

Oracle Commerce Guided Search

Glossary

Version 11.1 • July 2014



Copyright and disclaimer

Copyright © 2003, 2014, Oracle and/or its affiliates. All rights reserved.

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

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

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

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

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

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

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

This software or hardware and documentation may provide access to or information on content, products, and services from third parties. Oracle Corporation and its affiliates are not responsible for and expressly disclaim all warranties of any kind with respect to third-party content, products, and services. Oracle Corporation and its affiliates will not be responsible for any loss, costs, or damages incurred due to your access to or use of third-party content, products, or services.

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

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

A

Access Control List

See [Endeca Access Control List](#).

Access Control System

See [Endeca Access Control System](#).

Analytics API

See [Endeca Analytics API](#).

ancestors

The dimension values between the navigation descriptor (the current location in the dimension tree) and the dimension root (the beginning of the dimension tree). Although a parent is technically an ancestor, in general the term is used for those values above the parent in the dimension tree. See also [child dimension value](#) and [parent dimension value](#).

aggregated Endeca record

Multiple Endeca records grouped by a dimension or property and treated as a single record. Aggregated Endeca records are commonly used to minimize the effect of duplication in the data display. See [Endeca record](#).

aggregated record query

An aggregated record query is a query used to obtain a single aggregated record from the MDEX Engine.

Application Controller

See [Endeca Application Controller](#).

Assembler

The Oracle Commerce Assembler enables a Web application to query external data sources based on a user's navigation state or other triggers. The Assembler returns content item objects that encapsulate page configuration specified by a content administrator. All the content for a page, including the results of any additional queries needed for spotlighting or merchandising, are wrapped in the content item object, simplifying the logic in the front-end application by reducing the need to manage sub-queries in the application layer.

The Assembler distribution includes a Guided Navigation JAR file that provides support for querying an MDEX Engine, although users may choose to extend the Assembler to interact with other required data sources.

B

base SSL

See SSL.

baseline update

A baseline update (also called a full update) produces a complete re-indexing of the entire data set. It runs the update process for the whole data set.

In your baseline update pipeline, you can add, change or remove records, dimensions, dimension values and properties. In addition, configuration changes, such as dimension reordering or stop word changes require a baseline update.

You can run baseline updates nightly and use this method as your update strategy. For small to medium-sized data sets, baselines can be run frequently, as often as every few minutes. Alternatively, you can run as many partial updates as needed for those changes that can be done through partial updates, and periodically run baseline updates for those changes that require a baseline update.

C

cartridge

A cartridge is a functional component that a content administrator can place on a dynamic page using Experience Manager. Examples of cartridges may include a Guided Navigation cartridge, results list cartridge, or rotating banner cartridge. A cartridge is comprised of several parts:

- A configuration file (XML template) that defines the content structure and the editing interface in Experience Manager
- Zero or more Experience Manager editors that allow the content administrator to configure cartridge content
- One or more Content Assembler tag handlers to do additional query processing before returning results to the application.
- Rendering code to display the content in the Web application

CAS Console for Oracle Commerce Workbench

The CAS Console for Oracle Commerce Workbench is a Web-based application used to crawl various data sources including file systems, content management systems, and custom data source extensions. During the Content Acquisition System installation, the CAS Console is installed as an extension to Oracle Commerce Workbench.

CAS Server

A component of the Content Acquisition System that gathers source data by managing file system and CMS crawls. The output of a crawl is a set of Endeca records that are used in a Developer Studio project.

CAS Service

The CAS Service is a servlet container that runs the CAS Server, the Component Instance Manager, and any number of Record Store instances (one per data source).

certificate

A security file that is used by all Guided Search clients and servers to specify their identity when using SSL. The certificate, which is created by the Endeca enecerts utility, contains information about the certificate owner and a public key used to encrypt and decrypt data that is being exchanged between SSL-enabled components. The certificate file should be thought of as the identity of the Guided Search system, or as the identify of all components of the Guided Search system.

certificate authority file

A security file that is used in a Guided Search system that is configured for mutually authenticated SSL. The CA file, which is created by the Endeca `enecerts` utility, is used by all Guided Search clients and servers to authenticate the other endpoint of a communication channel.

Charting API

See Endeca Charting API.

child dimension value

Sub-dimension values are known as child dimension values. Child dimension values are always more specialized than their parents and help the users to further refine their navigation query. A dimension value may have multiple child values. Because these values refine query results, child dimension values are often referred to as refinement values or refinements. See also parent dimension value.

Cluster Discovery

See Endeca Cluster Discovery.

collapsible dimension hierarchy

A dimension hierarchy in which some or all of the internal (non-root and non-leaf) dimension values are flagged as collapsible. The MDEX Engine automatically removes, or collapses, these dimension values when there are only a few leaves available for refinement, creating a more streamlined, user-friendly navigation experience for your users.

Component Instance Manager

The Component Instance Manager manages Record Store instances. Management means that the Component Instance Manager is responsible for creating, deleting, and listing all Record Store instances.

compound dimension query

A combination of one or more dimension values (or, more specifically, dimension value IDs). A compound dimension query instructs the MDEX Engine to return the set of records that represents the intersection of all the dimension values that it contains.

Content Acquisition System

The Content Acquisition System is a set of components that add, configure, and crawl data sources for use in a Guided Search application. Data sources include file systems, Content Management System crawls, Web servers, and custom data sources. The Content Acquisition System crawls data sources, converts documents and files to Endeca records and stores them for use in a Guided Search pipeline.

D

Data Foundry

See Information Transformation Layer.

data pipeline

The data pipeline, as displayed in Developer Studio, is the graphical representation of how the Endeca Information Transformation Layer transforms source data into tagged Endeca records.

data pipeline component

The pipeline is composed of a collection of components. Each component performs a specific function during the transformation of your source data into Endeca records, such as loading data, standardizing properties, and tagging records with dimension values. Components are linked together, by means of cross-references, providing a sequential flow and a pipeline feel.

dead end query

When an intersection of common records doesn't exist between all of the dimension values in a navigation query, that query is considered a dead end. The MDEX Engine automatically eliminates such dead end queries in the way it structures the follow-on query information that it returns in its query results. See also Guided Navigation.

delta update

A full baseline index, similar to the baseline update, that is performed by joining a smaller extraction of source data (only the added, updated, or removed source records) with the output from the previous baseline update. A delta update is ideal when you need to reduce the time required for loading the source data repository and for extracting the data.

See also partial update and baseline update.

Deployment Template

The Deployment Template is a utility that you run to create a new Endeca application with the complete directory structure required for deployment, including Endeca Application Controller (EAC) control scripts, configuration files, and batch files or shell scripts that wrap common script functionality.

deprecated

Deprecated means that an item (platform, feature, flag, etc.) is fully supported, as in prior releases, but that Oracle intends to discontinue support in a subsequent release.

descriptor

See dimension descriptor and navigation descriptor.

Developer Studio

A Windows application used to define all aspects of the instance configuration including pipeline components, Endeca properties and dimensions, precedence rules, dynamic business rules, and user profiles. Developer Studio uses a project file, with an .esp extension, that contains pointers to the XML files that support an instance configuration. See also Oracle Commerce Workbench.

Dgidx

The name of the program that indexes tagged Endeca records and produces indices in the MDEX Engine format.

Dgraph

The name of the program that launches the MDEX Engine. You can run more than one Dgraph process on a single computer.

Dgraph Cluster

Dgraph clusters are used in the configuration of an application installed by the Deployment Template utility to apply actions to an entire cluster of Dgraphs, rather than manually iterating over a number of Dgraphs. They also contain logic associated with Dgraph restart strategies.

differential crawl

See incremental crawl.

dimension

A collection of related dimension values, organized into a navigable tree with a dimension root at the top. See also externally created dimension, and externally managed dimension.

dimension adapter

A pipeline component that loads dimension information. Dimension adapters can load dimension information in either XML or delimited format.

dimension descriptor

The dimension that has been used in a navigation query to generate a set of query results. Every descriptor dimension has one or more descriptors (dimension values) associated with it. A descriptor dimension provides a handle for accessing the dimension value(s) used to refine the index records to create the current query results. See also navigation descriptor.

dimension group

An implicit or explicit grouping of dimensions used to organize dimensions. An implicit group consisting of a single dimension is automatically generated for each dimension that is not included in an explicit, user-defined grouping (which generally contains multiple dimensions).

dimension hierarchy

A logical (though not physical) tree structure with a dimension root at the top and related dimension values below. Dimension hierarchies allow you to exercise a higher level of control over the number of follow-on choices that are presented to users as they navigate. See also parent dimension value and child dimension value.

dimension root

The first dimension value in a dimension tree. A dimension root generally has the same name and ID as its first child dimension.

dimension search

A search that finds all of the dimension values that have names that contain terms the user provides. The result of a dimension search is a set of dimension values, organized by dimension.

dimension server

Working in conjunction with dimension adapters, dimension servers serve as a centralized source of dimension information for all other pipeline components.

dimension value

Member of a dimension; used as a tag, or label, to classify a record in your data set. Tagging a record with a dimension value identifies that record as a valid result when a user queries for the dimension value. Dimension values can be auto-generated, defined explicitly, or imported from external taxonomy systems, for example Stratify.

dynamic ranking

The ordering of dimension values according to how frequently they appear within a data set, rather than in the MDEX Engine's default alphabetical order. Dynamic ranking is particularly useful in large data sets.

E

Endeca

Older name for Guided Search. The term Endeca is used in many places in the software, for example in installation paths and APIs, and it is also used in some areas of the documentation. It is synonymous with Guided Search.

Endeca Access Control List

A property on an Endeca record that lists user and group permissions, such as the right to view the record. In an implementation using the Endeca Access Control System, the MDEX Engine reads the Endeca ACL property and uses the user entitlement filter to determine if the user that made the query is authorized to view the record.

Endeca Access Control System

A security infrastructure that authenticates a user's identity against an external directory (such as an LDAP directory) and creates a user entitlement filter that limits access to only those records that the user is authorized to see. One or more login modules are configured to perform the actual authentication. See also file-based login module, LDAP login module, and user entitlement filter.

Endeca Analytics API

Available with Endeca Analytics, the Analytics API extends the Endeca Presentation API to enable interactive applications that allow users to explore aggregate and statistical views of information using a Guided Navigation interface. See also Endeca Presentation API.

Endeca Analytics module

Endeca Analytics builds on the core capabilities of the Endeca MDEX Engine to enable applications that examine aggregate information such as trends, statistics, analytical visualizations, comparisons, and so on, all within the Guided Navigation interface.

Endeca APIs

The collection of APIs, such as the Presentation API, Logging API, and Forge API, that provide an interface to a Guided Search implementation.

Endeca Application Controller

The Endeca Application Controller is the interface you use to control, manage, and monitor your Guided Search implementations. It provides the infrastructure to support Guided Search projects from design through deployment and runtime. It replaces the Control Interpreter, while leaving the Endeca tools (Developer Studio, Oracle Commerce Workbench) largely intact.

Endeca CAS Extension API

The CAS Extension API provides interfaces and classes to build extensions to the Content Acquisition System such as data sources and manipulators. An extension developer packages extensions into a JAR, and a CAS application developer installs the JAR and any additional JARs (for third-party dependencies) into the Content Acquisition System. After installation, the extensions are available and configurable using the CAS Console, the CAS Server API, and the CAS Server Command-line Utility.

Endeca Charting API

Available with Endeca Analytics, the Charting API extends the Endeca Presentation API to support graphical visualization of Endeca analytics results. See also Endeca Presentation API.

Endeca Cluster Discovery

A feature of the MDEX Engine that uses the salient terms created by Term Discovery and intelligently clusters them into groupings that represent distinct and coherent sets of records. These clusters are similarly leveraged as intuitive ways for end users to refine or broaden their Guided Search queries.

Endeca Data Foundry

The component of the Information Transformation Layer that aggregates information from disparate sources, transforms it into Endeca records, and produces indices for use by the MDEX Engine.

Endeca HTTP Service

The Endeca HTTP Service is an application server that runs the Endeca Application Controller.

Endeca Information Transformation Layer

The Endeca Information Transformation Layer (ITL) is the component of Oracle Commerce Guided Search that reads in your source data and manipulates it into an index for the MDEX Engine. It consists of the Content Acquisition

System, Forge, and the indexing processes. The ITL components are offline processes that you run at intervals that are appropriate for your business needs.

Endeca Logging and Reporting System

The Endeca Logging and Reporting System provides an application-level logging solution, using a stand-alone logging server. See also Logging API, Log Server and Report Generator.

Endeca Presentation API

An interface between client browsers and the MDEX Engine. The Endeca API must be available online, on a Web or application server, as long as you want clients to be able to access data.



Note: The Assembler API, rather than the Presentation API, is recommended for new front-end application development.

See also Endeca APIs.

Endeca record

The individual items that the user is trying to navigate to in a Guided Search application. Endeca records generally correspond to traditional records in a source database. Unlike source records, however, Endeca records have been standardized for consistency, and then classified with dimension values. The production of Endeca records is a primary result of the pipeline development process.

Endeca Record Store

The Endeca Record Store provides persistent storage for generations of records that can later be accessed by Forge for baseline and incremental updates. The CAS Server writes crawl output from each data source to a unique Record Store instance.

Endeca Relationship Discovery

An optional feature of Oracle Commerce Guided Search that includes Endeca Term Discovery and Endeca Cluster Discovery. It enables users to discover relationships in their source data records by extracting salient terms from records and using them to build sets (clusters) of related records.

Endeca Term Discovery

A feature of the Information Transformation Layer that identifies and extracts salient noun phrases (terms) from documents, and subsequently tags them to the associated records. The terms, which are identified through natural

language processing and statistical inference techniques, constitute words or groups of words that are significant in both the corpus and in individual records. These terms, when presented in the application's front end, allow users to refine and broaden their queries.

Endeca Tools Service

The Endeca Tools Service is an application server that runs Oracle Commerce Workbench.

Endeca Web Crawler

A component of the Content Acquisition System that gathers source data by crawling HTTP and HTTPS Web sites and outputs Endeca records for use in a Developer Studio project.

Experience Manager Editor SDK

The Experience Manager Editor SDK enables application developers to introduce new functionality into Experience Manager via custom editors. The SDK consists of the Experience Manager Editor API, a sample editor project, and associated documentation.

Experience Manager

The Experience Manager is a powerful template-based tool that enables rapid creation of rich, dynamic landing pages. Experience Manager gives merchandisers and content administrators unprecedented control over site content without the need for IT intervention.

expression

Analogous to functions, expressions tell the Forge program which records, properties, or dimensions to affect, and how to affect them. See also Forge.

externally created dimension

A logical hierarchy of a data set that is transformed from its source format to Guided Search compatible XML outside of Developer Studio. The logical hierarchy of the dimension conforms to Guided Search's external interface for describing a data hierarchy before you import the dimension into your project.

externally managed dimension

A logical hierarchy for a data set that is built and managed using a third-party tool. Once you include an externally managed taxonomy in your Developer Studio project, it becomes a dimension whose hierarchy is managed by the third-party tool that created it. In Developer Studio, you cannot add or remove dimension values from it. If you want to modify a dimension or its dimension values, you have to edit the taxonomy using the third-party tool and then update the taxonomy in your project.

F

file-based login module

A component of the Endeca Access Control System that authenticates a user's identity and group membership against information in a local directory file.

filter

Query settings that limit what the query can return. See also dimension search, record search, and user entitlement filter.

flat dimension

A dimension that has only one level of hierarchy underneath the dimension root. See also dimension hierarchy.

Forge

A component of the Data Foundry that processes your source data records or documents into tagged Endeca records.

full crawl

A crawl in which CAS retrieves all the source documents that it is configured to access. See also incremental crawl.

full update

See baseline update.

G

Guided Navigation

Guided Navigation is the presentation of valid follow-on refinement choices to the user. You can also think of Guided Navigation as the elimination of invalid refinement queries, or dead ends.

Guided Search application

The end-user-facing portion of a Guided Search implementation.

Guided Search implementation

An entire Guided Search system, comprising the Guided Search instance (or back end), the Guided Search application (or front end), and any machines included in the resource collection. See also Guided Search instance, Guided Search application, and resource collection.

Guided Search instance

The back end of a Guided Search implementation, which is not visible to the end user. This typically includes Developer Studio project files and system configuration files.

Guided Search system

See Guided Search implementation.

H

hidden dimension

A dimension that is not accessible from a Guided Search application. If a dimension is marked as hidden on the Dimension editor of Developer Studio, the MDEX Engine will not return the dimension or any of its values as refinement options. Although hidden dimensions are not rendered in the user interface, records are still tagged with relevant values from these dimensions. Therefore, a user is able to search for records based on dimension values within hidden dimensions.

hierarchical dimension

See dimension hierarchy.

I

incremental crawl

A crawl in which CAS retrieves only the source documents that have changed since the last crawl. Incremental crawls are also known as differential crawls. See also full crawl.

indexer adapter

A pipeline component that saves data that is ready to be indexed.

inert dimension value

A dimension value that is not included in the navigation state. When a user selects an inert dimension value, the navigation state is not changed, but the children of the dimension value are displayed for selection. See also hidden dimension.

Information Transformation Layer

See Endeca Information Transformation Layer.

instance configuration

A set of project files that configure the back-end processes (Forge, Dgidx, Dgraph) of a Guided Search implementation.

intermediate dimension value

A non-navigable dimension value that changes the query parameters but does not affect the current Endeca record set. For example, A-D, E-H, I-L, and so forth.

J

Java manipulator

A pipeline component that you can write in Java and add to your pipeline, to manipulate records. A Java manipulator that you create must adhere to the Content Adapter Development Kit (CADK) Java Adapter interface classes. A Java manipulator transforms source records and Endeca records as part of data processing in the Endeca Information Transformation Layer. For example, Java manipulators are used by the Endeca Term Discovery feature to extract terms from source records, provide filtering and scoring mechanisms, and tag the terms to the associated Endeca records.

Join

A join combines records from two or more tables in a relational database. In the Structured Query Language (SQL), there are three types of joins: inner, outer, and cross. Outer joins are subdivided further into left outer joins, right outer joins, and full outer joins.

K

key properties

Available with Endeca Analytics, key properties are property- and dimension-level metadata that allow customized application behavior.

keyword search

A query that returns results based on a user-specified text string (keyword). See also record search and dimension search.

L

LDAP login module

A component of the Endeca Access Control System that authenticates a user's identity and group membership against a lightweight directory access protocol (LDAP) directory.

leaf dimension value

The bottom-most dimension value in a dimension tree. A leaf dimension value has no children. See also child dimension value.

Log Server

The Log Server translates log requests into log files. These log files can be used by the Report Generator to create HTML-based, human-readable log reports. Log requests are passed from the application modules to the Log Server via the Logging API. See also Logging API and Report Generator.

Logging API

The Logging API receives log requests from a Guided Search application's modules, and passes them to the Log Server via HTTP.

M

MDEX Engine

The query engine that is the core of Oracle Commerce Guided Search. The MDEX Engine consists of the indexer (Dgidx) and the Dgraph. The MDEX Engine stores the indices generated by the Data Foundry. After the indices are stored, the MDEX Engine receives queries, executes them against the stored indices, and returns the results. The MDEX Engine is an online process that must remain running as long as you want clients to be able to access data. See also Endeca APIs.

MDEX Engine query request

A client browser request that has been altered to use MDEX Engine-specific parameters. Not to be confused with navigation query, which is a specific type of query. See also navigation query.

mutually authenticated SSL

See SSL.

N

navigation descriptor

The specific dimension value for a dimension that is specified in an MDEX Engine query. The navigation descriptor describes the user's current dimension value choice within a dimension; this is the user's current location within that dimension's tree. See also dimension descriptor.

navigation query

A query that returns a set of records based on user-selected characteristics along with any follow-on information. See also record query.

navigation reference

A collection of dimension values that can be used to create a navigation query. Navigation references, in essence, are navigation queries waiting to happen.

navigation state search query

A query to obtain the set of valid navigation references whose constituent dimension values have names that match all of the search terms.

O

Oracle Commerce Guided Search

The platform upon which Guided Navigation solutions are built. Oracle Commerce Guided Search is made up of the Information Transformation Layer, the MDEX Engine, and the Assembler.

Note that in earlier versions of the software, Oracle Commerce Guided Search was called Oracle Endeca Commerce. The latter term is used in some areas of the documentation. The terms are synonymous.

P

parent dimension value

A dimension value that has sub-dimension values. Each dimension value can have only one parent value. See also child dimension value and ancestors.

partial update

A partial update is a change in the overall data set that does not require restarting the MDEX Engine. Partial updates allow you to update only those portions of the MDEX Engine index that have changed since the last baseline update.

A partial update lets you implement a number of the source data changes. For project configuration changes, run a baseline update.

Even if you are only making source data changes, keep in mind that some configuration information that is derived from the data, such as dictionary or wildcarding information, can become outdated. Therefore, to keep dictionaries up-to-date, periodically run baseline updates.

Perl manipulator

A pipeline component that uses Perl to efficiently manipulate source records and Endeca records as part of data processing performed in the Information Transformation Layer. The Forge API for Perl provides the means to plug a perl manipulator into Forge with four plug-in methods that you write. From the methods that you write, you can also call methods in the EDF name space that Endeca provides to perform data manipulation.

pipeline

See data pipeline.

precedence rule

A trigger that causes dimension values that were not previously displayed to become available. Precedence rules reduce information overload, because the user is not presented with too many, or inappropriate, navigation choices.

Presentation API

See Endeca Presentation API.

property

A name-value pair (a combination of property name and property value) containing information about a record. Generally, a source record is nothing more than a set of properties. Properties are intended for display once the end user has searched or navigated to a record set or an individual record. Properties can also be used to automatically generate dimensions and dimension values.

property mapper

A pipeline component used to specify what your Guided Search application should do with individual properties in your source data: explicitly map them to existing Endeca dimensions or properties, create new Endeca dimensions or properties to map them to, or ignore them.

property name

Property names are mapped to dimensions. See [property](#).

property value

Property values are mapped to dimension values. See [property](#).

R

range dimension

Dimension that can be organized into discrete ranges, for example, \$10-\$20, \$21-\$50, \$51-\$100, and so forth. Range dimensions are useful for data, such as price, that should be navigated as discrete values.

rapid update

See partial update.

raw data

See source record.

record adapter

A pipeline component used to read source data into the Endeca Data Foundry.

record assembler

A pipeline component used to join source records originating from different files.

record cache

A pipeline component that stores a temporary copy of record data that has been read in by a record adapter. Record caches are generally used in conjunction with record assemblers and are set up to contain data from secondary data sources.

record manipulator

A pipeline component that changes the data that is associated with a record. This can include changing properties, property values, and dimension value tagging.

record page

The displayed representation of a single Endeca record. Most Guided Search-enabled applications are designed so that when an end-user selects an item (that is, an Endeca record) in their navigation or search results page, a record page appears displaying all or most of the record's properties. See also Endeca record.

record query

A query to obtain a single specific Endeca record from the MDEX Engine.

Record Relationship Navigation

Record Relationship Navigation is an extension to Endeca Query Language for relational (normalized) data that enables more complex Guided Navigation at larger scale. Record Relationship Navigation is an optional module that is intended for use with complex relational data. It allows users to logically connect records at query time and navigate based on the connected relationships. It is based on EQL. Interrelationships between logical entities are not only preserved by the indexing process, but also allow for simultaneously navigating different types of entities based on attributes of any subsets of the data.

record search

A keyword search that filters the record set to include only those Endeca records that have at least one property or dimension whose value matches a specified search term (keyword). The result of a record search is a set of records based on the user-defined keyword(s), plus any follow-on query information. See also dimension search.

reference implementation

A sample Guided Search web application that provides skeleton examples of typical pages (navigation, record, and aggregated record pages) and the components that make up these pages (for example, navigation controls, navigation descriptors, and a record set). It is intended only as a guide for creating MDEX Engine queries and building pages from the query results. See also Guided Search application.

refinement dimension

A dimension whose dimension values may be used to reduce or refine the current query's record set.

refinement dimension values

The next set of dimension value choices the user can make to refine his or her navigation query. You refine a navigation query by navigating from a dimension value to one of its children. See also child dimension value.

Relationship Discovery

See Endeca Relationship Discovery.

relevance ranking module

A search interface feature that lets the developer control the order in which record or dimension search results are displayed to the end user. A relevance ranking module assigns ranking scores to results based on its predetermined criteria, such as the frequency of a user's query terms in the result text. Modules can be combined to produce a complex ranking strategy for a search interface.

Report Generator

The Report Generator uses the log files created by the Log Server to generate HTML-based reports.

resident set size (RSS)

Resident set size (RSS) is the amount of physical memory currently allocated and used by the MDEX Engine process. As the MDEX Engine process runs, the active executable code and data are brought into RAM, becoming part of the RSS for the MDEX Engine.

For more information, see the *Oracle Commerce MDEX Engine Performance Tuning Guide*.

resource collection

Represents the set of resources used to host a Guided Search application, which generally consists of a set of computers with various application-specific roles (such as data processing, MDEX Engine servers, log collection and analysis, and so forth).

resource configuration

Specifies the configuration of the resources on which the application will run (such as machines in the resource pool, number of segments, number of replicas, and so on).

results page

A page summarizing multiple records returned as a result of a navigation or search query.

rollup key

The property or dimension name by which an aggregated Endeca record is consolidated. Records that have the same value for the property or dimension are rolled up into an aggregated record.

S

Search Engine Optimization Module

See URL Optimization API.

search interface

A named collection of properties and/or dimensions, each of which is enabled for record search. The search interface may include features that control the search behavior, such as relevance ranking modules and partial match. Search interfaces thus allow end users to search multiple properties and/or dimensions simultaneously.

search query

See dimension search, record search, and navigation state search query.

source record

The data that is input into the Guided Search system. Guided Search supports source records in a variety of formats.

SSL

The Secure Sockets Layer protocol that protects the privacy and integrity of data being transferred across a network. Base SSL provides security by encrypting communications among Guided Search components. In addition to using encryption, mutually authenticated SSL provides a higher level of security by forcing each component to use a certificate to confirm its identity to other components. See also certificate and certificate authority file.

stacked authentication

A feature of the Endeca Access Control System that allows multiple login modules to be used to authenticate a user's identity. See also file-based login module and LDAP login module.

static ranking

The order in which dimensions and dimension values appear in the user interface, overriding MDEX Engine defaults. See also dynamic ranking.

style

See dynamic business rule.

supplemental object

Optional object that accompanies a basic navigation query. Supplemental objects are the general data structure used to create merchandising or analytics displays. Each supplemental object may contain one or more navigation references, one or more records, and one or more object attributes.

synonym

Alternative terms used for the display and/or classification of dimension values. All dimension value definitions contain one or more synonyms. See also dimension value.

T

Term Discovery

See Endeca Term Discovery.

Tools and Frameworks

The Oracle Commerce Tools and Frameworks enable the development of Web applications that present dynamic and contextual content to end customers across multiple channels. The main components of Tools and Frameworks are:

- Oracle Commerce Workbench with Experience Manager, a Web-based tool that enables business users to create and manage dynamic content configurations across multiple channels or applications.
- Oracle Commerce Assembler, a system for collecting and sending queries to the MDEX Engine and returning the resulting data to an application as ready-to-render content items.
- The Discover Electronics reference application, an application that demonstrates best practices for Guided Search applications.

U

user entitlement filter

A record filter that is created by the Endeca Access Control System for an authenticated user. The filter defines the user's access rights to the data in the Guided Search implementation and is used by the MDEX Engine during queries to return only those records that the user is authorized to see.

URL Optimization API

The URL Optimization API enables you to create application URLs that are optimized for internet search engines. In particular, the API provides the capability to shorten and canonicalize URLs and add search and navigation keywords to URLs. The resulting URLs are more optimized for internet search engines and more understandable to front-end application users.



Note: This API is sometimes referred to as the Search Engine Optimization (SEO) Module.

V

virtual process size

The *virtual process size (or address space)* for the Dgraph is the total amount of virtual memory allocated by the operating system to the MDEX Engine process at any point in time. This includes the Dgraph code, the MDEX Engine data as represented on disk, the Dgraph cache and any temporary work space.

For more information, see the *Oracle Commerce MDEX Engine Performance Tuning Guide*.

See also resident set size (RSS) and working set size (WSS) of the Dgraph.

W

Workbench

Oracle Commerce Workbench is a Web-based tool that provides a way for business users and merchandisers to configure portions of their Guided Search application and provides system administrators with a means to configure and administer a Guided Search implementation.

Unlike Developer Studio, which provides a rich development environment for configuring all aspects of a Guided Search implementation, Workbench focuses on a smaller set of common, every day configuration and maintenance tasks. This reduced focus gives Workbench a smaller footprint that can exist within the bounds of a Web-based application.

working set size (WSS) of the Dgraph process

The working set of the Dgraph process is a collection of pages in the virtual address space of the process that is resident in physical memory. The pages in the working set have been most recently and frequently referenced. In other words, the Dgraph working set is the amount of memory a Dgraph process is consuming now. This is the amount of memory that is needed to avoid paging.

The following statements describe the working set size (WSS):

- The WSS cannot be computed, although it is always less than or equal to the amount of virtual process size for the MDEX Engine.
- Determine the WSS experimentally: if you notice that increasing RSS (by adding RAM or subtracting competing processes) improves performance of the MDEX Engine, this means that the WSS was previously larger than the RSS. This was likely the cause of the performance degradation.
- If the size of the WSS grows too close to the amount of RAM, or starts to exceed it, paging to disk begins and you will notice rapid decreases in performance.

For more information, see the *Oracle Commerce MDEX Engine Performance Tuning Guide*.

Z

zone

See dynamic business rule.

