

ORACLE SECURE ENTERPRISE SEARCH

VERSION 11G R2

KEY FEATURES

RELEASE 11.2.2.2

HIGHLIGHTS

- Facet Navigation
- Push-based Content Indexing
- Multi-tier Install
- Search Result Tagging
- Unified Microsoft Sharepoint Connector, certified for MOSS 2010 version
- New search result 'hard sort' option
- Auto Suggestions "As You Type"
- Sitemap.org Support

FACET NAVIGATION

- Full GUI support for facet creation, manipulation, and browsing. No programming required. Facet API support is also provided
- Facet navigation is secure. Facet values and counts are computed using only those documents that the search user is authorized to see
- Hierarchical- and range facets
- Date, number and string types
- Update facets without re-crawling

PUSH-BASED CONTENT INDEXING

- Allows customers to push their application content to an SES endpoint via HTTP POST
- Documents and data to be posted need to be wrapped as RSS-type 'feed'
- RSS feeds can also contain access control info and meta data for each document to be indexed. File attachments to each document are also supported
- Once posted, documents are immediately indexed without the need to wait for a crawl schedule

Oracle Secure Enterprise Search 11g (SES), a standalone product from Oracle, enables a high quality, secure search across all enterprise information assets. Key SES features include:

- *The ability to search and locate public, private and shared content across intranet web content, databases, files on local disk or file-servers, IMAP email, document repositories, applications, and portals*
- *Excellent search quality, with the most relevant items for a query spanning diverse sources being shown first*
- *Sub-second query performance*
- *Highly secure crawling, indexing, and searching*
- *Integration with Desktop Search tools*
- *Ease of administration and maintenance – a 'no-DBA' approach to Search.*

Information Uplift for the Intranet

As a result of search engines on the Internet, the power of effective search technologies has become clear to everyone. Using the World Wide Web, consumers have become their own information retrieval experts. But search within enterprises differs radically from public Internet search. The information that businesses store and use for decision making is scattered across millions of documents and data fragments. It resides on a wide range of, often incompatible, IT systems. Many document

s are confidential or private -- access to these must be controlled to ensure that only documents that the user is authorized to see are returned by the search engine.

Internet search engines such as Google use the links that URLs provide between web pages to deduce the importance or relevance of a document in a given search. This mechanism, called Page Rank, is akin to a citation count, or a 'vote' for one page by another. Unfortunately, Intranet resources do not vote for each other in the same way: a document authored in PDF may not URL-link to the database record of

MULTI-TIER INSTALL

- Three install modes allow for more flexible deployment with smaller footprint.
- Appliance Mode Install – standalone install of all SES components. Same as before.
- Existing Database – Utilize existing Oracle 11.2.0.3 DB to store SES index and data tables. SES installer will create a new Fusion Middleware application server.
- Software-only – Install into existing Oracle 11.2.0.3 DB and 13.3.6 WebLogic App Server/Fusion Middleware 11.1.1.6 mid tier.

SEARCH RESULT TAGGING

- Provide end users with a way to collectively improve relevancy by adding tags
- End users can add their own tags to search hits via new UI widgets
- Tags are immediately tied to the content of the tagged document, indexed, and become searchable. The creator of a tag, or others, can then easily retrieve the tagged document by simply searching for the tag
- The SES Administrator chooses who can tag: Everyone, all logged-in users, or users with special tagging privilege

MORE POWERFUL AND FLEXIBLE SORTING OF SEARCH RESULTS, INCL. NEW 'HARD SORT' OPTION

A new sort mechanism allows customers to influence the order in which search results are presented:

- 'Global hard sort' (new) – SES will sort all results by order of date or other attribute. Ideal for news, or other date-sensitive content
- 'Soft sort', combining date and relevancy score – SES will resort the top-N most relevant documents by order of date or other attribute of choice

a customer that it describes. In the Intranet, linkages specifying relatedness of information may be specified in different ways, or indeed be absent. Consequently, an algorithm that works well on the Internet will fare poorly in the intranet; different techniques are needed for high relevance when it comes to Intranet search.

Business users frequently need more sophisticated queries beyond the keyword model. An Intranet search engine must understand the query context better and provide analytic capabilities. 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. 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. It must be open enough to fit into existing management regimes, and its operations should be amenable to inspection and debugging. IT organizations must be able to ensure that the search engine is not a backdoor for collecting employees' personal data, showing them commercial messages from third parties, or a black box whose operations are impervious to audit.

Oracle's Database and Application Server technologies are widely used by businesses to store and access both data and content. For a number of years, Oracle has invested in building out technologies for secure access to data and content, text search and indexing, as well as connectors to different data sources and IT systems.

Secure Enterprise Search, a new standalone product from Oracle combining the above technologies, provides a comprehensive search over all intranet content including databases, files on local disk or file-servers, IMAP email, document repositories, applications, and portals. The search is secure and fully multilingual. Advanced search, including meta-data search, is supported. The product has been designed to be simple to use and administer. This Data Sheet summarizes the key features.

Customers familiar with Oracle Text will appreciate the enhanced search syntax ("query syntax"), which takes SES beyond the more simplistic Google-like syntax of earlier versions. A number of new operators allow for specifying thesaurus operations (e.g. synonym terms, narrower term, broader term), fuzzy spellings, and wildcard matching. Some of the new operators allow extending SES through Oracle Text (for example, user defined Thesauri).

A Document Service API provides developers with a hook into the SES crawler pipeline. They can obtain documents found by the crawler and manipulate them before they are passed to the indexer. This is useful if you would like to generate and insert your own meta data tags into your documents before they are being indexed. Or, use it to extract entities like addresses or phone numbers from your content for compliance or auditing purposes. Some of our customers even use this service to build their own customize search engine, while taking advantage of the

UNIFIED SHAREPOINT CONNECTOR

- Provides index capabilities for Microsoft Sharepoint Portal and Sharepoint Services
- Certified for Microsoft Sharepoint 2010, but also supports MOSS 2003, 2007, and Windows Sharepoint Services 2.0 and 3.0.
- Supports new Sharepoint 2010 document rating feature, library types, and objects
- Indexes HTTPS-enabled Sharepoint sites

OTHER NEW 11.2.2.2 FEATURES

- Auto-complete -- Google-like "type-ahead/auto complete/predictive" search where suggestions for the search term pop up as the user types into the search box
- Support for Sitemap.org

SES 11g FEATURE AREAS

SEARCHABLE REPOSITORIES

- HTML pages served up by a Web Server
- Database Tables - Search Oracle databases and any other databases that support the ODBC standard. Database tables can reside in Enterprise Search's own database instance, or they can be part of a remote database accessed over a network. Both full text columns and "fielded columns" can be crawled
- Files - Local or remote files can be made searchable through the file:// protocol
- Emails - Emails and mailing lists can be crawled via the IMAP protocol.
- Oracle WebCenter Suite – Crawls and indexes native content in Web Center Spaces (lists, pages, wikis, blogs, group spaces, people), document library- and discussions content.
- Starting with PeopleSoft 9.2

existing SES crawler, legacy software connectors, and infrastructure technology.

Developers can directly influence relevancy rankings by changing how document attributes like title and keywords are factored into ranking scores.

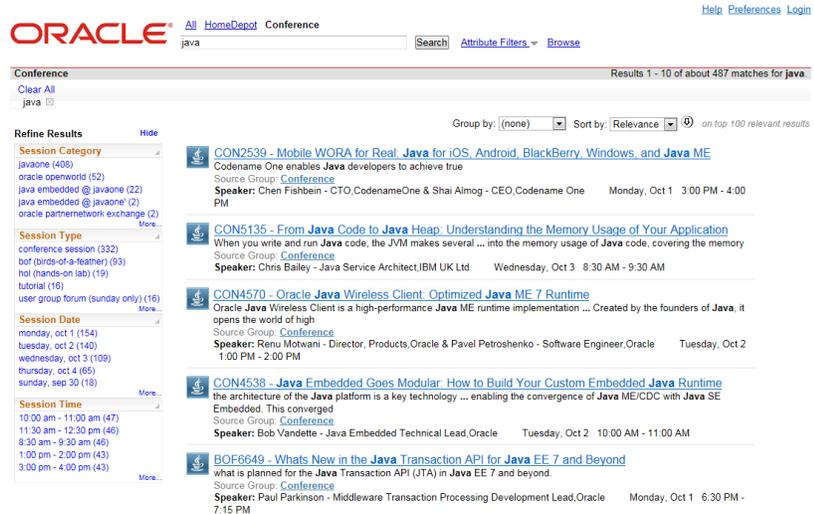


Figure 1: The Secure Enterprise Search query sample application, including the new facet navigation and some search results.

The SES 11gR2 Release

11.2.2.2 is the first release of the 11gR2 series and adds important new capabilities, including facet navigation, multi-tier install, search result tagging, global sorting of search results, push crawler, and a newly redeveloped connector to Microsoft Sharepoint that supports all versions up to, and including, Sharepoint 2010. Product Upgrade from SES 11R1 is supported.

Facet Navigation

Facet search is a new approach that combines text search with a progressive narrowing of choices in several dimensions. It uses a hierarchy structure to enable users to browse the information space by iteratively narrowing the scope of their searches in a predetermined order. It has become the prevailing user interaction mechanism in e-commerce sites and in Oracle SES it is being extended to also deal with semi-structured data. Facets are first created either in the SES administrator interface or per API. Administrators can decide which facets are deployed for which source group. The values of facets are collected during crawling, and are available immediately when a crawl completes. End users can then start navigating facets via SES search application. Facet definitions can be modified after crawling - there is only a small latency for introducing new nodes, for example.

(PeopleTools 8.5.3), SES replaces Verity as the built-in search engine for PeopleSoft.

PRE-BUILT CONNECTORS

Shipped with the product. Utilize document ACLs (user and group-membership):

- Microsoft Windows NT Filesystems (NTFS)
- Oracle WebCenter Content (formerly Stellent Content Server/UCM)
- EMC Documentum Content Server DocBases
- IBM Lotus Notes/Domino databases and Email
- Connectivity to Oracle E-Business Suite, Siebel, and PeopleSoft. Many modules offered out of the box. E-Business Suite Search Modeler tools allow customers to customize searchable objects for modules in EBS R12.

SECURE SEARCH

- Search non-public sources, but see only what you are authorized to see.
- Security Plug-in API allows for directly accessing groups and users in non-Oracle LDAP-based identity management software.
- Support for secure crawling, including HTTPS (secure socket layer/encrypted data stream), cookies to maintain a session, HTTP digest authentication, and HTML form login.
- Multiple ways of getting document visibility Access Control Lists (ACLs) from data sources: Applied by the crawler based on administrator input ("ACL stamping"), supplied by the sources (e.g. Document ACLs), or supplied by the search user ("self-service")
- Integrated with Oracle Access Manager (OAM) LDAP-based enterprise security. Share LDAP-users and groups

Push-Based Content Indexing

There are now two general strategies for crawling data sources, 'pull-' and 'push-', which allow to crawl- or extract data by accessing external sources/systems.

Previous releases of SES were limited to pull-crawlers which have the disadvantage that the index might not be accurate, since SES only learns about data changes when a crawl is executed. With SES 11.2.2.2, Oracle supports a mechanism for 'pushing' content into its search engine index. The new 'push' mechanism is implemented on the side of the client system which includes the data to be indexed. It proactively updates the index via an RSS-based feed mechanism. After wrapping their documents or data records in RSS-fashion, customers can simply use HTTP POST to post the feed to an endpoint and it will be indexed with a very small latency. Meta data and access control info can be included with the feed. The new push mechanism allows for content, which is not linked and can therefore may not easily be found by the SES crawler, to be sent directly to SES for indexing.

Multi-tier Install

Oracle recognizes the need to give its customers the power to deploy SES more flexibly and with smaller footprint. SES now supports three install options:

- Software Appliance -- Same as in earlier releases. Standalone install of the whole SES software stack.
- Existing Database -- Customers can use an existing Oracle database for the SES search engine index and data tables, while the SES installer creates a Fusion Middleware Application Server instance to run the SES application.
- Software only -- Install SES software into both preexisting Oracle database and preexisting Fusion Middleware Application Server.

Version requirements apply: The database can be single instance or RAC and must be Oracle release 11.2.0.3. The Application Server can be single node or cluster and must be Oracle WebLogic release 10.3.6/Fusion Middleware release 11.1.1.6.

Search Result Tagging

With a new tagging feature, SES 11gR2 provides search end users with the power to influence search result rankings directly. End users can log into the search UI and enter tags for a search result of their choice. A match for a tag makes a difference in relevancy – search requests that match tag text will have their relevancy scores boosted.

More Flexible and Powerful Sorting of Search Results, 'Hard' Sort

A new UI drop-down feature offers 'global' re-sorting of the search results by user selectable attribute. 'Global' means that sorting now extends over all matches of a given search term. In prior SES versions, only the top-N search results deemed most relevant by the relevancy ranking algorithm of SES would be re-sorted by attribute.

The administrator can configure which attributes are available in the drop-down list for sorting and can also choose a default ordering.

The new 'hard sort' is ideal for date-sensitive content such as news items, where an

between your data sources and Enterprise Search. Restrict search results based on LDAP-group membership or user id

Generic Security Infrastructure with plugs into:

- Oracle Access Manager – Allows for OAM to SSO-protect the Oracle SES search interface. Web crawling support for OAM protected sources is planned for SES 11.2.2.3 (contact SES Product Management for details)
- Microsoft's Active Directory
- Windows Native Authentication via Kerberos. A type of Single Sign-on where Users who are already logged into their Windows workstation are automatically signed into the SES search page. Reduces time and password fatigue.
- Oracle Virtual Directory (OVD) planned for next point release, 11.2.2.3 (contact SES Product Mgmt. for details)

PERFORMANCE & SCALE

- Internal query parallelization to leverage multiple cores and disk spindles. Exploits today's high RPM disk drives
- Supports high churn applications/very low latency of index updates by automating de-fragmentation of search engine index

150 SUPPORTED DOCUMENT FORMATS

- Microsoft Office® Suite
- Adobe Acrobat® PDF
- Sun StarOffice
- Other popular desktop formats – Lotus 1-2-3®, Lotus Freehand®, Corel Word Perfect® etc.
- Autovue 2D & 3D CAD files. Open CAD files directly inside Autovue
- Most other common document formats
- Documents in these formats can be searched whether they

ordering by recency often trumps relevancy requirements.

Newly Unified Sharepoint Connector Covers all Popular MOSS Versions

Oracle has unified Sharepoint connectors into one single plug-in, which supports both legacy Sharepoint versions 2003, 2007, as well as the new 2010 version. The connector supports the popular new features of Sharepoint 2010 including the new Sharepoint 2010 document rating feature and can index objects new to Sharepoint such as records- and asset library, blogs, pages, and team sites.

Other SES 11gR2 Enhancements

The Search UI now has Auto Complete, the google-like "type-ahead/auto-complete/predictive" search feature (where the suggestions for the search term pop up as the user types into the search box).

Support for sitemap.org is also provided.

Oracle SES 11g Feature Overview

Comprehensive Search over Intranet Content

Oracle SES can search not only Web documents, but any kind of documents found on an organization's servers, such as PDF, Word, Excel, etc. It can also create a common index for all documents of whatever filetype relevant to a particular term. .

SES features a family of built-in 'connectors' (Connectors are Java classes implementing the SES plug-in API) which securely extend SES to other enterprise systems, providing knowledge workers with a single point of access to information, people, and expertise across the enterprise. With cost-effective out-of-the-box deployment, comprehensive security-mapping, and Active Directory integration, SES connectors maximize the ROI of an organization's ERP, HCM and ECM systems by unlocking their content and making it searchable from a single, Google-like search interface or from a taskflow provided in WebCenter Portal.

Scale and Performance

Oracle SES 11g features internal parallelization and fine-tuned search engine index structures.

Parallel query and index partitioning algorithms leverage multiple CPU cores and disk spindles and can significantly improve search query performance and facilitate searches of very large data sources. Parallel query is automatically used for querying large, partitioned data sources. Each partition may reside on a separate disk. To make best use of this feature, you should run Oracle SES 11g on a server with a 64-bit multi core CPU. The server should have at least 8GB of RAM and multiple fast disk drives, either directly attached or via SAN with sufficient available bandwidth.

SES 11g fetches and caches index blocks from disk in much larger, contiguous chunks and buffers, minimizing the number of times the engine has to go to disk when serving search requests. Further changes in the layout of the index improve

are files on disk, email attachments, served by webservers, stored in a database or part of another supported repository.

WEB QUERY LANGUAGE

- Basic keyword search.
- Advanced, parametric search and Query Syntax, including Thesaurus (synonym terms, narrower term, broader term), fuzzy spellings, wildcard matching, NEAR, and nested search expressions like (x AND y) OR z.
- Dynamic page summaries. Search Keywords shown bold within title and excerpt of search-result page.
- Cached pages.
- Highlight query term on cached pages.
- Customize search result page look-and-feel using Freemarket templates

SEARCH ANALYTICS & METRICS

- Reports: Most popular queries, documents not found, click-throughs, and many other reports.
- Document Relevancy boosting. Administrator can boost document relevancy and customize result sets.

WEB BASED ADMINISTRATION CONSOLE AND SEARCH UI

- Administer crawl activities
- Configure secure search
- Insight into user search behavior: Monitor crawl and search query statistics; most popular end user searches, searches with no results
- Very flexible presentation of search UI. Incorporates the Freemarket templating engine, a Java library which makes customization of the default UI very easy without the need for any coding

HIGH PERFORMANCE,

the performance of single word searches.

Comprehensive Security Support

Secure Enterprise Search comprehensively addresses secure search. All your resources – intranet web pages, rows in database tables, Oracle Application Server Portal pages, emails, and documents in files or special repositories – can be protected. Search results are filtered so that the results page only shows links for documents to which the user has access. Oracle provides three ways through which this secure results filtering can be accomplished:

- Secure Enterprise Search can work from a centralized authentication scheme, like a Unix-, Microsoft-, or Oracle Internet Directory LDAP login, to identify which sources a given user can access. If you can't access Oracle Portal per your LDAP privileges, for example, Secure Enterprise Search will not show you any results coming from your Oracle Portal system.
- For more granular security, Secure Enterprise Search can store Access Control List (ACL) information associated with each document as part of the search engine index. Both 'ACL Crawling' and 'ACL Stamping' are supported. In ACL Crawling, Secure Enterprise Search obtains ACLs for each document directly from your repositories. This is not always possible, as when the source does not have a document model – e.g. an application generating dynamic content. In ACL Stamping, the search administrator specifies authorization roles directly in the Enterprise Search administrative console by entering a "grant" list of LDAP users, and groups, which are allowed to search a particular source. For example, all documents retrieved during a crawl can be marked searchable by anyone belonging to groups G1, G2 and G3. Or, all users U1, U2 and U3 can be granted permission to search all documents of a source.
- Lastly, SES can filter search requests through the authorization mechanism of the source system. This is called Query time authentication. After the search engine index has returned a hit list, for each item in the list the search engine re-accesses each store, and passes to the store the user's credentials to check if the user is (still) authorized to see the item. Thus, even if the access privileges have changed for the user since the last crawl, security is not compromised.

Security Plug-In Architecture

To avoid duplication of access control information, Secure Enterprise Search is integrated with Oracle's own corporate identity management solution, Oracle Internet Directory (OID), and it can be synchronized with other Identity Management solutions like Microsoft's Active Directory and Novell's eDirectory products. 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-

SCALEABLE CRAWLING

- Multi-threaded Java crawler.
- Gather from multiple Web sites and other data sources, each on a different schedule.
- Crawler can be adapted to crawl new repositories through the Secure Crawler SDK
- Page link relationships are mapped and analyzed for better relevancy.
- Limit crawling to specific sections of your Intranet by setting 'inclusion' and 'exclusion' domains.

FLEXIBLE APIS

- Integrate Oracle Secure Enterprise Search with your application, or create new custom search applications, by means of a Web Service-based Query API. Sample source code available via download
- Manage SES from your program or scripts via Administration APIs and Command line Interface
- Secure Crawler SDK for custom crawling sources or accessing custom sources.
- Document Service API
- Real-Time data access and suggested content API with ability to format output using XSLT stylesheet.
- Extensible authentication and authorization framework allow for building applications with dynamic security.
- A Display-URL feature lets you render database search results according to an associated database Web application

FEDERATED SEARCH APIS

- Federate queries to other Enterprise Search servers, or internet search engines, combining results in one display
- Suggested content feature lets you index and display real time content together with the search results. SES retrieves

LDAP) security schemes.

The Web Service Query API and query samples allow for specifying a search user. Only users known to your corporate LDAP server can submit searches.

A number of optimizations have been engineered to provide good performance. ACLs are cached in memory and evaluated at run time during query invocation. Authorization automatically turned off if there are no secure documents in an Enterprise Search instance.

Higher Search Quality

In a typical Internet search, hundreds of thousands of hits are returned. Similarly, with a comprehensive intranet search, the number of repositories increase, the 'noise' of poor matches in the search hit list can quickly overwhelm the user. As we have discussed, the Intranet is also handicapped by not being able to use URL references to determine the popularity of pages. New approaches are needed to be able to find the needle in the intranet haystack.

Secure Enterprise Search has built a new relevance model optimized for the Intranet.

- It incorporates award-winning relevance ranking of Oracle Text to ensure that users consistently find very precise information, with its algorithms re-calibrated for typical intranet loads
- It builds separate internal indexes for data and metadata, and uses a unique weighting mechanism to prioritize metadata over data.
- It performs URL link analysis, where this is possible and useful.
- It includes de-aliasing and disambiguation. A significant amount of the 'noise' in an Intranet search comes from nearly identical documents. If you created a presentation, mailed it to a colleague as an attachment, who copied a dozen other people who stored it in different places, should the search engine return dozens of identical hits simply because the location of each is different? Secure Enterprise Search introduces detection of identical content available at multiple URLs or locations.

Crawling and Filtering

Secure Enterprise Search searches a variety of repositories. It gathers information by crawling your corporate intranet and looking through all the information that exists in the various repositories. Secure Enterprise Search features:

- Multi-threaded Java crawlers
- Ability to dispatch crawler on multiple machines
- Security in the crawl framework, including HTTP basic and digest authentication, as well as ACL gathering as described in the Comprehensive Security Support section above.
- For binary document formats – Office, PDF etc. -- *filters* are needed to extract

data from your applications and applies a stylesheet to create an HTML fragment

- Google OneBox providers can be configured as Oracle SES content providers

SIMPLE TO ADMINISTER

- Simple one-touch Install
- No-DBA approach to administration
- Browser-based interface to manage search configuration.

SUPPORTED OS PLATFORMS

- Linux platforms (Oracle, RedHat and SUSE, 32-bit and 64 bit)
- Microsoft Windows 64bit 2008 R2
- Solaris SPARC 64 bit
- Will index any HTTP-compliant server including Netscape Enterprise Server, Microsoft IIS and Apache.

READY FOR GLOBAL MARKETS

- Supports all major languages, including double-byte languages like Japanese, Chinese, and Korean

EMBEDDING SES AS A SEARCH SERVICE

- Invoke searches via Query API
- Perform administrative options like starting and stopping schedules
- Silent install with your software
- Push custom metadata into SES for searching
- Tune relevancy of search results based on application specific characteristics

plain text from the binary document. Filtering technology that automatically identifies document type and invokes the correct filter to produce textual data and metadata suitable for indexing. Filters are provided for most (150+) popular file types.

By using the crawl-index-search model, SES insures data gathering is non-intrusive - documents are analyzed, but are physically left in their original location under their own name. Crawls can be scheduled for low-load times.

Metadata

Metadata is one of the key enablers of effective enterprise search.

- Secure Enterprise Search extracts metadata fields from documents, database tables, email and other repositories.
- It provides a flexible metadata mapping methodology, mapping the extracted metadata information to query page attributes for a powerful combination of full text retrieval and "fielded" text retrieval.
- Many different attribute types, including date ranges and List of values (LOV) are supported
- Metadata search is automatically incorporated in Basic Search to determine which documents are the most relevant. It can also be explicitly invoked from Advanced Search.

Web-based Deployment & Administration Environment

Secure Enterprise Search is simple to use and deploy because it is based on the open standards of the Internet and a 'no-DBA' philosophy. Users can install the product, initiate crawling, and get some search results in a very short period of time. All the components needed internally by the search engine – text index, web server, metadata store, crawlers – are all bundled and need no separate installation or configuration. For secure search, in addition to installing SES you have to point it to the appropriate LDAP installation that manages user identity.

Flexible, Easy-to-Integrate Query Framework

Secure Enterprise Search lets you create custom search applications that work with any type of information by means of a set of query API's. These API's can be used from Web applications to retrieve and display query results. A Web service query API provides for an easy embedding of Enterprise Search into your pages. Even Search Administrator functions, such as data source creation and scheduling can be driven from Web Services.

Web-style Search Interface

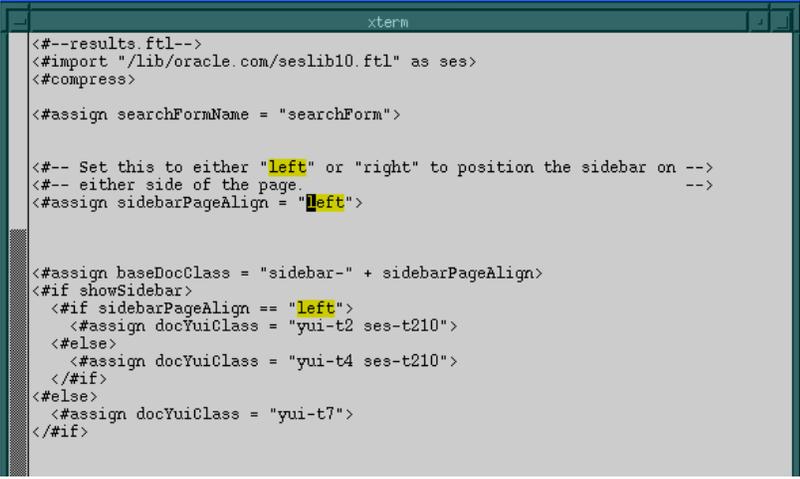
To shorten your development cycles, Secure Enterprise Search includes a fully functional search application for users to query and display search results. The search application comes in Java (JSP).

See figure 1 above for a screenshot of the search application, which also incorporates the Freemarker templating engine, a Java library which makes

customization of the default UI very easy. The idea behind Freemarker is that you separate UI design from the actual program code. This allows for changing the appearance of a UI page without the need for changing or recompiling code, because the application logic (the SES Java programs) and the query page design (Freemarker templates) are separated.

See figure 2 for an example how Freemarker can be used to customize and reconfigure the SES search page without the need for any Java coding.

SES provides Freemarker elements – templates, skins, macros – to customize all GUI elements of the search application.



```
xterm
<!--results.ftl-->
<#import "/lib/oracle.com/seslib10.ftl" as ses>
<#compress>

<#assign searchFormName = "searchForm">

<!-- Set this to either "left" or "right" to position the sidebar on -->
<!-- either side of the page. -->
<#assign sidebarPageAlign = "left">

<#assign baseDocClass = "sidebar-" + sidebarPageAlign>
<#if showSidebar>
  <#if sidebarPageAlign == "left">
    <#assign docYuiClass = "yui-t2 ses-t210">
  <#else>
    <#assign docYuiClass = "yui-t4 ses-t210">
  </#if>
<#else>
  <#assign docYuiClass = "yui-t7">
</#if>
```

Figure 2: Freemarker greatly eases the process of customizing the look and feel of the default SES interface. Here, changing the word “left” (highlighted) to “right” will have the effect of moving the “Filter Results By” sidebar of Figure 1 from the left side of the page to the right side.

Intuitive, Web-based Administration Environment

Secure Enterprise Search provides a simple browser-based administration tool that centralizes its maintenance. Wizard-like pages and dialog boxes simplify the tasks of identifying data sources, specifying security rules, scheduling crawls, monitoring search effectiveness, and tuning relevance.

Administrative functions: Search engine administrators can quickly and easily define information sources to be crawled, schedule maintenance crawling, and define user accounts for administrative users. Users simply identify themselves as an authorized search administrator to the Secure Enterprise Search administration login Web page. Once the administrator logs in, the administration environment displays Wizard-like pages and dialog boxes for managing the search engine.

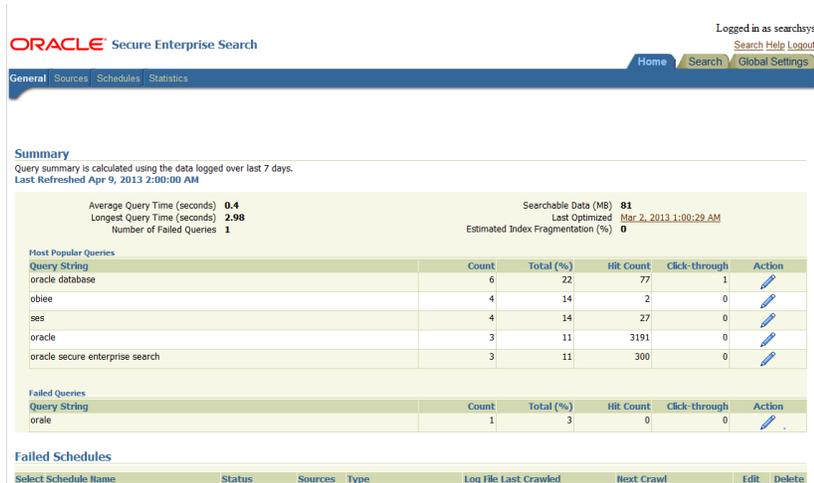


Figure 3: Part of the SES Administration Interface.

An Administration API supports the management of large-scale deployments by providing a command-line interface to administrative tasks previously only available through the SES Admin GUI:

- Create, change, or delete sources or schedules
- Start and stop schedules
- Configure SES crawlers
- Failed operations are automatically rolled back

Use the Administration API within an interactive session, or by executing commands from operating system prompt.

Search Once, Find Everything with Federated Search

Federated Search allows for searches both over Secure Enterprise Search crawled repositories and other heterogeneous data sources which do their own crawling and indexing. For several reasons a repository may not be suitable for crawling. The rate of updates may be too rapid for scheduled crawling. Opening a door for a crawler may lead to vulnerability for secure data sources, or crawling a database-backed website as a set of unstructured documents may be redundant if the data is already contained and indexed in structured form inside a database.

The ability to federate searches complements Secure Enterprise Search centralized crawling:

- Data sources keep their own full-text index or do their own crawling.
- SES accepts search terms from end users and brokers them to other data sources. The sources implement a web service using the SES query API to be able to send results back to SES for display.

Pricing

Oracle Secure Enterprise Search is licensed both by processor and by named users.

Please consult <http://www.oracle.com> for the latest pricing information.

Summary

Oracle's Secure Enterprise Search product allows you to reduce the time spent finding relevant documents on your company's information repositories. It crawls, indexes and makes searchable your corporate intranet through a Web-style search. It eliminates the need for coding against hard-to-use low-level APIs as it exists with open source systems such as Apache Lucene. It organizes content from multiple repositories by extracting valuable metadata that can be used in portal applications. It provides effective search by returning more relevant hits - the best relevance ranking in the industry - and finds what you want. And it provides the best database integration and secure searching in the industry.