

VOXOZ

Vladimir Kirillov [@darkproger](#)

ops



DESIGNED FOR CLOUDS



ERLANG



ERLANG ON XEN



docker

DOCKER

E2MP PLATFORMS

let's reimagine

Hack

Deploy

Scale

Hack
Deploy
Scale

**life's too short to
spend time on the
environment**

Hack

quickly

Hack reproducibly

Hack

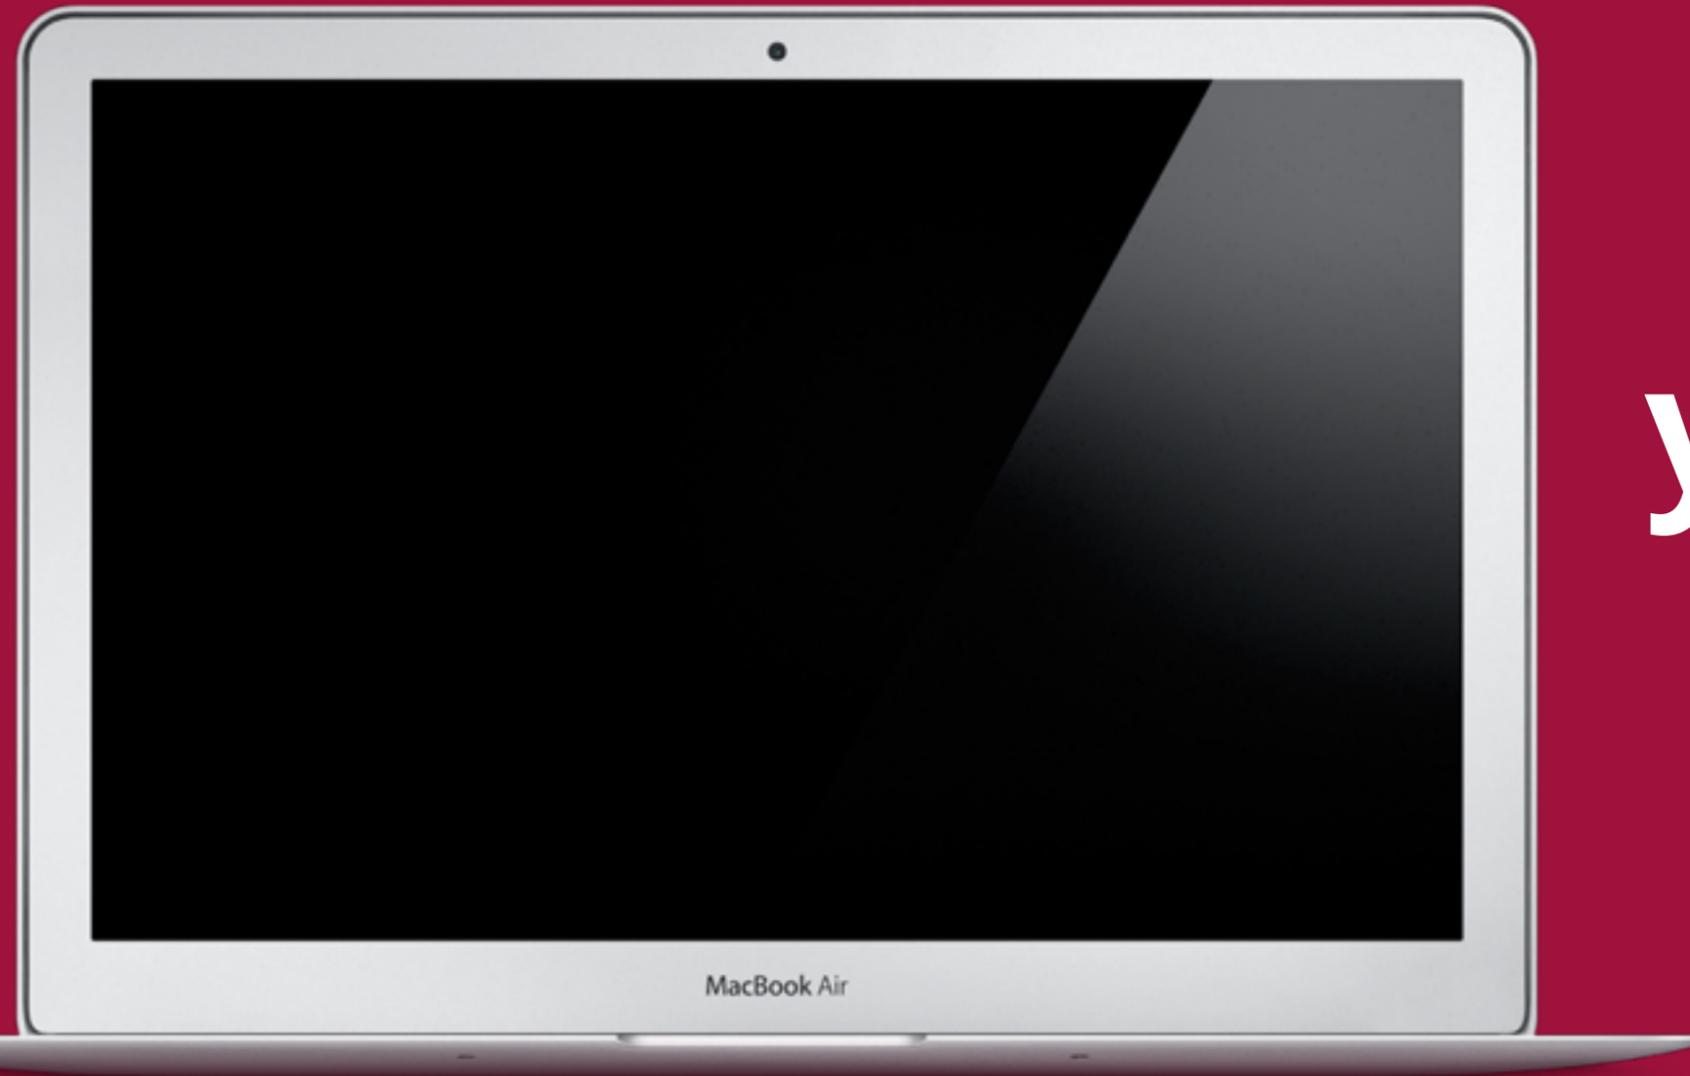
no frictions

Hack

no “works on
my machine”

excuses

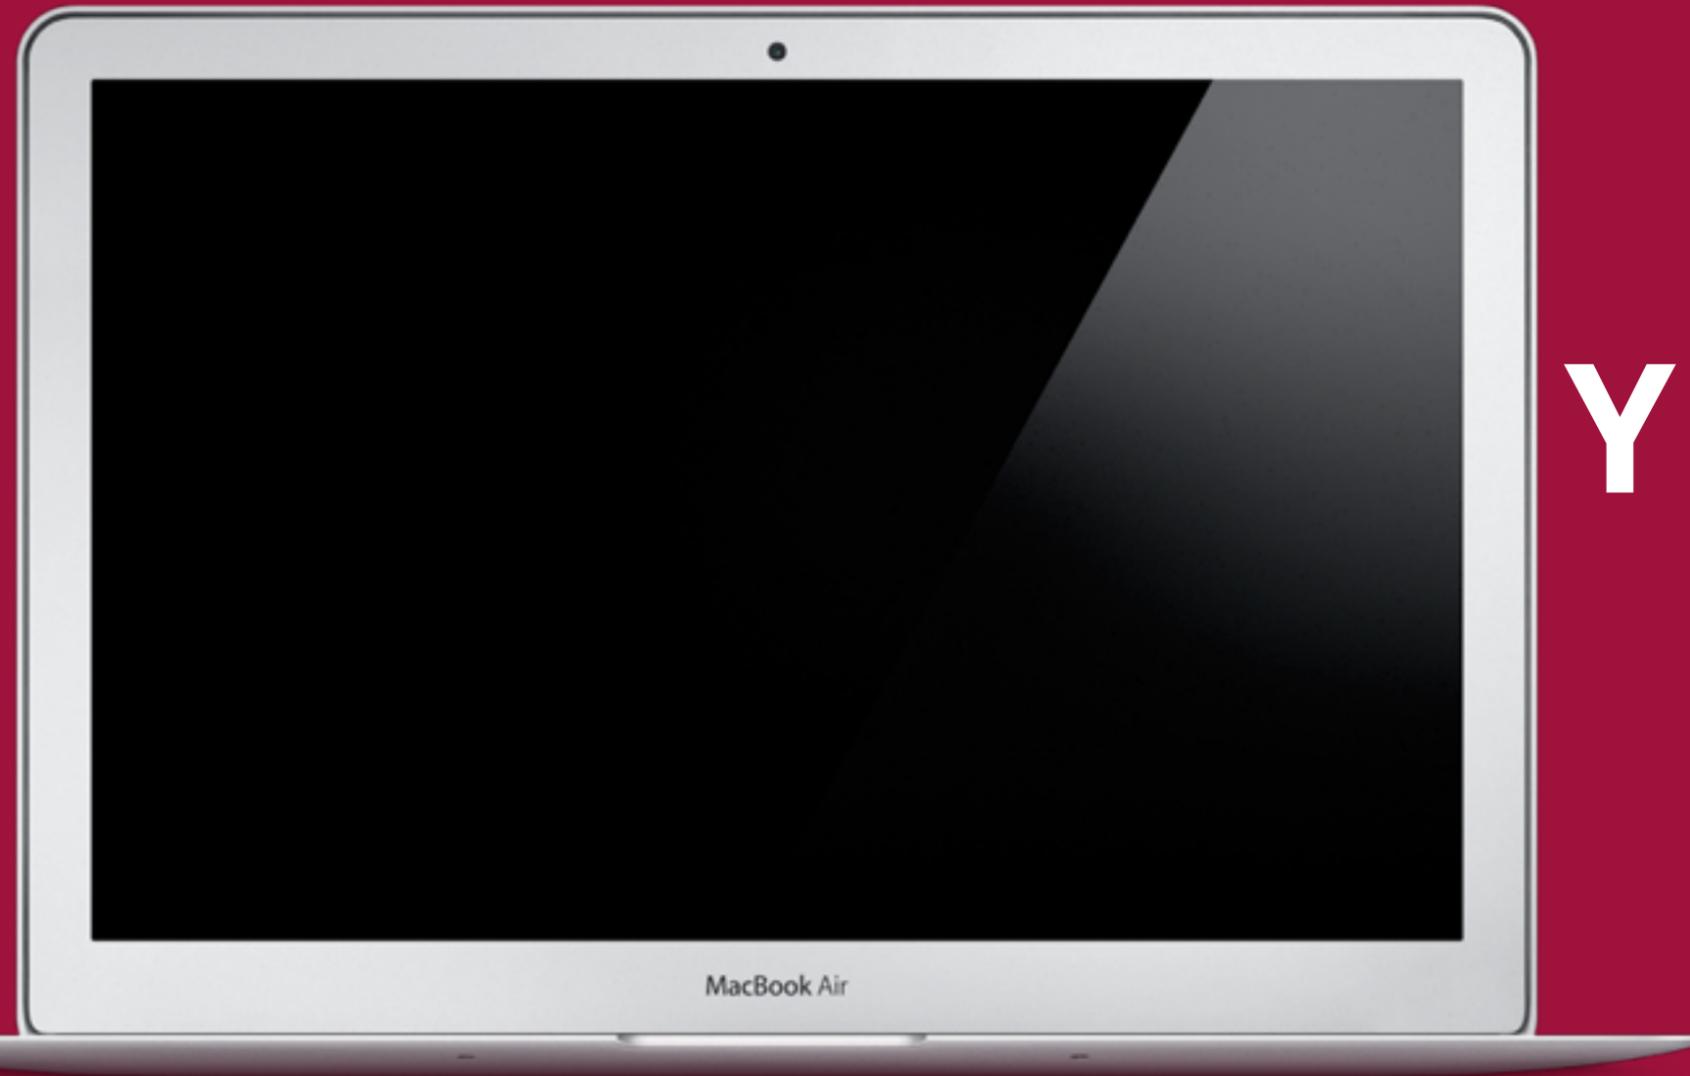
**your epoll into
my kqueue**



your inotify into
my fsevents



**Y UR CentOS runs
R14B03?**



**so I compiled this
port and copied it
to your Linux box
and it doesn't run**

:(



Deploy

Deploy no moving parts

**Moving
parts**

3rd-party software:

OS

runtime/VM

libraries

**3-rd party
software**

**using = being a
maintainer**

**3-rd party
software**

**isolate breakage
as much as
possible**

state

Deploy = mutating state

Linux

**too many
abstractions**

Linux

users

ACLs

SELinux

networking stack

package managers

filesystems

```
erts-5.9.1/bin/escrpt:  
/lib/libc.so.6: version GLIBC_2.14 not found
```

Problems

half-deploys

speed

rollbacks

chef / puppet

```
# apt-get install linux-image-3.2.0-4-amd64
```

```
...
```

```
Setting up linux-image-3.2.0-4-amd64 (3.2.51-1) ...
```

```
Running depmod.
```

```
cp: cannot stat `/boot/initrd.img-3.2.0-4-amd64': No such file  
or directory
```

```
Failed to copy /boot/initrd.img-3.2.0-4-amd64 to /initrd.img .
```

```
dpkg: error processing linux-image-3.2.0-4-amd64 (--configure):  
 subprocess installed post-installation script returned error  
exit status 1
```

```
Errors were encountered while processing:
```

```
 linux-image-3.2.0-4-amd64
```

```
E: Sub-process /usr/bin/dpkg returned an error code (1)
```

**Breaking
abstractions**

zerocopy userspace

Breaking abstractions

netmap
/dev/blk
paravirt

Scaling

Scaling spawn X instances of
\$that

Scaling

spawn X instances of
\$that

doesn't always work

Scaling

riak_core

rafter

OTP-like for cloud ops

OTP for cloud

instances \approx processes

Scaling

solution:

spawn X instances of
\$that

... and pray

Want

topology awareness

hot migration

instant bursting

topology awareness

are two nodes sharing
one disk?

are two nodes hogging
the same CPU cache?

hot aka double buffering
migration or RCU

**instant
bursting**

```
instance_supervisor  
: create_child
```

query cache

database b-tree

VM

VFS cache

filesystem b-tree

block cache

query cache

database b-tree

VM

VFS cache

filesystem b-tree

block cache

OS
kernel

query cache

database b-tree

VM

XXX

query cache

database b-tree

VM

Erlang on Xen

A Shiba Inu dog is sitting on a bed, looking towards the left. The image has a purple overlay. In the background, there is a wooden shelf with a framed picture and some pink flowers. The text 'Erlang on Xen' is written in white, bold, sans-serif font in the upper right corner.

Erlang on Xen

WOW



Erlang on Xen

wow

lol no OS

much bare vm

very not SMP

so 30 ms boot

nice elastic

such fast



import exectime.d

example:

```
clu# dtrace -s exectime.d
```

```
~
```

```
dtrace: script 'exectime.d' matched 6 probes
```

CPU	ID	FUNCTION:NAME
1	1074	fork:create ksh pid 5112 ppid 4983
2	1035	__mac_execve:exec-success ls pid 5123 ppid 5112 creation: 1208 us
2	1119	exit1_internal:exit ls pid 5123 ppid 5112 lifetime: 4319 us

master



proger authored 11 days ago

1 parent [b13d570](#) commit [ed09232afe7bac24507b715c5af85e](#)

Showing 1 changed file with 15 additions and 0 deletions.

Show

Short ring

kiev::fprog2013

Long ring

The message "kiev::fprog2013"
within 613 seconds travelled a ring
of 7,777,001 processes running on 1001 nodes

An ephemeral Erlang cluster demo

ling

build.erlangonxen.org

xlates beam modules to `ling`
yields a kernel binary

ling

mostly R16

actively developed

issues.erlangonxen.org

ling

9p for external filesystems

lwIP for networking

GooFS for persistence

no NIFs

no net_kernel (yet)

no SSL (WIP)

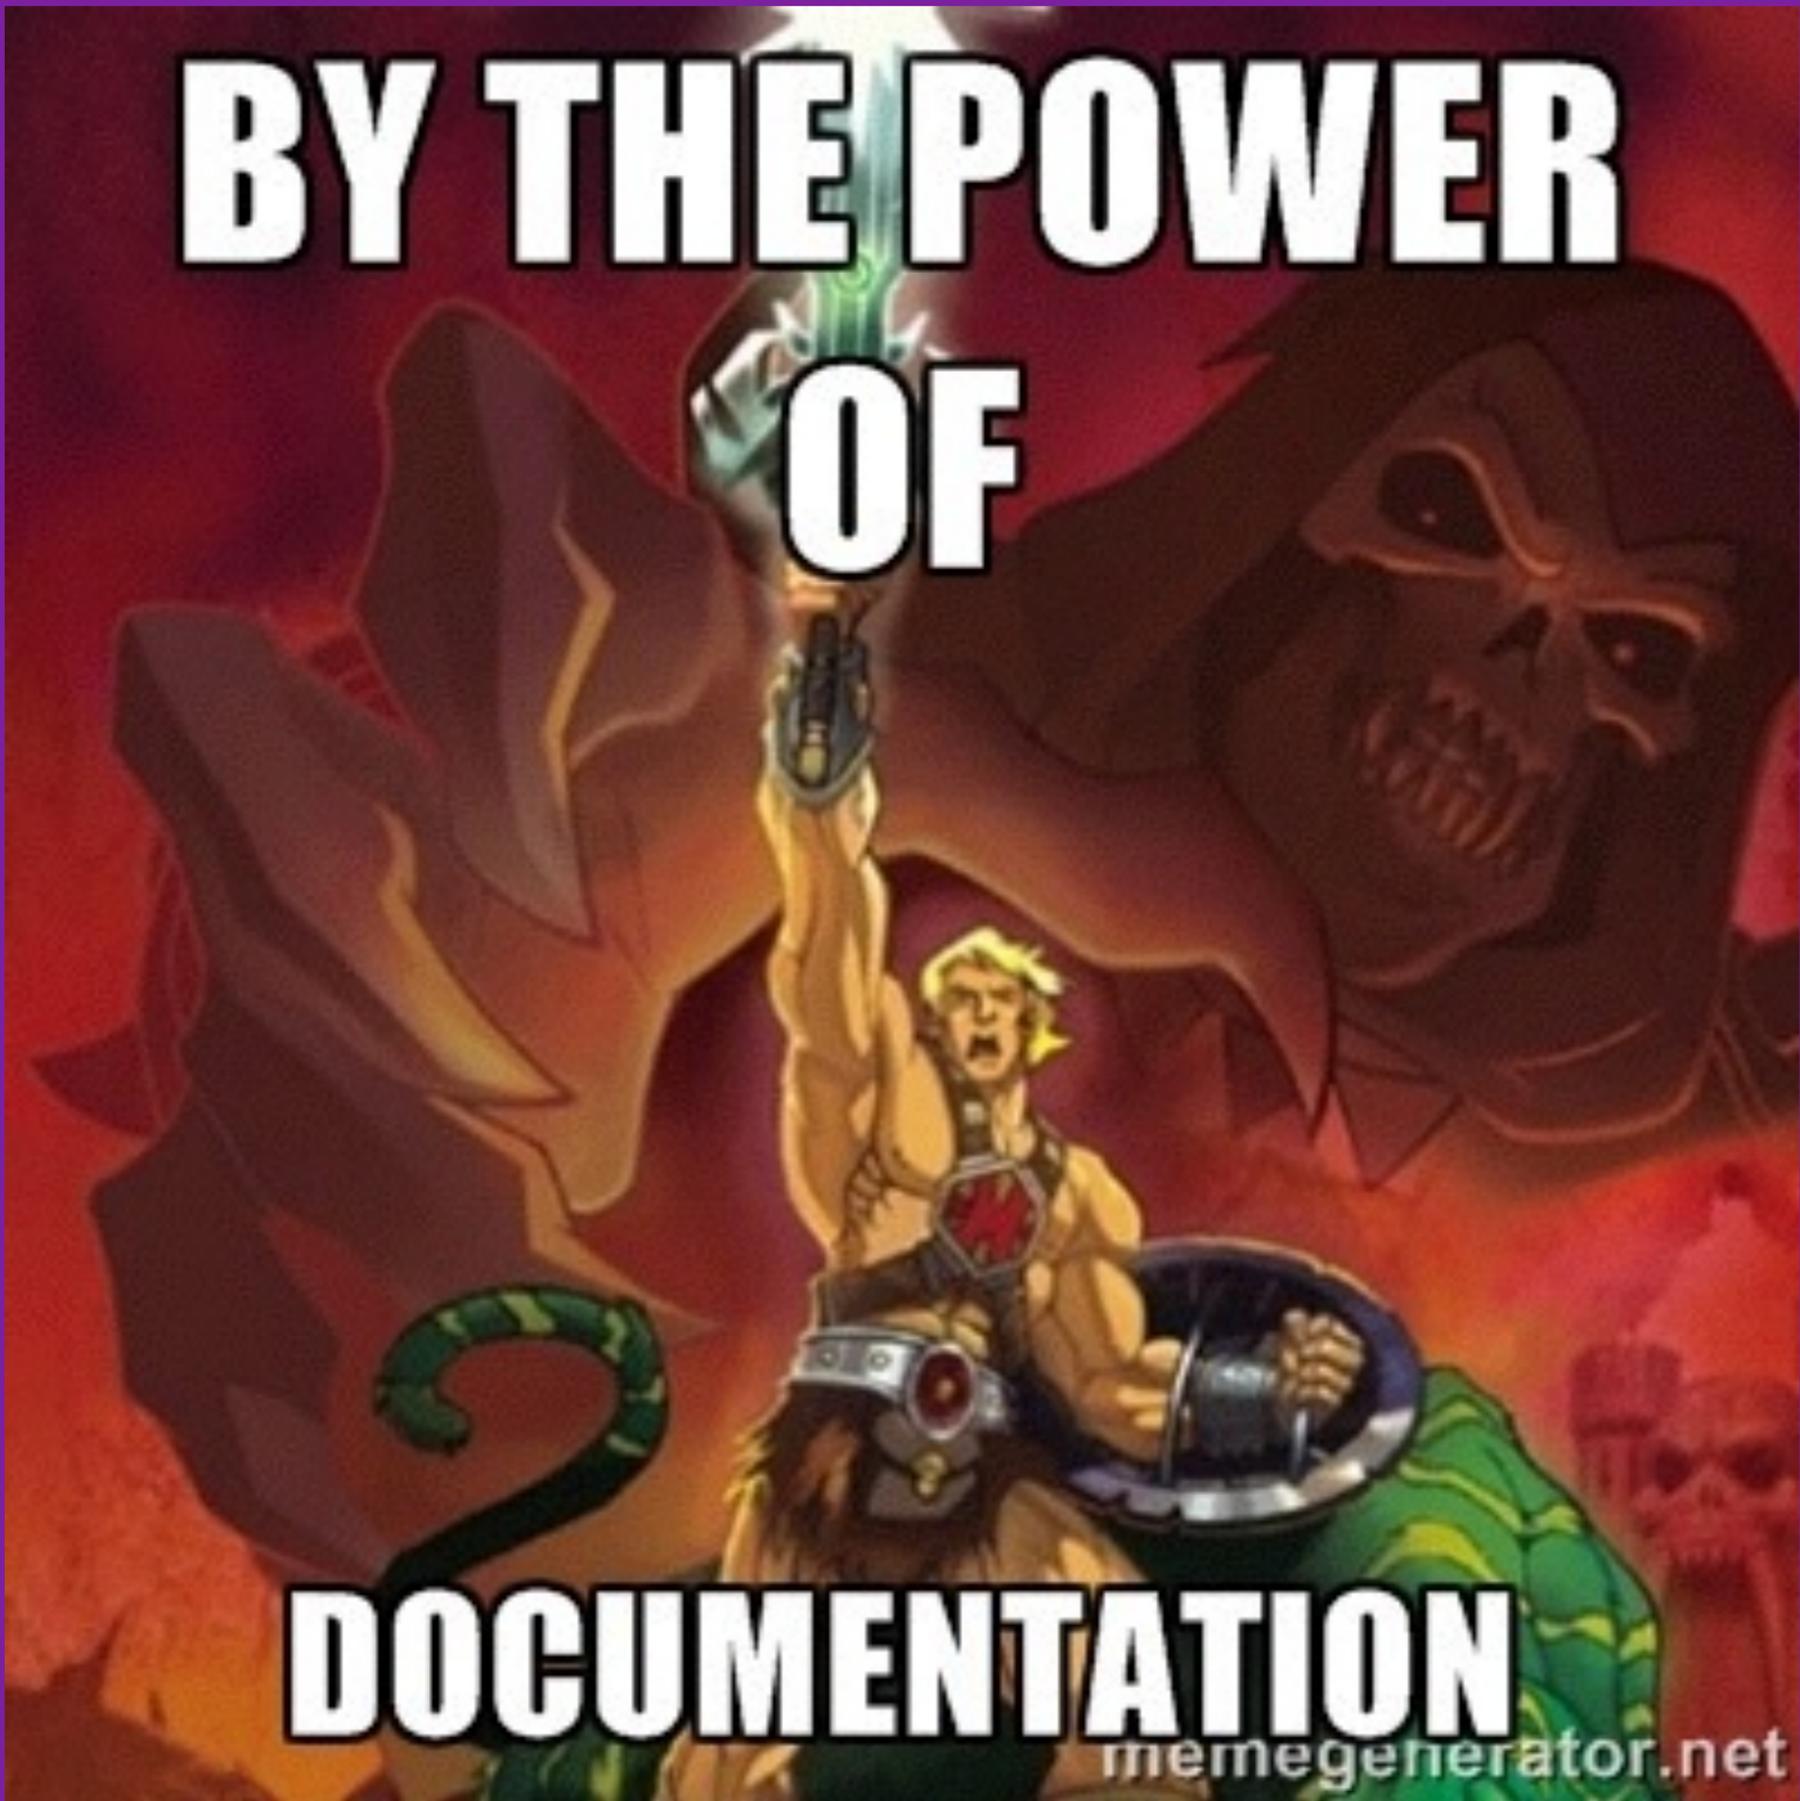
no erlang:trace

ling

{build_config, debug}

ling:trace

ling:profile



... or not

ports

Elixir (x.erlangonxen.org)

ejabberd

LINC (in progress)

riak_core / riak_pipe 1.4 *

n2o

xen 4.[1-3]

VOXOZ
& ling

xapi (too enterprise)

too much sh

libxl (fast)

gator

plugs

OpenMirage

9p ipc with unix

rump

Containers

Containers

LXC / Docker

Zones / Jails

NixOS

code vs data

Containers instant deploys

no configuration

ZFS

copy-on-write

snapshots

cheap volumes



lambda-ops





single-image deploys



app defines resources



copy-update-switch
(double buffering thing)



code is configuration



distributed by default



**instrumented by
default**

Tools



git://github.com/proger/erldocker.git build 1378375870

```
{remote_url,<<"git://github.com/proger/erldocker.git">>}{rev,<<"9caaad00730796a319fd959fcfe5058a7a615"
```

```
>>> {{2013,9,5},{10,11,10}} /usr/bin/git clone --no-checkout git://github.com/proger/erldocker.git .  
Cloning into '.'...
```

```
>>> {{2013,9,5},{10,11,16}} exit status: 0
```

```
>>> {{2013,9,5},{10,11,16}} /bin/sh -c git fetch
```

```
>>> {{2013,9,5},{10,11,17}} exit status: 0
```

```
>>> {{2013,9,5},{10,11,17}} /bin/sh -c git clean -dx
```

```
>>> {{2013,9,5},{10,11,17}} exit status: 0
```

```
>>> {{2013,9,5},{10,11,17}} /usr/bin/git checkout HEAD
```

```
>>> {{2013,9,5},{10,11,17}} exit status: 0
```

```
>>> {{2013,9,5},{10,11,17}} /bin/sh -c rebar get-deps
```

```
==> git%3A%2F%2Fgithub.com%2Fproger%2Ferldocker.git (get-deps)
```

```
Pulling hackney from {git,"git@github.com:benoitc/hackney",{tag,"0.4.4"}}
```

```
Cloning into 'hackney'...
```

```
Pulling jiffy from {git,"git@github.com:davisp/jiffy",{tag,"0.8.5"}}
```

```
Cloning into 'jiffy'...
```

```
Pulling erlsh from {git,"git@github.com:proger/erlsh","HEAD"}
```

```
Cloning into 'erlsh'...
```

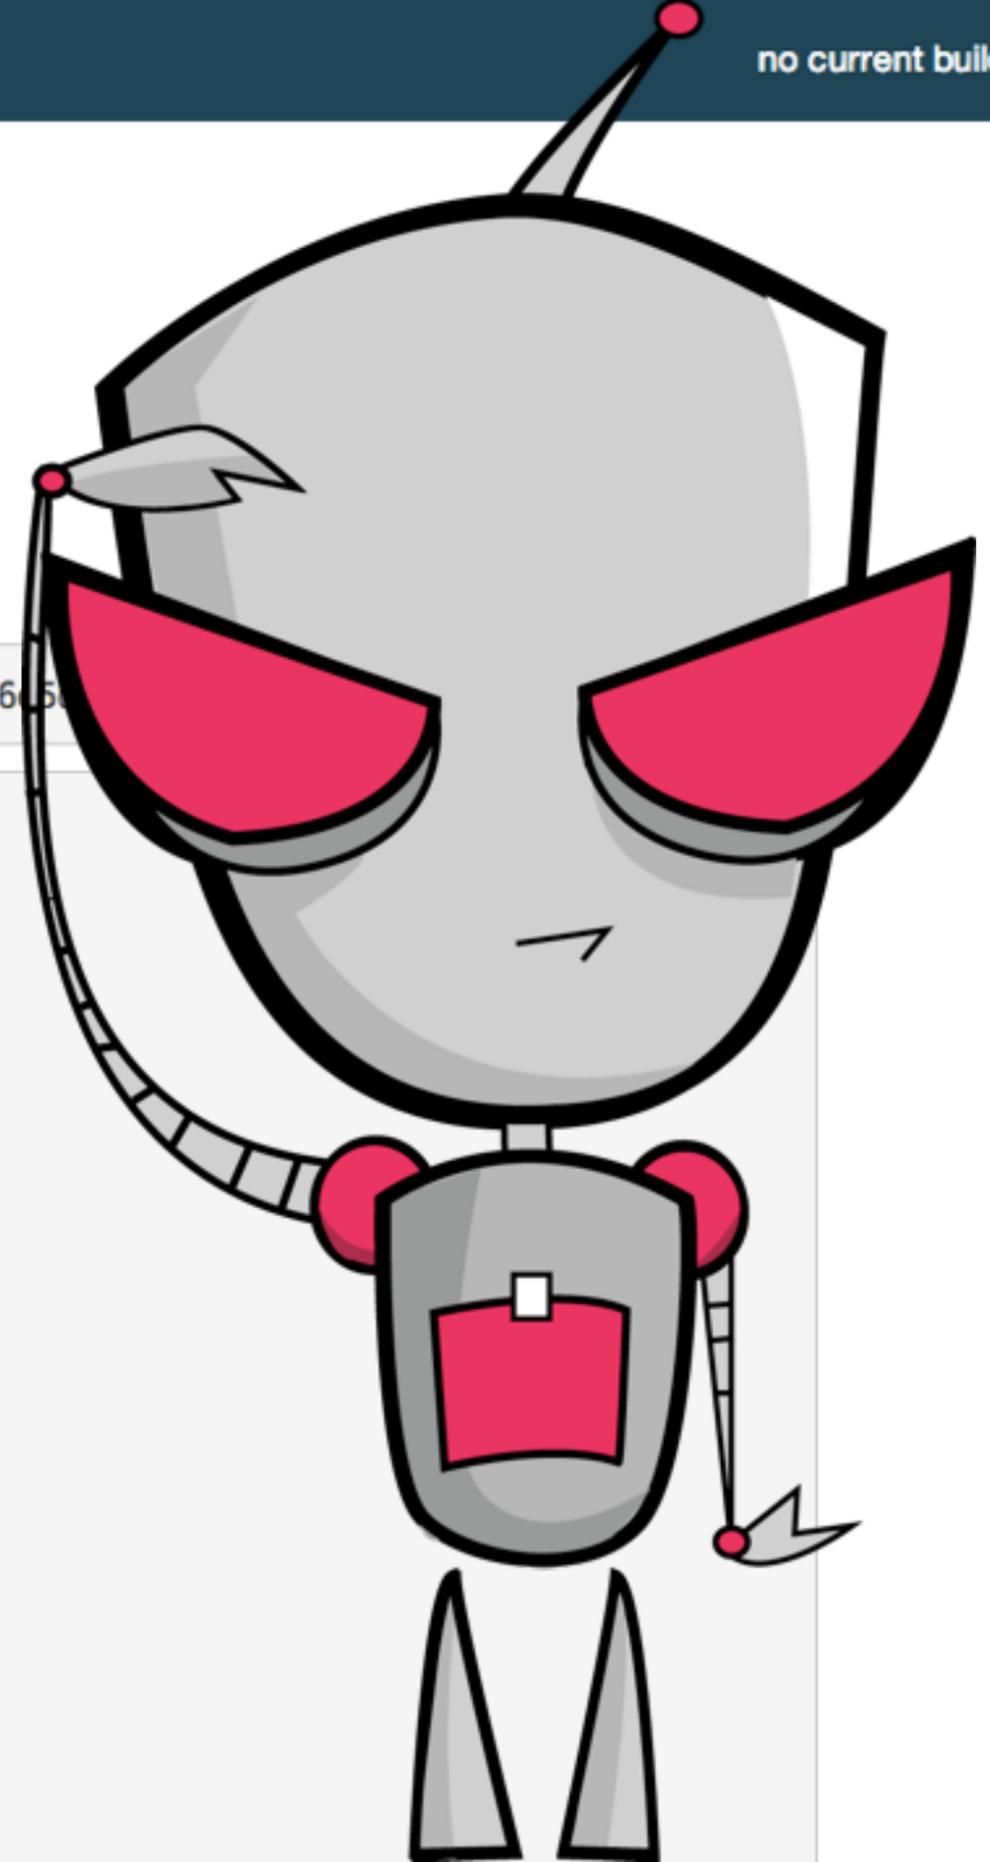
```
Pulling fn from {git,"git@github.com:reiddraper/fn",{tag,"0.3.0"}}
```

```
Cloning into 'fn'...
```

```
Pulling lager from {git,"git@github.com:basho/lager",{tag,"2.0.0"}}
```

```
Cloning into 'lager'...
```

```
Pulling active from {git,"git@github.com:proger/active","HEAD"}
```



github.com/voxoz

proger/active

proger/erlfsmon

proger/erldocker

proger/eflame

proger/erlsh

5HT/n2o

synrc/kvs

synrc/mqs

synrc/skyline

make software, not war.

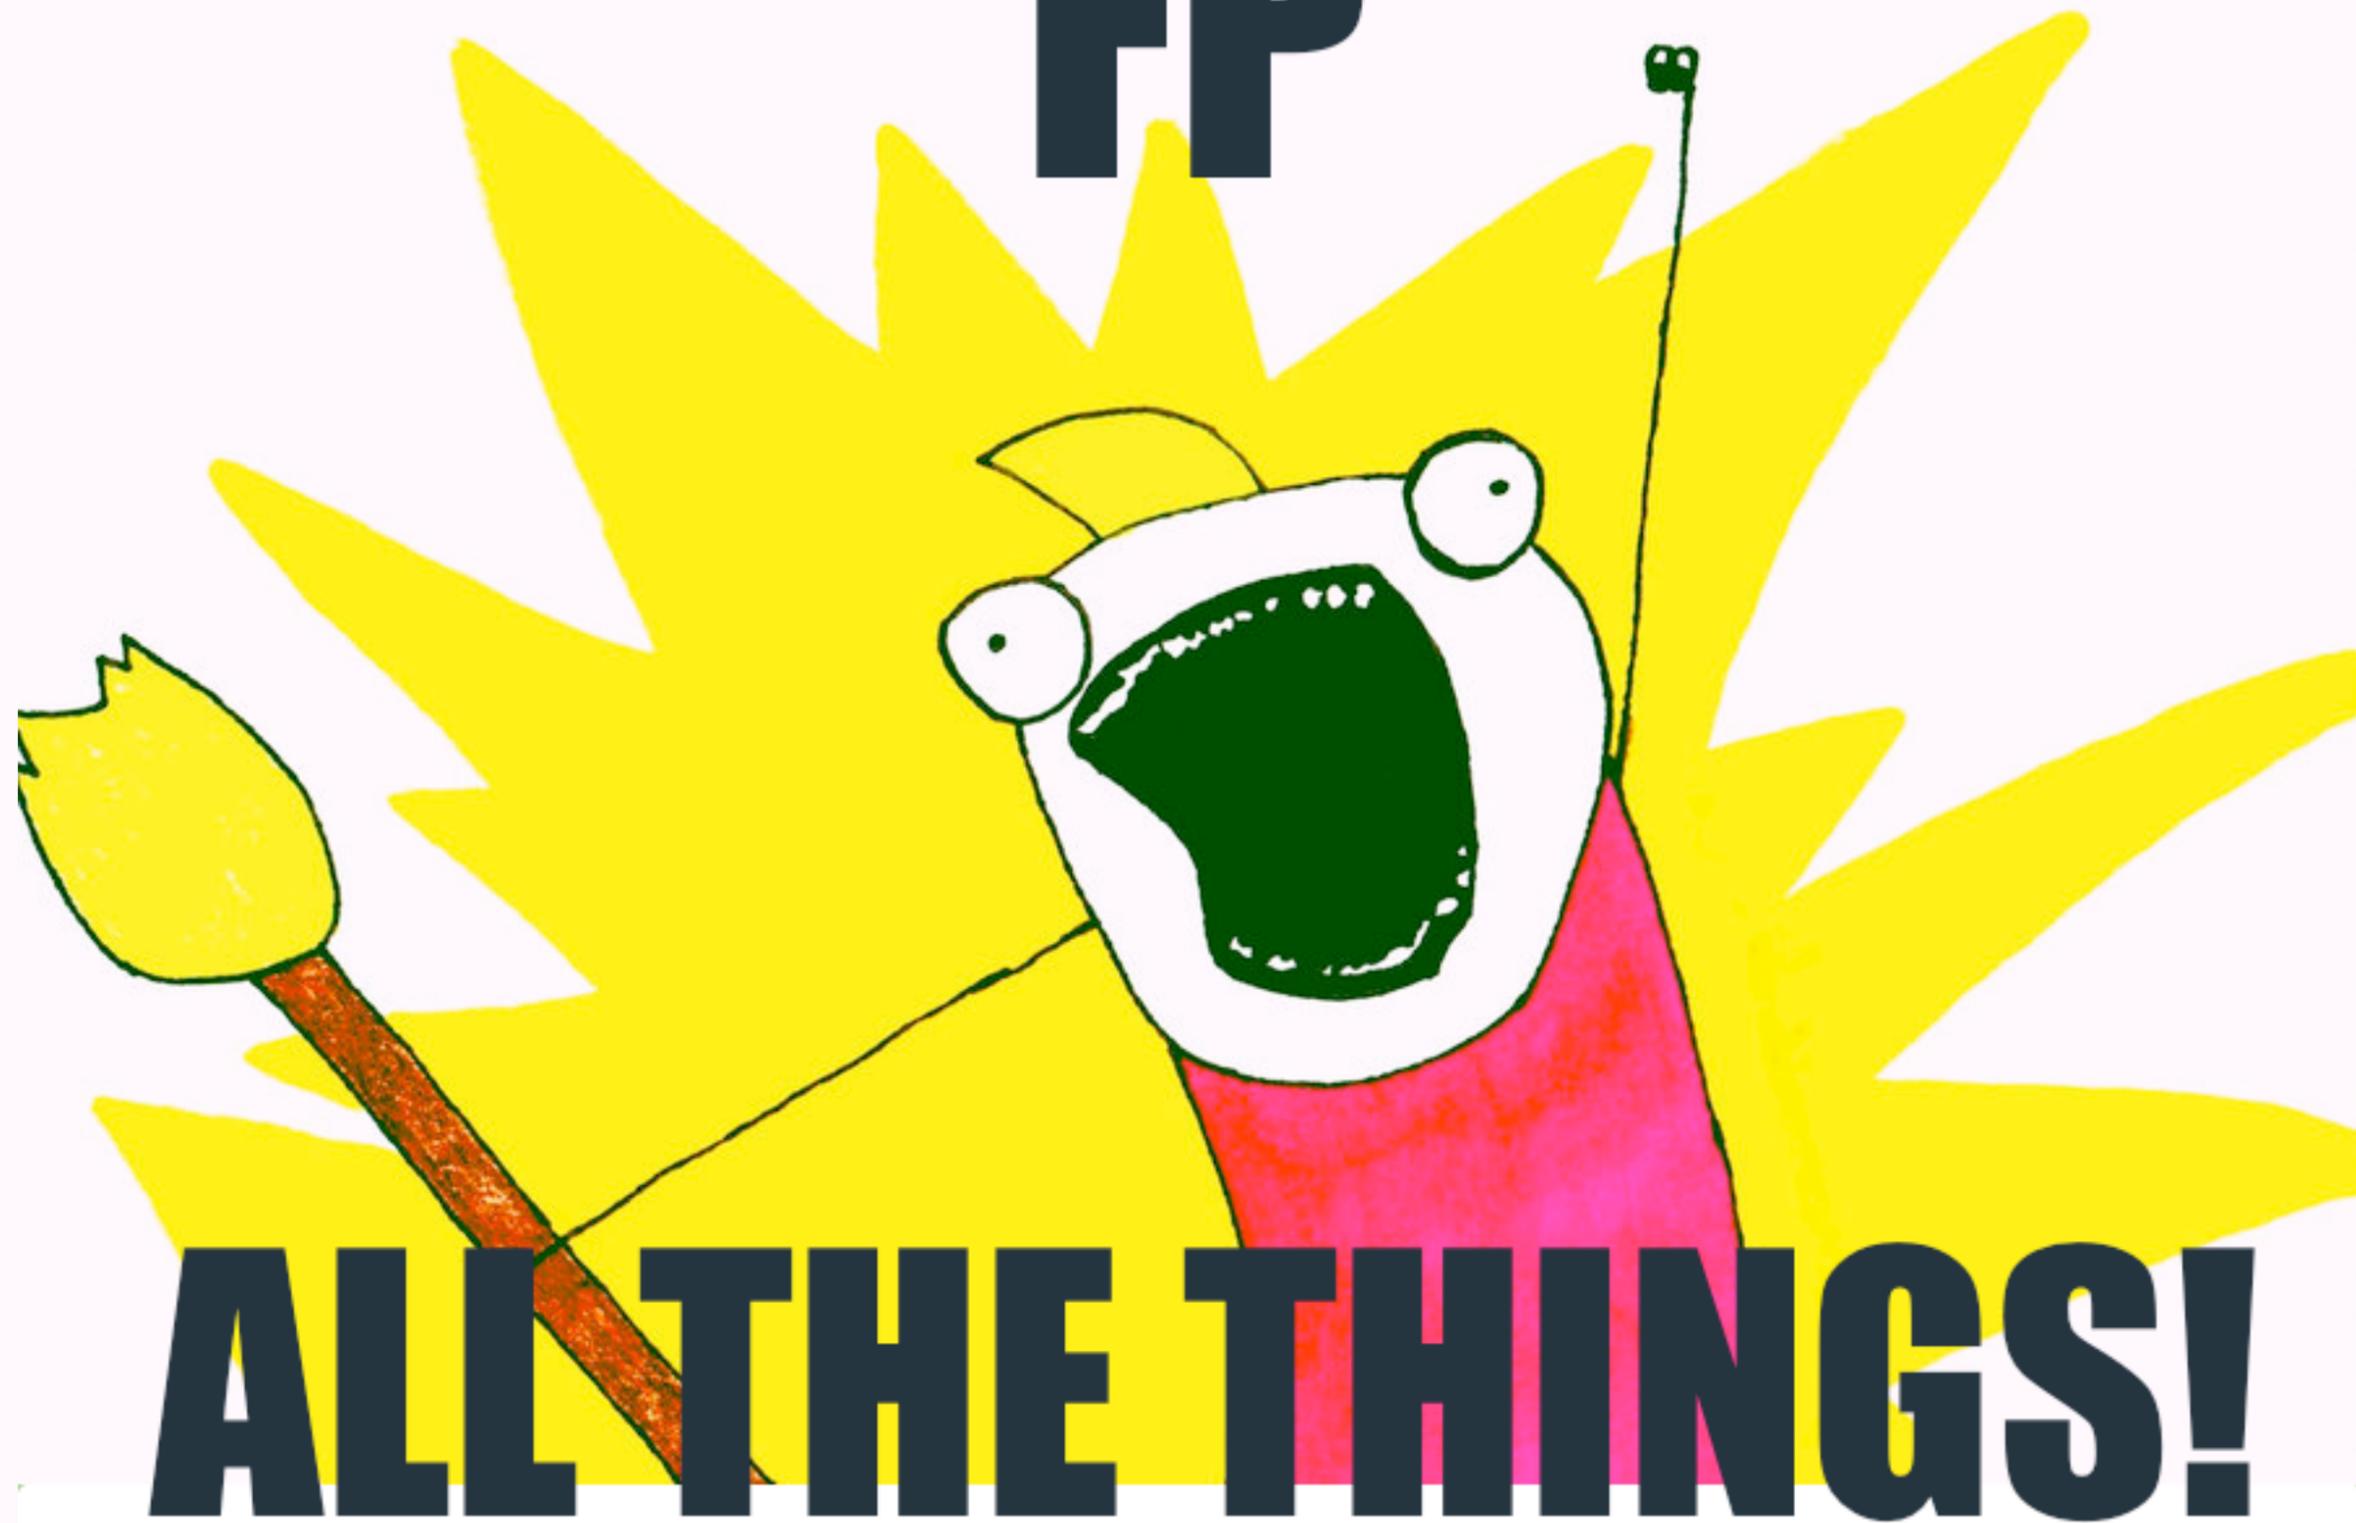
stay tuned

@darkproger

voxoz.com

kirillov.im

FP



ALL THE THINGS!

this slide is intentionally left almost blank

this slide is intentionally left almost blank

some images in this document are copyrighted by their
respective owners

[http://captainvendetta.deviantart.com/art/Gir-DUTY-
MODE-188982723](http://captainvendetta.deviantart.com/art/Gir-DUTY-MODE-188982723)