Administration & Configuration

MASCOT

MATRIX SCIENCE

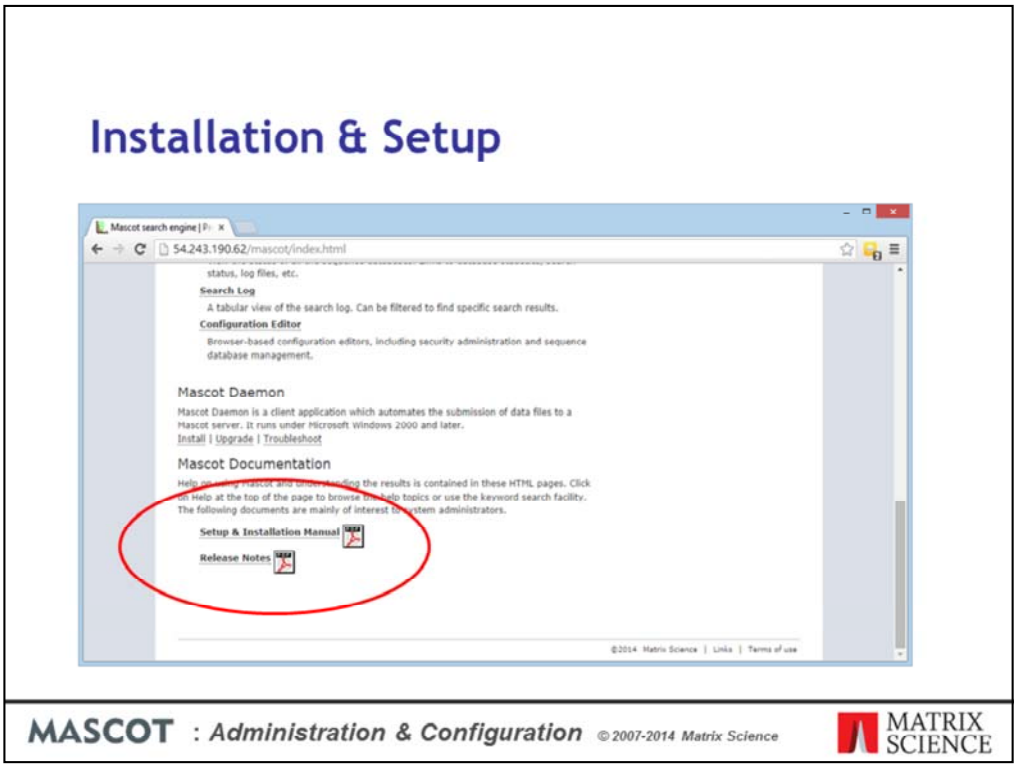

The Mascot Installation & Setup manual is linked from your local Mascot home page. If you need detailed information on any aspect of Mascot installation or configuration, this is the place to look

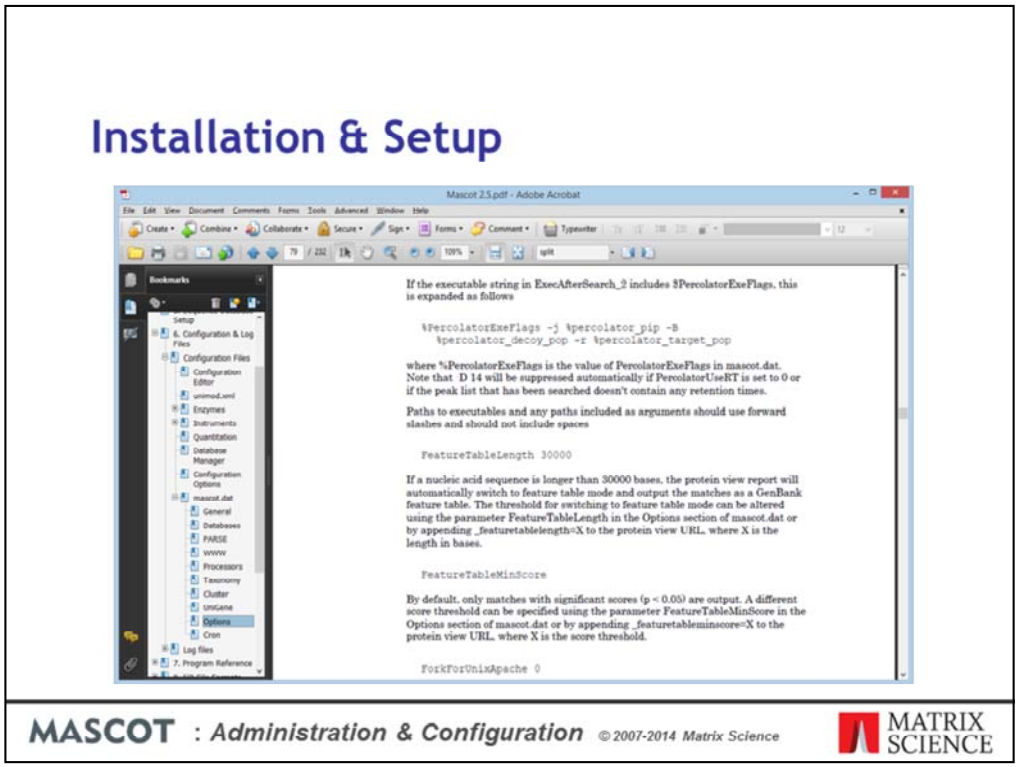

This is an administrator's manual, not a user manual. You won't find much relating to how to submit a search or how to interpret the results. User help is in the Mascot HTML pages.

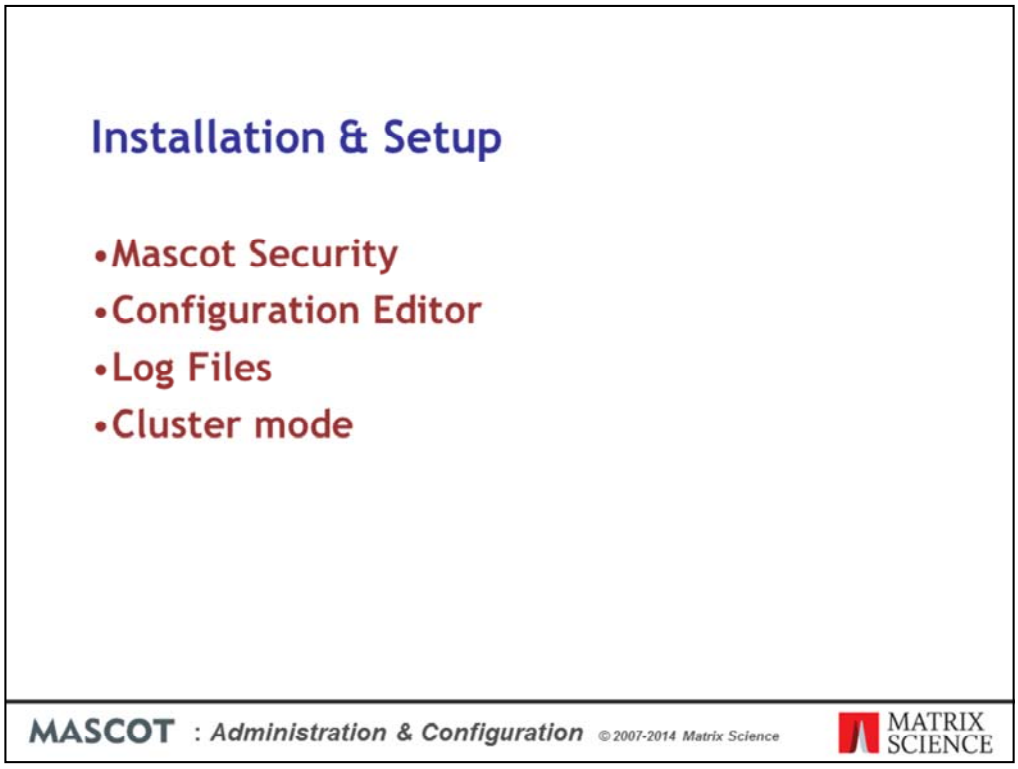

These are the topics we will cover in this presentation

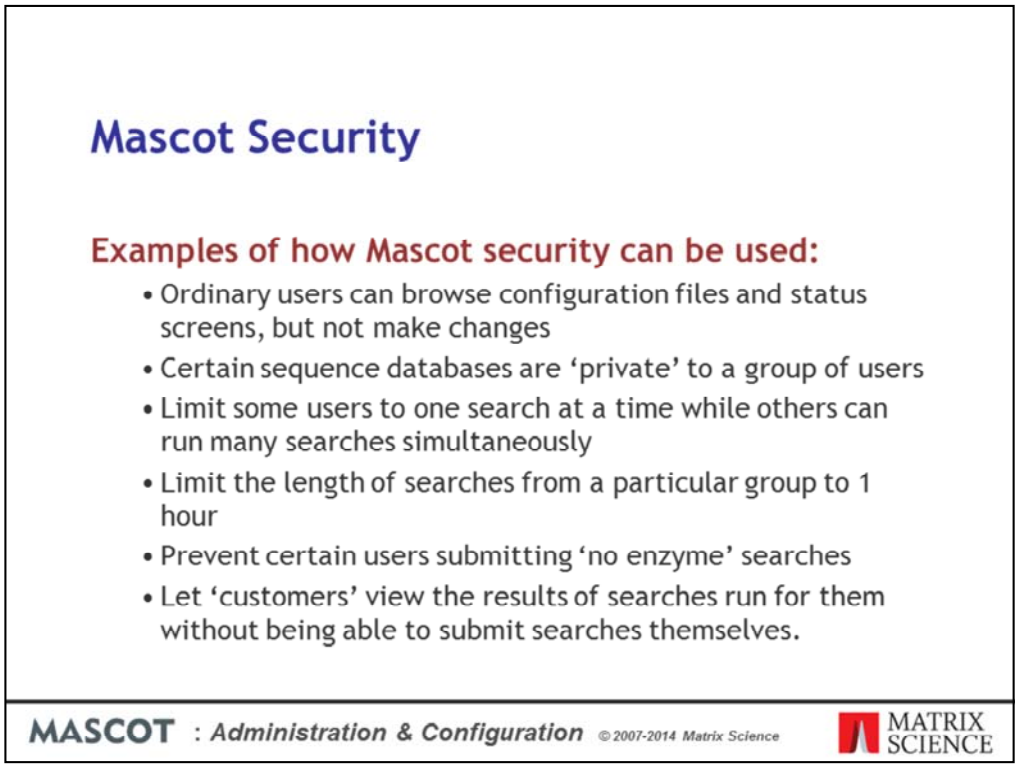

Mascot security is not a substitute for a firewall. It won't stop your server being hacked or infected by a virus. It is a way of managing and allocating the Mascot Server resources.

Mascot Security - enable / disable

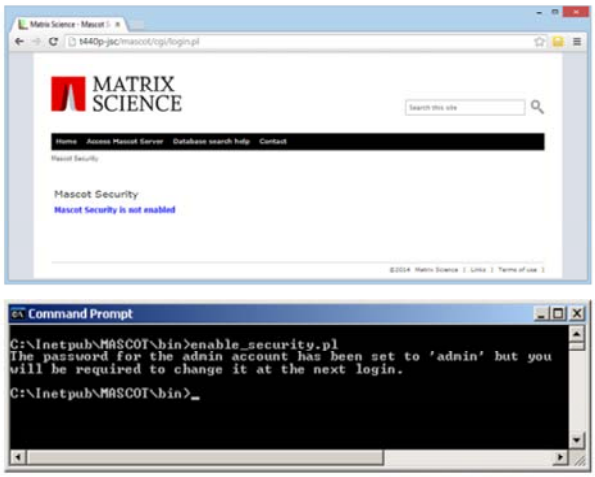

When Mascot is first installed, Mascot security is disabled. So, if you try to log in you will see this message.

To enable security, open a command prompt or a shell on the Mascot server, and change to the mascot/bin directory. Type enable_security.pl (or ./enable_security.pl if Linux).

This takes a few seconds. If you forget the administrator password, it can always be reset to ‘admin’ by running this script again.

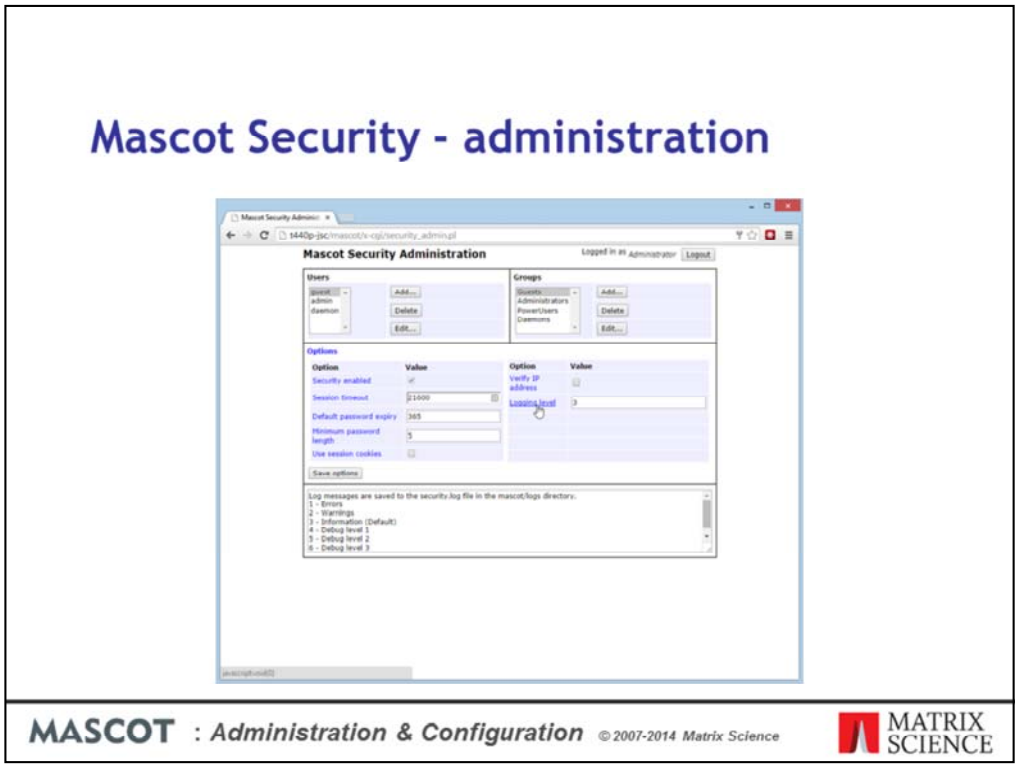

All security configuration is browser based. The top level page of the security administration utility looks like this.

There is a list of users, a list of groups and a few options.

If you hold the mouse over any of the blue text, help will appear at the bottom of the screen.

The options shown on this page are global options, which apply to all users.

Reference:

You can’t disable security from here - it has to be done on the server using the disable security command

The session timeout is in seconds. After this period of inactivity, the user will be required to login again.

The password expiry time is in days. After this period of time, the user will be required to enter a new password. Set to 0 to allow passwords to be permanent

Any new password must be at least this length.

Session cookies are automatically destroyed when the browser is closed. With some browsers, session cookies are not shared when a new instance of the browser is opened, which might mean that a user has to login again for each new window opened.

If verify the IP address is set, then any request to perform a privileged action will compare the IP address that the request is coming from with the one originally used to login.

Logging level should normally be left at 3

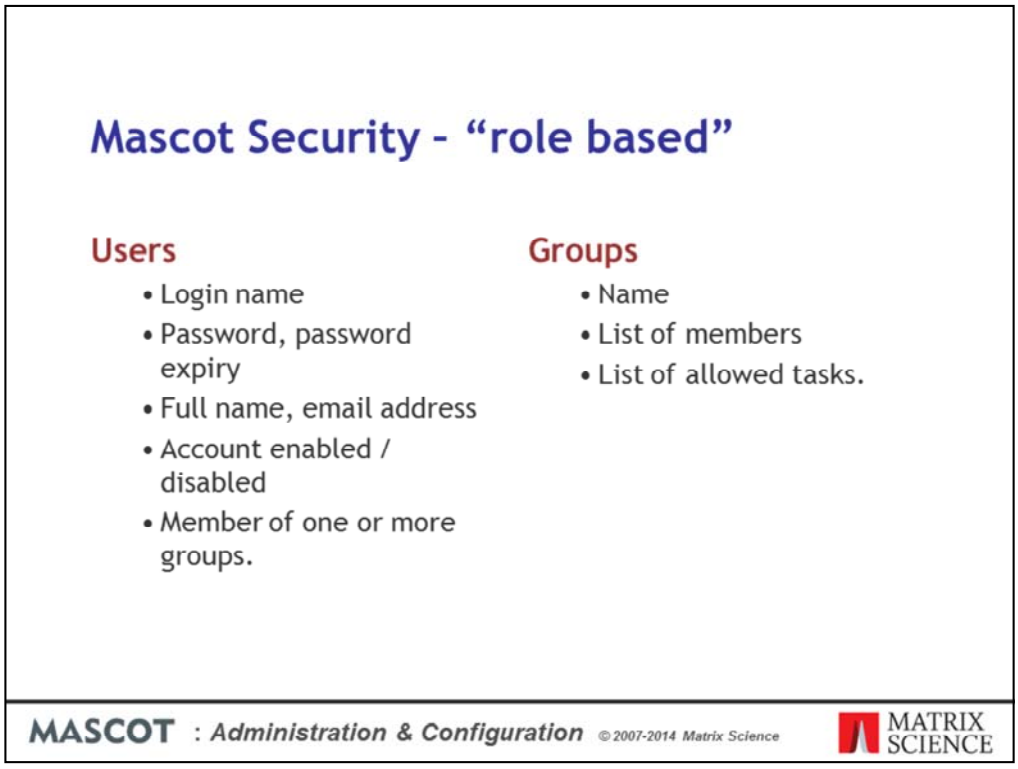

Mascot security is "role based". This means that privileges, known as tasks, are assigned to groups, not individual users.

Users gain these privileges by being members of one or more groups.

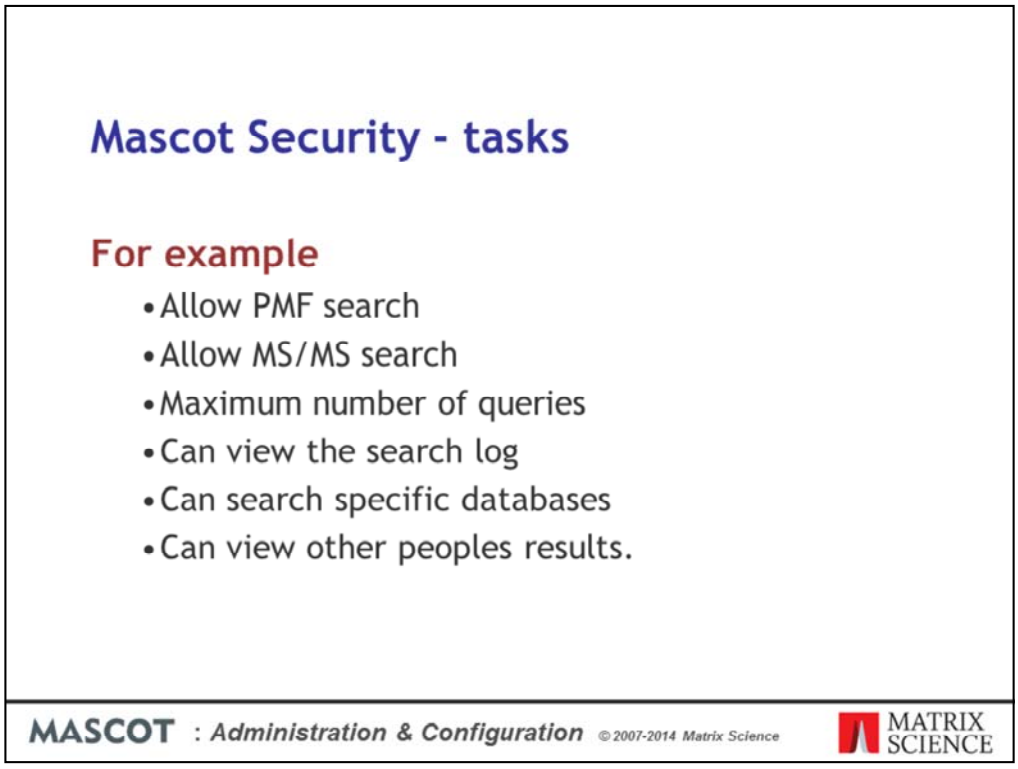

There are 30 different tasks that members of a group can be allowed to perform - for example:

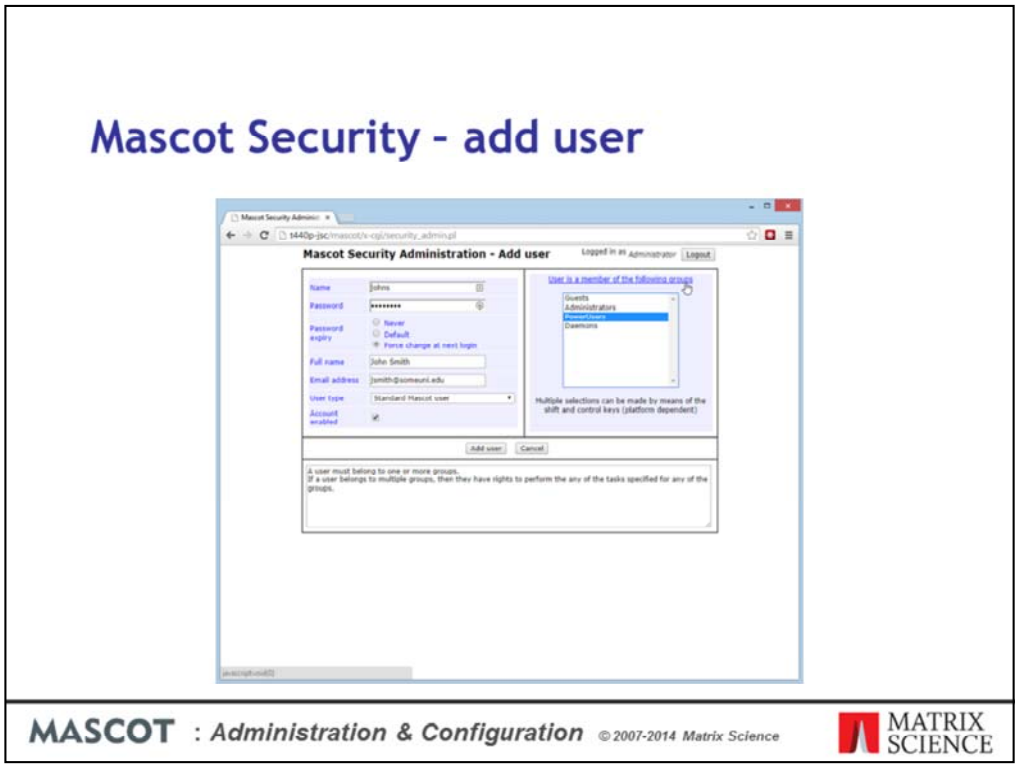

Adding a new user is very simple.

The Administrator must enter a username and password and it is usually a good idea to force the user to enter a new password when they first login

Enter their full name and email address. The user will be able to change this

I'll return to the choice of user types later - most users should just be standard Mascot users.

Make sure that the account is enabled, and then select one or more groups for the user to belong to.

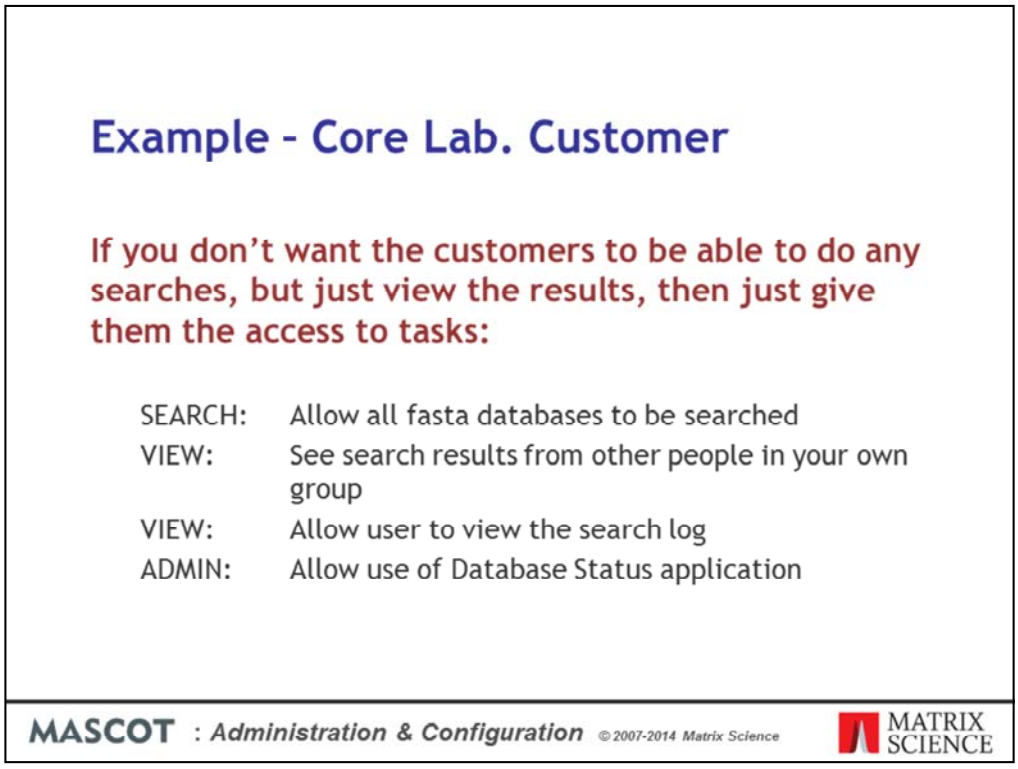

In a core lab., you may want a group that enables customers to view their results. In this case, only give them rights to perform these tasks:

You must allow all fasta databases to be searched, otherwise they won't be able to view reports. However, they can't perform PMF or MS/MS searches, because this task is missing.

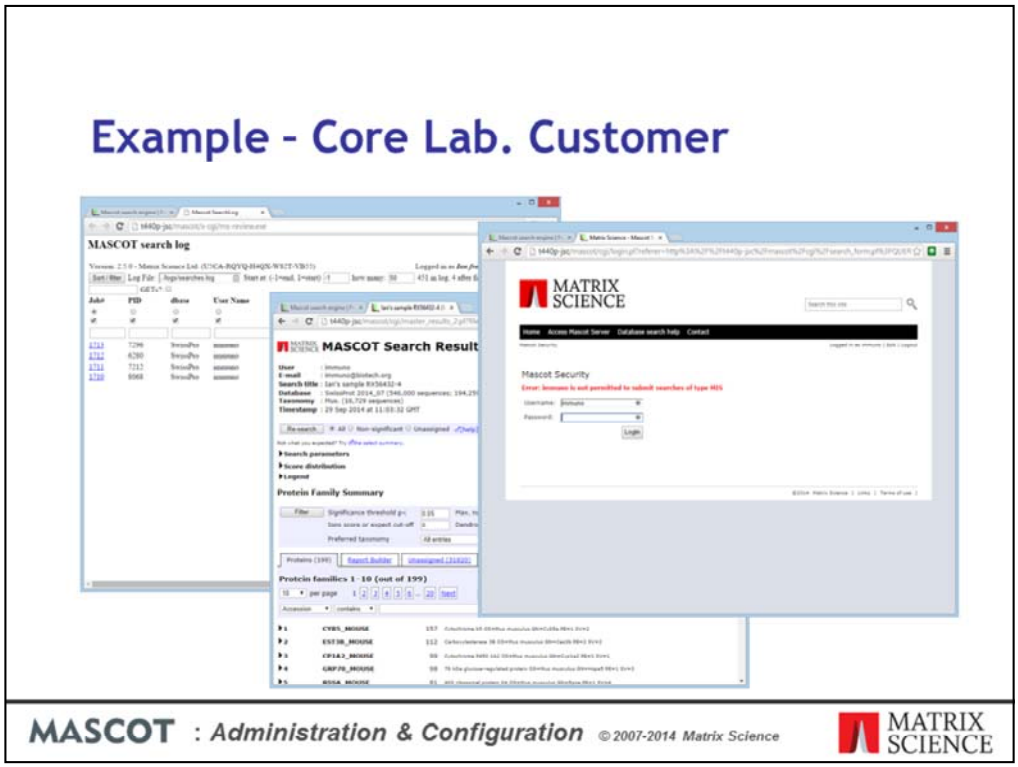

A group member will be able to view the search log, see their results, but when they try to do a repeat search, access is denied

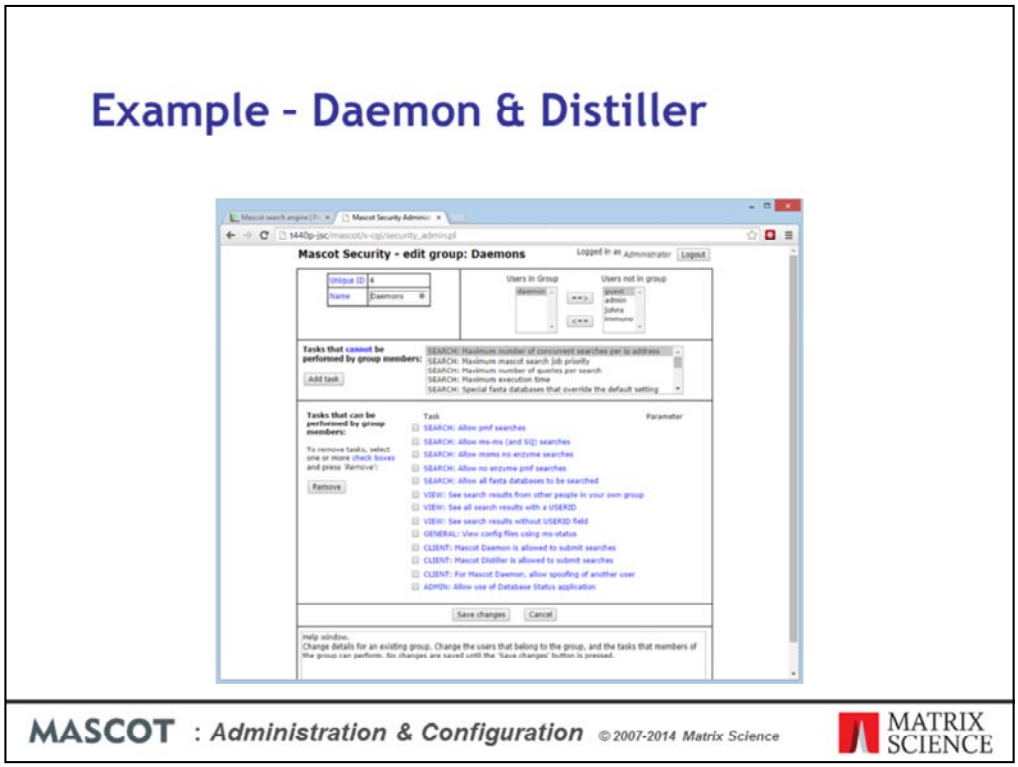

Client software, like Mascot Distiller and Mascot Daemon, requires special privileges. The Mascot Daemon user must have the security token 'Mascot Daemon is allowed to submit searches'. The default daemons group settings shown here are appropriate for either a Daemon or Distiller client.

View config files is required because both clients need to retrieve configuration information from the server, like a list of the databases that are available.

Mascot users can be given the privilege to submit searches under other user names. This is particularly useful in a core lab, when customers only have privileges to see their own search results, so the instrument operator needs to submit their searches under individual customer log-in names. The security task is 'For Mascot Daemon, allow spoofing of another user'.

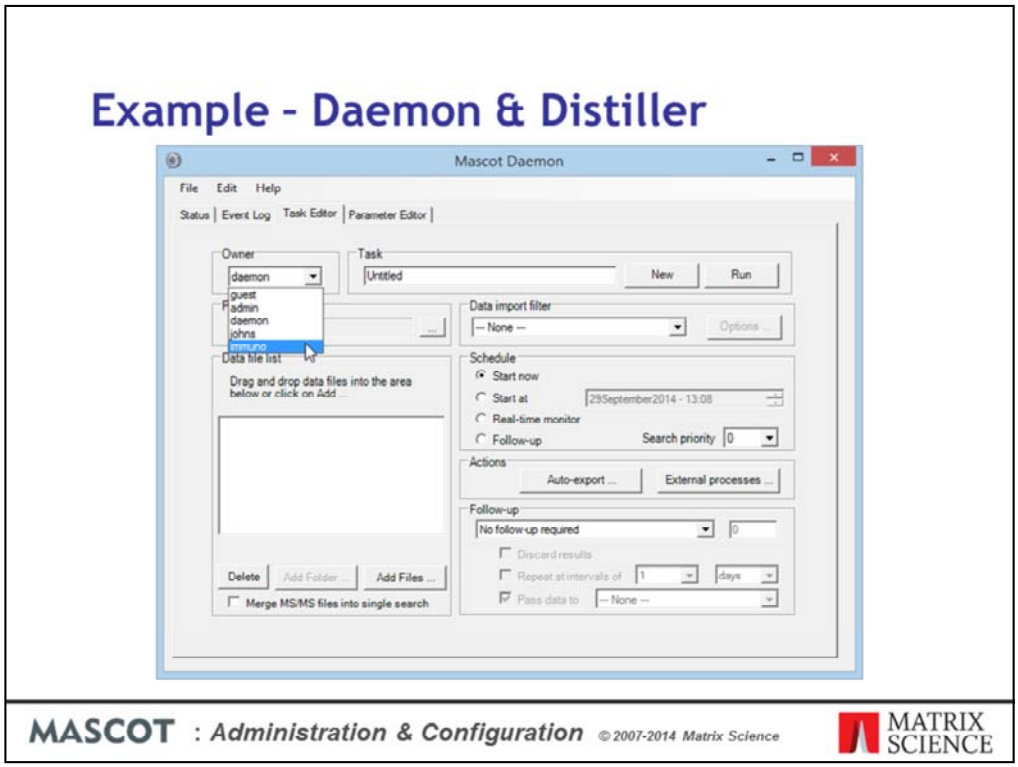

If this has been enabled for the user who is running Mascot Daemon, then a drop down list of other Mascot user names will be enabled on the Task Editor tab.

Example - 3rd party applications

Older applications that submit searches to Mascot won't have their own login

Try logging in using Internet Explorer before running the application - cookies

Use one of the 'special' user types:

There are potential issues with older, legacy applications that interface with Mascot but don't have code to support the security system.

Since session ids are saved as cookies, and since most Windows applications that access web sites use Microsoft Internet Explorer libraries, it may be sufficient just to login from an Internet Explorer browser window before starting the application.

Alternatively, you can use one of the special user types.

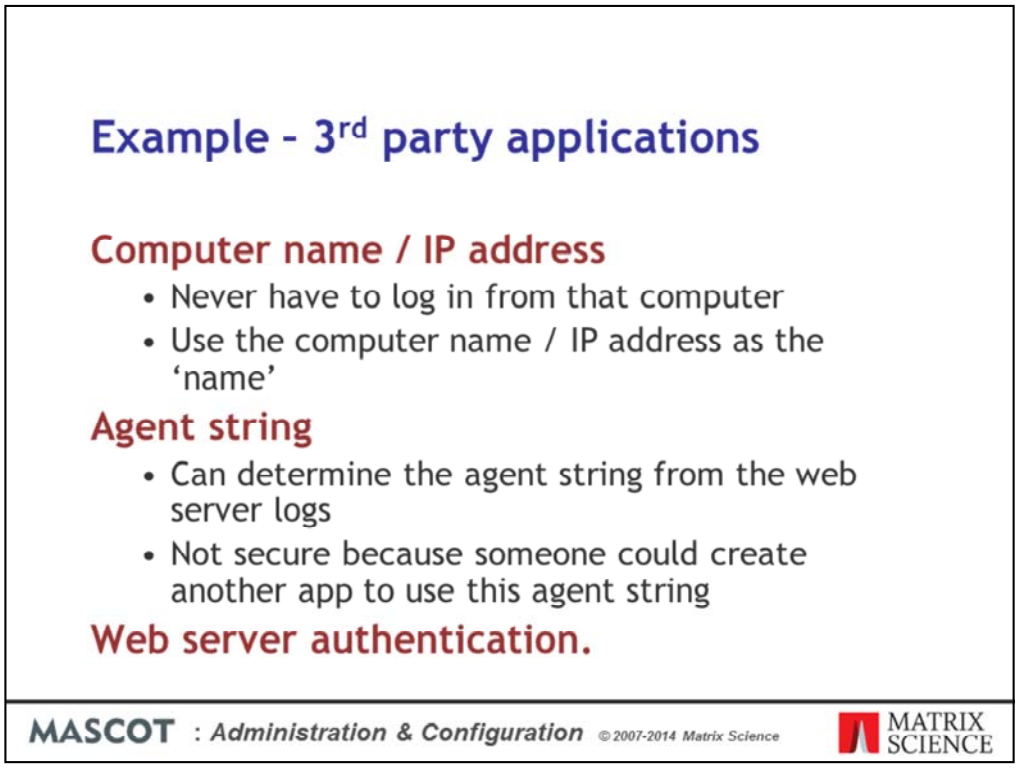

These methods are less secure than a password protected login, but ensure that all applications are able to connect somehow.

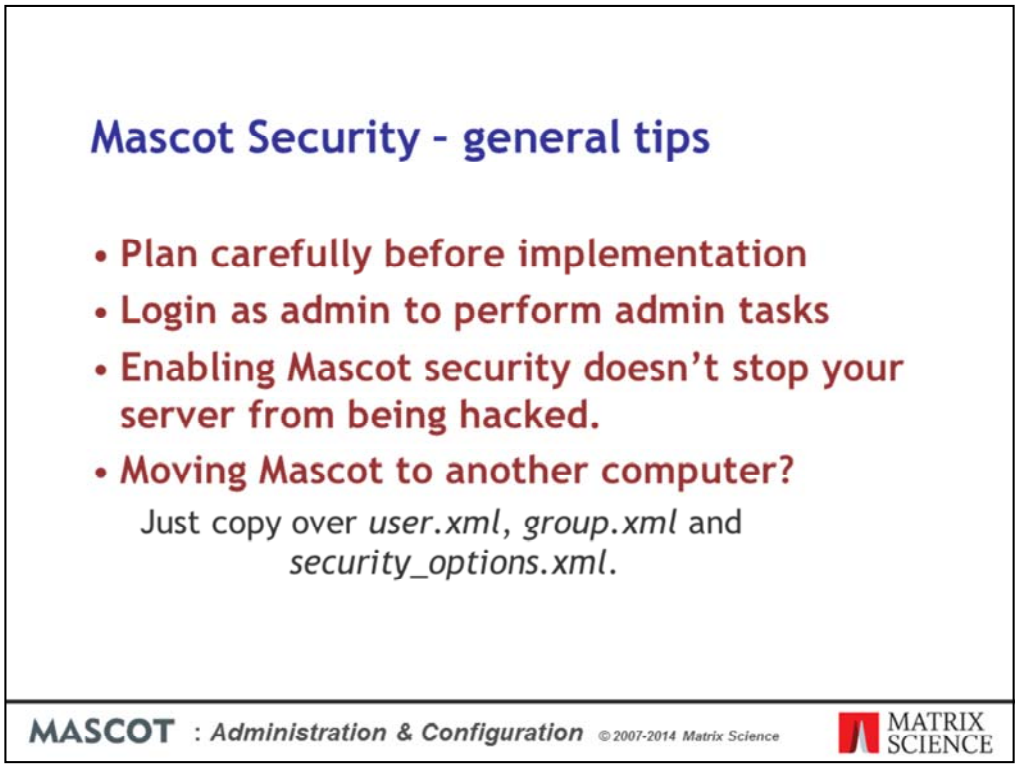

I can't stress enough that you should plan what you intend to do before you start. Think carefully about what groups you want to create

As any Unix administrator will tell you, it's always best to separate your administration and user tasks. However, we can't force you, but that is why the default admin user cannot submit searches.

It is very important to understand that Mascot security does not provide protection against a malicious hacker. Hackers don't attack a server through Mascot … they've probably never heard of Mascot. They attack through weaknesses in the operating system and through flaws in well known applications, like the web server. It is still essential to have a dedicated firewall between the Mascot server and the Internet

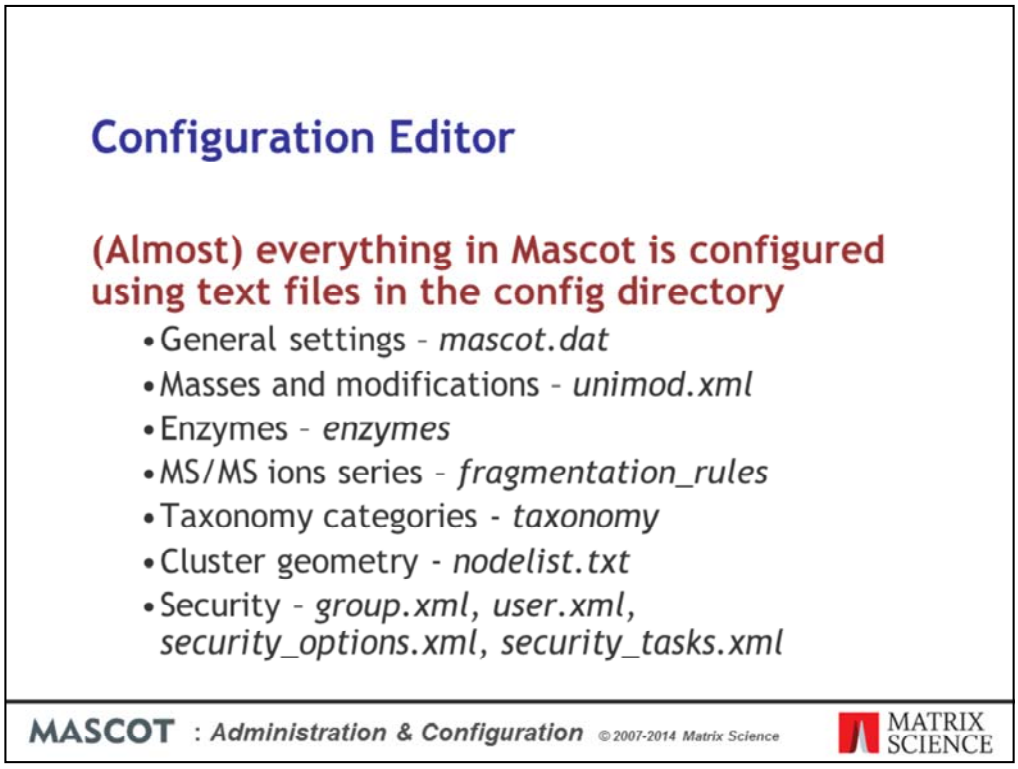

(Almost) everything in Mascot is configured using these text files in the config directory. The syntax for each of these files is described in the manual.

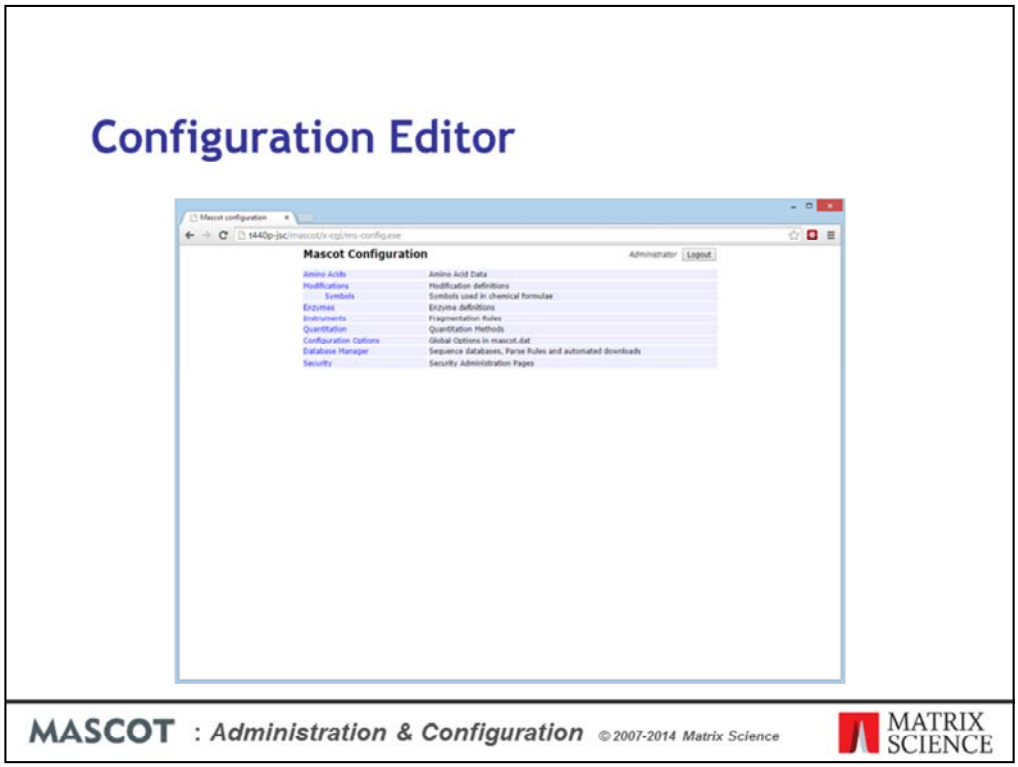

You can edit these files in a text editor, but it is easier and safer to use the browser-based Configuration Editor. The first three entries are interfaces to different sections of unimod.xml.

Configuration Editor

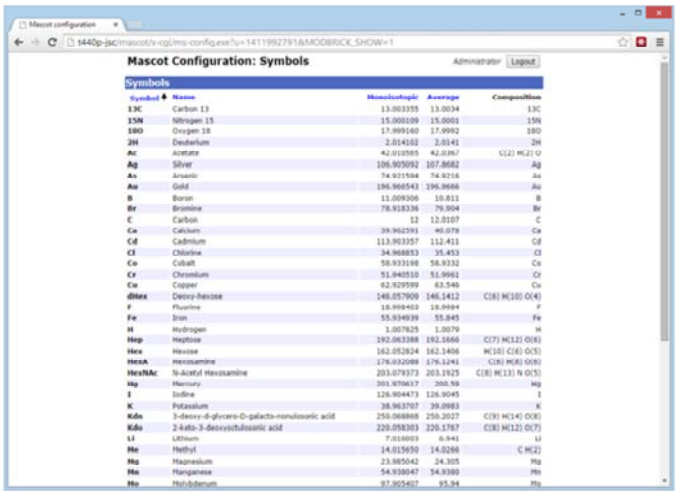

MASCOT : Administration & Configuration MATRIX SCIENCE

You may need to add or edit an element or an amino acid, but its fairly unlikely.

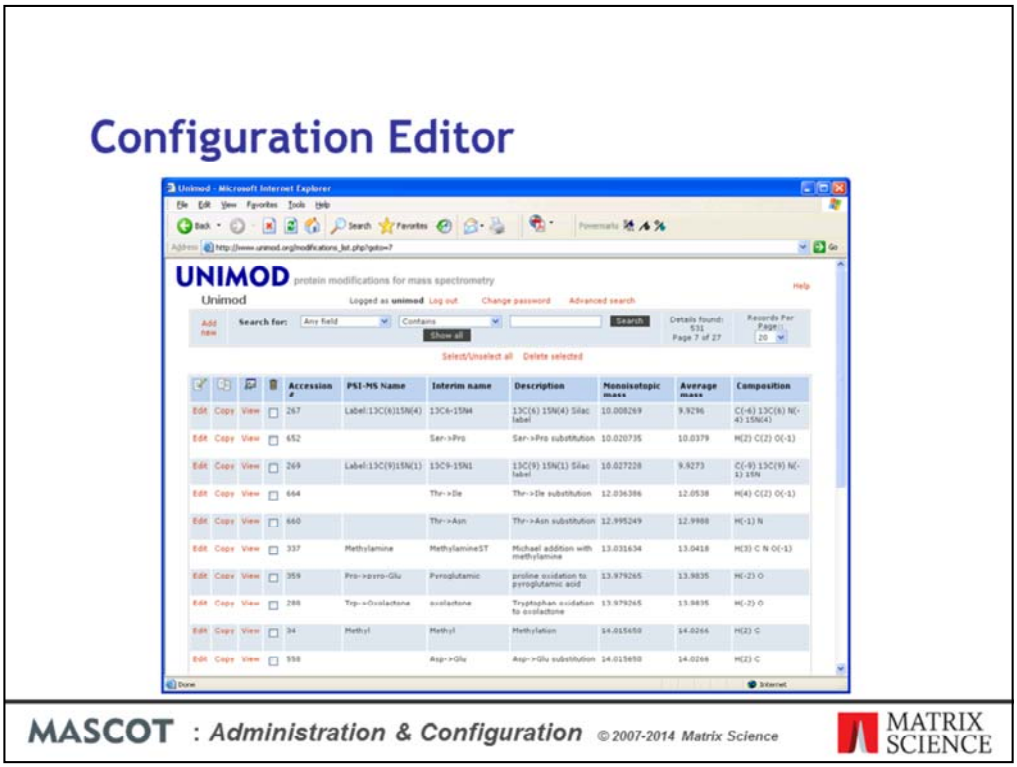

The modifications file, unimod.xml, is an XML representation of a public database called Unimod. This is the interface to the public database. From time to time, you should update your local file by downloading the latest file using the links in the Unimod help

Configuration Editor

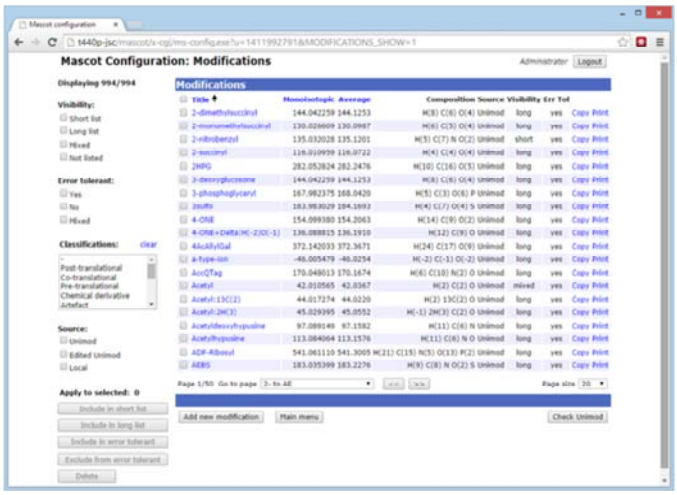

MASCOT : Administration & Configuration © 2007-2014 Matrix Science MATRIX SCIENCE

If you need to add a new modification that will be of interest to others, it's best to add it to the public Unimod database, so that the information is available to all Mascot users. If you want to create a local definition, because the modification is confidential or because you are just experimenting. If so, use the Modifications section of the configuration editor

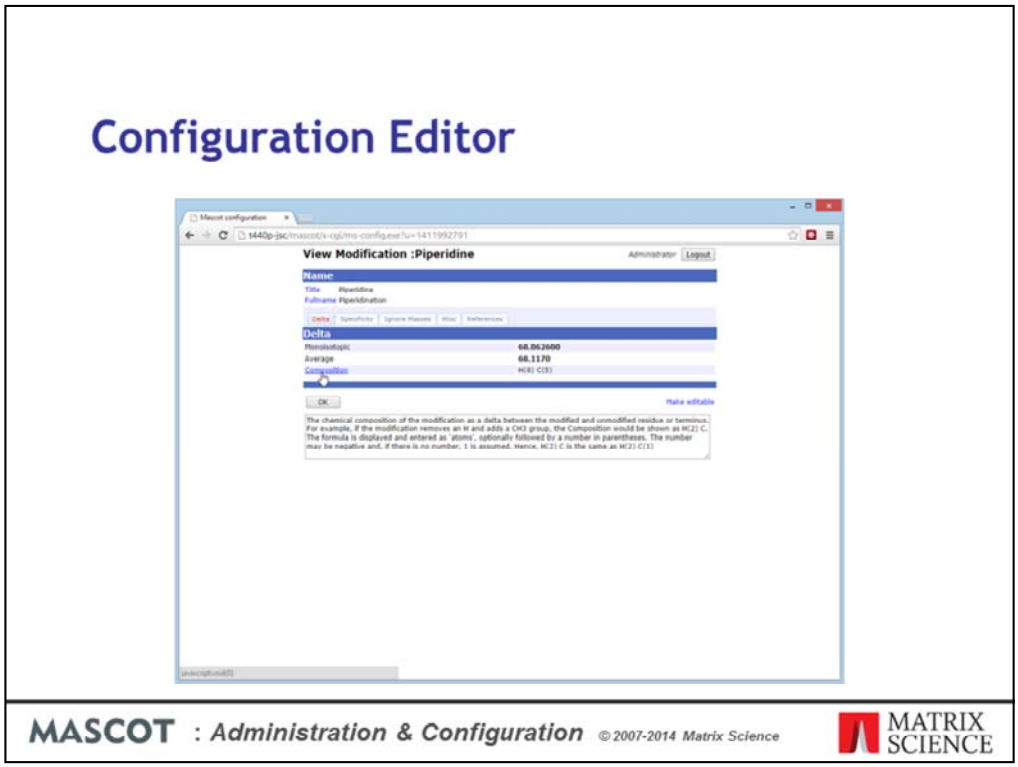

The user interface is fairly self explanatory, and help is displayed for each field when you mouse over the label. If the modification is from the downloaded unimod.xml file, and you want to make changes, click on the ‘Make editable’ link

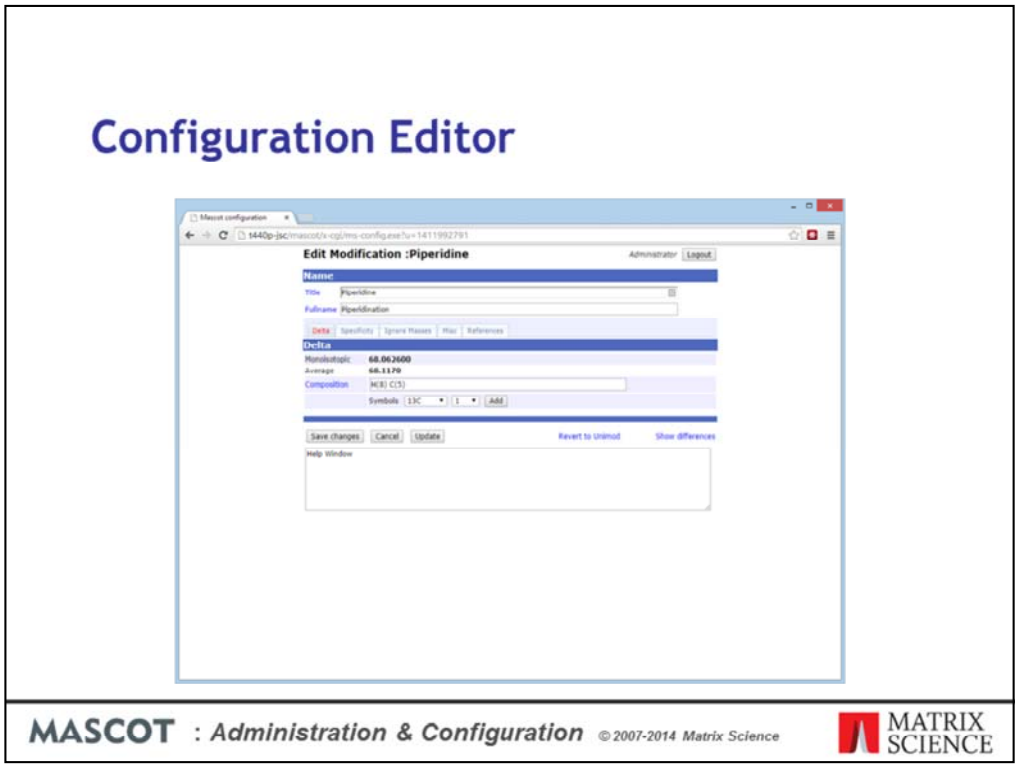

All aspects of the modification become editable. If you want to see the differences between the original Unimod entry and the local entry, choose ‘Show differences’. If you want to discard your changes, choose ‘Revert to Unimod’

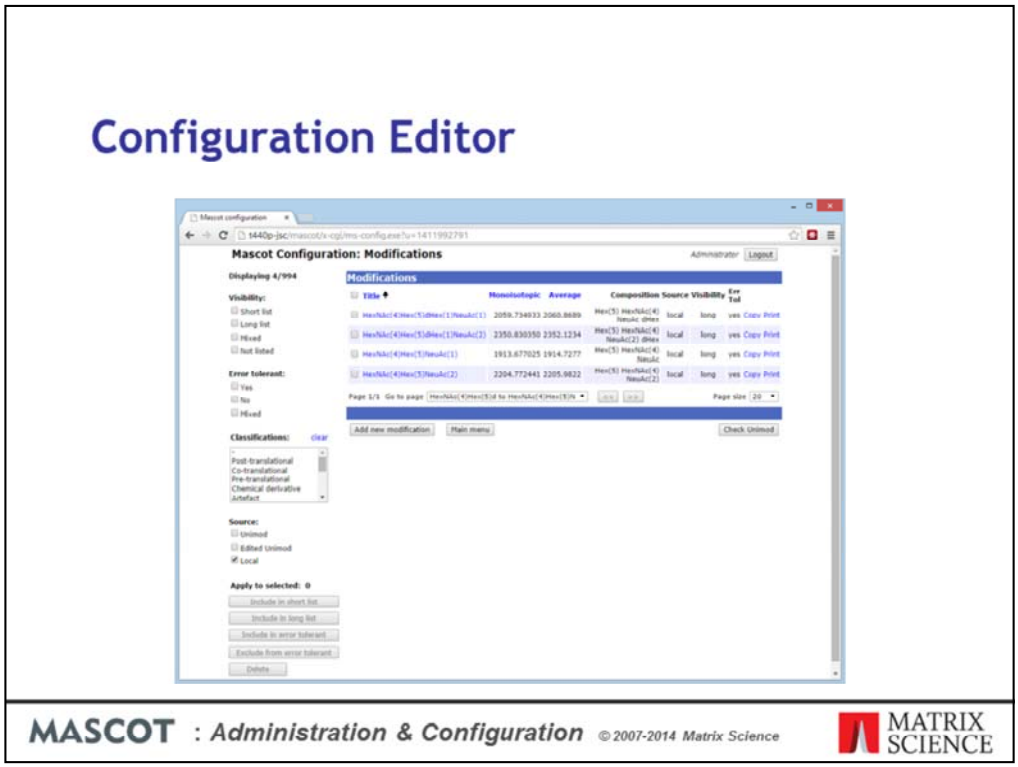

The top level page allows you to filter the listing. Here, we have checked the box to show only local definitions

Configuration Editor

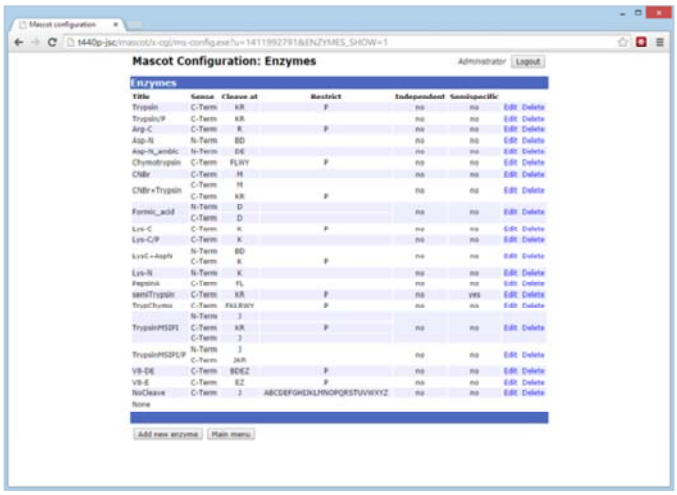

MASCOT : *Administration & Configuration* MATRIX SCIENCE

Similarly for Enzymes and Instruments. The other sections: Quantitation, Database Manager, and Security, have been touched on in earlier presentations.

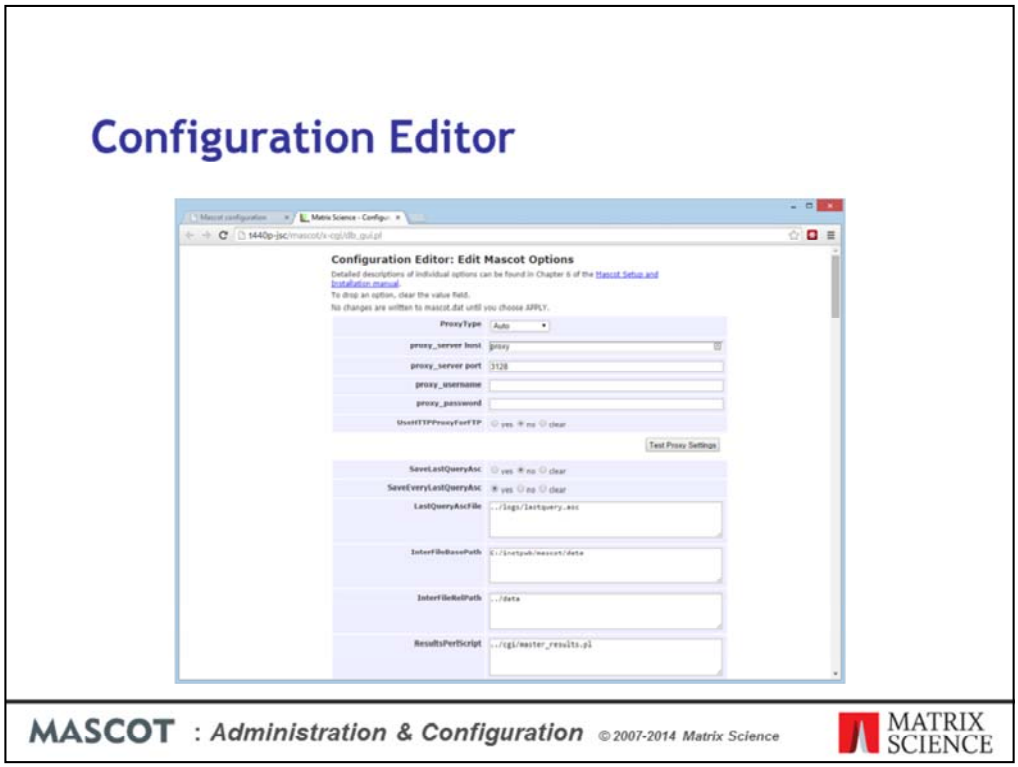

The Configuration Options section deals with various global settings stored in mascot.dat. Remember that all the sections of mascot.dat that deal with sequence database configuration may be re-written at any time by Database Manager. If you've been in the habit of editing mascot.dat in a text editor, only the Options and Cluster sections can be modified safely unless you decide never to use Database Manager.

There are two configuration files that are missing from the configuration editor: Taxonomy categories (taxonomy) and Cluster geometry (nodelist.txt). If you need to make changes to these, you still have to use a text editor. You can find full details of the file formats in the manual.

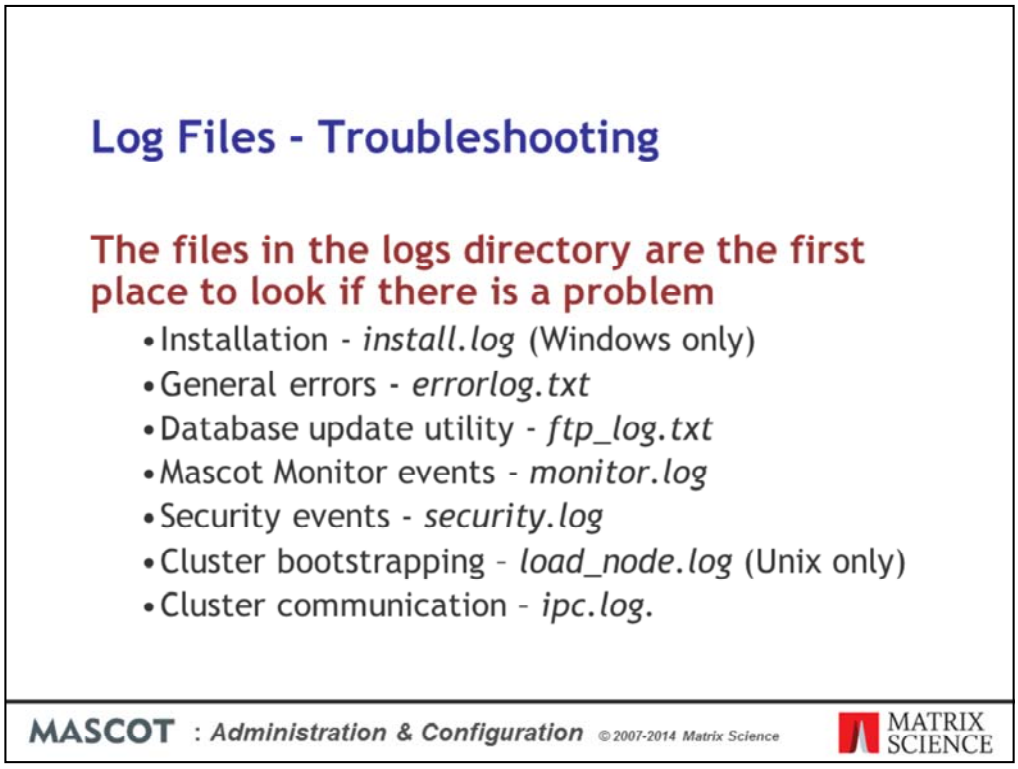

The log files are the first place to look if there is a problem. Most of these files can be accessed from a browser via links in the database status utility, which we'll come to in a minute

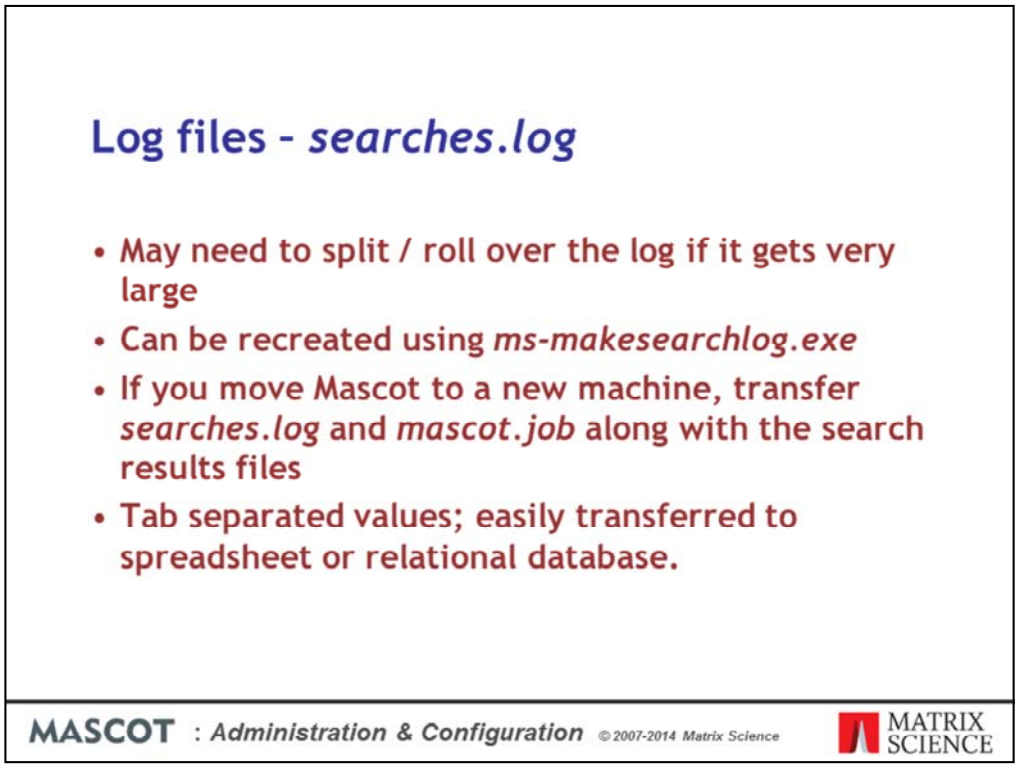

Every Mascot search adds a new line to the searches log. If the server is busy, the log files can get very large. For most of the logs, you'll probably want to delete the file from time to time. The searches log is different, and you'll probably want to keep all entries indefinitely. Best idea is to rename the file periodically. For example, searches.log.20060301. The system will then create a new one automatically.

If the search log is accidentally damaged, a new one can be created by scanning all the result files on the disk

If you move Mascot to a new machine, you'll probably want to transfer or rebuild *searches.log.* The other important file is *mascot.job*, which contains the "next" job number. If you don't copy this across, your job numbers will start all over and you'll have to sort the search log by descending start time to locate the most recent searches.

Some people import the search log into a database application so as to make it easier to search or filter, especially if it gets very large

Log files - database status

The database status page has links to many of the log files mentioned earlier.

This is the starting point for troubleshooting problems, especially those connected with database updates

When Mascot Monitor first starts, there is a single entry per database. Once a database has been updated, there are two entries, one for the “old” file and one for the ‘new’. We need this double view, because of the way databases can be updated in the background, without interrupting searches. During database exchange, while the new file is compressed and tested, searches can still be submitted to run against the old file. Once the new file is ready, all new searches are run against the new file. So, on a busy server, there may be a period when you have searches running against both the old and new files. As soon as the last search on the old file has completed, the old database can be taken out of use. Mascot monitor takes care of all of this automatically.

Log files - database status

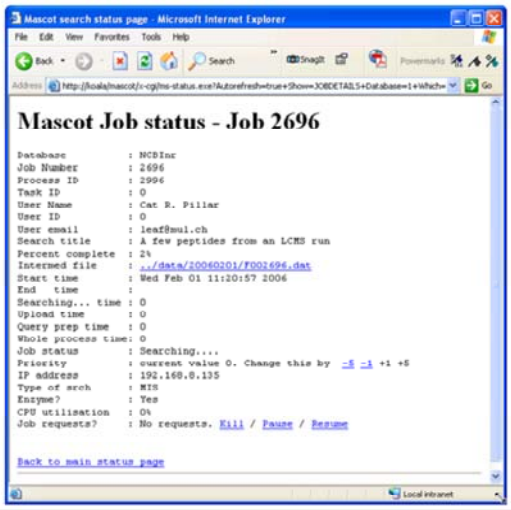

You can drill down from the database status page to a list of searches for each database, and then down to the details of an individual search, as shown here.

If necessary, you can kill or pause a search from this page, or change its priority

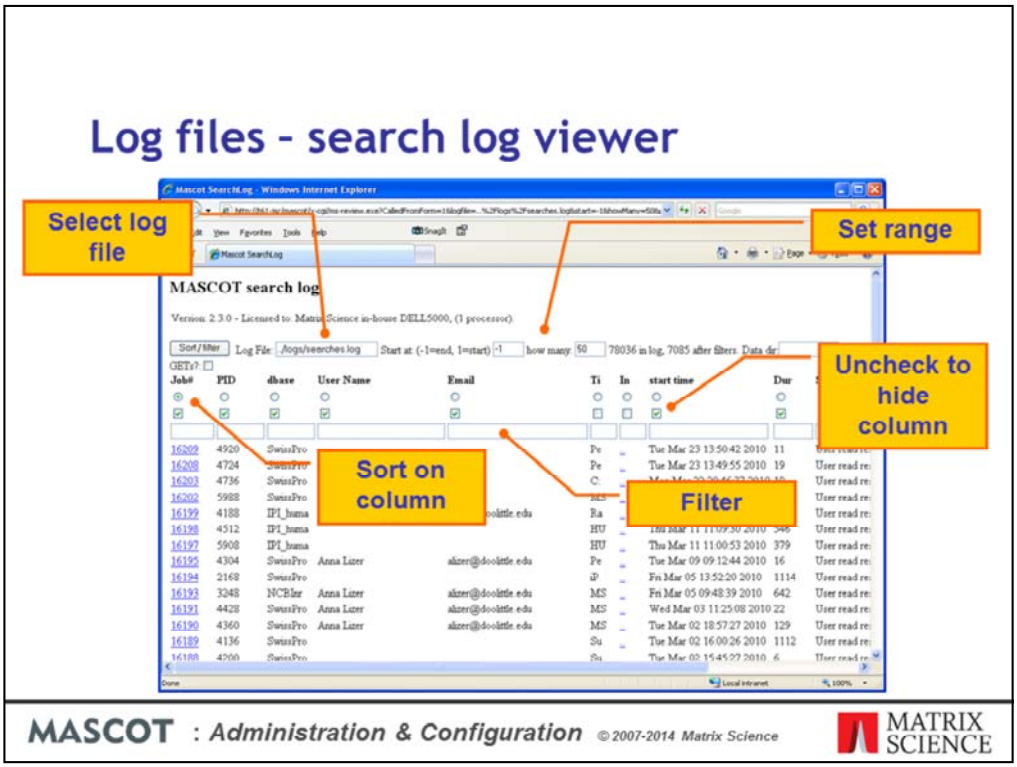

This is the search log viewer. You can sort on any of the columns using the radio buttons and then pressing Sort / filter.

The checkboxes determine whether a column is shown or hidden. The two hidden columns are Ti for title and In for intermediate file. This is a hyperlink to display the raw mascot result file, but you can only see the first two characters of the file name, which are "..". The hyperlink in the first column displays the formatted result report.

You can set the number of lines to display, and whether to start at the beginning or end of the file.

You can also filter the display by entering text into the edit boxes. For example, your user name or a word from the search title

If you have split the search log into sections, then you can choose which file to display by entering the path

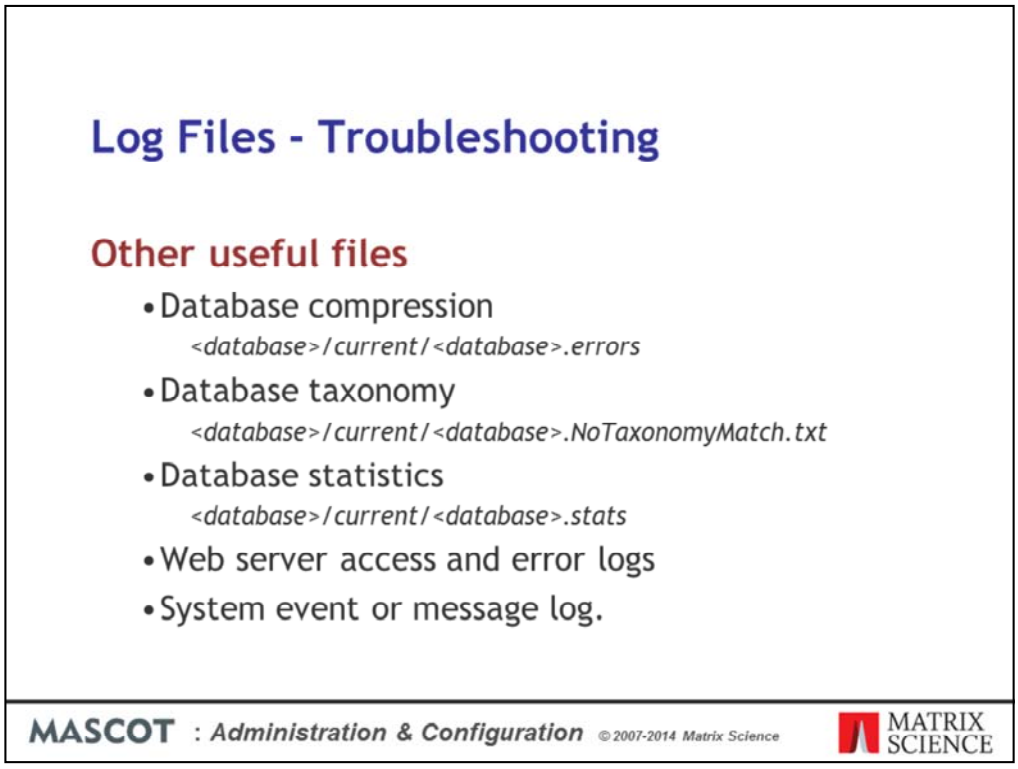

The first three files are most easily viewed using the links in database status.

Sometimes, there will be nothing in the Mascot logs because the problem is external to Mascot. Sometimes, it can be useful to look at the web server logs or even the operating system logs.

The Windows Event Viewer (control panel; administrative tools) allows you to browse Windows system messages

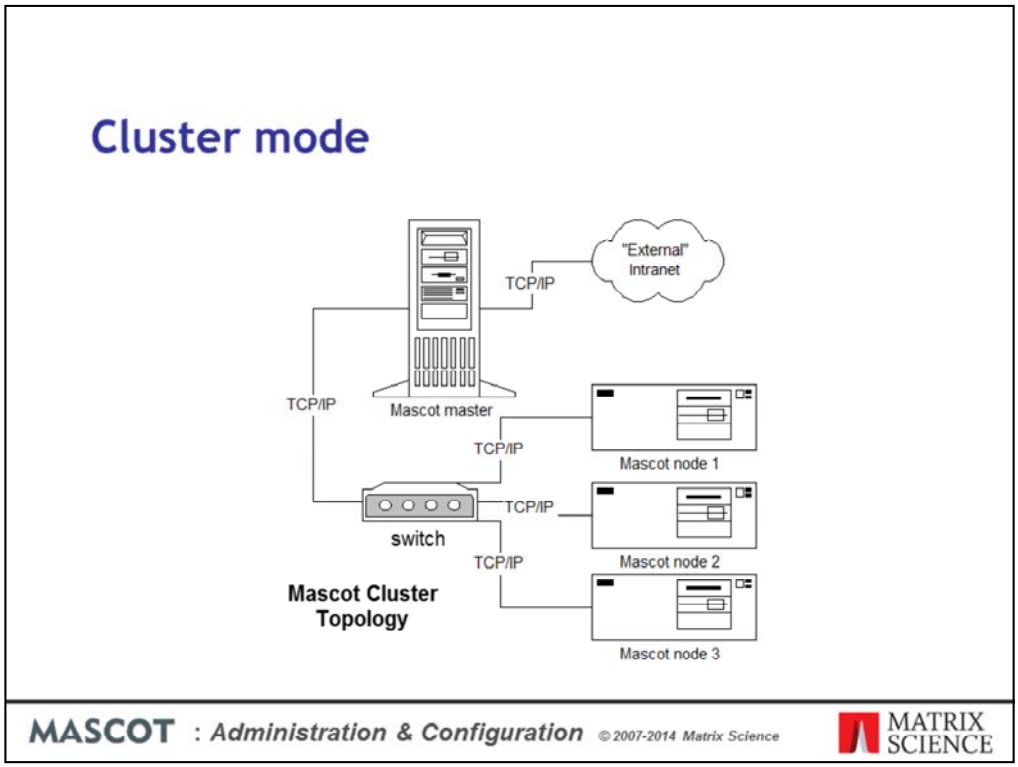

Just a few words about Mascot cluster mode

Mascot supports cluster operation using a Beowulf-like topology. Mascot supports cluster mode as standard, whenever the licence is for 4 CPU's or more. You just have to hook up an appropriate number of PC's on a local LAN

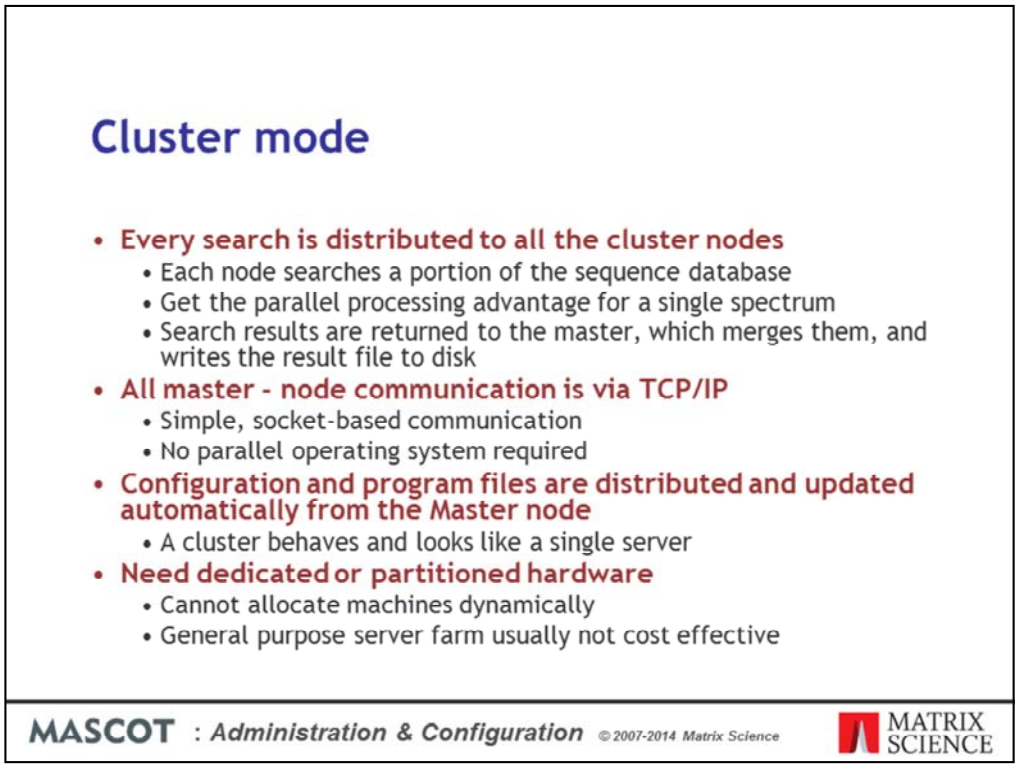

In cluster mode, every search is distributed to all the cluster nodes, and each node searches a portion of the sequence database. This means that you get the parallel processing advantage for a single spectrum. Search results are returned to the master, which merges them, and writes the result file to disk

All master - node communication is via TCP/IP. This uses simple, socket-based communication. A parallel operating system is not required

Configuration and program files are distributed and updated automatically from the Master node, so the cluster behaves and looks like a single server

Sometimes, we are asked about running Mascot on a general purpose server farm, that is being used for other applications.

For standard PC hardware, the cost of the Mascot licence is greater than that of the hardware. For a given capacity, the lowest cost route will always be a minimum number of licences on dedicated hardware. If you want to run Mascot on machines that are part of a larger cluster, best to partition off a number of nodes, either for exclusive Mascot use, or where Mascot has absolute priority over other processes.

Mascot cluster nodes cannot be dynamic because of the size of the database files. The time taken to move these files between machines, or even just in and out of memory, would always be unacceptable

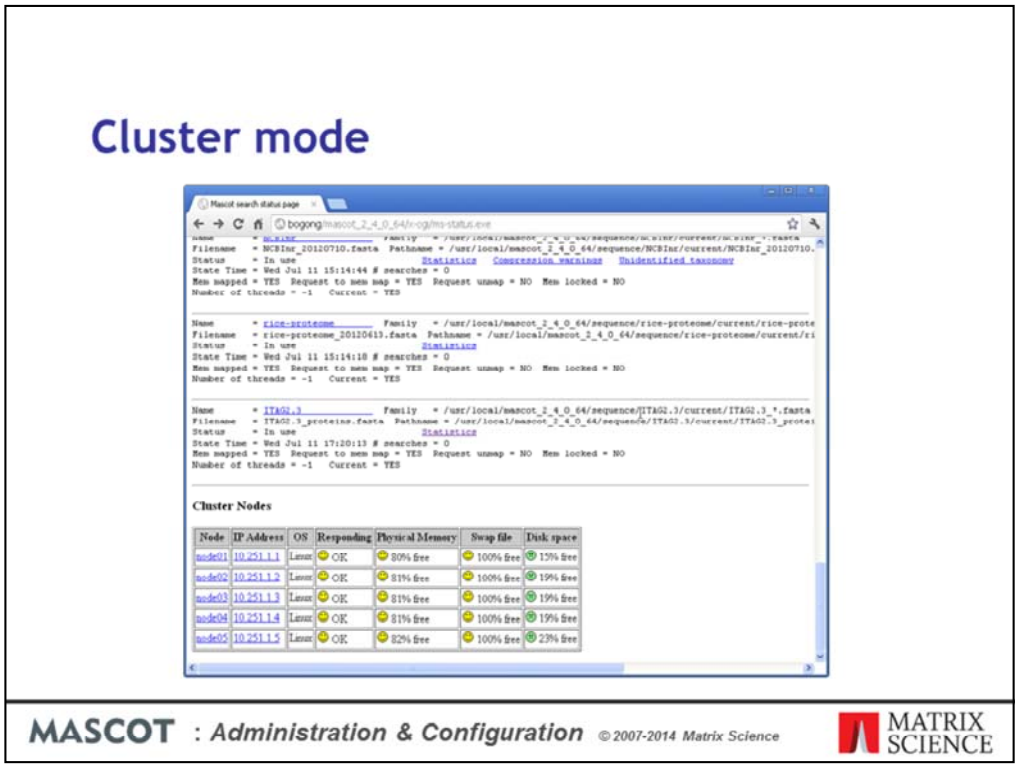

If you have a Mascot cluster, it is almost as easy to administer as a single machine. Everything is reported and controlled through the master node. The Cluster Nodes table provides a an overview of the search nodes; smiley faces or green unhappy faces. If you need a closer look at one of the search nodes, follow the links to view the search node log files.

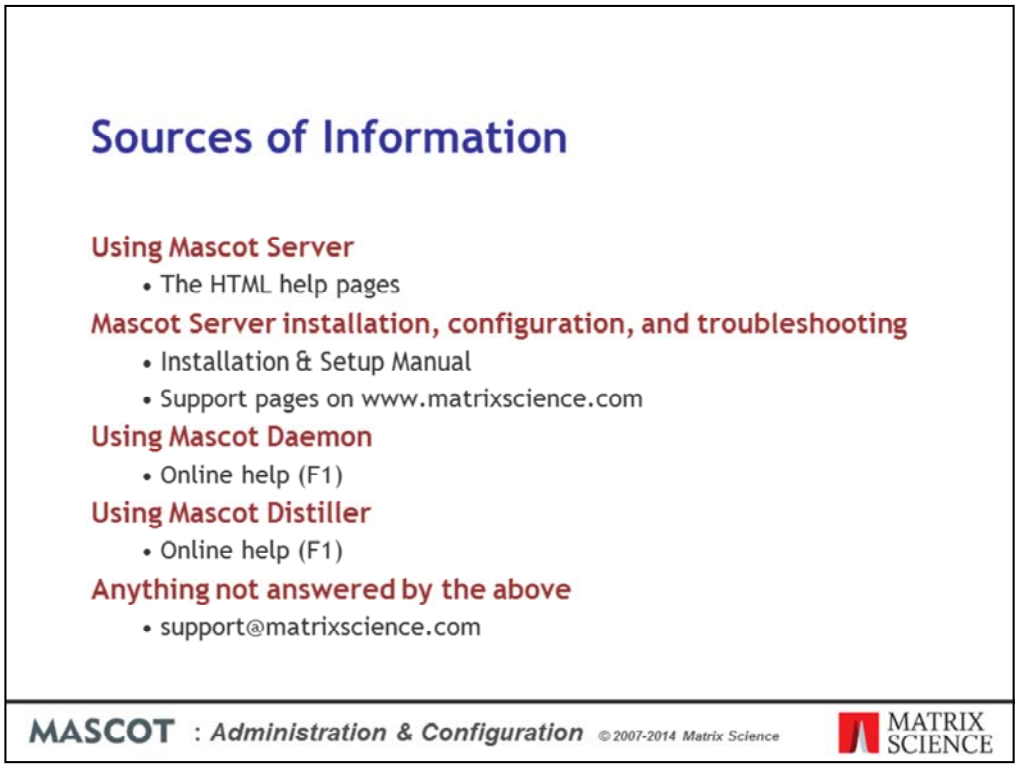

Finally, a reminder of where you can find technical information about Mascot