



June 26, 2017

The Paved PaaS to Microservices

Yunong Xiao,
Principal Software Engineer, Netflix
yunong@netflix.com, [@yunongx](https://twitter.com/yunongx), <http://yunong.io>

NETFLIX

100 million customers

in over **190** countries

streaming **125** million hrs/day

What is a Platform as a Service (Paas), Anyway?

“Platform as a service (PaaS)... allows customers to develop, run, and manage applications without the complexity of building and maintaining the infrastructure and platform...”

–Wikipedia

Our Use Case



NETFLIX ORIGINAL

MARVEL DAREDEVIL

★★★★★ 2016 2 Seasons **ULTRA HD 4K** 5.1

Blinded as a young boy, Matt Murdock fights injustice by day as a lawyer and by night as the Super Hero Daredevil in Hell's Kitchen, New York City.

Charlie Cox, Deborah Ann Woll, Elden Henson
TV Shows, Crime TV Shows

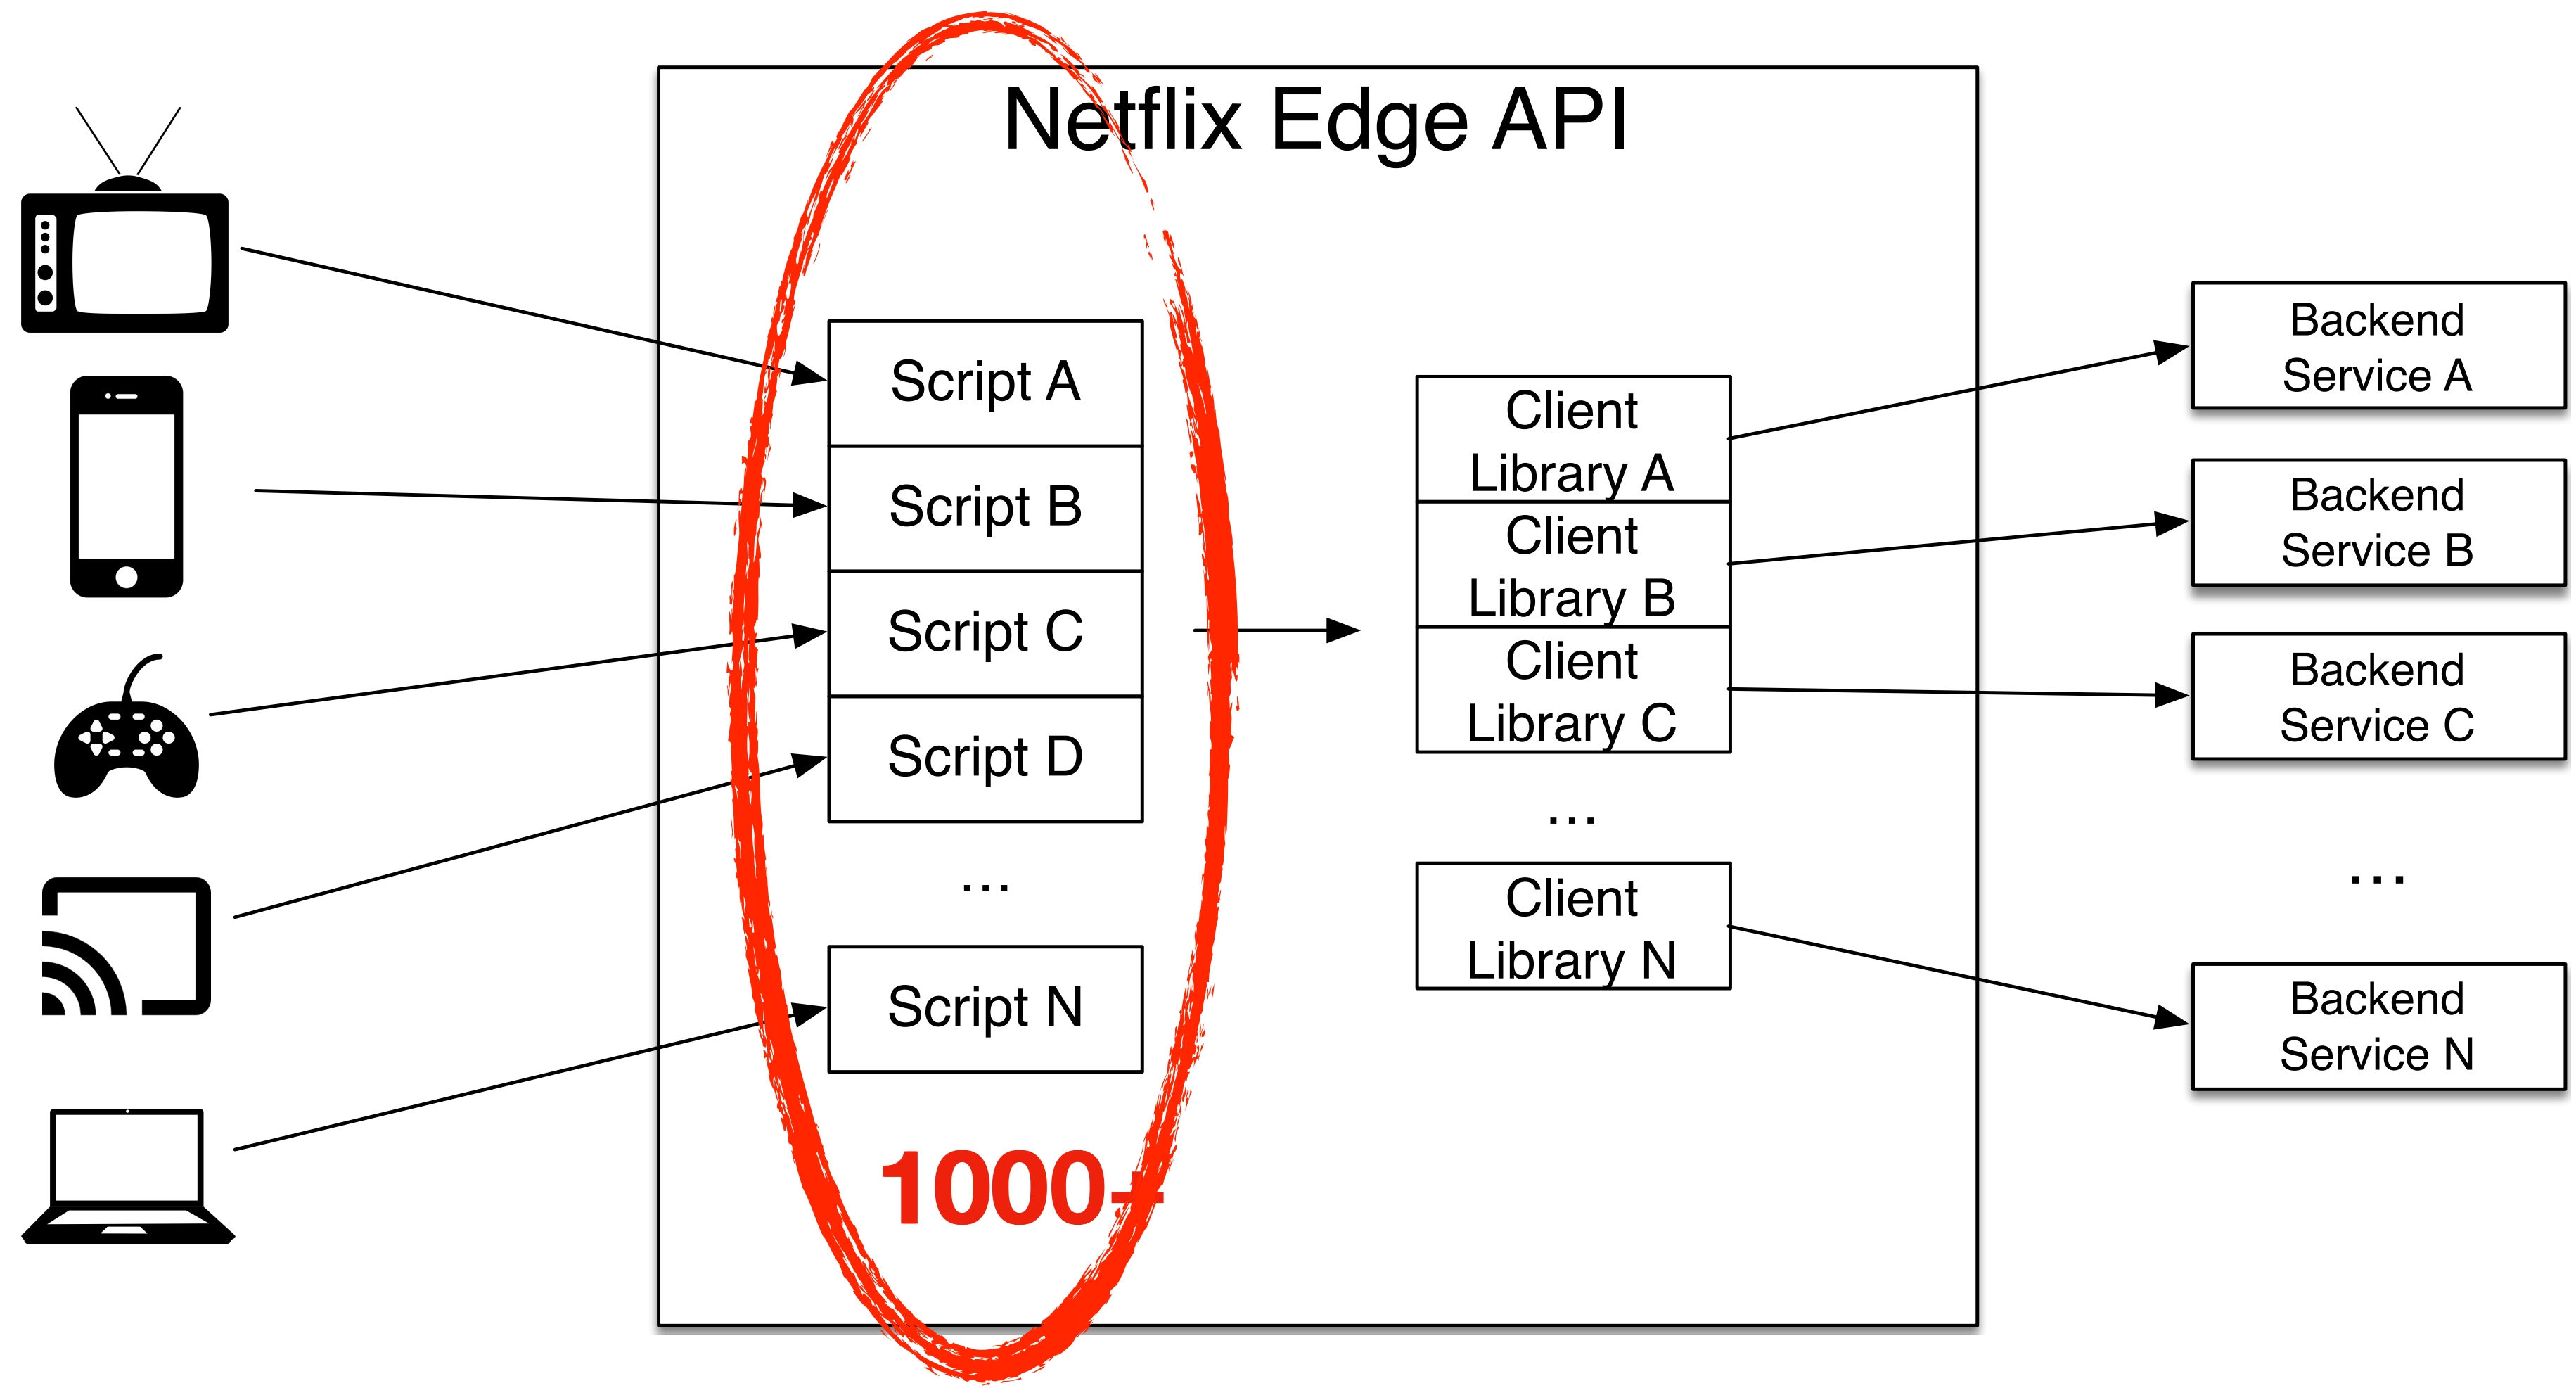


My List








Recently Viewed





Owned by client teams

Clients



JS Mostly

Standalone Services


TV

iOS

Android

Windows

Browsers



Edge API

Discovery
Playback
Non-member
...

Backend Services

Backend Service A

Backend Service B

Backend Service C

...

Backend Service N



Know Your Customer



Goals

Velocity

Reliability



Our PaaS to Achieving Both

1. Standardized components
2. Preassembled platform
3. Automation and tooling




1. Standardized Components

1 CUBIC CENTIMETER OR 1000 CUBIC

MILLIMETERS

=

1 MILLILITER.

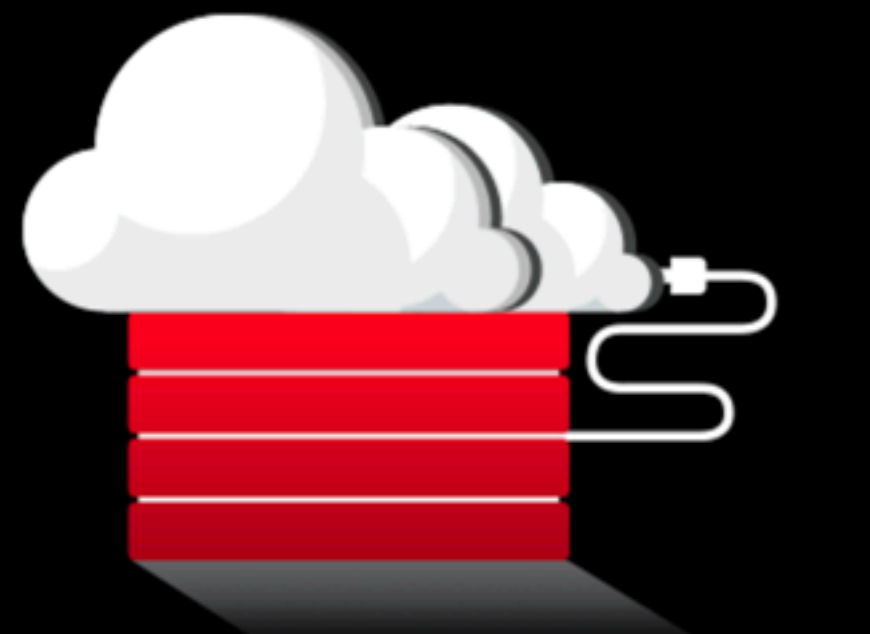


LITER.

10 CUBIC CENTIMETERS = 10 MILLILITERS AND OF WATER WEIGHS 1 DEKAGRAM.



What's Inside a Microservice?



Common Runtime Services & Libraries

Runtime containers, libraries and services that power microservices

The cloud platform is the foundation and technology stack for the majority of the services within Netflix. The cloud platform consists of cloud services, application libraries and application containers. Specifically, the platform provides service discovery through [Eureka](#), distributed configuration through [Archaius](#), resilient and intelligent inter-process and service communication through [Ribbon](#). To provide reliability beyond single service calls, [Hystrix](#) is provided to isolate latency and fault tolerance at runtime. The previous libraries and services can be used with any JVM based container.

The platform provides JVM container services through [Karyon](#) and [Governator](#) and support for non-JVM runtimes via the [Prana](#) sidecar. While Prana provides proxy capabilities within an instance, [Zuul](#) (which integrates Hystrix, Eureka, and Ribbon as part of its IPC capabilities) provides dynamically scriptable proxying at the edge of the cloud deployment.

The platform works well within the EC2 cloud utilizing the Amazon autoscaler. For container applications and batch jobs running on Apache Mesos, [Fenzo](#) is a scheduler that provides advanced scheduling and resource management for cloud native frameworks. Fenzo provides plugin implementations for bin packing, cluster autoscaling, and custom scheduling optimizations can be implemented through user-defined plugins.

Why Standards?

BBC ONLINE NETWORK [HOMEPAGE](#) | [SITEMAP](#) | [SCHEDULES](#) | [BBC INFORMATION](#) | [BBC EDUCATION](#) | [BBC WORLD SERVICE](#)

BBC NEWS


News in Audio News in Video Newyddion Новости Noticias أخبار 国际新闻 粵語廣播

Thursday, September 30, 1999 Published at 18:53 GMT 19:53 UK

[Front Page](#)
[World](#)
[UK](#)
[UK Politics](#)
[Business](#)
[Sci/Tech](#)
[Health](#)
[Education](#)
[Sport](#)
[Entertainment](#)
[Talking Point](#)
[In Depth](#)
[On Air](#)
[Archive](#)

Sci/Tech

Confusion leads to Mars failure



The Mars Climate Orbiter: Now in pieces on the planet's surface

The Mars Climate Orbiter Spacecraft was lost because one Nasa team used imperial units while another used metric units for a key spacecraft operation.

[Feedback](#)
[Low Graphics](#)
[Help](#)

Sci/Tech Contents

Relevant Stories

24 Sep 99 | Sci/Tech
[Scientist fights Mars setback](#)

23 Sep 99 | Sci/Tech
[Mars probe feared destroyed](#)

23 Sep 99 | Sci/Tech
[What the loss of Mars Climate Orbiter means](#)

17 Jul 99 | Sci/Tech
[Astronauts call for Mars mission](#)

Internet Links

[Mars Climate Orbiter](#)

The BBC is not responsible for the content of external internet sites.

Mélange of RPC

Apache Thrift TM



Wangle

C++ networking library

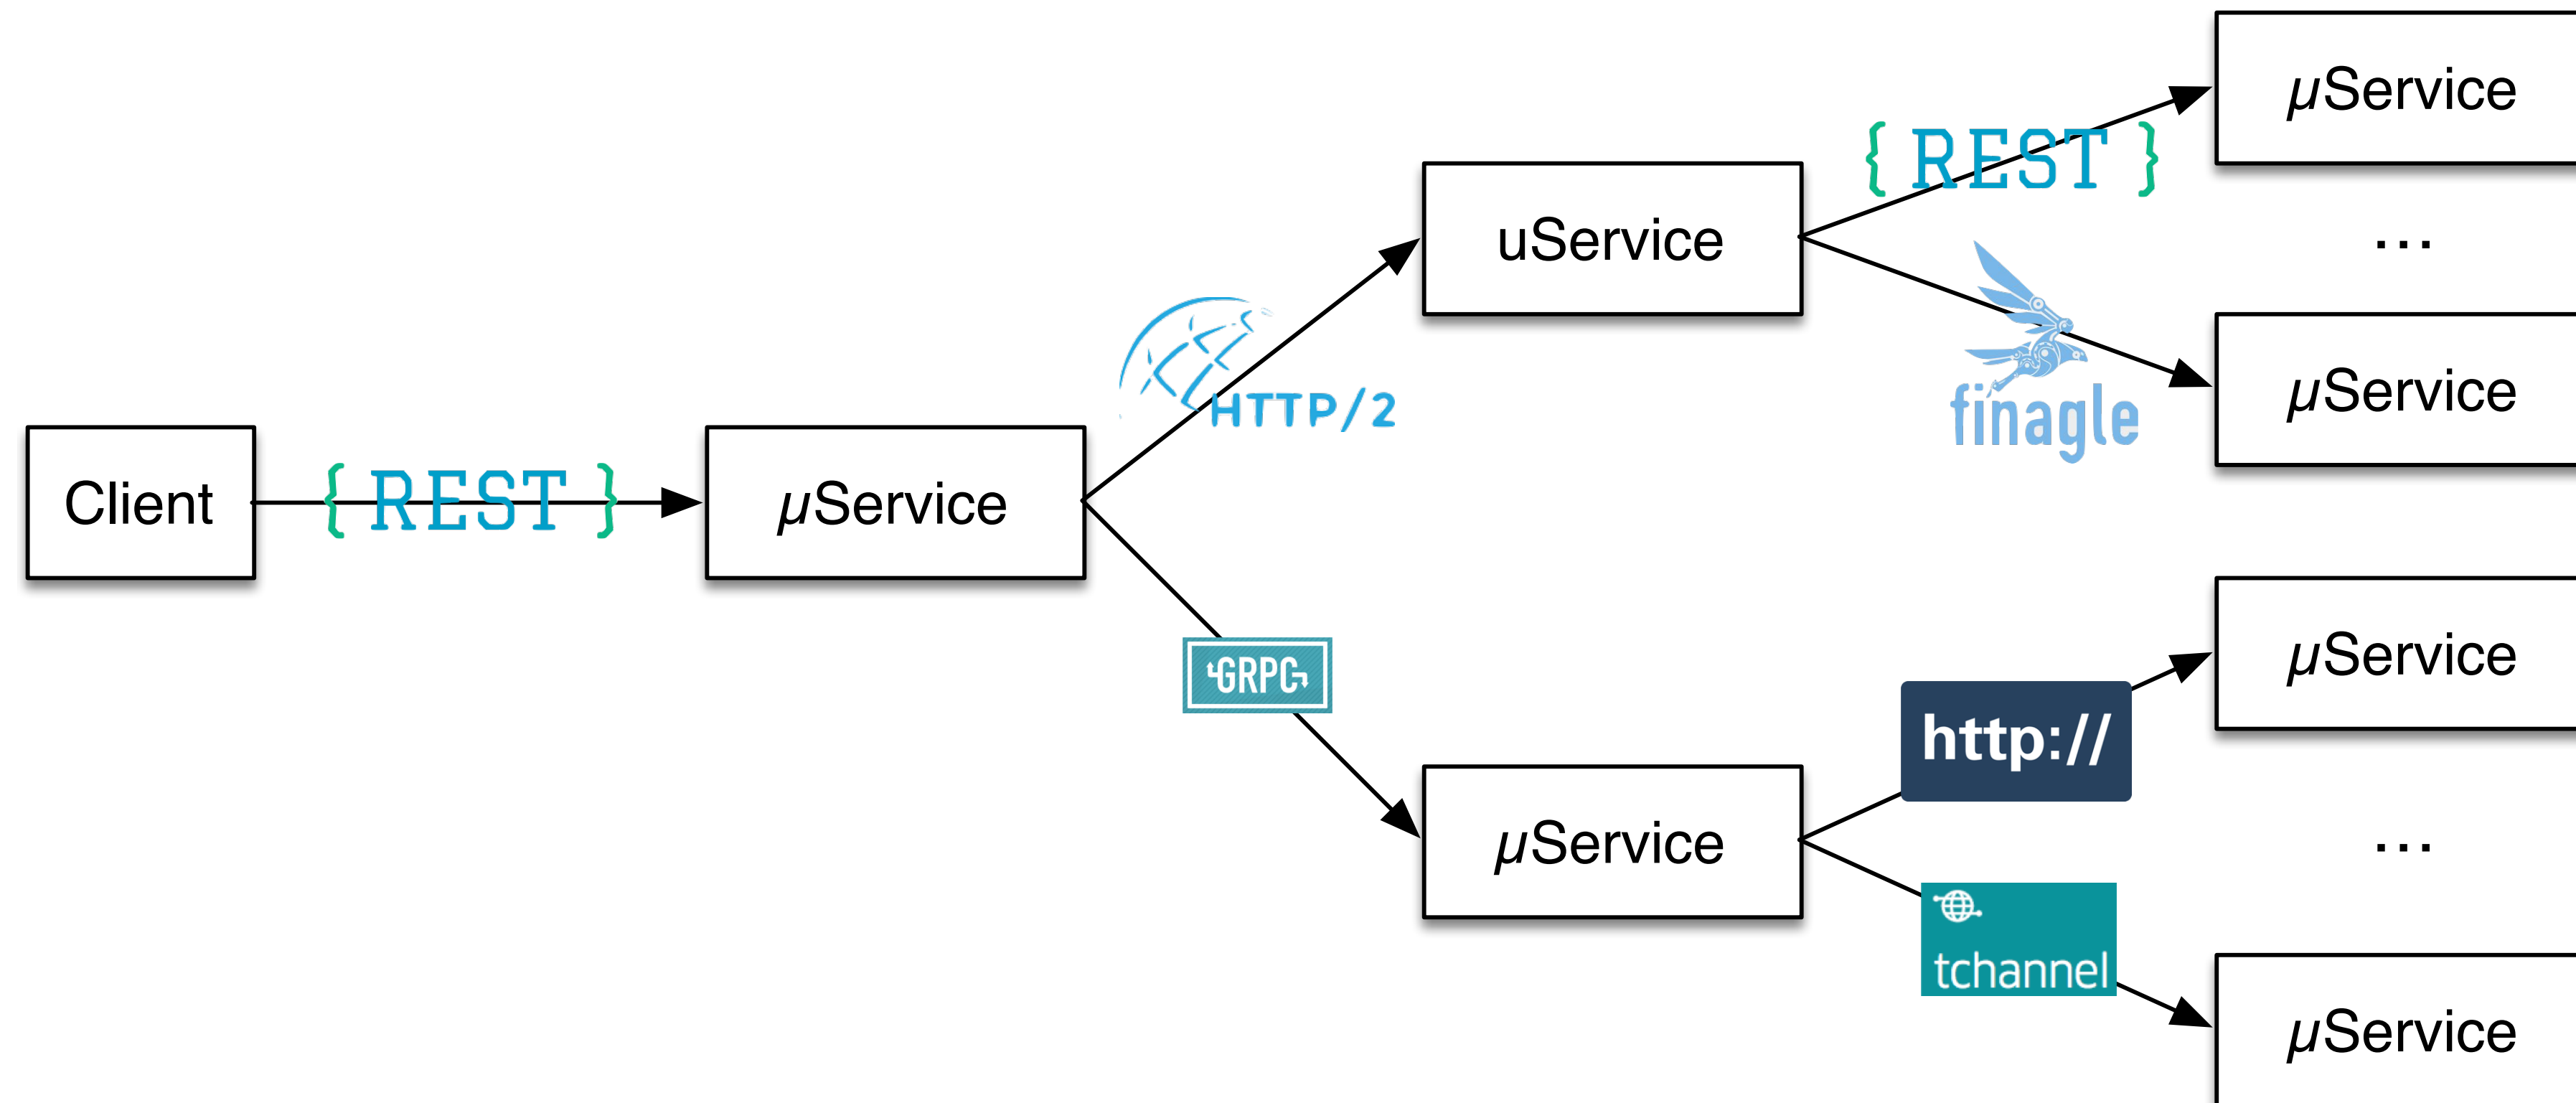


Q uick
U DP
I nternet
C onnections

{ REST }



Microservice Interactions



Failure: **When** not If



I don't get to chill. Can I at least Netflix? Fix you're sh!t!

4

RETWEETS

1

FAVORITES



Netflix is down. I have to parent now. **#UNFAIR**

121

RETWEETS

18

FAVORITES



Netflix is down and I feel like my best friend died.

#becauseitdid

19

RETWEETS

4

FAVORITES



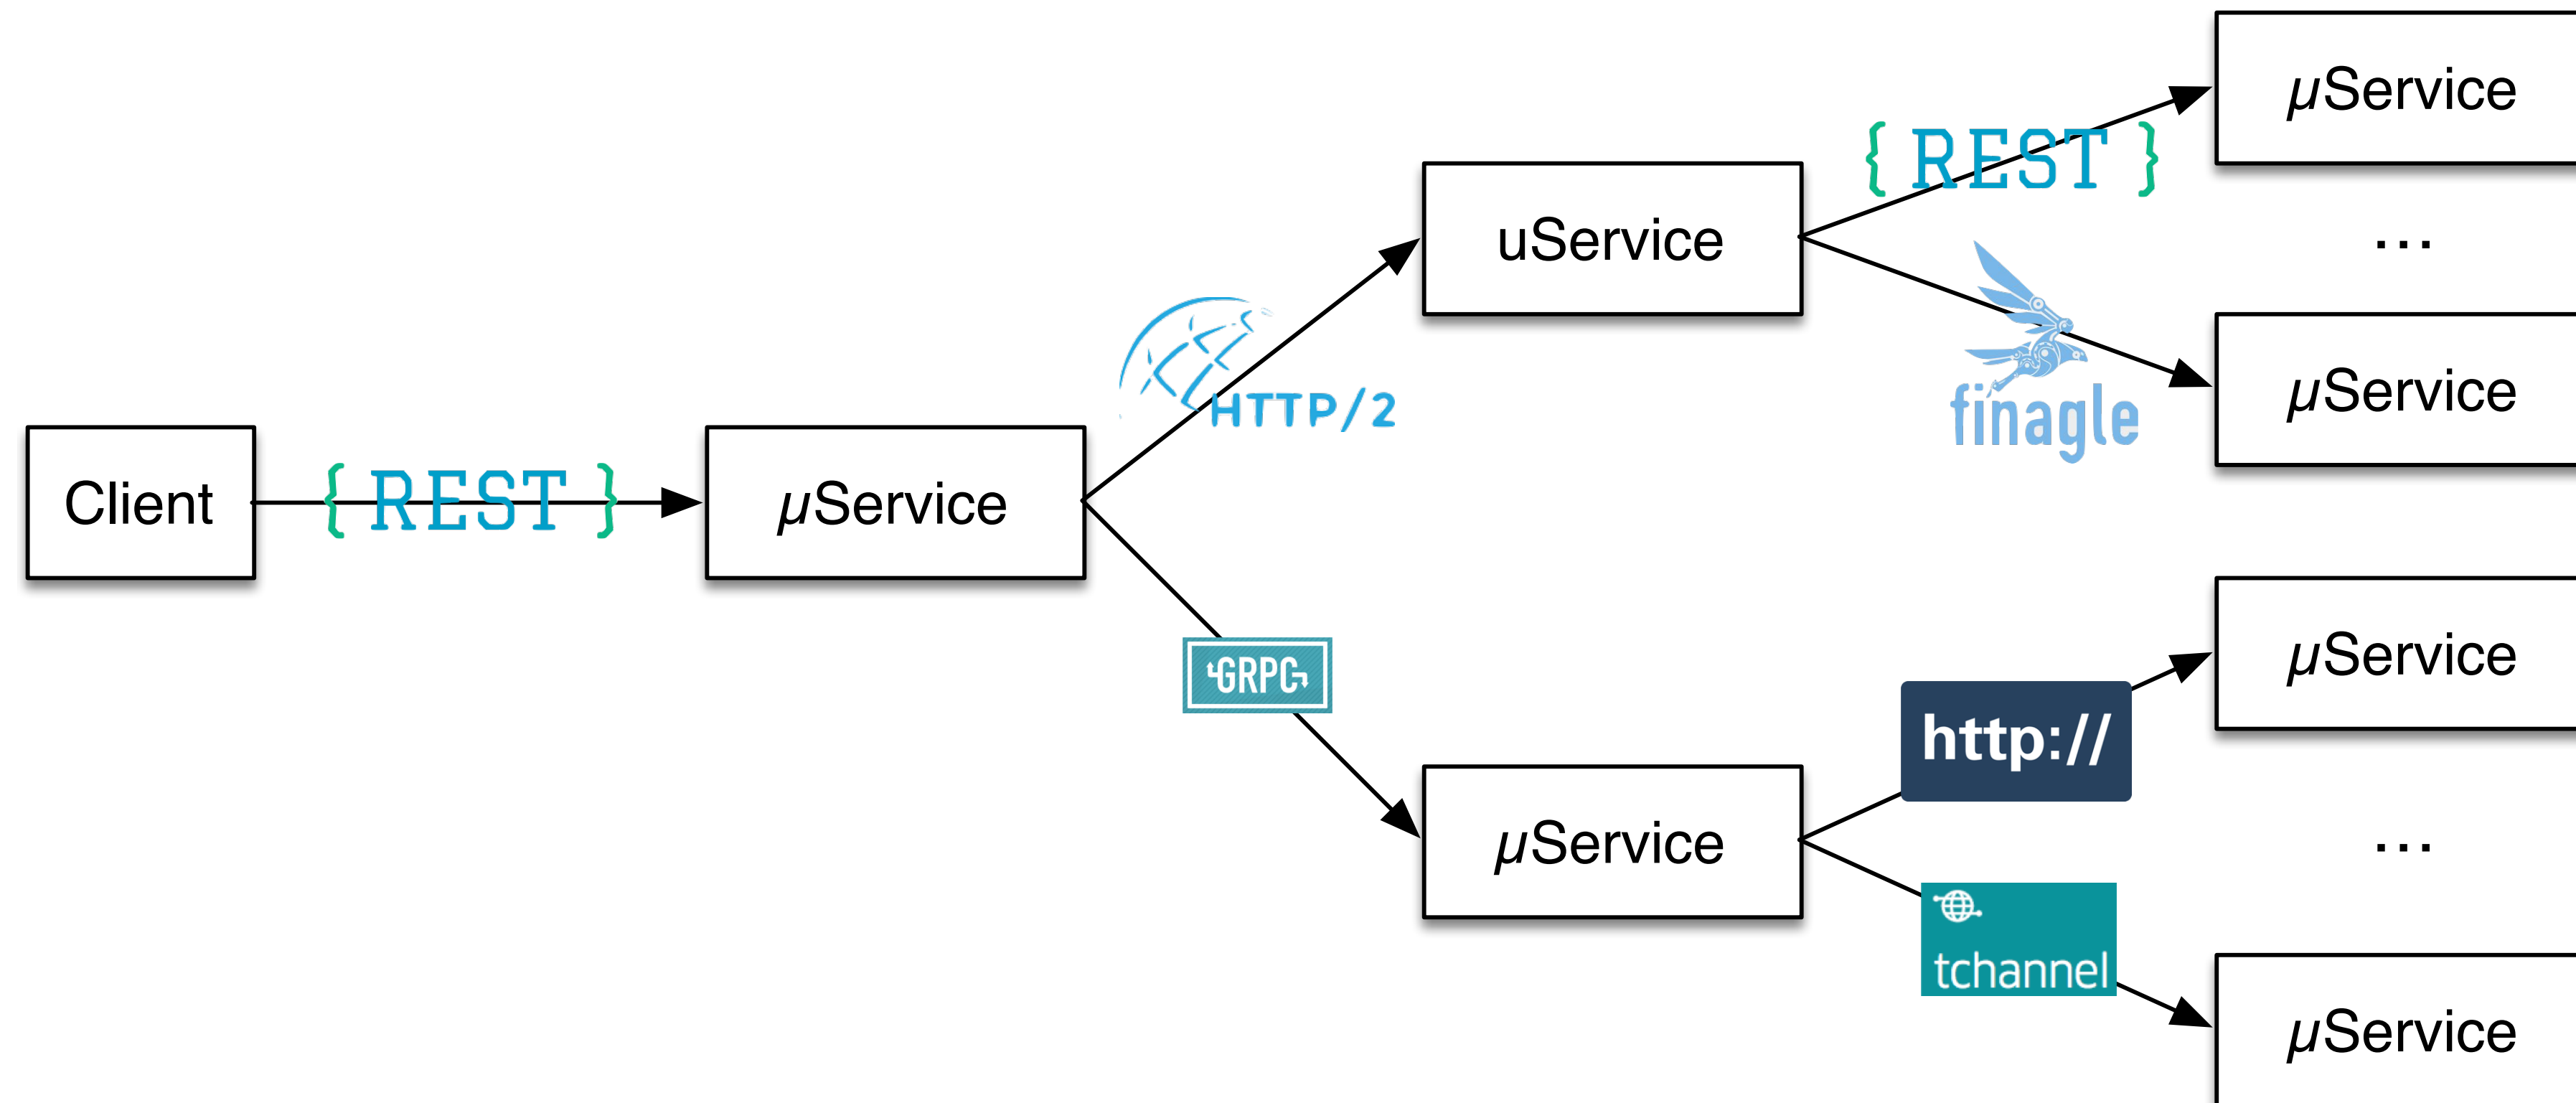
Mean Time to Detection (MTTD)

Mean Time to Repair (MTTR)

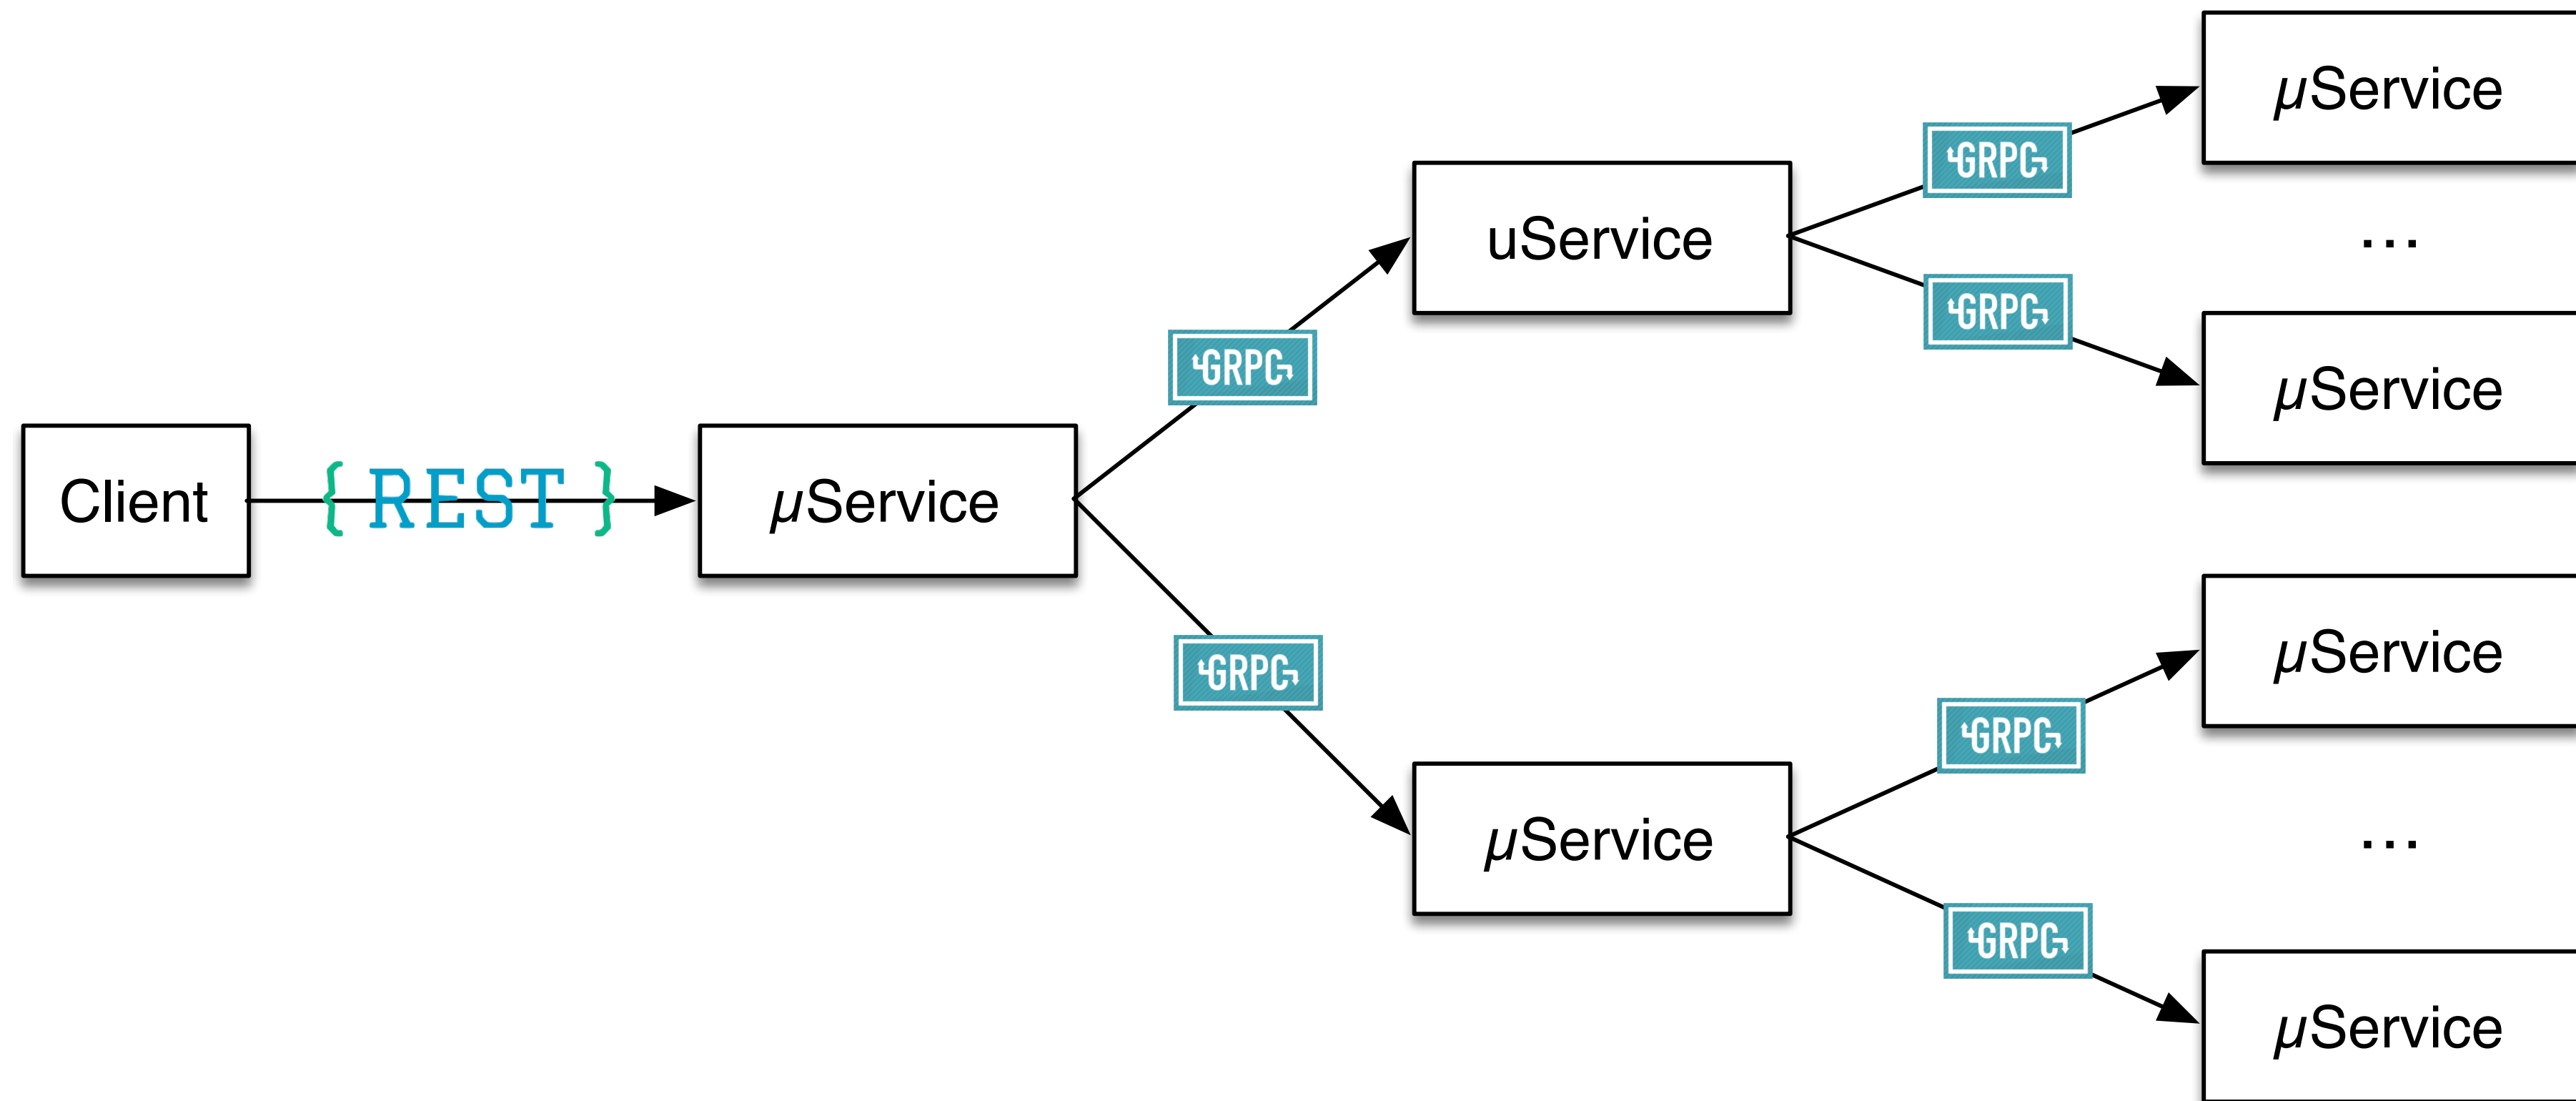
“Is it fixed yet?”

– Managers everywhere


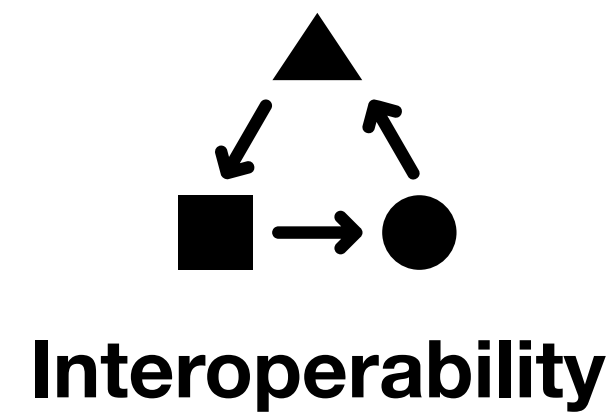

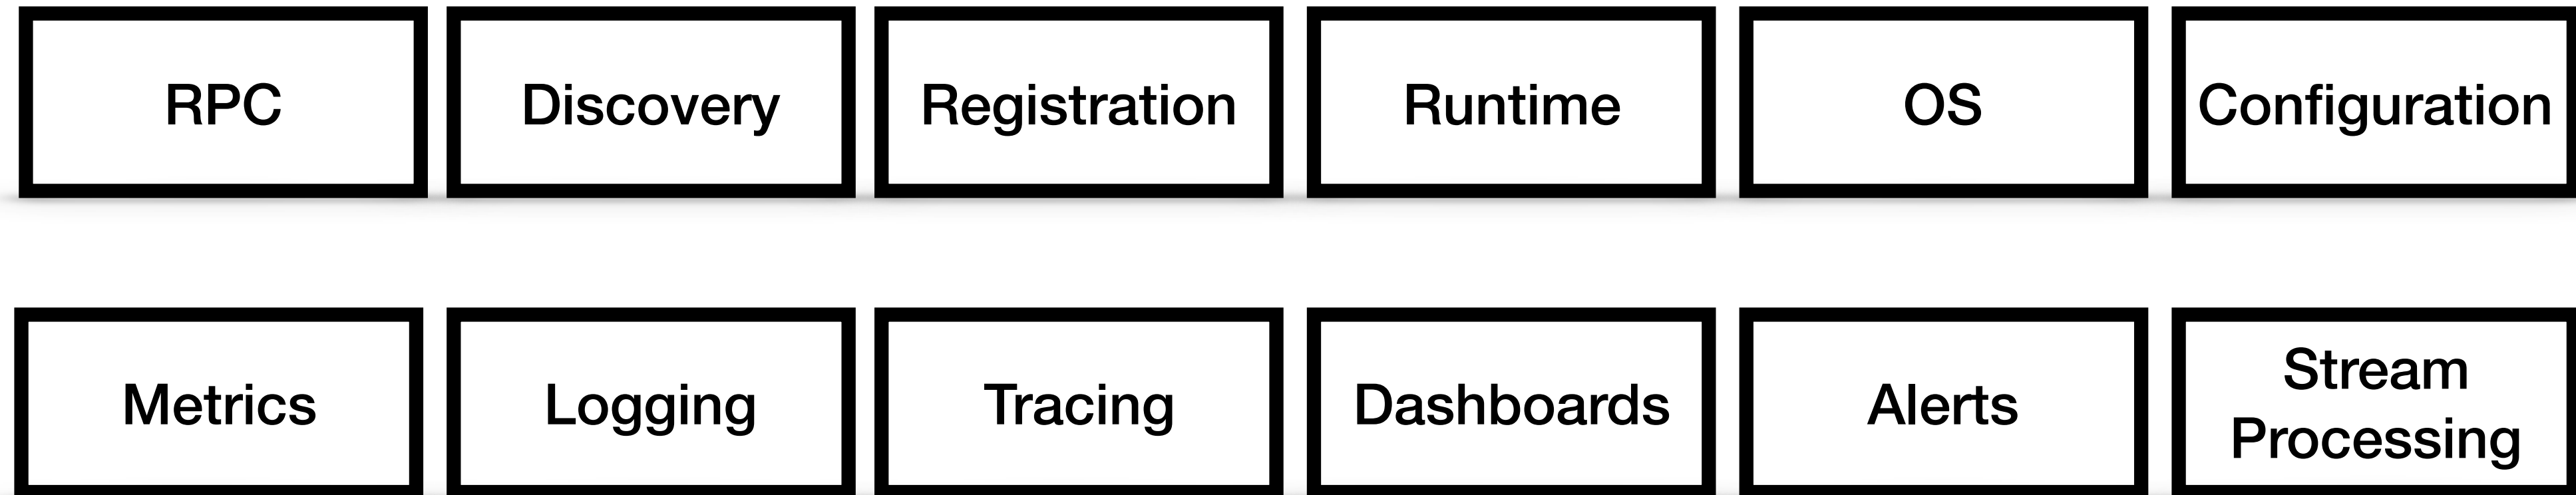
N Flavors of RPC










One Standard RPC



Benefits of Standardizing




But I'm a Snowflake...




Freedom

&

Responsibility



Off-road




Innovation



Reintegrate



New



Vacuum



Burden

A collection of various mechanical parts and tools scattered on a light-colored surface. The parts include gears, springs, shafts, bearings, and other mechanical components. The text "2. Preassembled Platform" is overlaid in the center of the image.

2. Preassembled Platform

Assembly Required

Discovery

Runtime

OS

Configuration

RPC

Registration

Logging

Dashboards

Stream
Processing

Metrics

Tracing

Alerts

Getting Out of the Blocks

Docs

Copy/paste

Which versions?

Initialization

Configuration

Missing components

Days or weeks



Netflix Lets Viewers Pick the Plot

By JOHN KOBLIN JUNE 20, 2017



Product Innovation



Velocity



Reliability

Not a single line of business logic!

Assembly Required

Discovery

Runtime

OS

Configuration

RPC

Registration

Metrics

Logging


Dashboards

Stream
Processing

Tracing

Alerts

Preassembled Platform



RPC

Discovery

Registration

Runtime

OS

Configuration

Metrics

Logging


Tracing

Dashboards


Alerts



Stream
Processing

Out of the Box




Component Management









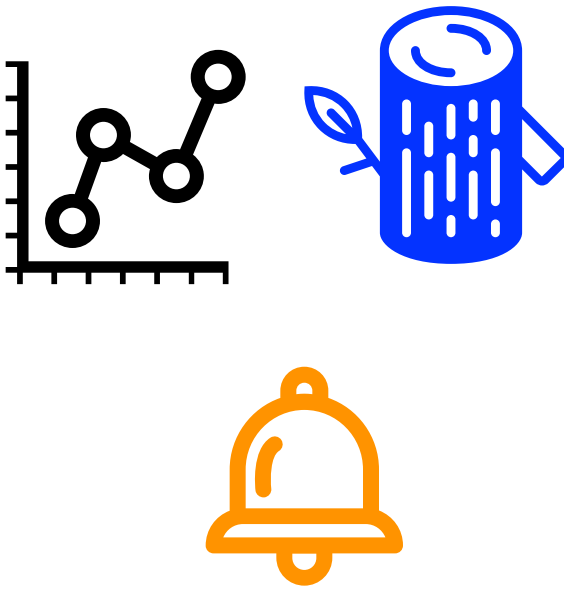


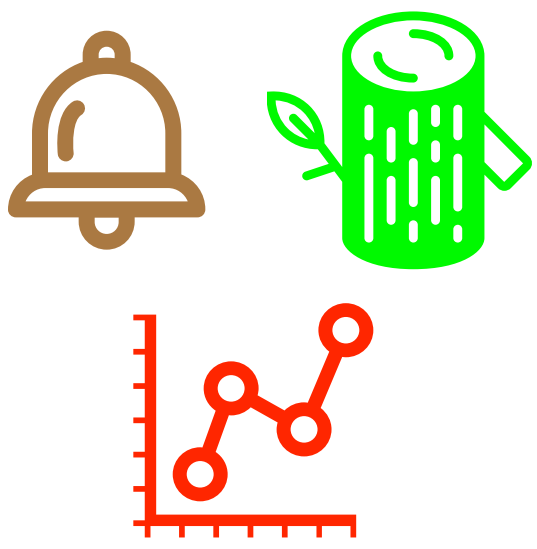
MISSING







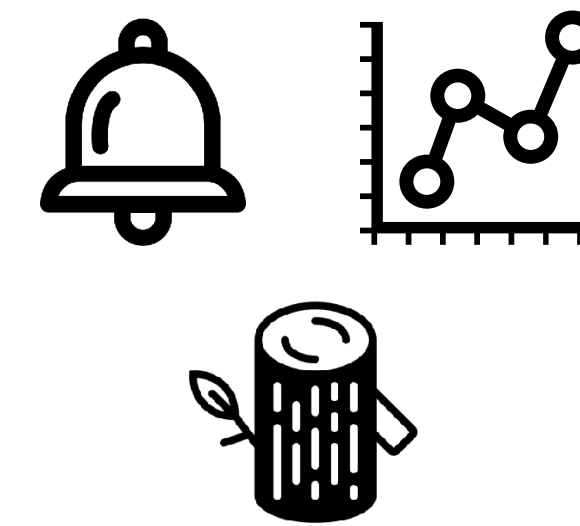


Insights



Application, system, and runtime metrics & logs



Consistent application, system, and runtime metrics & logs



Reduces MTTD & MTTR

Configures and Initializes Correctly

BOEING 747-400 NORMAL PROCEDURES CHECKLIST

POWER UP / SAFETY CHECK

First Officer	Captain
CIRCUITBREAKERS.....	CHECKED
BATTERY.....	ON
STANDBY POWER.....	AUTO
HYDRAULIC DEMAND PUMPS.....	OFF
WINDSHIELD WIPERS.....	OFF
ALTERNATE FLAPS AND GEAR.....	OFF
GEAR LEVER.....	DOWN
FLAPS.....	CHECKED
APU.....	RUNNING
ELECTRICAL SYSTEM.....	SET/APU AVAIL ON
APU BLEED AIR.....	ON
ISOLATION VALVES.....	OPEN
PACKS.....	NORMAL


PREFLIGHT











First Officer	Captain
EMERGENCY EQUIPMENT.....	CHECKED
FIRE PROTECTION.....	CHECKED
INTERRUPT SWITCHES.....	ON
PASSENGER OXYGEN.....	NORMAL
STAB TRIM CUTOUT SWITCHES.....	AUTO




TRANSPONDER.....	SET
SOURCE SELECTORS.....	SET
CLOCKS.....	SET
CRT SELECTORS.....	NORMAL
PFD.....	CHECKED
ND.....	CHECKED
AUTOBRAKES.....	RTO
EIU SELECTOR.....	AUTO
HDG REFERENCE SWITCH.....	NORMAL
FMC MASTER SELECTOR.....	LEFT
GROUND PROX SYSTEM.....	CHECKED

BEFORE STARTING

First Officer	Captain
HYDRAULIC DEMAND PUMPS.....	AUTO, AUX (1 AND 4)
BRAKE PRESSURE.....	NORMAL
FUEL QUANTITY.....	___ KG
FUEL SYSTEM.....	SET
X-FEEDS.....	OPEN (1 & 4 CLOSED)
SEAT BELTS SIGN.....	ON
NOTOC.....	CHECKED
SHIPS PAPERS.....	ON BOARD
PERFORMANCE DATA.....	CHECKED AND SET



3 days ago	v1.4.0 <small>...</small> - 869e373  zip  tar.gz
18 days ago	v1.4.0-pre1 <small>...</small> - 52bcda9  zip  tar.gz
19 days ago	v1.3.9 <small>...</small> - 28aee69  zip  tar.gz
25 days ago	v1.3.7 <small>...</small> - e680329  zip  tar.gz
on May 25	v1.3.6 <small>...</small> - ef2c853  zip  tar.gz



Versions

Updates

Compatibility

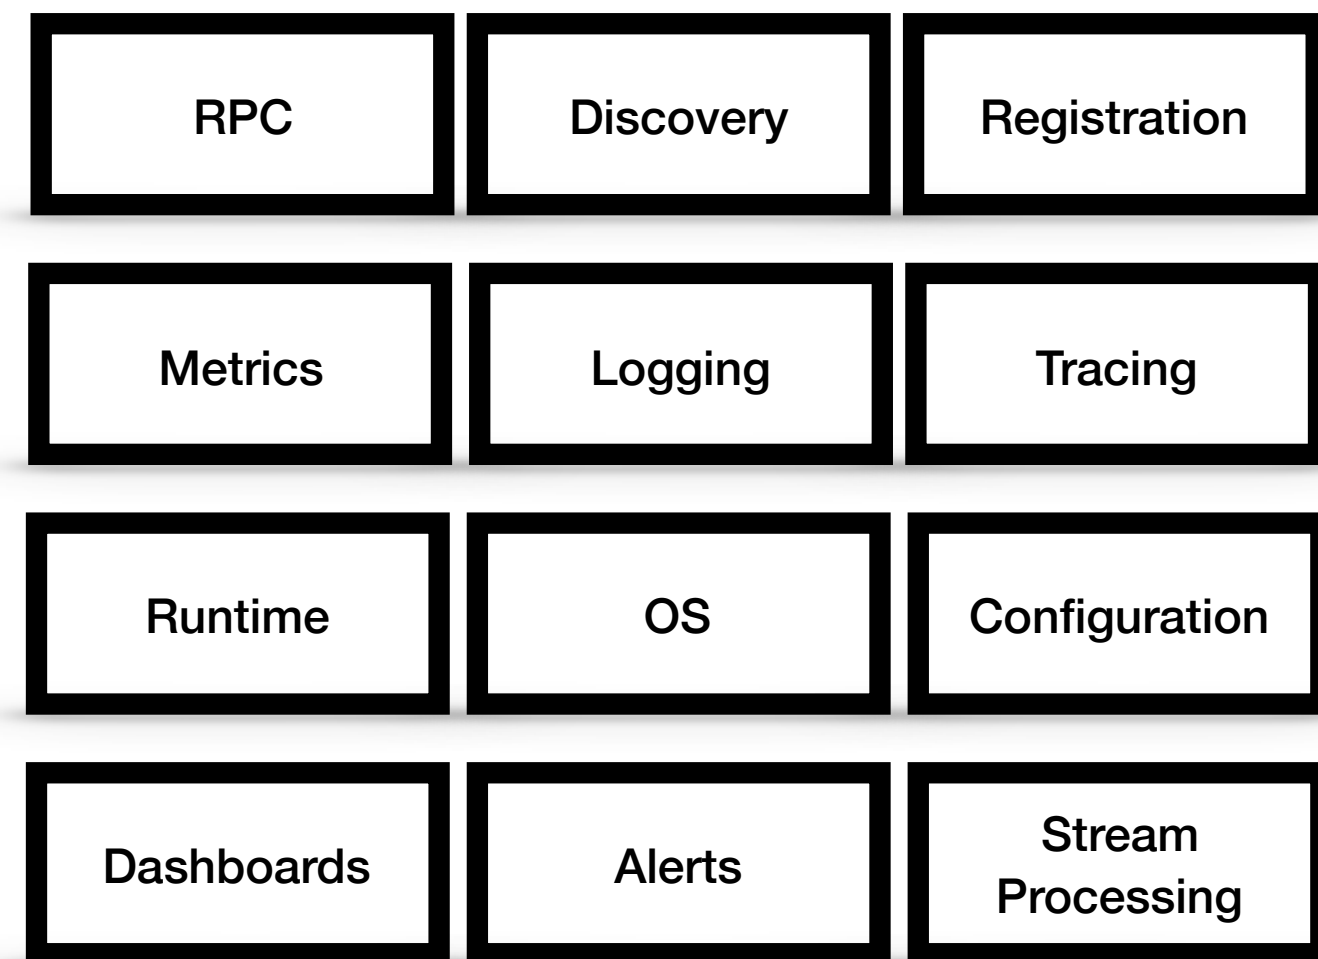
Important Questions





What's in and out?



Maintenance vs Convenience



Base Platform



Solution: Layers & Flavors



Base platform

Data access

Rendering

Backend


...



How to Ensure Platform Correctness?



Test, Test, Test

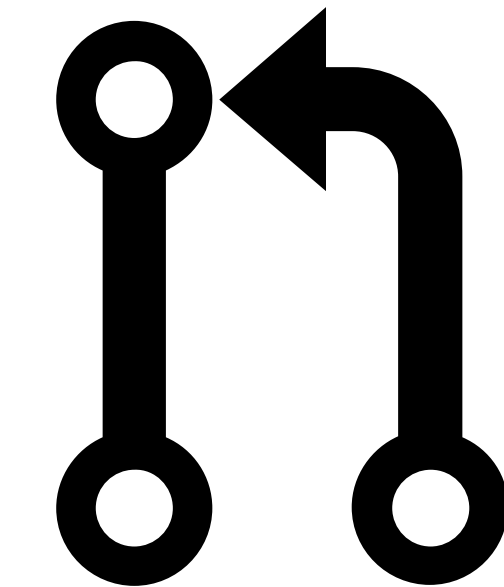


Unit

Integration

Functional

Cloud



Every PR






Dog Food with your own Service



Component Correctness



Customers




Gate Keeper


RPC	Discovery	Registration
Metrics	Logging	Tracing
Runtime	OS	Configuration
Dashboards	Alerts	Stream Processing



Lock down versions



Test components




Updates require PRs



How Locked Down is it?



Tradeoffs



Flexibility

vs



Reliability



Consistency



Support



Stay on paved path!



Season to Taste



Config overrides

Startup & shutdown hooks

Access to 3rd party libs

Swap, disable, or configure components

Raw component access



Platform Versioning?



API Semantic Versioning

1.2.3

^1.0.0

~1.3.0





Use Conventional Changelog



3. Automation and Tooling

Ship a Feature



Netflix Lets Viewers Pick the Plot

By JOHN KOBLIN JUNE 20, 2017



Steps



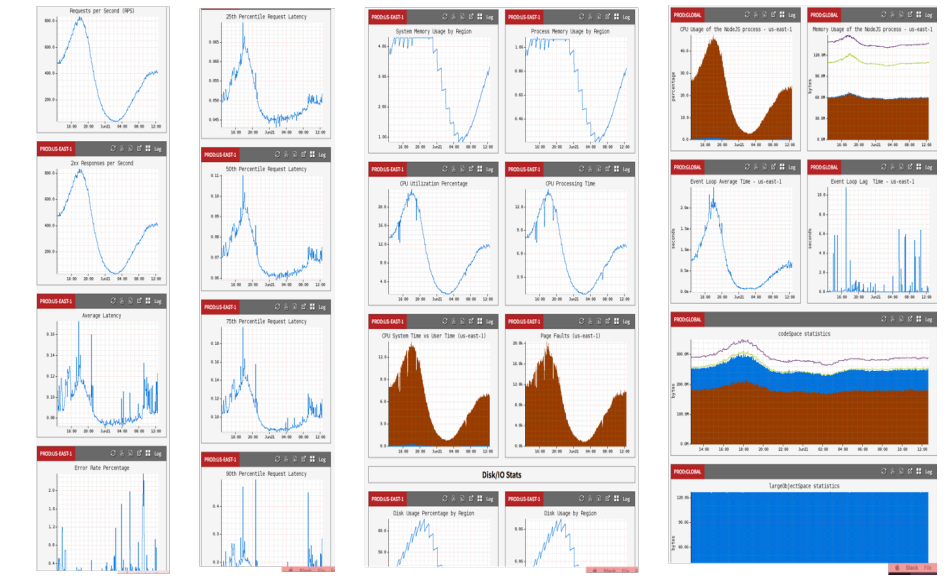
Development




Testing



Deployment





Operations



Development





CLI for common dev experience



Env bootstrap

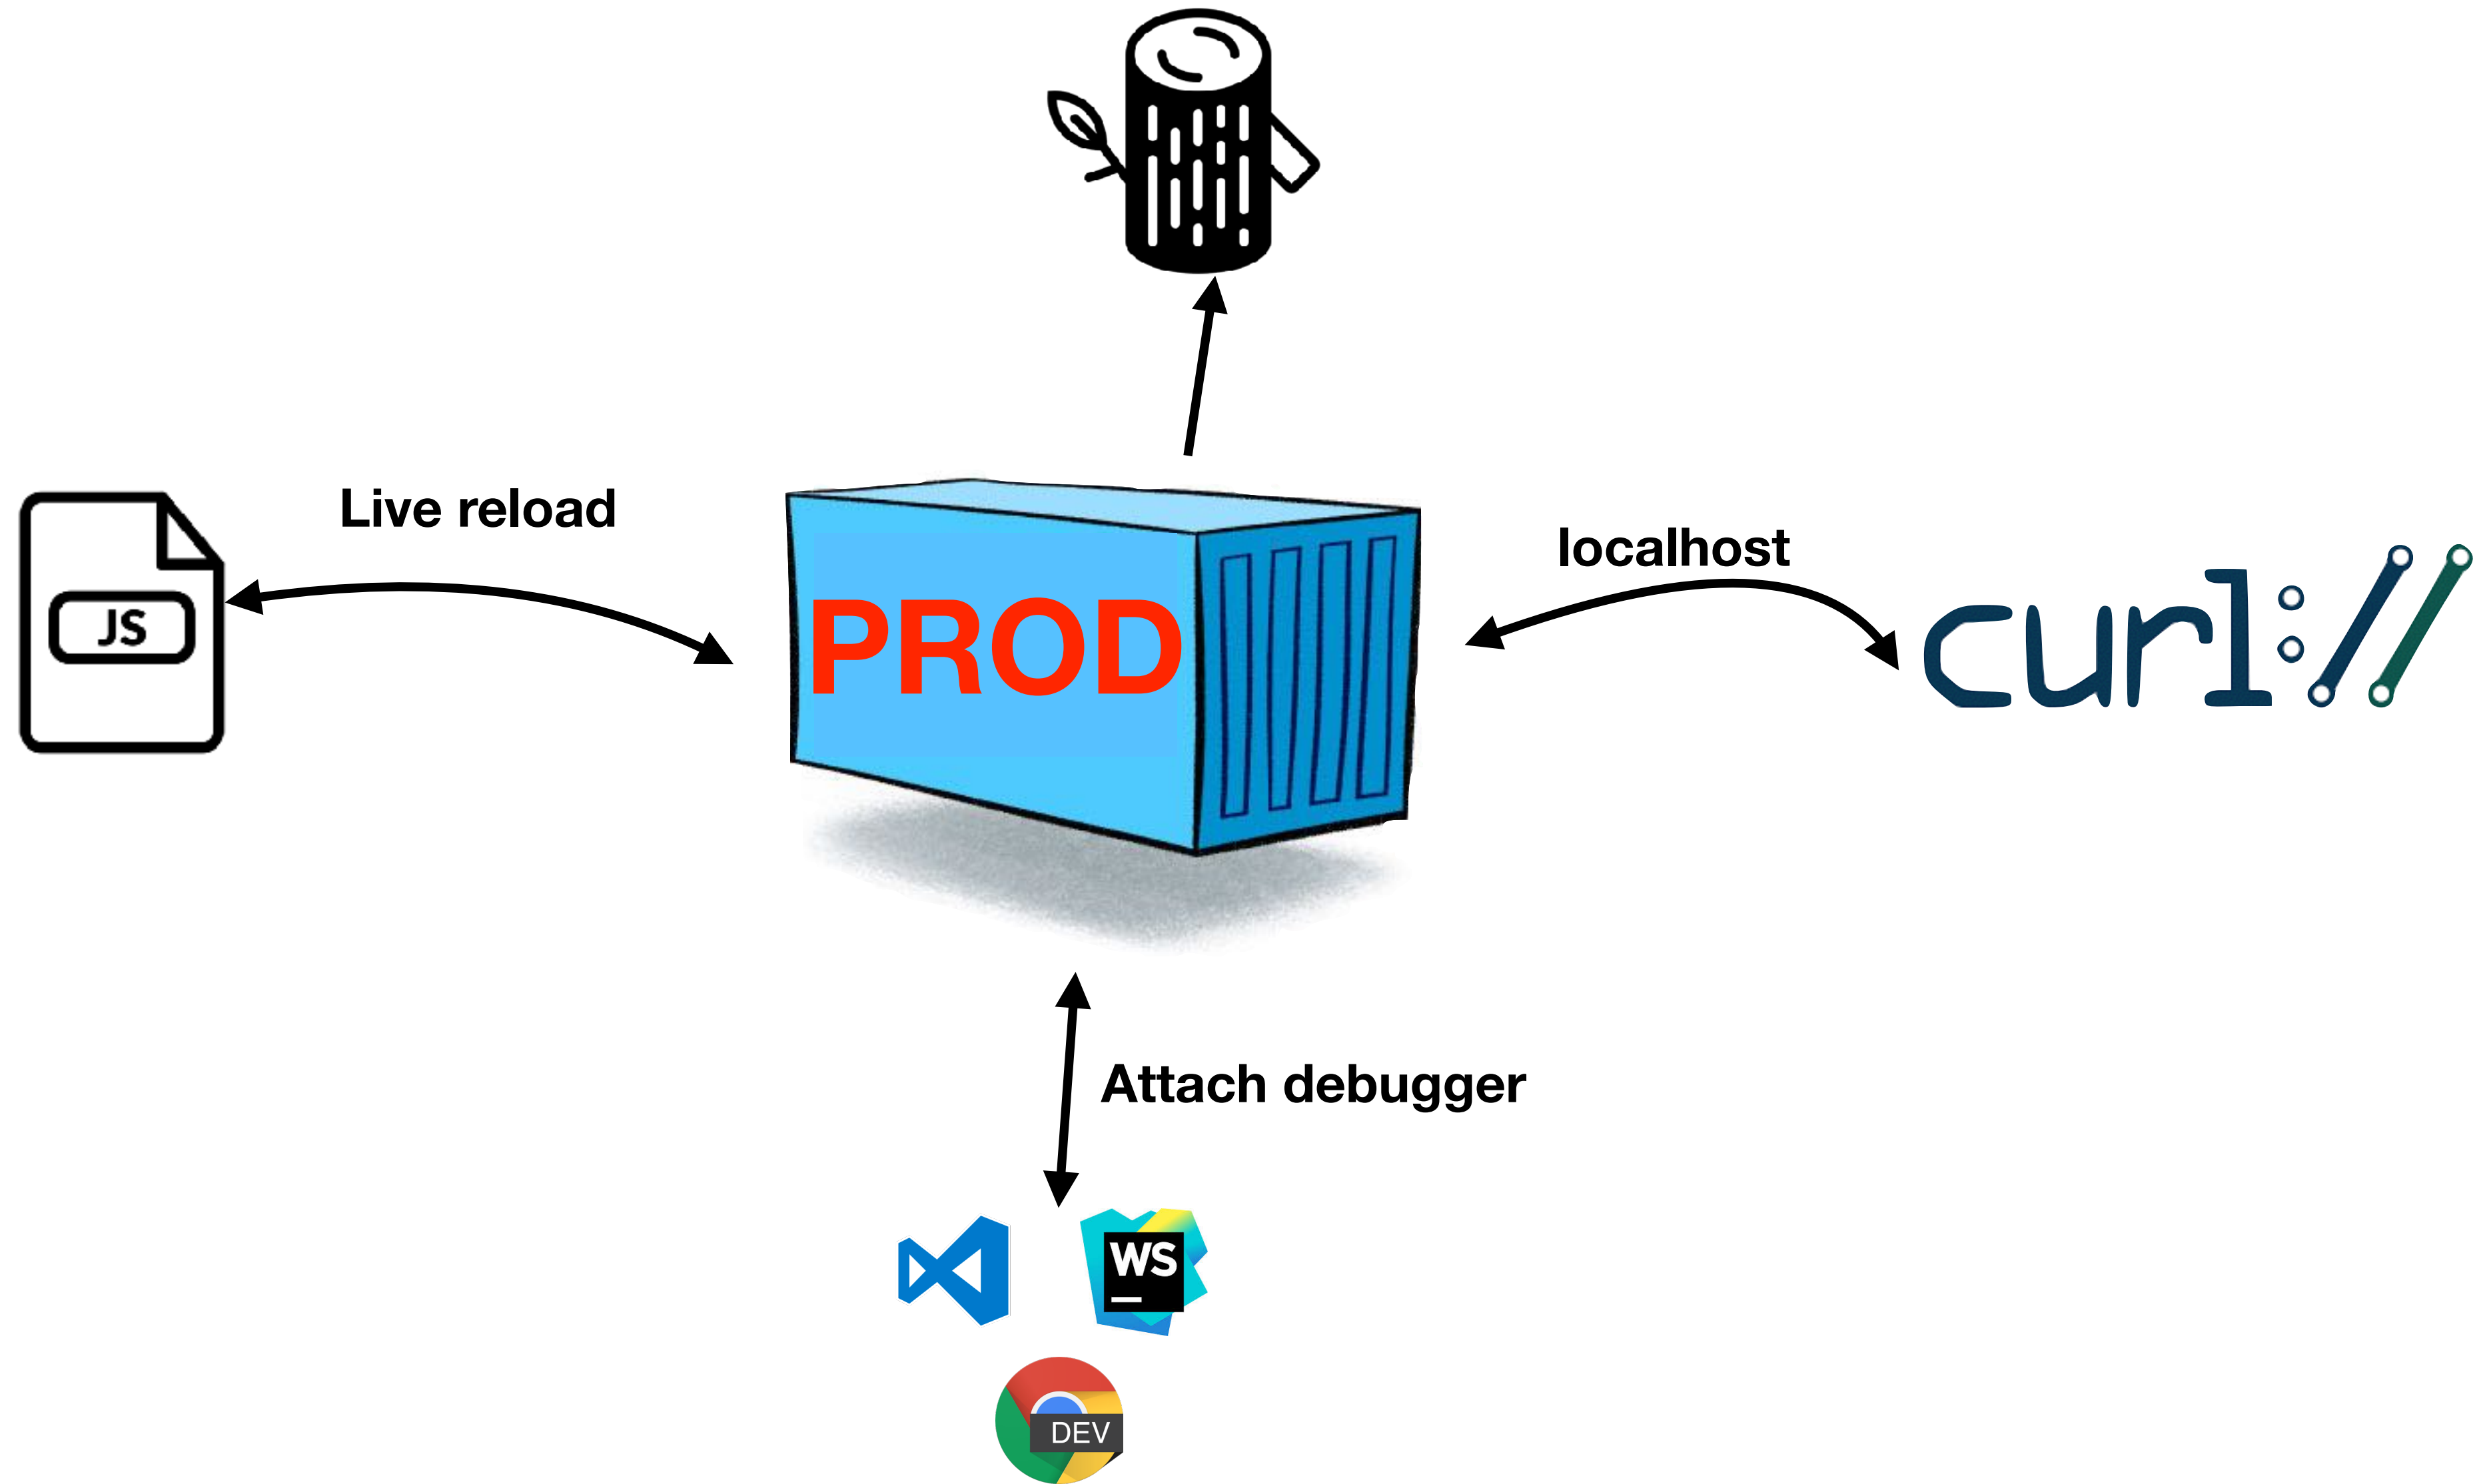




Integrate tooling & services



Run local & cloud


Local Development





Testing

Testing




Netflix Lets Viewers Pick the Plot


By JOHN KOBLIN JUNE 20, 2017



Testing



Preassembled Platform



RPC

Discovery

Registration

Runtime

OS

Configuration

Metrics


Logging

Tracing


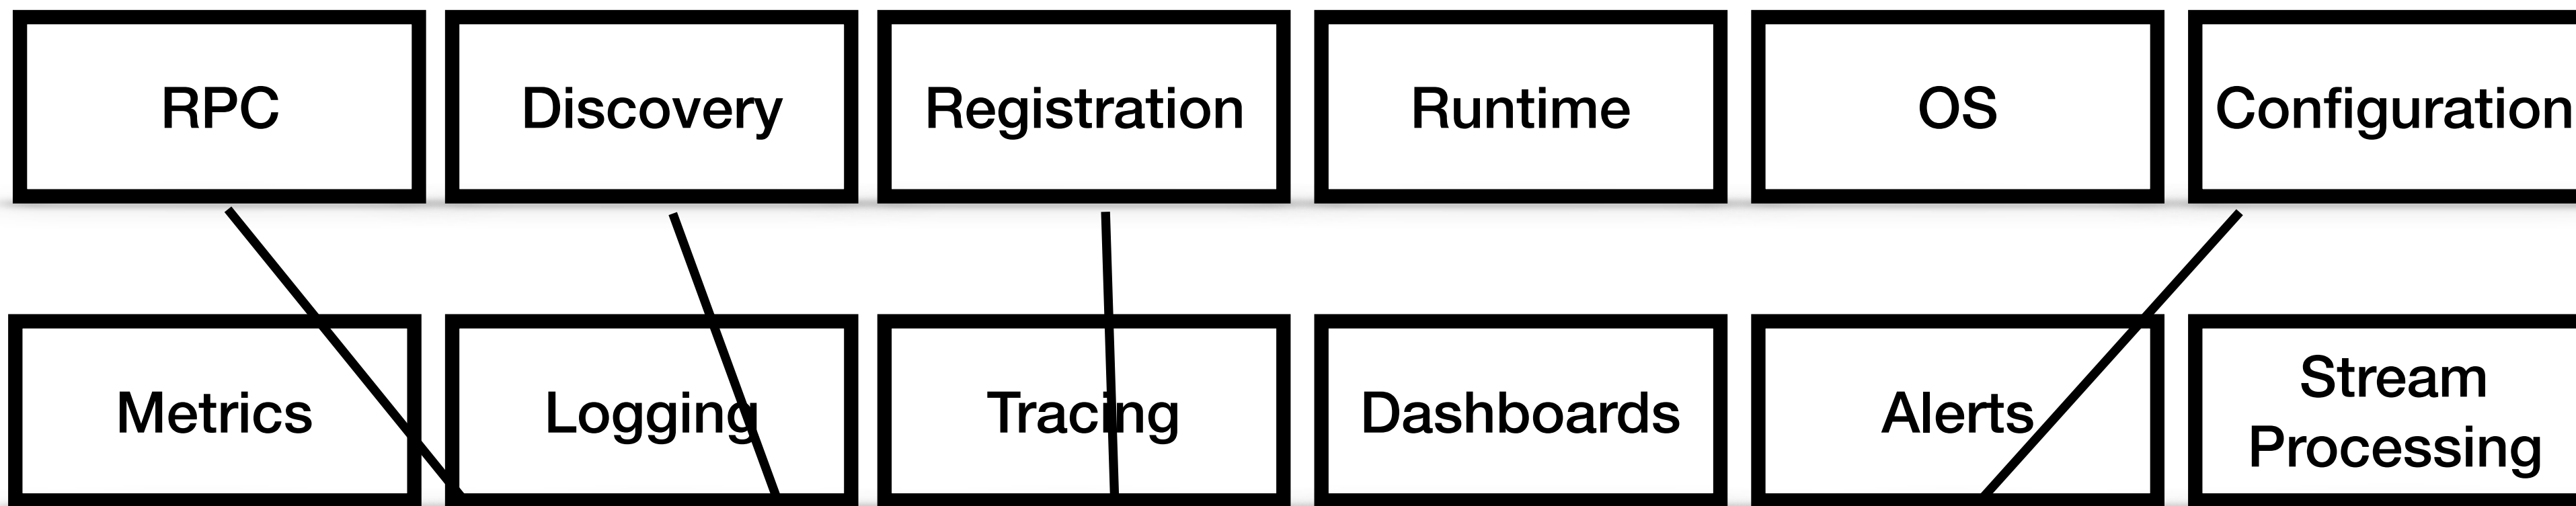

Dashboards

Alerts


Stream
Processing




Preassembled Platform



Provide First Class Mocks



Preassembled Platform



RPC

Discovery

Registration

Runtime

OS

Configuration

Metrics

Logging


Tracing

Dashboards


Alerts

Stream
Processing

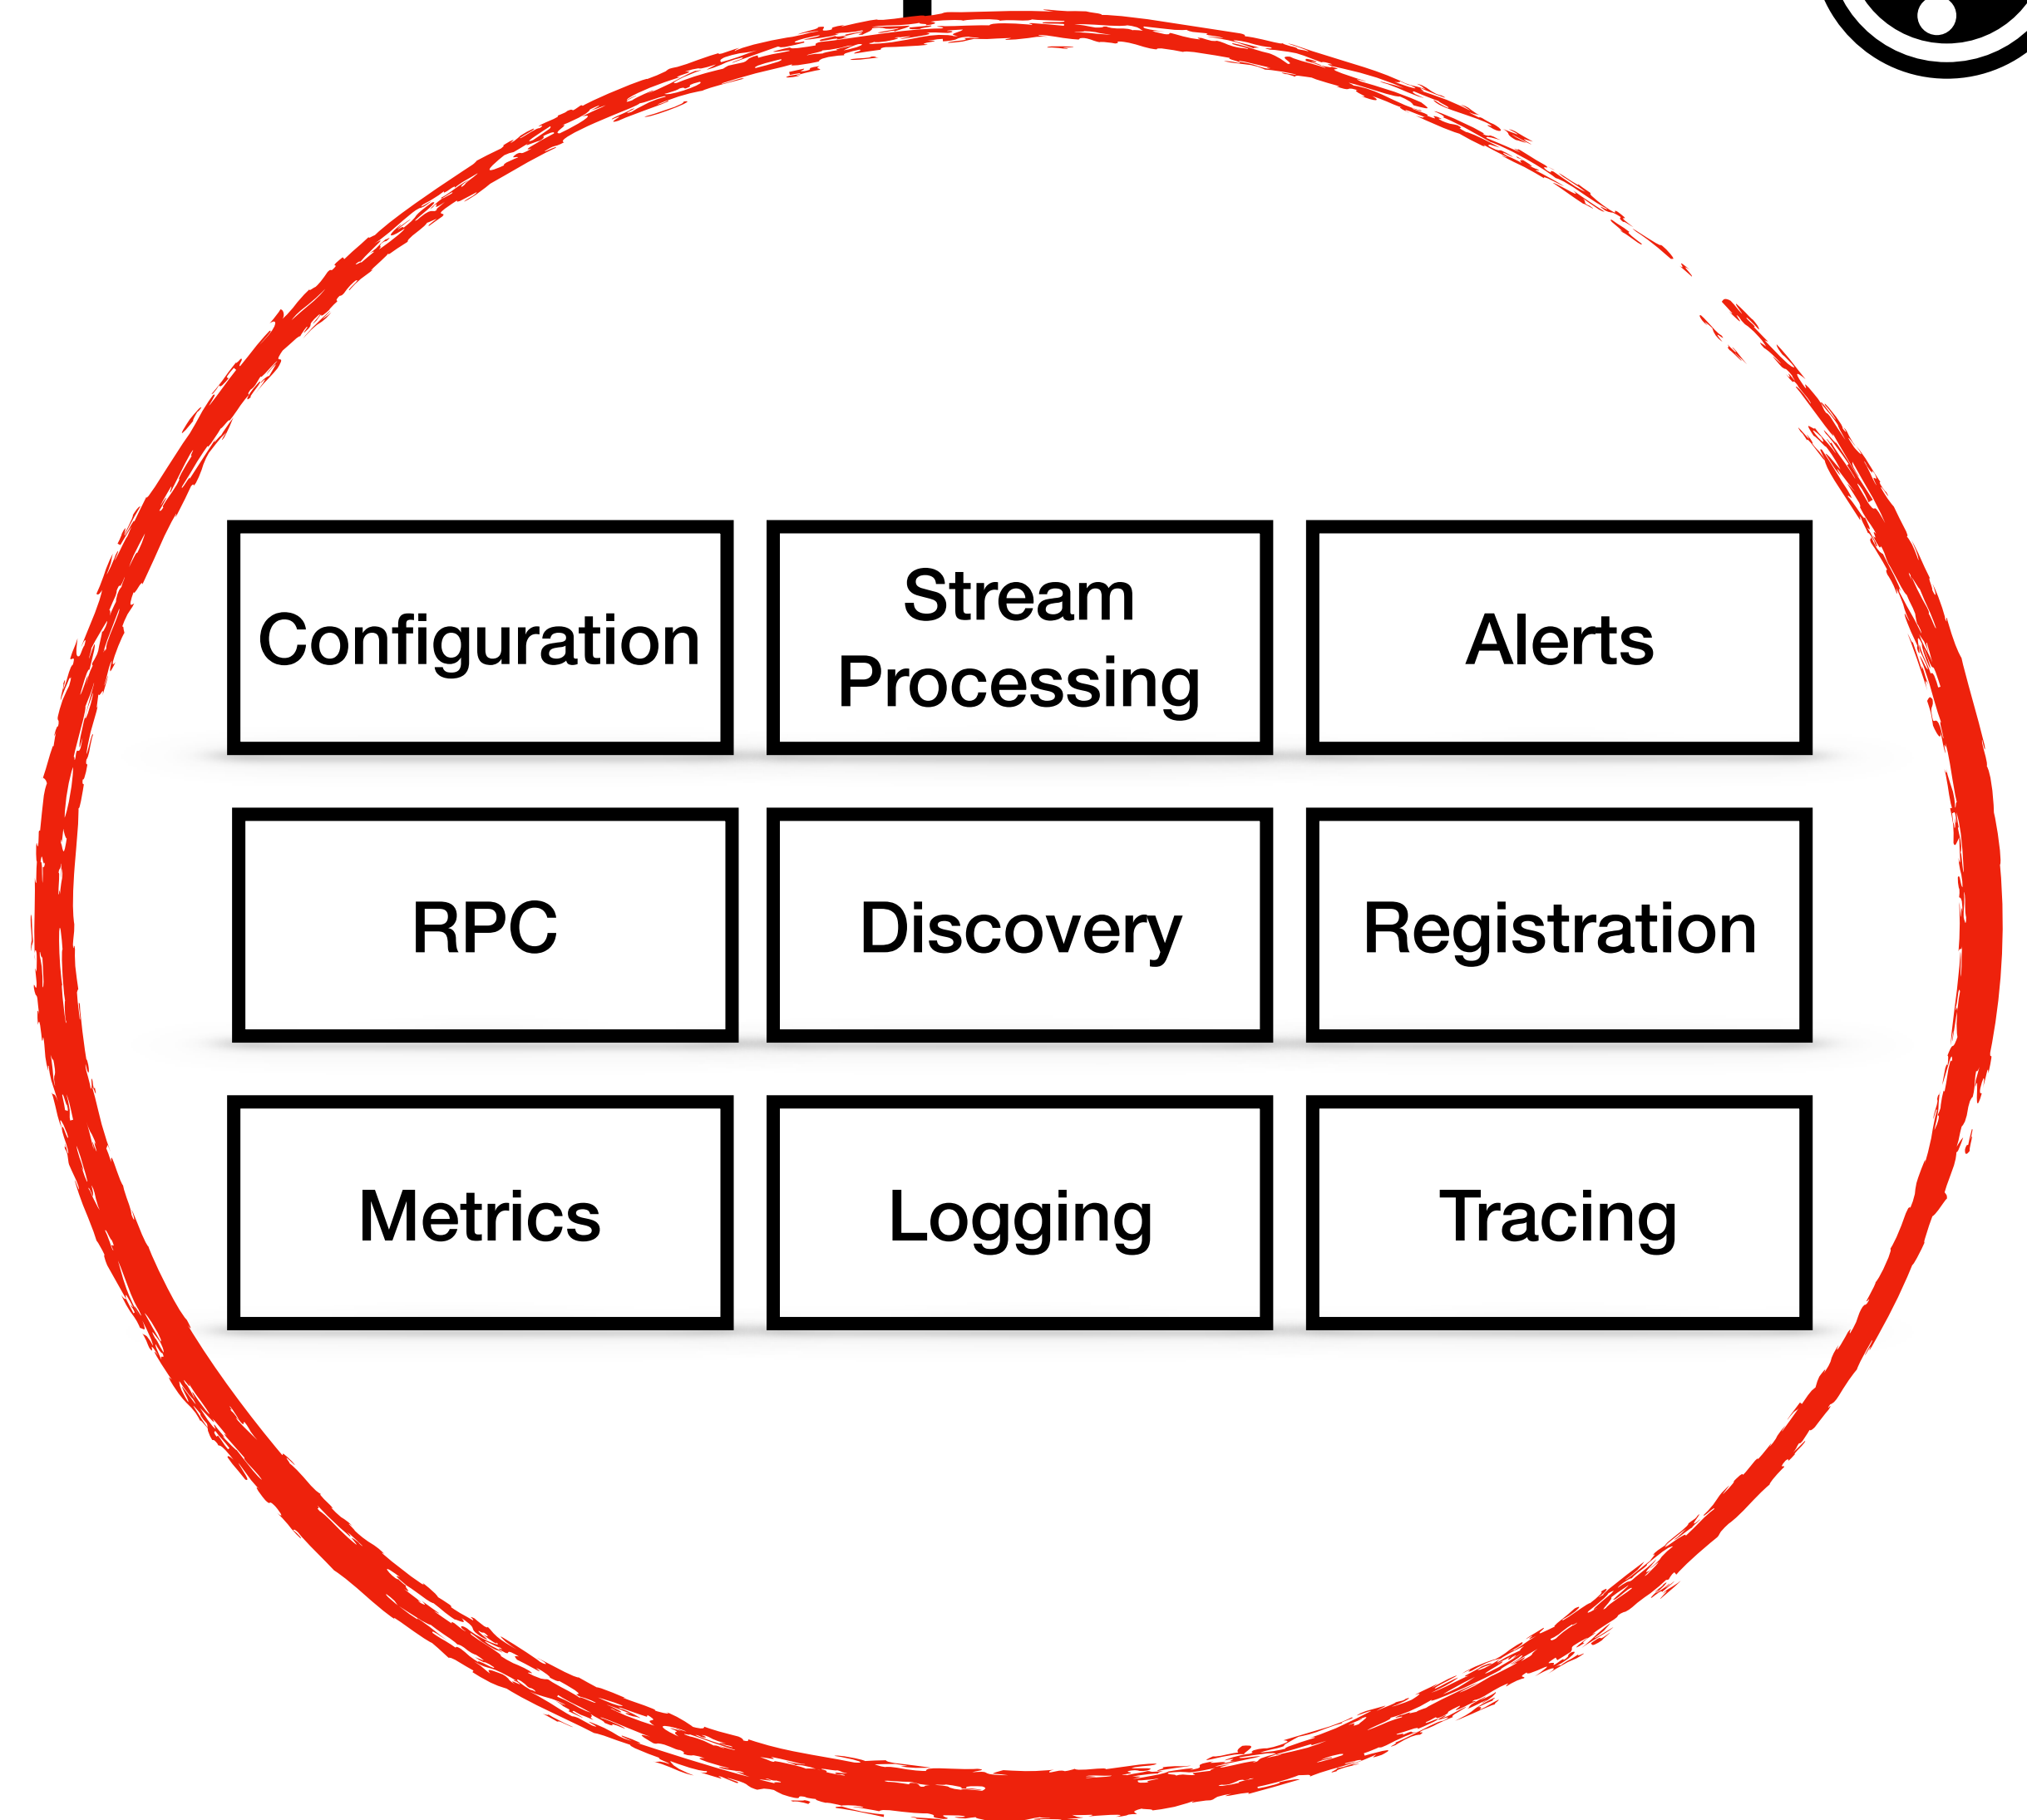


Mock Data Generation




+



Mock Ownership



Platform Testing API



- Just like a runtime API, need a testing API
- Provide mocks interface for components
- Gets platform out of the loop for providing mocks




Deployment



“Production is war!”




Experience Differences

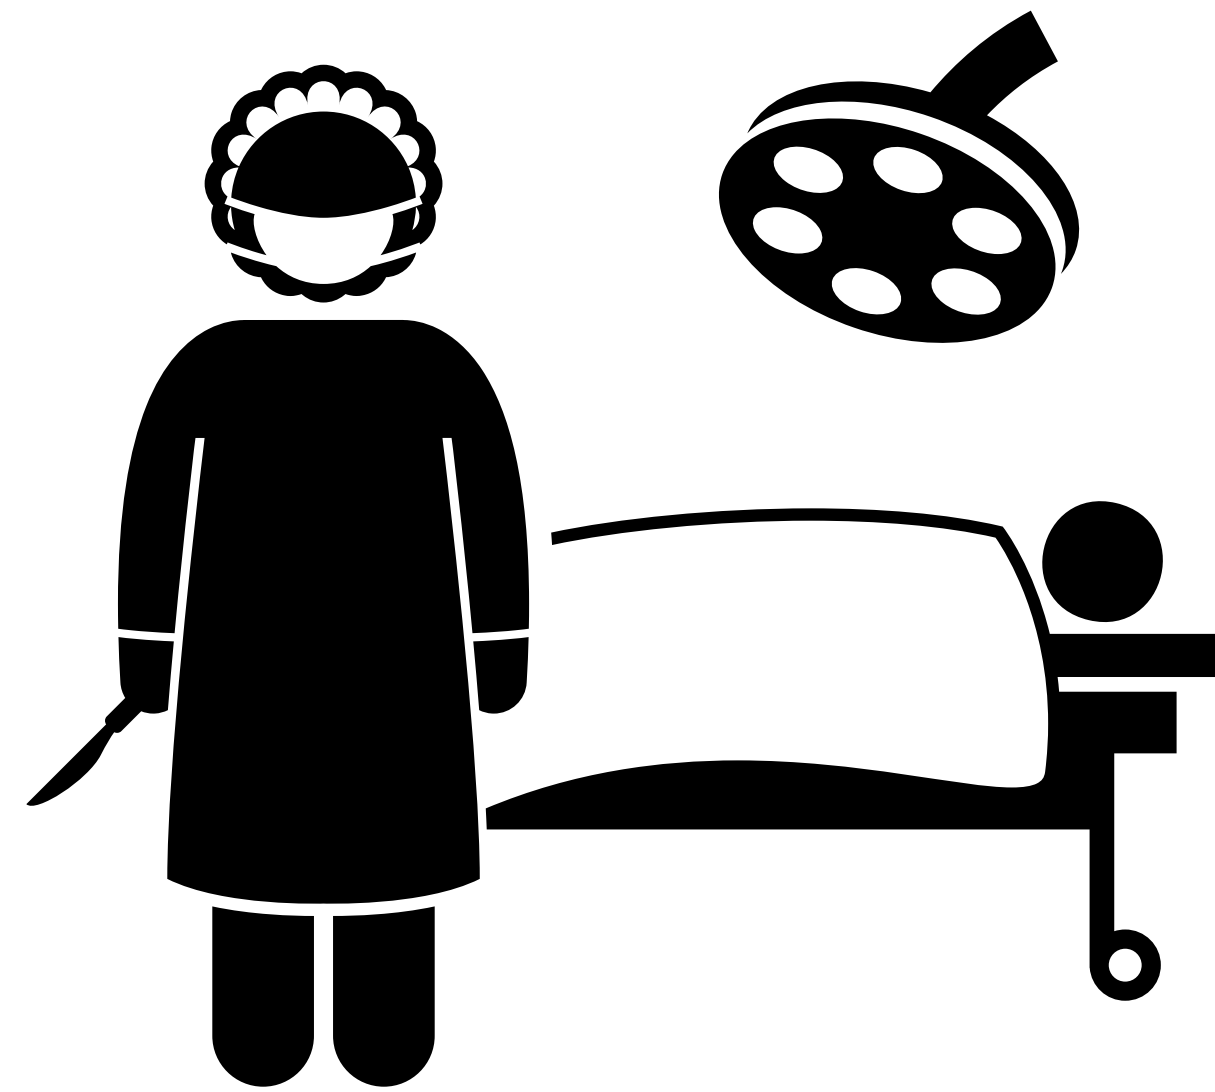


Deploy and Manage Services




- Pre-configured pipelines for deployment and rollback
- Single command deploy to any stack
- Integration for automated canary analysis
- Pre-configured autoscaling





Operations

Consolidated View



SERVICES > CHROMECAST



Stash Spinnaker Dashboard Alerts

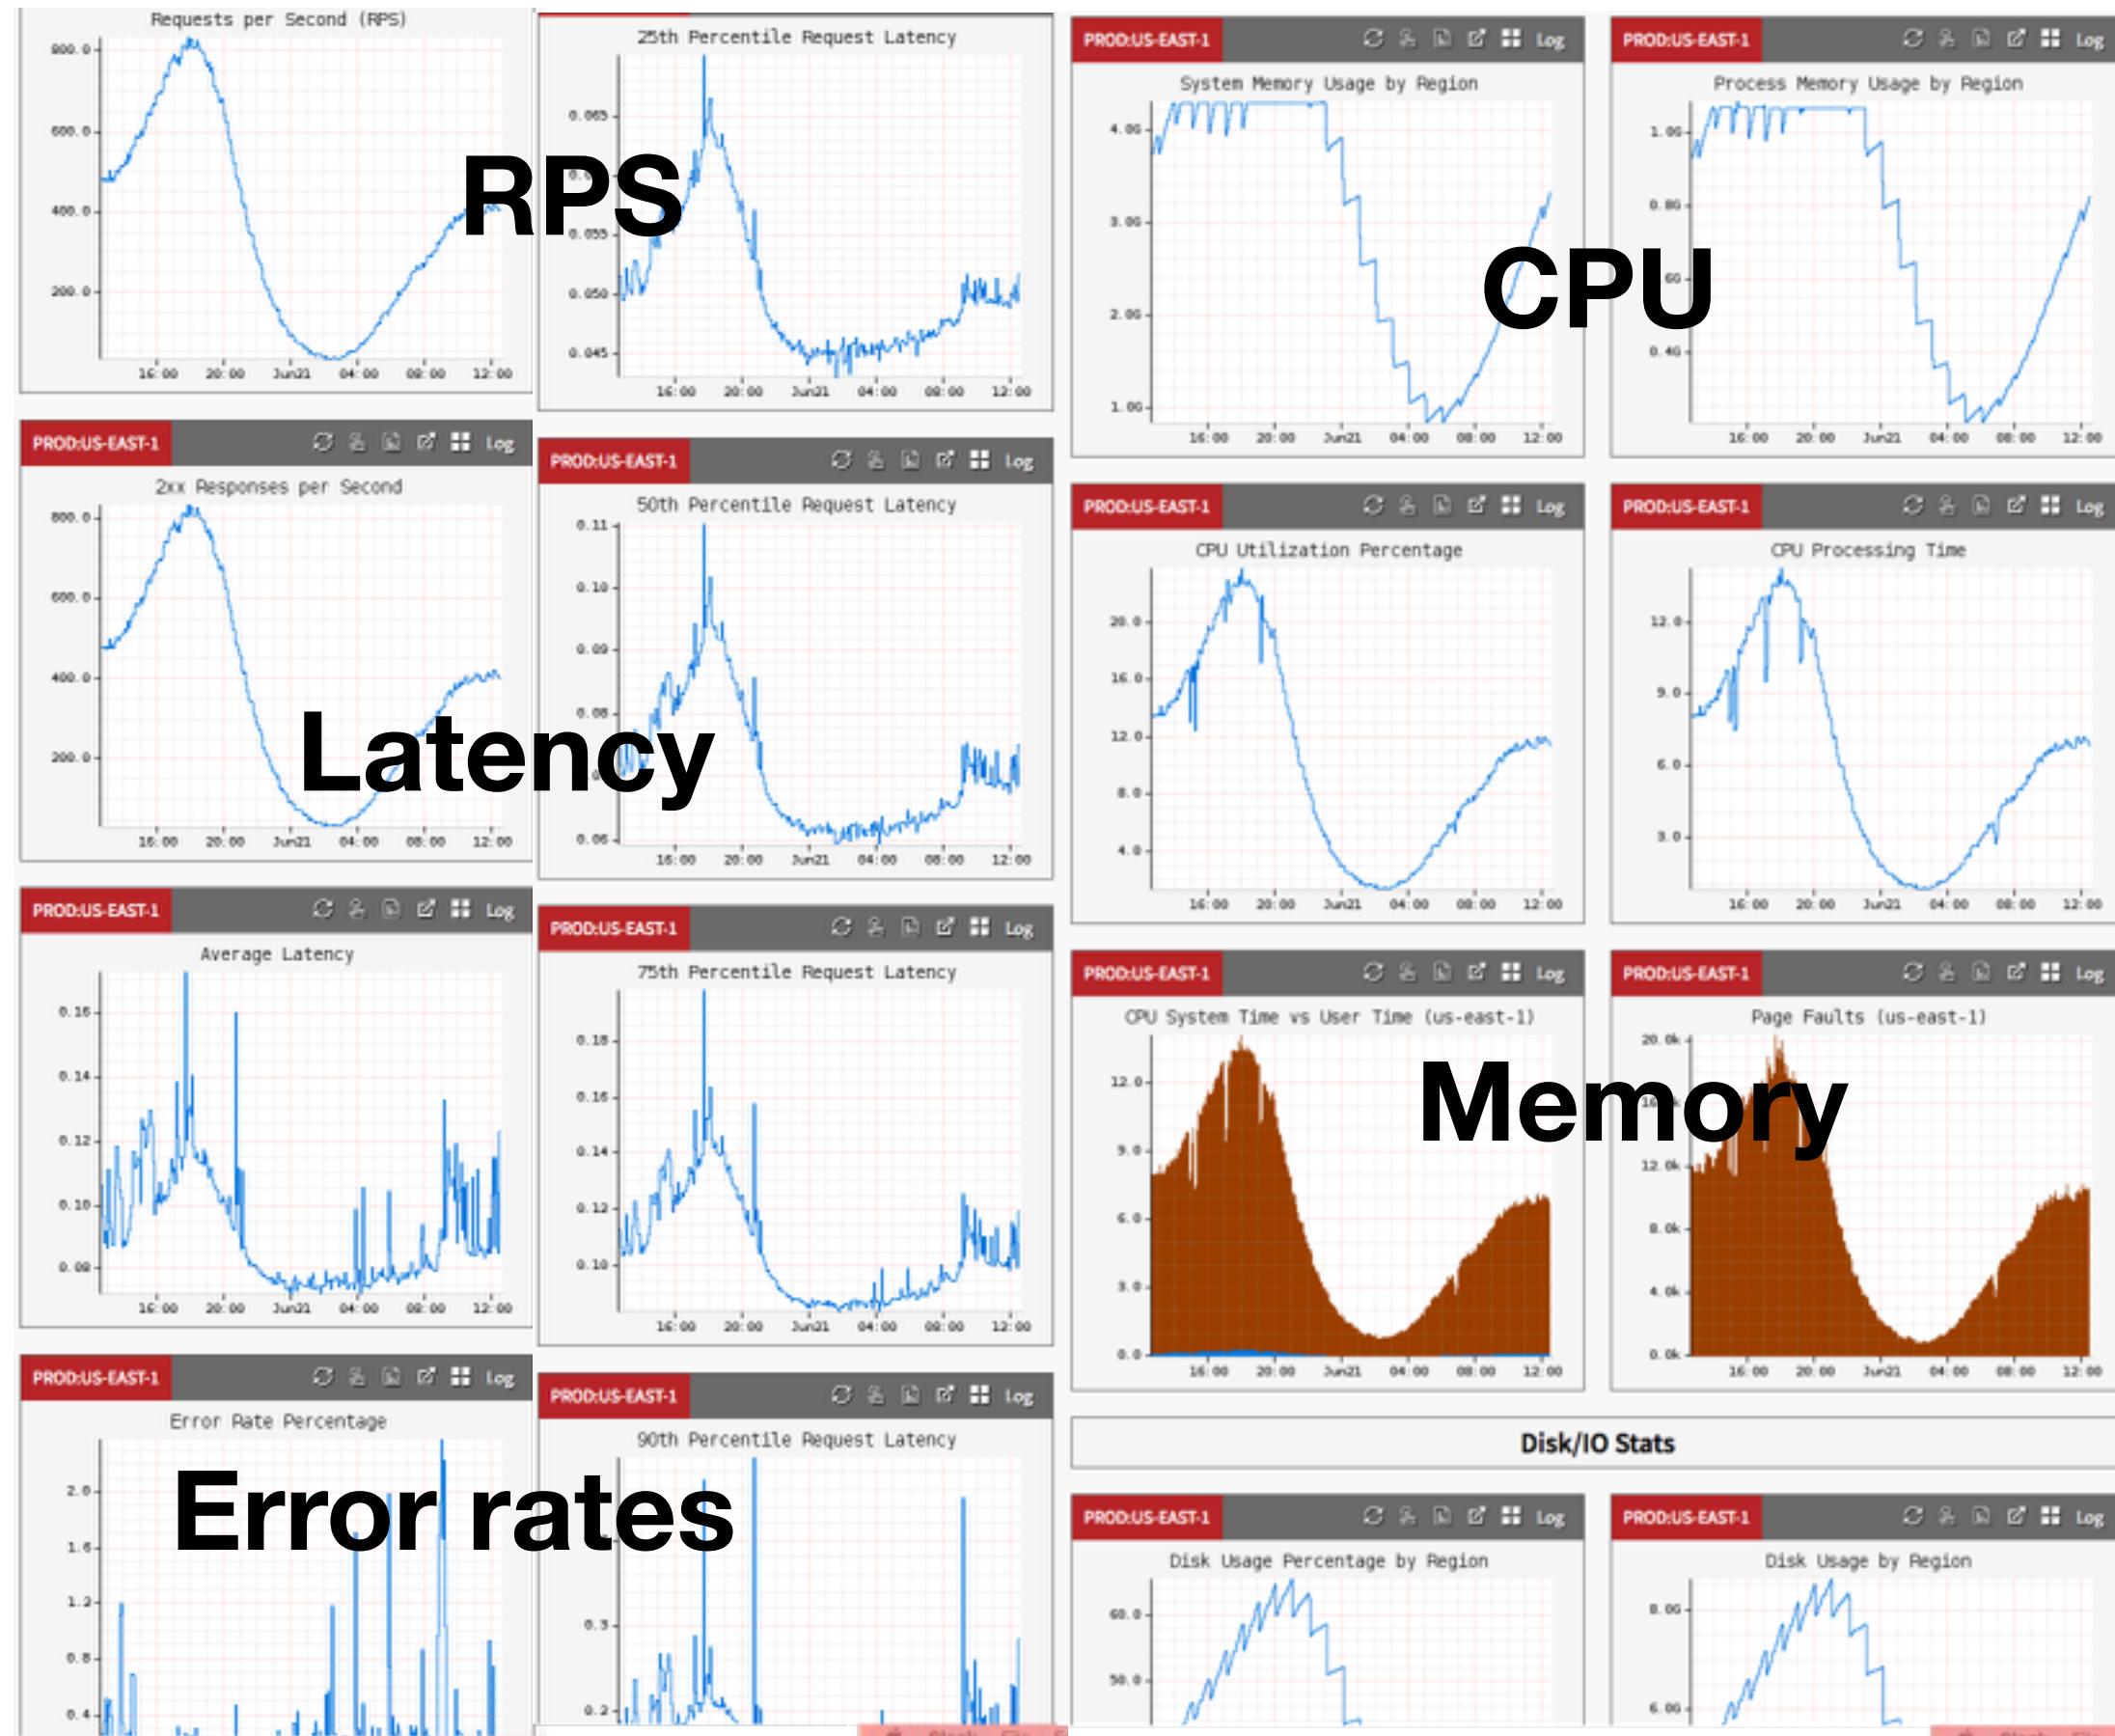

DEPLOYMENTS PUBLICATIONS CANARY EXECUTIONS RECENT ACTIVITY

STACK PROD 2 STAGING 14 INT 8 TEST 9

PROD	Version Constraint	Image Version	Region	Instance Counts	Created	Traffic (e-2w)	Dashboard	Stream Logs	Actions
^3.0.0		3.0.0	us-east-1	1▲ 0▼ 0◆ 0◇	2017-06-20 08:58:54 PDT				
		3.0.0	us-west-2	1▲ 0▼ 0◆ 0◇	2017-06-20 08:58:54 PDT				
		3.0.0	eu-west-1	1▲ 0▼ 0◆ 0◇	2017-06-20 08:58:55 PDT				
^1.0.0		1.0.32	us-east-1	30▲ 0▼ 0◆ 0◇	2017-06-07 10:28:12 PDT				
		1.0.32	us-west-2	30▲ 0▼ 0◆ 0◇	2017-06-07 11:09:15 PDT				
		1.0.32	eu-west-1	30▲ 0▼ 0◆ 0◇	2017-06-07 15:00:07 PDT				

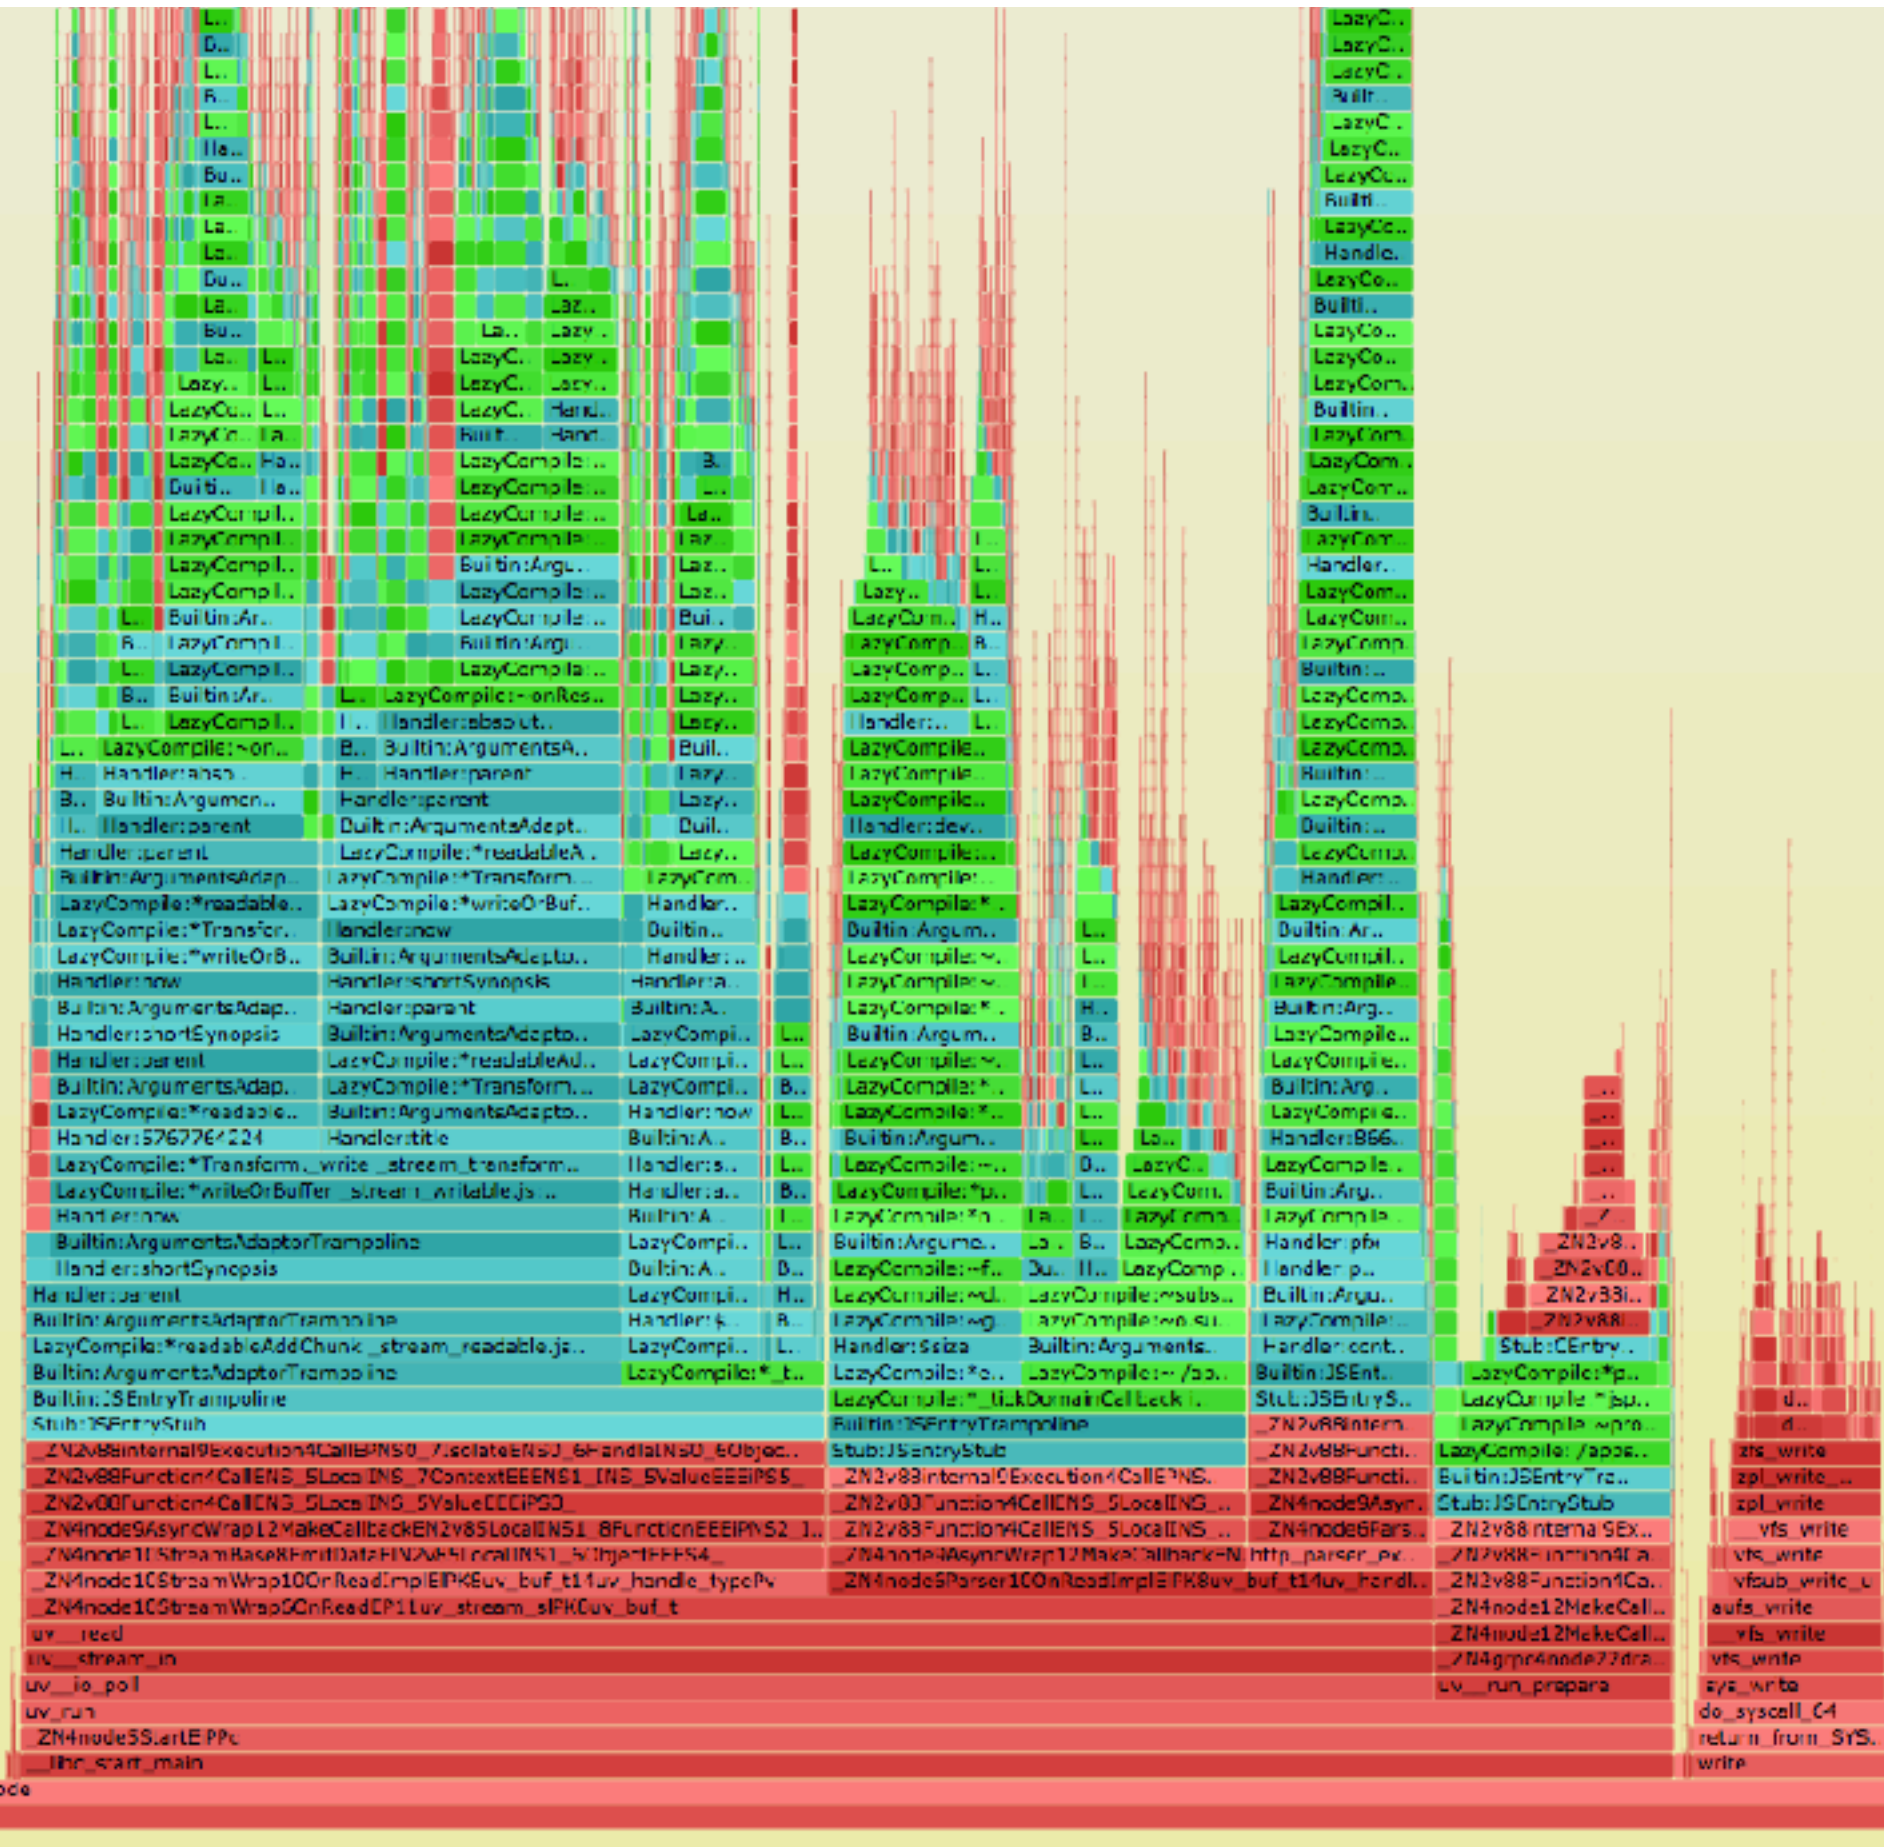
▲ Up (Taking traffic) ▼ Down ◆ Starting ◇ Disabled/Out of Service

Generated Dashboards & Alerts

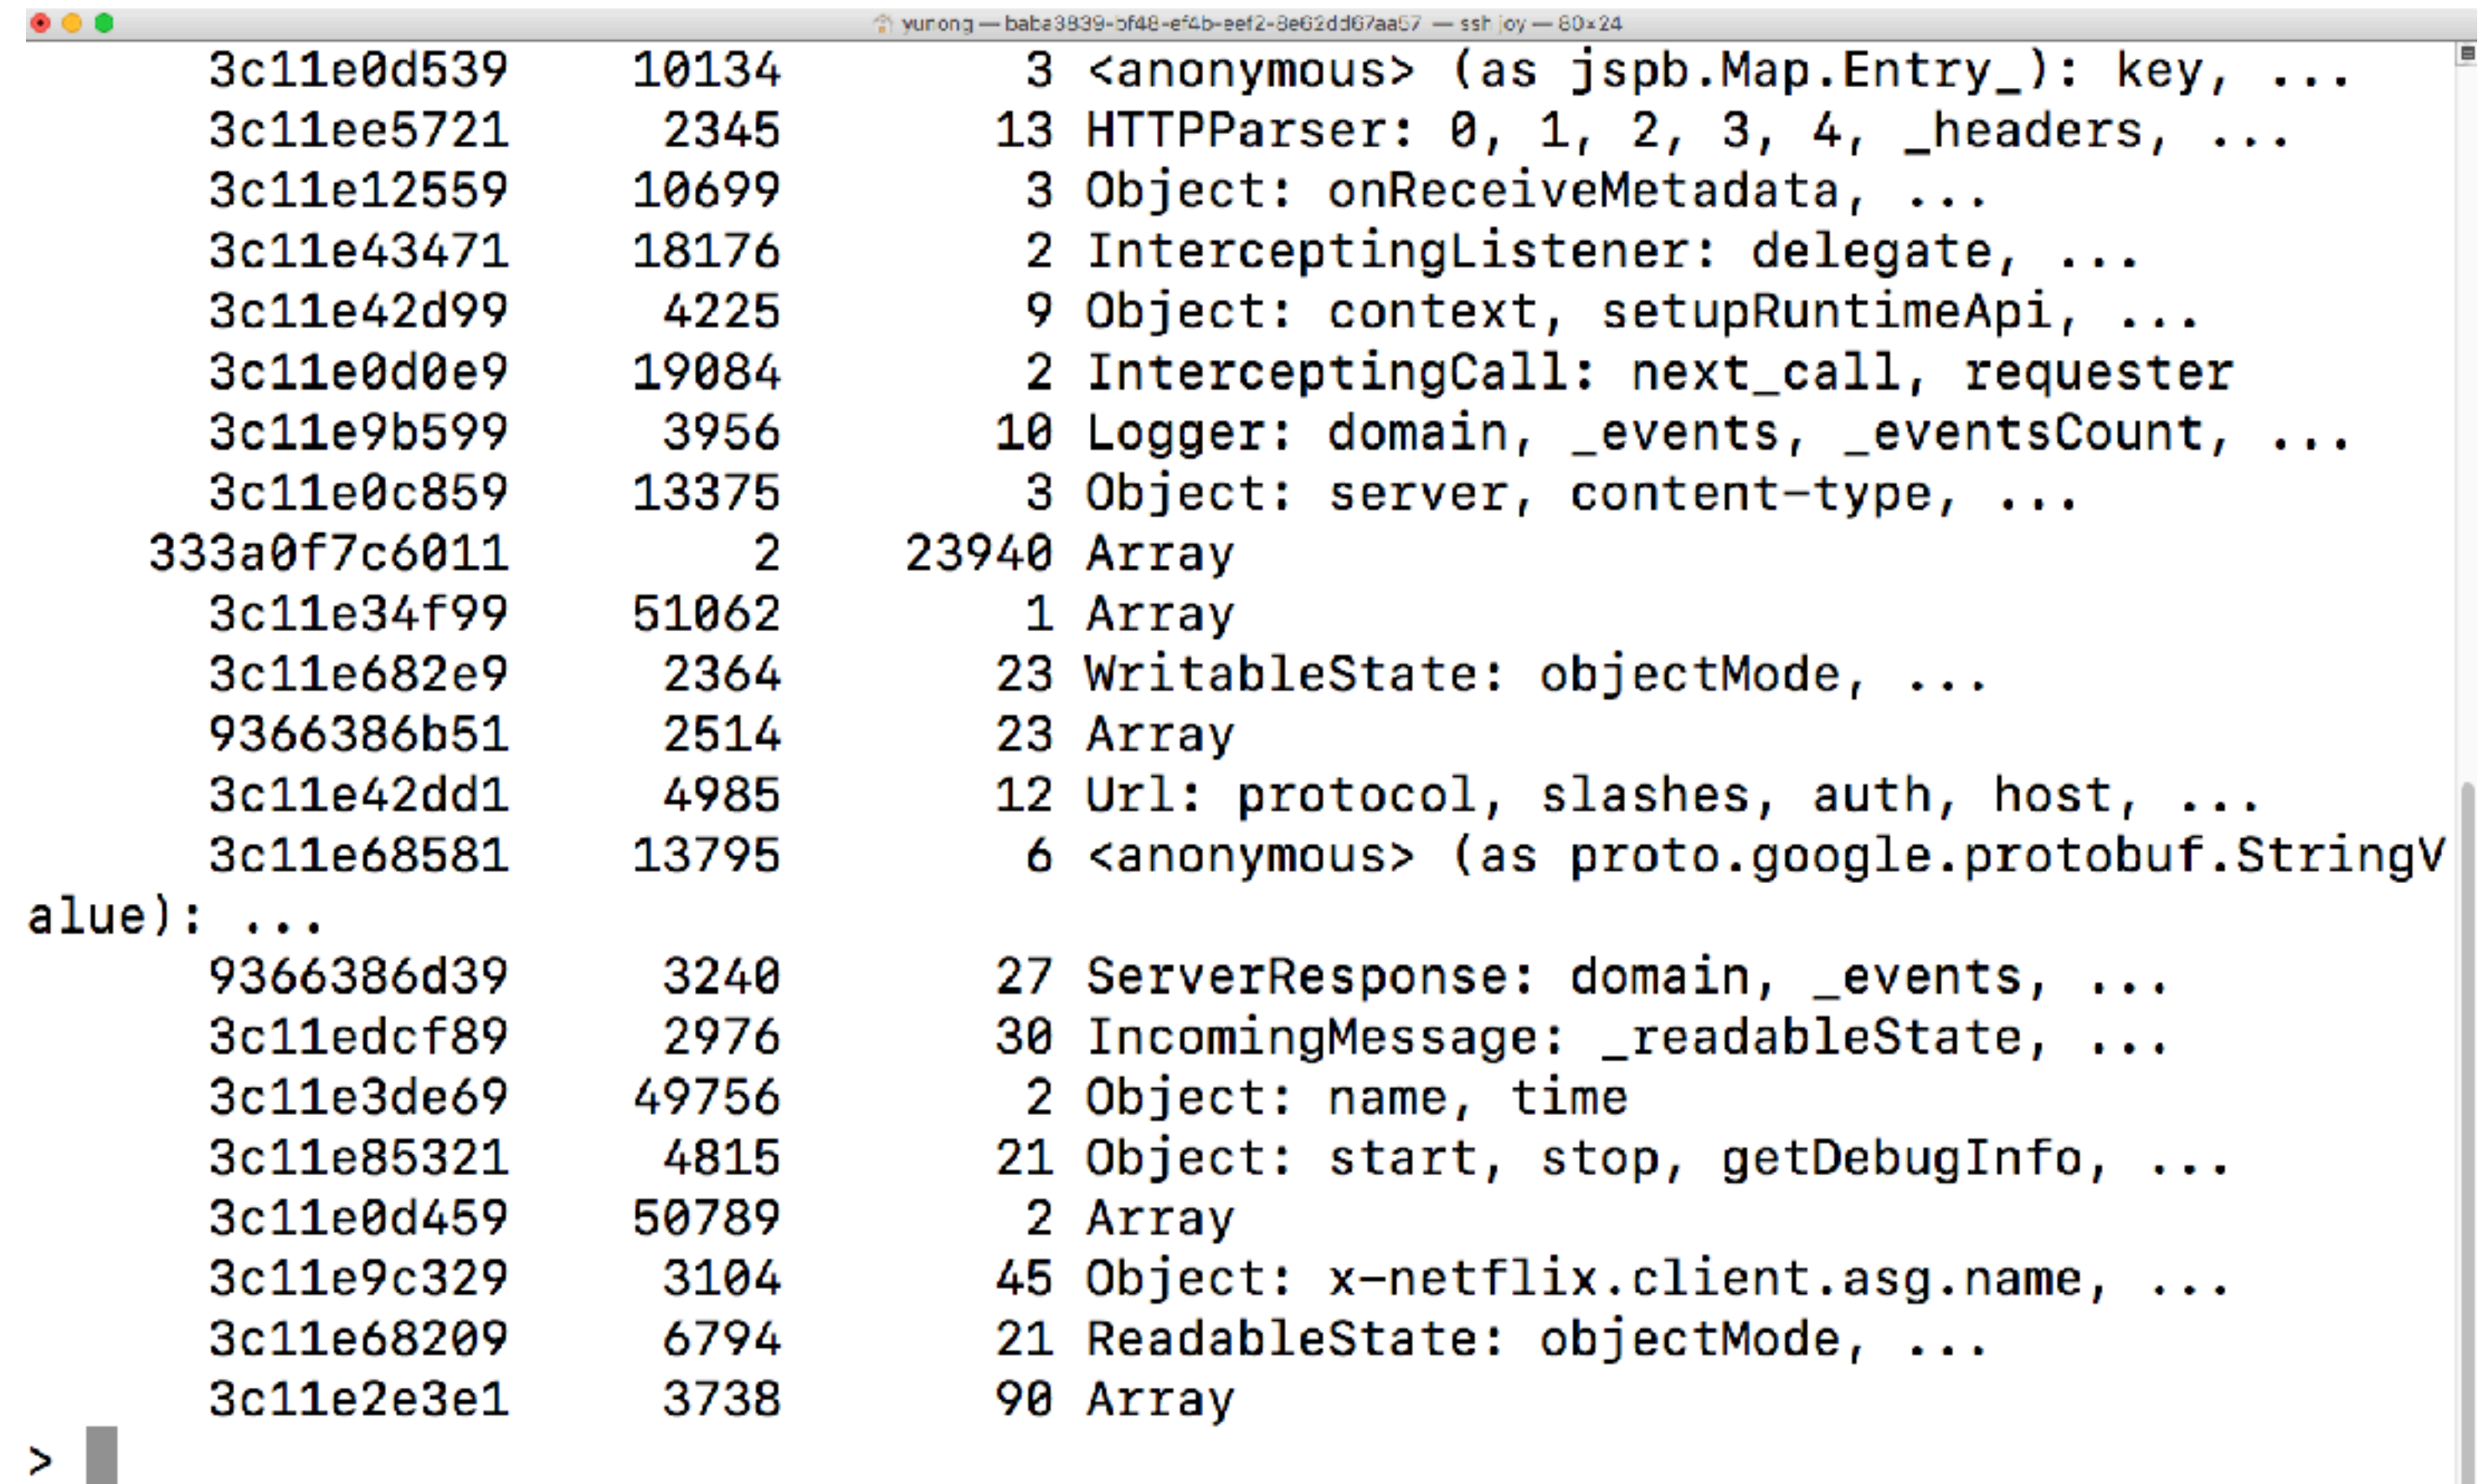


- alert-do-not-reply** Jun 22
 Average latency is more than 1s for r [redacted] in eu-west-1
Prod eu-west-1 High_Latency Summary: Average latency is more than 1s for r [redacted] (pr
- alert-do-not-reply (4)** Jun 22
 High Process Starts for r [redacted] in us-west-2
Prod us-west-2 High_Process_Starts Summary: High Process Starts for r [redacted] (prod) ir
- alert-do-not-reply** Jun 22
 Average latency is more than 1s for [redacted] n us-west-2
Prod us-west-2 High_Latency Summary: Average latency is more than 1s for [redacted] pr
- alert-do-not-reply (4)** Jun 22
 High Process Starts [redacted] in us-east-1
Prod us-east-1 High_Process_Starts Summary: High Process Starts for r [redacted] (prod) in
- alert-do-not-reply (4)** Jun 22
 High Process Starts for [redacted] (prod) in eu-west-1
Prod eu-west-1 High_Process_Starts Summary: High Process Starts for [redacted] (prod) ir
- alert-do-not-reply (2)** Jun 22
 Disk Usage is over 80% for r [redacted] (prod) in us-east-1
Prod us-east-1 High_Disk_Usage Summary: Disk Usage is over 80% for r [redacted] (prod) i
- alert-do-not-reply** Jun 22
 Disk Usage is over 80% for [redacted] (prod) in us-west-2
Prod us-west-2 High_Disk_Usage Summary: Disk Usage is over 80% for [redacted] (prod) i
- alert-do-not-reply (2)** Jun 21
 [redacted] Availability dropped - Action Required
Prod us-west-2 PBO_Availability Summary: F [redacted] availability dropped - Action Required Check time

Automated Analytics & Tooling



CPU profiling



Core dump analysis

Our PaaS to Velocity & Reliability

1. Standardized components
2. Preassembled platform
3. Automation and tooling



June 26, 2017

The Paved PaaS to Microservices

Yunong Xiao,
Principal Software Engineer, Netflix
yunong@netflix.com, [@yunongx](https://twitter.com/yunongx), <http://yunong.io>