
PeopleTools 8.57: Search Technology

March 2020

PeopleTools 8.57: Search Technology

Copyright © 1988, 2020, Oracle and/or its affiliates. All rights reserved.

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

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

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

U.S. GOVERNMENT END USERS: Oracle programs, including any operating system, integrated software, any programs installed on the hardware, and/or documentation, delivered to U.S. Government end users are "commercial computer software" pursuant to the applicable Federal Acquisition Regulation and agency-specific supplemental regulations. As such, use, duplication, disclosure, modification, and adaptation of the programs, including any operating system, integrated software, any programs installed on the hardware, and/or documentation, shall be subject to license terms and license restrictions applicable to the programs. No other rights are granted to the U.S. Government.

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

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

Intel and Intel Xeon 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, Opteron, the AMD logo, and the AMD Opteron 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 <http://www.oracle.com/pls/topic/lookup?ctx=acc&id=docacc>.

Access to Oracle Support

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

Contents

Preface: Preface.....	ix
Understanding the PeopleSoft Online Help and PeopleBooks.....	ix
PeopleSoft Hosted Online Help.....	ix
Locally Installed Help.....	ix
Downloadable PeopleBook PDF Files.....	ix
Common Help Documentation.....	x
Field and Control Definitions.....	x
Typographical Conventions.....	x
ISO Country and Currency Codes.....	xi
Region and Industry Identifiers.....	xi
Translations and Embedded Help.....	xii
Using and Managing the PeopleSoft Online Help.....	xii
Understanding PeopleSoft Search Technology.....	xii
PeopleTools Related Links.....	xii
Contact Us.....	xiii
Follow Us.....	xiii
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.....	27
PeopleSoft Search Features.....	29
Global Search.....	29
Search Pages.....	29
Chapter 3: Working with Elasticsearch Clusters.....	31
Understanding Clusters in Elasticsearch.....	31
Setting Up a Cluster During the Elasticsearch Installation.....	32
Naming the Cluster.....	32
Minimum Number of Master Nodes.....	32
List of Host Names where Elasticsearch will be Installed or is Installed.....	33
High-level Steps to Set Up a Cluster.....	33
Adding a New Node to an Existing Cluster.....	34
Verifying the Cluster and its Nodes.....	35
Using PeopleSoft Health Center.....	35
Using Elasticsearch Interact Page.....	35
Setting the Number of Replicas.....	36
Starting and Stopping a Cluster.....	36
Starting and Stopping a Node.....	36
Removing a Node from a Cluster.....	37

Managing Large Indexes.....	38
Chapter 4: Defining Search Definition Queries.....	39
Understanding Search Definition Queries.....	39
Defining Search Definition Queries with PeopleSoft Query.....	40
Adding a “Last Updated” Field to Records.....	40
Specifying a List of Fields to Index.....	41
Creating a Drilling URL.....	41
Creating a Prompt for the Last Modified Field.....	43
Defining Prompt Criteria.....	44
Defining Search Definition Queries with Connected Query.....	45
Defining a Deletion Query.....	45
Working with Images.....	45
Testing Your Search Definition Query.....	46
Chapter 5: Creating Query and Connected Query Search Definitions.....	49
Creating Search Definitions.....	49
Specifying General Settings.....	49
Adding Text for the Title and Summary.....	52
Inserting Bind Variables into Title\Summary Text.....	53
Mapping Search Attributes.....	53
Working with Attachment Properties.....	57
Understanding Attachment Processing in Elasticsearch.....	57
Specifying Attachments in a Search Definition.....	58
Specifying Attachment URL Properties.....	59
Specifying Static URL ID Attachment Properties.....	60
Specifying Dynamic URL ID Attachment Properties.....	62
Specifying Inline Attachment Properties.....	63
Managing Search Definition Security.....	63
Setting Source Level Security.....	63
Setting Document Level Security.....	64
Working With Advanced Settings.....	67
Mapping Components to Search Definitions.....	71
Viewing Search Attributes.....	72
Chapter 6: Creating File Source Search Definitions.....	75
Understanding File Source Search Definitions.....	75
Specifying File Source General Settings.....	75
Specifying Document Types.....	78
Chapter 7: Creating Web Source Search Definitions.....	83
Understanding Web Source Search Definitions.....	83
Specifying Web Source General Settings.....	83
Specifying Document Types for Web Source.....	85
Specifying Proxy Parameters in Elasticsearch.....	88
Chapter 8: Creating Search Categories.....	89
Understanding Search Categories.....	89
Specifying General Search Category Settings.....	89
Selecting Advanced Search Field Settings.....	91
Selecting Facet Settings.....	92
Creating Associated Facets.....	94
Defining Advanced Settings for String Facets.....	95
Defining Advanced Settings for Number Facets.....	96
Defining Advanced Settings for Date Facets.....	99
Identifying Custom Search Pages.....	102

Working With Display Fields.....	103
Chapter 9: Administering PeopleSoft Search Framework.....	105
Understanding PeopleSoft Search Framework Administration.....	105
Setting Up Local Nodes.....	105
Specifying Integration Broker Gateway.....	105
Setting Up Target Locations.....	105
Working With Search Instances.....	106
Creating Search Instances.....	106
Modifying Search Instances.....	111
Deleting Search Instances.....	111
Sharing a Search Instance Among Multiple PeopleSoft Application Systems.....	111
Administering Elasticsearch Using the Elasticsearch Interact Page.....	112
Setting Namespace Aliases.....	115
Administering Search Instances.....	117
Assigning Users to a Search Instance for Validating Indexes.....	118
Creating a Search Instance in Automated Configuration Manager.....	118
Enabling Global Search.....	119
Managing General Search Options.....	120
Administering Search Definitions and Search Categories.....	126
Understanding Search Definition Administration.....	126
Working with Search Definition Deployment.....	126
Working with Search Categories.....	132
Working with Search Indexes.....	135
Understanding the Index Build Process.....	135
Building Search Indexes.....	136
Viewing Search Index Build Process Details.....	139
Viewing the Exception Details in Attachment Processing.....	141
Partitioning Application Data in the Search Index Build Process.....	142
Using DirectTransfer.....	143
Understanding DirectTransfer.....	143
Configuration.....	144
System Considerations.....	145
Error Handling.....	146
Using DirectTransfer with SSL Enabled Elasticsearch.....	148
Verifying Domain Status.....	149
Verifying Integration Broker Queue Status.....	150
Managing Search Context.....	150
Understanding Search Contexts.....	150
Defining Search Contexts.....	150
Viewing Search Contexts.....	151
Viewing Node in Network.....	152
Importing Remote Search Groups.....	152
Creating the Attachment URL ID List.....	154
Implementing Report Repository Search.....	155
Working with Search Framework Definitions During Upgrades.....	156
Boosting Score of Search Results.....	156
Testing Score Boosting.....	158
Chapter 10: Working with PeopleSoft Search Framework Utilities.....	159
Using the Search Test Page.....	159
Understanding the Search Test Page.....	159
Testing Search Categories.....	161

Testing Custom Search Attributes.....	161
Testing Filter Settings.....	161
Testing Facet Requests.....	162
Testing Grouping and Sorting Options.....	162
Testing Additional Parameters.....	162
Downloading Search Data.....	163
Running Diagnostics.....	164
Viewing Event Logs.....	166
Chapter 11: Working with PeopleSoft Search Framework Security Features.....	169
Understanding Search Framework Security.....	169
Applying PeopleSoft Permissions.....	169
Working with Authentication and Authorization.....	170
Authentication and Authorization in Elasticsearch.....	171
Configuring SSL between PeopleSoft and Elasticsearch.....	172
Configuring SSL for Elasticsearch.....	172
Setting Up Role-Based Search Group Access.....	174
Chapter 12: Working with PeopleSoft Search.....	175
Understanding PeopleSoft Search.....	175
Working with Search Pages.....	176
Understanding Search Pages.....	176
Enabling Keyword Search.....	178
Developing for Search Pages.....	179
Working with Keyword Search Modes.....	182
Working with Global Search.....	183
Understanding Global Search.....	184
Working with Search Groups and Search Contexts.....	184
Working With the Portal Registry Search.....	185
Working with Global Search in Fluid User Interface.....	187
Working with Search Operators in PeopleSoft Search.....	188
Working with the Search Results.....	190
Working With Facets.....	191
Viewing Search Results With Grid Format and List Format.....	193
Working with Related Actions.....	194
Working with Search Results in Fluid User Interface.....	196
Working with Search Results in Global Search.....	196
Working with Real-time Component Search Results.....	201
Working with Keyword Component Search Results.....	204
Working with Search Results on a Touch Friendly Device.....	205
Chapter 13: Configuring High Availability in PeopleSoft Search Framework.....	211
Understanding High Availability Feature in PeopleSoft Search Framework.....	211
Fail-over Mechanism.....	211
Implementing Fail-over Mechanism.....	211
Implementing Fail-over Mechanism Using Multiple Nodes Feature of Elasticsearch.....	212
Implementing Fail-over Mechanism Using PeopleSoft Search Framework’s Multiple Search Instance Feature.....	212
Chapter 14: Monitoring PeopleSoft Search Framework and Elasticsearch Using Kibana.....	215
Understanding Monitoring of Search Framework and Elasticsearch.....	215
Kibana Terminology.....	215
Setting Up Kibana in PeopleSoft Search Framework.....	216
Setting Up Kibana.....	216
Using the Elasticsearch Monitoring Server Page.....	217

Using Dashboards and Visualizations in Kibana.....	220
Understanding Dashboard and Visualizations.....	220
Delivered Dashboards and Visualizations.....	220
Common Elements on a Visualization.....	221
Using the Indexing Metrics Dashboard and Visualizations.....	222
Using the Indexing Summary Dashboard and Visualizations.....	223
Using the System Metrics Dashboard and Visualization.....	224
Creating a Visualization.....	225
Cloning a Dashboard.....	226
Creating a Visualization.....	226
Chapter 15: Monitoring Elasticsearch Cluster Using PeopleSoft Health Center.....	229
Understanding Elasticsearch Cluster.....	229
Viewing the Elasticsearch Cluster Dashboard.....	229
Application Dashboard.....	229
Cluster Details Page.....	231
Viewing the Node Dashboard.....	234
Summary Section.....	234
JVM Section.....	235
File System Section.....	236
Index Activity Section.....	237
Cache Activity Section.....	238
Network Section.....	239
Appendix A: Users and Roles in Elasticsearch.....	241
Creating Users and Assigning Roles in Elasticsearch.....	241
Appendix B: Troubleshooting.....	243
Troubleshooting Elasticsearch Configuration.....	243
Performing Administrative Tasks.....	243
Handling Common Errors.....	244

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.

PeopleSoft Hosted Online Help

You access the PeopleSoft Online Help on Oracle's PeopleSoft Hosted Online Help website, which enables you to access the full help website and context-sensitive help directly from an Oracle hosted server. The hosted 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, because that documentation is now incorporated into the hosted website content. The Hosted Online Help website is available in English only.

Note: Only the most current release of hosted online help is updated regularly. After a new release is posted, previous releases remain available but are no longer updated.

Locally Installed Help

If you are setting up an on-premises PeopleSoft environment, and your organization has firewall restrictions that prevent you from using the Hosted Online Help website, you can install the PeopleSoft Online Help locally. If you install the help locally, you have more control over which documents users can access and you can include links to your organization's custom documentation on help pages.

In addition, if you locally install the PeopleSoft Online Help, you can use any search engine for full-text searching. Your installation documentation includes instructions about how to set up Elasticsearch for full-text searching. See *PeopleSoft 9.2 Application Installation* for your database platform, "Installing PeopleSoft Online Help." If you do not use Elasticsearch, see the documentation for your chosen search engine.

Note: See [Oracle Support Document 2205540.2 \(PeopleTools Elasticsearch Home Page\)](#) for more information on using Elasticsearch with PeopleSoft.

Note: Before users can access the search engine on a locally installed help website, you must enable the Search field. For instructions, go to your locally installed PeopleSoft Online Help site and select About This Help >Managing Locally Installed PeopleSoft Online Help >Enabling the Search Button and Field in the Contents sidebar.

Downloadable PeopleBook PDF Files

You can access downloadable PDF versions of the help content in the traditional PeopleBook format. 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.

Typographical Convention	Description
& (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

Click the Help link in the universal navigation header of any page in the PeopleSoft Online Help to see information on the following topics:

- What's new in the PeopleSoft Online Help.
- PeopleSoft Online Help accessibility.
- Accessing, navigating, and searching the PeopleSoft Online Help.
- Managing a locally installed PeopleSoft Online Help website.

Understanding PeopleSoft Search Technology

This document describes the administration of the Search Framework as well as the development tasks associated with enabling and implementing PeopleSoft Search features. These features are provided by PeopleTools and are designed to run against the Elasticsearch and the Oracle Secure Enterprise (SES) search engines.

PeopleTools Related Links

[PeopleTools 8.57 Home Page](#)

[PeopleTools Elasticsearch Home Page](#)

"PeopleTools Product/Feature PeopleBook Index" (PeopleTools 8.57: Getting Started with PeopleTools)

[PeopleSoft Hosted 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 PSOFT-INFODEV_US@ORACLE.COM. Please include release numbers for the PeopleTools and applications that you are using.

Follow Us



[Facebook.](#)



[YouTube](#)



[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 8.57 is integrated with Elasticsearch version 6.1.2 and Kibana 6.1.2.

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.

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 8.57 hardware and software requirements, see My Oracle Support, Certifications.
Elasticsearch	See <i>PeopleSoft Deployment Packages for Elasticsearch Installation on My Oracle Support (Doc ID 2205540.2)</i> . See Elasticsearch Reference [6.1] (www.elastic.co).

Required Item	Documentation
Kibana	See <i>PeopleSoft Deployment Packages for Elasticsearch Installation on My Oracle Support (Doc ID 2205540.2)</i> . See Kibana User Guide [6.1] (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" (PeopleTools 8.57: 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.
- Monitoring interface (PeopleSoft Health Center) that provides diagnostic data of the Elasticsearch cluster.
- 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 elasticsearch.yml 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>
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>

Elasticsearch Terminology	Description
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.

Item	Elasticsearch Functionality
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.
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.

Item	Elasticsearch Functionality
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 [6.1].</p>
<p>Custom Scoring</p>	<p>In Elasticsearch, custom scoring is supported through Search Framework APIs.</p>

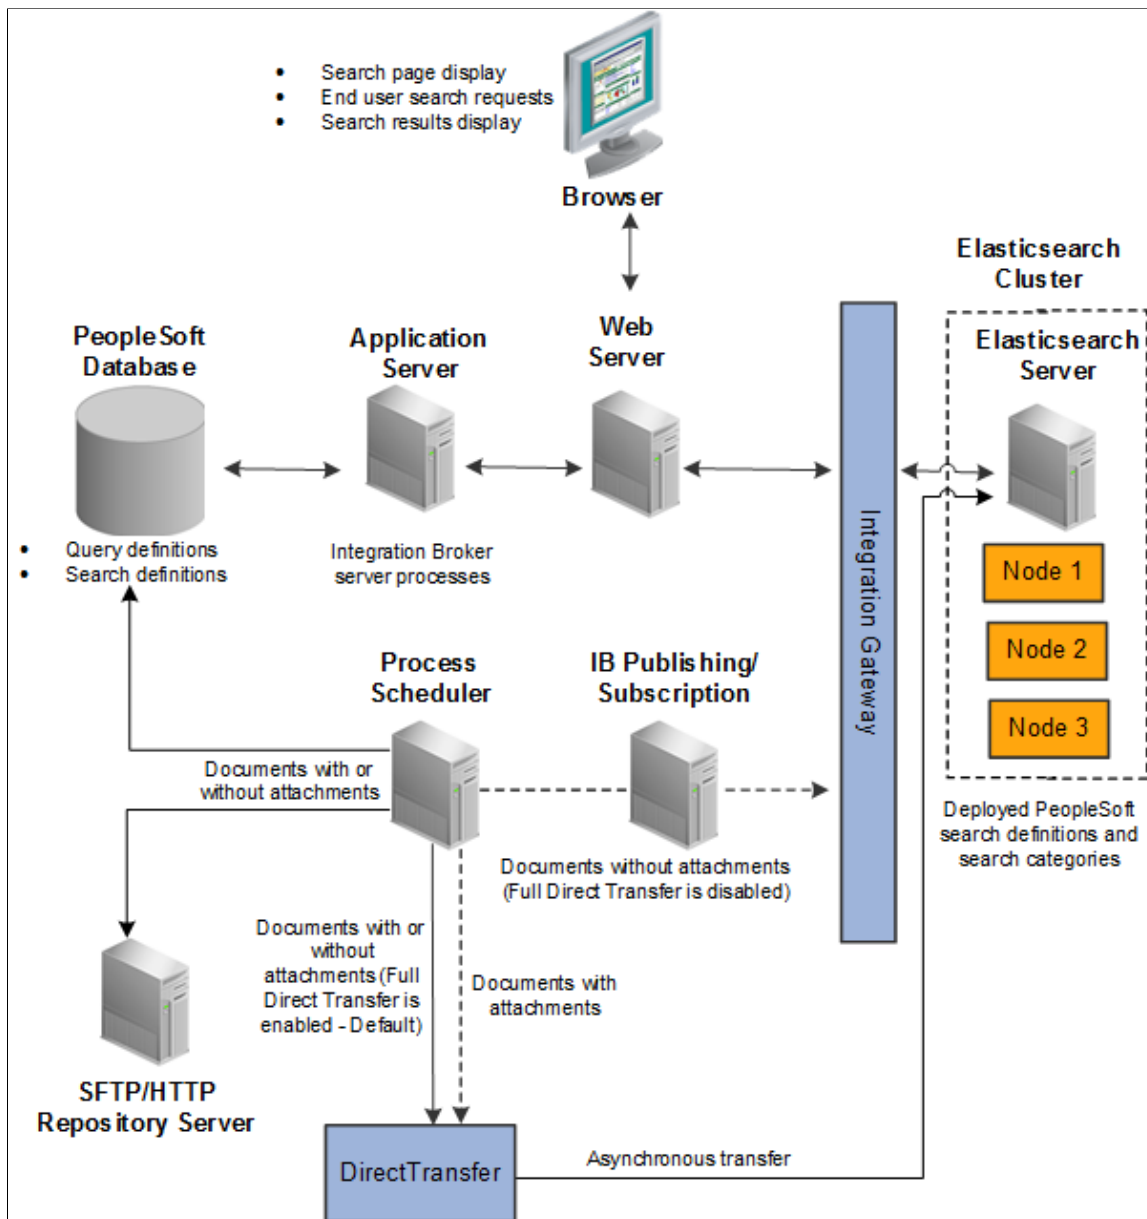
Item	Elasticsearch Functionality
Synonyms	<p>Elasticsearch uses a file (synonyms_grammar.txt) to store synonyms or any abbreviations that you want to include. If synonyms provided are not adequate for your installation, you can modify this file. The file is located at config/analysis/synonym_grammar.txt.</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.

Image: PeopleSoft server architecture with Elasticsearch

The following illustration depicts the PeopleSoft server architecture connected to Elasticsearch using Integration Broker and DirectTransfer.

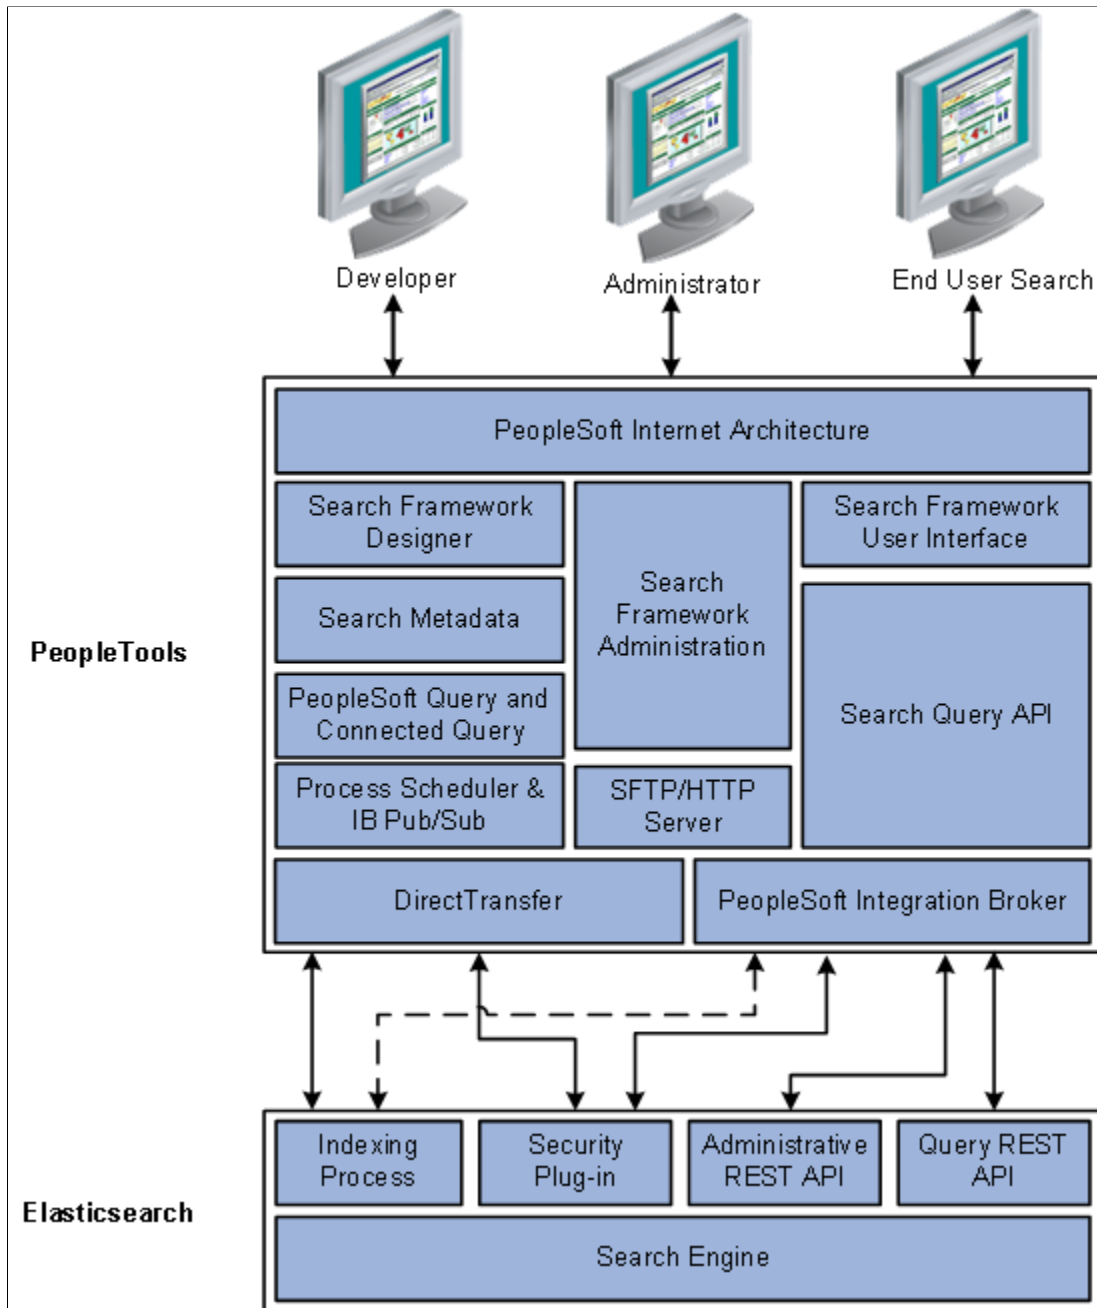


The web server and application server (the PeopleSoft Internet Architecture) provide the interface for the design, administration, and end user access. The PeopleSoft database stores the PeopleSoft queries that define the search data as well as the search definition meta data. PeopleSoft uses the Process 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.

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

Image: PeopleTools elements and Elasticsearch elements

The following illustration depicts the PeopleTools elements and the Elasticsearch elements interacting to enable the PeopleSoft Search Framework features.



The following table describes the PeopleTools elements.

PeopleTools Element	Description
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, and scheduling search index builds.
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.

The following table describes the Elasticsearch elements.

Elasticsearch Element	Description
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>

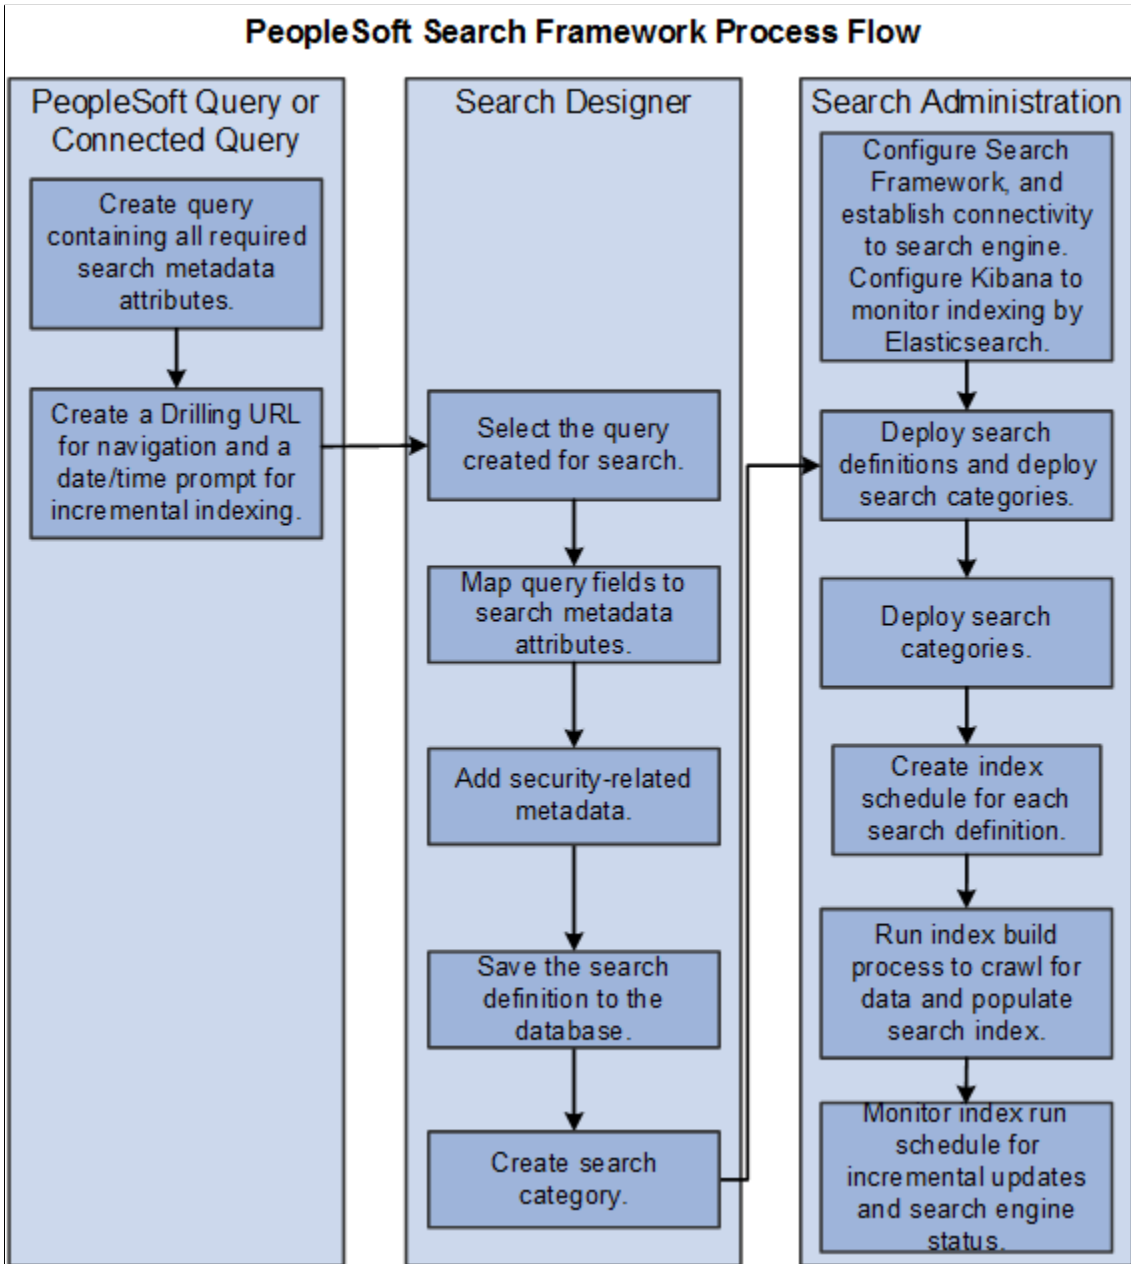
Elasticsearch Element	Description
Security Plug-in	This module gets invoked from the search engine for a user who initiates a search from an application. The search security framework is responsible for authentication, search user validation, and authorization of the search query request. The search 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 attribute 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.
Administrative REST APIs	The administrative tasks include various processes, such as, deploying, undeploying, report synchronization, and updating. The http methods of GET, PUT, POST, and DELETE are used while sending the REST API requests to Elasticsearch.
Query REST API	The Query plug-in works in conjunction with the PeopleCode Search API for application developers to compose search queries and execute them. The Query plug-in exposes various search features like filtering, sorting, grouping, restricting search scope, pagination, and so on.
Search Engine	PeopleSoft Search Framework works with Elasticsearch search engine. Elasticsearch is an open-source search engine, which is based on Apache Lucene™.

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.

Image: PeopleSoft Search Framework process flow

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](#).

Image: Global Search bar

This example illustrates the search icon and 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.

Image: Keyword Search page

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

Enter any information you have and click Search. Leave fields blank for a list of all values.

Find an Existing Value | Keyword Search

Note: Keyword Search will return data updated less than 1 hour ago (08/15/2011 1:57:19PM)

Search Criteria

Keywords clerk

Search | [Advanced Search](#)

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 (total number of master-eligible nodes / 2 + 1).

Note: This feature of specifying multiple nodes during the installation of Elasticsearch are contained in the DPKs (ESK-DPK-LNX-6.1.2_01 for Linux and ESK-DPK-WIN-6.1.2_01 for Windows).

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.zen.ping.unicast.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](#).

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.

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

discovery.zen.minimum_master_nodes

Enter the value that you provided in Step 1.

discovery.zen.ping.unicast.hosts

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](#).

3. 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.

4. 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 in two ways:

- Using PeopleSoft Health Center.
- Using the Elasticsearch Interact page.

Using PeopleSoft Health Center

On the Application Dashboard, the ES Cluster Summary section displays the name of the cluster and the number of nodes that are present in the cluster. You can view further details of the cluster including information on shards, replicas, documents on the Cluster Details dashboard.

See [Viewing the Elasticsearch Cluster Dashboard](#).

If you require instructions on how to configure PeopleSoft Health Center, see "Configuring PeopleSoft Health Center" (PeopleTools 8.57: Performance Monitor).

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.

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 [Viewing the Elasticsearch Cluster Dashboard](#).

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"
}
```

- 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.
- 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
}
```

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

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

- 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.

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 you use 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" (PeopleTools 8.57: Query)

"Creating New Queries" (PeopleTools 8.57: Query)

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 is used as an unique ID in Elasticsearch documents created by PeopleSoft. 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.

- 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.

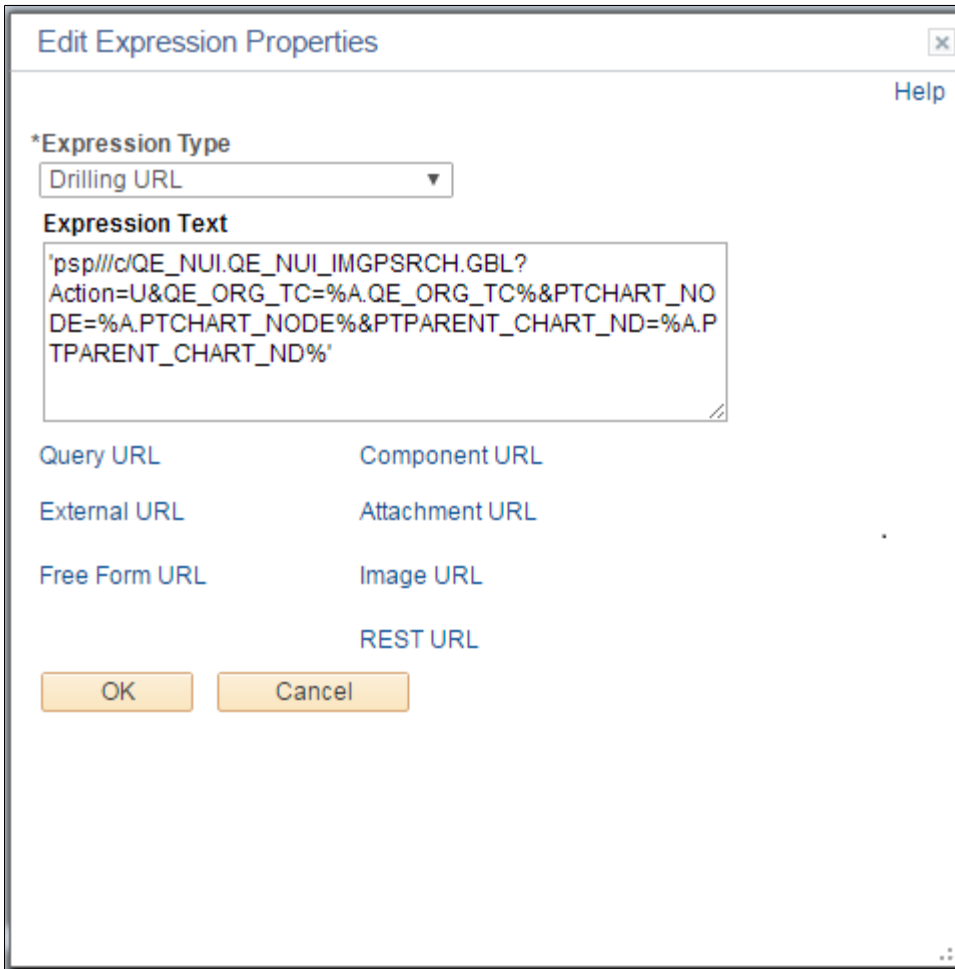
Note: If you do not specify the correct search keys, the system will launch the page you have specified, but the user would still need to add the keys manually to view the information associated with a search result, which is not the desired behavior.

- Click OK.
- 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.

Image: Edit Expression Properties dialog box

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.

Image: Edit Prompt Properties dialog box

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.

The screenshot shows the 'Edit Prompt Properties' dialog box with the following fields and controls:

- Field Name:** ROW_LASTMANT_DTTM
- *Heading Type:** RFT Short
- *Type:** Datetime
- *Format:** None
- Length:** 26
- Decimals:** (empty)
- *Edit Type:** No Table Edit
- *Unique Prompt Name:** BIND1
- Prompt Table:** (empty)
- Default Value:** (empty)
- Buttons:** OK, Cancel

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
Expression 2	<ul style="list-style-type: none"> Prompt Your prompt

For example:

Image: Edit Criteria Properties dialog box

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

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.

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.

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.

Image: Query Manager Run tab

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

Records	Query	Expressions	Prompts	Fields	Criteria	Having	Dependency	Transformations	View SQL	Run
View All Rerun Query Download to Excel Download to XML										
										First 1-100 of 16694 Last
	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, Search Definition.
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.
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, Search Definition.

Image: General Settings page

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.

The screenshot shows the 'General Settings' page for a search definition. At the top, there are tabs for 'General', 'Map Search Attributes', 'Attachment Properties', 'Security', 'Advanced', and 'Component Mapping'. The 'General' tab is selected. Below the tabs, the search definition name is 'PTPORTALREGISTRY', the description is 'Navigator', and the object owner ID is 'PeopleTools'. The 'Source Information' section shows the source type as 'Connected Query' and the source name as 'PORTALREGISTRY_CQ'. The 'Search Result Field Mapping' section includes fields for '*Last Modified Date Time' (A.LASTUPDDTTM), '*URL Link' (PORTAL_DRILLURL), and '*Image URL' (IMG_DRILLURL). It also has text areas for 'Title' and 'Summary' with corresponding buttons. At the bottom are 'Save' and 'Save As' buttons.

Search Definition

Displays the search definition’s name as specified when adding a new value.

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.

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	Select the query field containing the drilling URL defined for the query.
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>
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>
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 meta-data associated with the existing search definition is saved in the cloned search definition.</p>

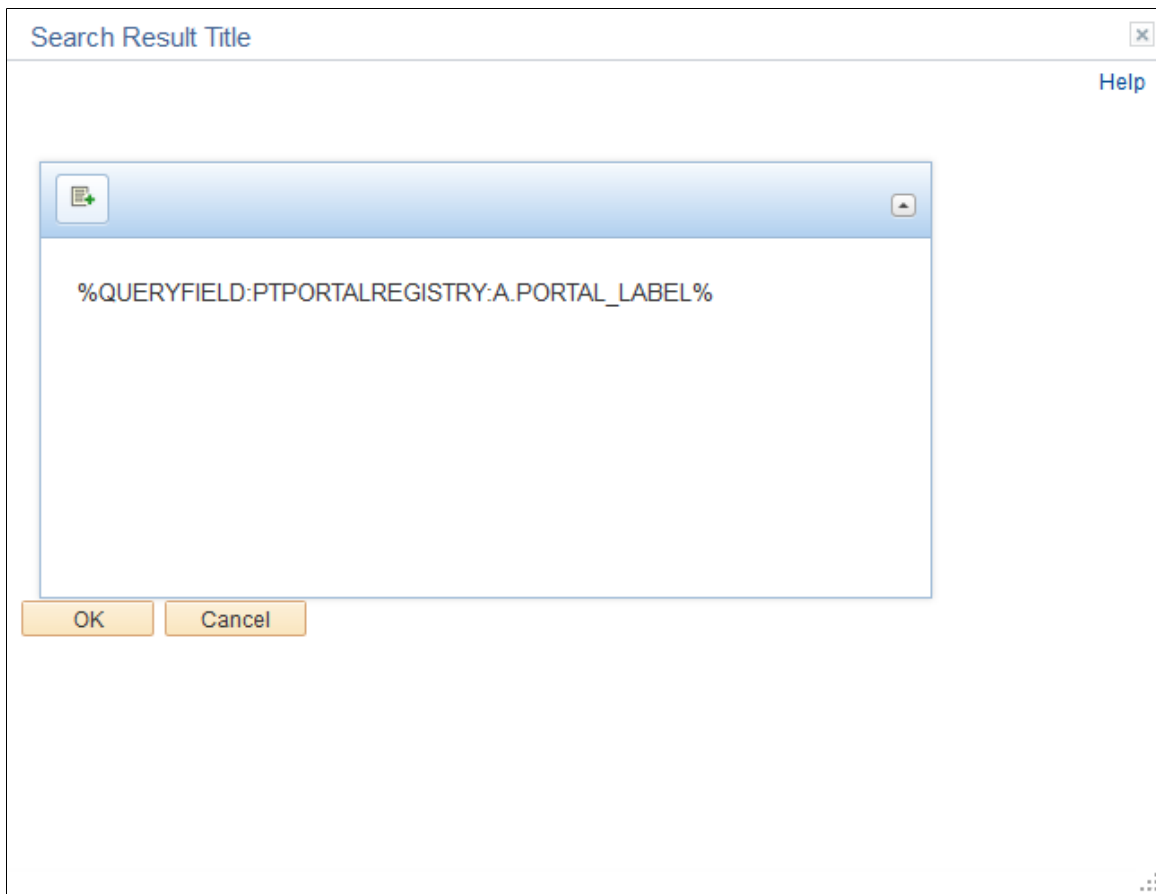
Note: The Save As button is not available when you create or define a new search definition.

Adding Text for the Title and Summary

In the Title\Summary dialog box 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.

Image: Search Result Title edit box

This example illustrates the fields and controls on the Search Result Title edit box.



Insert Bind Variables



Click to launch the Insert Bind Variables dialog box and insert bind variables of the following types:

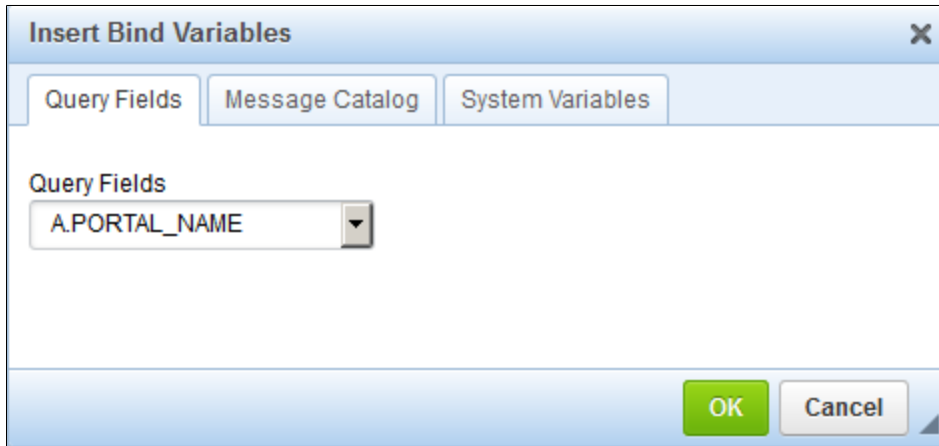
- Query fields
- Message catalog entries
- System variables

Inserting Bind Variables into Title\Summary Text

When inserting bind variables into the title or summary text, you use the icon at the top left corner to open the Insert Bind Variables dialog box, enabling you to select the bind variable type and the specific bind variable. You can insert multiple bind variables and bind variable types within the title\summary text, as needed.

Image: Insert Bind Variables dialog box

This example illustrates the fields and controls on the Insert Bind Variables dialog box. You can find definitions for the fields and controls later on this page.



Query Fields

Use the Query Fields drop-down list to select the desired fields from the underlying query of the search definition.

Message Catalog

Specify the Message Set and Message Number to identify the specific message to display. To determine the text that displays use these options:

- Message Catalog Text: Displays the Message Text edit box of the message catalog entry.
- Message Catalog Explanation: Displays the Description edit box of the message catalog entry.

System Variables

Use the System Variables drop-down list to select the desired variable to display.

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.

Image: Search Attribute Mapping page

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.

The screenshot shows the 'Search Attribute Mapping' page for search definition 'PTPORTALREGISTRY'. It features a table with columns: Query Name, Query Field Name, Field to Index, Attribute Name, Attribute Display Name, Is Faceted, Is Hierarchy, and Hierarchy Path. The 'Field to Index' column contains checkboxes, with the first four rows (rows 1-4) checked. The 'Attribute Name' column contains values like 'PORTAL_NAME', 'PORTAL_LABEL', 'DESCR254', etc.

Query Name	Query Field Name	Field to Index	Attribute Name	Attribute Display Name	Is Faceted	Is Hierarchy	Hierarchy Path
PTPORTALREGISTRY	A.PORTAL_NAME	<input checked="" type="checkbox"/>	PORTAL_NAME	Portal Name	<input type="checkbox"/>	<input type="checkbox"/>	
PTPORTALREGISTRY	A.PORTAL_OBJNAME	<input type="checkbox"/>			<input type="checkbox"/>	<input type="checkbox"/>	
PTPORTALREGISTRY	A.PORTAL_LABEL	<input checked="" type="checkbox"/>	PORTAL_LABEL	Portal Label	<input type="checkbox"/>	<input type="checkbox"/>	
PTPORTALREGISTRY	A.DESCR254	<input checked="" type="checkbox"/>	DESCR254	Long Description	<input type="checkbox"/>	<input type="checkbox"/>	
PTPORTALREGISTRY	A.OPRID	<input type="checkbox"/>			<input type="checkbox"/>	<input type="checkbox"/>	
PTPORTALREGISTRY	A.PORTAL_EFFDT	<input type="checkbox"/>			<input type="checkbox"/>	<input type="checkbox"/>	
PTPORTALREGISTRY	A.PORTAL_CREATION_DT	<input type="checkbox"/>			<input type="checkbox"/>	<input type="checkbox"/>	
PTPORTALREGISTRY	A.PORTAL_EXPIRE_DT	<input type="checkbox"/>			<input type="checkbox"/>	<input type="checkbox"/>	
PTPORTALREGISTRY	A.PORTAL_EXPIRES	<input type="checkbox"/>			<input type="checkbox"/>	<input type="checkbox"/>	
PTPORTALREGISTRY	A.LASTUPDDTTM	<input type="checkbox"/>			<input type="checkbox"/>	<input type="checkbox"/>	
PTPORTALREGISTRY	A.PORTAL_CREF_URLRT	<input type="checkbox"/>			<input type="checkbox"/>	<input type="checkbox"/>	
PTPORTALREGISTRY	A.PORTAL_URLTEXT	<input type="checkbox"/>			<input type="checkbox"/>	<input type="checkbox"/>	
PTPORTALREGISTRY	A.PORTAL_BASE_URI	<input type="checkbox"/>			<input type="checkbox"/>	<input type="checkbox"/>	
PTPORTALREGISTRY	A.PORTAL_CNTPRV_NAM	<input type="checkbox"/>			<input type="checkbox"/>	<input type="checkbox"/>	
PTPORTALREGISTRY	A.NODE_ITEMNAME	<input type="checkbox"/>			<input type="checkbox"/>	<input type="checkbox"/>	
PTPORTALREGISTRY	A.PT_URLTEXT	<input type="checkbox"/>			<input type="checkbox"/>	<input type="checkbox"/>	
PTPORTALREGISTRY	A.PORTAL_NODENAME	<input type="checkbox"/>			<input type="checkbox"/>	<input type="checkbox"/>	
PTPORTALREGISTRY	A.PORTAL_URLTEXT	<input type="checkbox"/>			<input type="checkbox"/>	<input type="checkbox"/>	
PTPORTALREGISTRY	PORTAL_DRILLURL	<input type="checkbox"/>			<input type="checkbox"/>	<input type="checkbox"/>	

Field to Index

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.

When selecting fields also consider fields that may help to group results or help distinguish security access.

Consider the following when you select fields for indexing:

- 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.

Note: This does not impact single value fields.

Attribute Name

Displays the pre-defined attribute name that identifies the query field in the search engine.

Note: The attribute name must be unique in a search definition.

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`.

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.

Attribute Display Name

Determines how the attribute will display to the end user in the search results.

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.

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.

Is Faceted

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.

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.

If a facet does not have a value, then the query is dynamically modified to return *No Value* as the title of the facet. This processing is performed only on character attributes; not on long and XLAT attributes.

Note: Facets based on XLAT fields should be changed to use XLAT Long/Short in the query.

Note: There is a data size limit of 2000 characters for faceted fields. The system truncates anything beyond 2000 characters.

Important! Values with data types of datetime cannot be used as facets.

Is Hierarchy

(Appears only for Connected Query.)

This parameter is used to define a “soft hierarchy” which is programmed in the query.

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.

This is different from the "hard" hierarchy discussed in Hierarchy Path.

Hierarchy Path

(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.

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.

This is considered a "hard" (predefined) hierarchy.

Example: Define Hierarchy Path

This is an example of defining a hierarchy path.

Image: Define Hierarchy Path page

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

Query Name: QE_PERS_DATA

Source Field: A.QE_BIRTHCOUNTRY

Hierarchy Separator: /

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.

Image: Example: Hierarchy Path

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](#).

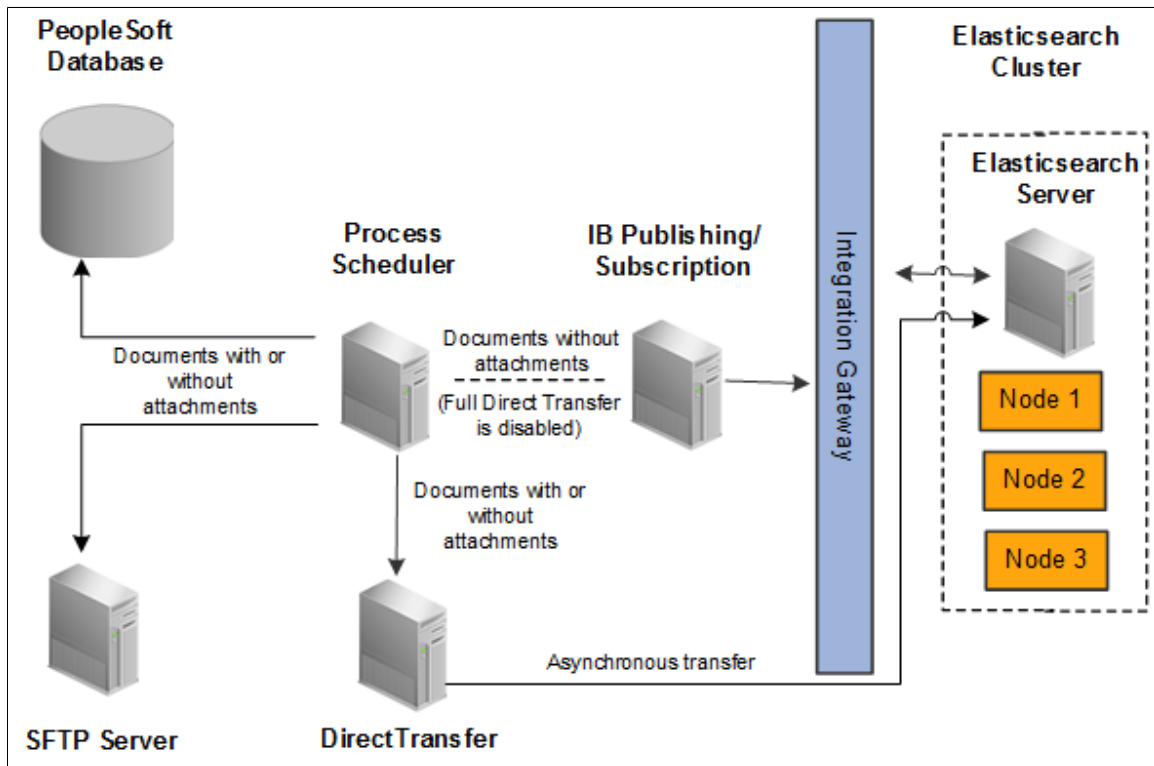
The search engine uses the ingest-attachment plug-in to extract the attachment contents. The attachment data is then indexed in the search engine.

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

- PeopleSoft database
- FTP
- SFTP
- HTTP

Image: PeopleSoft Search Framework architecture including file attachment repository

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.

Image: Attachment Properties page

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.

Attachment Type	Attachment URL	URL Identifier	Query Name	URLID Query Field	File Reference Path	File Path Query Field	File Name Query Field	Edit
1 Dynamic URLID			QE_PIASRCH_JOB_OPENINGS	DEPTNAME				Edit

Attachments

The Attachments grid contains attachments you specify to be indexed as part of your search definition. These are the types of attachments supported:

- Attachment URL
- Static URLID
- Dynamic URLID

The options for each attachment type are described later in this topic where each attachment type is discussed.

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.

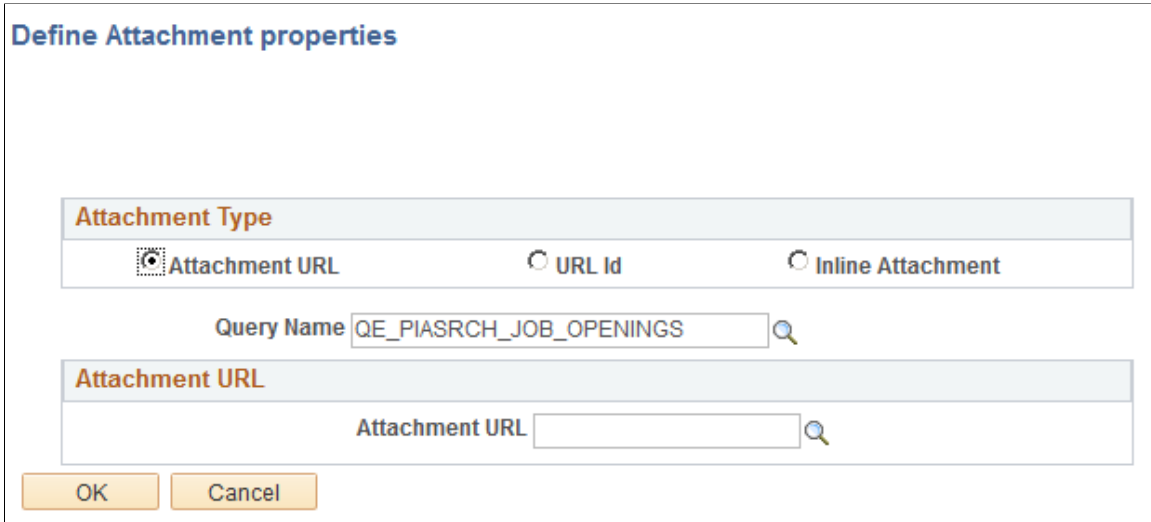
When using the Attachment URL type of attachment, the system uses the call-back credentials, defined for the search engine instance, to access the location of the attachment files. For example, if your attachments are stored in an FTP server, the FTP server needs to have the search engine call-back credentials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

Image: Define Attachment Properties: Attachment URL

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.



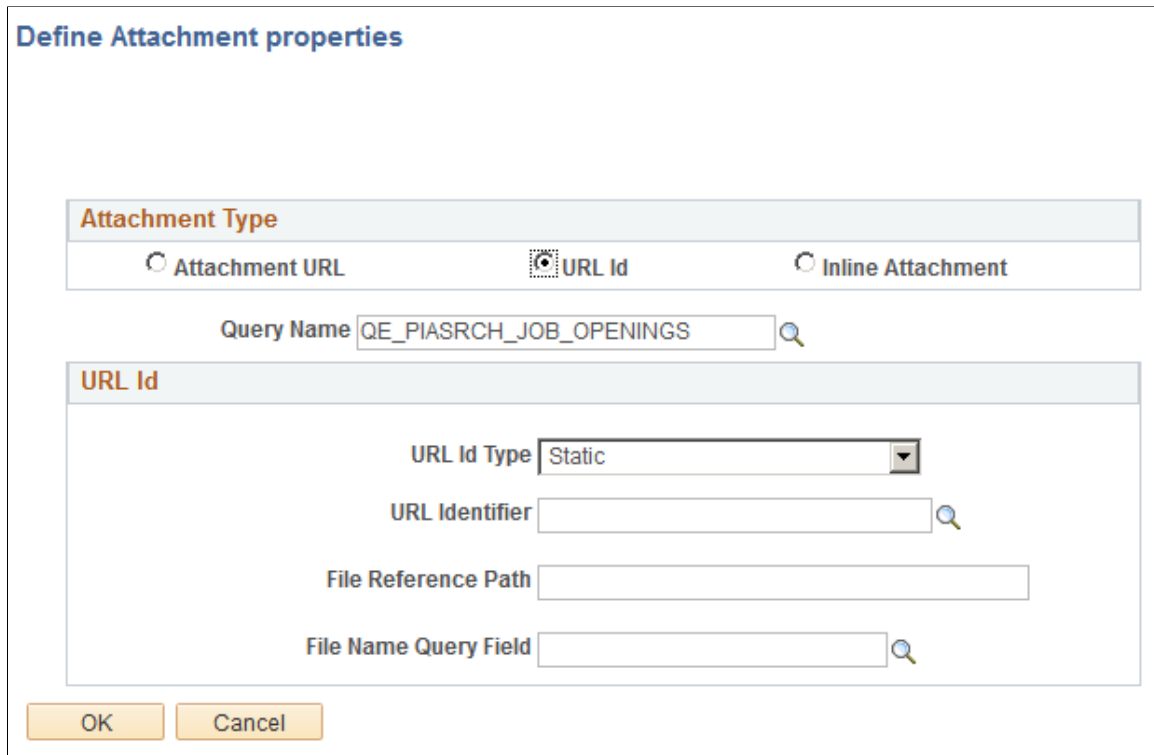
- Attachment URL** Select if the attachment is linked to the search definition through a drilling URL in the underlying query as specified on the Expression tab in PeopleSoft Query.
- Query Name** Select the query in which the Attachment URL drilling URL is defined.
- Attachment URL** Select the query 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.

Image: Define Attachment Properties: Static URL

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.



- 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 *Static*.

- 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.

Image: Define Attachment Properties: Dynamic URL

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.

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 *Dynamic*.

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.

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.

Image: Security page

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 screenshot shows the 'Security' tab selected in a navigation bar. Below the tabs, the 'Security Settings' section displays 'Search Definition' as 'ITEMDEFN' and 'Description' as 'Item Definition'. A 'Security Option' section contains three radio buttons: 'No Security' (which is selected), 'Source Level Security', and 'Document Level Security'.

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:

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.

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.

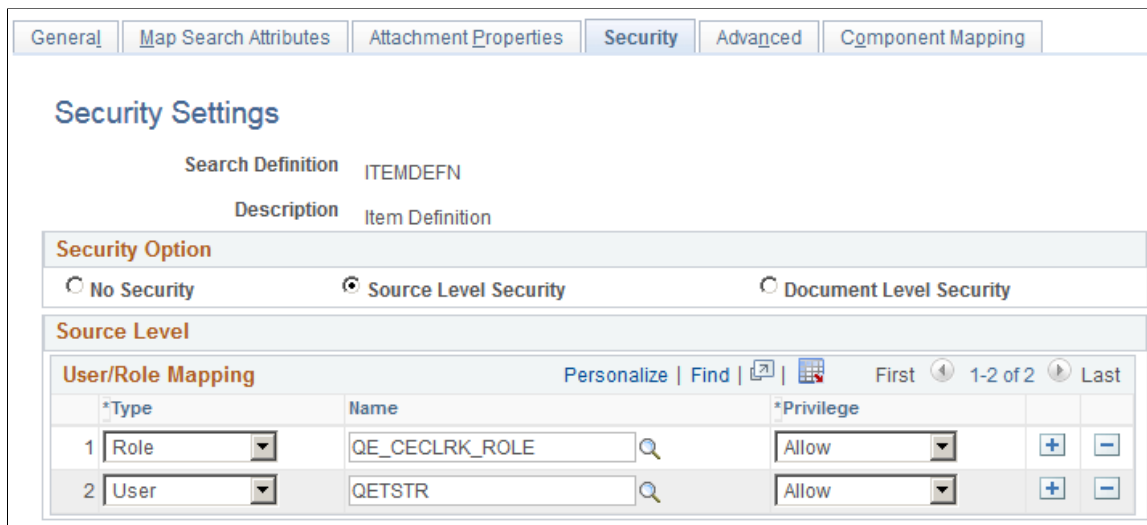
Note: This is generally referred to as row-level security in PeopleSoft applications.

Setting Source Level Security

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

Image: Security page - Source Level Security options

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.



Type Select *Role* or *User* depending on the scope of your intended access restriction.

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.

- 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.

Image: Security page: Document Level Security option

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.



General | Map Search Attributes | Attachment Properties | **Security** | Advanced | Component Mapping

Security Settings

Search Definition PTPORTALREGISTRY
Description Navigator

Security Option

No Security Source Level Security Document Level Security

Document Level

Query Field Mapping Personalize | Find | First 1 of 1 Last

Query Name	Source Field	Attribute Name	Privilege
1 PTPORTALREG_SEC	A.PORTAL_SECTYPE	PORTAL_SECTYPE	Allow

Document Filter App Class

Package Name: PTPORTALREGSEC
Path: :
Class ID: UserPermission

Document Level

Query Name

Select the name of the query or connected query containing the fields you want to use to restrict access.

Source Field

Select the field which will identify security values which will determine access.

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.

Privilege

Define the access privilege or restriction.

- 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.

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.

Note: If multiple attributes appear in the grid the system effectively inserts an AND clause between the items in the grid.

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.

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.

Image: Advanced Properties page

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.

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 type="text"/>		+	-

Full Indexing Criteria

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

Alternate search keywords

Define Query to Delete SBO

Query Name

Select the query that you've defined to remove orphaned search documents. This is your deletion query.

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.

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.

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.

See [Defining a Deletion Query](#).

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.

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.

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

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 \leftrightarrow Days field is 0, which is the default value, the Index Start Date is set to 01/01/1900.

Alternate Search Keywords

Alternate search keywords

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.

This field is translatable. If supporting multiple languages, the keywords need to be translated.

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.

Image: Component Mapping page

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.

Component Mapping																			
Search Definition		ITEMDEFN																	
Description		Item Definition																	
<table border="1"> <thead> <tr> <th colspan="4">Component Mapping</th> </tr> <tr> <td>Personalize</td> <td>Find</td> <td>View All</td> <td>1 of 1</td> </tr> <tr> <td>*Market</td> <td>*Component Name</td> <td>Description</td> <td>Search Criteria</td> </tr> </thead> <tbody> <tr> <td>1 Global</td> <td>MERCHANTCATBL</td> <td></td> <td></td> </tr> </tbody> </table>				Component Mapping				Personalize	Find	View All	1 of 1	*Market	*Component Name	Description	Search Criteria	1 Global	MERCHANTCATBL		
Component Mapping																			
Personalize	Find	View All	1 of 1																
*Market	*Component Name	Description	Search Criteria																
1 Global	MERCHANTCATBL																		

Market

Select the Market for the component.

Component Name	Select the component.
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 ATTRIBUTENAME:VALUE 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> <hr/>

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.

Image: Search Attributes page

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.

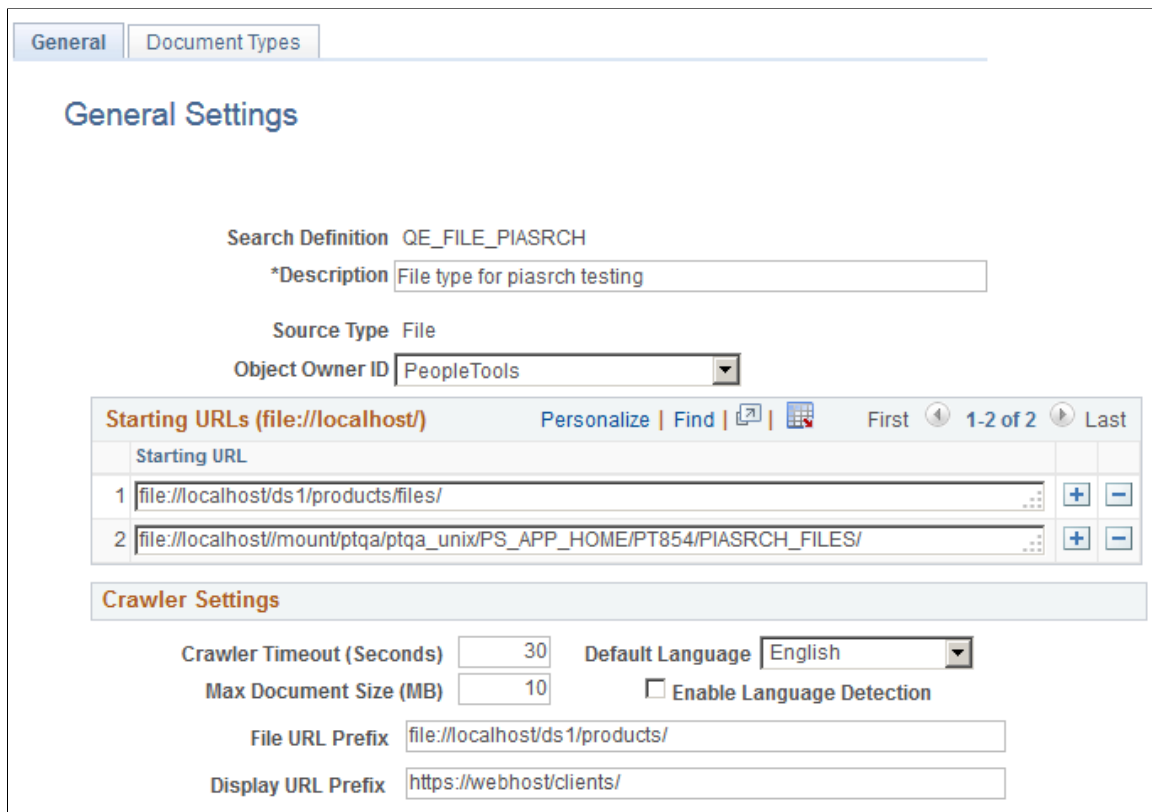
Specifying File Source General Settings

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.

Image: File Source: General Settings page

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.



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.

Starting URLs 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.

Note: The starting URL is not case sensitive.

Note: For the search engine crawler to access the URL, the file starting URL must be fully qualified, as in file://localhost/.

On UNIX the starting URL format is:

- To crawl local files use file://localhost/<directory structure>. For example:

file://localhost/recruitment/resume/

- To crawl from mounted file systems use `file://localhost/<mounted_dir_path>`. For example:

file://localhost//dfs/recruitment/resume/

On Windows the starting URL format is:

- To crawl local files use `file://localhost/<Directory_Path>`. For example:

file://localhost/D:/recruitment/resume/

- To crawl from a mapped drive use `file://localhost/<machinename>/<shared_folder_path>/`. For example:

file://localhost//RTDC78067TLSBLD/recruitment/resume/

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.

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.

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.

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.

Enable Language Detection

By Enabling Language Detection, the search engine automatically identifies the language of the document content and assigns the language code automatically.

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.

Note: In the Elasticsearch implementation, automatic language detection is supported for all languages that are supported by PeopleSoft.

File URL Prefix

The part of the access URL the system will not display in the search results due to security reasons.

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.

Display URL Prefix

The URL the system displays instead of the actual URL. For example, if the file URL is:

file://localhost/home/operation/doc/file.doc

and you want the display URL to appear as:

https://webhost/client/doc/file.doc

then specify the File URL Prefix as:

file://localhost/home/operation

and the Display URL Prefix as:

https://webhost/client

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.

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.

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.

Image: File Source: Document Types page

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 the 'Document Types' configuration page. At the top, there are tabs for 'General' and 'Document Types'. The page title is 'Document Types'. Below the title, the search definition is 'QE_FILE_PIASRCH' and the description is 'File type for piasrch testing'.

The 'Document types to crawl' section has four radio buttons: 'Include default types', 'Include all types', 'Add document types', and 'Exclude document types'. The 'Exclude document types' option is selected.

Below this is a 'Select' table with a search bar and navigation controls. The table contains two rows of document types:

Document Type		
1 text/plain	<input type="text"/>	<input type="button" value="+"/> <input type="button" value="-"/>
2 application/msword	<input type="text"/>	<input type="button" value="+"/> <input type="button" value="-"/>

At the bottom, there are two sections for 'URL Boundary Rules': 'Inclusion Rules' and 'Exclusion Rules', both of which are currently empty.

Document Types to Crawl

Include default types

Crawl only on the default document types, as defined by the search engine.

The search engine default document types for crawling are:

- PDF
- HTML
- TXT (plain text)
- Microsoft Word
- Microsoft Excel
- Microsoft PowerPoint

If no other document types are added to the Select grid, then the search engine considers only the default document types.

Include all types

Crawl all document types supported by the search engine.

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.

Add these types

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.

Exclude these types

Excludes specific document types added to the Document Type grid.

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.

URL Boundary Rules**Inclusion Rules**

Specify an inclusion rule that a URL must contain. For example:

es_xml

In this case, search engine crawls all documents with *es_xml* in the name.

Specify an inclusion rule that URL must start with. For example:

*file://localhost/ds1/product/ES_ADD/es_doc**

In this case, the search engine crawls all files starting with *file://localhost/ds1/product/ES_ADD/es_doc*.

Exclusion Rules

Specify an exclusion rule that a URL can't contain. For example:

**.xml*

In this case, the search engine does not crawl anything with a *.xml* extension.

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_lvl[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_lvl</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.

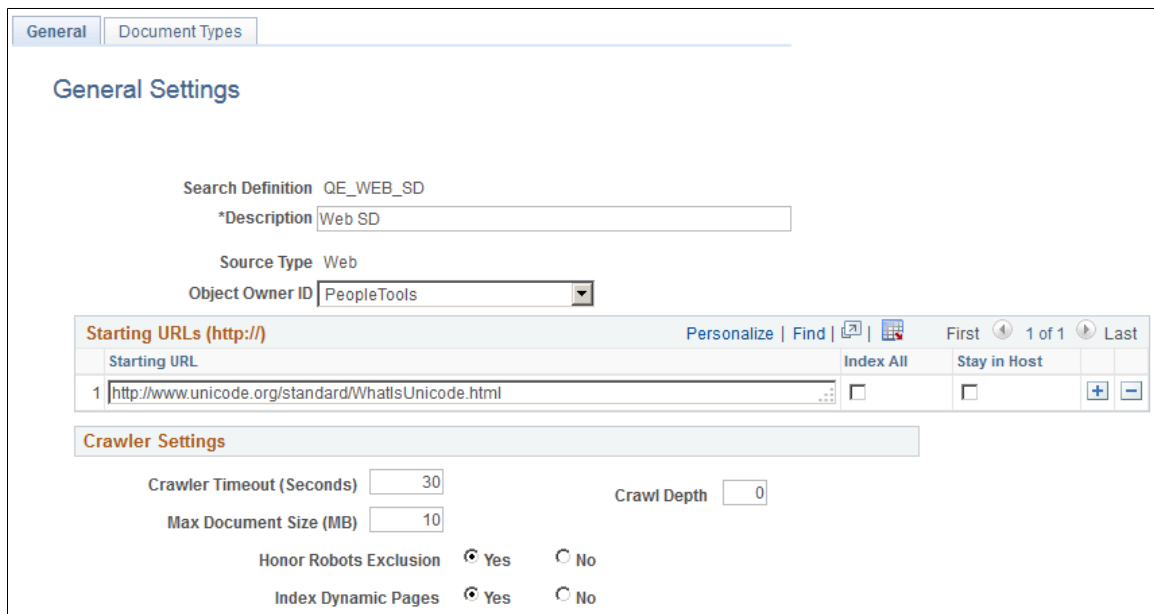
Specifying Web Source General Settings

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.

Image: Web Source — General Settings page

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.



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.
Starting URLs	Contains the URL of the web address. The search engine uses the URL as an entry point for starting to crawl a website. <hr/> 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. <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	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. <hr/> 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. <hr/>
Crawler Timeout	Indicates the maximum allowed time to retrieve a file for crawling.
Crawl Depth	The number of nested links the crawler follows, with the initial URL, or home page, residing at a <i>depth of 0</i> . With a crawling depth of 1, the crawler also fetches any document linked to from the starting URL. 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. 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.
Max Document Size	The maximum document size in megabytes that the system will crawl. Larger documents are not crawled.

Honor Robots Exclusion

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.

- *Yes.* 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.
- *No.* The crawler ignores any specified robot policy defined on the web server.

Index Dynamic Pages

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.

- *Yes.* The search engine crawls dynamic pages.
- *No.* The search engine does not crawl dynamic pages.

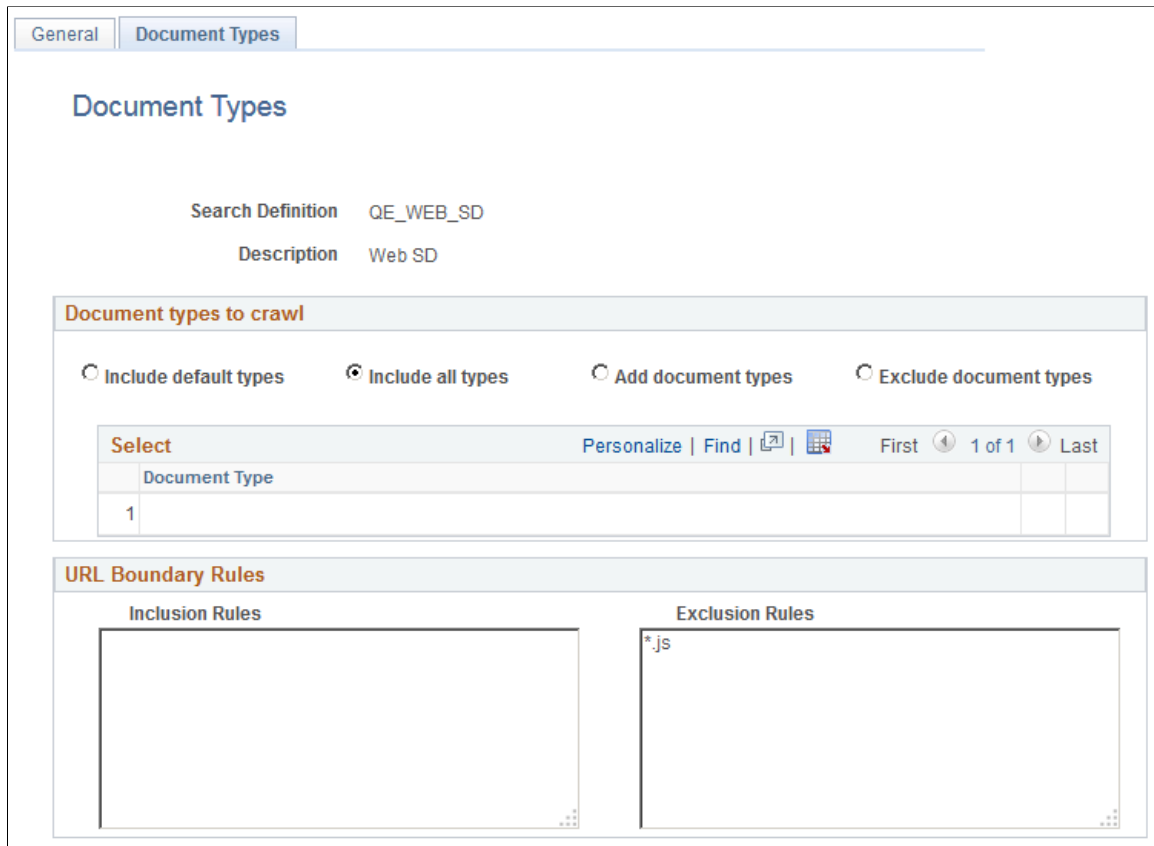
Specifying Document Types for Web Source

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.

Image: Web Source — Document Types page

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.



Document Types to Crawl

Include default types

Crawl only on the default document types, as defined by the search engine.

The search engine default document types for crawling are:

- PDF
- HTML
- TXT (plain text)
- Microsoft Word
- Microsoft Excel
- Microsoft PowerPoint

Note: When you select this option, notice that you cannot add or remove items from the Document Type grid.

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.

Include all types

Crawl all document types supported by the search engine.

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.

Note: When you select this option, notice that you cannot add or remove items from the Document Type grid.

Add document types

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.

Exclude document types

Excludes specific document types added to the Document Type grid.

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.

URL Boundary Rules

Inclusion Rules

Specify an inclusion rule that a URL must contain. For example:

www..example.com*

In this case, the search engine crawls all content within www.*.example.com.

Exclusion Rules

Specify an exclusion rule that a URL can't contain. For example:

www..uk.example.com*

In this case, the search engine does not crawl anything within www.*.uk.example.com.

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

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.

Chapter 8

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.

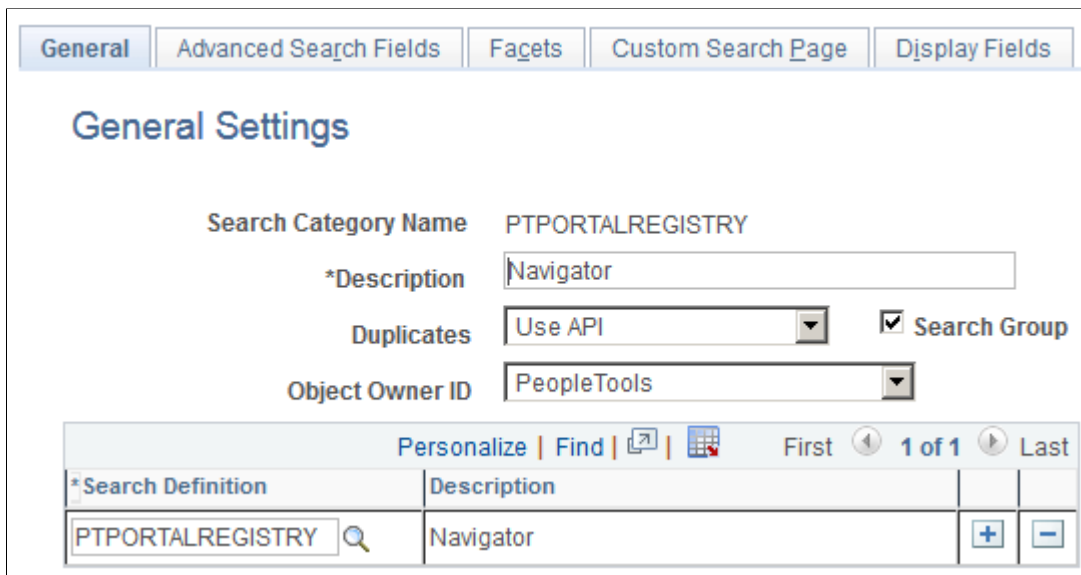
Note: Search queries cannot be run directly against search definitions. Search queries can only be run against search categories.

Specifying General Search Category Settings

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

Image: General page

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



Search Category Name

Displays the search category name.

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 *same* 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 *must* belong to one with the same name.

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 *not* 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.

Description

Add any additional information to distinguish search categories.

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.

Note: The description is translatable.

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	Click to clone the search category. In the subsequent page, you need to enter a name for the cloned search category. When you clone a search category, all the meta-data associated with the existing search category is saved in the cloned search category.
	Note: The Save As button is not available when you create or define a new search category.

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.

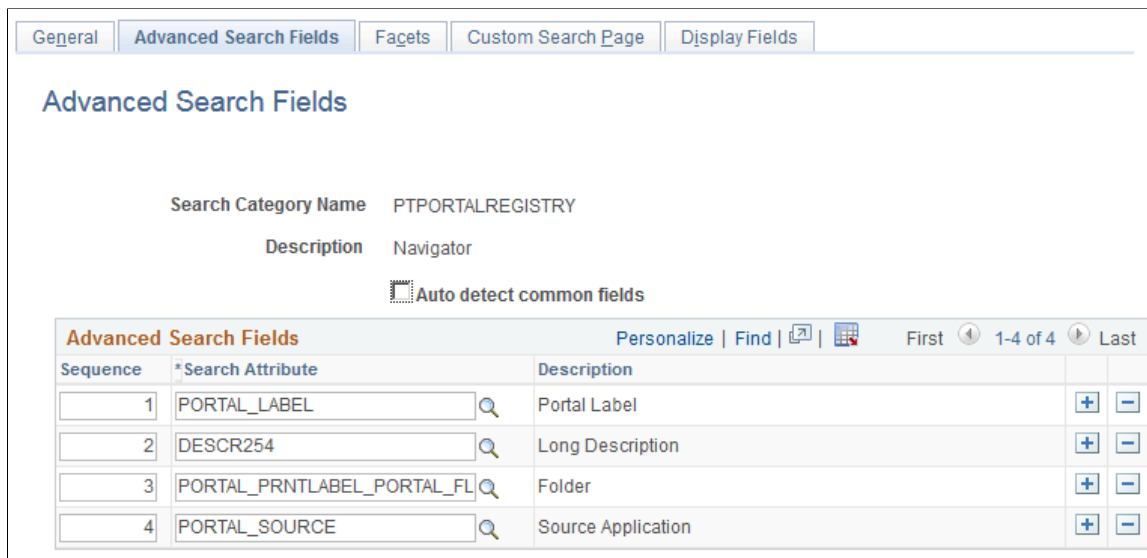
The Advanced Search Fields page enables you to view and modify attributes to show on the advanced search page during end user searches to add more criteria to the basic search mode.

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.

Image: Advanced Search Fields page

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.



Auto detect common fields

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.

Select to avoid manually listing the common fields from the joined records. (default)

Deselect to delete some or to add additional attributes which are not common to all search definitions.

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.

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.

Image: Facets page

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

Auto detect facets

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.

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).

Sequence

Control the sequence in which the faceted fields display for narrowing search results in an intuitive fashion for users.

For example, perhaps you might want to ensure that these faceted fields display in this order: Country, State, County, City.

Description

Displays the value of the Attribute Display Name column in the grid on the Map Search Attributes page for the search definition.

Is Associated

Select if you intend to create a user-specific association with the facet.

Note: Creating associated facets applies only to certain situations and expects additional requirements to be in place.

For more information, refer to the topic below, “Creating Associated Facets.”

Facet Association

Select how the value to which the facet should be associated will be determined. Options are:

- *Application Class:* 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.
- *Employee ID:* Associates the facet value with the employee ID value (%EmployeeID).

- *User ID*: Associates the facet value with the currently signed on user (%UserID).

Advanced Settings

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.

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.

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.

Important! Setting facet ranges applies only to date and number data types. Values with data types of datetime cannot be used as facets.

See [Defining Advanced Settings for String Facets](#).

See [Defining Advanced Settings for Number Facets](#).

See [Defining Advanced Settings for Date Facets](#).

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

Image: Advanced Facet Settings page (for string data types)

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.

The screenshot shows the 'Advanced Facet Settings' dialog box. At the top, it displays 'Facet Name' as 'PORTAL_PRNTLABEL_PORTAL_FLDR' and 'Data Type' as 'String'. A checkbox labeled 'Allow Multi Select' is checked. Below this, there is an 'Override default settings' section with 'Minimum Document Count' set to 0 and 'Maximum Facet Children' set to 0. A 'Disable Facet Count' checkbox is also present. A 'Sort Order' table is shown with columns 'Level', 'Sort as', and 'Sort order'. The table contains one row with 'Level' 1. Navigation buttons 'OK' and 'Cancel' are at the bottom left. A 'Help' link is at the top right.

Allow Multi Select

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 except for hierarchical facets. The default behavior for hierarchical facets is single select.

Select the Allow Multi Select check box to enable multiple facet-values selection in hierarchical facets.

Note: Selecting multiple facet values is supported only on fluid search pages; not on classic or classic plus search pages.

See [Working with Search Results in Global Search](#).

Override default settings	Select to override the values entered on the Search Options page (PeopleTools, Search Framework, Search Admin Activity Guide, 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.
Minimum Document Count	Set the minimum number of search results (documents) required to display a facet. Default is <i>1</i> .
Maximum Facet Children	Set the maximum number of children to be returned for each facet node. Default is <i>100</i> . Options to enter are: <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	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. Enter <i>Y</i> for Yes or <i>N</i> for No.
Level	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).
Sort as	Select one of these options to specify how the facet should be treated: <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	Specify the sort order that best suits your selected Sort as option and how users expect to see the values displayed, ascending or descending.

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.

Image: Advanced Facet Settings page (for number data types)

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

Advanced Facet Settings

Facet Name: QE_JOB_OPENING_ID Data Type: Number Allow Multi Select

Override default settings

Minimum Document Count: 0 Disable Facet Count

Maximum Facet Children: 0

Sort Order: Personalize | Find | View All | First 1 of 1 Last

Level	Sort as	Sort order
1	Predefined	

Number Range

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

OK Cancel

Allow Multi Select

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 except for hierarchical facets. The default behavior for hierarchical facets is single select.

Select the Allow Multi Select check box to enable multiple facet-values selection in hierarchical facets.

Note: Selecting multiple facet values is supported only on fluid search pages; not on classic or classic plus search pages.

See [Working with Search Results in Global Search](#).

Override default settings

Select to override the values entered on the Search Options page (PeopleTools, Search Framework, Search Admin Activity Guide, 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.

Minimum Document Count

Set the minimum number of search results (documents) required to display a facet. Default is 1.

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).

Maximum Facet Children

Set the maximum number of children to be returned for each facet node. Default is *100*.

Options to enter are:

- An integer greater than 0 to limit the number of children displayed.
- *0* to indicate no children should be displayed.
- *-1* to return all children (no limit).

Disable Facet Count

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.

Enter *Y* for Yes or *N* for No.

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**Sequence**

Order in which the ranges are defined.

Node Name

Display name of the facet range. This is a display only field. 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

Begin Operator

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:

- 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

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:

- 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.

Image: Advanced Facet Settings page (for date data types)

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' dialog box. At the top, it displays 'Facet Name: OPEN_DT', 'Data Type: Date', and a checked 'Allow Multi Select' checkbox. Below this, there are options to 'Override default settings' (unchecked) and 'Disable Facet Count' (checked). Input fields for 'Minimum Document Count' (0) and 'Maximum Facet Children' (0) are present. A 'Sort Order' table is shown with one level. The 'System' section contains various date range checkboxes like 'Today', 'This Year', 'Past 7 Days', etc. The 'Date Range' section contains a table for defining date ranges with columns for Sequence, Node Name, Range Type, Begin Operator, Begin Date, Begin No, Begin Range, End Operator, End Date, End No, End Range, Message Set Number, and Message Number. At the bottom are 'OK' and 'Cancel' buttons.

Allow Multi Select

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 except for hierarchical facets. The default behavior for hierarchical facets is single select.

Select the Allow Multi Select check box to enable multiple facet-values selection in hierarchical facets.

Note: Selecting multiple facet values is supported only on fluid search pages; not on classic or classic plus search pages.

See [Working with Search Results in Global Search](#).

Override default settings

Select to override the values entered on the Search Options page (PeopleTools, Search Framework, Search Admin Activity Guide, 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.

Minimum Document Count

Set the minimum number of search results (documents) required to display a facet. Default is *1*.

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).

Maximum Facet Children

Set the maximum number of children to be returned for each facet node. Default is *100*.

Options to enter are:

- An integer greater than 0 to limit the number of children displayed.
- *0* to indicate no children should be displayed.
- *-1* to return all children (no limit).

Disable Facet Count

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.

Enter *Y* for Yes or *N* for No.

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.

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.
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.

Sequence

Defines the order in which the facets appear.

Node Name

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.

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.

If the range is relative, then it is required to provide a message catalog entry for the facet node name.

Range Type

Enables you to create ranges on fixed dates or dates relative to today. Options are:

- *Fixed*. Select to define static beginning and end dates.
- *Relative*. Select to enable the Begin Range drop-down list to select dates relative to the current date.

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.

Begin Operator

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:

- 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

Set the span of the dates to be included in the range (past or future). Select in the form of:

- days/months/years before.
- days/months/years from now.

End Operator	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: <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.
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.

Image: Custom Search Page

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.



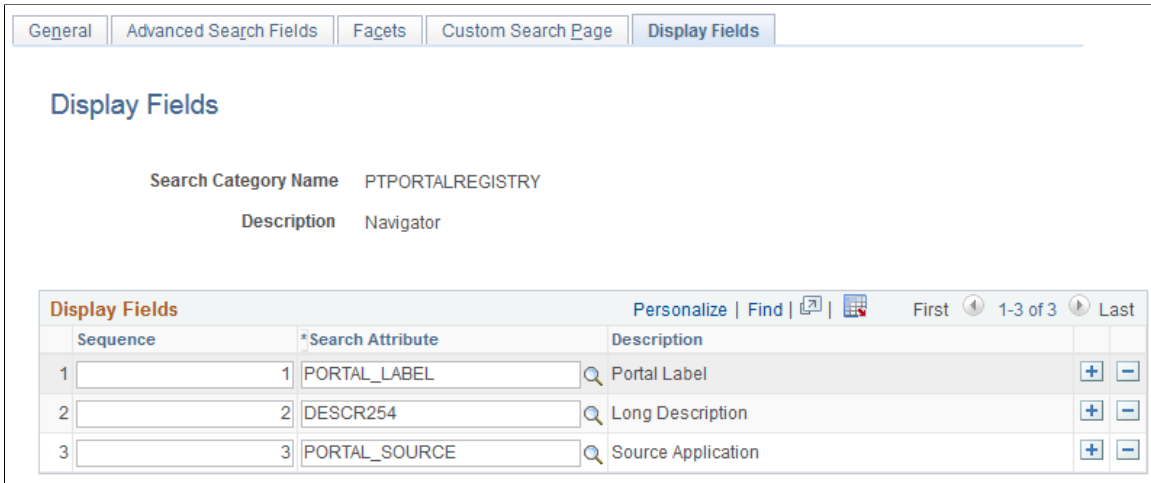
Menu Name	Select the menu to which the search page belongs.
Component Name	Select the search page's component name.
Market	Select the appropriate market.
Search results page	Select the search results page to use.
Advanced search page	Select the advanced search page to use.

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.

Image: Display Fields page

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.



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.

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" (PeopleTools 8.57: Integration Broker Administration) and "Understanding Managing Integration Gateways" (PeopleTools 8.57: 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" (PeopleTools 8.57: 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" (PeopleTools 8.57: 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](#).

Image: Search Instance Properties page

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

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

Search Instance

Enter a name for the search instance.

Note: PeopleSoft Search Framework enables you to create multiple search instances for the purpose of fail-over mechanism.

See [Understanding High Availability Feature in PeopleSoft Search Framework](#) and [Implementing Fail-over Mechanism](#) for more information.

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](#).

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](#).

Host Name

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.

For example:

example.com

Port

Enter the port on which the search engine listens for requests.

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.

SSL Option

Select one of the options for SSL:

- Select DISABLE in the following scenarios:
 - 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:

- 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.

See [Configuring SSL between PeopleSoft and Elasticsearch](#).

Ping

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.

Ping Test Result: Success. Elasticsearch Version 2.3.2 (262,613)

Admin User

Enter the user name to connect to the search engine.

Admin Password

Enter the password associated with the user name.

Test Login

Click to test the credentials you have specified against the search engine. A positive result displays a success message.

Login Success. (262, 615)

Proxy User

Enter a trusted entity.

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.

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.

Note: In Elasticsearch, Proxy User acts as an authentication user for every search request.

Proxy Password

Enter the password associated with the trusted entity.

Proxy Login

Click to test the credentials you have specified against the search engine. A positive result displays the success message.

Login Success (262,615)

+ (plus sign)

Click the + (plus) sign in the grid to enter details of another node if you plan to implement fail-over mechanism.

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.

See [Understanding Elasticsearch Cluster](#).

Elasticsearch Interact

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.

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

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.

URL

Enter the URL for the PeopleSoft system listening connector, using the following syntax.

```
http://<server>:<port>/PSIGW/RESTListeningConnector/
<node>
```

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 `elasticsearch.yml` configuration file.

If Elasticsearch is behind a load balancer and the setting is to terminate SSL at load balancer, the `orclssl.callback` property should be set to false in the `elasticsearch.yml` configuration file.

User Name

Enter the PeopleSoft user name granted the permission list PTPT3300. This user must have the Search Server and Search Query Administrator roles.

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.

Password\Confirm Password

Enter and confirm the password associated with the PeopleSoft user name.

Validate

Click to validate the URL, service operations, and credentials you have entered. A positive result displays the following message on the Validation Report page:

All validations are successful.

The following service-operation validations are performed to confirm:

- 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

Click this link to invoke a bulk update for all definitions currently deployed to the search engine.

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.

Set Namespace Aliases

Click this link to set namespace aliases when you clone an environment.

See [Setting Namespace Aliases](#) for more information.

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 signon 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.

Image: Elasticsearch Interact page (initial)

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.

The screenshot shows a web browser window titled "Elasticsearch Interact". The page contains the following elements:

- Search Instance:** A text field containing "PTSF_DEFAULT".
- Service Type:** A dropdown menu.
- Suggestions:** A dropdown menu.
- URL:** A text field containing "http://example.com:9200/".
- Buttons:** A "Submit" button in the top right and a "Return" button in the bottom left.
- Help:** A "Help" link in the top right corner.

Image: Elasticsearch Interact page (after selecting an option from the Service Type drop-down list)

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.

The screenshot shows the 'Elasticsearch Interact' page. At the top right is a 'Help' link. Below it are three input fields: 'Search Instance' with the value 'PTSF_DEFAULT', 'Service Type' with a dropdown menu showing 'Cluster', and 'Cluster API' with a dropdown menu showing '._cat/indices?v'. To the right of these fields is a 'Submit' button. Below the input fields is a 'URL' field containing 'http://example.com:9200/_cat/indices?v'. The main area of the page is a table of search results. The table has columns for 'Result', 'health', 'status', 'index', 'uuid', 'pri', 'rep', 'docs.count', 'docs.deleted', 'store.size', 'pri.store.size', and 'size'. The results are listed in a scrollable area.

Result	health	status	index	uuid	pri	rep	docs.count	docs.deleted	store.size	pri.store.size	size
green	open	qe_piasrch_portalname_t57r907w	4.7mb	AQ1zCJqKQGw9fpPg1F_qFA	5	0	10447	0	4.7mb		
green	open	hc_hpy_revw_garn_h92mf907	1.6mb	qLc5daUGRwmSWAN_OSHe9A	5	0	1318	33	1.6mb		
green	open	ptsearchreports_p91909rr	1.1kb	gu23rRZoTvmxw7iuYNaEjw	5	0	0	0	1.1kb	1.1kb	
green	open	qe_cl_itm_m_t57r907w	865.4kb	rx49Ov_QES_9b3ygC3pzQ	5	0	796	0	865.4kb	865.4kb	
green	open	hc_hrs_job_posting_h92911tg	2mb	WMA5Q-LuSky7Ah1-7NhgJA	5	0	543	0	2mb	2mb	
green	open	qe_nui_itemdefin_t57568p5	949.3kb	AIDYskvRRieUdHGwLcOP1Q	5	0	796	0	949.3kb		
green	open	hc_hrs_job_content_h92cpsrc	603.7kb	gt1H3L69S7Wa87JFpa9mdg	5	0	794	63	603.7kb		
green	open	qe_piasrch_sort_t57jd910	3.4mb	N0sppjDmSJWx9w7HHgwmSw	5	0	7852	0	3.4mb		
green	open	qe_piasrch_portalname_t57904ww	4.7mb	xzLANzZDRLa_hvFA-3hX3w	5	0	10435	0	4.7mb		
green	open	papp_resource_finder_ext_component_p91909rr	1.1kb	qtgyvXyURuGKLadZCM0a4Q	5	0	0	0	1.1kb		
green	open	hc_jpm_person_profile_h92mf907	15mb	Qv5vjiABTa-DilyEwd7KyQ	5	0	1031	0	15mb	15mb	
green	open	ptwl_gen_msg_wl_h92cpsrc	1.2kb	2NTwJw_FQUOM4afJR_cyw	5	0	0	0	1.2kb	1.2kb	
green	open	papp_content_p91909rr	1.1kb	nk1jDh4SVKUF94WoyupiQ	5	0	0	0	1.1kb	1.1kb	
green	open	qe_nui_img_cq_t57846c6	459.2kb	O_S5T0NIS_qXdyLcrUuX2w	5	0	294	0	459.2kb	459.2kb	
green	open	hc_hr_company_directory1_h92909rh	1.1kb	Ki2Cep3FRXIEG3tb43bpaw	5	0	0	0	1.1kb	1.1kb	

Search Instance

The search instance name.

Service Type

Displays a list of service type groups.

The following service type groups are available:

- 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.

Suggestions

The Suggestions field displays on the initial Elasticsearch Interact page and it changes based on the service type you select.

The Suggestions field changes to the following:

- When you select Cluster as service type, the Suggestions field changes to Cluster API, which provides cluster related options:

- 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 `/_cluster/health` 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.

For a description of the color codes used for cluster health in Elasticsearch, see [Viewing the Elasticsearch Cluster Dashboard](#).

- 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).
- Nodes statistics - provides statistics of all nodes in a cluster. By default, all statistics of a node are displayed.

For a description of the node statistics, see [Viewing the Node Dashboard](#).

- 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

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.

Valid options are:

- GET
- POST

URL

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:

```
http://example.com:9200/_cat/indices?v
```

Note: The URL cannot be modified. To modify the URL, you need to change the options you have selected.

Body

The Body field appears when you select POST as a method. This field must contain a JSON request.

The contents of the Body field are passed along with the request to Elasticsearch.

Result

The Result area displays the response from Elasticsearch for each request. PeopleSoft Search Framework displays the response from Elasticsearch as is.

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.

Image: Namespace Alias Settings page

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.

Namespace Alias Settings

Map the cloned and original database names and their corresponding node names

Database Name Aliases

Database Name	Database Alias Name
1 H9201CHC	H9201CHA

Node Name Aliases

Node Name	Node Alias Name	Alias Node URI
1 H9201CHC	H9201CHA	http://myserver.company.com:8000/psc/h9201ch

Save Cancel Clear Data

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.

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:

http or https://<hostname>:<port>/psc/<sitename>/

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.

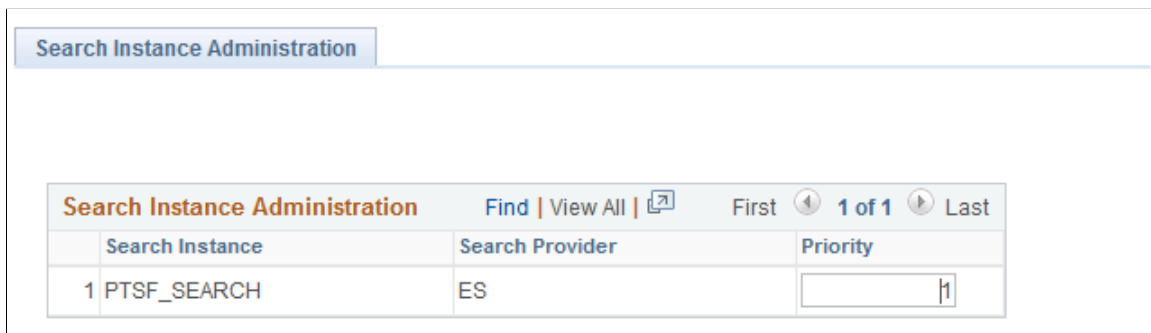
Administering Search Instances

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

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

Image: Search Instance Administration page

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.



Search Instance

Lists the search instances created in PeopleSoft Search Framework.

Search Provider

Lists the search provider associated with a search instance.

Priority

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.

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.

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.

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 Instance Administration page by selecting PeopleTools >Search Framework >Administration >Search Instance/User.

Image: Search User Instance page

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.

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

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 doesn't 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.

For descriptions of each of the properties, see "PTSFConfigureSrchInstance" (PeopleTools 8.57: 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" (PeopleTools 8.57: 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

To enable Global Search, access the General Settings page by selecting PeopleTools, Search Framework, Search Admin Activity Guide, Configuration, Enable Global Search.

On the General Settings page, in the Portal Search Options group box, select Display global search in header.

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.

In a fluid user interface, 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.

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

Managing General Search Options

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

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

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

Image: Search Options 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.

The screenshot shows the 'Search Options' page. At the top, there is a 'Search Instance' dropdown menu with 'PTSF_DEFAULT' selected. Below this is a section titled 'Default Values' containing a list of parameters and their current values in input fields:

Parameter	Value
Number of Shards	5
Number of Replicas	0
Minimum Document Count	1
Maximum Facet Children	100
Disable Facet Count	N
Min Match Percent	100
Cache Interval (Min)	120
Index Segment Size (MB)	10
Log ES requests for User ID	
Log ES metrics for User ID	
Enable Attachment Trace	N
Attachment handlers	20
Max Sub Queue Size	20
Max Attachment Error Count	100
Full DirectTransfer	Y
Maximum Tree Facet Nodes	1000

Search Instance

Enables you to select the search instance for which you want to specify the global settings.

Number of Shards

Elasticsearch enables you to distribute the indexed data into partitions. Default value is 5.

The value entered here is applicable to all search definition deployed from this PeopleSoft machine to Elasticsearch server unless it is overridden.

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).

Note: The value of this parameter cannot be changed after an index is created.

See [Elasticsearch Concepts and Terminology](#).

Number of Replicas

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.

Set the number of copies you want to maintain of the primary shards.

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 N , that is, the number of nodes in a cluster.)

The value entered here is applicable to a search instance and to the security index (orcl_es_acl).

See [Setting the Number of Replicas](#).

Note: After the index is created, you may change the number of replicas dynamically anytime.

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).

See [Elasticsearch Concepts and Terminology](#).

Minimum Document Count

Set the minimum number of search results (documents) required to display a facet. Default is 1 .

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).

Maximum Facet Children

Set the maximum number of children to be returned for each facet node. Default is 100 .

Options to enter are:

- An integer greater than 0 to limit the number of children displayed.
- 0 to indicate no children should be displayed.
- -1 to return all children (no limit).

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).

Disable Facet Count

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.

Enter *Y* for Yes or *N* for No.

- *N*: Returns the document count. Default is *N*.
- *Y*: Shows no document count.

Min Match Percent

Enables controlling the display of search results based on the percentage of search keywords found in a document. The default value is 100.

A value of 100 will display a document in search results only when all the search keywords are found in the document.

A value of 50 will display a document in search results if at least 50% of the search keywords are found in the document.

Cache Interval (Min)

Set the duration of time (in minutes) for storing the security cache of a user. The default value is 120 minutes.

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.

The *acl.cache.delete.interval* property in the *elasticsearch.yml* 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.

Index Segment Size (MB)

Set the size, in megabytes, by which the system segments (or chunks) the indexing data feed. Default is 10 megabytes.

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.

Log ES requests for User ID

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.

Note: Only one user ID can be entered at a time.

Log ES metrics for User ID

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.

Note: Only one user ID can be entered at a time.

Enable Attachment Trace

Enables controlling the detailed logging of each task on the trace file. You should consider setting this property to *Y* 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.

Enter *Y* (yes) if you want to enable detailed logging of each task on the trace file.

Enter *N* (no) to disable logging of each task on the trace file.

Attachment handlers

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.

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.

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).

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 more number of parallel threads can cause congestion in Elasticsearch queue, which directly reduces the heap availability.

See [Configuration](#).

Max Sub Queue Size

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.

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.

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.

Max Attachment Error Count

This property is associated with the direct transfer of search data from PeopleSoft to the search engine, which bypasses the Integration Gateway

The maximum number of exceptions that should be displayed on the Attachment Transfer Exception Details page.

The default value is 100.

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.

See [Viewing Search Index Build Process Details](#).

See [Configuration](#).

Full Direct Transfer

The default value is *Y*.

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

If you set the Full Direct Transfer setting to *N* (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.

See [Configuration](#).

Maximum Tree Facet Nodes

Set the maximum number of tree facet nodes to be returned for each facet. Default is *1000*.

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)

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.

Image: Deploy Search Definition page

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.

Deploy Search Definition							Personalize Find [Print] [Refresh]
Definition	Description	Deploy Status	Crawl Status	Run Date/Time	User/Run Control ID	Index Settings	
1	PTPORTALREGISTRY	Navigator	Deployed			Index Settings	
2	PTSEARCHREPORTS	Reports	Deployed	09/03/2016 9:47:09AM	QEDMIO / PTSEARCHREPORTS	Index Settings	
3	PTWL_GEN_MSG_WL	Worklist Note	Deployed			Index Settings	

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.

Deploy Search Definition (grid)

Displays all existing search definitions defined in the Search Framework. Columns in the grid are:

- **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 (*Success, Fail*).
- **Audit Report.** If the Audit Status is *Fail*, the Audit Report column appears providing a link to the audit report for review.

Index Settings

Click to display the Advanced Index Settings dialog box where you can override some of the settings specified on the Search Options page.

If you click the Index Settings link for a deployed search definition, you can change only the number of replicas.

If you click the Index Settings link for a search definition that is not deployed, you can update the following settings:

- **Override Settings** - Select this option if you want to override the settings specified on the Search Options page.
- **Index Name** - Enter a different index name than the one specified on the Search Options page, if desired.

Note: If the data volume to be indexed is large, it is recommended to map a search definition to a separate index other than the default index. By mapping a search definition to a separate index, you can improve the overall performance for the search definition.

- Number of Shards - Specify a new number of shards, if desired.
- Number of Replicas - Specify a new number of replicas, if desired.

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](#).

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.

Select All

Click to select all the rows within the grid.

Note: If all rows in the grid have been selected, the button title changes to Deselect All.

Report Sync Issues

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.

Audit

Click to run a series of audits on the selected search definitions to ensure they have all requirements in place.

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 *Success* or *Fail*. If the result is *Fail*, 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.

See the section below: “Running the Search Definition Audit Process.”

Deploy

Deploy the PeopleSoft search definitions selected to the search engine, where it becomes a searchable data source.

Undeploy

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.

Note: When you deploy a search definition, the search category of the same name is automatically deployed.

Click to remove the representation of the selected search definitions from the search engine.

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.

Note: When you undeploy a search definition, the search category of the same name is automatically undeployed.

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.

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. 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.

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

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.

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.

Important! Oracle PeopleSoft recommends that you do *not* delete delivered search definitions.

Warning! Delete is not reversible! Do not use this function unless you are certain it is what you want.

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.
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:

Image: Search Definition Compare Report page

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.

Compare report			Personalize	Find	First	1-4 of 4	Last
Property Type	Peoplesoft	Search Server					
1 Search Attribute	DESCR254	Absent					
2 Search Attribute	PORTAL_LABEL	Absent					
3 Search Attribute	PORTAL_PRNTLABEL_PORTAL_FLDR	Absent					
4 Search Attribute	PORTAL_SOURCE	Absent					

Return

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.

Image: Search Definition audit report columns


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.

Image: Search Definition Audit Report page

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.

Image: Deploy Search Category page

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
<input type="checkbox"/>	PTPORTALREGISTRY	Navigator	Deploy Related Definition	Auto Deployed
<input type="checkbox"/>	PTSEARCHREPORTS	Reports	Deploy Related Definition	Auto Deployed
<input type="checkbox"/>	PTWL_GEN_MSG_WL	Worklist Note	Deploy Related Definition	Auto Deployed
<input type="checkbox"/>	QE_ALL_CAT_CQ_SD1	QE All cat CQ SD1	Deploy Related Definition	Undeployed

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.

Deploy Search Category (grid)

Displays all existing search categories defined in the Search Framework. Columns in the grid are:

- **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

Click to select all the rows within the grid.

Note: If all rows in the grid have been selected, the button label changes to Deselect All.

Report Sync Issues

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.

Deploy

Deploy the PeopleSoft search category to the search engine.

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.

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.

Undeploy

Click to remove the search category from the search engine.

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.

Note: The search category will be removed from the search engine, but it will still exist within the PeopleSoft system.

Note: This does not undeploy the search definitions associated with a search category.

Update

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.

Delete

Deletes the search category from the Search Framework.

Note: This does not delete the search definitions associated with a search category.

Warning! This is not reversible! Oracle PeopleSoft does *not* recommend deleting delivered search categories.

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.
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 the Delete Query.

Note: Each partition 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

[Using Dashboards and Visualizations in Kibana](#)

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.)

Image: Build Search Index page

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.

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.

Search Instance

Select the search instance in which you want to build the search index.

Search Definition

Select the search definition for which you are building the search index.

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.

Indexing Type

Select one of the following:

- *Full index.* 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.
- *Incremental index.* 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.

Items are also removed from the index as designated by the delete query on the search definition.

Note: It is recommended to 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.

Language Options

Select one of the following:

- *All Languages.* Builds an index for each language enabled on the database.
- *Base Language.* Builds an index only for the base language defined for the database.

Your selection depends on the languages enabled for your database and the languages through which you anticipate end users using to perform searches.

Full indexing Criteria

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.

For more information on this option see, [Working With Advanced Settings](#).

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](#).

Date Range

The Date Range option appears when the selected search definition is based on a connected query. The Date Range option enables you to partition the search data based on a date range. You need to specify:

- *Start Date.* Specify the start date of the date range.
- *End Date.* Specify the end date of the date range.
- *Partition.* Enter the number of partitions you want the search index build process to create.

Note: Save the settings before you select the Run button in the Date Range option.

Note: PeopleSoft does not recommend using Date Range partitioning with search definitions that use preprocessing unless the preprocessing logic has been modified for use with partitioning. If you select search definitions that use preprocessing logic, then the indexing process may run preprocessing for each partition, which may yield undesirable results and may take considerable time to complete the indexing process.

For more details on using the Date Range option, see [Partitioning Application Data in the Search Index Build Process](#).

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.

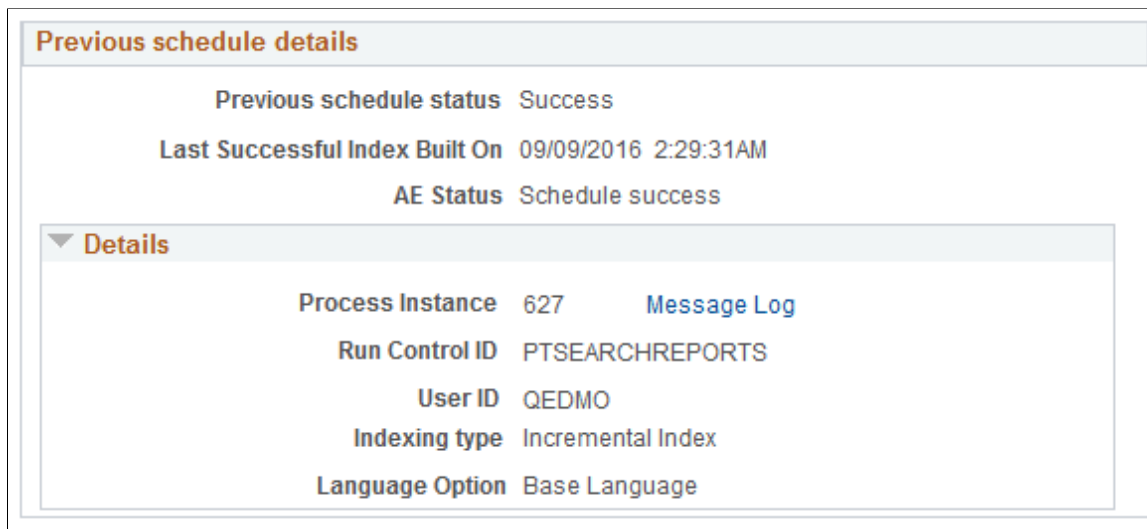
Viewing Search Index Build Process Details

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.

Image: Previous schedule details

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 status

Displays the status of the most recently executed index build process.

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.

See [Viewing the Exception Details in Attachment Processing](#).

Last Successful Index Built On

Displays the date when the index was built successfully. If an index was never built, it displays *Never*.

AE Status

Indicates whether the most recently executed PTSF_GENFEED Application Program ran to completion.

If the index build fails, this field displays the step where the process failed.

Resume schedule

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.

If the build fails, see [Handling Common Errors](#) to determine the cause of the failure to build an index.

Process Instance

Indicates whether the system successfully created the output that is required to populate the index.

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.

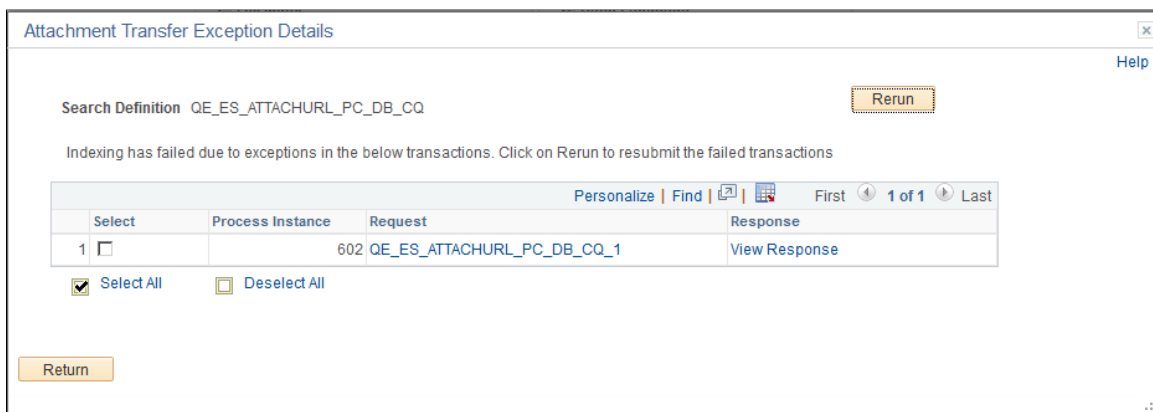
Indexing Type	Displays the index type (Full Index or Incremental Index).
Language Option	Displays the language options selected for the index build process (All Languages or Base Language).

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.

Image: 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.



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 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.

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.

Related Links

[Working with Attachment Properties](#)

Partitioning Application Data in the Search Index Build Process

Indexing of search definitions that are based on connected query and 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 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 connected 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 date span for the number of partitions specified by you.

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.

Note: When you specify a date range to implement partitioning of search data, the date span entered in the Full Indexing Criteria option is overridden.

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

Image: Date Range grid

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"/>

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 DirectTransfer

DirectTransfer 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.

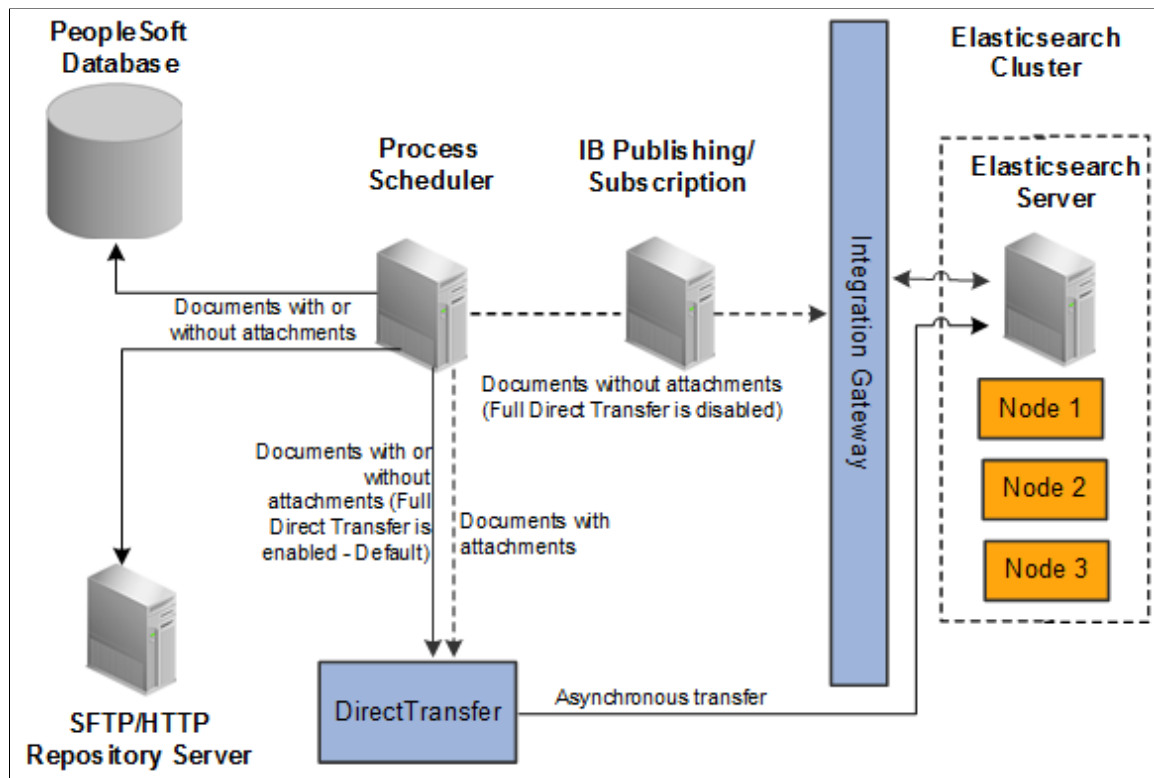
Understanding DirectTransfer

PeopleSoft Search Framework implements DirectTransfer 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.

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

Image: Full Direct Transfer architecture

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



To use DirectTransfer, you need to:

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

Configuration

DirectTransfer requires the following configurations:

- Specify the number of attachment 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 Attachment Handlers

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

The number of Attachment 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 Attachment 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 Attachment Handlers can be reduced to 10.

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

Note: The number of Attachment 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 DirectTransfer:

- Memory usage on PeopleSoft
- Memory usage on Elasticsearch

Memory Usage on PeopleSoft

The amount of memory used by DirectTransfer 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, DirectTransfer would utilize an additional 200 MB on the PeopleSoft Batch Server during data send.

Memory Usage on Elasticsearch

When DirectTransfer 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 DirectTransfer, 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 DirectTransfer, 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, DirectTransfer 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.

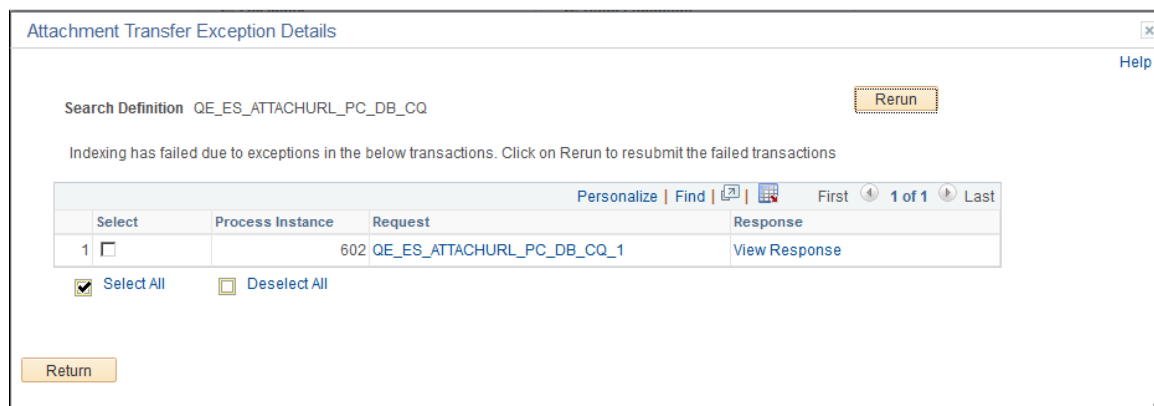
During the indexing of search documents without attachments (which are transmitted through the Integration Gateway), if any error occurs in IB transactions, an Error link may not be displayed on the Build Search Index page. You need to check the Process Monitor whether the PTSF_GENFEED Application Engine 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.

Image: 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.



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 DirectTransfer with SSL Enabled Elasticsearch

For DirectTransfer 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 DirectTransfer 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 >Digital Certificates.

Image: Digital Certificates page

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.

Image: Certificate bundle in PEM format

This example illustrates the certificate bundle in PEM format.

Name	Date modified	Type	Size
.adm	6/7/2017 12:02 AM	File folder	
.ExitProc	6/8/2017 7:57 AM	File folder	
Archive	6/7/2017 12:02 AM	File folder	
CACHE	6/7/2017 12:02 AM	File folder	
diagnostics	6/7/2017 12:01 AM	File folder	
files	6/7/2017 12:01 AM	File folder	
log_output	6/8/2017 8:06 AM	File folder	
LOGS	6/8/2017 8:06 AM	File folder	
security	6/7/2017 12:01 AM	File folder	
access.060717	6/7/2017 3:25 AM	060717 File	1 KB
curl-ca-bundle.pem	6/8/2017 8:06 AM	PEM File	25 KB
psprcs.cfg	6/7/2017 12:02 AM	CFG File	30 KB

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" (PeopleTools 8.57: 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.

Image: Define Search Context page

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.

The screenshot shows the 'Define Search Context' page for the 'Home Page' context. It features a table with the following columns: Default, Sequence, *Search Category Name, and Description. A single row is visible with the following values: Default (checkbox), Sequence (1), *Search Category Name (PTPORTALREGISTRY), and Description (Navigator). There are plus and minus buttons at the end of the row.

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

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

Default

Select the category that should be the default.

Only one default search category can be selected for a search context.

Note: If no categories are marked as Default, then All is the default.

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 *not* be removed.

Sequence

Specify a numerical sequence to define the order in which the search categories will appear in the Global Search Bar drop-down list.

Search Category Name

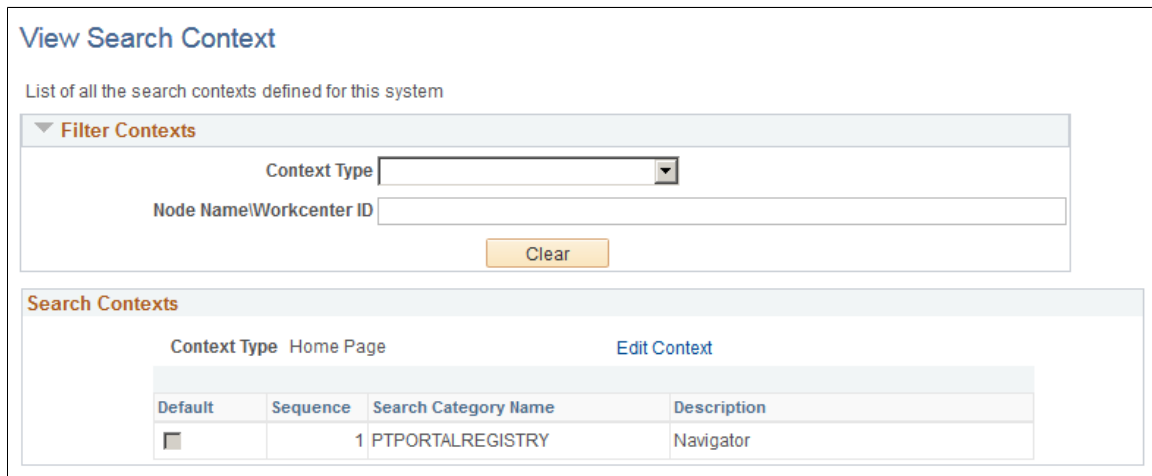
Select the search category to include for this search context.

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.

Image: View Search Context page

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.



Filter Contexts

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

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.

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" (PeopleTools 8.57: 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 signon 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.

Image: Remote Search Groups page

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.

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" (PeopleTools 8.57: Integration Broker Administration)

"Setting Portal Node Characteristics" (PeopleTools 8.57: 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, 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.

Image: Attachment URL IDs page

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.

The screenshot shows the 'Attachment URLIDs' page with the title 'URLID List used by Search Framework Attachments'. Below the title is a table with columns for 'URL Identifier' and 'Description'. The table contains six rows of data, each with a search icon and plus/minus buttons. The table also includes navigation controls like 'First', '1-6 of 19', and 'Last'.

URLID List	Personalize	Find	View All	First	1-6 of 19	Last
*URL Identifier	Description					
1 PSSESATT_1	SES URL from File Attachments	+	-			
2 PSSESATT_2	SES URL from File Attachments	+	-			
3 PSSESATT_3	SES URL from File Attachments	+	-			
4 QE_SES_ATTACH_S	QE_SES_ATTACH_STATIC_FTP	+	-			
5 QE_SES_ATTACH_S	QE_SES_ATTACH_STATIC_PATH_FTP	+	-			
6 QE_SES_DYN1	QE_SESDYN1	+	-			

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" (PeopleTools 8.57: 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, System 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)

Image: Search Reports page

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

The screenshot displays the 'Search Reports' interface. At the top, there is a search input field containing the text 'reports' and a search icon. Below the search field, the results are organized into a table with two main columns: 'Filter by' and 'Search Results'. The 'Filter by' column lists various filters such as 'Created By', 'Process Type', 'Folder Name', and 'Created On', each with a count in parentheses. The 'Search Results' column lists the search results, including a link to 'abcd', 'Detailed Financial Statement', and 'Cross Reference Window Listing'. Each result includes the 'Last Updated Date' and the 'File Name'. Navigation controls like 'View All', 'First', '1-3 of 3', and 'Last' are visible at the top right of the results section.

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" (PeopleTools 8.57: 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" (PeopleTools 8.57: 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", "speeks~");  
  &fg.BoostFilterScore(0.02);  
  &fg.EndMatchAny();  
End-If;
```

Related Links

"SearchFilterGenerator Class Methods" (PeopleTools 8.57: PeopleCode API Reference)

Testing Score Boosting

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

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

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.

Image: 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 shows the 'Search Administration Activity guide' interface. At the top, there are tabs for 'Search' and 'Results'. Below this, the '*Search Instance' is set to 'PTSF_DEFAULT'. The 'Search Text' field is empty. To the right is a 'Search' button. Below the search text are 'Start Index' (set to 1) and 'Docs Requested' (set to 10). The 'Language' is set to 'English'. There is a 'Refresh Security Cache' button and a 'Duplications' section with 'Remove' and 'Mark' checkboxes. Below these are 'Search Category to Search in' (set to 'PTPORTALREGISTRY') and a 'Navigator' field. There are buttons for 'Get Facets', 'Search Fields', and 'Display Fields'. At the bottom, there are several expandable sections: 'Custom Search Attributes', 'Filter Settings', 'Facet Request', 'Grouping /Sorting Options', and 'Additional Parameters'.

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	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. When complete, you should see a message similar to: <i>Security cache cleared successfully (262,1300)</i>
Duplicates	Enables you to control how duplicates are displayed: <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.
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.

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.
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, 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.

Image: Source XML Data

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\?

```

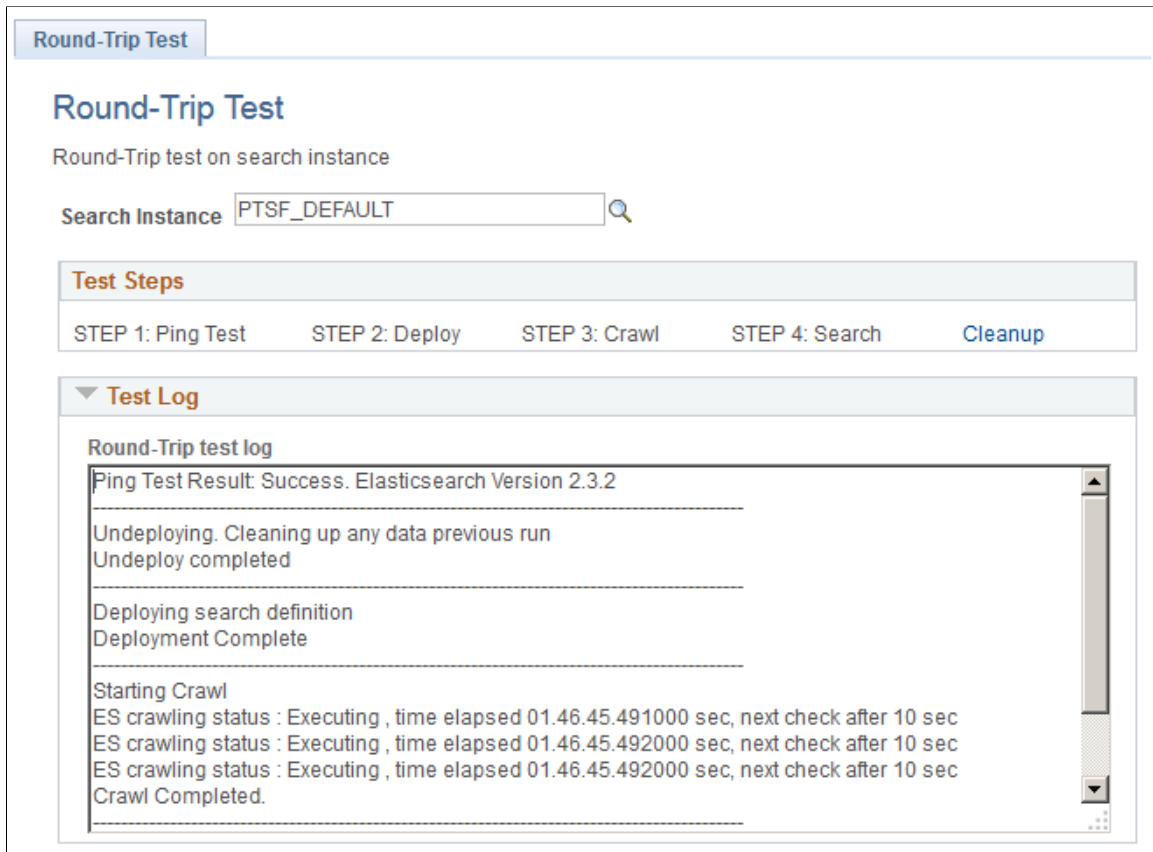
Running Diagnostics

Access the Round-Trip Test page by selecting PeopleTools, Search Framework, Search Admin Activity Guide, Testing and Utilities, 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.

Image: Round-Trip Test page

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.



Search Instance

Select the appropriate search instance defined for this PeopleSoft system.

Test Steps

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.

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.

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.

Image: Event Logs page

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.

The screenshot shows the 'Event Logs' page with the following elements:

- Event Name:** A dropdown menu set to 'All'.
- Event Severity:** A dropdown menu set to 'All'.
- Buttons:** 'View Events' and 'Purge Logs'.
- Time Period:** A section with 'From' and 'To' date pickers (05/11/2016 and 08/09/2016) and 'From Time' and 'To Time' time pickers (12:00:01AM and 11:59:59PM).
- Events Table:** A table with columns: Created, Event Name, Message Description, Message Text, and Severity. It shows one row with the number '1' in the 'Created' column.
- Page Controls:** 'Personalize | Find | View All | [Icons] | First | 1 of 1 | Last'.

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:

- **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.

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>This user ID needs to have access to all the queries (records) on all search definitions in order for the search engine to download data.</p> <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.

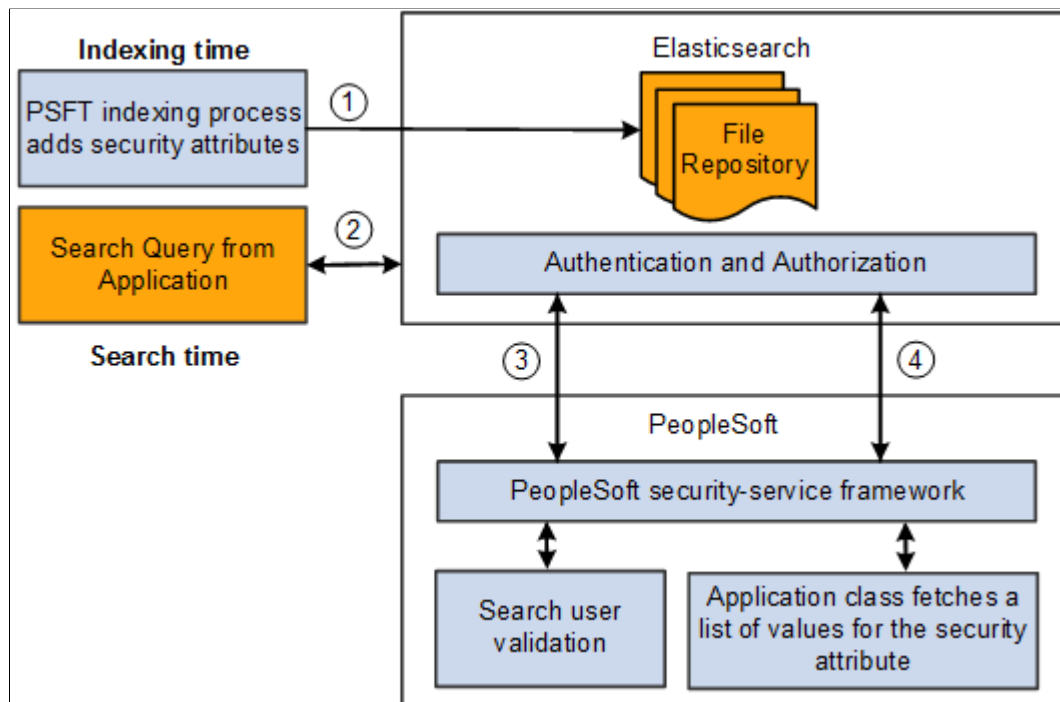
If applying any level of security, you should define an application class for every search definition. The application class is responsible for fetching a list of runtime values based on the security attributes. You associate the application class with a search definition on the Security tab of the search definition.

Authentication and Authorization in Elasticsearch

The following diagram illustrates the elements involved with authorization and authentication and the interaction between them.

Image: PeopleTools and Elasticsearch interacting to authenticate users and authorize user access to search results

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).

Step	Description
2	An application end user runs a search query.
3	Elasticsearch security framework using the <i>orcl-security-plugin</i> invokes the PeopleSoft security service for user authentication.
4	<p>Based on the authenticated user, Elasticsearch performs either a source-level or a document-level authorization using the <i>orcl-security-plugin</i>.</p> <p>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.</p>

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" (PeopleTools 8.57: 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/mykeystore.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 /home/certs/signed.cer -alias node_alias
```

7. In the `elasticsearch.yml` configuration file, add these properties to set up SSL:

orclssl.http.ssl	true
orclssl.transport.ssl	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.
orclssl.keystore	<path to keystore>
orclssl.keystore_password	<keystore password>
orclssl.truststore	<path to truststore>
orclssl.truststore_password	<truststore password>
orclssl.callback	To set this property, refer to the Call Back Properties section in Creating Search Instances .

8. Use `elasticsearchuser` script to encrypt the keystore and truststore password.

```
ES_HOME/bin> elasticsearchuser encrypt [password]
```

9. 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.

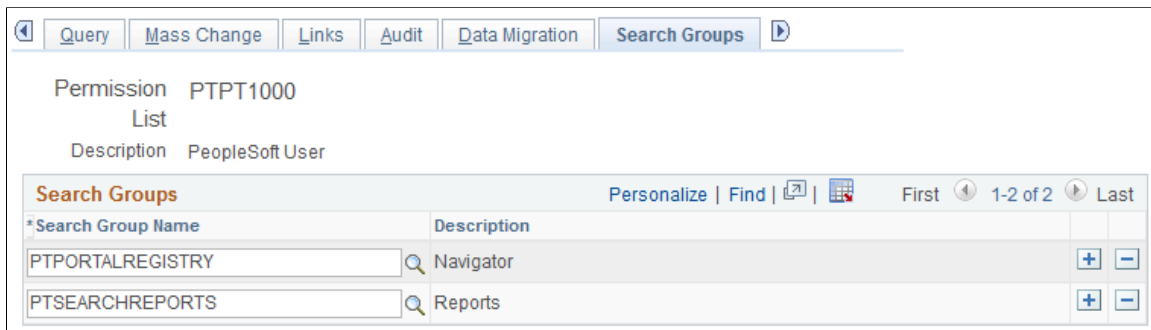
Setting Up Role-Based Search Group Access

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.

Image: PeopleTools Permission Lists: Search Groups page

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.



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](#)

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 the Elasticsearch search server. The PeopleSoft Search interfaces depend on:

- the Search Framework being completely configured.
- the PeopleSoft search definitions and categories deployed to the Elasticsearch instance.
- the search indexes are populated.

These are the PeopleSoft Search features:

Feature	Description
Global Search	Provides users a Global Search icon appearing in the header of the PeopleSoft browser session. Clicking the Global Search icon launches the Global Search Bar 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 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.

Feature	Description
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

The traditional search pages for PeopleSoft application components provided these options:

- Find an Existing Value.
- Add a New Value.

In fluid user interface, fluid search pages are rendered differently than the classic search pages. For fluid-enabled components, you can use search pages that are specific to the fluid user interface.

See "Working with Search Pages" (PeopleTools 8.57: Fluid User Interface Developer's Guide).

Image: Classic PeopleSoft Application Component 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.

The screenshot shows the search interface for the 'Department' component. At the top, the title 'Department' is displayed in blue. Below it, a instruction reads: 'Enter any information you have and click Search. Leave fields blank for a list of all values.' There are two tabs: 'Find an Existing Value' (active) and 'Add a New Value'. A 'Search Criteria' dropdown menu is visible. The search criteria section includes a 'Search by:' label, a dropdown menu with 'SetID' selected, and the text 'begins with' followed by an empty input field. Below this are two checkboxes: 'Include History' and 'Correct History'. At the bottom of the search area, there is a yellow 'Search' button and a blue link for 'Advanced Search'. At the very bottom of the interface, there are two blue links: 'Find an Existing Value' and 'Add a New Value'.

While the Find an Existing Value tab 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. The Find an Existing Value tab works in conjunction with the Keyword search, and it can be effective to provide both options to end users.

Using the Keyword search page, which you enable through the Search Framework, 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 tab 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.

Image: Find an Existing Value and Keyword Search pages

This example illustrates the fields and controls on the Find an Existing Value and Keyword Search pages. You can find definitions for the fields and controls later on this page.

The screenshot shows a web interface for 'Job Data'. At the top, it says 'Enter any information you have and click Search. Leave fields blank for a list of all values.' Below this are two tabs: 'Find an Existing Value' (selected) and 'Keyword Search'. A note states: 'Note: Keyword Search will return data updated less than 1 hour ago (08/15/2011 1:57:19PM)'. Under a 'Search Criteria' header, there is a 'Keywords' field with the text 'clerk' entered. At the bottom, there are two buttons: 'Search' and 'Advanced Search'.

The search pages that display after you've enabled Keyword search for a component are:

- Find an Existing Value Search: This is the same interface as the Find an Existing Value page, that you see for previous PeopleSoft releases.
- 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 Elasticsearch 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 both the classic search options and the Search Framework search options. That is, users can initiate a Find an Existing Value search against level 0 values, using the default PeopleTools search capabilities run directly against the database, and, they can also take advantage of the Keyword search page.

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 tab 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.
- Maintaining a consistent user experience.

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:

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 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 *if* Keyword Search is enabled.

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 *if* 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.

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" (PeopleTools 8.57: 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 "Working with Search Pages" (PeopleTools 8.57: Fluid User Interface Developer's Guide).

Working with Find an Existing Value and Keyword Search Behavior

For a component that has a search index associated with it, the user can perform these types of searches, Find and Existing Value and Keyword. Both search types will 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, where the Find an Existing Value search cannot.
- When switching between Find an Existing Value and Keyword search, the system replicates any text you've entered in one tab to the other.
- 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. 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 displays a message indicating how fresh the search results are. For example:

Note: Keyword Search will return data updated over 18 hours ago (2011-06-22-16.25.48.000000)

Note: The date format will appear according to the user's personalization setting for date display.

Maintaining a Consistent User Experience

Since both types of searches, Find an Existing Value and Keyword, form the same component search interface, these requirements ensure a consistent user experience:

- 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 in both cases.

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.

Working with Keyword Search Modes

The Keyword search page has these modes:

- Keyword-only
- Basic
- Advanced

Working with Keyword-Only Mode

The keyword-only mode can be selected in the Component Properties dialog box, and if selected only the Keywords edit box will appear as the default search mode when an end user accesses the component.

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.

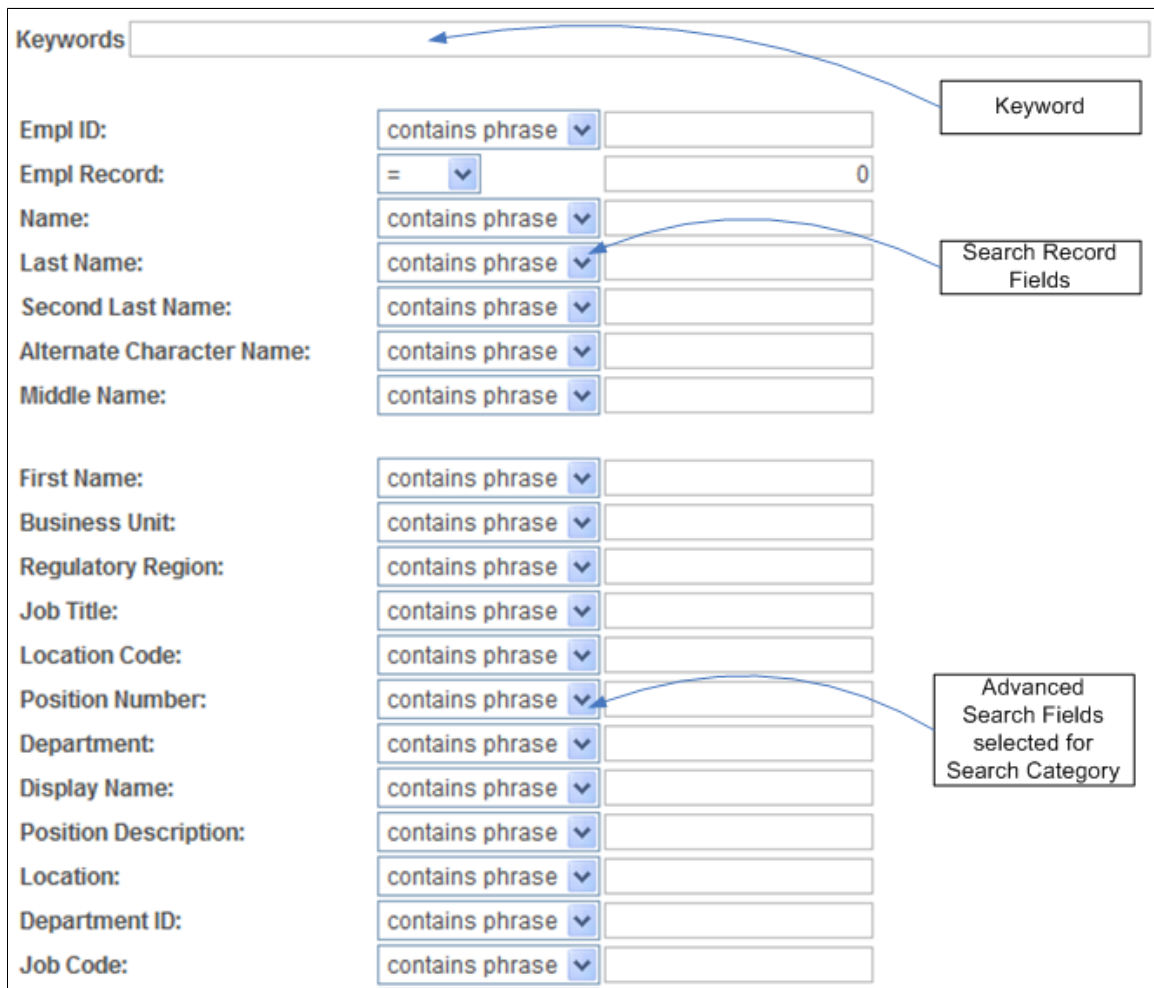
Working with Advanced Mode

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.

Image: Advanced mode search options

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 engine 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 classic component search page and entering search criteria there. In an Interaction Hub environment, Global Search can search across indexes from multiple content systems.

Image: Global Search bar

This example illustrates the search icon and the Global Search bar. You can find definitions for the fields and controls later on this page.



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 cannot be used. 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 menu search is named Navigator.

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 Navigator 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 Elasticsearch server.

For information on using ACM plug-ins, see "Running the PSRUNACM Script" (PeopleTools 8.57: 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" (PeopleTools 8.57: Portal Technology)

Working with Global Search in Fluid User Interface

In a fluid user interface, end users submit a Global Search using the Global Search icon, which is available in the header throughout the user's session irrespective of the content the user is accessing in the target frame. A user clicks the Global Search icon to display the Global Search bar, where you select a search category and enter search criteria in the search edit box.

In a fluid user interface, 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.

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.

The portal registry search category is named Navigator.

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

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.
~	maintenance~	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. maintenance~ returns search results containing words such as maintenance.

Working with the Search Results

After running a search using Keyword Search, the system presents the search results in an intuitive display, enabling interactive filtering and further navigation.

Image: Search Results

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

Personal Information

Enter any information you have and click Search. Leave fields blank for a list of all values.

Find an Existing Value **Keyword Search** Add a New Value

Note: Keyword Search will return data updated less than 1 hour ago (10/26/2012 1:14:23PM) [Search Tips](#)

Search Criteria

Keywords

Search Basic Search Advanced Search

Search Results

Filter by

City

- Alamo (1)
- Antioch (1)
- Buttontown (1)
- Canora (1)
- Concord (3)
- More...

Country

- USA (89)

Note: Keyword Search will return data updated less than 1 hour ago (10/26/2012 1:14:23PM)

50 of 89 results are displayed.

View as: [list icon] [grid icon]

1 2 3 4 5

[Brown, Alvin](#)
Last Updated Date:2012-07-01 08:01:59
This student is a BSA from Walnut Creek, CA.

[Piercy, Matthew](#)
Last Updated Date:2012-07-01 08:01:59
This student is a DEV from Walnut Creek, CA.

Personal Information
Student Enrollment

Filter by

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.

If searching the All category, you initially see the various search categories represented in the search results.

By drilling into the category, you view the facets associated with attributes in the search index.

Title

The title, which is the clickable link, enables the user to identify the search result and navigate to the associated component. The title is defined using the Title edit box on the General tab of the Search Framework, Search Designer Activity Guide, Search Definition component.

Summary

Displays a general overview of the target data.

The summary is defined using the Summary edit box on the General tab of the Search Framework, Search Designer Activity Guide, Search Definition component.

Related Actions Indicator

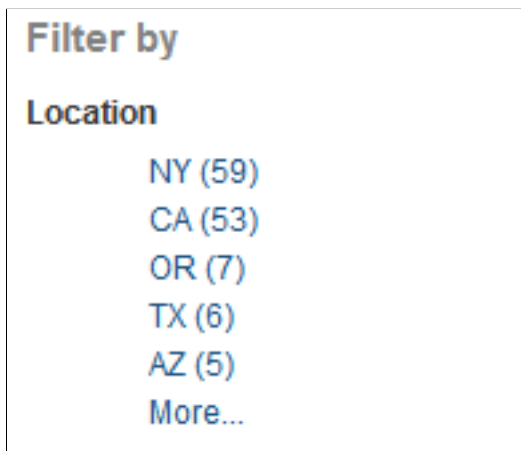
If the search definition is mapped to a component, and if there are any defined related actions associated with that component, the related actions indicator appears to the right of the search result. Clicking the related actions indicator displays a popup containing the defined related actions.

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.

Image: Filter by: Facets

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.

Image: Facet displaying More... control

This example illustrates the *More...* control used to reveal more facet values.



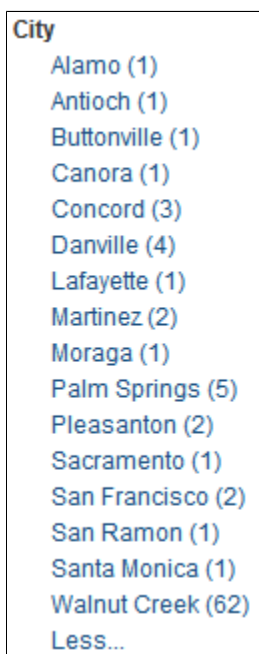
Click *More...* to expand the list.

Note: For any facet, the maximum number of values displayed is restricted to 100 for performance reasons.

Image: Facet displaying Less... control

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.



As you drill into a set of facets, your path is expressed in the form of facet breadcrumbs.

Image: Facet filter 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 View As control to toggle between the grid format and the list format for viewing search results.



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:

Image: Search Results in 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.

Note: Keyword Search will return results last updated over 1 day ago (07/25/2011 11:54:50AM)

View as:  

[View All](#) First  1-100 of 199  Last

Empl ID	Empl Record	Name	First Name	Last Name	Second Last Name	Alternate Character Name	Middle Name
HOWS001	0	Melanbacher, Mary Jane	Mary Jane	Melanbacher	(blank)	(blank)	(blank)
MUET219	0	McGuinness, Randy	Randy	McGuinness	(blank)	(blank)	(blank)
HOWS007	0	Collins, Genie	Genie	Collins	(blank)	(blank)	(blank)
HOWS002	0	Rainnie, Jennifer	Jennifer	Rainnie	(blank)	(blank)	(blank)
HOWS005	0	McDougall, Perry	Perry	McDougall	(blank)	(blank)	(blank)
HOWS003	0	Donaldson, Terry Anne	Terry Anne	Donaldson	(blank)	(blank)	(blank)
HOWS006	0	Lavoie, Michael	Michael	Lavoie	(blank)	(blank)	(blank)

To navigate to the PIA page loading the desired data, click the key column with the link.

The Related Actions link does not appear for search results in the grid format.

Note: The default search result display format for Keyword Search is the grid format.





Working with the List Format

The following example displays the list format:

Image: Search Results in 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.

Note: Keyword Search will return results last updated 19 days ago (11/06/2012 8:21:36PM)
50 of 89 results are displayed.

View as:    **1** 2 3 4 5 

[Brown, Alvin](#)
Last Updated Date:2012-07-01 08:01:59
This student is a BSA from Walnut Creek, CA.

[Piercy, Matthew](#)
Last Updated Date:2012-07-01 08:01:59
This student is a DEV from Walnut Creek, CA.

[Martin, Peter](#)
Last Updated Date:2012-07-01 08:01:59
This student is a DBA from Palm Springs, CA.

[Snyder, Joyce](#)
Last Updated Date:2012-07-01 08:01:59
This student is a APP from Palm Springs, CA.

[Barker, Janet](#)
Last Updated Date:2012-07-01 08:01:59
This student is a DEV from Palm Springs, CA.

To navigate to the PIA page loading the desired data, click the Title line of the search result.

The Related Actions link appears only for search results in the list format.

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. For every search result that has related actions defined for it, you see a related actions indicator to the right of the search result. The related actions indicator is a greybar with a white triangle within it pointing to the search result.

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.

Image: Related actions indicator

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.



Note: The related action indicator appears only in Global Search results and in Keyword Search results when viewing the list format.

By clicking on the related actions indicator, a related actions popup appears, showing the possible related actions for that user, according to security access.

Image: Related action indicator popup


This example illustrates the fields and controls on the Related action indicator popup. You can find definitions for the fields and controls later on this page.



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.

Image: Related action from a search result

This example illustrates the user accessing the related action from the search result, with no navigation.



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 icon available on the header to launch the Global Search bar 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.

Image: Global Search bar

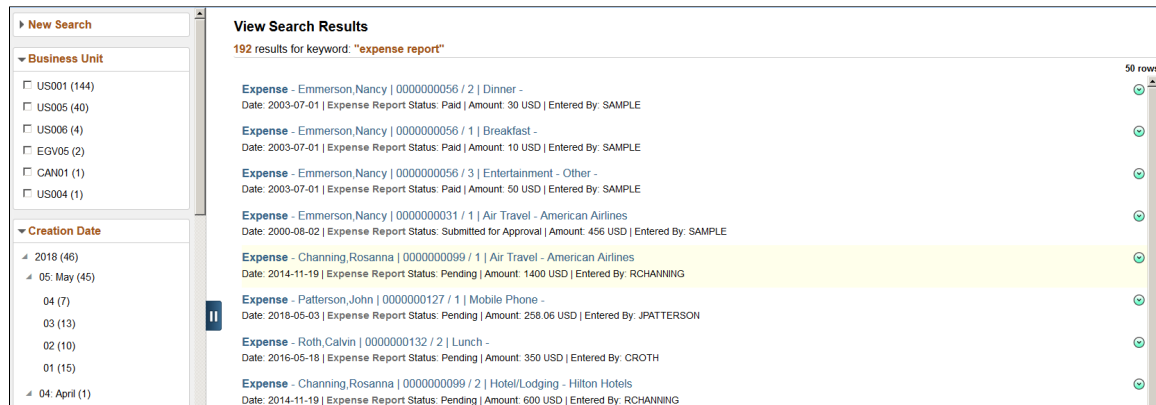
This example illustrates the Global Search bar.



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.

Image: Global Search results page

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.



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.

New Search

The new search area enables you to initiate a new search by selecting a search category and entering keywords.

Note: A blank search is not supported. That is, you must enter valid search criteria.

Facets

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.



Use the Hide Filters button to collapse the left panel so that the right panel occupies the entire real estate of the page.

Toggle the button to expand the left panel.

Number of returned rows

The total number of rows that match a keyword in the selected search category is displayed at the top of the search page.

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.

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.

If any bulk related actions are defined for a component, the bulk related actions indicator appears at the top of the search results page.



Images

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.

See [Working with Images](#).

Related Actions Indicator

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.

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.

Image: Advanced search fields - Global Search

This example illustrates the advanced search fields.

▼ New Search

Category
Navigator ▾

Keywords

With this Exact Phrase

With any of these Words

Exclude

Portal Label (contains phrase)

Long Description (contains phras

Folder (contains phrase)

Source Application (contains phra

< >

[Fewer Options](#) [Show Operators](#)

Using Facets

Beginning with PeopleTools 8.57, 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 except for hierarchical facets. 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, you can also choose multiple facet nodes, but you should enable the multiple facet-values selection option in a search category for a particular facet. In the case of a hierarchical facet, the facet nodes are displayed as a tree, and a green check mark indicates the selected facet node.

Note: The Category facet is set for single select.

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, a single breadcrumb is displayed with the number of selected nodes displayed (in parenthesis) in the breadcrumb.

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*.

Image: Global Search results showing multiple facet-values selection

This example illustrates selecting multiple facet values.

View Search Results
36 results for keyword: "administration"

Finance and Administration (36) Administrative Support (6) Clear All

Job Posting - Assistant Administrator - facet2 - 300901 | External Posting
Recruiting Location: Delaware Operations | Department: Finance and Administration | Job Family: Administrative Support

Job Posting - Administration Assistant - facet3 - 300902 | Internal Posting
Recruiting Location: Delaware Operations | Department: Finance and Administration | Job Family: Administrative Support

Job Posting - Assistant Administrator - facet2 - 300901 | Internal Posting
Recruiting Location: Delaware Operations | Department: Finance and Administration | Job Family: Administrative Support

Job Posting - Administration Assistant - facet3 - 300902 | External Posting
Recruiting Location: Delaware Operations | Department: Finance and Administration | Job Family: Administrative Support

Job Posting - Administration Assistant - Multiple - 300911 | External Posting
Recruiting Location: Delaware Operations | Department: Finance and Administration | Job Family: Administrative Support

Job Posting - Manager-Employee Relations - 10240 | Internal Posting
Recruiting Location: Delaware Operations | Department: Finance and Administration | Job Family: Administrative Support

Job Posting - Administration Assistant - Multiple - 300911 | Internal Posting
Recruiting Location: Delaware Operations | Department: Finance and Administration | Job Family: Administrative Support

Job Posting - Manager-Employee Relations - 10240 | External Posting
Recruiting Location: Delaware Operations | Department: Finance and Administration | Job Family: Administrative Support

Job Posting - Send Documents JO - 300235 | External Posting
Recruiting Location: Delaware Operations | Department: Finance and Administration | Job Family: Administrative Support

Job Posting - Management - Staff Assistant - facet1 - 300900 | External Posting
Recruiting Location: Delaware Operations | Department: Finance and Administration | Job Family: Administrative Support

Job Posting - Management - Staff Assistant - facet1 - 300900 | Internal Posting
Recruiting Location: Delaware Operations | Department: Finance and Administration | Job Family: Administrative Support

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.

Image: Related Actions indicator


This example illustrates the Related Actions indicator.



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.

Image: Bulk mode icon

This example illustrates the Bulk related actions indicator.



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" (PeopleTools 8.57: Fluid User Interface Developer's Guide) and "Setting Component Properties for Fluid Components" (PeopleTools 8.57: Fluid User Interface Developer's Guide).

Image: Real-time Component Search Results page

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.

SetID	Actions	Item	Item Group	Category ID	Desc	Family
CRM01	<input checked="" type="checkbox"/>	10000	REFRIGERATOR	00004	Refrigerator, Plastic Bins	REF ACCSRS
CRM01	<input checked="" type="checkbox"/>	10001	REFRIGERATOR	00004	Refrigerator, Shelves	REF ACCSRS
CRM01	<input checked="" type="checkbox"/>	10002	REFRIGERATOR	00004	Refrigerator, Ice Trays	REF ACCSRS
CRM01	<input checked="" type="checkbox"/>	10003	REFRIGERATOR	00004	Refrigerator, Ice Dispenser As	REF ACCSRS
CRM01	<input checked="" type="checkbox"/>	10004	REFRIGERATOR	00004	Refrigerator, Compressor	REF REP PT

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.

Grid view

Displays the search results in a grid format.

In a grid view, related actions, if defined, always appears as the second column of the grid.

In a grid view, bulk related actions, if defined, always appear as the first column of the grid.

List view

Displays the search results in a list.

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.

When you click Bulk mode, the bulk related actions check box appears to the right of the search result.

Bulk mode

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.

Use the Bulk mode to toggle between displaying bulk related actions and exiting the bulk mode.

Add

Enables you to add new rows to PeopleSoft tables.

The Add button is similar to the Add a New Value page in a classic environment.

Note: The Add button is not available on small form factor devices.

Chart

Enables you to view the search results in a chart format.

Options Menu

Use the Options Menu to:

- 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.

Related actions indicator

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.

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" (PeopleTools 8.57: 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.

Image: Master Detail Real-time Component Search page

This example illustrates the fields and controls on the Master Detail Real-time Component Search page. You can find definitions for the fields and controls later on this 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" (PeopleTools 8.57: Pivot Grid)

"Creating a Component Pivot Grid Model Using the Pivot Grid Wizard" (PeopleTools 8.57: 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" (PeopleTools 8.57: 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.

Image: Keyword Search page

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.

The screenshot shows a web browser window with a dark blue header. On the left, there is a navigation menu with a 'New Search' dropdown. The main content area is titled 'View ITEMS' and contains a search form. The form has the following fields: 'Keywords' (text input), 'SetID (=)' (text input with value 'CRM01' and a search icon), 'Item ID (contains phrase)' (text input with a search icon), 'Description (contains phrase)' (text input), and 'Item Group (contains phrase)' (text input). Below the form are two buttons: 'Search' (green) and 'Clear' (grey). At the bottom of the form, there are links for 'Fewer Options' and 'Show Operators'. The browser's address bar shows 'My Homepage' and 'ITEMS'.

Working with Search Results on a Touch Friendly Device

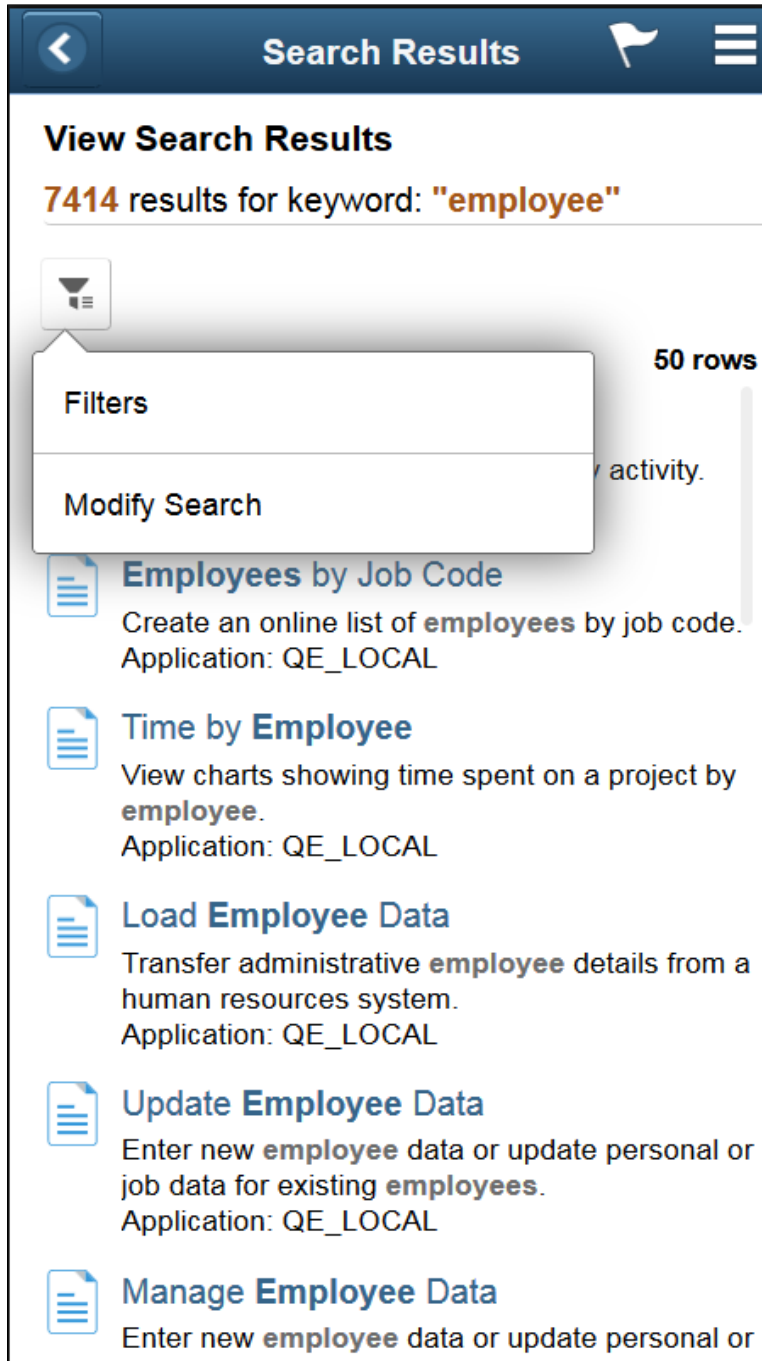
On small form factor devices, the search results display varies from desktop or tablet display.

The following search results features are available on small form factor devices:

- Initiate global search using the Search option in the Actions List menu; Global Search icon is not available on the header.
- 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.

Image: Global Search Results page

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.



Click the options icon to display the Filters and Modify Search options.

Filters

Displays the available facets and enables you to filter the search results.

Modify Search



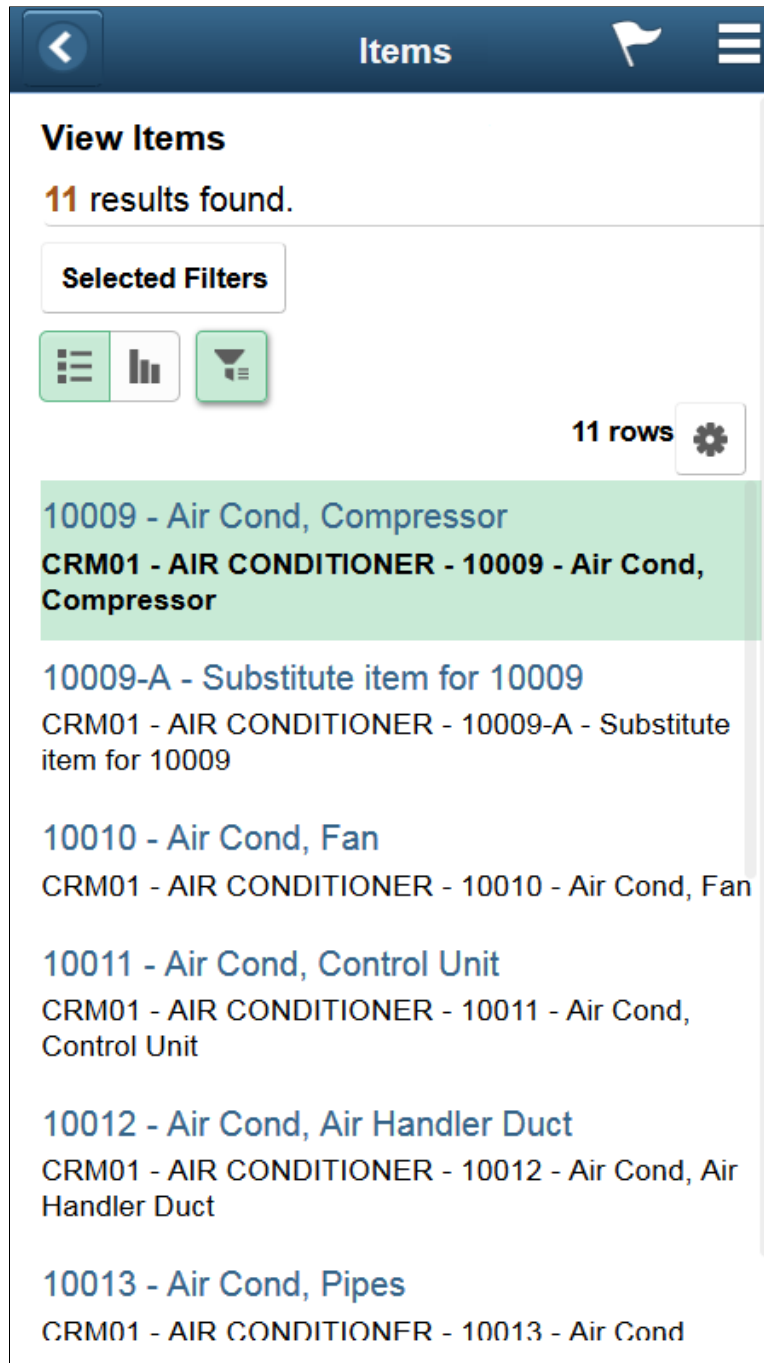
Displays the New Search panel that allows you to enter search criteria.

On the search results page, use the back button to return to the previous screen.

For example, when you drill into a search result, use the back button to return to the search results page.

Image: Component search - search result drill down

This example illustrates the fields and controls on the Component search - search result drill down. You can find definitions for the fields and controls later on this page.

**Selected Filters**

Click Selected Filters to display the filters that are applied. The pop-up allows you to clear the filters individually or all the filters.

Show Filters

On a master detail component, use the Show Filters button to display the search results. Toggle the Show Filters icon to collapse the filter pane.

Previous in List and Next in List

When you drill down to a search result, you can use the Previous in List and Next in List buttons to traverse the results.

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](#)

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. In PeopleTools 8.57, Search Framework uses Elasticsearch 6.1.2 and Kibana 6.1.2.

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. 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 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 User Guide [6.1], 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.
Index pattern	“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.” Kibana User Guide [6.1], Getting Started, Defining Your Index Patterns.
Attribute	An attribute in PeopleSoft Search Framework corresponds to a search field in Elasticsearch and Kibana.

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 8.57 and Elasticsearch 6.1.2.

Setting Up Kibana

To set up Kibana to provide data on Search Framework and Elasticsearch, 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 8.57 and Elasticsearch 6.1.2 are installed in your environment.

Installing Kibana

Search Framework supports the installation of Kibana on all platforms that are certified by PeopleSoft for Elasticsearch.

To install Kibana, you use the delivered Elasticsearch DPK. For instructions on installing Kibana, refer to the PeopleSoft Deployment Packages for Elasticsearch Installation (Doc ID 2205540.2) on My Oracle Support.

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.

Connecting Kibana to Elasticsearch

When you use the Elasticsearch 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.

After installing Kibana, you should manually set up the connection between the Elasticsearch search instance in PeopleSoft and the Kibana monitoring server using the Elasticsearch Monitoring Server page. For instructions on using the Elasticsearch Monitoring Server page, see [Using the Elasticsearch Monitoring Server Page](#).

When you are using 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 Elasticsearch Monitoring Server page.

For Kibana to be always available, you should ensure that the Elasticsearch server it connects to is always up and running.

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. When the service status displays as `Listening`, you can access the Kibana service in a browser using the URL displayed in the command prompt.

The default port for Kibana is 5601, but you can change it by modifying the value in the `kibana.yml` configuration file.

```
http://example.com:5601
```


The Kibana URL prompts you for an user name and password. You must enter the Elasticsearch administrator user name, that is, `esadmin` and the password, which you entered while installing Elasticsearch.

On the Linux platform, after you install Kibana, the Kibana service is automatically started. If the Kibana service is not started, you can start the service by running `kibana.sh`

An example of the Kibana home page:

Image: Kibana home page

This example illustrates the Kibana home page which you access using the URL.



Using the Elasticsearch Monitoring Server Page

Use the Elasticsearch Monitoring Server page to establish a connection between Kibana and the Elasticsearch server where you want to index the system and indexing metrics.

Access the Elasticsearch Monitoring Server page by selecting **PeopleTools >Search Framework >Administration >Monitoring Server**.

Image: Elasticsearch Monitoring Server page

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

Elasticsearch Monitoring Server

*Host Name

*Port

SSL Option

*User Name

*Password

*Scan Interval(Sec)

*Purge Retention(Days)

Enable Index Metrics Yes

Enable System Metrics Yes

Host Name

Enter the Elasticsearch server where you want to index the system and indexing metrics.

Note: This should be the Elasticsearch server 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.

Scan Interval(Seconds)

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.

Purge Retention Period(Days)

Specify the number of days to store the system and indexing metrics. For example, if you enter 7, the metrics data will be retained for 7 day.

Enable Index Metrics

Select Yes if you want indexing metrics to be collected for viewing and analyzing in the Indexing Metrics dashboard.

The default is No.

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: `ptsf_es_index_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.

Enable System Metrics

Select Yes if you want system metrics to be collected for viewing and analyzing in the System Metrics dashboard.

The default is No.

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.

Save

When you select Save:

- 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 (`.kibana`) gets created on the Elasticsearch server. The index contains the metadata required for Kibana.

Note: To save the configurations, the Elasticsearch server must be a running instance of Elasticsearch.

Using Dashboards and Visualizations in Kibana

In PeopleTools 8.57, 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 User Guide [6.1], which is available on www.elastic.co.

Understanding Dashboard 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 dashboards, and each of the dashboards display two or more visualizations. For a list of delivered dashboards and visualizations, see [Delivered Dashboards and Visualizations](#).

You may build dashboards based on your business requirements. For instructions on building dashboards, refer to the Kibana User Guide [6.1], Dashboard, Building a Dashboard, which is available on www.elastic.co.

Delivered Dashboards and Visualizations

In PeopleTools 8.57, 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

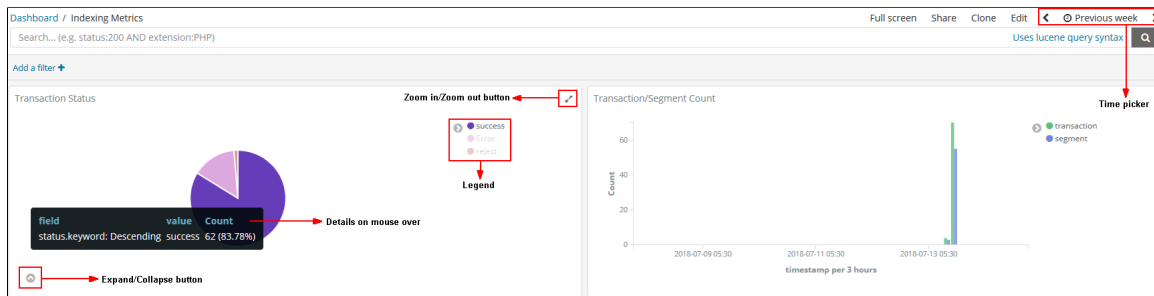
- System Metrics —provides an overview of the health of the system— and displays the following visualizations:
 - CPU (%) Utilization
 - Disk Space Utilization
 - JVM (%) Utilization
 - Total Indexed Documents

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 User Guide [6.1], which is available on www.elastic.co.

Image: Common Elements on Visualizations

This example illustrates the common elements on Visualizations. You can find definitions for these common elements later on this page.



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 User Guide [6.1], Discover, Setting the Time Filter.

Zoom in/Zoom out button

Use the Zoom out 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 of the slice or bar.

Expand/Collapse button

Use the Expand button to display the raw data behind a visualization. It also enables you to see the raw request sent to the Elasticsearch server, and the raw response from the server.

For more information, refer to Kibana User Guide [6.1], Visualize, Visualization Spy; Dashboard, Viewing Visualization Data.

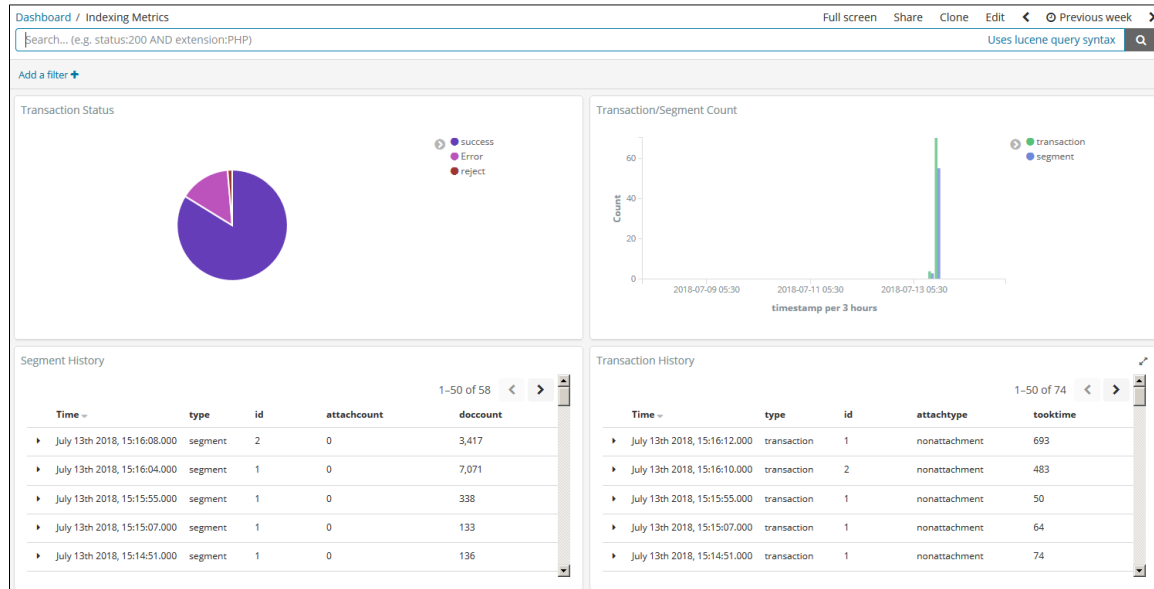
Using the Indexing Metrics Dashboard and Visualizations

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.

Image: Indexing Metrics dashboard

This example illustrates the visualizations on the Indexing Metrics dashboard. You can find definitions for the visualizations later on this page.



Transaction Status

Displays the status of transactions as a pie-chart. The transaction statuses are:

- 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.

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.

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:

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.

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.

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.

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.

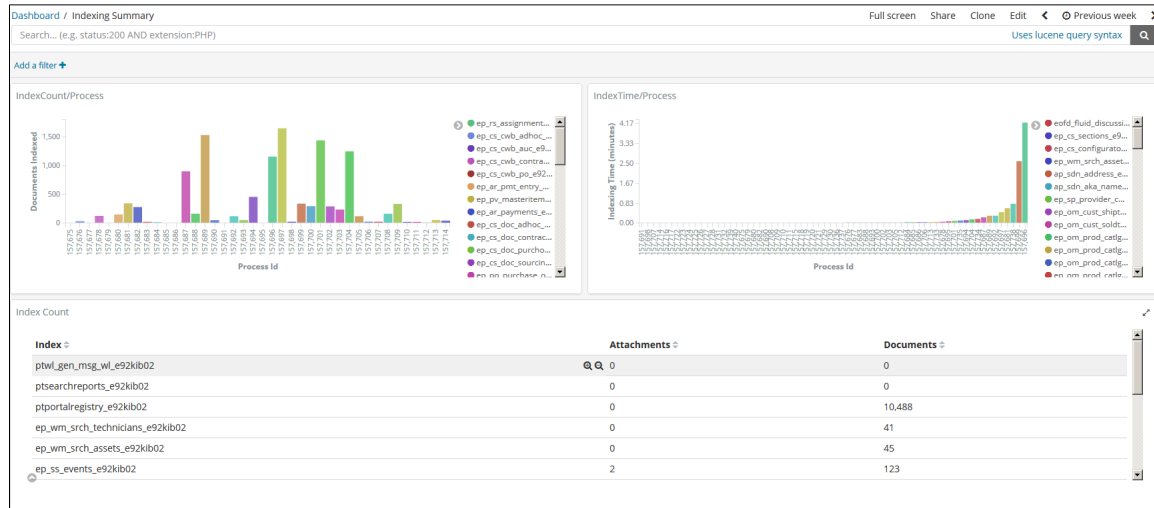
Using the Indexing Summary Dashboard and Visualizations

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.

Image: Indexing Summary dashboard

This example illustrates the visualizations on the Indexing Summary dashboard. You can find definitions for these visualizations later on this page.



IndexCount/Process

Displays the documents indexed for an index and the process instance ID as a vertical bar chart.

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 `indexname:ep_ap_payments*`.

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.

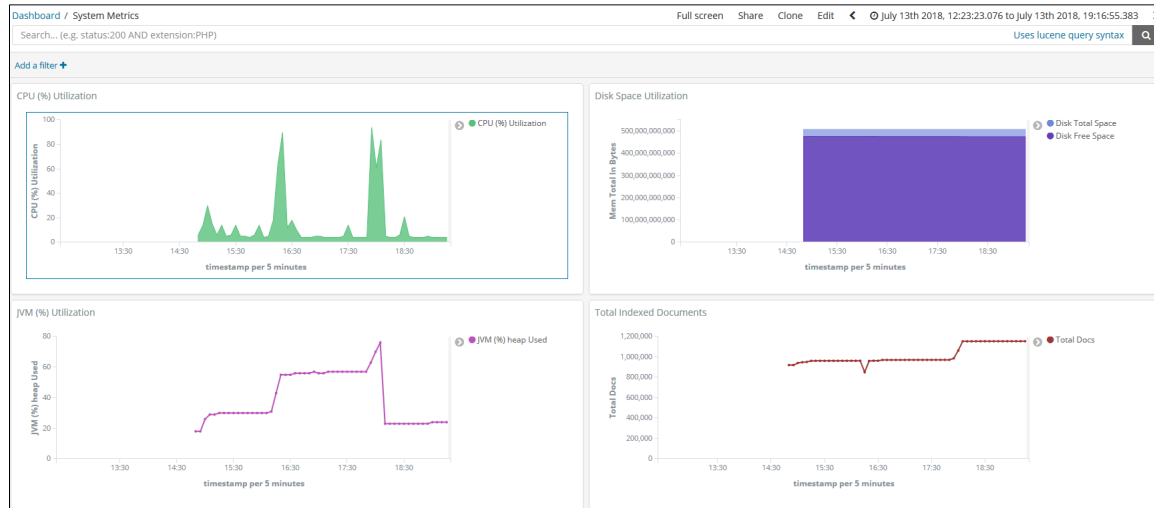
Using the System Metrics Dashboard and Visualization

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.

Image: System Metrics dashboard

This example illustrates the visualizations on the System Metrics dashboard. You can find definitions for the visualizations later on this page.



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.

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 Visualization

Oracle 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 a 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 [Setting Up 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 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 User Guide [6.1], 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 User Guide [6.1], Getting Started, Defining Your Index Patterns and for creating visualizations, refer to Kibana User Guide [6.1], Visualize, Creating a Visualization.

Monitoring Elasticsearch Cluster Using PeopleSoft Health Center

Understanding Elasticsearch Cluster

The PeopleSoft Health Center is the monitoring framework that provides monitoring and performance data of the domains within a PeopleSoft implementation, including the Elasticsearch cluster.

The monitoring and diagnostic dashboards for an Elasticsearch cluster enable users to understand the performance of the Elasticsearch cluster and the possible performance issues that may arise.

PeopleSoft Health Center provides a snapshot of the Elasticsearch cluster and the nodes in the cluster, however users cannot update the state or run tasks on the cluster. The data displayed includes the number of nodes, the number of indexes, the number of documents, among other performance data.

Related Links

"Understanding PeopleSoft Health Center" (PeopleTools 8.57: Performance Monitor)

Viewing the Elasticsearch Cluster Dashboard

In PeopleSoft Health Center, the Application Dashboard is the principal dashboard that provides a snapshot or summary of the health of the domains and the Elasticsearch cluster. Detailed information of the Elasticsearch cluster is available on a dashboard, which you access by clicking the Elasticsearch cluster.

Application Dashboard

This topic describes the snapshot of an Elasticsearch cluster available on the Application Dashboard.

Access the Application Dashboard (PeopleTools, Health Center, Health Center).

Image: Application Dashboard - ES Cluster Summary section

This example illustrates the fields and controls on the Application Dashboard - ES Cluster Summary section. You can find definitions for the fields and controls later on this page.

PeopleTools 8.56.00.000 PeopleTools 8.56-803-R1 Database: T56JD803		Ok Health Status	Light Load Status			
Alerts		Emergency	Urgent	Warning		
It's all good! You don't have any new alerts.						
Web Servers						
peoplesoft slc08aka:8000:8001	Health Status	Load Status	0.00 CPU %	0.00 Heap Size (MB)	0.00 Heap Used %	0 Threads
Application Server Domains						
T56JD803 slc08aka:9000	Ok Health Status	Light Load Status	22.79 CPU %	71.76 Mem Used %	30 Servers	0 Queued
Process Scheduler Domains						
T56JD803 slc08aka:	Ok Health Status	Light Load Status	18.72 CPU %	71.76 Mem Used %	10 Servers	0 Queued
ES Cluster Summary						
ES_CLUSTER2 Cluster Name	green Status	3 Nodes	1280572 Documents #	0 CPU (%)	0.09 Heap (%)	1398.83 Free Disk Space (GB)

In PeopleSoft Health Center, the Application Dashboard displays a summary of the Elasticsearch cluster that includes the following information.

Cluster Name

The name of the Elasticsearch cluster.

You specify a cluster name when you install Elasticsearch. To change the cluster name, you need to specify a new name in the *elasticsearch.yml* configuration file.

Status

The health 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.

Nodes

Total number of nodes in the cluster.

Documents #

Total number of documents in the cluster.

CPU (%)

CPU utilization (in percentage) by JVM.

Heap (%)

Total of average JVM percent used in the cluster.

Free Disk Space (GB)

Available disk space for the entire cluster.

Cluster Details Page

While the ES Cluster Summary section in the Application Dashboard provides a high-level status of the cluster, the Cluster Details page provides detailed information on the cluster, nodes and indices.

Access the Cluster Details page (PeopleTools, Health Center, Health Center and on the Application Dashboard, click any of the ES Cluster Summary values).

Image: Cluster Details page

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

Cluster Summary

Cluster Name: ES_CLUSTER2 Status: green Nodes: 3 Index: 36 Used Data: 1.28 GB Free Data: 1.38 TB CPU (%): 3
Documents #: 1097666 Up Time (Hrs): 62.83 Primary Shards: 180 Replica Shards: 0.08333333333333333

Nodes: 3 Cache:
Master Only: 0 Data Only: 0 Master-Data: 3 Client: 0 Field Data: 324.00 KB Query Cache: 41.60 KB

Nodes Summary

Node Details	Documents #	Disk Utilization	CPU (%)	Heap (%)	Up Time (Hrs)	More
Name: SLC10WRZ Host: 10.196.19.137 Address: 10.196.19.137:9300	346116	Used Data: 413.61 MB Free Data: 471.43 GB	1	13	62.83	More...

Indices Summary

Index Name

Index	Documents #	Primary Shards	Replica Shards
H21PBGL2 (ORCL_ES_DEFAULTINDEX)	29968	5	0
GP_PL_SRCH	11145		
PTPORTALREGISTRY	9944		
HC_HRS_JOB_OPENING	6190		
HC_HR_COMPANY_DIRECTORY1	795		
HC_JPM_PERSON_PROFILE	643		

The Cluster Details page provides information on the cluster, nodes, and indices.

Cluster Summary

Cluster Name	The name of the Elasticsearch cluster.
Status	Health status of the cluster is color coded - Green, Yellow, Red.
Nodes	The total number of nodes in the cluster.
Index	The total number of indexes in the cluster.
Used Data	Indicates the memory used by the cluster.
Free Data	Indicates the available memory in the cluster.
CPU%	The percentage of CPU used by JVM.
Documents #	Total number of documents in the cluster.
Up Time (Hrs)	Total time in hours since the cluster is up.
Primary Shards	Total number of primary shards in the cluster.

Replica Shards

Total number of replica shards in the cluster.

Nodes

Displays a breakdown of the nodes:

- **Master Only:** A master node controls a cluster.
- **Data Only:** A data node holds the shards that contain the documents you have indexed. A data node handles data related operations such as CRUD, search, and aggregations, which are memory and CPU intensive. It is important to monitor these resources and to add more data nodes if they are overloaded.
- **Master Data:** A master-data node acts both as a master node and data node. That is, it maintains the meta data of a cluster and it also stores data.
- **Client:** A client node only route requests, that is, routes cluster-level requests to the master node and data related requests to the data nodes. Essentially, a client node behaves as smart load balancers.

Cache

- **Field Data:** The field data cache is used mainly when sorting on or computing aggregations on a field. It loads all the field values to memory in order to provide fast document-based access to those values.
- **Query Cache:** The query cache is responsible for caching the results of queries. There is one query cache per node that is shared by all shards. The cache implements an LRU eviction policy, that is, when a cache becomes full, the least recently used data is evicted to make way for new data.

Nodes Summary**Nodes Details**

Contains details of the node name, the host information, and the address.

Documents #

Total number of documents in a node.

Disk Utilization

Contains information on used memory and available memory.

CPU (%)

The percentage of CPU utilization by the node.

Heap (%)

Percentage of JVM used by the node.

Up Time (Hrs)

Time in hours since the host is up and running.

Indices Summary**Index**

The name of the index.

Documents #	Total number of documents residing in an index and the number of documents in a search definition.
Primary Shards	Number of primary shards. You specify the number of primary shards in the General Options page. See Managing General Search Options for more information on the use of primary shards.
Replica Shards	Number of replica shards. You specify the number of replica shards in the General Options page. See Managing General Search Options for more information on the use of replica shards.

Viewing the Node Dashboard

Elasticsearch provides an array of statistics of each node in a cluster. Depending on your implementation of the search functionality, you can have multiple nodes within the Elasticsearch cluster.

Access the Node Details page (PeopleTools, Health Center, Health Center, and on the Cluster Details page, click the More link for a node).

For debugging purposes, you may refer to [Elasticsearch: The Definitive Guide \[2.x\], Administration, Monitoring, and Deployment](#) on the Elasticsearch website (www.elastic.co).

Summary Section

The Summary section of the node dashboard provides general information on the Elasticsearch server, which are useful for debugging purposes.

Image: Summary section

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

Summary	
Node Name:	SLC10WRZ
IP Address:	slc10wrz.us.oracle.com/10.196.19.137:9300
Host:	10.196.19.137
Up Time:	2.54
Data Path:	/ds1/products/es_data/data/ES_CLUSTER2/nodes/0
Logs:	/ds1/products/es_logs/logs

Node Name The name of the node in the Elasticsearch cluster.

IP Address IP address of the node.

Host Host address.

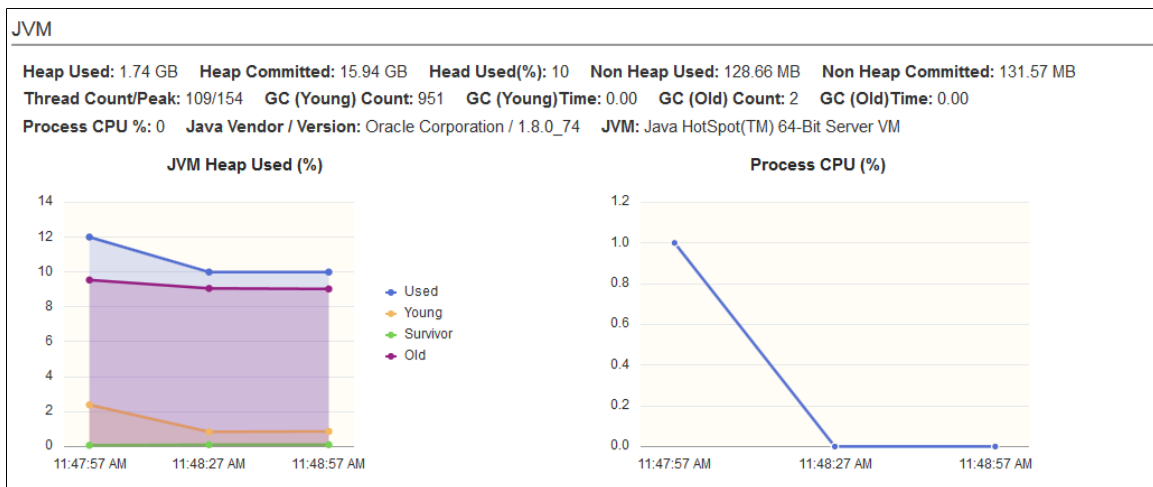
Up Time	Number of hours the node is up and running in the cluster.
Data Path	Location of the cluster data.
Logs	Location of the Elasticsearch logs.

JVM Section

The JVM section of the node dashboard provides critical information about the JVM that is running the Elasticsearch server.

Image: JVM section

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



Heap Used

The amount of heap memory used by JVM.

Heap Committed

The amount of memory actually allocated to the JVM process.

Heap Used (%)

The amount of heap memory used by JVM in percentage.

Elasticsearch initiates garbage collections when the heap percentage is 75%.

Note: If the heap used percentage is consistently greater than or equal to 75%, it is a warning sign that slow garbage collections may occur.

If the heap used percentage is consistently greater than or equal to 85%, you need to increase the memory (maximum of 30 GB per node) or add nodes to the cluster.

If heap used percentage is greater than 90–95%, risk of very poor performance that may result in very slow garbage collections and out of memory exceptions.

For more information on heap used percentage, see [Elasticsearch: The Definitive Guide \[2.x\], Monitoring Individual Node on the Elasticsearch website \(www.elastic.co\)](#).

Non Heap Used

Non heap memory is where classes are loaded and is controlled by JVM application automatically. This parameter shows how much non-heap memory is used.

Non Heap Committed

The amount of non heap memory allocated.

Thread Count/Peak

JVM Thread count and peak count.

GC (Young) Count

Shows the garbage collections count for the Young generation. This is usually large.

GC (Young) Time

Shows the garbage collections cumulative time in milliseconds for the Young generation.

GC (Old) Count

Shows the garbage collections count for the Old generation. This value should be small.

For more details on garbage collections, see [Elasticsearch: The Definitive Guide \[2.x\], Monitoring Individual Node on the Elasticsearch website \(www.elastic.co\)](#).

GC (Old) Time

Shows the garbage collections cumulative time in milliseconds for the Old generation. The garbage collection time should be small.

Note: Search administrators must periodically collect garbage collection counts and time.

Process CPU %

CPU usage in percentage. A chart also displays the CPU usage percentage.

Java Vendor/Version

Displays the vendor and version for Java.

JVM

Displays generic information about the JVM.

JVM Heap Used (%)

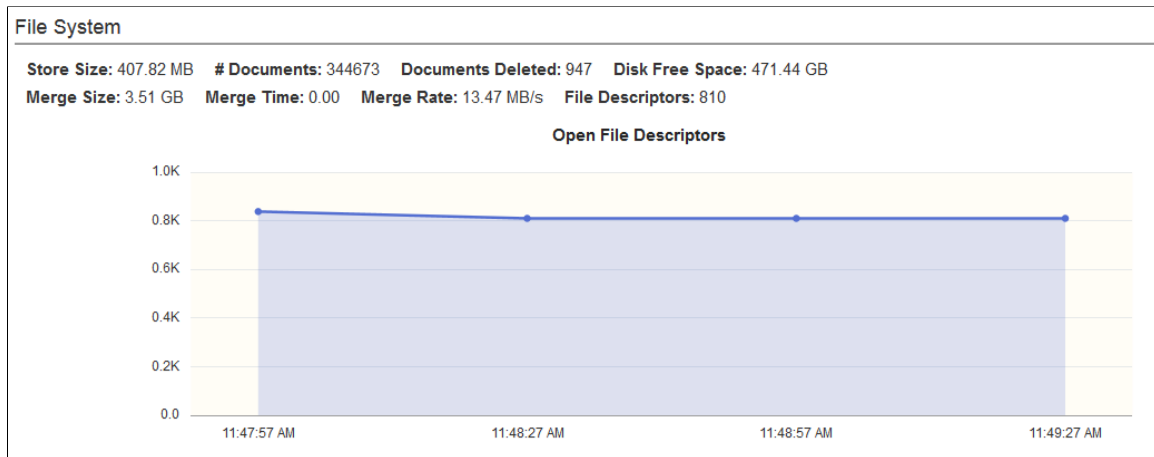
The chart displays breakdown of memory usage by each generation — Young, Survivor, Old.

File System Section

The File System section provides statistics on the file system in the node.

Image: File System section

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




- Store Size** Physical storage consumed by the node. It includes both primary and replica shards.
- # Documents** Total number of documents that reside on the node.
- Documents Deleted** Number of deleted documents that are not purged from the segments yet.
- Disk Free Space** Free space available in the node.
- Merge Size** Shows the cumulative size of segments being merged.
- Merge Time** Shows the time taken to complete the merge.
- Merge Rate** Shows the rate at which the merge is completed.
- File Descriptors** Shows the number of opened file descriptors associated with the current process.

Index Activity Section

The Index Activity section provides aggregate statistics for all indices on the node.

Image: Index Activity section

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



Indexing - Index

Shows the average time taken for indexing in milliseconds.

Indexing - Delete

Shows the average time taken (in milliseconds) for deleting documents from an index.

Search - Query

Shows the average time spent on queries (per request) since the node was started.

If query time is high, you may need to tune or optimize your search server and optimize your queries too.

For more details, see *Elasticsearch: The Definitive Guide [2.x]*, *Monitoring Individual Node* on the Elasticsearch website (www.elastic.co).

Search - Fetch

Shows the average time taken to fetch data as part of the query process.

If the fetch time is more than the query time, it may indicate slow disks or that large documents are fetched.

Get - Total

Shows the average time taken (in milliseconds) for Get requests.

Refresh

Shows the average time taken (in milliseconds) to refresh requests.

Flush


Shows the average time taken (in milliseconds) for each of the flush requests.

Cache Activity Section

The Cache Activity section displays details of the cache in the node.

Image: Cache Activity section

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



Field Size

Shows the memory used by field data, which is used for aggregations, sorting, and so on.

Field Evictions

Shows the number of times a field has been evicted. This value should be zero or a value close to zero.

Filter Cache Size

Shows the amount of memory used by cached filter bitsets.

Filter Evictions

Shows the number of times a filter has been evicted.

Cache Size

The allocated memory for cache.

Cache Evictions


Shows the number of times a query cache has been evicted to make way for new data.

Network Section

The Network section provides statistics of the HTTP port (usually 9200).

Image: Network section

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

**HTTP Current Open**

Shows the number of HTTP clients currently open.

HTTP Total Opened

Shows the total number of HTTP clients opened. If this value keeps increasing constantly, you need to be aware that the HTTP clients may not be configured correctly, that is, the HTTP clients may be not using keep-alive connections, which is necessary for better performance.

Users and Roles in Elasticsearch

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](#)

Appendix B

Troubleshooting

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>\websrv\<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
<code>{"error":{"root_cause":[{"type":"orcl_auth_plugin_exception","reason":"Invalid user/password"}],"type":"orcl_auth_plugin_exception","reason":"Invalid user/password"},"status":401}</code>	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.
<code>":{"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]"}}</code>	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
<code>ConnectTransportException[[host] [ipaddress:9300] connect _timeout[30s]]; nested: ConnectTimeoutException[connection timed out: /ipaddress:9300]; at org.elasticsearch.transport.netty.NettyTransport.connectToChannels (NettyTransport.java:987)</code>	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.
<code>org.elasticsearch.transport.BindTransportException: Failed to resolve host</code>	Elasticsearch network host is set to a hostname or IP address that is not resolvable.	Set it to the IP address of the machine where Elasticsearch is running.

Error	Possible Cause	Resolution
RemoteTransportException- EsRejectedExecutionException		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
Integration Gateway - External System Contact Error 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.	This error might be due to no connectivity to the Elasticsearch server.	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.
Search Instance Ping not working.	This error might be due to the following reasons: <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).

Error	Possible Cause	Resolution
Deploy is throwing errors.	This error might be due to one or more search definitions in deployed state, but unavailable in Elasticsearch.	<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.
Index schedule status shows success, but search results are not displayed.	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>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"> 1. Ensure that some documents are indexed in Elasticsearch. Use the Elasticsearch Interact page to check whether documents are indexed in Elasticsearch. 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. 3. If the security response is empty or does not contain the required security values, then check the security applied for the user.

