copenhagen goto; conference

A Leap from Agile to DevOps

Jeffrey Fredrick
TIM Group

INTERNATIONAL
SOFTWARE DEVELOPMENT
CONFERENCE

gotocon.com

TIM GROUP

UNLOCK YOUR ALPHA.

TIM IDEAS

The most comprehensive global alpha capture network for trade idea distribution, evaluation and management

BUYSIDE

Identify exclusive, market-beating opportunities quickly and manage brokers wisely.

SELLSIDE

Establish your value, broaden your network of connections and increase commissions.

TIM FUNDS

Superior analytics and portfolio management for funds of hedge funds directly or working with administrators and custodians.

OVERVIEW

TIM Funds transforms the data that administrators and custodians already hold into powerful, accurate, analytics.

CONNECT

TIM Funds Connect is a series of APIs that enable users to interact seamlessly with in-house or third party technology or data systems.

2005

TIM IDEAS

The most comprehensive global alpha capture network for trade idea distribution, evaluation and management

BUYSIDE

Identify exclusive, market-beating opportunities quickly and manage brokers wisely.

SELLSIDE

Establish your value, broaden your network of connections and increase commissions.

2006

TIM FUNDS

Superior analytics and portfolio management for funds of hedge funds directly or working with administrators and custodians.

OVERVIEW

TIM Funds transforms the data that administrators and custodians already hold into powerful, accurate, analytics.

CONNECT

TIM Funds Connect is a series of APIs that enable users to interact seamlessly with in-house or third party technology or data systems.

2007 2008 2009

- Adopt pair programming
- Add extensive testing including TDD
- Support with elaborate continuous integration

2005 2006

2010 2011

- Managed hosting limiting evolution of architecture
- Bring in consultants, create migration plan
- Hire sysadmin team, build collocation sites

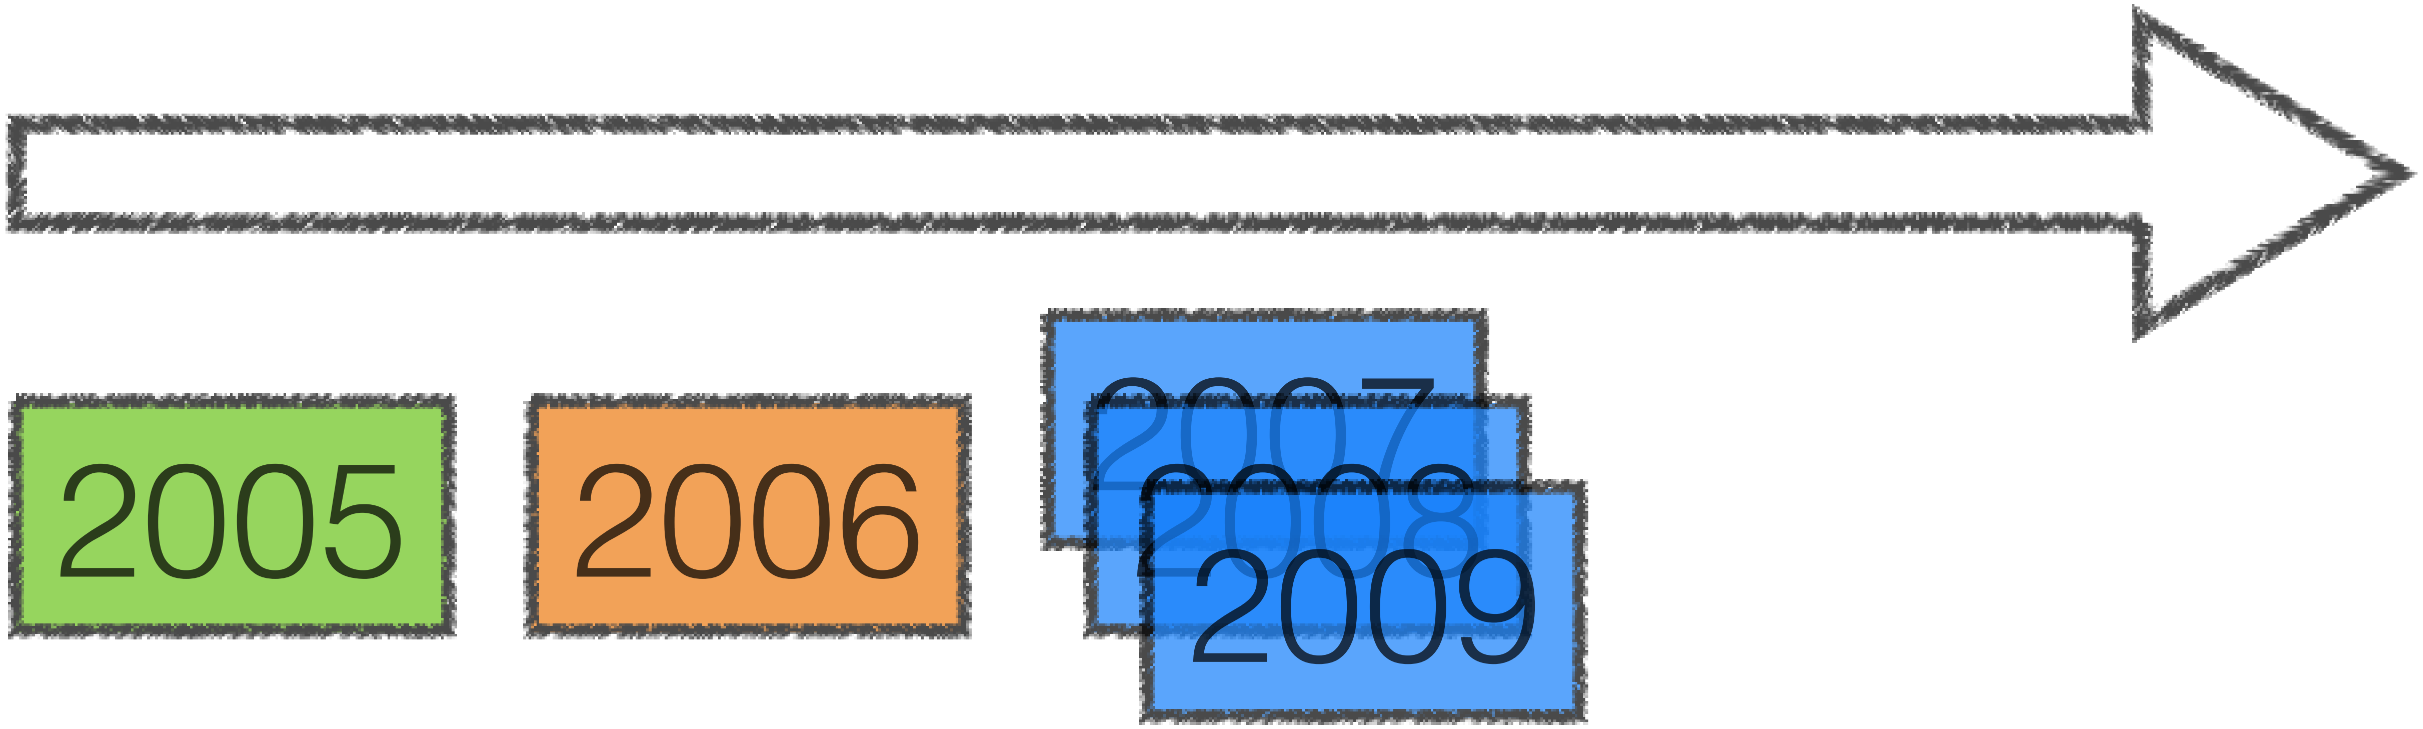

2012

- At migration complete merged sysadmin project team with Ops and Tools team to create new Infrastructure team

- Went live from our new data centers in January

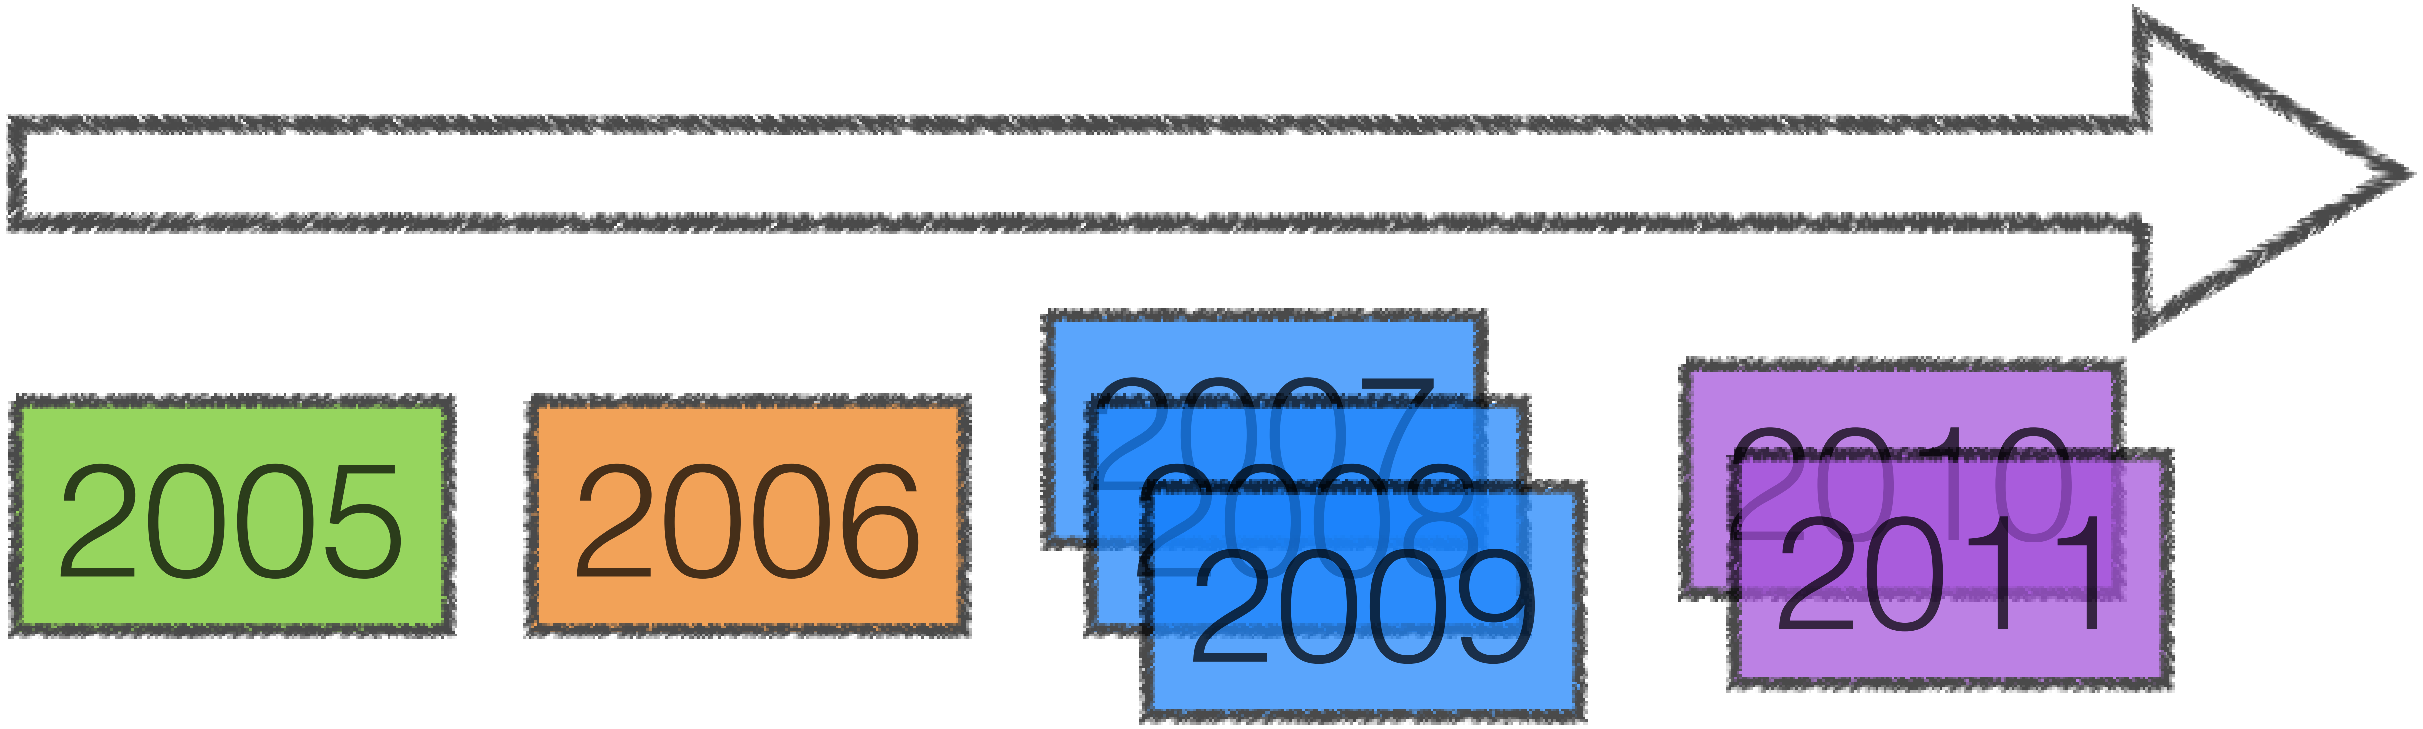

2005

2006

2007

2008

2009

2010

2011

- 20May12 3:18am: pg-timapp-002 TIM Tomcat Direct alert, resolved itself within 5 minutes. OoH effort: Bob: 5m.
- 18May12 8:15am: portal on macau alerting because it had hit the open files limit. Restarted portal to resolve the immediate issue. OoH effort: Bob: 45m → passed to Kei/Mehul.
- 18May12 5:20am: mail queue on pg-timapp-005 over threshold, 78/50 messages in queue. Resolved itself. OoH effort: Bob: 10m
- 17May12: 1:18am: Nagios alerts for portal returning http 500. Investigated, restarted macau. OoH effort: Bob: 1h.
- *18May12 8:00am: installing MBA Test Environment TC certificate: OoH effort: Kei: ?*
- 15May12 3:50am: HK office reported problems with delays in pricing for Japanese/Korean stocks→ended up being a TRKD problem. OoH effort: Bob: 1h, Jason: 30m
- *14May12 8:00pm: MBA TC certificate cutover: OoH effort: Kei: 1*
- *12May12 8:00am: Standard TIM Funds and TIM Ideas release: OoH effort: Kei: 1*
- 11May12: 0059: TRKD down. Resolved itself. Adam: ?
- 9May12: 0014: ApiRankingCyclicProcess failure. Resolved itself. Adam: 30m.
- 8May12: 0400: TradeRepricerOfOddSiblingEventsCyclicProcess commit failure: StaleStateException. Resolved itself. Adam: 30m.
- 7May12: 0850: Ideas Stop Loss cyclic process failure due to StaleObjectStateException Jtf: 15 min
- 6May12: 0640: TF Log Check failure due to Fx CP. Jtf: ? Tony: ? Isaac: ? Graham A.: ? Ian: ? JoeS: 1h
- *5May12 ?: Scheduled Maintenence* Jtf: ?
- 5May12 01:00, 02:40, 06:30 : Merc and TF FX CP failures. Kei: ?
- *4May12 1700: prep for Scheduled Maintenence. Jtf: 0.5*
- 28-29Apr12 ?: Ideas upgrade process broke replication. Bob: 10h Gary: 2h Adam: Kei: Jeff: Waseem: Joel: Jason:
- *28Apr12 1600: Release of Funds & Ideas. Long release because trying to run both simultaneously botched the Ideas release. Jtf: 2.5 hrs, Waz: ? Kei: ?*
- 27Apr12 05:55pm: Gavin enabled CS restricted lists entitlement and wanted us to check if it was running. Bob: 30m
- 25Apr12 05:07am: TradeRepricerOfOddSiblingEventsCyclicProcess Failed: StaleStateException. Resolved itself. Bob: 30m
- 22Apr22 ?: TIM Connect failures, turned out to be remote end, not us. Adam: ? Waz: ? Jeff: ?
- 21Apr12 16:13: ExchangeRateUploader failure (Could not open a connection to SQL Server). Resolved itself. Adam: 30m
- 20Apr12 08:10am: High number of processes alert on proxy server. Adam: 2m
- 19Apr12 04:45am: PG TIM FW failure leading to lost packets and TIM outage. Adam: 1h
- *17Apr12 07:00pm: TIM Funds maintenance release: OoH effort: Kei: 0.5*
- *15Apr12 09:50am: TIM Ideas healthcheck: OoH effort: Kei: 1*
- *14Apr12 7:00am: TIM Ideas database archiving: OoH effort: Kei: 4*
- *14Apr12 7:00am: Standard TIM Funds and TIM Ideas release: OoH effort: Kei: 2*
- *14Apr12 21:00: Brixton upgrade OoH Effort: Richard: 2*
- *10Apr12 10:30pm: TIM Ideas maintenance release : OoH effort: Kei: 0.5*
- 9Apr12 Midnight: TIMConnect outages. OoH effort: Bob: 4h Adam: 1h Kei: 4 Jeff: 4 Jason: 2h Gavin: 2h
- 4Apr12 00:56: Cyclic Process failure for hourly idea rankings. OoH effort: Bob 4h (+ Brian Roberts 30m, Gavin 1h)
- 31Mar12 04:45: TradeRepricerOfEvenSiblingEventsCyclicProcess failure. Resolved itself. OoH effort: Adam: 10m
- *30Mar12 8:00pm: TIM Funds maintenance release : OoH effort: Kei: 0.5*
- 29Mar12 04:48: TradeRepricerOfOddSiblingEventsCyclicProcess failure. Resolved itself. OoH effort: Adam: 10m
- *25Mar12 9:00am: Standard TIM Funds and TIM Ideas healthcheck: OoH effort: Kei: 1*
- *24Mar12 07:00: Standard TIM Funds and TIM Ideas releases + other maintenance: OoH effort: Kei: 3 GaryR: 2*
- 20Mar12 07:00: Abacus outage: OoH effort: Kei: 1
- 19Mar12 02:28: Today returns generation issue: OoH effort: Kei: 5h
- 19Mar12 00:54: Today returns generation issue: OoH effort: Adam: 4h
- *15Mar12 22:30: TIM Ideas maintenance release: OoH effort: Kei: 0.5*
- *14Mar12 22:30: bounce TC/tomcat for Marshall Wace: OoH effort: Kei: 0.5*
- *11Mar12 9:00am: Standard TIM Funds and TIM Ideas healthcheck: OoH effort: Kei: 1*
- *10Mar12 9:00am: Standard TIM Funds and TIM Ideas releases: OoH effort: Kei: 1*
- 06Mar12 1730: Email backups failing due to wrong VMware disk size. Implemented recovery solution after discussing with Paul Howard/Department 7. OoH effort: Craig + Richard: 2.5h
- 29Feb12 11pm: bos-f17-cloud-2 (fcsentiment + shortel) same issue as bos-f17-cloud-1. OoH effort: Bob: 0.5h
- 28Feb12 12am: bos-f17-cloud-1 (fcdomain + others) 208+ day uptime bug (cpu soft lockup/"divide by zero"). OoH effort: Bob: 1.5h
- *25Feb12 8:30am: Standard TIM Funds and TIM Ideas releases: OoH effort: Kei: 0.5*
- 20Feb12 0530: broken replication on dominica... denmark running with old data. took site down. restore backup from tuesday and replayed binlog. Lost data since Saturday. OoH effort: Adam: 3h30m Gary: Kei:
- 19Feb12 2300 : broken replication on dominica. OoH effort: Adam: 1h
- 18Feb12 2001: broken replication on dominica. OoH effort: Adam: 2h
- 18Feb12 1311: Merc problem with interday pricing. Reuters IP addressed changed, making it unreachable due to our firewall rules. OoH effort: Adam: 1h30m
- *18Feb12 08:00am: TIM Funds fail back to Powergate : OoH effort: Jarek: tbd , Kei: tbd*
- *16Feb12 10:30pm: TC feature release : OoH effort: Kei: 0.5*
- 12Feb12 8:30am: TIM Funds production database server (Denmark) went down. After receiving alerts for an hour ForLinux called the Support number and reached Jarek. Denmark was not reachable over iLo (why?) and Kei followed the documented directions and switched Funds to DR site. Gary was able to connect to iLo and reboot Denmark but it did not come up cleanly. Jarek fixed some partition stuff (clean this part up) and Denmark recovered. Though we did not have sync-binlog enabled (fix this?) the positions matched and Jarek was able to start replicating from OY back to Denmark. Other work: allowing mail out from OY; changing Nagios alerts. Root Cause: ? Outage duration: 2.5 hrs, 0800-1030 OoH effort: Jarek: 4, Kei: 1, Gary: , Craig: 1
- *11Feb12 9:00am: Standard TIM Funds and TIM Ideas releases: OoH effort: Kei: 3*
- *10Feb12 10:00pm: Scheduled SugarCRM and Drupal upgrades. OoH effort: Craig 1.5*

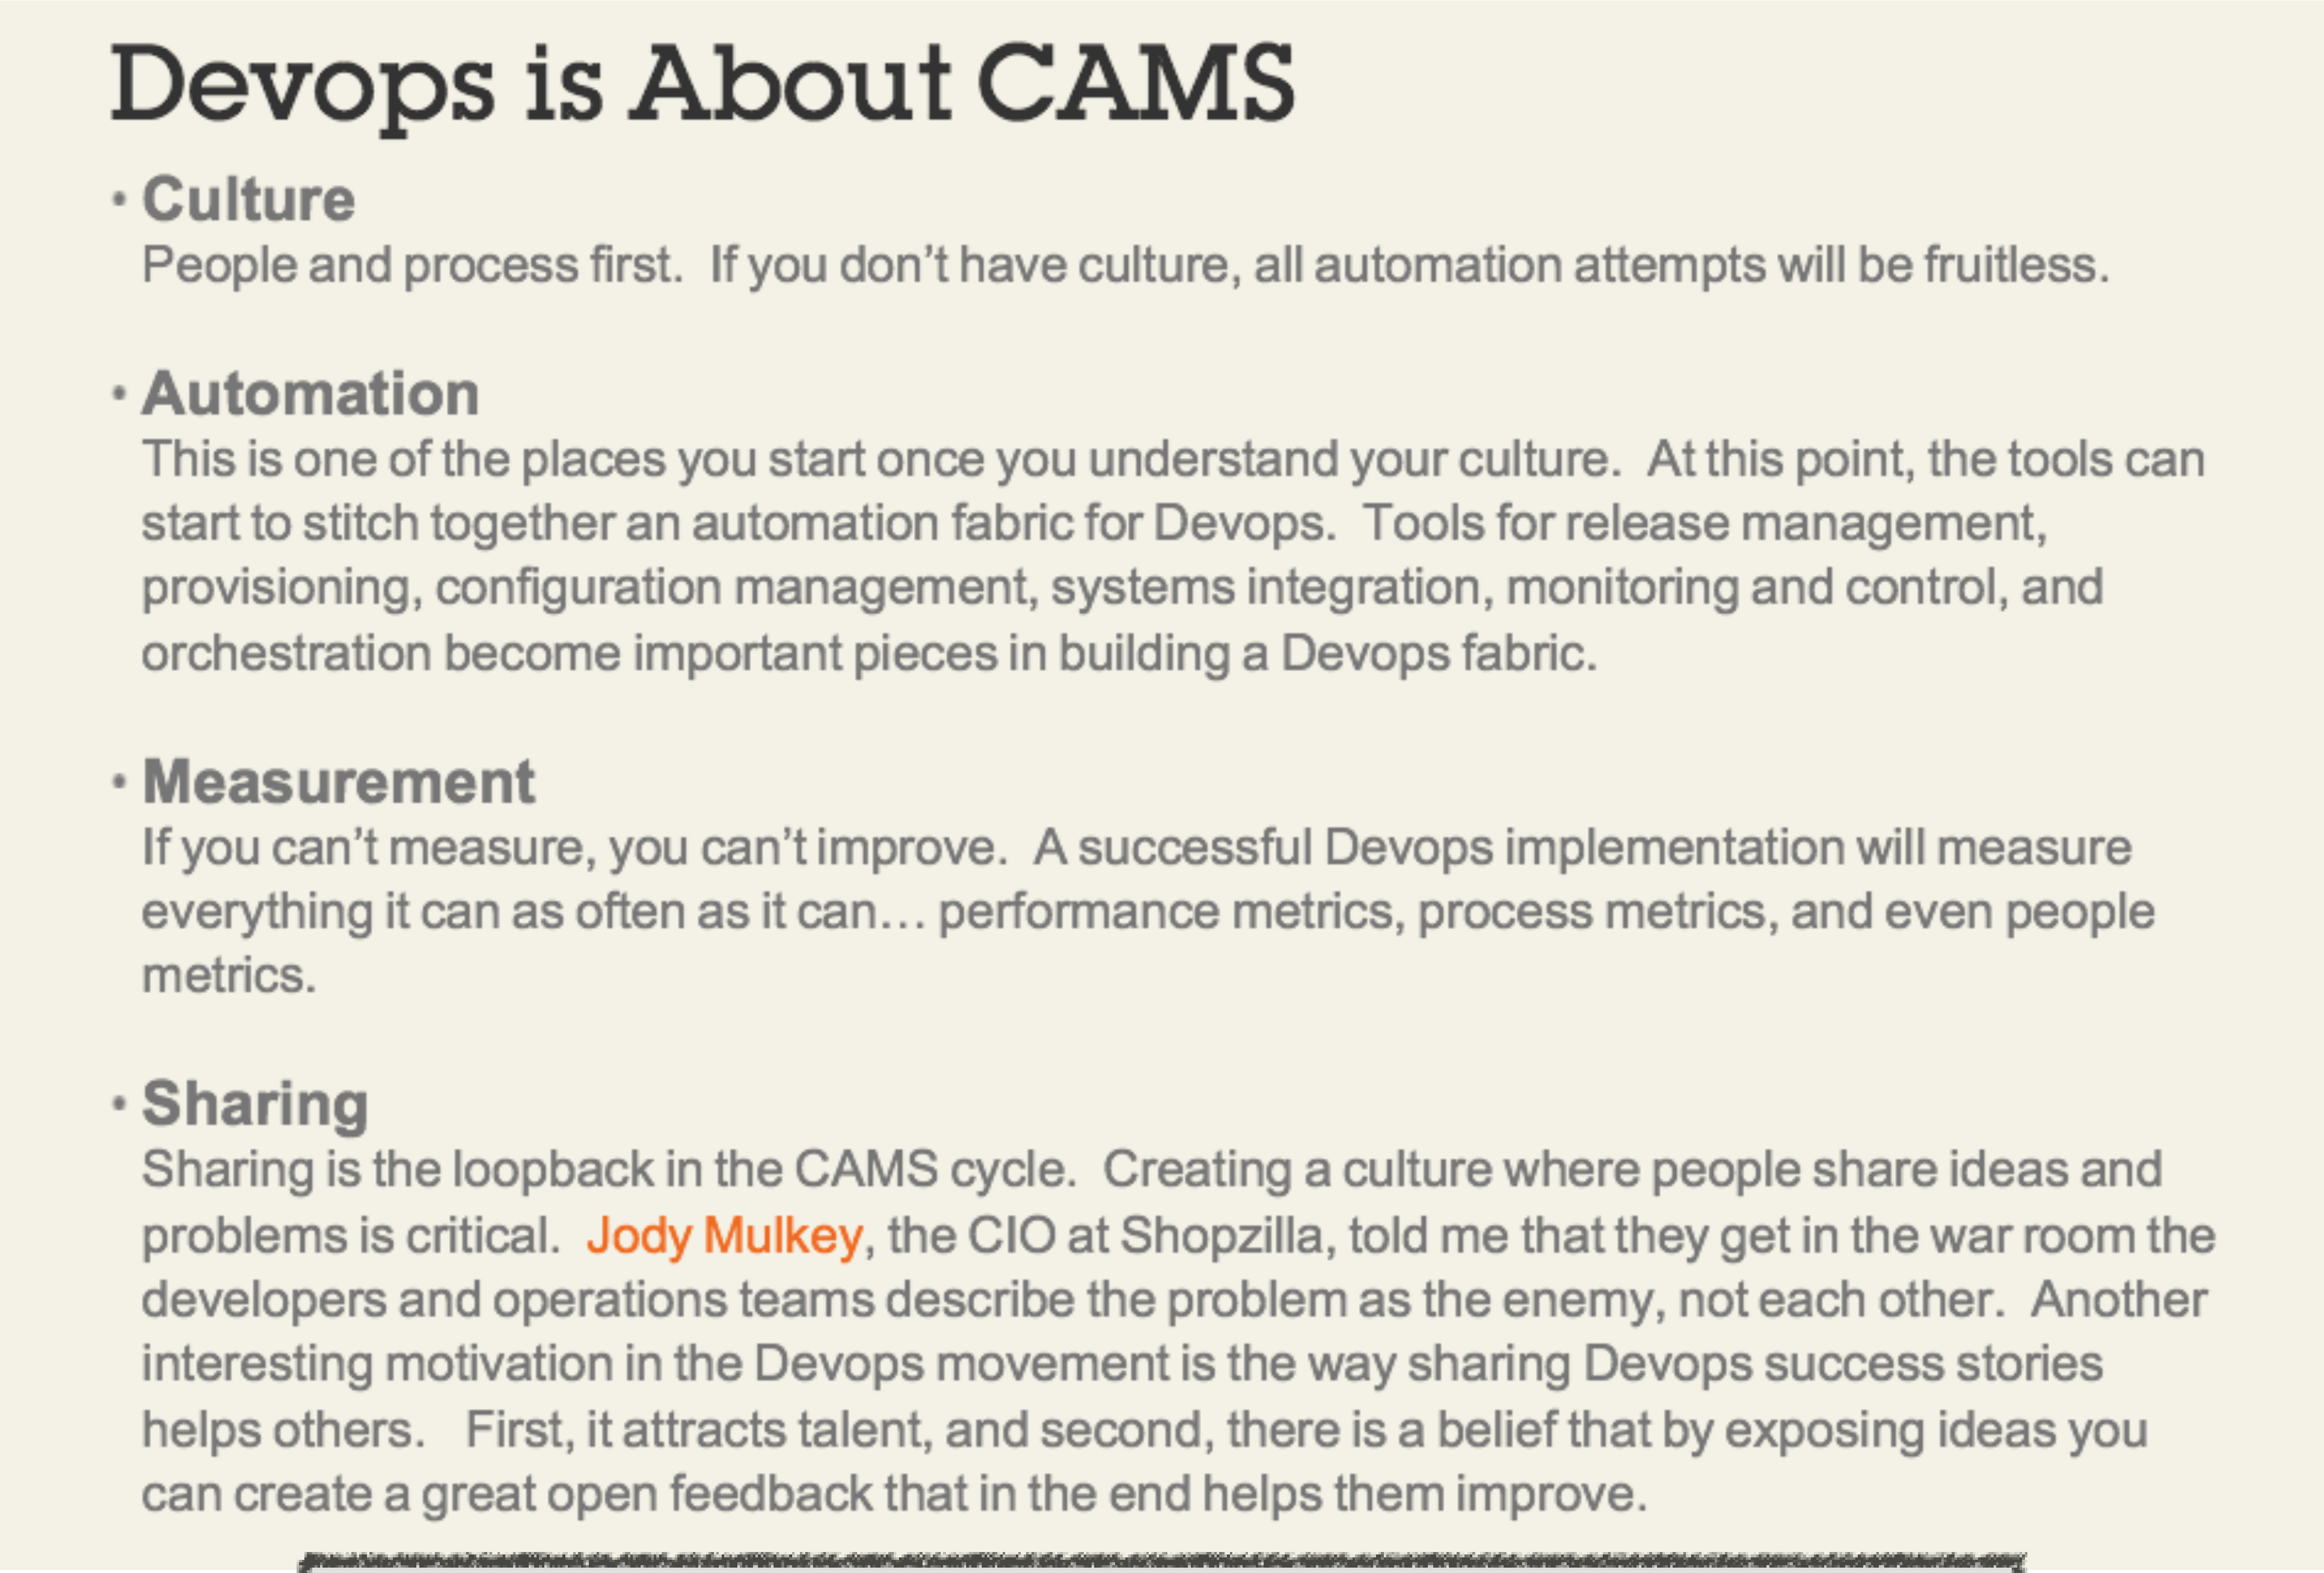

Devops is About CAMS

- **Culture**
  People and process first. If you don't have culture, all automation attempts will be fruitless.

- **Automation**
  This is one of the places you start once you understand your culture. At this point, the tools can start to stitch together an automation fabric for Devops. Tools for release management, provisioning, configuration management, systems integration, monitoring and control, and orchestration become important pieces in building a Devops fabric.

- **Measurement**
  If you can't measure, you can't improve. A successful Devops implementation will measure everything it can as often as it can… performance metrics, process metrics, and even people metrics.

- **Sharing**
  Sharing is the loopback in the CAMS cycle. Creating a culture where people share ideas and problems is critical. Jody Mulkey, the CIO at Shopzilla, told me that they get in the war room the developers and operations teams describe the problem as the enemy, not each other. Another interesting motivation in the Devops movement is the way sharing Devops success stories helps others. First, it attracts talent, and second, there is a belief that by exposing ideas you can create a great open feedback that in the end helps them improve.

John Willis: What Devops Means to Me
http://is.gd/B6z7E7

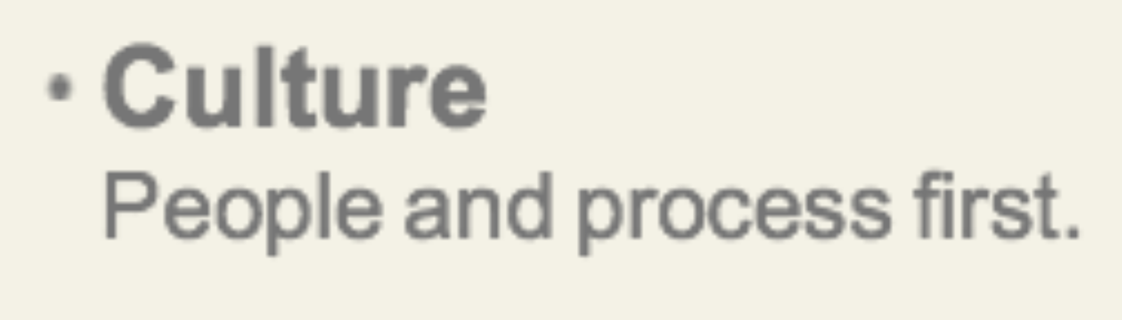

- **Culture**
  People and process first.

- Ops was application support, the semi-technical buffer between production & users and the developers
- Company bottleneck seen as developer time
- Enabling architectural change was the motivating factor

John Willis: What Devops Means to Me
http://is.gd/B6z7E7

· Automation

- Continuous deployment for new components put in place by development

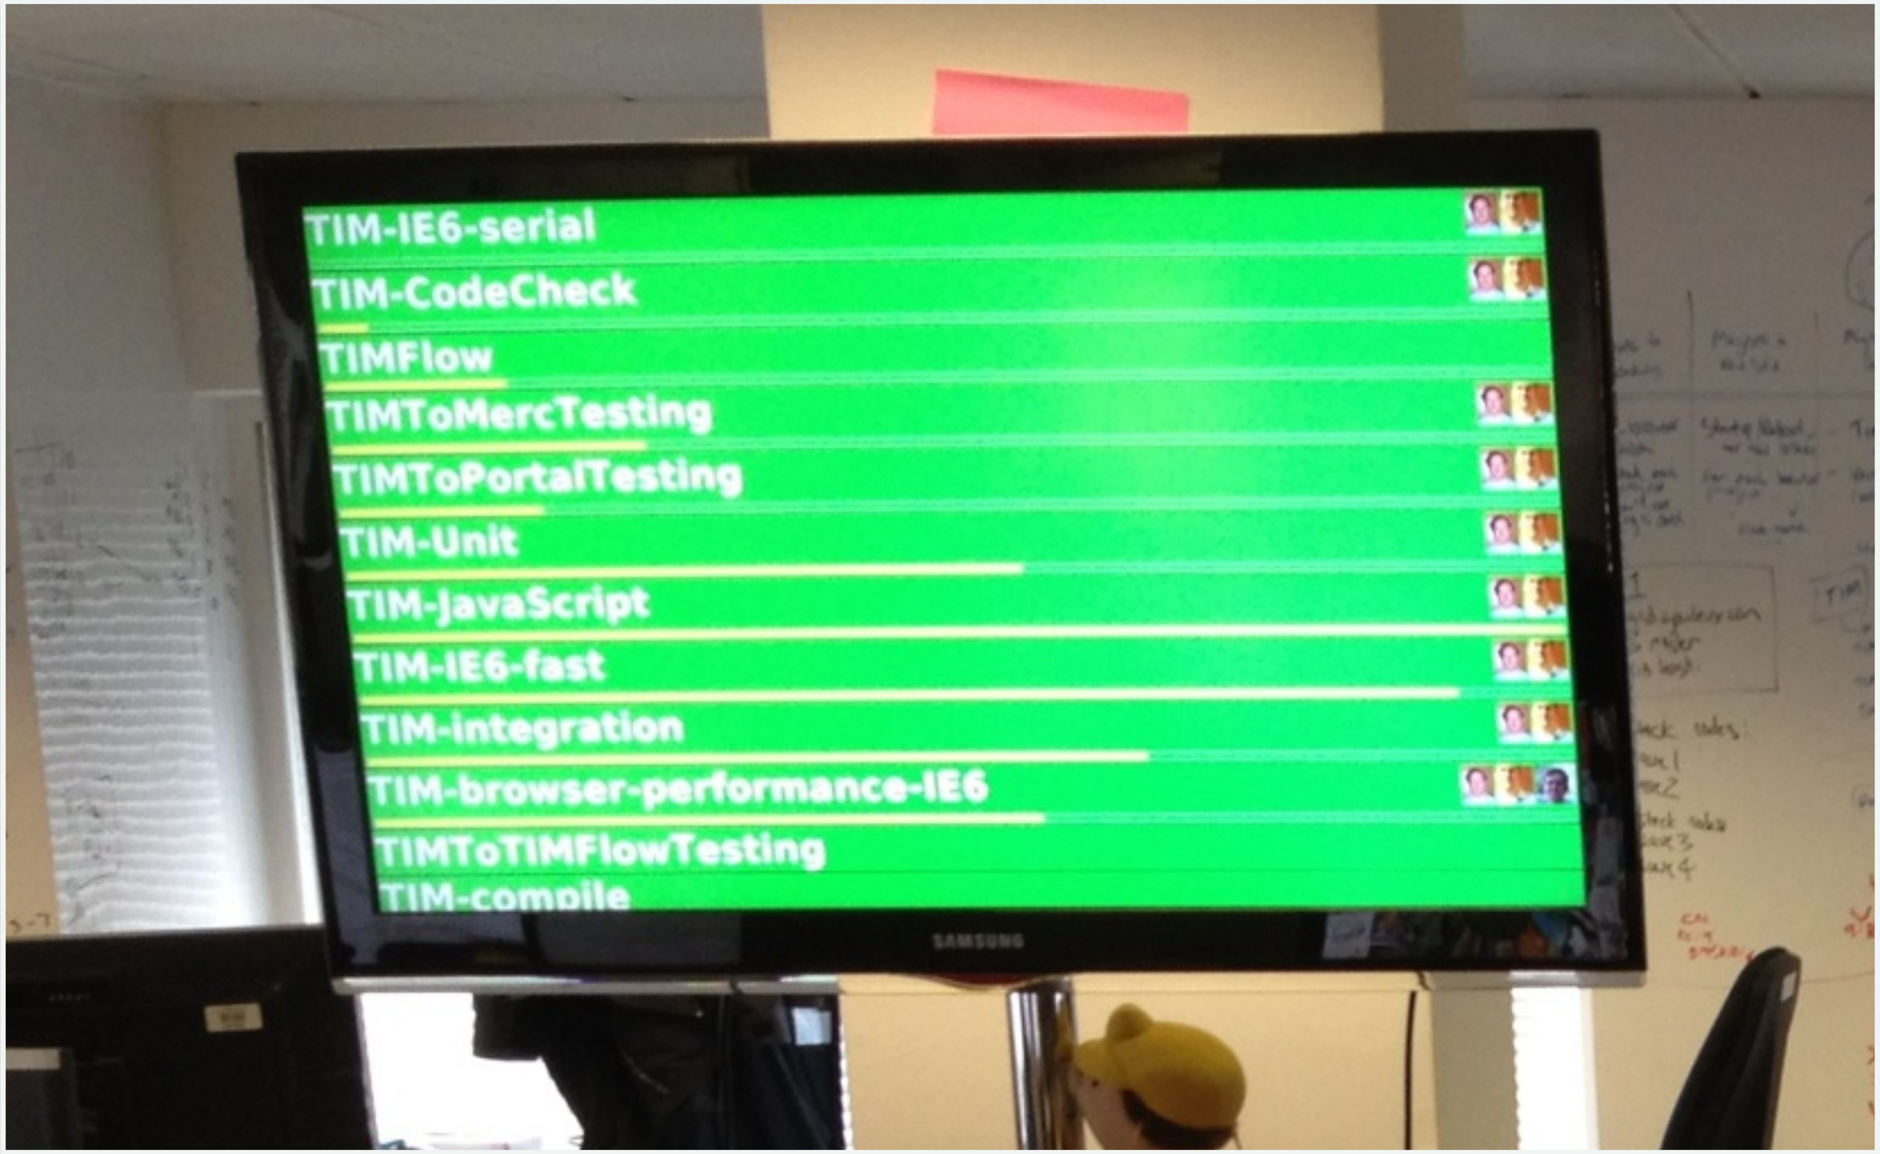

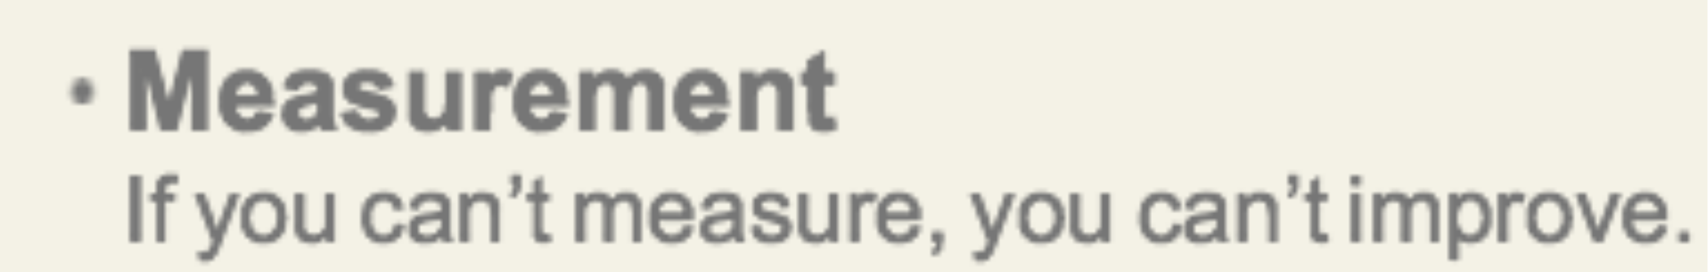

- **Measurement**

  If you can’t measure, you can’t improve.

- "Metrics": a two week block on the Gantt chart
- Application metrics are (almost) all technical
- Business metrics come from the same nightly "ops project" scripts as before the transition

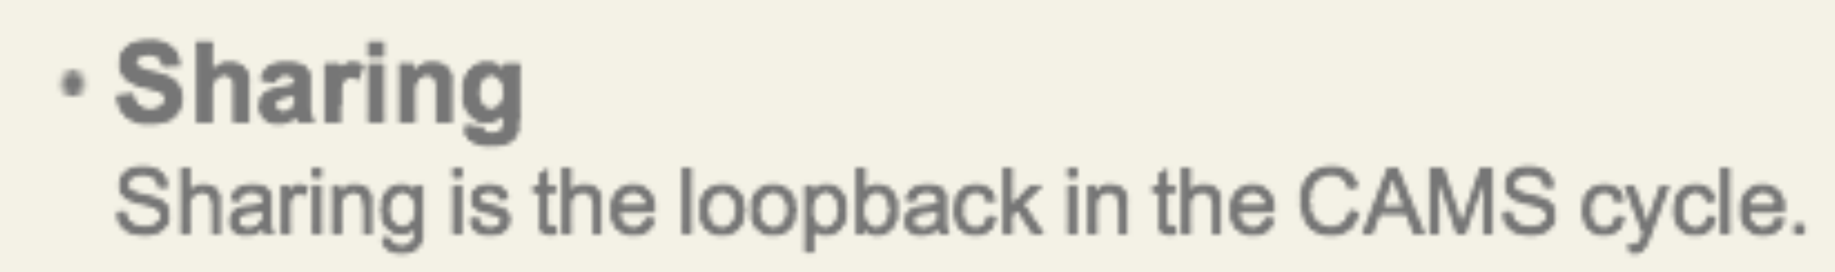

- **Sharing**
  Sharing is the loopback in the CAMS cycle.

- "Shared problems are the basis of teamwork"
- Root Cause Analysis is our established mechanism for addressing problems
- Success on shared problems have been creating a virtuous cycle

The wake up call

Fault tolerant infrastructure

+

Fault tolerant architecture

=

Outage within 15 minutes of going live

- New component introduced to have calculation code against a slave database
- API code transparently calling new component; fallback to old code on failure
- Two component instances behind a load balancer

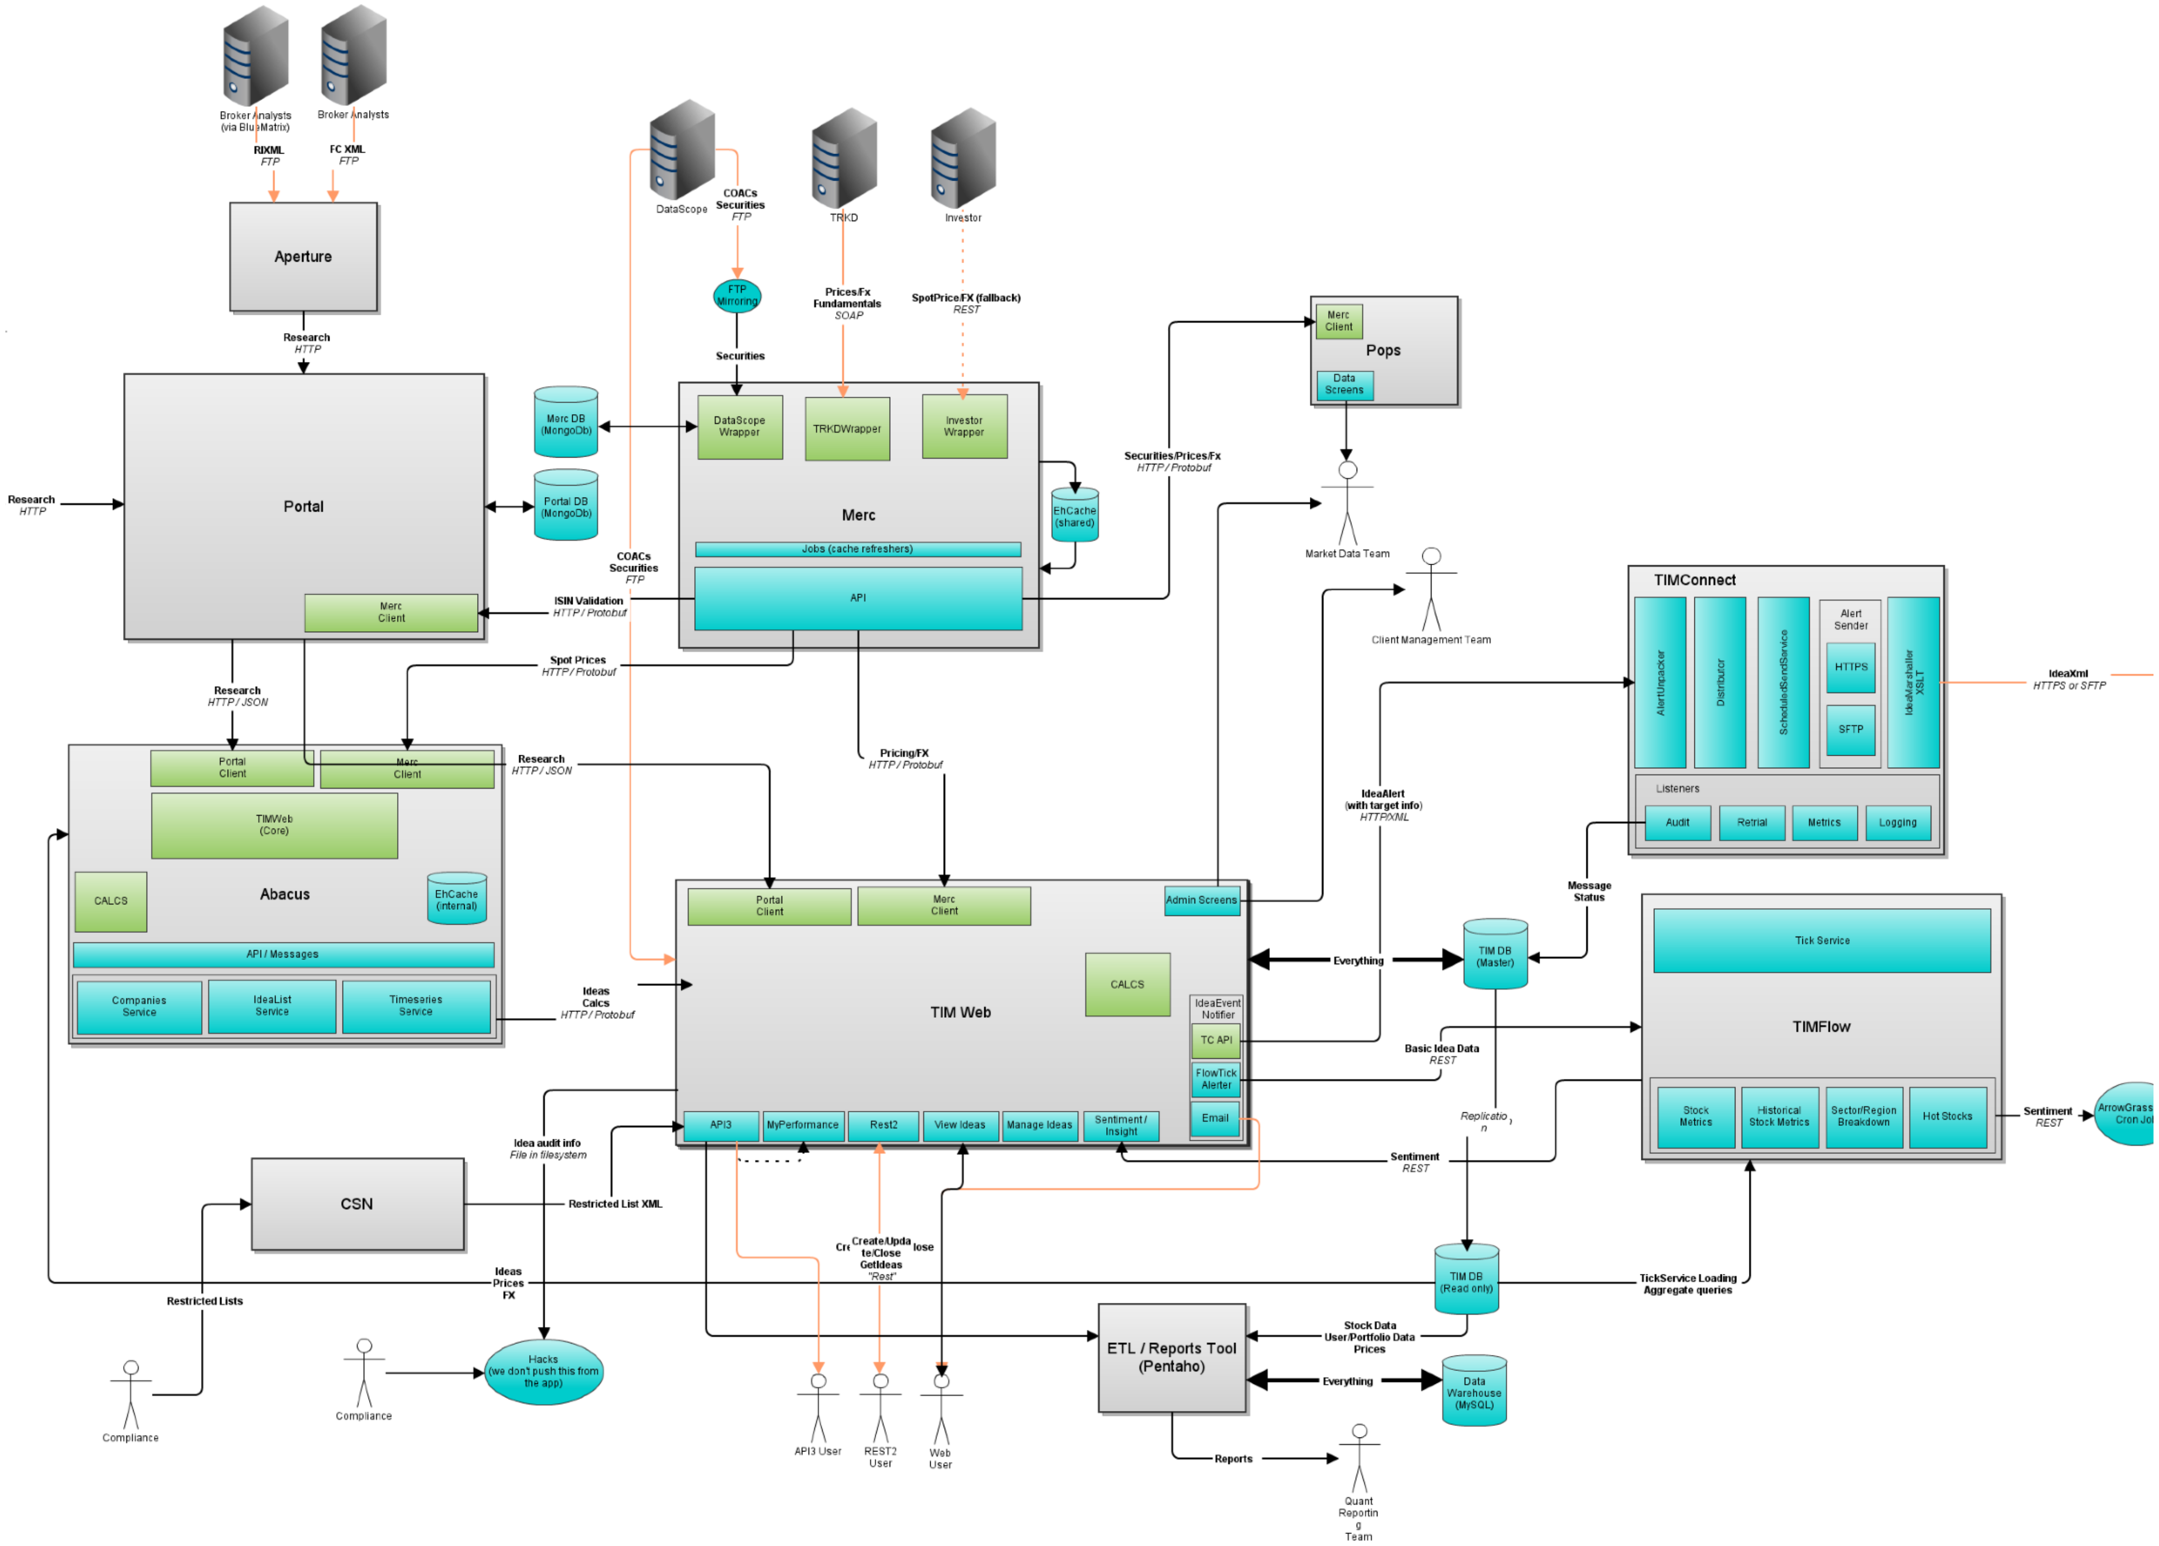

grams at gliffy.com

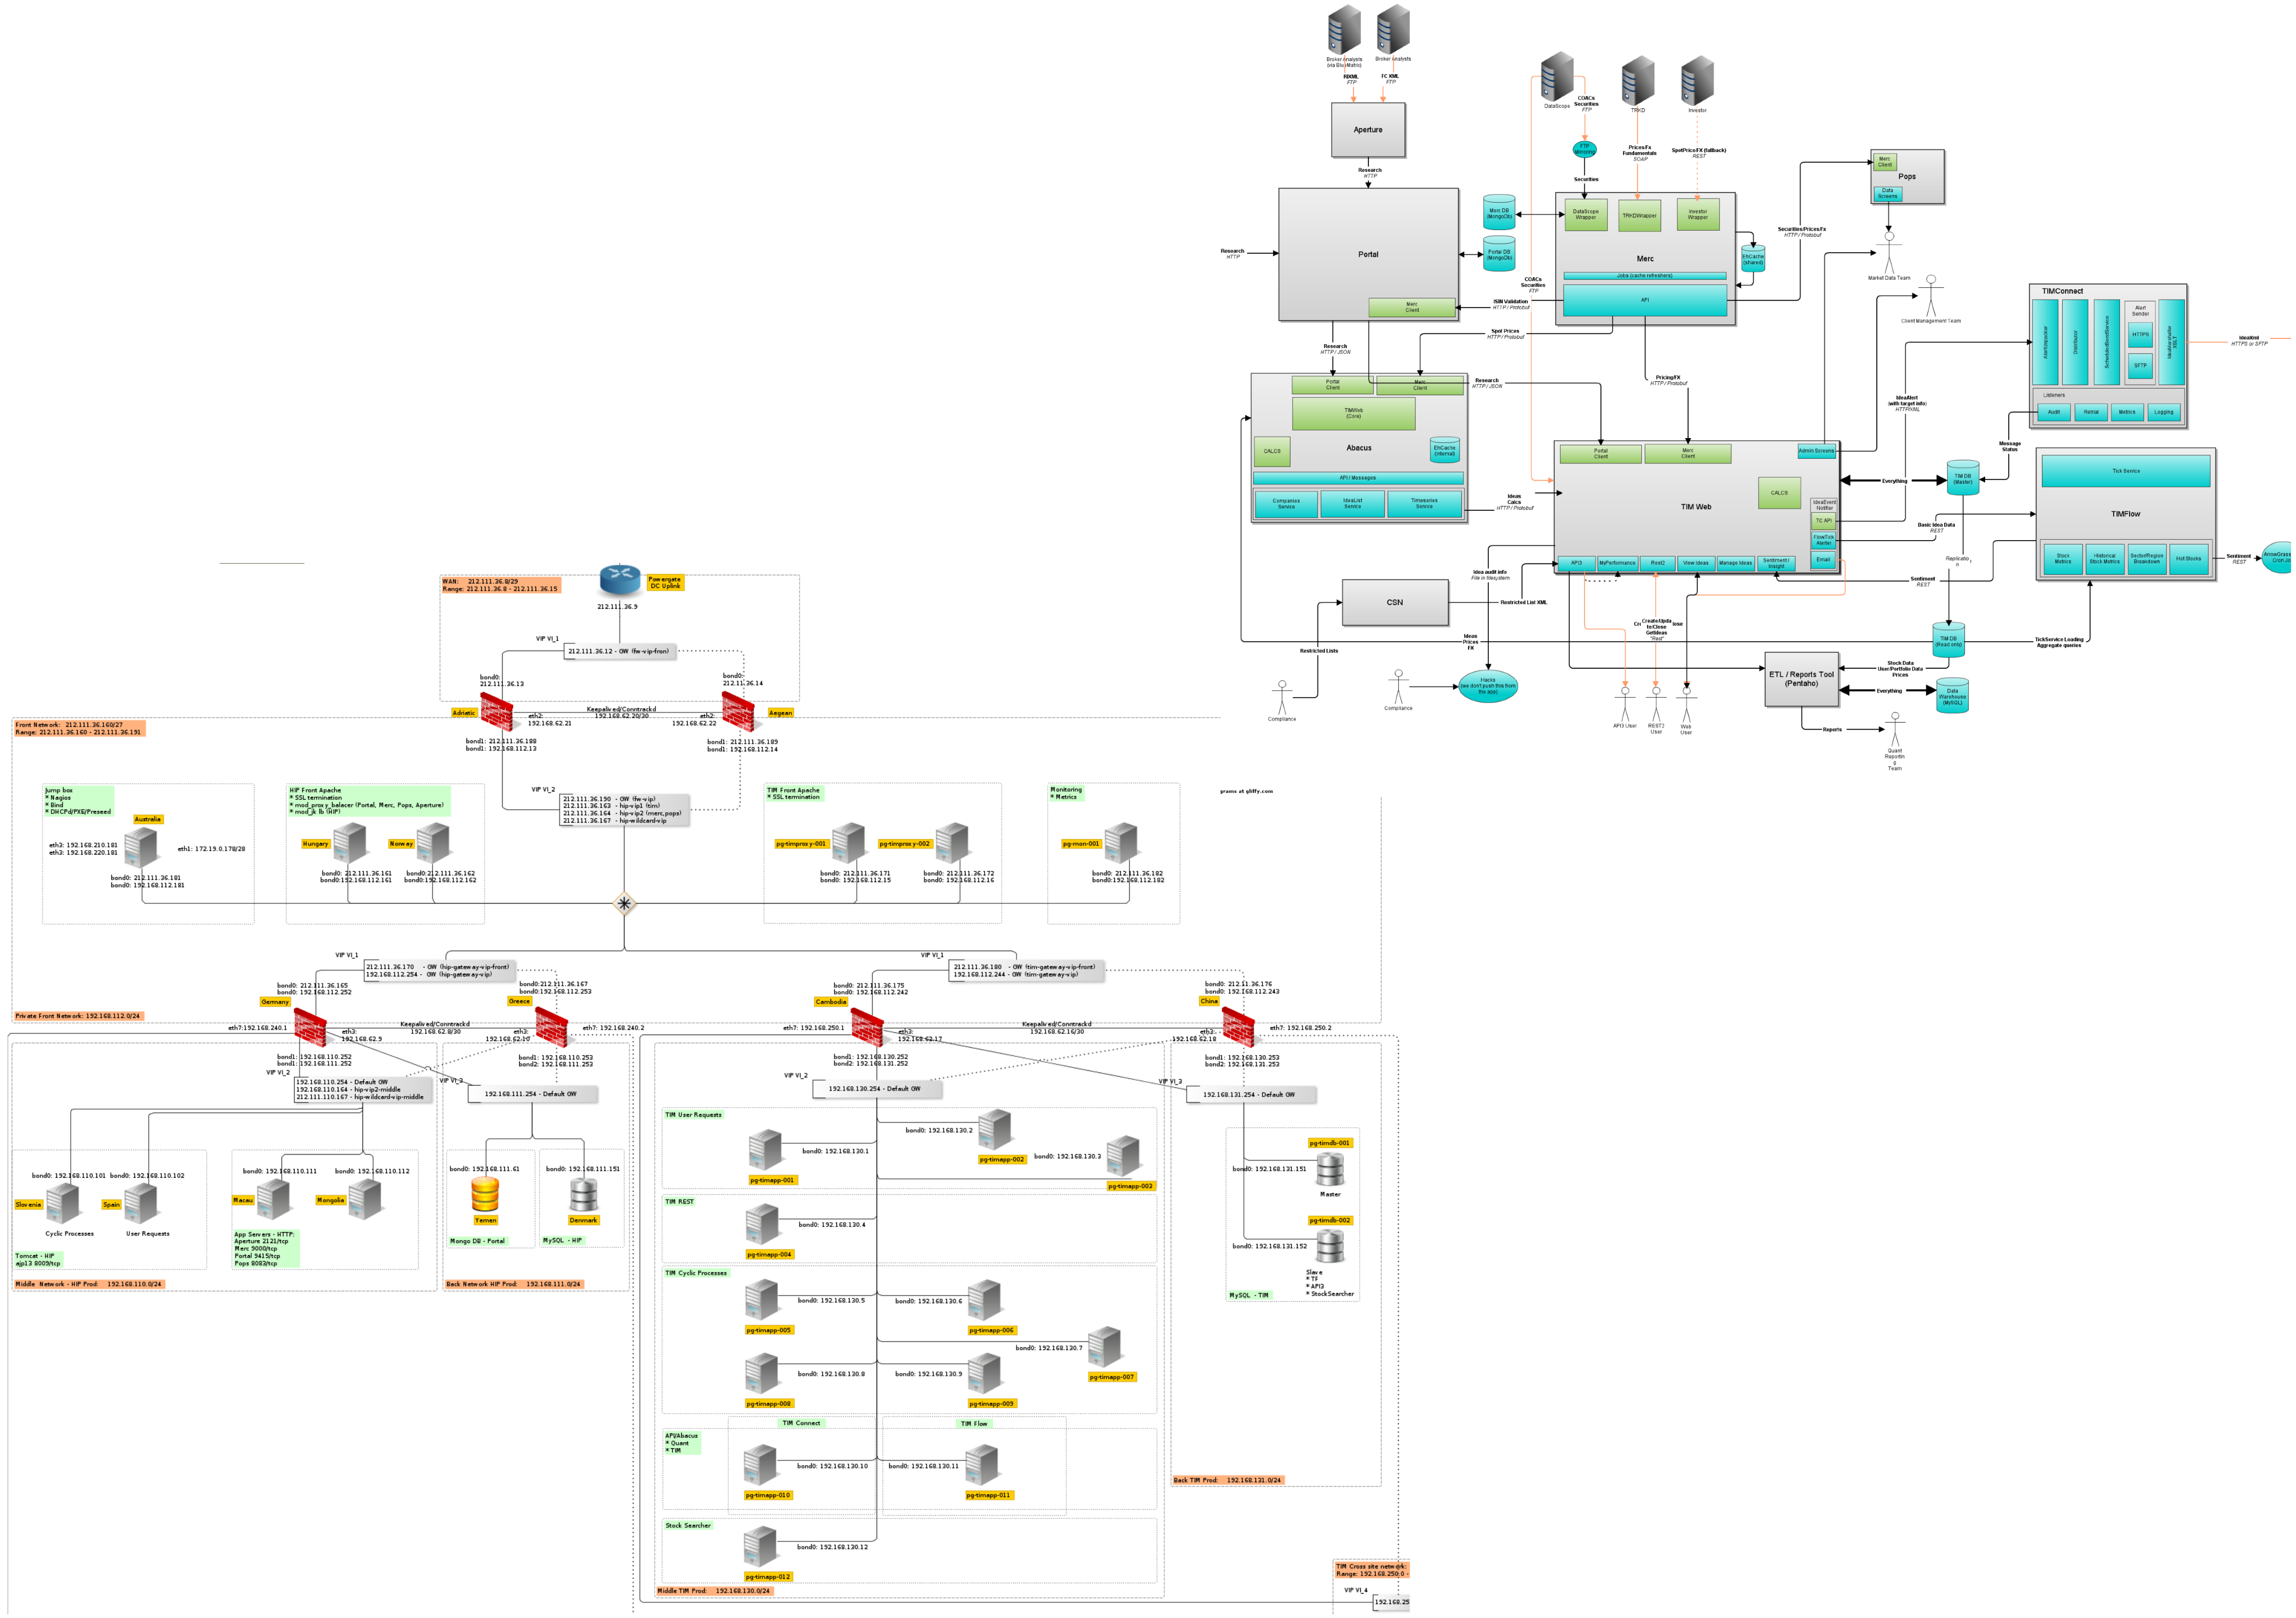
Broker Analysts (via BlueMatrix)
Broker Analysts
RIXML
FTP
FC XML
FTP
Aperture
Research
HTTP
Research
HTTP
Portal
Merc Client
Merc DB (MongoDb)
Portal DB (MongoDb)
DataScope
COACs
Securities
FTP
FTP Mirroring
Securities
TRKD
Prices/Fx
Fundamentals
SOAP
Investor
SpotPrice/FX (fallback)
REST
DataScope Wrapper
TRKDWrapper
Investor Wrapper
Merc
EhCache (shared)
Jobs (cache refreshers)
API
COACs
Securities
FTP
ISIN Validation
HTTP / Protobuf
Spot Prices
HTTP / Protobuf
Securities/Prices/Fx
HTTP / Protobuf
Merc Client
Pops
Data Screens
Market Data Team
Client Management Team
Research
HTTP / JSON
Research
HTTP / JSON
Pricing/FX
HTTP / Protobuf
Portal Client
Merc Client
TIMWeb (Core)
CALCS
Abacus
EhCache (internal)
API / Messages
Companies Service
IdeaList Service
Timeseries Service
Ideas
Calcs
HTTP / Protobuf
Portal Client
Merc Client
Admin Screens
CALCS
TIM Web
IdeaEvent Notifier
TC API
FlowTick Alerter
Email
API3
MyPerformance
Rest2
View Ideas
Manage Ideas
Sentiment / Insight
TIMConnect
AlertUnpacker
Distributor
ScheduledSendService
Alert Sender
HTTPS
SFTP
IdeaMarshaller XSLT
IdeaXml
HTTPS or SFTP
Listeners
Audit
Retrial
Metrics
Logging
Message Status
IdeaAlert (with target info)
HTTP/XML
Everything
TIM DB (Master)
Basic Idea Data
REST
Tick Service
TIMFlow
Stock Metrics
Historical Stock Metrics
Sector/Region Breakdown
Hot Stocks
Sentiment
REST
ArrowGrass Cron Jo
Replication
Sentiment
REST
Idea audit info
File in filesystem
CSN
Restricted List XML
Restricted Lists
Compliance
Compliance
Hacks (we don't push this from the app)
Ideas
Prices
FX
Create/Update/Close
GetIdeas
"Rest"
API3 User
REST2 User
Web User
TIM DB (Read only)
TickService Loading
Aggregate queries
ETL / Reports Tool (Pentaho)
Stock Data
User/Portfolio Data
Prices
Everything
Data Warehouse (MySQL)
Reports
Quant Reporting Team
grams at gliffy.com
WAN: 212.111.36.8/29
Range: 212.111.36.8 - 212.111.36.15
Powergate DC Uplink
212.111.36.9
VIP VI_1
212.111.36.12 - GW (fw-vip-fron)
bond0: 212.111.36.13
bond0: 212.111.36.14
Adriatic
Aegean
Keepalived/Conntrackd
192.168.62.20/30
eth2: 192.168.62.21
eth2: 192.168.62.22
bond1: 212.111.36.188
bond1: 192.168.112.13
bond1: 212.111.36.189
bond1: 192.168.112.14
Front Network: 212.111.36.160/27
Range: 212.111.36.160 - 212.111.36.191
VIP VI_2
212.111.36.190 - GW (fw-vip)
212.111.36.163 - hip-vip1 (tim)
212.111.36.164 - hip-vip2 (merc,pops)
212.111.36.167 - hip-wildcard-vip
Jump box
* Nagios
* Bind
* DHCPd/PXE/Preseed
Australia
eth3: 192.168.210.181
eth3: 192.168.220.181
eth1: 172.19.0.178/28
bond0: 212.111.36.181
bond0: 192.168.112.181
HIP Front Apache
* SSL termination
* mod_proxy_balacer (Portal, Merc, Pops, Aperture)
* mod_jk lb (HIP)
Hungary
Norway
bond0: 212.111.36.161
bond0:192.168.112.161
bond0:212.111.36.162
bond0:192.168.112.162
TIM Front Apache
* SSL termination
pg-timproxy-001
pg-timproxy-002
bond0: 212.111.36.171
bond0: 192.168.112.15
bond0: 212.111.36.172
bond0: 192.168.112.16
Monitoring
* Metrics
pg-mon-001
bond0: 212.111.36.182
bond0:192.168.112.182
VIP VI_1
212.111.36.170 - GW (hip-gateway-vip-front)
192.168.112.254 - GW (hip-gateway-vip)
bond0: 212.111.36.165
bond0: 192.168.112.252
bond0:212.111.36.167
bond0:192.168.112.253
Germany
Greece
Private Front Network: 192.168.112.0/24
eth7:192.168.240.1
eth7: 192.168.240.2
Keepalived/Conntrackd
192.168.62.8/30
eth3: 192.168.62.9
eth3: 192.168.62.10
bond1: 192.168.110.252
bond1: 192.168.111.252
bond1: 192.168.110.253
bond2: 192.168.111.253
VIP VI_2
192.168.110.254 - Default GW
192.168.110.164 - hip-vip2-middle
212.111.110.167 - hip-wildcard-vip-middle
VIP VI_3
192.168.111.254 - Default GW
bond0: 192.168.110.101
bond0: 192.168.110.102
Slovenia
Spain
Cyclic Processes
User Requests
Tomcat - HIP
ajp13 8009/tcp
bond0: 192.168.110.111
bond0: 192.168.110.112
Macau
Mongolia
App Servers - HTTP:
Aperture 2121/tcp
Merc 9000/tcp
Portal 9415/tcp
Pops 8083/tcp
Middle Network - HIP Prod: 192.168.110.0/24
bond0: 192.168.111.61
bond0: 192.168.111.151
Yemen
Denmark
Mongo DB - Portal
MySQL - HIP
Back Network HIP Prod: 192.168.111.0/24
VIP VI_1
212.111.36.180 - GW (tim-gateway-vip-front)
192.168.112.244 - GW (tim-gateway-vip)
bond0: 212.111.36.175
bond0: 192.168.112.242
bond0: 212.11.36.176
bond0: 192.168.112.243
Cambodia
China
eth7: 192.168.250.1
eth7: 192.168.250.2
Keepalived/Conntrackd
192.168.62.16/30
eth3: 192.168.62.17
eth3: 192.168.62.18
bond1: 192.168.130.252
bond2: 192.168.131.252
bond1: 192.168.130.253
bond2: 192.168.131.253
VIP VI_2
192.168.130.254 - Default GW
VIP VI_3
192.168.131.254 - Default GW
TIM User Requests
bond0: 192.168.130.1
pg-timapp-001
bond0: 192.168.130.2
pg-timapp-002
bond0: 192.168.130.3
pg-timapp-003
TIM REST
bond0: 192.168.130.4
pg-timapp-004
TIM Cyclic Processes
bond0: 192.168.130.5
pg-timapp-005
bond0: 192.168.130.6
pg-timapp-006
bond0: 192.168.130.7
pg-timapp-007
bond0: 192.168.130.8
pg-timapp-008
bond0: 192.168.130.9
pg-timapp-009
TIM Connect
TIM Flow
API/Abacus
* Quant
* TIM
bond0: 192.168.130.10
pg-timapp-010
bond0: 192.168.130.11
pg-timapp-011
Stock Searcher
bond0: 192.168.130.12
pg-timapp-012
Middle TIM Prod: 192.168.130.0/24
pg-timdb-001
bond0: 192.168.131.151
Master
pg-timdb-002
bond0: 192.168.131.152
Slave
* TF
* API3
* StockSearcher
MySQL - TIM
Back TIM Prod: 192.168.131.0/24
TIM Cross site network:
Range: 192.168.250.0 -
VIP VI_4
192.168.25

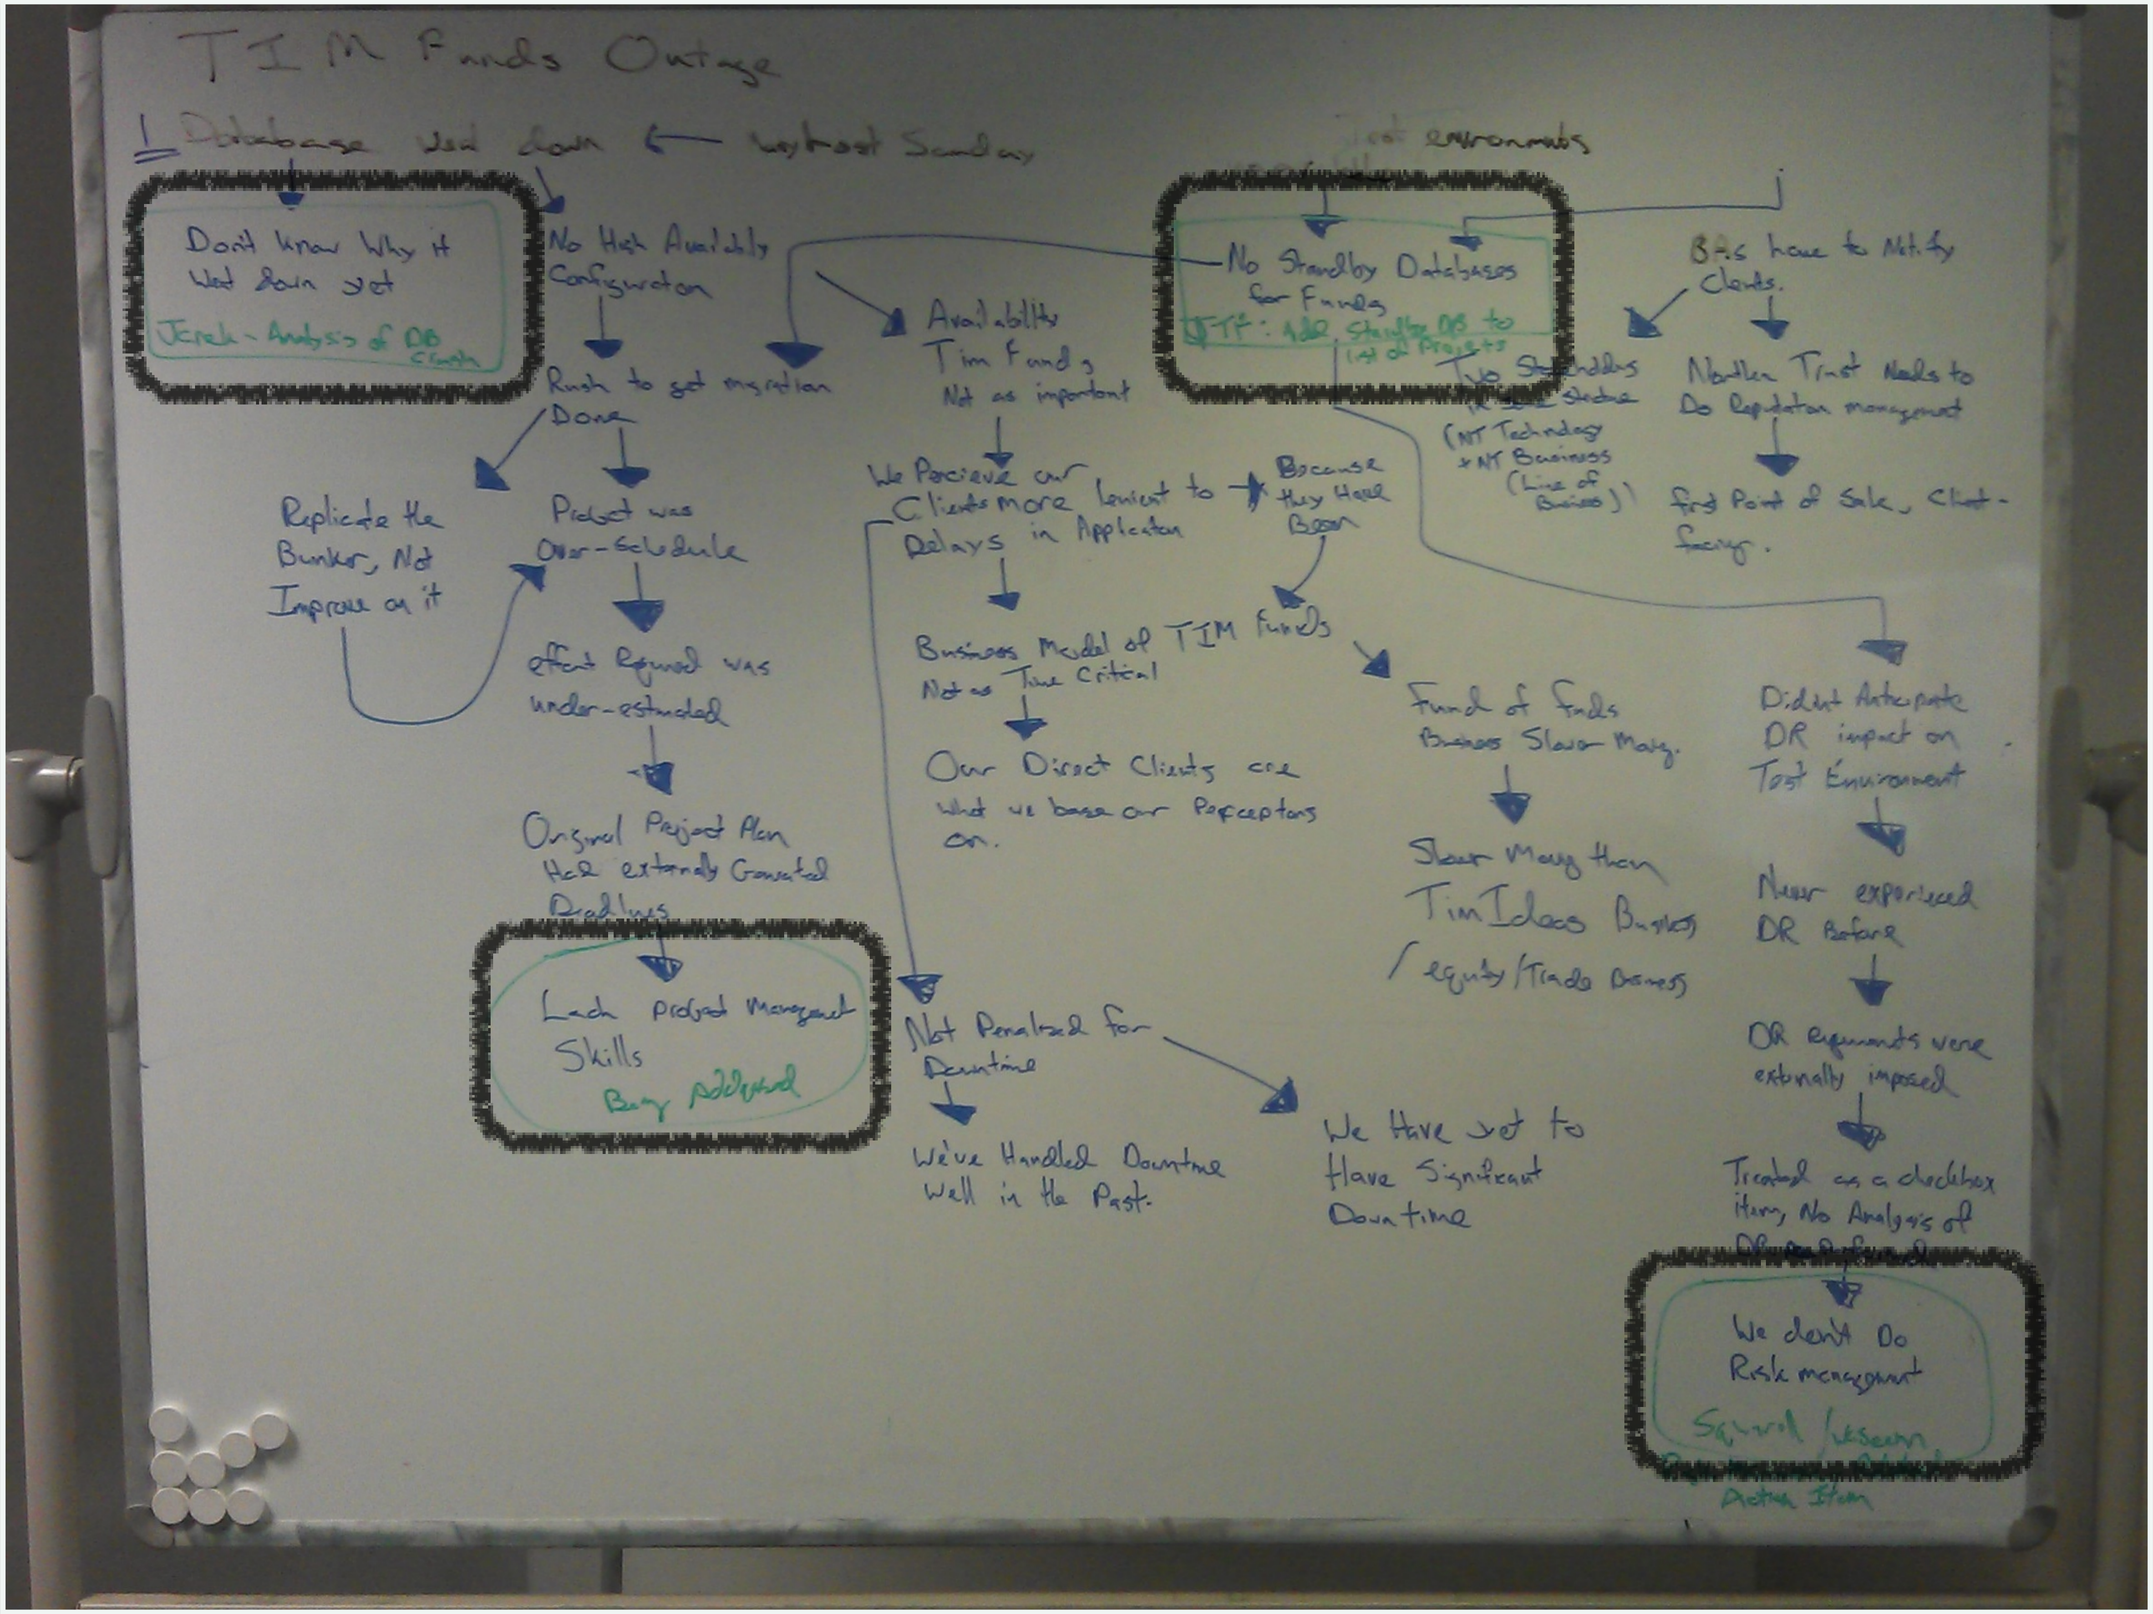
TIM Funds Outage
1 Database Went down ← Wayfrost Sunday
Test environments
Don't know Why it Went down yet
Jerek - Analysis of DB Crash
No High Availability Configuration
Rush to get migration Done
Replicate the Bunker, Not Improve on it
Project was Over-Schedule
effort Required was under-estimated
Original Project Plan Had externally Generated Deadlines
Lack project Management Skills
Being Addressed
Availability Tim Funds Not as important
We Percieve our Clients more lenient to Delays in Application
Because they Have Been
Business Model of TIM Funds Not as Time Critical
Our Direct Clients are what we base our Perceptions on.
Not Penalized for Downtime
We've Handled Downtime Well in the Past.
We Have yet to Have Significant Downtime
No Standby Databases for Funds
TF: Add Standby DB to List of Projects
Two Stakeholders (NT Technology + NT Business (Line of Business))
BAs have to Notify Clients.
Northern Trust Needs to Do Reputation management
first Point of Sale, Client-facing.
Fund of funds Business Slower Moving.
Slower Moving than Tim Ideas Business / equity / Trade Business
Didn't Anticipate DR impact on Test Environment
Never experienced DR Before
DR requirements were externally imposed
Treated as a checkbox item, No Analysis of
We don't Do Risk management
Squirrel / Weseem
Action Item

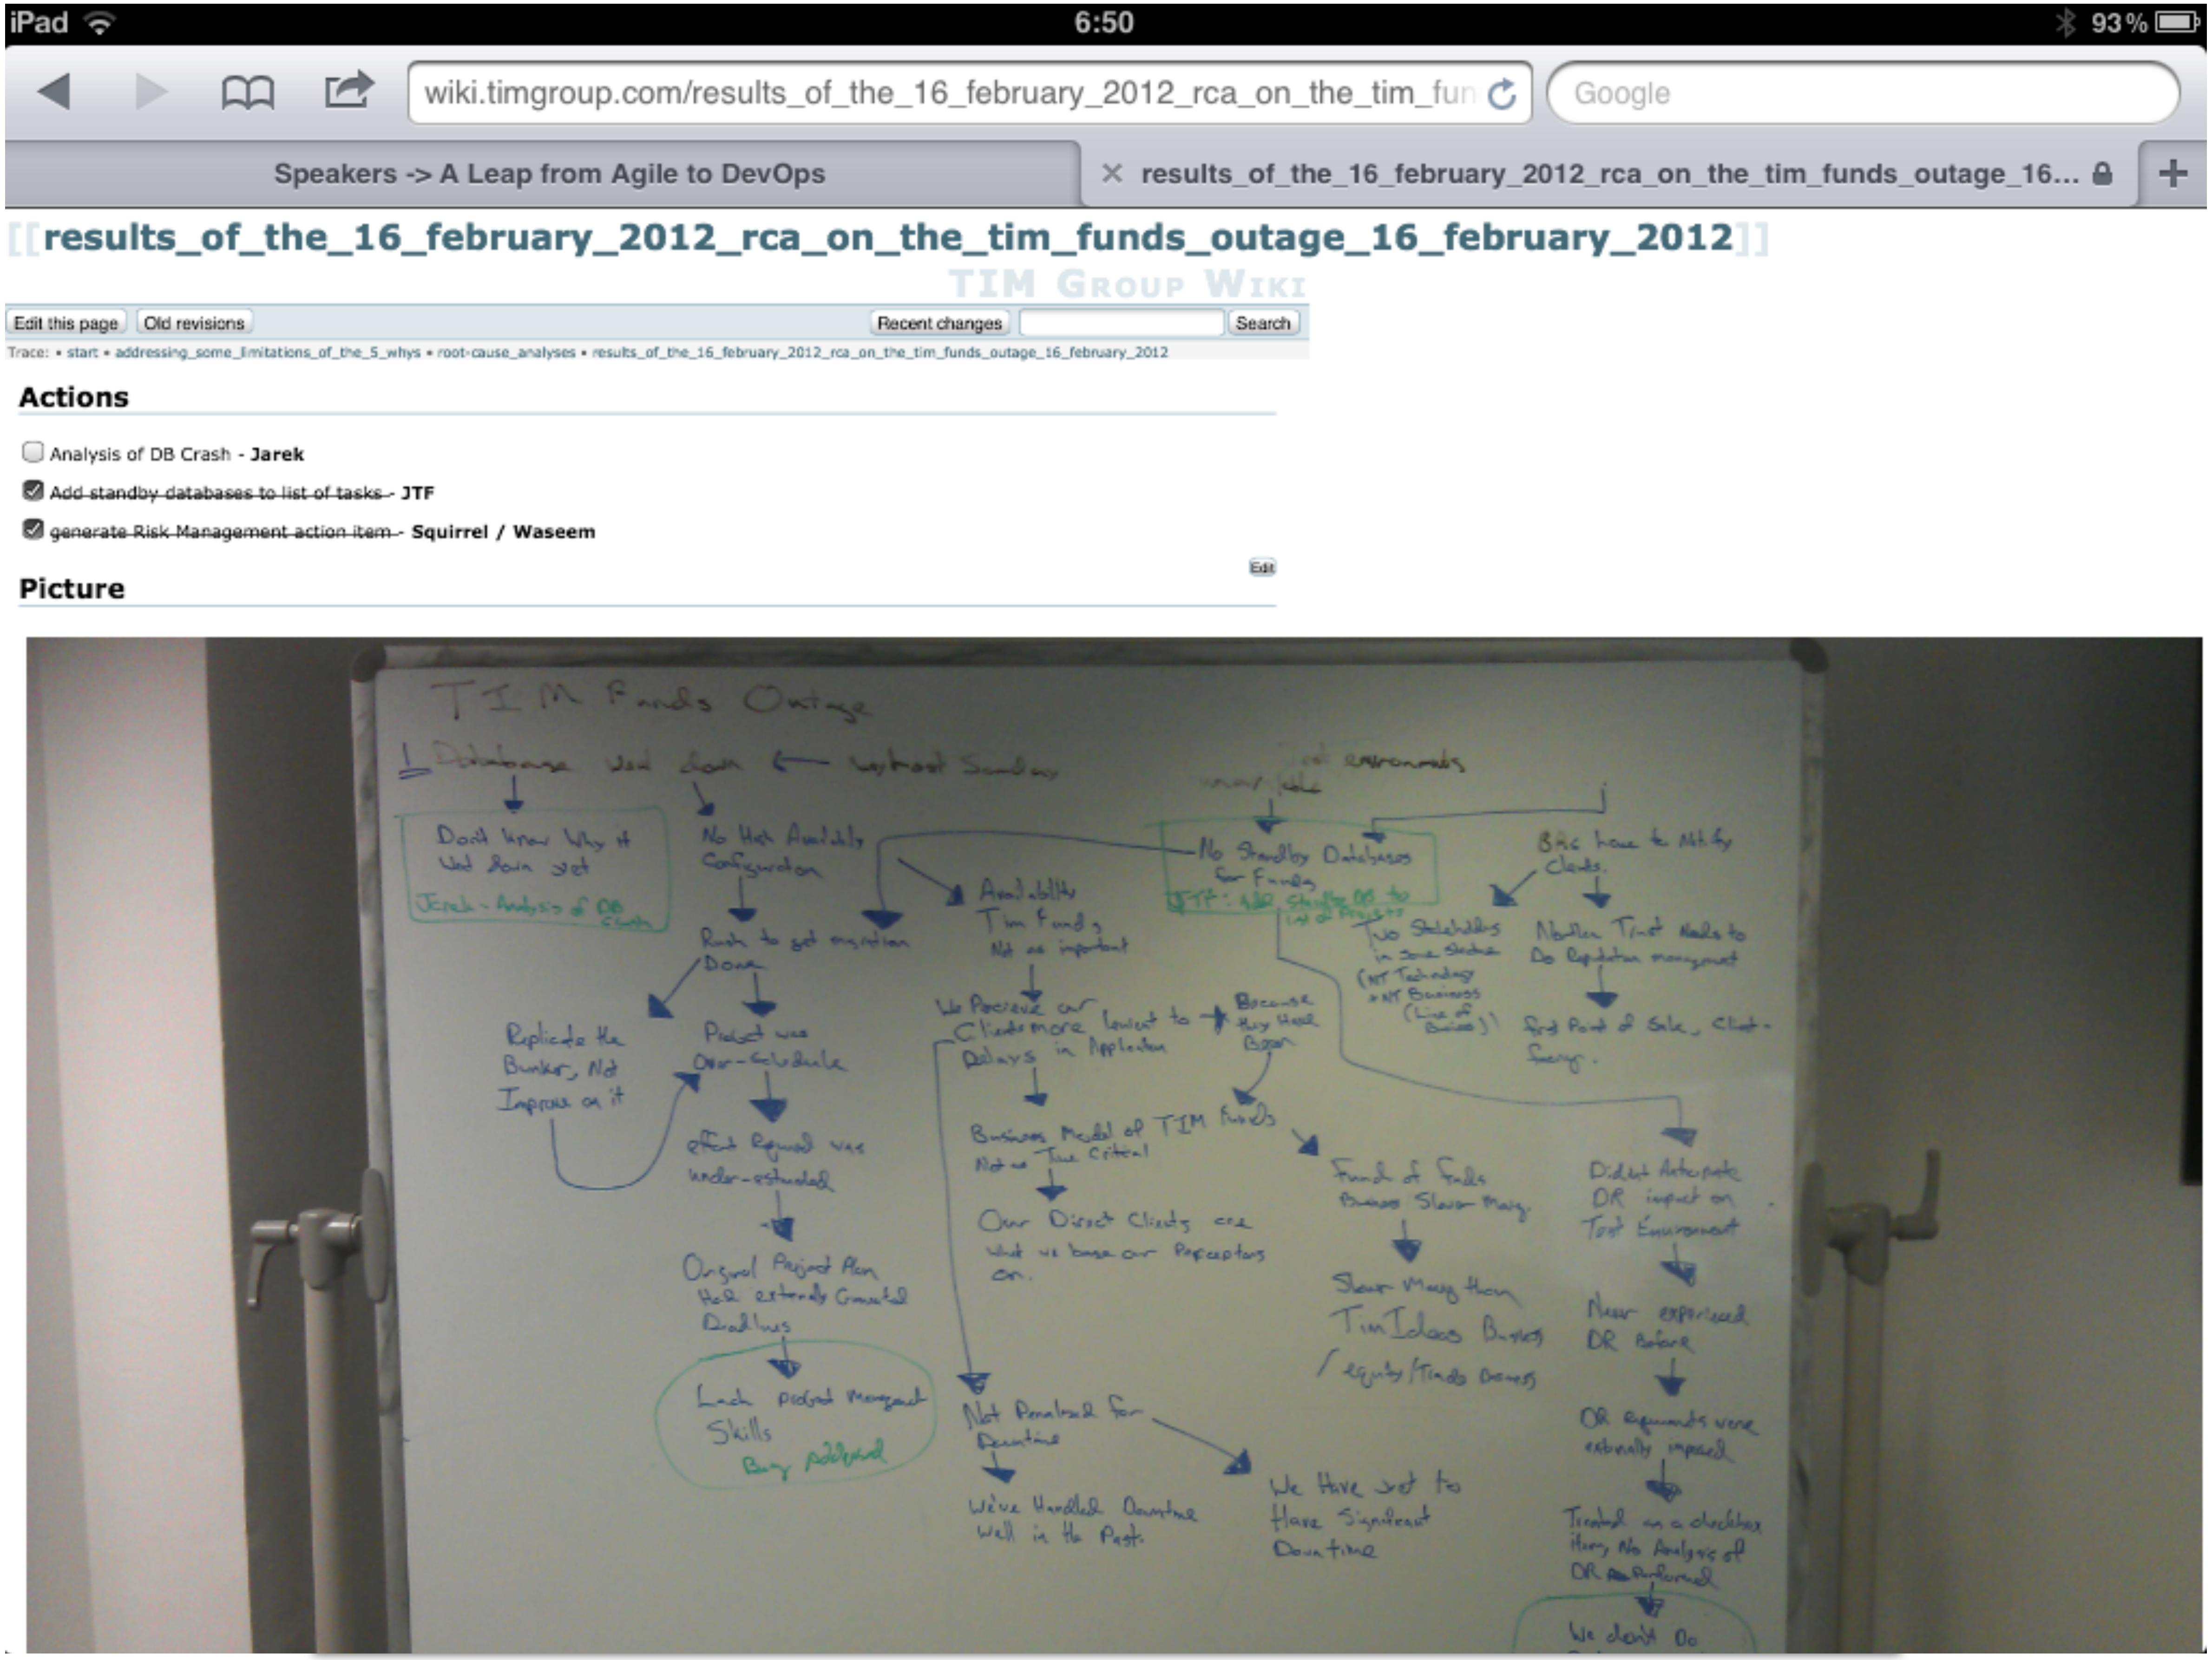
iPad
6:50
93 %
wiki.timgroup.com/results_of_the_16_february_2012_rca_on_the_tim_fun
Google
Speakers -> A Leap from Agile to DevOps
results_of_the_16_february_2012_rca_on_the_tim_funds_outage_16...
+

[[results_of_the_16_february_2012_rca_on_the_tim_funds_outage_16_february_2012]]

TIM GROUP WIKI

Edit this page
Old revisions
Recent changes
Search

Trace: • start • addressing_some_limitations_of_the_5_whys • root-cause_analyses • results_of_the_16_february_2012_rca_on_the_tim_funds_outage_16_february_2012

Actions

- ☐ Analysis of DB Crash - **Jarek**
- ☑ ~~Add standby databases to list of tasks~~ - **JTF**
- ☑ ~~generate Risk Management action item~~ - **Squirrel / Waseem**

Picture

TIM Funds Outage
1 Database Went down ← without Standby
Don't know Why it Went down yet
Jarek - Analysis of DB Crash
No High Availability Configuration
Rush to get migration Done
Replicate the Bunker, Not Improve on it
Product was Over-Schedule
effort Required was under-estimated
Original Project Plan Had externally Generated Deadlines
Lack product Management Skills
Being Addressed
Availability Tim Funds Not as important
We Perceive our Clients more lenient to Delays in Application
Because they Have Been
Business Model of TIM Funds Not as Time Critical
Our Direct Clients are what we base our Perceptions on.
Not Penalised for Downtime
We've Handled Downtime Well in the Past.
We Have yet to Have Significant Downtime
Test environments
unavailable
No Standby Databases for Funds
JTF: Add Standby DB to list of Projects
No Stakeholders in Some Sense (NT Technology + NT Business (Line of Business))
BAs have to Notify Clients.
Northern Trust Needs to Do Reputation management
First Point of Sale, Client-Facing.
Fund of Funds Business Slower Moving
Slower Moving than TimIdeas Buyers / equity Funds Owners
Didn't Anticipate DR impact on Test Environment
Never experienced DR Before
DR Requirements were externally imposed
Treated as a checkbox item, No Analysis of DR Performed
We don't Do

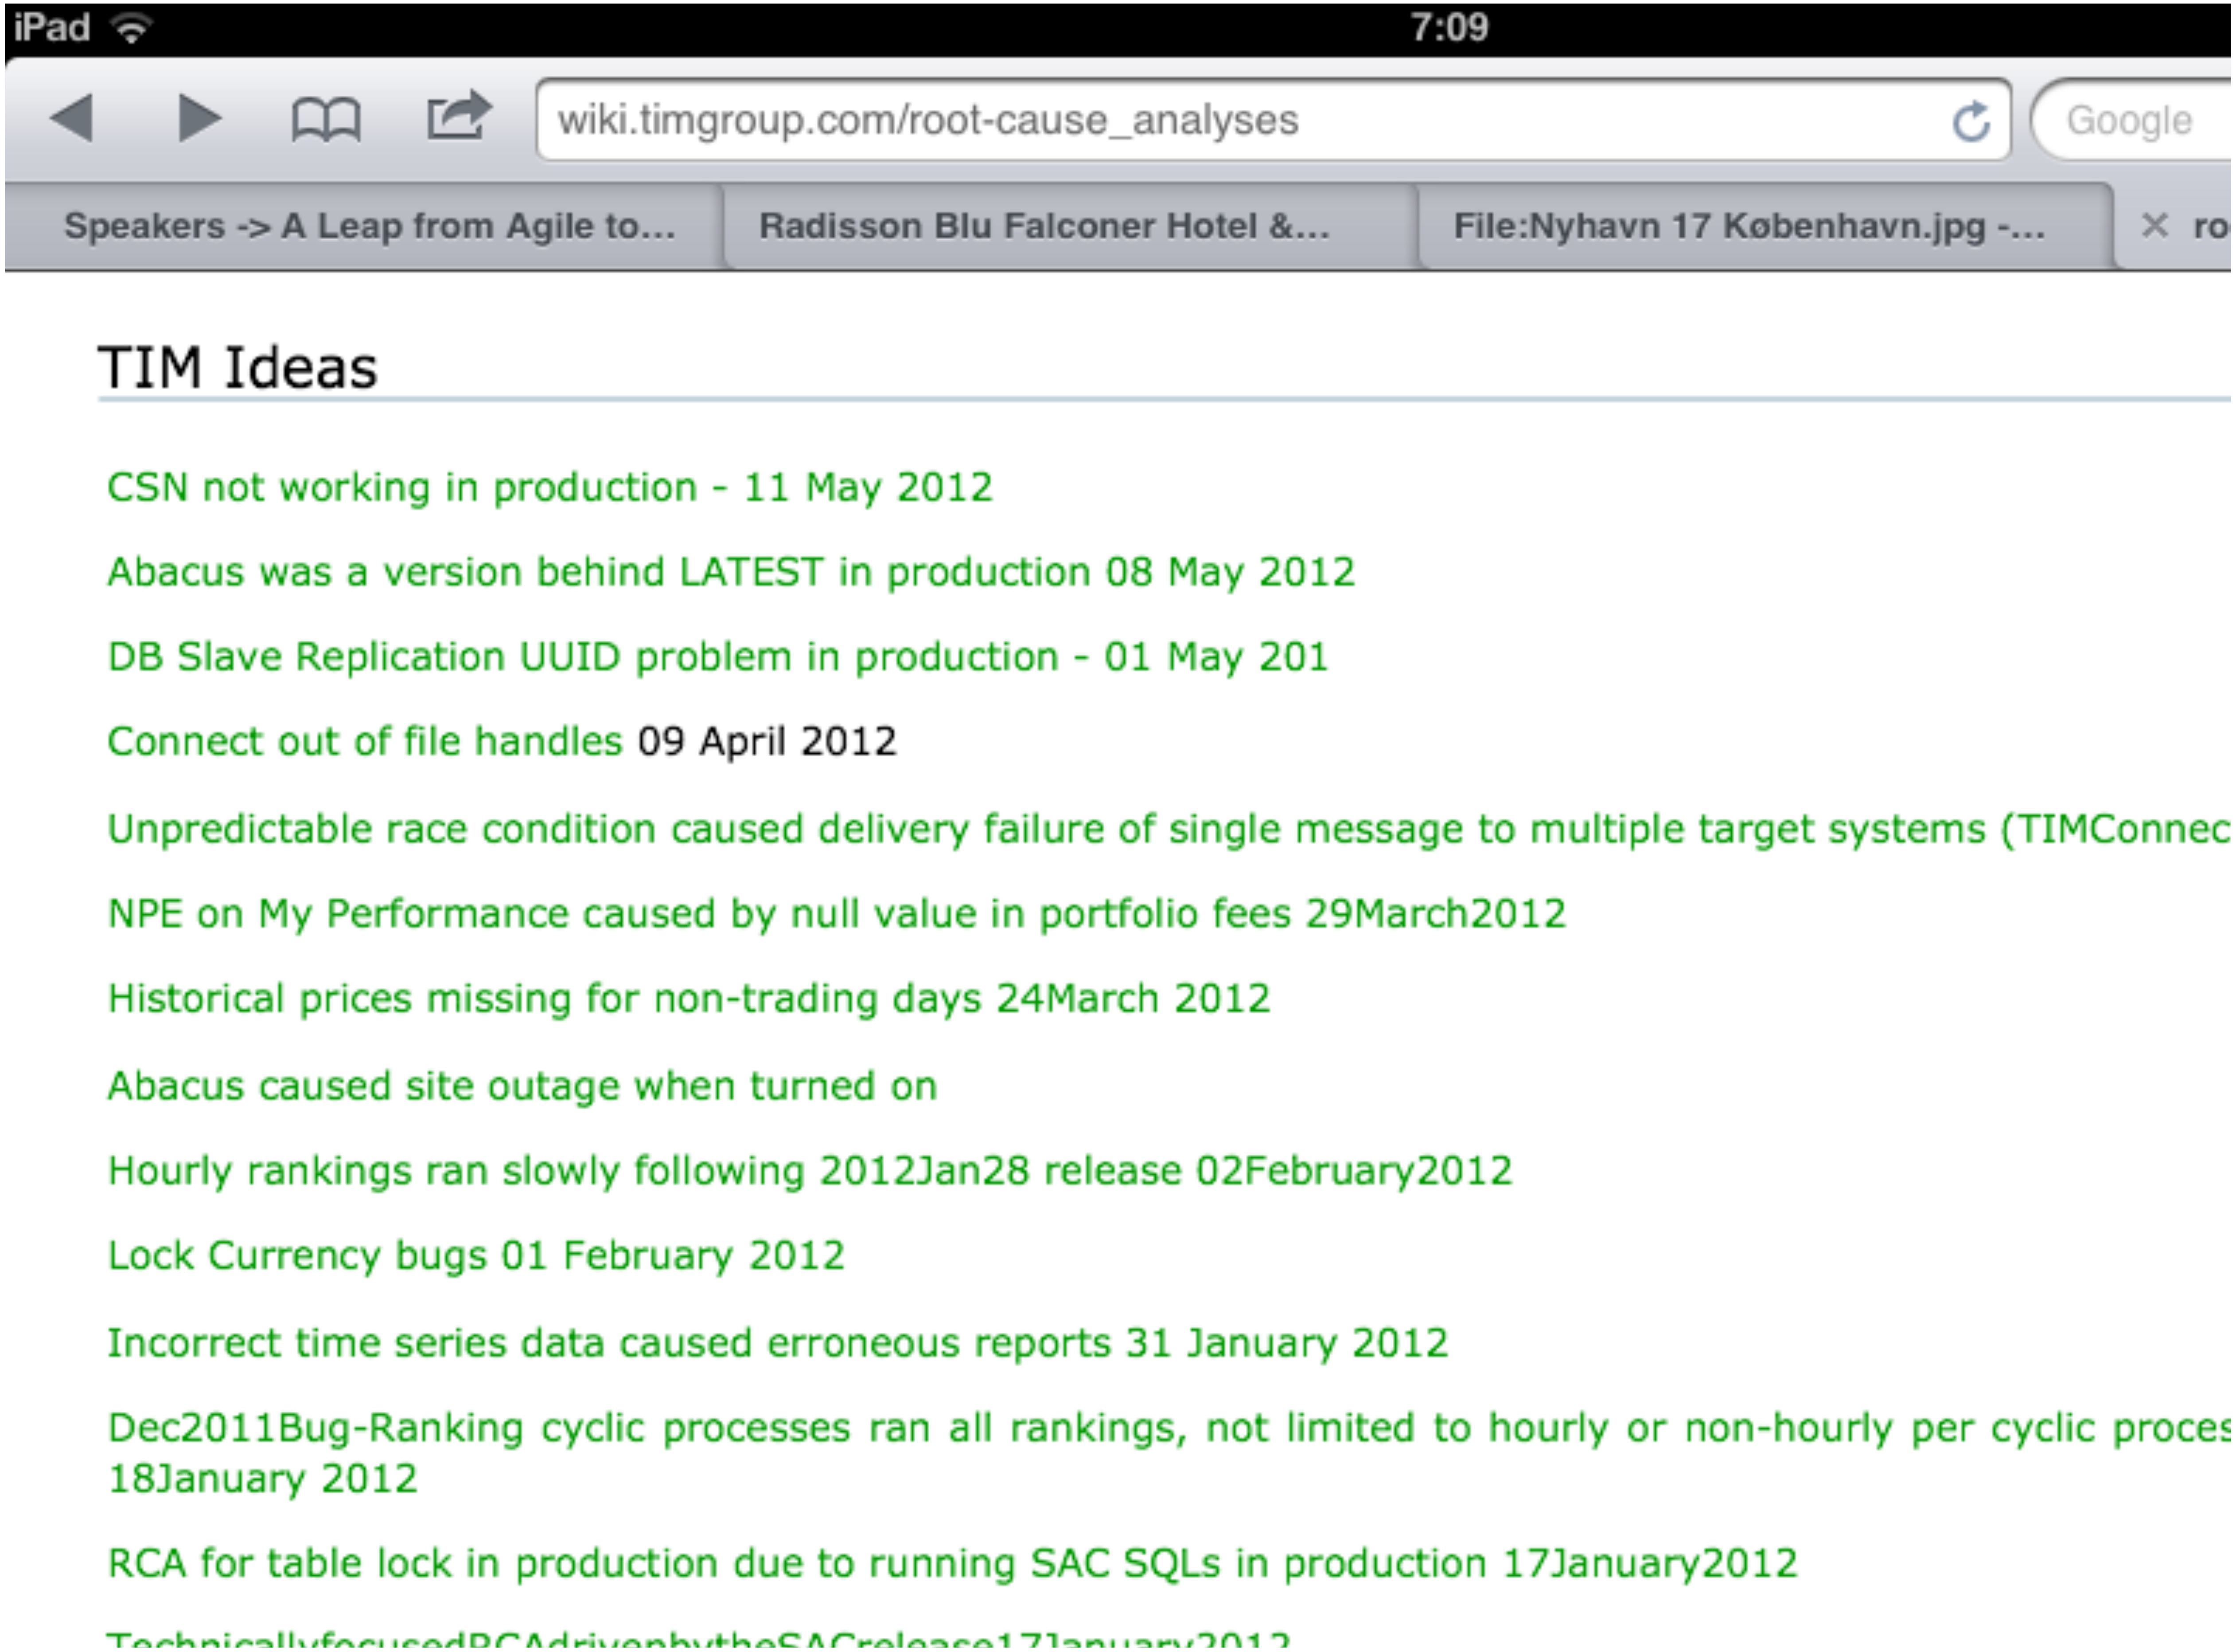
iPad
7:09
wiki.timgroup.com/root-cause_analyses
Google
Speakers -> A Leap from Agile to...
Radisson Blu Falconer Hotel &...
File:Nyhavn 17 København.jpg -...
ro

TIM Ideas

CSN not working in production - 11 May 2012

Abacus was a version behind LATEST in production 08 May 2012

DB Slave Replication UUID problem in production - 01 May 201

Connect out of file handles 09 April 2012

Unpredictable race condition caused delivery failure of single message to multiple target systems (TIMConnec

NPE on My Performance caused by null value in portfolio fees 29March2012

Historical prices missing for non-trading days 24March 2012

Abacus caused site outage when turned on

Hourly rankings ran slowly following 2012Jan28 release 02February2012

Lock Currency bugs 01 February 2012

Incorrect time series data caused erroneous reports 31 January 2012

Dec2011Bug-Ranking cyclic processes ran all rankings, not limited to hourly or non-hourly per cyclic proces 18January 2012

RCA for table lock in production due to running SAC SQLs in production 17January2012

TechnicallyfocusedRCAdrivenbytheSACrelease17January2012

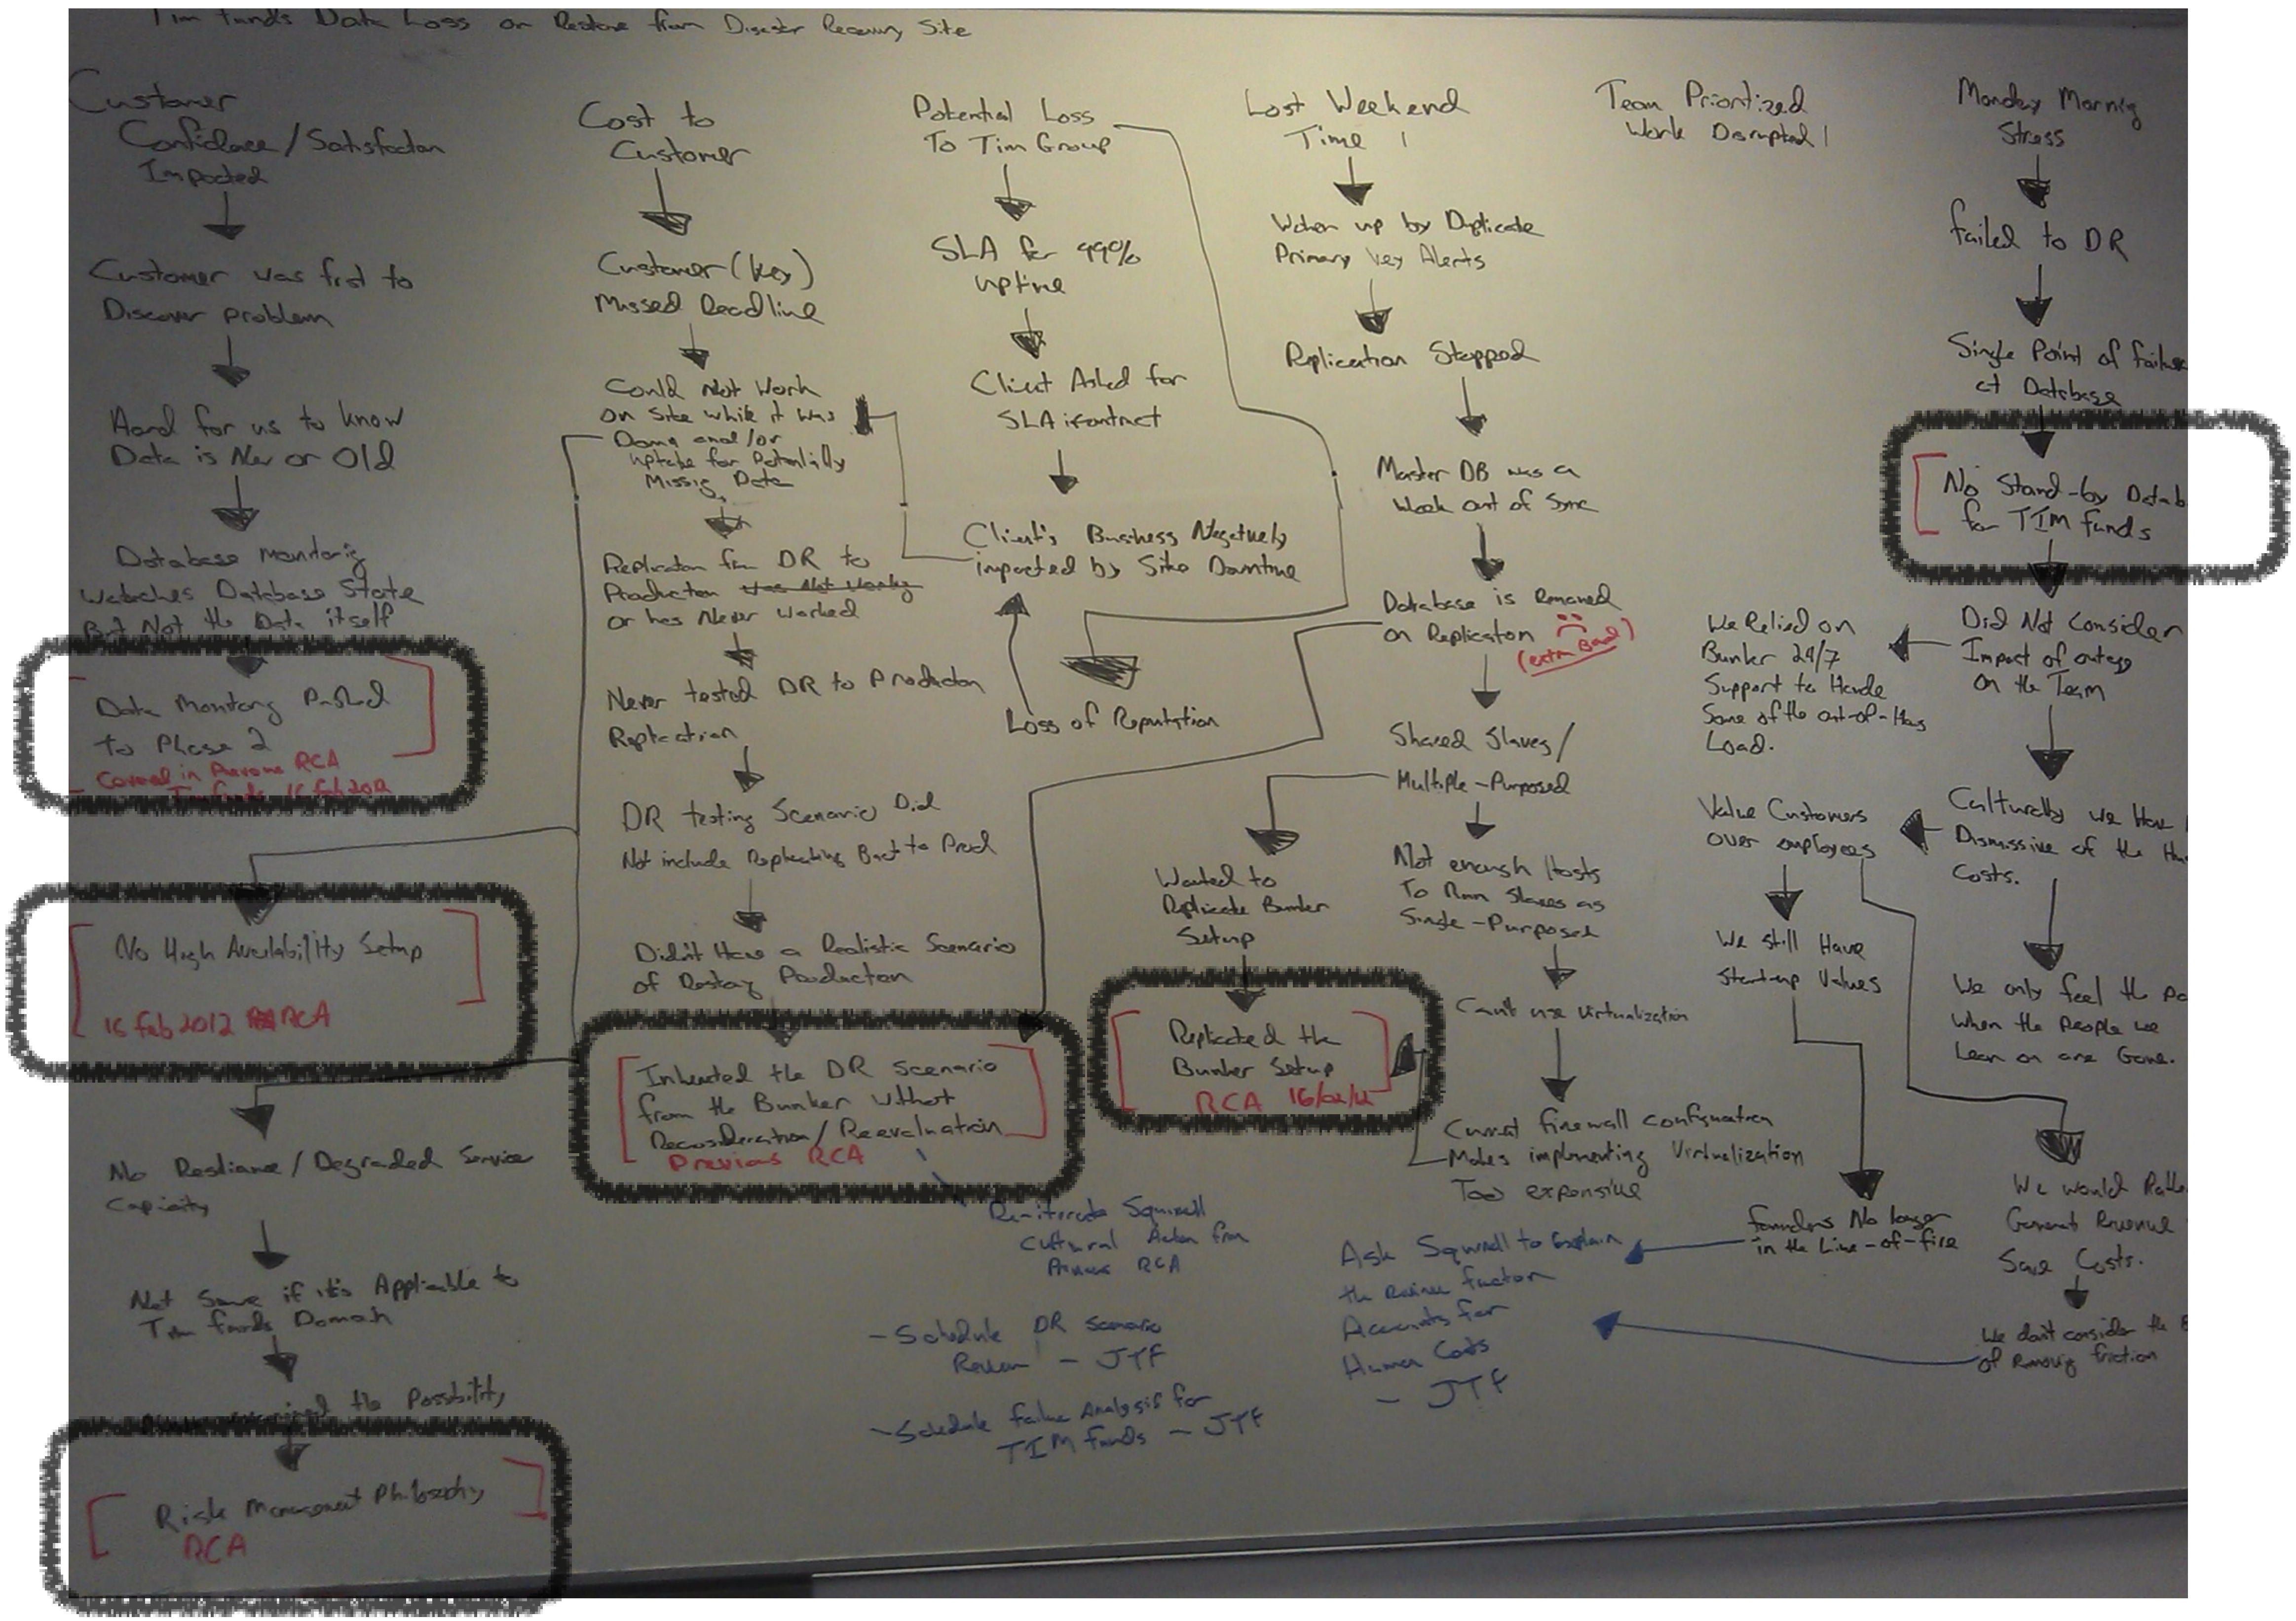

Tim funds Data Loss on Restore from Disaster Recovery Site

Customer Confidence / Satisfaction Impacted

Customer was first to Discover problem

Hard for us to know Data is New or Old

Database Monitoring Watches Database State But Not the Data itself

Data Monitoring Pushed to Phase 2

Covered in Previous RCA Tim funds 16 Feb 2012

No High Availability Setup

16 Feb 2012 RCA

No Resiliance / Degraded Service Capacity

Not Sure if it's Applicable to Tim funds Domain

[illegible] the Possibility

Risk Management Philosophy

RCA

Cost to Customer

Customer (key) Missed Deadline

Could Not Work On Site while it was Down and/or Update for Potentially Missing Data

Replication from DR to Production ~~has Not Working~~ or has Never Worked

Never tested DR to Production Replication

DR testing Scenario Did Not include Replicating Back to Prod

Didn't Have a Realistic Scenario of Restoring Production

Inherited the DR Scenario from the Bunker without Reconsideration / Reevaluation

Previous RCA

Reinforce Squirrell Cultural Action from Previous RCA

- Schedule DR Scenario Review – JTF

- Schedule failure Analysis for TIM funds – JTF

Potential Loss To Tim Group

SLA for 99% uptime

Client Asked for SLA in contract

Client's Business Negatively impacted by Site Downtime

Loss of Reputation

Lost Weekend Time !

Woken up by Duplicate Primary Key Alerts

Replication Stopped

Master DB was a Week out of Sync

Database is Renamed on Replication ☹ (extra Bad)

Shared Slaves / Multiple-Purposed

Wanted to Replicate Bunker Setup

Replicated the Bunker Setup

RCA 16/02/12

Not enough Hosts To Run Slaves as Single-Purposed

Can't use Virtualization

Current firewall configuration Makes implementing Virtualization Too expensive

Ask Squirrell to Explain the Business function Accounts for Human Costs – JTF

Team Prioritized Work Disrupted !

Monday Morning Stress

failed to DR

Single Point of failure at Database

No Stand-by Datab for TIM funds

Did Not Consider Impact of outages On the Team

We Relied on Bunker 24/7 Support to Handle Some of the out-of-Hours Load.

Culturally we Have Dismissive of the Hu Costs.

Value Customers over employees

We still Have Start-up Values

founders No longer in the Line-of-fire

We only feel the pa When the People we Lean on are Gone.

We would Rathe Generate Revenue Save Costs.

We don't consider the of Removing friction

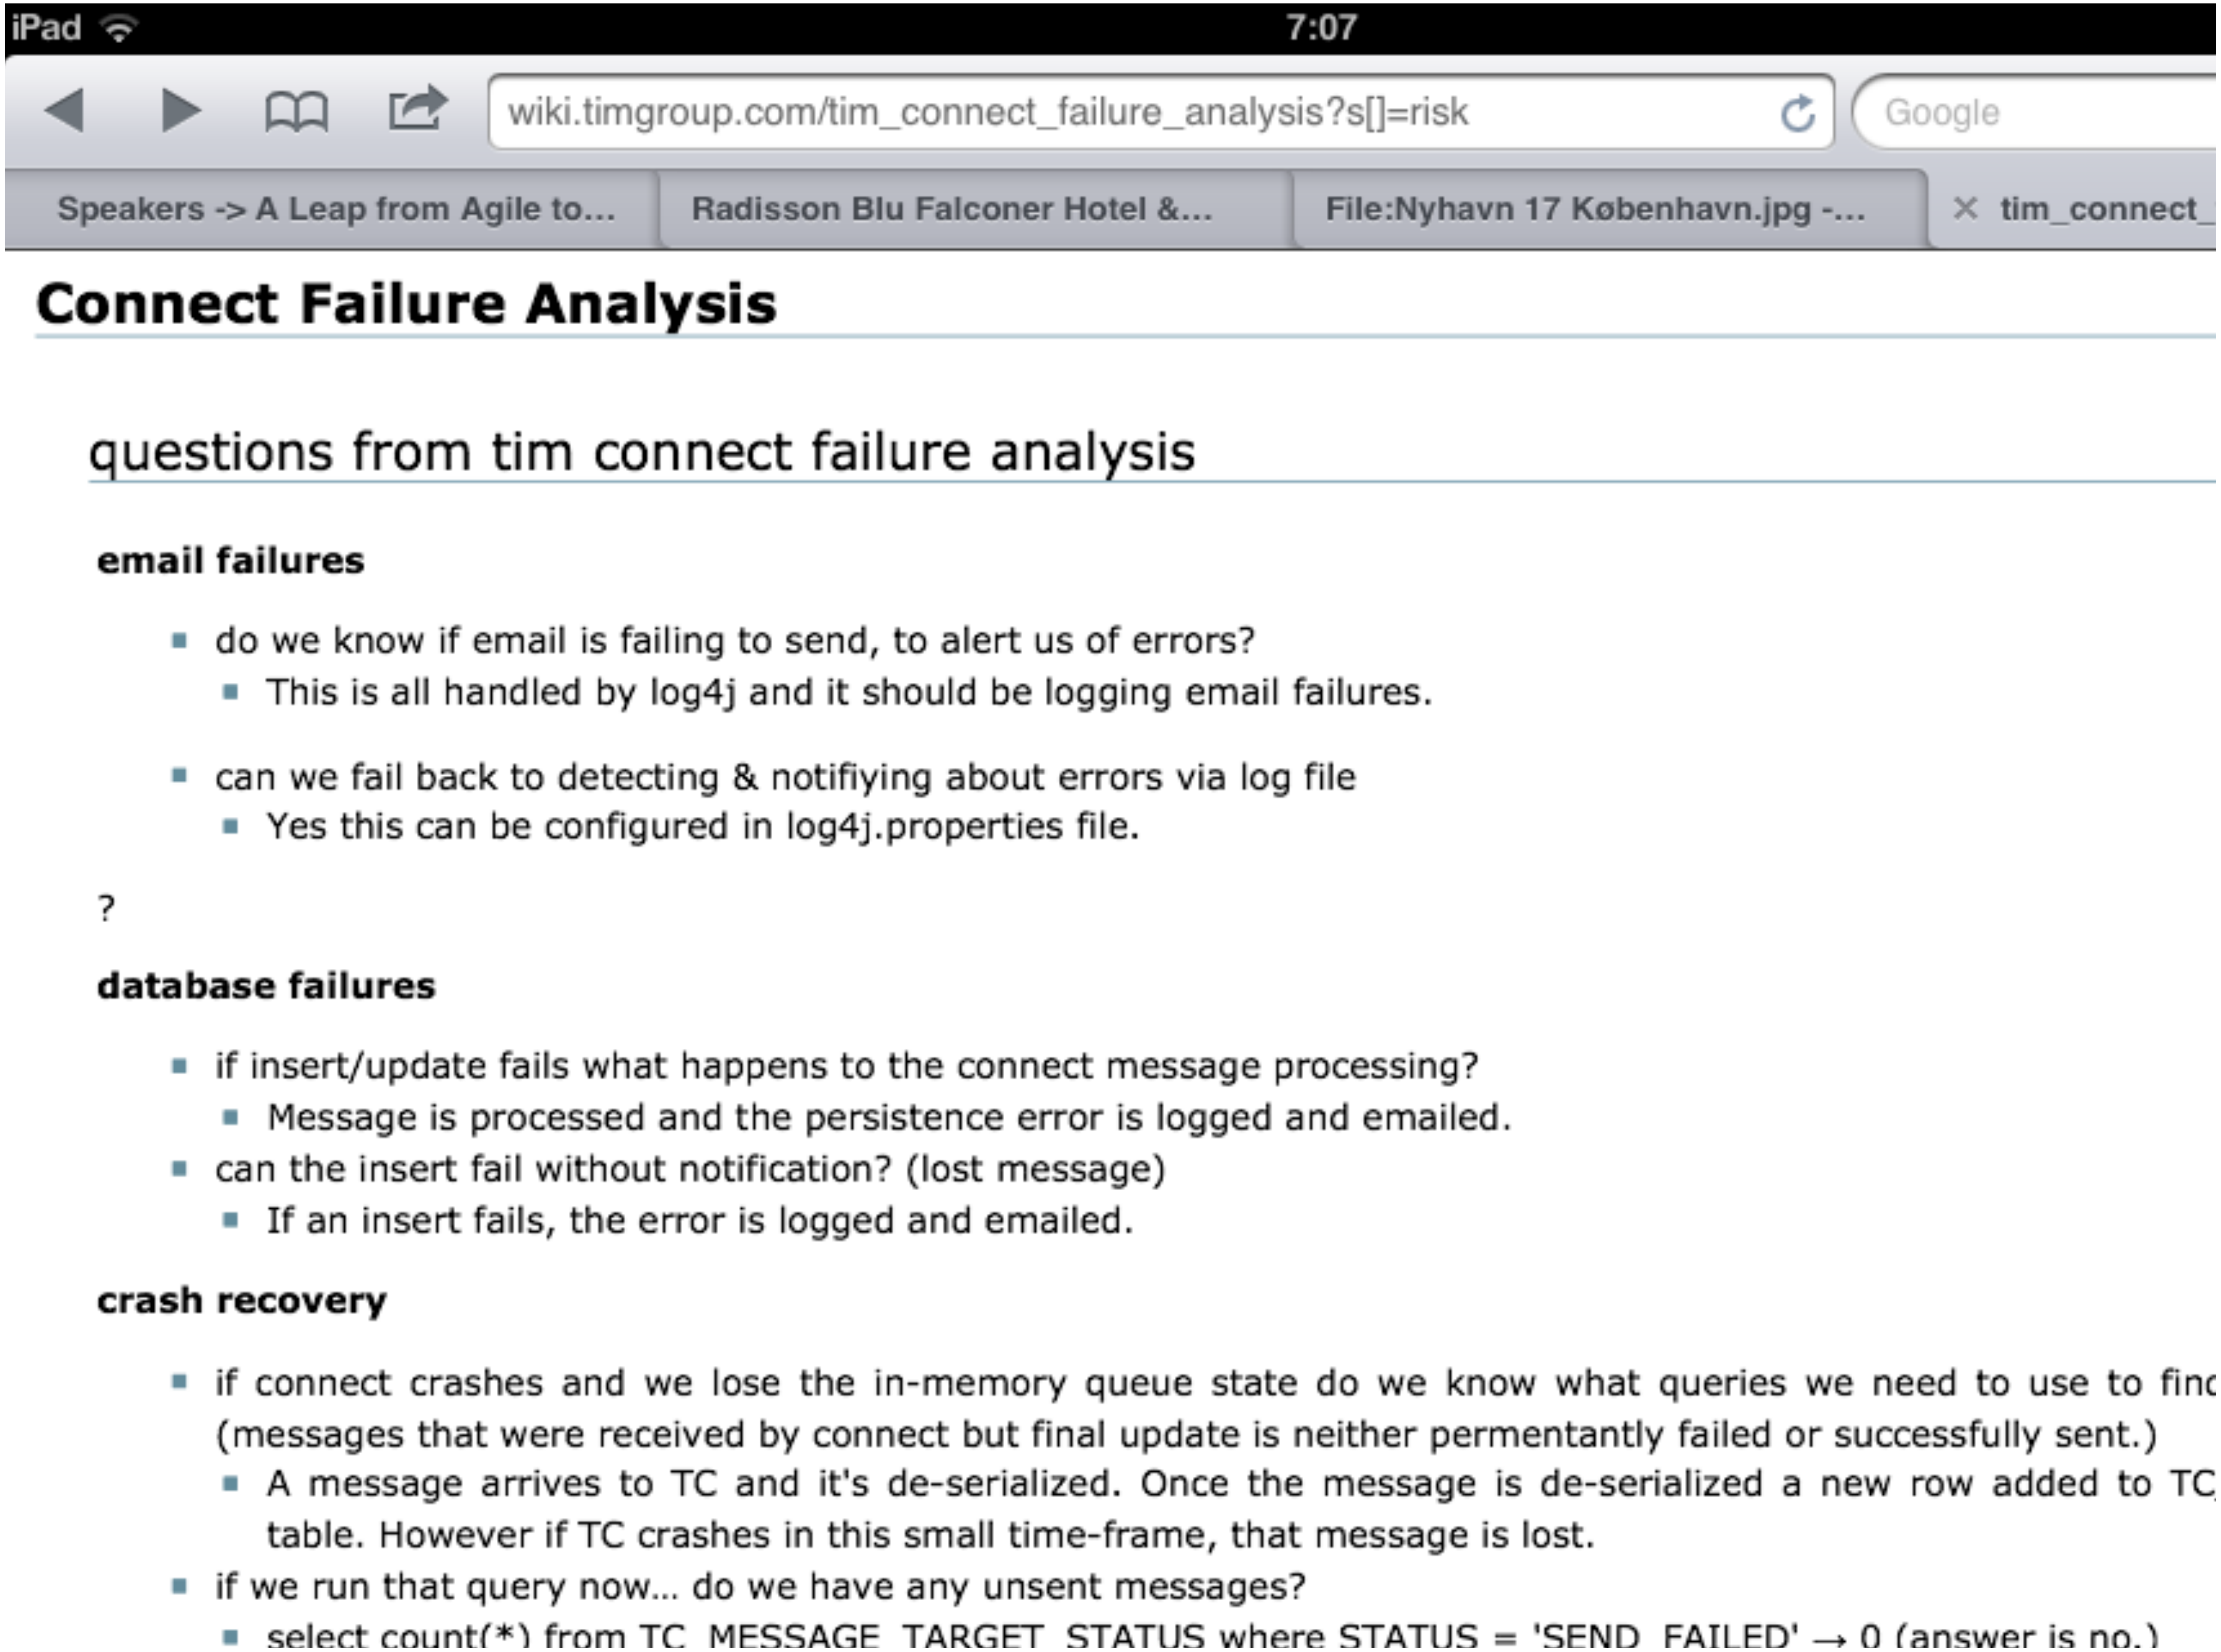
iPad
7:07
wiki.timgroup.com/tim_connect_failure_analysis?s[]=risk
Google
Speakers -> A Leap from Agile to...
Radisson Blu Falconer Hotel &...
File:Nyhavn 17 København.jpg -...
tim_connect_

Connect Failure Analysis

questions from tim connect failure analysis

email failures

- do we know if email is failing to send, to alert us of errors?
  - This is all handled by log4j and it should be logging email failures.

- can we fail back to detecting & notifiying about errors via log file
  - Yes this can be configured in log4j.properties file.

?

database failures

- if insert/update fails what happens to the connect message processing?
  - Message is processed and the persistence error is logged and emailed.
- can the insert fail without notification? (lost message)
  - If an insert fails, the error is logged and emailed.

crash recovery

- if connect crashes and we lose the in-memory queue state do we know what queries we need to use to finc (messages that were received by connect but final update is neither permentantly failed or successfully sent.)
  - A message arrives to TC and it's de-serialized. Once the message is de-serialized a new row added to TC table. However if TC crashes in this small time-frame, that message is lost.
- if we run that query now... do we have any unsent messages?
  - select count(*) from TC_MESSAGE_TARGET_STATUS where STATUS = 'SEND_FAILED' → 0 (answer is no.)

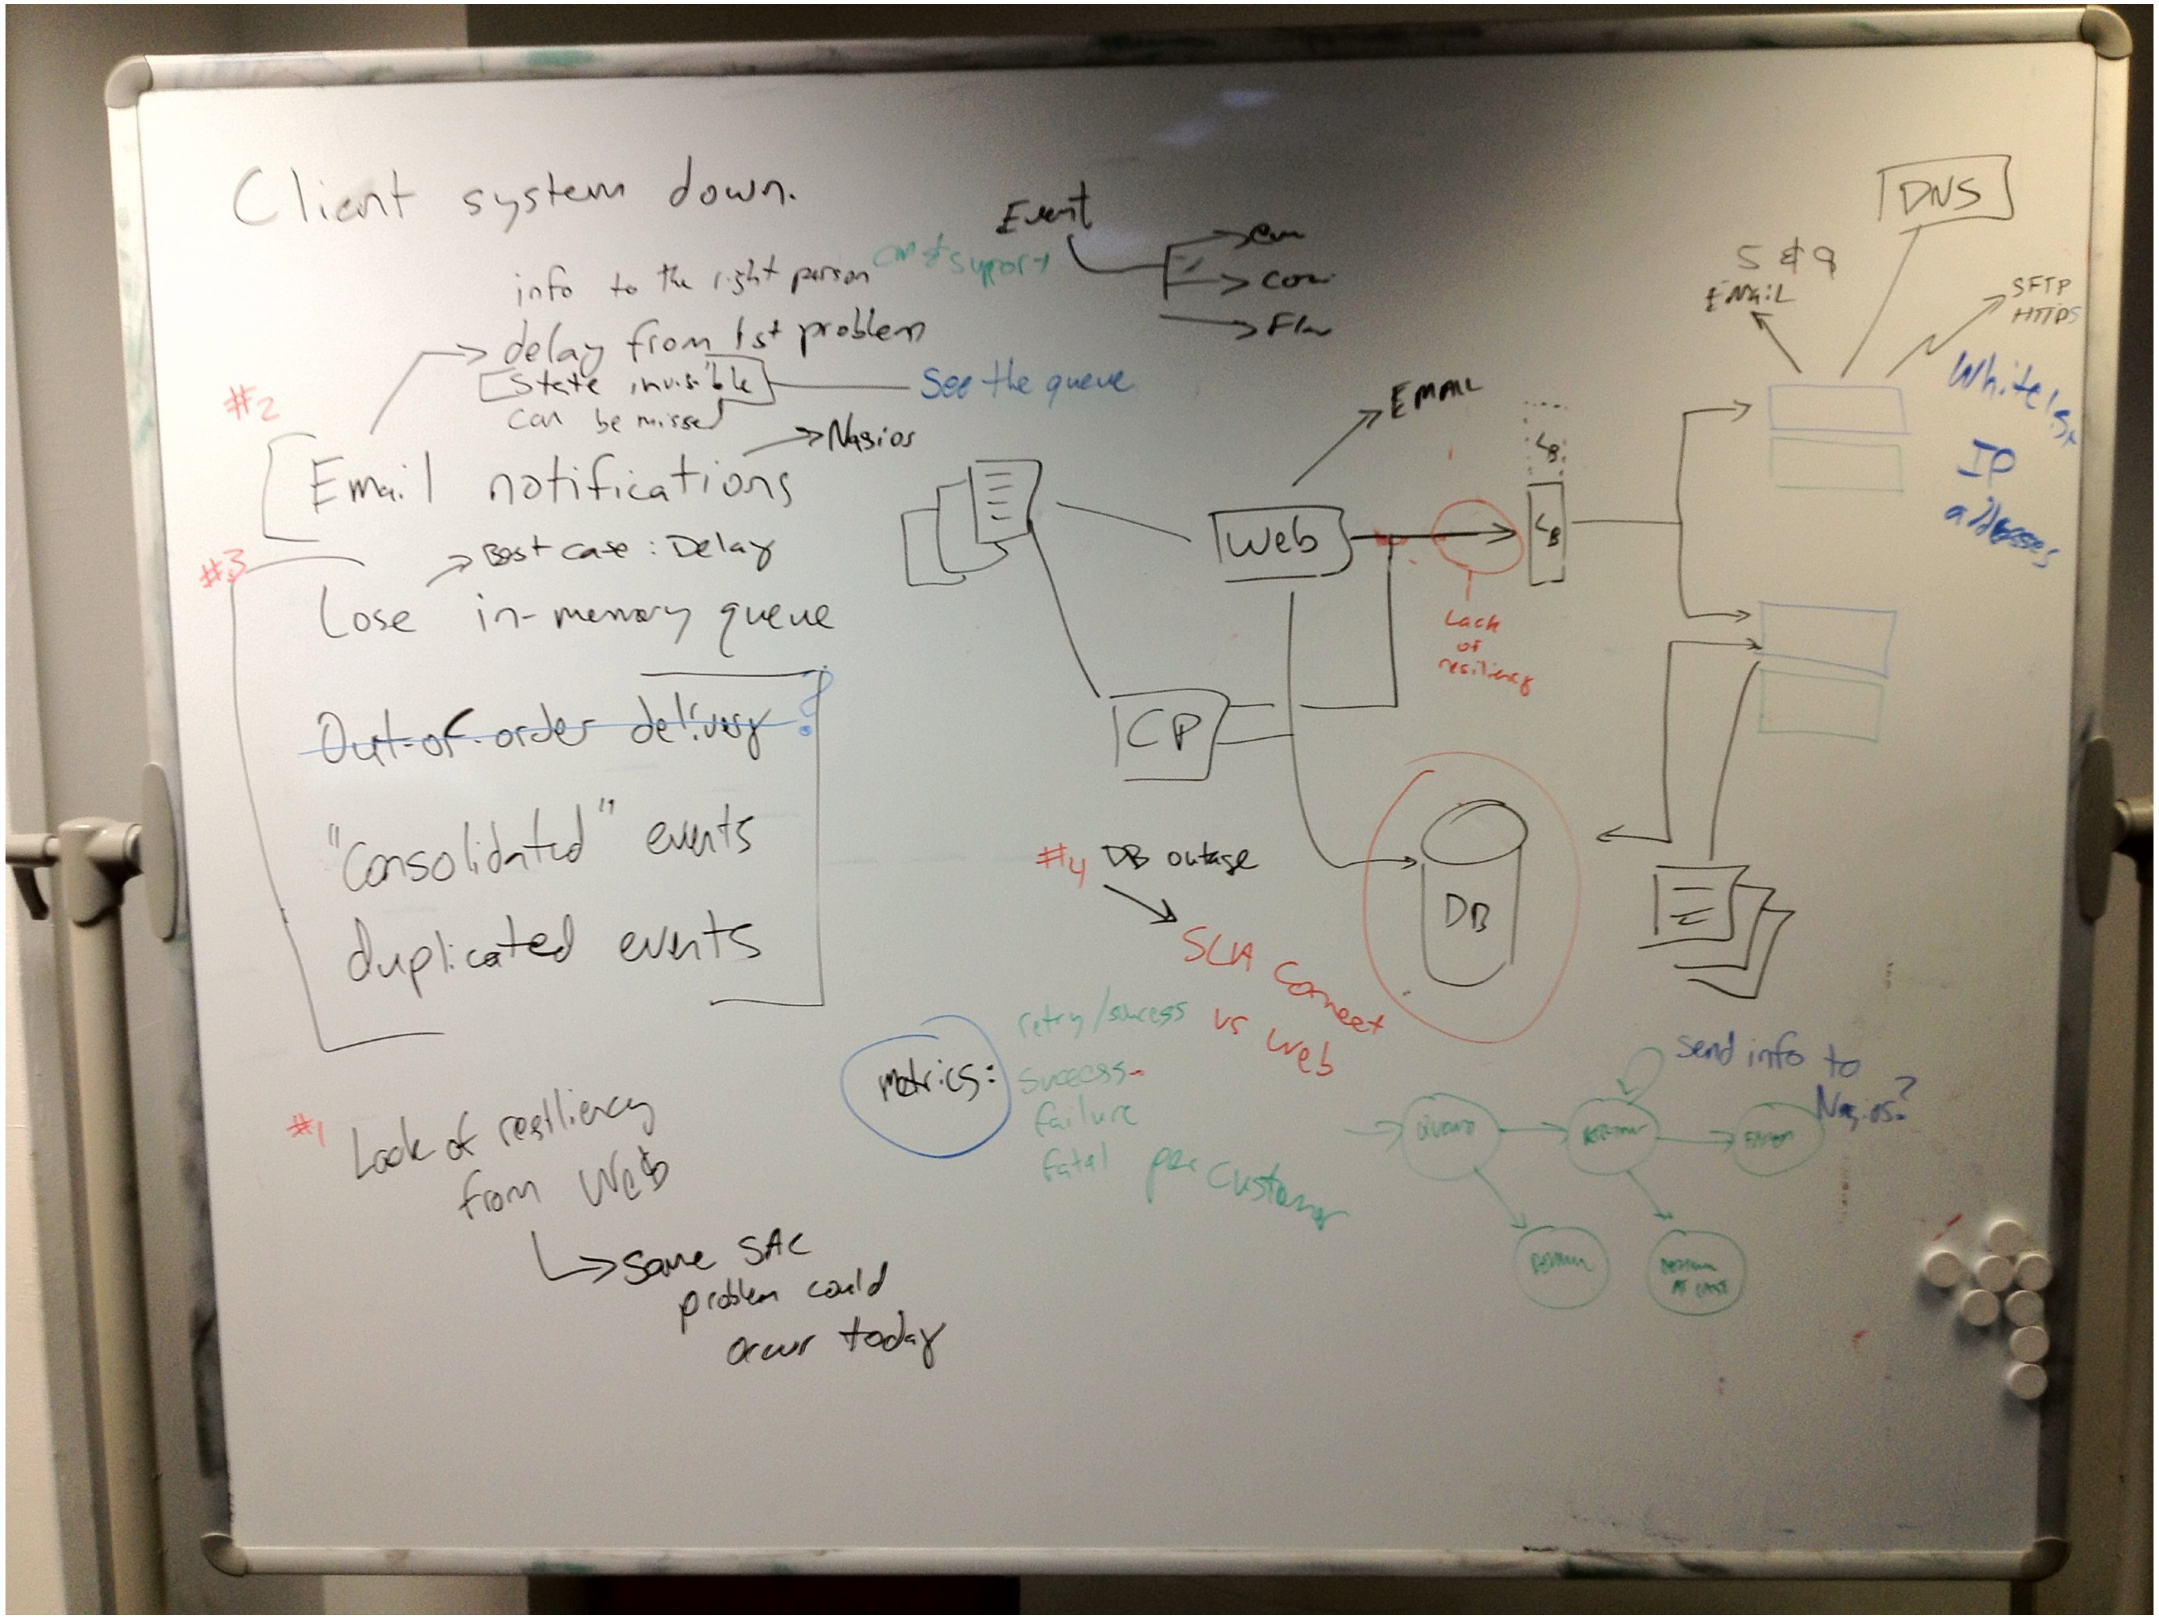

Client system down.

info to the right person

delay from 1st problem

State invisible

can be missed

#2 Email notifications

#3 Lose in-memory queue

Best case : Delay

~~Out-of-order delivery~~ ?

"Consolidated" events

duplicated events

#1 Lack of resiliency from Web

Same SAC problem could occur today

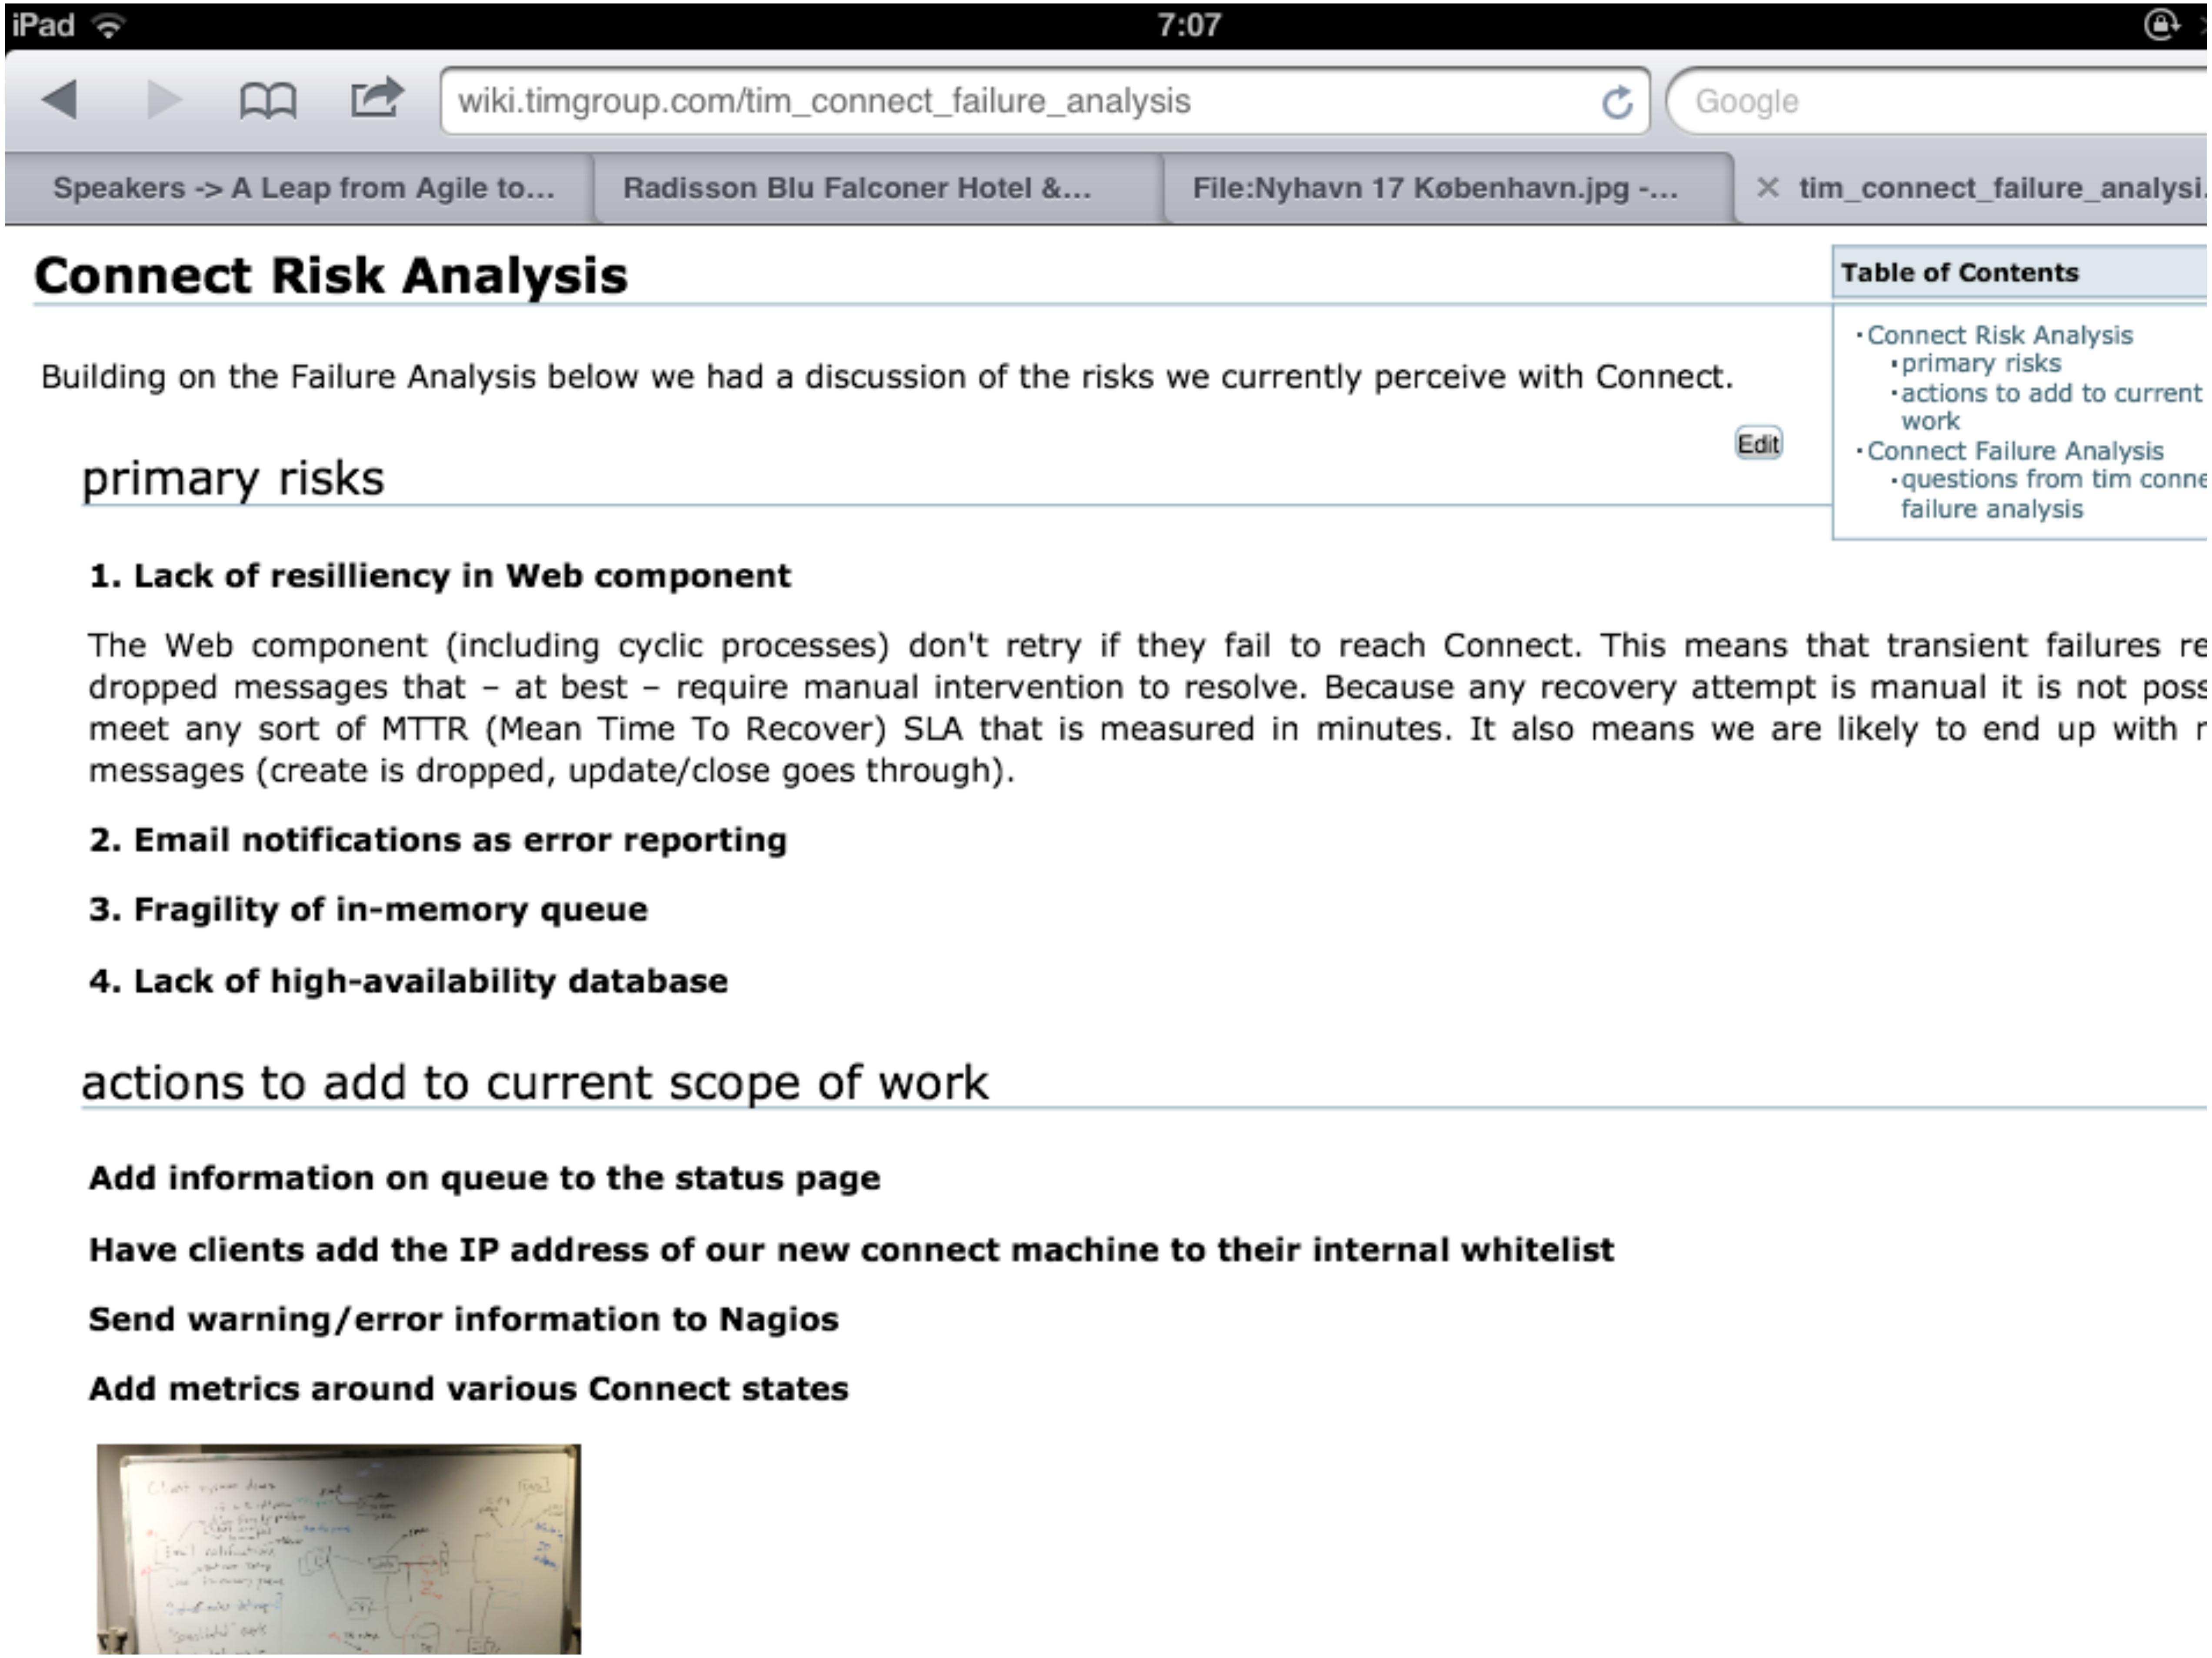
iPad
7:07
wiki.timgroup.com/tim_connect_failure_analysis
Google
Speakers -> A Leap from Agile to...
Radisson Blu Falconer Hotel &...
File:Nyhavn 17 København.jpg -...
tim_connect_failure_analysi
Edit

Connect Risk Analysis

Building on the Failure Analysis below we had a discussion of the risks we currently perceive with Connect.

primary risks

1. Lack of resilliency in Web component

The Web component (including cyclic processes) don't retry if they fail to reach Connect. This means that transient failures re dropped messages that – at best – require manual intervention to resolve. Because any recovery attempt is manual it is not poss meet any sort of MTTR (Mean Time To Recover) SLA that is measured in minutes. It also means we are likely to end up with r messages (create is dropped, update/close goes through).

2. Email notifications as error reporting

3. Fragility of in-memory queue

4. Lack of high-availability database

actions to add to current scope of work

Add information on queue to the status page

Have clients add the IP address of our new connect machine to their internal whitelist

Send warning/error information to Nagios

Add metrics around various Connect states

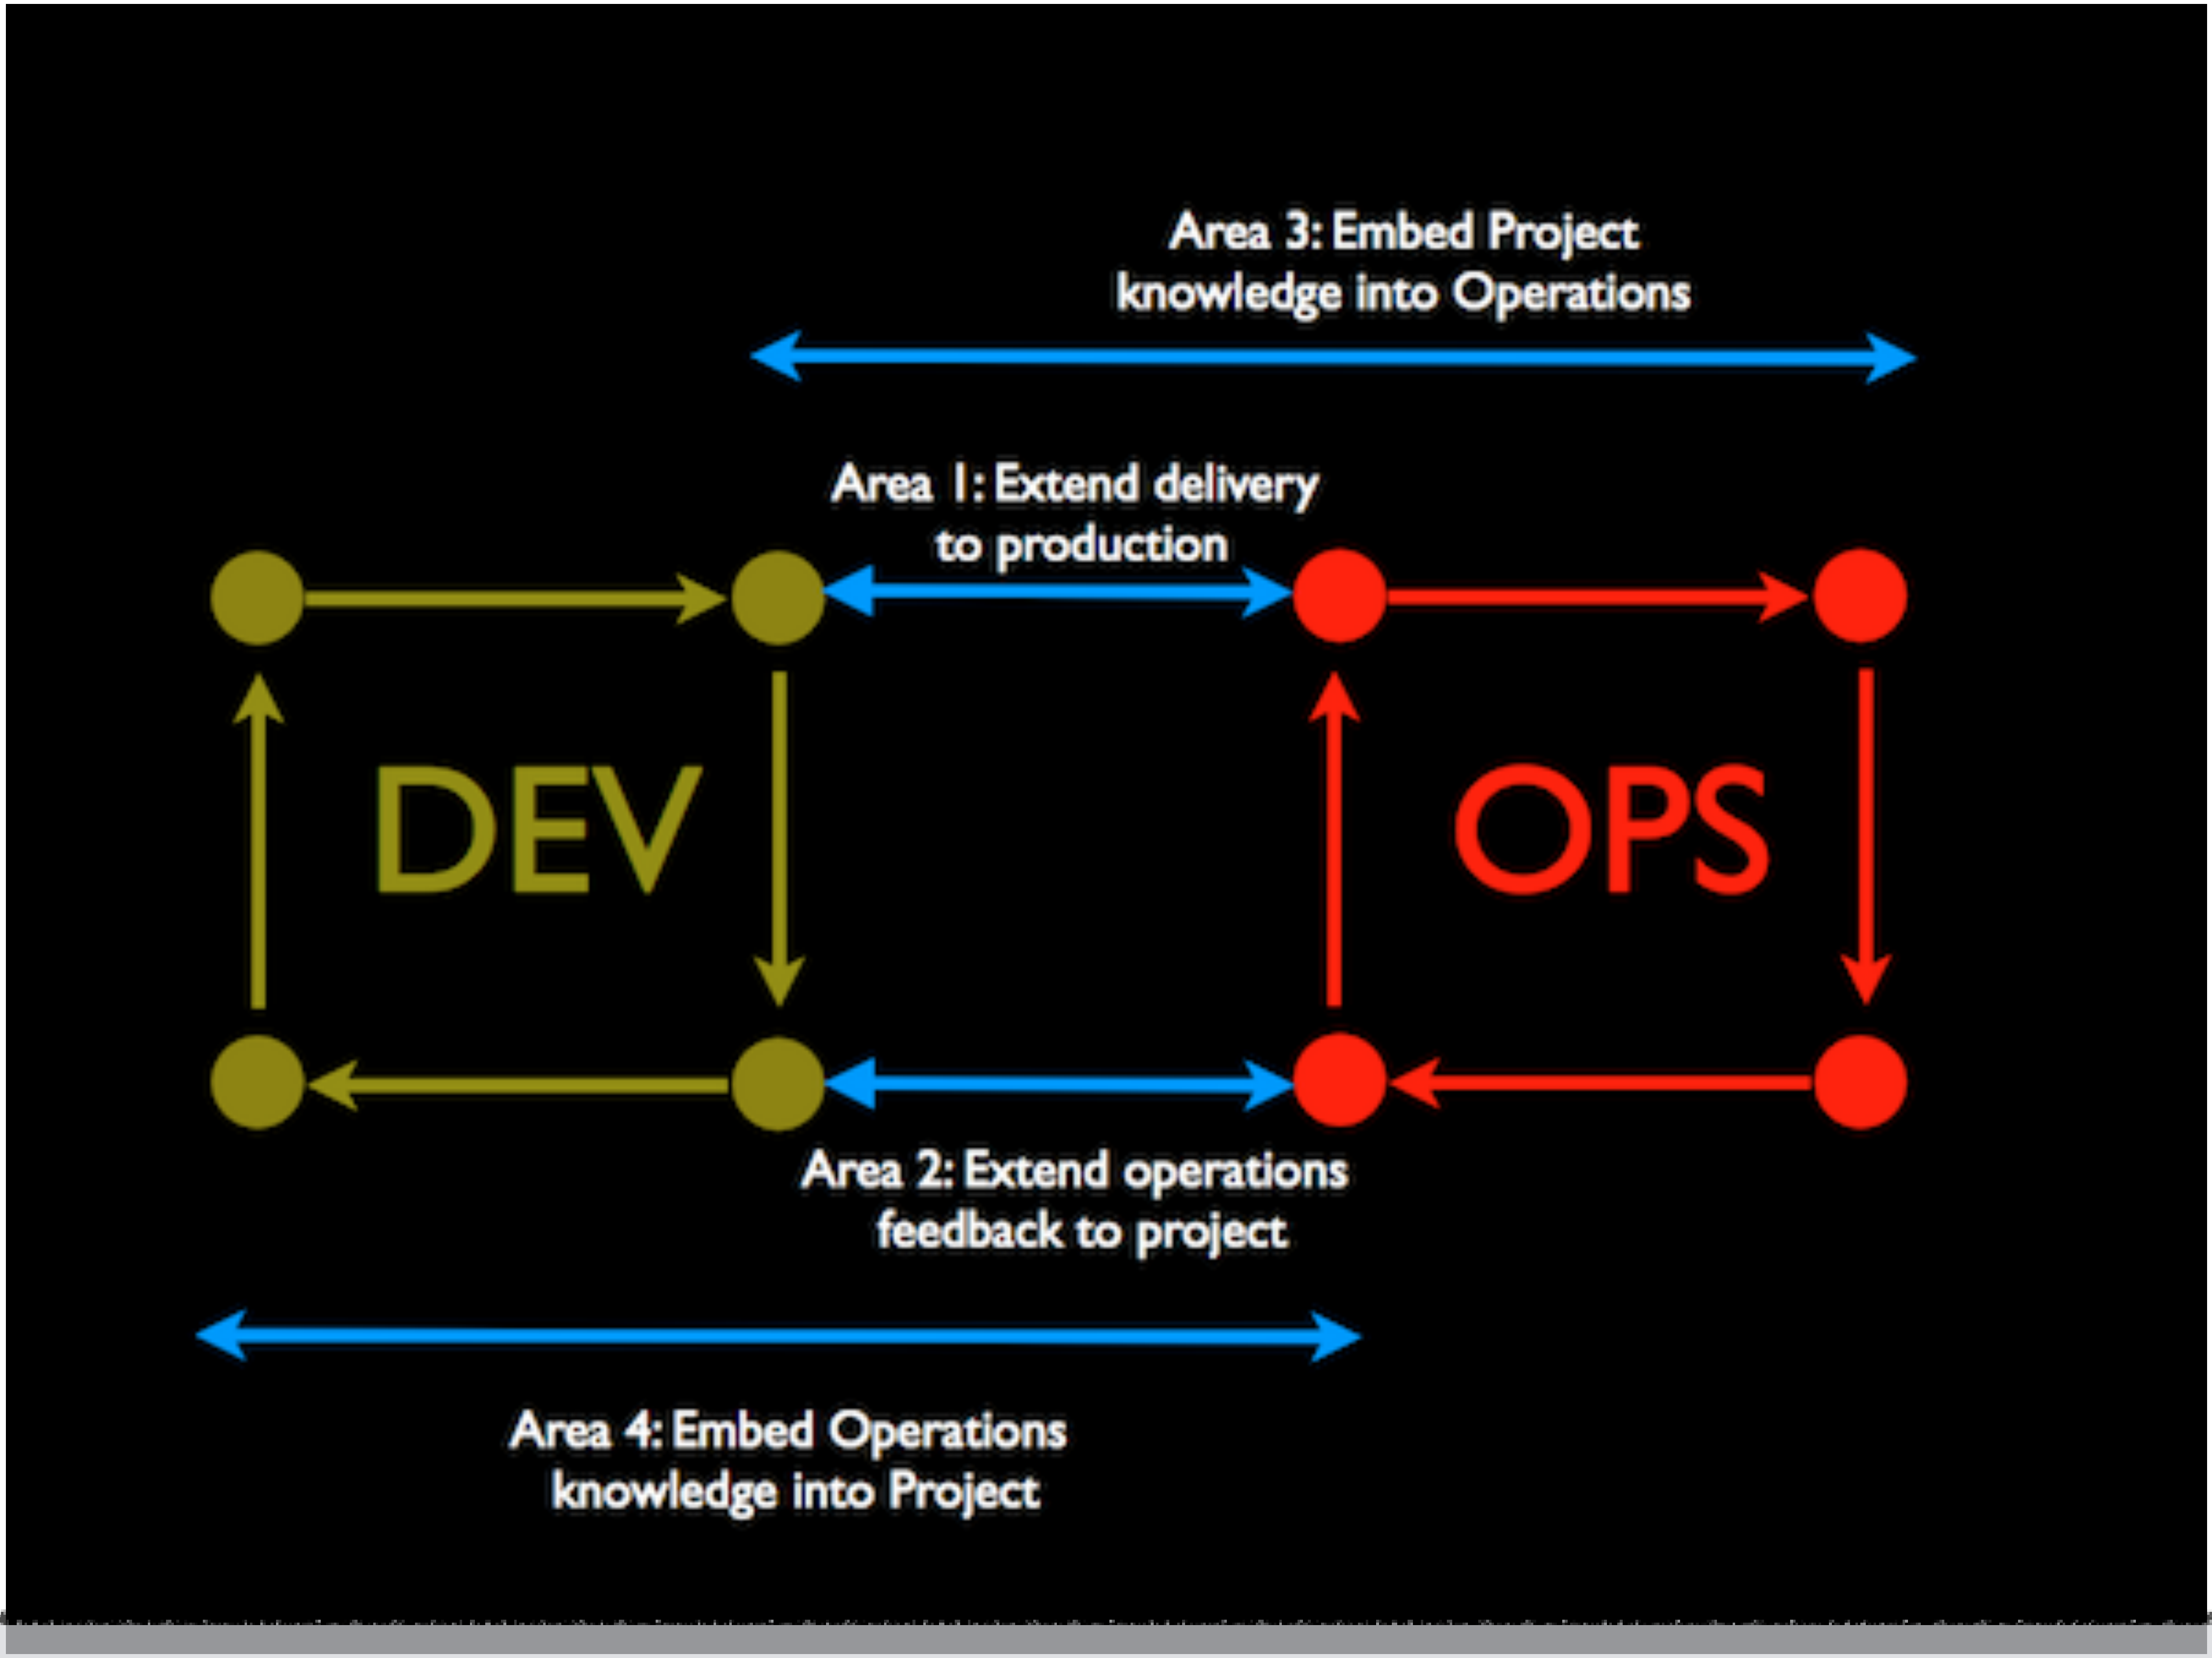

Area 3: Embed Project knowledge into Operations

Area 1: Extend delivery to production

DEV

OPS

Area 2: Extend operations feedback to project

Area 4: Embed Operations knowledge into Project

Patrick Debois: Devops Areas - Codifying devops practices
http://is.gd/IyicB3

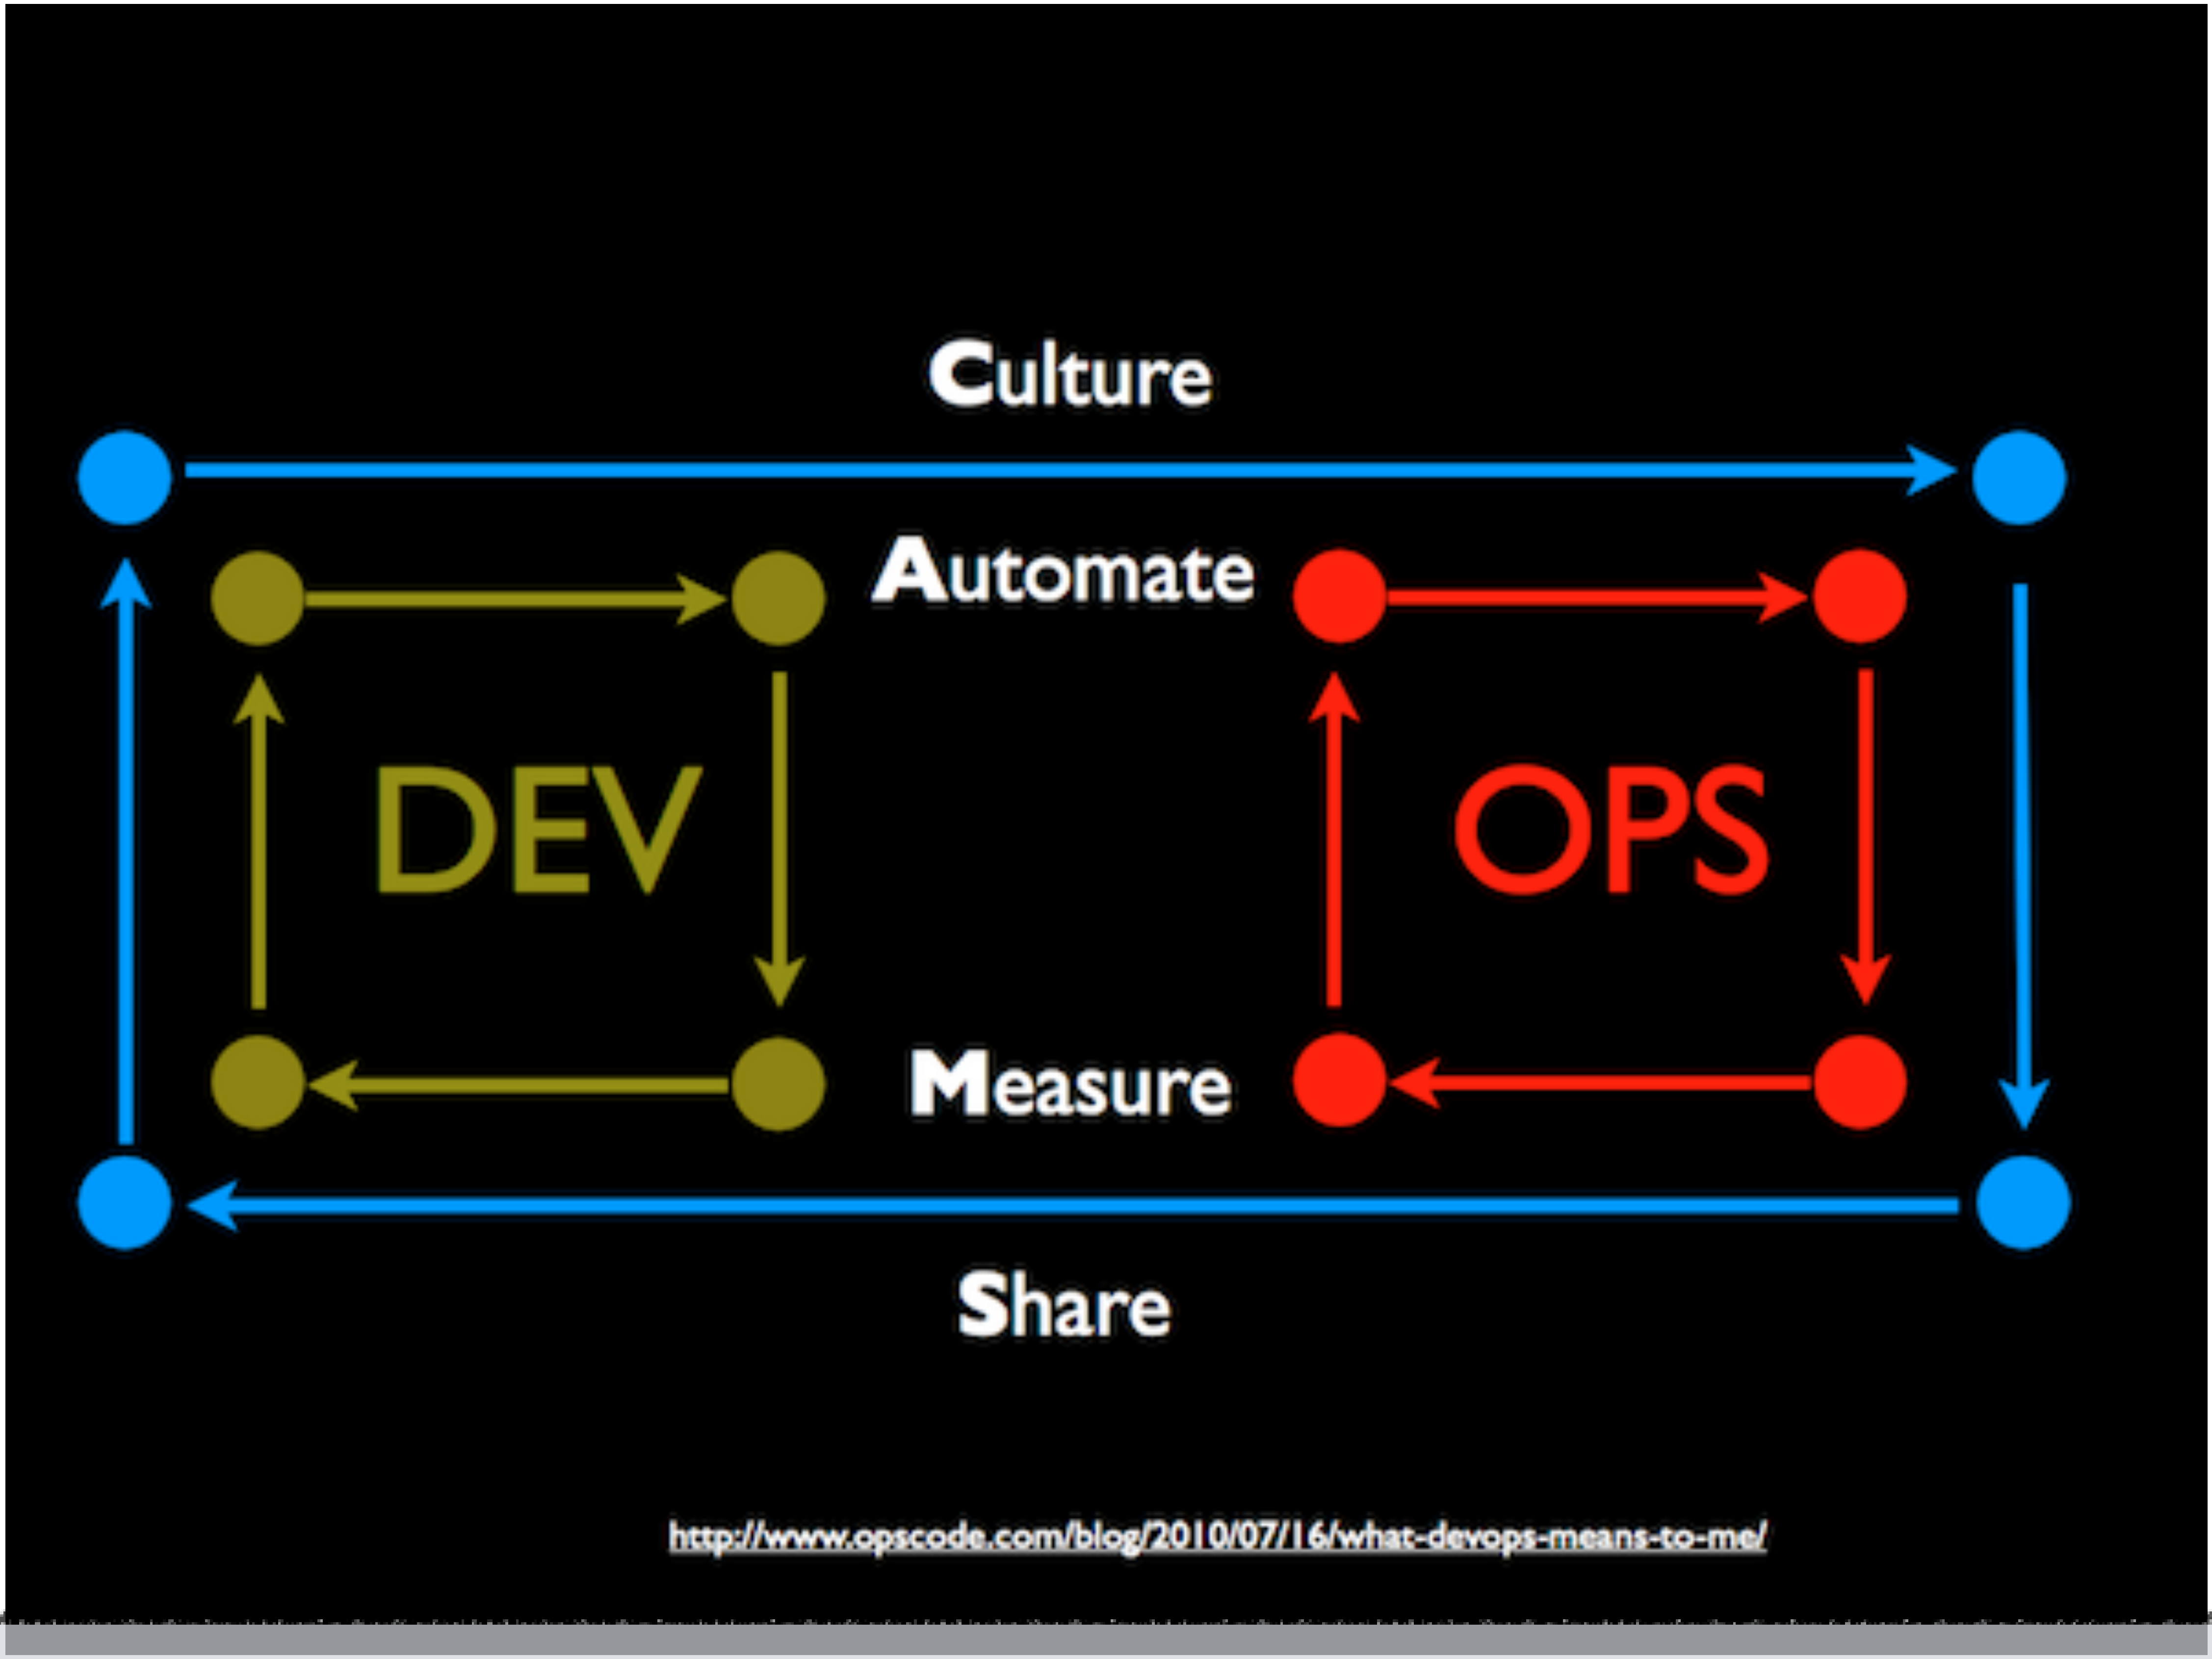

http://www.opscode.com/blog/2010/07/16/what-devops-means-to-me/

Patrick Debois: Devops Areas - Codifying devops practices
http://is.gd/IyicB3

TIM-IE6-serial
TIM-CodeCheck
TIMFlow
TIMToMercTesting
TIMToPortalTesting
TIM-Unit
TIM-JavaScript
TIM-IE6-fast
TIM-integration
TIM-browser-performance-IE6
TIMToTIMFlowTesting
TIM-compile
SAMSUNG

46p

Production Nagios Dashboard - Mozilla Firefox
Last Checked
Host
Status Info
Service
11:11:42
australia.pgldn.youdevise.com
apt-cacher-ng
Idea_Creation_By_Region
Mixed Paper
Yes please
No thank you
1237191

2012

Questions?

jtf@jeffreyfredrick.com

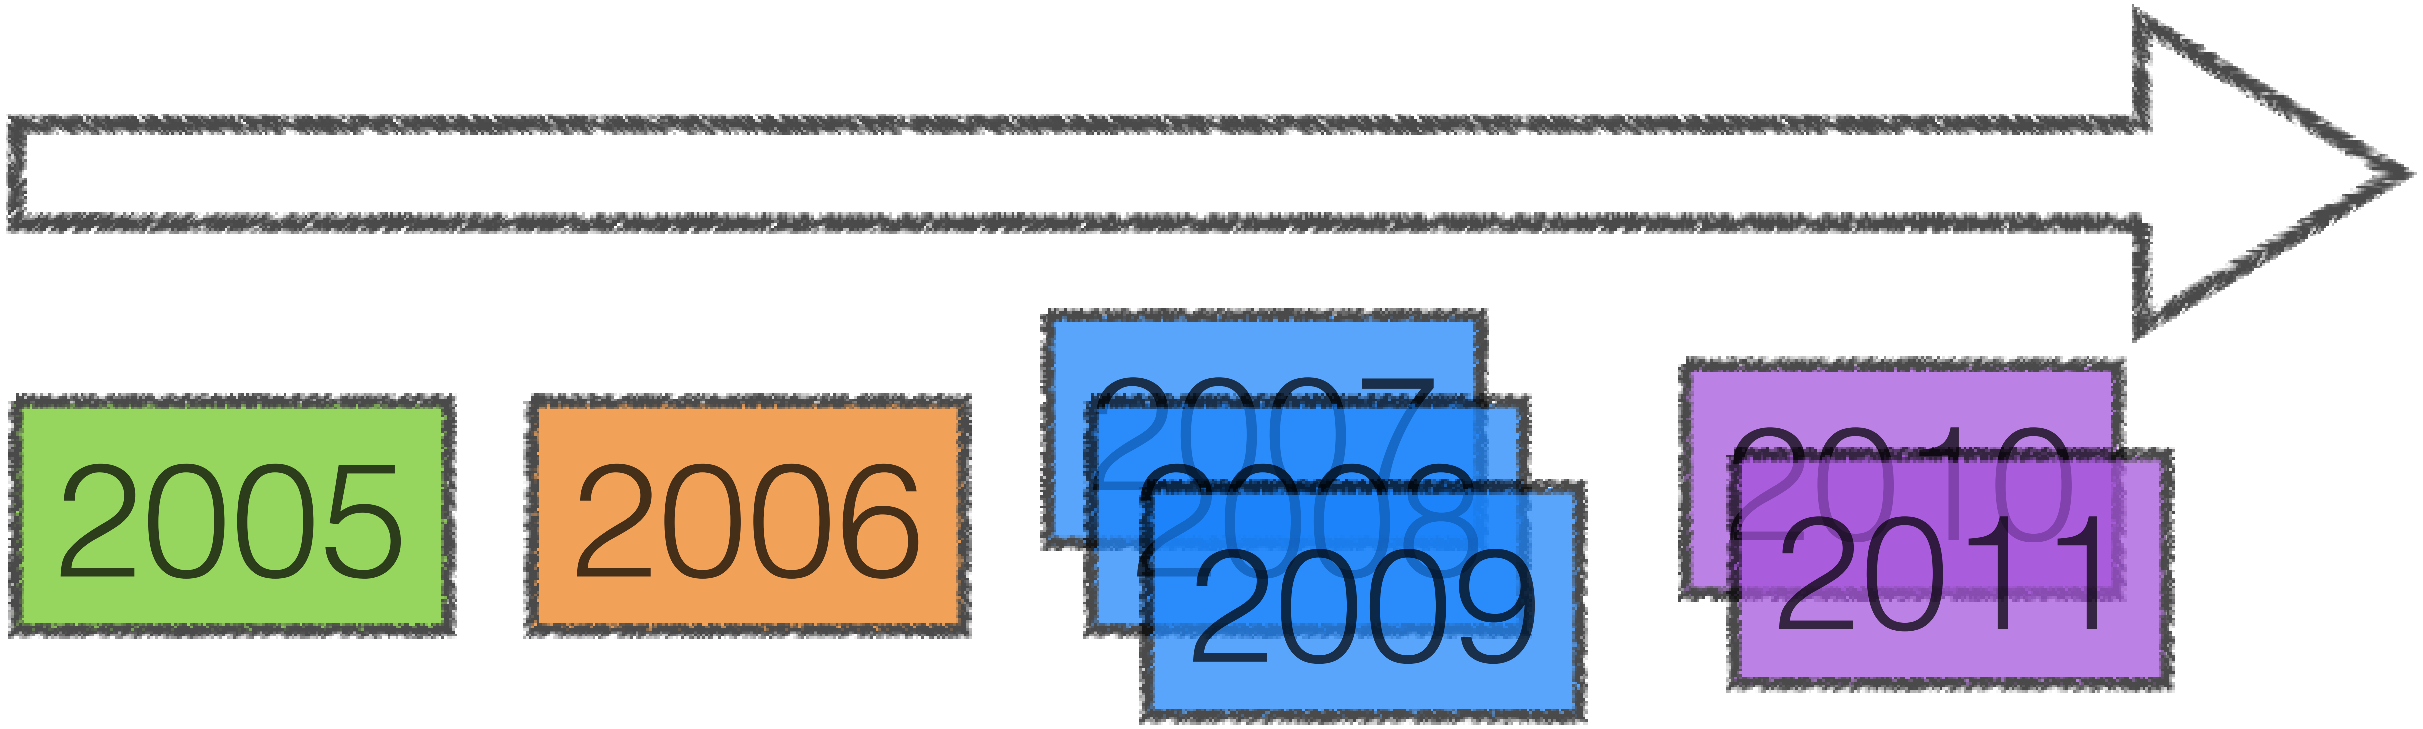

2005

2006

2007

2008

2009

2010

2011

2012

A Leap from Agile to DevOps

jtf@jeffreyfredrick.com

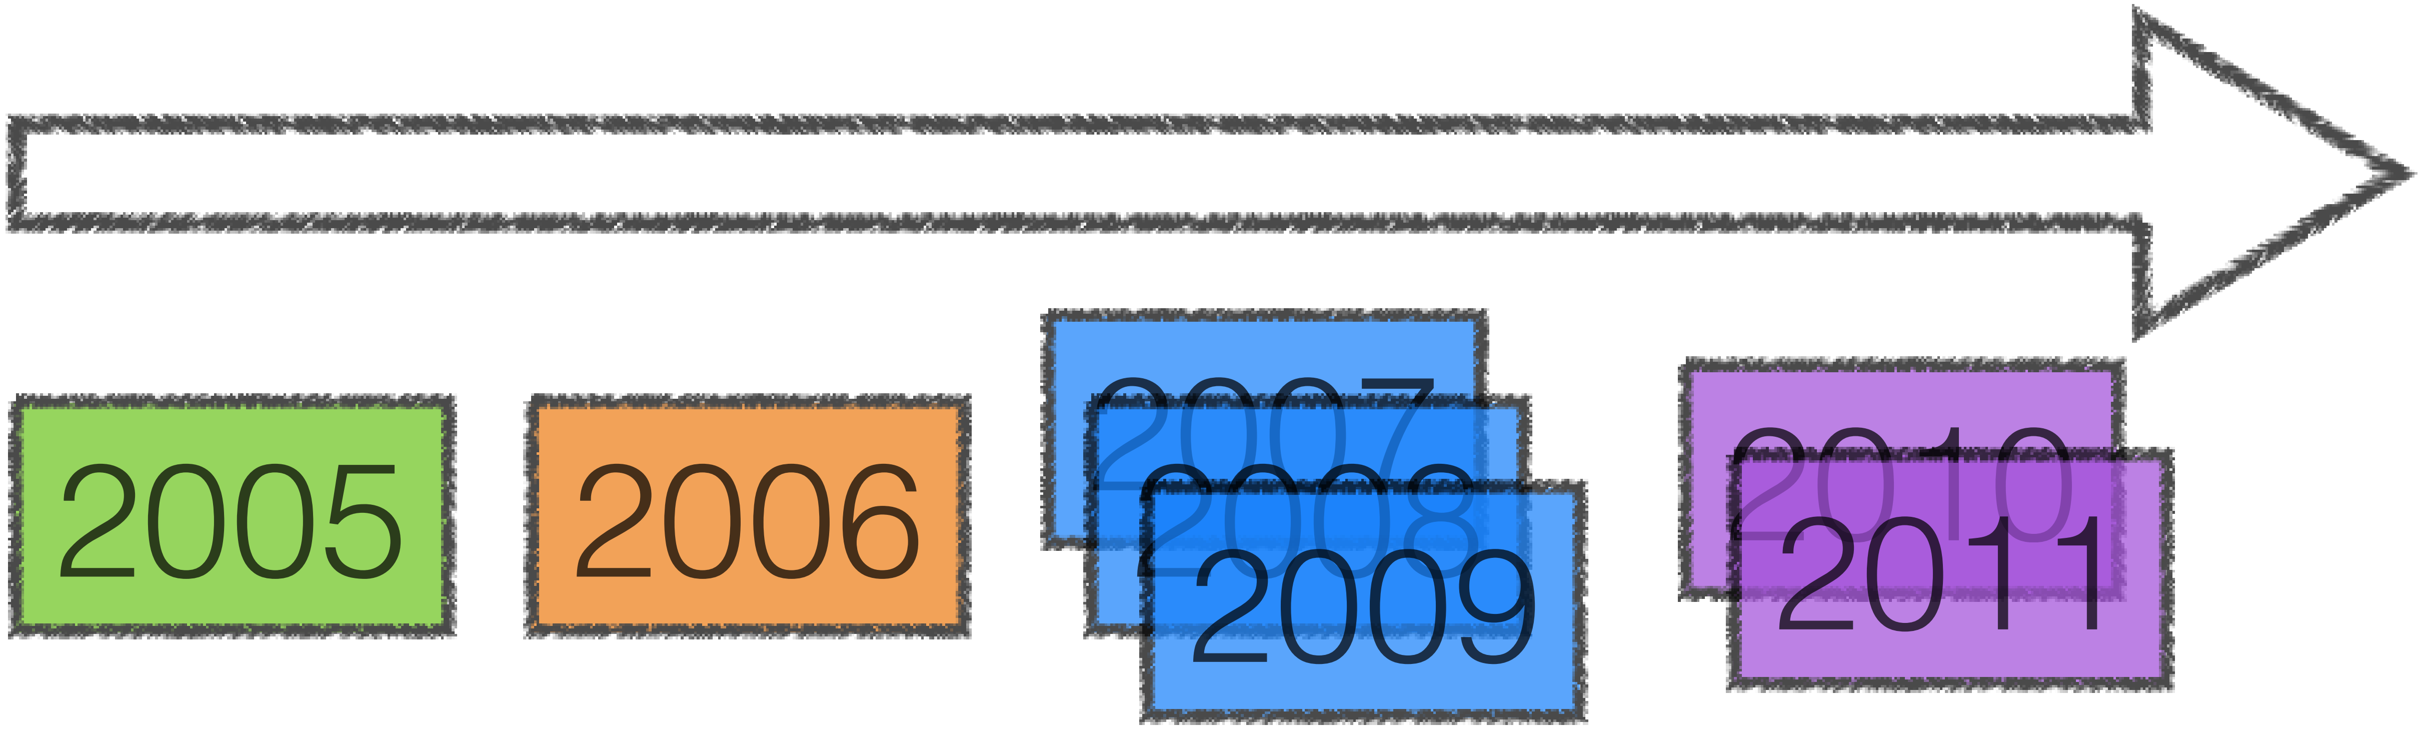
2005
2006
2007
2008
2009
2010
2011