

Secure Enterprise Search

Version 1 (10.1.8)

*An Oracle Technical White Paper
January 2007*

Secure Enterprise Search

Executive Overview	3
Introduction	3
Oracle's SES	5
SES Architecture.....	6
Crawler	8
Crawler Plug-In API.....	10
Web services API	10
Administration.....	12
Search quality.....	13
Secure search	15
Admin Based Authentication.....	16
Self Service Authentication.....	16
Custom Agent or Crawler.....	16
SES Methodology.....	16
The Gather Step	16
The Analyze Step	17
Making crawling results searchable.....	17
The Maintain Step	18
New In SES Release 10.1.8.....	18
Security Plug-In Architecture	20
Other Improvements.....	20
Conclusion	21
Further Reading	21

EXECUTIVE OVERVIEW

Secure Enterprise Search (SES) is Oracle's standalone search solution with web-like quality, which securely covers all enterprise sources, and is easy to use and deploy. SES provides:

- Excellent search quality
- Sub-second query performance
- Highly secure crawling, indexing, and searching
- Out-of-the-box user experience

INTRODUCTION

Internet searches have shown that significant information uplift can accrue from search technology. Without search engines, the Internet would still have billions of web pages, but surfers would have to know URLs a priori, or navigate through directories, to locate pages of interest. Clearly it is Search that makes Google popular, and the Internet more useful even as the amount of information on it grows at a rapid pace.

Proliferation of information also exists in the enterprise; however enterprises have so far not benefited from the information uplift that good search provides. This has been largely due to the differences between the Intranet and the Internet. For example:

- Information on the Internet consists overwhelmingly of web pages. In the Intranet, information – data and content – is spread across web pages, databases, mail servers or other collaboration software, document repositories, file servers, and desktops. An Intranet search engine must be able to search an organization's web content, its applications, databases, and mail through the same interface. Comprehensiveness, across structured and unstructured sources, and ability to reach every corner (the deep Intranet) is the key to Intranet search.

Many organizations have larger volumes of text than structured data, yet consistently underutilize the textual data.

- Unlike the Internet, where all information is publicly visible, Intranet information needs to be secure. Different users have different access rights to information, and information resources are often password protected. An Intranet search engine must be able to enforce security. If a user is not authorized to see a document, email message or record – then even the existence of the record should not be visible to him. The access rights can change, and access-changes made to the different underlying information-stores have to be propagated to the search-engine quickly.
- Internet search engines like Google use the links that URLs provide between web pages to deduce the importance or relevance of a document in a given search. Unfortunately, Intranet resources do not invariably vote for each other by URL links: a document authored in PDF may not link to the database record of a customer that it describes. Consequently, different techniques are needed for high relevance when it comes to Intranet search.
- While keyword search can provide information uplift, an organization will frequently need more sophisticated queries dictated by its own business model. Rather than show dozens of hits, the search engine for an Intranet should be answer-oriented, and provide analytic capabilities that permit users to go beyond keyword search and ask sophisticated questions, untapping the hidden intelligence in the deep Intranet.
- Different Intranet users not only have different access-control rights to resources, but they also have different information needs based on job function. Search results have to be personalized to meet those needs.
- Intranet search must be multilingual. Corporations have offices all over the world with content in local languages.
- Higher service level expectations exist for the Intranet, and the robustness of an Intranet search product must match that of mission-critical enterprise software.
- Intranet search software must be simple to use and administer.

SES solves the problem of finding relevant information across your company's many disparate repositories of information, providing a very intuitive interface to search and administer.

ORACLE'S SES

Oracle has developed text and information retrieval technology for over 15 years. The base underlying capabilities of Oracle Text (a comprehensive API) have been long available with the Oracle Database. Oracle Ultra Search was introduced with Oracle9i to enable a portal search across different repositories, and is available with the Oracle Database, Application Server, and Collaboration Suite. Building on these products, Oracle's latest Enterprise Search technology adds several key capabilities.

- **Security.** The ability to search password protected sources securely. Oracle's search technology provides single-sign-on (SSO) based security where available, and can also employ application-specific security where SSO is not available.
- **Simplicity.** A simple out-of-the-box web user interface, for both search and administration - that has both the clean look-and-feel and ease of use that users prefer on Internet searches.
- **Comprehensiveness.** The ability to search across all your sources – web pages, files in file servers or desktop drives, databases, applications, mail servers and groupware, and more.
- **Connectivity to Legacy Repositories.** SES allows companies to access their most valuable assets – information about its specific business, its processes, products, customers, and documents that previously resided in proprietary repositories. Connectors include interfaces for EMC Documentum, FileNet, OpenText LiveLink, Microsoft SharePoint, IBM Lotus Notes, Oracle's E-Business Suite, among others.
- **High quality search results.** Brings for the Intranet a high level of relevance that users associate with Internet searches.
- **Going beyond keywords.** As the volume of information grows, even with high relevance the paradigm of keyword search starts reaching a plateau of diminishing returns. Users need advanced search techniques like the ability to categorize and cluster search results for iterative navigation.

SES is fully globalized and can search in all major languages, including Western European, Chinese, Japanese, Korean, Arabic, Hebrew and many more.

SES is robust and enterprise hardened. Many hundreds of queries per second can typically be served off a modest Linux machine. Typical

enterprise Intranets typically run into terabytes, and Oracle's search infrastructure has been repeatedly deployed for multi-terabyte loads¹.

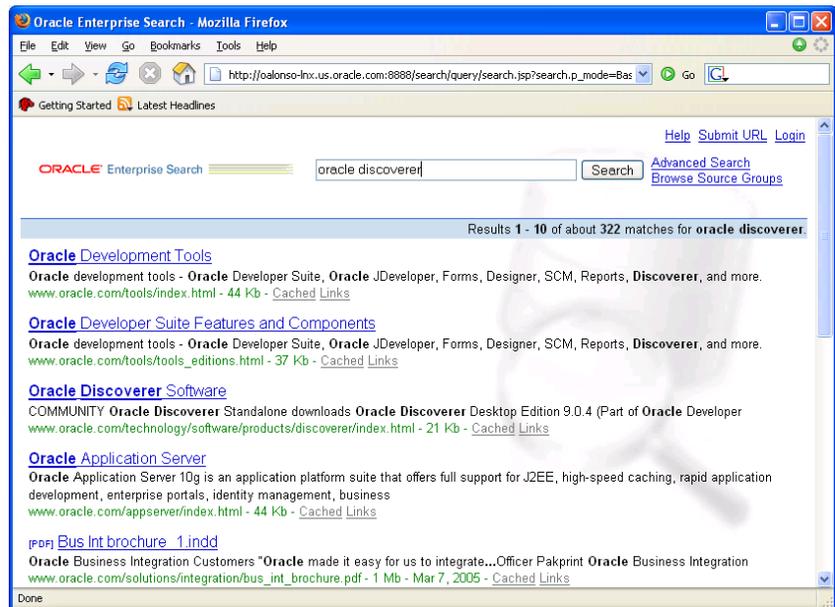


Figure 1 shows an example of SES in action.

SES ARCHITECTURE

Oracle's Secure Enterprise Search is a standalone, self-contained server for search; it operates as a "black box" that indexes information from the crawler and serves up the results. It comes with its own user-interface and administration; it does not, for example, need you to program using SQL or administer as a DBA.

Architecturally, as presented in Figure 2, the product is made up of five distinct components.

- **Crawler.** The SES Crawler is a Java process activated by your Oracle server according to a set schedule. When activated, the crawler spawns a configurable number of processor threads that fetch documents from various data sources. The crawler maps

¹ IronMountain Case Study
(http://www.oracle.com/technology/products/text/pdf/oracle_ironmountain.pdf)

link relationships and analyzes them to avoid going in circles and taking wrong turns. Whenever the crawler encounters embedded, non-HTML documents during the crawling it uses filters to automatically detect the document type and to filter and index the document.

- Database. An Oracle10g database contains the SES-repository, which stores information about the repositories indexed by SES and the search engine 'index' (information collected by the crawler, filtered and indexed by Oracle10g Text).
- Search UI & API. SES provides an out-of-the-box user interface to the Server. It also provides a web services API for building custom applications for querying indexed data, and contains interfaces for Basic Search Form, Advanced Search Form, Query Result Display, URL registration, authentication and authorization, and so on.
- Administration Tool and Interface. The SES administration tool is a browser-based application that you use to configure and schedule the crawler, configure the server, run several reporting features, and other similar tasks.
- Federator. Via the Suggested Content feature, SES also provides the ability to federate queries to other engines that implement their own search – mail servers, Internet search engines, and specific applications. Additionally, SES provides the ability to federate queries to other SES instances. These results can be combined and displayed together along with those results served off the internal index of SES Server.

SES is based on an Oracle10g database and a “stand-alone” version of the Oracle J2EE container “OC4J” used to implement a web server to serve up HTTP. During installation, the SES OC4J ‘application’—consisting of search and administration environments -- is deployed within this J2EE runtime environment. The Oracle database is configured and tuned for the special needs of the search engine.

Both SES database and web server are installed on the same machine. Moving SES database and search/ administration applications to different machines can be done in principle, but is not officially supported by Oracle today.

SES connects to Oracle’s SSO-infrastructure (OID) without the need for any customization – only simple connection parameters to Oracle Internet Directory (OID) must be specified. Users and groups required to secure searches are stored and managed within OID. SES also supports other Identity Management Servers, such as, Active Directory.

If other Oracle software is found during install, the SES installer creates a new listener with its own network configuration/port.

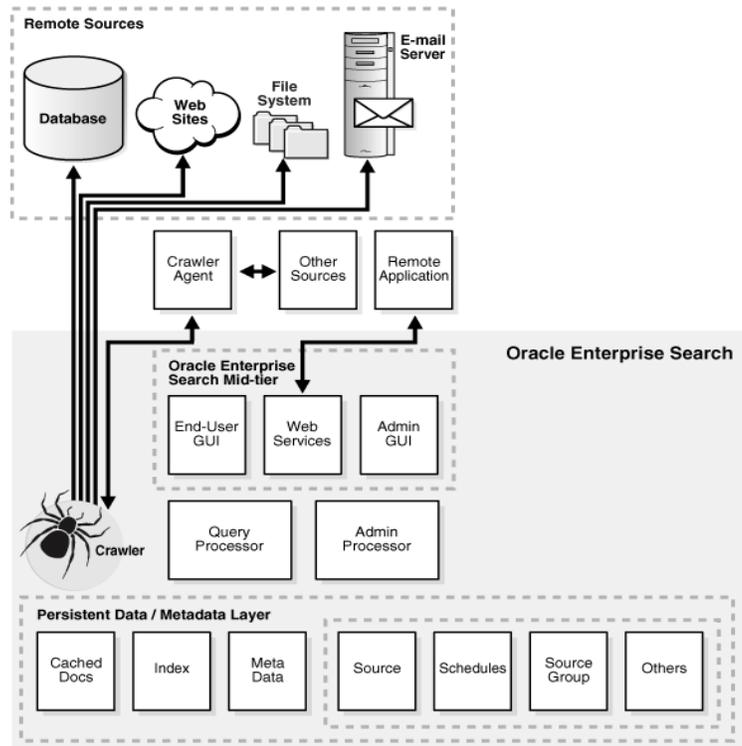


Figure 2. SES Architecture

The following sections describe some of the components and other aspects in more detail.

Crawler

The SES crawler is a multi-threaded Java application responsible for gathering documents from the data sources you specify during configuration. To crawl different repositories, the SES crawler allows you to define specific ‘data sources’. A data source is a logical construct identifying a repository. You can take a single physical repository, such as a database, and map it to multiple data sources (A data source is also the granularity at which you define metadata). SES

knows the following standard types of data sources (more data sources are available via plug-ins, see below):

- Web Sites – Define web sites as a data source with the HTTP protocol like www.oracle.com.
- Database Tables – SES can crawl Oracle databases and other relational databases that support the ODBC/JDBC standard. SES allows the crawling of both full text columns and “fielded text” columns. Fielded text columns allow you to map a database column to an SES attribute (e.g. author, title), creating a set of indexes tuned to the content of your database.
- Files – Files must be directly accessible by the crawling machine. Remote files may be crawled so long as they can be crawled through the file:// protocol. Files must be accessible by each crawler machine either locally or remote over the network.
- Emails - SES can connect to an IMAP email server and index all the emails for a user. To index mailing lists, you might choose to create a specific IMAP account, which is subscribed to the mailing lists of interest. Connectors to Microsoft Exchange and IBM Lotus Notes are also available.
- Portals - SES connectors can crawl OracleAS Portal instances, Documentum eRoom and Microsoft Sharepoint Server.
- Content Management Systems – SES connectors are also available to crawl and search EMC Documentum Content Engine, Opentext Livelink, Hummingbird DM, Filenet, IBM DB2 Content Manager.
- Applications – SES can connect to an Application, crawl its business objects and make them searchable. Application specific authorization models are supported via plug-ins. SES provides connectors to certain specific modules of Oracle’s EBusiness Suite and Siebel CRM.
- Custom – SES comes with an open set of APIs that allows a custom crawler to be built to enable crawling of any proprietary or home-grown information system. Suggested Content capability expands SES’ ability to connect to other information systems.

SES uses 3rd party filters to extract text and metadata from documents and automatically identifies document types. The filters handle popular document formats like PDF and MS-Office. There is also support for filtering documents that have been compressed with ZIP utilities.

To maintain fresh, comprehensive search results, SES uses synchronization schedules. Email search results, for example, can be updated continuously, while published content is gathered on a less frequent schedule. Each synchronization schedule can have one or more data sources attached to it.

To limit the crawling to a specific section of your corporate network or to ensure that crawling does not take wrong turns and follow link relationships that point outside your Intranet, SES lets you specify so-called ‘inclusion’ and ‘exclusion’ domains for crawls. SES supports the ability to backup the metadata that preserves all your crawler and scheduler settings and restore them if and when necessary.

The SES crawler can be instructed to collect URLs without indexing them. This data harvesting mode allows you to examine document URLs and their status, remove unwanted documents, and start indexing.

Crawler Plug-In API

The crawler can be extended through ‘plug-ins’ (Plug-ins are Java classes supplied by Oracle, or developed by the customer). Out-of-the-box, SES ships a family of plug-ins (also called connectors) to a number of enterprise content sources like EMC Documentum, Lotus Notes, and Microsoft Sharepoint among others.

Plug-ins are deployed via the SES Administration GUI – they are listed as new data sources after being defined. After configuration, plug-ins supply data to the crawler and can be indexed just like other data sources.

Technically speaking, plug-ins are responsible for collecting URLs pointing to the documents to be indexed – they pass these URLs to the crawler for indexing.

Web services API

Search engines like SES are usually integrated in existing customer web- and portal sites. Ideally, end users invoke searches from a search mask and don’t even realize that Oracle SES handles their search requests in the background. “Look and feel” of the result list must correspond with the Portal site, which invokes the search. To achieve this, SES provides a webservice interface, based on standards like SOAP and WSDL.

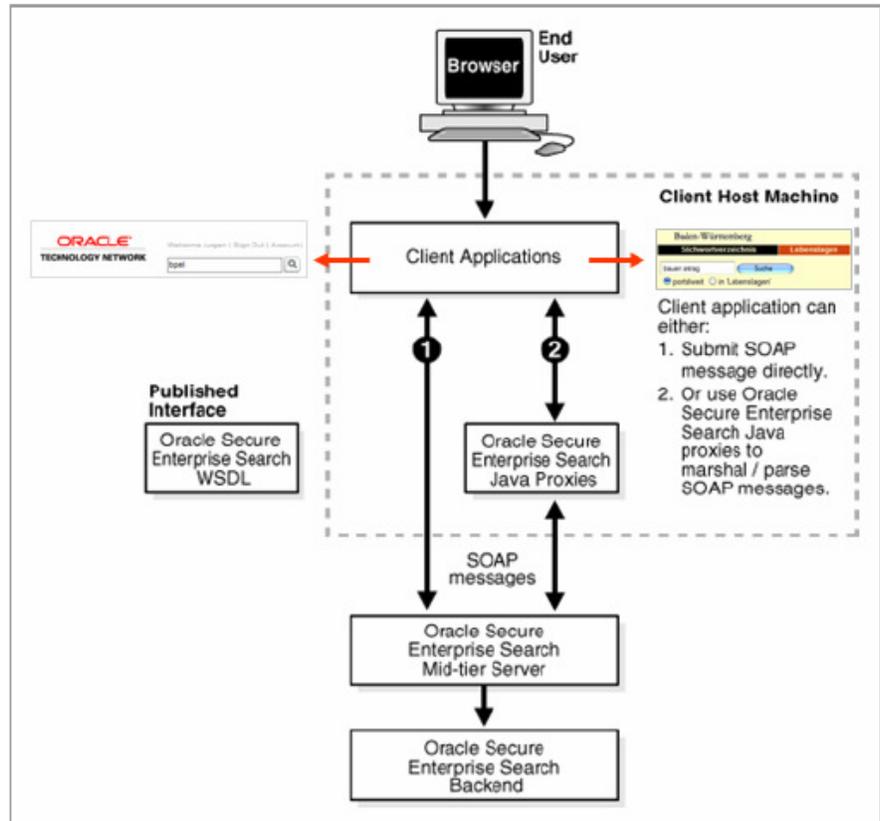


Figure 3: SES Webservice Architecture

In the above example, the end user enters their search term into an input field (“search box”). The search request is sent from the CMS application server directly to the SES web service. SES executes the search and sends results back to the calling application in the form of XML via SOAP. Results are displayed in embedded fashion – within the application.

SES uses no UDDI-repository --the WSDL-description can directly be obtained from the server.

The following code snippet shows how easy is to use the API:

```
import oracle.soap.transport.http.OracleSOAPHTTPConnection;
import oracle.soap.encoding.soapenc.EncUtils;
import oracle.search.query.webservice.client.*;

public class TestWS
{
```

```

public static void main (String[] argv)
{
    try
    {
        OracleSearchService search = new OracleSearchService();

        // Add your own code here, for example to populate
        // the query string.

        // Set SOAP URL. The URL is
        http://<host>:<port>/search/query/OracleSearch
        stub.setSoapURL("http://oes-serv-
        example:7777/search/query/OracleSearch");

        String queryString = "oracle";

        //
        // Do a simple search for the queryString we set up above
        //
        OracleSearchResult result = stub.doOracleSimpleSearch(
            queryString, // query
            new Integer(1), // startIndex
            new Integer(3), // docsRequested
            Boolean.FALSE, // dupRemoved
            Boolean.FALSE, // dupMarked
            Boolean.TRUE); // returnCount

        // Get the result set

        ResultElement[] resElemArray = result.getResultElements();

        // Loop through the results displaying the document title

        for (int i=0; i<resElemArray.length; i++)
        {
            System.out.println("Document Title:
                               "+resElemArray[i].getTitle());
        }
    }
    catch(Exception ex)
    {
        ex.printStackTrace();
    }
}

```

Administration

The administration tool is a web application that allows the administrator to:

- Define and crawl data sources.
- Define crawler parameters like URL boundary rules, crawling depth, language and proxy settings, etc.
- Create and modify schedules for the crawler.
- Set query options - Query options allow users to limit their searches. Searches can be limited to document attributes (e.g. title, author) and data groups. Data source groups are logical entities exposed to the search engine user. When entering a

query, the search engine user is asked to select one or more data groups to search from. Each data group consists of one or more data sources.

- Adjust relevancy ranking of the search hit list – SES allows administrators to influence the order that documents are ranked in the search hit list. Use this to promote important documents to higher scores and make them easier to find.
- Define suggested links for specific search terms.
- Define alternative words for specific search terms.
- Setup authentication mechanisms for certain data sources.
- Manage the backup and recovery of search metadata.

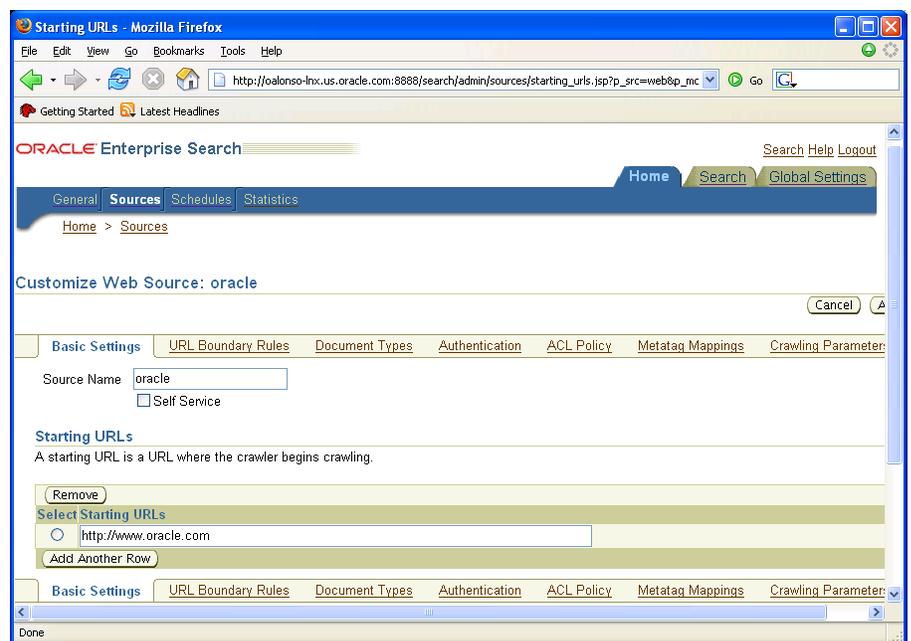


Figure 4: SES Admin interface

Search quality

Search quality or the ability to find relevant information is one of the most important features of any search engine. SES uses a wide range of techniques for providing excellent search quality.

The following techniques are used at different stages of the crawling and index processes for enhancing the overall search quality:

- Metadata processing. Metadata from pages and documents like title, author, description, headline, email header, and anchor text are all used to improve relevance.
- Duplicate elimination. There is a lot of content in a corporate Intranet that are duplicates. Copies of same presentations, web pages, text documents are all over the place. Sometimes people produce multiple files and sometimes the servers duplicate the content for mirroring. Near duplicates result from different versions, formats, HTML style, site-specific links, contact information, etc. To better match end-user expectations, duplicates are normalized to a single document.
- Complete duplicate elimination identifies and removes duplicates at the crawling stage before the document is even indexed.
- Link analysis. One of the most widely used techniques for improving relevancy is link analysis. Briefly, the idea is to discover authoritative pages by performing analysis on the link structure of the web collection. A page that is linked by many pages is important. A page that is linked by a high link score page is also important. A number of algorithms exist today like HITS and PageRank. SES has its own algorithm implementation.

The administrator can also control the relevancy using a couple of extra features: alternative links and suggested words. Alternative links is a useful feature for registering a well-known authoritative page against keywords. These links will then be displayed at the top of the search result page when the user searches for these keywords. Suggested words can map user search terms to synonyms. For example, cellular phones for cell phones or wireless phones.

In case users have trouble spelling query terms, the spell checker feature suggests corrections based on data available from a dictionary and crawled data.

Apart from all the searching features of SES, it is possible to combine browsing and searching at the same time. You can click on the browse link to navigate all the directories that SES has created automatically after the crawl. This is a good entry point when you are trying to

explore all the content that is available to search. Of course, you can search within a directory at any time in the search box.

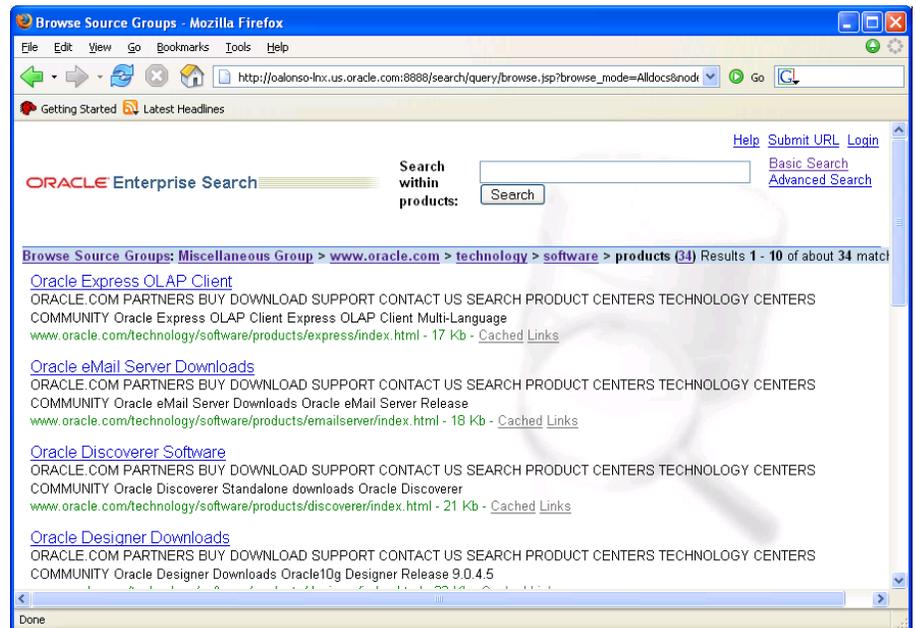


Figure 5: Browsing content directories

Secure search

SES features secure searching – the ability for users to log in and find documents which are not publicly available. To do this, SES has secure crawling capabilities, and the capability to store Access Control List (ACL) information alongside data sources

SES integrates with the Oracle Internet Directory (OID) – a Lightweight Directory Access Protocol (LDAP) server which provides network authentication capabilities. Users log in to OID, and then SES allows them to only find documents for which they are authorized.

The crawler can handle secure sources in a variety of ways. While the sources may themselves be protected by the same (or a different) OID server, this is not a requirement. Any source can be crawled in a secure manner, so long as it is protected by one of the following:

1. OID Authentication
2. HTTP Basic Authentication

3. Form-based Authentication

There are several different ways of providing access credentials to a secure crawler.

Admin Based Authentication

When a data source is defined, the administrator can enter an authorized password (either a user password or “superuser” password). This will be used to collect the information from the source. An ACL may be defined for the source that defines who can search the information.

Self Service Authentication

When a data source is defined, the administrator sets up the source but does not provide any username or password. Users are then able to log in and provide their own access credentials. A data source is then created which is specific to that user, and they are the only user who can search that information.

Custom Agent or Crawler

A custom agent is a Java module that can be used to crawl any user-specified data. The agent passes back a pointer (URL) to the information to be indexed, and optionally specifies an ACL for each document. This allows great flexibility in access control.

SES also provides the QTA (Query Time Authorization) API that allows customer to have fine control on search results at query time. SES uses this technique as the main interface to filter documents based on authentication access.

SES METHODOLOGY

What steps do you need to follow for using SES? The SES search engine follows four logical steps to provide universal search – gather, analyze, make queryable, and maintain. These steps are not novel, and are indeed found in most organizations’ business process.

The Gather Step

Gathering refers to information that exists in structured relational databases and in unstructured files, Word processing documents, spreadsheets, presentations, e-mail, news feeds, Adobe Acrobat files, and Web pages. SES gathers this information by “crawling” your corporate Intranet and looking through all the information that exists in the various repositories of your company – databases, Web pages, IMAP mail servers and others.

During the gathering process, link relationships are analyzed to avoid going in circles and taking wrong turns. As a result, SES administrators have an easier time keeping search results complete and up-to-date.

The Analyze Step

In the analyze phase SES looks at the meaning and structure of gathered information. In order for information to be searched, it must be indexed. During the analyze phase, SES uses the Oracle Text engine to extract both meaning and structure from the gathered information by creating an integrated index, effectively “normalizing” both structured and unstructured data. Oracle Text indexes contain a complete wordlist along with other information.

During indexing, text and metadata are extracted from documents by third party filtering software. This filtering technology automatically identifies document type, invokes the correct filter and produces indexable text and data. Several predefined metadata fields are supported, including author, date, and title. The filters include the most popular file types like MS Office and PDF.

Unlike some document management systems, SES gathering and analyzing is non-intrusive. Instead of physically moving documents, information and documents are analyzed but reside in their original location under their own name.

In typical Web search technologies, hundreds of hits are returned. As the number of repositories increase, the ability to rank relevance of documents decreases. SES uses the award winning relevance ranking of Oracle Text to ensure that users consistently find the needle in the haystack.

Making crawling results searchable

“Make Searchable” is the function of providing access to all the information that has been indexed in a programmatic fashion. Oracle SES provides a web services API for this purpose. Passing a search term into the query API locates all relevant documents, whether they are stored on Web servers, databases, or in applications. Customers can use SES APIs to integrate universal search into their own Web pages or applications.

The Maintain Step

The maintain step ensures that search results are updated continuously. SES lets you gather from multiple Web sites and repositories, each on a different schedule. IMAP messaging servers, for example, can be updated continuously, while published content is gathered on a less frequent schedule. SES maintains content by providing easy, intuitive utilities that provide Administrators with an easy way to keep up with new content that is added through growth or acquisition.

NEW IN SES RELEASE 10.1.8

Robust Connector Framework

Consumer search engines, like Google and Yahoo, index and search mainly HTML pages on web server. Enterprise Search Engines must also index Portals, Document Management Systems, custom applications and other software applications and systems. Oracle SES ships a family of built-in ‘connectors’ (Connectors are Java classes based on the SES plug-in API) for unlocking stored content in the most popular of these systems in use today.

The new connector family provides access to documents that reside in the following proprietary systems and applications:

- Windows NT Filesystems (NTFS) -- NT fileshares can be indexed over a network connection and don’t have to be located on the SES host machine. SES provides strong access control by reading group and user access information and storing it in its search engine index.

For SES installed on Unix operating systems, a small Agent process is installed in the same AD domain where the NT filesystem to be indexed is located. The agent sends content,

metadata, and access control information to the connector in the SES machine (agent protocol is based on HTTP and can be encrypted via HTTPS). Microsoft IIS must be present for the agent to work.

- EMC Documentum Content Server – Indexes files in cabinets and folders of ContentServer DocBases. A native identity plug-in allows SES to show only those documents that a user has access to according to permissions within Documentum. Efficient recrawls are supported – documents are only re-indexed if changed or moved within a Documentum.
- IBM Lotus Notes – Notes databases on IBM Lotus Notes Domino server instances (Notes Mail and custom applications planned for future release). The connector automatically navigates through all Notes databases on a Notes server instance. SES provides a Notes identity plug-in to use the Notes directory for authentication & validation of Notes-native users and groups.
- FileNet Content Engine – Indexes object stores and documents in Content Engine Instances. Automatically navigates object stores to retrieve all documents.
- FileNet Image Services – Folders, documents and image annotations containing text (stamps, transparent text, and sticky note) in Image Services libraries. Native identity plug-in.
- Open Text Livelink – Folders, files, projects, and tasklists in LiveLink Enterprise Server workspaces. Identity plug-in for Livelink provided.
- Microsoft Exchange – Indexes emails, attachments, calendar items and related metadata attributes in Exchange 2000 and 2003 stores. Efficient incremental recrawls are supported. Requires Microsoft IIS and ‘Agent’ software from Oracle (Agent, included with 10.1.8 release, sends content and metadata between Exchange host and SES host machines) from Oracle to be installed on the same Windows domain as the Exchange server.

All connectors are pre-configured (additional licensing may be required for some of the above connectors) and feature access control integration between SES and the legacy repository served by the connector.

Oracle SES also searches across a number of Oracle-internal sources:

- OracleAS Portal page group, pages, and items.

- Oracle Collaboration Suite ContentServices and Calendar
- Oracle ContentDB – Folders, documents, and categories. Supports efficient re-crawls: Only documents with changed content, changed metadata/category metadata, and moved documents are re-indexed during incremental crawls.
- Oracle E-Business Suite 11i – Allows for crawling views, or queries, in Oracle database underlying 11i. Each record in the view or query is considered a separate document.

Security Plug-In Architecture

Secure Enterprise Search is now directly integrated with third-party access control- and identity management solutions, including Microsoft's Active Directory. No synchronization of users or groups with Oracle Internet Directory is necessary. SES can directly access Active Directory (no extra coding required) through a new authorization API and identity 'plug-in' architecture. SES ships plug-ins for Oracle's Internet Directory and Microsoft's Active Directory, among others. The new architecture even allows customers to build their own 'identity plug-ins' (supplies user and group information) for crawling sources with proprietary (non-LDAP) security schemes.

Other Improvements

A new Suggested Content feature lets you display real time content or federated search results from other information sources in the search results screen. A stylesheet can be applied to the suggested content before it is displayed in the search result list.

A simplified method for configuring secure search with OracleAS Single Sign-On, a "title fallback" feature to override default document titles picked up during crawling with a more meaningful title later, a more simple configuration of federated sources, and case insensitive relevancy boosting (documents with "Oracle" are boosted when you enter "oracle".)

Upgrade from SES Release 1 (10.1.6) is supported.

CONCLUSION

The Enterprise Intranet is different from the Internet -- the information in it comes from many different types of sources; searches need to access password protected content; determining the importance of Intranet documents requires different techniques than on the Internet, and effective answers must often go beyond result hit lists. Secure Enterprise Search is built to bring to the Intranet the information uplift users get on the Internet. By deploying Oracle's search solution, you can not only find information securely and effectively, mitigating information over-load, but also unlock the hidden intelligence that lies untapped in the deep Intranet.

FURTHER READING

- [1] SES Home page:
<http://www.oracle.com/technology/products/ses/index.html>
- [2] SES Business White Paper
- [3] SES Features Overview



Oracle Enterprise Search White Paper
January 2007

Author: Stefan Buchta

Contributing Authors: Vishu Krishnamurthy, Meeten Bhavsar

Oracle Corporation
World Headquarters
500 Oracle Parkway
Redwood Shores, CA 94065
U.S.A.

Worldwide Inquiries:

Phone: +1.650.506.7000

Fax: +1.650.506.7200

www.oracle.com

Oracle Corporation provides the software
that powers the internet.

Oracle is a registered trademark of Oracle Corporation. Various
product and service names referenced herein may be trademarks
of Oracle Corporation. All other product and service names
mentioned may be trademarks of their respective owners.

Copyright © 2007 Oracle Corporation
All rights reserved.