

■ 9/1 DevOps企業實戰    ■ 9/2 DevOps技能精進    集思台大會議中心

# DEVOPS 2015

開發敏捷與維運高效的IT新典範

9/1 DevOps企業實戰

# DEVOPS 2015

開發敏捷與維運高效的IT新典範


## DevOps In Public Cloud Service

Story of 2PB data migration

# WHOAMI

陳彥宏 Eric Chen

- MediaTek (Mobile Phone Developer)
- Wiwynn/Wistron (Server System Engineer)
- Trendmicro (DevOps Engineer)
  - Process & Tool Chain
  - New technology evaluation and adoption
  - Capacity planning and performance measurement
  - AWS porting and huge data migration



[Eric\\_chen@trend.com.tw](mailto:Eric_chen@trend.com.tw)

[Pjack.chen@gmail.com](mailto:Pjack.chen@gmail.com)

## Agenda

- Overview
- Architecture Redesign
- Operational Works
- Migration Approach
- D-Day
- Summary

## Online Storage Delivery Platform

- File sync (iOS/Android, Mac/Windows)
- Web console
- RESTFul API
- WebDAV protocol



## Why Migration?

- Out of warranty (4+ years)
- TrendMicro Deep Security Platform in AWS



## The Scale

- 2 data centers, 1 in US, 1 in Europe
- 1,000,000,000 files
- 2PB data
- Over 1 year



**2PB = 1000 \* 2T disk**

**2PB / 100 Mbps = 2400 days**

**2 RD , 2QA, 2 Ops**



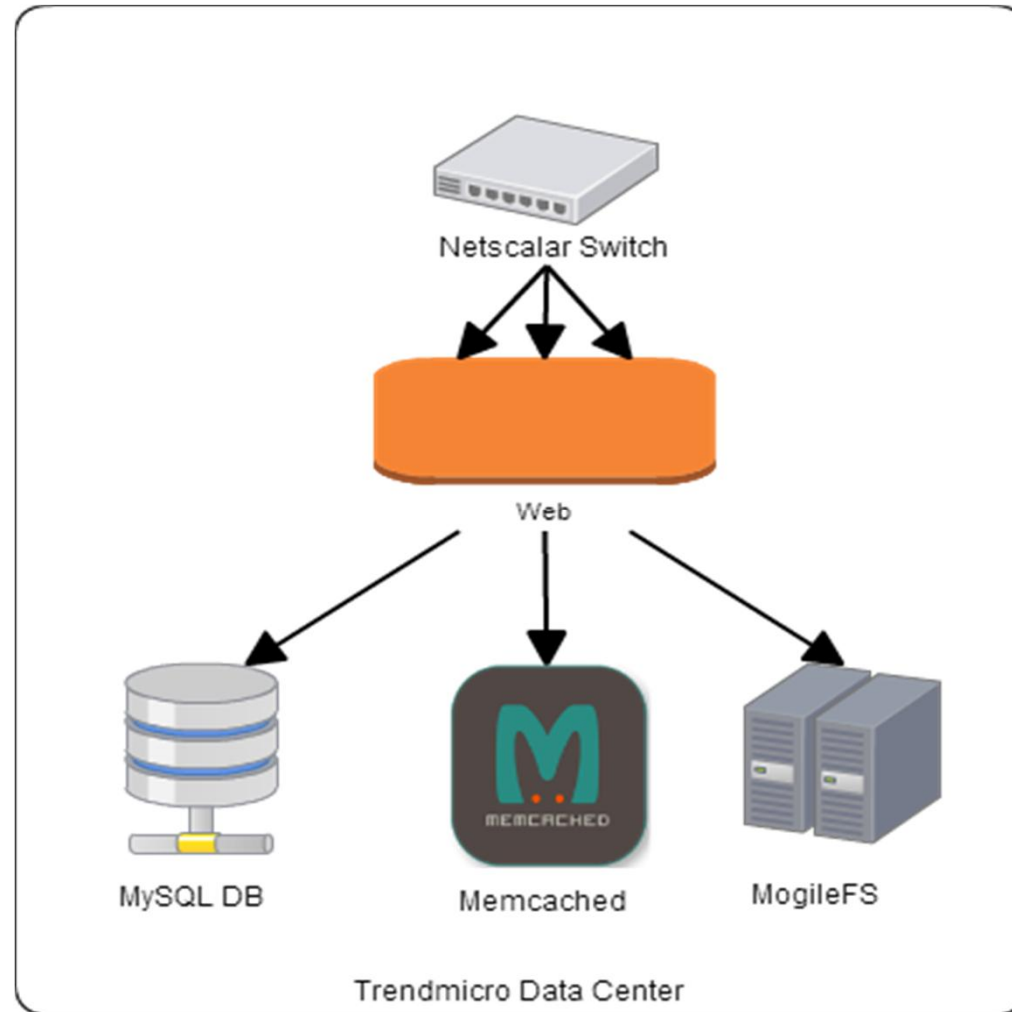


# DEVOPS 2015

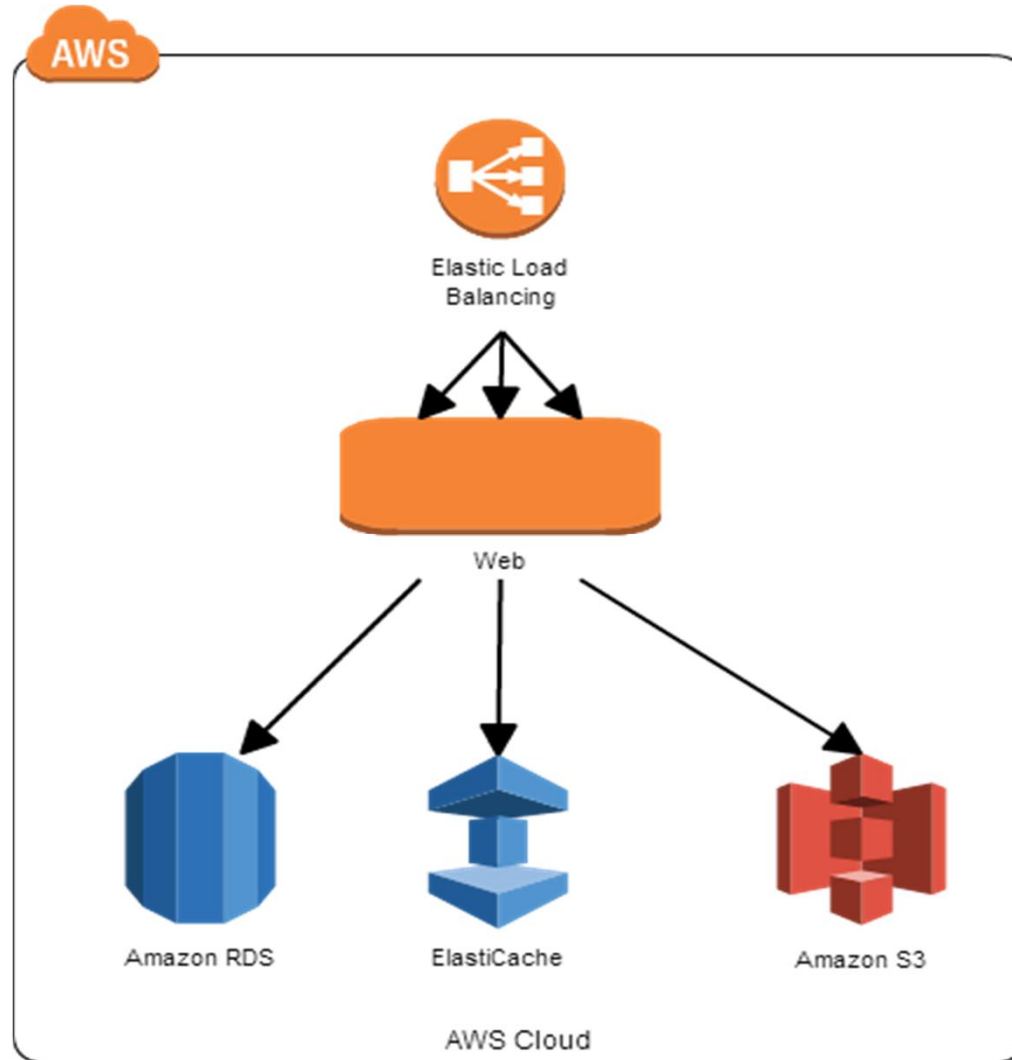
開發敏捷與維運高效的IT新典範

## Architecture Redesign

# Architecture – TrendMicro DC



# Architecture – AWS

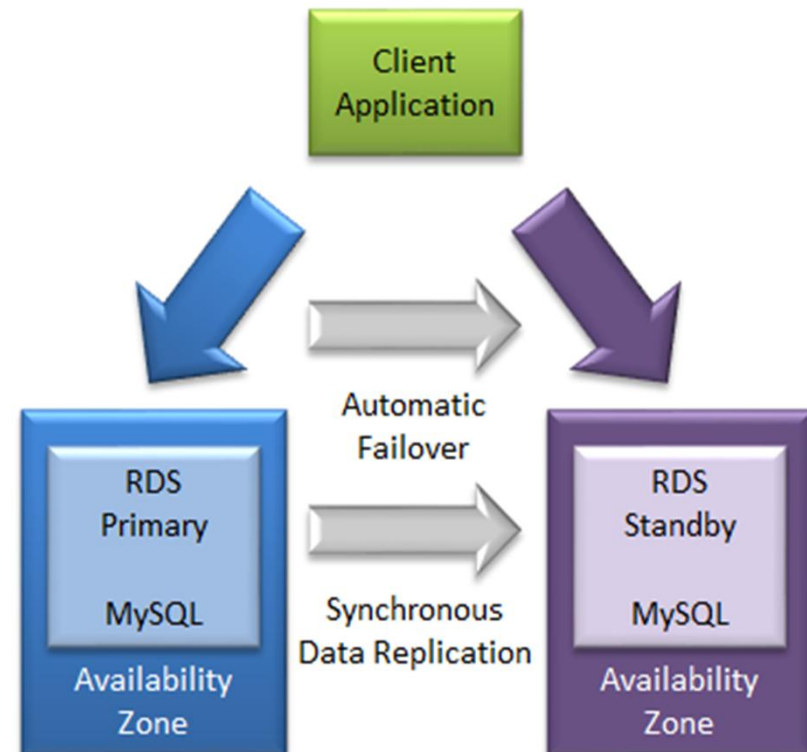


# S3 as Storage Backend

- Application and Operation impact
  - Access latency
  - Network bandwidth
  - Additional storage space and I/O loading
- Security review
  - IAM-Role instead of Access Key
  - Support VPC End-point (May, 2015)

# RDS Overview

- Compatible with MySQL
- Multi-AZ (HA)
- Backups
  - Daily snapshots
  - Point in time recovery
- SLA 99.95%



# RDS or Hosted MySQL?

Pros

## RDS

Easy Start  
Little Management  
Low Operation Effort

Cons

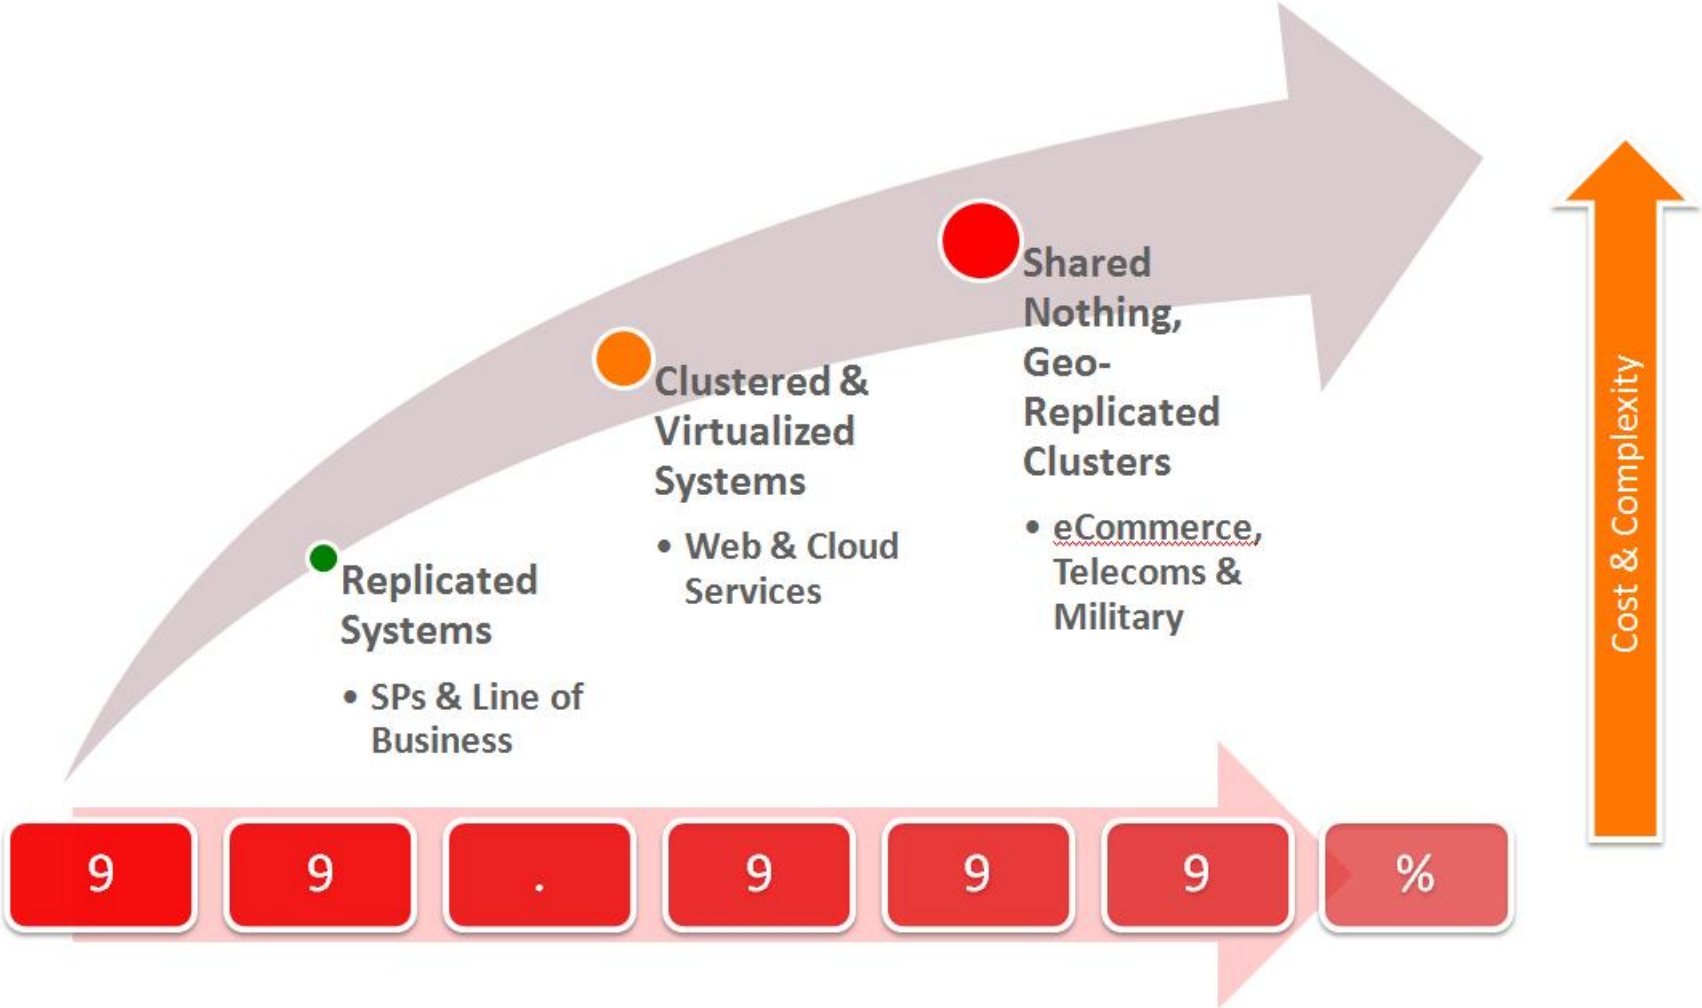
Longer Mgnt. Time  
Less Control  
Migration Compatibility

## Hosted MySQL

Full Control  
App Compatibility  
Familiar by Ops


HA Solution  
More Operation Effort

# SLA v.s. Cost and Complexity



# Hosted Mysql - HA Solution

- No L2 protocol allowed
  - Keepalived + multicast + VRRP **NOT** work
- Distributed HAproxy





# Architecture with HA

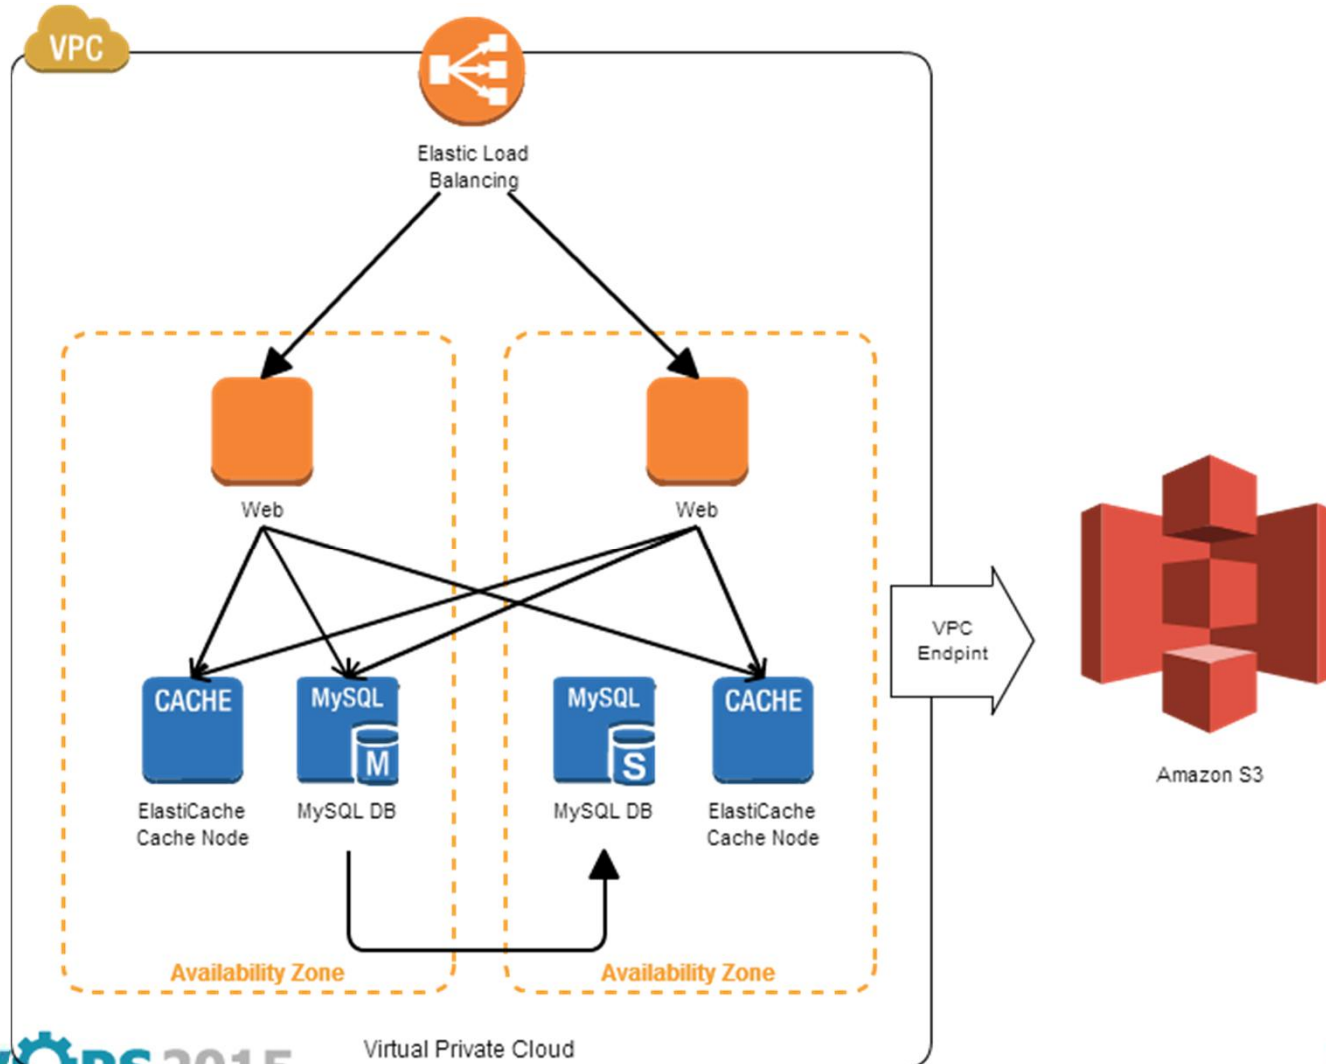



圖 9/1 DevOps 企業實戰

# DEVOPS 2015

開發敏捷與維運高效的IT新典範

## Operational Works



# Infra Monitoring & Metrics

- AWS component monitoring
- AWS CloudWatch integration
- Performance data collection (ex: RDS)

osdp-memc-proddc1.ldrqy5.0001	CPU_UTILIZATION	OK	2015-08-23 08:35:29	58d 15h 25m 10s	1/5	CPUUtilization = 3.0% OK
	MEMCACHE_HITRATE	OK	2015-08-23 08:34:49	58d 15h 25m 13s	1/5	MEMCACHED OK - OK, Hit checked: OK - at osdp-memc-proddc1.ldrqy5.0001.usw2.cache.amazonaws.com:11211
	MEMCACHE_RESPONSE	OK	2015-08-23 08:34:49	2d 17h 1m 26s	1/5	MEMCACHED OK - OK, Time checked: OK
osdp-memc-proddc1.ldrqy5.0002	CPU_UTILIZATION	OK	2015-08-23 08:35:02	13d 6h 42m 51s	1/5	CPUUtilization = 3.0% OK
	MEMCACHE_HITRATE	OK	2015-08-23 08:34:53	75d 5h 36m 13s	1/5	MEMCACHED OK - OK, Hit checked: OK - at osdp-memc-proddc1.ldrqy5.0002.usw2.cache.amazonaws.com:11211
	MEMCACHE_RESPONSE	OK	2015-08-23 08:34:50	1d 10h 56m 12s	1/5	MEMCACHED OK - OK, Time checked: OK
osdp-storage-elb-proddc1	ELB_CERTIFICATE	OK	2015-08-23 08:34:53	54d 13h 58m 6s	1/5	OK: 461 days remaining for 'storage.trendmicro.com'. Certificate Expires: 'Nov 26 14:54:01 2016 GMT'. Hash algorithm: sha256WithRSAEncryption
	ELB_HEALTHY_STATUS	OK	2015-08-23 08:36:00	1d 7h 16m 46s	1/5	OK:
	HEALTHY_HOST_COUNT	OK	2015-08-23 08:35:54	1d 7h 17m 4s	1/5	HealthyHostCount = 4.0 OK
	LATENCY	OK	2015-08-23 08:35:44	1d 7h 17m 19s	1/5	Latency = 0.20206005503s OK
	UNHEALTHY_HOST_COUNT	OK	2015-08-23 08:34:55	1d 1h 0m 10s	1/5	UnHealthyHostCount = 0.0 OK

# Backup & DR Drill

- Daily Backup
  - R53
  - EBS
  - Database
- DR Drill, Monkey Test
  - One node crash
  - One zone crash
  - One region crash
  - Database gone

# Capacity Planning (DB Node)

- Measure Point: Transactions Per Second
- Factors
  - Instance type
  - Disk I/O requirement
  - Read/Write workload
  - Database size
  - Concurrent connection number
  - Memory requirement

# Security Policy & Design

- Two-way authentication (MFA)
- Password policy
- Trust Advisor
- IAM-role
- No Public IP
- Encryption


## Others

- Log management
  - ELB log
  - S3 log
  - CloudTrail
  - CloudConfig
- DNS Zone migration
- Cost estimation

# DEVOPS 2015

開發敏捷與維運高效的IT新典範

## Data Migration






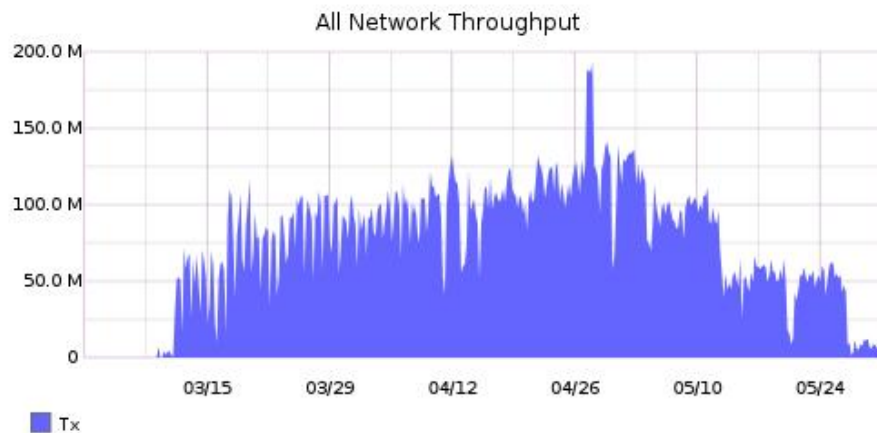
# How to migrate PB data?



# Migration Approach

- Disk health status
- Copy 99.9% files before D-day
- Replicate DB to AWS
- Read-Only mode for D-Day
- Shorter DNS TTL

Minimum Custom Impact



# D-DAY

---

NORMANDY 1944






## Time-out Period

- ELB warm-up
- DB warm-up
- People standby
- Final Healthy Check



## Operation (120 min)

- Deploy Read-Only code
- DB sync up
- Upload last batch file



## Verification (120 min)

- Sanity test
- Dog feeding


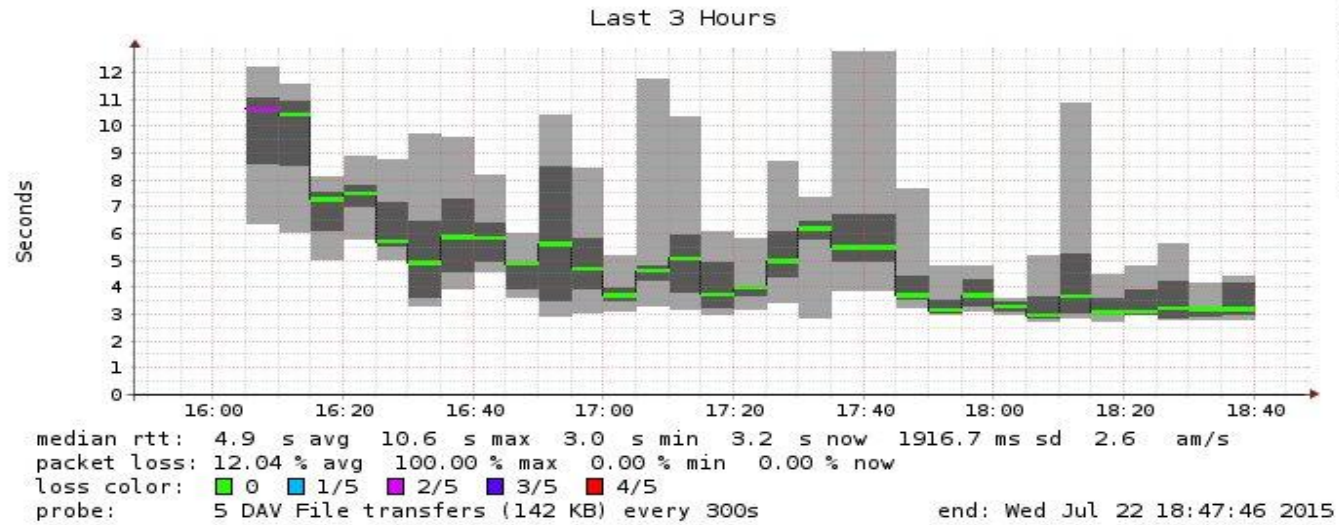


## Online

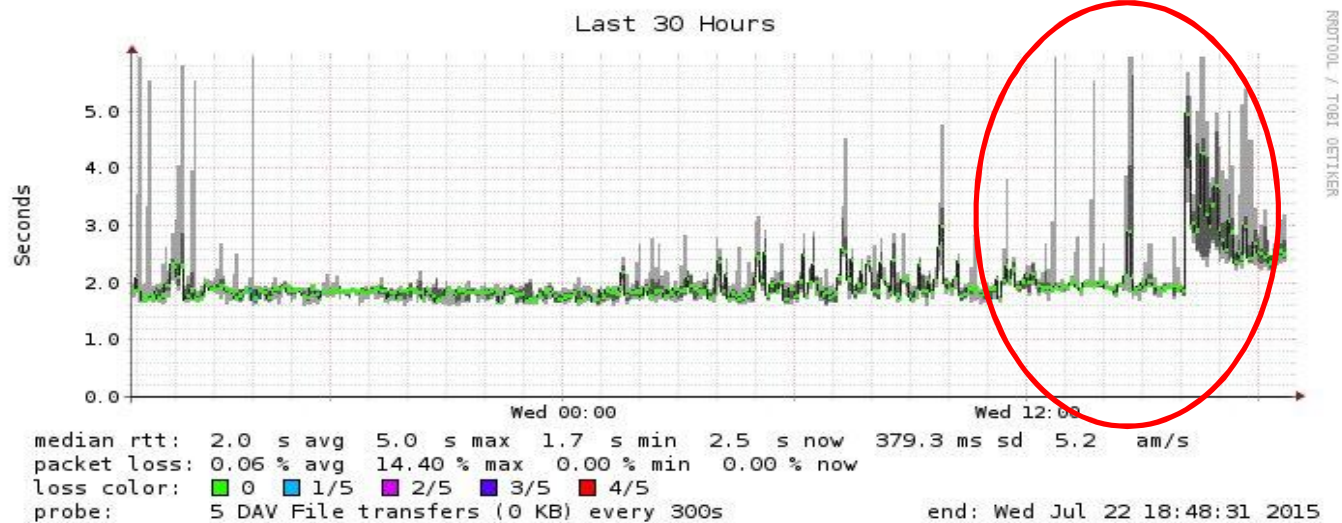
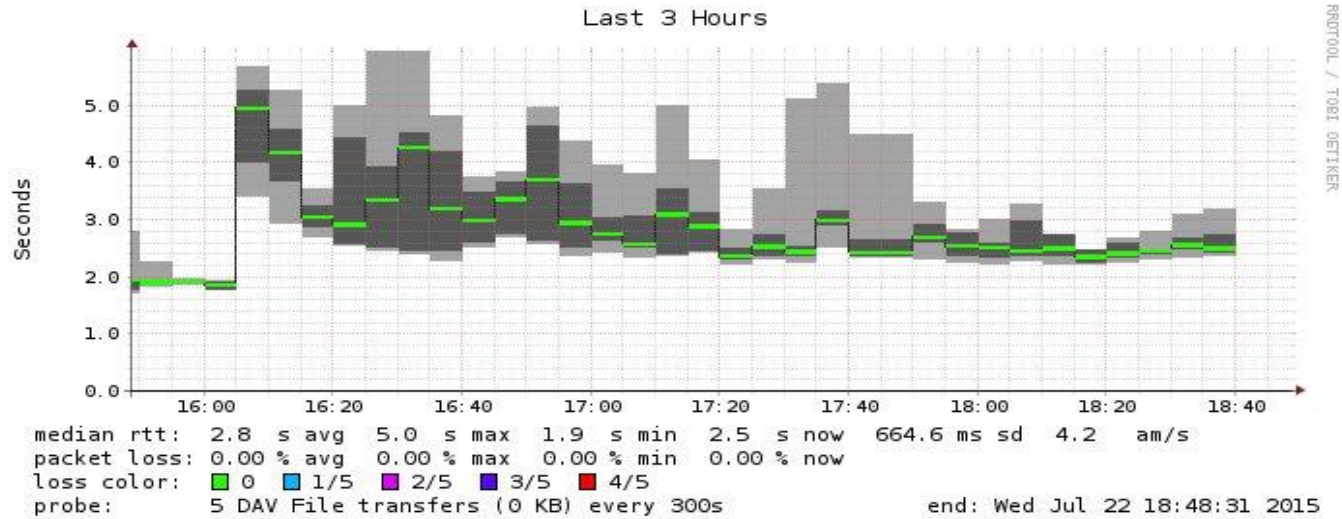
- DNS switch

4 hours

## WebDAV PUT Global domain latency



## WebDAV GET latency



## Summary

- Cowork with Dev team (RD+QA) closely
- Understand the infrastructure
- Application is white box
- Full control the system
- With the same goal: Customer Insight

```
h){deferfunc(){//Ensureoldpathisremovedifnoerroroccurrediferr==nil{err=os.Remove(oldPath)}else{logrus.Warnf("Keymigrationfailed,keyfilenotremovedat%s",oldP
(0644));err=nil{returnfmt.Errorf("Unabletocreatedaemonconfigurationdirectory:%s",err)}newFile,err:=os.OpenFile(newPath,os.O_RDWR|os.O_CREATE|os.O_TRUNC,06
,newPath,err)}defernewFile.Close()}oldFile,err:=os.Open(oldPath)iferr!=nil{returnfmt.Errorf("Unabletoopenoldfile:%s",err)}iferr:=os.Rename(oldFile.Close()if
"errorcopyingkey:%s",err)}logrus.Infof("Migratedkeyfrom%sto%s",oldPath,newPath)}return}funcmain(){logrus.SetFormatter(&logrus.TextFormatter{TimestampFormat:timeutils.RFC3339NanoFixed})varpfile*pidfile.PidFileifdaemonCfg.Pidfile!="{pf,err:=pidfile.New(da
```

# DEVOPS 2015

開發敏捷與維運高效的IT新典範





# DEVOPS 2015

開發敏捷與維運高效的IT新典範

