RED HAT :: CHICAGO :: 2009
SUMMIT

Infrastructure is Development

Michael Stahnke

(stahnma)

02-SEP-2009

presented by

The ideas presented today are not representative of my employer, business success, past jobs and do not offer endorsement to any particular products brands or companies. Heck, a lot of this stuff isn't even my idea to begin with.

http://www.flickr.com/photos/jjze/726682393/

http://www.flickr.com/photos/sea-turtle/3049443478/

"There are only two hard problems in Computer Science: cache invalidation and naming things."

--Phil Karlton

Who Am I?

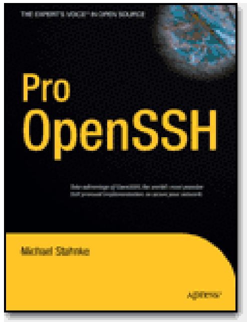
THE EXPERT'S VOICE® IN OPEN SOURCE
Pro
OpenSSH
Michael Stahnke
Apress®

Who Am I?

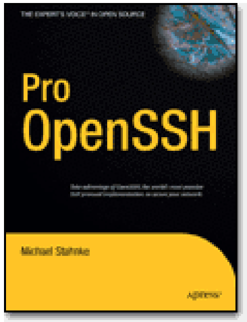

THE EXPERT'S VOICE® IN OPEN SOURCE

Pro OpenSSH

Michael Stahnke

Apress®

fedora

Who Am I?

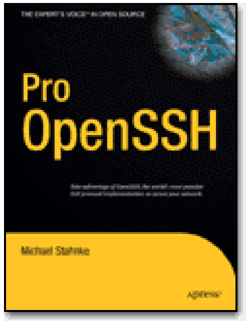

THE EXPERT'S VOICE® IN OPEN SOURCE

Pro OpenSSH

Michael Stahnke

Apress®

Who Am I?

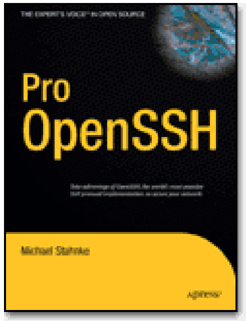

THE EXPERT'S VOICE® IN OPEN SOURCE

Pro OpenSSH

Michael Stahnke

Apress®

fedora

EPEL

OG Consulting™

Who Am I?

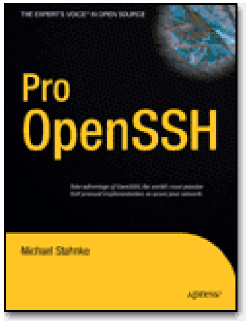
Pro OpenSSH

Michael Stahnke

Apress

<Insert Large Company Name Here>

fedora

EPEL

OG Consulting™

Where Am I?

http://stahnma.fedorapeople.org

http://www.stahnkage.com

@stahnma on Twitter

@stahnma on identi.ca

stahnma on github

stahnma@fedoraproject.org

stahnma on Freenode IRC

Where Am I?

If you google *stahnma*, it's probably me.

Baseline

Infrastructure – the collection of all components that make up the non-external facing IT realm. Primarily, in this case, servers connected to storage and network. It can include network/SAN devices in some cases.

What's Coming Up

The Tao of the Agile Infrastructure

What's Coming Up

Inherit Problems with System Administration

What's Coming Up

Some solutions to problems with System Administration

Three Domains

1. Infrastructure Practices

Three Domains

1. Infrastructure Practices
2. People

Three Domains

1. Infrastructure Practices
2. People
3. Technology Choices

The Problem

System Administration isn't a science

You can't get a degree in System Administration

Heterogeneous Systems

No Clear Job Description

Expectations not Clear

System Admins are @ holes

Can't keep up with demand

We need to stop Rock Paper Scissors
from being such an elitist sport.
--Douglas Walker

Admins fight management

http://www.flickr.com/photos/jjze/726682393/

Solutions

Solutions

Are delivered

Solutions

Through Ideas

Solutions

Put together by Good People

Solutions

Using some awesome technologies

Topic 1

Infrastructure Practices

Why is your infrastructure special?

Steal Ideas

http://www.flickr.com/photos/purplemattfish/3508761485/

Open your infrastructure

Opening your infrastructure includes your issues

Open your infrastructure and your solutions

Well, what do you mean?

To Have a High Performing Team :

You **must** know

what you manage.

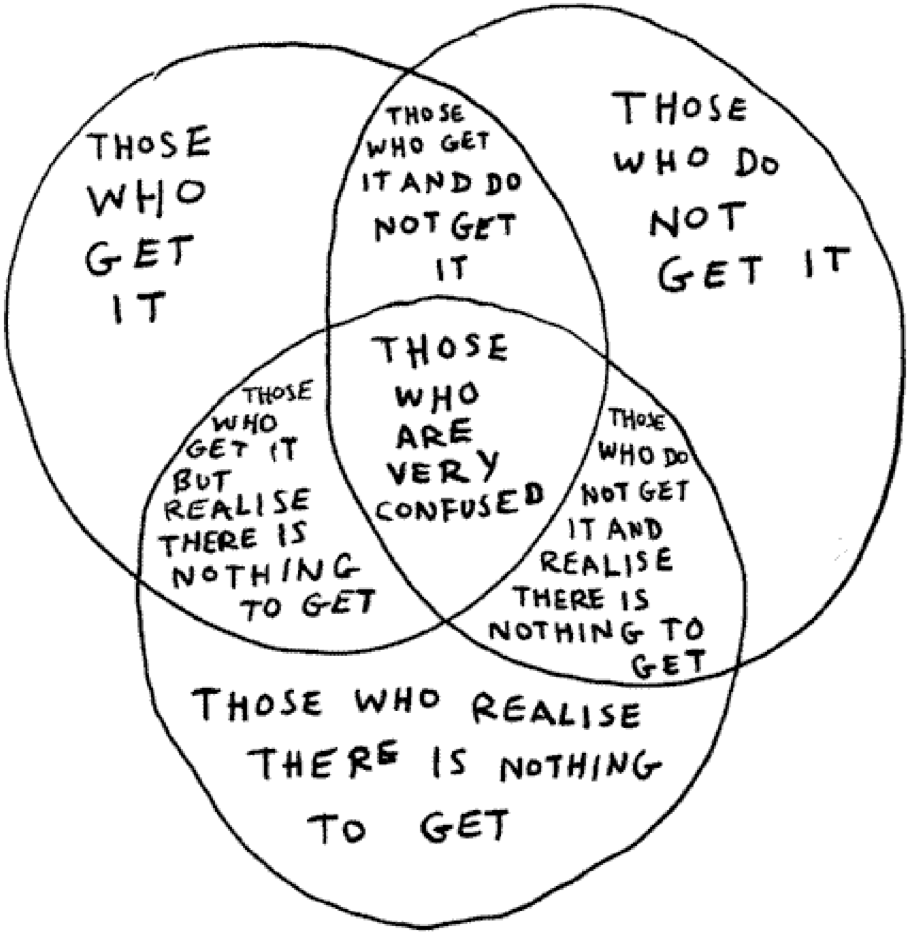
THOSE WHO GET IT
THOSE WHO GET IT AND DO NOT GET IT
THOSE WHO DO NOT GET IT
THOSE WHO ARE VERY CONFUSED
THOSE WHO GET IT BUT REALISE THERE IS NOTHING TO GET
THOSE WHO DO NOT GET IT AND REALISE THERE IS NOTHING TO GET
THOSE WHO REALISE THERE IS NOTHING TO GET

What do you have?

Asset Database

LDAP Directory

Hardware Management Tools

Power Management Tools

Monitoring Tools

Provisioning Tools

Storage Management Tools

Backup Tools

Policy Engines

Patch Tools

Security Scanning Tools

Virtualization Management Tools

Log Management Tools

Common Example:

Asset Database – Who knows?

LDAP Directory – RHDS, 389, SunOne, AD, OpenLDAP

Hardware Management Tools – IBM Director, DRAC

Power Management Tools – APC

Monitoring Tools – Nagios, Tivoli, OpenView

Provisioning Tools – Cobbler, Vmware,

Storage Management Tools – IBM Whatever, Some Custom Stuff

Backup Tools – Netbackup, Tivoli, Networker, Tar, Gzip, Rsync, Cron

Policy Engines – Cfengine, Puppet, Scripts

Patch Tools – RHN, NIM, Custom Repos

Security Scanning Tools – Lots of stuff

Virtualization Management Tools – vCenter, Virt-Manager, Spacewalk, RHN

Log Management Tools – Syslog Server

DNS – Bind

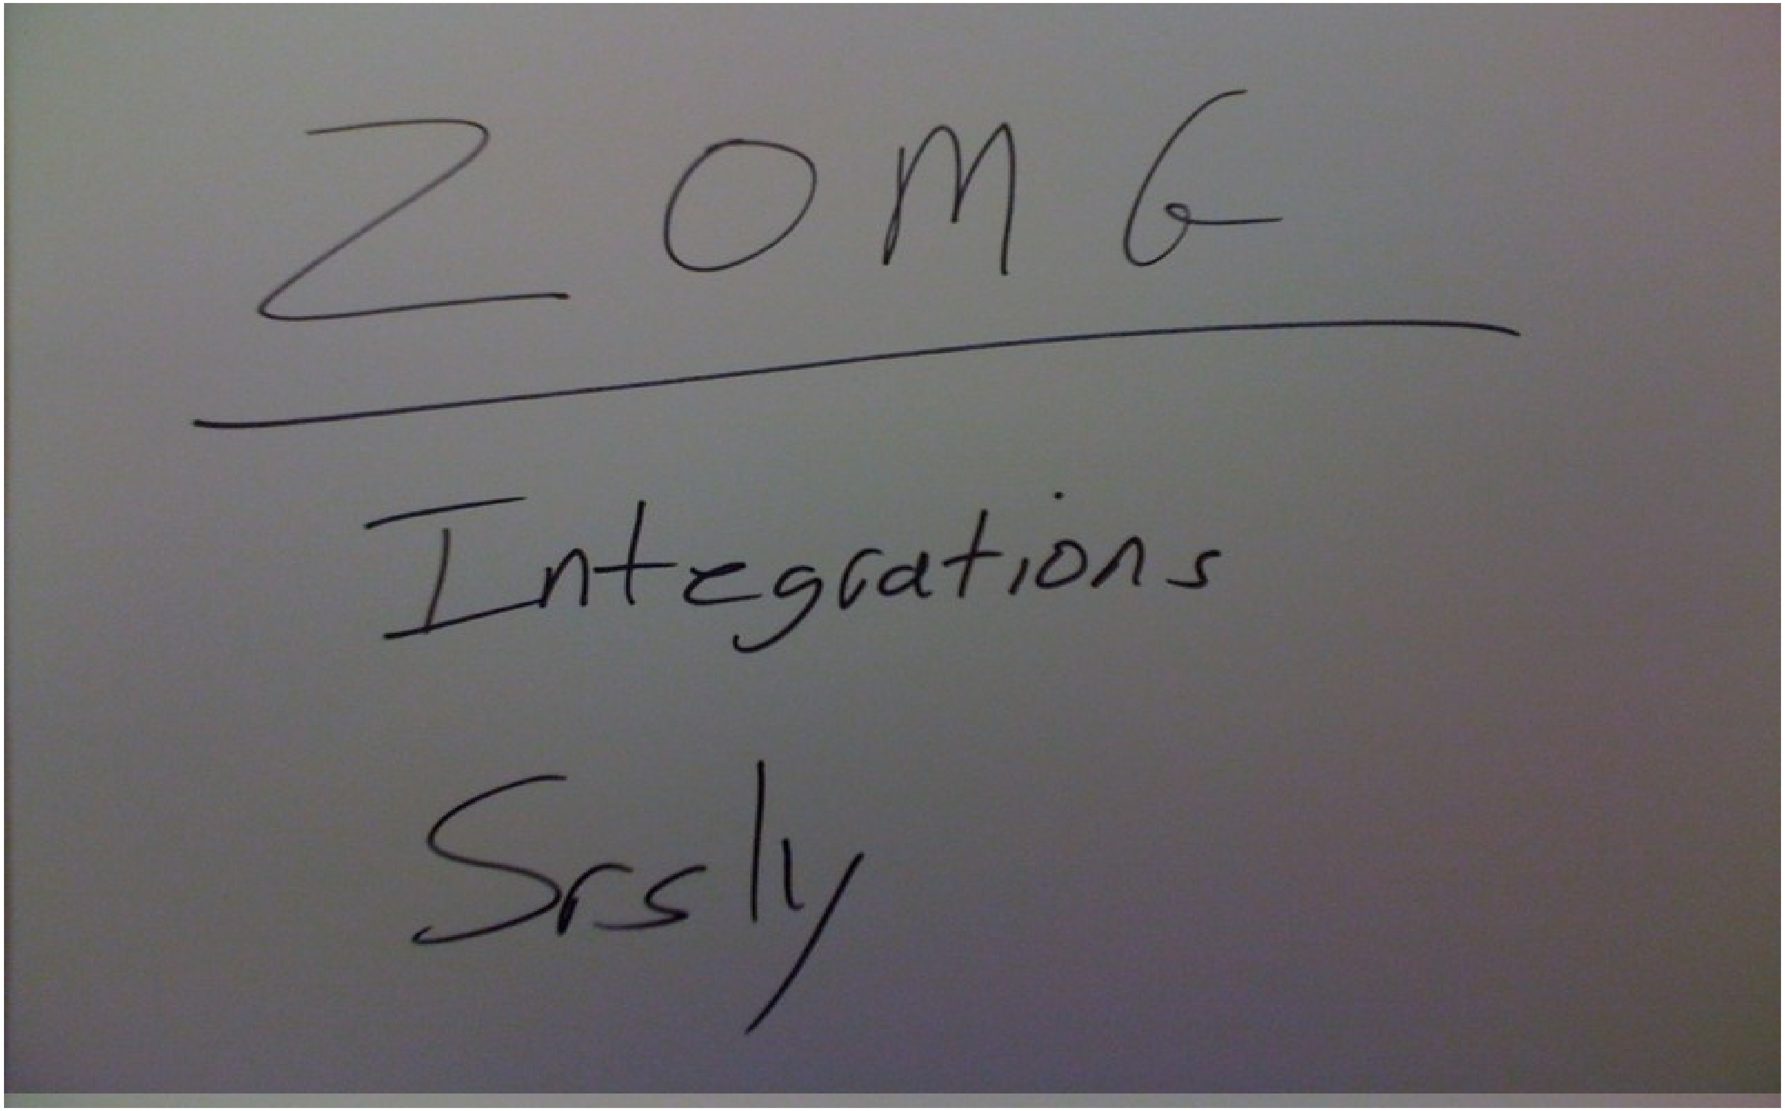

ZOMG

Integrations

Srsly

What now?

A. You can cry about it

What now?

A. You can cry about it

B. You can remove data sources

What now?

A. You can cry about it

B. You can remove data sources

C. You can integrate/federate them

What now?

A. You can cry about it

B. You can remove data sources

C. You can integrate/federate them

D. All of the Above

What now?

A. You can cry about it

B. You can remove data sources

C. You can integrate/federate them

D. All of the Above

The correct answers are both B and C; however it is very likely you will encounter A, so the likely answer is, in fact, D.

Infrastructure is Development

The Infrastructure is the Application

Application == Infrastructure

Infrastructure Goals

Deliver results to the business

Infrastructure Goals

Deliver results to the business

Make the infrastructure an investment, not a cost

Infrastructure Direction

Solving Problems

Infrastructure Direction

Solving Problems

Automation

Infrastructure Direction

Solving Problems

Automation

The right mix of people

Infrastructure Direction

Solving Problems

Automation

The right mix of people

The right decision processes

Infrastructure Direction

Solving Problems

Automation

The right mix of people

The right decision processes

Vision

Infrastructure Direction

Solving Problems

Automation

The right mix of people

The right decision processes

Vision

Known State

Where do we start?

Axiom 1

Reuse before building or purchasing

Your Infrastructure isn't a secret

Your Infrastructure isn't the secret sauce

Your Infrastructure isn't differentiating

Your Infrastructure isn't a secret

- Everybody has servers, a network, some storage
- Somebody has probably solved this problem
- Check some common places:
  - Google
  - IRC
  - Sourceforge
  - Ohloh.net
  - Amazon book selection
- What solutions do you see?

So when I say re-use???

Reuse the code and tools you have

So when I say re-use???

Enable features you are currently not utilizing

So when I say re-use???

Search for a solution built on an open platform

So when I say re-use???

Find other organizations successes on an open infrastructure

You can't reuse?

Ask why.

You can't reuse?

Is your problem new?

You can't reuse?

Is your problem special?

You can't reuse?

Ok, you can look at building or purchasing...

http://www.flickr.com:80/photos/iambrad/289462494/

I have/need a purchased proprietary solution for Problem X

I have/need a purchased proprietary solution for Problem X

So do I, and I'm sorry.

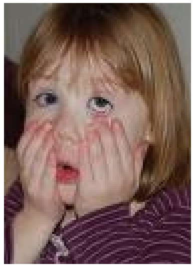

http://www.stahnkage.com/outage

Software/Tool Selection

Purchased Software Selection Criteria

There a few things to evaluate

Purchased Software Selection Criteria

There a few things to evaluate

Price

Purchased Software Selection Criteria

There a few things to evaluate

Price

Performance

Purchased Software Selection Criteria

There a few things to evaluate

Price

Performance

Functionality

Purchased Software Selection Criteria

There a few things to evaluate

Price

Performance

Functionality

But really....

Meatcloud Manifesto

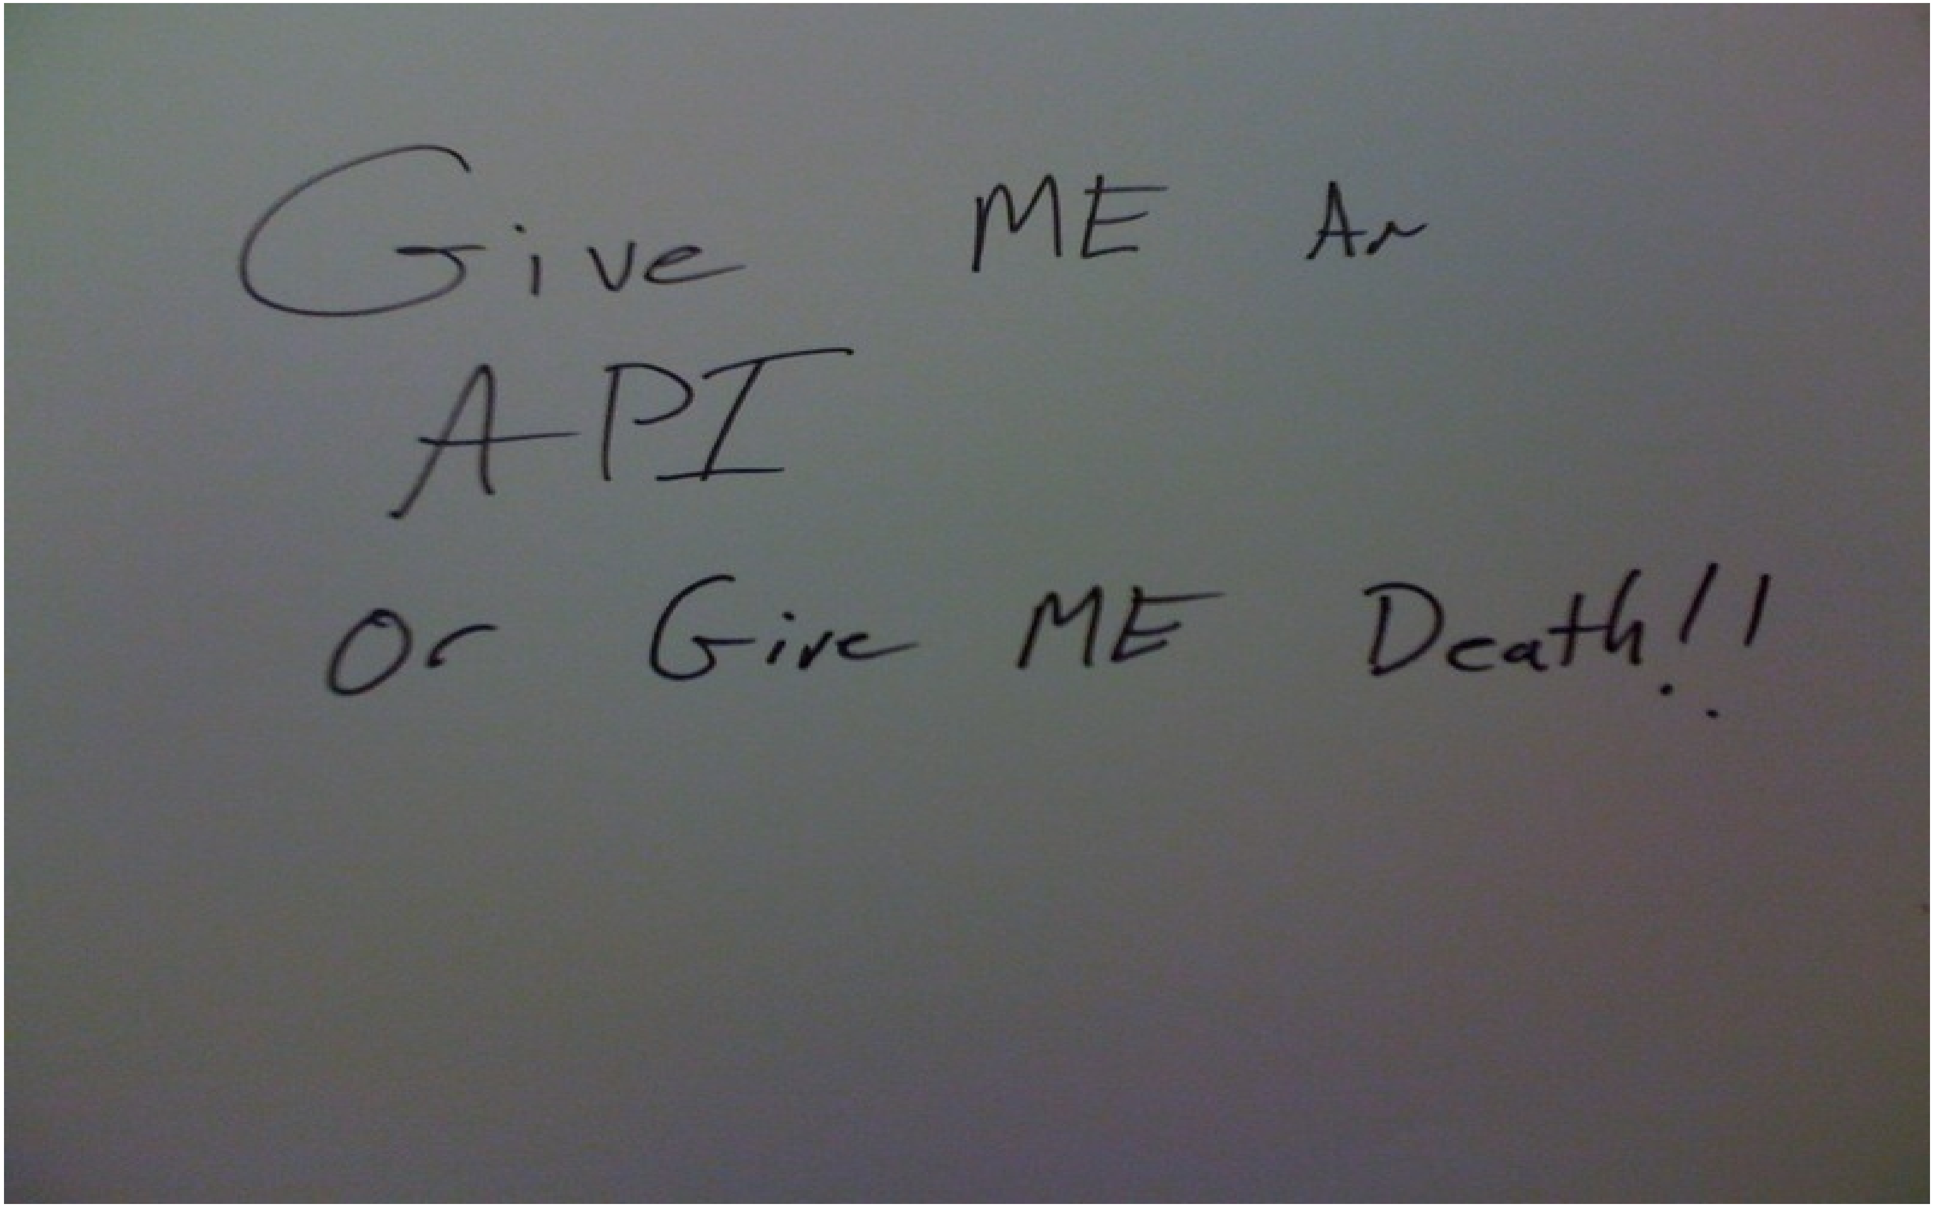

Give ME An
API
Or Give ME Death!!

Cloud Computing is all the rage

Cloud Computing is all the rage

Scaling through people is not

The GUI is for what some user interface designer thought you wanted to do. The CLI is for what you actually need to get done.

The GUI is for what some user interface designer thought you wanted to do. The CLI is for what you actually need to get done.

-- Mike Stahnke

Axiom 2

Don't Leverage the Meatcloud

Choose your technologies wisely

Very wisely

Rules for Software Evaluation

Do not implement any product that does not provide an API.

Rules for Software Evaluation

Do not implement any product that does not provide an API.

The provided API must have all functionality that the application provides.

Rules for Software Evaluation

Do not implement any product that does not provide an API.

The provided API must have all functionality that the application provides.

The provided API must be tailored to more than one language and platform.

Rules for Software Evaluation

Do not implement any product that does not provide an API.

The provided API must have all functionality that the application provides.

The provided API must be tailored to more than one language and platform.

Source code counts as an API, and may be restricted to one language or platform.

Rules for Software Evaluation

Do not implement any product that does not provide an API.

The provided API must have all functionality that the application provides.

The provided API must be tailored to more than one language and platform.

Source code counts as an API, and may be restricted to one language or platform.

The API must include functional examples and not require someone to be an expert on the product to use.

Rules for Software Evaluation

Do not use any product with configurations that are not machine parseable and machine writable

Rules for Software Evaluation

Do not use any product with configurations that are not machine parseable and machine writable

All data stored in the product must be machine readable and writable by applications other than the product itself.

Rules for Software Evaluation

Do not use any product with configurations that are not machine parseable and machine writable

All data stored in the product must be machine readable and writable by applications other than the product itself.

Writing work-arounds to cover the deficiencies in a product should be less work than writing the product’s designed functionality.

But I don't have automation specialists

That's a problem.

Do you have developers?

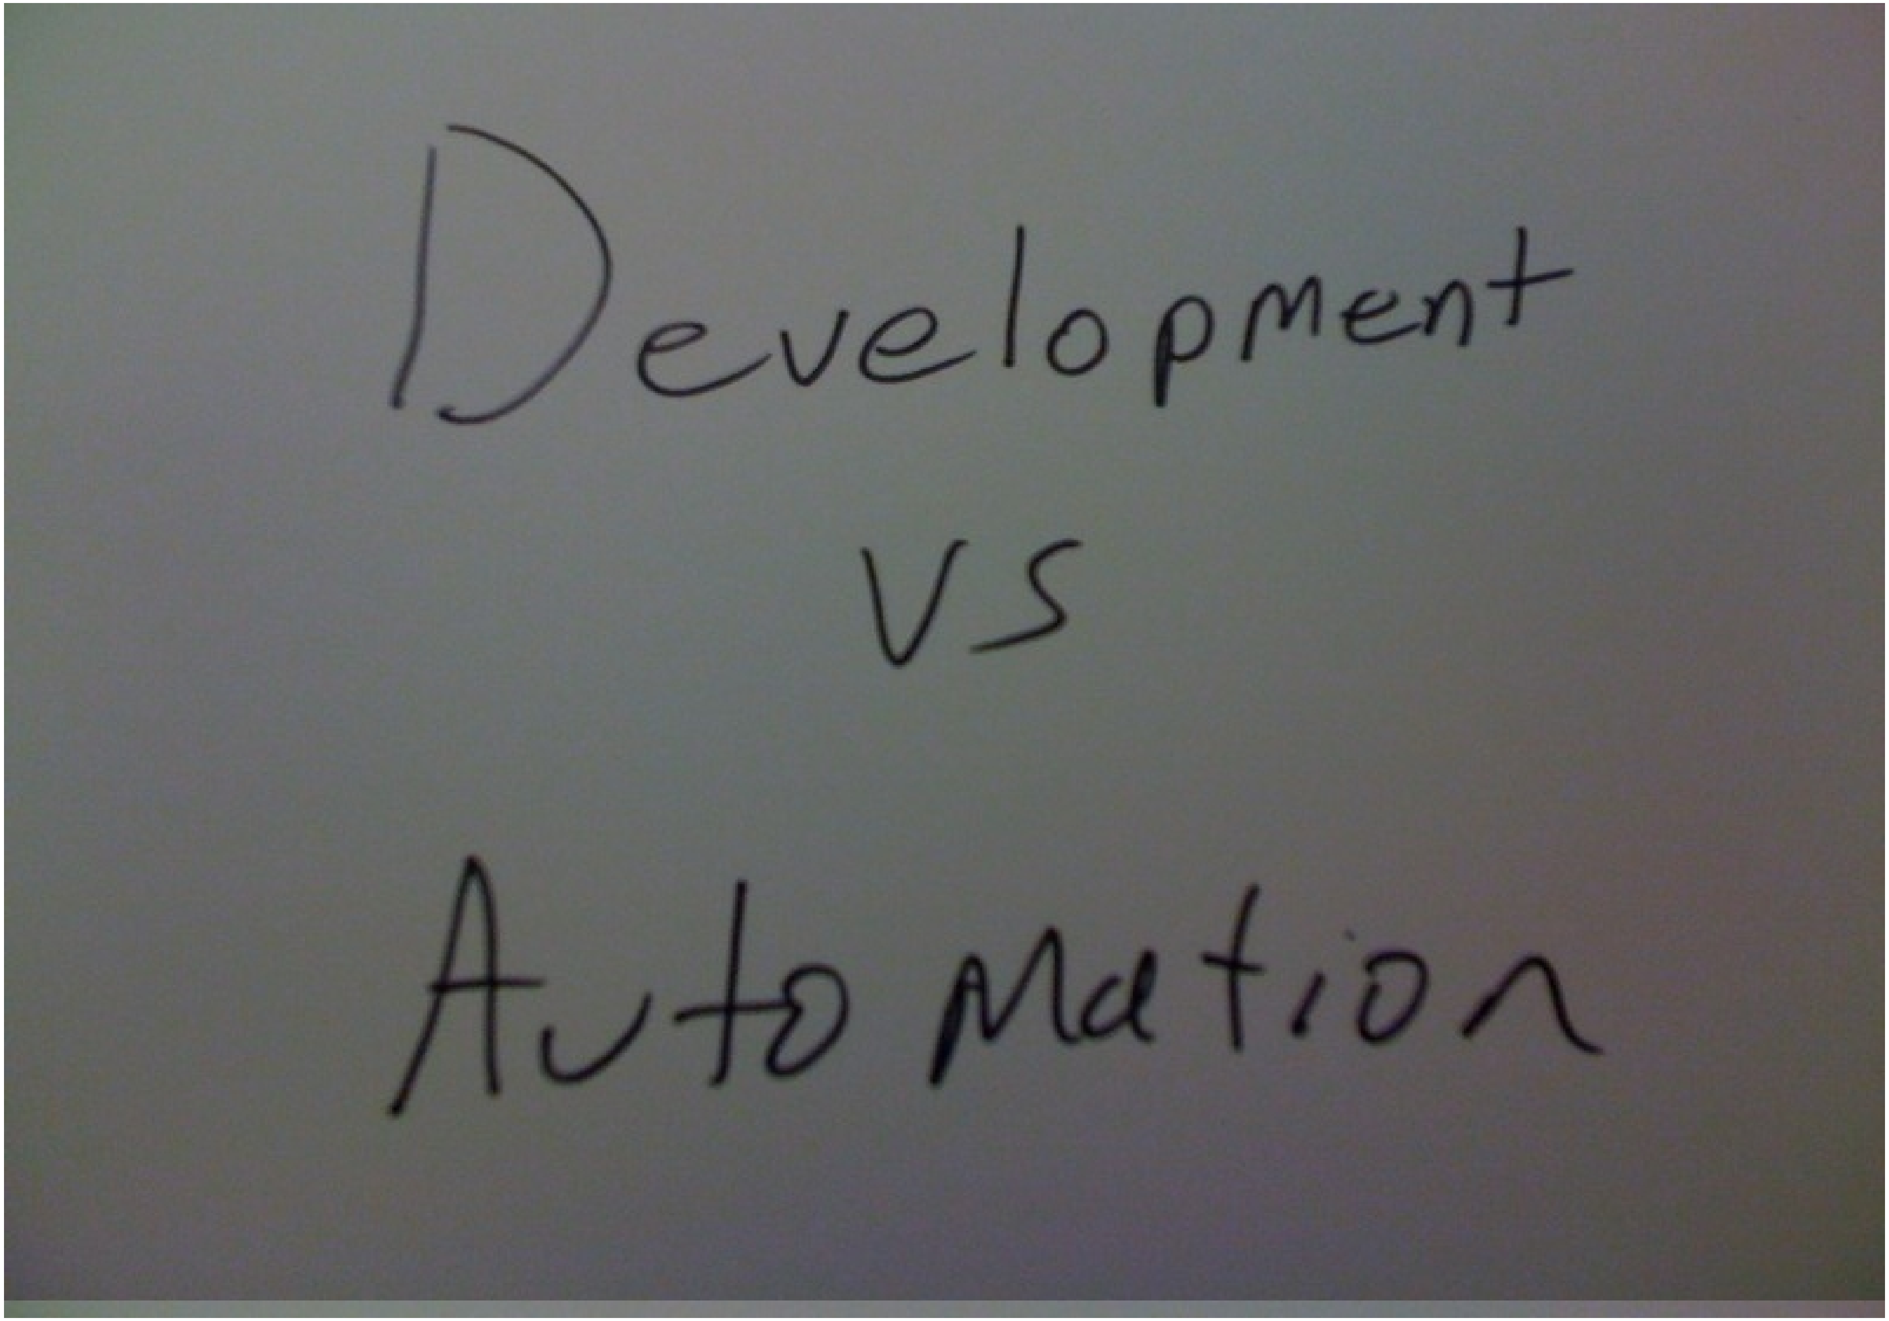

Development

vs

Automation

What's the difference between development and automation?

Process?

Requirements Gathering?

Cost?

Testing?

Egos?

Funky Lava Lamps?

Design Patterns?

Code portability?

What's the difference between development and automation?

Developers write code for use by somebody else

http://www.flickr.com:80/photos/davesag/71441707/

System Admins write code for use by themselves, and hopefully somebody else

http://www.flickr.com:80/photos/herzogbr/2274372747/

People make the Infrastructure Agile

People

Attract Talent

People

Retain Talent

People

Hire rock stars

People

Treat them like rock stars

People

Observe their organizational behavior

People

“...organizations which design systems ... are constrained to produce designs which are copies of the communication structures of these organizations.”

-- Conway's Law (1968)

People

Test their knowledge

Hiring Questions

You want rock stars remember?

Some helpful hiring questions

Describe your home network setup.

Some helpful hiring questions

What cool technology projects didn't make your resume?

Some helpful hiring questions

Do you know what version control is?

Some helpful hiring questions

Do you know what version control is?

Are you beyond CVS?

Some helpful hiring questions

Security or Availability?

Some helpful hiring questions

Apt or Yum?

Some helpful hiring questions

Provide three uses for a towel.

Some helpful hiring questions

Pirates or Ninjas?

Some helpful hiring questions

Icanhascheezburger or ihasahotdog ?

Two Types of People

Breadth vs Depth

Breadth vs Depth

Breadth

Systems Automation People.

Breadth vs Depth

Breadth

Systems Automation People.

Big Picture thinkers.

Breadth vs Depth

Breadth

Systems Automation People.

Big Picture thinkers.

Visionaries.

Breadth vs Depth

Breadth

Systems Automation People.

Big Picture thinkers.

Visionaries.

Holistic View.

Breadth vs Depth

Depth

Solve the problem for this exact situation.

Breadth vs Depth

Depth

Solve the problem for this exact situation.

Tune like it's going out of style.

Breadth vs Depth

Depth

Solve the problem for this exact situation.

Tune like it's going out of style.

Maximum ROI.

Breadth vs Depth

Depth

Solve the problem for this exact situation.

Tune like it's going out of style.

Maximum ROI.

Deep understanding of the technology.

Breadth vs Depth

Turns out, you need both.

People

Specialist vs Generalist

People

Specialist vs Generalist.^H^H^H^H^...not so much.

How many people is the right amount?

50:1?

100:1?

200:1?

2:1?

Gartner Says:

Gartner Says: “Yes”

There is one rule for Server/Admin Ratio:

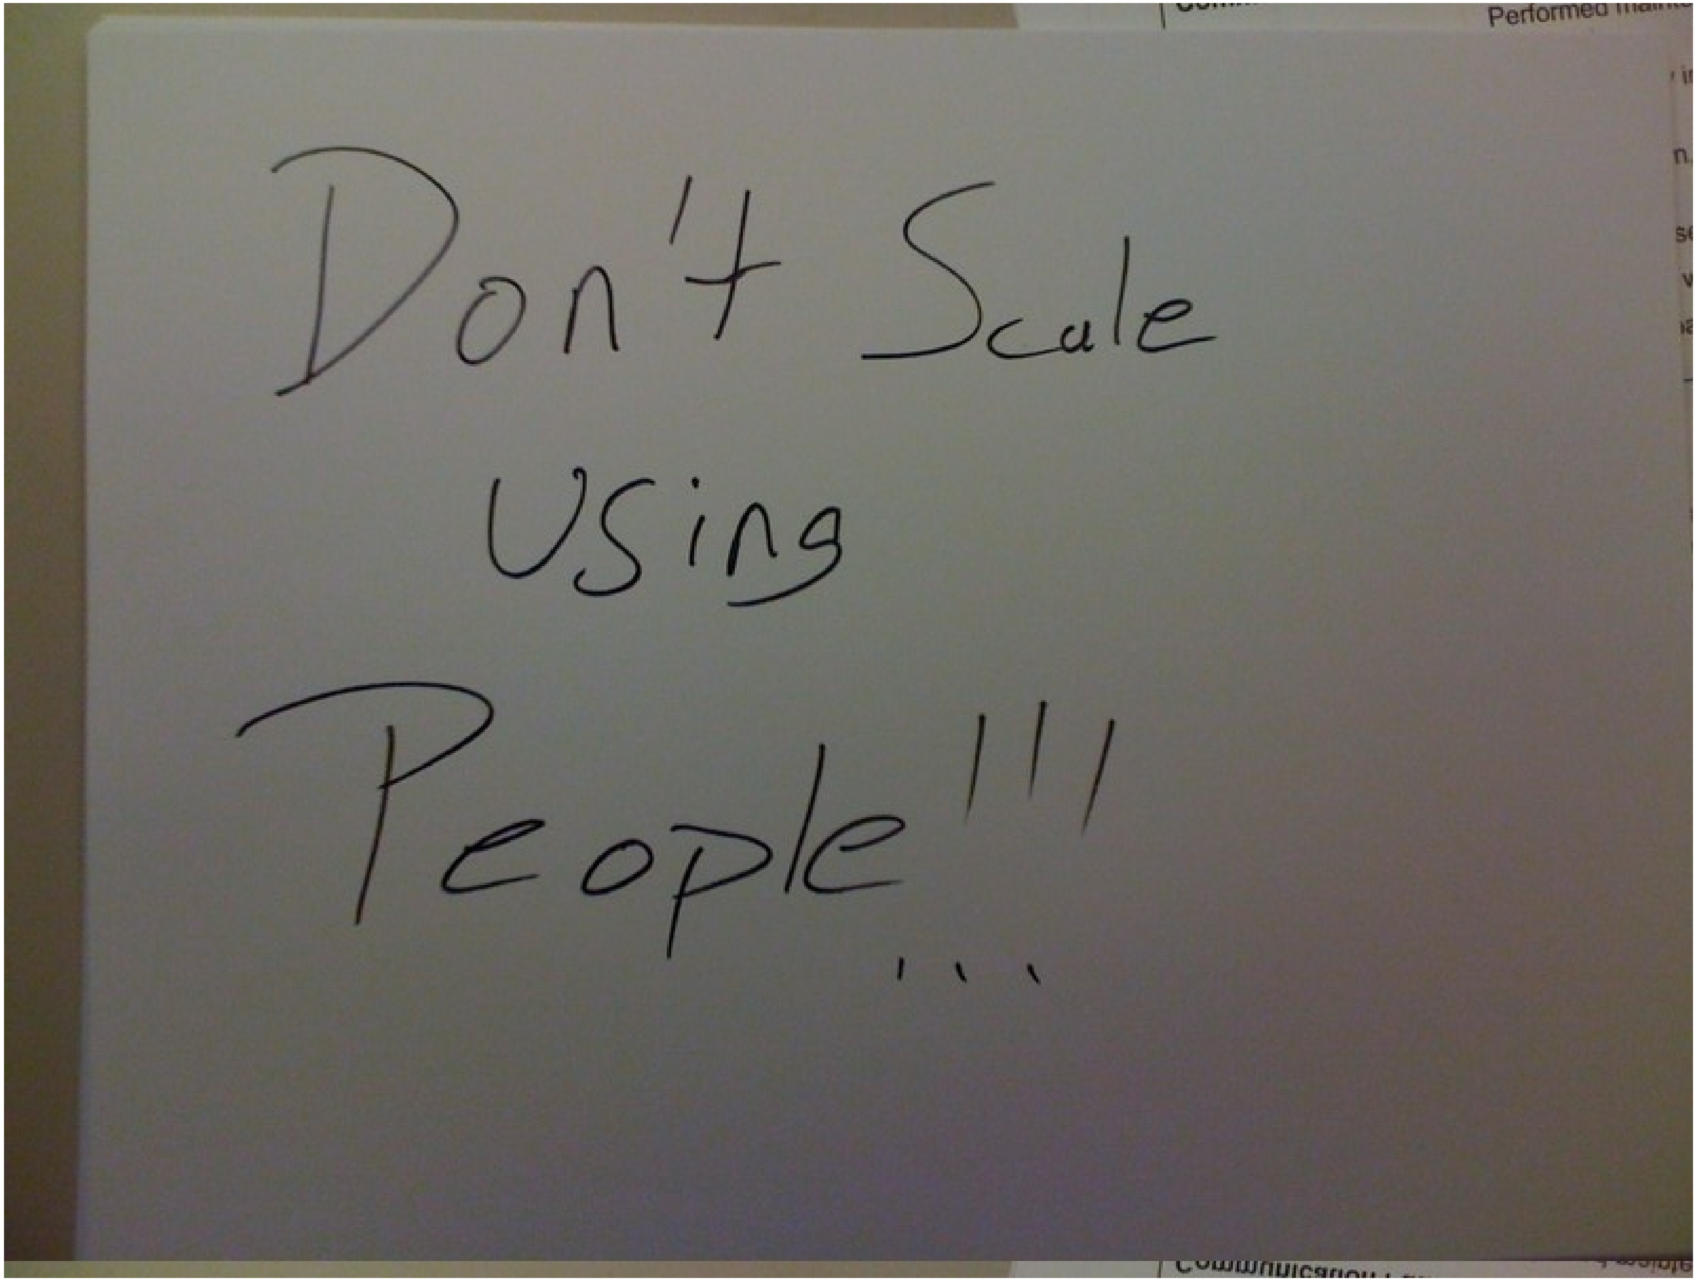

Don't Scale
Using
People!!!
...

Axiom 2

Don't Leverage the Meatcloud

We've Covered Technology Selection

We've Covered People Decisions

Let's cover Community

Axiom 3

Decouple your Infrastructure.

Code for the generic case

Stop reinventing the wheel, srsly.

Moving Forward

Can you implement an Open Strategy to solve this issue?

Moving Forward

If you have to solve it yourself, can it be Open for others?

Moving Forward

Have an Open Infrastructure

Moving Forward

Driving on hexagonal wheels isn't fun

http://www.flickr.com:80/photos/andrewmbutler/2762480367/

Moving Forward

Driving on hexagonal wheels isn't fun

Ask this guy.

http://www.flickr.com:80/photos/andrewmbutler/2762480367/

Moving Forward

Driving on hexagonal wheels isn't fun

Quit reinventing the wheel...poorly

Steal it all

Practical Examples of Open Infrastructure

Steal it all

Practical Examples

The Fedora Infrastructure Project

https://fedoraproject.org/wiki/Infrastructure

Steal it all

Practical Examples

The Community Services Infrastructure Standards

http://infrastructure.fedoraproject.org/csi/free-software-policy/en-US/html-single/

https://fedorahosted.org/csi/

De-coupled Infrastructure Benefits

Commonly accepted solutions to problem

http://www.flickr.com:80/photos/thomas-merton/255204957

De-coupled Infrastructure Benefits

Commonly accepted solutions to problem

Portability of solutions

http://www.flickr.com:80/photos/thomas-merton/255204957

De-coupled Infrastructure Benefits

Commonly accepted solutions to problem

Portability of solutions

Ability to hire knowledgeable individuals

http://www.flickr.com:80/photos/thomas-merton/255204957

De-coupled Infrastructure Benefits

Commonly accepted solutions to problem

Portability of solutions

Ability to hire knowledgeable individuals

Able to retain rock stars

De-coupled Infrastructure Benefits

Commonly accepted solutions to problem

Portability of solutions

Ability to hire knowledgeable individuals

Able to retain rock stars

Community built around tools and support

http://www.flickr.com:80/photos/thomas-merton/255204957

3 Axioms

1. Reuse before building or purchasing
2. Don't leverage the meatcloud
3. Decouple your infrastructure

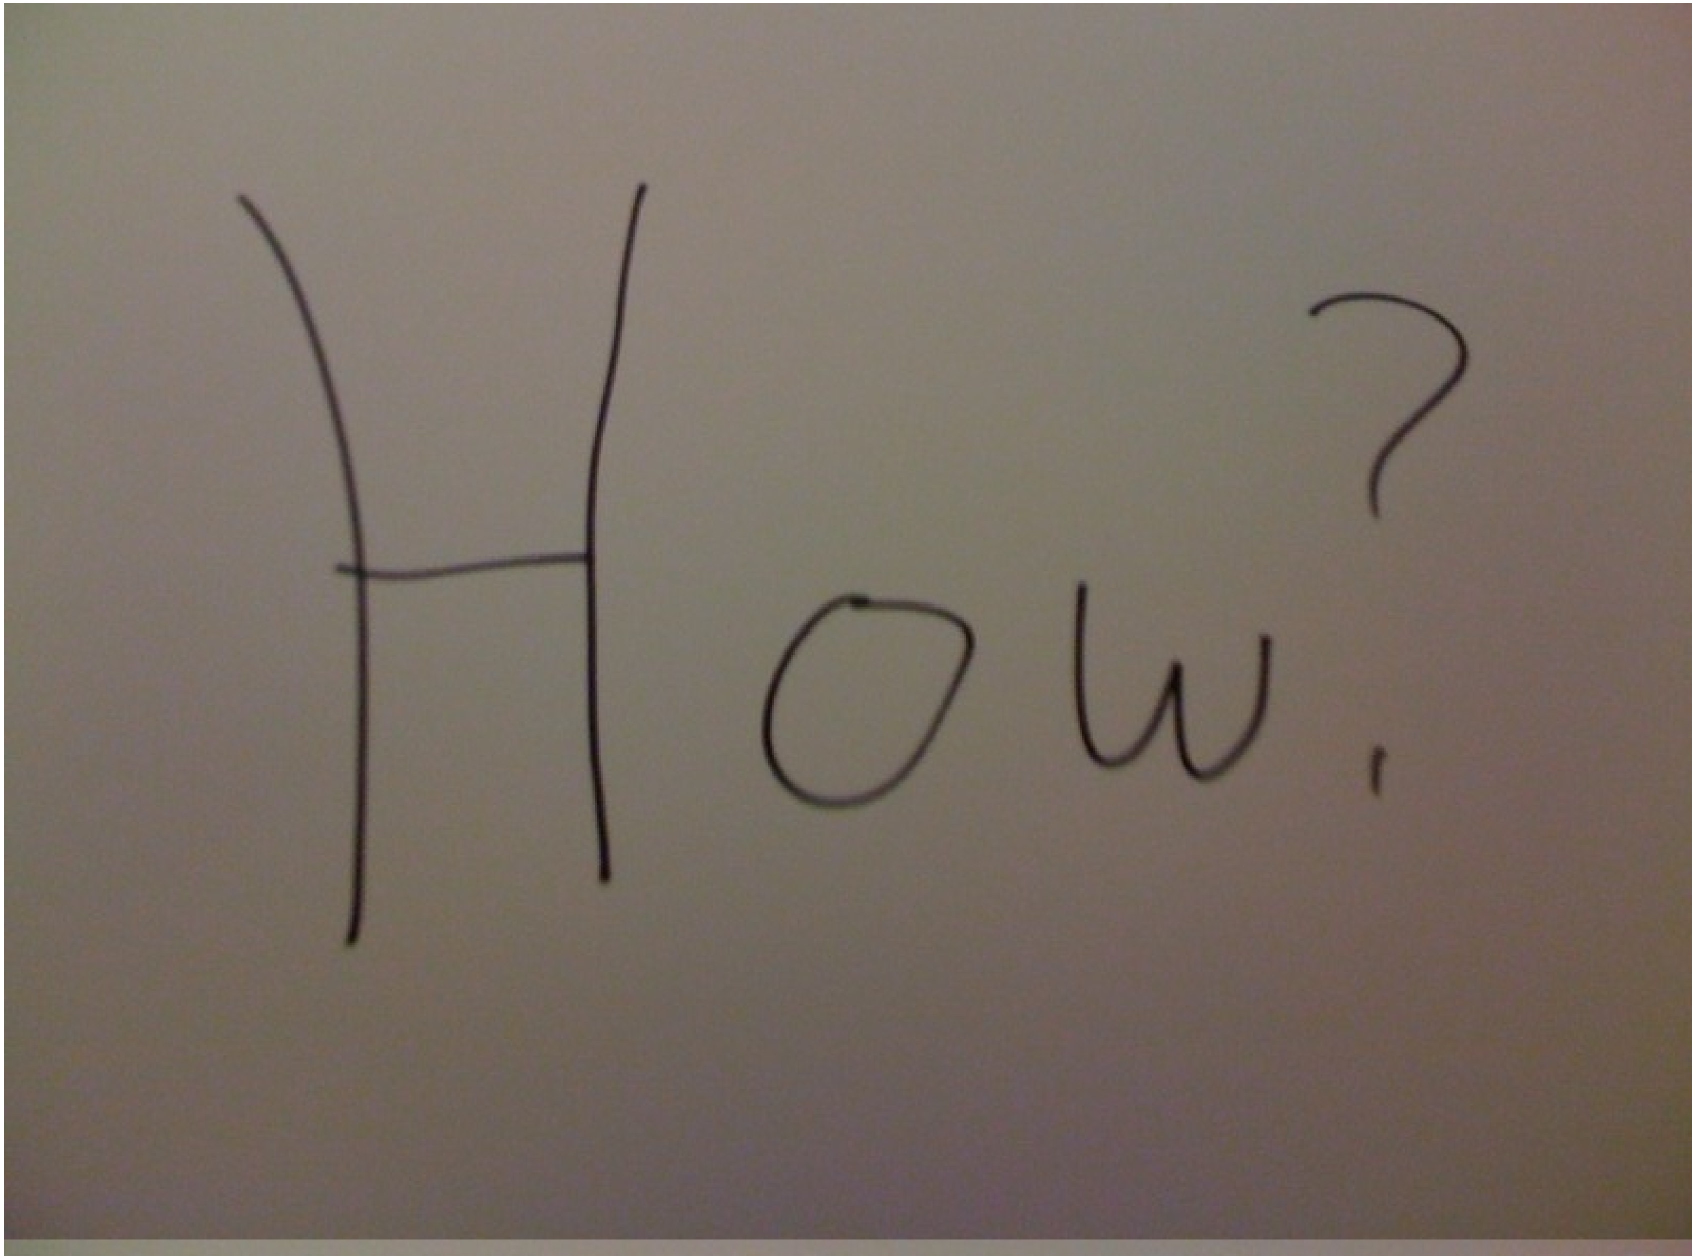

How?

Your Time is Valuable

Waste less of it

Your team's time is valuable

Manage it

Your Time is Valuable

Identify tasks on which the team spends the most time.

Commonly

- (Growth) Deployment
- Account Management
- Ad-Hoc File Transfer type activity
- Patches

http://www.flickr.com:80/photos/denaldo/3326738675/

Time Evaluation

- Evaluate your team's time spent on “**Displacement Activities**”
  - Fund Raising
  - Parties/Showers
  - Volunteer Stuff
  - Charity
- This is normally done because the person can't actually meet the requirements of $DAYJOB

See Also: People

Pick a Task

Focus on one thing

Focus on one thing only

Seriously, this will help

You're ready to automate.

OK, maybe you're not ready to automate.

Don't automate bad process.

http://www.flickr.com:80/photos/houseofcards/94935538/in/photostream/

You get bad results.....faster.

Bad processes deliver bad results.

You don't need a panflute.

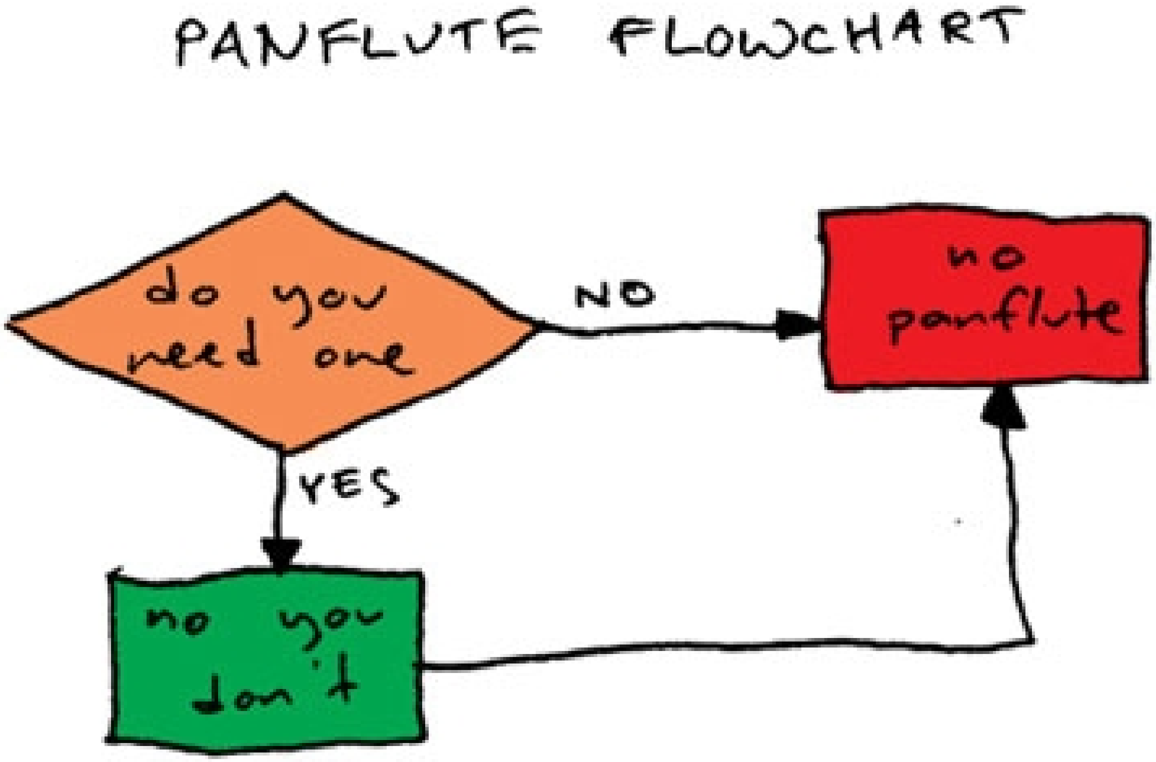
PANFLUTE FLOWCHART
do you need one
NO
no panflute
YES
no you don't

It starts with requirements

Pick a Task

- Let's start with system Deployment
  - You have some great options for tools
  - Do you need a tool?
  - Do you need it automated?

Pick a Task

- Let's start with system Deployment
  - You have some great options for tools
  - Do you need a tool?
  - Do you need it automated?
  - Do you need a pan-flute?

Process Evaluation

Who consumes the process?

Process Evaluation

What are the outputs?

Process Evaluation

What are the inputs?

Process Evaluation

What's unknown?

Process Evaluation

What's that rare case that you should just leave out?

Process Evaluation

What's the gotcha you didn't put in because it was hard?

Gather Requirements

Ask your process consumers about their needs

Gather Requirements

Ask them to think about outputs, not the process

Gather Requirements

How do they provide the information?

Requirements

Formal Specification Document?

Requirements

Formal Specification Document?

Awesome.

Requirements

Formal Specification Document?

- Awesome.
    - I doubt it

Requirements

Hundreds of tickets, IM conversations, phone calls, hallway conversations and just plain old complaining?

Requirements

Hundreds of tickets, IM conversations, phone calls, hallway conversations and just plain old complaining?

Yeah, I thought so.

Requirements

Let's Get Agile.

Requirements

Let's Get Agile. Development is agile.

Requirements

Let's Get Agile. ~~Development~~ Infrastructure is agile.

Requirements

Stories

Requirements

Use Cases

Requirements

Test Cases

Requirements

Stores, Use Cases, Test Case --> Track them

Requirements

Use a tracker,

Requirements

Use a tracker, use a wiki,

Requirements

Use a tracker, use a wiki, use a notebook,

Requirements

Use a tracker, use a wiki, use a notebook, a whiteboard

Requirements

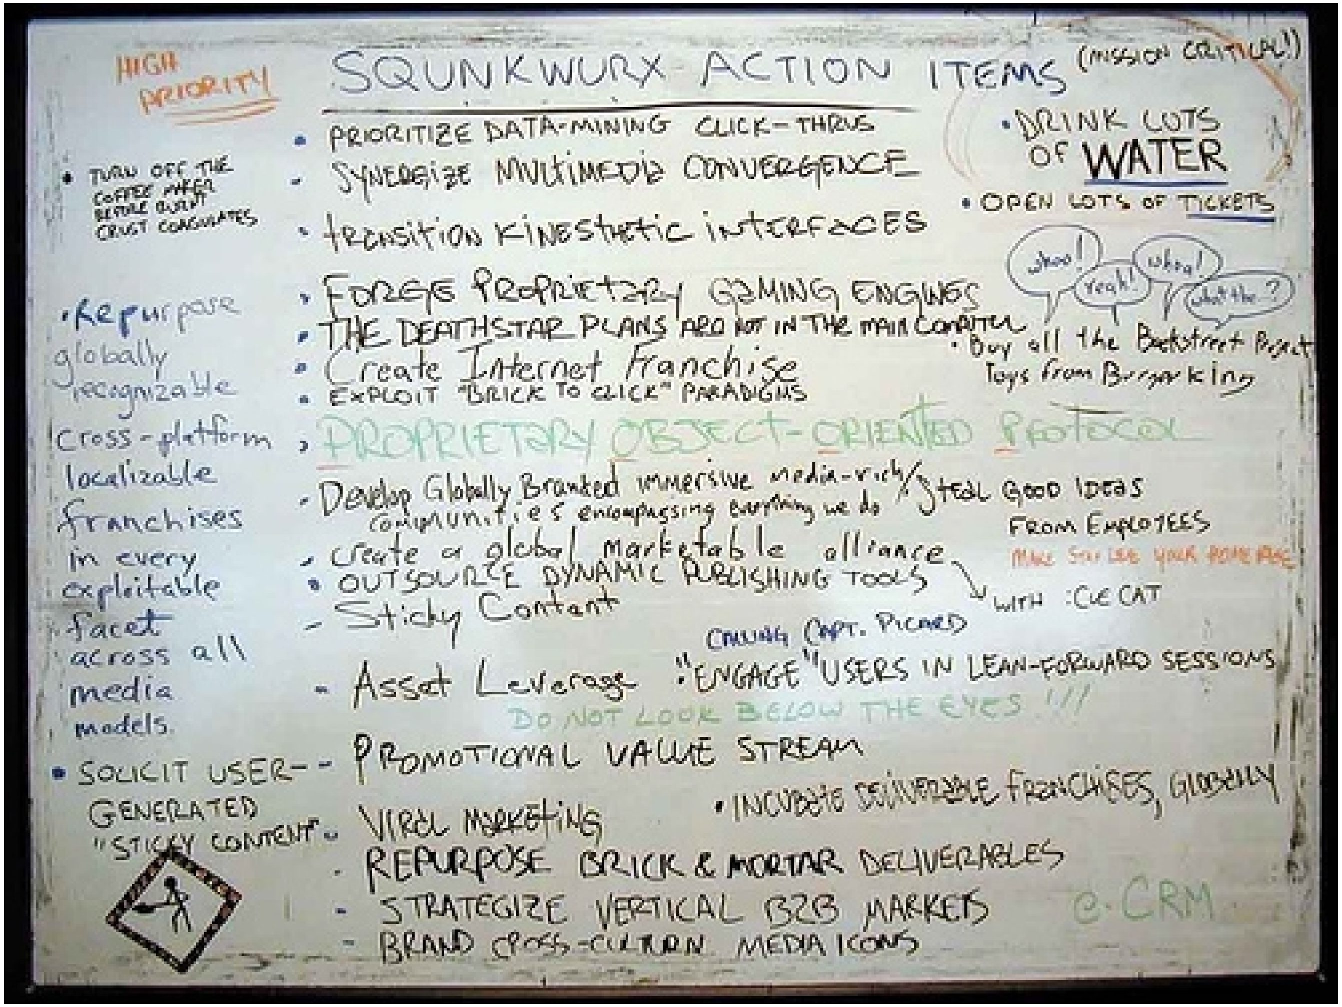

SQUNKWURX ACTION ITEMS (MISSION CRITICAL!)

HIGH PRIORITY

- TURN OFF THE COFFEE MAKER BEFORE BURNT CRUST COAGULATES

- Repurpose globally recognizable cross-platform localizable franchises in every exploitable facet across all media models.

- SOLICIT USER-GENERATED "STICKY CONTENT"

- PRIORITIZE DATA-MINING CLICK-THRUS
- SYNERGIZE MULTIMEDIA CONVERGENCE
- TRANSITION KINESTHETIC INTERFACES
- FORGE PROPRIETARY GAMING ENGINES
- THE DEATHSTAR PLANS ARE NOT IN THE MAIN COMPUTER
- Create Internet Franchise
- EXPLOIT "BRICK TO CLICK" PARADIGMS
- PROPRIETARY OBJECT-ORIENTED PROTOCOL
- Develop Globally Branded immersive media-rich communities encompassing everything we do
- create a global marketable alliance with ICE CAT
- OUTSOURCE DYNAMIC PUBLISHING TOOLS
- Sticky Content
- Asset Leverage
- PROMOTIONAL VALUE STREAM
- VIRAL MARKETING
- REPURPOSE BRICK & MORTAR DELIVERABLES
- STRATEGIZE VERTICAL B2B MARKETS
- BRAND CROSS-CULTURAL MEDIA ICONS

- DRINK LOTS OF WATER
- OPEN LOTS OF TICKETS

whoo!

yeah!

whoa!

what the...?

- Buy all the Backstreet Project Toys from Burger king

- Steal GOOD IDEAS FROM EMPLOYEES

MAKE [illegible] YOUR HOME PAGE

CALLING CAPT. PICARD

- "ENGAGE" USERS IN LEAN-FORWARD SESSIONS

DO NOT LOOK BELOW THE EYES !!!

- INCUBATE DELIVERABLE FRANCHISES, GLOBALLY

e-CRM

http://www.flickr.com/photos/usonian/106005435/

Requirements

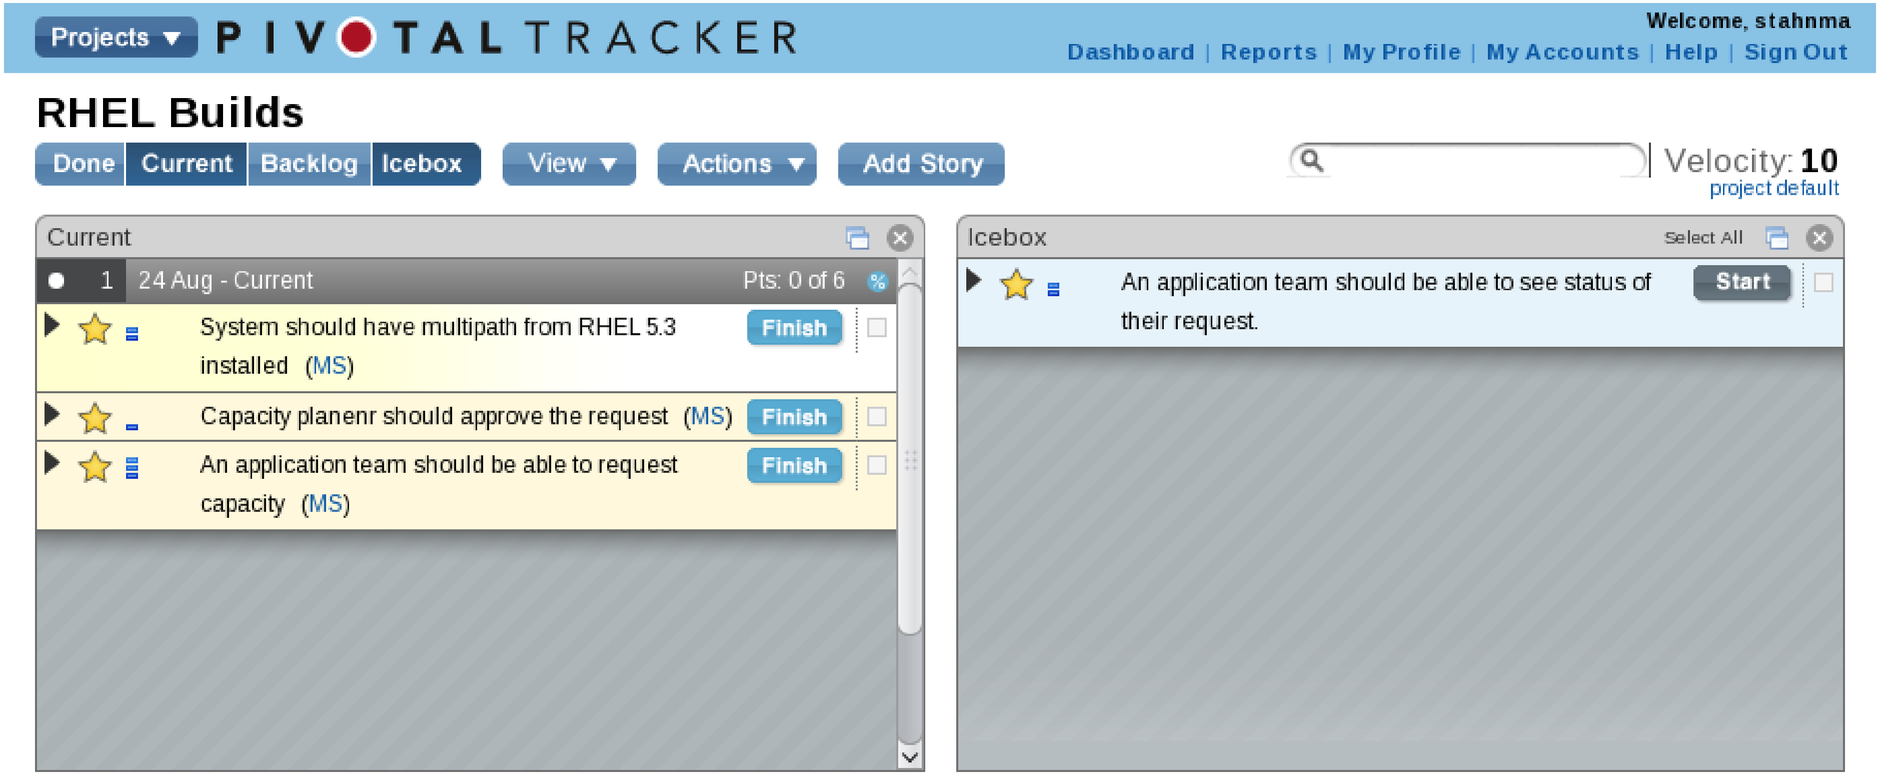

RHEL Builds

Current

● 1 24 Aug - Current Pts: 0 of 6

- System should have multipath from RHEL 5.3 installed (MS) □
- Capacity planenr should approve the request (MS) □
- An application team should be able to request capacity (MS) □

Icebox

- An application team should be able to see status of their request. □

http://pivotaltracker.com

Example Case

Application team needs an instance

2 CPUs

4 GB RAM

100G Disk Space

RHEL 5

http://www.flickr.com:80/photos/jimfrazier/1480809458/

Requirements

Requirements Story:

An application team should be able to request capacity.

Requirements

Requirements Story:

- An application team should be able to request capacity.
- An application team should be able to see the status of their requests.

Requirements

- Requirements Story:
  - An application team should be able to request capacity.
  - An application team should be able to see the status of their requests.
  - A capacity planner should approve the request

Requirements

Requirements Story:

- An application team should be able to request capacity.
- An application team should be able to see the status of their requests.
- A capacity planner should approve the request
- A system should have multipath from RHEL 5.3 installed

Requirements

Requirements Story:

- An application team should be able to request capacity.
- An application team should be able to see the status of their requests.
- A capacity planner should approve the request
- A system should have multipath from RHEL 5.3 installed
- A physical system should have a bonded network connection

Design

Review Stories

Design

Look for common abstractions and workflows

Design

Abstract the process as an interface.

Design

Review the process

Design

Track issues.

Design

Do not automate JUNK

Design

Break work into small units

Design

Don't be afraid to say 'good enough for version 1'

Design

Don't be afraid to say 'good enough for version 1'

- Limit Scope.

Design

Progress over Perfection

Design

A good enough system today is better than a perfect one that hasn't been invented.

Design

You will iterate again.

You will iterate again.

Design

What points of your process require measurement?

Design

Can that be automated?

Design

Can you test for it?

Design Your Tests

Am I testing directly against my stories? (requirements)

Design Your Tests

Does it have the right version of Multipath?

Design Your Tests

Is the network connection bonded?

Testing

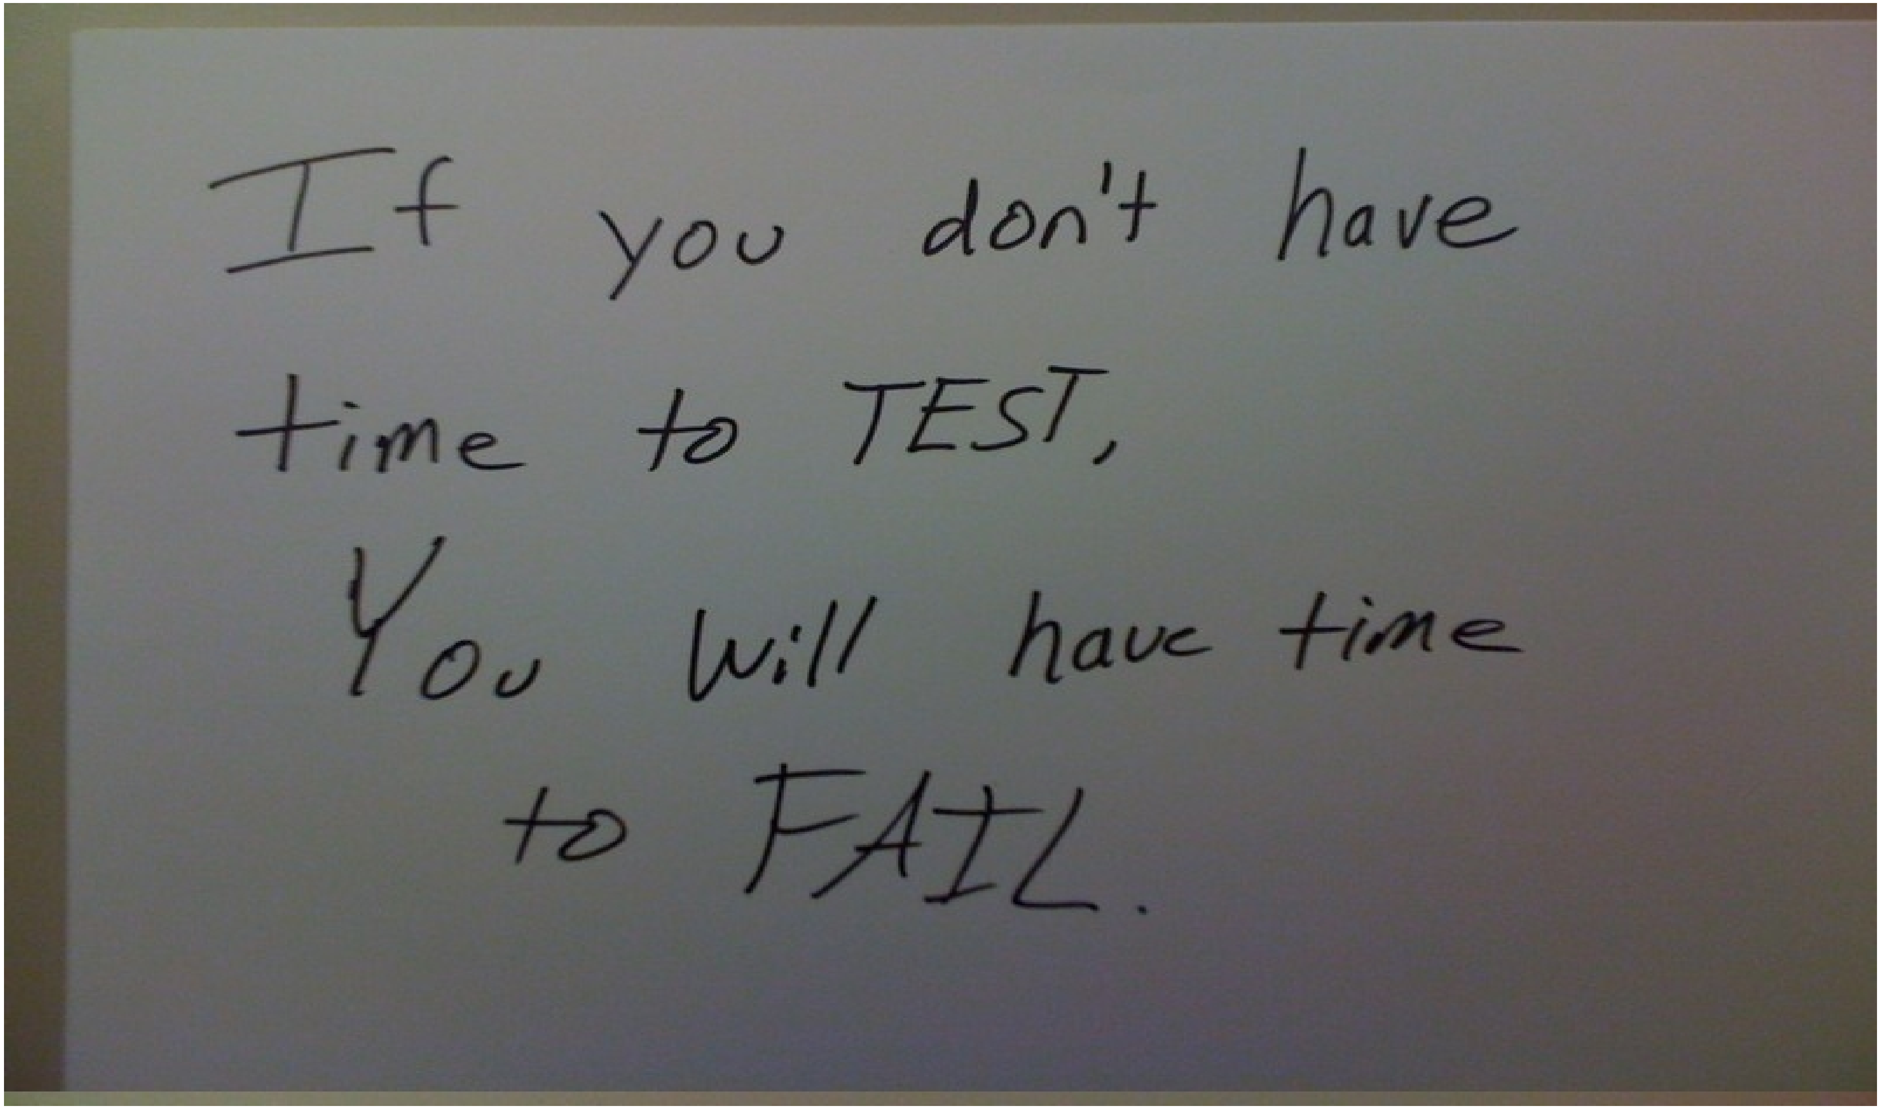

If you don't have time to TEST, You will have time to FAIL.

Testing

Provides quality measurements

Testing

Provides metrics for reporting.

Testing

Your management team will love it

Testing

You will have facts.

Testing

You will have facts. Your team will love it.

Specifications for Testing

Write Simple Tests

Specifications for Testing

Unit Test your Configuration

Specifications for Testing

Watch it FAIL.

Specifications for Testing

Watch it FAIL. Good tests FAIL often.

Specifications for Testing

http://www.flickr.com/photos/fireflythegreat/2845637227/

Specifications for Testing

Write a policy/method/procedure/function to fix it

Specifications for Testing

Watch it PASS

Specifications for Testing

Store Results.

Test Case

Network configuration should be bonded.

Test Case

Network configuration should be bonded.

Can I script this?

Test Case

Network configuration should be bonded.

Can I script this?

- Given enough Time and Money, it's all scriptable.

Test Case

Network configuration should be bonded.

Can I script this?

- Given enough Time and Money, it's all scriptable.

Can I store the results?

Test Case

- Network configuration should be bonded.
- Can I script this?
  - Given enough Time and Money, it's all scriptable.
- Can I store the results?
  - Yes

Test Case

Network configuration should be bonded.

Can I script this?

- Given enough Time and Money, it's all scriptable.

Can I store the results?

- Yes

How Often should I test this?

Test Case

```
#!/bin/bash

. config_vars.sh

rc=0

# First, if we a VM exit all good
if ( lspci | grep -i vmware &> /dev/null ) ; then
    rc=0
else
    if ( ! ifconfig -a | grep -i bond0 &> /dev/null ) ; then
       echo "FAIL: Network interfaces not bonded."
       rc=1
    fi
fi

exit $rc
```

Test Case

Site Specific Data

```
#!/bin/bash

. config_vars.sh

rc=0

# First, if we a VM exit all good
if ( lspci | grep -i vmware &> /dev/null ) ; then
    rc=0
else
    if ( ! ifconfig -a | grep -i bond0 &> /dev/null ) ; then
       echo "FAIL: Network interfaces not bonded."
       rc=1
    fi
fi

exit $rc
```

Test Case

```
#!/bin/bash

. config_vars.sh

rc=0

# First, if we a VM exit all good
if ( lspci | grep -i vmware &> /dev/null ) ; then
    rc=0
else
    if ( ! ifconfig -a | grep -i bond0 &> /dev/null ) ; then
       echo "FAIL: Network interfaces not bonded."
       rc=1
    fi
fi

exit $rc
```

Site Specific Data

Decoupled Logic

Test Case

I have results, where do I put them?

Pick One

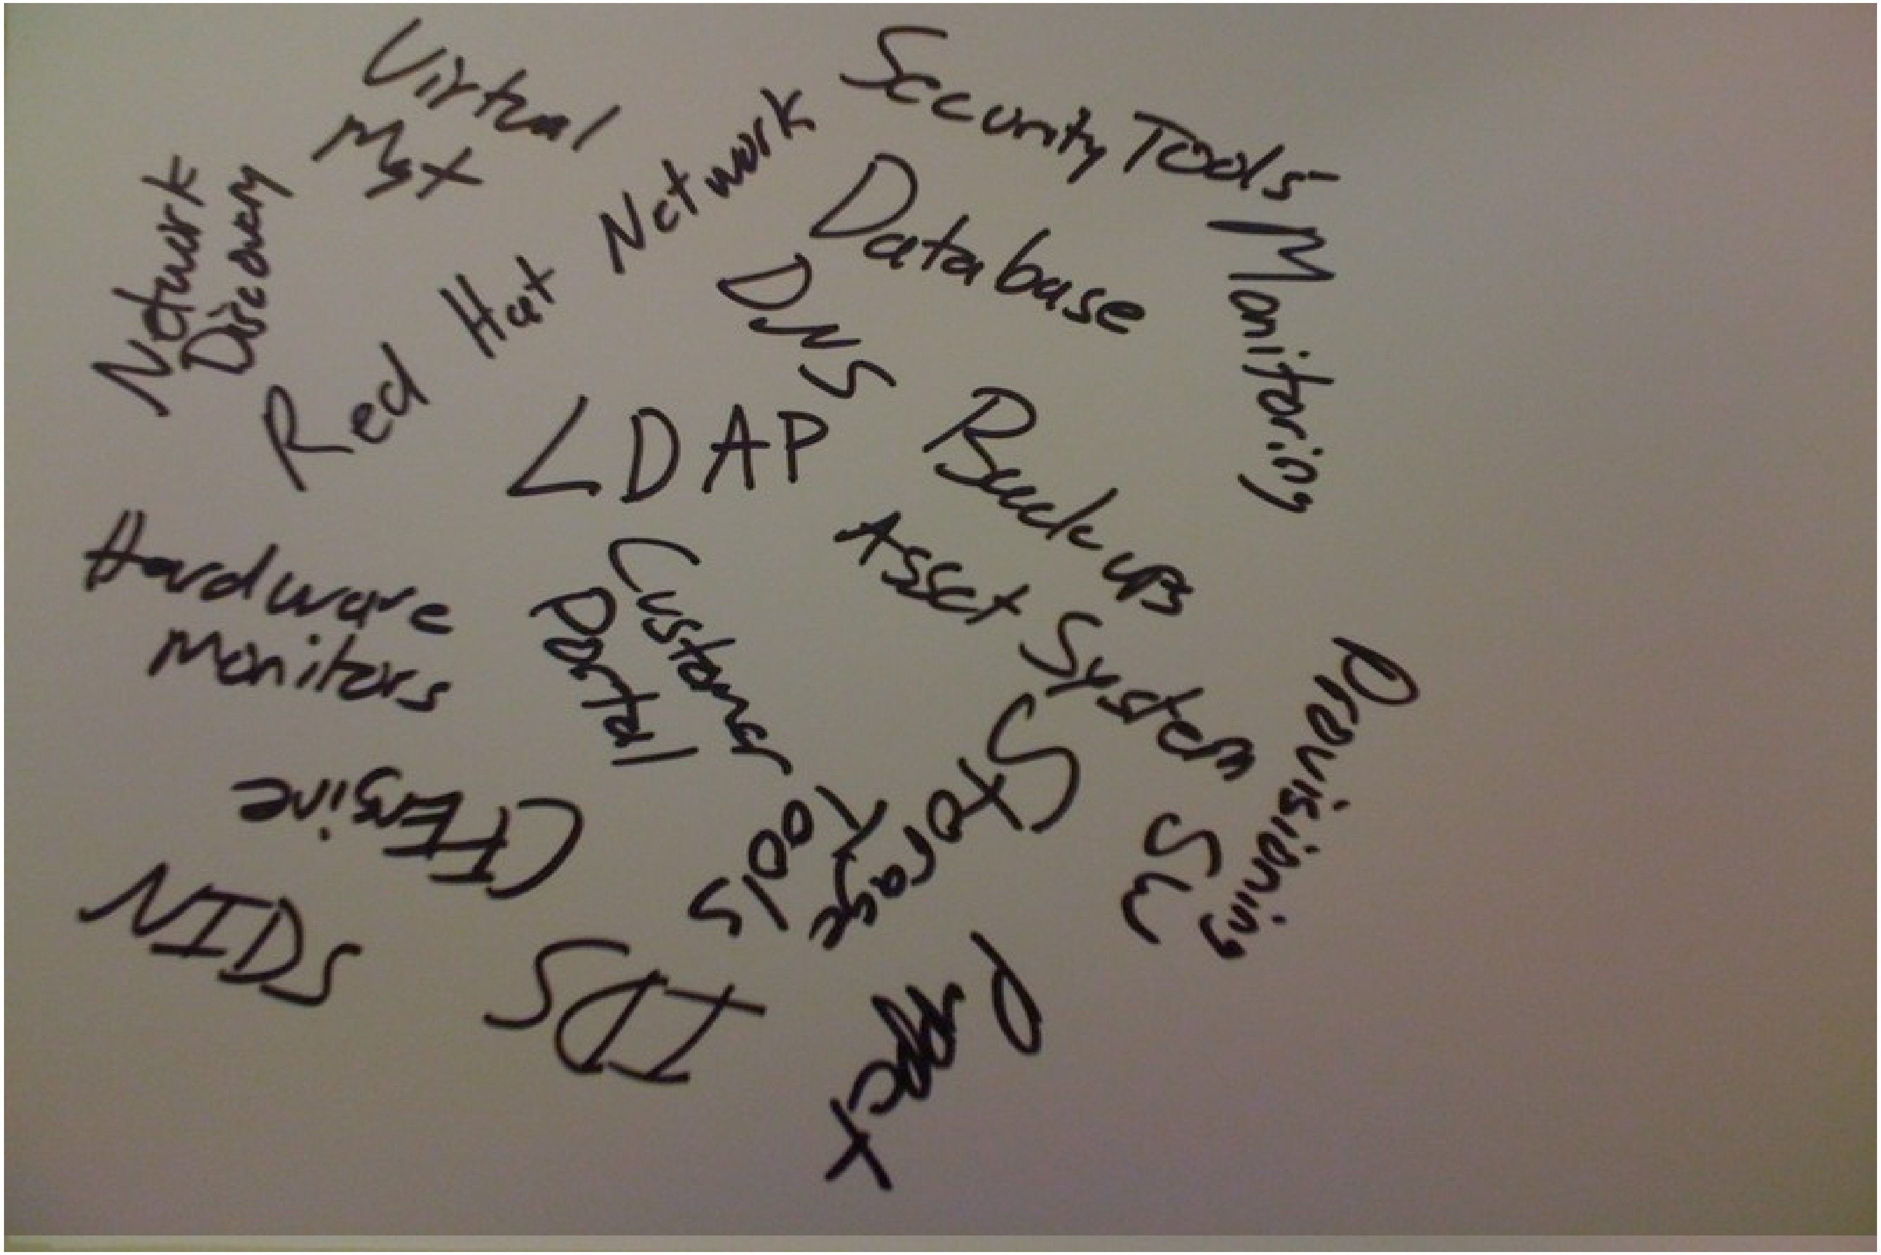
Virtual Mgt
Network Discovery
Red Hat Network
Security Tools
Database
Monitoring
DNS
LDAP
Backups
Asset System
Hardware monitors
Customer Portal
Provisioning
Offensive
Storage Tools
Su
NIDS
IDS
Puppet

Example Storage

Ruby using RHN API

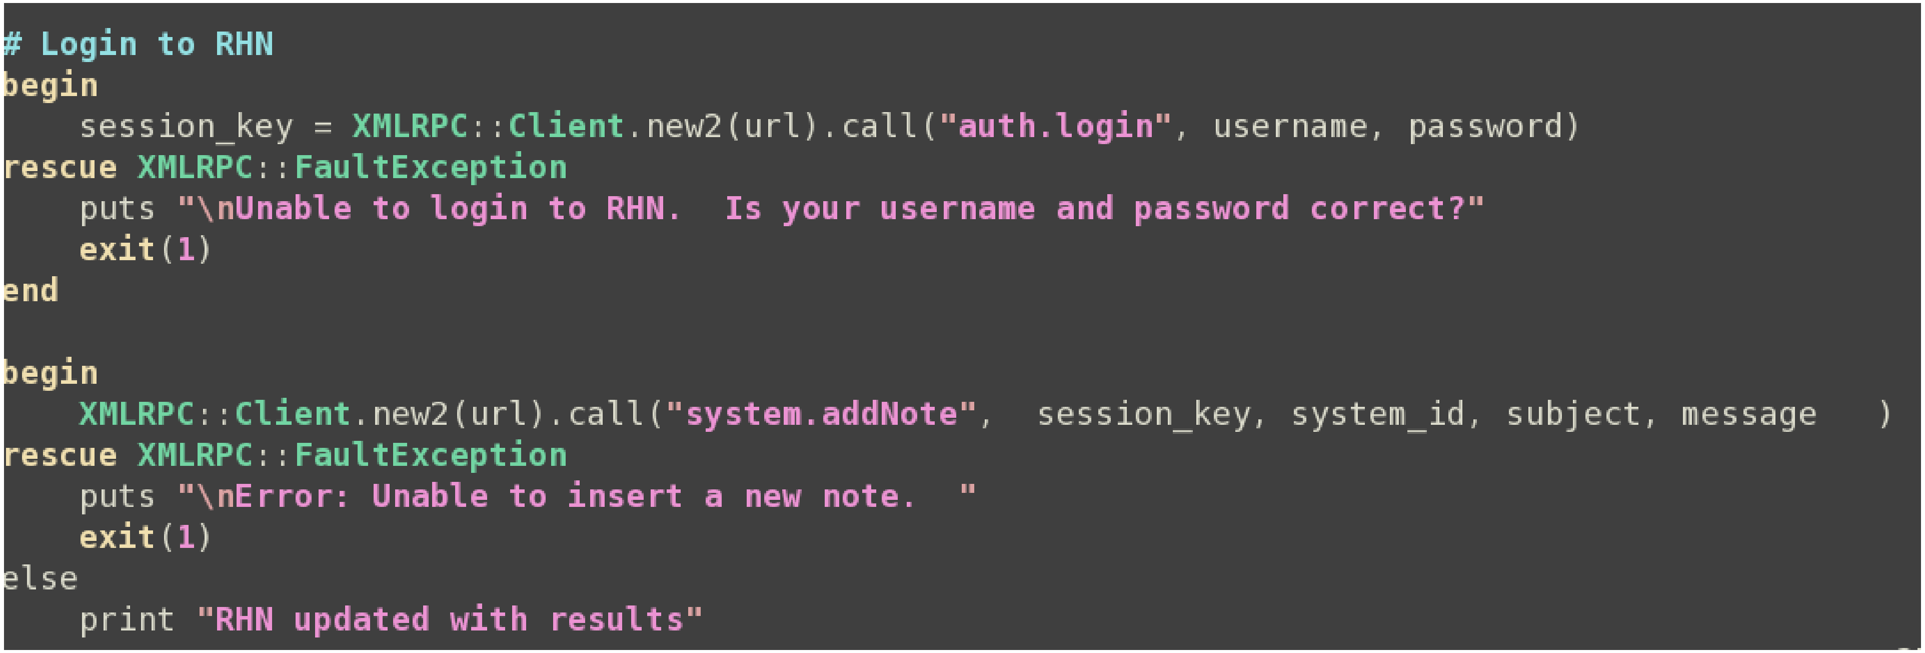

```
# Login to RHN
begin
    session_key = XMLRPC::Client.new2(url).call("auth.login", username, password)
rescue XMLRPC::FaultException
    puts "\nUnable to login to RHN.  Is your username and password correct?"
    exit(1)
end

begin
    XMLRPC::Client.new2(url).call("system.addNote",  session_key, system_id, subject, message   )
rescue XMLRPC::FaultException
    puts "\nError: Unable to insert a new note.  "
    exit(1)
else
    print "RHN updated with results"
```

http://www.redhat.com/spacewalk

Implementation

Run through the process a couple times **manually**

Implementation

Think about re-use

- Inventory Queries
- Asking for/updating configuration info
- Authentication Modules

Implementation

Are existing tools close to what you need?

Implementation

Can you reuse existing code?

Test Again

After implementation, test again.

Test Again

After implementation, test again.

- It's automated right?

Test Again

- After implementation, test again.
  - It's automated right?
    - So it's quick and painless.

Close the loop

Maintain state of your process

Close the loop

Maintain state of your components

Close the loop

If the version of package X changes, and you have a policy for package X, report.

Close the loop

If the version of package X changes, and you have a policy for package X, report^H^H^H^H^H fix it.

Close the loop

If don't have a network bond, report it or fix it.

Maintain

Most fun step

Deployment Example

Host is RHEL 5.3

Host has the correct Version of Multipath

Host should have 4 GB RAM

Host should have 2 CPUs

Host must have a network bond if it's physica

Maintain

Offers most improvement

Maintain

Customers offer ideas

Maintain

Other admins complain (offer ideas)

Maintain

Iterate

Maintain

Lather, rinse, repeat.

CI

You need CI.

Continuous Integration

CI is used in the best development shops in the world

Continuous Integration

Run upon commits

Continuous Integration

You are using a VCS right?

Continuous Integration

You are using a VCS right?

- Please

Continuous Integration

Track Failures

Continuous Integration

Any new bugs, require Tests

Continuous Integration

Automate Builds

Self-test the Build

Commit early, commit often

Have a test environment

Report Progress for everybody (laconica?)

Automate Deployment – (usage packages)

Can you have a 'release' of infrastructure?

Yes

Plan change-sets

Plan your work-life balance

Work with setting expectations for consumers of your infrastructure

Work on hot/rolling releases

Communicate like a politician, except tell the truth

Some common areas for improvement

A few things to avoid when crafting solutions.

Things to Avoid

Meatcloud

Things to Avoid

Automation of Junk

Things to Avoid

Talking about IT rather than doing IT

Things to Avoid

Thinking there is one solution

Things to Avoid

UberTools – Silver Bullets

Things to Avoid

New Hotness

Things to Avoid

NIH Syndrome

Axiom 1

Reuse before building or purchasing

Axiom 2

Don't Leverage the Meatcloud

Axiom 3

Decouple your Infrastructure.

Go From This:

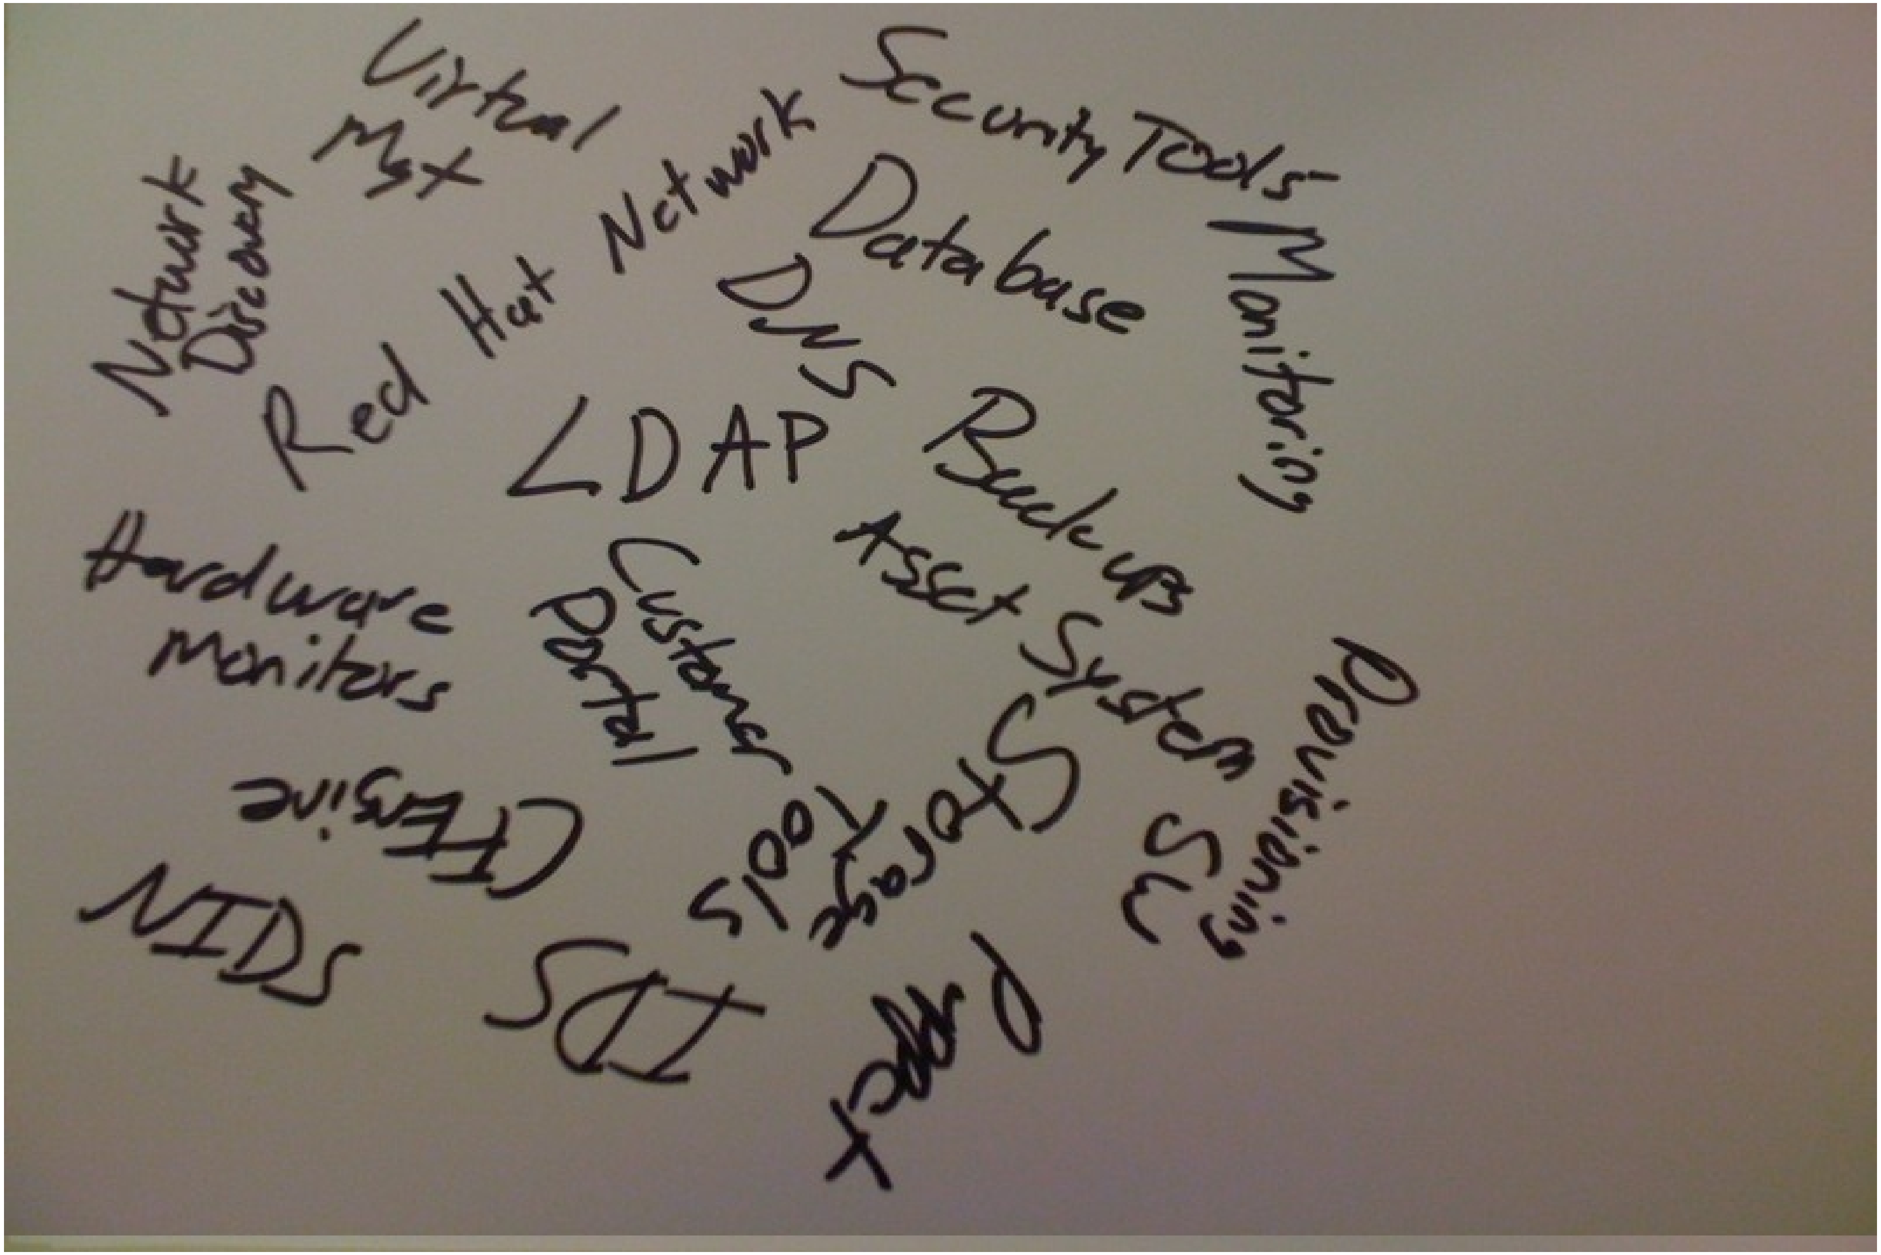
Virtual Mgt
Network Discovery
Red Hat Network
Security Tools
Database
DNS
Monitoring
LDAP
Backups
Asset System
Hardware Monitors
Customer Portal
Offensive
Storage Tools
Provisioning
SW
NIDS
SDI
Puppet

To This:

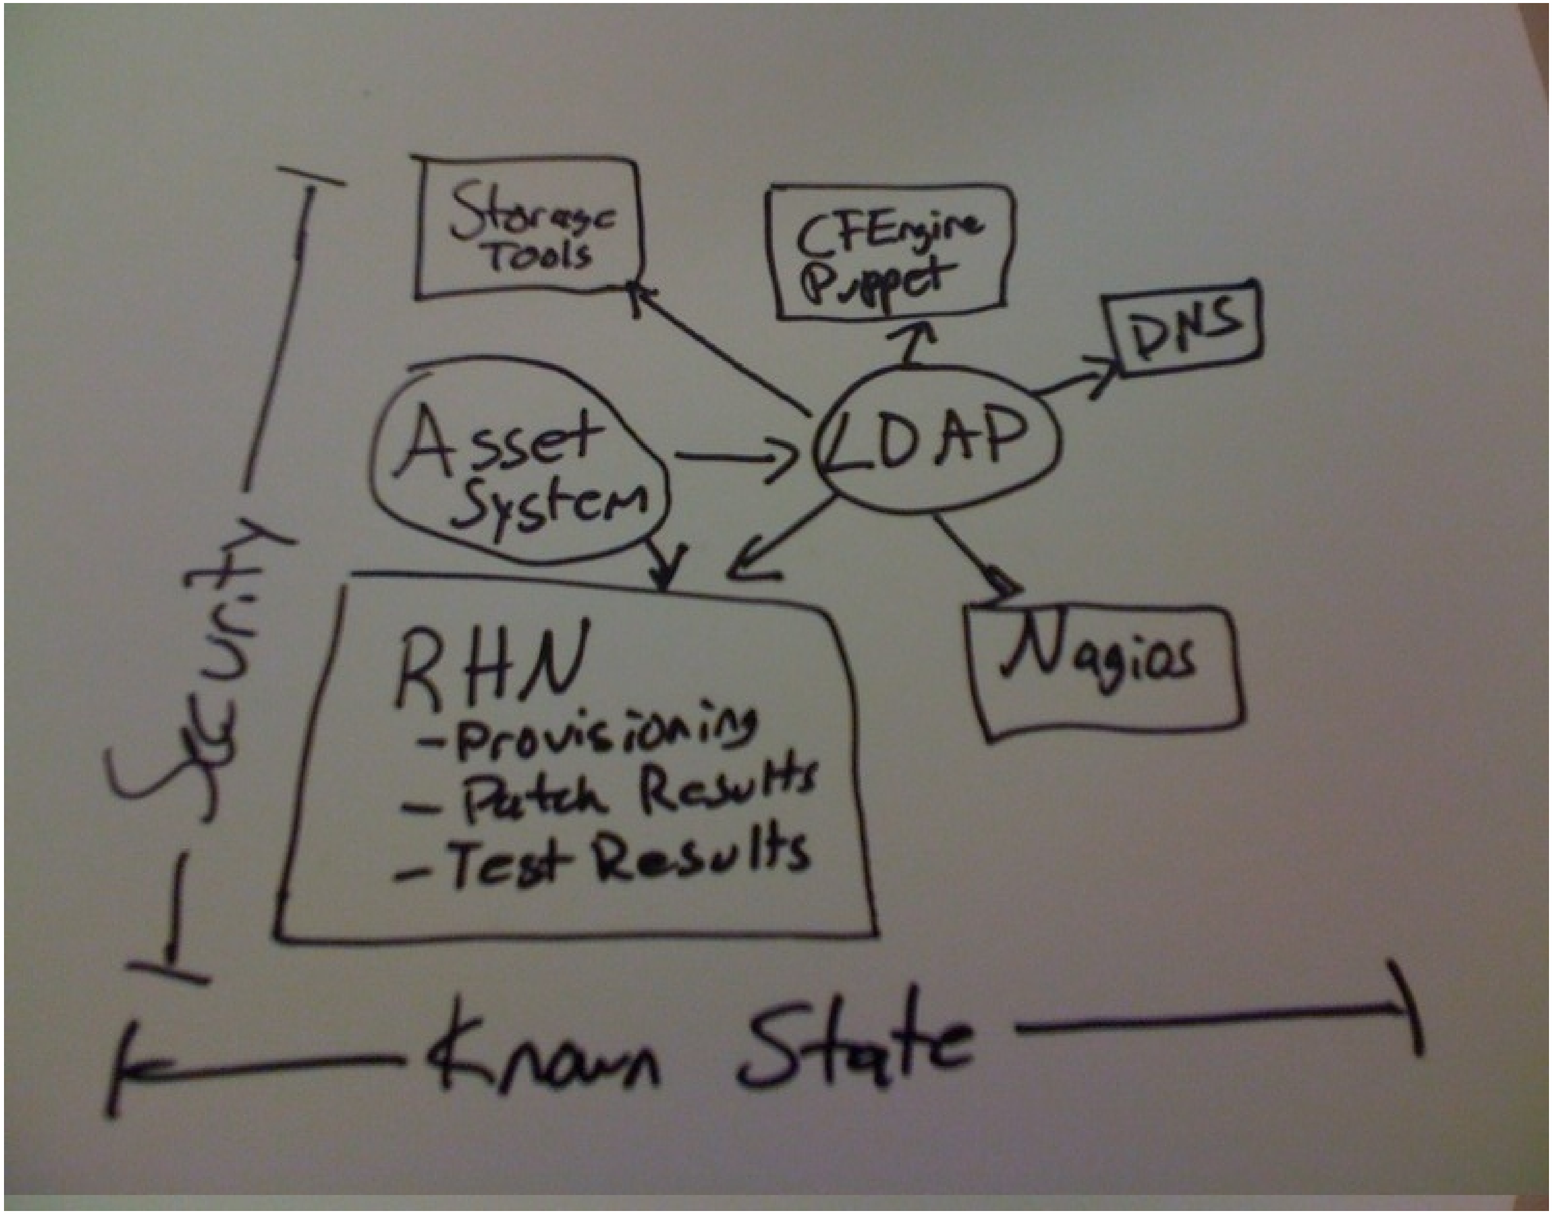
Storage Tools
CFEngine Puppet
DNS
Asset System
LDAP
Nagios
RHN
- Provisioning
- Patch Results
- Test Results
Security
Known State

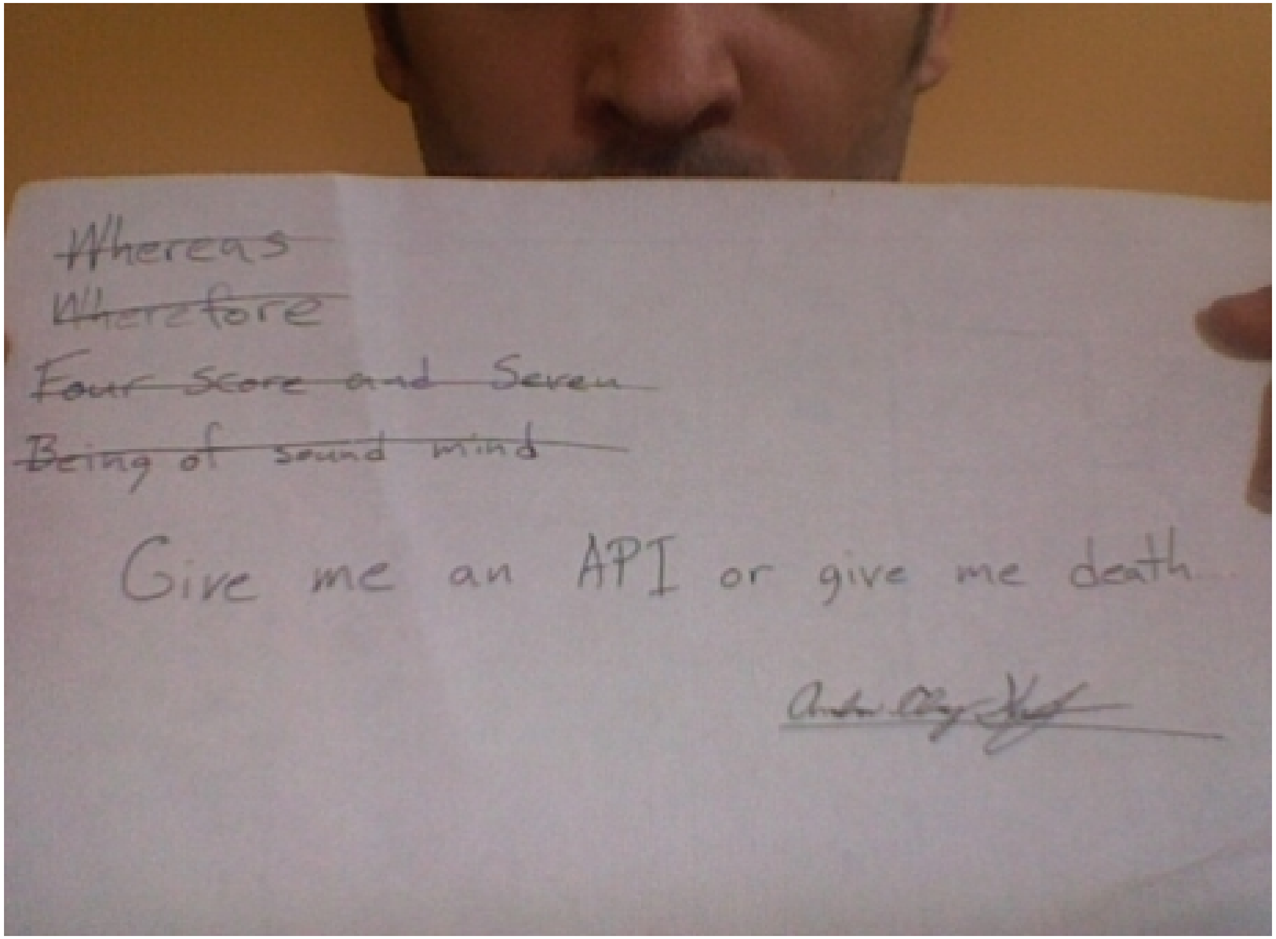

~~Whereas~~

~~Wherefore~~

~~Four Score and Seven~~

~~Being of sound mind~~

Give me an API or give me death.

PuT
a CoRK
IN
IT

Sources

http://stochasticresonance.wordpress.com/2009/04/01/mea

http://www.boingboing.net/2007/08/22/panflute-flowchart.ht

http://gilesbowkett.blogspot.com/2007_05_01_archive.htm

http://www.verber.com/mark/sysadm/how-many-admins.htm

Suggested Reading

http://meetronome.com

http://gist.github.com/161265

https://fedoraproject.org/wiki/Infrastructure

https://fedorahosted.org/csi/

Ship It – Pragmatic Press

Automating Unix and Linux Administration – Apress

Pro Unix Administration – Apress

Pulling Strings with Puppet - Apress

Release IT – Pragmatic Press

More Reading

http://groups.google.com/group/agile-system-administratio

http://www.melconway.com/research/committees.html

http://status.net/?source=laconica

http://reductivelabs.com/trac/puppet/

https://fedorahosted.org/cobbler/

http://www.mockobjects.com/book/

Infrastructure is Development