

AVAILABILITY, THE CLOUD AND EVERYTHING


JOE WILLIAMS



CLOUDANT

ME

- Joe Williams
 - Infrastructure Engineer
 - Cloudant
 - @williamsjoe
 - joeandmotorboat.com



CLOUDANT

- Distributed database built on CouchDB
- Real-time Search and Analytics
- Sign Up for our Beta!
- cloudant.com

BIAS

- Distributed Databases
- Amazon EC2
- Chef
- Erlang

AVAILABILITY



AVAILABILITY

- What is *Availability*?

AVAILABILITY

- Uptime / Downtime
 - Perceived
 - Actual

AVAILABILITY

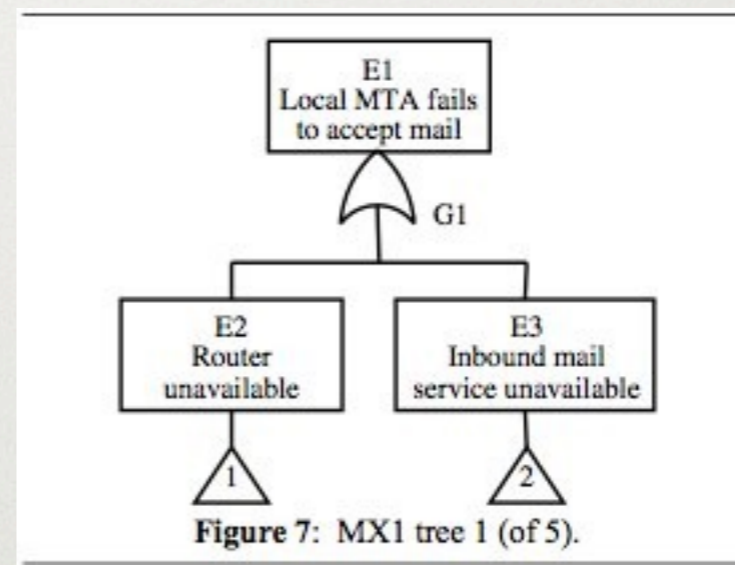
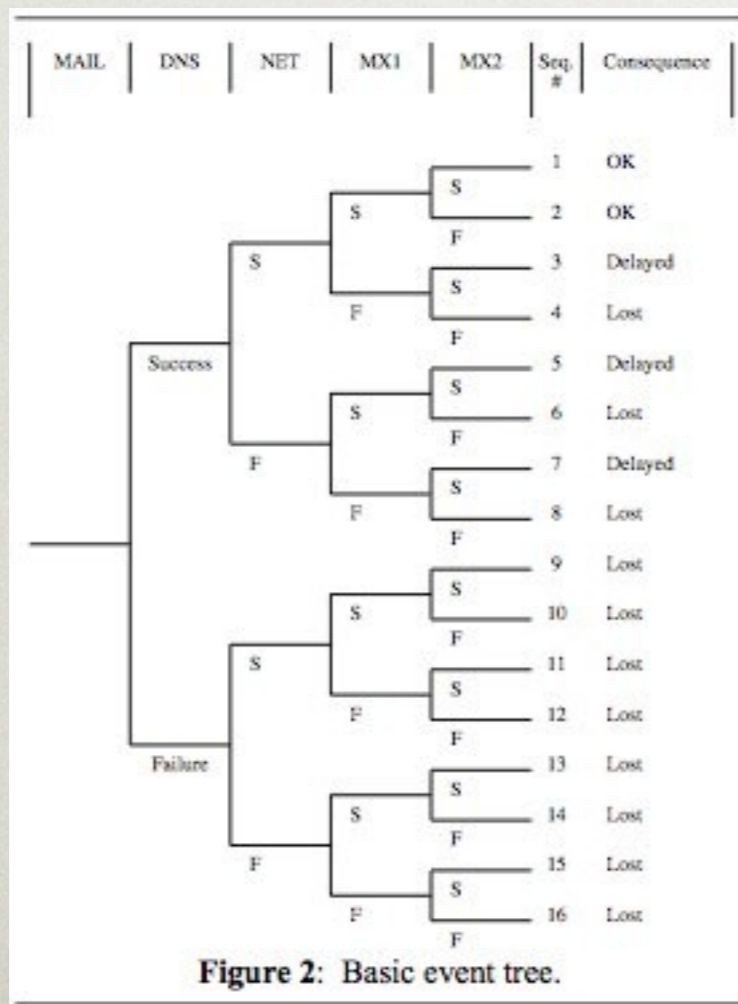
- Reliability
 - Failure rate / Mean time between failures
 - Failure Resistance / Fault Tolerance
 - Durability

AVAILABILITY

- Probabilistic Risk Assessment
 - Event Tree Analysis
 - Fault Tree Analysis

AVAILABILITY

- Future Work
- Erlang Supervision Trees and FTA / ETA



AVAILABILITY

- Performance, Logging and Monitoring
 - Baracus, Rsyslog, Munin

AVAILABILITY



The image is a screenshot of a Twitter post. At the top left is the Twitter logo. To the right are navigation links: Home, Profile, Find People, Settings, Help, and Sign out. The main text of the tweet reads: "Blaming a vendor for your outage denies the cardinal rule of #webops: you own your availability." The text "#webops" is in red. To the right of the text is a yellow star icon. Below the text, it says "3:57 PM Dec 11th, 2009 from Twitterrific" and "Retweeted by 8 people". To the right of this are "Reply" and "Retweet" icons. Below the tweet is the user's profile information: a profile picture of Benjamin Black, the name "benjaminblack" in red, and "Benjamin Black" below it. At the bottom of the page is a footer with copyright information and various links: © 2009 Twitter, About Us, Contact, Blog, Status, Goodies, API, Business, Help, Jobs, Terms, Privacy.

twitter Home Profile Find People Settings Help Sign out

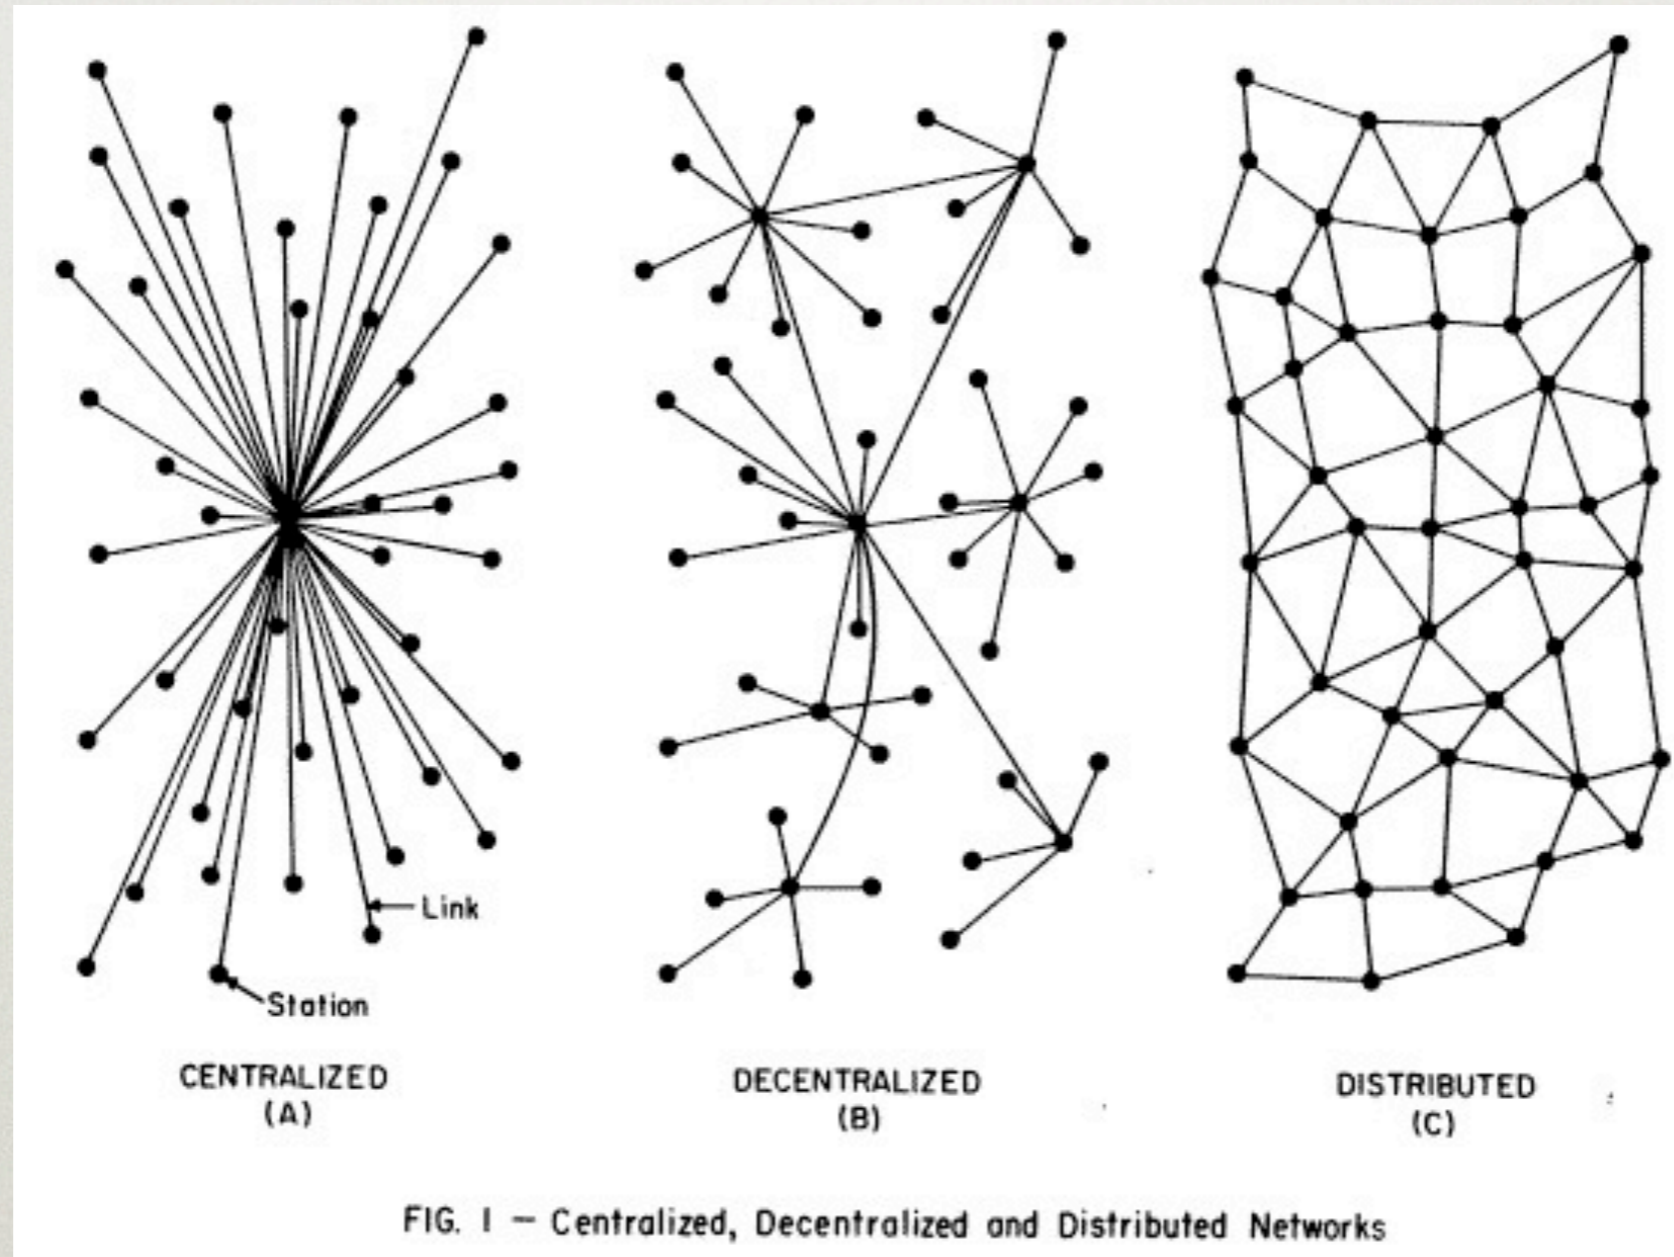
Blaming a vendor for your outage
denies the cardinal rule of **#webops**: you
own your availability. ★

3:57 PM Dec 11th, 2009 from Twitterrific
Retweeted by 8 people Reply Retweet

 **benjaminblack**
Benjamin Black

© 2009 Twitter About Us Contact Blog Status Goodies API Business Help Jobs Terms Privacy

DISTRIBUTED SYSTEMS



Baran (http://www.rand.org/pubs/research_memoranda/RM3420/)

DISTRIBUTED SYSTEMS

- Abstraction

DISTRIBUTED SYSTEMS

- Redundancy
 - Duplication
 - Distribution

DISTRIBUTED SYSTEMS

- RAID ain't as redundant as it used to be.

Leventhal (<http://queue.acm.org/detail.cfm?id=1670144>)

DISTRIBUTED SYSTEMS

- Alphabet Soup
 - ACID, CAP, BASE, 2PC, MVCC
 - Vector Clocks, Eventual Consistency
 - Dynamo, Paxos, Chandra, Byzantine

DISTRIBUTED SYSTEMS

- CAP == Availability

DISTRIBUTED SYSTEMS

- Effects on *Availability*

THE CLOUD



THE CLOUD

- Abstraction
- Commoditization

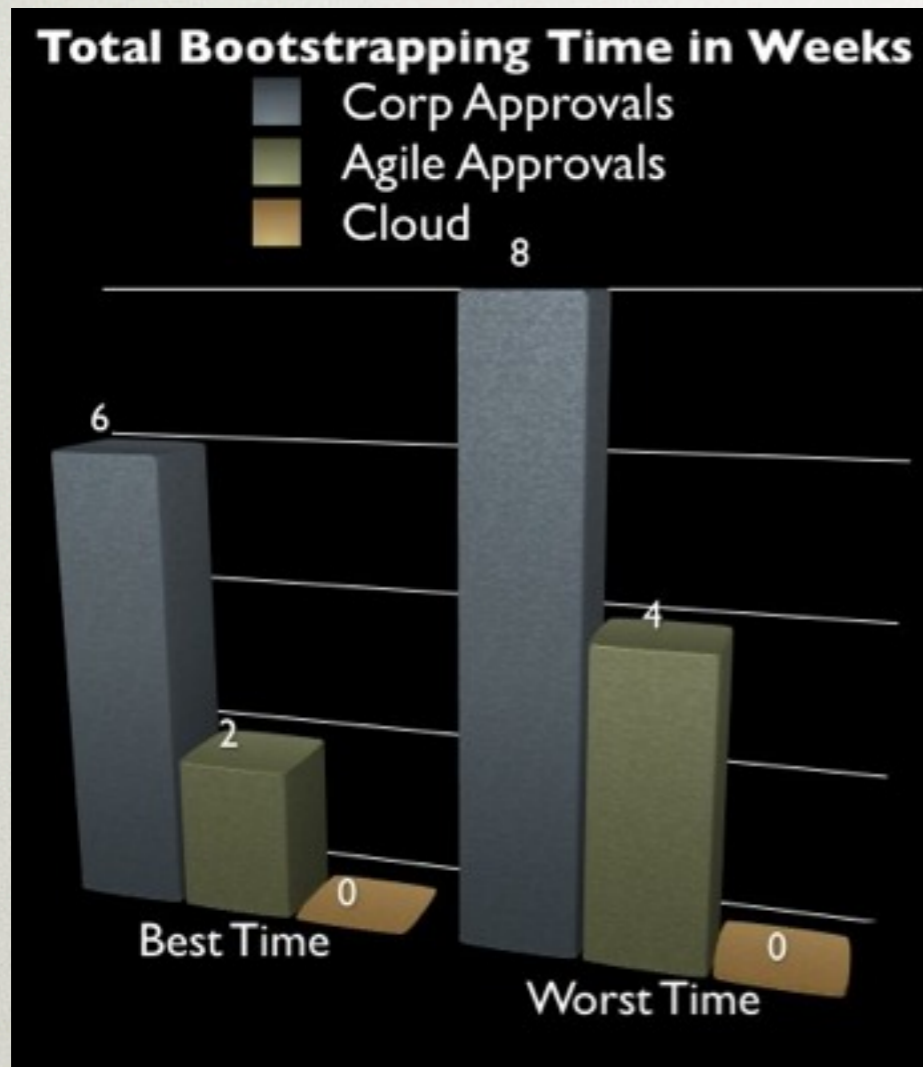
THE CLOUD

- Costs
 - Loss of Control
 - Single Points of Failure
 - Network Partitions / Data Locality
 - Unreliable
 - Performance

THE CLOUD

- Benefits
 - API to everything
 - Fast and Flexible Resource Mgmt
 - “Unlimited” Resources

THE CLOUD



- Bootstrapping
- Time and Effort

Adam Jacob and Ezra Zygmuntowicz (<http://blip.tv/file/2285124/>)

THE CLOUD

- Nodes are stateless and disposable.

THE CLOUD

- Compute vs Data
- Performance

THE CLOUD



- Costs (\$\$) - EC2
 - Different prices in different zones
 - Plus data transfer, EBS, and extras
 - Base Instance Price Range:
 - m1.small ~ \$62 / mo
 - m2.4xlarge ~ \$1750 / mo

<http://aws.amazon.com/ec2/#pricing>

THE CLOUD

- Effects on *Availability*

THE CLOUD

"Clouds are systems ... and with systems, you have to think hard and know how to deal with issues in that environment. The scale is so much bigger, and you don't have the physical control.

But we think people should be optimistic about what we can do here. If we are clever about deploying cloud computing with a clear-eyed notion of what the risk models are, maybe we can actually save the economy through technology."

~ Security in the Ether By David Talbot - MIT Technology Review Jan/Feb 2010

AUTOMATION AND CONFIGURATION MGMT



ETTEETOODE-A

Etteetoode-a is a leettle-a theeng that mekes a beeg
deeffference-a. Bork Bork Bork!

AUTOMATION AND CONFIGURATION MGMT

- Abstraction

AUTOMATION AND CONFIGURATION MGMT

- Implementations
 - Chef, Puppet, bcfg2, cfengine

AUTOMATION AND CONFIGURATION MGM

- Chef
 - Cookbooks / Recipes
 - File Templates
 - Resources / Providers
 - Provides a DSL to every part of your system

AUTOMATION AND CONFIGURATION MGMT

- Erlang + Chef (v0.8)
- erl_call Provider

```
erlang.rb #  
  
1  version = "1.0"  
2  
3  erl_call "erlang hot upgrade" do  
4    cookie "erlangcookie"  
5    node_name "upgrade@hostname"  
6    name_type "sname"  
7    code <<-EOH  
8    release_handler:unpack_release("yourapp_#{version}.tar.gz"),  
9    release_handler:install_release(version),  
10   release_handler:make_permanent(version).  
11   EOH  
12  end
```


AUTOMATION AND CONFIGURATION MGMT

- Erlang
 - Hot Code Upgrades
 - Distributed Upgrades are HARD

AUTOMATION AND CONFIGURATION MGMT

- Effects on *Availability*

SOURCES AND FURTHER READING

Availability

http://en.wikipedia.org/wiki/Fault_Tree_Analysis
http://en.wikipedia.org/wiki/Probabilistic_risk_assessment
<http://www.usenix.org/events/lisa01/tech/apthorpe/apthorpe.ps>

The Cloud

<http://cseweb.ucsd.edu/~hovav/dist/cloudsec.pdf>
<http://cloudscaling.com/blog>
<http://open.eucalyptus.com/wiki/presentations>
<http://www.morganclaypool.com/doi/pdf/10.2200/S00193ED1V01Y200905CAC006>

Configuration Mgmt


<http://wiki.opscode.com/>
<http://kallistec.com/2010/01/23/the-chef-way/>
<http://onlamp.com/pub/a/onlamp/2004/04/15/cfengine.html>
<http://blog.loftninja.org/2010/01/22/configuration-management-vs-meatcloud-5-reasons-cm-wins/>
http://cfwiki.org/cfwiki/index.php/Singlecopy_Nirvana

Distributed Systems

http://en.wikipedia.org/wiki/Byzantine_agreement
http://en.wikipedia.org/wiki/Paxos_algorithm
http://en.wikipedia.org/wiki/Multiversion_concurrency_control
<http://en.wikipedia.org/wiki/ACID>
<http://www.julianbrowne.com/article/viewer/brewers-cap-theorem>
[http://en.wikipedia.org/wiki/Dynamo_\(storage_system\)](http://en.wikipedia.org/wiki/Dynamo_(storage_system))
<http://en.wikipedia.org/wiki/2PC>
<http://queue.acm.org/detail.cfm?id=1394128>
<http://it.toolbox.com/blogs/oracle-guide/acid-vs-base-25938>
<http://devblog.streamy.com/2009/08/24/cap-theorem/>
<http://www.infoq.com/presentations/availability-consistency>
<http://citeseer.ist.psu.edu/652189.html>
<http://labs.google.com/papers/bigtable-osdi06.pdf>
<http://database.cs.brown.edu/sigmod09/benchmarks-sigmod09.pdf>
http://www.allthingsdistributed.com/2007/10/amazons_dynamo.html
<http://www.ctlab.org/documents/How%20Complex%20Systems%20Fail.pdf>
<http://labs.google.com/papers/mapreduce-osdi04.pdf>
http://www.rand.org/pubs/research_memoranda/RM3420/
http://www.cs.jhu.edu/~jak/docs/paxos_for_system_builders.pdf
<http://citeseer.ist.psu.edu/652189.html>
<http://blog.basho.com/2010/01/29/why-vector-clocks-are-easy/>

QUESTIONS?

Joe Williams - @williamsjoe



CLOUDANT