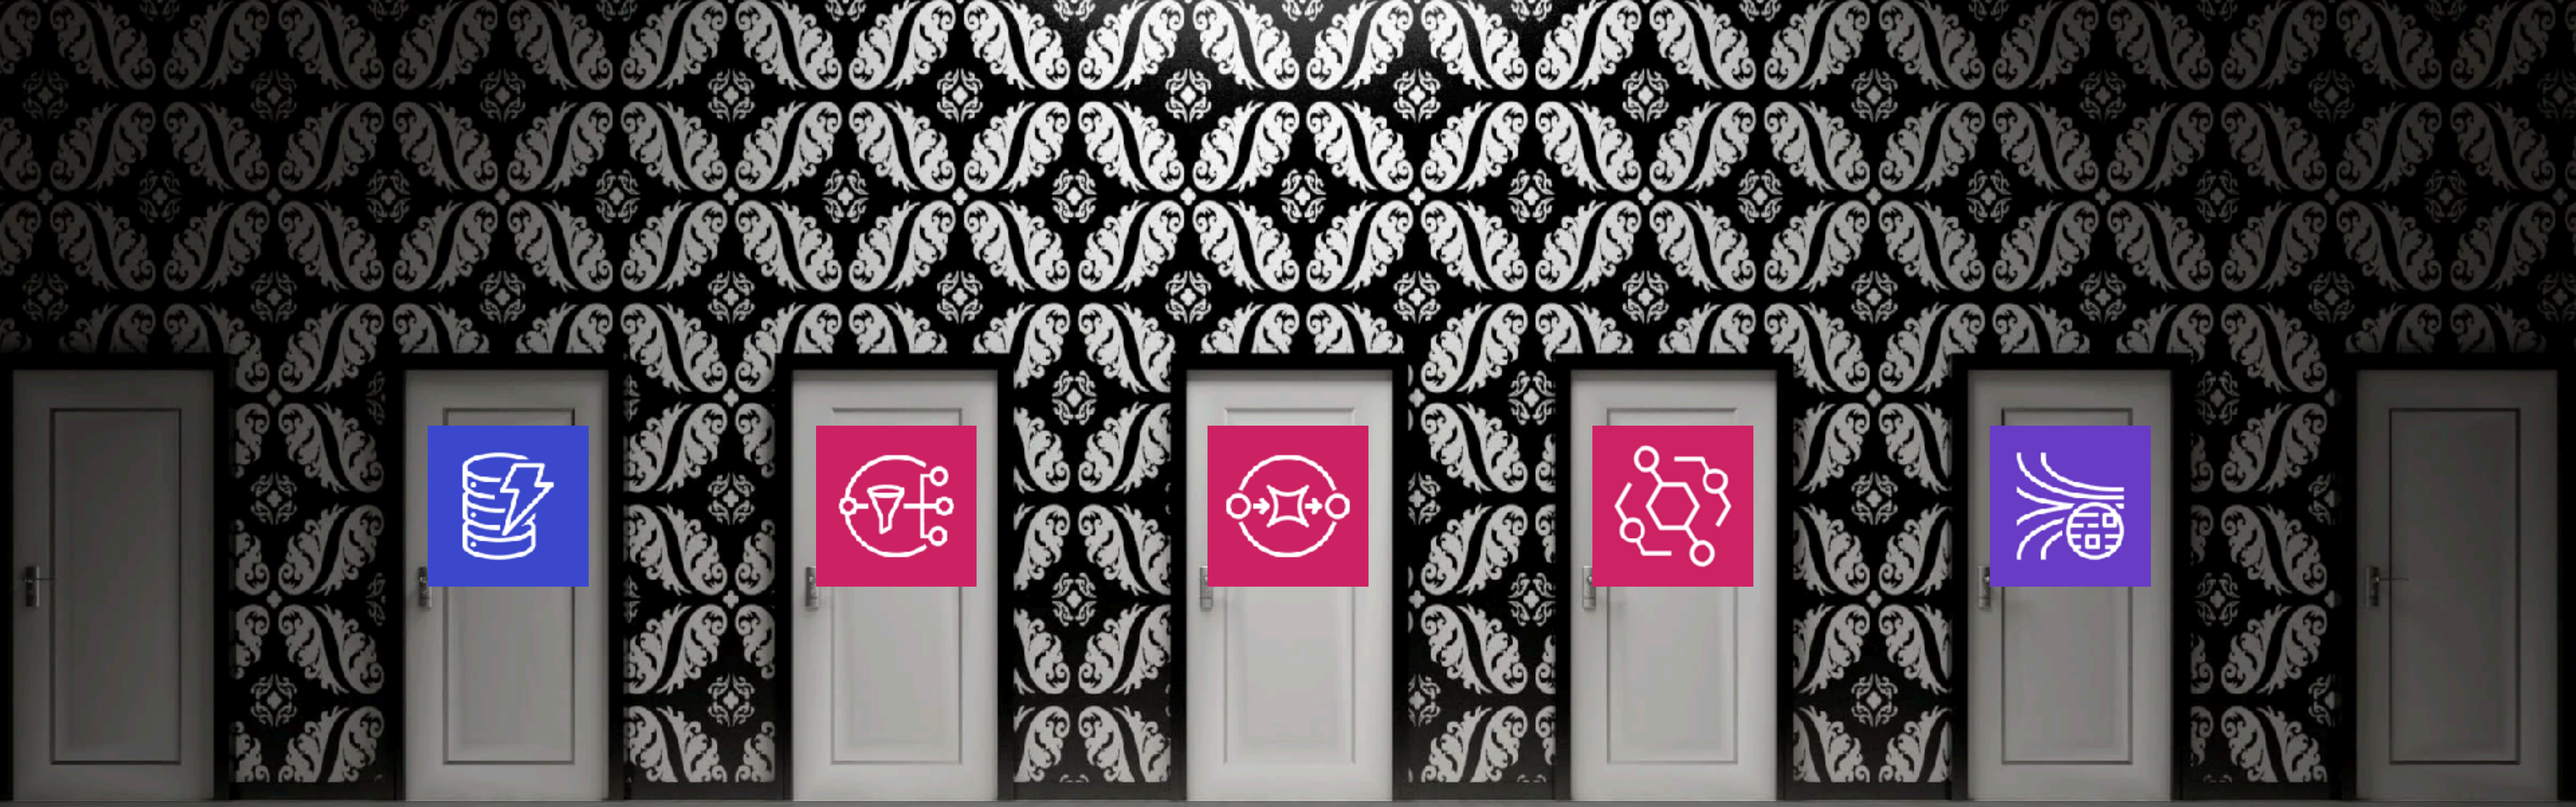

How to choose the right messaging service for your workload

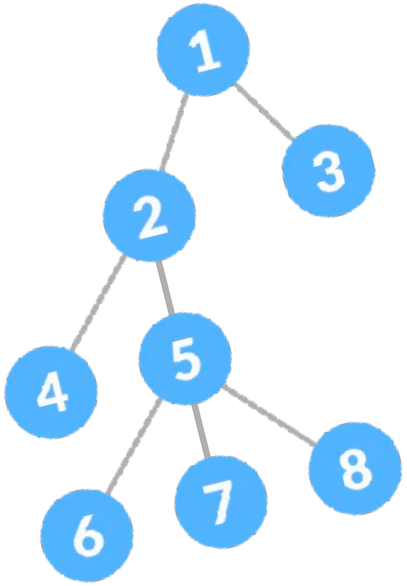

O(*n*)

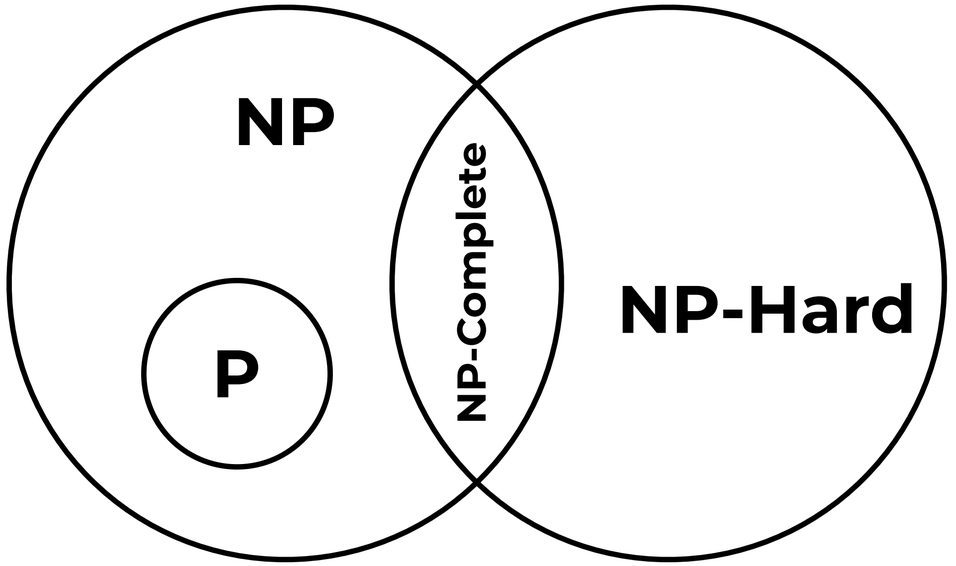

Yan Cui

@theburningmonk

http://theburningmonk.com

AWS user since 2010

Yan Cui

@theburningmonk

http://theburningmonk.com

Developer Advocate @ lumigo

aws serverless HERO

Yan Cui

@theburningmonk

http://theburningmonk.com

aws serverless HERO

Independent Consultant

JustGiving · TOYOTA · stedi · homegate.ch · SPARKOL

HARGREAVES LANSDOWN · Klarna. · CYBERARK · Zava · aufeminin

TACO BELL · YNAP · Intelliflo · Driver & Vehicle Licensing Agency · SB Simply Business

training

advise

delivery

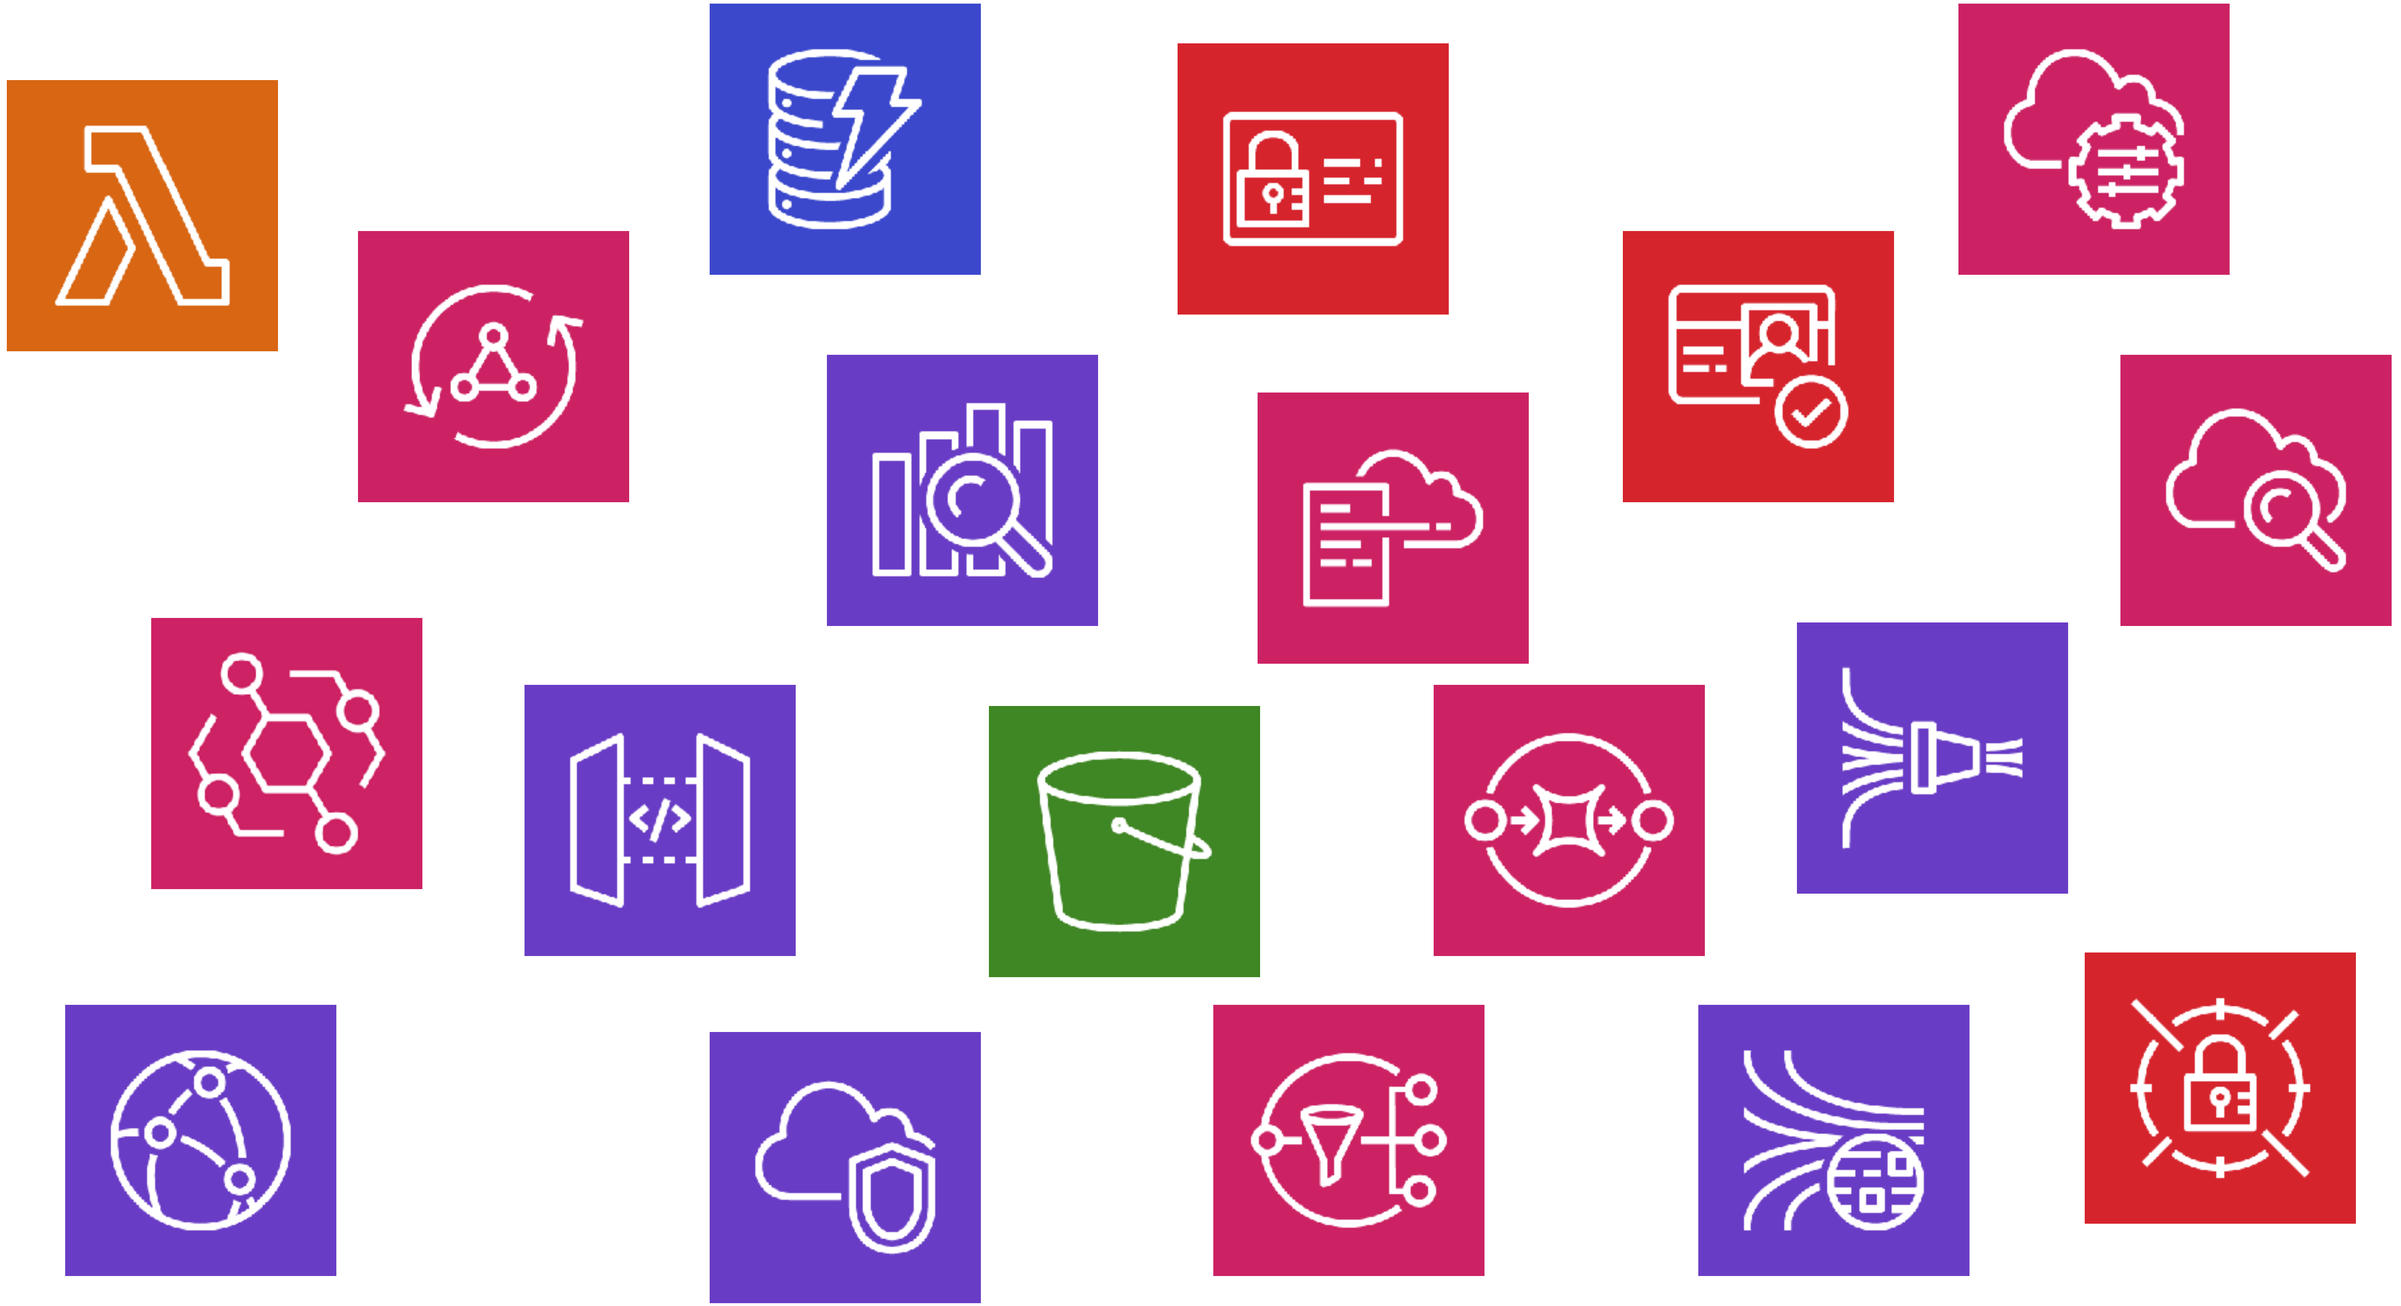

AWS services are the new data structures.

Understanding their **trade-offs** is the new big O notation.

Event-Driven Architecture

Event-Driven Architecture

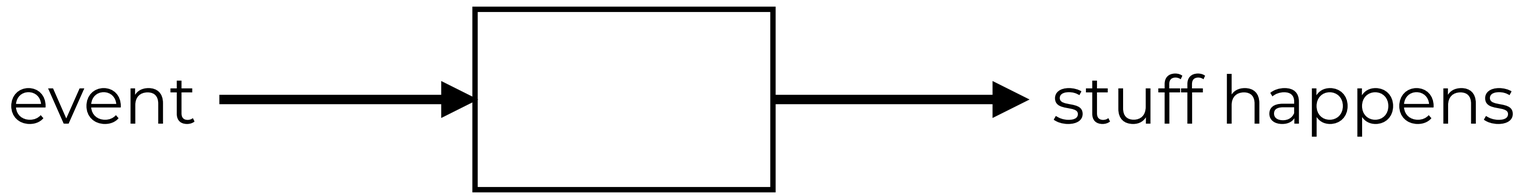
event
stuff happens

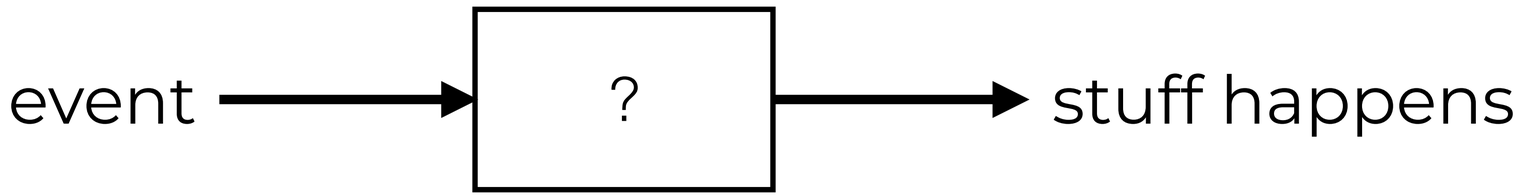
event
?
stuff happens

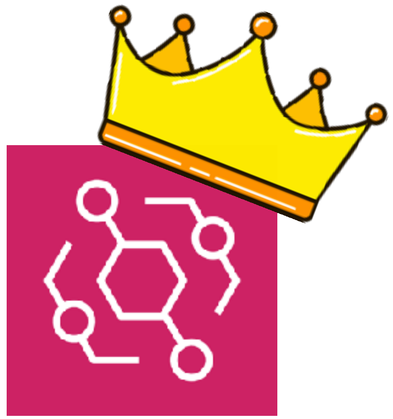

EventBridge

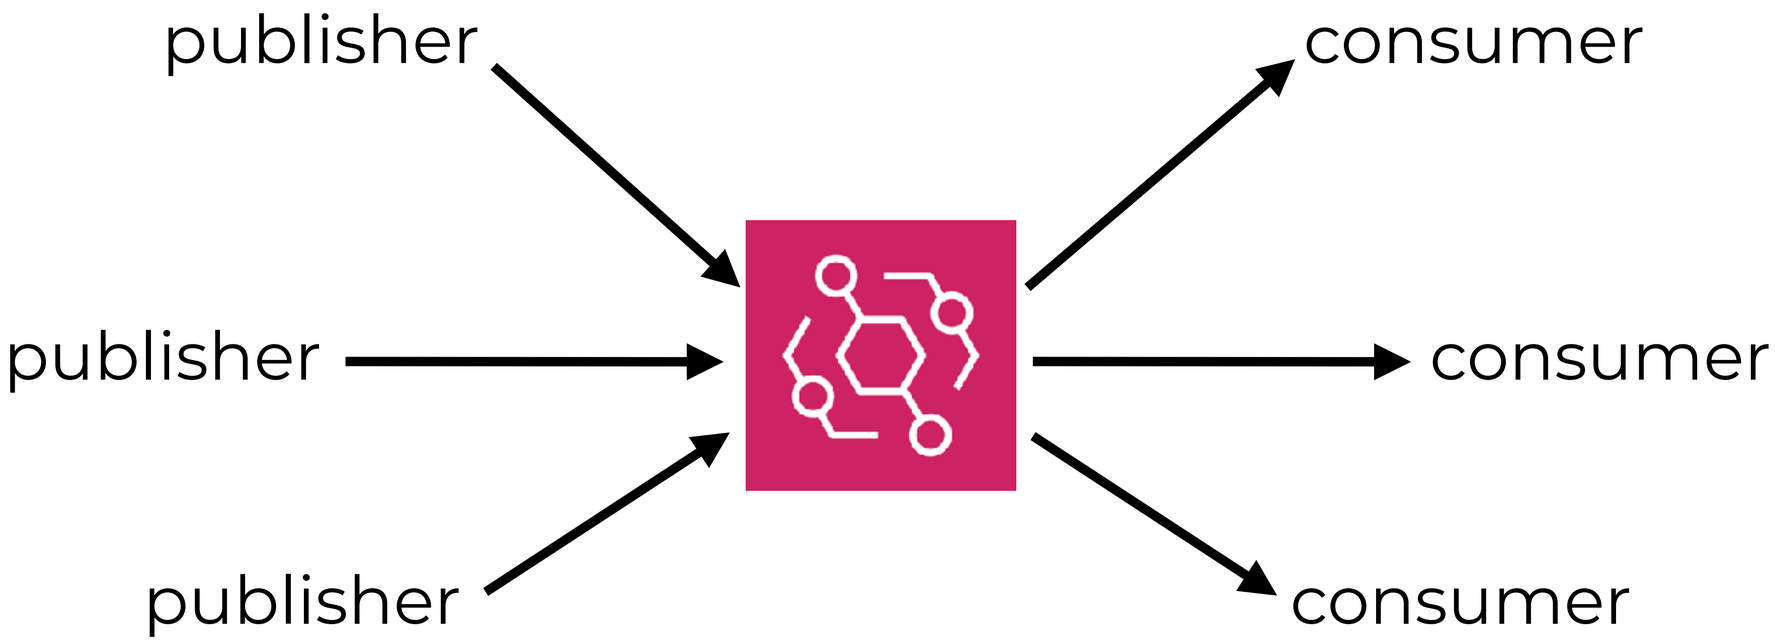
publisher
publisher
publisher
consumer
consumer
consumer

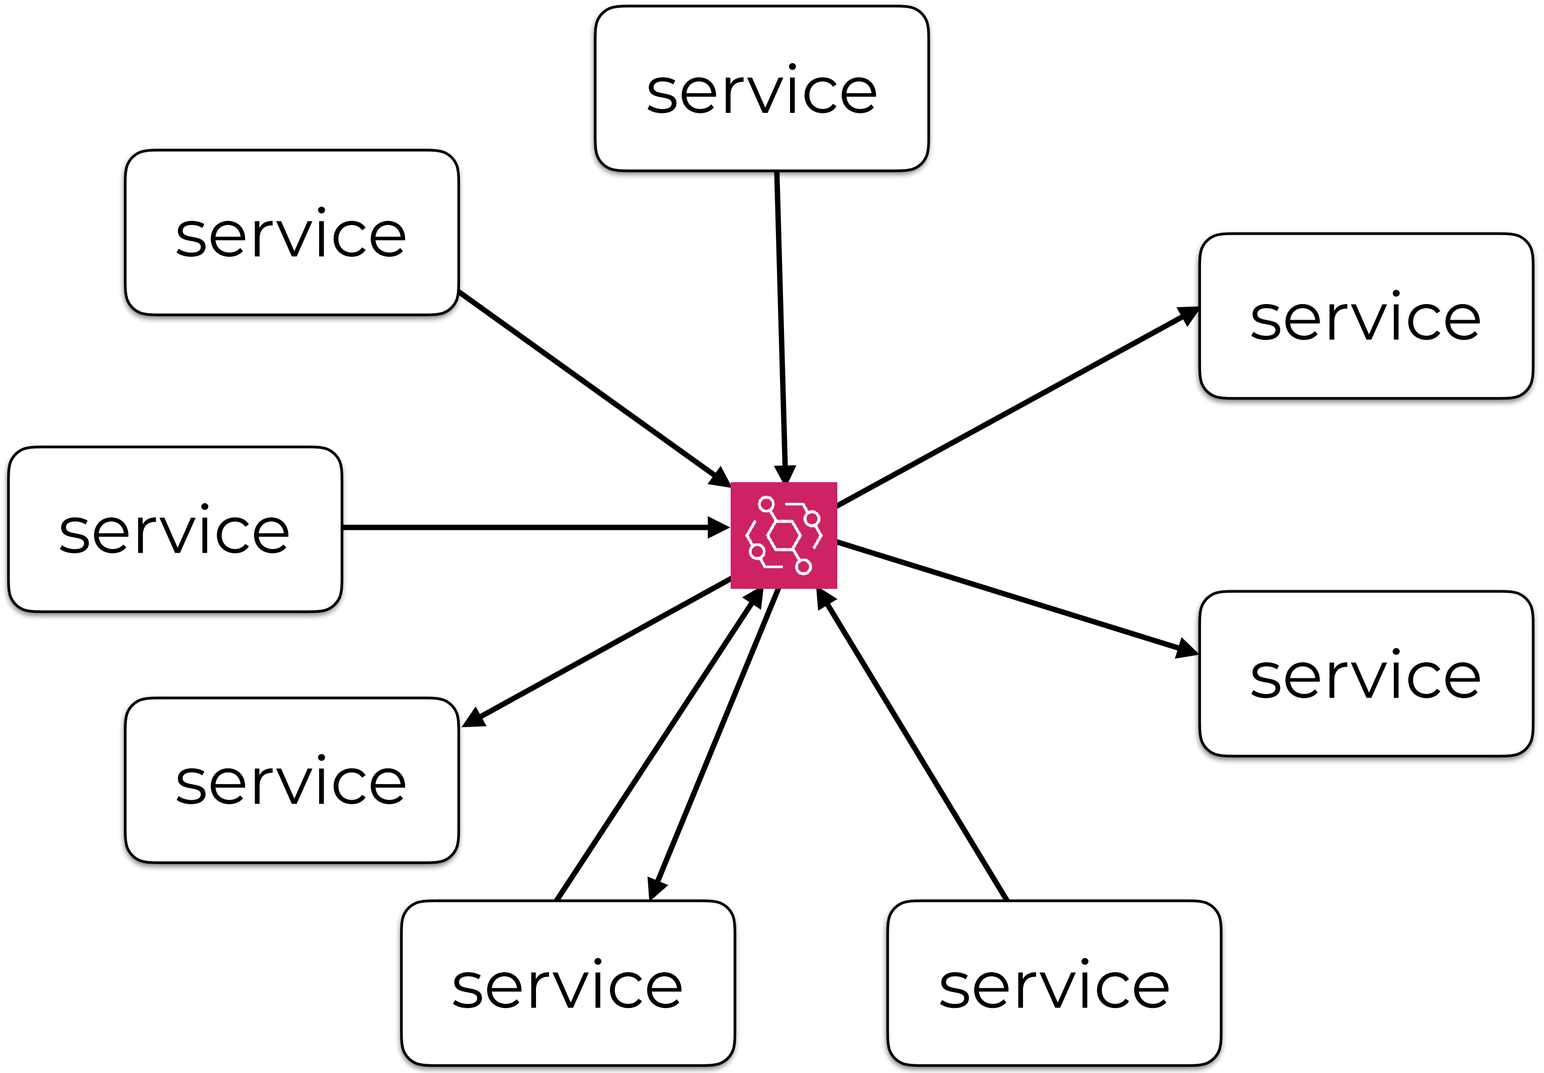
service
service
service
service
service
service
service
service
service

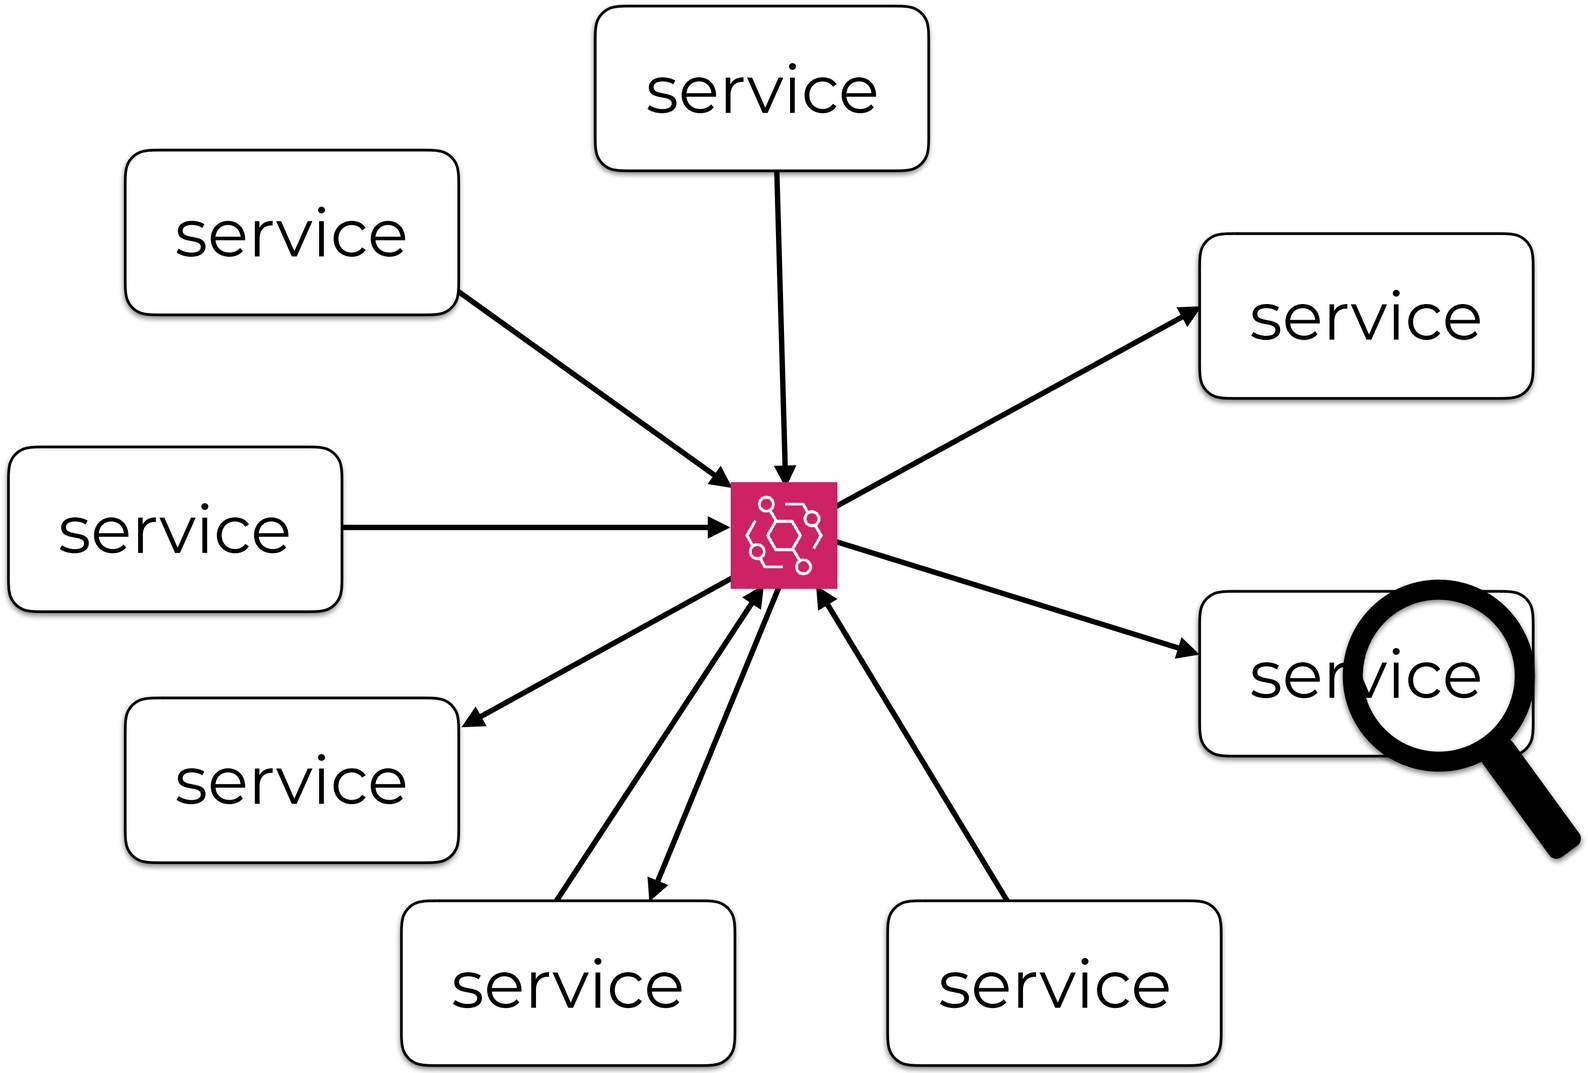
service
service
service
service
service
service
service
service
service

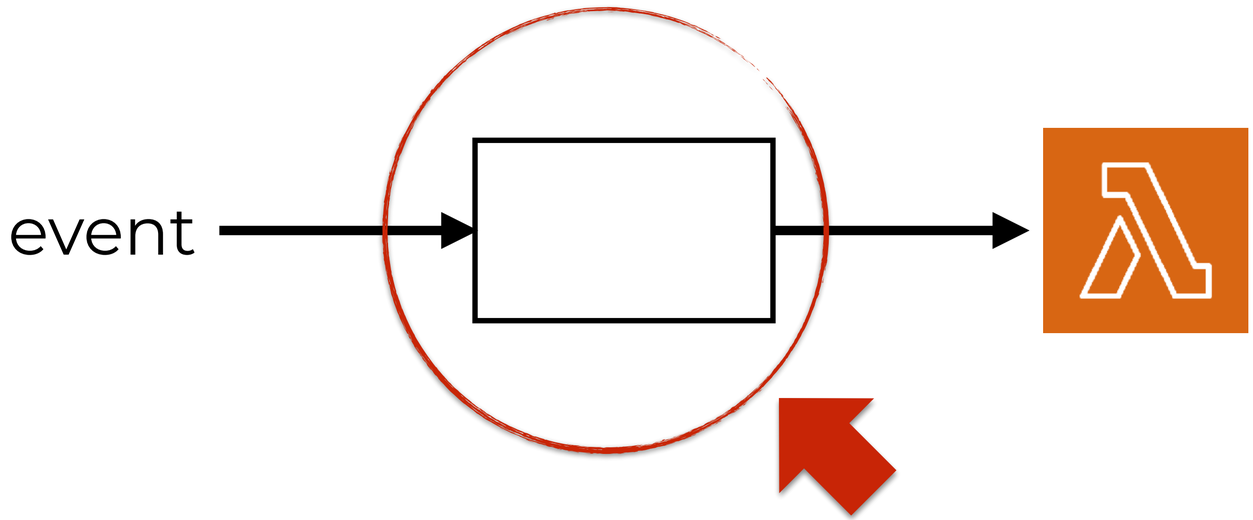
event

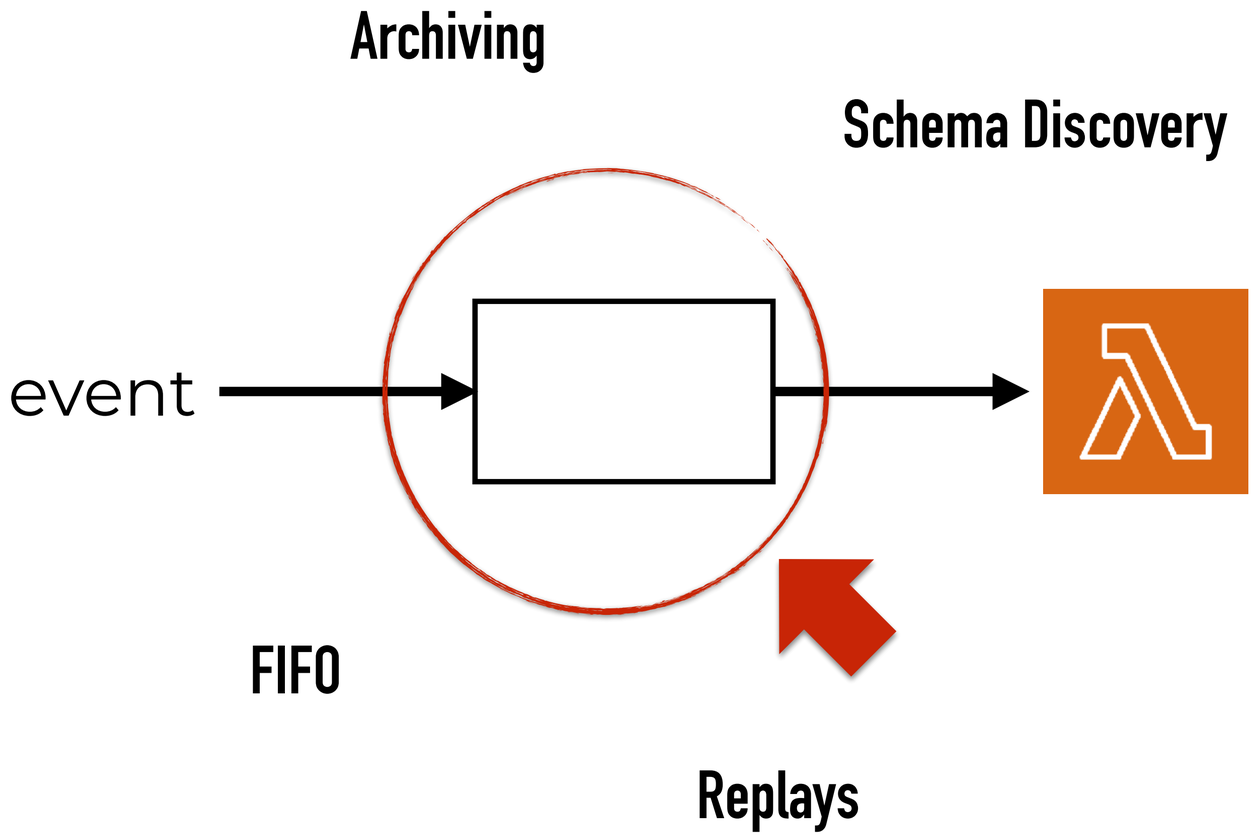
Archiving
Schema Discovery
event
FIFO
Replays

Error Handling

Performance

Archiving

Cost

Schema Discovery

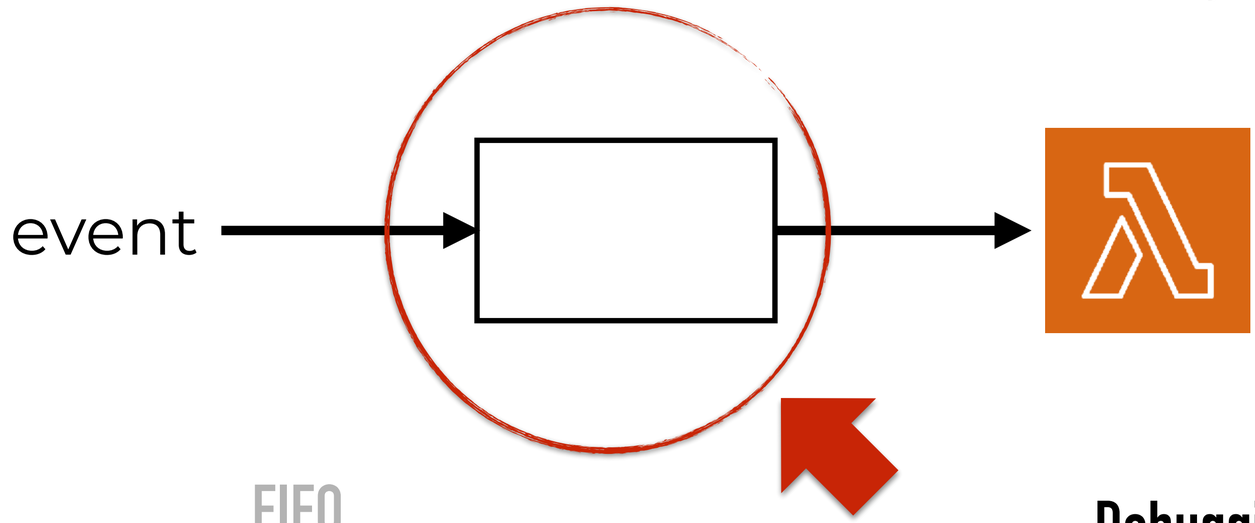

FIFO

Debugging

Replays

Scalability

SNS

SQS

EventBridge

Kinesis

DynamoDB Stream

IOT Core

...

picking the right AWS service is a valuable skill

SNS

SQS

EventBridge

Kinesis

DynamoDB Stream

IOT Core

...

Scaling

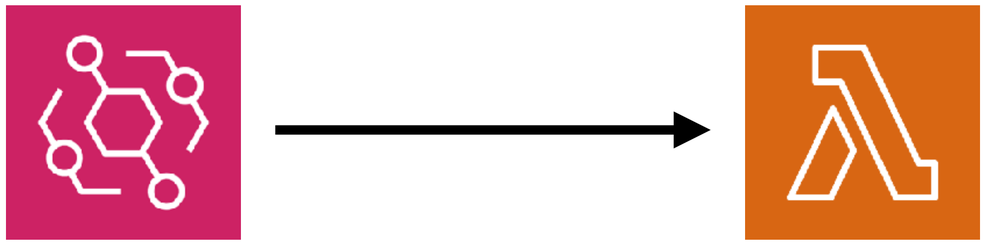

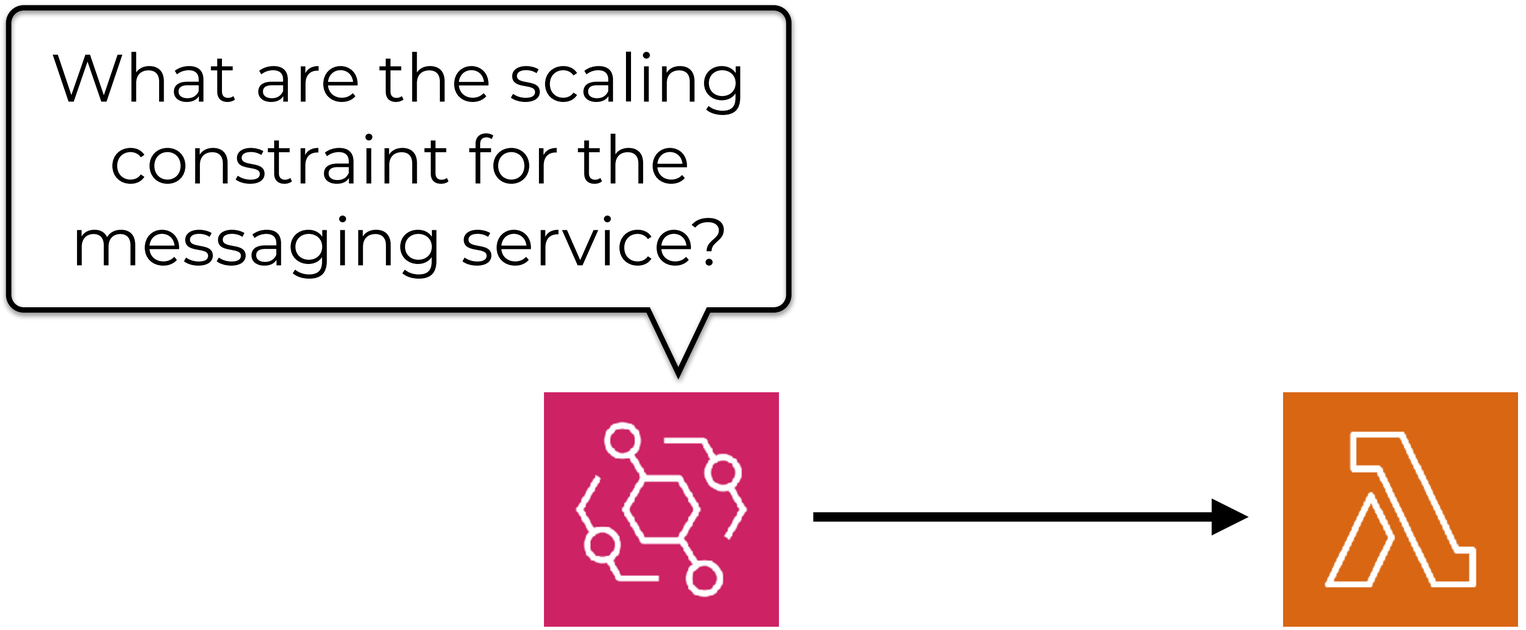
What are the scaling constraint for the messaging service?

Google “{service name} quotas”

Amazon EventBridge quotas

PDF RSS

There are quotas for most aspects of EventBridge.

Topics

EventBridge quotas

EventBridge has the following quotas.

Resource	Default limit
	The Service Quotas console provides information about EventBridge quotas. Along with viewing the default quotas, you can use the Service Quotas console to request quota increases for adjustable quotas.
ListRuleNamesByTarget	Maximum 100 results per page for requests.
ListRules	Maximum 100 results per page for requests.
ListTargetsByRule	Maximum 100 results per page for requests.
PutEvents entry size	Maximum 256KB
PutTargets	10 entries per request. Up to 5 targets per rule.
RemoveTargets	10 entries per request.
Rules	300 per event bus. The Service Quotas console provides information about EventBridge quotas. Along with viewing the default quotas, you can use the Service Quotas console to request quota increases for adjustable quotas. Before increasing quotas, examine your rules. If you have multiple rules that match very specific events, consider broadening their scope by using fewer identifiers in your Amazon EventBridge events. In addition, consider adding

PutEvents quotas by Region

The Service Quotas console provides information about EventBridge quotas. Along with viewing the default quotas, you can use the Service Quotas console to request quota increases for adjustable quotas. For quotas above 100,000 TPS, our service team will host a call to best support you

> **Note**
>
> Events sent to a different Region using PutEvents count towards the PutEvents quota of the destination Region for the account that owns the role used to send the events.

Regions	Transactions per second
• US East (N. Virginia) • US West (Oregon) • Europe (Ireland)	PutEvents has a soft limit of 10,000 requests per second by default in these Regions.
• US East (Ohio) • Europe (Frankfurt)	PutEvents has a soft limit of 2,400 requests per second by default in these Regions.
• US West (N. California) • Europe (London) • Asia Pacific (Sydney) • Asia Pacific (Tokyo)	PutEvents has a soft limit of 1,200 requests per second by default in these Regions.

Amazon EventBridge (CloudWatch Events)

Service quotas

Request quota increase

Find quotas

1

Quota name	Applied quota value	AWS default quota value	Adjustable
Api destinations	3,000	3,000	Yes
Connections	3,000	3,000	Yes
CreateEndpoint throttle limit in transactions per second	5 per second	5 per second	No
DeleteEndpoint throttle limit in transactions per second	5 per second	5 per second	No
Endpoints	100	100	Yes
Invocations throttle limit in transactions per second	18,750	18,750	Yes
Number of rules	300	300	Yes
PutEvents throttle limit in transactions per second	10,000	10,000	Yes
Rate of invocations per API destination	300	300	Yes
Targets per rule	5	5	No
Throttle limit in transactions per second	50	50	Yes
UpdateEndpoint throttle limit in transactions per second	5 per second	5 per second	No

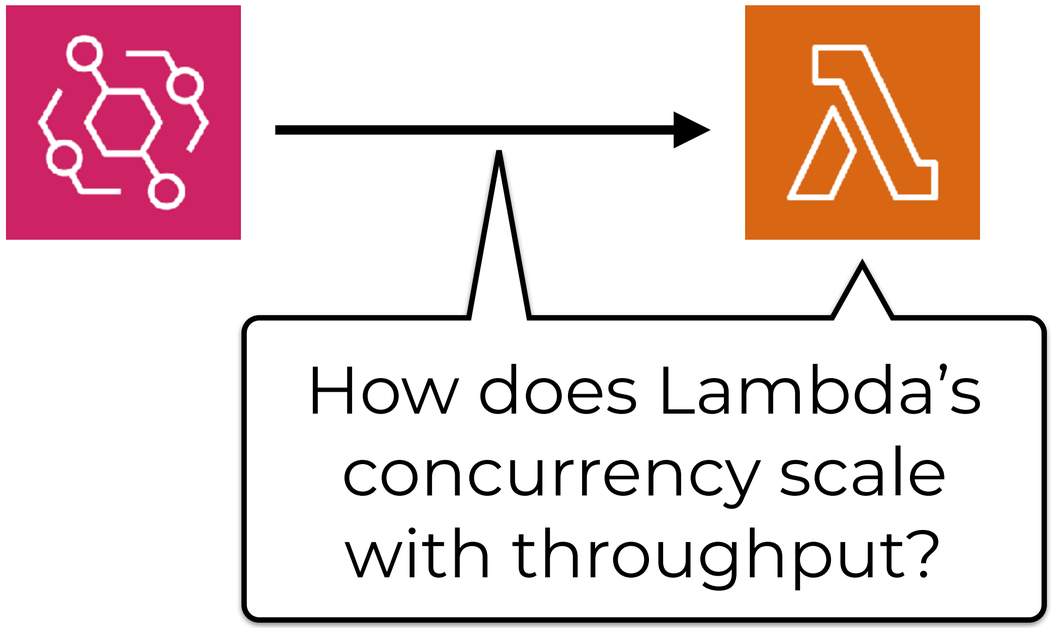
How does Lambda's concurrency scale with throughput?


Concurrency
Msgs/s

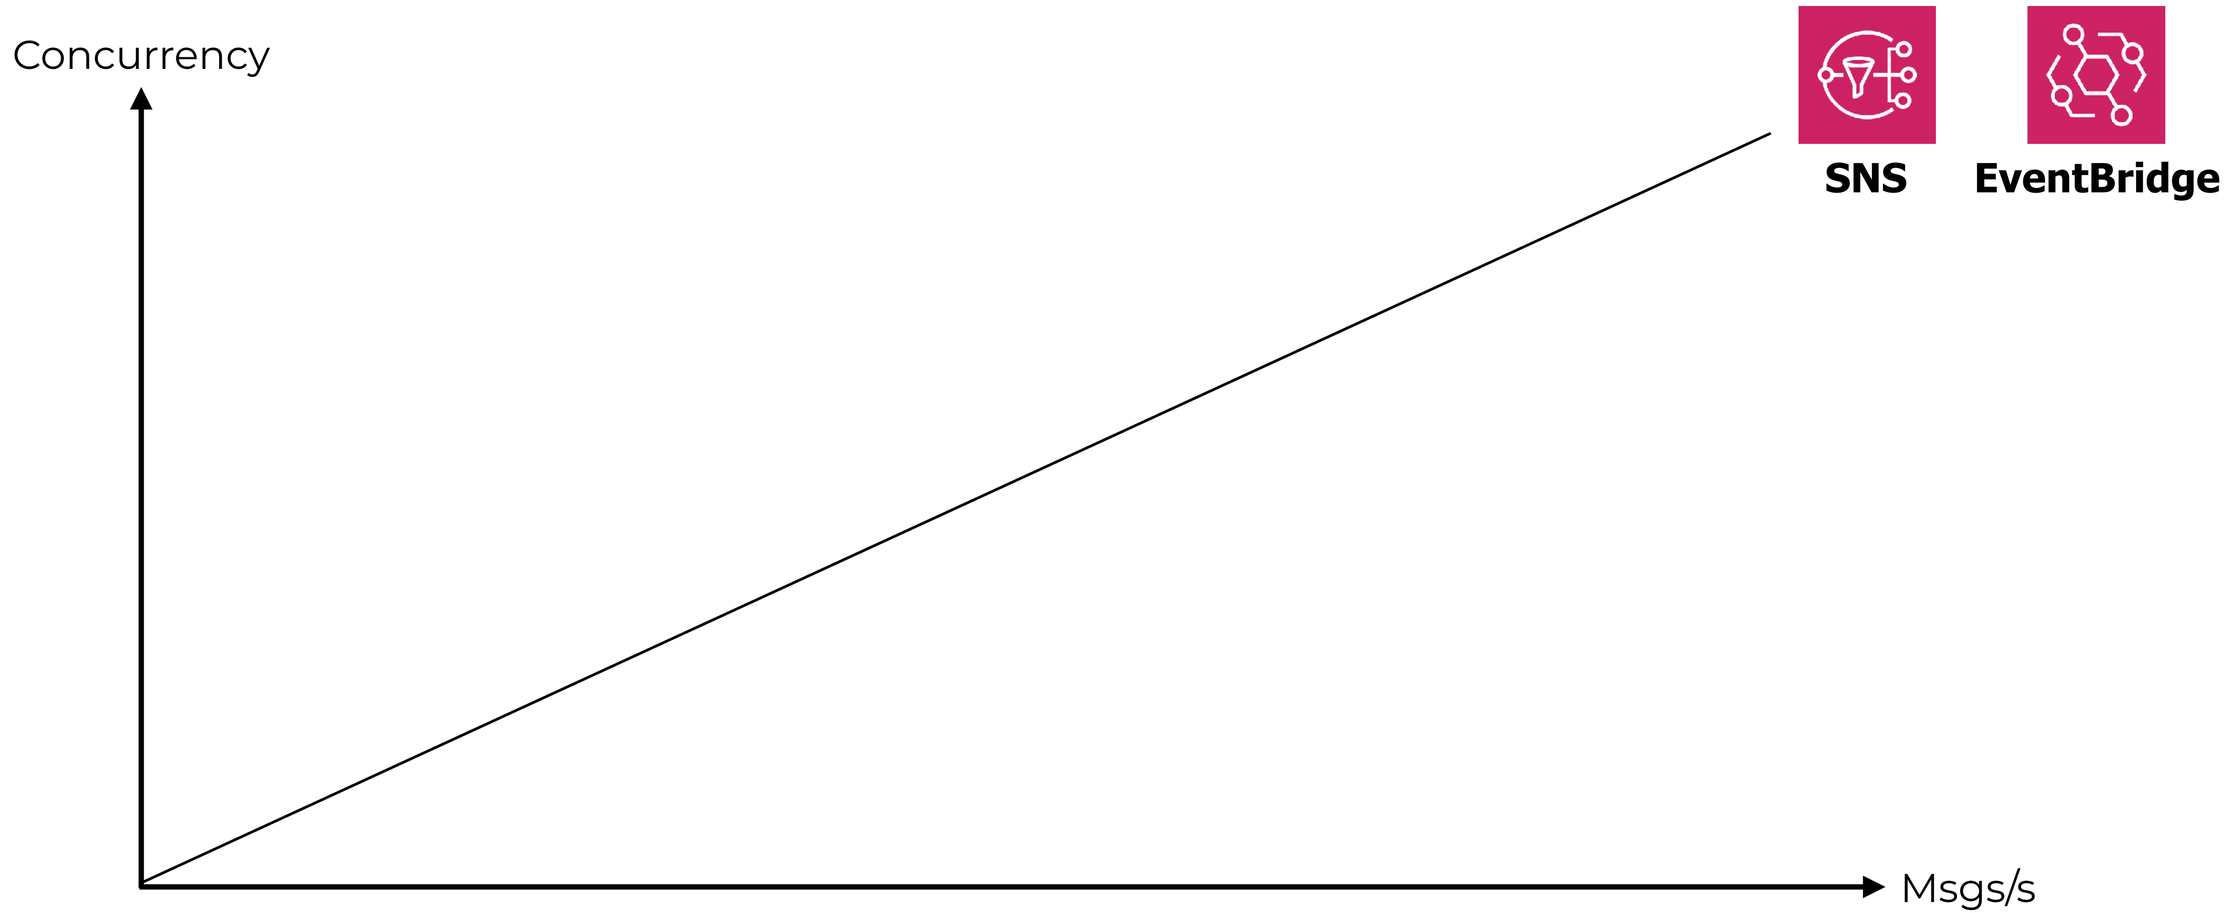
Concurrency
SNS
EventBridge
Msgs/s

Scaling and processing

For standard queues, Lambda uses long polling to poll a queue until it becomes active. When messages are available, Lambda reads up to five batches and sends them to your function. If messages are still available, Lambda increases the number of processes that are reading batches by up to 60 more instances per minute. The maximum number of batches that an event source mapping can process simultaneously is 1,000.

For FIFO queues, Lambda sends messages to your function in the order that it receives them. When you send a message to a FIFO queue, you specify a message group ID. Amazon SQS ensures that messages in the same group are delivered to Lambda in order. Lambda sorts the messages into groups and sends only one batch at a time for a group. If your function returns an error, the function attempts all retries on the affected messages before Lambda receives additional messages from the same group.

Your function can scale in concurrency to the number of active message groups. For more information, see SQS FIFO as an event source on the AWS Compute Blog.

https://docs.aws.amazon.com/lambda/latest/dg/with-sqs.html

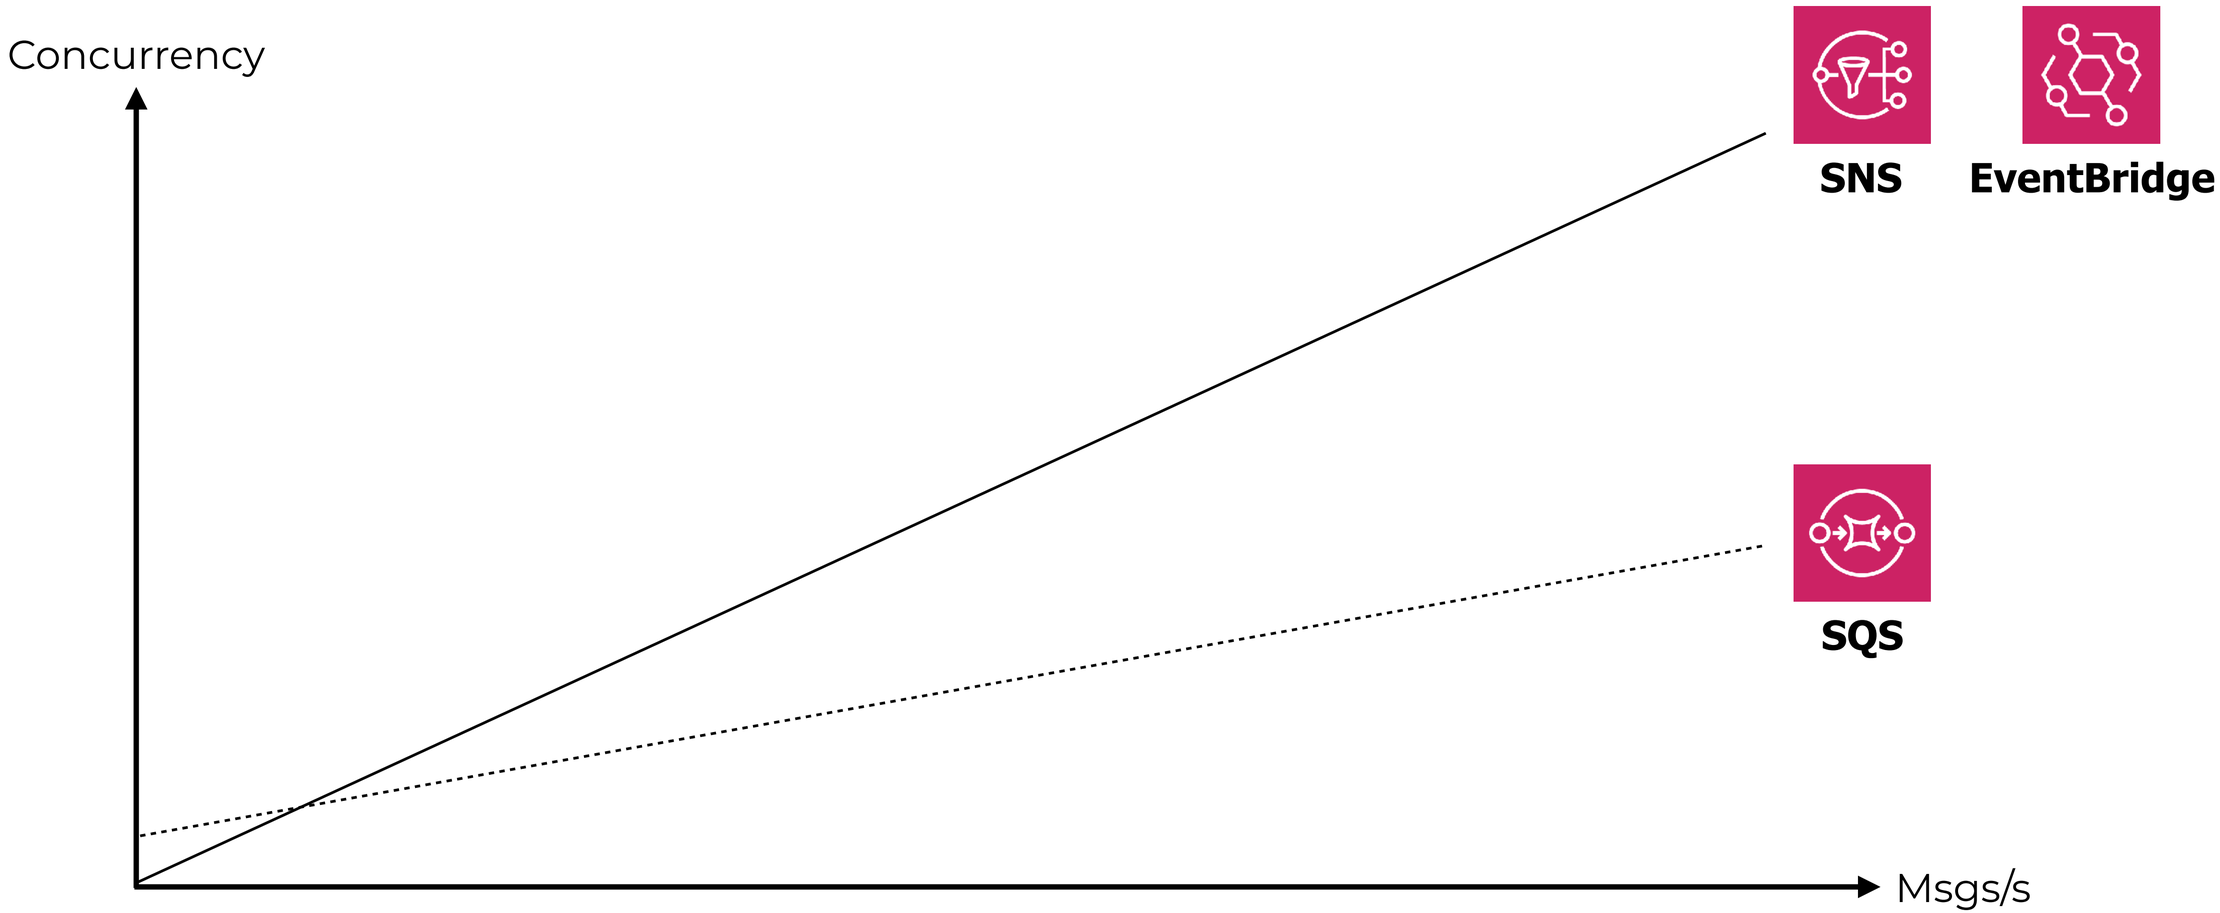
Concurrency
SNS
EventBridge
SQS
Msgs/s

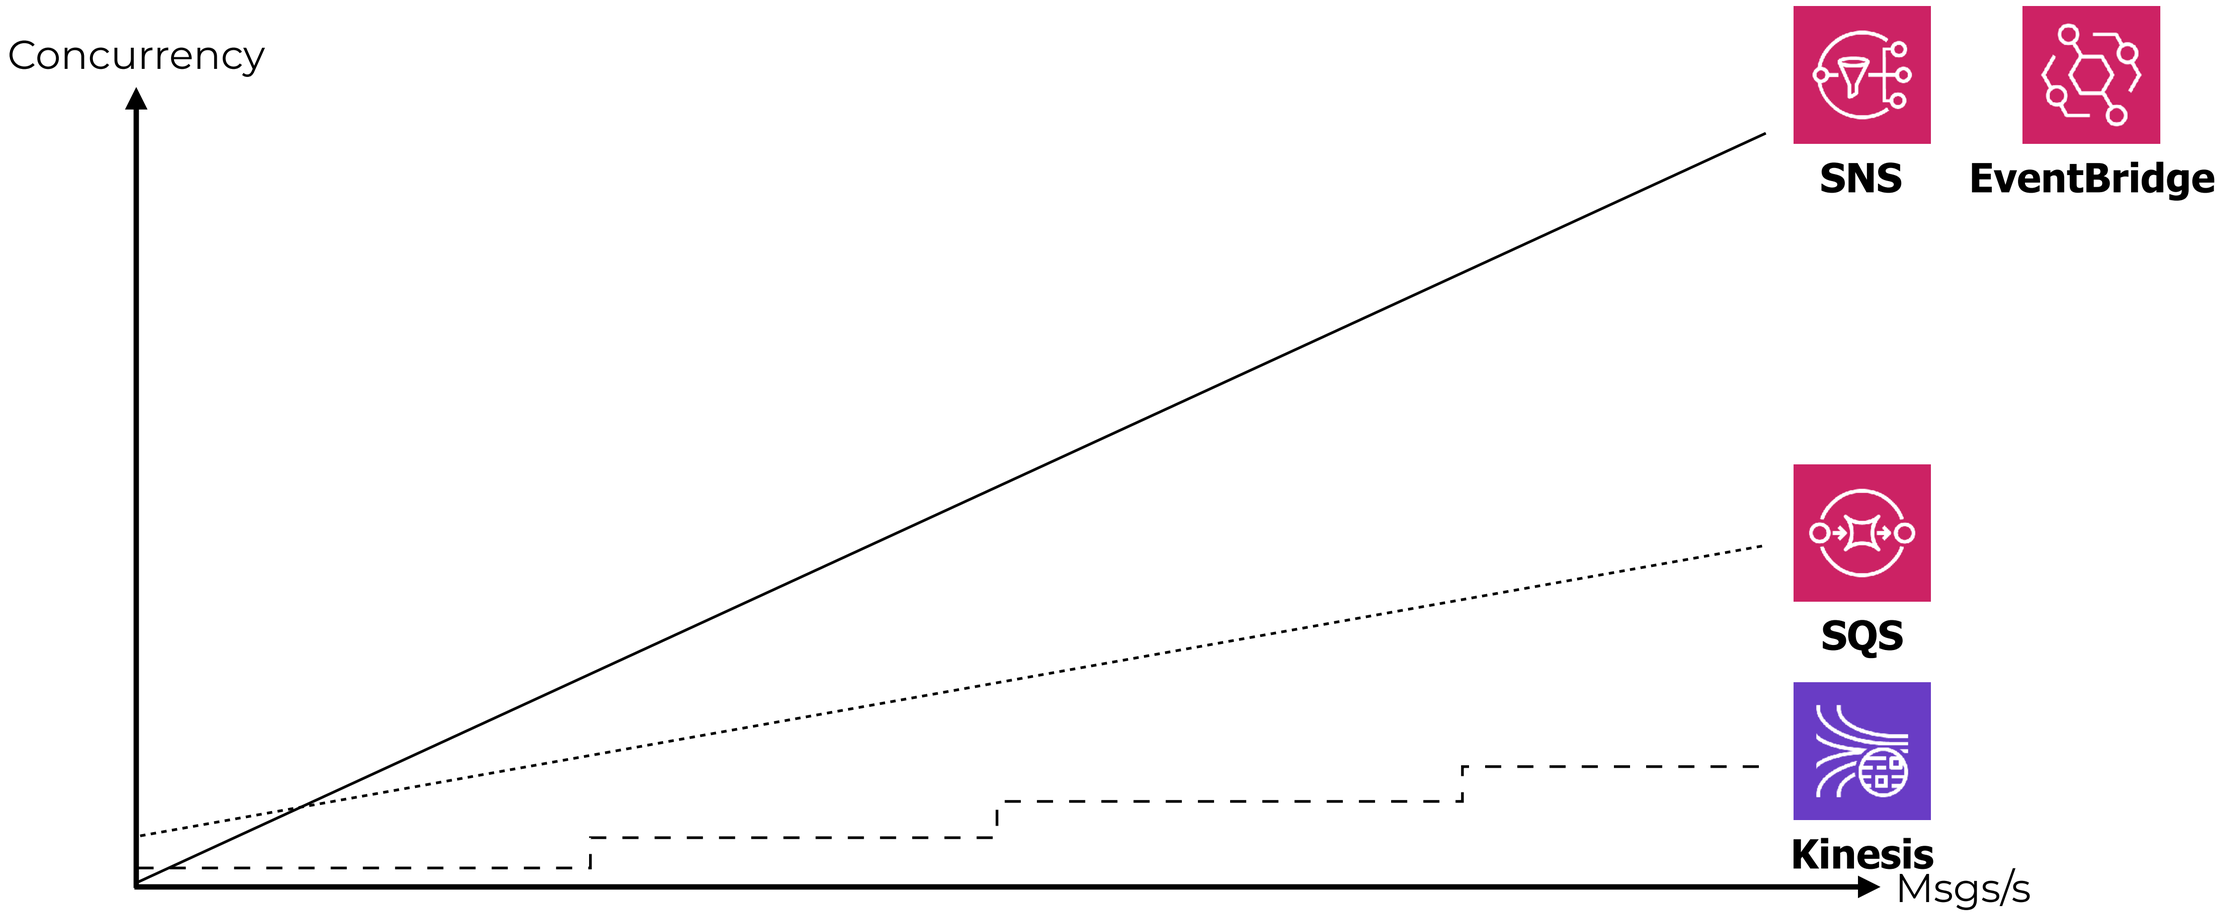
Concurrency
SNS
EventBridge
SQS
Kinesis
Msgs/s

If your function returns an error, Lambda retries the batch until processing succeeds or the data expires. To avoid stalled shards, you can configure the event source mapping to retry with a smaller batch size, limit the number of retries, or discard records that are too old. To retain discarded events, you can configure the event source mapping to send details about failed batches to an SQS queue or SNS topic.

You can also increase concurrency by processing multiple batches from each shard in parallel. Lambda can process up to 10 batches in each shard simultaneously. If you increase the number of concurrent batches per shard, Lambda still ensures in-order processing at the partition-key level.

Sections

- Configuring Your Data Stream and Function

https://amzn.to/2RudmGV

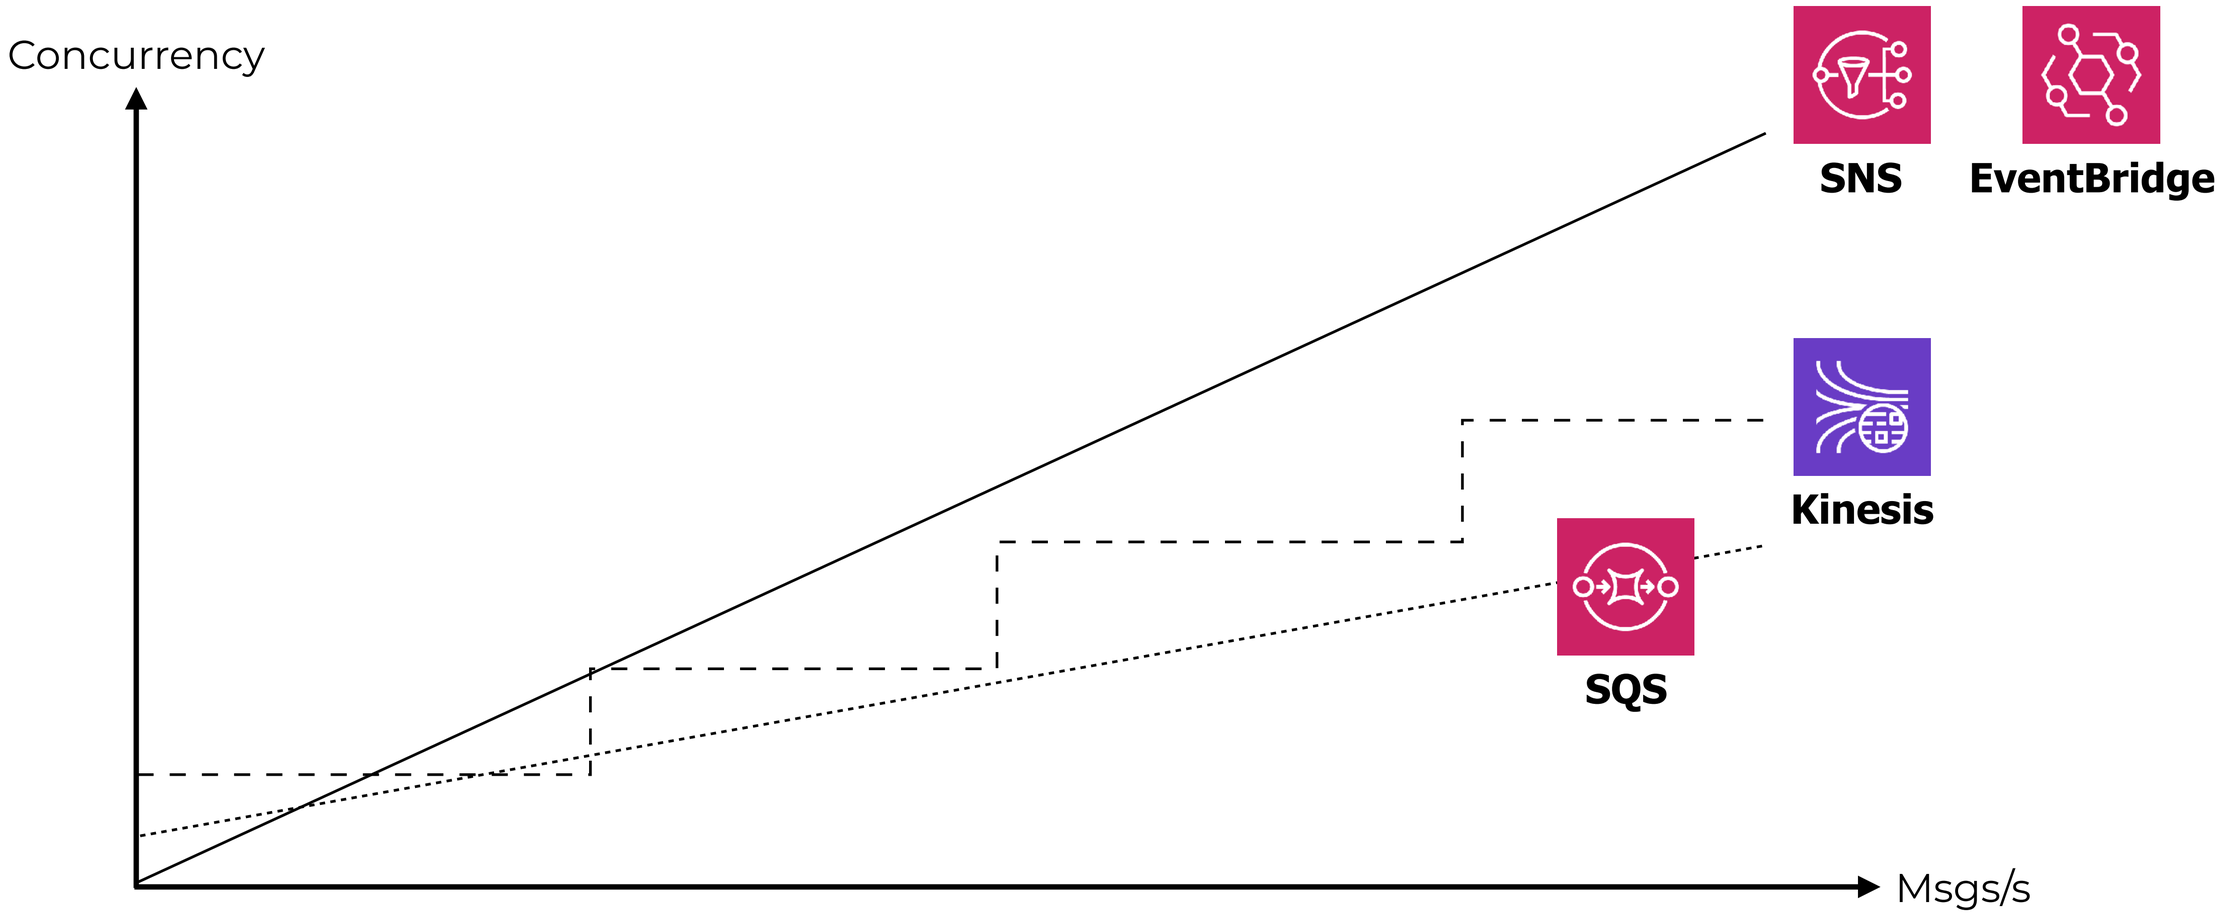
Concurrency
SNS
EventBridge
Kinesis
SQS
Msgs/s

more concurrency is not always better...

if you want...

maximum throughput

SNS

EventBridge

precise control over throughput

Kinesis Provisioned

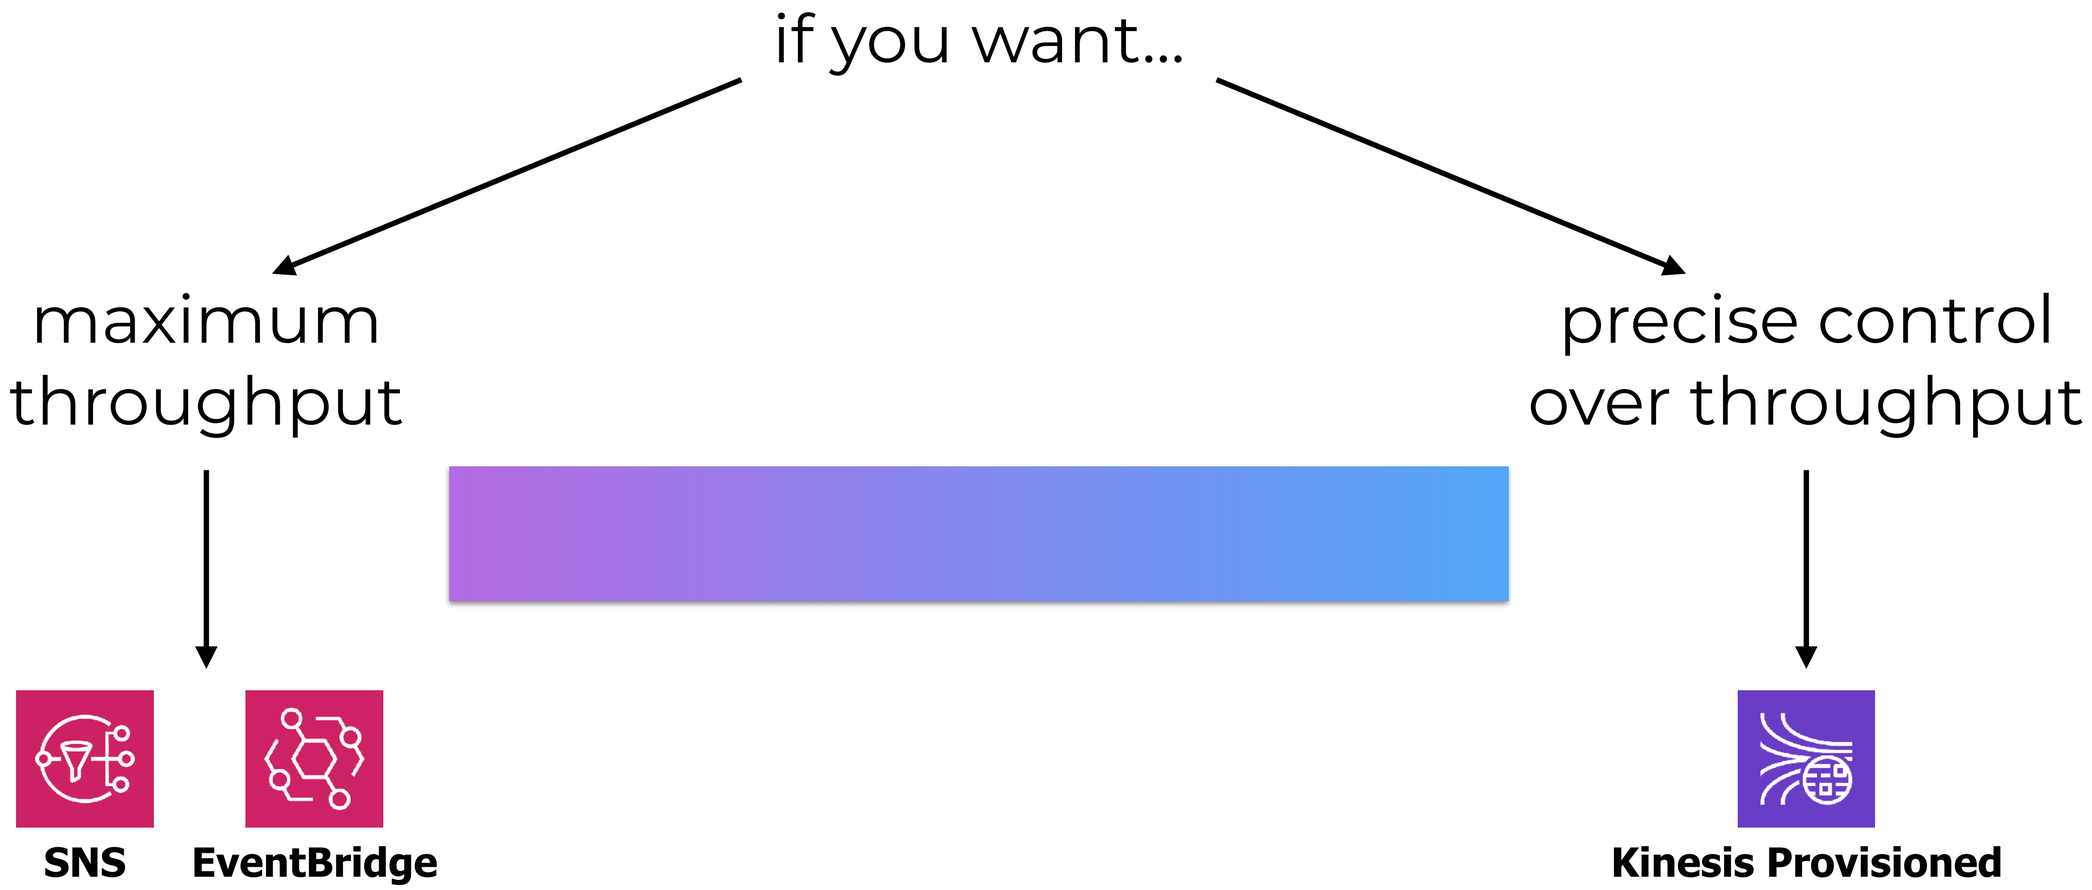
if you want...
maximum throughput
precise control over throughput
SNS
EventBridge
Kinesis Provisioned

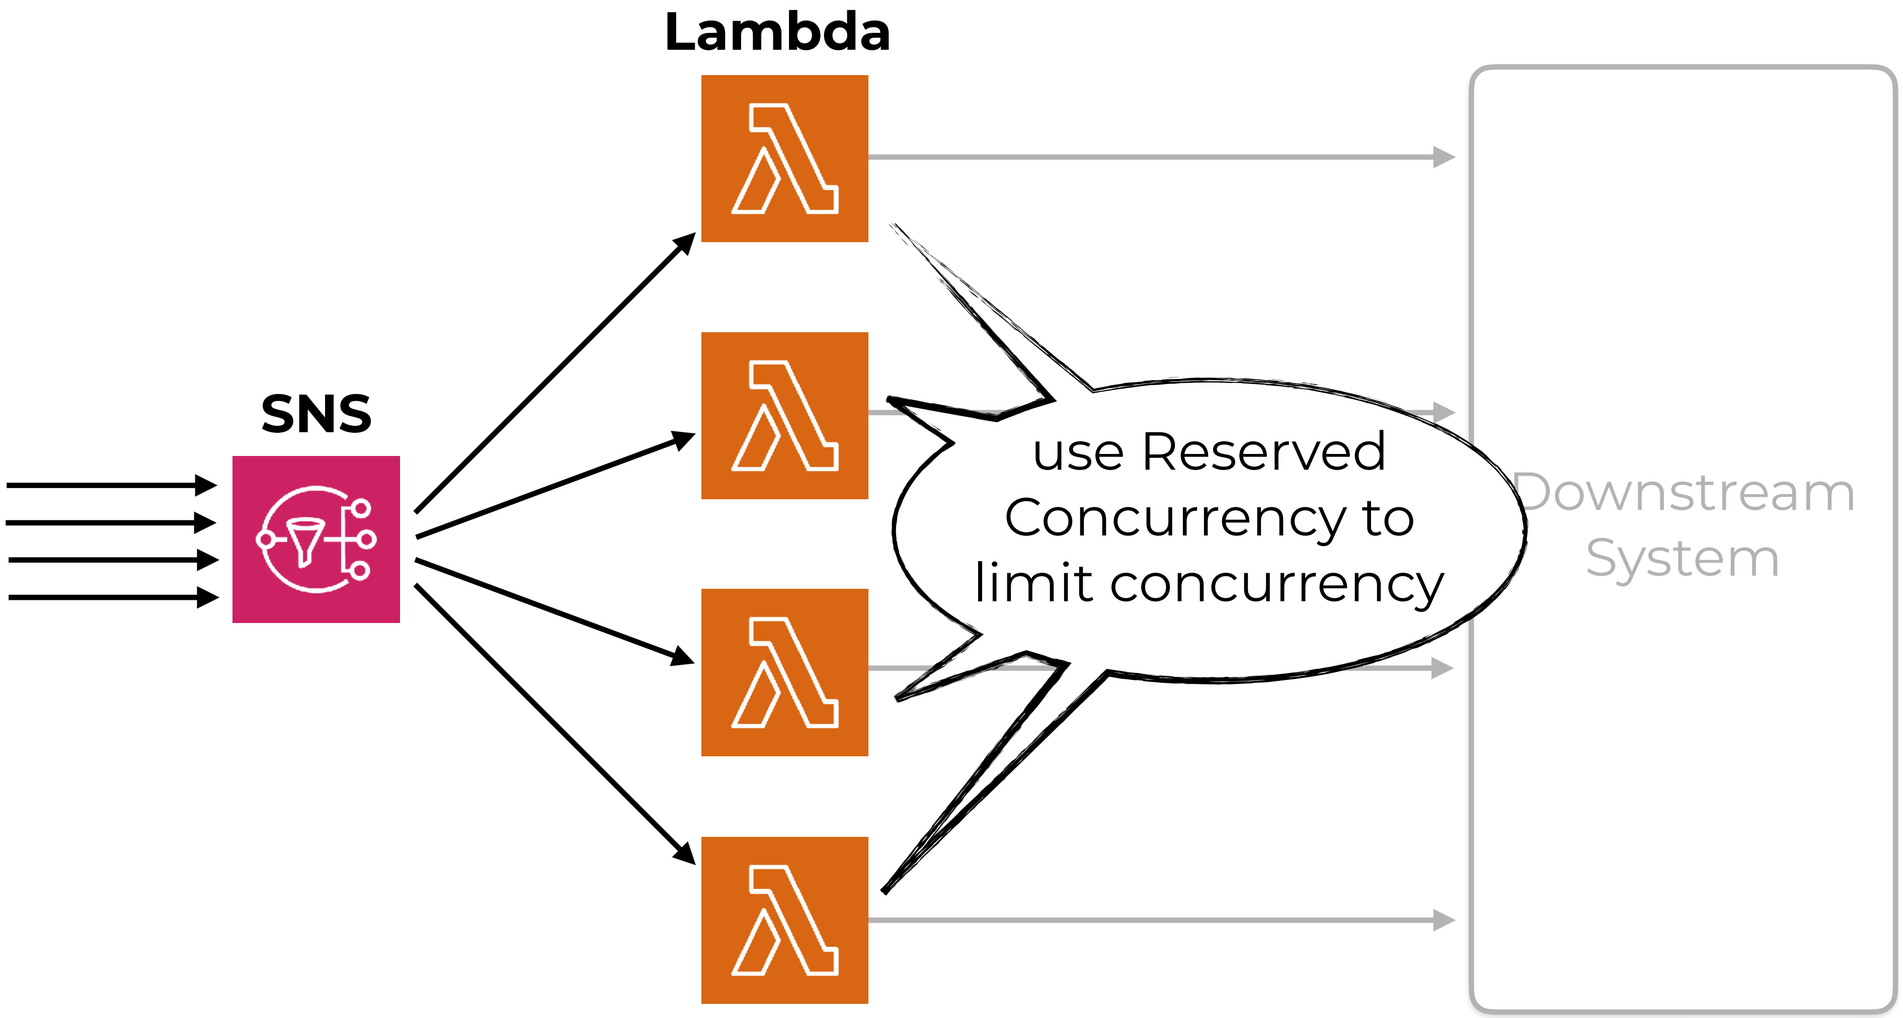
Lambda
SNS
use Reserved Concurrency to limit concurrency
Downstream System

Reserved Concurrency are taken out of the available regional Lambda concurrency

managing Reserved Concurrency for many functions is difficult and error prone, easy to create more problems than it solves

Costs

always factor **scale** into the equation

1 msg/s for a month, 1KB per msg

SNS

1 x 60s x 60m x 24hr x 30days
@ $0.50 per mil

$1.296

SQS

1 x 60s x 60m x 24hr x 30days
@ $0.40 per mil

$1.037

EventBridge

1 x 60s x 60m x 24hr x 30days
@ $1.00 per mil

$2.592

Kinesis

1 x 60s x 60m x 24hr x 30days
@ $0.014 per mil
+
24hrs x 30days
@ $0.015 per shard per hr

$10.836

Kinesis on-demand mode pricing

Pricing by Region

Region: US East (Ohio)

	Pricing
Per stream, per hour	$0.04
Data ingested, per GB (Includes 24-hour retention)	$0.08
Data retrievals, per GB	$0.04
Enhanced fan-out data retrievals, per GB	$0.05
Data stored, per GB-month (beyond 24 hours, up to 7 days)	$0.10
Data stored, per GB-month (beyond 7 days)	$0.023

1 msg/s for a month, 1KB per msg

SNS

1 x 60s x 60m x 24hr x 30days
@ $0.50 per mil

$1.296

SQS

1 x 60s x 60m x 24hr x 30days
@ $0.40 per mil

$1.037

EventBridge

1 x 60s x 60m x 24hr x 30days
@ $1.00 per mil

$2.592

Kinesis Provisioned

1 x 60s x 60m x 24hr x 30days
@ $0.014 per mil
+ 24hrs x 30days
@ $0.015 per shard per hr

$10.836

Kinesis On-Demand

1kb x 60s x 60m x 24hr x 30days
@ $0.08 per GB ingested
+ 24hrs x 30days
@ $0.04 per stream per hr

$28.998

1,000 msg/s for a month, 1KB per msg

SNS

1000 x 60s x 60m x 24hr x 30days
@ $0.50 per mil

$1296.00

SQS

1000 x 60s x 60m x 24hr x 30days
@ $0.40 per mil

$1036.80

EventBridge

1000 x 60s x 60m x 24hr x 30days
@ $1.00 per mil

$2592.00

Kinesis Provisioned

1000 x 60s x 60m x 24hr x 30days
@ $0.014 per mil
+ 24hrs x 30days
@ $0.015 per shard per hr

$47.088

Kinesis On-Demand

1000kb x 60s x 60m x 24hr x 30days
@ $0.08 per GB ingested
+ 24hrs x 30days
@ $0.04 per stream per hr

$226.55

services that charge by uptime are **order(s) of magnitude** cheaper when running at scale

National broadcaster of Finland

National broadcaster of Finland

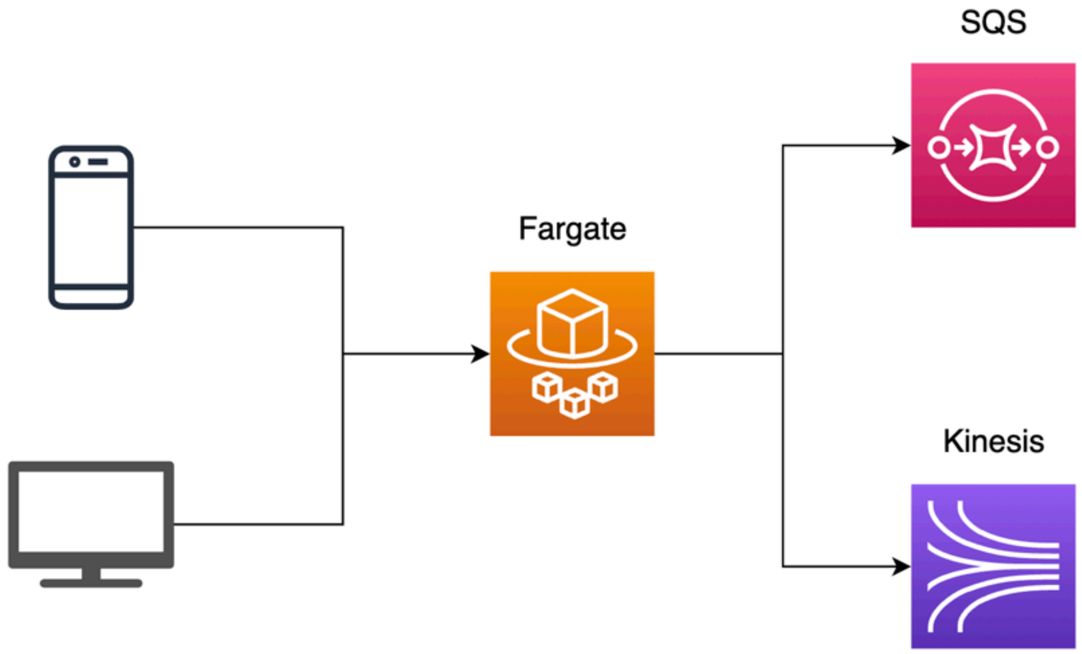

500+ millions events per day, peaks at 600K+ messages/min

yle

National broadcaster of Finland

Anahit Pogosova

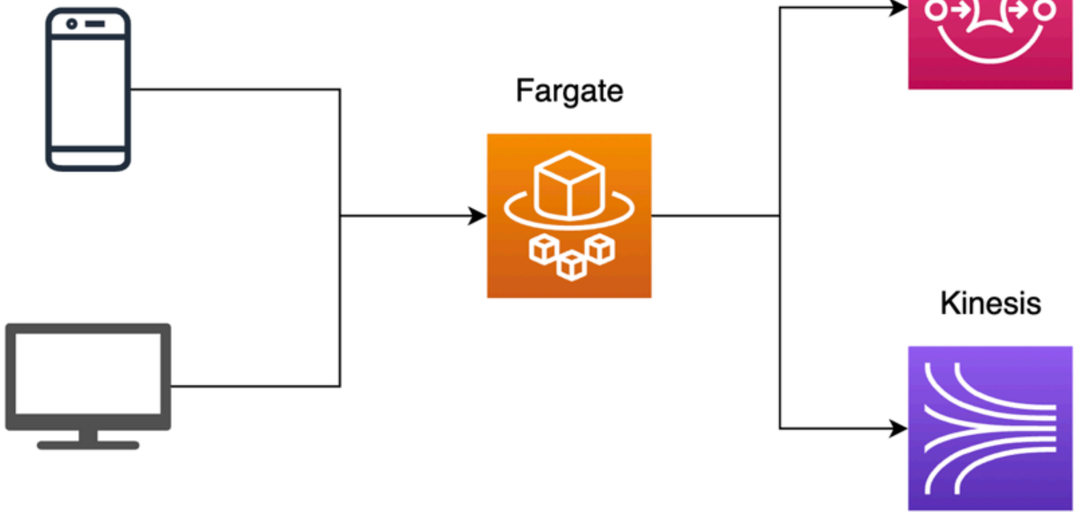

500+ millions events per day, peaks at 600K+ messages/min

yle

National broadcaster of Finland

Anahit Pogosova

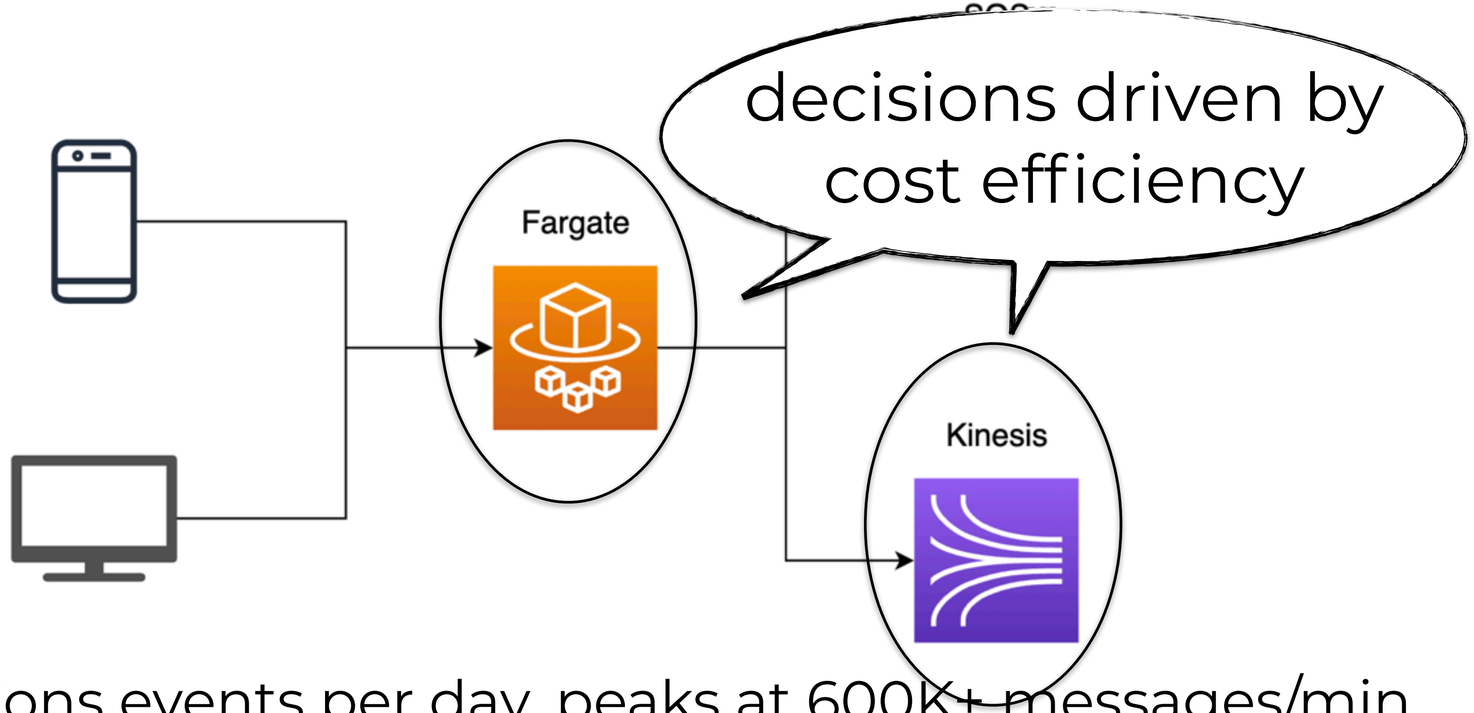

500+ millions events per day, peaks at 600K+ messages/min

National broadcaster of Finland

Anahit Pogosova

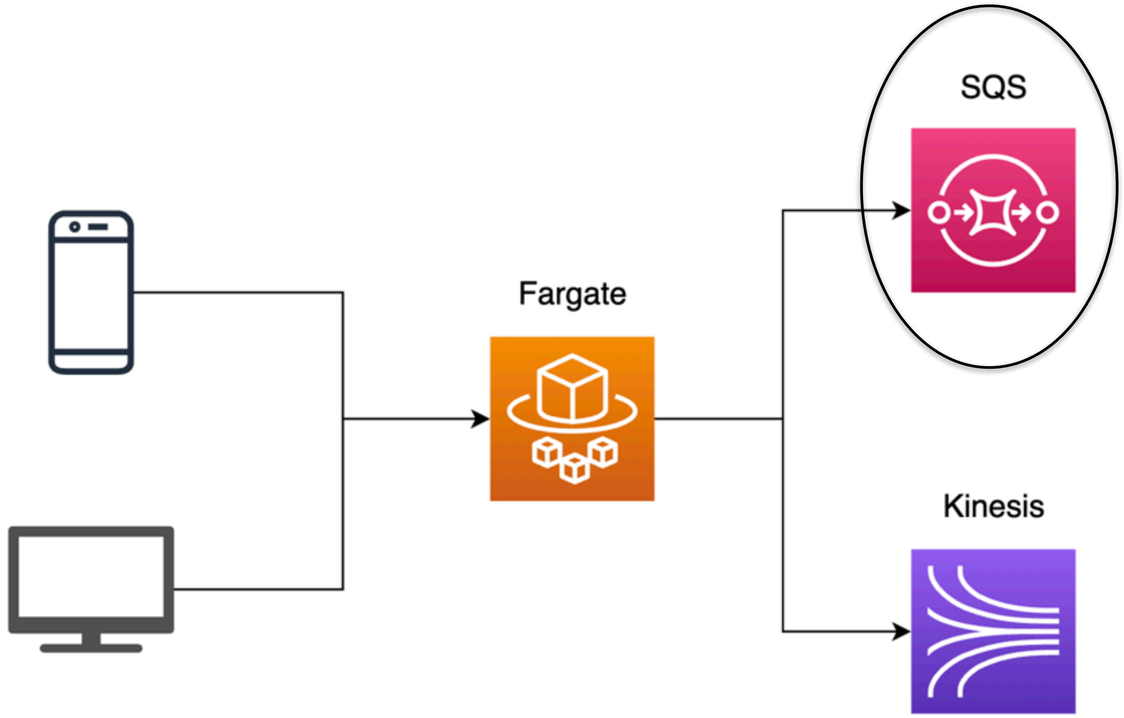

500+ millions events per day, peaks at 600K+ messages/min

yle

National broadcaster of Finland

Anahit Pogosova

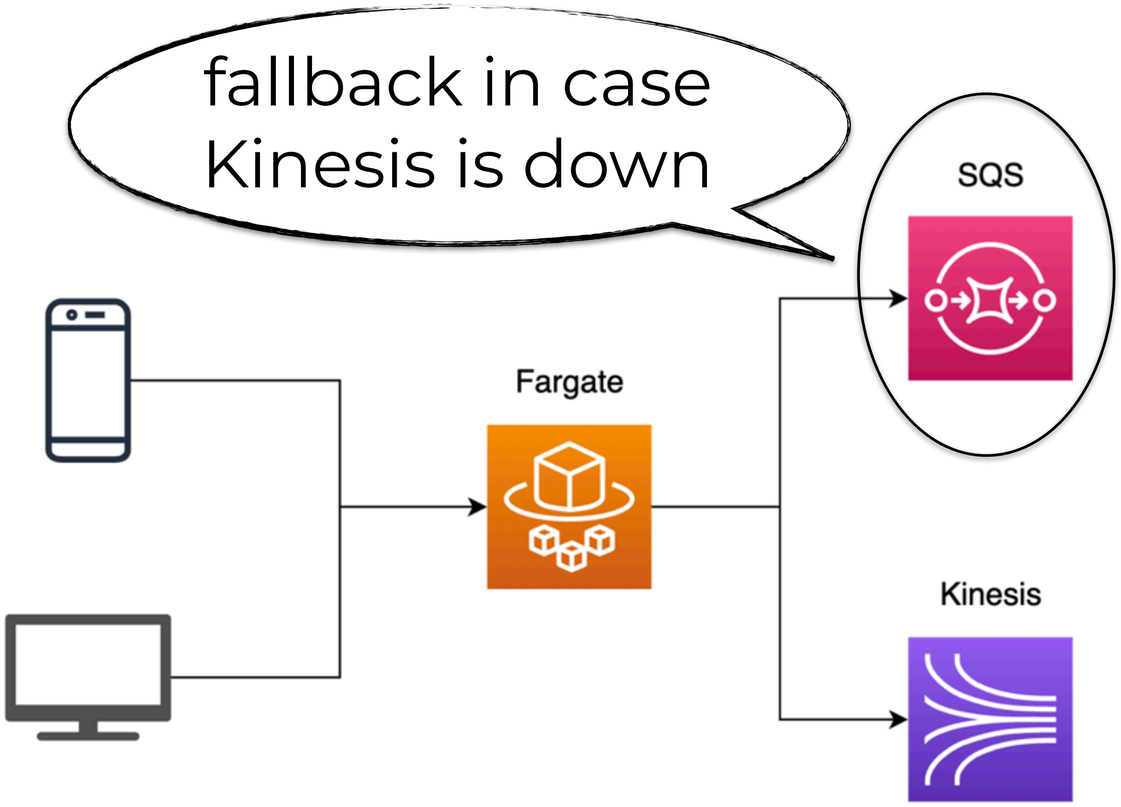

500+ millions events per day, peaks at 600K+ messages/min

Error Handling

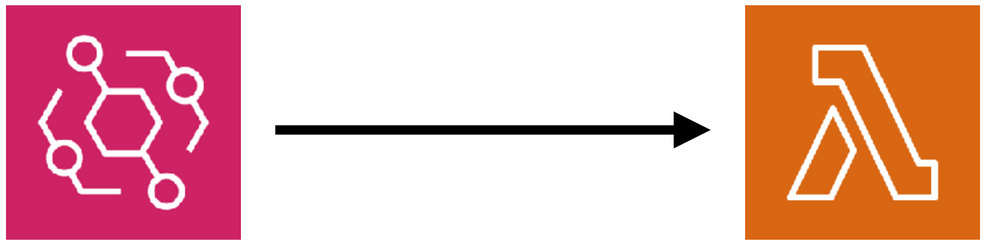

Edit asynchronous configuration

Retries Info

Maximum age of event

The maximum amount of time to keep unprocessed events in the queue.

6 h 0 min 0 sec

0, 1, 2

Retry attempts

The maximum number of times to retry when the function returns an error.

2

Dead-letter queue Info

Dead-letter queue service

You can send unprocessed events from an asynchronous invocation to an Amazon SQS queue or an Amazon SNS topic.

None

None

Amazon SNS

Amazon SQS

Cancel Save

DLQs only capture the failed event

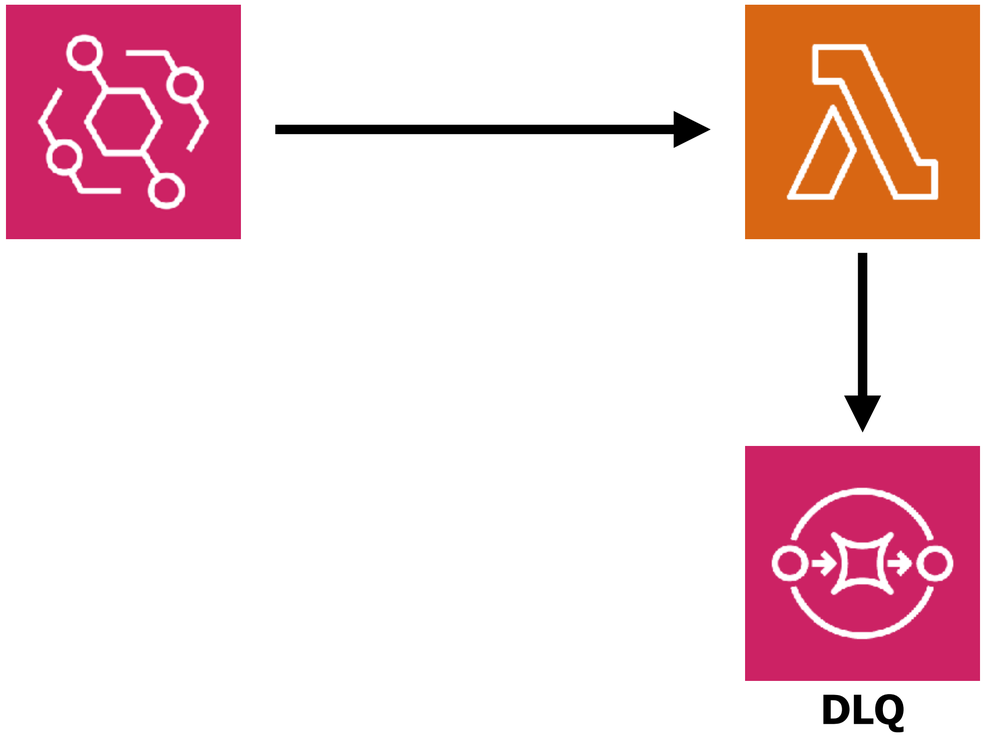
DLQ

Lambda > Functions > workshop-yancui-dev-get-index > Add destination

Add destination

Destination configuration

Send invocation records to a destination when your function is invoked asynchronously, or if your function processes records from a stream.

Source

The type of invocation that maps to the destination.

- (•) Asynchronous invocation
- ( ) Stream invocation

Condition

The condition for using the destination.

- (•) On failure
- ( ) On success

Destination type

An SQS queue, SNS topic, Lambda function, or EventBridge event bus.

SNS topic ▼

Destination

▼ 

Cancel Save

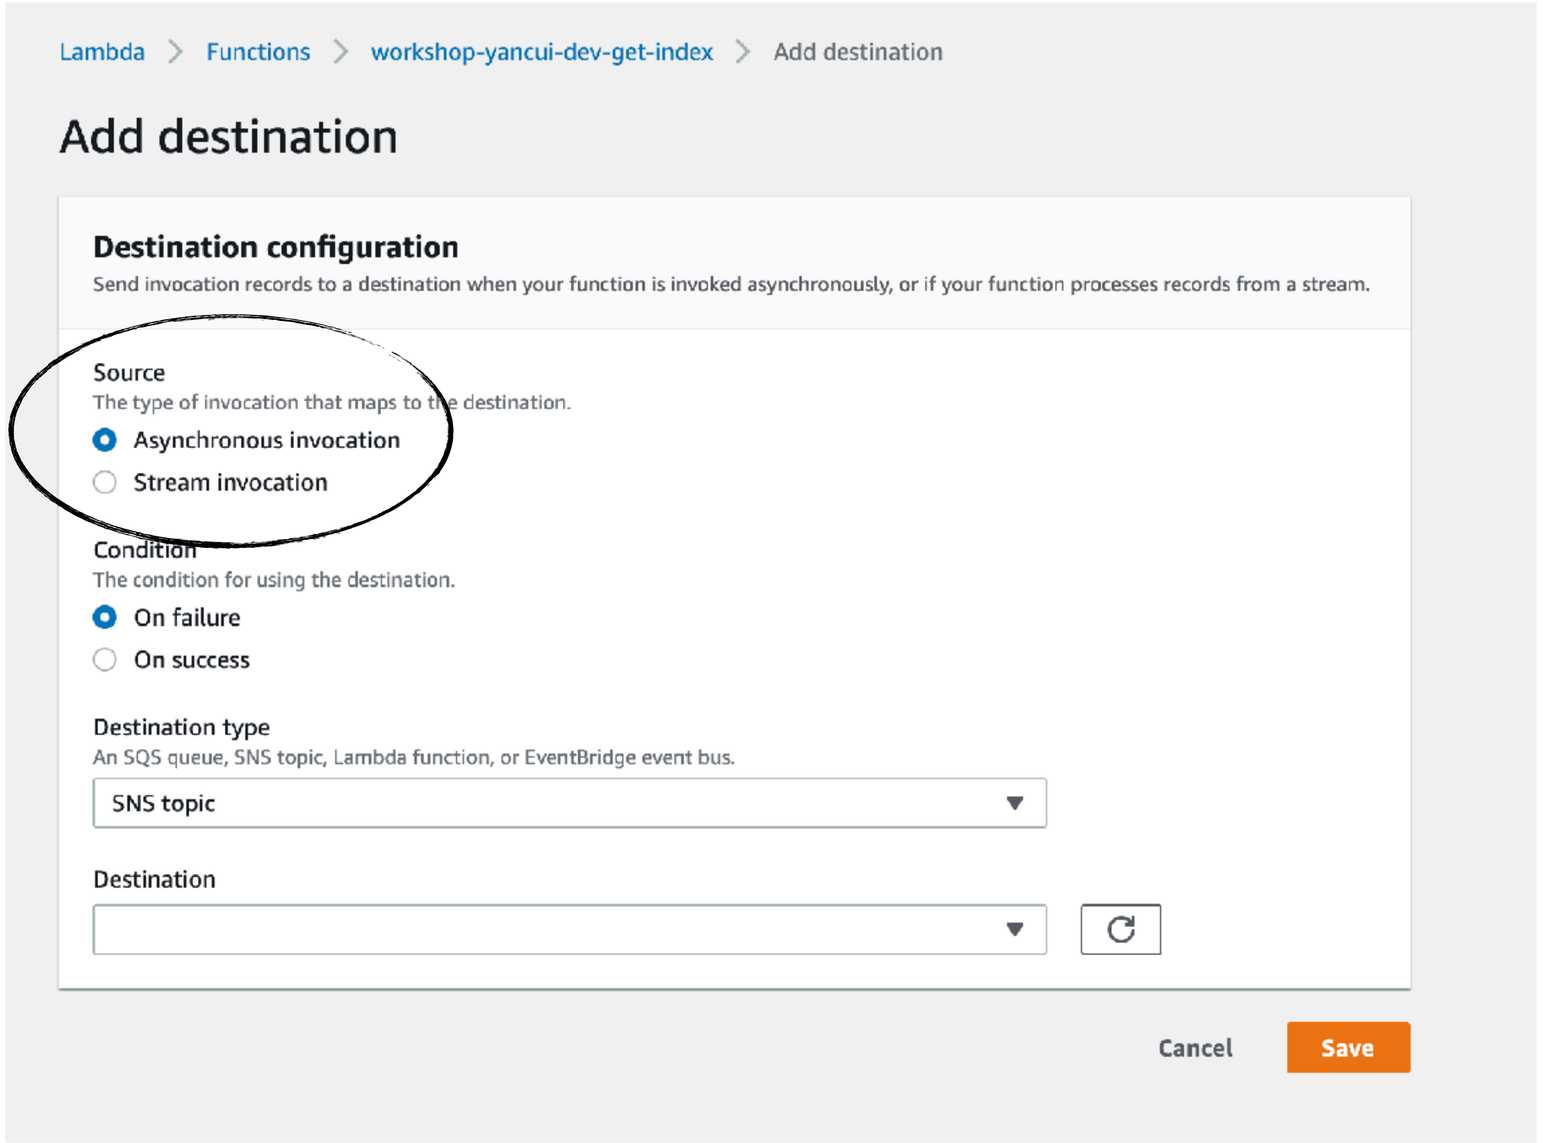
Lambda > Functions > workshop-yancui-dev-get-index > Add destination
Add destination
Destination configuration
Send invocation records to a destination when your function is invoked asynchronously, or if your function processes records from a stream.
Source
The type of invocation that maps to the destination.
Asynchronous invocation
Stream invocation
Condition
The condition for using the destination.
On failure
On success
Destination type
An SQS queue, SNS topic, Lambda function, or EventBridge event bus.
SNS topic
Destination
Cancel
Save

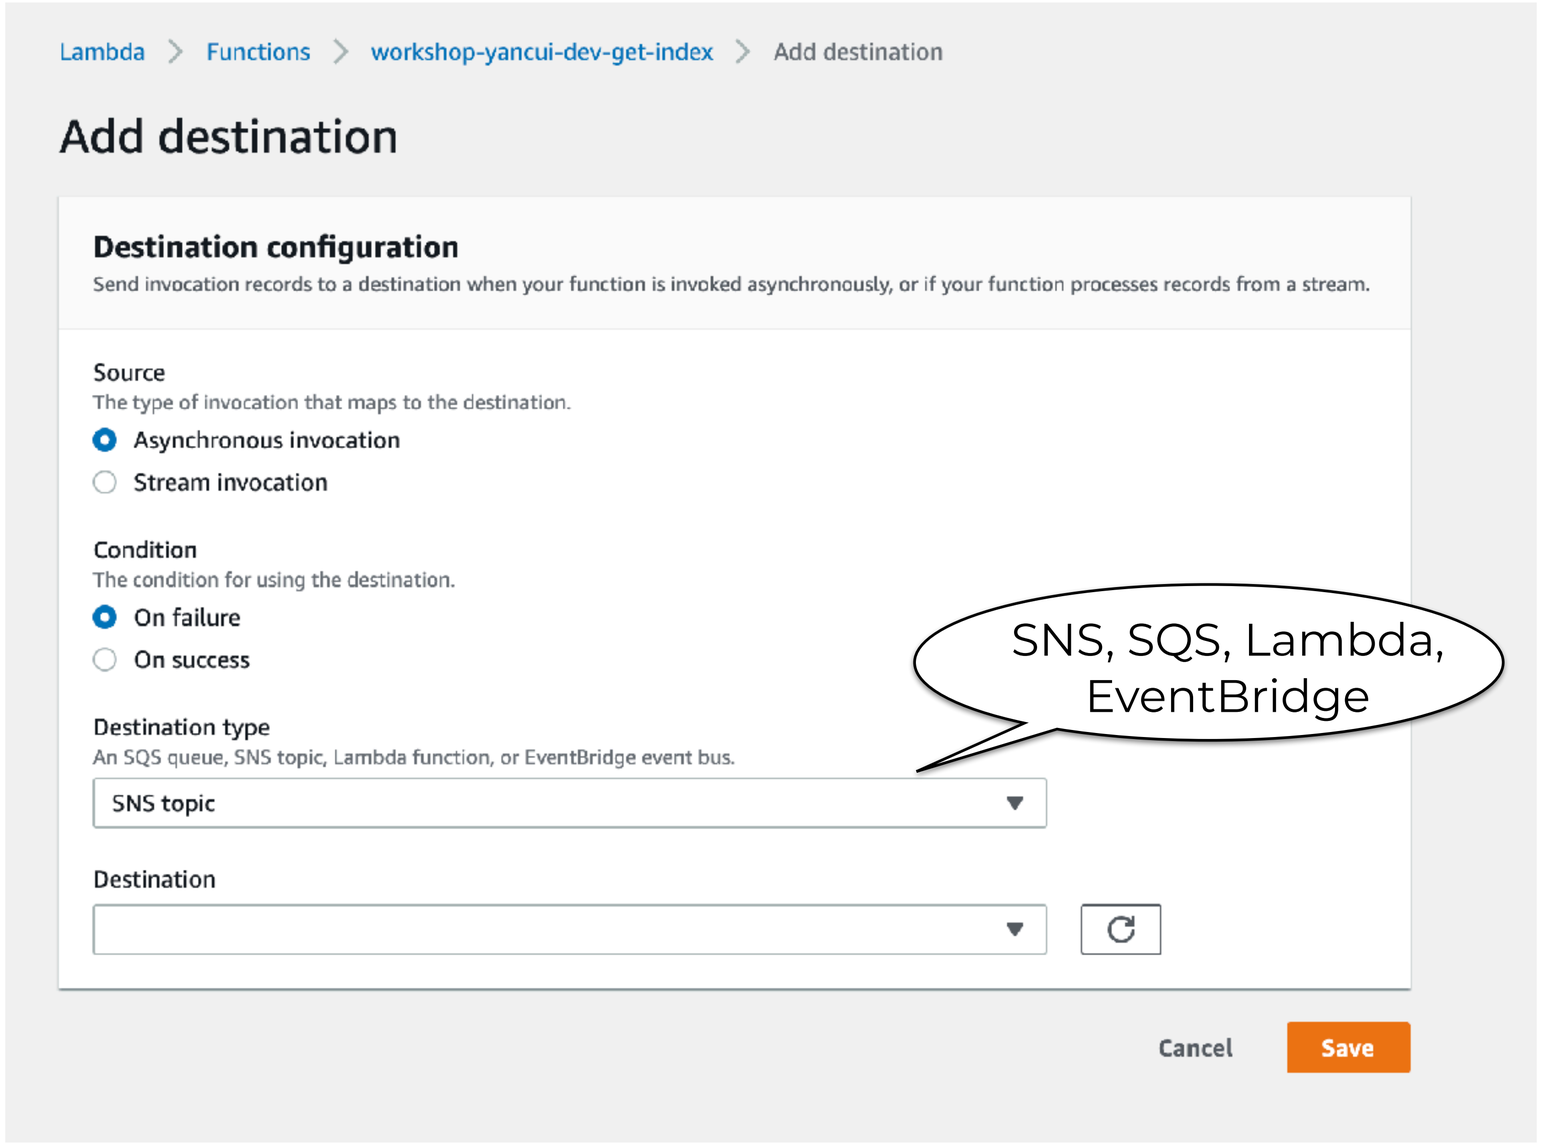
Lambda > Functions > workshop-yancui-dev-get-index > Add destination
Add destination
Destination configuration
Send invocation records to a destination when your function is invoked asynchronously, or if your function processes records from a stream.
Source
The type of invocation that maps to the destination.
Asynchronous invocation
Stream invocation
Condition
The condition for using the destination.
On failure
On success
Destination type
An SQS queue, SNS topic, Lambda function, or EventBridge event bus.
SNS topic
Destination
SNS, SQS, Lambda, EventBridge
Cancel
Save

```
{
  "version": "1.0",
  "timestamp": "2019-11-24T21:52:47.333Z",
  "requestContext": {
    "requestId": "8ea123e4-1db7-4aca-ad10-d9ca1234c1fd",
    "functionArn": "arn:aws:lambda:sa-east-1:123456678912:function:event-destinations:$LATEST",
    "condition": "RetriesExhausted",
    "approximateInvokeCount": 3
  },
  "requestPayload": {
    "Success": false
  },
  "responseContext": {
    "statusCode": 200,
    "executedVersion": "$LATEST",
    "functionError": "Handled"
  },
  "responsePayload": {
    "errorMessage": "Failure from event, Success = false, I am failing!",
    "errorType": "Error",
    "stackTrace": [
      "exports.handler (\/var\/task\/index.js:18:18)"
    ]
  }
}
```

```
{
  "version": "1.0",
  "timestamp": "2019-11-24T21:52:47.333Z",
  "requestContext": {
    "requestId": "8ea123e4-1db7-4aca-ad10-d9ca1234c1fd",
    "functionArn": "arn:aws:lambda:sa-east-1:123456678912:function:event-destinations:$LATEST",
    "condition": "RetriesExhausted",
    "approximateInvokeCount": 3
  },
  "requestPayload": {
    "Success": false
  },
  "responseContext": {
    "statusCode": 200,
    "executedVersion": "$LATEST",
    "functionError": "Handled"
  },
  "responsePayload": {
    "errorMessage": "Failure from event, Success = false, I am failing!",
    "errorType": "Error",
    "stackTrace": [
      "exports.handler (\/var\/task\/index.js:18:18)"
    ]
  }
}
```


```
{
  "version": "1.0",
  "timestamp": "2019-11-24T21:52:47.333Z",
  "requestContext": {
    "requestId": "8ea123e4-1db7-4aca-ad10-d9ca1234c1fd",
    "functionArn": "arn:aws:lambda:sa-east-1:123456678912:function:event-destinations:$LATEST",
    "condition": "RetriesExhausted",
    "approximateInvokeCount": 3
  },
  "requestPayload": {
    "Success": false
  },
  "responseContext": {
    "statusCode": 200,
    "executedVersion": "$LATEST",
    "functionError": "Handled"
  },
  "responsePayload": {
    "errorMessage": "Failure from event, Success = false, I am failing!",
    "errorType": "Error",
    "stackTrace": [
      "exports.handler (\/var\/task\/index.js:18:18)"
    ]
  }
}
```

prefer Lambda destination over DLQs

(you can use both together, but there's no clear reason for doing so)

```
{
  "Body": "{
    \"requestContext\":{
      \"requestId\":\"0cdf90e8-6290-4834-a46c-cffdd6277c4f\",
      \"functionArn\":\"arn:aws:lambda:us-east-1:374852340823:function:aws-sdk-example\",
      \"condition\":\"RetryAttemptsExhausted\",
      \"approximateInvokeCount\":11
    },
    \"responseContext\":{
      \"statusCode\":200,
      \"executedVersion\":
      \"$LATEST\",
      \"functionError\":
      \"Unhandled\"
    },
    \"version\":\"1.0\",
    \"timestamp\":\"2020-02-08T19:53:20.998Z\",
    \"KinesisBatchInfo\":{
      \"shardId\":\"shardId-000000000000\",
      \"startSequenceNumber\":\"496040450608013792680370773552622891186968974110626938 90\",
      \"endSequenceNumber\":\"496040450608013792680370773552622891186968974110626938 90\",
      \"approximateArrivalOfFirstRecord\":\"2020-02-08T19:47:05.582Z\",
      \"approximateArrivalOfLastRecord\":\"2020-02-08T19:47:05.582Z\",
      \"batchSize\":1,
      \"streamArn\":\"arn:aws:kinesis:us-east-1:374852340823:stream/yc-test\"
    }
  }"
}
```

```
{
  "Body": "{
    \"requestContext\":{
      \"requestId\":\"0cdf90e8-6290-4834-a46c-cffdd6277c4f\",
      \"functionArn\":\"arn:aws:lambda:us-east-1:374852340823:function:aws-sdk-example\",
      \"condition\":\"RetryAttemptsExhausted\",
      \"approximateInvokeCount\":11
    },
    \"responseContext\":{
      \"statusCode\":200,
      \"executedVersion\":
      \"$LATEST\",
      \"functionError\":
      \"Unhandled\"
    },
    \"version\":\"1.0\",
    \"timestamp\":\"2020-02-08T19:53:20.998Z\",
    \"KinesisBatchInfo\":{
      \"shardId\":\"shardId-000000000000\",
      \"startSequenceNumber\":\"496040450608013792680370773552622891186968974110626938 90\",
      \"endSequenceNumber\":\"496040450608013792680370773552622891186968974110626938 90\",
      \"approximateArrivalOfFirstRecord\":\"2020-02-08T19:47:05.582Z\",
      \"approximateArrivalOfLastRecord\":\"2020-02-08T19:47:05.582Z\",
      \"batchSize\":1,
      \"streamArn\":\"arn:aws:kinesis:us-east-1:374852340823:stream/yc-test\"
    }
  }"
}
```

Lambda > Functions > workshop-yancui-dev-get-index > Add destination

Add destination

Destination configuration

Send invocation records to a destination when your function is invoked asynchronously, or if your function processes records from a stream.

Source

The type of invocation that maps to the dest

- (•) Asynchronous invocation
- ( ) Stream invocation

Condition

The condition for using the destination.

- (•) On failure
- ( ) On success

Destination type

An SQS queue, SNS topic, Lambda function, or EventBridge event bus.

SNS topic

Destination

Cancel Save

SNS ✓

SQS ✗

EventBridge ✓

Kinesis ✓

DynamoDB Stream ✓

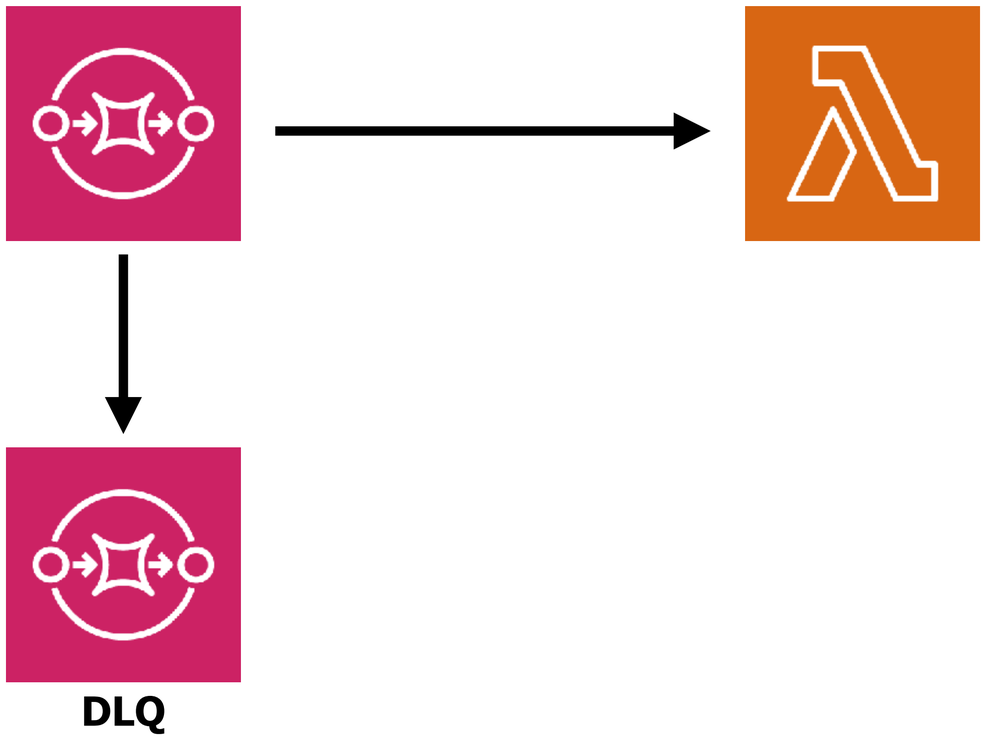
DLQ

Observability

A measure of how well the internal state of an application can be inferred from its external outputs

X-Ray

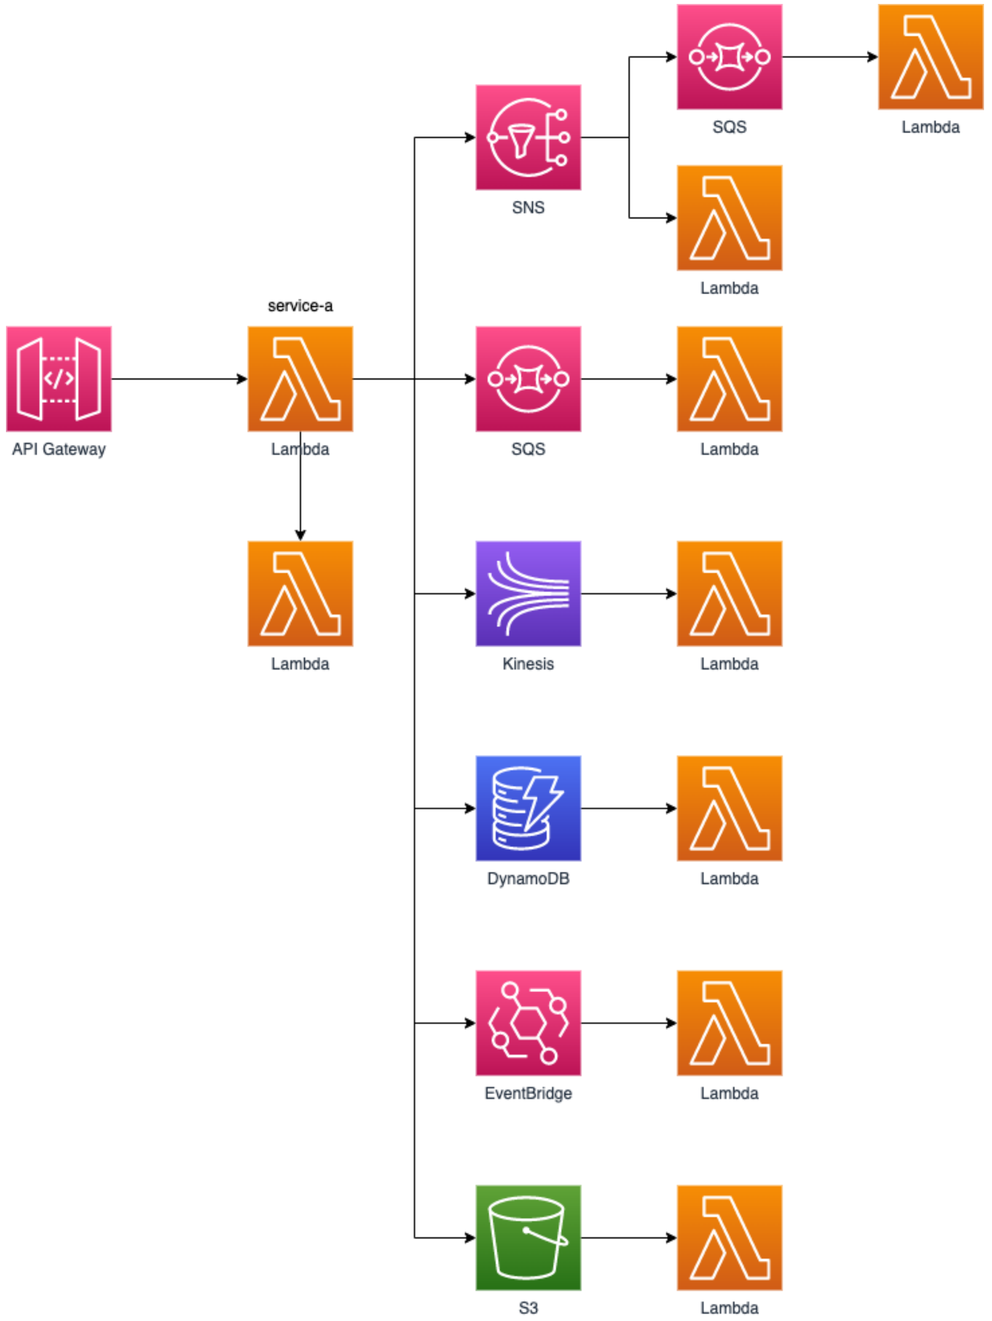
API Gateway
service-a
Lambda
Lambda
SNS
SQS
Lambda
Lambda
SQS
Lambda
Kinesis
Lambda
DynamoDB
Lambda
EventBridge
Lambda
S3
Lambda

Trace overview

Group by: URL

Done 100% scanned (f

URL	AVG RESPONSE TIME	% OF TRACES	RESPONSE
-	461 ms	83.33%	5 OK, 0 Throttled, 0 Errors, 0 Faults
https://a4ut5pds5m.execute-api.us-east-1.amazona...	1.8 sec	16.67%	1 OK, 0 Throttled, 0 Errors, 0 Faults

Trace list

ID	AGE	METHOD	RESPONSE	RESPONSE TIME	URL	CLIENT IP	ANNOTATIONS
...703630b9c4d5	7.8 sec		200	725 ms			0
...41fd6d8940b7	7.8 sec		200	558 ms			0
...68aa45b0d4d2	8.8 sec			-			0
...f842d677cd4b	8.8 sec		200	497 ms			0
...ccfc0fc875b00	8.8 sec		200	523 ms			0
...68aa45b0d4d2	9.8 sec	GET	200	1.8 sec	https://a4ut5pds5m....	84.85.234.72	2

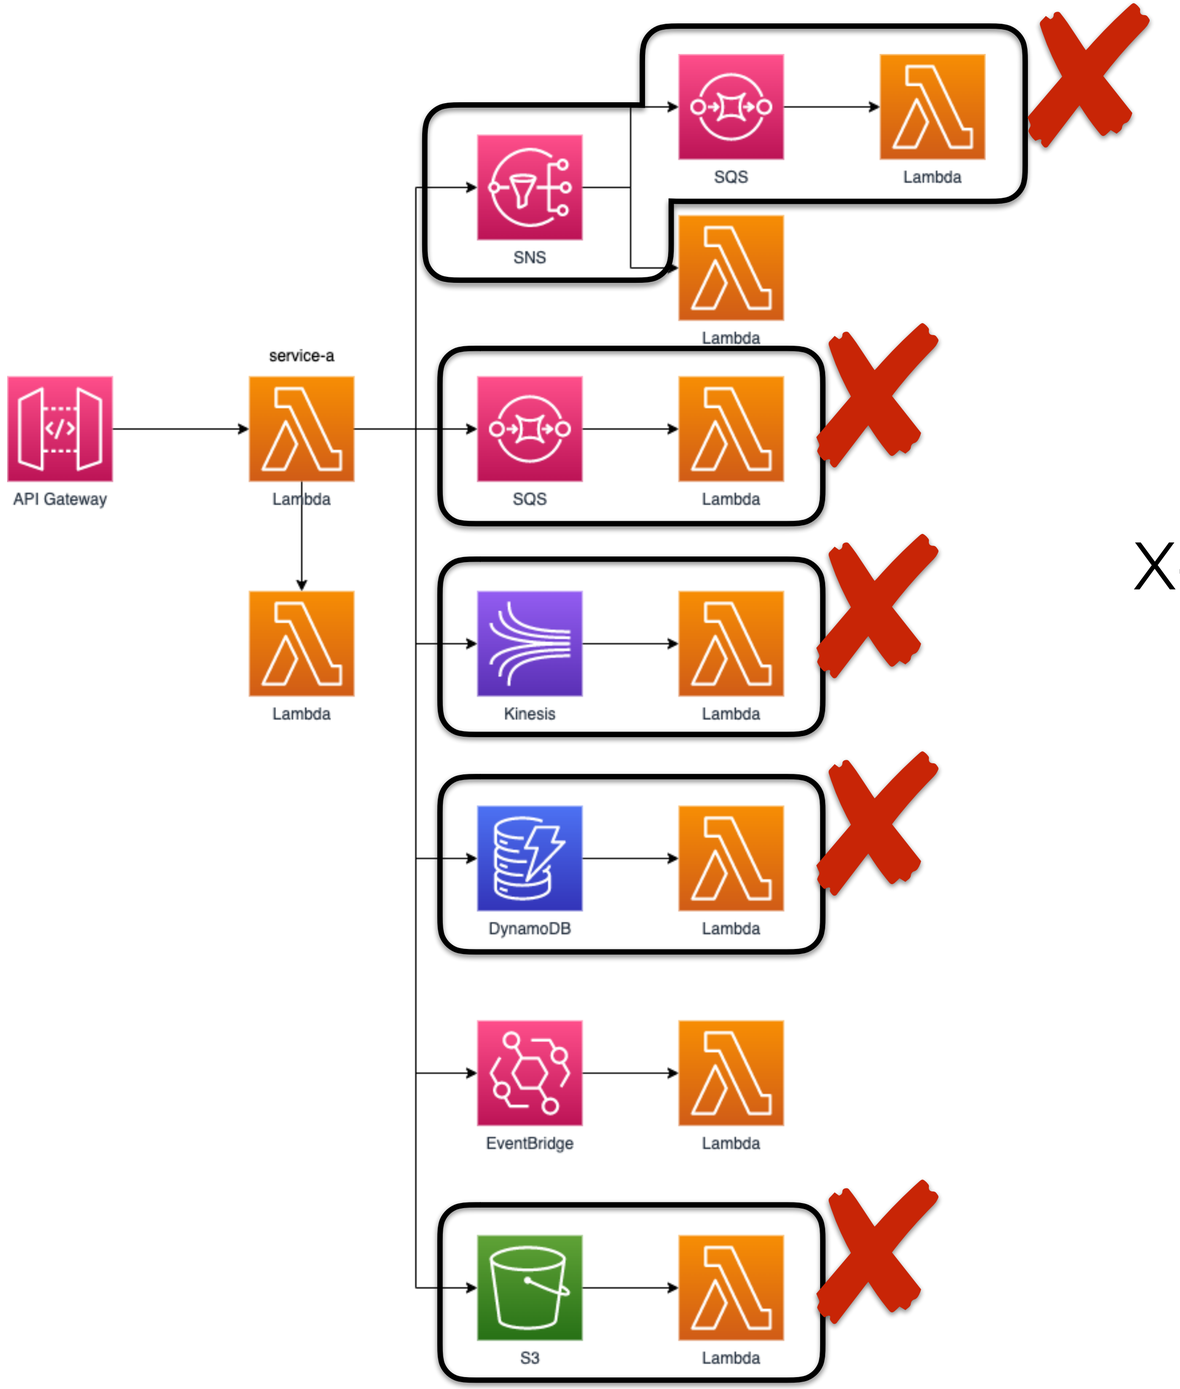

X-Ray doesn't trace through many popular messaging services.

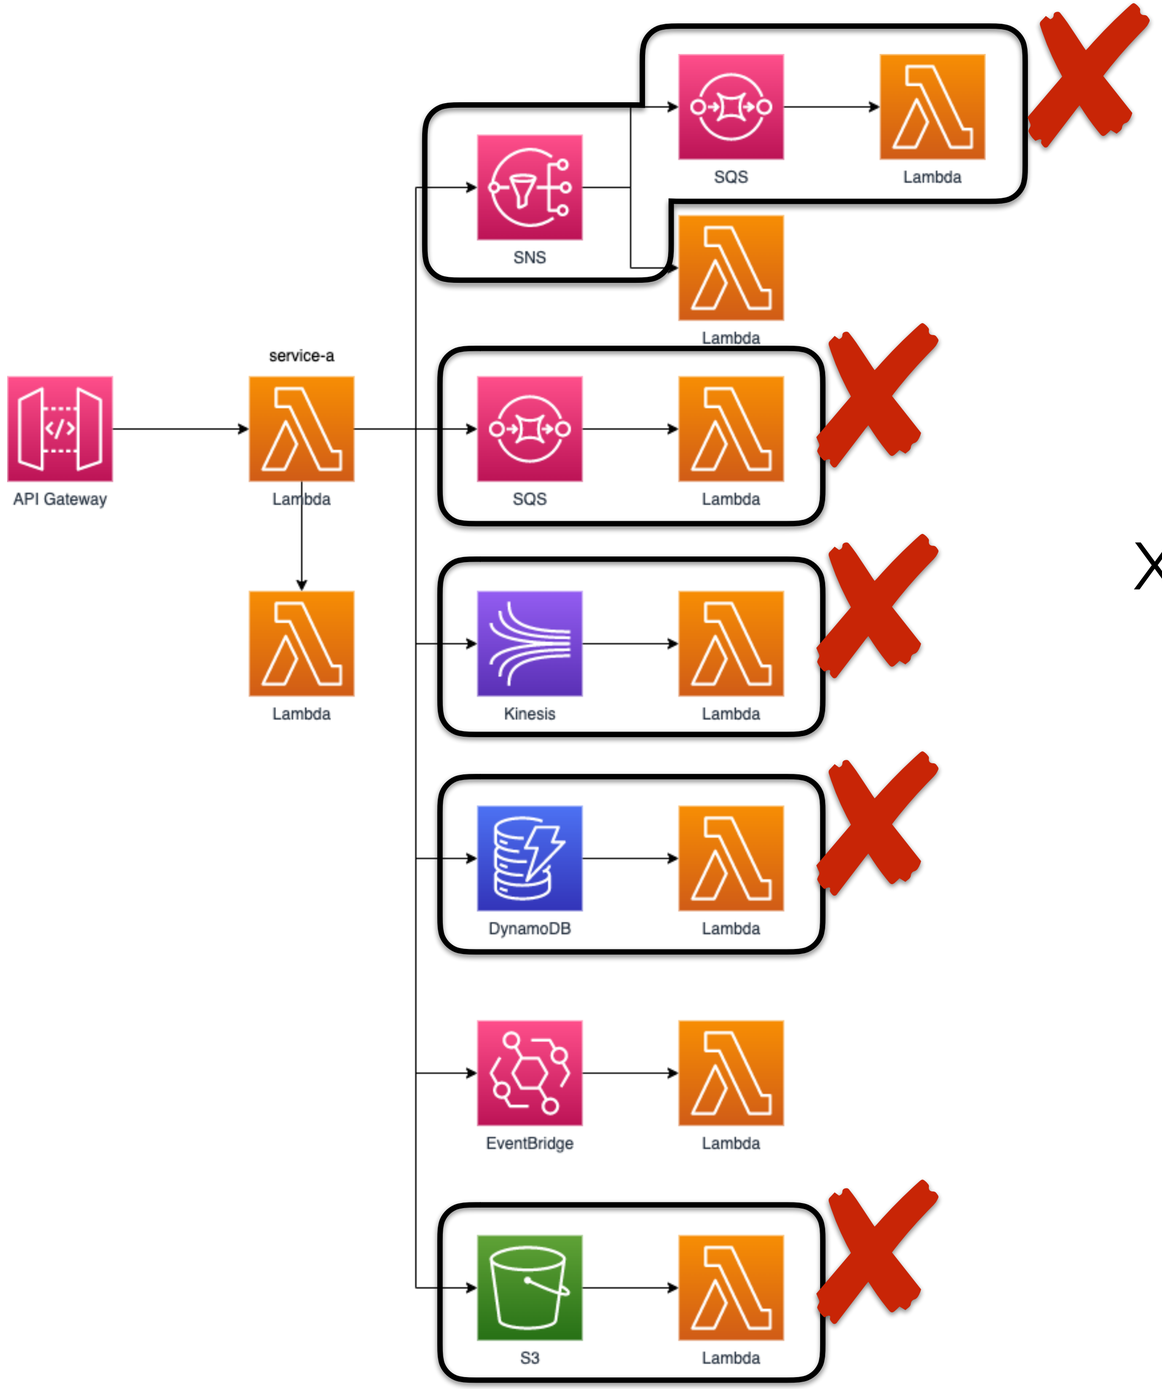

X-Ray doesn't trace through many popular messaging services.

Need separate solution (e.g. correlation IDs) for Lambda logs.

Transaction ID: 1c6eabf86b6b64b920a19321

Start Time	Duration	Cost	Region	Issues
3:34:53 PM	2,350 ms	< $0.01	us-east-1	0

Show Similar Transactions

GRAPH TIMELINE LOGS

Expand

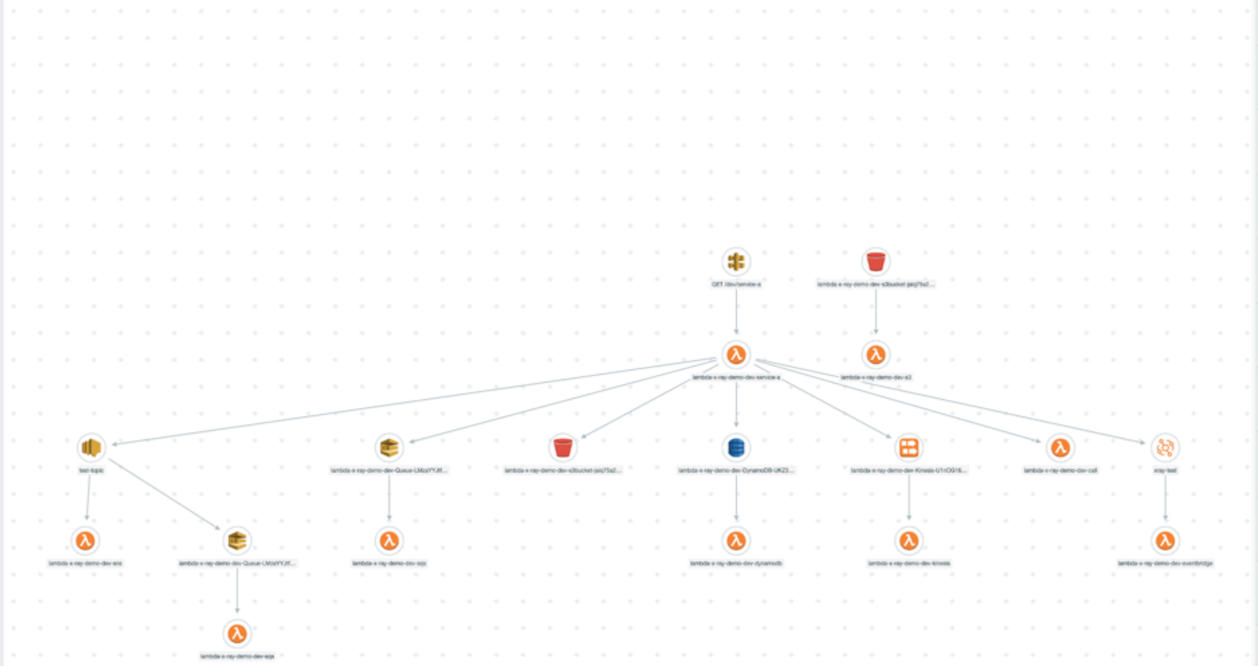

Log Entries (30)

Search in logs

Timestamp	Name	Message
3:34:53 PM...	lambda-x-ray-demo-dev-s...	START RequestId: 74db077a-6...
3:34:54 PM...	lambda-x-ray-demo-dev-s...	INFO Call "eventbridge.putEven...
3:34:54 PM...	lambda-x-ray-demo-dev-s...	INFO Call "s3.putObject" is not ...
3:34:54 PM...	lambda-x-ray-demo-dev-s...	INFO Call "kinesis.putRecord" i...
3:34:54 PM...	lambda-x-ray-demo-dev-sqs	START RequestId: daa6ee23-5...
3:34:54 PM...	lambda-x-ray-demo-dev-call	START RequestId: c4267b0c-8...
3:34:54 PM...	lambda-x-ray-demo-dev-sqs	START RequestId: 4d5ab0f8-a4...
3:34:54 PM...	lambda-x-ray-demo-dev-sqs	END RequestId: daa6ee23-5d6...
3:34:54 PM...	lambda-x-ray-demo-dev-sqs	REPORT RequestId: daa6ee23-...
3:34:54 PM...	lambda-x-ray-demo-dev-sns	START RequestId: 94de0437-f6...
3:34:54 PM...	lambda-x-ray-demo-dev-call	END RequestId: c4267b0c-8ec...
3:34:54 PM...	lambda-x-ray-demo-dev-call	REPORT RequestId: c4267b0c-...
3:34:54 PM...	lambda-x-ray-demo-dev-s...	END RequestId: 74db077a-6b6...

lumigo

Transaction ID: 1c6eabf86b6b64b920a19321

Start Time	Duration	Cost	Region	Issues
3:34:53 PM	2,350 ms	< $0.01	us-east-1	0

Show Similar Transactions

GRAPH TIMELINE LOGS

Expand

all the indirect invocations are accounted for

Log Entries (30)

Search in logs

Timestamp	Name	Message
3:34:53 PM...	lambda-x-ray-demo-dev-s...	START RequestId: 74db077a-6...
3:34:54 PM...	lambda-x-ray-demo-dev-s...	INFO Call "eventbridge.putEven...
3:34:54 PM...	lambda-x-ray-demo-dev-s...	INFO Call "s3.putObject" is not ...
3:34:54 PM...	lambda-x-ray-demo-dev-s...	INFO Call "kinesis.putRecord" i...
3:34:54 PM...	lambda-x-ray-demo-dev-sqs	START RequestId: daa6ee23-5...
3:34:54 PM...	lambda-x-ray-demo-dev-call	START RequestId: c4267b0c-8...
3:34:54 PM...	lambda-x-ray-demo-dev-sqs	START RequestId: 4d5ab0f8-a4...
3:34:54 PM...	lambda-x-ray-demo-dev-sqs	END RequestId: daa6ee23-5d6...
3:34:54 PM...	lambda-x-ray-demo-dev-sqs	REPORT RequestId: daa6ee23-...
3:34:54 PM...	lambda-x-ray-demo-dev-sns	START RequestId: 94de0437-f6...
3:34:54 PM...	lambda-x-ray-demo-dev-call	END RequestId: c4267b0c-8ec...
3:34:54 PM...	lambda-x-ray-demo-dev-call	REPORT RequestId: c4267b0c-...
3:34:54 PM...	lambda-x-ray-demo-dev-s...	END RequestId: 74db077a-6b6...

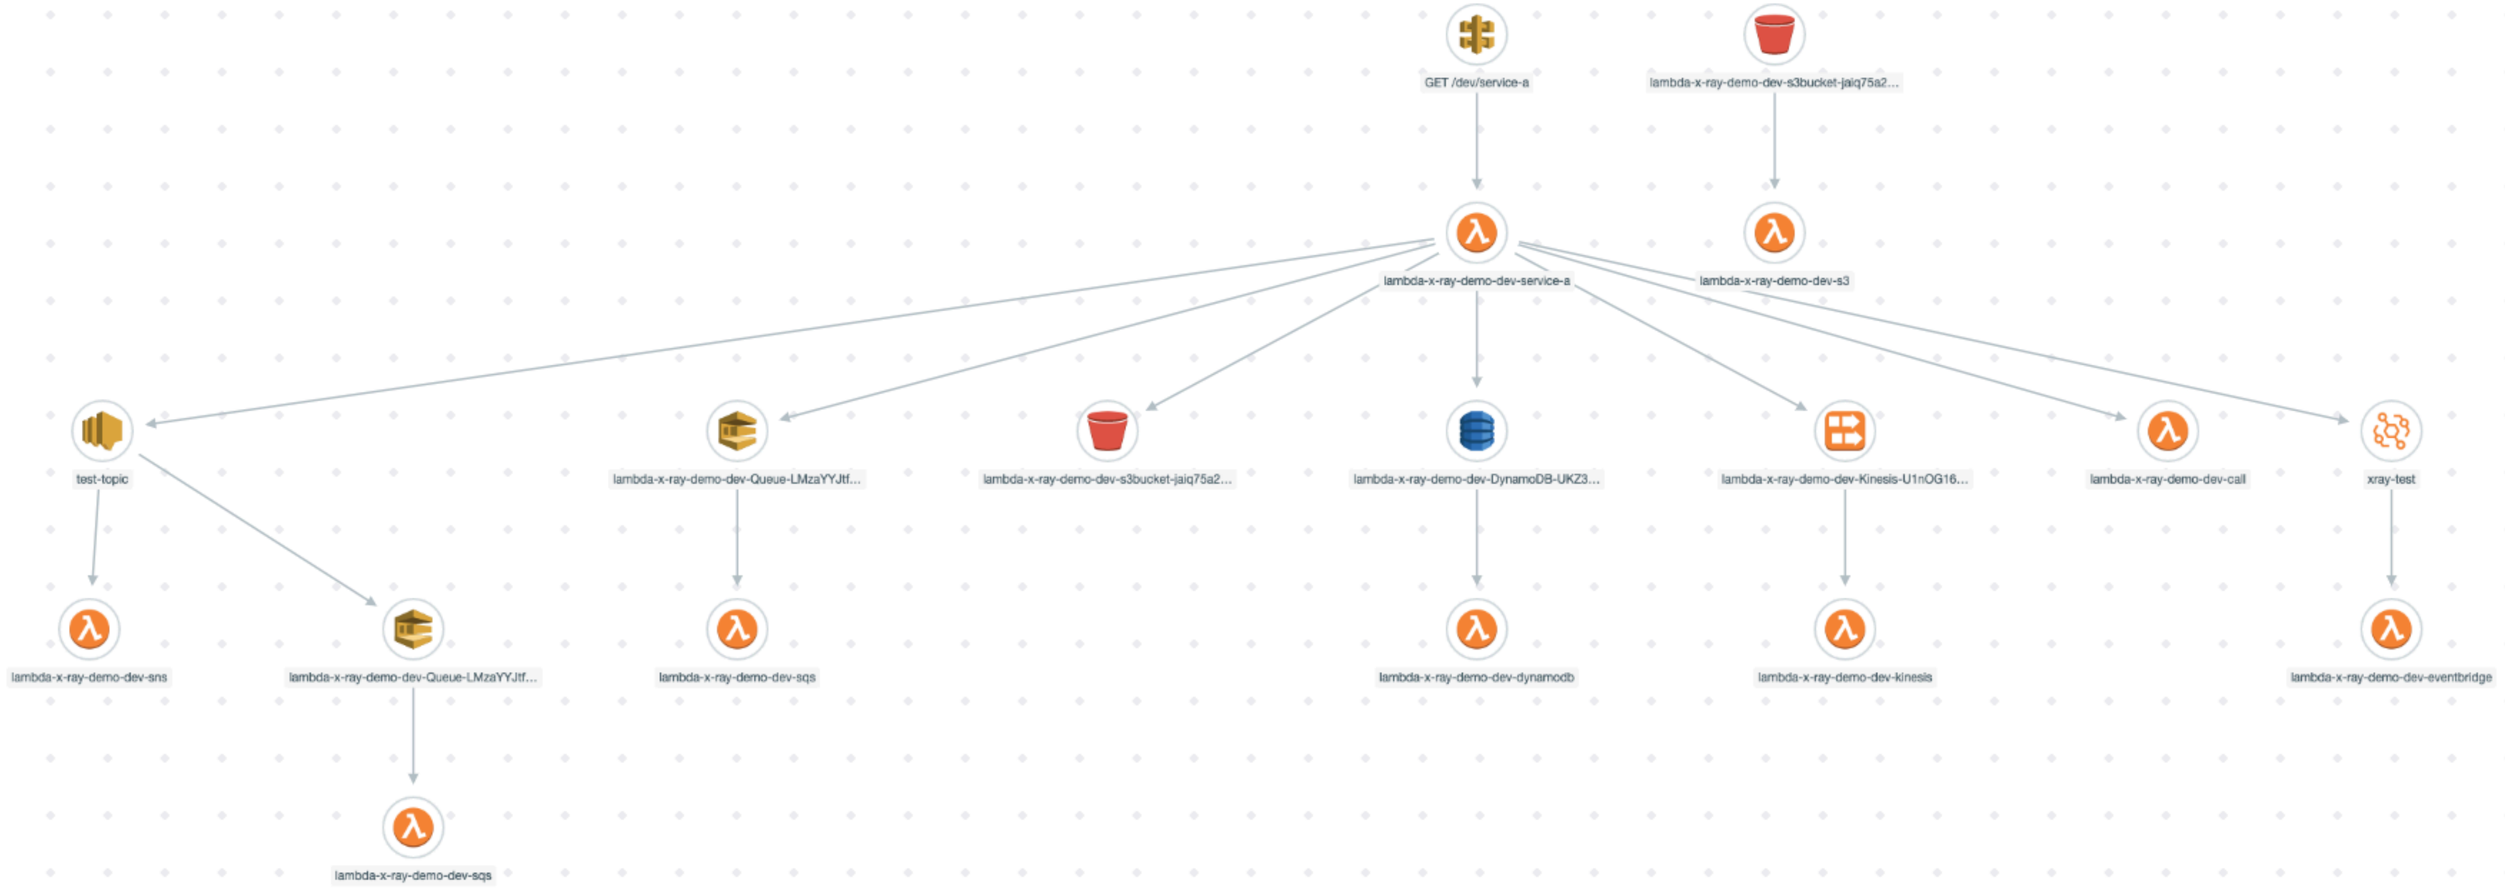
GET /dev/service-a
lambda-x-ray-demo-dev-s3bucket-jaiq75a2...
lambda-x-ray-demo-dev-service-a
lambda-x-ray-demo-dev-s3
test-topic
lambda-x-ray-demo-dev-Queue-LMzaYYJtf...
lambda-x-ray-demo-dev-s3bucket-jaiq75a2...
lambda-x-ray-demo-dev-DynamoDB-UKZ3...
lambda-x-ray-demo-dev-Kinesis-U1nOG16...
lambda-x-ray-demo-dev-call
xray-test
lambda-x-ray-demo-dev-sns
lambda-x-ray-demo-dev-Queue-LMzaYYJtf...
lambda-x-ray-demo-dev-sqs
lambda-x-ray-demo-dev-dynamodb
lambda-x-ray-demo-dev-kinesis
lambda-x-ray-demo-dev-eventbridge
lambda-x-ray-demo-dev-sqs

lumigo

Transaction ID: 1c6eabf86b6b64b920a19321

Start Time	Duration	Cost	Region	Issues
3:34:53 PM	2,350 ms	< $0.01	us-east-1	0

Show Similar Transactions

GRAPH TIMELINE LOGS

Expand

all the Lambda logs from the transaction in chronological order

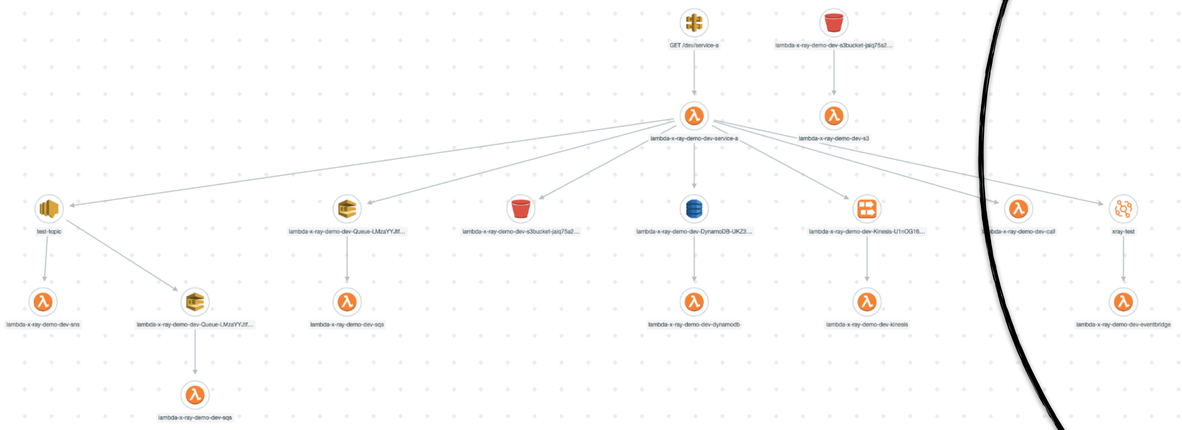

Log Entries (30)

Search in logs

Timestamp	Name	Message
3:34:53 PM...	lambda-x-ray-demo-dev-s...	START RequestId: 74db077a-6...
3:34:54 PM...	lambda-x-ray-demo-dev-s...	INFO Call "eventbridge.putEven...
3:34:54 PM...	lambda-x-ray-demo-dev-s...	INFO Call "s3.putObject" is not ...
3:34:54 PM...	lambda-x-ray-demo-dev-s...	INFO Call "kinesis.putRecord" i...
3:34:54 PM...	lambda-x-ray-demo-dev-sqs	START RequestId: daa6ee23-5...
3:34:54 PM...	lambda-x-ray-demo-dev-call	START RequestId: c4267b0c-8...
3:34:54 PM...	lambda-x-ray-demo-dev-sqs	START RequestId: 4d5ab0f8-a4...
3:34:54 PM...	lambda-x-ray-demo-dev-sqs	END RequestId: daa6ee23-5d6...
3:34:54 PM...	lambda-x-ray-demo-dev-sqs	REPORT RequestId: daa6ee23-...
3:34:54 PM...	lambda-x-ray-demo-dev-sns	START RequestId: 94de0437-f6...
3:34:54 PM...	lambda-x-ray-demo-dev-call	END RequestId: c4267b0c-8ec...
3:34:54 PM...	lambda-x-ray-demo-dev-call	REPORT RequestId: c4267b0c-...
3:34:54 PM...	lambda-x-ray-demo-dev-s...	END RequestId: 74db077a-6b6...

Issues
0

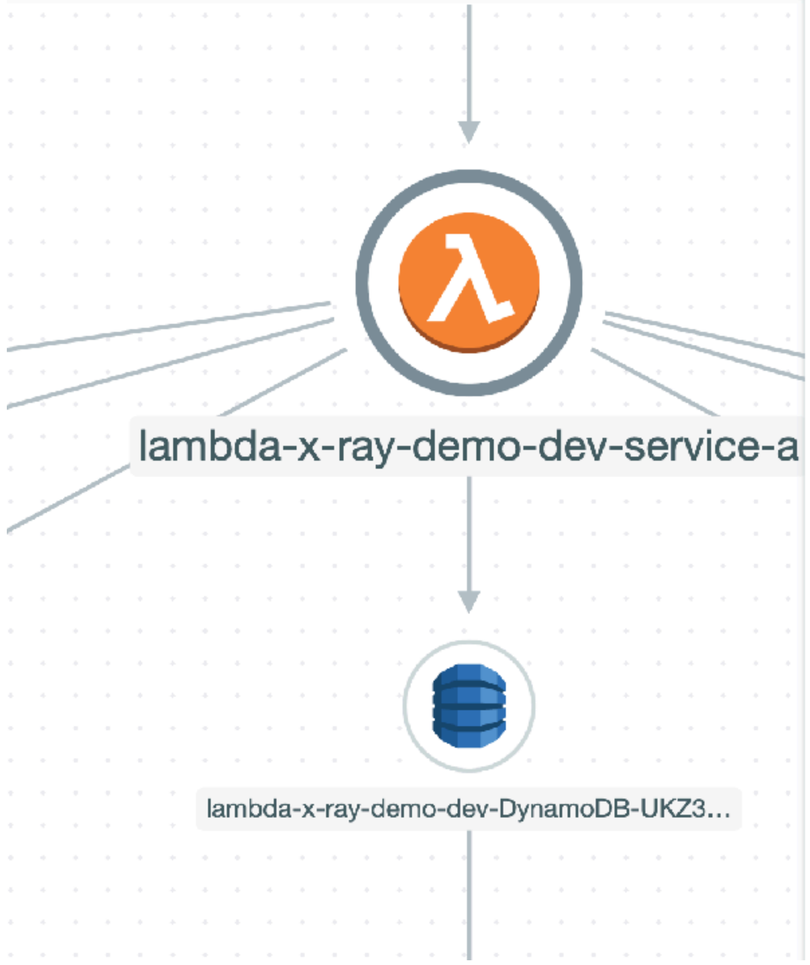

lambda-x-ray-demo-dev-service-a

us-east-1 | nodejs14.x

Request id	Lumigo Tag	Memory Limit
74db077a-6b64-404c-...	Not Set	1,024MB

Start Time	Duration	Cold Start
3:34:53 PM	762 ms	Yes

Execution Tags | Return Value | Event | Environment Variables | Log Entries

Execution Tags

path /service-a | method GET | apigw-request-id ae24e4f1-0cde-46a0-b499-d5886625f0d2

Return Value

Body


```
{
  statusCode: 200
  body: "{"message":"done"}"
}
```


Event

Body


```
{
  resource: "/service-a"
  path: "/service-a"
  httpMethod: "GET"
▶ requestContext: { 1 prop }
▶ headers: { 3 props }
  stageVariables: null
  isBase64Encoded: false
}
```

b6b64b920a19321

Cost	Region	Issues
< $0.01	us-east-1	0

LOGS

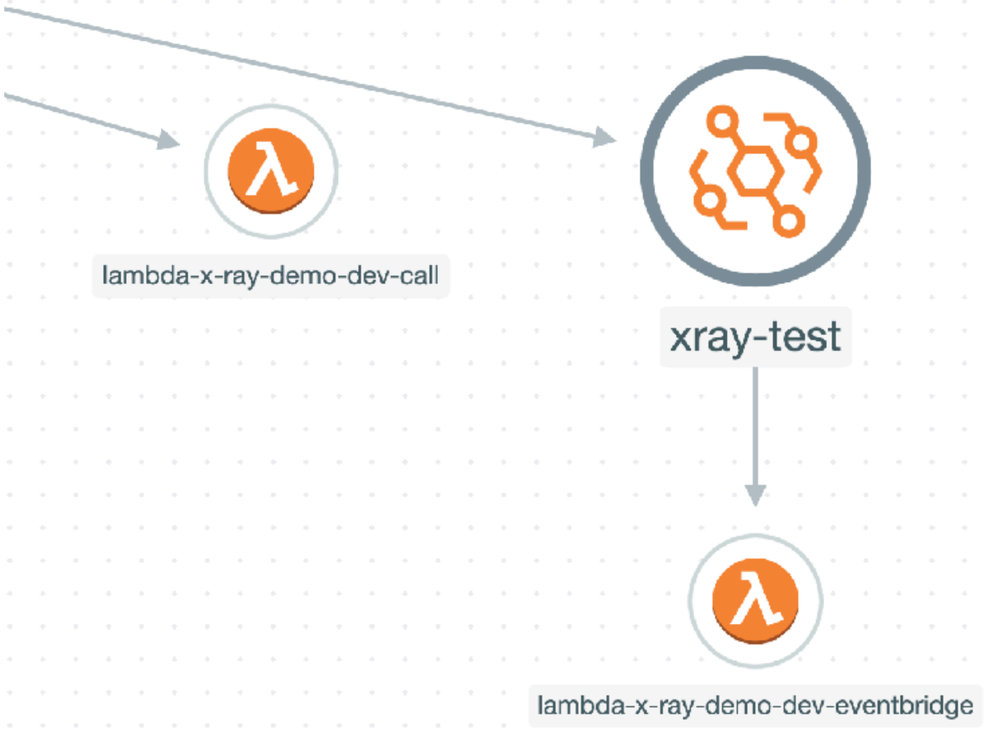

xray-test us-east-1

Start Time	Duration	HTTP Method	URI	HTTP Status Code
3:34:53 PM	65 ms	POST	events.us-east-1.amaz...	200

Request | Response

Request

Copy request as cURL

Body


```
{
▶ Entries: [ 1 item ]
}
```


Headers


```
{
  user-agent: "aws-sdk-nodejs/2.1055.0 linux/v14.20.0 exec-env/AWS_Lambda_nodejs14.x
              promise"
  content-type: "application/x-amz-json-1.1"
  x-amz-target: "AWSEvents.PutEvents"
  x-amzn-trace-id: "Root=1-630f637d-
                   1c6eabf86b6b64b920a19321;Parent=4068306606aa9ec1;Sampled=1"
  x-amz-content-sha256: "f93cab3226f9f15320d381775c14d30dc5c8386b750e7601e3c4e9d844e55e
                        66"
  content-length: 120
  host: "events.us-east-1.amazonaws.com"
  x-amz-date: "20220831T133453Z"
  x-amz-security-token: Hidden Information
  authorization: Hidden Information
}
```

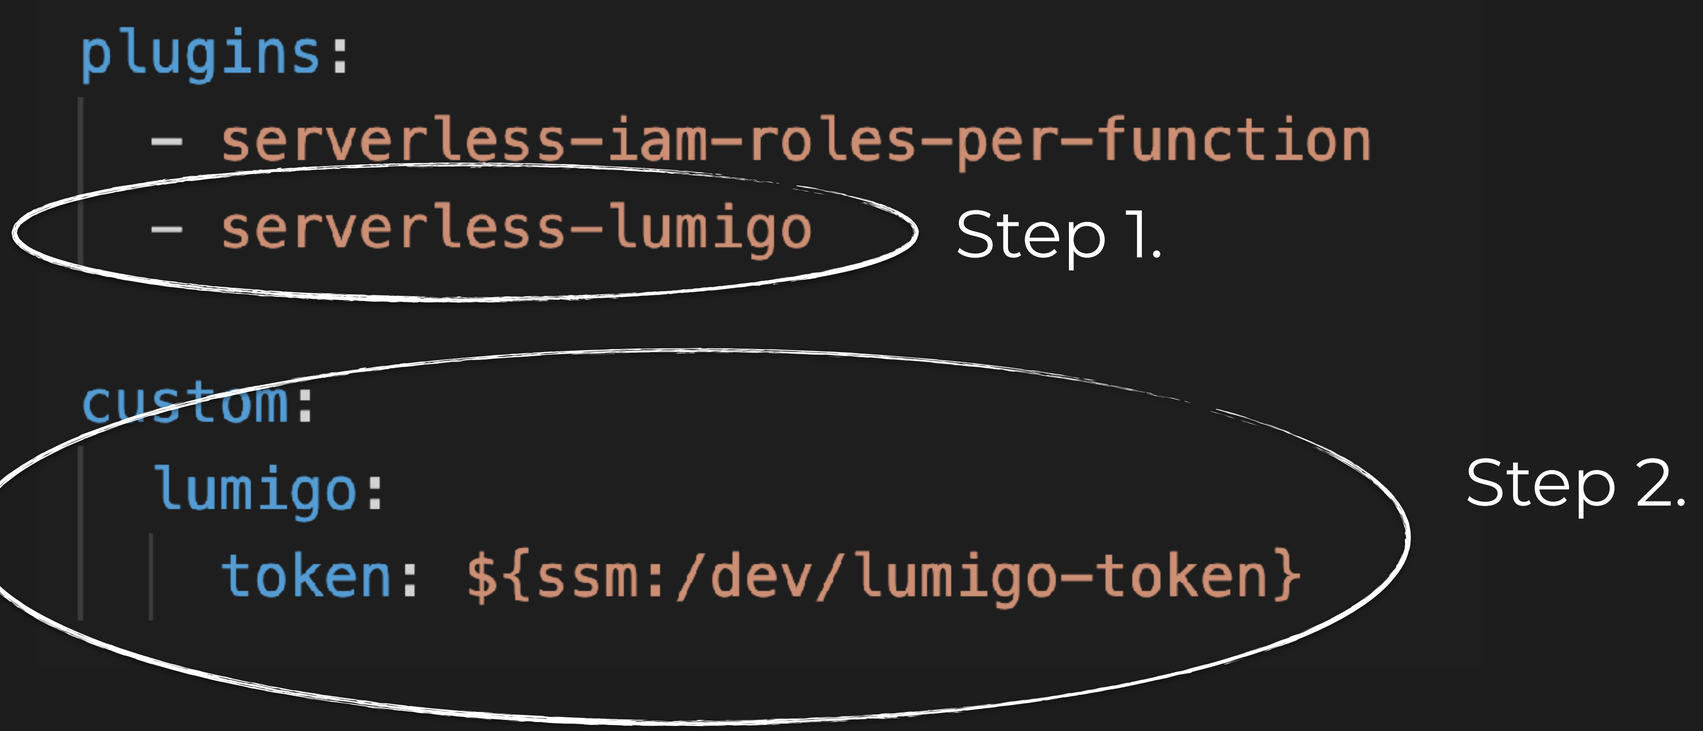
plugins:
- serverless-iam-roles-per-function
- serverless-lumigo
Step 1.
custom:
lumigo:
token: ${ssm:/dev/lumigo-token}
Step 2.

No manual instrument.

Better support for messaging services.

Supports containers and Lambda.

“which messaging service should I use?”

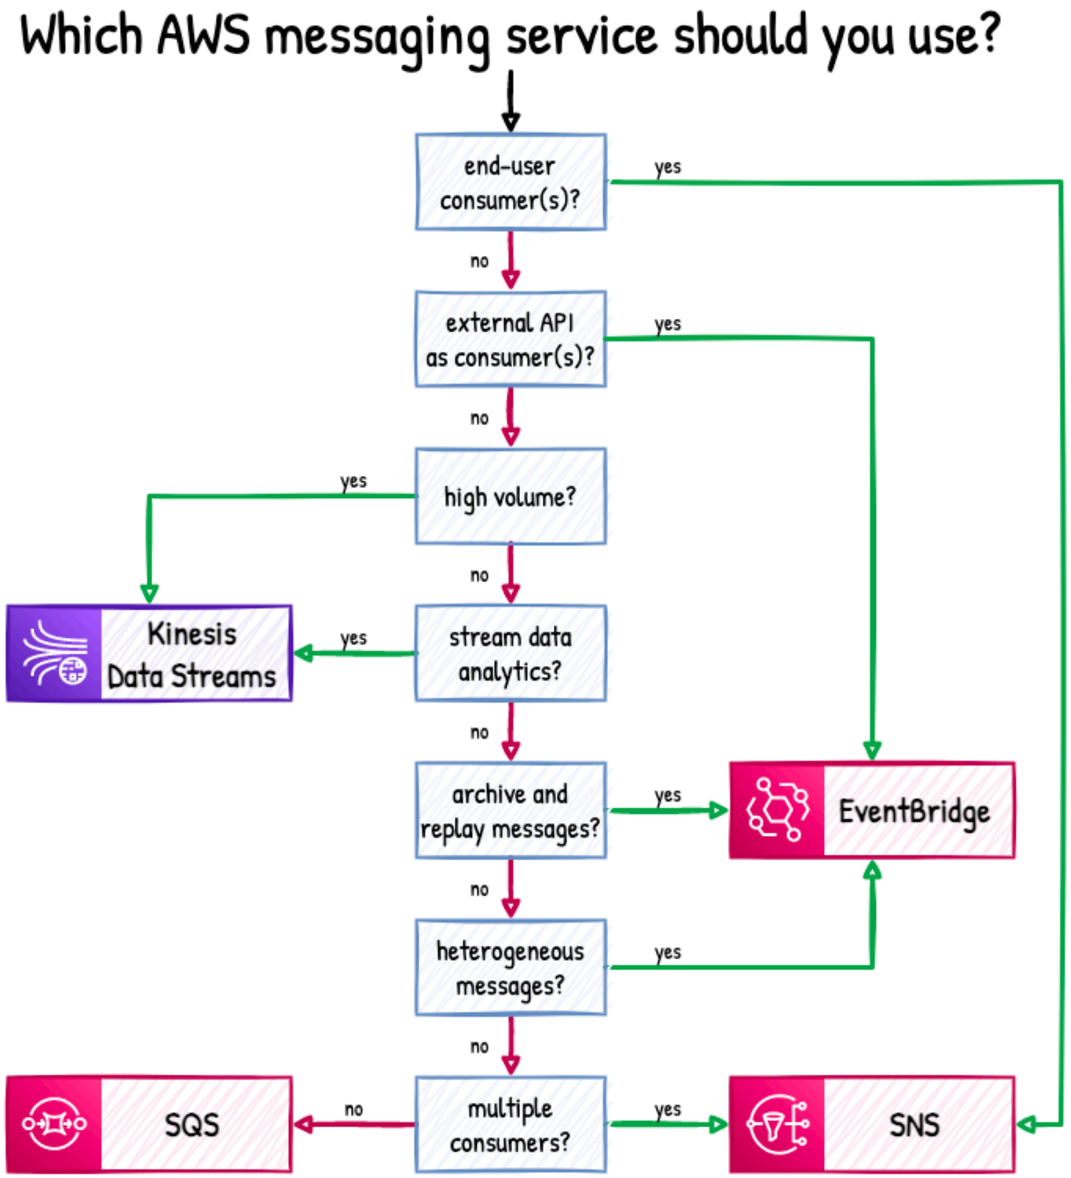

Maciej Radzikowski
@radzikowski_m

A handy Decision Tree for choosing the right messaging service on AWS.

As per my calculations, following it gives you a 90% chance of making the right choice.

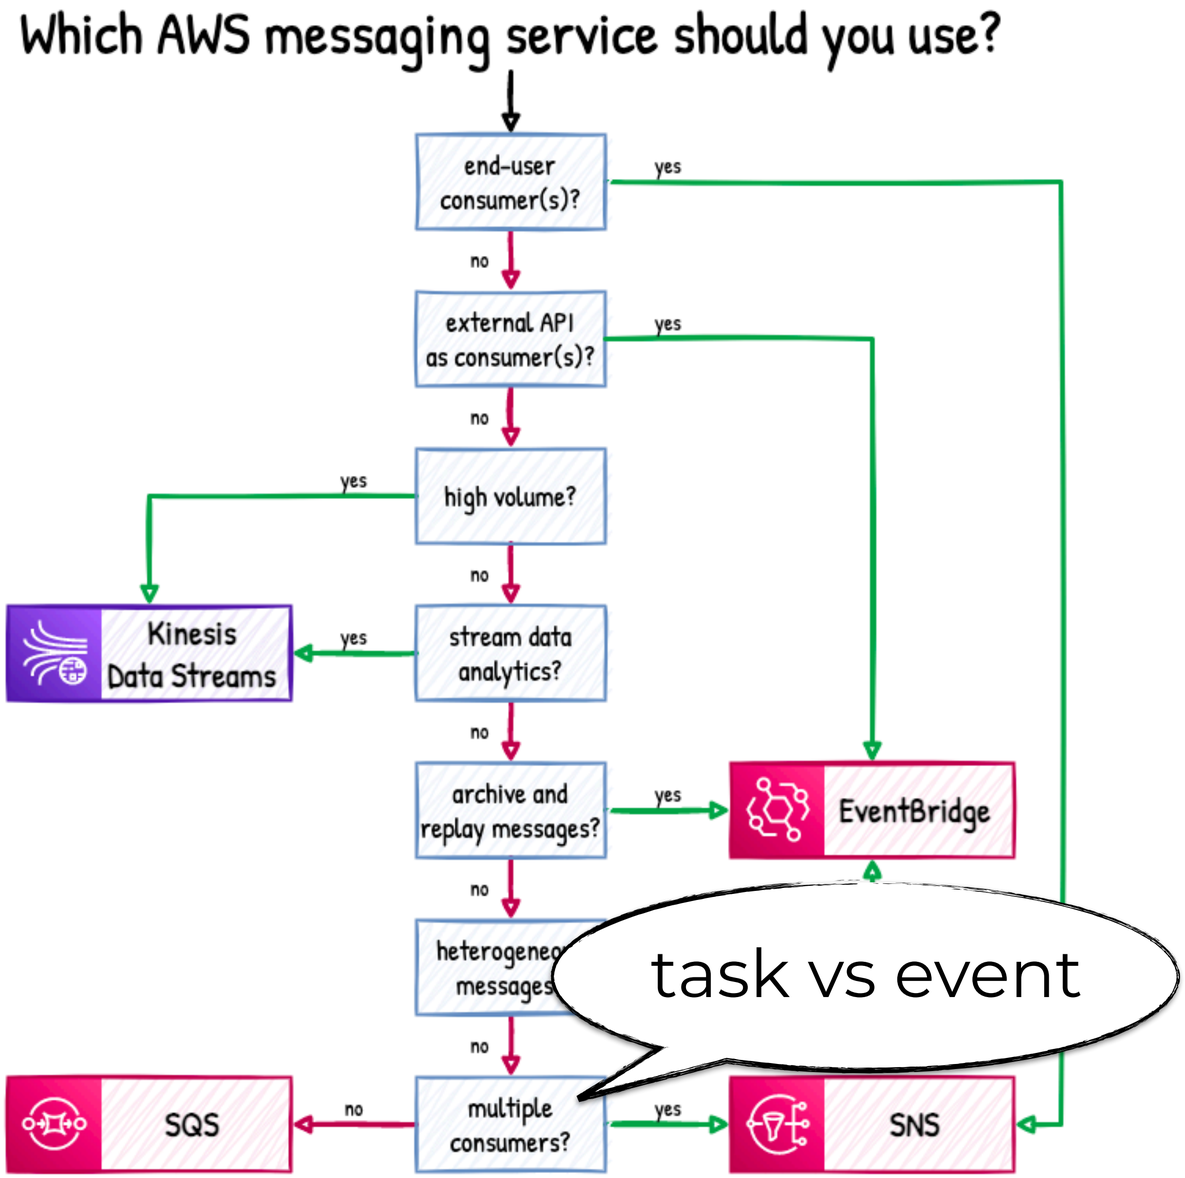

Maciej Radzikowski
@radzikowski_m

A handy Decision Tree for choosing the right messaging service on AWS.

As per my calculations, following it gives you a 90% chance of making the right choice.

Task

“Go do a thing”

Event

“Something happened”

Task

“Go do a thing”

Intended for a target receiver.

Event

“Something happened”

Task

“Go do a thing”

Intended for a target receiver.

Often expects an answer.

Event

“Something happened”

Task

“Go do a thing”

Intended for a target receiver.

Often expects an answer.

Event

“Something happened”

Publishers are obvious to subscribers.

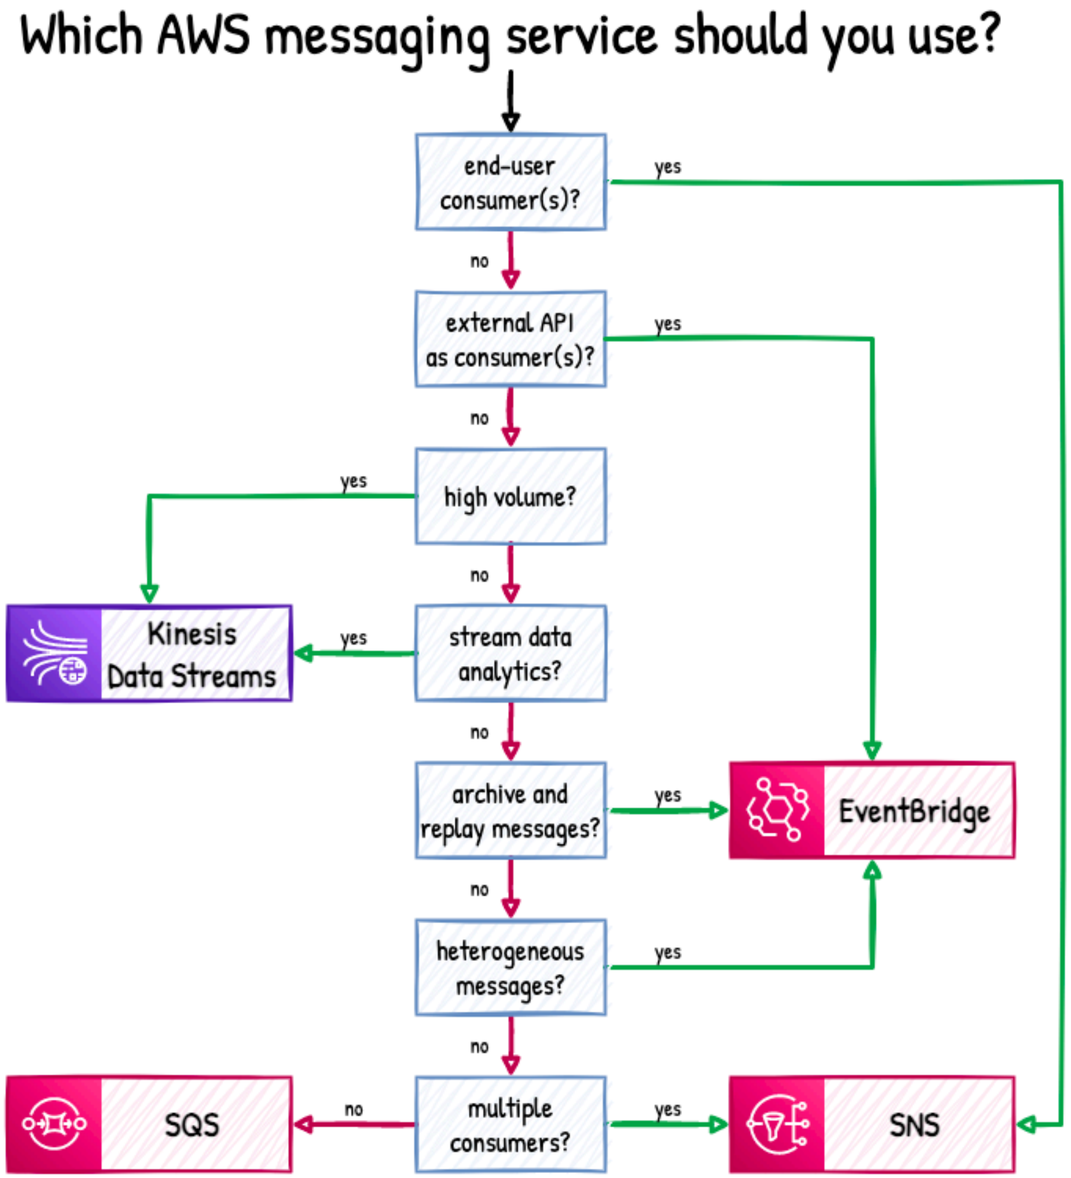

Maciej Radzikowski
@radzikowski_m

A handy Decision Tree for choosing the right messaging service on AWS.

As per my calculations, following it gives you a 90% chance of making the right choice.

	Kinesis Provisioned	Kinesis On-Demand	SQS	SQS FIFO	SNS	SNS FIFO	EventBridge
Subscribers	one-to-many	one-to-many	one-to-one	one-to-one	one-to-many	one-to-many	one-to-many
Ordering	by shard	by shard	none	by group ID	none	by group ID	none
Archive	up to 1 year	up to 1 year	none	none	none	none	up to indefinite
Replay	yes	yes	no	no	no	no	yes
Batching	yes	yes	yes	yes	no	no	no
Retry Behaviour	configurable	configurable	configurable	configurable	up to 2 retries	up to 2 retries	up to 2 retries
DLQ	yes (no body)	yes (no body)	yes (redrive policy)	yes (redrive policy)	yes	yes	yes
Concurrency	1-10 per shard	1-10 per shard	5 + 60/min (max 1000)	5 + 60/min (max 1000)	fan-out	fan-out	fan-out
Cost	shard hour + PUTs	stream hour + data in/out	per request	per request	per request	per request	per request

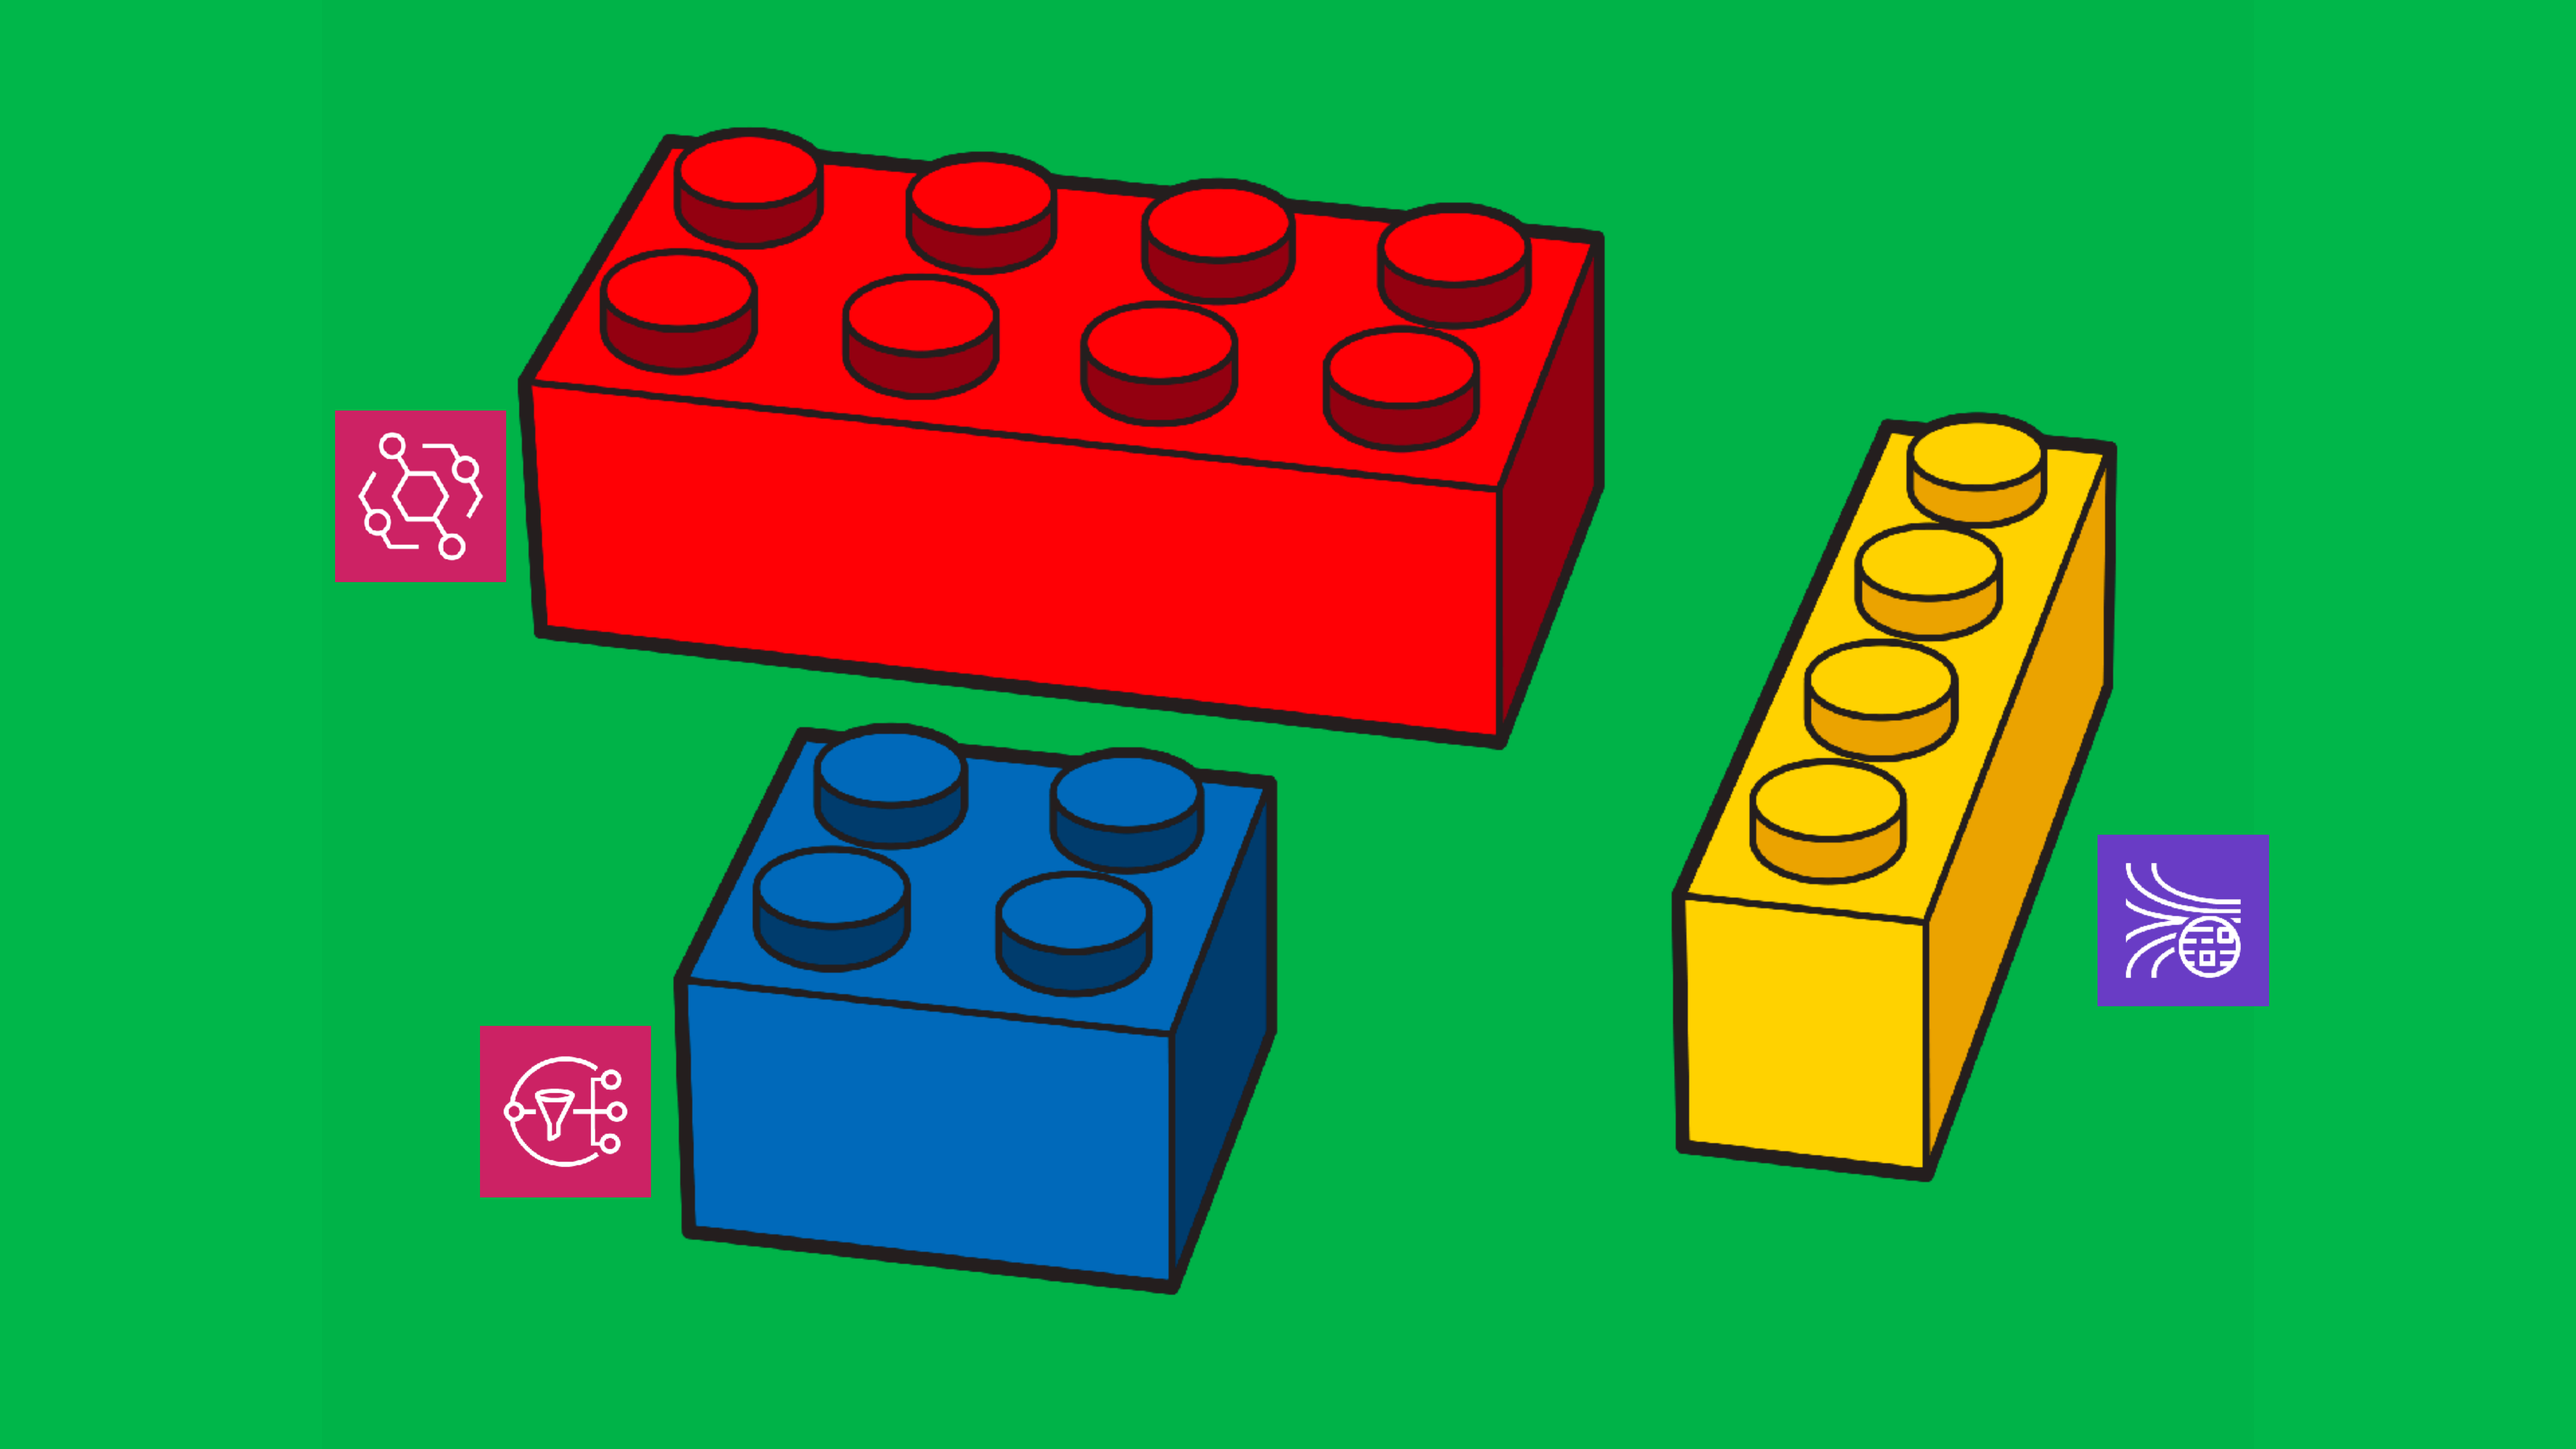

	Kinesis Provisioned	Kinesis On-Demand	SQS	SQS FIFO	SNS	SNS FIFO	EventBridge
Subscribers	one-to-many	one-to-many	one-to-one	one-to-one	one-to-many	one-to-many	one-to-many
Ordering	by shard	by shard	none	by group ID	none	by group ID	none
Archive	up to 1 year	up to 1 year	none	none	none	none	up to indefinite
Replay	yes	yes	no	no	no	no	yes
Batching	yes	yes	yes	yes	no	no	no
Retry Behaviour	configurable	configurable	configurable	configurable	up to 2 retries	up to 2 retries	up to 2 retries
DLQ	yes (no body)	yes (no body)	yes (redrive policy)	yes (redrive policy)	yes	yes	yes
Concurrency	1-10 per shard	1-10 per shard	5 + 60/min (max 1000)	5 + 60/min (max 1000)	fan-out	fan-out	fan-out
Cost	shard hour + PUTs	stream hour + data in/out	per request	per request	per request	per request	per request

	Kinesis Provisioned	Kinesis On-Demand	SQS	SQS FIFO	SNS	SNS FIFO	EventBridge
Subscribers	one-to-many	one-to-many	one-to-one	one-to-one	one-to-many	one-to-many	one-to-many
Ordering	by shard	by shard	none	by group ID	none	by group ID	none
Archive	up to 1 year	up to 1 year	none	none	none	none	up to indefinite
Replay	yes	yes	no	no	no	no	yes
Batching	yes	yes	yes	yes	no	no	no
Retry Behaviour	configurable	configurable	configurable	configurable	up to 2 retries	up to 2 retries	up to 2 retries
DLQ	yes (no body)	yes (no body)	yes (redrive policy)	yes (redrive policy)	yes	yes	yes
Concurrency	1-10 per shard	1-10 per shard	5 + 60/min (max 1000)	5 + 60/min (max 1000)	fan-out	fan-out	fan-out
Cost	shard hour + PUTs	stream hour + data in/out	per request	per request	per request	per request	per request

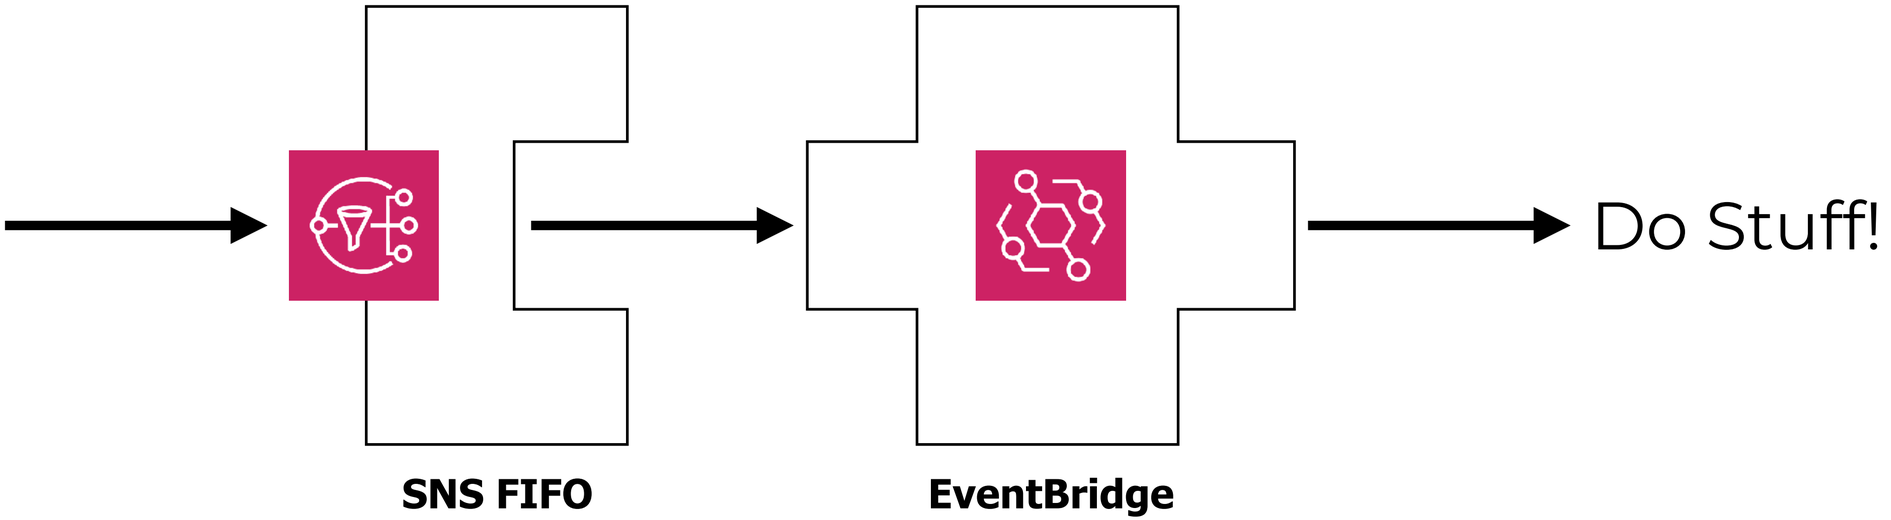
Do Stuff!
SNS FIFO
EventBridge

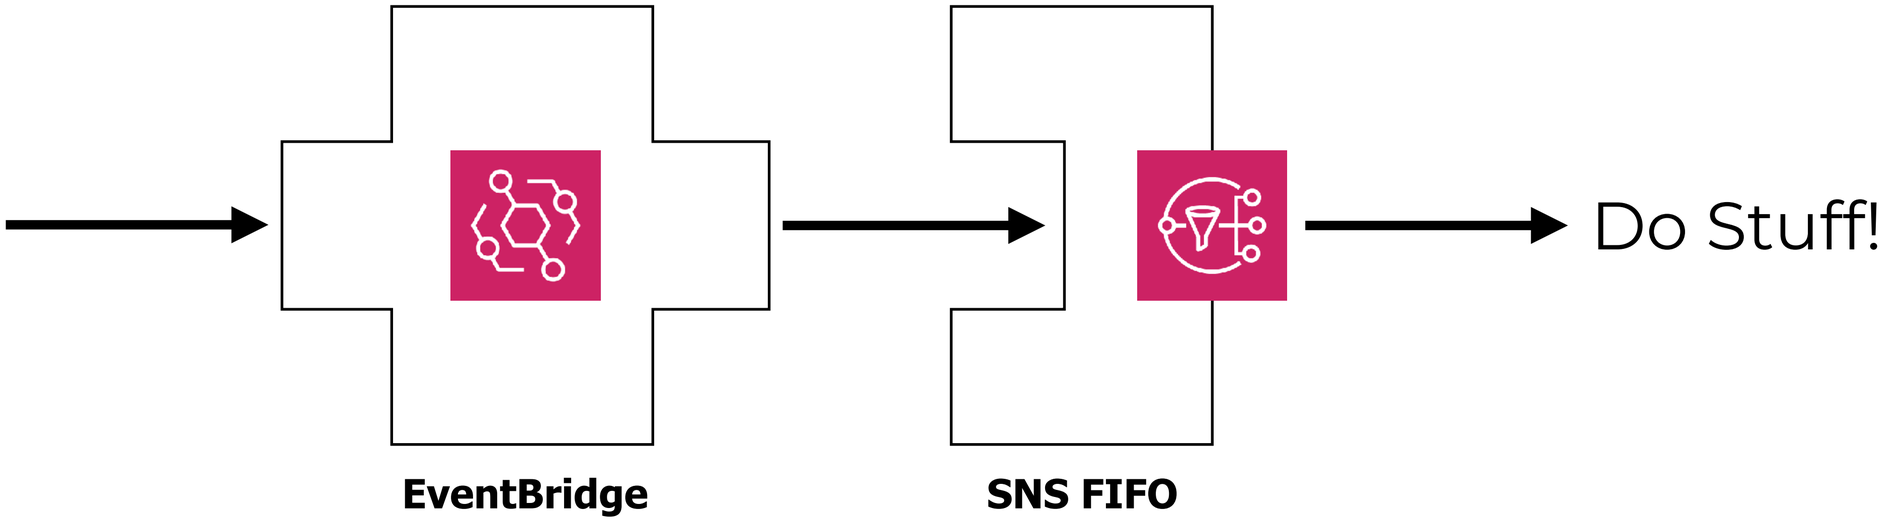
Do Stuff!
EventBridge
SNS FIFO

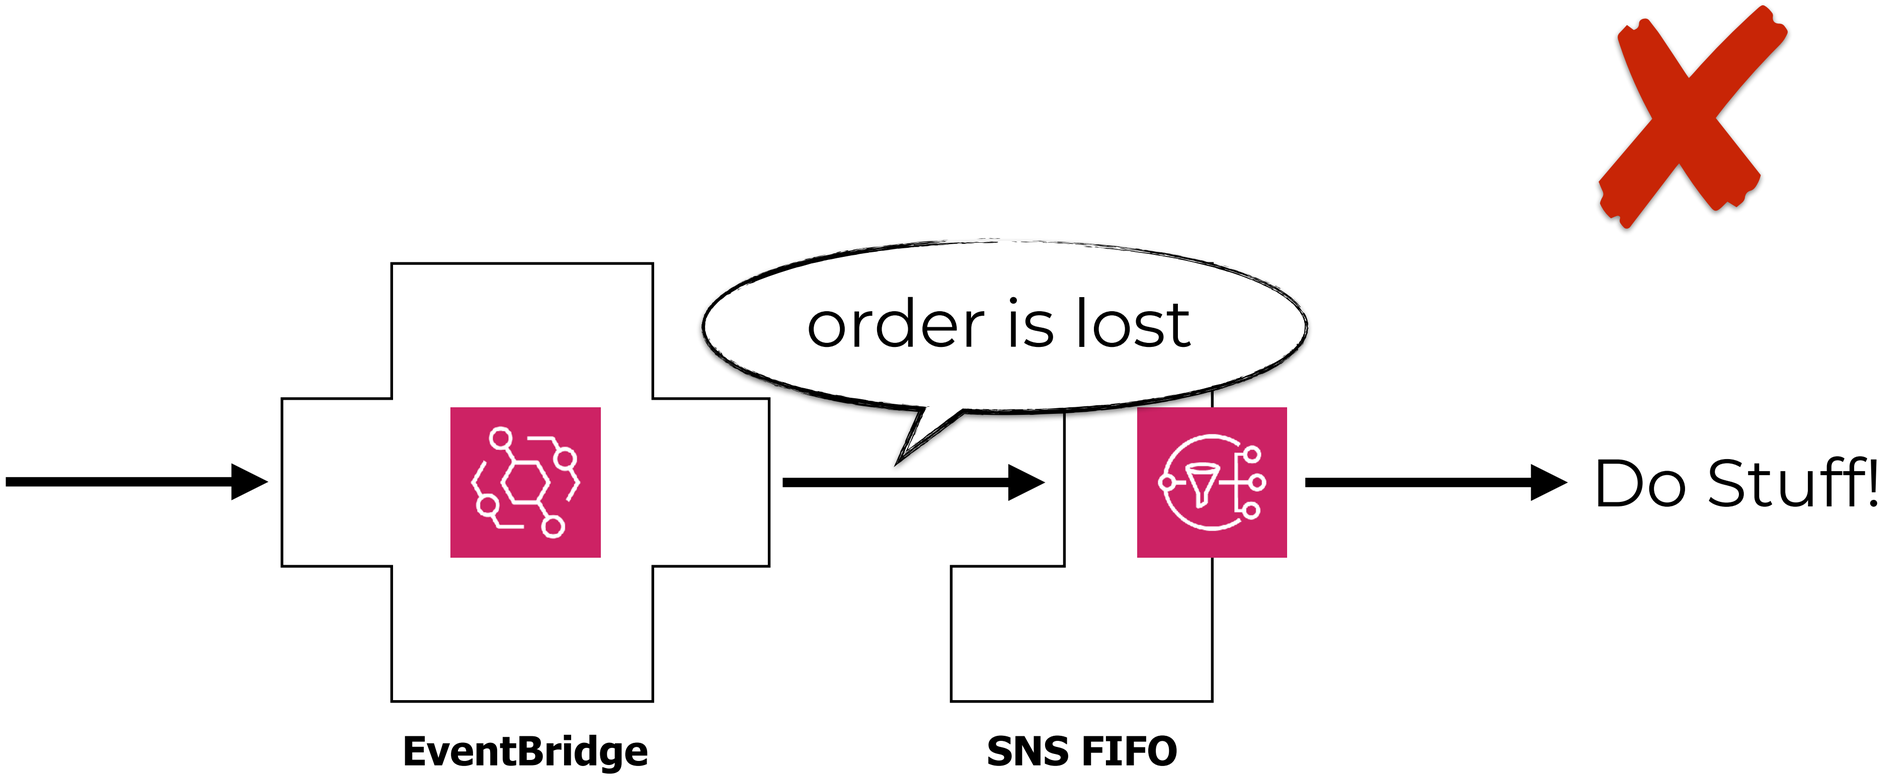
order is lost
Do Stuff!
EventBridge
SNS FIFO

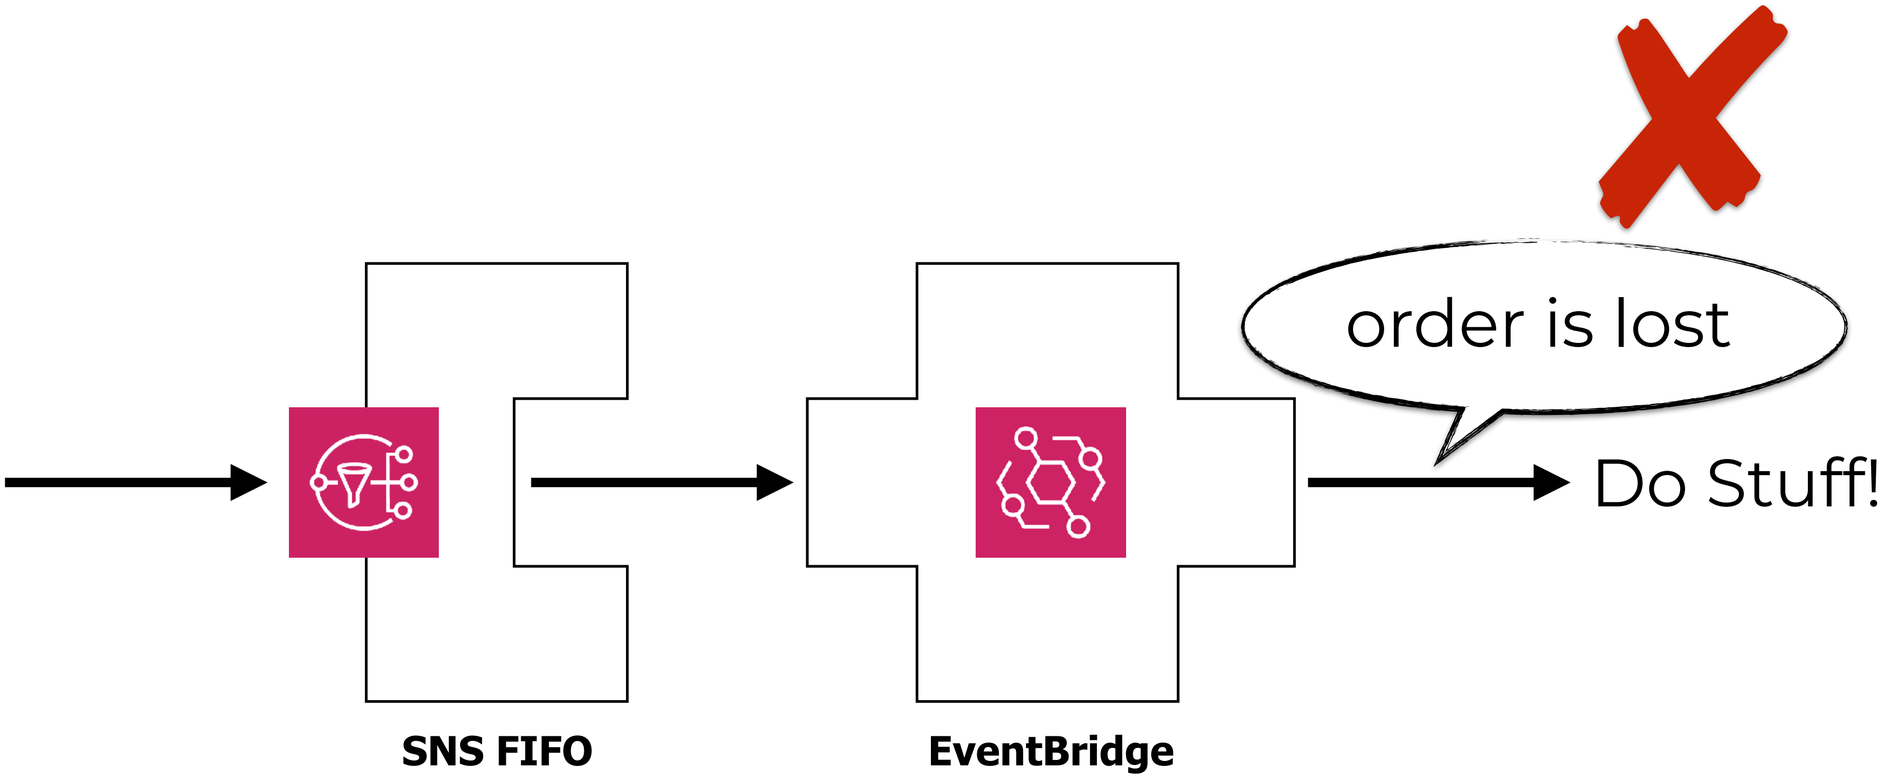
order is lost
Do Stuff!
SNS FIFO
EventBridge

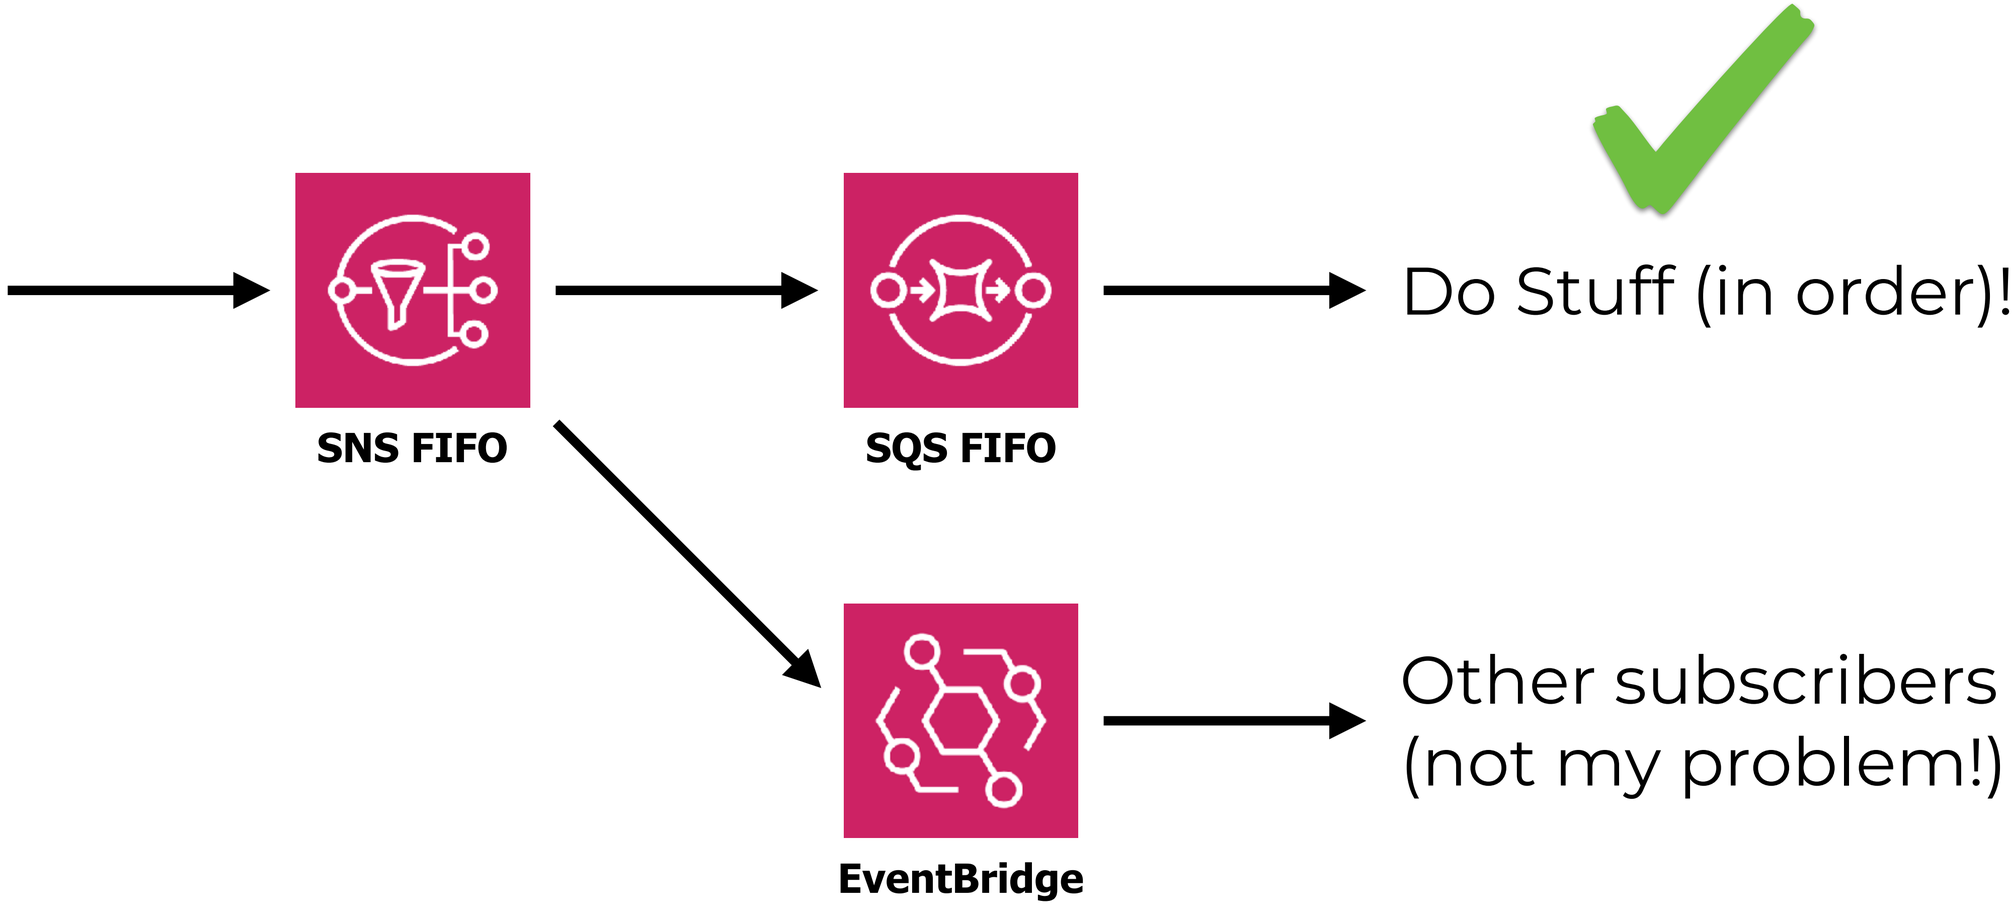
SNS FIFO
SQS FIFO
Do Stuff (in order)!
EventBridge
Other subscribers
(not my problem!)

	Kinesis Provisioned	Kinesis On-Demand	SQS	SQS FIFO	SNS	SNS FIFO	EventBridge
Subscribers	one-to-many	one-to-many	one-to-one	one-to-one	one-to-many	one-to-many	one-to-many
Ordering	by shard	by shard	none	by group ID	none	by group ID	none
Archive	up to 1 year	up to 1 year	none	none	none	none	up to indefinite
Replay	yes	yes	no	no	no	no	yes
Batching	yes	yes	yes	yes	no	no	no
Retry Behaviour	configurable	configurable	configurable	configurable	up to 2 retries	up to 2 retries	up to 2 retries
DLQ	yes (no body)	yes (no body)	yes (redrive policy)	yes (redrive policy)	yes	yes	yes
Concurrency	1-10 per shard	1-10 per shard	5 + 60/min (max 1000)	5 + 60/min (max 1000)	fan-out	fan-out	fan-out
Cost	shard hour + PUTs	stream hour + data in/out	per request	per request	per request	per request	per request

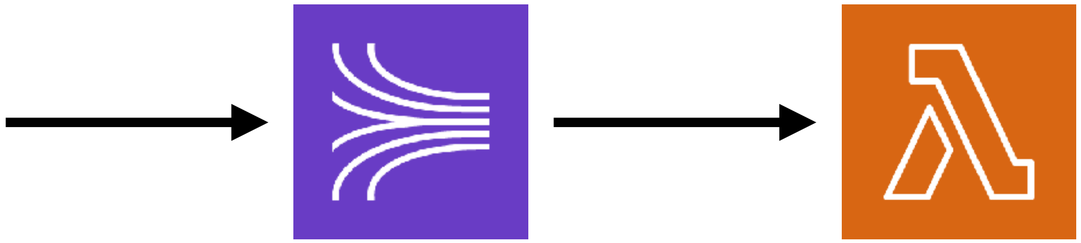
Kinesis
Lambda

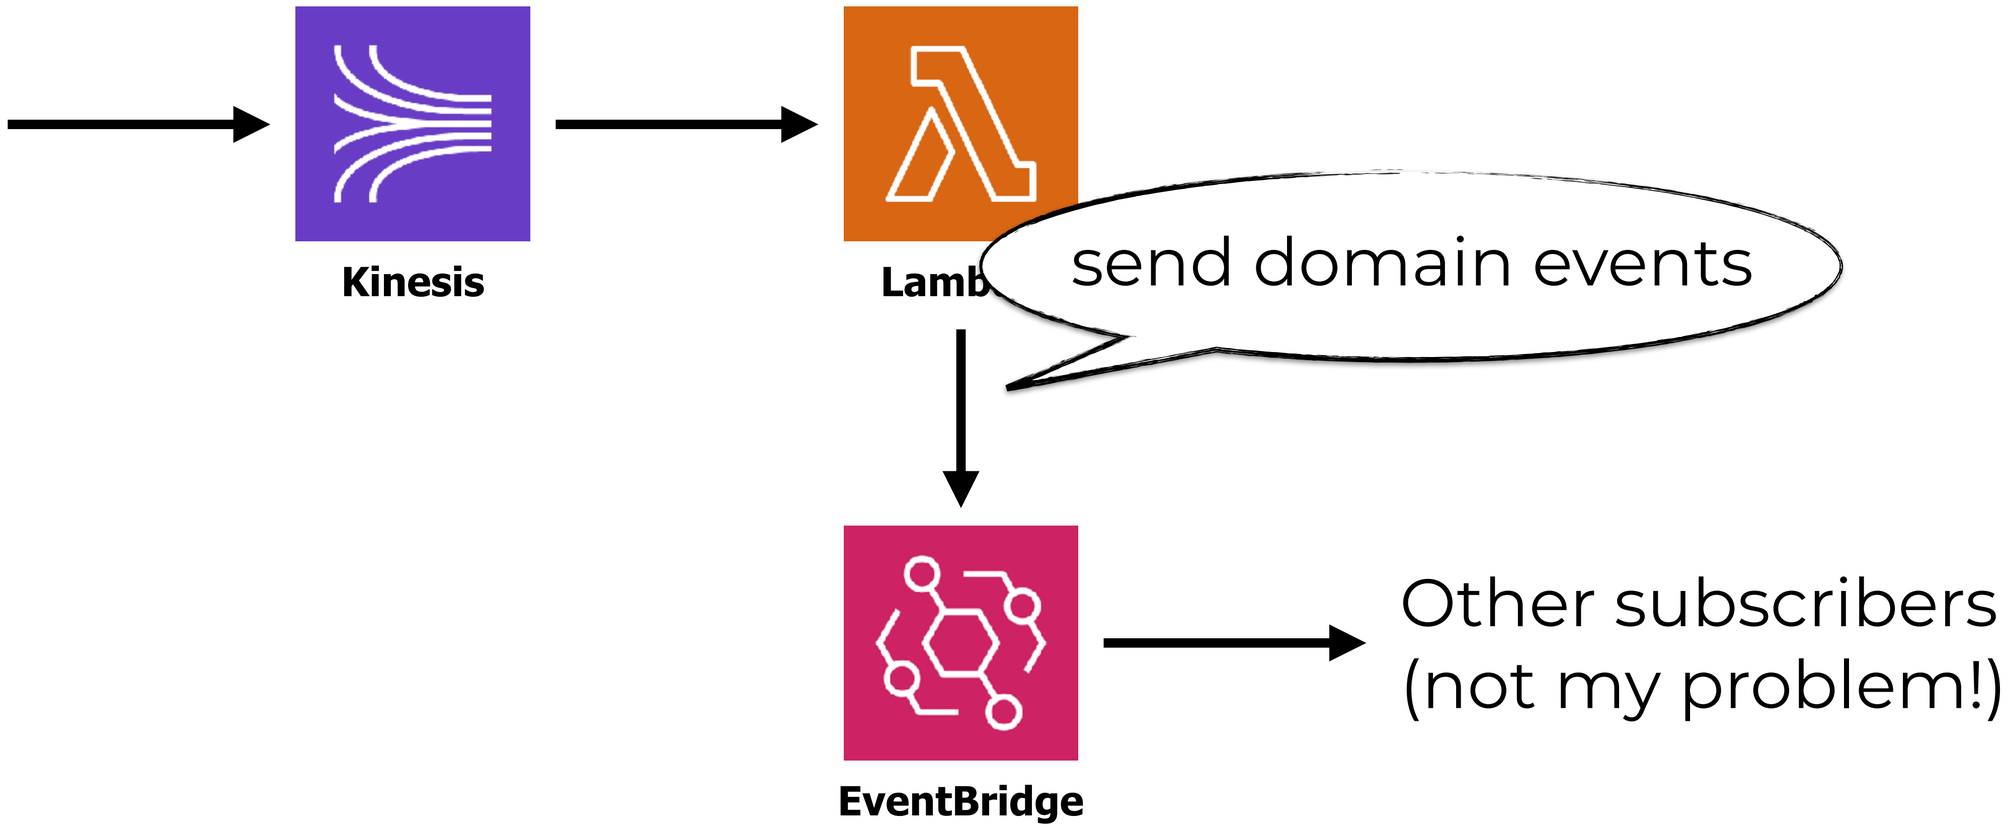
Kinesis
Lamb
send domain events
EventBridge
Other subscribers
(not my problem!)

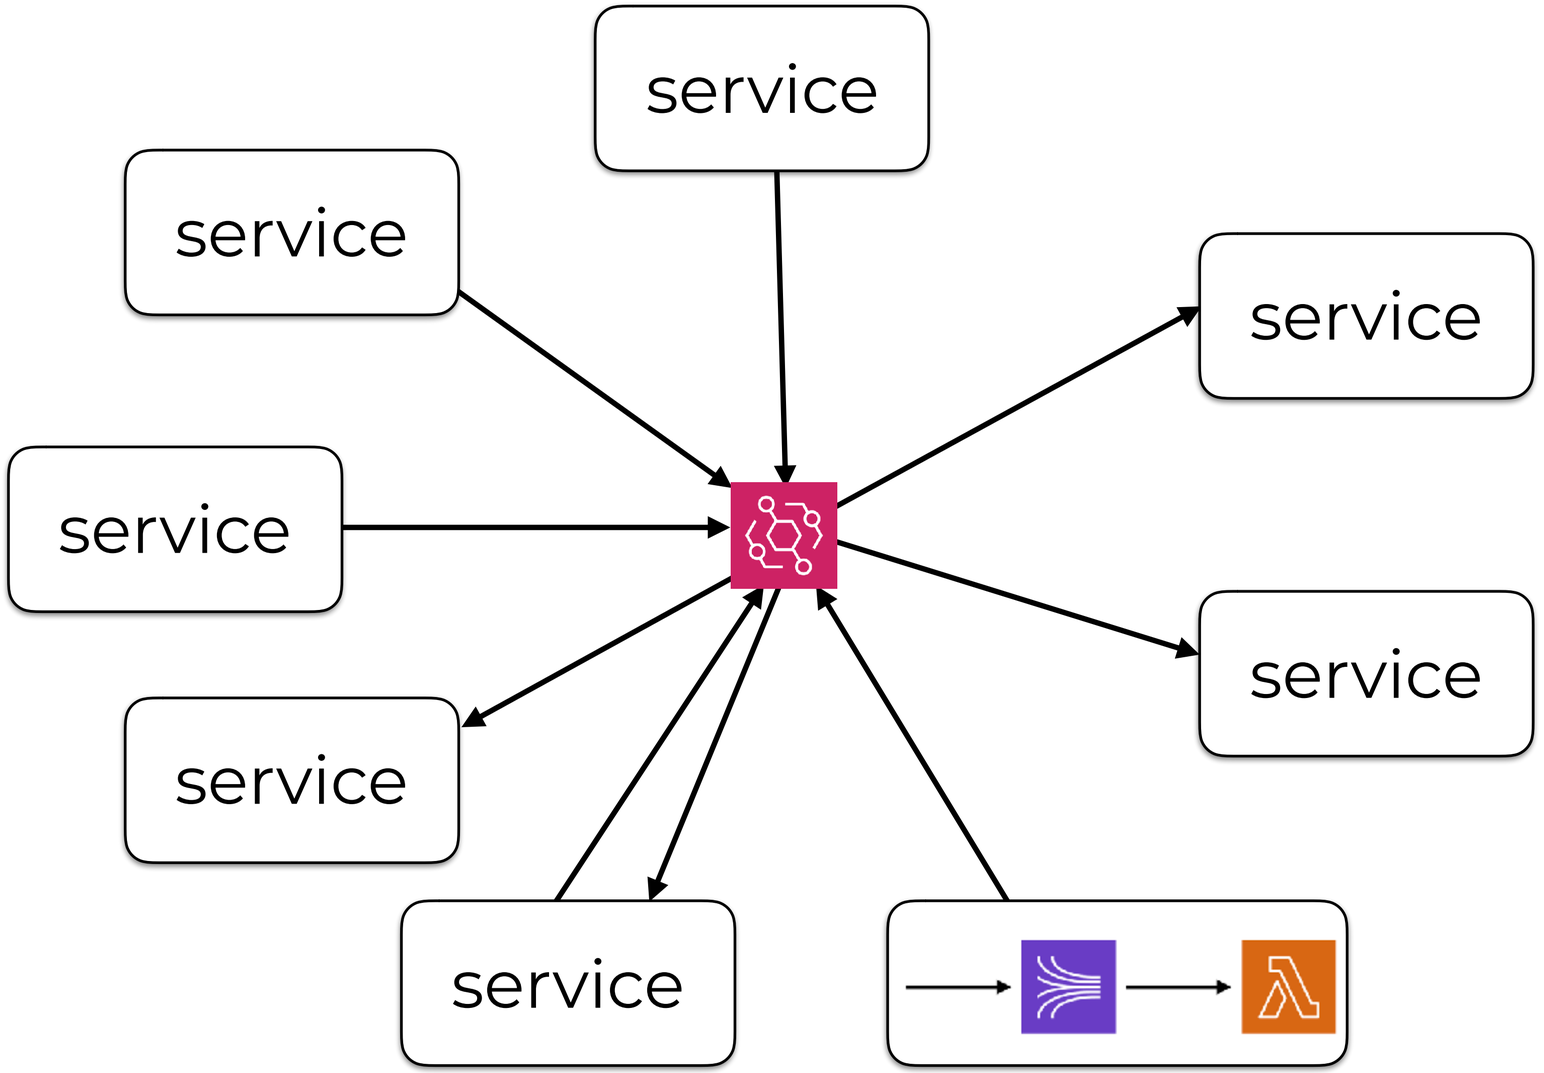
service
service
service
service
service
service
service
service

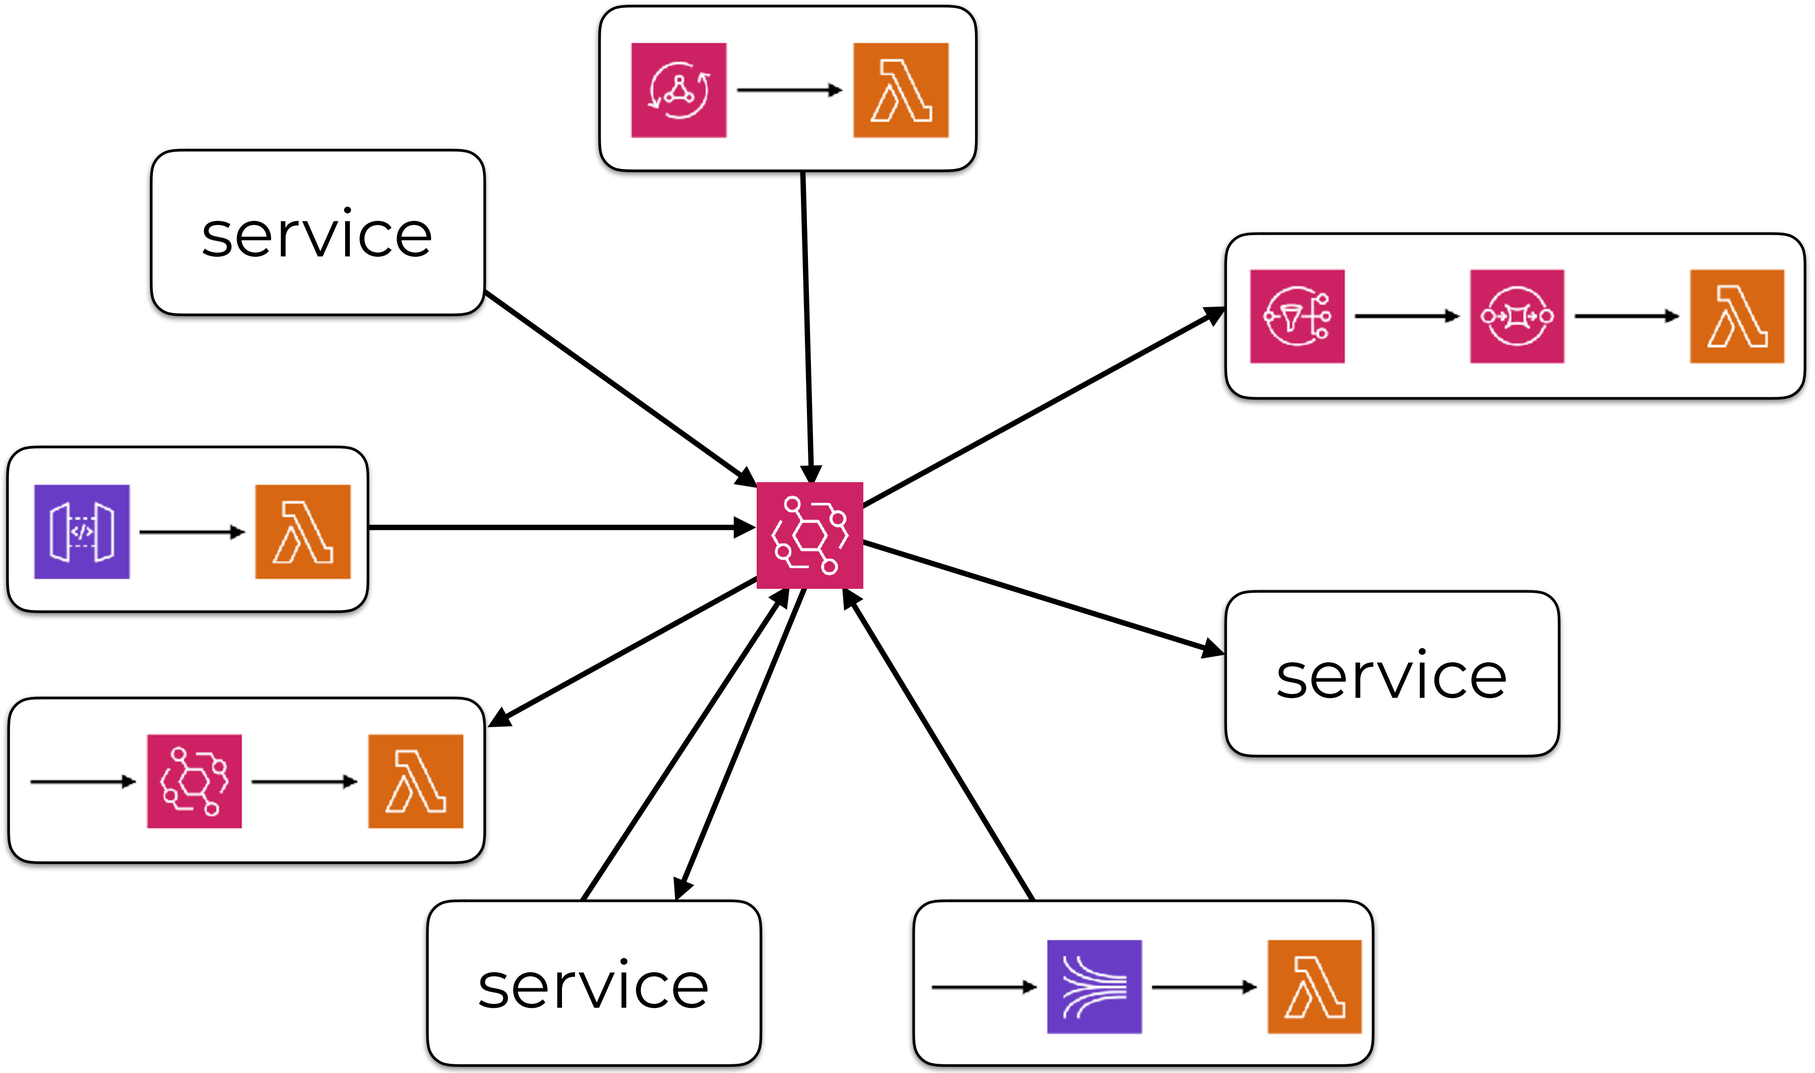
service
service
service

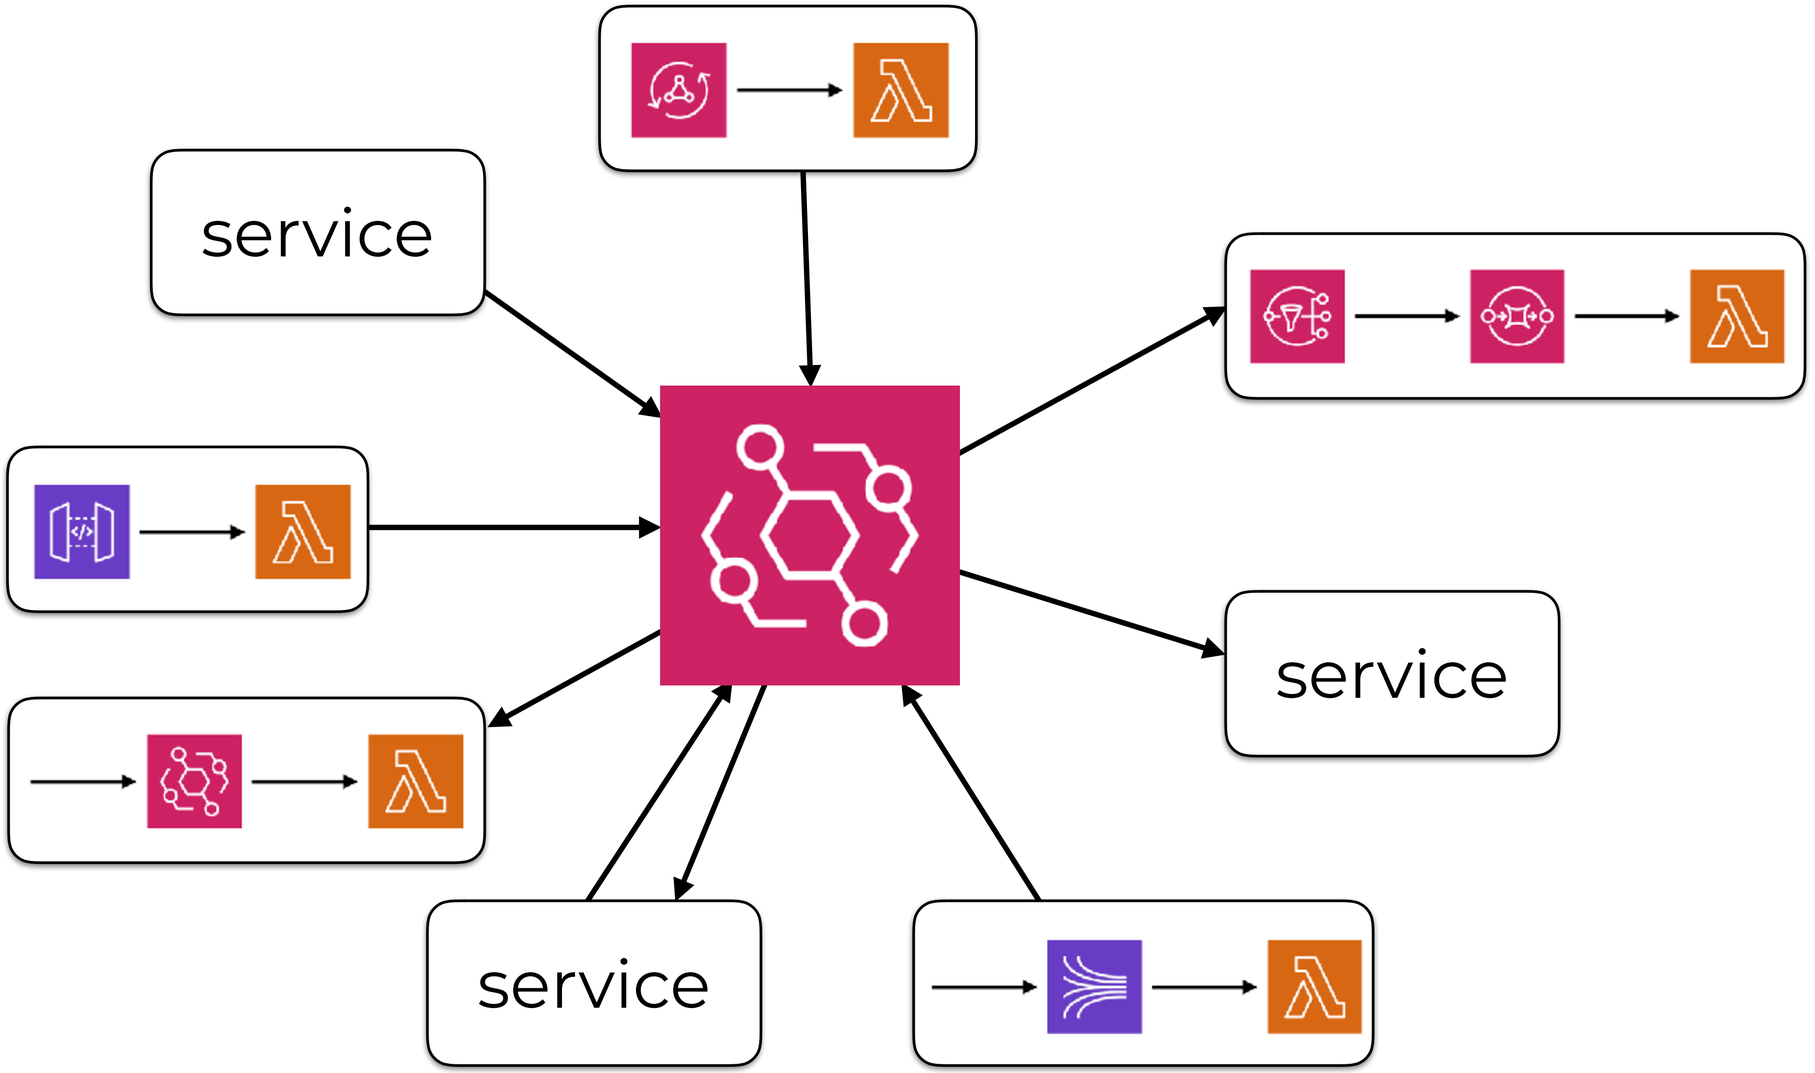
service
service
service


"That's all Folks!"

Training **Advise** **Delivery**

https://theburningmonk.com/hire-me

Nick Blair

Tech Lead TotallyMoney

"Fundamentally, Yan has improved our team by increasing our ability to derive value from AWS and Lambda in particular."

@theburningmonk

theburningmonk.com

github.com/theburningmonk