
AWS
re:Invent

ANT325

Guide to Apache Kafka replication and migration with Amazon MSK

Jeremiah O'Connor
Solutions Architect
AWS

Priyanka Chaudhary
Associate Big Data Specialist
AWS

AWS re:Invent

© 2020, Amazon Web Services, Inc. or its affiliates. All rights reserved.

Agenda

Why Amazon MSK?

Apache Kafka mirroring tools and architectures

Migration using MirrorMaker 2.0

Migration using Apache Flink

Amazon MSK migration and capacity planning best practices



© 2020, Amazon Web Services, Inc. or its affiliates. All rights reserved.

Why Amazon Managed Streaming for Apache Kafka (MSK)?



AWS re:Invent

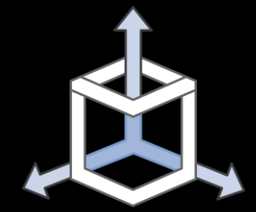
© 2020, Amazon Web Services, Inc. or its affiliates. All rights reserved.

Challenges operating Apache Kafka

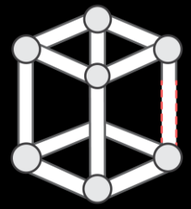
Difficult to set up



Difficult to scale



Hard to achieve high availability



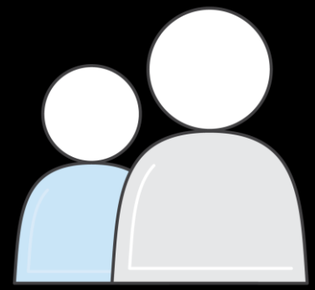
AWS integrations = development



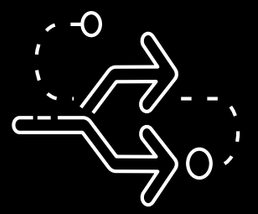
No console, no visible metrics



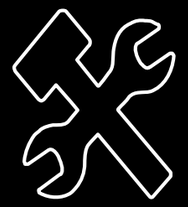
$$f(kafka_{usage}) = \sum_{n=1}^{\infty}(SRE)$$



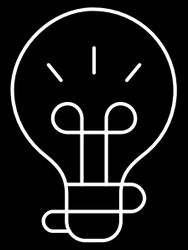
What Amazon MSK does for you



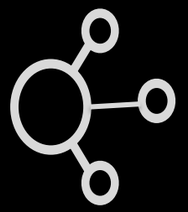
Fully managed, production-ready Apache Kafka clusters



Drives best practices through design, defaults, and automation



Security, scalability, and high availability for enterprise workloads



Amazon MSK is committed to improving open-source Apache Kafka

Apache Kafka mirroring tools and architectures



AWS re:Invent

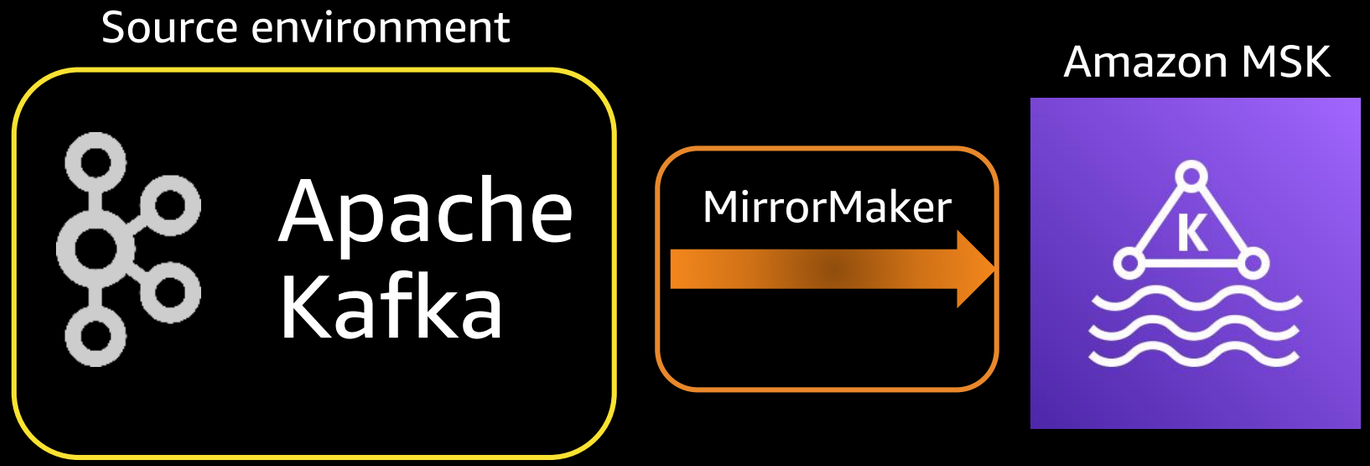
© 2020, Amazon Web Services, Inc. or its affiliates. All rights reserved.

Apache Kafka mirroring tools

- Apache Kafka mirroring tools help in maintaining a replica of an existing Apache Kafka cluster

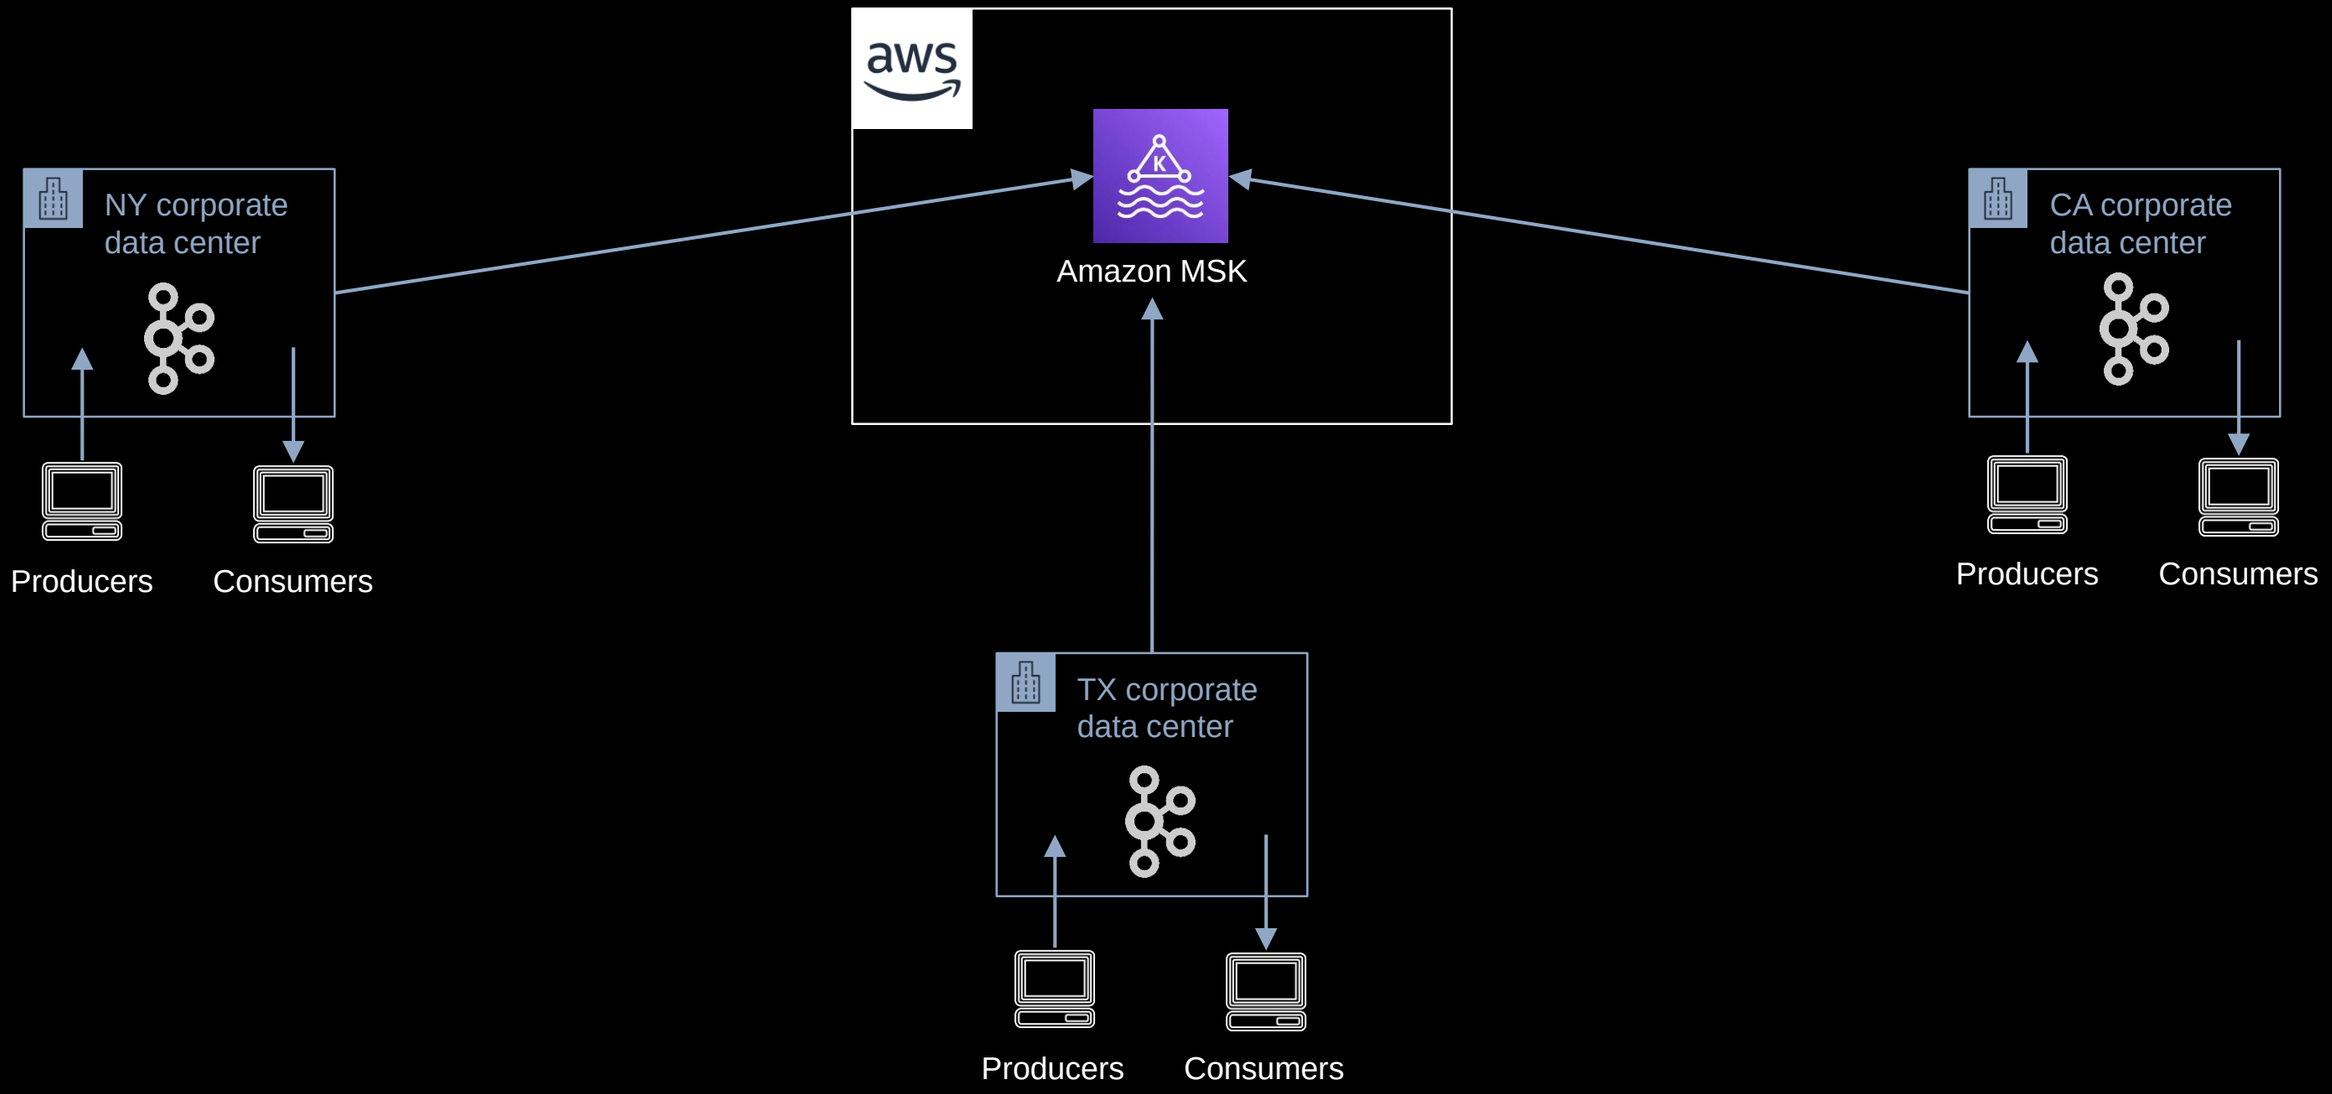
- Common uses of Apache Kafka mirroring tools
  - Disaster recovery
  - Aggregation
  - Cloud migration
  - Data isolation
  - Geoproximity

- Apache Kafka mirroring tools
  - MirrorMaker 2.0
  - MirrorMaker 1.0



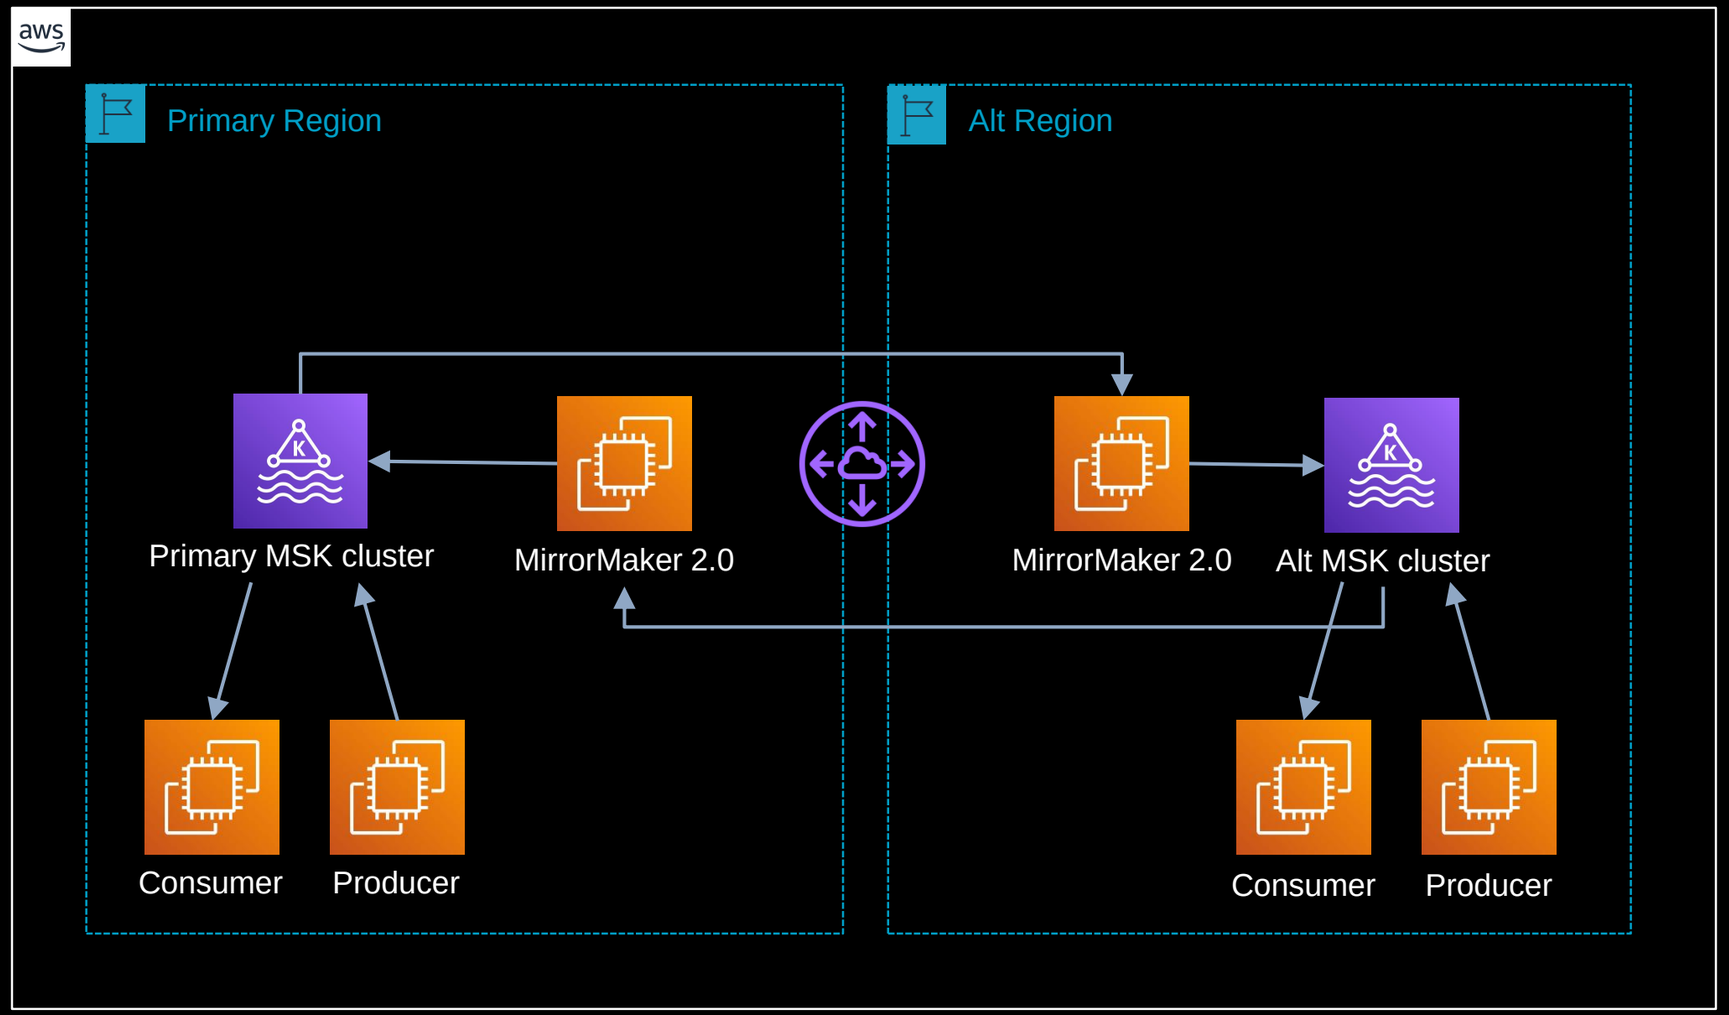
Hub-and-spoke mirror architecture

AGGREGATION



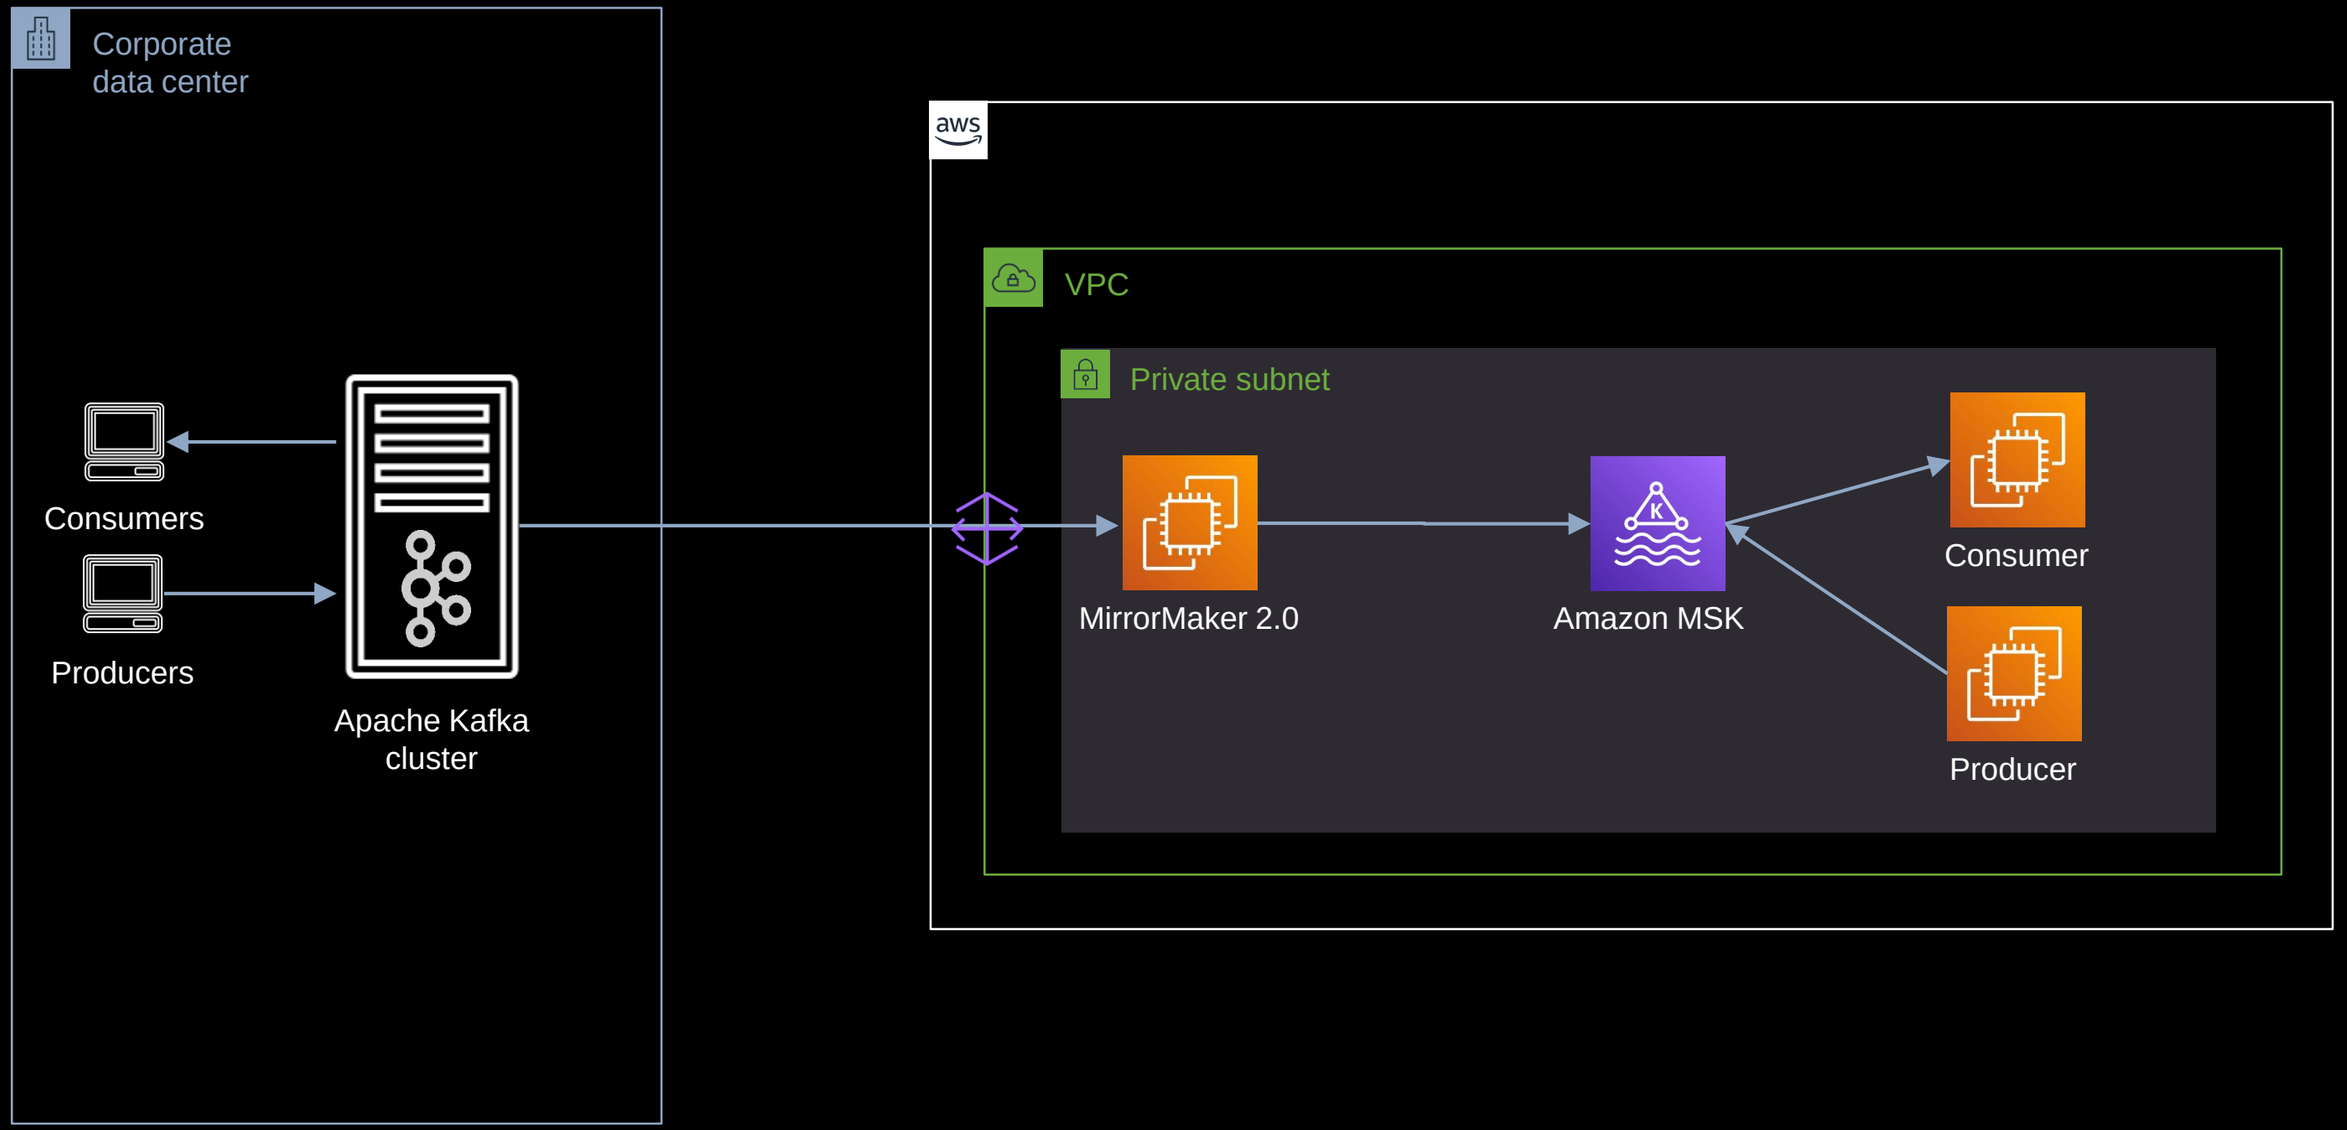
Active-Active mirror architecture

HIGH AVAILABILITY: BI-DIRECTIONAL REPLICATION



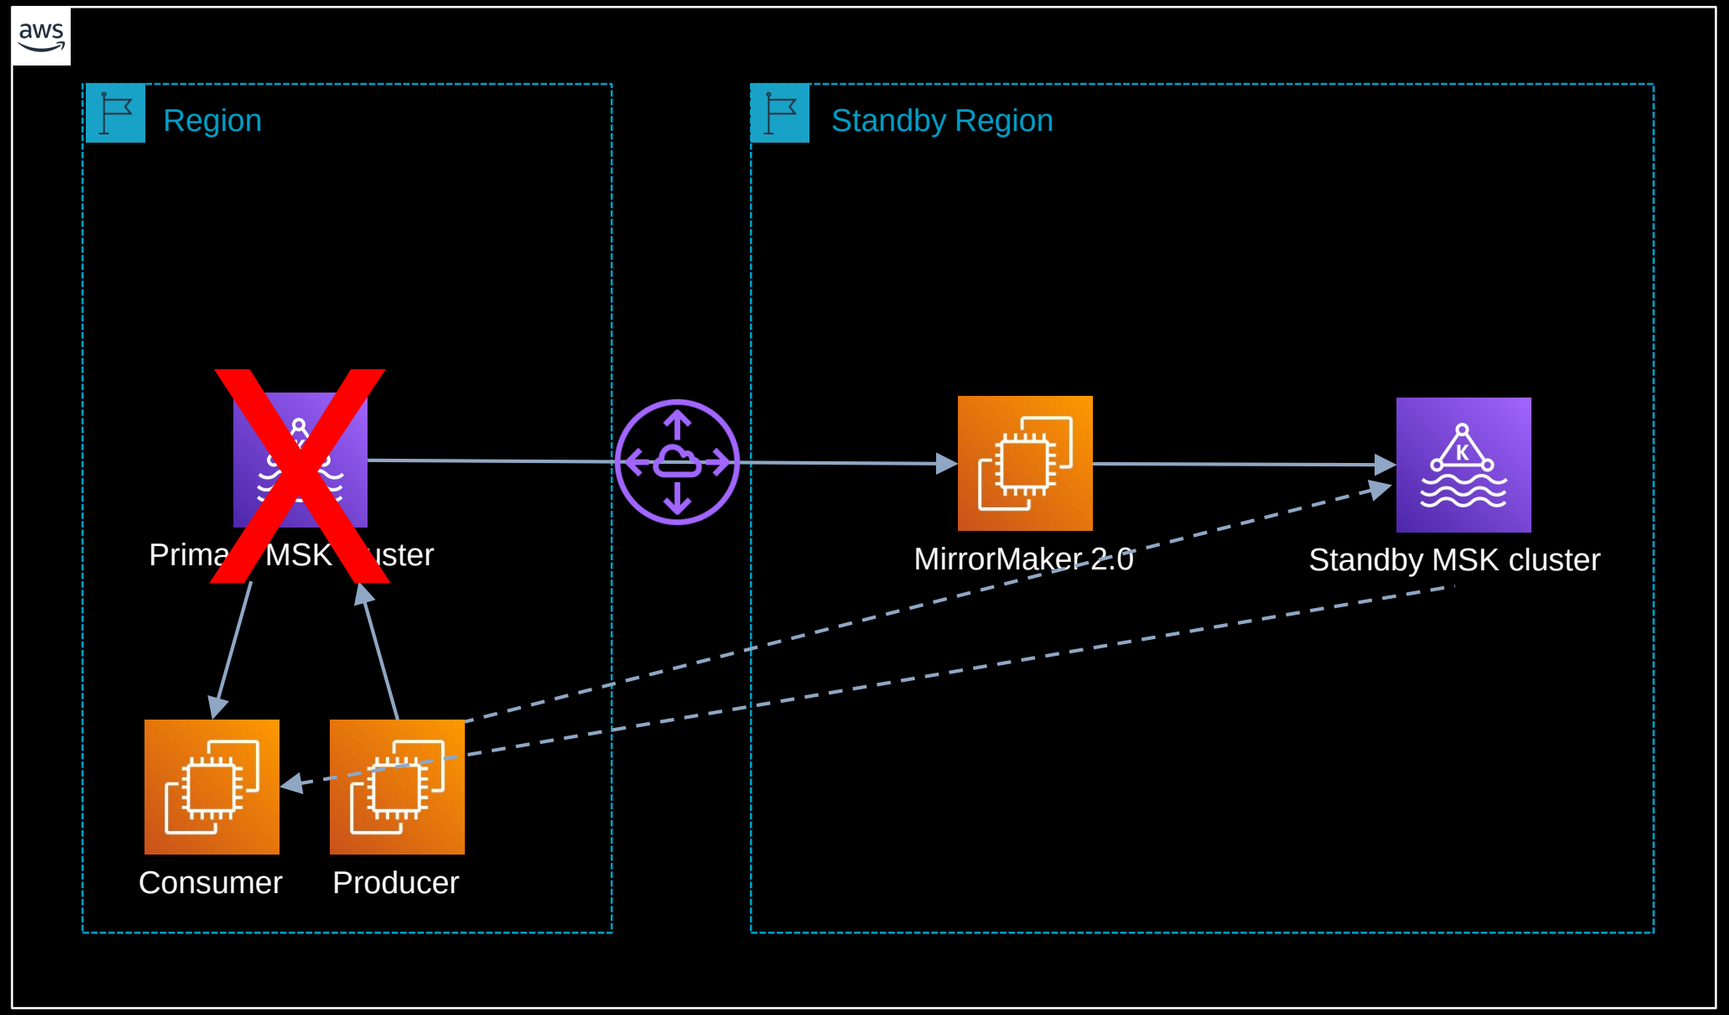
Active-Active mirror architecture

GEOPROXIMITY: 1-WAY REPLICATION



Active-passive mirror architecture

DISASTER RECOVERY



Migration using MirrorMaker 2.0



AWS re:Invent

© 2020, Amazon Web Services, Inc. or its affiliates. All rights reserved.

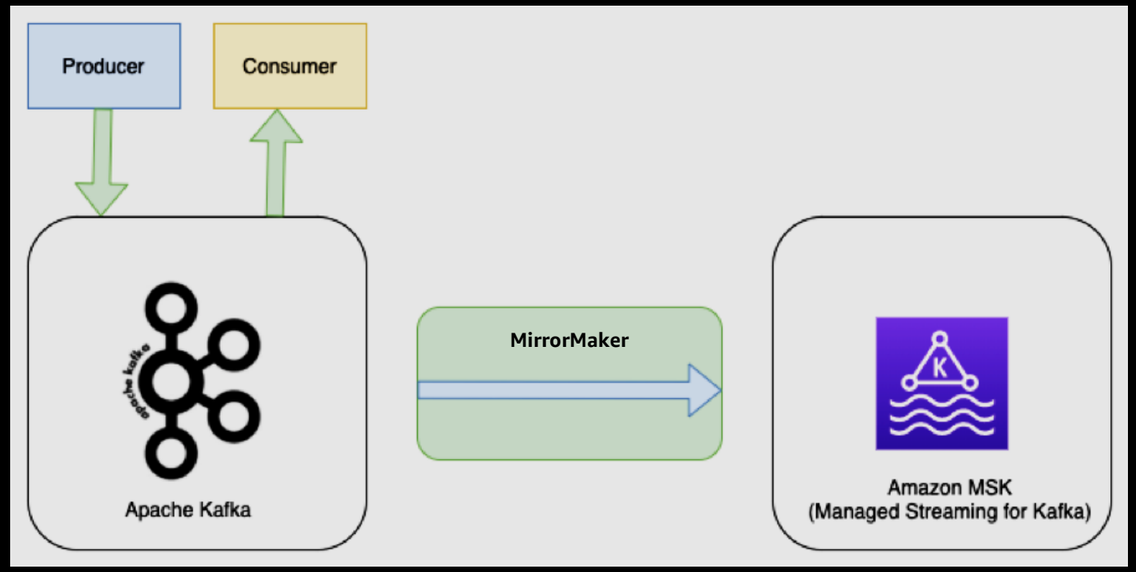
MirrorMaker 2.0 vs. MirrorMaker 1.0

MM2

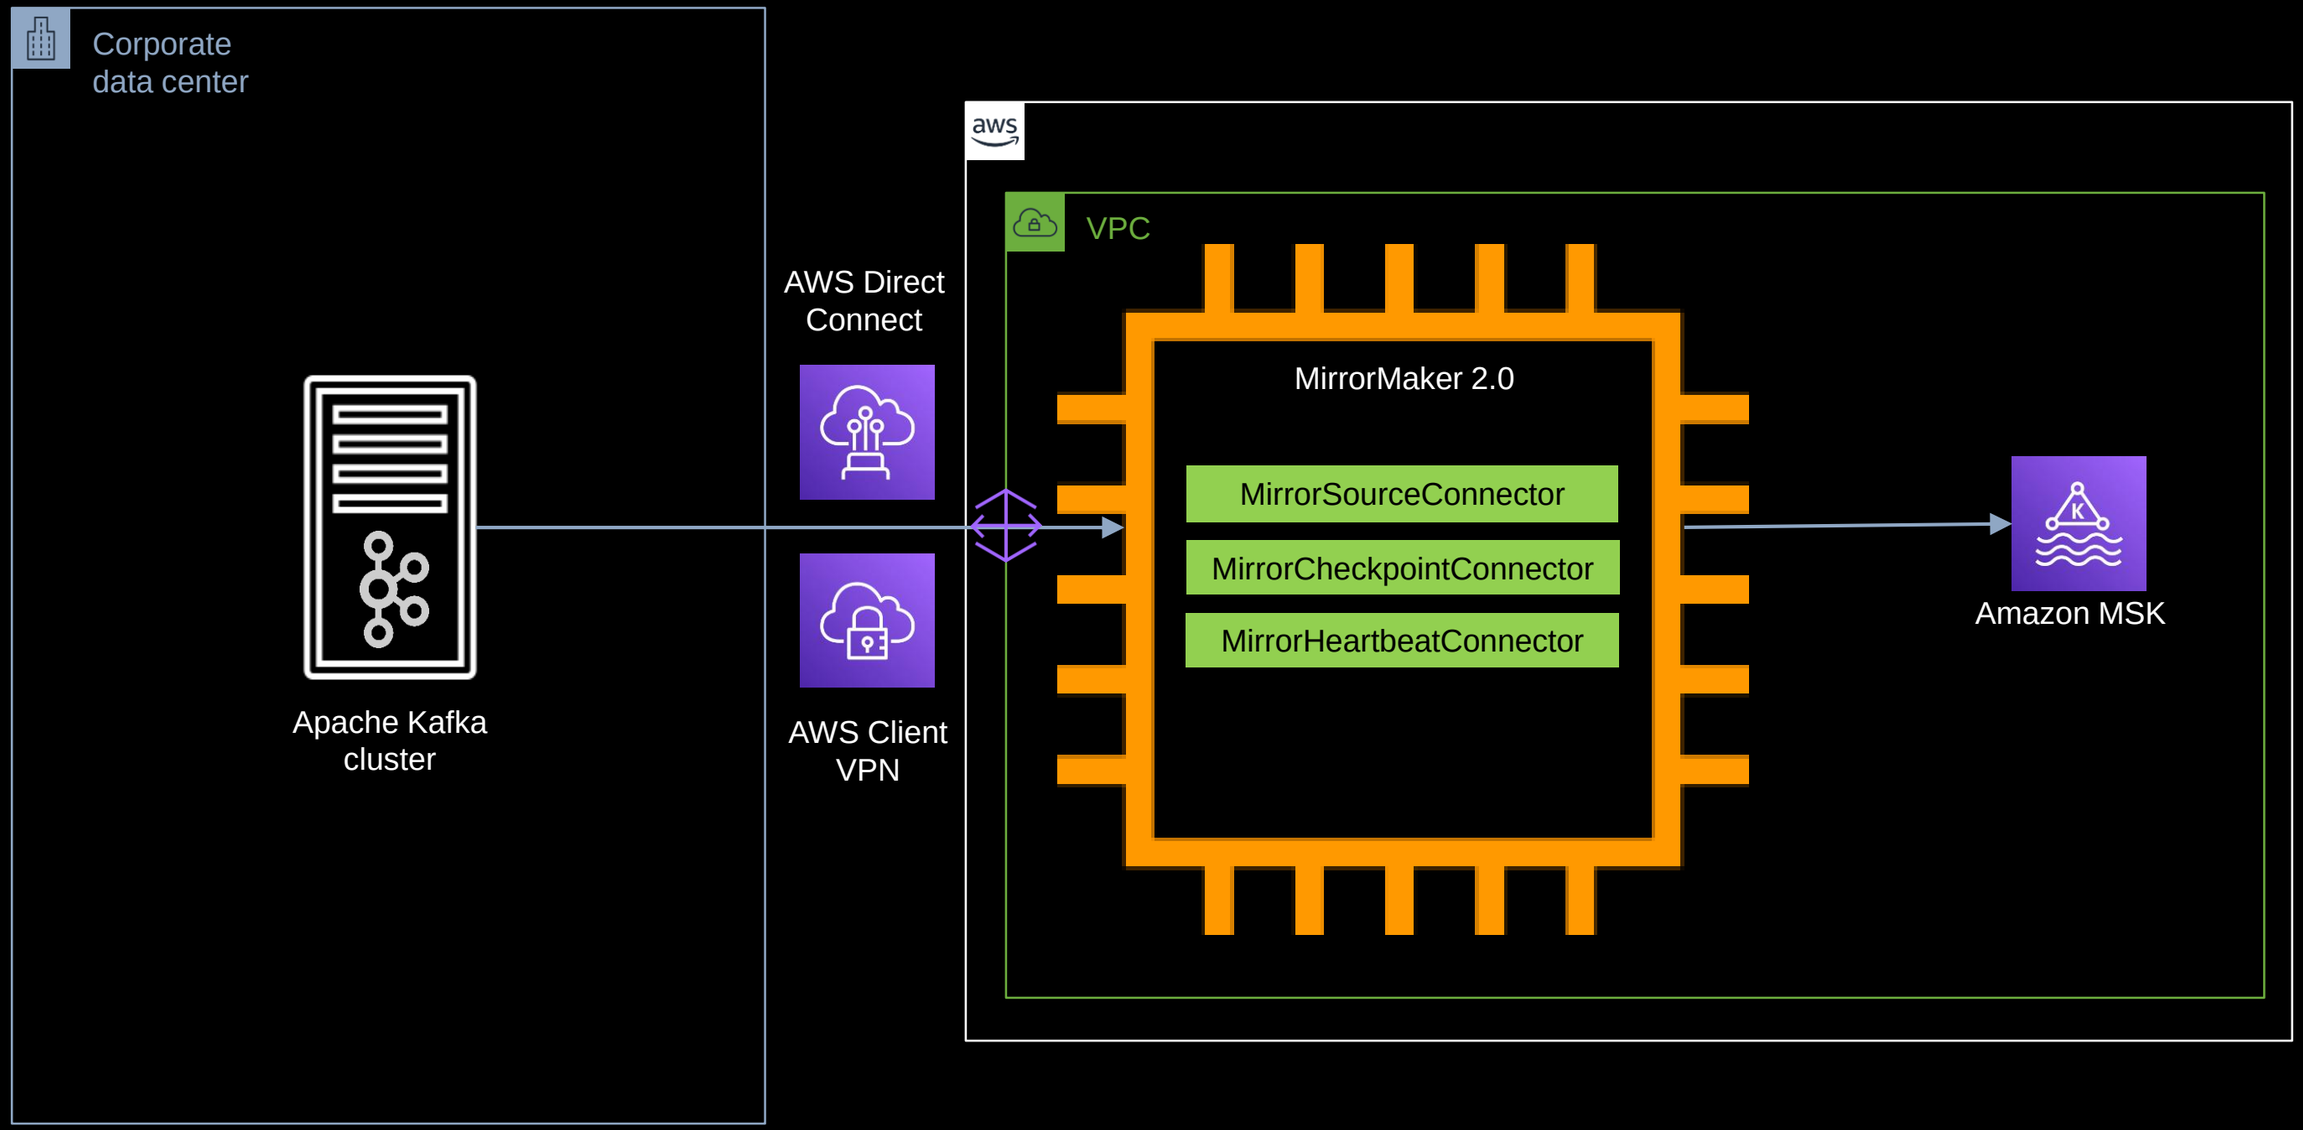
- Multi-cluster data replication engine based on Kafka Connect framework
- Detects new topics and partitions in self-discovery mode
- Automatically synchronizes topic configuration
- Supports “active/active” cluster pairs, as well as any number of active clusters
- Provides new metrics including end-to-end replication latency across multiple clusters
- Emits the offsets required to migrate consumers between clusters and tooling for offset translation
- Configuration file for specifying multiple clusters and replication flows in one place

MM1

- Tool to mirror data between two Kafka clusters
- Data is read from topics in source and written with same name in destination cluster
- Uses the cluster default config to create topics in destination cluster
- Source and destination clusters are independent of each other



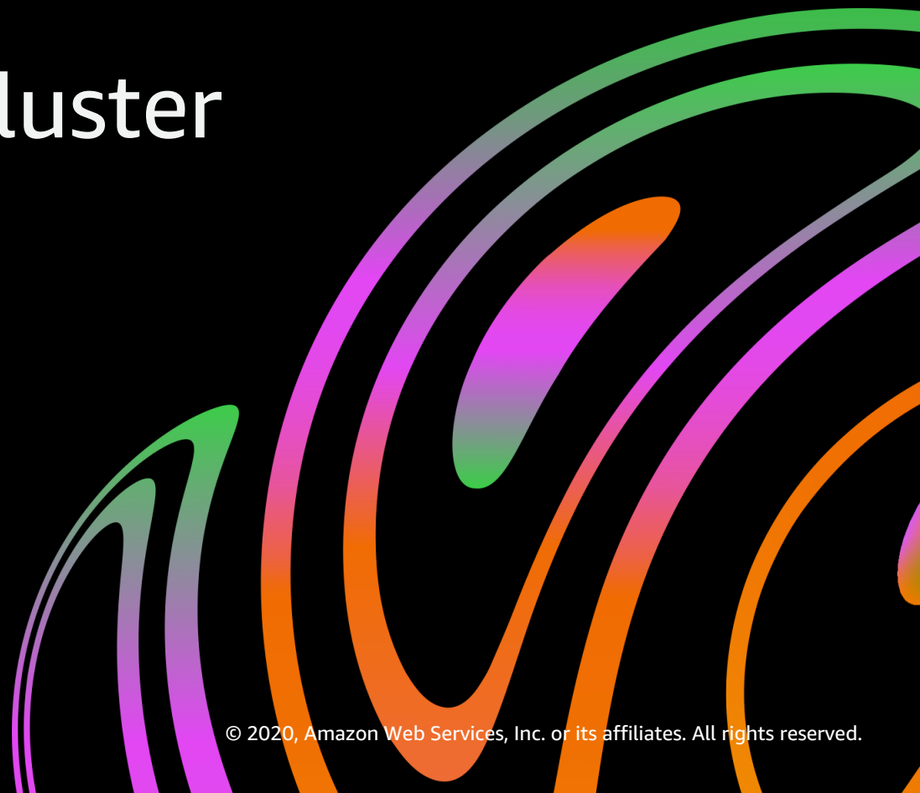
Anatomy of MirrorMaker 2.0



MirrorMaker 2.0 deployment methods

MirrorMaker 2.0 can be deployed and run four ways:

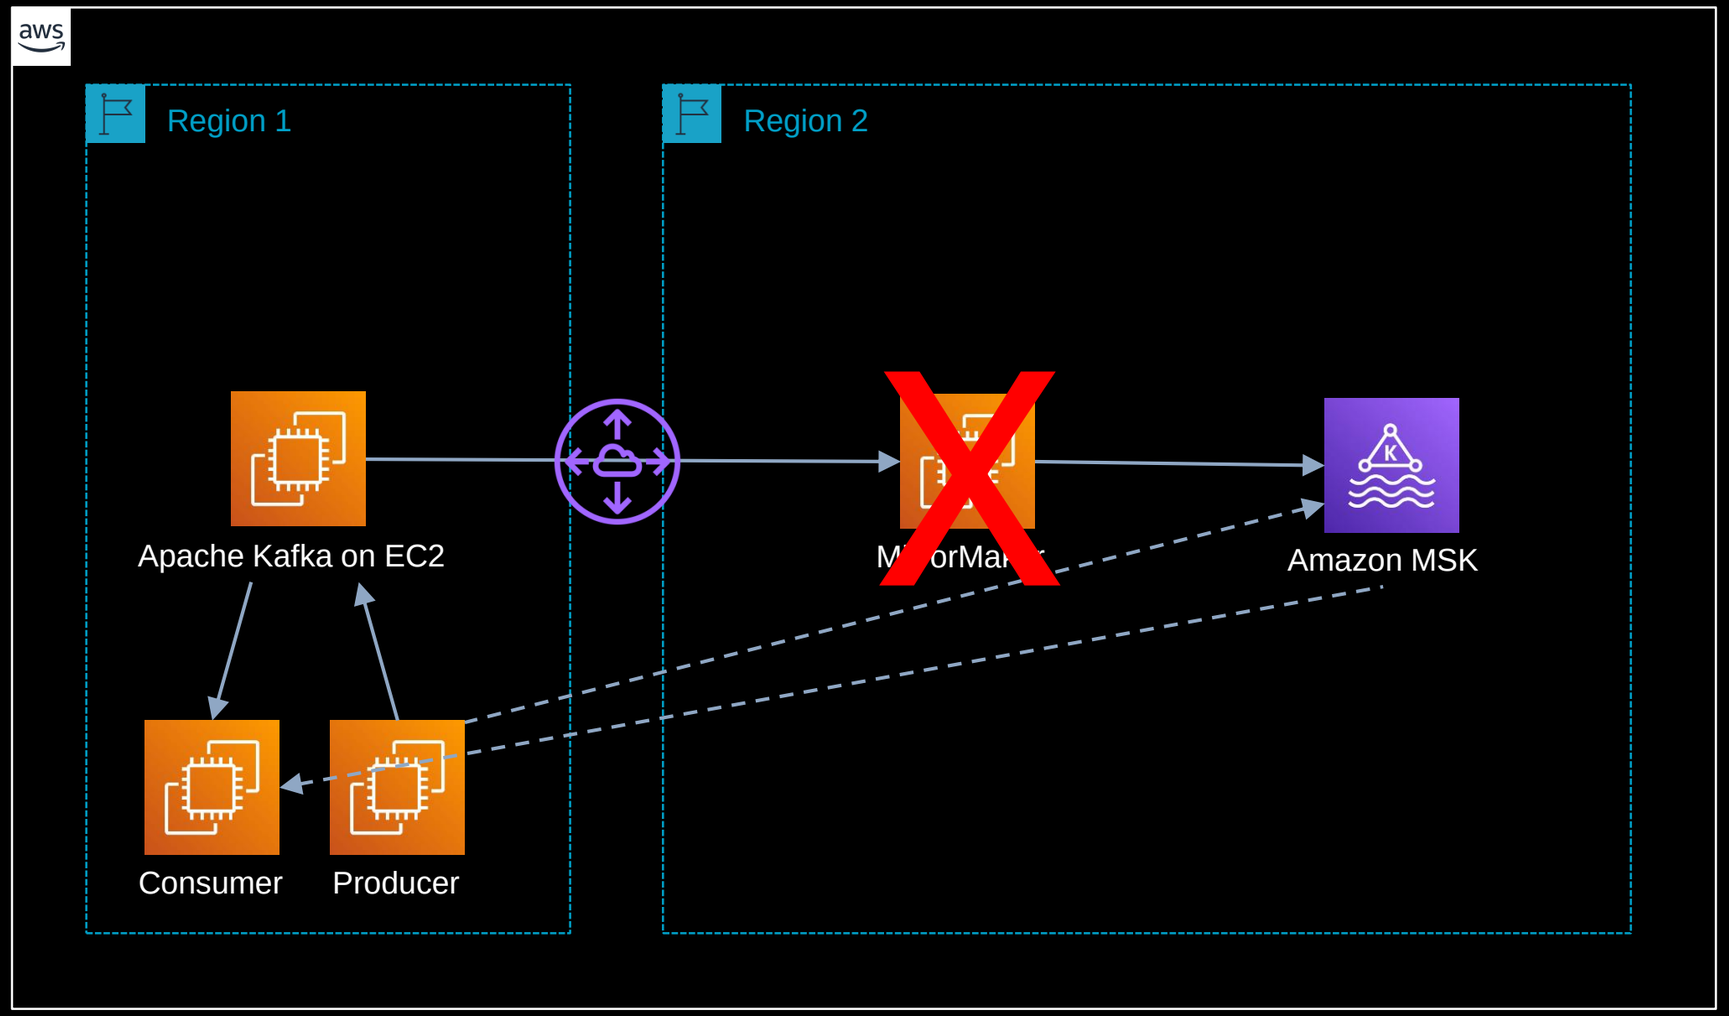
- Dedicated MirrorMaker 2.0 cluster
- Connector in a distributed Connect cluster
- As a standalone Connect worker
- In legacy mode



© 2020, Amazon Web Services, Inc. or its affiliates. All rights reserved.

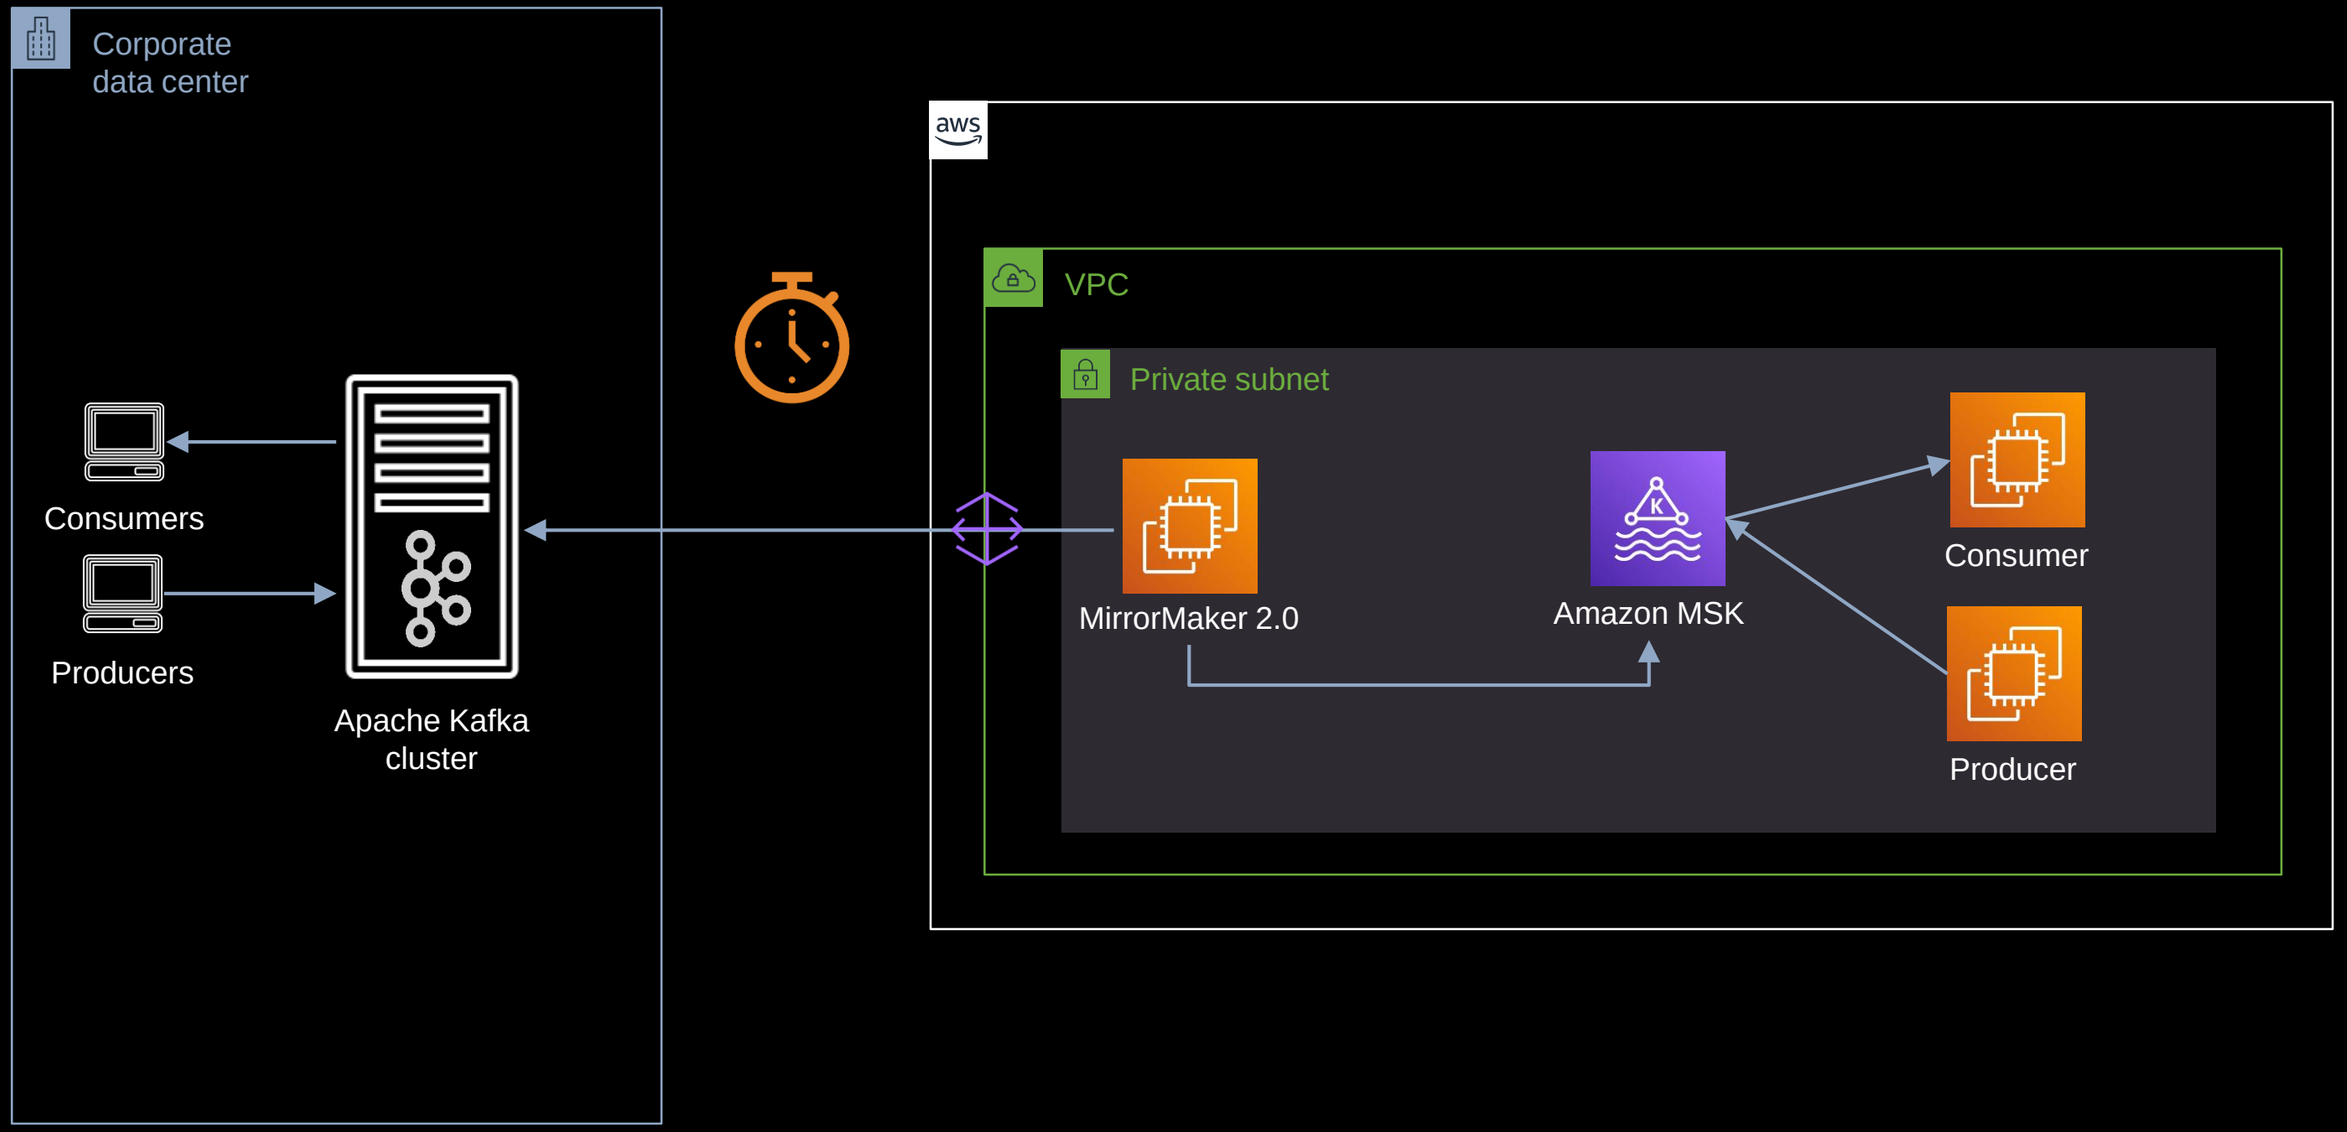
No retention migration

WITH CONSUMER AND PRODUCER MIGRATION



Migration with minimal downtime

WITH CONSUMER AND PRODUCER MIGRATION



Migration using Apache Flink



AWS re:Invent

© 2020, Amazon Web Services, Inc. or its affiliates. All rights reserved.

Apache Flink vs. MirrorMaker

APACHE FLINK

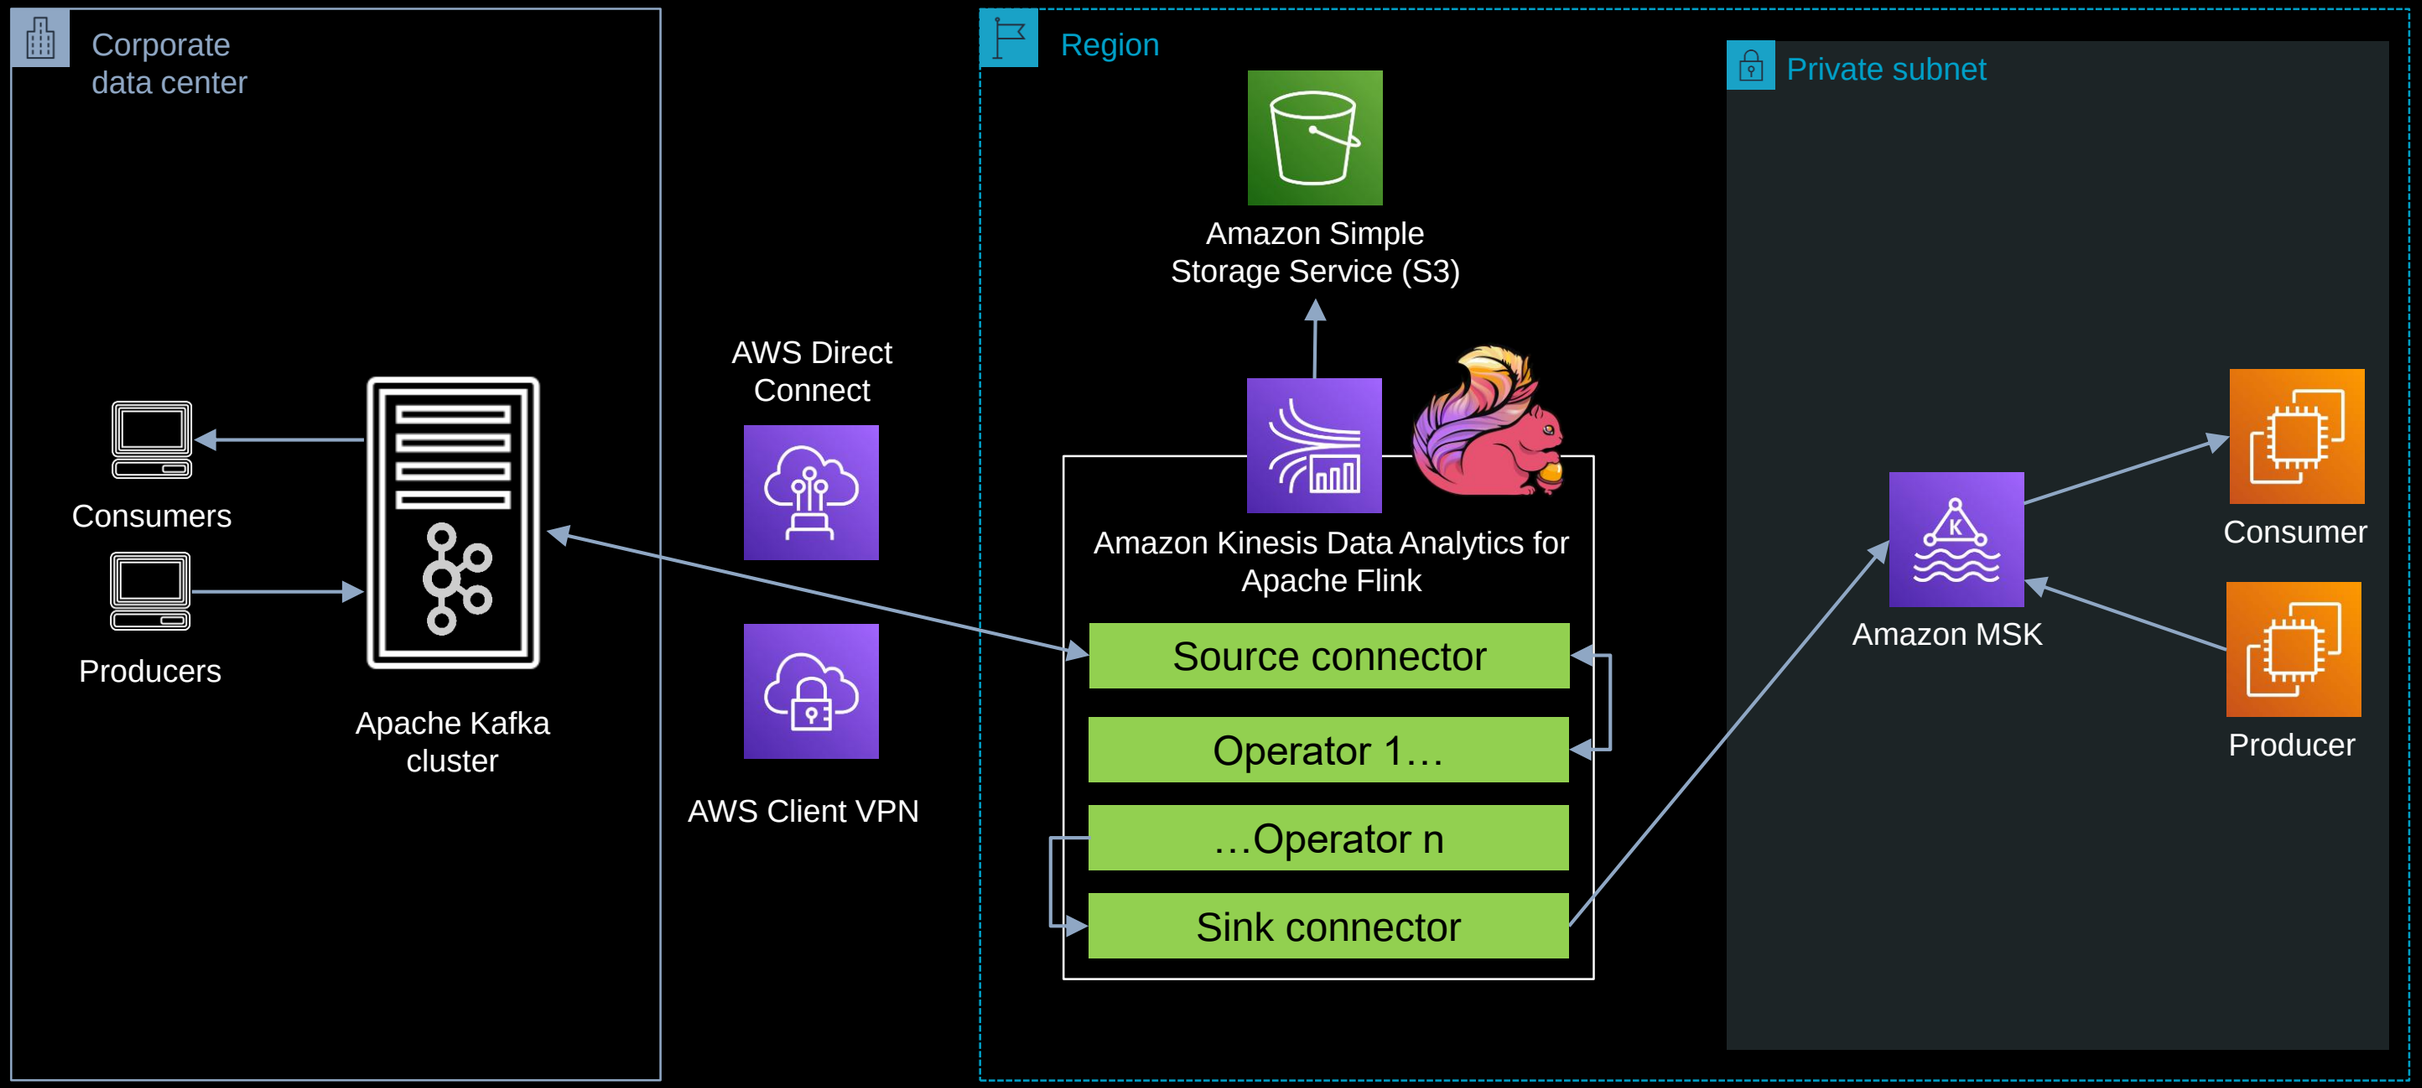
- Streaming data is able to be enriched and/or transformed between your source and destination clusters
- On Amazon Kinesis Data Analytics, Apache Flink applications are serverless so teams don’t have to manage underlying infrastructure
- Apache Flink on Data Analytics has the ability to auto-scale to achieve increased throughput requirements of your migration
- Support for multiple different types of sinks and connectors outside of Apache Kafka
- Supports exactly-once semantics
- External tools required for migrating consumer offsets properly, schema, etc.

MIRRORMAKER

- MirrorMaker 2.0 can be run in different clusters, can use Apache Kafka Connect via connector and run it in legacy mode (MM1)
- MirrorMaker 2.0 supports at-least-once semantics
- In addition to data migration, MirrorMaker 2.0 supports migration of ACLs, topic configs, and offset translation
- MirrorMaker 2.0 supports bi-directional replication

Migration using Apache Flink

POWERED BY AMAZON KINESIS DATA ANALYTICS



Fully managed Apache Flink

POWERED BY AMAZON KINESIS DATA ANALYTICS

Quickly build and easily run sophisticated streaming applications



Real-time processing
Millisecond latency



Fully managed Java and SQL apps
Near-zero ops overhead



Automatic elasticity
Auto and provisioned scaling



Easy to use
Build POCs in minutes using open-source libraries



Integrated applications
Easily integrate with AWS services



Pay for what you use
Fine-grained billing closely follows what you use

Amazon MSK migration and capacity planning best practices



AWS re:Invent

© 2020, Amazon Web Services, Inc. or its affiliates. All rights reserved.

Capacity planning best practices

- Two primary decisions required in terms of capacity planning
  - Instance type of the broker
    - Amazon MSK supports M-family and T-family of Amazon EC2 instances
    - M5 brokers can have more partitions per broker than T3
    - M5 can be used if running larger production-grade workloads or if you require a greater number of partitions
  - Capacity planning for Amazon MSK
    - Use the MSK sizing and pricing spreadsheet to determine the correct number of brokers for your MSK cluster
    - After you've decided on the number of brokers, test to validate assumptions

MSK sizing and pricing guide

https://bit.ly/2m3hlwz (downloads spreadsheet)

	Input Cells
CLUSTER DESIGN INPUT PARAMETERS	
Ingestion Rate, MB/s	1.0
Replication Factor	3
Total Data Out, MB/s	1.0
Retention, Hrs	24
Encryption in-transit	FALSE
EC2 Comparison Inputs	
Engineering Cost	$300,000
Number of Engineers	1
EC2 Reserved Instances (Years) **	1
Advanced parameters	
Number of AZs (3 Recommended)	3
Nearest Replica Fetching (Kafka 2.4.x)	FALSE
Lagging Consumer Percentage	20%
Broker Safety Factor	1
EBS Disk Utilization	60%
Ram Data Cached Seconds	90

Note: Hover over the input (red triangle) to get a description of the field.

MSK CLUSTER SIZING AND COSTS

	Recommended # of	Hourly Cost				Monthly Cost			
Instance Type	Brokers	Broker	Storage	Data Transfer*	Total	Broker+Storage	AZ Transfer	**Total**	$/GB
t3.small	**3**	$0.14	$0.06	$0.09	$0.19	$142	$68	**$210**	$0.080
m5.large	**3**	$0.63	$0.06	$0.09	$0.69	$502	$68	**$571**	$0.217
m5.xlarge	**3**	$1.26	$0.06	$0.09	$1.32	$962	$68	**$1,030**	$0.392
m5.2xlarge	**3**	$2.52	$0.06	$0.09	$2.58	$1,882	$68	**$1,950**	$0.742
m5.4xlarge	**3**	$5.04	$0.06	$0.09	$5.10	$3,721	$68	**$3,790**	$1.442
m5.8xlarge	**3**	$10.08	$0.06	$0.09	$10.14	$7,401	$68	**$7,469**	$2.842
m5.12xlarge	**3**	$15.12	$0.06	$0.09	$15.18	$11,080	$68	**$11,148**	$4.242
m5.16xlarge	**3**	$20.16	$0.06	$0.09	$20.22	$14,759	$68	**$14,827**	$5.642
m5.24xlarge	**3**	$30.24	$0.06	$0.09	$30.30	$22,117	$68	**$22,186**	$8.442

us-east-1 pricing *Cross AZ Costs

EC2 SELF MANAGED KAFKA COSTS

	Recommended # of	Hourly Cost						Hourly		Monthly		$/GB	
Instance Type	Brokers	Broker	Storage	Data Transfer*	Replication	Zookeeper	Engineers	Total	+ Devs	Total	+ Devs	Total	+ Devs
t3.small	**3**	$0.04	$0.06	$0.09	$0.14	$0.37	$34.25	$0.71	$34.95	$515	$25,515	0.196	9.709
m5.large	**3**	$0.18	$0.06	$0.09	$0.14	$0.37	$34.25	$0.85	$35.09	$618	$25,618	0.235	9.748
m5.xlarge	**3**	$0.36	$0.06	$0.09	$0.14	$0.37	$34.25	$1.03	$35.28	$752	$25,752	0.286	9.799
m5.2xlarge	**3**	$0.73	$0.06	$0.09	$0.14	$0.37	$34.25	$1.39	$35.64	$1,017	$26,017	0.387	9.900
m5.4xlarge	**3**	$1.45	$0.06	$0.09	$0.14	$0.37	$34.25	$2.12	$36.37	$1,547	$26,547	0.589	10.102
m5.8xlarge	**3**	$2.90	$0.06	$0.09	$0.14	$0.37	$34.25	$3.57	$37.82	$2,607	$27,607	0.992	10.505
m5.12xlarge	**3**	$4.36	$0.06	$0.09	$0.14	$0.37	$34.25	$5.02	$39.27	$3,667	$28,667	1.395	10.908
m5.16xlarge	**3**	$5.81	$0.06	$0.09	$0.14	$0.37	$34.25	$6.47	$40.72	$4,725	$29,725	1.798	11.311
m5.24xlarge	**3**	$8.71	$0.06	$0.09	$0.14	$0.37	$34.25	$9.38	$43.62	$6,845	$31,845	2.604	12.117

us-east-1 pricing *Cross AZ Costs

SUMMARY

Selected Cluster Instance Type	m5.large	Drop down selection. Select instance based on the # of brokers column in MSK CLUSTER SIZING AND COSTS table
Cluster Size (Brokers)	3	We recommend referencing the MSK default limits and requesting a limit increase if needed.
Provisioned Storage per Broker (GB)	141	

Cluster	Monthly Cost	$/GB	% Difference to MSK
Amazon MSK	$571	$0.217	-
Self Managed EC2	$618	$0.235	8%
Self Managed EC2 with Engineers	$25,618	$9.748	4390%

Sizing guidance in minutes



Up to 40% lower cost compared to self-managed Apache Kafka

The calculator estimates your workload needs

Costs are normalized to $/GB ingested to make pricing comparisons easy

Your latency requirements will impact how much you should provision, adjust accordingly

We recommend that you test your workloads and scale accordingly

Special thanks

Rajeev Chakrabarti

Principal Developer Advocate

Thank you!



AWS re:Invent

© 2020, Amazon Web Services, Inc. or its affiliates. All rights reserved.



Please complete the session survey

AWS re:Invent

© 2020, Amazon Web Services, Inc. or its affiliates. All rights reserved.