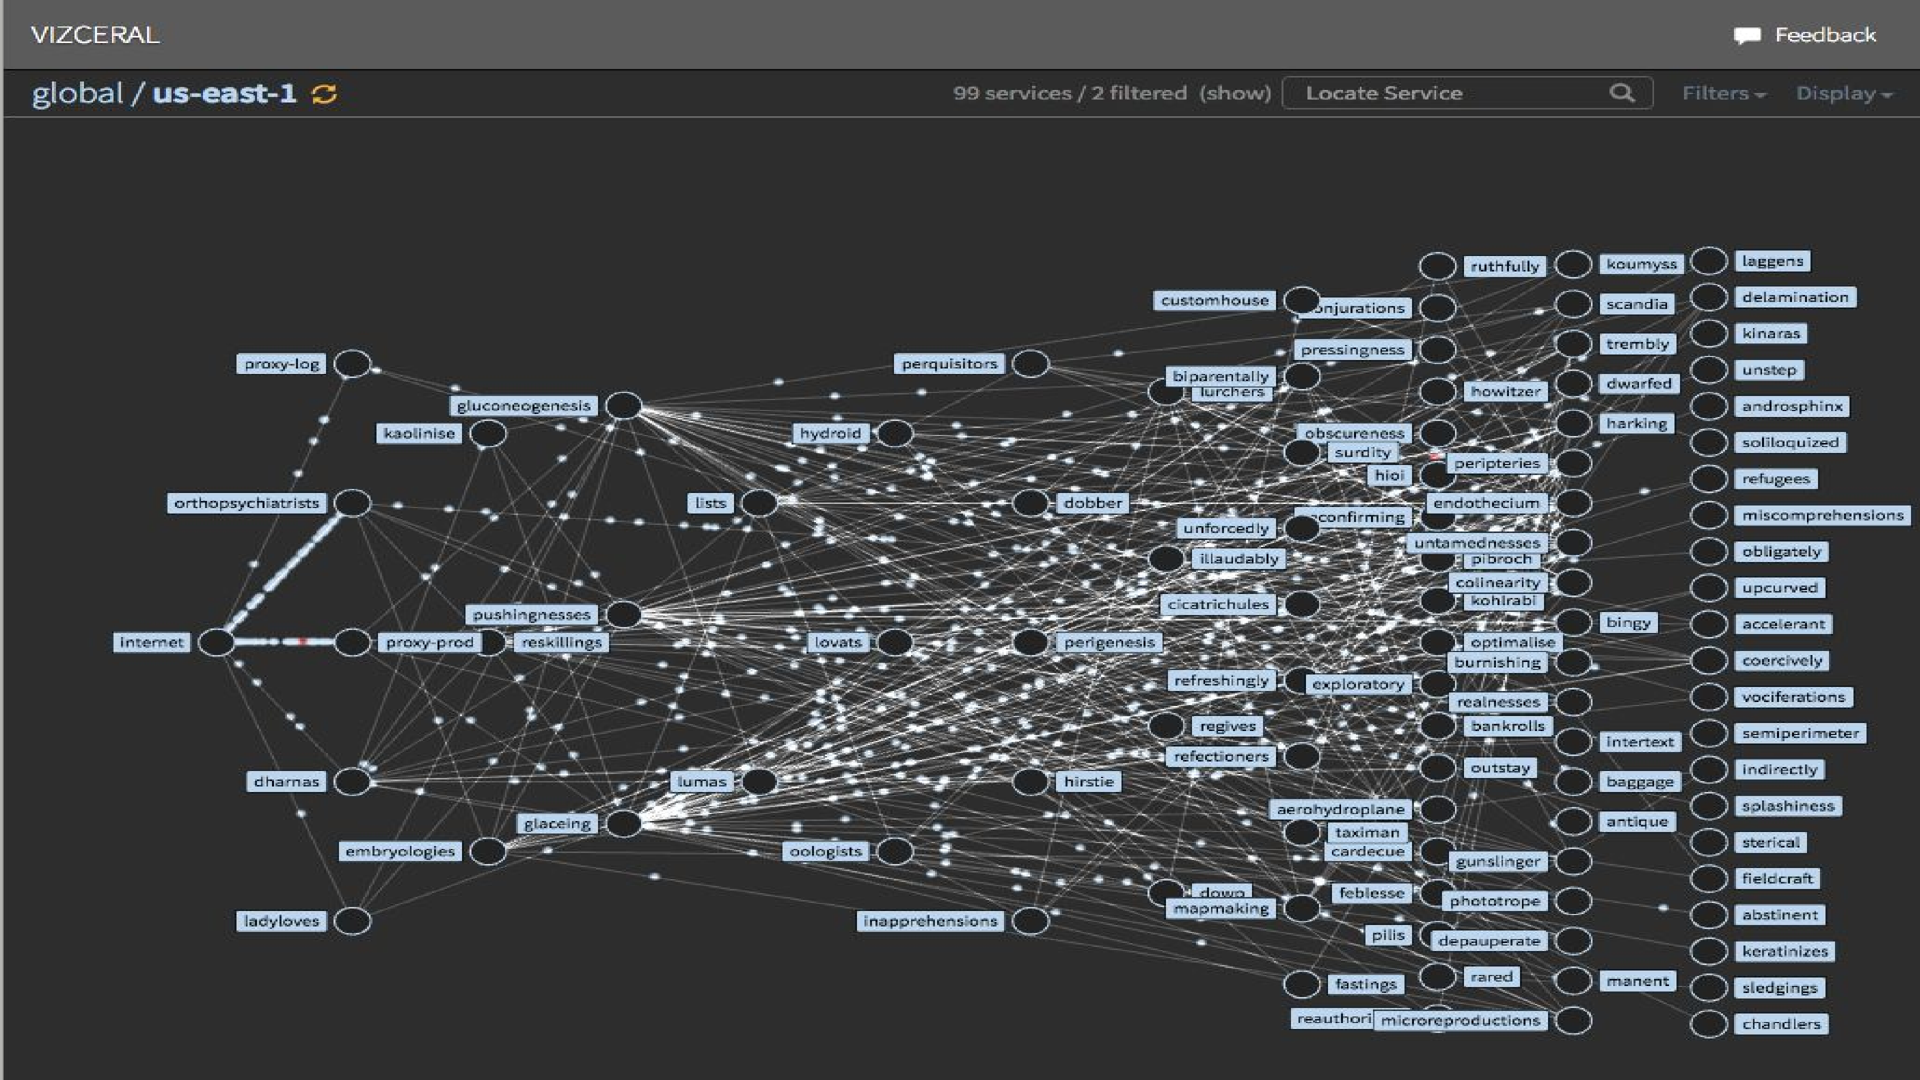
VIZCERAL
Feedback
global / us-east-1
99 services / 2 filtered (show)
Locate Service
Filters
Display
internet
proxy-log
orthopsychiatrists
proxy-prod
dharnas
ladyloves
kaolinise
gluconeogenesis
pushingnesses
reskillings
glaceing
embryologies
hydroid
lists
lovats
lumas
oologists
perquisitors
dobber
perigenesis
hirstie
inapprehensions
customhouse
biparentally
lurchers
unforcedly
illaudably
cicatrichules
refreshingly
regives
refectioners
dowp
mapmaking
onjurations
pressingness
obscureness
surdity
hioi
econfirming
untamednesses
exploratory
aerohydroplane
taximan
cardecue
feblesse
pilis
fastings
reauthori
ruthfully
howitzer
peripteries
endothecium
pibroch
colinearity
kohlrabi
optimalise
burnishing
realnesses
bankrolls
outstay
gunslinger
phototrope
depauperate
rared
microreproductions
koumyss
scandia
trembly
dwarfed
harking
bingy
intertext
baggage
antique
manent
laggens
delamination
kinaras
unstep
androsphinx
soliloquized
refugees
miscomprehensions
obligately
upcurved
accelerant
coercively
vociferations
semiperimeter
indirectly
splashiness
sterical
fieldcraft
abstinent
keratinizes
sledgings
chandlers

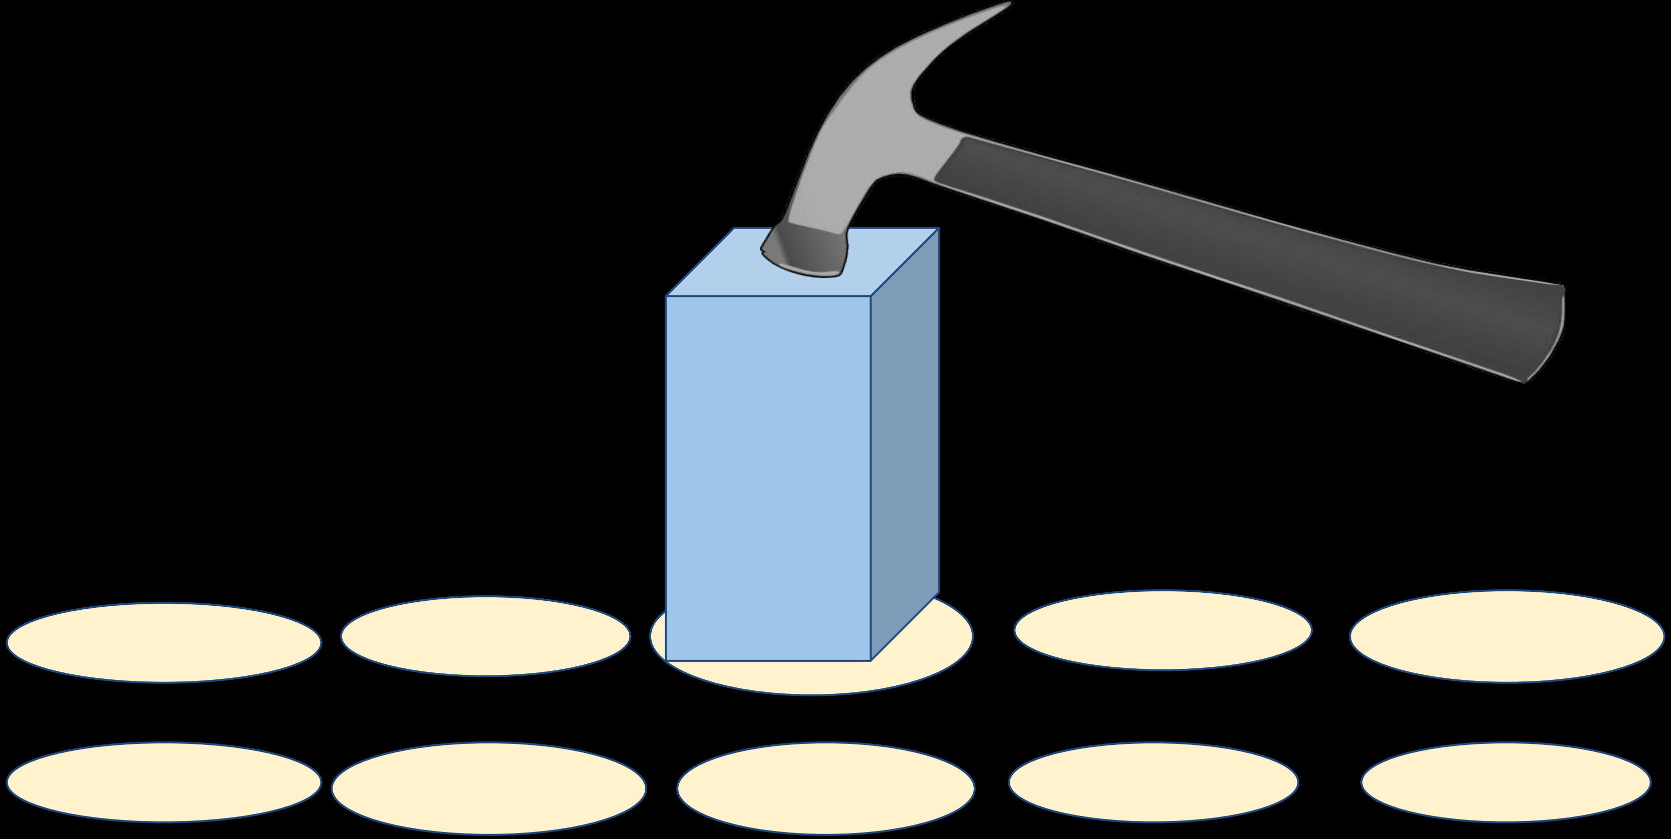

Polyglot Persistence Powering Microservices

Roopa Tangirala
Engineering Manager
Netflix

Agenda

- 5 Use Cases
- Challenges
- Current Approach
- Takeaway

About Netflix

Netflix has been leading the way for digital content since 1997

NETFLIX
NETFLIX
STRANGER THINGS
NETFLIX
NARCOS
NETFLIX
MARVEL
LUKE CAGE

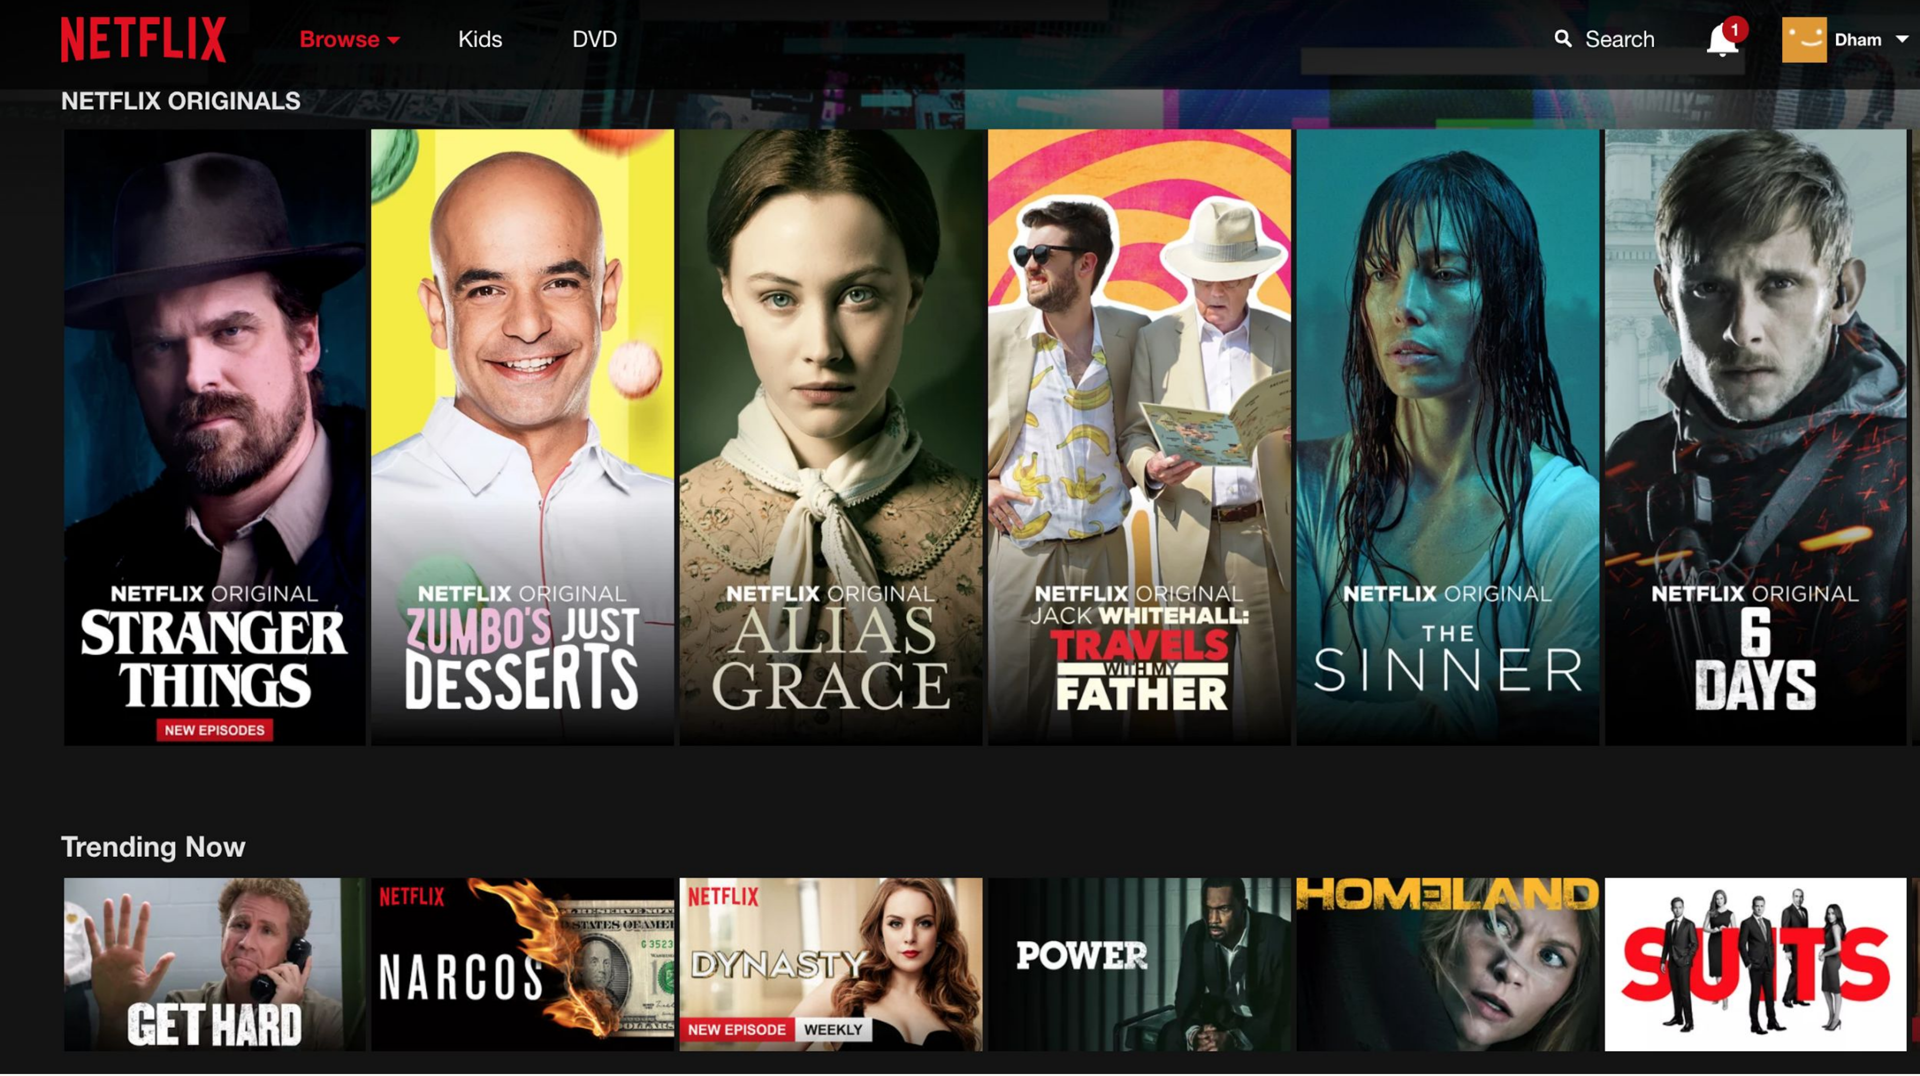
NETFLIX
Browse
Kids
DVD
Search
1
Dham
NETFLIX ORIGINALS
NETFLIX ORIGINAL
STRANGER THINGS
NEW EPISODES
NETFLIX ORIGINAL
ZUMBO'S JUST DESSERTS
NETFLIX ORIGINAL
ALIAS GRACE
NETFLIX ORIGINAL
JACK WHITEHALL: TRAVELS WITH MY FATHER
NETFLIX ORIGINAL
THE SINNER
NETFLIX ORIGINAL
6 DAYS
Trending Now
GET HARD
NETFLIX
NARCOS
NETFLIX
DYNASTY
NEW EPISODE
WEEKLY
POWER
HOMELAND
SUITS

elastic
DYNOMITE
EVCache
AWS S3
cassandra
RDS
ZOO
Apache
Zookeeper
MySQL®

CDE

elastic	Search, Analyze and visualize in near real time
EVCache	Distributed in-memory caching solution based on memcached
cassandra	Distributed NOSQL database to handle large datasets providing high availability.
DYNOMITE	Distributed dynamo layer for different storage engines and protocols supporting Redis, memcached, RocksDB
	TitanDB is scalable graph database optimized for storing and querying graph datasets.

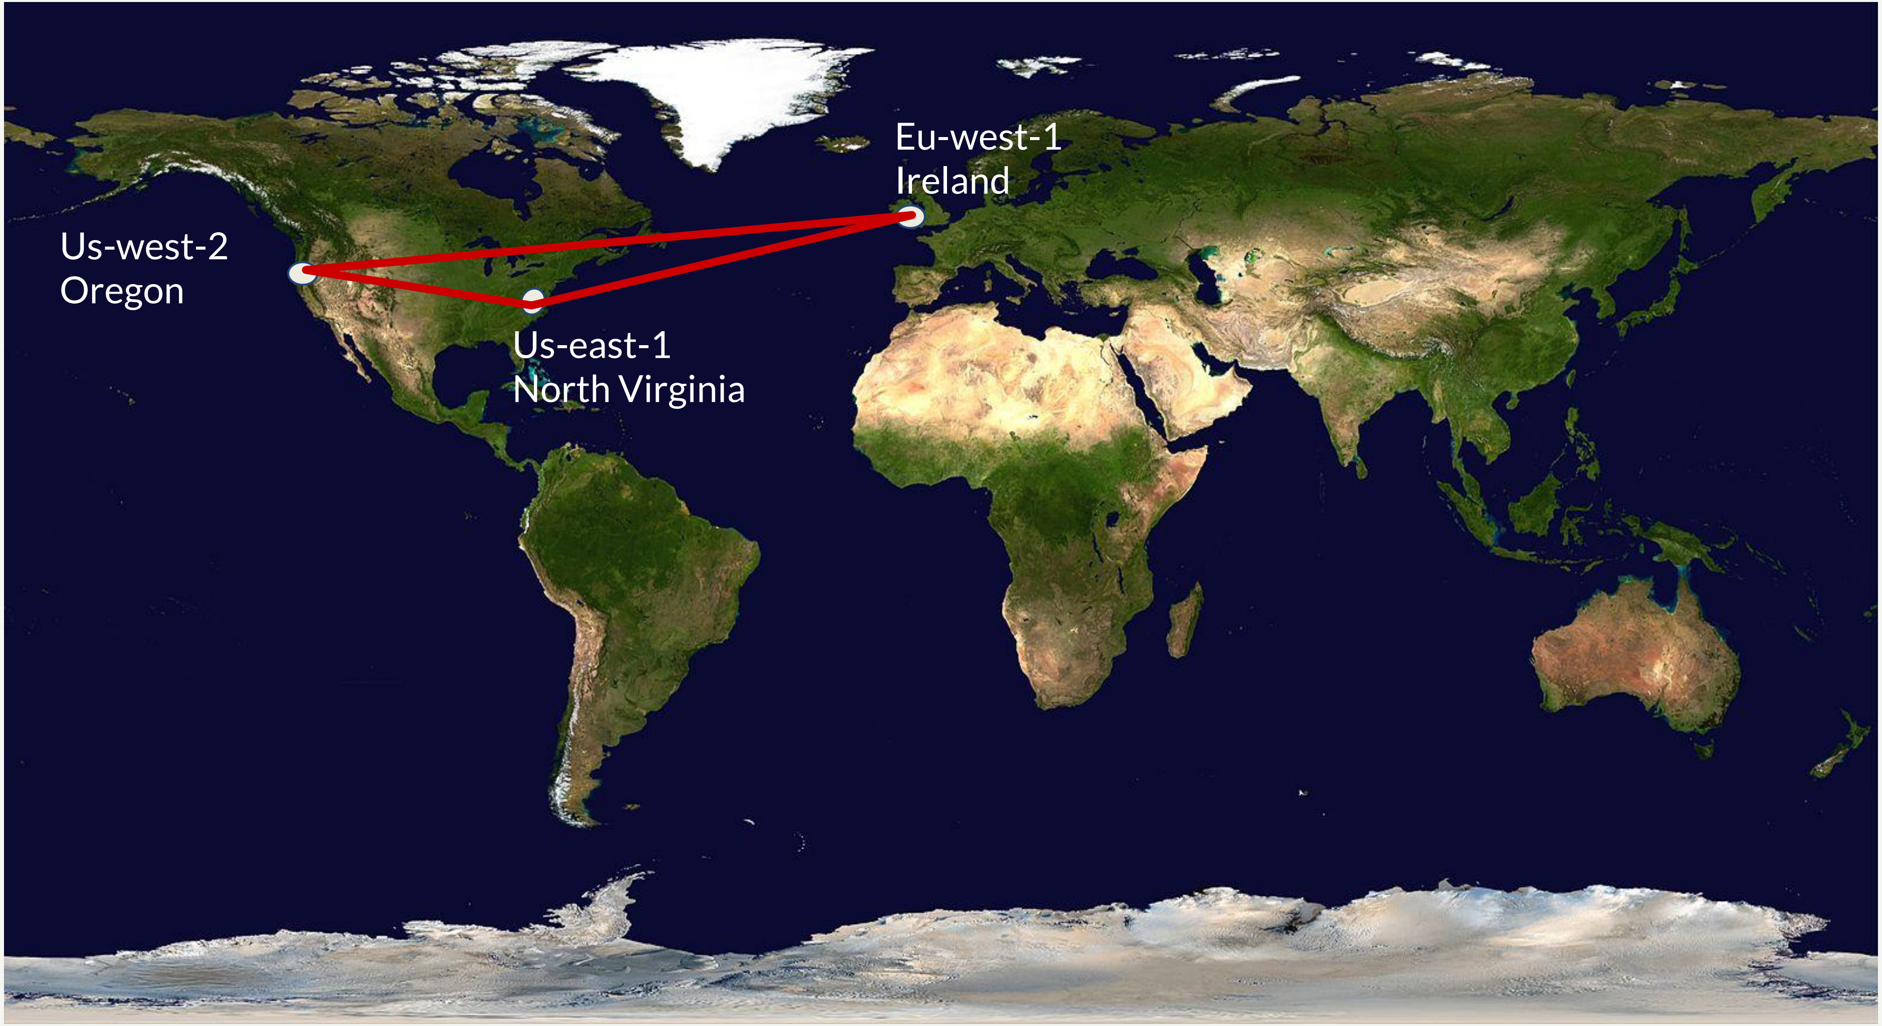
Eu-west-1
Ireland
Us-west-2
Oregon
Us-east-1
North Virginia

★ CDN URL

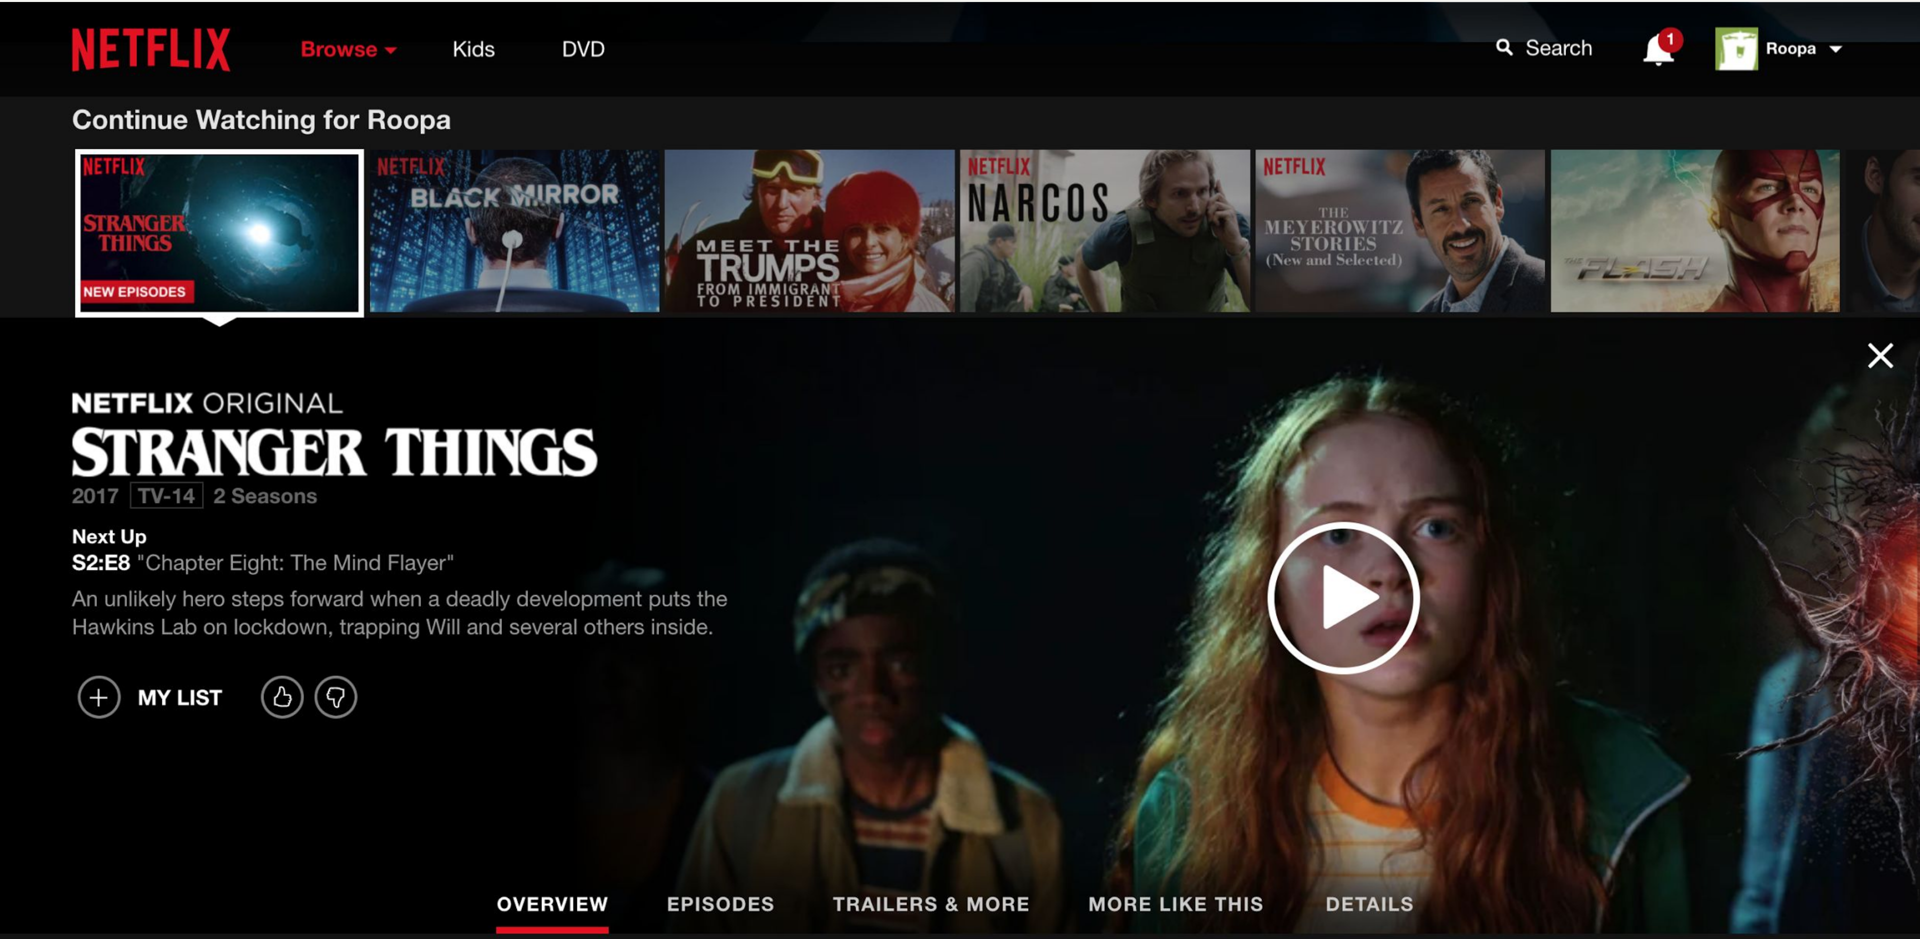

Next Up

S2:E8 "Chapter Eight: The Mind Flayer"

An unlikely hero steps forward when a deadly development puts the Hawkins Lab on lockdown, trapping Will and several others inside.

NETFLIX
47:25
Stranger Things
S2:E8
Chapter Eight: The Mind Flayer

Requirements - CDN URL

- High availability
- Very low latency reads/writes (less than 1ms)
- High Throughput per node

DYNOMITE
elastic
EVCache
cassandra

GUESS?

EVCache

- Distributed In Memory
- Very low Latency responses

★ Playback Error

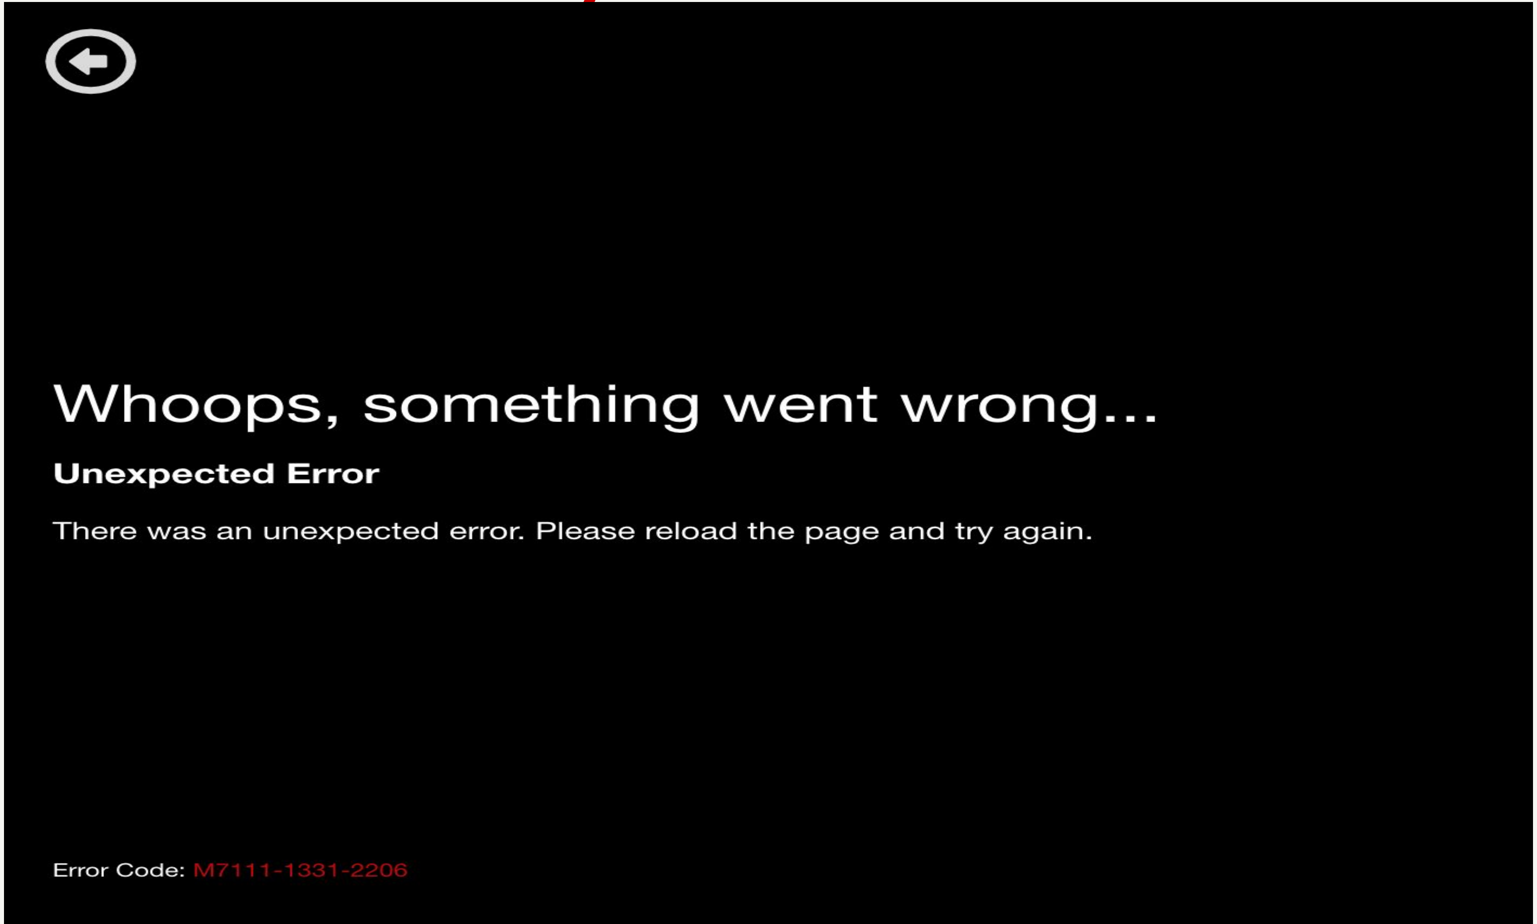

Whoops, something went wrong...

Unexpected Error

There was an unexpected error. Please reload the page and try again.

Error Code: M7111-1331-2206

PLAYBACK CONTEXT
(Tracks + Track Urls)

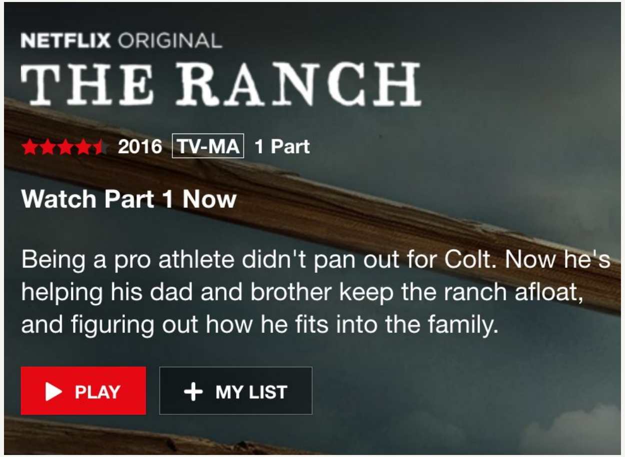

NETFLIX ORIGINAL

THE RANCH

★★★★★ 2016 TV-MA 1 Part

Watch Part 1 Now

Being a pro athlete didn't pan out for Colt. Now he's helping his dad and brother keep the ranch afloat, and figuring out how he fits into the family.

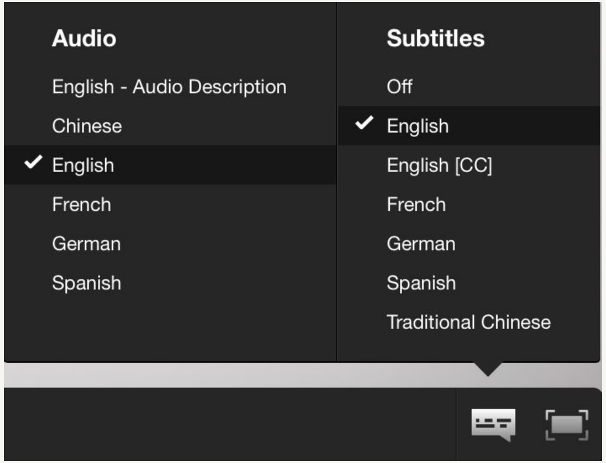

Audio

- English - Audio Description
- Chinese
- ✓ English
- French
- German
- Spanish

Subtitles

- Off
- ✓ English
- English [CC]
- French
- German
- Spanish
- Traditional Chinese

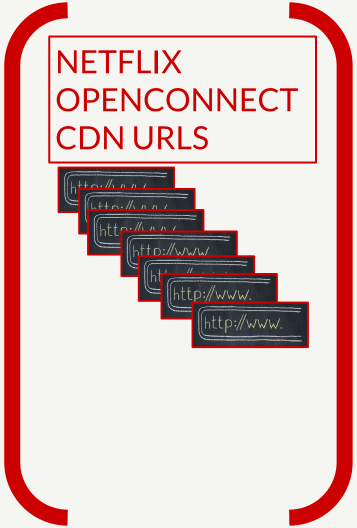

NETFLIX OPENCONNECT CDN URLS

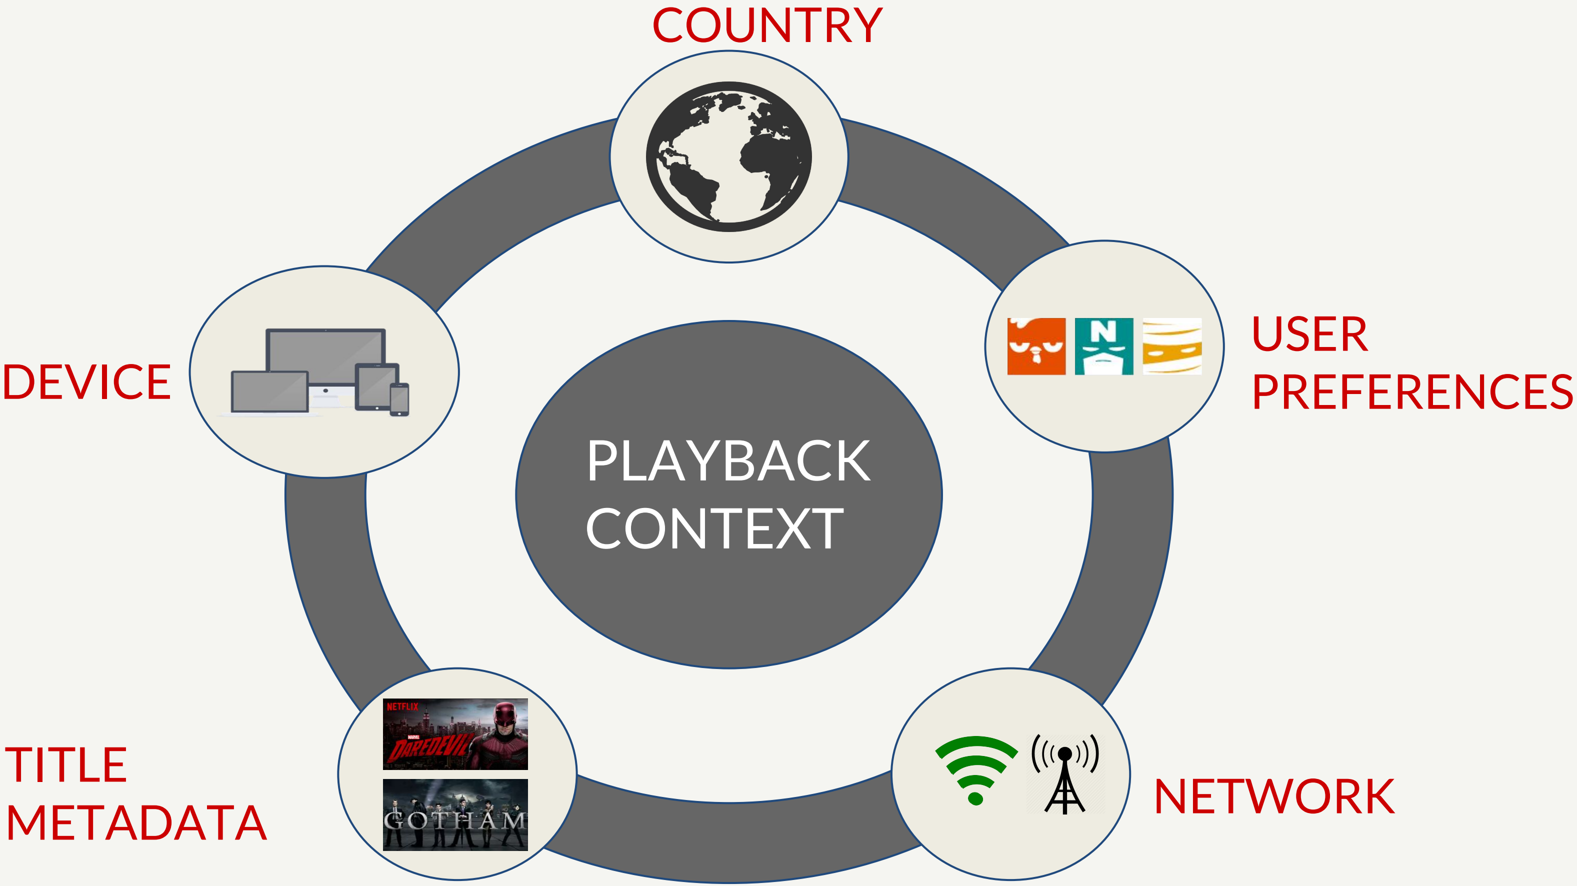

COUNTRY

USER PREFERENCES

NETWORK

TITLE METADATA

DEVICE

PLAYBACK CONTEXT

Requirements - Playback Error

- Quick Incident Resolution
- Interactive Dashboards
- Near realtime Search
- Ad Hoc Queries

DYNOMITE
elastic
EVCache
cassandra

GUESS?

elastic

⇨ Powerful Search & Analytics

⇨ Interactive Dashboards

Interactive Exploration

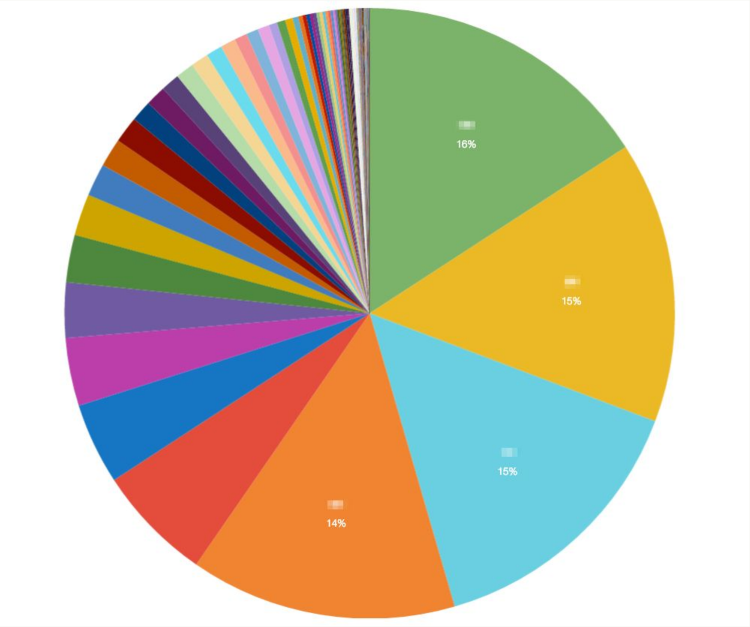

Top N queries

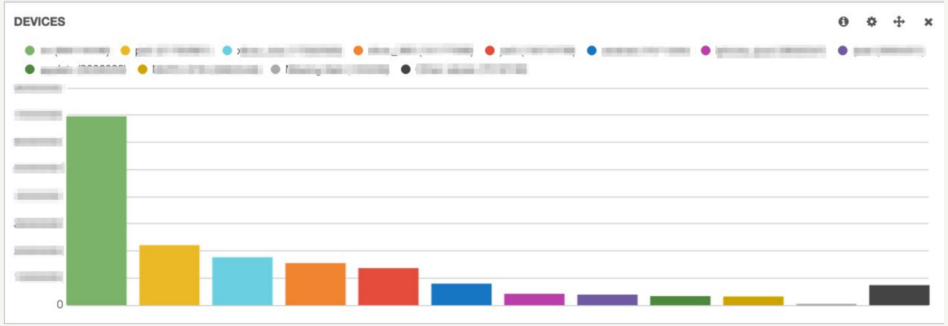

Incident To Resolution Time

2+ Hours → Under 10 Minutes

★ Viewing History

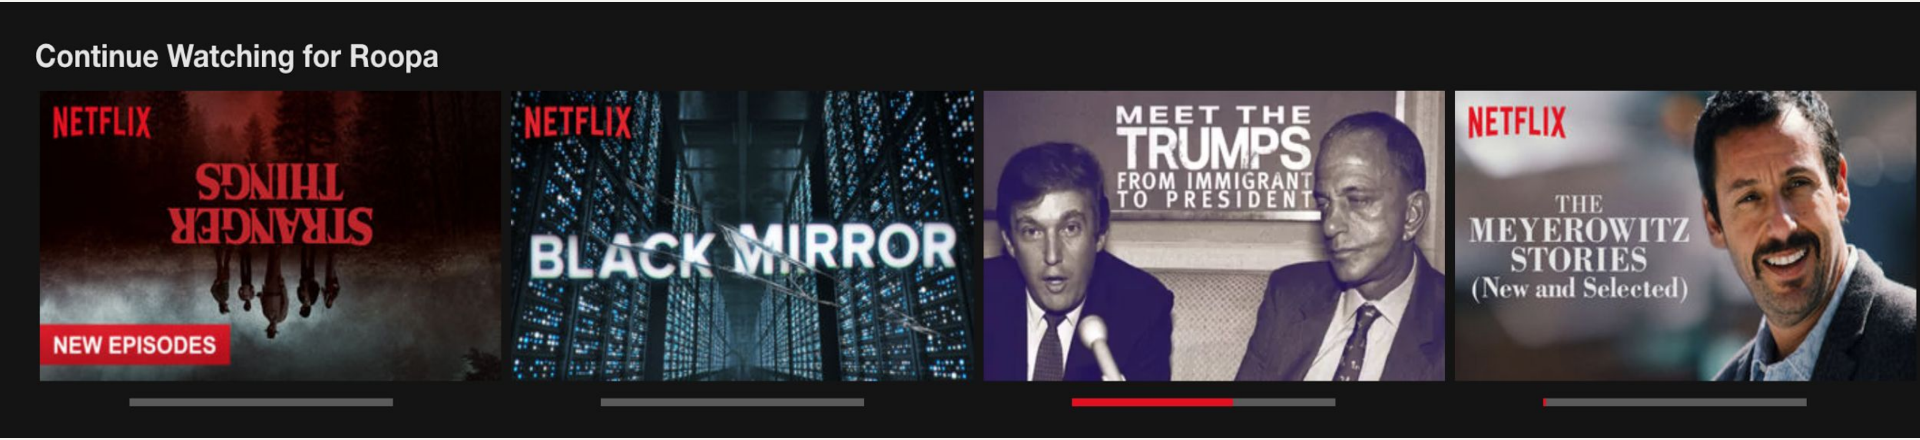
Continue Watching for Roopa
NETFLIX
STRANGER THINGS
NEW EPISODES
NETFLIX
BLACK MIRROR
MEET THE TRUMPS
FROM IMMIGRANT TO PRESIDENT
NETFLIX
THE MEYEROWITZ STORIES (New and Selected)

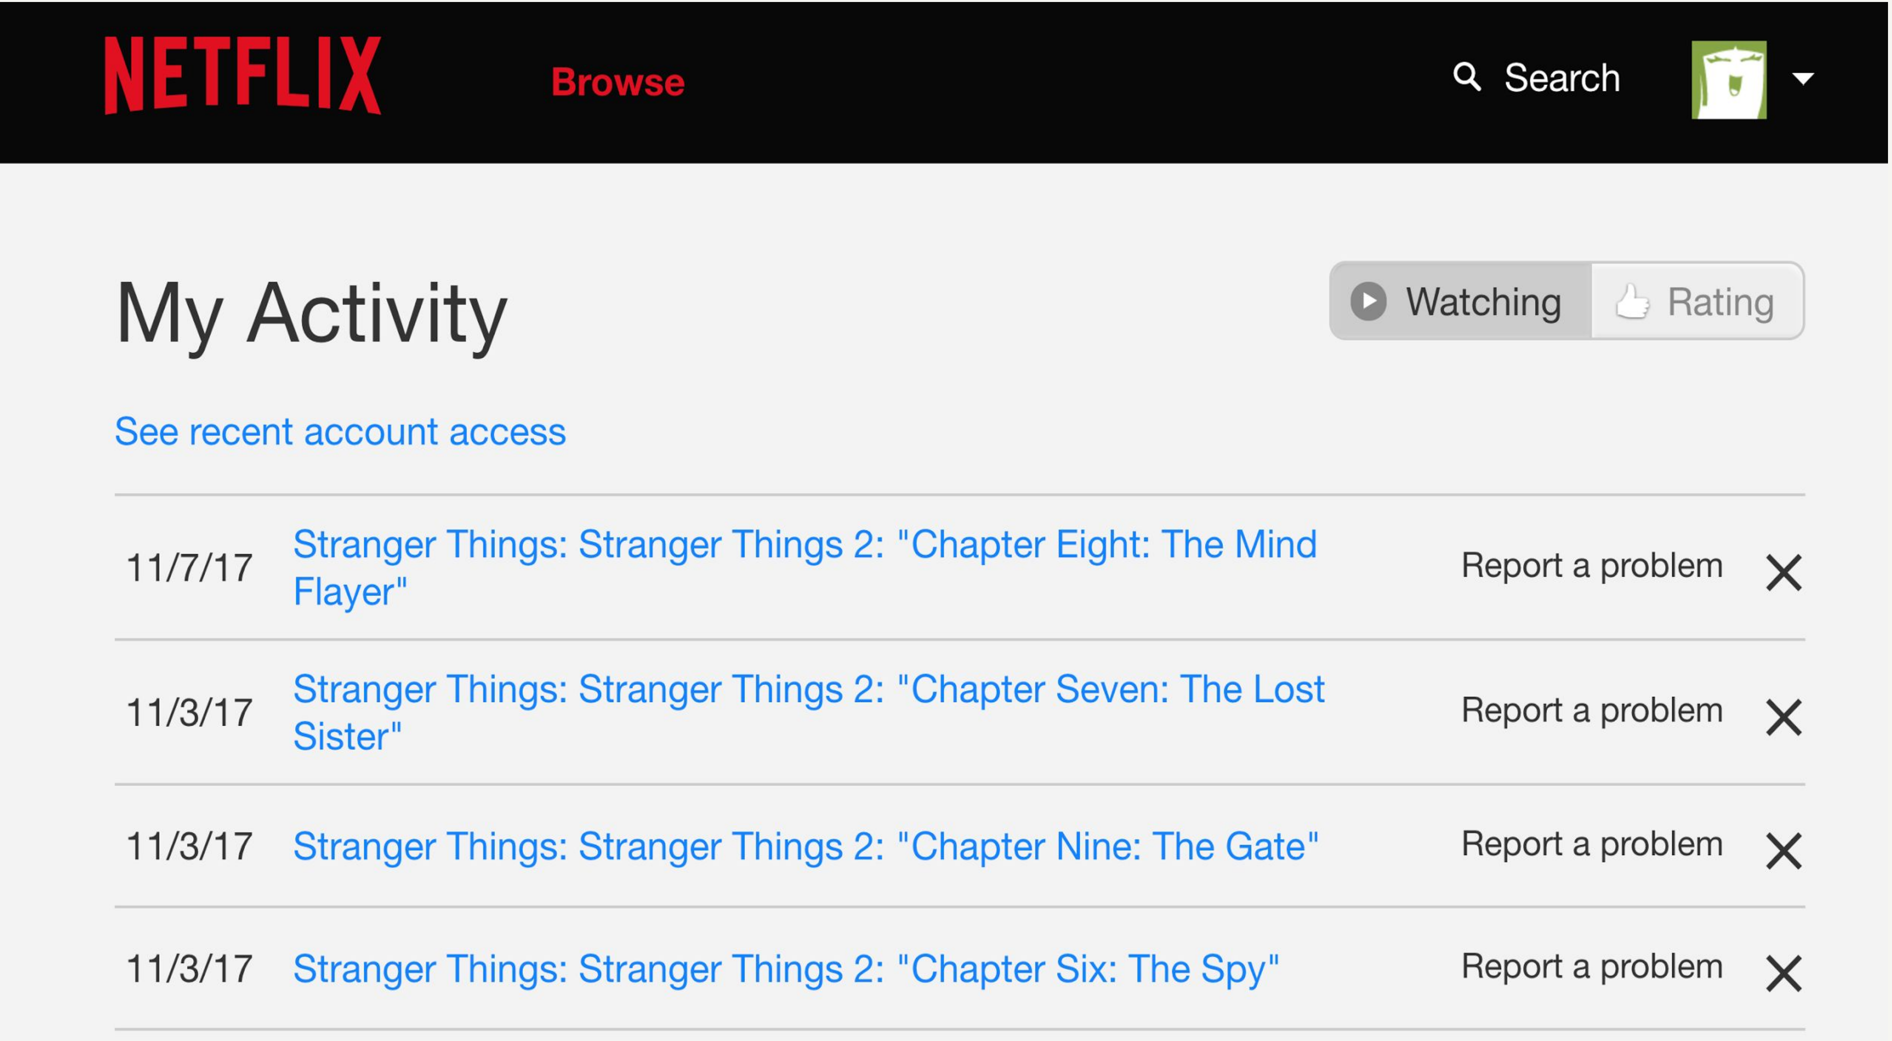

My Activity

See recent account access

11/7/17	Stranger Things: Stranger Things 2: "Chapter Eight: The Mind Flayer"	Report a problem	✕
11/3/17	Stranger Things: Stranger Things 2: "Chapter Seven: The Lost Sister"	Report a problem	✕
11/3/17	Stranger Things: Stranger Things 2: "Chapter Nine: The Gate"	Report a problem	✕
11/3/17	Stranger Things: Stranger Things 2: "Chapter Six: The Spy"	Report a problem	✕

Requirements - Viewing History

- Time series dataset
- Support high writes
- Cross region replication
- Large dataset

Growth of Viewing History

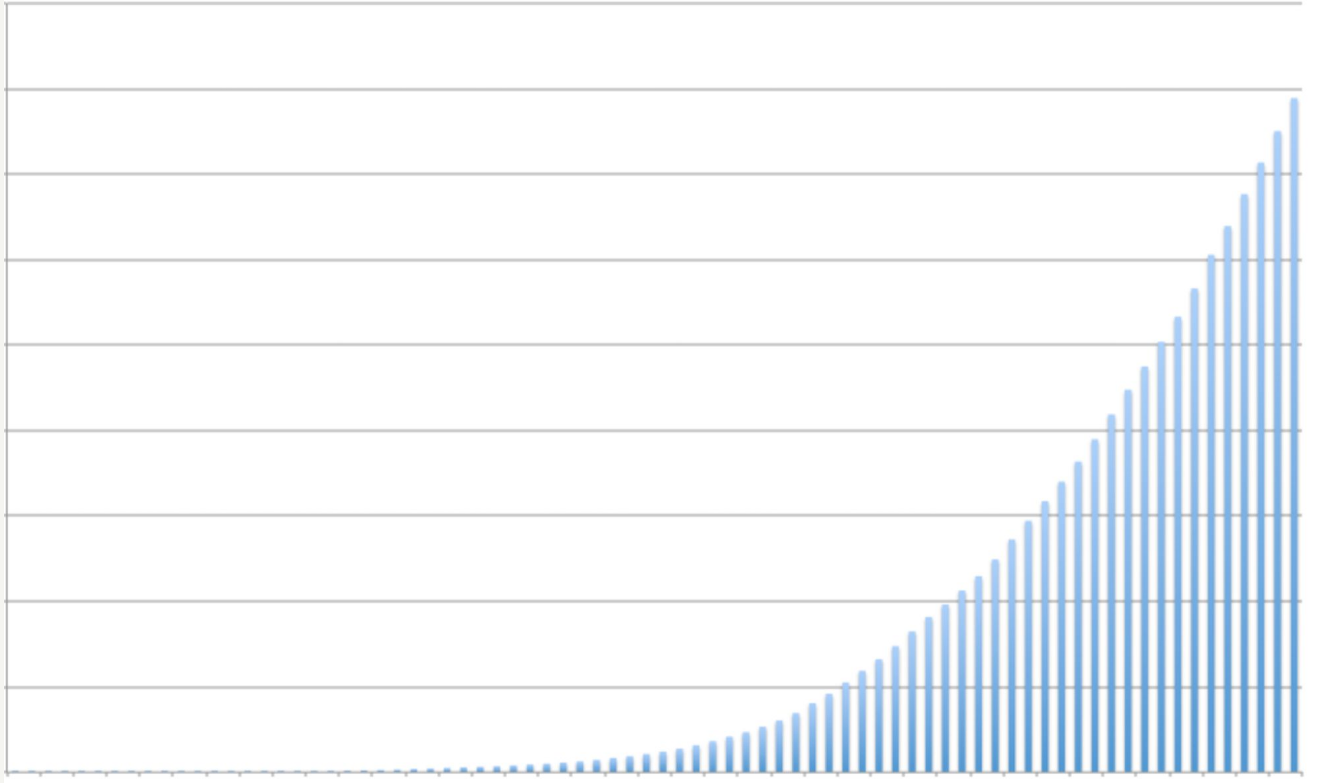

elastic
DYNOMITE
EVCache
cassandra

GUESS?

⇨ Multi-datacenter, multi-directional replication

⇨ Highly availability and scalability

Data Model

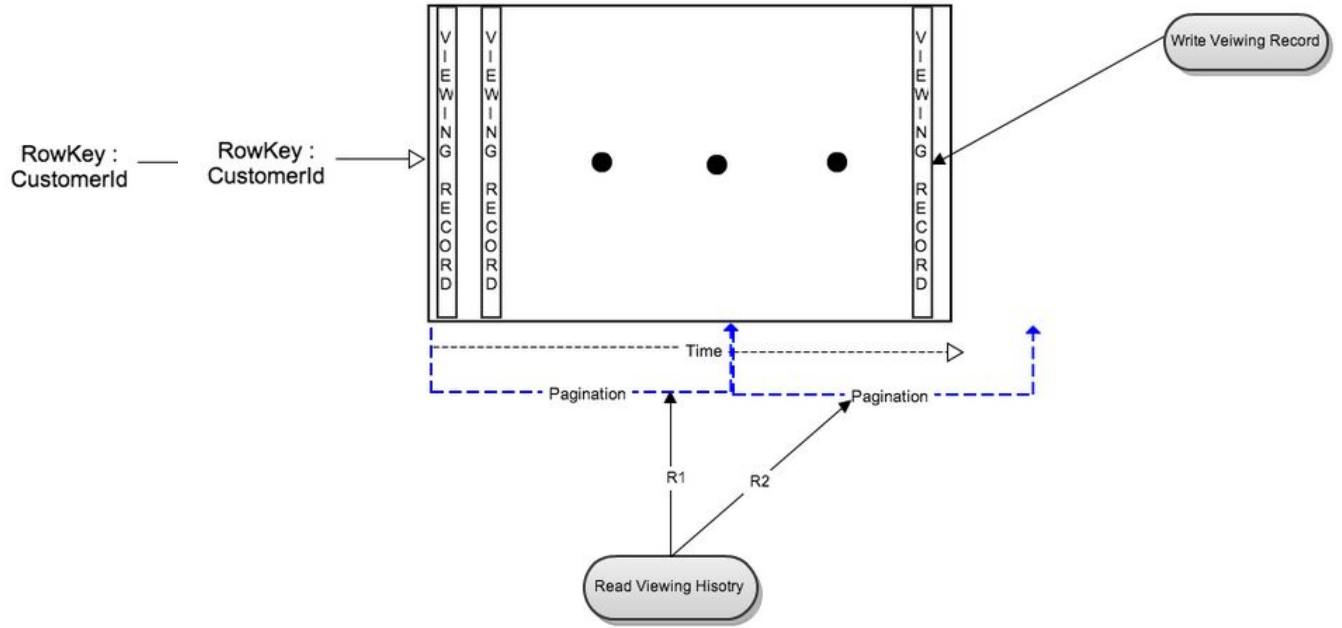
RowKey : CustomerId
RowKey : CustomerId
VIEWING RECORD
VIEWING RECORD
VIEWING RECORD
Write Veiwing Record
Time
Pagination
Pagination
R1
R2
Read Viewing Hisotry

New Data Model

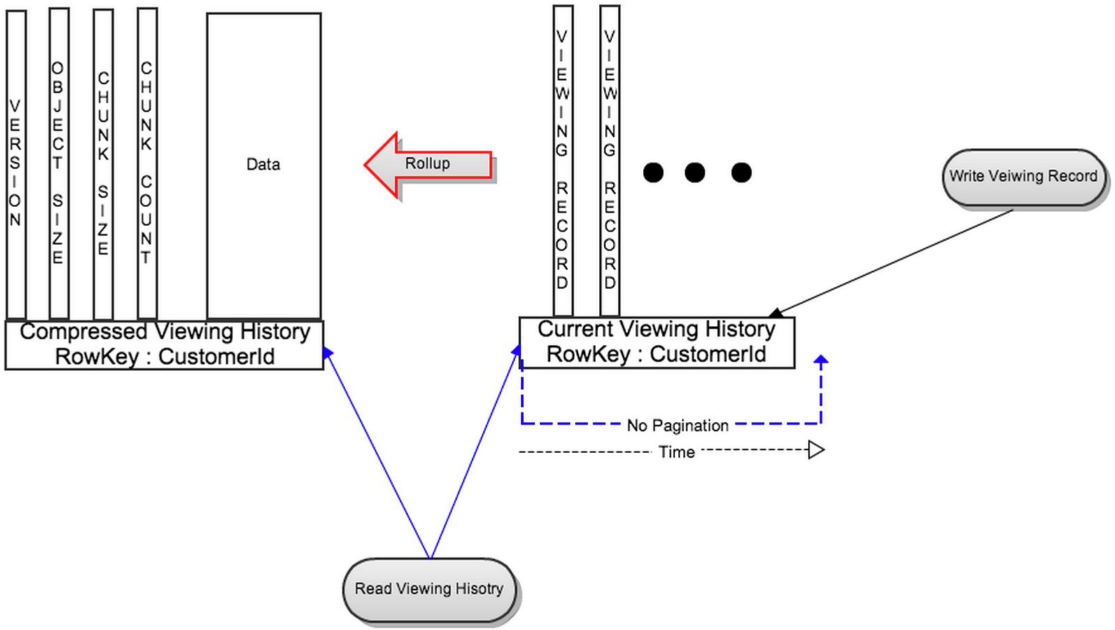

★ Digital Asset Management

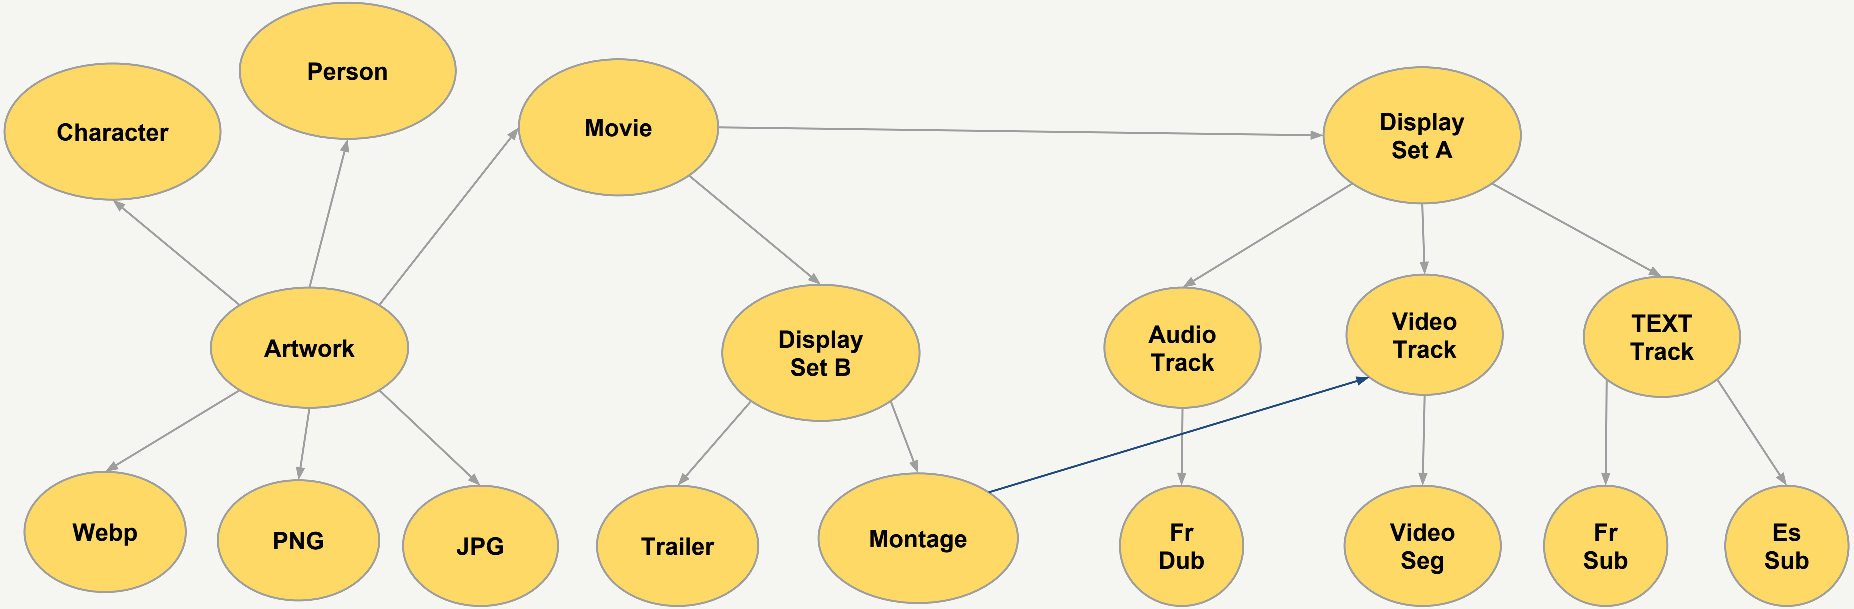

Requirements - DAM

- One backend plane for all asset metadata
- Storage of relationships/connected data
- Searchable

elastic
DYNOMITE
EVCache
cassandra

GUESS?

- Distributed GraphDB
- Support for various storage backends

★Distributed Delayed Queues

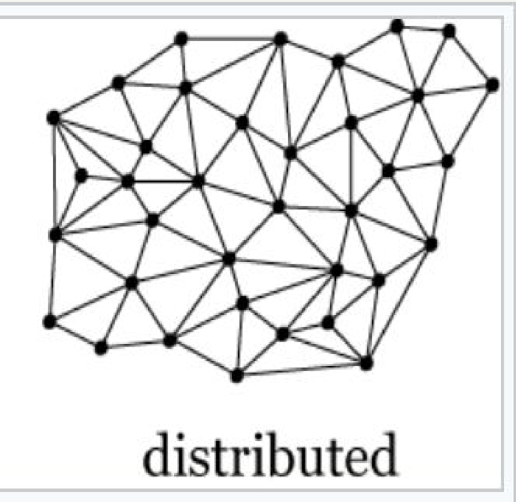

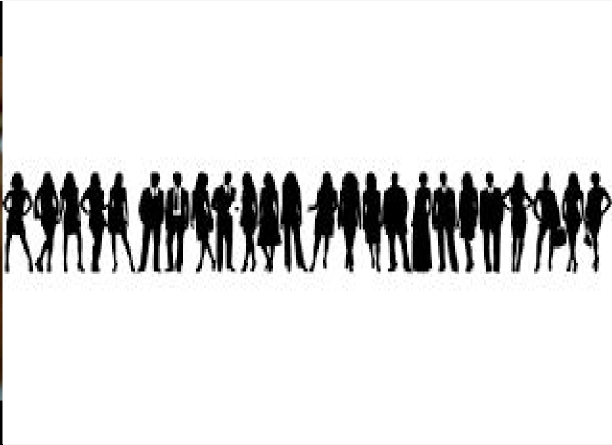

Requirements - Delayed Queues

- Distributed
- Highly concurrent
- At-least-once delivery semantics
- Delayed queue
- Priorities within the shard

elastic
DYN MITE
EVCache
cassandra

GUESS?

- Pluggable datastore supporting Redis
- Multi-datacenter replication

Data Model

For each queue three set of Redis data structures are maintained:

1. A Sorted Set containing queued elements by score.
2. A Hash set that contains message payload, with key as message ID.
3. A Sorted Set containing messages consumed by client but yet to be acknowledged. Un-ack set.

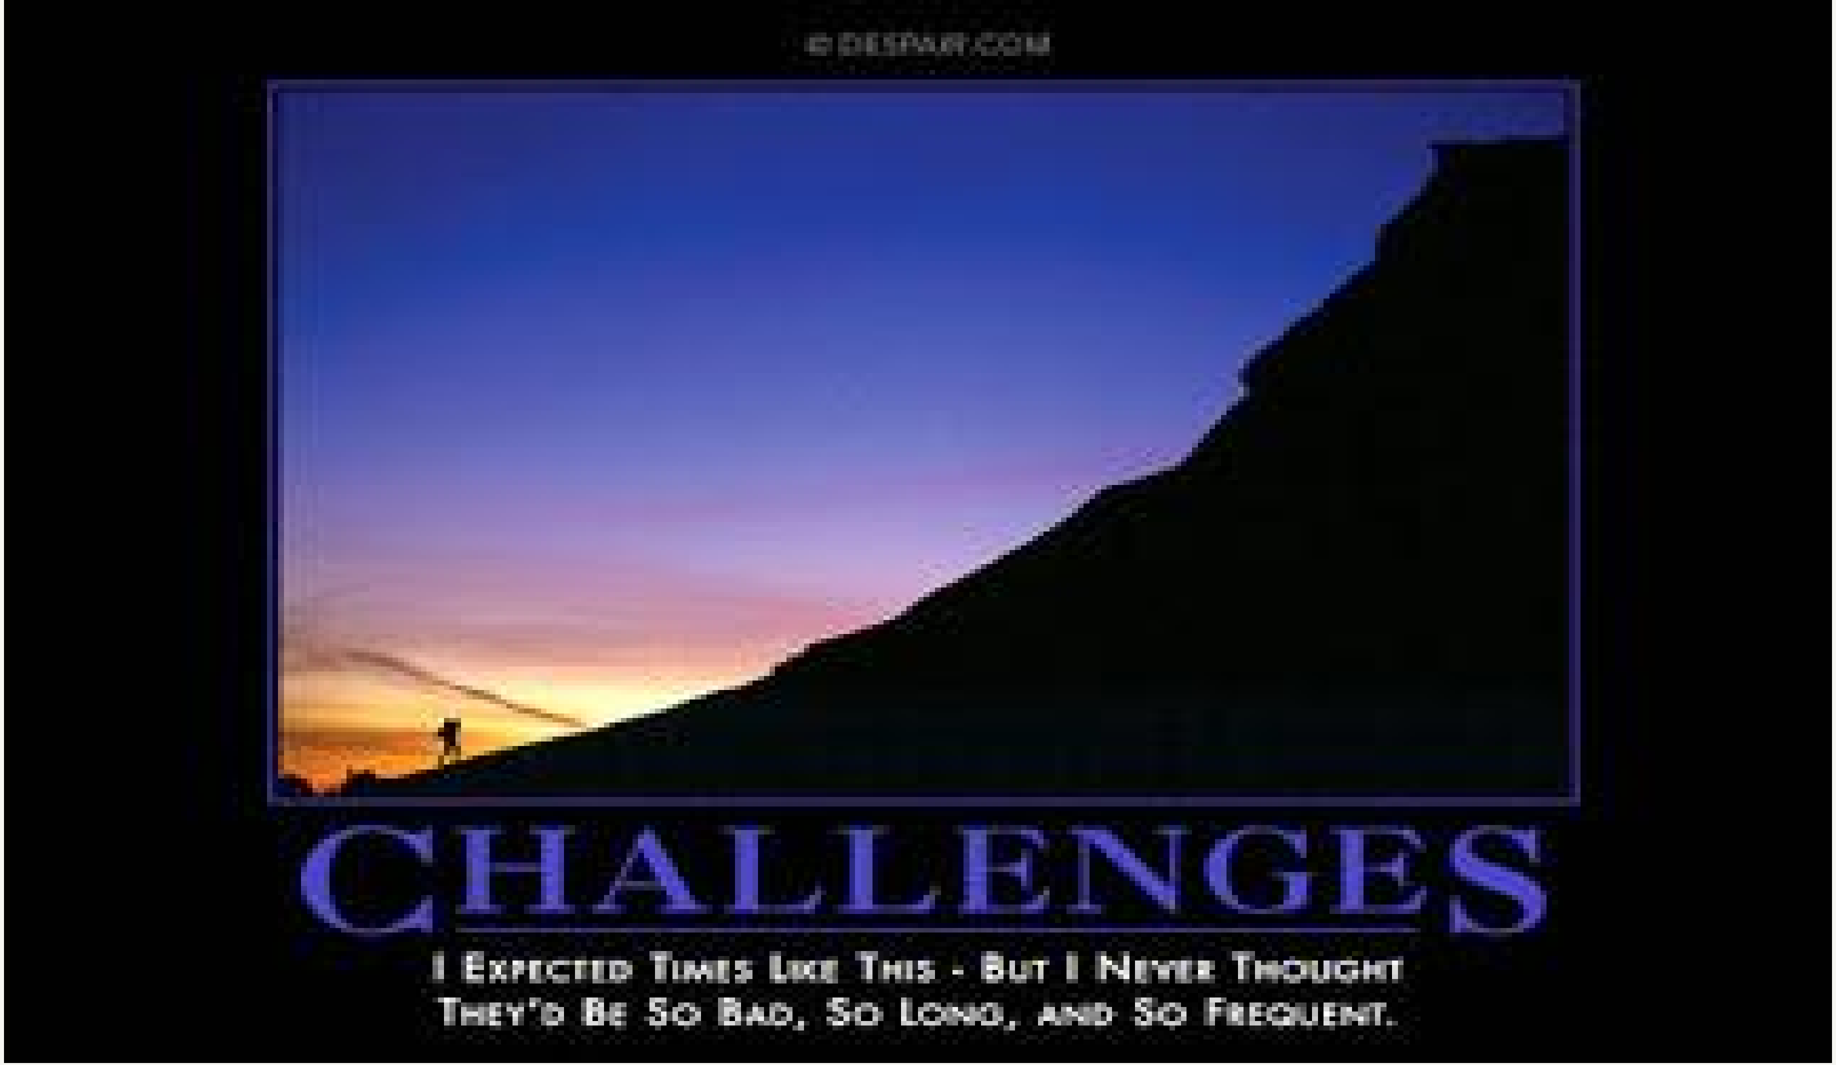

© DESPAIR.COM
CHALLENGES
I Expected Times Like This - But I Never Thought
They'd Be So Bad, So Long, and So Frequent.

MAINTENANCE

25%
50%
75%
100%
Oct
Nov
Dec
100 %
60 %
35 %
20 %
30%
20%
35%
5%
25%
50%
75%
100%
87%
36%
74%
24%
ON
OFF

Current Approach

Subject Matter Expert

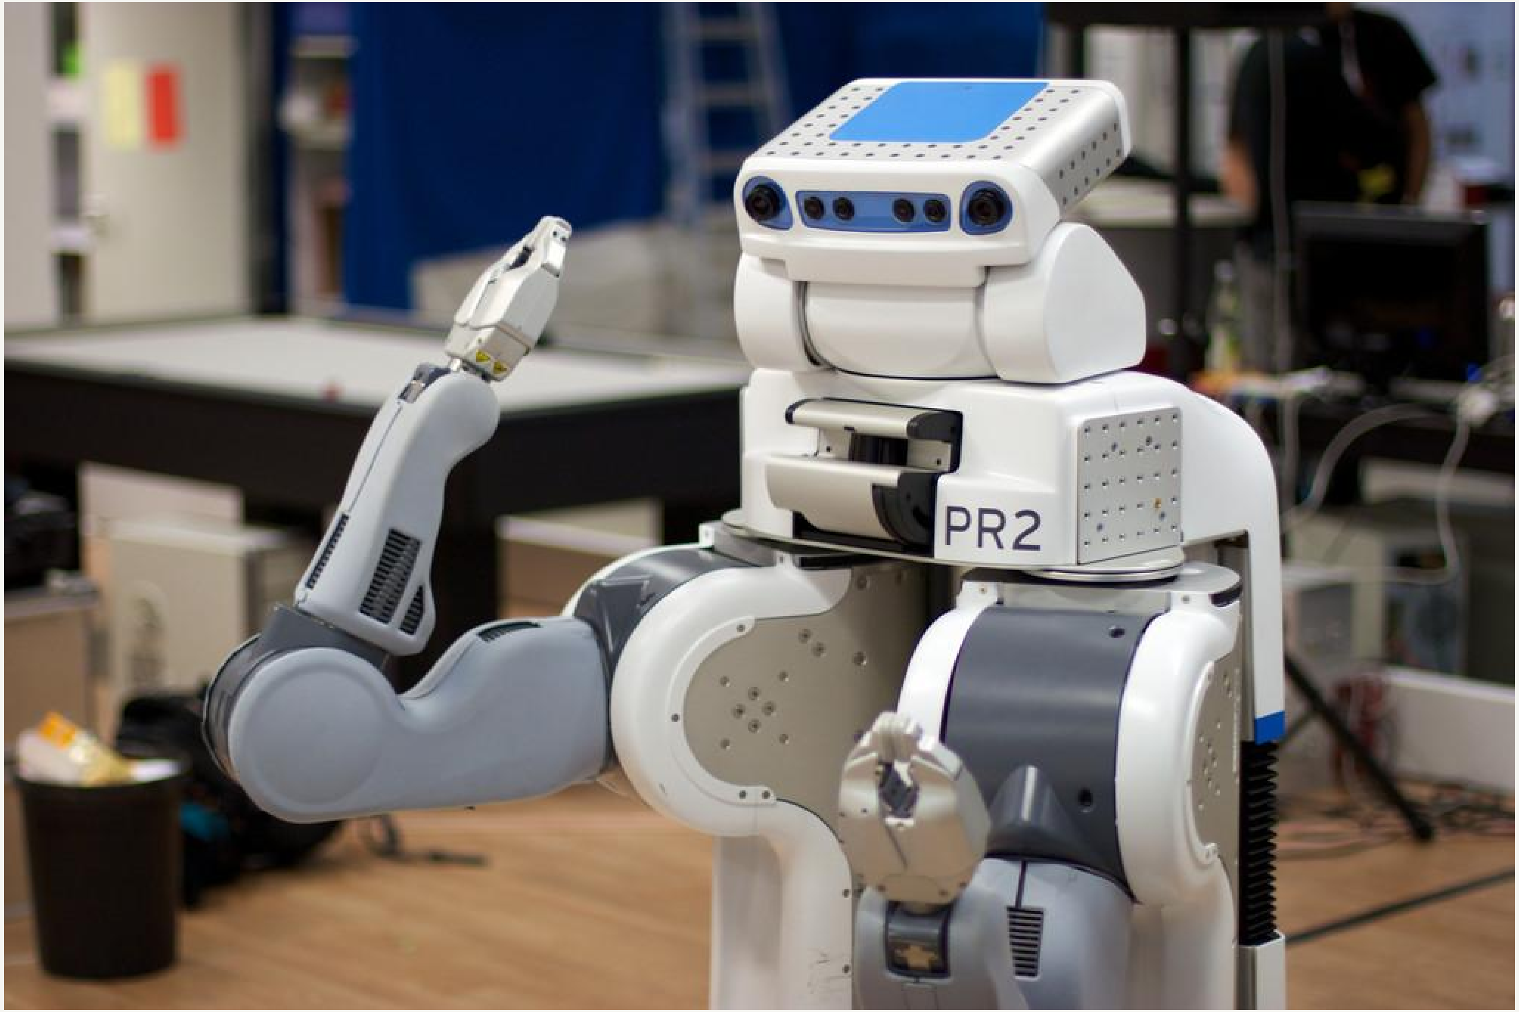
PR2

CDE Service

“Empowering CDE to provide datastores as a service”

CDE Service

- Thresholds/SLAs
- Cluster metadata
- Self Service
- Contact information
- Maintenance windows

Architecture

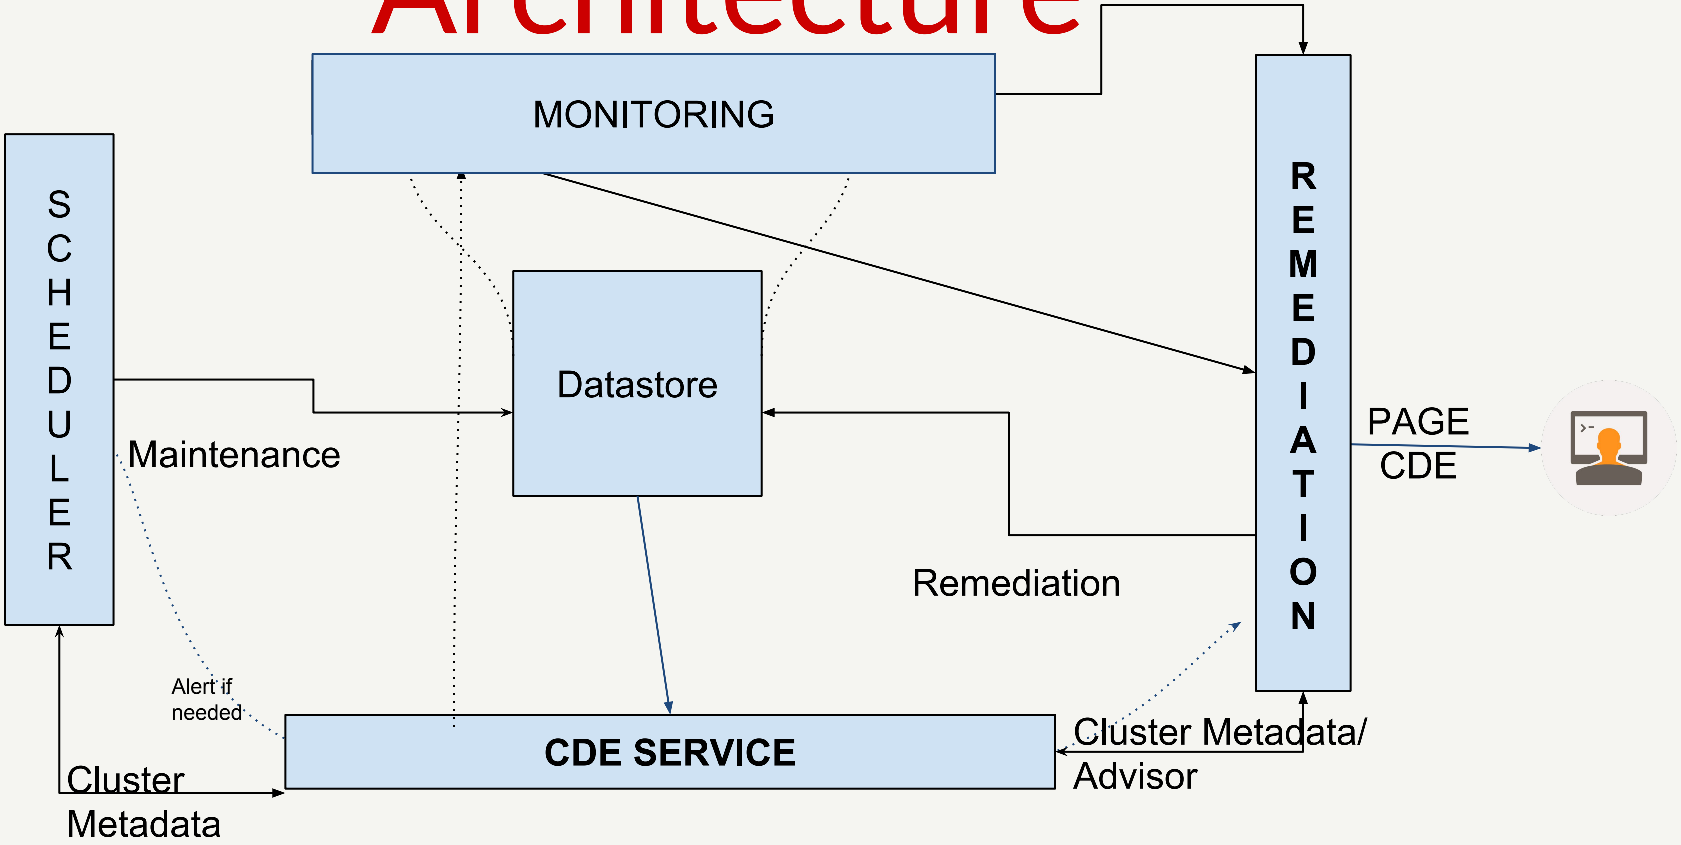

SLA

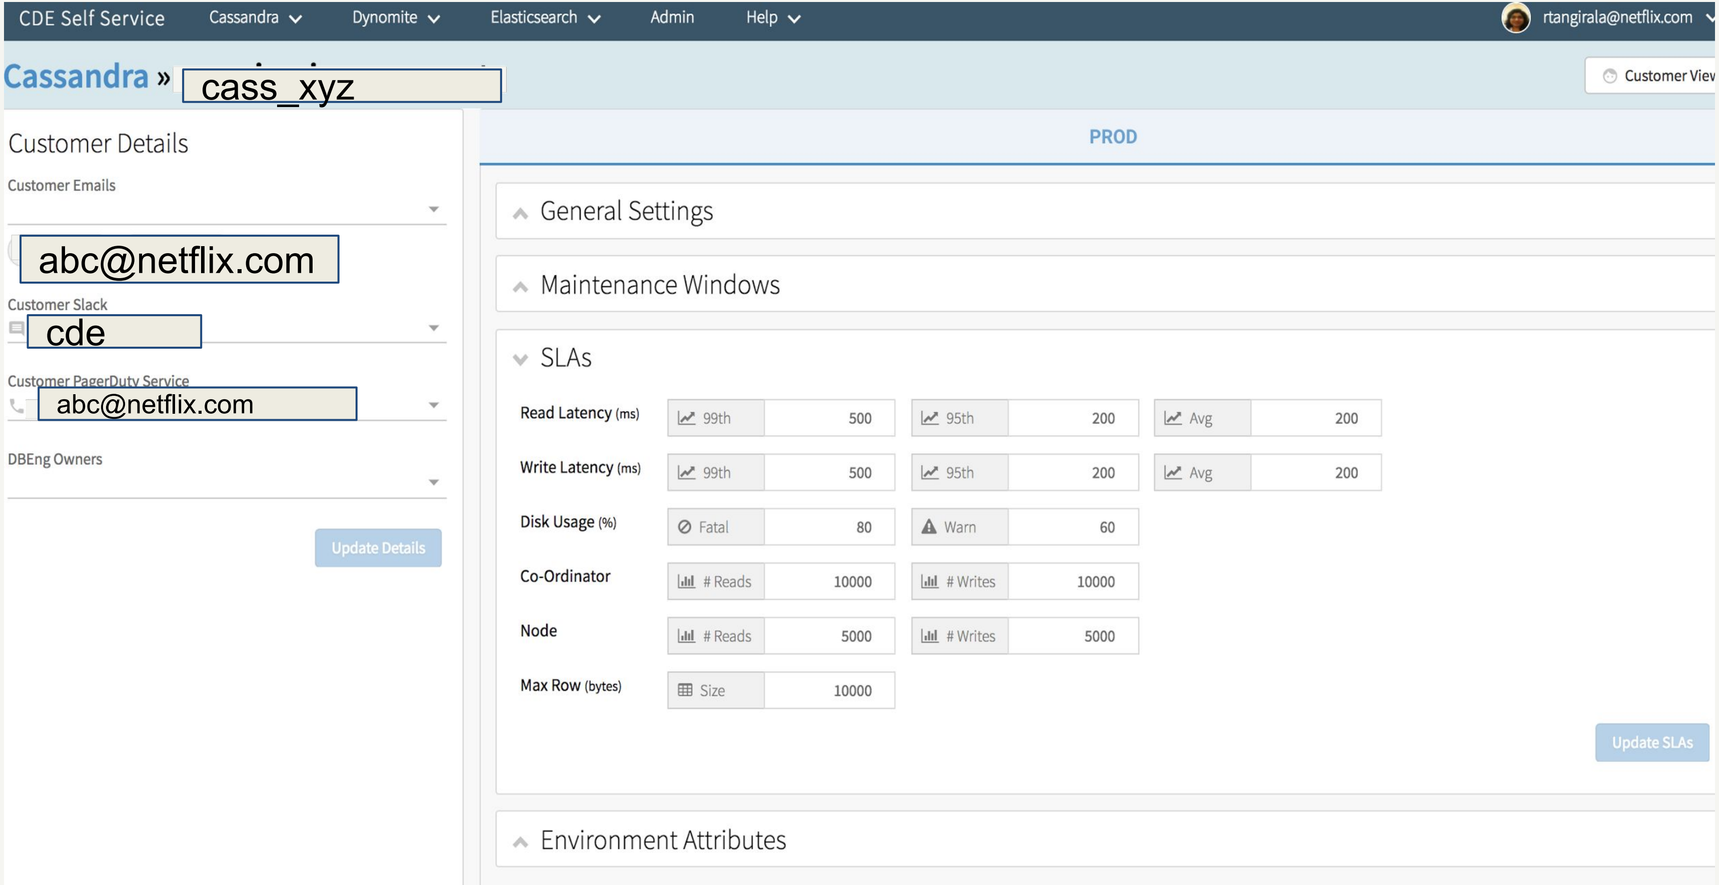

Cassandra » cass_xyz

Customer Details

Customer Emails: abc@netflix.com

Customer Slack: cde

Customer PagerDuty Service: abc@netflix.com

DBEng Owners:

SLAs

Read Latency (ms)	99th	500	95th	200	Avg	200
Write Latency (ms)	99th	500	95th	200	Avg	200
Disk Usage (%)	Fatal	80	Warn	60		
Co-Ordinator	# Reads	10000	# Writes	10000		
Node	# Reads	5000	# Writes	5000		
Max Row (bytes)	Size	10000				

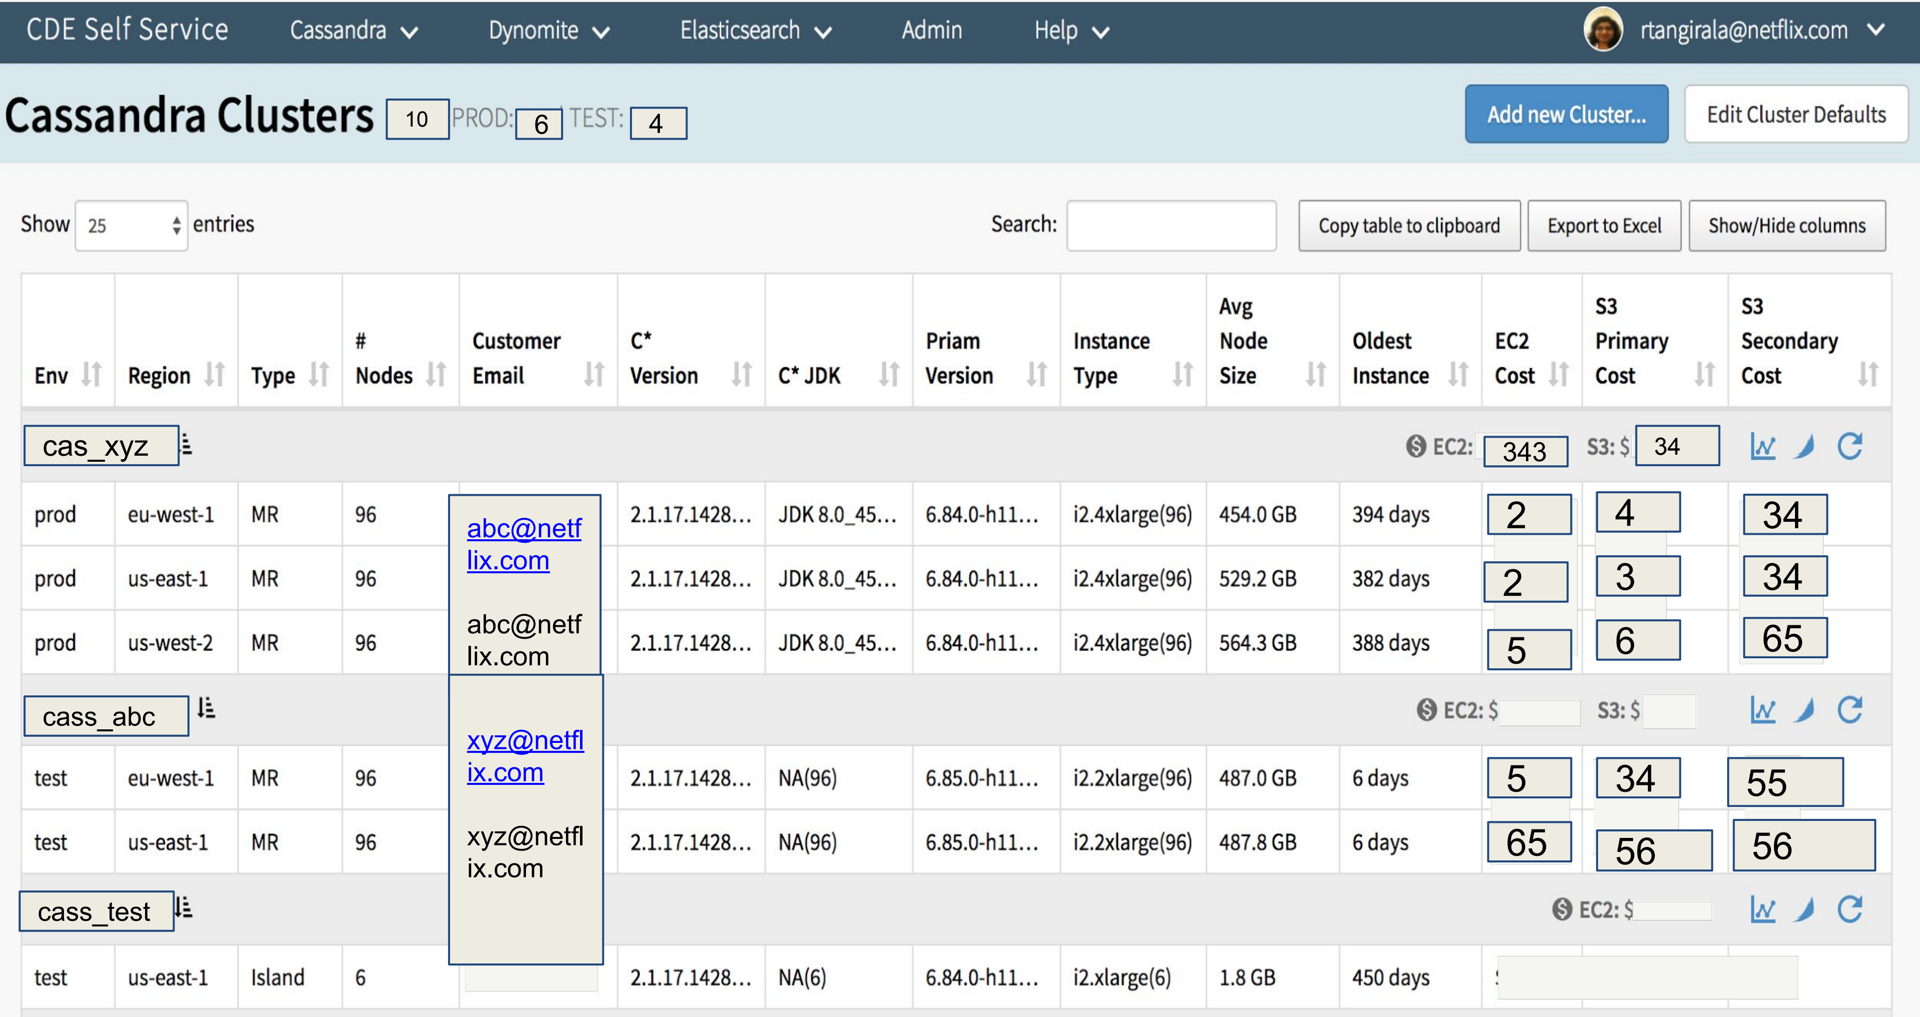

Cassandra Clusters 10 PROD: 6 TEST: 4

Env	Region	Type	# Nodes	Customer Email	C* Version	C* JDK	Priam Version	Instance Type	Avg Node Size	Oldest Instance	EC2 Cost	S3 Primary Cost	S3 Secondary Cost
cas_xyz											EC2: 343	S3: $ 34	
prod	eu-west-1	MR	96	abc@netflix.com	2.1.17.1428...	JDK 8.0_45...	6.84.0-h11...	i2.4xlarge(96)	454.0 GB	394 days	2	4	34
prod	us-east-1	MR	96		2.1.17.1428...	JDK 8.0_45...	6.84.0-h11...	i2.4xlarge(96)	529.2 GB	382 days	2	3	34
prod	us-west-2	MR	96	abc@netflix.com	2.1.17.1428...	JDK 8.0_45...	6.84.0-h11...	i2.4xlarge(96)	564.3 GB	388 days	5	6	65
cass_abc											EC2: $	S3: $	
test	eu-west-1	MR	96	xyz@netflix.com	2.1.17.1428...	NA(96)	6.85.0-h11...	i2.2xlarge(96)	487.0 GB	6 days	5	34	55
test	us-east-1	MR	96	xyz@netflix.com	2.1.17.1428...	NA(96)	6.85.0-h11...	i2.2xlarge(96)	487.8 GB	6 days	65	56	56
cass_test											EC2: $		
test	us-east-1	Island	6		2.1.17.1428...	NA(6)	6.84.0-h11...	i2.xlarge(6)	1.8 GB	450 days			

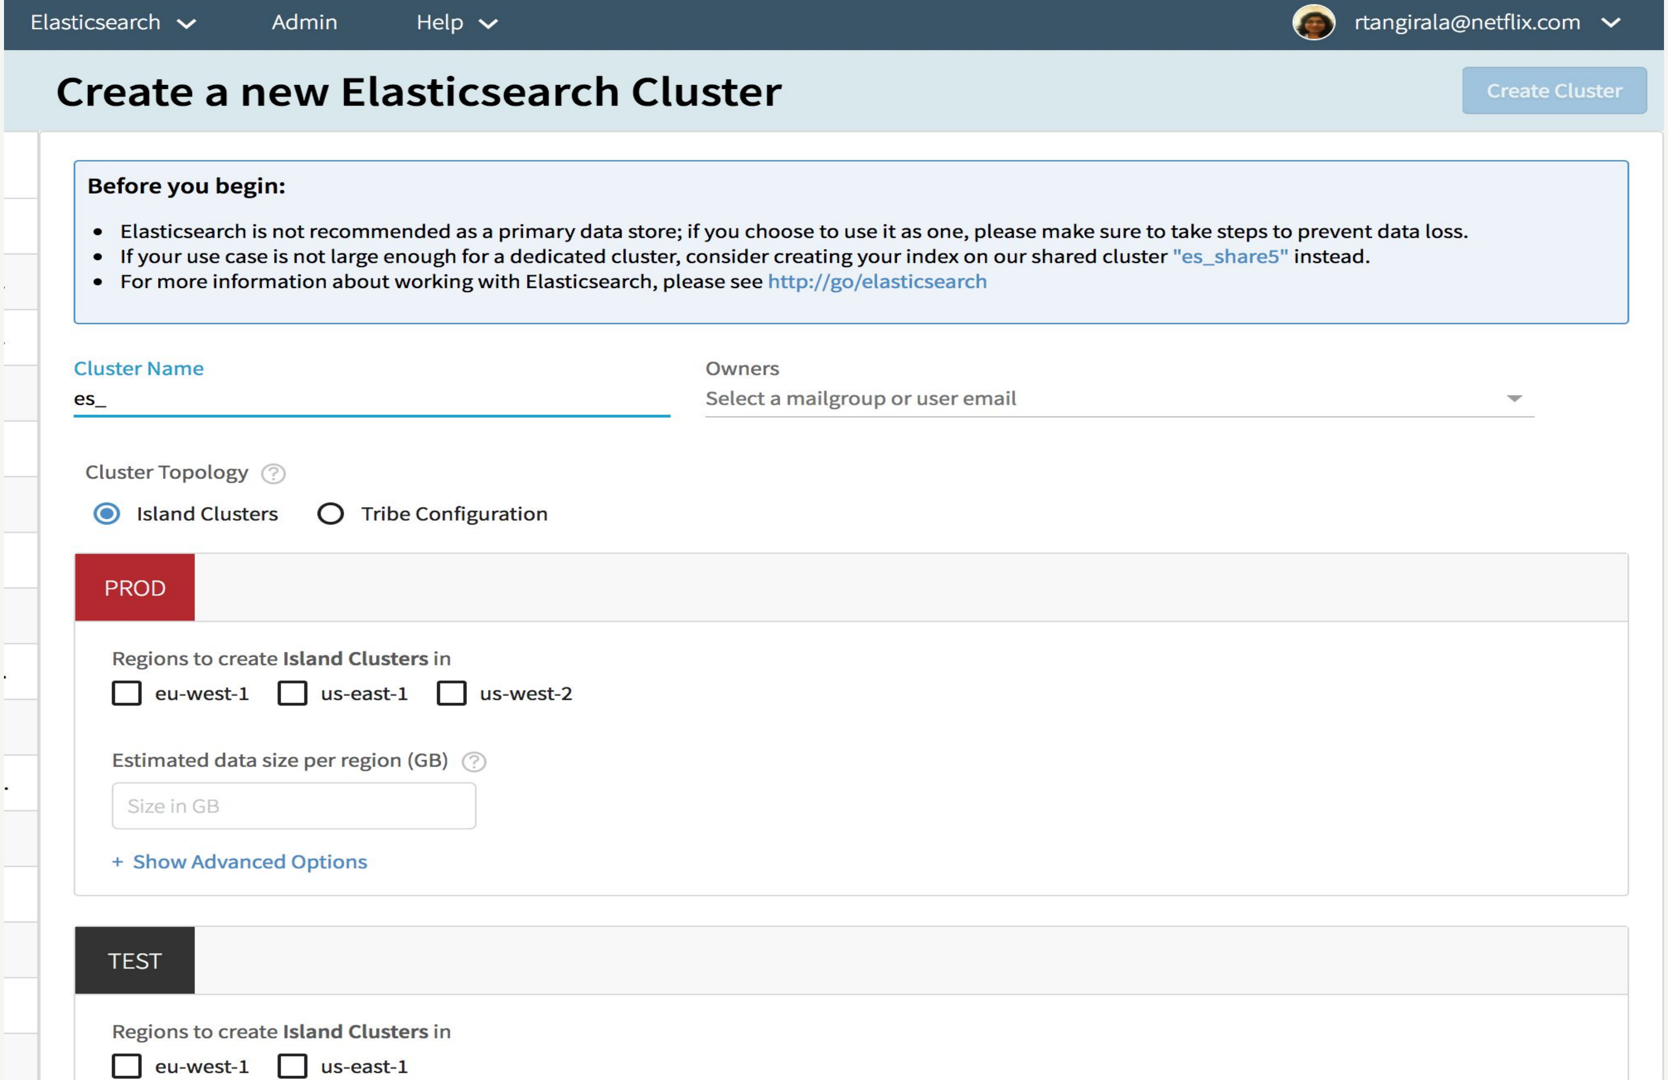

Create a new Elasticsearch Cluster

Before you begin:

- Elasticsearch is not recommended as a primary data store; if you choose to use it as one, please make sure to take steps to prevent data loss.
- If your use case is not large enough for a dedicated cluster, consider creating your index on our shared cluster "es_share5" instead.
- For more information about working with Elasticsearch, please see http://go/elasticsearch

Cluster Name: es_

Owners: Select a mailgroup or user email

Cluster Topology

● Island Clusters ○ Tribe Configuration

PROD

Regions to create **Island Clusters** in

□ eu-west-1 □ us-east-1 □ us-west-2

Estimated data size per region (GB)

Size in GB

+ Show Advanced Options

TEST

Regions to create **Island Clusters** in

□ eu-west-1 □ us-east-1

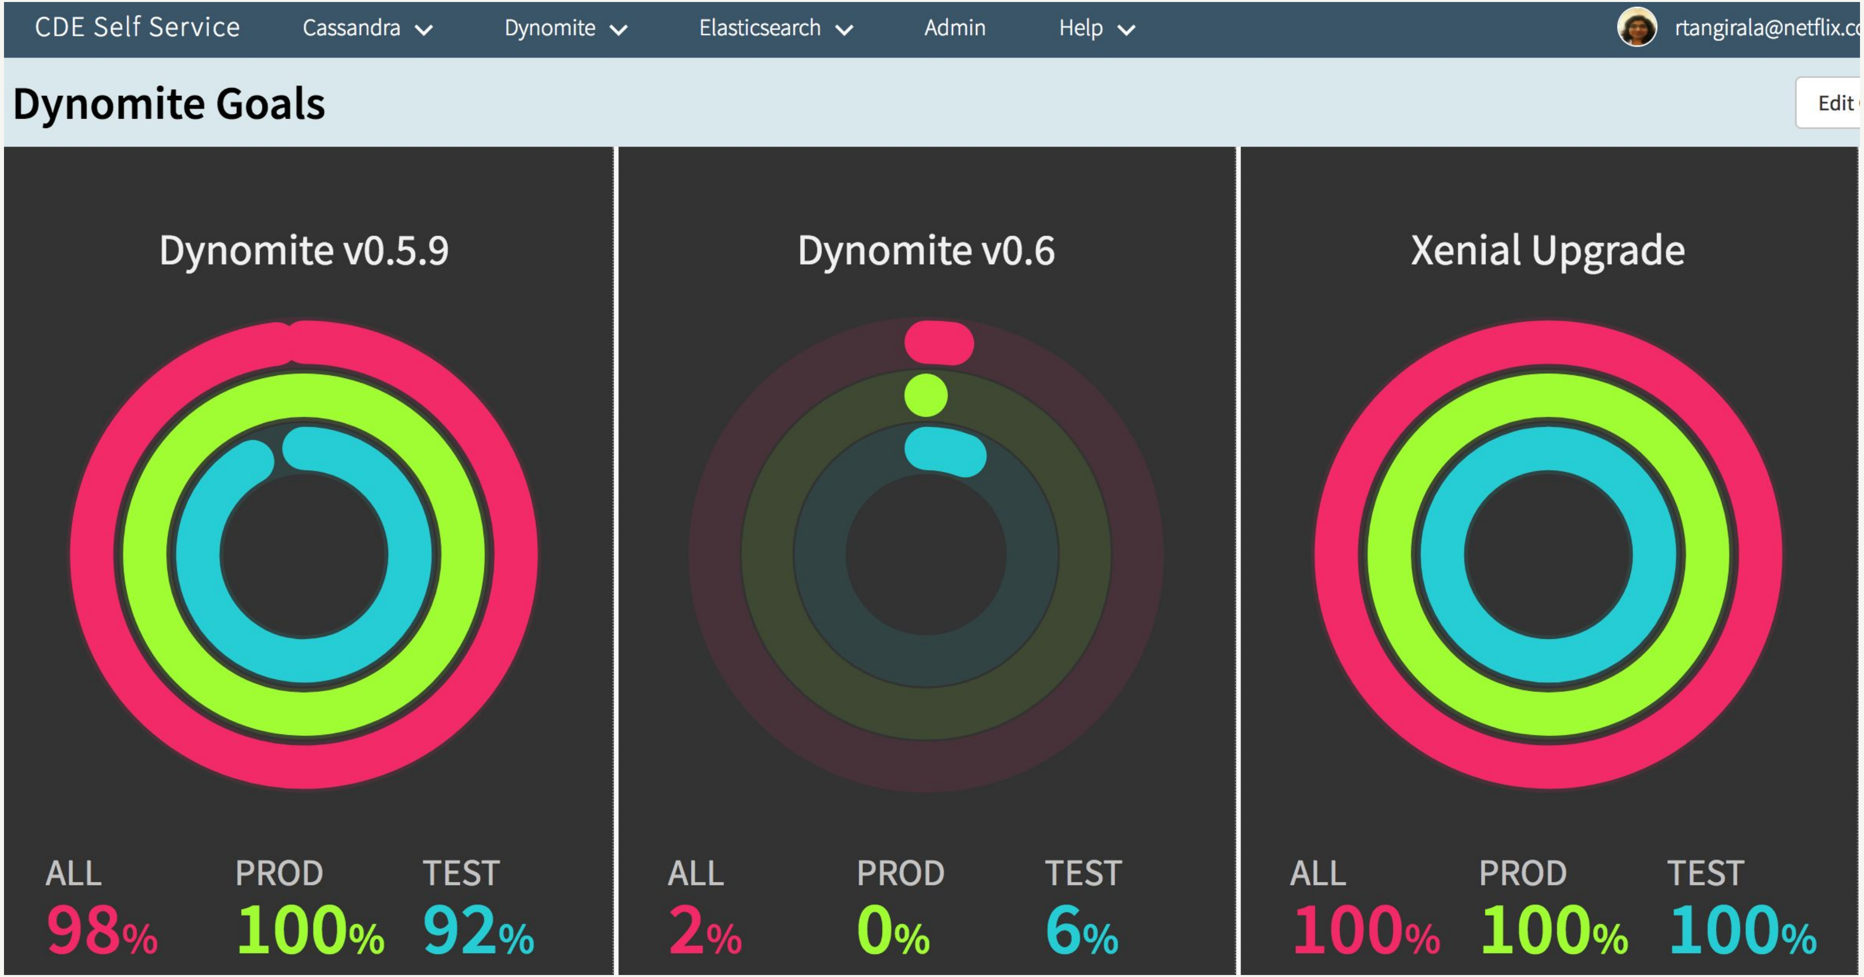

Dynomite Goals

Dynomite v0.5.9

ALL 98%

PROD 100%

TEST 92%

Dynomite v0.6

ALL 2%

PROD 0%

TEST 6%

Xenial Upgrade

ALL 100%

PROD 100%

TEST 100%

Machine learning

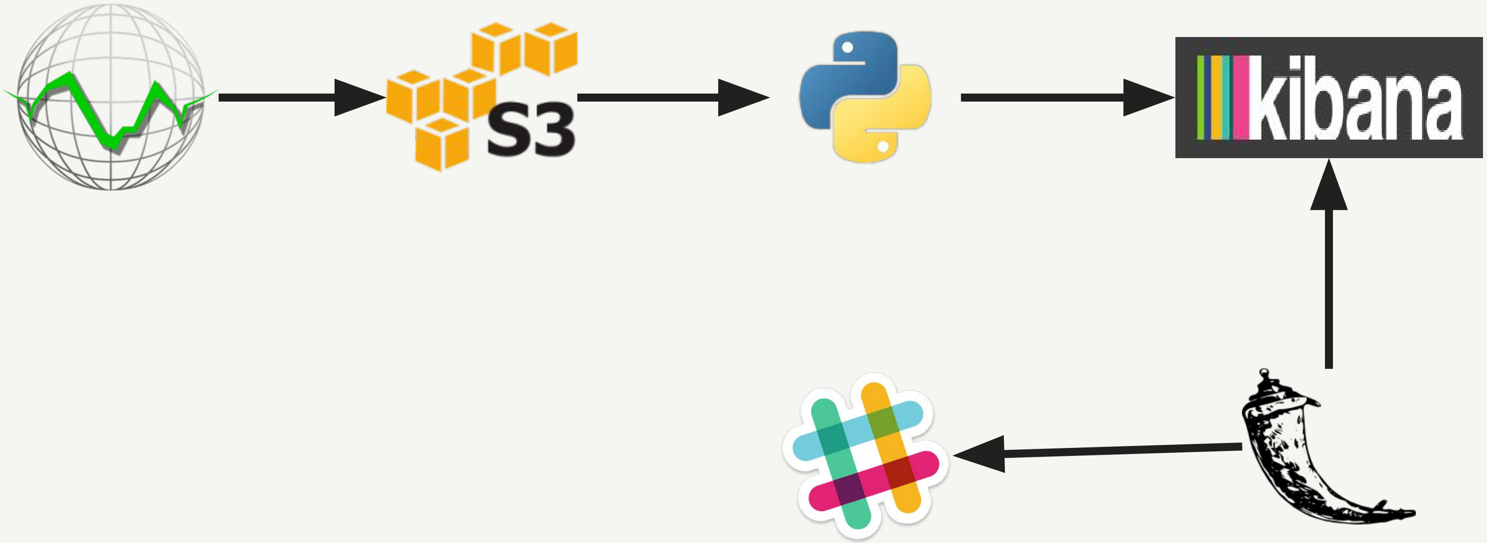

Pattern in Disk usage

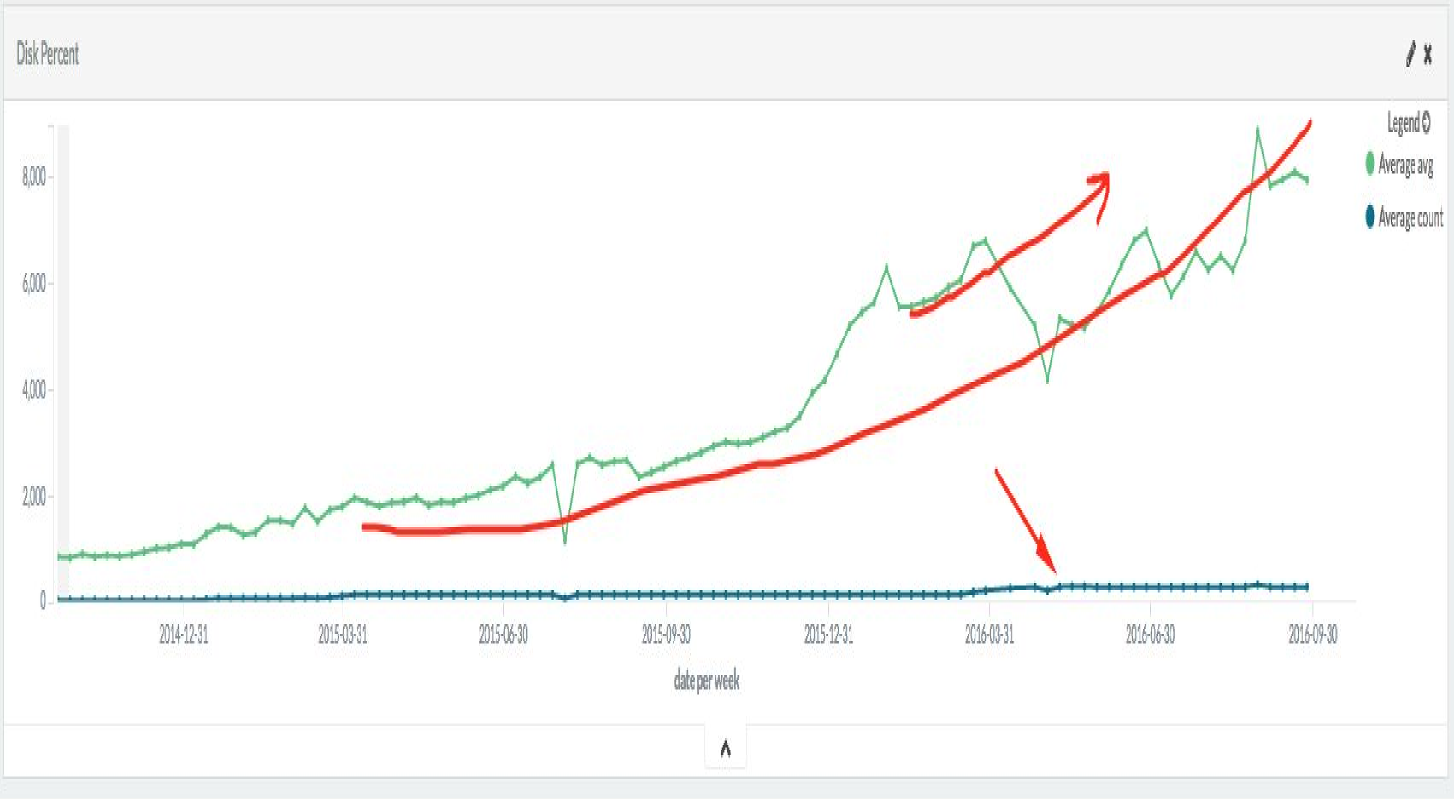

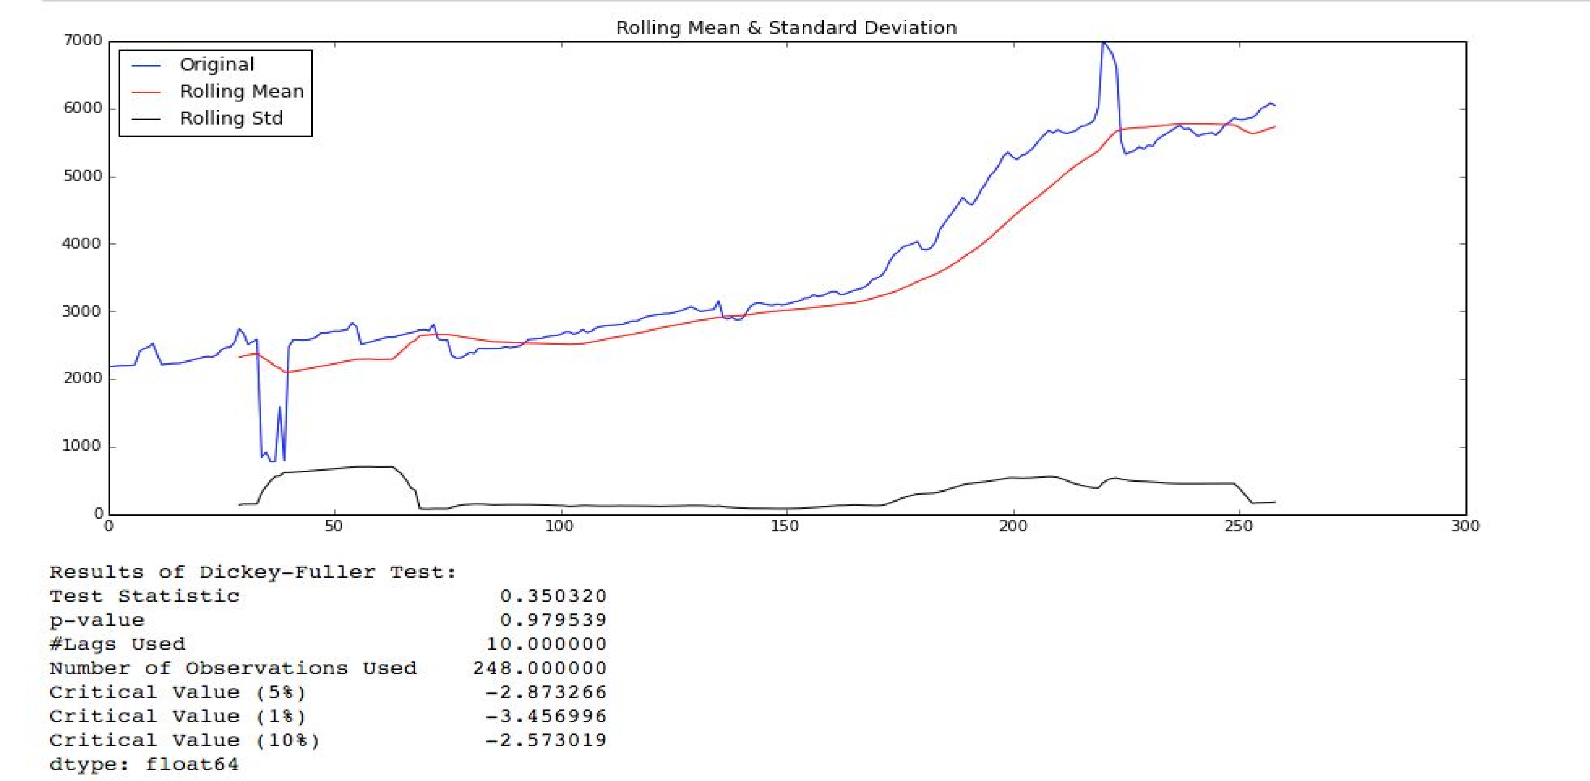

```
Results of Dickey-Fuller Test:
Test Statistic                   0.350320
p-value                          0.979539
#Lags Used                      10.000000
Number of Observations Used    248.000000
Critical Value (5%)             -2.873266
Critical Value (1%)             -3.456996
Critical Value (10%)            -2.573019
dtype: float64
```

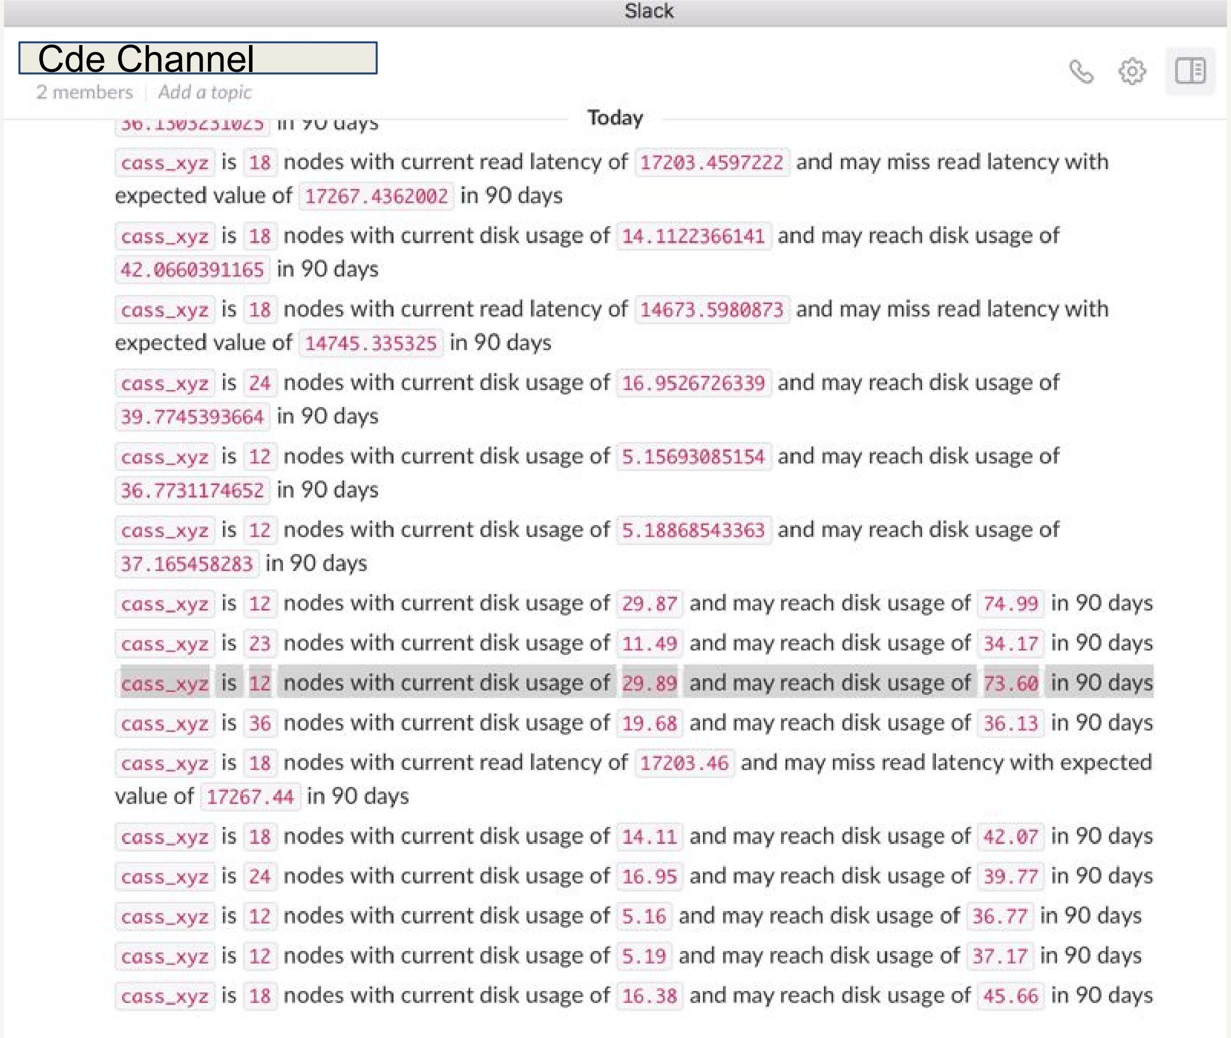

Cde Channel

2 members | Add a topic

Today

36.1303231025 in 90 days

cass_xyz is 18 nodes with current read latency of 17203.4597222 and may miss read latency with expected value of 17267.4362002 in 90 days

cass_xyz is 18 nodes with current disk usage of 14.1122366141 and may reach disk usage of 42.0660391165 in 90 days

cass_xyz is 18 nodes with current read latency of 14673.5980873 and may miss read latency with expected value of 14745.335325 in 90 days

cass_xyz is 24 nodes with current disk usage of 16.9526726339 and may reach disk usage of 39.7745393664 in 90 days

cass_xyz is 12 nodes with current disk usage of 5.15693085154 and may reach disk usage of 36.7731174652 in 90 days

cass_xyz is 12 nodes with current disk usage of 5.18868543363 and may reach disk usage of 37.165458283 in 90 days

cass_xyz is 12 nodes with current disk usage of 29.87 and may reach disk usage of 74.99 in 90 days

cass_xyz is 23 nodes with current disk usage of 11.49 and may reach disk usage of 34.17 in 90 days

cass_xyz is 12 nodes with current disk usage of 29.89 and may reach disk usage of 73.60 in 90 days

cass_xyz is 36 nodes with current disk usage of 19.68 and may reach disk usage of 36.13 in 90 days

cass_xyz is 18 nodes with current read latency of 17203.46 and may miss read latency with expected value of 17267.44 in 90 days

cass_xyz is 18 nodes with current disk usage of 14.11 and may reach disk usage of 42.07 in 90 days

cass_xyz is 24 nodes with current disk usage of 16.95 and may reach disk usage of 39.77 in 90 days

cass_xyz is 12 nodes with current disk usage of 5.16 and may reach disk usage of 36.77 in 90 days

cass_xyz is 12 nodes with current disk usage of 5.19 and may reach disk usage of 37.17 in 90 days

cass_xyz is 18 nodes with current disk usage of 16.38 and may reach disk usage of 45.66 in 90 days

PROACTIVE

UPGRADE

NDBench

Common Approach

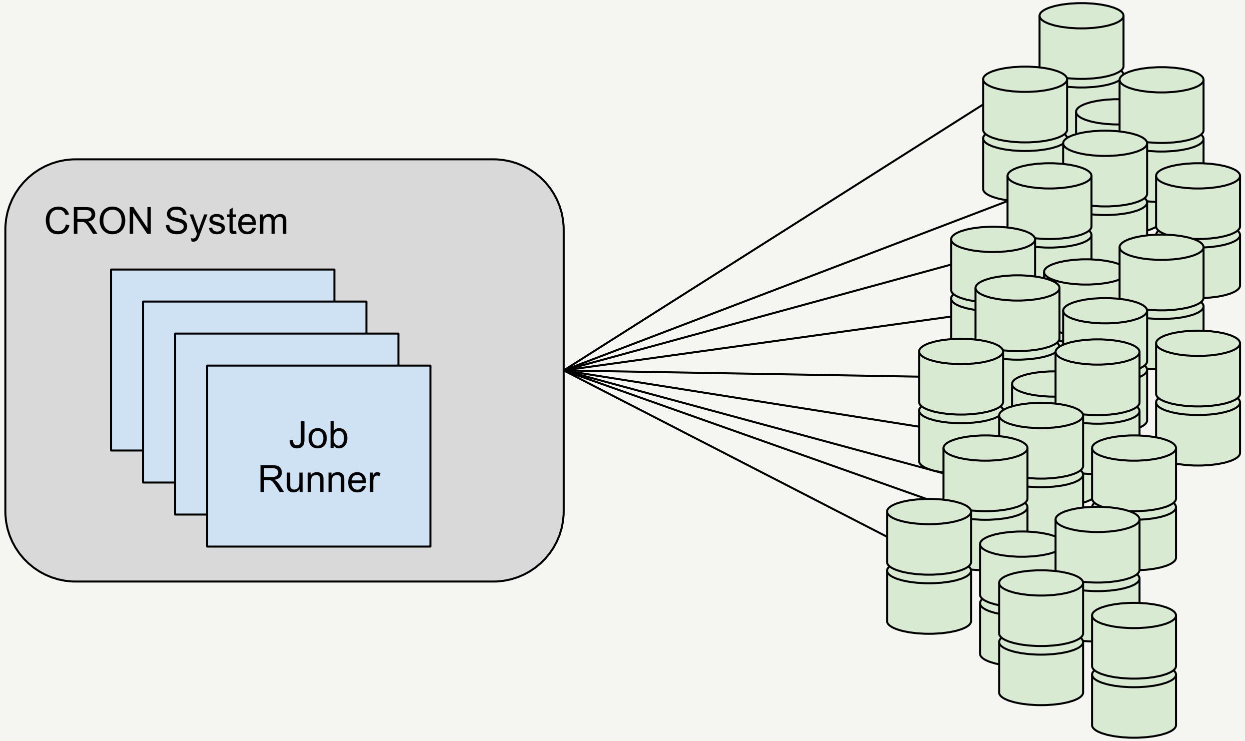

Streaming micro-services

Source - input, handles backpressure
Stage - business logic
Sink - output, handles backpressure

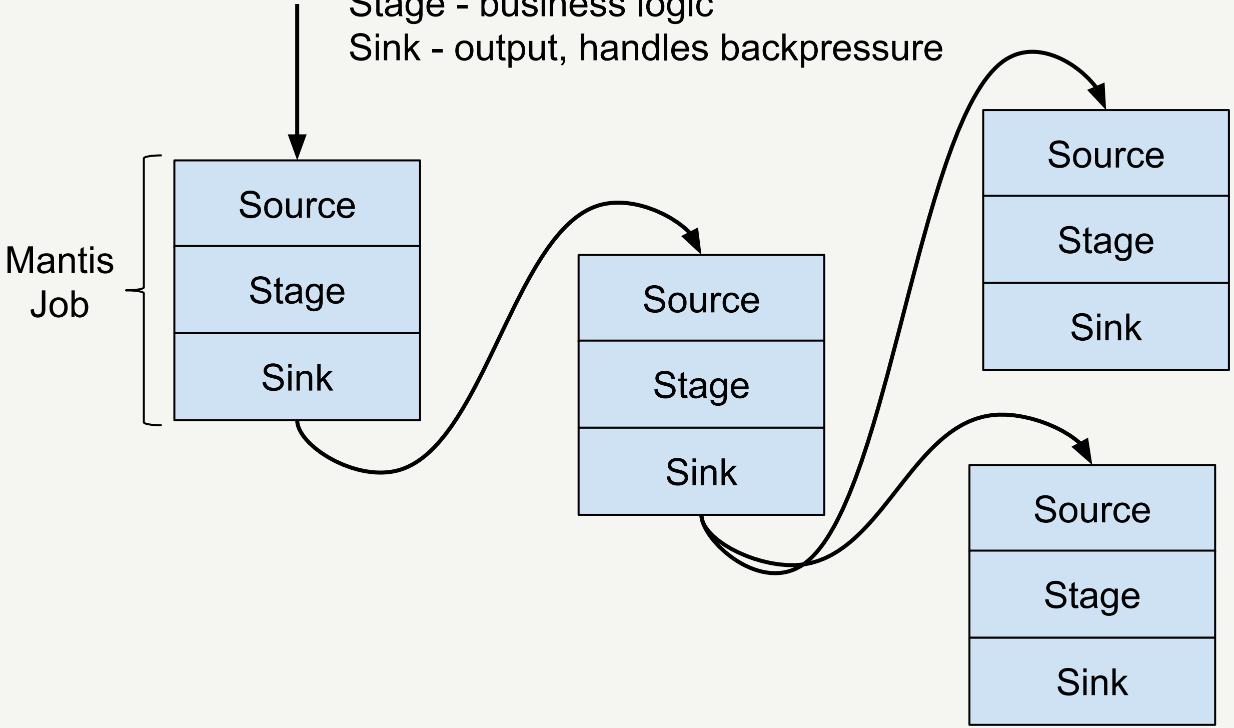

Real Time Dash (Macro View)

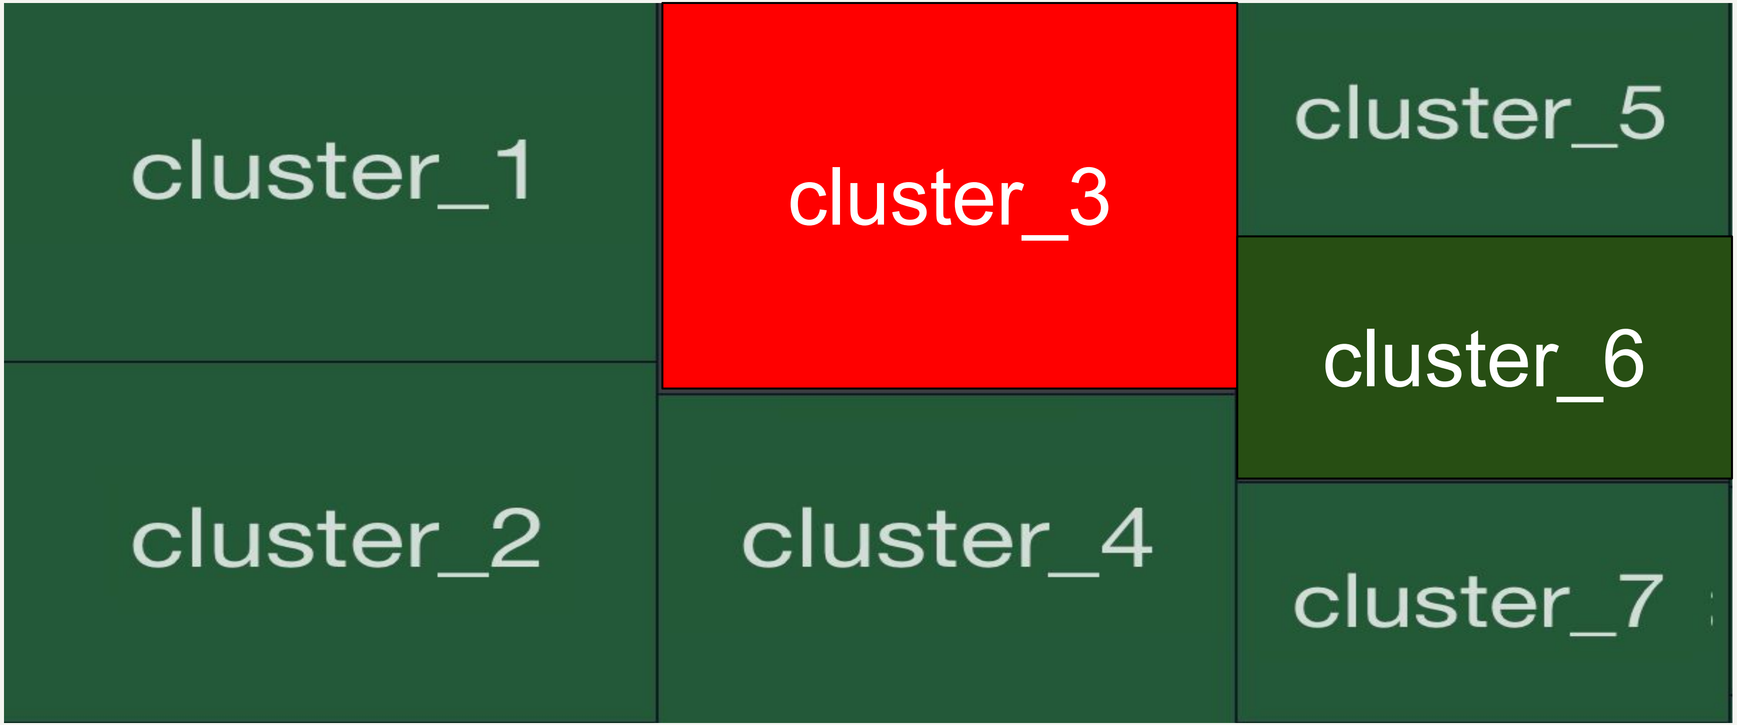

Real Time Dash (Cluster View)

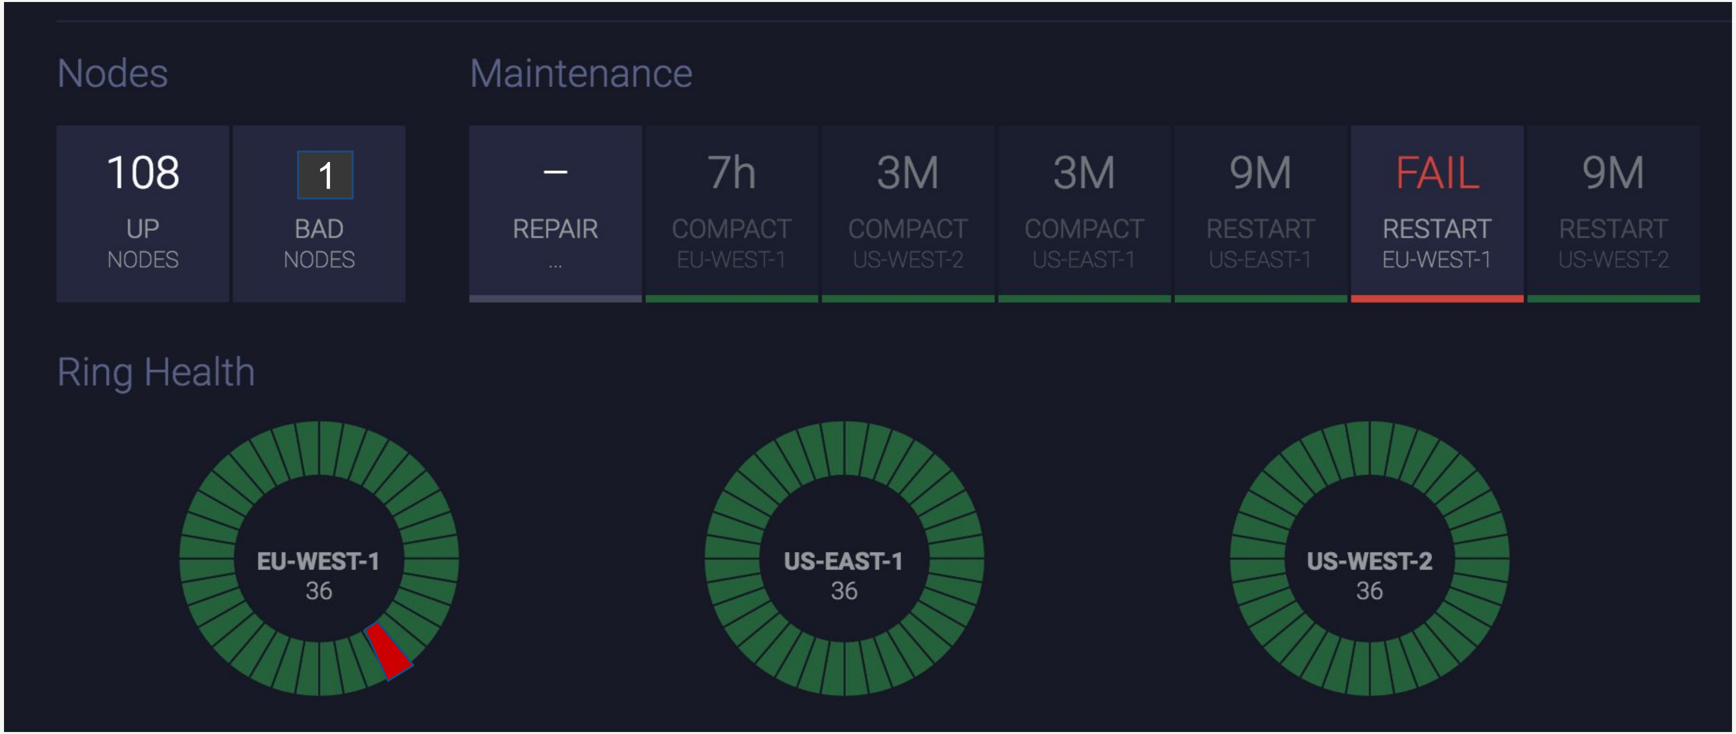

Nodes

108 UP NODES	1 BAD NODES

Maintenance

– REPAIR ...	7h COMPACT EU-WEST-1	3M COMPACT US-WEST-2	3M COMPACT US-EAST-1	9M RESTART US-EAST-1	FAIL RESTART EU-WEST-1	9M RESTART US-WEST-2

Ring Health

Takeaway

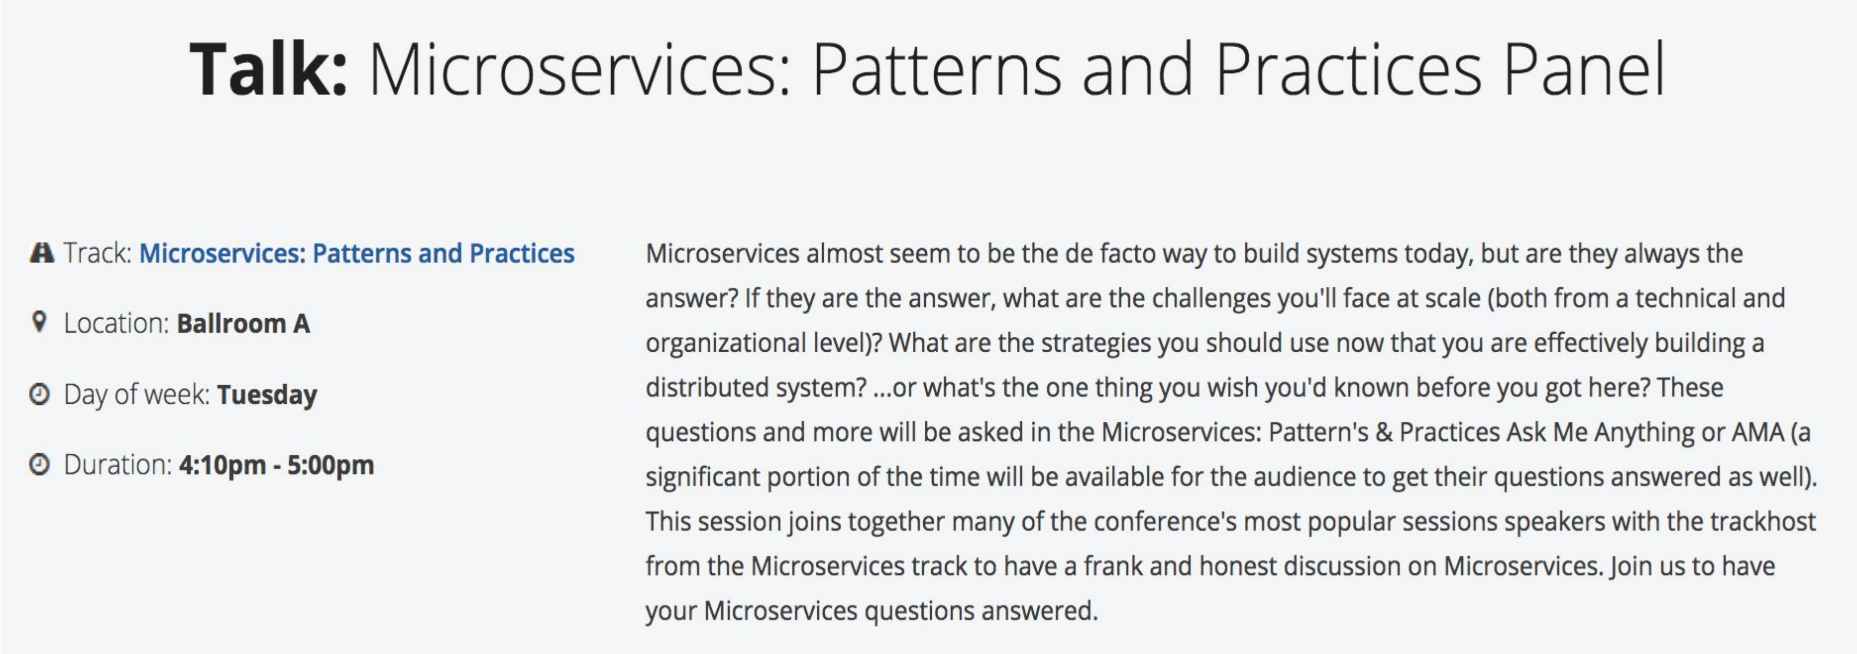

Talk: Microservices: Patterns and Practices Panel

Track: **Microservices: Patterns and Practices**

Location: **Ballroom A**

Day of week: **Tuesday**

Duration: **4:10pm - 5:00pm**

Microservices almost seem to be the de facto way to build systems today, but are they always the answer? If they are the answer, what are the challenges you'll face at scale (both from a technical and organizational level)? What are the strategies you should use now that you are effectively building a distributed system? ...or what's the one thing you wish you'd known before you got here? These questions and more will be asked in the Microservices: Pattern's & Practices Ask Me Anything or AMA (a significant portion of the time will be available for the audience to get their questions answered as well). This session joins together many of the conference's most popular sessions speakers with the trackhost from the Microservices track to have a frank and honest discussion on Microservices. Join us to have your Microservices questions answered.