ANTI-FRAGILE CLOUD ARCHITECTURES

AGIM EMRULI - @AEMRULI - MIMACOM

“Antifragility is beyond resilience or robustness. The resilient resists shocks and stays the same; the antifragile gets better.”
Nasim Nicholas Taleb

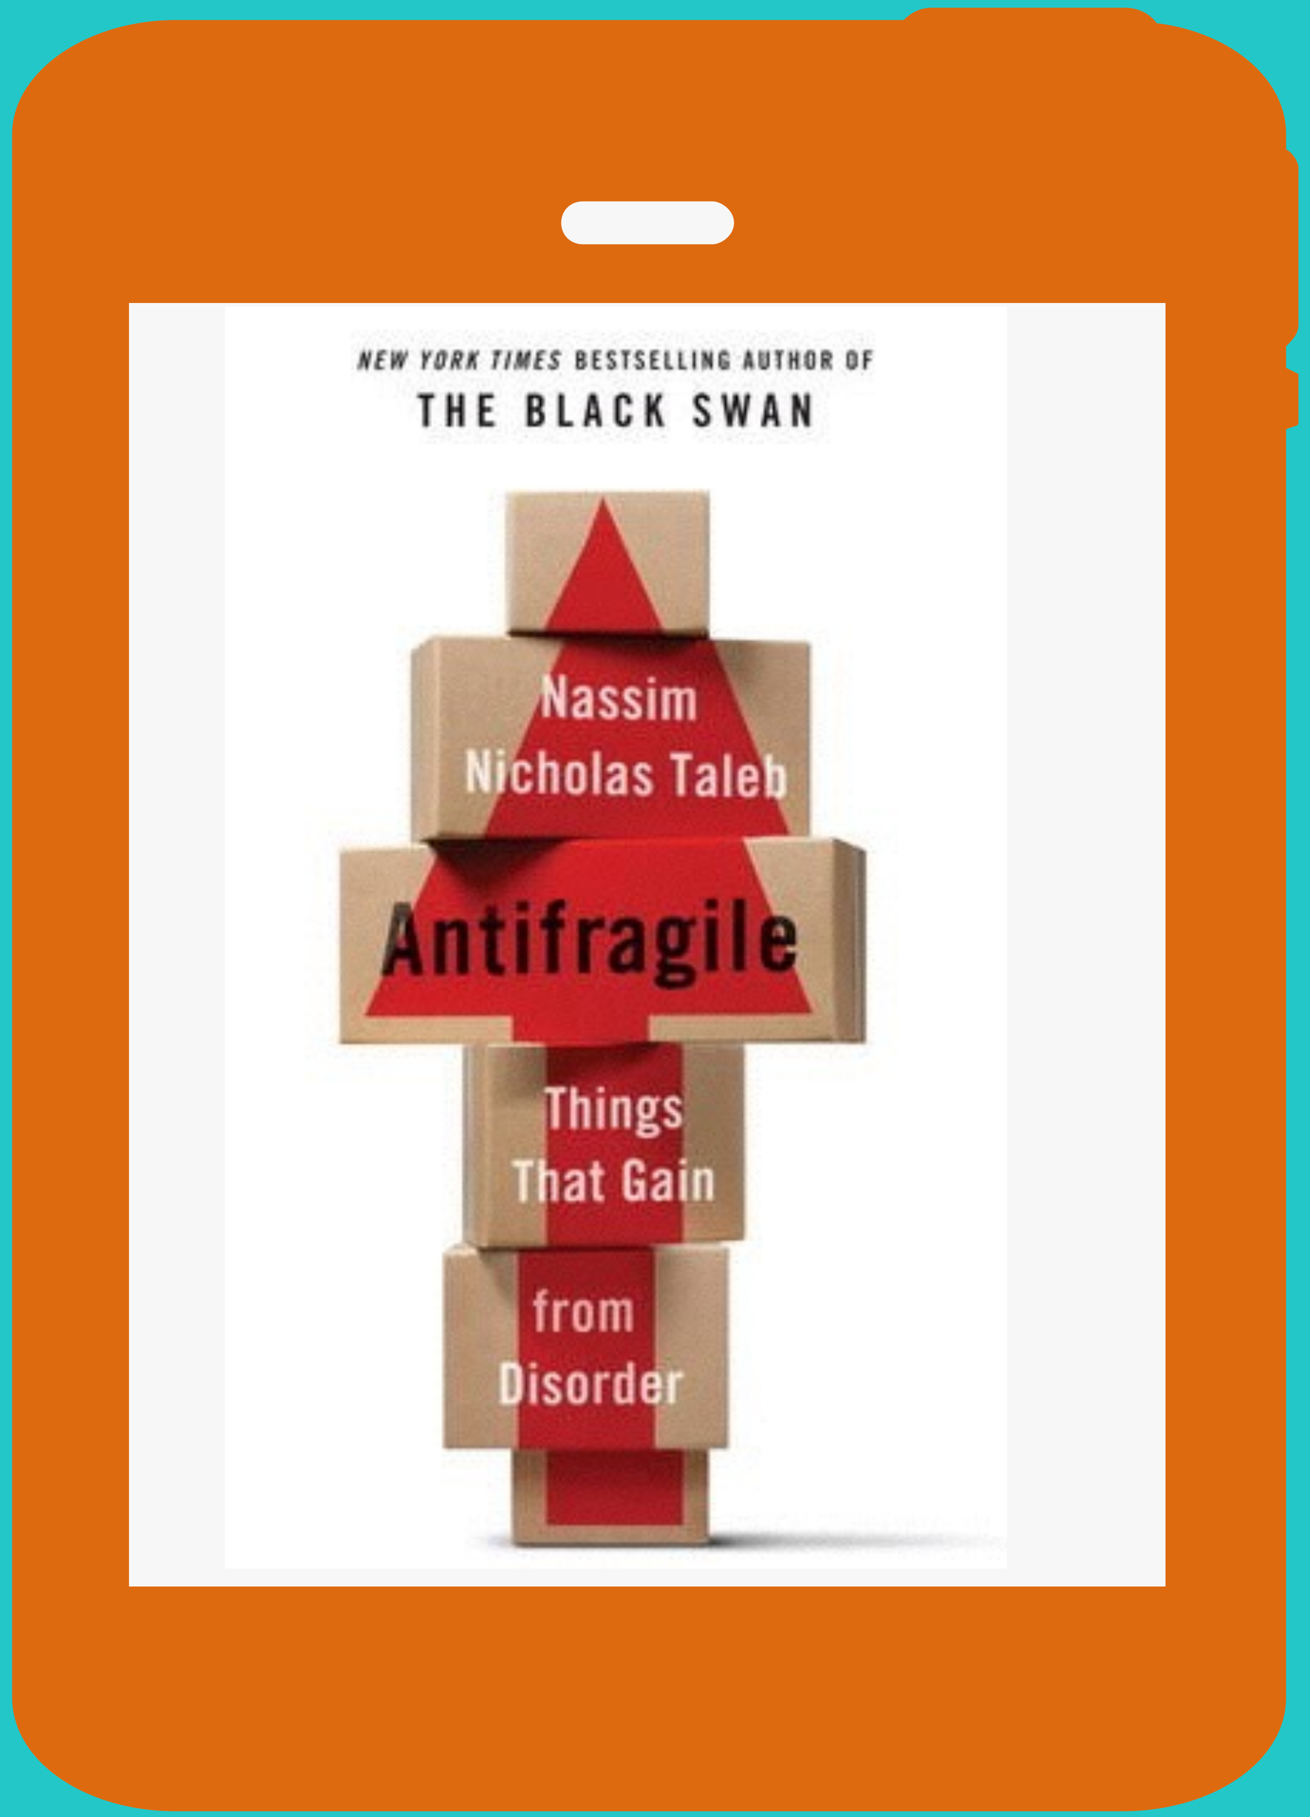

FRAGILE	ROBUST	ANTI-FRAGILE
NON-LINEAR (KONKAV)	LINEAR	NON-LINEAR (KONVEX)
POST-TRAUMATIC SYNDROM		POST-TRAUMATIC GROWTH
CENTRALIZED		DECENTRALIZED

FRAGILE

ROBUST

FRAGILE

ANTI-FRAGILE

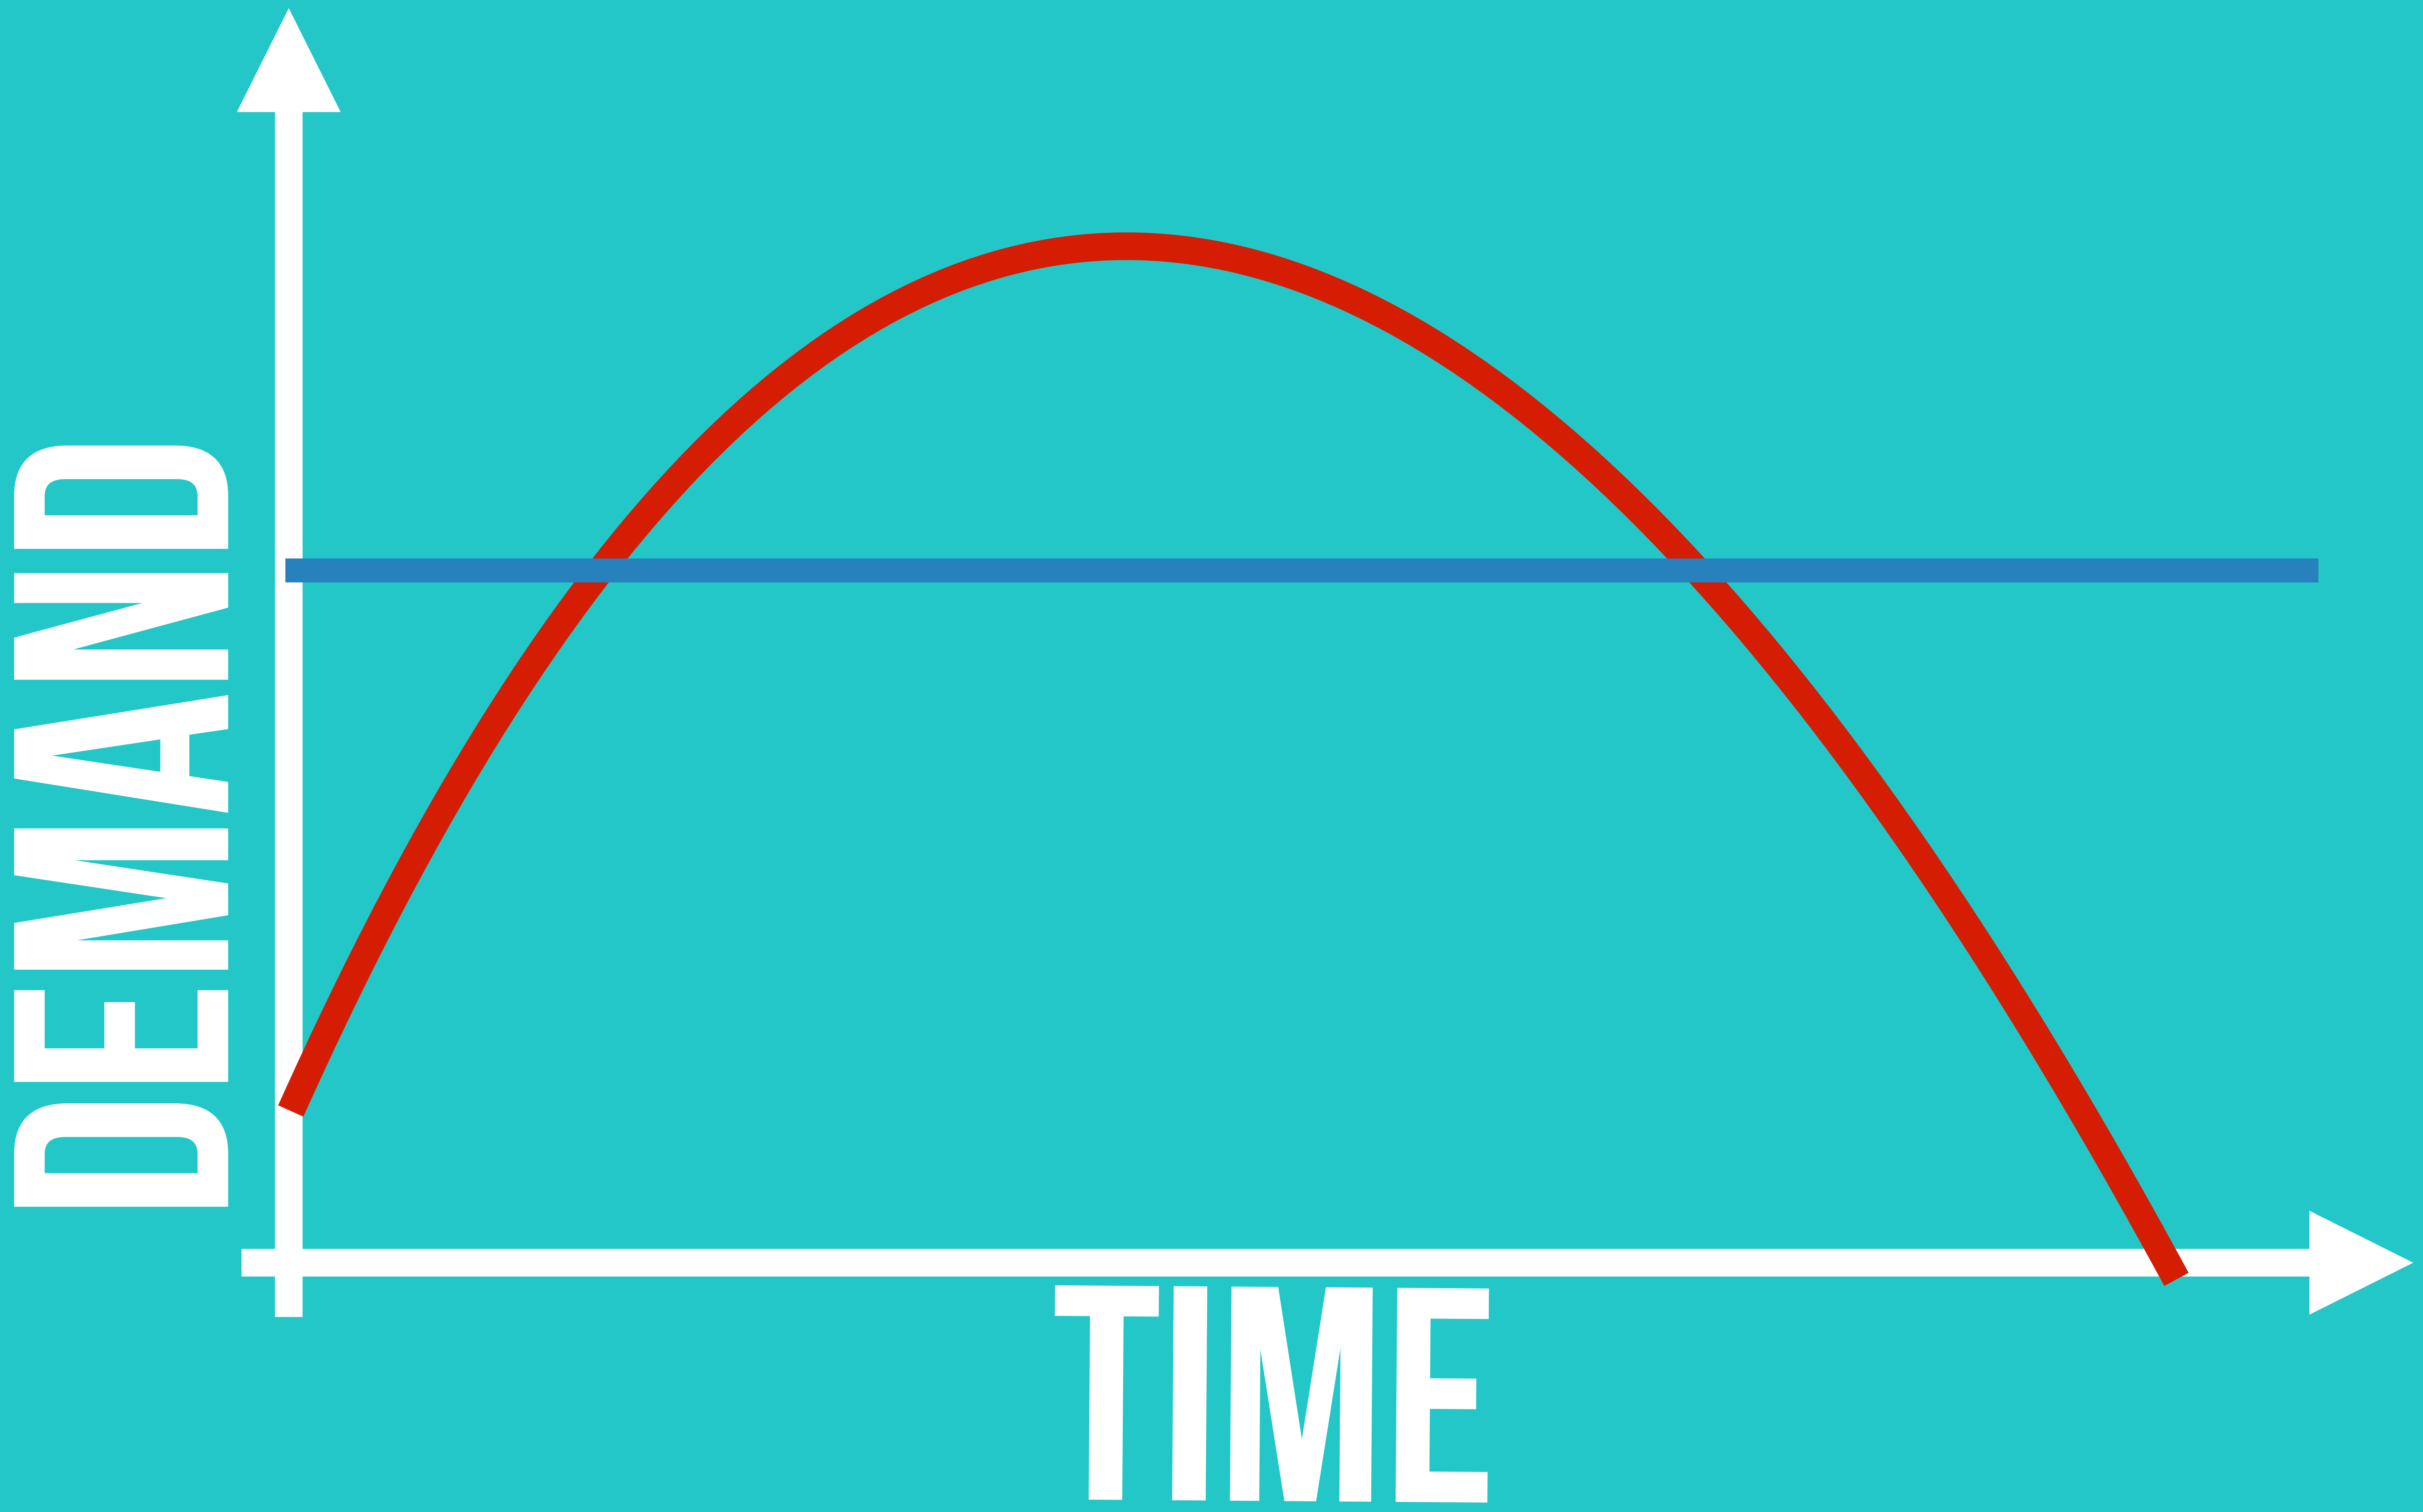
DEMAND
TIME

SCALE CUBE

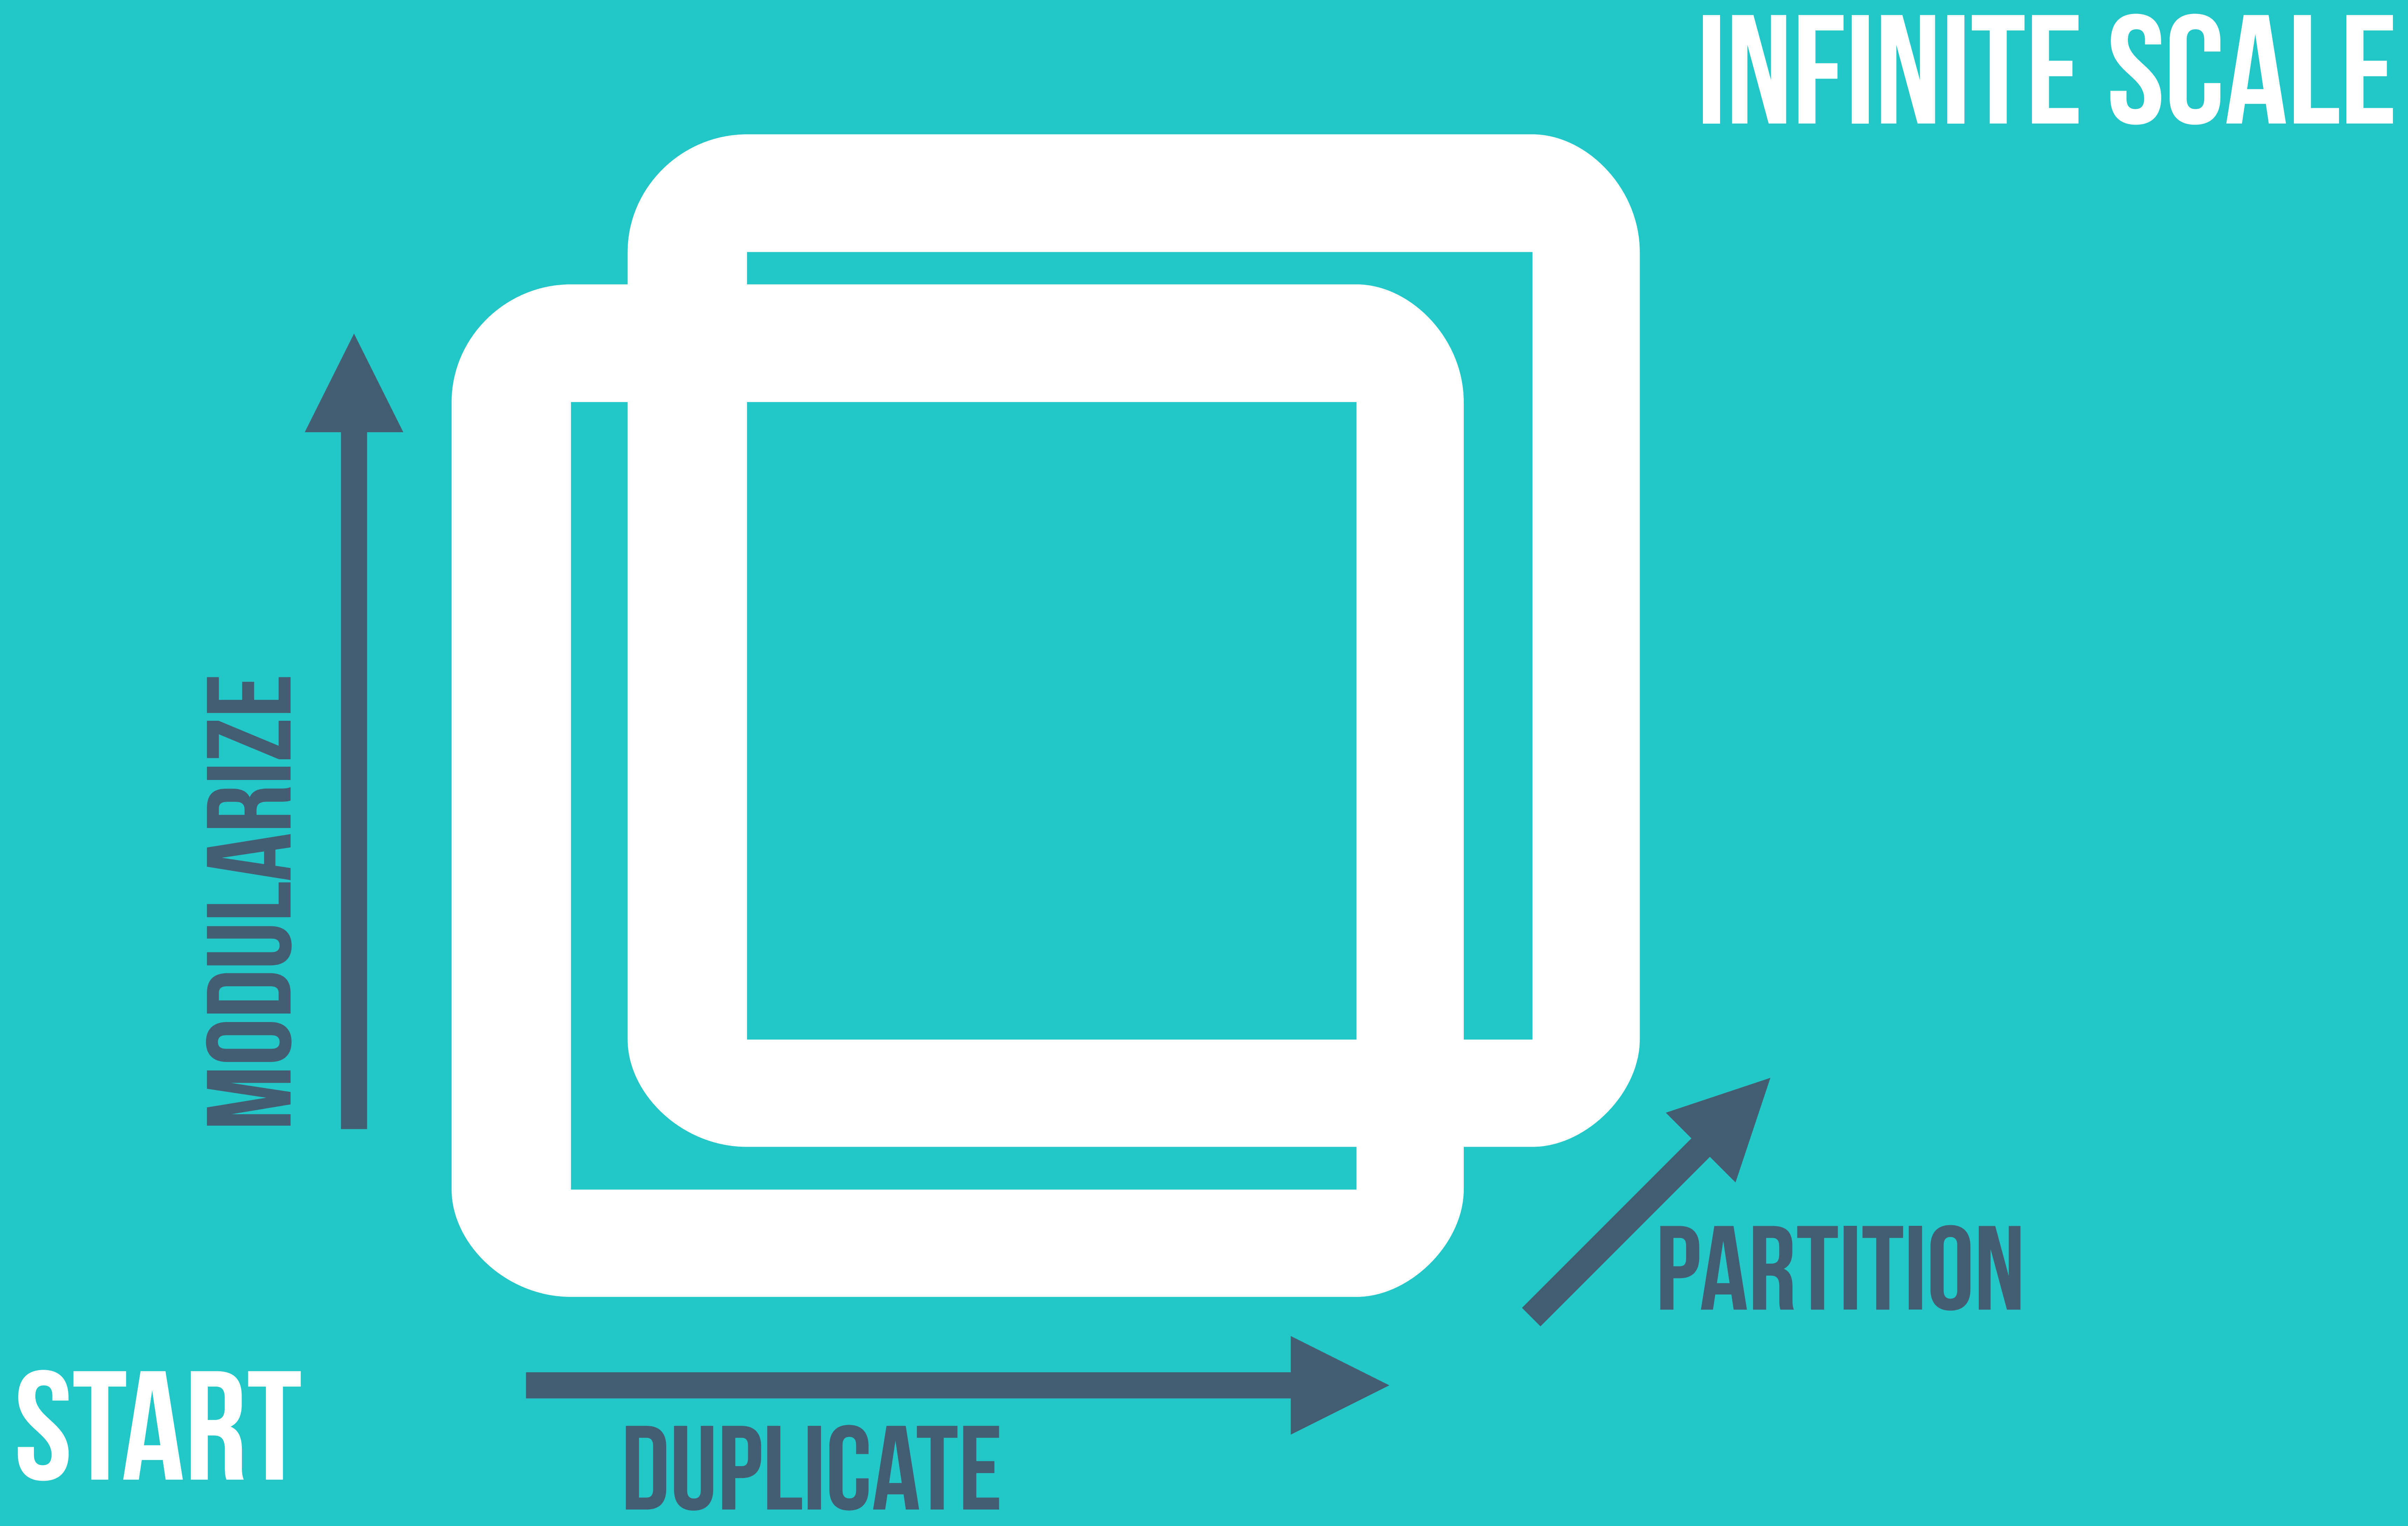

START

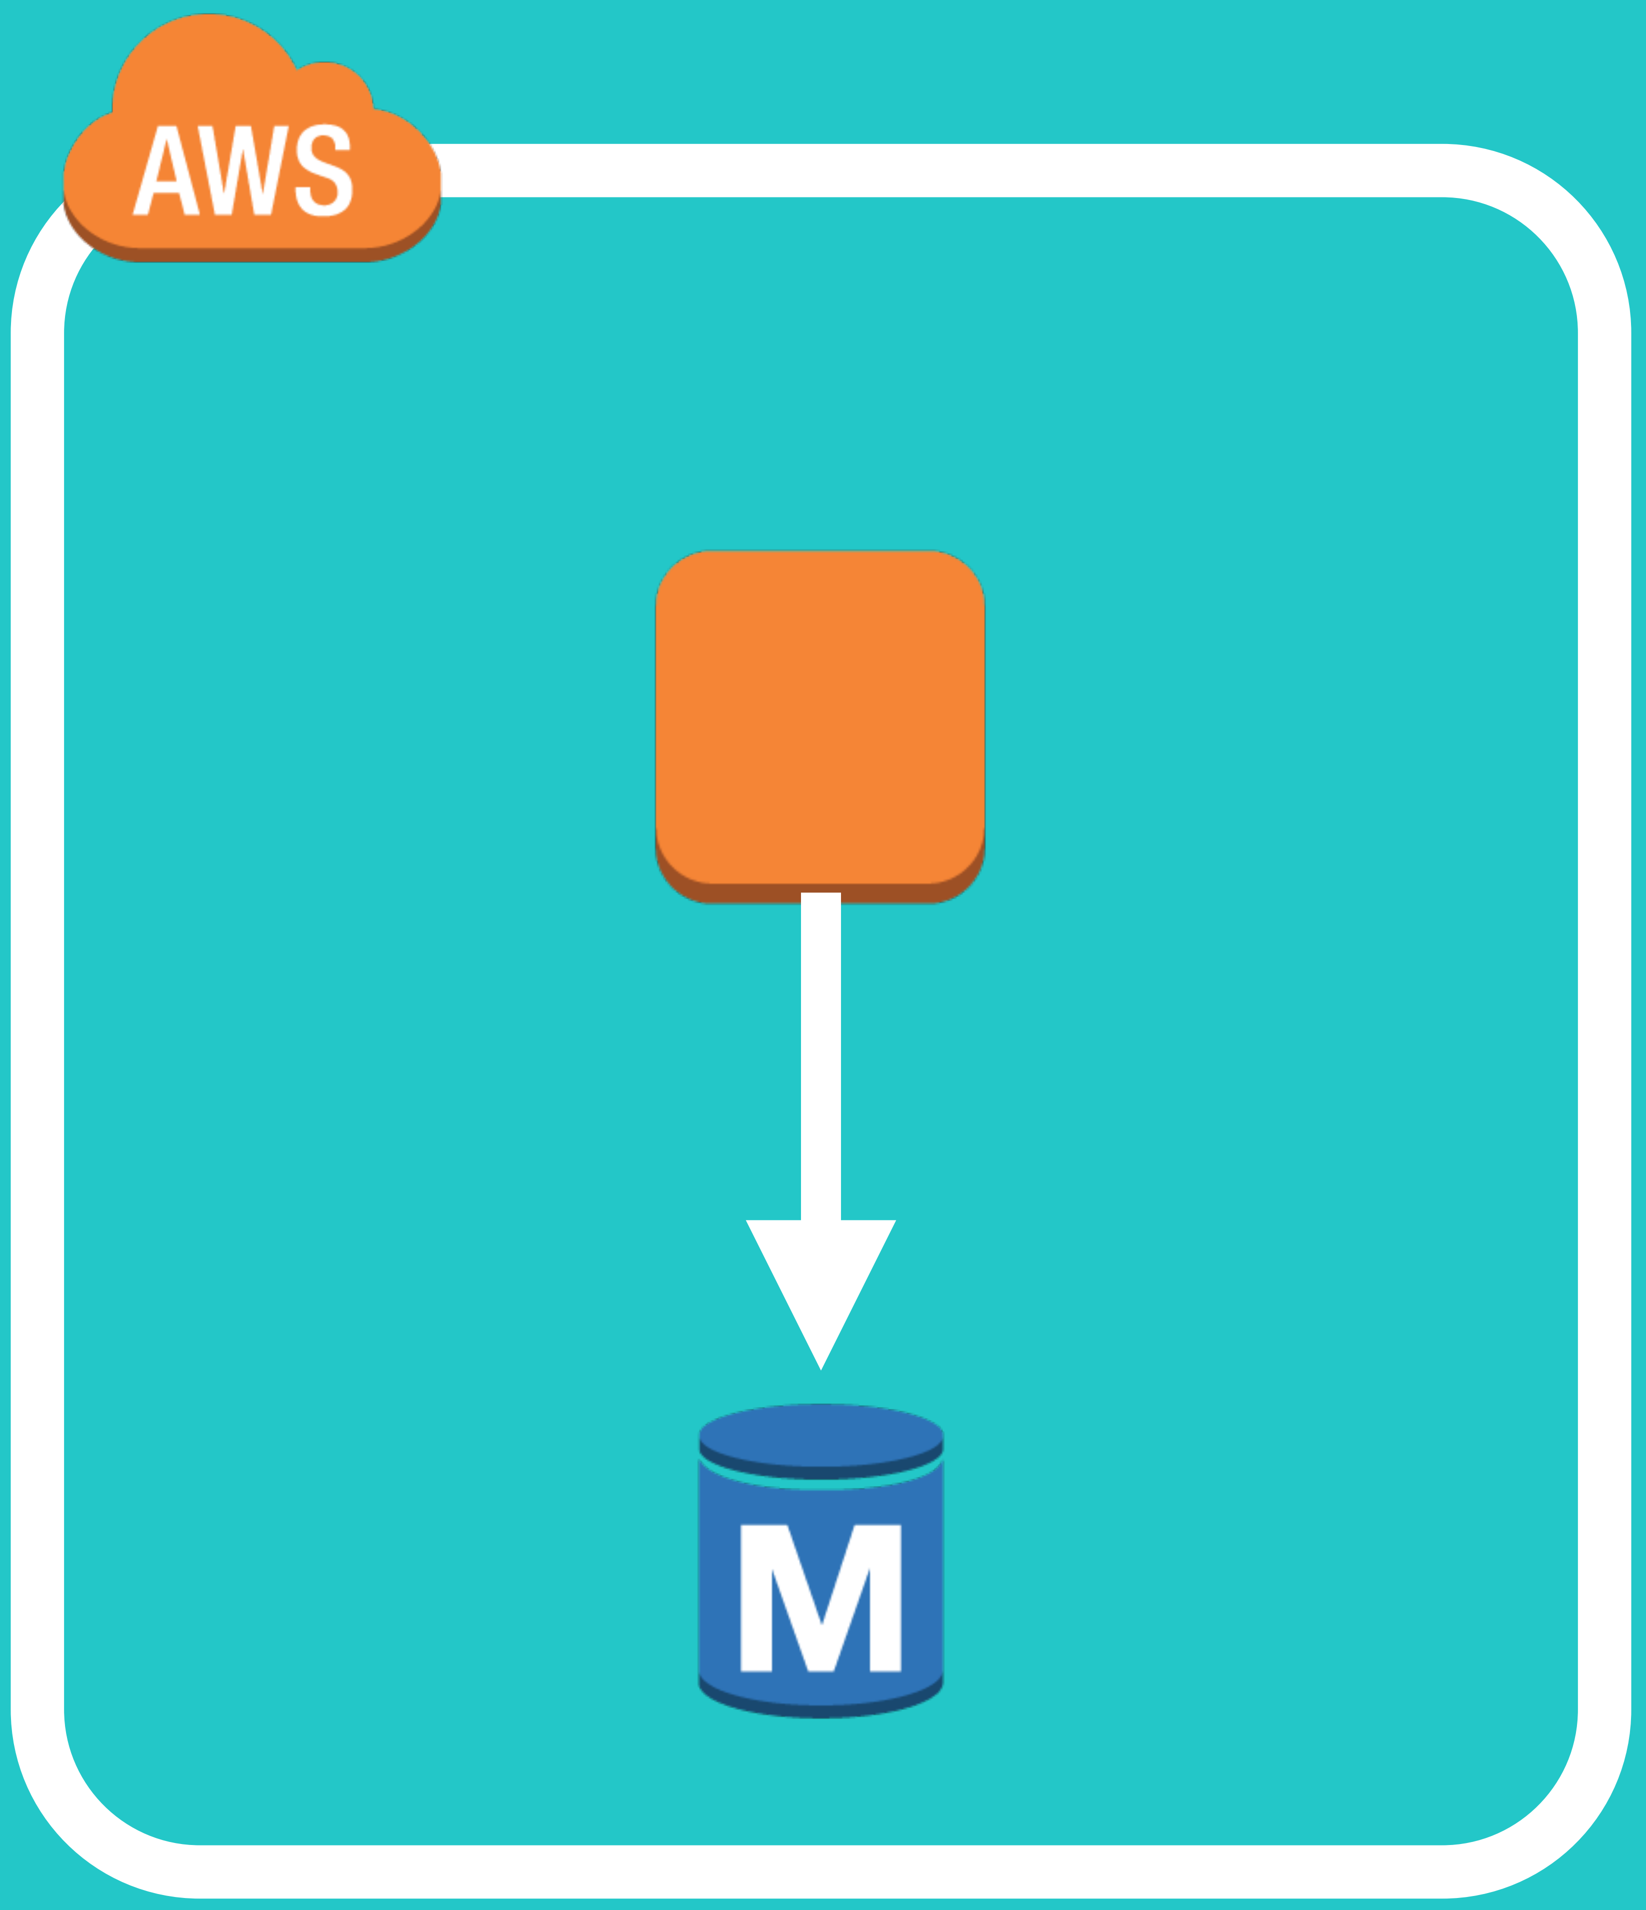

DUPLICATION

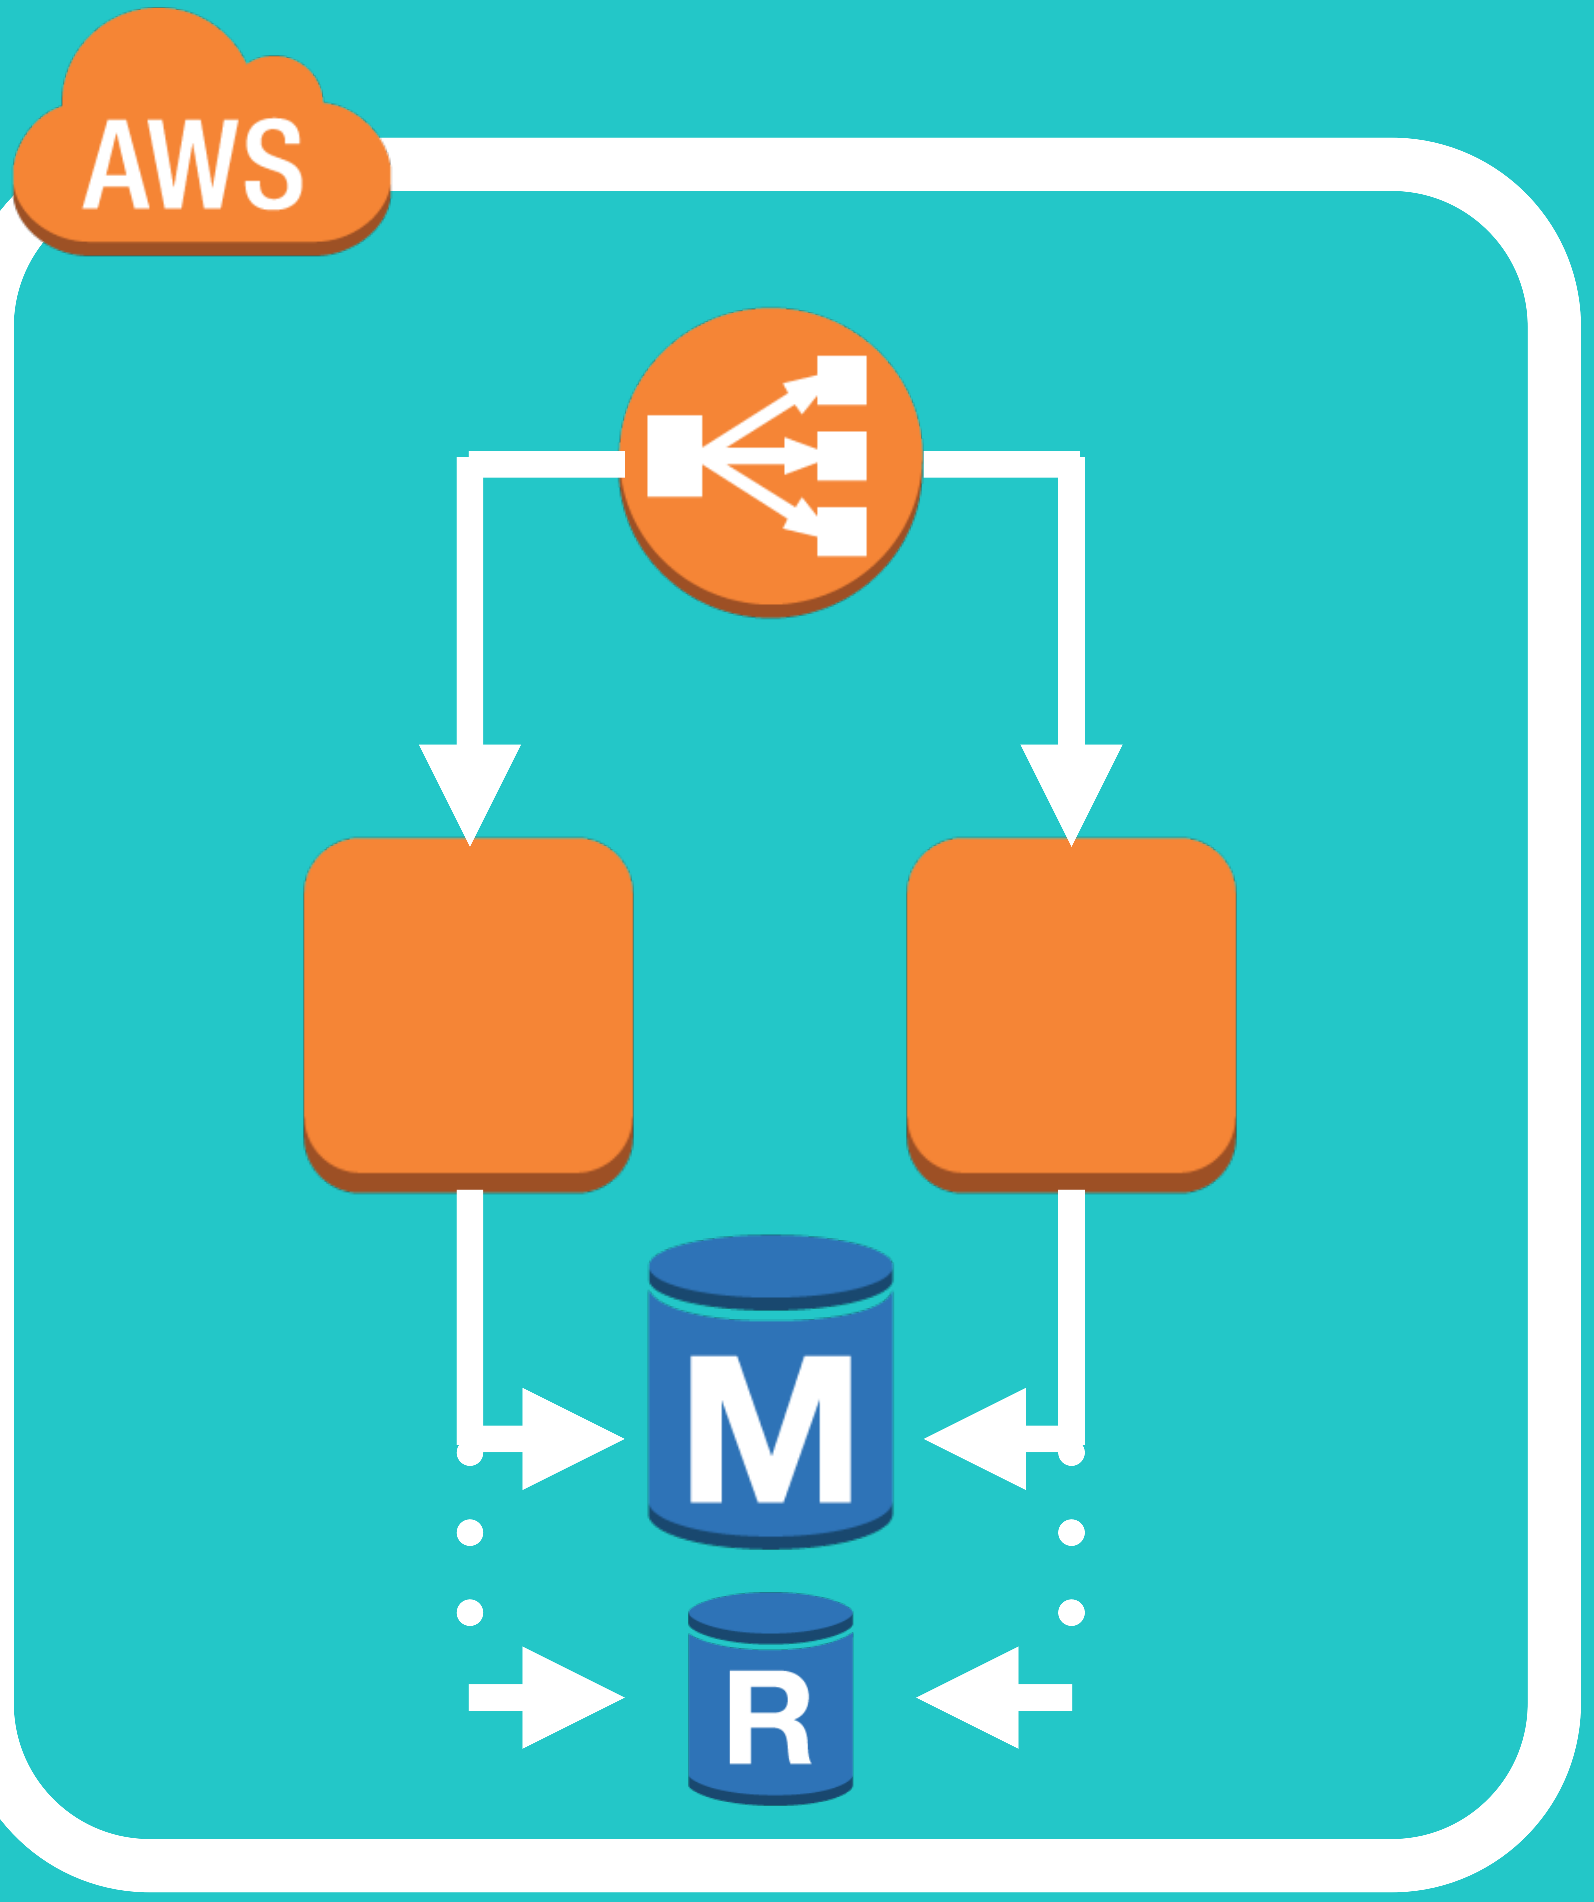

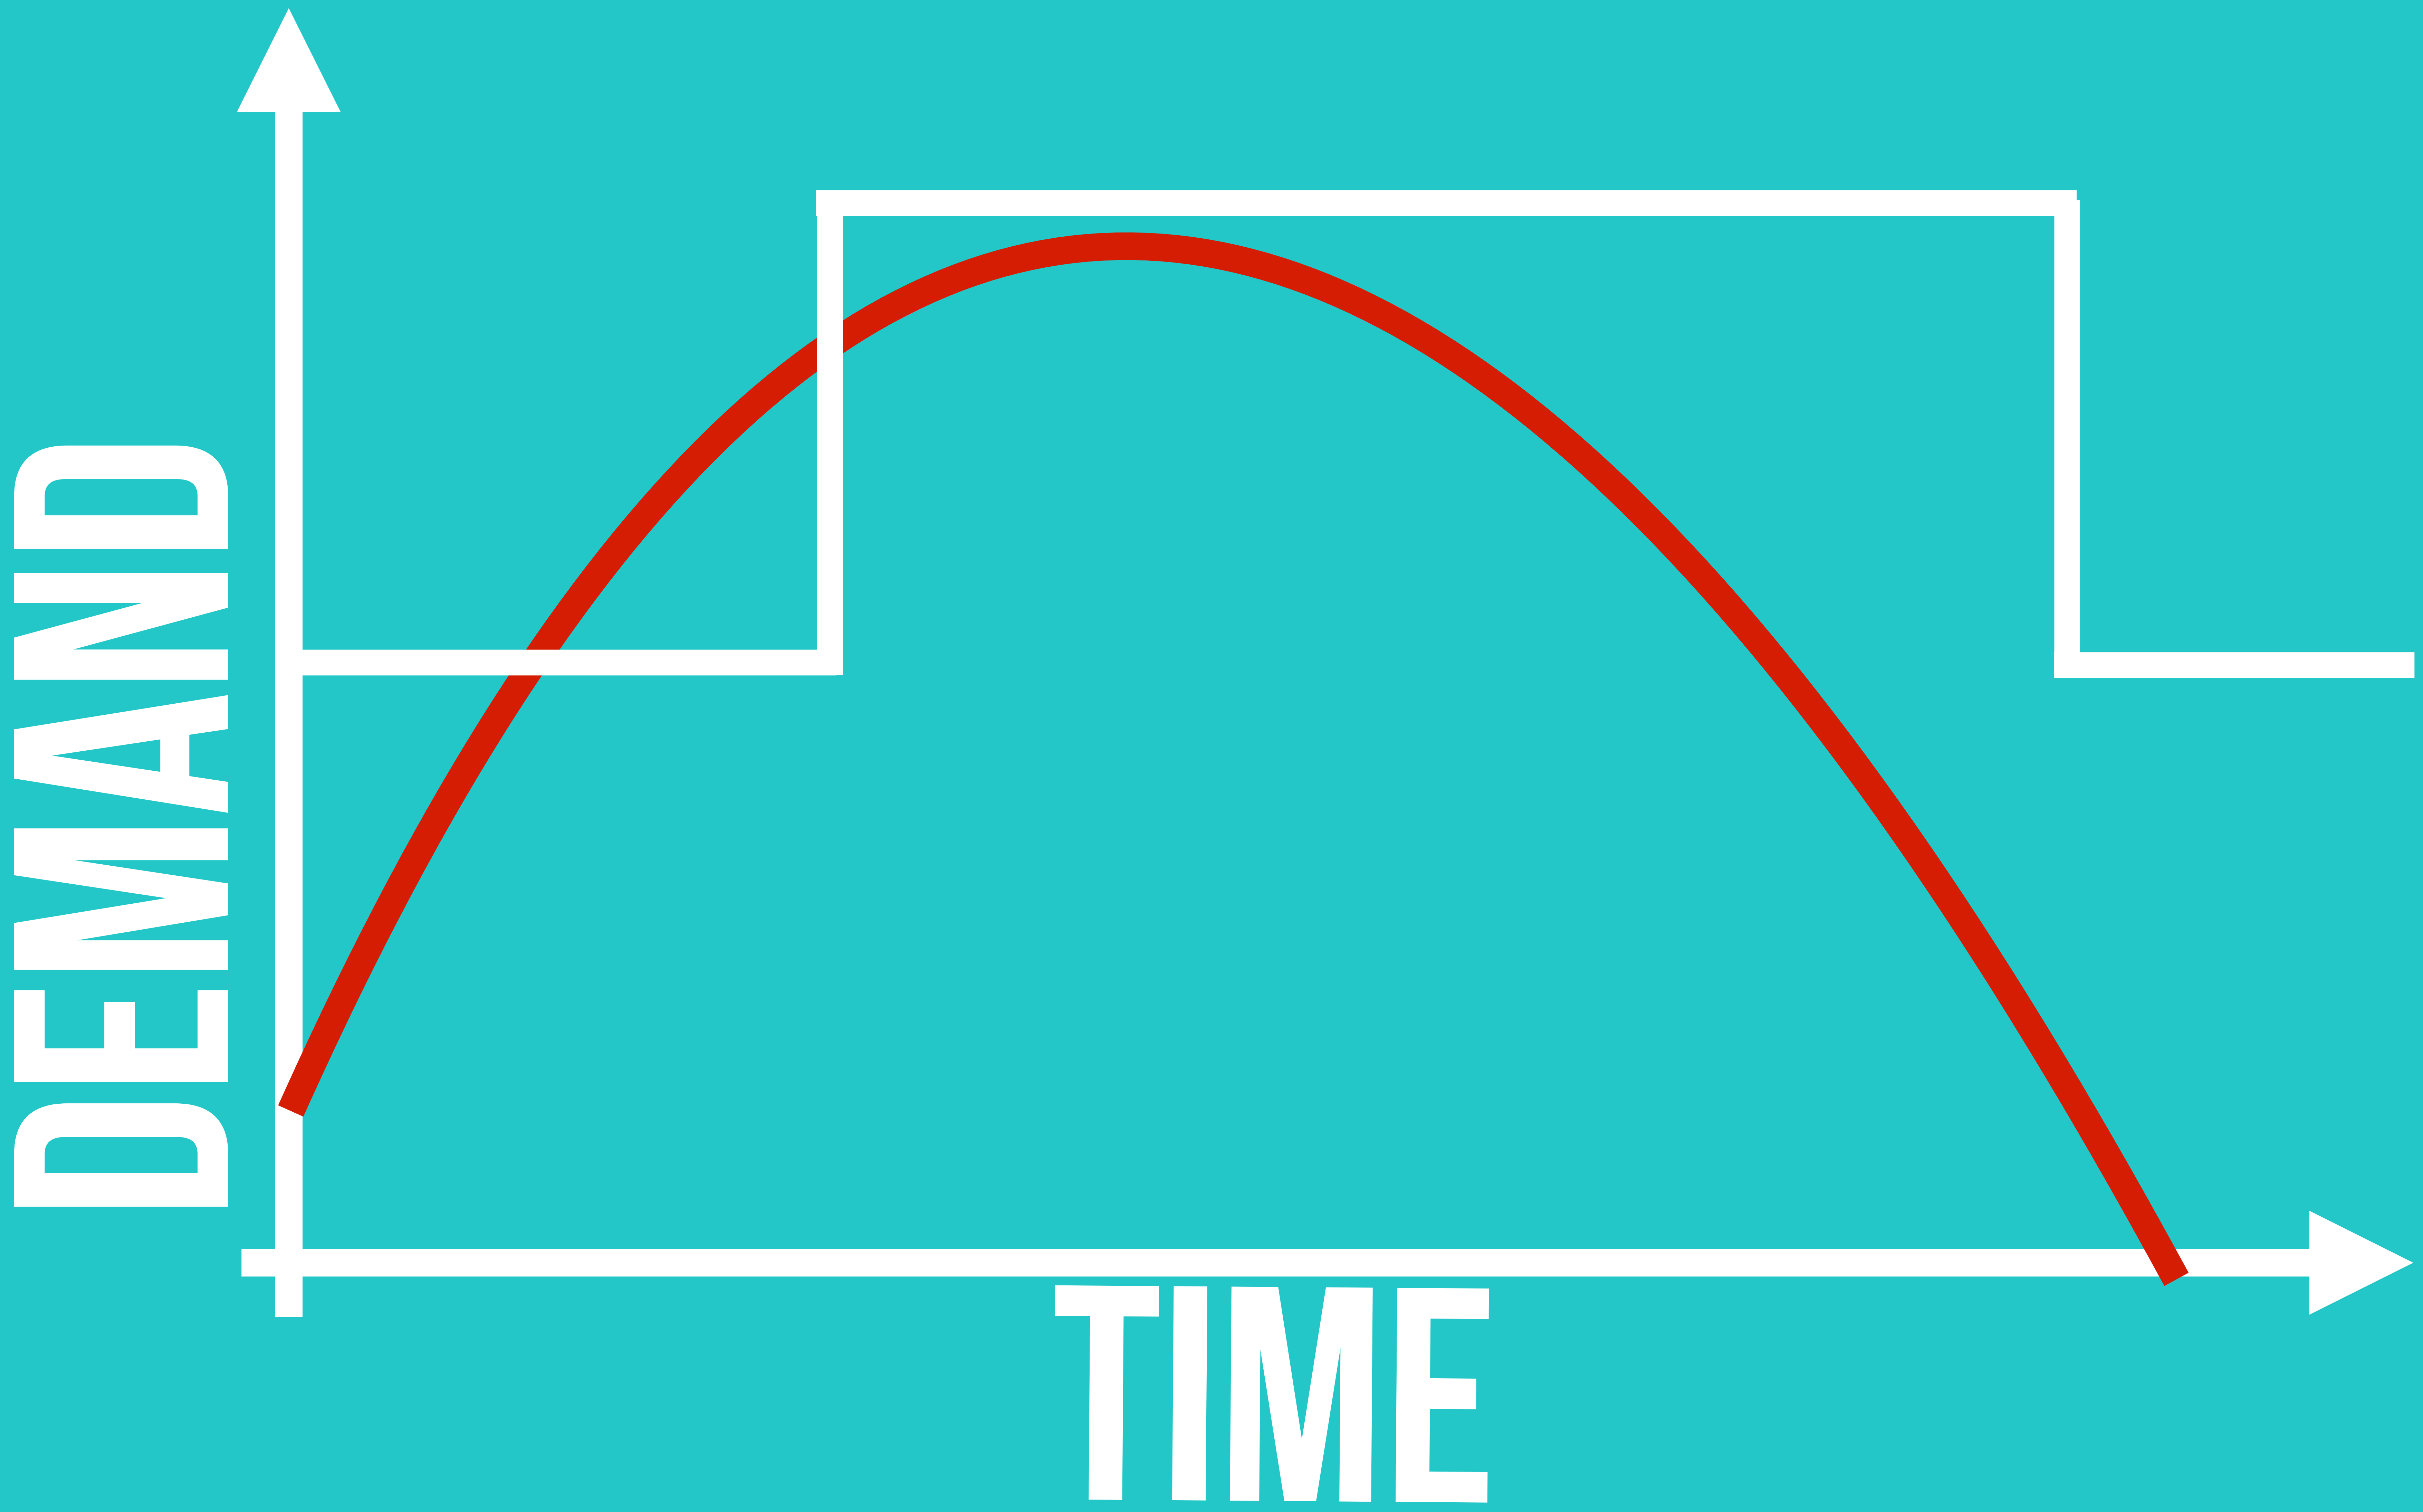
DEMAND
TIME

MODULARIZATION

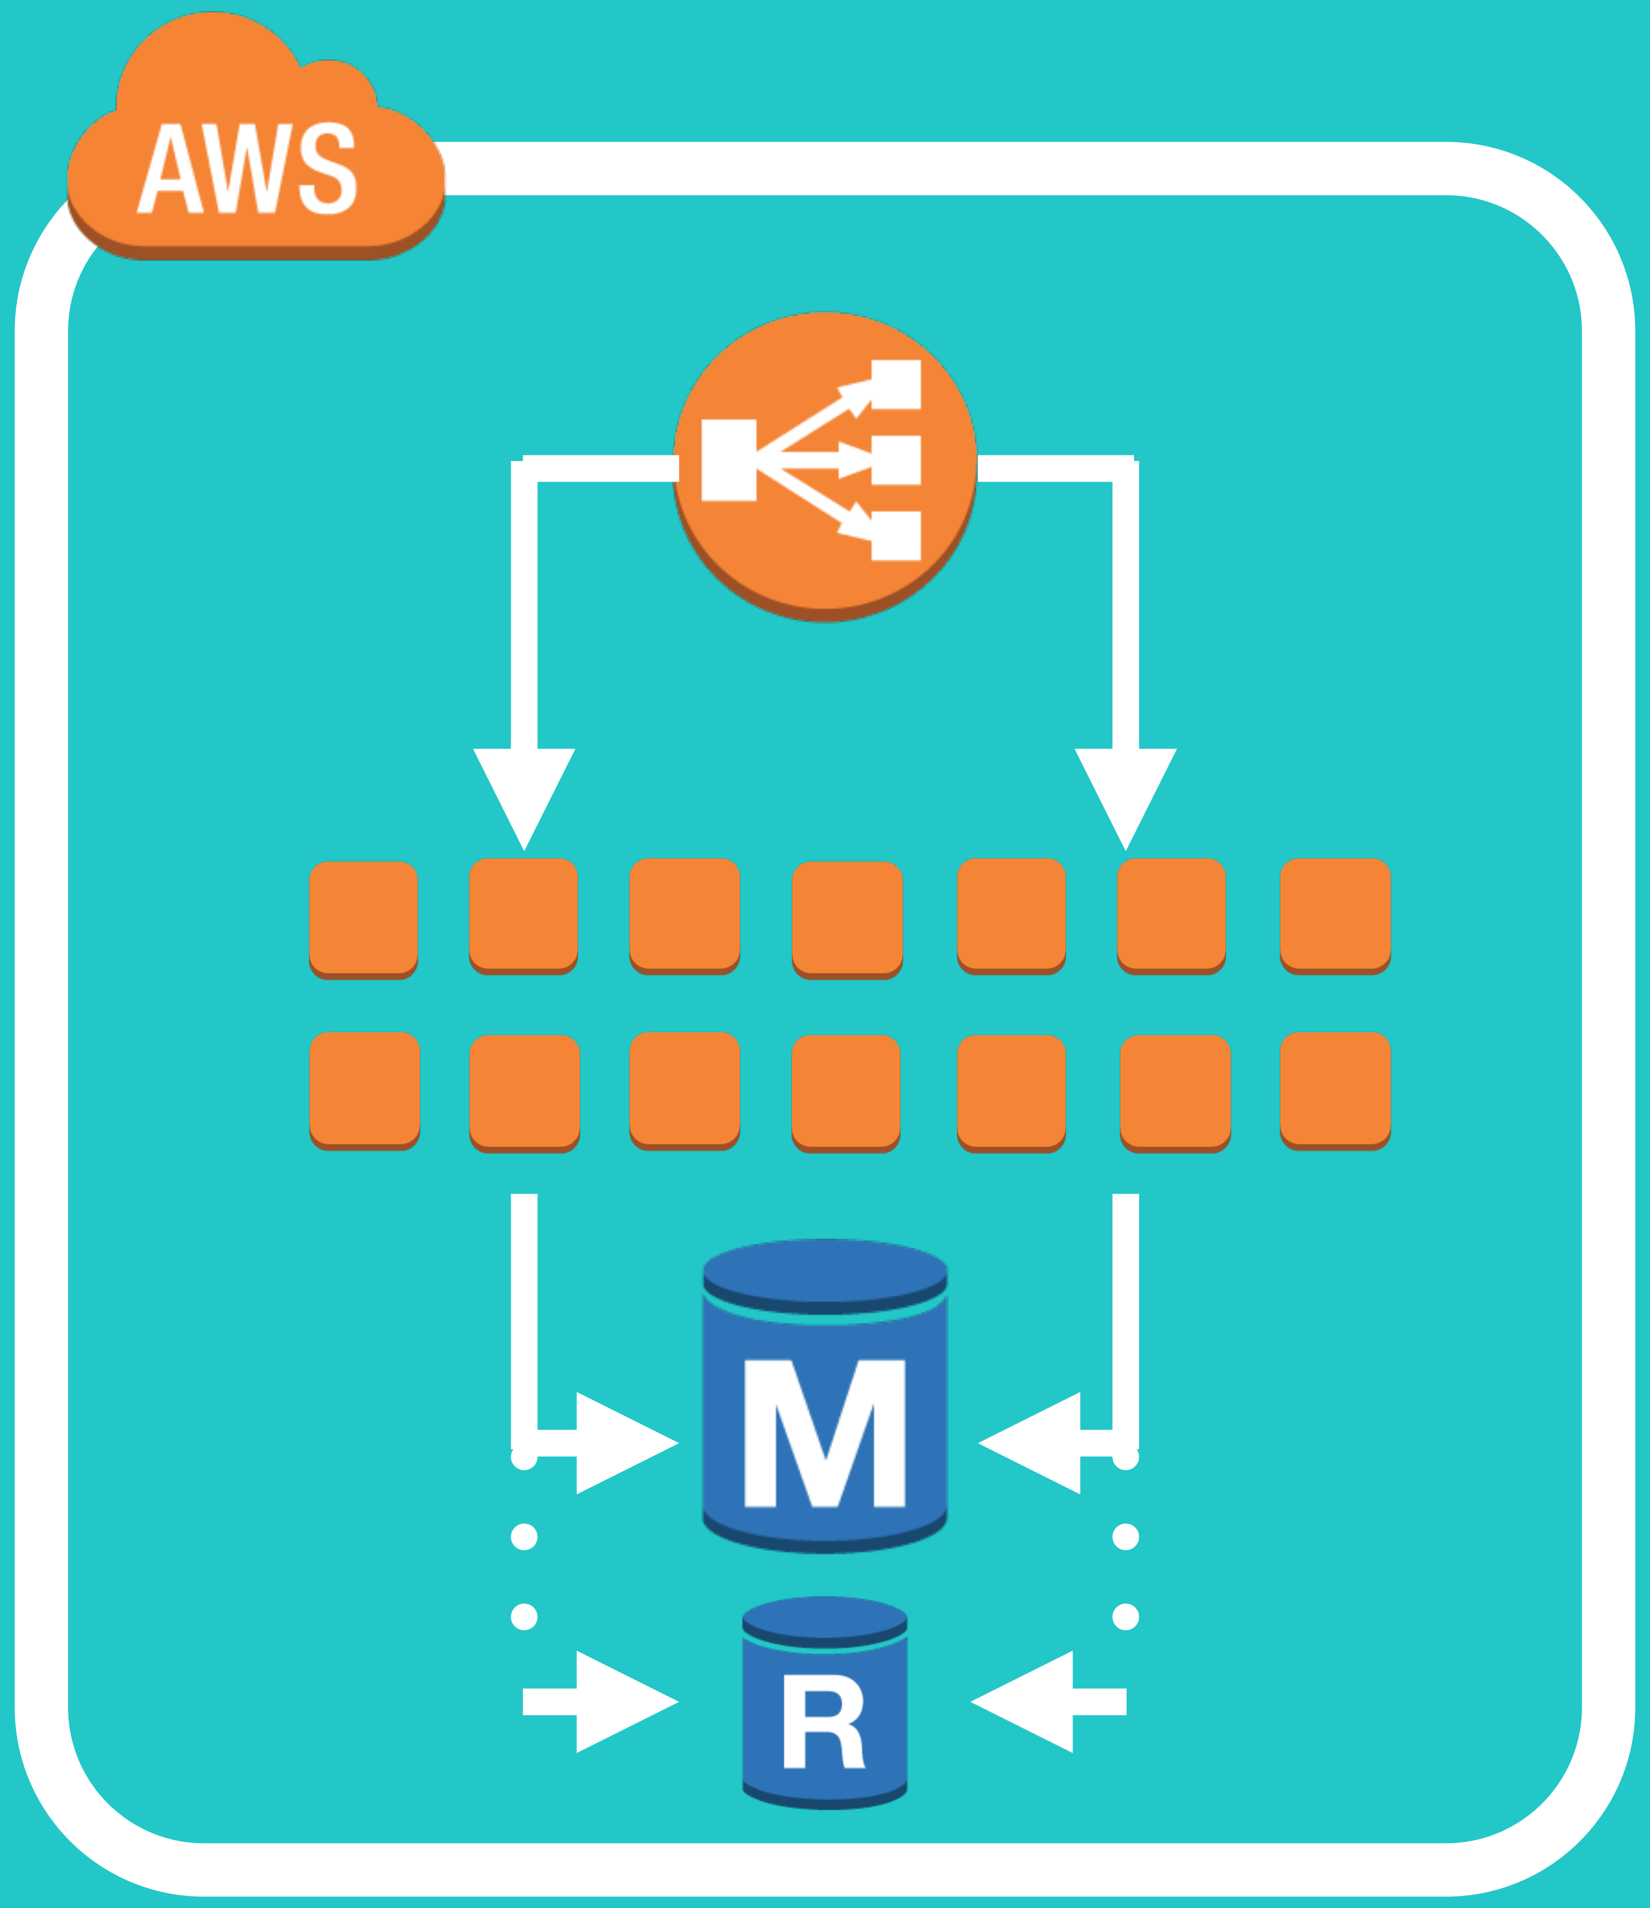

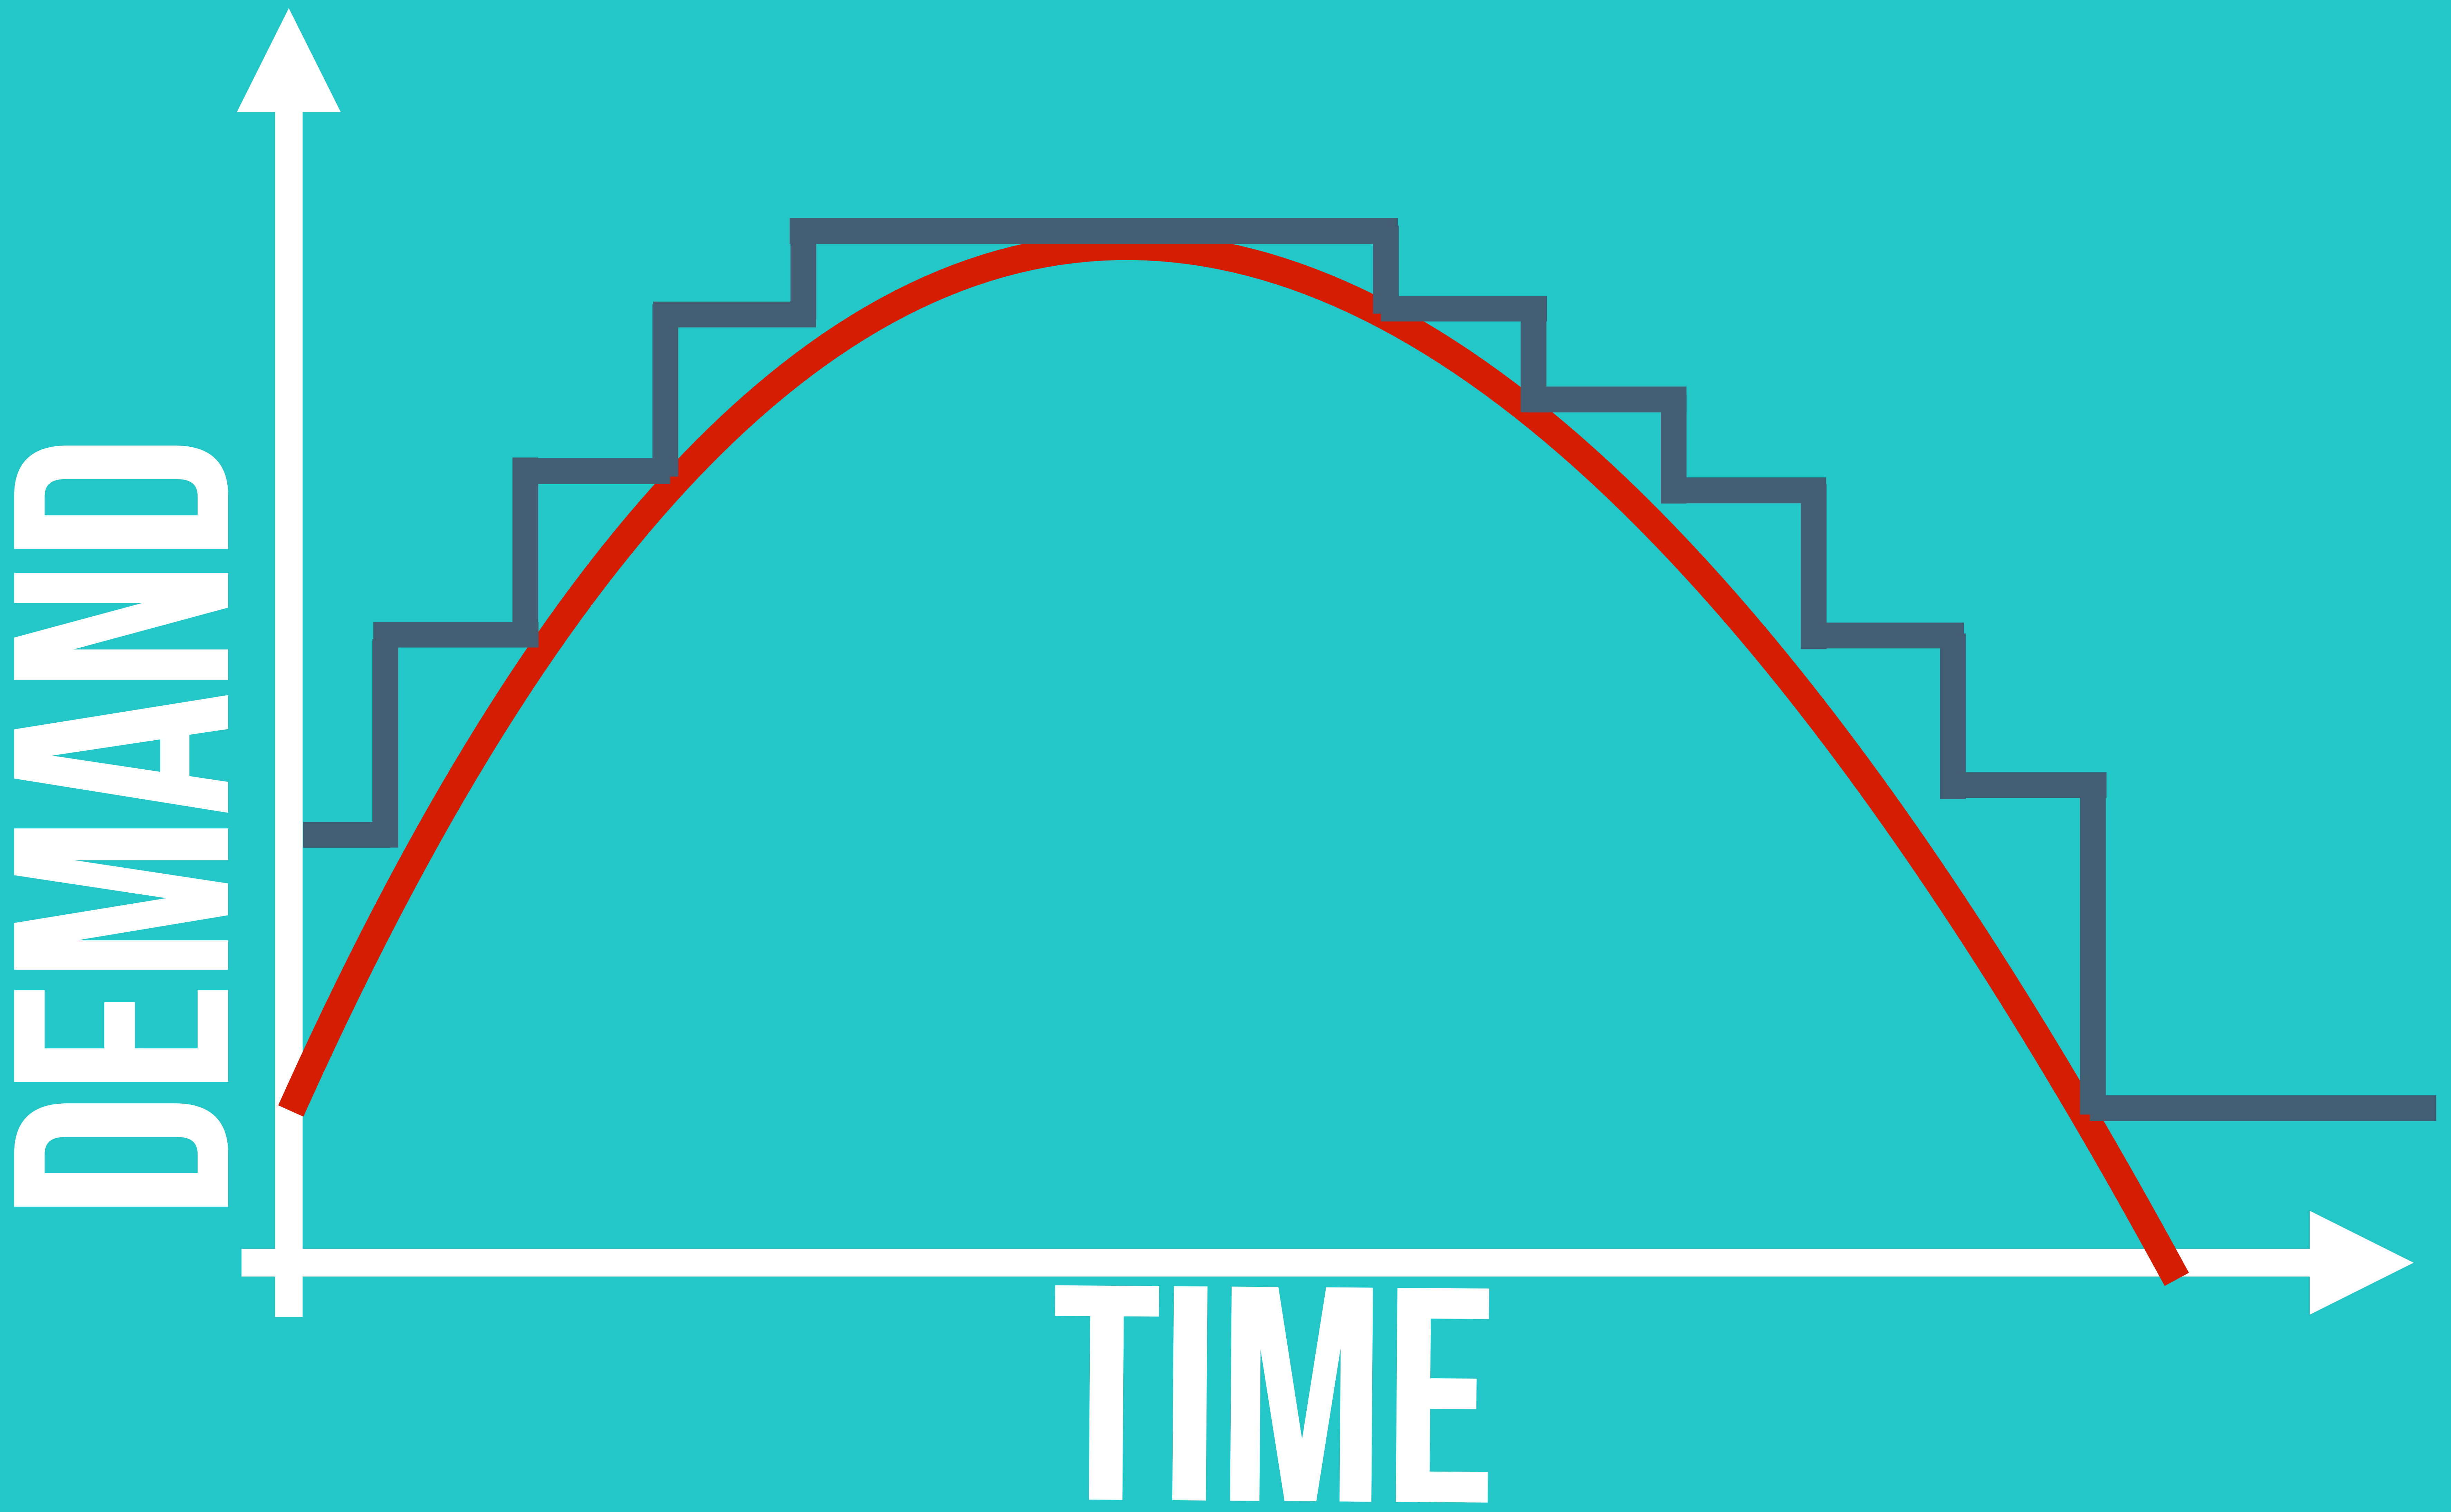
DEMAND
TIME

PARTIONING

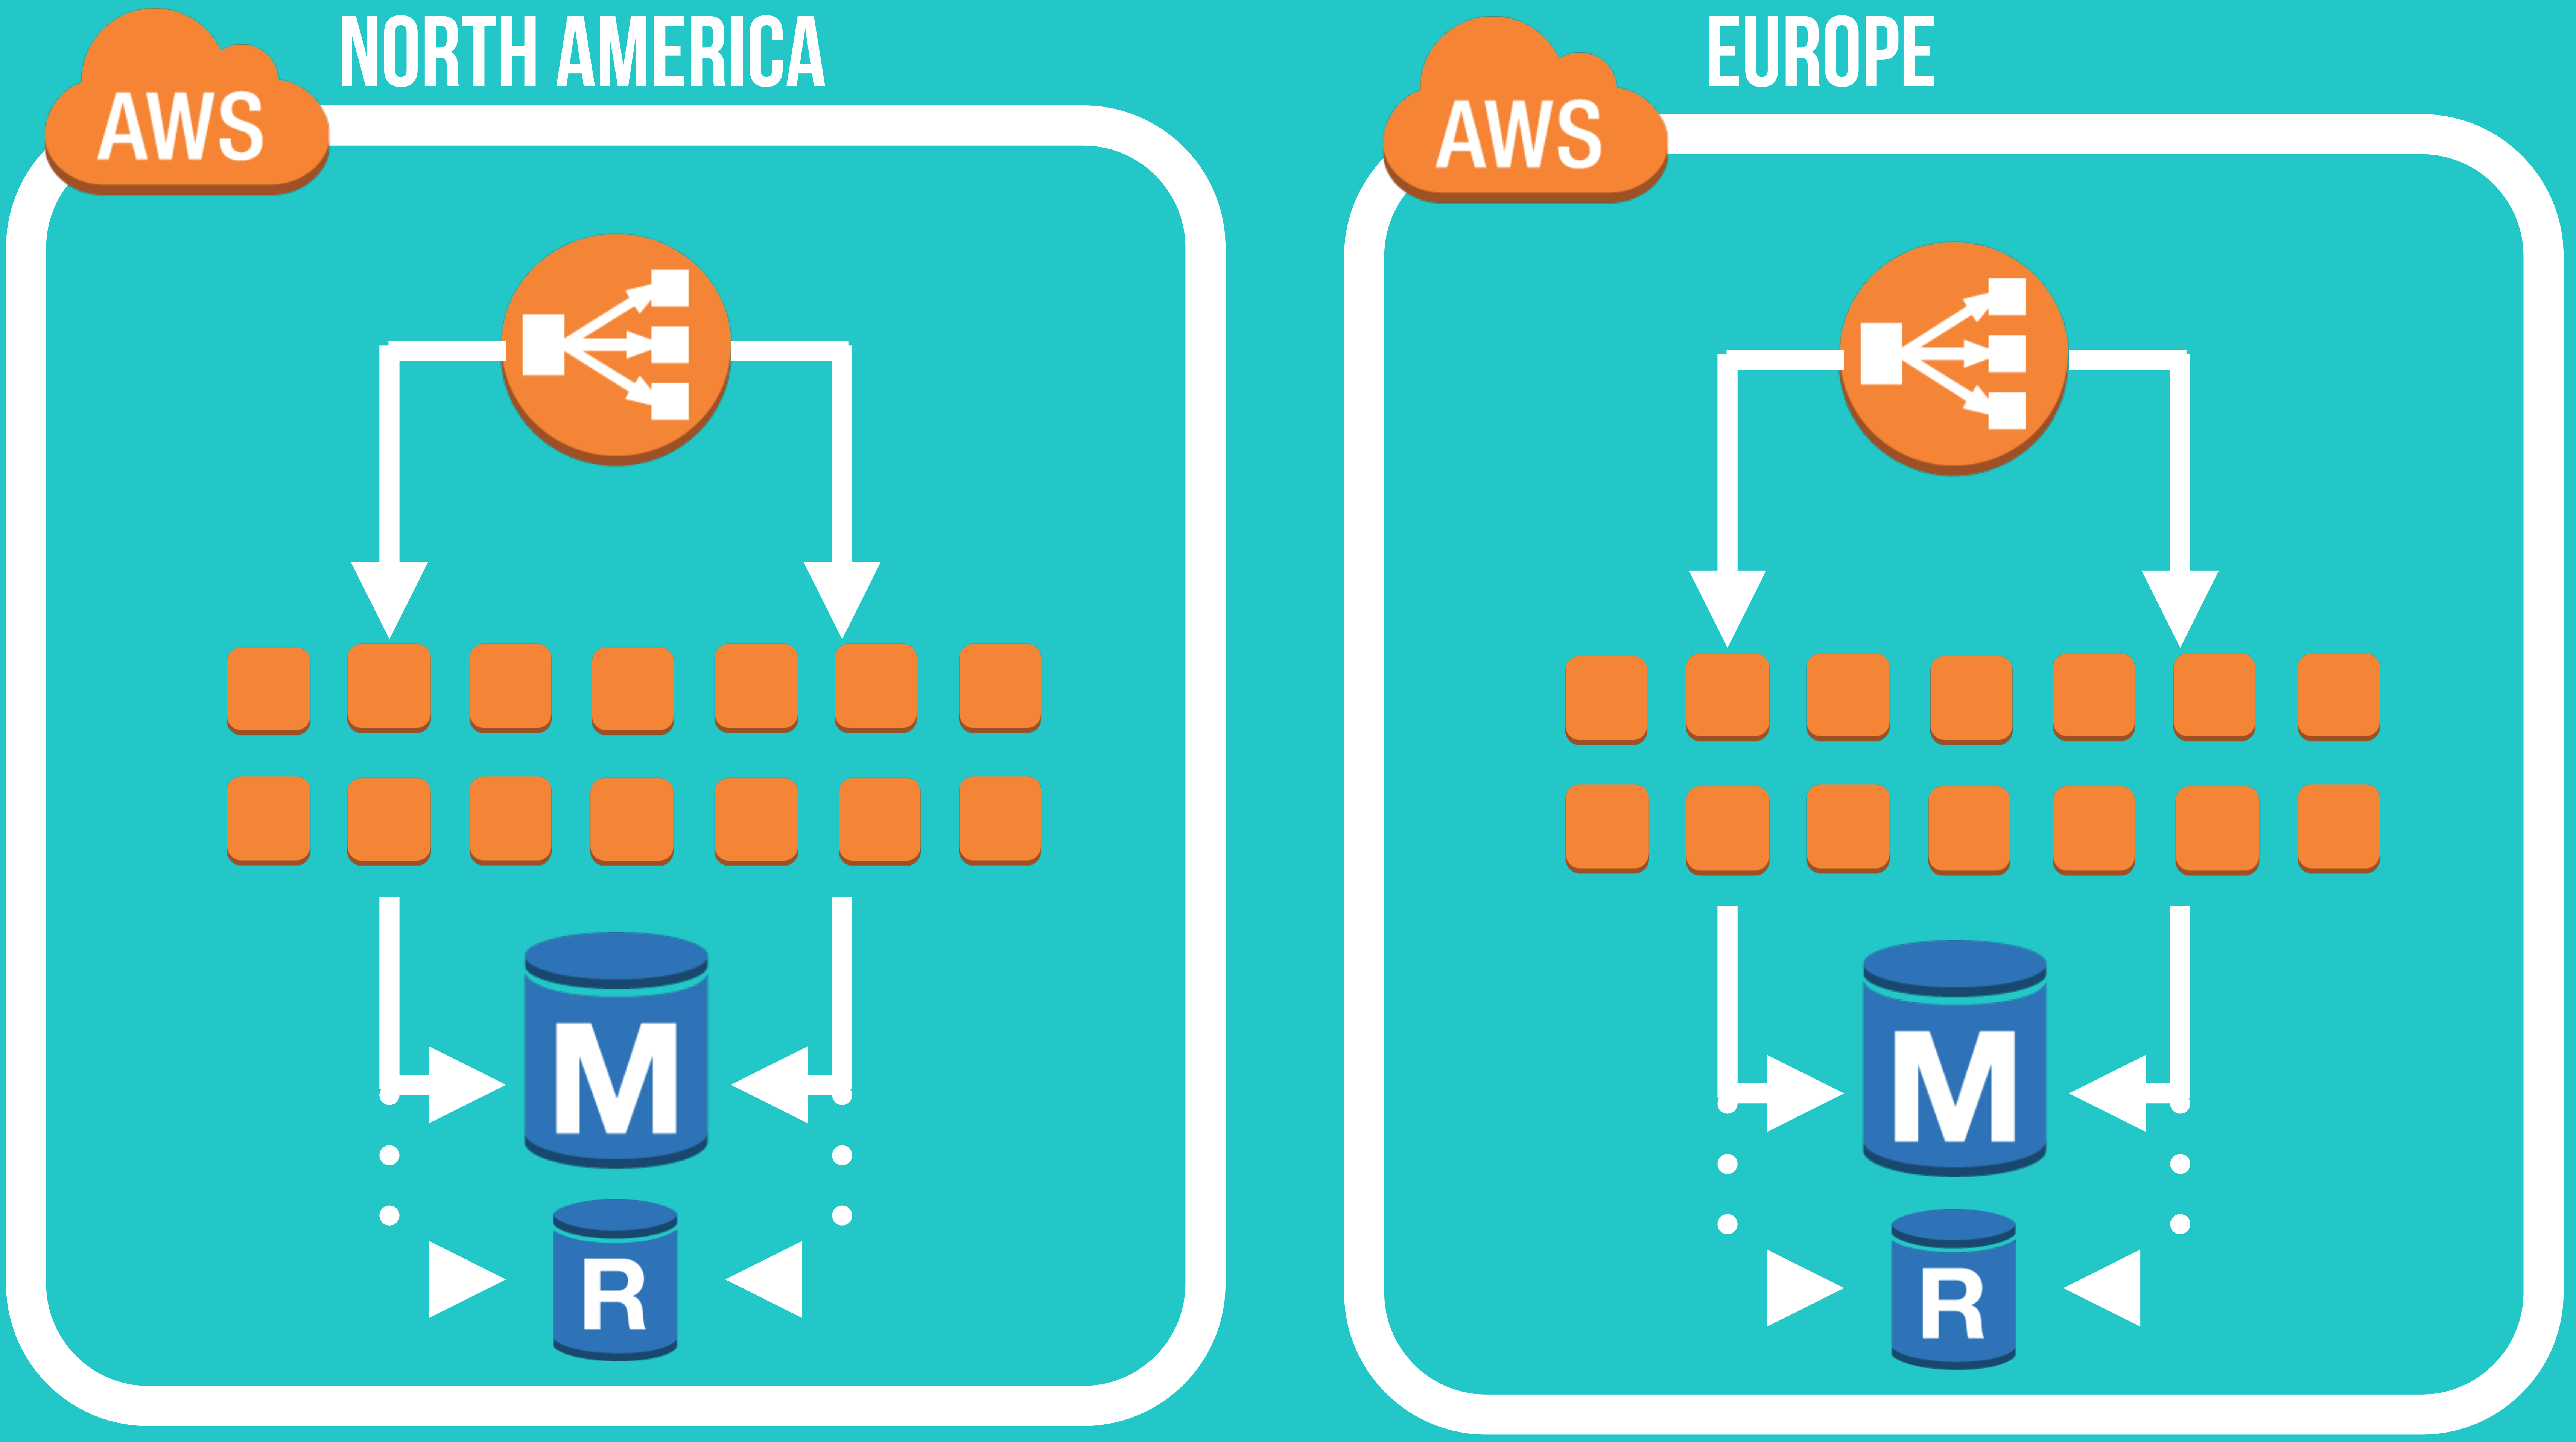

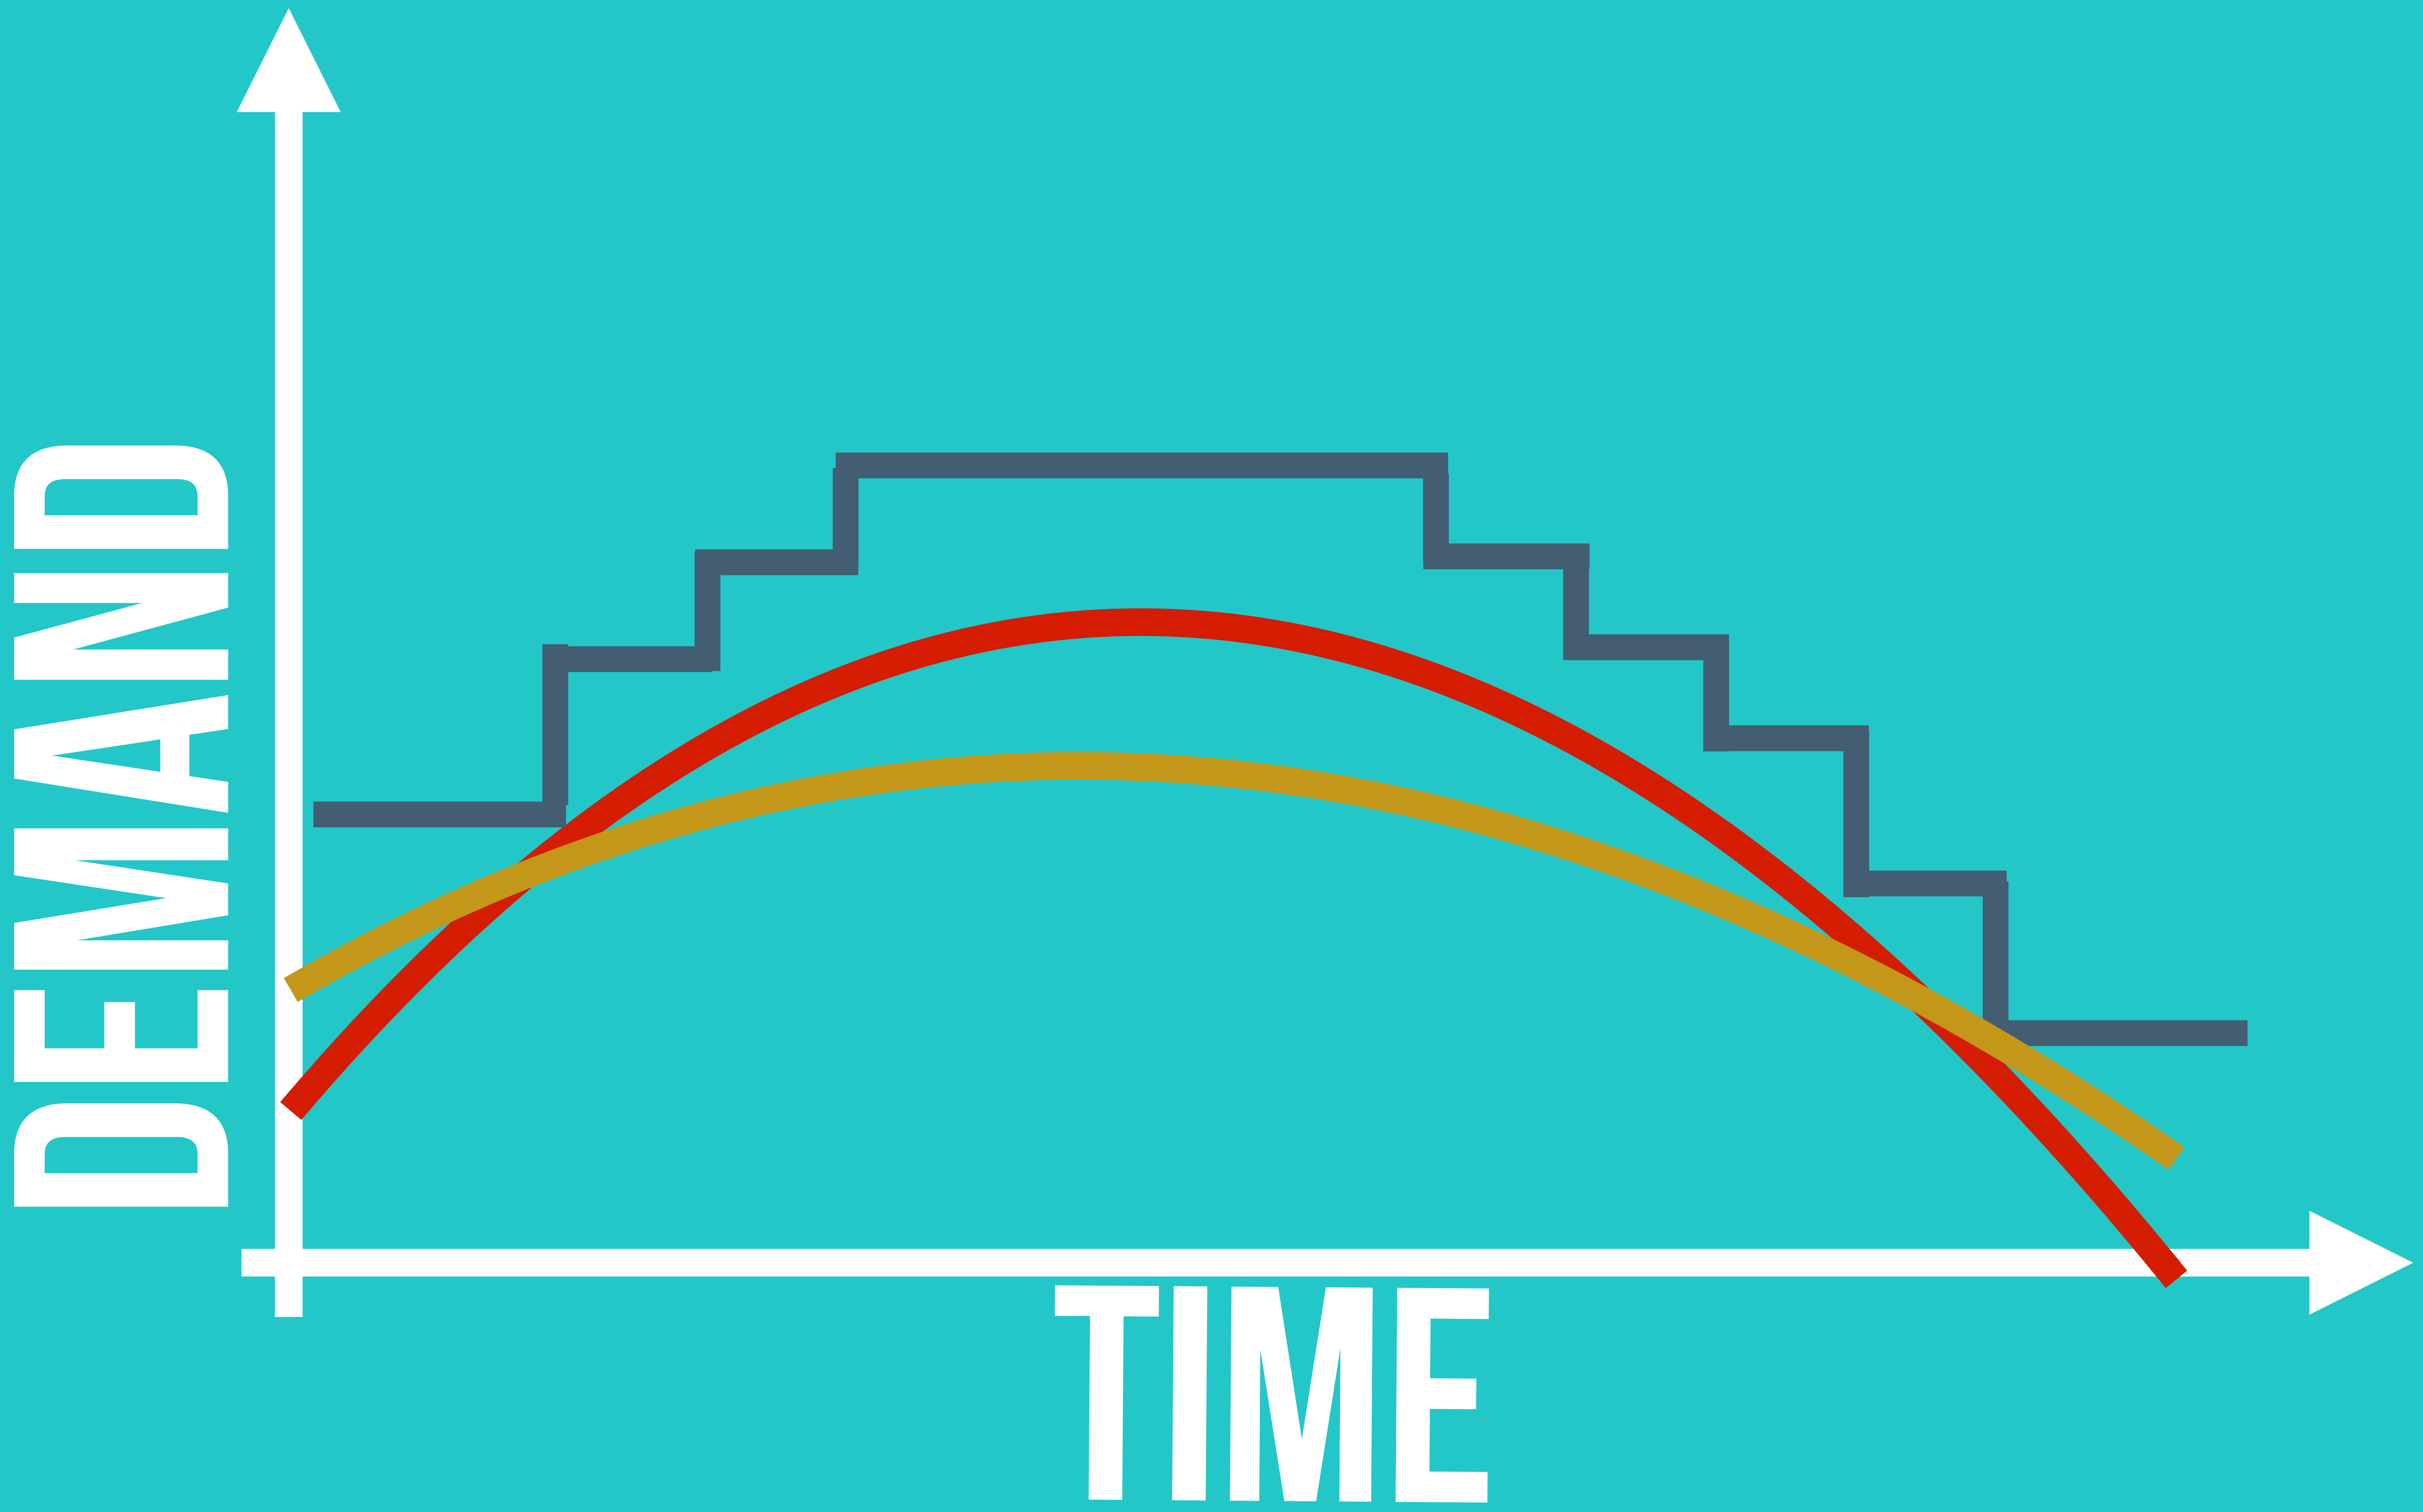
DEMAND
TIME

{YOU NAME IT}
SERVICE

A MICROSERVICES ARCHITECTURE IS A SERVICE-ORIENTED ARCHITECTURE COMPOSED OF LOOSELY COUPLED ELEMENTS THAT HAVE BOUNDED CONTEXTS

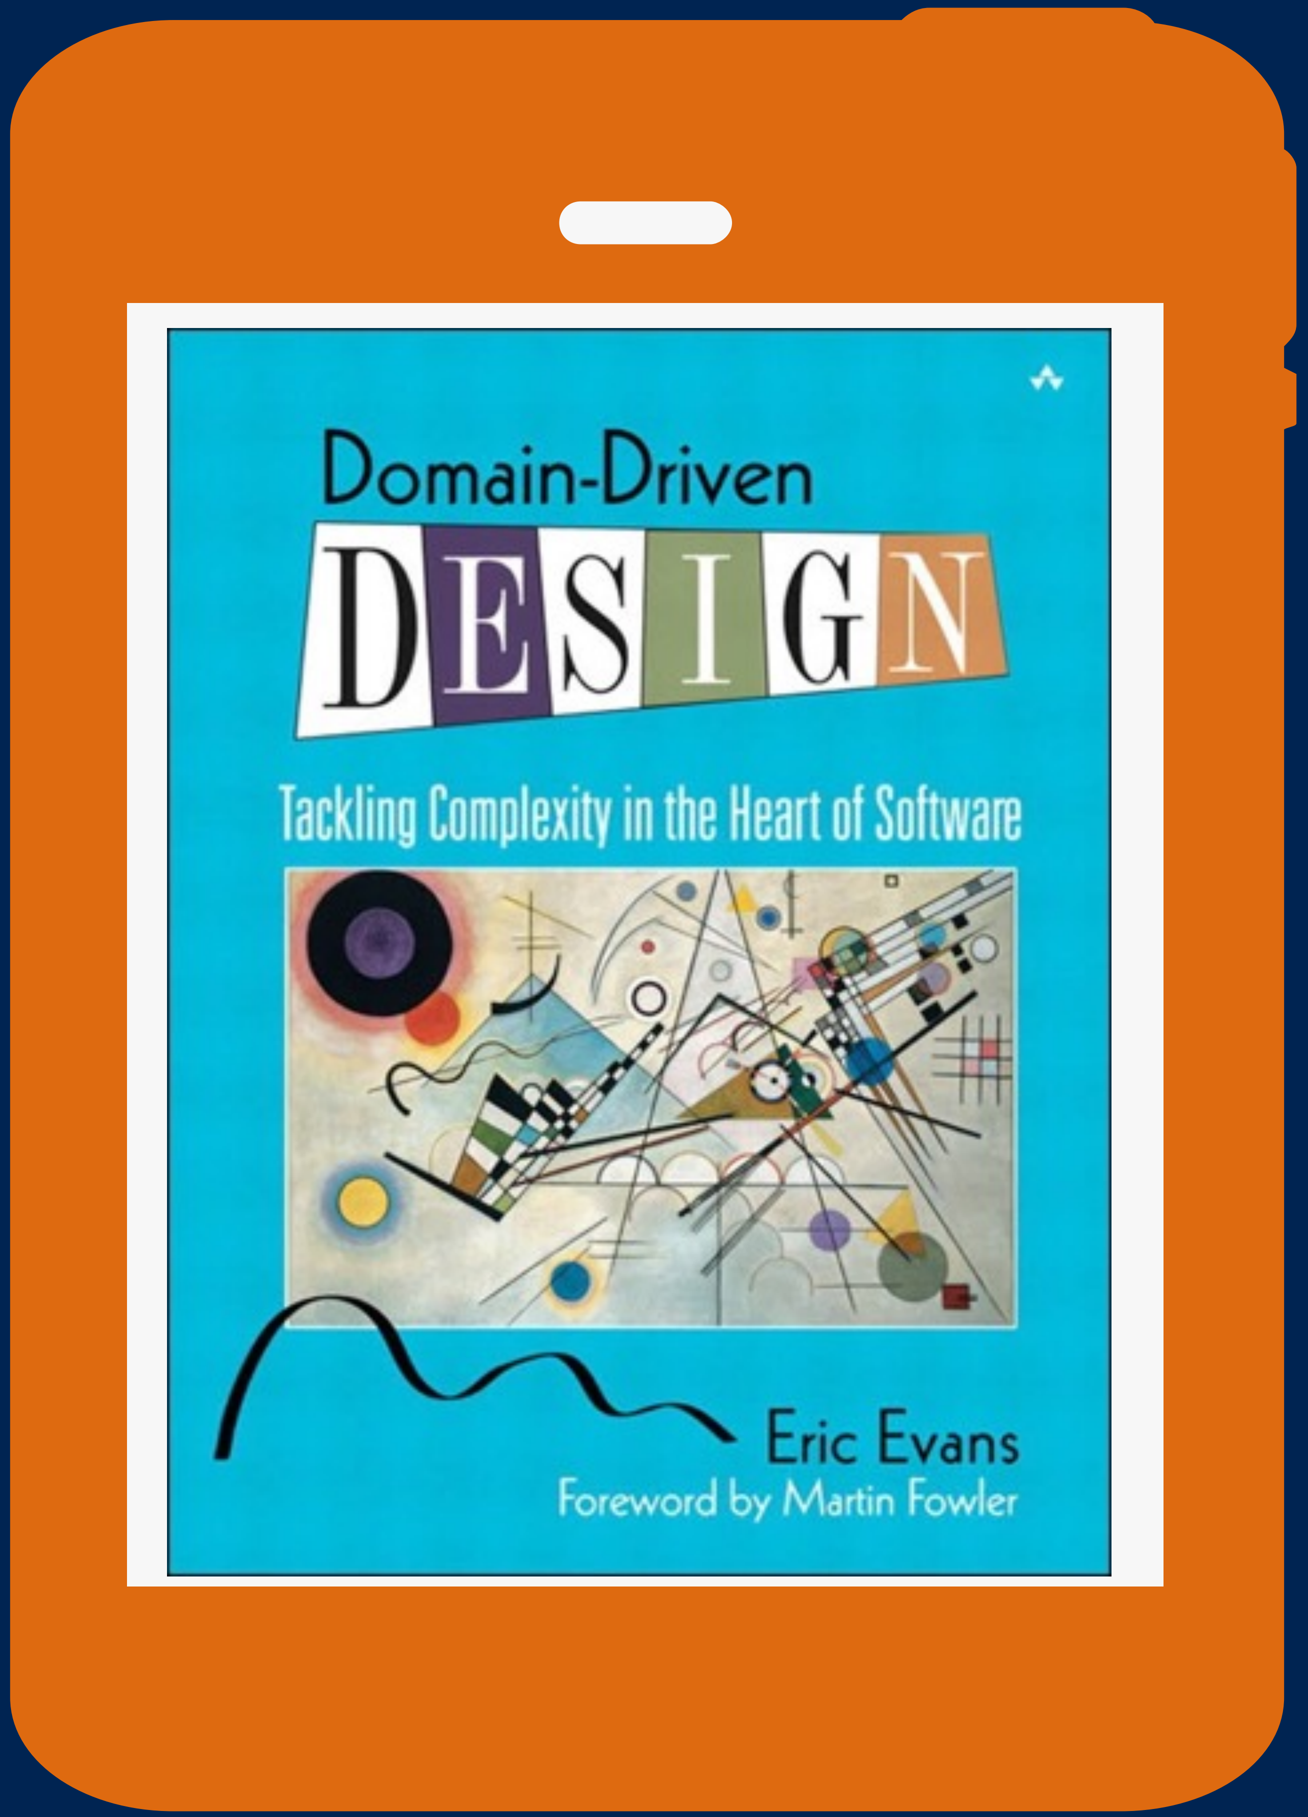
Domain-Driven

DESIGN

Tackling Complexity in the Heart of Software

Eric Evans

Foreword by Martin Fowler

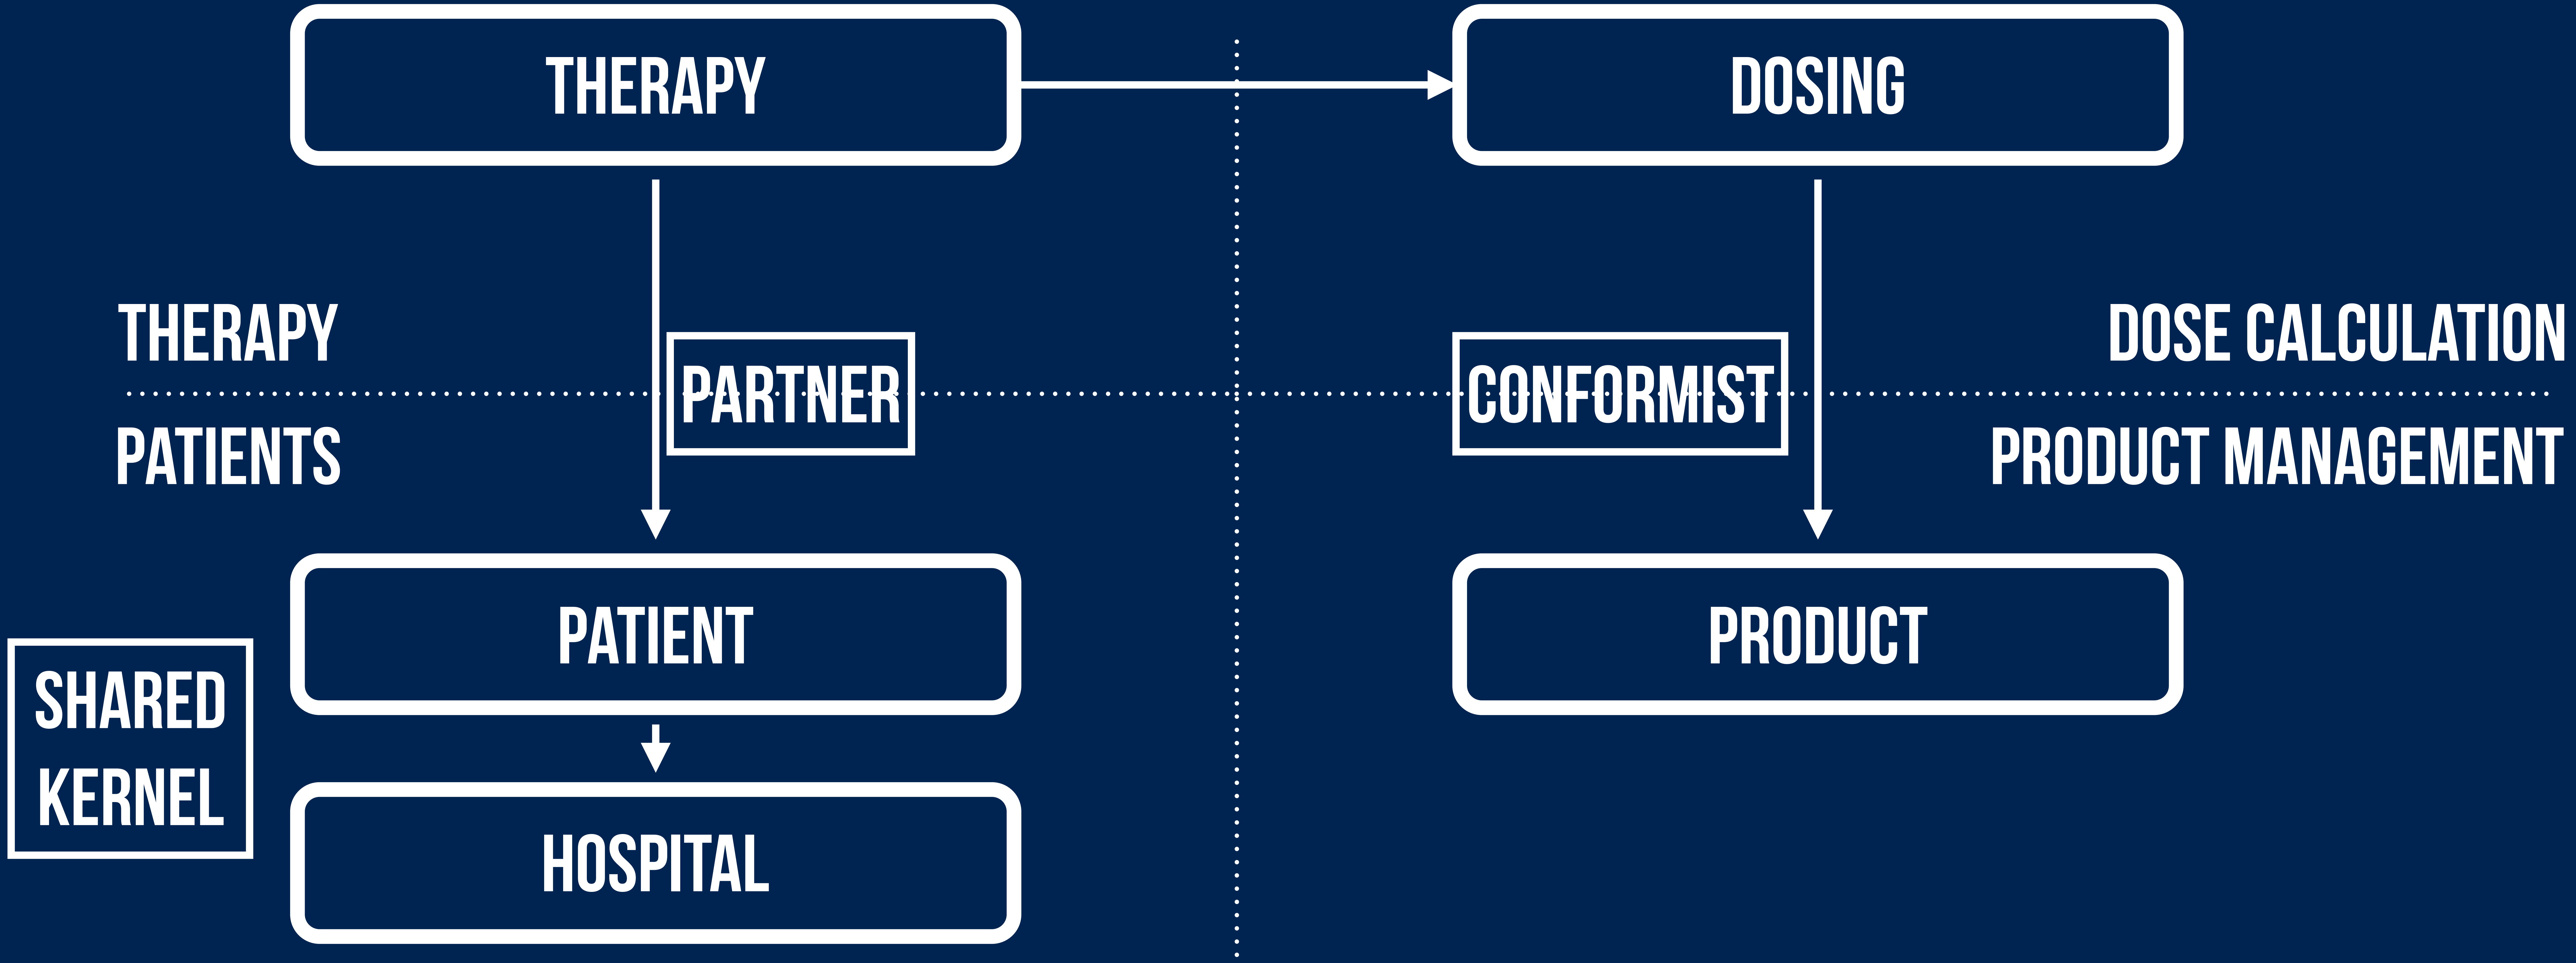
THERAPY
DOSING
THERAPY
PATIENTS
PARTNER
CONFORMIST
DOSE CALCULATION
PRODUCT MANAGEMENT
PATIENT
PRODUCT
SHARED KERNEL
HOSPITAL

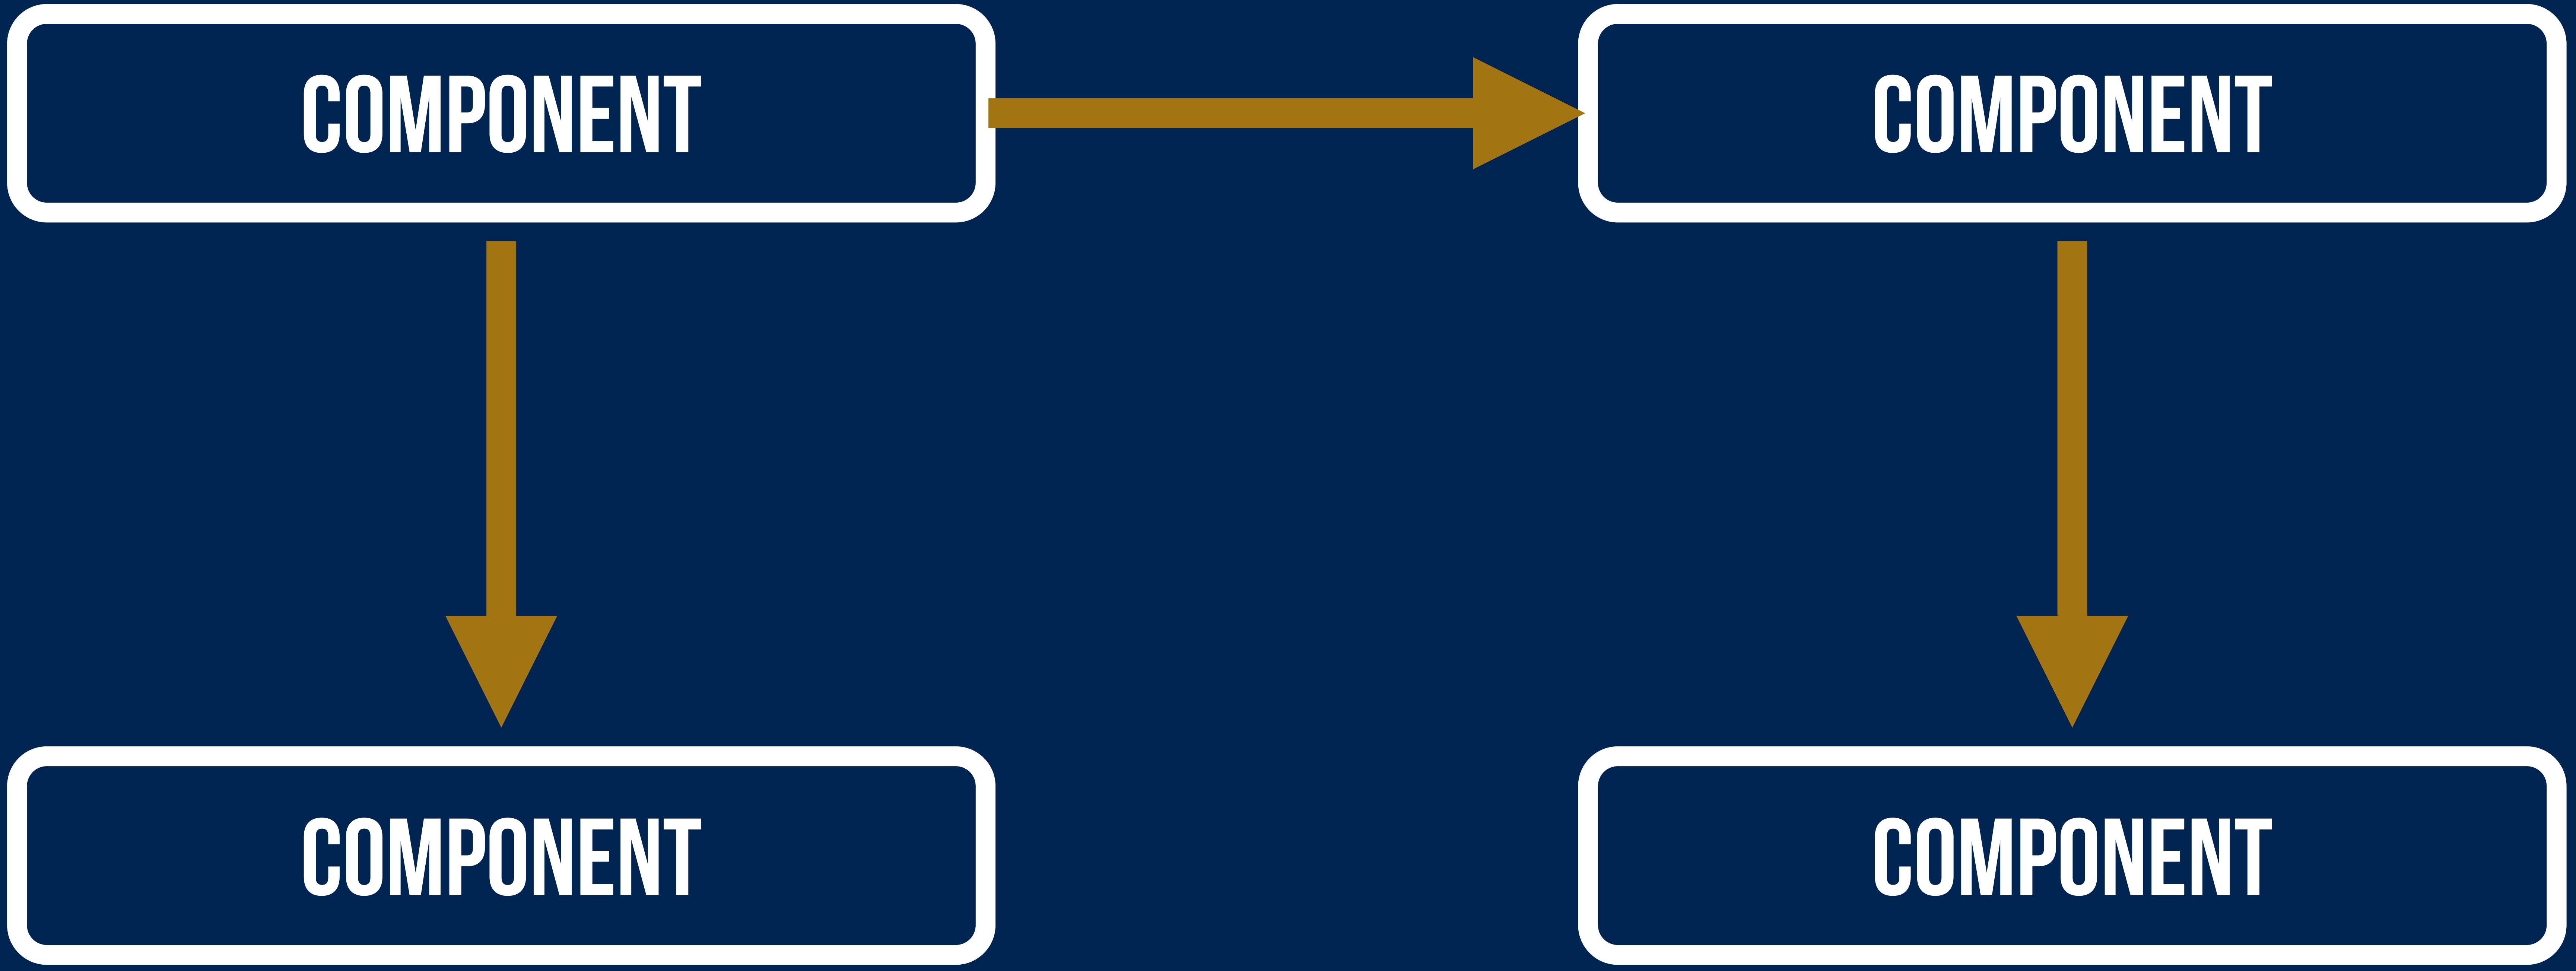
COMPONENT
COMPONENT
COMPONENT
COMPONENT

- ☐ THE NETWORK IS RELIABLE
- ☐ LATENCY IS ZERO
- ☐ BANDWITH IS INFINITE
- ☐ THE NETWORK IS SECURE
- ☐ TOPOLOGY DOESN'T CHANGE
- ☐ THERE IS ONE ADMINISTRATOR
- ☐ TRANSPORT COST IS ZERO
- ☐ THE NETWORK IS HOMOGENEOUS

- ✗ THE NETWORK IS RELIABLE
- ✗ LATENCY IS ZERO
- ✗ BANDWITH IS INFINITE
- ✗ THE NETWORK IS SECURE
- ✗ TOPOLOGY DOESN'T CHANGE
- ✗ THERE IS ONE ADMINISTRATOR
- ✗ TRANSPORT COST IS ZERO
- ✗ THE NETWORK IS HOMOGENEOUS

THE NETWORK FALLACIES

HTTP://WWW.RGOARCHITECTS.COM/FILES/FALLACIES.PDF

TIMEOUT

```java
public class MessageReceiver {

    private Session session;
    private Destination destination;

    public void doReceive() throws Exception{
        MessageConsumer consumer =
             session.createConsumer(destination);

        consumer.receive();
    }
}
```

```java
public class MessageReceiver {

    private Session session;
    private Destination destination;

    public void doReceive() throws Exception{
        MessageConsumer consumer =
             session.createConsumer(destination);

        consumer.receive(20L);
    }
}
```

```java
public class HttpReceiver {

    public String getResource() throws IOException {
        URL url = new URL("http://www.google.de");
        InputStream inputStream = url.openStream();
        return “…”;
    }

}
```

```java
public class HttpReceiver {

    public String getResource() throws IOException {
        URL url = new URL("http://www.google.de");
        URLConnection urlConn = url.openConnection();
        urlConn.setConnectTimeout(10);
        urlConn.setReadTimeout(10);
        InputStream inputStream = url.getInputStream();
        return “…”;
    }

}
```

```java
public class DataSourceConfig {

    public void DataSource setupDataSource(){
        BasicDataSource basicDataSource =
                  new BasicDataSource();
        basicDataSource.setMaxWait(30L);
    }
}
```

ADD LATENCY

```
tc qdisc add dev eth0 root latency delay 1000ms 500ms
```

CORRUPT PACKAGES

```
tc qdisc add dev eth0 root netem corrupt 5%
```

DROP PACKAGES

```
tc qdisc add dev eth0 root netem loss 7% 25%
```

BLOCK DNS

```
iptables -A INPUT -p tcp -m tcp --dport 53 -j DROP
```

PATTERNS

STABILITY

CAPACITY

TRANSPARENCY

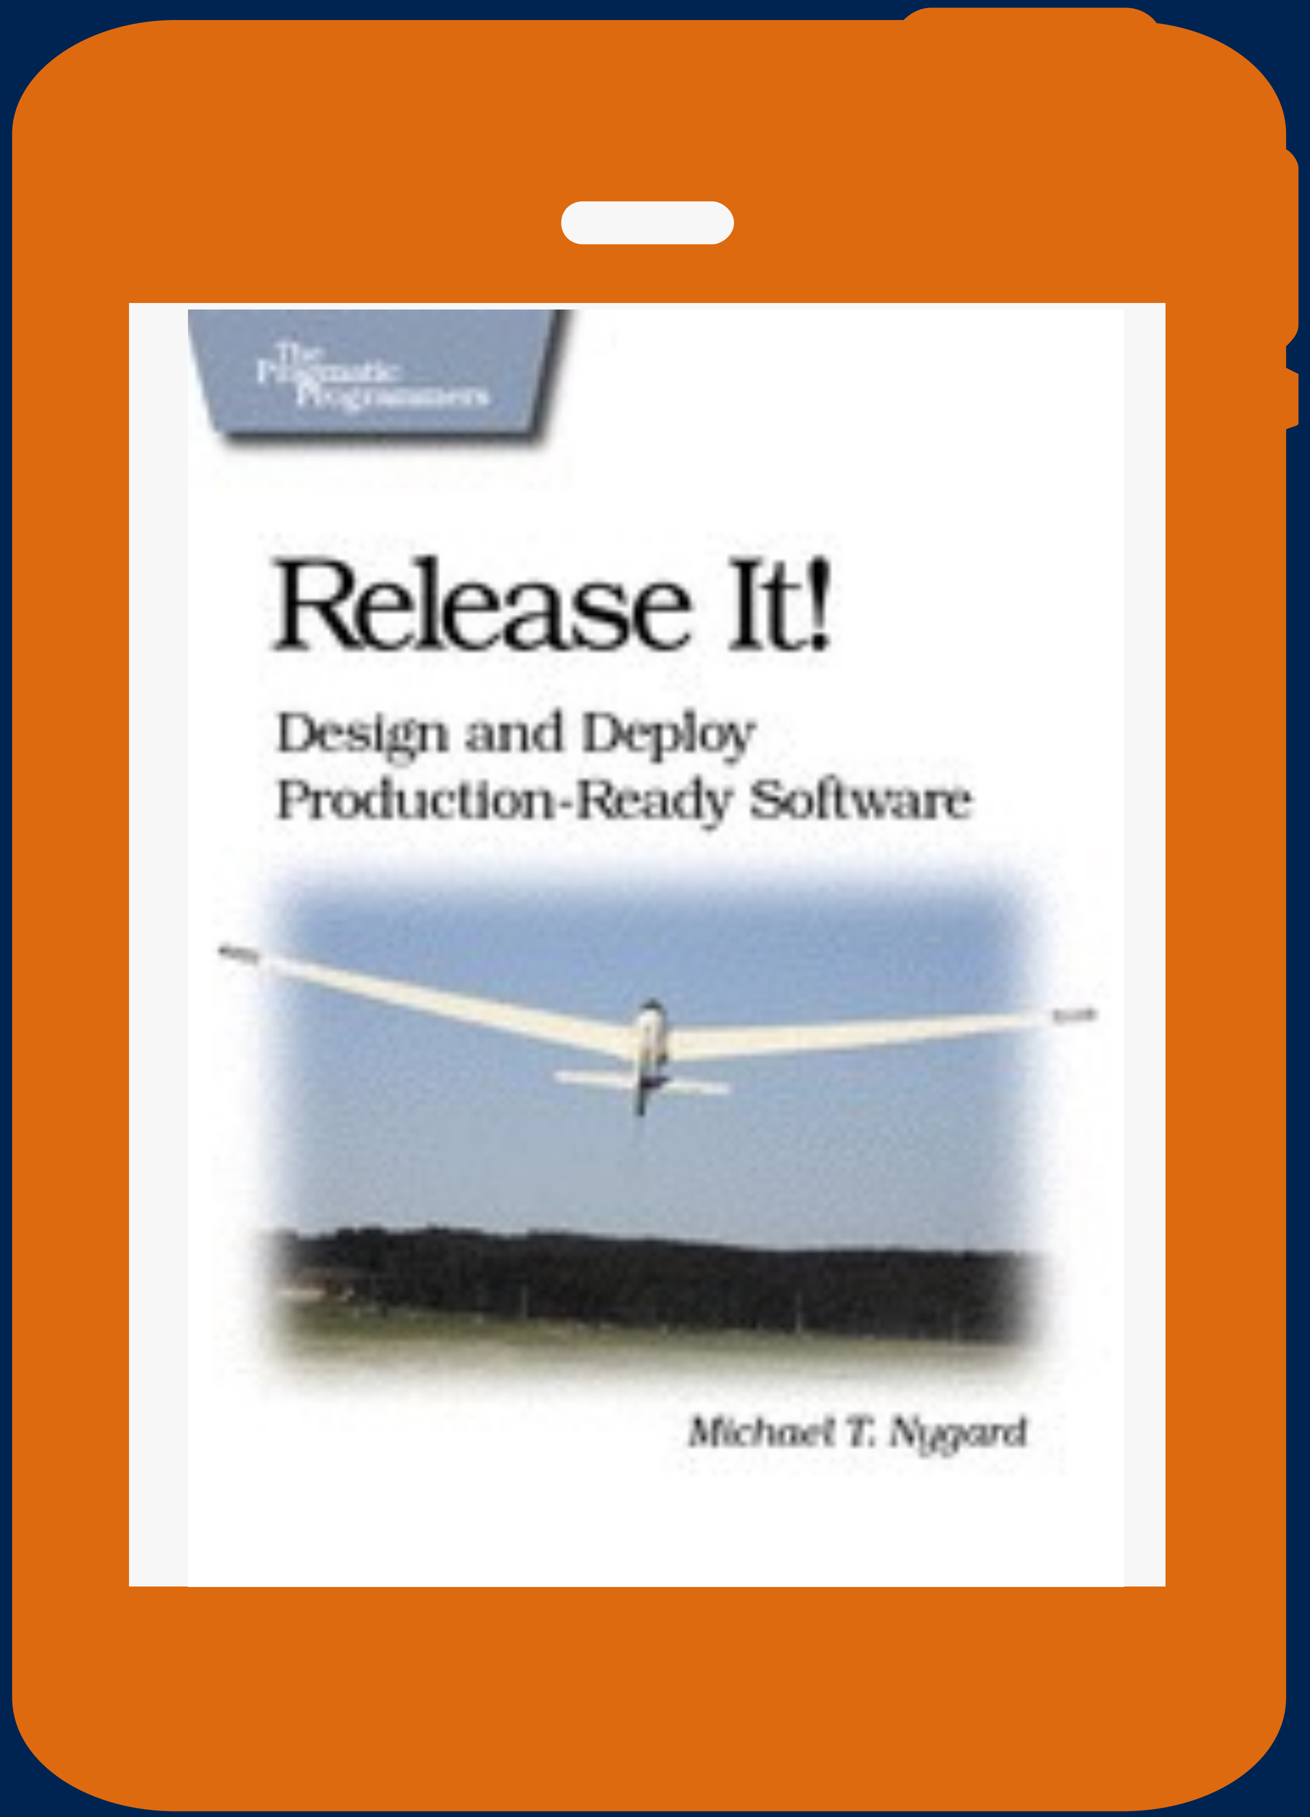
The Pragmatic Programmers
Release It!
Design and Deploy Production-Ready Software
Michael T. Nygard

37%

INTERNET

TRAFFIC

NETFLIX

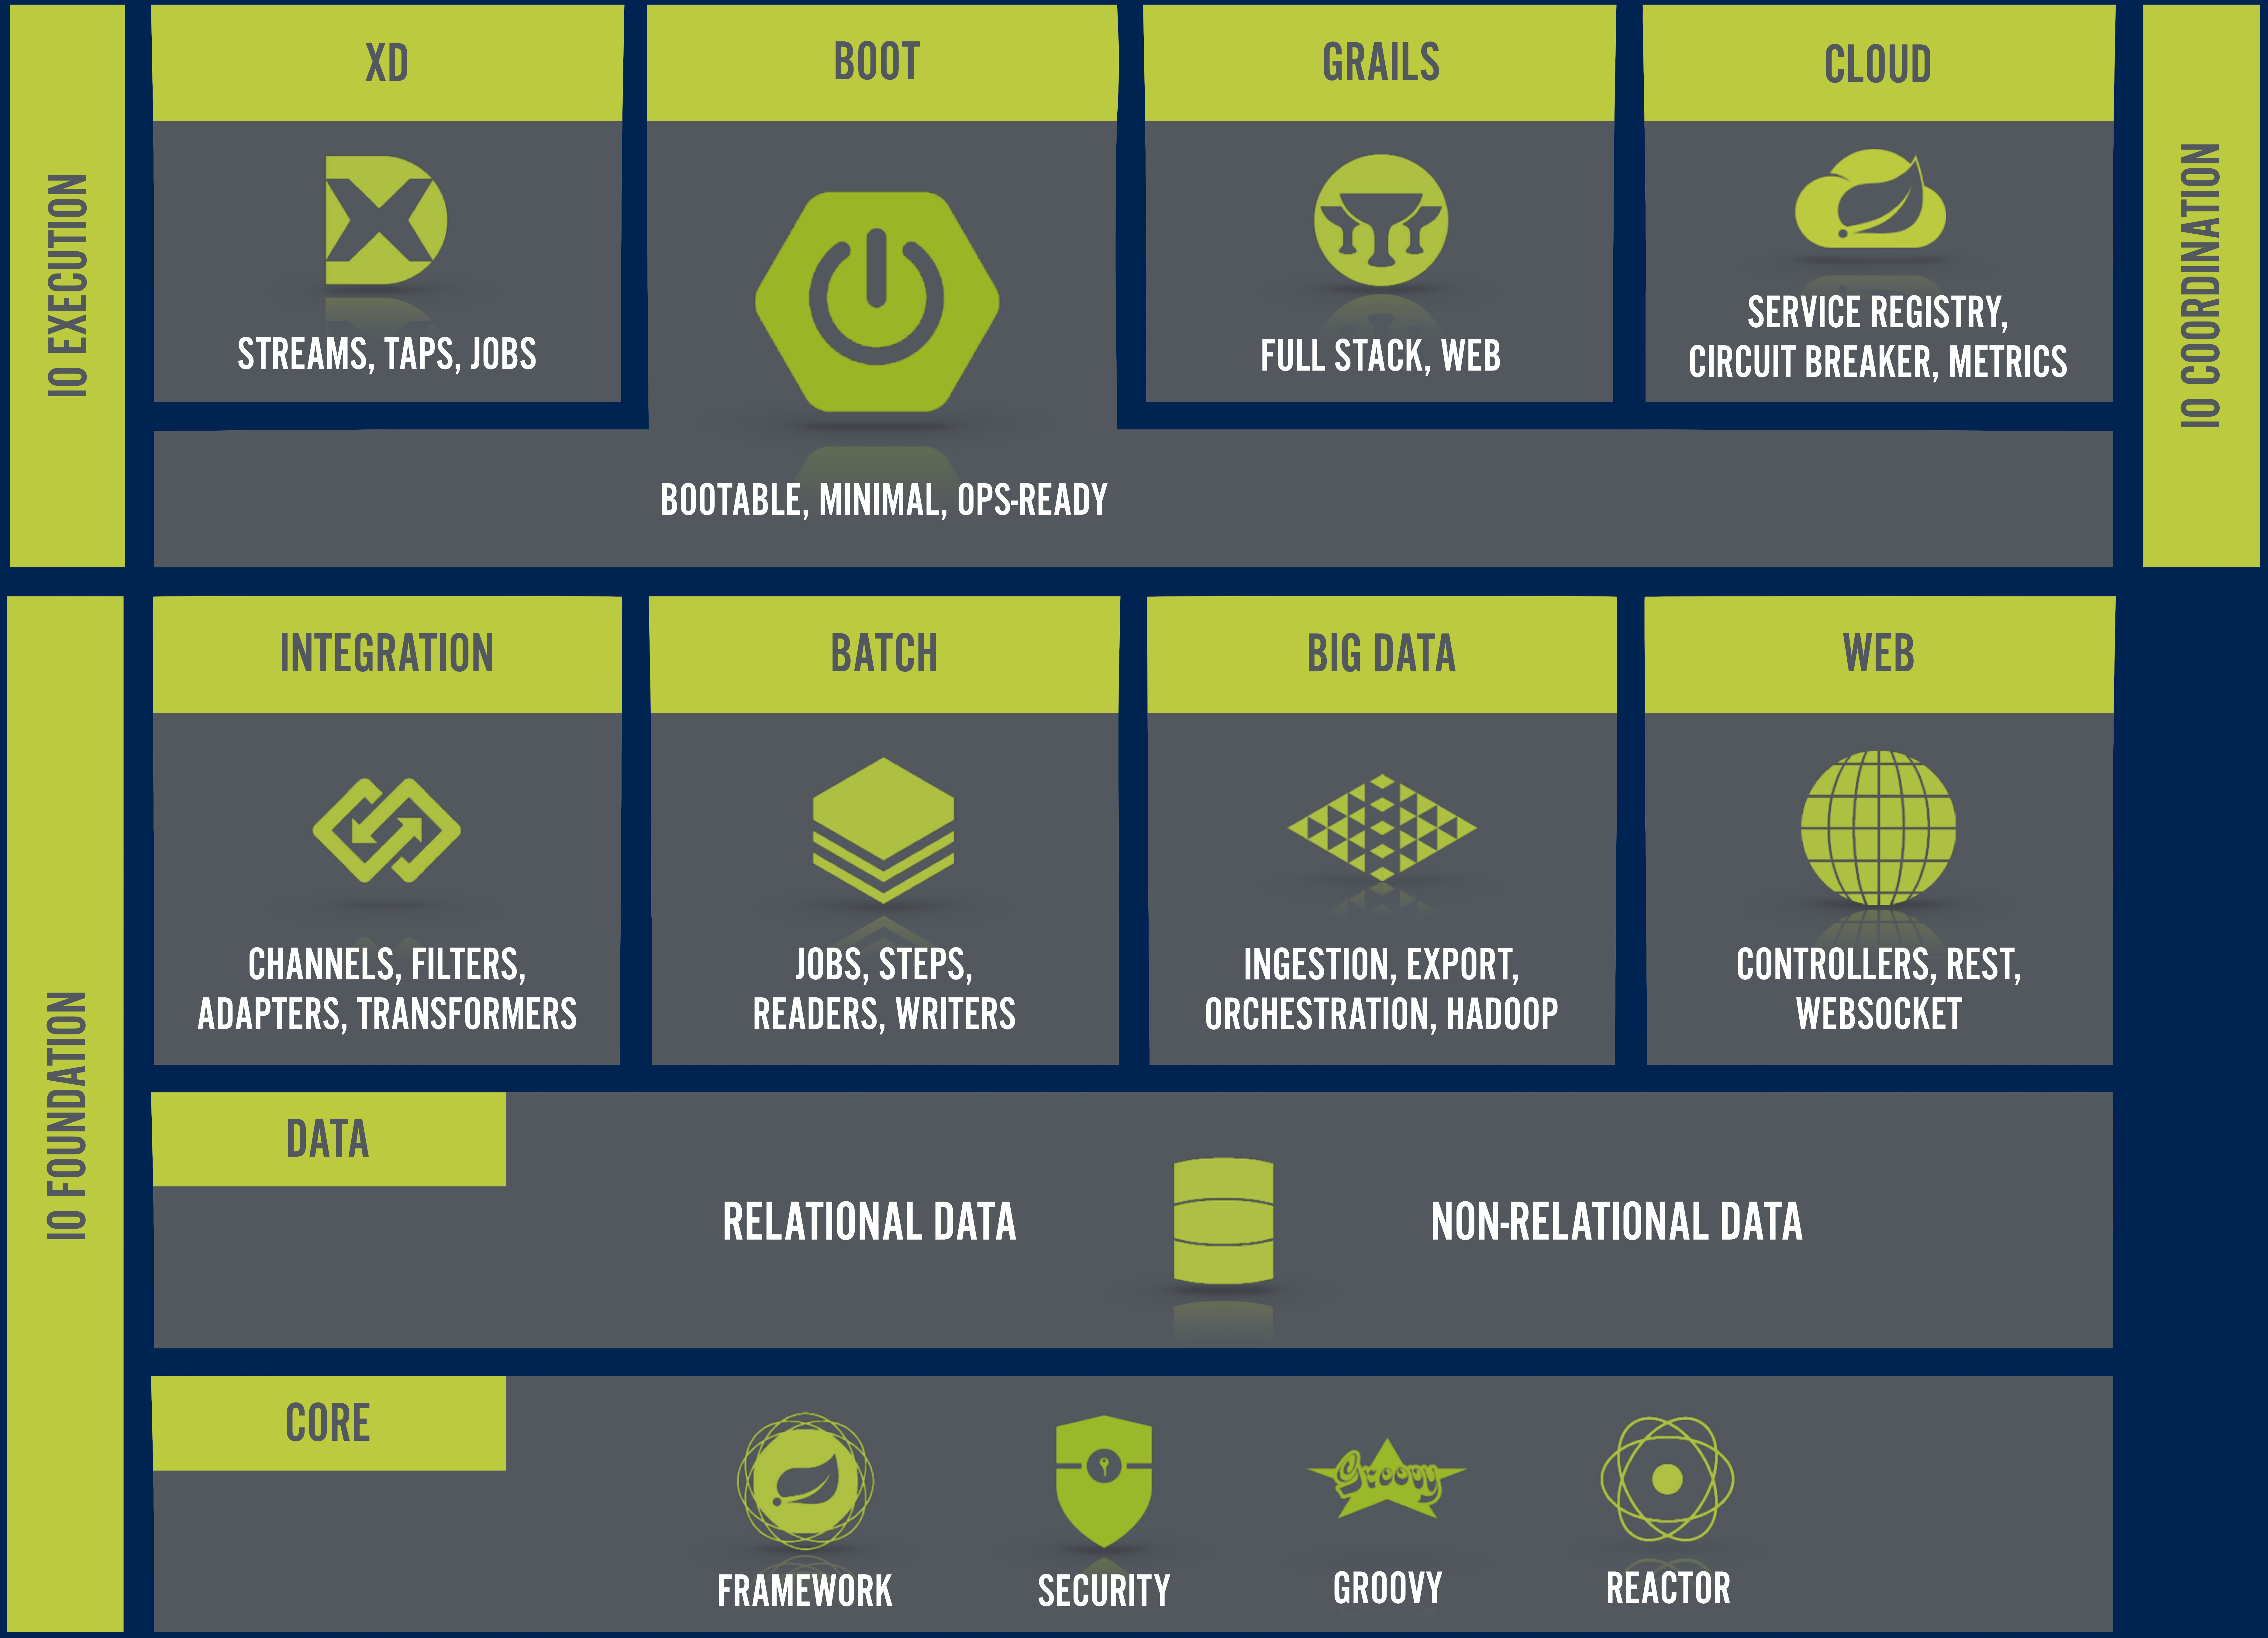
IO EXECUTION
XD
STREAMS, TAPS, JOBS
BOOT
BOOTABLE, MINIMAL, OPS-READY
GRAILS
FULL STACK, WEB
CLOUD
SERVICE REGISTRY, CIRCUIT BREAKER, METRICS
IO COORDINATION
IO FOUNDATION
INTEGRATION
CHANNELS, FILTERS, ADAPTERS, TRANSFORMERS
BATCH
JOBS, STEPS, READERS, WRITERS
BIG DATA
INGESTION, EXPORT, ORCHESTRATION, HADOOP
WEB
CONTROLLERS, REST, WEBSOCKET
DATA
RELATIONAL DATA
NON-RELATIONAL DATA
CORE
FRAMEWORK
SECURITY
GROOVY
REACTOR

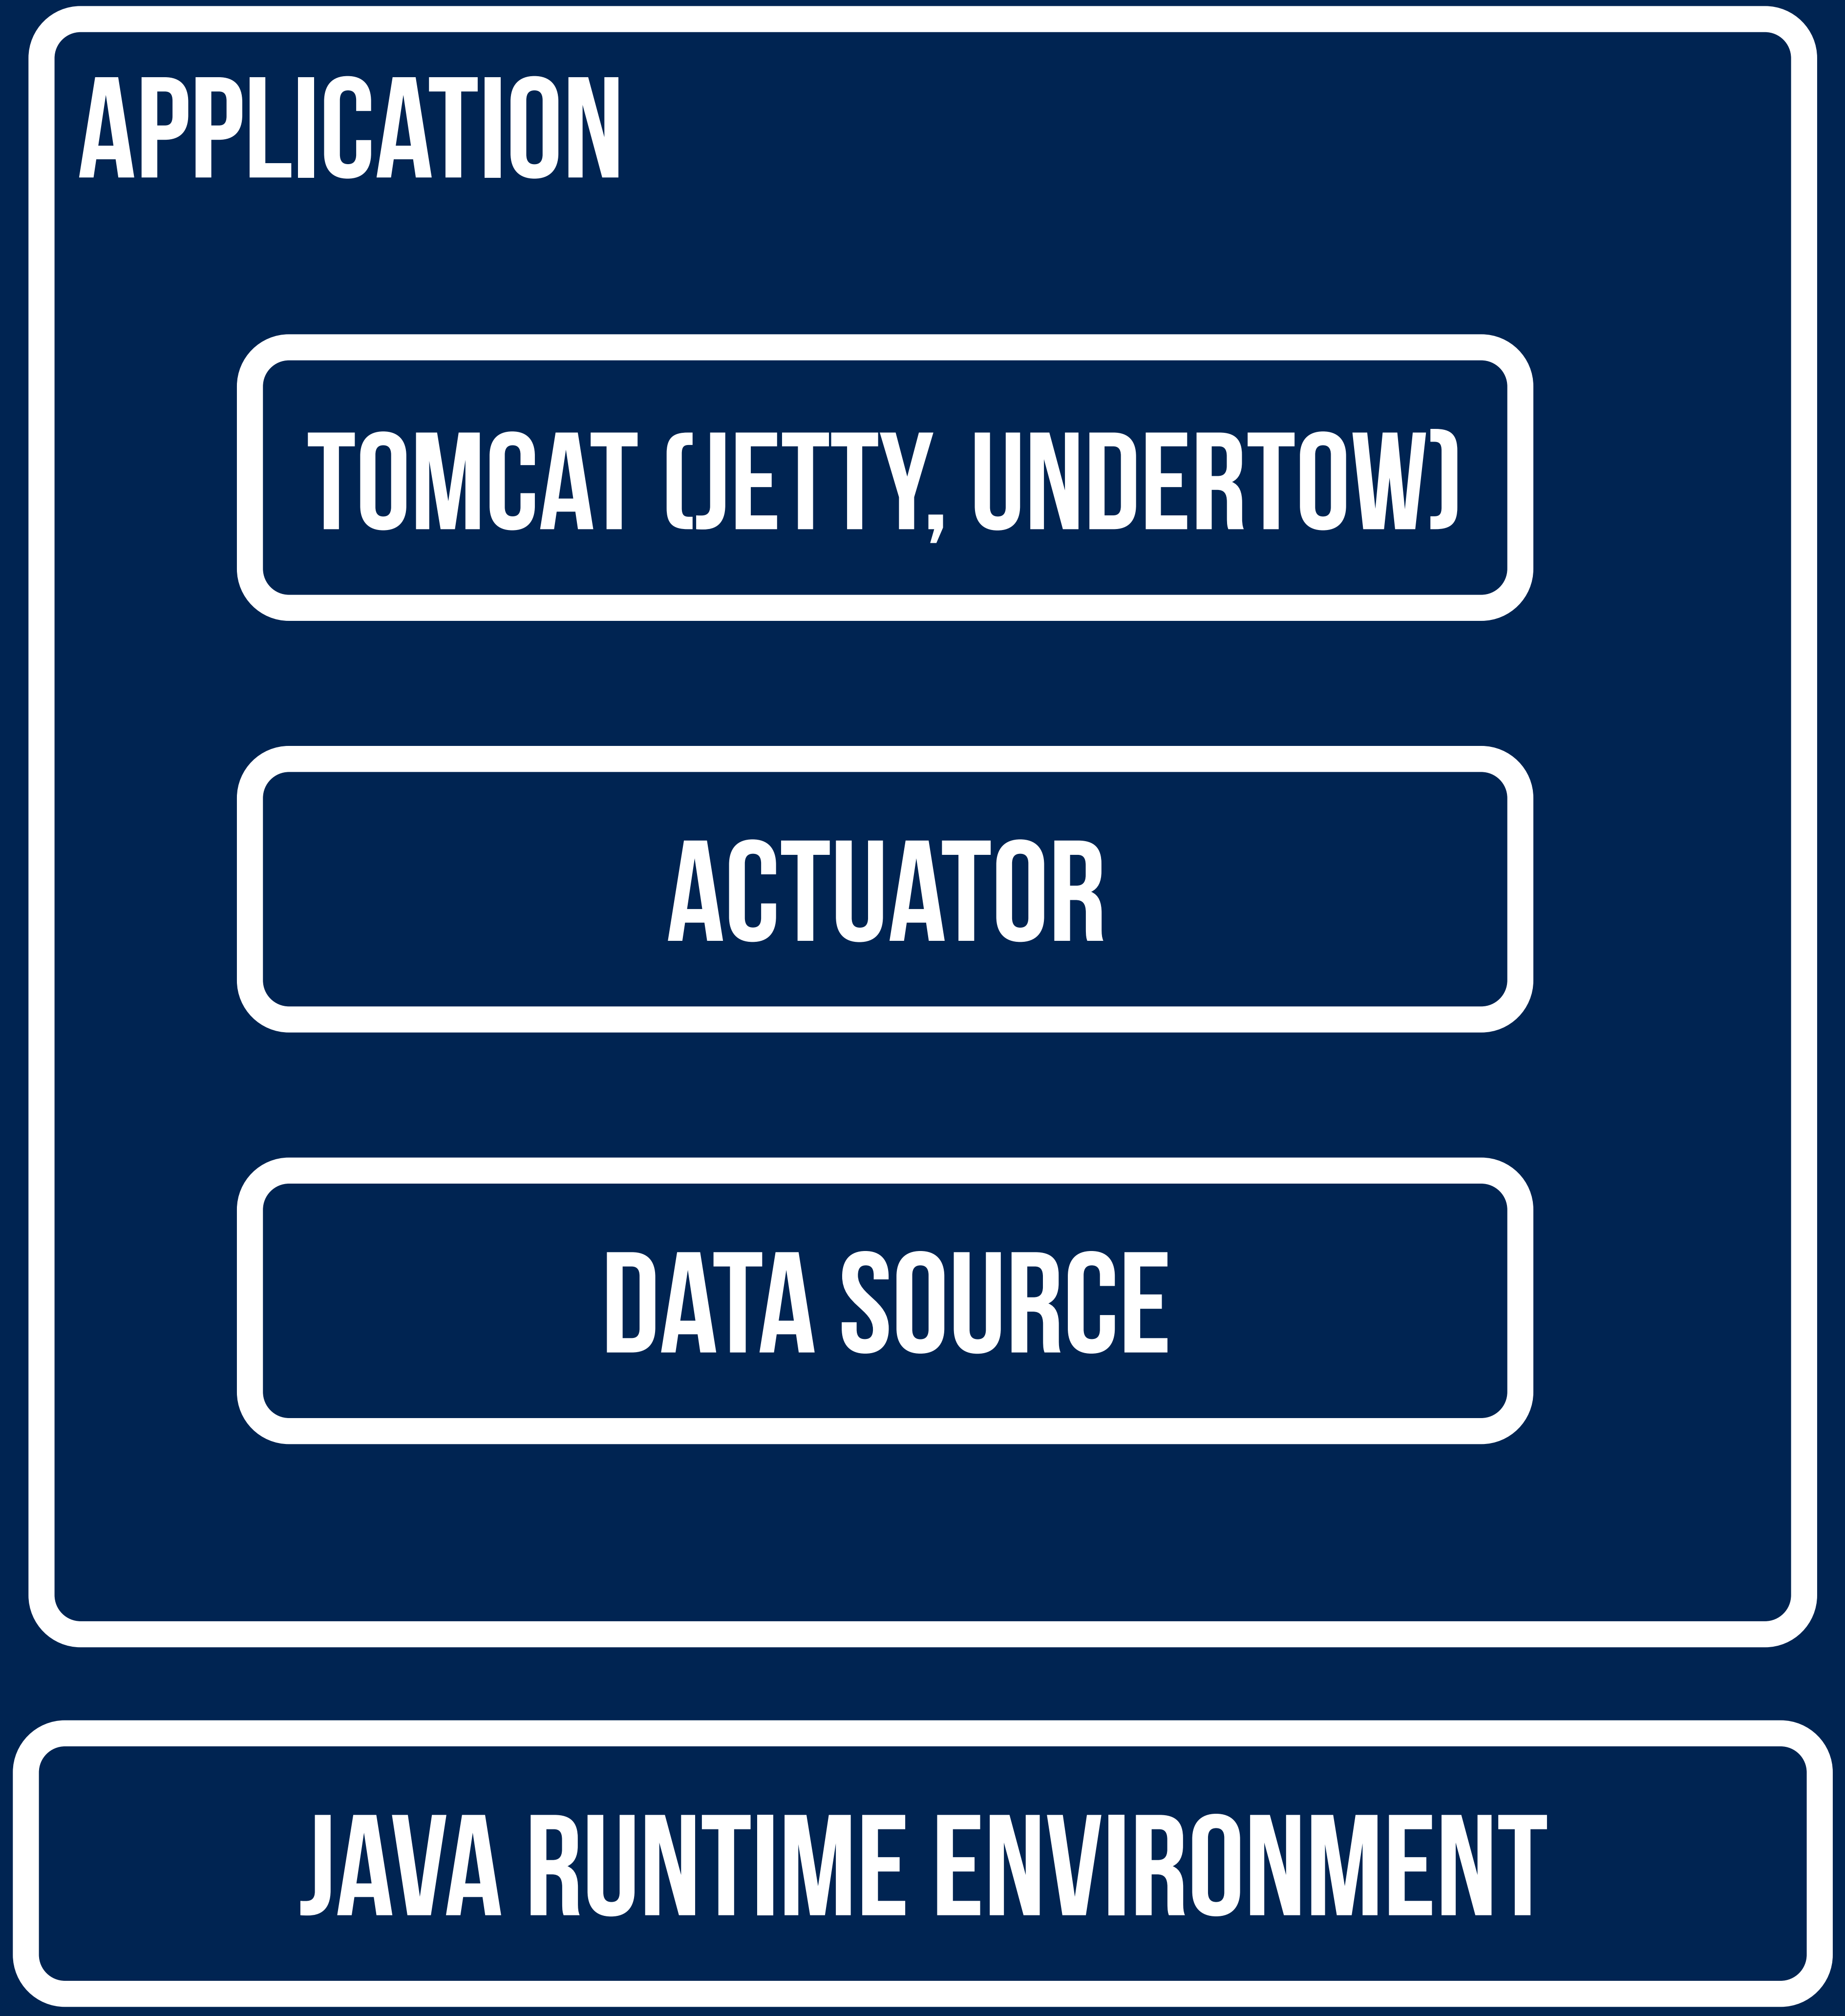

```
JAVA -JAR MYAPPLICATION.JAR
```

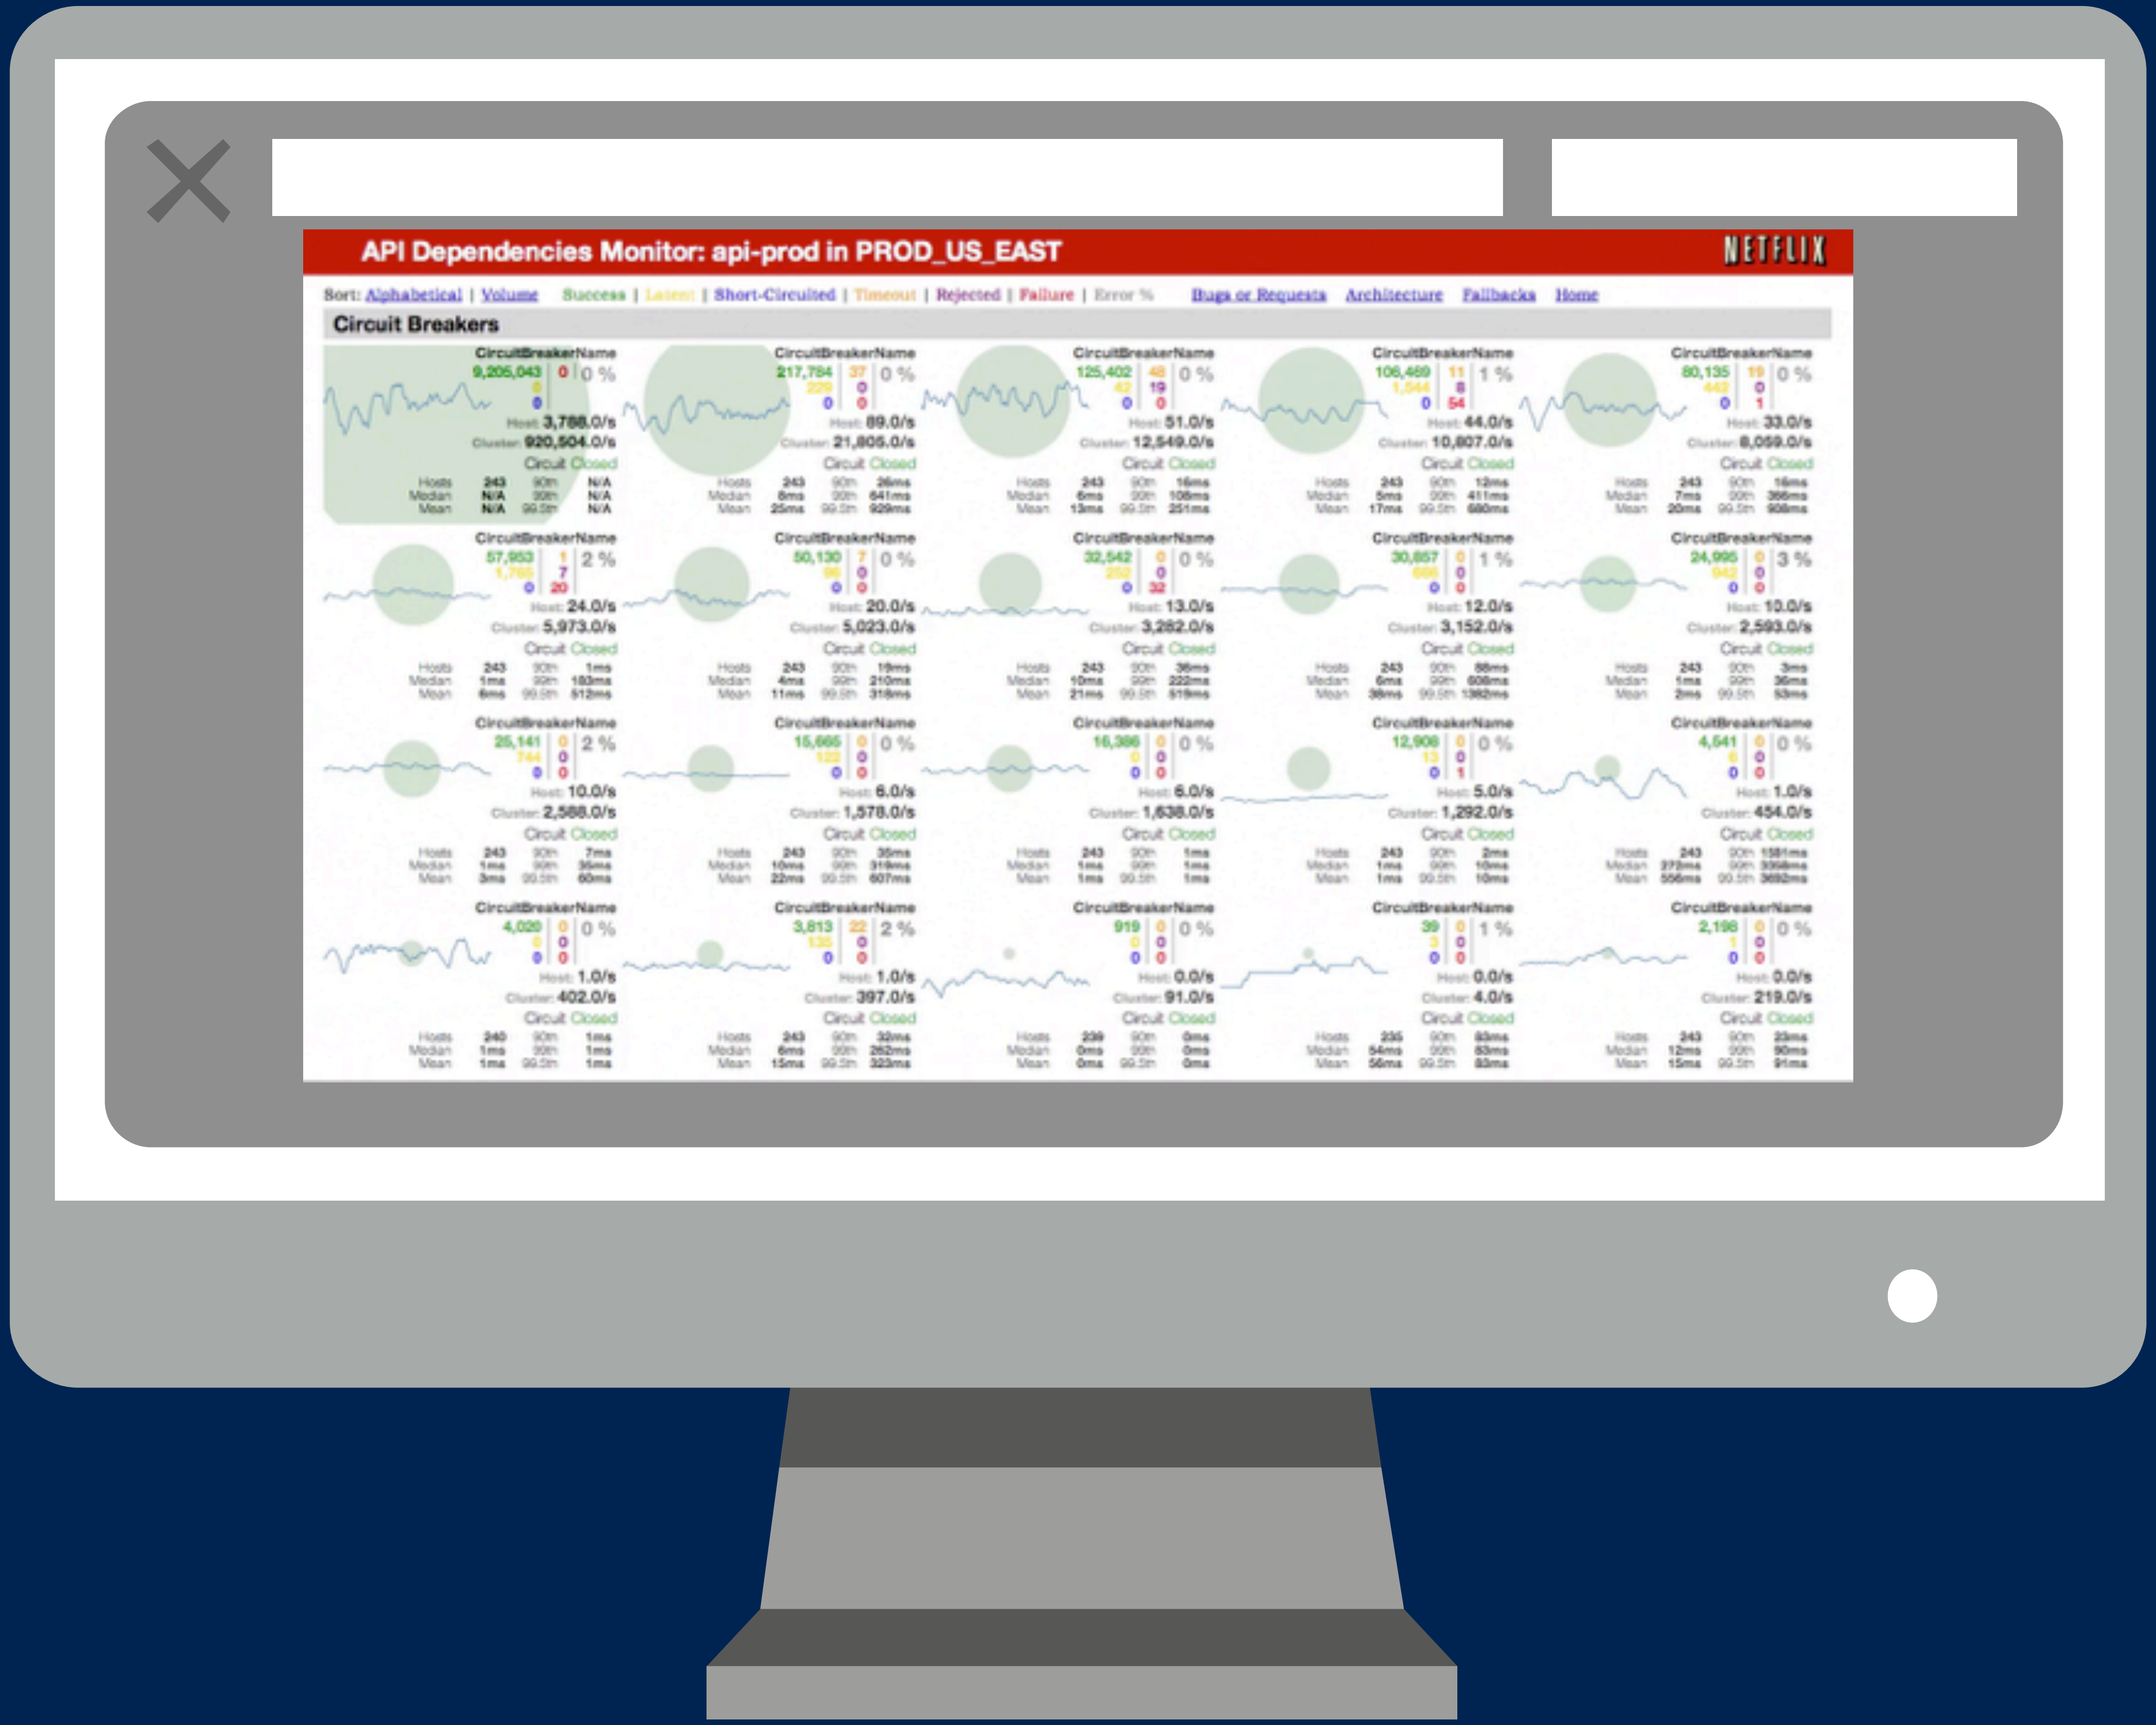
API Dependencies Monitor: api-prod in PROD_US_EAST
NETFLIX
Sort: Alphabetical | Volume | Success | Latent | Short-Circuited | Timeout | Rejected | Failure | Error %
Bugs or Requests
Architecture
Fallbacks
Home
Circuit Breakers

CIRCUIT BREAKER

FRAGILE	ROBUST	ANTI-FRAGILE
NO TIMEOUT	TIMEOUT	CIRCUIT-BREAKER

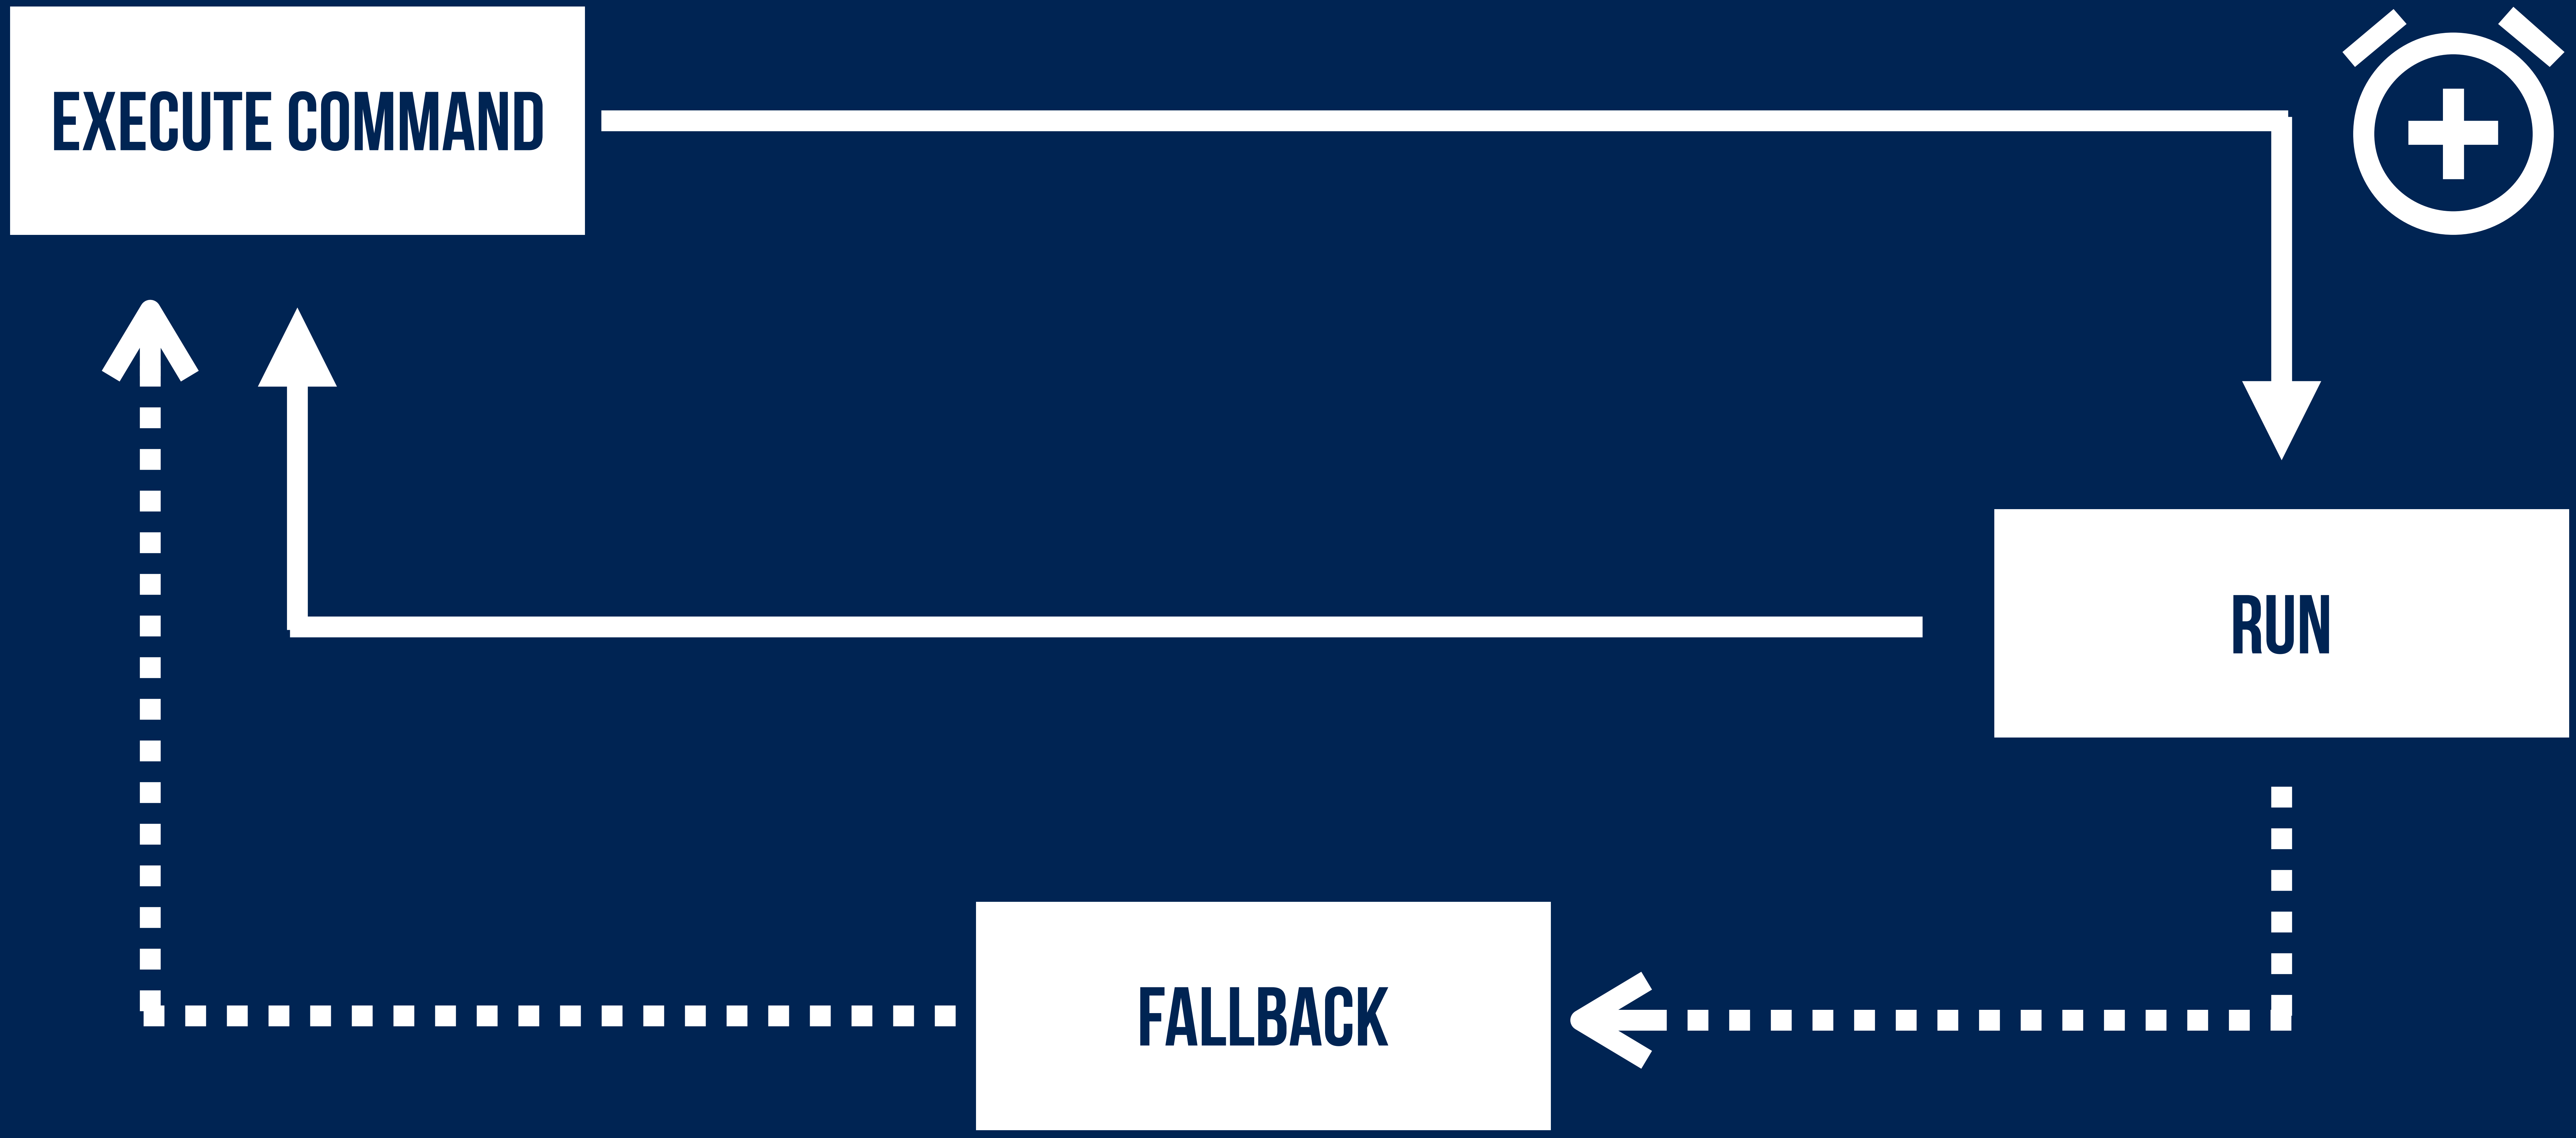
EXECUTE COMMAND
RUN
FALLBACK

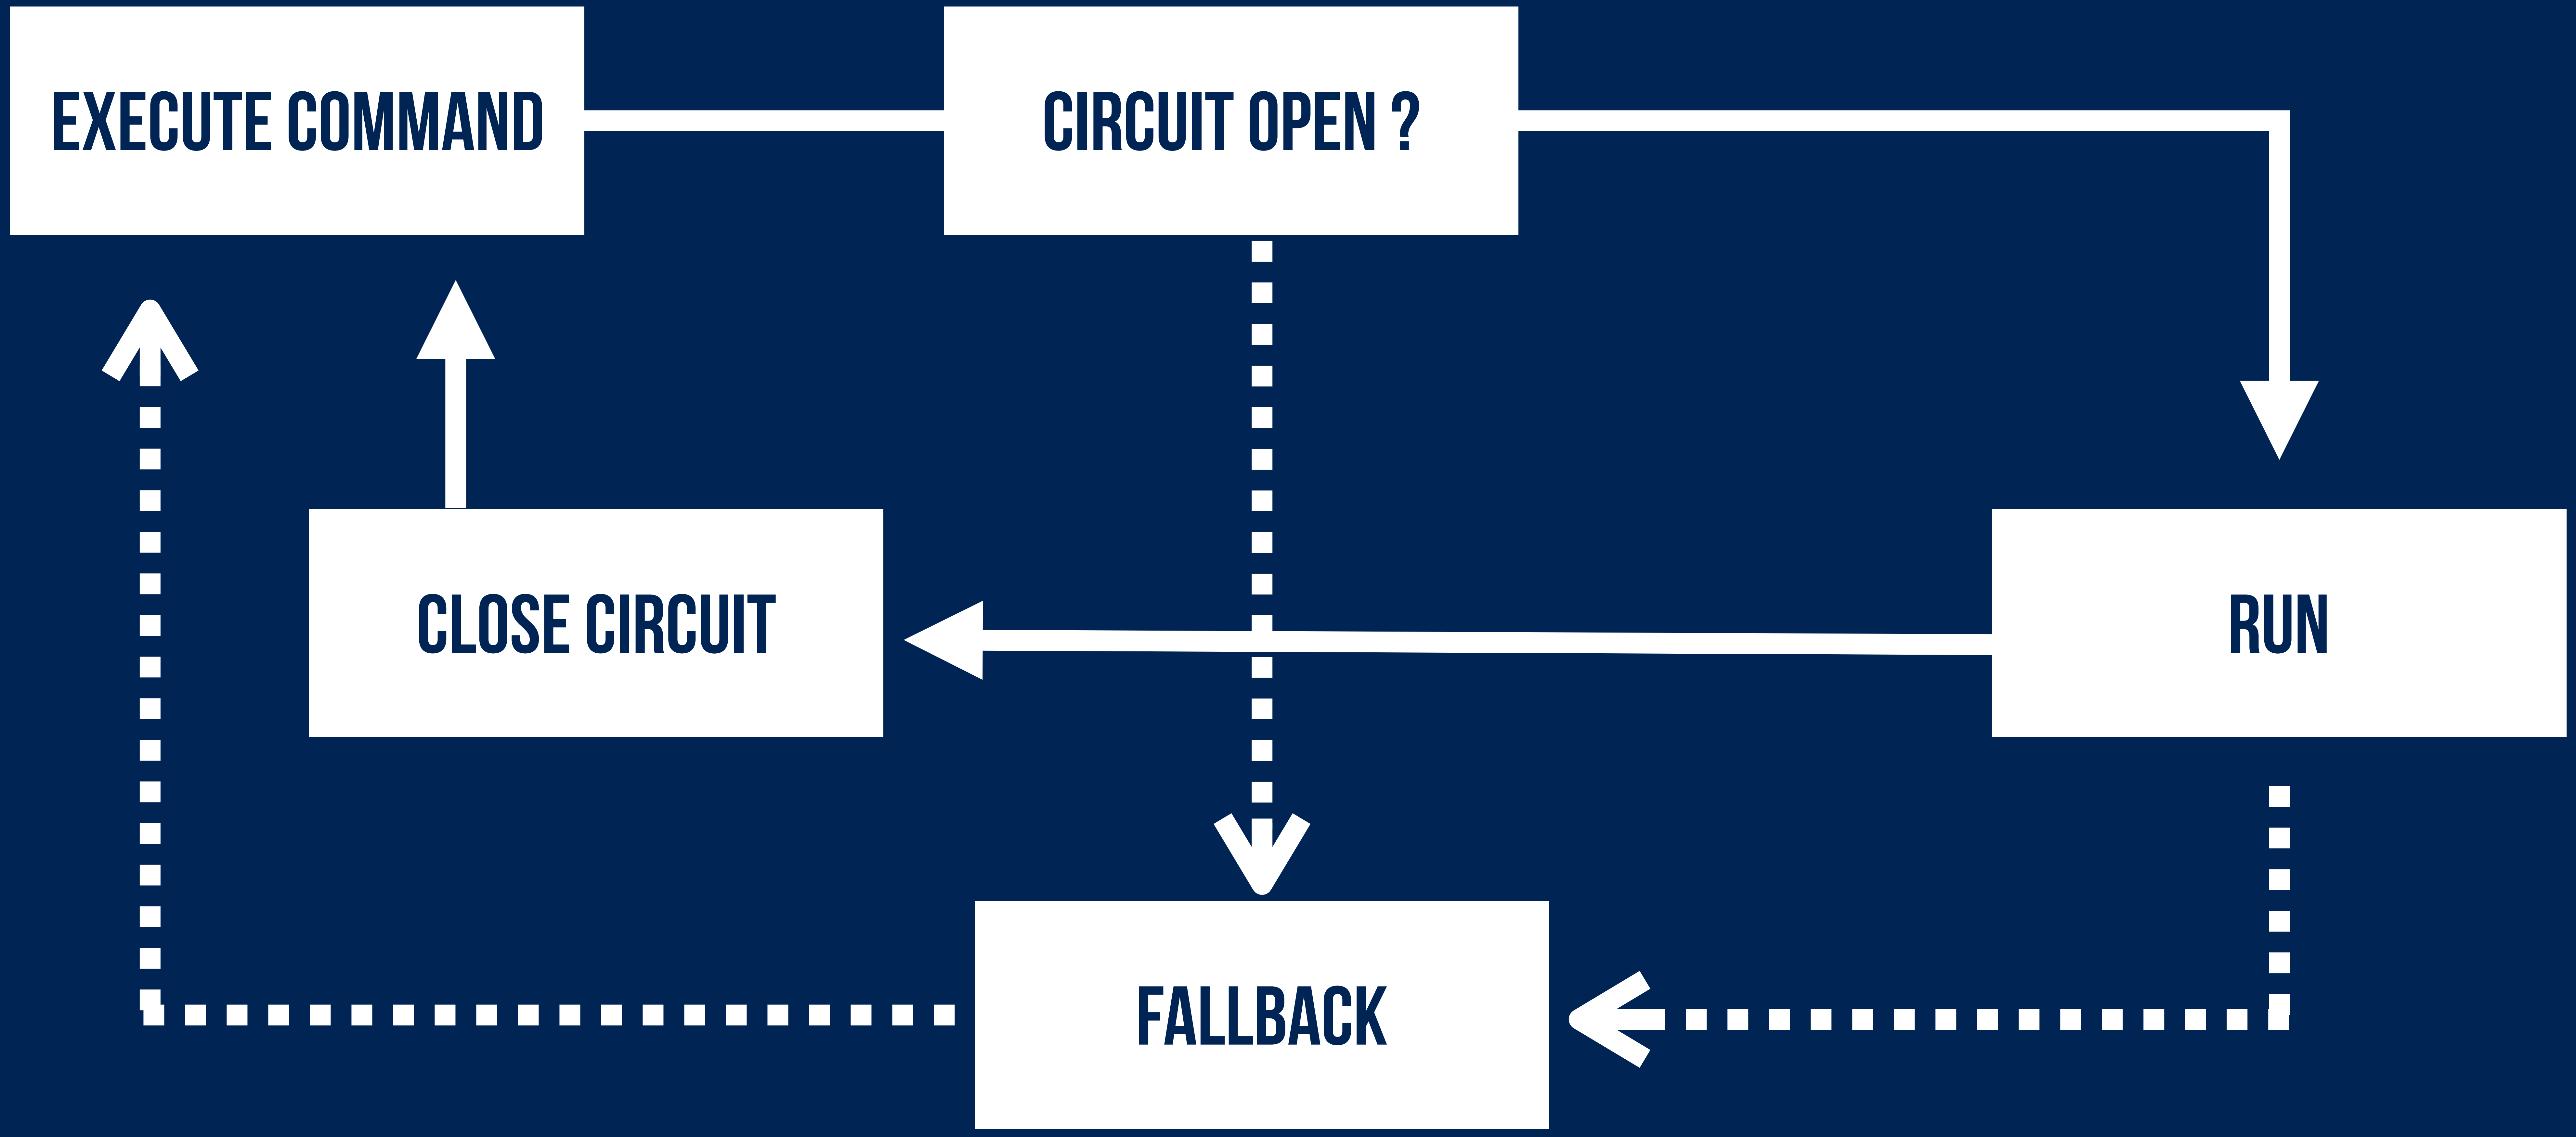
EXECUTE COMMAND
CIRCUIT OPEN ?
CLOSE CIRCUIT
RUN
FALLBACK

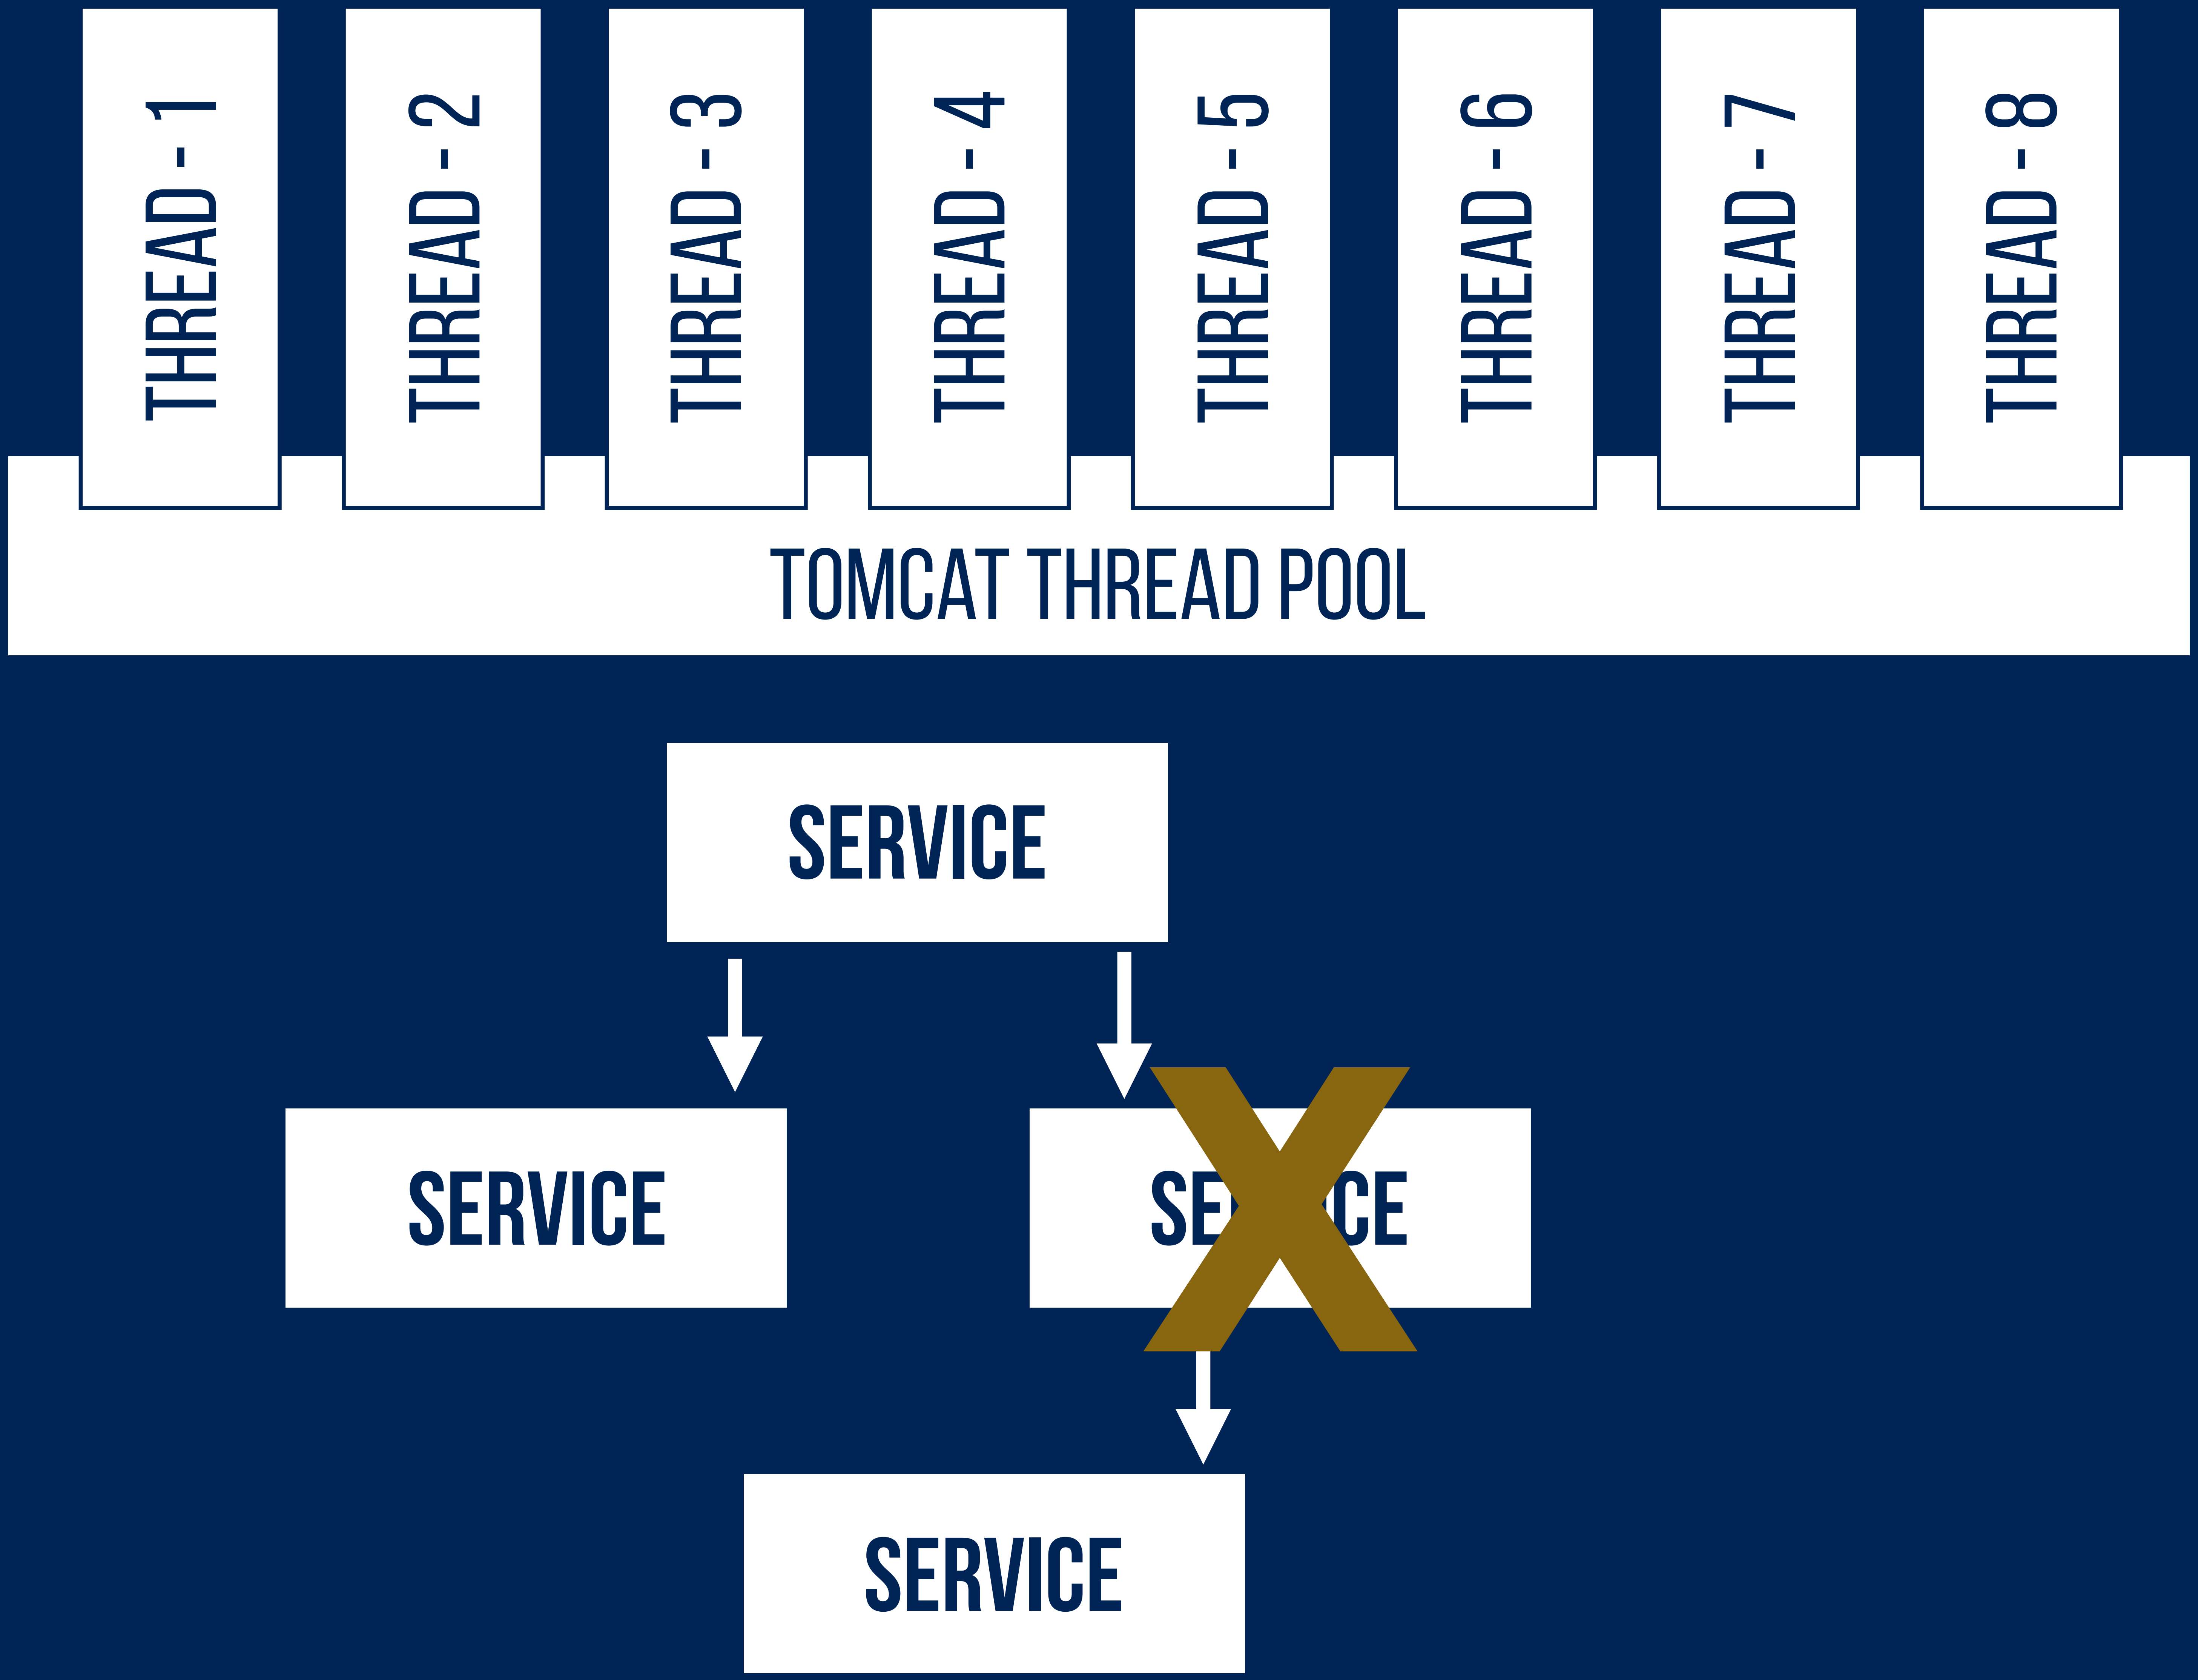
THREAD - 1
THREAD - 2
THREAD - 3
THREAD - 4
THREAD - 5
THREAD - 6
THREAD - 7
THREAD - 8
TOMCAT THREAD POOL
SERVICE
SERVICE
SERVICE
X
SERVICE

■ COMMAND THREAD

■ COMMAND THREAD

■ COMMAND THREAD

```
TRACE --- [HTTP-NIO-AUTO-1-EXEC-10] OUTSIDE COMMAND
TRACE --- [HYSTRIX-RESTCURRENCYEXCHANGE-10] INSIDE COMMAND
```

THREAD LOCALS

```java
@SpringCloudApplication
public class SearchGateway {

  @HystrixCommand(fallbackMethod = "fallback")
  public List<SearchHit> search(String query) {
    return …;
  }
  public List<SearchHit> fallback() {
    return Collections.emptyList();
  }
}
```

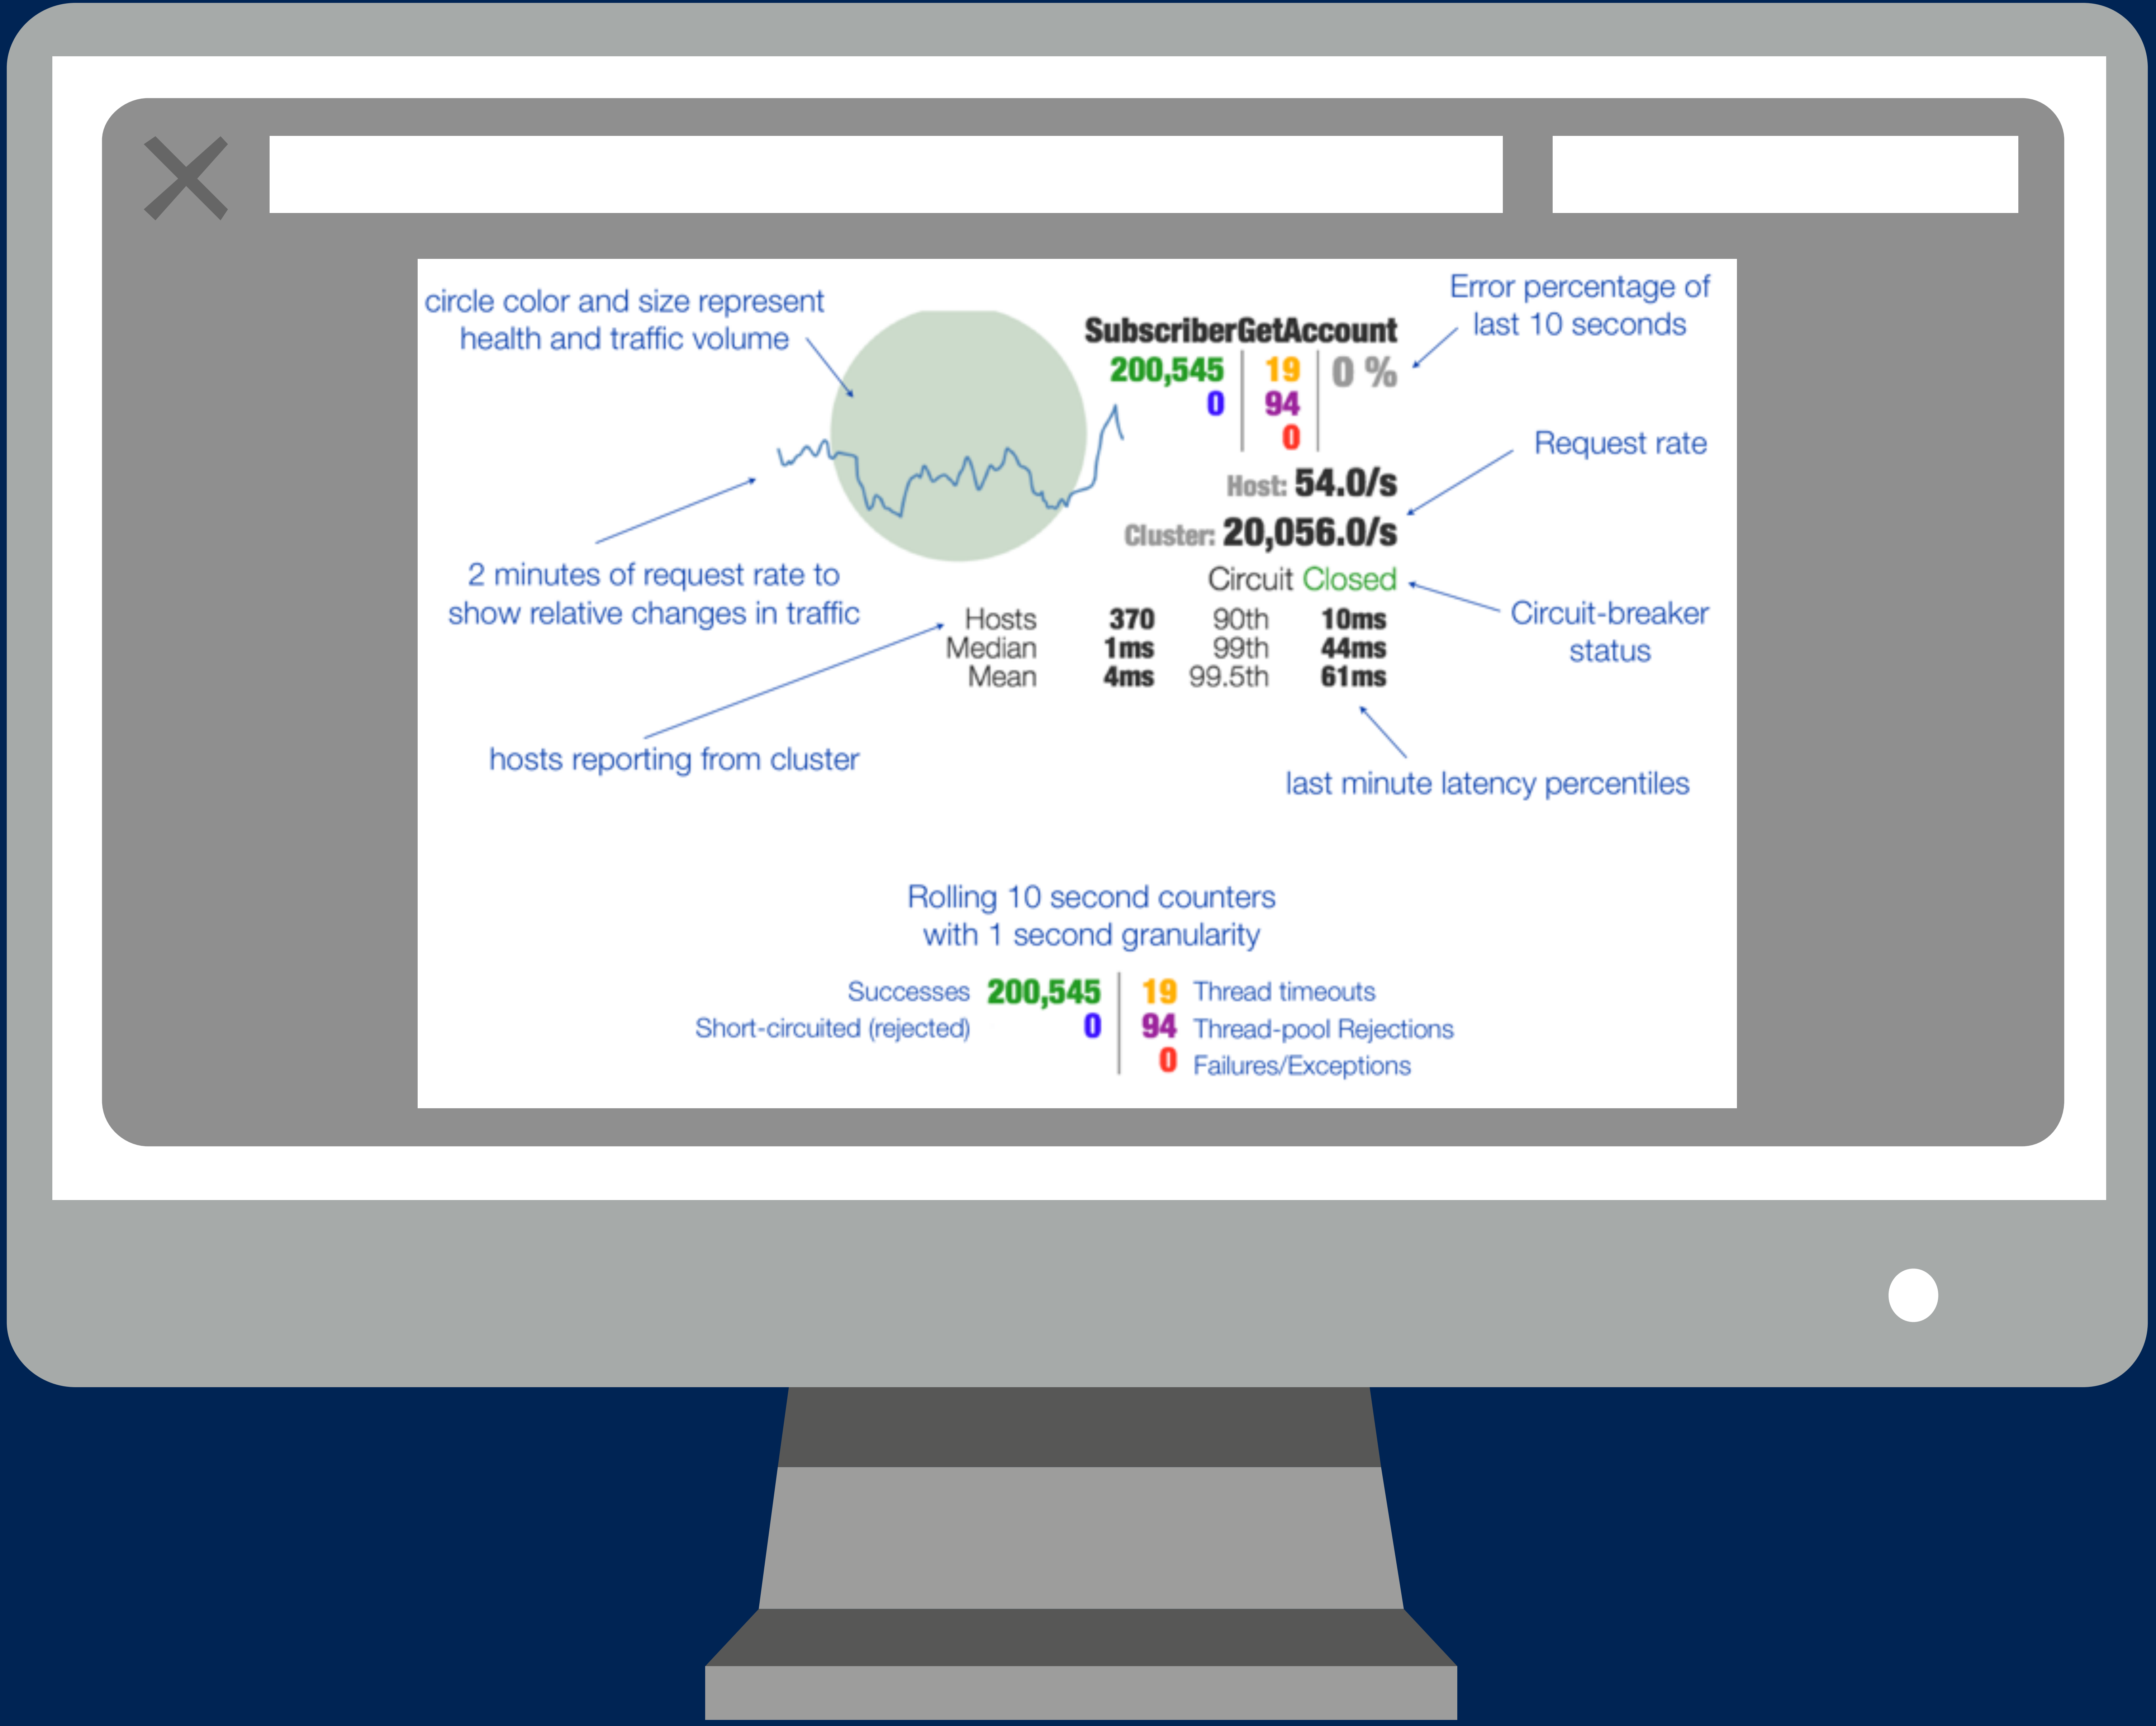
circle color and size represent health and traffic volume
SubscriberGetAccount
200,545
0
19
94
0
0 %
Error percentage of last 10 seconds
Request rate
Host: 54.0/s
Cluster: 20,056.0/s
Circuit Closed
Circuit-breaker status
2 minutes of request rate to show relative changes in traffic
Hosts 370
Median 1ms
Mean 4ms
90th 10ms
99th 44ms
99.5th 61ms
hosts reporting from cluster
last minute latency percentiles
Rolling 10 second counters with 1 second granularity
Successes 200,545
Short-circuited (rejected) 0
19 Thread timeouts
94 Thread-pool Rejections
0 Failures/Exceptions

SERVICE REGISTRY

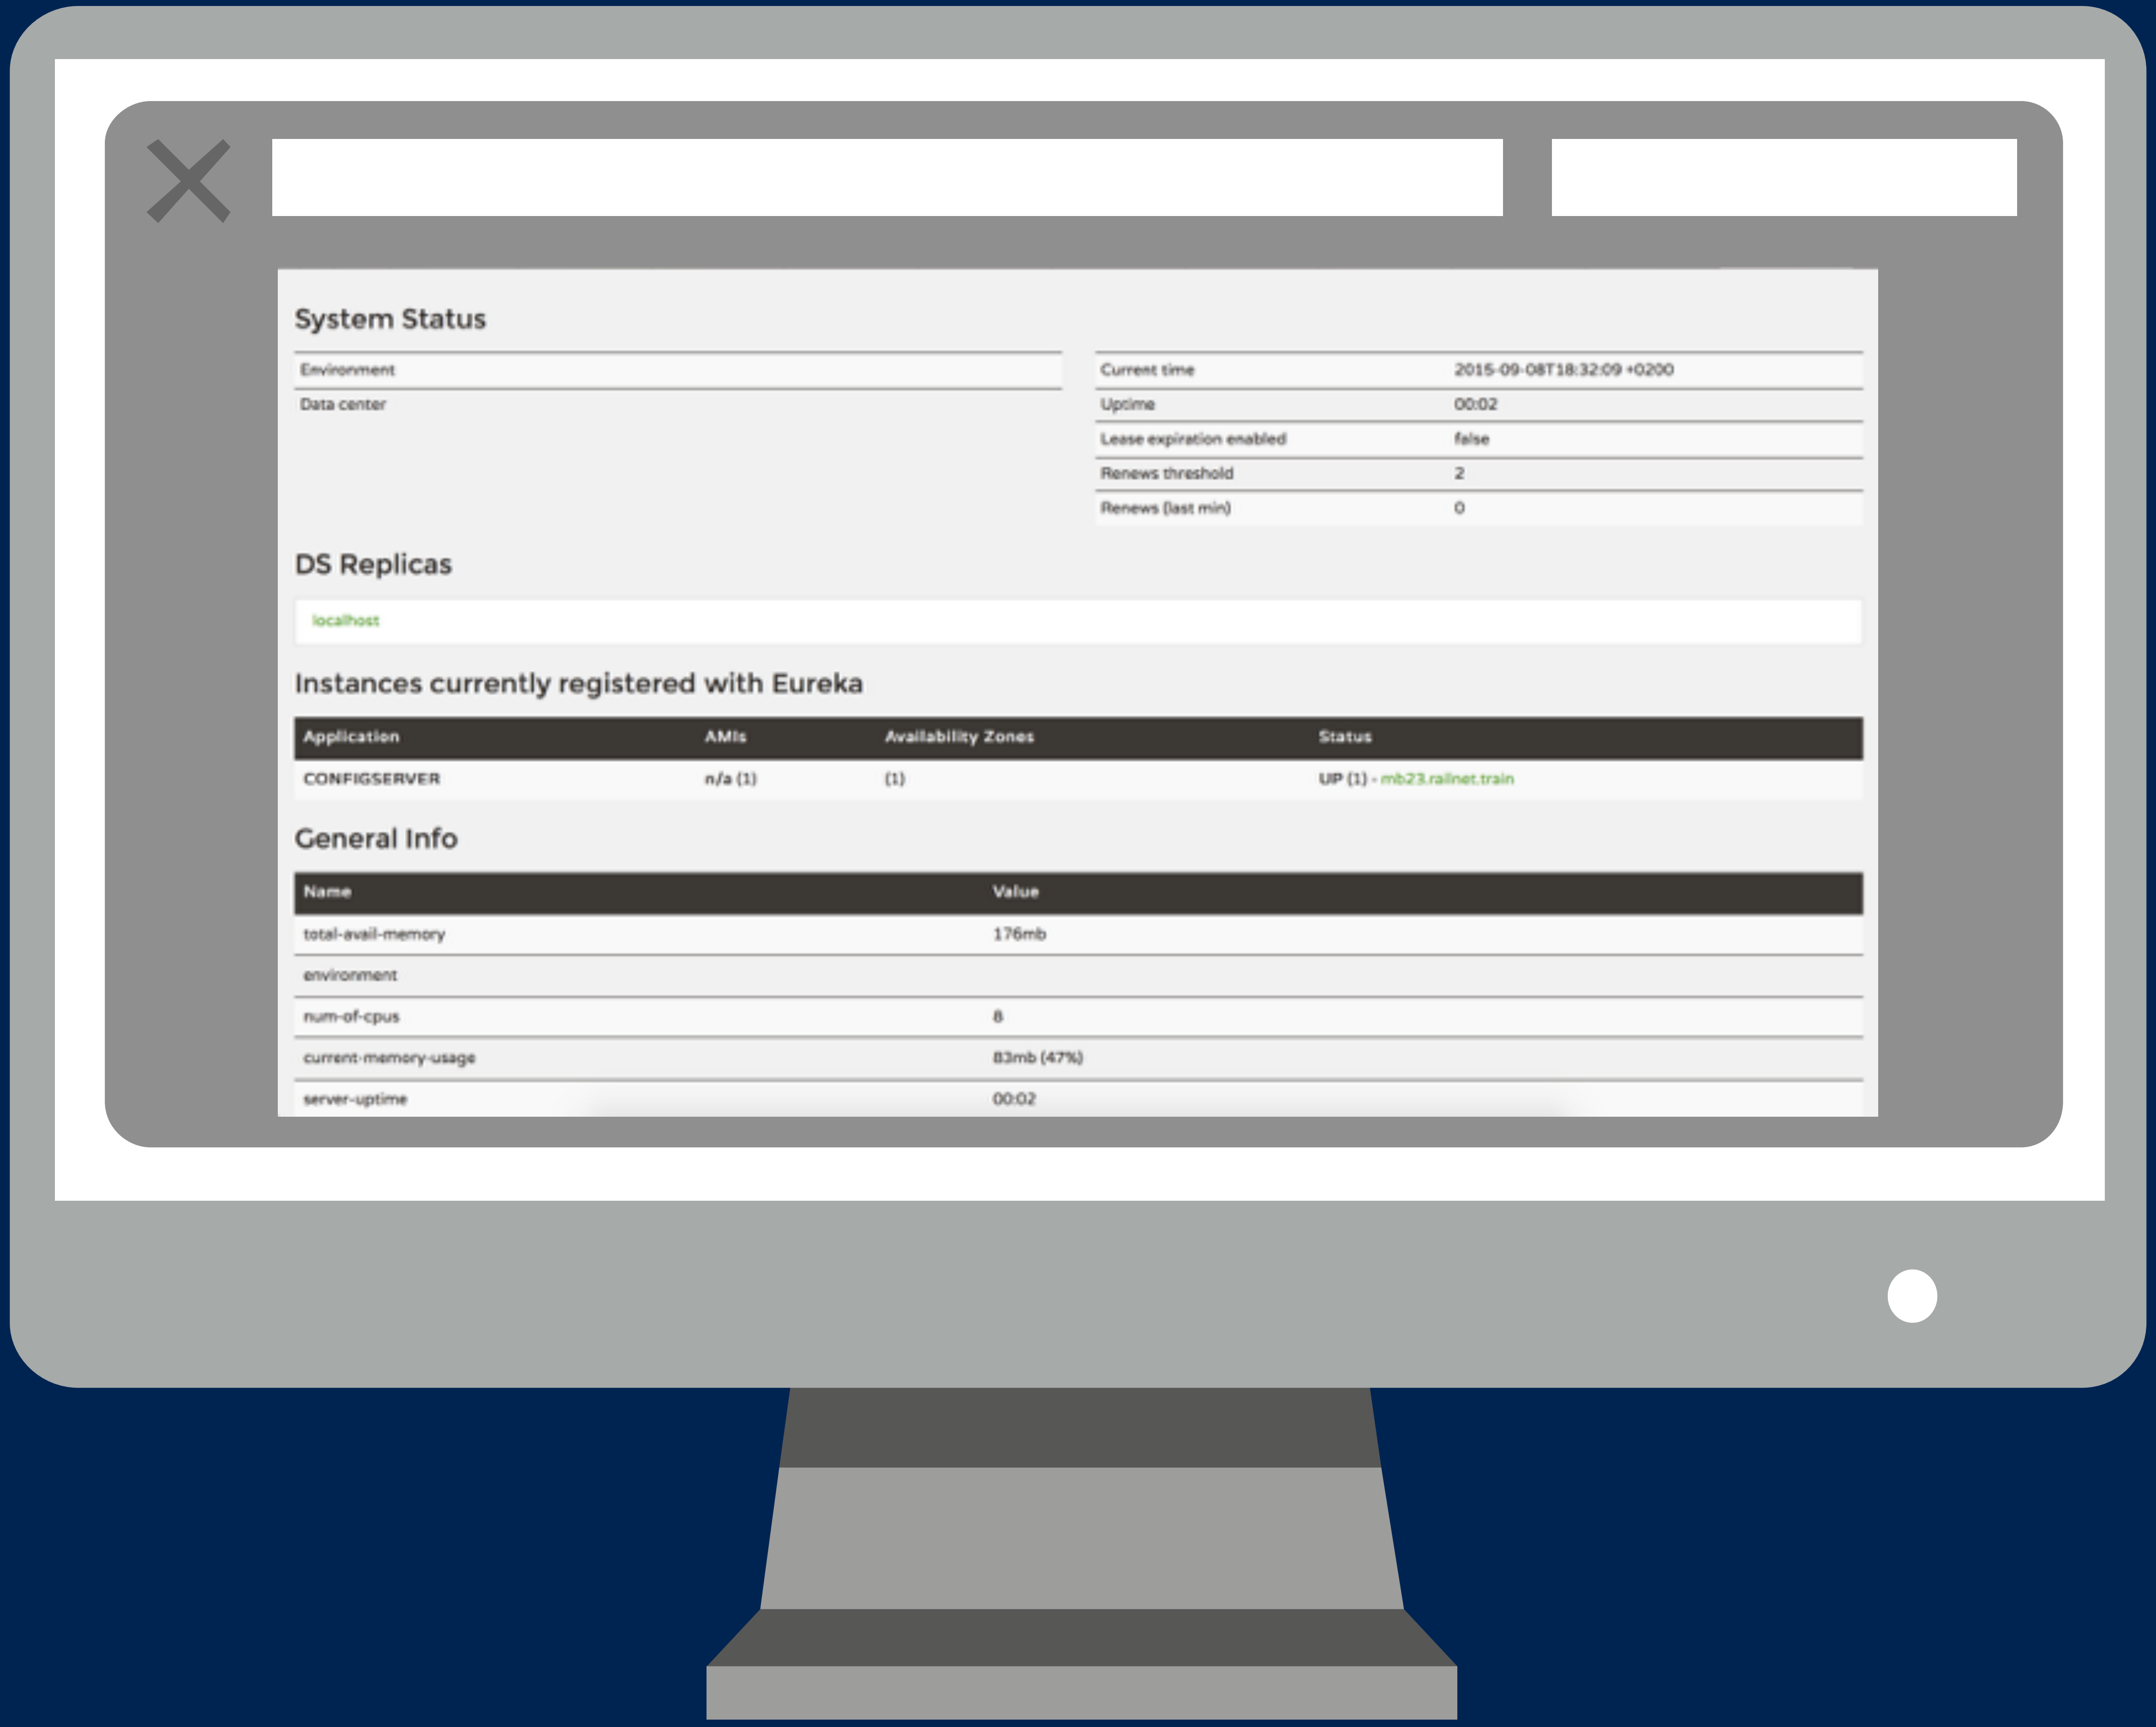

System Status

Environment	
Data center	

Current time	2015-09-08T18:32:09 +0200
Uptime	00:02
Lease expiration enabled	false
Renews threshold	2
Renews (last min)	0

DS Replicas

localhost

Instances currently registered with Eureka

Application	AMIs	Availability Zones	Status
CONFIGSERVER	n/a (1)	(1)	UP (1) - mb23.railnet.train

General Info

Name	Value
total-avail-memory	176mb
environment	
num-of-cpus	8
current-memory-usage	83mb (47%)
server-uptime	00:02

FRAGILE	**ROBUST**	**ANTI-FRAGILE**
POINT-TO-POINT		SERVICE-REGISTRY

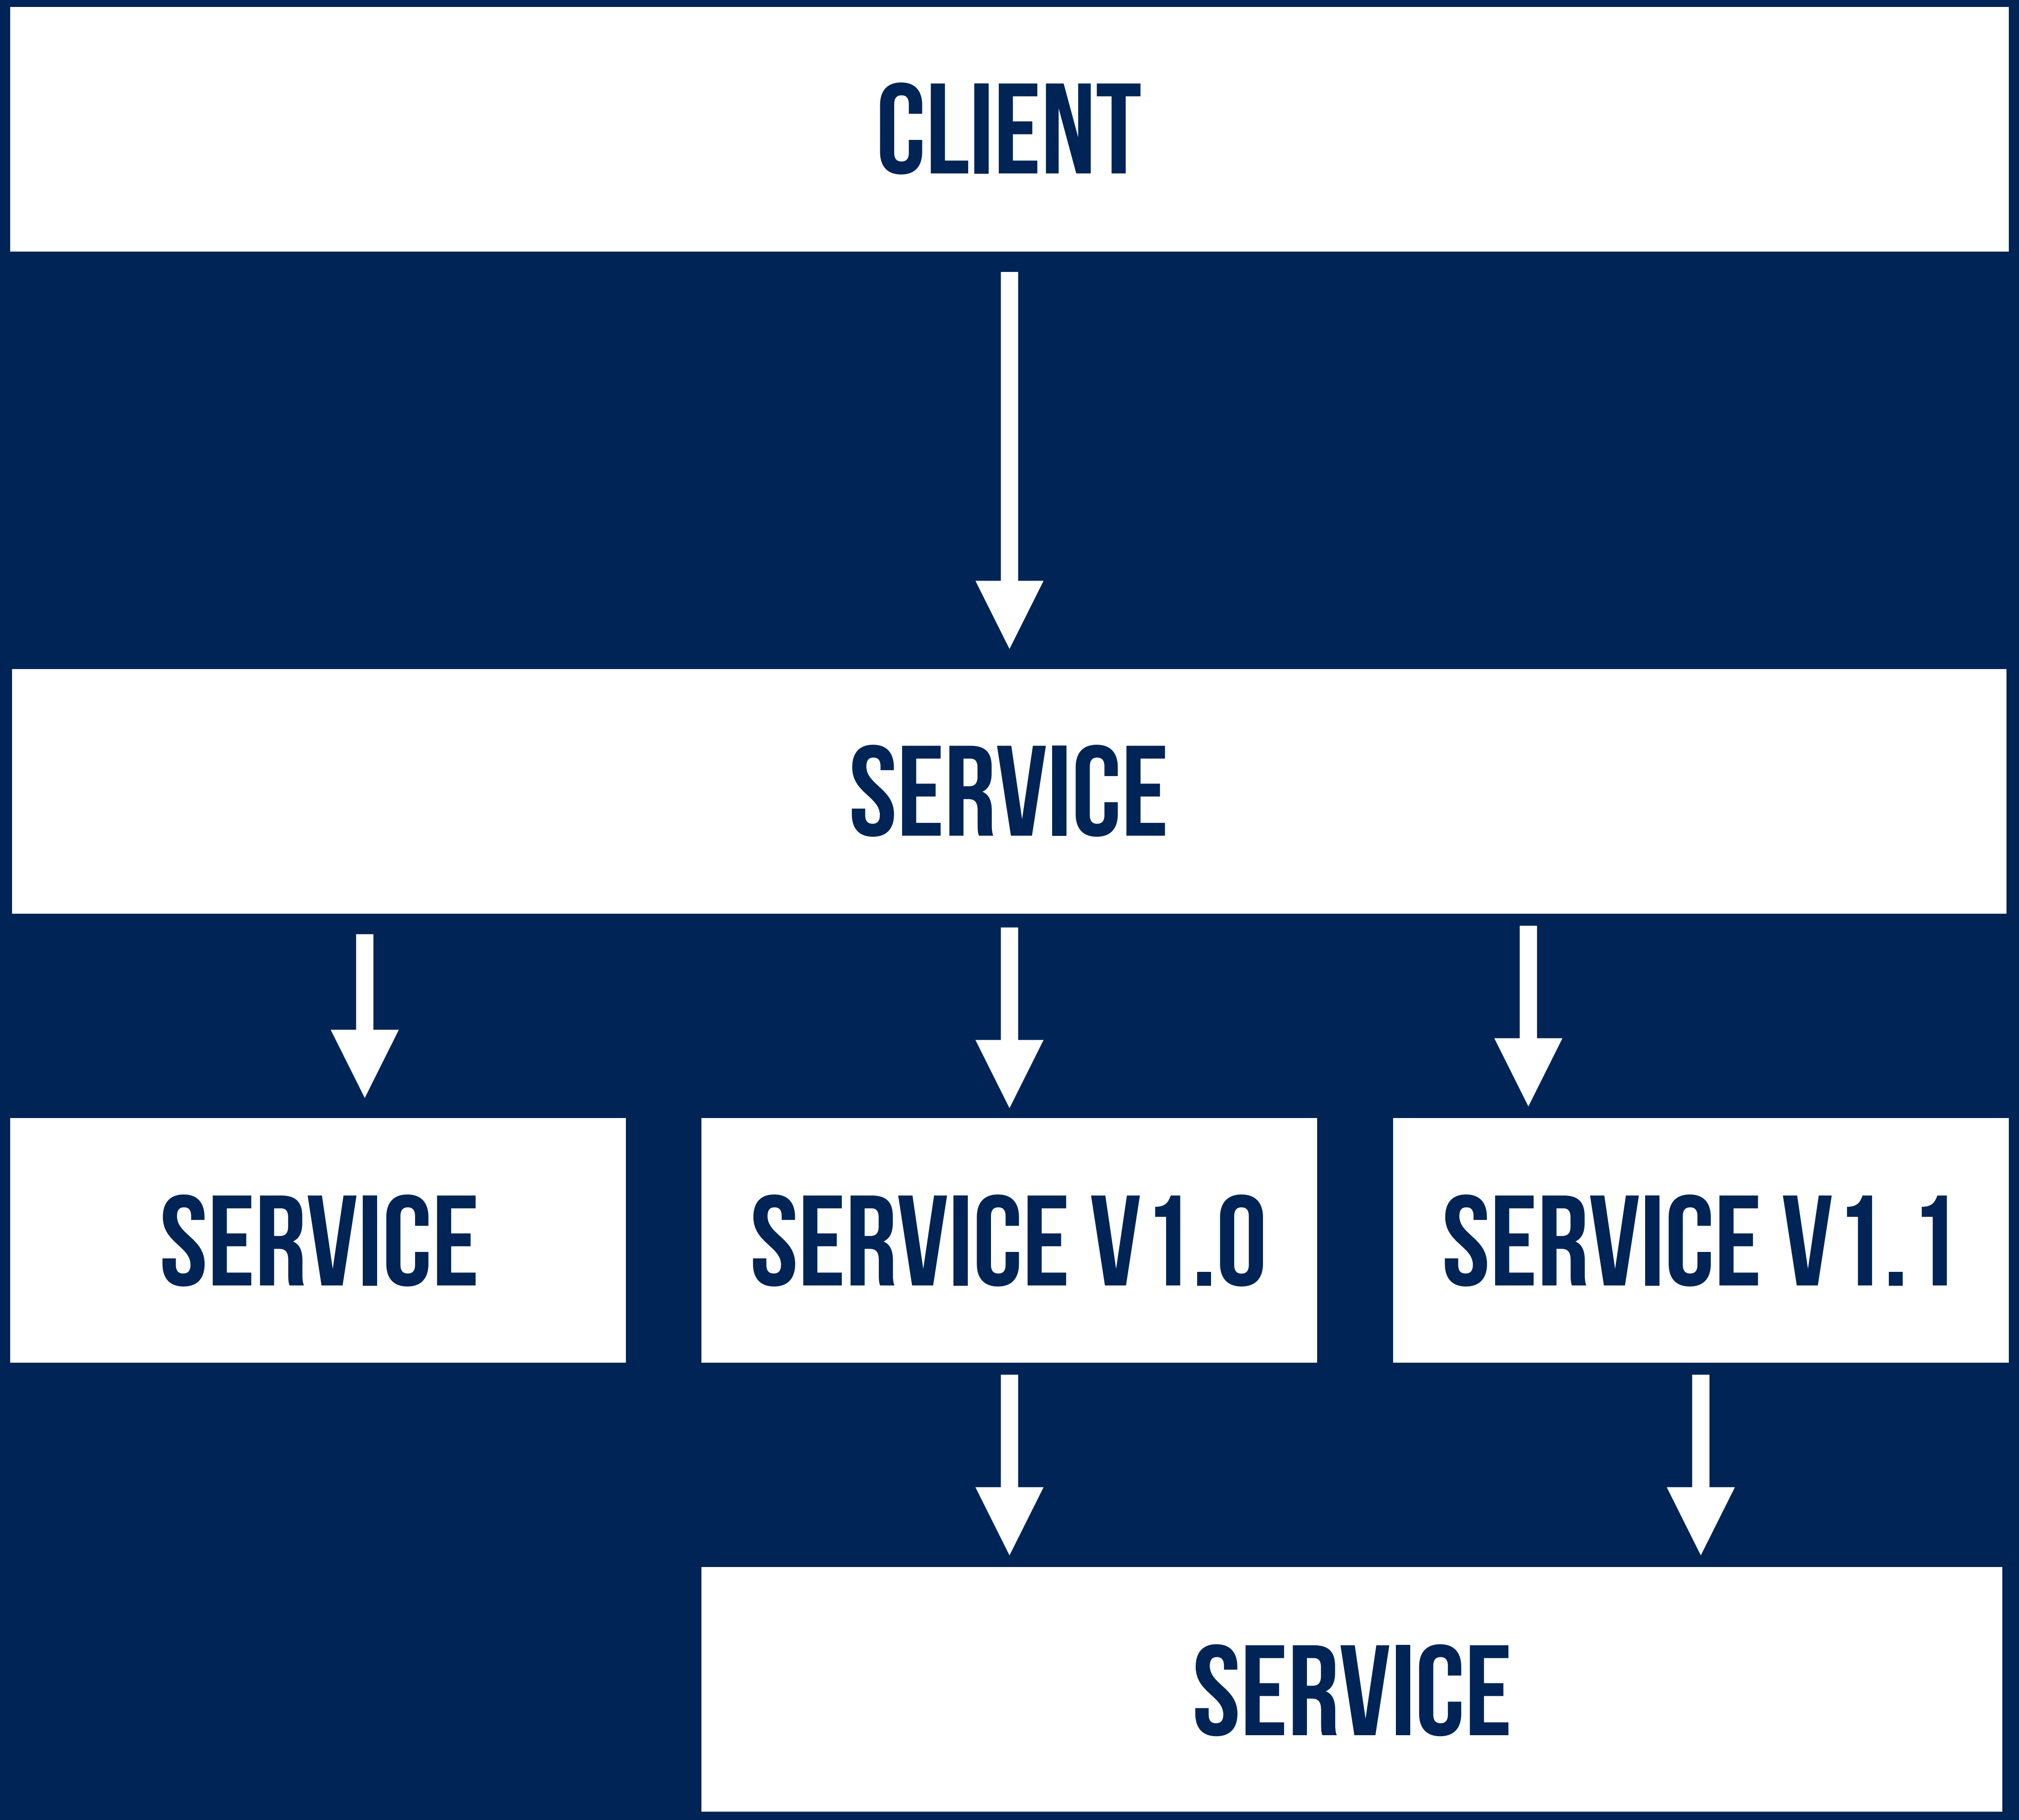
CLIENT
SERVICE
SERVICE
SERVICE V1.0
SERVICE V1.1
SERVICE

NETFLIX

EUREKA

AVAILABILITY
PARTITIONING

CONSUL

CONSISTENCY
AVAILABILITY

ZOOKEEPER

CONSISTENCY
AVAILABILITY

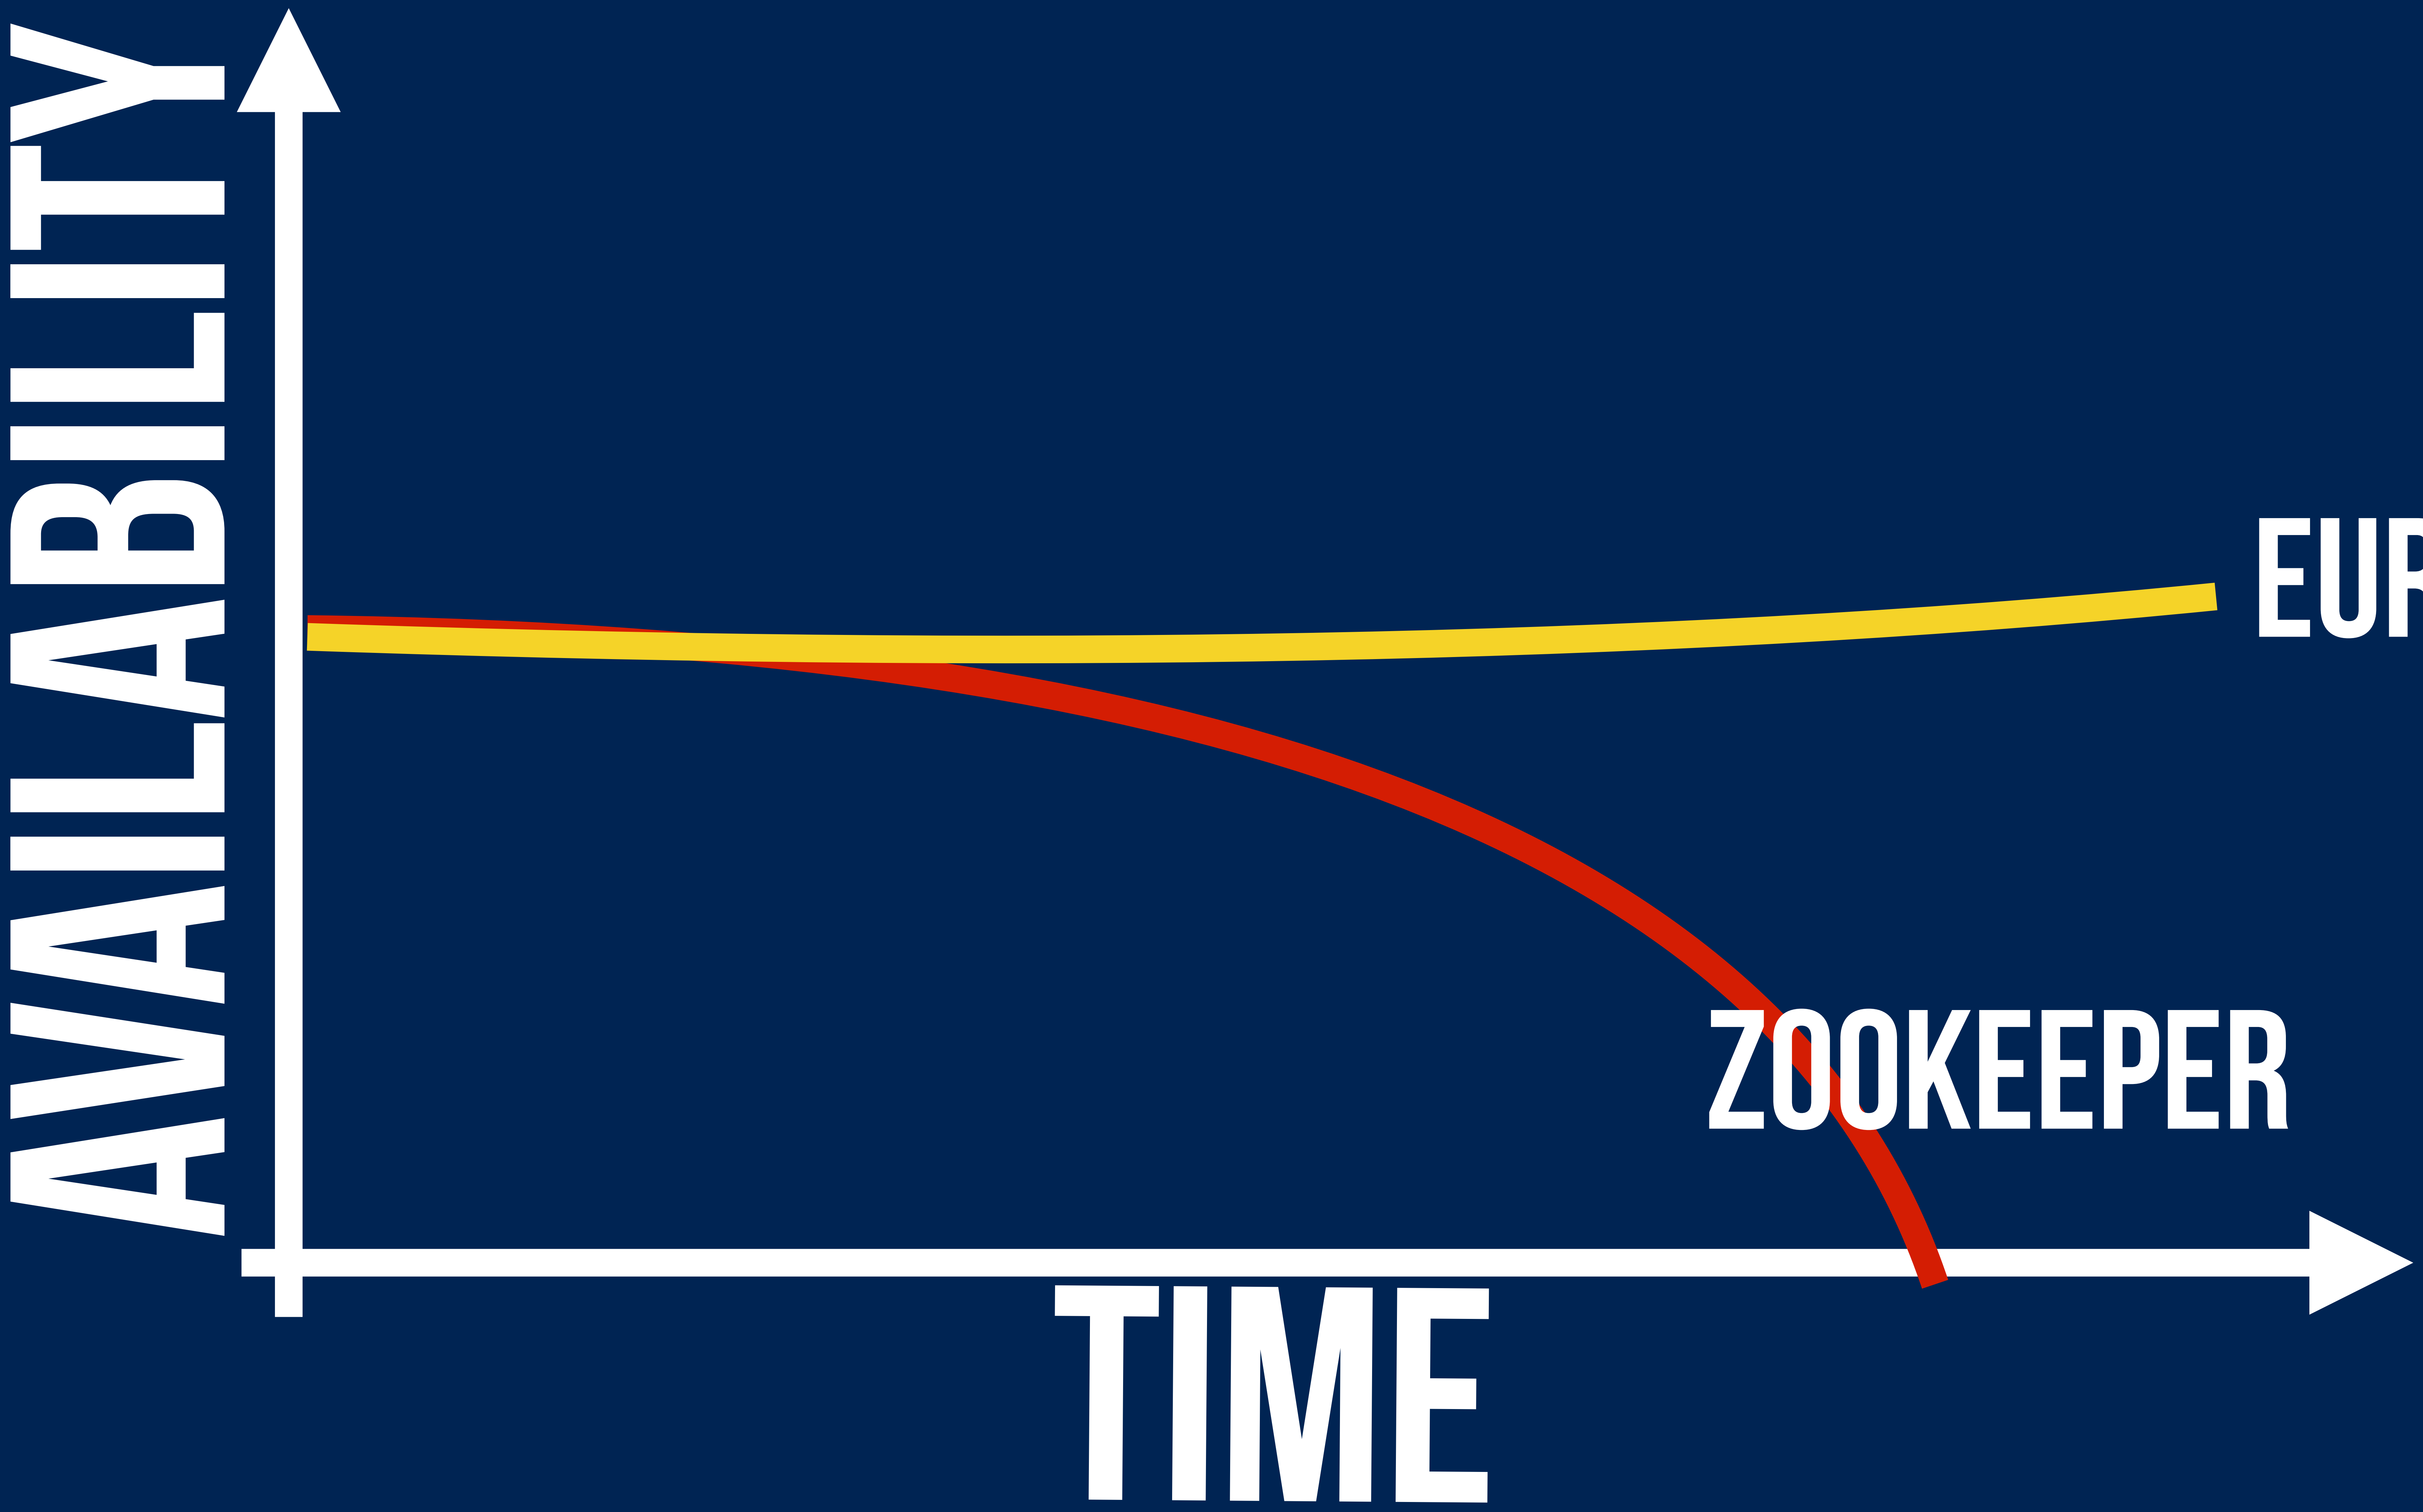
AVAILABILITY
TIME
EUREKA
ZOOKEEPER

PARTIONING

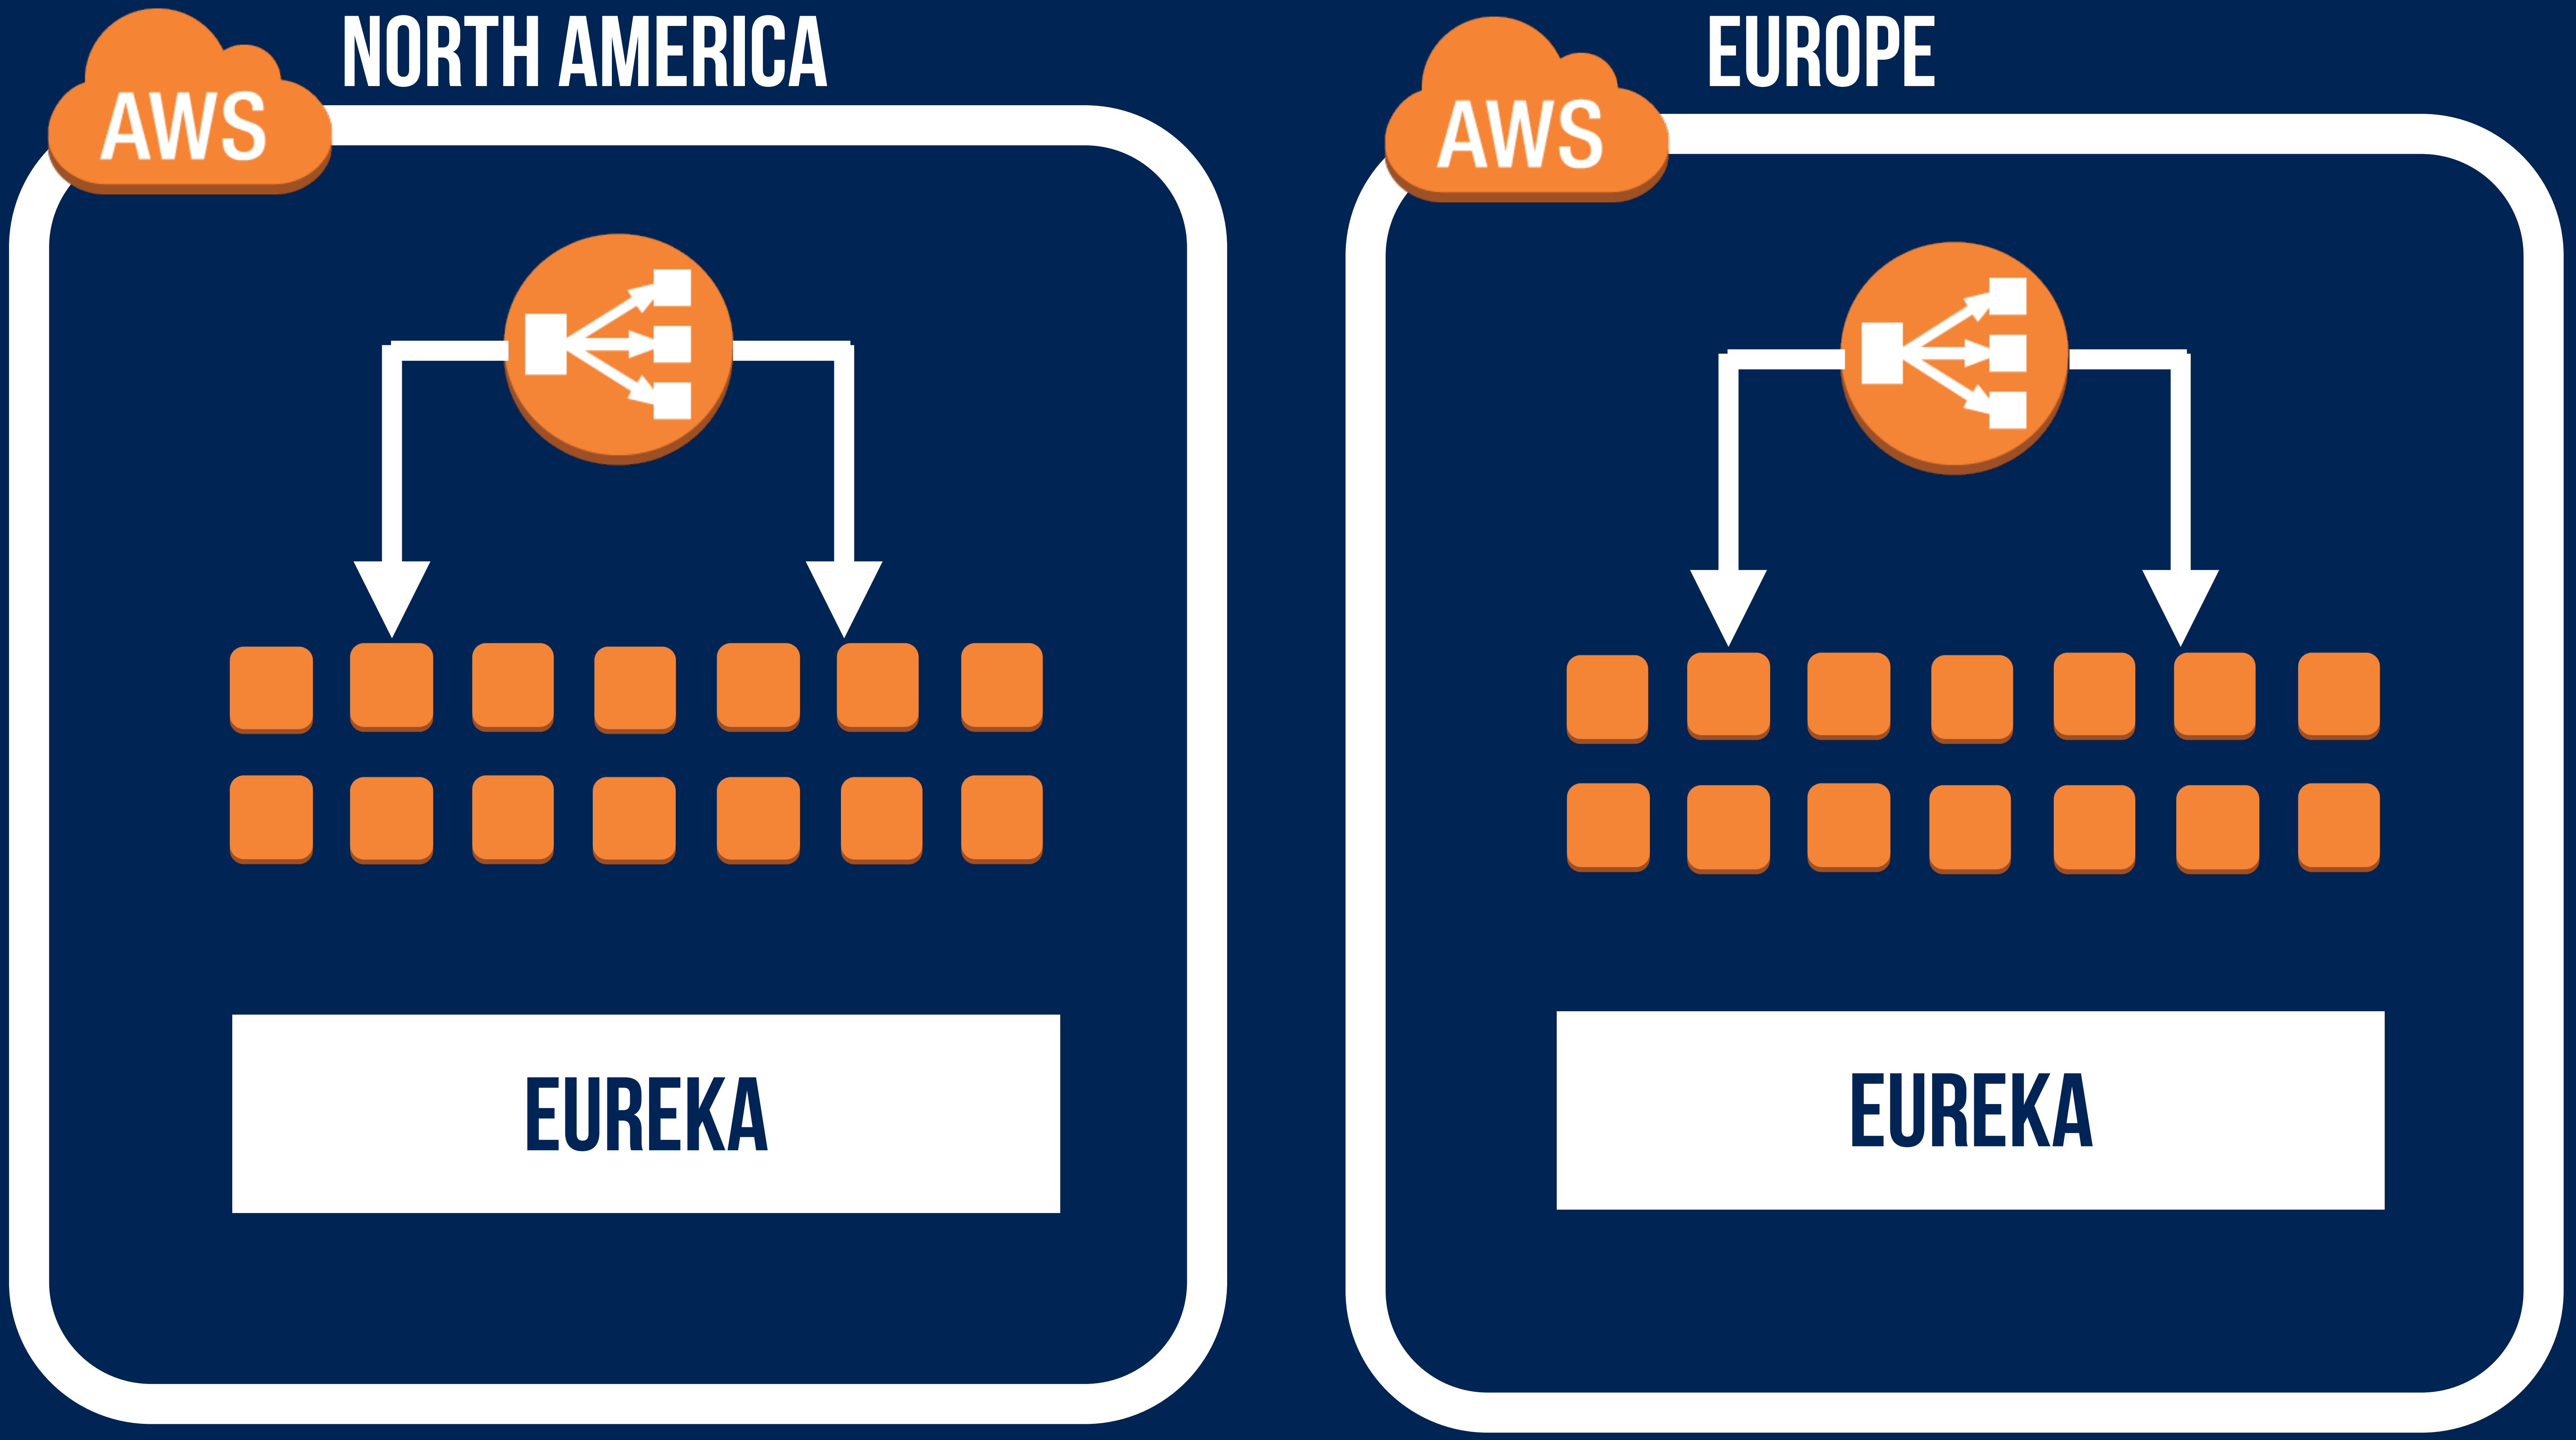

DISCOVERY

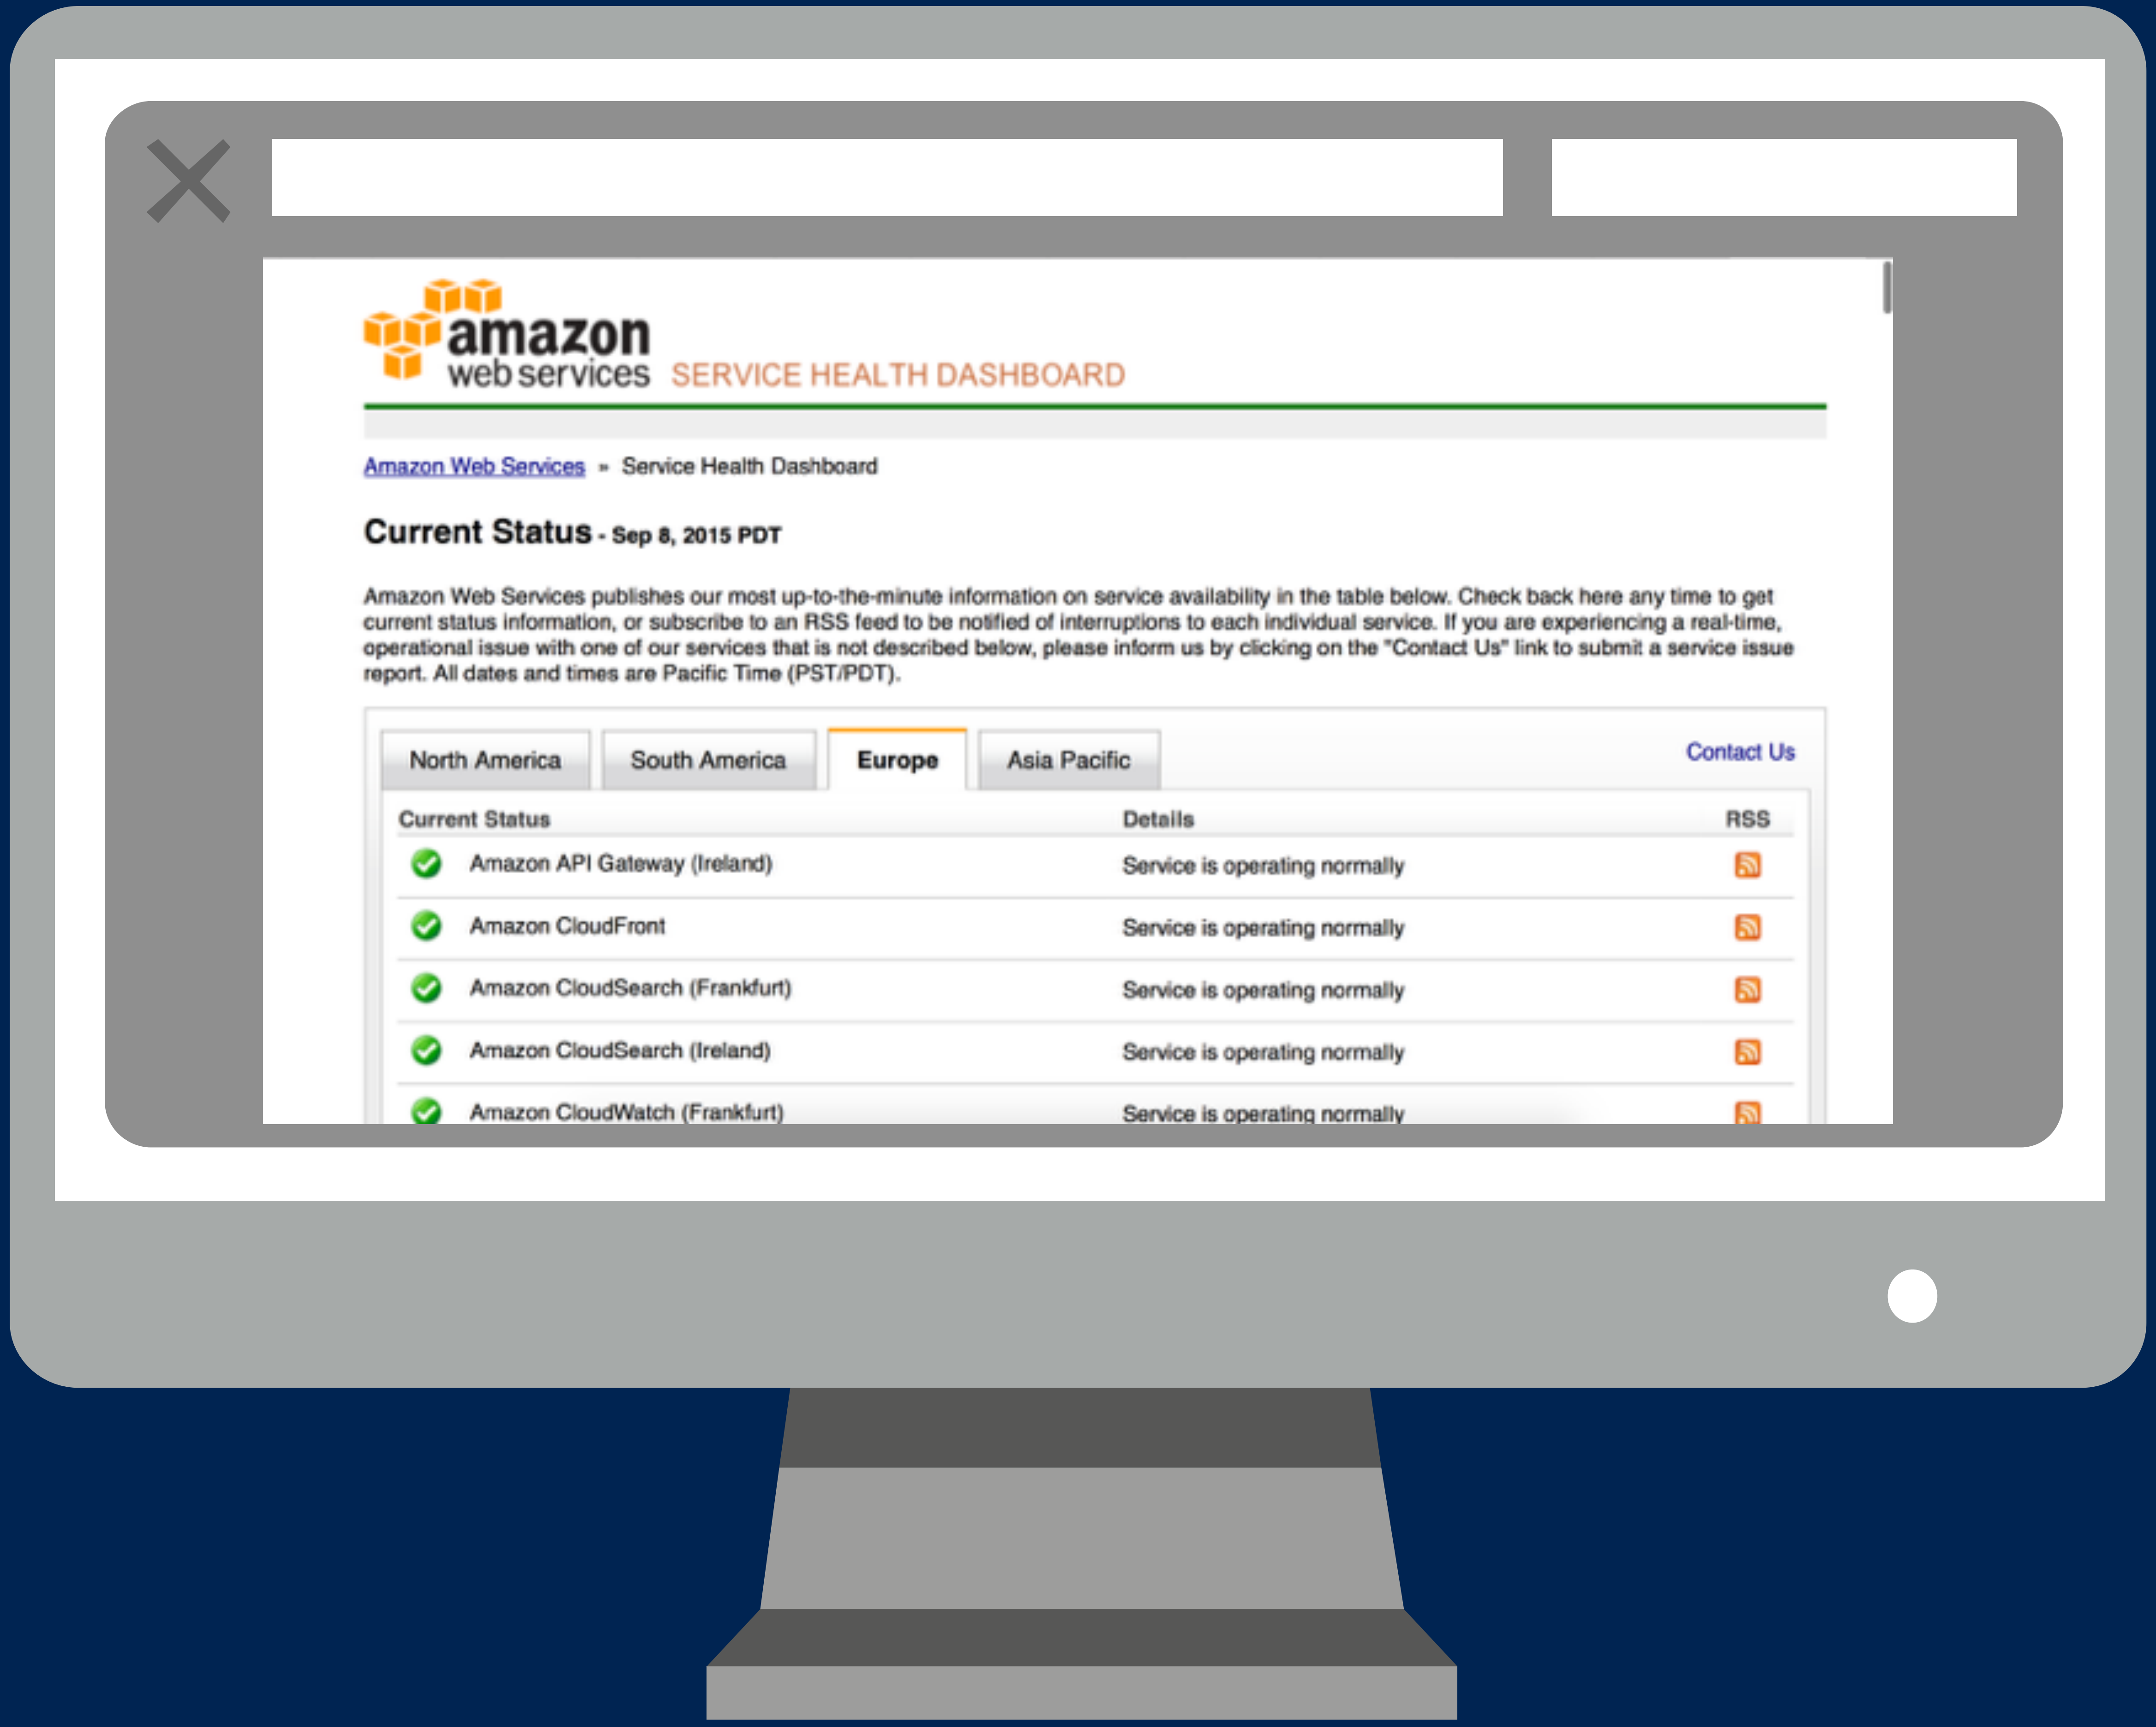
amazon web services SERVICE HEALTH DASHBOARD

Amazon Web Services » Service Health Dashboard

Current Status - Sep 8, 2015 PDT

Amazon Web Services publishes our most up-to-the-minute information on service availability in the table below. Check back here any time to get current status information, or subscribe to an RSS feed to be notified of interruptions to each individual service. If you are experiencing a real-time, operational issue with one of our services that is not described below, please inform us by clicking on the "Contact Us" link to submit a service issue report. All dates and times are Pacific Time (PST/PDT).

North America | South America | **Europe** | Asia Pacific | Contact Us

Current Status	Details	RSS
✓ Amazon API Gateway (Ireland)	Service is operating normally	▨
✓ Amazon CloudFront	Service is operating normally	▨
✓ Amazon CloudSearch (Frankfurt)	Service is operating normally	▨
✓ Amazon CloudSearch (Ireland)	Service is operating normally	▨
✓ Amazon CloudWatch (Frankfurt)	Service is operating normally	▨

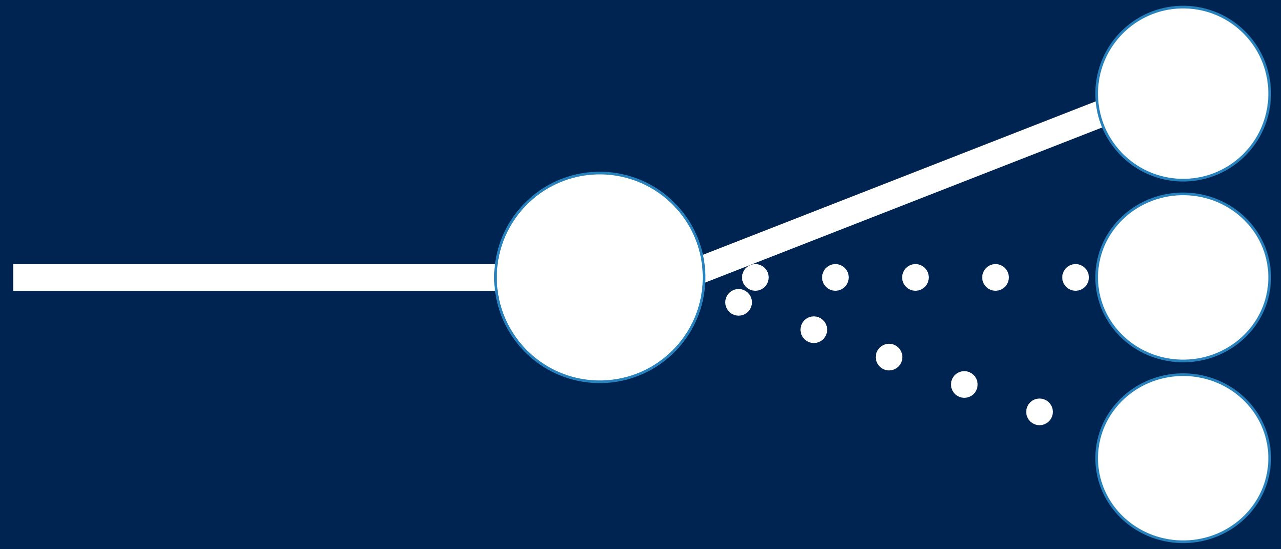

CLIENT DISCOVERY

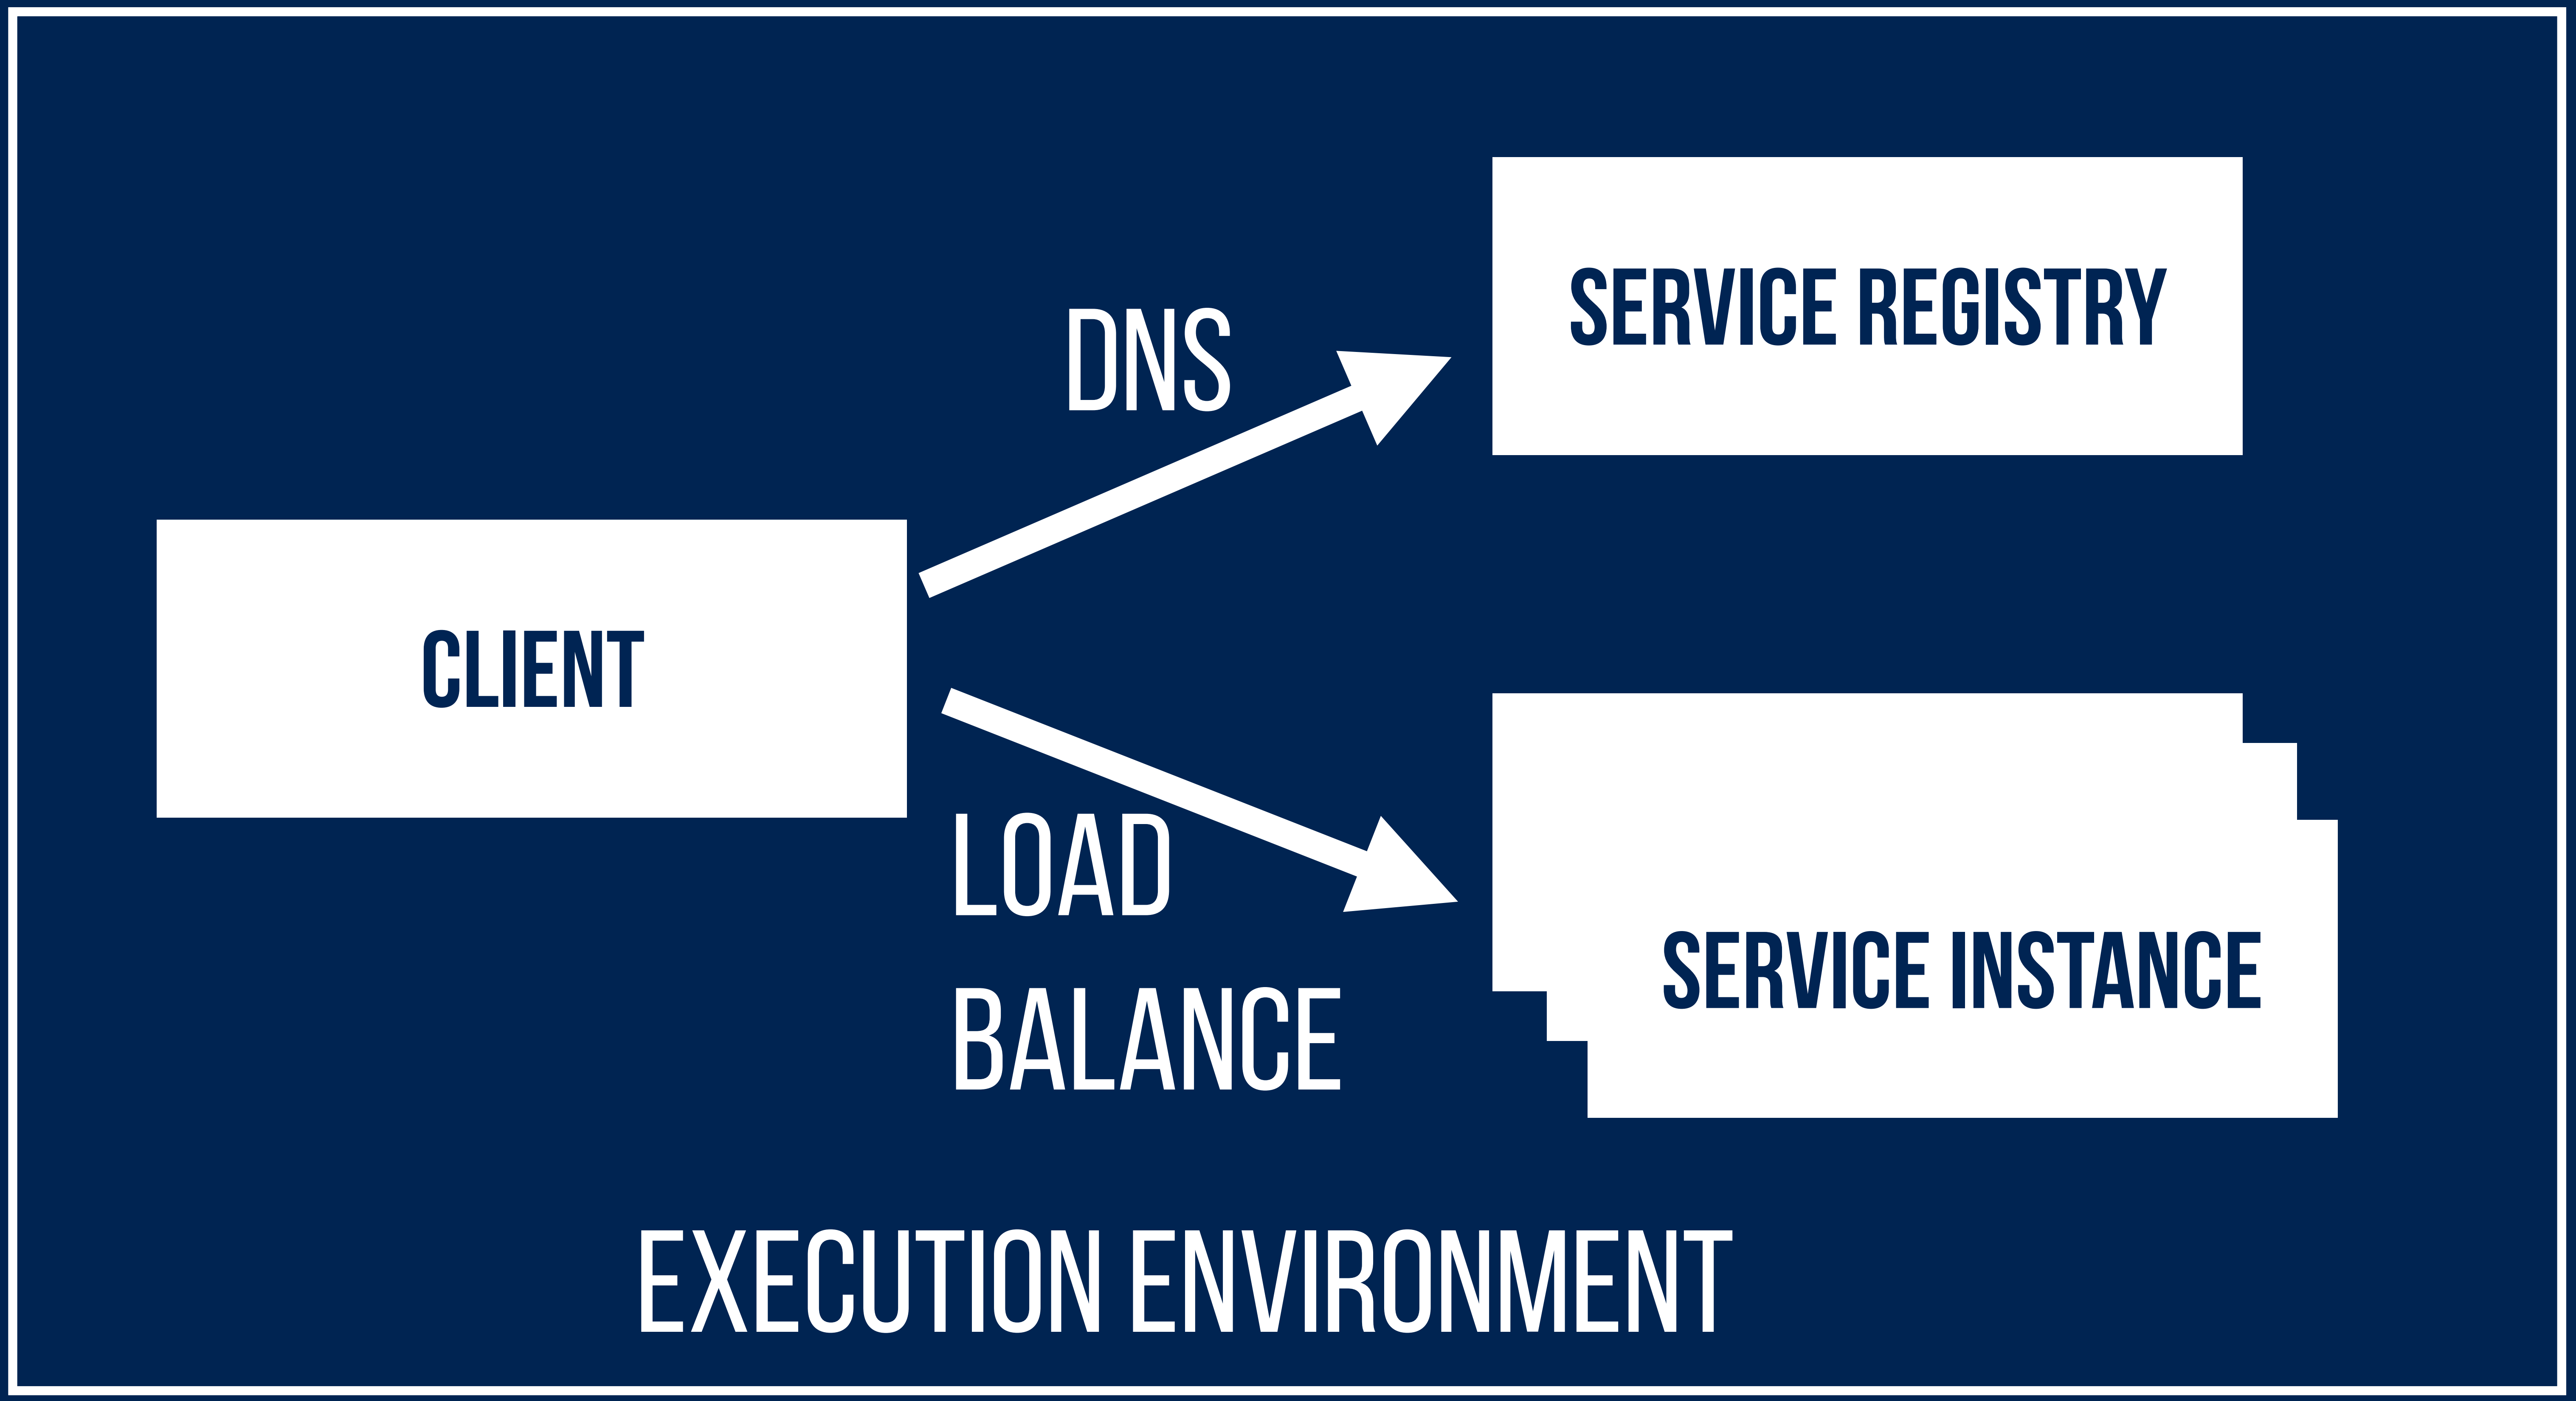

CONSUL

KUBERNETES

CLIENT DISCOVERY

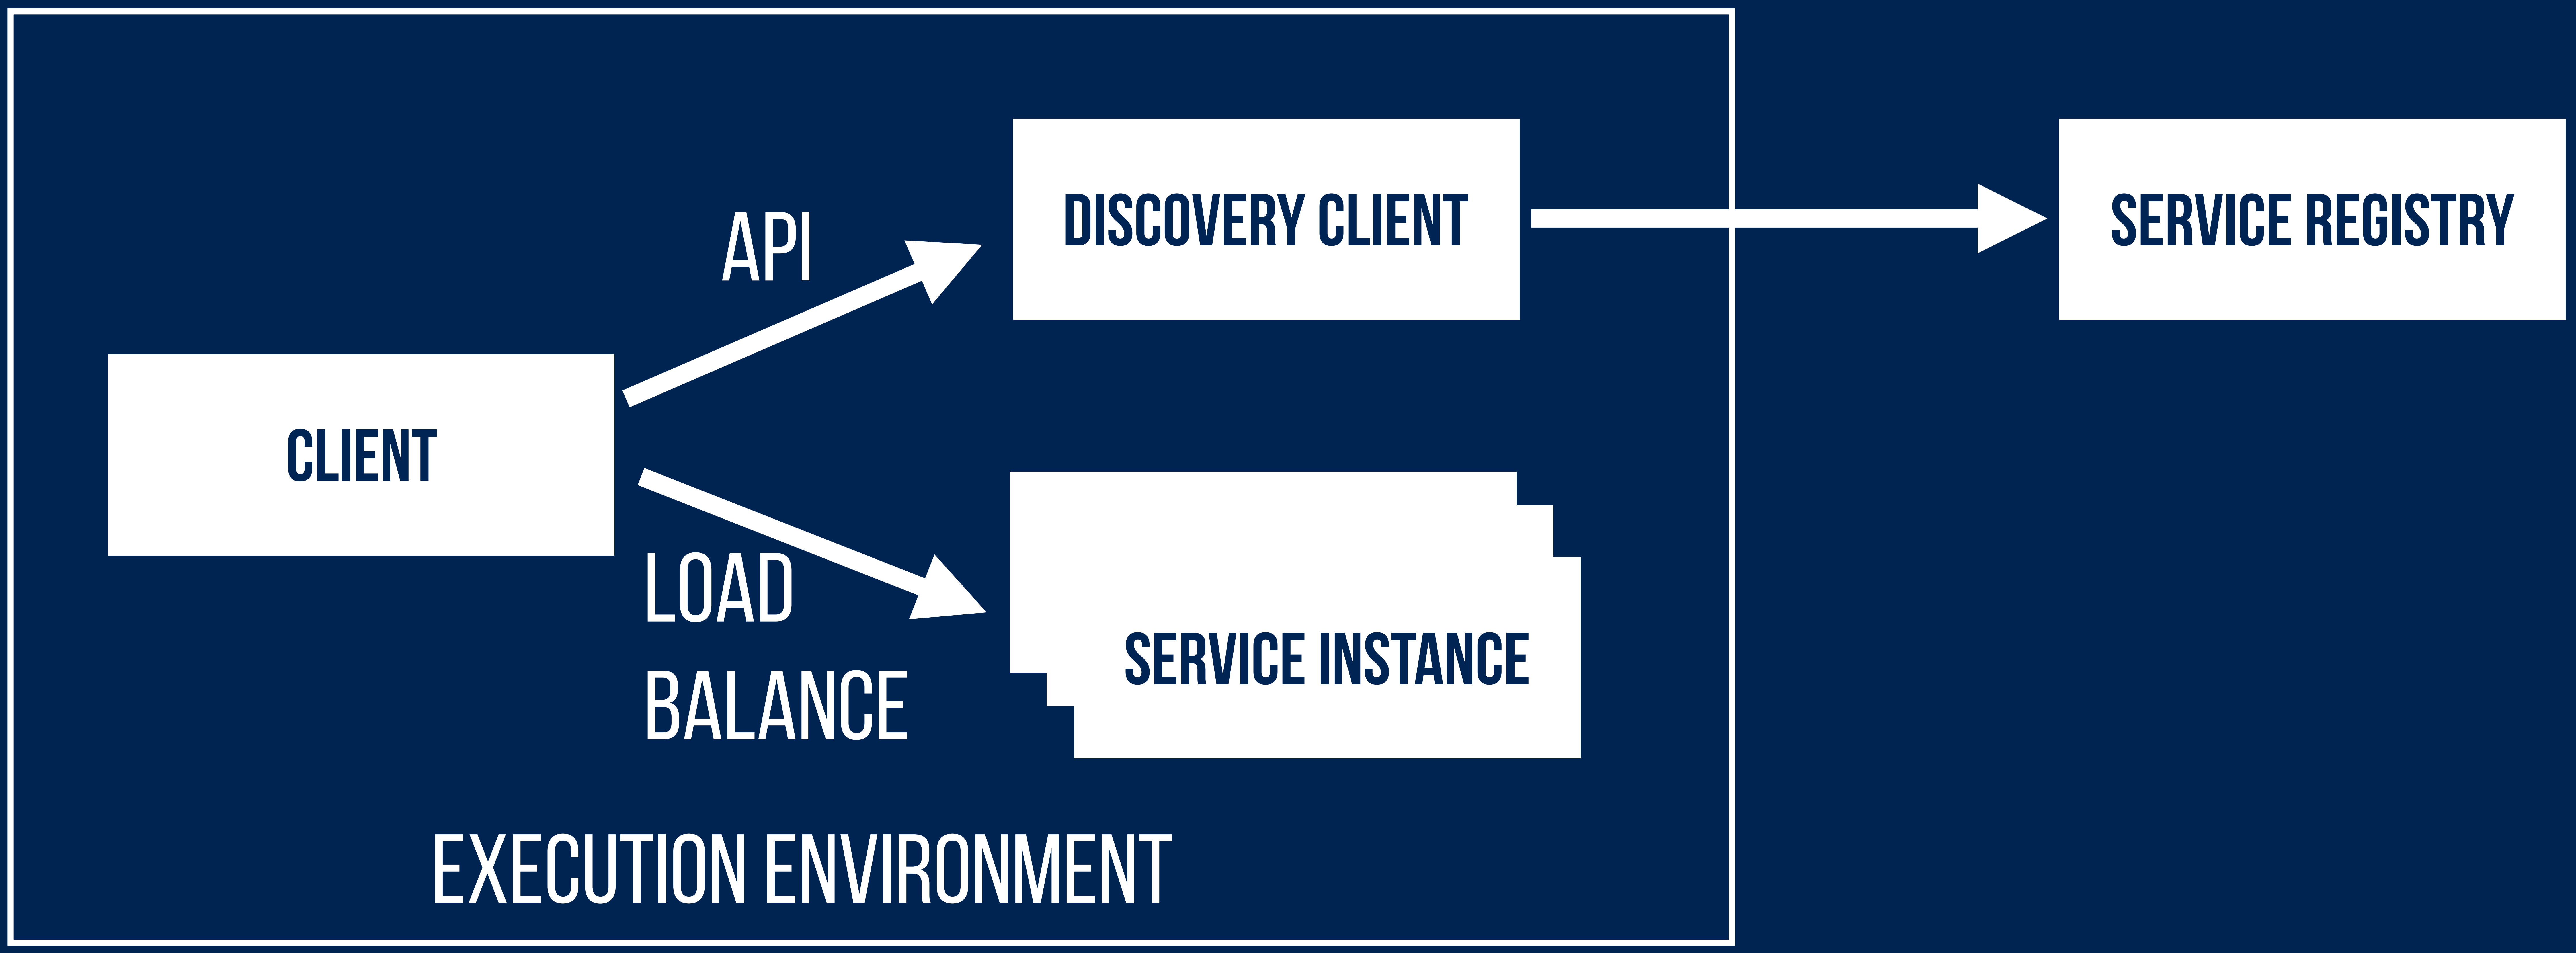

SERVER DISCOVERY

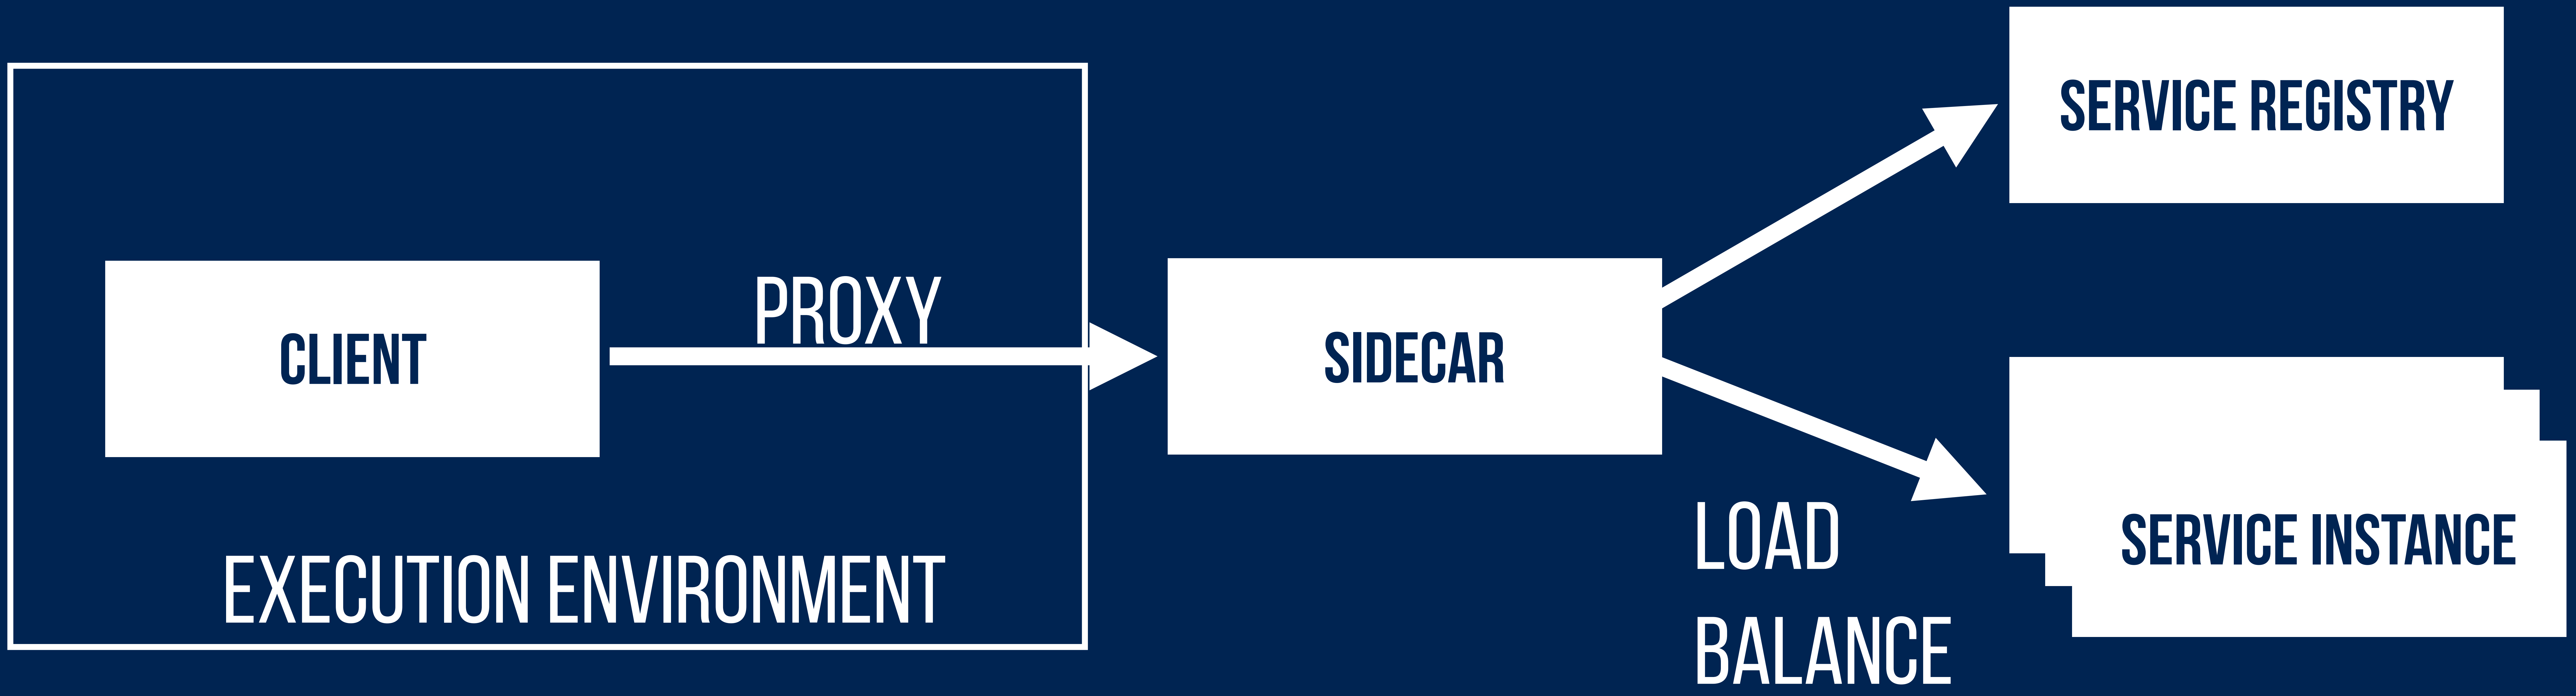

SERVER DISCOVERY

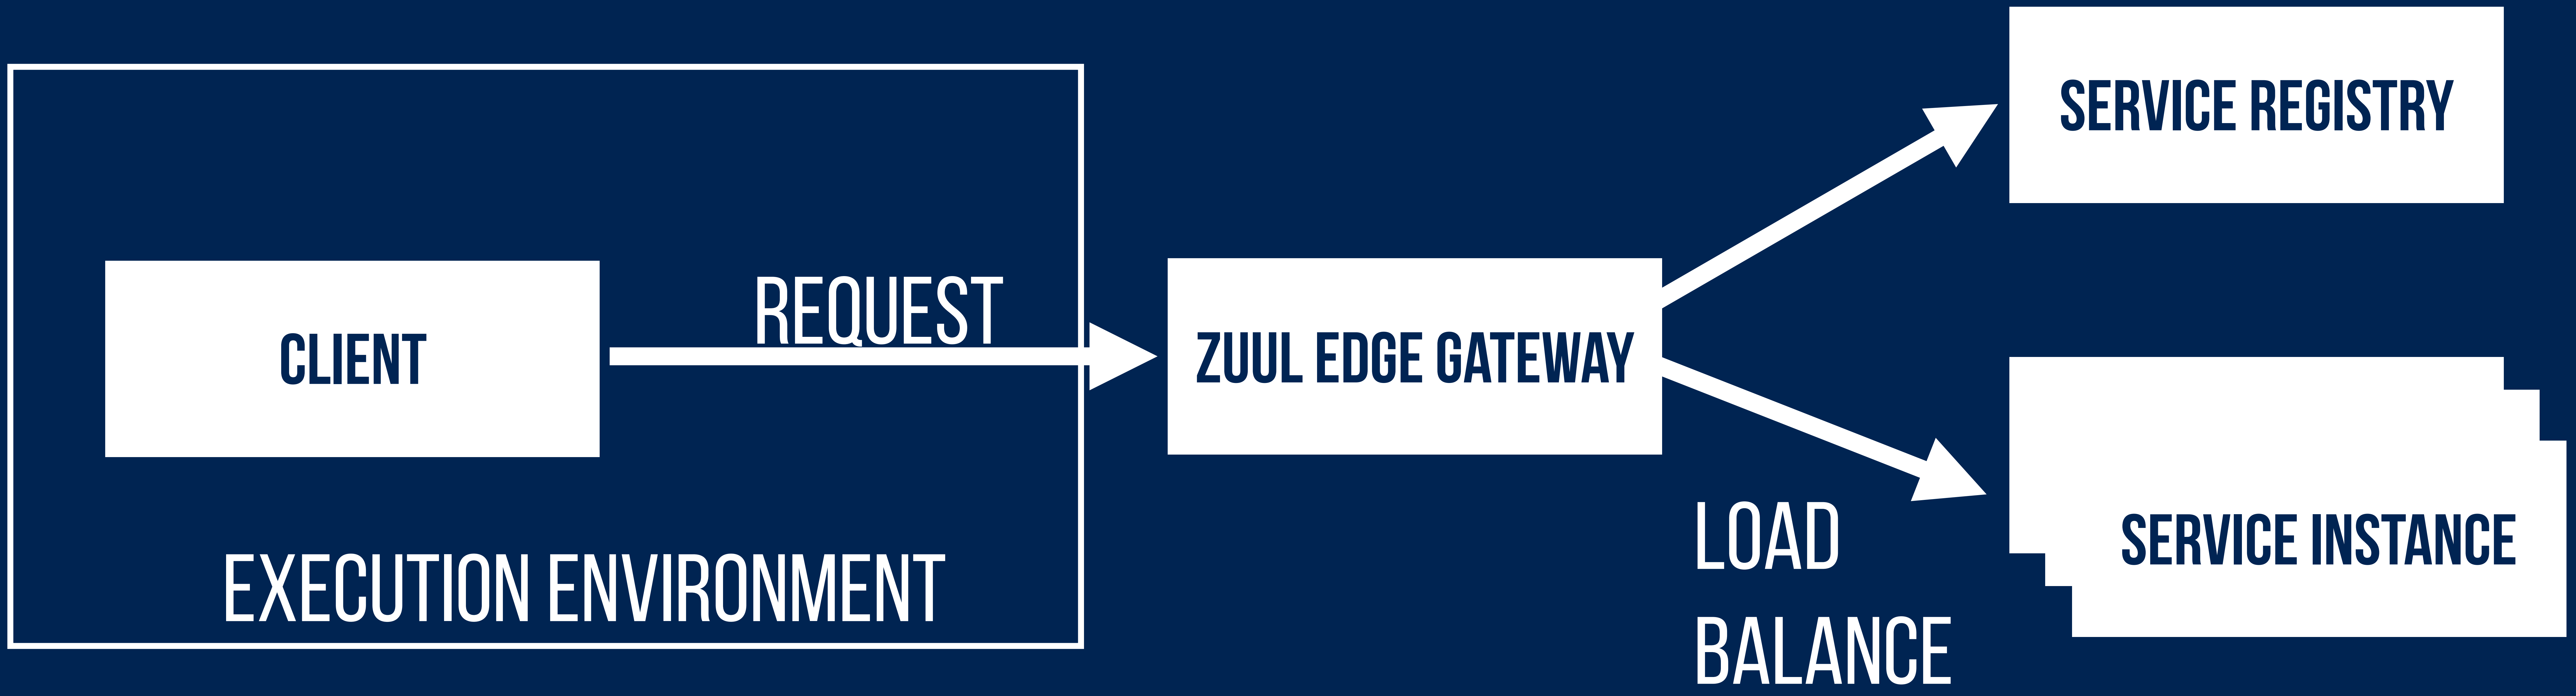

LOAD BALANCING

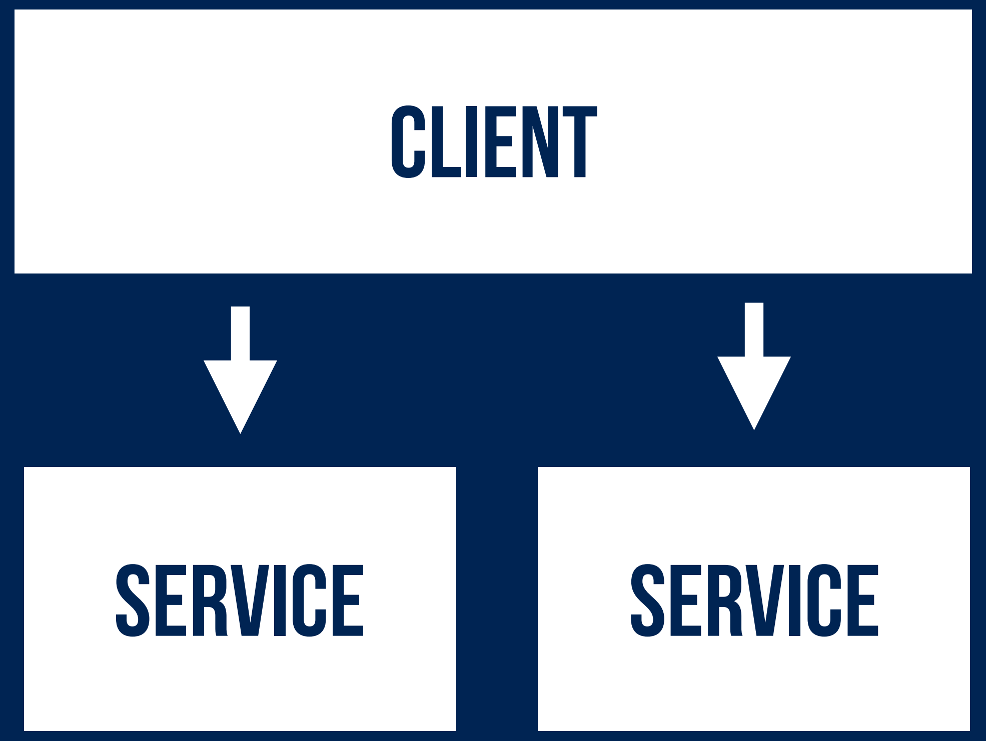

ROUND ROBIN

LINEAR

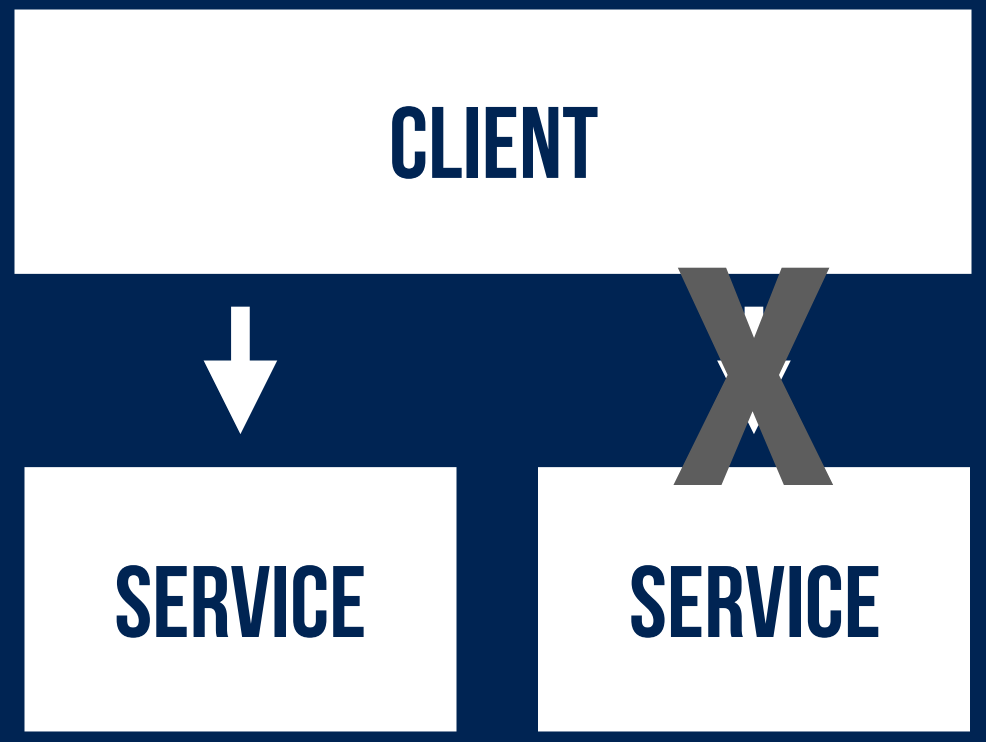

AVAILABILITY FILTERING

NON-LINEAR

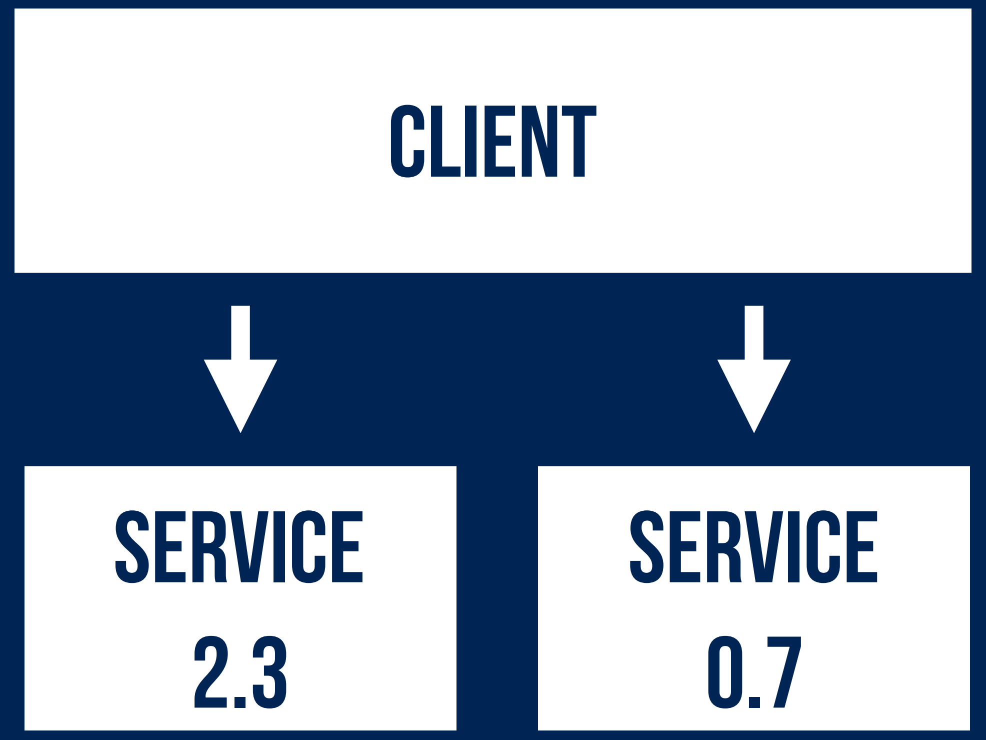

WEIGHTED RESPONSE TIME

NON-LINEAR

ADD LATENCY

```
tc qdisc add dev eth0 root latency delay 1000ms 500ms
```

CORRUPT PACKAGES

```
tc qdisc add dev eth0 root netem corrupt 5%
```

DROP PACKAGES

```
tc qdisc add dev eth0 root netem loss 7% 25%
```

BLOCK DNS

```
iptables -A INPUT -p tcp -m tcp --dport 53 -j DROP
```

SIMULATE HEAVY IO

```
while true;

do
    dd if=/dev/urandom of=/burn bs=1M count=1024 iflag=fullblock
done
```

BURN CPU

```
while true;
    do openssl speed;
done
```

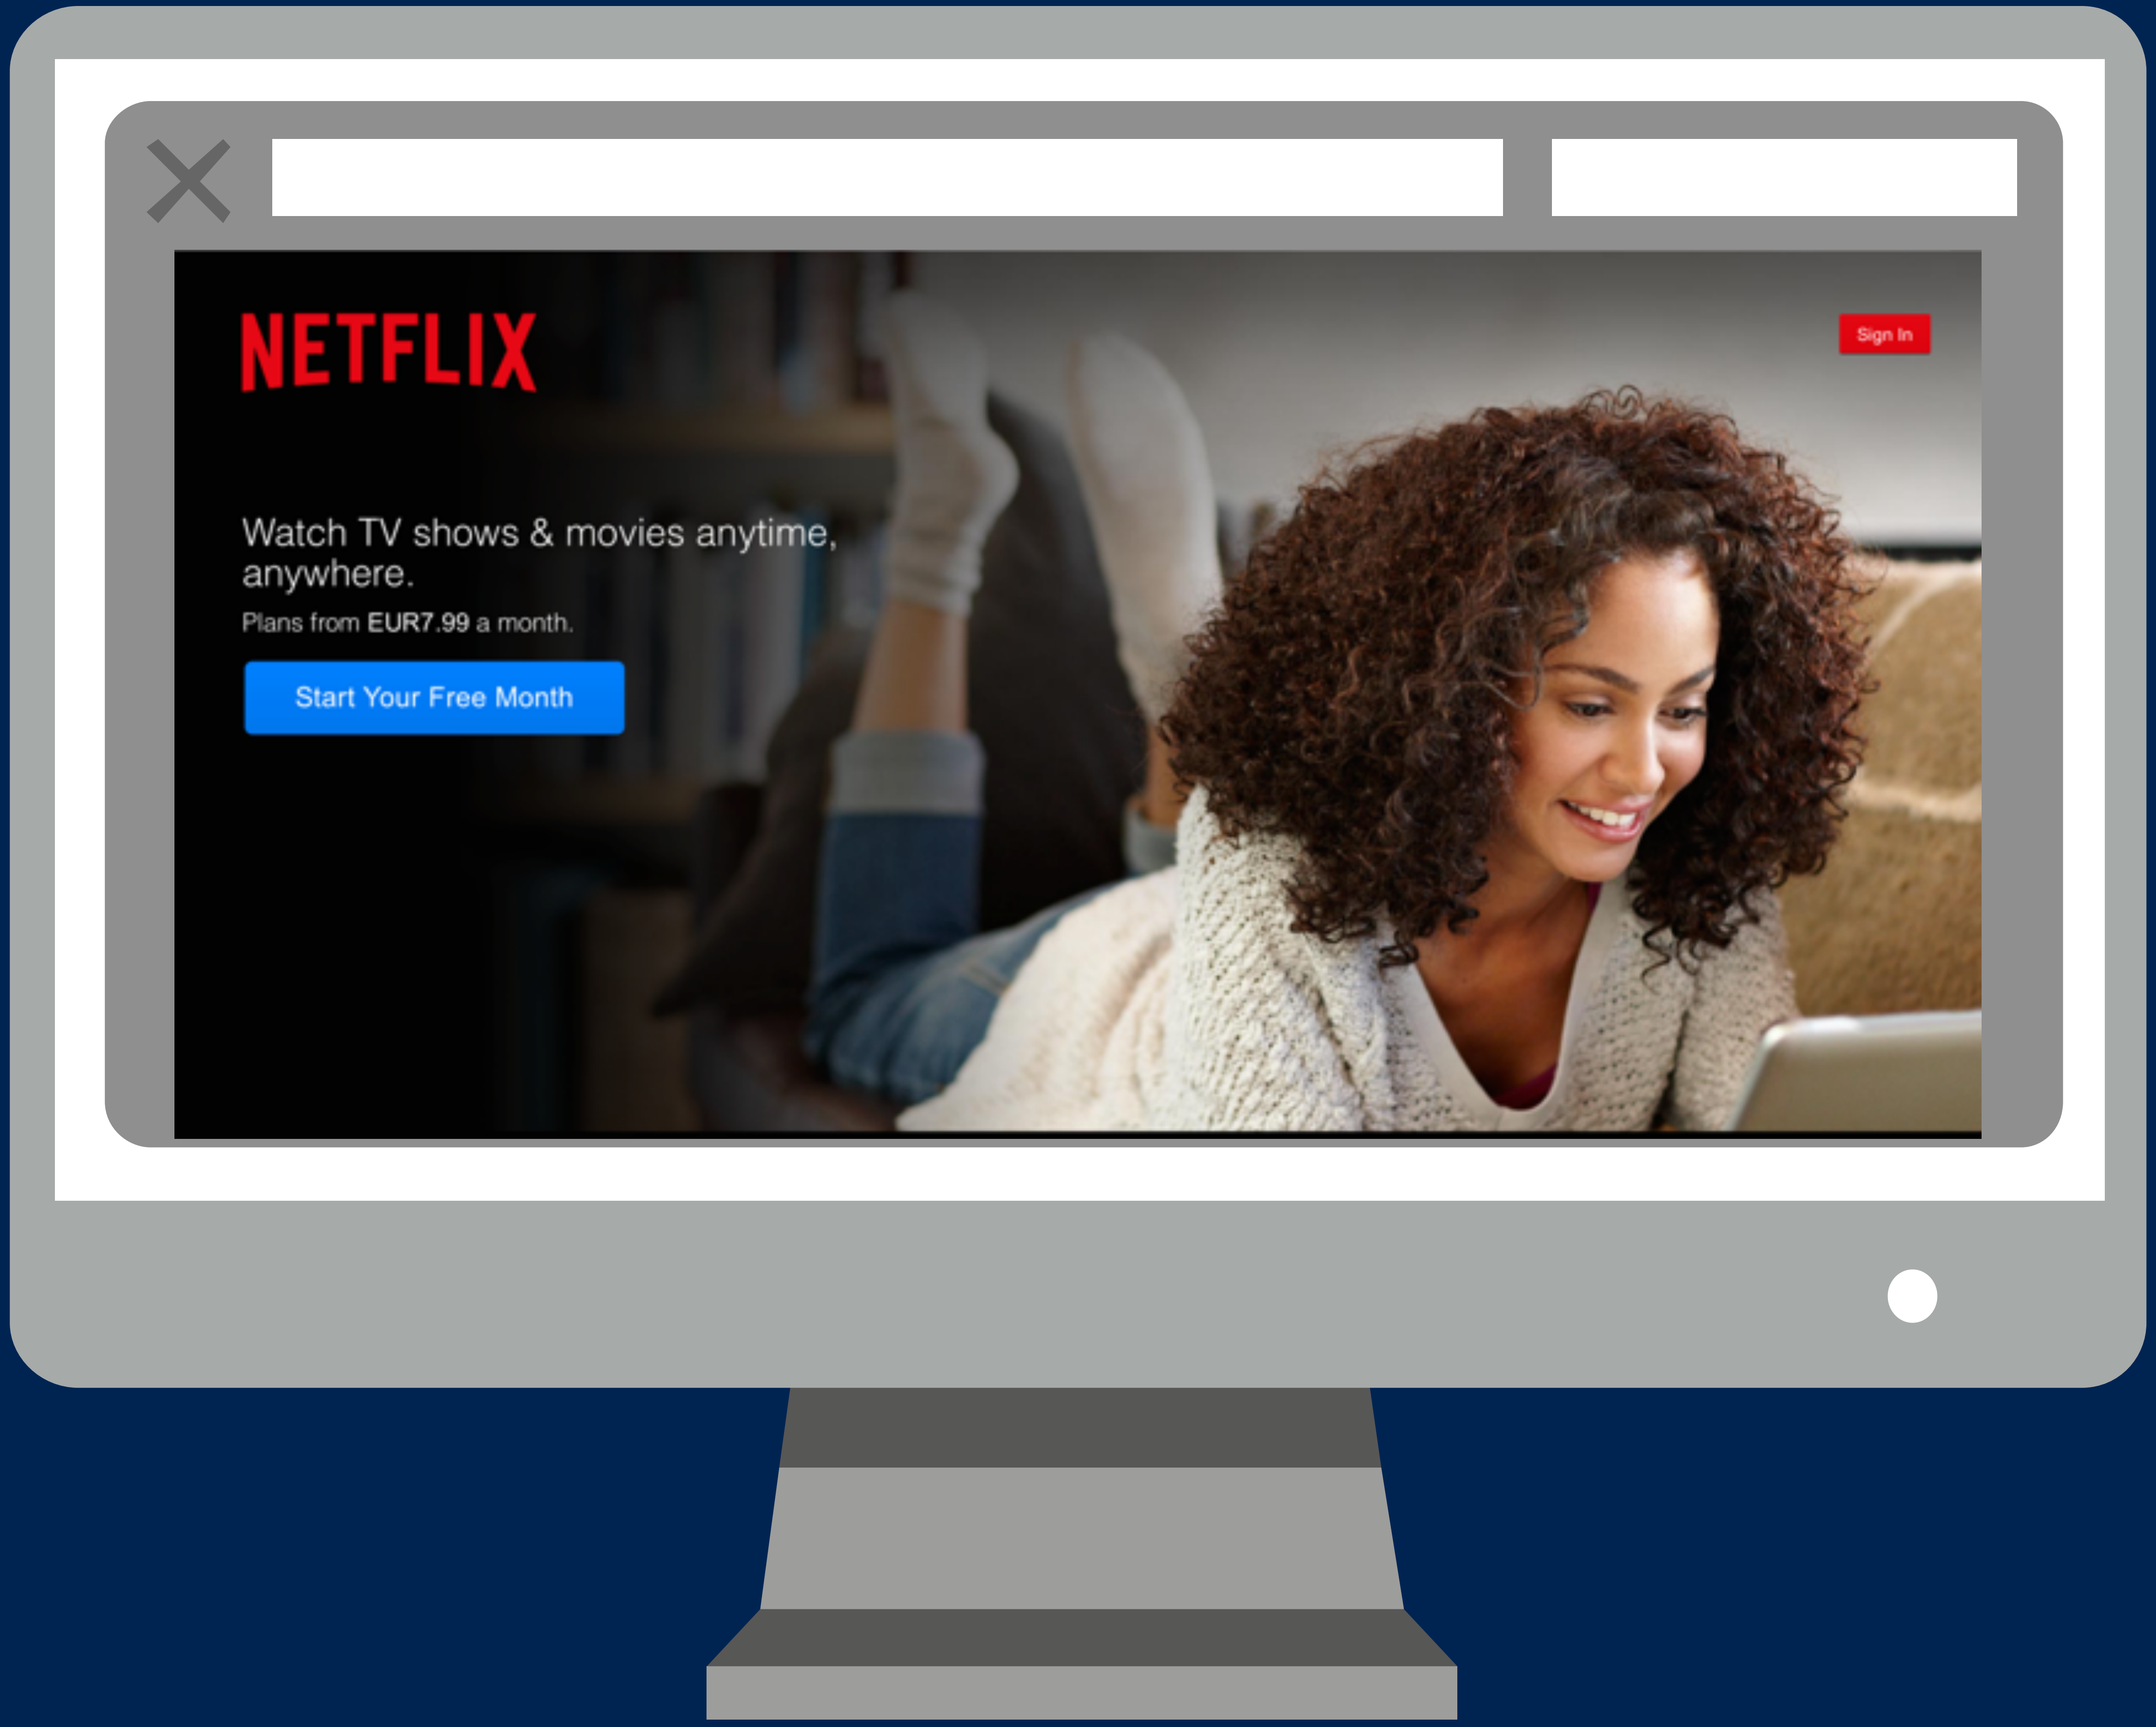
NETFLIX
Sign In
Watch TV shows & movies anytime, anywhere.
Plans from EUR7.99 a month.
Start Your Free Month

API GATEWAY

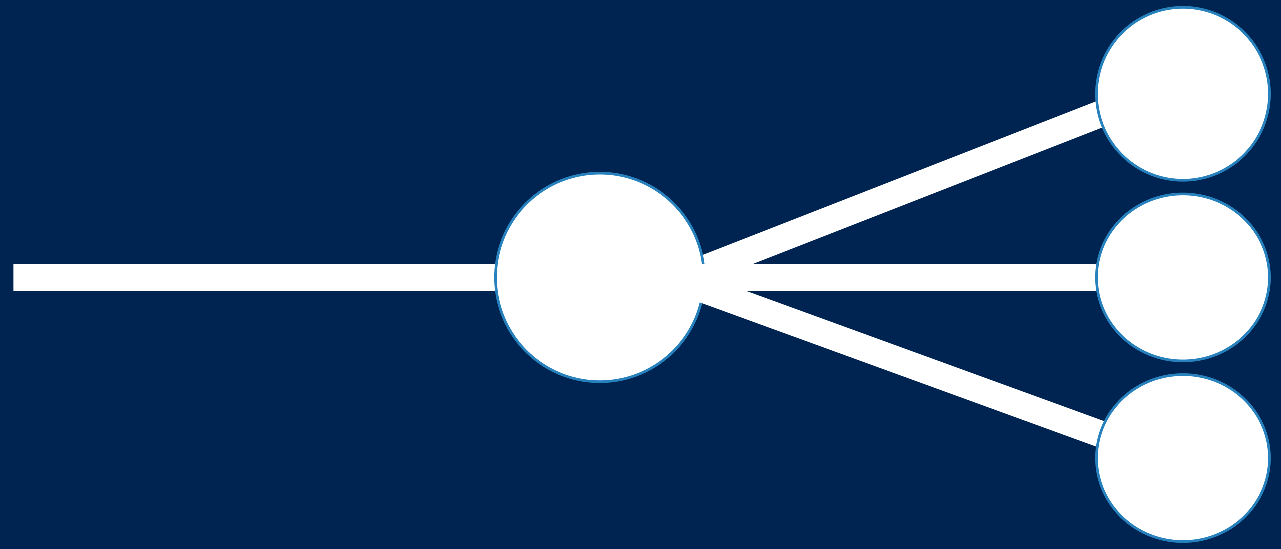

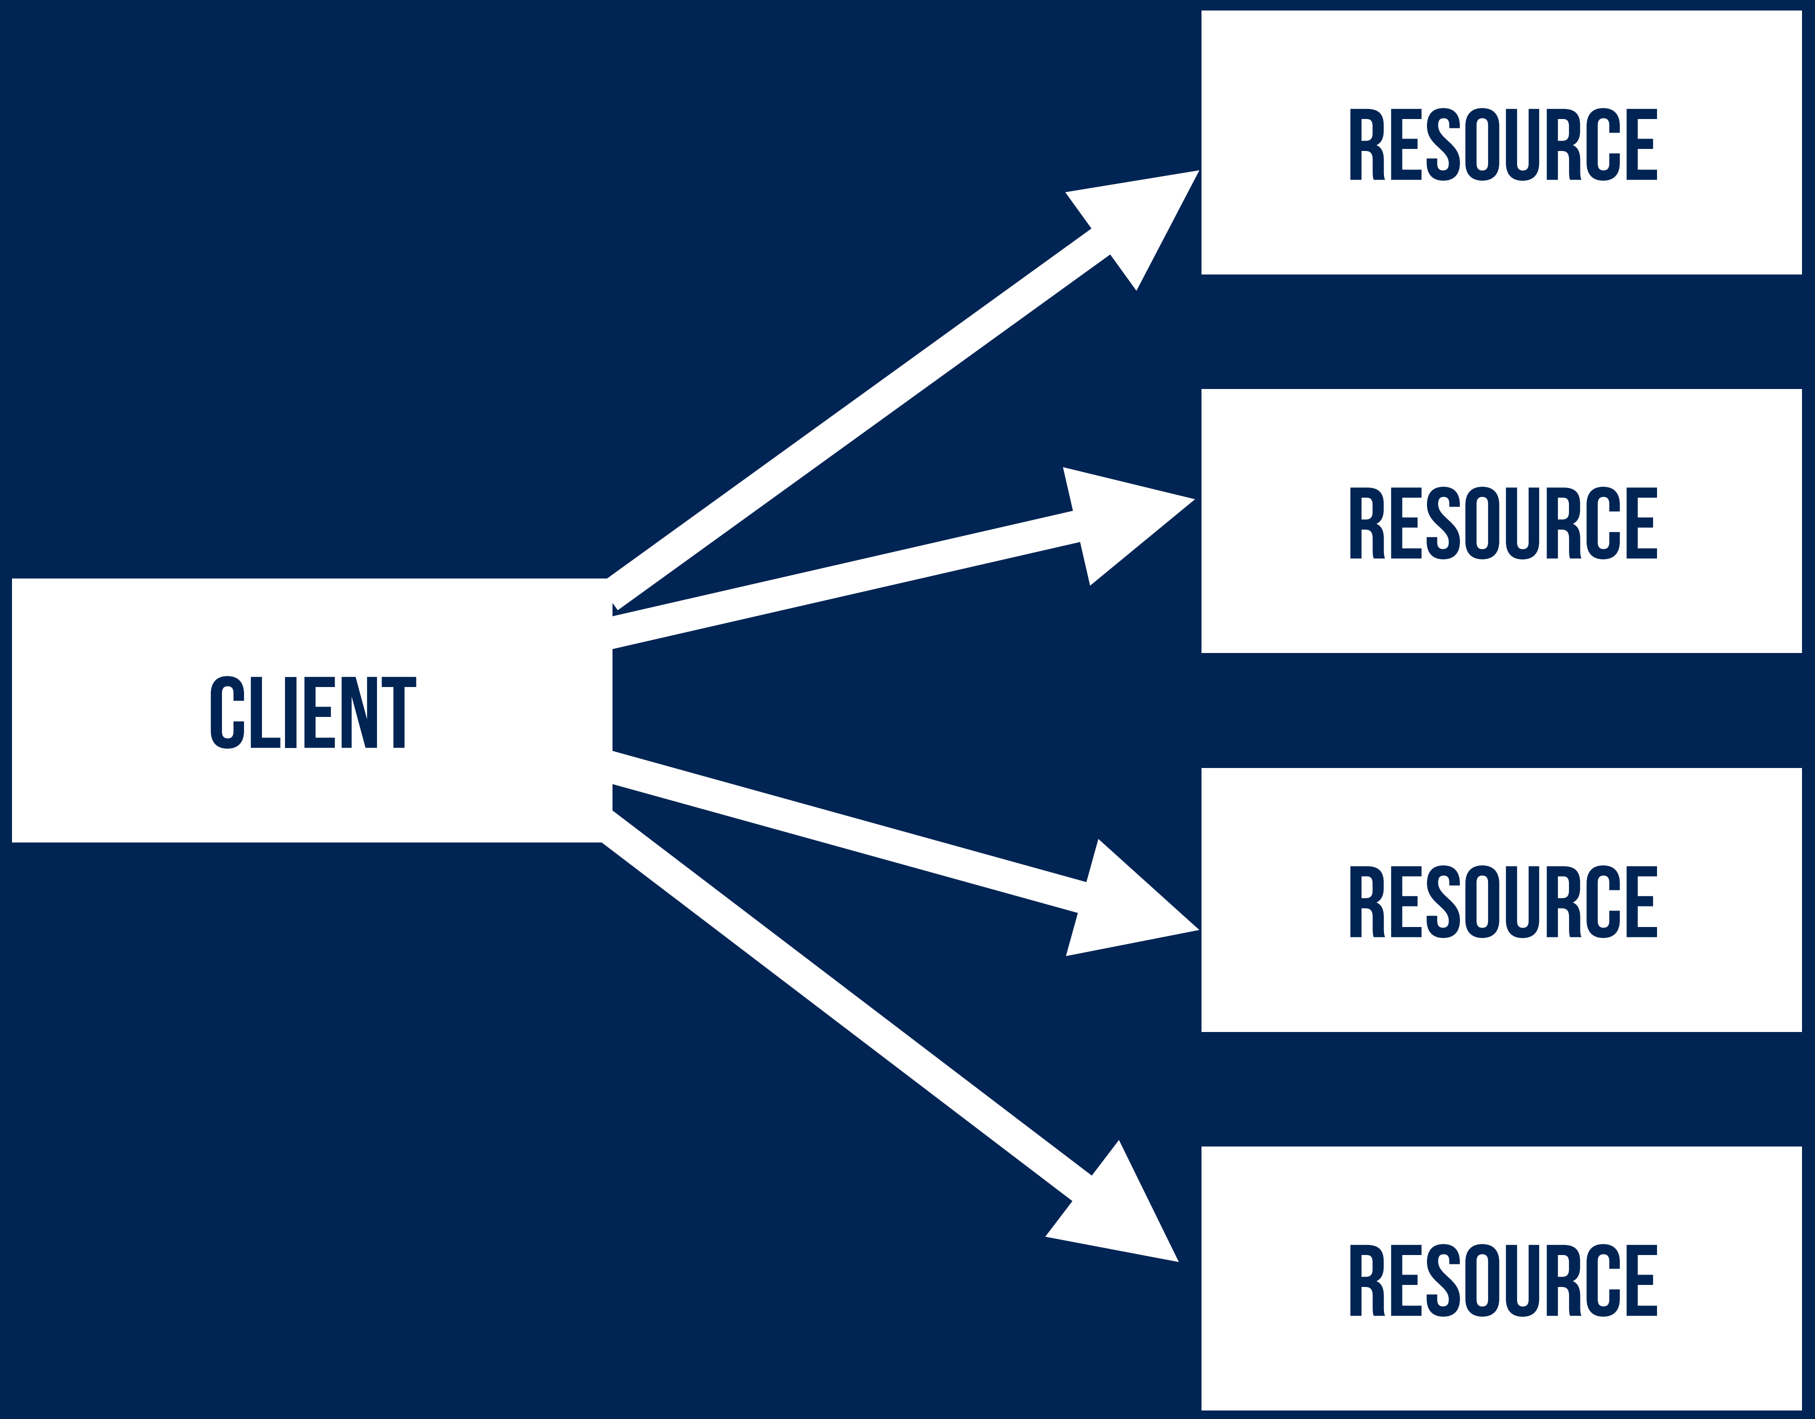

(C OR R1 OR R2 OR R3 OR R4)

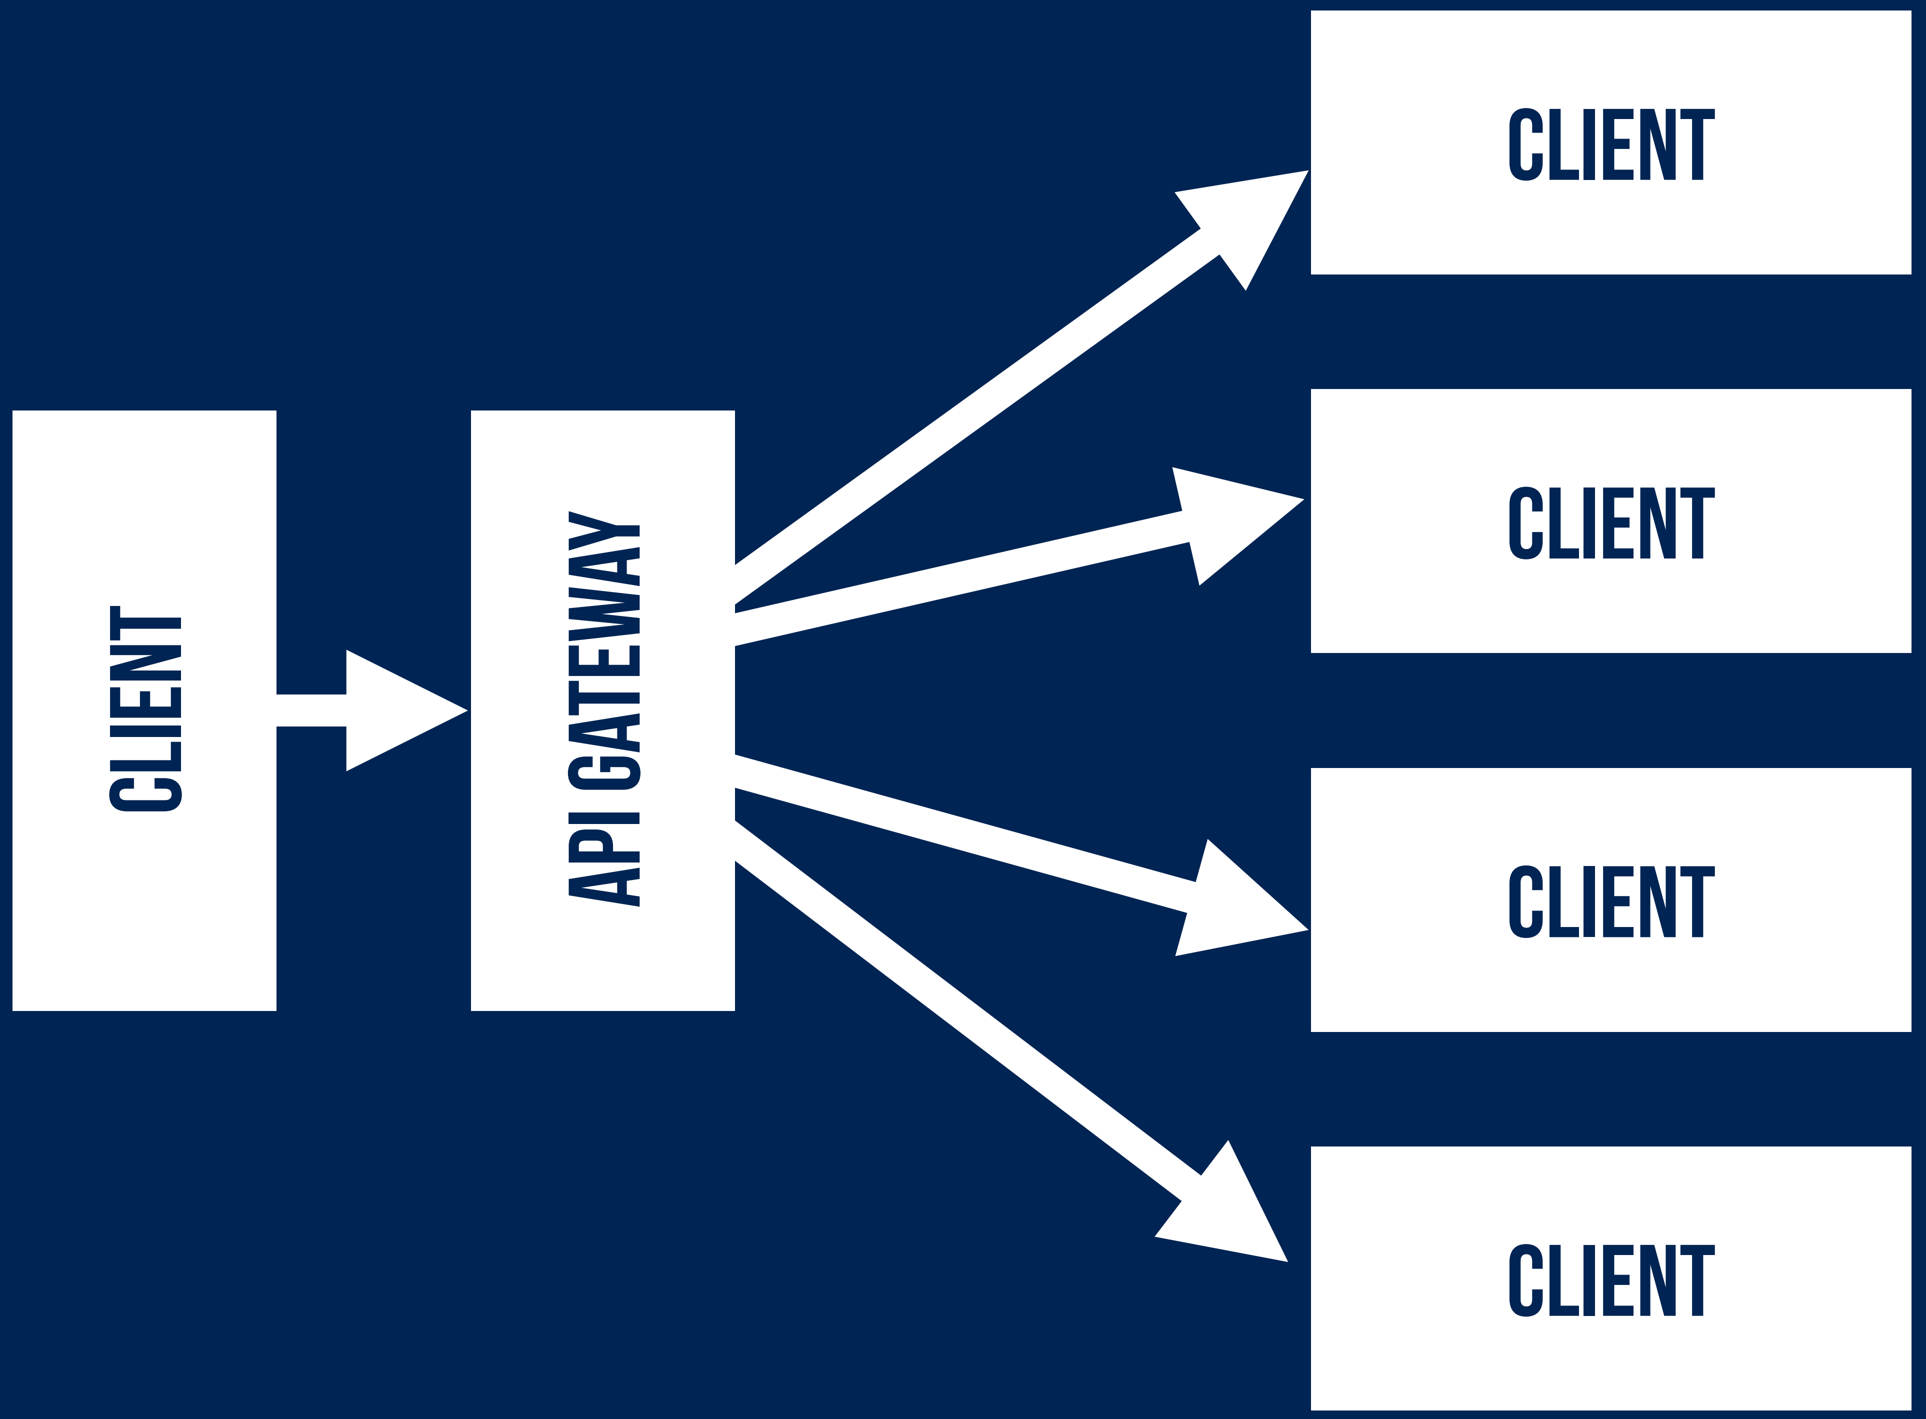

(C OR A) AND (R1 OR R2 OR R3 OR R4)

RUNNING DISTRIBUTED ARCHITECTURES

LOAD

PERFORMANCE

STRESS

ROBUST

AT 12:24 PM PACIFIC TIME ON DECEMBER 24 **NETWORK TRAFFIC** STOPPED ON A FEW ELBS….. AT AROUND 3:30 PM ON DECEMBER 24, NETWORK TRAFFIC STOPPED ON ADDITIONAL ELBS …..NETFLIX IS DESIGNED TO HANDLE FAILURE OF ALL OR PART OF A SINGLE AVAILABILITY ZONE IN A REGION AS WE RUN ACROSS THREE ZONES AND OPERATE WITH NO LOSS OF FUNCTIONALITY ON TWO. WE ARE **WORKING ON WAYS OF EXTENDING OUR RESILIENCY** TO HANDLE PARTIAL OR COMPLETE REGIONAL OUTAGES.

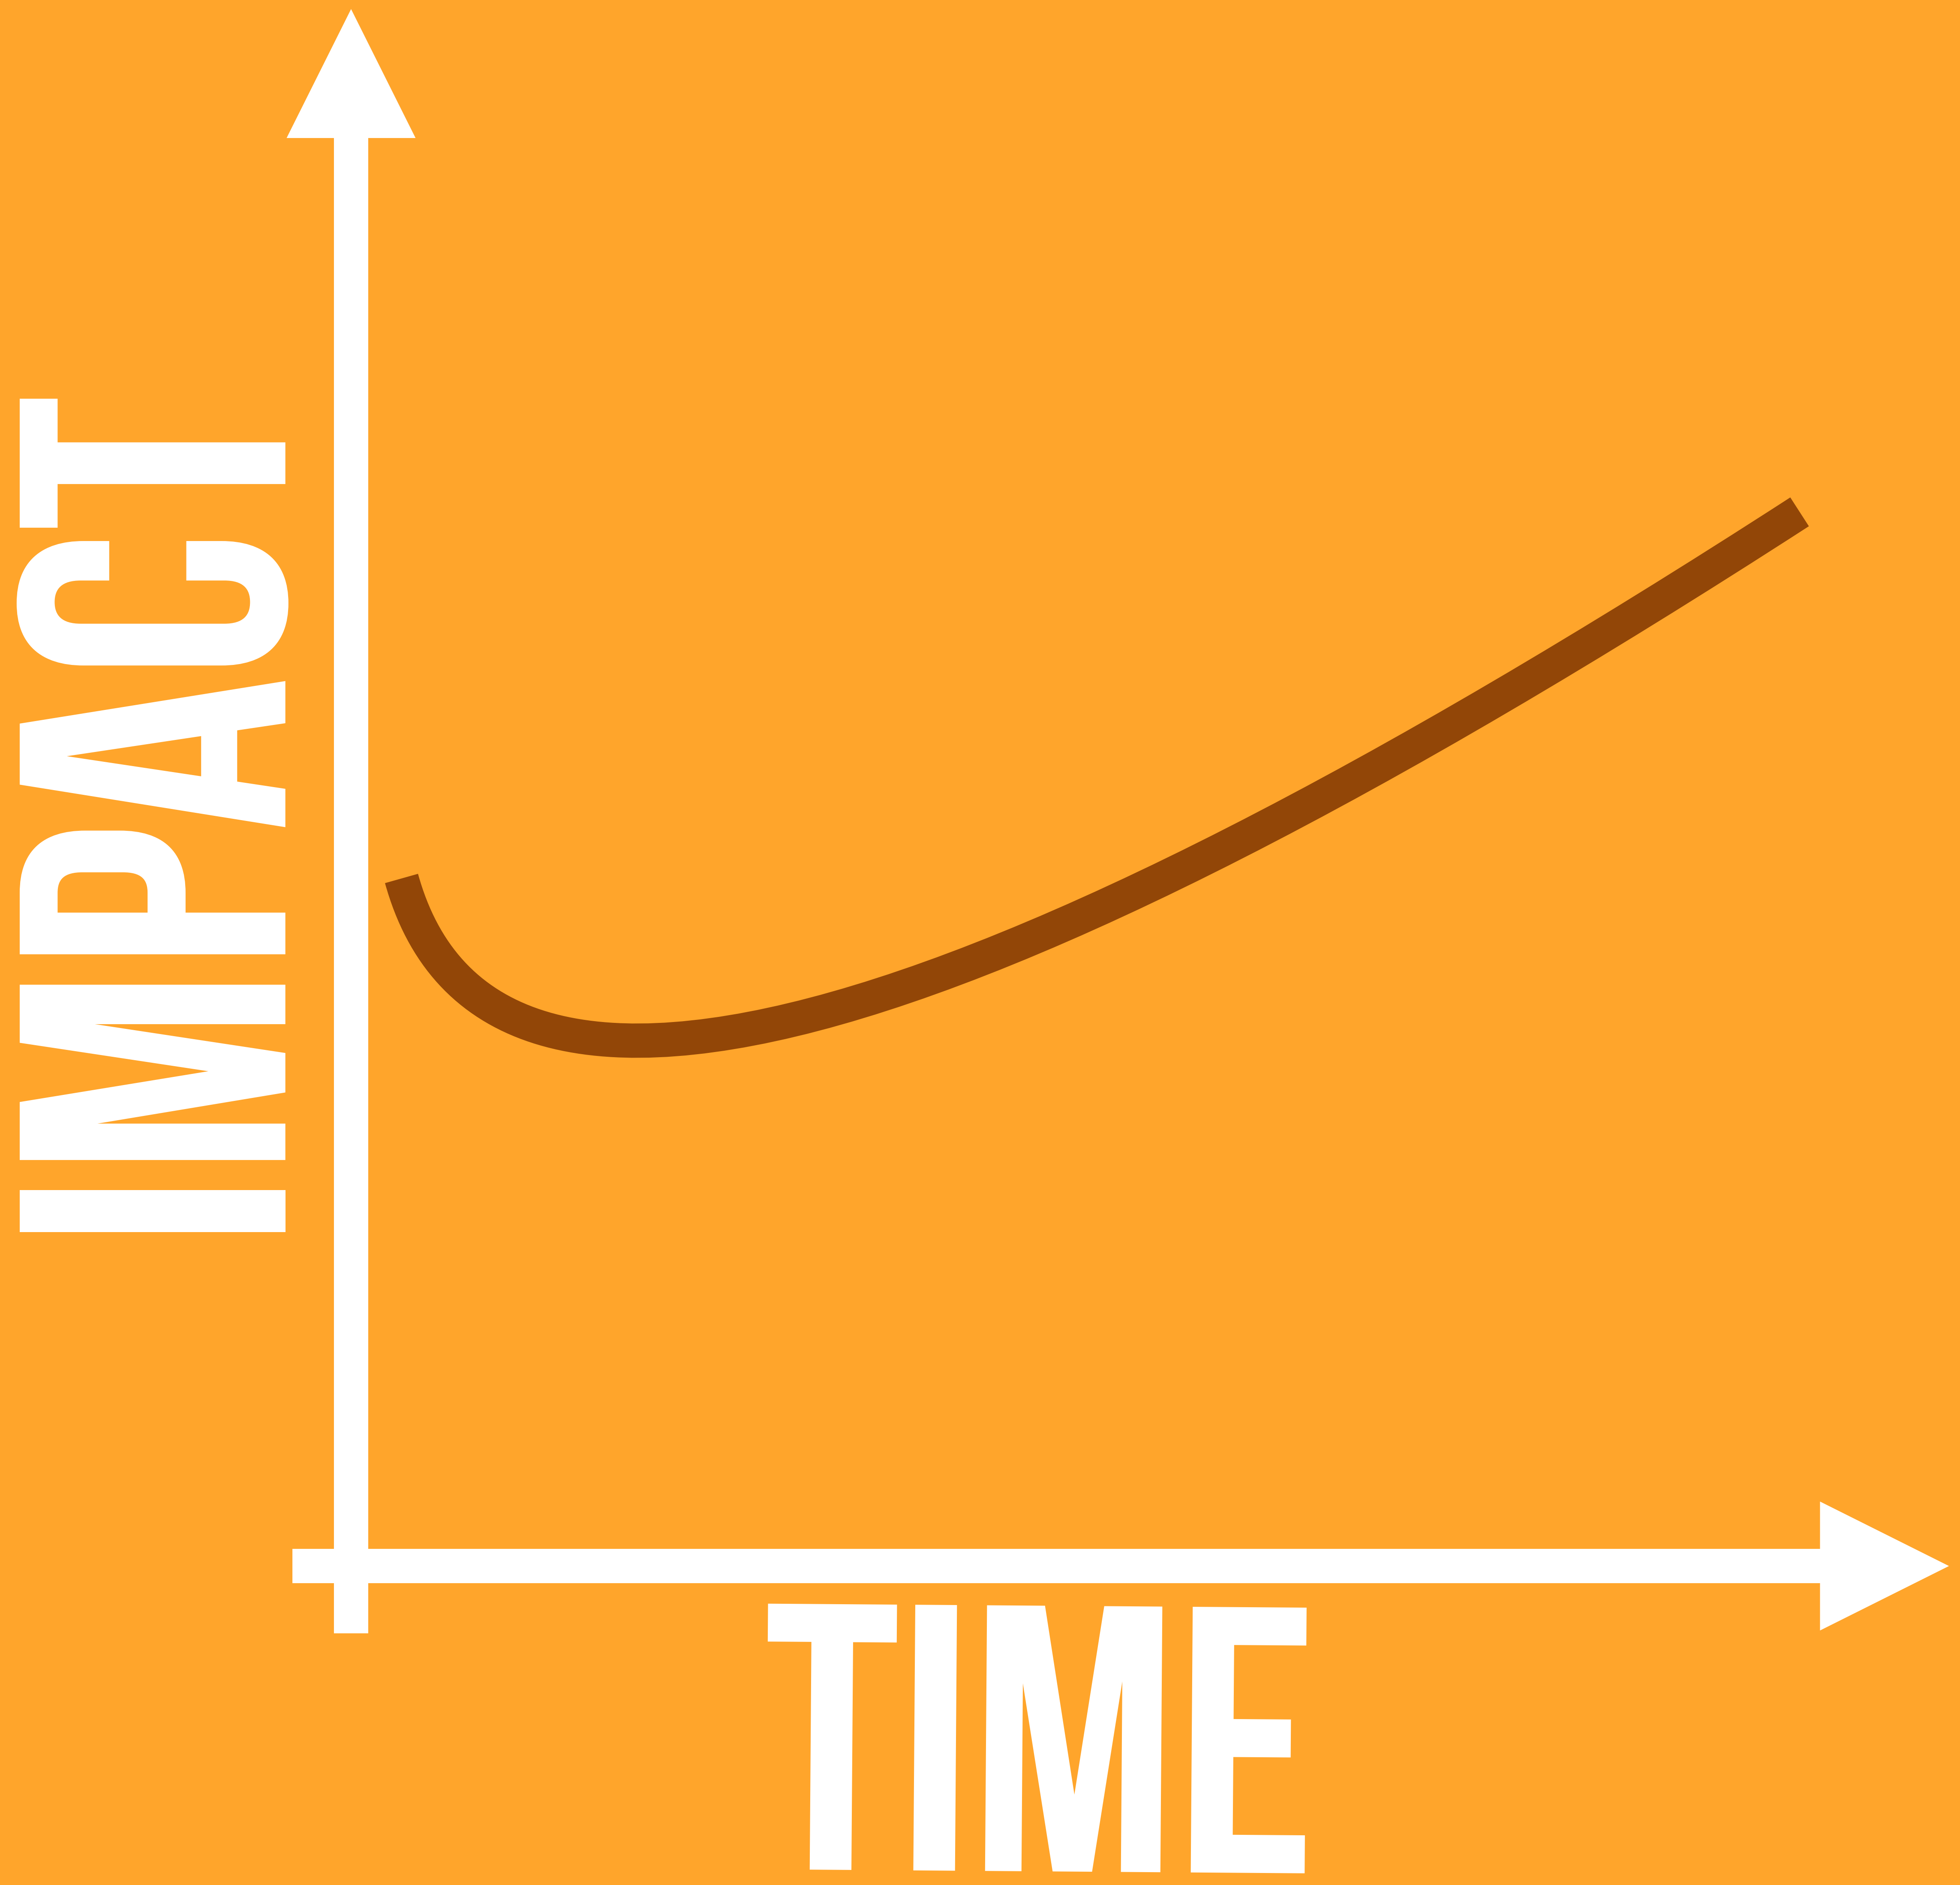

HORMESIS

NETFLIX
SIMIAN
ARMY

ANTI-FRAGILE

ON SUNDAY, AT 2:19AM PDT, THERE WAS A BRIEF **NETWORK DISRUPTION** THAT IMPACTED........
SO, WHEN THE **NETWORK DISRUPTION** OCCURRED ON SUNDAY MORNING, AND A NUMBER OF STORAGE SERVERS SIMULTANEOUSLY REQUESTED THEIR MEMBERSHIP DATA,.....
BY 2:37AM PDT, THE **ERROR RATE** IN CUSTOMER REQUESTS TO DYNAMODB HAD RISEN FAR BEYOND ANY LEVEL EXPERIENCED IN THE LAST 3 YEARS......
AFTER SEVERAL FAILED ATTEMPTS AT ADDING CAPACITY, AT 5:06AM PDT, WE DECIDED TO **PAUSE REQUESTS** TO THE METADATA SERVICE.

ANTI-FRAGILE

DESPITE BEING RUN ENTIRELY FROM AWS' CLOUD PLATFORM THE ONLINE STREAMING GIANT NETFLIX REPORTS A QUICK RECOVERY FROM SUNDAY'S DISRUPTION - **DEMONSTRATING THE IMPORTANCE OF ITS APPROACH OF BUILDING CLOUD-BASED SYSTEMS TO "FAIL".**

AWS:REBOOT

CASSANDRA

2700+ NODES

218 REBOOTED

22 DEAD

THANKS

AGIM EMRULI - MIMACOM

@AEMRULI