



Hewlett Packard
Enterprise

IDOL

Software Version: 11.5

Getting Started Guide

Document Release Date: October 2017

Software Release Date: October 2017

Legal notices

Warranty

The only warranties for Hewlett Packard Enterprise Development LP products and services are set forth in the express warranty statements accompanying such products and services. Nothing herein should be construed as constituting an additional warranty. HPE shall not be liable for technical or editorial errors or omissions contained herein.

The information contained herein is subject to change without notice.

Restricted rights legend

Confidential computer software. Valid license from HPE required for possession, use or copying. Consistent with FAR 12.211 and 12.212, Commercial Computer Software, Computer Software Documentation, and Technical Data for Commercial Items are licensed to the U.S. Government under vendor's standard commercial license.

Copyright notice

© Copyright 2017 Hewlett Packard Enterprise Development LP

Trademark notices

Adobe™ is a trademark of Adobe Systems Incorporated.

Microsoft® and Windows® are U.S. registered trademarks of Microsoft Corporation.

UNIX® is a registered trademark of The Open Group.

This product includes an interface of the 'zlib' general purpose compression library, which is Copyright © 1995-2002 Jean-loup Gailly and Mark Adler.

Documentation updates

The title page of this document contains the following identifying information:

- Software Version number, which indicates the software version.
- Document Release Date, which changes each time the document is updated.
- Software Release Date, which indicates the release date of this version of the software.

To check for recent software updates, go to <https://downloads.autonomy.com/productDownloads.jsp>.

To verify that you are using the most recent edition of a document, go to <https://softwaresupport.hpe.com/group/softwaresupport/search-result?doctype=online help>.

This site requires that you register for an HPE Passport and sign in. To register for an HPE Passport ID, go to <https://hpp12.passport.hpe.com/hppcf/login.do>.

You will also receive updated or new editions if you subscribe to the appropriate product support service. Contact your HPE sales representative for details.

Support

Visit the HPE Software Support Online web site at <https://softwaresupport.hpe.com>.

This web site provides contact information and details about the products, services, and support that HPE Software offers.

HPE Software online support provides customer self-solve capabilities. It provides a fast and efficient way to access interactive technical support tools needed to manage your business. As a valued support customer, you can benefit by using the support web site to:

- Search for knowledge documents of interest
- Submit and track support cases and enhancement requests
- Access product documentation
- Manage support contracts
- Look up HPE support contacts
- Review information about available services
- Enter into discussions with other software customers
- Research and register for software training

Most of the support areas require that you register as an HPE Passport user and sign in. Many also require a support contract.

To register for an HPE Passport ID, go to <https://hpp12.passport.hpe.com/hppcf/login.do>.

To find more information about access levels, go to <https://softwaresupport.hpe.com/web/softwaresupport/access-levels>.

To check for recent software updates, go to <https://downloads.autonomy.com/productDownloads.jsp>.

Contents

Part 1: IDOL Systems	9
Chapter 1: Introduction to IDOL	11
Key IDOL Components	11
IDOL Server	11
OEM Certification	11
Connectors	11
Security Overview	11
Interfaces	12
Distributed Systems	12
Multimedia	13
IDOL Server Operations	13
Agents	13
Alerts	14
Automatic Query Guidance	14
Categorization	14
Category Matching	14
Channels	14
Cluster Information	14
Collaboration	14
Dynamic Clusters	15
Dynamic Thesaurus	15
Education	15
Expertise	15
Hyperlinks	15
Email Users	15
Profiles	16
Search and Retrieval	16
Spell Check	16
Summarization	17
Taxonomy Generation	17
View Documents	17
IDOL System Architecture	17
Actions	18
Index and Query	18
Security	19
Text Queries	19
Agent, Profile and Category Queries	20
Community Queries	20
Related Documentation	20
Chapter 2: Types of IDOL Systems	25
IDOL Core Components	25

IDOL System Setups	27
Unified Setup	28
Component Setup	29
Distributed Setup	30
Mirror Mode	31
Non-Mirror Mode	31
Chain Distribution Servers	32
Distributed Component Setup	32
Other IDOL Components	34
License Server	34
Connectors and Connector Framework Server	35
Multimedia	35
Secured Search	36
Query Manipulation Server	37
Statistics Server	39
Service Control	39
Knowledge Graph	39
Front-End Applications	40
IDOL Installation and Setup	40
Unified Installation	41
Component Installation	41
Component ZIP Packages	41
Other Component Installations	42
Chapter 3: Applications	43
Front-End Applications	43
Find	43
Custom Front-End Applications	44
Administrative Applications	44
IDOL Admin	44
IDOL Site Admin	44
IDOL Data Admin	45
Chapter 4: Security in IDOL	47
Security Overview	47
User Authentication	47
Document Security	48
Mapped Security	48
Unmapped Security	49
Mapped Security Example	50
Secure Communications	51
SSL Communications	51
Part 2: Install and Run IDOL Server	53
Chapter 5: Install IDOL	55
System Requirements	55
Basic Requirements	55

Supported Platforms	55
Recommended Minimum Hardware Specifications	55
Windows System Recommendations	56
TCP Port Requirements	56
Linux System Recommendations	56
Open File Descriptors	56
Transparent Huge Pages	57
Tune Virtual Memory	58
Memory Map Counts	58
Configure the I/O Scheduler	59
Install IDOL	60
Install IDOL on Windows	60
Install a Unified IDOL Server Setup on Windows	62
Windows Services	63
Install an IDOL Component as a Service on Windows	64
Install IDOL on UNIX	65
Install an IDOL Component as a Service on Linux	68
Install a Component as a Service for a systemd Boot System	68
Install a Component as a Service for a System V Boot System	69
Use Unattended Installation	70
Install Separate IDOL Components	71
IDOL Component Dependencies	71
Install IDOL Admin	72
Supported Browsers	72
Access IDOL Admin	72
Licenses	73
Display License Information	74
Revoke a Client License	75
Troubleshoot License Errors	75
Find the Number of Licensed Documents	77
Chapter 6: Run IDOL Server	79
Start and Stop IDOL	79
Start IDOL	79
Start an IDOL Component on Microsoft Windows	79
Start an IDOL Component on UNIX	80
Start IDOL from the Command Line	80
Stop IDOL	82
Send Actions to IDOL	82
Verify IDOL Runs Correctly	83
GetRequestLog	83
GetLicenseInfo	83
GetStatus	84
GetVersion	84
Display Online Help	84
Chapter 7: Configure IDOL Server	87
IDOL Configuration Files	87

Unified Configuration	87
Modify Configuration Parameter Values	87
Include an External Configuration File	88
Include the Whole External Configuration File	89
Include Sections of an External Configuration File	89
Include a Parameter from an External Configuration File	89
Merge a Section from an External Configuration File	90
Chapter 8: IDOL Performance	93
Performance Overview	93
Schedule Index and Query Operations	93
Use a Component Setup	93
Optimize the IDOL Content component	94
Use a Stop Word List	94
Index Numbers	94
Optimize the Index Process	94
Delayed Synchronization	94
Distribute IDOL Content component Data Across Multiple Disks	95
Optimize Query Operations	96
Optimize IDOL Fields	97
Index Fields	97
Match Fields	98
Numeric Fields	98
Numeric Date Fields	98
Count Fields	98
Sort Fields	98
Bit Fields	98
Parametric Fields	98
Field Check Fields	99
Hardware Considerations	99
Optimize Distributed Systems	99
Distributed Action Handler Performance	100
Fast Mirror Mode	100
Simple Combinator Mode	100
Distributed Index Handler Performance	100
Preserve DREADD	100
Distribute On Batch	101
Advanced Distribution Modes	101
Round Robin Mode	101
Chapter 9: Tutorial: Index Data into IDOL	103
Overview	103
Configure the IDOL Content component Index	103
Configure Databases	104
Configure Fields	104
Choose Field Properties	104
Set up Field Processes	106
Configure Languages	108

Define Language Types	108
Associate Language Types with Documents	109
Create Documents	111
Pre-Process Documents	111
Index Documents	111
Use DREADD to Index IDX and XML Files Directly	112
Use DREADDDATA to Index Data Over a Socket	112
Index Data in a Distributed IDOL Setup	113
Chapter 10: Query IDOL	115
IDOL ACI Responses	115
The Query ACI Action	116
Search Parameters	116
Text	116
FieldText	117
Combine	118
DatabaseMatch	118
Start and MaxResults	119
MinScore	119
StoreState, StateMatchID, and StateDontMatchID	119
UserMetaFields	120
Response Parameters	120
Predict and TotalResults	120
Print and Printfields	121
QuerySummary	121
ResponseFormat	121
Sort	122
Highlight, StartTag, and EndTag	122
Summary, Sentences, and Characters	122
Cluster	122
Query Response	123
Glossary	125
Send documentation feedback	129

Part 1: IDOL Systems

This section describes IDOL server and describes the various set ups that you can use in your IDOL system.

- [Introduction to IDOL](#)
- [Types of IDOL Systems](#)
- [Applications](#)
- [Security in IDOL](#)

Chapter 1: Introduction to IDOL

The Intelligent Data Operating Layer (IDOL) server integrates unstructured, semi-structured, and structured information from multiple repositories through an understanding of the content. It delivers a real time environment to automate operations across applications and content, removing all the manual processes involved in getting information to the right people at the right time.

- [Key IDOL Components](#) 11
- [IDOL Server Operations](#) 13
- [IDOL System Architecture](#) 17
- [Related Documentation](#) 20

Key IDOL Components

The following section describes the core components that IDOL systems use.

IDOL Server

The *Intelligent Data Operating Layer* (IDOL) gathers and processes unstructured, semi-structured, and structured information in any format from multiple repositories using IDOL connectors and a global relational index. It can automatically form a contextual understanding of the information in real time, linking disparate data sources together based on the concepts contained within them.

For example, IDOL can automatically link concepts contained in an e-mail message to a recorded phone conversation, which can be associated with a stock trade. This information is then imported into a format that is easily searchable, adding advanced retrieval, collaboration, and personalization to an application that integrates the technology.

OEM Certification

IDOL works in OEM licensed environments.

Connectors

Connectors enable automatic content aggregation from any type of local or remote repository (for example, a database, a Web site, a real-time telephone conversation and so on). Connectors form a unified solution across all information assets within the organization.

Security Overview

HPE provides the software infrastructure that automates operations on unstructured information. This software infrastructure is based on IDOL Server.

Document security protects the information that you index into IDOL Server.

Your organization is likely to store information in many repositories. Many of these repositories have security features that apply permissions to files, so that they only authorized personnel can view them. Document security ensures that when you index information into IDOL Server, IDOL continues to enforce these permissions. In response to a query, IDOL returns only documents that a user is permitted to view.

IDOL includes the following features to protect your data:

- **User authentication.** At the front end, users must log on before they can query IDOL. IDOL can authenticate users against an existing directory service, such as Microsoft Active Directory.
- **Document security.** IDOL Server checks each document that is returned in response to a query. If the user who submitted the query does not have permission to view the document, the document is removed from the query results before the results are returned.
- **Secure communications.** You can encrypt the communications between ACI servers and any applications that use the HPE IDOL ACI API.

Interfaces

- **IDOL Admin** allows you to administer an IDOL component. It provides a user interface for many common operations for various IDOL components. IDOL Admin is installed by the IDOL Server installer.
- **IDOL Site Admin** allows you to administer a wider IDOL installation, with multiple servers and components.
- **IDOL Data Admin** allows you to manage the content of data indexed in IDOL Servers. By creating and modifying different business projects, you can predefine which results are returned to users, and how the user views them. You can also view statistical information to help you refine your business projects and make them more effective.
- **Find** provides a basic end-user search application for IDOL. You can install Find by using the IDOL Server installer.
- **ACI API** uses HTTP to allow custom-built applications to communicate with IDOL ACI servers. Several IDOL SDKs are available to allow you to develop applications with the ACI API.

Related Topics

- [Applications](#)

Distributed Systems

IDOL distribution solutions facilitate linear scaling of systems through faster action execution and reduction of processing time.

- **DAH** (Distributed Action Handler) enables the distribution of ACI (Autonomy Content Infrastructure) actions to multiple IDOL Servers or components, providing failover and load balancing.
- **DIH** (Distributed Index Handler) enables distributed indexing of documents into multiple IDOL Servers or IDOL Content components, providing failover and load balancing.

Multimedia

IDOL includes components that allow you to incorporate information from video, image, and audio sources.

- **Media Server** analyzes video files and streams, to extract information about their content. Media Server can run analysis operations such as face recognition, speech-to-text, and number plate recognition.
- **IDOL Speech Server** processes audio files, performing tasks such as speech-to-text transcription, speaker identification, and audio search.

IDOL Server Operations

IDOL Server can perform the following intelligent operations across structured, semistructured, and unstructured data.

Agents	Expertise
Alerts	Hyperlinks
Automatic Query Guidance	Mailing
Categorization	Profiles
Channels	Retrieval
Cluster Data	Spelling Correction
Collaboration	Summarization
Dynamic Clusters	Taxonomy Generation
Dynamic Thesaurus	Viewing
Education	

NOTE:

Your license determines which of these operations your IDOL Server installation can perform.

Agents

Users can create agents in IDOL Server to find and monitor information that is relevant to their interests. The agents collect this information from a configurable list of Internet and intranet sites, news feeds, chat streams, and internal repositories.

Alerts

IDOL Server analyzes data when it receives new documents, and compares the concepts in documents with user agents. If new data matches a user agent, it immediately notifies the user by email or a third-party system (for example by SMS or a pager).

Automatic Query Guidance

IDOL Server finds the most salient terms and phrases in query results, and automatically clusters these terms and phrases. It uses the clustered phrases to provide a hierarchical set of queries that guide users to the result area that they are looking for.

Categorization

IDOL Server can automatically categorize data. IDOL categorization allows you to derive categories from the concepts found in unstructured text. This process ensures that IDOL Server accurately classifies all data in the correct context. IDOL categorization is a scalable solution capable of handling high volumes of information accurately and consistently.

Category Matching

IDOL Server accepts a category or piece of content and returns categories ranked by conceptual similarity. This ranking determines the most appropriate categories for the piece of content, so that IDOL Server can subsequently tag, route, or file the content accordingly.

Channels

IDOL Server can automatically provide users with a set of hierarchical channels with highly relevant information pertinent to the respective channel. Channels are similar to agents, aggregating information that is relevant to the channel concept. Usually, administrators set up channels that are available to all users.

Cluster Information

IDOL Server automatically clusters information. Clustering takes a large repository of unstructured data, agents, or profiles and automatically partitions the data to cluster similar information together. Each cluster represents a concept area in the knowledge base, and contains a set of items with common properties.

Collaboration

IDOL Server automatically matches users with common explicit interest agents or similar implicit profiles. This information creates virtual expert knowledge groups.

Dynamic Clusters

When it processes queries, IDOL Server automatically clusters the query results, and then clusters the first set of clusters to produce subclusters. This process allows you to generate a hierarchy of clusters that allows users to navigate quickly to their area of interest.

Dynamic Thesaurus

When it processes queries, IDOL Server can automatically suggest alternative queries, allowing users to quickly produce a variety of relevant result sets.

Eduction

Eduction is a tool that you can use to extract an entity (a word, phrase, or block of information) from text, based on a pattern you define. The pattern can be a dictionary of names such as people or places. The pattern can also describe what the sequence of text looks like without listing it explicitly, for example, a telephone number. The entities are contained inside grammar files.

When you use Eduction with IDOL Server, Eduction extracts the entities while the document is indexed and adds them into fields for easy retrieval.

The Eduction capability of IDOL Server is described in the *Eduction User Guide*.

Expertise

IDOL Server accepts a natural language or Boolean search string and returns users who own matching agents or profiles. This process allows instant identification of experts in a subject, eliminating time-consuming searches for specialists, and unnecessary researching of subjects for which expert knowledge is already available.

Hyperlinks

You can automatically generate hyperlinks in real time. These link to contextually similar content, for example to recommend related articles, documents, affinity products or services, or media content that relates to textual content.

IDOL Server automatically inserts these links when it retrieves the document. This process means that new documents can reference older documents, and that archived documents can link to the latest news or material on the subject.

Email Users

IDOL Server matches the agents and profiles against its document content at regular intervals, and automatically notifies users of documents that match their agents or profiles by emailing them.

Profiles

IDOL Server automatically creates interest and expertise profiles for users, in real time.

You can create interest profiles by tracking the content that a user views and extracting a conceptual understanding of it. IDOL Server then uses this understanding to keep user interest profiles up to date. You can use interest profiles to:

- Target information to users.
- Recommend content to users.
- Alert users to the existence of content.
- Put users in touch with other users who have similar interests.

You can create expertise profiles by tracking the content that a user creates and extracting a conceptual understanding of it. IDOL Server uses this understanding to keep user expertise profiles up to date. You can use expertise profiles to trace users who are experts in particular subject areas.

Search and Retrieval

IDOL Server offers a range of retrieval methods, from simple legacy keyword search to sophisticated conceptual querying.

Examples of the kind of query that is possible with IDOL server include:

- Conceptual Matching
- Advanced Keyword Search
- Boolean/Bracketed Boolean Search
- Exact Phrase Search
- Field Restrictions
- Field Text Search
- Fuzzy Search
- Parametric Search
- Proper Names Search
- Proximity Search
- Soundex Keyword Search
- Synonym Search

Spell Check

IDOL Server can automatically spell check the query text it receives and suggest correct spelling for terms that its dictionary does not contain.

Summarization

IDOL Server accepts a piece of content and returns a summary of the information. IDOL Server can generate different types of summary.

- **Conceptual Summaries.** Conceptual summaries contain the most salient concepts of the content.
- **Contextual Summaries.** Contextual summaries relate to the context of the original query. They provide the most applicable dynamic summary in the results of a particular query.
- **Quick Summaries.** Quick summaries include a few sentences of the result documents.

Taxonomy Generation

Automatic taxonomy generation can automatically understand and create deep hierarchical contextual taxonomies of information. You can use clustering, or any other conceptual operation, as a seed for the process.

The resulting taxonomy can:


- Provide insight into specific areas of the information.
- Provide an overall information landscape.
- Act as training material for automatic categorization, which then places information into a formally dictated and controlled category hierarchy.

View Documents

IDOL Server uses IDOL KeyView filters to convert documents into HTML format for viewing in a Web browser.

IDOL System Architecture

IDOL server uses the ACI (Autonomy Content Infrastructure) Client API to communicate with custom-built applications that retrieve data using HTTP requests. It implements this communication over HTTP using XML and can adhere to SOAP.



Actions

When communicating with IDOL Server, there are two main types of action:

- ACI (Autonomy Content Infrastructure) actions
- Index actions

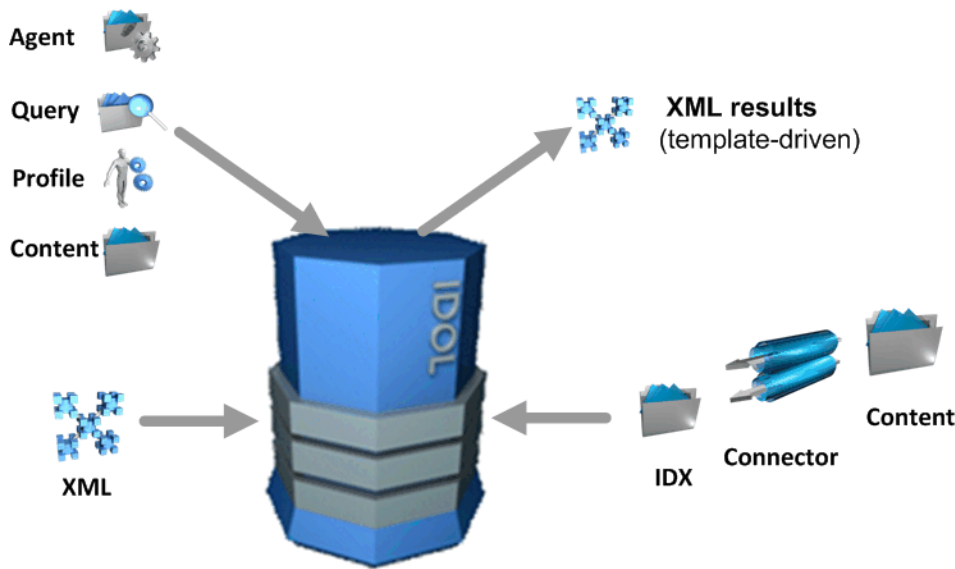
ACI actions request information or perform operations and return results. Different IDOL components accept different ACI actions. For example, the `Query` action requests results from the IDOL Content component, while the `UserRead` action requests user information from the IDOL Community component.

Index actions maintain the IDOL Server data index. You use them to index data, delete data, and perform operations on the data in IDOL Server.

Index and Query

You index documents into IDOL in IDOL IDX format or in XML format (directly or using a Connector). IDOL stores the concepts of the document. In response to queries, agents, profiles or content, it returns a link to the result document. IDOL also returns a percentage weighting, which indicates how relevant the result document is to the original query.

IDOL can return results as XML (even if the document was not in XML format when it was indexed) or other formats, such as plain text, using XSLT:



Security

It is often necessary to ensure that users can access or retrieve only data that they are authorized to view. IDOL enables you to set permissions for users or groups of users, to protect information, and ensure only the correct people can access it.


For more information on security, refer to the *IDOL Document Security Administration Guide*.

Text Queries

IDOL contains data that has been aggregated from one or more repositories. In this example each of the repositories has its own group server that stores the repository user names and the groups that these users belong to. IDOL aggregates this security information from the group servers.

When users log onto a client, their authentication details are sent to IDOL. IDOL returns the user security details to the client, which stores them until the client logs off or the session times out. Every time users send a text query from a client, the client attaches their security details to the query string that it sends to IDOL.

IDOL uses the security information in the query string to check the user permissions. It matches the security string against the document access control lists (ACLs). IDOL returns documents that match the query and that the user has permission to see.



Agent, Profile and Category Queries

IDOL contains data aggregated from one or more repositories. In this example, each of the repositories has its own group server that stores the repository user names and the groups that these users belong to. IDOL aggregates this security information from the group servers.

The client sends an agent, profile, or category query to IDOL. IDOL (which stores all agents and profiles) matches this agent, profile or category against the documents it contains.


IDOL uses the information that it receives from the group servers to check the user permissions. It returns documents that match the query and that the user has permission to see.



Community Queries

IDOL stores user agents and profiles, to match them against community queries (that is, any type of query that returns agents or profiles).

When a client sends a community query to IDOL, IDOL matches it against the agents and profiles it stores. It returns matching agents, profiles, or both to the client.



Related Documentation

The following documents provide more details on IDOL.

Document	Description
<i>IDOL Server Release Notes</i>	Provides information about the new features and resolved issues in a release of IDOL Server.
<i>IDOL Expert</i>	Provides conceptual overviews and expert knowledge of IDOL and its features and functionality.
<i>IDOL Server</i>	Describes how to configure and use IDOL server.

<i>Administration Guide</i>	
<i>IDOL Server Reference</i>	Provides information about all the actions and configuration parameters that you can use in IDOL.
<i>Content Component Reference Category Component Reference Community Component Reference IDOL Proxy Component Reference View Server Reference</i>	Provides information about the actions and configuration parameters that you can use in each IDOL subcomponent.
<i>IDOL Admin User Guide</i>	Describes how to set up and use the IDOL Admin interface to administer the IDOL Content component.
<i>IDOL Site Admin Installation Guide</i>	Describes how to install set up IDOL Site Admin for monitoring and maintaining your IDOL installations.
<i>IDOL Site Admin User Guide</i>	Describes how to use IDOL Site Admin to monitor and maintain your IDOL installations.
<i>Distributed Action Handler Administration Guide</i>	Describes how to configure and use the Distributed Action Handler (DAH) with IDOL server.
<i>Distributed Action Handler Reference</i>	Provides information about all the actions and configuration parameters that you can use in the DAH.
<i>Distributed Index Handler Administration Guide</i>	Describes how to configure and use the Distributed Index Handler (DIH) with IDOL server.
<i>Distributed Index Handler Reference</i>	Provides information about all the actions and configuration parameters that you can use in the DIH.
<i>Query Manipulation Server Administration Guide</i>	Describes how to configure and use the Query Manipulation Server (QMS) to modify queries and results in IDOL server.
<i>Query Manipulation</i>	Provides information about all the actions and configuration parameters that

<i>Server Reference</i>	you can use in QMS.
<i>License Server Administration Guide</i>	Describes how to configure and use the License Server to manage IDOL services and licences.
<i>License Server Reference</i>	Provides information about all the actions and configuration parameters that you can use in the License Server.
<i>Controller Reference</i>	Provides information about all the actions and configuration parameters that you can use in the Controller component, which is used by IDOL Site Admin to monitor other IDOL components on a host.
<i>Coordinator Reference</i>	Provides information about all the actions and configuration parameters that you can use in the Coordinator component, which is used by IDOL Site Admin to monitor other hosts in an IDOL system.
<i>IDOL Document Security Administration Guide</i>	Describes how to set up and use security for IDOL server. This guide has been renamed from <i>Intellectual Asset Protection System (IAS) Administration Guide</i> .
<i>OmniGroupServer Reference</i>	Provides information about all the actions and configuration parameters that you can use in the OmniGroupServer.
<i>IDOL Education User Guide</i>	Describes how to set up and use education with IDOL server to extract entities such as names, statistics, and locations from unstructured content.
<i>Connector Framework Server Administration Guide</i>	Describes how to configure and use Connector Framework Server (CFS) to convert data from connectors into index files, process content, and index data into IDOL server.
<i>Connector Framework Server Reference</i>	Provides information about all the actions and configuration parameters that you can use in the Connector Framework Server.
<i>File System Connector Administration Guide</i> <i>Web Connector Administration Guide</i> <i>HTTP Connector Administration Guide</i> Other connector guides, as needed.	Describes how to configure and use IDOL connectors to aggregate data from different sources to index into IDOL server.
<i>File System Connector</i>	Provides information about all the actions and configuration parameters that you can use in the connectors.

<p><i>Reference</i> <i>Web Connector</i> <i>Administration</i> <i>Guide</i> <i>HTTP Connector</i> <i>Reference</i> Other connector references, as needed.</p>	
<p><i>IDOL Speech</i> <i>Server</i> <i>Administration</i> <i>Guide</i></p>	<p>Describes how to configure and use IDOL Speech server to process audio data.</p>
<p><i>IDOL Speech</i> <i>Server Reference</i></p>	<p>Provides information about all the actions and configuration parameters that you can use in IDOL Speech Server.</p>
<p><i>Knowledge Graph</i> <i>Technical Note</i></p>	<p>Describes how to configure and use the IDOL Knowledge Graph component to create and use a graph.</p>
<p><i>Knowledge Graph</i> <i>Component</i> <i>Reference</i></p>	<p>Provides information about all the actions and configuration parameters that you can use in the Knowledge Graph component.</p>
<p><i>Find Administration</i> <i>Guide</i></p>	<p>Provides information about how to install and set up the Find user interface.</p>
<p><i>IDOL Data Admin</i> <i>Installation Guide</i></p>	<p>Describes how to install and set up the IDOL Data Admin to manage the content of data indexed in IDOL servers.</p>
<p><i>IDOL Data Admin</i> <i>User Guide</i></p>	<p>Describes how to use IDOL Data Admin to manage the content of data indexed in IDOL servers.</p>

Chapter 2: Types of IDOL Systems

You can set up IDOL systems in different ways to suit the requirements of your organization. This chapter discusses some of the possible IDOL setups, and some of the considerations involved when you choose a setup.






- [IDOL Core Components](#)25
- [IDOL System Setups](#)27
- [IDOL Installation and Setup](#)40

IDOL Core Components

This section describes the common IDOL components.

The following components form the core IDOL text indexing and processing functions. Together, they are often collectively known as *IDOL Server*. For more information, refer to the *IDOL Server Administration Guide*.

IDOL Text Processing







Service	Description
 Content	Indexes, collects, manipulates and stores unstructured and semi-structured text data.
 AgentStore	Indexes, collects, manipulates and stores agent and category information. Agentstore is a special configuration of the Content component.
 Category	Stores predefined or customized categories in which data is organized.
 Community	Stores information about users, roles, and permissions.
 View	Converts documents to HTML format for viewing in a Web browser.

The following components form part of a wider IDOL infrastructure. They allow you to license, distribute, monitor, and control your IDOL components.

Infrastructure



Service	Description
License	Distributes and manages licenses for IDOL products and components.

Infrastructure, continued

Service	Description
Server 	Refer to the <i>License Server Administration Guide</i> for more information.
IDOL Proxy 	Routes actions to the correct components in a unified IDOL Server, and starts, restarts and stops IDOL components.
DAH 	Distributed Action Handler. Distributes action requests across IDOL Servers or IDOL components. Refer to the <i>DAH Administration Guide</i> for more information.
DIH 	Distributed Index Handler. Distributes indexing requests across IDOL Servers or IDOL Content components. Refer to the <i>DIH Administration Guide</i> for more information.
Controller 	Monitors the IDOL services on a host machine, and communicates with Coordinator to report status information. This component is primarily used by IDOL Site Admin. Refer to the <i>IDOL Site Admin User Guide</i> for more information.
Coordinator 	Manages status information from Controllers, and acts as a central point for viewing logs and monitoring the status of the IDOL system. This component is primarily used by IDOL Site Admin. Refer to the <i>IDOL Site Admin User Guide</i> for more information.







The following components allow you to automatically retrieve content and index into IDOL Server.

Connectors

Service	Description
CFS 	Connector Framework Server. Aggregates data from IDOL connectors and indexes the data into IDOL Server. CFS can also manipulate the index content (for example, to additional fields relevant to the content), and index into other systems, such as Knowledge Graph and Vertica databases. Refer to the <i>Connector Framework Server Administration Guide</i> for more information.
Connectors 	Gather data from different sources for indexing into IDOL Server. The IDOL Server installer includes the File System Connector and the Web Connector (on Microsoft Windows and Linux platforms) or HTTP Connector (on Solaris platforms). The File System Connector gathers data from system files, while the Web or HTTP Connector gathers data from Web servers (internet or intranet). Refer to the Connector Administration guides for more details.

The following components perform additional operations and analysis on unstructured data. You can use many of these components with the IDOL Server text processing, or as part of an independent system.

Additional Services

Service	Description
 <p>QMS</p>	<p>Query Manipulation Server. Manages promotions, modifies queries to IDOL Server, and manipulates results from IDOL Server.</p> <p>Refer to the <i>Query Manipulation Server Administration Guide</i> for more information.</p>
 <p>Statistics Server</p>	<p>Accumulates events from client applications or from a script that reads IDOL log files, then uses that data to report statistics.</p> <p>Refer to the <i>Query Manipulation Server Administration Guide</i> for more information.</p>
 <p>Media Server</p>	<p>Analyzes video files and streams, to extract information about their content. Media Server can run analysis operations such as face recognition, speech-to-text, and number plate recognition.</p> <p>Refer to the <i>Media Server Administration Guide</i> for more information.</p>
 <p>IDOL Speech Server</p>	<p>Processes audio files, performing tasks such as speech-to-text transcription, speaker identification, and audio search.</p> <p>Refer to the <i>IDOL Speech Server Administration Guide</i> for more information.</p>
 <p>Knowledge Graph</p>	<p>Uses connections between IDOL documents to create a set of relationships called a <i>graph</i>, and allows you to explore these connections in your data.</p> <p>Refer to the <i>Knowledge Graph Technical Note</i> for more information.</p>
 <p>OGS</p>	<p>Omni Group Server. Collects security information from security repositories and uses it to ensure that users can access only documents that they have permission to access.</p> <p>Refer to the <i>IDOL Document Security Administration Guide</i> for more information.</p>

IDOL System Setups

The flexibility of IDOL allows you to create systems that suit your environment and the unique needs of your organization. The following sections describe some of the ways that you can set up your IDOL system.

- [Unified Setup](#)28
- [Component Setup](#)29
- [Distributed Setup](#)30
 - [Mirror Mode](#)31
 - [Non-Mirror Mode](#)31

- Chain Distribution Servers 32
- Distributed Component Setup 32
- Other IDOL Components 34
 - License Server 34
 - Connectors and Connector Framework Server 35
 - Multimedia 35
 - Secured Search 36
 - Query Manipulation Server 37
 - Statistics Server 39
 - Service Control 39
 - Knowledge Graph 39
- Front-End Applications 40

The installation procedures to support these setups are described in [IDOL Installation and Setup](#), on page 40.


Unified Setup

A unified setup is the most basic setup, for training and test environments. It uses an integrated IDOL Server, rather than separate IDOL components. The integrated IDOL Server contains the core IDOL text indexing and processing functionality.

The unified IDOL Server includes the Content, Category, Community, AgentStore, View, and IDOL Proxy components. It does not include Query Manipulation Server (QMS), or the multimedia components, such as Media Server.

By default, it does not include the distribution components, but you can add these to a unified configuration if required.

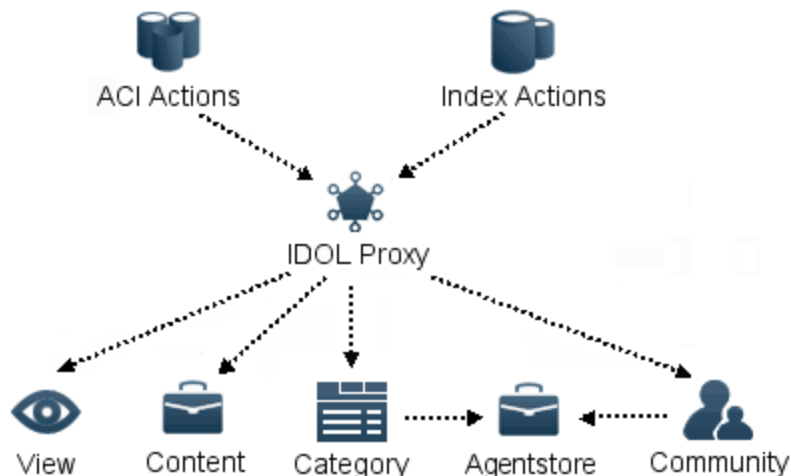
In a unified setup, you configure all the operations for the IDOL Server components using a single configuration file (except for AgentStore, which has a separate configuration file). You send all actions and index actions to a single host and port, and IDOL Server stores and processes the data and actions centrally.



The unified setup uses the IDOL Proxy component. IDOL Proxy acts as a single point of contact for all your IDOL requests. It forwards all requests to the appropriate component. For example, it sends a

Query action or indexing request to the Content component, and it sends a UserRead action to the Community component. IDOL Proxy also starts and stops the components in IDOL Server, and restarts components if they become inactive for any reason.

The following diagram shows the unified setup architecture in more detail:



This setup is useful for training and testing, because it allows you to easily modify the configuration for the server, and you do not have to understand the whole IDOL architecture. However, for most production environments, the unified IDOL setup is too small and restrictive. For example, it has limited scalability and no failover, it is difficult to add additional Content components to increase the size of the index, and you cannot include QMS.

When you want to set up a full IDOL system, you usually move to a component-based setup, where you configure and optimize components separately, and often on different hardware. For more information, see [Component Setup, below](#).

Related Topics

- [Unified Installation, on page 41](#)

Component Setup

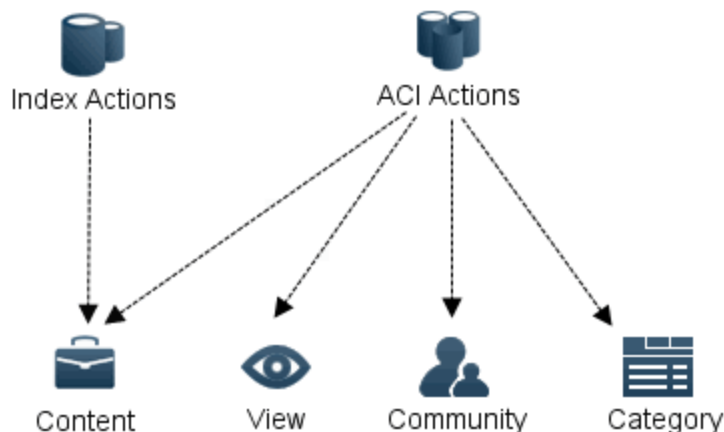
A component setup uses a combination of IDOL components, such as Content, Community, View, and so on, rather than the unified IDOL Server. This more advanced setup is for administrators who are already familiar with IDOL architecture.

HPE recommends that you use a component-based setup in most production environments.

The component setup is more flexible than a unified IDOL Server. It allows you to:

- use only the IDOL components that you need.
- configure each component separately, which can be useful for optimization and tuning, as well as troubleshooting.
- set the components up on separate hardware or with dedicated resources to enhance the performance for different operations, and to allow you to scale resources for different components in a more flexible way.

- simplify component management and maintenance. In particular, it is easier to start, stop, and reinitialize individual components.
- design highly scalable, fault tolerant IDOL systems.



In a component setup, you configure each component with its own configuration file. For example, you configure the `content.exe` in the `content.cfg` file, you configure the `category.exe` in the `category.cfg` and so on.

Each component installation must also include any extra files or modules that you require. These additional files are included as part of the installation for the component if you use the IDOL Server installer, or you can use the zip package downloads, which also include the required files and modules.

NOTE:

There are also other dependencies between components, which you must configure in the component configuration files. For example, the Community and Category configuration files must contain the host and port details for the Agentstore component.

In general, HPE does not recommend that you use the IDOL Proxy component in a component setup. However, if required, you can use a stand-alone IDOL Proxy to forward requests to the appropriate component. In this case, IDOL Proxy can only forward actions; it cannot perform any of the component maintenance (such as starting and stopping components).

If you use a stand-alone IDOL Proxy, the IDOL Proxy configuration file must contain the host IP address and ACI port for each of the IDOL components that you want it to forward requests to. It dynamically configures other ports (such as the index port) when the components start up.

Related Topics

- [Distributed Setup, below](#)
- [Component Installation, on page 41](#)
- [Component ZIP Packages, on page 41](#)

Distributed Setup

A distributed setup involves using a Distributed Index Handler (DIH) and Distributed Action Handler (DAH) to route actions to multiple instances of an IDOL component.

This kind of setup is effective for load-balancing among components, as well as for expanding indexes that no longer fit on one machine. DIH and DAH balance indexing and action requests among the appropriate components. You can either set up the distributed system in mirror mode or non-mirror mode.

- In mirror mode, each instance of the component that you distribute to is identical. You can use this option for load-balancing or failure tolerance.
- In non-mirror mode, each instance of the component is different. This option usually applies only to the IDOL Content component, to allow you to expand the size of your IDOL index.

The DIH distributes index actions (for example, to add or remove documents from your IDOL index). DIH can only distribute to components that have an index port.

The DAH distributes ACI actions. In mirror mode, it can distribute any action to any component. In non-mirror mode, it can distribute most IDOL Content component actions and combine the results of queries from multiple components.

The following sections provide more information about mirror mode and non-mirror mode for the DIH and DAH when distributing to multiple instances of the IDOL Content component.

For more information, refer to the *Distributed Index Handler Administration Guide*, and the *Distributed Action Handler Administration Guide*.

Mirror Mode

In a mirrored setup, you store the same set of data in each instance of the IDOL Content component. Each Content is identical to the others, and you must configure them in the same way.

Run the DIH in mirror mode to ensure uninterrupted service if one of the IDOL Content components becomes inoperable. While one Content is unavailable, the DIH continues to index data into its identical copies, which are also still available to return data for queries. DIH queues the actions for the inoperable Content, and sends them when the Content becomes available again, so that the servers do not become inconsistent.

DIH sends all index actions to all connected IDOL Content components.

In mirror mode, you can configure the DAH to distribute ACI actions in one of two ways:

- **Load Balancing.** The DAH assigns each incoming action to just one of the connected IDOL Content components (using a cumulative predictive algorithm that spreads the action load efficiently).
- **Failover.** The DAH forwards incoming actions to the first Content that you list in the DAH configuration file. If this IDOL Content component stops responding for any reason, the DAH marks it as down and switches to the next Content.

Non-Mirror Mode


In a non-mirrored system, you distribute the data equally among the Content components.

Run the DIH in non-mirror mode if the amount of data to index is too large for a single Content. You can also separate the resources for each part of the index by setting the Content components up on different machines. This approach can improve the indexing time.

In non-mirror mode, the DAH sends ACI actions to all connected IDOL Servers. You can configure the DAH to combine the results in different ways when it returns them.

Chain Distribution Servers

You can set up multiple DIH and DAH instances in a chained configuration. For example, a parent DIH or DAH distributes actions to child DIH or DAH servers, which in turn distributes to child Content components.



In this configuration, the parent DIH and DAH distributes actions to child DIH and DAH servers in the same way as it distributes to child Content components. Each child DIH or DAH accepts all Content actions and forwards them.

Some actions may have a different effect when you send them to a child DIH or DAH server rather than a Content component, because the actions goes to multiple Content components.

Chaining provides an extra level of redundancy both at the DIH or DAH, and the Content level. It also distributes network traffic and system load over a larger number of computers. A chained configuration provides a pool of IDOL Servers that are both fault-tolerant for maximum availability and distributed for the best performance.


For more information about chaining distribution servers, and the architectural considerations, refer to *IDOL Expert*.

Distributed Component Setup

A distributed system with stand-alone components uses any combination of IDOL components with the DIH and DAH. You configure the IDOL components, DIH, and DAH separately with individual component configuration files.

This setup is highly flexible. You can distribute as many or as few of the IDOL components as you need to, and you can scale each component with additional instances as required by your usage. For example, if your system has a high load for categorization, you can increase the number of Category components without adding any other components.

The following diagram shows one example scenario. In this case, there are two Content components, two Community components, and two Category components.




In this example:

- You send index actions to the DIH, which distributes them between the two Content components.
- You route ACI actions between the three DAHs (for example, by using a front-end application, or an IDOL Proxy).
 - You send actions for Content (such as `Query`) to Content DAH, which distributes actions between the two Content components.
 - You send actions for Community (such as `UserRead`) to Community DAH, which distributes actions between the two Community components.
 - You send actions for Category (such as `CategoryQuery`) to Category DAH, which distributes actions between the two Category components.

DAH cannot distribute all ACI actions in non-mirror mode, so the Community and Category components in this example must be mirrored.

Agentstore is not shown in this example, but you could use a single Agentstore for all the Category and Community components, or you can have Category and Community connect to an Agentstore DIH and Agentstore DAH. Agentstore indexes are not usually large, so you might not need to split the index into multiple Agentstore components, but you can create mirrored Agentstores for redundancy.

If you require only the basic indexing and retrieval functionality of IDOL Server, a more minimalist setup may be suitable. In this case, you can set up a number of Content servers, with DIH and DAH servers distributing index and ACI actions between them.:



Depending on the size or requirements of the system, you can either use mirror mode for fault tolerance, or non-mirror mode for performance. You can also use chained DIH and DAH servers to allow for both.

Related Topics

- [Component Setup](#)
- [Component Installation, on page 41](#)

Other IDOL Components

The previous sections have mainly considered the core IDOL Server text processing components, and distribution. IDOL also includes a large number of other components that allow you to analyze and process unstructured data, which you can use in combination with the text processing, or independently. This section describes the main components, and how you might want to include them in your IDOL system.

License Server

License Server is an essential part of any IDOL installation. It manages the licenses for your IDOL products and components.

When you purchase IDOL, HPE Big Data support provide a license key. The license key is bound to a particular MAC address and port, which determines where you must install your License Server. The machine where you install License Server must be accessible over the network to the machines that contain the IDOL services that you want to license.

The only limit to the number of services that a single License Server can manage is your license capacity. In many cases you need only one License Server for your whole setup. However, in some cases you might want to use multiple license servers.

You might want to have multiple License Servers to license separate environments. For example, if you have a production and a test environment, these might have different License Servers. Alternatively, you might want to have a system where each machine that hosts IDOL services has its own License Server.

In these cases, each License Server must have its own license key, with a fixed MAC address and port.

NOTE:

If you want to change or update your system, for example to include new components or features that change your license key, you must update every affected License Server. Similarly, if you need to change the host or port of a License Server, you must request a new license key from HPE Big Data support.

License Server is available in the IDOL Server installer. For more information, see [Install IDOL, on page 60](#), and the *License Server Administration Guide*.

Connectors and Connector Framework Server

IDOL Connectors and the Connector Framework Server (CFS) allow you to collect data from different repositories to import into IDOL.

- Connectors access a particular repository, and extract the content to send to CFS. Each connector is specific to a particular repository, rather than a particular file type. For example, you must use File System Connector to access a file system, and Web Connector to access files from the Internet, but both connectors can extract any available file content.
- Connector Framework Server processes the files from your connectors. CFS uses IDOL KeyView and the IDOL multimedia components to extract text from files, which it then sends to your IDOL index (the Content or DIH component). CFS can also perform data enrichment before you index to make it easier to retrieve information.

The number of connectors and CFS components that you need depends on the number of repositories you want to collect data from, and the amount of processing that you want to do.

CFS can process data from multiple connectors. For example, if you want to index content from a few small repositories, and the repositories do not change very often, you might set all the connectors up to send content to the same CFS, which indexes into IDOL Server.

In other cases, you might want to set up multiple CFS instances to process data from different connectors. If your repositories change regularly (generating a large number of new files and updates), or if they are very large, you can use one CFS for each connector.

The CFS is available in the IDOL Server installer. For Linux and Microsoft Windows platforms, the IDOL Server installer also includes File System Connector and Web Connector. For Solaris platforms, the installer includes File System Connector and HTTP Connector. Other connectors are available as separate installers.

For more information, see [Install IDOL, on page 60](#), and refer to the *Connector Framework Server Administration Guide*, as well as the administration guides for the individual connectors.

Multimedia

The IDOL multimedia components allow you to process and analyze images, video, and audio files.

- **Media Server** analyzes video files and streams, to extract information about their content. Media Server can run analysis operations such as face recognition, speech-to-text, and number plate recognition.

- **IDOL Speech Server** processes audio files, performing tasks such as speech-to-text transcription, speaker identification, and audio search.

You can use these components in combination with the IDOL text processing components. For example, you might want to use Speech Server with CFS to transcribe audio files that you extract from a repository, and then index the text output into the Content component. You can also use multimedia components independently of other IDOL components. For example, you might use Media Server for face recognition from video files.

Each of these components has its own configuration and setup.

When you use a multimedia component with CFS, you must configure CFS with details of the component to use for processing. If necessary, you can use these components with a DAH for load-balancing and failover.

These components are all available in the IDOL Server installer. For more information, see [Install IDOL, on page 60](#). For more information about the components, refer to the *Media Server Administration Guide*, and the *IDOL Speech Server Administration Guide*.

Secured Search

To add secured search for IDOL Server, you might need to install Omni Group Server (OGS). OGS manages security for documents that you index in the Content component, and ensures that users can access only documents that they have permission to see. OGS can collect security information from several different security repositories and store this information, updating it at regular intervals.

When a user starts a session with the IDOL Community component, the client application requests security information. Community retrieves the user security details from OGS. The client application adds the security information to every subsequent query that it sends to the IDOL Content component. Content can then compare the user security details to the document access control lists (ACLs) to determine what type of access to grant.



Omni Group Server is available as a separate installer. For more information, see [Security in IDOL, on page 47](#), and refer to the *IDOL Document Security Administration Guide*.

Setup for secured search requires configuration in the Content and Community components, and your connectors. For more information, refer to *IDOL Expert*.

Related Topics

- [Mapped Security Example](#)

Query Manipulation Server


Query Manipulation Server (QMS) is an additional component that you can use with a Content index to modify user queries and results. For example, you can use QMS to remove certain terms from the query text, set up promotions to return when a user searches for a particular item, or to return a promoted item to a particular position in the search results.

QMS communicates with a Promotion Agentstore, which stores the rules that QMS uses to determine how to modify queries.

In an IDOL architecture, you send the Content actions that you want to modify to QMS, rather than directly to Content. QMS performs any additional processing, and then forwards the actions to Content.


When you have configured distribution with your Content components, you must decide how to include QMS in the distributed architecture.

- You can use one QMS component for each Content, and direct the actions to a DAH that distributes to the QMS components.



In a mirror mode configuration with multiple identical Content components, this architecture ensures that QMS is part of the load-balancing or failover that you use for your system.

- You can use one QMS component, and then send the Content actions to the DAH after the QMS processing.



In a non-mirror mode configuration where your index is split over multiple Content components, this architecture reduces the number of additional components you need to configure. In a mirror mode configuration, you can use this architecture if you do not need to include load-balancing or failover for the QMS part of the system.

You can use QMS alone or as part of IDOL Data Admin. IDOL Data Admin provides a front-end application for manipulating and monitoring your queries. For more information, see [IDOL Data Admin, on page 45](#).

QMS is available as part of the IDOL Server installer. For more information, see [Install IDOL, on page 60](#), and refer to the *Query Manipulation Server Administration Guide*.

Statistics Server

Statistics Server monitors IDOL logs and collects information on different events.

You can use Statistics Server on its own, but typically you install it as part of an IDOL Data Admin installation. IDOL Data Admin uses Statistics Server to collect information about the queries and terms that you use in your IDOL installation. For more information about IDOL Data Admin, see [IDOL Data Admin, on page 45](#).

Statistics Server is available as part of the IDOL Server installer. For more information, see [Install IDOL, on page 60](#), and refer to the *Query Manipulation Server Administration Guide*.

Service Control

The Controller and Coordinator components form part of a wider monitoring architecture. Usually, you install these components as part of an IDOL Site Admin installation to monitor and control IDOL services. See [IDOL Site Admin, on page 44](#).

- Controller monitors the IDOL services on a single host. In the typical case where you have your IDOL system distributed across multiple machines, you include a Controller server on each machine. Each Controller reports information to the central Coordinator server.
- Coordinator collects the monitoring information from your Controller servers. You have one Coordinator for the IDOL system that you want to monitor. You can then use the IDOL Site Admin application.

Controller and Coordinator are available as part of the IDOL Server installer. For more information, see [Install IDOL, on page 60](#).

NOTE:

The Controller component is selected by default in the installer, so that it is automatically deployed when you install components on a new host machine.

For information on how to install IDOL Site Admin, refer to the *IDOL Site Admin Installation Guide*. For information about how to use Controller and Coordinator as part of IDOL Site Admin, refer to the *IDOL Site Admin User Guide*.

Knowledge Graph

Knowledge Graph can index content to create a graph of the relationships between concepts and entities in your data. You can use Knowledge Graph to provide additional tools for manipulating your unstructured data, and to model connections in a different way to the normal IDOL text processing methods.

You can configure CFS to index data into Knowledge Graph, so that you can create a graph from a set of data. Other than this, Knowledge Graph does not interact directly with other IDOL components in your architecture.

Front-End Applications

When you have multiple front-end applications, it can be beneficial to use separate IDOL installations for each application.

For example, you might have one application for collaboration in your organization, and another for data mining. You can set up each application to use a separate IDOL system, with any required components, distribution, and security.

Separating the applications in this way has several advantages:

- It is straightforward to use different data sets for each application, if required. Each different installation can have its own set of connectors, indexing data from the required sources. When the applications have different uses, this approach allows you to tailor the content for these different uses.
- It can be easier to manage security. For example you can separate an application for use on an intranet from another application that is accessible over the Internet.
- You can configure each IDOL system individually, to optimize performance and functionality for its associated front-end application. In particular, if the different applications have different usage and availability requirements, you can manage these requirements more easily in a separated system.
- You can optimize and scale resources more efficiently.
- It is straightforward to update your application or IDOL installation without causing interruptions to other applications.
- It is easier to test the effects of new configurations.

IDOL Installation and Setup

After you have decided on an IDOL architecture, you must decide on how to set up and distribute your system across multiple machines.

For the simple unified IDOL Server installation for testing, you can usually use a single machine for all components. However, in a more advanced system with a component-based setup, you usually separate some of the components on to separate disks or hardware.

In systems where you have large processing requirements for certain tasks, you can separate the IDOL components onto their own servers. This separation allows you to isolate and scale resources independently for different functions, according to their individual requirements. You can also place the components on their own hardware or virtual machine.

For more information about sizing and scaling your IDOL system, refer to *IDOL Expert*.

For full details about installing IDOL Server and the components, see [Install IDOL, on page 60](#).

Unified Installation

The IDOL Server installer **Unified** installation option allows you to install a simple unified IDOL Server installation. This option has a very simple installation procedure, and installs:

- the main IDOL Server components (Content, Category, Community, Agentstore, and View)
- the IDOL Proxy component
- the unified IDOL Server configuration file, which contains standard configuration options for the main components
- the Agentstore configuration file, which contains standard configuration options for the IDOL Agentstore component
- all standard modules and libraries for the main IDOL Server components (such as language files, security modules, and KeyView filters)
- the IDOL Server `help.dat` file, which contains the *IDOL Server Reference*, *IDOL Server Administration Guide*, and *IDOL Expert*
- the `admin.dat` file that contains the IDOL Admin interface
- License Server, and the License Server configuration file and `help.dat` file

Component Installation

The IDOL Server installer **Custom** installation option allows you to individually select and install components, or sets of components. You can run the installer on each machine where you want to install a selection of IDOL components, and choose the appropriate components.

For each component, the installer provides:

- the component executable file
- a configuration file, which contains standard configuration options for the component
- all standard modules and libraries for the component
- the documentation for the component (in the `help.dat` file, which you can access by using the `Help` action)
- the IDOL Admin interface (in the `admin.dat` file, which you can access by using the `Admin` action)

When you run the installation process, you must select the components and any sub-components that you want to install.

By default, the Controller component is selected. For IDOL Site Admin monitoring, you must include a Controller on every host machine where you install IDOL services. Clear the Controller installation option only if you have already installed Controller on a particular host.

Component ZIP Packages

IDOL components are also available to install as separate ZIP packages.

Each component ZIP package provides a configuration file, and the required modules for the component. It also includes a component-specific Reference, which contains the actions and configuration parameters relevant to that component.

To install the component, you extract the ZIP package to the host where you want to run the component. You must also update the [License] section of the configuration file to include appropriate license information. If you have more than one version of a component running on the same host, you might also need to update the port information to ensure that each component uses a different port.

After you install each of the components, you must modify the configuration files to ensure that the components can contact any other necessary components. For example, you must configure:

- the IDOL Category component with the details of the data index (IDOL Content component), and the category index (IDOL Agentstore component).
- the IDOL Community component with the details of the agent index (IDOL Agentstore component).
- the IDOL View component with the details of an IDOL Content component that it can use for highlighting.

Other Component Installations

The IDOL Server installer includes most of the common IDOL components and options. However, most IDOL connectors, some of the front-end applications, and Omni Group Server, are available in separate installers.

For more information about these components, refer to the documentation for the component.

Chapter 3: Applications

There are many applications that you can use to access, search, and customize the data in your IDOL system. This section describes some of the available applications that you can use to access and administer your IDOL system.

- [Front-End Applications](#)43
- [Administrative Applications](#)44

Front-End Applications

There are a variety of applications that provide a user interface through which you can access, customize, search, and modify data in your IDOL system.

Find and BIFHI are provided as part of the IDOL Server installer. Find provides a basic end-user search interface for IDOL, and BIFHI provides the search interface with some additional functionality. There are also several IDOL applications that use IDOL for different business purposes, such as archiving, data mining, and eDiscovery. For more information, contact your Account Manager.

In addition, you can create custom front-end applications to use your IDOL data.

Find

Find is a basic end-user search interface for IDOL. Find also includes Business Information for Human Intelligence (BIFHI).

Find supports the following functionality:

- **Advanced search.** Refine searches by database, date, or parametric values.
- **Document preview.** Find uses the IDOL View component to render near-native views of the source documents for your search results in a web browser. The Document Preview feature works with View server and documents in source repositories.
- **Query manipulation.** Find applies synonym and blacklist rules when you use the IDOL Query Manipulation Service as a back end for your search. You can use IDOL Data Admin to create and manage these rules.
- **Results manipulation.** Find applies pin-to-position promotions when you use the IDOL Query Manipulation Service as a back end for your search. You can use IDOL Data Admin to create and manage these promotions.
- **Results augmentation.** Find applies spotlight promotions when you use the IDOL Query Manipulation Service as a back end for your search. You can use IDOL Data Admin to create and manage these promotions.
- **Document security.** You can set up document security so that only users with appropriate permissions can access the documents, or even see them in search results.
- **Visualizations.** You can view results as topic maps, sunburst charts, maps, or in table format, as well as in list format.

- **Save searches.** You can save searches and run them again later, save a result set as a snapshot, or compare two or more saved searches.

NOTE:

Visualizations, saved searches, and comparisons are available only to users with the FindBI role. For more information, refer to the *Find Installation Guide*.

Custom Front-End Applications

The ACI (Autonomy Content Infrastructure) Client API allows you to create your own custom front-end applications.

The ACI API enables easy communication between custom-built applications and IDOL ACI servers, as well as simple manipulation of the results sets. IDOL SDKs are available for several programming languages, including C, Java, and .NET.

For more information, refer to the *ACI API Programming Guide*.

Administrative Applications

There are several IDOL applications that allow you to manipulate and modify the data in IDOL, and to monitor the IDOL services and applications.

IDOL Admin

IDOL Admin allows you to administer the IDOL Content component. You can use the Web-based application interface to:

- monitor IDOL performance
- search the Content configuration file
- retrieve information about data in the index
- index documents
- modify data fields in index documents
- export indexed documents
- back up and restore the data index
- export diagnostics information
- control the IDOL service

For more information, refer to the *IDOL Admin User Guide*.

IDOL Site Admin

IDOL Site Admin is an administration component used to monitor and manage IDOL services. You can use IDOL Site Admin to start and stop services, and monitor individual services and the overall health of your IDOL services. You can also view all log files and configuration files.

For information on how to install and access IDOL Site Admin, refer to the *IDOL Site Admin Installation Guide*.

For information about how to use IDOL Site Admin to monitor your services, refer to the *IDOL Site Admin User Guide*.

IDOL Data Admin

NOTE:

Before IDOL version 11.3, IDOL Data Admin was known as IDOL Search Optimizer.

IDOL Data Admin allows you to manage data indexed in IDOL servers to optimize the search process for your end users. You can create and modify promotions or keywords to predefine the results that are returned to users, and control how users view them. You can also view statistical information to help you to refine the rules that you created, and make them more effective.

You can also use IDOL Data Admin administer an Answer Bank system for Answer Server, to set up your question and answer store.

For information on how to install and access IDOL Data Admin, refer to the *IDOL Data Admin Installation Guide*.

For information about how to use IDOL Data Admin to manage your data, refer to the *IDOL Data Admin User Guide*.

Chapter 4: Security in IDOL

This chapter provides an overview of the Intellectual Asset Protection System. For more information on security, refer to the *IDOL Document Security Administration Guide*.

• Security Overview	47
• User Authentication	47
• Document Security	48
• Mapped Security Example	50
• Secure Communications	51

Security Overview

HPE provides the software infrastructure that automates operations on unstructured information. This software infrastructure is based on IDOL Server.

Document security protects the information that you index into IDOL Server.

Your organization is likely to store information in many repositories. Many of these repositories have security features that apply permissions to files, so that they only authorized personnel can view them. Document security ensures that when you index information into IDOL Server, IDOL continues to enforce these permissions. In response to a query, IDOL returns only documents that a user is permitted to view.

IDOL includes the following features to protect your data:

- **User authentication.** At the front end, users must log on before they can query IDOL. IDOL can authenticate users against an existing directory service, such as Microsoft Active Directory.
- **Document security.** IDOL Server checks each document that is returned in response to a query. If the user who submitted the query does not have permission to view the document, the document is removed from the query results before the results are returned.
- **Secure communications.** You can encrypt the communications between ACI servers and any applications that use the HPE IDOL ACI API.

User Authentication

To access front-end applications, users must log on by providing their credentials. If you build your own front-end applications, you can use the IDOL API to customize the logon process.

IDOL supports authentication against industry standard systems, including:

- Windows NT logon.
- LDAP authentication.
- Other third-party authentication (for example, Lotus Notes).

A single logon allows a user to access all the systems for which they have permission. For example, if a user submits a query they receive all documents that they are permitted to view, even though these documents originated from different data repositories. The IDOL Server Community component enables you to store and update user security details for this purpose.

In addition, you can use IDOL in a single sign-on environment that uses Kerberos for authentication.

Document Security

Document security ensures that users can access only those documents for which they have the necessary permissions.

When a user logs on to a front-end application and is authenticated successfully, IDOL returns an encrypted security string to the front-end application. This string identifies the user and contains information about their group memberships. The front-end application must include this security string in every subsequent query it sends to IDOL.

When a user submits a query, IDOL compares the user's information with the permissions on each document. IDOL offers two ways of doing this, *Mapped Security* and *Unmapped Security*.

NOTE:

HPE strongly recommends using Mapped Security, because IDOL can return query responses significantly faster than when Unmapped Security is used.

Mapped Security

In the *Mapped Security* architecture, IDOL determines whether a user is permitted to view a document by comparing the user's security details against an Access Control List (ACL) that has been added to the document.

With Mapped Security, when connectors fetch information from data repositories they add an encrypted Access Control List (ACL) to a metadata field in each document. The ACL contains information about which users and groups are permitted to access the document. The documents, and therefore the ACLs, are indexed into IDOL Server.

A user might be allowed or denied permission to view a document because they are a member of a security group. This means that IDOL must consider group memberships, in addition to permissions, before it can determine whether a user can view a document. OmniGroupServer collects user and group information, and stores it, so that IDOL Server can access this information.

When a user queries IDOL Server, IDOL sends the user's security string and the documents that match the query to the Mapped Security Plug-In. The Mapped Security Plug-In compares the user's details to the ACL in each document, and determines which documents the user is permitted to view. IDOL only returns those documents in its response.

The advantage of this process is that IDOL Server can quickly respond to a query, because it does not need to connect to the original data repository to check the ACL for each document. The disadvantage is that there might be a delay between the security settings changing in the original data repository and the information being updated in IDOL Server.

Mapped Security is suitable for most environments, particularly where the security settings for documents do not change often.

Unmapped Security

In the *Unmapped Security* architecture, IDOL determines whether a user is permitted to view a document by connecting directly to data repositories and checking the user's security details against the entitlement information of the documents that matched the query.

The connection between IDOL and the data repositories is made possible by IDOL's unmapped security libraries.


The advantage of unmapped security is that the security information is current. The disadvantage is that IDOL Server has to connect to the original data repositories to check permissions for each result document. This means that there can be a significant delay between a user submitting a query, and IDOL returning its response. For this reason, HPE strongly recommends using [Mapped Security](#).

Unmapped security is suitable for environments where the security settings for documents change frequently.

In some cases, you must also configure a group server. A group server is required because it is impossible for IDOL to retrieve group information from the repositories in a reasonable time. OmniGroupServer extracts user and group information and stores it so that it is available to IDOL immediately.

Mapped Security Example

The following diagram shows the components involved in a Mapped Security architecture:



Connectors extract information from third party repositories so that the information can be indexed into IDOL Server. The connector adds an Access Control List (ACL) to a metadata field in each document. The ACL describes which users and groups are permitted to view the document. CFS indexes the document into IDOL Server.

At the same time, OmniGroupServer retrieves group memberships from the third party repositories and from directories such as Active Directory. OmniGroupServer stores this information until it is needed.

In some cases, where a repository uses its own system for storing users and groups, OmniGroupServer queries a connector to retrieve group information.

If the permissions set on a file in a repository are changed, the connector sends the update to CFS and the document's ACL is updated in IDOL Server. If a user's group memberships change, the group information is updated in OmniGroupServer the next time the group server synchronizes with the repository.

To use the front-end application, a user must log on. After authentication is successful, the front-end application sends a query to Community, to retrieve the user's security information. Community returns an encrypted `securityinfo` string that contains the names of the groups the user is a member of. The front-end stores the string, because it must be sent with all queries to IDOL server.

When a user does something in the front-end application that requires information from IDOL server (for example, starting a search), the front-end sends a query to IDOL server. IDOL Server (Content) runs the operation and sends the resulting documents and the user's security information to the Mapped Security plug-in. This compares the user's security information with the ACL of each document and returns the documents that the user is permitted to view. IDOL Server then returns these documents to the front-end.

Secure Communications

You can configure IDOL Server and other ACI servers to communicate using encryption and Secure Socket Layer (SSL) communications.

SSL Communications

You can configure Secure Socket Layer (SSL) communications for IDOL Server, as well as for other ACI servers, front-end applications, and connectors.

The exact configuration you use depends on the component you are configuring. For more information, refer to the relevant component guide.

Part 2: Install and Run IDOL Server

This section describes how to install, run, and configure IDOL Server.

- [Install IDOL](#)
- [Run IDOL Server](#)
- [Configure IDOL Server](#)
- [IDOL Performance](#)
- [Tutorial: Index Data into IDOL](#)

Chapter 5: Install IDOL

This section describes how to install IDOL and the software and hardware requirements to run IDOL.

• System Requirements	55
• Install IDOL	60
• Install IDOL Admin	72
• Licenses	73

System Requirements

This section describes the software and hardware requirements to run IDOL.

Basic Requirements

- To install IDOL services on Microsoft Windows operating systems, you must run the IDOL Server installer as a user with administrative privileges.

NOTE:

You do not require special privileges to run the IDOL Server installer on Linux or Solaris operating systems, or to run IDOL Server on any platform.

- You cannot run IDOL with restricted file system permissions (for example: disk quotas, file handle limits or memory limits).
- Your file system must permit file locking (for example, you cannot run IDOL on an NFS mount).
- Your network must support TCP/IP.
- If you are running anti-virus software on the machine hosting IDOL server, ensure it does not monitor the IDOL Server directories, which can have a serious impact on IDOL Server performance.

In addition some advanced anti-virus software can scan the network and might block some IDOL traffic, which can cause errors. Where possible, exempt the IDOL processes from this kind of network traffic analysis.

Supported Platforms

IDOL runs on a variety of Windows and UNIX platforms. For details of supported platforms, refer to the *IDOL 11.5 Release Notes*.

Recommended Minimum Hardware Specifications

HPE recommends the following minimum hardware specifications for IDOL Server.

- a dedicated SCSI disk
- 4 GB RAM

- 100 GB Disk
- a minimum of 2 dedicated CPU - Intel Xeon or AMD Opteron or above

Windows System Recommendations

This section describes recommendations and requirements for IDOL for Microsoft Windows operating system platforms.

TCP Port Requirements

On Microsoft Windows platforms, if your IDOL service is receiving a large number of actions, or you have a large number of services on the same machine, your machine might run out of available TCP ports.

If your system does not have enough ports, you can edit the following Windows registry parameters:

- TCPTimedWaitDelay
- MaxUserPort

These parameters are located in:

```
HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Services\Tcpip\Parameters\
```

For more information about these registry parameters, refer to the Microsoft Windows documentation.

Linux System Recommendations

This section describes some Linux operating system settings that you might want to modify to improve IDOL performance and function. You might need additional system privileges to make these changes. In all cases, HPE recommends that you test on your own systems to find the optimal settings for your setup.

NOTE:

For many of the settings listed in this section, values and example commands are given for Red Hat Enterprise Linux (RHEL). You might need to check the default values and suggested commands for other Linux distributions.

Open File Descriptors

The IDOL Content component can use a large number of files for certain tasks, such as indexing, particularly for large indexes. On Linux systems, the number of file descriptors is limited to prevent them from using too much memory. The memory usage of a single file descriptor is low, so it is usually safe to set the limit to a high value. HPE recommends that you increase the limit value to at least 65536 for the user that the IDOL Content component runs as.

You can find an estimate of the number of files that a process (with a specified process ID) is using by using the following command:

```
ls -l /proc/ProcessID/fd | wc -l
```

You can find the current limit for the number of files by using the following commands:


```
Ulimit -Sn  
Ulimit -Hn
```

There are two limits to the number of files that processes can use:

- The soft limit is the default limit at the start of a user session, and you can use `ulimit` commands to increase it (up to the hard limit).
- The hard limit is the absolute limit for your system.

To modify the individual process file limit for the user that the IDOL Content component runs as (*UserName*), open the `/etc/security/limits.conf` configuration file, and add or modify the following lines to modify the soft and hard file limits (where *NewValue* is the limit that you want to set, in MB of RAM):

```
UserName hard nofile NewValue  
UserName soft nofile NewValue  
UserName hard nproc NewValue  
UserName soft nproc NewValue
```

TIP:

On Ubuntu, when you use the `start-stop-daemon`, you might need to run the `ulimit` command in the init script to update the hard and soft limits when you start the service.

If you start IDOL component executable files directly (rather than using the init script) by using SSH and `su`, you might also need to modify the `/etc/pam.d/su` file to uncomment the following line:

```
# session    required    pam_limits.so
```

Transparent Huge Pages

Transparent Huge pages (THP) is the use of an abstraction layer that automates the creation, management and use of Huge Pages for memory management. In RHEL 6 and later, this feature is turned on by default, but Red Hat does not recommend the use of THP for database workloads. For IDOL, turning this setting off can improve the performance for many query operations.

You can change this setting in a terminal by using the following commands:

```
echo never > /sys/kernel/mm/redhat_transparent_hugepage/enabled  
echo never > /sys/kernel/mm/redhat_transparent_hugepage/defrag
```

TIP:

On Ubuntu, you can use the following commands:

```
echo never > /sys/kernel/mm/transparent_hugepage/enabled  
echo never > /sys/kernel/mm/transparent_hugepage/defrag
```

You can also turn this setting off at boot time, by adding the following text to the bottom of the `/etc/rc.local` file:

```
#disable THP at boot time  
if test -f /sys/kernel/mm/redhat_transparent_hugepage/enabled; then  
echo never > /sys/kernel/mm/redhat_transparent_hugepage/enabled  
fi
```

```
if test -f /sys/kernel/mm/redhat_transparent_hugepage/defrag; then  
echo never > /sys/kernel/mm/redhat_transparent_hugepage/defrag  
fi
```

TIP:

On Ubuntu, the `/etc/rc.local` ends with the line `exit 0`, and you must make any changes to the file above that line.

Tune Virtual Memory

HPE recommends tuning the virtual memory usage for Linux systems for database workloads to improve IDOL performance. In particular, you might want to modify the following settings:

- `vm.swappiness`. This setting controls how much the system favors swapping out runtime memory. The default value on Linux systems is 60, but for database workloads, the recommended value is 10. HPE recommends that you set this value to 10 or lower. A value of 0 or 1 configures the system to only use the minimum amount of swapping, but does not turn it off.
- `vm.dirty_ratio`. This setting controls the maximum percentage of memory that can contain pages that have not yet been written to disk. The default value is 20, but for database workloads the recommended value is 15.
- `vm.dirty_background_ratio`. This setting controls the percentage of memory that can contain pages that have not yet been written to disk before the system starts to write the data in the background. The default is 10, but for database workloads the recommended value is 3.

You can check the current values for these settings by running the following command in the terminal:

```
# sysctl -a | grep "vm.Parameter"
```

To change these settings permanently, update the `/etc/sysctl.conf` configuration file, to add or modify the following lines:

```
vm.swappiness = 10  
vm.dirty_ratio = 15  
vm.dirty_background_ratio = 3
```

Memory Map Counts

The IDOL Content component can use a large number of memory map areas for certain tasks, such as indexing, particularly for large indexes. On Linux systems, the number of memory mapped areas that a process can use is limited, to a value controlled by the `max_map_count` kernel parameter. To prevent IDOL from reaching this limit, HPE recommends that you increase the value of the `max_map_count` parameter to the amount of memory on the system in KB /16, but no less than 65536.

NOTE:

This recommendation applies for a single IDOL Server (or Content component) running on the host machine. If you have multiple Content components running on the same host, you might need to limit this value further.

You can find an estimate of the number of memory mapped areas that a process (with a specified process ID) is using by using the following command:

```
cat /proc/processID/maps | wc -l
```

You can find the current value of `max_map_count` by using the following command:

```
sysctl vm.max_map_count
```

You can change the value of `max_map_count` by using the following command:

```
sysctl vm.max_map_count=NewValue
```

Configure the I/O Scheduler

The Linux kernel can use several different I/O schedulers to prioritize disk input and output. By default, many Linux distributions use the Completely Fair Queuing (CFQ) scheme for I/O scheduling, which gives input and output requests equal priority. For IDOL systems, HPE recommends that you set I/O scheduling to either `deadline` or `noop`:

- The Deadline scheduler gives priority to read requests over write requests. It also imposes a deadline on all requests. After reaching the deadline, these requests gain priority over all other requests. This scheduling method helps prevent processes from becoming starved for I/O access. The Deadline scheduler is recommended for physical media drives (HDD), because it attempts to group requests for adjacent sectors on disk, reducing the amount of time that the disk spends seeking.
- The NOOP scheduler uses a simple first-in first-out approach, placing all input and output requests into a single queue. This scheduler is recommended for solid state drives (SSDs). Because SSDs do not have a physical read head, there is no performance penalty for accessing non-adjacent sectors.

HPE recommends that you always use one of these schedulers, rather than the CFQ scheduler. However, the exact method to use depends on your environment and usage, so you might need to test both of these methods on your system to see which provides the best performance.

You can check whether your disks are SSD or HDD by using the following command:

```
for i in /sys/block/sd*; do cat $i/queue/rotational; done
```

This command returns 1 for HDD, and 0 for SSD.

You can find the scheduler for your mounted devices by using the following command:

```
for i in /sys/block/sd*; do cat $i/queue/scheduler; done
```

NOTE:

On RHEL 6 and above, for multipath devices, you must also change the `dm-*` block device.

This command returns a result of the following type, with the current option surrounded with square brackets:

```
noop deadline [cfq]
```

You can change the scheduler by using a command of the following type:

```
for i in /sys/block/sd*; do echo "noop" > $i/queue/scheduler; done
```

This example changes the system to use the NOOP scheduler.

Install IDOL

The IDOL installation is a simple, single-pass installer that installs IDOL Server and some other common components. This section describes how to use the installer to install IDOL server.

Install IDOL on Windows

Use the following procedure to install IDOL on Microsoft Windows operating systems, by using the IDOL Server installer.

The IDOL Server installer provides the major IDOL components. It also includes License Server, which IDOL requires to run.

To install IDOL

1. Double-click the appropriate installer package:

IDOLServer_VersionNumber_Platform.exe

where:

VersionNumber is the product version.

Platform is your software platform.

The Setup dialog box opens.

2. Click **Next**.

The License Agreement dialog box opens.

3. Read the license agreement. Select **I accept the agreement**, and then click **Next**.

The Installation Directory dialog box opens.


4. Specify the directory to install IDOL (and optionally other components such as License Server) in. By default, the system installs on C:\HewlettPackardEnterprise\IDOLServer-

VersionNumber. Click  to choose another location. Click **Next**.

The Installation Mode dialog box opens.

5. Select **Custom**, and then click **Next**.

The License Server dialog box opens. Choose whether you have an existing License Server.

- To use an existing License Server, click **Yes**, and then click **Next**. Specify the host and ACI port of your License Server, and then click **Next**.
- To install a new instance of License Server, click **No**, and then click **Next**. Specify the ports that you want License Server to listen on, and then type the path or click  and navigate to the location of your IDOL license key file (*licensekey.dat*), which you obtained when you purchased IDOL. Click **Next**.

The Component Selection dialog box opens.

6. Click **Next**.

7. Select the check boxes for the components that you want to install, and specify the port

information for each component, or leave the fields blank to accept the default port settings. For each component, you can optionally provide port information to specify the ports that the component must use. The port that you choose must not be used by any other service. The installer allows you to set the following port types:

- ACI Port** The port that client machines use to send ACI actions to the component.
- Index Port** The port that client machines use to send index actions to the component. This port is available only for the Content, Agentstore, and DIH components.
- Service Port** The port that client machines use to send service requests to the component. This port is also the port that the component uses for License Server communication.
- Binary Data Port** The port to use for streaming audio data. This port is available only for the IDOL Speech Server component.

You can also optionally choose to install the components as a Windows service. In this case, select **Create Windows service**. You can specify a name to use for the service, or accept the default name.

If you do not want to install the component as a service at installation time, you can choose to add the services later. See [Install an IDOL Component as a Service on Windows, on page 64](#).

The following table lists the components and subcomponents that you can install, and the default port that the installer uses if you do not provide an alternative value. You must change the port information if the default port value is already in use.

Component	Default port information
Content	ACI Port: 9100 Index Port: 9101 Service Port: 9102
Agentstore	ACI Port: 9050 Index Port: 9051 Service Port: 9052
Category	ACI Port: 9020 Service Port: 9022
Community	ACI Port: 9030 Service Port: 9032
View	ACI Port: 9080 Service Port: 9082
QMS	ACI Port: 16000 Service Port: 16002
Statistics Server	ACI Port: 19870 Service Port: 19872

Controller	ACI Port: 41200 Service Port: 41202
Coordinator	ACI Port: 40200 Service Port: 40202
DAH	ACI Port: 9060 Service Port: 9062
DIH	ACI Port: 9070 Index Port: 9071 Service Port: 9072
CFS	ACI Port: 7000 Service Port: 17000
File System Connector	ACI Port: 7002 Service Port: 17002
Web Connector	ACI Port: 7006 Service Port: 17006
Media Server	ACI Port: 14000 Service Port: 14001
Speech Server	ACI Port: 15000 Service Port: 15001 Binary Data Port: 15002
Knowledge Graph	ACI Port: 10300 Service Port: 10302
Find	HTTP Port: 8080
IDOL Site Admin	HTTP Port: 8080

For more information about these components, see [IDOL Core Components, on page 25](#).

Click **Next** or **Back** to move between components.

8. After you have specified your settings, the Summary dialog box opens. Verify the settings you made and click **Next**.

The Ready to Install dialog box opens.

9. Click **Next**.

The Installing dialog box opens, indicating the progress of the installation. If you want to end the installation process, click **Cancel**.

10. After installation is complete, click **Finish** to close the installation wizard.

Install a Unified IDOL Server Setup on Windows

The unified IDOL Server setup is the most basic IDOL setup, for training and test environments. For more information, see [Unified Setup, on page 28](#).

The following procedure describes how to install the unified IDOL Server setup on the Microsoft Windows operating system platform.

To install unified IDOL Server

1. Double-click the appropriate installer package:

`IDOLServer_VersionNumber_Platform.exe`

where:

VersionNumber is the version of the installer.

Platform is your software platform.

The Setup dialog box opens.


2. Click **Next**.

The License Agreement dialog box opens.

3. Read the license agreement. Select **I accept the agreement**, and then click **Next**.

The Installation Directory dialog box opens.


4. Specify the directory to install IDOL and License Server in. By default, the system installs on

`C:\HewlettPackardEnterprise\IDOLServer-VersionNumber`. Click  to navigate to another location. Click **Next**.

The Installation Mode dialog box opens.

5. Select **Unified**, and then click **Next**.

The License Server dialog box opens.

6. Specify the ports that you want License Server to listen on, and then type the path or click  and navigate to the location of your IDOL license key file (`licensekey.dat`), which you obtained when you purchased IDOL. Click **Next**.

7. The Summary dialog box opens. Verify the settings you made and click **Next**.

The Ready to Install dialog box opens.

8. Click **Next**.

The Installing dialog box opens, indicating the progress of the installation. If you want to end the installation process, click **Cancel**.

9. After installation is complete, click **Finish** to close the installation wizard.

Windows Services

When you install a unified IDOL Server, the installer automatically installs Windows Services for the License Server and the IDOL Proxy component (which automatically starts up other IDOL components when it starts). You can use these Windows Services to easily start and stop your IDOL Server while you test configurations.

You can also manually install components as Windows Services. See [Install an IDOL Component as a Service on Windows, on the next page](#).

Install an IDOL Component as a Service on Windows

On Microsoft Windows operating systems, you can install any IDOL component as a Windows service. Installing a component as a Windows service makes it easy to start and stop the component, and you can configure a component to start automatically when you start Windows.

Use the following procedure to install IDOL as a Windows service from a command line.

To install a component as a Windows service

1. Open a command prompt with administrative privileges (right-click the icon and select **Run as administrator**).
2. Navigate to the directory that contains the component that you want to install as a service.
3. Send the following command:

```
Component.exe -install
```

where *Component.exe* is the executable file of the component that you want to install as a service.

The `-install` command has the following optional arguments:

<code>-start {[auto] [manual] [disable]}</code>	The startup mode for the component. Auto means that Windows services automatically starts the component. Manual means that you must start the service manually. Disable means that you cannot start the service. The default option is Auto .
<code>-username <i>UserName</i></code>	The user name that the service runs under. By default, it uses a local system account.
<code>-password <i>Password</i></code>	The password for the service user.
<code>-servicename <i>ServiceName</i></code>	The name to use for the service. If your service name contains spaces, use quotation marks (") around the name. By default, it uses the executable name.
<code>-displayname <i>DisplayName</i></code>	The name to display for the service in the Windows services manager. If your display name contains spaces, use quotation marks (") around the name. By default, it uses the service name.
<code>-depend <i>Dependency1</i> [,<i>Dependency2</i> ...]</code>	A comma-separated list of the names of Windows services that Windows must start before the new service. For example, you might want to add the License Server as a dependency.

For example:

```
Component.exe -install -servicename ServiceName -displayname "Component Display Name" -depend LicenseServer
```

After you have installed the service, you can start and stop the service from the Windows Services manager.

When you no longer require a service, you can uninstall it again.

To uninstall an IDOL Windows Service

1. Open a command prompt.
2. Navigate to the directory that contains the component service that you want to uninstall.
3. Send the following command:

```
Component.exe -uninstall
```

where *Component.exe* is the executable file of the component service that you want to uninstall.

If you did not use the default service name when you installed the component, you must also add the `-servicename` argument. For example:

```
Component.exe -uninstall -servicename ServiceName
```

Install IDOL on UNIX

Use the following procedure to install IDOL in text mode on UNIX platforms.

To install IDOL on UNIX

1. Open a terminal in the directory in which you have placed the installer, and enter the following command:

```
./IDOLServer_VersionNumber_Platform.exe --mode text
```

where:

VersionNumber is the product version

Platform is the name of your UNIX platform

NOTE:

Ensure that you have execute permission for the installer file.

The console installer starts and displays the Welcome screen.

2. Read the information and then press the `Enter` key.
The license information is displayed.
3. Read the license information, pressing `Enter` to continue through the text. After you finish reading the text, type `y` to accept the license terms.
4. Type the path to the location where you want to install the servers, or press `Enter` to accept the default path.
The Installation Mode screen is displayed.
5. Press `2` to select the Custom installation mode.
The License Server screen opens. Choose whether you have an existing License Server.
 - To use an existing License Server, type `y`. Specify the host and port details for your License Server (or press `Enter` to accept the defaults), and then press `Enter`. Go to Step 7.
 - To install a new instance of License Server, type `n`.
6. If you want to install a new License Server, provide information for the ports that the License

Server uses.

- a. Type the value for the ACI Port and press `Enter` (or press `Enter` to accept the default value).

ACI Port The port that client machines use to send ACI actions to the License Server.

- b. Type the value for the Service Port and press `Enter` (or press `Enter` to accept the default value).

Service Port The port by which you send service actions to the License Server. This port must not be used by any other service.

- c. Type the location of your IDOL license key file (`licensekey.dat`), which you obtained when you purchased IDOL. Press `Enter`.

7. The Component Selection screen is displayed. Press `Enter`. When prompted, type `y` for the components that you want to install. Specify the port information for each component, and then press `Enter`. Alternatively, leave the fields blank and press `Enter` to accept the default port settings.

For each component, you can optionally provide port information to specify the ports that the component must use. The port that you choose must not be used by any other service. The installer allows you to set the following port types:

ACI Port The port that client machines use to send ACI actions to the component.

Index Port The port that client machines use to send index actions to the component. This port is available only for the Content, Agentstore, and DIH components.

Service Port The port that client machines use to send service requests to the component. This port is also the port that the component uses for License Server communication.

Binary Data Port The port to use for streaming audio data. This port is available only for the IDOL Speech Server component.

The following table lists the components and subcomponents that you can install, and the default port that the installer uses if you do not provide an alternative value. You must change the port information if the default port value is already in use.

Component	Default port information
Content	ACI Port: 9100 Index Port: 9101 Service Port: 9102
Agentstore	ACI Port: 9050 Index Port: 9051 Service Port: 9052
Category	ACI Port: 9020 Service Port: 9022

Community	ACI Port: 9030 Service Port: 9032
View	ACI Port: 9080 Service Port: 9082
QMS	ACI Port: 16000 Service Port: 16002
Statistics Server	ACI Port: 19870 Service Port: 19872
Controller	ACI Port: 41200 Service Port: 41202
Coordinator	ACI Port: 40200 Service Port: 40202
DAH	ACI Port: 9060 Service Port: 9062
DIH	ACI Port: 9070 Index Port: 9071 Service Port: 9072
CFS	ACI Port: 7000 Service Port: 17000
File System Connector	ACI Port: 7002 Service Port: 17002
Web Connector (Linux)	ACI Port: 7006 Service Port: 17006
HTTP Connector (Solaris)	ACI Port: 7004 Service Port: 17004
Media Server	ACI Port: 14000 Service Port: 14001
Speech Server	ACI Port: 15000 Service Port: 15001 Binary Data Port: 15002
Knowledge Graph	ACI Port: 10300 Service Port: 10302
Find	HTTP Port: 8080
IDOL Site Admin	HTTP Port: 8080

For more information about these components, see [IDOL Core Components, on page 25](#).

NOTE:
These ports must not be used by any other service.

The Init Scripts screen is displayed.

8. Type the user that the server should run as, and then press `Enter`.

NOTE:
The installer does not create this user. It must exist already.

9. Type the group that the server should run under, and then press `Enter`.

NOTE:
If you do not want to generate init scripts for installed components, you can simply press `Enter` to move to the next stage of the installation process without specifying a user or group.

The Summary screen is displayed.

10. Verify the settings that you made, then press `Enter` to begin installation.

The Installing screen is displayed.

This screen indicates the progress of the installation process.

The Installation Complete screen is displayed.

11. Press `Enter` to finish the installation.

Install an IDOL Component as a Service on Linux

On Linux operating systems, you can configure a component as a service to allow you to easily start and stop it. You can also configure the service to run when the machine boots. The following procedures describe how to install IDOL as a service on Linux.

NOTE:
To use these procedures, you must have `root` permissions.

NOTE:
When you install IDOL on Linux, the installer prompts you to supply a user name to use to run the server. The installer populates the init scripts, but it does not create the user in your system (the user must already exist).

The procedure that you must use depends on the operating system and boot system type.

- For Linux operating system versions that use `systemd` (including CentOS 7, and Ubuntu version 15.04 and later), see [Install a Component as a Service for a `systemd` Boot System](#), below.
- For Linux operating system versions that use System V, see [Install a Component as a Service for a System V Boot System](#), on the next page.

Install a Component as a Service for a `systemd` Boot System

NOTE:
If your setup has an externally mounted drive that IDOL uses, you might need to modify the init

script. The installed init script contains examples for an NFS mount requirement.

To install an IDOL component as a service

1. Run the appropriate command for your Linux operating system environment to copy the init scripts to your `init.d` directory.

- Red Hat Enterprise Linux (and CentOS)

```
cp IDOLInstalLDir/scripts/init/systemd/componentname
/etc/systemd/system/componentname.service
```

- Debian (including Ubuntu):

```
cp IDOLInstalLDir/scripts/init/systemd/componentname
/lib/systemd/system/componentname.service
```

where *componentname* is the name of the init script that you want to use, which is the name of the component executable (without the file extension).

For other Linux environments, refer to the operating system documentation.

2. Run the following commands to set the appropriate access, owner, and group permissions for the component:

- Red Hat Enterprise Linux (and CentOS)

```
chmod 755 /etc/systemd/system/componentname
chown root /etc/systemd/system/componentname
chgrp root /etc/systemd/system/componentname
```

- Debian (including Ubuntu):

```
chmod 755 /lib/systemd/system/componentname
chown root /lib/systemd/system/componentname
chgrp root /lib/systemd/system/componentname
```

where *componentname* is the name of the component executable that you want to run (without the file extension).

For other Linux environments, refer to the operating system documentation.

3. (Optional) If you want to start the component when the machine boots, run the following command:

```
systemctl enable componentname
```

Install a Component as a Service for a System V Boot System

To install an IDOL component as a service

1. Run the following command to copy the init scripts to your `init.d` directory.

```
cp IDOLInstalLDir/scripts/init/systemv/componentname /etc/init.d/
```

where *componentname* is the name of the init script that you want to use, which is the name of the component executable (without the file extension).

2. Run the following commands to set the appropriate access, owner, and group permissions for the

component:

```
chmod 755 /etc/init.d/componentname  
chown root /etc/init.d/componentname  
chgrp root /etc/init.d/componentname
```

3. (Optional) If you want to start the component when the machine boots, run the appropriate command for your Linux operating system environment:

- Red Hat Enterprise Linux (and CentOS):

```
chkconfig --add componentname  
chkconfig componentname on
```

- Debian (including Ubuntu):

```
update-rc.d componentname defaults
```

For other Linux environments, refer to the operating system documentation.

Use Unattended Installation

For many IDOL setups, you might want to install multiple sets of the same components with the same settings, on different machines. In this case, you can use the *unattended* installation mode.

NOTE:

Unattended installation is a feature of the BitRock InstallBuilder installers. For more information, refer to the InstallBuilder documentation.

You can use this procedure to install any IDOL component that uses an InstallBuilder installer (for example, connectors).

For an unattended installation mode, you provide all the installation options at the command line, and the installer runs without any further user interaction. You can also supply the installation options in an options file.

To retrieve a full list of installation options for the installer

1. Open a command prompt.
2. Navigate to the directory that contains the installer file.
3. Run the following command:

```
InstallerName.exe --help
```

where *InstallerName* is the full name of the installer file, for example `IDOLServer_11.0.0_WINDOWS_X86_64`.

This option displays the help page for the installer, which includes all the options that you can set on the command line.

To run the installer in unattended mode

1. Open a command prompt.
2. Navigate to the directory that contains the installer file.
3. Run the installer with the `--mode` command-line option set to `unattended`. For example:

```
InstallerName.exe --mode unattended
```

where *InstallerName* is the full name of the installer file, for example `IDOLServer_11.0.0_WINDOWS_X86_64`.

The default option installs a unified IDOL Server. To install a custom set of components, you can specify additional installation parameters on the command line, or you can set the `--optionfile FileName` option to specify a file that contains the installation options. For example:

```
InstallerName.exe --mode unattended --optionfile IDOLinstall.options
```

The option file must contain a list of *key=value* pairs, with one parameter on each line. For example:

```
param1=value1  
param2=value2
```

For more information, refer to the BitRock InstallBuilder documentation.

Install Separate IDOL Components

For a component setup, where you install and run each IDOL component separately, you can download and install individual component ZIP packages, if you do not want to use the IDOL Server installer.

Each component ZIP package provides a configuration file, and the required modules for the component. It also includes a component-specific Reference, which contains the actions and configuration parameters relevant to that component.

To install the component, you extract the ZIP package to the host where you want to run the component. You must also update the `[License]` section of the configuration file to include appropriate license information. If you have more than one version of a component running on the same host, you might also need to update the port information to ensure that each component uses a different port.

Component packages are available for the core IDOL Server components, as well as the Filesystem and HTTP CFS Connectors.

On Microsoft Windows operating systems, you can optionally also add components as a Windows service. See [Install an IDOL Component as a Service on Windows, on page 64](#).

IDOL Component Dependencies

When you install a unified IDOL Server, the installer automatically configures the components with associated dependencies. When you use a custom installation, or install components using ZIP packages, you must ensure that you configure these dependencies in the component configuration files.

For example, you must configure:

- the Category component with the details of the data index (Content component), and the category index (Agentstore component).
- the Community component with the details of the agent index (Agentstore component).
- the View component with the details of a Content component that it can use for highlighting.
- the QMS component with the details of the promotion Agentstore component.

NOTE:

If you install Agentstore with the IDOL Server installer, there are two available configuration files, `agentstore.cfg` and `agentstore_qms.cfg`. The default `agentstore.cfg` file is for use with Category and Community. To use Agentstore with QMS, rename `agentstore_qms.cfg` to `agentstore.cfg` (and rename or remove the default file), or run Agentstore from the command line to choose the configuration file. See [Start IDOL from the Command Line, on page 80](#).

- the Media Server component with details of a Speech Server component to use for speech processing and analysis.
- the Connector and CFS with the details of the component to send data to.
- the DIH and DAH components with details of the components to distribute to.

This list is not exhaustive. For more information about component dependencies, refer to the component documentation.

Install IDOL Admin

IDOL Admin is an administration interface for performing IDOL administration tasks, such as gathering status information, retrieving information about data in the index, and sending index actions to administer the index.

The IDOL Server installer includes the `admin.dat` file when you install any IDOL component. For each component that you install, the installer sets:

- the `AdminFile` configuration parameter in the `[Paths]` section of the configuration file to the path to the installed `admin.dat` file.
- the `Access-Control-Allow-Origin` configuration parameter in the `[Service]` section of the configuration file to allow the local host and the component port to access the IDOL Admin interface.

You might need to modify the `Access-Control-Allow-Origin` configuration parameter to allow access to IDOL Admin from other hosts. For more information, refer to the Reference for your component.

Supported Browsers

IDOL Admin supports the following browsers:

- Internet Explorer 11 and later
- Edge
- Chrome (latest version)
- Firefox (latest version)

Access IDOL Admin

You access IDOL Admin from a web browser. You can access the interface only through URLs that are set in the `Access-Control-Allow-Origin` parameter in the ACI server or component configuration file.

To access IDOL Admin from the host that it is installed on

- Type the following URL into the address bar of your web browser:

`http://localhost:port/action=admin`

where *port* is the ACI server or component ACI port.

To access IDOL Admin from a different host

- Type the following URL into the address bar of your web browser:

`http://host:port/action=admin`

where:

host is the name or IP address of the host that IDOL Admin is installed on.

port is the ACI server or component ACI port of the IDOL Admin host.

Licenses

To use IDOL solutions, you must have a running License Server, and a valid license key file for the products that you want to use. Contact HPE Big Data Support to request a license file for your installation.

License Server controls the IDOL licenses, and assigns them to running components. License Server must run on a machine with a static, known IP address, MAC address, or host name. The license key file is tied to the IP address and ACI port of your License Server and cannot be transferred between machines. For more information about installing License Server and managing licenses, see the *License Server Administration Guide*.

When you start IDOL, it requests a license from the configured License Server. You must configure the host and port of your License Server in the IDOL configuration file.

You can revoke the license from a product at any time, for example, if you want to change the client IP address or reallocate the license.

CAUTION:

Taking any of the following actions causes the licensed module to become inoperable.

You **must not**:

- Change the IP address of the machine on which a licensed module runs (if you use an IP address to lock your license).
- Change the service port of a module without first revoking the license.
- Replace the network card of a client without first revoking the license.
- Remove the contents of the `license` and `uid` directories.

All modules produce a `license.log` and a `service.log` file. If a product fails to start, check the contents of these files for common license errors. See [Troubleshoot License Errors, on page 75](#).

Display License Information

You can verify which modules you have licensed either by using the IDOL Admin interface, or by sending the `LicenseInfo` action from a web browser.

To display license information in IDOL Admin

- In the **Control** menu of the IDOL Admin interface for your License Server, click **Licenses**. The **Summary** tab displays summary information for each licensed component, including:
 - The component name.
 - The number of seats that the component is using.
 - The total number of available seats for the component.
 - (Content component only) The number of documents that are currently used across all instances of the component.
 - (Content component only) The maximum number of documents that you can have across all instances of the component.

The **Seats** tab displays details of individual licensed seats, and allows you to revoke licenses.

To display license information by sending the `LicenseInfo` action

- Send the following action from a web browser to the running License Server.

```
http://LicenseServerHost:Port/action=LicenseInfo
```

where:

LicenseServerHost is the IP address of the machine where License Server resides.

Port is the ACI port of License Server (specified by the `Port` parameter in the `[Server]` section of the License Server configuration file).

In response, License Server returns the requested license information. This example describes a license to run four instances of IDOL Server.

```
<?xml version="1.0" encoding="UTF-8" ?>
<autnresponse xmlns:autn="http://schemas.autonomy.com/aci/">
  <action>LICENSEINFO</action>
  <response>SUCCESS</response>
  <responsedata>
    <LicenseDiSH>
      <LICENSEINFO>
        <autn:Product>
          <autn:ProductType>IDOLSERVER</autn:ProductType>
          <autn:TotalSeats>4</autn:TotalSeats>
          <autn:SeatsInUse>0</autn:SeatsInUse>
        </autn:Product>
      </LICENSEINFO>
    </LicenseDiSH>
  </responsedata>
</autnresponse>
```

```
</responsedata>  
</autnresponse>
```

Revoke a Client License

After you set up licensing, you can revoke licenses at any time, for example, if you want to change the client configuration or reallocate the license. The following procedure revokes the license from a component.

NOTE:

If you cannot contact the client (for example, because the machine is inaccessible), you can free the license for reallocation by sending an `AdminRevokeClient` action to the License Server. For more information, see the *License Server Administration Guide*.

To revoke a license

1. Stop the IDOL solution that uses the license.
2. At the command prompt, run the following command:

```
InstalLDir/ExecutableName[.exe] -revokelicense -configfile cfgFilename
```

This command returns the license to the License Server.

You can send the `LicenseInfo` action from a web browser to the running License Server to check for free licenses. In this sample output from the action, one IDOL Server license is available for allocation to a client.

```
<autn:Product>  
  <autn:ProductType>IDOLSERVER</autn:ProductType>  
  <autn:Client>  
    <autn:IP>192.123.51.23</autn:IP>  
    <autn:ServicePort>1823</autn:ServicePort>  
    <autn:IssueDate>1063192283</autn:IssueDate>  
    <autn:IssueDateText>10/09/2003 12:11:23</autn:IssueDateText>  
  </autn:Client>  
  <autn:TotalSeats>2</autn:TotalSeats>  
  <autn:SeatsInUse>1</autn:SeatsInUse>  
</autn:Product>
```

Troubleshoot License Errors

The table contains explanations for typical licensing-related error messages.

License-related error messages

Error message	Explanation
Error: Failed to update license from the license server. Your license cache details do not match the current service configuration. Shutting the service down.	The configuration of the service has been altered. Verify that the service port and IP address have not changed since the service started.

License-related error messages, continued

Error message	Explanation
Error: License for <i>ProductName</i> is invalid. Exiting.	The license returned from the License Server is invalid. Ensure that the license has not expired.
Error: Failed to connect to license server using cached licensed details.	Cannot communicate with the License Server. The product still runs for a limited period; however, you should verify whether your License Server is still available.
Error: Failed to connect to license server. Error code is SERVICE: <i>ErrorCode</i>	Failed to retrieve a license from the License Server or from the backup cache. Ensure that your License Server can be contacted.
Error: Failed to decrypt license keys. Please contact Autonomy support. Error code is SERVICE:<i>ErrorCode</i>	Provide HPE Big Data Support with the exact error message and your license file.
Error: Failed to update the license from the license server. Shutting down	Failed to retrieve a license from the License Server or from the backup cache. Ensure that your License Server can be contacted.
Error: Your license keys are invalid. Please contact Autonomy support. Error code is SERVICE:<i>ErrorCode</i>	Your license keys appear to be out of sync. Provide HPE Big Data Support with the exact error message and your license file.
Failed to revoke license: No license to revoke from server.	The License Server cannot find a license to revoke.
Failed to revoke license from server <i>LicenseServer Host:LicenseServerPort</i>. Error code is <i>ErrorCode</i>	Failed to revoke a license from the License Server. Provide HPE Big Data Support with the exact error message.
Failed to revoke license from server. An instance of this application is already running. Please stop the other instance first.	You cannot revoke a license from a running service. Stop the service and try again.
Failed to revoke license. Error code is SERVICE:<i>ErrorCode</i>	Failed to revoke a license from the License Server. Provide HPE Big Data Support with the exact error message.
Your license keys are invalid. Please contact Autonomy Support. Error code is ACISERVER:<i>ErrorCode</i>	Failed to retrieve a license from the License Server. Provide HPE Big Data Support with the exact error message and your license file.
Your product ID does not match the generated ID.	Your installation appears to be out of sync. Forcibly revoke the license from the License Server and rename the <code>license</code> and <code>uid</code> directories.
Your product ID does not match this	The service port for the module or the IP

License-related error messages, continued

Error message	Explanation
configuration.	address for the machine appears to have changed. Check your configuration file.

Find the Number of Licensed Documents

You can send the LicenseInfo action to your License Server to find out how many documents your IDOL installation is licensed for.

The LicenseInfo response includes a list of the components that you have a license for. For the Content component license information, the following tags show your licensed number of documents:

Tag	Description
<autn:AggregateLimit name="maxdocuments">	The maximum number of documents that you can have across all instances of the Content component that connect to the License Server.
<autn:AggregateUsage name="maxdocuments">	The number of documents that are currently used across all instances of the Content component that connect to the License Server.

In addition, the <autn:Product> section of the response has an <autn:Client> subsection that provides details of the services that are using the license. In this section, the <autn:TrackedLimit name="maxdocuments"> shows how many documents the particular instance of the content component is using.

For example:

```
<autn:Product>
  <autn:ProductType>SUIR</autn:ProductType>
  <autn:ComponentName>content</autn:ComponentName>
  <autn:Client>
    <autn:IP>12.34.56.78</autn:IP>
    <autn:MAC>01.23.45.67.89.AB</autn:MAC>
    <autn:HostName>MyHost.example.com</autn:HostName>
    <autn:ServicePort>9993</autn:ServicePort>
    <autn:IssueDate>1379588409</autn:IssueDate>
    <autn:ProductUID>1706984408</autn:ProductUID>
    <autn:TrackedLimit name="maxdocuments">10001?
  </autn:TrackedLimit>
    <autn:IssueDateText>19/09/2013 11:00:09</autn:IssueDateText>
  </autn:Client>
  <autn:TotalSeats>100</autn:TotalSeats>
  <autn:SeatsInUse>1</autn:SeatsInUse>
  <autn:AggregateLimit name="maxdocuments">50000000</autn:AggregateLimit>
  <autn:AggregateUsage name="maxdocuments">10001</autn:AggregateUsage>
</autn:Product>
```


Chapter 6: Run IDOL Server

When you have installed IDOL server, you are ready to run it. This section describes the basics of using IDOL and verifying that it is running correctly.

- [Start and Stop IDOL](#) 79
- [Send Actions to IDOL](#) 82
- [Verify IDOL Runs Correctly](#) 83
- [Display Online Help](#) 84

Start and Stop IDOL

This section describes how to start and stop your IDOL service.

Start IDOL

The following sections describe the different ways that you can start IDOL and IDOL components.

Before you can start any IDOL component, you must start the License Server. When the License Server is running, you can start your other IDOL components.

TIP:

If you install a unified IDOL Server, you start IDOL Server by starting the IDOL Proxy component, which starts up the other IDOL components. The License Server must be running when you start the IDOL Proxy component.

In a unified installation on Microsoft Windows platforms, License Server and IDOL Server are automatically installed as services.

Start an IDOL Component on Microsoft Windows

- Double-click the *componentname.exe* file in your component installation directory.
- Start the component Service from a system dialog box. The component must be installed as a Windows Service. See [Install an IDOL Component as a Service on Windows, on page 64](#).
 1. Display the Windows **Services** dialog box.
 2. Select the **ServiceName** service for the component, and click **Start** to start the component.
 3. Click **Close** to close the **Services** dialog box.

TIP:

You can also configure the Windows Service to run automatically when you start the machine.

- Start a component from the command line. See [Start IDOL from the Command Line, on the next page](#).

Start an IDOL Component on UNIX

- Start the IDOL component service from the command line. The component must be installed as a service. See [Install an IDOL Component as a Service on Linux, on page 68](#). You can use one of the following commands to start the service:

- On systemd Linux platforms:

```
systemctl start ServiceName
```

- On System V Linux platforms:

```
service ServiceName start
```

- On Solaris platforms (using System V):

```
/etc/init.d/ServiceName start
```

TIP:

You can also configure the service to run automatically when you start the machine.

- Start the IDOL component from the command line. See [Start IDOL from the Command Line, below](#).
- (Advanced) Use the start script (`start.sh`).

A template start script is included with the component installation. However, you must update the template script with details for the component that you want to start. In most cases, HPE recommends that you use the provided init scripts instead.

Start IDOL from the Command Line

You can start IDOL components from the command line. This advanced functionality allows you to adjust certain parts of the component configuration when you start up. You might want to use this method to install components as Windows services, or to start a component with a configuration file that is in a different directory or has a different name.

NOTE:

On UNIX platforms, if you start components from the command line (rather than using the init script), you might need to set the `LD_LIBRARY_PATH` to include the `InstallDir/common` and `InstallDir/common/runtimes` directories, to ensure that the component can access the installed shared libraries.

You can also copy the shared libraries to the component working directory.

This section describes the command line arguments available to specify the component configuration files and ports, and arguments that are available for specific components to specify dependencies.

All command line arguments are optional. By default, the component attempts to read a configuration file in the same directory, and with the same name as the executable file (for example `content.exe` and `content.cfg`). The components attempt to read all the other values from the configuration file.

The following command line arguments are accepted by all the IDOL Server text processing components (Content, Category, Community, DAH, DIH, IDOL Proxy, and View).

Argument	Description	Type
-aciport	The port to use to receive ACI actions.	Number
-configfile	The path to the configuration file to use.	FilePath
-idolcomponent	Specifies that the component runs under the control of an IDOL Proxy component.	None
-install	Installs the component as a Windows service. See Install an IDOL Component as a Service on Windows, on page 64.	None
-revokelicense	Revokes the component license from the License Server.	None
-serviceport	The port to use to receive Service actions.	Number
-uninstall	Uninstalls the component Windows service. See Install an IDOL Component as a Service on Windows, on page 64.	None
-version	Displays the version and build information for the component.	None

The following command line arguments are accepted by the specified components.

Argument	Description	Type	Components
-agentstoreport	The ACI port to use to contact the Agent index (IDOL Agentstore component).	Number	Category, Community
-communityhost	The host name or IP address of the IDOL Community component.	Hostname	Category
-communityport	The ACI port to use to contact the IDOL Community component.	Number	Category
-contentport	The ACI port to use to contact the IDOL Content component.	Number	View
-datahost	The host name or IP address of the Data index (IDOL Content component).	Hostname	Category, Community
-dataport	The ACI port to use to contact the Data index (IDOL Content component).	Number	Category, Community
-indexport	The port to use to receive Index actions.	Number	Content, DIH, IDOL Proxy

The following example starts the IDOL Content component as a child component of an IDOL Proxy component. It specifies the ports to use, and the host and port for the Agent and Data indexes:

```
Community.exe -idolcomponent -configfile C:\IDOL\IDOLServer\IDOLServer.cfg -aciport 9030 -serviceport 9032 -datahost localhost -dataport 9010 -agentstoreport 9050
```

Stop IDOL

You can stop an IDOL component from running in several different ways.

- (All Platforms) Send the `Stop` service action to the component service port:

```
http://host:servicePort/action=stop
```

where *host* is the name or IP address of the host on which the IDOL component is running, and *servicePort* is the component service port (which is specified in the [Service] section of the IDOL configuration file).

To stop a unified IDOL Server, you can send the `Stop` service action to the IDOL Proxy component, and it automatically stops the other components.

NOTE:

It can take a short while for the Content process to stop after you have sent the stop action to IDOL.

- On Windows platforms, when the component is installed as a service, you can use the system dialog box to stop the service:
 1. Display the Windows **Services** dialog box.
 2. Select the **IDOLServer** service, and click **Stop** to stop IDOL.
 3. Click **Close** to close the **Services** dialog box.
- On UNIX platforms, when the component is installed as a service, you can run one of the following commands to stop the service:
 - On systemd platforms:

```
systemctl stop ServiceName
```
 - On system V platforms:

```
service ServiceName stop
```
 - On Solaris platforms (using System V):

```
/etc/init.d/ServiceName stop
```

Send Actions to IDOL

IDOL actions are HTTP requests, which you can send, for example, from your web browser. The general syntax of these actions is:

```
http://host:port/action=action&parameters
```

where:

host is the IP address or name of the machine where IDOL is installed.

port is the IDOL ACI port. The ACI port is specified by the `Port` parameter in the [Server] section of the IDOL configuration file. For more information about the `Port` parameter, see the *IDOL Reference*.

action is the name of the action you want to run.

parameters are the required and optional parameters for the action.

NOTE:

Separate individual parameters with an ampersand (&). Separate parameter names from values with an equals sign (=). You must percent-encode all parameter values.

For more information about actions, see the *IDOL Server Reference*.

Verify IDOL Runs Correctly

When you have installed IDOL and are using it, you can run actions to verify that IDOL is running correctly.

GetRequestLog

Send a GRL (or GetRequestLog) action to any IDOL component to return a log of the requests that have been made to it, including:

- the date and time that a request was made.
- the client IP address that made the request.
- the internal thread that handled the action.

For example:

```
http://IDOLhost:port/action=GRL
```

For further details on the GRL action, refer to the *IDOL Server Reference*.

Related Topics

- [Send Actions to IDOL, on the previous page](#)

GetLicenseInfo

You can send a GetLicenseInfo action to IDOL to return information on your license. This action allows you to check whether your license is valid, which IDOL operations your license includes, and which actions you can run.

For example:

```
http://IDOLhost:port/action=GetLicenseInfo
```

The following result indicates that your license is valid:

```
-<autn:license>  
  <autn:validlicense>true</autn:validlicense>  
</autn:license>
```

The following result indicates that your license includes the IDOL Agent operation:

```
-<autn:section>  
  <autn:name>Agent</autn:name>  
  <autn:licensed>>true</autn:licensed>  
</autn:section>
```

The following result indicates that you are permitted to run Query actions:

```
-<autn:section>  
  <autn:name>Query</autn:name>  
  <autn:licensed>>true</autn:licensed>  
</autn:section>
```

GetStatus

Use the GetStatus action to check whether the IDOL service is running.

For example:

```
http://IDOLhost:port/action=GetStatus
```

GetVersion

Use the GetVersion action to check the version number of an IDOL service.

For example:

```
http://IDOLhost:port/action=GetVersion
```

Display Online Help

You can display *IDOL Expert*, the *IDOL Server Reference*, and the *IDOL Server Administration Guide* by sending an action from your Web browser.

For IDOL to display help, the help data file (`help.dat`) must be available in the same directory as the service instance.

To display help for IDOL

1. Start IDOL server.
2. Send the following action from your Web browser:

```
http://IDOLhost:port/action=Help
```

where,

<i>IDOLhost</i>	is the IP address or name of the machine on which IDOL is installed.
<i>port</i>	is the ACI port by which you send actions to IDOL (set by the <code>Port</code> parameter in the IDOL configuration file [Server] section).

For example:

<http://12.3.4.56:9000/action=help>

3. On the help landing page, click one of the following options to open the relevant help:

Admin Guide	The <i>IDOL Server Administration Guide</i> describes how to configure and use IDOL Server.
Reference	The <i>IDOL Server Reference</i> contains details of all the actions and configuration parameters that you can set in IDOL Server.
IDOL Expert	<i>IDOL Expert</i> provides conceptual overviews and expert knowledge of IDOL server and its features and functionality.

If you are new to IDOL, you can use IDOL Expert to find out more about the different functions that IDOL can perform. Use the *IDOL Server Administration Guide* to get specific information about configuration and administration procedures. Use the *IDOL Server Reference* to look up specific configuration parameters and actions.

The navigation panel for the IDOL Server Reference lists the following options to display reference information:

Tab	Description
Actions	Describes the actions you can send to IDOL. Actions allow you to query IDOL, and to instruct it to perform a variety of operations.
Configuration Parameters	Describes the parameters that determine how the IDOL operates. Configuration parameters are set in the IDOL configuration file.
Index Actions	Describes the index actions you send to IDOL. Index actions allow you to index content into IDOL, and to administer the IDOL Data index.
Service Actions	Describes service actions. Service actions allow you to return data about the IDOL service, and to control IDOL.

Related Topics

- [Send Actions to IDOL, on page 82](#)

Chapter 7: Configure IDOL Server

This section describes how to configure IDOL server by editing the configuration file.

- [IDOL Configuration Files](#)87

IDOL Configuration Files

When you install a component setup, each IDOL component that you install has a default configuration file. You configure IDOL components by modifying these configuration files.

The configuration file for a component has the same name as the main component executable file. For example, the IDOL Content component is `content.exe` and the configuration file is `content.cfg`.

You can modify the configuration parameters to customize your IDOL configuration.

Unified Configuration

In a unified IDOL Server, you configure IDOL by manually editing the IDOL Server configuration file.

The IDOL configuration file contains the parameters that determine how the main IDOL components operate, except for the IDOL Agentstore component which has a separate configuration file. You can modify the configuration parameters to customize IDOL.

The IDOL Server configuration file is in the following directory of your IDOL installation:

```
InstallDir\idolserver.cfg
```

where *InstallDir* is the main installation directory.

Modify Configuration Parameter Values

You modify IDOL configuration parameters by directly editing the parameters in the configuration file. When you set configuration parameter values, you must use UTF-8.

CAUTION:

You must stop and restart IDOL for new configuration settings to take effect.

This section describes how to enter parameter values in the configuration file.

Enter Boolean Values

The following settings for Boolean parameters are interchangeable:

TRUE = true = ON = on = Y = y = 1

FALSE = false = OFF = off = N = n = 0

Enter String Values

To enter a comma-separated list of strings when one of the strings contains a comma, you can indicate the start and the end of the string with quotation marks, for example:

```
ParameterName=cat,dog,bird,"wing,beak",turtle
```

Alternatively, you can escape the comma with a backslash:

```
ParameterName=cat,dog,bird,wing\,beak,turtle
```

If any string in a comma-separated list contains quotation marks, you must put this string into quotation marks and escape each quotation mark in the string by inserting a backslash before it. For example:

```
ParameterName="<font face=\"arial\" size=\"+1\"><b>\",<p>
```

Here, quotation marks indicate the beginning and end of the string. All quotation marks that are contained in the string are escaped.

Include an External Configuration File

You can share configuration sections or parameters between ACI server configuration files. The following sections describe different ways to include content from an external configuration file.

You can include a configuration file in its entirety, specified configuration sections, or a single parameter.

When you include content from an external configuration file, the `GetConfig` and `ValidateConfig` actions operate on the combined configuration, after any external content is merged in.

CAUTION:

You cannot use shared configuration files if you use any actions or index actions that modify the configuration file (for example, the `DRECREATEDBASE` index action writes the new database configuration to the file). If IDOL updates the configuration file in this way, it overwrites the link to the shared configuration file with the imported content. This behavior prevents possible conflicts if the configuration change updates any shared configuration.

In the procedures in the following sections, you can specify external configuration file locations by using absolute paths, relative paths, and network locations. For example:

```
../sharedconfig.cfg  
K:\sharedconfig\sharedsettings.cfg  
\\example.com\shared\idol.cfg  
file://example.com/shared/idol.cfg
```

Relative paths are relative to the primary configuration file.

NOTE:

You can use nested inclusions, for example, you can refer to a shared configuration file that references a third file. However, the external configuration files must not refer back to your original configuration file. These circular references result in an error, and IDOL does not start.

Similarly, you cannot use any of these methods to refer to a different section in your primary configuration file.

Include the Whole External Configuration File

This method allows you to import the whole external configuration file at a specified point in your configuration file.

To include the whole external configuration file

1. Open your configuration file in a text editor.
2. Find the place in the configuration file where you want to add the external configuration file.
3. On a new line, type a left angle bracket (<), followed by the path to and name of the external configuration file, in quotation marks (""). You can use relative paths and network locations. For example:

```
< "K:\sharedconfig\sharedsettings.cfg"
```

4. Save and close the configuration file.

Include Sections of an External Configuration File

This method allows you to import one or more configuration sections from an external configuration file at a specified point in your configuration file. You can include a whole configuration section in this way, but the configuration section name in the external file must exactly match what you want to use in your file. If you want to use a configuration section from the external file with a different name, see [Merge a Section from an External Configuration File, on the next page](#).

To include sections of an external configuration file

1. Open your configuration file in a text editor.
2. Find the place in the configuration file where you want to add the external configuration file section.
3. On a new line, type a left angle bracket (<), followed by the path to and name of the external configuration file, in quotation marks (""). You can use relative paths and network locations. After the configuration file name, add the configuration section name that you want to include. For example:

```
< "K:\sharedconfig\extrasettings.cfg" [License]
```

NOTE:

You cannot include a section that already exists in your configuration file.

4. Save and close the configuration file.

Include a Parameter from an External Configuration File

This method allows you to import a parameter from an external configuration file at a specified point in your configuration file. You can include a section or a single parameter in this way, but the value in the external file must exactly match what you want to use in your file.

To include a parameter from an external configuration file

1. Open your configuration file in a text editor.
2. Find the place in the configuration file where you want to add the parameter from the external configuration file.
3. On a new line, type a left angle bracket (<), followed by the path to and name of the external configuration file, in quotation marks (""). You can use relative paths and network locations. After the configuration file name, add the name of the configuration section name that contains the parameter, followed by the parameter name. For example:

```
< "license.cfg" [License] LicenseServerHost
```

To specify a default value for the parameter, in case it does not exist in the external configuration file, specify the configuration section, parameter name, and then an equals sign (=) followed by the default value. For example:

```
< "license.cfg" [License] LicenseServerHost=localhost
```

4. Save and close the configuration file.

Merge a Section from an External Configuration File

This method allows you to include a configuration section from an external configuration file as part of your IDOL configuration file. For example, you might want to specify a standard SSL configuration section in an external file and share it between several servers. You can use this method if the configuration section that you want to import has a different name to the one you want to use.

To merge a configuration section from an external configuration file

1. Open your configuration file in a text editor.
2. Find or create the configuration section that you want to include from an external file. For example:

```
[SSLOptions1]
```

3. After the configuration section name, type a left angle bracket (<), followed by the path to and name of the external configuration file, in quotation marks (""). You can use relative paths and network locations. For example:

```
[SSLOptions1] < "../sharedconfig/ssloptions.cfg"
```

If the configuration section name in the external configuration file does not match the name that you want to use in your configuration file, specify the section to import after the configuration file name. For example:

```
[SSLOptions1] < "../sharedconfig/ssloptions.cfg" [SharedSSLOptions]
```

In this example, IDOL uses the values in the [SharedSSLOptions] section of the external configuration file as the values in the [SSLOptions1] section of the IDOL configuration file.

NOTE:

You can include additional configuration parameters in the section in your file. If these parameters also exist in the imported external configuration file, IDOL uses the values in the local configuration file. For example:

```
[SSLOptions1] < "ssloptions.cfg" [SharedSSLOptions]  
SSLCACertificatesPath=C:\IDOL\HTTPConnector\CACERTS\
```

4. Save and close the configuration file.

Chapter 8: IDOL Performance

This section discusses methods you can use to improve the performance of your IDOL system.

For further recommendations about improving IDOL performance, refer to *IDOL Expert*.

- [Performance Overview](#) 93
- [Optimize the Index Process](#) 94
- [Optimize Query Operations](#) 96
- [Optimize Distributed Systems](#) 99

Performance Overview

IDOL Server performance generally refers to the performance of the following IDOL Content component operations:

- **Index Data.** Documents in IDX or XML file format are added to the IDOL Content component.
- **Query IDOL server.** Users send natural language, keyword or Boolean queries to the IDOL Content component, which analyzes the concept of the query and returns documents that are conceptually similar.

To optimize IDOL Server performance, you must consider the relationship between these operations. The appropriate setup for your system depends on whether your priority is achieving fast query speeds or making new information available to users as quickly as possible.

Schedule Index and Query Operations

The IDOL Content component provides the fastest query responses when it is not indexing, and indexes fastest when it is not being queried.

One way to improve performance for both indexing and querying is to schedule indexing so that it occurs when the query load is at a minimum. For example, if your users send most queries during office hours, you can schedule indexing tasks to run during the night.

When you consider this approach, you must also consider how important up-to-date information is in your system. In some systems, users must be able to access the latest information immediately. For other systems, weekly updates to your IDOL data index are enough.

Use a Component Setup

You can improve the performance of IDOL operations by using a component setup. In this setup, you can place each IDOL component on its own hardware, maximizing the available resources.

Related Topics

- [Types of IDOL Systems, on page 25.](#)

Optimize the IDOL Content component

There are many ways to optimize the content that you store in the IDOL Content component. You can improve both index and query performance by using a stop word list, and configuring the IDOL Content component to not index numbers.

Use a Stop Word List

A stop word list is a list of common words that do not convey much meaning to sentences. For example, in English, words such as *the* and *that* do not add extra meaning in a sentence.

By default, the IDOL Content component does not index any stop words, and does not use them to search documents. This option reduces the number of terms that the IDOL Content component must index. It also means that the IDOL Content component retrieves documents based only on the relevant terms in the query.

To optimize index and query performance, ensure that your stop word list includes all common words for your system.

Index Numbers

You can determine how the IDOL Content component treats numbers when indexing documents. By default, the IDOL Content component does not index any numbers. This option reduces the number of terms that the IDOL Content component must index and search.

You can configure the IDOL Content component to index numbers, if you want users to be able to search for numeric terms. You can also configure the IDOL Content component to index numbers only if they occur in a mixed alphanumeric word (for example, IPv6).

Optimize the Index Process

The speed of the indexing process is often less critical than the speed of the query process. However, when indexing large amounts of data into IDOL, it is important to improve the efficiency of the process where possible. In addition, the way you configure the indexing process can affect the efficiency of the query process.

Delayed Synchronization

The indexing process has the following stages:

1. The IDOL Content component creates a representation of the new data in the index cache.
2. The IDOL Content component synchronizes the cache with data that it currently contains, and stores the new data on disk, removing it from the index cache.

Delayed synchronization allows you to select how the IDOL Content component synchronizes the index cache with the IDOL data. This process is useful in systems where you schedule indexing tasks at times when IDOL is also handling queries.

The `DeLayedSync` parameter in the `[Server]` section of IDOL Content component configuration file allows you to specify whether the indexing process uses delayed synchronization.

If the `DeLayedSync` parameter is not set, IDOL synchronizes the cache as soon as the representation of data is made. New data is available to the user (as query results) quickly. Use this setting in systems where up-to-date data is the priority.

Synchronization uses resources that the IDOL Content component could otherwise use for querying. Delayed synchronization reduces this effect by collecting multiple data representations in the index cache and then synchronizing them all with IDOL data at the same time. This process is useful in systems where query speed is more important than having the most up-to-date data.

NOTE:

HPE recommends that you use delayed synchronization if you index a large number of small files (files that are smaller than 100MB).

To configure delayed synchronization

1. Open the IDOL Content component configuration file in a text editor.
2. In the `[Server]` section, set the `DeLayedSync` parameter to `True`.
3. In the `[Server]` section, set the `MaxSyncDeLay` parameter to the maximum length of time between synchronization operations.
4. In the `[IndexCache]` section, set the `IndexCacheMaxSize` parameter to the maximum size that the IDOL Content component index cache can grow to before synchronizing.

NOTE:

When the index cache reaches this maximum size, the IDOL Content component synchronizes the cache, but data is not available for searching until the `MaxSyncDeLay` has passed, or you send a `DRESYNC` index action.

5. Save and close the configuration file.
6. Restart the IDOL Content component for your changes to take effect.

Distribute IDOL Content component Data Across Multiple Disks

If your IDOL data becomes too big to store on one volume (as the stored terms, references, content, and so on increase in size), you can store the data files across multiple disk partitions.

To distribute IDOL data

1. Open the IDOL Content component configuration file in a text editor.
2. Find the `[Paths]` section, or create one if it does not exist.
3. Set the path parameters to the full paths to the directories where you want to store the corresponding part of the index. The following paths have the largest impact on performance:

Parameter	Description
<code>DynTermPath</code>	The directory that contains conceptual data.
<code>NodeTablePath</code>	The directory that contains content and structured data.

Parameter	Description
StatusPath	The directory that contains status files, and data that is streamed over the network before indexing.

For a full list of available parameters, refer to the *IDOL server Reference*.

4. Save and close the configuration file.
5. Restart the IDOL Content component for your changes to take effect.

For example:

```
[Paths]
DyntermPath=C:\idolserver\dynterm
NodetablePath=D:\idolserver\nodetable
StatusPath=E:\idolserver\status
RefIndexPath=F:\idolserver\refindex
MainPath=G:\idolserver\main
TagPath=H:\idolserver\tagindex
```

Related Topics

- [Display Online Help, on page 84](#)

Optimize Query Operations

Query operations include any action that requests data from IDOL Server (particularly the IDOL Content component), for example:

- Query
- Suggest
- SuggestOnText
- GetQueryTagValues

Query speed can be particularly important because users require a real-time response to any query. This section describes how to optimize the performance of query operations.

The IDOL Content component Query action allows you to search for documents using a number of restrictions and criteria. The main types of query are:

- **Text.** The query contains a natural language expression, which can also contain Boolean or Proximity operators. The IDOL Content component finds documents that contain concepts that are similar to the given text, and that match any additional operations.
- **FieldText.** The query contains restrictions that specify values that must occur in specified fields. For example, you can restrict a query to documents that contain the term *Dog* in the *Animal* field.

Simple Text queries are generally quicker than FieldText queries.

Optimize IDOL Fields

You can optimize `FieldText` queries by configuring field processes. These processes define how the IDOL Content component stores fields, to optimize the retrieval of the data that the field contains. For example, if the field contains a number, configure the field as a numeric field to allow IDOL server to quickly process and retrieve values for that field.

HPE recommends that you use field processes to identify types of fields appropriately. For details of how to set up field processes, refer to the *IDOL Server Administration Guide*.

The following table lists the field properties that optimize `FieldText` specifiers, and the specifiers that they optimize.

Field Property	Optimized Field Specifiers	
MatchType	BIASVAL EMPTY EXISTS MATCH MATCHALL	NOTMATCH NOTSTRING NOTWILD STRING WILD
NumericType	BIAS EMPTY EQUAL EQUALALL EQUALCOVER	EXISTS GREATER LESS NOTEQUAL NRANGE
NumericDateType	BIASDATE GTNOW	LTNOW RANGE
CountType	EQUALCOVER	MATCHCOVER
SortType	ARANGE	
BitFieldType	BITSET	

Index Fields

The IDOL Content component processes index fields linguistically. It removes stop words, and stems each term before storing the terms. This process allows IDOL to return documents that match a conceptual query or contain keyword search terms.

Define fields that contain document content as index fields. For example, the document title and body. When performing `Text` queries, IDOL checks Index fields for matching terms and concepts.

HPE recommends that you do not store URLs or content that you are unlikely to query as Index fields. HPE also recommends that you use Match fields, rather than index fields, for fields where you query only the whole value of the field.

Match Fields

The IDOL Content component stores the value of match fields in a fast look-up structure in memory. You can define fields as match field when you frequently want to retrieve documents using the whole value of this field. For example, you might define the *Author* field of a document as a Match field so that users can search for the author.

Numeric Fields

The IDOL Content component stores the value of numeric fields in a fast look-up structure in memory. You can use numeric fields to allow users to search for values or ranges of values. For example, you might make a *Price* field a numeric field so that users can restrict results to products within a price range.

Numeric Date Fields

The IDOL Content component stores the value of numeric date fields in a fast look-up table in memory. You can use numeric date fields to allow users to search for a date or range of dates. For example, you might make a *Date* field a numeric date field so that users can restrict results to those between two dates.

Count Fields

The IDOL Content component stores the number of occurrences of count fields. You can use count fields to allow users to search for more than one string in a given field.

Sort Fields

You can use sort fields to allow users to search for values in an alphabetical range. For example, you might make the *Author* field a sort field so that users can search for books alphabetically by author.

Sort fields also optimize the `Sort` action parameter when sorting on a the value of a field.

Bit Fields

Typically, bit fields contain information about the sets that a document belongs to. You can use bit fields to allow users to search for documents that only occur in a particular set. For example, you might make a *Workbook* field a bit field so that users can search for documents that occur in a particular workbook.

Parametric Fields

You can use parametric fields to allow users to restrict a search to certain parameters. For example, you might make a *Model* field a parametric field so that users can restrict a search to products of a

particular model.

You might also want to make a field both `ParametricType` and `MatchType` to allow you to list the available values and then restrict by those values.

NOTE:

You can configure a field as simultaneously `MatchType` and `ParametricType` only if you set `ParametricNumericMapping` to `false`.

Field Check Fields

The IDOL Content component stores a checksum hash of the value of a field check fields. This option allows especially fast retrieval when you restrict results to the exact value of the field. For example, you might make a *Category* field a field check field so that you can restrict a search to a particular category.

Field Check fields are also useful for the *Combine* operation. In a *Query* action, you can combine results that have the same value of the field check field. The IDOL Content component then returns only one document for this field check value.

NOTE:

Each document that you index into the IDOL Content component must contain only one `FieldCheckType` field.

Hardware Considerations

To maximize IDOL performance, ensure that your IDOL components have adequate resources assigned:

- Ensure that IDOL server has enough memory.
- Monitor disk usage to check input and output.

Related Topics

- [Recommended Minimum Hardware Specifications, on page 55](#)

Optimize Distributed Systems

Distributed IDOL systems use a Distributed Action Handler (DAH) and Distributed Index Handler (DIH) to distribute actions and index actions respectively. In these systems, you can use one of two modes:

- **Mirror mode** uses several identical child servers to process client queries. You can use this system when you have large numbers of queries, to minimize the query load on individual child servers.
- **Non-mirror mode** distributes all data across several child servers, which can improve the speed of indexing for the system, and increases the amount of data that child server contains. You can use this system when you have a large amount of data to index, or to minimize the index load on individual servers. As servers might respond more slowly while they are indexing data, this method can also improve query response speed.

In both modes, you can consider other modes in the DIH and DAH to distribute ACI and index actions most efficiently.

Related Topics

- [Types of IDOL Systems](#)

Distributed Action Handler Performance

The Distributed Action Handler distributes actions between child servers. It also performs some additional processing on actions. For example, in mirror mode it can apply templates, and recognize and distribute administrative actions to all child servers.

Depending on your system, you can use Distributed Action Handler modes to reduce the amount of additional processing that Distributed Action Handler performs, which can improve the performance of the system.

Fast Mirror Mode

In fast mirror mode, Distributed Action Handler distributes actions between mirrored child servers but does not perform any additional processing to actions. You can use this option to increase the performance of your IDOL system in mirror mode when you use Distributed Action Handler only for requesting information from child servers.

When running the Distributed Action Handler in fast mirror mode:

- You cannot send state-changing actions to the Distributed Action Handler (for example, actions including the `State` and `StoreState` action parameters).
- You cannot request additional processing on actions (for example, actions that use the `Template`, `EncryptResponse` or `Output` action parameters).
- You cannot use document IDs in action parameters.
- You cannot use distributed actions (actions to send to all child servers) or asynchronous actions.

Simple Combinator Mode

In simple combinator mode, Distributed Action Handler does not use virtual databases to distribute and combine actions in non-mirror mode. Distributed Action Handler forwards the `DatabaseMatch` action parameter to child servers, and combines results from all child servers.

This option reduces the amount of processing that Distributed Action Handler must perform to distribute and combine virtual database.

Distributed Index Handler Performance

The DIH distributes index actions to its child servers.

In standard mirror and non-mirror mode, DIH forwards the `IDX` or `XML` index file to all its child servers, with instructions to the child servers about which documents it must index. You can use different modes in the DIH to reduce the amount of data that the DIH, or the child servers, must process.

Preserve DREADD

The DIH configuration parameter `PreserveDREADD` allows it to distribute index actions more quickly.

By default, DIH accepts an IDX or XML file input and converts the original DREADD index action into a DREADDATA index action, which contains all the IDX or XML data. If you set `PreserveDREADD` to `true` in the configuration file, DIH forwards the DREADD action directly. This option reduces the network load.

To use this parameter, each child IDOL Content component must be able to access the original IDX or XML file.

You can use `PreserveDREADD` only in mirror mode, simple non-mirror mode or `DistributeOnBatch` mode. For example, you cannot use `PreserveDREADD` if you have set `DistributeByReference` to `true`. In advanced distribution modes, DIH must parse the IDX or XML file to distribute the documents correctly.

Distribute On Batch

In distribute on batch mode, DIH indexes data to alternating child servers. This method is an efficient way of distributing the data to child servers in non-mirror mode. You can use this method if:

- you index a steady stream of IDX or XML files, with similar numbers of documents. This mode might not be appropriate if you infrequently index very large IDX or XML files.
- you do not need to remove duplicate documents.

Advanced Distribution Modes

DIH has several distribution modes that determine how it distributes documents between child servers in non-mirror mode. Unlike standard non-mirror mode, these methods send each child server only the documents that it must index. This option reduces the amount of data that each child server receives.

NOTE:

If you use these methods, the number of DIH child servers is fixed. To add child servers, you must clear all engines, add the new servers, and then re-index using the DIH.

- **Distribute by reference.** DIH distributes documents to child servers based on the document reference. This method ensures that you can prevent duplicate documents that have the same reference.
- **Distribute by date.** DIH distributes documents to child servers according to the document date. You configure a date range for each child server.
- **Distribute by fields or field values.** DIH distributes documents to child servers based on the value of a certain field.
 - In distribute by fields mode, DIH sends documents with the same value of the specified field to the same child server. DIH determines which child server receives a given field value.
 - In distribute by field values mode, you can specify which child servers receive documents that contain certain field values.

Round Robin Mode

The Round-Robin indexing mode rotates indexing between several child servers, so that indexing and querying are performed by different child servers at different times.

The IDOL Content component provides the fastest queries when it is not indexing, and indexes fastest when it is not being queried. When you configure indexing for round-robin mode, DIH suspends query

handling for a specific child server. It then indexes only into this child server, which has optimal indexing performance.

After the specified time period, DIH starts indexing to a different child server and makes the previous server available for querying.

Chapter 9: Tutorial: Index Data into IDOL

This chapter describes how to configure the IDOL Content component index and index data.

- [Overview](#) 103
- [Configure the IDOL Content component Index](#) 103
- [Create Documents](#) 111
- [Pre-Process Documents](#) 111
- [Index Documents](#) 111
- [Index Data in a Distributed IDOL Setup](#) 113

Overview

The IDOL Content component data index is a representation of all the data that you can search for in IDOL Server. When the IDOL Content component indexes data, it processes the terms and fields in the new documents. This process allows it to retrieve the documents quickly when you send queries.

Before you index documents, you must configure the data index to process documents according to your requirements. You can then create and index documents, either by using an IDOL connector, or by creating an IDX or XML file and indexing them manually.

Related Topics

- [Configure the IDOL Content component Index](#)
- [Create Documents](#)
- [Pre-Process Documents](#)
- [Index Documents](#)

Configure the IDOL Content component Index

Before you begin to index documents into the IDOL Content component, you must configure the data index to process the documents correctly.

You can configure:

- IDOL databases to store your data.
- field processes to correctly identify and process the fields in the documents you index.
- language types to correctly process different languages and encodings.

Related Topics

- [Configure Databases](#)
- [Configure Fields](#)
- [Configure Languages](#)

Configure Databases

The IDOL Content component can store documents in different databases. You can configure up to 65534 databases in an IDOL Content component.

Indexing data into different databases allows you to restrict queries to data in a particular database. This can improve the performance of your queries.

You can add databases to the IDOL Content component by editing the configuration file or by sending an index action. For details of adding a database by editing the configuration file, refer to the *IDOL Server Administration Guide* and the *IDOL Server Reference*.

To add IDOL Content component databases

- Send a DRECREATEDBASE action (case sensitive) from your Web browser:

```
http://ContentHost:indexPort/DRECREATEDBASE?DREdbname=databaseName
```

where,

<i>ContentHost</i>	is the IP address or host name of the machine on which the IDOL Content component is installed.
<i>indexPort</i>	is the IDOL Content component index port (specified as <code>IndexPort</code> in the IDOL Content component configuration file [Server] section).
<i>databaseName</i>	is the name of the database that you want to create.

For example:

```
http://12.3.4.56:20001/DRECREATEDBASE?DREdbname=Archive
```

This action uses port 20001 to create a new database named *Archive* on the IDOL Content component installed on the machine with the IP address 12.3.4.56.

You can also add other index action parameters to this action, for example to make the new database internal. For details, refer to the *IDOL Server Reference*.

Configure Fields

The documents in the IDOL Content component store data in fields. These fields can contain many different types of information. For example, one field might contain the bulk of the document content, and another field might contain only the name of the author.

When you index data into the IDOL Content component, it is important to set up field processes so that the IDOL Content component treats each field correctly.

Choose Field Properties

Before you set up fields in IDOL, you must consider the fields in the documents that you want to index.

There are a large number of different field properties that you can apply to IDOL fields. Fields can contain information about the document, or values that you want to retrieve in queries.

For a complete list of the types of fields that you can store in the IDOL Content component, refer to the *IDOL Server Administration Guide* and *IDOL Server Reference*.

Fields that contain information about the document	
ReferenceType	Reference fields contain a unique document reference, which you can use to remove duplicate documents, and to retrieve a specific document. Each document must contain at least one reference field. The IDOL Content component also uses the reference field to populate the <code>autn:reference</code> metadata field.
DateType	Date fields contain the document date, which the IDOL Content component uses to populate the <code>autn:date</code> metadata field. If a document does not have a date field, Content uses the date that the document was indexed.
TitleType	Title fields contain the document title, which the IDOL Content component uses to populate the <code>autn:title</code> metadata field.
DatabaseType	Database fields contain the IDOL database that the IDOL Content component must index the document into. If the document does not contain a database field, you must specify the database in the index action.
LanguageType	Language fields contain the language type of the document, which the IDOL Content component uses to find the appropriate language configuration. The IDOL Content component also uses this field to populate the <code>autn:language</code> metadata field.
SecurityType	Security fields contain the security type of the document, which the IDOL Content component uses to index the document according to the specified security configuration.
ACLType	ACL fields contain document access control lists (ACLs), which determine the access restrictions for that document.
ExpireDateType	Expire date fields contain the date that the document expires. On this date, IDOL processes the document according to the expiry rules for the database, for example to delete the document or move it to an archive database. The IDOL Content component also uses this field to populate the <code>autn:expiredate</code> metadata field.
SectionBreakType	Section break fields contain the section number for a document section, which the IDOL Content component uses to populate the <code>autn:section</code> metadata field.
ParametricType	Parametric fields contain values that you want to use to restrict queries. In a parametric query, you can return all values that occur in a certain parametric field in all documents.
ParametricRangeType	Parametric fields contain values that you want the IDOL Content component to dynamically analyze to generate numeric ranges in parametric queries.

Fields that contain document content	
Index	<p>Index fields contain document content. The IDOL Content component processes these fields linguistically and stores terms to allow fast data retrieval. Typically, you store the document content and title field as index fields. Do not store data as index fields if:</p> <ul style="list-style-type: none"> • you do not need to query the field contents. • you query only the whole value of the fields. It is more efficient to store these fields as <code>MatchType</code> fields and query the fields using a <code>FieldText</code> specifier such as <code>MATCH</code>.
SourceType	<p>Source fields contain document content that the IDOL Content component uses to create document summaries and suggest conceptually similar documents. If you do not configure source fields, the IDOL Content component uses the configured Index fields.</p>
LangDetectType	<p>Language Detection fields contain document content that the IDOL Content component uses to automatically detect the document language. If you do not configure a language detection fields, the IDOL Content component uses the configured source fields.</p>
HighlightType	<p>Highlight fields contain document content that IDOL server can highlight when you send the <code>Highlight</code> action (or a query action with the <code>Highlight</code> parameter set to <code>true</code>).</p>
Field properties that optimize FieldText query operators	
MatchType	See Match Fields, on page 98
NumericType	See Numeric Fields, on page 98
NumericDateType	See Numeric Date Fields, on page 98
SortType	See Sort Fields, on page 98
Field properties that determine how to display fields	
HiddenType	Hidden fields do not appear
PrintType	Print fields are displayed in a <code>Query</code> action when the <code>Print</code> parameter is set to <code>Fields</code> (the default value).

Set up Field Processes

To identify properties for different fields, you must define field processes. In a field process, you define:

- the set of fields that the field process applies to
- the property that applies to this process

You can have multiple field processes that share the same property. You must then create a configuration section for each property that you use, and define the field properties.

NOTE:

Use the following formats to identify fields:

- `/FieldName` to match root-level fields
- `*/FieldName` to match all fields except root-level
- `/Path/FieldName` to match fields that the specified path points to.

Field names must not contain spaces nor accents, and they must not start with a number. For IDX documents, the IDOL Content component converts these text elements to underscores (`_`) when it indexes the fields. You must also change any queries that reference these field names to use the modified field name.

To apply processes to fields or documents that contain specific fields

1. Open the IDOL Content component configuration file in a text editor.
2. In the `[FieldProcessing]` section, list the processes to apply to fields. For example:

```
[FieldProcessing]
0=IndexFields
1=DateFields
2=DatabaseFields
3=SetReferenceFields
```

3. Create a configuration section for each process listed.
 - Set `PropertyFieldCSVs` to a comma-separated list of fields that this process applies to.
 - Set `Property` to the name of the property configuration section.

NOTE:

Each property must have a unique configuration section name.

- (Optional) Set `PropertyMatch` to a comma-separated list of values that the field must have to be processed. For example, you can use this parameter to set up a process that identifies security or language fields.
- (Optional) Set `PropertyNegativeFieldCSVs` to a comma-separated list of fields that this process does not apply to. For example, if you use a wildcard value in `PropertyFieldCSVs`, you can define any exceptions in `PropertyNegativeFieldCSVs`.

For example:

```
[IndexFields]
Property=Index
PropertyFieldCSVs=*/DRECONTENT,*/DRETITLE
```

```
[DateFields]
Property=Date
PropertyFieldCSVs=*/DREDATE,*/harvest_time
```

```
[DatabaseFields]
Property=Database
PropertyFieldCSVs=*/DREDBNAME,*/Database*
PropertyNegativeFieldCSVs=*/DatabaseNumber
```

```
[SetReferenceFields]  
Property=Reference  
PropertyFieldCSVs=*/DRREFERENCE,*/DRETITLE
```

4. Create a section for each property. Specify appropriate configuration parameters for each property. These configuration parameters define the processes that apply to all the fields (or all documents that contain the fields) that you previously associated with the processes.

For example:

```
[Index]  
Index=TRUE
```

```
[Date]  
DateType=TRUE
```

```
[Database]  
DatabaseType=TRUE
```

```
[Reference]  
ReferenceType=TRUE  
TrimSpaces=TRUE
```

NOTE:

For some properties, you must reindex your content if you want to change the property value after you have indexed content into the IDOL Content component.

For details, refer to the *IDOL Server Reference*.

5. Save and close the configuration file.
6. Restart the IDOL Content component for your changes to take effect.

Related Topics

- [Display Online Help](#)

Configure Languages

The IDOL Content component can process documents in multiple languages and encodings. For each language that you want to use, you must define the language types in the IDOL Content component configuration file. You must also configure the IDOL Content component to classify documents, either by automatically detecting the language and encoding, or by reading the language type from a field.

Define Language Types

To run the IDOL Content component in multiple languages, specify the language types you want the IDOL Content component to process. A language type is a combination of the language and encoding.

NOTE:

You must specify languages and language types before you index data into the IDOL Content component.

To define language types

1. Open the IDOL Content component configuration file in a text editor.
2. Find the [LanguageTypes] section. List the languages that you want the IDOL Content component to process. You must use ASCII characters to specify the language names.

For example:

```
[LanguageTypes]
0=English
1=Afrikaans
2=General
```

3. For each language, create a configuration section that matches the name you defined in the [LanguageTypes] section.
4. In this section, specify appropriate settings that determine how the IDOL Content component handles this language. For details on the configuration parameters you can use, refer to the *IDOL Server Reference*.
5. For each section, set the Encodings parameter to a list of the encodings and corresponding language types used by the language. List each encoding and language in the format *encoding:language*. Separate multiple language types with commas.

For example:

```
[english]
Encodings=ASCII:englishASCII,UTF8:englishUTF8
Stoplist=english.dat
IndexNumbers=1
```

```
[afrikaans]
Encodings=ASCII:afrikaansASCII,UTF8:afrikaansUTF8
IndexNumbers=1
```

```
[general]
Encodings=UTF8:generalUTF8,ASCII:generalASCII,CYRILLIC:generalCYRILLIC
IndexNumbers=1
```

6. Save and close the configuration file.
7. Restart IDOL server for your changes to take effect.

You can now configure the IDOL Content component to associate the language types you defined with documents.

Associate Language Types with Documents

After you define all the language types you want the IDOL Content component to process, set up a field process that allows the IDOL Content component to associate these language types with documents.

The way you configure the field process depends on the documents that you want to index:

- If all the documents contains a field that exactly specifies the language type, configure a field process to define this field as a LanguageType field. The language types that appear in this field

must exactly match the language types that you define in the [LanguageTypes] configuration section. See [Configure Fields, on page 104](#).

- If the documents contain a field that specifies the language, but does not exactly specify the language type, you can configure field processes to detect the language from this field data.

To configure the IDOL Content component to detect the language type from field data

1. Open the IDOL Content component configuration file in a text editor.
2. In the [FieldProcessing] section, define a field process for each language that you want to detect.

For example:

```
[FieldProcessing]
0=DetectArabic
1=DetectEnglish
2=DetectFrench
```

3. Create a configuration section with the same name as each of the field processes you defined in the [FieldProcessing] section.
4. In this section:
 - Set Property to the name of the property for the specified language type.
 - Set PropertyFieldCSVs to a comma-separated list of fields that can contain the language data.
 - Set PropertyMatch to a comma-separated list of values that this field might contain to identify the specified language type.

For example:

```
[DetectArabic]
Property=SetArabicProperty
PropertyFieldCSVs=*/DRELANGUAGETYPE,*/LANG
PropertyMatch=arabic
```

```
[DetectEnglish]
Property=SetEnglishProperty
PropertyFieldCSVs=*/DRELANGUAGETYPE,*/LANG
PropertyMatch=*eng*,uk,*british
```

```
[DetectFrench]
Property=SetFrenchProperty
PropertyFieldCSVs=*/DRELANGUAGETYPE,*/LANG
PropertyMatch=*fre*,fran*
```

5. Create a configuration section for each property that you define in the field processing sections.
6. In the property configuration section, set the LanguageType parameter to the language type to use to define documents that match this property (that is, that contain a field with a matching value for the field process). This language type must match one of the language types you configure in the [LanguageTypes] configuration section.

For example:

```
[SetArabicProperty]  
LanguageType=Arabic
```

```
[SetEnglishProperty]  
LanguageType=English
```

```
[SetFrenchProperty]  
LanguageType=French
```

7. Save and close the configuration file.
8. Restart the IDOL Content component for your changes to take effect.

Create Documents

Unless your data is in XML format, you must first import it. You can import data using one of two methods.

- **Import with a connector.** The IDOL connectors (for example, File System Connector, HTTP Connector, Database Connector, and so on) retrieve documents from different repositories and import them into IDX or XML file format. For details of how to configure a connector to import data and index it into IDOL, refer to the appropriate Connector Administration Guide.
- **Import manually.** You can create a text file in either XML or IDX format, which contains the information that you want to index into your IDOL Content component in specific IDOL fields. For details of how to create a manual IDX document, refer to the *IDOL Server Administration Guide*.

Pre-Process Documents

You can use the Connector Framework Server (CFS) to process documents before you index them. For example, you can run Education, modify document fields, categorize documents, or write the documents to disk.

For more information about how to use CFS to process your documents, refer to the *Connector Framework Server Administration Guide*.

TIP:

CFS has most of the processing functionality that is available in the deprecated IDOL IndexTasks component. HPE recommends that you use CFS to process your documents before indexing where possible, instead of the IndexTasks IDOL component.

For information about how to convert an existing IndexTasks installation to use CFS, refer to the *IndexTasks to CFS Technical Note*.

Index Documents

You can index only files in XML or IDX format into the IDOL Content component. If the data that you want to index is in XML format, you can index it directly into the IDOL Content component, without having to first *import* it (convert its content and metadata to IDX).

IDOL connectors use the DREADD and DREADDATA index actions to index data into IDOL Server. You can also use these actions to directly index data into the IDOL Content component.

Use DREADD to Index IDX and XML Files Directly

The DREADD index action (case sensitive) directly indexes an IDX or XML file that is located on the same machine as the IDOL Content component. For example:

```
http://ContentHost:indexPort/DREADD?filename&optionalParams
```

where,

<i>ContentHost</i>	is the IP address or host name of the machine on which the IDOL Content component is installed.
<i>indexPort</i>	is the indexing port of the IDOL Content component (specified in the <i>IndexPort</i> parameter in the [Server] section of the IDOL Content component configuration file.
<i>filename</i>	is the file name or location of the IDX or XML file to index.
<i>optionalParams</i>	are any optional index action parameters that you want to add. NOTE: You must set the DREDbName parameter if your IDOL Content component configuration does not contain a field process to read the database from a field.

For a full list of available parameters for the DREADD action, refer to the *IDOL Server Reference*.

Use DREADDATA to Index Data Over a Socket

The DREADDATA index action (case sensitive) allows you to directly index data over a socket into the IDOL Content component.

NOTE:

This index action requires a POST request method. For more information, refer to the *IDOL Server Administration Guide*.

For example:

```
http://  
ContentHost  
:indexPort/DREADDATA?OptionalParamsData#DREENDDATAKillDuplicatesOption\n\n
```

where,

<i>ContentHost</i>	is the IP address or host name of the machine on which the IDOL Content component is installed.
<i>indexPort</i>	is the indexing port of the IDOL Content component (specified in the <i>IndexPort</i> parameter in the [Server] section of the IDOL Content

	component configuration file.
<i>optionalParams</i>	are any optional index action parameters. DREADDATA accepts the same optional parameters as the DREADD index action, except for <i>KillDuplicatesOption</i> .
<i>Data</i>	is the content of the IDX or XML document to index. You must add #DREENDDATA to the end of your data. You can use gzipped documents, but #DREENDDATA must not be compressed. This parameter is required.
<i>KillDuplicatesOption</i>	is the option to use for removing duplicate documents.

For details of all the parameters available in the DREADDATA index action, refer to the *IDOL Server Reference*.

Index Data in a Distributed IDOL Setup

Configuration and indexing in a distributed system is very similar to using an individual IDOL Content component, except that you must configure multiple child servers, and send your index actions to the DIH, instead of Content.

In most cases, you use an identical or very similar configuration for each of your child IDOL Content components, even in non-mirror mode, to ensure that data is processed and indexed consistently. For example, you usually configure the same fields, languages, and databases, in all your child servers (although you might need to configure different port settings, and paths).

After you configure your child servers, you configure the DIH to distribute data to those child servers. For more information, refer to the *DIH Administration Guide*.

You send index actions (such as DREADD and DREADDATA) to the DIH, which distributes the data to its child servers. Depending on your DIH configuration, DIH might also perform some preprocessing.

Related Topics

- [Distributed Setup, on page 30](#)

Chapter 10: Query IDOL

This section describes the IDOL action responses, and the basics of query operations in the IDOL Content component.

- [IDOL ACI Responses](#)115
- [The Query ACI Action](#)116
- [Search Parameters](#)116
- [Response Parameters](#) 120
- [Query Response](#)123

IDOL ACI Responses

Interactions with IDOL components are structured according to the Autonomy Content Infrastructure (ACI) architecture. The ACI architecture defines how clients communicate with IDOL components and how the IDOL components respond. In these interactions, clients submit ACI requests to IDOL components as GET or POST HTTP requests, and receive ACI responses in the form of XML or JSON.

An ACI request consists of an ACI action and a collection of ACI parameters. See [Send Actions to IDOL, on page 82](#).

For example, the following ACI request sends an action to a local IDOL Server on port 9000. The action is Query, and there are two ACI parameters, Text with a value of IDOL Server, and MaxResults with a value of 20.

```
http://localhost:9000/action=Query&text=IDOL Server&MaxResults=20
```

The ACI action determines the collection of valid ACI parameters. For a given ACI action, some ACI parameters might be required while others are optional.

The ACI responses in XML use a uniform structure for the top two levels of the document tree. For example:

```
<autnresponse xmlns:autn="http://schemas.autonomy.com/aci/">  
  <action>QUERY</action>  
  <response>SUCCESS</response>  
  <responsedata>...</responsedata>  
</autnresponse>
```

Every ACI response returns the <autnresponse> node, defines the autn namespace as shown, and has three children of the document root:

<action>	contains the ACI action of the request.
<response>	contains either SUCCESS or ERROR indicating whether the ACI request succeeded or failed.
<responsedata>	contains the rest of the response, the structure of which depends on the ACI

	action.
--	---------

The Query ACI Action

The query action is an ACI action supported by the IDOL Content component that searches the corpus of indexed documents for those matching the restrictions of the query. Query action ACI parameters come in two types:

- Search parameters define restrictions on the result set of documents.
- Response parameters change the information that returns with the result set of documents.

Search Parameters

This section describes many of the most important parameters for the Query action that you can use to restrict the document result set.

For more information about each of these parameters, refer to the *IDOL Server Reference*.

Text

The `Text` parameter defines restrictions on the content of matching documents. You can specify restrictions in plain text, or by using Boolean and proximity operators.

If `Text` contains plain text, for example *airplane leases*, the IDOL Content component treats the query as a conceptual query. It *tokenizes* the query text, removes *stop words*, and *stems* the remaining tokens to produce a collection of terms.

NOTE:

These functions are language-specific. For more information about language configuration in IDOL, refer to the *IDOL Server Administration Guide*, and *IDOL Expert*.

The IDOL Content component finds documents that contains these terms, and uses statistical algorithms to quantify the relevance of the match to the original query. It assigns a relevancy score to each result document. By default, Content also uses this relevancy score to sort the results of the query.

For example:

```
action=query&text=the airplane leases
```

The IDOL Content component tokenizes this query text into *the*, *airplane*, and *leases*. The term *the* is a stop word, so Content ignores it. It stems the other two terms to produce the terms *airplan* and *leas*. Stemming allows Content to match documents that contain alternate forms of the words airplane and leases (for example, *lease*, *leasing*, and *leased* also stem to *leas*).

If `Text` contains Boolean or proximity operators, such as `AND`, or `OR`, Content treats the query as a Boolean query. As with conceptual queries, Content stems the search terms and ignores stop words, but documents match the query only if they satisfy the Boolean or proximity expression.

FieldText

The `FieldText` parameter defines restrictions on the fields of documents that match the query. `FieldText` restrictions have the following format:

```
FIELDOPERATOR{Value1[,Value2,...]}:FieldName1[:FieldName2:...]
```

where,

<code>FIELDOPERATOR</code>	is the type of comparison or test to perform on the value of the specified fields, against the specified values.
<code>Value1</code> <code>[,Value2,...]</code>	is a comma-separated list of field values.
<code>FieldName1</code> <code>[:FieldName2:...]</code>	is a colon-separated list of field names.

There are more than 40 different field operators. The following examples demonstrate the `FieldText` restriction format and some of the most common field operators. For a full list of field operators, refer to the *IDOL Server Reference*. For more information about fields and optimizations for the `FieldText` queries, see [Configure Fields, on page 104](#).

NOTE:

By default, `FieldText` matches are not case sensitive. For more information about case sensitive search, refer to *IDOL Expert*.

The following `FieldText` restriction matches documents whose `DOCUMENT_TYPE` field contains the value `record`:

```
FieldText=MATCH{record}:DOCUMENT_TYPE
```

The following `FieldText` restriction matches documents whose `FILETYPE` field contains one of the values `pdf`, `doc`, or `txt`:

```
FieldText=MATCH{pdf,doc,txt}:FILETYPE
```

The following `FieldText` restriction matches documents whose `DIR` or `DIRECTORY` field exactly matches one of `docs`, `archive`, or `work`:

```
FieldText=MATCH{docs,archive,work}:DIR:DIRECTORY
```

You can combine `FieldText` restrictions by using Boolean operators, such as `AND`, `OR`, and `NOT`, and group the restrictions by using parentheses. For example:

```
FieldText=MATCH{dog}:ANIMAL+AND+(MATCH{celery}:VEGETABLE+OR+NOT+MATCH{quartz}:MINERAL)
```

The IDOL Content component also includes field operators for numerical comparison. The following `FieldText` restriction matches documents for which `SCORE1` is greater than `.95` and `SCORE2` is less than `.85`:

```
FieldText=GREATER{.95}:SCORE1+AND+LESS{.85}:SCORE2
```

The following `FieldText` restriction matches documents that have a `COUNT` field, regardless of its value:

```
FieldText=EXISTS{}:COUNT
```

The following `FieldText` restriction matches documents whose `BRAND` field contains *tech* as a substring:

```
FieldText=STRING{tech}:BRAND
```

The following `FieldText` restriction matches documents whose `EMOTION` field contains a string that has the same stem as the word *loving*.

```
FieldText=TERM{loving}:EMOTION
```

For example, if the `EMOTION` field has any of the values *love*, *loving*, *loves*, or *lover*, the document matches.

The following `FieldText` restriction matches documents whose position (stored in the `LAT` and `LONG` fields) is within a distance of 5 kilometers from the point with latitude 52.2°N and longitude 0.12° E.

```
FieldText=DISTSPHERICAL{52.2,0.12,5}:LAT:LONG
```

The following `FieldText` restriction matches documents whose `DUE_DATE` is before the current time.

```
FieldText=LTNOW{}:DUE_DATE
```

Combine

The `Combine` parameter contains instructions on how to treat certain groups of documents as a single result for the `Query` action.

The most important value for `Combine` is `Simple`. By default, when you index long documents, the IDOL Content component splits the documents into shorter sections. When Content processes a query, it can return each of these sections as a separate search result. If you set `Combine` to `Simple`, Content returns only the best matching section from each document.

It is a best practice to set `Combine` to `Simple` in every Content Query unless there is a specific reason not to.

DatabaseMatch

Every document in the IDOL Content component belongs to exactly one database. The `DatabaseMatch` parameter specifies a comma-separated list of databases that resulting documents can come from.

By default (if `DatabaseMatch` is missing or empty), Content retrieves results from any database. Content applies database restrictions early in the query process, and these restrictions are generally quick, so using a `DatabaseMatch` can improve the performance of your query.

It is a best practice to set `DatabaseMatch` in every Content Query to be as restrictive as possible.

NOTE:

A `DatabaseMatch` restriction is equivalent to the following `FieldText` operation:

```
MATCH{DatabaseNames}:autn_database
```

Start and MaxResults

Together, `Start` and `MaxResults` control the paging of query results.

`MaxResults` indicates the total maximum number of results to retrieve from the index. By default, the IDOL Content component returns the results in order of the relevance score, which indicates how well the document matches the query. You can change the order of the results by setting the `Sort` parameter (see [Sort, on page 122](#)). Content returns the top results according to your sort, up to the value of `MaxResults`.

`Start` indicates the first result to return from the set. This option allows you to return results in pages. For example, the following query returns up to 50 results. If there are more than 50 results, it returns the best 50.

```
action=Query&Text=bananas&MaxResults=50
```

For the following query, Content finds the best 50 results that match the query, and then returns these results from the eleventh best onwards. The response therefore contains 40 results (results 11 to 50).

```
action=Query&Text=bananas&MaxResults=50&Start=11
```

MinScore

The `Minscore` parameter specifies the smallest relevancy score that a document must have to the query for the IDOL Content component to return it as a result. For example:

```
action=Query&Text=aircraft&MinScore=60
```

If this query has 10 matching documents, with scores of 100, 90, 80, 70, and so on down to 10, the query returns five results (the documents with relevancy scores of 100, 90, 80, 70, and 60).

StoreState, StateMatchID, and StateDontMatchID

The `StoreState` parameter allows you to save an IDOL Content component search. When you set `StoreState` to `True`, Content returns an alphanumeric token with the query. This token is a reference to the original set of search results.

You can set this token as the value of the `StateMatchID` parameter to restrict the results of a query to only those in the original stored search. In this way, you can use `StoreState` and `StateMatchID` together to implement 'search within a search' functionality in IDOL.

The `StateDontMatchID` parameter is the opposite of `StateMatchID`. When you set `StateDontMatchID` to the value of a stored state token, the query excludes any result that was a result in the original search. Effectively, `StoreState` and `StateDontMatchID` implement 'search outside a search' functionality.

UserMetaFields

The `UserMetaFields` parameter allows you to combine the field values in result documents into a new field (a *metafield*) dynamically at query time. This method allows you to create on-the-fly analytics for your documents. The format of `UserMetaFields` is:

```
UserMetaFields=OPERATION{UserFieldName}:DocumentFieldName
```

The following examples demonstrate both the `UserMetaFields` format, and some of the existing operators. For a full list of operators, refer to the *IDOL Server Reference* or the *IDOL Server Administration Guide*.

In the following example, the `TOTAL` metafield for a matching document is the sum of all the `PRICE` fields in that document:

```
UserMetaFields=ADD{TOTAL}:PRICE
```

You can create multiple metafields with one `UserMetaFields` parameter. In the following example, the `TOTALCOST` metafield for a matching document contains the sum of all its `COST` fields, and the `AVGREVENUE` field contains the average (mean) of the `REVENUE` fields:

```
UserMetaFields=ADD{TOTALCOST}:COST+MEAN{AVGREVENUE}:REVENUE
```

You can use metafields in the `Sort` and `FieldText` parameters of the `Query` where you create them. In the following example, the `IDOL Content` component creates a `MAXPRICE` metafield for each document that matches the initial query stages, which is the maximum of the `PRICE` fields in the document. The query returns only documents for which the `MAXPRICE` metafield is greater than 500.

```
UserMetaFields=MAX{MAXPRICE}:PRICE&FieldText=GREATER{500}:autn_user_MAXPRICE
```

Response Parameters

This section describes many of the most important parameters for the `Query` action that you can use to modify the query response.

For more information about each of these parameters, refer to the *IDOL Server Reference*.

Predict and TotalResults

The `TotalResults` parameter allows you to find out how many query results exist in the `IDOL Content` component (when this value is higher than `MaxResults`). To return this information, set `TotalResults` to `True`.

For performance reasons, when there are a large number of results, the `IDOL Content` component does not normally compute the exact number. Instead, `Content` makes a guess at the total number of

results, based on the distribution of the results that it has already found. This guess is almost always of the right order of magnitude, and is usually within 20% of the exact number.

To make Content compute the exact number of results, you can set the `Predict` parameter to `False`. However, this might incur a significant performance penalty. When you set `TotalResults` to `True` and `Predict` to `False`, Content also returns a breakdown of the available query results by database.

Print and Printfields

The `Print` and `Printfields` parameters control the information from result documents that returns. The most important values for `Print` are `Fields`, `All`, `None`, and `NoResults`.

The default value for `Print` is `Fields`. In this case, the IDOL Content component returns metadata, and either any fields in the document that are configured as `PrintType`, or the fields that you list in the `PrintFields` parameter. If you set `PrintFields`, only the specified fields return.

For example, the following query returns the `PrintType` fields of the results:

```
action=Query&Text=pillows&Print=Fields
```

The following query returns the `FIRMNESS`, `WEIGHT`, and `STUFFING` fields of the results, along with the document metadata. No other fields return.

```
action=Query&Text=pillows&Print=Fields&PrintFields=FIRMNESS,WEIGHT,STUFFING
```

When you set `Print` to `All`, Content returns all fields in the results documents (except for hidden fields).

When you set `Print` to `None`, Content prints metadata such as the reference, document ID, and date for each result, but not the document content or fields.

When you set `Print` to `NoResults`, Content does not print any information about results at all; it returns only the number of results. This option is often used with the `QuerySummary` parameter (see [QuerySummary](#), below) to return just the summary terms and phrases.

QuerySummary

The `QuerySummary` parameter allows you to extract important common words and phrases from result documents. When you set `QuerySummary` to `True`, the IDOL Content component returns a list of these summary words and phrases at the top of the response (above the results).

You can use the collection of words and phrases as query guidance, where each phrase is a potential search that you can use to refine the original query. The query summary results might also expose aspects of the result set that you might not have known about originally.

For more information about using query summaries, refer to *IDOL Expert*, and the *IDOL Server Administration Guide*.

ResponseFormat

The `ResponseFormat` parameter is available for all ACI actions. By default, all IDOL responses are in XML format. You can set `ResponseFormat` to `JSON` to return JSON, or to `config` to return results in IDOL configuration file format (a series of lines in `field=value` format).

Sort

The `Sort` parameter controls the order in which results return. There are many options for sorting, but the most important ones are `Date`, `Relevance`, and field sort. For field sort, you specify the name of the field whose value you want to sort by, and a sorting method (for example, `alphabetical`, or `numberdecreasing`).

For example:

```
Sort=relevance+date
```

This example sorts the query results by relevance. If two or more documents have the same relevance score, the IDOL Content component sorts them by the document date.

For a full list of the available sort and field sort options, refer to the *IDOL Server Reference*.

Highlight, StartTag, and EndTag

The `Highlight` parameter controls how the IDOL Content component highlights match query terms and phrases in the result documents. For example, you can set `Highlight` to `Terms` to use highlight tags to surround each instance of a query term in the content of the result documents. You can use the `StartTag` and `EndTag` parameters to specify the tags that surround matching terms.

For example:

```
action=Query&Text=quickly  
jumping&Highlight=Terms&StartTag=<strong>&EndTag=</strong>
```

If this query matches a document with the content *the quick brown fox jumps over the lazy dog*, the IDOL Content component returns the following modified content in the query response:

```
the <strong>quick</strong> brown fox <strong>jumps</strong> over the lazy dog
```

For information about other available highlight options, refer to the *IDOL Server Reference*.

Summary, Sentences, and Characters

The `Summary` parameter allows you to automatically create a summary of result documents. The main options for `Summary` are `Quick`, `Concept`, and `Context`. A *quick* summary is just the first few sentences of the result document. A *concept* summary contains sentences that are typical of the content of the result document. A *context* summary contains sentences from the document that contain words and phrases that match the query text.

You can use the `Sentences` and `Characters` parameters to control the maximum length of a summary.

Cluster

The IDOL Content component can dynamically cluster the results of a query, grouping similar results together, by including a cluster id with each result document. To turn on dynamic clustering, set `Cluster` to `True`.

Query Response

This section describes some of the response tags that a query returns.

For example:

```
action=Query&Text=spiders&Combine=Simple&Print=all&MaxResults=2
```

This query might return the following ACI response:

```
<autnresponse xmlns:autn="http://schemas.autonomy.com/aci/">
  <action>QUERY</action>
  <response>SUCCESS</response>
  <responsedata>
    <autn:numhits>2</autn:numhits>
    <autn:hit>
      <autn:reference>430377369900437504</autn:reference>
      <autn:id>128534</autn:id>
      <autn:section>0</autn:section>
      <autn:weight>85.27</autn:weight>
      <autn:links>SPIDER</autn:links>
      <autn:database>Twitter</autn:database>
      <autn:content>
        <DOCUMENT>
          <DREREFERENCE>430377369900437504</DREREFERENCE>
          <AUTN_GROUP>Connector</AUTN_GROUP>
          <AUTN_IDENTIFIER>PGLkIHM9ILRXViIgcj0iNDMw</AUTN_IDENTIFIER>
          <CREATED_AT>Mon Feb 03 16:28:55 +0000 2014</CREATED_AT>
          <DOCTRACKINGID>256a3b295232c32c4a3ad4caa5c05d3f</DOCTRACKINGID>
          <DREDBNAME>Twitter</DREDBNAME>
          <DRECONTENT>If spiders can sit on the web all day then so can
I</DRECONTENT>
        </DOCUMENT>
      </autn:content>
    </autn:hit>
    <autn:hit>
      <autn:reference>430369534923374592</autn:reference>
      <autn:id>22023</autn:id>
      <autn:section>0</autn:section>
      <autn:weight>85.27</autn:weight>
      <autn:links>SPIDER</autn:links>
      <autn:database>Twitter</autn:database>
      <autn:content>
        <DOCUMENT>
          <DREREFERENCE>430369534923374592</DREREFERENCE>
          <AUTN_GROUP>Connector</AUTN_GROUP>
          <AUTN_IDENTIFIER>PGLkIHM9ILRXViIgcj0iNDM</AUTN_IDENTIFIER>
          <CREATED_AT>Mon Feb 03 15:57:47 +0000 2014</CREATED_AT>
          <DOCTRACKINGID>2a6e7bbc51001b815efa9b30acf6704e</DOCTRACKINGID>
```

```

    <DREDBNAME>Twitter</DREDBNAME>
    <DRECONTENT>eek! spiders!</DRECONTENT>
  </DOCUMENT>
</autn:content>
</autn:hit>
</respondedata>
</autnresponse>

```

For a query action, the <respondedata> element contains the <autn:numhits> tag, which contains the number of results returned, and a series of <autn:hit> tags, which each contain a query result.

A query result consists of metadata for the result document, and the returned portions of the result document under the <autn:content> tag (in this example, each document is in a <DOCUMENT> subtree).

The important result document metadata tags are described in the following table.

<autn:reference>	The identifier for the document in the IDOL index.
<autn:id>	A unique number assigned to the document by the Content component where the document is indexed. The ID counts up from the first document indexed, and is unique only in a single Content component. In a distributed index, ID is not a reliable way to refer to a document in the index. As a best practice, always use references to refer to documents in IDOL Server.
<autn:section>	The section number of the result.
<autn:weight>	The relevancy score of the document to the query.
<autn:links>	The matching terms from the query that were found in the document.
<autn:database>	The database the result document came from.

The <DOCUMENT> subtree consists of field values for the document. The tag is the name of the document field, and the value of the tag is the value of the field in the document. For documents that you indexed in IDX format, this subtree is flat. For documents that you indexed in XML format, this subtree might have additional field structure, corresponding to the original XML.

Commonly, the reference for a document is stored in a field called DRREFERENCE, the content of a document is stored in a field called DRECONTENT, and the database name is stored in a field called DREDBNAME.

Glossary

A

ACI (Autonomy Content Infrastructure)

A technology layer that automates operations on unstructured information for cross-enterprise applications. ACI enables an automated and compatible business-to-business, peer-to-peer infrastructure. The ACI allows enterprise applications to understand and process content that exists in unstructured formats, such as email, Web pages, Microsoft Office documents, and IBM Notes.

ACI API

The programming interface for interacting with IDOL ACI Servers. Several IDOL SDKs are available to allow you to develop applications with the ACI API.

ACI Server

A server component that runs on the Autonomy Content Infrastructure (ACI).

ACL (access control list)

An ACL is metadata associated with a document that defines which users and groups are permitted to access the document.

action

A request sent to an ACI server.

active directory

A domain controller for the Microsoft Windows operating system, which uses LDAP to authenticate users and computers on a network.

Agentstore

An IDOL component that indexes, collects, manipulates, and stores agent and category information. Agentstore is a special configuration of the Content component.

C

Category component

The IDOL Server component that manages categorization and clustering.

Community component

The IDOL Server component that manages users and communities.

connector

An IDOL component (for example File System Connector) that retrieves information from a local or remote repository (for example, a file system, database, or Web site).

Connector Framework Server (CFS)

Connector Framework Server processes the information that is retrieved by connectors. Connector Framework Server uses KeyView to extract document content and metadata from over 1,000 different file types. When the information has been processed, it is sent to an IDOL Server or Distributed Index Handler (DIH).

Content component

The IDOL Server component that manages the data index and performs most of the search and retrieval operations from the index.

Controller

An IDOL component that monitors IDOL services on a host machine, and communicates with Coordinator to report status information. This component is primarily used by the IDOL Site Admin administrative front-end application.

Coordinator

An IDOL component that manages status information from Controller components and acts as a central point for viewing logs and monitoring the status of the IDOL system. This component is primarily used by the IDOL Site Admin administrative front-end application.

D

DAH (Distributed Action Handler)

DAH distributes actions to multiple copies of IDOL Server or a component. It allows you to use failover, load balancing, or distributed content.

DIH (Distributed Index Handler)

DIH allows you to efficiently split and index extremely large quantities of data into multiple copies of IDOL Server or the Content component. DIH allows you to create a scalable solution that delivers high performance and high availability. It provides a flexible way to batch, route, and categorize the indexing of internal and external content into IDOL Server.

F

fields

The parts of a document that store different sections of the content.

Find

A basic search user interface for IDOL.

G

graph

A set of relationships that describes the connections in a set of data. The IDOL Knowledge Graph component uses graphs to explore connections in IDOL documents.

IDOL

The Intelligent Data Operating Layer (IDOL) Server, which integrates unstructured, semi-structured and structured information from multiple repositories through an understanding of the content. It delivers a real-time environment in which operations across applications and content are automated.

IDOL Proxy component

An IDOL Server component that accepts incoming actions and distributes them to the appropriate subcomponent. IDOL Proxy also performs some maintenance operations to make sure that the subcomponents are running, and to start and stop them when necessary.

IDOL SDK

Software Development Kits that allow you to develop applications that use the ACI API to use IDOL component ACI Servers.

IDOL Speech Server

An IDOL component that processes audio files, performing tasks such as speech-to-text transcription, speaker identification, and audio search.

index

The store of content in IDOL Server.

indexing

The process of adding content to IDOL Server.

Intellectual Asset Protection System (IAS)

An integrated security solution to protect your data. At the front end, authentication checks that users are allowed to access the system that contains the result data. At the back end, entitlement checking and authentication

combine to ensure that query results contain only documents that the user is allowed to see, from repositories that the user has permission to access. For more information, refer to the IDOL Document Security Administration Guide.

K

KeyView

The IDOL component that extracts data, including text, metadata, and subfiles from over 1,000 different file types. KeyView can also convert documents to HTML format for viewing in a Web browser.

Knowledge Graph

An IDOL component that uses connections between IDOL documents to create a set of relationships called a graph, and allows you to explore these connections in your data.

L

LDAP

Lightweight Directory Access Protocol. Applications can use LDAP to retrieve information from a server. LDAP is used for directory services (such as corporate email and telephone directories) and user authentication. See also: active directory, primary domain controller.

License Server

License Server enables you to license and run multiple IDOL solutions. You must have a License Server on a machine with a known, static IP address.

M

mapped security

A method for configuring document security to ensure that only permitted users can

access a document. In mapped security, IDOL compares user security details against an access control list (ACL), which is stored in the document.

Media Server

An IDOL component that analyzes video files and streams to extract information about their content. Media Server can run analysis operations such as face recognition, speech-to-text, and number plate recognition.

O

OmniGroupServer (OGS)

A server that manages access permissions for your users. It communicates with your repositories and IDOL Server to apply access permissions to documents.

P

primary domain controller

A server computer in a Microsoft Windows domain that controls various computer resources. See also: active directory, LDAP.

Q

QMS (Query Manipulation Server)

An IDOL component that manages promotions, modifies queries to IDOL Server, and manipulates results from IDOL Server.

query

A search to retrieve documents from IDOL Server.

S

stemming

The process of extracting the morphological root of a word. In languages, some words

have a common morphological root. HPE IDOL provides stemming algorithms that reduce words to this form. This process allows IDOL Server to match concepts regardless of the grammatical use of words. In English, for example, the words 'help', 'helpful', 'helping', and 'helped' all reduce to their stem 'help' without significant loss of meaning. HPE IDOL provides as standard a set of stemming algorithms for the most commonly used languages. IDOL Server applies stemming after stop words have been discarded, both at index time (when content is stored in IDOL Server), and at query time (query text is stopped and stemmed before it is matched).

stop word

A very common word that occurs too frequently to be useful for searching. Stop words include articles (for example, the) and prepositions (for example, to or from). Stop words are language-specific. You can use a stop word list in IDOL Server to allow it to discard these words at index and query time to save index space and improve retrieval performance.

T

tokenization

The process of splitting a query string up into individual search terms.

U

unmapped security

A method for configuring document security to ensure that only permitted users can access a document. In unmapped security, IDOL connects to the original repository to check the user permissions for a document. This method is usually a lot slower than mapped security.

V

View

An IDOL component that converts files in a repository to HTML formats for viewing in a Web browser.

W

Wildcard

A character that stands in for any character or group of characters in a query.

X

XML

Extensible Markup Language. XML is a language that defines the different attributes of document content in a format that can be read by humans and machines. In IDOL Server, you can index documents in XML format. IDOL Server also returns action responses in XML format.

Send documentation feedback

If you have comments about this document, you can [contact the documentation team](#) by email. If an email client is configured on this system, click the link above and an email window opens with the following information in the subject line:

Feedback on Getting Started Guide (IDOL 11.5)

Add your feedback to the email and click **Send**.

If no email client is available, copy the information above to a new message in a web mail client, and send your feedback to AutonomyTPFeedback@hpe.com.

We appreciate your feedback!