JESSICA TAI / MARCH 4, 2019 / QCON LONDON

Airbnb's Great Migration: Building Services at Scale

2014

@jessicamtai

2015+

@jessicamtai


Hi, I'm Jessica.
I pair program with my corgi.
@jessicamtai

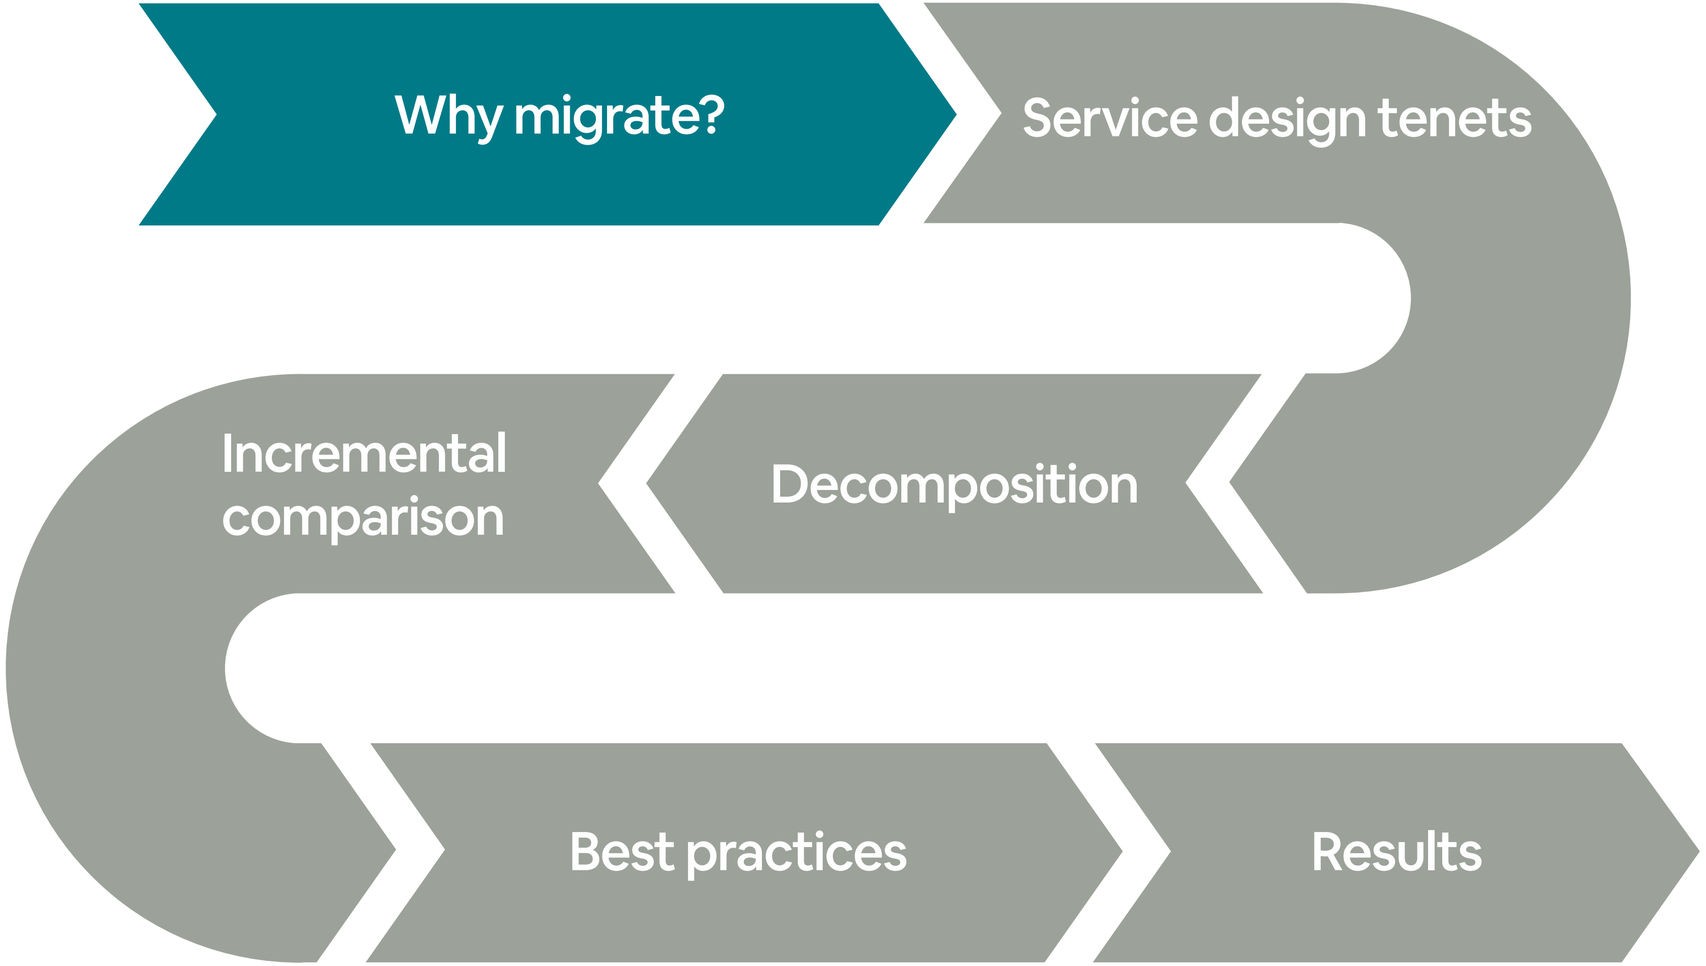
Why migrate?
Service design tenets
Decomposition
Incremental comparison
Best practices
Results

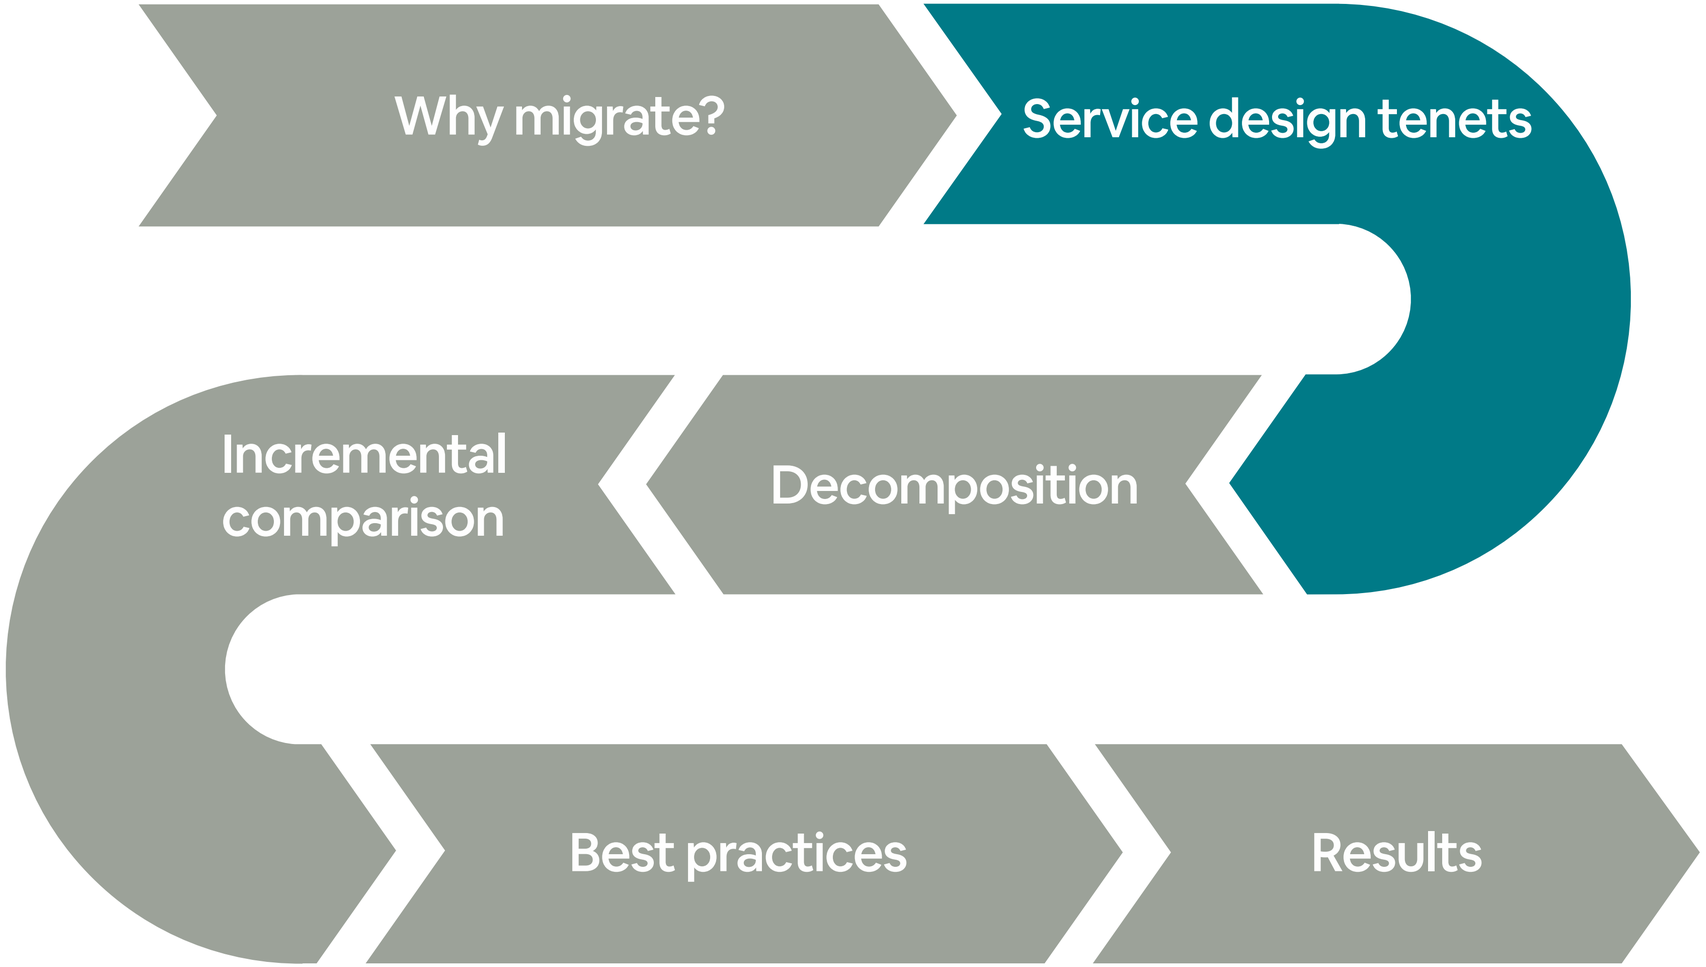
Why migrate?
Service design tenets
Decomposition
Incremental comparison
Best practices
Results

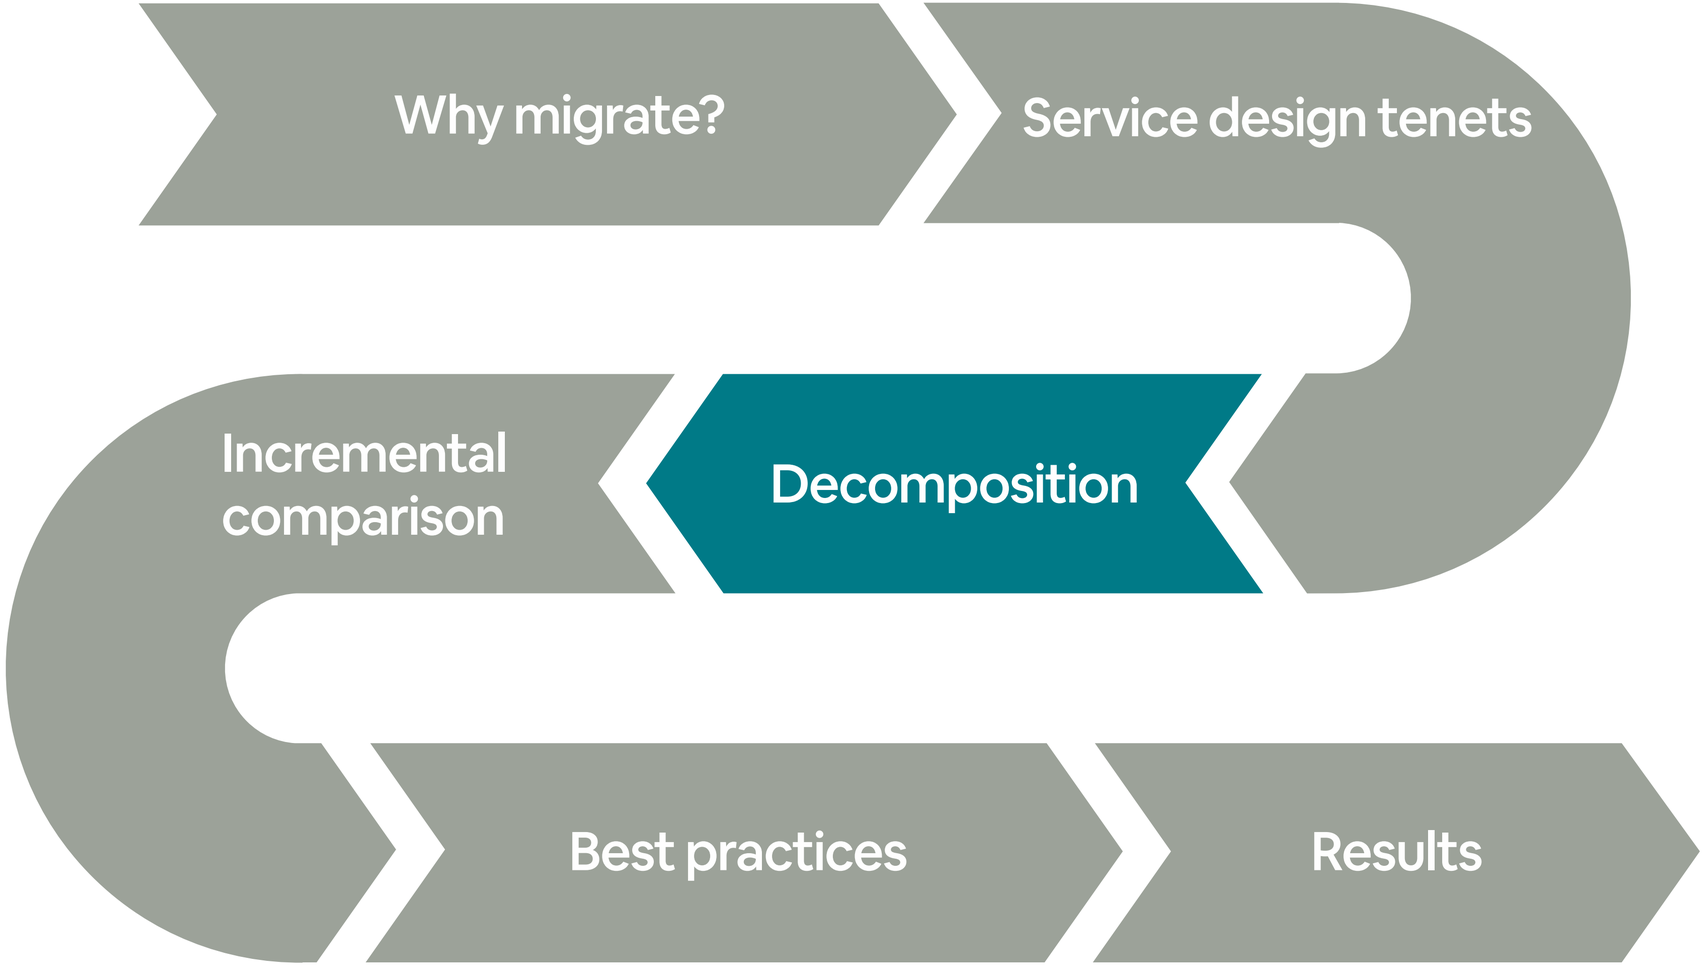
Why migrate?
Service design tenets
Incremental comparison
Decomposition
Best practices
Results

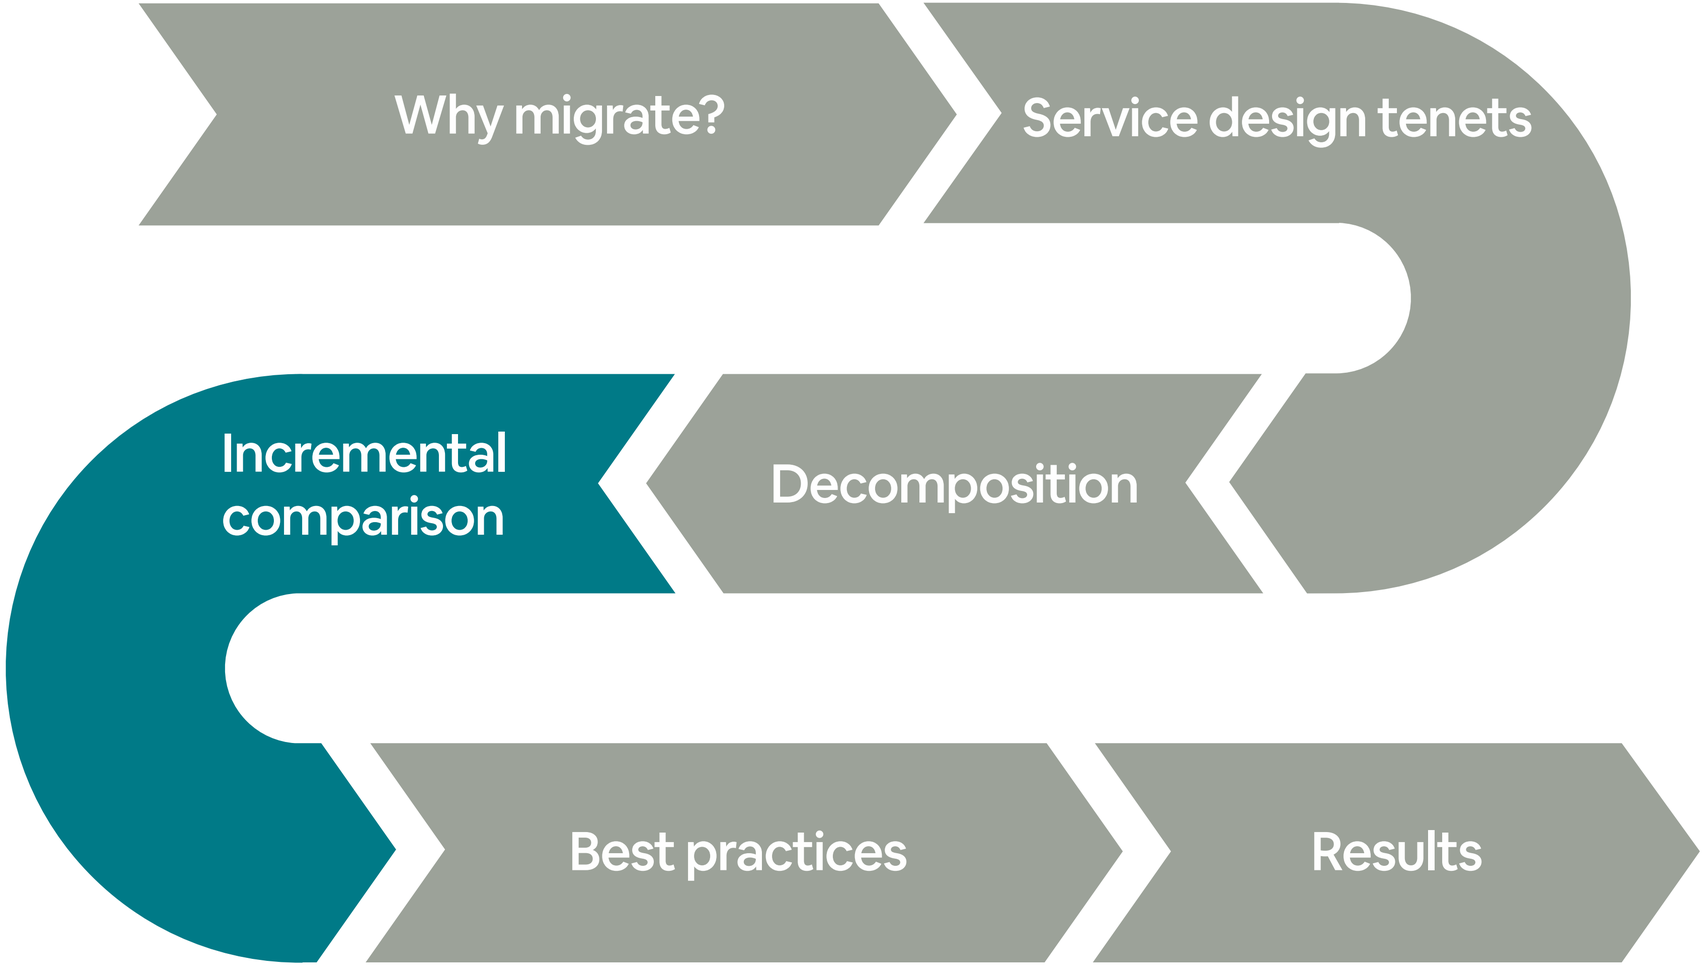
Why migrate?
Service design tenets
Decomposition
Incremental comparison
Best practices
Results

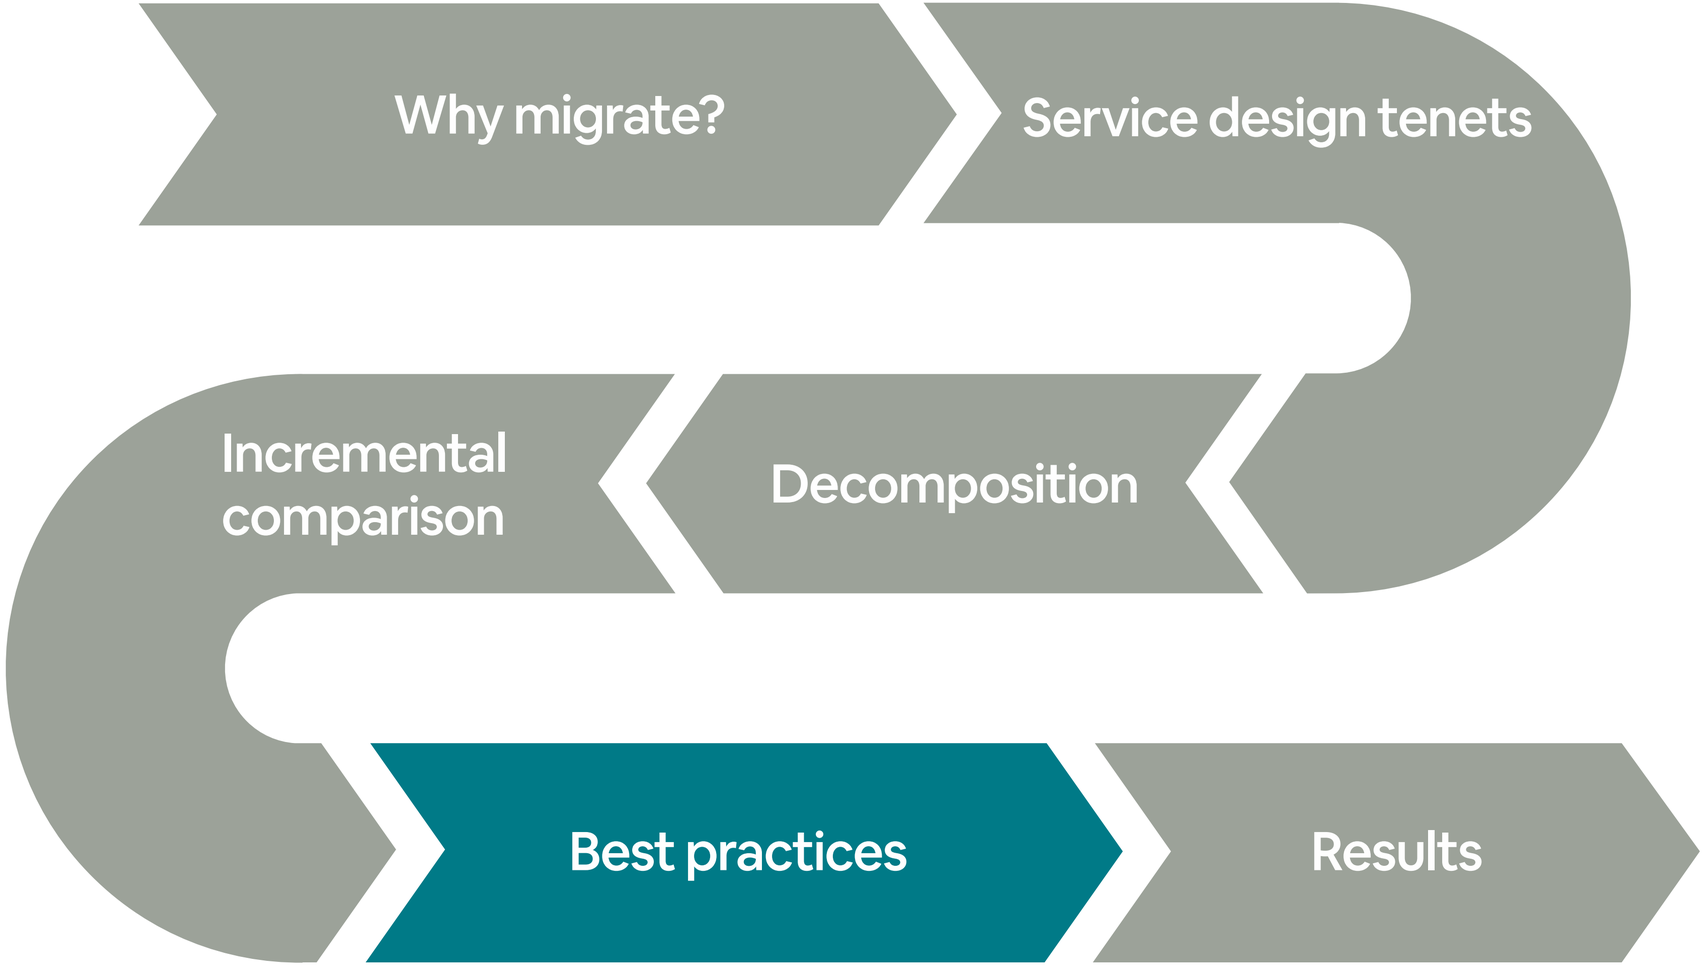
Why migrate?
Service design tenets
Decomposition
Incremental comparison
Best practices
Results

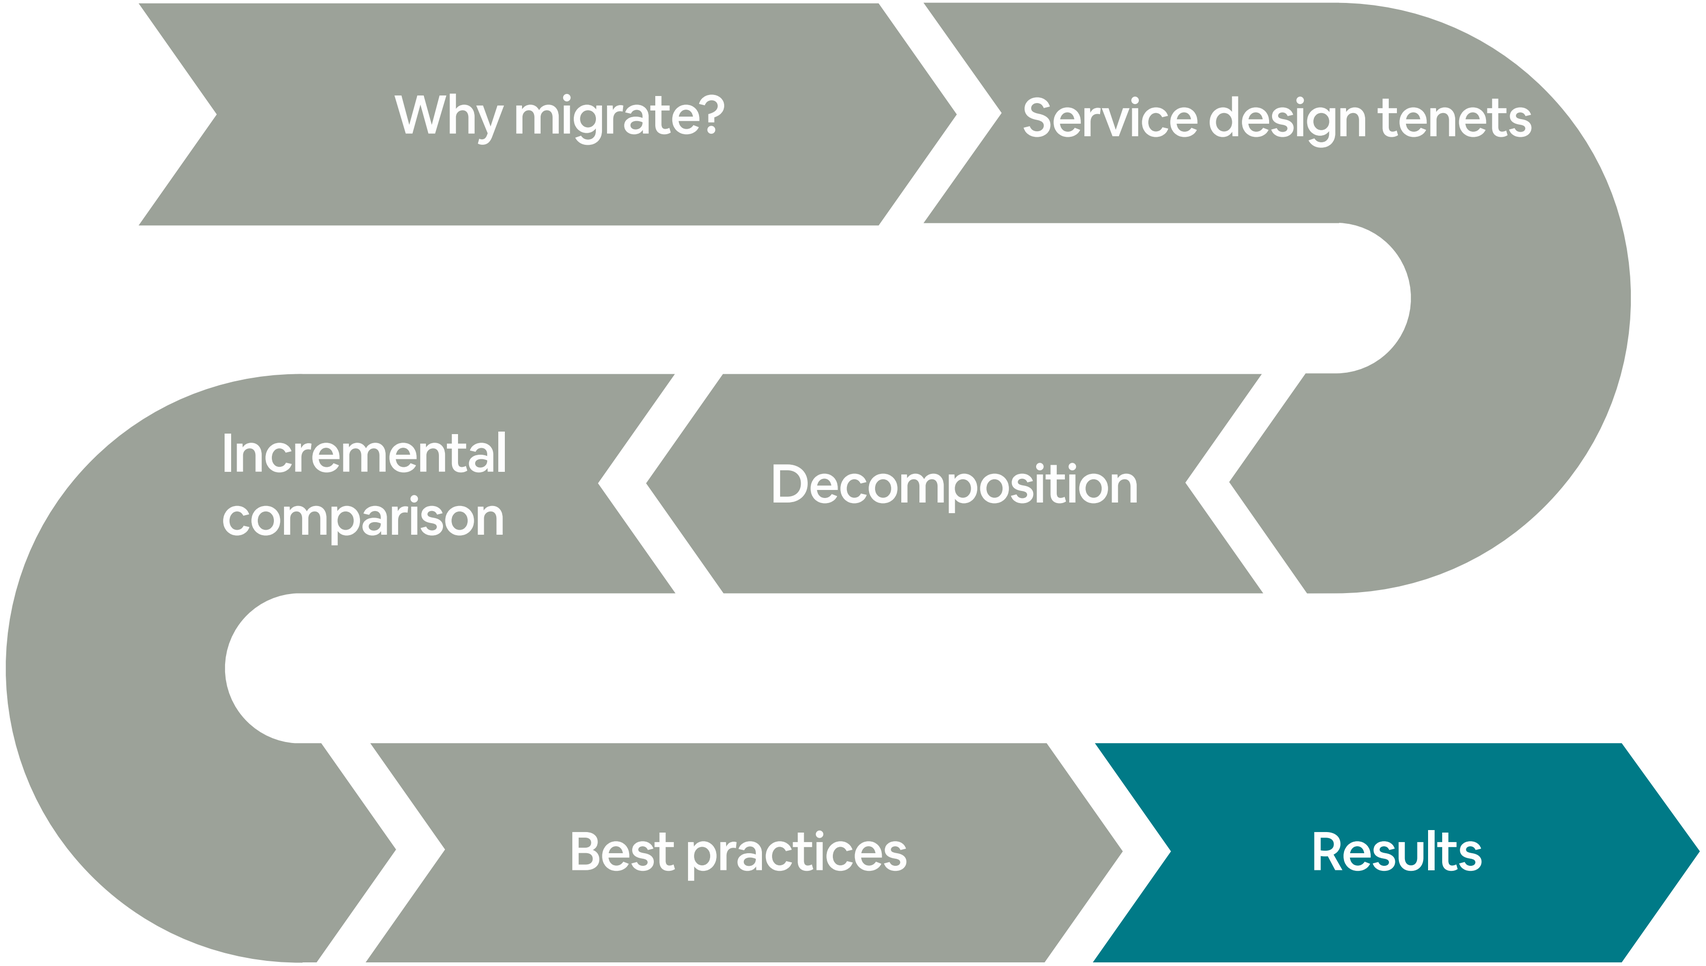
Why migrate?
Service design tenets
Decomposition
Incremental comparison
Best practices
Results

Monorail, our Ruby on Rails monolith

@jessicamtai

Easy start with monoliths

EARLY AIRBNB

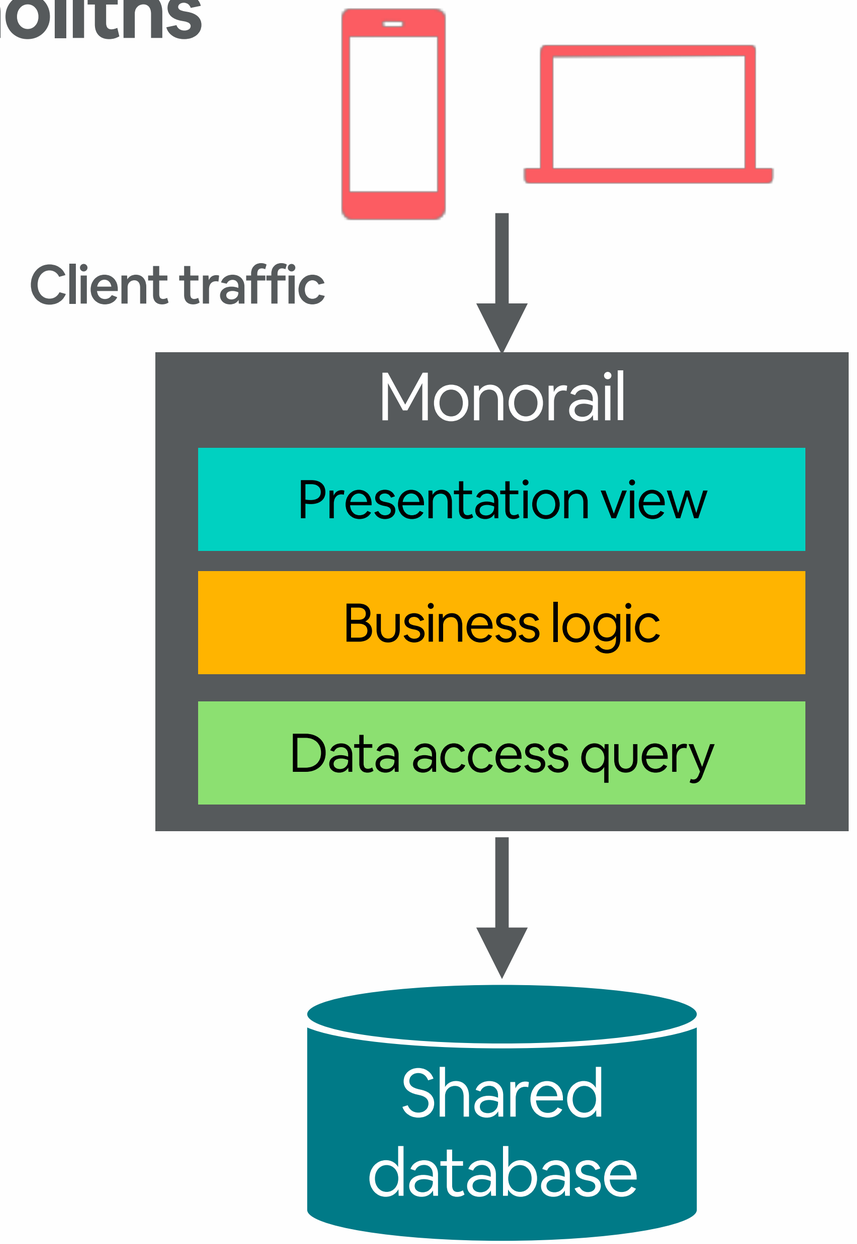

@jessicamtai

DOME HOUSE

Mushroom Dome Cabin: #1 on airbnb in the world

Aptos

Kitty

This home is on people's minds.
It's been viewed 500+ times in the past week.

3 guests **1 bedroom** **2 beds** **1 bath**

Great location
100% of recent guests gave the location a 5-star rating.

Kitty is a Superhost
Superhosts are experienced, highly rated hosts who are committed to providing great stays for guests.

Great check-in experience

$130 per night
★★★★★ 1263

Request to Book

rules > 2. Who's coming? > 3. Confirm and pay

Who's coming?

This is a rare find. Kitty's place is usually booked.

Guests

1 guest

What's the main purpose of this trip?

- Personal travel
- Business travel

Say hello to your host

Let Kitty know a little about yourself and why you're coming.

Hello Kitty! Can't wait to spend 1 night in your home.

Continue

Mushroom
on airbnb in
Entire home/
★★★★★ 12

1 guest
Mar 4, 2

$130.00 x 1
Cleaning fe
Service fee
Occupancy
Gift Credit
Total

Due now (U
Due on Feb

Moderate –
Cancel within
get a full refu

DOME HOUSE

Mushroom Dome Cabin: #1 on airbnb in the world

Aptos

Kitty

This home is on people's minds.
It's been viewed 500+ times in the past week.

3 guests 1 bedroom 2 beds 1 bath

Great location
100% of recent guests gave the location a 5-star rating.

Kitty is a Superhost
Superhosts are experienced, highly rated hosts who are committed to providing great stays for guests.

Great check-in experience

$130 per night
★★★★★ 1263

Request to Book

This is a rare find. Kitty's place is usually booked.

Guests

1 guest

What's the main purpose of this trip?

- Personal travel
- Business travel

Say hello to your host

Let Kitty know a little about yourself and why you're coming.

Hello Kitty! Can't wait to spend 1 night in your home.

Continue

@jessicamtai


WHY DECIDE TO MIGRATE?
@jessicamtai


2
Europe,
Middle East
& Africa
LONDON
@jessicamtai

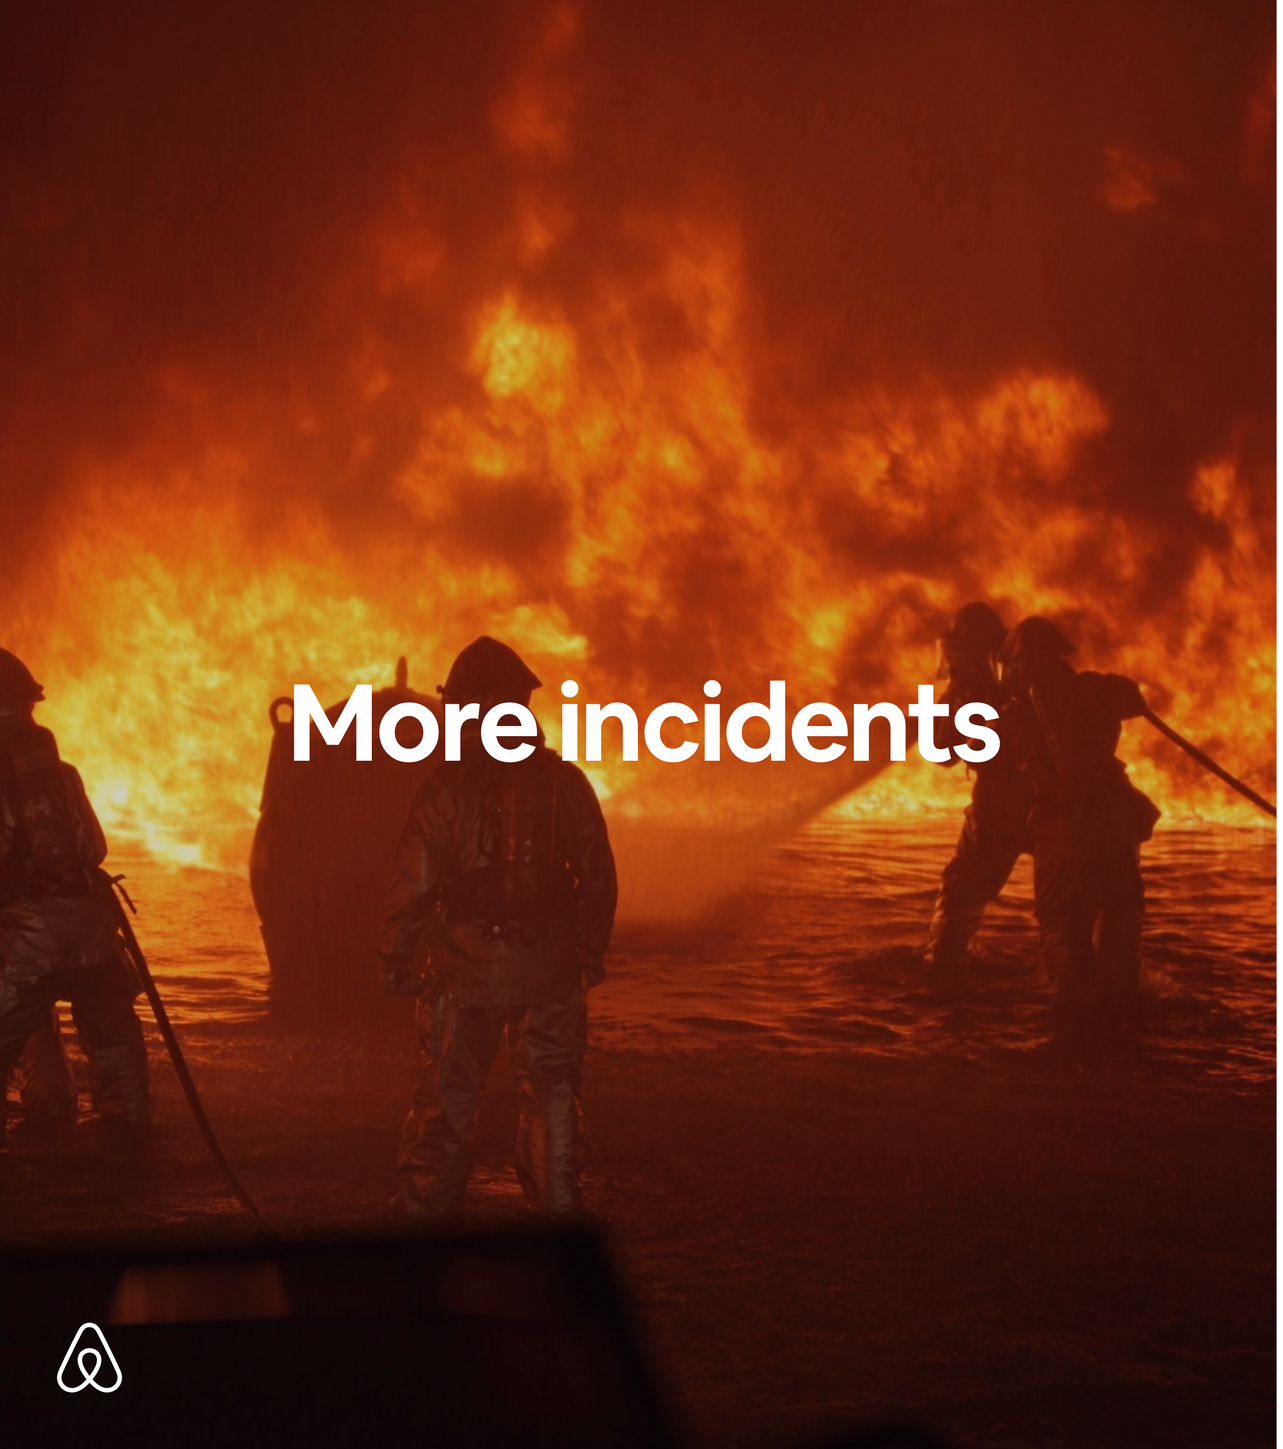
More incidents

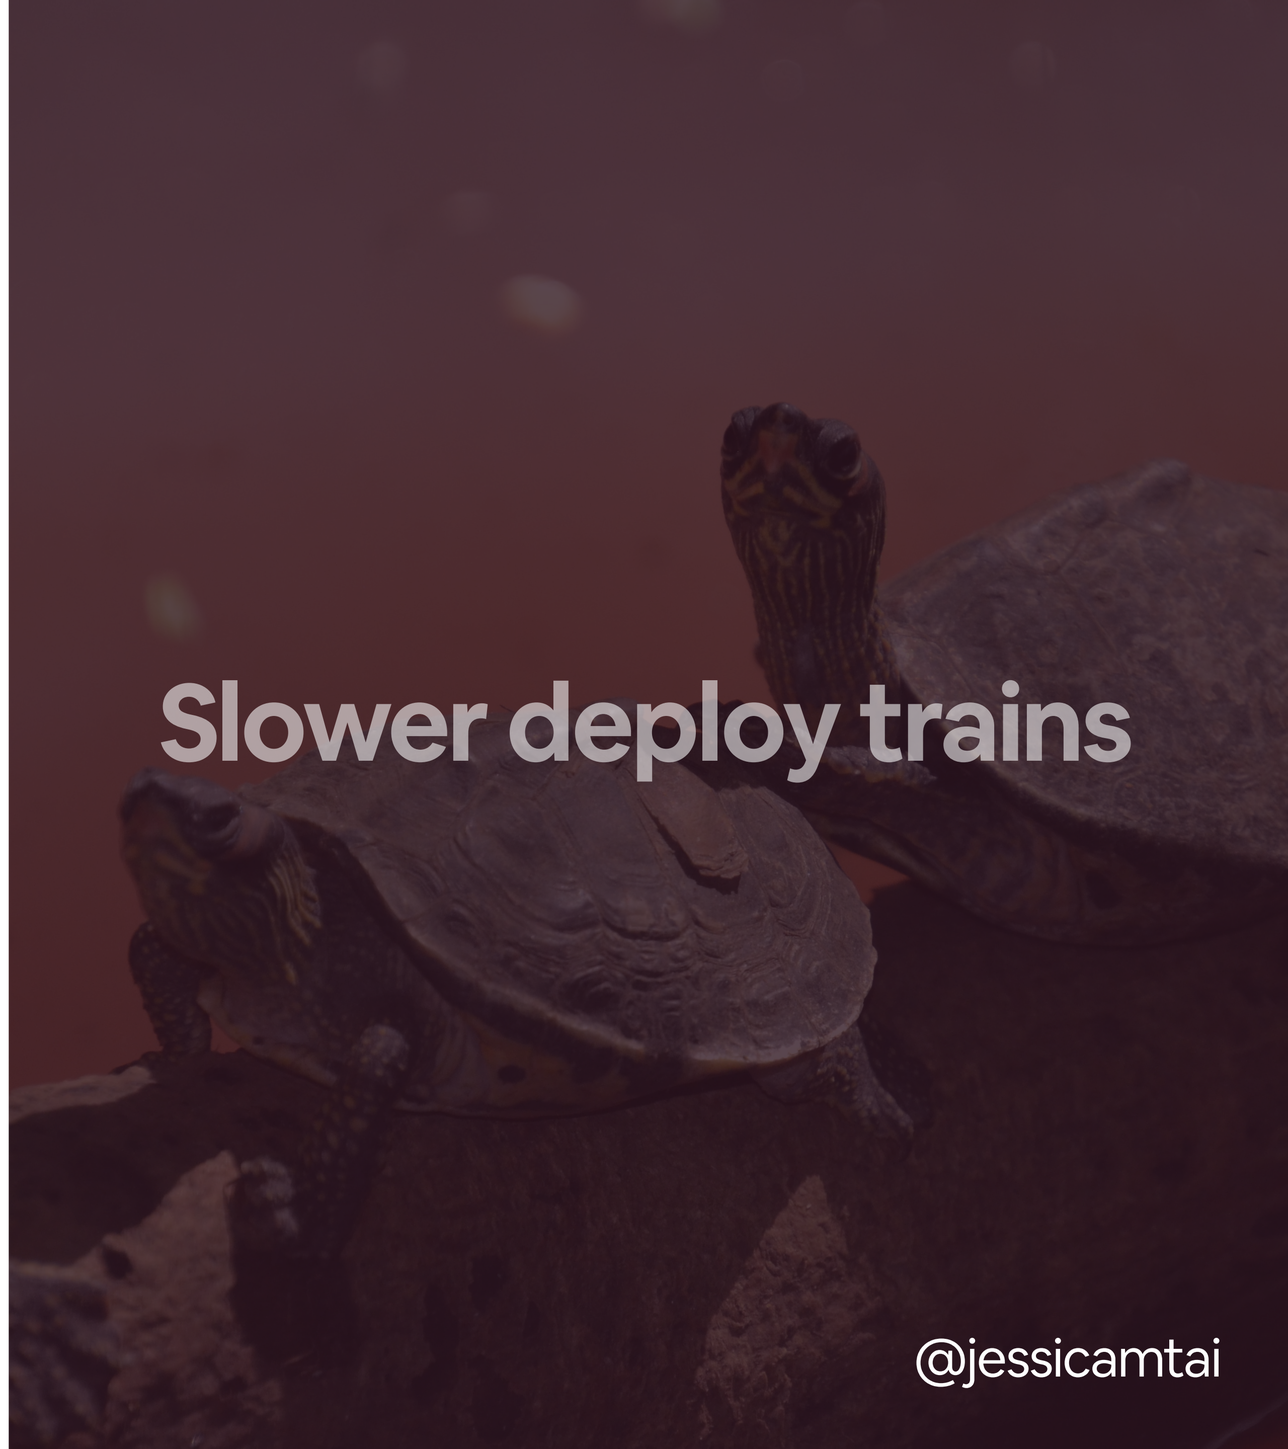
Slower deploy trains
@jessicamtai

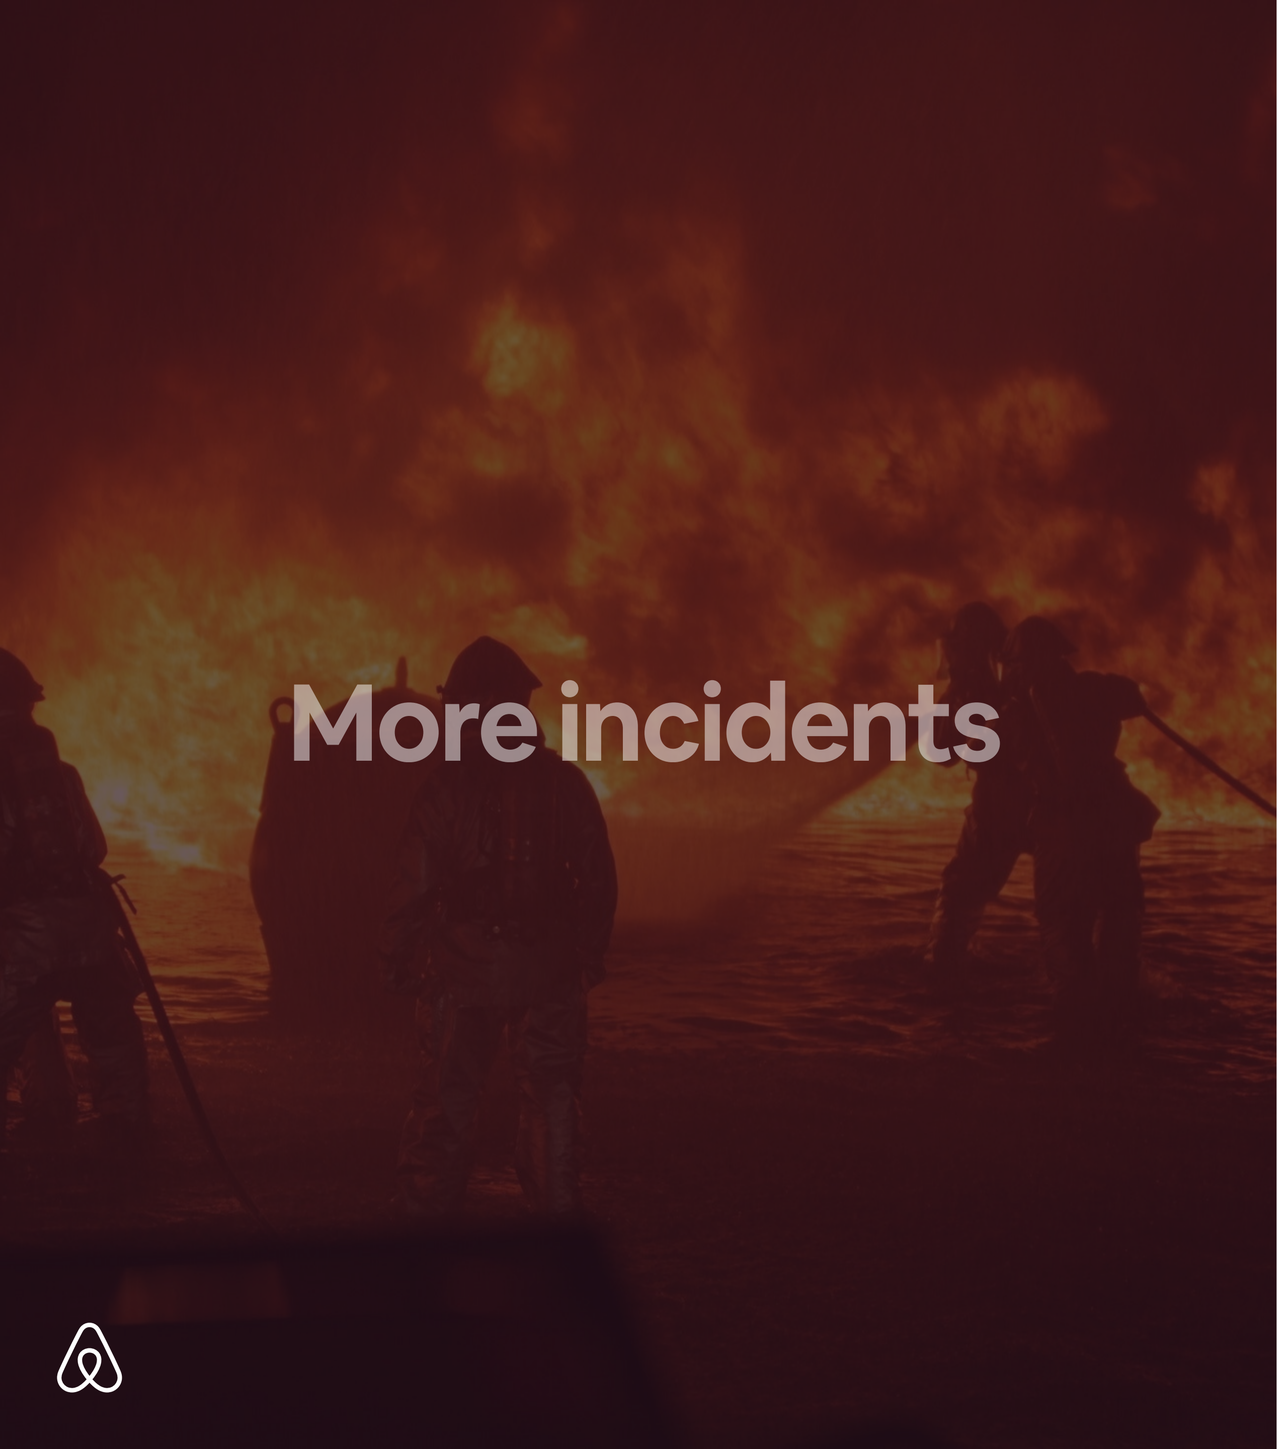
More incidents

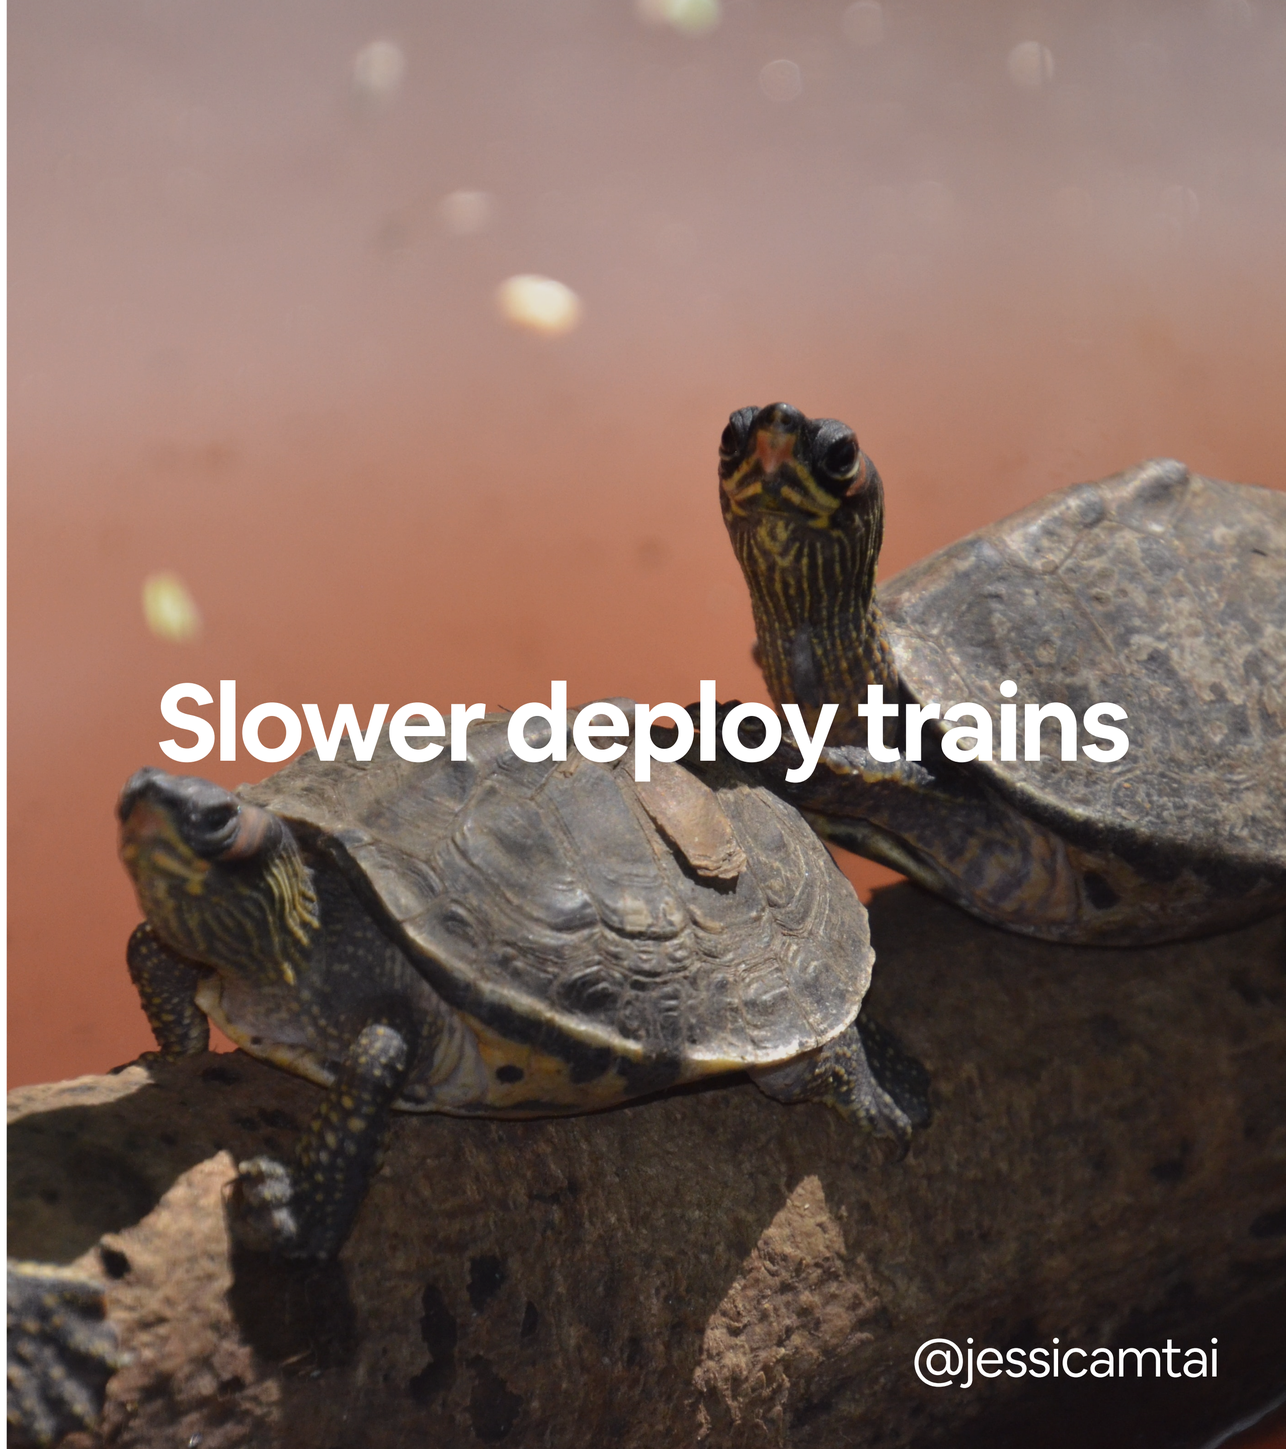
Slower deploy trains
@jessicamtai

Our solution: Service-oriented architecture (SOA)

NETWORK OF LOOSELY-COUPLED SERVICES

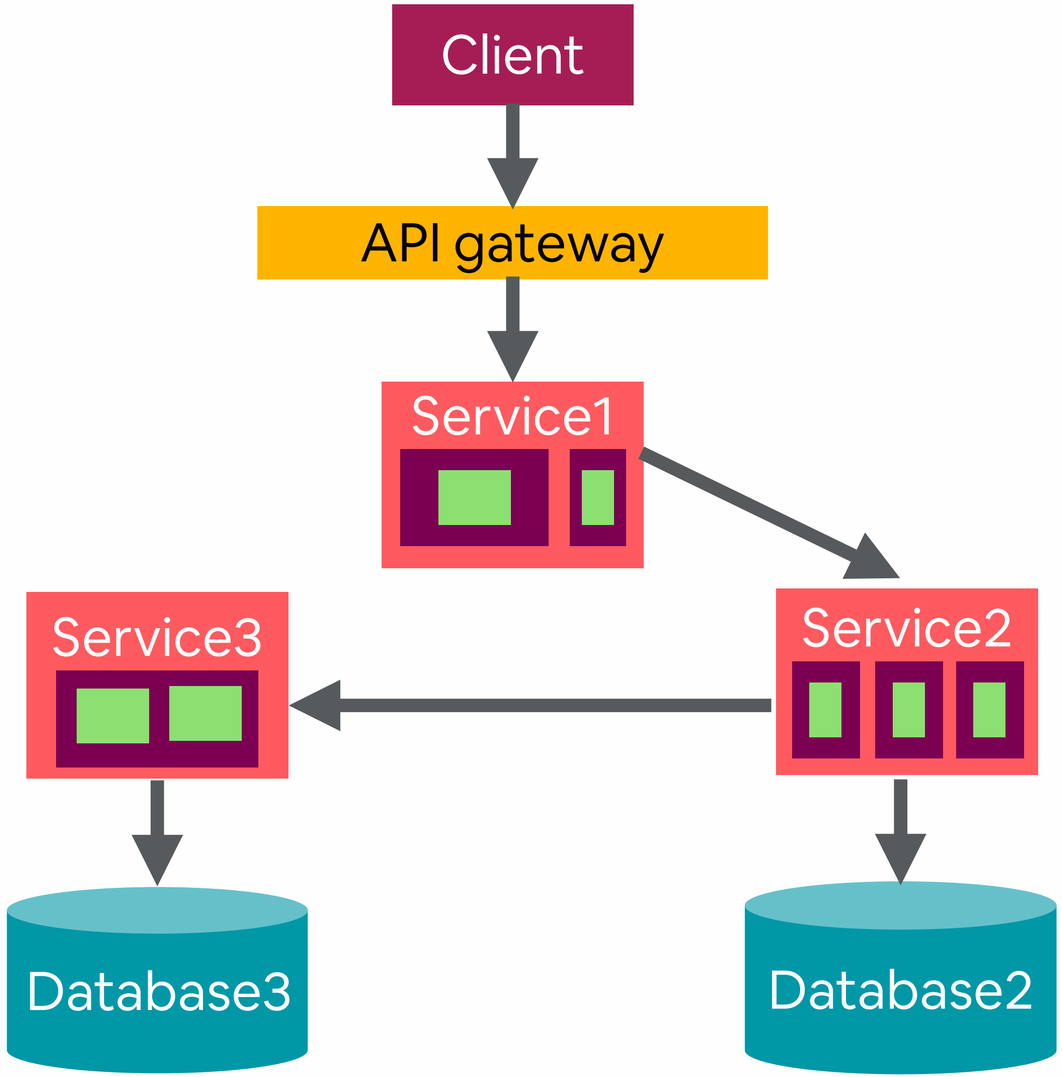

Checkout page in SOA

Home demand service

Reservation service

Business travel service

Messaging service

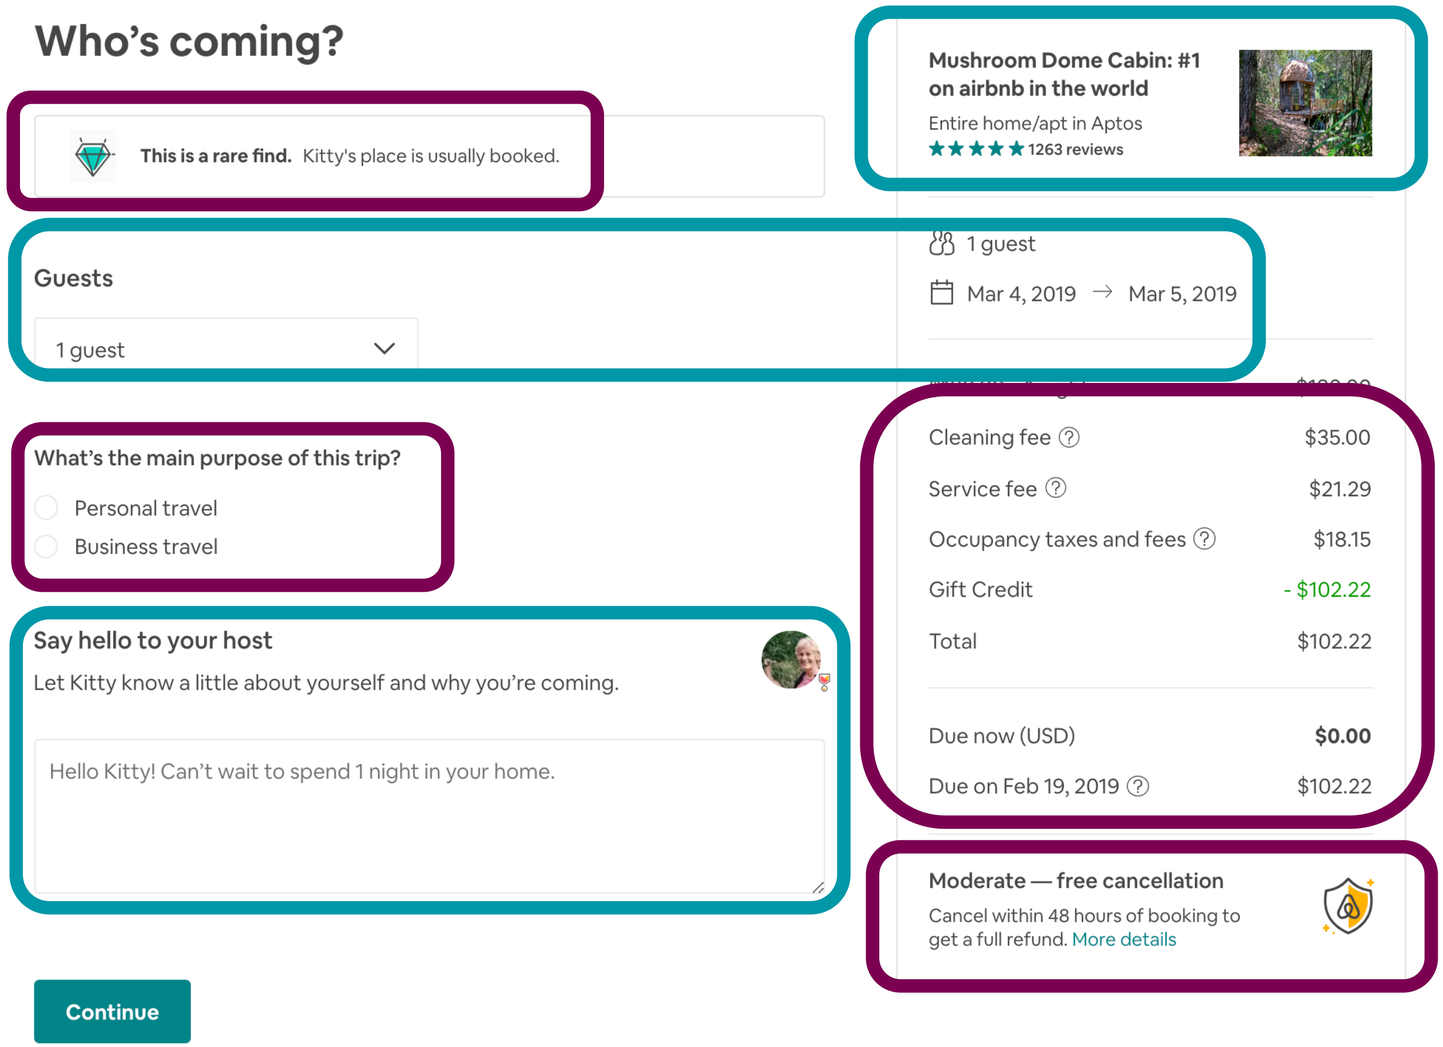

Home service

Pricing service

Cancellation service

@jessicamtai


SERVICE DESIGN TENETS
@jessicamtai

Services own reads & writes to their data

@jessicamtai

Services address a specific concern

Avoid duplicate functionality

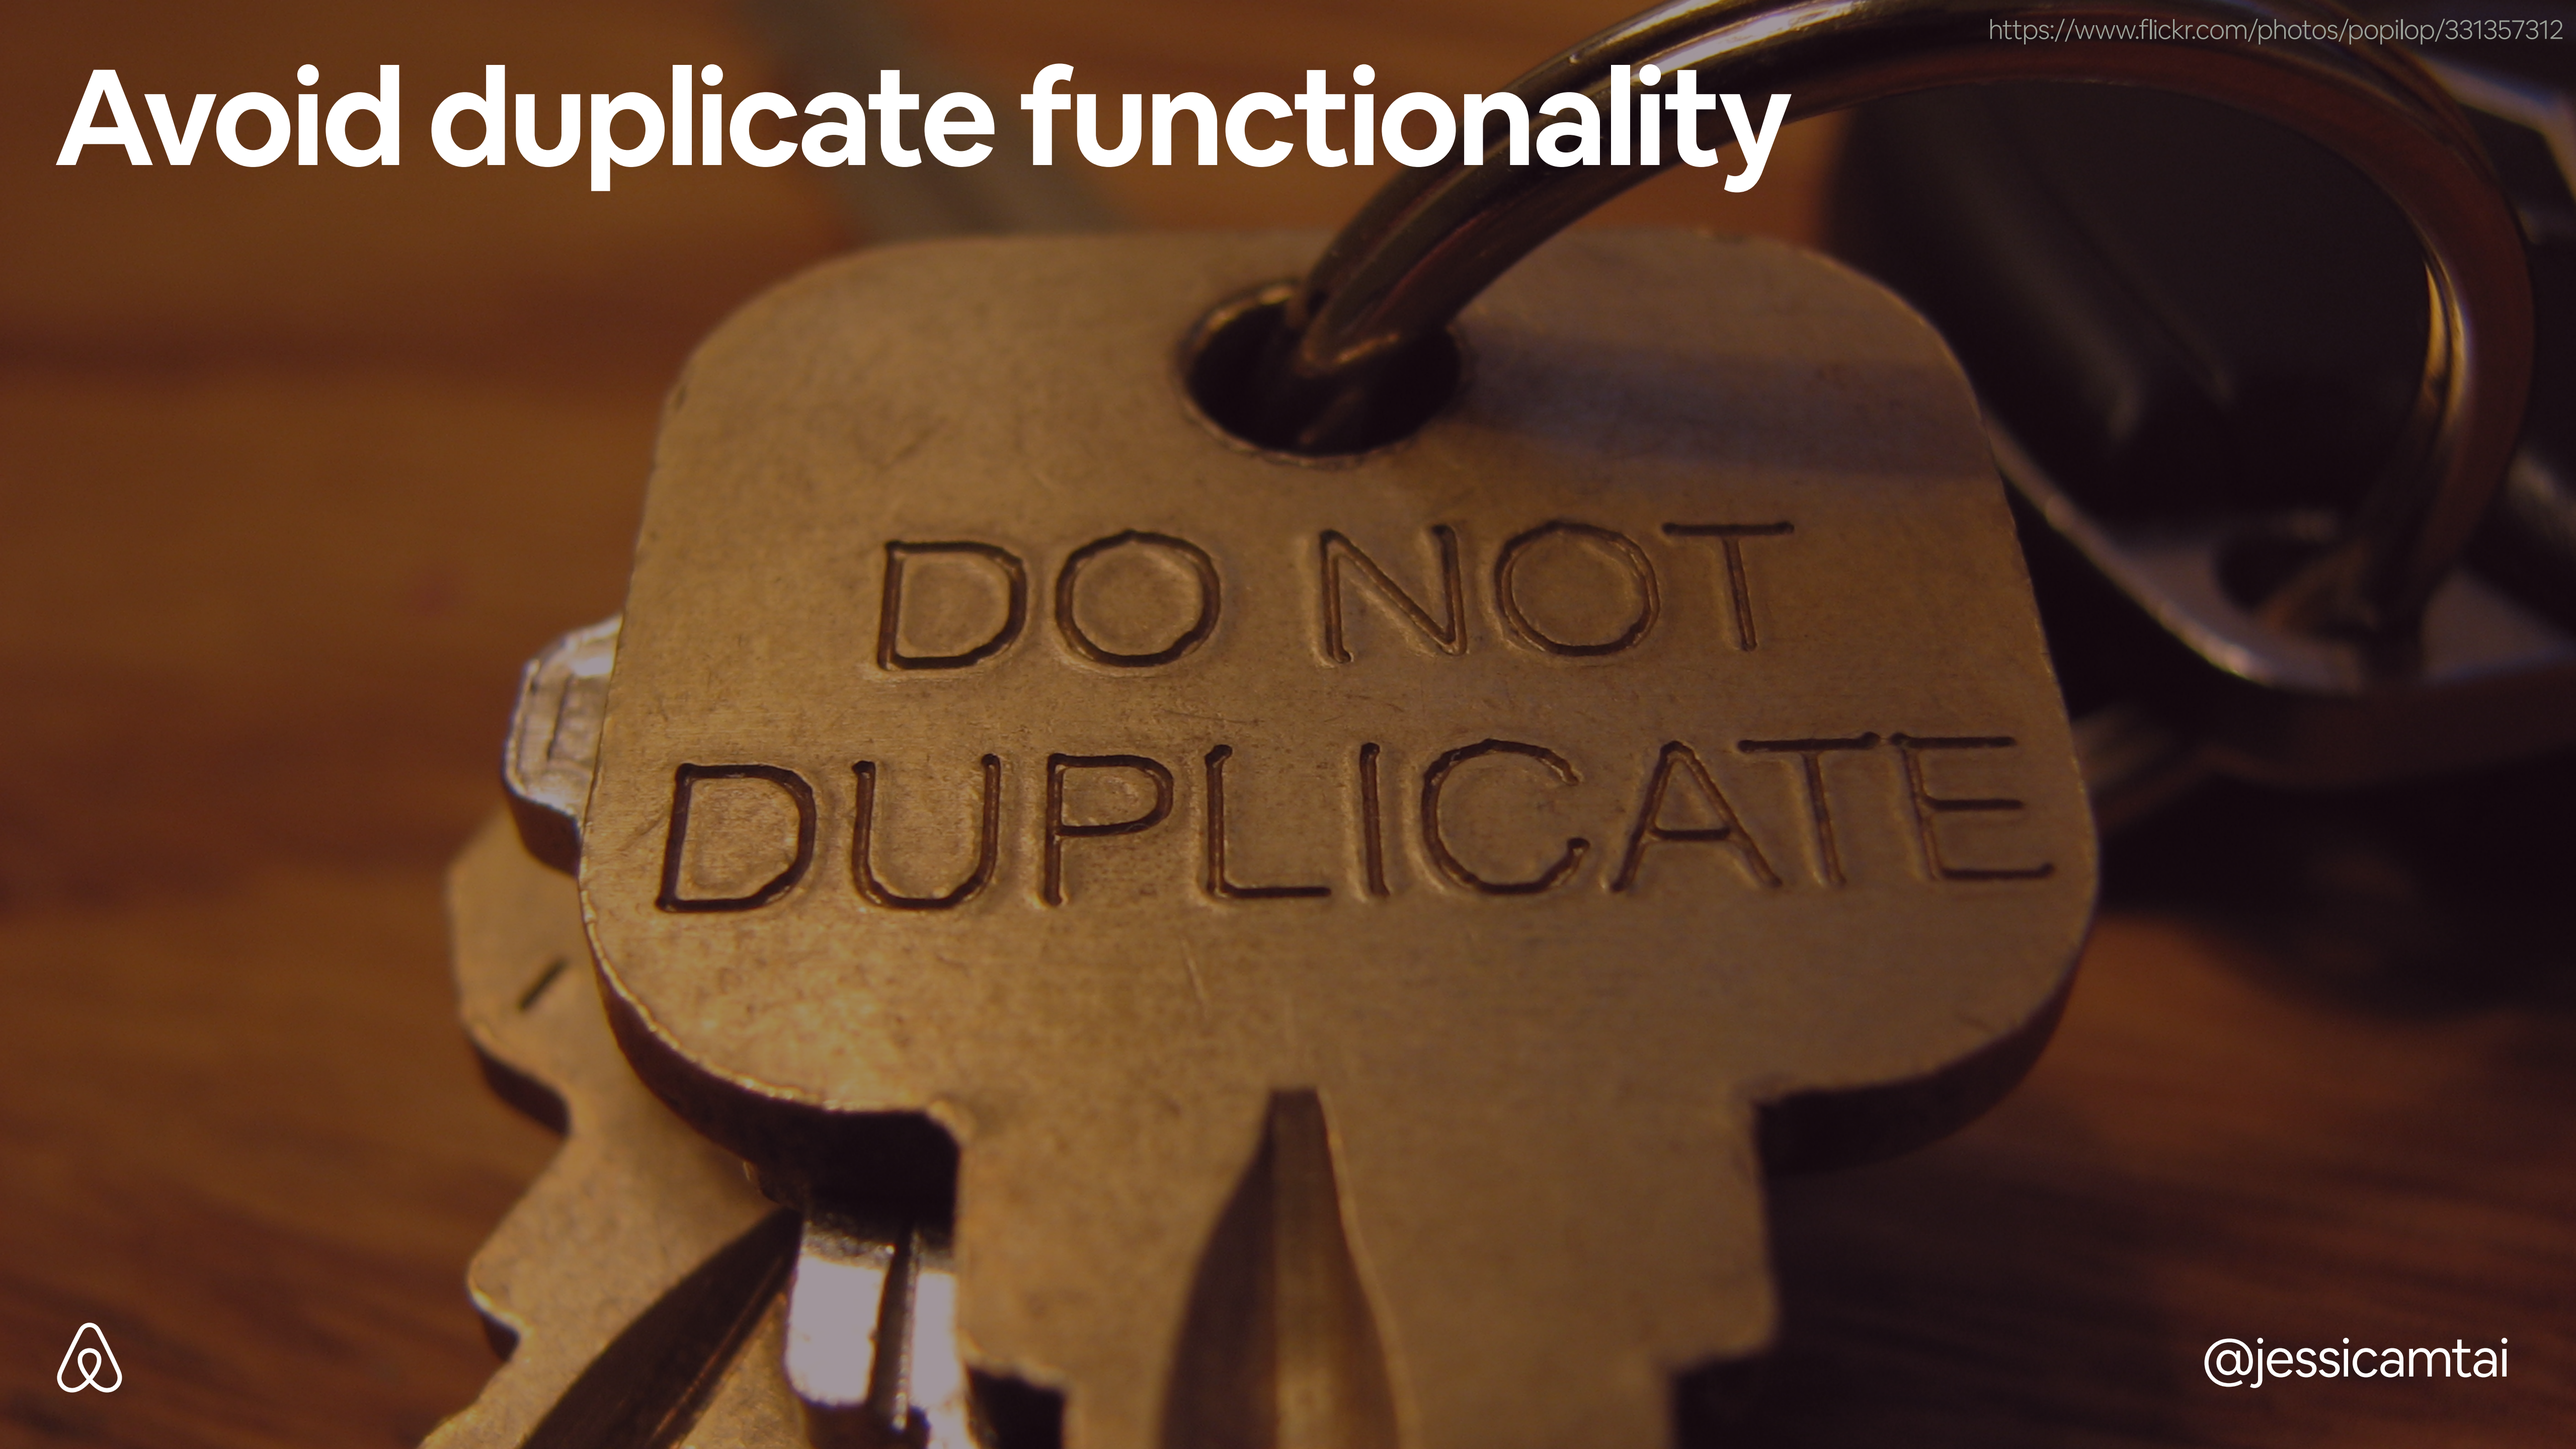

https://www.flickr.com/photos/popilop/331357312

@jessicamtai

Data mutations propagate via standard events

@jessicamtai


Build for production


DECOMPOSE BY
REQUEST LIFE CYCLE
@jessicamtai

Request life cycle

V1: MONORAIL

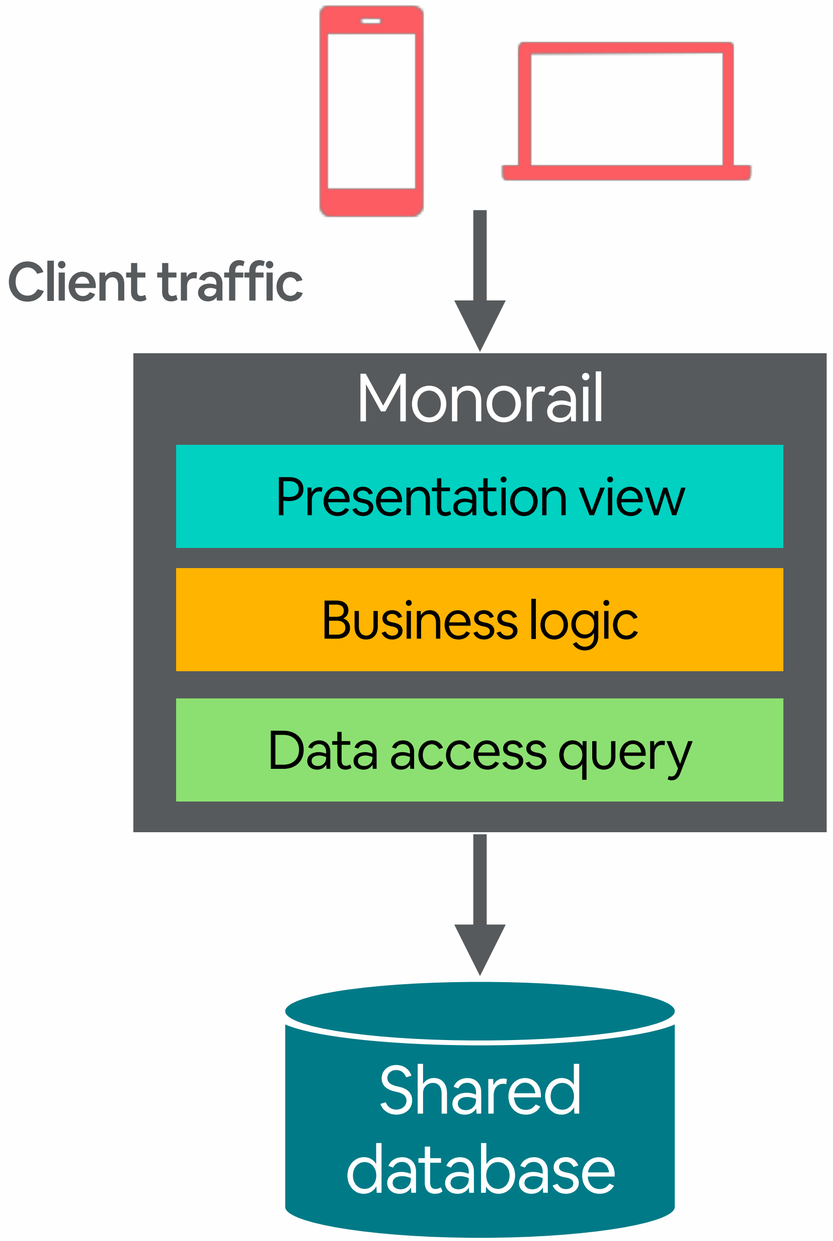

Request life cycle

V2: MONORAIL & SERVICES

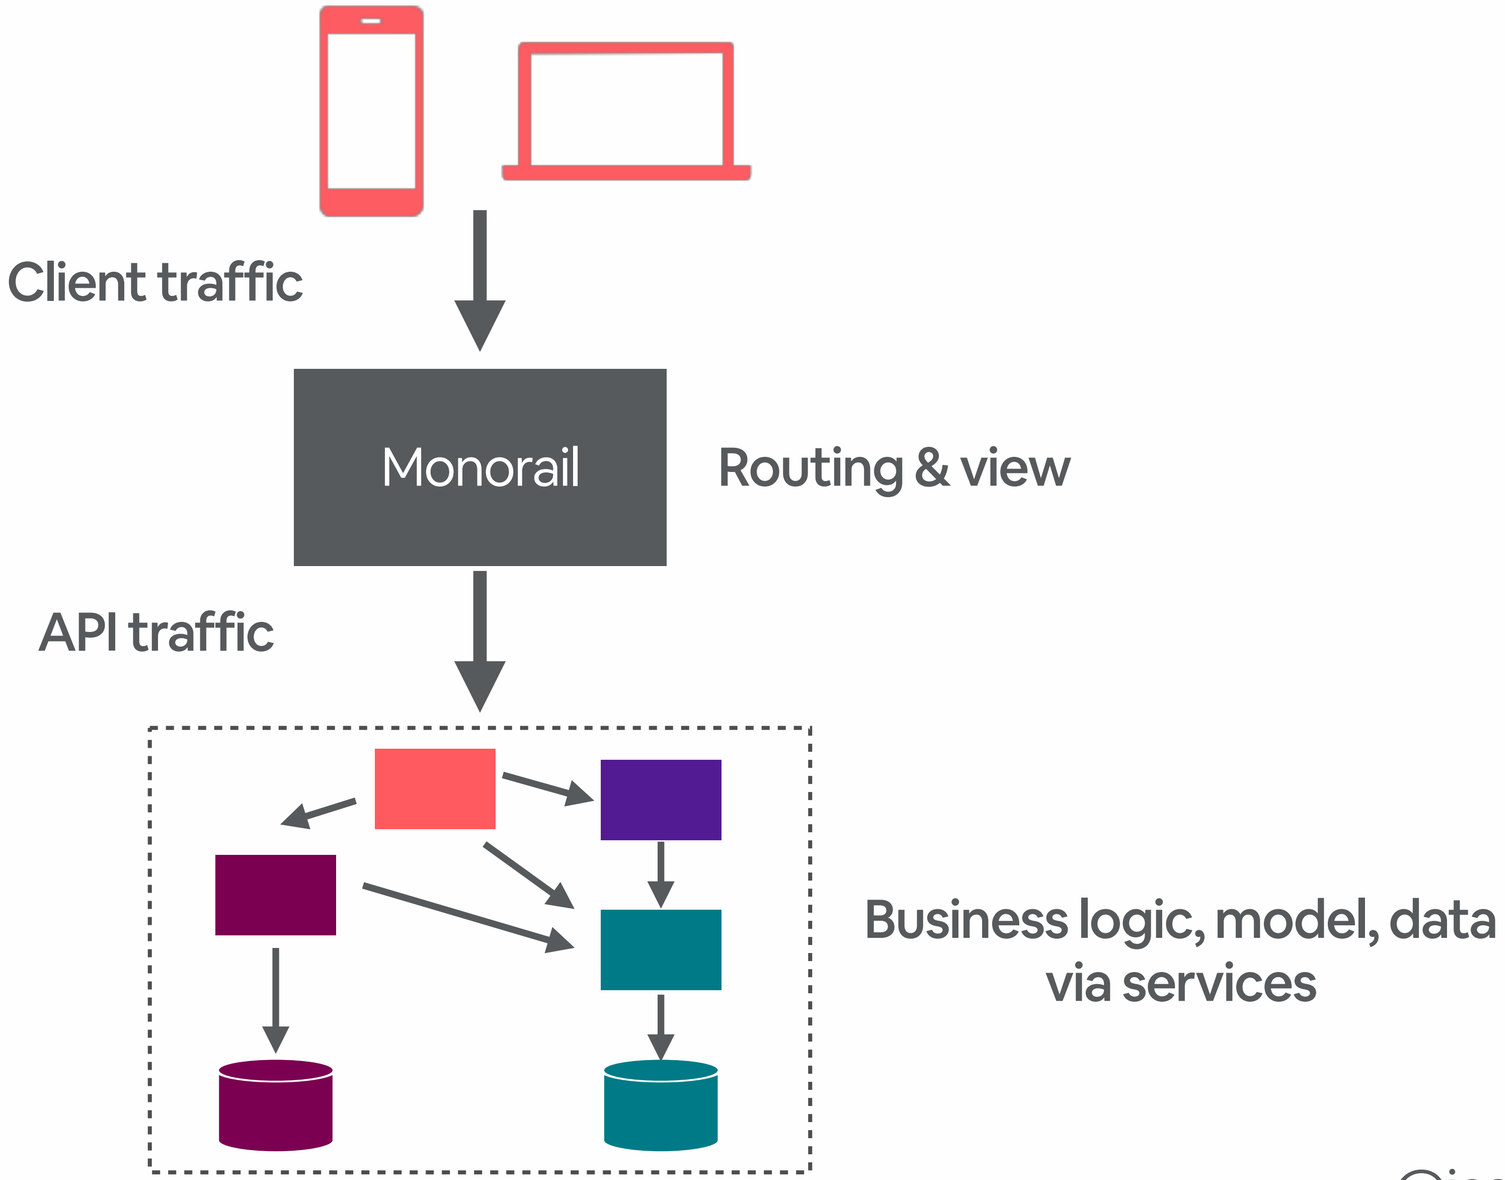

Service types

STRICT FLOW OF DEPENDENCIES

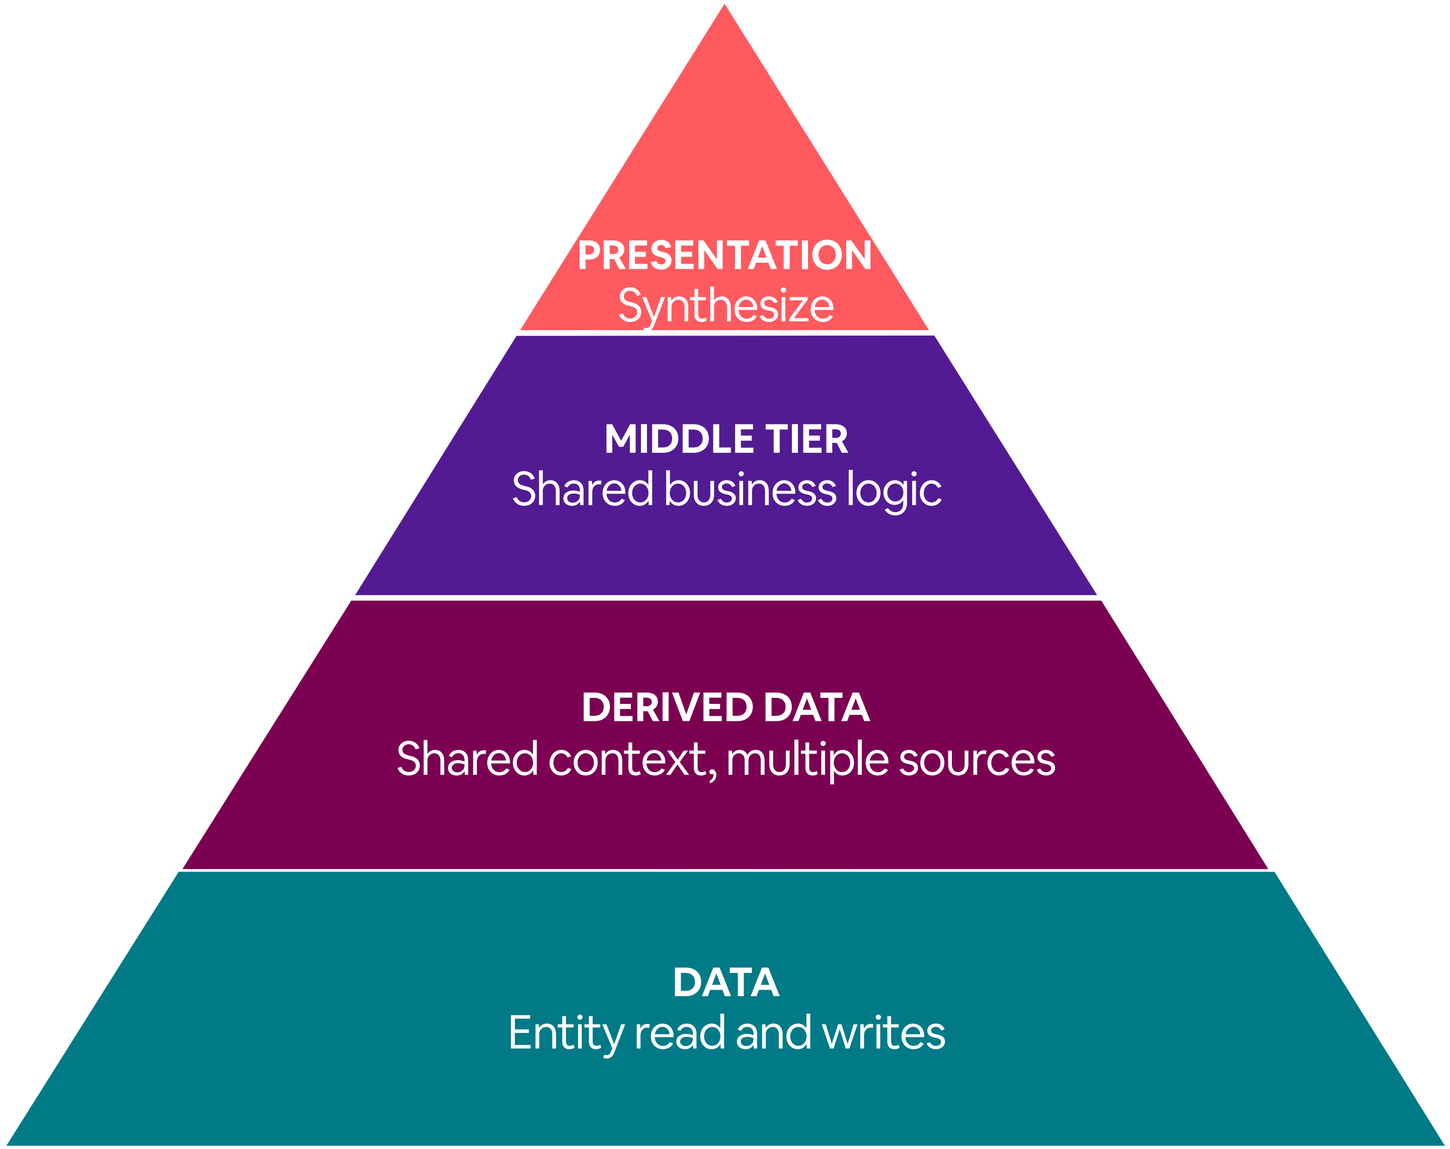

Service types

STRICT FLOW OF DEPENDENCIES

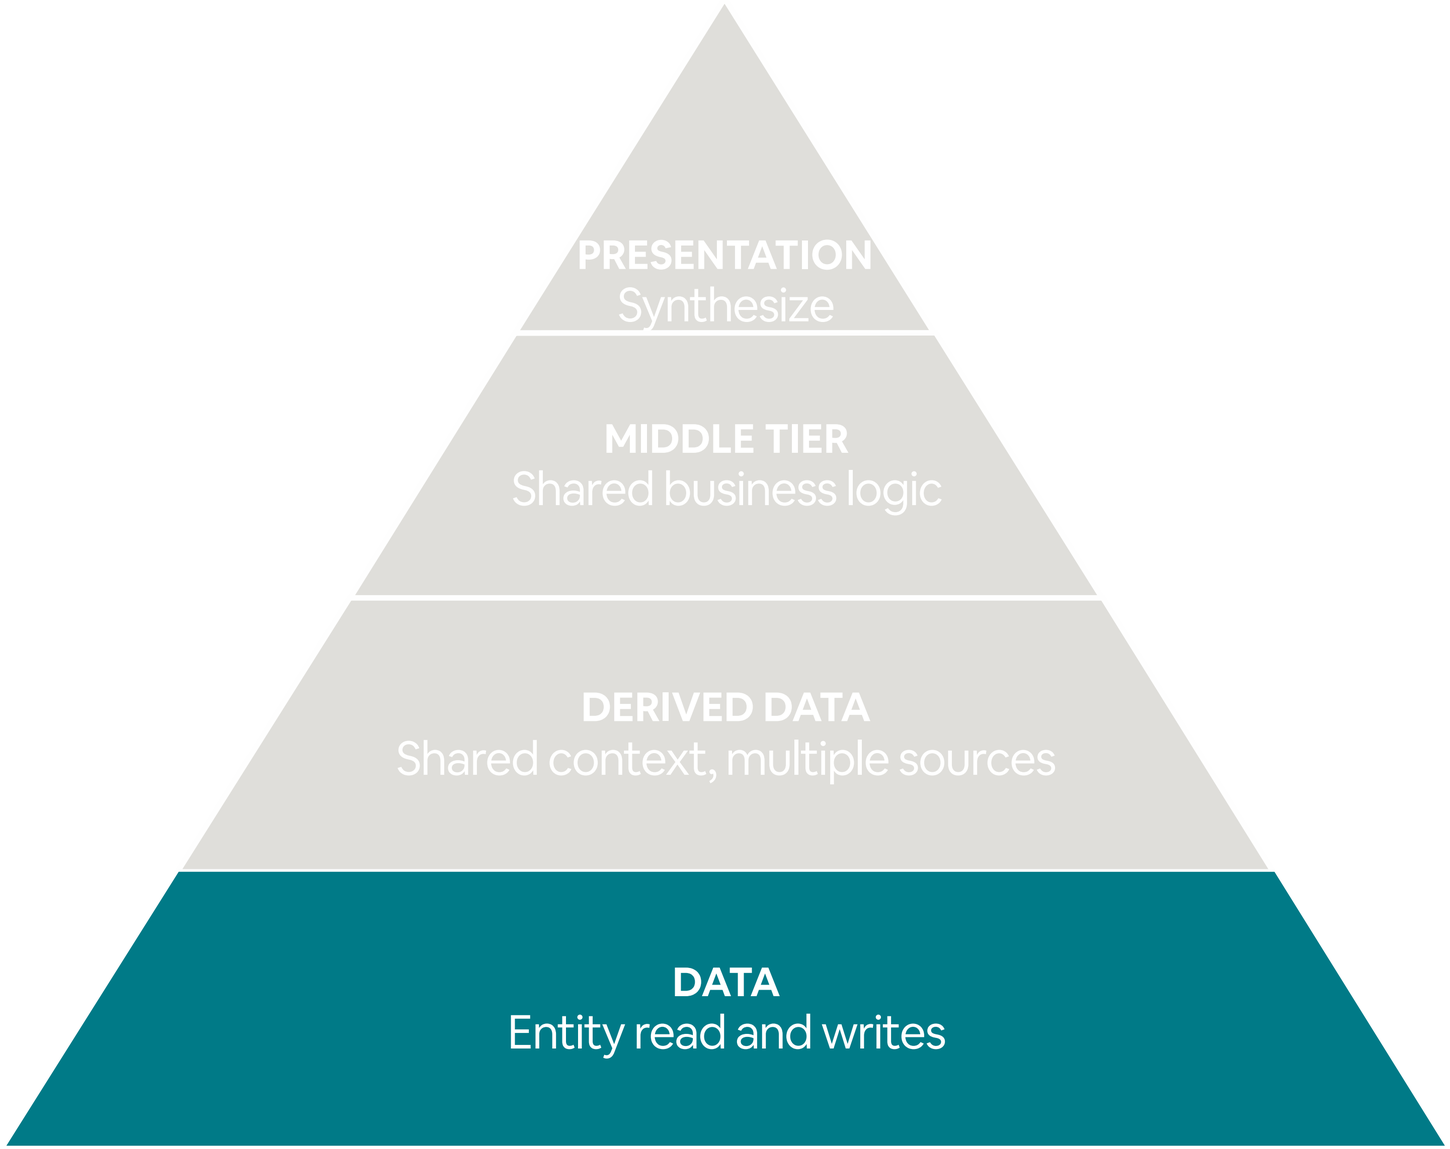

Service types

STRICT FLOW OF DEPENDENCIES

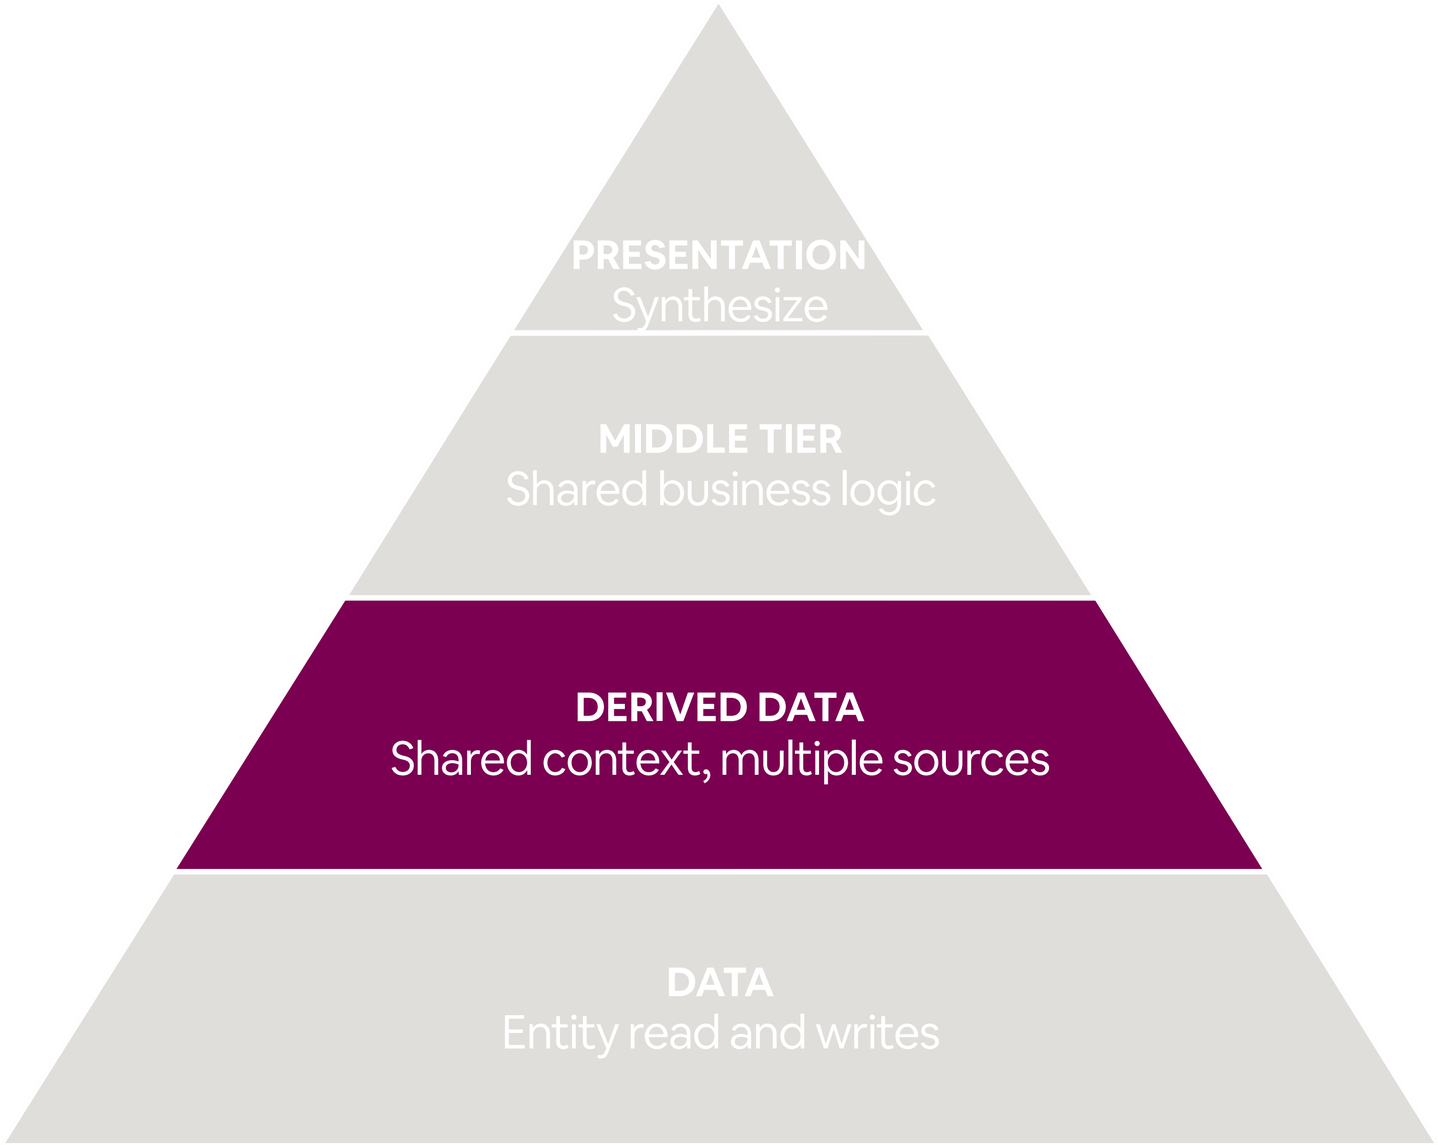

Service types

STRICT FLOW OF DEPENDENCIES

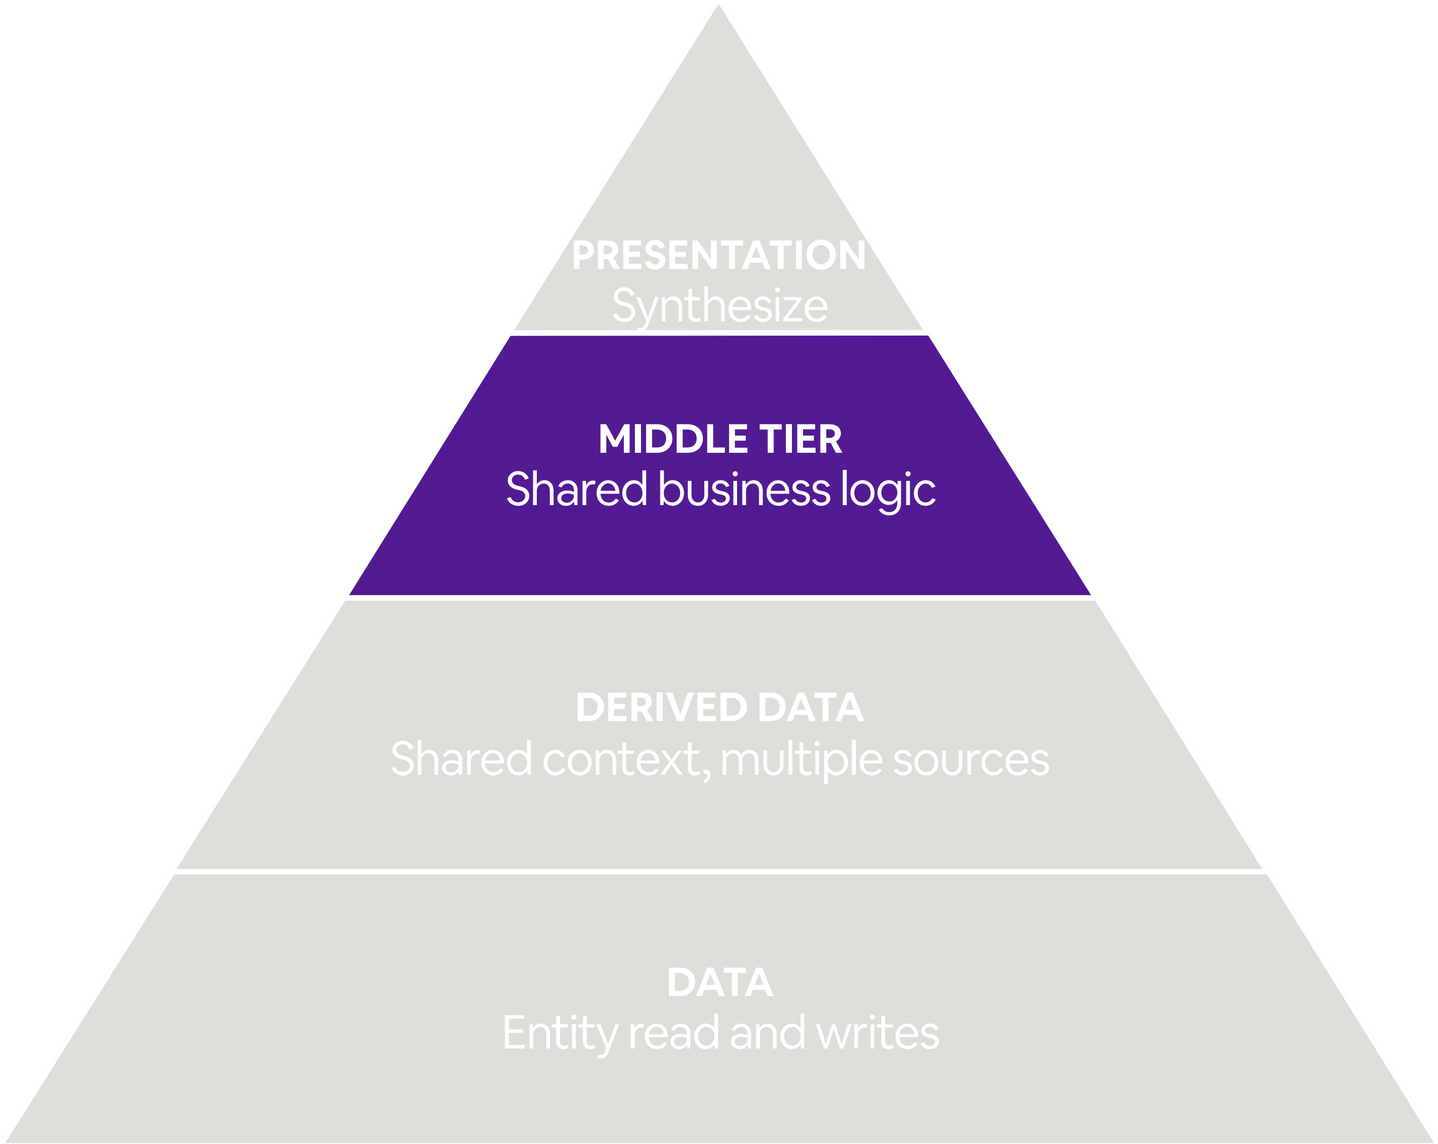

Service types

STRICT FLOW OF DEPENDENCIES

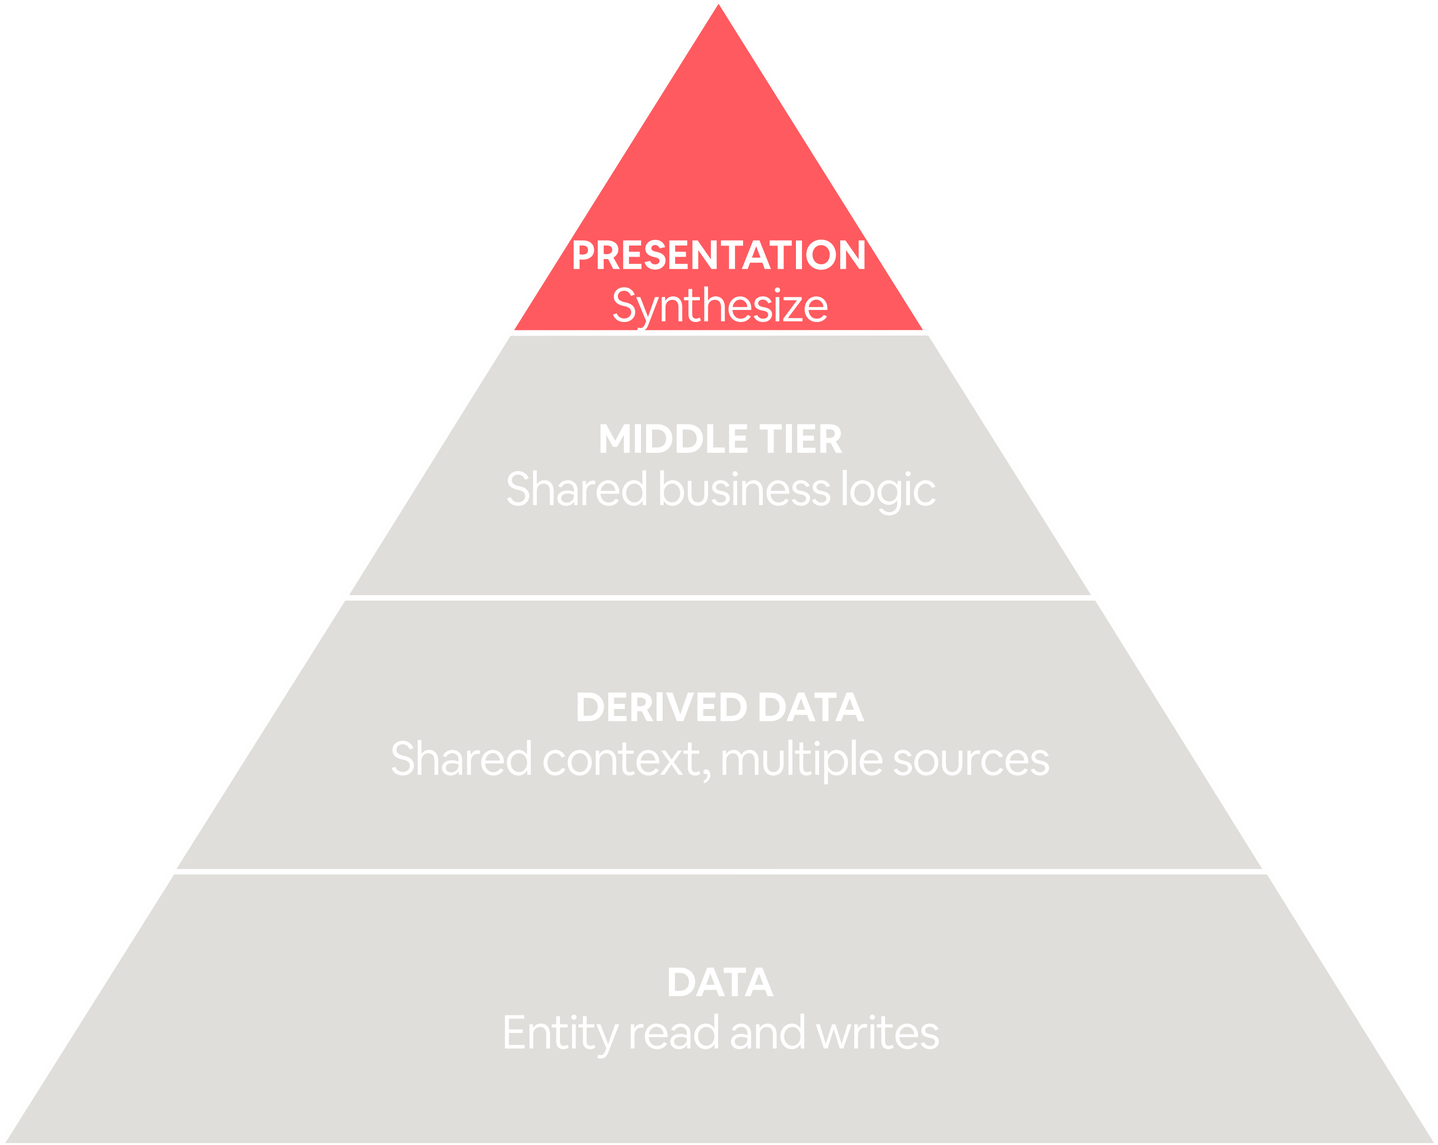

Homes

Mushroom Dome

HOSTED BY KITTY · APTOS, CALIFORNIA

@jessicamtai

1. Migrating core data models

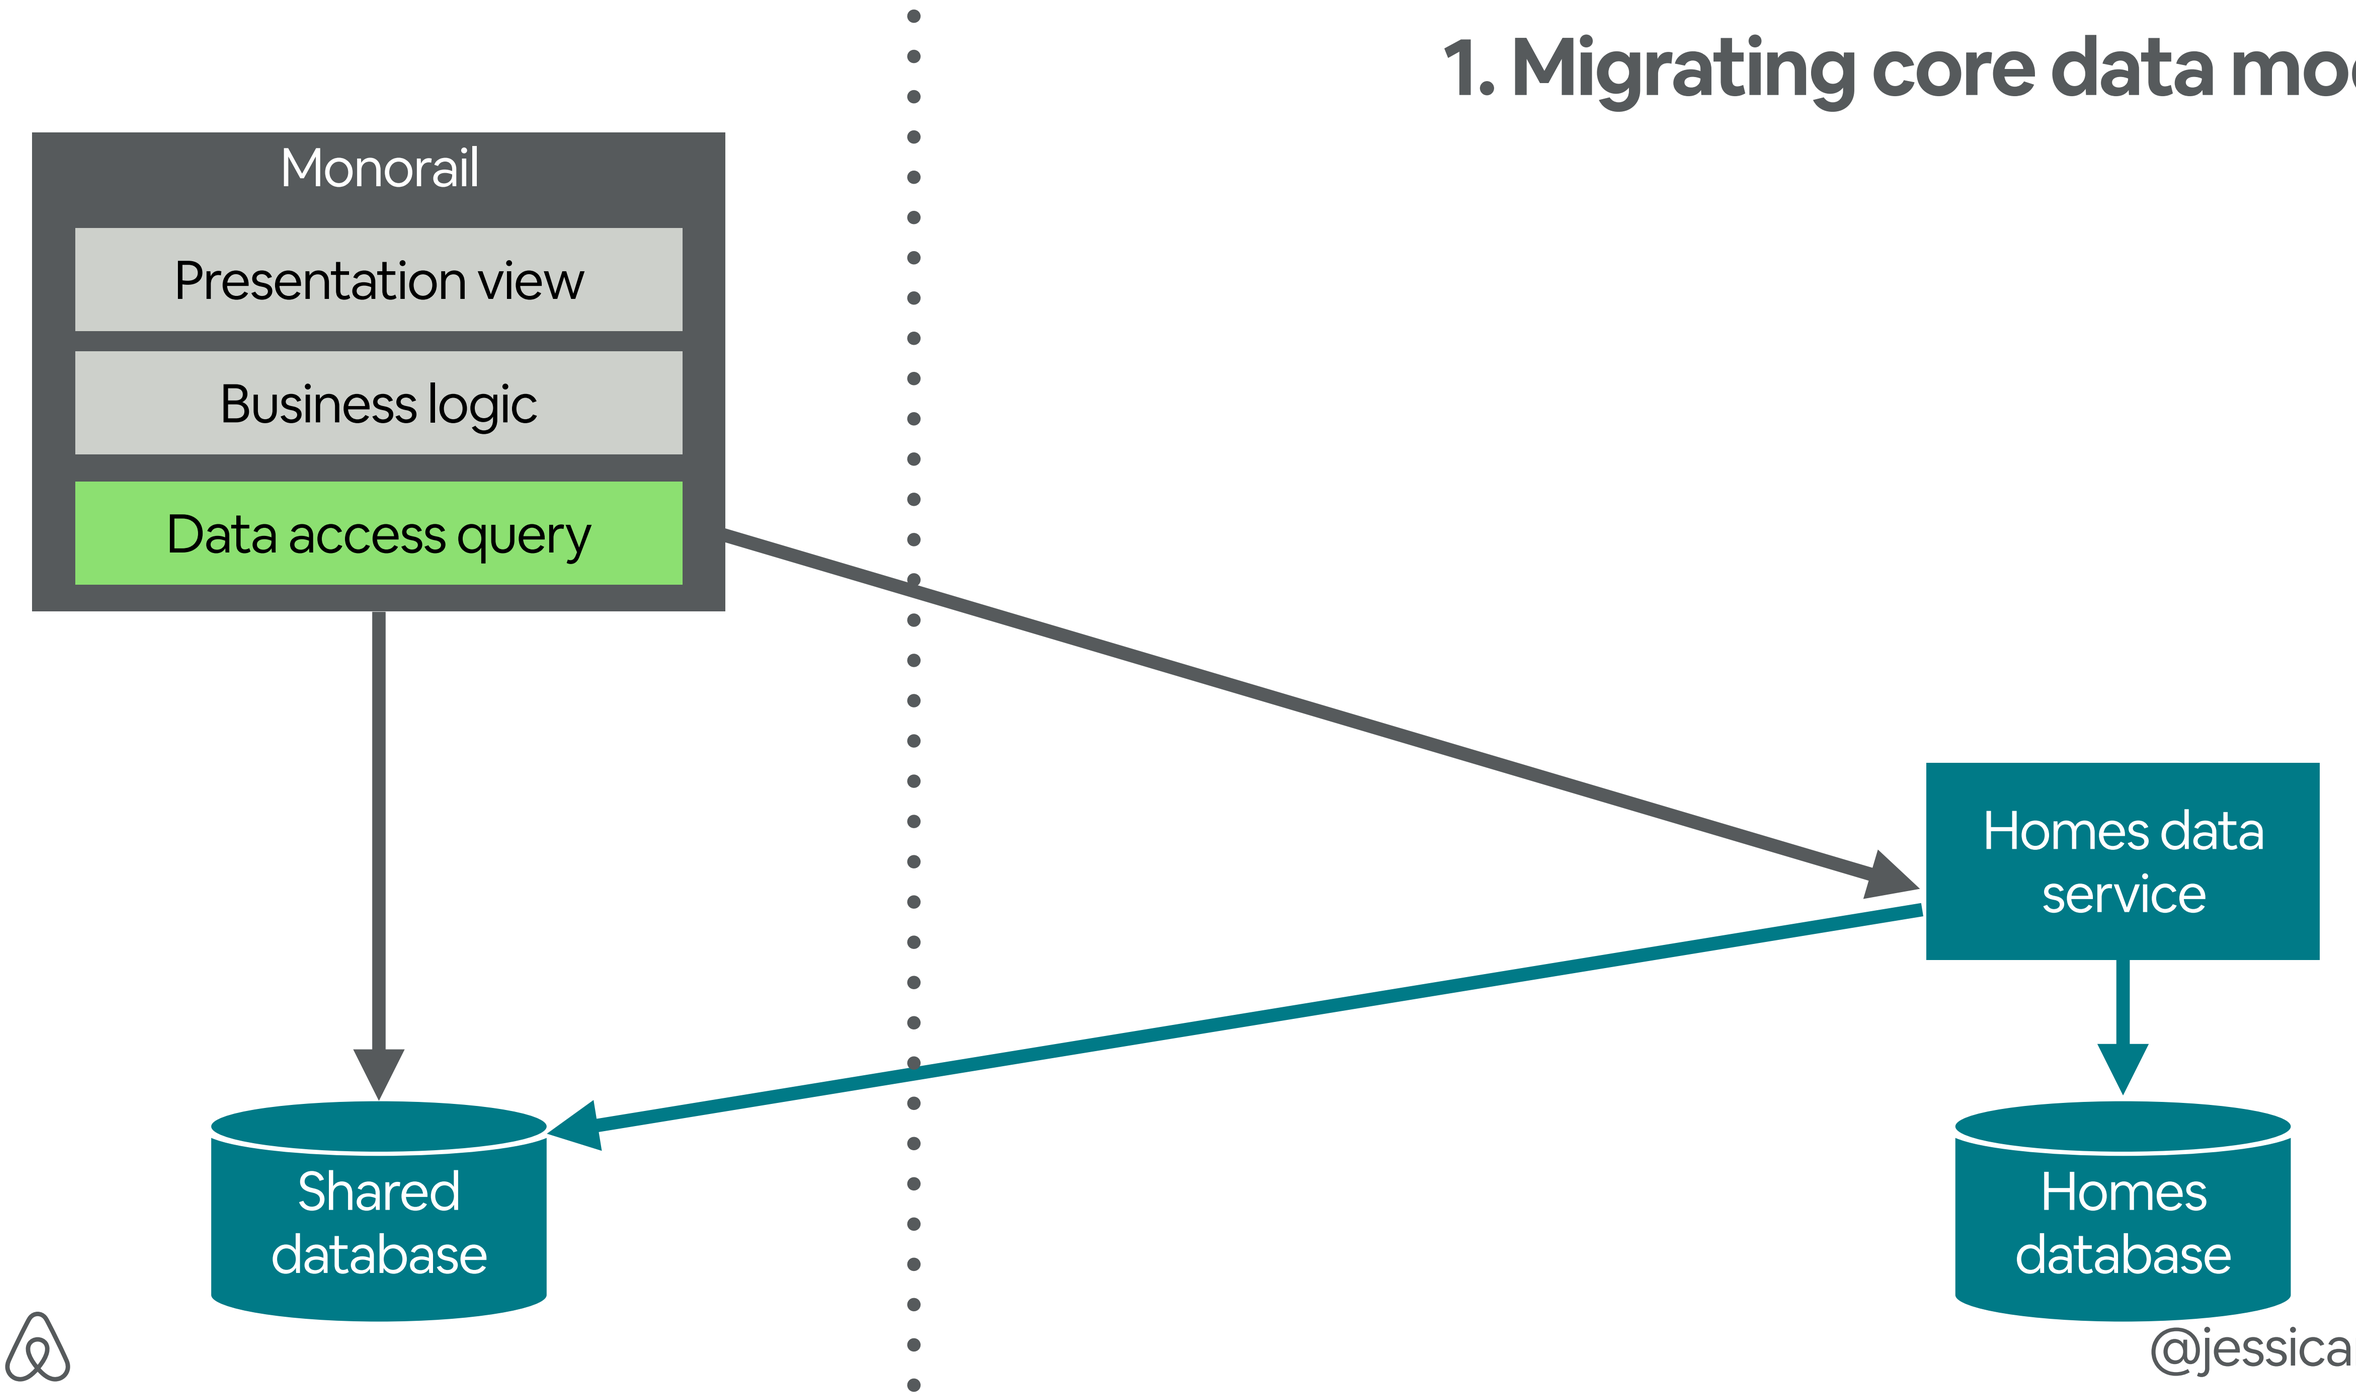

@jessicamtai

2. Migrating core business logic

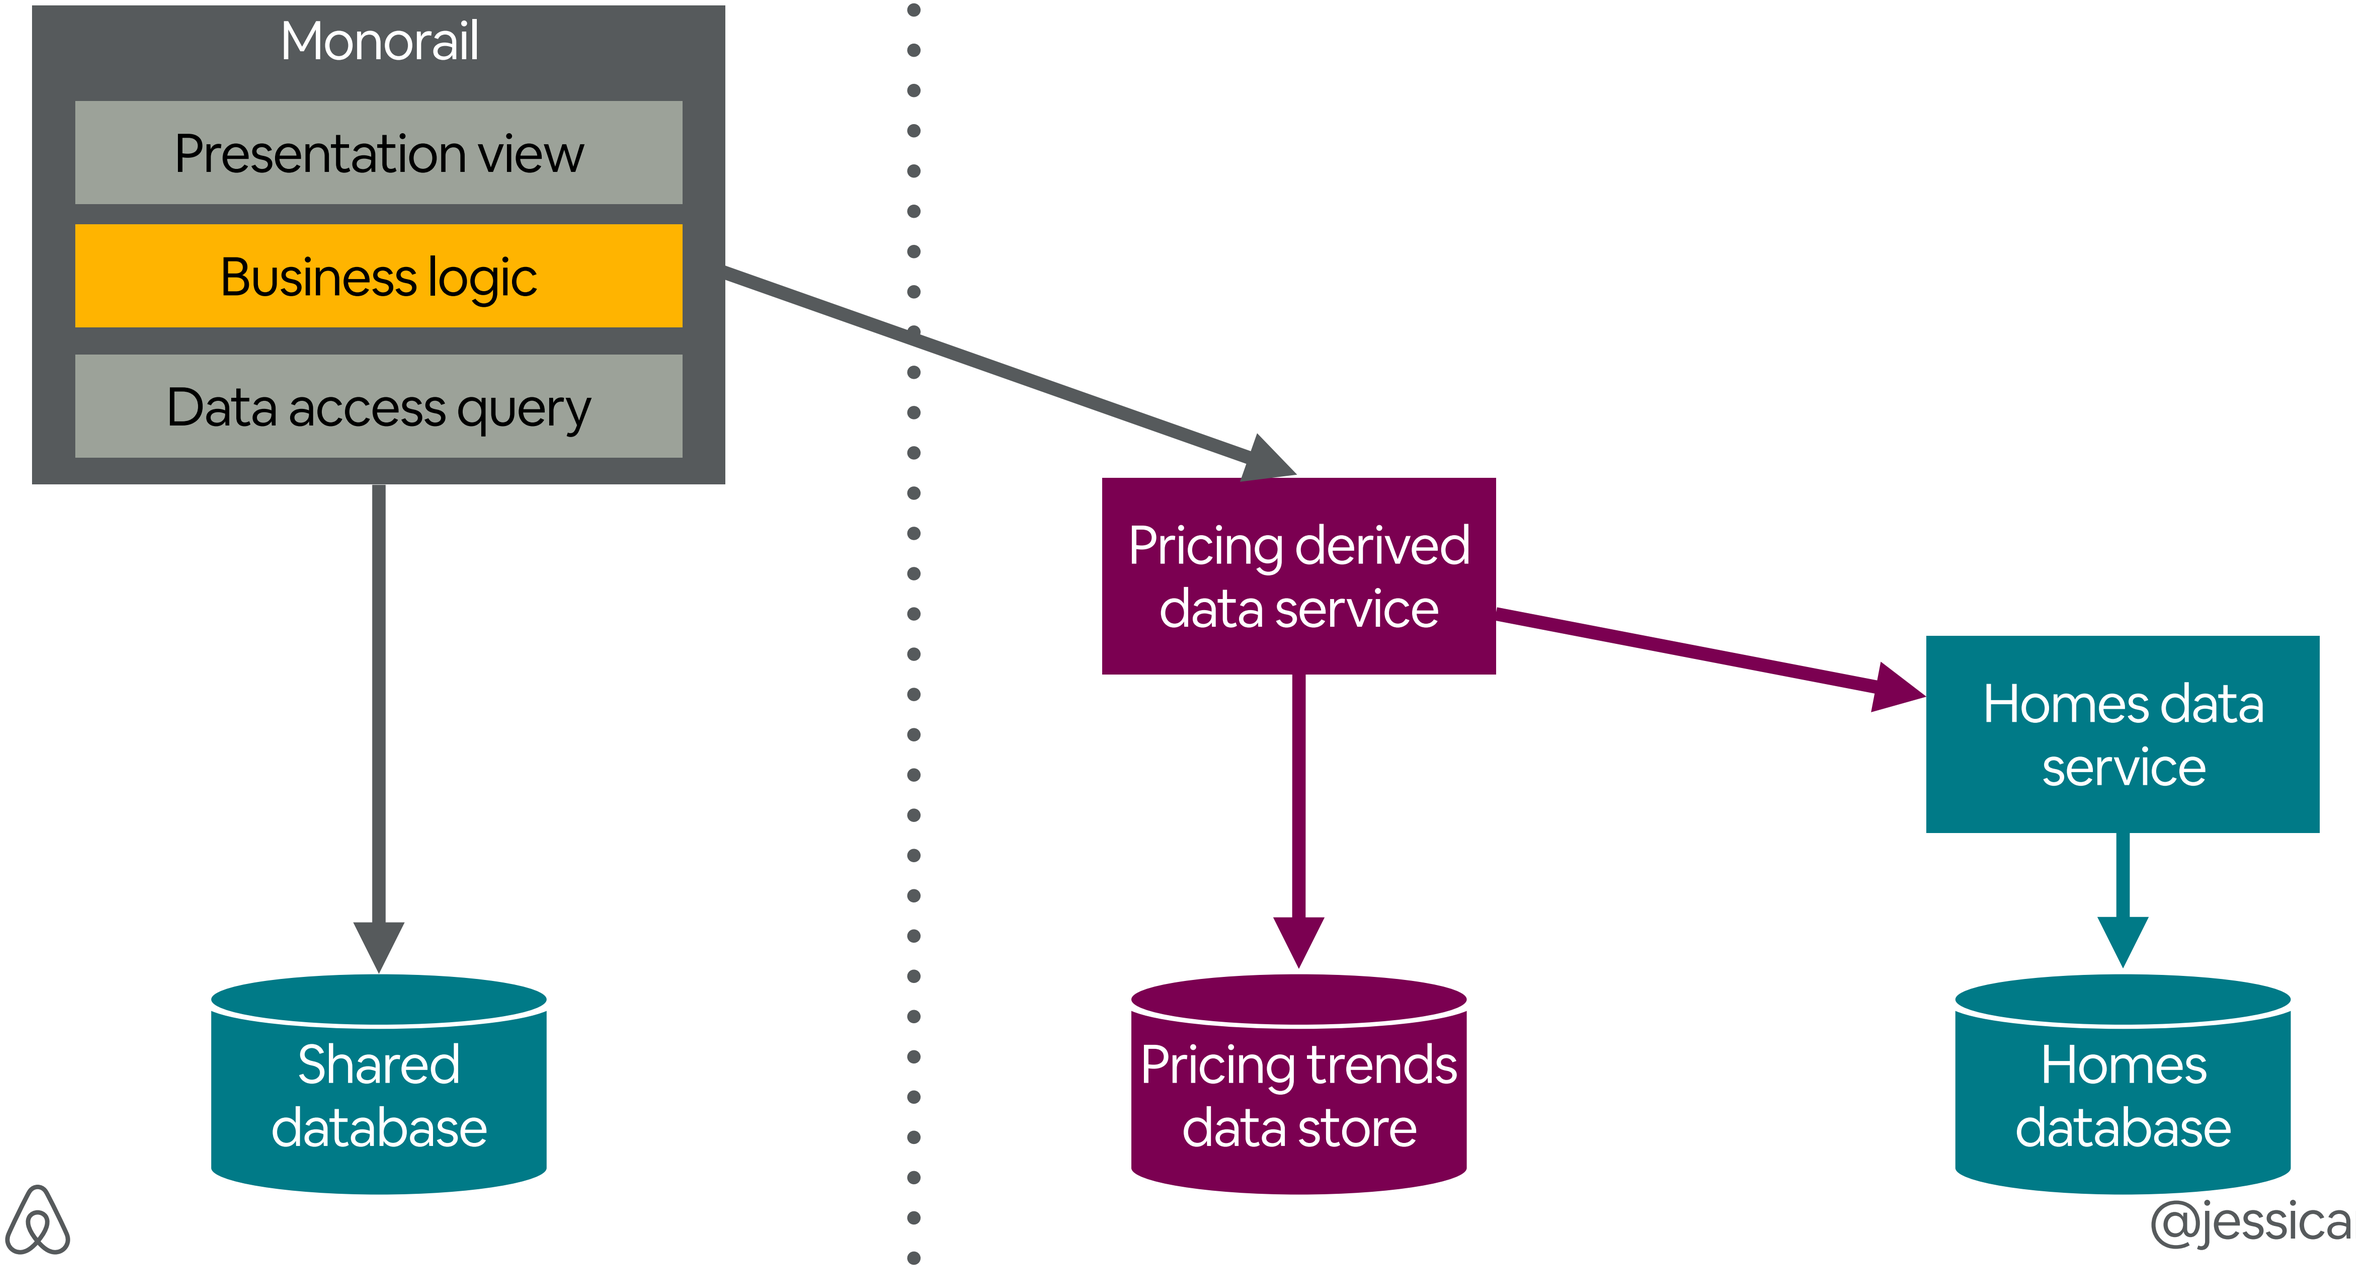

@jessicamtai

3. Migrating core product views

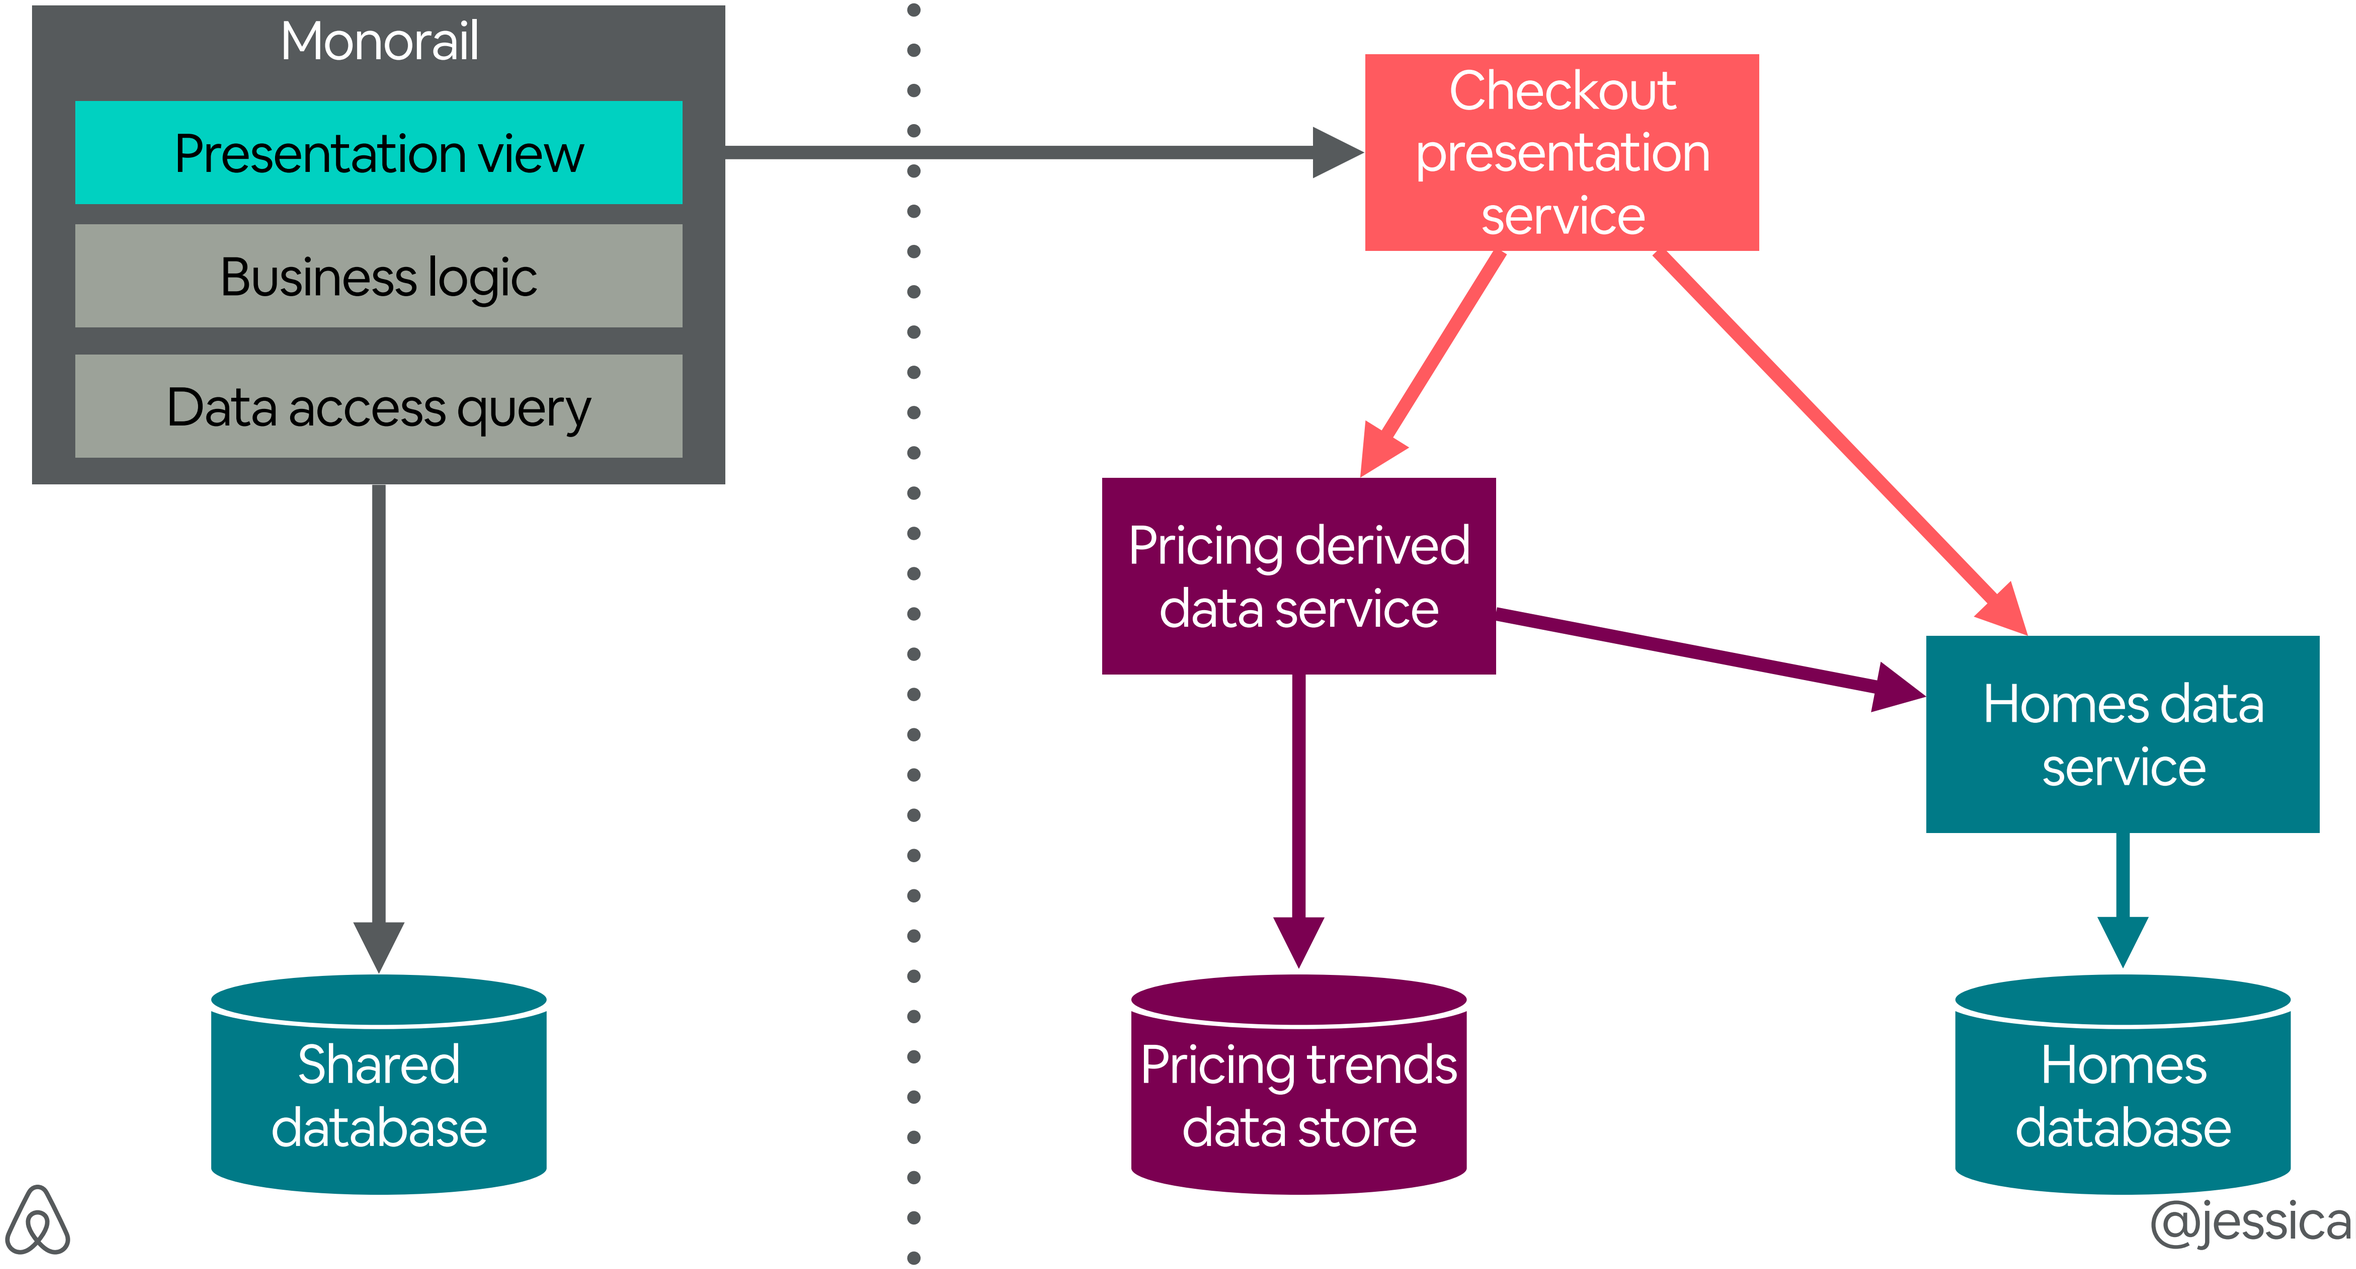

@jessicamtai

4. Migrating core product writes

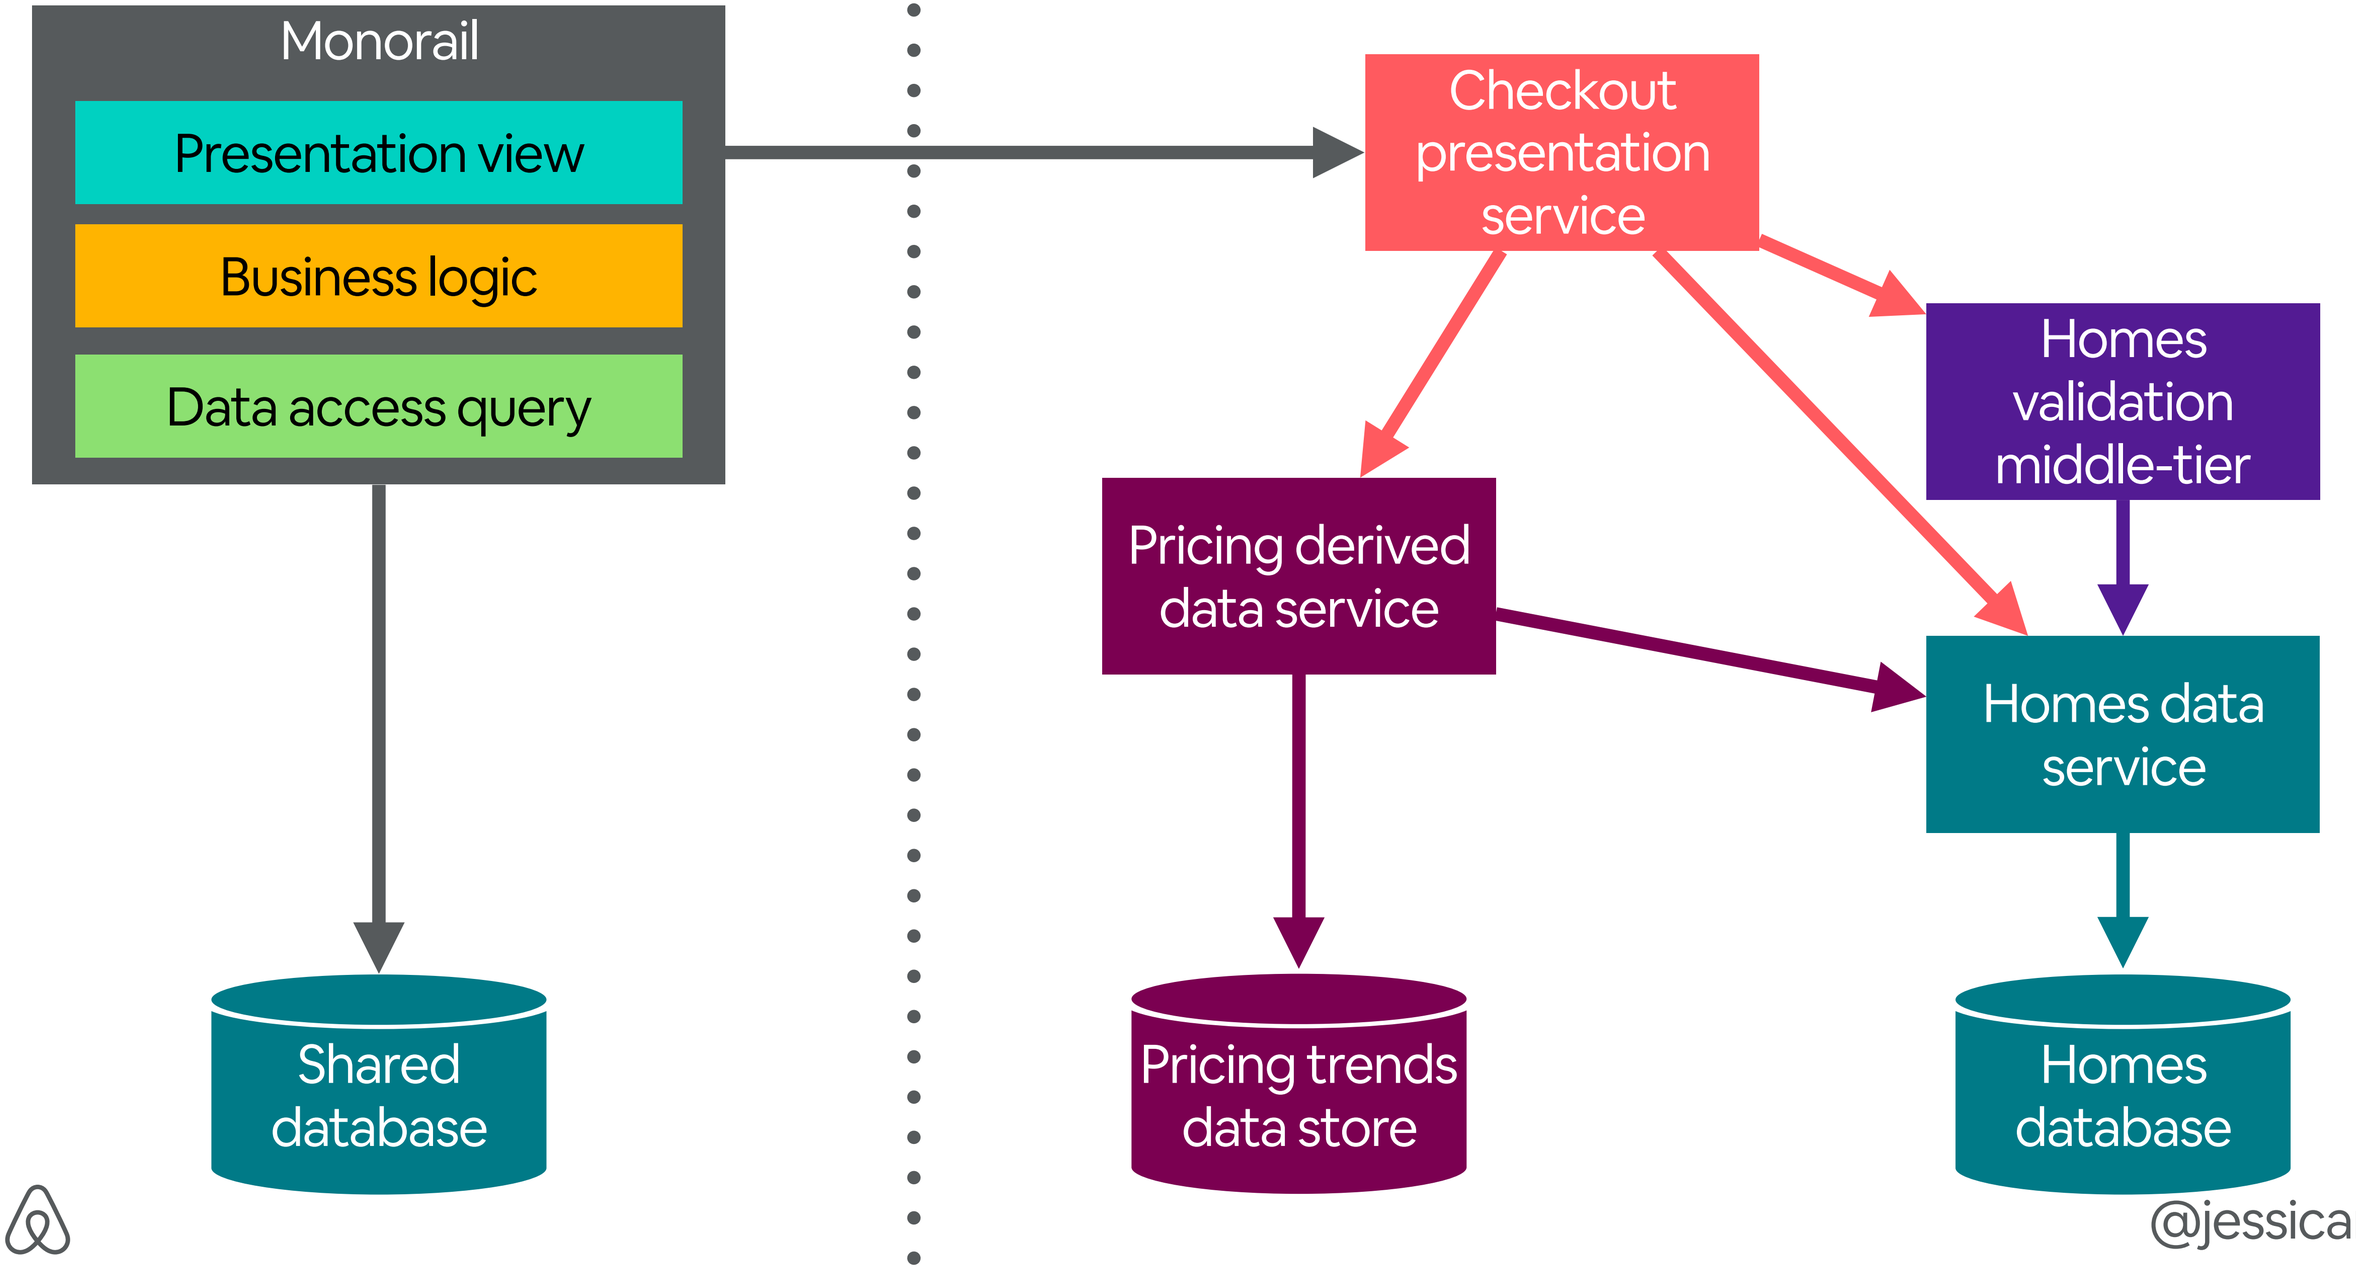

Request life cycle

V2: MONORAIL & SERVICES

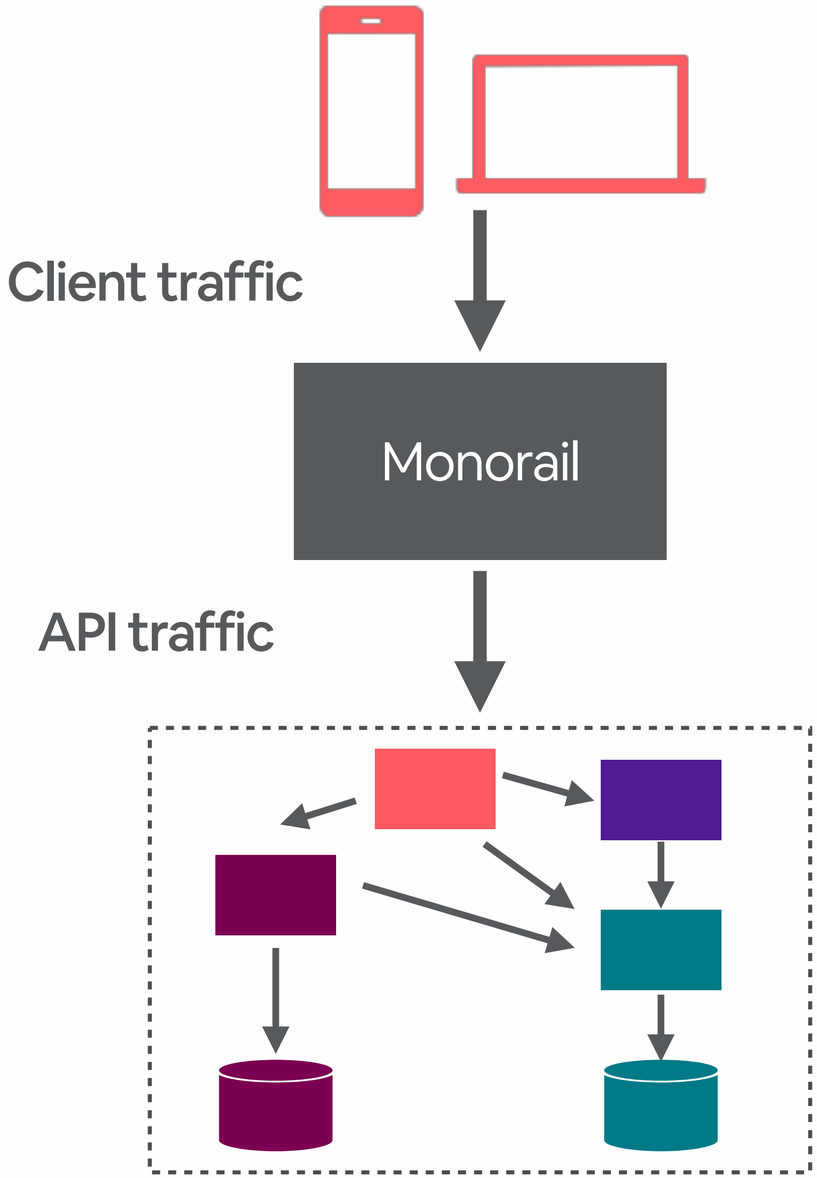

Routing & view

Business logic, model, data
via services

Request life cycle

V3: SOA & API GATEWAY

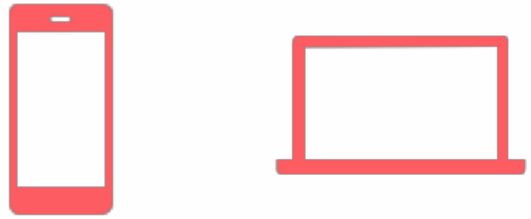

Routing

API gateway

Request context

Middleware

Session data service

Authentication data service

Oauth data service

Risk derived data service

...

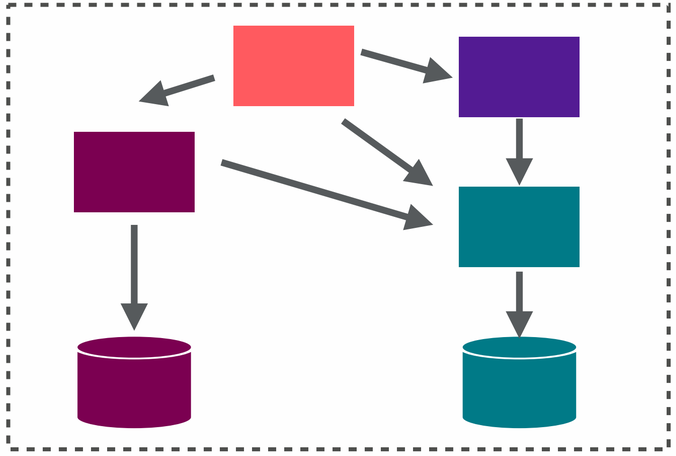

Presentation, logic, data

@jessicamtai

Request life cycle

V3: SOA & API GATEWAY

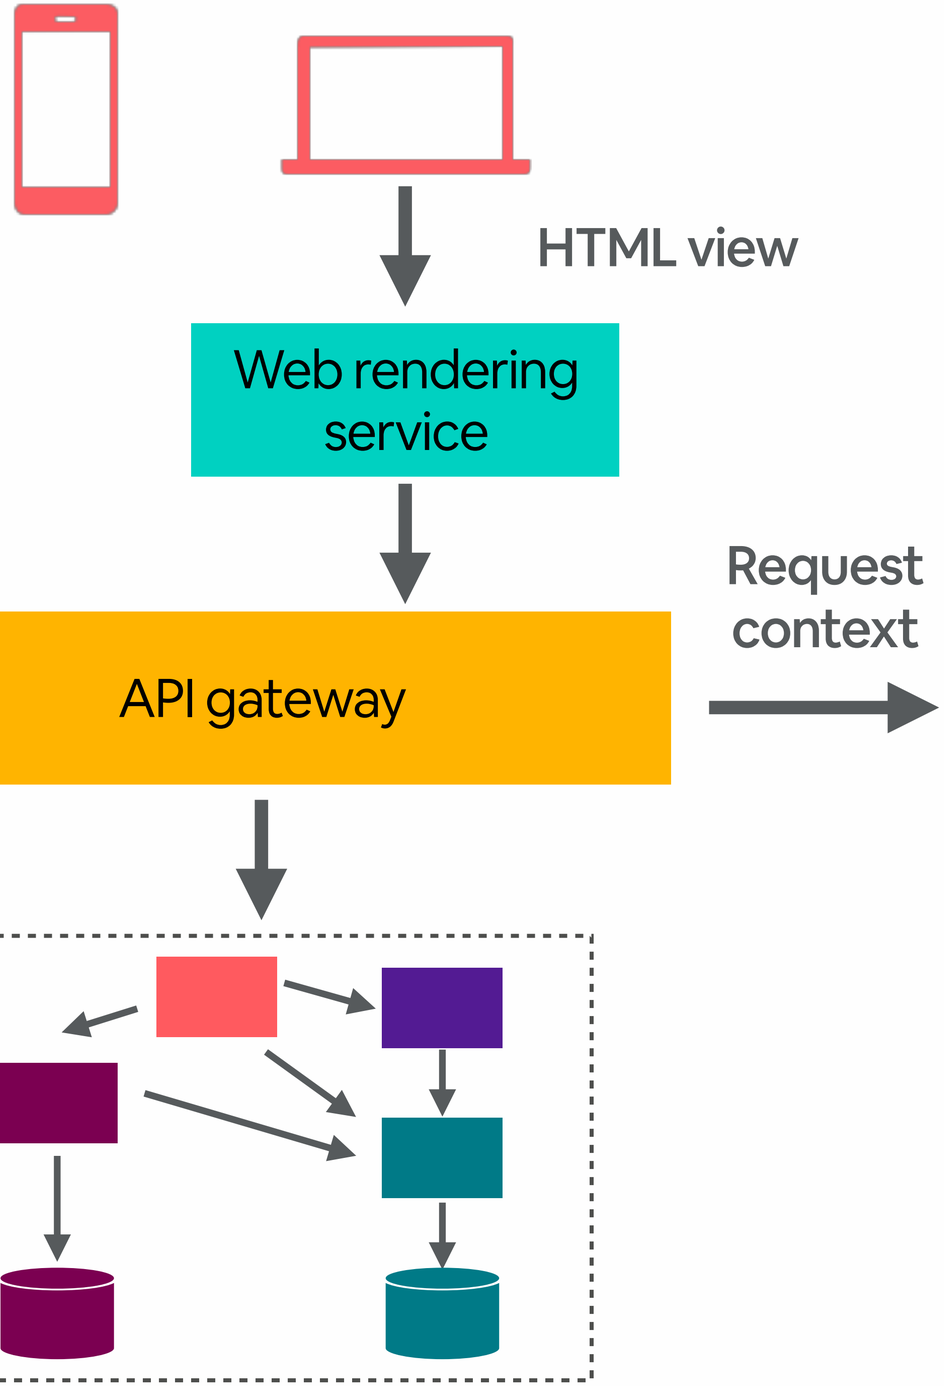

@jessicamtai

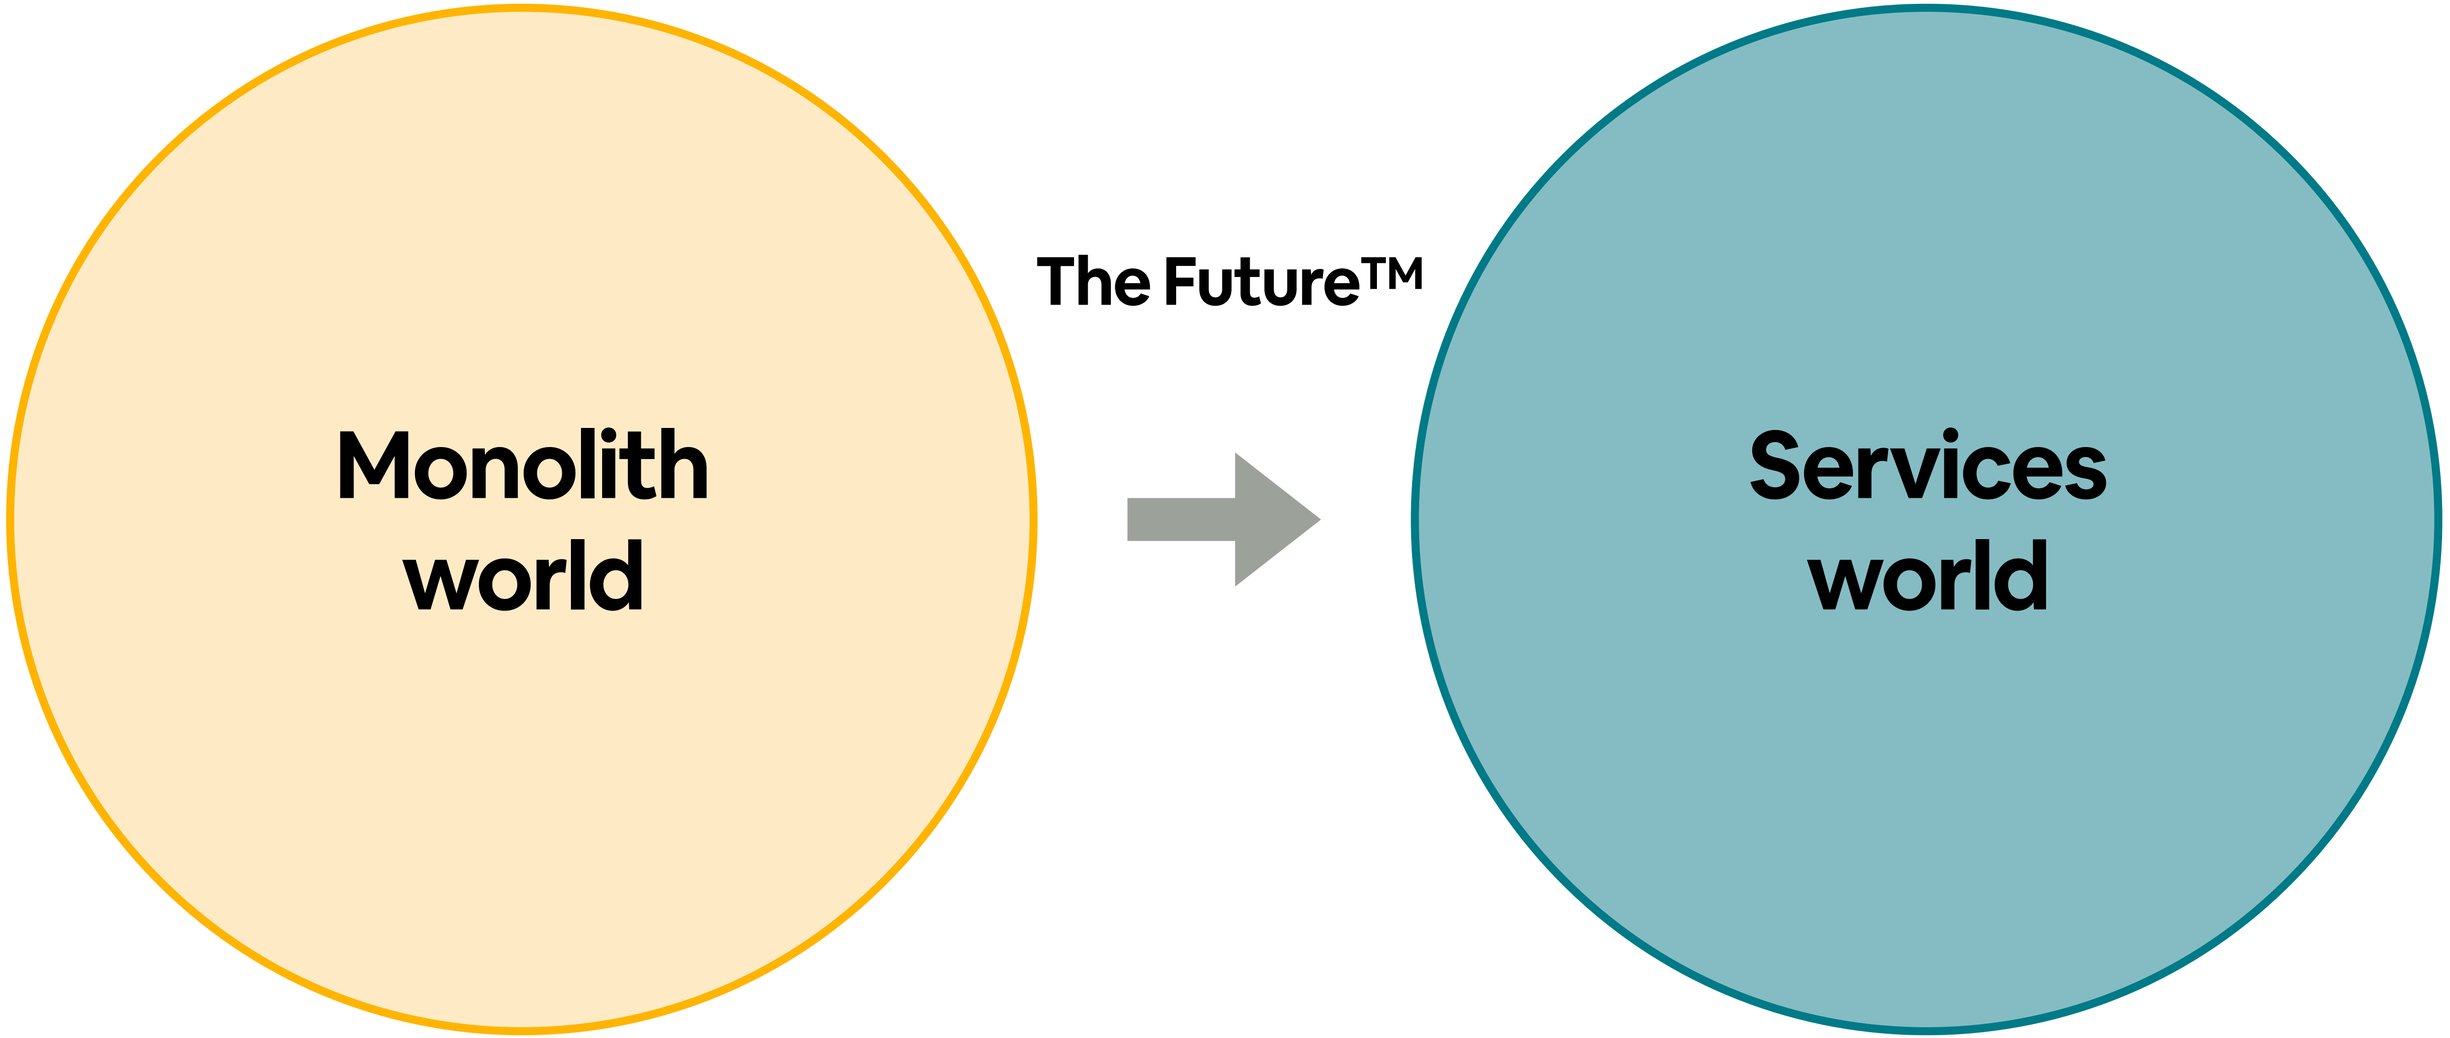
Monolith world
The Future™
Services world

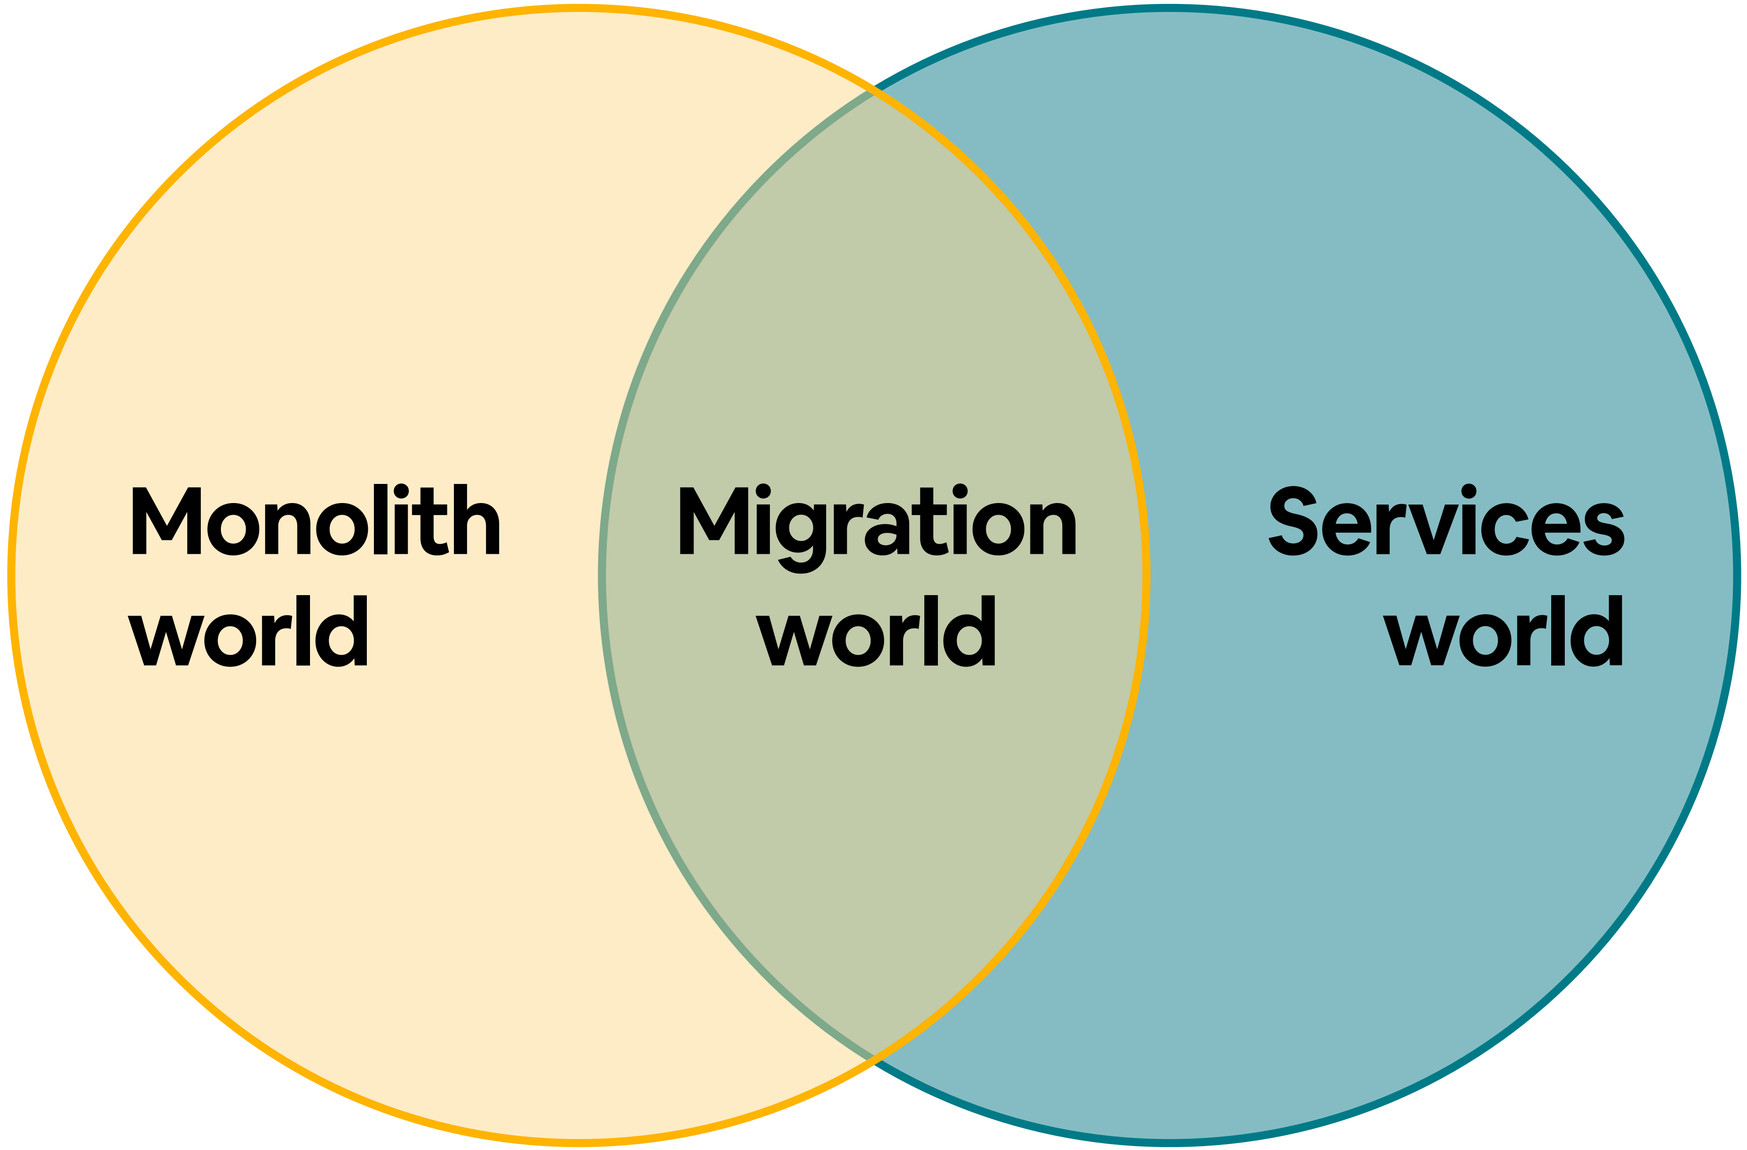
Monolith world
Migration world
Services world


COMPARE FOR DIFFERENCES
@jessicamtai

Dual read comparison

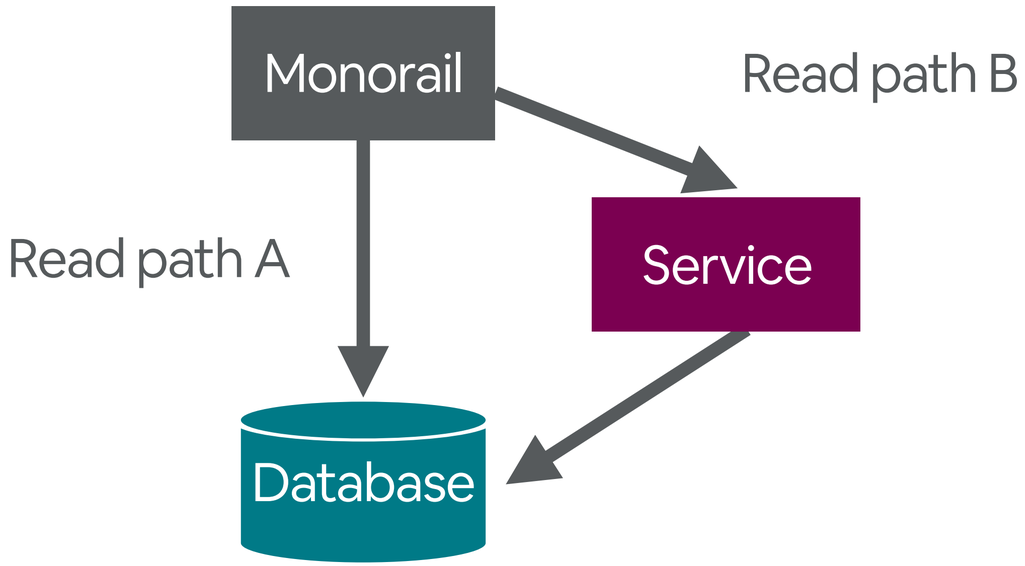

Gate

Admin web UI configuration

Compare

1% traffic

Ramp & wait

Gradual increments

Wait

Gather traffic patterns

Switch

All traffic through service only

Dual read comparison

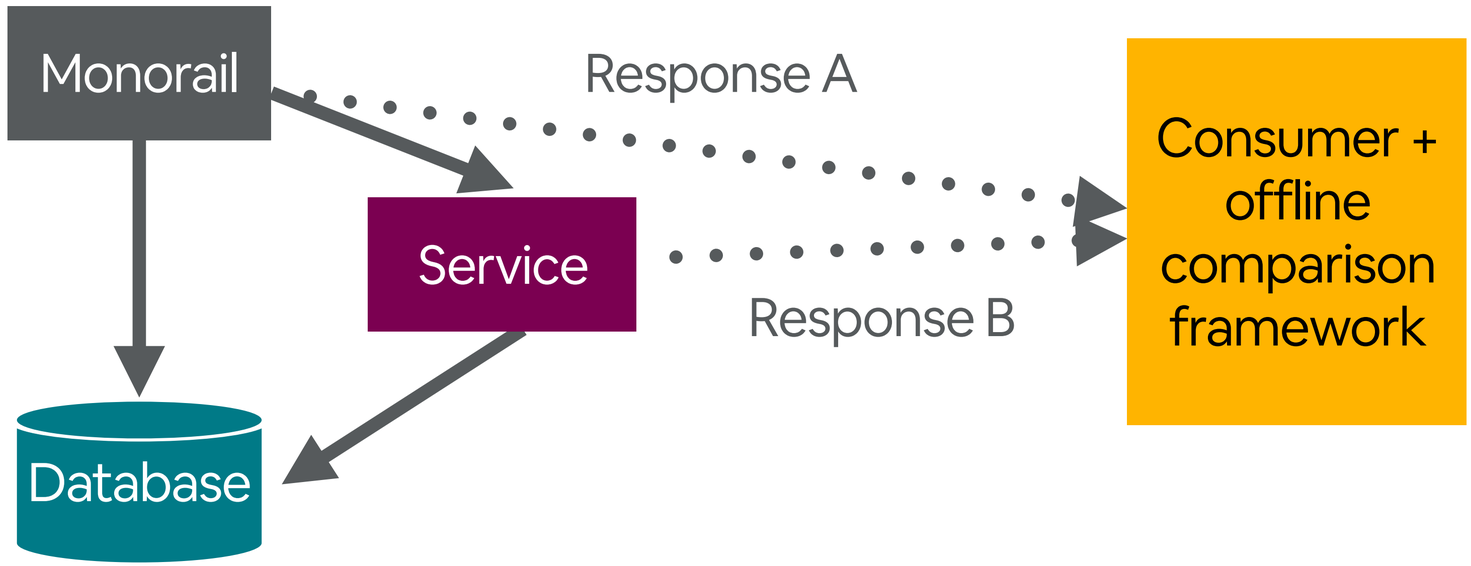

Gate

Admin web UI configuration

Compare

1% traffic

Ramp & wait

Gradual increments

Wait

Gather traffic patterns

Switch

All traffic through service only

Dual read comparison

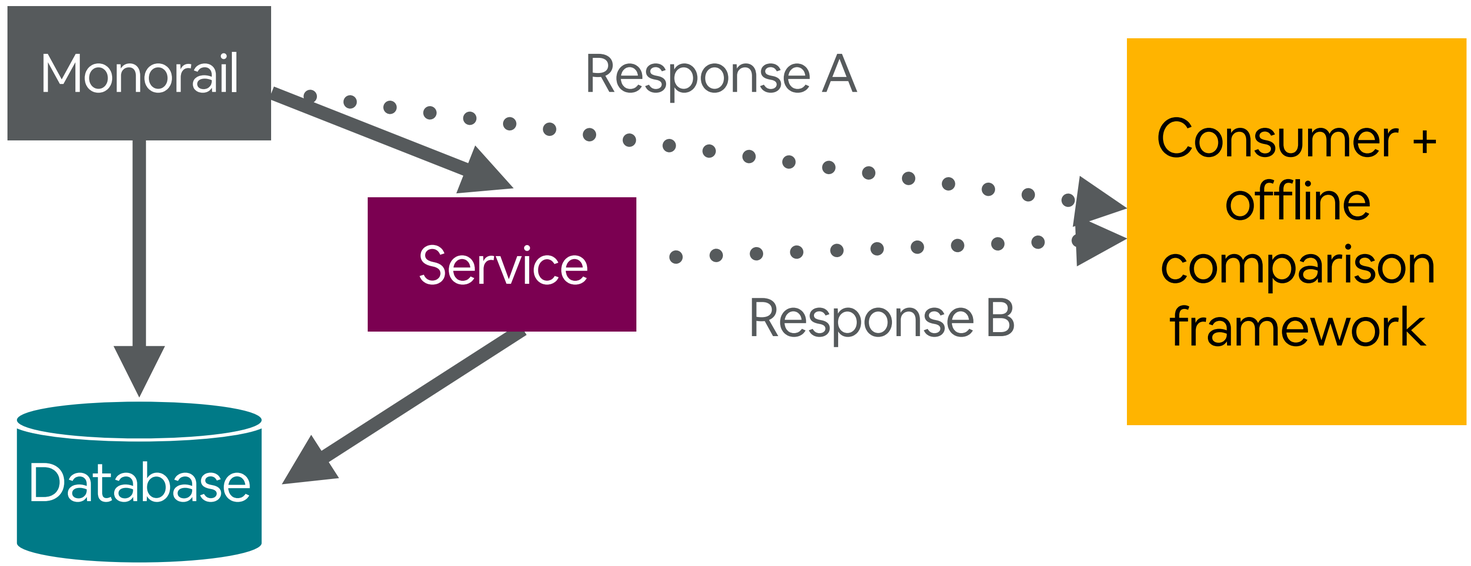

Gate — Admin web UI configuration

Compare — 1% traffic

Ramp & wait — Gradual increments

Wait — Gather traffic patterns

Switch — All traffic through service only

Dual read comparison

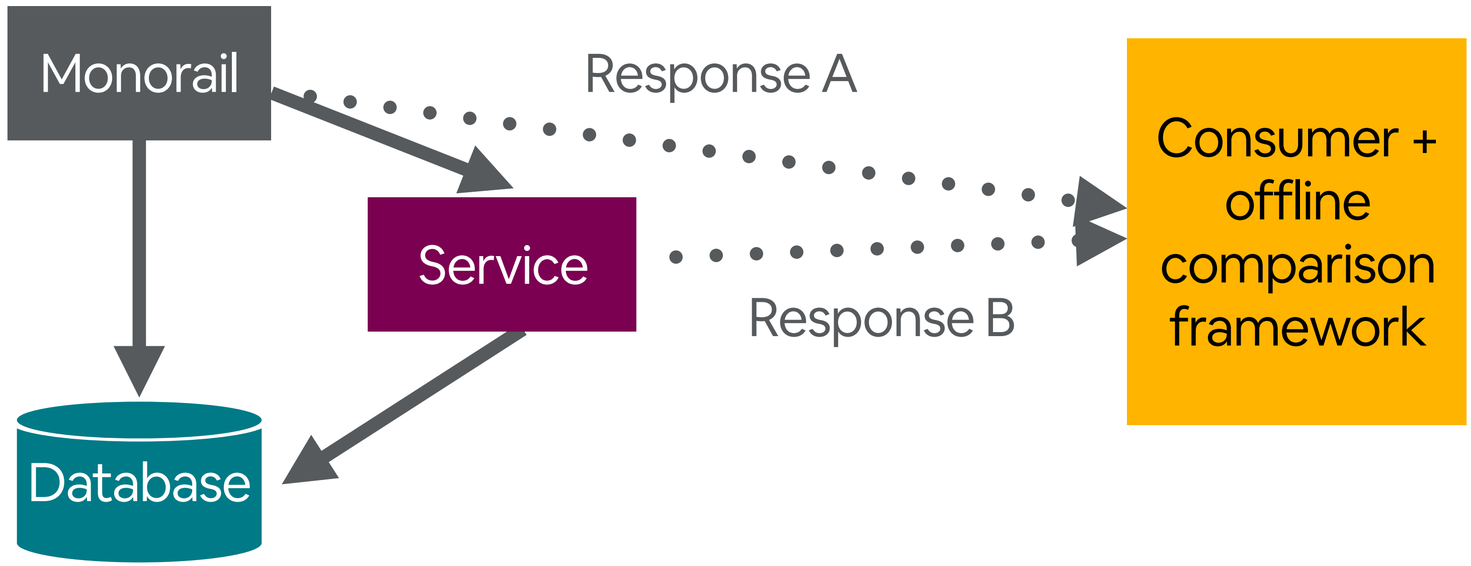

Gate — Admin web UI configuration

Compare — 1% traffic

Ramp & wait — Gradual increments

Wait — Gather traffic patterns

Switch — All traffic through service only

Dual read comparison

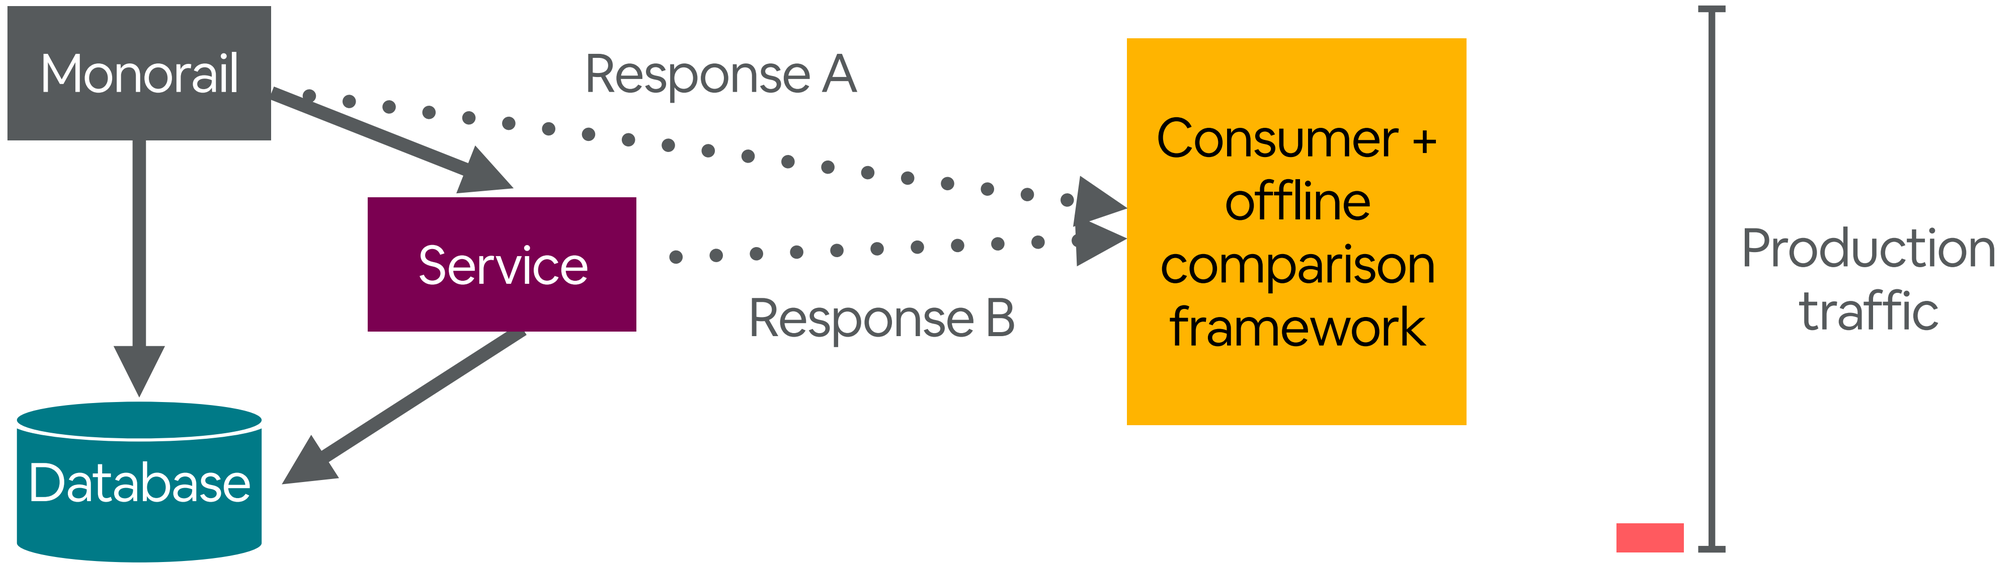

Gate
Admin web UI configuration

Compare
1% traffic

Ramp & wait
Gradual increments

Wait
Gather traffic patterns

Switch
All traffic through service only

Dual read comparison

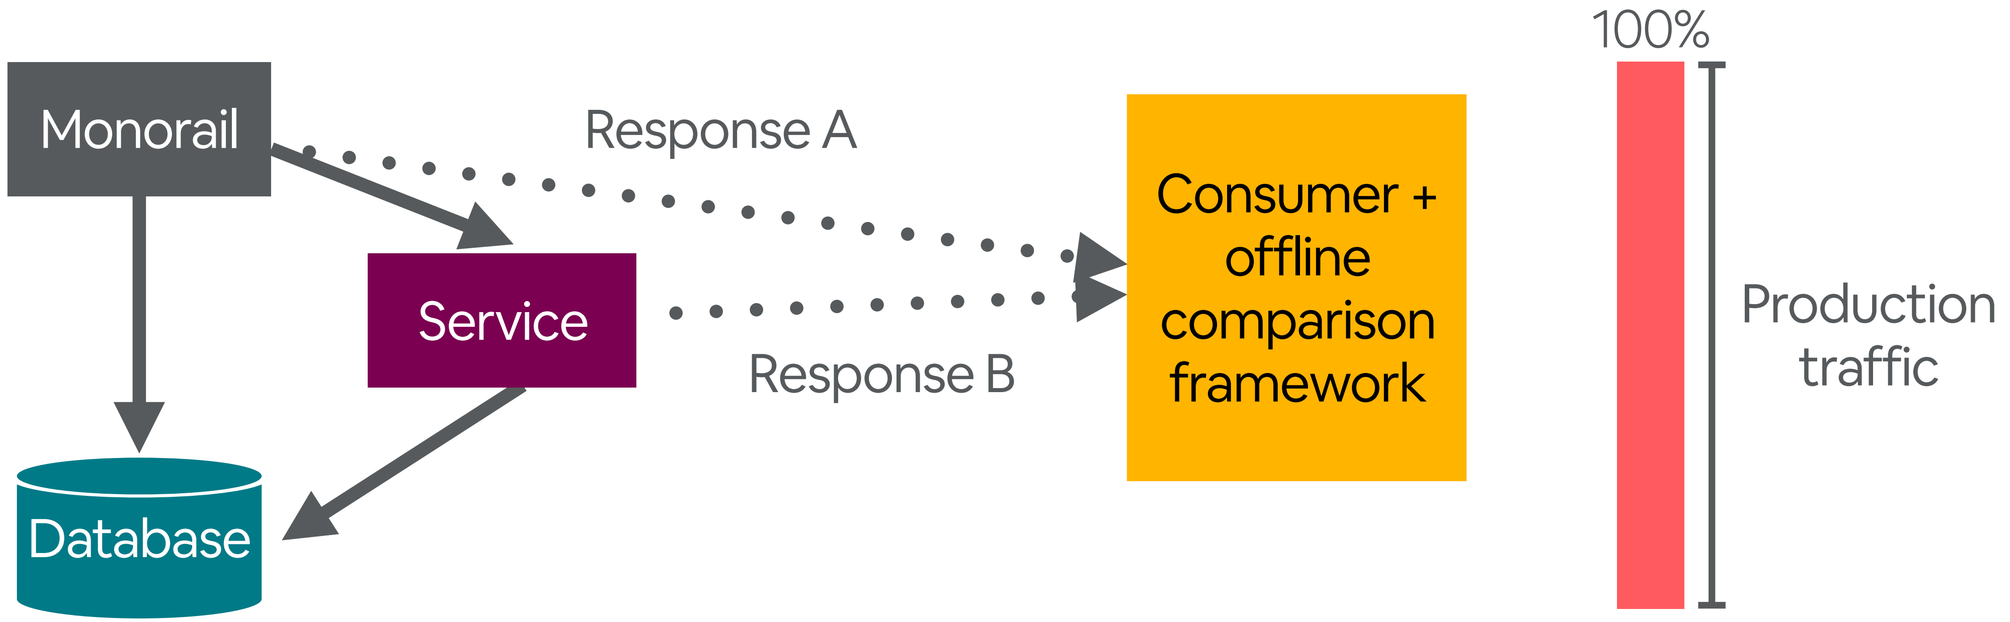

Gate
Admin web UI configuration

Compare
1% traffic

Ramp & wait
Gradual increments

Wait
Gather traffic patterns

Switch
All traffic through service only

Dual read comparison

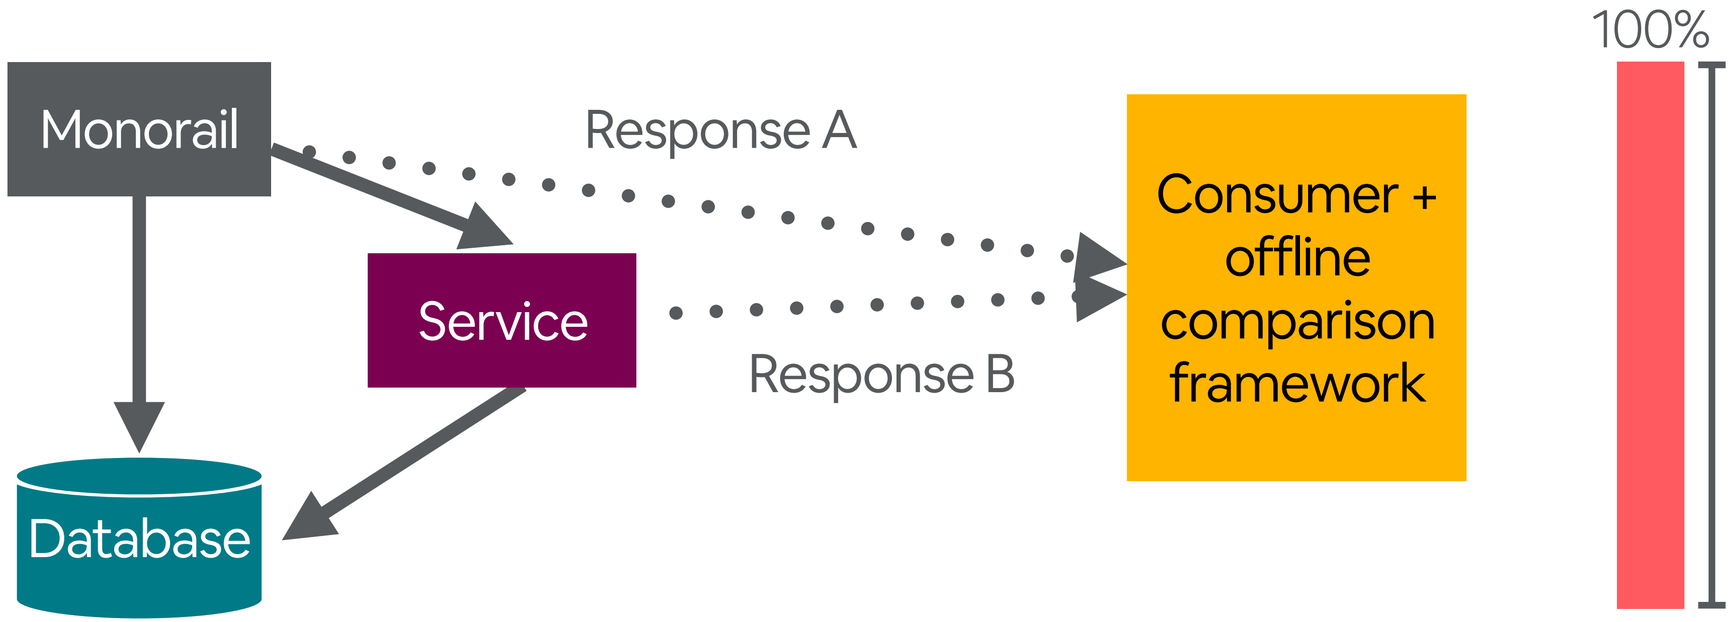

Gate — Admin web UI configuration

Compare — 1% traffic

Ramp & wait — Gradual increments

Wait — Gather traffic patterns

Switch — All traffic through service only

Dual read comparison

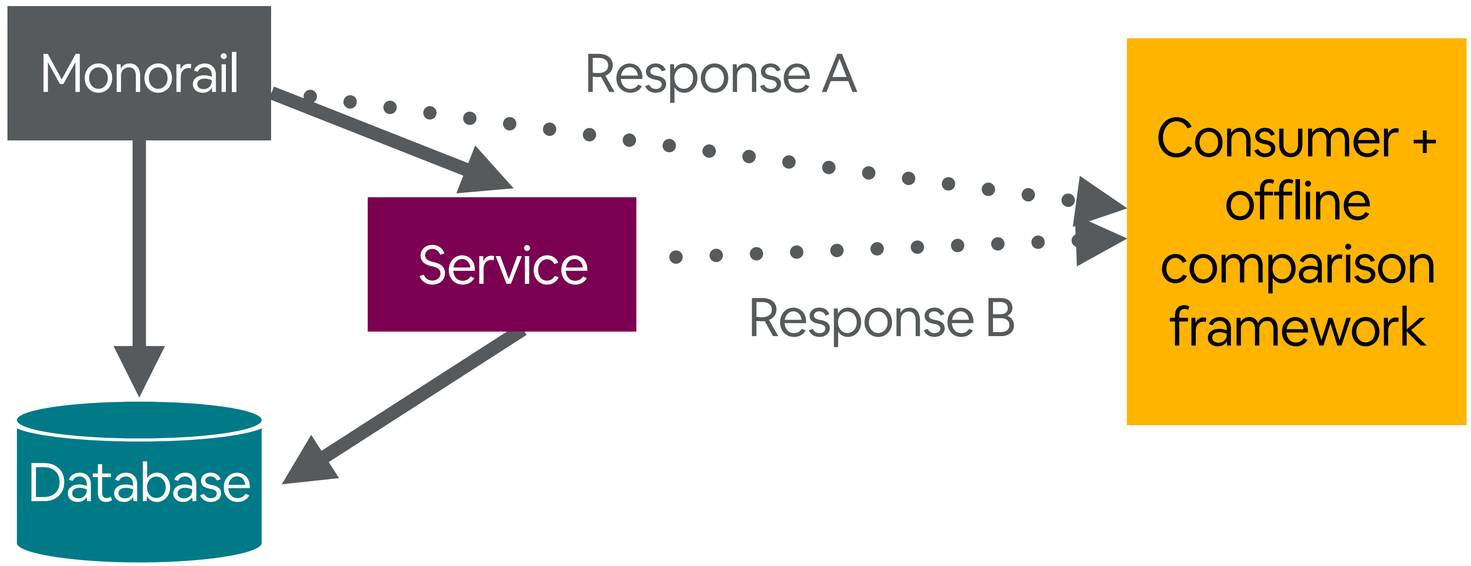

Gate

Admin web UI configuration

Compare

1% traffic

Ramp & wait

Gradual increments

Wait

Gather traffic patterns

Switch

All traffic through service only

Dual read comparison

The only read path

Gate

Admin web UI configuration

Compare

1% traffic

Ramp & wait

Gradual increments

Wait

Gather traffic patterns

Switch

All traffic through service only

Write comparison

DUAL WRITE TO SEPARATE DATABASES

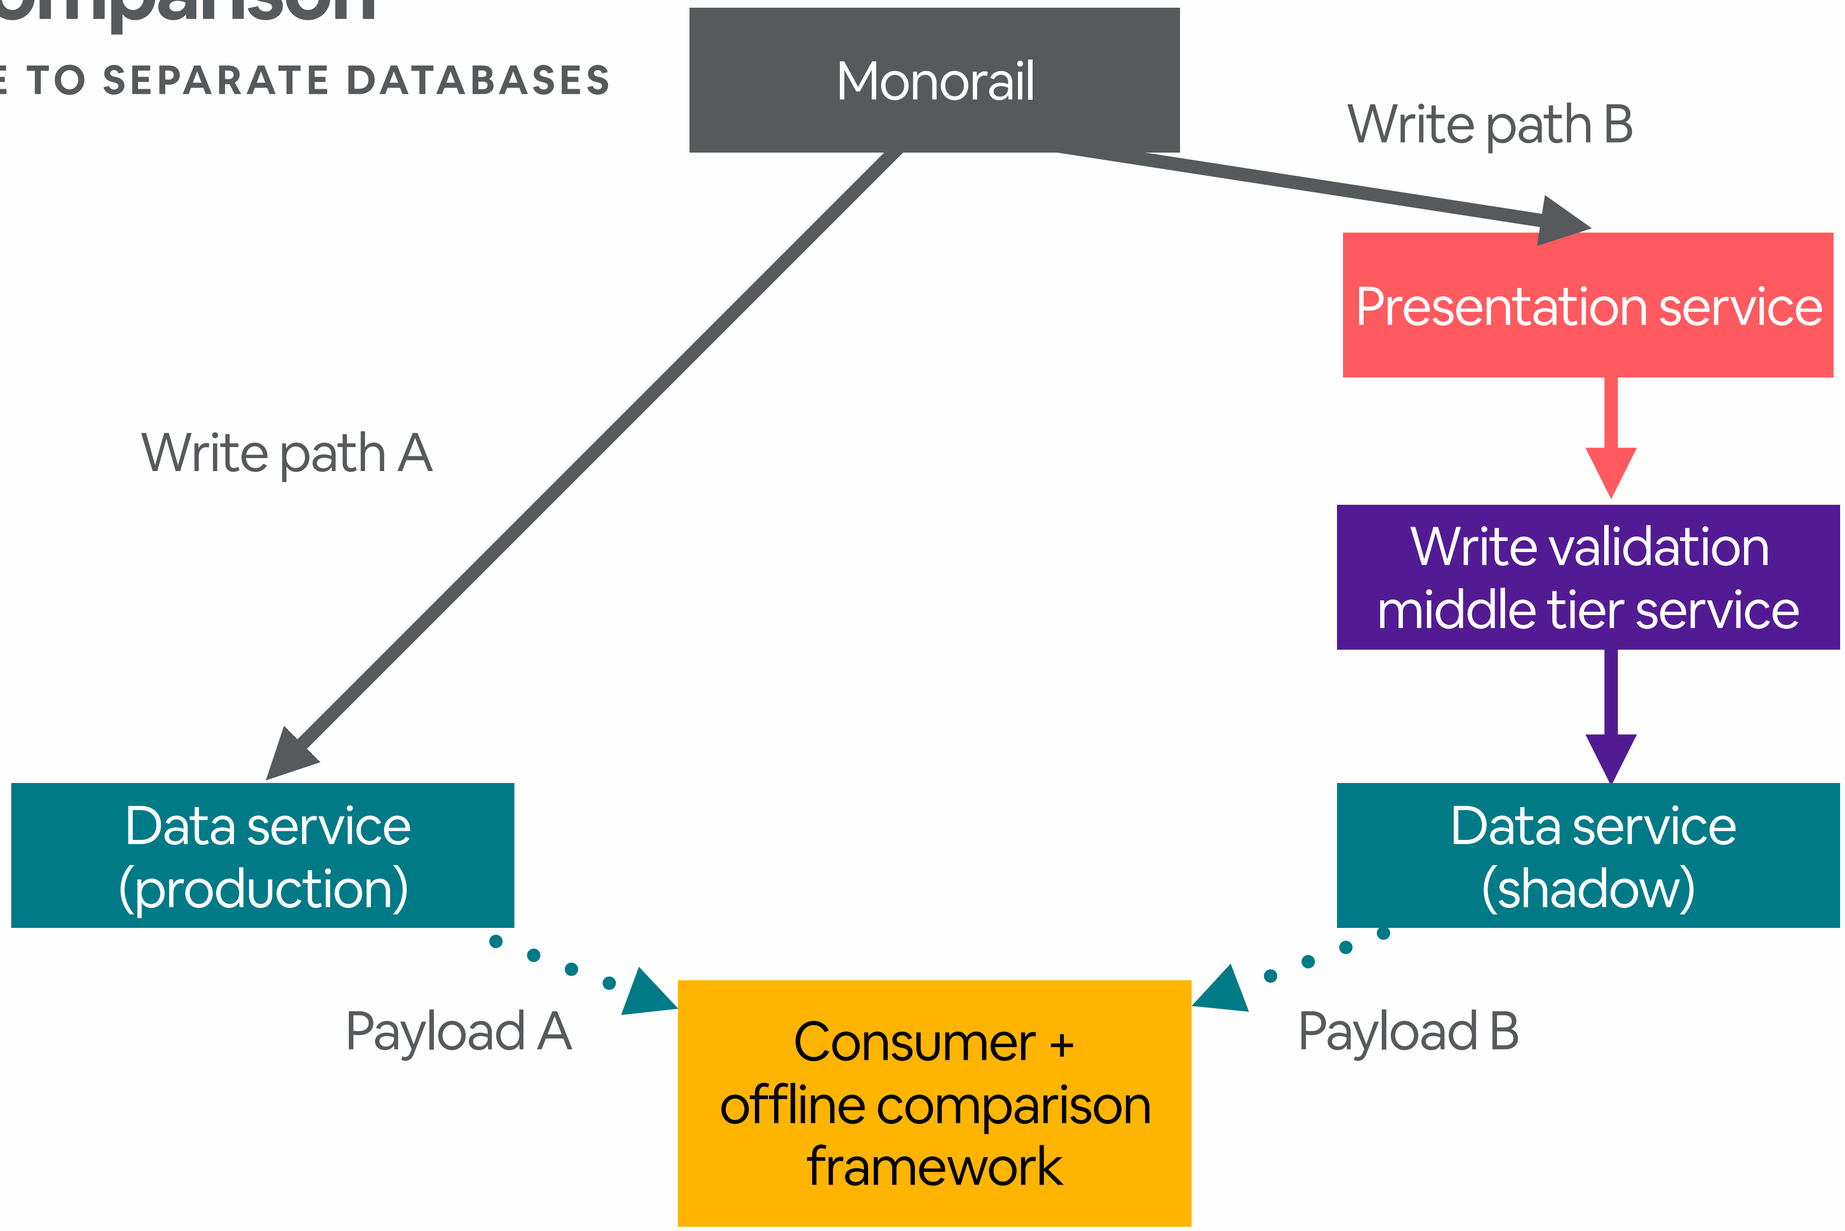

@jessicamtai

Write comparison

DUAL WRITE TO SEPARATE DATABASES

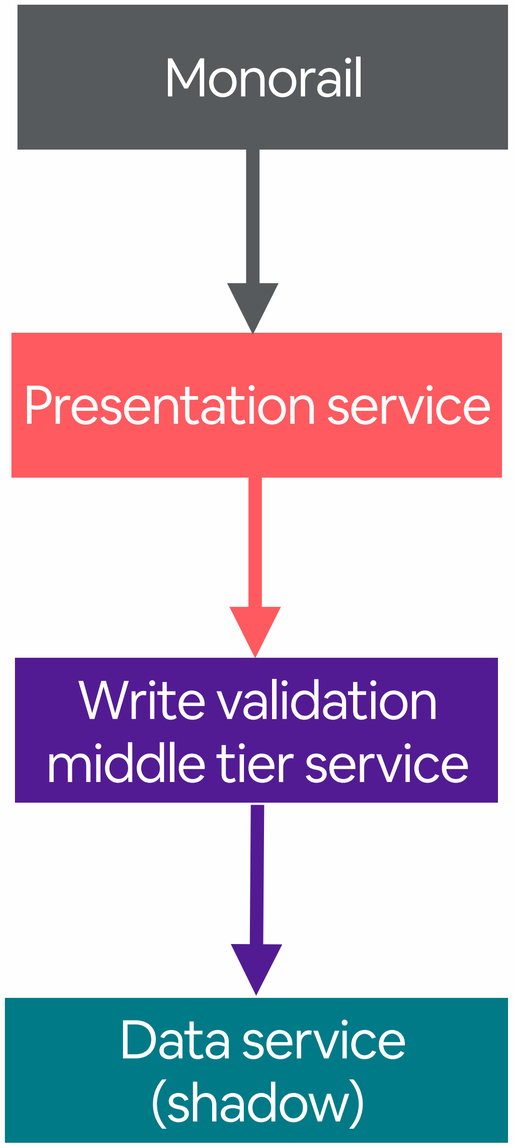

The only write path

API gateway comparison

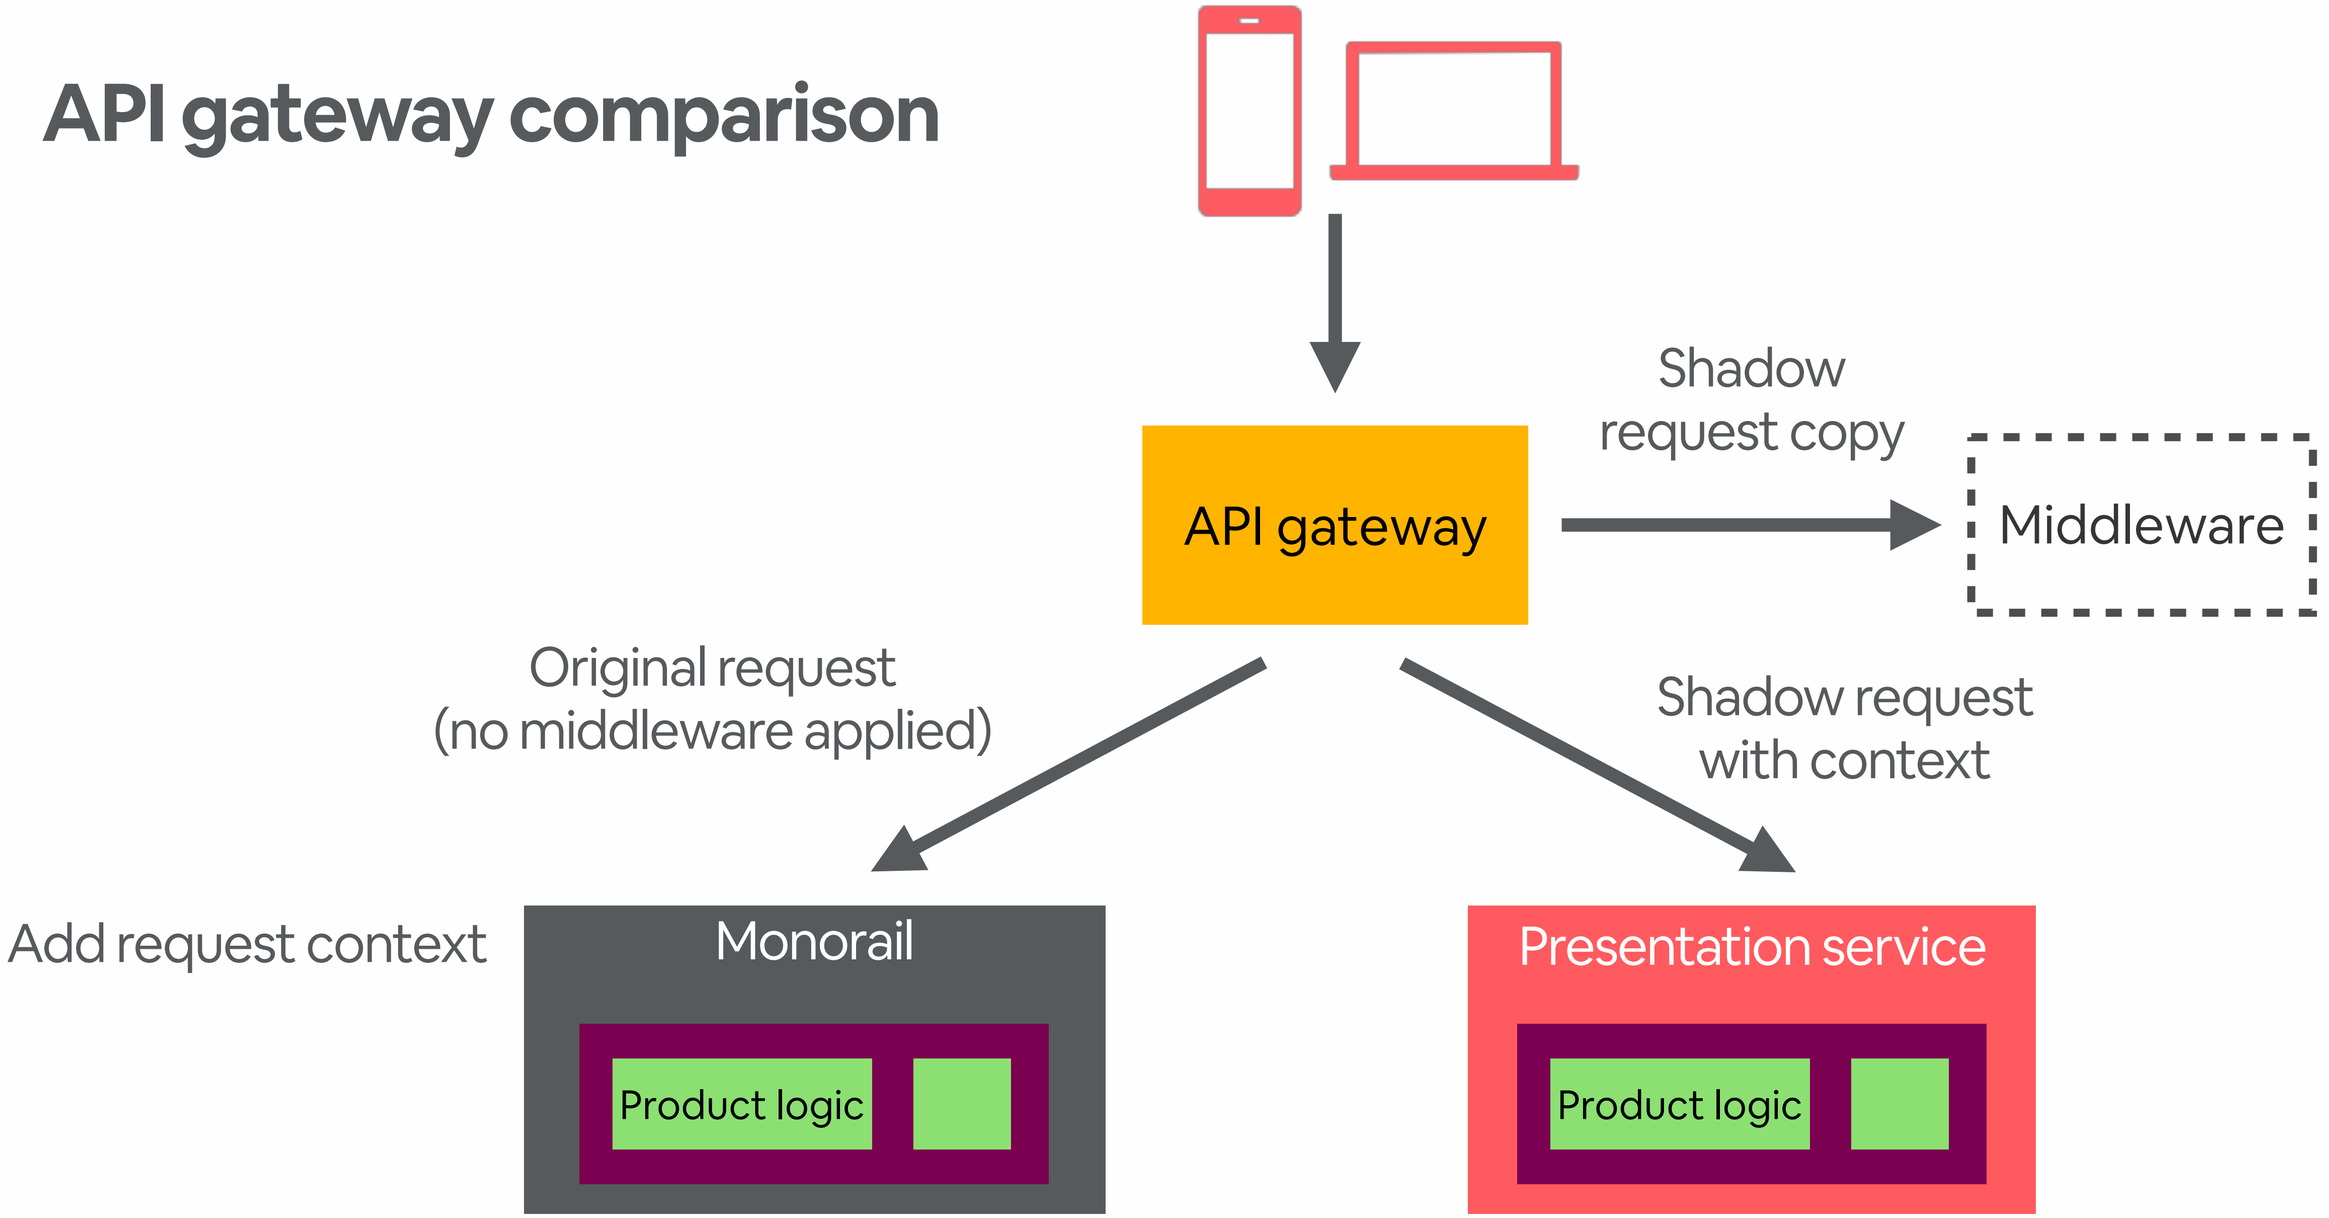

API gateway comparison

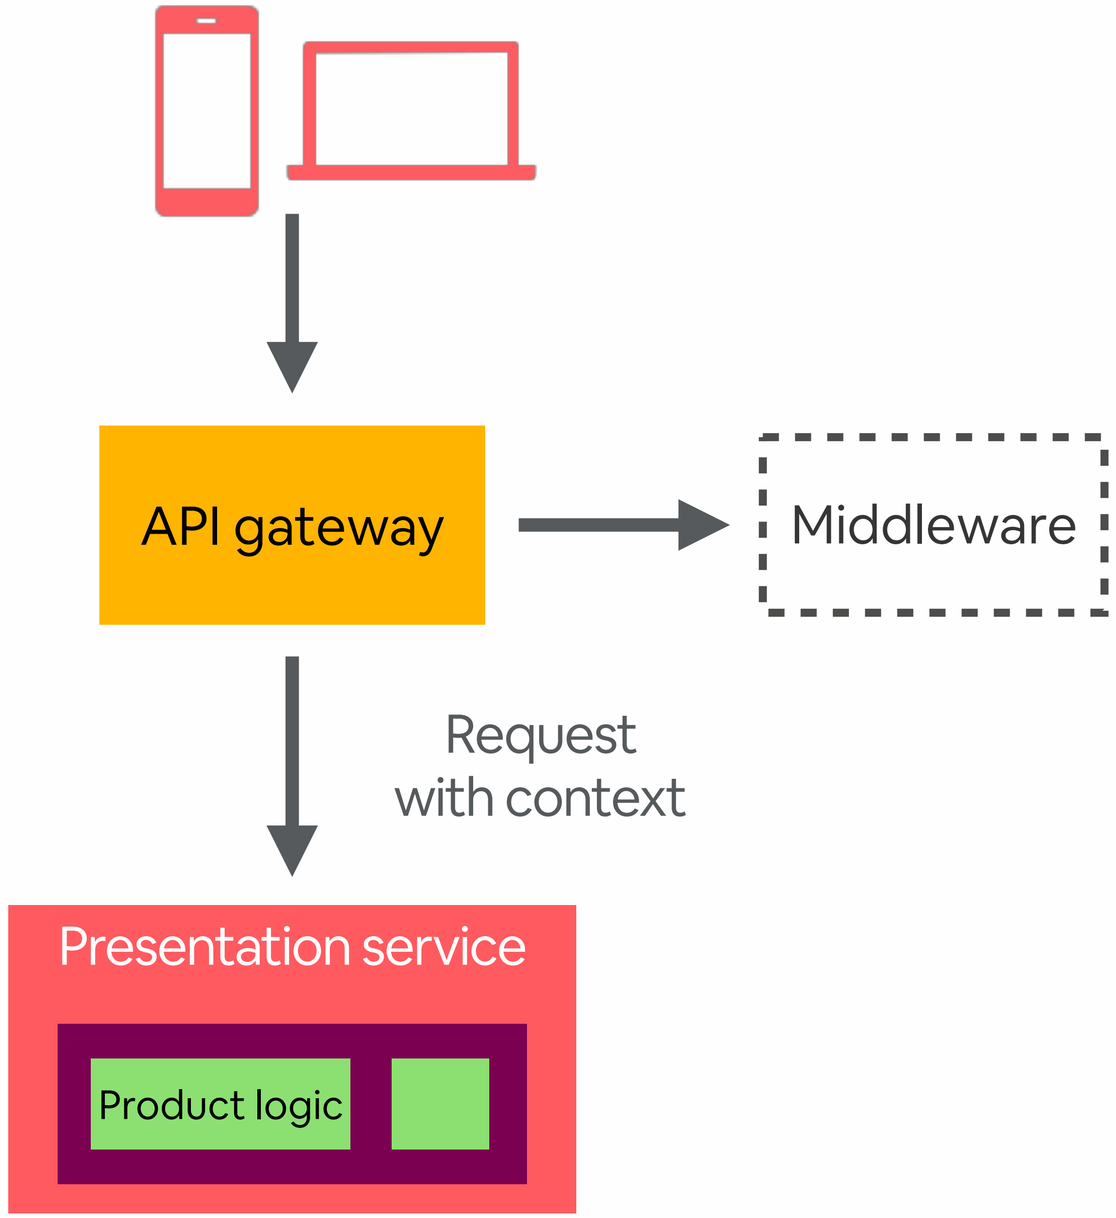


Incremental migration
@jessicamtai

Migrate by endpoint

- Production traffic with partially complete service
  - e.g. batch API `/loadUsers`
  - Fetch users only by user id
- Unblock clients

Migrate by attribute

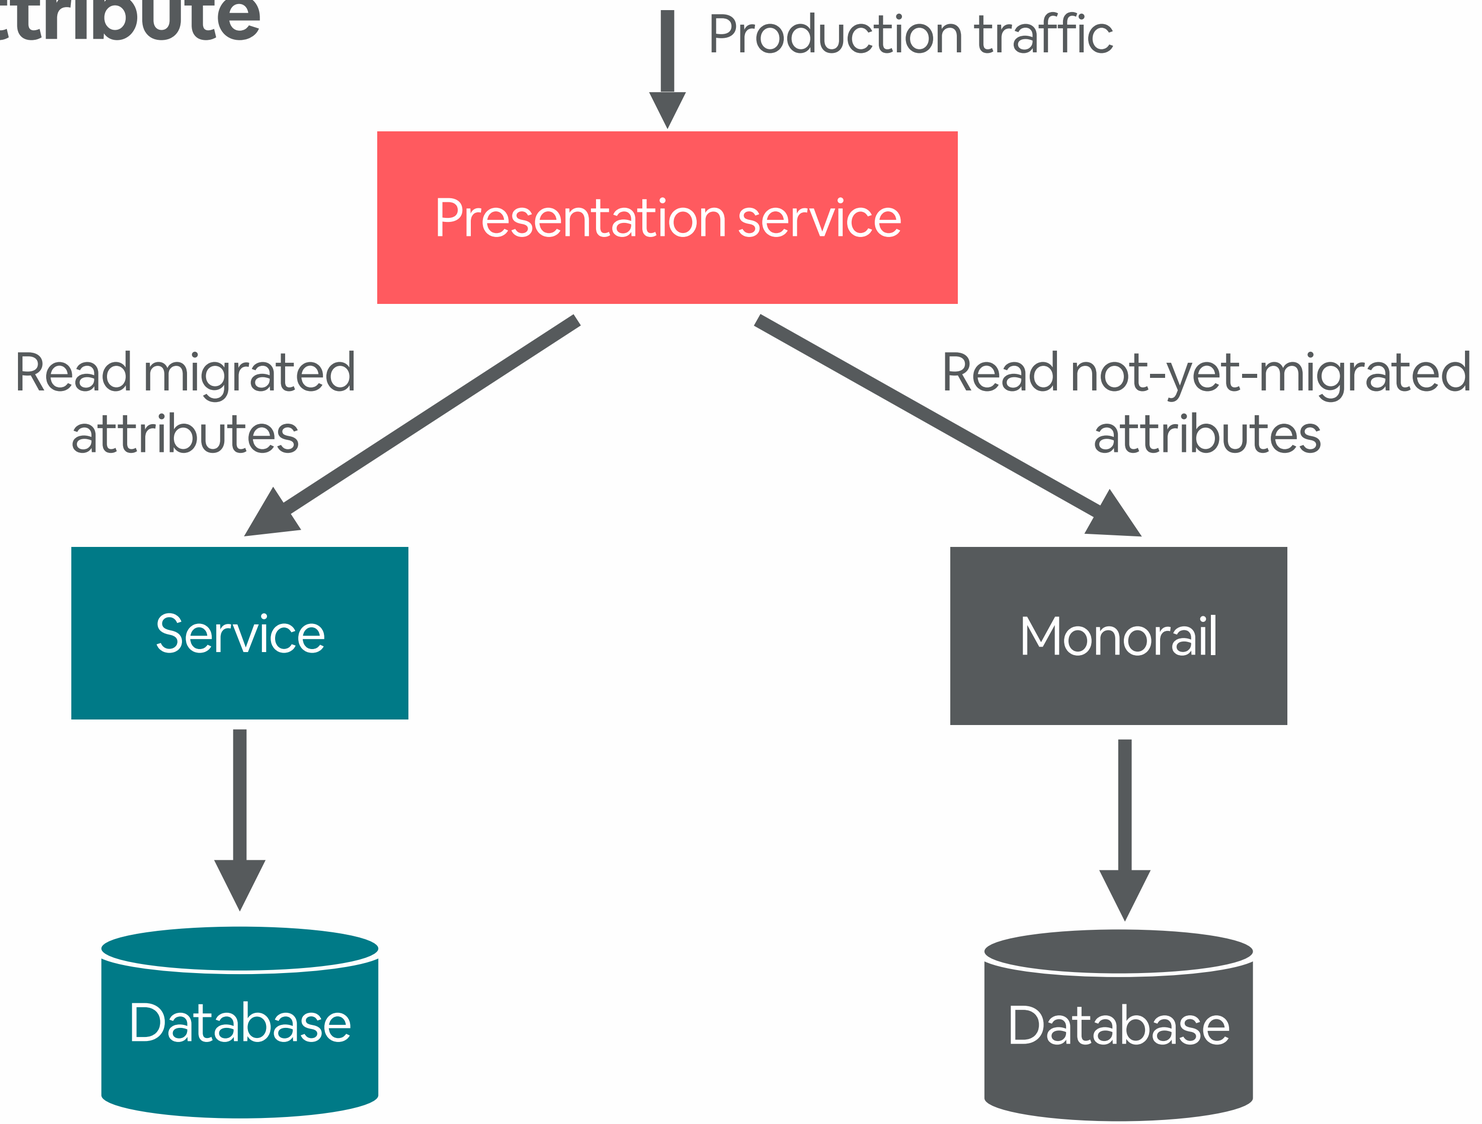


SOA BEST PRACTICES
@jessicamtai

Standardize service building

SCALE WITH CONSISTENCY

Frameworks

Auto-generate code

Testing & deploying

Replay production traffic

Observability

Standard templates

Service & client setup

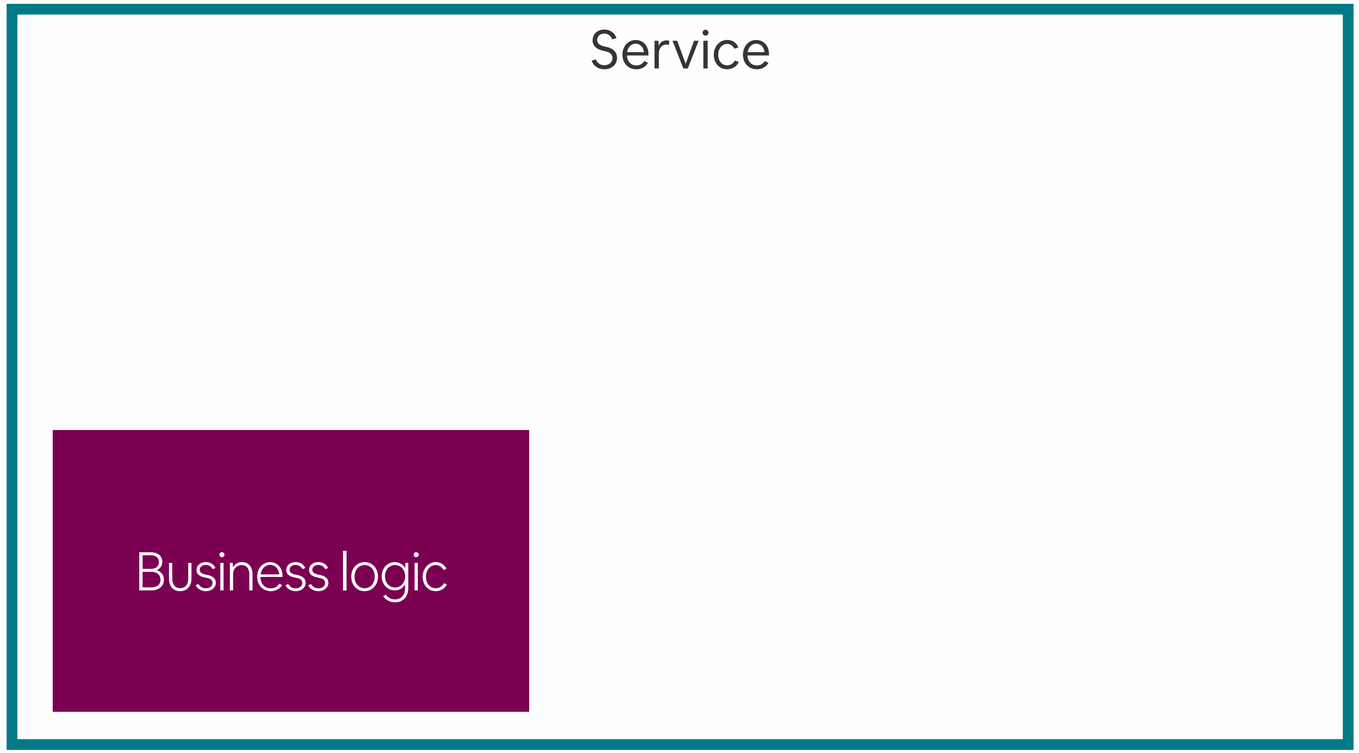

Service & client setup

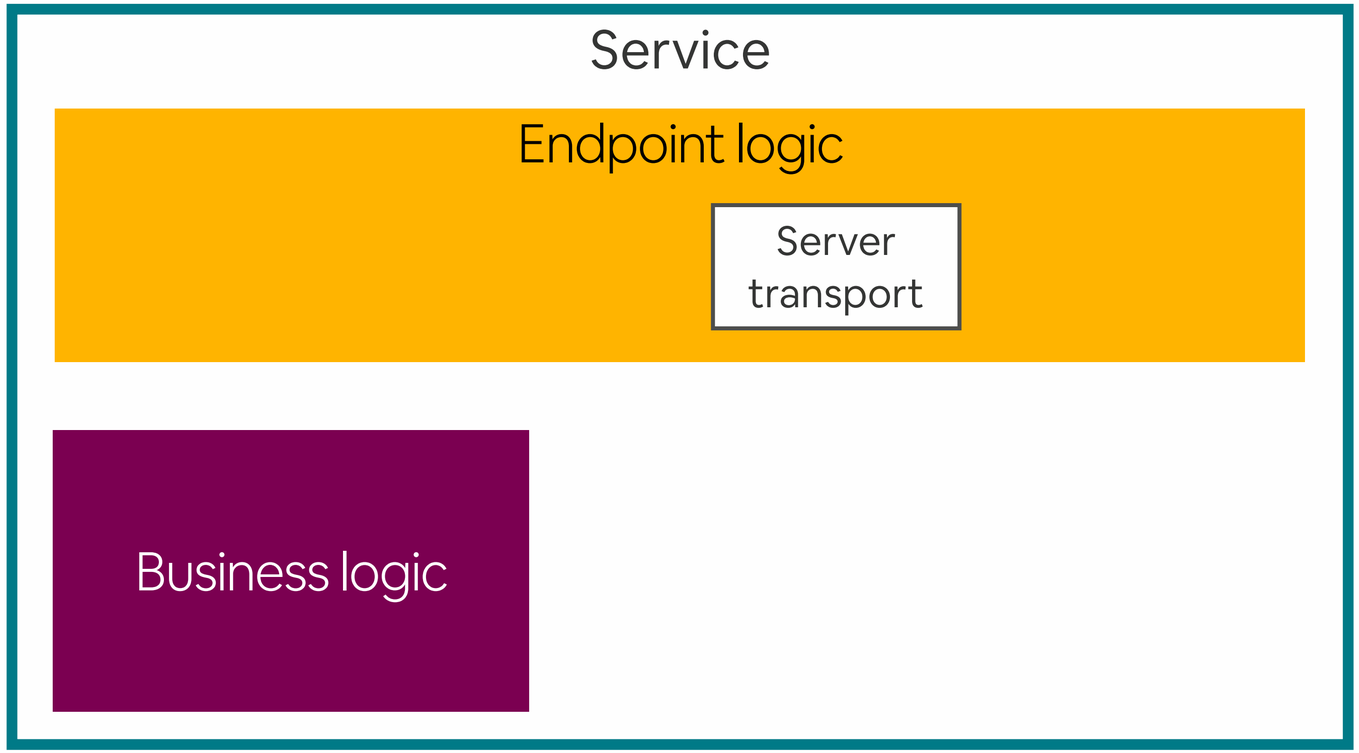

Service & client setup

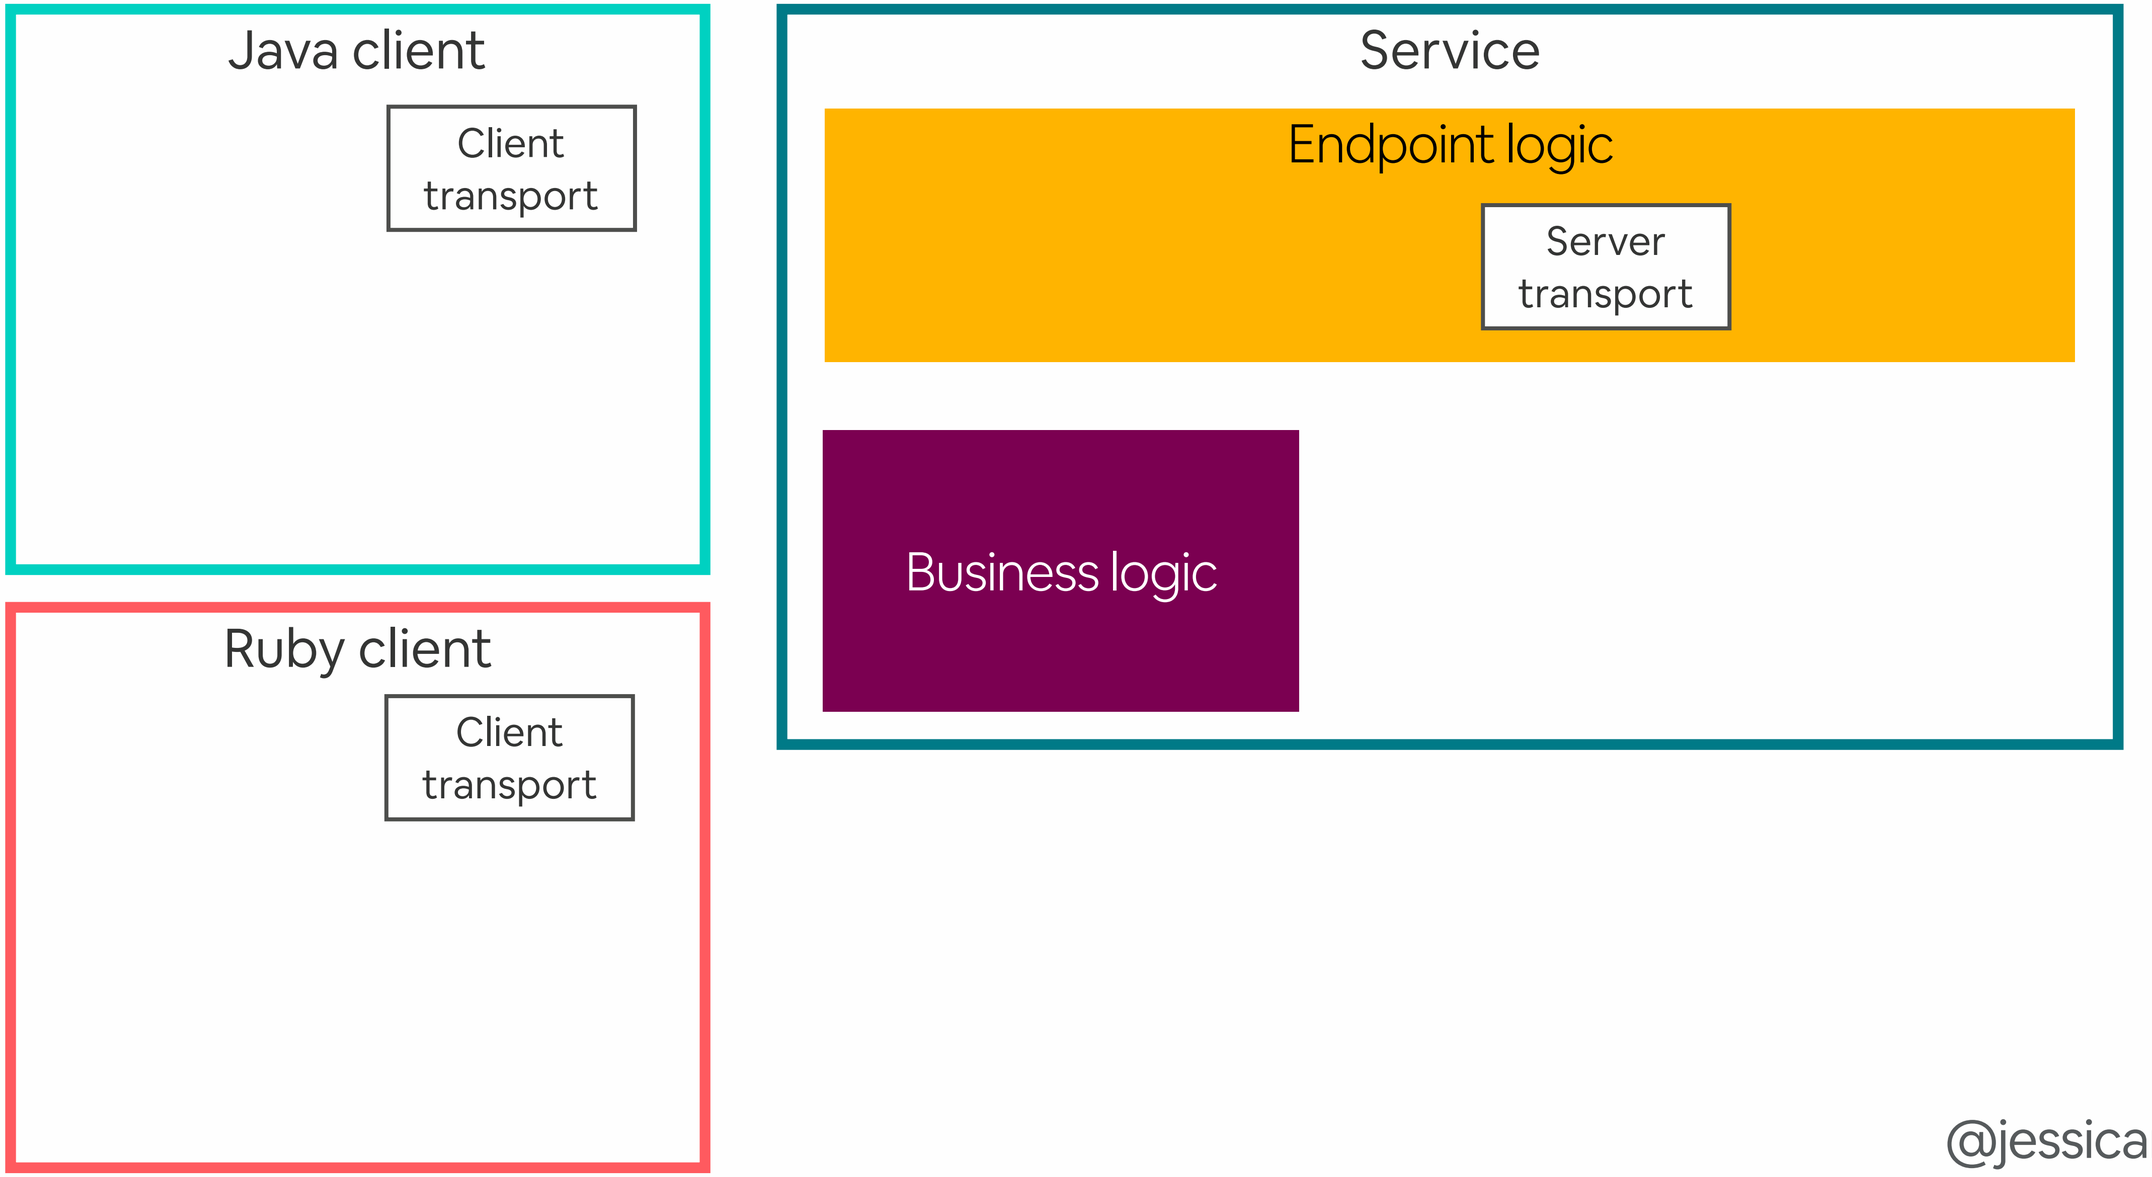

Service & client setup

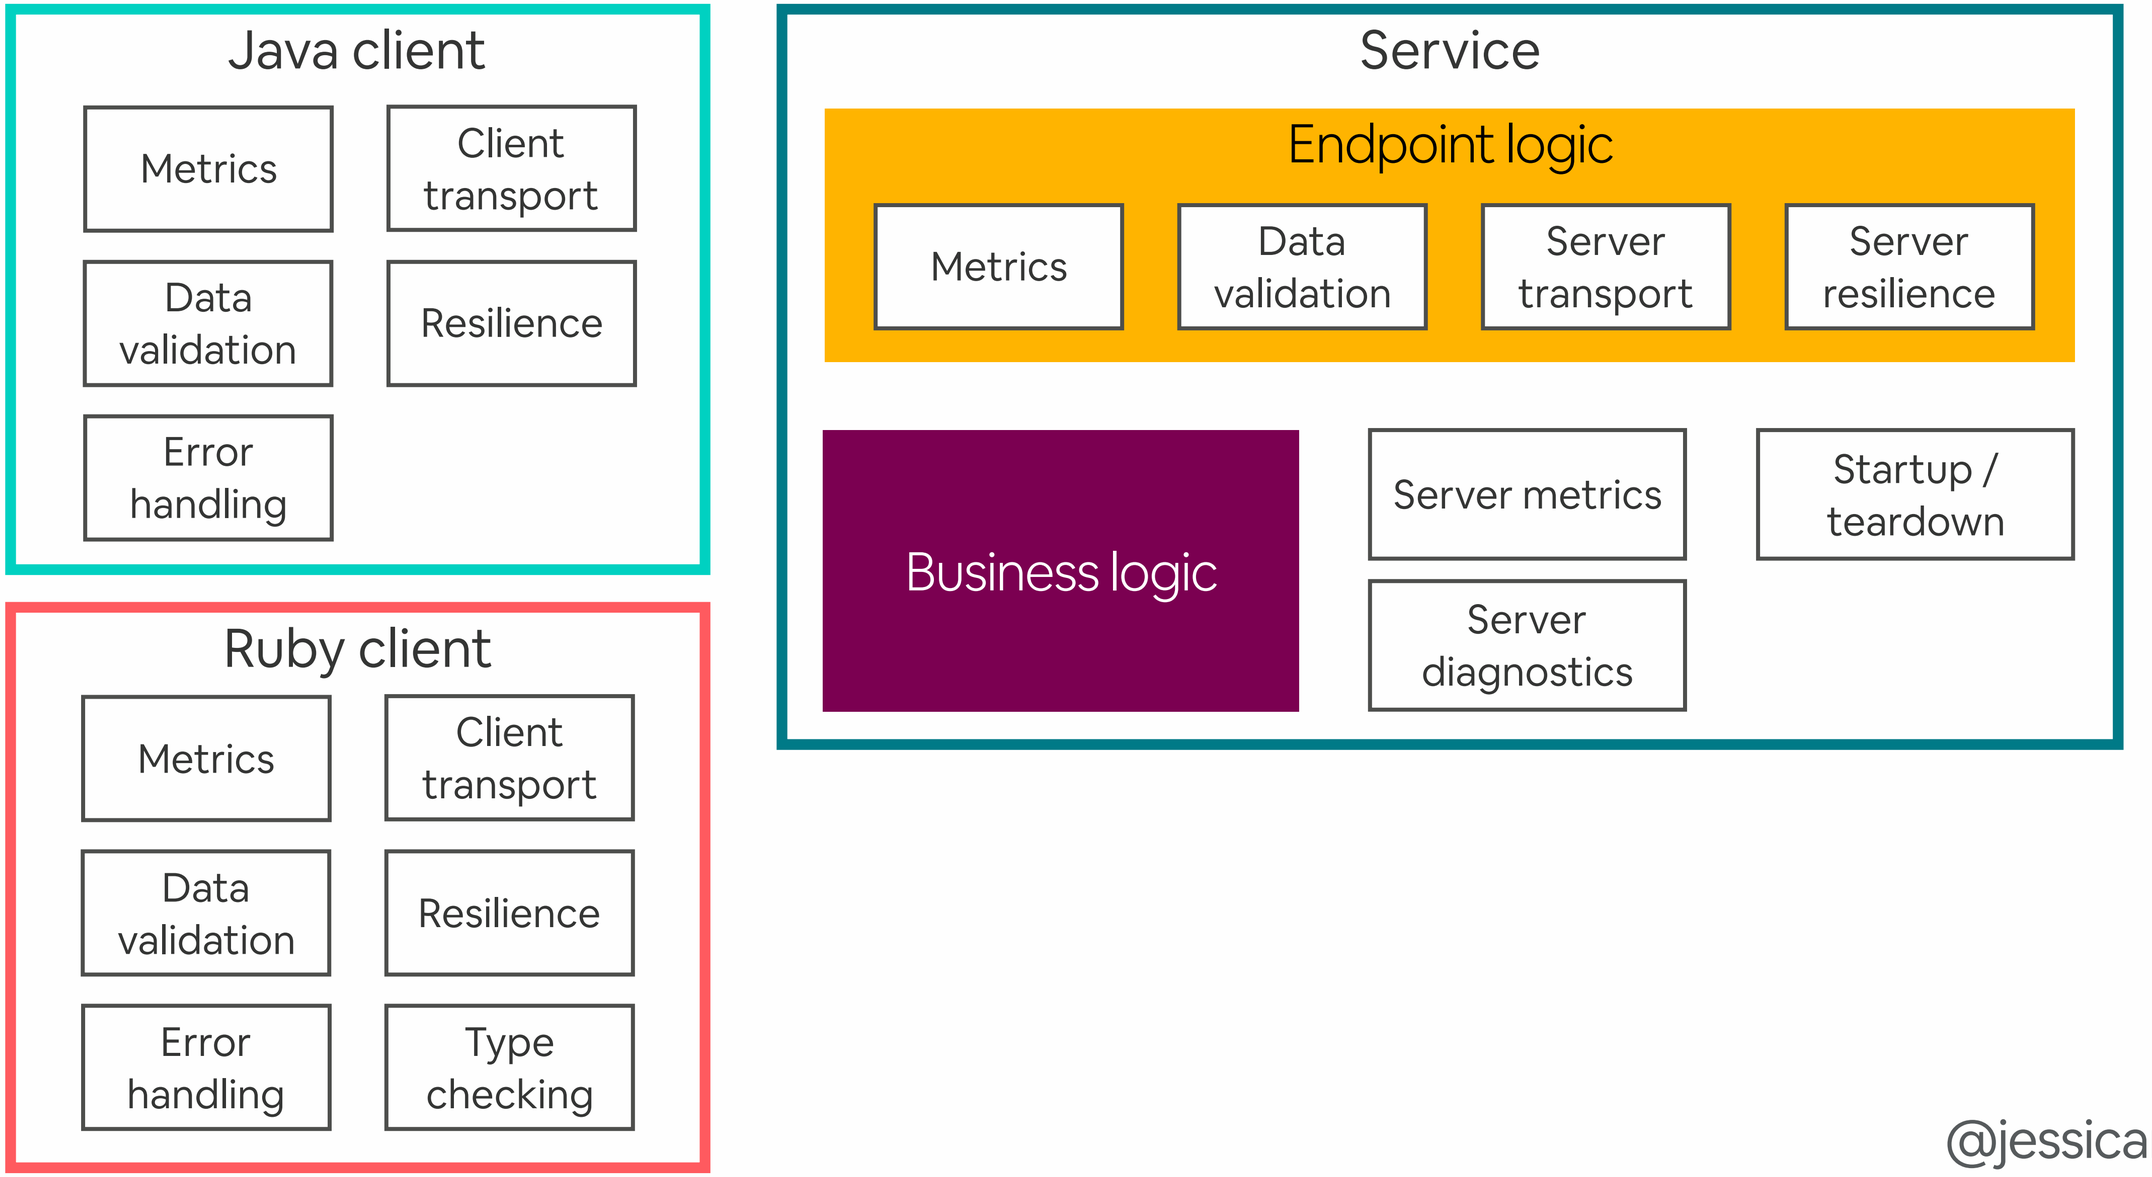

Service & client setup

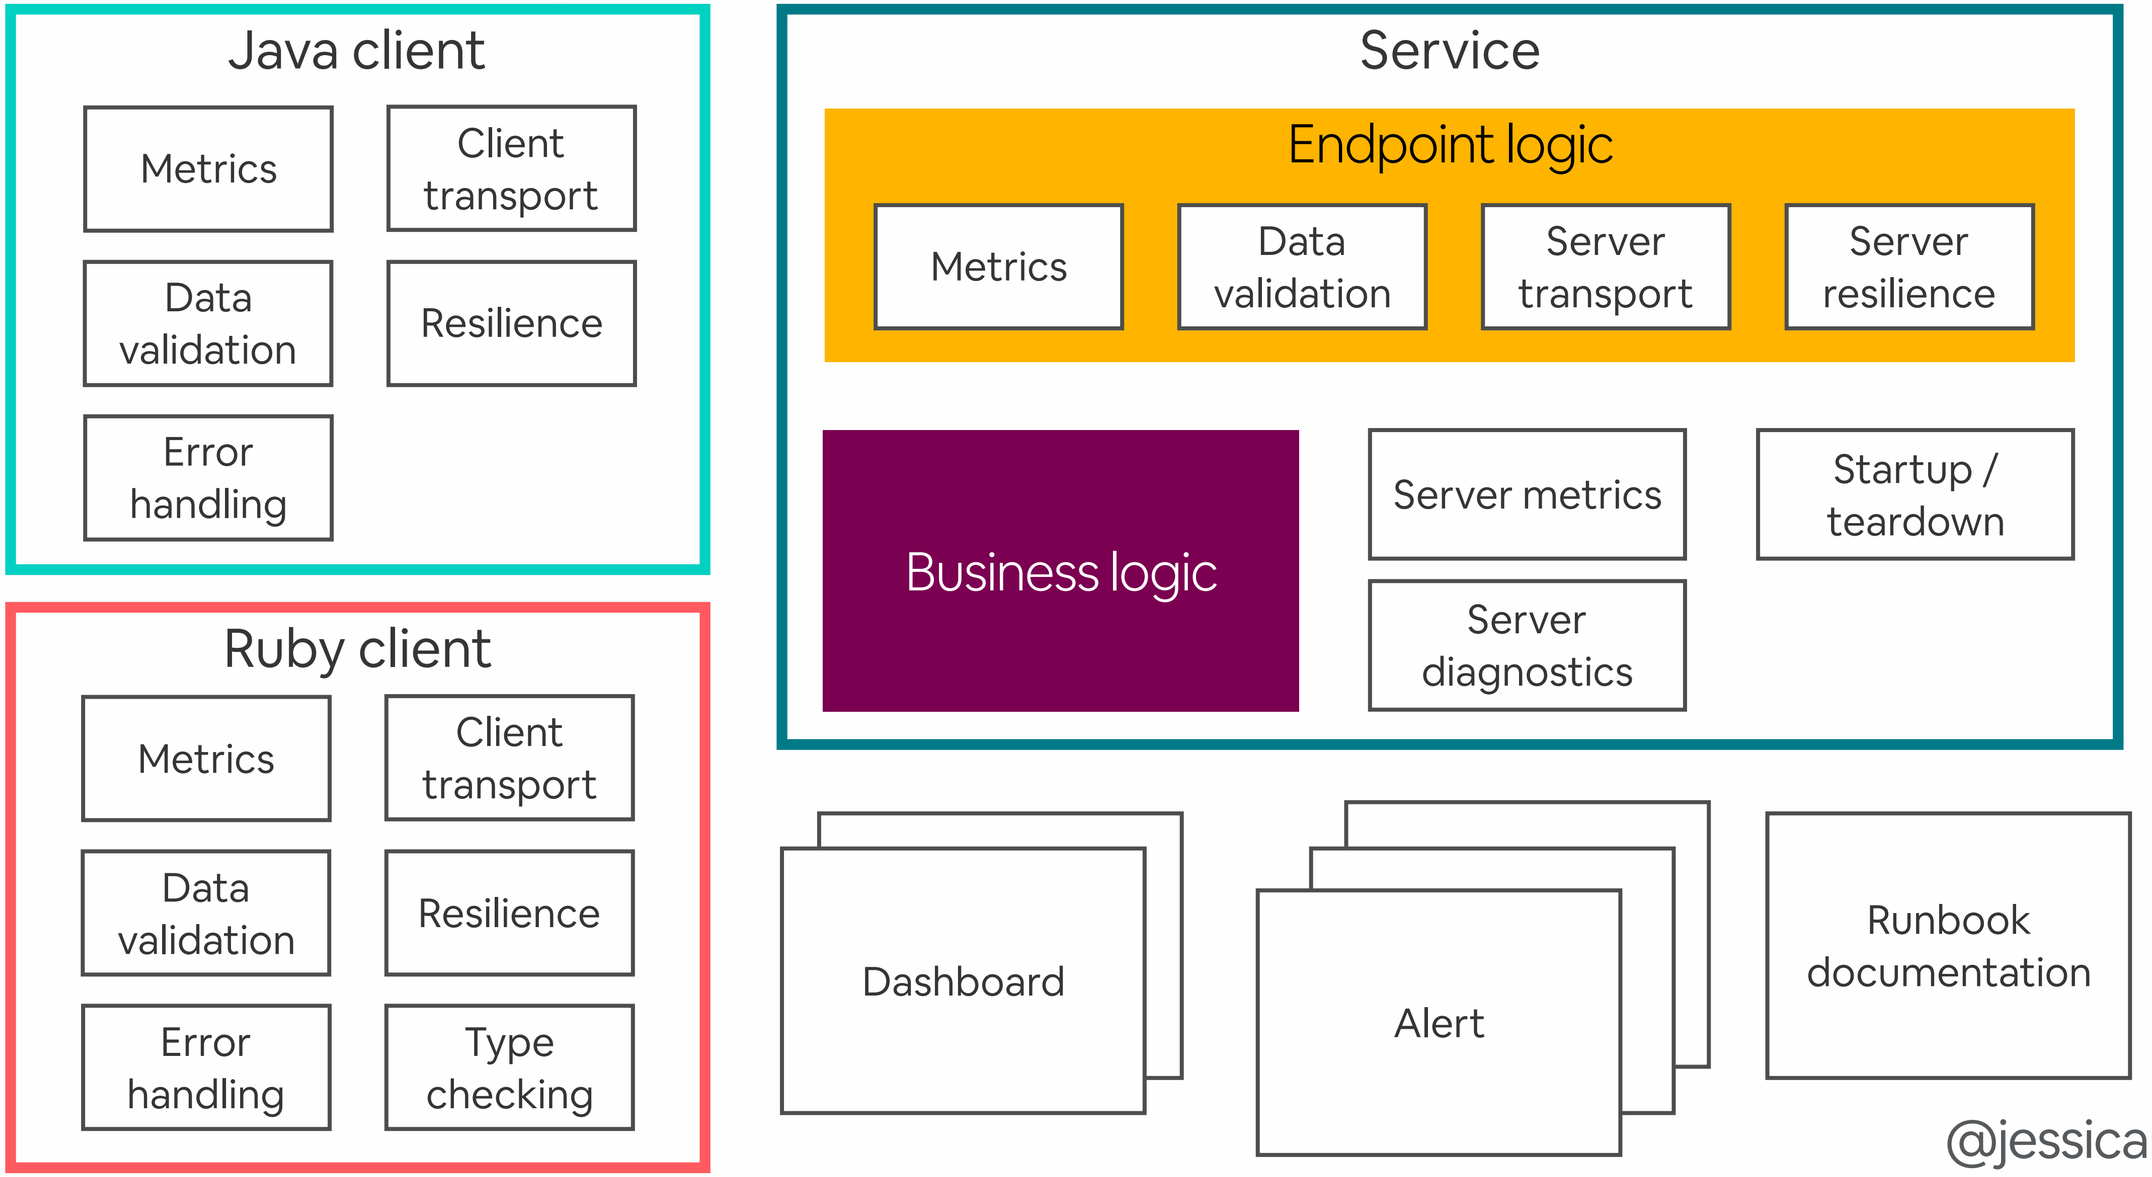

Service & client setup

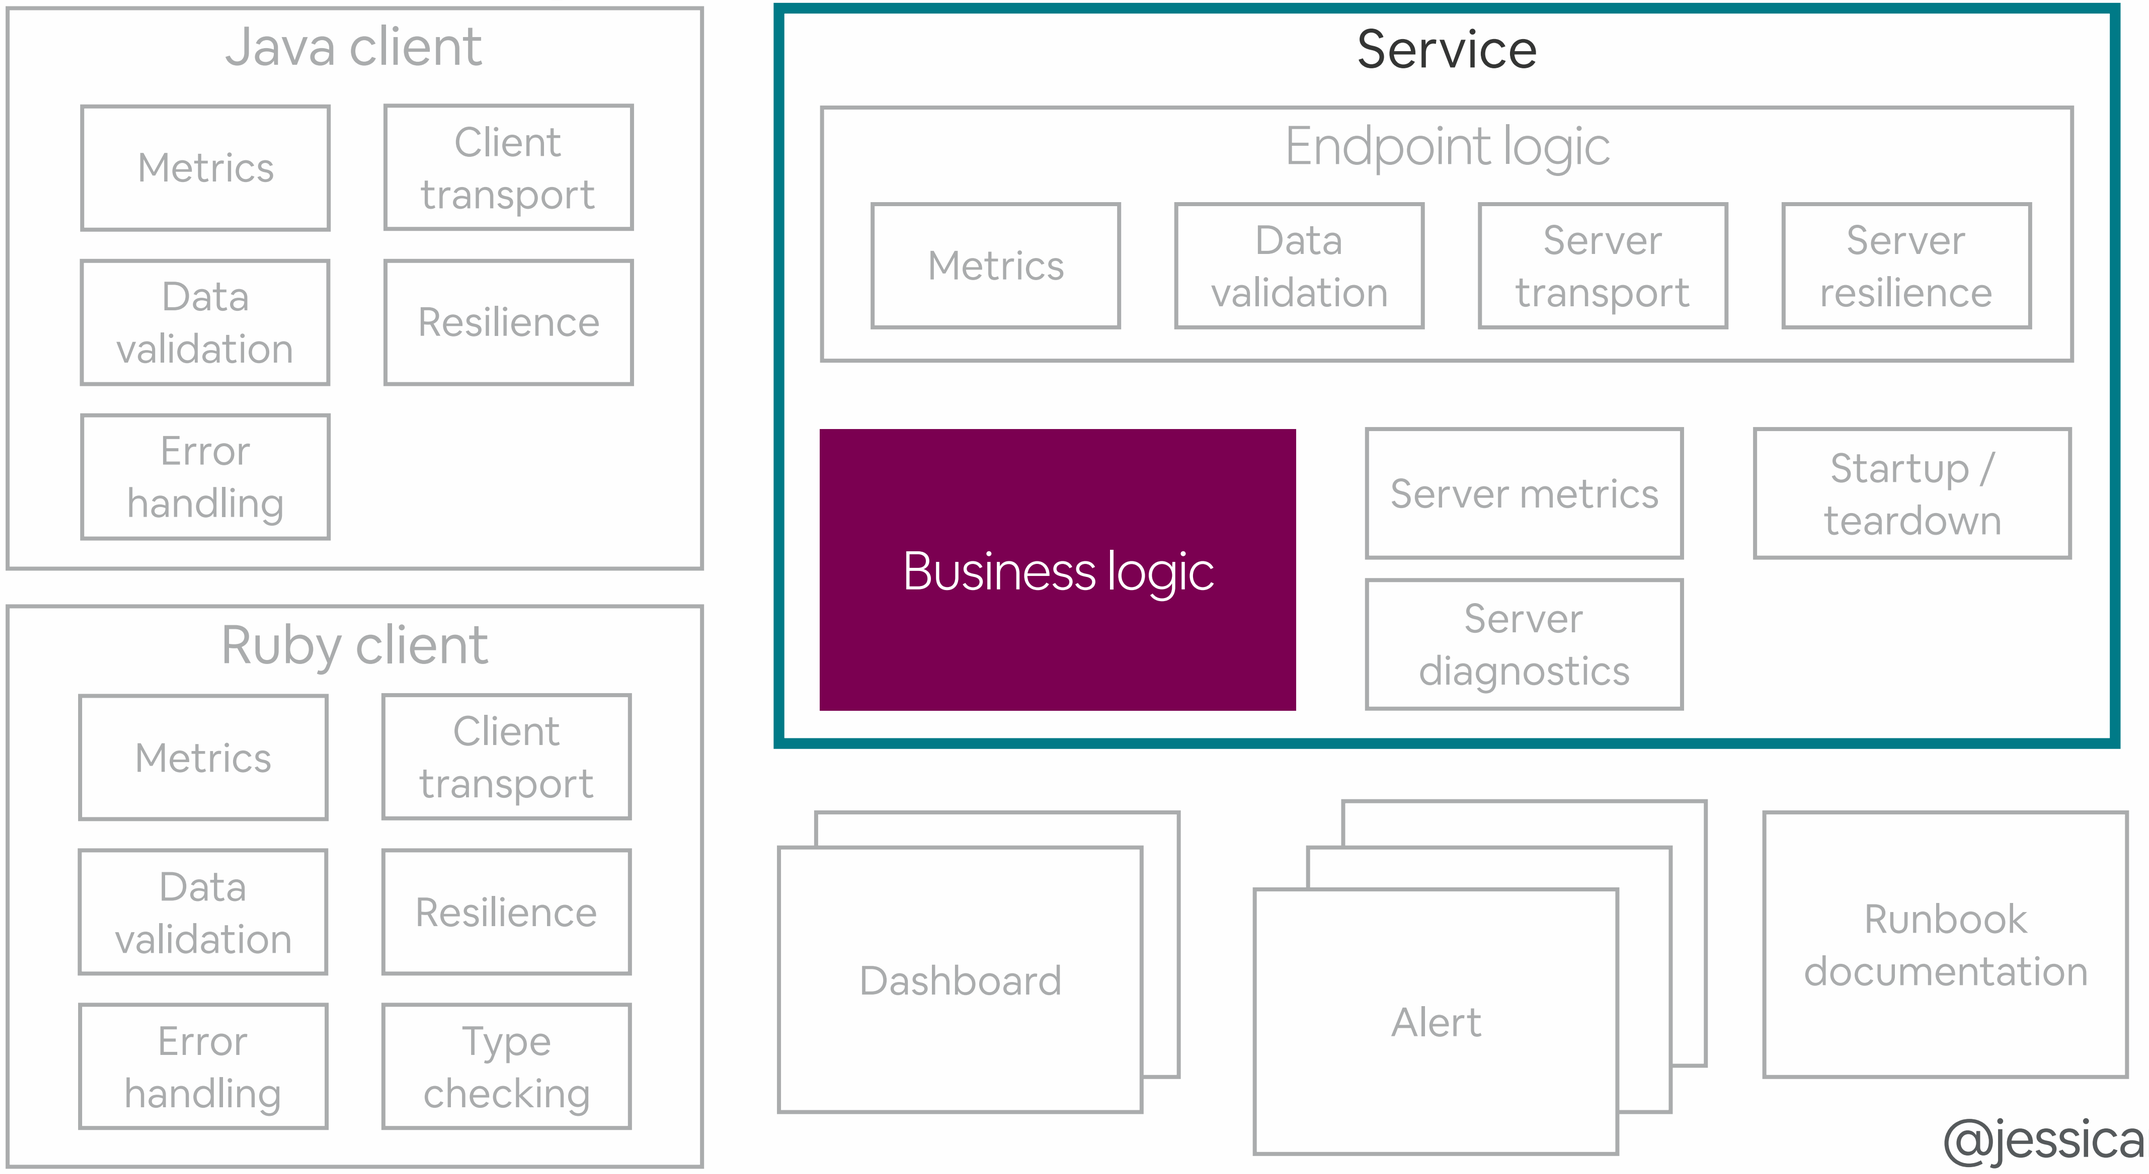

Frameworks using Thrift IDL

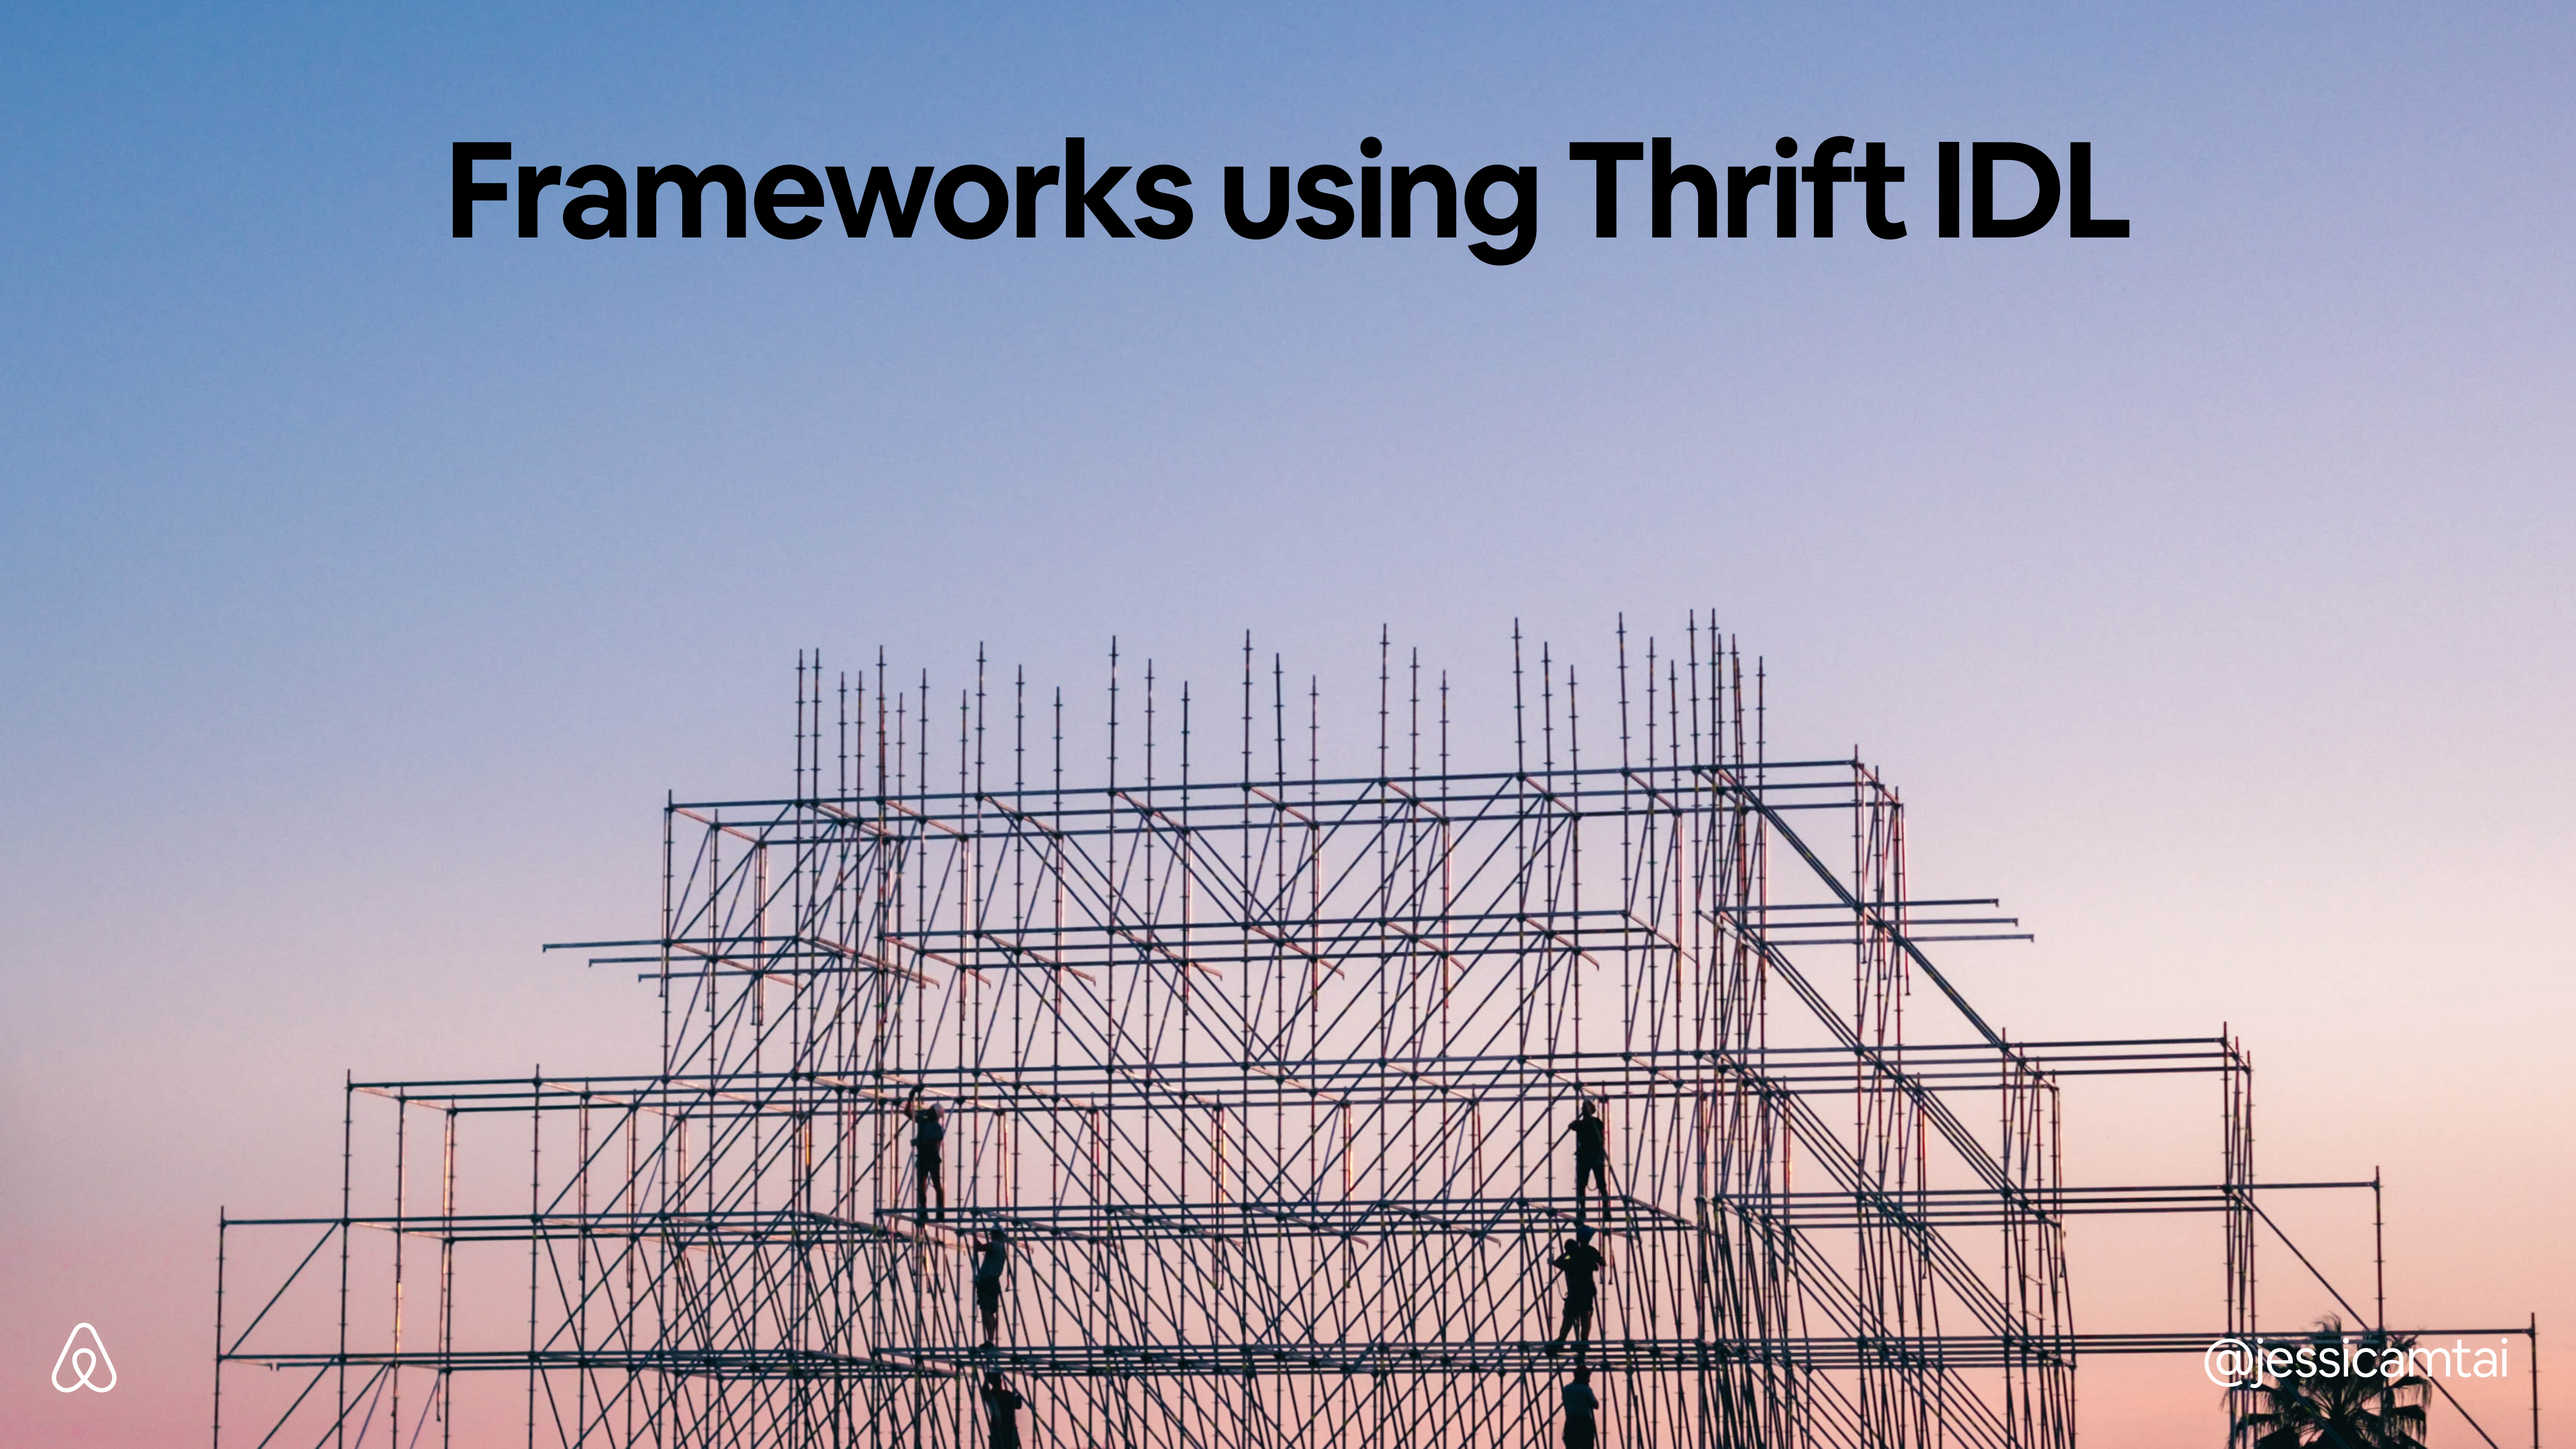

IDL

Java client

- Metrics
- Client transport
- Data validation
- Resilience
- Error handling

Ruby client

- Metrics
- Client transport
- Data validation
- Resilience
- Error handling
- Type checking

Service

Endpoint logic

- Metrics
- Data validation
- Server transport
- Server resilience

Business logic

- Server metrics
- Server diagnostics
- Startup / teardown

- Dashboard
- Alert
- Runbook documentation

@jessicamtai

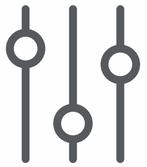

Thrift IDL

API FRAMEWORK

```
/** Batch request for demo data */
struct LoadSomeDataRequest {
  1: optional set<i64> ids
  /** Some extra context baz */
  2: optional bool fooBar (personal)
}
```

Thrift IDL

API FRAMEWORK

```
/** Batch request for demo data */
struct LoadSomeDataRequest {
  1: optional set<i64> ids (personal)
  /** Some extra context baz */
  2: optional bool fooBar
}

/** id to data response */
struct LoadSomeDataResponse {
  1: optional map<i64, SomeData> data
}
```

Thrift IDL

API FRAMEWORK

```
/** /loadSomeData batch endpoint */
LoadSomeDataResponse loadSomeData
(1: LoadSomeDataRequest request)
throws (1: SomeException exception1)
(accept_replay = "true", rate_limit = "true")
```

Service API Explorer

RENDERED | RAW JSON

demo

Demo Service

JAVA

slacks: #demo-goalie

dashboards: Standard Service Metrics, Standard Client Metrics, Demo Dashboard

doc link: Demo tech design doc

description: The Demo Service owns test data in MySQL and Redis.

Thrift annotations

Search Demo schema (51 structs, 478 fields) ...

loadSomeData RPC

/loadSomeData batch endpoint

LoadSomeDataRequest → LoadSomeDataResponse

@jessicamtai

demo.data.LoadSomeDataRequest STRUCT

Demo Service

JAVA

slacks: #demo-goalie

dashboards: Standard Service Metrics, Standard Client Metrics, Demo Dashboard

doc link: Demo tech design doc

description: The Demo Service owns test data in MySQL and Redis.

Search Demo schema (51 structs, 478 fields) ...

loadSomeData RPC

/loadSomeData batch endpoint

LoadSomeDataRequest → LoadSomeDataResponse

demo_data.LoadSomeDataRequest

STRUCT

Batch request for demo data

ID	Req.	Name	Type	Document
1	*	ids	set<i64>	self-explanatory
2		fooBar	boolean	Some extra context baz

Block comments from .thrift file

@jessicamtai

FAIL FAST

@jessicamtai

Separate async worker thread pools

Graceful degradation

@jessicamtai

Testing & deploying

TIMELINE

Local dev

Staging

Diffy

Canary

Production

Testing & deploying

TIMELINE

Dev environment supports shared dev services

Local dev — Staging — Diffy — Canary — Production

@jessicamtai

Testing & deploying

TIMELINE

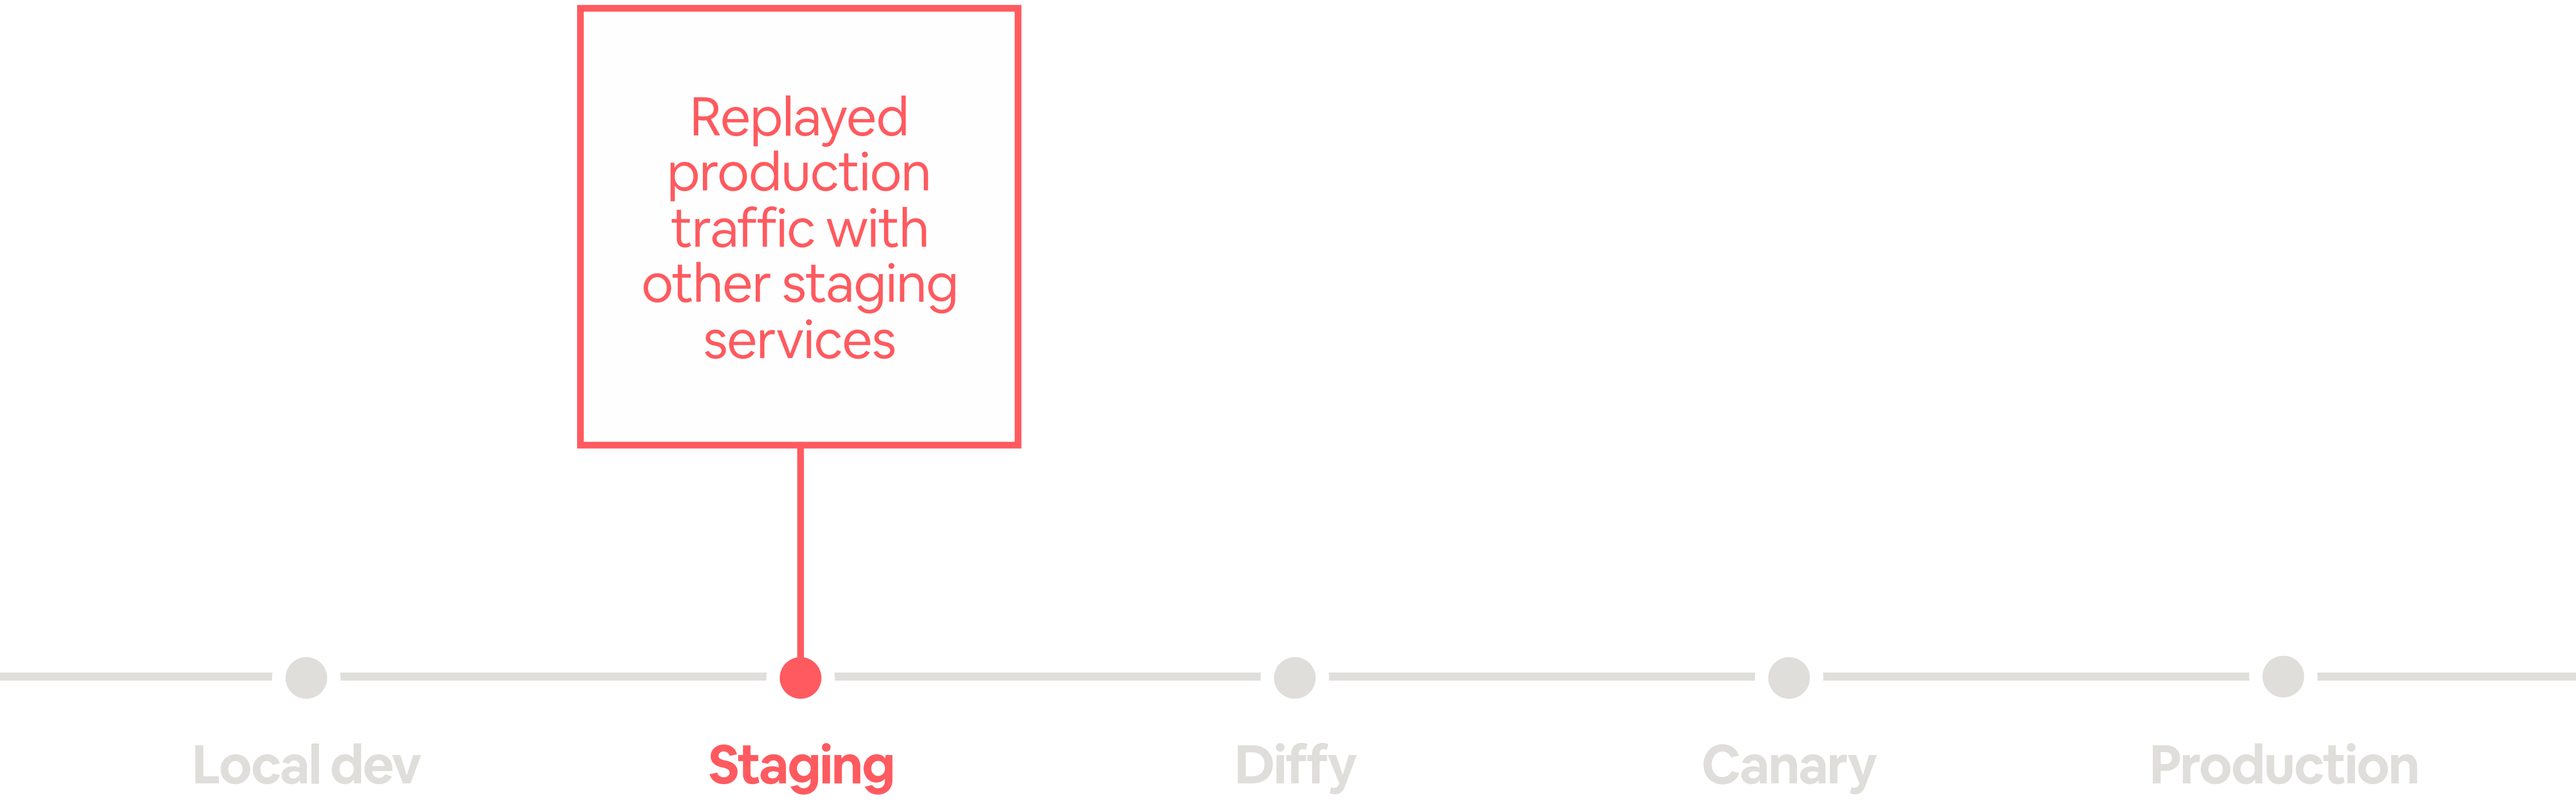

Testing & deploying

TIMELINE

Compare responses from staging against production

Local dev · Staging · **Diffy** · Canary · Production

@jessicamtai

Regression Testing

DIFFY

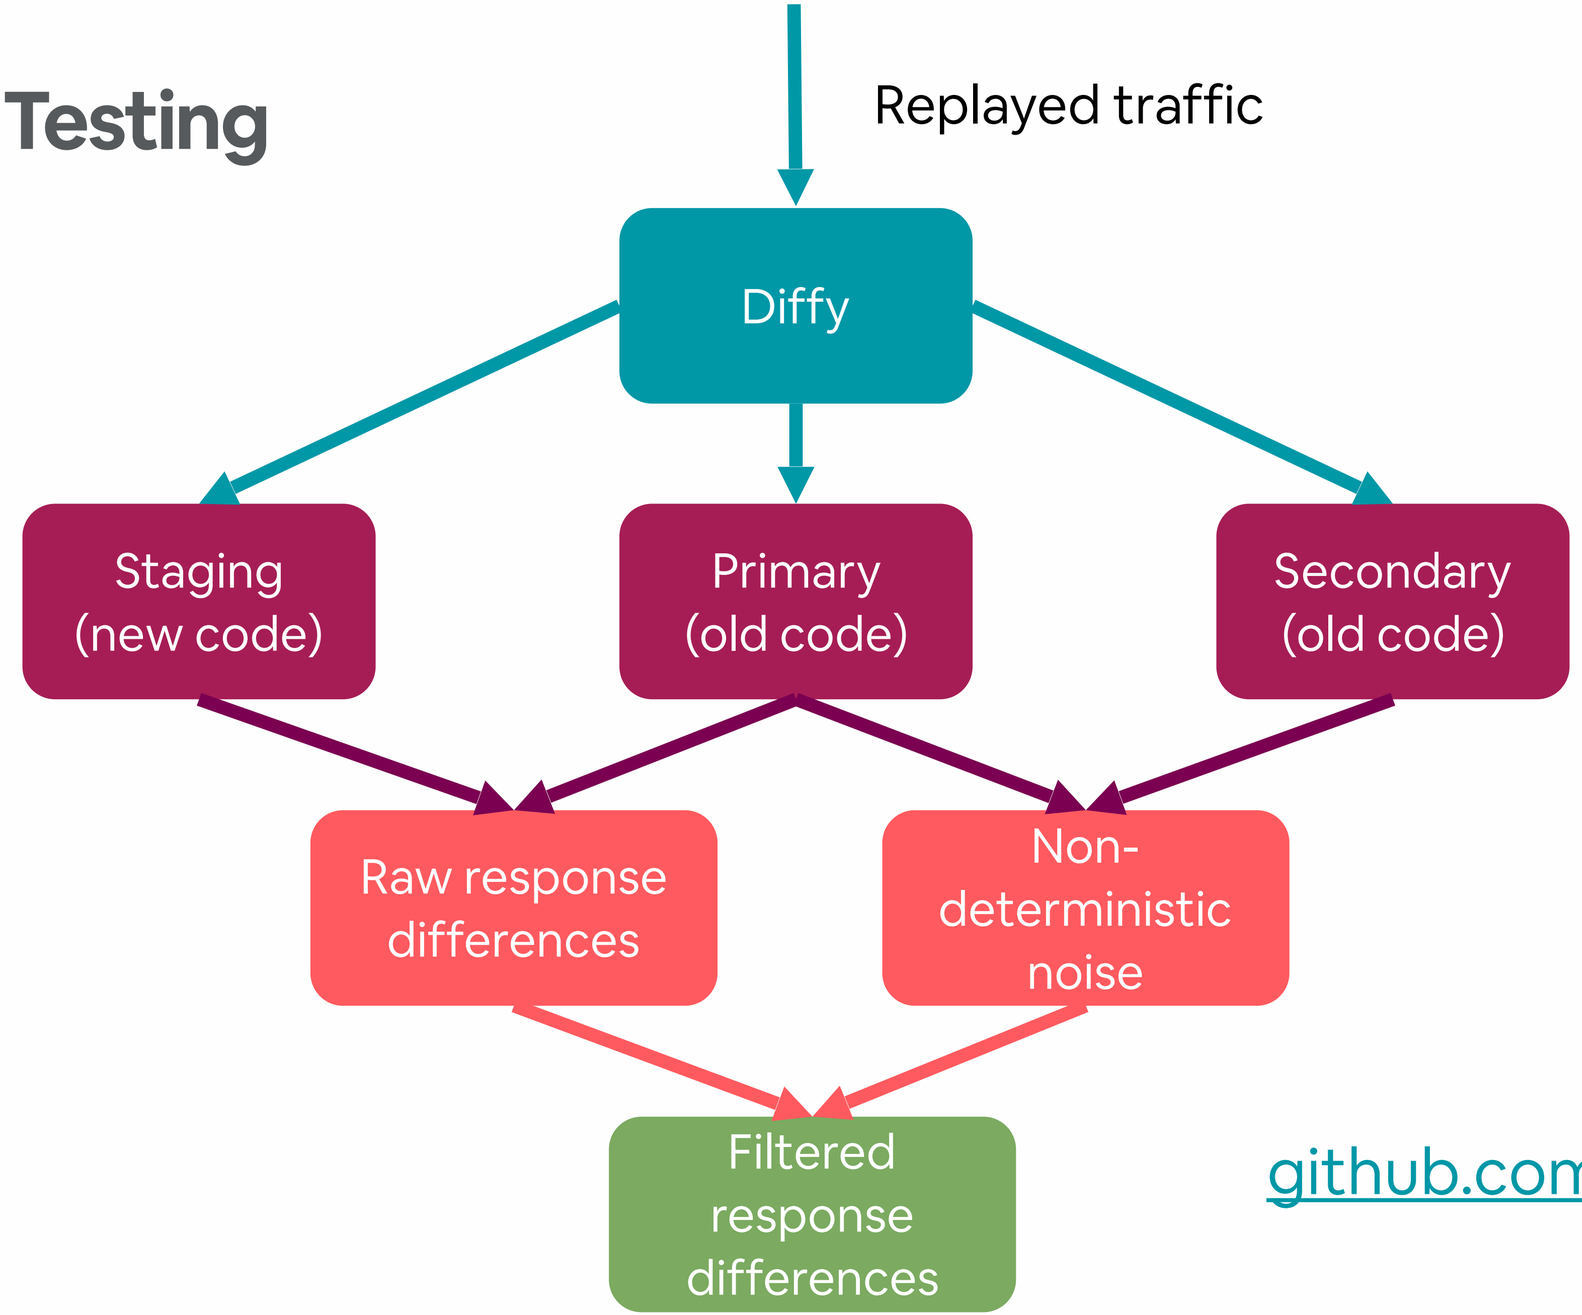

github.com/twitter/diffy

@jessicamtai

Testing & deploying

TIMELINE

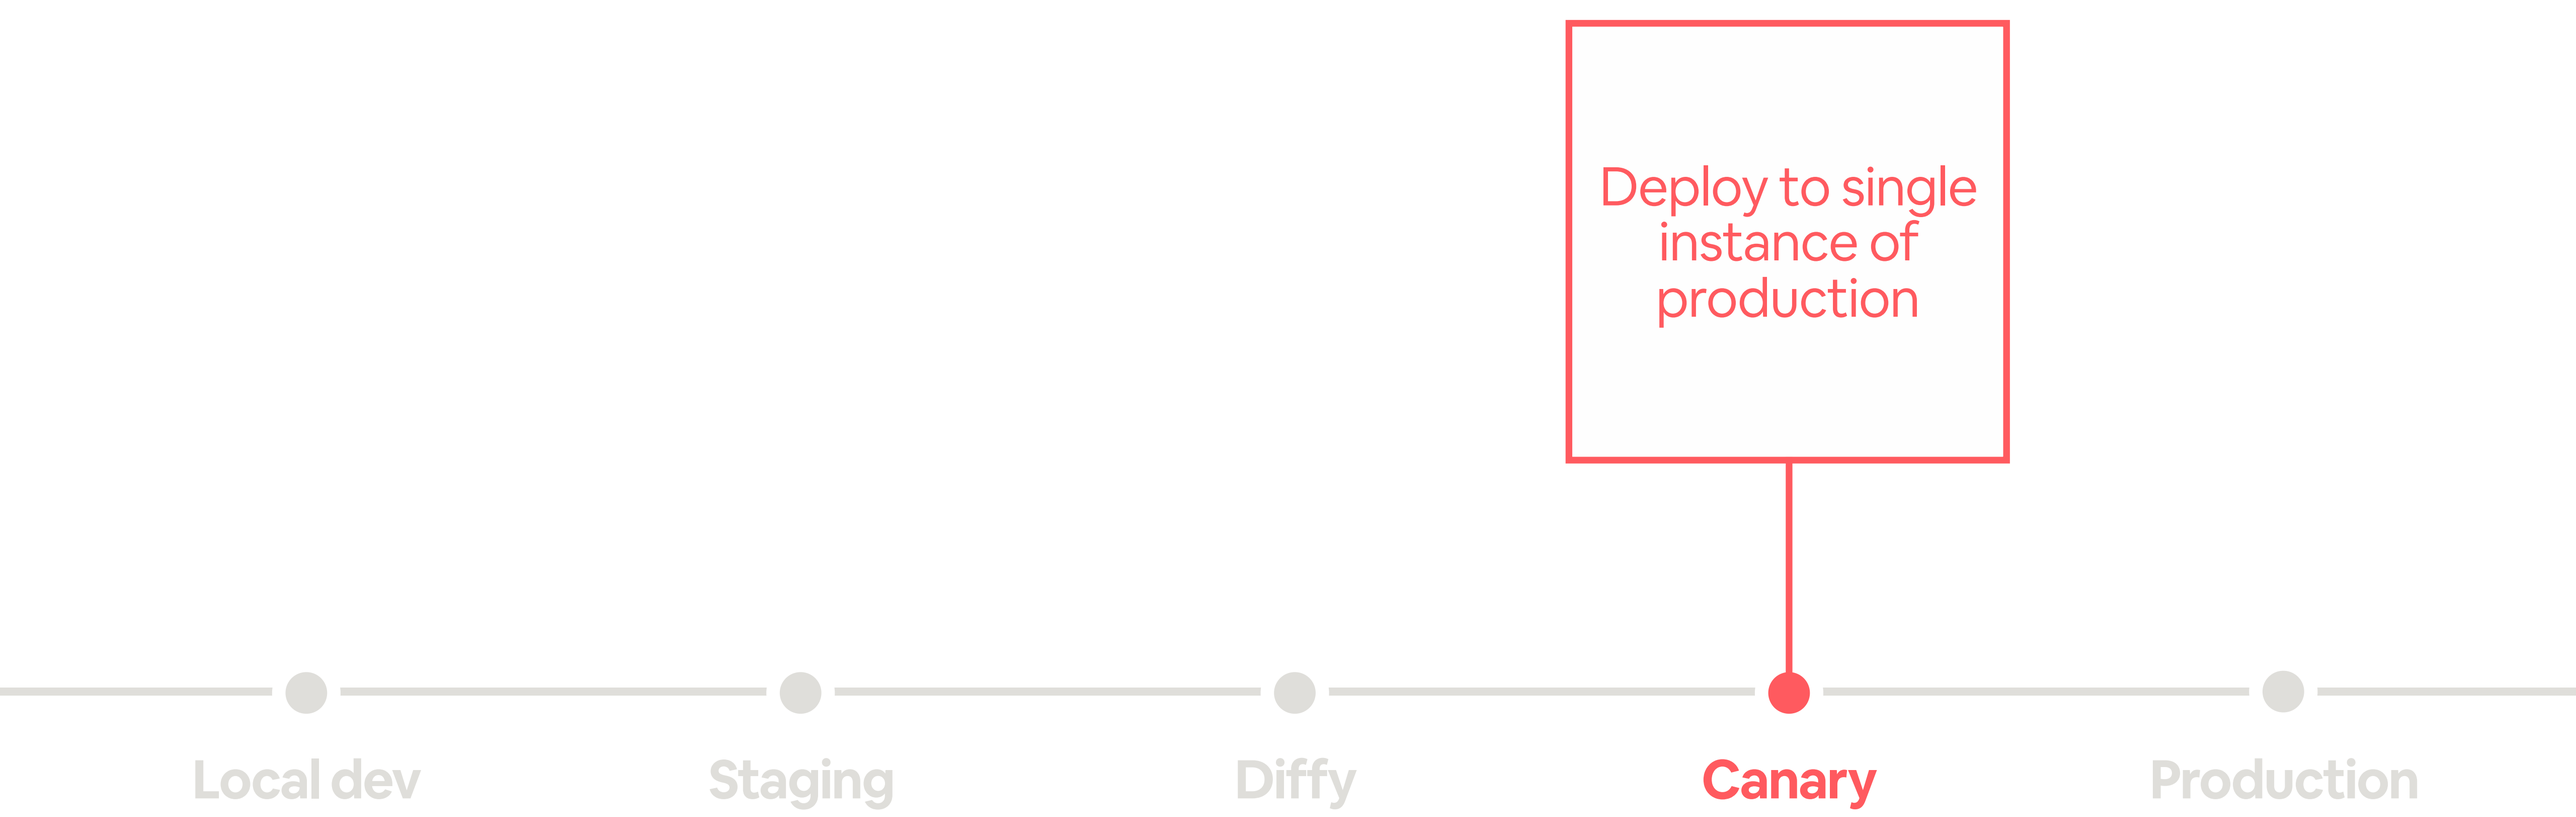

Testing & deploying

TIMELINE

Confidently deploy to prod

Local dev · Staging · Diffy · Canary · **Production**

@jessicamtai

Standard Client-Side Metrics

Add Graphs +

$service banana | $caller_chef_role * | $caller_k8s_app * | $caller_host * | $method * | $tie

Search events to overlay...

Show 1h The Past Hour

Request Rate (QPS) by Method

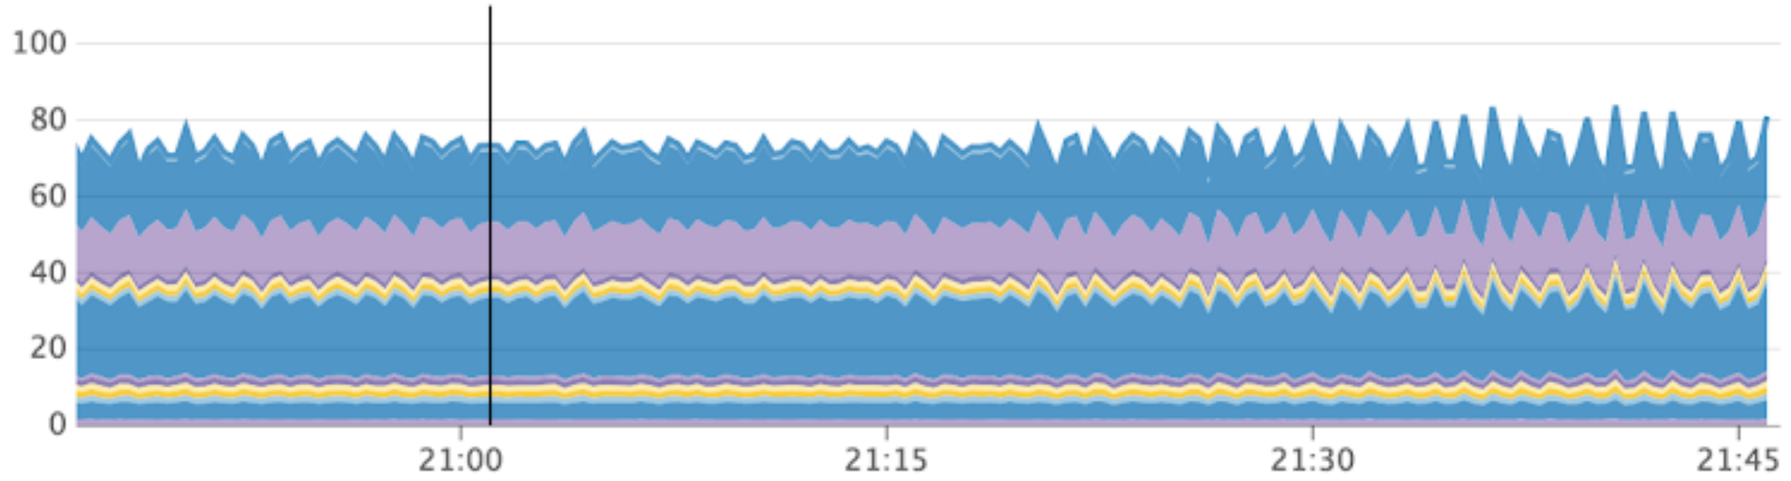

Error Rate Week-over-week

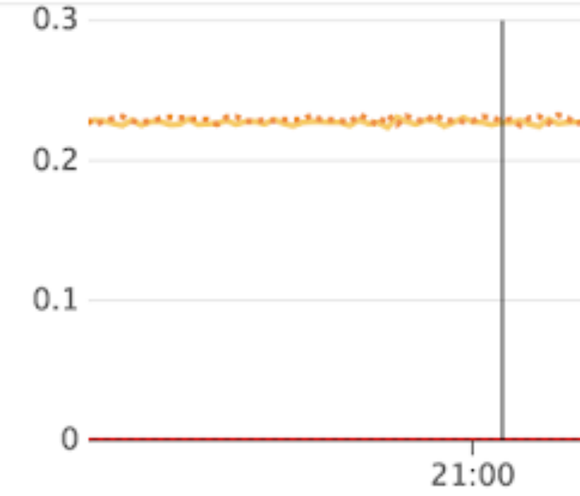

Requests by Method

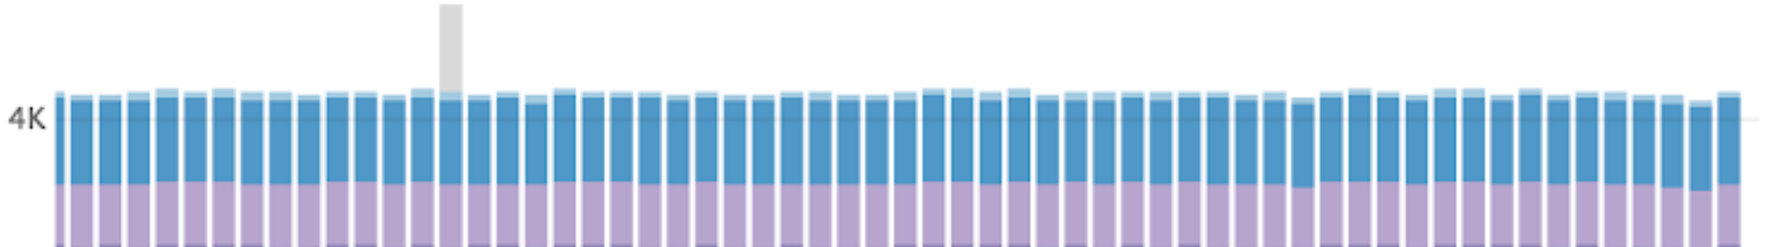

Requests by Caller

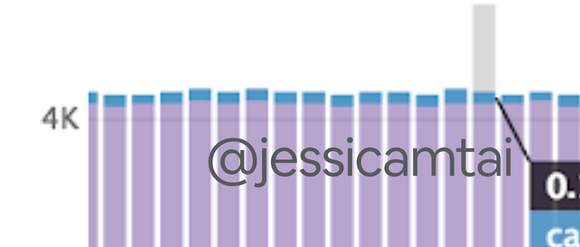

@jessicamtai

Standard Client-Side Metrics

Add Graphs +

$service banana

$caller_chef_role *

$caller_k8s_ap

Search events to overlay...

Request Rate (QPS) by Method

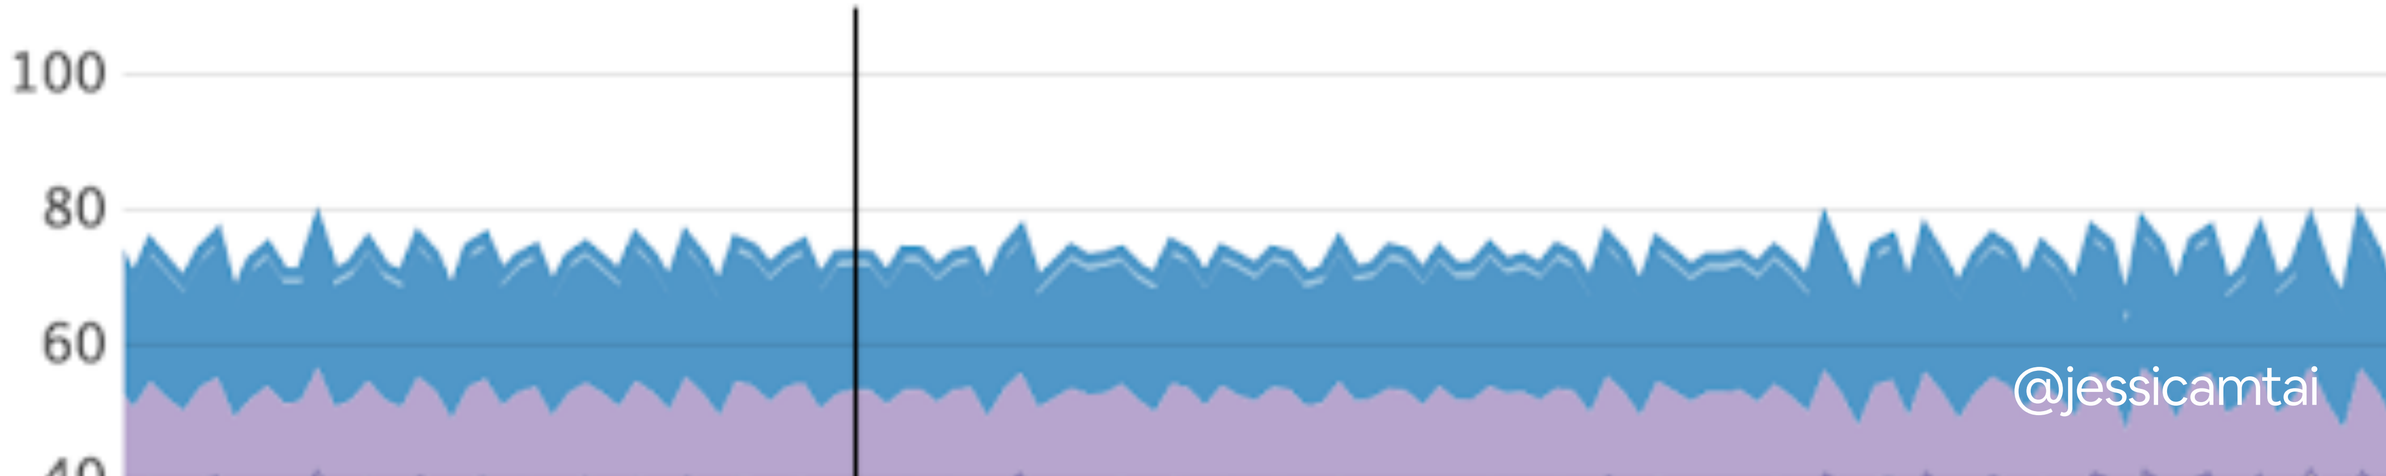


EMPOWERING THE MIGRATION
@jessicamtai

2016:
One small infra team

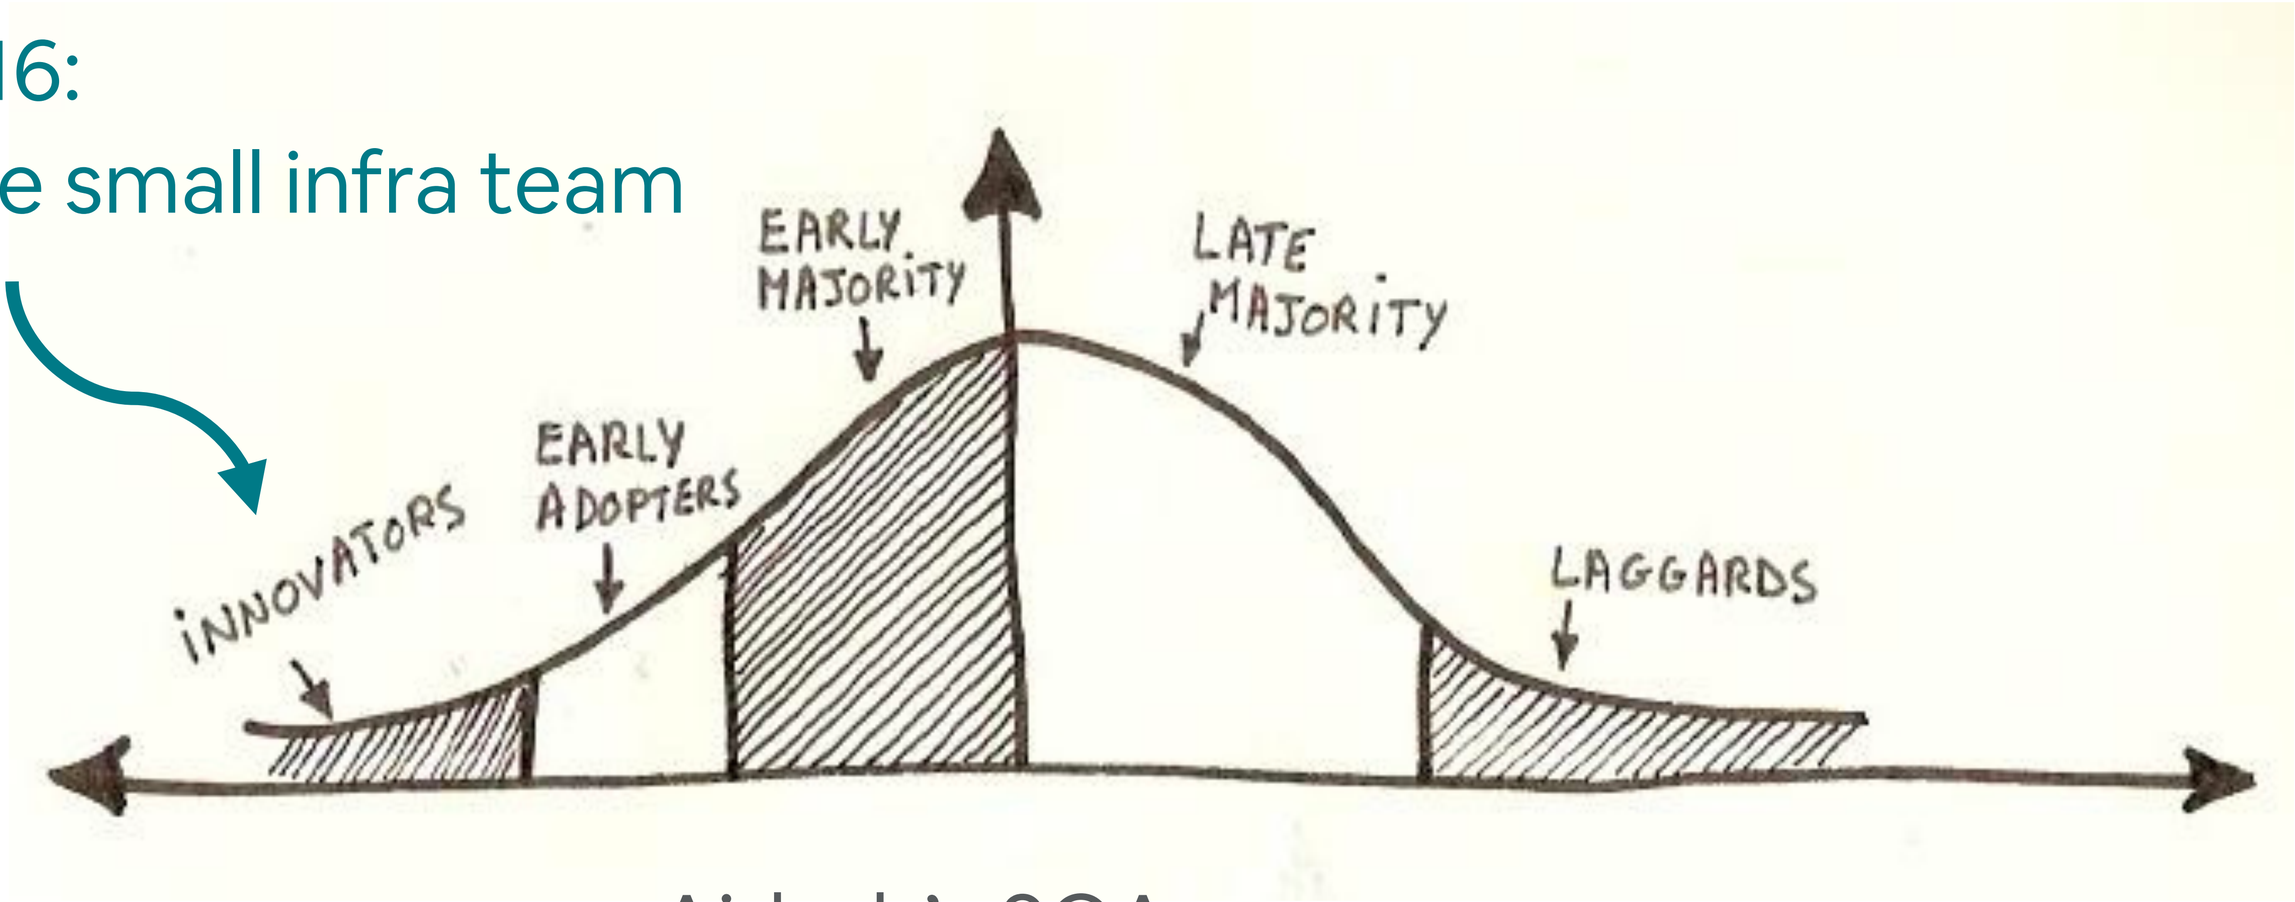

Airbnb's SOA progress


Product culture:
ship things quickly
@jessicamtai


Org challenges:
service building in parallel
@jessicamtai

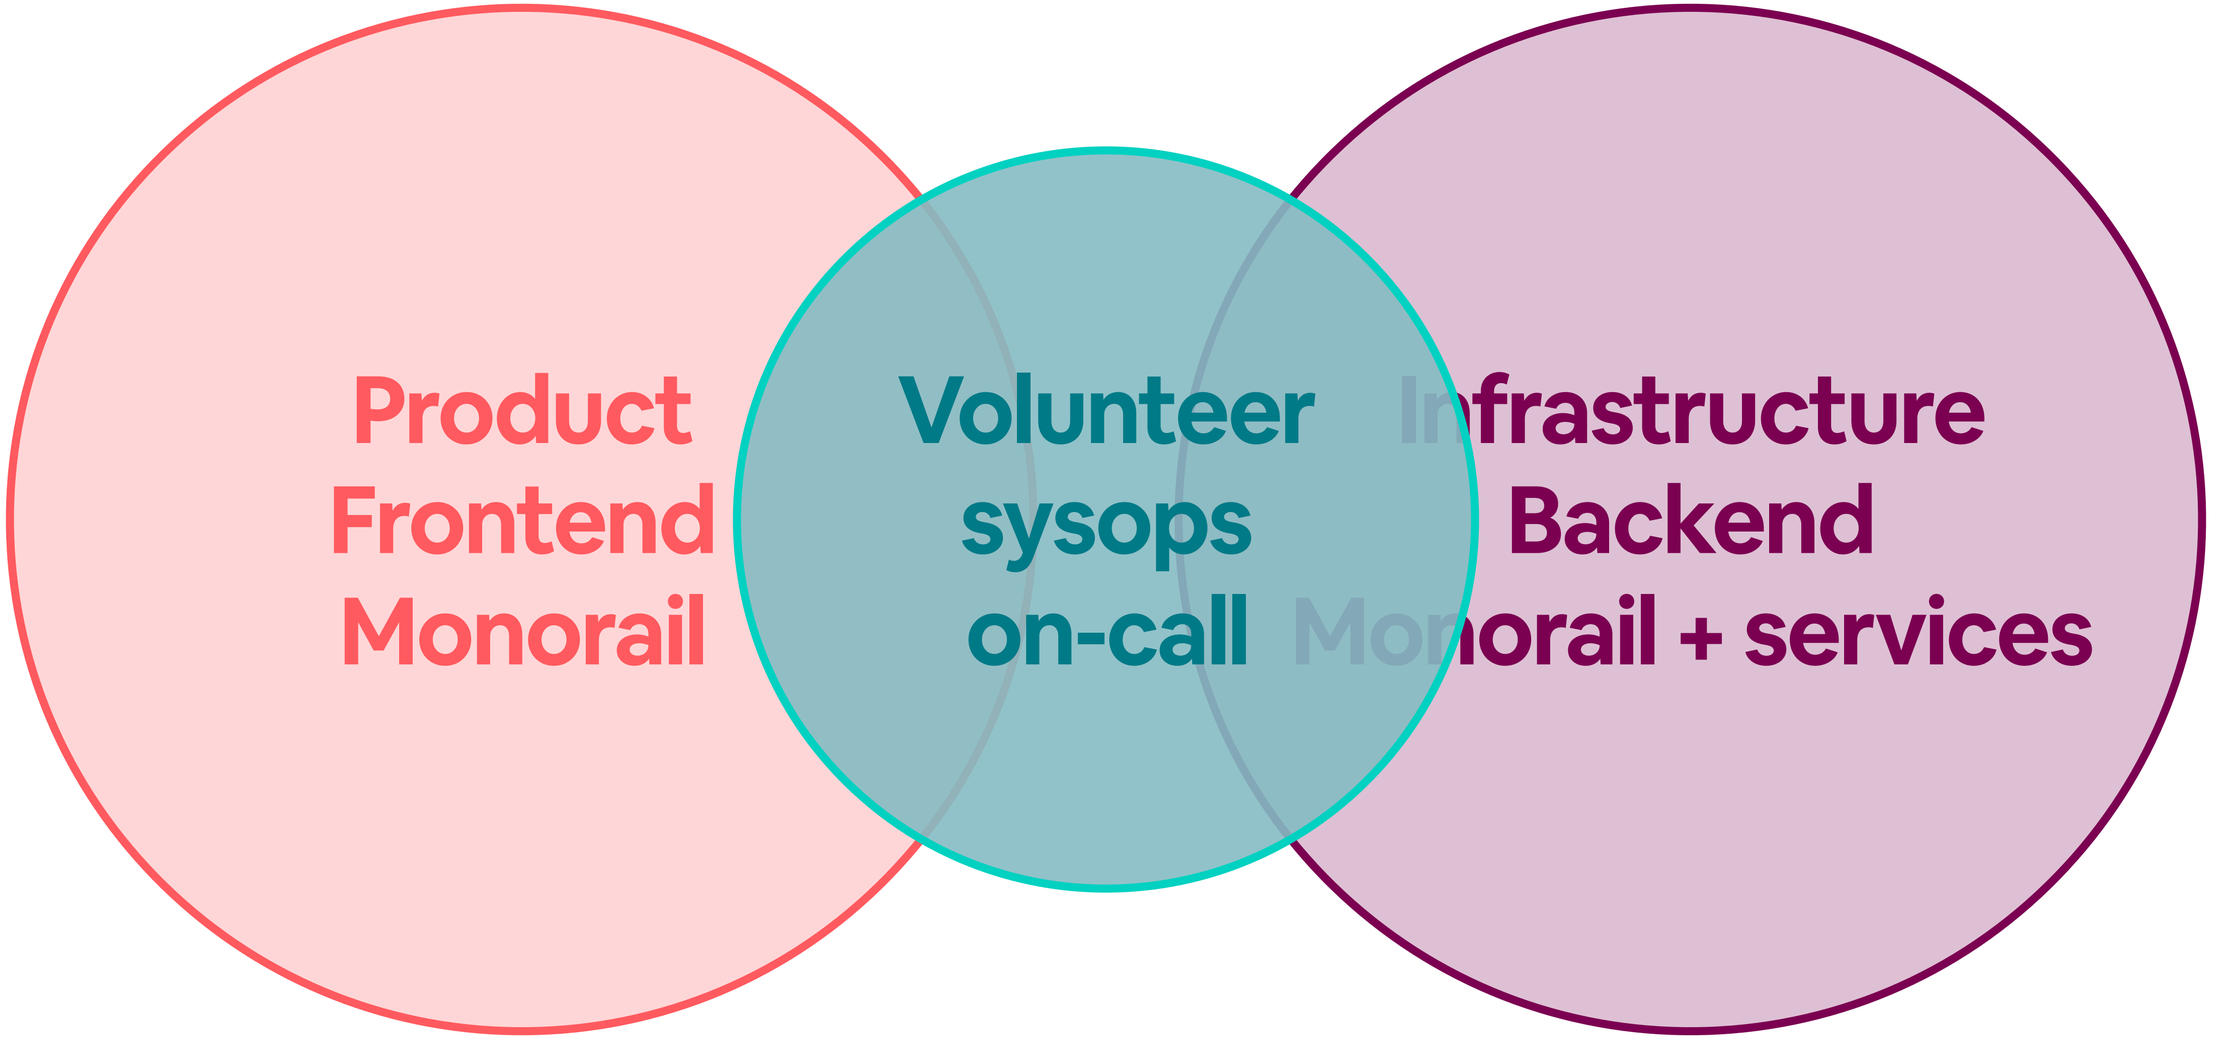
Product
Frontend
Monorail
Volunteer
sysops
on-call
Infrastructure
Backend
Monorail + services

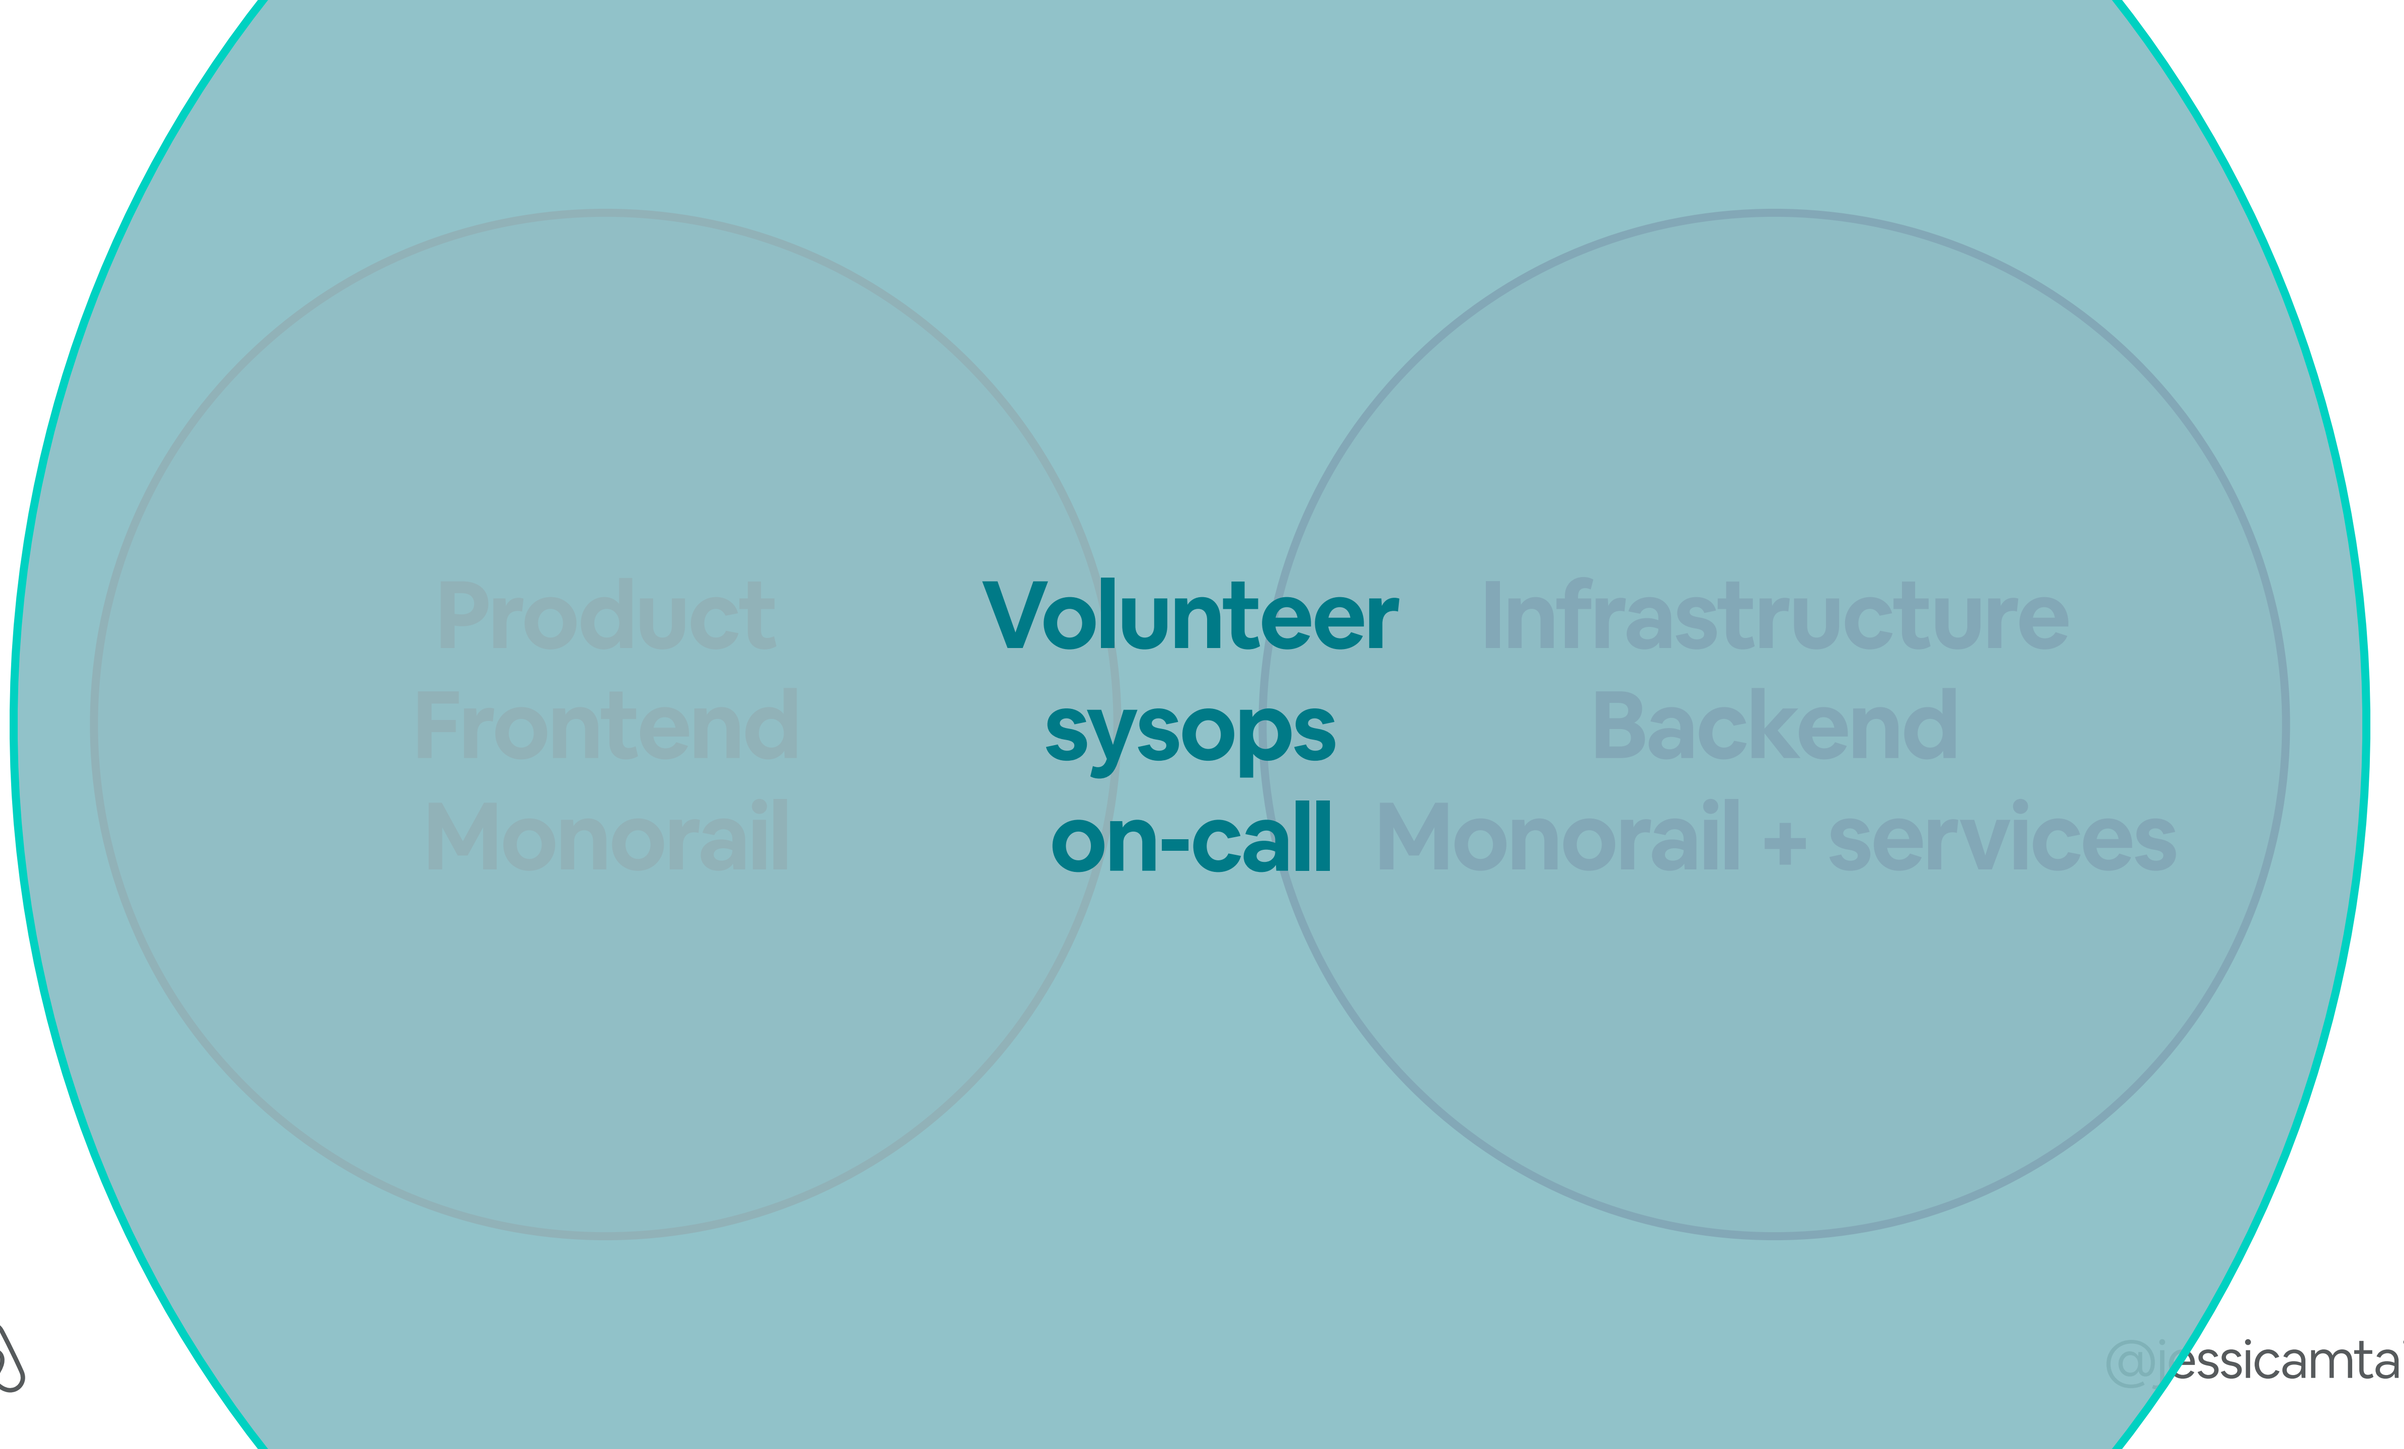
Product
Frontend
Monorail
Volunteer
sysops
on-call
Infrastructure
Backend
Monorail + services

On-call rotation per team

SERVICE OWNERSHIP

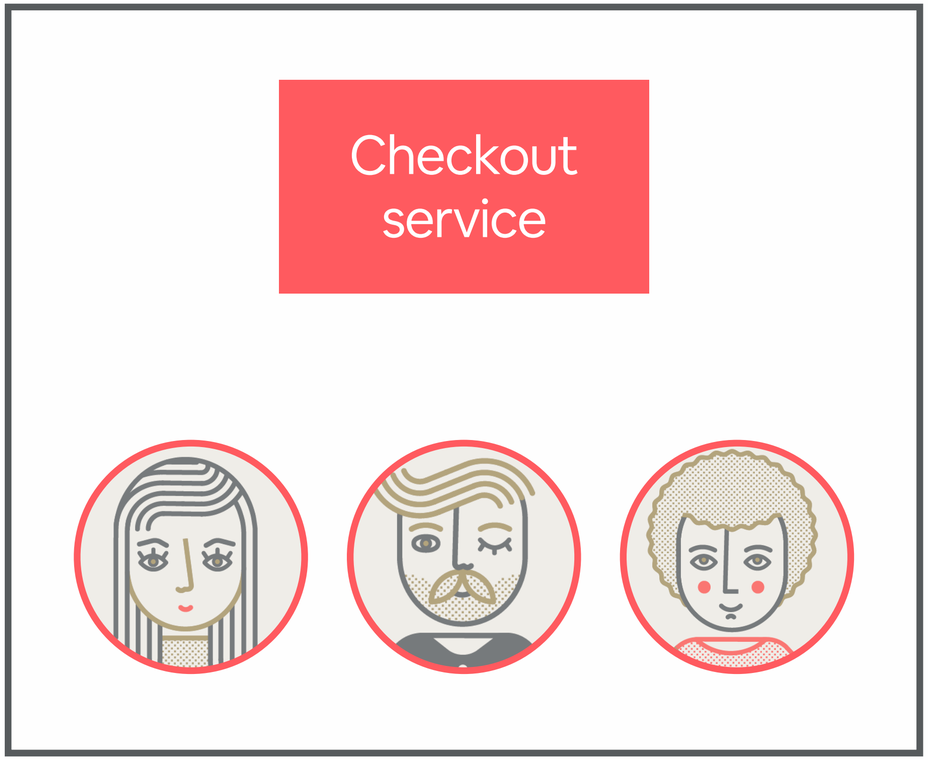

Product Team

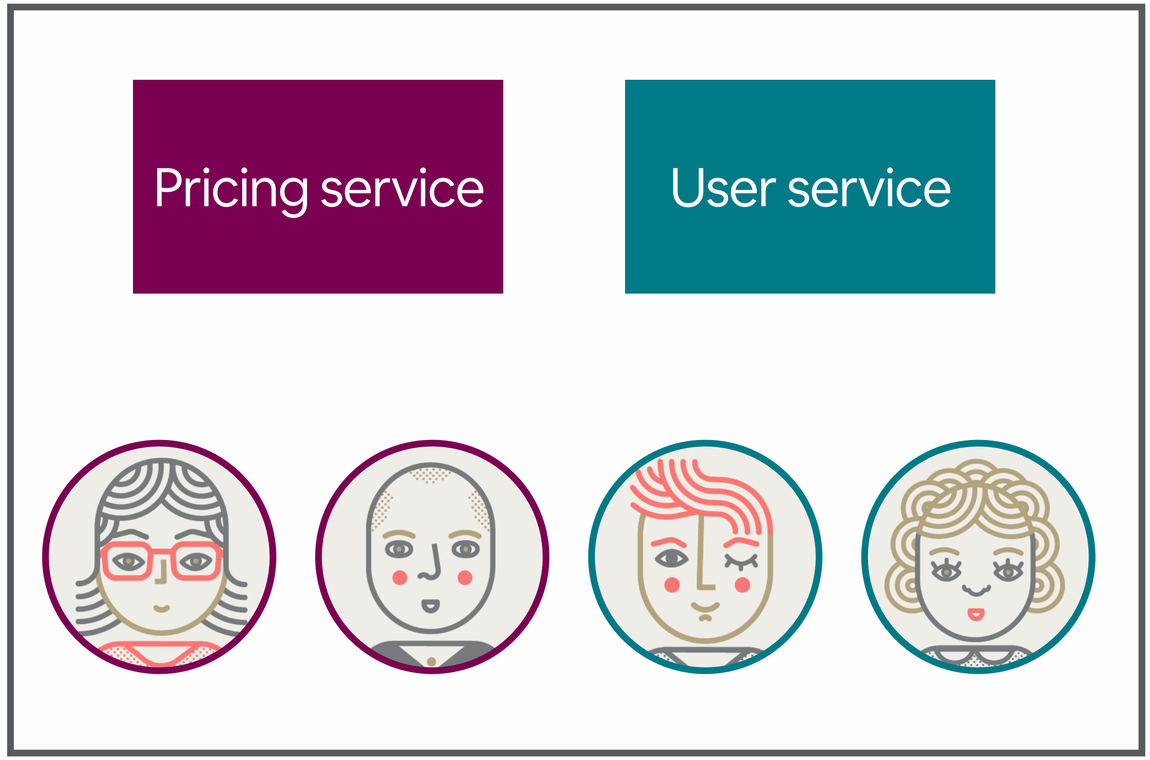

Infrastructure Team


PROGRESS SO FAR?
@jessicamtai

2019:
Whole engineering org

2016:
One small infra team

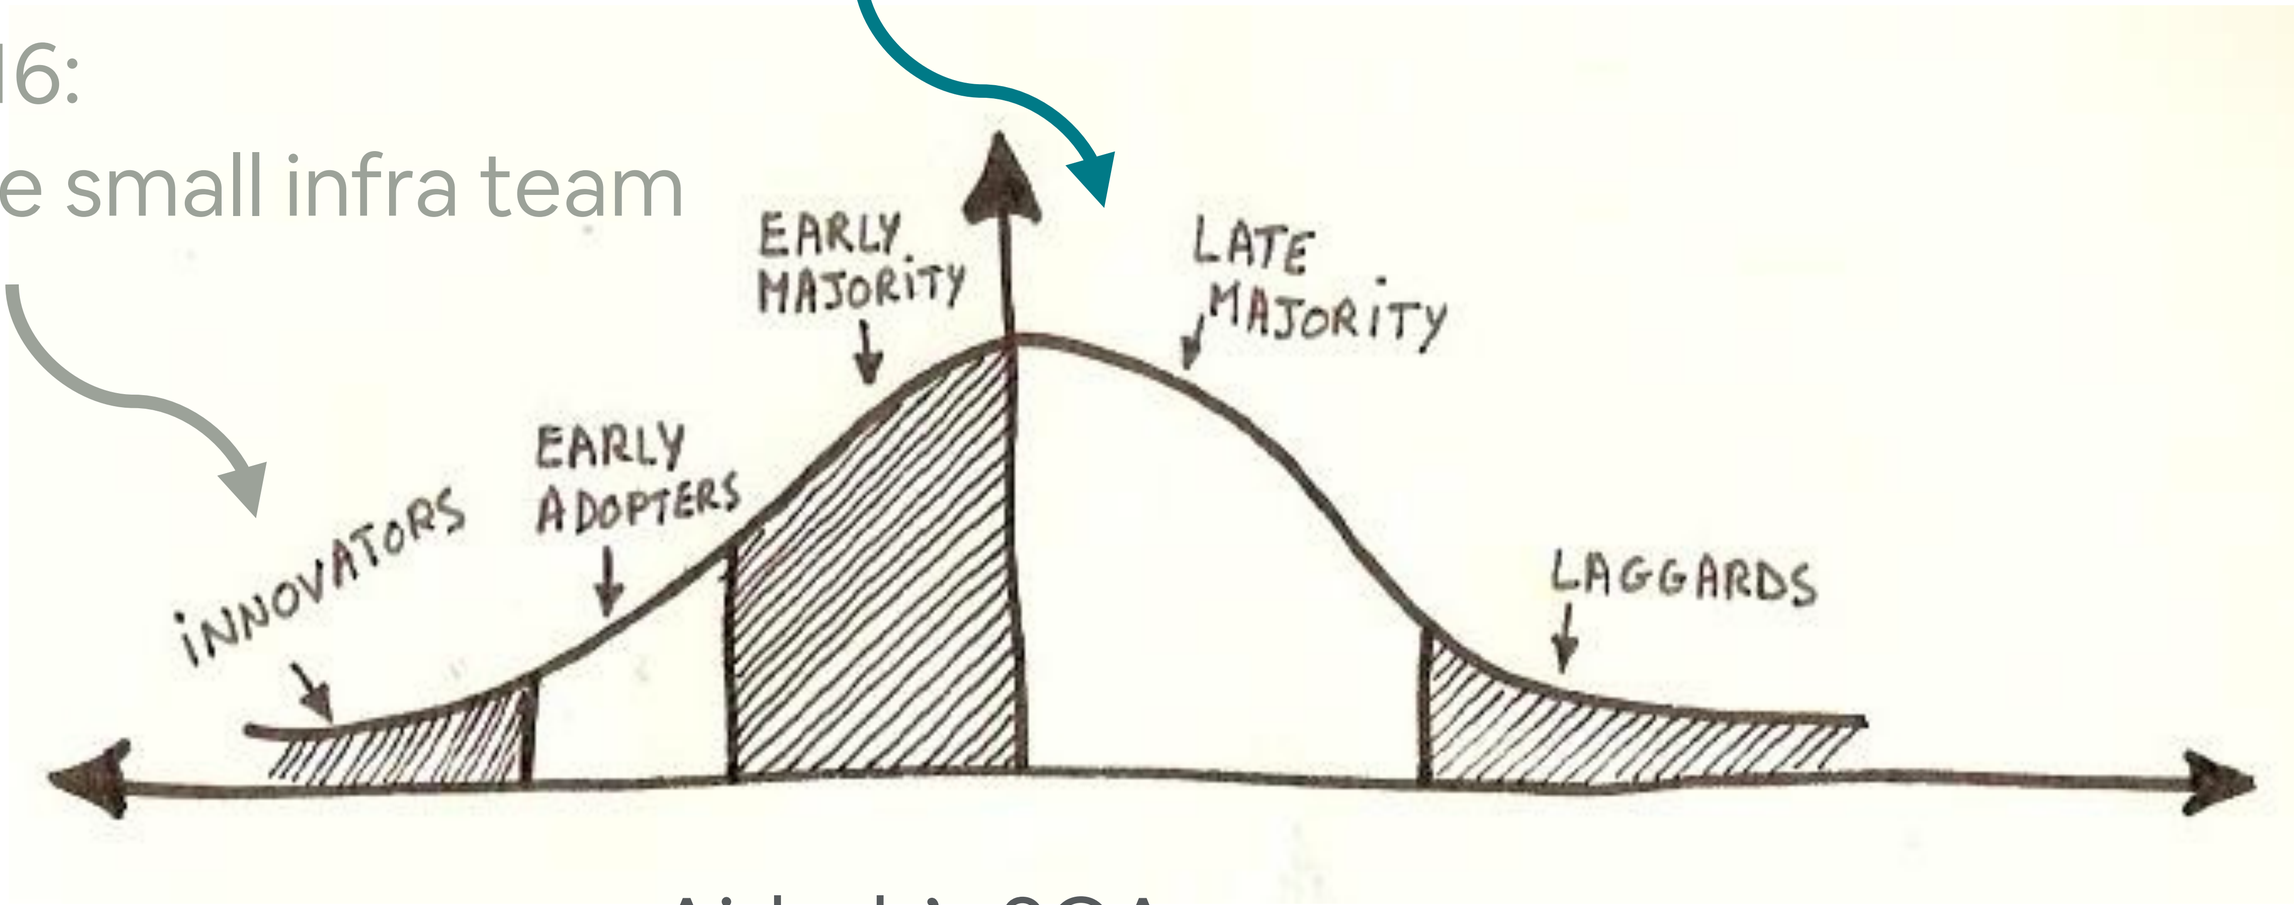

Airbnb's SOA progress

Promising initial results

SUCCESS

- Faster build & deploy times
  - Hours (Monorail) to minutes (service)
  - Fewer reverts
- Quicker bug fixes
- Increased developer productivity happiness

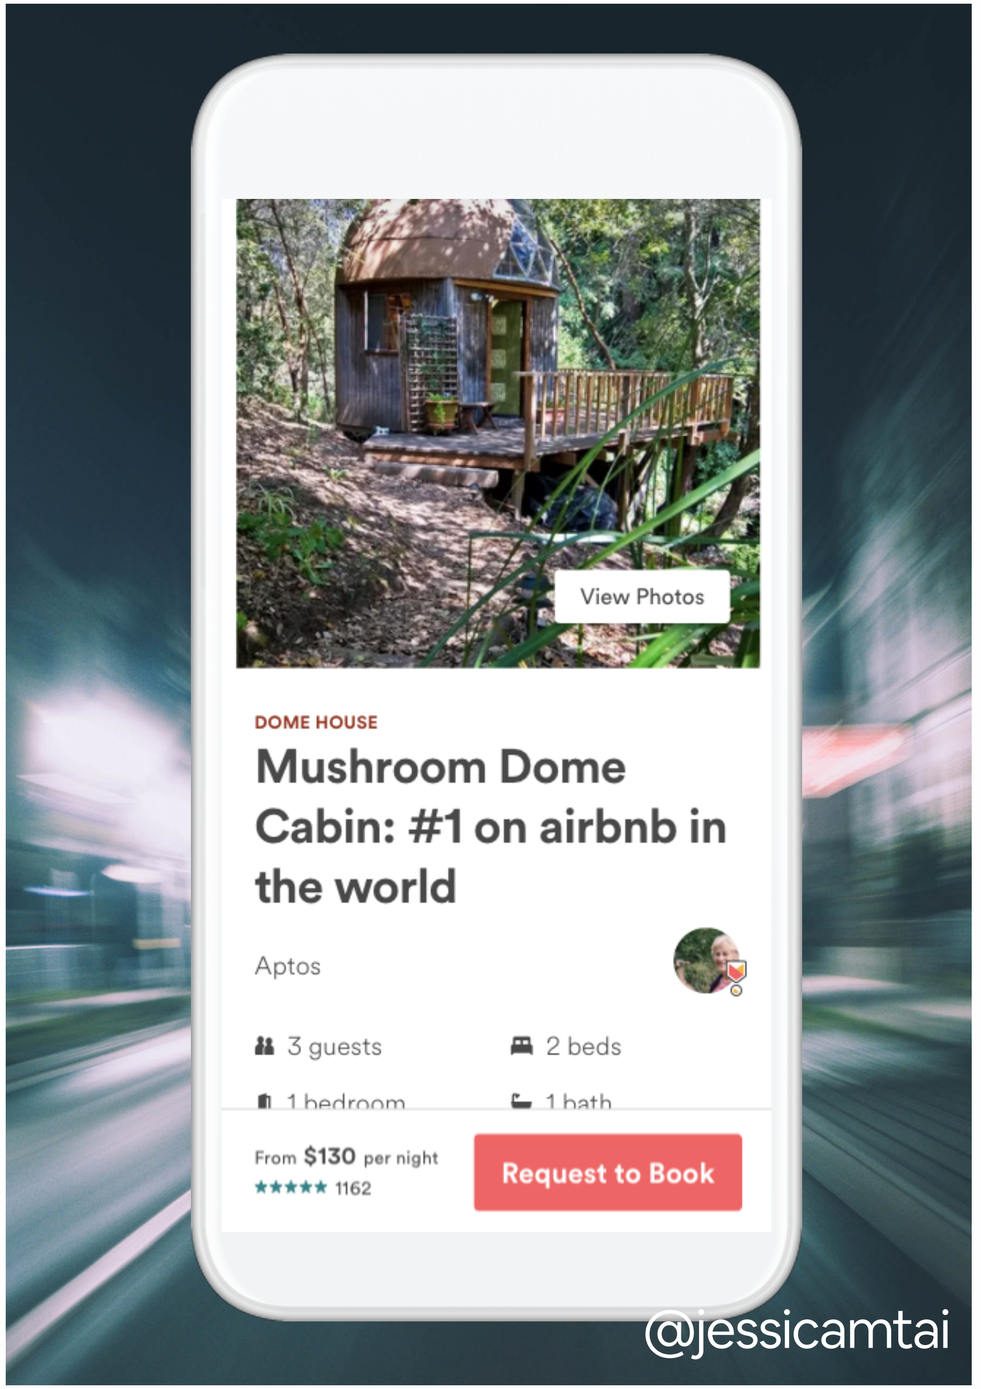

Latency results

SUCCESS

- Lower latency from parallelization
  - Ruby monorail single-threaded
  - Java services multi-threaded
- Search results page **3x faster**
- Homes description page **10x faster!**

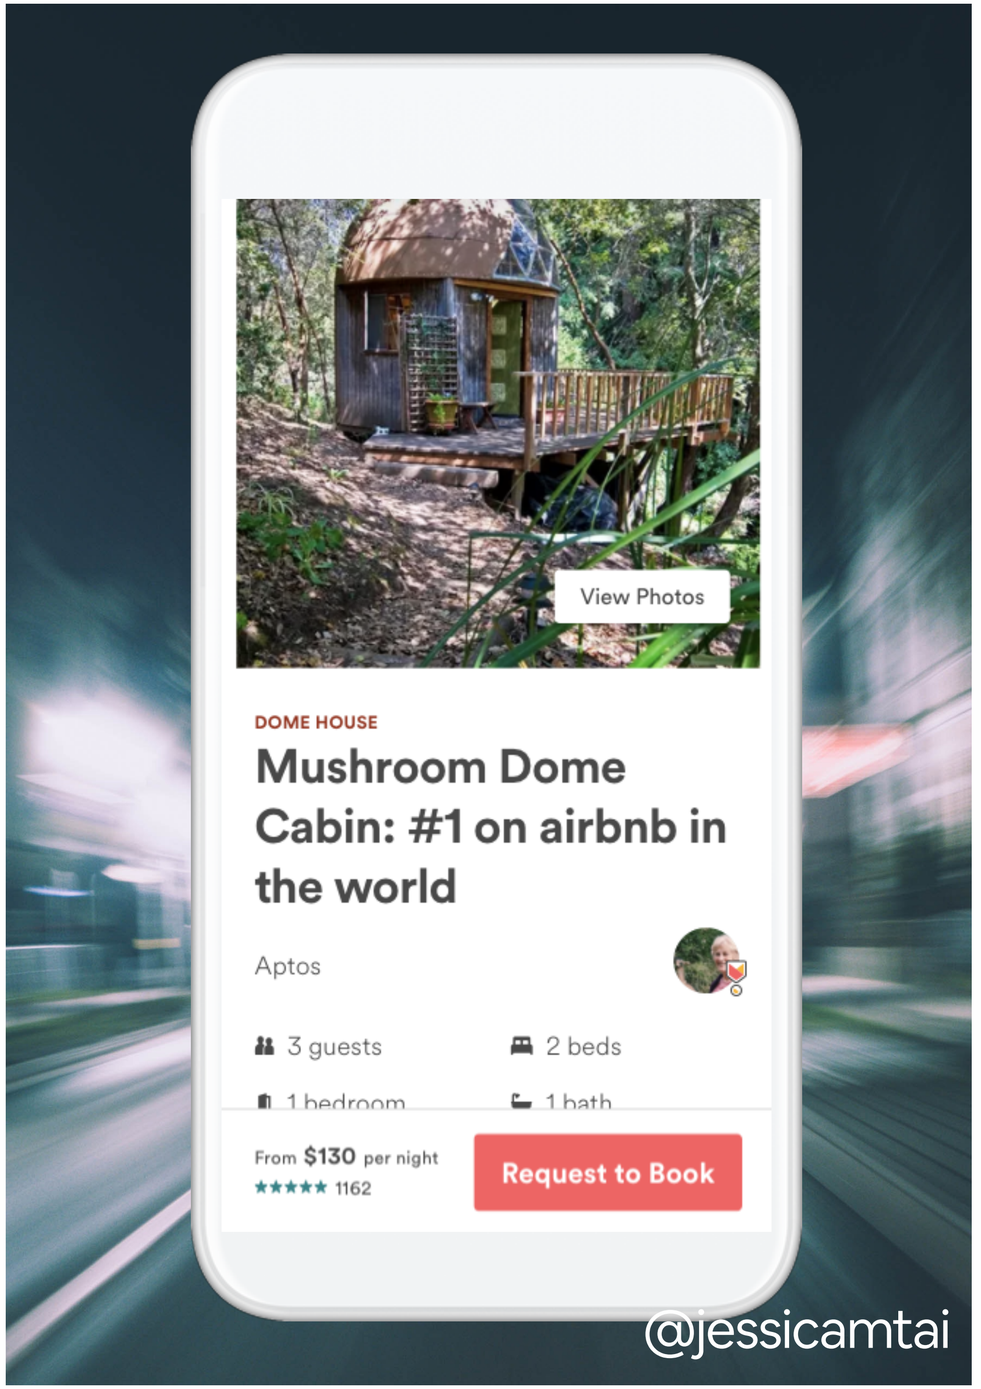

Monorail freeze

@jessicamtai

2016

500+

Engineers

67%

Deploys in Monorail

@jessicamtai

2016 500+

2019 1200+

Engineers

67%

7%

Deploys in Monorail

@jessicamtai

40%

Production traffic via API Gateway

400+

IDL services in production

Checkout page required message in SOA

Checkout **presentation**

1. Review trip details > 2. Who's coming? > 3. Confirm and Pay

Home demand **derived data**

This is a rare find. Kitty's place is usually booked.

Reservation **data**

Who's coming?

Guests

1 guest

Business travel **derived data**

What's the main purpose of this trip?

- Personal travel
- Business travel

Messaging **data**

Say hello to your host

Let Kitty know a little about yourself and why you're coming.

Hello Kitty! Can't wait to spend 1 night in your home.

Continue

Home **data**

Mushroom Dome Cabin: #1 on airbnb in the world

Entire home/apt in Aptos

★★★★★ 1165 reviews

1 guest

Jun 21, 2018 → Jun 22, 2018

Pricing **derived data**

$130.00 x 1 night	$130.00
Cleaning fee	$35.00
Service fee	$21.55
Occupancy Taxes	$18.15
Gift Credit	-$204.70
Total (USD)	**$0.00**

The quoted fees include any applicable exchange rate.

Cancellation **derived data**

Moderate — free cancellation

Cancel within 48 hours of booking to get a full refund. More details

Review **data**

Tay

Stayed on April 2018

Featured review

"Staying at the Mushroom Dome was

@jessicamtai


CAUTION
SOA has its challenges


CAUTION
Distributed services


CAUTION
Multiple, isolated databases

@jessicamtai

SOA migration

TAKEAWAYS

- Prepare for a long commitment
- Decompose incrementally
- Scale with auto-generating frameworks, tools
- Shift development culture


Look both ways
during your Great Migration
@jessicamtai

QCON LONDON 2019

@JESSICAMTAI