



by
[Susan Florio](#)
(with David Kajmo)

Level: Intermediate
Works with: Domino 5.0
Updated: 05/03/99

Inside this article:

[The benefits of Domain Search](#)
[How do you use Domain Search?](#)
[How do you set up Domain Search?](#)

Related links:

[Knowledge Management in R5](#)
[Domino R5 Technical Overview](#)

Get the PDF:

Many consider the search for knowledge a lifelong quest. Using the new Domain Search capabilities in Notes and Domino R5, you can shorten that journey substantially. Now you can get the information you need when you need it, so your search for knowledge really isn't much of a search at all. The new Domain Search capabilities allow you to search every server in your domain for any information located in virtually any document, whether it's in a Domino database or a public file system. You can search the information from a Web browser or the Notes client. Domain Search combines the ease of searching on the Web with the power and security of Notes.

Domain Search in R5 is a key element in the Lotus knowledge management strategy. In addition to Domain Search, knowledge management tools in R5 include the Content Catalog, Headlines, and TeamRoom. For an overview of Domain Search and how it fits into the bigger picture of knowledge management, see "[Knowledge Management in R5](#)."


As an administrator, you might think "Great! I'm glad to have these new capabilities, but I'll bet it's hard to set up and administer." The good news is that Domain Search is simple to set up and once you get it going, you can centrally control what users can search.

Now you probably want details. We'll start by giving you an overview of Domain Search and talk about its advantages. Then we'll give you a brief run-through of how Domain Search works in the Notes client, and get into the details of how you can set up Domain Search for your organization.

First, be assured that any applications you designed in R4.6 with custom searches will work when you upgrade to R5. R5 search is backwards compatible, meaning that whatever you could do in R4.x, you can do in R5.

Getting the big picture

Before we get going, we'll define some terminology. The Domino server that brings you all of the Domain Search capabilities is called the *Domain Catalog server*, because the Domain Catalog plays a fundamental role in Domain Search. The Domain Catalog server holds the *Domain Catalog* and the *Domain Index*. The Domain Catalog collects an abundance of information about every database on every server in your domain, including information on whether or not it should include the database in the Domain Index. The Domain Index is a central index of all the text in all the databases and all the file systems you choose to allow users to search. When users conduct a Domain Search, they search the Domain Index stored on this server. The following illustration shows a Domain Catalog server with a Domain Catalog and a Domain Index:



If you are wondering how Domain Search works in a mixed environment (with both R5 and R4.x servers), you will be glad to hear that you can have both R5 and R4.x servers in your environment. Users will need the Notes R5 client or a Web browser to execute a Domain Search.

Domain Search is different from [Domino Extended Search](#), which works with Notes and Domino Release 4.5.x , 4.6.x and 5.0, and allows you to search across multiple Notes domains. Extended Search is a separate product that you can purchase in addition to Notes. Domino Extended Search is a search broker application that accepts user queries, translates them, and brokers them out, in native language, to various repositories, which can include Domino databases, relational databases, or ERP systems that already have their own full text indices. These various repositories then perform their own searches and return their results to the Extended Search broker, which aggregates the results and presents them to the user through a dedicated Notes application.

The benefits of Domain Search

The benefits of Domain Search are:

- It's easier to search - you don't need to know where the data you're looking for is located to be able to execute a search. You can search everything from one place.
- It's secure - it uses the ACL of the database to ensure that users receive only the results that they're allowed to see. In addition, if a document has a Reader field, the user must be listed in that field in order to gain access to the search result.
- It's centrally administered - you control what gets indexed.
- It's browser accessible - you can use a Web browser to execute the same searches, using the same user interface as you would in the Notes client.

Ease of use

Domain Search overcomes the database-specific searching limitations of R4.x and allows you to search any and all databases in your domain. You don't need to know what database contains the information you are looking for. You can click the arrow that appears next to the magnifying glass at the top-right corner of your screen and search everything in your domain.

Security

One of the biggest advantages of using Domain Search to search Domino databases is that it is completely secure. This is because in addition to using a central index, Domain Search collects the information contained in the ACL of each individual database. When you submit a query to the R5 Domain Catalog server, it finds all the index hits. However, before it displays the links to you, it figures out whether you have access to the document and only returns the links for documents you have access to. It does this by looking at the ACL information in the Domain Catalog. In addition, R5 also supports Reader fields. Reader fields list groups and individual users allowed to see a particular document. These fields further refine the ACL for a specific document. If you are in the ACL as Reader, but not in the Reader field of a document, you can't access that document through a search.

Central administration

Using the new Domino Administrator, you can centrally control and administer the databases you want indexed and therefore, determine what users can search. Since you can control what databases get indexed, you also control the size of the Domain Index. There may be applications that you don't want to include in the Domain Index because you don't want users to search for that information. For example, you may not want to include salary, review, benefits or other sensitive information for your organization. Or you may have information that is only used by a small or specific audience that you want to index in individual databases, but not include in the Domain Index.

Browser accessibility

Web application and site developers need to use full text search capabilities in their designs because searching is how most users navigate the Web. Recognizing this, most of the new Domain Search features render in HTML within a Web browser. In addition, the query form and the results form are documents. As a result of this, if you use a browser to access Notes databases, you can conduct a Domain Search just like you would in the Notes client. To activate this functionality, Web developers can use a URL syntax, JavaScript, Java, or the usual C/C++ and LotusScript APIs. For example, a developer could include a link within the database that points to the query form in the Domain Catalog. This link has the following syntax:

<http://servername/catalog.nsf/domainquery>

A developer can place this link on any Web page to allow browser users to search the domain.


How do you use Domain Search?

Domain Search allows you to set up a single, central index for all the content in your domain. Then, Notes R5 users or browser users can securely search all the databases in that domain. As an administrator, you need to have a basic understanding of how end users execute a Domain Search.

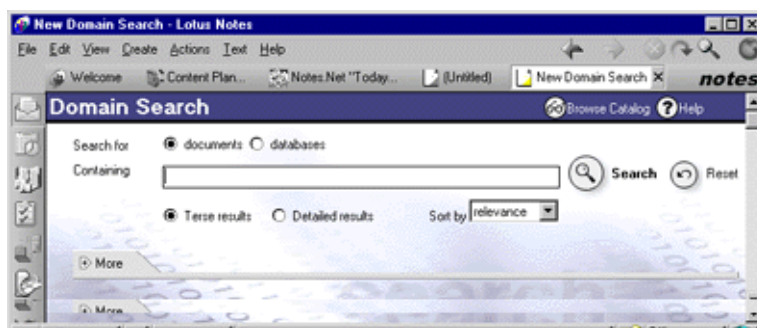
Before users can search, the Notes clients in your organization must know where the Domain Catalog server is. When Notes R5 users use the Notes client for the first time, the client uses the default setup profile on the server to populate the Catalog/domain Search server field in their Location document with the name of their designated Domain Catalog server. When you register new users, you can use a setup profile to automatically enter their search server. When users execute a search, Notes goes to the designated Domain Catalog server, looks in the specified Domain Catalog and retrieves a search form, which is stored in the Domain Catalog.

The following steps show you how to use Domain Search:

1. Click the arrow that appears next to the magnifying glass at the top-right corner of your screen, as follows:



2. Choose Domain Search.
3. Enter your search query in the Domain Search form:

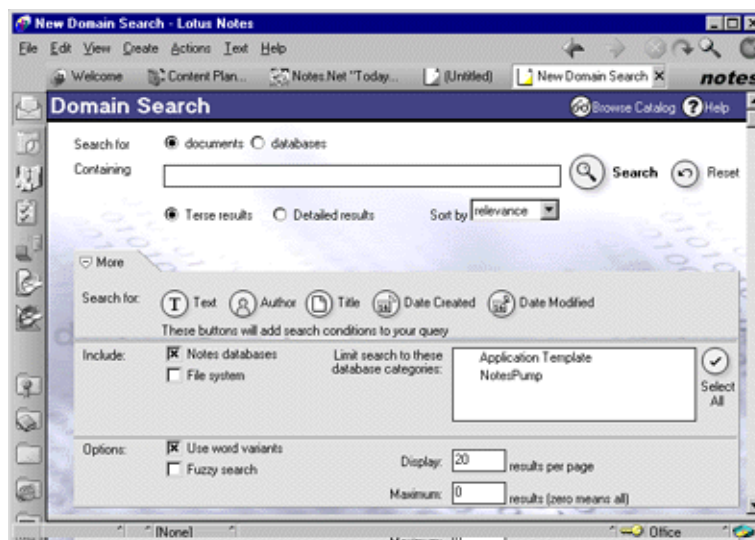


You can choose to search for documents or databases. You can search for any word, phrase, or number by entering it in the Containing field. You can specify whether your results are detailed or terse and how you want Notes to sort them. If you want to execute a more advanced search query, choose More.

4. When you finish defining your query, click Search to execute your search.

Advanced query options

If you choose More to execute an advanced query and you chose to search for a document, you can add more conditions to your query, as shown in the following screen:



You can specify that Notes should search for the specified words in one or more of the following types of fields:

- Text
- Author
- Title
- Date created
- Date modified

For example, you could specify that you are looking for the words "Domino" and "failover" in Text fields, or that you are looking for "John Doe" in an Author field.

You can also change the scope of your search. By default when you use Domain Search, you search all the databases in the Domain Index. However, you can target only databases that fall within a certain category. To categorize a database, open the database and choose File - Database - Properties. In the Design tab, you can include the database under a category. For example, you could categorize certain databases under Human Resources and others under Marketing. These categories appear in the More tab and you can choose to search only Human Resources.

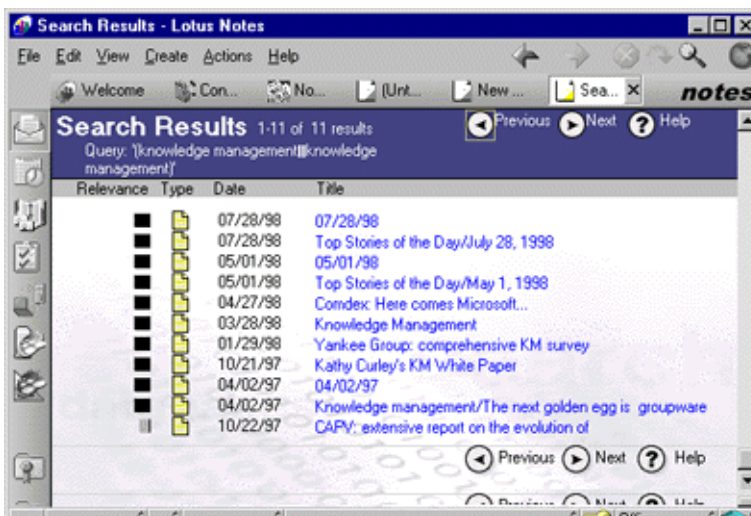
Another advantage of R5 Domain Search is that you can also search an external file system. You can specify this by selecting File System in the More tab. We will discuss how to set this up later in the article. If you enable both the Notes Database and File system options, then the Domain Search query searches all information indexed in databases and file systems.

You can specify that you want Notes to search for word variants (also called *stemming*). This finds common variations of the word you search for. For example, if you search for the word "swim," Notes returns results that also contain "swimming," "swims," "swimmer," and so on.

You can also do a fuzzy search. If you choose this option and then misspell a word that you want to search for, Notes will usually find the word in spite of the misspelling. Even if you spell the word correctly, a fuzzy search finds misspelled words in documents and returns them in your search results. Fuzzy search also works well if you want to find information by searching for only a few characters as opposed to an entire word.

Lastly, the More tab allows you to control the number of results displayed per page when the results of your search come back. R5 now returns results to you page by page. If you specify 20 in this field and execute a search, Notes quickly returns the first 20 hits. It no longer calculates all of your search results before displaying the first 20. When you choose to view the next 20, it calculates and returns only those 20. This is a substantial gain in terms of performance.

The results appear after you execute your search. Results display with a relevance rank, the date of the document, the document type (Notes or file system), and the title as a link. If you specify that you want details, Notes returns a title, short description of the document or database, and the location of the document (which database it's in).



You can edit a page of results and forward it to anyone you think would be interested in the information. When you have a Search Results page displayed in the Notes client, choose Actions - Forward and the results appear in the body of a new mail message. You can then edit the results or send them off to co-workers. When your co-workers open the e-mail, they can click any of the links in your search results and retrieve the information associated with it.

You can also cut and paste your search results into a separate document, or store the results somewhere else, such as in your group's TeamRoom database. This makes it easy to share information with others and to point people to documents you've found that may be helpful to them.

Behind the scenes -- How do you set up Domain Search?

To set up Domain Search, you must have an R5 server that is your Domain Catalog server. This server contains the Domain Catalog. The Domain Catalog is one of the central components necessary for Domain Search to work.

The Domain Catalog and the Database Catalog -- what's the difference?

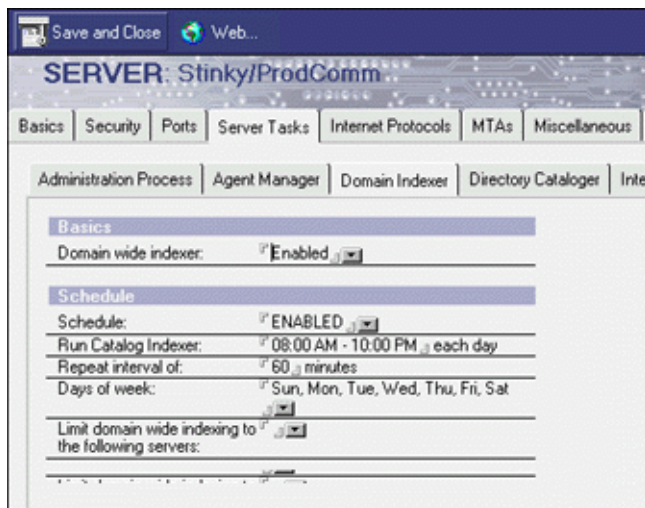
A Domain Catalog and a Database Catalog are based on the same R5 Catalog template (catalog.ntf), and they collect similar information, but the Database Catalogs in R5 don't replicate as they did in R4.x. If you designate an R5 server as a Domain Catalog server, the Administration Process puts that server's name in the LocalDomainCatalogServers group, which has sufficient access to allow replication between Domain Catalogs. If the server is not a Domain Catalog server, the Administration Process puts the server in the LocalDomainServers group, which has only Reader rights in the Database Catalog ACL. This prevents replication of the Database Catalog and prevents it from collecting information from every server, growing to a large size, and taking up a lot of disk space on each server.

Database Catalogs exist on all new R5 servers by default (the NOTES.INI files have the line "ServerTasksAt1=Catalog", so the Catalog task will run every night at 1AM). Any given Database Catalog only contains information about the databases on its particular server. A server can have either a Domain Catalog or a Database Catalog, but not both.

Creating a Domain Catalog server

To create a Domain Catalog server, you have to install Domino R5, and then check that the design of the Domino Directory (formerly known as the Public Address Book) updated during installation. Next, you must edit the Server document to enable the "Domain wide indexer" and its Schedule. To enable the Domain wide indexer:

1. Open the Domino Directory and find the Server document for the server you want to become your Domain Catalog server.
2. Select the Server Tasks tab.
3. Select the Domain Indexer tab.
4. Enable the "Domain wide indexer" field option and Schedule field.



5. The Domain Indexer runs according to the schedule you specify in the next three fields. By default, the Domain Indexer runs every 60 minutes. You will probably need to experiment with the schedule to meet the needs of your organization. The more often the indexer runs, the fresher the index will be, but the more CPU it requires.
6. In the "Limit domain wide indexing to the following servers" field, you can further restrict which databases are indexed. If the field is empty, all databases are indexed. If the field is not empty, the Domain Indexer only indexes those databases that match the specified criteria. The field itself is a Names type. Therefore, you can use various combinations of text including wildcard characters, or multiple server names to specify which databases you want to index. For example, Lotus databases are hosted on servers identified by a hierarchy such as /CAM/A/LOTUS and /CAM/M/LOTUS. Leaving this field blank causes the Domain Indexer to index all databases with the multi-database bit enabled on any of these servers. However, if you want to limit Domain Indexer activity to just one of the servers, for example mail servers designated by /M/LOTUS, enter */CAM/M/LOTUS in this field. This makes it easy to exclude mail servers from the index. **Note:** The "Limit domain wide indexing to the following servers" field is not available in the R5 beta releases.

Now instead of just a regular R5 server, you have a designated Domain Catalog server.

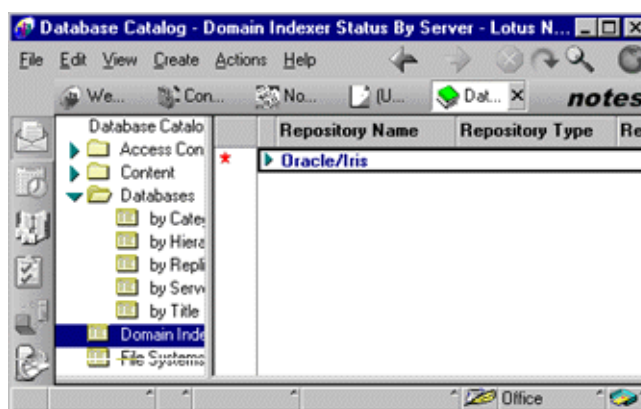
The Catalog task knows that if it runs on a Domain Catalog server, it should catalog databases on itself and then collect information from other Domino servers in its domain. It adds the information collected from databases to the Domain Catalog. If a Domain Catalog doesn't exist, the Catalog task creates the database from the catalog.ntf template.

Both the R5 Domain Catalog and the R5 Database Catalog collect considerably more information than the R4.x Database Catalog. They collect the database's:

- Title

- The server it is on
- File name
- Size (in bytes)
- Full ACL
- Replica ID
- Percentage used
- Creation date
- Number of documents
- Whether or not it is full-text indexed
- Whether or not it should be part of the multi-database index, or Domain Index

They also collect most of the information found in the Database Properties infobox. To see the Domain Catalog (catalog.nsf), open the database on the server that is your Domain Catalog server. The following screen shows the views within the Domain Catalog:



You can also use the Domain Catalog to get information about database ACLs. For example, if someone leaves the company, you can look in the Domain Catalog and see all the databases where that person had Manager or Author access.

When the Catalog task runs, it collects the information listed above from the server it is on and creates one entry for each database in the Domain Catalog. It then looks for other Server documents in its Domino Directory. The Catalog task tries to connect to each server it sees in the directory. If it can connect, it first looks for a Database Catalog. If it finds one, it does a pull replication and takes in all of the information contained in the Database Catalog. (Remember, the Database Catalog contains the same information as the Domain Catalog.) If the task doesn't find a Database Catalog, it manually collects the necessary information from each individual database in turn. Then it creates entries in the Domain Catalog for each database and ends up with a listing of all the databases on all the servers in your domain. Lotus recommends that you have a Database Catalog on each server, because pull replication is faster than spidering each database individually to collect the required information.

But do you need more than one Domain Catalog server? If you have only one server in your organization, you can let it handle serving mail, applications, and search. As a rough rule of thumb, if you have eight or more Domino servers in your domain, you may want to dedicate a server for search. When the user load becomes too great for one Domain Catalog server, you can have multiple Domain Catalog servers. You set them up the same way you set up the first one. The second server simply replicates with the first server. The Domain Catalog is a database so it replicates with other Domain Catalog

servers, however the actual binary index needs to build on each server. The second Domain Catalog server builds its own index by establishing its own sessions with the target servers and indexing their content. You should take this activity into account when scheduling your Domain indexing. If you have multiple Domain Catalog servers, you can cluster them and get the benefits of failover and load balancing.

Controlling what gets indexed

After you set up the Domain Catalog server, you will probably want to control the content included in the Domain Index. The Domain Indexer looks through all the entries in the Domain Catalog for a list of databases that have a specific bit ("Include in multi-database indexing") enabled. This tells the Domain Indexer to include that database in your Domain Index. The "Include in multi-database Indexing" option appears on the Design tab of the Database Properties box. If the Domain Indexer finds this bit enabled, it indexes the database. Then, users executing a Domain Search can find information contained in this database.

You can set the Domain Indexer to run as often as you like. When it runs, it looks in the Domain Catalog to see which databases it should index and includes the index information in the central index. The central index can be large, depending on the size of your domain.

You can use the Domino Administrator to control indexing centrally. You can select a group of databases and centrally enable this option. To do this:

1. Open the Domino Administrator.
2. Select the desired server.
3. Select the Files tab.
4. Shift and click to select the desired databases.
5. Under Tools on the right side of the screen, expand the Database menu.
6. Choose Tools - Multi-database Index.

If you want to include an individual database in the Domain Index:

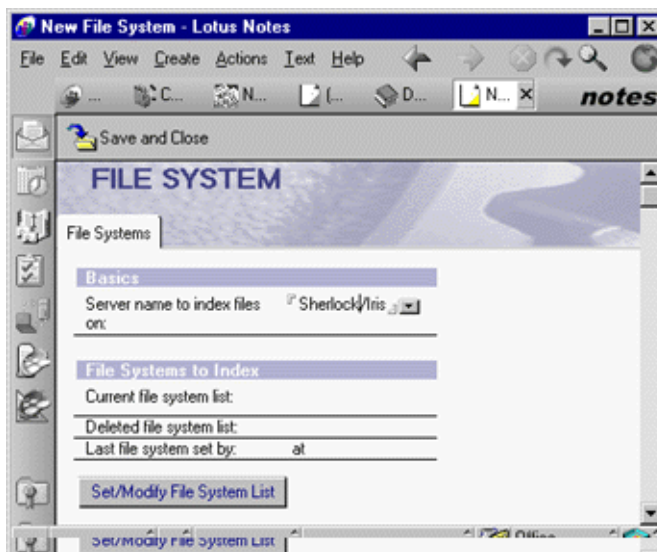
1. Open the database and choose File - Database - Properties.
2. Select the Design tab.
3. Select the "Include in multi-database indexing" option.

Indexing file systems

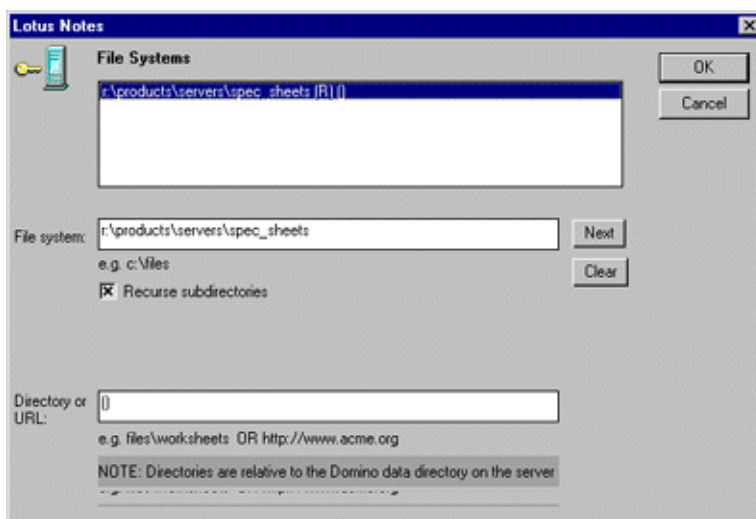
You can also index and search file systems. As long as the file system services of the operating system on which the Domain Catalog server runs can map a local drive or resource to the network volume, Domino can index the file system.

To set up file system indexing, use the File System document in the Domain Catalog. The File System document allows you to specify the file system directories and the file types to include in the index.

1. Open the Domain Catalog (catalog.nsf).
2. Choose Create - File System.



3. Specify the name of the Domain Catalog server in the Server name field.
4. Click the Set/Modify File System List button.
5. Enter the specific directory in the File System field and specify whether or not to recurse (or move down the hierarchy of) subdirectories.



6. Enter the Directory or URL for the files. This path appears in front of the file name and helps retrieve the file when users click on the link in the Search Results page. If the files reside on the Domain Catalog server or a mapped drive, you can enter the directory path (or directory links) here. If the files reside on a server with an active HTTP server, you can enter the URL for the files. The remote HTTP server then serves the files to users when they click the link in the Search Results page. This is recommended since the remote HTTP server off-loads the file serving chore from Domino.
7. Click the Next button, and enter another directory if there are other directories on this server you want to index.
8. Click OK when you finish with all the directories.

HTML files are indexed through the file system. This means that the Domain Indexer won't spider a Web site. For example, if there are three HTML files on

your indexed file system and one file has six links to other Web pages, the indexer won't follow those links and index information on those other pages.

In addition to creating File System documents to specify which directories to index, you need to create a Web Configuration document to map URLs to directories. This allows the Domino server to return the actual file to users when they click on the link on the Search Results page. To create a Web Configuration document:

1. Open the Domino Directory and select the Server document for the server that is your Domain Catalog server.
2. Choose Actions - Web - Create URL Mapping/Redirection.
3. On the Basics tab, choose URL - Directory.
4. On the Mapping tab, enter the directory or URL that you entered in the File System document you just created (step 7, above) in the Incoming URL string field. (for example, /spec_sheets)
5. Enter the path to the indexed file directory in the Target server directory field (for example, r:\products\servers\spec_sheets)

This allows the Domino server to retrieve the files and deliver them to the user. For better performance, consider storing files on a file server that has its own HTTP services enabled. By creating a URL Redirection document pointing to this remote server, you'll off-load the file serving chores to the remote server, freeing up the Domain Catalog server to handle the next user query.

Indexing a file system is not as secure as indexing a Domino database. There is no Domino ACL for a file system as there is for a Notes database. If you want to index a file system, we recommend you don't put sensitive data in it. If you choose to index a file system, be aware that security only comes into play when a user attempts to open a document from the result list, not just when a user executes a query. Users can query the file system and get back index hits for documents that they don't have access to. When they try to open a document they don't have access to, they won't be successful, but they may discern the contents of a file by repeated queries.

Domain indexing in R5 automatically indexes file attachments in Domino databases. If you have an indexed document, and it has a .PDF file (or another popular file format) attached to it, Domino indexes that .PDF file. When users execute a Domain Search, the search engine can produce hits to the attached .PDF file as well as the other contents of that document.

Customizing the forms

Both the search form that appears when users execute a search and the results form that appears with the search results are customizable. An application developer can customize either form to include a corporate logo or bitmap or re-arrange the fields.

You can create additional, customized query forms and create a setup profile in the Domino Administrator to push bookmarks to the new forms to users. The bookmark should appear in the Search folder under the user's bookmarks folder. For example, you can set up a form that allows users to search just Human Resources databases, or a form users can use to do stored queries (search for the same thing repeatedly).

Conclusion

Domain Search in R5 combines the ease and scope of Web searching with the power and security of Notes. This makes it easy to find the knowledge you're looking for anywhere in your domain. Lotus will continue to enhance Domain Search in Notes and Domino knowing that the ability to search for needed information is a critical aspect in the future of knowledge management.

ABOUT DAVID KAJMO

David Kajmo has worked for Lotus about 7 1/2 years. Currently, he's a product manager

focusing on knowledge management. Outside of work, Dave recently developed a fondness for skydiving.

What do you
think about
this article?

Register
Here!

[About this Site](#) | [Feedback](#)
[Lotus Home](#) | [IBM Home](#) | [Iris Home](#)
Copyright 1999 Iris Associates Inc.

