

Scaling Event Sourcing for Netflix Downloads

Phillipa Avery
Robert Reta

NETFLIX

Presenters



Phillipa Avery

Senior Software Engineer

Project Technical Lead and Engineer
Downloads License Accounting

pavery@Netflix.com /  @PhillipaAvery



Robert Reta

Senior Software Engineer

Event Sourcing System Architect
Downloads License Accounting

rreta@netflix.com /  @rreta04

11/30/2016

THE VERGE TECH · SCIENCE · CULTURE · CARS · REVIEWS · LONGFORM · VIDEO · MORE ·

APPS \ HOME \ ENTERTAINMENT


Netflix finally lets you download shows and movies to watch offline

No streaming necessary

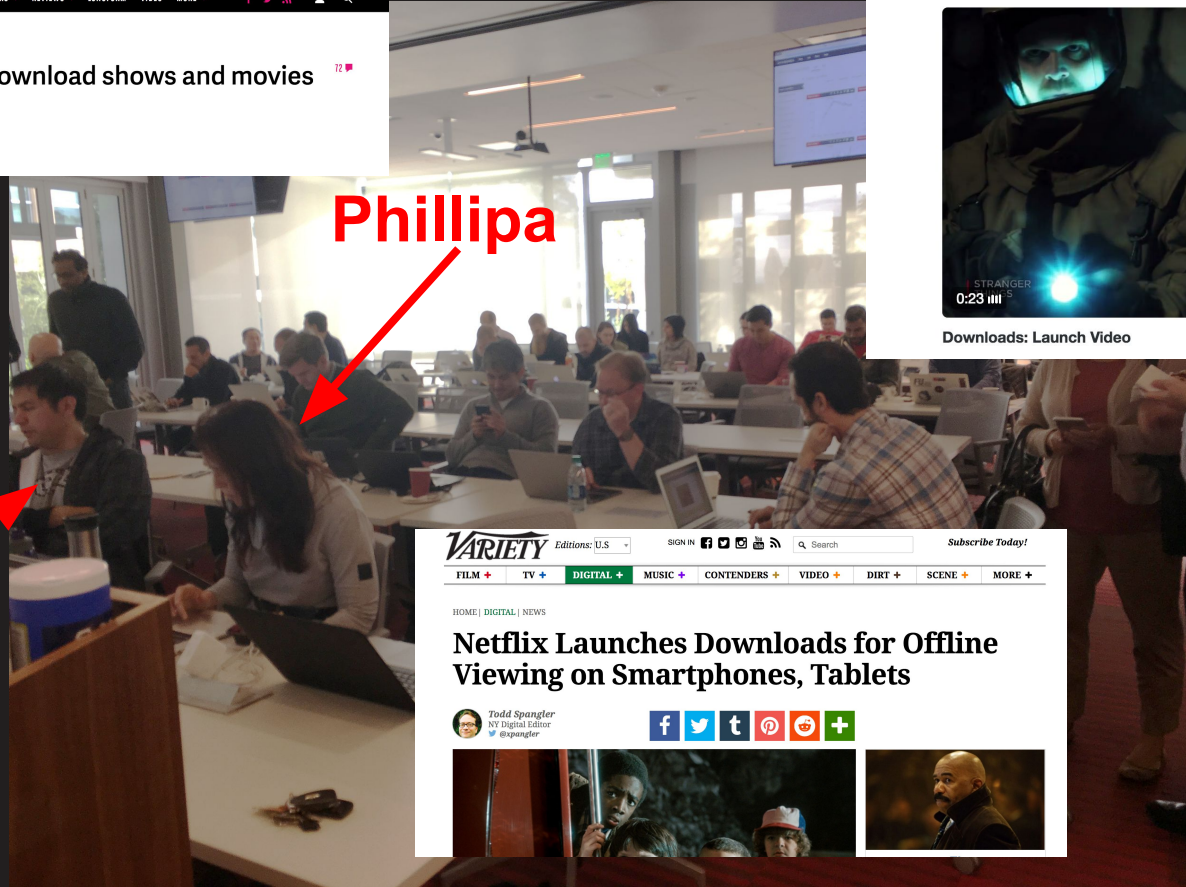
by Chris Welch | Nov 30, 2016, 9:14am EST

N **Netflix US** @netflix · 26m

Airplane mode. Road trip mode. Stuck-in-the-subway-for-20-minutes mode. Your favorite stories are now available for download any time.



Downloads: Launch Video



Phillipa



Robert





VARIETY Editions: U.S · SIGN IN Search [Subscribe Today!](#)

FILM + TV + **DIGITAL +** MUSIC + CONTENTENDERS + VIDEO + DIRT + SCENE + MORE +

HOME | DIGITAL | NEWS

Netflix Launches Downloads for Offline Viewing on Smartphones, Tablets

Todd Spangler
NY Digital Editor
@xspangler

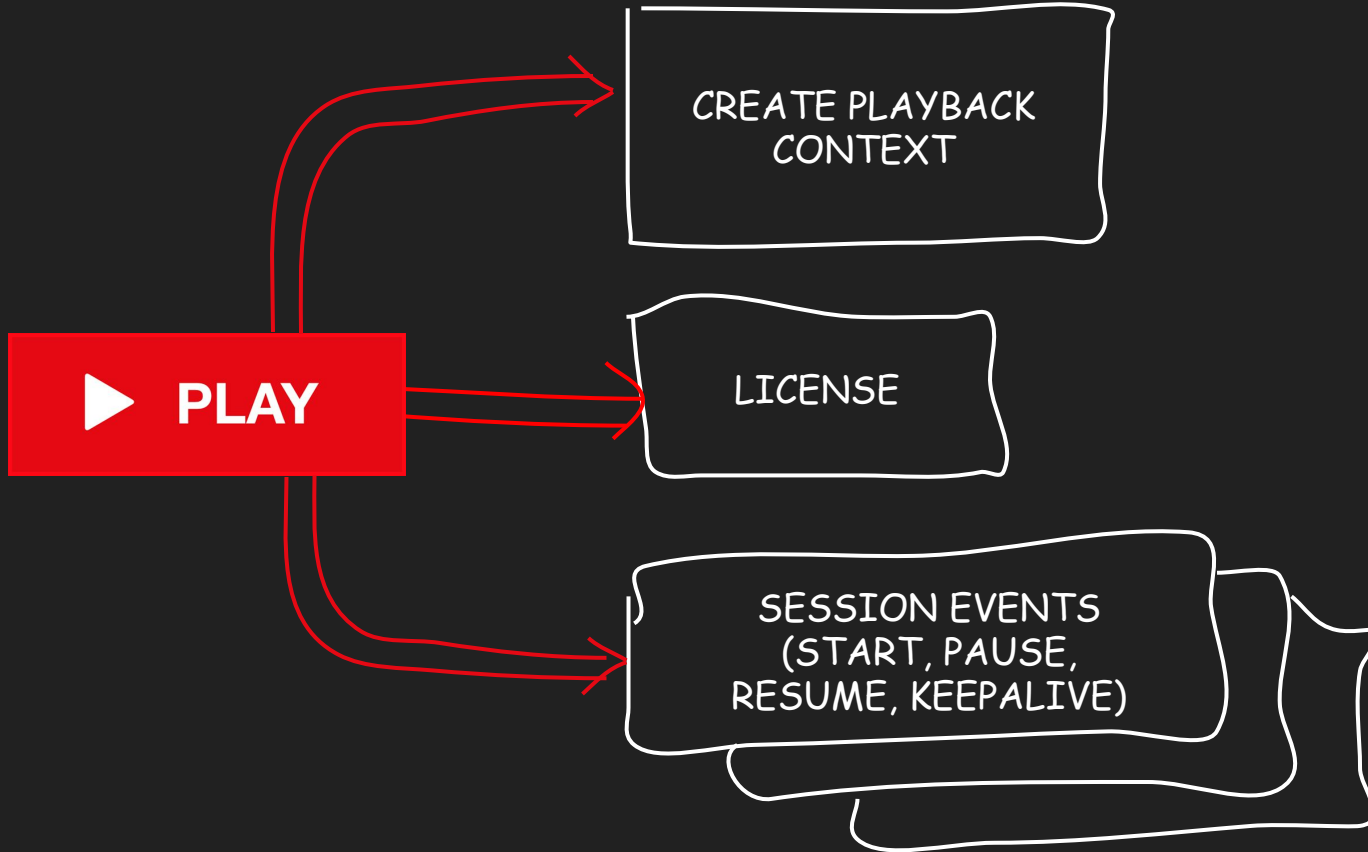


NETFLIX


Overview

1. Why do we need a Downloads License Accounting Service?
2. Event Sourcing overview.
3. Deep dive into the Event Sourcing Architecture.
4. What's it like working with the License Accounting Service after release?

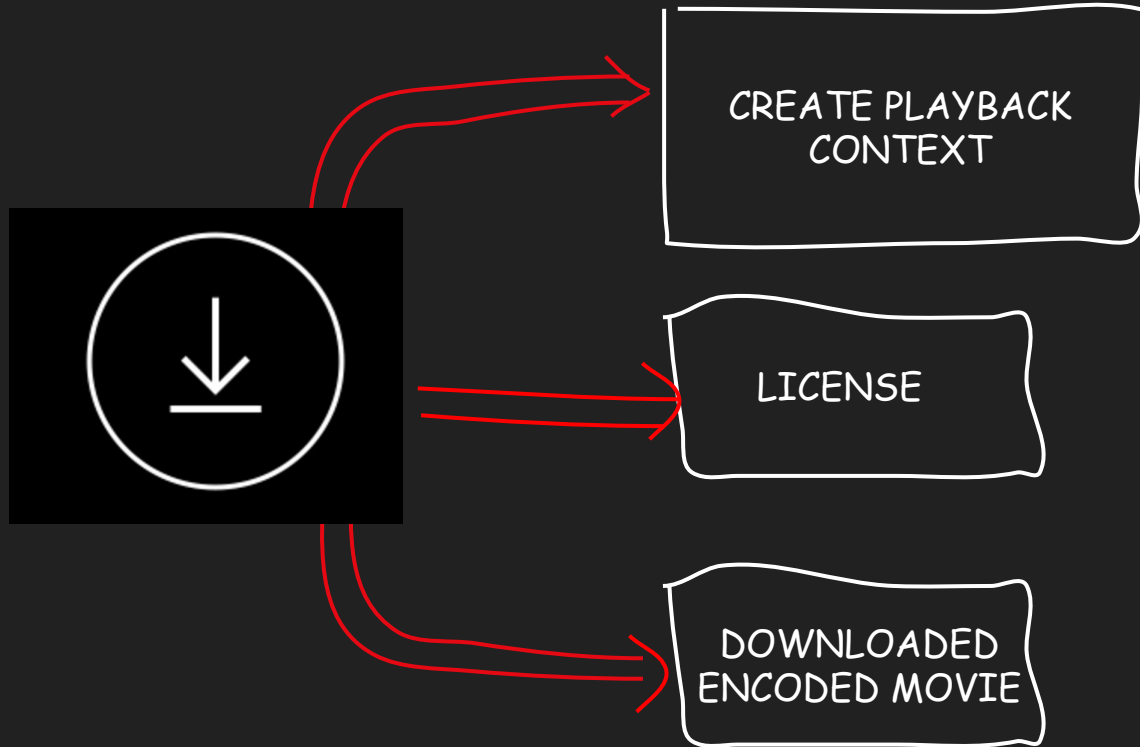
Streaming Playback Lifecycle



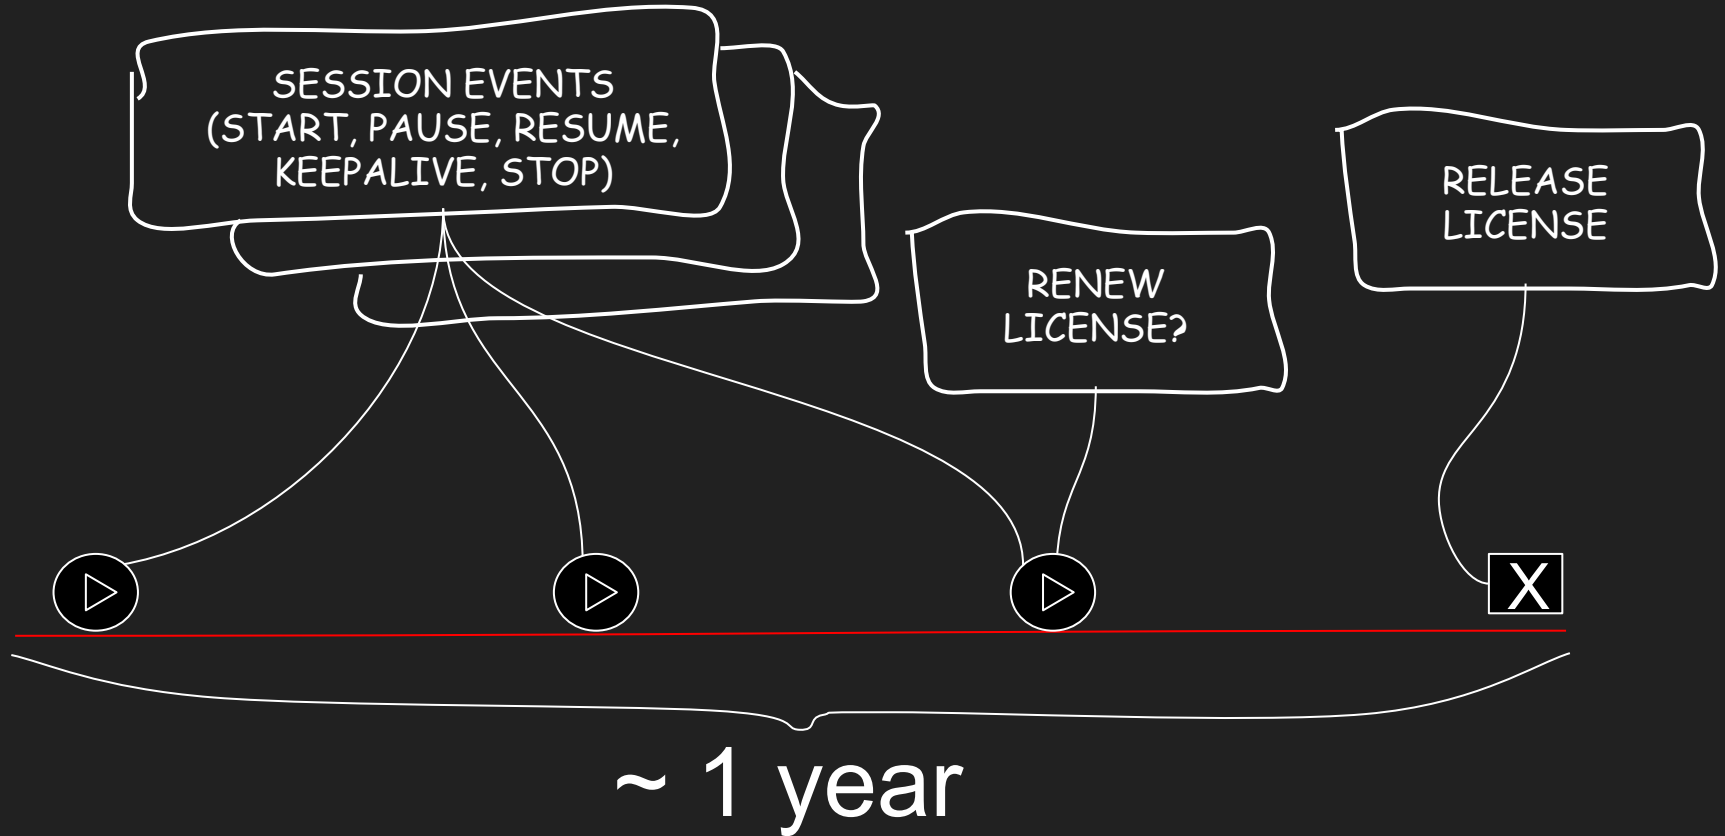
Streaming Playback Lifecycle



Downloads Playback Lifecycle

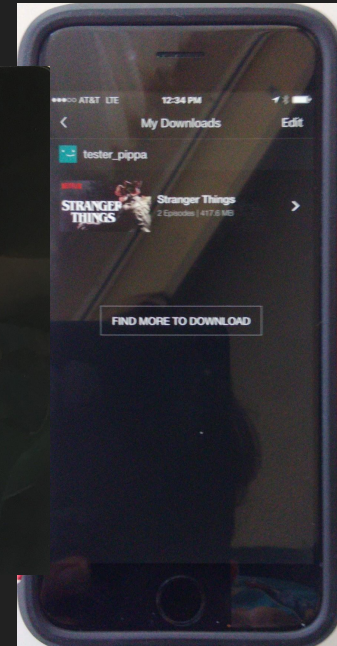



Downloads Playback Lifecycle




Download Business Requirements

- Devices with downloads



Download Business Requirements

- Devices with downloads
- Downloads per studio



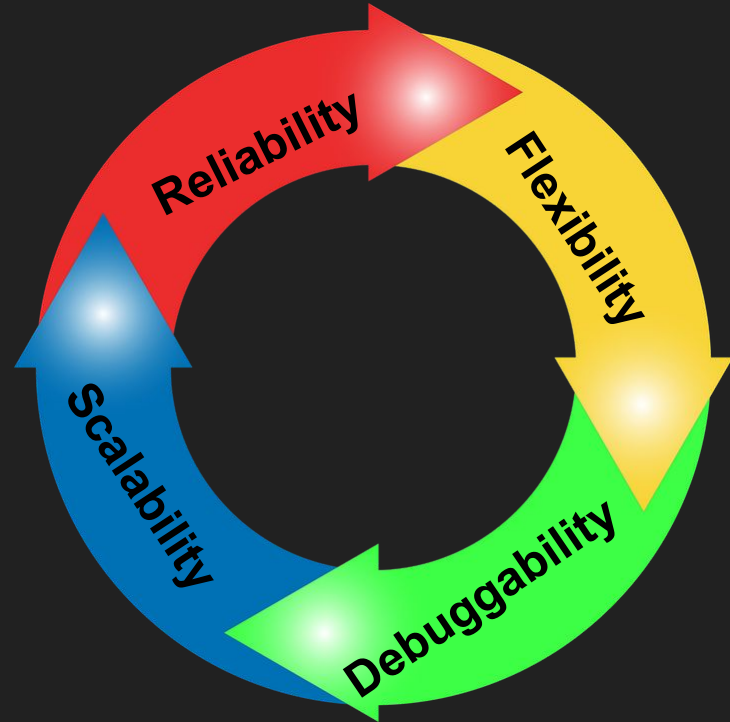
Download Business Requirements

- Devices with downloads
- Downloads per studio
- Movie downloads or playback over year



License Accounting Service Requirements

- ☐ Flexible
- ☐ Debuggable
- ☐ Reliable
- ☐ Scalable



License Accounting Service Requirements...

✓ Flexible: data model can be changed

- ❑ RDBMS

- ❑ Document Model



License Accounting Service Requirements...

- ✓ Flexible: data model can be changed
 - ❑ RDBMS
 - ❑ Document Model
- ✓ Debuggable
 - ❑ Event Sourcing



License Accounting Service Requirements...

✓ Flexible: data model can be changed

- ❑ RDBMS

- ❑ Document Model

✓ Debuggable

- ❑ Event Sourcing

✓ Reliable

- ❑ Fallbacks



License Accounting Service Requirements...

✓ Flexible: data model can be changed

- ❑ RDBMS

- ❑ Document Model

✓ Debuggable

- ❑ Event Sourcing

✓ Reliable

- ❑ Fallbacks

✓ Scalable





Event Sourcing

Domain Model






Commands

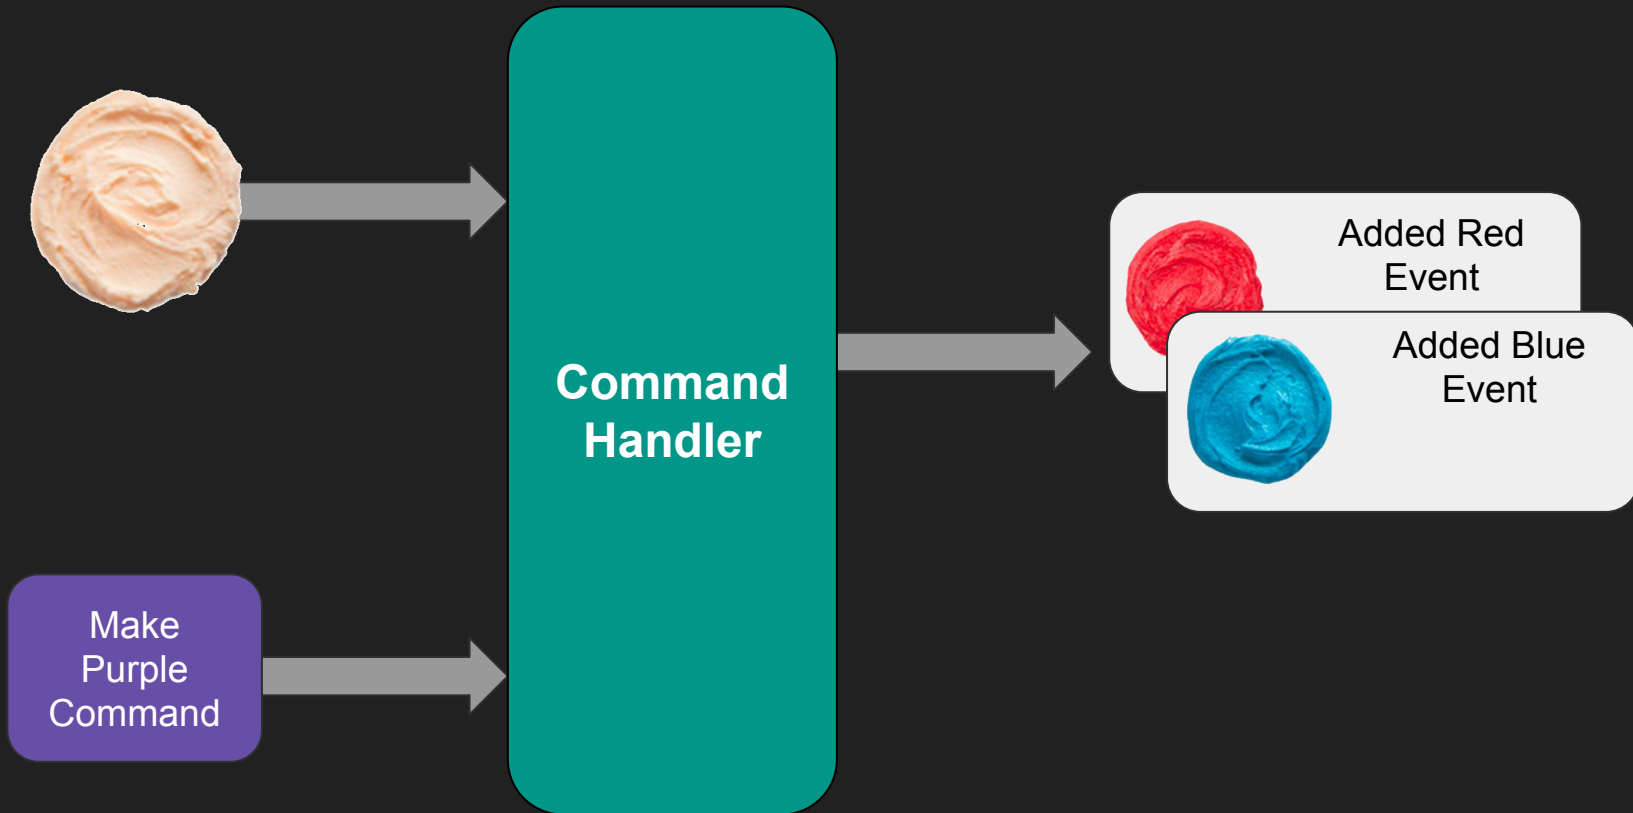


Make
Purple
Command

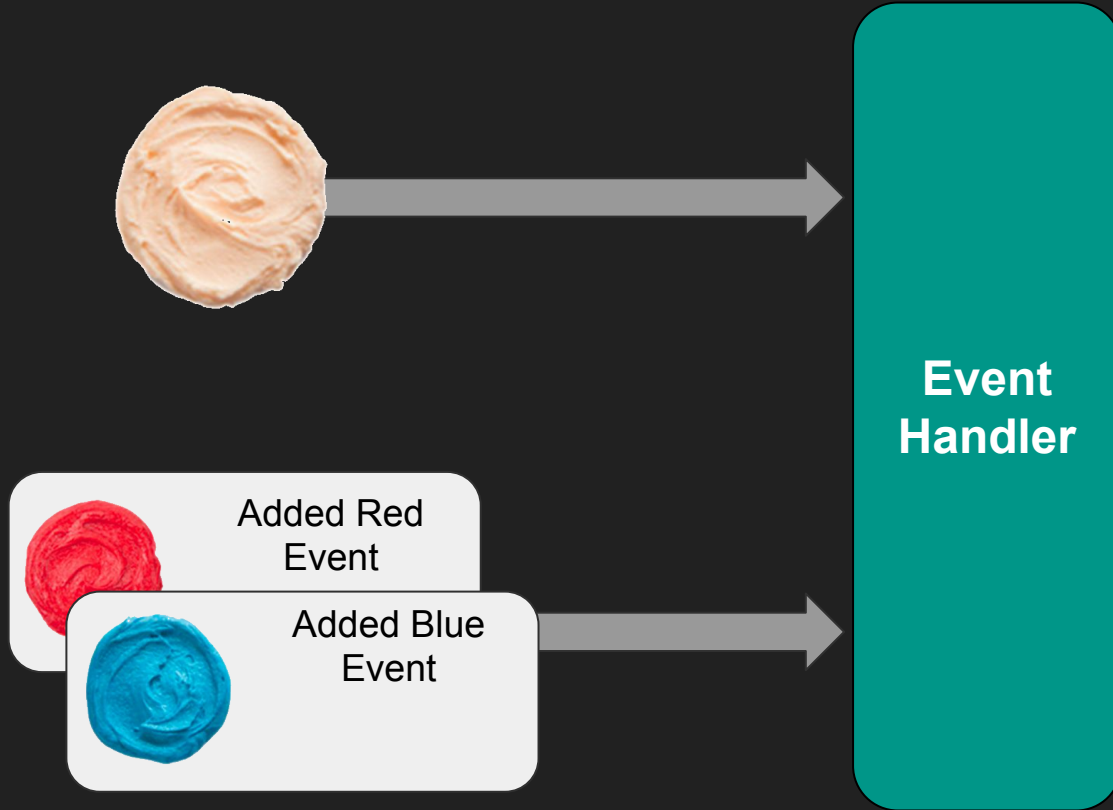
Commands




Commands




Events



Events



Added Blue Event

A blue circular icon with a textured, marbled appearance, representing an event.


Domain Model



Aggregate




Commands






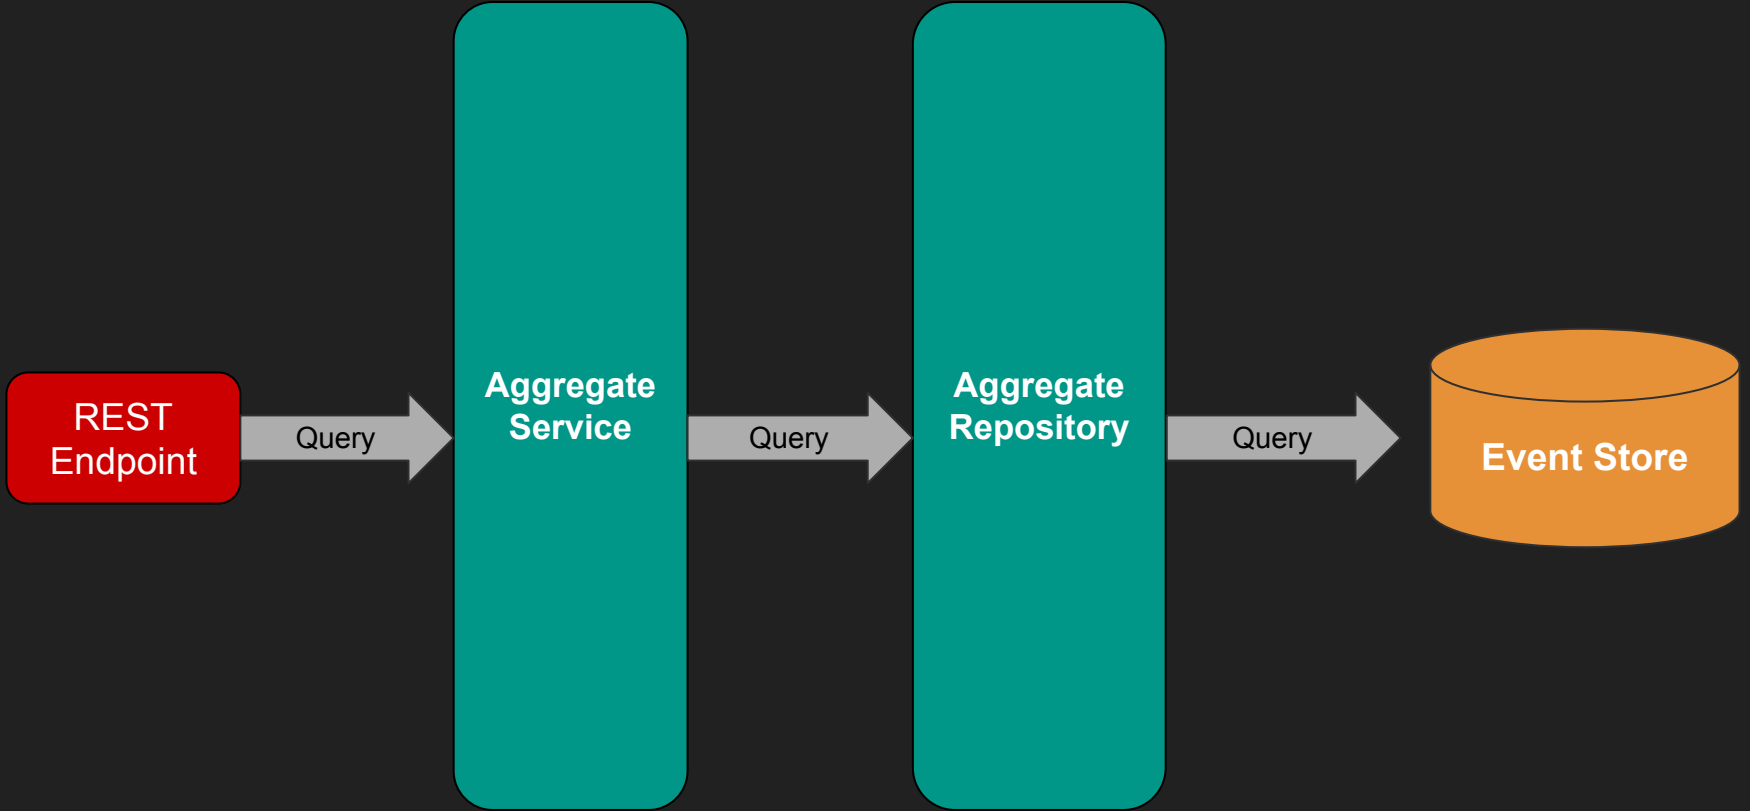
**Command
Handler**

**Make Red
Command**

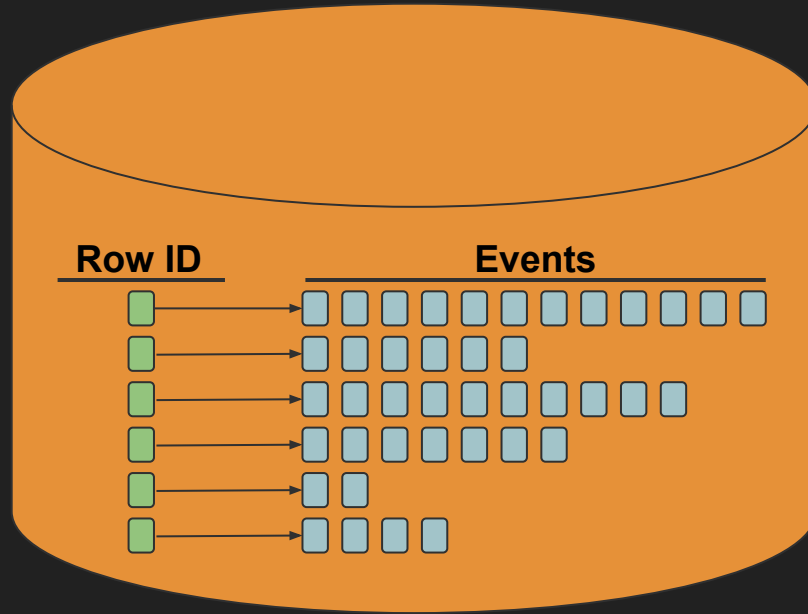


Commands






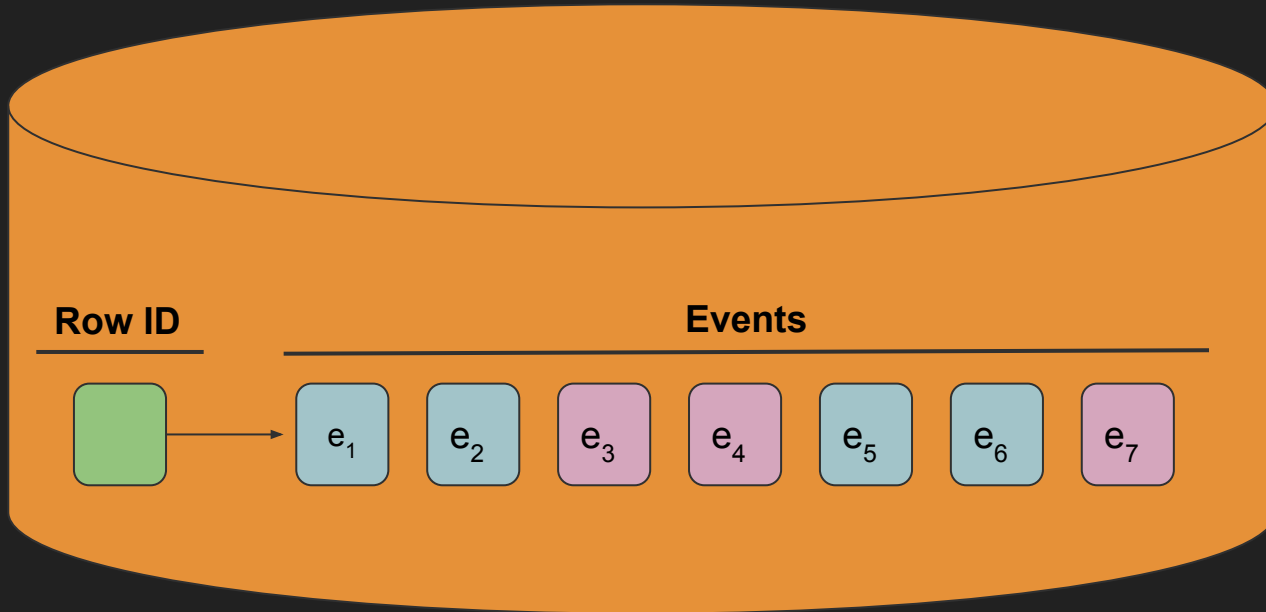
Event Store



Event Store




Event Store




Aggregate id 1



Aggregate id 2





Aggregate
Repository




```
select * from events  
where rowId = █
```

Event Store






Aggregate Repository



Aggregate id 1



Aggregate id 2



Aggregate Repository

Aggregate 1

Aggregate id 1

e_1

e_2

e_5

e_6

Aggregate 2


Aggregate id 2

e_3



e_4

e_7

Aggregate Repository




Aggregate id 1



Aggregate id 2



Aggregate Repository



Aggregate id 1




Aggregate id 2




Aggregate Repository

Aggregate id 1




Aggregate id 2



Aggregate Repository

Aggregate id 1



Aggregate id 2



Aggregate Repository

Aggregate 1

Aggregate 2

Aggregate Repository

```
graph LR; A[Aggregate 1] --> CH[Command Handler]; C[Command] --> CH; subgraph AR [Aggregate Repository]; CH; end
```

The diagram illustrates the Aggregate Repository pattern. A large teal rounded rectangle represents the repository, containing a light blue rounded rectangle labeled 'Command Handler'. To the left, a red rounded rectangle labeled 'Aggregate 1' and a purple rounded rectangle labeled 'Command' both have grey arrows pointing into the 'Command Handler' box.

Aggregate 1

Command
Handler

Command

Aggregate Repository

Command
Handler



e_8

e_9

Aggregate Repository

Aggregate 1


e_8

e_9

Aggregate Repository




Aggregate Repository



Aggregate Repository

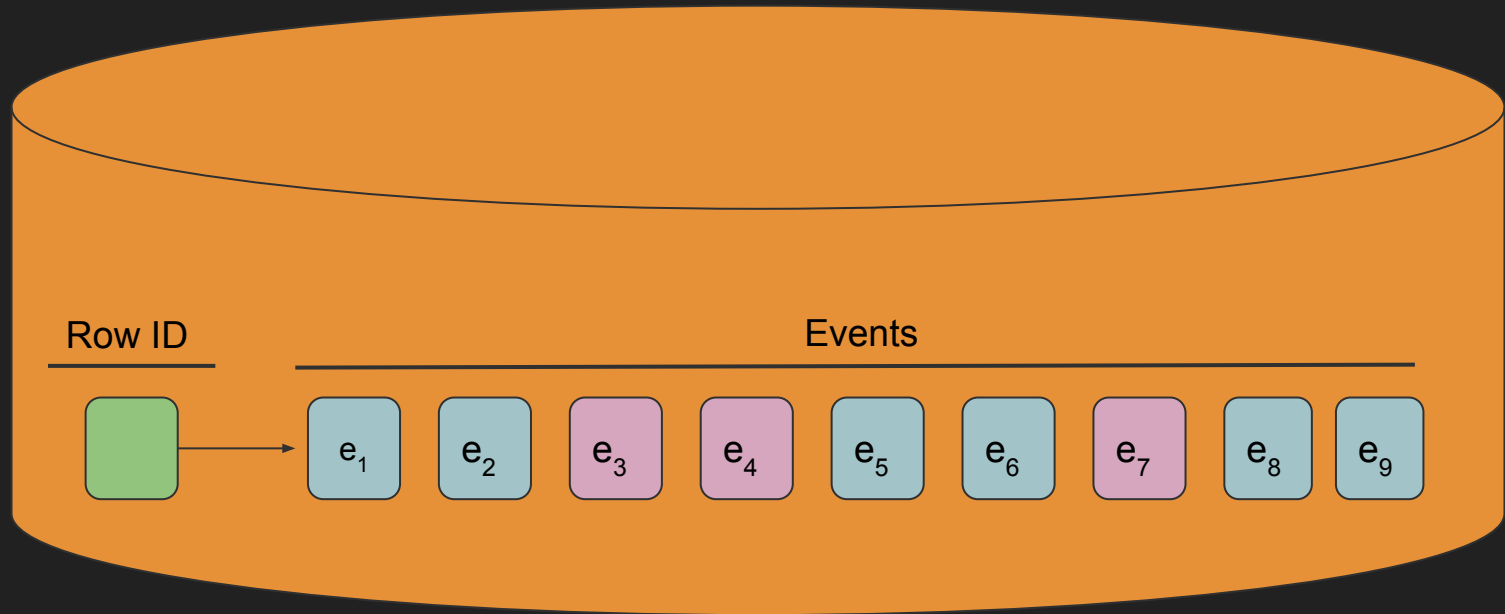
Aggregate 1



append events e_8 e_9

where rowId =

Event Store



Aggregate id 1



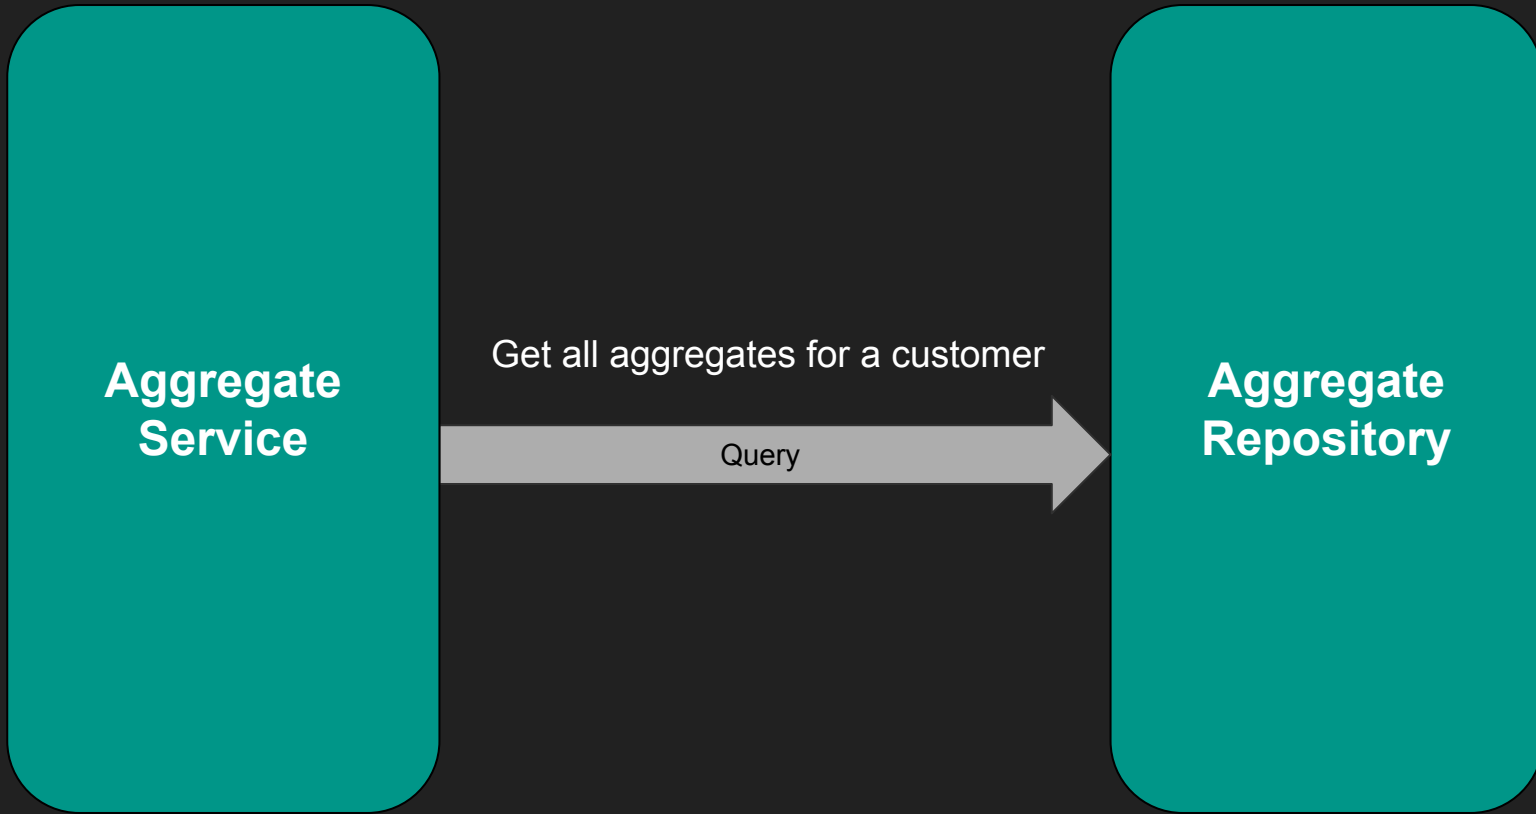
Aggregate id 2

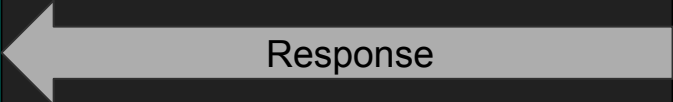
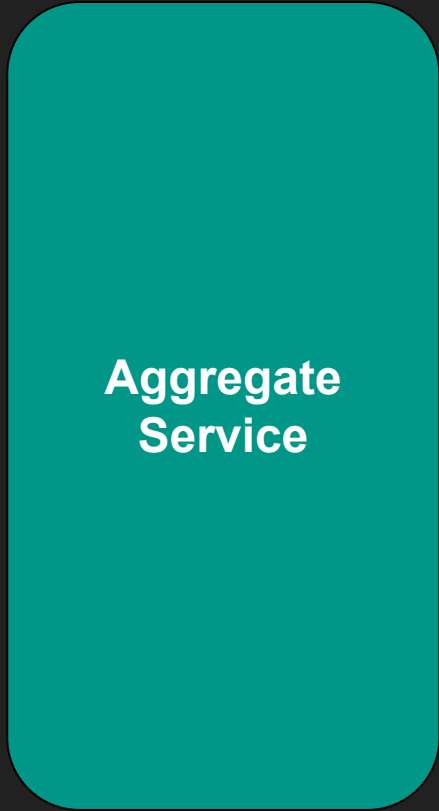
**Aggregate
Service**

Get all aggregates for a customer

Query

**Aggregate
Repository**

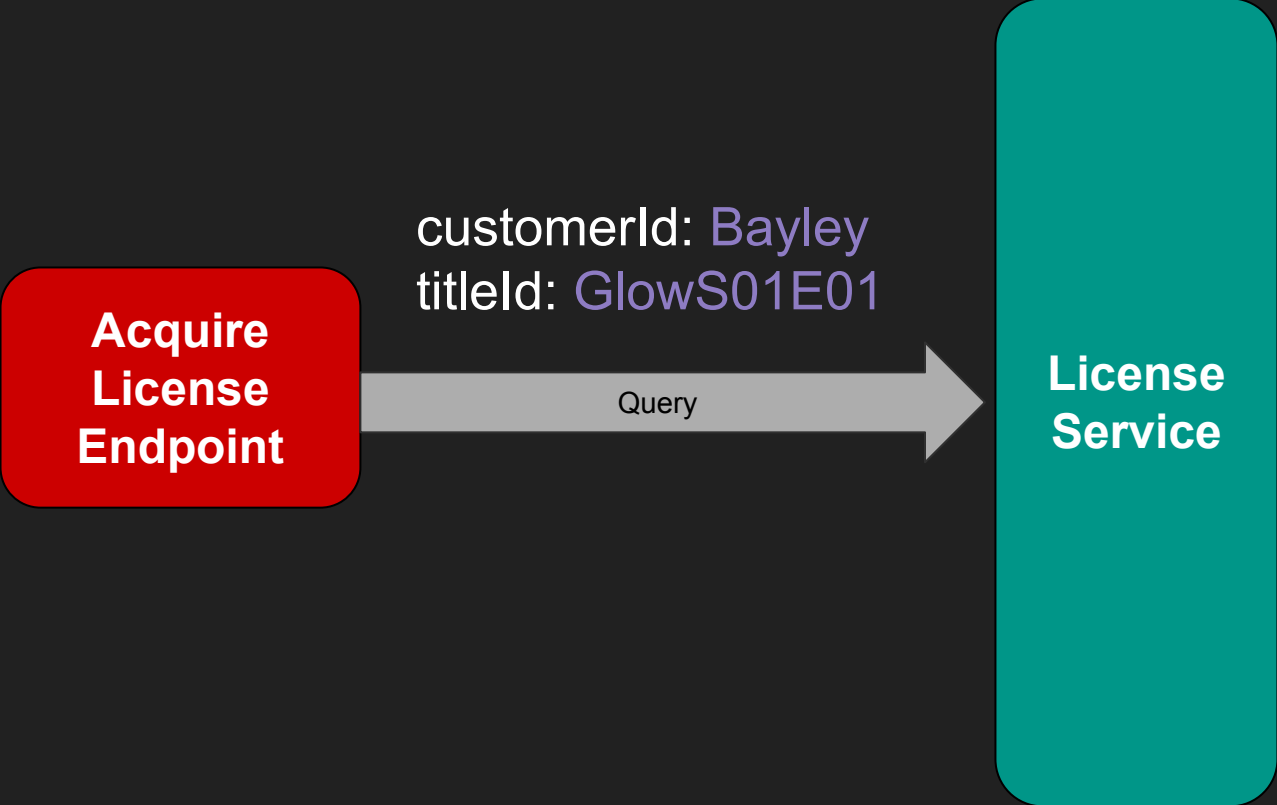





Aggregate


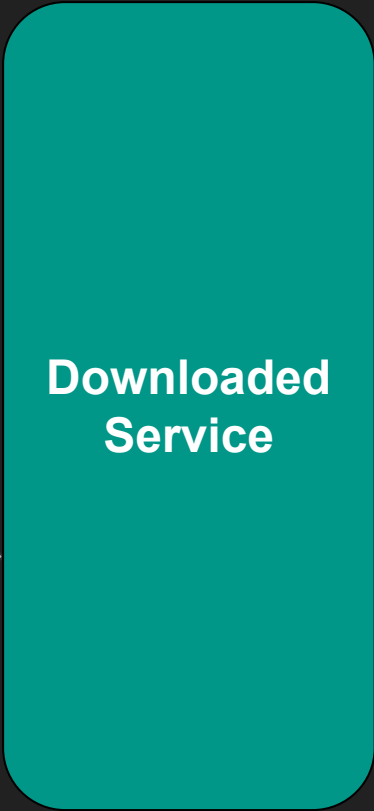
**License
Aggregate**

**Downloaded
Aggregate**





isAllowed?
customerId: Bayley
titleId: GlowS01E01



**Downloaded
Service**

getAggregates

customerId: Bayley
titleId: GlowS01E01
since: 6/27/2016

**Downloaded
Repository**

Query

```
graph LR; A[Downloaded Service] -- Query --> B[Downloaded Repository];
```

The diagram illustrates a query flow from a service to a repository. On the left is a teal rounded rectangle labeled 'Downloaded Service'. On the right is another teal rounded rectangle labeled 'Downloaded Repository'. A grey arrow points from the service to the repository, with the word 'Query' written below it. Above the arrow, the text 'getAggregates' is underlined, followed by three lines of query parameters: 'customerId: Bayley', 'titleId: GlowS01E01', and 'since: 6/27/2016'.

**Downloaded
Service**

Downloaded
Aggregate

customerId: Bayley
titleId: GlowS01E01
date: 2/15/2017


Downloaded
Aggregate

customerId: Bayley
titleId: GlowS01E01
date: 5/25/2017

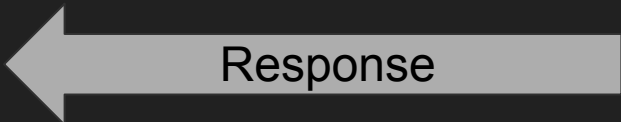
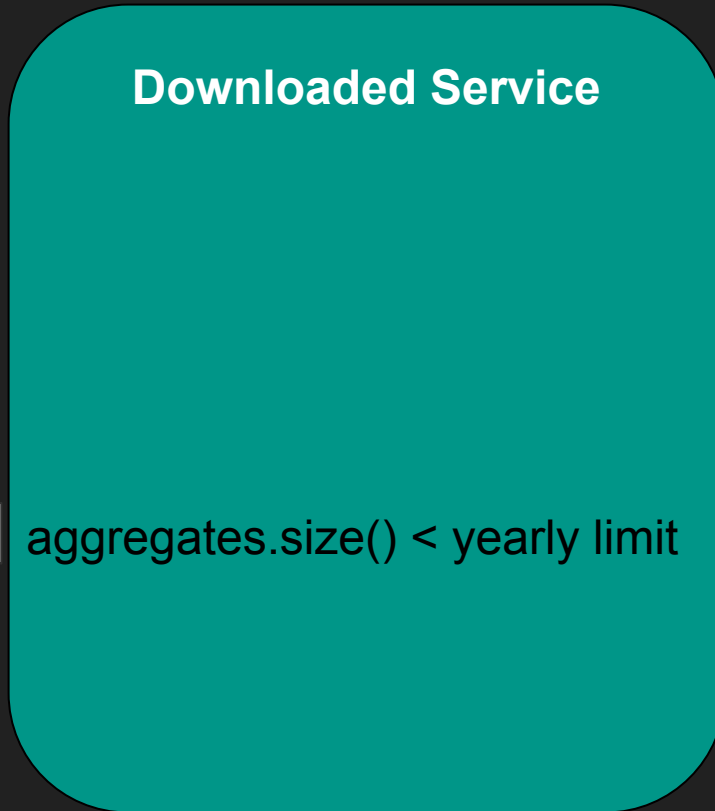
**Downloaded
Repository**


Response

```
graph LR; RS[Downloaded Repository] -- Response --> DS[Downloaded Service]; subgraph Aggregates; direction TB; A1["Downloaded Aggregate  
customerId: Bayley  
titleId: GlowS01E01  
date: 2/15/2017"]; A2["Downloaded Aggregate  
customerId: Bayley  
titleId: GlowS01E01  
date: 5/25/2017"]; end; RS --- Aggregates;
```





isAllowed?





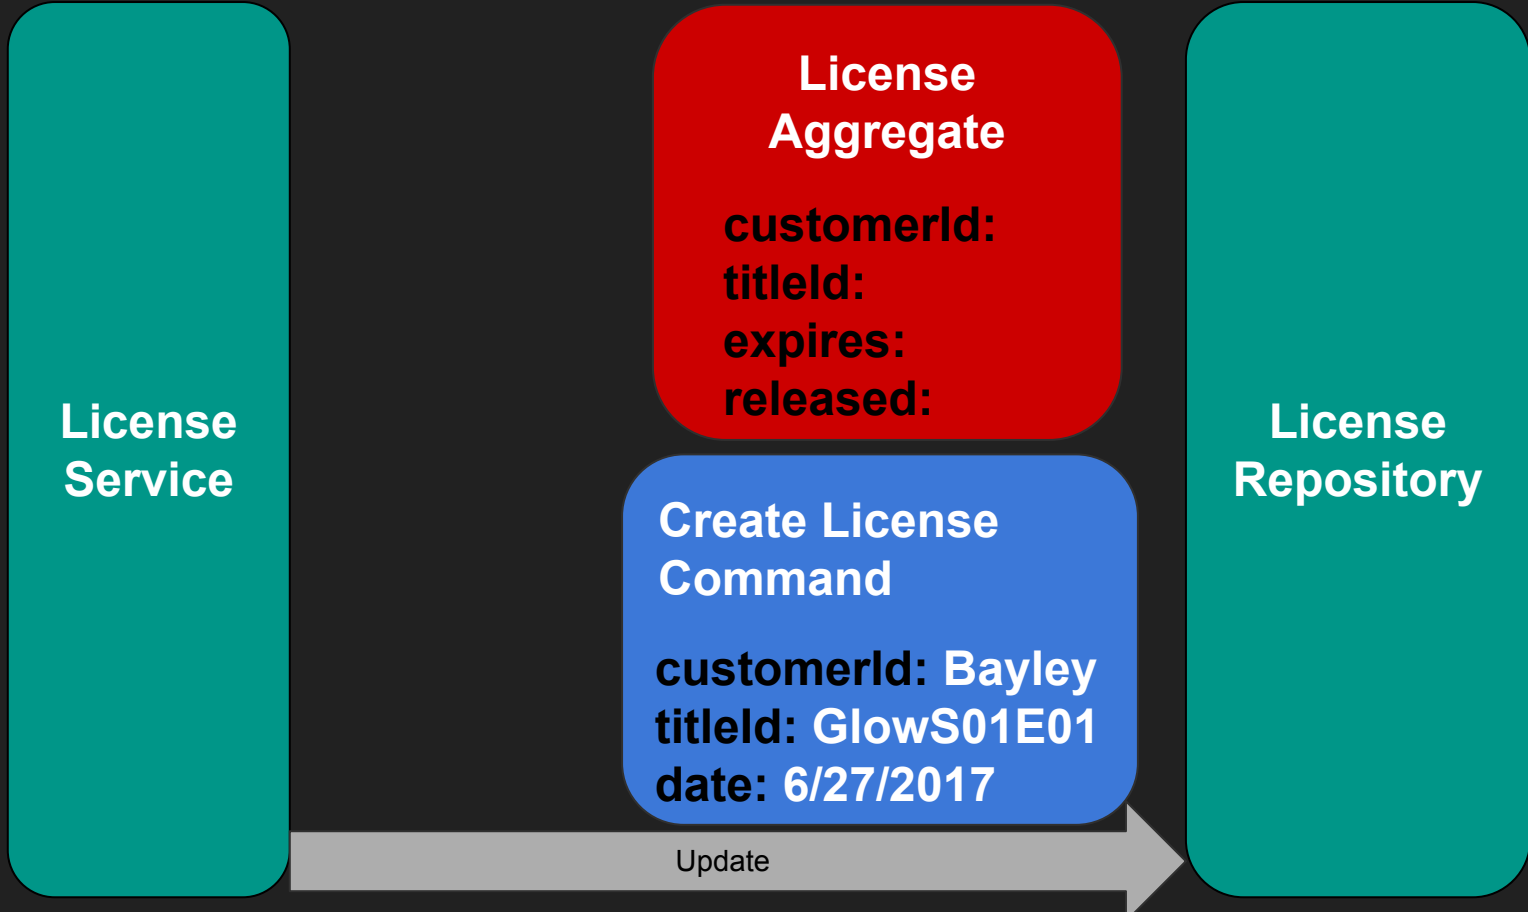
isAllowed?



True



Create Aggregate



License Repository

License Aggregate

customerId:
titleId:
expires:
released:

Create License Command

customerId: Bayley
titleId: GlowS01E01
date: 6/27/2017

**Command
Handler**



License Repository

License Aggregate

customerid:
titleid:
expires:
released:


Create License Command

customerid: Bayley
titleid: GlowS01E01
date: 6/27/2017

Command Handler

License Created Event

customerid: Bayley
titleid: GlowS01E01
date: 6/27/2017



License Repository

License Aggregate

customerId:
titleId:
expires:
released:

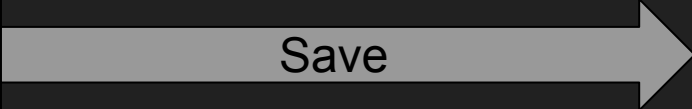

License Created Event

customerId: Bayley
titleId: GlowS01E01
date: 6/27/2017

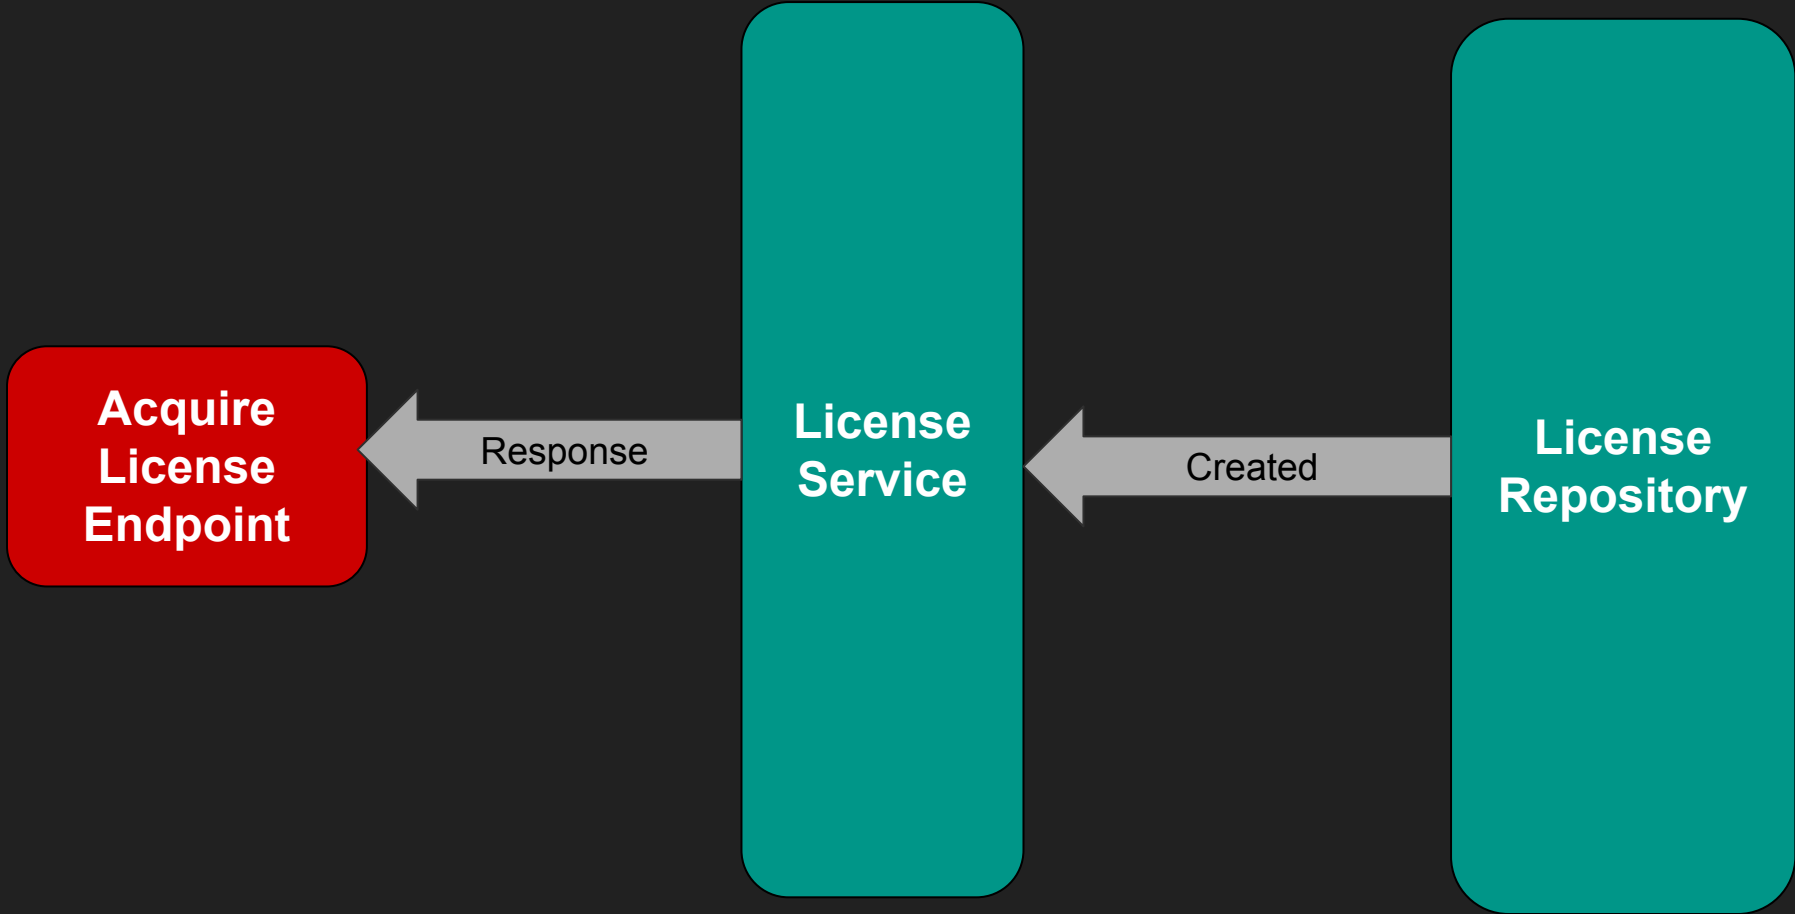
Event Handler

License Aggregate

customerId: Bayley
titleId: GlowS01E01
expires: 7/27/2017
released: false



Create Aggregate



isAllowed?

License Service

```
aggregates.filter(GlowS01E01)
.size()
< yearly limit
```

License Repository

Get all licenses

Response



Table

Partition Keys

Clustering Columns

Columns

Event Table

rowId
text

Aggregate Id
text

Event Time
timestamp

Event Data
blob

Event Mapper
String



- A fast and efficient object graph serialization framework for Java
 - <https://github.com/EsotericSoftware/kryo>
- Able to preconfigure the library to your data models for extra compaction
- Custom serializers

| rowId text | AggregateId text | Event Time timestamp | Event Data blob |
|------------------------------|------------------------|-------------------------|--------------------|
| Matt:LicenseAggregate | House of Cards | 0 | LicenseCreated |
| | Buddy Thunderstruck | 1 | LicenseCreated |
| Jeff:LicenseAggregate | The Hardy Boys | 2 | LicenseCreated |

| rowId text | AggregateId text | Event Time timestamp | Event Data blob |
|-----------------------|------------------------|-------------------------|--------------------|
| Matt:LicenseAggregate | House of Cards | 0 | LicenseCreated |
| | House of Cards | 3 | LicenseReleased |
| | Buddy Thunderstruck | 1 | LicenseCreated |
| Jeff:LicenseAggregate | The Hardy Boys | 2 | LicenseCreated |

| rowId text | AggregateId text | Event Time timestamp | Event Data blob |
|-----------------------|------------------------|-------------------------|--------------------|
| Matt:LicenseAggregate | House of Cards | 0 | LicenseCreated |
| | House of Cards | 3 | LicenseReleased |
| Jeff:LicenseAggregate | Buddy Thunderstruck | 1 | LicenseCreated |
| | The Hardy Boys | 2 | LicenseCreated |

rowId
text

AggregateId
text

Event Time
timestamp

Event Data
blob

Matt:LicenseAggregate

House of Cards

0

LicenseCreated

House of Cards

3

LicenseReleased

Buddy
Thunderstruck

1

LicenseCreated

Jeff:LicenseAggregate

The Hardy Boys

2

LicenseCreated

Snapshotting

| | | | |
|-----------------------|---------------------|---|-----------------|
| Matt:LicenseAggregate | House of Cards | 0 | LicenseCreated |
| | House of Cards | 3 | LicenseReleased |
| | Buddy Thunderstruck | 1 | LicenseCreated |
| | Master of None | 0 | LicenseCreated |
| | Master of None | 3 | LicenseReleased |
| | Bojack | 1 | LicenseCreated |
| | Lucha Underground | 0 | LicenseCreated |
| | Lucha Underground | 3 | LicenseReleased |
| | Grace & Frankie | 1 | LicenseCreated |
| | House of Cards | 0 | LicenseCreated |

Snapshot Table

rowId
text

version
int

Snapshot Data
blob

rowId
text

Version
int

Event Data
blob

Matt:LicenseAggregate

1

binaryData

rowId
text

AggregateId
text

Event Time
timestamp

Event Data
blob

Matt:LicenseAggregate:0

House of Cards

0

LicenseCreated

Jessica Jones

1

LicenseCreated

Matt:LicenseAggregate:1

House of Cards

3

LicenseReleased

Jeff:LicenseAggregate:0

The Hardy Boys

2

LicenseCreated



Working With Event Sourcing

Flexibility in Practice

- Changes to the Data Model are trivial!
 - Device deactivation

Manage Download Devices

Your plan allows downloads on 4 devices. When you remove a device, all downloads on that device will no longer be available.

Apple iPhone 6s

Last download on June 1, 2017

► [Show downloads](#)

Remove device

[Back To Account](#)

New Device Deactivation Requirement

Deactivate
Device
Endpoint

Device
Service

Device
Repository

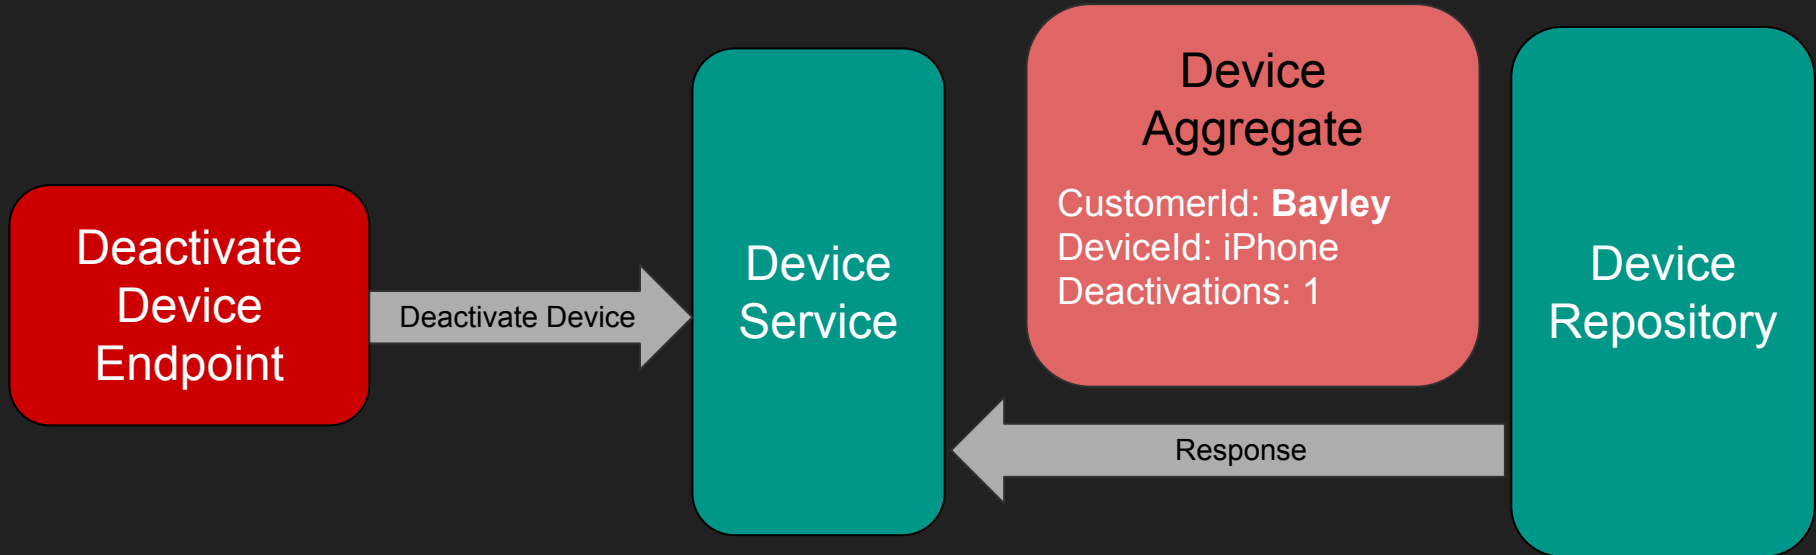
Device
Aggregate

Deactivate
Device
Command

Device
Deactivated
Event


New Device Deactivation Requirement...

Can Deactivate?

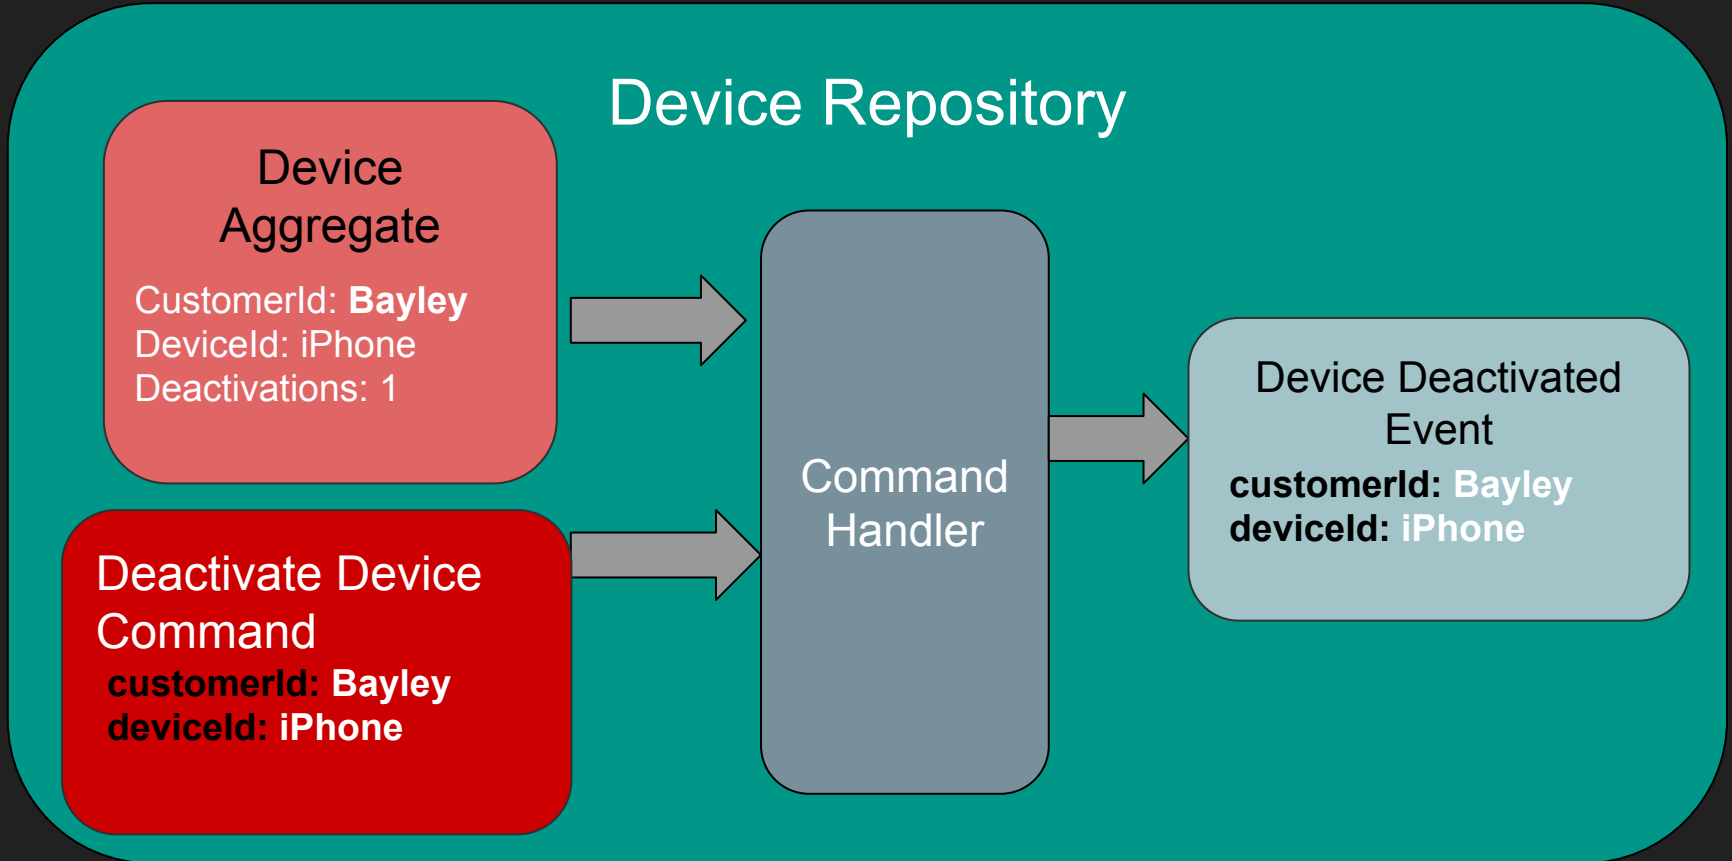


New Device Deactivation Requirement...

Can Deactivate ✓



New Device Deactivation Requirement...



Debugging

- Possibly the biggest win
- Current tooling is rudimentary



```
{  
  "Event" : "LicenseAcquiredEvent",  
  "Event Time" : "2017-06-25 05:23:00 PM",  
  "aggregateld" : "aggregateld1",  
  "committed" : true,  
  "eventTimeStamp" : 1498411380903,  
  "eventOrderNumber" : 0  
}
```

Debugging...

Acquire License
Event

customerId: Bayley
deviceId: iPhone
date: 01/05/2017

Release License
Event

customerId: Bayley
deviceId: iPhone
date: 01/05/2017

Renew License
Event

customerId: Bayley
deviceId: iPhone
date: 01/06/2017

License
Aggregate

customerId: Bayley
expires: 01/07/2017
released: True

Reliability: Fallbacks




How Did it Scale?



Scalability with Cassandra SSD (I2) Nodes

| Model | vCPU | Memory (GiB) | Storage | Read IOPS | Write IOPS |
|------------|------|--------------|-------------|-----------|------------|
| i2.xlarge | 4 | 30.5 | 1 x 800 SSD | 35,000 | 35,000 |
| i2.2xlarge | 8 | 61 | 2 x 800 SSD | 75,000 | 75,000 |
| i2.4xlarge | 16 | 122 | 4 x 800 SSD | 175,000 | 155,000 |
| i2.8xlarge | 32 | 244 | 8 x 800 SSD | 365,000 | 315,000 |

Scalability with Cassandra SSD Nodes...



- Could handle much higher loads
- Storage use ramping up very quickly

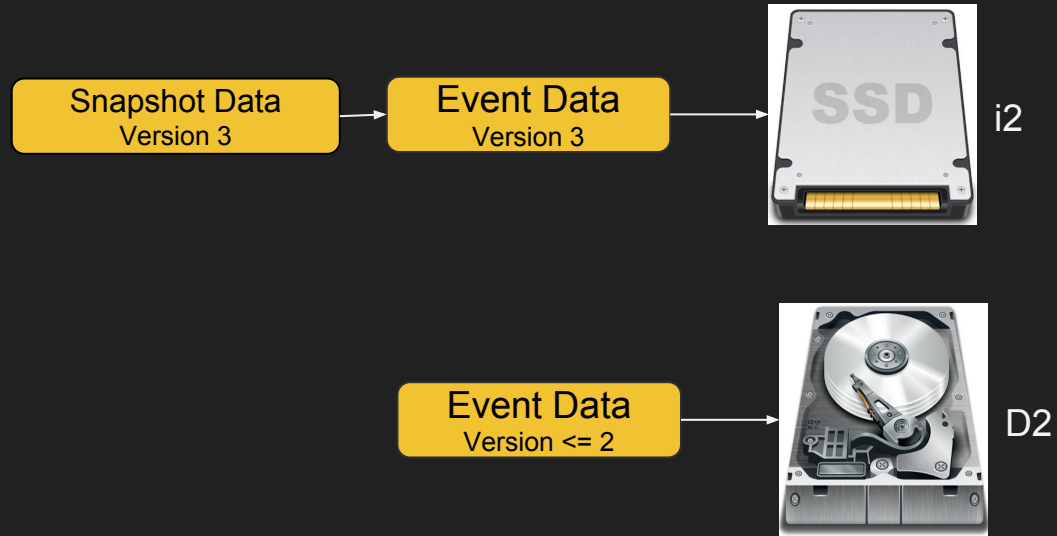
Storage Optimizations

- TTLs
- HDD (D2) clusters:
 - More storage
 - Higher latency (up to 1 second)

| Model | vCPU | Memory (GiB) | Storage | Read Throughput (2 MiB Block size) |
|------------|------|--------------|-----------|------------------------------------|
| d2.xlarge | 4 | 30.5 | 3 x 2 TB | 438 MB/s |
| d2.2xlarge | 8 | 61 | 6 x 2 TB | 875 MB/s |
| d2.4xlarge | 16 | 122 | 12 x 2 TB | 1,750 MB/s |
| d2.8xlarge | 36 | 244 | 24 x 2 TB | 3,500 MB/s |

Storage Optimizations...

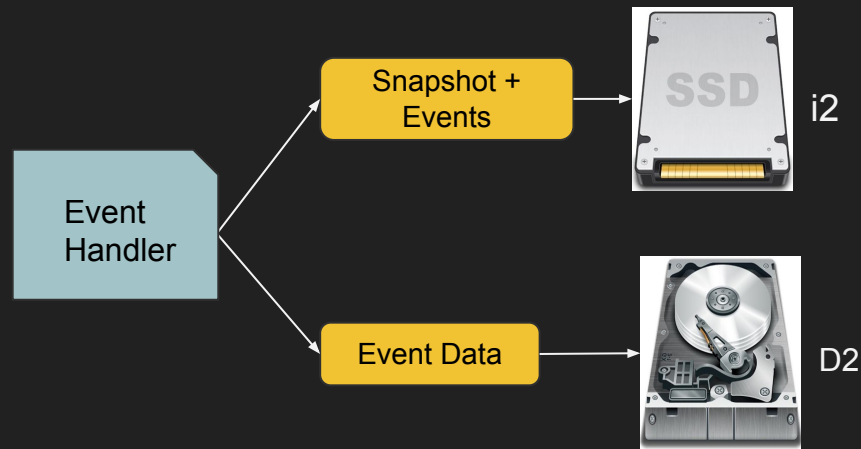
- Partitioned approach
 - Snapshot and subsequent events saved to i2 (SSD) cluster
 - Archive events to D2 (HDD) cluster



Storage Optimizations...

Partitioned approach using CQRS (Command Query Responsibility Segregation)

- Write segregation
 - Event handlers that determine partitioned usage
- Query segregation
 - Query SSD or HDD (or both)
 - Uninterrupted event stream



Key Event Sourcing Take-Aways

- Flexible: Adapting to change can be simple

Key Event Sourcing Take-Aways

- Flexible: Adapting to change can be simple
- Debugging: Debugging data / state transitions greatly simplified

Key Event Sourcing Take-Aways

- Flexible: Adapting to change can be simple
- Debugging: Debugging data / state transitions greatly simplified
- Reliable: Fallbacks provide service reliability

Key Event Sourcing Take-Aways

- Flexible: Adapting to change can be simple
- Debugging: Debugging data / state transitions greatly simplified
- Reliable: Fallbacks provide service reliability
- Scalable: Service scaled well, but good architectural solutions for data storage solutions should be considered.

Questions?



Phillipa Avery

Senior Software Engineer

Project Technical Lead and Engineer
Downloads License Accounting

pippa@netflix.com /  @PhillipaAvery



Robert Reta

Senior Software Engineer

Event Sourcing System Architect
Downloads License Accounting

rreta@netflix.com /  @rreta04