

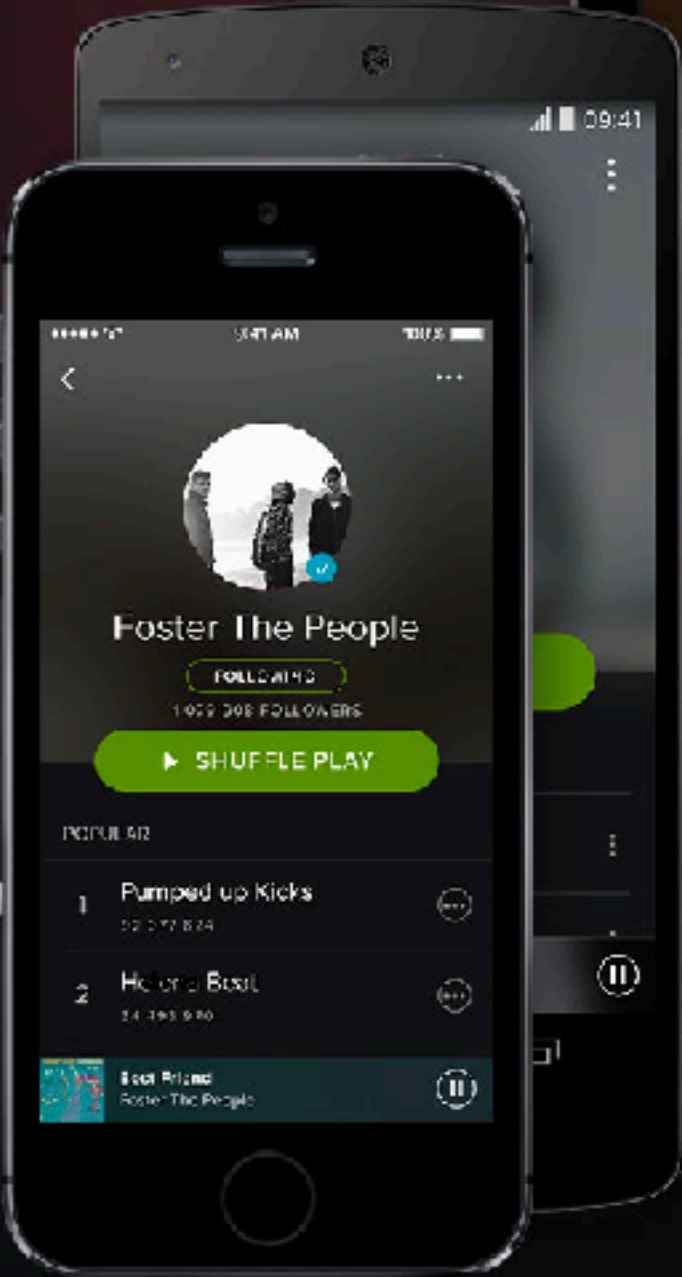
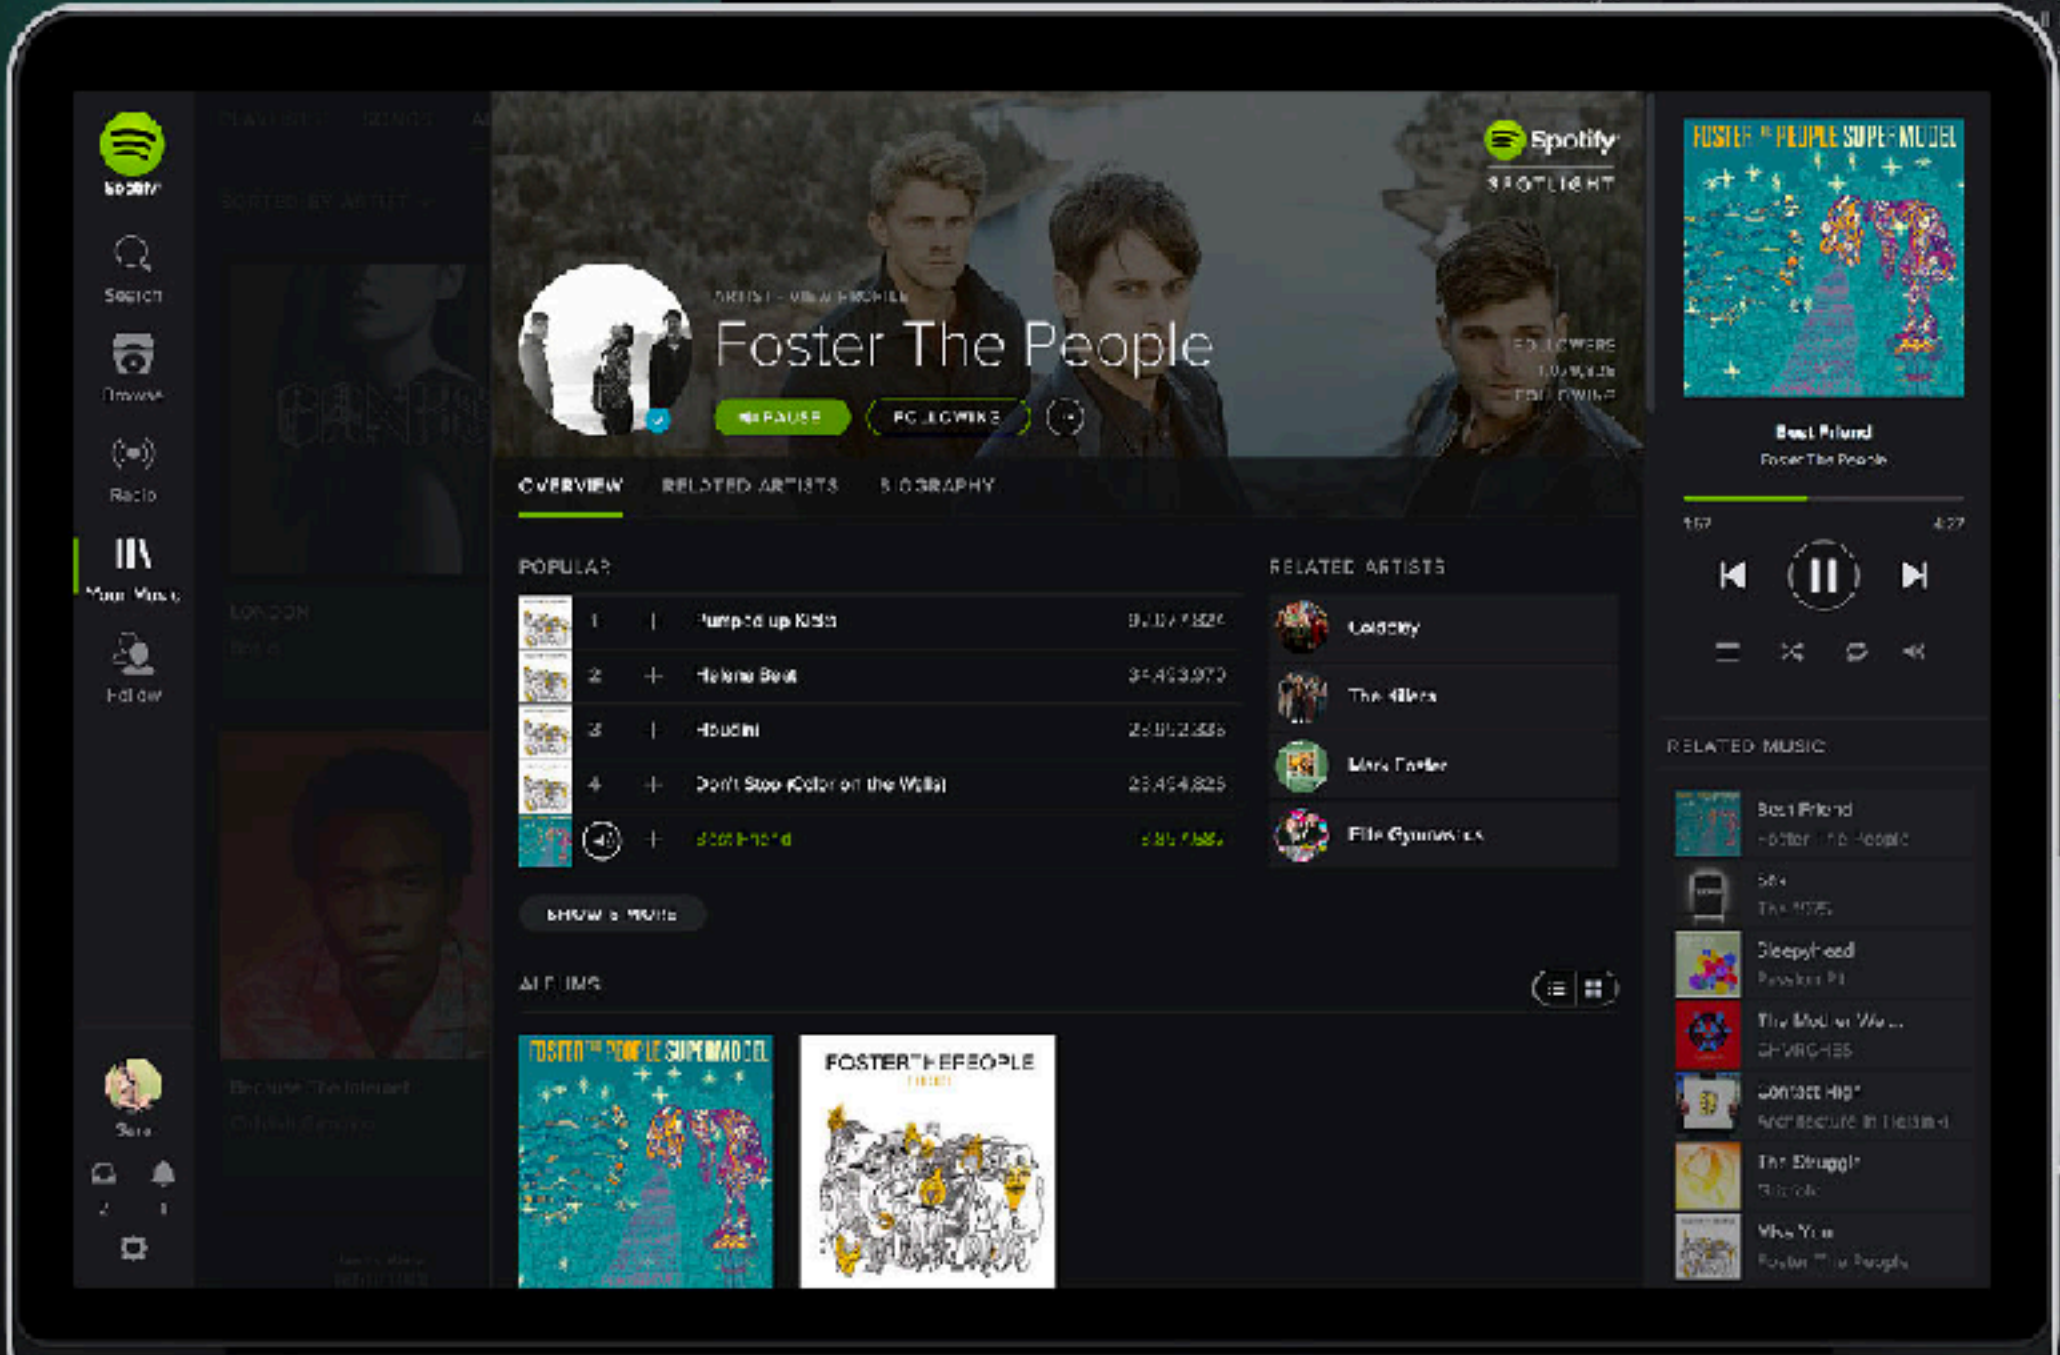
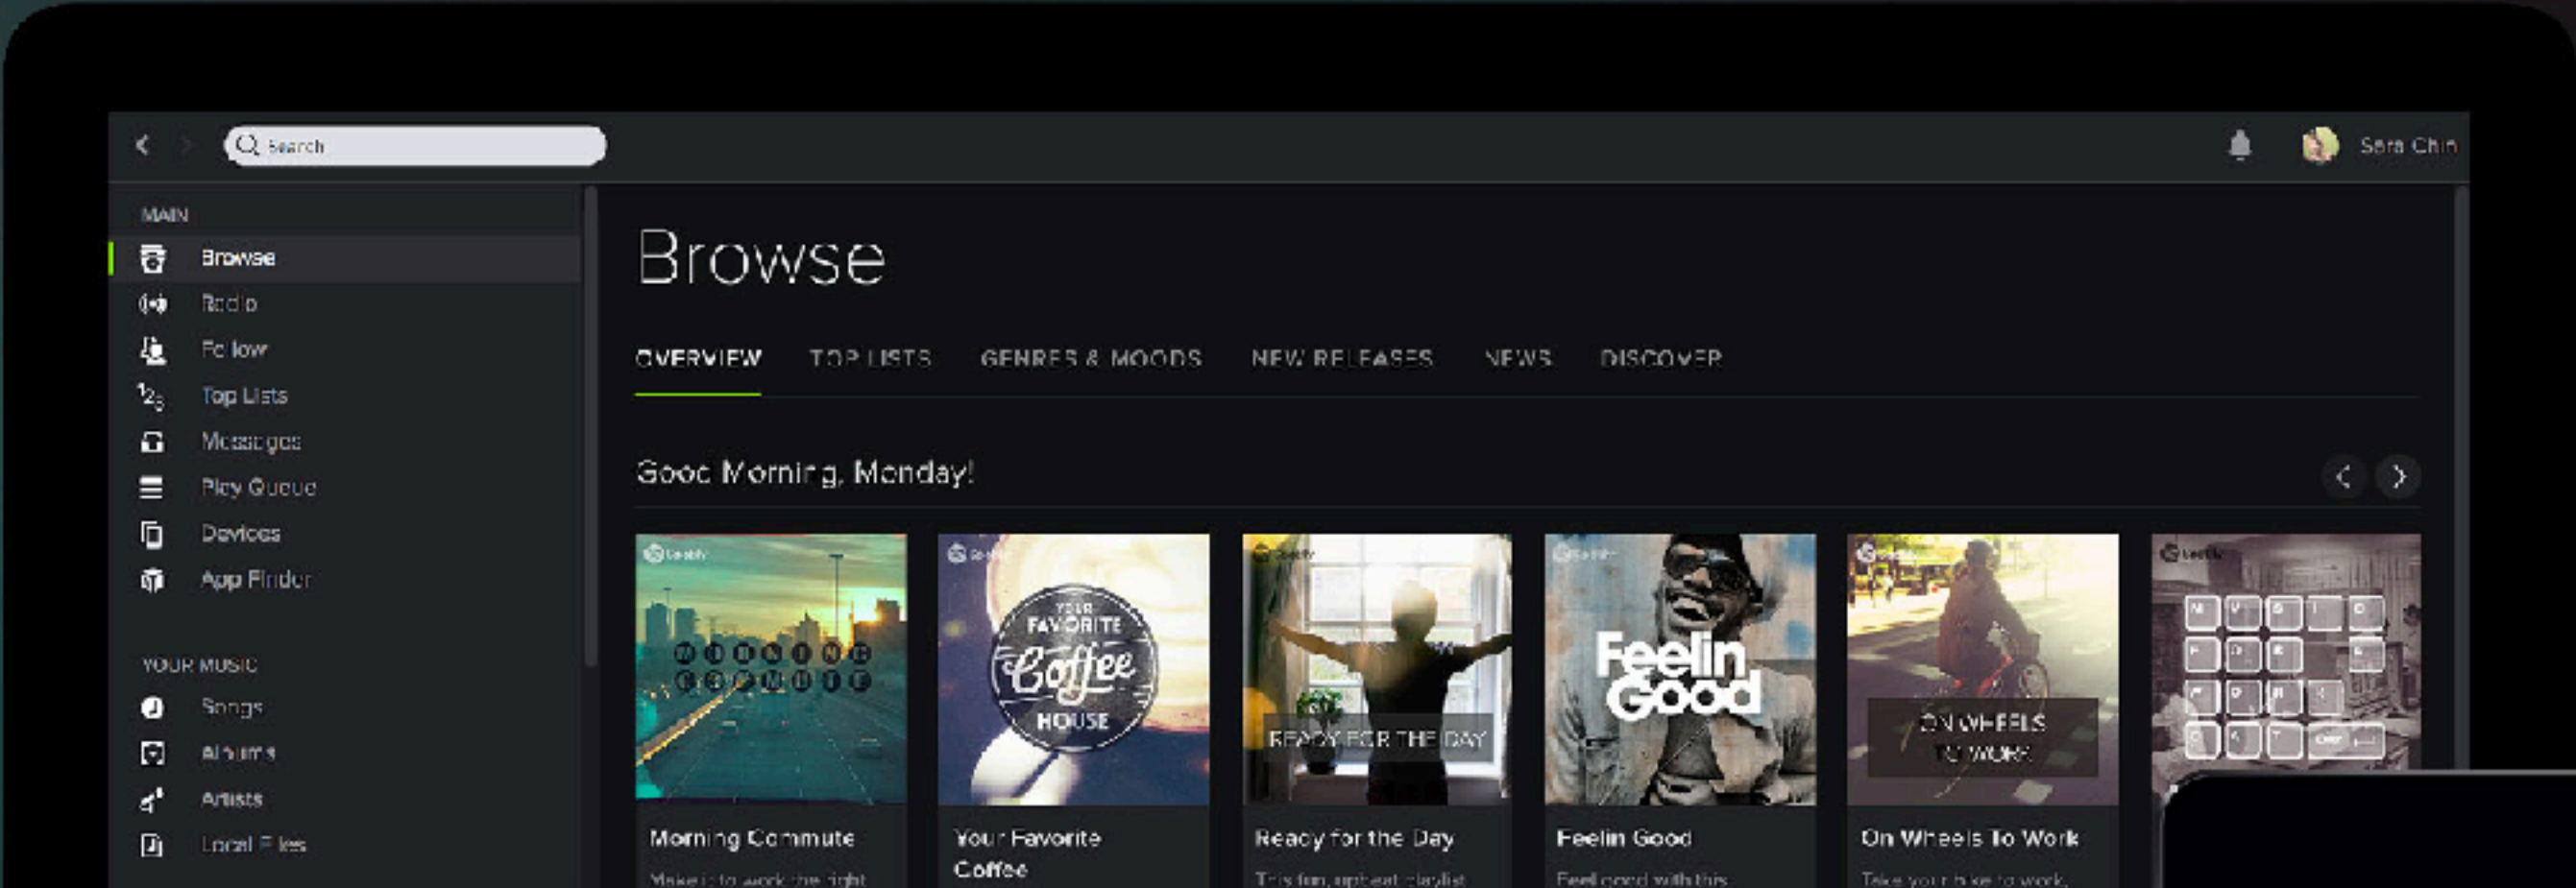
# Spotify Lessons: Learning to Let Go of Machines

James Wen, Site Reliability Engineer at Spotify  
ALF Squad, Infrastructure & Operations Tribe



IO Tribe





**Let's control how  
feature developers think  
about what their code is  
actually running on.**



# Takeaways

- **Feature developers = happiest with feature work**
- **Find out developer machine concerns and mitigate**
- **Migrating to cloud or hybrid? Start embracing ephemeral service design and infrastructure**

# Agenda

- **Why?**
- **Journey**
  - **Hybrid Cloud**
  - **Ops in Squads**
- **Future**
- **Learnings**

# Why?

**Why don't we want feature devs to care too much about infrastructure and machines?**

# Why?

**Time taken on infrastructure tasks = time taken away from feature work**

**Feature devs = focused on features**

# Spotify Scale Stats

- **140 Million+ Monthly Active Users**
  - **50 Million+ Subscribers**
    - **30 Million+ Songs**
    - **2 Billion+ Playlists**
  - **Available in 60 markets**





# Spotify Dev Scale Stats

**~900 Devs**

**~100 Tech Teams**

**~2000 Services**

# Spotify Machine Scale Stats

~10,000 Bare Metal Hosts

~13,000 Hosts on GCP

46 Hardware/VM Types

# Example: Capacity Planning

$$TIME\_SPENT = \frac{w \text{ devs}}{x \text{ teams}} * \frac{y \text{ hours}}{\text{service}} * z \text{ services}$$

Avg # devs on a team      Capacity Planning

$$\frac{900 \text{ devs}}{100 \text{ teams}} * \frac{1 \text{ hour}}{\text{service}} * 2000 \text{ services} = 18,000 \text{ dev hours}$$

# Scale doesn't really matter

- **Smaller companies/teams = developer time is more valuable**
- **Larger companies/teams = wasted infra time scales as well**

# Other Infrastructure Tasks

- Machine provisioning
- Failure planning
- Security updates
- Machine maintenance

# Dedicated Ops?



# Dedicated Ops?

**~2000 Services**

**74 Infrastructure and Operations  
Engineers**

**If all IO engineers → dedicated ops**

**27:1 service:engineer ratio**



# Ops In Squads

**Feature teams handle their own ops and provisioning**


**Using the services and tooling the Infrastructure and Operations tribe has written**



**We control the level of context feature teams need to operate their services.**

- **Developer Happiness**
- **Developer effectiveness and context**

# Journey



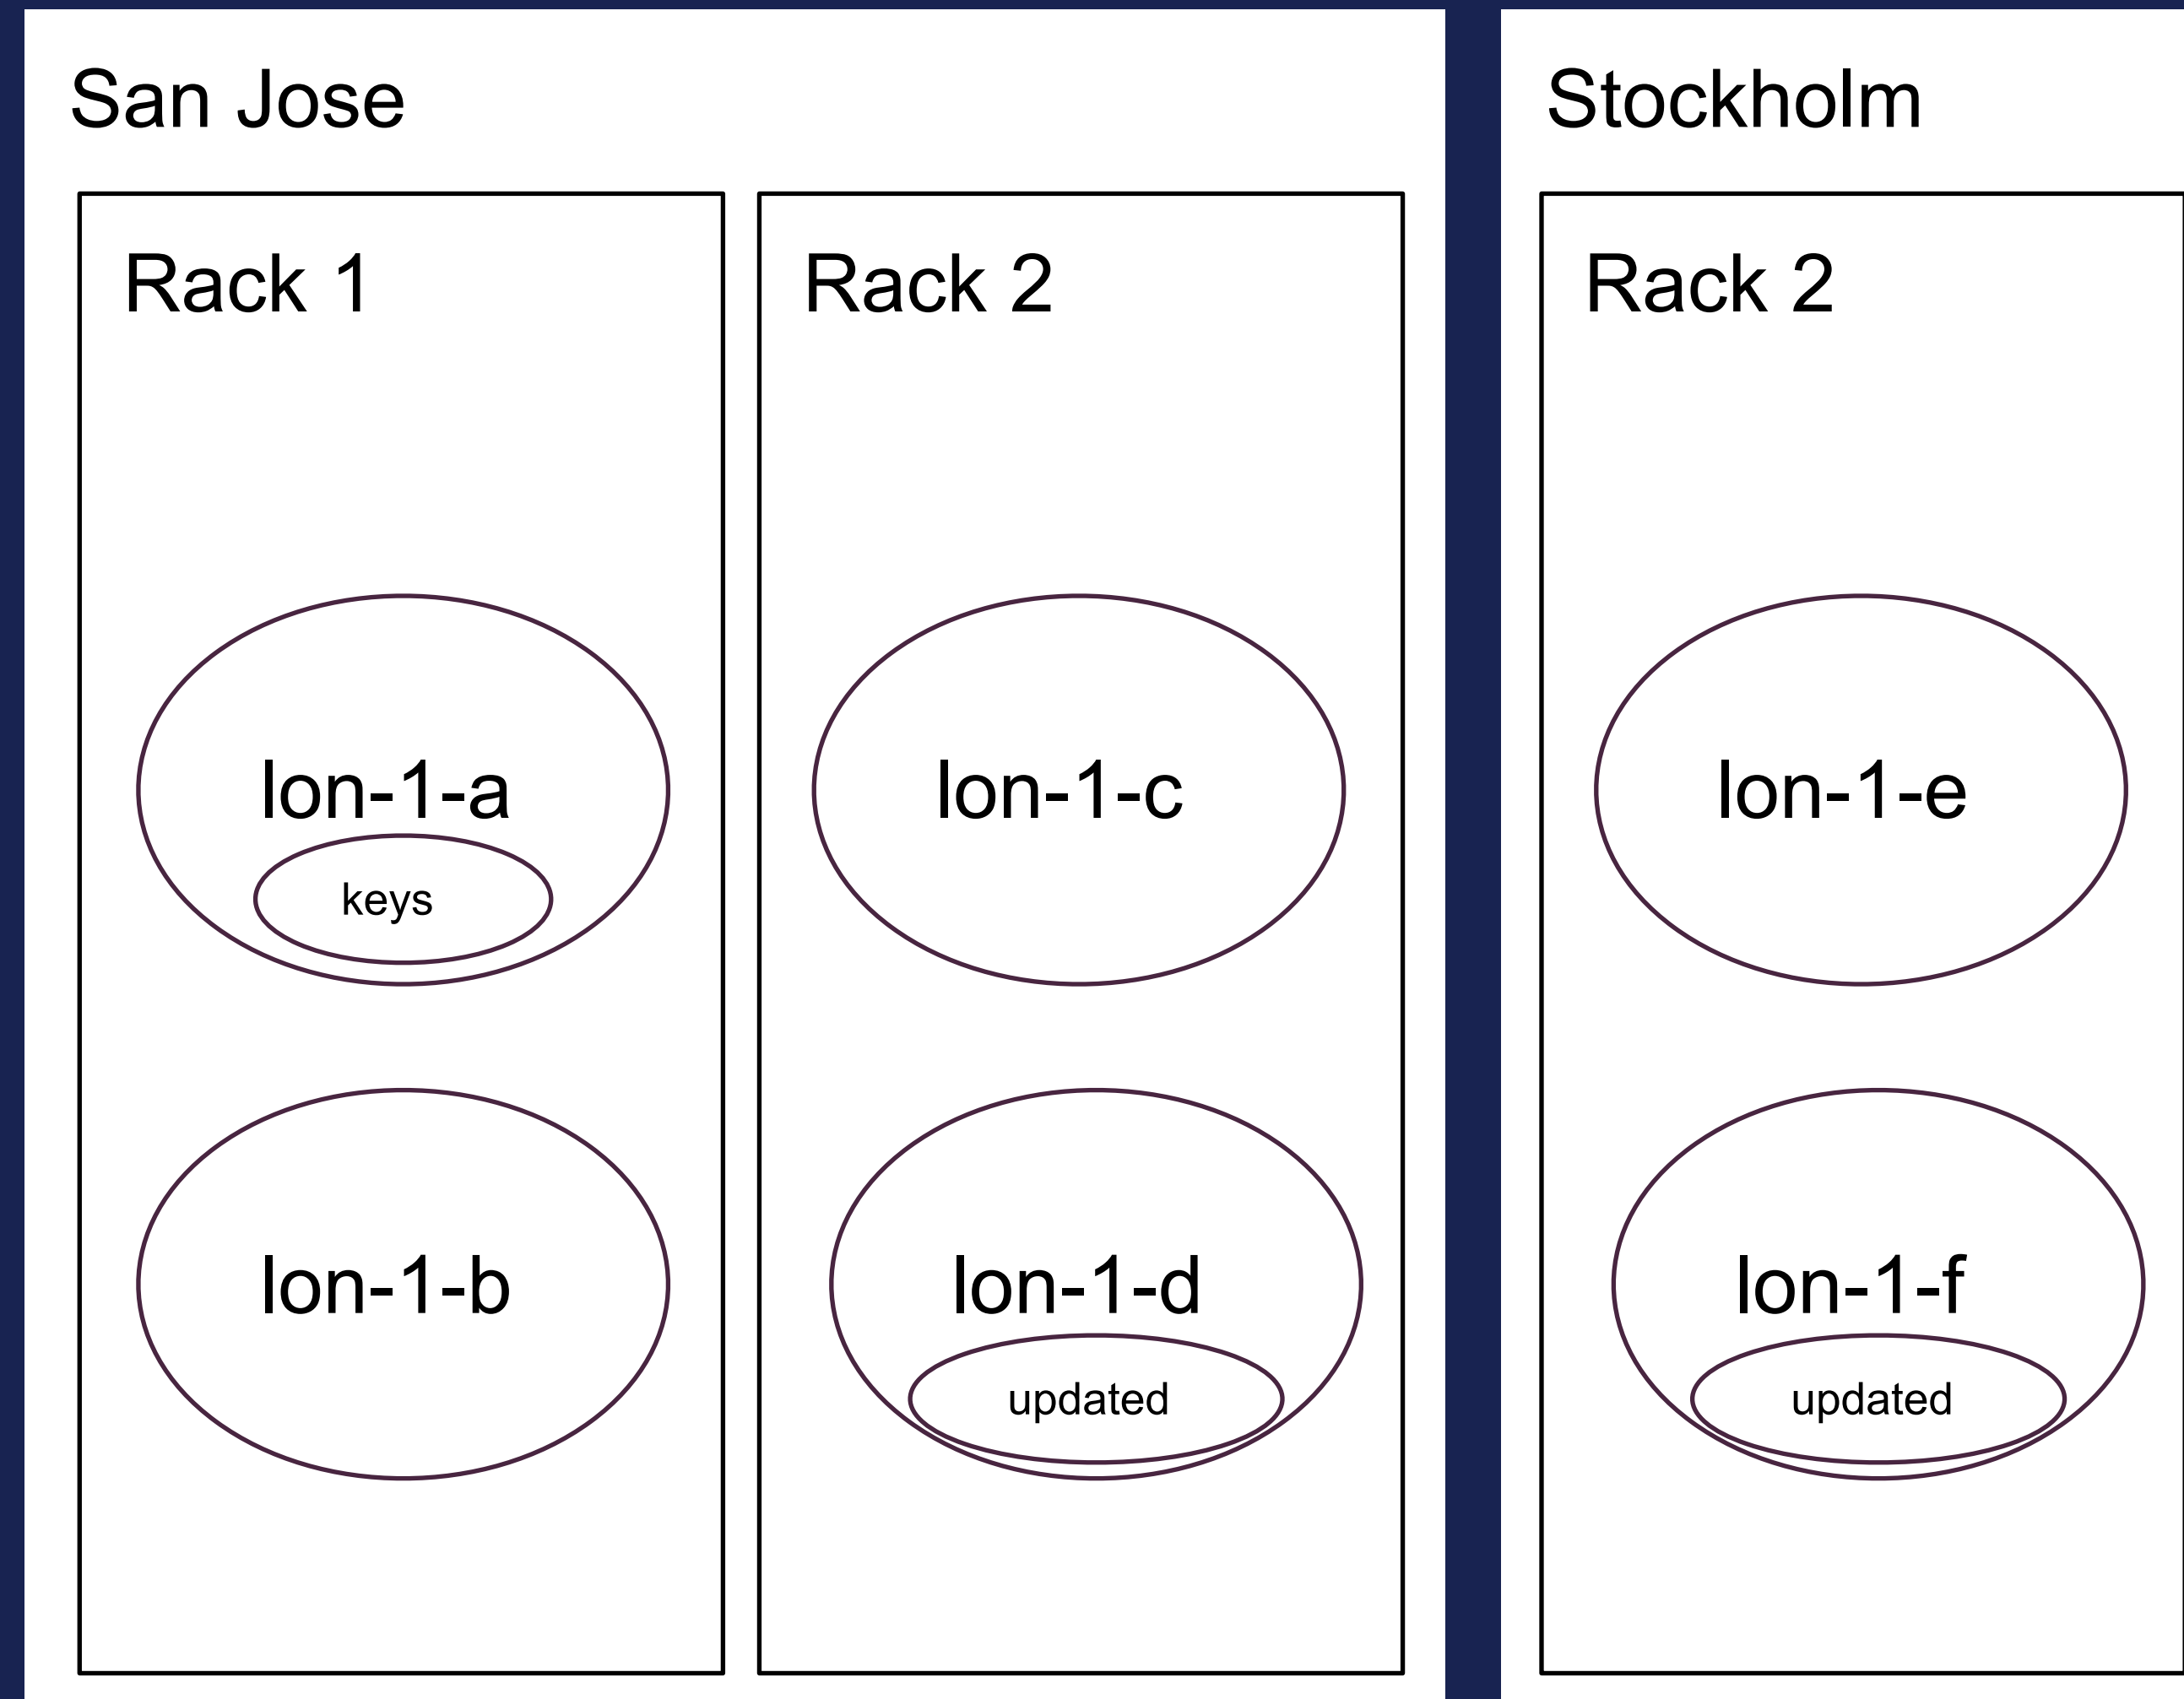
**- Ops in Squads**

**- Hybrid Cloud  
(Ephemerality)**

# Starting Out



# Historical: Feature Developer's Context for Service's Capacity



# Machine Context

Unbound  
v1.6.3

Openssl  
v1.0.0f

2 Cores

8 GB  
RAM

3 Years

In Virginia

Tarred Logs

ash2-metadata-a.ash2.spotify.net

- Packages
- Hostname
- Machine specs (CPU, RAM, disk, etc.)
- Uptime and service duration
- Location
- Local state (files on disk, info in memory)

# Feature Developer Concerns



How to track?

How to get?

How to talk to it?

Maintenance?

What tools on it?

What to put on it?

Where?

Up to date?

How long?

Available?

Specs?

How many?

Service + Business







# Feature Developer Concerns



How to track?

How to get?

How to talk to it?

Maintenance?

What tools on it?

What to put on it?

Where?

Up to date?

How long?

Available?

Specs?

How many?

Service + Business



# ServerDB

Overview	Facts	Properties	Changelog
<b>Hardware name</b>	72GDF62		
<b>Name</b>	72GDF62.lon6.spotify.net		
<b>State</b>	available		
<b>Order ID</b>	<a href="#">lon6-expansion-2015-08</a>		
<b>Pod / Site</b>	lon6 / lon		
<b>Rack</b>	ms4f11		
<b>Hardware Type</b>	HighCPU/64GB:201505		
<b>Chassis</b>			
<b>Comment</b>			
<b>Puppet classes</b>			
<b>Network interfaces</b>			
NIC	MAC	IP-addresses	
eth0	ec:f4:bb:dd:a5:00	10.243.45.243 (lon6-srv-ms4f11)	
eth1	ec:f4:bb:dd:a5:01		
eth2	ec:f4:bb:dd:a5:02		
eth3	ec:f4:bb:dd:a5:03		
lom	74:e6:e2:fc:f1:22	10.244.140.94 (lon6-lom)	

# Feature Developer Concerns



How to track?

How to get?

How to talk to it?

Maintenance?

What tools on it?

What to put on it?

Where?

Up to date?

How long?

Available?

How many?

Specs?

Service + Business



# ProvGun/ProvCannon

## Activity

All


Comments

History


Activity

Transitions



✓  The provisioning cannon added a comment - 2014-02-10 17:26

Hej, I'm the provisioning cannon. I'm about to attempt to provision 3 metadata machines in sto3 for you. Hold tight!  
Please check [https://wiki.spotify.net/wiki/Provcannon\\_faq](https://wiki.spotify.net/wiki/Provcannon_faq) if you have any questions.

✓  The provisioning cannon added a comment - 2014-02-10 18:02

Kablam! I provisioned 2 of 3 remaining machines:

sto3-metadata-a2 (██████)

sto3-metadata-a1 (██████)

They should be ready in a few minutes. Enjoy!

# Feature Developer Concerns



How to track?

How to get?

How to talk to it?

Maintenance?

What tools on it?

What to put on it?

Where?

Up to date?

How long?

Available?

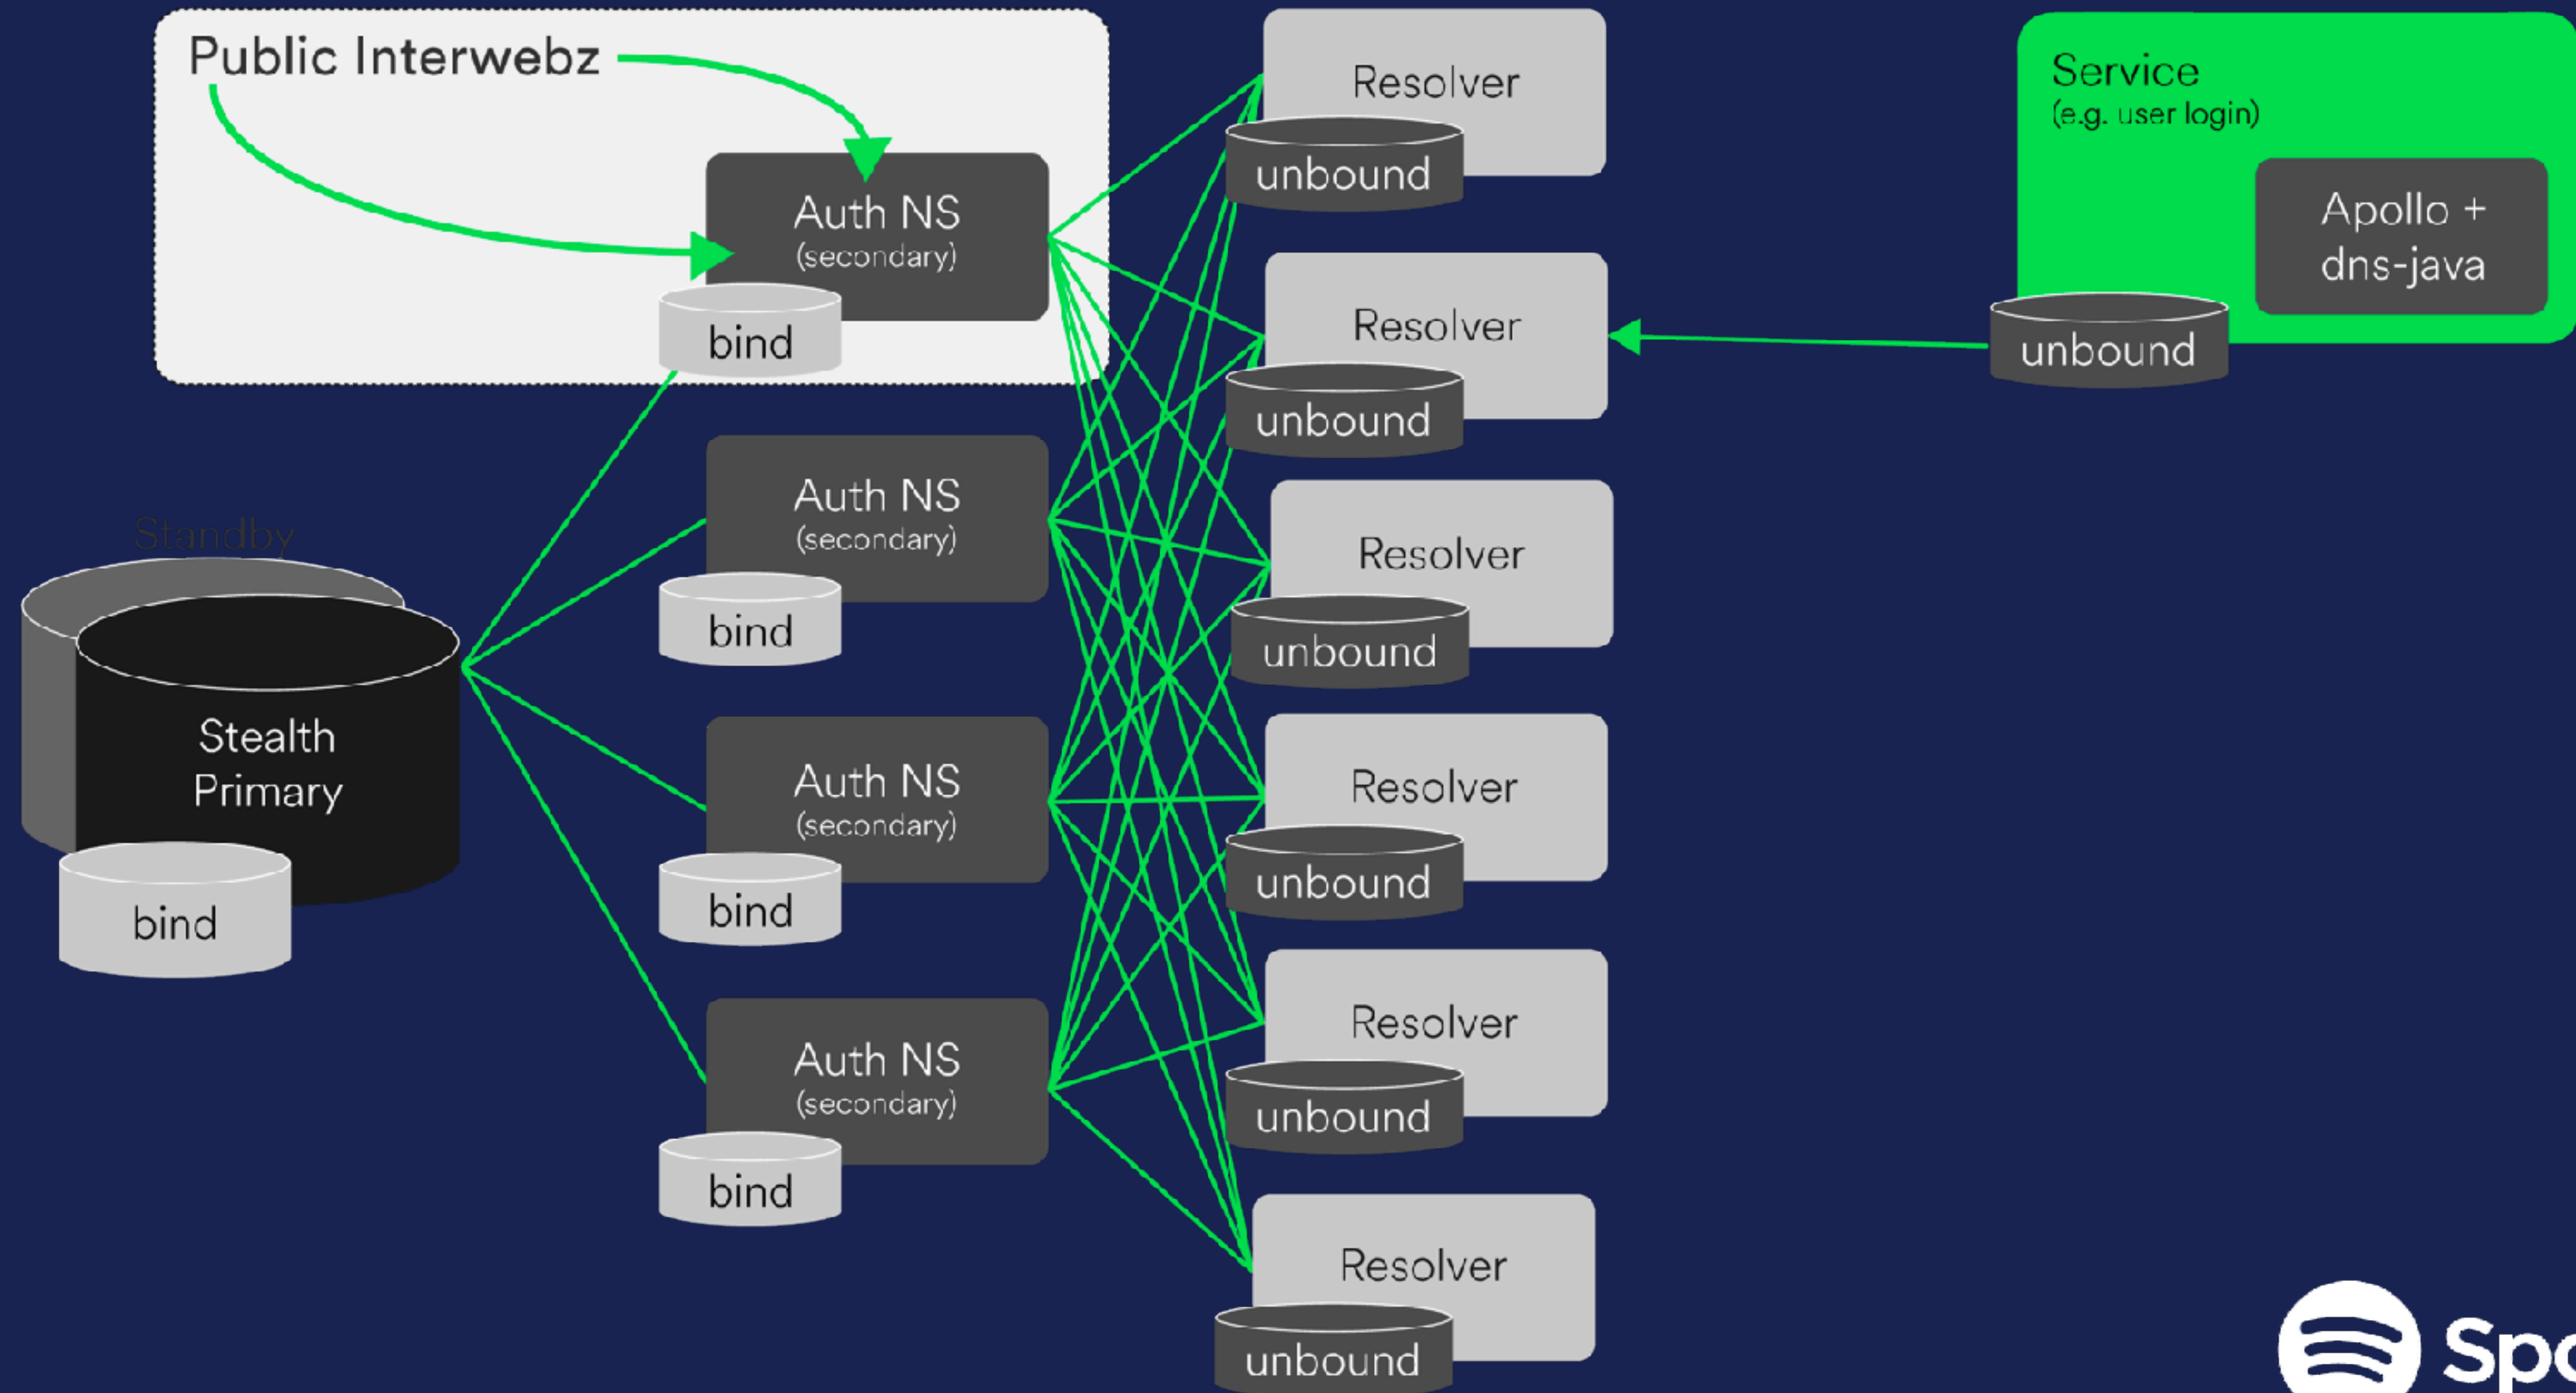
How many?

Specs?

Service + Business



# DNS



# Feature Developer Concerns



How to track?

How to get?

How to talk to it?

Maintenance?

What tools on it?

What to put on it?

Where?

Up to date?

How long?

Available?

How many?

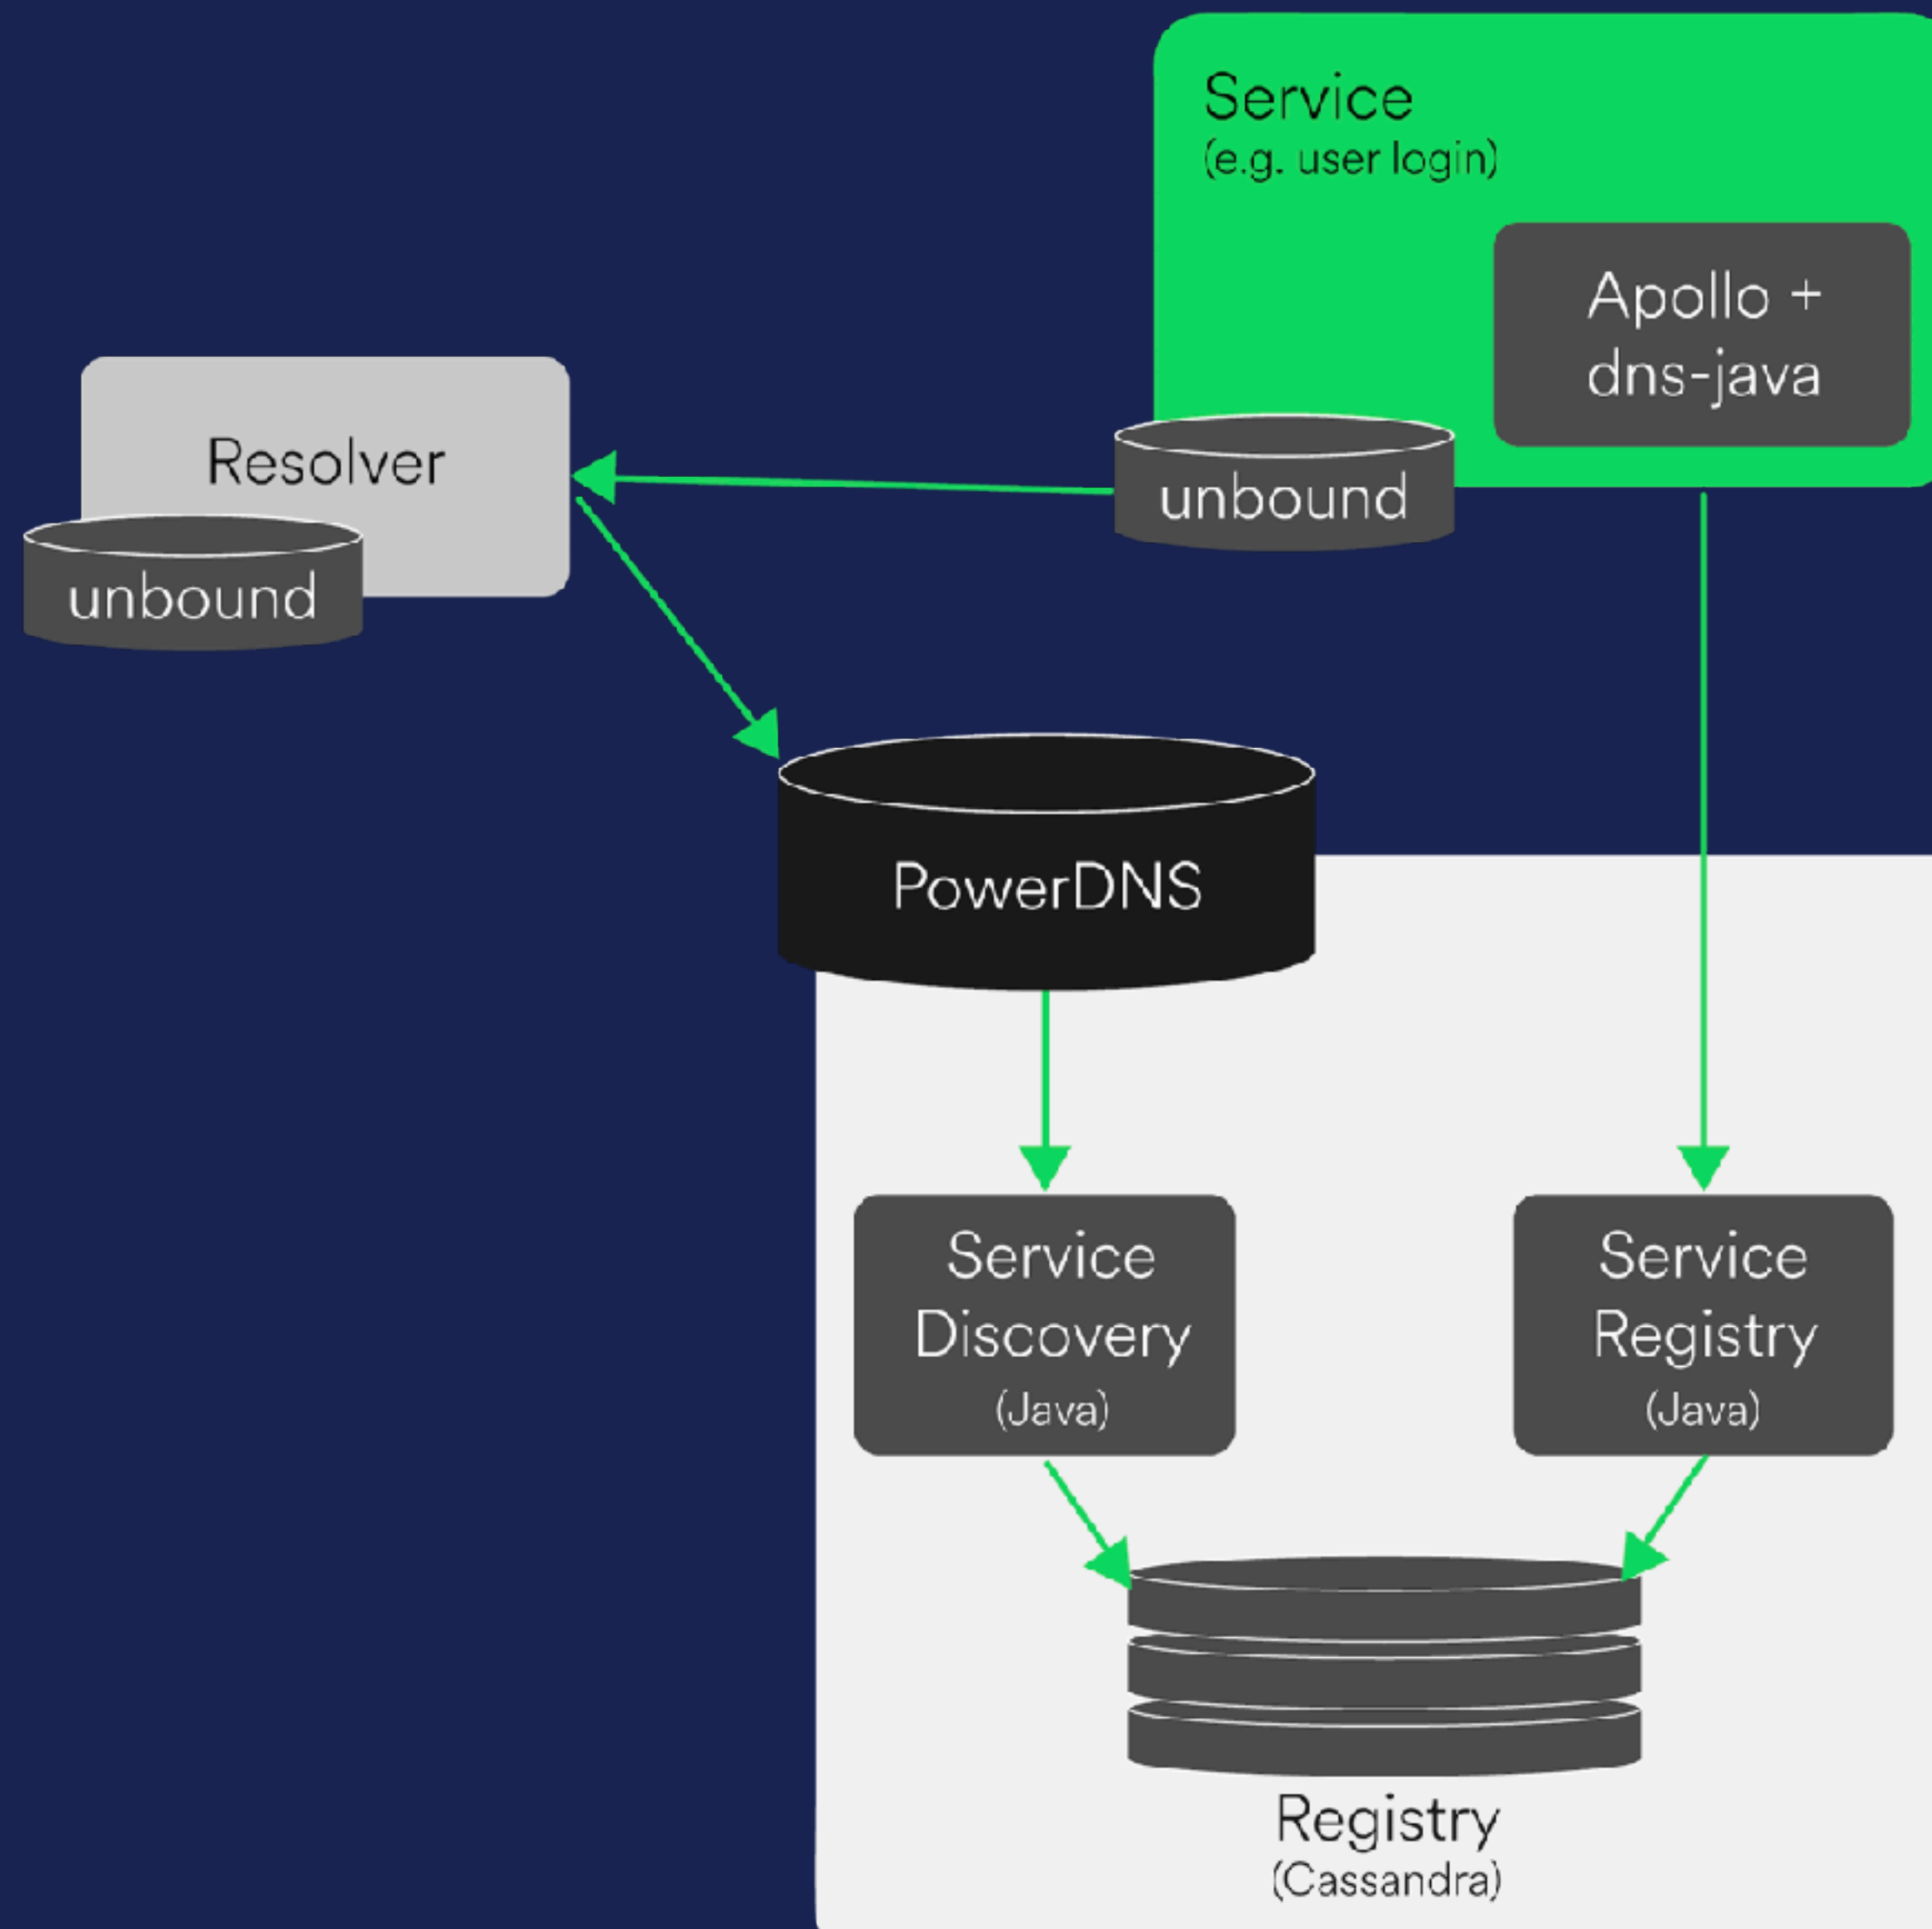
Specs?

Service + Business





# Nameless



# Feature Developer Concerns



How to track?

How to get?

How to talk to it?

Maintenance?

What tools on it?

What to put on it?

Where?

Up to date?

How long?

Available?

How many?

Specs?

Service + Business



# Cortana

## Machines

Filter by Hostname or Pod

Only display machines from jameswen's groups

Found 133 servers



Page 1/3



NAME

POD

ash2-alfjenkins-a1

ash2

ash2-alfjenkins-a2

ash2

ash2-alfzookeeper-a1

ash2

ash2-alfzookeeper-a2

ash2

# Cortana

Request hardware

Hardware Type <sup>?</sup> **HIGHCPU/32GB** HIGHCPU/64GB DATABASE HIGHIOPS STORAGE

POD <sup>?</sup> ASH **ASH2**  
LON **ANY** LON2 LON3 **LON6**  
SJC **SJC1**  
STO **STO3**

Pool <sup>?</sup> n

Operating System <sup>?</sup> **UBUNTU TRUSTY**

Role <sup>?</sup> awesomeservice

Justification <sup>?</sup> Machines are cool.

RAID Setup <sup>?</sup> **DEFAULT** JBOD RAID 5 RAID 6 RAID 10 TRISTRIPE MIRROR PRODSTORAGE

Use this 'default' option unless you know what you're doing.

Quantity <sup>?</sup> 50

Rack Diversity <sup>?</sup> 5

The maximum machines of this role per rack, including existing machines.

**⚠ Hold your horses!** There is low to medium availability for HighCPU/HighMemory in lon6. Please consider using a different location.

Chill out, I know what I'm doing

**SUBMIT REQUEST**

## Provisioning Request Details

Info:	
ID	f650a19a-daa3-4273-a9f6-2a7446bbfb91
Status	fulfilled
Hardware Type	HighCPU/HighMemory:201406
Pod	sto3
Pool	a
Operating System	trusty
Role	welcomeimages
Quantity	3
Rack Diversity	1
Requester	cpm
Justification	Automatically created by Cortana Pool Manager
Machines	<a href="#">View associated machines</a>

### Created Jobs:

- Job #1** finished (install)
- Job #2** finished (install)
- Job #3** finished (install)

# Feature Developer Concerns



How to track?

How to get?

How to talk to it?

Maintenance?

What tools on it?

What to put on it?

Where?

Up to date?

How long?

Available?

How many?

Specs?

Service + Business



# Helios and Containers



docker

# Feature Developer Concerns



How to track?

How to get?

How to talk to it?

Maintenance?

What tools on it?

What to put on it?

Where?

Up to date?

How long?

Available?

How many?

Specs?

Service + Business



# Google Compute Platform





**ash2-cortana-a1.ash2**



Zone



Service



Group

Sequential #


**gew1-cortana-a-l33t.gew1**



Zone



Service



Pool



Random 4 Chars

# Cortana Pool Manager

POOLS						GROUP BY	ROLE ▾	
PLATFORM ▾	POD ▾	POOL ▾	ACTUAL ▾	DESIRED ▾	STATUS ▾			
▼ metadata								
gcp ( <i>xpn-metadata-1</i> )	gew1	d	23	23	STABLE			
gcp ( <i>xpn-metadata-1</i> )	gae2	d	14	14	STABLE			
gcp ( <i>xpn-metadata-1</i> )	guc3	d	30	30	STABLE			
physical	lon6	d	32	32	STABLE			
physical	ash2	d	39	39	STABLE			
physical	sjc1	d	24	24	STABLE			
physical	sto3	d	23	23	STABLE			

# Feature Developer Concerns



How to track?

How to get?

How to talk to it?

Maintenance?

What tools on it?

What to put on it?

Where?

Up to date?

How long?

Available?




How many?

Specs?

Service + Business



# Regional Managed Instance Groups

<i>gcp (xpn-provisioning-1)</i>	gew1 ⓘ	This pool's instances will be deployed across any/all of europe-west1's zones in GCP	4	STABLE			
---------------------------------	--------	--	---	--------	---	---	---

# Feature Developer Concerns



How to track?

How to get?

How to talk to it?

Maintenance?

What tools on it?

What to put on it?

Where?

Up to date?

How long?

Available?

How many?

Specs?

Service + Business



# MBMI: Minimal Base Machine Image



# Feature Developer Concerns



How to track?

How to get?

How to talk to it?

Maintenance?

What tools on it?

What to put on it?

Where?

Up to date?

How long?

Available?


How many?

Specs?

Service + Business



# Phoenix





# Feature Developer Concerns



How to track?

How to get?

How to talk to it?

Maintenance?

What tools on it?

What to put on it?

Where?

Up to date?

How long?

Available?

How many?

Specs?

Service + Business



# Current: Feature Developer's Context for Service's Capacity

Stockholm

Pool:  
4 instances x (High Mem)

GCP - europe-west-1

Pool:  
2 instances x (n1-standard-32)

# Feature Developer Concerns



How to track?

How to get?

How to talk to it?

Maintenance?

What tools on it?

What to put on it?

Where?

Up to date?

How long?

Available?

How many?

Specs?

Service + Business



# Future



# Gordon (Cloud DNS)



# Feature Developer Concerns



How to track?

How to get?

How to talk to it?

Maintenance?

What tools on it?

What to put on it?

Where?

Up to date?

How long?

Available?

How many?

Specs?

Service + Business



# Autoscaling



# Feature Developer Concerns



How to track?

How to get?

How to talk to it?

Maintenance?

What tools on it?

What to put on it?

Where?

Up to date?

How long?

Available?

How many?

Specs?

Service + Business





# Right Sizing

## Resize instance

This instance has had low CPU and memory utilization recently. You can save an estimated \$17 per month by switching to the machine type: g1-small (1 vCPU, 1.7 GB memory). [Learn more](#)


### Current machine type

n1-standard-1 (1 vCPU, 3.75 GB memory)

### New machine type

g1-small (1 vCPU, 1.7 GB memory) Recommended

[Customize](#)

 This instance is running. Changing the machine type will restart the VM.

[Apply](#)

[Dismiss recommendation](#)

[Cancel](#)

# Feature Developer Concerns



How to track?

How to get?

How to talk to it?

Maintenance?

What tools on it?

What to put on it?

Where?

Up to date?

How long?

Available?

How many?

Specs?

Service + Business



# Future Feature Developer's Context for Service's Capacity

GCP - asia-east-1

Service Pool

GCP - europe-west-1

Service Pool

GCP - us-central-1

Service Pool

# Feature Developer Concerns



How to track?

How to get?

How to talk to it?

Maintenance?

What tools on it?

What to put on it?

Where?

Up to date?

How long?

Available?

How many?

Specs?

Service + Business



# Learnings

# Why Pets to Cattle was Difficult:

- Manual/tedious setup
- Wait times for machine becoming ready (packages, DNS)
- Non-automatic security updates
- A fixed, reliable hostname
- SSH Access
- Always up/present unless team tears down

# Ephemerality Learnings

- **Monitoring**
- **Logging**
- **Service Design**
- **Incidents**

# Hybrid Learnings

- Replicate bare metal functionality, then iterate
- When in doubt, devs provision up and many
- Migration = great time to influence dev paradigms
- Don't need to DIY



# DevEx Learnings

- Feature devs need carrots, sledgehammers, and/or limos to change
- Edge Cases: REST API + CLI = provide enough for feature teams to handle the edge cases

# Recap

- Decrease necessary infrastructure context
- Increase reliability
- Save \$\$\$
- Increase dev happiness and productivity

**Let's strategically  
control and limit how  
feature developers  
think about  
infrastructure.**

**James Wen**

**Email: [jameswen@spotify.com](mailto:jameswen@spotify.com)**

**Twitter/Github: [@rochesterinnyc](https://twitter.com/rochesterinnyc)**

**LinkedIn: [jamesrwen](https://www.linkedin.com/in/jamesrwen)**

**Spotify is hiring! [spotifyjobs.com](https://spotifyjobs.com)**



IO Tribe

