

Building a Culture of Observability Within your Organization

CNCF Webinar Series - May 7
Grant Schofield - Humio

Greetings.

Who am I?



<https://humio.com>

Overview

What is observability?

Observability's Past, Present, and Future

How can my organization make their systems observable?

What Observability isn't....

New concept

Monitoring

A tool or product

Just your operations team's concern

Observability is...

The intersection systems, applications, and **users**

An attribute of a system than encompasses many others such as: functionality, performance, testability, maintainability, monitorability

The ability to know how your system is working from the outside and being able to contextualize any events within

Observability is...(to me)

Having the right systems, and implementation, to know how your system is feeling about any request

Having the ability to know why the needle is oriented a particular way in the haystack

Having what I need, as an engineer, to debug a problem with minimal steps

Most importantly..

A cultural facet of your organization, like DevOps
Responsibility of everyone in the organization
Different for every organization



The Past

Control Theory

Control theory in control systems engineering is a subfield of mathematics that deals with the control of continuously operating dynamical systems in engineered processes and machines. The objective is to develop a control model for controlling such systems using a control action in an optimum manner without delay or overshoot and ensuring control stability.

Rudolf E. Kálmán

In control theory, observability is a measure of how well internal states of a system can be inferred from knowledge of its external outputs. The observability and controllability of a system are mathematical duals.

Kálmán's ground breaking work in the 50s and 60s that, lead to Kálmán filters, was used by Nasa for the Apollo and shuttle programs as well as a vast array of other applications.

Observing Systems of the Past

SaaS didn't exist, built everything

Basic Open-source Software - Big Brother, RRDTool, MRTG, Cacti, NetSaint (Nagios)

Expensive (ugly) Enterprise tools - HP Openview, CA, etc

Script everything, or die trying

CREAM - *Cron Rules Everything Around Me*

World ran on rsync

What did we build?

Monitoring

big brother last update
Tue Jan 30 17:27:45 EST 2001

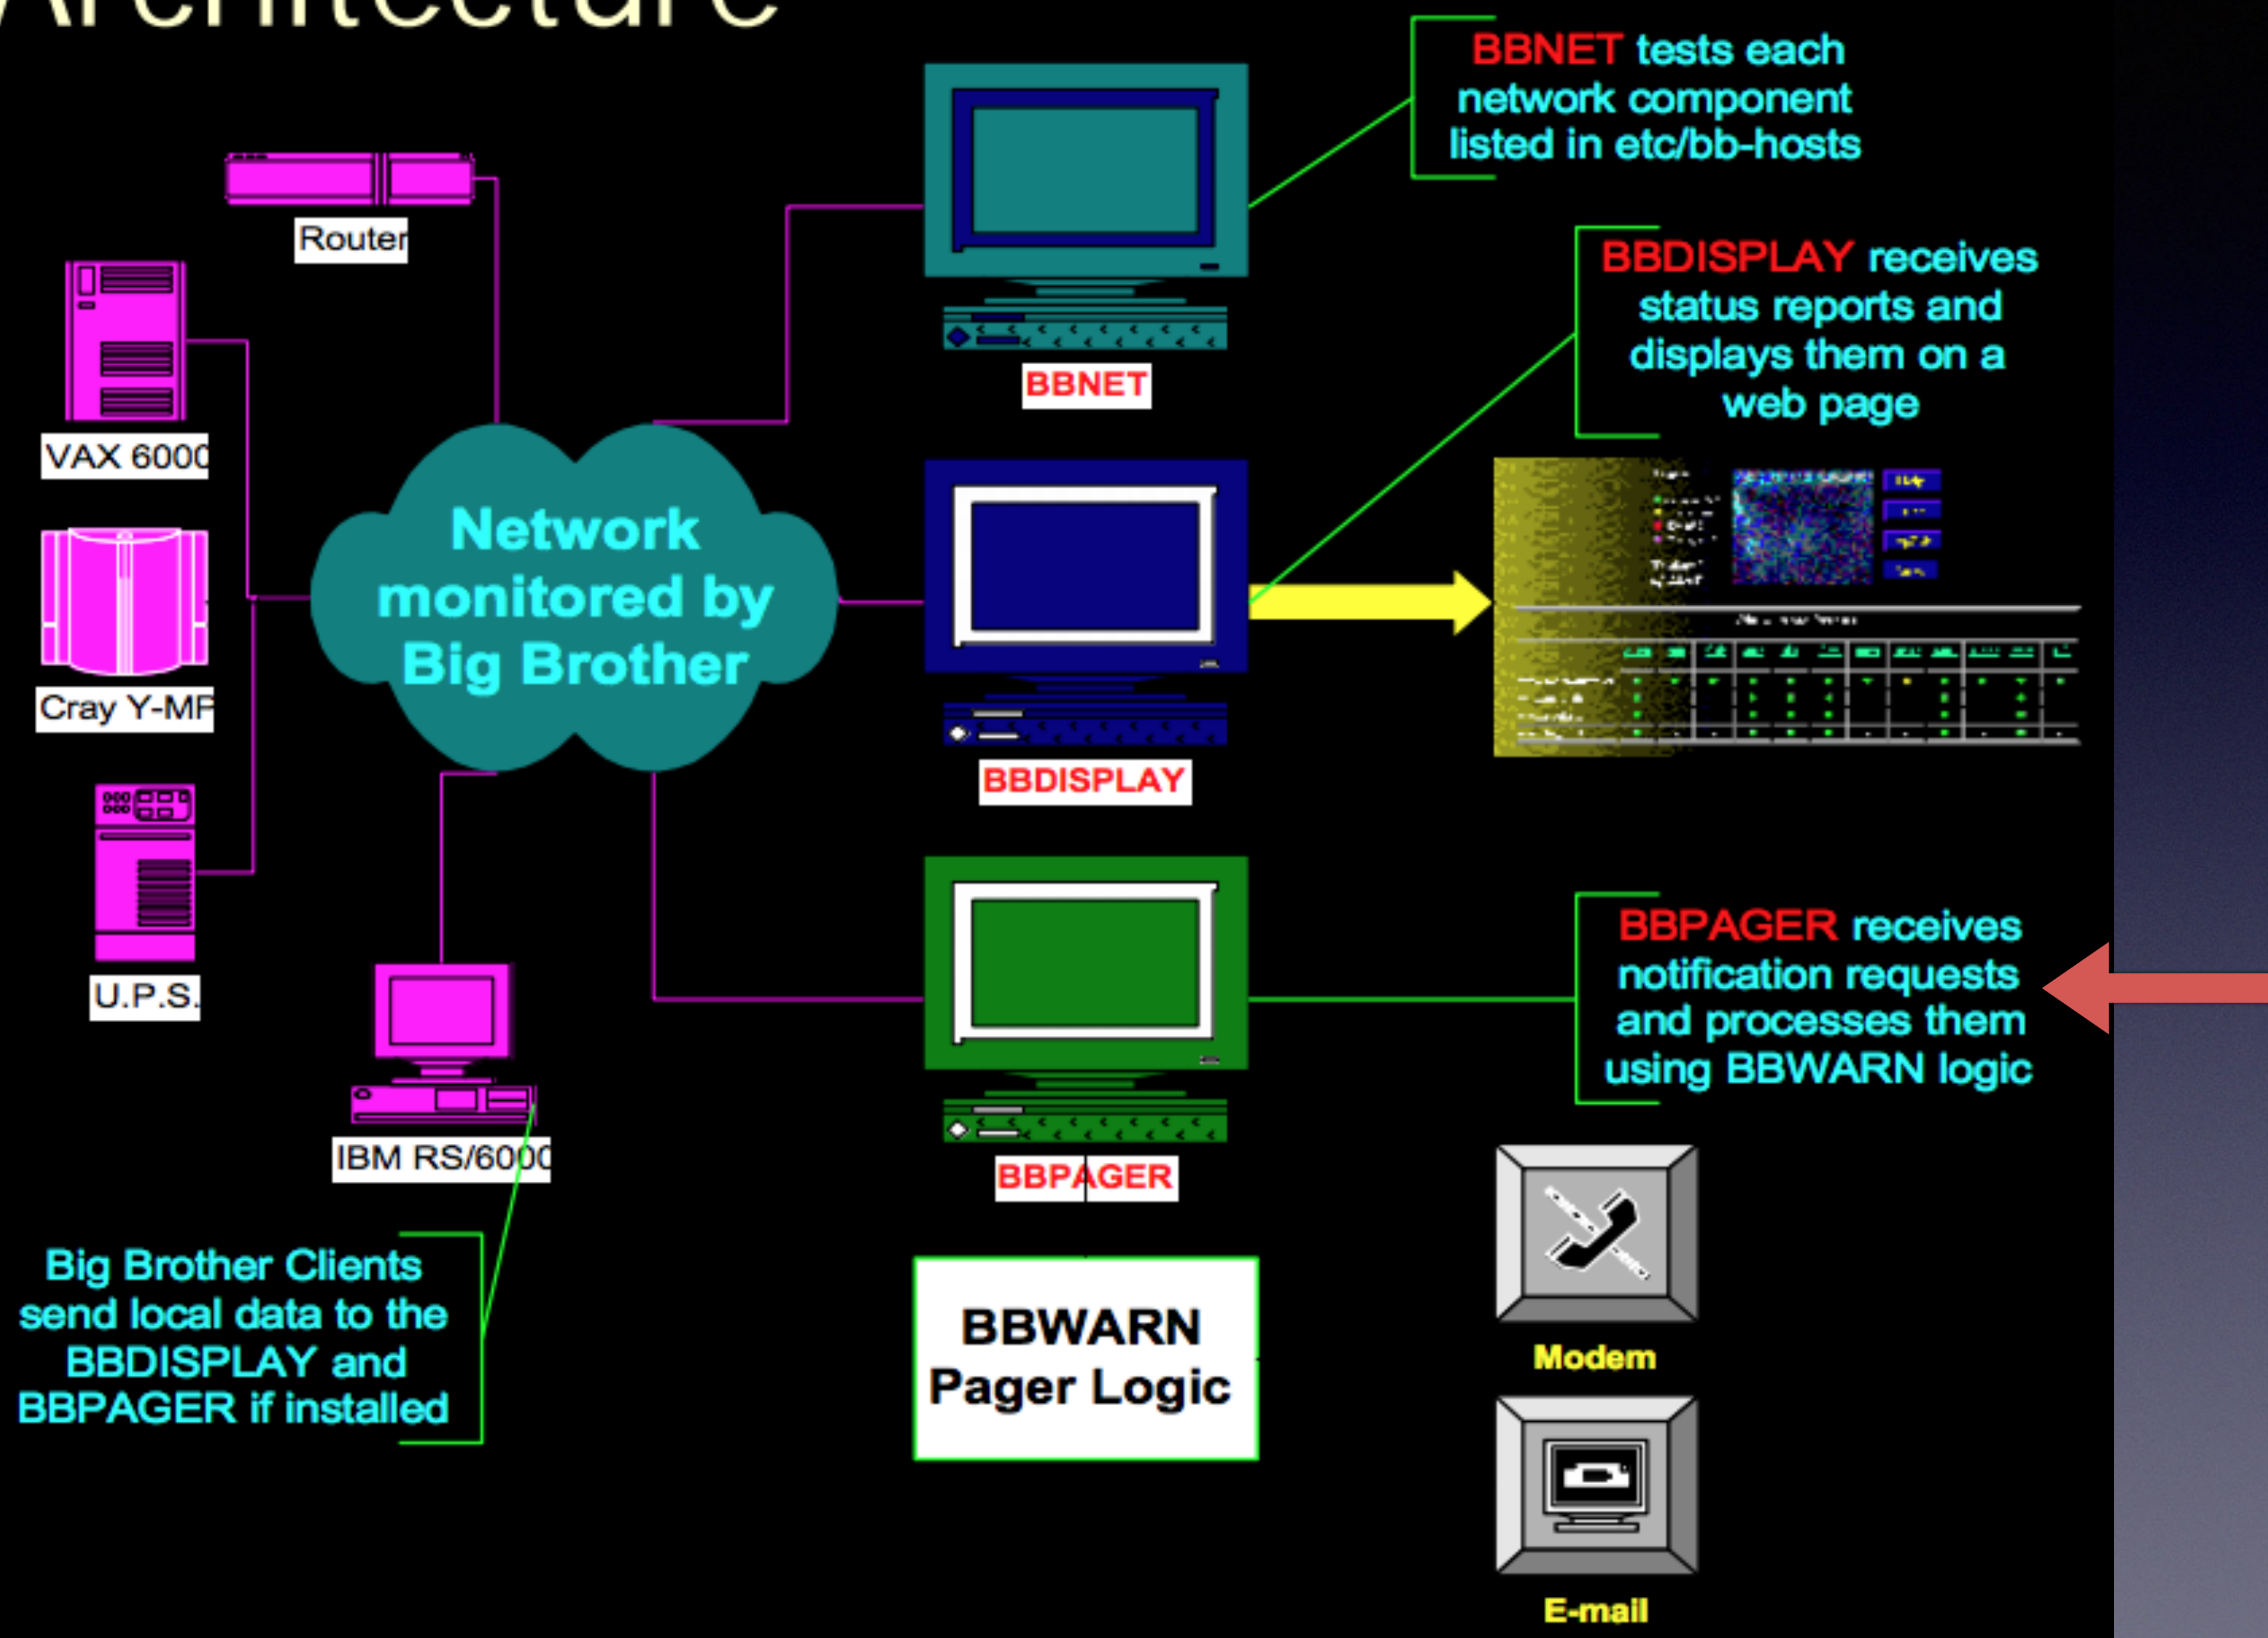
  

<i>Zoo net</i>	<u>conn</u>	<u>cont</u>	<u>cpu</u>	<u>disk</u>	<u>dns</u>	<u>ping</u>	<u>http</u>	<u>msgs</u>	<u>procs</u>
tower.bordercontrol.com	■	■	■	■	-	■	■	■	■
www.zoo.net	■	-	-	-	-	■	■	-	-
ns1.zoo.net	■	-	-	-	■	■	-	-	-
ns2.zoo.net	■	-	-	-	■	■	-	-	-
cougar.zoo.net	■	-	-	-	-	■	-	-	-
somehost.bb4.com	-	-	-	-	-	-	-	-	-
pm3a.zoo.net	■	-	-	-	-	■	-	-	-
hippo.zoo.net	■	-	-	-	-	■	-	-	-

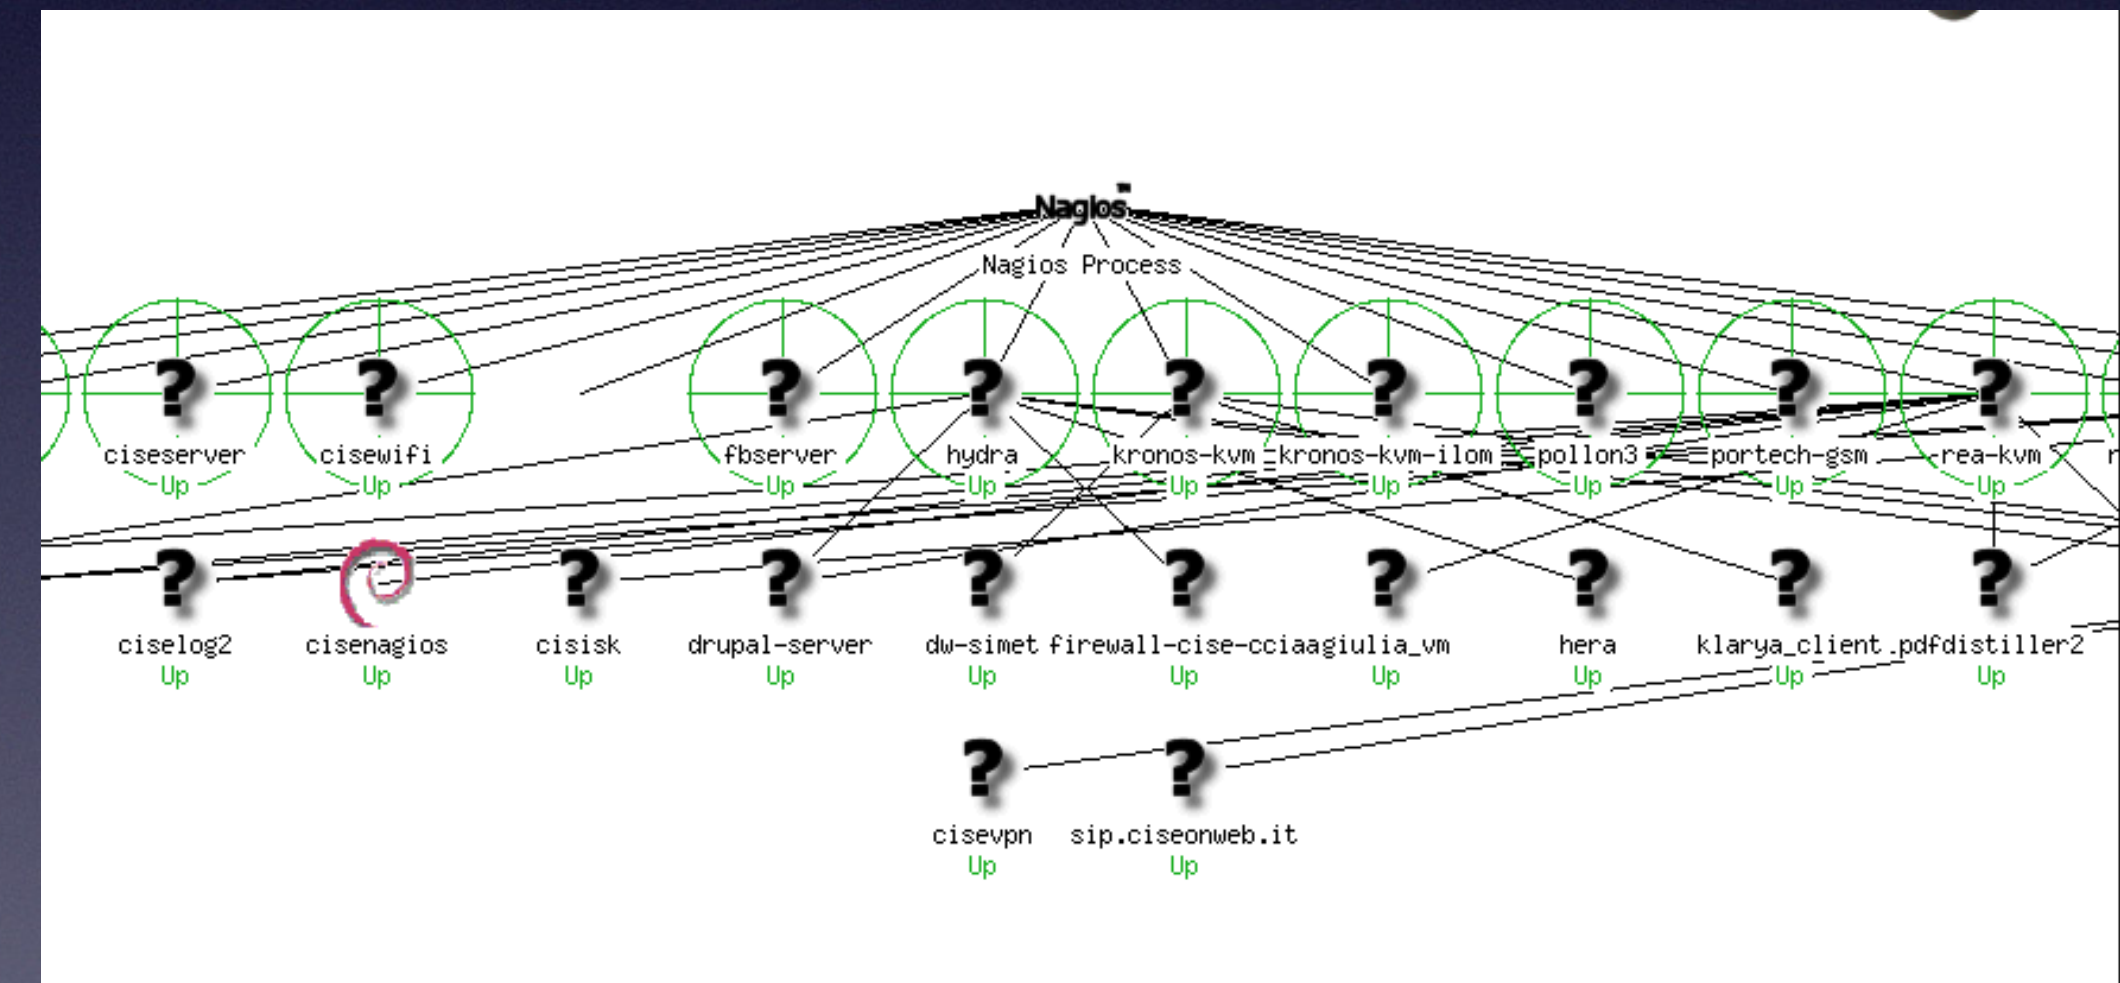
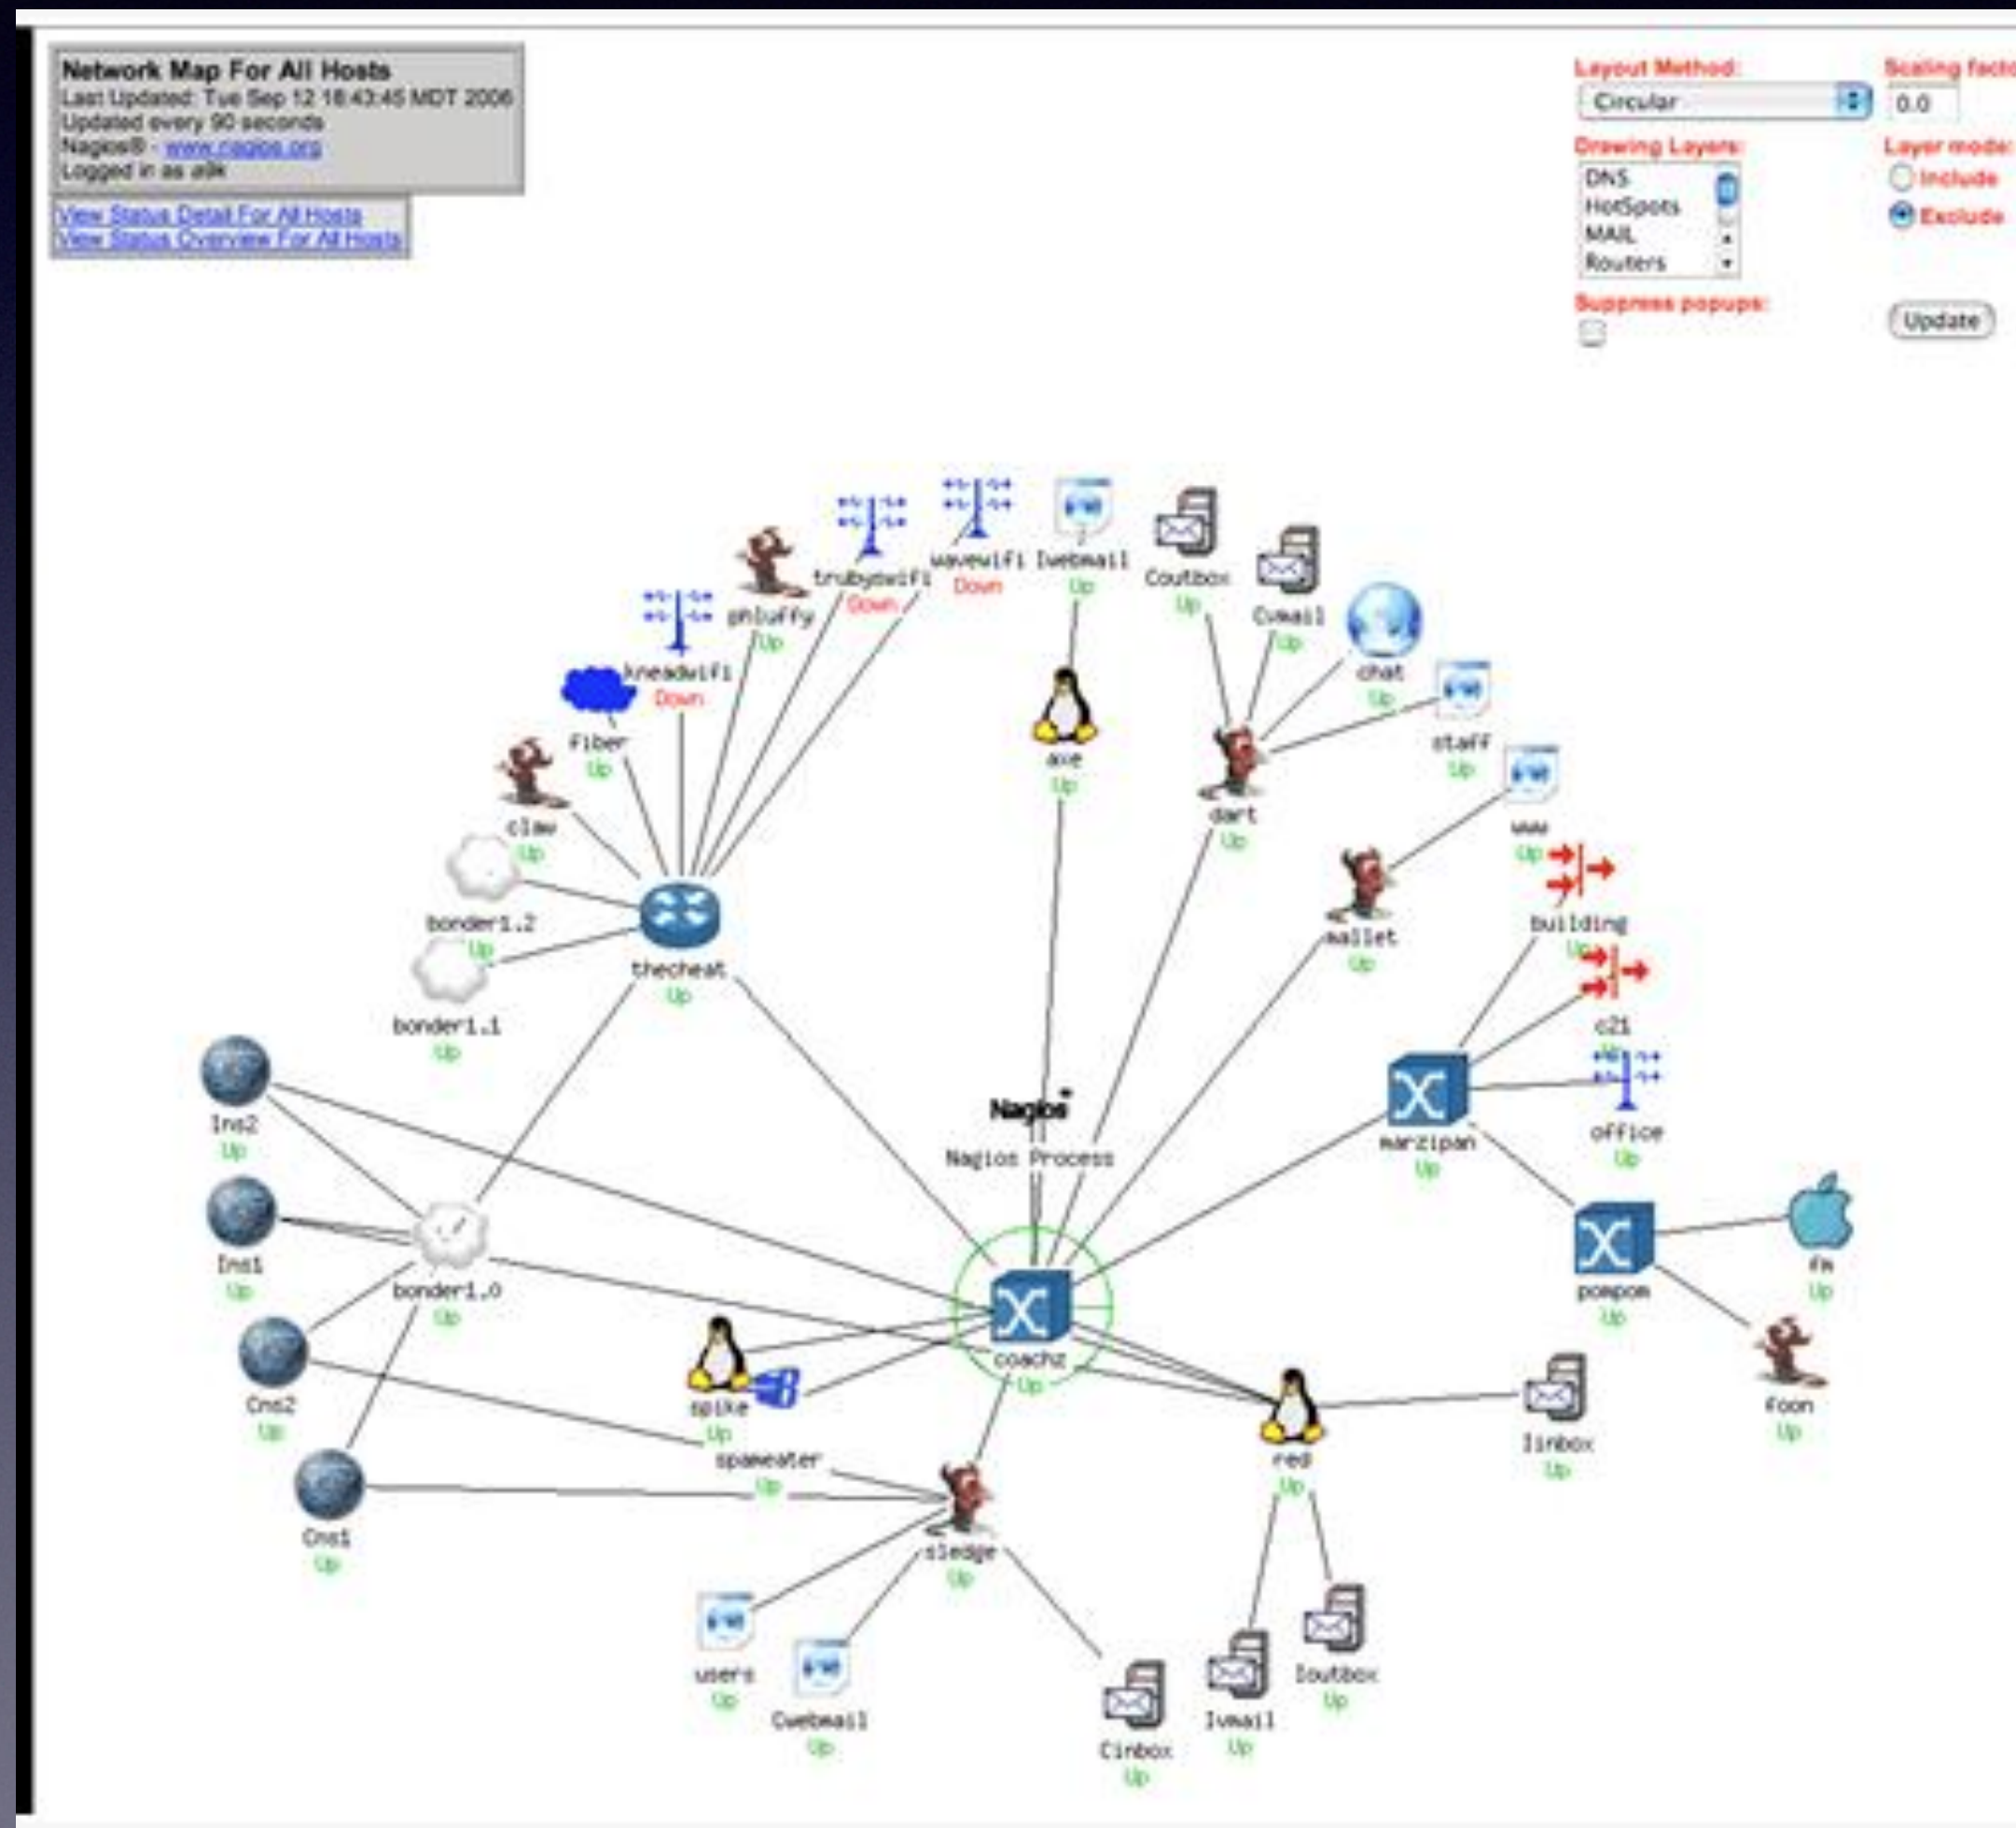
■ ok ■ attention ✗ trouble ■ no report ○ unavailable ● offline

Questionable Bash and Perl execed via Cron to feed Big Brother



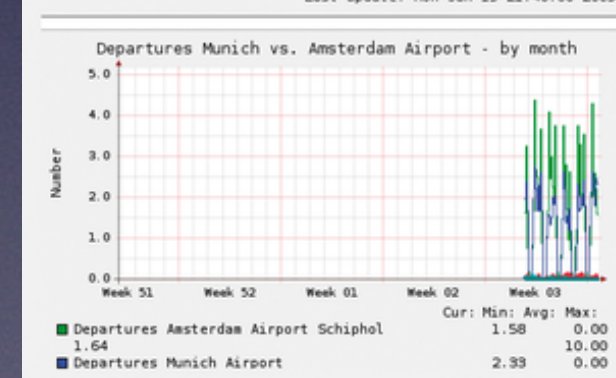



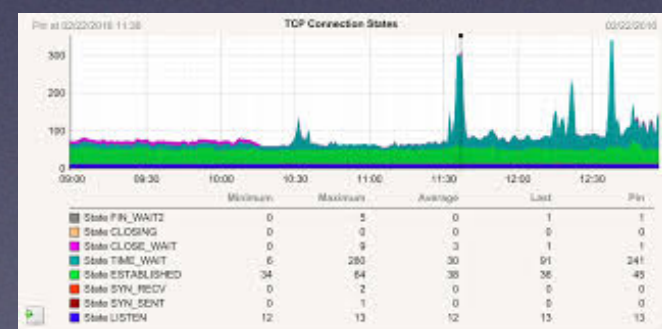




Architecture



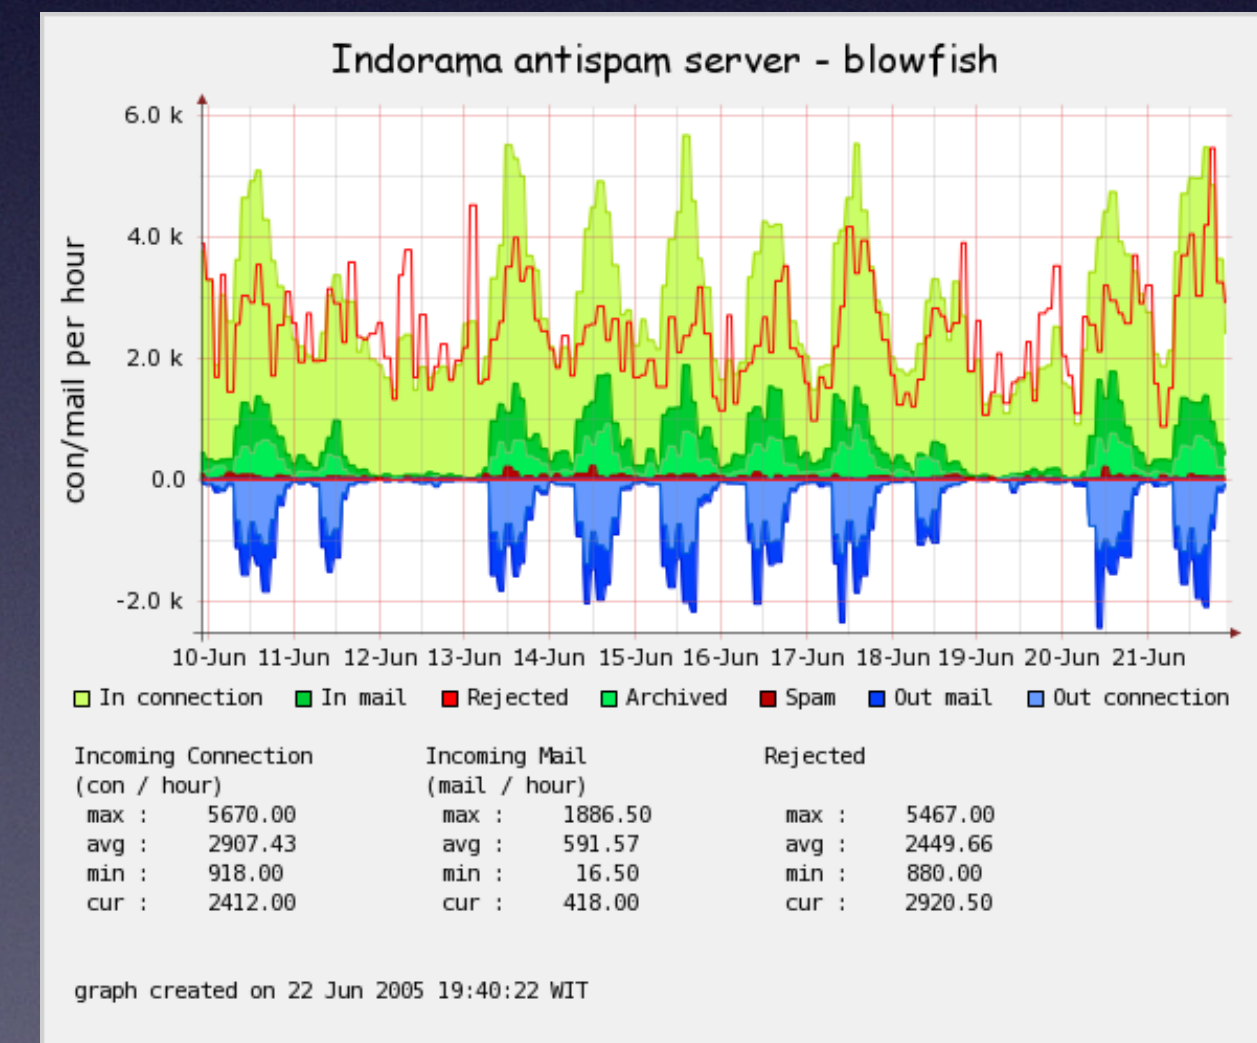

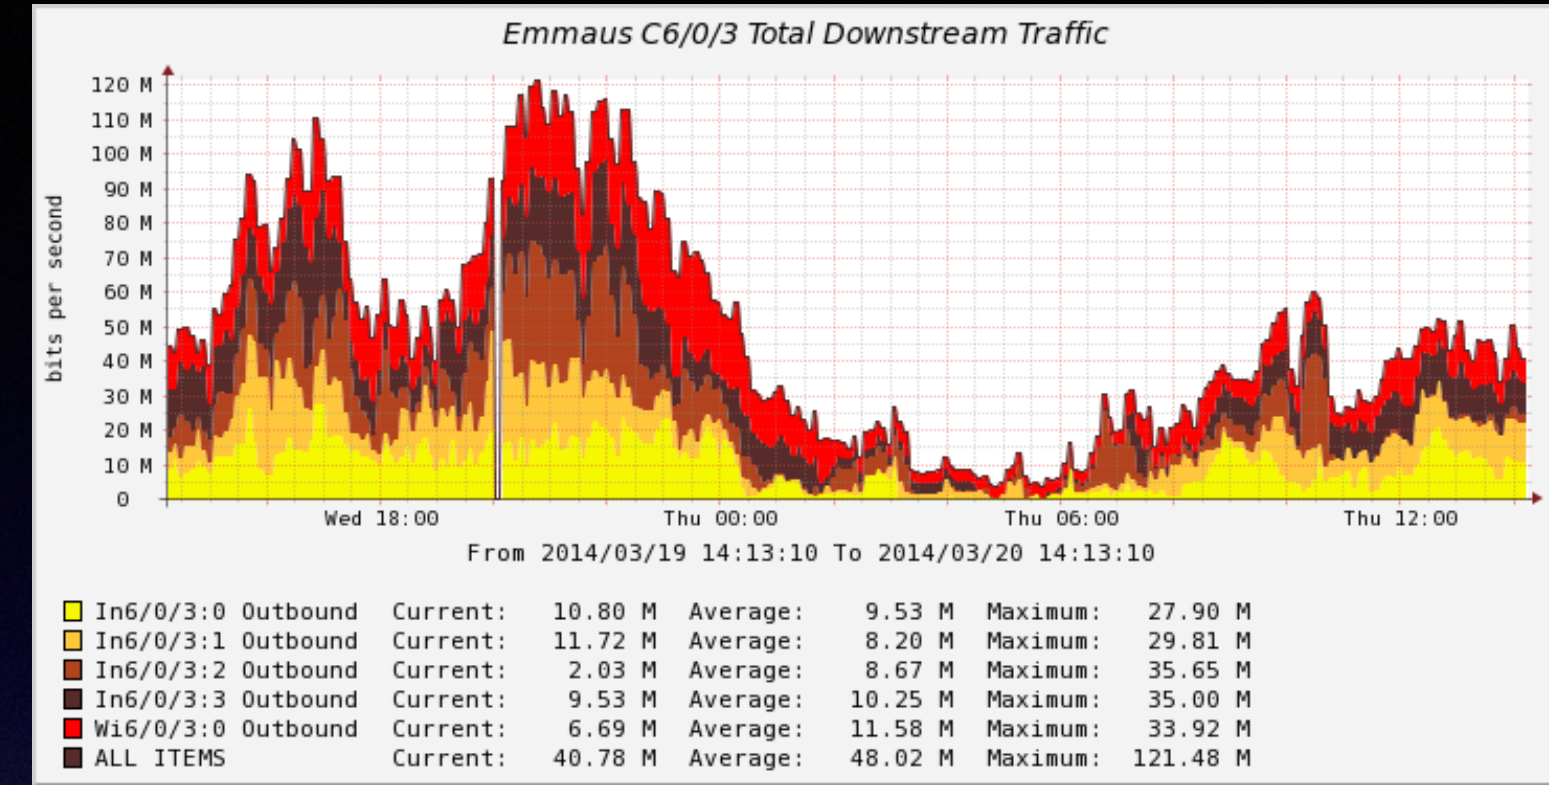

Do sysadmins have VRML colored dreams?



Metrics - Mostly RRDTool Based



Rub more Perl in it



Things improved with Cacti

Were our systems
observable?

Kind of?

Intuition

Limiting Factors

Cost of storage

Cost of compute

Grepping lots of logs

No TSDBs or NoSQL - Filesystems and RDBMSs

No context from previous failures except
gzipped collections of logs and MRTG graphs
maybe screenshots

Limiting Factors

Your sysadmins problem

Not quite configuration management

Awful technologies (SNMP)

Mostly statically generated assets from logs

Tools were macro world view

Sysadmin Squad Goals of Yore

Not letting engineers access production

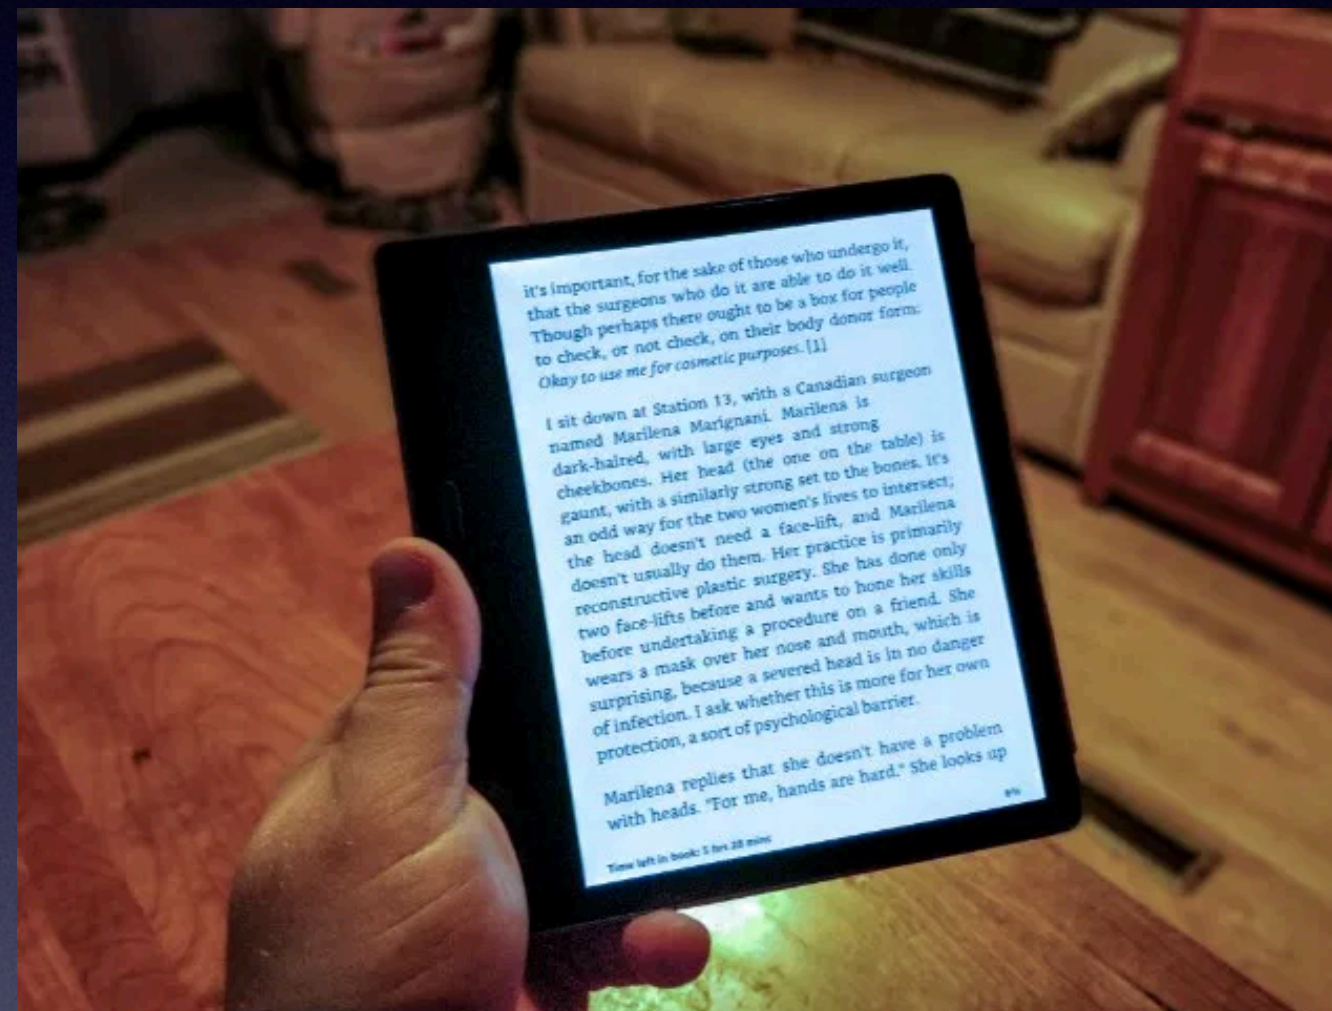
Having monitoring at all

Getting alerted reliably

Avoiding breaking things

Uptime, preferably **5 Nines**

Cross training, teaching intuition



The Present

SSD AWS Cloud Docker

Schedulers (Kubernetes) Grafana SaaS SRE

NoSQL Rest TSDBs APM

Azure 15 Years of Moore's Law DevOps

AIOT.

Log Management Containers Containers GCP

Open-source SDN VMs

Microservices CNCF Distributed Systems

Key Factors Driving Observability's Adoption

Democratization of technology (DBs, Log aggregation, VMs)

Tradeoffs are much different today

Monitoring is about predicting failure, harder to predict today due to the complexity

Clouds are ephemeral

Microservices

Key Factors Driving Observability's Adoption

Multi-Region / Multi-cloud concerns

Complexity has increased exponentially

Monitoring and metrics systems are far more advanced, less brittle, many times built in

Many more disparate systems / APIs (internal and external)

Scale has increased dramatically for even simple applications

4 Pillars of Observability

Metrics

Logging

Distributed Tracing

UX (Alerting/ Visualization)

Metrics

Choose your own adventure, there are so many to choose from, beware of cardinality concerns (don't tag metrics with UUIDs)

More predictable overhead than logs

Lower integration costs for applications than in the past, usually just a library or API call

You don't always know what metrics you need

Logs

Contextualizes the lifecycle of a requests

The least standardized component of most environments

Generation is the easy part, storing is not

Finite amount

Sampling for the long term if necessary

Distributed Tracing

Newest component to modern infrastructure

Correlation IDs are a good place to start

UI not required, but more valuable if you do

Not APM

Difficult to retrofit unless you use a service mesh

Requires engineering wide adoption

Alerting / Monitoring / UX

Alert fatigue, make actionable alerts, not just your SREs

Alerts must include context

The solution is never, simply, to make a alert and dashboard for X thing

Duplication and one-off dashboards

Manage the entire stack with configuration management

Current Groundwork

Engineers want APMish things

The rise of SaaS Infrastructure

Companies..... Scaling is Hard

Log aggregation (at scale) is easier: Humio,
ELK, SaaS

Many more options than the past

Current Groundwork

Open and closed source options abound

OpenTracing

Systems capable of high throughput and ability to query data that has high cardinality

Integration of disparate systems simpler

How can I make my
systems observable?

Start with Empathy for the User

The end goal is to make the best user experience possible

Users don't care about your CPU Load

Users care about errors and latency

All of your observability goals should be focused on answering questions about end user experience

How?

Gather as much data as you can from any path a user takes and align the context from disparate systems around individual users / geographies (CDNs, Load Balancers, etc)

Instrument your clients focusing on key experiences in your applications as well as basic interactions such as DNS

How?

Adopt a DevOps approach in your organization with observability concerns

Standardize your engineering organizations requirements for observability, give engineers a carrot

Your goals will be congruent with other departments such as business and marketing

Don't keep it a secret

New Approaches to Consider

AI / Machine Learning

Adopting stream processing infrastructure

Chaos Engineering

Testing in production

Sampling

Challenges

Percentiles are misleading at scale

Log amounts and sizes will increase

Scaling and spiky work loads

Complexity of tools when self managed

Many, especially edge, integrations are asynchronous

Challenges

Shoehorning into existing tools can lead to technical debt, substandard solutions

SaaS Cost at Scale

SaaS / Application Lock-in

Once SaaS cost is an issue, so is rolling your own

Not adopting standard approaches in your organization

Challenges

Intuition is still a powerful, necessary, tool

Uptime and MTTR are important, quality user experience encompasses both

Uptime is easier to come by than ever before

Best Practices

Not one size fits all, or even many, be mindful of your specific tradeoffs

Integrate systems, such as your CRM, Zendesk, etc for additional context

Balance cost and usage tradeoffs

Have your observability config live with your application code, deploy it with your CD system, bonus points for CI

Best Practices

Start small, one application end to end

Maximize context at every step of your request path

Be wary edge aggregation, you lose context

Canary deployments that are easy to see in your UX, handles real (your) user traffic

Developers on call



The Future

The next 10 Years

Cheaper, faster, better systems and clouds

Less complexity to the tooling

Less tool sprawl

Less SaaS

More real-time

Streaming native infrastructure, less HTTP APIs, log shipping

The next 10 Years

Observability Engineers and Orgs

Users are first class concerns for engineers

Sample on and off dynamically, by user?

Synchronous Edge Observability

Making everything observable, your team, your repos, etc will be commonplace

Building Your Observability Culture

It's never over, there are always new questions

Promote the power of the ethos outside of engineering

Use context from your business to inform your focus on what is important to observe

Your observability will provide business value, push it up and out

Building Your Observability Culture

Foster a culture of accountability, you own your uptime and errors, as an organization

Empower engineers to make it easy to input and utilize the data, everyone instruments

Make it easy for first time users to take actionable measures

Culture of empathy for users, and each other

Brian Knox, Digital Ocean:

The goal of an Observability team is not to collect logs, metrics, or traces. It is to build a culture of engineering based on facts and feedback, and then spread that culture within the broader organization.

Acknowledgements

The Infrastructure Team at Vevo

Humio @meethumio (<https://humio.com>)

Cindy Sridharan, Apple, @copyconstruct
(<https://medium.com/@copyconstruct>)

Charity Majors et. al, honeycomb.io
@mipsytipsy



Visit us at Kubecon EU at
Booth #SE54

Try Humio for Free
<https://humio.com/getting-started>
30 Day Free Trial - Free Cloud Tier

Thank you!

@schofield - grant@humio.com

Notes and Sources

<https://www.humio.com/chaos-observability>

<https://distributed-systems-observability-ebook.humio.com/>

<https://medium.com/humio/data-driven-observability-logs-52e98e27a83b>

<https://distributed-systems-observability-ebook.humio.com/>

<https://www.honeycomb.io/resources/white-papers/>

RRDTool <https://calomel.org/rrdtool.html>

Cacti - Many google image searches

https://en.wikipedia.org/wiki/Control_theory

<https://www.vividcortex.com/blog/monitoring-isnt-observability>