CISearch.NET Setup Guide

Table of Contents **Page No.**

Installing - Introduction	2
Installing - Planning	2
Installing - Prerequisites	3
Installing - Running the setup routine	4
Installing - Setting up CISearch.NET in IIS	7
Installing - Trusted for delegation setting in Active Directory	7
Installing - Important note for Windows 2003	8
Installing - Setting up SQL Express (if needed)	8
Installing – CISearch and Microsoft Exchange	9
Configuring - Tuning Indexing Service Performance	11
Configuring - Quick Start Guide	13
Trouble shooting - Known Issues and resolutions	14

Introduction

Welcome to the setup guide for CI Search. CI Search is in our opinion the best document search engine in the world. CI Search is designed to give users a powerful centralized search facility where they can search through the contents of millions of documents on their computer system in a matter of seconds. It is capable of cataloging and indexing both content stored in NTFS file systems and within MS Exchange server.

It is perfectly placed for organizations that want a powerful document search and retrieval system but do not want to pay thousands for a solution like Share Point (and we actually think CI Search beats Share Point's searching hands down anyway).

CI Search is designed to be installed on any Windows 2000 or 2003 server based operating system. That includes the 64bit variants. It is not currently certified for Vista or 2008.

CI Search fully respects NTFS file and folder AND exchange mailbox permissions, only showing users the files they have the rights to open.

It also features unique search optimization features that make it one of the most accurate search engines around. Pretty much all other Microsoft Indexing service based products just submit a user's query to Indexing Services and expect to get accurate results back.

CI Search is very different, it actually re-works a users search analyzing in depth what the user has entered and pre-processing it before the search is submitted. The full breadth of the Indexing Service query language is used to optimize user searches without users even knowing.

If however you want to manually use some of the Indexing Service language in your query CI Search will detect this and automatically turn off its optimization features for that search.

The most important step, planning your install

How you deploy CI Search heavily depends on your environment. You need to decide what is best for you and your end users, and due to its flexibility there are actually many ways that CI Search can be deployed.

The easiest and most common installation is to have one central server nominated as your indexing and cataloging server. CISearch is then installed on that machine and configured to catalog the content on all of your other file and Exchange servers.

This method has the added advantage that in large installations you can have a dedicated machine tuned for high responsiveness to user searches, while not putting any extra load on your existing servers.

This doesn't mean that you have to deploy CI Search like that though. If you only have a few servers CISearch can normally reside happily on one of them without the requirement for a dedicated box of its own.

There are a few things that need to be borne in mind though when installing CI Search…

1) Disk space – The catalogs for CI Search will take up a small amount of space. As a guide, the location that you install CI Search to (Re. volume) should have 750mb free for every 500,000 files (Word and Excel type files) that you wish to index.

2) Hardware – Try to use reasonably up to date hardware to run CI Search on. Again as a guide, a single processor PIII 1Ghz server with 1GB of ram, a suitable raid array with sufficient disk space and gigabit network connection will suffice for average use by 200 users.

3) Cataloging files takes bandwidth – CI Search is fully tunable through Indexing Service itself, but we would not recommend trying to catalog the contents of a server over a WAN link. It would be far better to install an extra copy of CI Search on the remote server.

4) Domain Controller local security policies – The only place that you cannon install CISearch is on a domain controller. The default security policies will prevent it from working properly.

For further guidance on configuring and performance tuning Indexing Service once you have installed CI Search please see the Configuration and Tuning guide.

Prerequisites

Before starting the install you need to ensure the following prerequisites have been met. The server you are installing CISearch on must have:-

1) The .NET Framework Version 2.0 installed
2) Indexing Service installed and running
2) IIS (Internet Information Services) installed and running
3) Windows 2000 Server SP4 or later
4) An available SQL 2000 or later instance (SQL 2005 Express recomended).

The Microsoft .NET framework 2.0 package is included as part of the install package. If is needed you will be prompted to install it. You may also wish to install SQL Express 2005 if you do not have an SQL server

You must also check that Indexing Service and IIS are installed, if they are not installed then you can add them through the “Add/Remove Programs“ icon in the Control Panel.

Double click on “Add/Remove Programs“ (start -> settings -> control panel), click on “Add/Remove Windows Components”, and then ensure “Indexing Service” and “Internet Information Services (IIS)” are both ticked. Highlight “Internet Information Services (IIS)” and click on the “Details” button.

Ensure that at least the following options are ticked, "Internet Services Manager" (In windows 2000 this is called "Internet Information Services snap-in”), "Common files", and "World Wide Web Server". Click on “OK”, then “Next”. If any components were missing then windows will now install them. You may be asked for your Windows CD.

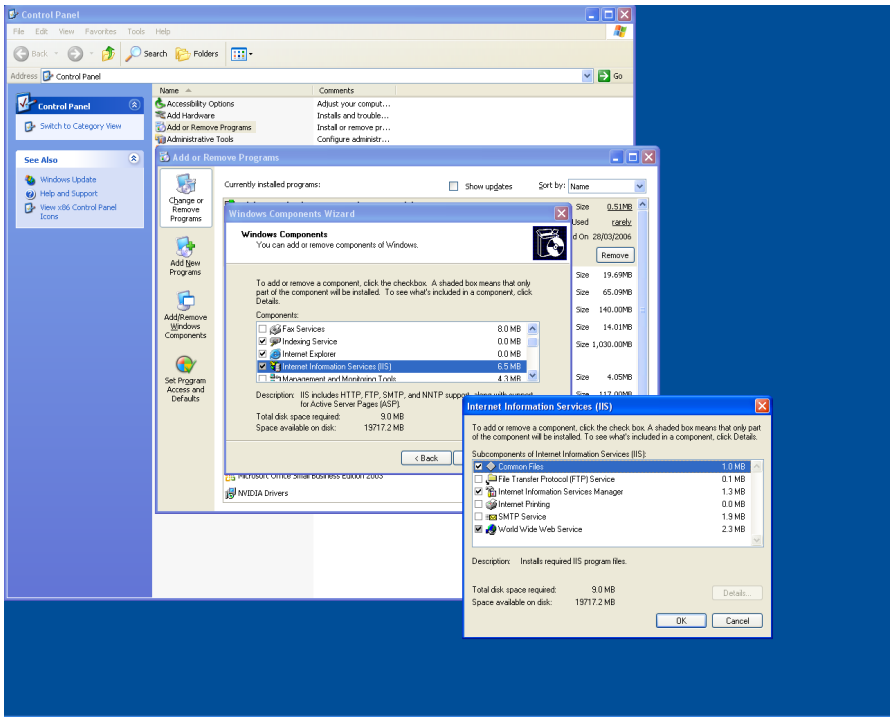

Running the setup routine

The setup routine is very straightforward and is split into several sections. First run the setup.exe file to begin the installation routine, and then follow the onscreen prompts to complete the install.

One point that you have to consider though is where you install CISearch too. This is because the Indexing Service catalogs are installed to the same location as the application and they can take up a noticeable amount of space.

As a rule of thumb you should allow 750mb on the drive CISearch is installed on for every 500,000 files that you will be cataloging.

Configuration Database Setup:

All of CISearch's configuration information is stored in a small SQL database. The database setup screen will display after the main part of the setup has run.

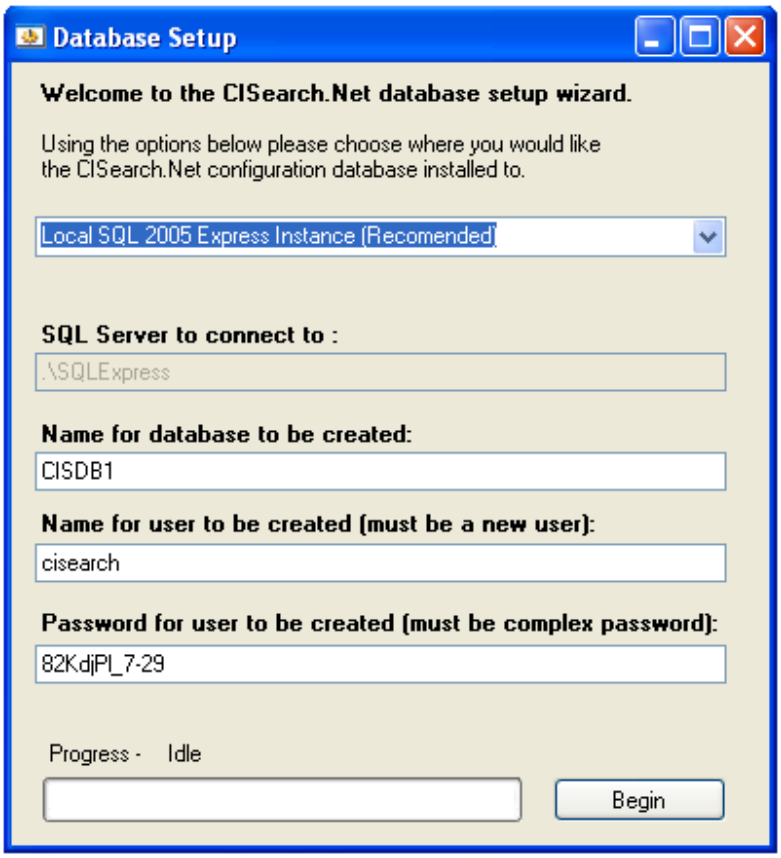

Setup Wizard:

The next stage of the setup process is the CISearch Setup Wizard. The wizard automates the process of setting up the basics, prompting you for information in a step by step way.

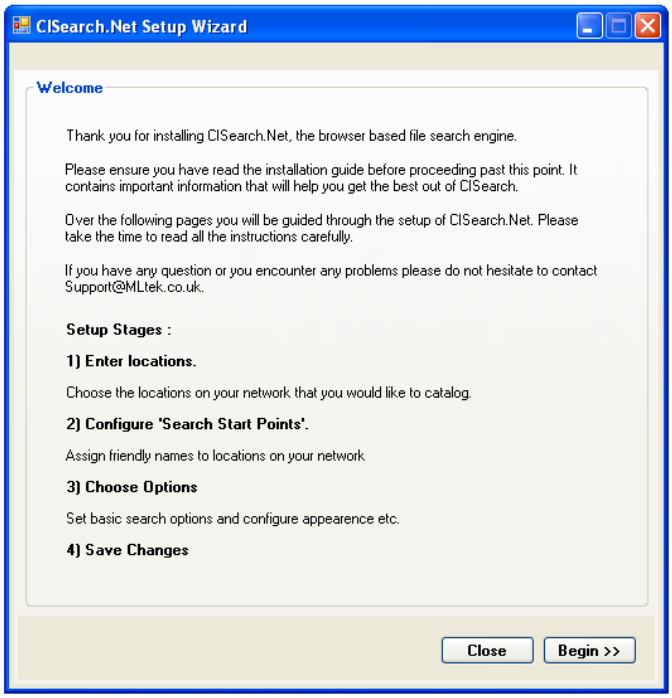

First you must enter details of the locations on your network that you wish to make searchable. This is done on the screen below. Use to browse button to select the locations via their UNC path (e.g. \\YourServer\Share). For each location you must also enter the details of the account you wish Indexing Service to use when accessing the location (e.g. YourDomain\Username).

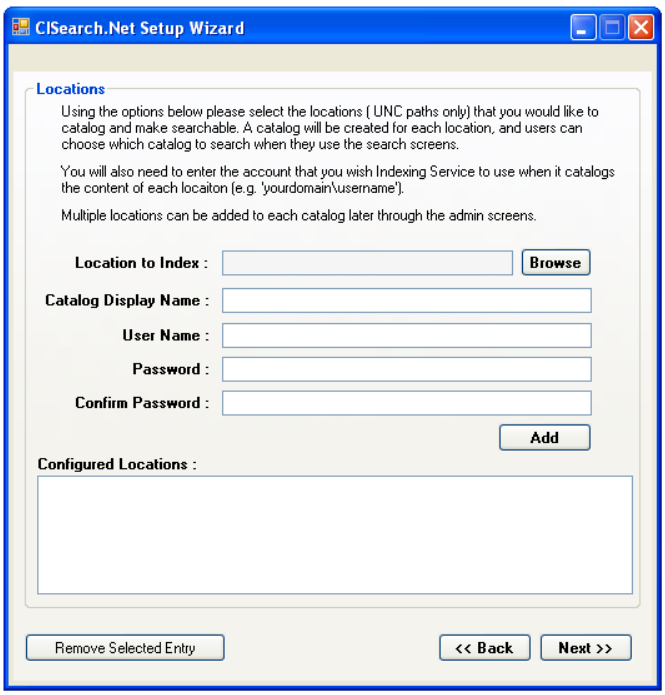

After specifying the locations that you would like to make searchable, you are then given the opportunity to set up what we call ‘search start points’, a feature unique to CISearch.

When user chooses a catalog to search, the locations indexed by that catalog will be compared to the stored list of Search Start Points.

If one or more Search Start Points are found that relate to the chosen catalog then they will be displayed to the end user via a second drop down box. The user can then choose to restrict their search to that location.

For example, if you had a share called ‘Files’ on a server called ‘Server1’ (UNC path = \\Server1\Files), and folders called ‘Sales’, ‘Support’ and ‘Billing’ in that share….

You would add the UNC path ‘\\Server1\Files’ to the locations you wanted to make searchable, and then add the following Search Start Points.

1- Path = ‘\\Server1\Files\Sales’ Name = ‘Sales’
2- Path = ‘\\Server1\Files\Support’ Name = ‘Support’
3- Path = ‘\\Server1\Files\Billing’ Name = ‘Billing’

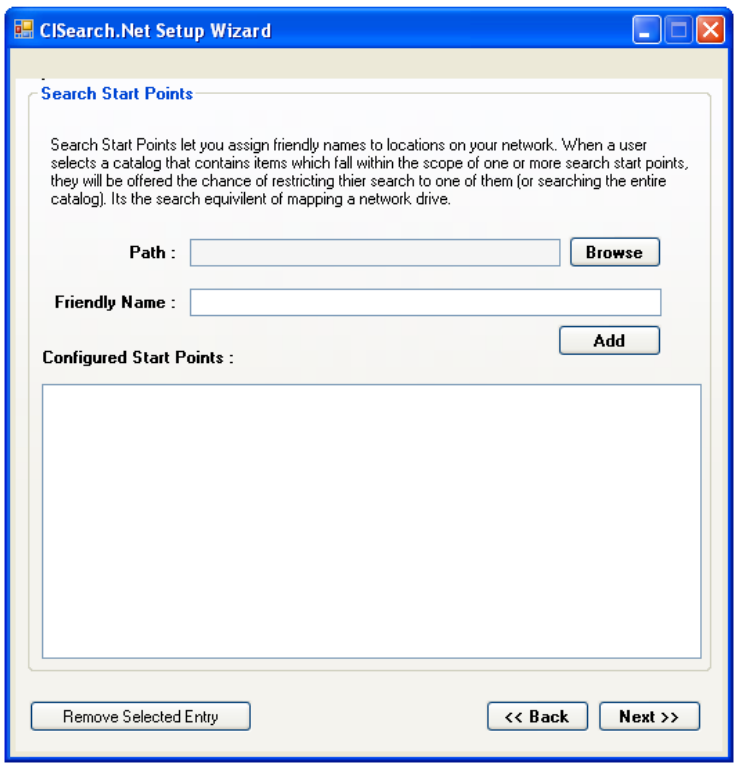

General Options:

The last screen of the setup wizard is the general options screen. Here you can choose from various options that let you control the look of CISearch and effect its default behavior.

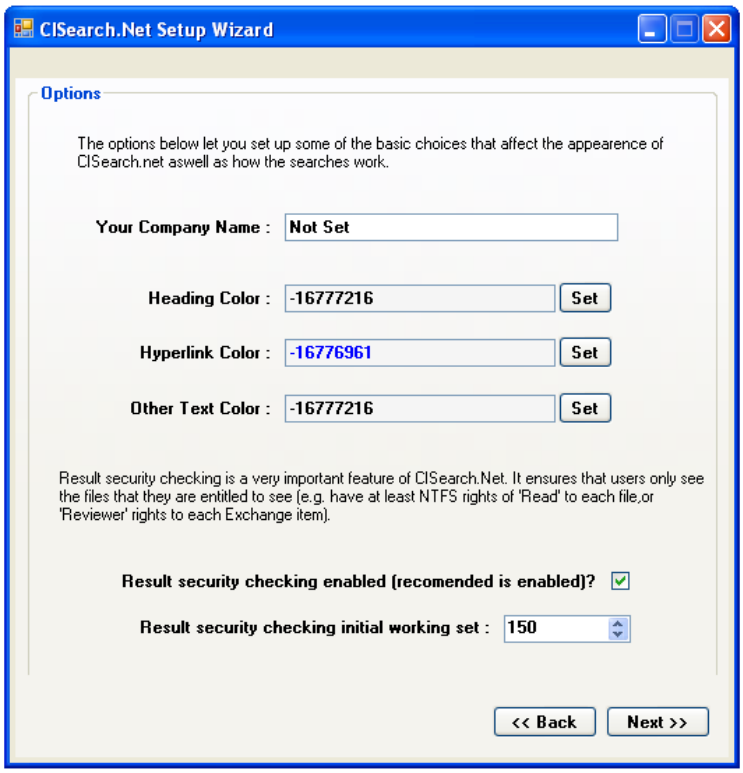

Setting up CISearch.NET in IIS

The install routine will automatically configure CISearch.NET for you as a virtual directory and application in IIS. It is worth checking however that the following settings are in place after the install.

1) The CISearch.Net virtual directory should be set to use ASP.NET version 2.0.50727.
2) Only ‘Windows Integrated’ authentication should be configured.
3) The CISearch virtual directory is configured as an application

Trusted for delegation setting in Active Directory

CISearch makes extensive use of impersonation and delegation. To do this the computer that CISearch is installed on has to be set as ‘Trusted for Delegation’ in Active Directory.

Open Active Directory Computers and Users and locate the computer account for the machine that CISearch is installed on. Right click on it and then click ‘Properties’.

Ensure that the 'Trust this computer for delegation’ option is ticked.

You may need to wait for a few hours for this change to replicate around your domain controllers. Until it has taken effect you will see such side effects as not getting any results with result security checking enabled, and paths being marked as unreachable when you try to set them up in ‘Search Start Point’ admin.

Important note for Windows 2003

By default Windows 2003 doesn't allow .asp files to run under IIS, and so you have to tell it to allow them. After running the CISearch.net install routine :-

1) Right click on the "My Computer" icon on the desktop
2) Left click on "Manage"
3) Expand out "Services and Applications"
4) Expand out "Internet Information Services (IIS) Manager"
5) Click on "Web Service Extensions"
6) Click on "Active Server Pages" and click on the "Allow" button, then if you have an entry for 'ASP.NET v2.0.50727' select 'Allow' for that too.*

*If the .Net framework 2.0 was installed before IIS then the 'ASP.NET v2.0.50727' entry may not be listed. If this is the case then the CISearch.Net installation routine will correct the problem, but you will have to allow the 'ASP.NET v2.0.50727' extensions after the installation has completed.

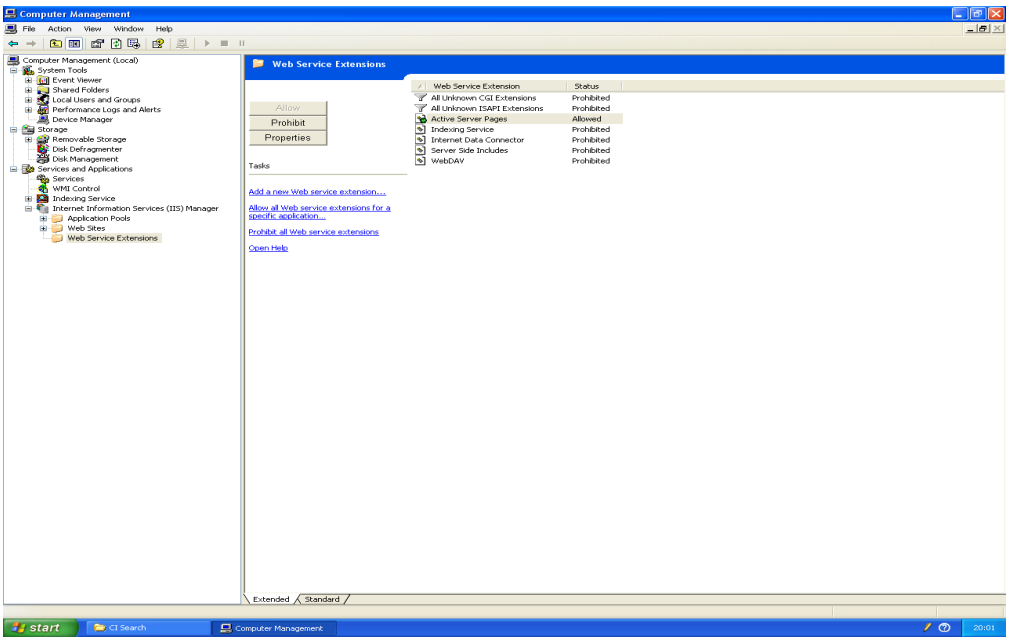

Setting up SQL Express (if needed)

SQL Express 2005 is Microsoft's free lightweight version of the full SQL 2005 product. It features a lot of the functionality of the full version and replaces the previous MSDE utility.

If an application wishes to use a local database on the machine it is installed on and be compatible with Microsoft's future OS releases then it has to move to a SQL Express. Microsoft have not provided a 64 bit compatible version of its client network access libraries (MDAC) which forces developers to use SQL Express instead of the ever popular MS Access database / MDAC combination.

SQL Express handles security slightly differently from MS Access. As well as specified username/password combinations (SQL authentication), SQL Express also supports windows authentication.

By default only windows authentication is turned on which is not very practical. The CISearch.Net setup routine will detect if this is the case and if so ask if you want to enable SQL authentication as well. You should say 'Yes' or the install will fail.

CISearch and Microsoft Exchange

Not only can CISearch catalog and index NTFS file systems, it can also catalog and index mail items in both Exchange mailbox and public folder stores.

Up until now we have just dealt with cataloging and indexing an NTFS file system because before you catalog and index your exchange environment you need to perform a few steps.

Firstly you need to enable EXIFS on your Exchange server. When you do this it is very important that you follow all of the steps below.

Step 1 - Enabling EXIFS :

EXIFS was introduced by Microsoft with Exchange 2000. It enables Exchange to present the information in its stores as NTFS files and folders. By default it is enabled on Exchange 2000 but not on Exchange 2003 and above.

Unfortunately, people didn't appreciate how to use it properly and there were some problems with certain Antivirus software products which all led to a significant number of support calls logged with Microsoft by people with corrupt exchange stores after enabling EXIFS.

As long as you follow the steps below it is perfectly safe to enable and utilize EXIFS for CISearch.

EXIFS normally presents itself as the 'M:\' drive, but if the M:\ drive is in use already it will use the next available drive letter) you can enable it through the following registry key after which you should restart the Microsoft Exchange Information Store service..

Location: HKLM\System\CurrentControlSet\Services\ExIFS\Parameters
Value: DriveLetter
Type: REG_SZ
Value: *Drive letter for ExIFS (EG 'M', do not include the ':\', just enter the letter)*

Once you have enabled EXIFS you must make sure that the M:\ drive (We will assume you used the letter 'M') has been excluded from ALL virus scanning, including scheduled and on access scanning.

Also, do not try to compress, encrypt or change the security on any folders/files through the M:\ drive. It should be thought of as a read only window into the Exchange store.

Step 2 – Creating the read only shares :

Open Windows Explorer and browse to the M:\ drive. At the root you will see a folder who's name is the same as you internal mail domain. Inside that folder you will see two more folders, one called 'MBX' and one called 'PUBLIC FOLDERS'.

Both folders should be shared out with their default names, but you MUST set the permissions on the share (not the folder) to 'Everyone – Read Only'. By setting the permission on the share to read only you are ensuring that this is the highest level of permissions anyone will get when accessing the share over the network.

See the screenshot on below for an illustration of how to set the permissions on each share.

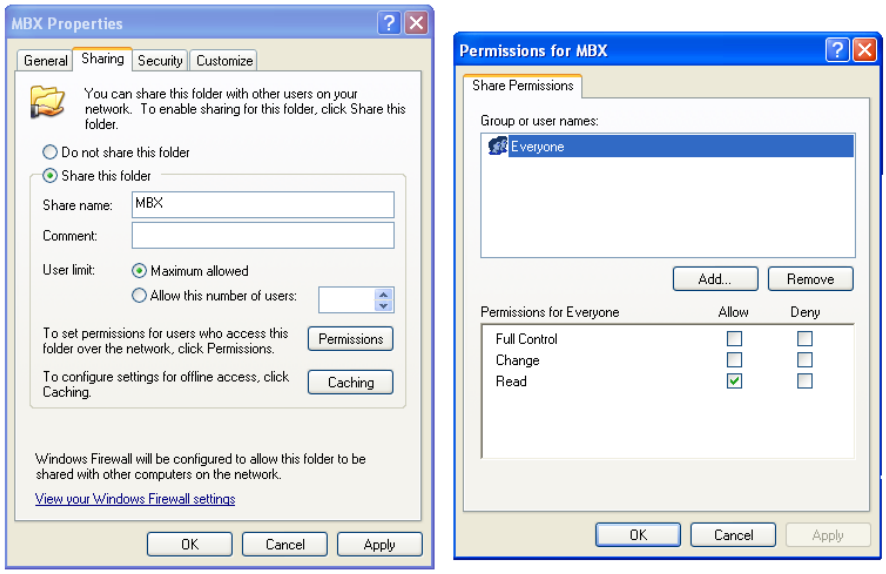

Step 3 - Running the CISearch Client Setup on all client computers

A recently released Microsoft security update effectively blocked the ability of Internet explorer to open hyperlinks that point to .eml files. See http:// support.microsoft.com / default.aspx?scid=kb;EN-US;825803 for further details.

The CISearch Client Setup routine changes the way that Internet Explorer handles .eml files and overcomes the limitations imposed by this security update.

Microsoft Exchange server presents email content as .EML files through EXIFS, and as such for clients to be able to open Emails indexed by CISearch through the web interface the Client Setup routine needs to be run on them.

The setup routine is supplied as a single silent installing .MSI file so that it can easily be deployed through Active Directory group policies or be run manually/through a login script. The CISearch Client Setup.MSI can be found in the CISearch program folder.

Step 4 – Internet Explorer 7

As the EXIFS drive is not an actual NTFS governed volume you can get .eml files with characters in there name that are normally reserved and this can cause problems when you try to open them.

Internet Explorer 7 can overcome his problem, but earlier versions of Internet Explorer cannot. Therefore we recommend that if you do index and catalog your Exchange server any client machines that use CISearch have Internet Explorer 7 installed.

Step 5 - Finally

Once the 3 steps above have been completed you can add your ‘MBX’ and ‘PUBLIC FOLDER’ shares on your exchange server to CISearch just like you would any other network share.

Tuning the performance of Indexing Service

Indexing service itself can be tuned to set how vigorously it keeps up with the creation of new files and the changing of existing files in shares it is watching.

If you are running with your hardware already quite stressed it is a good idea to start at the lower settings, and then work your way over to the higher ones, watching your disk queues with performance monitor over a period of time.

To tune the performance of Indexing Service right click on your “My computer” Icon and then Click on “Manage”. Expand out “Services and Applications” and then right click on “Indexing Service”. Click on “All Tasks” and the click on “Tune Performance”.

If “Tune Performance” is greyed out you will need to stop Indexing Service before you can access it. You can do that by right clicking on “Indexing service” and clicking on stop. Once you have finished making adjustments you will need to right click on “Indexing Service” and click on “Start” to start Indexing Service again.

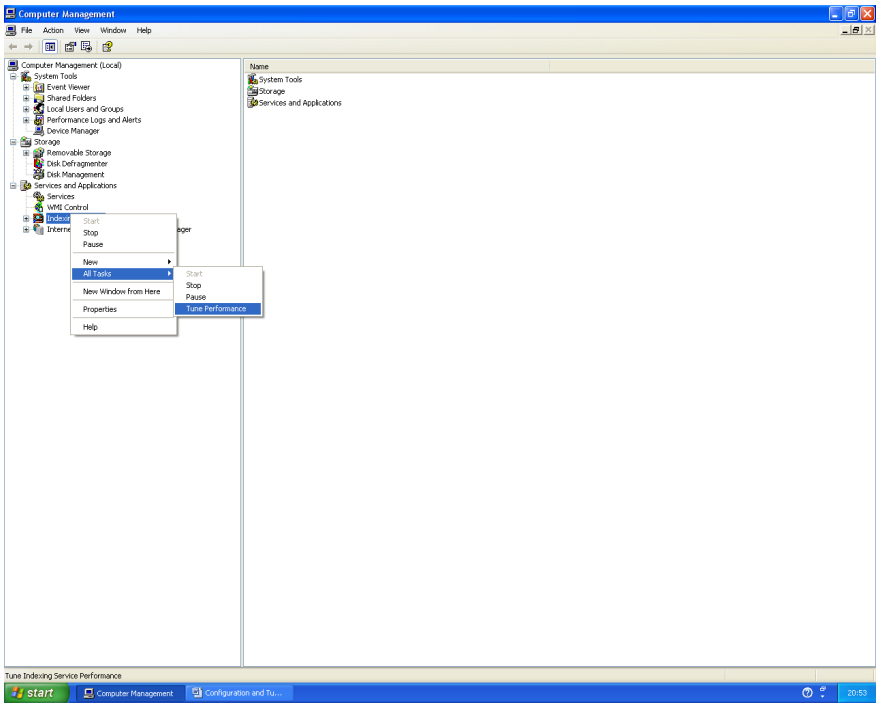

You will now be presented with the “Indexing Service Usage” box. It depends on your hardware and the number of queries being submitted, but as a rough guide:

For 850,000 documents or more choose ‘Dedicated’
For 350,000 -> 849,999 documents choose ‘Used Often’
For up to 349,999 documents choose ‘Used Occasionally’

To fine tune the service choose ‘Customize’

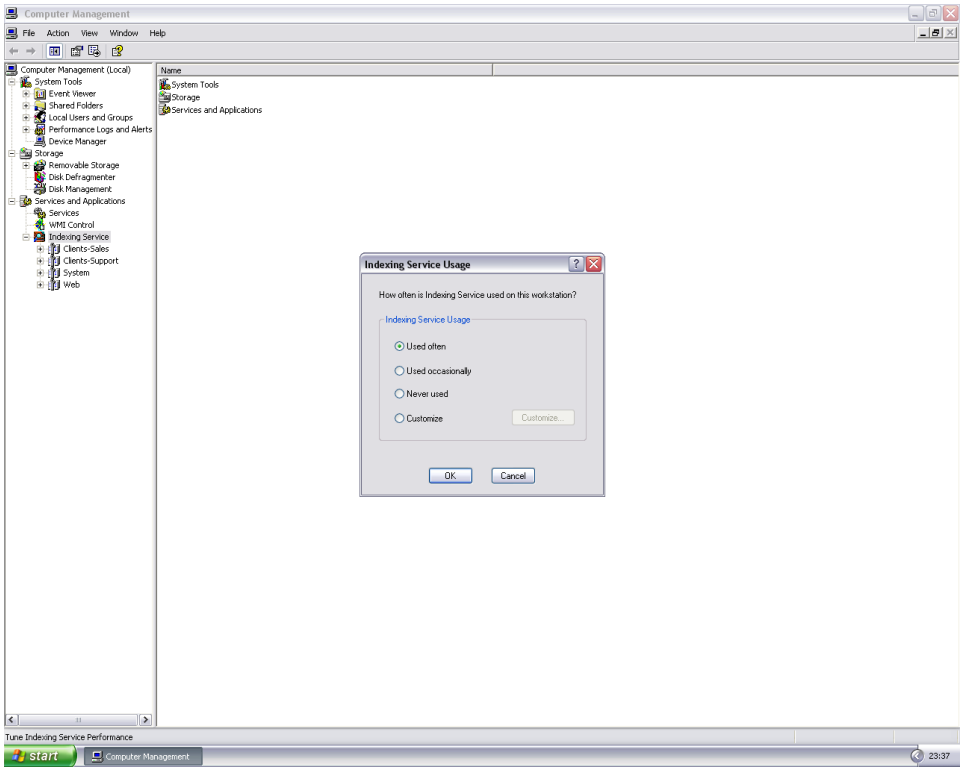
Computer Management
File Action View Window Help
Computer Management (Local)
System Tools
Event Viewer
Shared Folders
Local Users and Groups
Performance Logs and Alerts
Device Manager
Storage
Removable Storage
Disk Defragmenter
Disk Management
Services and Applications
Services
WMI Control
Indexing Service
Clients-Sales
Clients-Support
System
Web
Name
System Tools
Storage
Services and Applications
Indexing Service Usage
How often is Indexing Service used on this workstation?
Indexing Service Usage
Used often
Used occasionally
Never used
Customize
Customize...
OK
Cancel
Tune Indexing Service Performance
start
Computer Management
23:37

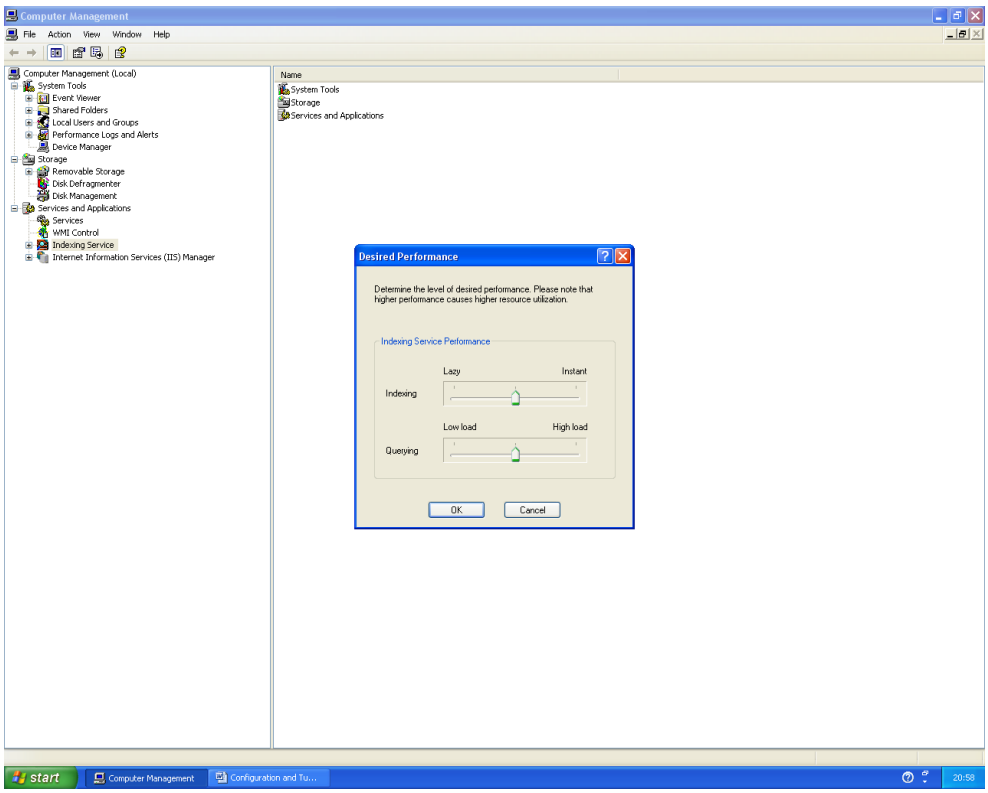
Computer Management
File Action View Window Help
Computer Management (Local)
System Tools
Event Viewer
Shared Folders
Local Users and Groups
Performance Logs and Alerts
Device Manager
Storage
Removable Storage
Disk Defragmenter
Disk Management
Services and Applications
Services
WMI Control
Indexing Service
Internet Information Services (IIS) Manager
Name
System Tools
Storage
Services and Applications
Desired Performance
Determine the level of desired performance. Please note that higher performance causes higher resource utilization.
Indexing Service Performance
Lazy
Instant
Indexing
Low load
High load
Querying
OK
Cancel
start
Computer Management
Configuration and Tu...
20:58

Quick Start Guide.

Getting up and running with CISearch.Net is both quick and simple. We will assume that you have already downloaded the setup file (http://www.mltek.co.uk/downloads/cisearch/setup.exe) and have had a quick read of the preceding instructions.

Below you will find a simple bulleted list of steps to undertake to get up and running, starting with the Install itself

First however you must ensure that you are not installing CISearch.Net on a domain controller and that the computer account for the server that is going to host CISearch.net is set as 'Trusted for delegation' in Active Directory.

- Double click on the 'Setup.exe' file to begin the installation routine.
- After any prerequisites are installed if needed proceed through the installation.
- When prompted to choose an install location ensure you choose some where suitable (750mb free space for every 500,000 office documents as a rough guide)
- After the install completes the 'Admin' screen will open
- Set an admin password
- Log in
- Go to 'Catalog Admin' and create a new catalog
- Add locations on your network to the new catalog using their UNC paths. Don't forget to specify an account for Indexing Service to Index each location with as you add each location.
- If you desire go into 'Search Start Point Admin' and add in 'start points'. For example if you are adding a location called '\\server1\data\clients' to a catalog you may wish to add :

  Path1: '\\server1\data\clients\new clients' *Desc:* 'New Clients'
  Path2: '\\server1\data\cliients\ex clients' *Desc:* 'Ex-Clients'

  Your users would then be able to choose to search 'All available locations' *(in the catalog)*, 'New Clients' and 'Ex-Clients'.

There are many more things you can do, like alter the appearance of CI Search.Net on the options page and more, but that is in the simplest possible deployment all you need to do to have your own network wide search facility.

Known Issues and Resolutions.

The vast majority of problems with CISearch.Net can be traced to problems with the .net framework version 2 itself.

The known issues are:

1) **'*The application has attempted to perform an action not allowed by the security policy*'** *(error message on load)*

This error message is displayed because of a miss-configuration in the .net framework. The easiest way to resolve this is to uninstall the .Net framework version 2, delete the contents of the 'C:\WINNT\Microsoft.NET\Framework\v2.0.50727' folder, then re-install the .Net framework version 2.

After re-installing the .Net framework you will need to make sure that CISearch.Net's virtual site in IIS is set to use the .net framework version 2 as it may have reverted to the .Net framework version 1.0 or 1.1 .

2) '*Documents are not appearing in the results of CI Search*' *(generic issue)*

This can be for several reasons including:

a) The user you are logged in as does not have at least NTFS 'read' rights to the file you are searching for and security checking is enabled.
b) The criteria you are searching with are wrong
c) The file has not been indexed yet.

The first step in troubleshooting is to establish if the document you are looking for is truly in the catalog being searched or not.

First you need to find a document that you believe should be returned by CISearch.Net but isn't appearing in the results. Once you have done this open the document and find a sentence that is fairly unique to that document.

Turn off security checking in CISearch.Net's options screen and then search for the document you are looking for using a simple search, choosing the catalog you believe it to be in and entering the sentence from the document enclosed by quotes as the term to search for.

This will search for documents that contain that exact phrase.

If the document does not appear it is either because.

a) Indexing Service has not reached the document yet to index it
b) The location the document is in has not been set to be indexed as part of the chosen catalog

In the case of option 'a' you may wish to increase the allowable utilization of Indexing Service.

See the 'Tuning Indexing Service Performance' section for more details.

If the any catalog has more than 650,000 files in it, you may benefit from manually triggering a 'Merge' on the affected catalog. To do this follow the Steps in the 'Tuning Indexing Service Performance' section up to the point where you click on 'Tune Performance'. Instead right click on the catalog in question and click on 'Merge'.

3) No documents are returned for any search with security checking enabled, but they do appear with security checking disabled.

The Active Directory computer account of the server that CISearch is installed on MUST be set to 'Trusted for delegation'. If this is not done then the above situation will arise.

This situation can also arise if CISearch has been installed on a domain controller. Domain controller default local security policies will prevent impersonation and delegation working properly.

4) Other issues

Due to the almost infinite range of potential configurations of a Windows based server it is perfectly possible that you may encounter issues which we have not yet seen.

If this is the case the first recommendation is to try a repair of the .Net framework v2.0.

This can be done via Add/Remove programs in the control panel on the server hosting CISearch.Net.

A) Open the control panel and double click on ‘Add/Remove Programs’
B) Click on ‘Microsoft .Net Framework Version 2.0’
C) Click on the ‘Change/Remove’ button
D) When Prompted select ‘Repair’ and click ‘Next’ to begin the repair.

You may need to re-boot the system afterwards and please be aware that any software that uses the .Net framework version 2.0 on the server in question might suffer interruptions while the repair is taking place/until the server is re-booted.