

**NAME**

search – SWISH++ searcher

**SYNOPSIS**

**search** [ *options* ] *query*

**DESCRIPTION**

**search** is the SWISH++ searcher. It searches a previously generated index for the words specified in a query. In addition to running from the command-line, it can run as a daemon process functioning as a “search server.”

**QUERY INPUT****Query Syntax**

The formal grammar of a query is:

```

query:          query optional_relop meta
                  meta

meta:           meta_name = primary
                  primary

meta_name:     word

primary:       ( query )
                  not meta
                  word
                  word*

optional_relop: and
                  or
                  (empty)

```

In practice, however, the query is the set of words sought after, possibly restricted to meta data, and possibly combined with the Boolean operators “and,” “or,” and “not.” The asterisk (\*) can be used as a wildcard character at the end of words. Queries are evaluated in left-to-right order, i.e., “and” has the same precedence as “or.” See the EXAMPLES.

**Character Mapping and Word Determination**

The same character mapping and word determination heuristics used by **index**(1) are used on queries prior to searching.

**RESULTS OUTPUT****Result Components**

The results are output either in “classic” or XML format. In either case, the components of the results are:

<i>rank</i>	An integer from from 1 to 100.
<i>path-name</i>	The relative path to where the file was originally indexed.
<i>file-size</i>	The file’s size in bytes.
<i>file-title</i>	If the file is of a format that can have titles (HTML, XHTML, LaTeX, mail, or Unix manual pages) and the title was extracted, then <i>file-title</i> is its title; otherwise, it is its filename.

**Classic Results Format**

The “classic” results format is plain text as:

```
rank path-name file-size file-title
```

It can be parsed easily in Perl with:

```
($rank,$path,$size,$title) = split( / /, $_, 4 );
```

(The separator can be changed via the **-R** or **--separator** options or the **ResultSeparator** variable.)

Prior to results lines, comment lines may also appear containing additional information about the query results. Comment lines are in the format of:

```
# comment-key: comment-value
```

The keys and values are:

<b>ignored:</b> <i>stop-words</i>	The list of stop-words (separated by spaces) ignored in the query.
<b>not found:</b> <i>word</i>	The word was not found in the index.
<b>results:</b> <i>result-count</i>	The total number of results.

### XML Results Format

The XML results format is given by the DTD:

```
<!ELEMENT SearchResults (IgnoredList?, ResultCount, ResultList?)>
<!ELEMENT IgnoredList (Ignored+)>
<!ELEMENT Ignored (#PCDATA)>
<!ELEMENT ResultCount (#PCDATA)>
<!ELEMENT ResultList (File+)>
<!ELEMENT File (Rank, Path, Size, Title)>
<!ELEMENT Rank (#PCDATA)>
<!ELEMENT Path (#PCDATA)>
<!ELEMENT Size (#PCDATA)>
```

and by the XML schema located at:

```
http://homepage.mac.com/pauljlucas/software/swish/SearchResults/SearchResults.xsd
```

For example:

```
<?xml version="1.0" encoding="us-ascii"?>
<!DOCTYPE SearchResults SYSTEM
  "http://homepage.mac.com/pauljlucas/software/swish/SearchResults.dtd">
<SearchResults
  xmlns="http://homepage.mac.com/pauljlucas/software/swish/SearchResults"
  xmlns:xsi="http://www.w3.org/2001/XMLSchema-instance"
  xsi:schemaLocation="http://homepage.mac.com/pauljlucas/software/swish/SearchRes
    SearchResults.xsd">
  <IgnoredList>
    <Ignored>stop-word</Ignored>
  ...
  </IgnoredList>
  <ResultCount>42</ResultCount>
  <ResultList>
    <File>
      <Rank>rank</Rank>
      <Path>path-name</Path>
      <Size>file-size</Size>
      <Title>file-title</Title>
    </File>
  ...
  </ResultList>
</SearchResults>
```

## RUNNING AS A DAEMON PROCESS

### Description

**search** can alternatively run as a daemon process (via either the **-b** or **--daemon-type** options or the **SearchDaemon** variable) functioning as a “search server” by listening to a Unix domain socket (specified by either the **-u** or **--socket-file** options or the **SocketFile** variable), a TCP socket (specified by either the **-a** or **--socket-address** options or the **SocketAddress** variable), or both. Unix domain sockets are preferred for both performance and security. For search-intensive applications, such as a search engine on a heavily used web site, this can yield a large performance improvement since the start-up cost (**fork(2)**, **exec(2)**, and initialization) is paid only once.

If the process was started with root privileges, it will give them away immediately after initialization and before servicing any requests.

### Clients and Requests

Search clients connect to a daemon via a socket and send a query in the same manner as on the command line (including the first word being “search”). The only exception is that shell meta-characters *must not* be escaped (backslashed) since no shell is involved. Search results are returned via the same socket. See the EXAMPLES.

### Multithreading

A daemon can serve multiple query requests simultaneously since it is multi-threaded. When started, it “pre-threads” meaning that it creates a pool of threads in advance that service an indefinite number of requests as a further performance improvement since a thread is not created and destroyed per request.

There is an initial, minimum number of threads in the thread pool. The number of threads grows dynamically when there are more requests than threads, but not more than a specified maximum to prevent the server from thrashing. (See the **-t**, **--min-threads**, **-T**, and **--max-threads** options or the **ThreadsMin** or **ThreadsMax** variables.) If the number of threads reaches the maximum, subsequent requests are queued until existing threads become available to service them after completing in-progress requests. (See either the **-q** or **--queue-size** options or the **SocketQueueSize** variable.)

If there are more than the minimum number of threads and some remain idle longer than a specified timeout period (because the number of requests per unit time has dropped), then threads will die off until the pool returns to its original minimum size. (See either the **-O** or **--thread-timeout** options or the **Thread-Timeout** variable.)

### Restrictions

A single daemon can search only a single index. To be able to search multiple indices concurrently, multiple daemons can be run, each searching its own index and using its own socket file. An index *must not* be modified or deleted while a daemon is using it.

## OPTIONS

Options begin with either a ‘-’ for short options or a “--” for long options. Either a ‘-’ or “--” by itself explicitly ends the options; however, the difference is that ‘-’ is returned as the first non-option whereas “--” is skipped entirely. Either short or long options may be used. Long option names may be abbreviated so long as the abbreviation is unambiguous.

For a short option that takes an argument, the argument is either taken to be the remaining characters of the same option, if any, or, if not, is taken from the next option unless said option begins with a ‘-’.

Short options that take no arguments can be grouped (but the last option in the group can take an argument), e.g., **-lrv4** is equivalent to **-l -r -v4**.

For a long option that takes an argument, the argument is either taken to be the characters after a ‘=’, if any, or, if not, is taken from the next option unless said option begins with a ‘-’.

**-?**

**--help** Print the usage (“help”) message and exit.

**-aa**

**--socket-address=a** When running as a daemon, the address, *a*, to listen to for TCP requests. (Default is all IP addresses and port 1967.) The address argument is of the form:

[ *host* : ] *port*

that is: an optional host and colon followed by a port number. The *host* may be one of a host name, an IP address, or the \* character meaning “any IP address.” Omitting the *host* and colon also means “any IP address.”

- bt**  
**--daemon-type=*t*** Run as a daemon process. (Default is not to.) The type, *t*, is one of:
- `none` Same as not specifying the option at all. (This does not purport to be useful, but rather consistent with the types that can be specified to the **SearchDaemon** variable.)
  - `tcp` Listen on a TCP socket (see the **-a** option).
  - `unix` Listen on a Unix domain socket (see the **-u** option).
  - `both` Listen on both.
- By default, if executed from the command-line, **search** appears to return immediately; however, it has merely detached from the terminal and put itself into the background. There is no need to follow the command with an ‘&’.
- B**  
**--no-background** When running as a daemon process, do not detach from the terminal and run in the background. (Default does.)
- The reason not to run in the background is so a wrapper script can see if the process dies for any reason and automatically restart it.
- cf**  
**--config-file=*f*** The name of the configuration file, *f*, to use. (Default is `swish++.conf` in the current directory.) A configuration file is not required: if none is specified and the default does not exist, none is used; however, if one is specified and it does not exist, then this is an error.
- d**  
**--dump-words** Dump the query word indices to standard output and exit. Wildcards are not permitted.
- D**  
**--dump-index** Dump the entire word index to standard output and exit.
- fn**  
**--word-files=*n*** The maximum number of files, *n*, a word may occur in before it is discarded as being too frequent. (Default is infinity.)
- Ff**  
**--format=*f*** The format, *f*, search results are output in. The format is either `classic` or `XML`. (Default is `classic`.)
- Gs**  
**--group=*s*** The group, *s*, to switch the process to after starting and only if started as root. (Default is `nobody`.)
- if**  
**--index-file=*f*** The name of the index file, *f*, to use. (Default is `swish++.index` in the current directory.)
- mn**  
**--max-results=*n*** The maximum number of results, *n*, to return. (Default is 100.)
- M**  
**--dump-meta** Dump the meta-name index to standard output and exit.

<b>-os</b>	
<b>--socket-timeout=<i>s</i></b>	The number of seconds, <i>s</i> , a search client has to complete a query request before the socket connection is closed. (Default is 10.) This is to prevent a client from connecting, not completing a request, and causing the thread servicing the request to wait forever.
<b>-Os</b>	
<b>--thread-timeout=<i>s</i></b>	The number of seconds, <i>s</i> , until an idle spare thread dies while running as a daemon. (Default is 30.)
<b>-pn</b>	
<b>--word-percent=<i>n</i></b>	The maximum percentage, <i>n</i> , of files a word may occur in before it is discarded as being too frequent. (Default is 100.) If you want to keep all words regardless, specify 101.
<b>-Pf</b>	
<b>--pid-file=<i>f</i></b>	The name of the file to record the process ID of <b>search</b> if running as a daemon. (Default is none.)
<b>-qn</b>	
<b>--queue-size=<i>n</i></b>	The maximum number of socket connections to queue. (Default is 511.)
<b>-rn</b>	
<b>--skip-results=<i>n</i></b>	The initial number of results, <i>n</i> , to skip. (Default is 0.) Used in conjunction with <b>-m</b> or <b>--max-results</b> , results can be returned in “pages.”
<b>-Rs</b>	
<b>--separator=<i>s</i></b>	The classic result separator string. (Default is " ".)
<b>-s</b>	
<b>--stem-words</b>	Perform stemming (suffix stripping) on words during the search. Words that end in the wildcard character are not stemmed. (Default is no.)
<b>-S</b>	
<b>--dump-stop</b>	Dump the stop-word index to standard output and exit.
<b>-tn</b>	
<b>--min-threads=<i>n</i></b>	Minimum number of threads to maintain while running as a daemon.
<b>-Tn</b>	
<b>--max-threads=<i>n</i></b>	Maximum number of threads to allow while running as a daemon.
<b>-uf</b>	
<b>--socket-file=<i>f</i></b>	The name of the Unix domain socket file to use while running as a daemon. (Default is /tmp/search.socket.)
<b>-Us</b>	
<b>--user=<i>s</i></b>	The user, <i>s</i> , to switch the process to after starting and only if started as root. (Default is nobody.)
<b>-V</b>	
<b>--version</b>	Print the version number of <b>SWISH++</b> to standard output and exit.
<b>-wn[<i>c</i>]</b>	
<b>--window=<i>n</i>[,<i>c</i>]</b>	Dump a “window” of at most <i>n</i> lines around each query word matching <i>c</i> characters. Wildcards are not permitted. (Default for <i>c</i> is 0.) Every window ends with a blank line.

## CONFIGURATION FILE

The following variables can be set in a configuration file. Variables and command-line options can be mixed, the latter taking priority.

<b>Group</b>	Same as <b>-G</b> or <b>--group</b>
--------------	-------------------------------------

<b>IndexFile</b>	Same as <b>-i</b> or <b>--index-file</b>
<b>PidFile</b>	Same as <b>-P</b> or <b>--pid-file</b>
<b>ResultSeparator</b>	Same as <b>-R</b> or <b>--separator</b>
<b>ResultsFormat</b>	Same as <b>-F</b> or <b>--format</b>
<b>ResultsMax</b>	Same as <b>-m</b> or <b>--max-results</b>
<b>SearchBackground</b>	Same as <b>-B</b> or <b>--no-background</b>
<b>SearchDaemon</b>	Same as <b>-b</b> or <b>--daemon-type</b>
<b>SocketAddress</b>	Same as <b>-a</b> or <b>--socket-address</b>
<b>SocketFile</b>	Same as <b>-u</b> or <b>--socket-file</b>
<b>SocketQueueSize</b>	Same as <b>-q</b> or <b>--queue-size</b>
<b>SocketTimeout</b>	Same as <b>-o</b> or <b>--socket-timeout</b>
<b>StemWords</b>	Same as <b>-s</b> or <b>--stem-words</b>
<b>ThreadsMax</b>	Same as <b>-T</b> or <b>--max-threads</b>
<b>ThreadsMin</b>	Same as <b>-t</b> or <b>--min-threads</b>
<b>ThreadTimeout</b>	Same as <b>-O</b> or <b>--thread-timeout</b>
<b>User</b>	Same as <b>-U</b> or <b>--user</b>
<b>WordFilesMax</b>	Same as <b>-f</b> or <b>--word-files</b>
<b>WordPercentMax</b>	Same as <b>-p</b> or <b>--word-percent</b>

## EXAMPLES

### Simple Queries

The query:

```
librar*
```

will return all documents that contain “library,” “libraries,” or “librarian.” The query:

```
mouse and computer
```

will return only those documents regarding the kind of mice attached to a computer and not the rodents. The query:

```
cat or kitten or feline
```

will return only those documents regarding cats. The query:

```
mouse or mice and not computer
```

will return only those documents regarding mice (the rodents) and not the kind attached to a computer. The query:

```
mouse and computer or keyboard
```

is the same as:

```
(mouse and computer) or keyboard
```

in that they will both return only those documents regarding either mice attached to a computer or any kind of keyboard. However, neither of those is the same as:

```
mouse and (computer or keyboard)
```

that will return only those documents regarding mice and either a computer or a keyboard.

**Queries Using Meta Data**

The query:

```
author = carroll
```

will return only those documents whose author attribute contains “carroll.” The query:

```
author = stevenson treasure
```

will return only those documents whose author attribute contains “stevenson” and also regarding treasure. The query:

```
author = (lewis carroll)
```

will return only those documents whose author is Lewis Carroll. The query:

```
author = (lewis carroll) or wonderland
```

will return only those documents whose author is Lewis Carroll or that contain the word “wonderland” anywhere in the document regardless of the author.

**Sending Queries to a Search Daemon**

To send a query request to a search daemon using Perl, first open the socket and connect to the daemon (see [Wall], pp. 439-440):

```
use Socket;

$SocketFile = '/tmp/search.socket';
socket( SEARCH, PF_UNIX, SOCK_STREAM, 0 ) or
    die "can not open socket: $!\n";
connect( SEARCH, sockaddr_un( $SocketFile ) ) or
    die "can not connect to \"$SocketFile\": $!\n";
```

Autoflush *must* be set for the socket filehandle (see [Wall], p. 781), otherwise the server thread will hang since I/O buffering will wait for the buffer to fill that will never happen since queries are short:

```
select( (select( SEARCH ), $| = 1)[0] );
```

Next, send a query request (beginning with the word “search” and any options just as with a command-line) to the daemon via the socket filehandle making sure to include a trailing newline since the server reads an entire line of input (so therefore it looks and waits for a newline):

```
$query = 'mouse and computer';
print SEARCH "search $query\n";
```

Finally, read the results back and print them:

```
print while <SEARCH>;
close( SEARCH );
```

**EXIT STATUS**

Exits with one of the values given below:

0 Success.

- 1 Error in configuration file.
- 2 Error in command-line options.
- 40 Unable to read index file.
- 50 Malformed query.
- 51 Could not write to PID file.
- 52 Host or IP address is invalid or nonexistent.
- 53 Could not open a TCP socket.
- 54 Could not open a Unix domain socket.
- 55 Could not **unlink**(2) a Unix domain socket file.
- 56 Could not **bind**(3) to a TCP socket.
- 57 Could not **bind**(3) to a Unix domain socket.
- 58 Could not **listen**(3) to a TCP socket.
- 59 Could not **listen**(3) to a Unix domain socket.
- 60 Could not **select**(3).
- 61 Could not **accept**(3) a socket connection.
- 62 Could not **fork**(2) child process.
- 63 Could not change directory to /.
- 64 Could not create thread.
- 65 Could not detach thread.
- 66 Could not initialize thread condition.
- 67 Could not initialize thread mutex.

## CAVEATS

1. Stemming can be done **only** when searching through and index of files that are in English because the Porter stemming algorithm used only stems English words.
2. When run as a daemon using a TCP socket, there are no security restrictions on who may connect and search. The code to implement domain and IP address restrictions isn't worth it since such things are better handled by routers.
3. XML output can currently only be obtained for actual search results and not word, index, meta-name, or stop-word dumps.

## FILES

swish++.conf        default configuration file name  
 swish++.index     default index file name

## SEE ALSO

**index**(1), **perlfunc**(1), **exec**(2), **fork**(2), **unlink**(2), **accept**(3), **bind**(3), **listen**(3), **select**(3), **swish++.conf**(4), **searchmonitor**(8)

Tim Bray, et al. *Extensible Markup Language (XML) 1.0*, February 10, 1998.

Bradford Nichols, Dick Buttlar, and Jacqueline Proulx Farrell. *Pthreads Programming*, O'Reilly & Associates, Sebastopol, CA, 1996.

M.F. Porter. "An Algorithm For Suffix Stripping," *Program*, 14(3), July 1980, pp. 130-137.

W. Richard Stevens. *Unix Network Programming, Vol 1, 2nd ed.*, Prentice-Hall, Upper Saddle River, NJ, 1998.

Larry Wall, et al. *Programming Perl, 3rd ed.*, O'Reilly & Associates, Inc., Sebastopol, CA, 2000.

## AUTHOR

Paul J. Lucas <paullucas@mac.com>