for the pages of a component, the related actions indicator appears to the right of the search result. Clicking the related actions indicator displays a popup containing the defined related actions.</p>

Using New Search

Users select a search category and enter a search criteria to search for content. Advanced users can leave the search text box blank and click the Search button to display the advanced search options.

Use the More Options control to display additional fields to specify the search criteria. If a search category is not selected, the general set of search fields are displayed. If a search category is selected, the general set of search fields and a set of search fields specific to the selected search category are displayed. For example, if you select Navigator as the search category, you see the general set of search fields and additional fields that are specific to the Navigator search category. Clicking the Fewer Options control collapses the additional search fields.

The additional search fields display default operators for each search field. Use the Show Operators link to display additional operators for each search field. Clicking the Hide Operators control collapses the additional operators.

Note: When you use the contains operator, enter a complete word or a set of words, for example desktop or desktop computer. You can also use wild card characters, but the wild card characters can be used only within a word or at the end of a word. Examples of valid search text containing wild cards: t*st, test*.

The Clear button allows you to clear the search criteria that you entered. However, along with the search criteria, the search results are also cleared enabling you to create a fresh search.

This example illustrates the advanced search fields.

The image shows a 'New Search' form with the following fields and controls:

- Category:** A dropdown menu.
- Keywords:** A text input field.
- With this Exact Phrase:** A text input field.
- With any of these Words:** A text input field.
- Exclude:** A text input field.
- Menu Item Label (contains):** A text input field.
- Description (contains):** A text input field.
- Folder Label (contains):** A text input field.
- Source Application (contains):** A text input field.
- Controls:** 'Fewer Options' and 'Show Operators' links.
- Buttons:** 'Search' and 'Clear' buttons.

Using Facets

By default, on Global Search results page and the component search results page, you can choose multiple facet values to filter or narrow down your search results. You can use the check boxes under a facet to select multiple facet values. For example, in the Department facet, you can choose multiple departments, such as Finance, Human Resources, and so on.

In hierarchical facets too, multi-facet selection is enabled by default. you can choose multiple facet nodes. In hierarchical facets, the facet nodes are displayed as a tree, and a blue check mark indicates the selected facet node. When you select a parent node, the child nodes are not selected by default. If you select a child node after selecting a parent node, then the parent node is deselected automatically. The tree nodes are collapsed by default. If a child node is selected, then the parent node is expanded.

Under a facet (including hierarchical facet), when you select a facet filter, the selected facet filter moves to the top of the facet filters, that is, the selected facet filter is displayed above the non-selected facets. The movement of selected facet filters to the top affects the order of the facet filters. Each time you select a facet filter, the order of facet filters is rearranged to display the selected facet filter above the non-selected facet filters.

Note: The Category facet is set for single select. Multi-facet selection is always turned off for the Category facet.

Note: Only fluid search pages support selecting of multiple facet values or tree nodes. Classic or classic plus search pages do not support this option.

For information on enabling multiple facet-value selection, see [Selecting Facet Settings](#) and [Managing General Search Options](#).

As you drill into a set of facets, your path is expressed in the form of facet breadcrumbs.

The selected facets appear as breadcrumbs in the order you choose the facets. For a hierarchical facet, individual breadcrumb is displayed for each selected tree node.

Facets display 10 values by default. If a particular facet has more than 10 facet values, *More* appears at the bottom of the list. To collapse the expanded facet list, select *Less*.

This example illustrates search results based on multiple facet values and illustrates the display of breadcrumbs.

The screenshot displays the 'Search Results' interface. On the left, there are several facet categories: 'My Association', 'Department', 'Job Family', 'Location', and 'Job Posted In'. The 'Department' facet is highlighted with a red box and includes 'Human Resources (12)', 'All Depts for City of Nelson (4)', and 'Benefits (2)'. The 'Job Family' facet is also highlighted with a red box and includes 'Human Resources (8)', 'Administrative Support (2)', 'Environmental Health & Safety (2)', and 'RS Administrative Job Family (12)'. The main area shows 'View Search Results' for the keyword 'specialist', with 12 results. Breadcrumbs at the top indicate the path: Human Resources > Human Resources > Administrative Support > Environmental Health & Safety. The search results list various job postings, including 'Employment Specialist', 'HRIS Specialist', 'Labor Relations Specialist', and 'Health & Safety Specialist', each with details like Job ID, Type, and Recruiting Location.

If a facet contains sensitive or PII data, the facet is not displayed. For more information on masking sensitive data, refer to [Masking of Data in Search Results](#).

Using Related Actions

For components, related actions can be defined at component level or at page level. End users can navigate to related actions defined for that component from search results without accessing the main result document.

When component-level and page-level related actions are available for a search result, the page-level related actions take precedence and you see the related actions indicator alongside a search result and the bulk mode indicator at the top of the search results.

By clicking the related actions indicator, a related actions popup appears, showing the possible related actions for that user, according to security access.

After clicking a related action, the system displays the target related action as configured by the application (as in, in a modal window, a new browser window, and so on). The system displays the related action interface with the information associated with the search result displayed.

This example illustrates the Related Actions indicator.

The screenshot shows a single search result: 'Job Posting: HRIS Specialist | Job ID: 30017 | Type: External Posting'. Below the title, it lists 'Recruiting Location: Corporation Headquarters | Department: Human Resources | Job Family: Human Resources'. A small circular icon with a checkmark is visible in the top right corner of the result box, indicating related actions are available.

The bulk related actions indicator appears when a search result has bulk related actions defined at the component level. When you click the Bulk mode icon, you see an Actions button with the related actions indicator and check boxes to the right of the search results. You perform a bulk mode action by selecting a search result or search results and by selecting the required bulk action from the related actions list.

Working with Real-time Component Search Results

Fluid user interface enables users to create tiles to access a component. In a similar manner, tiles can be created for a real-time component search enabling you to access the component search without having to navigate to a component.


To configure real-time component search, you must create a pivot grid. In Application Designer, you need to specify the component properties for fluid search page. A component configured for real-time search can be set to use either a master detail or a standard search type page.

Note: In fluid user interface, the real-time component search results are meant to work only with tiles, not with classic pages.

See "Setting up Real-Time Search for Fluid Components" (Fluid User Interface Developer’s Guide) and "Setting Component Properties for Search Options" (Fluid User Interface Developer’s Guide).

You can use configurable search to configure standard or keyword search configuration for a component. See [Guidelines on Using Configurable Search for Fluid Components](#).

This example illustrates the fields and controls on the Real-time Component Search Results page. You can find definitions for the fields and controls later on this page.



Field or Control	Description
New Search	The new search area enables you to initiate a new search by selecting a search category and entering keywords.
Facets	Facets are attributes used to filter and narrow down a set of search results. You can choose to view the facets as a chart.
Facet Sort	Use facet sort to alphabetically sort the facet attributes in an ascending or descending order.

Field or Control	Description
Grid view	<p>Displays the search results in a grid format.</p> <p>In a grid view, related actions, if defined, always appears as the second column of the grid.</p> <p>In a grid view, bulk related actions, if defined, always appear as the first column of the grid.</p>
List view	<p>Displays the search results in a list.</p> <p>In a list view, a title and a summary is displayed for each search result. The related actions indicator is always displayed on the right of the search result.</p> <p>When you click Bulk mode, the bulk related actions check box appears to the right of the search result.</p> <p>If a list view contains sensitive or PII data, the list view is not displayed. For more information on masking sensitive data, refer to Masking of Data in Search Results.</p>
Bulk mode	<p>If the search definition is mapped to a component, and if there are any defined bulk related actions associated with that component, the bulk related actions indicator appears at the top of the search results page.</p> <p>Use the Bulk mode to toggle between displaying bulk related actions and exiting the bulk mode.</p>
Add	<p>Enables you to add new rows to PeopleSoft tables.</p> <p>The Add button is similar to the Add a New Value page in a classic environment.</p> <hr/> <p>Note: The Add button is not available on small form factor devices.</p> <hr/>
Chart	<p>Enables you to view the search results in a chart format.</p>
Options Menu	<p>Use the Options Menu to:</p> <ul style="list-style-type: none"> • Update filters. • Change chart options. • Export data to a spreadsheet. • Save the search with the applied filters so you can access the same search again. The search is saved for a user ID. • Change the sort options.

Field or Control	Description
Related actions indicator	<p>If any related actions are defined for the pages of the component, the related actions indicator appears to the right of the search result. Clicking the related actions indicator displays a popup containing the defined related actions.</p> <p>If you use configurable search to configure search for a component, you can choose whether to display related actions or not on the search results page. Administrators set this option on the Search Configuration page. See Working with Configurable Search.</p>

When you click a real-time component search tile, the search results are automatically displayed without you having to enter any search criteria. Real-time component search uses default search criteria, which is defined in the following:

- Record field properties.
- PeopleCode - SearchInit.
- Pivot Grid.

For more information on defining default search criteria, see "Setting up Real-Time Search for Fluid Components" (Fluid User Interface Developer's Guide).

A real-time component search page can display a maximum of 100 rows of search results. This value can be configured.

Images can be displayed along with the search results if you configure to display images.

Master Detail Real-time Component Search

If you choose master detail as the search type for a component, when you create the pivot grid for the component, you must ensure that you select either Grid and List or List. List view is mandatory for a master detail search page because the search results are displayed in list view in the left panel when you drill into a search result.

This example illustrates the Master Detail Real-time Component Search page.

	Group Box
10009 - Air Cond, Compressor	
10009-A - Substitute item for 10009	SetID CRM01
10010 - Air Cond, Fan	Item ID 10011
10011 - Air Cond, Control Unit	Description Air Cond, Control Unit
10012 - Air Cond, Air Handler Duct	Short Description Air Cond,
10013 - Air Cond, Pipes	Item Group AIR CONDITIONER
10014 - Air Cond, Ducts	Family AC REP PRT
SR1001 - 6600 BTU Room Air (Light Beige)	Category ID 00001
SR1002 - 6000 BTU Room Air (Grey)	Date Added 03/16/2001
	User ID DVP1
	Approval Date 03/16/2001

To traverse through the search results, you can use either the search results in the left pane or the Previous in List and Next in List buttons (guided mode buttons) at the top-right corner of the screen.

To return to the search results page, use the Back (Search Results) button.

Standard Real-time Component Search

If you select the standard search type for a component, selecting List view is optional.

When you drill into a search result, component displays on the page. The search results are not displayed in the left panel. To traverse through the search results, you use the Previous in List and the Next in List (Guided mode) buttons.

Accessing Keyword Component Search

If a component is configured for both real-time search and keyword search, you can access keyword search using the Global Search icon when you are on the real-time search results page. The search category drop-down automatically displays the component from where you launched the keyword search and allows you to search for content within the component.

Related Links

"Understanding Component Real Time Search" (Pivot Grid)

"Creating a Component Pivot Grid Model Using the Pivot Grid Wizard" (Pivot Grid)

Working with Keyword Component Search Results

You can enable Keyword Search on a component. Keyword search is a PeopleTools Search Framework search.

See "Setting up Keyword Search for Fluid Components" (Fluid User Interface Developer's Guide).

The fluid user interface enables you to create a tile to launch a component that is configured for keyword search.

Consider the following when you configure a component for keyword search:

- If you configure a component for real time and keyword search, you can set real-time search to be the default, so you can access the keyword search from the real-time search results page.
- If you set keyword search as the default, you must create separate tiles to launch real-time and keyword search on the component.

A keyword search page enables you to enter search criteria to initiate your search.

You can use configurable search to configure standard or keyword search configuration for a component. See [Guidelines on Using Configurable Search for Fluid Components](#) and [Working with Configurable Search](#).


Working with Search Results on a Touch Friendly Device


On small form factor devices, the search results display varies from desktop or tablet display.


The following search results features are available on small form factor devices:

- Initiate global search using the Search option in the Actions menu.
- Search results are displayed in list and chart view. List is the default view.
- Filter options and search criteria options are available on the options icon.
- Selected filters are listed in the Selected Filters button.
- In real-time component search, advanced search options are not available.
- In component search, the Add button is not available.
- Related actions and bulk related actions are not available.

This example illustrates the fields and controls on the Global Search Results page. You can find definitions for the fields and controls later on this page.



<i>Field or Control</i>	<i>Description</i>
	Click the options icon to display the Filters and Modify Search options.

Field or Control	Description
Filters	Displays the available facets and enables you to filter the search results.
Modify Search	Displays the New Search panel that allows you to enter search criteria.
	<p>On the search results page, use the back button to return to the previous screen.</p> <p>For example, when you drill into a search result, use the back button to return to the search results page.</p>

Masking of Data in Search Results

PeopleSoft provides the ability to mask personally identifiable information (PII) or sensitive data in the search results on classic and fluid search pages (component keyword search results and component real-time search results), and includes classic and fluid prompt page search results. If list view of the keyword search results page contains sensitive data, the list view is removed and if a facet contains sensitive data, the facet is removed.

To mask PII or sensitive data and to remove list view and facets, you can use the PeopleCode methods and functions.

See "SetDisplayMask" (PeopleCode API Reference), "CopyDisplayMask" (PeopleCode API Reference), "SetFacetNamesToRemove" (PeopleCode Language Reference), "SetRemovelistView" (PeopleCode Language Reference).

Note the following when you mask sensitive data for search results or prompt page search results:

- Use the data masking methods in the SearchInit event PeopleCode of any SearchKey field.
- In a method when you specify the number of trailing characters to be unmasked, the unmasked characters are always on the right irrespective of the left to right or right to left language selection.
- If field type is Date, the mask is applied on the formatted date value (as per user preference).
- When masking is applied, it uses field length; not data length. For example, if field length is 8 characters, and you set the trailing unmasked characters as 4, and the field data is *abcdefg*, masking display as ****cdef*.
- If a masked field has translate values, it displays ten asterisks (*).
- In component real-time search, a chart is not displayed if any field marked for masking forms the x-axis or series of the chart. The facet filters and grid displays the data with masking applied on the field. However, in Pivot Grid Viewer, you can view the chart because PeopleCode is not executed.
- For data masking in fluid prompt pages, the prompt table search key record field property All Search Events for Prompt Dialogs should be selected. By default, the property is not selected. SearchSave and SearchInit PeopleCode gets executed only when this record field property is selected.

- After drilling into a search result, if any key field is masked, then the field is set to blank when you return to the search page by selecting the Return to Search button.
- In a grid, sorting is performed on the actual data, not the masked data.
- In search criteria, a field can be disabled if it is sensitive data. A disabled field is removed from grid view if it does not contain masking information.

If the SetDisplayMask method is called on a search record field, then the field is displayed in grid view with masking format applied, even if it is a disabled field.

This example illustrates a search results page with masked values for the National ID field.

Search Results										
View All										
1-43 of 43										
Empl ID	Reference Record	Company	Name	Alternate Character Name	Middle Name	National ID	First Name	Last Name	Second Last Name	Payroll Status
CFG0300	0	CFG	Jenny Dohms	(blank)	(blank)	*****8874	Jenny	Dohms	(blank)	Active
CFG0301	0	CFG	Jack Brent	(blank)	(blank)	*****8817	Jack	Brent	(blank)	Active
CFG0302	0	CFG	Tom Ardy	(blank)	(blank)	*****8825	Tom	Ardy	(blank)	Active
CFG0303	0	CFG	Jessica Snell	(blank)	(blank)	*****8833	Jessica	Snell	(blank)	Active
CFG0304	0	CFG	Bob Alan	(blank)	(blank)	*****3744	Bob	Alan	(blank)	Active
CFG0305	0	CFG	Andy Miller	(blank)	(blank)	*****8841	Andy	Miller	(blank)	Active
CFG0306	0	CFG	Holly Fright	(blank)	(blank)	*****8858	Holly	Fright	(blank)	Active
CFG0307	0	CFG	Joe McCauly	(blank)	(blank)	*****8866	Joe	McCauly	(blank)	Active
CFG0308	0	CFG	Angela Lee	(blank)	(blank)	*****9179	Angela	Lee	(blank)	Active
CFG0309	0	CFG	Ann Ratcliffe	(blank)	(blank)	*****8254	Ann	Ratcliffe	(blank)	Active

Note: Search results initiated from the Search Test page are not masked.

Masking of Data in Search Results for Global Search

PeopleSoft provides the ability to mask personally identifiable information (PII) or sensitive data for a restricted user in keyword search for global search.

Note the following when you mask sensitive data for restricted users:

- When a user enters keywords to search for data, and if the keywords are marked as PII or sensitive data, then search results will not be displayed.
- If a facet contains sensitive data, the facet is hidden for the restricted user.

Monitoring PeopleSoft Search Framework and Elasticsearch Using Kibana

Understanding Monitoring of Search Framework and Elasticsearch

Beginning with PeopleTools 8.57, PeopleSoft Search Framework uses Kibana, an open source analytics and visualizations platform, to monitor and provide analysis on Search Framework and Elasticsearch. Search Framework uses Elasticsearch 7.10 and Kibana 7.10 and enables application developers to monitor the PeopleSoft application data through the application indexes in Elasticsearch.

Kibana is a Web application and works with Elasticsearch. Kibana provides visual reports in the form of charts, tables, etc. based on queries that you set up on Elasticsearch indexes.

In the PeopleSoft implementation, Search Framework uses Kibana to provide data on system metrics and indexing metrics and PeopleSoft application data. It provides real-time data as well as historical data. In the context of search functionality, users need real-time data to understand how their system is performing and to track errors, if any, in indexing. Historical data provides an insight into the performance of the system over a period of time and in the context of indexing proves to be an useful tool to analyze indexing statistics.

This topic aims to provide you instructions and information on using Kibana only with respect to its implementation in PeopleSoft Search Framework. Therefore, some of the capabilities of Kibana may not be described in this topic. If you need more information on any of the capabilities of Kibana, you should refer to the Kibana Guide [7.10], which is available on www.elastic.co.

Kibana Terminology

Kibana uses the following terminology:

<i>Kibana Terminology</i>	<i>Description</i>
Dashboard	A collection of visualizations. You can arrange the visualizations within a dashboard and share a dashboard.
Visualization	A visualization is based on Elasticsearch queries on the Elasticsearch data. A visualization is in the form of charts, maps, tables, etc.

<i>Kibana Terminology</i>	<i>Description</i>
Index pattern	<p>“It tells Kibana which Elasticsearch indices you want to explore. An index pattern can match the name of a single index, or include a wildcard (*) to match multiple indices.”</p> <p>Kibana Guide [7.10], Management, Index patterns.</p>
Attribute	<p>An attribute in PeopleSoft Search Framework corresponds to a search field in Elasticsearch and Kibana.</p>

Understanding Dashboards and Visualizations

In Kibana, a dashboard is a collection of visualizations and searches. A visualization is based on Elasticsearch queries on the Elasticsearch data. These visualizations are in the forms of charts, maps, tables, etc. PeopleSoft Search Framework delivers three system-monitoring dashboards, and each of the dashboards display two or more visualizations. For a list of delivered dashboards and visualizations, see [Delivered Dashboards and Visualizations](#).

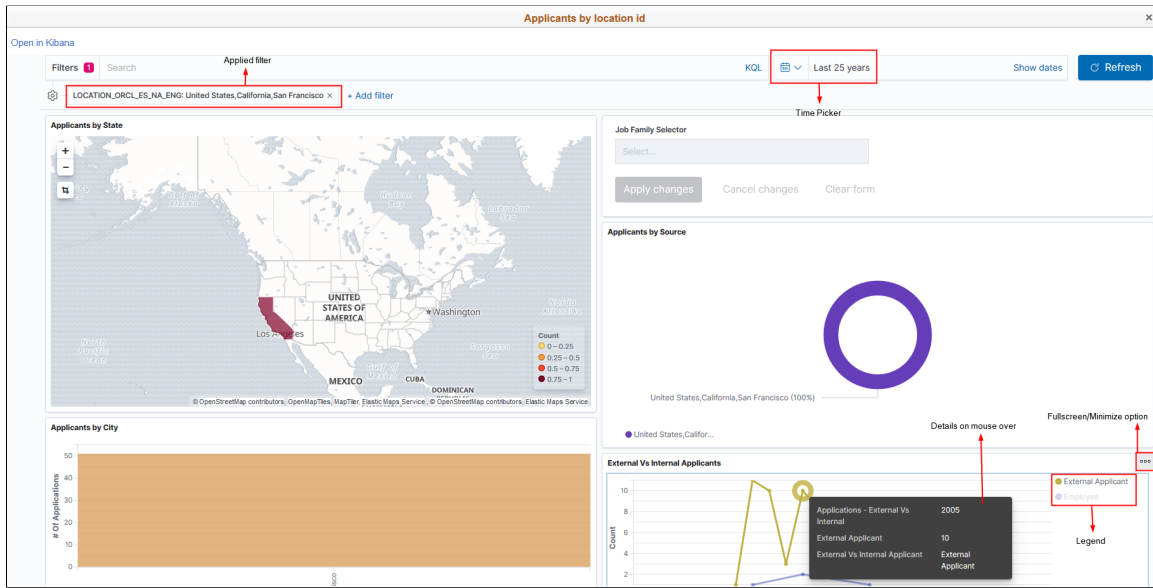
Search Framework enables you to create your own visualizations for PeopleSoft application data. Search Framework does not deliver any dashboard or visualization on application data, but in future PeopleSoft applications will implement some. Based on your business requirements, you may build your own dashboards and create visualizations. Typically, an application developer creates the dashboards and visualizations in Kibana and imports them into PeopleSoft. These dashboards can be configured as a tile or as related information in a PeopleSoft application component. An end-user needs to add this tile to a home page. A dashboard configured as related information appears based on the users access to the component.

For instructions on creating visualizations, refer to [Creating a Visualization for Application Data](#).

Common Elements on a Visualization

This topic discusses some of the common elements on the visualization. For more information on the other elements on the Kibana visualizations, refer to the Kibana Guide [7.10], which is available on www.elastic.co.

This example illustrates the common elements on Visualizations. You can find definitions for these common elements later on this page.



Term	Definition
Applied Filter	Use the filter option to filter data on the visualization.
Time Picker	Use the Time Picker area to restrict the search results to a specific time period. The time picker enables you to set the filter to view real time data and historical data. For more information on time picker, refer to Kibana Guide [7. 10], Discover, Setting the Time Filter.
Fullscreen/Minimize Option	Use the Fullscreen/Minimize Option button to enlarge the selected visualization.
Legend	Kibana randomly uses different colors in a legend. You can collapse the legend if you want more real estate on the selected visualization.
Details on mouse over	Hover the mouse over a chart to view details.

Setting Up Kibana in PeopleSoft Search Framework

Before you can begin to set up Kibana for use in Search Framework, you should complete some prerequisites, such as the installation of PeopleTools 8.59 and Elasticsearch 7.10.

To set up Kibana to provide data on Search Framework, Elasticsearch, and PeopleSoft application data, you need to complete the following steps:

1. Install Kibana.
2. Connect Kibana to the Elasticsearch server.
3. Start Kibana service.

Note: These instructions assume that PeopleTools 8.59 and Elasticsearch 7.10 are installed in your environment.

Installing Kibana

Search Framework supports the installation of Kibana on Windows and Linux.

To install Kibana, you use the delivered ELK DPK. For instructions on installing Kibana, refer to the PeopleSoft Deployment Packages for Elasticsearch Installation (Doc ID 2205540.2) on My Oracle Support.

Important! For the implementation of Kibana with PeopleSoft, Kibana is delivered as part of ELK DPK because a security module is built into it. Therefore, you must use ELK DPK to install Kibana. For Kibana implementation with PeopleSoft, do not install Kibana from the elastic.co Web site.

Kibana is a Web-application, and you can install it on the same server where Elasticsearch is installed or you can install it on a separate server.

Note: Oracle recommends that you install Kibana on one of the Elasticsearch nodes in the cluster to ensure better connectivity between Elasticsearch and Kibana.

Connecting Kibana to Elasticsearch

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

For Kibana to be always available, you should ensure that the Elasticsearch server it connects to is always up and running.

After installing Kibana, you should complete the following tasks:

- Set up Kibana for application data analytics.
- Set up Kibana for monitoring system and Elasticsearch metrics.

Setting Up Kibana for Application Data Analytics

To set up Kibana for application data analytics, you should update the Kibana section of the Search Instance Properties page. You enter the host name and port of the server where Kibana is installed.

For more information, see [Creating Search Instances](#).

Setting Up Kibana for Monitoring System and Elasticsearch Metrics

You should manually set up the connection between the Elasticsearch search instance in PeopleSoft and the Kibana monitoring server using the Configure Server page. For instructions on using the Configure Server page, see [Configuring the Monitoring Server](#).

When you use an Elasticsearch cluster with three nodes, Kibana is connected to only one node of the cluster. The system monitoring job that keeps track of the cluster statistics is activated in the first node provided in the search instance configuration. If there is a change in the primary node in the search instance, it is necessary to stop the monitoring job and then start the monitoring job by using the system monitoring flag on the Configure Server page.

Starting the Kibana Service

On the Windows platform, after you install Kibana, you need to start the Kibana service. Refer to the PeopleSoft Deployment Packages for Elasticsearch Installation (Doc ID 2205540.2) on My Oracle Support.

To start Kibana on Windows, perform the following steps:

1. Open a command prompt, and change directory to Kibana_Home\bin.
2. Run kibana.bat.
3. The service status displays as `Listening` indicating that the service has started

On the Linux platform, after you install Kibana, the Kibana service is automatically started. You can use the following command to check the status of the service:

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

If the Kibana service is not started, you can start the service by running `kibana.sh`

Configuring the Monitoring Server

Use the Configure Server page to establish a connection between Kibana and the Elasticsearch server where you want to index the system metrics, indexing metrics, and application data.

Access the Configure Server page by selecting **PeopleTools >Search Framework >Administration >Configure Monitoring Server**.

This example illustrates the fields and controls on the Configure Server page. You can find definitions for the fields and controls later on this page.

The screenshot shows the 'Configure Server' page with two main sections: 'Elasticsearch' and 'Kibana'.
Elasticsearch section:
 *Host Name: example.com
 *Port: 9200
 SSL Option: Disable
 *User Name: esadmin
 *Password: [masked]
Kibana section:
 *Host Name: example.com
 *Port: 5601
 SSL Option: Disable
 *Scan Interval(Sec): 5
 *Purge Retention(Days): 7
 Enable Index Metrics: Yes (radio button selected)
 Enable System Metrics: Yes (radio button selected)
 Save button

<i>Field or Control</i>	<i>Description</i>
Elasticsearch	
Host Name	Enter the Elasticsearch server where you want to index the system and indexing metrics. <hr/> Note: This should be the Elasticsearch server for monitoring purpose mapped to Kibana during installation of Kibana. You can use the same Elasticsearch server where the application data is indexed or you can use a standalone server for the sole purpose of metric analysis. You can also specify the server in the kibana.yml file.
Port	Enter the Elasticsearch server port that you specified during installation of Elasticsearch.
SSL Option	Select an option based on whether SSL is enabled or disabled on the Elasticsearch server. The default is Disable.
User Name	Enter the user ID for the Elasticsearch admin user. The default is esadmin.
Password	Enter the password associated with the user name.

Field or Control	Description
Kibana	
Host Name	<p>Enter the Kibana server including the domain, which you want to set up as a monitoring server.</p> <p>The host name value can be a DNS name or an IP address</p>
Port	<p>Enter the Kibana server port that you specified during installation of Kibana.</p>
SSL Option	<p>Select an option based on whether SSL is enabled or disabled on the Kibana server.</p> <p>The default is Disable.</p>
Scan Interval(Sec)	<p>Specify the time interval (in seconds) when system metrics will be retrieved from the Elasticsearch server. For example, if you enter 10, system metrics will be retrieved every 10 seconds. The lowest time interval is 5 seconds, that is, you cannot specify a time interval lower than 5 seconds.</p> <hr/> <p>Note: When you set the value for this property, you should consider the performance of the system.</p> <hr/>
Purge Retention Period(Days)	<p>Specify the number of days to store the metrics data. For example, if you enter 7, the metrics data will be retained for 7 days.</p>
Enable Index Metrics	<p>Select Yes if you want indexing metrics to be collected for viewing and analyzing in the Indexing Metrics dashboard.</p> <p>The default is No.</p> <p>When you set it to Yes, the indexing metrics are sent to the configured server. PeopleSoft Search Framework delivers the following index for indexing metrics: <code>ptsf_es_index_metric_yyyy_mm</code>, where <code>yyyy</code> and <code>mm</code> indicates the year and month of the indexed data. The naming convention of the index ensures that the files do not grow too large over a period of time and also ensures that you can easily clean up the index files, if required.</p>

Field or Control	Description
Enable System Metrics	<p>Select Yes if you want system metrics to be collected for viewing and analyzing in the System Metrics dashboard.</p> <p>The default is No.</p> <p>When you set it to Yes, the orcl-monitor-plugin starts indexing the system metrics data of the Elasticsearch server, which is configured as the primary search server on the Search Instance page. PeopleSoft Search Framework delivers the following index for system metrics: ptsf_es_system_metric_yyyy_mm, where yyyy and mm indicates the year and month of the indexed data. The naming convention of the index ensures that the files do not grow too large over a period of time and also ensures that you can easily clean up the index files, if required.</p>
Save	<p>When you select Save:</p> <ul style="list-style-type: none"> • If you set the Enable System Metrics to Yes, the system metrics flag on the Elasticsearch server is updated to enable the indexing of system metrics.. • On save for the first time, an index (orcl_es_auth) gets created on the Elasticsearch server. The index stores the callback service URL for each PeopleSoft database. <hr/> <p>Note: To save the configurations, the Elasticsearch and Kibana servers must be running instances.</p> <hr/>

Monitoring Elasticsearch System Metrics and Indexing Metrics

PeopleSoft Search Framework delivers three dashboards and each dashboard displays several visualizations that provide you real-time data as well as historical data. These visualizations enable you to analyze data and act upon any errors if any. For example, you can use the historical data to view the indexing process and analyze any failed transaction and rectify it, if needed.

This topic aims to provide you information on the dashboards that are delivered by PeopleSoft Search Framework. For information on the other capabilities of Kibana, such as building dashboards, using Discover, using Visualize, etc., refer to the Kibana Guide [7.10], which is available on www.elastic.co.

After configuring the monitoring server as discussed in the Configuring the Monitoring Server topic, you can deploy the delivered system monitoring dashboards and visualizations.

This topic discusses the following:

- Deploy the system-monitoring visualizations.
- View the dashboards and visualizations.
- Create a system-monitoring visualization.

Using the Monitoring Visualizations Page

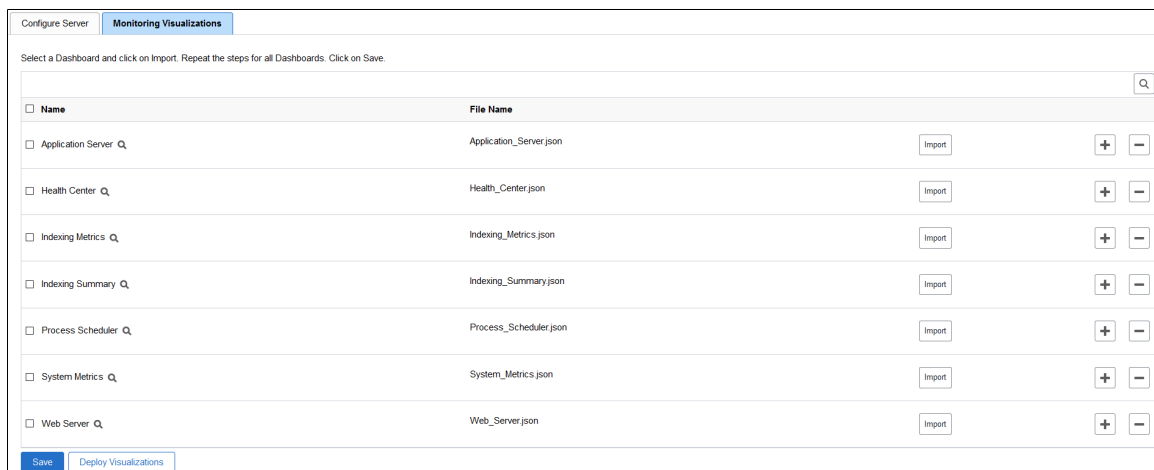
The Monitoring Visualizations page lists the dashboards that are delivered by PeopleSoft. Some of the delivered dashboards pertain to PeopleSoft Health Center that monitors the health and performance of PeopleSoft systems. Refer to "Understanding PeopleSoft Health Center" (Performance Monitor) for a description of the dashboards.

Use the Monitoring Visualizations page to deploy the delivered system-monitoring and indexing-metrics dashboards and visualizations so that you can view the dashboards and visualizations. See [Delivered Dashboards and Visualizations](#) for a list of delivered dashboards and visualizations.

You can also use the Monitoring Visualizations page to import a system-monitoring or indexing-metrics dashboard or visualization that you have created in Kibana. For more information on creating system-monitoring dashboards and visualizations, refer to [Creating a Monitoring Visualization](#).

Access the Monitoring Visualizations page by selecting **PeopleTools >Search Framework >Administration >Monitoring Server**, and then select the Monitoring Visualizations tab.

This example illustrates the Monitoring Visualizations page. You can find definitions for the fields and controls later on this page.



Field or Control	Description
Name	Select the Name check box if you want to select all the dashboards.

Field or Control	Description
Import	<p>If you create a dashboard and visualization in Kibana, the dashboard is listed on this page. Select the dashboard you created and select the Import button to import the dashboard with the associated visualization into PeopleSoft.</p> <p>You would use the Import option when you make any changes to the dashboards. New or modified dashboards need to be imported to PeopleSoft.</p> <p>The import option on the Monitoring Visualizations page should be used only to import the system-monitoring, indexing-metrics, and PeopleSoft Health Center dashboards. To import application data dashboards, see Importing a Dashboard.</p>
Deploy Visualizations	<p>To use the delivered system monitoring and indexing metrics dashboards, you should deploy the dashboards (and the associated visualizations) to Kibana.</p> <p>On successful deployment of dashboards, the following message is displayed: Visualizations indexed successfully.</p> <p>When you access Kibana using the Kibana URL, for example, example.com:5601/app/kibana, the selected dashboards are listed under Dashboard on Kibana.</p>
Save	<p>Select Save when you import a dashboard that you've created or modified.</p> <hr/> <p>Note: When you deploy a visualization or dashboard, you do not need to select Save.</p> <hr/>

Note: PeopleSoft Health Center delivers the Application Server, Health Center, Process Scheduler, and Web Server dashboards.

Delivered Dashboards and Visualizations

PeopleSoft Search Framework delivers the following dashboards and visualizations.

Note: Oracle recommends that you do not edit or delete the delivered dashboards and visualizations.

- Indexing Metrics dashboard provides statistics of the transactions and segments and displays the following visualizations:
 - Transaction Status
 - Transaction/Segment Count
 - Segment History
 - Transaction History

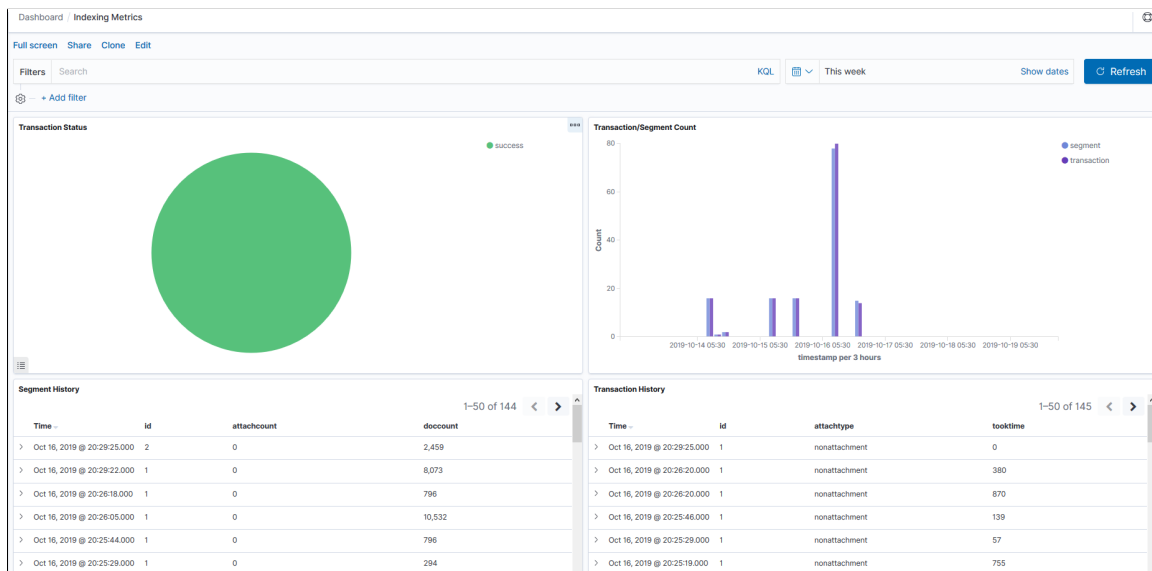
- Indexing Summary provides data of the indexing count and time taken for indexing and displays the following visualizations:
 - IndexCount/Process
 - IndexTime/Process
 - Index Count
- System Metrics provides an overview of the health of the system and displays the following visualizations:
 - CPU (%) Utilization
 - Disk Space Utilization
 - JVM (%) Utilization
 - Total Indexed Documents

Using the Indexing Metrics Dashboard and Visualizations

Access the Indexing Metrics dashboard by selecting Dashboard, and then selecting Indexing Metrics from the list of dashboards.

The Indexing Metrics dashboard provides statistics on the transactions and segments during indexing.

This example illustrates the visualizations on the Indexing Metrics dashboard. You can find definitions for the visualizations later on this page.



Field or Control	Description
Transaction Status	<p>Displays the status of transactions as a pie-chart. The transaction statuses are:</p> <ul style="list-style-type: none"> • Success • Error - The transactions in Error status. Refer to the example described later in this section on the steps to troubleshoot these transactions. • Reject - Search Framework attempts to re-run the indexing for the transactions in Reject status. If the number of transaction in Reject status are high, it may indicate that the Elasticsearch cluster has resource constraints. <p>Selecting any of the slices from the pie-chart enables you to view only the transactions in the selected status in the Transaction Status visualization and the Transaction History visualization displays only the transactions of the selected status, thus enabling you to drill down to each transaction and view details of the transaction, such as process instance, index name, etc.</p> <p>For example, to drill down to the transactions that are in Error status and to understand why the transaction is in Error status, you can perform the following steps:</p> <ol style="list-style-type: none"> 1. Select the Error status slice of the pie-chart. The Transaction History visualization displays only the transactions in the Error status. 2. Select a transaction. The transaction expands to display data in a tabular format. You can use the JSON tab to view it in raw JSON format. 3. Note down the process instance and the name of the index where the error occurred. <hr/> <p>Note: The name of the database forms the suffix for an index name. For example, ptportalregistry_xxxx, where xxxx refers to the database name. The database name is handy when you are monitoring multiple PeopleSoft applications, for example, HCM, FMS, etc. It is also handy when the same Elasticsearch server is used in shared test environments of an application, for example, HCM, HCMdev, HCMTEST, etc.</p> <hr/> <ol style="list-style-type: none"> 4. Log on to PIA, and navigate to the Build Search Index page (select PeopleTools, Search Framework, Search Admin Activity Guide, Administration, Schedule Search Index) where you can use the process instance number to obtain information on the transaction in error.
Transaction/Segment Count	Displays the count of transactions and segment at a specific time period in the form of a vertical bar chart.


Field or Control	Description
Segment History	Displays the segment history and includes the attachment count and document count in each segment.
Transaction History	Displays the transaction history and includes the attachment type (attachment or non attachment) in each transaction and the time taken for each transaction.

Using the Indexing Summary Dashboard and Visualizations

Access the Indexing Summary dashboard by selecting Dashboard, and then selecting Indexing Summary from the list of dashboards.

The Indexing Summary provides data on the indexing count and time taken for indexing.

This example illustrates the visualizations on the Indexing Summary dashboard. You can find definitions for these visualizations later on this page.



Field or Control	Description
IndexCount/Process	<p>Displays the documents indexed for an index and the process instance ID as a vertical bar chart.</p> <p>If you want to view a particular index, you can use the Search area (at the top of the page) to enter the index name. For example, if you want to obtain information on the ep_ap_payments index, you enter <code>indexname:ep_ap_payments*</code>.</p>

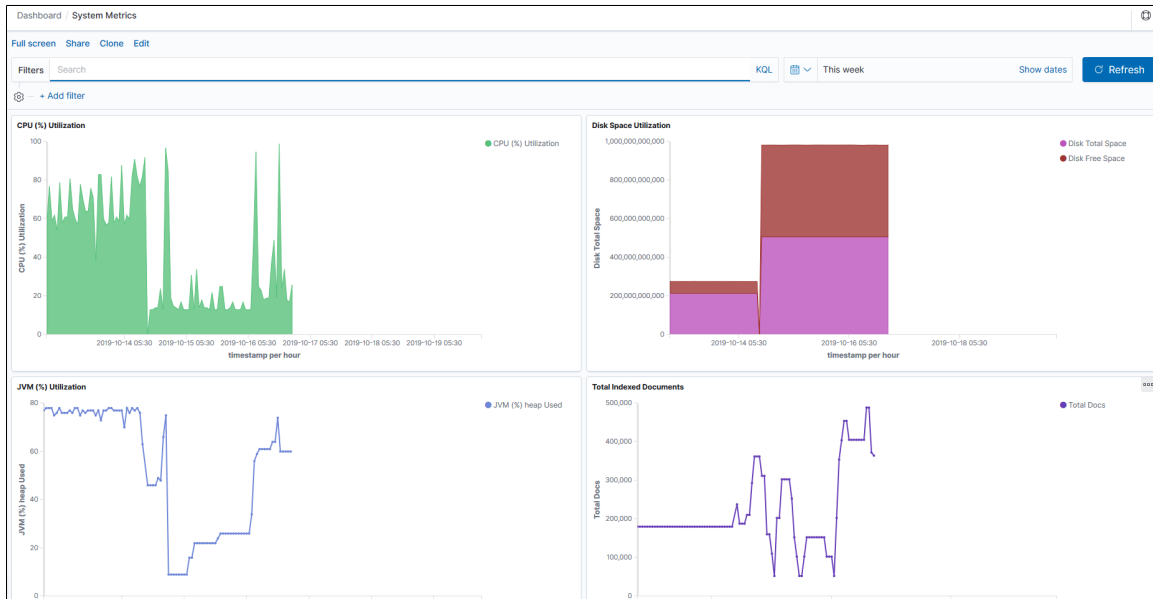
Field or Control	Description
IndexTime/Process	Displays the time taken for indexing documents for each index and the process instance ID as a vertical bar chart. If the document count is low, but the time taken is considerably high, then you should correlate with the System Metrics dashboard to observe whether the system was unstable.
Index Count	Displays the count of attachments and documents for each index in a tabular format.

Using the System Metrics Dashboard and Visualization

Access the System Metrics dashboard by selecting Dashboard, and then selecting System Metrics from the list of dashboards.

The System Metrics dashboard provides an overview of the system health, that is, you can view the real-time health of the system and also view the health of the system over a specific period of time in the past.

This example illustrates the visualizations on the System Metrics dashboard. You can find definitions for the visualizations later on this page.



Field or Control	Description
CPU (%) Utilization	Displays the CPU utilization over a specific period of time by the Elasticsearch server in the form of an area chart.
Disk Space Utilization	Displays the total disk space and the free space available for use in the form of an area chart. The total memory is displayed in bytes.

Field or Control	Description
JVM (%) Utilization	Displays the total JVM percent used by the Elasticsearch server in the form of a vertical bar chart.
Total Indexed Documents	Displays the total number of indexed documents over a specific period of time in the form of a line chart.

Creating a Monitoring Visualization

Search Framework delivers three dashboards and each dashboard contains two or more visualizations. Oracle recommends that you do not modify the delivered dashboards and visualizations. Each of these visualizations are built on index patterns, which must not be modified. The delivered visualizations monitor the important aspects of search, that is, indexing and system health. However, if your business requirements demand to include any other attributes (fields), you may use the delivered index patterns and build new visualizations.

This topic describes how you can use the delivered index patterns and build visualizations of your choice.

Creating a visualization involves the following high-level steps:

- Clone an existing dashboard.
- Create a visualization.

Cloning a Dashboard

To clone an existing dashboard, complete the following steps:

1. Access the Kibana home page using the URL, which is generated when you run the Kibana service.
See [Installing Kibana](#).
2. In the left pane, select Dashboard.
3. From the list of dashboards, select a dashboard.
4. From the top-right corner, select Clone.
Alternatively, select Edit, and then select Save.
5. Enter a name for the cloned dashboard, and then select Confirm Clone.

The selected dashboard is cloned and it displays the same visualizations as the original dashboard. Now, you need to create visualizations for the new dashboard.

Creating a Monitoring Visualization

To create a visualization, complete the following steps:

1. Select Visualize, and then select the plus (+) icon to add a visualization.
2. On the Select visualization type page, select a chart type.

Choose a chart type to display the visualization, such as, tabular form, pie chart, bar chart, and so on. To know more about the chart types, refer to Kibana Guide [7.10], Visualize.

3. Choose a delivered index pattern: `psft_es_index_metrics` or `psft_es_system_metrics`.
4. In the Data tab, specify the Y-axis and the X-axis.

The attributes that are indexed by PeopleSoft are displayed as aggregations in the drop-down list.

5. Select the Metrics & Axes tab, and select values as required.
6. Select the Panel Settings tab, and specify the display settings of the visualization.
7. After ensuring that you've selected or entered the appropriate values, select the Apply changes button, which is next to the Panel Settings tab.

The right pane displays the new visualization as a chart or in a tabular format depending upon the type that you selected.

8. From the top-right corner, select Save.

Enter a name for the new visualization, and then select Save.

9. Select Edit, and then select Add.
10. The Add Panels page displays a list of visualizations, which includes your newly created visualization.

Select the newly created visualization to add to your dashboard.

11. Use the Add Panels page to remove existing visualizations from your dashboard.

On the visualization panel that you want to remove, select the gear icon, and then select Delete from dashboard.

12. From the top-right corner, select Save.

In this manner, you can create more visualizations to add to your dashboard.

If you need to monitor attributes that are not covered by the delivered index patterns and visualizations, it is advisable that you create index patterns that suit your requirements and build visualizations on it. To know more about creating index patterns, refer to Kibana Guide [7.10], Getting Started, Building your own dashboard, Defining your index patterns and for creating visualizations, refer to Kibana Guide [7.10], Visualize, Creating a Visualization.

Working with Kibana Dashboards for PeopleSoft Application Data

Understanding Application Data and Kibana Dashboards

Search Framework enables you to use applications-delivered dashboards and create visualizations in Kibana based on PeopleSoft application data indexes in Elasticsearch. In order to be able to create visualizations and use them, you need to ensure that the following pieces are in place:

- PeopleTools 8.59 is installed.
- Elasticsearch 7.10 is installed and running.

The application search definitions are deployed and indexes are built.

- Kibana 7.10 is installed and running.

On the Search Instance Properties page, enter the Kibana server details.

This topic describes the following:

1. Access Kibana from PeopleSoft.
2. Manage an applications-delivered dashboard.
3. Create a visualization in Kibana.
4. Use a dashboard.
5. Copy Kibana dashboards using ADS definitions.

Delivered Kibana Dashboards

This table lists PeopleTools-delivered Kibana dashboards:

<i>Kibana Dashboard</i>	<i>Link to Product Documentation</i>
PeopleSoft Health Center	See "Viewing the Health Center Summary Dashboard" (Performance Monitor)

Accessing Kibana From PeopleSoft

Prior to PeopleTools 8.58, access to Kibana was restricted to the Elasticsearch Administrator user. The implementation was limited to view the monitoring dashboards and visualizations that were delivered by Search Framework. From PeopleTools 8.58, PeopleSoft Search Framework enables users to view the monitoring dashboards and visualizations and also enables users to create visualizations based on application indexes. Therefore, Search Framework integrates the access to Kibana from PeopleSoft. In the integrated access to Kibana from PeopleSoft, session management in Kibana is performed using callback authentication similar to global search with Elasticsearch.

User Authentication and Data Authorization

Search Framework delivers a custom plug-in that performs user authentication and data authorization using callback logic at runtime.

Kibana displays the application search definitions deployed to Elasticsearch, which enables application developers to create index patterns and create visualizations and dashboards.

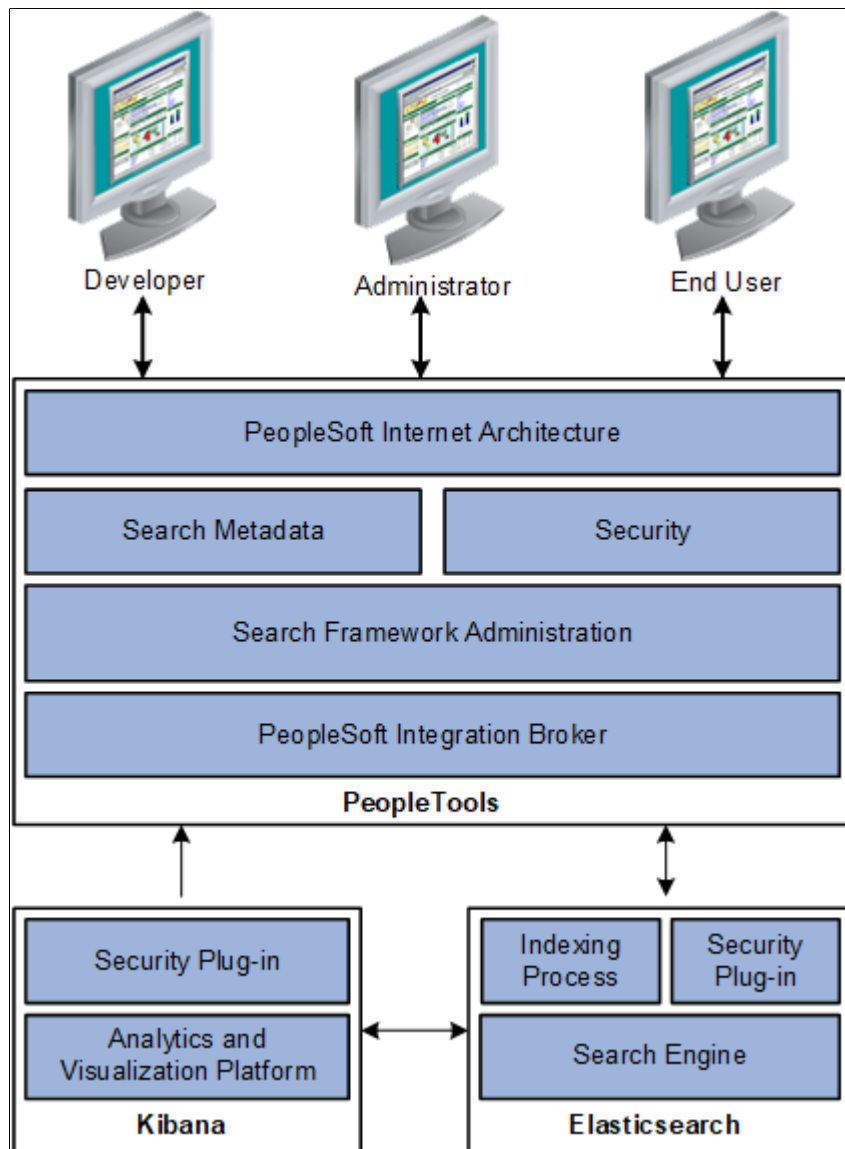
In Kibana, when visualizations and dashboards are created, PeopleSoft applies data authorization based on document-level security set up for the search definition. This ensures that a user has the same set of data access when they use global search and when they use the dashboards.

When a user accesses the dashboards using the tile on their homepage or by using the related information, data shown is restricted by the document-level security applied on the dashboard.

User Authentication and Data Authorization Process Flow

The following illustration describes the process flow involved in user authentication by Kibana and data authorization when a user opens a dashboard.

This illustration describes the elements of PeopleTools, Elasticsearch, and Kibana involved in user authentication and data authorization.



When a user logs on to Kibana, a request is triggered to the PeopleSoft database for authenticating the user using the `orcl_es_auth` index in Elasticsearch.

When a user opens a dashboard, Kibana obtains data authorization for the user and then displays the visualization data based on the security attributes for that user.

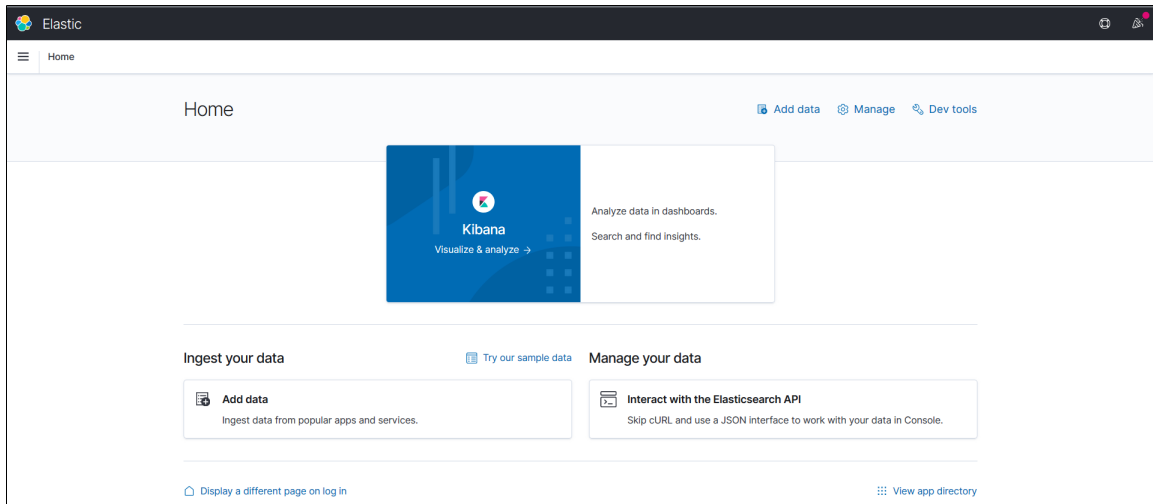
Accessing Kibana Using the Kibana Menu Under Reporting Tools

PeopleSoft provides administrators and business analysts direct access to the Kibana home page from PeopleSoft (PIA) using the Kibana menu under Reporting Tools. Access to the Kibana menu is controlled through the Kibana Administrator role.

Note: Only those users with the Kibana Administrator role can view the Kibana menu under Reporting Tools and access the Kibana home page.

Access the Kibana home page by selecting **Reporting Tools >Kibana**.

This example illustrates the Kibana home page.



For information on performing administrative tasks in Kibana, refer to the Kibana Guide [7.10], which is available on www.elastic.co.

Related Links

[Creating a Visualization for Application Data](#)

Managing a Delivered Dashboard

PeopleSoft applications deliver dashboards with associated visualizations through PUM images. The delivered dashboards are available in the PeopleSoft database. The delivered dashboards should be deployed to the Kibana server.

Note: Oracle recommends that you do not modify or edit a delivered dashboard or visualization. You may copy a delivered visualization and then modify it as per your business requirements.

Note: Kibana dashboards or visualizations that were created using Kibana 7.0 are compatible with Kibana 7.10, so the existing dashboards can be deployed and viewed using Kibana 7.10. However, if you create a new visualization or dashboard using Kibana 7.10, the visualization or dashboard cannot be deployed or viewed in Kibana 7.0.

Deploying a Dashboard

In order to use the delivered dashboards, a user with the Search Administrator role should deploy the dashboard. When you deploy a dashboard, you are copying the dashboard from the PeopleSoft database to your instance of Kibana. The deployed dashboard is then available on the Saved Objects page in the Management section (on the Kibana home page).

Note: Before you deploy a dashboard, you must ensure that the prerequisites of installing and setting up PeopleTools, Elasticsearch, and Kibana are completed. Also you must ensure that the required search definitions are deployed and indexed successfully because a visualization in Kibana is based on the indexes in Elasticsearch.

Use the Deploy Kibana Dashboards page to deploy or update a dashboard.

Access the Deploy Kibana Dashboards page by selecting **PeopleTools >Search Framework >Administration >Deploy Dashboards to Kibana**.

This example illustrates the fields and controls on the Deploy Kibana Dashboards. You can find definitions for the fields and controls later on this page.

Dashboard Name	Deploy Status
Contracts Kibana Dashboard	Undeployed
Payment	Deployed

<i>Field or Control</i>	<i>Description</i>
Dashboard Name	The delivered dashboards are listed on the Deploy Kibana Dashboards page.
Deploy Status	Displays the status of the dashboard, that is, as deployed or undeployed.
Deploy	<p>Deploys the dashboard to your instance of Kibana.</p> <p>You use the deploy option in these scenarios:</p> <ul style="list-style-type: none"> When PeopleSoft applications deliver a dashboard in a PUM image. When you migrate dashboards from one environment to another, for example, from a testing environment to the production environment.
Update	If you make any changes to a deployed dashboard, you should use the update option to synchronize the dashboard in the PeopleSoft database and in the Kibana server.

Assigning Roles for a Dashboard

PeopleSoft applications assign certain roles to the delivered dashboards that enable users to view and use the dashboards and visualizations. You may want to check whether the roles assigned to a dashboard are sufficient for your business requirements, and change the assigned roles, if required. To change the assigned roles, you should access the Kibana Privileges page by selecting PeopleTools > Search Framework > Administration > Assign Kibana Dashboard Roles.

Additionally, you should ensure that the roles assigned to a dashboard are also assigned to users who need to use that dashboard.

See [Specifying User Privileges](#) for more information on assigning roles to a dashboard.

After you deploy a dashboard and assign the required roles to a dashboard, your users can begin using the dashboards and its visualizations. For information on how to use dashboards, see [Using a Dashboard](#).

Working with Field Alias

This topic discusses the field alias implementation and its impact on existing Kibana dashboards.

Understanding Field Alias

A field alias is a meaningful field label that is shown as the Attribute Display Name on the Map Search Attributes page (navigate to **PeopleTools >Search Framework >Search Designer Activity Guide >Search Definition**, and then select the Map Search Attributes tab).

For information on attribute display name, refer to [Mapping Search Attributes](#).

Note: Field alias should be unique in a search definition. Duplicate field aliases within a search index are not copied to Elasticsearch or Kibana.

Understanding the Impact of Field Alias on Existing Dashboards

In PeopleTools 8.59.01, PeopleSoft Search Framework supports field aliases and technical field names, but in PeopleTools 8.59.03, PeopleSoft Search Framework supports only field aliases. That is, technical field names are not available when you create visualizations in Kibana. Therefore, when you access the existing dashboards, you may see errors, such as Field not found or Internal server error.

For details regarding support of aliases in different PeopleTools versions and PeopleSoft applications plans with respect to existing dashboards, refer to [Using Field Aliases in Kibana Visualizations in PeopleTools 8.58.13/8.59.03 and Later \(Doc ID 2787722.1\)](#).

ELK 7.10 DPK supports the field alias implementation in Kibana, so you must ensure that your PeopleSoft system is on the latest version of ELK 7.10 DPK. The ELK 7.10 DPK displays the field alias (and hides the technical field names) when you perform field selection for an index pattern, on visualizations and dashboards, and filter selection. However, the Discover page and saved searches in Kibana display the source fields from the Elasticsearch mapping and not the field alias. You should not use the Discover page to review fields.

Note: The ELK 7.10 DPK does *not* perform an automatic replacement of the technical field names with field aliases in the existing dashboards.

Considerations regarding PeopleSoft applications delivered dashboards:

- PeopleSoft applications will re-deliver the existing (delivered) dashboards with alias-based implementation in subsequent update images. Refer to the *PeopleSoft Search and Kibana Analytics Home Page on My Oracle Support (Doc ID 2205540.2)* for details.
- When you copy the re-delivered dashboards to Kibana, Oracle PeopleSoft recommends deleting any existing metadata from old dashboards to avoid the presence of duplicate names.

- The re-delivered dashboards are created using field aliases in English. Currently, PeopleSoft implementation of Kibana does not support multi language. If you want to use the multi-language fields, you can create new visualizations using the field names from the related language field aliases.

Recreating Existing Visualizations

If you created visualizations (without using aliases), you should recreate these visualizations to avoid any errors.

Complete these steps to recreate a visualization:

1. If you copied the existing visualizations to the newly installed Kibana, you should delete the visualizations from the Saved Objects page (Stack Management > Saved Objects) to avoid any conflict with the recreated visualization.
2. Create a new Index Pattern by choosing the field alias for the date time field.
3. Create a new visualization. You can create a new visualization with the same name as the existing visualization.
4. Either create a new dashboard with the same name as the existing one or you can add the new visualization to the existing dashboard.
5. In the case of scripted fields, you should add a similar, but unique name for the field name, if an alias with the same name already exists. Field names are case sensitive in Kibana.
6. Import the new dashboards from Kibana to PeopleSoft.

For instructions on importing dashboards, see [Importing a Dashboard](#).

Masking of Data in Kibana

PeopleSoft provides the ability to mask personally identifiable information (PII) or sensitive data for a restricted user in Kibana.

To mask PII or sensitive data, see the product documentation for *PeopleSoft Enterprise Components*, "Data Privacy Framework," Enabling Query Data Masking.

Note the following when you mask sensitive data for restricted users in Kibana:

- When an administrator creates a saved search from the Discover tab and it includes a sensitive field, then for the restricted user, the field is available but its values are not displayed.
- For restricted users, visualizations do not appear if PII or sensitive information is present. The PII or sensitive information is masked in the data table visualization.
- For restricted users, if a data table visualization contains PII or sensitive data, then the data table will show **No result found**.
- For restricted users, fields with PII or sensitive data are hidden in saved search.
- PII or sensitive information does not appear in filter criteria and KQL Search.

Creating a Visualization for Application Data

Search Framework enables you to create visualizations based on application data, which is indexed in Elasticsearch. You create a visualization and dashboard in Kibana.

This section discusses the steps that should be completed before you can use the newly created dashboards and its visualizations:

- Import a newly created dashboard.
- Configure a dashboard as a tile or related information.
- Specify user privileges for a dashboard.

After completing these steps, the dashboard is ready for use by an end user. For instructions on using a dashboard, see [Using a Dashboard](#).

Creating a Visualization

Creating a visualization from scratch involves creating an index pattern. In the PeopleSoft implementation of Kibana, an index pattern is associated with a search definition, so a visualization is based on a search definition. The index pattern identifies the mapping of the indexed fields. You could also use an index pattern that has already been created if you need to create a visualization based on the same search definition.

A user with the Search Administrator role or a user with create privilege specified on the Kibana Privileges page can create a visualization.

See [Specifying User Privileges](#).

Before you can create a visualization, you need to create an index pattern.

To create a new visualization, complete the following steps:

1. On the Kibana home page, select the Management icon.

The Management page displays existing index patterns, if any.

Note: You must specify one of the index patterns as the default index pattern. If you do not set a default search pattern, users may see error messages when they attempt to view any of the dashboards.

2. Select Create index pattern.

The Create index pattern page displays the application data indexes that are indexed in Elasticsearch. The indexes are suffixed with a database name.

3. For Step 1: Define index pattern, enter an indexed search definition and suffix it with a wild card character (*) to generalize the index pattern and to avoid dependency on the database. If the search definition name that you entered matches the indexed search definition, a success message is displayed.

For example, ep_ap_payments*

When you enter a name, take into consideration the characters that can be used.

If you prefer to create a new index pattern, based on a delivered index pattern, you can name it with an additional underscore (-) before the wildcard (*), for example, `ep_ap_payments_*`. Thus, you can avoid overwriting your customized index pattern when PeopleSoft delivers an update on the index pattern.

4. Select Next Step.
5. For Step 2: Configure settings, enter a time filter field name or you can select I don't want to use the Time Filter.

If you select a time filter field name, Kibana will use the field to filter the data by time.

Use the Show advanced options link to specify a custom name for the index pattern instead of the system-generated name.

6. Select Create index pattern.

On creating an index pattern, Kibana displays the mapping of all the fields in the index pattern.

In the PeopleSoft implementation of Kibana, you should use the field alias for aggregation to return the original text based on the required case-sensitive text.

If additional languages are enabled in the index, the field names display the corresponding field alias.

7. After creating the index pattern, select the Discover icon to view the indexed documents of the index pattern.

Select a date range to filter the documents.

Note: The Discover page uses the technical names for fields rather than the alias names displayed on the index pattern.

8. Next, select the Visualize icon. The page lists existing visualizations, if any.
 - a. Select the plus sign to create a new visualization.
 - b. On the New Visualization page, select a visualization type.

Hover over each type to read a description of the visualization type.
 - c. On the New Pie/Choose a source page, select the index pattern that you created.
 - d. Specify metrics for the visualization.

The options differ based on the visualization type.
 - e. Specify bucket aggregations as required.
 - f. Select Save, and then enter a name for the new visualization.
9. Select the Dashboard icon, and then select Create a new dashboard.

You may add the visualization to an existing dashboard.

You can add one or more visualizations to a dashboard.

- a. On the Editing New Dashboard page, select Add to choose a visualization.
- b. From the Add Panels page, select the required visualization.
- c. Select Save, and enter a name for the new dashboard.

Note: The dashboard name should be unique.

You may add a description for the dashboard.

The new dashboard is listed on the Dashboard page.

For details on creating visualizations and using the various features of Kibana, refer to the Kibana Guide [7.10] on www.elastic.co.

Importing a Dashboard

A newly created dashboard and visualization is an object in Kibana. The dashboard and visualization must be imported in to the PeopleSoft database in order to apply the needed user privileges and to configure it as a tile or related information.

Note: Only a PeopleSoft user with the Search Administrator role can perform the task of importing dashboards. In fact, the Search Administrator role is required to perform all the subsequent tasks..

Note: Kibana administration pages use the default search instance. If you configure more than one search instance, the search instance with the lowest number in priority is treated as the default search instance.

Use the Import Kibana Dashboards page to import the newly created visualization and dashboard in Kibana to the PeopleSoft database.

Access the Import Kibana Dashboards page by selecting **PeopleTools >Search Framework >Administration >Import Dashboards from Kibana**.

This example illustrates the fields and controls on the Import Kibana Dashboard page. You can find definitions for the fields and controls later on this page.

Select a Dashboard and click on Import. Repeat the steps for all Dashboards. Click on Save.

Dashboards

<input type="checkbox"/>	Name	File Name
<input type="checkbox"/>	<input type="text"/>	<input type="text"/>

Import

Field or Control	Description
Name	Choose a dashboard that you want to import. The dashboards that are created in Kibana are listed here. The Import Kibana Dashboard page also displays the applications-delivered dashboards that you have deployed. For information on applications-delivered dashboards, see Managing a Delivered Dashboard .
File Name	This field is populated when you select the Import button. This is the JSON file that Kibana exports. Typically, the file name is the name of the dashboard with .json as the extension.
Import	Imports the selected dashboard from Kibana into the PeopleSoft database. If you modify a dashboard after it has been imported to PeopleSoft, you should re-import the dashboard to PeopleSoft.
Save	Saves the imported dashboard in the PeopleSoft database.
Add to Project	Enables you to add data set instances on the page to an existing or new ADS project. See "Adding Data Set Instances from Object Designer Pages" (Lifecycle Management Guide).

Configuring a Dashboard as a Tile or Related Information

After importing the dashboard to the PeopleSoft database, the user with the Search Administrator role can configure the dashboard as a tile or related information in a component to enable end users to use the dashboard directly from PeopleSoft.

Note: In addition to the Search Administrator role, you must also have the Portal Administrator role, which is needed to create content references.

Use the Configure Kibana Dashboards page to configure dashboards as tiles or related information in a component.


Access the Configure Kibana Dashboards page by selecting **PeopleTools >Search Framework >Administration >Configure Kibana Dashboards**.

Configuring a Dashboard as a Tile

Search Framework provides the Configure Kibana Dashboards page to enable you to create a content reference for a tile. In order to use this dashboard, end users should add the tile to their PeopleSoft home page using the personalizations option.

For more information, see [Using a Dashboard](#).

This example illustrates the fields and controls on the Configure Kibana Dashboards page - Tile tab. You can find definitions for the fields and controls later on this page.



Field or Control	Description
Content Reference	<p>Enter a unique name for the content reference of the tile. For example, JOB_POSTINGS_TILE. This name is not displayed to the end user. Before you enter a name for a content reference, refer to "Defining Content References" (Portal Technology).</p> <p>The content reference is stored in the PeopleSoft Applications folder at this location: PeopleTools >Portal >Structure and Content >Fluid Structure Content >Fluid Pages >PeopleSoft Applications.</p>
Label	<p>Enter a label for the tile. For example, Job Postings. The label appears as the tile header.</p>
Full View	<p>Choose an imported Kibana dashboard from the list. You may configure a dashboard with multiple visualizations for full view. This is a mandatory field.</p> <p>When a dashboard displays in full view in PeopleSoft, the Open in Kibana link displays if the edit privilege for the dashboard is assigned to you on the Kibana Privileges page. The Open in Kibana link enables you to open the dashboard in Kibana and edit the dashboard.</p> <p>A user with the view privilege can open the dashboard in PeopleSoft where you have options to filter the data or use the time picker to view data based on a time range.</p> <p>See Specifying User Privileges for more information.</p>
Full View App Class	<p>Choose an application class that contains PeopleCode to apply filters based on context and display visualization data accordingly. For more information, see Specifying an Application Class section later in this topic.</p> <p>If you do not specify an application class to filter data, Kibana displays the default visualization data.</p>


Field or Control	Description
Select	<p>Select to launch the Full View page to choose an application class.</p> <p>For more information, see Specifying an Application Class section later in this topic.</p>
Tile View	<p>Choose an imported Kibana dashboard from the list. When you configure a dashboard as tile view, the configured dashboard displays on the loading of the tile.</p> <p>This is a mandatory field.</p> <hr/> <p>Note: Oracle recommends that you configure a dashboard that has a single visualization for tile view.</p> <hr/>
Tile View App Class	<p>Choose an application class that contains PeopleCode to apply filters based on context and display visualization data accordingly. For more information, see Specifying an Application Class section later in this topic.</p> <p>If you do not specify an application class to filter data, Kibana displays the default visualization data.</p>
Select	<p>Select to launch the Tile View page to choose an application class.</p> <p>For more information, see Specifying an Application Class section later in this topic.</p>
Height	<p>Specify the height of the tile.</p> <p>For more information on tile sizes, see "Step 4: Tile Layout Properties" (Portal Technology).</p>
Width	<p>Specify the width of the tile.</p>
Add to Project	<p>Enables you to add data set instances on the page to an existing or new ADS project.</p> <p>See "Adding Data Set Instances from Object Designer Pages" (Lifecycle Management Guide).</p>

Configuring a Dashboard as Related Information of a Component

Configuring a dashboard as related information of a component makes the dashboard appear in the right panel of the component. However, access to the component and dashboard data is controlled by the security privileges assigned to a user.

On the Configure Kibana Dashboards page, select the Component Level tab to configure a dashboard as related information of a component.

This example illustrates the fields and controls on the Configure Kibana Dashboards page - Component Level tab. You can find definitions for the fields and controls later on this page.



Field or Control	Description
Component Name	Choose the fluid component to which you want to configure a dashboard as related information.
Content Reference	Choose the content reference name associated with the selected component.
Full View	<p>Choose an imported Kibana dashboard from the list. You may configure a dashboard with multiple visualizations for full view. This is a mandatory field.</p> <p>When a dashboard displays in full view in PeopleSoft, the Open in Kibana link displays if the edit privilege for the dashboard is assigned to you on the Kibana Privileges page. The Open in Kibana link enables you to open the dashboard in Kibana and edit the dashboard.</p> <p>A user with the view privilege can open the dashboard in PeopleSoft where you have options to filter the data or use the time picker to view data based on a time range.</p> <p>See Specifying User Privileges for more information.</p>
Full View App Class	<p>Displays the selected application class, which contains PeopleCode to apply filters based on context and display visualization data accordingly.</p> <p>If you do not specify an application class to filter data, Kibana displays the default visualization data.</p>

Field or Control	Description
Select	<p>Select to launch the Full View page to choose an application class.</p> <p>For more information, see Specifying an Application Class section later in this topic.</p>
Tile View	<p>Choose an imported Kibana dashboard from the list.</p> <p>When you configure a dashboard as tile view, the related information tile displays the dashboard on loading of the component.</p> <hr/> <p>Note: Oracle recommends that you configure a dashboard that has a single visualization for tile view.</p> <hr/>
Tile View App Class	<p>Displays the selected application class, which contains PeopleCode to apply filters based on context and display visualization data accordingly.</p> <p>If you do not specify an application class to filter data, Kibana displays the default visualization data.</p>
Select	<p>Select to launch the Tile View page to choose an application class.</p> <p>For more information, see Specifying an Application Class section later in this topic.</p>
Content	<p>Enables you to preview the dashboard in the component.</p>
Add to Project	<p>Enables you to add data set instances on the page to an existing or new ADS project.</p> <p>See "Adding Data Set Instances from Object Designer Pages" (Lifecycle Management Guide).</p>

Specifying an Application Class

PeopleSoft provides you the flexibility to filter the data in a visualization based on a context by using an application class that contains PeopleCode to apply filters.

For more information on using the application classes, see ["FilterValues Class"](#) (PeopleCode API Reference) and ["DashboardFilters Class"](#) (PeopleCode API Reference).

Access the Full View or Tile View page by selecting the search icon on the row.

This example illustrates the fields and controls on the Full View page. You can find definitions for the fields and controls later on this page.

<i>Field or Control</i>	<i>Description</i>
Package	Specifies the root package where the application class is created.
Path	Species the path, represented by a colon (:).
Class Name	Choose the relevant application class. You must ensure that you choose the correct application class; if not an error message is displayed.

Related Links

"FilterValues Class Methods" (PeopleCode API Reference)

"DashboardFilters Class Methods" (PeopleCode API Reference)

Specifying User Privileges

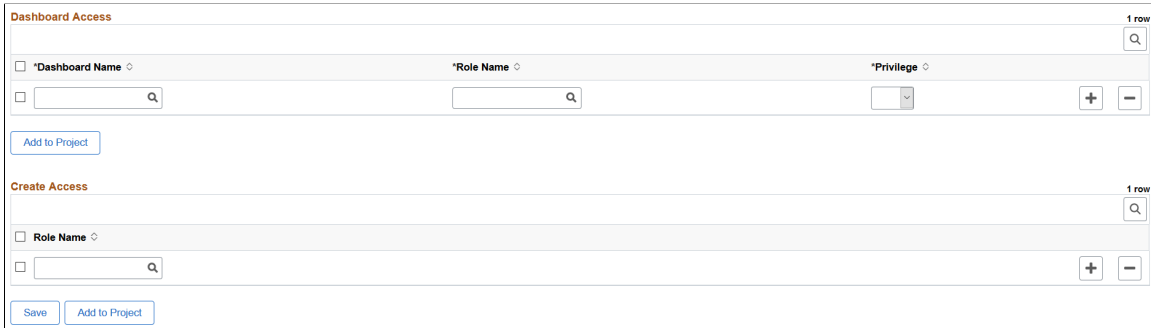
Kibana does not provide role level access to the dashboards and visualizations, so PeopleSoft provides the required security by mapping the dashboards to PeopleSoft roles.

Note: The Search Administrator role is considered as the super user in Kibana. Thus, a user with the Search Administrator role has access to all the dashboards and visualizations. However, the Search Administrator role does not have access to the data in Kibana. *Access to the data is controlled by the application-specific roles.* If a user does not have access to the data, the visualization displays the message: No results found.

Use the Kibana Privileges page to specify user privileges for the dashboards and to create visualizations.

Access the Kibana Privileges page by selecting **PeopleTools >Search Framework >Administration >Assign Kibana Dashboard Roles**.

This example illustrates the fields and controls on the Kibana Privileges page. You can find definitions for the fields and controls later on this page.



Field or Control	Description
Dashboard Access grid	
Dashboard Name	The drop-down list displays the imported dashboards. Choose a dashboard.
Role Name	Choose an application specific role that you want to map to the selected dashboard. The application-specific role enables a user to access and view data in Kibana.
Privilege	<p>Valid values are:</p> <ul style="list-style-type: none"> • Edit - A user with the edit privilege can edit a dashboard or visualization in Kibana and save the changes. On PIA, when a dashboard is displayed, the Open in Kibana link appears for a user with the edit privilege. The Open in Kibana link enables you to open the dashboard in Kibana and edit the dashboard. • View - A user with the view privilege can open a dashboard or visualization in PeopleSoft and view it. The user can use the filter options and time picker options. <p>With this security layer, only users with the specified role can view or edit a dashboard. Furthermore, if you want to personalize a tile (Kibana dashboard) on your home page, the tile will be listed only if you have the privilege, that is, you have the role specified on the Kibana Privileges page. Therefore, administrators must take into account the users to whom the dashboards should be available.</p>
Create Access grid	

Field or Control	Description
Role Name	<p>For a user to create a new dashboard, the user should have at least one of the access roles mapped in the create access list.</p> <p>A user who is the creator of the dashboard possesses the edit privilege irrespective of the role mapped in the dashboard create access grid. A user with the Search Administrator role can access all the dashboards.</p>
Add to Project	<p>Enables you to add data set instances on the page to an existing or new ADS project.</p> <p>See "Adding Data Set Instances from Object Designer Pages" (Lifecycle Management Guide).</p>

A dashboard can be mapped to multiple roles with different privileges for a user.

Using a Dashboard

As an end user, you can use the dashboards depending on your user privileges. The dashboards are made available to you in two ways:

- As a tile, which can be added to a home page.
- As related information of a component.

Using a Dashboard that is Configured as a Tile

The dashboards that are configured as a tile are listed when you select the Add Tile button on the Personalize Homepage page. Every user who wants to view the visualizations should use the Personalize Homepage page to add the dashboard to their home page.

To use a dashboard that is configured as a tile, complete the following steps:


1. On the PeopleSoft applications environment, select a home page that you want to add a dashboard.
2. From the Actions menu, select Personalize Homepage.
3. In personalization mode, select Add Tile.
4. Find the desired tile in the Add Tile dialog box.

Note: Due to the privileges specified on the Kibana Privileges page, the tiles are displayed only if you have the specified roles.

5. Select the desired tile.
6. Select Save on Personalize Homepage.
7. On the home page, select the tile to view the visualizations.

Typically, the dashboard opens in a larger window and displays all the visualizations associated with the dashboard. When the dashboard displays in full view, you can filter the data or use the time picker to view data based on a time range. If the edit privilege is assigned to you, then the Open in Kibana link displays, which enables you to open the dashboard in Kibana and edit it.

This example illustrates the Applicant dashboard configured as a tile on a homepage.



Note: The tile loads all the visualizations associated with the dashboard, so the rendering of the tile may be delayed based on the data volume and network speed.

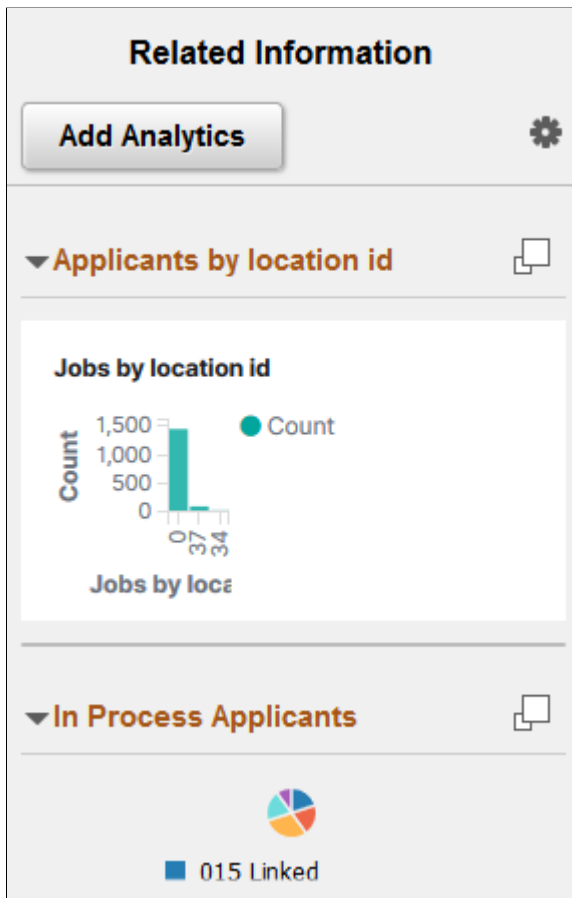
Related Links


"Working with Fluid Homepages and Dashboards" (Applications User's Guide)

Using a Dashboard that is Configured as Related Information of a Component

The dashboards that are configured as related information of a component display on the right panel of the component.

This example illustrates a Kibana dashboard configured as related information. You can find definitions for the fields and controls later on this page.



<i>Field or Control</i>	<i>Description</i>
	<p>Use the button to display the dashboard in full view.</p> <p>Typically, the dashboard opens in a larger window and displays all the visualizations associated with the dashboard. When the dashboard displays in full view, you can filter the data or use the time picker to view data based on a time range. If the edit privilege is assigned to you, then the Open in Kibana link displays, which enables you to open the dashboard in Kibana and edit it.</p>

Note: Kibana dashboards are rendered in real-time with data aggregations, which are collected at runtime. This may delay the rendering of the visualizations based on the data volume and network speed.

Copying Kibana Dashboards Using ADS Definitions

Oracle delivers ADS definitions to copy or migrate Kibana dashboards from one environment to another. For information on using ADS, see "Understanding ADS Projects" (Lifecycle Management Guide).

ADS Definitions for Elasticsearch

The following table lists the ADS definitions:

ADS Definition	Description
PTSF_KIBANA_DBOARD_ATT	Defines the JSON files of a dashboard.
PTSF_KIBANA_DBOARD_CREATE_ROLE	Defines the privileges assigned to a role with respect to a Kibana dashboard.
PTSF_KIBANA_DBOARD_DTL	Defines the dashboards that are deployed to Kibana.
PTSF_KIBANA_DBOARD_ROLE_MAP	Defines the roles mapped to a dashboard.
PTSF_KIBANA_SYSDDB	Defines the system monitoring dashboards that are imported and deployed.

Records Used in ADS Definitions

The following table lists the records that are used in the ADS definitions for Elasticsearch:

Record	Purpose	Merge Group	Custom Transform	Customization	Comments
PTSF_DBOARD_DTL	Stores the dashboard details.	Yes	No	No	The PTSF_KIB_ROLE record is it's child. Has a relation to the PTSF_KIBANA_DBOARD_ATT dataset.
PTSF_DBOARD_ATT	Attachment table that stores the dashboard JSON files.	No	No	No	
PTSF_KIB_ROLE	Maps dashboard to roles.	Yes	No	No	The customer's (target's) privilege (Edit/View access) is preserved.
PTSF_KIBCR_ROLE	Stores the roles assigned with Create access in Kibana.	No	No	No	
PTSF_KIB_SYSDDB	Stores the server metrics dashboard details.	No	No	No	Has a relation to the PTSF_KIBANA_DBOARD_ATT dataset.

Note: Copying of tile configuration using ADS requires purging of cache and restarting of the application server domain on the target environment.

Configuring High Availability in PeopleSoft Search Framework

Understanding High Availability Feature in PeopleSoft Search Framework

Elasticsearch uses distributed architecture and is built to be always available. Elasticsearch uses clusters and nodes to provide high availability.

Oracle PeopleSoft's recommended approach to high availability is to deploy an Elasticsearch cluster with at least 3 nodes on machines that are not on the same physical server and by following the configuration guidelines provided by PeopleSoft Search Framework. The high availability feature of Elasticsearch enables PeopleSoft Search Framework to provide a fail-over mechanism.

Fail-over Mechanism

PeopleSoft Search Framework uses Elasticsearch search engine to provide search functionality. Elasticsearch search engine is distributed by nature, that is, it knows how to manage multiple nodes and manages failover through primary shards and replicas. To enable users to take advantage of this fail-over mechanism, PeopleSoft Search Framework provides you the framework to configure multiple nodes in a search instance of your PeopleSoft implementation. Using the constructs of cluster and node of Elasticsearch, a search instance in PeopleSoft represents an Elasticsearch cluster that can contain multiple nodes.

See [Implementing Fail-over Mechanism Using Multiple Nodes Feature of Elasticsearch](#).

In addition to the built-in fail-over mechanism of Elasticsearch, PeopleSoft Search Framework provides you another fail-over mechanism by using its feature of configuring multiple search instances for the Elasticsearch search engine. Implementing this feature means that you are creating multiple Elasticsearch clusters.

See [Implementing Fail-over Mechanism Using PeopleSoft Search Framework's Multiple Search Instance Feature](#).

Implementing Fail-over Mechanism

This topic discusses the two ways of implementing fail-over mechanism in PeopleSoft Search Framework.

Implementing Fail-over Mechanism Using Multiple Nodes Feature of Elasticsearch

When you configure PeopleSoft Search Framework with the Elasticsearch search engine, PeopleSoft Search Framework enables you to implement a fail-over mechanism.

When you use the multiple-nodes option and when you deploy search definitions/categories and generate indexes, Elasticsearch automatically distributes the indexed data to the nodes based on its constructs of primary shards and replica shards.

You can create multiple nodes on the same machine or on different machines. If you are creating nodes on different machines, but want the nodes to be on the same cluster, refer to the Elasticsearch documentation for information on how you can complete the setup.

To create another node using the installation instructions, refer to *PeopleSoft Deployment Packages for Elasticsearch Installation on My Oracle Support (Doc ID: 2205540.2)*.

Complete these steps to set up multiple nodes in a search instance:

1. After Installing the Elasticsearch search engine, you need to configure the search functionality.
2. Use PIA to configure the PeopleSoft Search Framework by creating a search instance.
3. Enter information in the required fields and save the search instance.
4. Install another node on the same machine or on a different machine. Follow the instructions in the *PeopleSoft Deployment Packages for Elasticsearch Installation on My Oracle Support (Doc ID: 2205540.2)*.
5. Reopen the search instance and In the Search Instance Properties grid, add the newly created node.
6. Enter information in the required fields to point to the new node and save the search instance.

Note: In a multi-node Elasticsearch cluster, always ensure that the replica value is set to at least 1. For details, see [Managing General Search Options](#).

7. Deploy the search definitions/categories and generate indexes for the search instance.

Your search functionality is complete and ready for use. In the event that a node fails, PeopleSoft Search Framework automatically points searches to an active node of the search instance.

Implementing Fail-over Mechanism Using PeopleSoft Search Framework's Multiple Search Instance Feature

PeopleSoft Search Framework provides a mechanism for fail-over by enabling you to create multiple search instances using Elasticsearch as a search engine. The Search Framework enables you to specify priority for these multiple instances so that based on the priority, search will be directed to a different search instance when a search instance fails. If you want to implement this option, you should understand that you need to create another Elasticsearch cluster that you will configure as a search instance on the PeopleSoft system.

When you create multiple search instances, you must ensure that you deploy search definitions/categories and generate indexes on all the search instances. This is to ensure that your search instances contain the same searchable data for a user.

Complete these steps to set up multiple search instances:

1. After Installing the Elasticsearch search engine, you need to configure the search functionality.
2. Use PIA to configure the PeopleSoft Search Framework by creating a search instance.
3. Create or reuse the existing search definitions and search categories.
4. Deploy the search definitions/categories and generate indexes for the search instance.

Your search functionality is complete and ready for use.

5. Create another Elasticsearch cluster by following the instructions in the installation guide.
6. Use PIA to configure the PeopleSoft Search Framework by creating a search instance.
7. Deploy the same search definitions/categories and generate indexes on the new search instance, that is, you are creating another cluster with the same details.
8. On the Search Administration page (PeopleTools, Search Framework, Administration, Search Administration), specify the priority for the search instances.

In the event that a search instance fails, PeopleSoft Search Framework will automatically point searches to the next search instance as specified in the priority list.

Related Links

[Working With Search Instances](#)

Integrating External Data with PeopleSoft

Understanding the Integration of External Data with PeopleSoft

PeopleSoft enables you to integrate external data with PeopleSoft through an integration mechanism. Based on your business requirements, you may want to integrate external data with PeopleSoft to enable you to arrive at business decisions, use data for visualization and analytics, supplement existing data in PeopleSoft, and so on. For example, as a PeopleSoft HCM payroll user, you may want information about your company products that are purchased, but this information is in external systems. As a PeopleSoft FSCM user, you may prefer to perform cash discounting or dynamic discounting based on federal funds rate, and the federal funds rate information is in an external system.

If you plan to use the integration mechanism, you should ensure that the following requirements are met:

- Credentials for data retrieval from an external system.
- Network infrastructure that allows you to connect to the external system.
- Pre-defined roles and users who have the permissions or access to the external data.

PeopleSoft uses the following PeopleSoft products and third-party tools in the integration mechanism:

- PeopleSoft Search Framework
- Elasticsearch
- Logstash

Setting Up the Integration of External Data with PeopleSoft

Before you set up the integration of external data with PeopleSoft, you must decide whether SSL configuration for Logstash is required.

If you need to configure SSL for Logstash, see [Configuring SSL for Logstash](#).

To set up the integration of external data with PeopleSoft, you must complete the following steps:

1. Create a search definition for external data.
2. Deploy an external data search definition to Elasticsearch.
3. Specify a data source for the external data search definition.
4. Build an index for the external data search definition.

To complete these steps, you may either use the External Data Source Setup navigation collection or the individual pages.

Creating a Search Definition for External Data

Use the External Definition page to create a new search definition for external data and to view a list of existing search definitions for external data. You could create a search definition for each functional area, for example, recruiting, job data, and so on.

Access the External Definition page by selecting **PeopleTools >Search Framework >Deploy External Source Defns >External Data Source Setup**, and from the left panel, choose External Data Search Defns and select the Add Definition or plus button.

Alternatively, you may select **PeopleTools >Search Framework >Deploy External Source Defns >External Data Search Defns**.

To create a new definition, select the Add Definition or plus button that displays the External Definition page.

This example illustrates the fields and controls on the External Definition page. You can find definitions for the fields and controls later on this page.

<i>Field or Control</i>	<i>Description</i>
Definition	Specify a name for the search definition.
Description	Enter a description.

Attributes

You use the Attributes section to specify the attributes which are available in the external data and are functionally relevant in the external data index. Attributes are fields that form a part of each document in the Elasticsearch index of external data. You decide which attributes you want based on the functional information that you require in the index. These attributes can be obtained either from the external data or from PeopleSoft mapping to existing PeopleSoft fields.

Field or Control	Description
Target Attribute	Enter an attribute that is available in the external data that you want to include in the index.
Type	Choose a type for each attribute. The selected type is defined in Elasticsearch as part of the index mapping.
Attribute Description	Enter a description for an attribute.

PeopleSoft Data Mapping

This section enables you to map some of the defined attributes with existing PeopleSoft fields. The purpose of PeopleSoft mapping is to provide fixed values per data source for each document indexed from a data source. The PeopleSoft mapping also helps to group all the documents for each data source with a specific attribute value stored in PeopleSoft.

Field or Control	Description
Search Definition	Choose a search definition that has the attribute that you want to map with the external data.
Search Attribute	Choose the attribute that you want to map.

Security

Field or Control	Description
No Security	Select to define no security restriction for a search definition's data. Anyone with access to the application can view the data for a search definition set to No Security. That is, the data is public to all users.

Field or Control	Description
Source Level Security	<p>Select to allow or restrict access to the entire search definition as per the specified user or role. That is, only specified users and roles are able to view data for that search definition.</p> <p>Specify values for the following fields:</p> <ul style="list-style-type: none"> • Type - Select user or role, that is, set the scope of the intended access restriction. • Name - Choose a role or user name. • Privilege - Select Deny or Allow, that is, set the restriction or access privilege. <p>For more information on source level security, see Setting Source Level Security.</p>
Component Level	<p>Select to associate component level security with the external data search definition, that is, users who have access to a component can view the data of the external data search definition.</p> <p>Specify values for the following fields:</p> <ul style="list-style-type: none"> • Market - Select a Market for the component. • Component Name - Choose a component to associate with the external data search definition. • Search Criteria - Enter any additional custom search criteria to be added automatically to the search engine search query. The search criteria should take the form <code>ATTRIBUTENAME:VALUE</code> and can use keywords and brackets. For example: <code>PTSF_SBO_NAME:EP_AP_VENDOR & (-STATUS:PND)</code> <p>For more information on search criteria, refer to Mapping Components to Search Definitions.</p>

Deploying an External Data Search Definition to Elasticsearch

Use the Deploy External Search Defns page to deploy the external data search definition and create a placeholder in Elasticsearch.

Access the Deploy External Search Defns page by selecting **PeopleTools >Search Framework >Deploy External Source Defns >External Data Source Setup**, and from the left panel, choose Deploy External Search Defns.

Alternatively, you may select **PeopleTools >Search Framework >Deploy External Source Defns >Deploy External Search Defns**.

This example illustrates the fields and controls on the Deploy External Search Defns page. You can find definitions for the fields and controls later on this page.

Deploy External Search Defns					
Definition	Index Name	Created On	Count	Create/Delete Index	1 row
1	EXT_PRODUCT_INFO	ext_product_info_t5901jdc	03/07/21 9:18PM	0	<input type="button" value="Delete"/>

Field or Control	Description
Definition	Displays a list of external data search definitions.
Index Name	When you click the Create button, the name of the Elasticsearch index for the specified external data search definition is generated. The naming convention is the name of the search definition with the database name appended to it.
Created On	Displays the day and time when the Elasticsearch index was created for the specified search definition.
Count	Displays the count of documents in an index. On the Deploy External Search Defns page, the count is zero because the index for the search definition is not built. After an index is built, the number of documents in an index is displayed.
Create/Delete Index	<p>Select the Create button to create a placeholder for the index in Elasticsearch.</p> <p>Select the Delete button to delete the index. The delete action only removes the index placeholder from Elasticsearch; it does not delete the search definition.</p>

Specifying a Data Source for External Data Search Definition

Use the Data Source Definition page to specify the actual source of data. In this context, a data source represents the actual source of external data from which the required data is retrieved and indexed.

Access the Data Source Definition page by selecting **PeopleTools >Search Framework >Deploy External Source Defns >External Data Source Setup**, and from the left panel, choose Define External Data Sources.

Alternatively, you may select **PeopleTools >Search Framework >Deploy External Source Defns >Define External Data Sources**.

This example illustrates the fields and controls on the Data Source Definition page. You can find definitions for the fields and controls later on this page.

Cancel
Data Source Definition
Save

*Data Source Name

Search Definition

▼ **Input**

*Data Source Type ▼

*Method ▼

URL

▶ **Header**

Body

Proxy

Additional Options

▶ **Filters And Transformations**

▶ **Output**

▶ **Test External Data Source**

<i>Field or Control</i>	<i>Description</i>
Data Source Name	Specify a name for the data source.
Search Definition	Select an external data search definition to associate with the data source definition. If you're creating a data source for intermediate processing, you do not associate a search definition with it.

Input

In the Input and Header sections, the fields that appear are based on your selection of data source type.

<i>Field or Control</i>	<i>Description</i>
Data Source Type	Select a type for the data source. Valid options for type are: <ul style="list-style-type: none"> • File - A CSV file that contains the actual data. • URL - A REST end point to which a request is made and it provides data as a response.

Field or Control	Description
File Type	<p>The File Type field appears when you select File as data source type.</p> <hr/> <p>Note: CSV is the only supported file type.</p> <hr/>
File Name	<p>The File Name field appears when you select File as data source type.</p> <p>All external CSV files should be placed in the <i>psftappdata</i> folder under <LOGSTASH_HOME>. Any generated files are also saved to this folder.</p> <p>Specify the absolute path of the CSV file.</p>
Method	<p>The Method field appears when you select URL as data source type.</p> <p>Choose the type of HTTP request that needs to be issued.</p> <p>Valid options for method are:</p> <ul style="list-style-type: none"> • GET • POST • PUT
URL	<p>The URL field appears when you select URL as data source type.</p> <p>Specify the path of the external data. The path should not be absolute; it should be relative. Path should be relative to <LOGSTASH_HOME>/psftappdata.</p>

Header

Field or Control	Description
Body	<p>The Body field appears when you select URL as data source type.</p> <p>This field contains the JSON request for the POST and GET methods.</p>
Proxy	Specify the proxy server details for REST API connection.

Field or Control	Description
Additional Options	Specify any additional connection parameters.

Filters and Transformations

This example illustrates the fields and controls on the Filters and Transformations section of Data Source page. You can find definitions for the fields and controls later on this page.

Filters And Transformations

*Type

External Search Definition Attributes

Split Fields

Data Source Attribute Mappings 3 rows

Target Attribute	Type	Source Attribute	Fixed Value	Suggestions
EXT_BRAND	String	<input type="text"/>	<input type="text"/>	<input type="text" value=""/>
EXT_CATEGORY	String	<input type="text"/>	Sports	<input type="text" value=""/>
EXT_DESCRIPTION	String	<input type="text" value="[[itemsummaries]][title]"/>	<input type="text"/>	<input type="text" value=""/>

PeopleSoft Mappings 1 row

Search Definition	Search Attribute	Mapping Value	Suggestions
1 <input type="text" value="QE_CL_ITM_M"/>	<input type="text" value="QE_CATEGORY_ID"/>	<input type="text"/>	<input type="text" value=""/>

In the Filters and Transformations section, the fields that appear are based on your selection of type.

Field or Control	Description
Type	<p>Specify the type of data transformation that you want to achieve.</p> <ul style="list-style-type: none"> External - Select external if a data source is created for intermediate processing, that is, the data source is created to generate data which needs to be processed further before sending to Elasticsearch. PeopleSoft - Select PeopleSoft if you're mapping the data source with an external data index.

External Attributes

The External Attributes section appears when you select External as the type.

Field or Control	Description
Split Fields	<p>You need to specify split fields when you select URL as the data type.</p> <p>You indicate a field from the data source, which requires a split.</p> <p>When a request is made to the REST end point specified in the URL type, it returns a response containing JSON documents. These JSON documents can be nested or hierarchical in nature. If the data you require is under a parent node, then you need to use the split fields option to obtain the children under the parent node.</p> <p>For example, in a product data source, the information related to a product is under a parent node called itemSummaries. To obtain the information of sub-products (children) under the parent node for further mapping, you need to split the field.</p>
Allowlist	<p>Use this field to limit or filter the data from the data source that is passed to Elasticsearch. Enter comma separated values in this field.</p> <p>If you do not enter any values here, all data from the data source will be passed to Elasticsearch.</p>

External Search Definition Attributes

The External Search Definition Attributes section appears when you select PeopleSoft as the type.

Field or Control	Description
Split Fields	<p>You need to specify split fields when you select URL as the data type.</p> <p>You indicate a field from the data source, which requires a split.</p> <p>When a request is made to the REST end point specified in the URL type, it returns a response containing JSON documents. These JSON documents can be nested or hierarchical in nature. If the data you require is under a parent node, then you need to use the split fields option to obtain the children under the parent node.</p> <p>For example, in a product data source, the information related to a product is under a parent node called itemSummaries. To obtain the information of sub-products (children) under the parent node for further mapping, you need to split the field.</p>

Data Source Attribute Mappings

In this section, you map the attributes of the external data search definition to the attributes of the data source. The attributes of the external data search definition associated with this data source are displayed enabling you to map these attributes to the attributes of the data source.

Field or Control	Description
Target Attribute and Type	<p>Contains the attributes and its type from the external data search definition that you associated with the data source.</p> <p>These values cannot be edited here.</p>
Source Attribute	<p>As part of the mapping of the external data search definition with the attributes of data source, if you need to map any nested fields under a split attribute, the mapping can be done using the split field as described in the previous section. For example, to map the field, Title, under itemSummaries, you should enter it as [itemSummaries][title] in the Source field.</p>
Fixed Value	<p>The Fixed Value column enables you to specify a value that meets your requirements. For example, for the City source attribute, you want information only about San Francisco, so you should specify San Francisco in the Fixed Value column.</p>

Field or Control	Description
Suggestions	When you select the suggestions icon, it displays the attributes from the actual data source, which you can choose for the Source Attribute column.

PeopleSoft Mappings

Use the PeopleSoft Mappings section to map a specific value from the data source to the PeopleSoft attribute. The mapping from the external data search definition, which you associated with this data source, is displayed here.

Field or Control	Description
Search Definition	The search definition defined in the PeopleSoft Mapping section of the external data search definition is displayed here. This field is read-only.
Search Attribute	The search attribute defined in the PeopleSoft Mapping section of the external data search definition is displayed here. This field is read-only.
Mapping Value	Map a specific value from the data source to the PeopleSoft attribute. For example, the EP_AP_VENDOR search definition contains the VENDOR_NAME_SHORT attribute. For this attribute for the vendor, XYZ, the value stored in PeopleSoft is XYZ01. You should enter this value, XYZ01, in the Mapping value column. This value is added to all the Elasticsearch documents which are indexed and obtained from the XYZ data source.
Suggestions	Displays the values from the selected search attribute, which you can choose for the Mapping Value column.

If an external search definition, which is mapped to an external data source, is updated, then the updated data automatically reflects on this page.

Output

In the Output section, you specify where the data source output should be stored.

<i>Field or Control</i>	<i>Description</i>
Output Type	<p>Select an output type based on your requirement.</p> <ul style="list-style-type: none">• Elasticsearch - Select Elasticsearch to save the external data to Elasticsearch.• File - Select File for data sources created for intermediate processing. <p>If you select File, the following two fields —Codec and File— are displayed.</p>

Field or Control	Description
Codec	<p>The Codec field corresponds to the codec setting of the Logstash file output plugin.</p> <p>The external data that is sent as input requires to be transformed in order to ensure that it has all the necessary attributes to be used in PeopleSoft. The following sample code shows the transformed data.</p> <pre> %{[categories][id]},https://api.myproduc ts.io/newproducts/categories/%{[categori es][id]}?api-version=2018-10-18&Subscrip tion-Key=c455d00cb0f64e238a5282d75921f27 e </pre> <p>Output gets transformed from:</p> <pre> "id" : "6642", "name" : "Prepared Foods", "type" : "1", "_links" : [{ "href" : "/newproduc ts/categories/6642?api-version=2018-10-1 8", "rel" : "self", "type" : "GET" }] } </pre> <p>To:</p> <pre> 6642,https://api.myproducts.io/newproduc ts/categories/6642?api-version=2018-10-1 8&Subscription-Key=c455d00cb0f64e238a528 2d75921f27e </pre> <p>For more information on the codec setting, refer to Logstash Reference [7.10] in www.elastic.co.</p>
File	Specifies the location of the file.

Test External Data Source

This section enables you to check the REST end point request and response.

<i>Field or Control</i>	<i>Description</i>
Request	Specify a request to test the data source.
Response	Displays the REST end point response.
Submit	When you click Submit, the REST end point request is initiated.

Building an Index for an External Data Search Definition

Use the Manage External Search Index page to control the indexing of data from the defined data sources using Logstash. The page lists the existing data sources and enables you to schedule the index build process or immediately start or stop the indexing process.

Access the Manage External Search Index page by selecting **PeopleTools >Search Framework >Deploy External Source Defns >External Data Source Setup** and from the left panel choose Manage External Search Index.

Alternatively, you may select **PeopleTools >Search Framework >Deploy External Source Defns >Manage External Search Index**.

This example illustrates the fields and controls on the Manage External Search Index page. You can find definitions for the fields and controls later on this page.

The screenshot shows the 'Manage External Search Index' page. On the left is a navigation menu with options: 'External Data Search Defns', 'Deploy External Search Defns', 'Define External Data Sources', and 'Manage External Search Index' (which is highlighted). The main content area has a title 'Manage External Search Index' and a table with the following data:

Data Source Name	Definition	Index Name	Count	Created On	Schedule	Start/Stop
1 EXT_PRODUCT	EXT_PRODUCT_INFO	ext_product_info_15901jdc	0	03/07/21 9:18PM		<input type="button" value="Start"/>

<i>Field or Control</i>	<i>Description</i>
Data Source Name	Specifies the data sources that you added on the Data Source page.
Definition	Specifies the external data search definition that you associated with a data source.
Index Name	Displays the Elasticsearch index name for the specified external data search definition. The index name is generated when you deploy the search definition..
Count	Indicates the document count for each search definition.
Created On	Indicates the day and time when the index was built.

Field or Control	Description
Schedule	Enables you to set up the time and frequency of the indexing process. The Schedule Definition page is described later in this topic.
Start/Stop	When you select Start, Logstash initiates the indexing process as per the schedule you've set. When you select Stop, Logstash aborts the indexing process.

Scheduling the Indexing Process

Use the Schedule Definition page to specify the time and frequency of the indexing process, which is performed by Logstash. The scheduling information is stored in the Logstash configuration file. After setting the time and frequency schedule, you must select the Start button (on the Schedule Definition page) that initiates the indexing process as per the schedule .

Access the Schedule Definition page by selecting the calendar icon on the Manage External Search Index page.

This example illustrates the fields and controls on the Schedule Definition page. You can find definitions for the fields and controls later on this page.

Cancel

Schedule Definition

Done

Schedule At

Date

12/15/2020

📅

Minute

▼

Hour

▼

Repeat Hourly

Every

Repeat Weekly

Days

▼

<i>Field or Control</i>	<i>Description</i>
Schedule At	<p>Specify the date and time when you want the index process to commence. The Schedule page provides you the flexibility to specify the exact minute to commence the indexing process.</p> <p>The time and frequency options are mutually exclusive. When you schedule the date and time for the indexing process, the other two frequency options are not considered and vice versa.</p>
Repeat Hourly	Specify to repeat the indexing process by the number of hours, for example, you can set it to repeat every four hours.
Repeat Weekly	Specify to repeat the indexing process on a specific day of the week and the indexing process repeats every week on the specified day.

Note: The default schedule for indexing is one minute. If you do not set the time or frequency, the system runs the indexing process after one minute from the time you press the Start button. It runs only once after one minute and does not repeat.

Note: After you've set the schedule and saved it, ensure that you select the Start button. The Start button initiates the indexing process as per the schedule.

Troubleshooting

Creating Users and Assigning Roles in Elasticsearch

After you install Elasticsearch, you need to create users and assign roles to users in Elasticsearch. These users in Elasticsearch are necessary to validate the incoming requests from PeopleSoft Search Framework. While the user information is maintained in the `esusers.properties` file, the role information is maintained in the `esuserroles.properties` file at `ES_HOME/plugins/orcl-security-plugin/config/properties`.

The following roles are delivered to perform specific operations:

- `read` - to provide permission as read-only.
- `admin` - to perform read and administrative operations.
- `security` - to access the security index.

To create users and assign roles to users, use the `elasticsearchuser` script available in `ES_HOME/bin`. You can use the script for the following purposes:

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

```
elasticsearchuser adduser [user]
```

- To assign roles to an existing user:

```
elasticsearchuser addrole [user]
```

Use commas to separate the roles assigned to a 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, for example, a password:

```
elasticsearchuser encrypt [text]
```

This command is used for encrypting the keystore password while configuring SSL. The password needs to be encrypted in the `elasticsearch.yml` configuration file.

- To create a new key:

```
elasticsearchuser buildkey
```

This command creates a new key to encrypt password in text file.

For example, you may want to create a superuser with the roles of admin and security assigned to the superuser.

1. `elasticsearchuser adduser superuser`

Enter a password and confirm the password.

2. `elasticsearchuser addrole superuser`

Enter the password.

Enter comma separated list of roles, for example: admin,security.

Related Links

[Working with Authentication and Authorization](#)

Managing Elasticsearch Services

This appendix describes how you can manage Elasticsearch service.

Managing an Elasticsearch Service on Linux

After installing Elasticsearch, you can add an Elasticsearch service, verify that the service starts automatically, and remove the service.

Adding Elasticsearch as a Service

To add Elasticsearch as a service and verify that the service is created:

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 the `systemctl` command:

```
systemctl status elasticsearch
```


This should give an output with the status "active (running)," as shown in the following 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 the **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 the **service** command:

```
service elasticsearch status
```

This should give an output such as:

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

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 the following 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 the **service** command:

```
service elasticsearch status
```

This should give an output such as:

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

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 the **chkconfig** command. The command should give an empty output:

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

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

```
service elasticsearch status
```

This should give an output such as:

```
elasticsearch: unrecognized service
```

Troubleshooting Elasticsearch Configuration

This appendix describes some of the administrative functions you can perform on Elasticsearch using the command prompt and how to handle some common errors.

Performing Administrative Tasks

This section describes some of the administrative tasks you can perform in Elasticsearch using the command prompt.

Adding or Deleting a Node

When you add or delete a node, the shards and replicas are dynamically distributed across the available nodes.

Note: The number of replicas determine the number of nodes that can go down without any data loss. The maximum number of replicas should be less than the number of nodes, that is, one less than the number of nodes. Replicas that are equal to or greater than the number of nodes do not add any benefit to the cluster, rather indexing will fail in such a scenario.

Changing the Node Type

When you deploy the Elasticsearch search engine using the DPK that PeopleSoft delivers, by default the node type is set to master-data type.

If you want to change the node type, you need to update the *elasticsearch.yml* configuration file.

To specify a data node, for example, set

```
node.master: false
node.data: true
```

You should specify this setting before the node is brought up.

Viewing the List of Indexes in an Elasticsearch Cluster

After deploying search definitions and building indexes, you can find out the list of indexes present in your ES Cluster by executing the following command in the browser:

```
http(s)://host:port/_cat/indices?v
```

You can use a similar command to view the plug-ins in your Elasticsearch implementation:

```
https(s)://host:port?_cat/plugins?v
```

Handling Common Errors

If you want to troubleshoot the PeopleSoft Search Framework and Elasticsearch integration or any search related issues, Oracle PeopleSoft recommends, as an initial step, to check the following:

- Ensure that the integration of Elasticsearch with PeopleSoft Search Framework is correctly completed. You can check the integration using the Ping, Login Test, and Proxy Test and Validate buttons on the Search Instance Properties page in PIA.
- If you've installed Kibana and configured it, you can use the Indexing Metrics dashboard to view whether the indexing process is successful or not. If the Transaction Status visualization displays errors in indexing, click the error slice of the pie-chart, and obtain the process instance ID and the index name.
- Check the run control ID on the Schedule Search Index page for any errors.
 - In case of an error, the View Error link gives you a clue on the error.
 - The Schedule Search Index page also displays a link where you can view the IB transactions from asynchronous service monitoring.

- If the Schedule Search Index page has no details, click the Message Log link and follow the steps in it.
- Check the Integration Broker queues.

Note: If Full Direct Transfer is enabled, then Integration Broker is bypassed. You should check the Indexing Metrics dashboard for indexing errors. If Full Direct Transfer is not enabled, the Indexing Metrics dashboard does not capture data, so you should check the Integration Broker queues for indexing of documents without attachments.

- Navigate to the Asynchronous Services page (PeopleTools, Search Framework, Search Admin Activity Guide, Administration, Asynchronous Services) and in the Queue Level field, select Sub Contract.
- Verify the Queue status for the PTSF_ES_SEND_Q queue.
- If PTSF_ES_SEND_Q status is New, verify whether the IB Pub/Sub domain is active or not and activate if required (by selecting PeopleTools, Search Framework, Search Admin Activity Guide, Administration, Domain Status).

Ensure that there are no previous error transaction with the same subqueue name.

- If PTSF_ES_SEND_Q status is Error, read the error message.

If the error was due to unavailability of the Elasticsearch server, that is, no connectivity with IB, you may resubmit the transaction.

If you need more clarity on the error and need to troubleshoot further, activate the IB message log to level 5 and submit the transaction, and watch for the response from Elasticsearch.

- Check the PeopleSoft MsgLog to identify reported errors.

MsgLog is located in <PIA_HOME>\webserv\<DOMAIN>\applications\peoplesoft\PSIGW.war \WEB-INF folder.

- Check the Elasticsearch logs to identify reported errors.

Elasticsearch logs are logged by default in the ES_HOME/logs folder. However, you can specify the location of the log folder in the elasticsearch.yml configuration file.

Errors in MsgLog

Error	Possible Cause	Resolution
<pre>{"error":{"root_cause":[{"type":"orcl_auth_plugin_exception", "reason":"Invalid user/password"}], "type":"orcl_auth_plugin_exception", "reason":"Invalid user/password"}, "status":401}</pre>	This might be due to invalid proxy user and password.	In PIA, navigate to the Search Instance Properties page and in the Search Instance Properties section, ensure that the proxy user and password are entered correctly. Click the Proxy Login button and if the credentials are correct, you should see a login success message.
<pre>": {"type":"unavailable_shards_exception","reason":"[orcl_es_defaultindex_h9285505][0] Not enough active copies to meet write consistency of [QUORUM] (have 1, needed 2). Timeout: [1m], request: [BulkShardRequest to [orcl_es_defaultindex_h9285505] containing [304] requests]"}}</pre>	This message is displayed because the number of replicas are more than the number of nodes.	In PIA, navigate to the Search Options page and update the replicas to a number that is less than or equal to the number of nodes in the cluster.

Errors in Elasticsearch Logs

Error	Possible Cause	Resolution
<pre>ConnectTransportException[[host] [ipaddress:9300] connect _timeout[30s]]; nested: ConnectTimeoutException[connection timed out: /ipaddress:9300]; at org.elasticsearch.transport.netty.NettyTransport.connectToChannels (NettyTransport.java:987)</pre>	This is due to one node trying to connect other node to form the cluster. However, the IP address of the other node is not reachable.	Navigate to \$ES_HOME/config and ensure that the network.host and discovery.zen.ping.unicast.hosts values contain the correct machine details.
<pre>org.elasticsearch.transport.BindTransportException: Failed to resolve host</pre>	Elasticsearch network host is set to a host name or IP address that is not resolvable.	Set it to the IP address of the machine where Elasticsearch is running.
<pre>RemoteTransportException- EsRejectedExecutionException</pre>		Set the threadpool.bulk.queue_size value to a value higher than the number of concurrent bulk requests you want to send.

Errors on PIA

Error	Possible Cause	Resolution
<p>Integration Gateway - External System Contact Error</p> <p>Integration Gateway was not able to contact the external system. The network location specified may be incorrect, or the site is permanently or temporarily down.</p>	<p>This error might be due to no connectivity to the Elasticsearch server.</p>	<p>In PIA, navigate to the Search Instance Properties page, and ensure that the correct host name and port are entered. Click the Ping button and verify whether the Elasticsearch server is up and running.</p>
<p>Search Instance Ping not working.</p>	<p>This error might be due to the following reasons:</p> <ul style="list-style-type: none"> • Incorrect configuration of Integration Broker. • Exception in Elasticsearch. • Elasticsearch server is down or invalid network host. • Incorrect configuration of search instance. 	<ol style="list-style-type: none"> 1. Check if Integration Broker is configured correctly and ping the IB node to verify. If IB is correctly configured, check whether the Elasticsearch server is up and running. You can ping the Elasticsearch server from any browser by entering the server and port, for example, http://server:port. 2. If the ping from a browser is not successful, but the Elasticsearch server is up, ensure that the network.host property in the elasticsearch.yml configuration file found in ES_HOME/config/ contains the correct server name (if you host your own Elasticsearch server).
<p>Deploy is throwing errors.</p>	<p>This error might be due to one or more search definitions in deployed state, but unavailable in Elasticsearch.</p>	<ol style="list-style-type: none"> 1. Check if any other search definition is already deployed. If so, select the search definition and click Report Sync. If the definitions are unavailable in the server, reset the definitions so that you can deploy them again. 2. If you do not want to do Step 1 and you need to deploy only a few definitions, then in the Index Settings, click the override option and change the Index name to the search definition name and deploy.

Error	Possible Cause	Resolution
<p>Index schedule status shows success, but search results are not displayed.</p>	<p>This message may be due to security of the search user or the security call back may not be responding with the correct values.</p>	<p>Resolution: 1. 2. 3. If the security response is empty or does not contain the security values applicable then revisit the security given to the user.</p> <ol style="list-style-type: none"> <li data-bbox="1073 426 1464 590">1. Ensure that some documents are indexed in Elasticsearch. Use the Elasticsearch Interact page to check whether documents are indexed in Elasticsearch. <li data-bbox="1073 625 1464 789">2. If security is enabled on a search definition, then set the IB logging to 5 and then perform the test again. Verify whether a security call back took place. <li data-bbox="1073 825 1464 951">3. If the security response is empty or does not contain the required security values, then check the security applied for the user.

Error	Possible Cause	Resolution
<p>Unable to deploy a search definition due to validation_exception.</p> <p>Sample error message:</p> <pre> Unable to deploy\ index se arch definition <search de finition name> due to erro r Index creation Failed : => validation_exception (262,=> 2012) PTSF_ES.ESIMPL.AdminServic e.OnExecute Name:AddIndexS ettings PCPC:25544 State=> ment:345 Called from:PTSF_ES.ESIMPL .AdminService.OnExecute Na me:AddPSFTSource Statemen t:902 Called from:PTSF_ES.ESIMPL .AdminService.OnExecute Na me:AddPSFTSources Stateme nt:909 Called from:PTEM_CONFIG.PT SFAdministerSearch.OnExecu te Name:deployPeoplesoftSo urce Statement:519 Called from:PTEM_CONFIG.PT SFAdministerSearch.OnExecu te Name:deploySearchDefini ton Statem (0,0) PTSF_ES.=> ESIMPL.AdminService.OnExec ute Name:AddIndexSettings PCP C:25664 Statement:350 </pre>	<p>This error is due to the soft limit on the number of shards to 1000, which is introduced in Elasticsearch 7.0.</p>	<p>Resolution - increase the number of shards per node as required.</p> <p>To increase the number of shards per node, for example, to 1500, use this command:</p> <pre> CURL -XPUT http://<eshost>=> :<port>/_cluster/settings </pre> <p>and in the body type the following line:</p> <pre> {"persistent":{"cluster.ma x_shards_per_node":1500}} </pre>

Managing Kibana Services

This appendix describes how you can manage the Kibana service.

Managing Kibana Service on Linux

After you install Kibana using the delivered DPK, you can add a Kibana service, verify that the service is created, or delete the service.

Adding Kibana as a Service on Linux

To add a 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 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 the `systemctl` command:

```
systemctl status kibana
```

This should give an output with the status "active (running)," as shown in the following 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
```

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 the following sample:

```
kibana.service - Kibana
Loaded: loaded (/etc/systemd/system/kibana.service; enabled; vendorpreset: dis⇒
abled)
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
```

Removing the Kibana Service on Linux

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
```

Troubleshooting Kibana Configuration

This appendix describes how you can troubleshoot issues in the PeopleSoft implementation of Kibana.

Kibana Menu is not Displayed Under Reporting Tools

Access to the Kibana menu is controlled through the Kibana Administrator role. Therefore, only users with the Kibana Administrator role can view the Kibana menu and access Kibana.

See [Accessing Kibana From PeopleSoft](#).

Errors on Kibana Visualization Tiles

If you see the following errors on the Kibana visualization tiles, you should follow the instructions provided in the corresponding section to fix the issue:

- Access denied error message

- Internal error message
- Site not found error message
- Create new index pattern error message

Each of these errors and their solutions are described here:

Access Denied Error Message

If a user sees the Access Denied error message on a Kibana visualization tile, follow these instructions to resolve the issue:

- As a first step, refresh the browser or clear the cache on the browser.
- Verify that the call back information on the Search Instance page is set correctly.
 1. After modifying the call back information, save the page.
 2. Select the Validate button to ensure that the validation returns a positive message.

See [Creating Search Instances](#).

- Verify that the `orcl_es_auth` index exists in Elasticsearch.

Internal Error Message

If a user sees the Internal Error message on a Kibana visualization tile, follow these instructions to resolve the issue:

- Verify that the search definition associated with the tile is deployed and indexed.
- Verify that the Kibana dashboard is deployed.

See [Managing a Delivered Dashboard](#).

Site Not Found Error Message

If a user sees the Site Not Found error message on a Kibana visualization tile, follow these instructions to resolve the issue:

- Verify whether the Kibana server is down.
- Verify whether Kibana is configured correctly.

If the environment is SSL enabled, ensure that end to end SSL configuration is successful. However, it is advisable to verify the Kibana configuration in a non-SSL environment first before configuring Kibana in an SSL enabled environment.

Create New Index Pattern Error Message

If a user sees the Create New Index Pattern error message on a Kibana visualization tile, follow these instructions to resolve the issue:

- Verify that the Kibana dashboards are deployed.

- Verify whether an index pattern is specified as default on the Management page in Kibana.

See [Creating a Visualization](#).

Certificate Error When Launching Kibana Visualization Tiles in the Chrome Browser

The certificate error is displayed because the subject alternate name (SAN) attribute is missing in the SSL certificate of Elasticsearch and Kibana. You should add the SAN attribute when setting up the keystore and private key and when you create the certificate signing request.

Example of SAN attribute: `-ext SAN=dns:<domain_name>`. Note that `<domain_name>` should be the exact host name.

The SAN attribute should be included in the script as shown in the following examples:

Elasticsearch SSL Certificate

Set up keystore and private key:

```
keytool -genkey -alias node_alias -keystore ES_HOME/config/keystore/mykeystore.jks ⇒
-keyalg RSA -keysize 2048 -validity 90 -ext SAN=dns:<domain_name>
```

Create Certificate Signing request (CSR):

```
keytool -certreq -alias node_alias -keystore ES_HOME/config/keystore/mykeystore.jks⇒
-file /home/certs/mycsr.csr -keyalg rsa -ext SAN=dns:<domain_name>
```

Kibana SSL Certificate

Set up keystore and private key:

```
keytool -genkey -alias alias1 -keystore KIBANA_HOME/bin/mykeystore.jks -keyalg RSA ⇒
-keysize 2048 -validity 712 -ext SAN=dns:<domain_name>
```

Create Certificate Signing request:

```
keytool -certreq -alias alias1 -keystore KIBANA_HOME/bin/mykeystore.jks -file D:\ca⇒
\mycsr.csr -keyalg rsa -ext SAN=dns:<domain_name>
```

Related Links

[Configuring SSL between PeopleSoft and Elasticsearch](#)

[Configuring SSL for Kibana](#)

Unable to Save a Visualization

If you are creating a visualization and unable to save a visualization, you should ensure that you have either of the following privileges:

- The Search Administrator role is assigned to the user.

- The create access privilege is assigned to the user on the Kibana Privileges page. See [Specifying User Privileges](#).

