

Network Spotlight Best Practices

A Technical Best Practices White Paper

Group Logic White Paper | October 2010



Access.



Share.



Extend.™

What is Network Spotlight?

In addition to traditional filename search, ExtremeZ-IP's Network Spotlight support enables Mac OS X 10.5 or later clients to use the Mac's built-in ability to search the contents of files stored on ExtremeZ-IP shared volumes. Once Network Spotlight support is properly installed and configured on the ExtremeZ-IP server, the Mac client simply connects to a volume as they always have and the Finder will issue Network Spotlight searches automatically.

Server Requirements:

Windows 2003 Server, 2008 Server, XP Professional, Vista

ExtremeZ-IP 6.0 or later

Windows Search 4.0 or later - Windows Search is designed to handle up to approximately 1 million files. Performance degradation is likely on servers under regular high load or when more than 1 million files are being indexed.

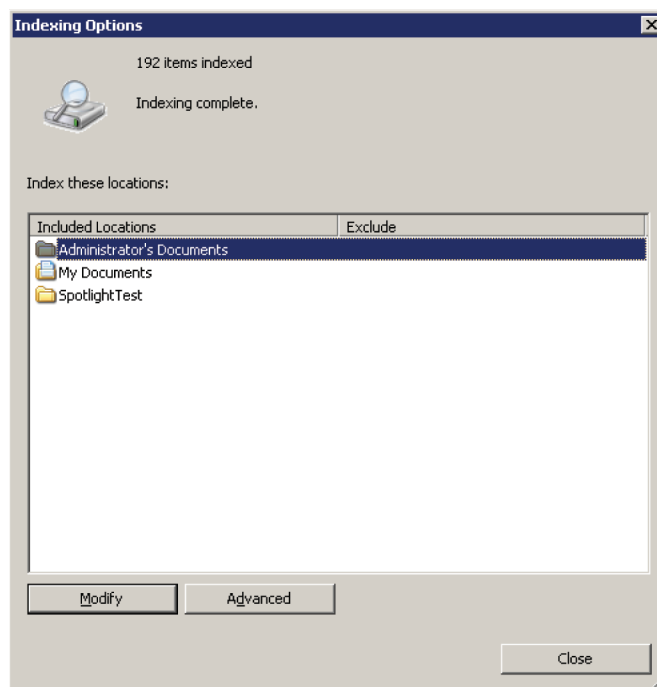
Macintosh Requirements:

Mac OS X 10.5 Leopard or later

Installing Windows Search on the File Server

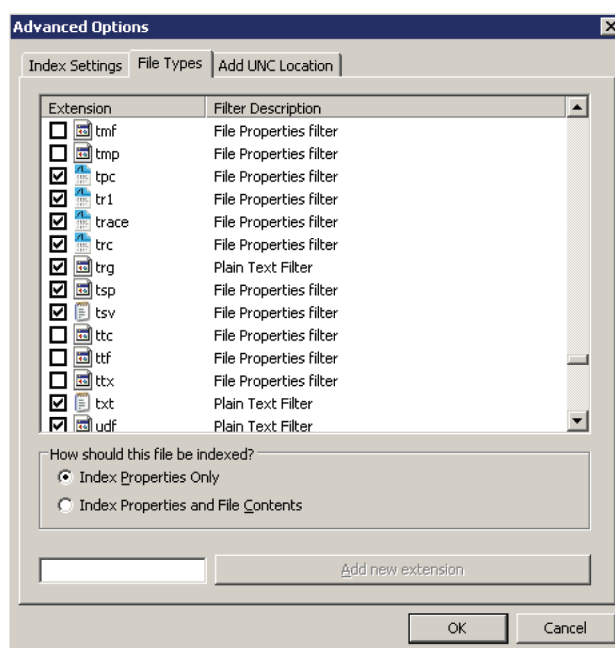
ExtremeZ-IP Network Spotlight support uses Microsoft Windows Search to satisfy search requests. In order for Network Spotlight support to function, Windows Search must be installed on your ExtremeZ-IP server. Windows Search is included and enabled by default on Windows Server 2008 and Windows Vista installations. If your server is running Windows 2003 or XP, Windows Search is available as an optional update through the Windows Update utility.

By default, Windows Search will index any My Documents folders and any folders shared using Windows File Sharing (SMB). The remainder of your server's storage will not be indexed and any ExtremeZ-IP shared volumes that lie outside the folders indexed by default will not be searchable with Network Spotlight until they are manually configured for indexing. To include additional folders in Windows Search indexing, go to the Windows [Control Panel](#) and select [Indexing Options](#).



The [Indexing Options](#) window will open and show a summary of all locations currently being indexed. To configure additional folders, click [Modify](#). The [Indexed Locations](#) window will open and allow you to select one or more drives or folders to be indexed. Once you've selected all the necessary locations, click [OK](#) and these new locations will begin to be indexed immediately.

Windows Search also provides a set of [Advanced](#) options. These allow for modifying the storage location of the search index database, manually rebuilding the index, configuring which file types are included in the index, and determining whether to index file contents and properties or just file properties for those files.



By default, there are many relatively uncommon file extensions that are disabled in the Windows Search advanced [File Types](#) options. It is important to note that, if a file extension is unchecked in this dialog, files with that extension will not be indexed in any way by Windows Search. If this is the case, Mac clients will be unable to find files with those extensions, even by searching on their file name. While the majority of the disabled file types are quite uncommon, it is recommended that they all be manually enabled in case your users require them in the future.

Based on Group Logic's own testing and recommendations from Microsoft, we suggest you limit the total number of files being indexed to 1 million or less. Indexing performance degradation and an increase in time required to return search results will become noticeable as the number of indexed items approaches 1 million. If your server contains over 1 million files that need to be indexed, we recommend you do not install and enable Network Spotlight support. The default ExtremeZ-IP indexed search functionality can be used in these cases and will continue to provide fast search results for filename searches. However, this search method does not do content searching.

When choosing locations to index, we recommend you separately select each ExtremeZ-IP shared volume folder that you would like to enable for Network Spotlight support. This will keep the number of items being indexed to a minimum and typically allow more headroom for volume growth. Simply indexing an entire physical disk is quick, and easy to configure, but if your disk contains a large number of unshared files or files that do not need to be searched by content, they will adversely affect search performance as the total number climbs past 1 million.

If the root of an ExtremeZ-IP shared volume is not configured for indexing by Windows Search, Network Spotlight support will not function for that volume and a message highlighting this condition will appear in the ExtremeZ-IP debug log. This situation can be remedied by adding that specific folder, or any parent of that folder, to the Windows Search indexed locations.

Additional details on configuring and maintaining Windows Search can be found in Microsoft's Windows Search IT Guides:
<http://technet.microsoft.com/en-us/library/cc771203.aspx>

Windows Search Optional File Type Support

Windows Search is capable of searching the contents of the majority of common file types. If and how content is indexed is determined by the file's extension.

Windows Search supports optional "iFilters", which allow additional file types to have their contents indexed.

A list of the file types supported by default and additional iFilter information is available at:
<http://www.microsoft.com/windows/products/winfamily/desktopsearch/technicalresources/filetypes.msp>

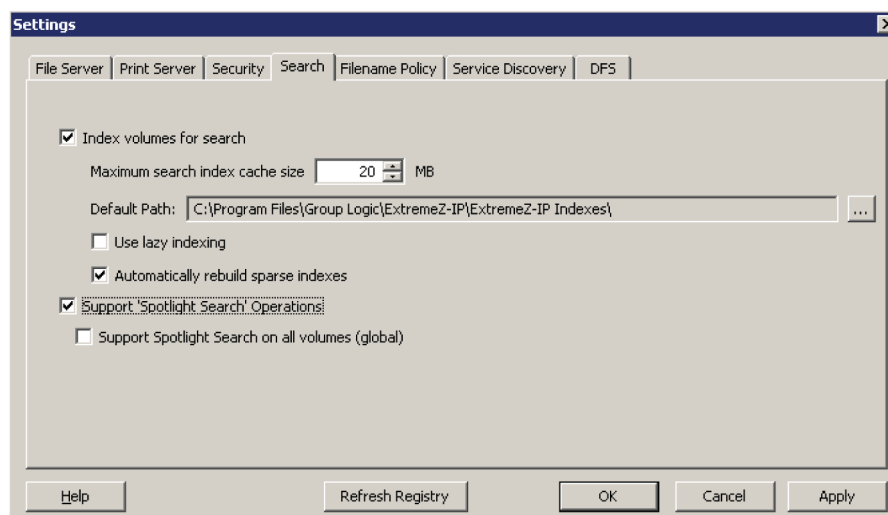
One common file type that is not supported for content indexing by default is PDF. This capability can be added by installing the Adobe PDF iFilter on your server. Ensure that the PDF extension is checked in File Types tab of the advanced Indexing Options after installation.

The Adobe PDF iFilter for 32-bit systems is installed along with the Adobe Reader application:
<http://get.adobe.com/reader/>

The Adobe PDF iFilter for 64-bit systems can be found here:
<http://www.adobe.com/support/downloads/detail.jsp?ftpID=4025>

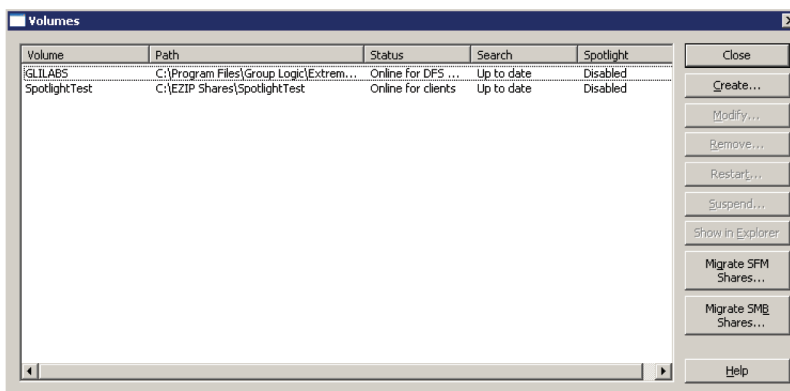
Enabling Network Spotlight in ExtremeZ-IP

Once Windows Search has been installed and configured to index the required ExtremeZ-IP shared volume locations, ExtremeZ-IP must be configured to support Network Spotlight

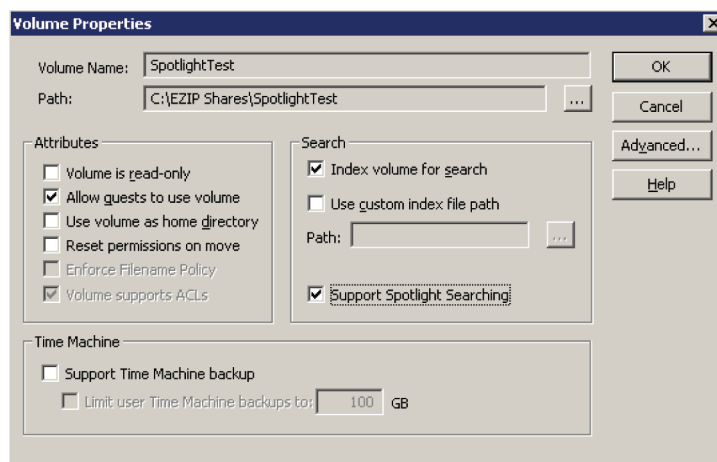


From the main ExtremeZ-IP Administrator window, select [Settings](#) and navigate to the [Search](#) settings tab. Check the [Support \'Spotlight Search\' Operations](#) option to enable support for Network Spotlight. This setting turns on Network Spotlight support but does not enable Network Spotlight for any shared volumes by default. If you would like all volumes on your server to support Network Spotlight, you can also check the [Support Spotlight Search on all volumes](#) option. If you only wish to enable Network Spotlight on a subset of your volumes, you can do so in those individual volumes' properties.

To enable Network Spotlight on a single volume, select **Volumes** from the main ExtremeZ-IP Administrator window.



The **Volumes** window displays all volumes being shared by ExtremeZ-IP. Each volume's current Spotlight search status is shown in the **Spotlight** column. Please note: The **Search** status column pertains to standard ExtremeZ-IP file search and not Spotlight. To enable Network Spotlight for a volume, select the volume and click **Modify**.



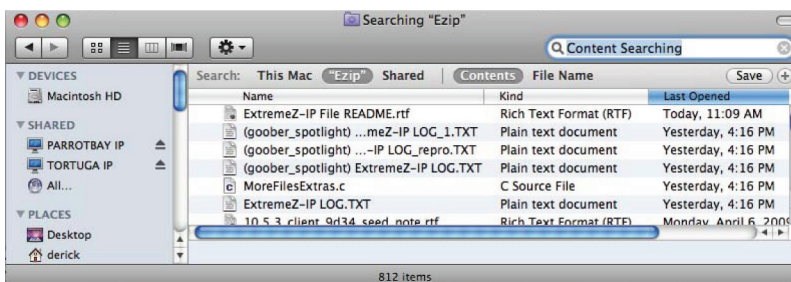
In the **Volume Properties** dialog, check **Support Spotlight Searching** and click **OK** to apply the change. The volume's **Spotlight** status will change to Online. If Windows Search is not installed or is disabled, the status will indicate that Windows Search is not running. If the volume folder is not included as a Windows Search indexing location, the status will continue to show as Online. Please ensure that the location was configured for indexing as described in the Installing Windows Search on the File Server section above.

Once the volume is Online for Spotlight searching, any new Mac OS X AFP client sessions that connect to the volume will default to Network Spotlight file search through the Finder. If a client was connected to the ExtremeZ-IP server prior to Network Spotlight being enabled, the client will have to disconnect from all volumes on that server and reconnect with a new session before Spotlight searches will be possible.

Known Issues and Limitations

Supported search terms

The search interface in the Macintosh Finder allows for a variety of search terms to be selected. This interface permits a search to be constrained using any combination of these search terms. ExtremeZ-IP Network Spotlight supports these search terms: Kind, Last Opened Date, Last Modified Date, Created Date, Name, Contents, and Size. The Finder provides an extensive list of “Other...” search terms, such as “Audio bit rate” and “Genre”. These, and all other search terms not listed above, are not supported by ExtremeZ-IP Network Spotlight and will simply be ignored when determining search results.



Files without .xxx extensions are not searchable by content

Because Windows Search uses file extensions to determine if a file's contents can be indexed, files that do not include an extension in their name will not be indexed for content. These files will continue to be searchable by filename and other properties. The creation of extension-less filenames by Macintosh applications is more common than with PC applications. If content search is important, Macintosh users should be encouraged to include file extensions when creating files.

Delay between file changes and file indexing

When files are added, moved, or deleted from a shared volume, Windows Search is notified of these events and proceeds to update its index to reflect the changes. Windows Search attempts to operate at a lower priority than critical Windows processes so that it has a minimal impact on server performance. If a server is busy, there can be a delay between the time a file is changed and the time it is updated in the Windows Search index. Under high server load, it could conceivably take significant time for a file change to be indexed. While this delay is usually a few seconds or less, Mac clients may observe that files they've modified are not immediately found by Network Spotlight. ExtremeZ-IP is only able to return these files as search results after Windows Search has had a chance to index them.

Configuring Windows Search for higher priority indexing behavior

By default, Windows Search will throttle down the aggressiveness of its indexing as general server load increases. If you find that search indexing is frequently not able to keep up with file changes, Windows Search allows its index throttling behavior to be modified. To configure Windows Search to maintain full speed indexing, regardless of server load, set the value of the registry key (HKEY_LOCAL_MACHINE\SOFTWARE\Policies\Microsoft\DisableBackoff) to 1.

By default, some file types are not included in Windows Search results

Indexing of some less common file types is disabled by default in Windows Search. Any files ending with the extension associated with these file types will not be indexed and will not be reported in Network Spotlight search results.

If Macintosh users need the ability to find files of these types, the file types need to be enabled in the File Types tab of the advanced Indexing Options. This is covered in greater detail in the Installing Windows Search on the File Server section above.

Differences in file type categories between the Mac OS and Windows OS

The Mac OS and Windows OS categorize types of files (i.e. Music, Documents, Images) in somewhat different ways. Because of small differences in the way these categories are defined, you may observe slightly different behavior when searching the local Macintosh hard drive vs. searching an ExtremeZ-IP volume with Network Spotlight. An example of this occurs with music files purchased from the Apple iTunes store. This file type is included in local Macintosh search results for files of Kind = "Music", but is not considered a Music file by Windows Search. For this reason, files of this type are not included in Network Spotlight search results for Kind = "Music".

Searching network-based Mac OS home directories results in local hard drive searches

As of Mac OS X 10.5.6, Spotlight search requests made for network-based home directories are directed by the Mac to the local hard drive, and not to the ExtremeZ-IP based home directory as would be expected. This is a limitation of the Mac OS.

Network Spotlight searches are not dynamically updated after completion

When using the Finder to do a standard search on a local Macintosh hard drive, search results are continuously updated as files meeting the search terms are added to or removed from the drive. This live updating of search results does not happen with ExtremeZ-IP Network Spotlight support. The results presented to the user reflect the state of the shared volume only at the time the search was issued. Any files added, moved, or deleted after the search completes will not be reflected until the client reissues the search request.

File size limitation to content indexing

When Windows Search indexes the content of a large file, it restricts the indexed portion of the content to approximately the first 1.5 MB of the file. This is done so that very large files do not overwhelm search indexing. For this reason, Macintosh clients using ExtremeZ-IP Network Spotlight search will not be able to find files by searching on content that exists beyond 1.5 MB into the file.

Some search term characters are interpreted differently by the Mac OS and Windows Search

Noise characters (period, underscore, percent, pound) act as word breaks in Windows Search queries. Searching on "file.doc" will find files whose first word contains "file" and whose second word contains "doc". In this case, the search might return these results: file.doc, filebackup.documentation, filed document, file_document. The search results should contain the file the user intended to locate but could contain additional, unexpected results that are considered valid by Windows Search.

About Group Logic

Group Logic, Inc. (GLI) is a leading provider of digital content-driven collaboration solutions for the enterprise and the cloud. With over 20 years of unmatched experience, Group Logic's emphasis on customer success is the very core of its business. More than 4,500 customers trust Group Logic every day to access, share, and extend their digital content investments around the world. For more information, visit Group Logic on the Web at www.grouplogic.com or call **800.476.8781 / +1.703.528.1555**.