

Squeezing Water from Stone:

KornShell in 2019

```
main() { printf(&unix["\021%six\012\0"], (unix) ["have"]+"fun"-0x60); }
```

- Winner of IOCCC (The International Obfuscated C Code Contest) one line entry in 1987.
- Termed as best one liner ever received.
- Written by David Korn.

What this talk is about ?

- It's about AT&T KornShell (and not mksh/pdksh).
- Brief status update from new upstream maintainers (Kurtis Rader and Siteshwar Vashisht).
- No deep dive in the source code.

/whois situ

- Siteshwar Vashisht.
- Based in Brno, Czech Republic.
- Maintainer of bash and ksh at Red Hat.
- Previously involved with fish shell and Sailfish OS.
- Current upstream maintainer of ksh.

What is KornShell ?

- Direct descendant of Bourne shell.
- Initial development based on original bourne shell.
- Roots back to late 70s/early 80s.

Advantages

- Superior language specification (Support for Discipline functions, namespaces, compound data type, enums etc.).
- Better POSIX compliance.
- Fastest POSIX shell.

A 20 year old bug

`whence -a` is broken if builtin for a command is enabled
#954

 Closed **siteshwar** opened this issue on Oct 14, 2018 · 10 comments

Edit

New issue



siteshwar commented on Oct 14, 2018

Collaborator



Description of problem:

whence -a is broken if builtin for a command is enabled

Ksh version:

2017.0.0-devel-1945-gc7752ec6-dirty

How reproducible:

Always

Steps to reproduce:

```
whence -a cat
builtin cat
whence -a cat
```

Actual results:

```
cat is a shell builtin
cat is /usr/bin/cat
cat is a tracked alias for /home/situ/upstream/ast/build/cat
```

Expected results:

```
cat is a shell builtin
cat is /usr/bin/cat
cat is a tracked alias for /usr/bin/cat
```

Additional info:

This bug is present in `ksh-20120801`.

Assignees



No one—assign yourself

Labels



bug

Projects



None yet

Milestone



No milestone

Notifications

Unsubscribe

You're receiving notifications because you modified the open/close state.

4 participants



Lock conversation

Pin issue

State of ksh in 2017

- Previous developers have left.
- Codebase has origins going back to the 80s.
- Old Tooling.
- No revision control history.
- Almost no comments in the source code.
- Spaghetti code.
- Bad test coverage.
- Lots of old bugs.
- And ...

How to maintain the unmaintainable ?

```
< Don't change anything! >
-----
\      ^ ^
 \    (oo)\_____
      (__) \         )\/\
           ||-----w ||
           ||         ||
```


One of my most productive days was throwing away 1,000 lines of code - Ken Thompson

We threw away more than 500,000 lines of code.

Lines of code (ksh93v-)

files	blank	comment	code
2478	66068	135004	665597

Lines of code (current)

files	blank	comment	code
629	13363	21047	137376

What did we drop ?

- Support for non-POSIX operating systems.
- Reimplementation of POSIX functions.
- Entire subsystems like the AST vmalloc and locale.
- Lots of code that was not used by ksh.
- All external commands (e.g., grep) other than ksh were dropped.

The project is now focused solely on ksh rather than providing replacements for commonly used commands found in BSD, Linux, and SysV.

Refactor Code

```
if (path[0] == '/' && path[1] == 'd' && path[2] == 'e' && path[3] == 'v' && path[4] == '/') {
    switch (path[5]) {
        case 'f': {
            if (path[6] == 'd' && path[7] == '/') {
                if (flags == 0 NONBLOCK) return (1);
                fd = (int)strtol(path + 8, &e, 10);
                if (*e) fd = -1;
            }
            break;
        }
        case 's': {
            if (path[6] == 't' && path[7] == 'd') {
                switch (path[8]) {
                    case 'e': {
                        if (path[9] == 'r' && path[10] == 'r' && !path[11]) fd = 2;
                        break;
                    }
                    case 'i': {
                        if (path[9] == 'n' && !path[10]) fd = 0;
                        break;
                    }
                    case 'o': {
                        if (path[9] == 'u' && path[10] == 't' && !path[11]) fd = 1;
                        break;
                    }
                }
            }
        }
    }
}
```

Refactor Code


```
if (strncmp(path, "/dev/", sizeof("/dev/") - 1) == 0) {  
    if (strncmp(path, "/dev/fd/", sizeof("/dev/fd/") - 1) == 0) {  
        if (flags == O_NONBLOCK) return 1;  
        fd = (int)strtol(path + sizeof("/dev/fd/") - 1, &e, 10);  
        if (*e) fd = -1;  
    } else if (strcmp(path, "/dev/stdin") == 0) {  
        fd = 0;  
    } else if (strcmp(path, "/dev/stdout") == 0) {  
        fd = 1;  
    } else if (strcmp(path, "/dev/stderr") == 0) {  
        fd = 2;  
    }  
}
```


Here Be Dragons

```
sh_ioinit(shp);
shp->pwdfd = sh_diropenat(shp, AT_FDCWD, e_dot);
#if 0 SEARCH
/* This should never happen, guaranteed by design and goat sacrifice */
if(shp->pwdfd < 0)
    errormsg(SH_DICT,ERROR_system(1), "Can't obtain directory fd.");
#endif

/* initialize signal handling */
sh_siginit(shp);
stakinstall(NIL(Stak_t*),nospace);
/* set up memory for name-value pairs */
shp->init_context = nv_init(shp);
/* read the environment */
if(argc>0)
{
    shgd->shtype= type = sh_type(*argv);
    if(type&SH_TYPE_LOGIN)
        shp->login_sh = 2;
}
```


Better CI

✓ # 2631.1	 </> Compiler: gcc C	 DISTRO_TYPE=fedora INSTALL_REQUIREME	 4 min 21 sec	
✓ # 2631.2	 </> Compiler: gcc C	 DISTRO_TYPE=opensuse INSTALL_REQUIRE	 2 min 45 sec	
✓ # 2631.3	 </> Compiler: gcc C	 DISTRO_TYPE=ubuntu INSTALL_REQUIREM	 3 min 12 sec	
✓ # 2631.4	 </> Compiler: gcc C	 DISTRO_TYPE=debian INSTALL_REQUIREM	 3 min	
✓ # 2631.5	 </> Compiler: gcc C	 DISTRO_TYPE=i386/ubuntu INSTALL_REQU	 3 min 28 sec	
✓ # 2631.6	 </> Compiler: clang C	 DISTRO_TYPE=macOS	 8 min 21 sec	
✓ # 2631.7	 </> C	 SCAN_TYPE="shellcheck" DISTRO_TYPE=fe	 53 sec	

Better CI

2 checks passed



FreeBSD Build passed

[Details](#)



continuous-integration/travis-ci/pr The Travis CI build passed

[Details](#)

Test Coverage

Directory: `../src/`






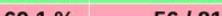


































Date: 2019-01-23 07:21:19

Legend: low: < 75.0 % medium: >= 75.0 % high: >= 90.0 %

Exec Total Coverage

Lines: 33221 51000 65.1 %

Branches: 28201 52229 54.0 %

File	Lines	Branches
<code>cmd/ksh93/bltins/cd_pwd.c</code>	 85.1 % 160 / 188	 74.1 % 126 / 170
<code>cmd/ksh93/bltins/cflow.c</code>	 89.3 % 50 / 56	 90.0 % 45 / 50
<code>cmd/ksh93/bltins/enum.c</code>	 82.2 % 125 / 152	 69.1 % 56 / 81
<code>cmd/ksh93/bltins/getopts.c</code>	 95.2 % 120 / 126	 81.1 % 60 / 74
<code>cmd/ksh93/bltins/hist.c</code>	 77.3 % 136 / 176	 69.2 % 83 / 120
<code>cmd/ksh93/bltins/math.c</code>	 23.6 % 13 / 55	 17.4 % 8 / 46
<code>cmd/ksh93/bltins/misc.c</code>	 83.9 % 239 / 285	 72.9 % 129 / 177
<code>cmd/ksh93/bltins/print.c</code>	 81.6 % 519 / 636	 73.4 % 350 / 477
<code>cmd/ksh93/bltins/read.c</code>	 84.7 % 415 / 490	 76.4 % 412 / 539
<code>cmd/ksh93/bltins/sleep.c</code>	 70.2 % 59 / 84	 60.5 % 49 / 81
<code>cmd/ksh93/bltins/test.c</code>	 73.0 % 259 / 355	 58.4 % 241 / 413
<code>cmd/ksh93/bltins/trap.c</code>	 79.4 % 173 / 218	 65.9 % 149 / 226
<code>cmd/ksh93/bltins/typeset.c</code>	 86.5 % 773 / 894	 79.3 % 732 / 923
<code>cmd/ksh93/bltins/ulimit.c</code>	 84.0 % 79 / 94	 77.9 % 67 / 86
<code>cmd/ksh93/bltins/umask.c</code>	 88.9 % 40 / 45	 92.0 % 23 / 25
<code>cmd/ksh93/bltins/whence.c</code>	 87.8 % 159 / 181	 77.8 % 123 / 158
<code>cmd/ksh93/edit/completion.c</code>	 57.8 % 189 / 327	 33.0 % 144 / 437
<code>cmd/ksh93/edit/edit.c</code>	 57.3 % 409 / 714	 41.9 % 265 / 633
<code>cmd/ksh93/edit/emacs.c</code>	 60.7 % 465 / 766	 50.3 % 260 / 517
<code>cmd/ksh93/edit/hexpand.c</code>	 33.6 % 115 / 342	 25.9 % 90 / 347

Test Coverage

✓ **application_compatibility** Trigger application compatibility t

↻ #2630 passed

↻ Restart build

↻ Commit 163e168 ↗

🕒 Ran for 2 min 42 sec

🔗 Compare 53d45eb...163e168 ↗

🕒 Total time 4 min 18 sec

🔗 Branch application_compatibility ↗

📅 about 8 hours ago

○ kshbugbot

Build jobs

View config

✓ # 2630.1



</> Ruby



PROJECT_NAME=kshdb GIT_URL="https://"

🕒 2 min 37 sec



✓ # 2630.2



</> Ruby



PROJECT_NAME=yash GIT_URL="https://"

🕒 1 min 41 sec



Test Coverage


All Red Hat internal tests are executed on every upstream build.

Call for Testers


You can be a crazy tester.

Coverity Defect Rate Graph

Outstanding vs Fixed defects over period of time



Defect Density over period of time



The graph compares the defect density of the project with the average defect density of open source projects that are similar in size (i.e. 100,000 to 499,999 lines of code)

It's simple to try

@ksh / **latest**

 Overview

 Packages

 Builds

 Modules

 Monitor

Description

Development builds from master branch of GitHub.

Installation Instructions

On Fedora:

```
dnf copr enable @ksh/latest
dnf install ksh
```

On RHEL/CentOS:

```
yum install yum-plugin-copr
yum copr enable @ksh/latest
yum install ksh
```

It's simple to try

openSUSE Build Service > Projects > shells:ksh:new:latest > ksh

[Sign Up](#) | [Log In](#)

[Overview](#) [Repositories](#) [Revisions](#) [Requests](#) [Users](#) [Advanced](#)

Development builds from master branch of GitHub.




 [Download package](#)

No description set

Source Files

Show entries

Search:

Filename	Size	Changed	Actions
ksh.spec	1.73 KB	about 2 hours ago	
ksh_2017.0.0+git.2244.gcbf6c339.dsc	282 Bytes	about 2 hours ago	
ksh_2017.0.0+git.2244.gcbf6c339.orig.tar.gz	1.97 MB	about 2 hours ago	

Showing 1 to 3 of 3 entries

[Previous](#)

1

[Next](#)

Latest Revision



[Siteshwar Vashisht \(siteshwar\)](#) committed [about 2 hours ago](#) (revision 455)

Build from cbf6c3392d50c79fa52b65cb5e662c988e7db1f9



[Files changed](#)



[Browse Source](#)

[Build Results](#)



[Rpmlint Results](#)

ksh

openSUSE_Leap_15.0	 x86_64	succeeded 
openSUSE_Leap_42.3	 x86_64	succeeded 
openSUSE_Tumbleweed	 i586	succeeded 
xUbuntu_17.10	 x86_64	succeeded 
xUbuntu_18.04	 x86_64	succeeded 

Is it too late to make a change ?

commit a10c3a13c3313242704efb78b94dfc801592b4a2

Author: Siteshwar Vashisht <svashisht@redhat.com>

Date: Wed Dec 5 13:48:03 2018 +0100

`whence -a` should print correct path for tracked aliases

`whence -a` should print path that a tracked alias refers to instead of associating it with current directory.

This bug was reported to previous maintainers around 2 decades ago, but it was never fixed. Original credits to @DavidMorano for finding and reporting this bug.

Better late than never!

I

Resolves: #954

Thank you!