

PROGRAMS ON LINUX COMPUTERS
FOR ANYONE WHO WRITES (OR RUNS!!)
A SMALL TOOL HANDBOOK
WIZARD



DEBBUGGING TOOLS
→ LINUX ←

You can print more!
for free or
<http://jvns.ca/zines>
liked this?

Hi! This is me:



JULIA EVANS
blog: jvns.ca ☺
twitter: @b0rk

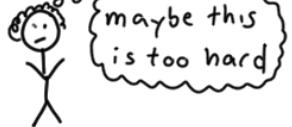
and in this zine I want to tell you about

how I got better
at debugging

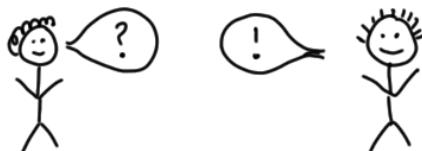
These are 5 ways I've changed how I think about
debugging:

1 Remember the bug is happening
for a logical reason. There's no magic.

2 Be confident I can fix it

before:  now: 

3 Talk to someone



phew

I hope you learned
something new.
Thanks for reading ♥

Thanks to my partner Kamal for
endless reviews, to the amazing
Monica Dinculescu (@notwaldorf)
for the cover art, and many others.

If you want to know more - my
site has a lot (jvns.ca) and
brendangregg.com does too.

But really you just need to experiment.
Try these tools everywhere. See
where they help and where they don't.
It takes a lot of practice to use these
tools to debug real problems.

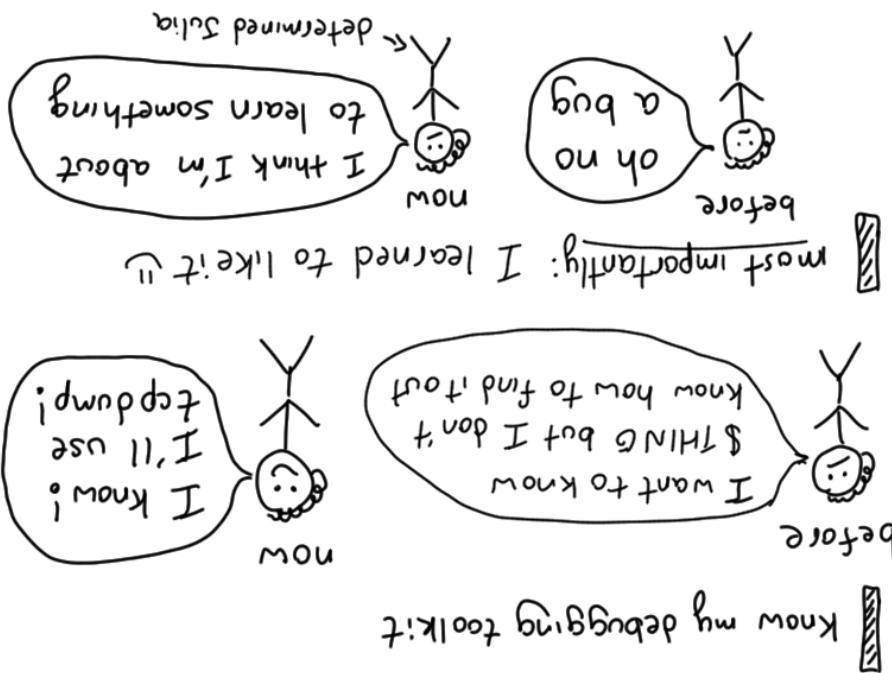
I've been learning them for 2 years, and
I've gotten pretty far, but there's a long
way to go. It's really fun ☺☺☺

you 1-2 new microscopes to use.
are doing. I hope at the end to have given
a question about what my programs

The tools I each for when I have

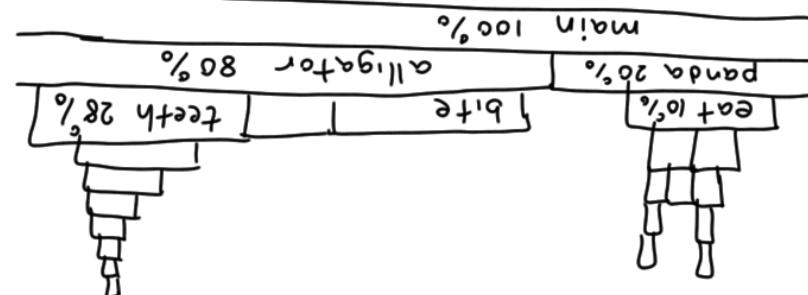
tools I toolkit:
show you some of my debugging:
confident or to my debugging. I can
(I'll try anything)
I can't teach you in 20 pages to be

What you'll learn



You can construct them from perf recordings
(look up "Brendan Gregg's flamegraph", for how), but lots of
other unrelated tools can produce them too. I do them.

and 10% with "main"
of the stack traces started with "alligator",
a program. This one above means 80%
thousands) of stack traces sampled from
They're constructed from collections (usually



Here's what they look like:

Flamegraphs are an awesome way to
visualize CPU performance, popularized by
Brendan Gregg's Flamegraph.pl tool.

flamegraphs

Section 1: I/O and * system calls *

Hello, dear reader! In this zine, there are 3 sections of tools that I love.

For each tool, I'll tell you why it's useful and give an example. Each one is either



or



Some of the most basic questions you might have when you log into a misbehaving machine are:

- is this machine writing to or reading from disk? The network?
- are the programs reading files? Which files?

So, we're starting with finding out which resources are being used and what our programs are doing. Let's go!

spy on your CPU!

Your CPU has a small cache on it (the L1 cache) that it can access in ~0.5 nanoseconds! faster than RAM!

* tip:
google "Latency Numbers every programmer should know!"

If you're trying to do an operation in microseconds, CPU cache usage matters!

: how do I know if my program is using those caches?

perf stat! : how

to use it : how to use it . perf stat -e L1-dcache-load-misses ; ls ;

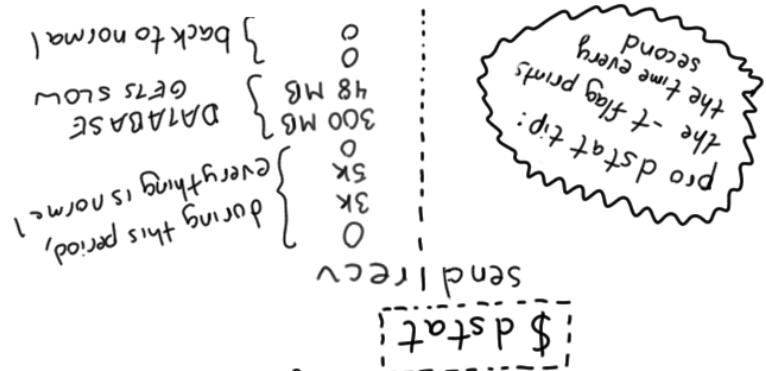
This runs 'ls' and prints a report at the end.

how it works

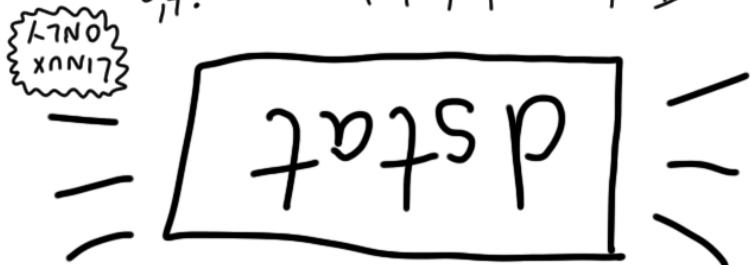
Your CPU can keep all kinds of counters about what it's doing. `perf stat` asks it to count things + then collects the result.

: Hardware is cool. knowing more about how your hardware works can really pay off *

helped us isolate the problem query.
This was an AWEsome clue that
= YES ð =
actually mean ... a 300MB data base query??
Could 300MB coming in over the



the output while monitoring database speed.
Once I had an interim intermittently slow database
server. I opened up dstat and stared at
the output while monitoring database speed.
disk your computer used that second.
prints out how much network and
super simple. Every second, it's
I love dstat because it's



Remember when I said perf only knows
functions? It's not quite true. node.js
and the JVM (java, scala, clojure...) have both
taught perf about their functions.
use the node = Java
perf-basic-profile
on GitHub and follow
look up 'perf-map-agent',
the directions
command line option

... especially Java and node devs
when it does help.
Far from it. But it's an easy
tool to try, and it's always helpful
perf top doesn't always help.
Ruby's internal "perf top"
"regexp matching function"
"Ruby's internal"
"when it does help."
process PID % function
ruby 1957 77 match-a+

\$ sudo perf top
in Ruby. How?
I knew it was doing regular expression matching
100% of its CPU. Within about 60 seconds,
One day, I had a server that was using
perf is for everyone

strace

LINUX ONLY

(I have a strace
sticker on my phone)

Strace is my favourite program. It prints every system call your program used. It's a cool way to get an overall picture of what your program is doing, and I ❤️ using it to answer questions like "which files are being opened?"

```
! $ strace python my_program.py:  
-----  
read(3, "the contents of the file") = 3  
-----  
a file!  
{open("/home/bork/.config_file") = 3  
... hundreds of lines...  
file descriptor  
-----  
read(3, "the contents of the file") = 3  
-----  
a file!  
{connect(5, "172.217.0.163")  
... hundreds of lines...  
networking!  
{sendto(5, "hi!!")
```



strace can make your program run 50x slower. Don't run it on your production database

I can't do justice to strace here, but I have a whole other zine about it at

jvns.ca/zines

perf

perf is not simple or elegant. It is a weird multitool that does a bunch of different, very useful things. First! It's a:

sampling profiler

Try running:

```
! $ sudo perf record python  
-----  
(press Ctrl+C after 2 seconds)  
-----  
saves a file "perf.data"
```

This records, every few milliseconds, what the python process is doing. Let's see the results!

```
! $ sudo perf report
```

Mine says it spent 5% of its time in the PyDict_GetItem function. Cool! We learned a tiny thing about the Python interpreter!

just C functions

If you're a Python/Ruby/Java/Node programmer, you might be getting antsy.

"I want to know which Ruby function is running! Not the C function!"

works everywhere

perf can be installed on pretty much any Linux machine. The exact features it has will depend a little on your kernel version.

Stick with me though. I get you.

SECTION 3: CPU + perf

Your programs spend a lot of time on the CPU! Billions of cycles.

This section is about using `perf` to answer that question, a Linux-only tool that is extremely useful and not as well-known as it should be!!

(in general, my aim in this zine is to showcase tools that I don't think get enough )

Some things I didn't have space for in this section but I want to mention anyway:

-valgrind
-the Java ecosystem's fantastic tools (jstacK, VisualVM, Mission Control, YourKit)
which your language is probably jealous of
`perftrace` (for Linux kernel perf problems)

-EBPF

= how it works =
are being started.
execsnoop tells you what programs
opennoop won't slow you down.
program run `lsof` slower.
strace can make your
program just use `strace -e open-p $PID`
and you should be right. But
being opened by a program. You might think...
it will print out `in real time` every file

When you run
`opennoop -p $PID`

OS X
too!

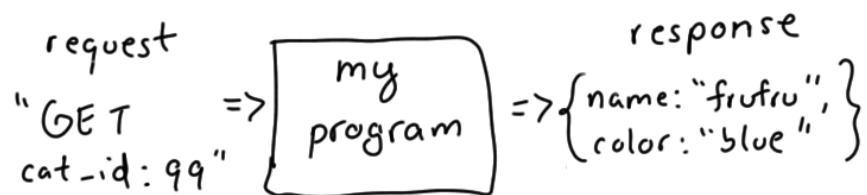
opennoop!
EBPF!
execsnoop!

That one is powered by DTrace. EBPF is fast!
feature called `EBPF`.
opennoop is a script that uses a new kernel or a ~4.4+ kernel version
Requires: Ubuntu 16.04+
Installations:
github.com/lcovisor/bcc-tools
That one is powered by DTrace. EBPF is fast!
There's also an opensnoop on OSX & BSD!
That one is powered by DTrace. EBPF is fast!
feature called `EBPF`.
opennoop is a script that uses a new kernel or a ~4.4+ kernel version
Requires: Ubuntu 16.04+
Installations:
[github.com/lcovisor/bcc-tools](http://gitgithub.com/lcovisor/bcc-tools)
This won't work on many servers today, but keep it in mind! One day you'll have servers running a newer kernel.

section 2: networking

I've devoted a lot of space in this zine to networking tools, and I want to explain why.

A lot of the programs I work with communicate over HTTP.



Every programming language uses the same network protocols! So the network is a nice language-independent place to answer questions like:

- was the request wrong, or was it the response?
- is my server even on?
- my program is slow. Whose fault is it?

Let's go!

wireshark

OS X
too!

Wireshark is an amazing GUI tool for network analysis. Here's an exercise to learn it! Run this:

sudo tcpdump port 80 -w http.pcap

While that's running, open metafilter.com in your browser. (or jvns.ca!). Then press Ctrl+C to stop tcpdump. Now we have a pcap!

Wireshark http.pcap

Explore the Wireshark interface!

Questions you can try to answer:

- ① What HTTP headers did your browser send to metafilter.com?

(hint: search |frame contains "GET"|)

- ② How long did the longest request take?

(hint: click Statistics \rightarrow Conversations)

- ③ How many packets were exchanged with metafilter.com's server?

(hint: search |ip.dst == 54.186.13.33|)
IP address from
ping metafilter.co

HTTP requests are fundamentally
really simple — they're just text! To see
that, let's make one by hand. First, make a
file that says

```
GET / HTTP/1.1
request.txt
Host: ask.metafilter.com
User-Agent: Zine
(2 new lines! important!!)
```

Then:

```
$ cat request.txt | nc metafilter.com 80
```

You should get a response back with a bunch
of HTML! You can also use netcat to send
huge files over a local network quickly:

```
$ hostname - I
at bigfile
$ nc 192.168.2.132 ...
$ nc 192.168.2.132 9931 > bigfile
$ nc -l 9931 > bigfile
```

This means on the port;

- I'm sending a request to a machine and
- I want to know whether it's even getting there
- (tcpdump port 80 will print every packet on port 80)
- I have some slow network connections and
- I want to know whether to blame the client or server. (we'll also need wireshark!)
- I just want to print out packets to
- see them (tcpdump -A)

Some situations where I'll use tcpdump:

a "pcap file" ("packet capture") is the standard for saving network traffic. Everything understands pcap.

sudo tcpdump port 8997, is actually a tiny program in the Berkley Packet Filter (BPF) language. BPF filters get compiled and they run really fast!

w service pcap —

analyze later!

I use it to save network traffic to

and it took me a while to do it.

networking tools we'll discuss here,

tcpdump is the most difficult



tcpdump !!

networking
art is an
handcuffed

NetCat

netstat

Every network request gets sent to a port (like 80) on a computer. To receive a request, a program (aka "server") needs to be "listening" on the port. Finding out which programs are listening on which ports is really easy. It's just

☆ "tuna, please!" ☆

also known as

`sudo netstat -tunap`

thanks to
@icco for
the tuna!

Here's what you'll see:

<u>proto</u>	<u>local address</u>	PID / program name
tcp	0.0.0.0:5353 port ↴	2993 / python

So! I ❤️ netstat because it tells me which processes are running on which ports.

On OSX, use `lsof -i -P` instead.

OS X
too!

ngrep

OS X
too!



ngrep is my favourite starter network spy tool! Try it right now! Run:

`-----
sudo ngrep -d any metafilter
-----`

Then go to <http://metafilter.com> in your browser. You should see matching network packets in ngrep's output! We are SPIES !!

Recently at work I'd made a change to a client so that it sent `{"special-id": ...}` with all its requests. I wanted to make sure it was working, so I ran

`-----
(sudo ngrep special_id)
-----`

I found out that everything was ok !!