

# What can Test Clutch teach us about curl tests?

Daniel Fandrich  
curl://up 2025



# What is Test Clutch?

- Tracks and analyzes automated regression test results
- Retrieves logs from CI services
- Comments on PR test results
- All Open Source



# Reports

- Test Clutch periodically generates static reports
- curl reports are public at <https://testclutch.curl.se/>
- Updated 4 times a day (currently)

# Test report for <https://github.com/curl/curl>

Report generated Wed, 30 Apr 2025 00:55:47 +0000 covering runs over the past 90 days

Hover over cells for more information.

successful test run    successful older test run

\*failed test run \*failed older test run

aborted test run

unknown test run

disabled job

new day

[illegible]

# Test failure counts for test runs on https://github.com/curl/curl

Report generated Tue, 29 Apr 2025 12:52:07 +0000 covering runs over the past 30 days.

Test	Result	Count	Examples
580	UNKNOWN	500	<a href="#">Log</a> <a href="#">Log</a> <a href="#">Log</a> <a href="#">Log</a> <a href="#">Log</a>
1632	FAILIGNORE	472	<a href="#">Log</a> <a href="#">Log</a> <a href="#">Log</a> <a href="#">Log</a> <a href="#">Log</a>
565	UNKNOWN	396	<a href="#">Log</a> <a href="#">Log</a> <a href="#">Log</a> <a href="#">Log</a> <a href="#">Log</a>
576	UNKNOWN	394	<a href="#">Log</a> <a href="#">Log</a> <a href="#">Log</a> <a href="#">Log</a> <a href="#">Log</a>
639	UNKNOWN	389	<a href="#">Log</a> <a href="#">Log</a> <a href="#">Log</a> <a href="#">Log</a> <a href="#">Log</a>
581	UNKNOWN	384	<a href="#">Log</a> <a href="#">Log</a> <a href="#">Log</a> <a href="#">Log</a> <a href="#">Log</a>
641	UNKNOWN	376	<a href="#">Log</a> <a href="#">Log</a> <a href="#">Log</a> <a href="#">Log</a> <a href="#">Log</a>
640	UNKNOWN	361	<a href="#">Log</a> <a href="#">Log</a> <a href="#">Log</a> <a href="#">Log</a> <a href="#">Log</a>
574	UNKNOWN	353	<a href="#">Log</a> <a href="#">Log</a> <a href="#">Log</a> <a href="#">Log</a> <a href="#">Log</a>
642	UNKNOWN	353	<a href="#">Log</a> <a href="#">Log</a> <a href="#">Log</a> <a href="#">Log</a> <a href="#">Log</a>
578	UNKNOWN	343	<a href="#">Log</a> <a href="#">Log</a> <a href="#">Log</a> <a href="#">Log</a> <a href="#">Log</a>
2004	FAILIGNORE	340	<a href="#">Log</a> <a href="#">Log</a> <a href="#">Log</a> <a href="#">Log</a> <a href="#">Log</a>
643	UNKNOWN	330	<a href="#">Log</a> <a href="#">Log</a> <a href="#">Log</a> <a href="#">Log</a> <a href="#">Log</a>
1631	FAILIGNORE	317	<a href="#">Log</a> <a href="#">Log</a> <a href="#">Log</a> <a href="#">Log</a> <a href="#">Log</a>
567	UNKNOWN	315	<a href="#">Log</a> <a href="#">Log</a> <a href="#">Log</a> <a href="#">Log</a> <a href="#">Log</a>
638	UNKNOWN	290	<a href="#">Log</a> <a href="#">Log</a> <a href="#">Log</a> <a href="#">Log</a> <a href="#">Log</a>
632	UNKNOWN	273	<a href="#">Log</a> <a href="#">Log</a> <a href="#">Log</a> <a href="#">Log</a> <a href="#">Log</a>
611	FAILIGNORE	251	<a href="#">Log</a> <a href="#">Log</a> <a href="#">Log</a> <a href="#">Log</a> <a href="#">Log</a>
626	FAILIGNORE	250	<a href="#">Log</a> <a href="#">Log</a> <a href="#">Log</a> <a href="#">Log</a> <a href="#">Log</a>
604	FAILIGNORE	249	<a href="#">Log</a> <a href="#">Log</a> <a href="#">Log</a> <a href="#">Log</a> <a href="#">Log</a>
606	FAILIGNORE	248	<a href="#">Log</a> <a href="#">Log</a> <a href="#">Log</a> <a href="#">Log</a> <a href="#">Log</a>
610	FAILIGNORE	248	<a href="#">Log</a> <a href="#">Log</a> <a href="#">Log</a> <a href="#">Log</a> <a href="#">Log</a>
634	FAILIGNORE	248	<a href="#">Log</a> <a href="#">Log</a> <a href="#">Log</a> <a href="#">Log</a> <a href="#">Log</a>
582	FAILIGNORE	247	<a href="#">Log</a> <a href="#">Log</a> <a href="#">Log</a> <a href="#">Log</a> <a href="#">Log</a>
613	FAILIGNORE	247	<a href="#">Log</a> <a href="#">Log</a> <a href="#">Log</a> <a href="#">Log</a> <a href="#">Log</a>
627	FAILIGNORE	247	<a href="#">Log</a> <a href="#">Log</a> <a href="#">Log</a> <a href="#">Log</a> <a href="#">Log</a>
637	FAILIGNORE	247	<a href="#">Log</a> <a href="#">Log</a> <a href="#">Log</a> <a href="#">Log</a> <a href="#">Log</a>

## Configured Test Job Features on <https://github.com/curl/curl>

Report generated Tue, 29 Apr 2025 12:39:47 +0000 covering jobs over the past 6 days.

- ✓ job has this feature
- ✓ job does not have this feature
- job does not have this mutually exclusive feature
- ? unable to determine if job has this feature
- ? unable to determine if job has this binary feature (but likely no)

The job links to the results of the latest run.

Job	arch: aarch64	arch: evbarm	arch: i386	arch: sparc	arch: x86_64	targetarch: aarch64	targetarch: armv7eb	targetarch: i386	targetarch: i686	targetarch: sparc	targetarch: x86_64	os: cygwin	os: darwin	os: freebsd	os: linux	os: msys	os: netbsd	os: openbsd	os: solaris	buildsystem: automake	buildsystem: cmake/make	buildsystem: cmake/msbuild	buildsystem: cmake/ninja	compiler: AppleClang	compiler: clang	compiler: gcc	compiler: IntelC
<a href="#">[circle] CircleCI / arm-openssl (curl)</a>	✓					✓									✓					✓						✓	
<a href="#">[circle] CircleCI / arm-openssl-c-ares (curl)</a>	✓					✓									✓					✓						✓	
<a href="#">[circle] CircleCI / openssl-c-ares (curl)</a>					✓						✓				✓					✓						✓	
<a href="#">[circle] CircleCI / openssl-libssh (curl)</a>					✓						✓				✓					✓						✓	
<a href="#">[circle] CircleCI / openssl-no-proxy (curl)</a>					✓						✓				✓					✓						✓	
<a href="#">[circle] CircleCI / x86-openssl (curl)</a>					✓						✓				✓					✓						✓	
<a href="#">[curlauto] Kirill Sokolovsky / Linux 4.1.15 x86_64 gcc 13.2.1 Gentoo (curl)</a>					✓						✓				✓					✓						✓	
<a href="#">[curlauto] Marcel Raad / MSYS2 MinGW32 WinThreads SSPI HTTP-only WinSSL (curl)</a>					✓				✓						✓					✓						✓	
<a href="#">[curlauto] Marcel Raad / MSYS2 MinGW64 WinThreads SSPI HTTP-only OpenSSL --disable-proxy --disable-crypto-auth (curl)</a>					✓						✓				✓					✓						✓	
<a href="#">[curlauto] Michael Kaufmann / NetBSD 10.1 ARM big-endian gcc 10.5.0 (curl)</a>		✓					✓										✓			✓						✓	
<a href="#">[curlauto] Dagobert / Solaris 10 I386 GCC 4.9.2 + OpenCSW (ares resolver) (curl)</a>			✓					✓											✓	✓						✓	
<a href="#">[curlauto] Dagobert / Solaris 10 I386 GCC 4.9.2 + OpenCSW (threaded resolver) (curl)</a>			✓					✓											✓	✓						✓	
<a href="#">[curlauto] Dagobert / Solaris 10 I386 GCC 4.9.2 + OpenCSW (curl)</a>			✓					✓											✓	✓						✓	
<a href="#">[curlauto] Dagobert / Solaris 10 SPARC GCC 4.9.2 + OpenCSW (ares resolver) (curl)</a>				✓						✓									✓	✓						✓	
<a href="#">[curlauto] Dagobert / Solaris 10 SPARC GCC 4.9.2 + OpenCSW (threaded resolver) (curl)</a>				✓						✓									✓	✓						✓	
<a href="#">[curlauto] Dagobert / Solaris 10 SPARC GCC 4.9.2 + OpenCSW (curl)</a>				✓						✓									✓	✓						✓	
<a href="#">[curlauto] Dagobert / Solaris 11 SPARC GCC 4.9.2 + OpenCSW (ares resolver) (curl)</a>				✓						✓									✓	✓						✓	
<a href="#">[curlauto] Dagobert / Solaris 11 SPARC GCC 4.9.2 + OpenCSW (threaded resolver) (curl)</a>				✓						✓									✓	✓						✓	
<a href="#">[curlauto] Dagobert / Solaris 11 SPARC GCC 4.9.2 + OpenCSW (curl)</a>				✓						✓									✓	✓						✓	
<a href="#">[gha] dist / maketgz-and-verify-in-tree (curl)</a>				✓							✓				✓					✓						✓	
<a href="#">[gha] dist / verify-out-of-tree-autotools-debug (curl)</a>				✓							✓				✓					✓						✓	
<a href="#">[gha] dist / verify-out-of-tree-docs (curl)</a>				✓							✓				✓					✓						✓	
<a href="#">[gha] Linux HTTP/3 / AM gnutls (curl)</a>				✓							✓				✓					✓						✓	
<a href="#">[gha] Linux HTTP/3 / AM gnutls (pytest)</a>				✓							✓				✓					✓						✓	
<a href="#">[gha] Linux HTTP/3 / AM openssl-quick (curl)</a>				✓							✓				✓					✓						✓	
<a href="#">[gha] Linux HTTP/3 / AM openssl-quick (pytest)</a>				✓							✓				✓					✓						✓	

# Metadata for test runs on <https://github.com/curl/curl>

Report generated Tue, 29 Apr 2025 18:48:59 +0000 covering runs over the past 90 days.

Expand each name to see all its values among recent test runs. Note that not all test runs expose all metadata.

## ▼ arch

- aarch64
- amd64
- arm64
- evbarm
- i386
- riscv64
- sparc
- x86\_64

- ▶ branch
- ▶ buildcode
- ▶ buildsystem
- ▶ buildsystemver
- ▶ checkrepo
- ▶ ciarch
- ▶ cicores
- ▶ cidef
- ▶ cijob
- ▶ ciname
- ▶ cinodetype
- ▶ cios
- ▶ ciresult
- ▶ cistep
- ▶ cistepid
- ▶ cistepresult
- ▶ commit
- ▼ compiler

- APPLECLANG
- AppleClang
- CLANG
- Clang

# Reports: Statistics

Test run statistics for test runs on <https://github.com/curl/curl>  
Report generated Tue, 29 Apr 2025 01:33:53 +0000 covering runs over the past 368 days.

Days of stats: 369

Most recent run: Fri, 25 Apr 2025 23:48:01 +0000

Oldest run used: Thu, 25 Apr 2024 07:46:46 +0000

Number of git commits tested: 2404

Number of unique configured test jobs: 401

Test runs: 260702

Runs per day: 707.0

Runs by test result: 260702 (100%)

- success: 252564 (97%)
- failure: 7084 (2.7%)
- truncated: 1054 (0.40%)

Tests run: 354846350

Tests executed per day: 962295.2

Tests considered: 372440284 (100%)

- PASS: 354485469 (95%)
- FAIL: 9029 (0.0024%)
- SKIP: 17593934 (4.7%)
- FAILIGNORE: 273845 (0.074%)
- UNKNOWN: 78007 (0.021%)

Total clock time spent running tests: 266754193 sec. (3087 days)

Time spent running tests per day: 723401 sec./day (8.4 days/day)

Average time spent running each test: 0.752 sec./test

Longest test runs:

- curl: 13118 sec.
- pytest: 243 sec.

Shortest test runs:

- curl: 16 sec.
- pytest: 72 sec.

Average test runs:

- curl: 420 sec.
- pytest: 125 sec.

Most number of unique tests attempted in one run by test format:

- curl: 1749
- pytest: 673

Average number of tests attempted in one run by test format:

- curl: 1480.4
- pytest: 451.4

Runs by CI system: 260702 (100%)

- appveyor: 1327 (0.51%)
- azure: 12127 (4.7%)
- circle: 19714 (7.6%)
- cirrus: 444 (0.17%)
- curlauto: 10650 (4.1%)

- gha: 216440 (83%)
- Runs by build system: 240593 (100%)
- automake: 171781 (71%)
  - cmake/make: 9741 (4.0%)
  - cmake/msbuild: 10370 (4.3%)
  - cmake/ninja: 48701 (20%)

Runs by test format: 260702 (100%)

- curl: 230469 (88%)
- pytest: 30233 (12%)

Runs by test mode: 230469 (100%)

- normal: 220198 (96%)
- torture: 10271 (4.5%)

Runs by operating systems: 260702 (100%)

- aix: 10 (0.0038%)
- cygwin: 7010 (2.7%)
- darwin: 69613 (27%)
- freebsd: 2120 (0.81%)
- linux: 138805 (53%)
- msys: 27284 (10%)
- netbsd: 2359 (0.90%)
- openbsd: 1988 (0.76%)
- solaris: 11513 (4.4%)






# Logs

- Batch jobs log erroneous output
- Logs are manually downloaded
- Job graph produced on demand
- Helps me make sure things are still working

Test Clutch for curl job durations



# Database

- Central store of test results
- Source of data for all reports
- Source of flakiness status for PR comments
- Can be exported in OpenMetrics format

# Open Metrics

- Born in the Prometheus monitoring project
- Attempt to become IETF standard was abandoned
- Still used by much software as “Prometheus format”
- Simple text-based format for exporting metrics
- Simple example:

```
# TYPE testclutch_tests_duration_seconds gauge
# UNIT testclutch_tests_duration_seconds seconds
# HELP testclutch_tests_duration_seconds How long the tests took to run
(wall clock time).
testclutch_tests_duration_seconds 2650.0 1742118441
```



# Open Metrics

```
testclutch_tests_duration_seconds{arch="sparc",buildcode="2222030376",buildsystem="automake",checkrepo="https://github.com/curl/curl",cijob="Solaris 10 SPARC GCC 4.9.2 + OpenCSW (ares resolver)",ciname="Dagobert",commit="4911e7af119c1b7efd46a742d47bca44832c3041",compiler="GNU_C",compilerversion="5.5.0",compilerversioncode="505",configureargs="--enable-ares' '--with-openssl' '--enable-debug' '--enable-curldebug' '--enable-symbol-hiding' '--with-gssapi=/opt/csw/lib' '--enable-ldap' '--enable-ldaps' 'CC=/opt/csw/bin/gcc' 'LDFLAGS=-L/opt/csw/lib'",curldeps="libcurl/8.13.0-DEV OpenSSL/1.1.1w zlib/1.3.1 brotli/1.0.2 zstd/1.4.5 c-ares/1.14.0 libidn2/2.0.4 libpsl/0.21.5 nghttp2/1.30.0 librtmp/2.3 libgsasl/1.8.0 OpenLDAP/2.4.40",curlprotocols="dict file ftp ftps gopher gophers http https imap imaps ipfs ipns ldap ldaps mqtt pop3 pop3s rtmp rtsp smb smbs smtp smtps telnet tftp ws wss",executor="testcurl",executorver="2024-11-28",features="alt-svc AsynchDNS brotli Debug gsasl GSS-API HSTS HTTP2 HTTPS-proxy IDN IPv6 Kerberos Largefile libz NTLM PSL SPNEGO SSL threadsafe TLS-SRP TrackMemory UnixSockets zstd",host="unstable10s",hostarch="sparc",hostos="solaris2.10",hosttriplet="sparc-sun-solaris2.10",hostvendor="sun",origin="curlauto",os="solaris",perlver="5.10.1",randomseed="223460",runid="build-20250316094727-315931.log",runtestsopts="-a -p -r",systemhost="unstable10s",systemos="SunOS",systemosver="5.10",targetarch="sparc",targetos="solaris2.10",targettriplet="sparc-sun-solaris2.10",targetvendor="sun",testformat="curl",testingver="8.13.0-DEV",testmode="normal",testresult="success",uniquejobname="Solaris 10 SPARC GCC 4.9.2 + OpenCSW (ares resolver) 2222030376!curl",url="https://curl.se/dev/log.cgi?id=20250316094727-315931",withduphandle="no",withevent="no",withvalgrind="no",job="testclutch",instance="test"} 2650.0 1742118441
```





> max(last\_over\_time(testclutch\_tests\_seconds\_count[24h])) by (origin, account, uniquejobname, testformat)

Execute

Table Graph Explain


17w

End time

High res.

Unstacked

Stacked



{origin="circle", testformat="curl", uniquejobname="arm-openssl!curl"}



```
> max(last_over_time(testclutch_tests_seconds_count[24h])) by (origin, account, uniquejobname)
```

1705

origin: gha  
testformat: curl  
uniquejobname: .github/workflows/macros.yml/AM clang OpenSSL libssh c-ares/curl

Execute

Table | Graph | Explain


- 17w +

< End time >

High res.

Unstacked

Stacked





> `sum(last_over_time(testclutch_tests_seconds_count{result="FAIL"}[2h]))`

Execute

Table

Graph

Explain


17w

End time

High res.

Unstacked

Stacked



■ {}



> max(last\_over\_time(testclutch\_tests\_seconds\_count[24h])) by (origin, account, uniquejobname, testformat, result)

Execute

 Table

 Graph

 Explain


7w

< End time >

High res.

Unstacked

Stacked





```
{origin="circle", result="FAIL", testformat="curl", uniquejobname="arm-openssl-c-ares!curl"}
{origin="circle", result="FAIL", testformat="curl", uniquejobname="openssl-c-ares!curl"}
{origin="circle", result="FAIL", testformat="curl", uniquejobname="openssl-no-proxy!curl"}
{origin="circle", result="PASS", testformat="curl", uniquejobname="arm-openssl!curl"}
{origin="circle", result="PASS", testformat="curl", uniquejobname="arm-openssl-c-ares!curl"}
{origin="circle", result="PASS", testformat="curl", uniquejobname="openssl-c-ares!curl"}
{origin="circle", result="PASS", testformat="curl", uniquejobname="openssl-libssh!curl"}
{origin="circle", result="PASS", testformat="curl", uniquejobname="openssl-no-proxy!curl"}
{origin="circle", result="PASS", testformat="curl", uniquejobname="x86-openssl!curl"}
{origin="circle", result="SKIP", testformat="curl", uniquejobname="arm-openssl!curl"}
{origin="circle", result="SKIP", testformat="curl", uniquejobname="arm-openssl-c-ares!curl"}
{origin="circle", result="SKIP", testformat="curl", uniquejobname="openssl-c-ares!curl"}
{origin="circle", result="SKIP", testformat="curl", uniquejobname="openssl-libssh!curl"}
{origin="circle", result="SKIP", testformat="curl", uniquejobname="openssl-no-proxy!curl"}
{origin="circle", result="SKIP", testformat="curl", uniquejobname="x86-openssl!curl"}
{origin="curlauto", result="FAIL", testformat="curl", uniquejobname="NetBSD 10.1 ARM big-endian gcc 10.5.0 127122733!curl"}
{origin="curlauto", result="FAIL", testformat="curl", uniquejobname="Solaris 10 I386 GCC 4.9.2 + OpenCSW (threaded resolver) 2057333316!curl"}
{origin="curlauto", result="FAIL", testformat="curl", uniquejobname="Solaris 10 SPARC GCC 4.9.2 + OpenCSW (ares resolver) 2222030376!curl"}
{origin="curlauto", result="FAIL", testformat="curl", uniquejobname="Solaris 10 SPARC GCC 4.9.2 + OpenCSW (threaded resolver) 2057333316!curl"}
{origin="curlauto", result="FAIL", testformat="curl", uniquejobname="Solaris 11 SPARC GCC 4.9.2 + OpenCSW (ares resolver) 2222030376!curl"}
{origin="curlauto", result="FAIL", testformat="curl", uniquejobname="Solaris 11 SPARC GCC 4.9.2 + OpenCSW (threaded resolver) 2057333316!curl"}
{origin="curlauto", result="FAIL", testformat="curl", uniquejobname="Solaris 11 SPARC GCC 4.9.2 + OpenCSW 2714428864!curl"}
{origin="curlauto", result="PASS", testformat="curl", uniquejobname="Linux 4.1.15 x86_64 gcc 13.2.1 Gentoo 606020064!curl"}
{origin="curlauto", result="PASS", testformat="curl", uniquejobname="MSYS2 MinGW32 WinThreads SSPI HTTP-only WinSSL 1952821939!curl"}
{origin="curlauto", result="PASS", testformat="curl", uniquejobname="MSYS2 MinGW64 WinThreads SSPI HTTP-only OpenSSL --disable-proxy --disable-crypto-auth 2503529752!curl"}
{origin="curlauto", result="PASS", testformat="curl", uniquejobname="NetBSD 10.1 ARM big-endian gcc 10.5.0 127122733!curl"}
{origin="curlauto", result="PASS", testformat="curl", uniquejobname="NetBSD 9.3 ARM gcc 7.5.0 127122733!curl"}
{origin="curlauto", result="PASS", testformat="curl", uniquejobname="Solaris 10 I386 GCC 4.9.2 + OpenCSW (ares resolver) 2222030376!curl"}
{origin="curlauto", result="PASS", testformat="curl", uniquejobname="Solaris 10 I386 GCC 4.9.2 + OpenCSW (threaded resolver) 2057333316!curl"}
{origin="curlauto", result="PASS", testformat="curl", uniquejobname="Solaris 10 I386 GCC 4.9.2 + OpenCSW 2714428864!curl"}
```




> max(last\_over\_time(testclutch\_tests\_seconds\_count[24h])) by (origin, account, uniquejobname, testformat, result)

Execute

Table Graph Explain

- 7w + < End time > High res.

Unstacked Stacked



> max(last\_over\_time(testclutch\_tests\_seconds\_count[24h])) by (testformat, result)

Execute

Table

Graph

Explain


- 17w +

< End time >

High res.

Unstacked

Stacked






> max(last\_over\_time(testclutch\_tests\_duration\_seconds[24h])) by (origin, account, uniquejobname, testformat)

Execute

Table Graph Explain

17w End time High res. Unstacked Stacked



> avg(last\_over\_time(testclutch\_tests\_duration\_seconds[24h])) by (origin, account, testformat)

Execute

Table

Graph

Explain


- 17w +

< End time >

High res.

Unstacked

Stacked



> avg(last\_over\_time(testclutch\_tests\_duration\_seconds{testformat="curl"}[6h])) by (origin, account, testformat, paralleljobs)

Execute

Table

Graph

Explain


- 17w +

< End time >

High res.

Unstacked

Stacked



> max(last\_over\_time(testclutch\_job\_duration\_seconds{systemos=~"(MSYS|CYGWIN|MINGW).\*"}[12h])) by (origin, account, systemos)

Execute

Table

Graph

Explain


- 17w +

< End time >

High res.

Unstacked

Stacked



■ {origin="gha", systemos="CYGWIN\_NT-10.0-20348"}



# PromQL Issues

- Best with continuously-pulled metrics
- Delay of >5m between samples causes discontinuities
- Expand window to avoid gaps, but can become misleading
- Average() or Max() over larger window can combine several commits & job runs
- Most queries here used 6h or 24h windows
- Missing labels can cause missing data if care isn't taken

# Questions?

