

# libcurl in the Real World

Dan Fandrich

May 2026  
curl://up 2026



# Using libcurl

- Designed to satisfy many applications
- Large API to make everyone happy
  - 100 functions
  - 351 CURLOPT\_ options
  - 123 CURLE\_ errors
  - Many more: 1270 symbols in total

# Using libcurl

- No application uses it all
- But what parts *do* they use?
- Many are optional
- Some are mutually exclusive or deprecated
  - `curl_form*` vs `curl_mime*`
  - `curl_multi_socket` vs `curl_multi_socket_action`
  - `CURLOPT_PROGRESSFUNCTION` vs `CURLOPT_XFERINFOFUNCTION`
  - `CURLE_OPERATION_TIMEOUTED` vs `CURLE_OPERATION_TIMEDOUT`

# Changing libcurl

- libcurl tries to be highly backward-compatible
- Bug fixes are actually behaviour changes
- Sometimes necessary to
  - Deprecate APIs
  - Change return codes
  - Change behaviour

# What parts are being used?

- How can curl developers know what parts are being used?
- Annual survey gives clues
- Tracking people's issues, feature requests, pull requests

# What parts are being used?

- libcurl is Open Source
- Many libcurl-using applications are as well
- Why not just look?
- Just download the source and see

# Finding Applications

- Just Google it
- Manual, doesn't scale
- Need to manually download hundreds/thousands of apps
- Distributions have done the hard work for us
- Limited to C/C++ applications using libcurl directly

# Choosing a Distro

- Restrict search to Linux only for maximum reach
- Distrowatch lists 1201 distros
- Big one to get as many apps as possible
- Fairly recent to get up-to-date versions
- Easy to find just libcurl-using applications
- Easy to download the source code

# Choosing a Distro

- Arch: 409
- Debian Testing: 446
- Fedora 43: 397
- Fedora 43 w/RPM Fusion: 419
- nixOS: 253
- openSUSE Tumbleweed: 306

# Choosing a Distro

- Settled on Debian Testing
- 441 rdepends packages
- 367 individual source tarballs
- 28.5 GiB source code
- 1,710,001 source files

# Type of Analysis

- Run-time
  - Limited to what features we can figure out
  - Not scalable
- Static binary
  - nm on the executable
  - Skips use of options, error codes, macros, etc.
- Compile of source
  - Too difficult with many different applications
  - Would need tooling to extract the data we want

# Type of Analysis

- Static source
  - Most flexible
  - Access to all features used by the program
  - Access to all build configurations
- Cons:
  - Misses obfuscated symbols, e.g. macros wrapping `CURLOPT_*`
  - Simplistic scanning means some references in e.g. docs count
  - Can include source files using non-C bindings


# Processing

- Some vendored libcurl sources found embedded in packages
- Some copies of `curl.h` found
- Some copies of `symbols-in-versions` found
- Removed (mostly) to avoid diluting results

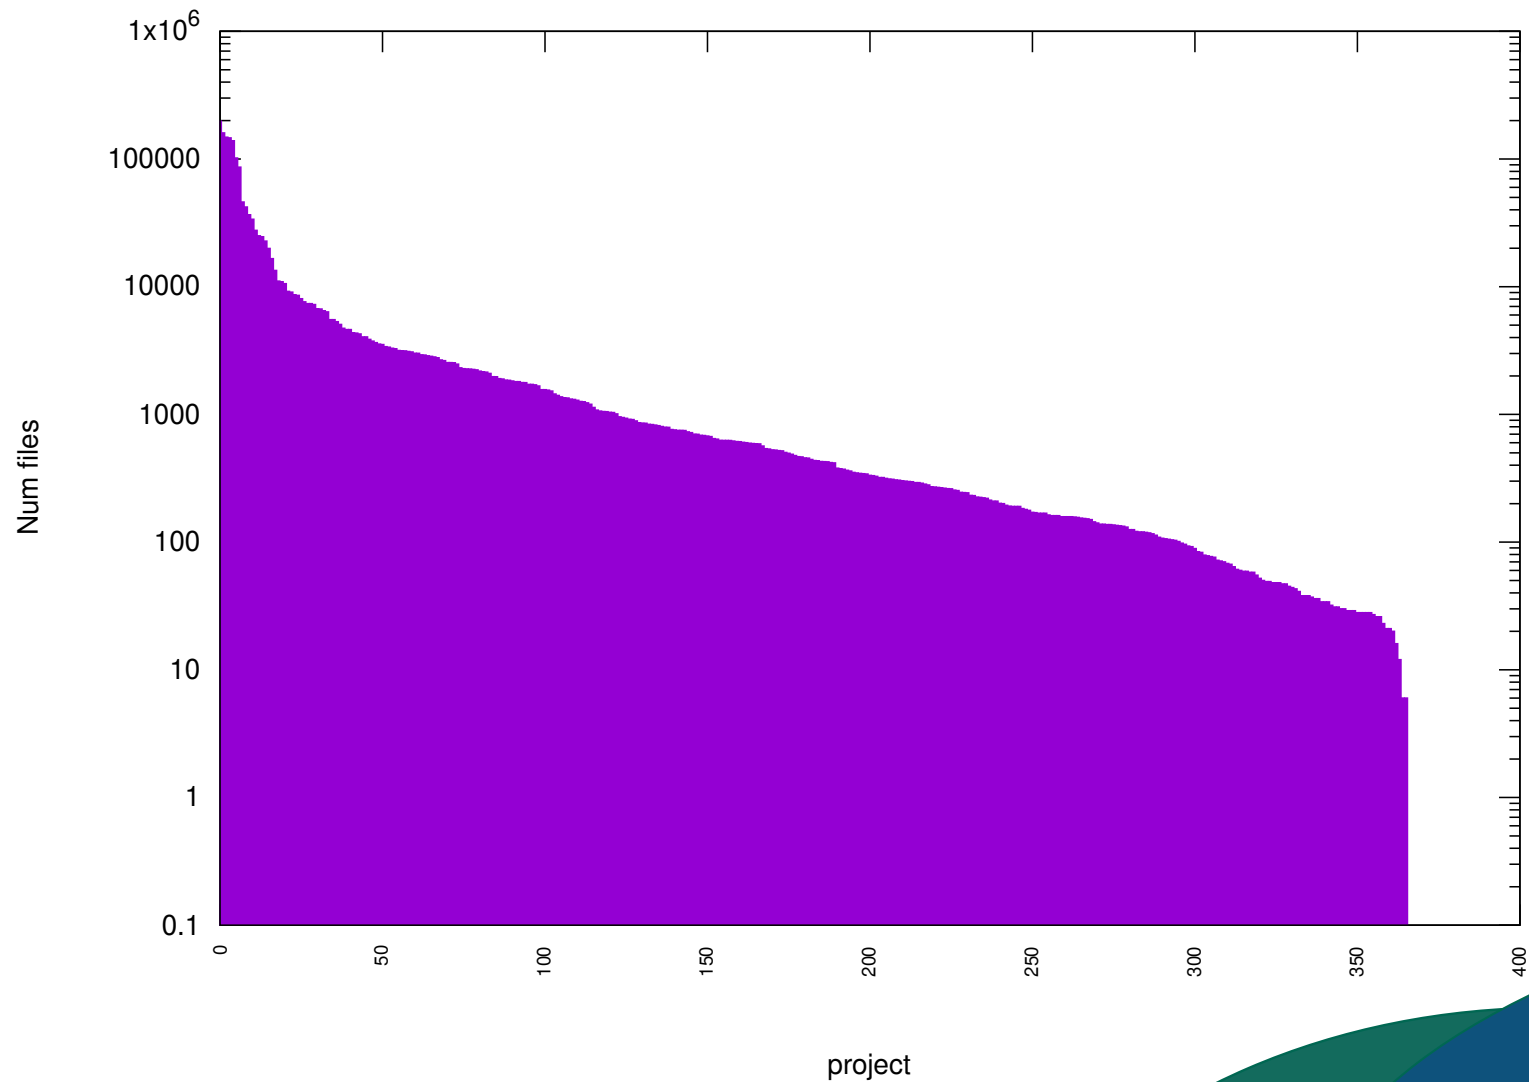
# Processing

- Extract all functions & symbols from curl source code
- Files like `symbols-in-versions`, `test1135`, `VERSIONS.md` made it easy(ish)
- Supplanted with version info from man page `.md` files
- Massive grep of all symbols over all source code
- Results entered into sqlite3 DB


Number of symbols available per version




Number of files in analyzed projects



Size of analyzed projects




Projects using libcurl from each file extension




extensions



Total instances of each symbol across all codebases, ordered by popularity




## Symbols used by most projects




individual symbols




Instances of each symbol, ordered by version of introduction




Instances of each symbol and trend, ordered by version of introduction



Projects using each category of libcurl features (top 50)




Projects using easy vs multi interface



interface




Projects using curl\_form vs curl\_mime interface




interface




# Projects using CURLOPT\_PROGRESSFUNCTION vs CURLOPT\_XFERINFOFUNCTION




Projects using curl\_multi\_socket vs curl\_multi\_socket\_action



Projects using CURLINFO\_PROTOCOL vs CURLINFO\_SCHEME




Projects using CURLOPT\_PUT vs CURLOPT\_UPLOAD



interface



CURLVERSION symbols by use



CURLVERSION symbols



# Unused Symbols

CURLFOLLOW\_ALL

CURLFOLLOW\_FIRSTONLY

CURLFOLLOW\_OBEYCODE

CURLINFO\_HTTPAUTH\_USED

CURLINFO\_PROXYAUTH\_USED

CURLMINFO\_NONE

CURLMINFO\_XFERS\_ADDED

CURLMINFO\_XFERS\_CURRENT

CURLMINFO\_XFERS\_DONE

CURLMINFO\_XFERS\_PENDING

CURLMINFO\_XFERS\_RUNNING

CURLMNOTIFY\_EASY\_DONE

CURLMNOTIFY\_INFO\_READ

CURLMNBC\_CLEAR\_CONNS

CURLMNBC\_CLEAR\_DNS

CURLMNBC\_NETWORK\_CHANGED

CURLMNBC\_NOTIFYDATA

CURLMNBC\_NOTIFYFUNCTION

CURLMNBC\_SSL\_SIGNATURE\_ALGORITHMS

CURLMNBC\_UPLOAD\_FLAGS

CURLLULFLAG\_ANSWERED

CURLLULFLAG\_DELETED

CURLLULFLAG\_DRAFT

CURLLULFLAG\_FLAGGED

CURLLULFLAG\_SEEN

CURLWS\_NOAUTOPONG

CURL\_HAS\_DECLSPEC\_ATTRIBUTE

curl\_easy\_ssls\_export

curl\_easy\_ssls\_import

curl\_multi\_get\_offt

curl\_multi\_notify\_disable

curl\_multi\_notify\_enable

curl\_ws\_start\_frame

# Symbols Used Exactly Once\*

|                          |                            |                              |                        |
|--------------------------|----------------------------|------------------------------|------------------------|
| CURLFORM_ARRAY_END       | CURL_HTTPPOST_CALLBACK     | CURL_PROGRESS_STATS          | curl_multi_waitfds     |
| CURLFORM_ARRAY_START     | CURL_HTTPPOST_FILENAME     | CURL_PUSH_ERROROUT           | curl_mvaprintf         |
| CURLOPT_NOTHING          | CURL_HTTPPOST_LARGE        | CURL_TRAILERFUNC_ABORT       | curl_mvfprintf         |
| CURLSSLBACKEND_AWSLC     | CURL_HTTPPOST_PTRBUFFER    | CURL_TRAILERFUNC_OK          | curl_mvprintf          |
| CURLSSLBACKEND_BORINGSSL | CURL_HTTPPOST_PTRCONTENTS  | CURL_UPKEEP_INTERVAL_DEFAULT | curl_mvsprintf         |
| CURLSSLBACKEND_LIBRESSL  | CURL_HTTPPOST_PTRNAME      | CURL_WIN32                   | curl_pushheader_byname |
| CURLSSLOPT_EARLYDATA     | CURL_HTTPPOST_READFILE     | LIBCURL_COPYRIGHT            | curl_strequal          |
| CURLU_DISALLOW_USER      | CURL_LOCK_TYPE_CONNECT     | LIBCURL_HAS                  | curl_strequal          |
| CURLWS_OFFSET            | CURL_LOCK_TYPE_COOKIE      | curl_easy_nextheader         |                        |
| CURL_EASY_NONE           | CURL_LOCK_TYPE_DNS         | curl_global_trace            |                        |
| CURL_EASY_TIMEOUT        | CURL_LOCK_TYPE_NONE        | curl_mfprintf                |                        |
| CURL_HET_DEFAULT         | CURL_LOCK_TYPE_SSL_SESSION | curl_mprintf                 |                        |
| CURL_HTTPPOST_BUFFER     | CURL_PROGRESS_BAR          | curl_msprintf                |                        |

\* But likely not *really* used



# Future

- What other data would be useful to have?
- Can we use these data to inform future curl decisions?
- What are we in the mood for deprecating now?

Thank-you!

Questions?