AWS re:Invent

API312

How to select the right application-integration service

Christian Müller
Solutions Architect, Amazon Web Services
cmr@amazon.de | @ChristianM

Dirk Fröhner
Solutions Architect, Amazon Web Services
froehner@amazon.de | @dirk_f5r

AWS re:Invent

aws

Agenda

Introduction – bringing everybody onto the same page

Open discussion – what are your app-int challenges?

Related breakouts

ARC314 – Decoupled microservices: Building scalable applications

API304 – Scalable serverless event-driven apps using Amazon SQS & Lambda

API306 – Building event-driven architectures

API307 – Build efficient and scalable distributed applications using Amazon MQ

API309 – Durable serverless architecture: Working with dead-letter queues

API311 – Managing business processes using AWS Step Functions

API315 – Application integration patterns for microservices

API316 – Building serverless workflows using AWS Step Functions

API318 – Deep dive on event-driven development with Amazon EventBridge

Introduction

AWS re:Invent

“If your application is **cloud-native**, or **large-scale**, or **distributed**, and **doesn’t include** a **messaging** component, that’s probably a **bug**.”

Tim Bray

Distinguished Engineer

AWS Messaging, Workflow Management

AWS re:Invent

aws

Messaging vs streaming

Message processing vs stream processing

Message processing

Each individual message is unit of work
Computation / processing per message
Message occurrence can vary

DLQ functionality built-in
Messages deleted after consumption
No need to track position of a message

Stream processing

Message stream is unit of work
Complex computation on many messages
Constant stream of messages

No built-in DLQ functionality
Messages available after consumption until expiration
Each client must track current position in the stream

Integration services - overview

Integration services

IOT

AWS IoT Core

Messaging

Amazon SNS

Amazon SQS

Managed

Amazon MQ

Amazon MSK

Amazon ElastiCache (Redis)

Streaming

Amazon Kinesis

Events

Amazon CloudWatch

Amazon EventBridge

Which service should I recommend?

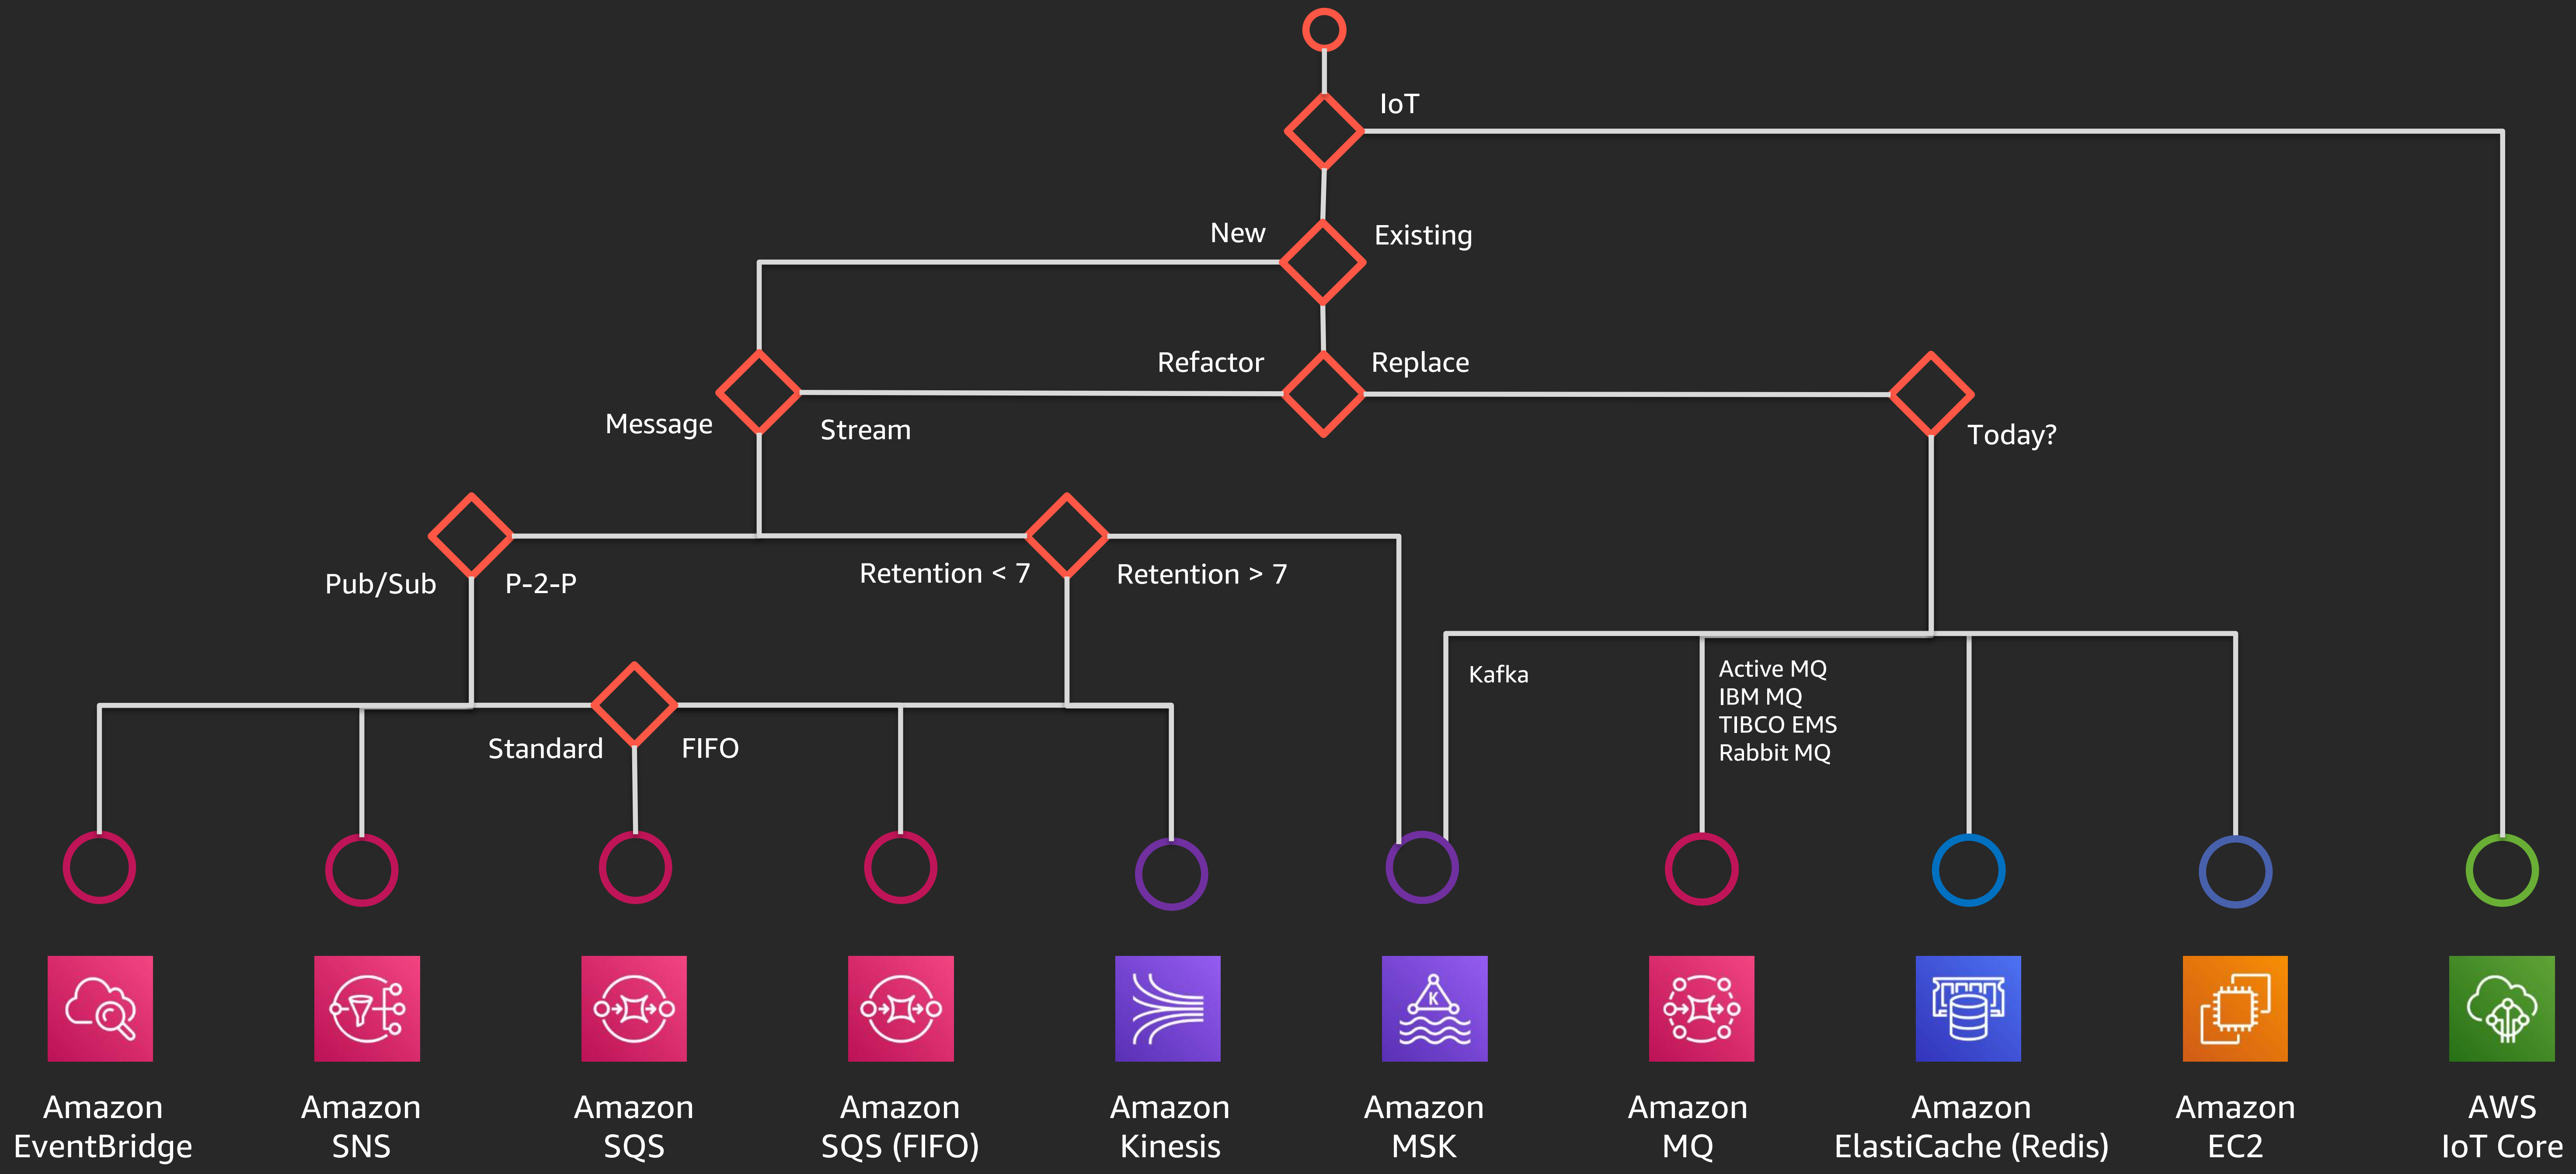

Which service should I recommend?

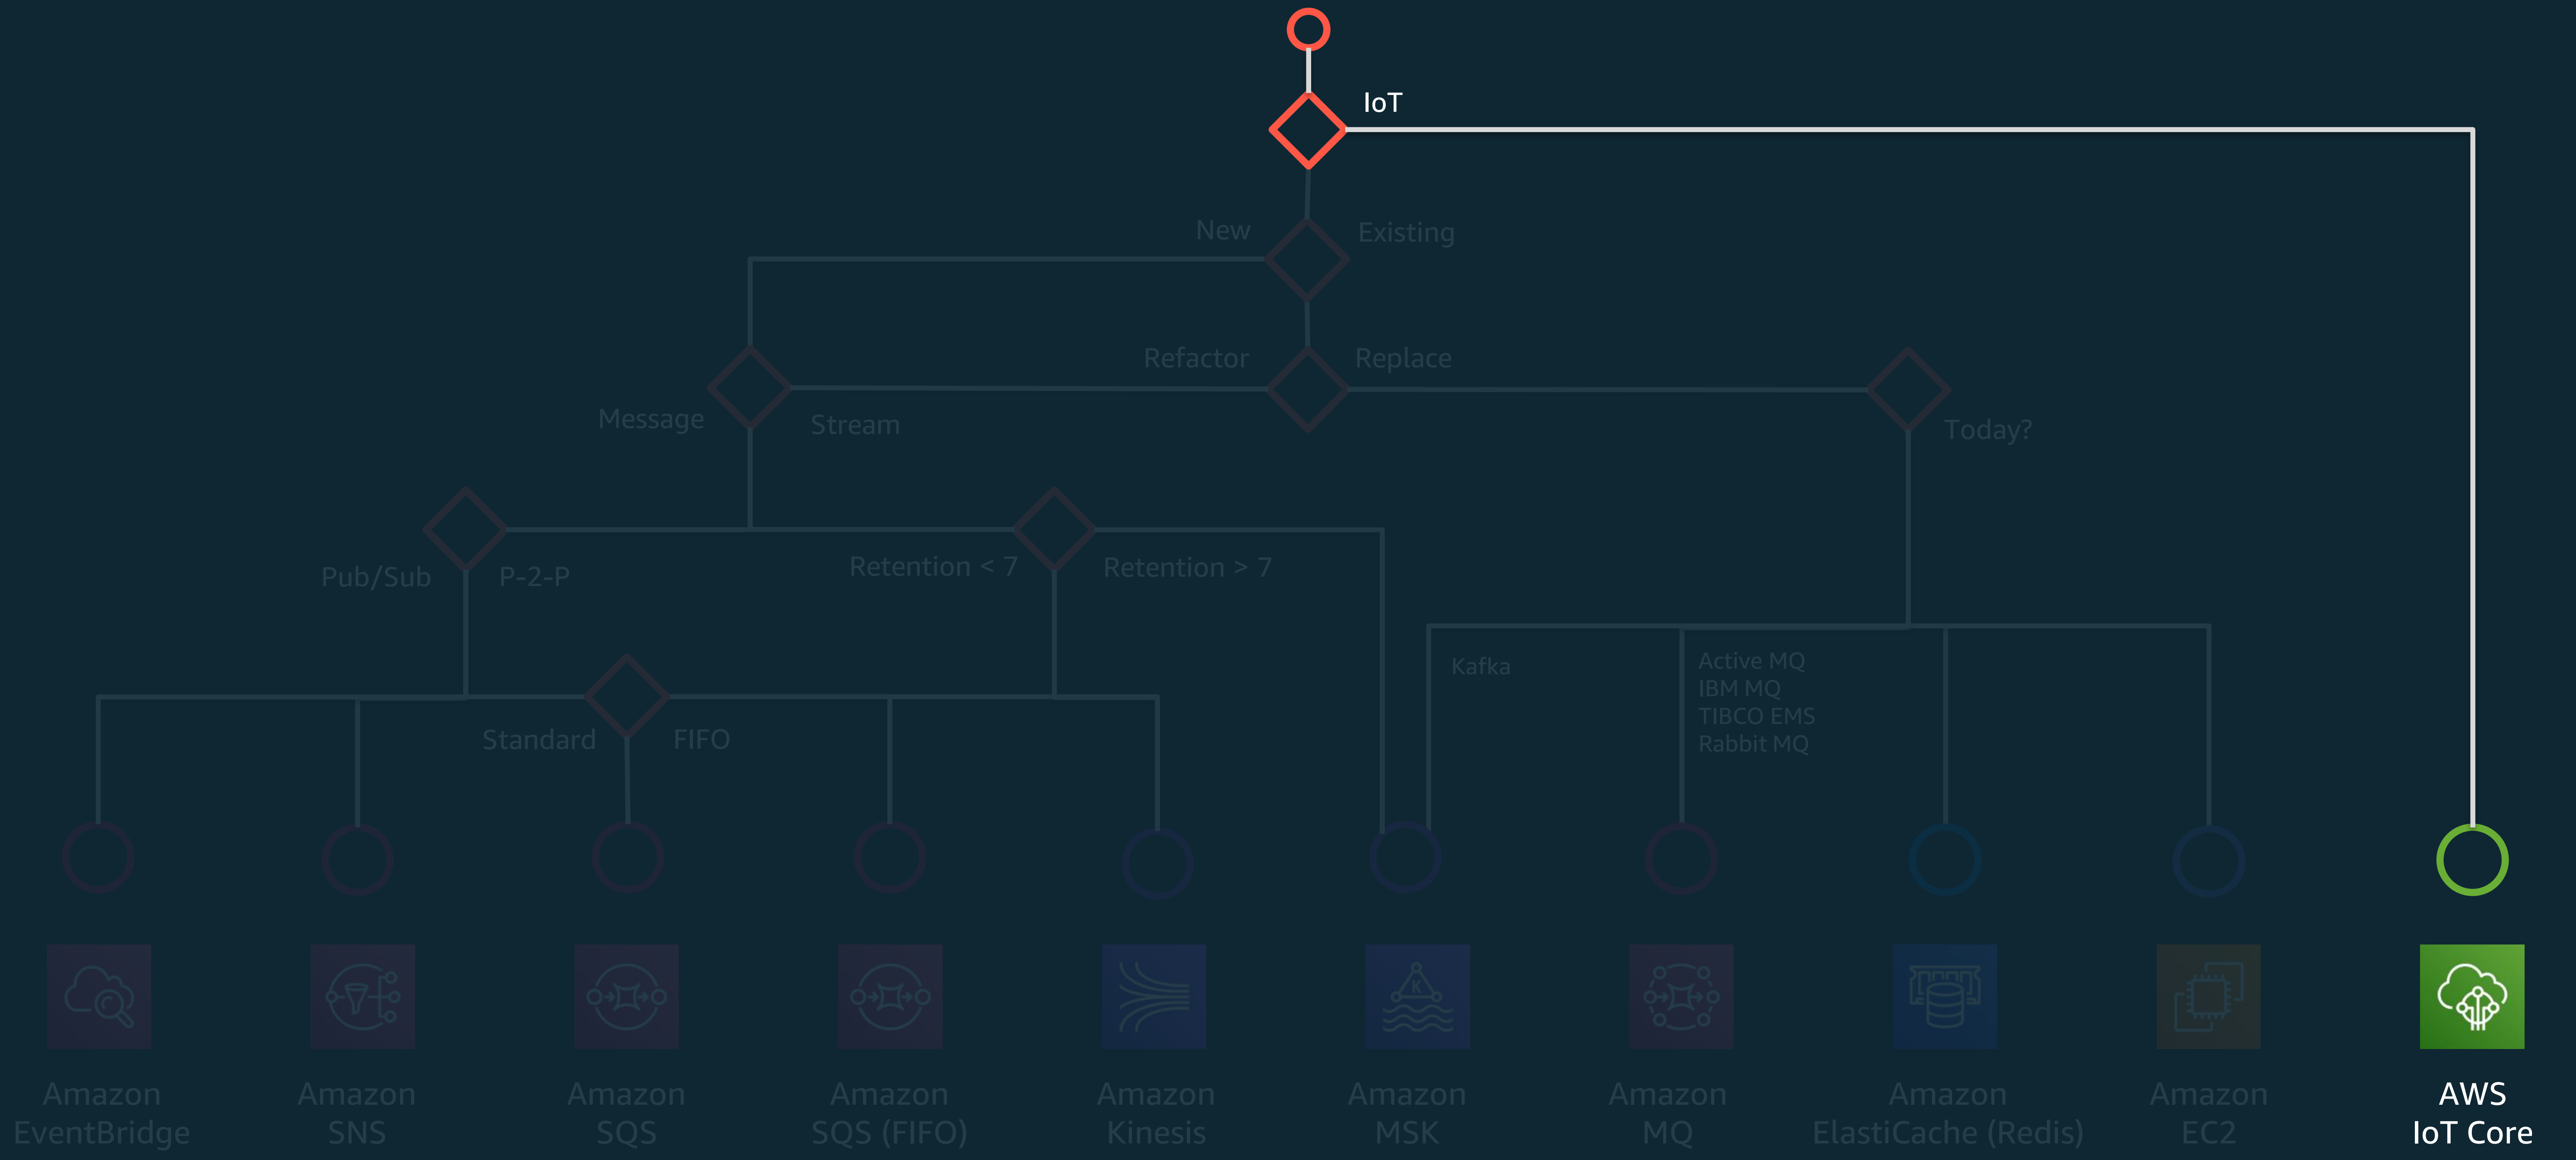

Which service should I recommend?

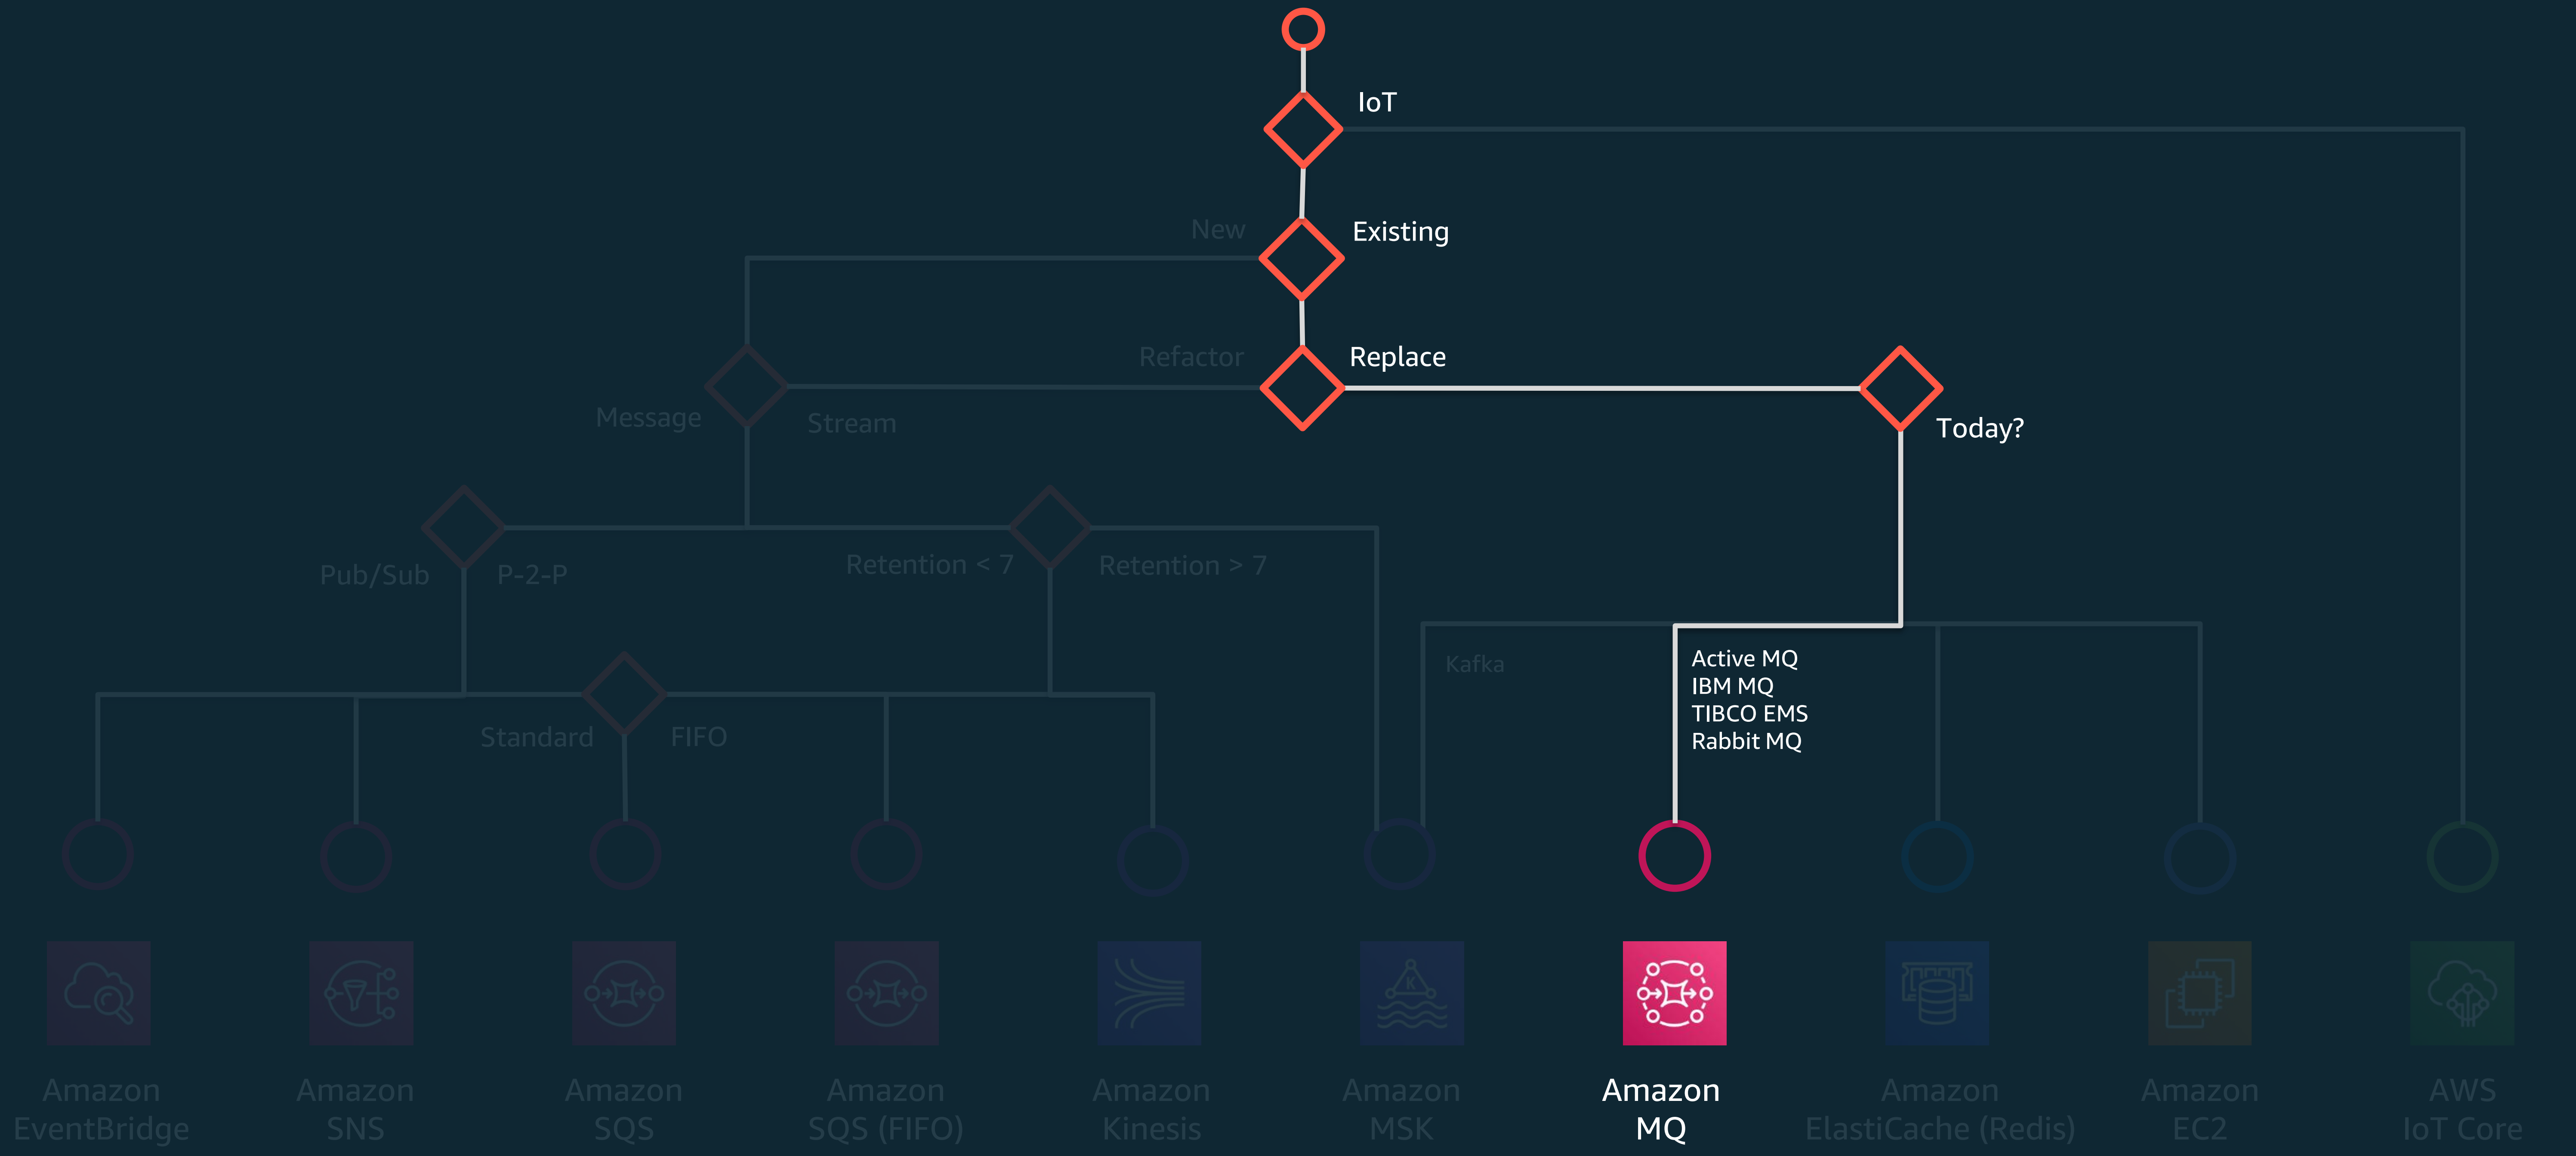

Which service should I recommend?

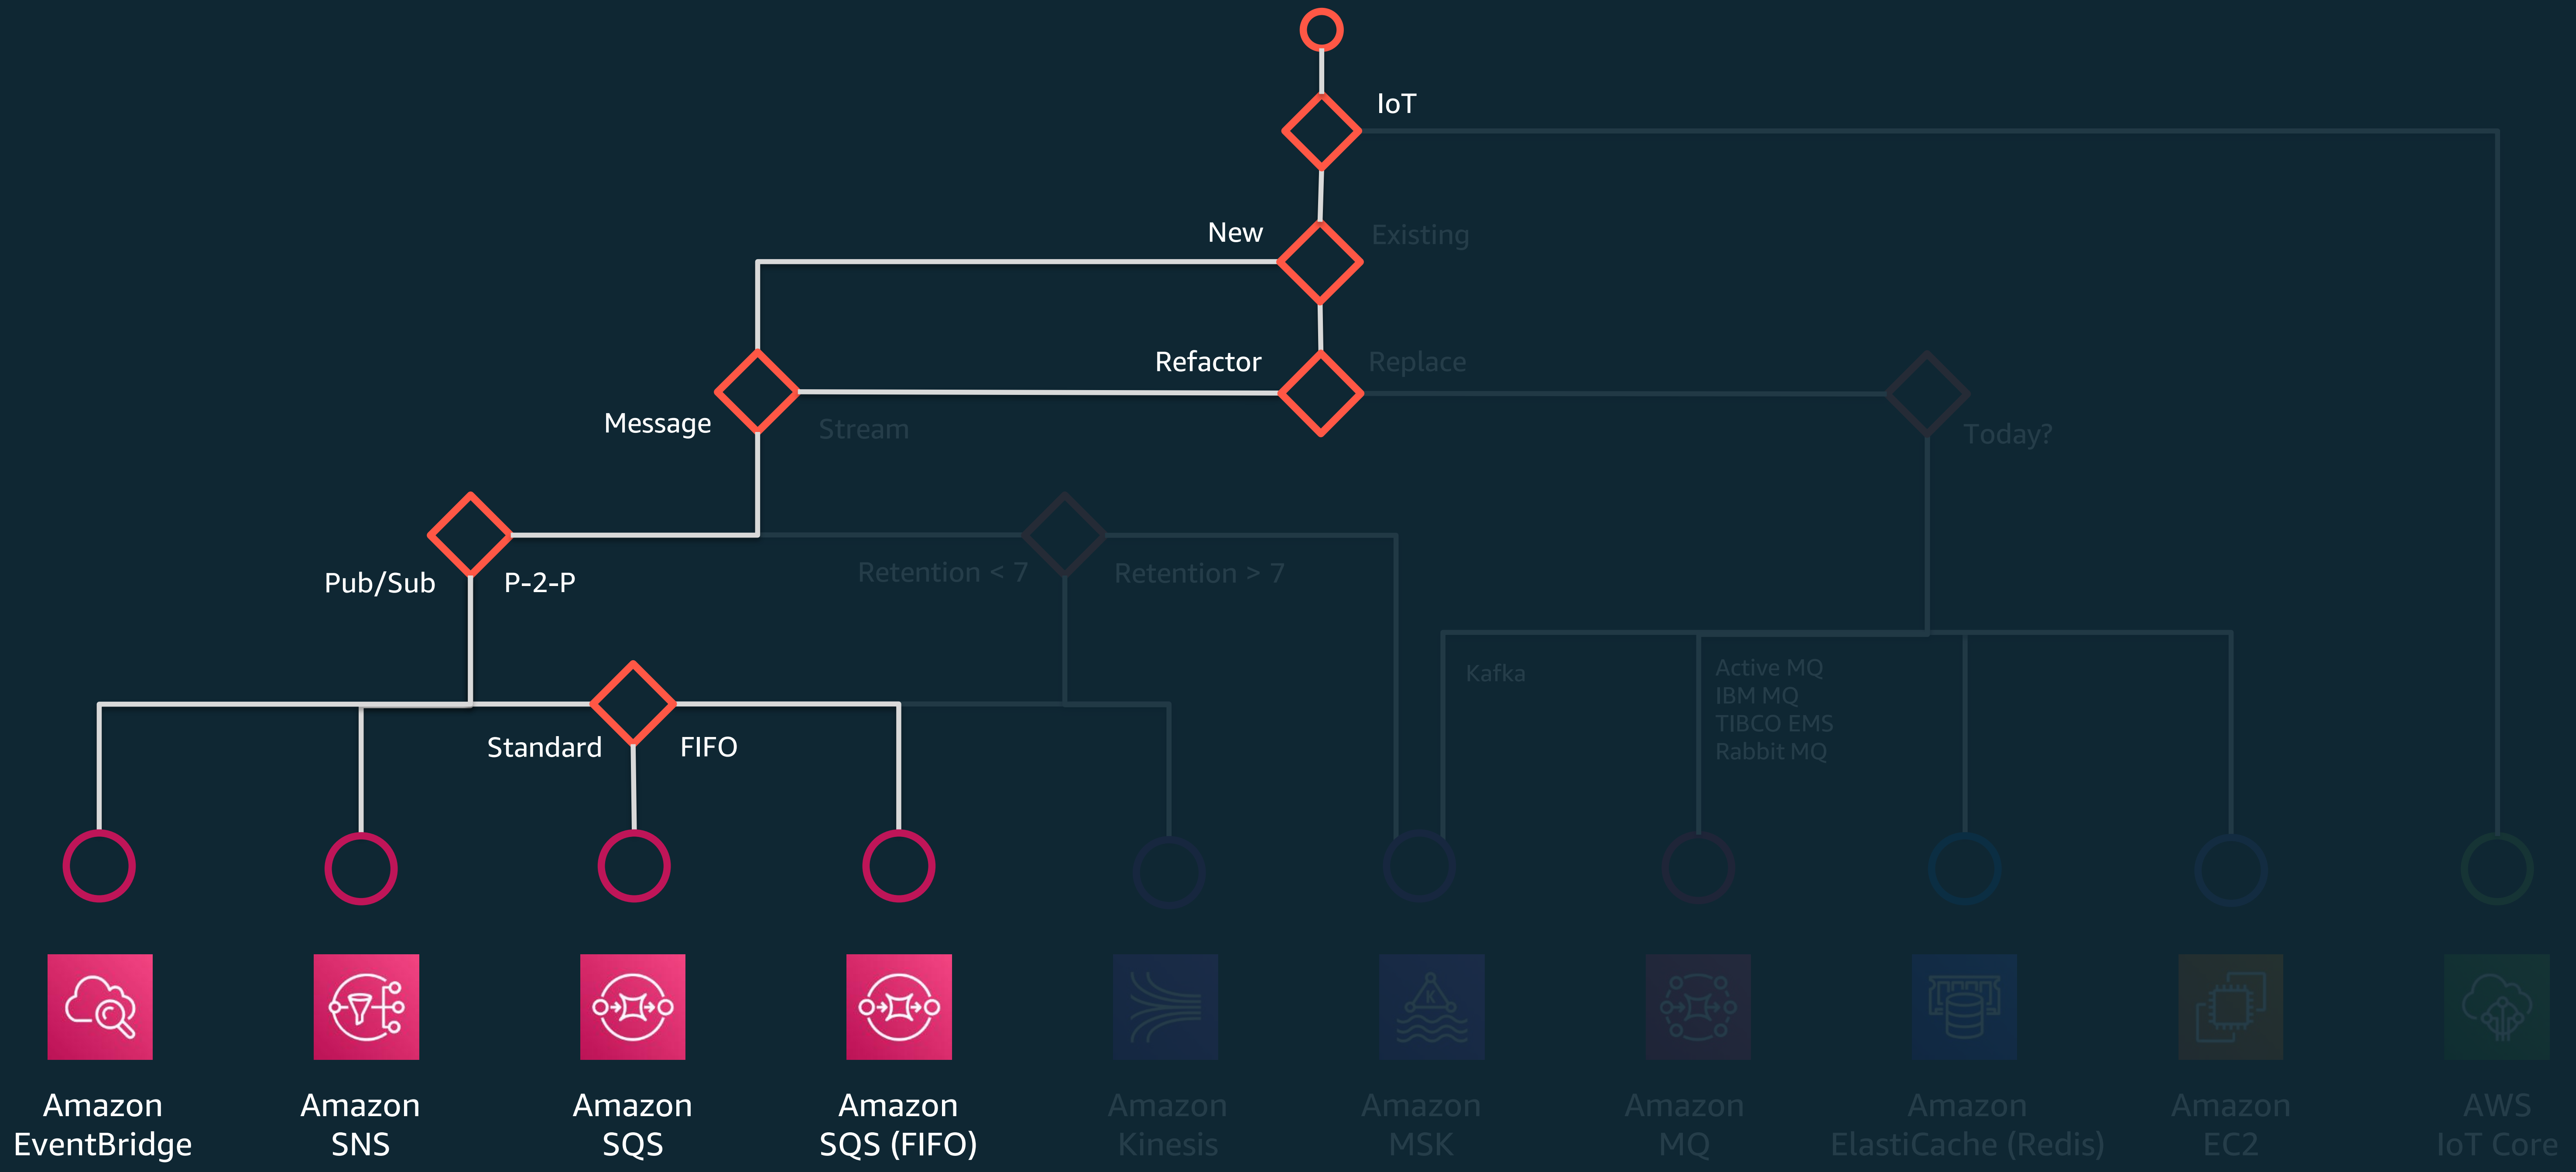

Which service should I recommend?

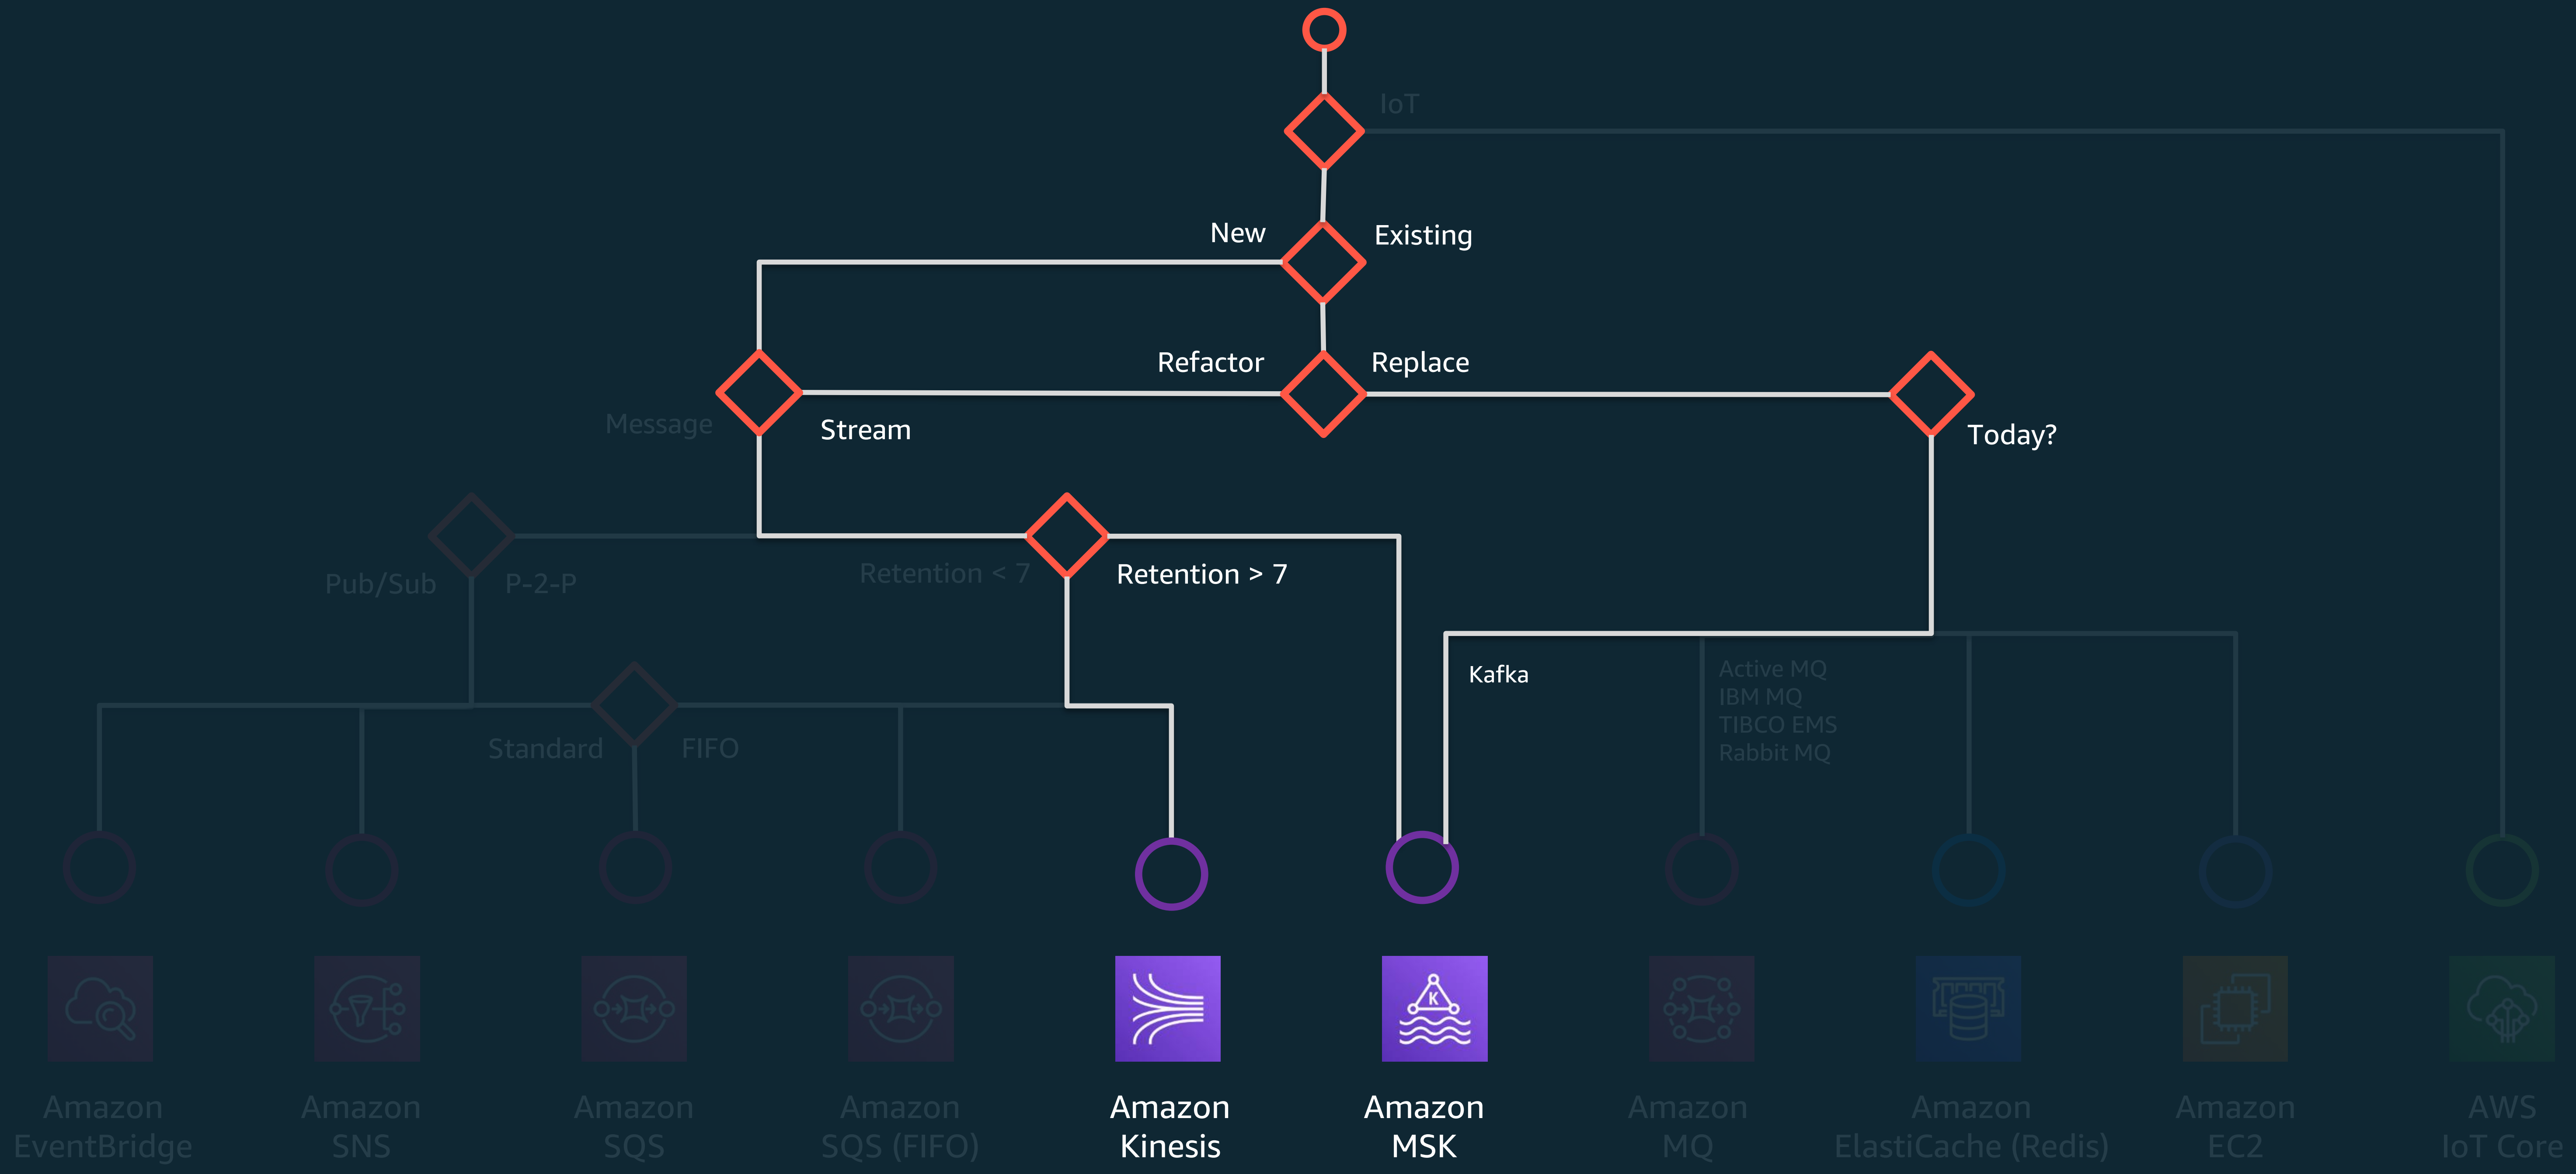

IoT

New

Existing

Refactor

Replace

Message

Stream

Today?

Pub/Sub

P-2-P

Retention < 7

Retention > 7

Kafka

Active MQ
IBM MQ
TIBCO EMS
Rabbit MQ

Standard

FIFO

Amazon EventBridge

Amazon SNS

Amazon SQS

Amazon SQS (FIFO)

Amazon Kinesis

Amazon MSK

Amazon MQ

Amazon ElastiCache (Redis)

Amazon EC2

AWS IoT Core

Which service should I recommend?

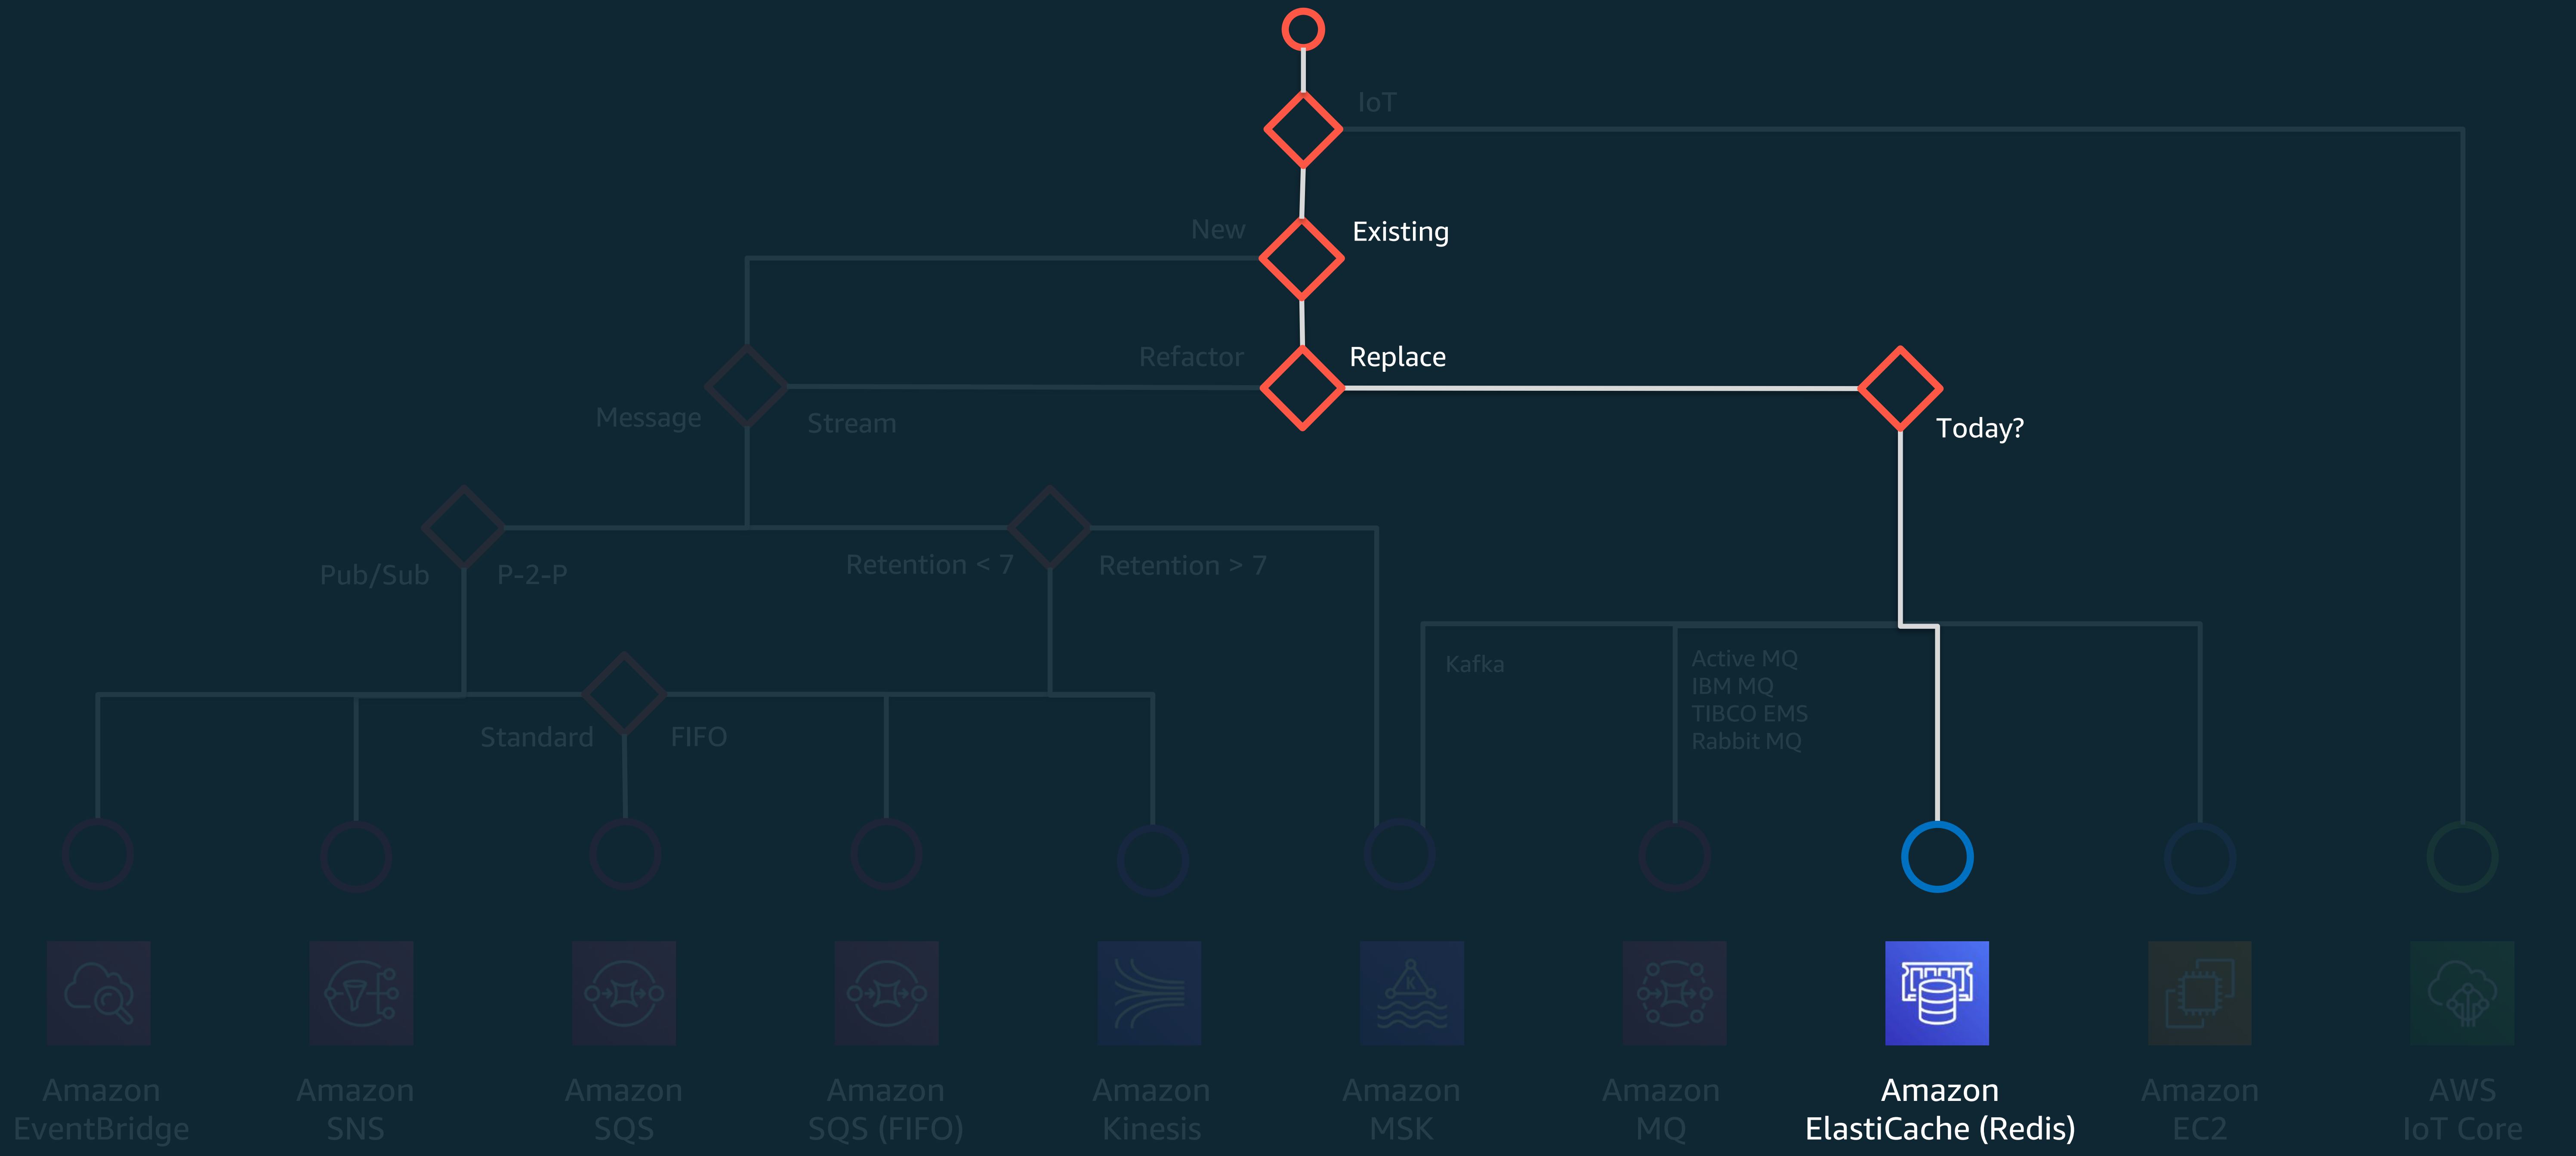

Which service should I recommend?

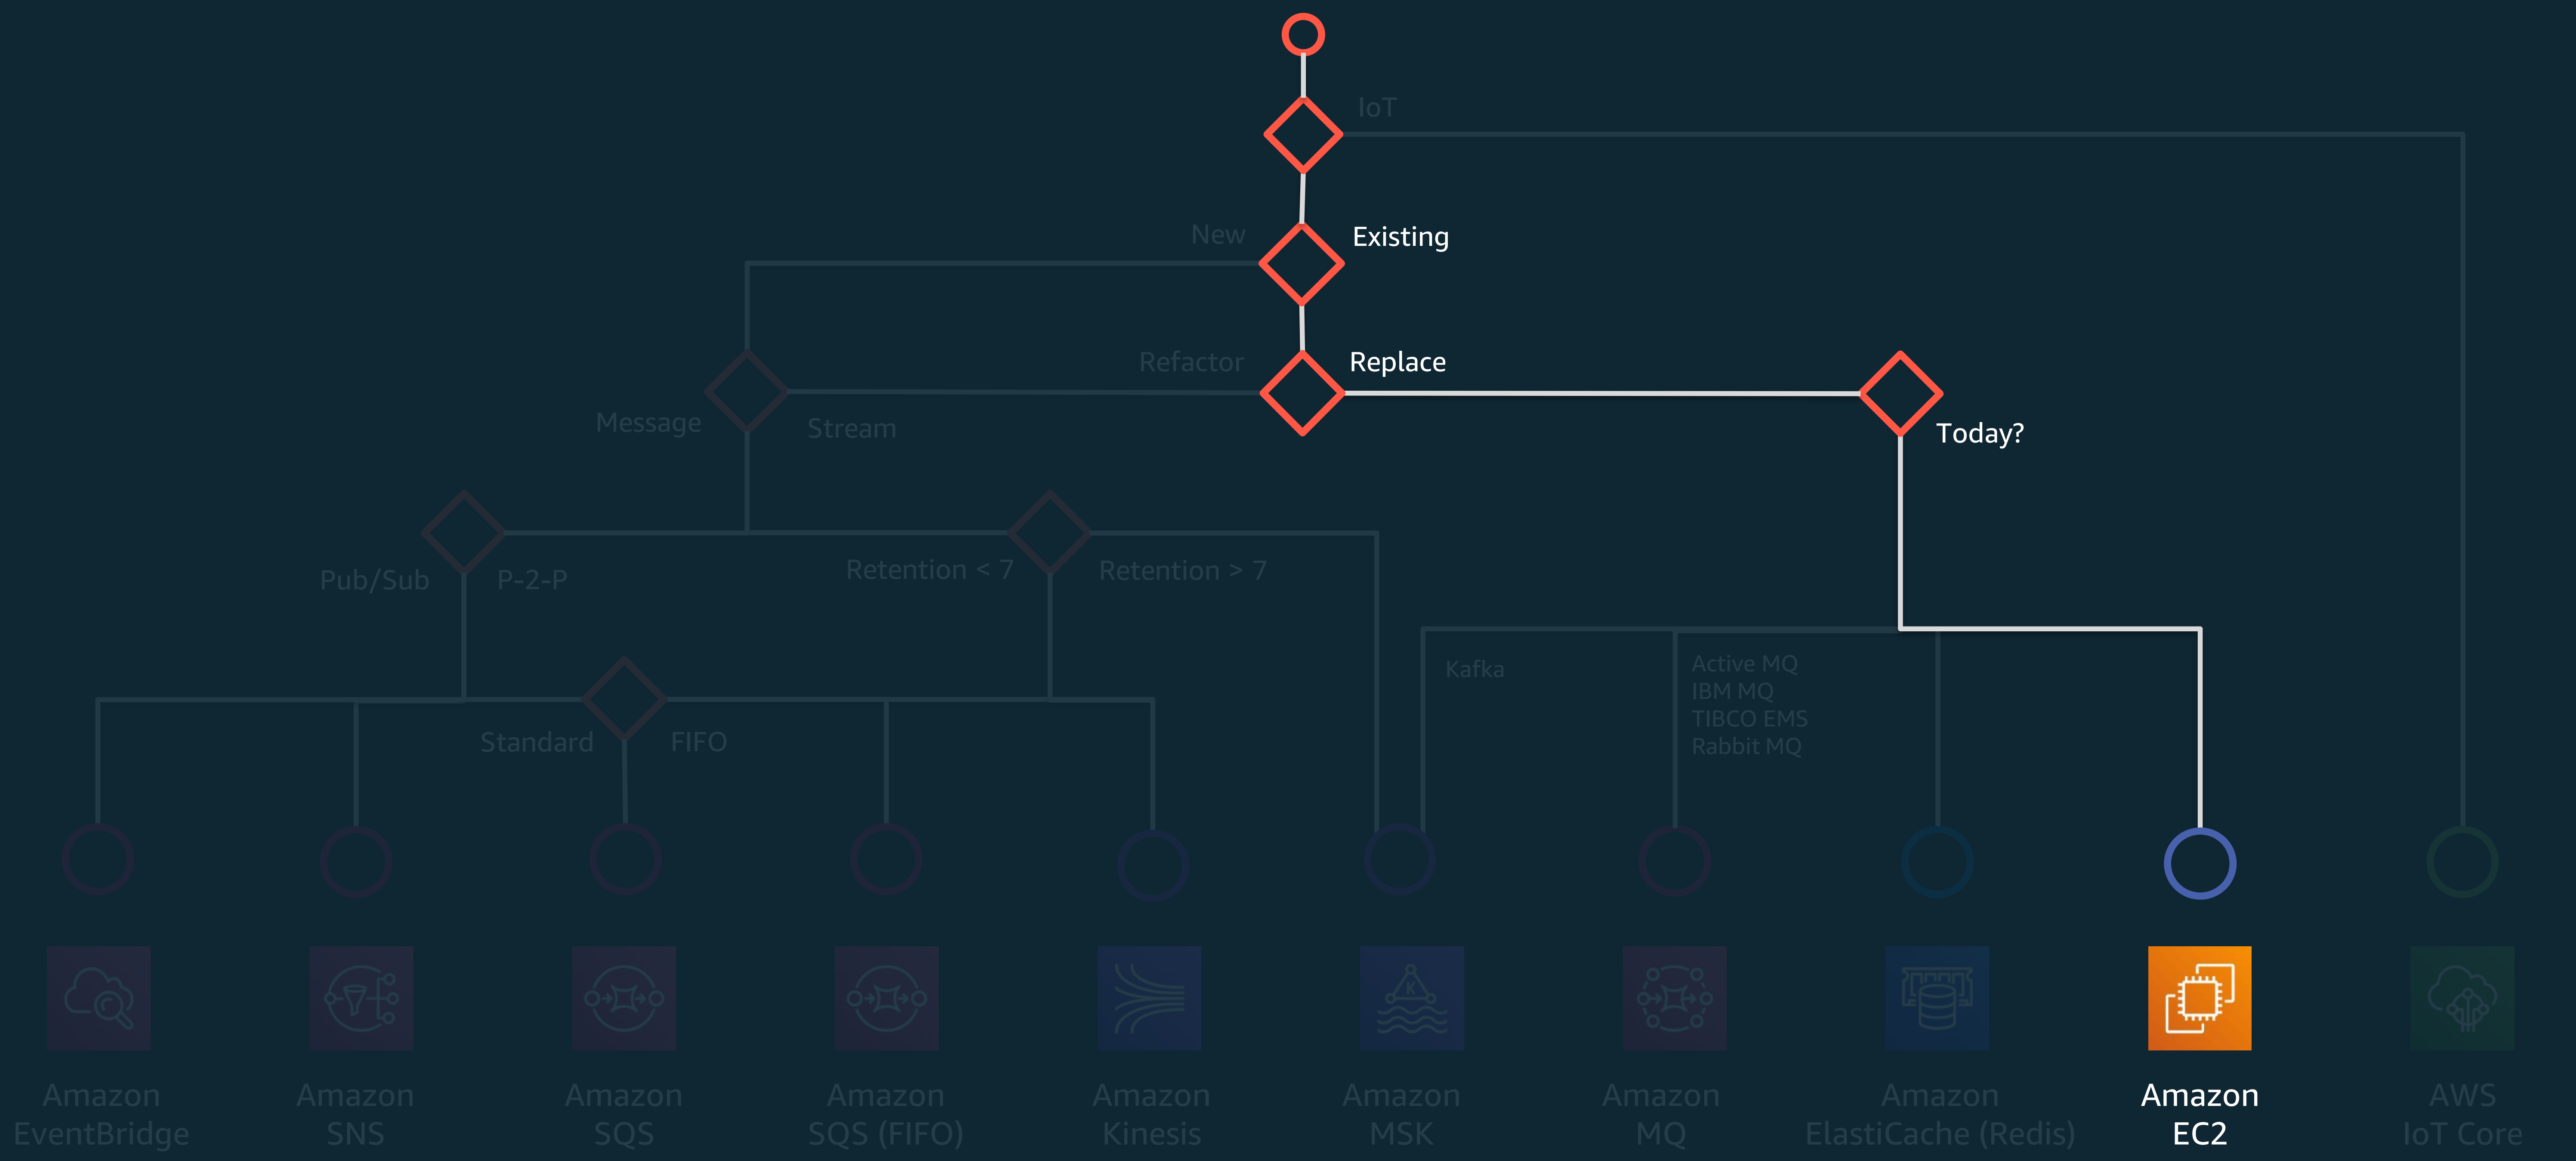

Integration services - differentiation

Integration services - differentiation

- Scalability
- Durability
- Persistence
- Consumption models
- Retries
- Pricing

Scalability

Service	How do you scale the service?
SNS	automatically
EventBridge	automatically
SQS	automatically
Kinesis Streams	adding shards to a stream one shard provides ingress capacity of 1MB/sec or 1000 records/sec, up to 2MB/sec of egress
Amazon MQ	vertically and horizontally (forming a network of brokers)
Amazon MSK	vertically (adding brokers to the cluster)

Durability

Service	Durability of requests “in flight”
SNS	replicated across multiple AZ’s
EventBridge	replicated across multiple AZ’s
SQS	replicated across multiple AZ’s
Kinesis Streams	replicated across 3 AZ’s
Amazon MQ	replicated across multiple AZ’s (when using persisted messaging)
Amazon MSK	Kafka: data is replicated to multiple nodes in a cluster (default 3) EBS: data is replicated across multiple servers in an AZ

Persistence

Service	Persistence of requests “in flight”
SNS	no formal persistence model - delivery retry up to potentially 23 days (different per protocol) supports DLQ’s
EventBridge	no formal persistence model - delivery retry up to potentially 24 hours
SQS	up to 14 days - default 4 days
Kinesis Streams	up to 7 days (with additional cost) - default 24 hours
Amazon MQ	storage capacity per broker is 200 GB (20 GB for mq.t2.micro)
Amazon MSK	as provisioned (between 1000 GB and 16384 GB)

Consumption

Service	Invocation model
SNS	push based integrated with AWS Lambda, Amazon SQS, HTTP, SMTP, mobile push
EventBridge	push based Integrated with 17 AWS services, including AWS Lambda, Amazon SQS, Amazon SNS, Amazon ECS, AWS Batch and AWS Step Functions
SQS	Pull based integrated with AWS Lambda
Kinesis Streams	Pull based integrated with AWS Lambda clients has to maintain a checkpoint/cursor in the stream
Amazon MQ	pull/push based typically with protocols like JMS, NMS, AMQP or MQTT
Amazon MSK	pull based clients has to maintain a checkpoint/cursor in the stream

Retry / Failure handling

Service	Retry/failure capabilities
SNS	retry delivery with back-off up to 23 days Lambda & SQS: 23 days SMS & SMTP & mobile push: over 6 hours HTTP: custom
EventBridge	retry delivery with back-off up to 24 hours
SQS	messages remain in the queue until deleted (or expired) consumers responsibility to retry
Kinesis Streams	messages remain in the queue until expired consumers responsibility to move on with the checkpoint/cursor
Amazon MQ	messages remain in the queue until deleted consumers responsibility to retry
Amazon MSK	messages remain in the queue until expired consumers responsibility to move on with the checkpoint/cursor

Pricing

Service	Model/Cost
SNS	per request $ 0.50 per 1 mio. requests (64 KB chunk)
EventBridge	per events published $ 1.00 per 1 mio. events (256 KB chunk) AWS events in same account are FREE
SQS	per request $ 0.40 per 1 mio. requests (64 KB chunk) up to 10 messages per request
Kinesis Streams	per shard hour & put payload requests $ 0.015 per shard hour & $ 0.014 put payload requests (256 KB chunk) additional cost for Enhanced Fanout and Extended Data Retention
Amazon MQ	per instance hour & storage used $ 0.03 - $ 2.304 per instance hour & $ 0.30 per GB-month storage used
Amazon MSK	per instance hour & storage provisioned $ 0.21 - $ 10.08 per instance hour & $ 0.10 per GB-month storage provisioned

Selected new features from 2019

Amazon MQ

- CloudTrail and CloudFormation support
- Compliance programs (SOC-1,2,3, PCI, ISO)
- Region expansion
- Enhanced monitoring
- Network of brokers
- Resource-level and tag-based permissions
- CMK-KMS support
- Enhanced broker configuration

Amazon SQS

- SLAs
- VPC endpoint policies
- SQS FIFO in all regions
- X-Ray
- 1-minute CW Metrics
- Tag on Create
- FIFO – Lambda

Amazon SNS

- SLAs
- Improved Console
- Fork Design Patterns
- VPC endpoint policies
- Cost-allocation Tags
- X-Ray
- Dead Letter Queues

Amazon EventBridge

- Launched!

Amazon Kinesis

- Regional expansion
- VPC endpoint for Data Firehose
- Tag on Create for Data Firehose
- SLA 99.9% (data streams & firehose)

Amazon MSK

- GAed!
- Two-AZ deployment (in addition to 1 and 3)
- PCI DSS

Open discussion

Resources / Call-to-action

Resources/Call-to-action

AWS blogs and other content about application integration

- http://bit.ly/2019-api312

Resources/Call-to-action

AWS blogs and other content about application integration

- https://aws.amazon.com/blogs/compute/tag/messaging/

Talk about app-int patterns for microservices

- API315 during this re:Invent (Monday + Tuesday + Wednesday + Thursday)

Resources/Call-to-action

AWS blogs and other content about application integration

- https://aws.amazon.com/blogs/compute/tag/messaging/

Talk about app-int patterns for microservices

- API315 during this re:Invent (Monday + Tuesday + Wednesday + Thursday)

Hands-on workshop on implementing the patterns from this talk

- ARC314 during this re:Invent (Monday + Tuesday)
- Ask your AWS SA for an application integration immersion day

Resources/Call-to-action

AWS blogs and other content about application integration

- https://aws.amazon.com/blogs/compute/tag/messaging/

Talk about app-int patterns for microservices

- API315 during this re:Invent (Monday + Tuesday + Wednesday + Thursday)

Hands-on workshop on implementing the patterns from this talk

- ARC314 during this re:Invent (Monday + Tuesday)
- Ask your AWS SA for an application integration immersion day

Keep in mind

- Loose coupling is better than lousy coupling

Please complete the session survey in the mobile app.

AWS re:Invent

aws

Thank you! Go build!

Christian Müller

cmr@amazon.de | @ChristianM

Dirk Fröhner

froehner@amazon.de | @dirk_f5r