Iterative dashboards & monitors

CARMEL HINKS | SOFTWARE ENGINEER | ATLASSIAN

Us, with our hands down

You build it, you run it

You build it, you run it

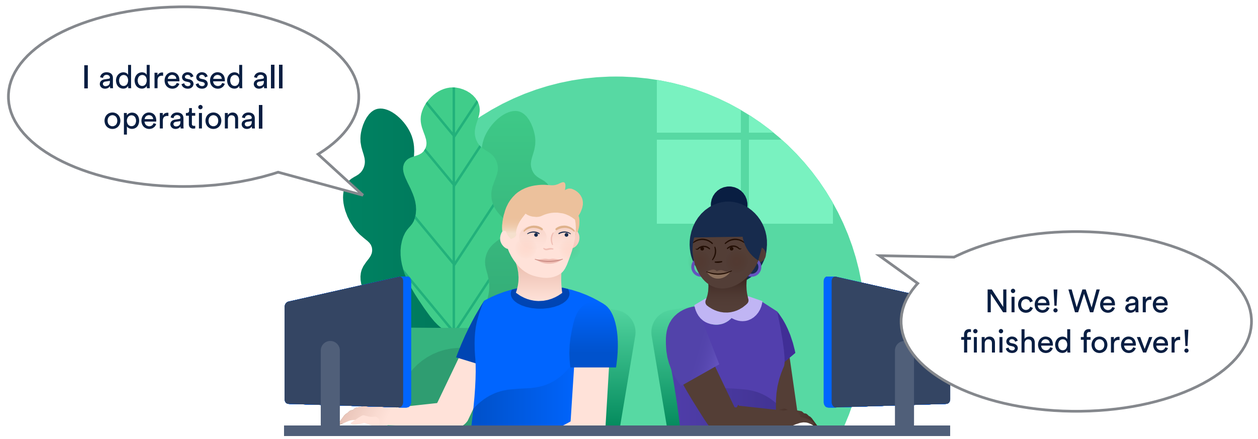

Past

Present

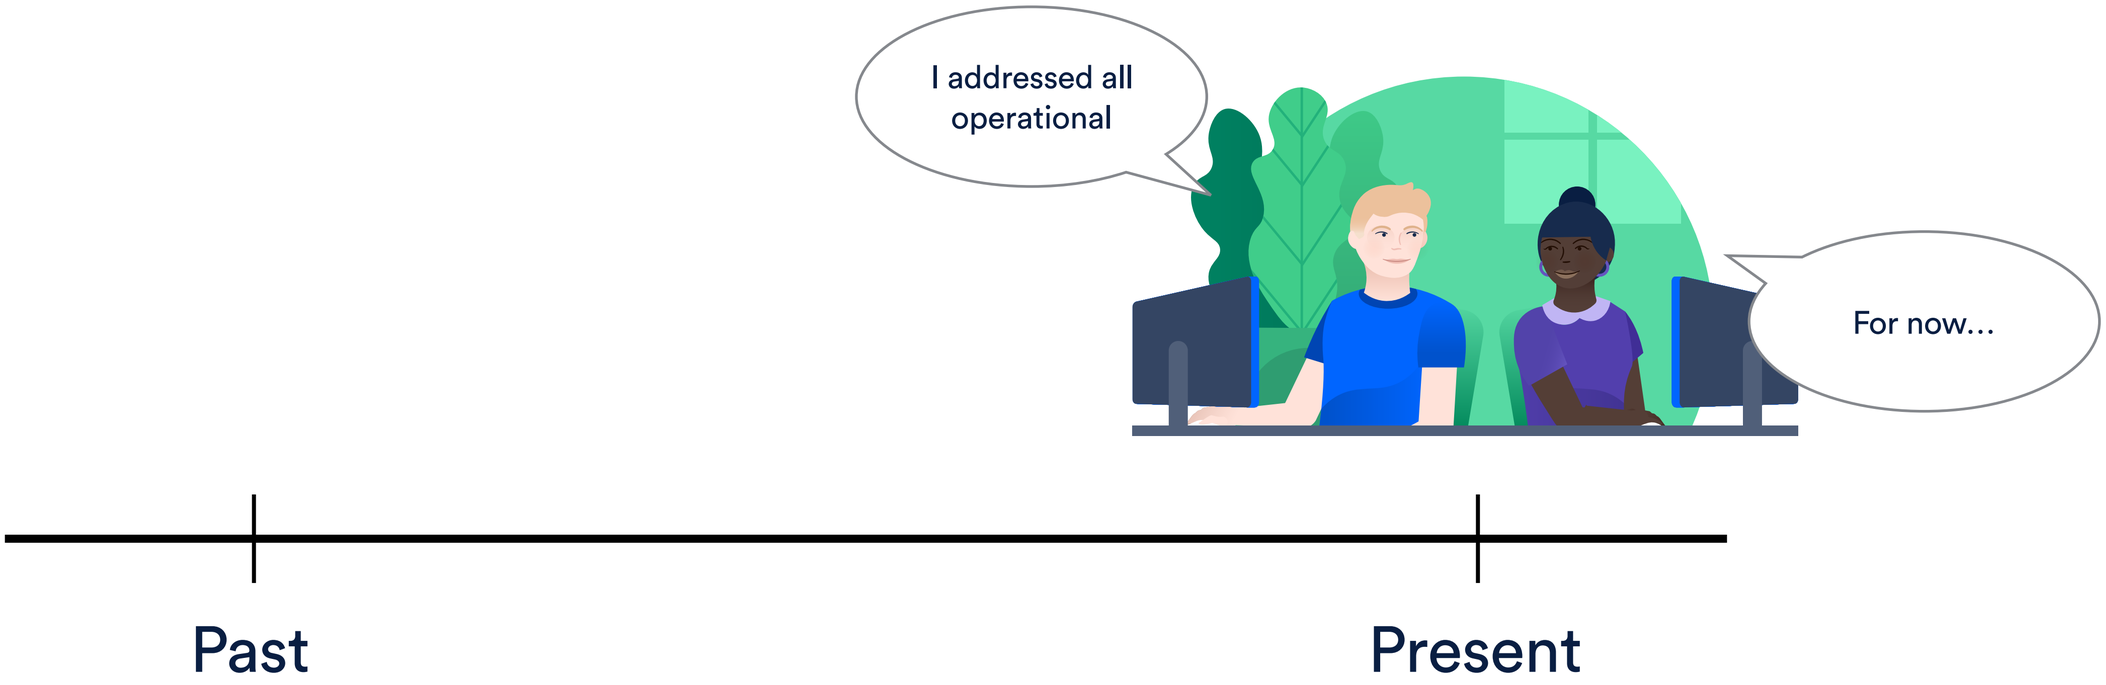
I addressed all operational
For now...
Past
Present

Agenda

Iterative… what?

Setting some context

Deciding what to measure

Verifying your metrics

Keeping up with change

Summary

Agenda

Iterative... what?

Setting some context

Deciding what to measure

Verifying your metrics

Keeping up with change

Summary

Metrics

Dashboards

Monitors

Metric

A measure of a software characteristic

Analytic..?

Metric

A measure of a software characteristic

Analytic

What are our **users** doing?

Metric

What are our **systems** doing?

Dashboard

A visualisation of your metrics

Monitor

An alert against one or more metrics

Operational health

What went wrong?

Operational health

When did it go wrong?

Operational health

Why did it go wrong?

Operational health

Operational health

Agenda

Iterative... what?

Setting some context

Deciding what to measure

Verifying your metrics

Keeping up with change

Summary

Agenda

Iterative... what?

Setting some context

Deciding what to measure

Verifying your metrics

Keeping up with change

Summary


ATLASSIAN

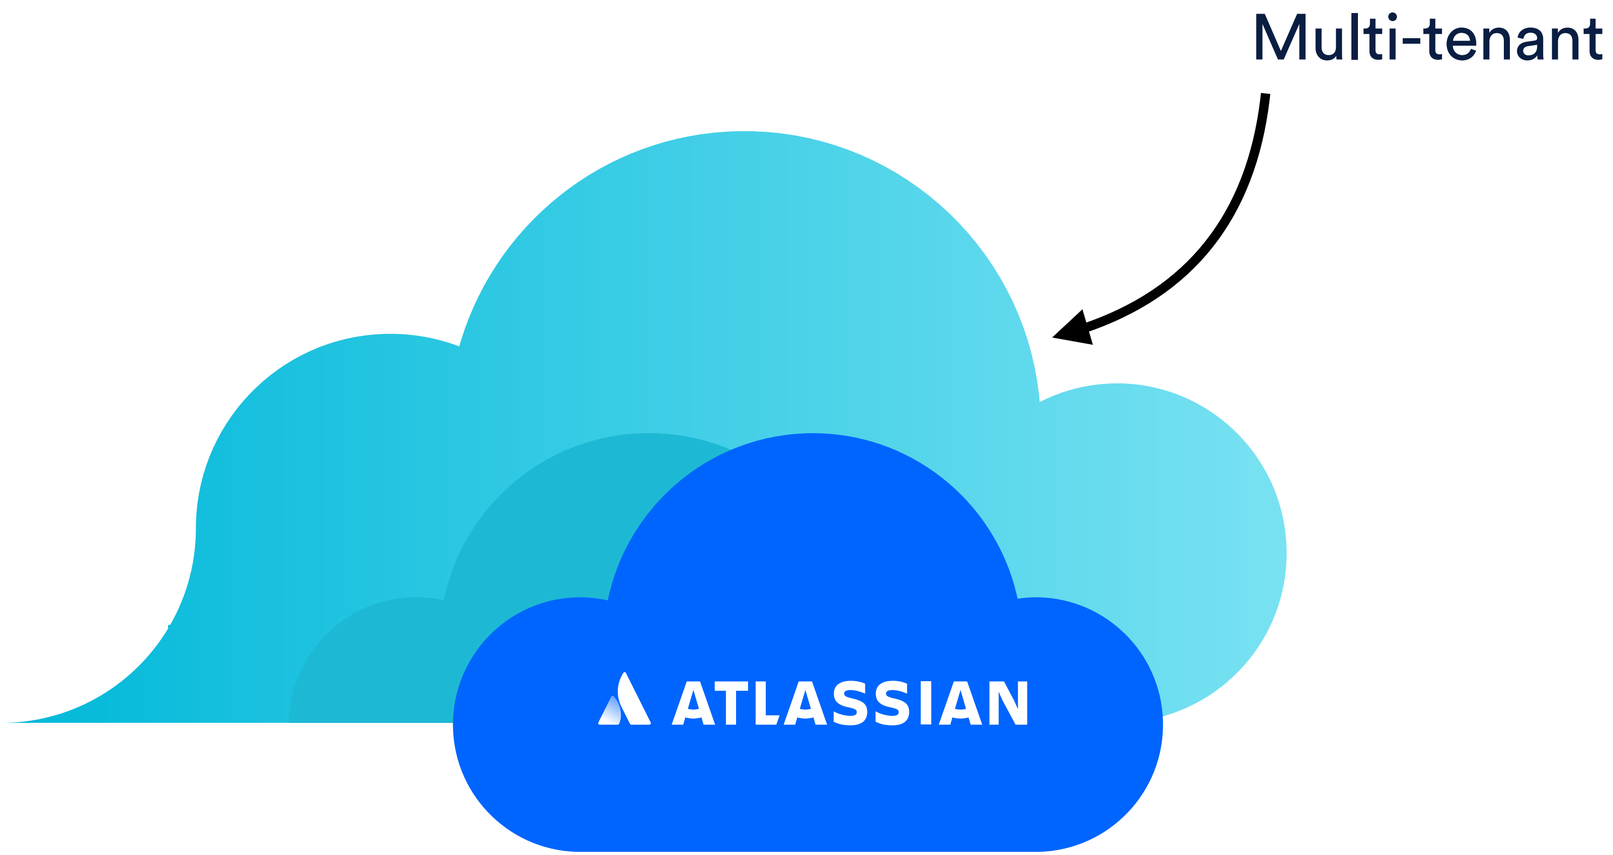
Multi-tenant
ATLASSIAN

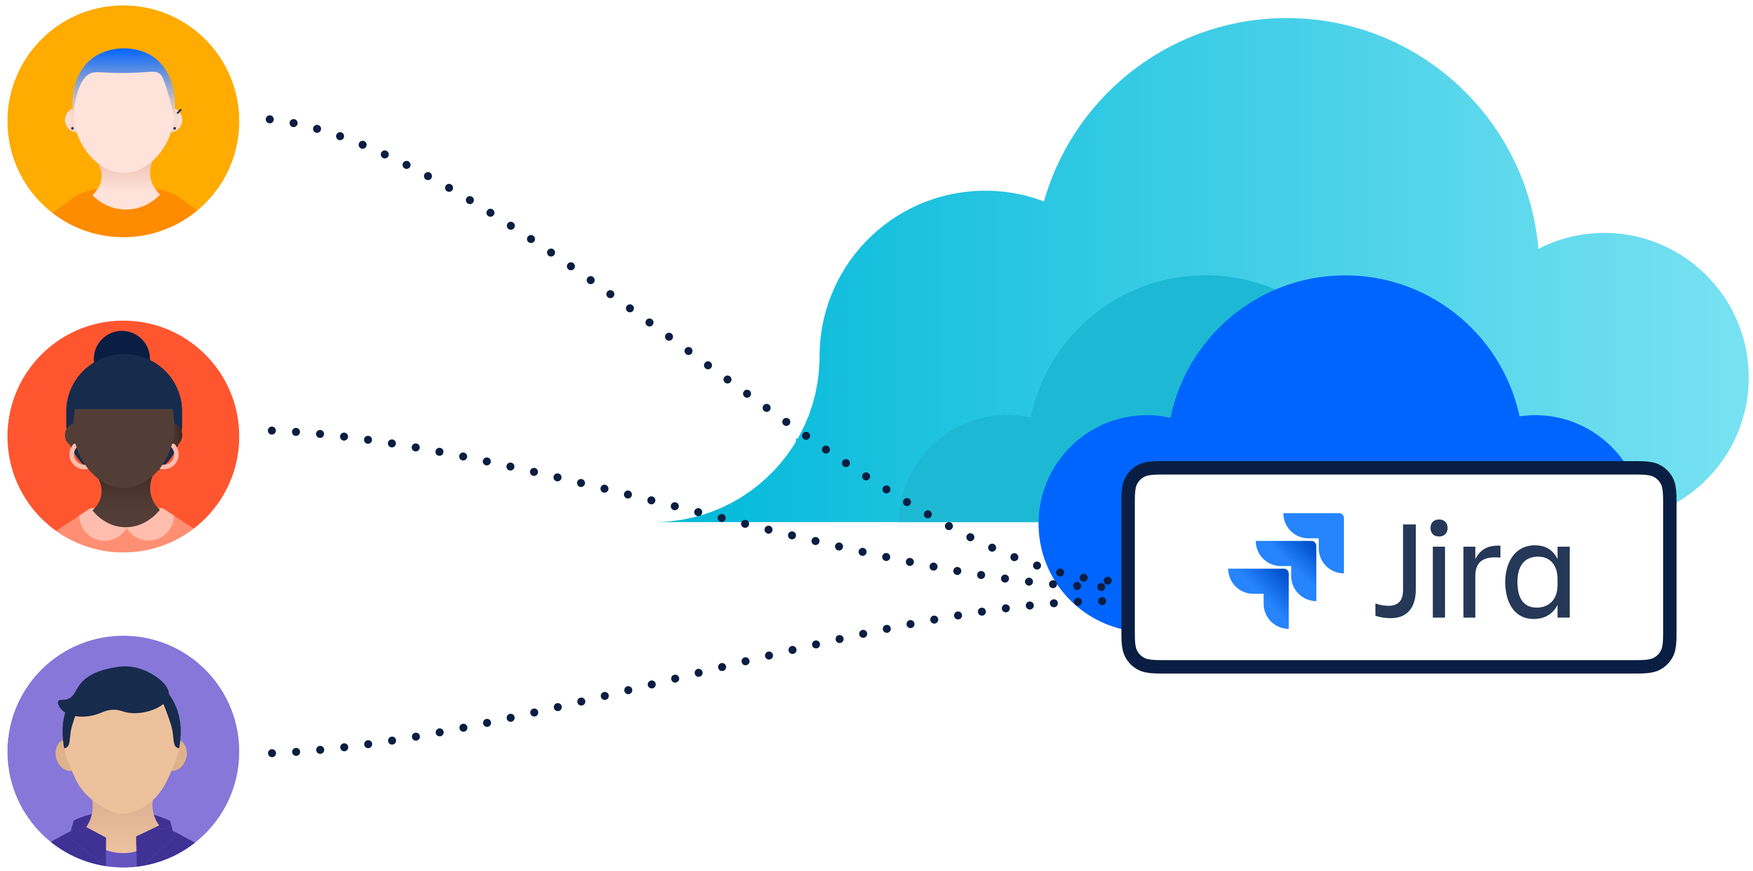
Jira

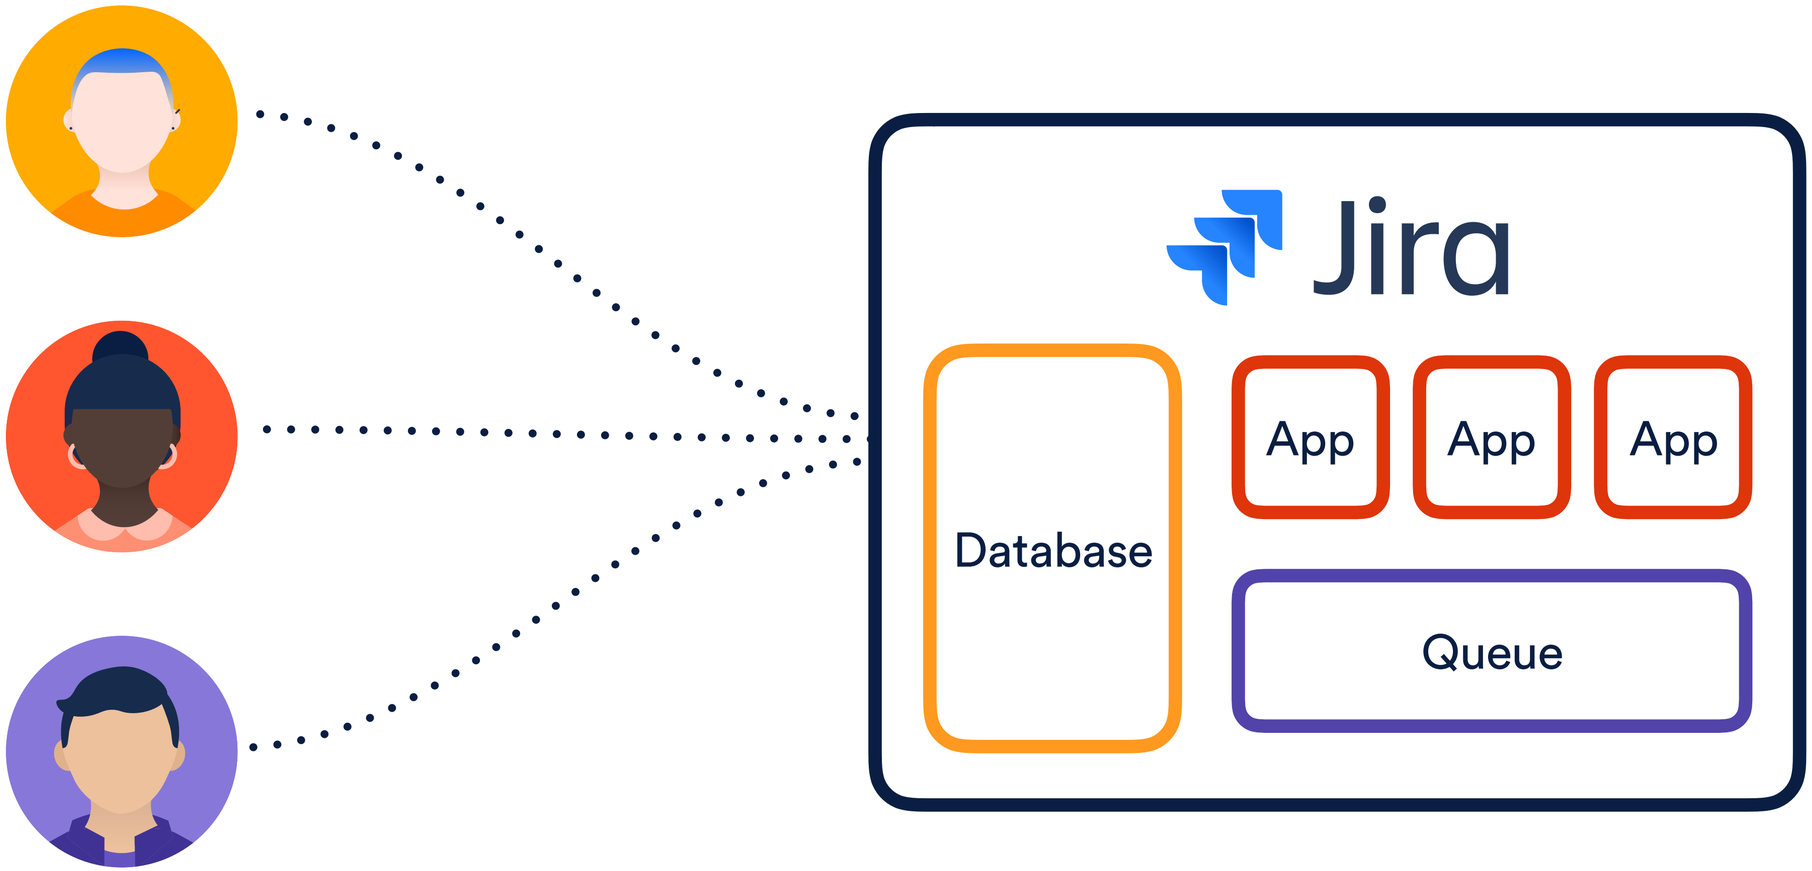
Jira
Database
App
App
App
Queue

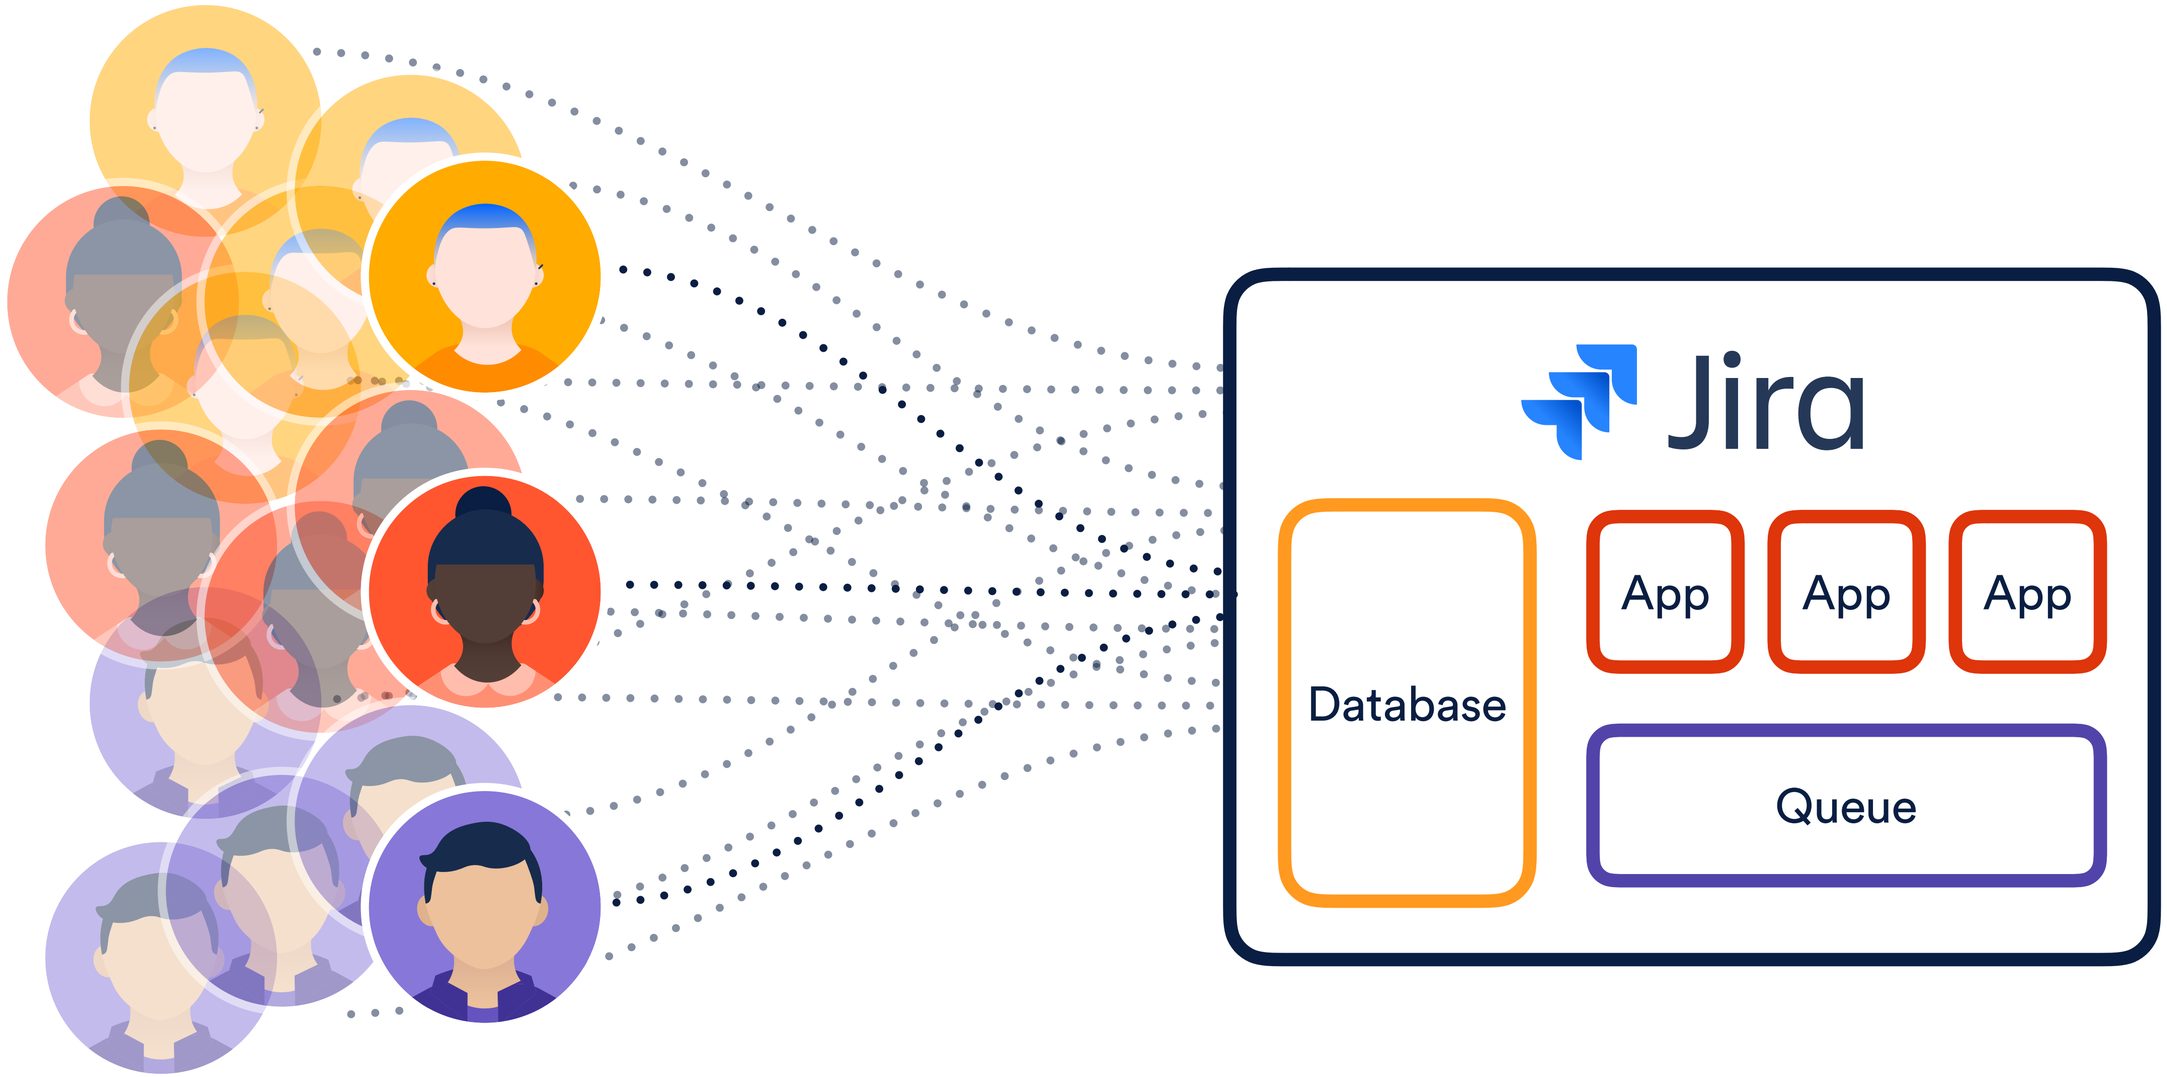
Jira
Database
App
App
App
Queue

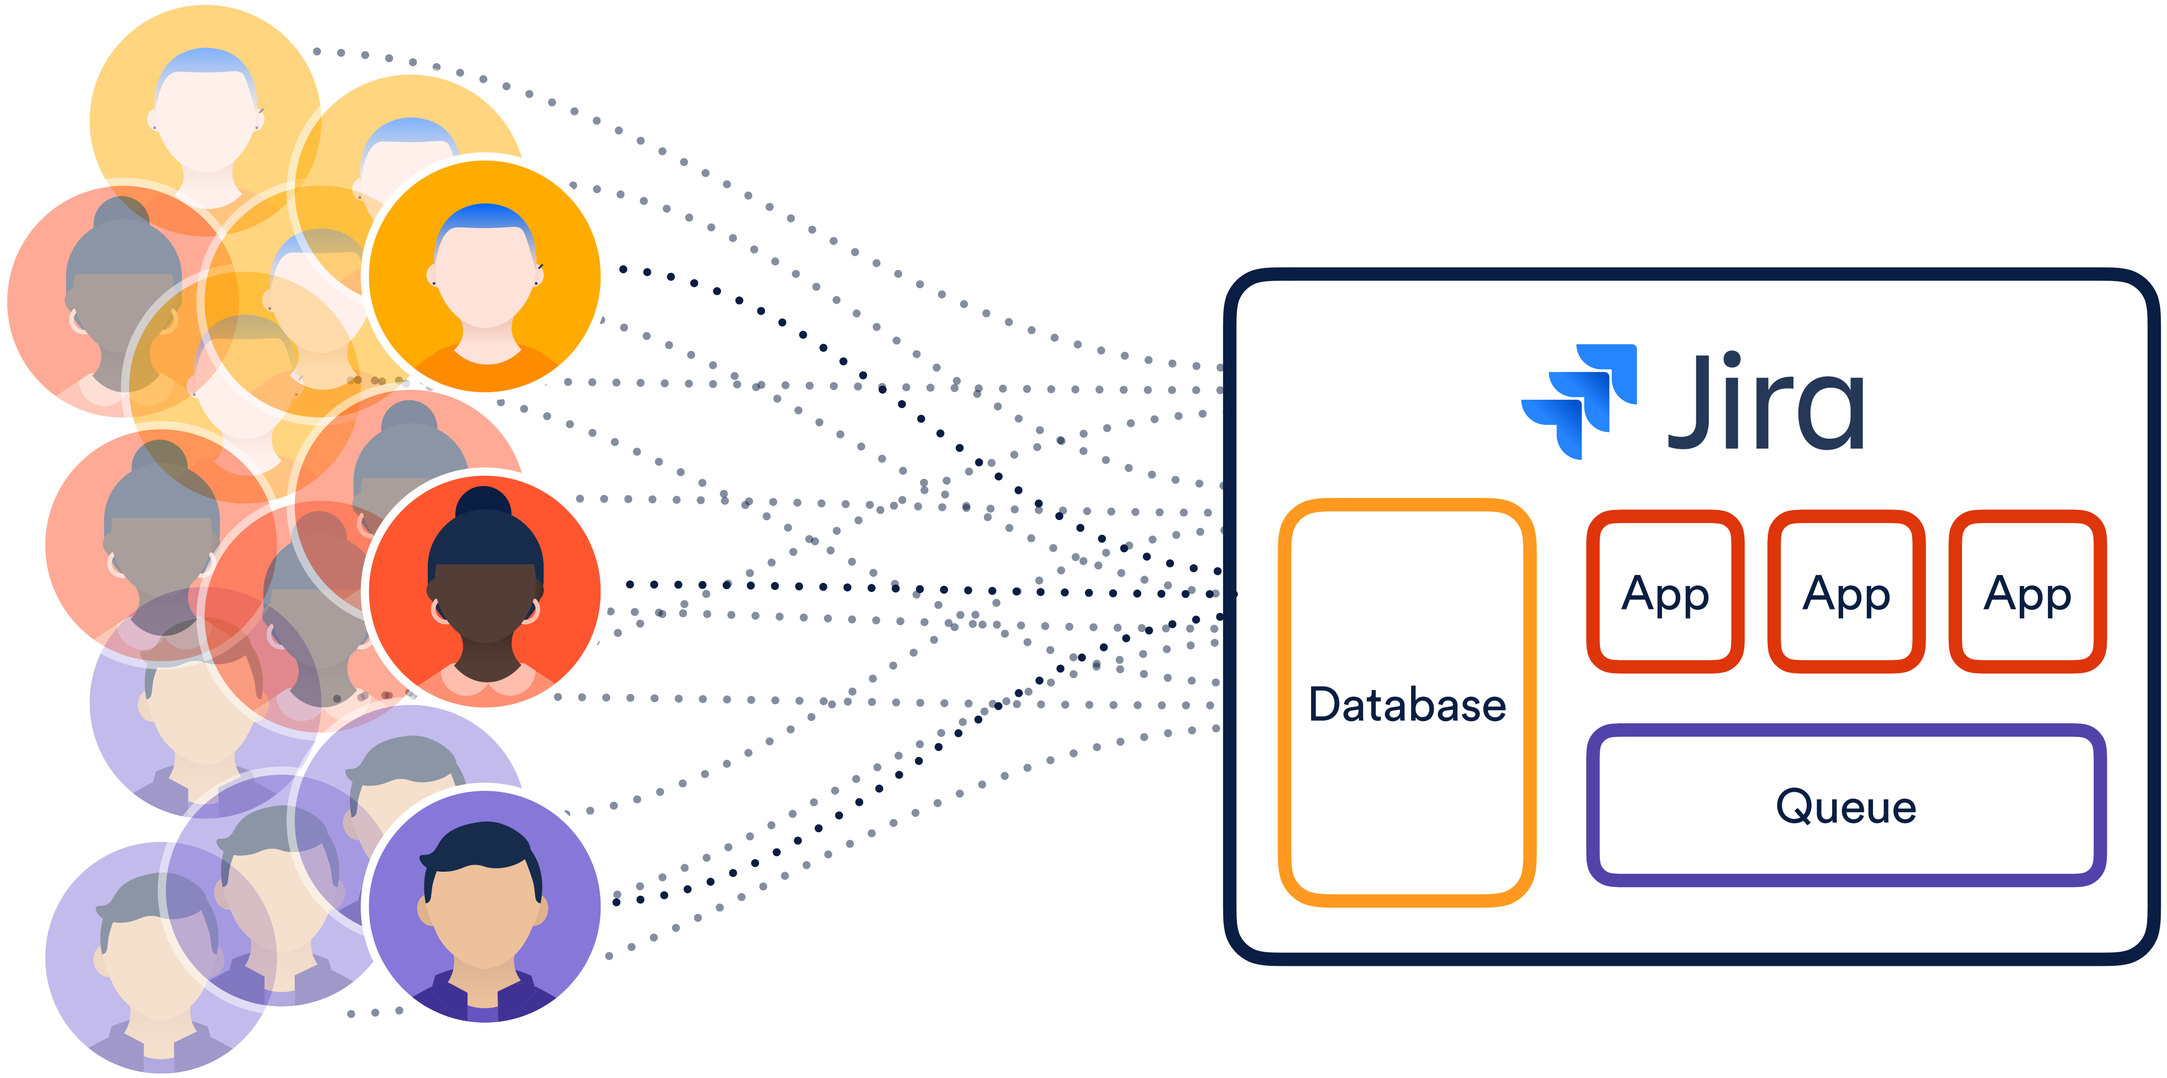

What about **cross-region latency?**

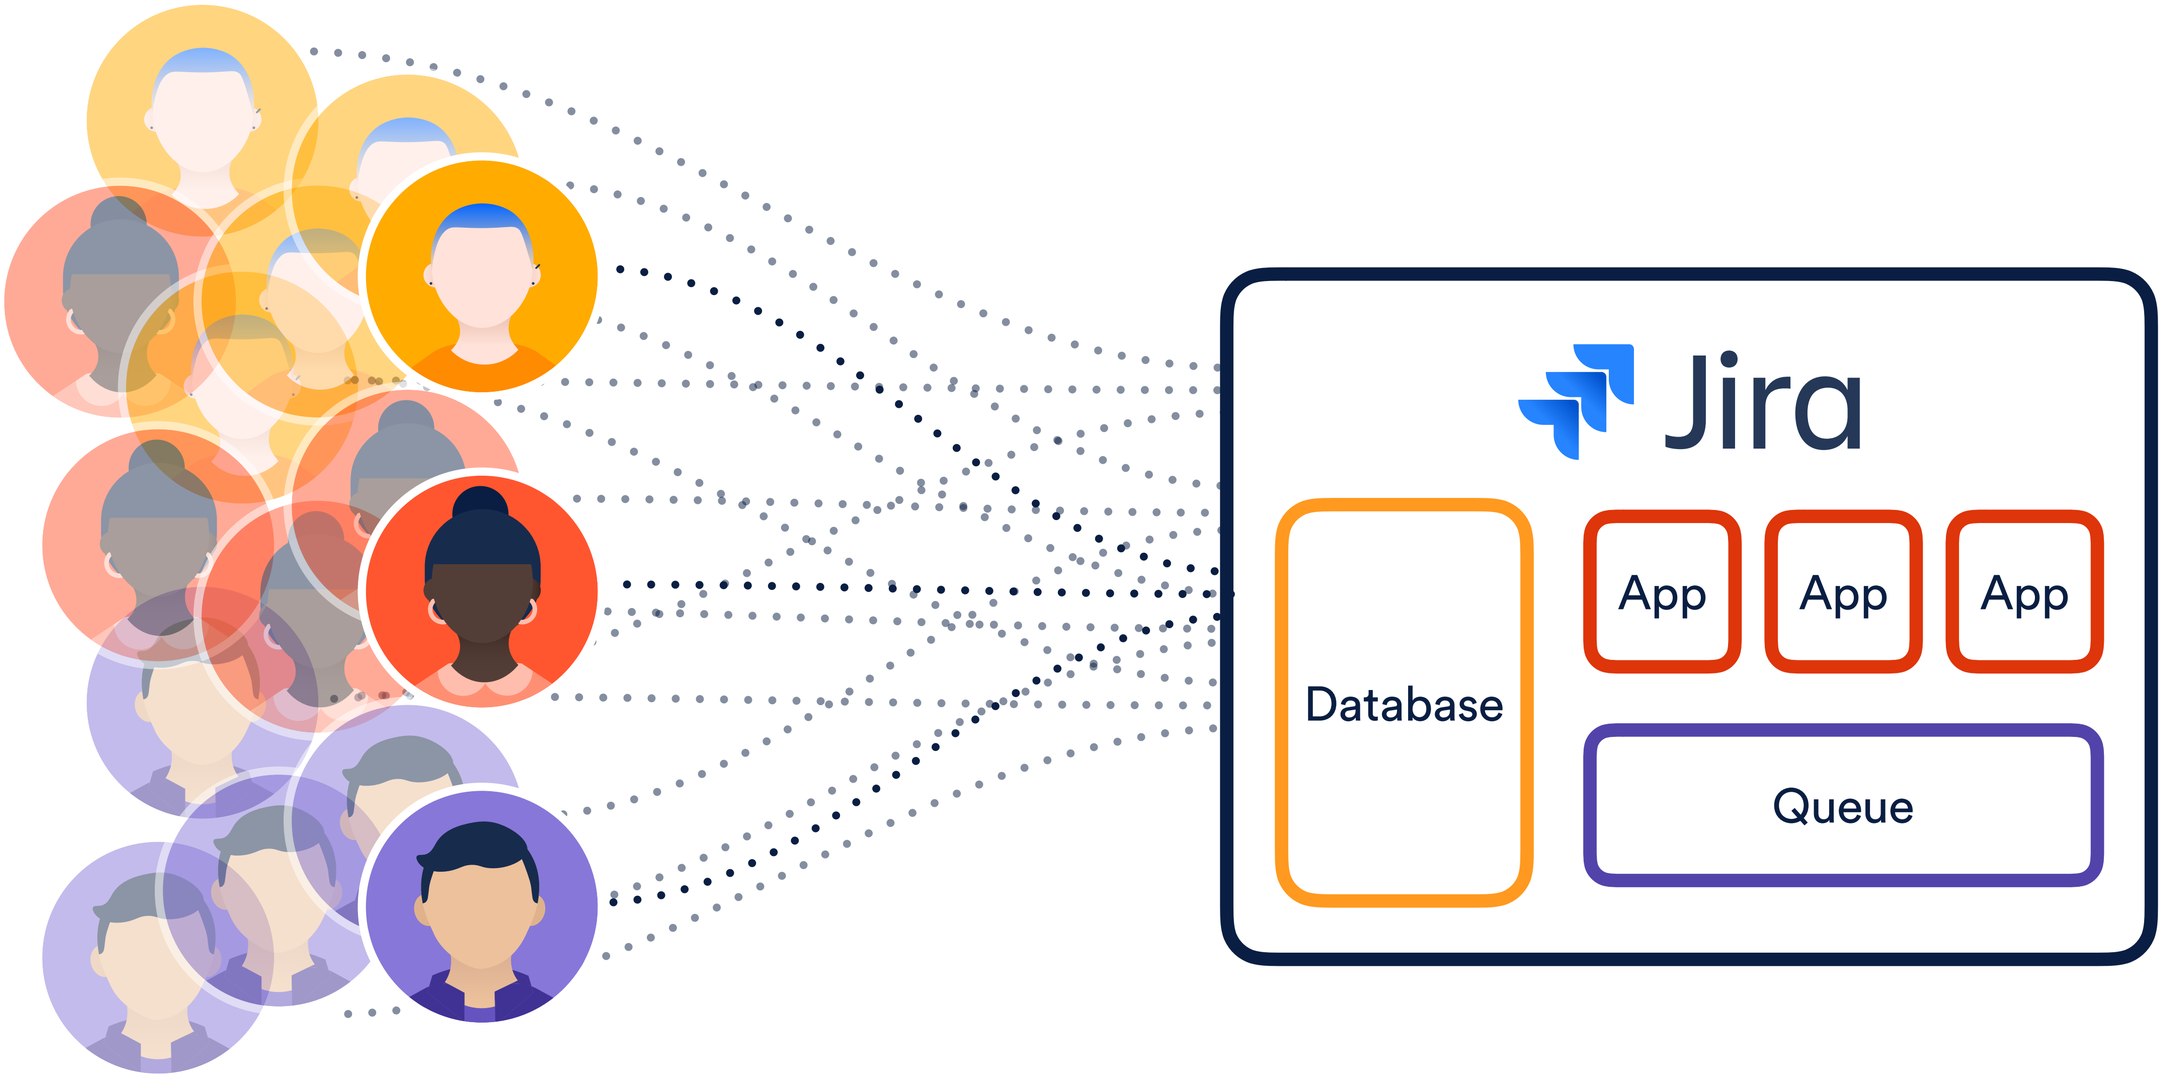

What about **cross-region latency?**

What about **scale?**

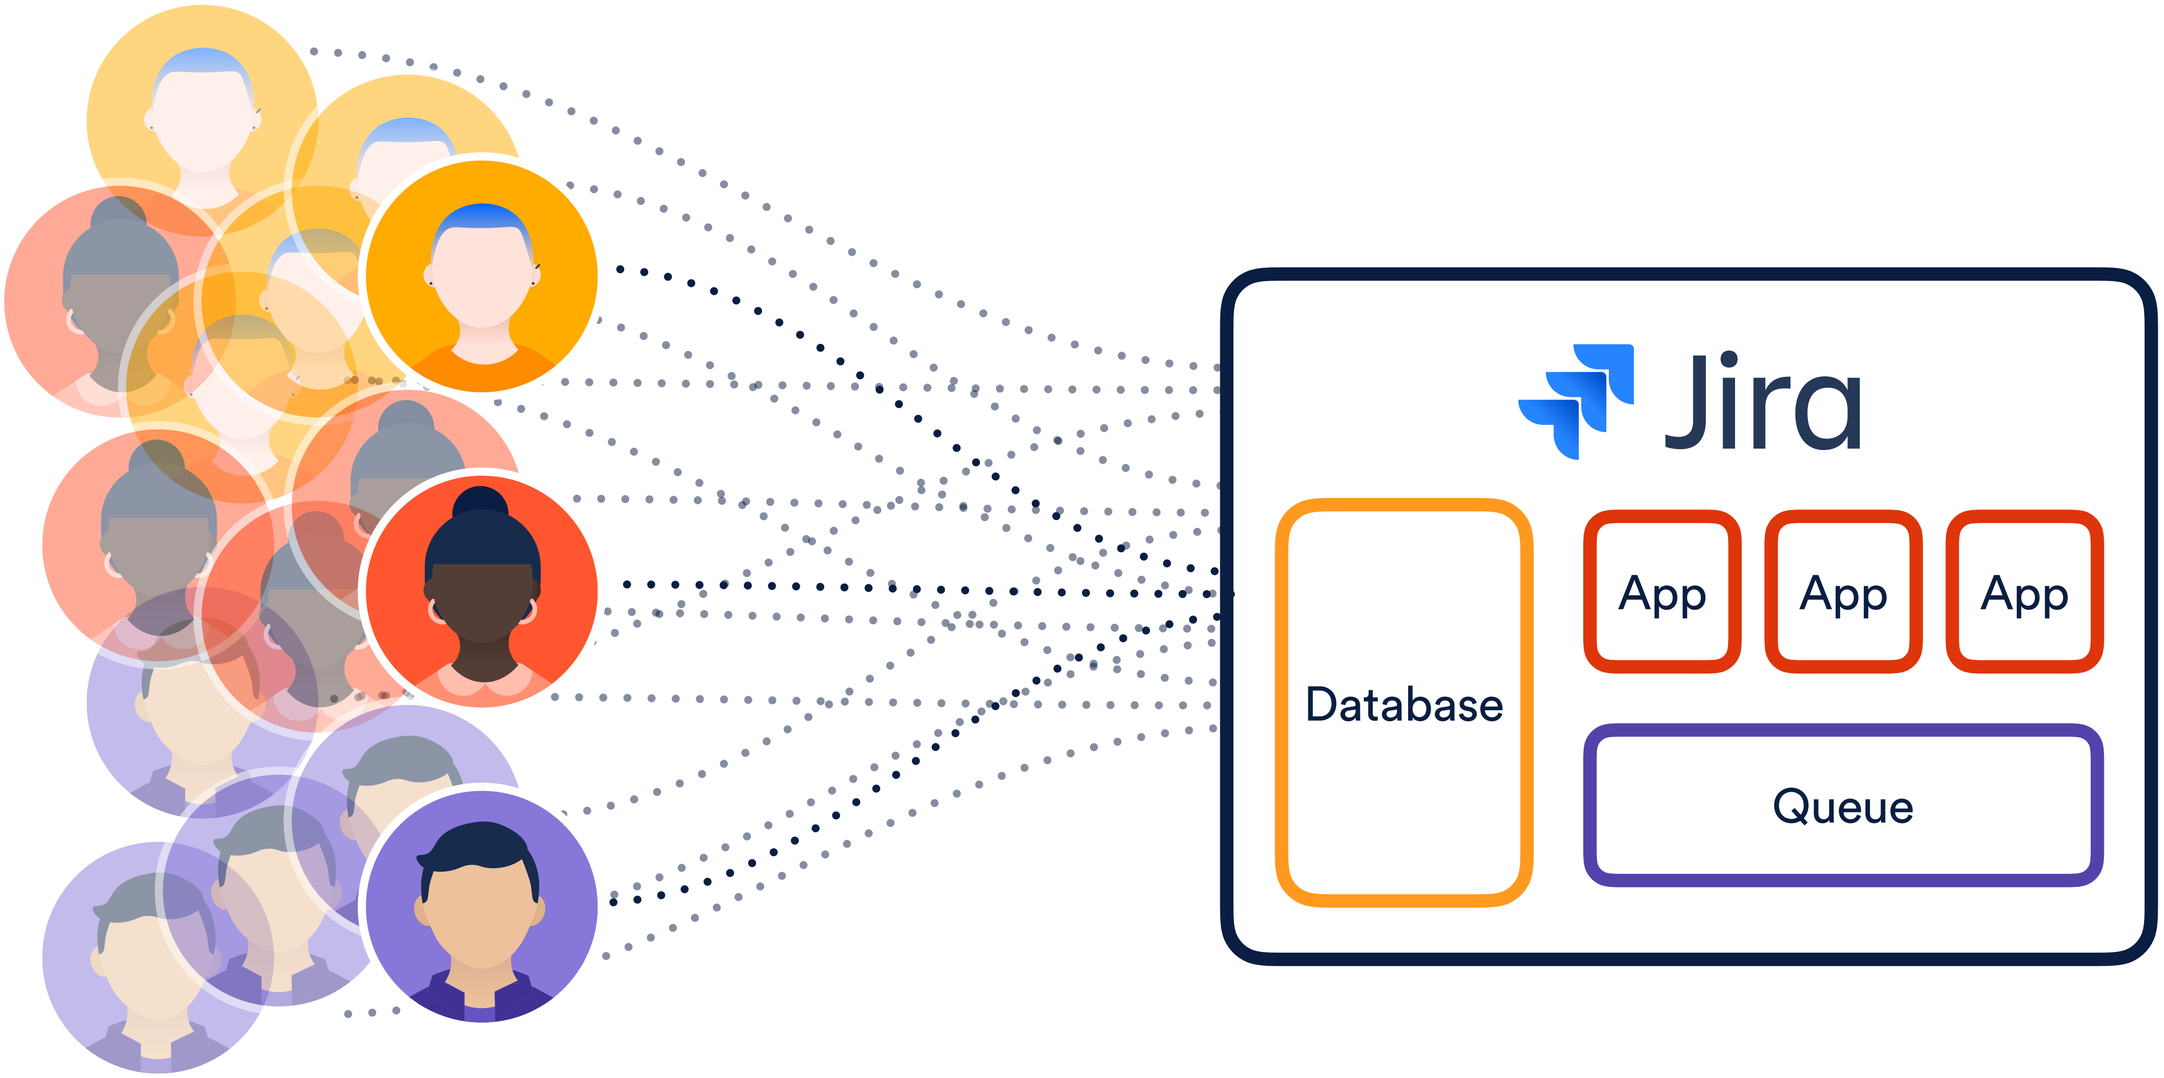

What about **cross-region latency?**

What about **scale?** What about **progressive rollouts?**

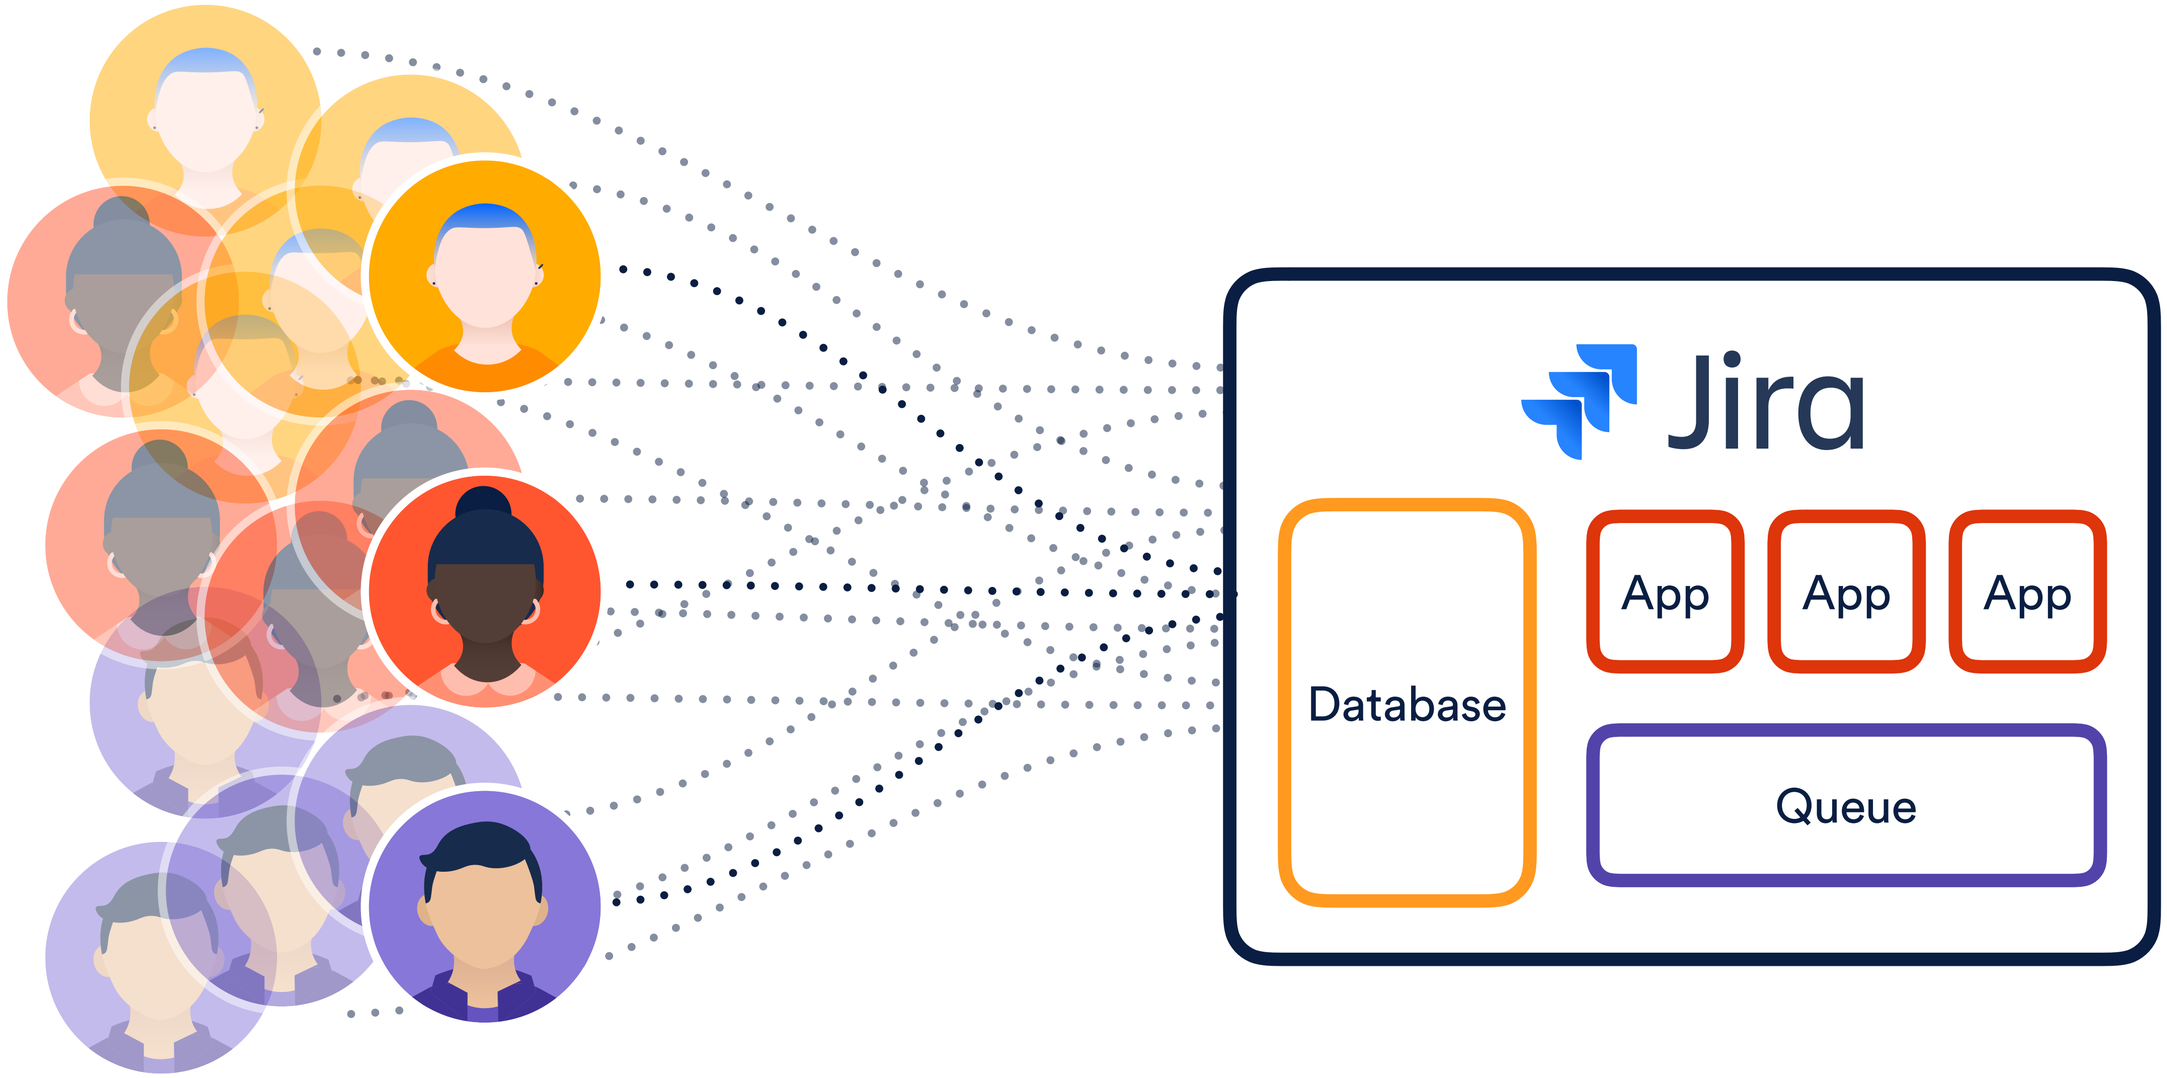

What about **cross-region latency?**

What about **outage blast radius?**

What about **scale?** What about **progressive rollouts?**

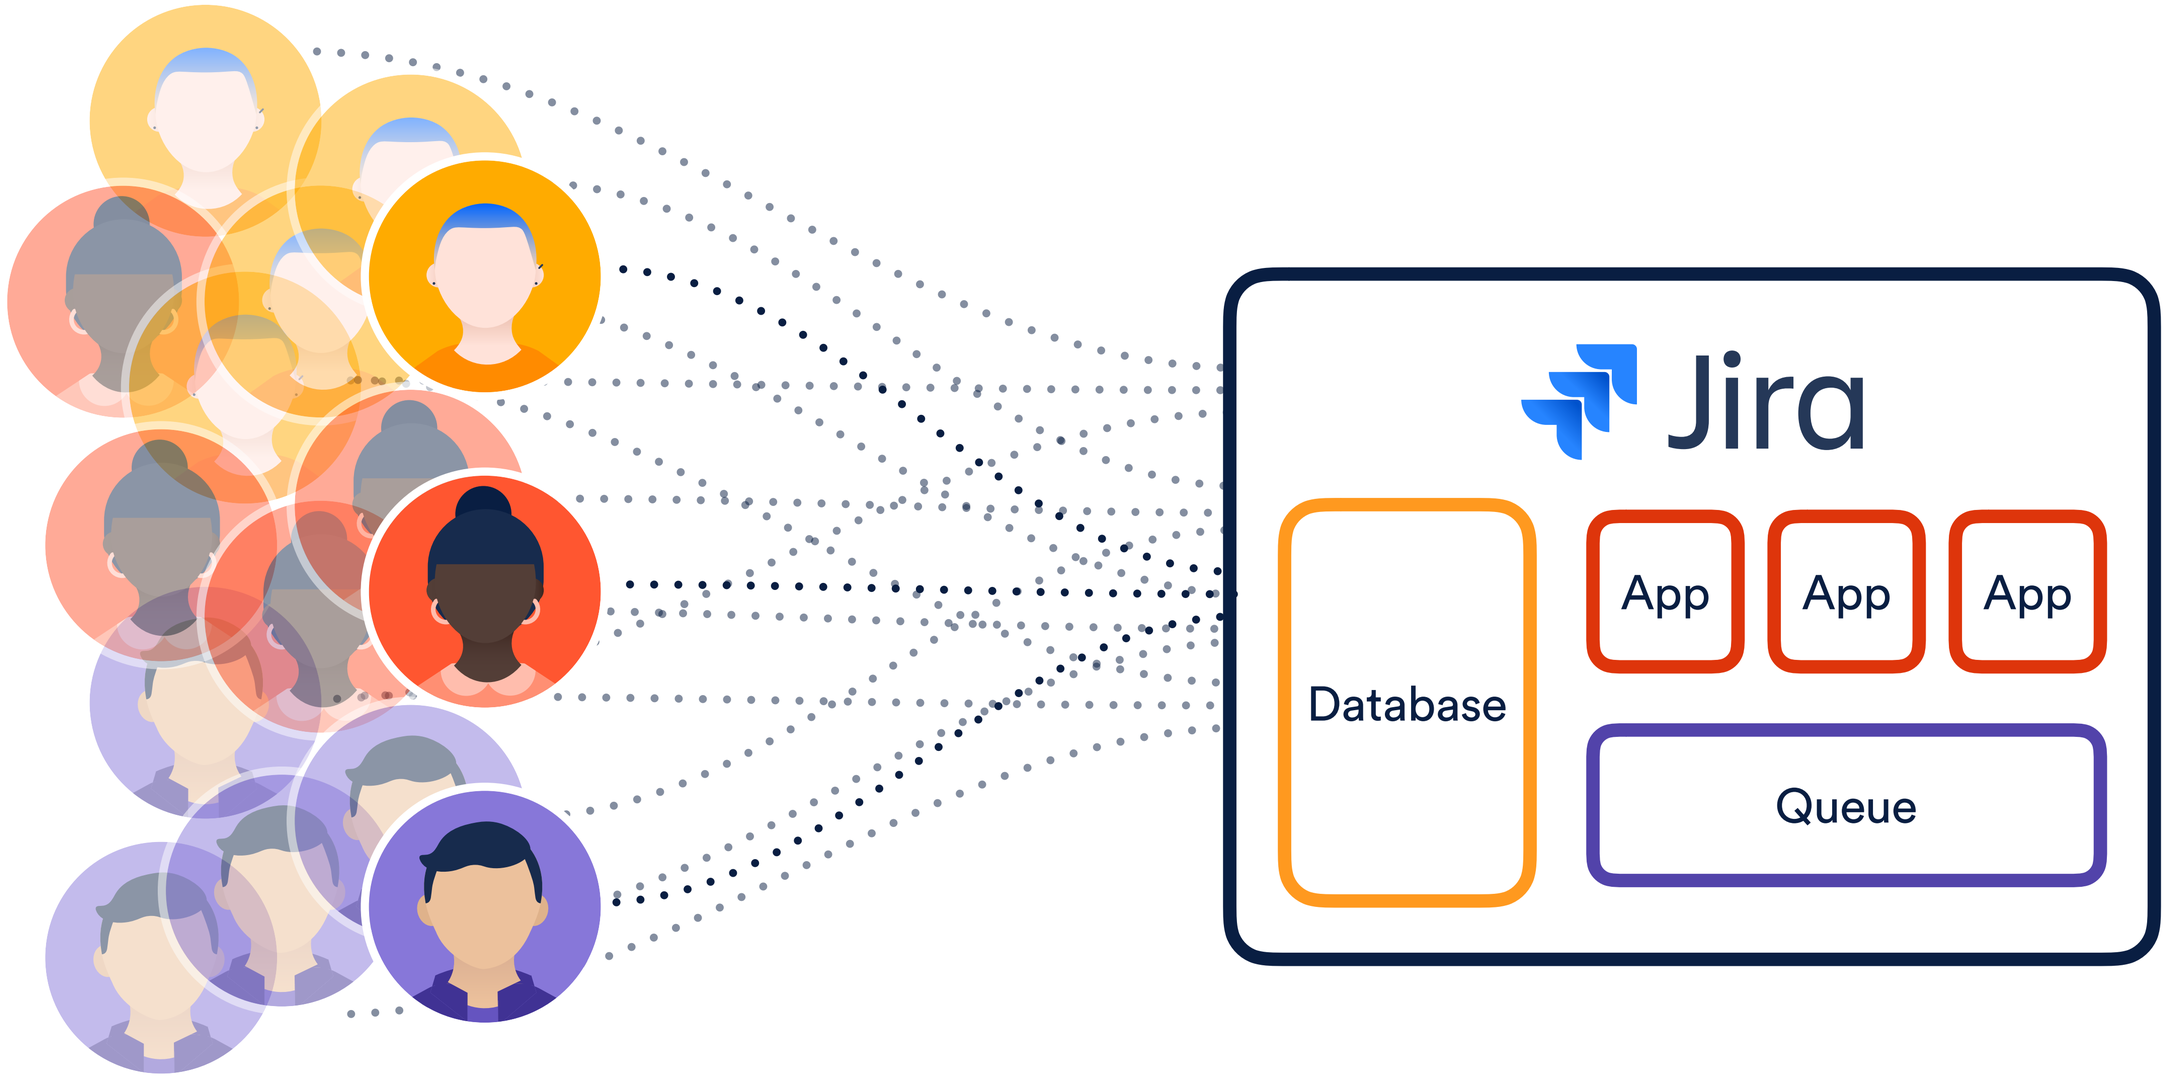

What about **cross-region latency?**
What about **outage blast radius?** What about **data sovereignty?**
What about **scale?** What about **progressive rollouts?**

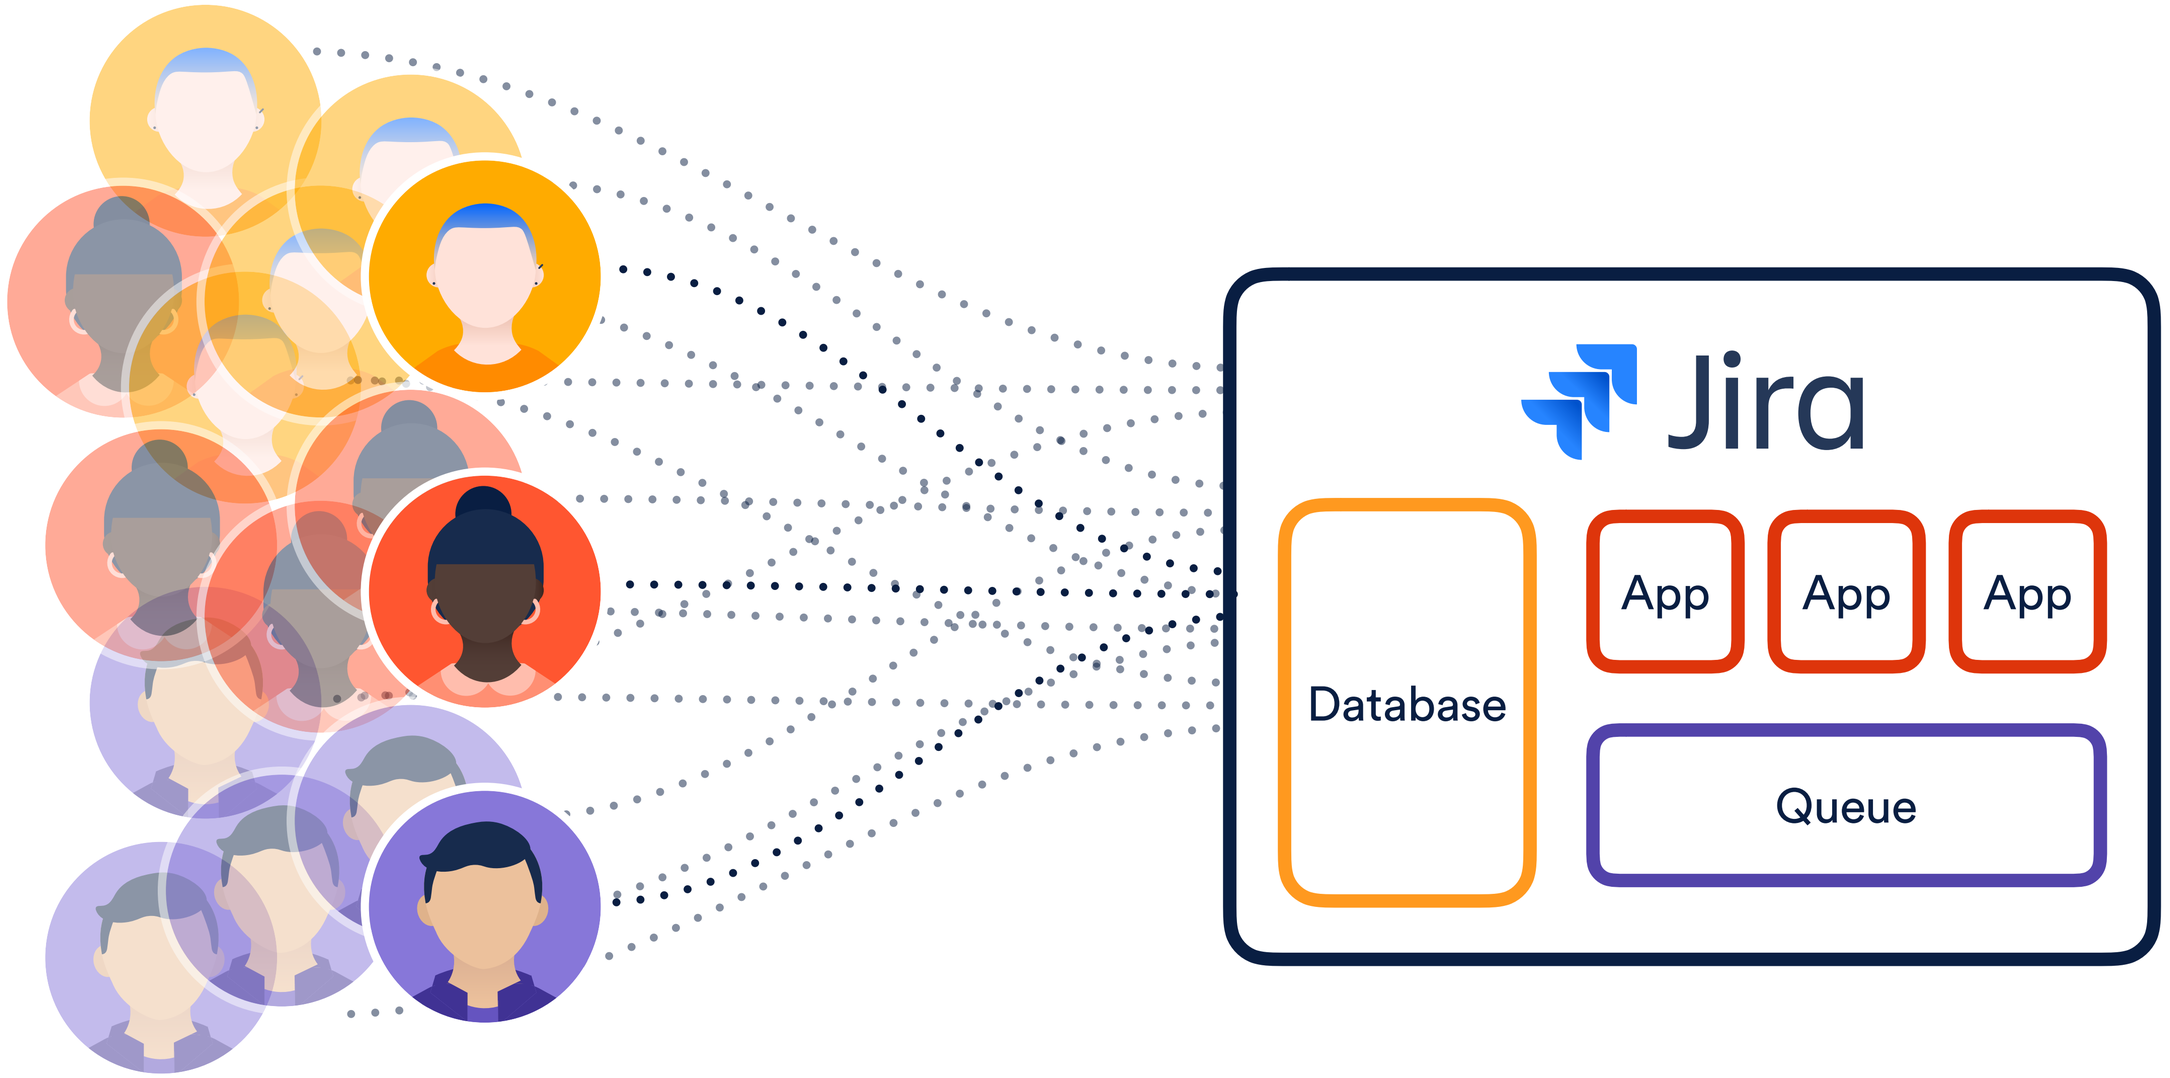

What about **cross-region latency?**

What about **outage blast radius?** What about **data sovereignty?**

What about **scale?** What about **progressive rollouts?**

What about **noisy neighbours?**

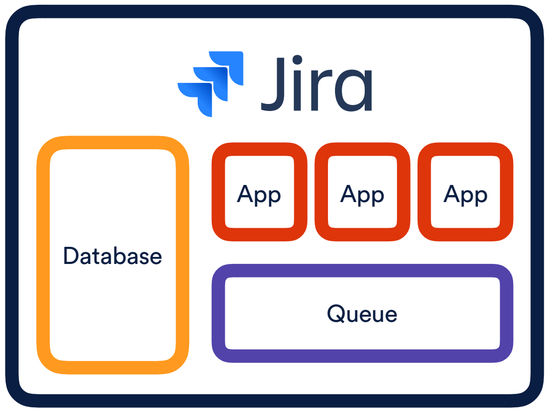
Jira
Database
App
App
App
Queue

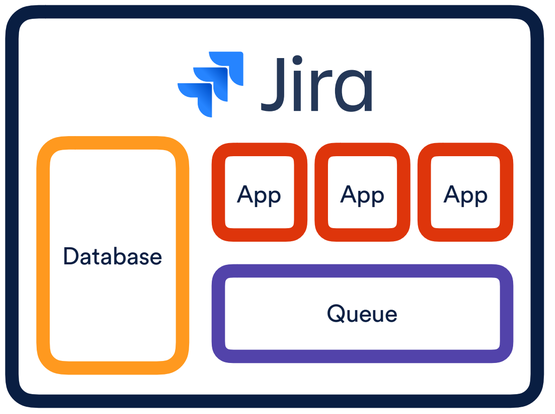
Jira
Database
App
App
App
Queue

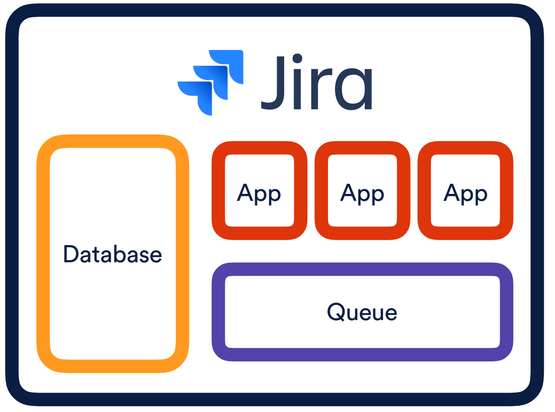
Jira
Database
App
App
App
Queue

Shard

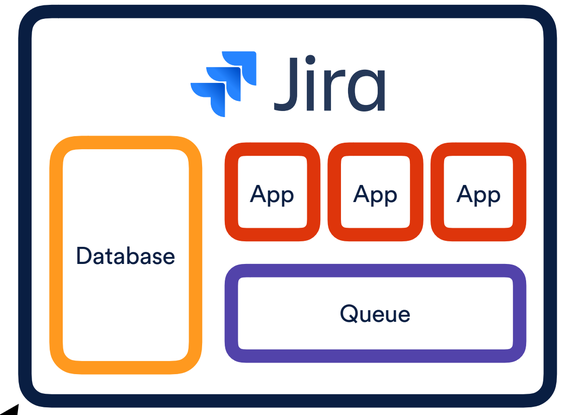
Jira
Database
App
App
App
Queue

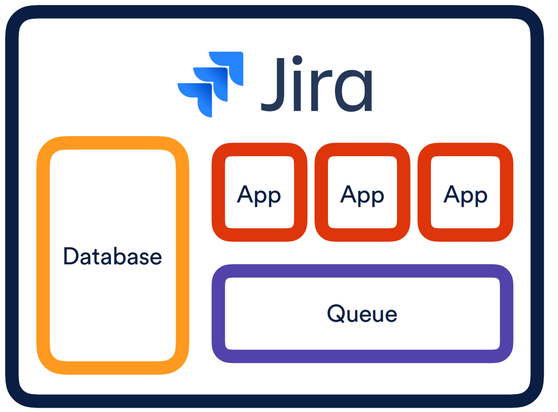
Jira
Database
App
App
App
Queue

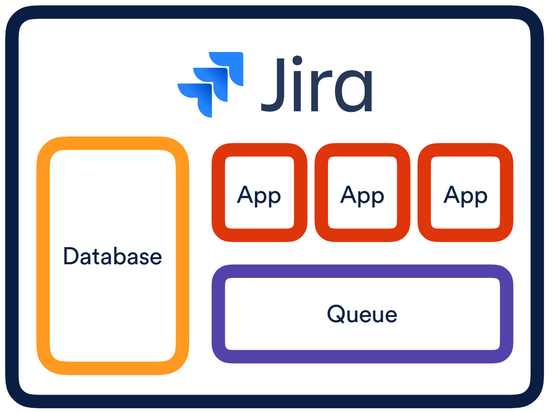
Jira
Database
App
App
App
Queue

Provisioning pipeline

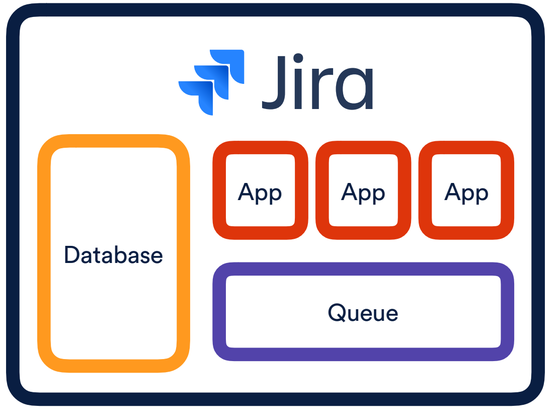

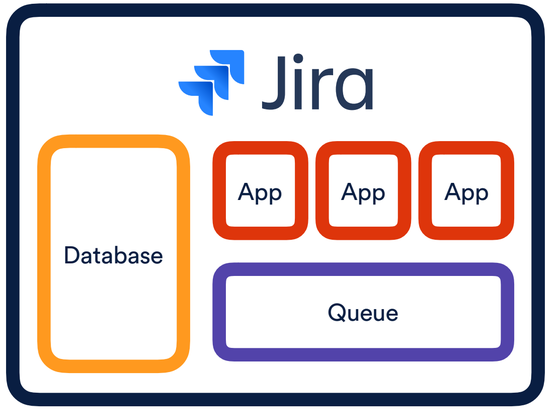

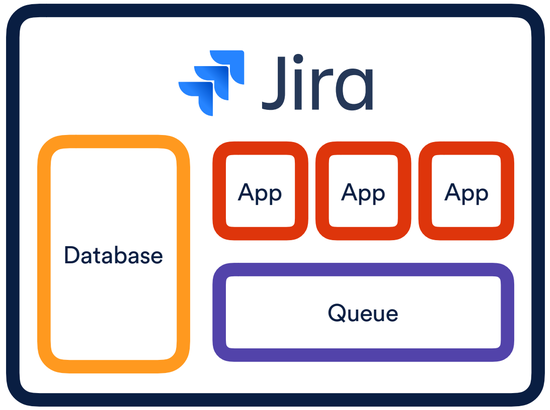

Provisioning pipeline

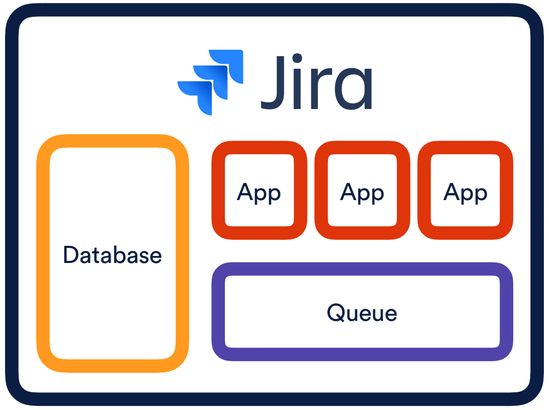

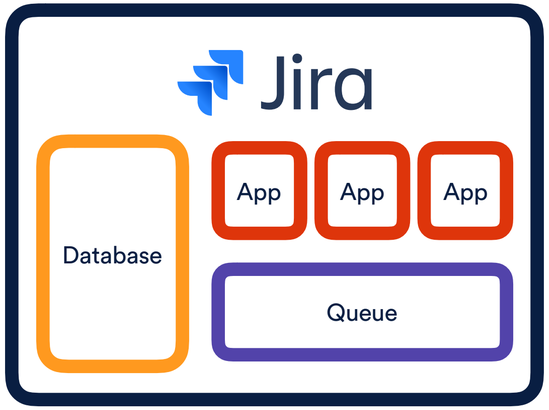

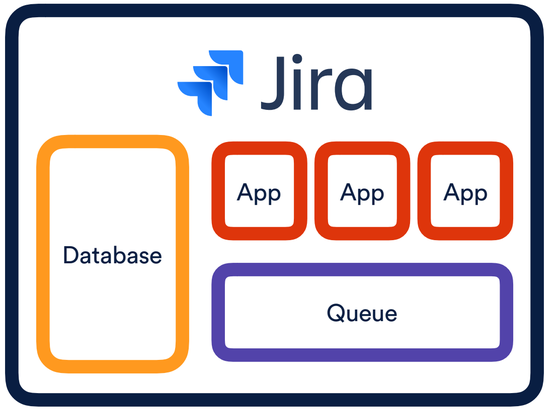

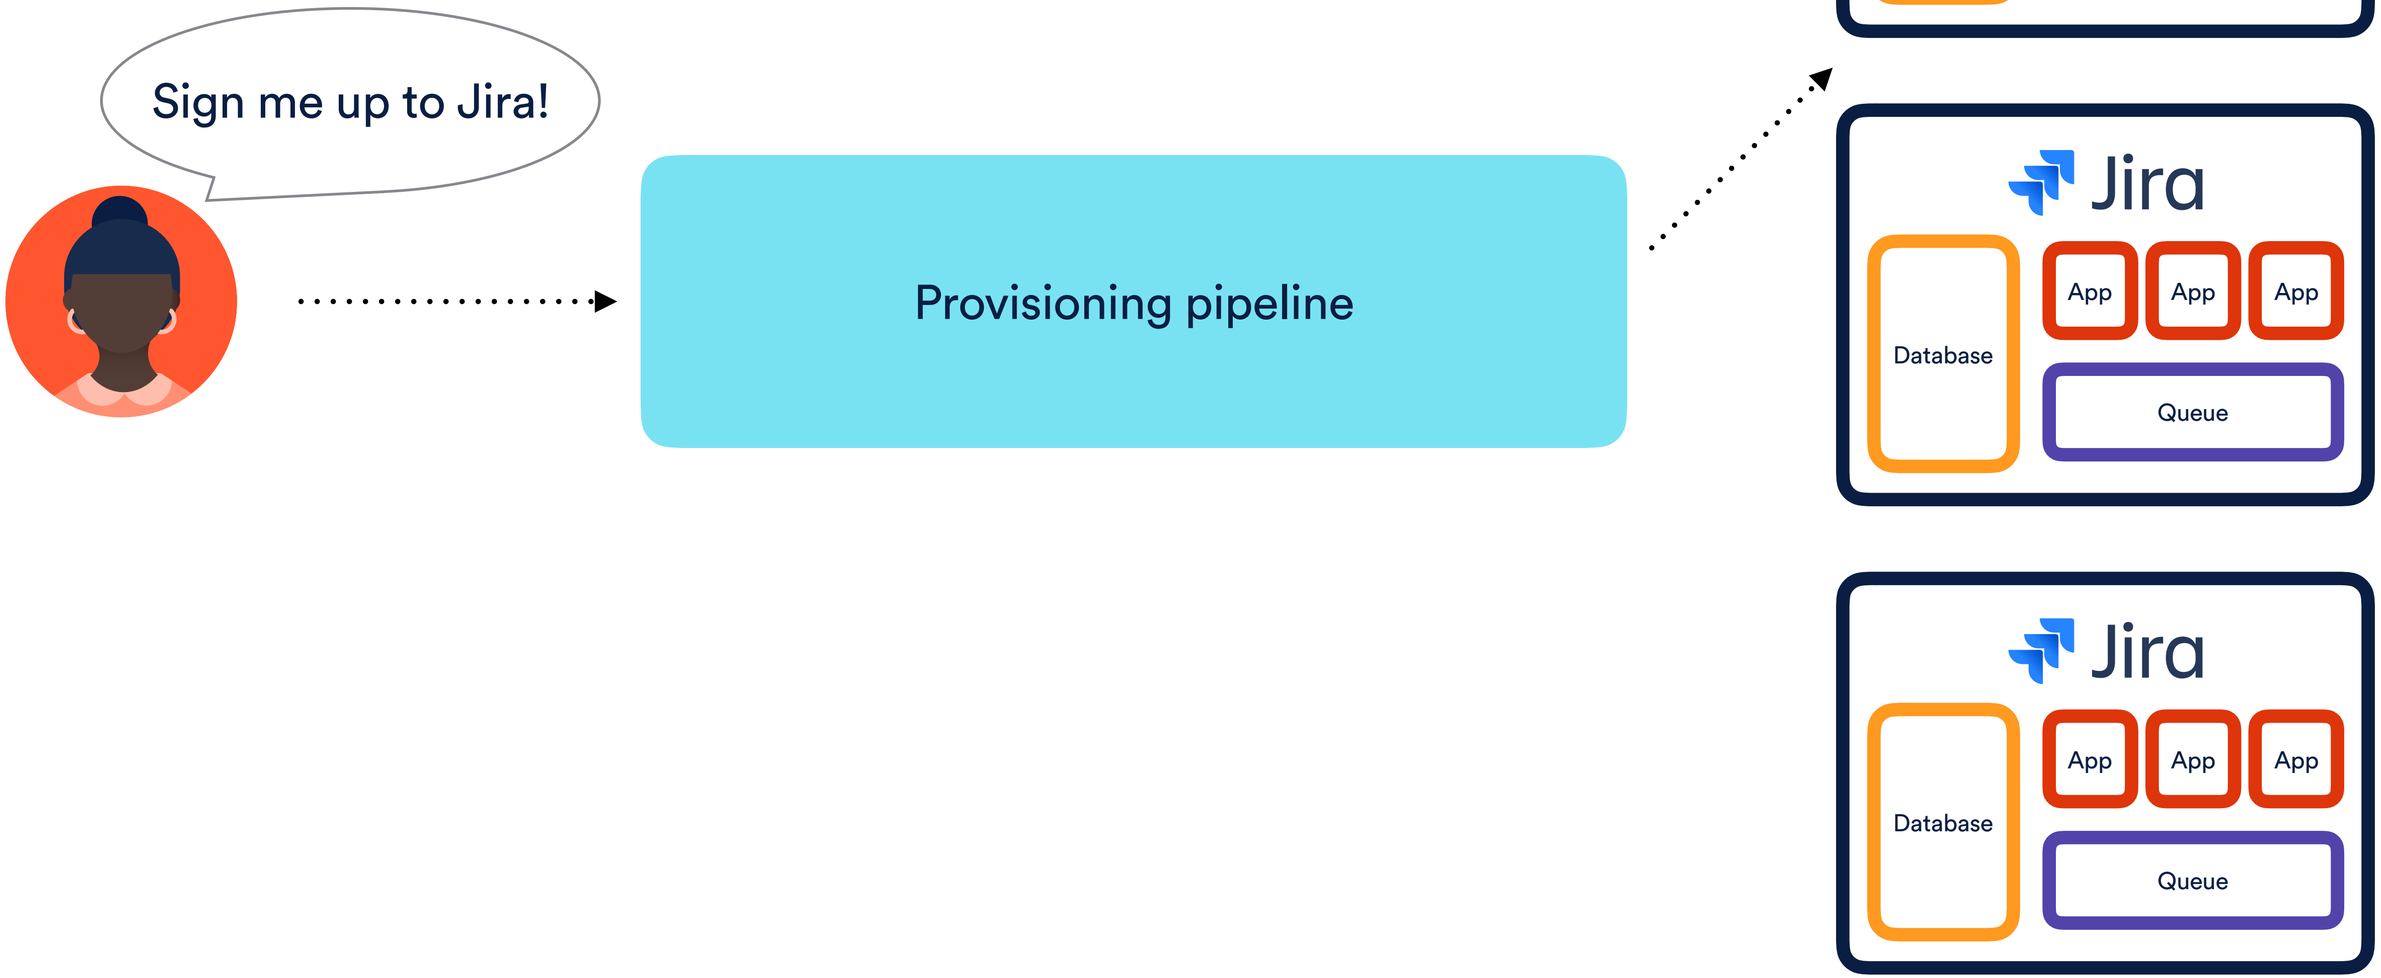
Sign me up to Jira!
Provisioning pipeline
Jira
Database
App
App
App
Queue
Jira
Database
App
App
App
Queue
Jira
Database
App
App
App
Queue

Provisioning pipeline

Jira

Database

App

App

App

Queue

Jira

Database

App

App

App

Queue

Jira

Database

App

App

App

Queue

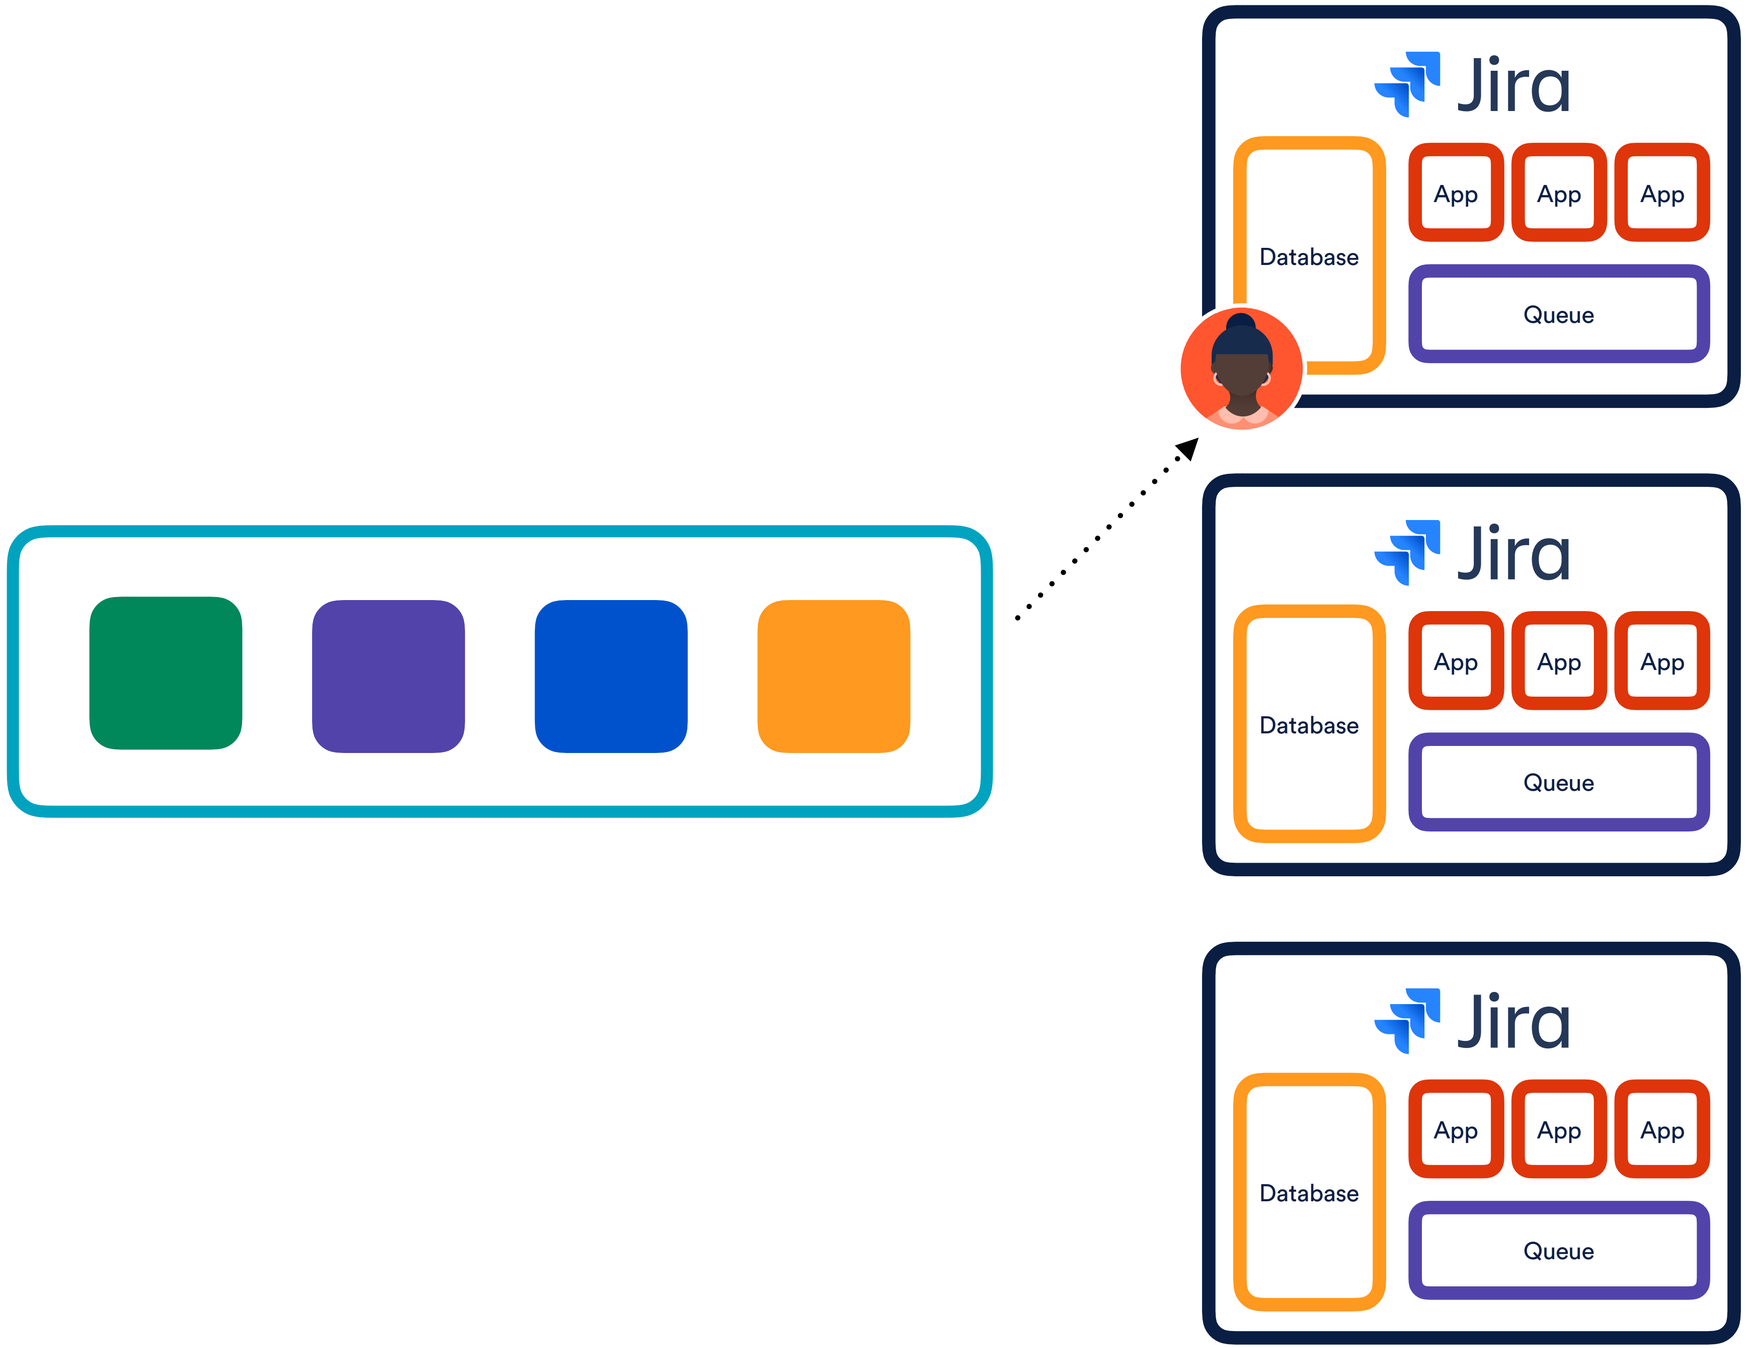
Jira
Database
App
App
App
Queue
Jira
Database
App
App
App
Queue
Jira
Database
App
App
App
Queue

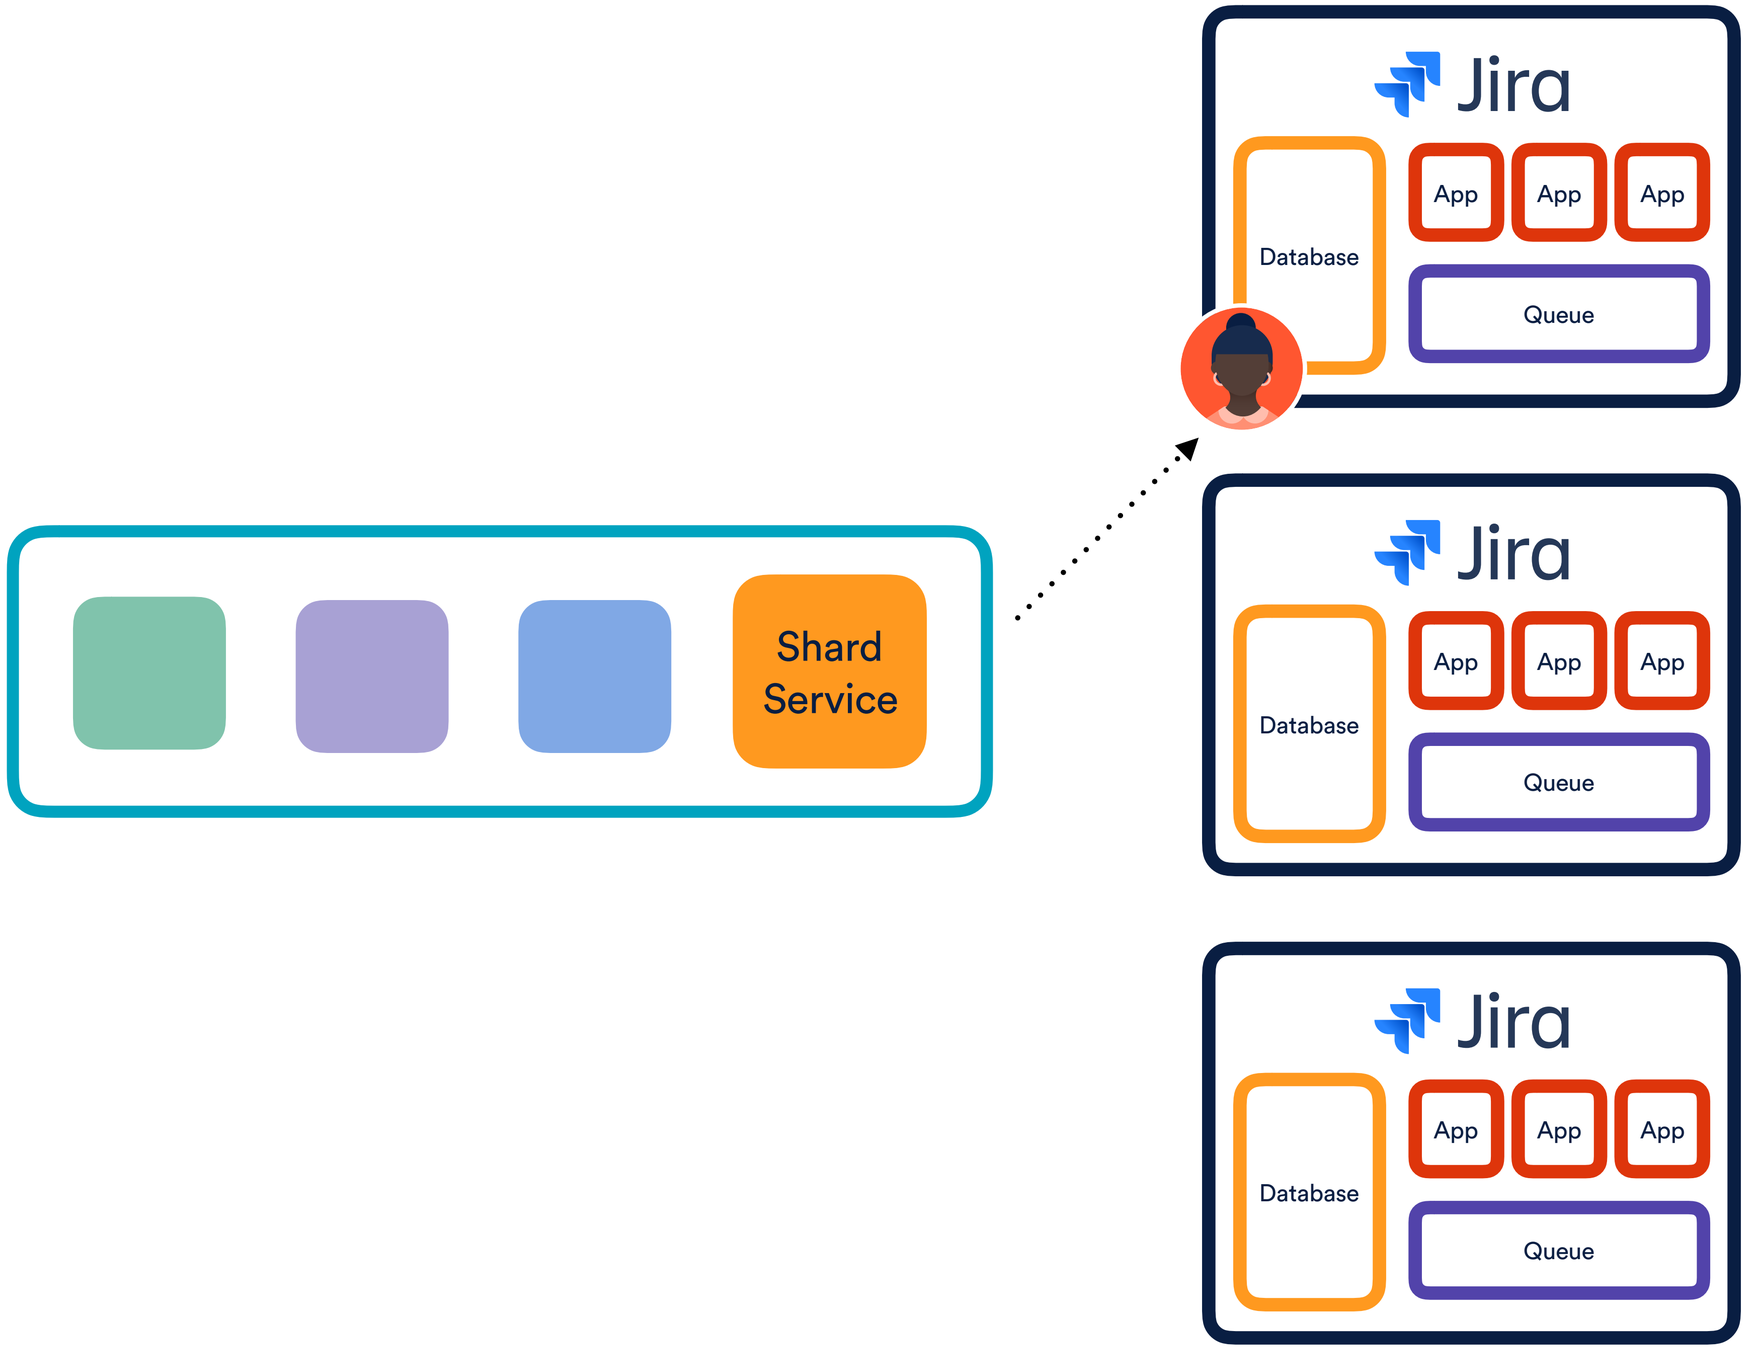
Shard Service
Jira
Database
App
App
App
Queue
Jira
Database
App
App
App
Queue
Jira
Database
App
App
App
Queue

Shard
Service

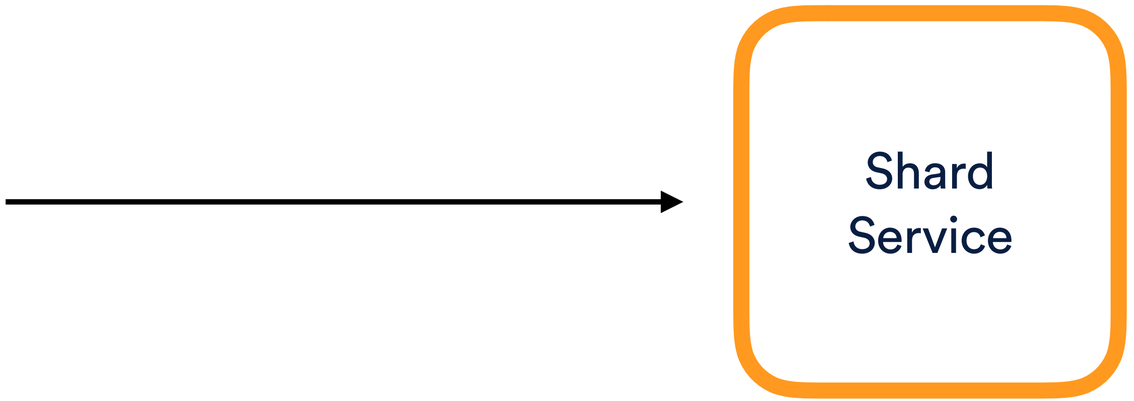
Shard
Service


Jira
App
App
App
Database
Queue
Jira
App
App
App
Database
Queue
Jira
App
App
App
Database
Queue
Jira
App
App
App
Database
Queue
Jira
App
App
App
Database
Queue
Jira
App
App
App
Database
Queue
Jira
App
App
App
Database
Queue
Jira
App
App
App
Database
Queue
Jira
App
App
App
Database
Queue
Jira
App
App
App
Database
Queue

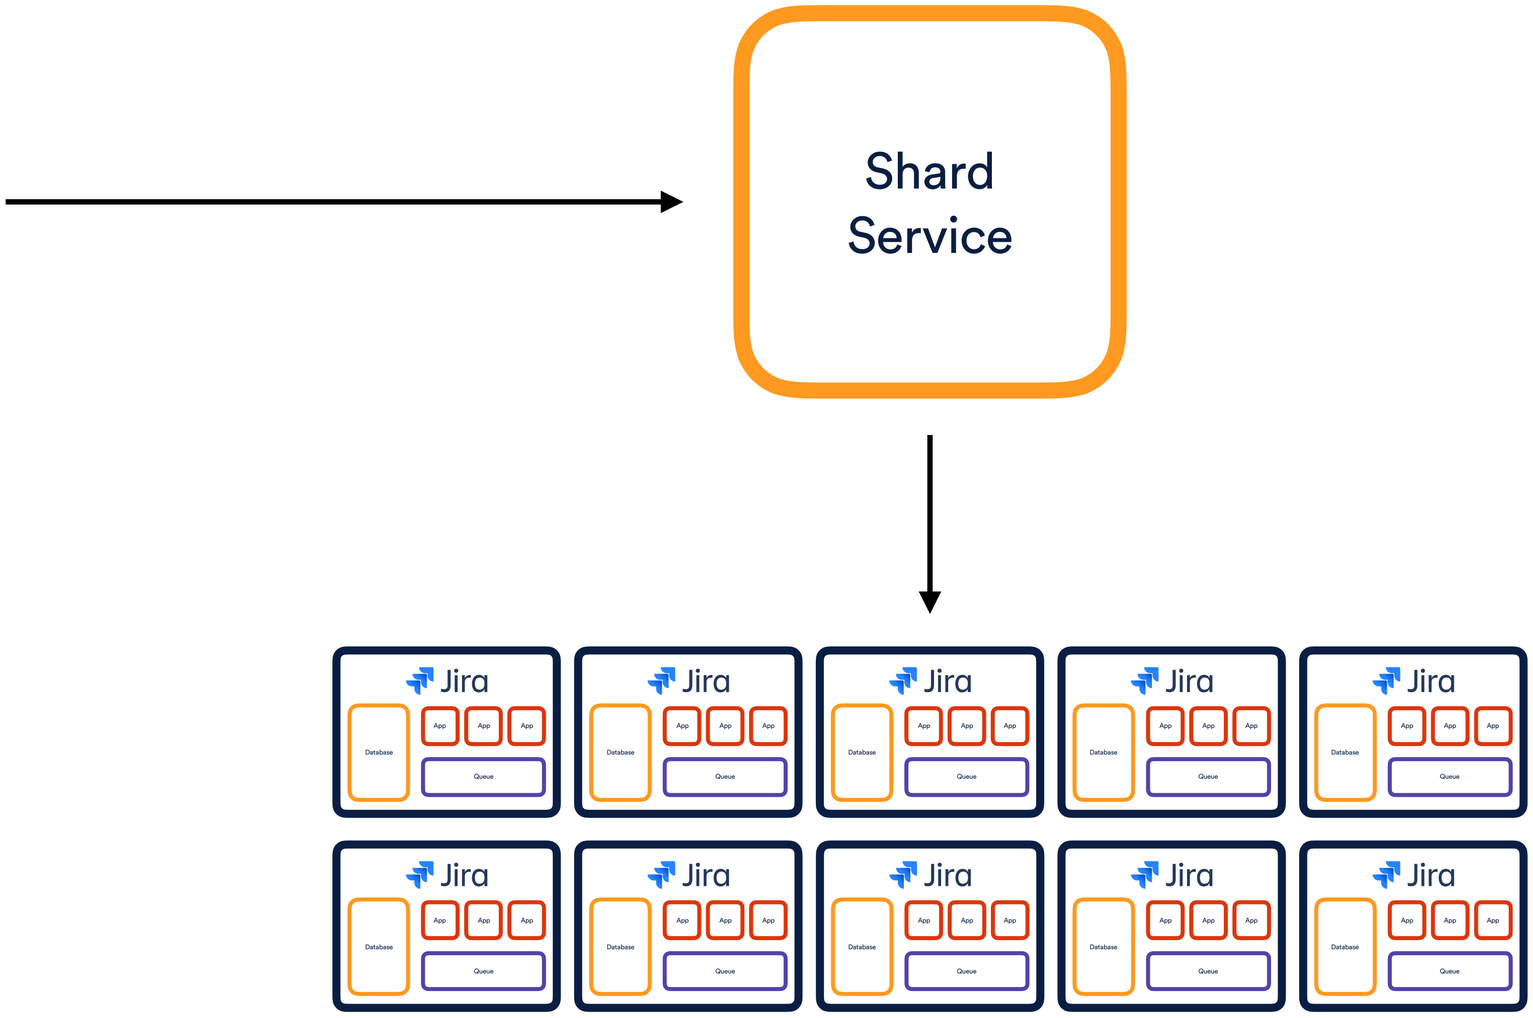
Shard
Service
Jira
App
App
App
Database
Queue
Jira
App
App
App
Database
Queue
Jira
App
App
App
Database
Queue
Jira
App
App
App
Database
Queue
Jira
App
App
App
Database
Queue
Jira
App
App
App
Database
Queue
Jira
App
App
App
Database
Queue
Jira
App
App
App
Database
Queue
Jira
App
App
App
Database
Queue
Jira
App
App
App
Database
Queue

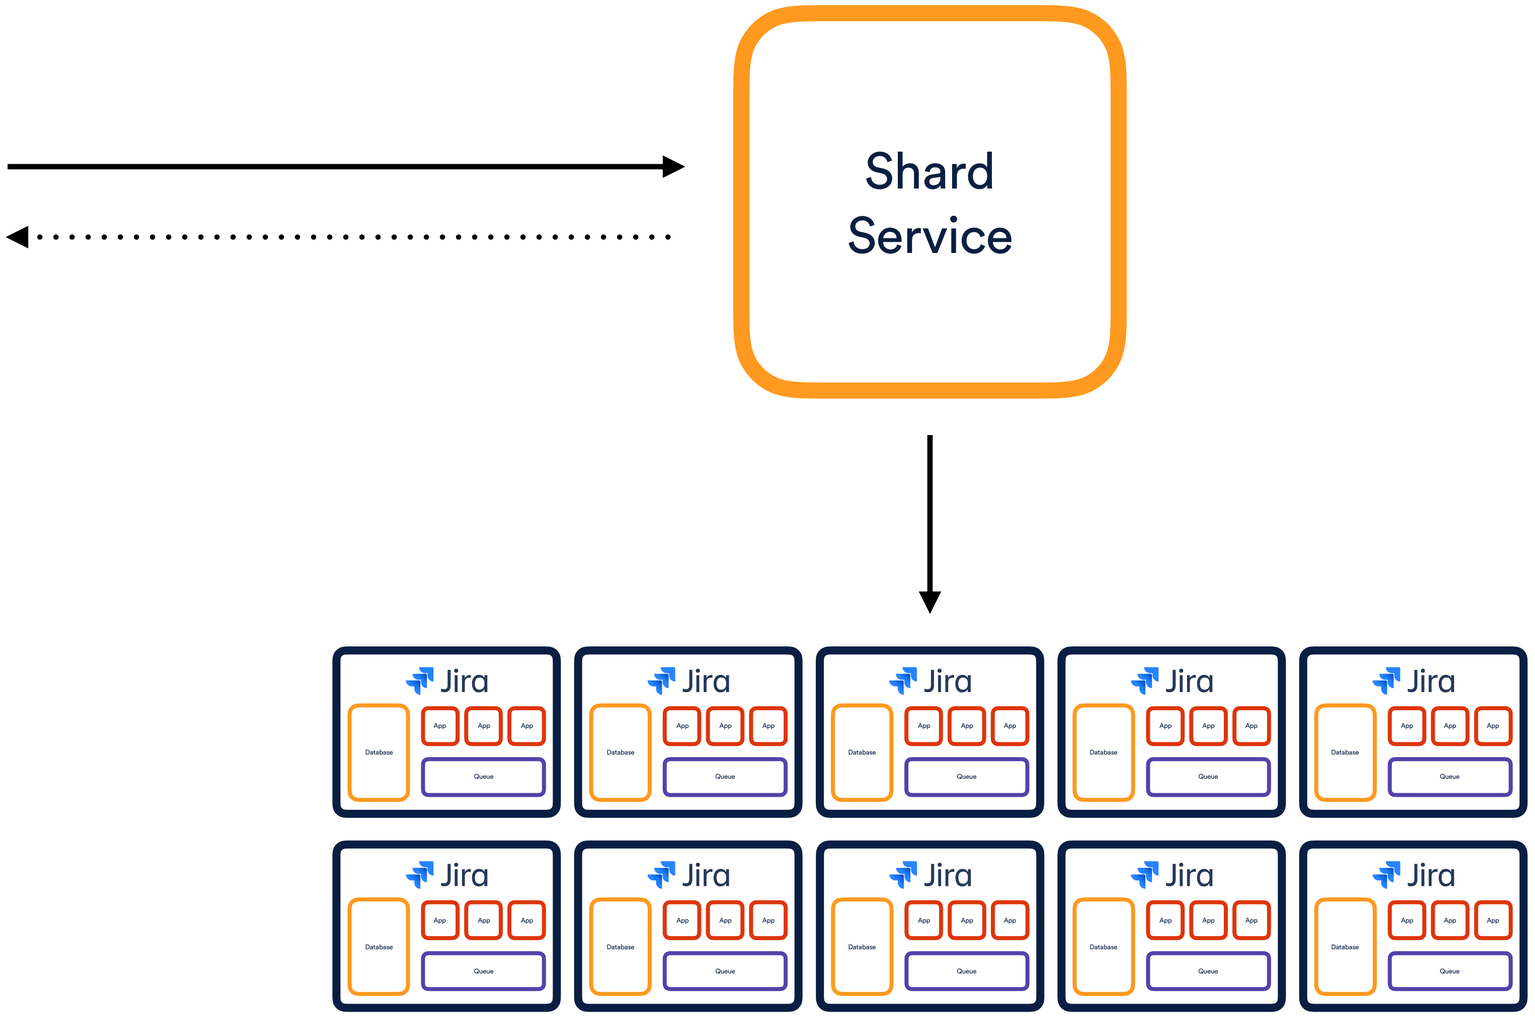
Shard
Service
Jira
App
App
App
Database
Queue

Shard Service

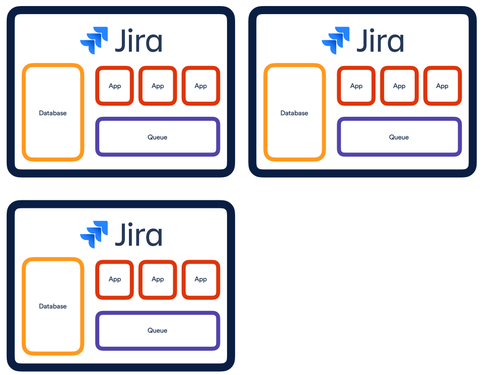

Europe

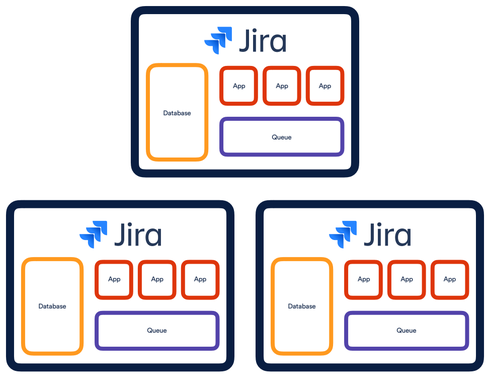

Australia

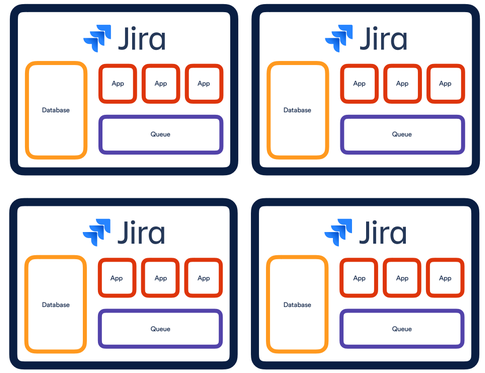

USA

Shard Service

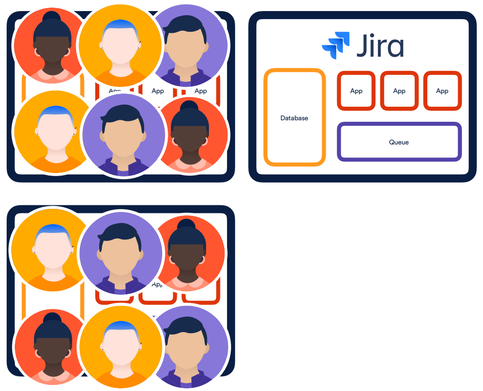

Europe

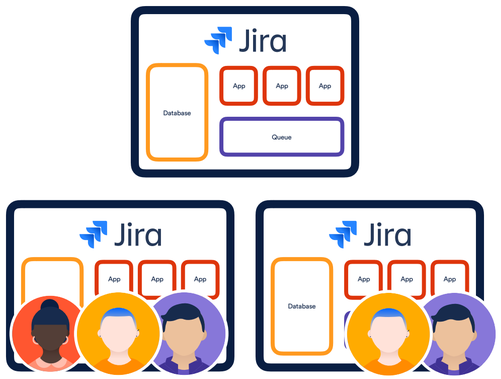

Australia

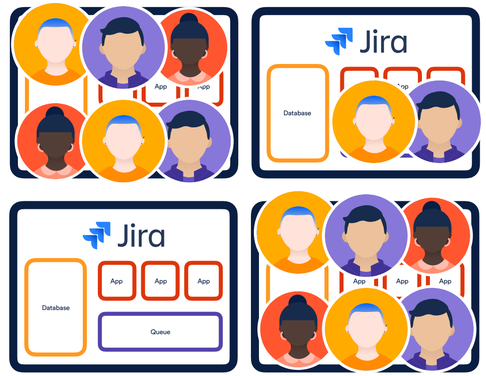

USA

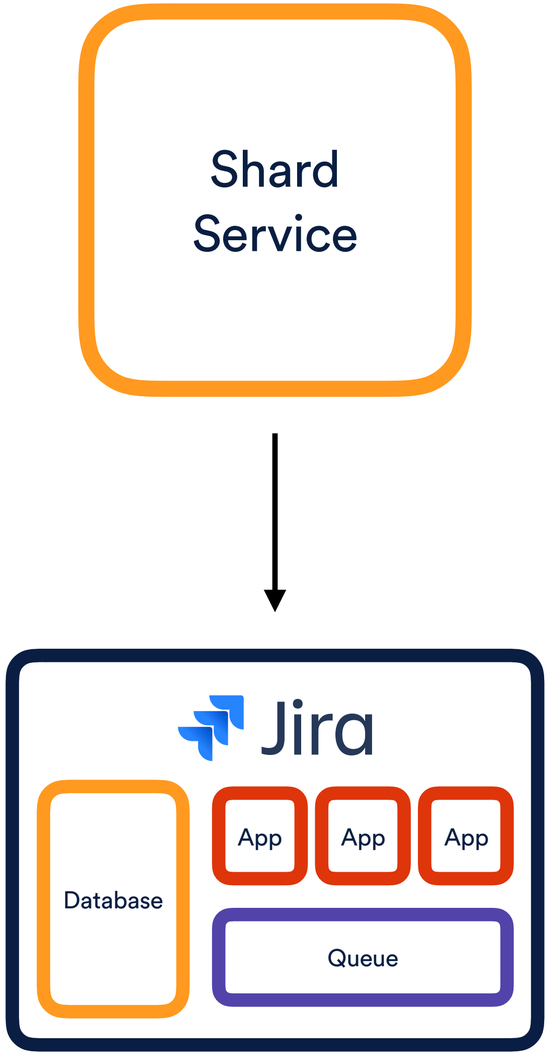
Shard
Service
Jira
Database
App
App
App
Queue

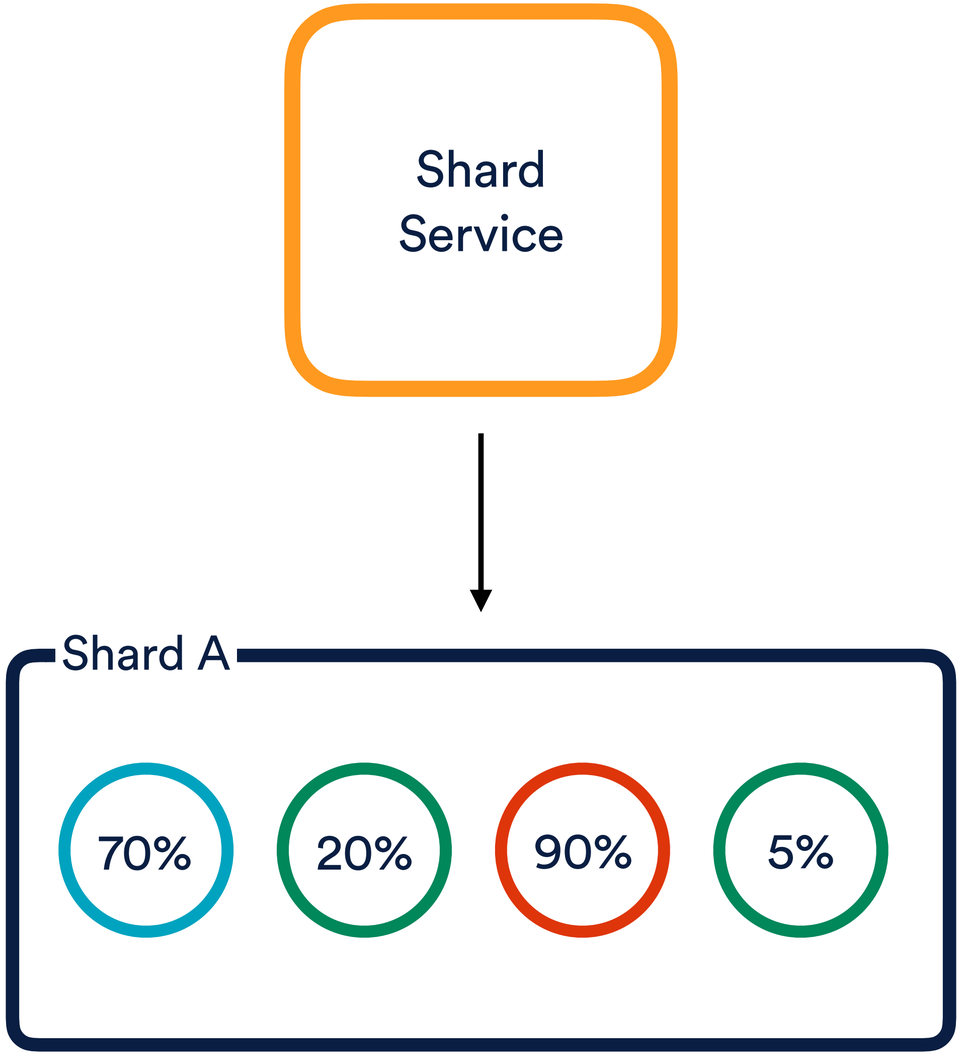
Shard
Service
Shard A
70%
20%
90%
5%

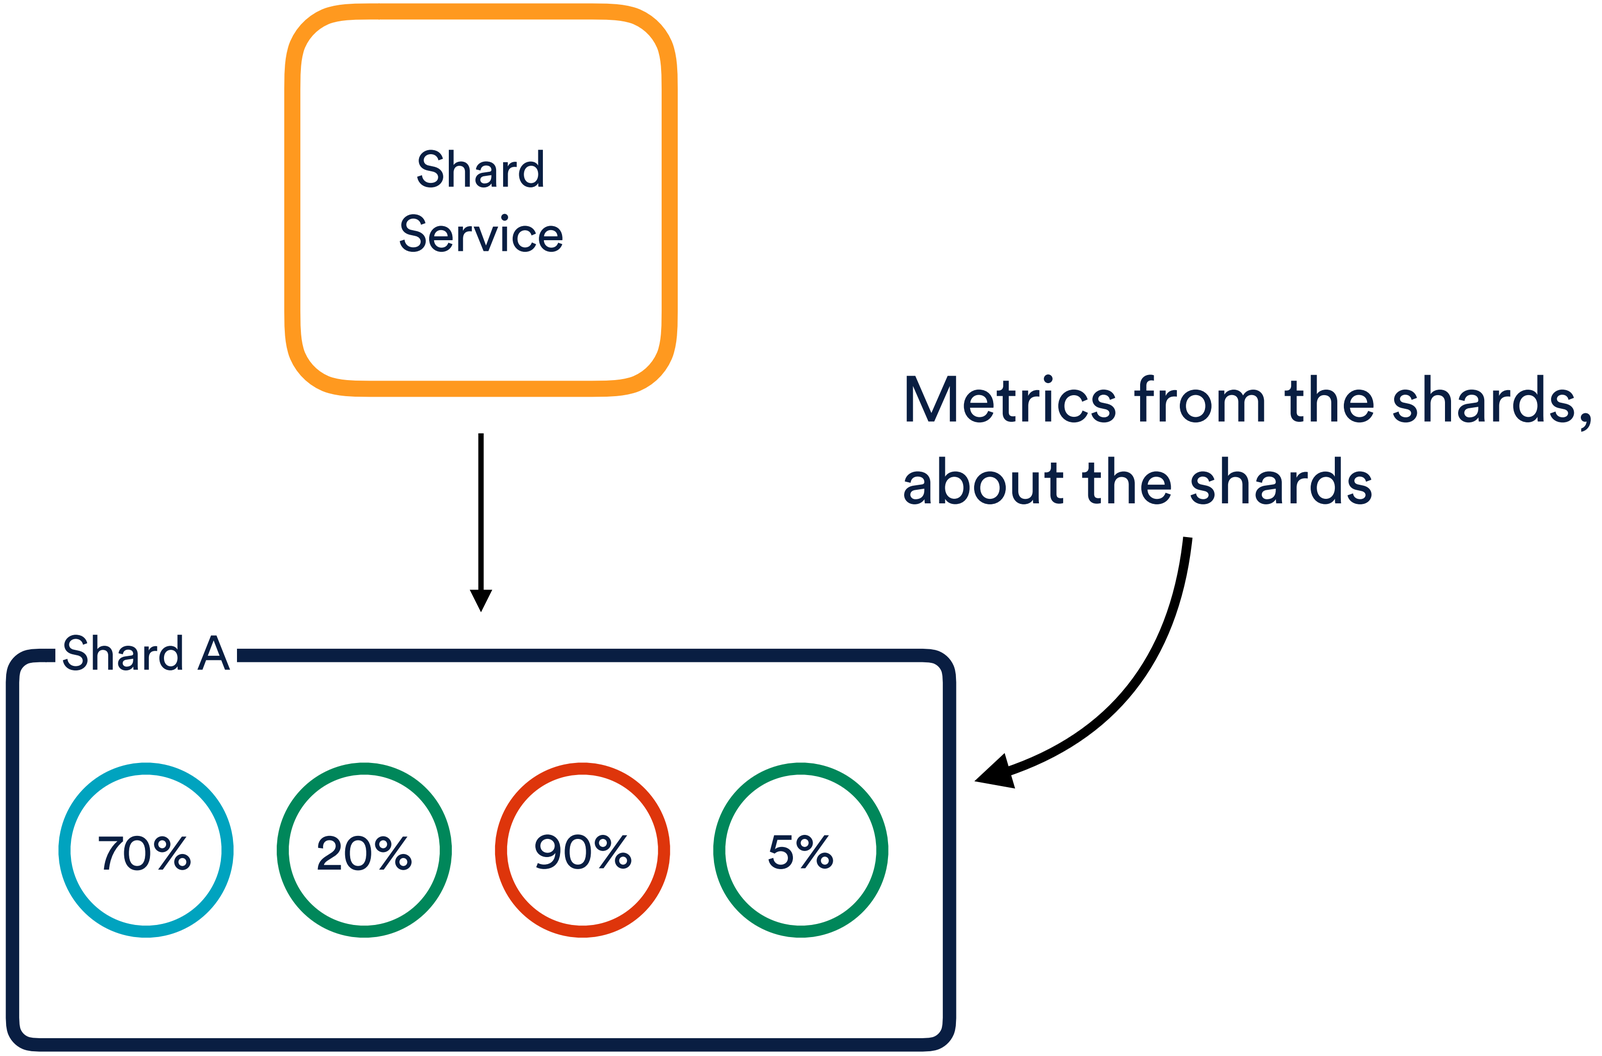
Shard
Service
Metrics from the shards,
about the shards
Shard A
70%
20%
90%
5%

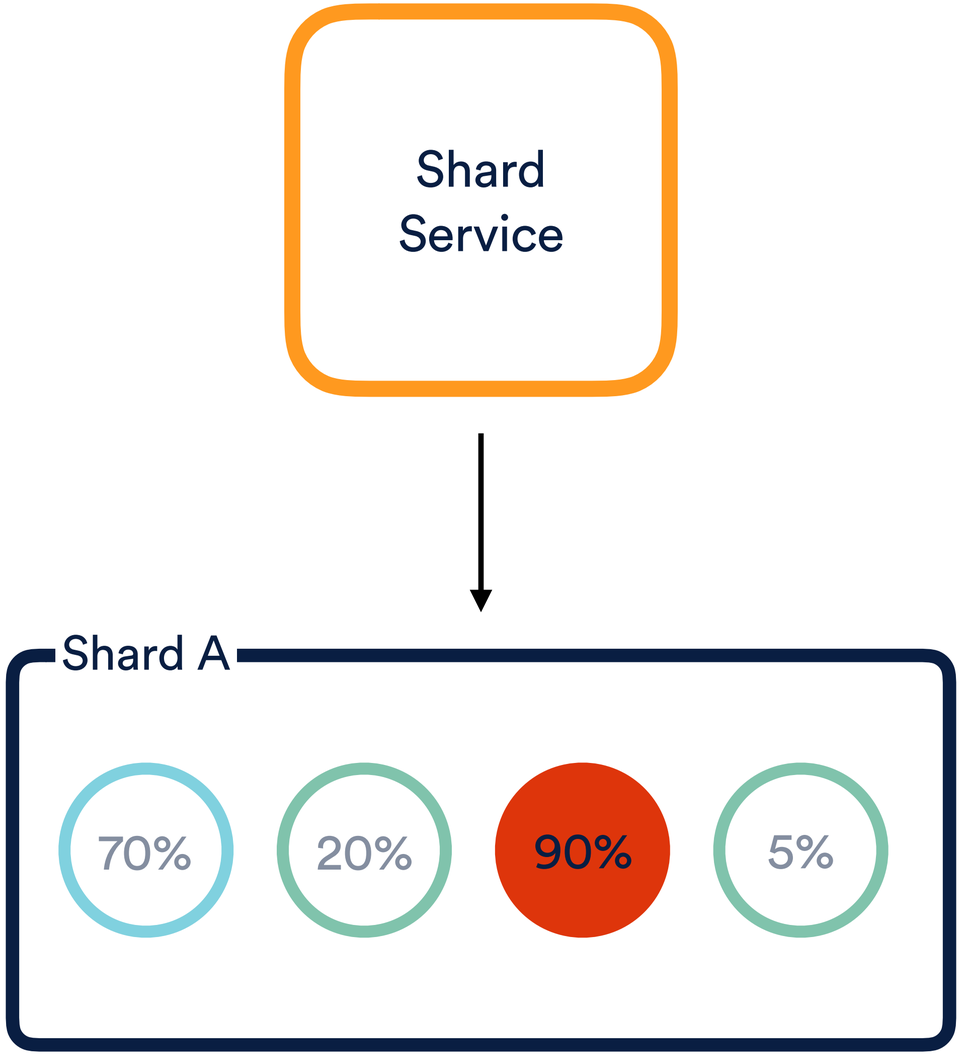
Shard
Service
Shard A
70%
20%
90%
5%

Wait...

Understanding what to measure is hard

What worked today, may not work tomorrow

Understanding what to measure is hard

Keeping everyone & everything up to date is hard

What worked today, may not work tomorrow

Understanding what to measure is hard

Agenda

Iterative... what?

Setting some context

Deciding what to measure

Verifying your metrics

Keeping up with change

Summary

Agenda

Iterative... what?

Setting some context

Deciding what to measure

Verifying your metrics

Keeping up with change

Summary

“

It is a capital mistake to theorise before one has data

SHERLOCK HOLMES

Measure nothing

Measure nothing

- Metrics aren’t verified before going live

Measure nothing

- Metrics aren’t verified before going live
- First incident is going to SUCK

Measure nothing

- Metrics aren’t verified before going live
- First incident is going to SUCK
- This isn’t a solution, it’s a deferral

Measure everything

Measure everything

- Expensive (time, money & resources)

Measure everything

- Expensive (time, money & resources)
- Lots of noise

Measure everything

- Expensive (time, money & resources)
- Lots of noise
- Does not scale

Measure stuff from out the box

SHARD SERVICE DASHBOARD

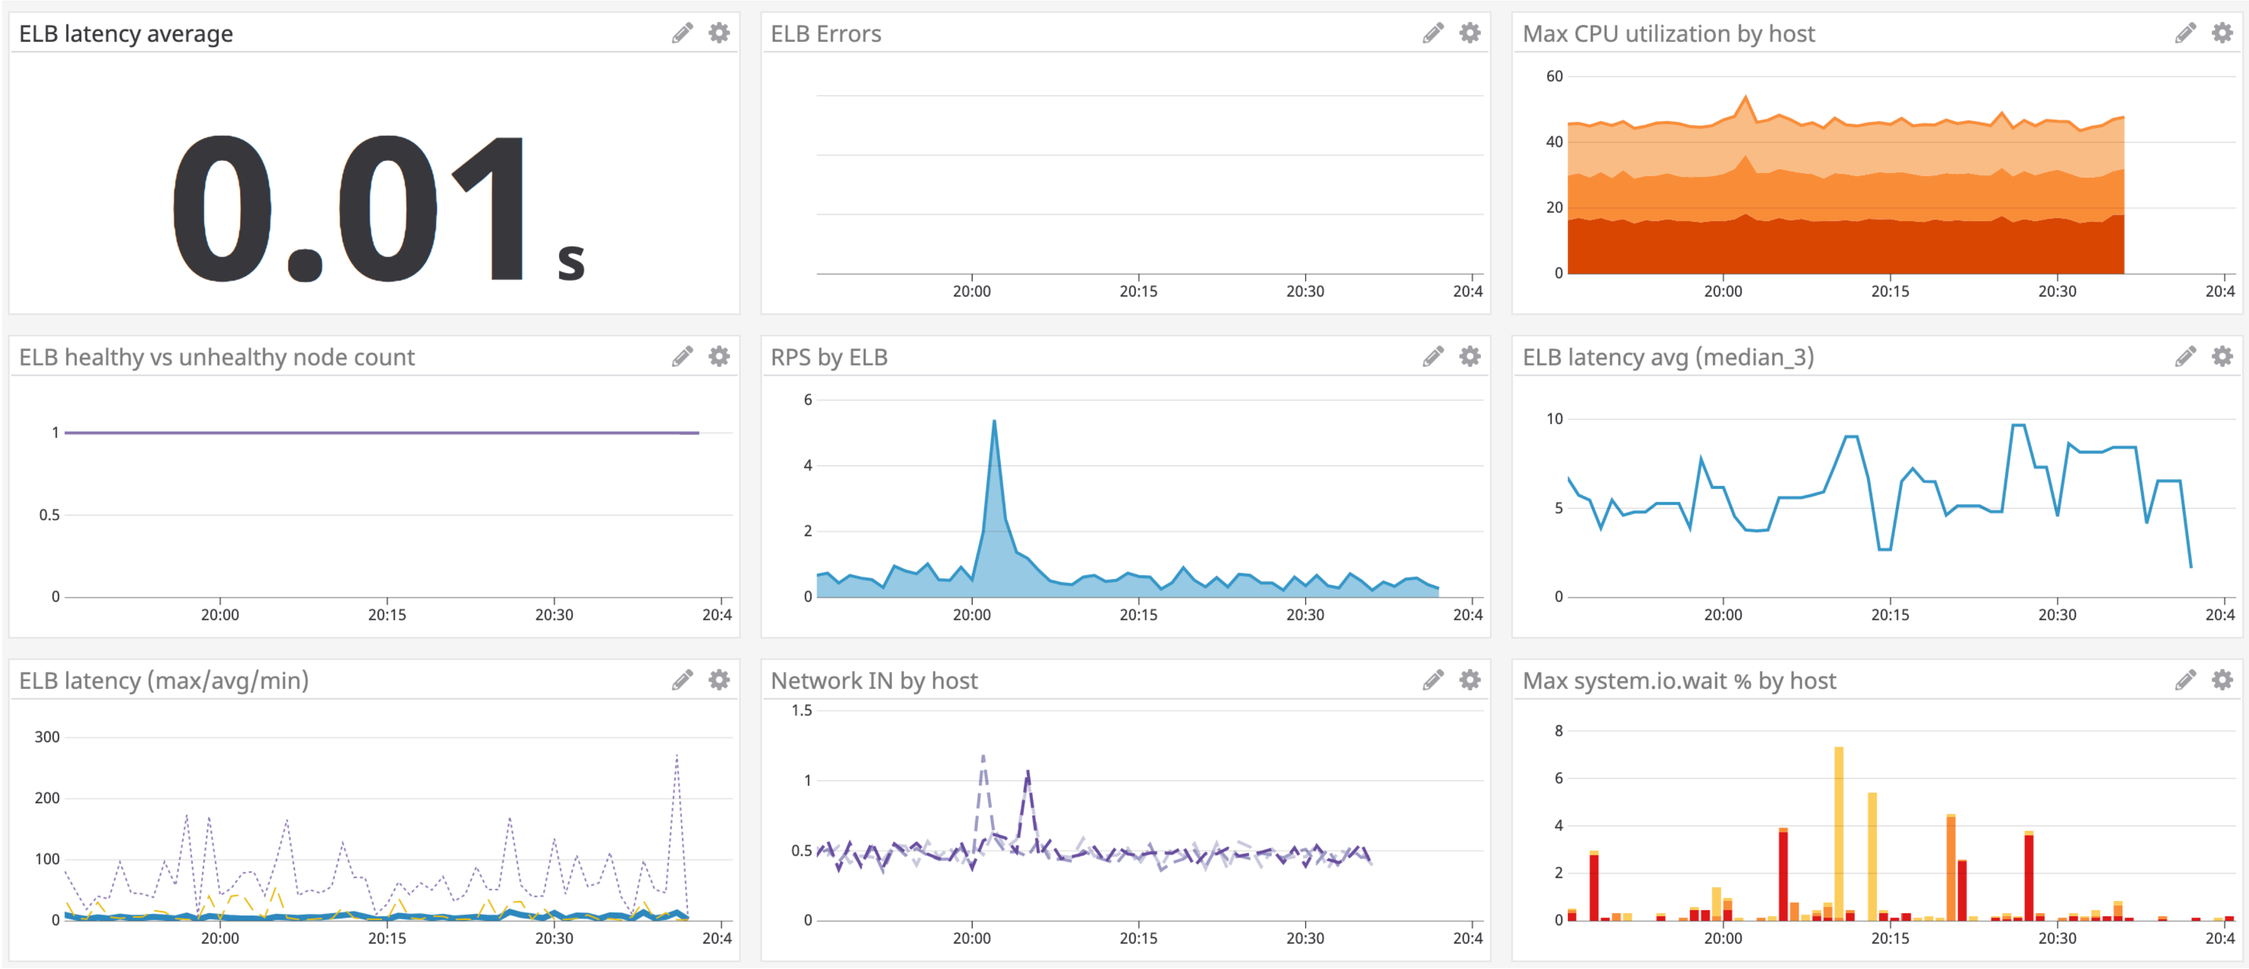

SHARD SERVICE DASHBOARD

ELB latency average

0.01 s

ELB Errors

Max CPU utilization by host

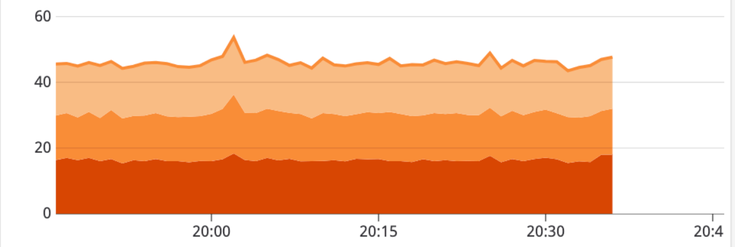

ELB healthy vs unhealthy node count

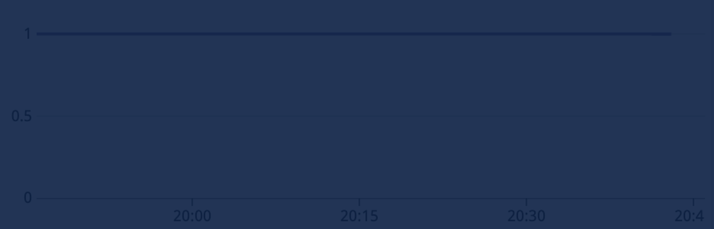

RPS by ELB

ELB latency avg (median_3)

ELB latency (max/avg/min)

Network IN by host

Max system.io.wait % by host

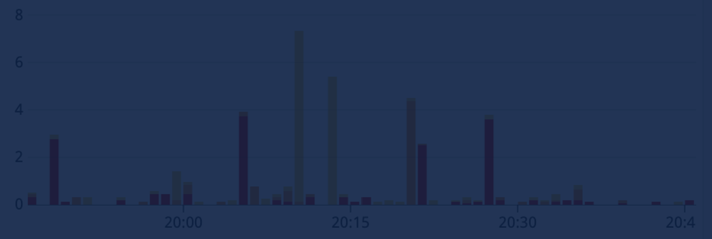

SHARD SERVICE DASHBOARD

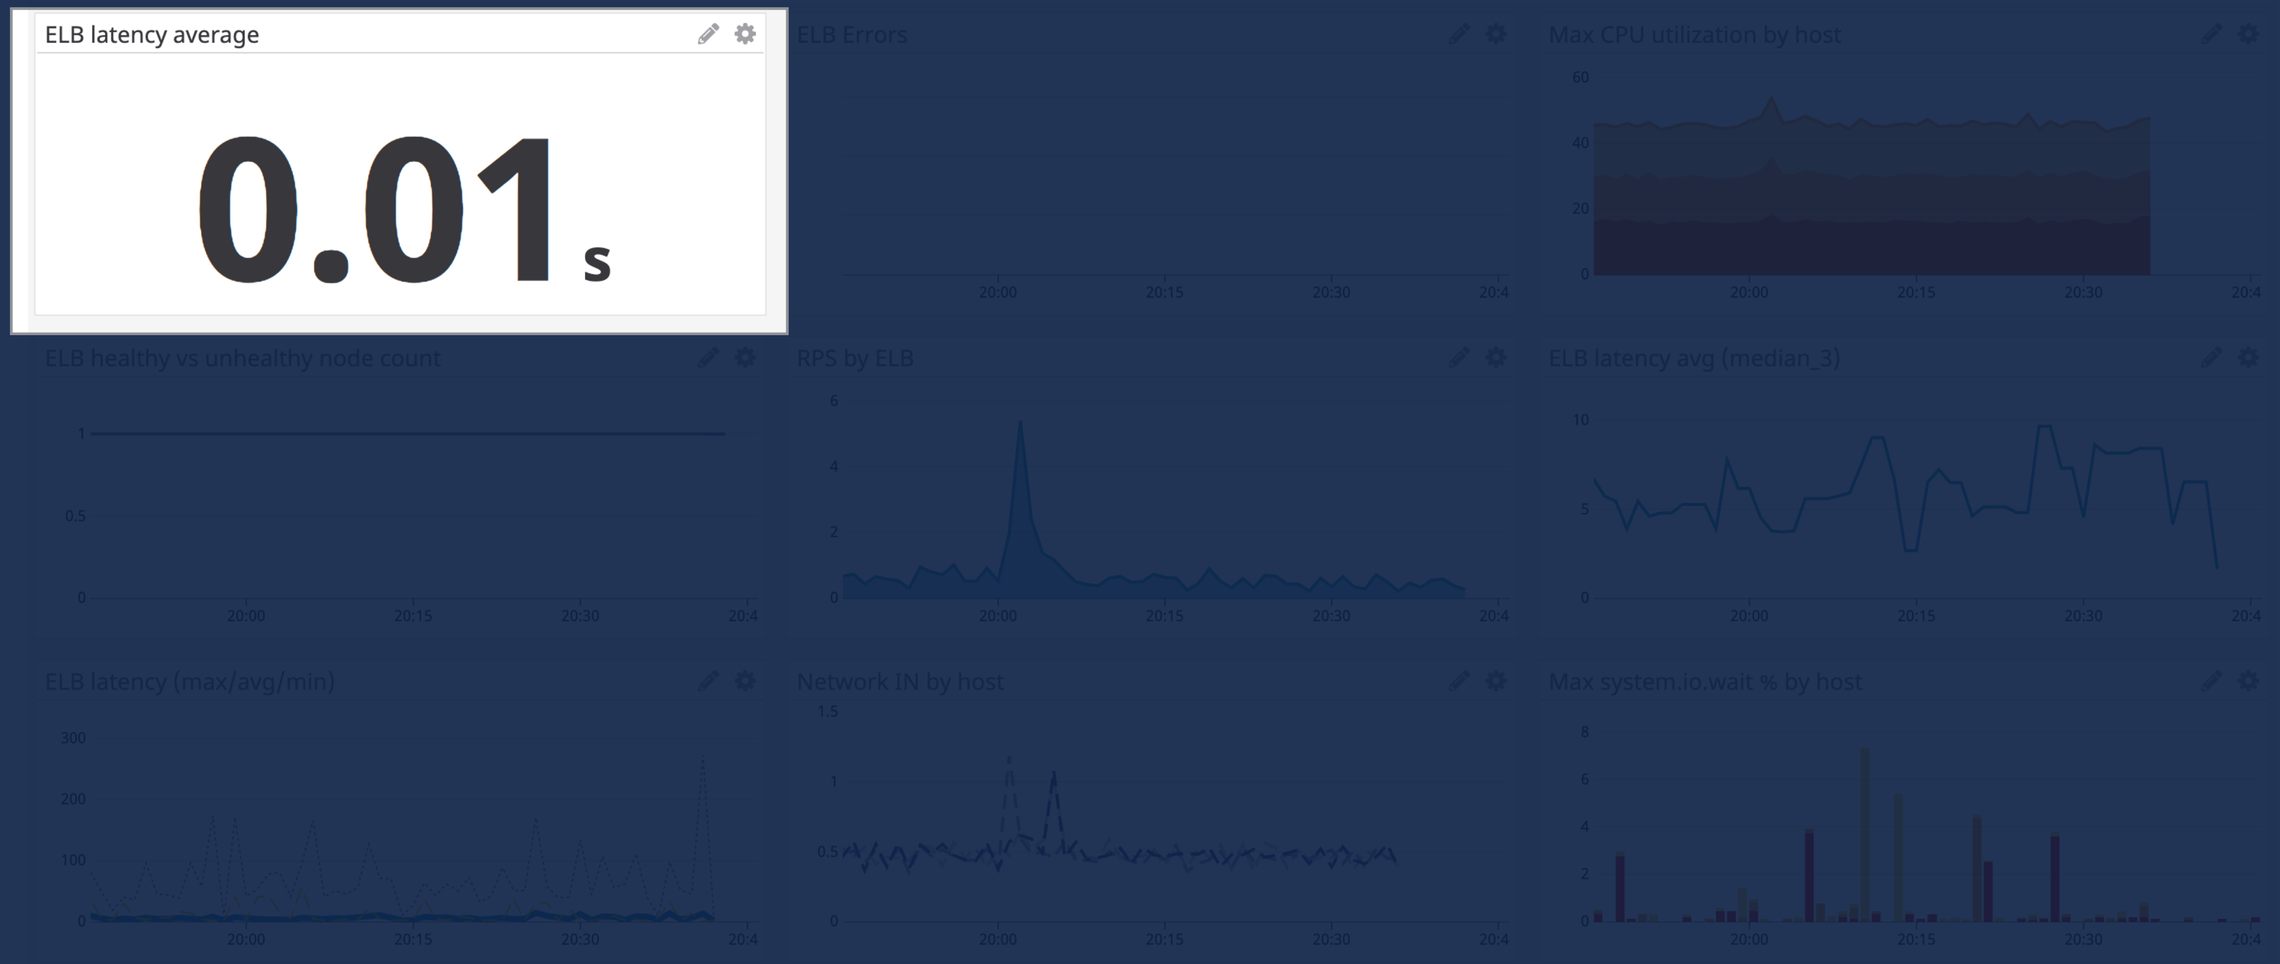

SHARD SERVICE DASHBOARD

ELB latency average

0.01 s

ELB Errors

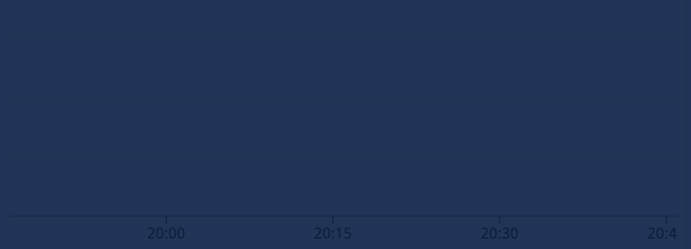

Max CPU utilization by host

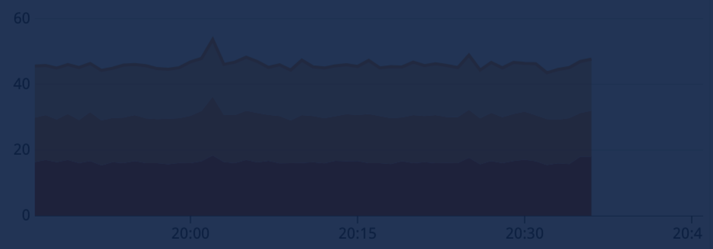

ELB healthy vs unhealthy node count

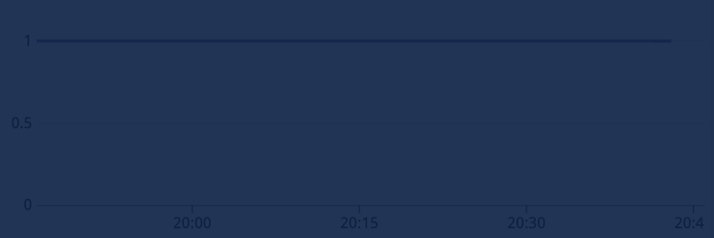

RPS by ELB

ELB latency avg (median_3)

ELB latency (max/avg/min)

Network IN by host

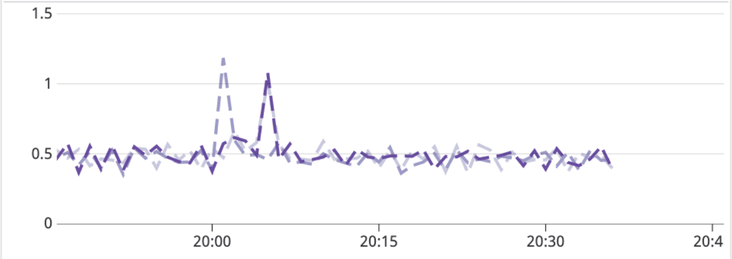

Max system.io.wait % by host

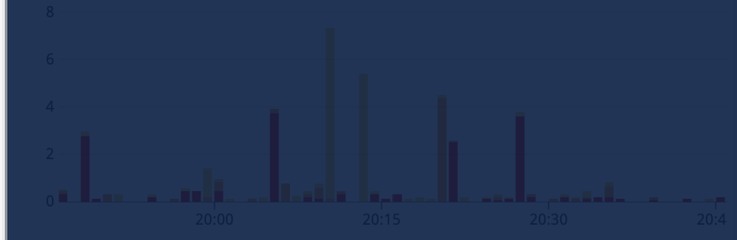

IF CPU > 80% for over 5 minutes THEN page

Think of idea

Design service

Build service (MVP)

Reach operational maturity

Release

Iterate service

Think of idea

Design service

Build service (MVP)

Reach operational maturity

Release

Iterate service

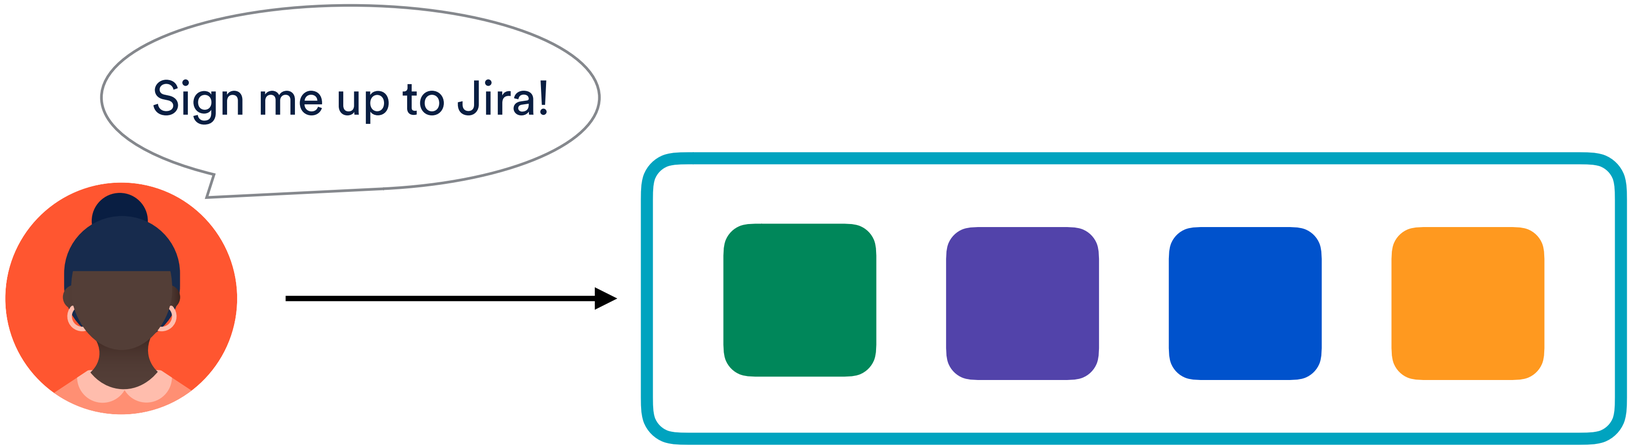
Sign me up to Jira!

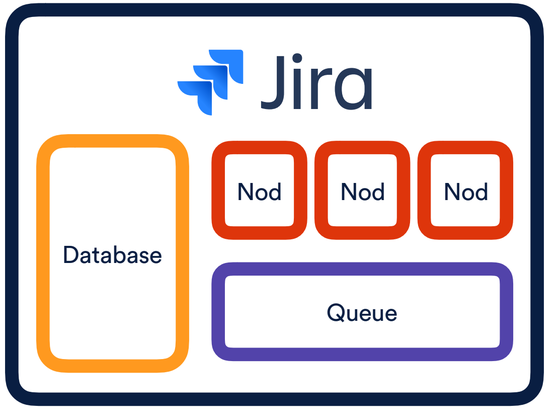
Jira
Database
Nod
Nod
Nod
Queue

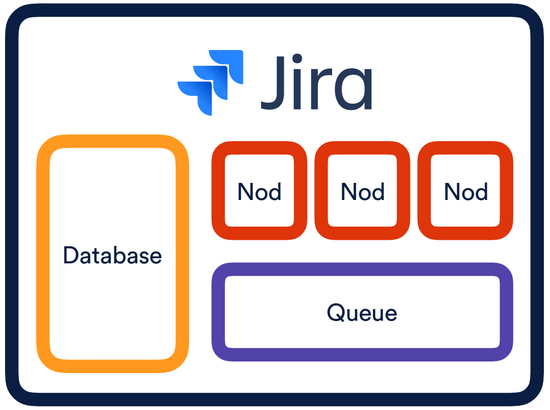
Jira
Database
Nod
Nod
Nod
Queue

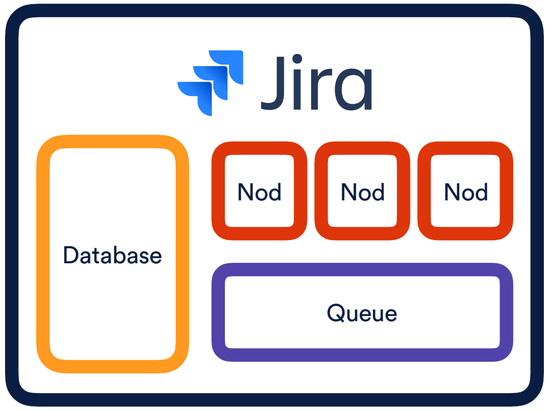
Jira
Database
Nod
Nod
Nod
Queue

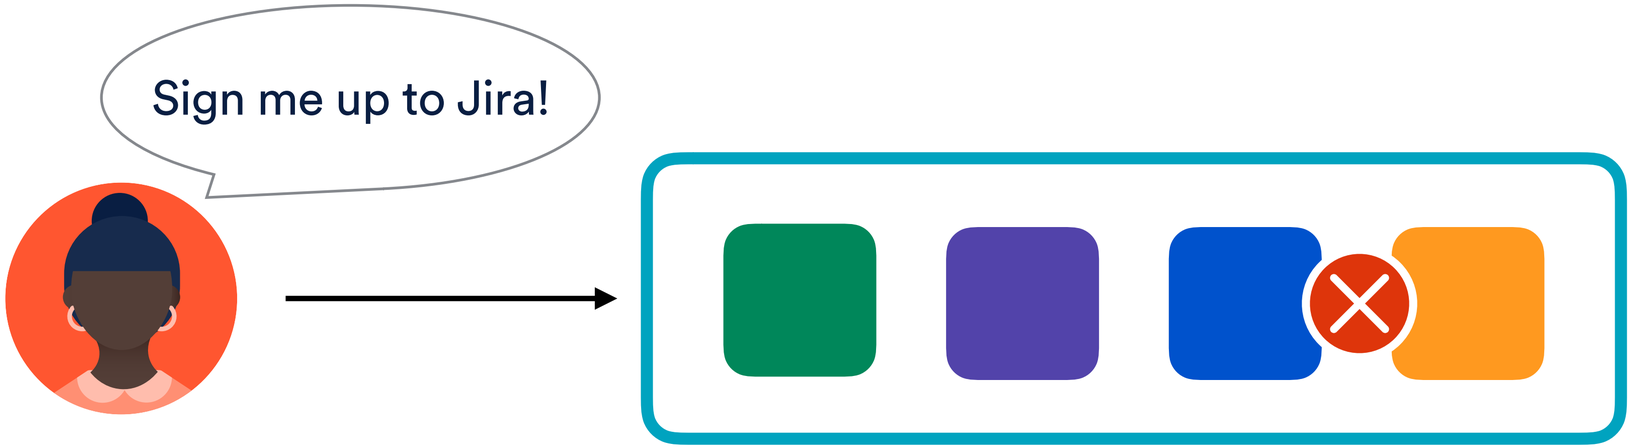
Sign me up to Jira!

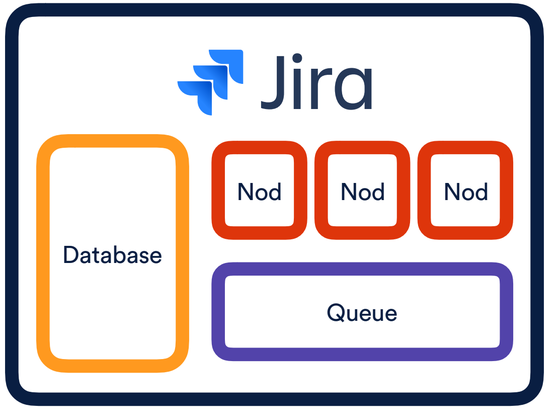
Jira
Database
Nod
Nod
Nod
Queue

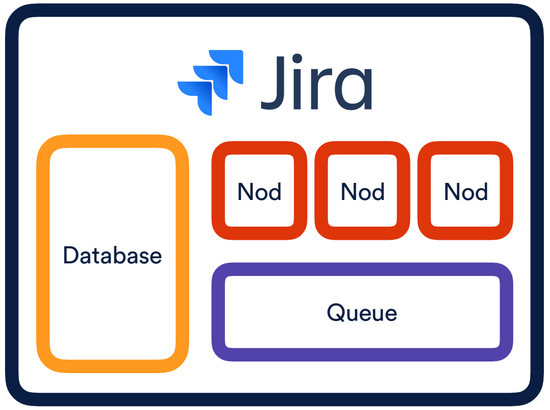
Jira
Database
Nod
Nod
Nod
Queue

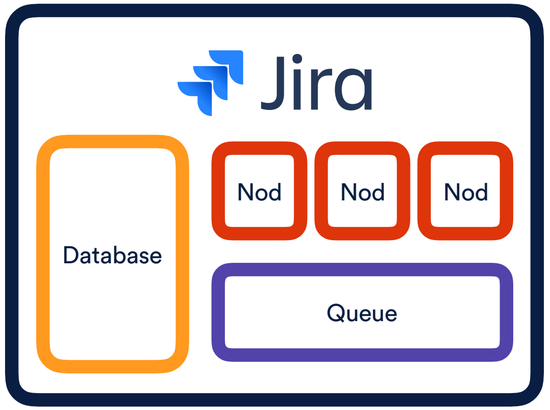
Jira
Database
Nod
Nod
Nod
Queue

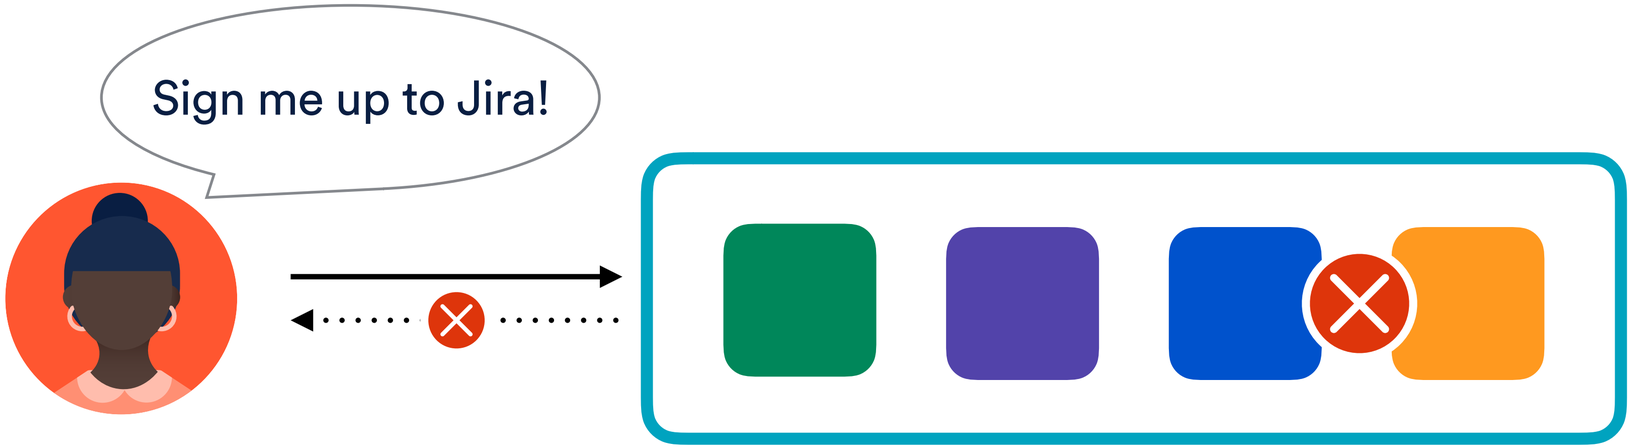
Sign me up to Jira!

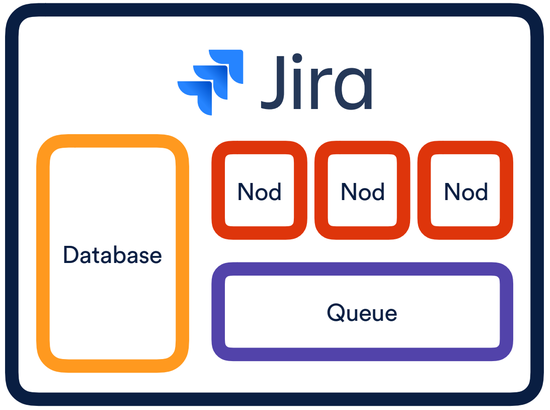
Jira
Database
Nod
Nod
Nod
Queue

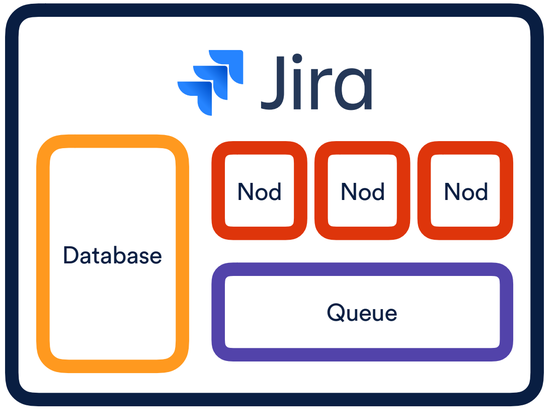
Jira
Database
Nod
Nod
Nod
Queue

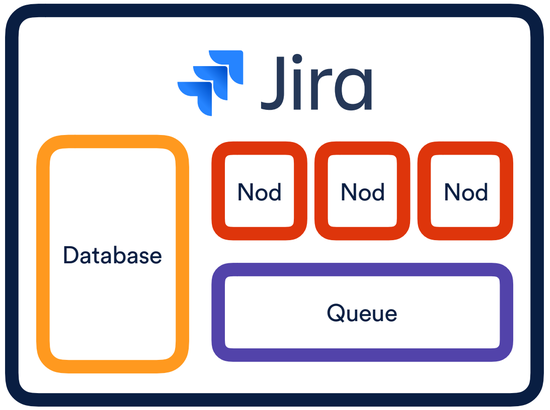
Jira
Database
Nod
Nod
Nod
Queue

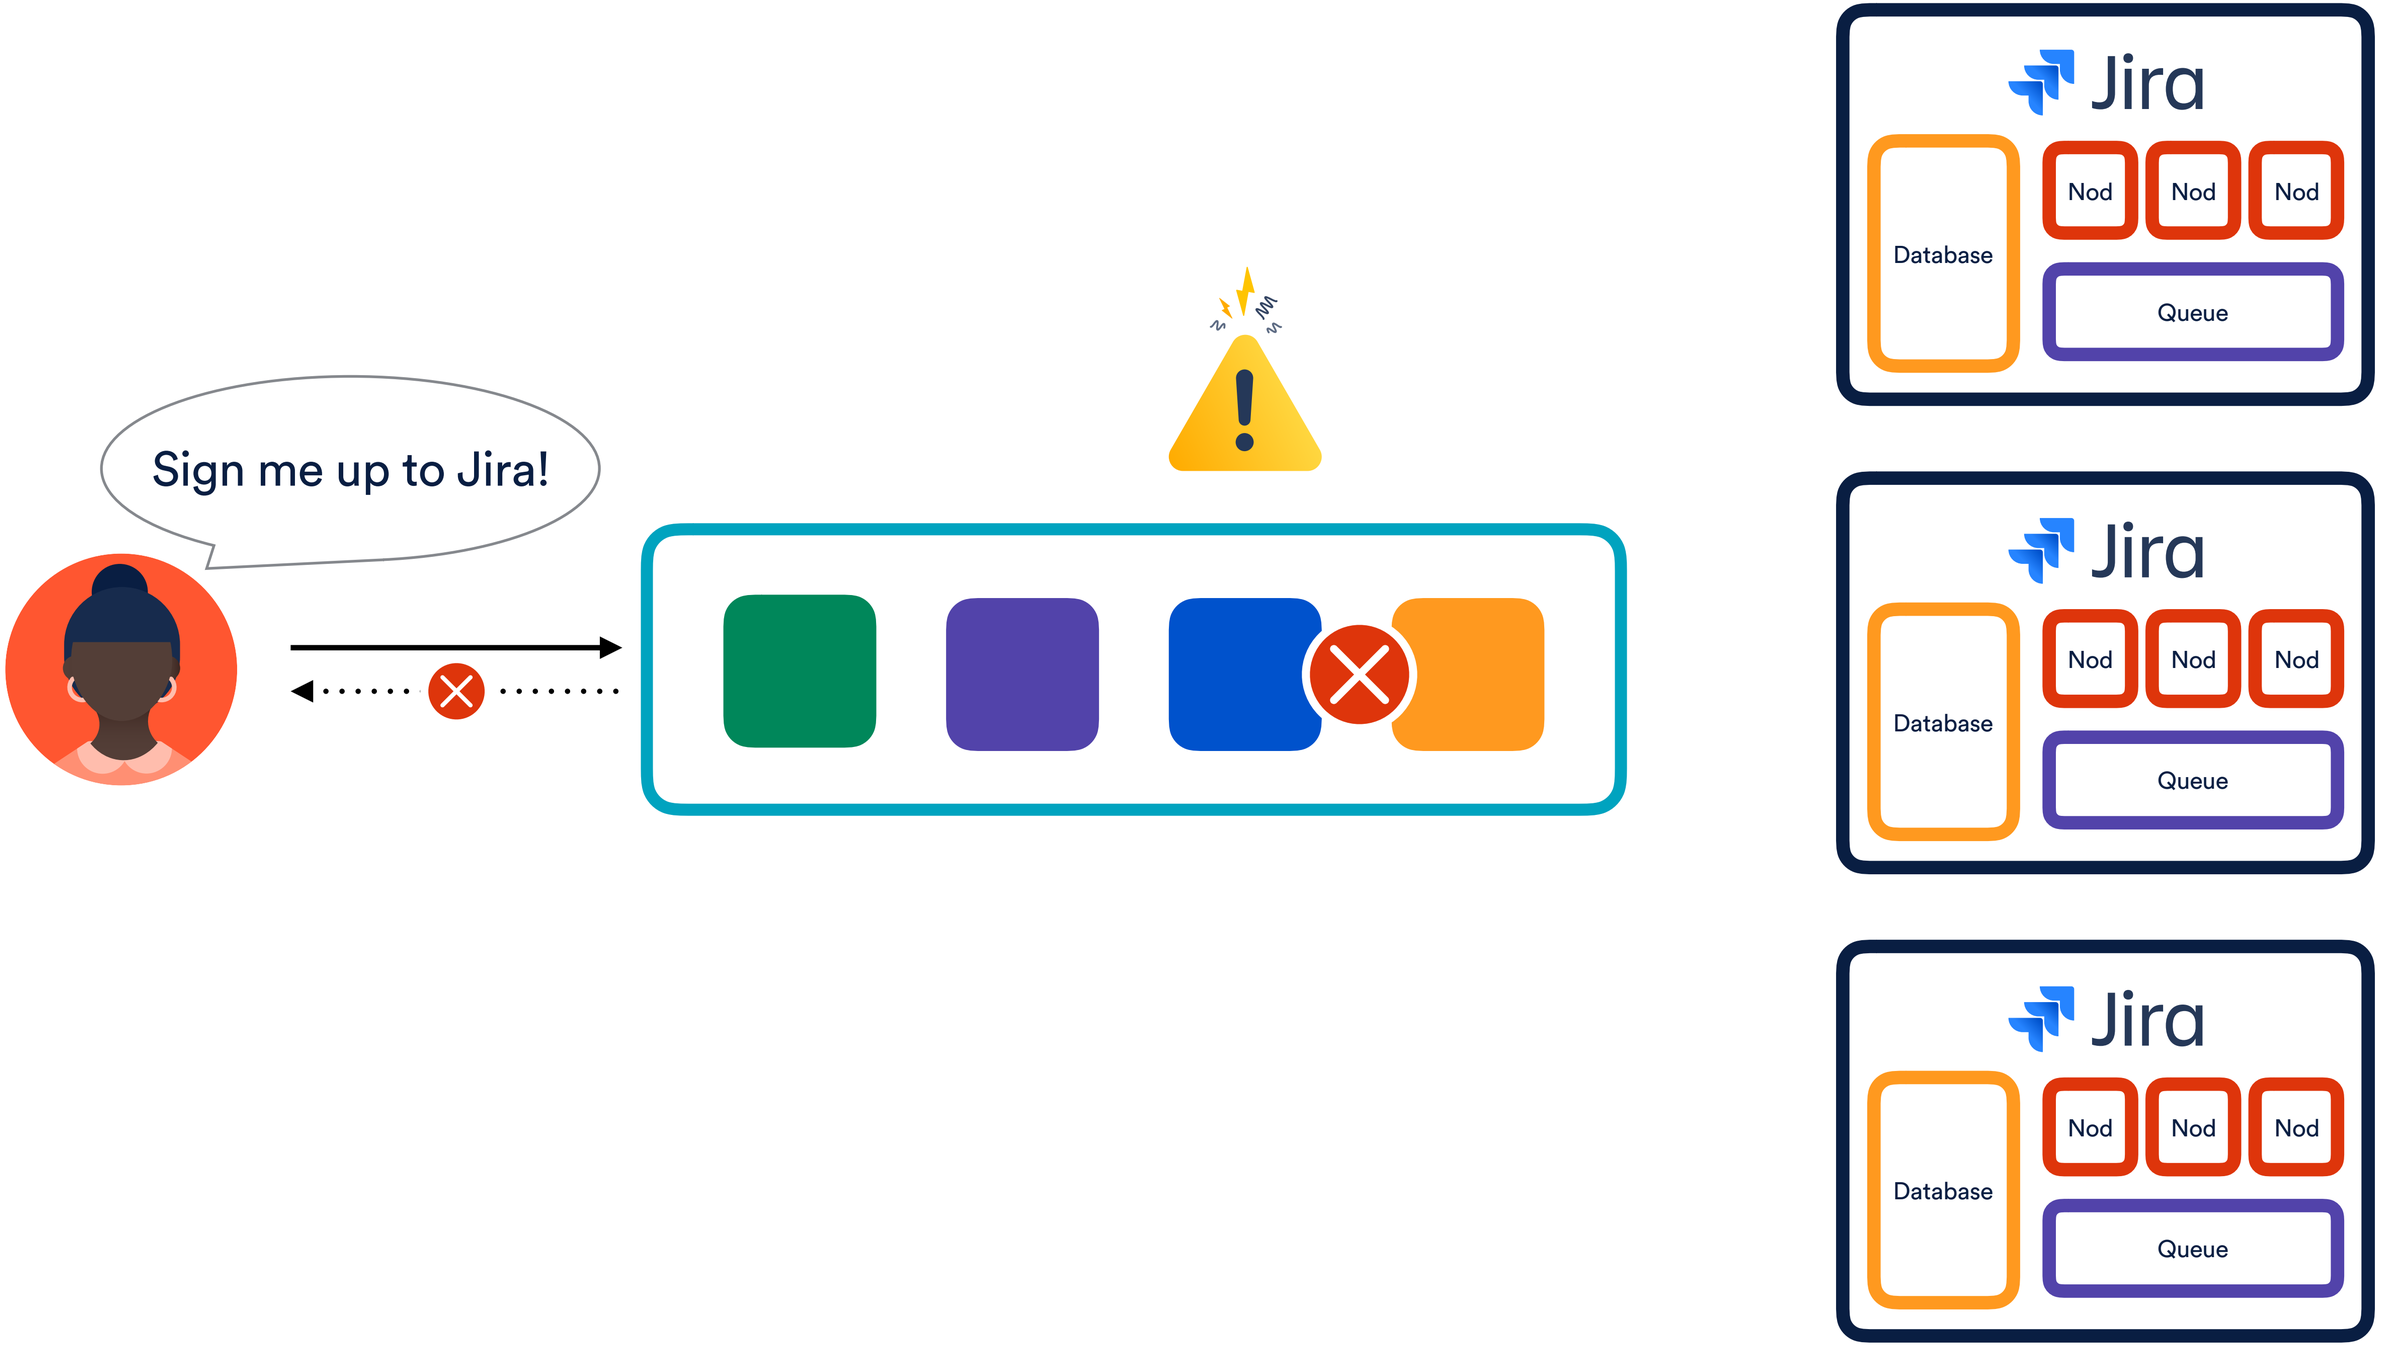
Sign me up to Jira!
Jira
Database
Nod
Nod
Nod
Queue
Jira
Database
Nod
Nod
Nod
Queue
Jira
Database
Nod
Nod
Nod
Queue

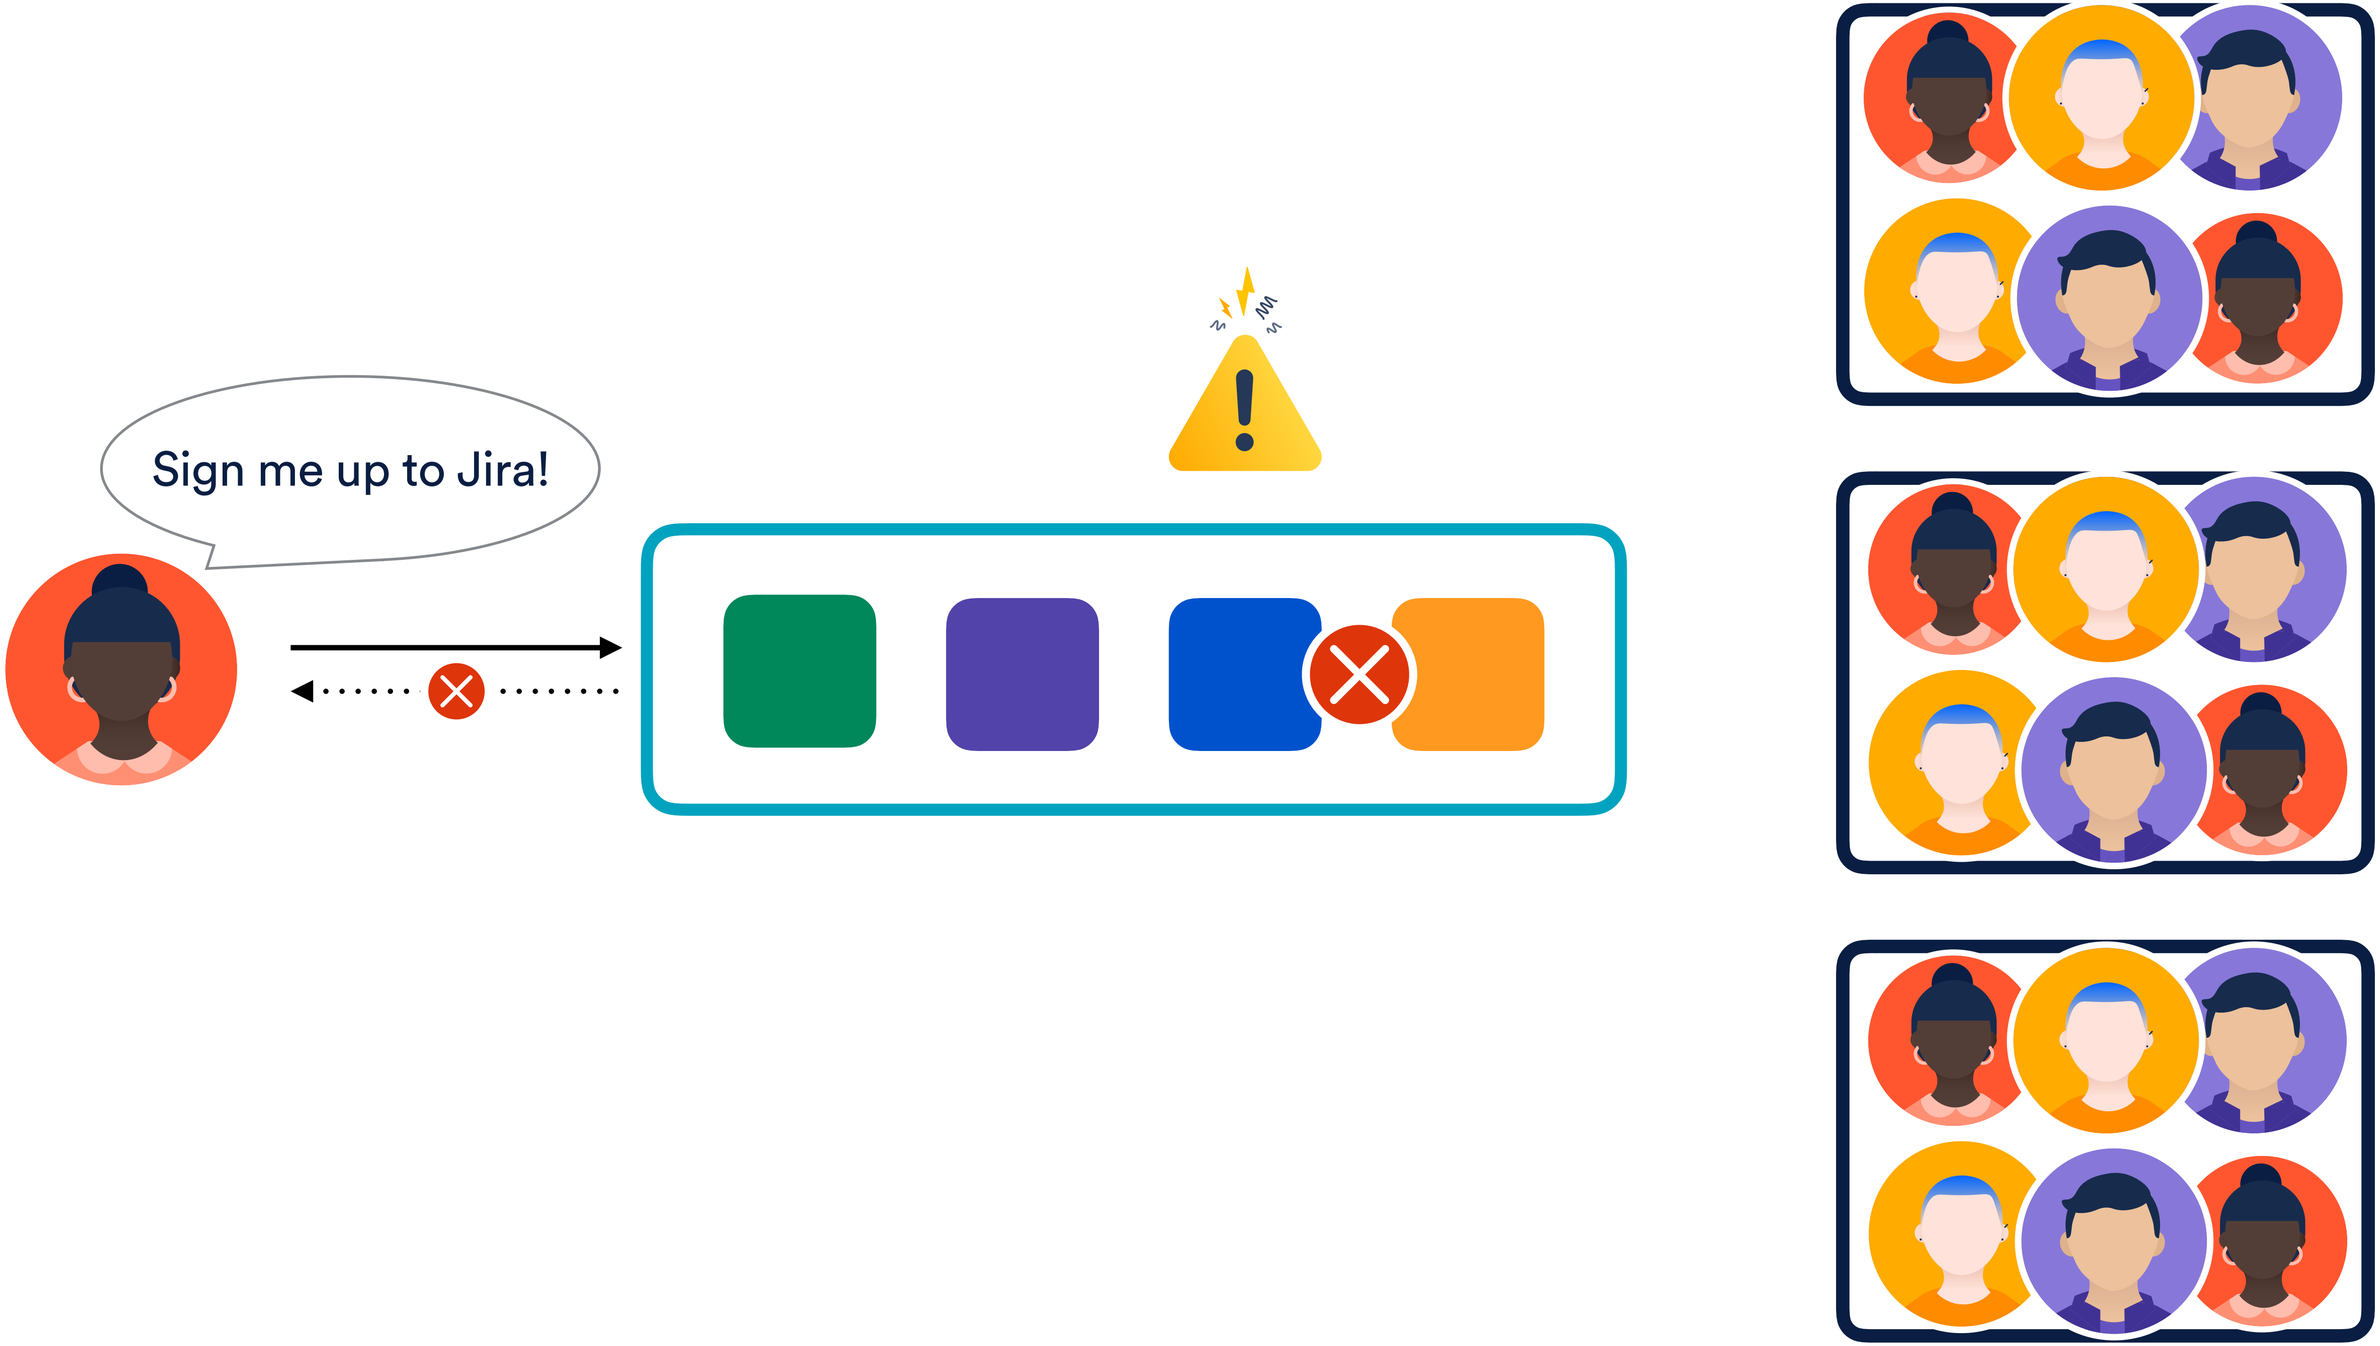
Sign me up to Jira!

TAKING A STEP BACK

What questions do we want to be able to answer with our operational resources?

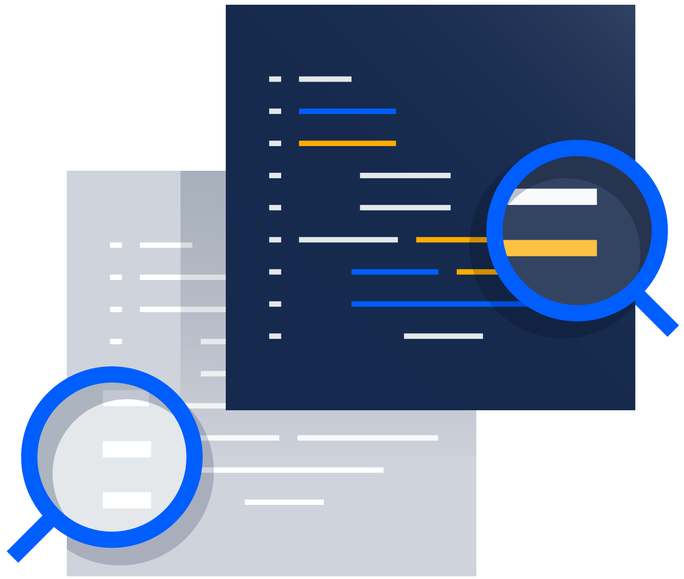

Shard Service

Performs the selection of a suitable shard based on geographical location and dynamic capacity metrics

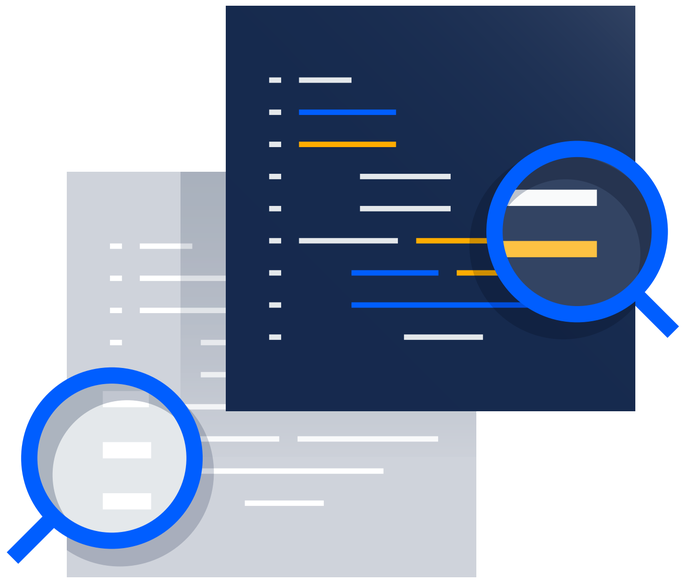

Synchronous, http-facing

Shard Service

Performs the selection of a suitable shard based on geographical location and dynamic capacity metrics

Shard Service

- Requests slow down significantly

Shard Service

- Requests slow down significantly
- Requests are accepted, but then fail

Shard Service

- Requests slow down significantly
- Requests are accepted, but then fail
- Requests start being rejected

Shard Service

- Requests slow down significantly
- Requests are accepted, but then fail
- Requests start being rejected
- There are no suitable shards

Shard Service

- Requests slow down significantly
- Requests are accepted, but then fail
- Requests start being rejected
- There are no suitable shards
- Incorrect shards were selected

Shard Service

- Requests slow down significantly
- Requests are accepted, but then fail
- Requests start being rejected
- There are no suitable shards
- Incorrect shards were selected
- There is insufficient data to make decisions

Shard Service

Infrastructure metrics

- Requests slow down significantly
- Requests are accepted, but then fail
- Requests start being rejected
- There are no suitable shards
- Incorrect shards were selected
- There is insufficient data to make decisions

Application metrics

Shard Service

Infrastructure metrics

- Requests slow down significantly
- Requests are accepted, but then fail
- Requests start being rejected
- There are no suitable shards
- Incorrect shards were selected
- There is insufficient data to make decisions

Application metrics

Infrastructure health

Useful metrics tied to components in your techstack.

Shard Service

Infrastructure metrics

- Requests slow down significantly
- Requests are accepted, but then fail
- Requests start being rejected
- There are no suitable shards
- Incorrect shards were selected
- There is insufficient data to make decisions

Application health

Useful metrics tied to the domain of your application

Application metrics

Infrastructure metrics + Application metrics = 

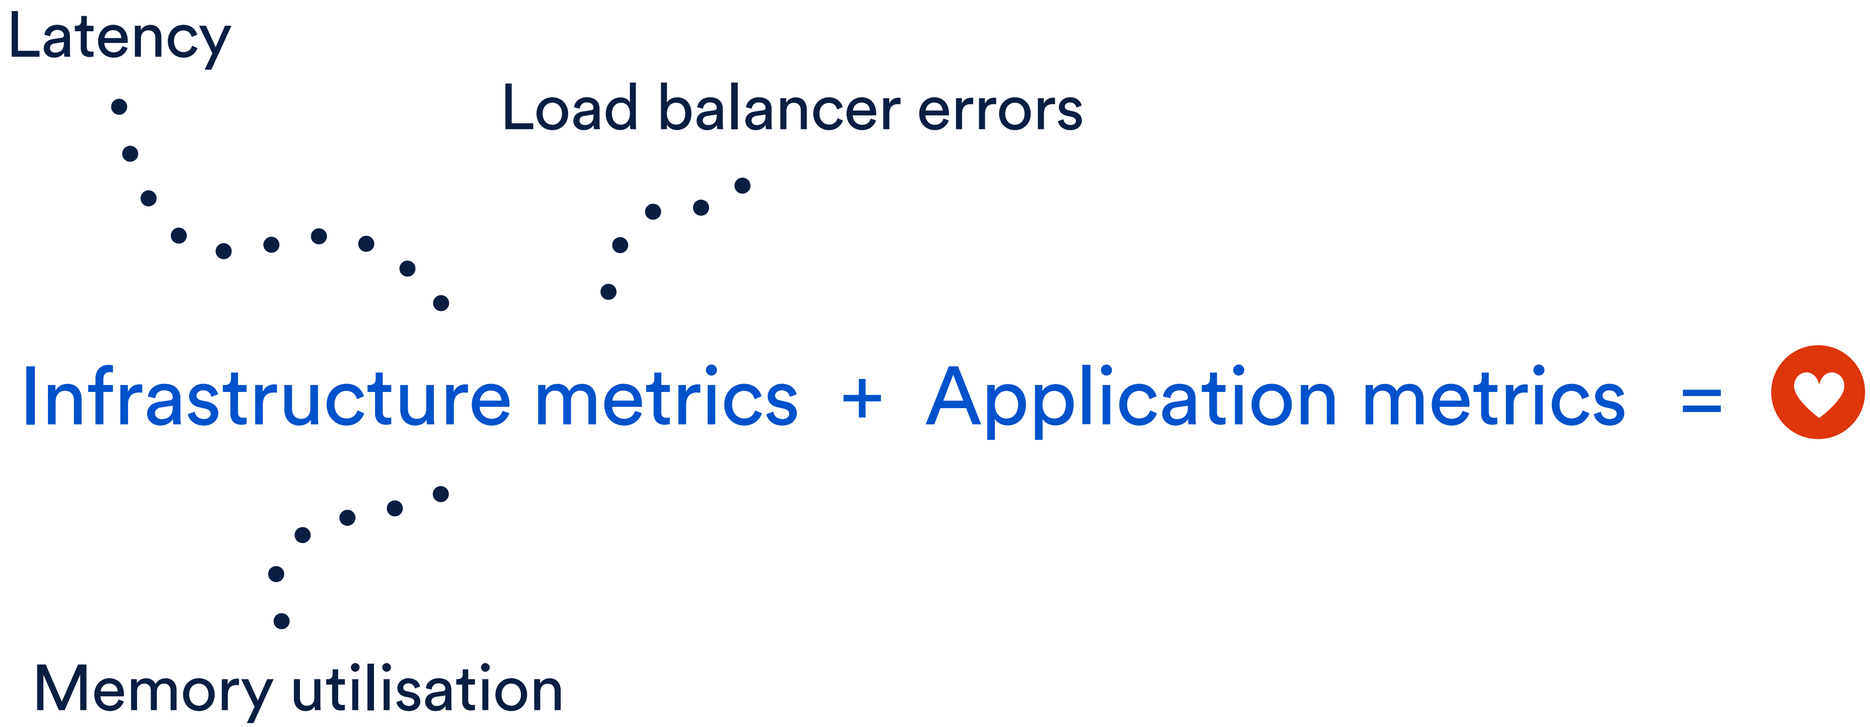
Latency
Load balancer errors
Infrastructure metrics + Application metrics =
Memory utilisation

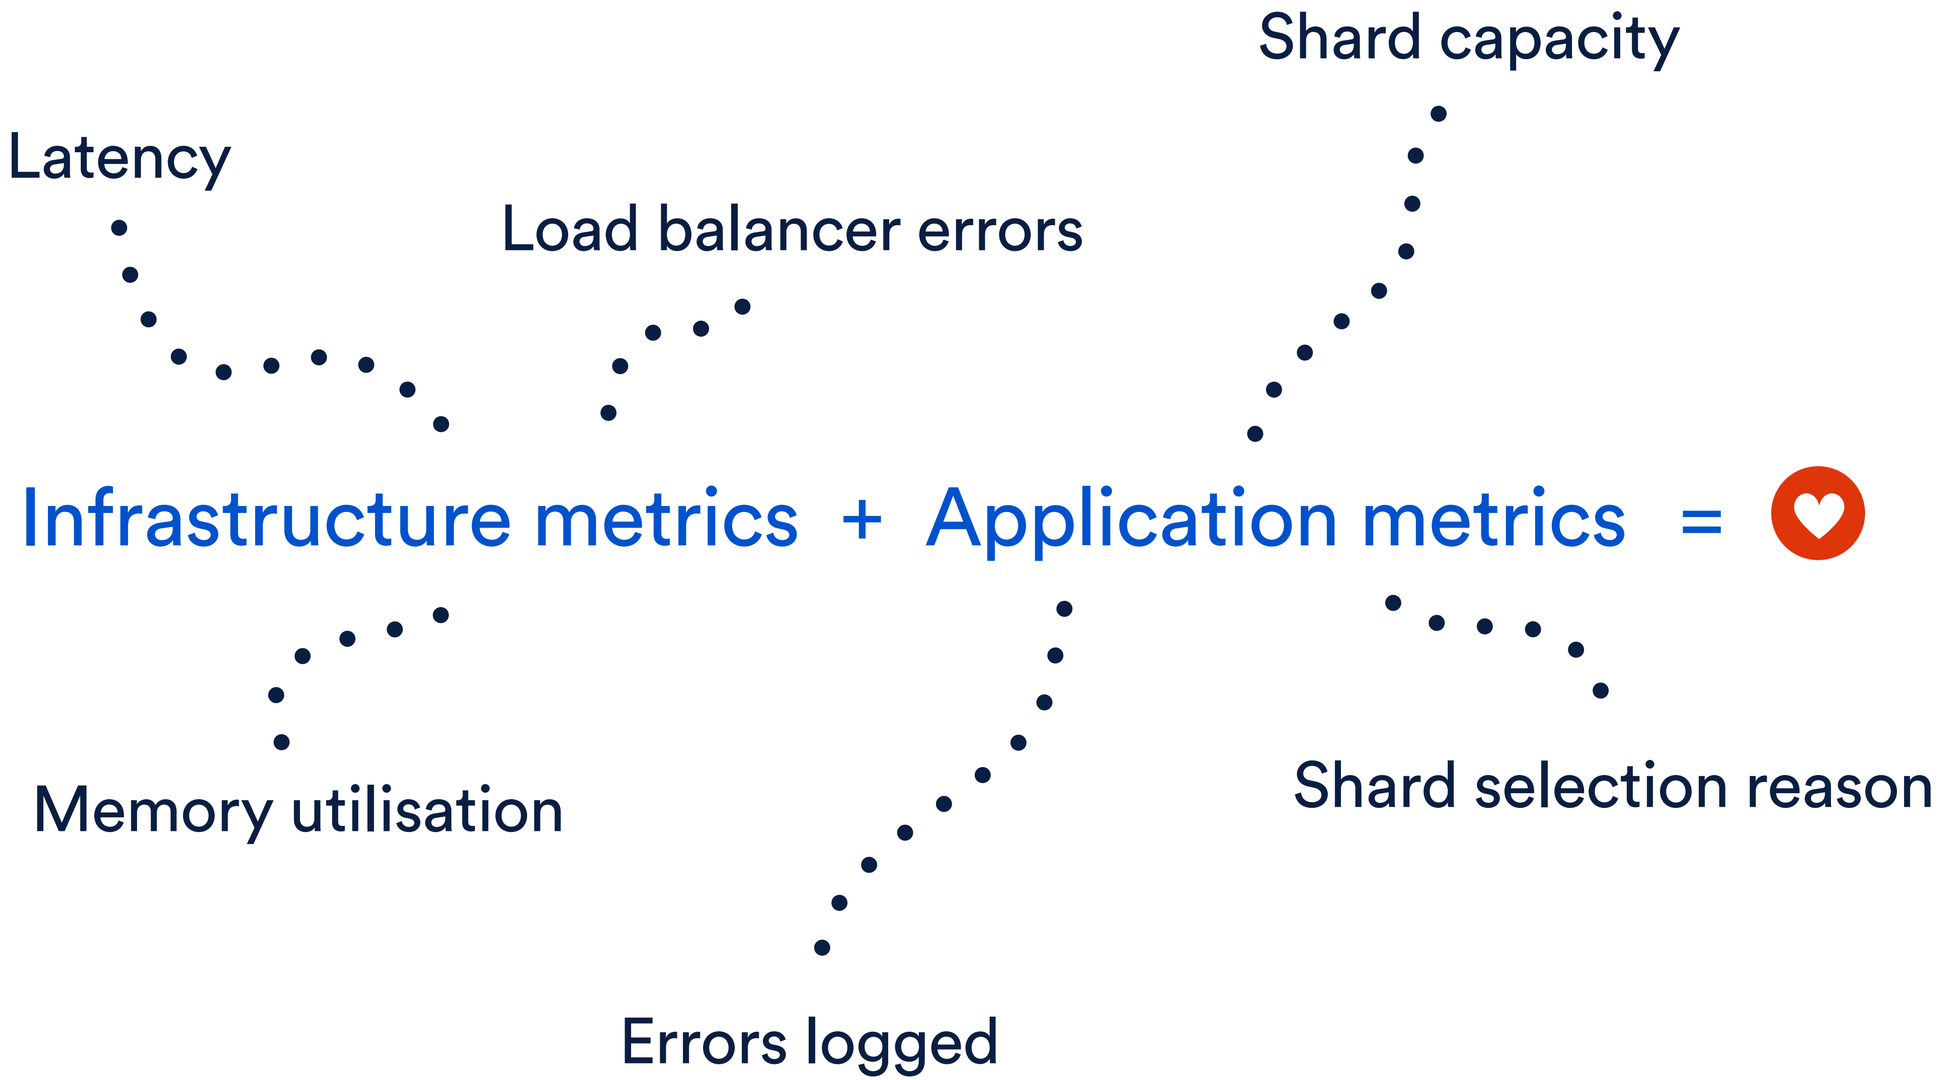
Shard capacity
Latency
Load balancer errors
Infrastructure metrics + Application metrics =
Memory utilisation
Shard selection reason
Errors logged

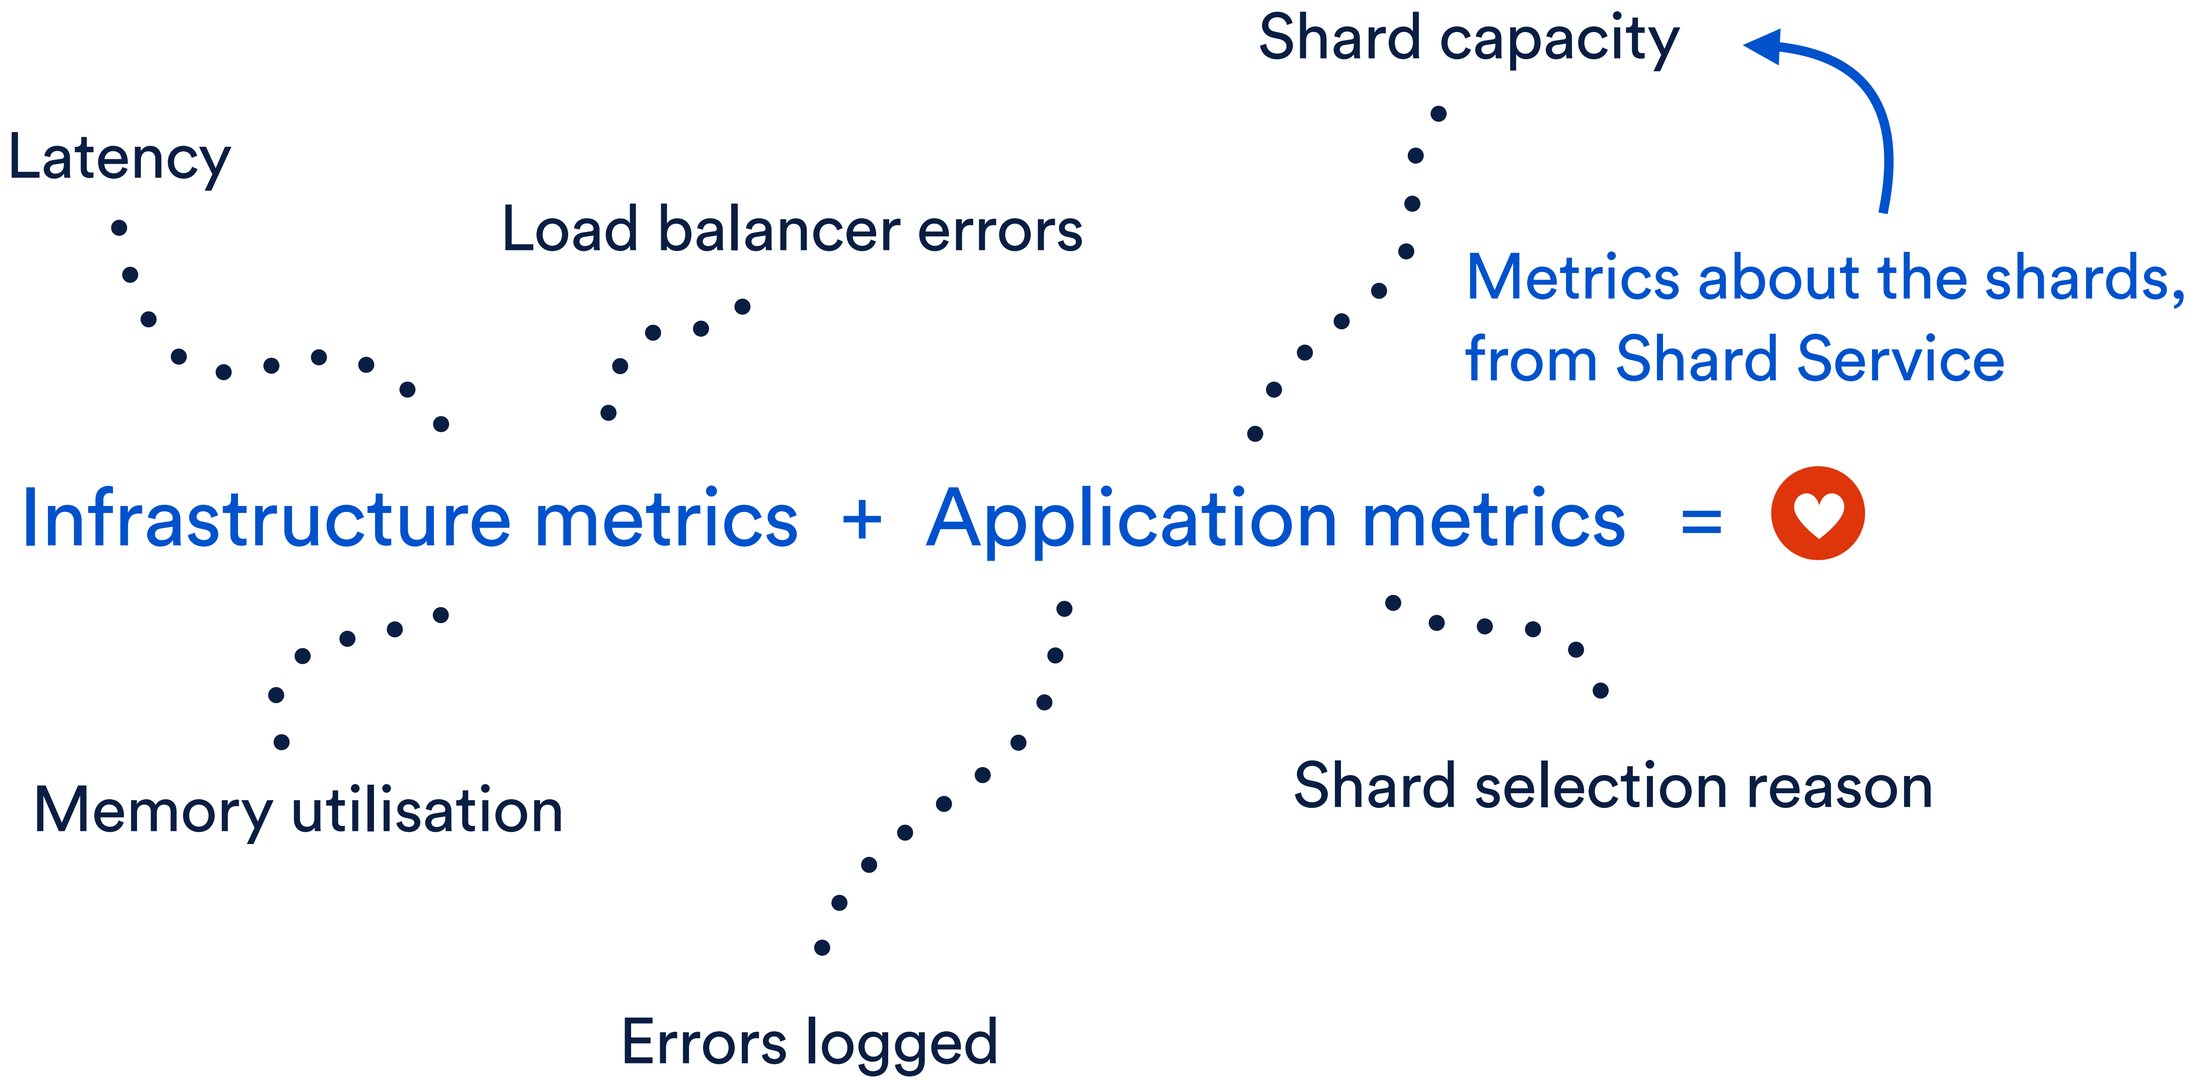
Shard capacity
Latency
Load balancer errors
Metrics about the shards,
from Shard Service
Infrastructure metrics + Application metrics =
Memory utilisation
Shard selection reason
Errors logged

Shard Service: capacity

Edit Board

$shard_status default_provisioning | $shard_purpose standard | $shard_type confluence | $shard_aws_region * | $shard_deployment_group production | $timezone_range * | $shard_name *

Important notice This dashboard in maintained in the code. Any changes made directly in UI should be reflected in the code or eventually they will be overwritten.

Effective Capacity vs. week-1

5m

11.98K	confluence-prod...	-4.37K
10.42K	confluence-prod...	-3.06K
10.02K	confluence-prod...	+2.22K
7.24K	confluence-prod...	-7.34K
7.18K	confluence-prod...	-83.43
2.92K	confluence-prod...	-4.82K
1.56K	confluence-prod...	-4.32K

Monarch Capacity vs. week-1

5m

7.24K	confluence-pr...	-7.34K
13.76K	confluence-pr...	-5.8K
2.92K	confluence-pr...	-4.82K
4.04K	confluence-pr...	-1.85K
13.27K	confluence-pr...	-587
7.18K	confluence-pr...	-83.4
10.02K	confluence-pr...	+2.22K

Change in Capacity vs. week-1

15.02K	confluence-pro...	-8.11K
1.56K	confluence-pro...	-7.2K
16.18K	confluence-pro...	-4.66K
10.22K	confluence-pro...	-4.41K
11.98K	confluence-pro...	-4.37K
10.42K	confluence-pro...	-3.06K
12.67K	confluence-pro...	-464.4

IR eu-west-1 capacity 5m
22401

FR eu-central-1 capac... 5m
1561

us-east-1 capacity 5m
2919

us-west-2 capacity 5m
7239

SYD ap-southeast-2 c... 5m
7180

SG ap-southeast-1 ca... 5m
10022

Shard Service: selection

Edit Board

1m The Past Month

Search...

$shard_type * | $shard_status * | $shard_name * | $shard_purpose * | $shard_aws_region * | $shard_deployment_group production | $shard_select_reason * | $shard_micros_region * | $timezone_range *

236062

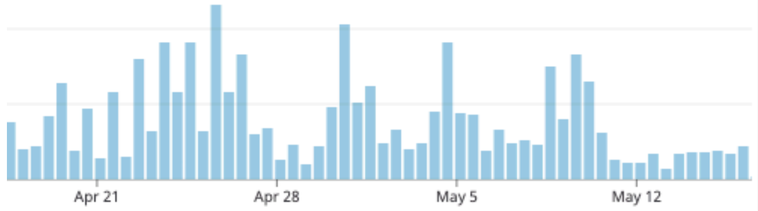

Top selected shards

- jira-prod-us-7
- jira-prod-eu-8
- confluence-prod-us-9
- jira-prod-eu-7
- jira-prod-us-19
- jira-prod-ap-2
- jira-prod-us-2
- jira-prod-us-11
- confluence-prod-us-8
- jira-prod-ap-1
- jira-prod-us-1
- confluence-prod-us-3

Top selection reasons

- specified_shard_in_request
- in_offset_and_available
- has_country_and_available
- cp_state
- in_region_and_available
- specified_shard_in_cloud_name
- cloud_name_security

Top selected region (AWS)

- us-east-1
- us-west-2
- eu-west-1
- eu-central-1
- ap-southeast-1
- ap-southeast-2

Top selected region (internal)

- prod_east
- prod_west2
- prod_euwest
- prod_eucentral
- prod_apse
- prod_apse2

Database utilisation metric: current vs target (by shard)

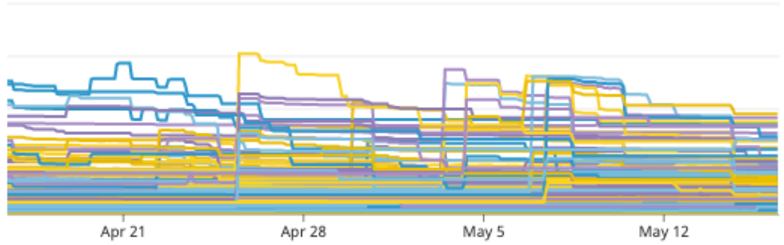

Jira Apdex Count metric: current vs target (by shard)

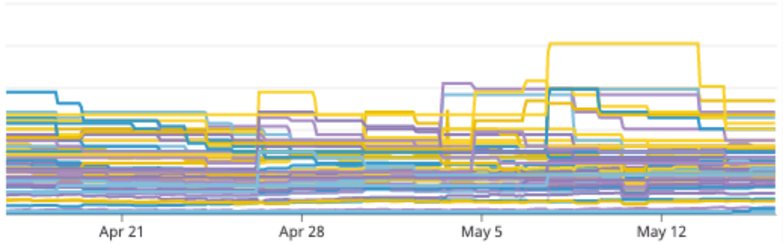

Confluence Apdex Count metric: current vs target (by shard)

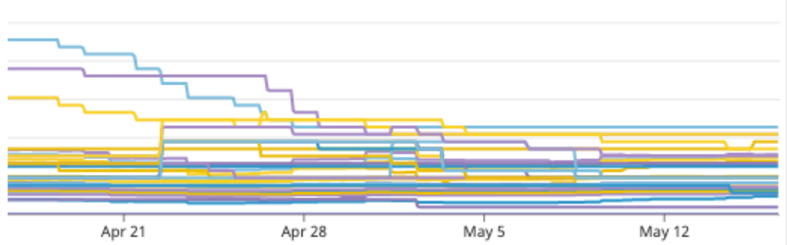

Provisioning failures metric: current vs target (by shard)

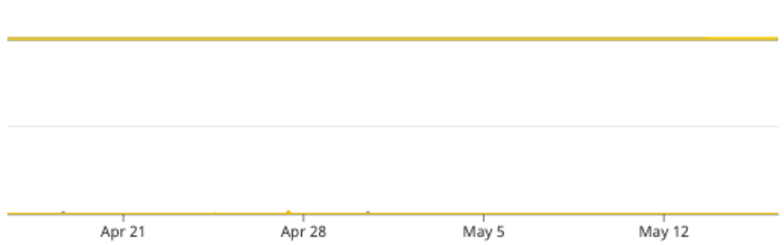

Average remaining capacity by shard (Jira Apdex)

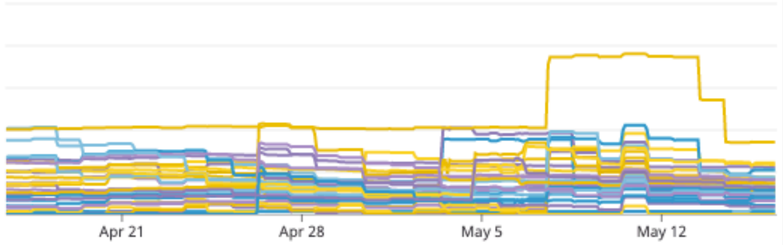

Average remaining capacity by shard (Confluence Apdex)

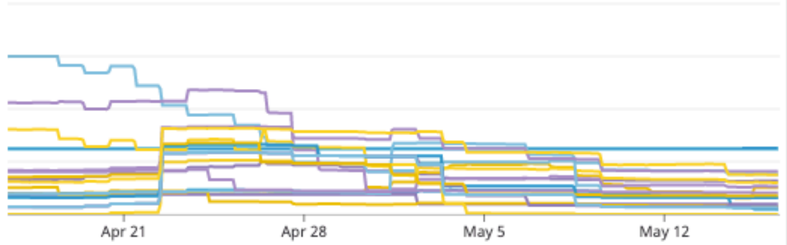

SHARD SERVICE INFRASTRUCTURE

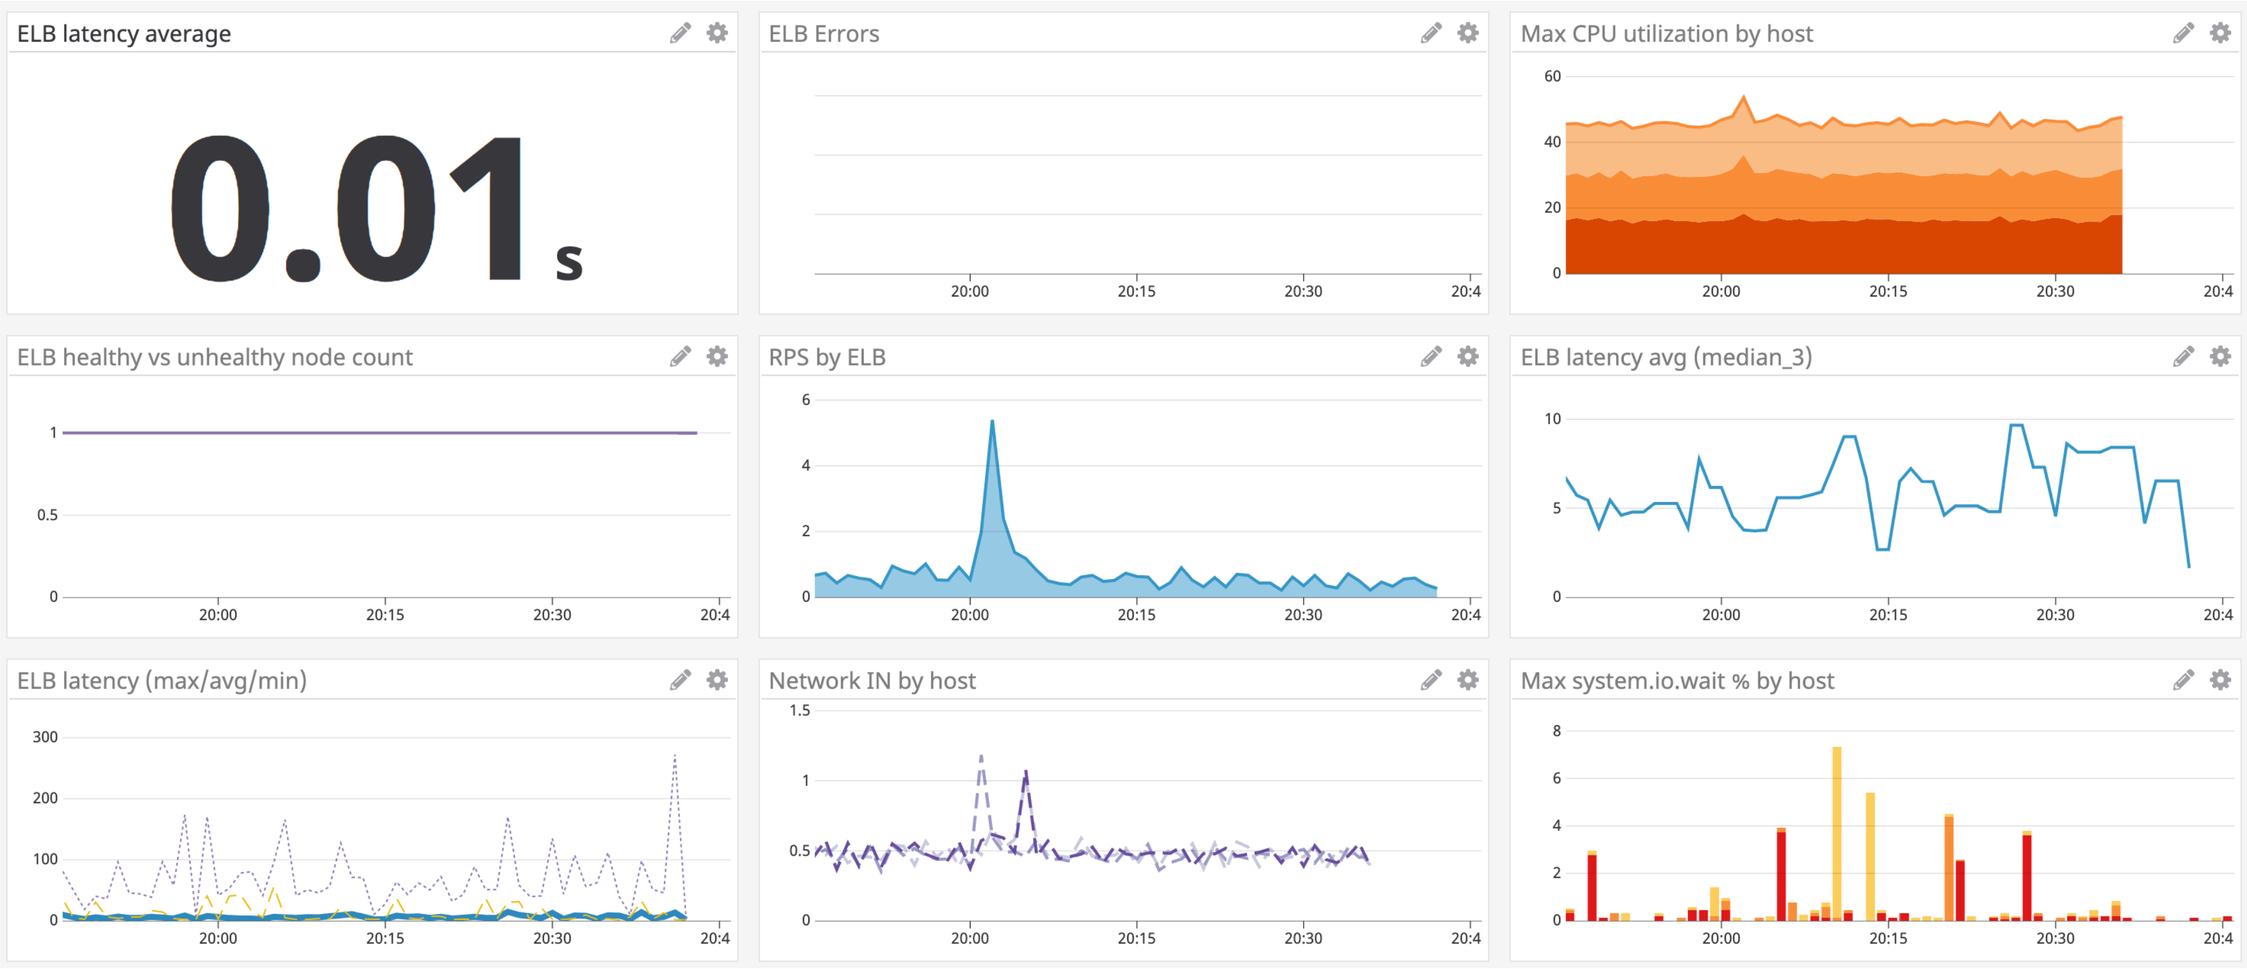

Monitors

Shard capacity exhausted

Region capacity exhausted

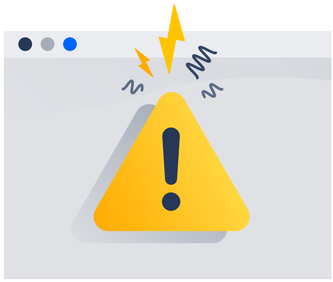

Surge in errors logged

How can you…

How can you...
Figure out what to measure?

What questions do you want to answer?

What questions do you want to answer?

Why

does your service exist (what are its roles and responsibilities)?

What

does it look like for those roles and responsibilities to degrade?

How

can you verify whether or not such a degradation is occurring?

Agenda

Iterative... what?

Setting some context

Deciding what to measure

Verifying your metrics

Keeping up with change

Summary

Agenda

Iterative... what?

Setting some context

Deciding what to measure

Verifying your metrics

Keeping up with change

Summary

Everything is right.

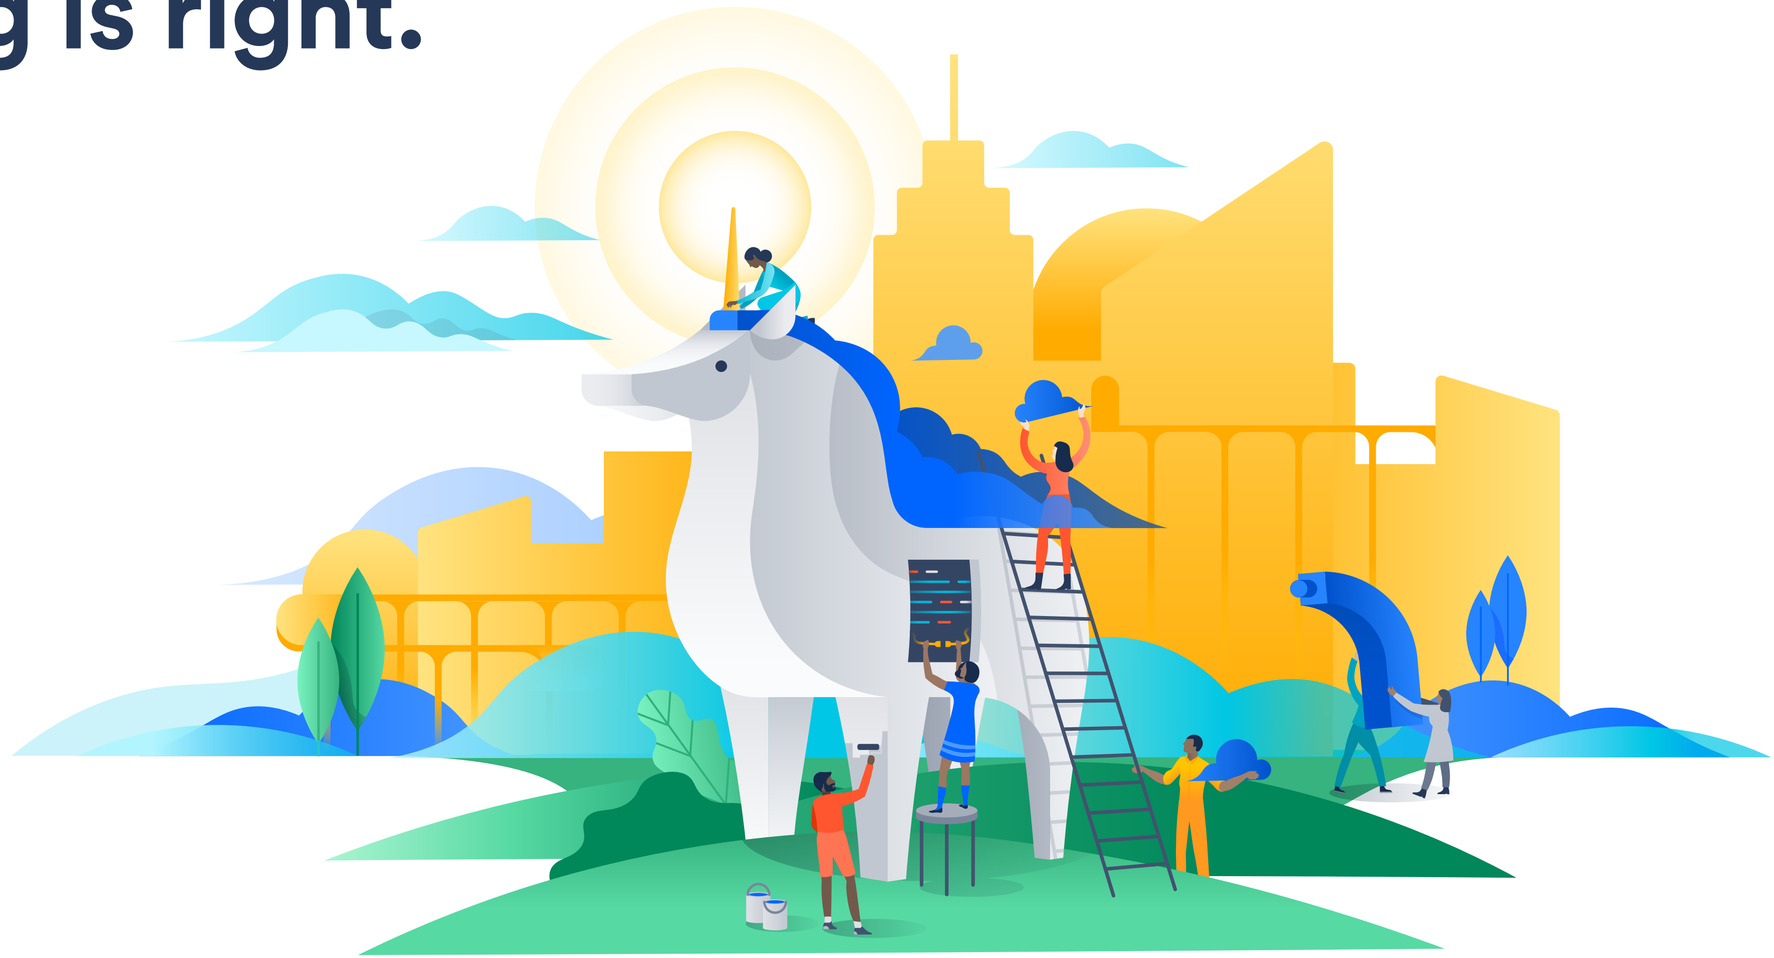

As time went on...

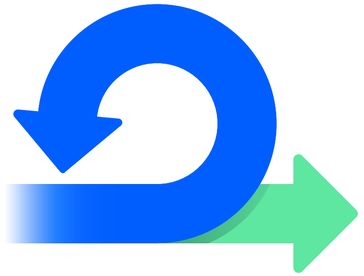

Things changed

Because, you know, agile

Noisy alerts

Frequent & un-actionable

Fine... or exploding

Never checked operational health unless it was on fire

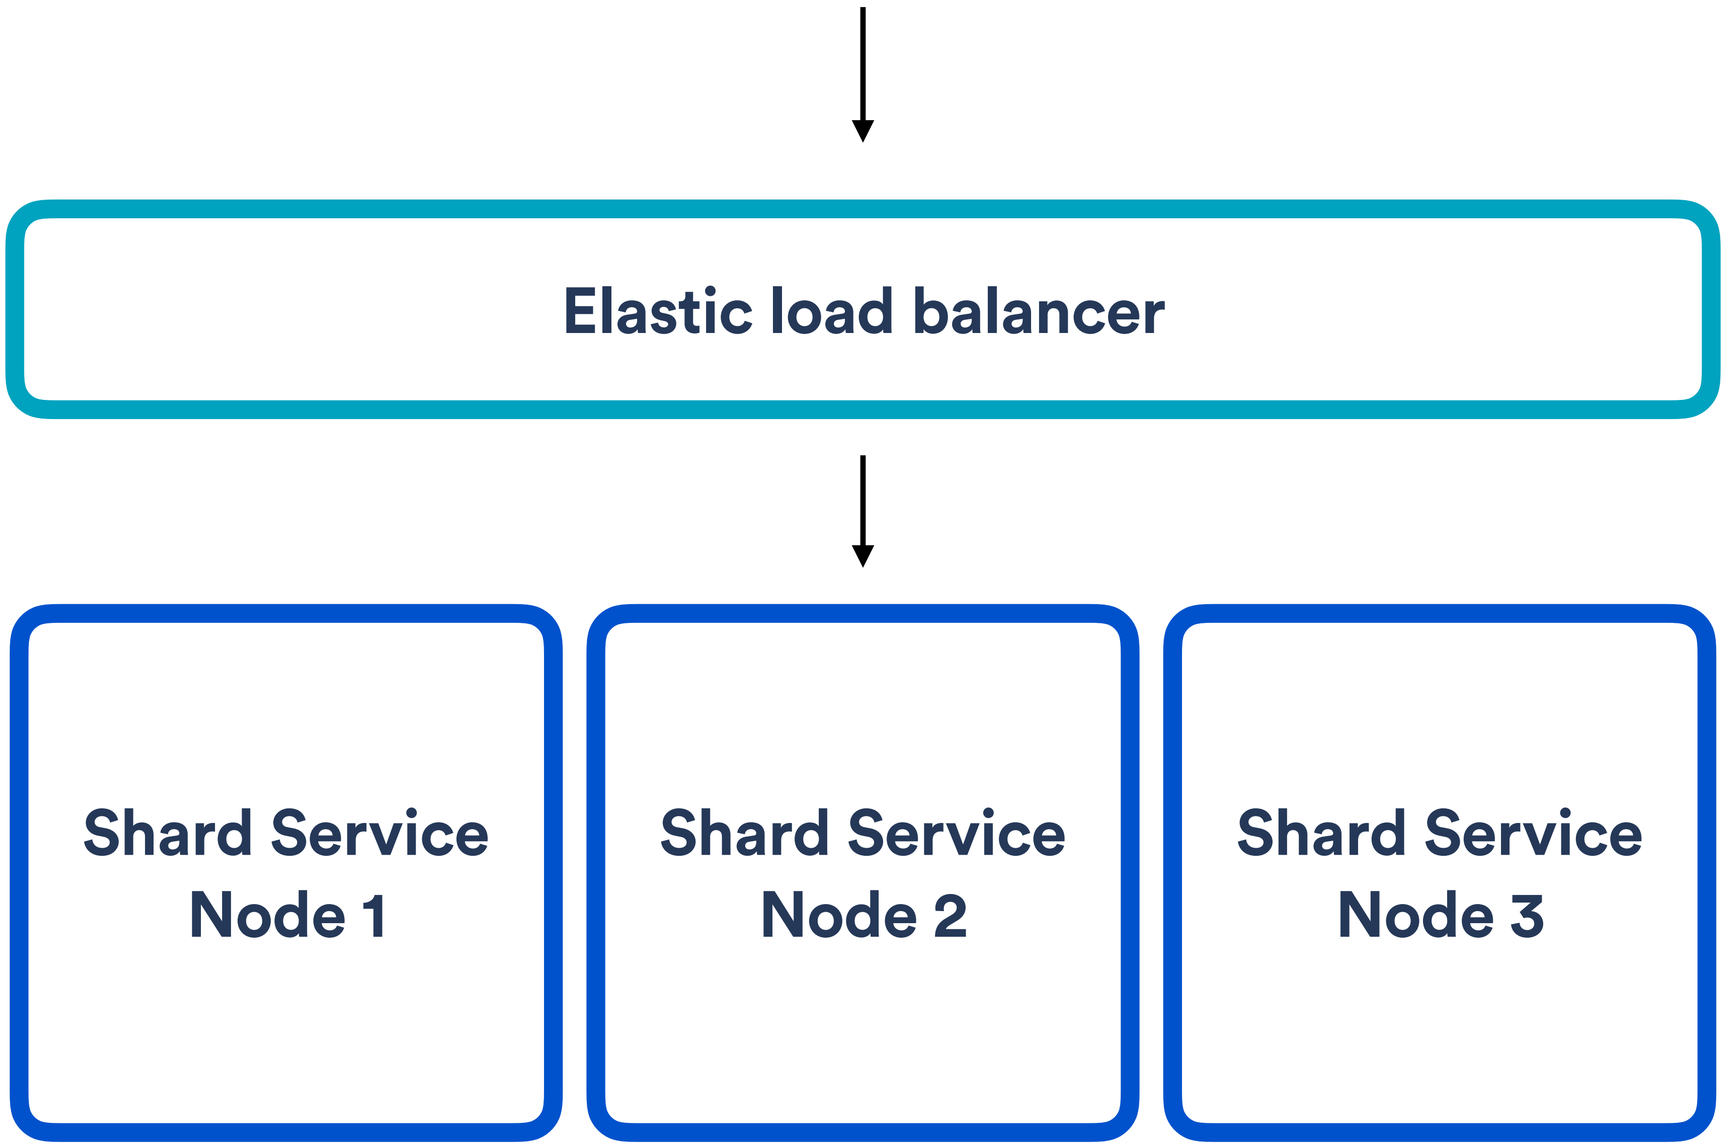
Elastic load balancer
Shard Service Node 1
Shard Service Node 2
Shard Service Node 3

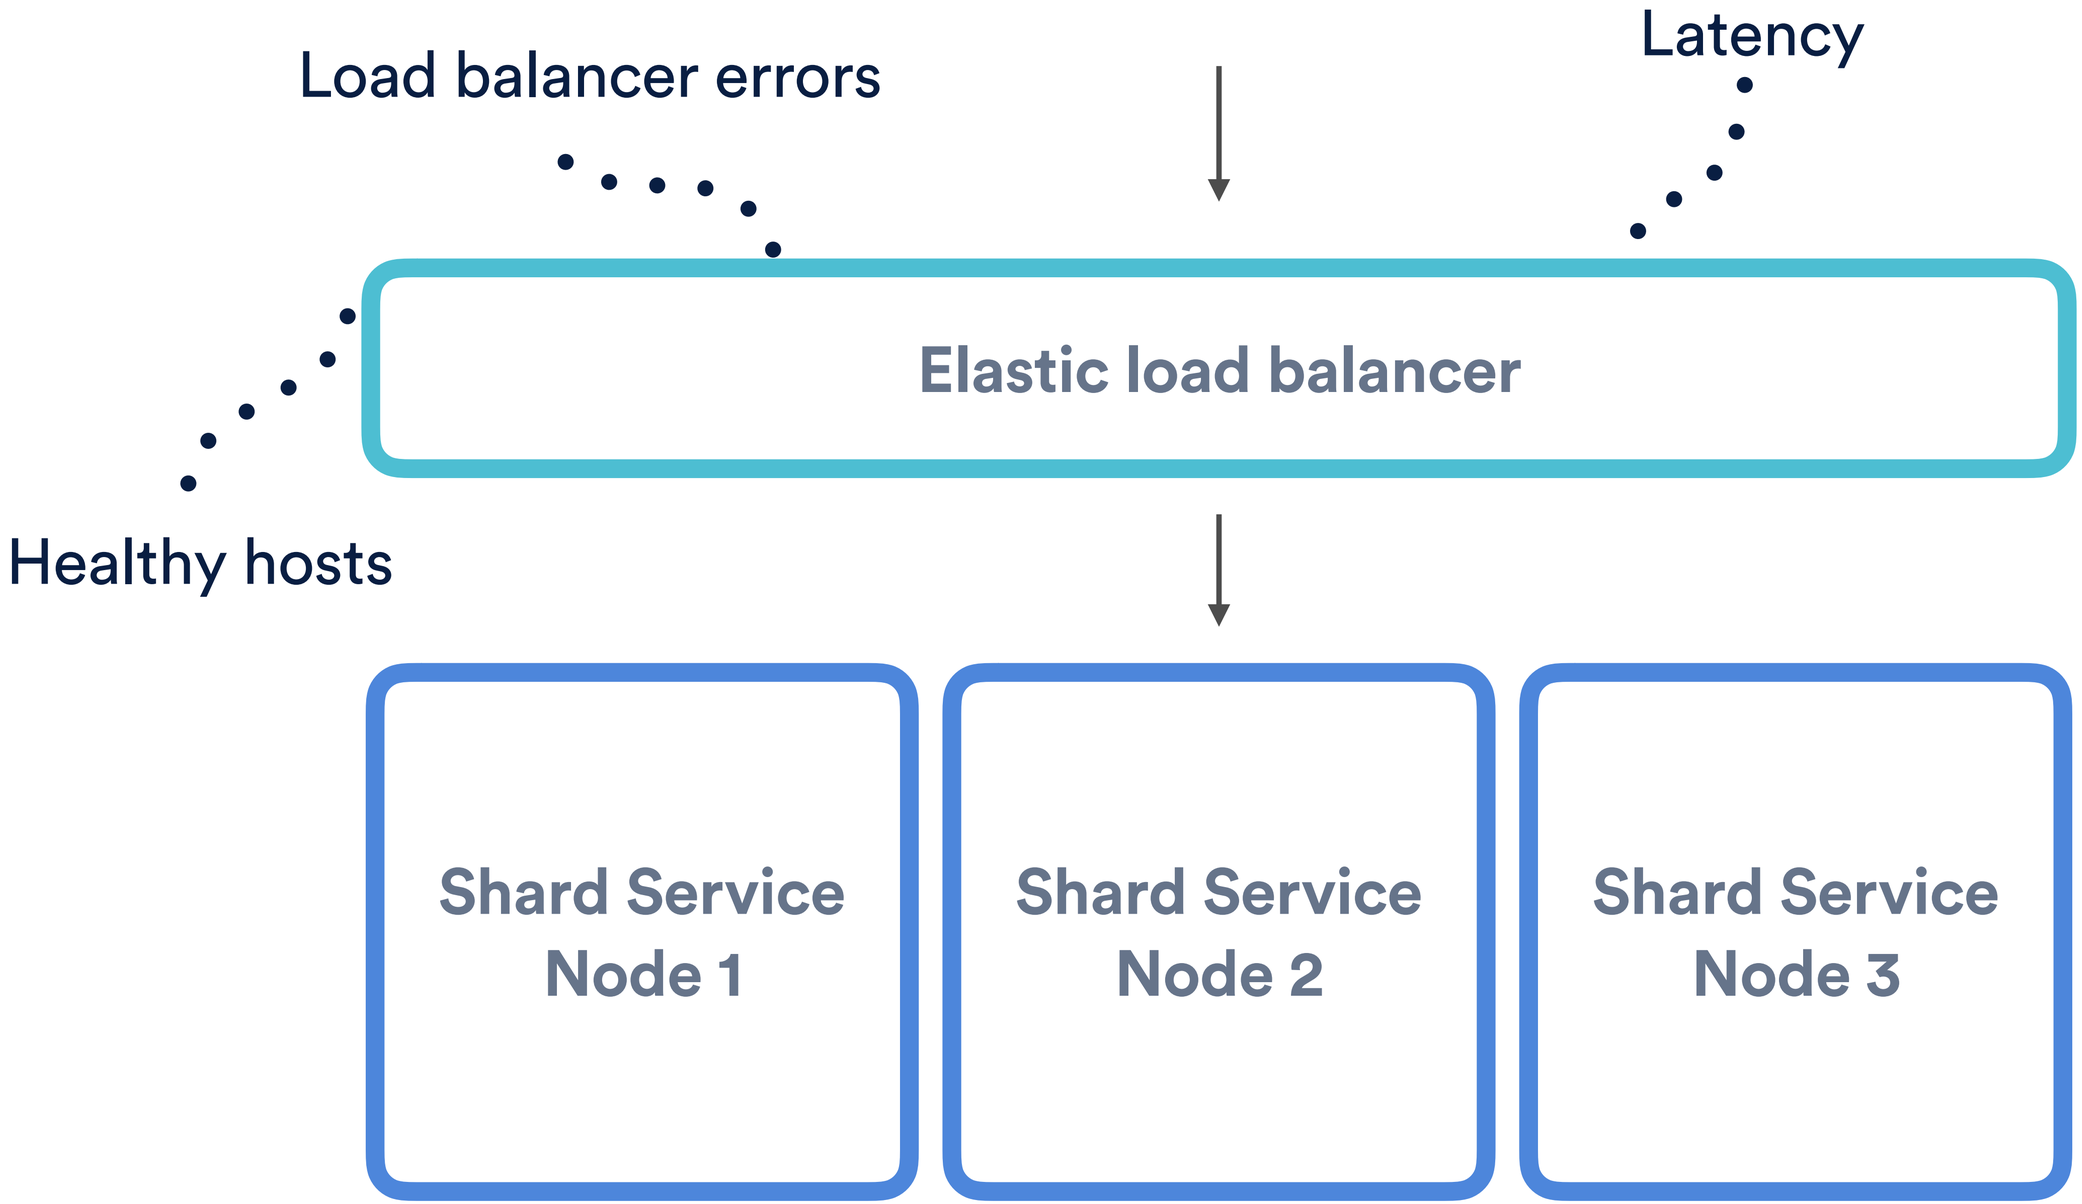
Load balancer errors
Latency
Elastic load balancer
Healthy hosts
Shard Service Node 1
Shard Service Node 2
Shard Service Node 3

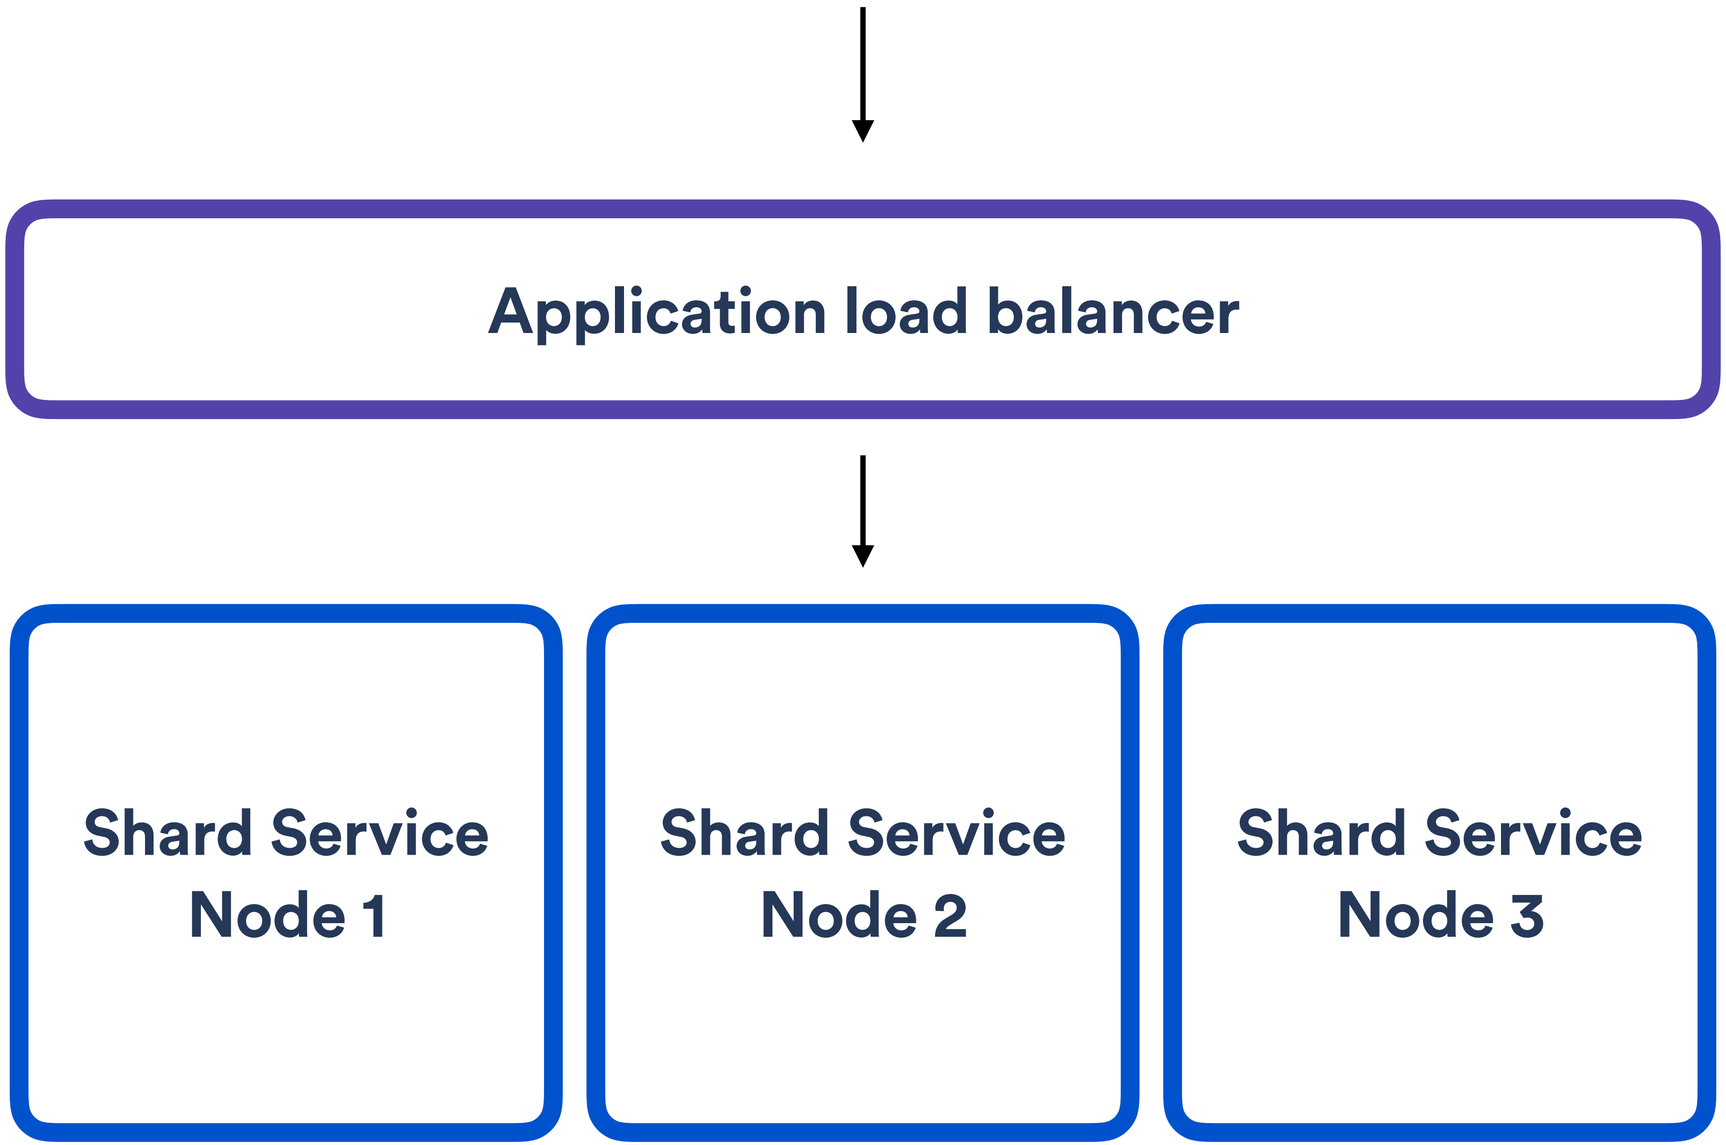
Application load balancer
Shard Service Node 1
Shard Service Node 2
Shard Service Node 3

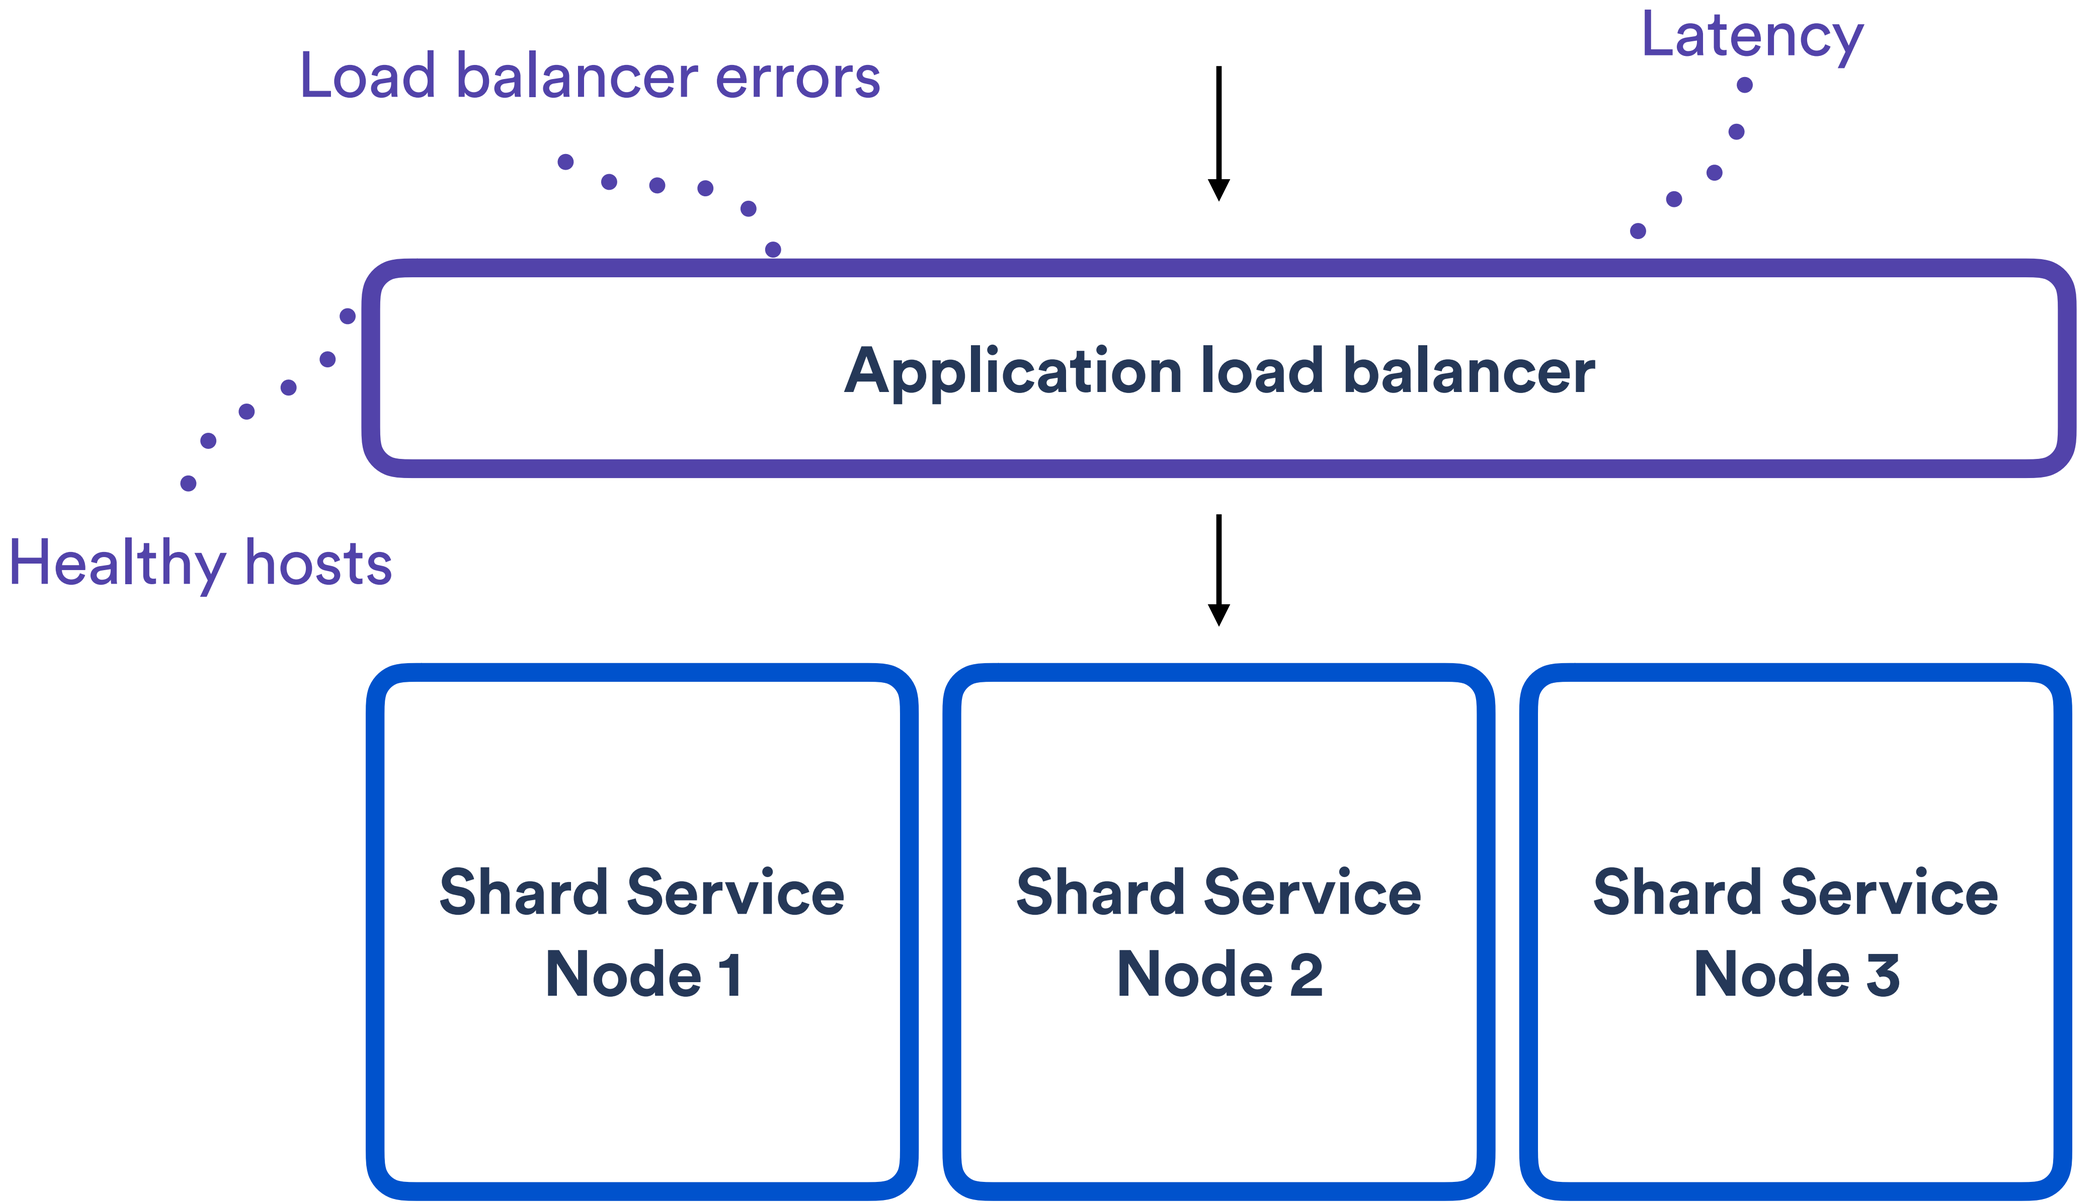
Load balancer errors
Latency
Application load balancer
Healthy hosts
Shard Service Node 1
Shard Service Node 2
Shard Service Node 3

As time went on...

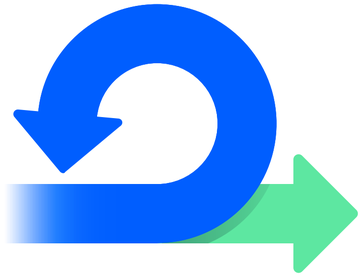

Things changed

Because, you know, agile

Noisy alerts

Frequent & un-actionable

Fine... or exploding

Never checked operational health unless it was on fire

As time went on...

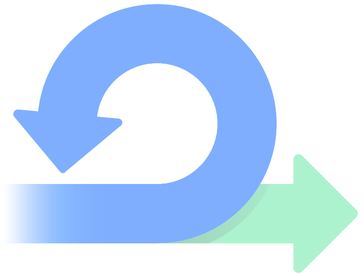

Things changed

Because, you know, agile

Noisy alerts

Frequent & un-actionable

Fine... or exploding

Never checked operational health unless it was on fire

Our team

Team who could actually fix the problem

As time went on...

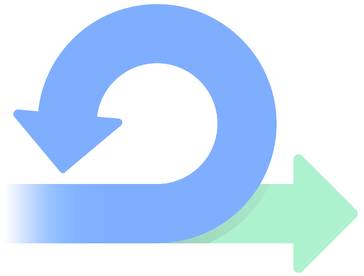

Things changed

Because, you know, agile

Noisy alerts

Frequent & un-actionable

Fine... or exploding

Never checked operational health unless it was on fire

As time went on...

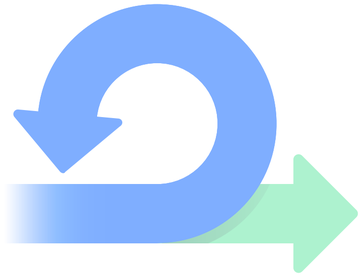

Things changed

Because, you know, agile

Noisy alerts

Frequent & un-actionable

Fine... or exploding

Never checked operational health unless it was on fire

SERVICE LEVEL OBJECTIVE

What level of service you can commit to offer

SERVICE LEVEL OBJECTIVE

What level of service you can commit to offer

E.g. 99.99% requests should succeed

We were not alone

TECHOPS

Process dedicated to regularly reviewing, discussing and iterating on operational health

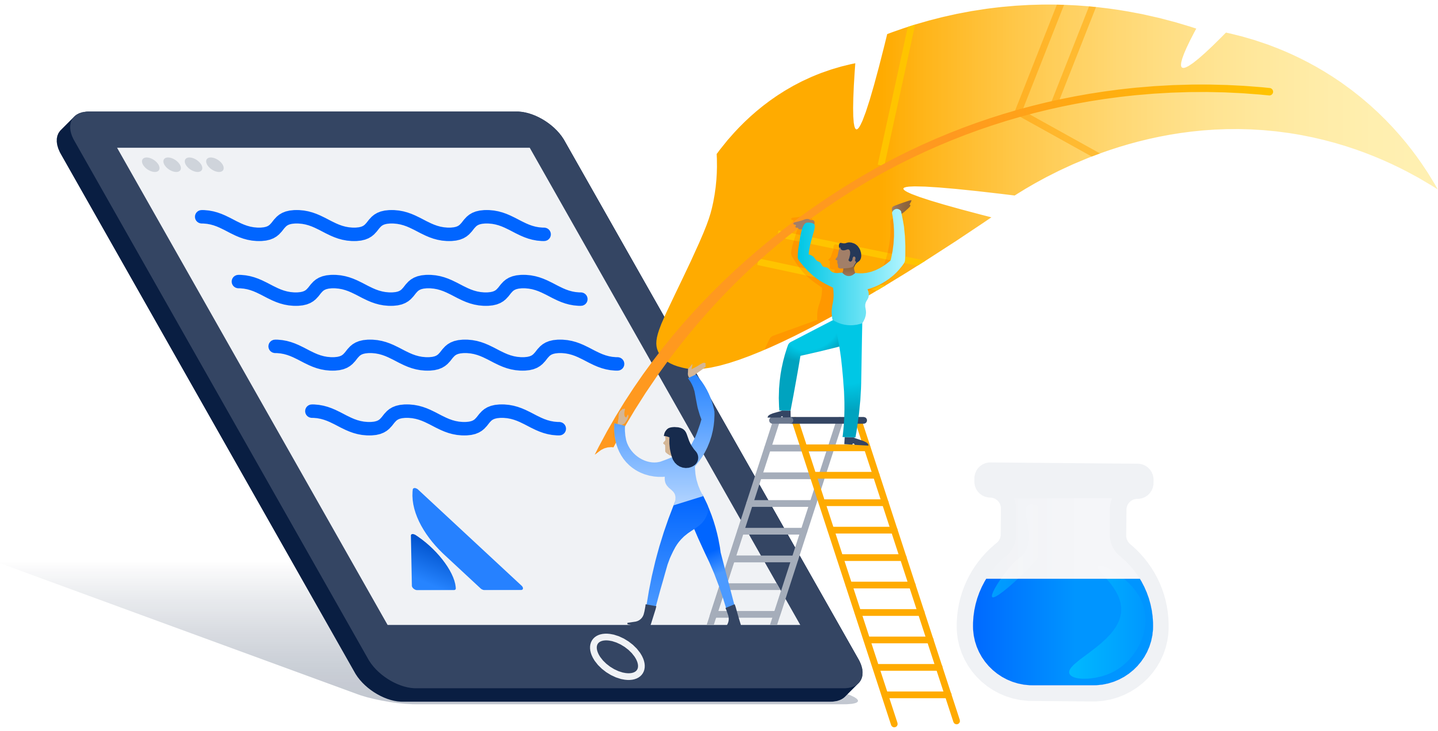

TechOps

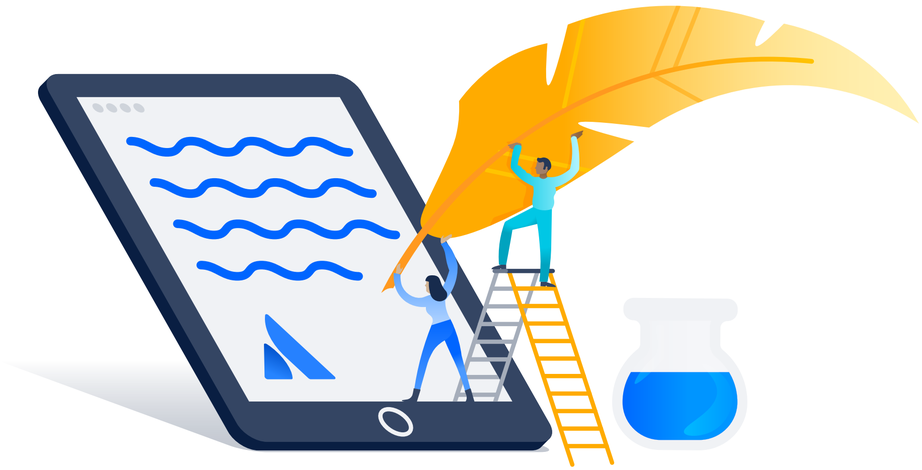

Develop measurable goals

Collect data

Prepare a report

Meet and discuss

Repeat and iterate

TechOps

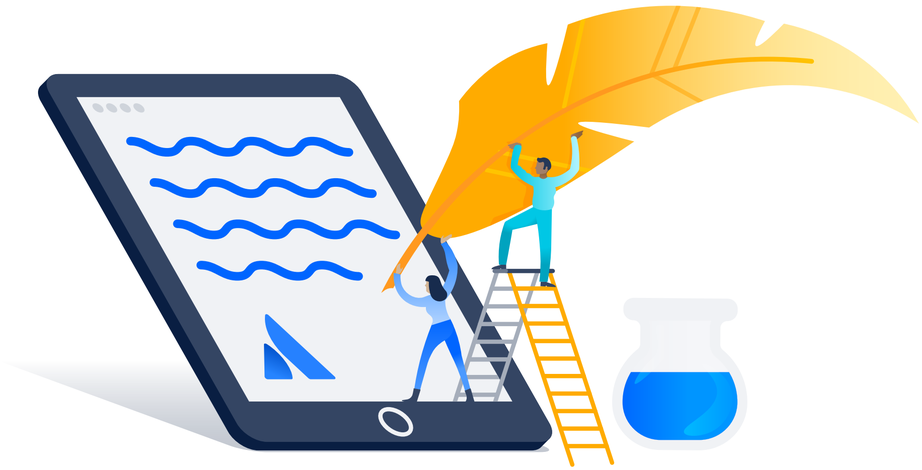

Develop measurable goals

Collect data

Prepare a report

Meet and discuss

Repeat and iterate

TechOps

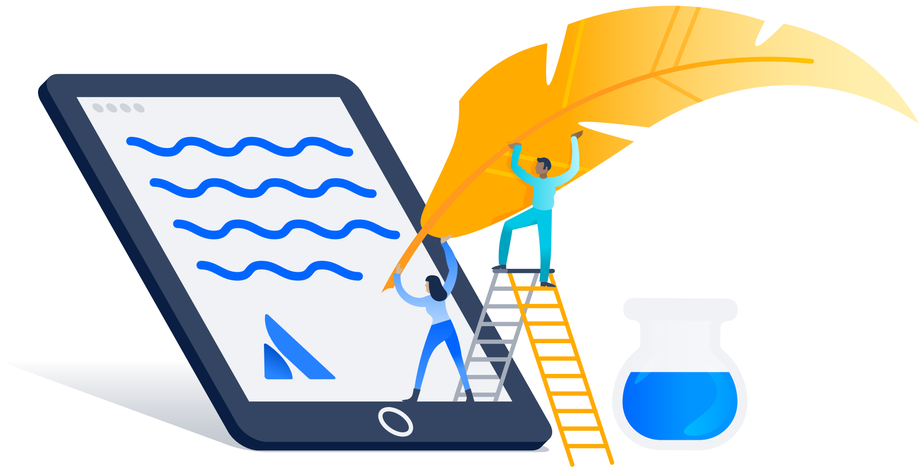

Develop measurable goals

Collect data

Prepare a report

Meet and discuss

Repeat and iterate

TechOps

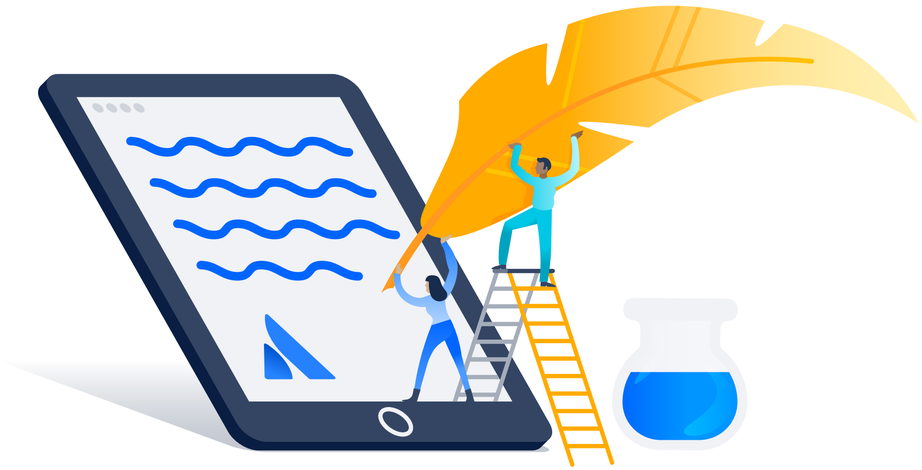

Develop measurable goals

Collect data

Prepare a report

Meet and discuss

Repeat and iterate

TechOps

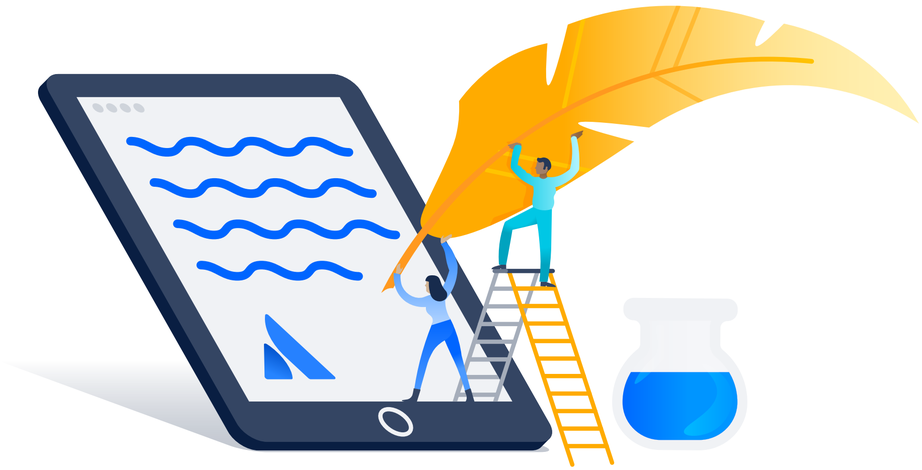

Develop measurable goals

Collect data

Prepare a report

Meet and discuss

Repeat and iterate

TechOps for everyone!

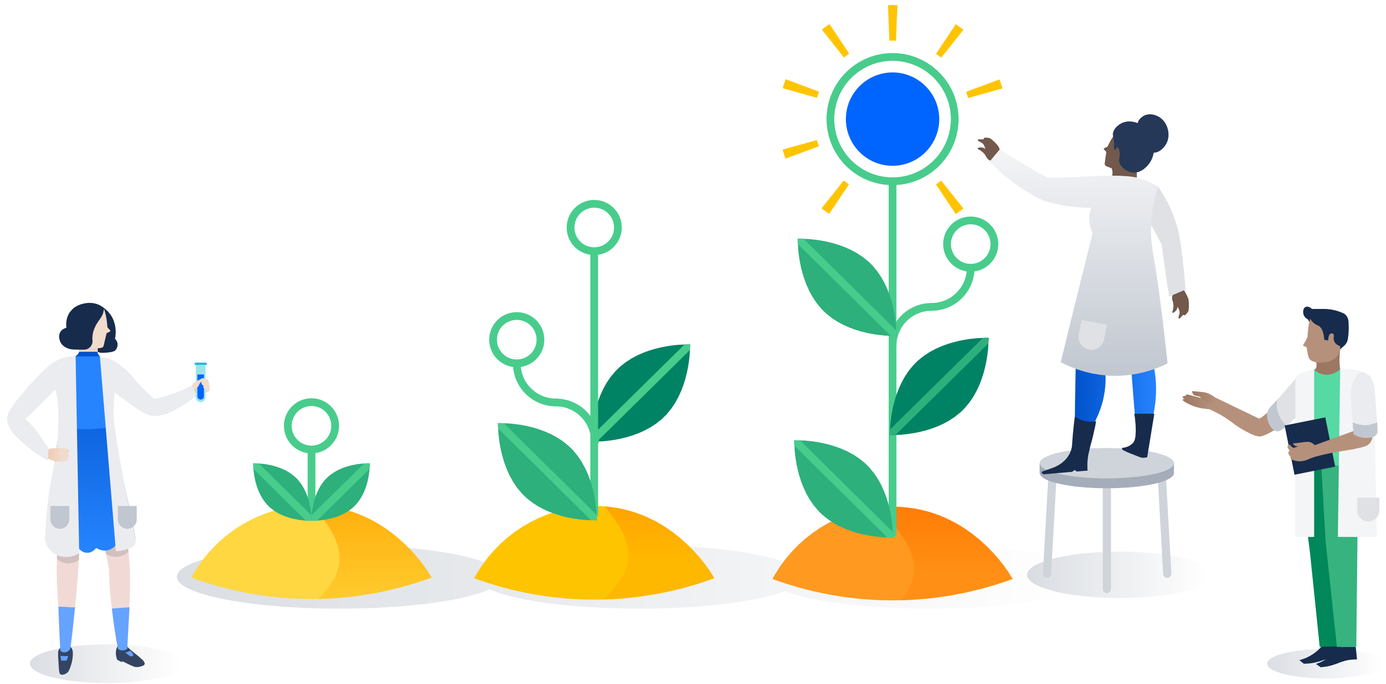

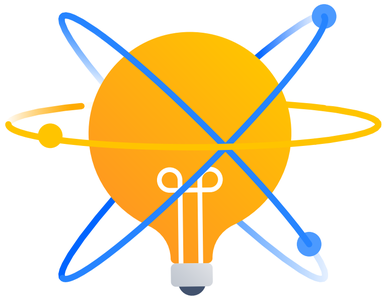

Goal

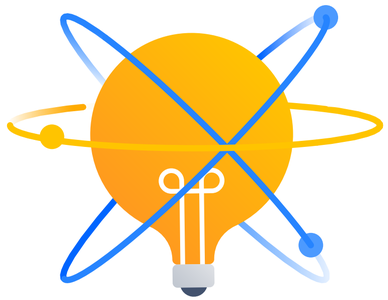

Goal

Reduce the number of noisy alerts

Data

Data

Alerts received in the past week

87

87

Low priority alerts

Report

Report

Alerts, dashboard screenshots, incidents...

Meet & discuss

Meet & discuss

Actionable? Discoverable? Useful?

Meet & discuss

Actionable? Discoverable? Useful?

Meet & discuss

Actionable? **Discoverable?** Useful?

Meet & discuss

Actionable? Discoverable? **Useful?**

ALERTS (ALL SERVICES, STAGING + PRODUCTION)

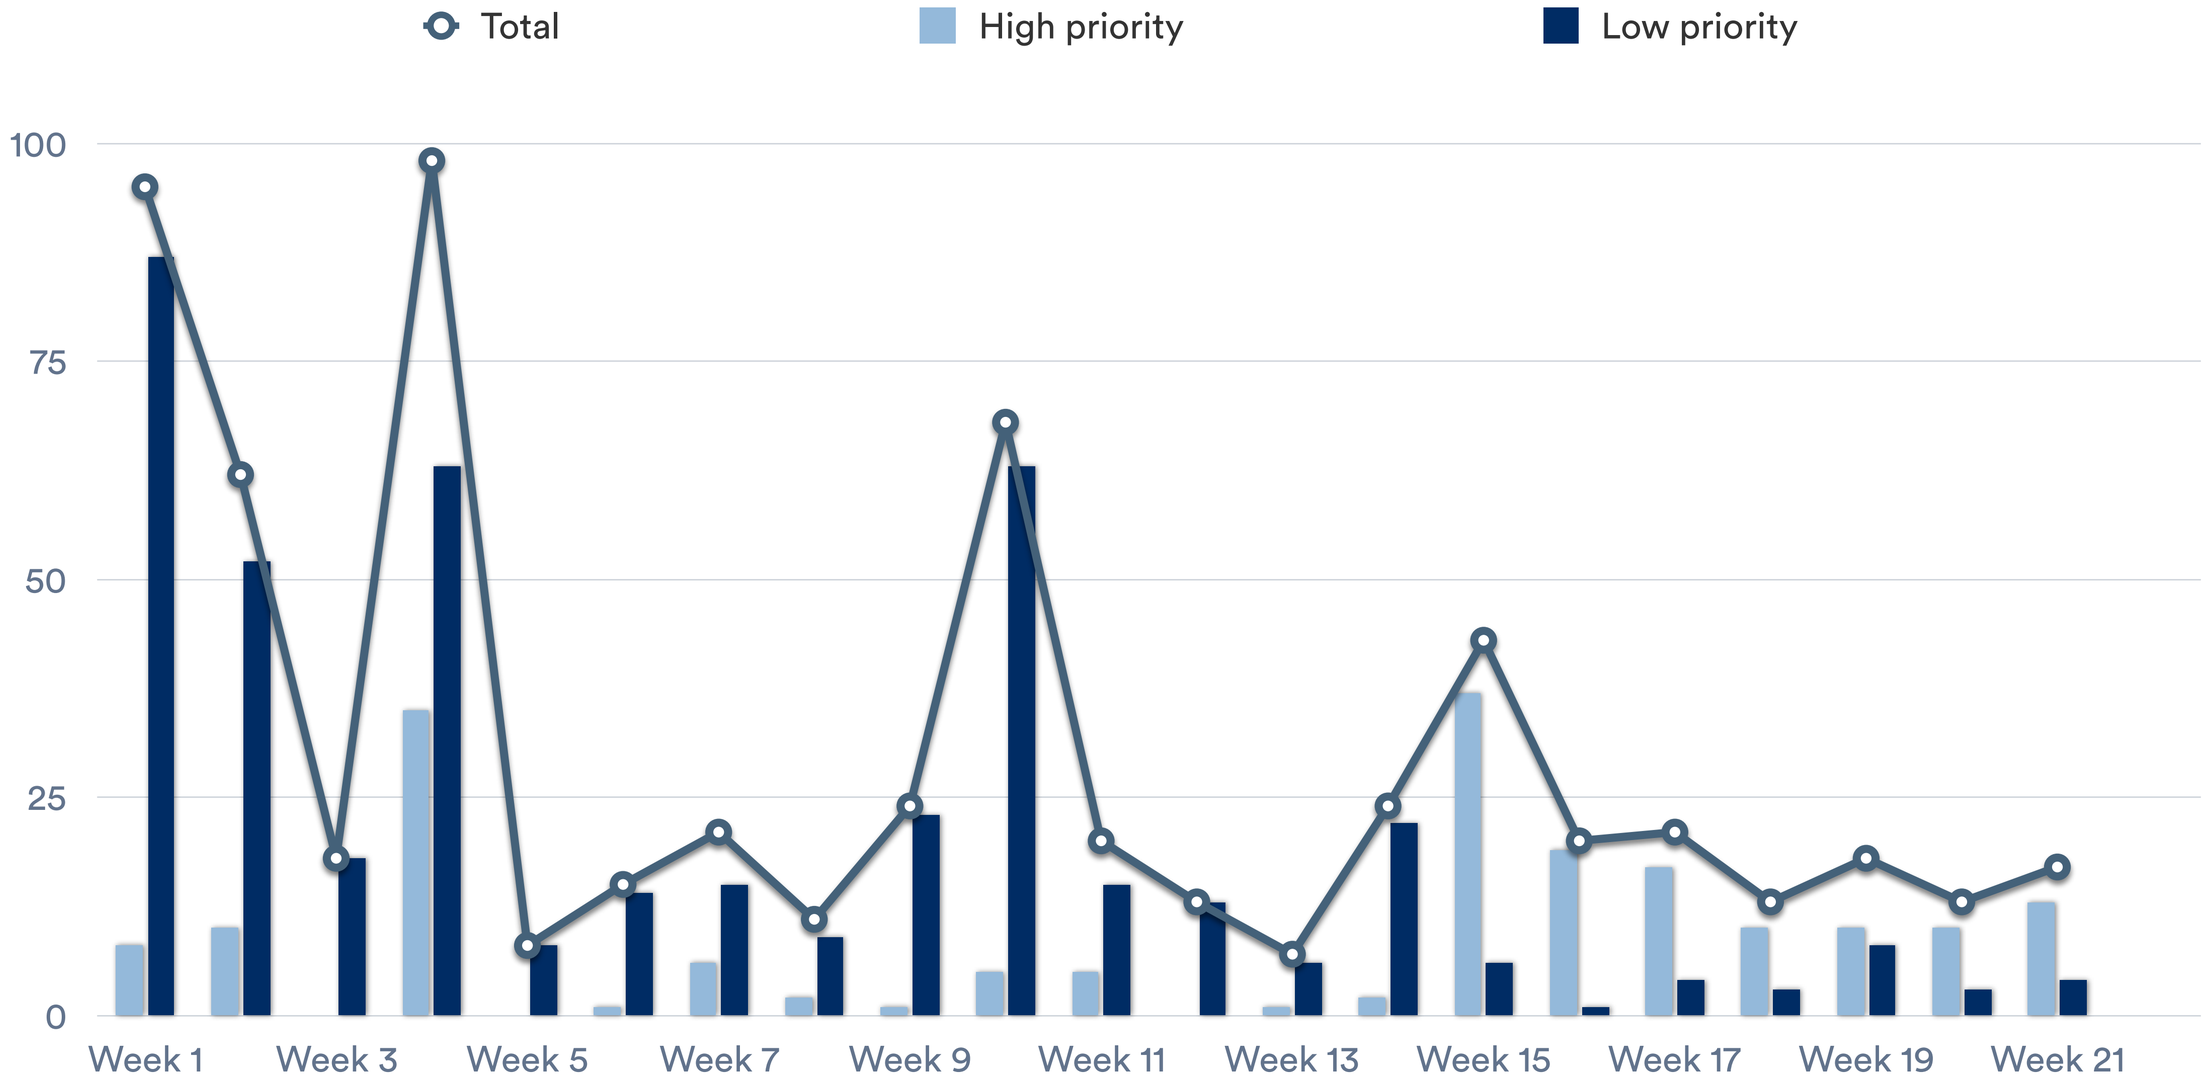

CASE #2 - RELIABILITY INCREASE

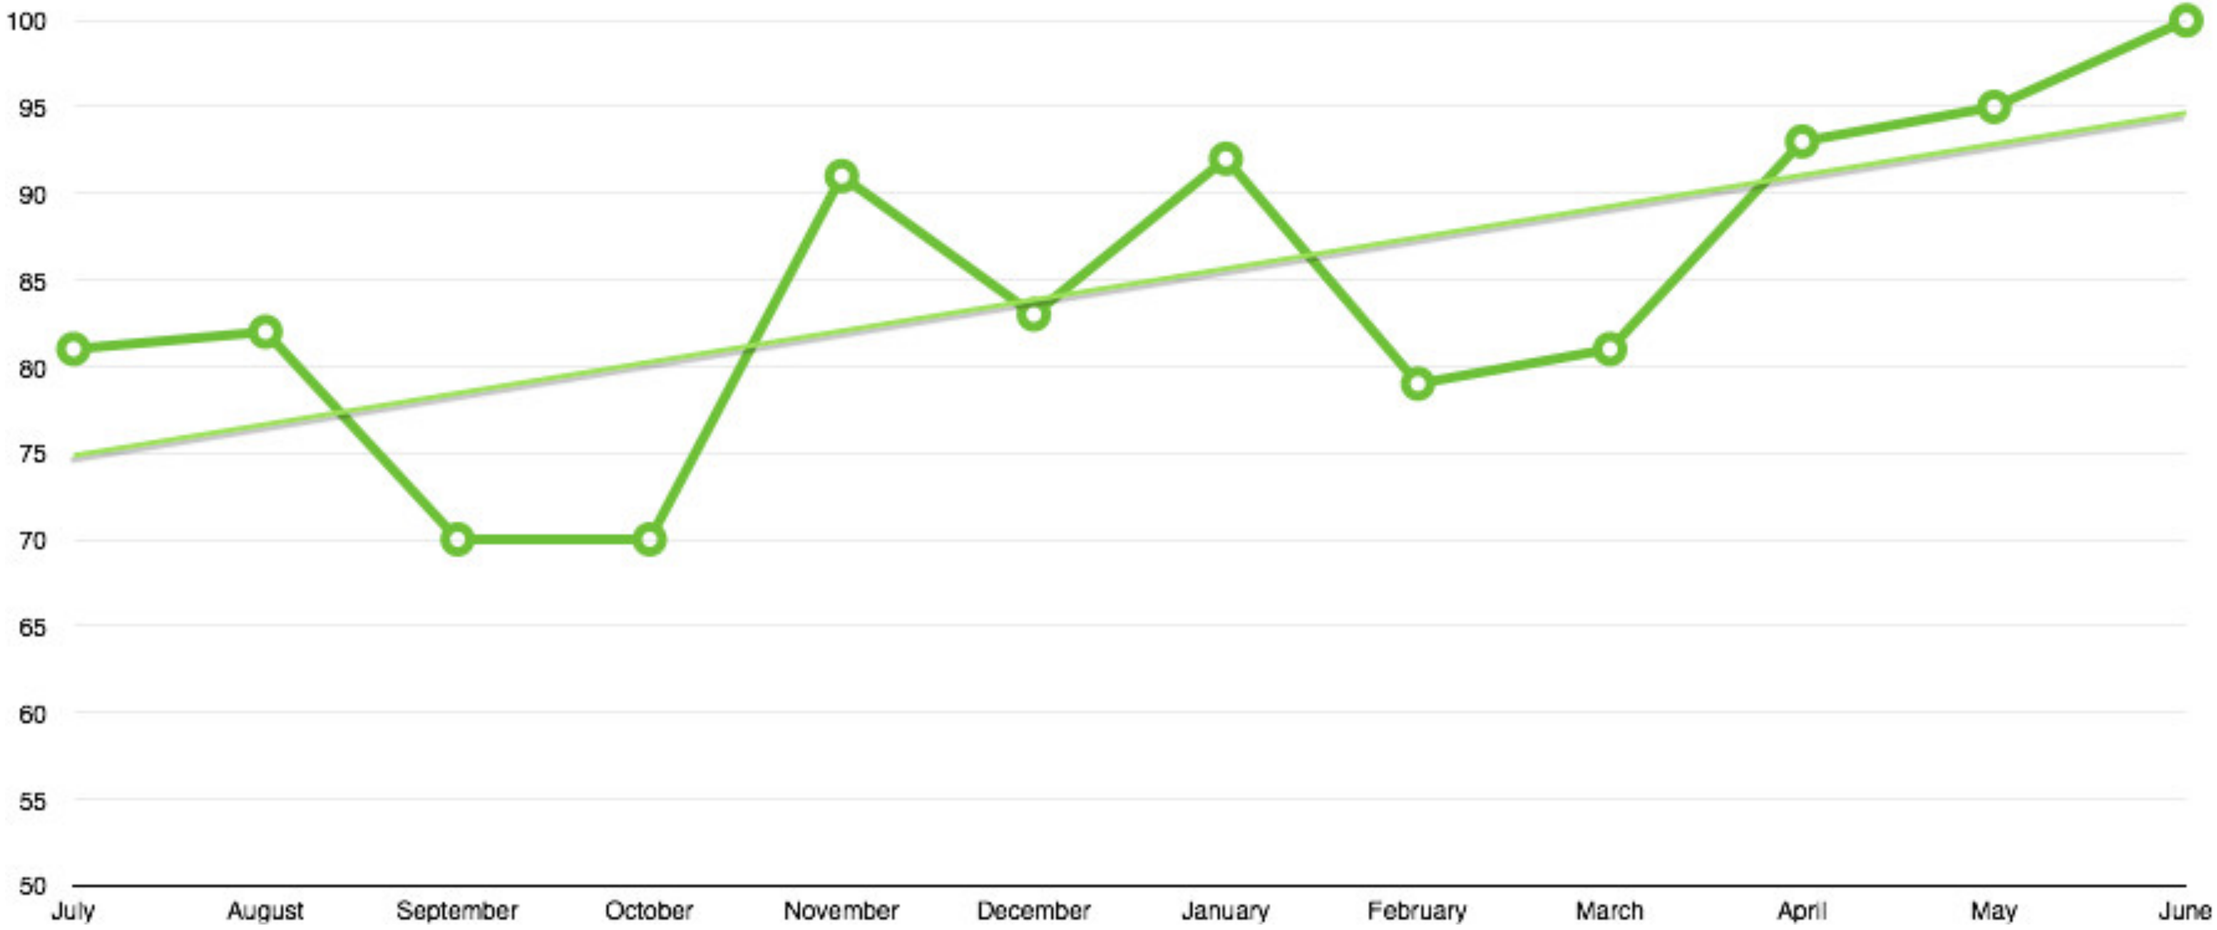

CASE #3 - ALERT REDUCTION

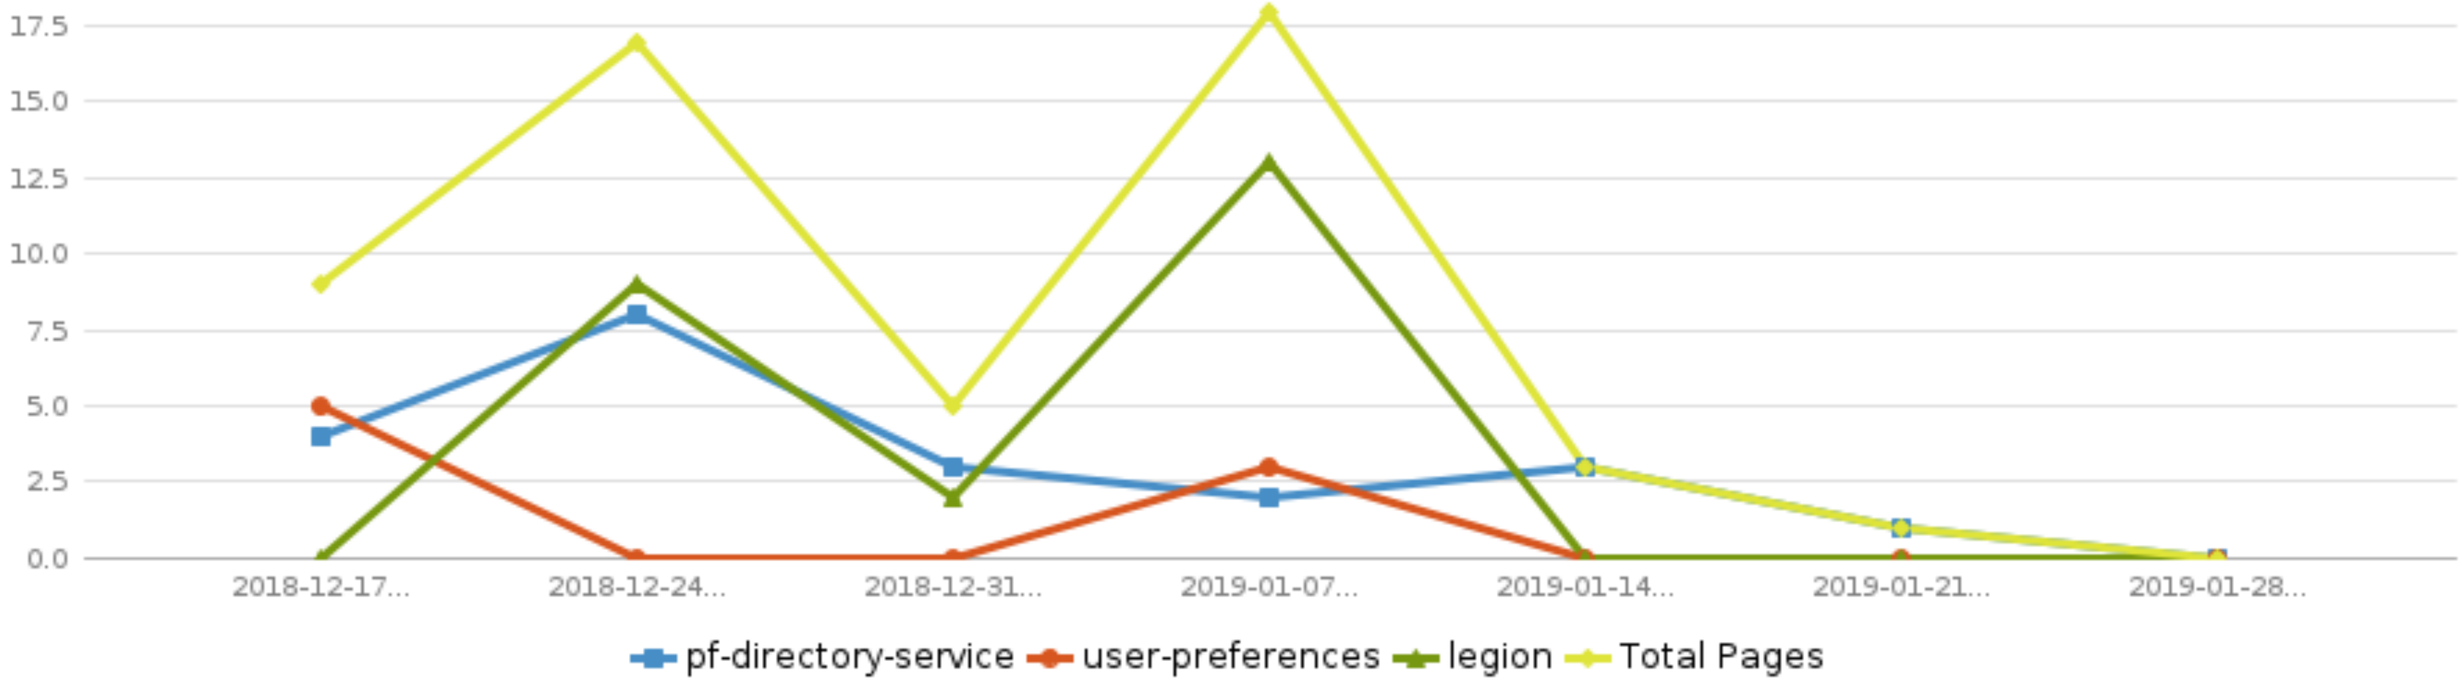

How can you...

How can you…
Verify you’re measuring the right things?

Review your operational resources!

Review your operational resources!

...frequently

Agenda

Iterative... what?

Setting some context

Deciding what to measure

Verifying your metrics

Keeping up with change

Summary

Agenda

Iterative... what?

Setting some context

Deciding what to measure

Verifying your metrics

Keeping up with change

Summary

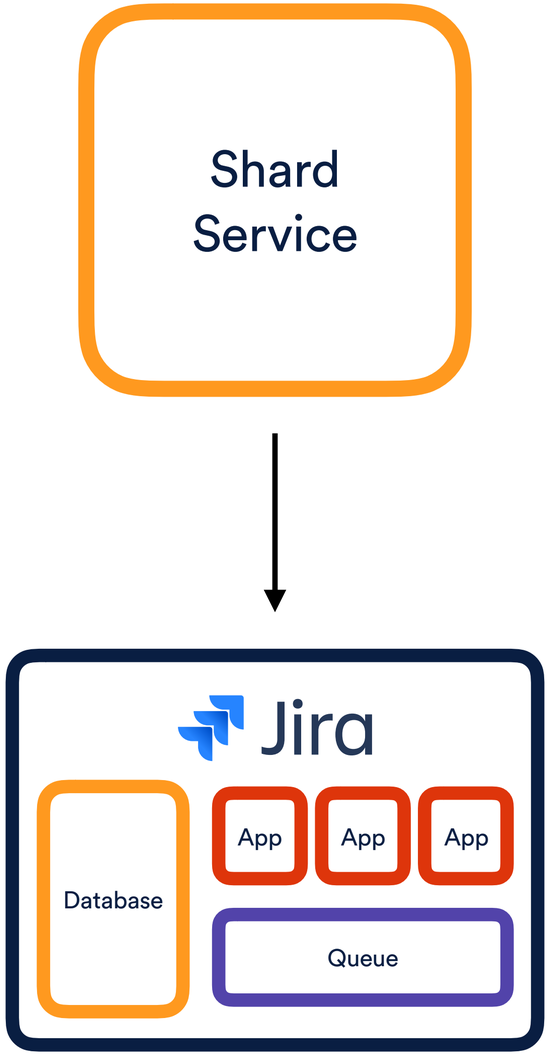
Shard
Service
Jira
Database
App
App
App
Queue

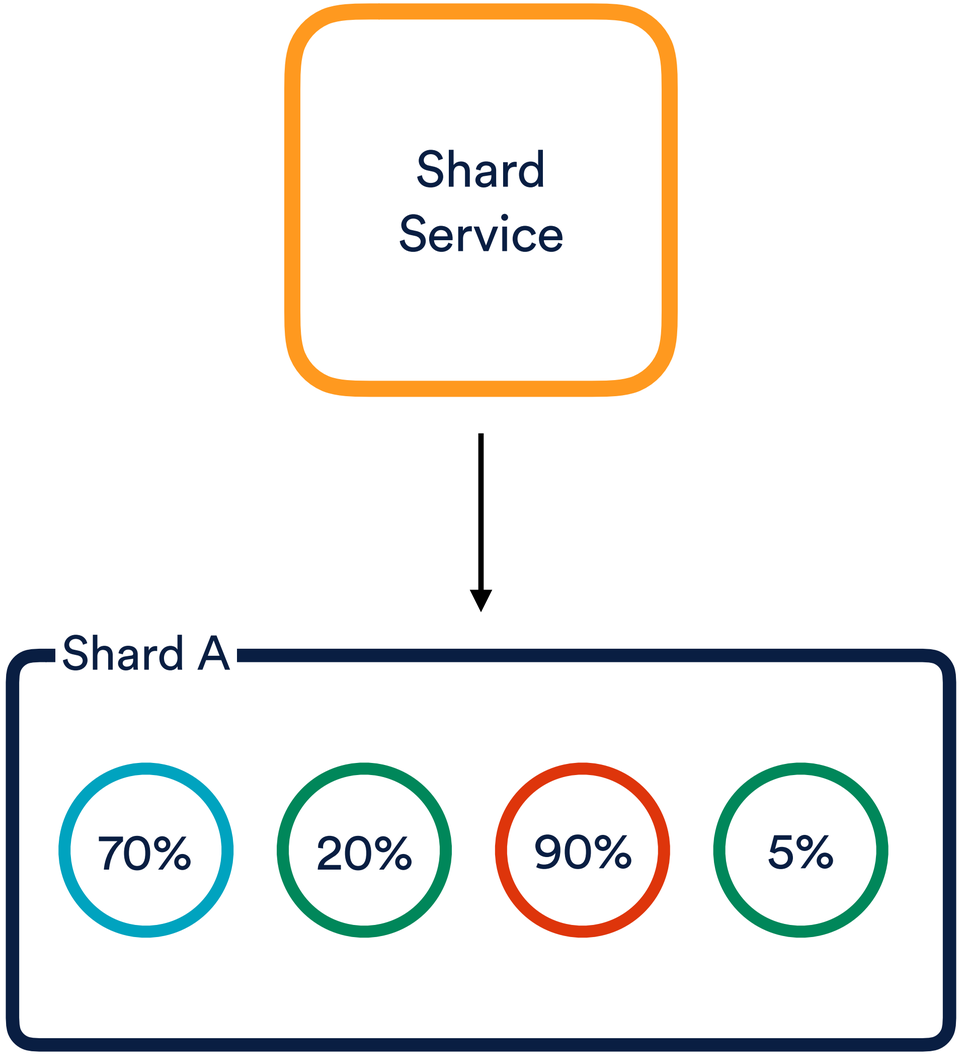
Shard
Service
Shard A
70%
20%
90%
5%

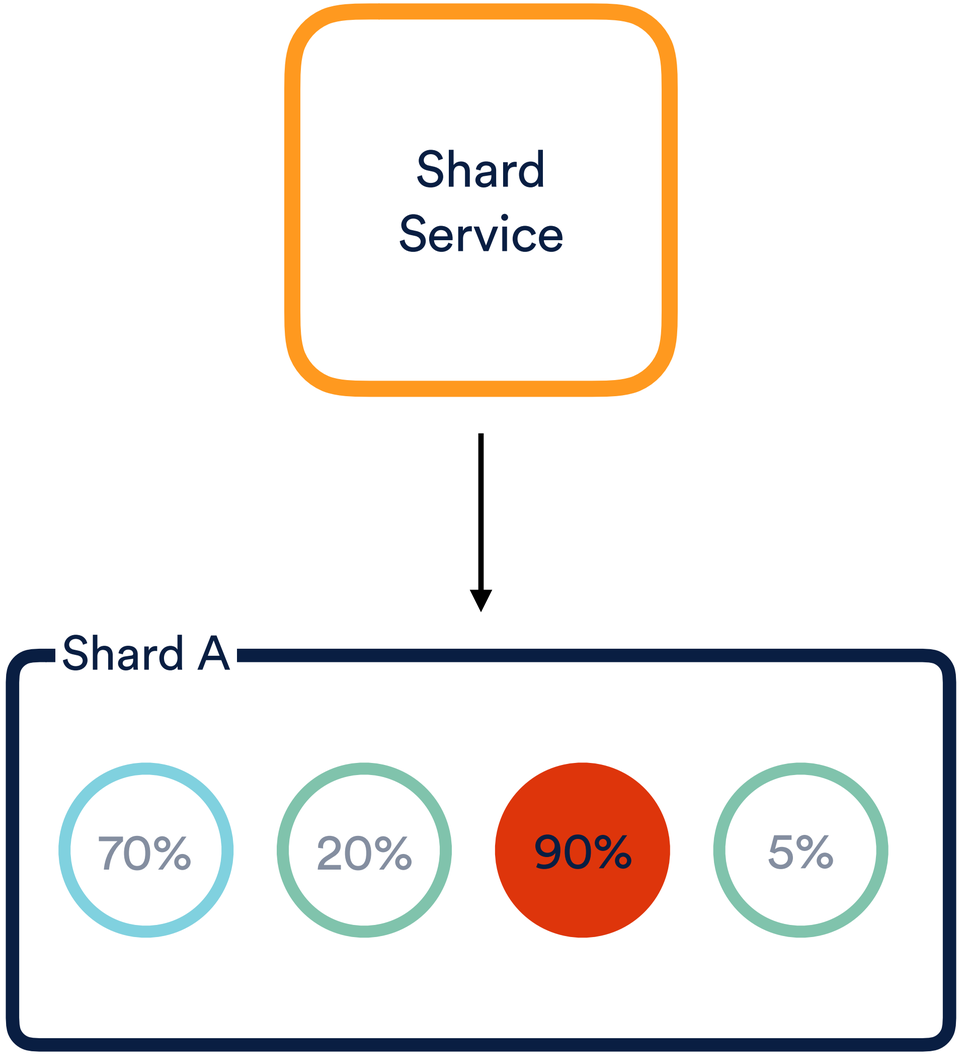
Shard
Service
Shard A
70%
20%
90%
5%

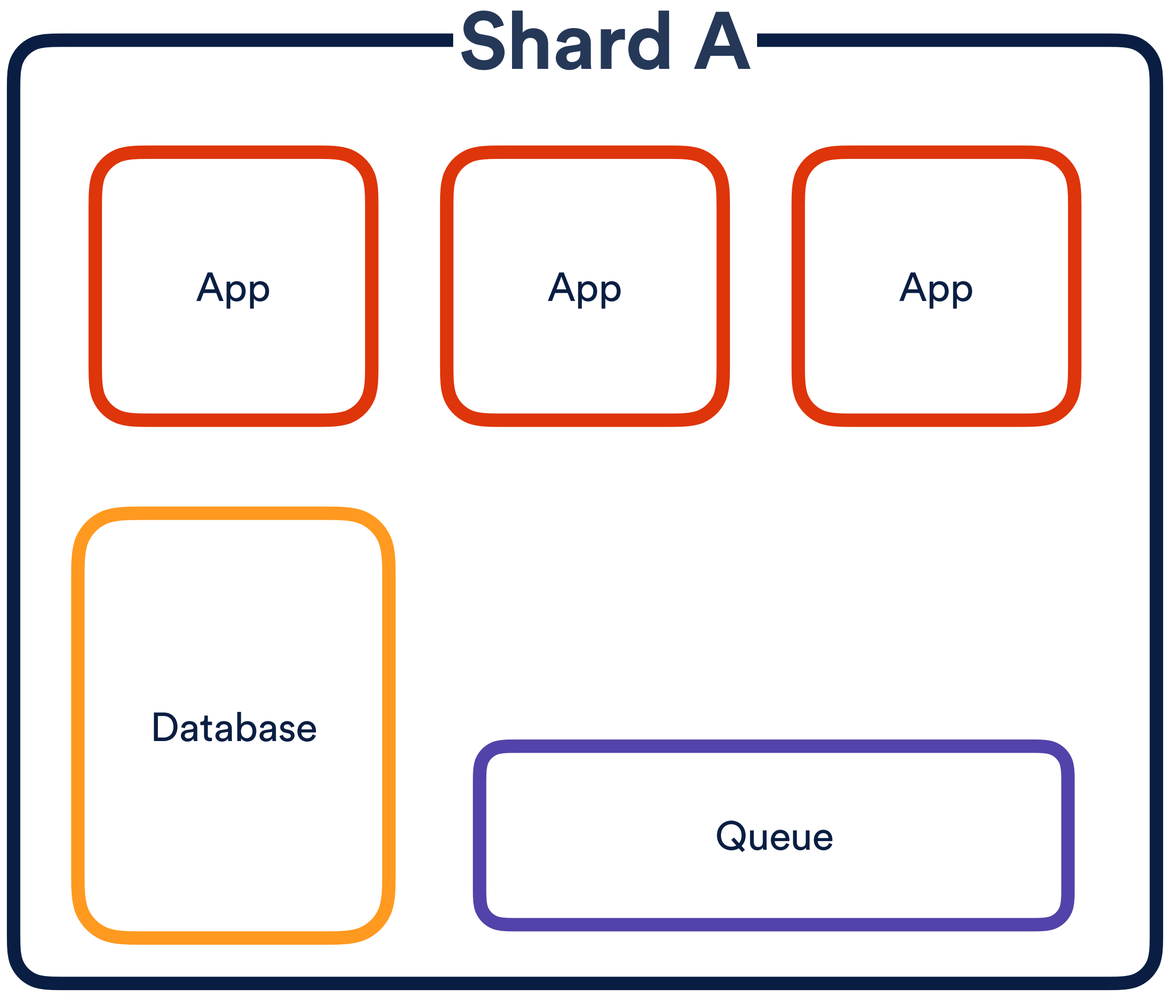
Shard A
App
App
App
Database
Queue

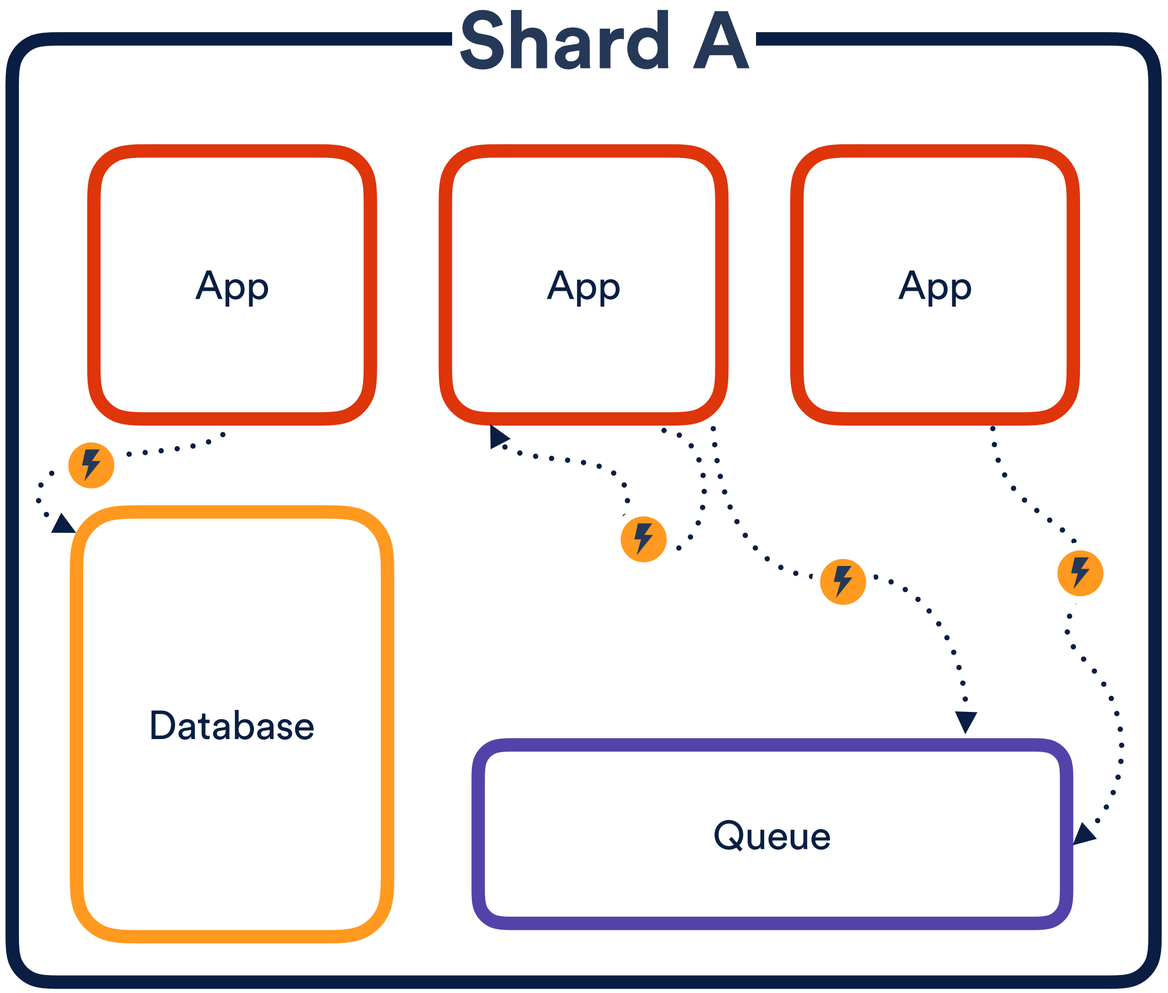

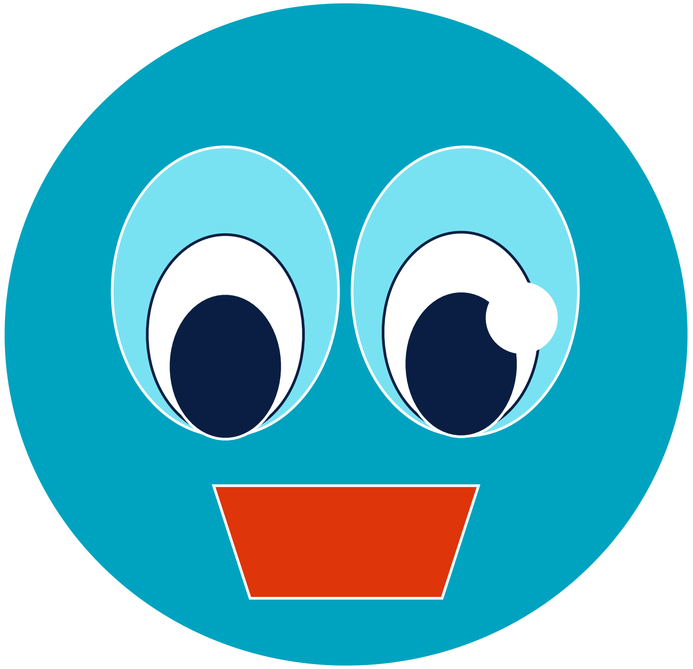

Shard Sevice

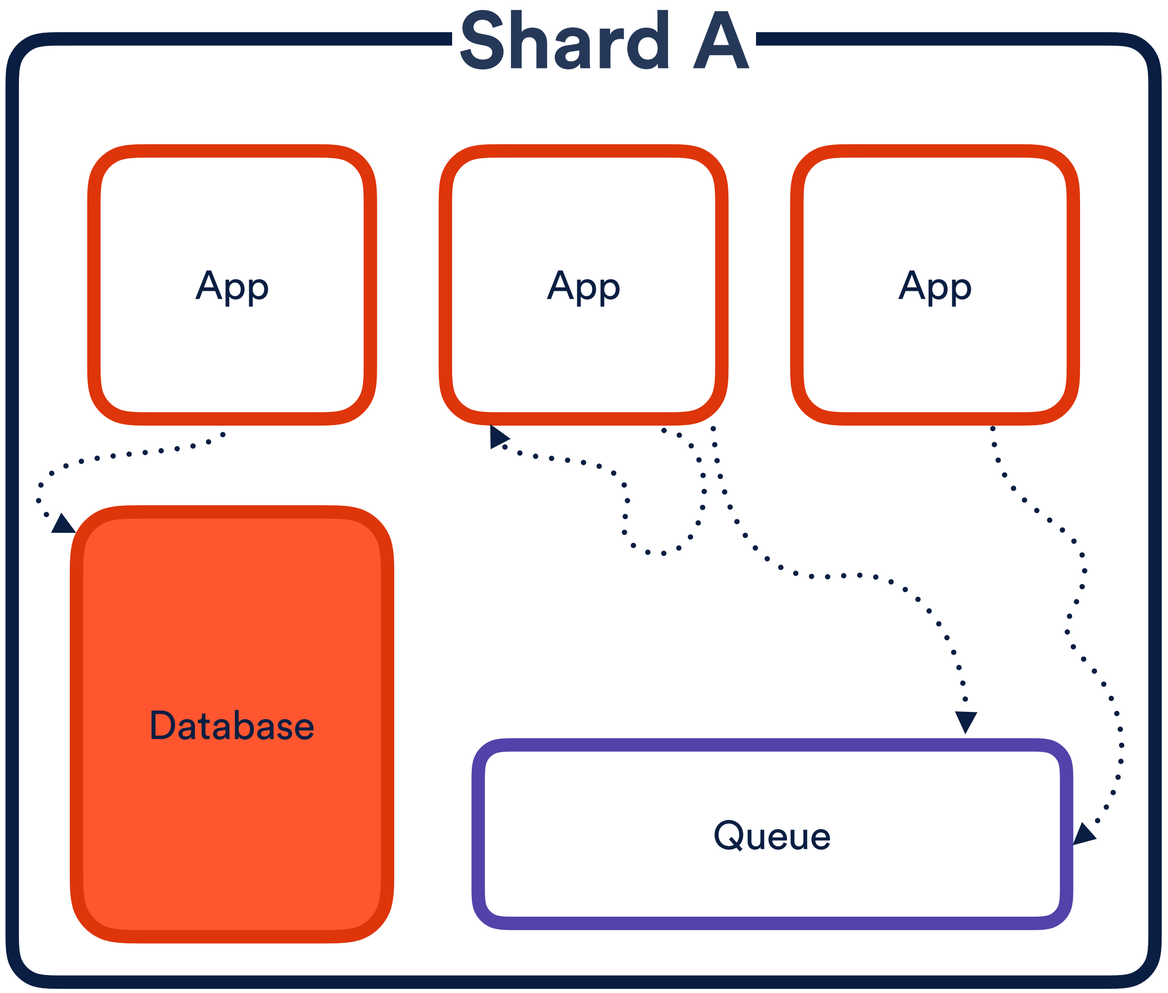

Shard Sevice

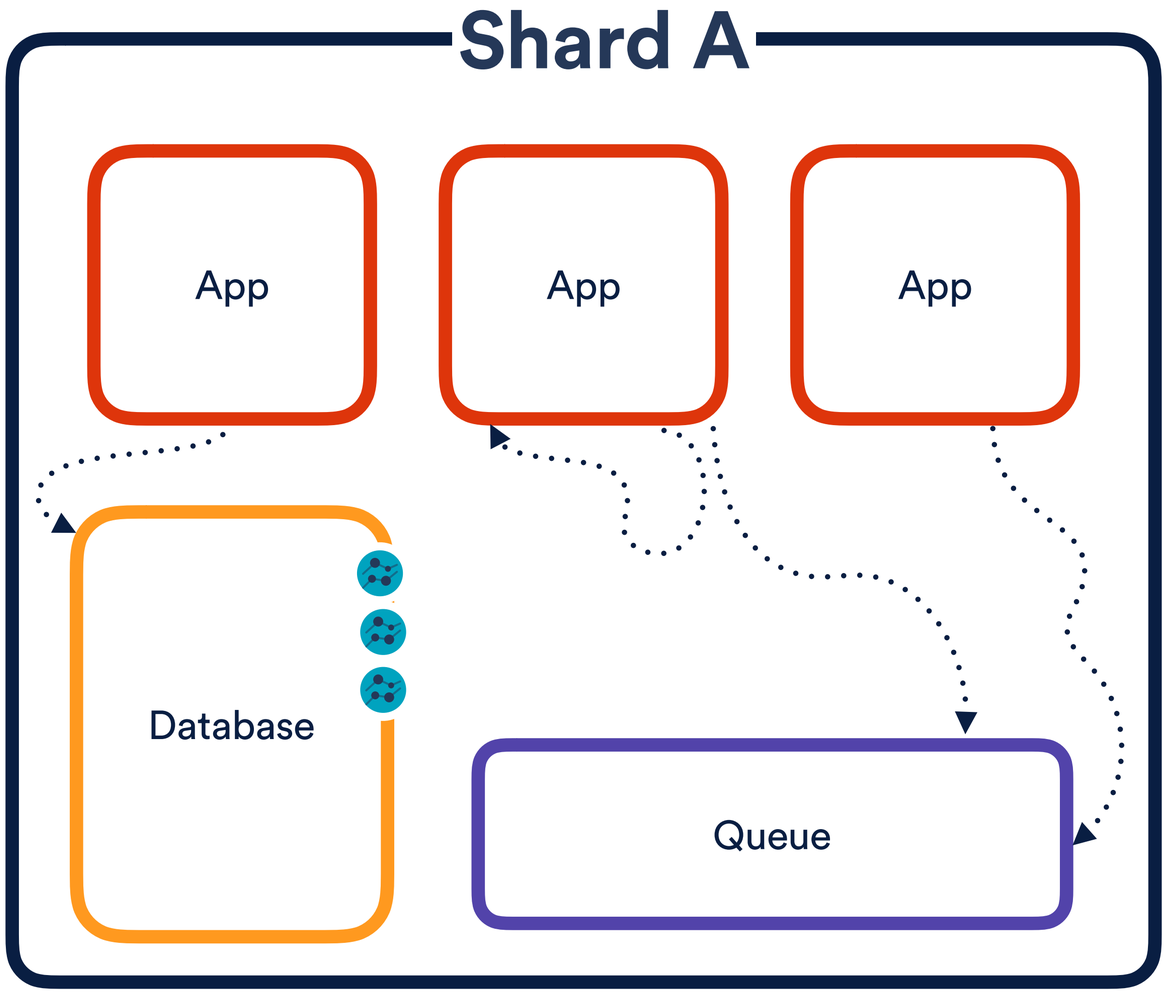

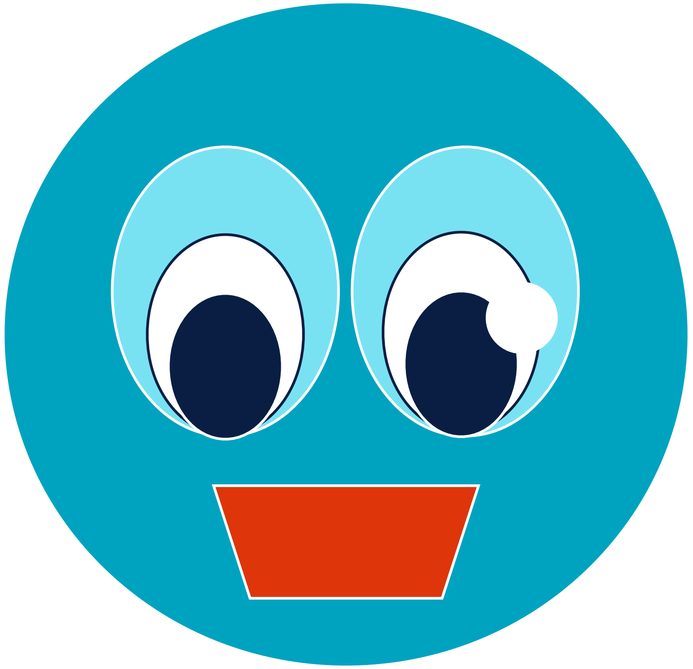

Shard Sevice

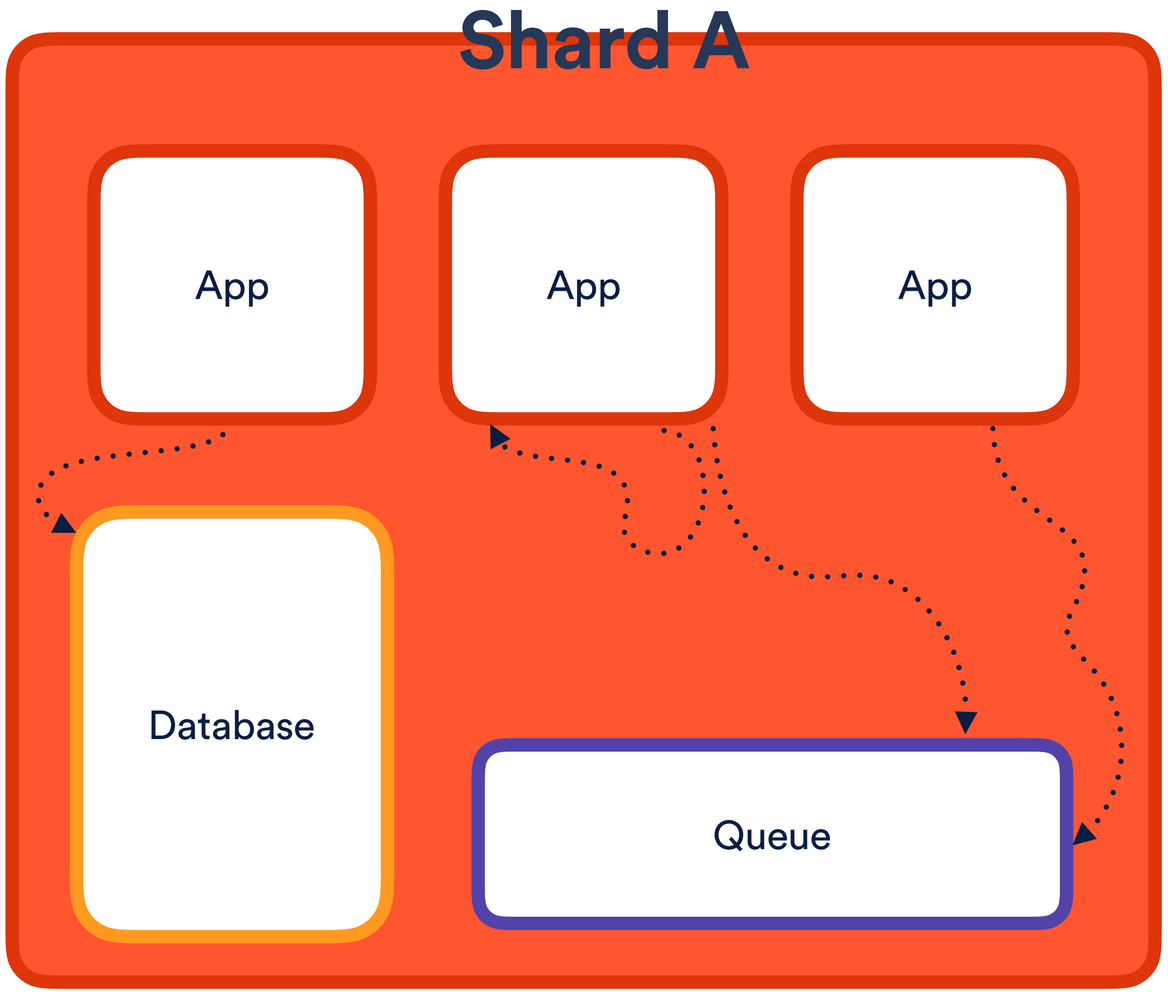

Shard Sevice

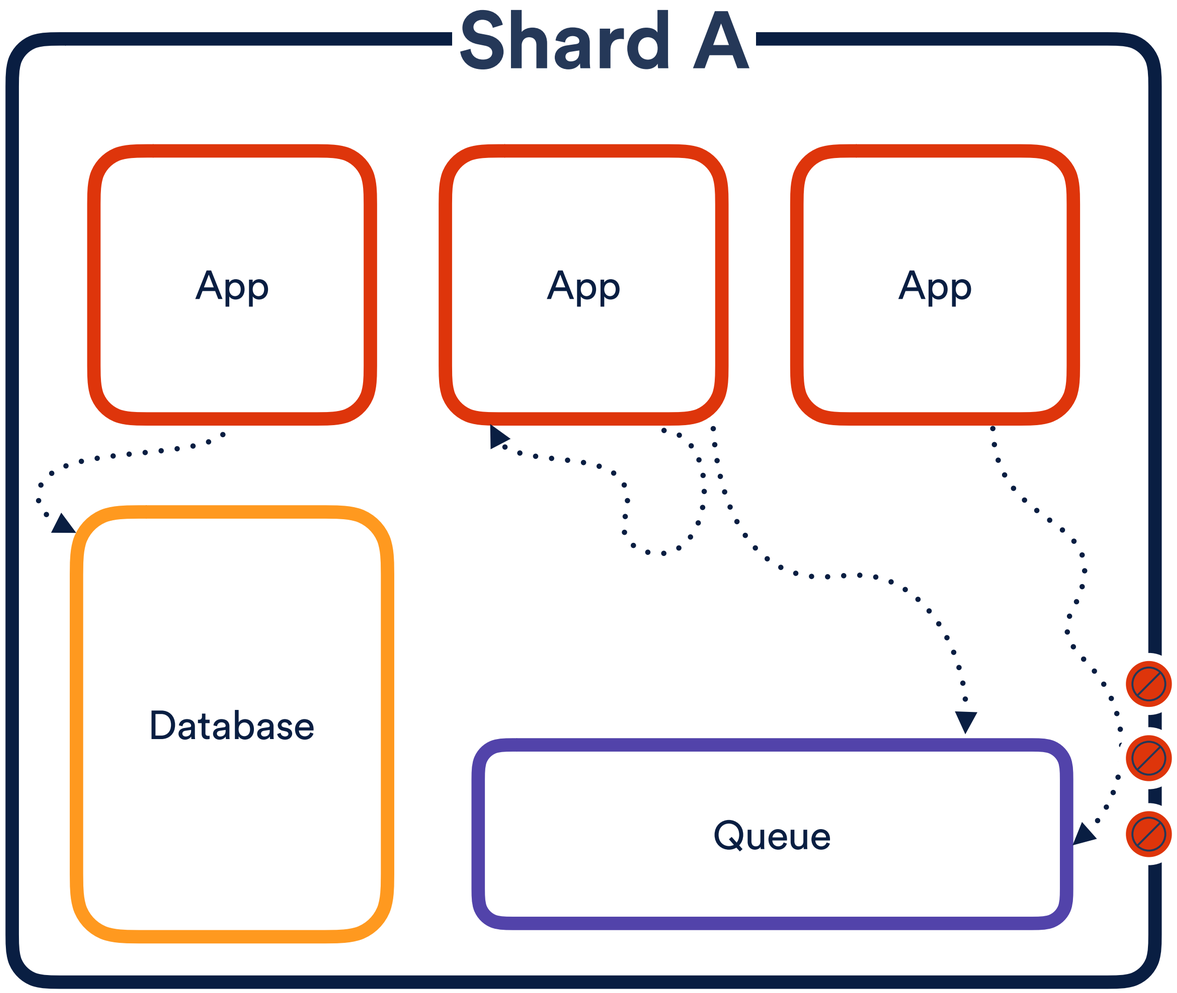

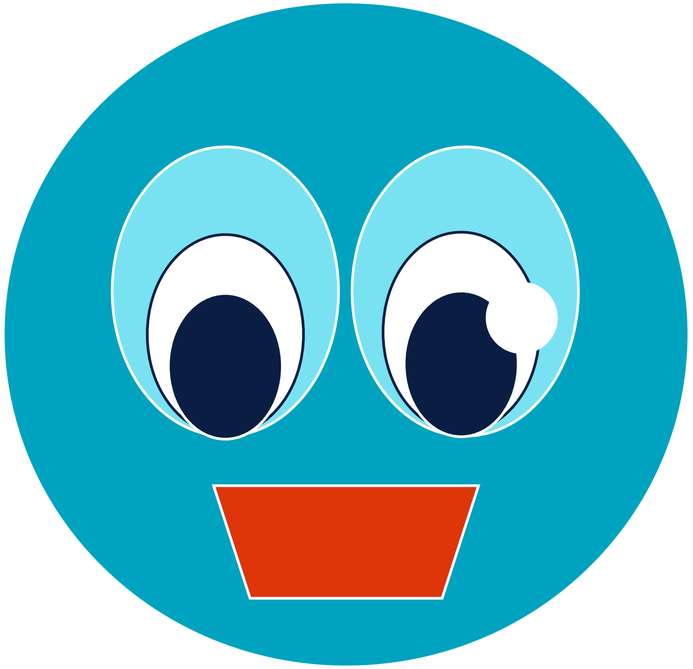

Shard Sevice

What’s the big deal?

Shard Service: selection

Edit Board

1m The Past Month

Search...

$shard_type * | $shard_status * | $shard_name * | $shard_purpose * | $shard_aws_region * | $shard_deployment_group production | $shard_select_reason * | $shard_micros_region * | $timezone_range *

236062

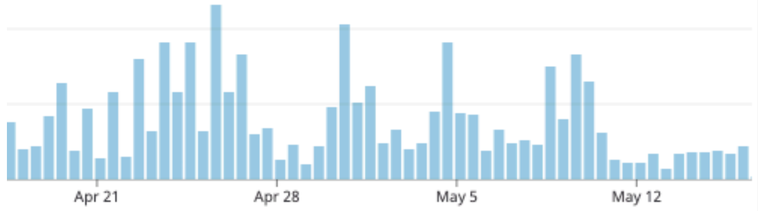

Top selected shards

- jira-prod-us-7
- jira-prod-eu-8
- confluence-prod-us-9
- jira-prod-eu-7
- jira-prod-us-19
- jira-prod-ap-2
- jira-prod-us-2
- jira-prod-us-11
- confluence-prod-us-8
- jira-prod-ap-1
- jira-prod-us-1
- confluence-prod-us-3

Top selection reasons

- specified_shard_in_request
- in_offset_and_available
- has_country_and_available
- cp_state
- in_region_and_available
- specified_shard_in_cloud_name
- cloud_name_security

Top selected region (AWS)

- us-east-1
- us-west-2
- eu-west-1
- eu-central-1
- ap-southeast-1
- ap-southeast-2

Top selected region (internal)

- prod_east
- prod_west2
- prod_euwest
- prod_eucentral
- prod_apse
- prod_apse2

Database utilisation metric: current vs target (by shard)

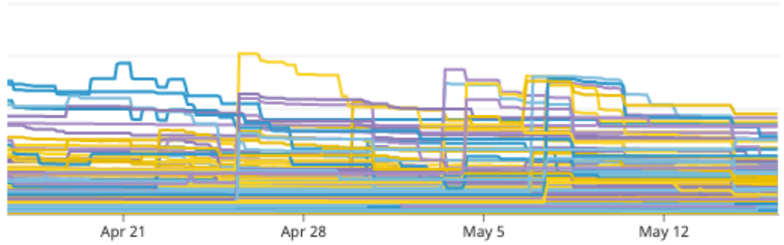

Jira Apdex Count metric: current vs target (by shard)

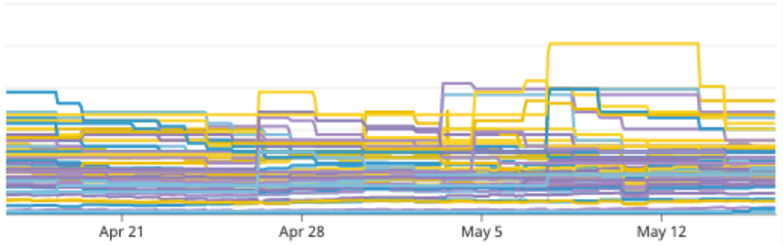

Confluence Apdex Count metric: current vs target (by shard)

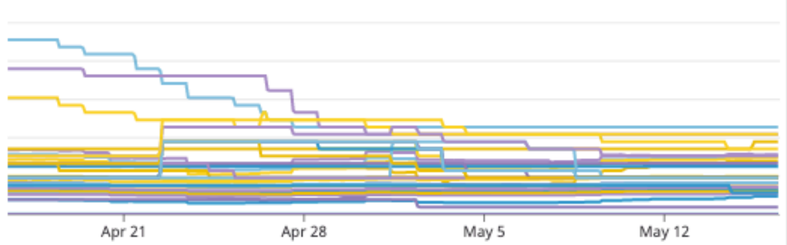

Provisioning failures metric: current vs target (by shard)

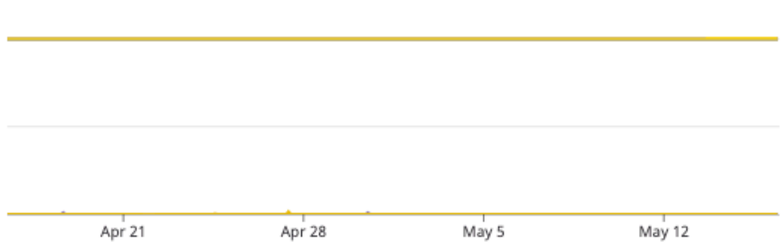

Average remaining capacity by shard (Jira Apdex)

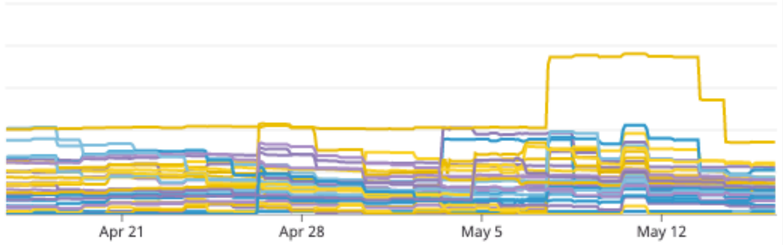

Average remaining capacity by shard (Confluence Apdex)

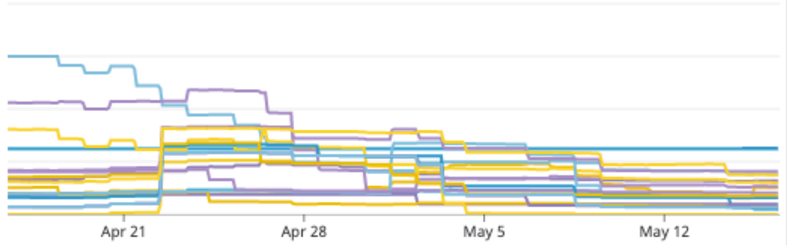

Shard Service: selection

Edit Board

1m The Past Month

Search... $shard_type * $shard_status * $shard_name * $shard_purpose * $shard_aws_region * $shard_deployment_group production $shard_select_reason * $shard_micros_region *

$timezone_range *

23606

Apr 21 Apr 28 May 5 May 12

Top selected shards

jira-prod-us-7
jira-prod-eu-8
confluence-prod-us-9
jira-prod-eu-7
jira-prod-us-19
jira-prod-ap-2
jira-prod-us-2
jira-prod-us-11
confluence-prod-us-8
jira-prod-ap-1
jira-prod-us-1
confluence-prod-us-3

Top selection reasons

specified_shard_in_request
in_offset_and_available
has_country_and_available
cp_state
in_region_and_available
specified_shard_in_cloud_name
cloud_name_security

Top selected regio

us-east-1
us-west-2
eu-west-1
eu-central-1
ap-southeast-1
ap-southeast-2

Top selected region (internal)

prod_east
prod_west2
prod_euwest
prod_eucentral
prod_apse
prod_apse2

Panel per metric

Database utilisation metric: current vs target (by shard)

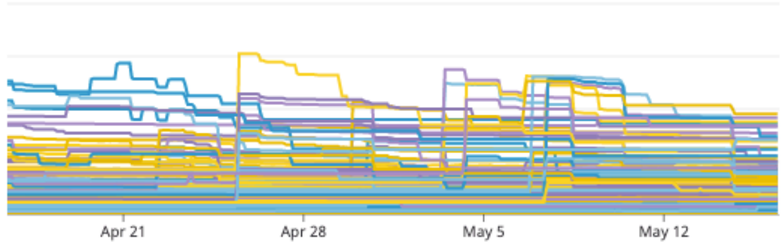

Jira Apdex Count metric: current vs target (by shard)

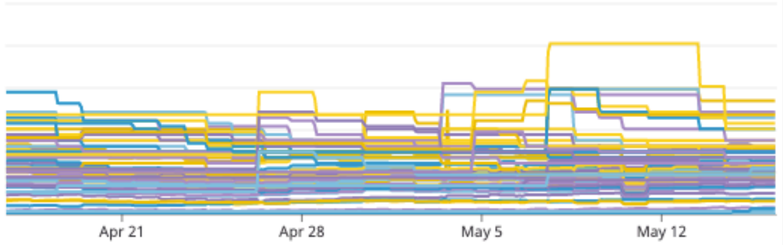

Confluence Apdex Count metric: current vs target (by shard)

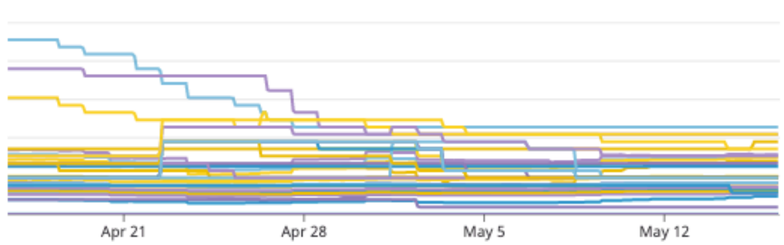

Provisioning failures metric: current vs target (by shard)

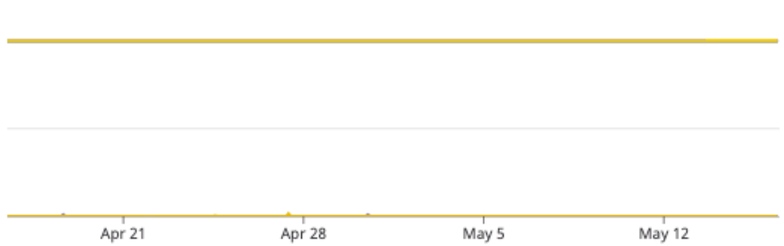

Average remaining capacity by shard (Jira Apdex)

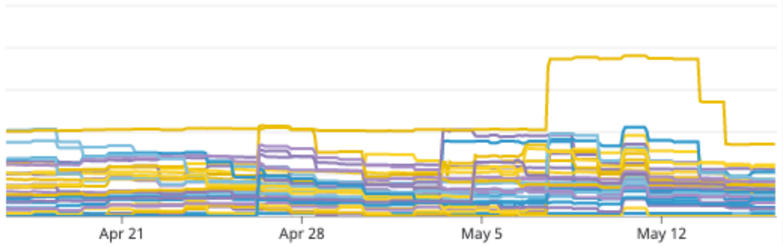

Average remaining capacity by shard (Confluence Apdex)

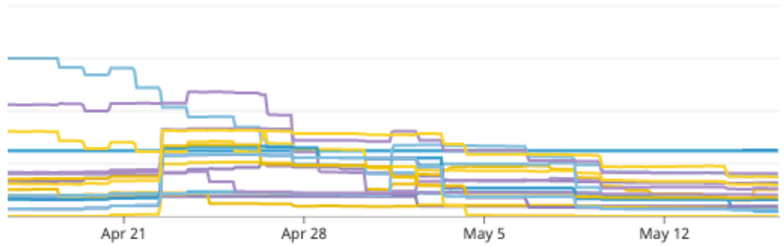

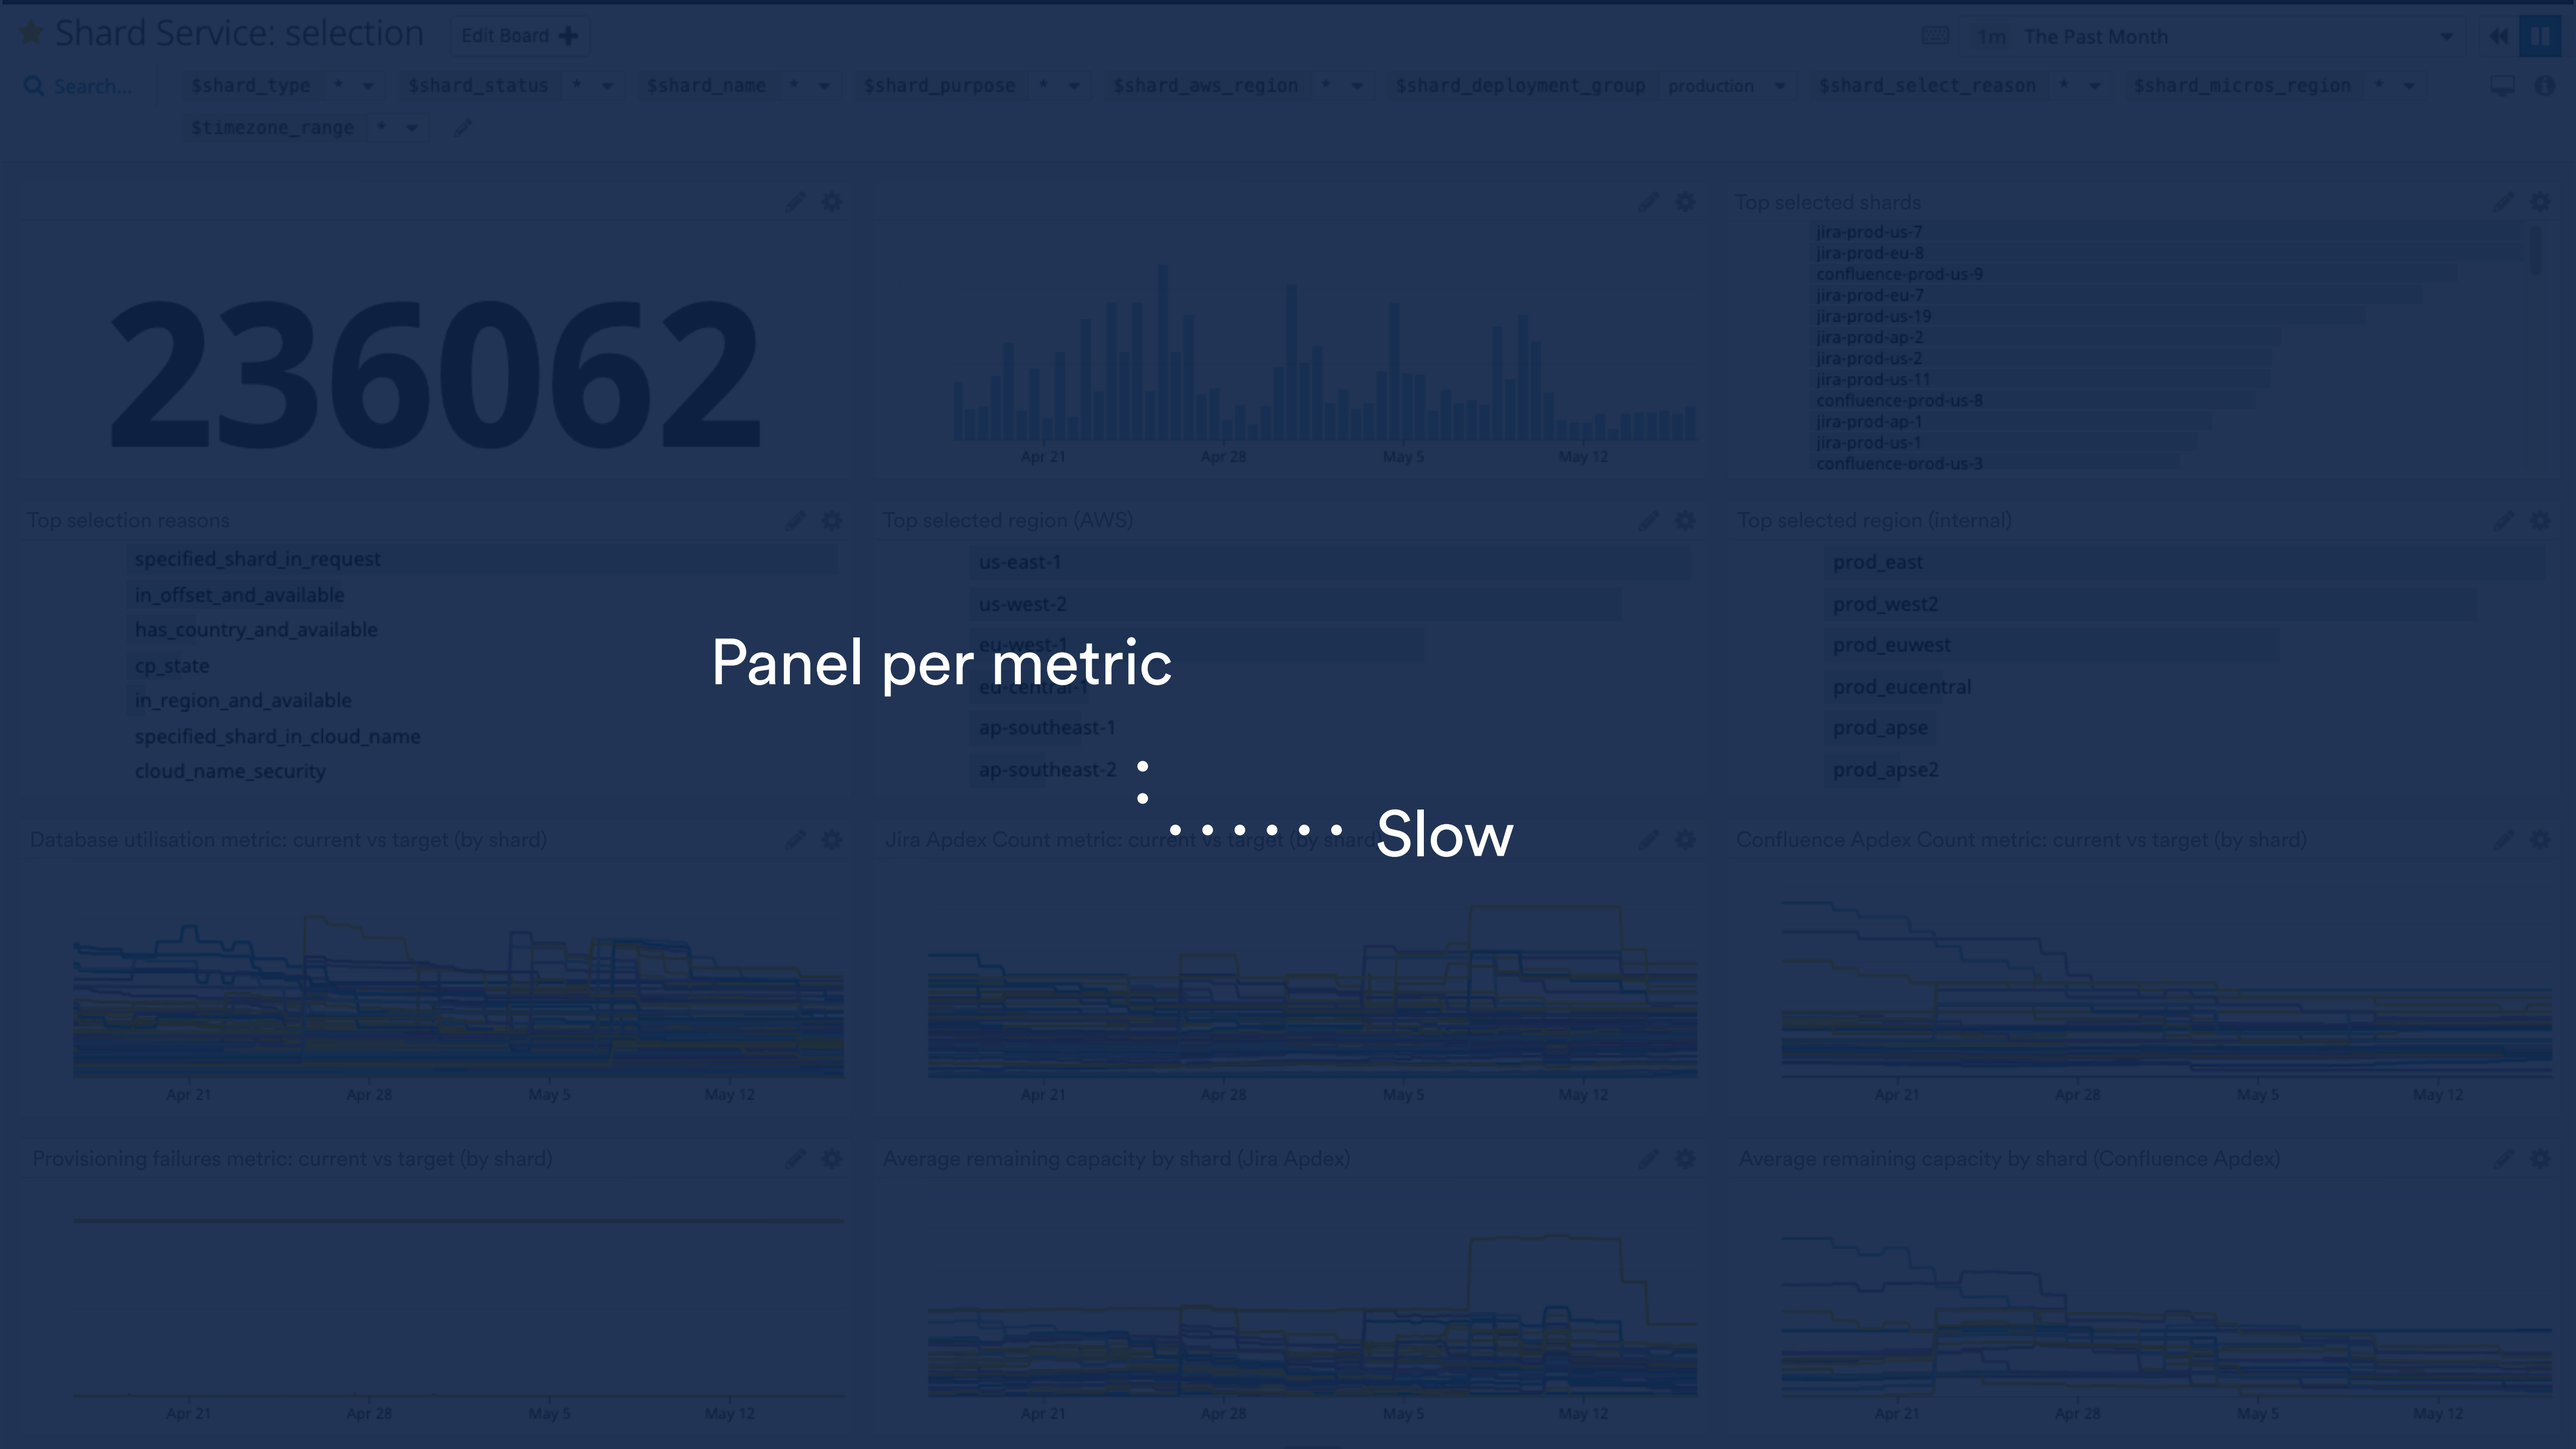
Shard Service: selection
Edit Board
1m The Past Month
Search...
$shard_type
$shard_status
$shard_name
$shard_purpose
$shard_aws_region
$shard_deployment_group
production
$shard_select_reason
$shard_micros_region
$timezone_range
236062
Apr 21
Apr 28
May 5
May 12
Top selected shards
jira-prod-us-7
jira-prod-eu-8
confluence-prod-us-9
jira-prod-eu-7
jira-prod-us-19
jira-prod-ap-2
jira-prod-us-2
jira-prod-us-11
confluence-prod-us-8
jira-prod-ap-1
jira-prod-us-1
confluence-prod-us-3
Top selection reasons
specified_shard_in_request
in_offset_and_available
has_country_and_available
cp_state
in_region_and_available
specified_shard_in_cloud_name
cloud_name_security
Top selected region (AWS)
us-east-1
us-west-2
ap-southeast-1
ap-southeast-2
Top selected region (internal)
prod_east
prod_west2
prod_euwest
prod_eucentral
prod_apse
prod_apse2
Panel per metric
Slow
Database utilisation metric: current vs target (by shard)
Jira Apdex Count metric: current vs target (by shard)
Confluence Apdex Count metric: current vs target (by shard)
Provisioning failures metric: current vs target (by shard)
Average remaining capacity by shard (Jira Apdex)
Average remaining capacity by shard (Confluence Apdex)

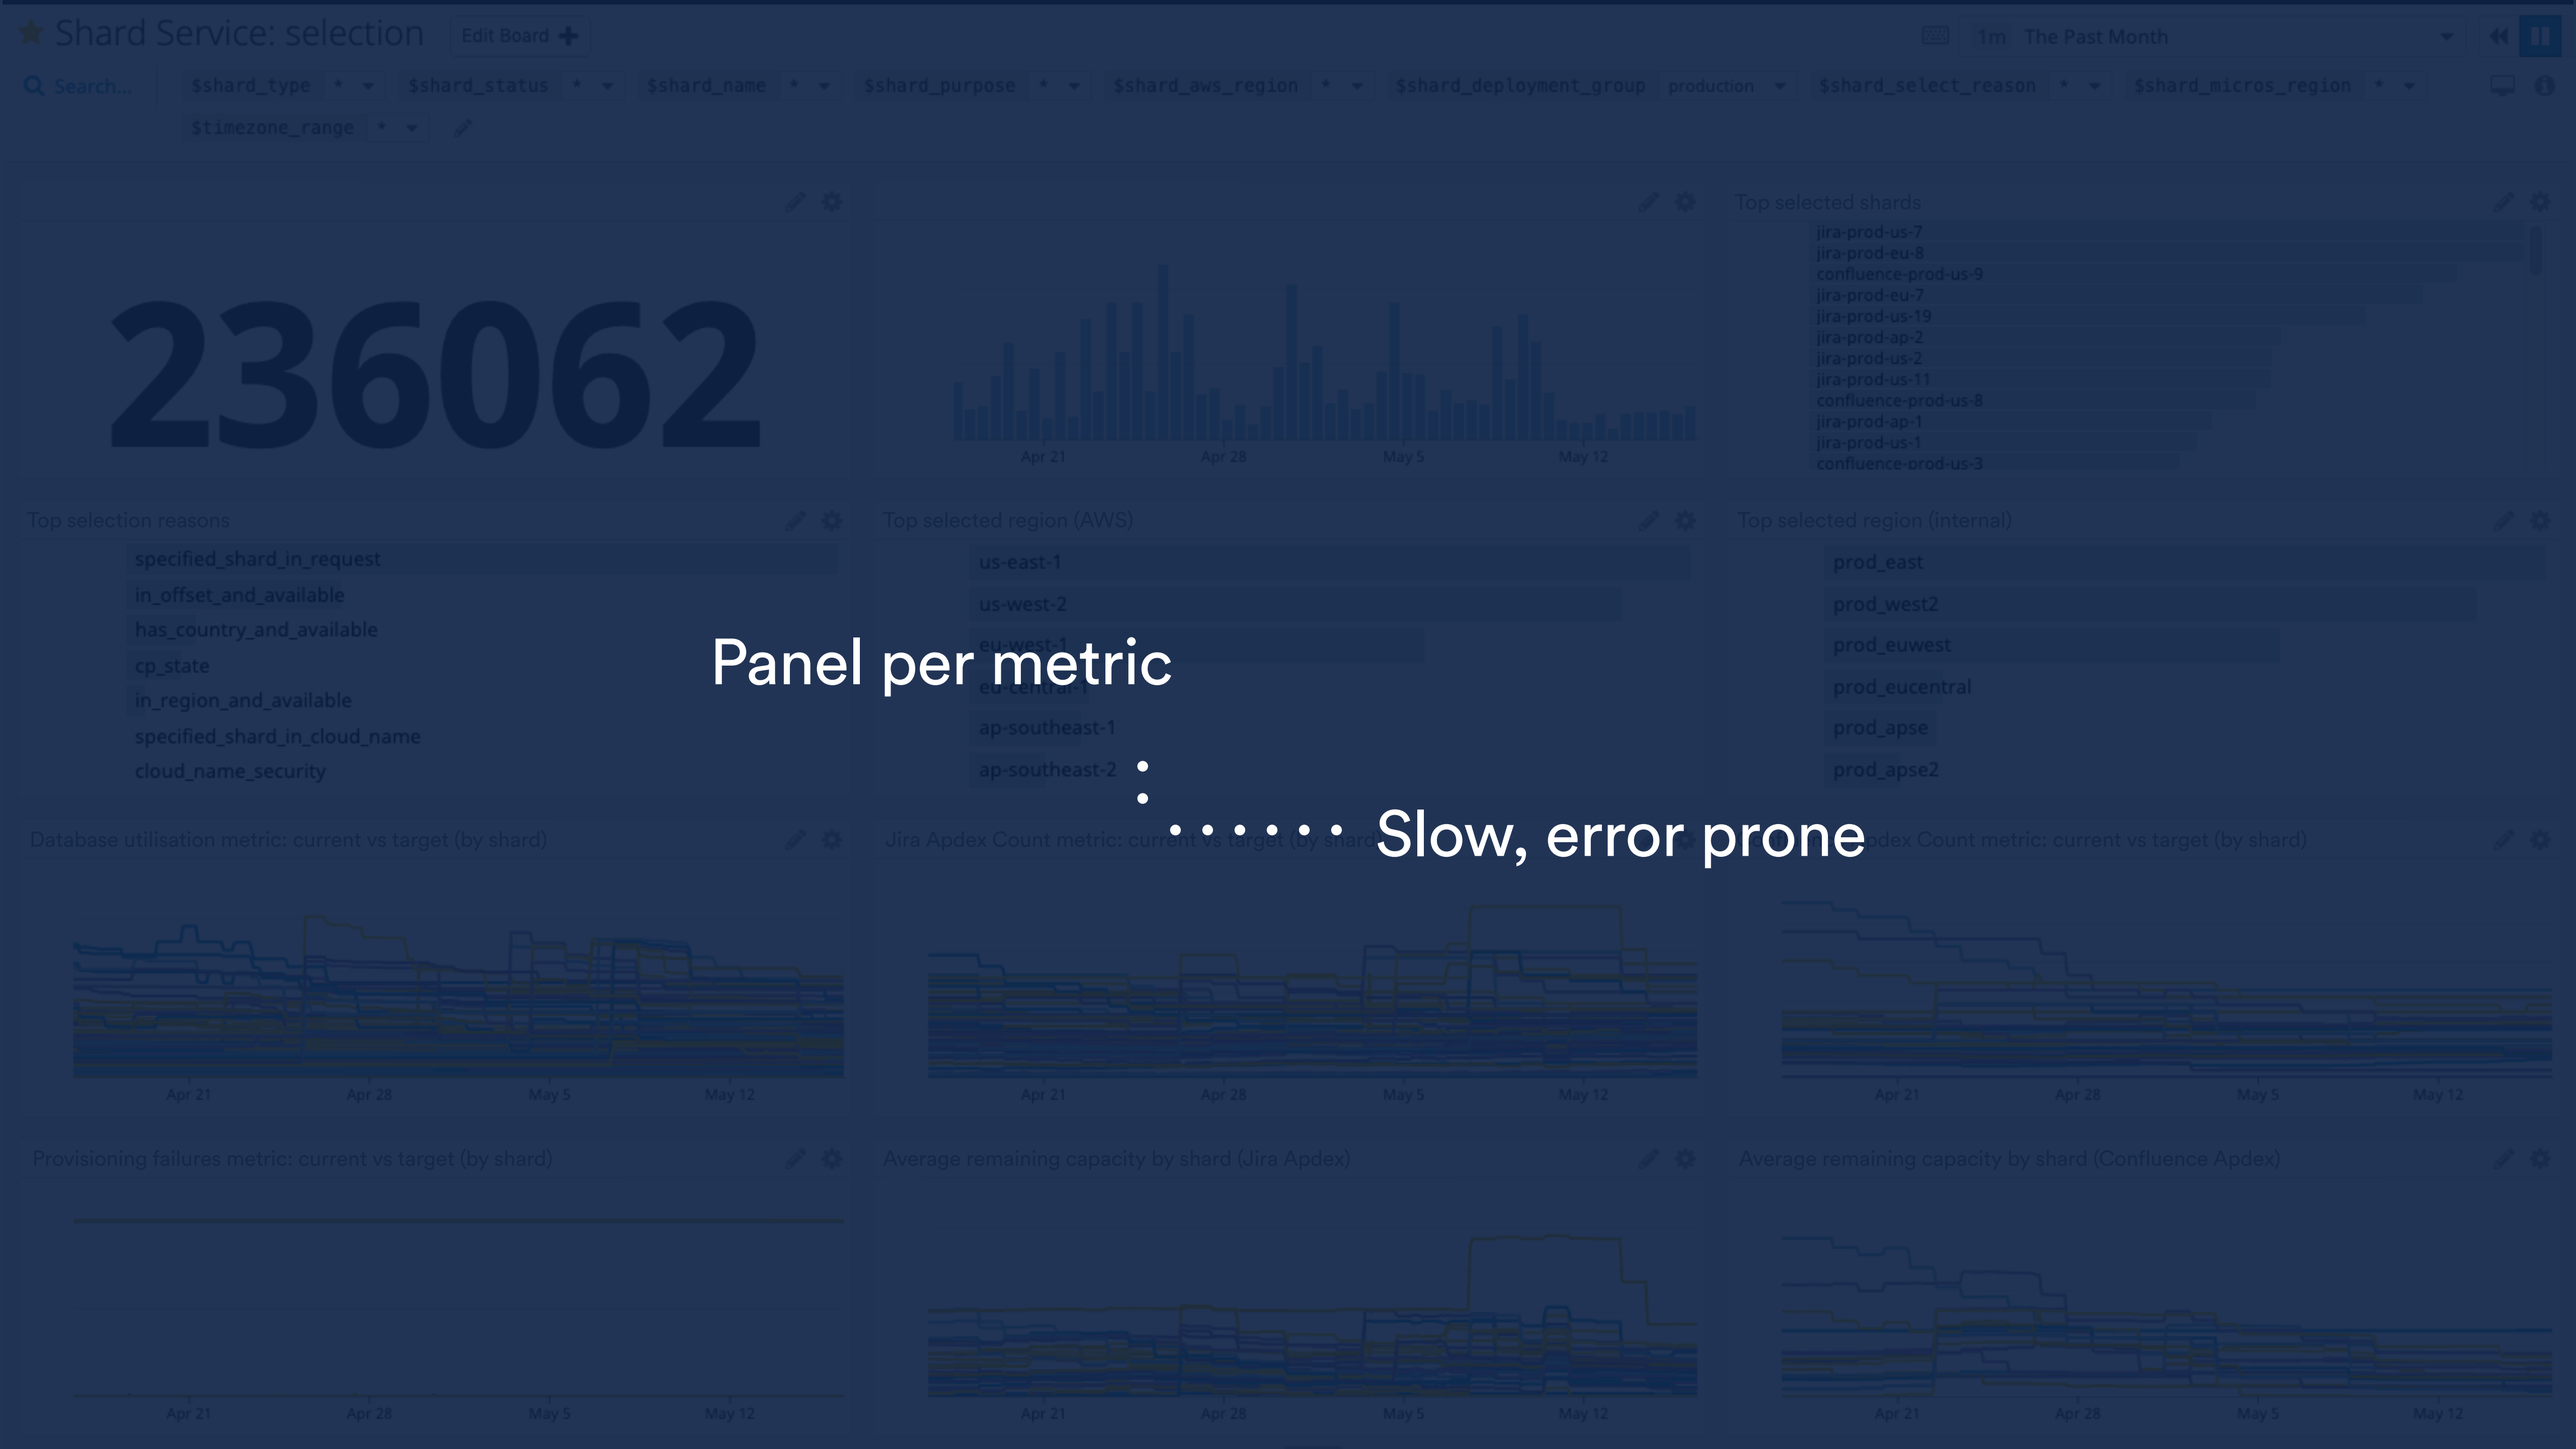
Shard Service: selection
Edit Board
1m
The Past Month
Search...
$shard_type
$shard_status
$shard_name
$shard_purpose
$shard_aws_region
$shard_deployment_group
production
$shard_select_reason
$shard_micros_region
$timezone_range
236062
Apr 21
Apr 28
May 5
May 12
Top selected shards
jira-prod-us-7
jira-prod-eu-8
confluence-prod-us-9
jira-prod-eu-7
jira-prod-us-19
jira-prod-ap-2
jira-prod-us-2
jira-prod-us-11
confluence-prod-us-8
jira-prod-ap-1
jira-prod-us-1
confluence-prod-us-3
Top selection reasons
specified_shard_in_request
in_offset_and_available
has_country_and_available
cp_state
in_region_and_available
specified_shard_in_cloud_name
cloud_name_security
Top selected region (AWS)
us-east-1
us-west-2
ap-southeast-1
ap-southeast-2
Top selected region (internal)
prod_east
prod_west2
prod_euwest
prod_eucentral
prod_apse
prod_apse2
Panel per metric
Slow, error prone
Database utilisation metric: current vs target (by shard)
Jira Apdex Count metric: current vs target (by shard)
Provisioning failures metric: current vs target (by shard)
Average remaining capacity by shard (Jira Apdex)
Average remaining capacity by shard (Confluence Apdex)

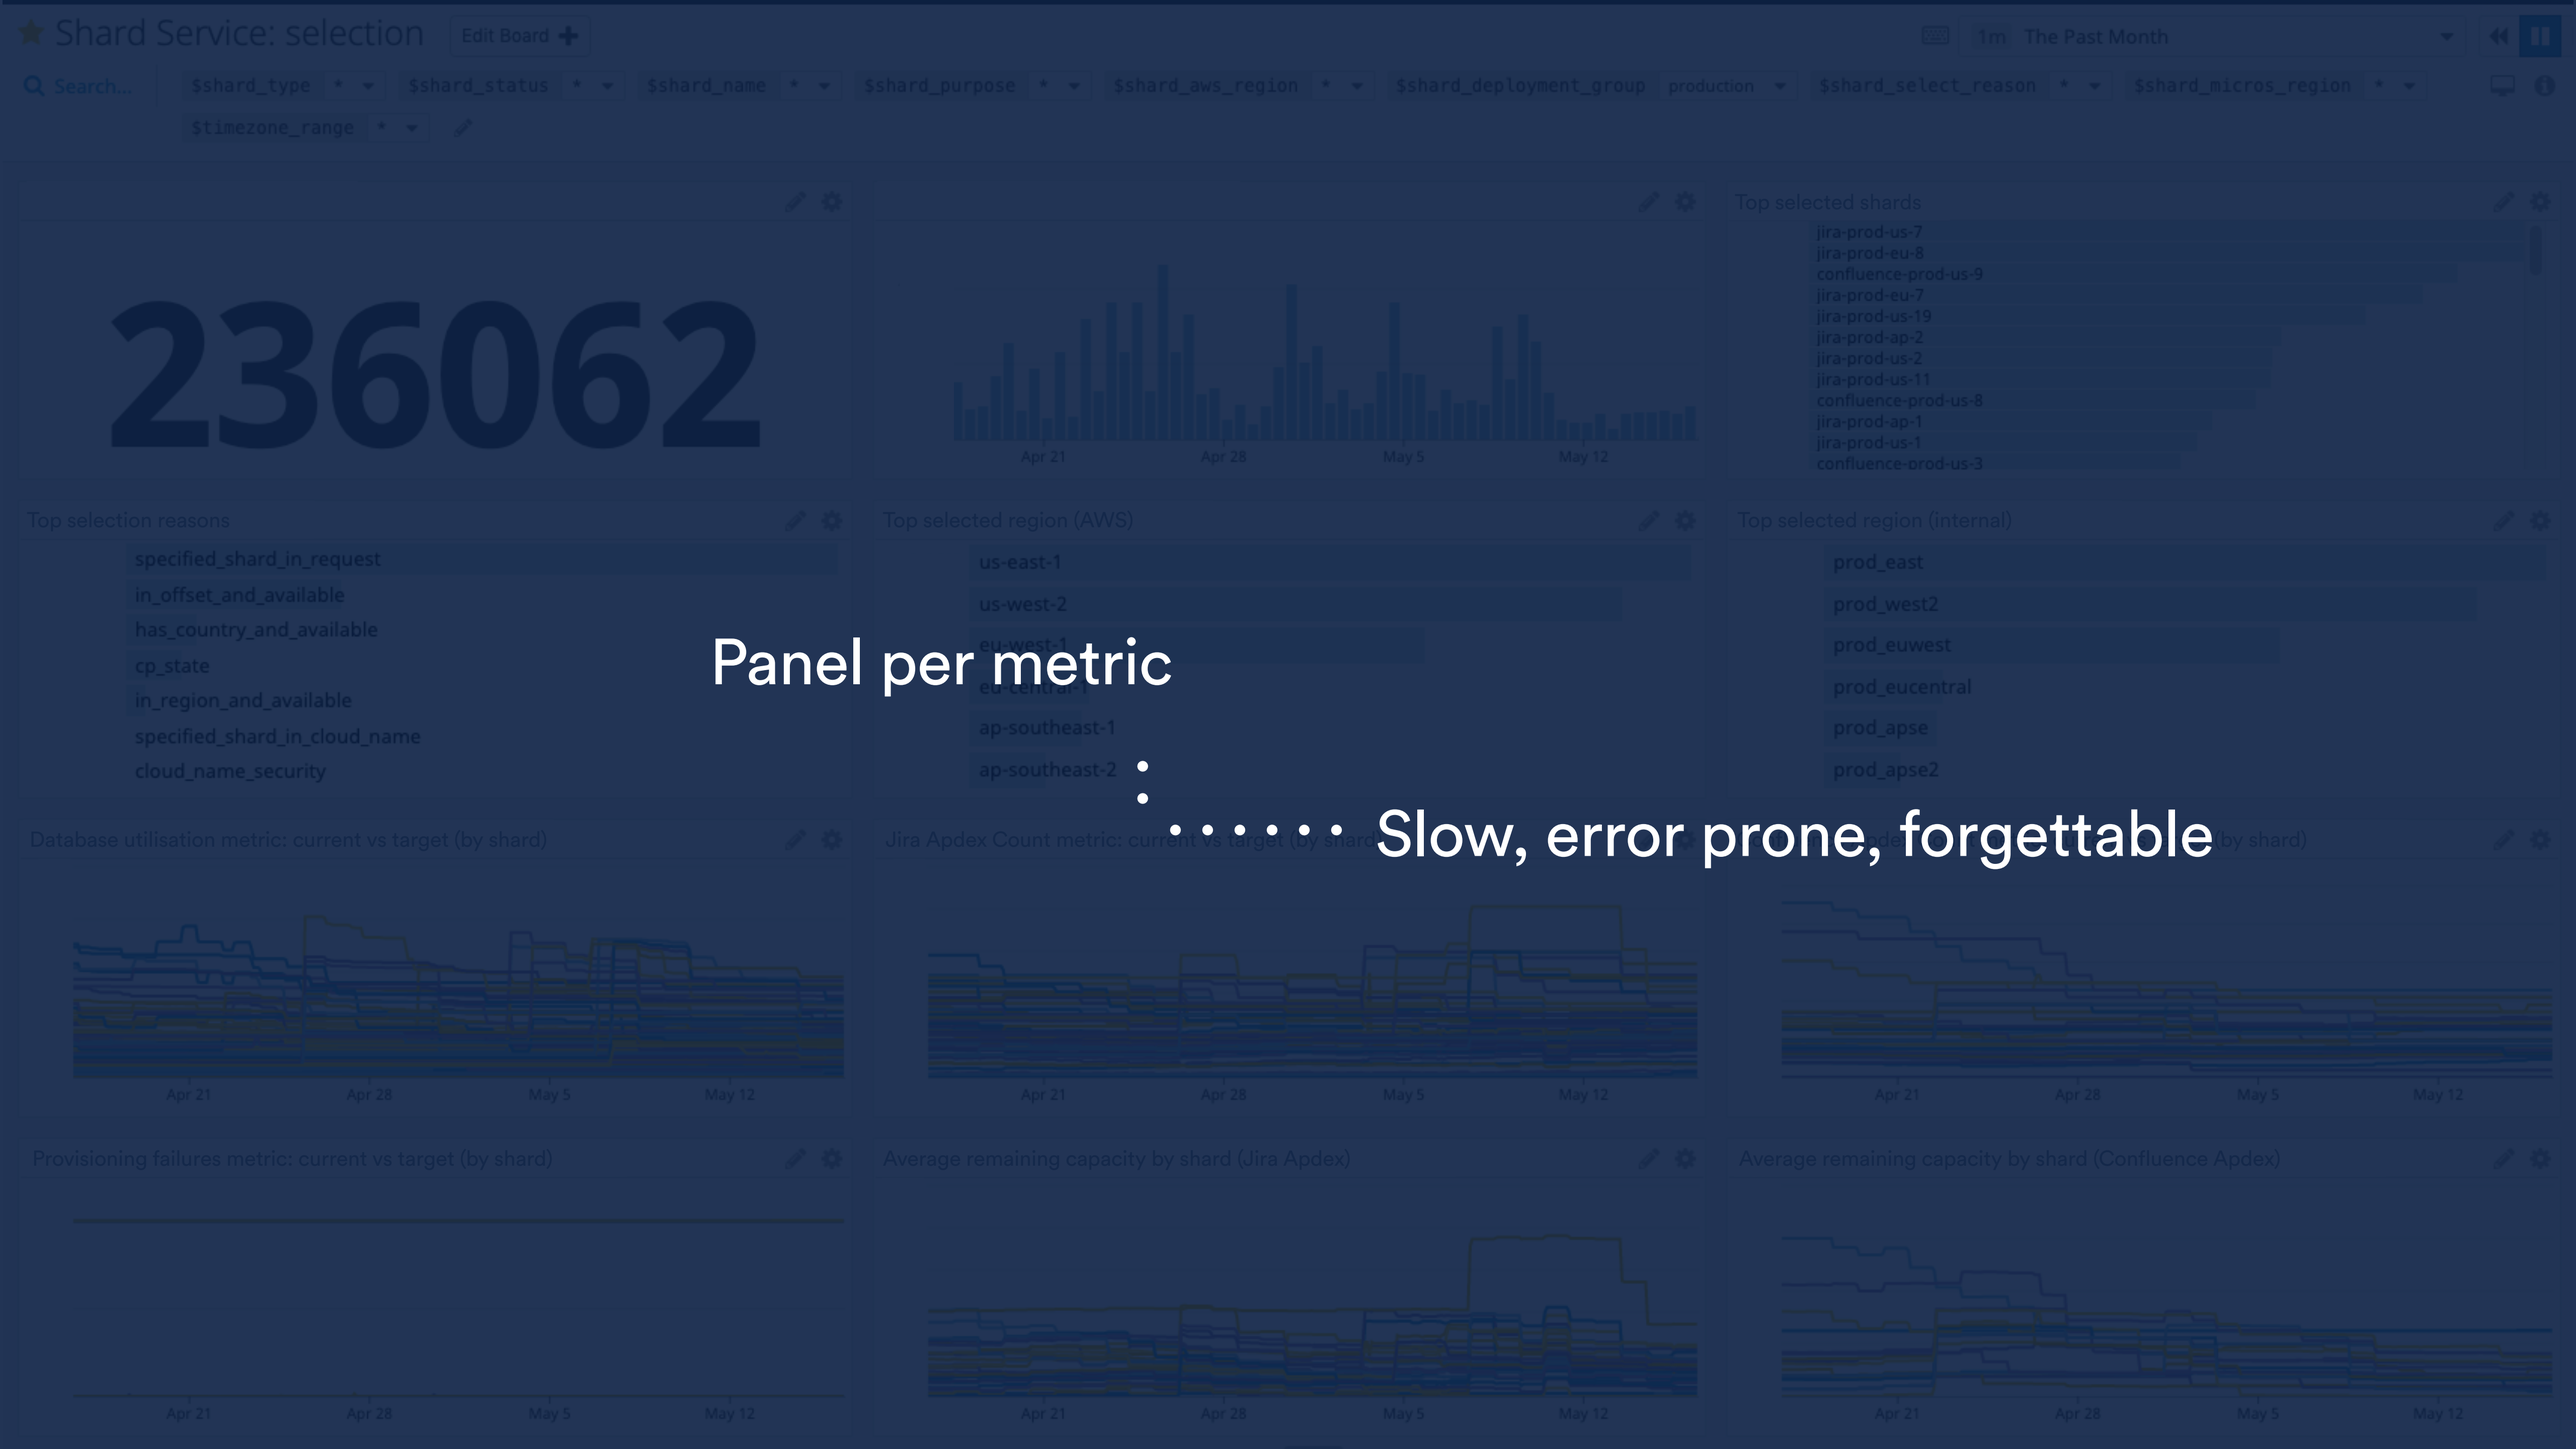
Shard Service: selection
Edit Board
1m The Past Month
Search...
$shard_type
$shard_status
$shard_name
$shard_purpose
$shard_aws_region
$shard_deployment_group
production
$shard_select_reason
$shard_micros_region
$timezone_range
236062
Apr 21
Apr 28
May 5
May 12
Top selected shards
jira-prod-us-7
jira-prod-eu-8
confluence-prod-us-9
jira-prod-eu-7
jira-prod-us-19
jira-prod-ap-2
jira-prod-us-2
jira-prod-us-11
confluence-prod-us-8
jira-prod-ap-1
jira-prod-us-1
confluence-prod-us-3
Top selection reasons
specified_shard_in_request
in_offset_and_available
has_country_and_available
cp_state
in_region_and_available
specified_shard_in_cloud_name
cloud_name_security
Top selected region (AWS)
us-east-1
us-west-2
ap-southeast-1
ap-southeast-2
Top selected region (internal)
prod_east
prod_west2
prod_euwest
prod_eucentral
prod_apse
prod_apse2
Panel per metric
Slow, error prone, forgettable
Database utilisation metric: current vs target (by shard)
Jira Apdex Count metric: current vs target (by shard)
Provisioning failures metric: current vs target (by shard)
Average remaining capacity by shard (Jira Apdex)
Average remaining capacity by shard (Confluence Apdex)

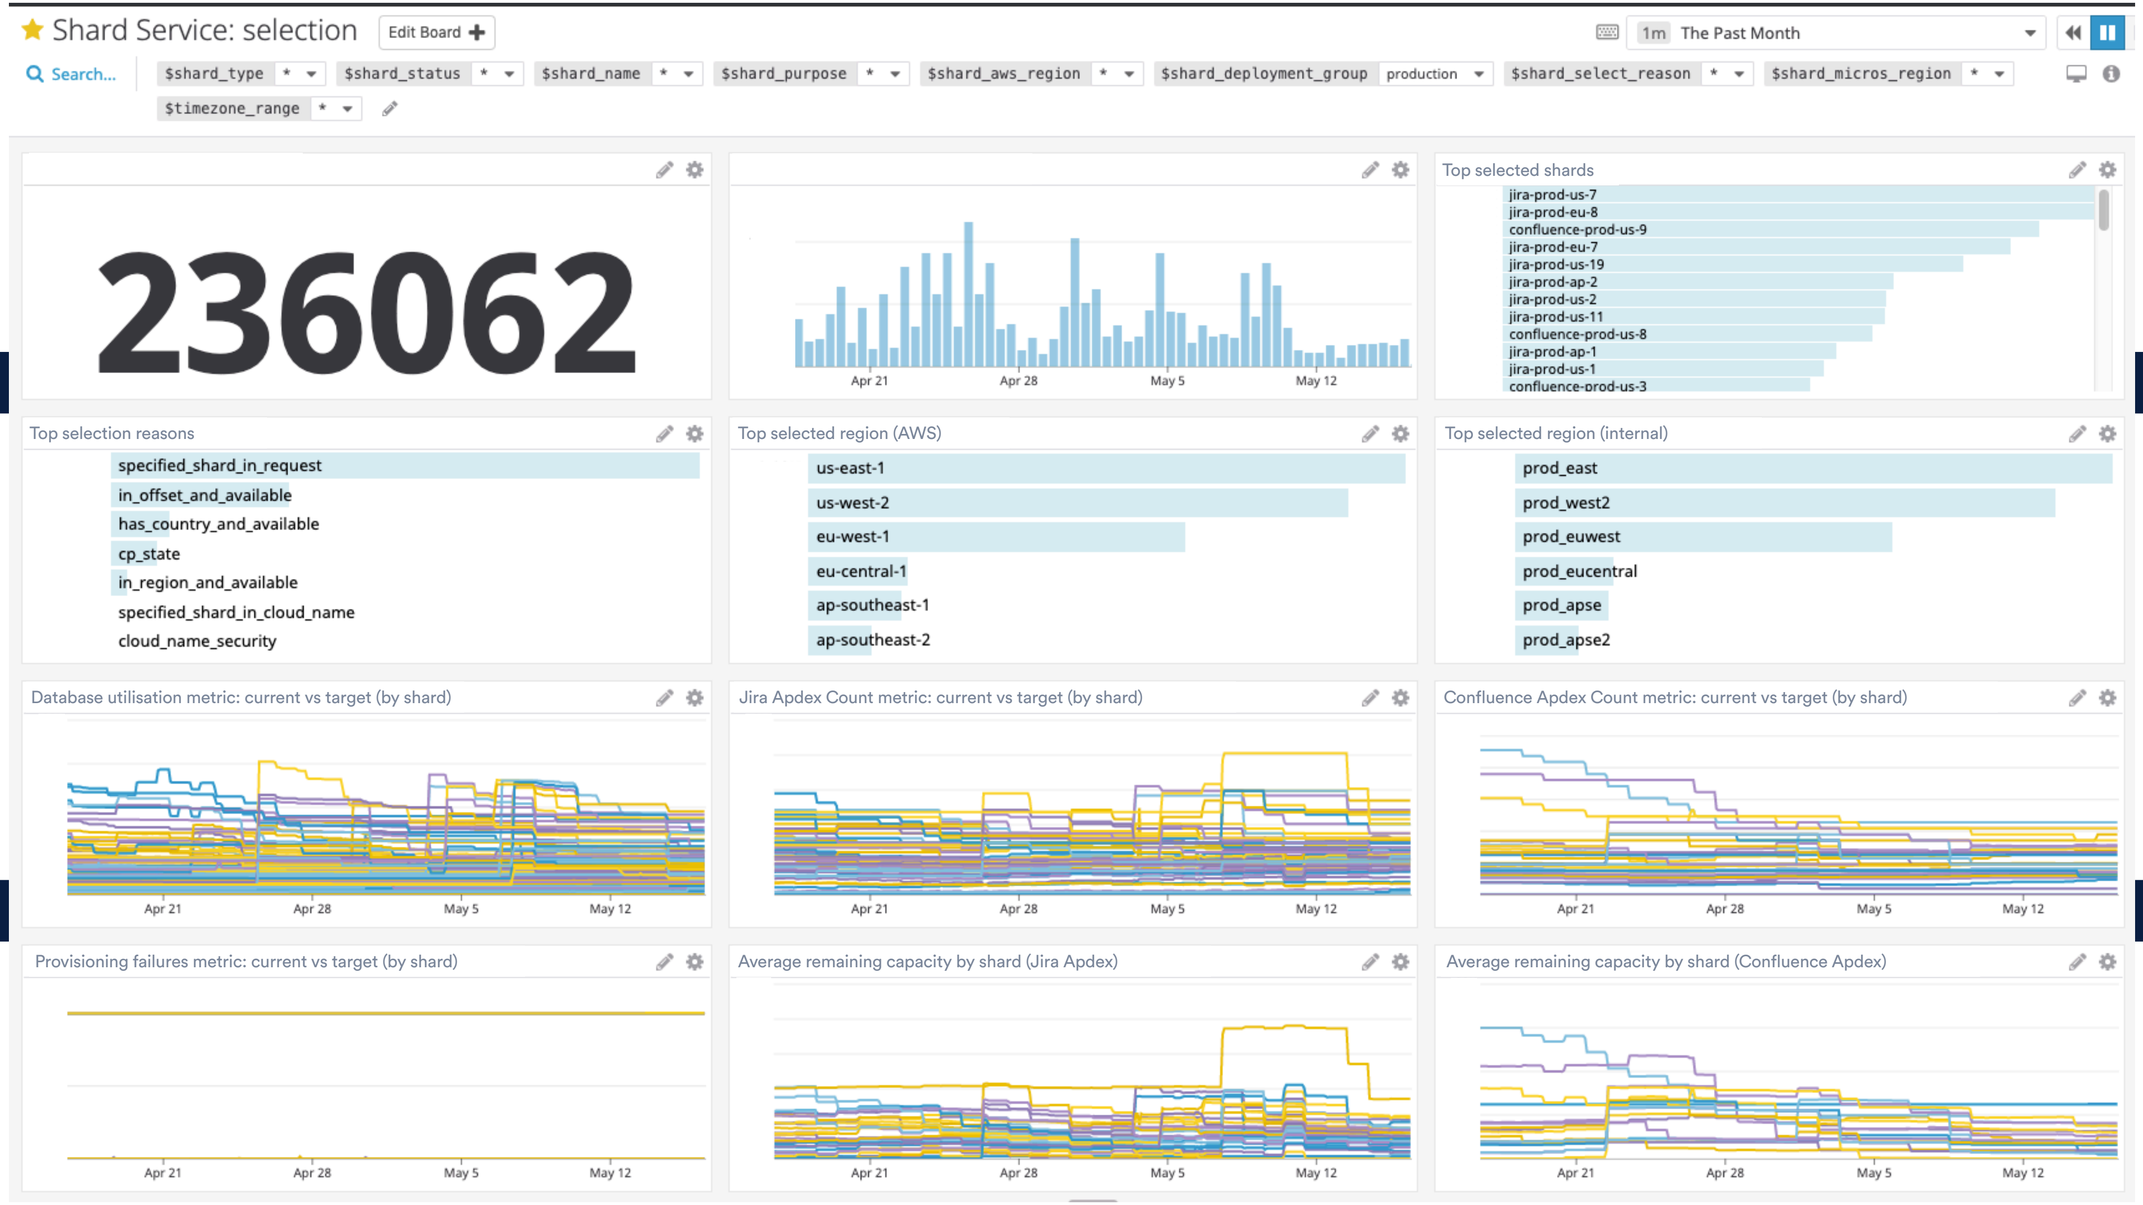
Shard Service: selection
Edit Board
1m The Past Month
Search...
$shard_type *
$shard_status *
$shard_name *
$shard_purpose *
$shard_aws_region *
$shard_deployment_group production
$shard_select_reason *
$shard_micros_region *
$timezone_range *
236062
Apr 21
Apr 28
May 5
May 12
Top selected shards
jira-prod-us-7
jira-prod-eu-8
confluence-prod-us-9
jira-prod-eu-7
jira-prod-us-19
jira-prod-ap-2
jira-prod-us-2
jira-prod-us-11
confluence-prod-us-8
jira-prod-ap-1
jira-prod-us-1
confluence-prod-us-3
Top selection reasons
specified_shard_in_request
in_offset_and_available
has_country_and_available
cp_state
in_region_and_available
specified_shard_in_cloud_name
cloud_name_security
Top selected region (AWS)
us-east-1
us-west-2
eu-west-1
eu-central-1
ap-southeast-1
ap-southeast-2
Top selected region (internal)
prod_east
prod_west2
prod_euwest
prod_eucentral
prod_apse
prod_apse2
Database utilisation metric: current vs target (by shard)
Jira Apdex Count metric: current vs target (by shard)
Confluence Apdex Count metric: current vs target (by shard)
Provisioning failures metric: current vs target (by shard)
Average remaining capacity by shard (Jira Apdex)
Average remaining capacity by shard (Confluence Apdex)

Shard Service repository

{ Application }

{ Test }

{ 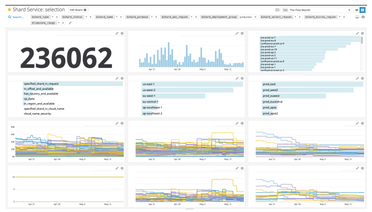 }

Shard Service repository

{ Application }

{ Test }

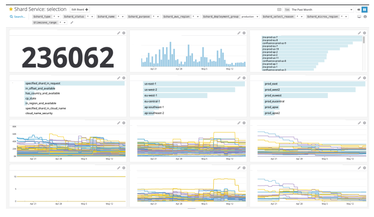

Dashboard tool

Marge's Service

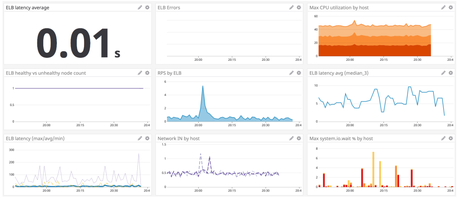

Homer's Service

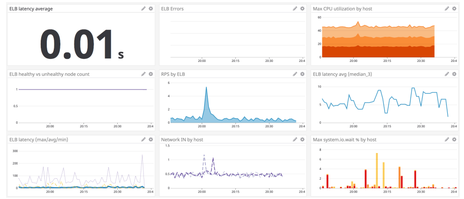

Shard Service

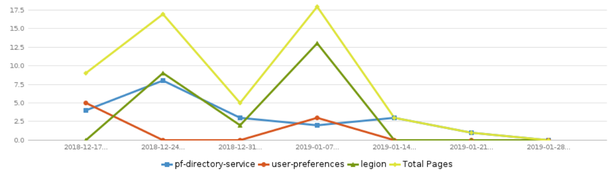

Shard Service repository

{ Application }

{ Test }

{ 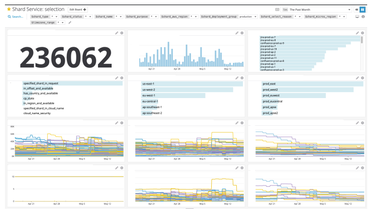 }

Dashboard tool

Marge's Service

Homer's Service

Shard Service

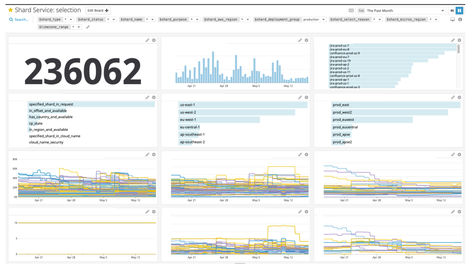

Operational resources as code

Front of mind

Version control

Discoverable

Operational resources as code

Front of mind

Version control

Discoverable

Operational resources as code

Front of mind

Version control

Discoverable

Going a step further...

286

286

Dashboards

JIRA SHARD DASHBOARD (SUBSET)

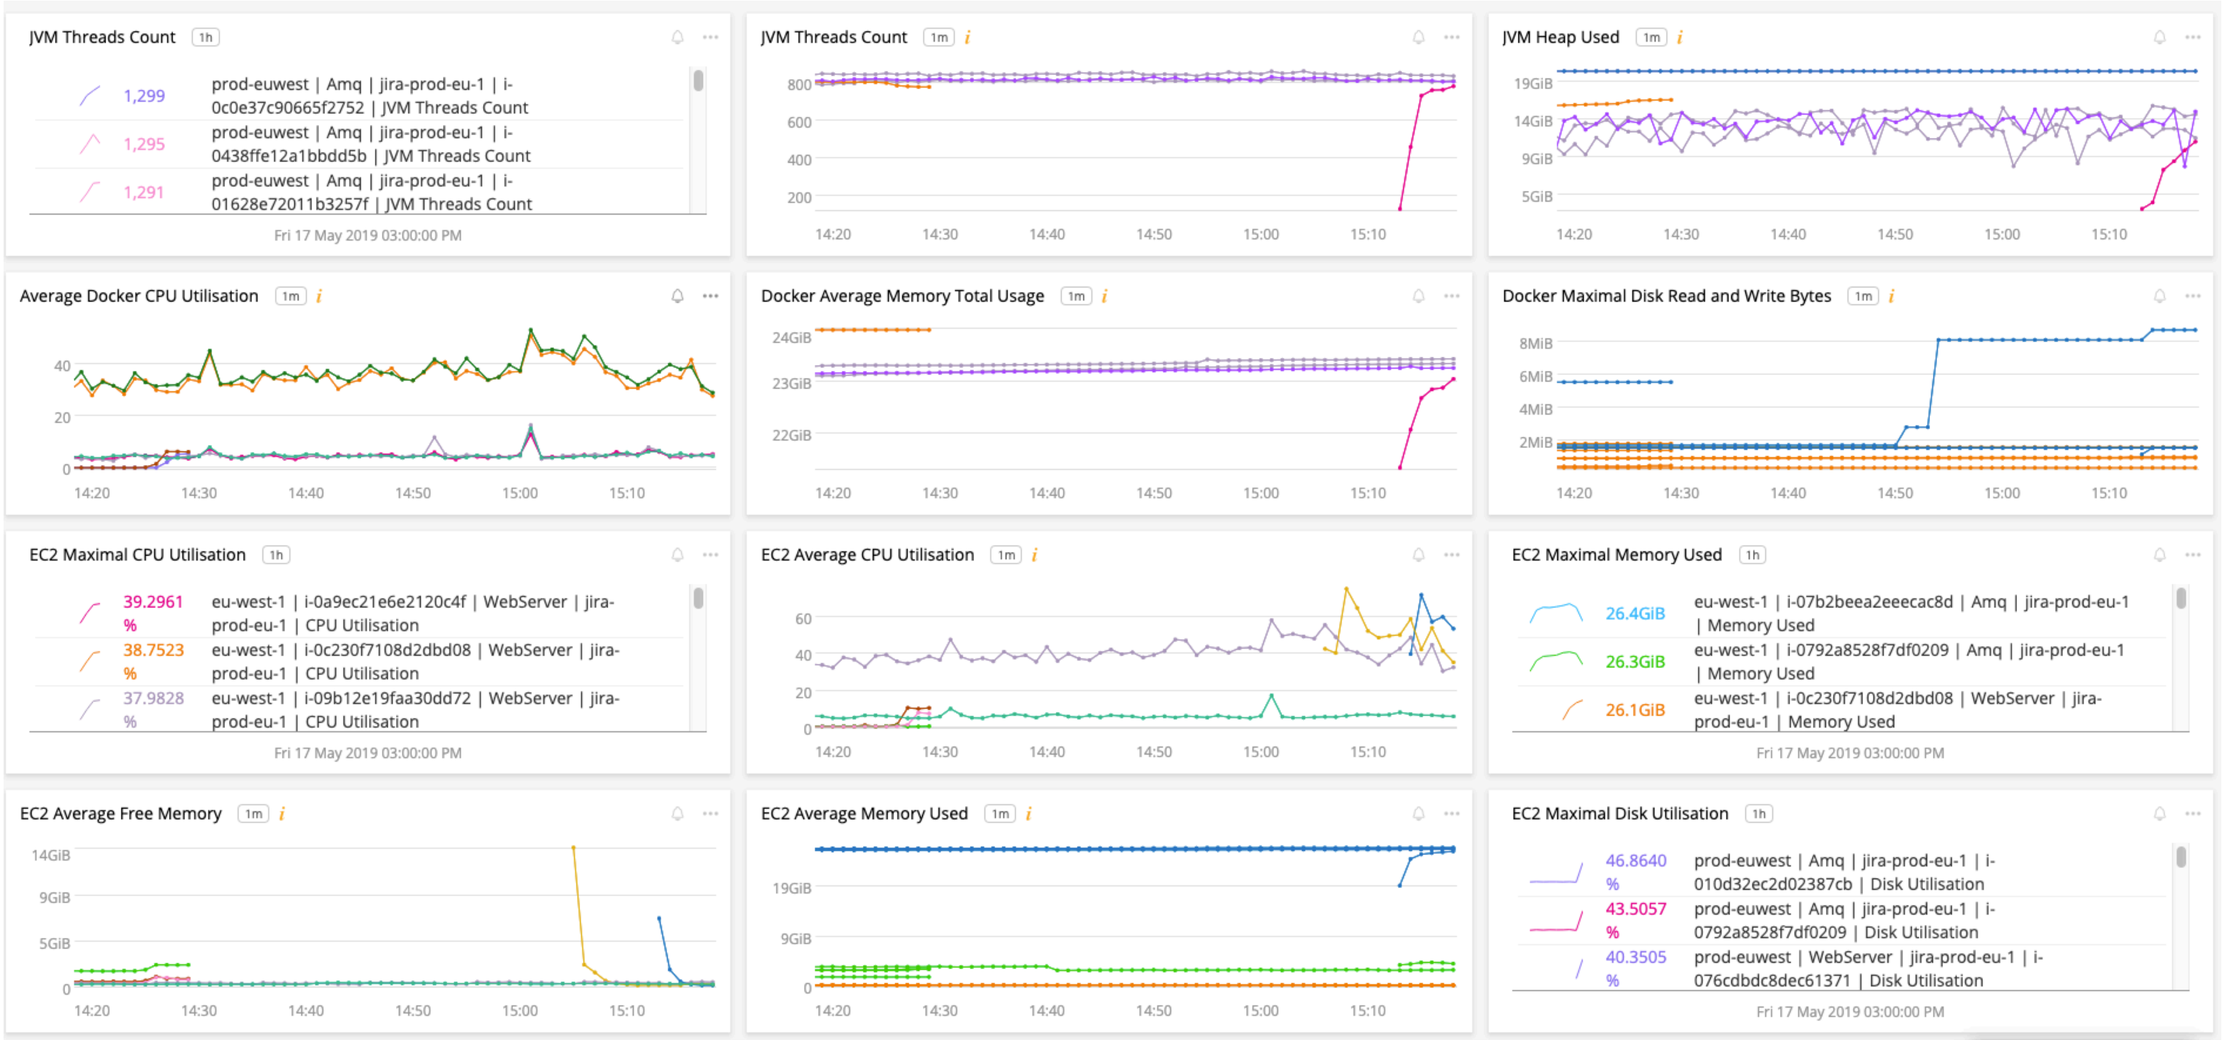

Cue, templates

This can be solved at the platform level

SERGEJS SINICA, ATLASSIAN SENIOR DEVELOPER

91

91

respondents who owned 1 to 10+ services

83.5%

83.5%

maintained operational resources through the UI

67%

67%

kept their team up to date via “tribal knowledge”

23.1%

23.1%

satisfied with their existing process

“We've been saying for eons that we should put our monitors and dashboards into code, but the task is too big to start so we don't do it. Over time the job just gets bigger and bigger and less likely to get done :P”

ATLASSIAN DEVELOPER

Introducing, Sauron

Introducing, Sauron

The “all seeing eye” for dashboards & monitors

Sauron

Dashboard tool

Marge's Service

Homer's Service

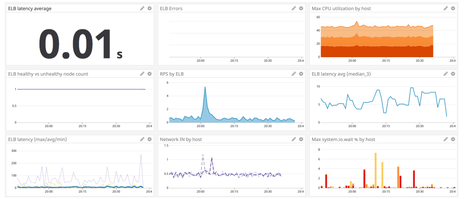

Shard Service

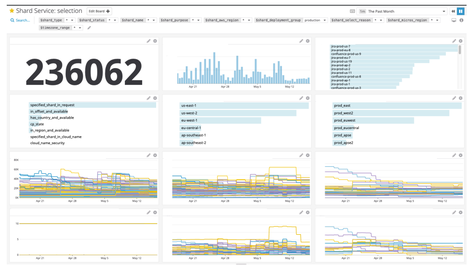

Export my dashboard!

Sauron

Dashboard tool

Marge's Service

Homer's Service

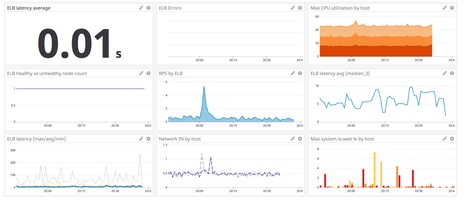

Shard Service

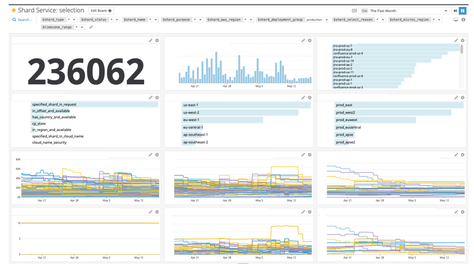

Export my dashboard!

Sauron

Dashboard tool

Marge's Service

Homer's Service

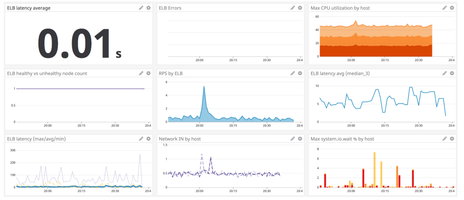

Shard Service

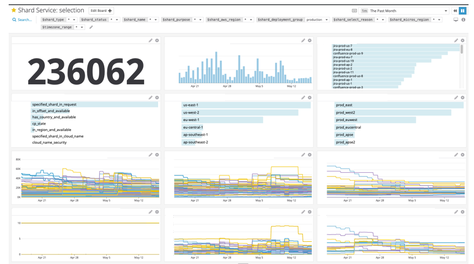

Export my dashboard!

Sauron

Dashboard tool

Marge's Service

Homer's Service

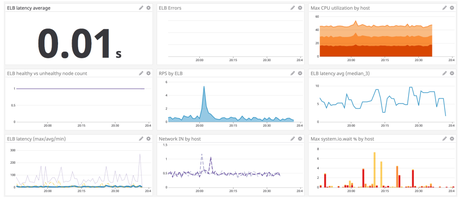

Shard Service

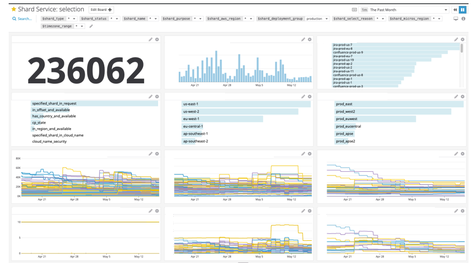

Export my dashboard!

Sauron

Dashboard tool

Marge's Service

Homer's Service

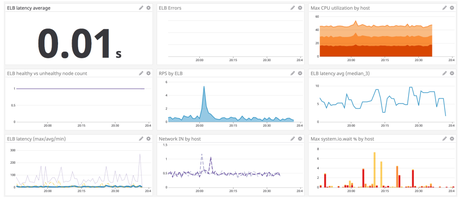

Shard Service

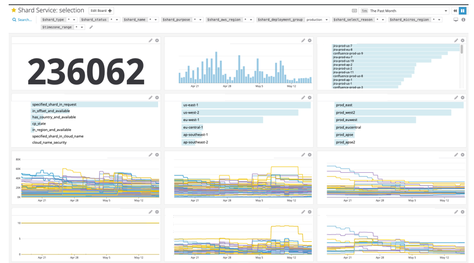

Dashboard tool

Sauron

Marge's Service

Homer's Service

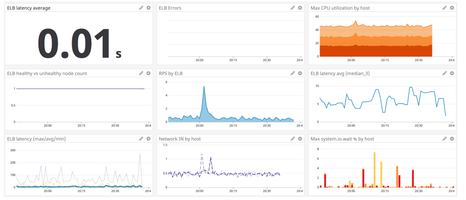

Shard Service

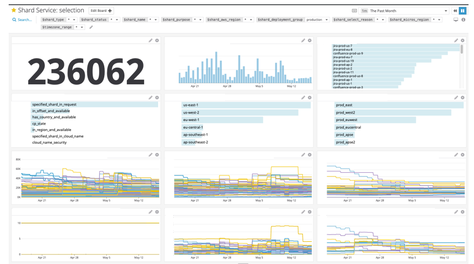

Sauron

Application repository

{  }

Dashboard tool

Marge's Service

Homer's Service

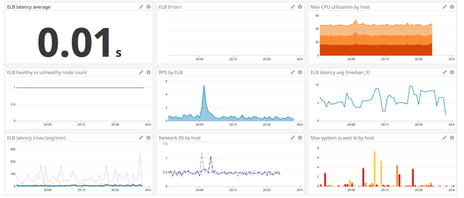

Shard Service

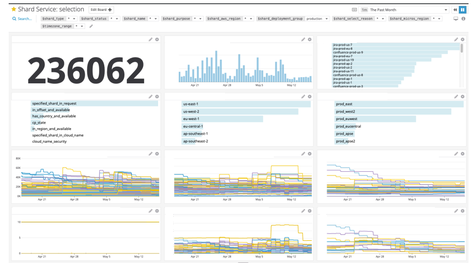

Sauron

Application repository

{  }

Dashboard tool

Marge's Service

Homer's Service

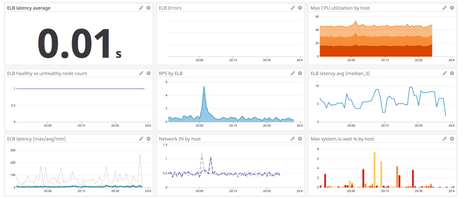

Shard Service

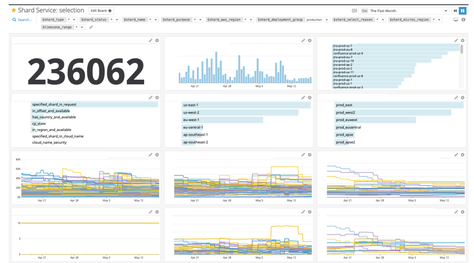

Sauron

Application repository

Dashboard tool

Marge's Service

Homer's Service

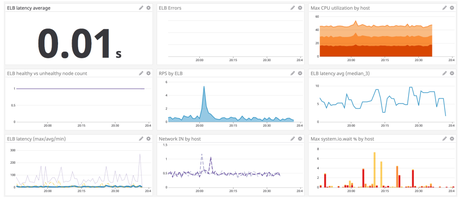

Shard Service

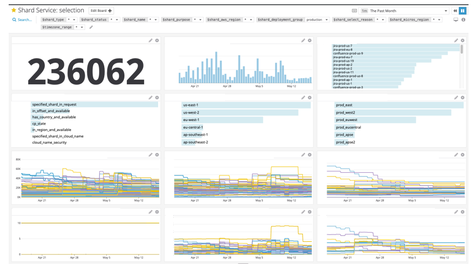

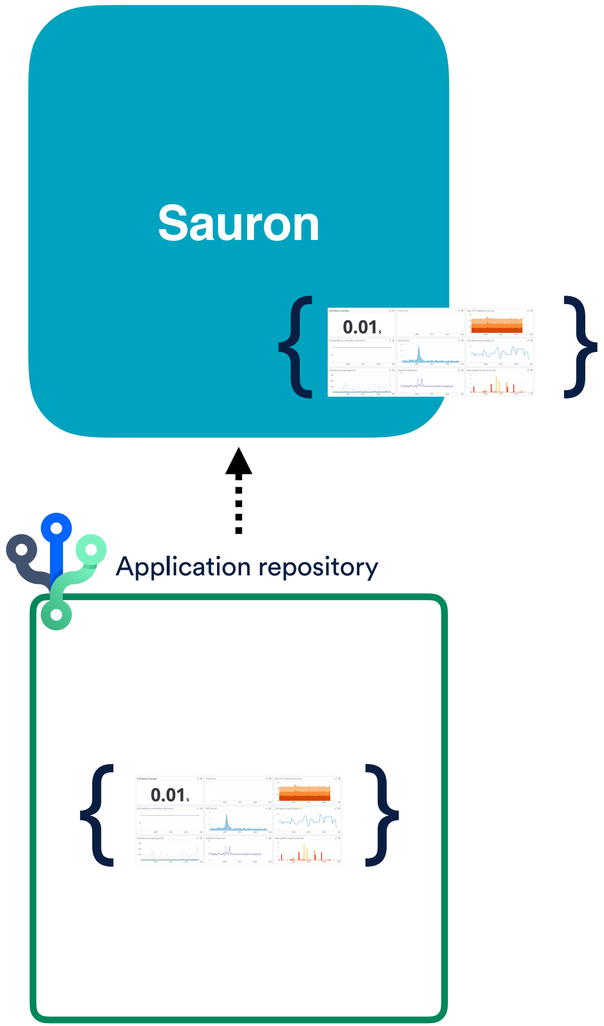

Dashboard tool

Marge's Service

Homer's Service

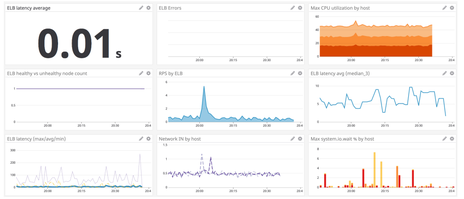

Shard Service

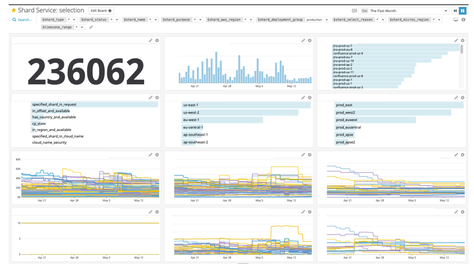

Sauron

Application repository

{  }

Dashboard tool

Marge's Service

Homer's Service

Shard Service

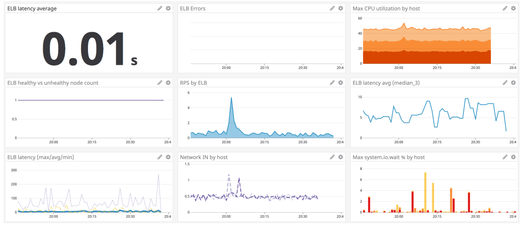

shard-service operations

Monitors

- Metric is overridden for too long
- Metric is overridden for too long
- Shard Service (prod): Errors
- Shard Service (prod): Incorrect Selections detected
- Shard Service (prod): Warnings
- Shard Service (prod): remaining capacity low by shardtype, zonerange
- Shard Service (prod): remaining capacity low for {{shardtype.name}} {{shardawsregion.name}}
- Shard Service (prod): shard capacity below threshold {{shardtype.name}} {{shardawsregion.name}}

Dashboards

- Equalizers - Shard Service Utilization Over Time
- Shard Service: selection

Screenboards

- Shard Service: capacity

> 50

> 50

Services adopted Sauron

How can you...

How can you…

Help your team keep up to date with change?

Define operational resources in code

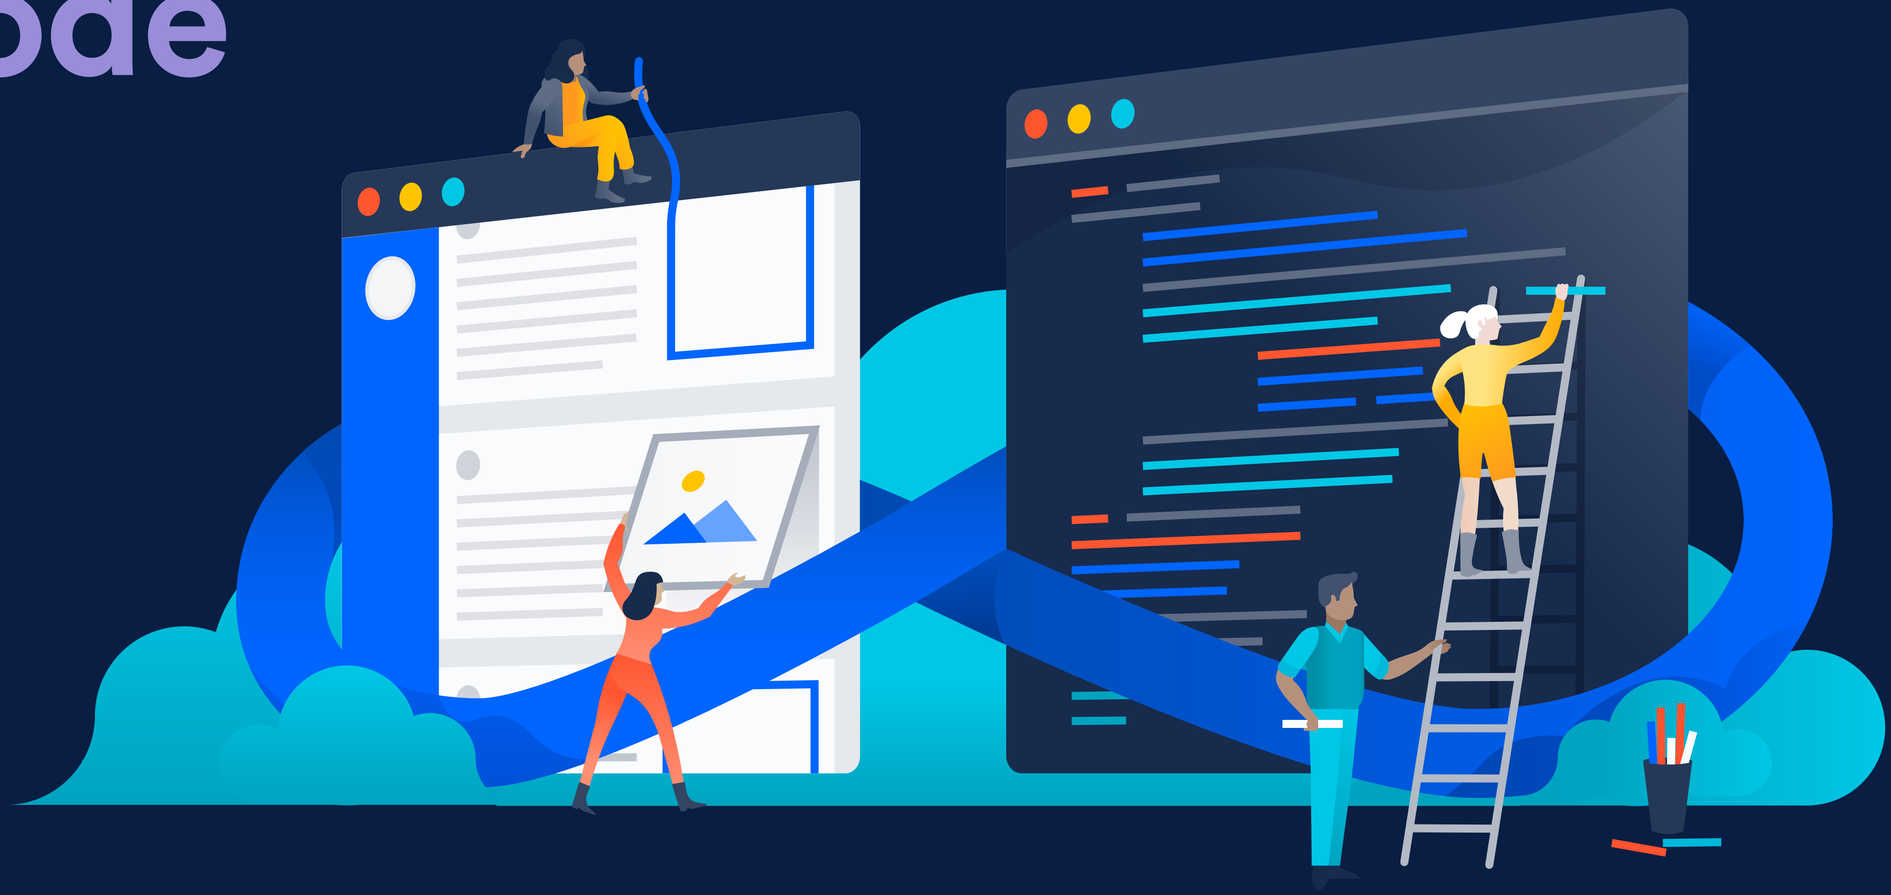

Agenda

Iterative... what?

Setting some context

Deciding what to measure

Verifying your metrics

Keeping up with change

Summary

Agenda

Iterative... what?

Setting some context

Deciding what to measure

Verifying your metrics

Keeping up with change

Summary

Select

Learn what questions you want to answer

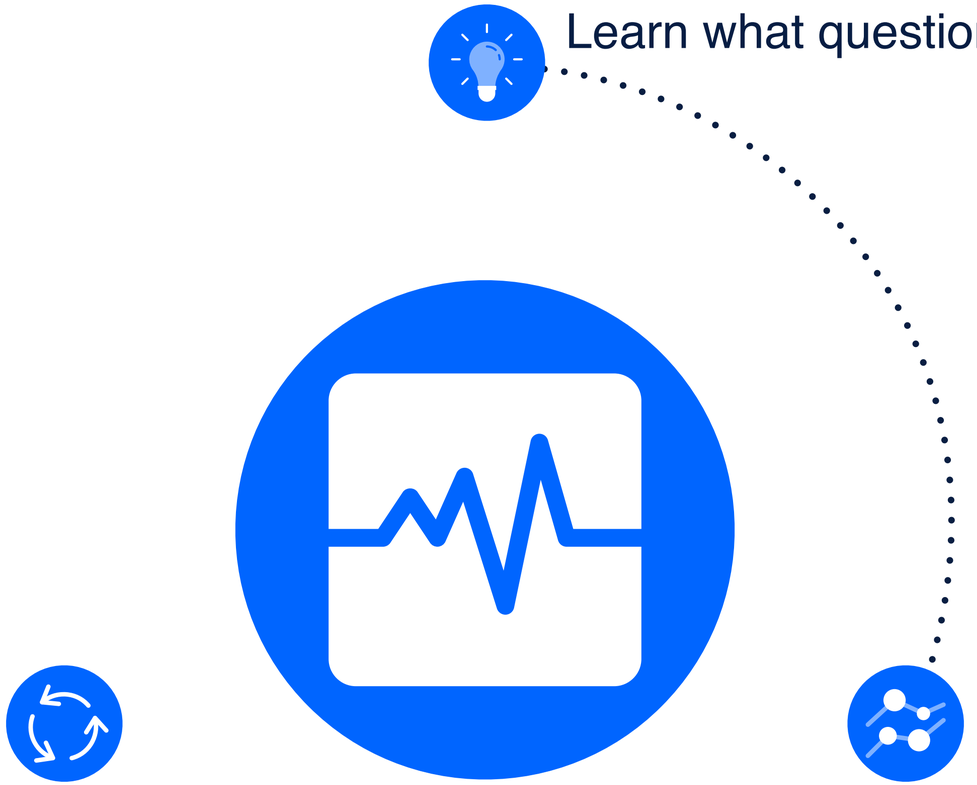

Select

Learn what questions you want to answer

Verify

Review, review, review

Select

Learn what questions you want to answer

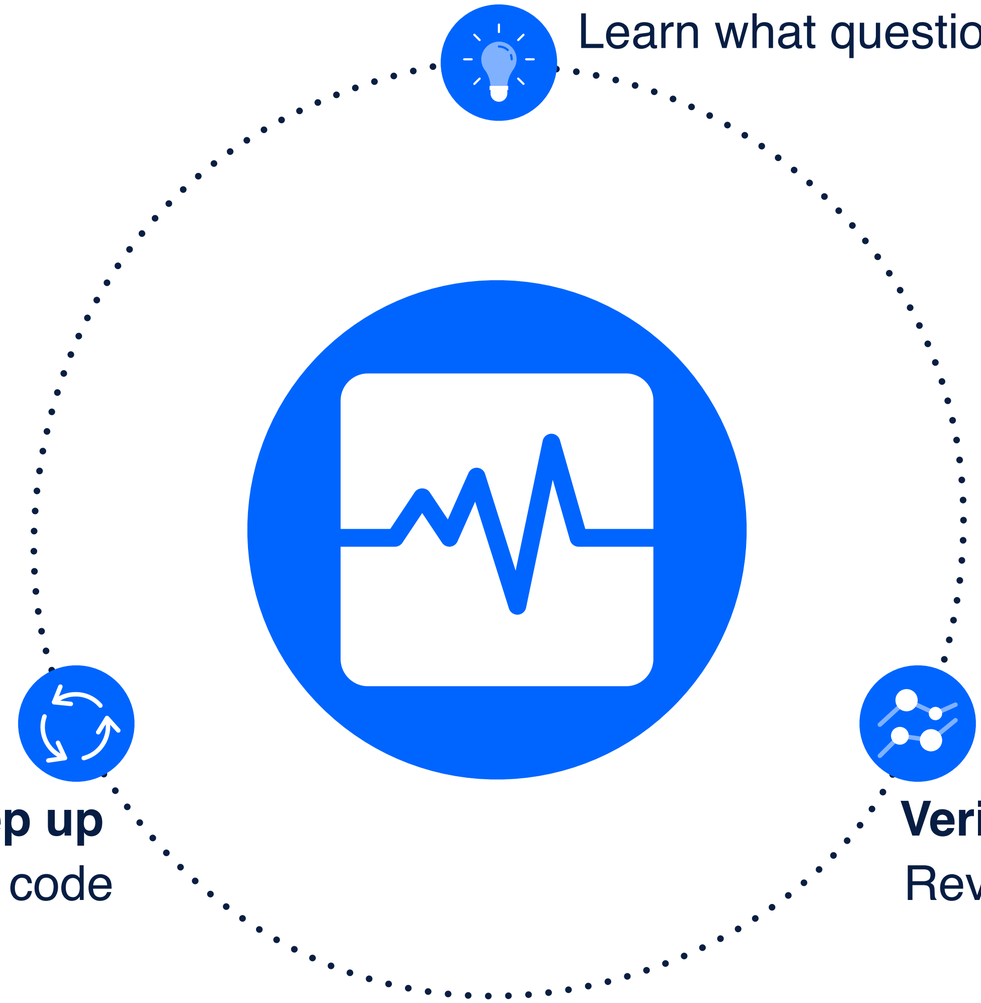

Verify

Review, review, review

Keep up

Define all the things in code

Thank you!