

The background is a vibrant, multi-colored gradient. It features a diagonal split between a blue-purple gradient on the left and an orange-yellow gradient on the right. The text 'AWS re:Invent' is rendered in white, with 'AWS' in a smaller font size above 're:Invent'.

AWS
re:Invent

MFG202

Building on AWS: The architecture of the Siemens MindSphere platform

Kathleen DeValk

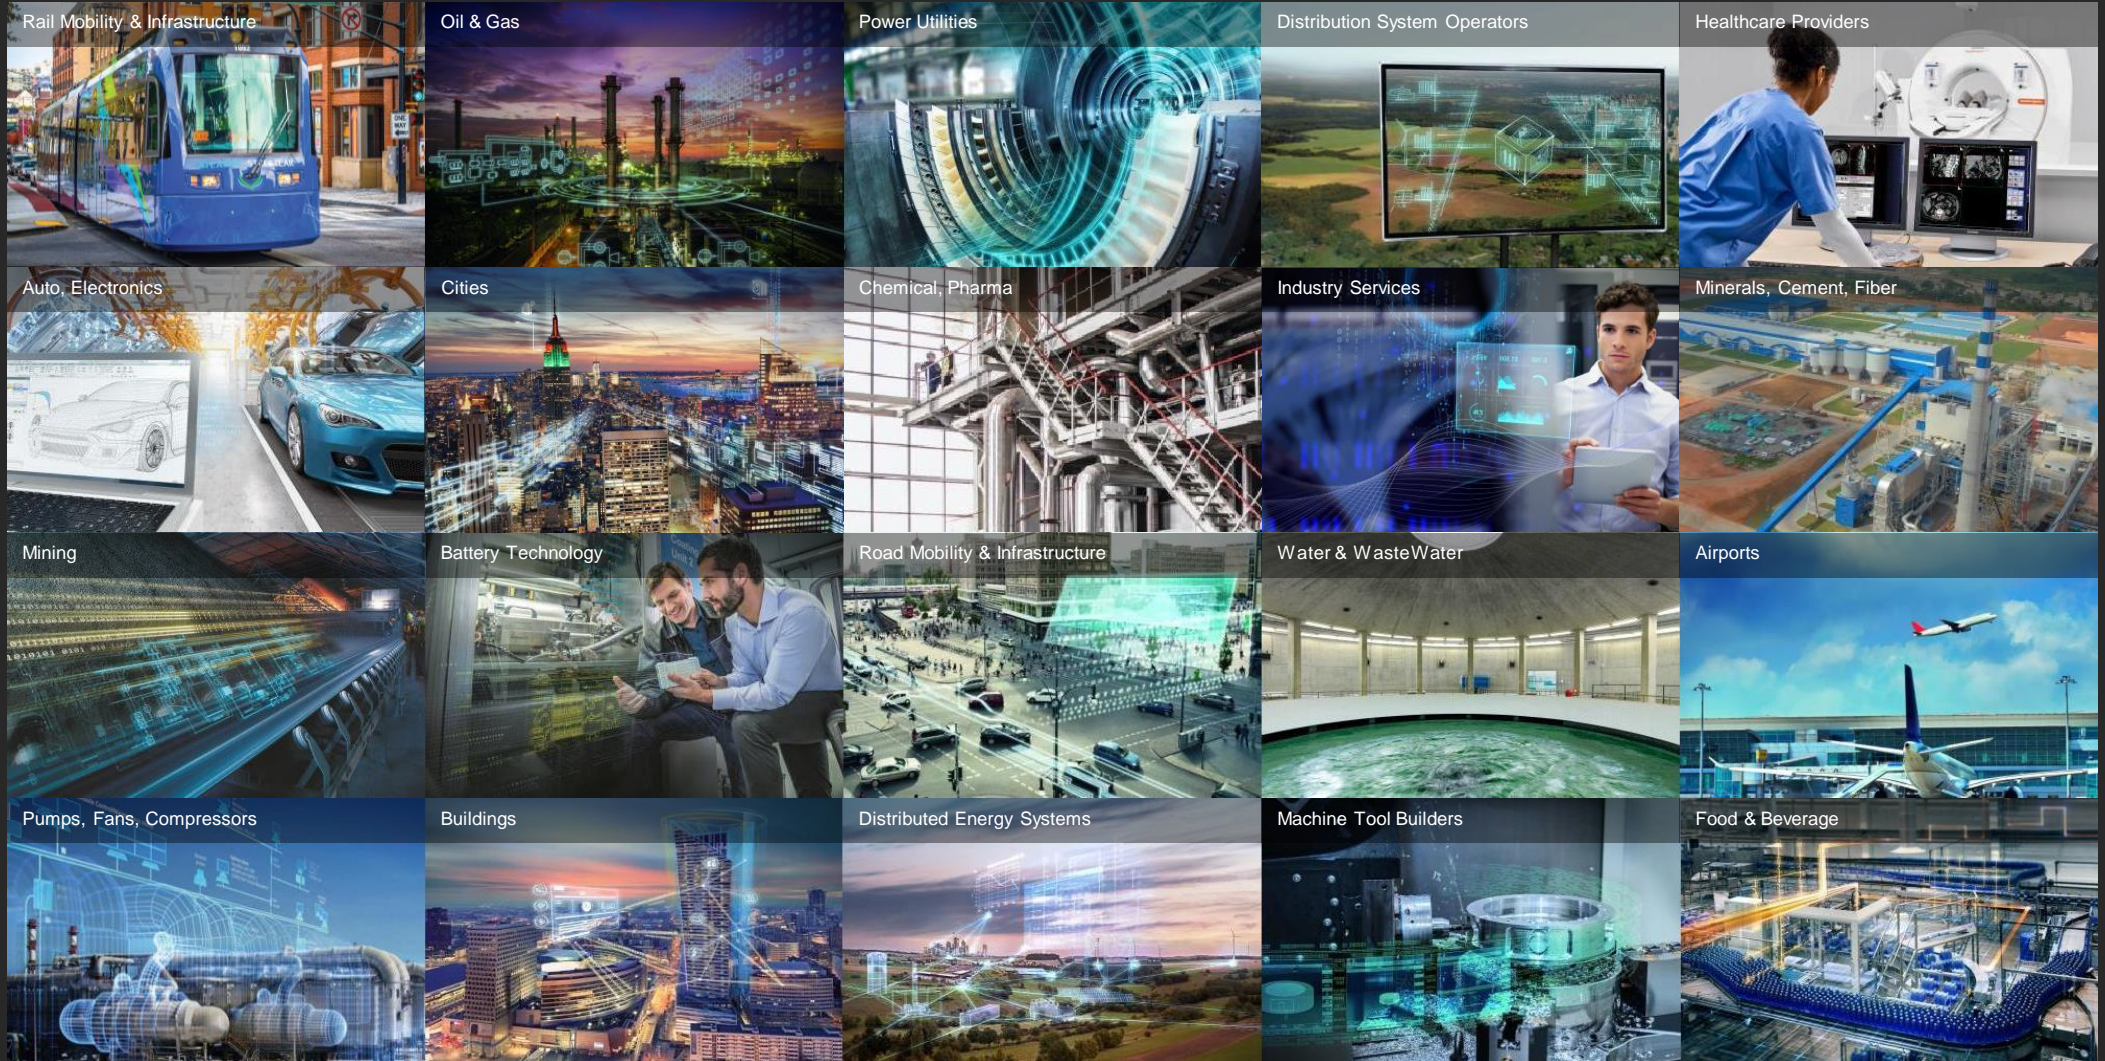
Chief Architect
Siemens MindSphere
Siemens

Frank Kovacs

AWS Principal PDM – Siemens
Amazon Web Services

Siemens

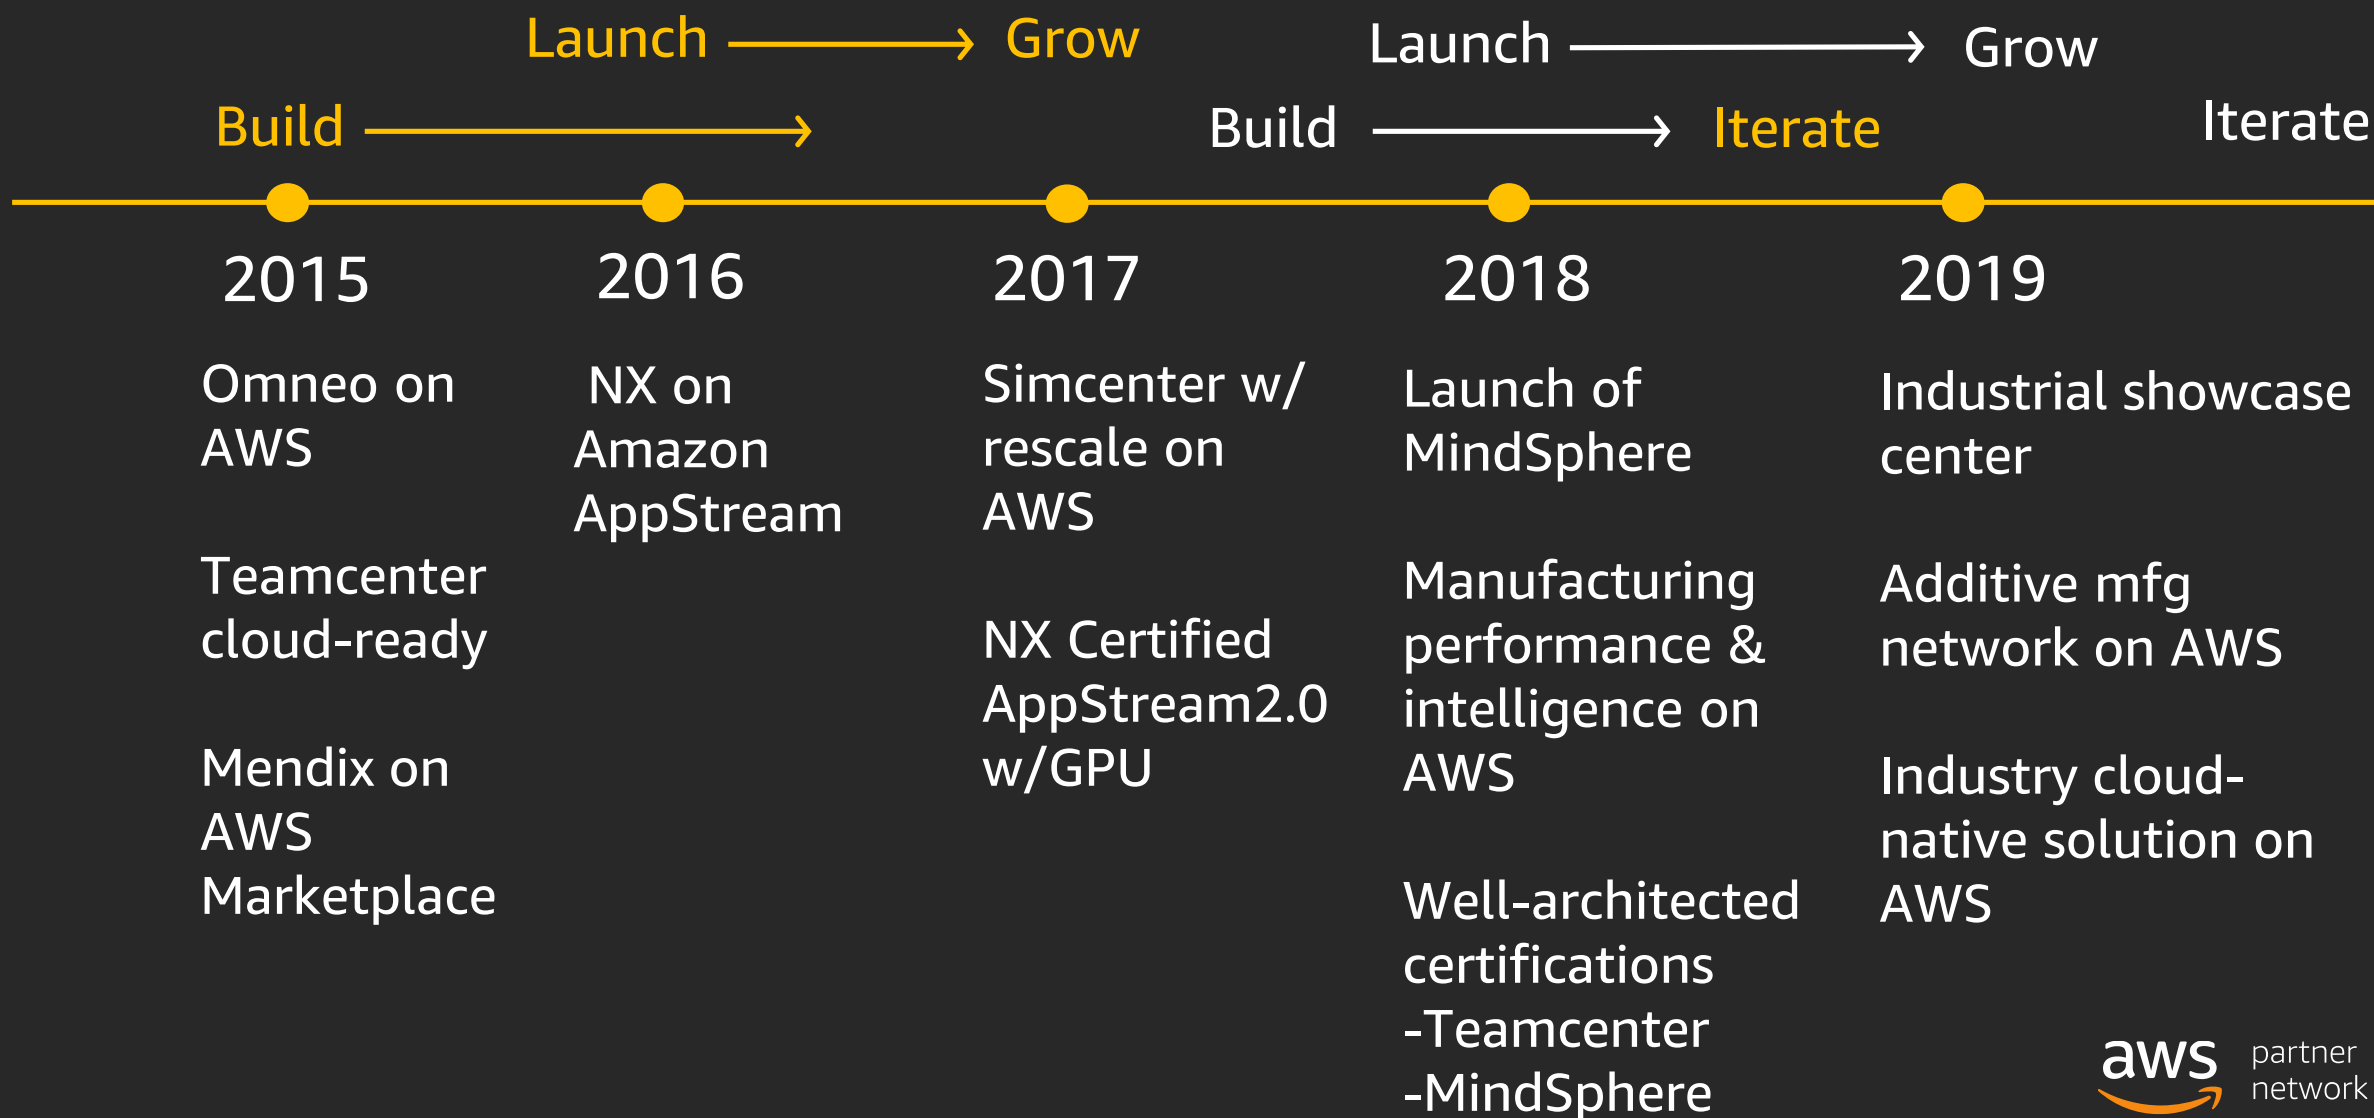
The world leader in industrial automation



Customers move from ideas to implementation easier, faster, more securely, and at a lower cost with AWS



Siemens DISW and AWS Cloud journey




AWS Industrial Software Competency



- Partners who have demonstrated technical expertise through rigorous technical validation and proven customer success
- Siemens MindSphere and Teamcenter were both recognized for gaining Industrial Competency status
- Siemens is the only company to receive competency status for two solutions



Converging IT & OT to maximize value



**Accelerate
decision making**



**Operations
transparency**



**Rapid applications
bring new products**



**Optimize
products**

Over 475 partners




Millions of devices

400 applications



Hundreds of GB/
sec



Thousands of
tenants



Hundreds of TBs
of data

What is MindSphere?

Applications

MindSphere applications provide asset transparency and analytical insights into machines, plants, fleets, and systems

Open PaaS

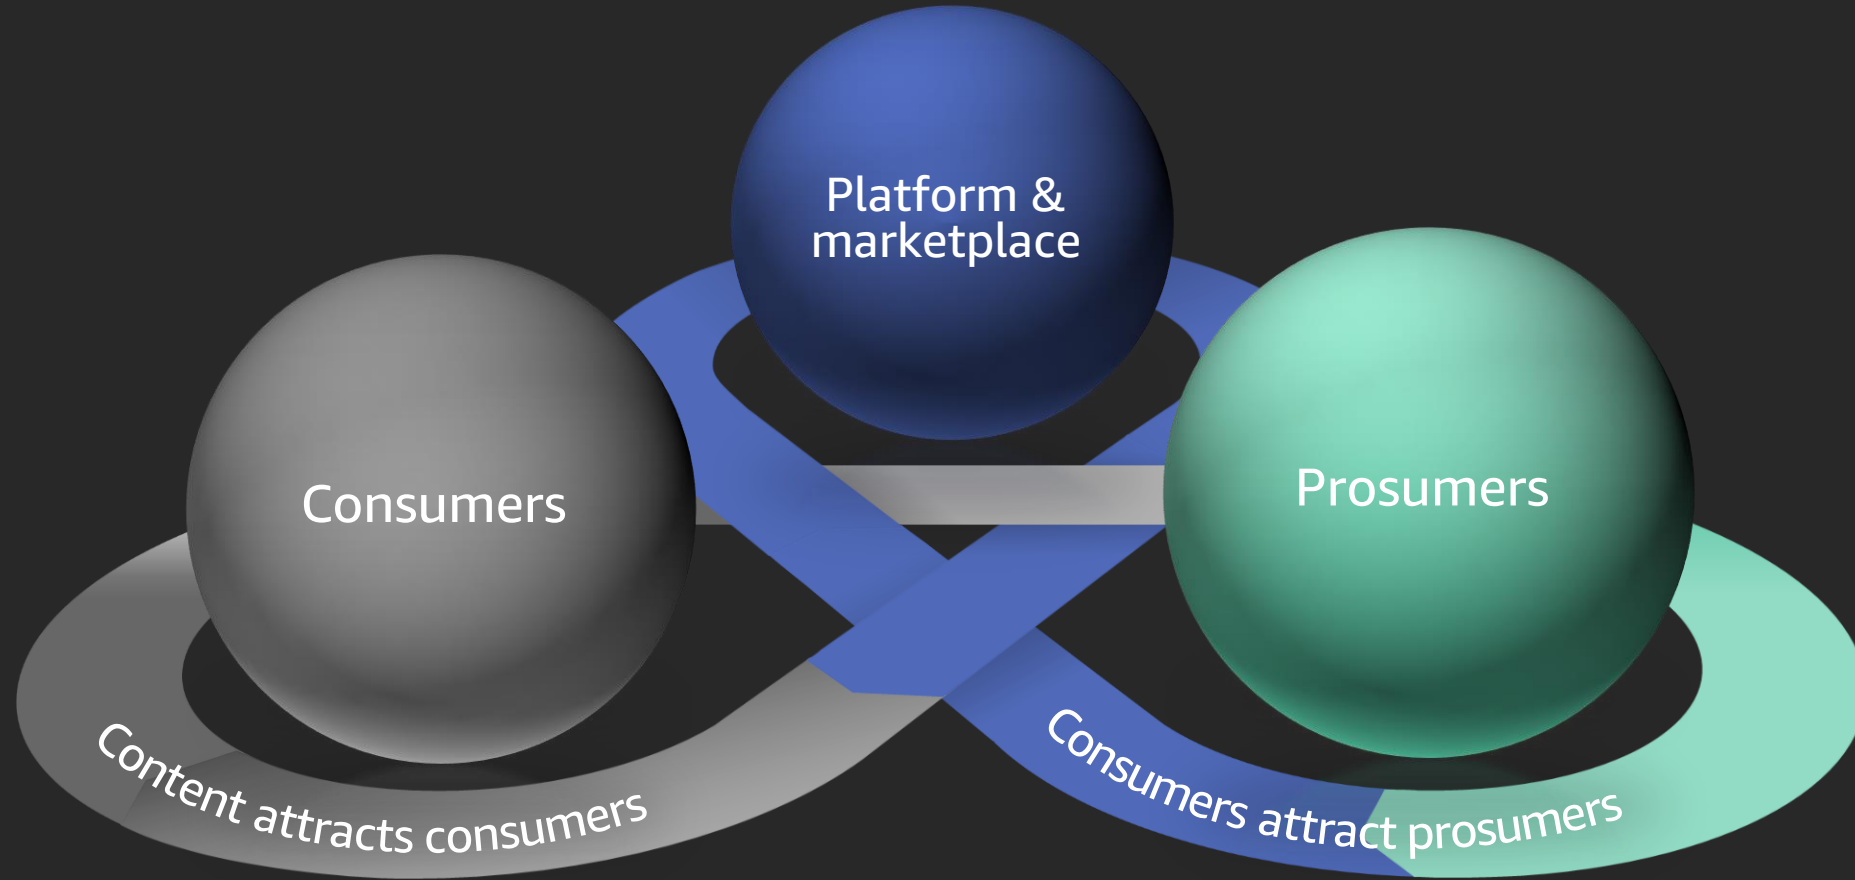
MindSphere, the open PaaS, gives you scalable, global IoT connectivity and application development with native cloud accessibility

Connectivity

MindConnect connects products, plants, systems, machines, enterprise applications, and legacy DB via a secured plug-and-play collection of products and equipment



Industrial ecosystem



MindSphere ecosystem: Customers, partners, developers

MindSphere world

Growing worldwide



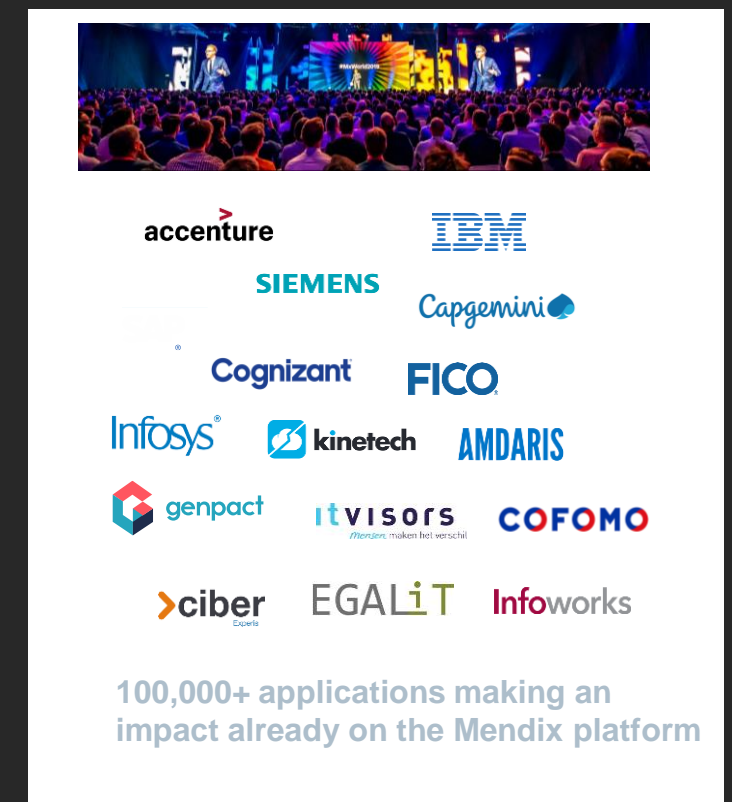
Global partner ecosystem

Over 475 and growing

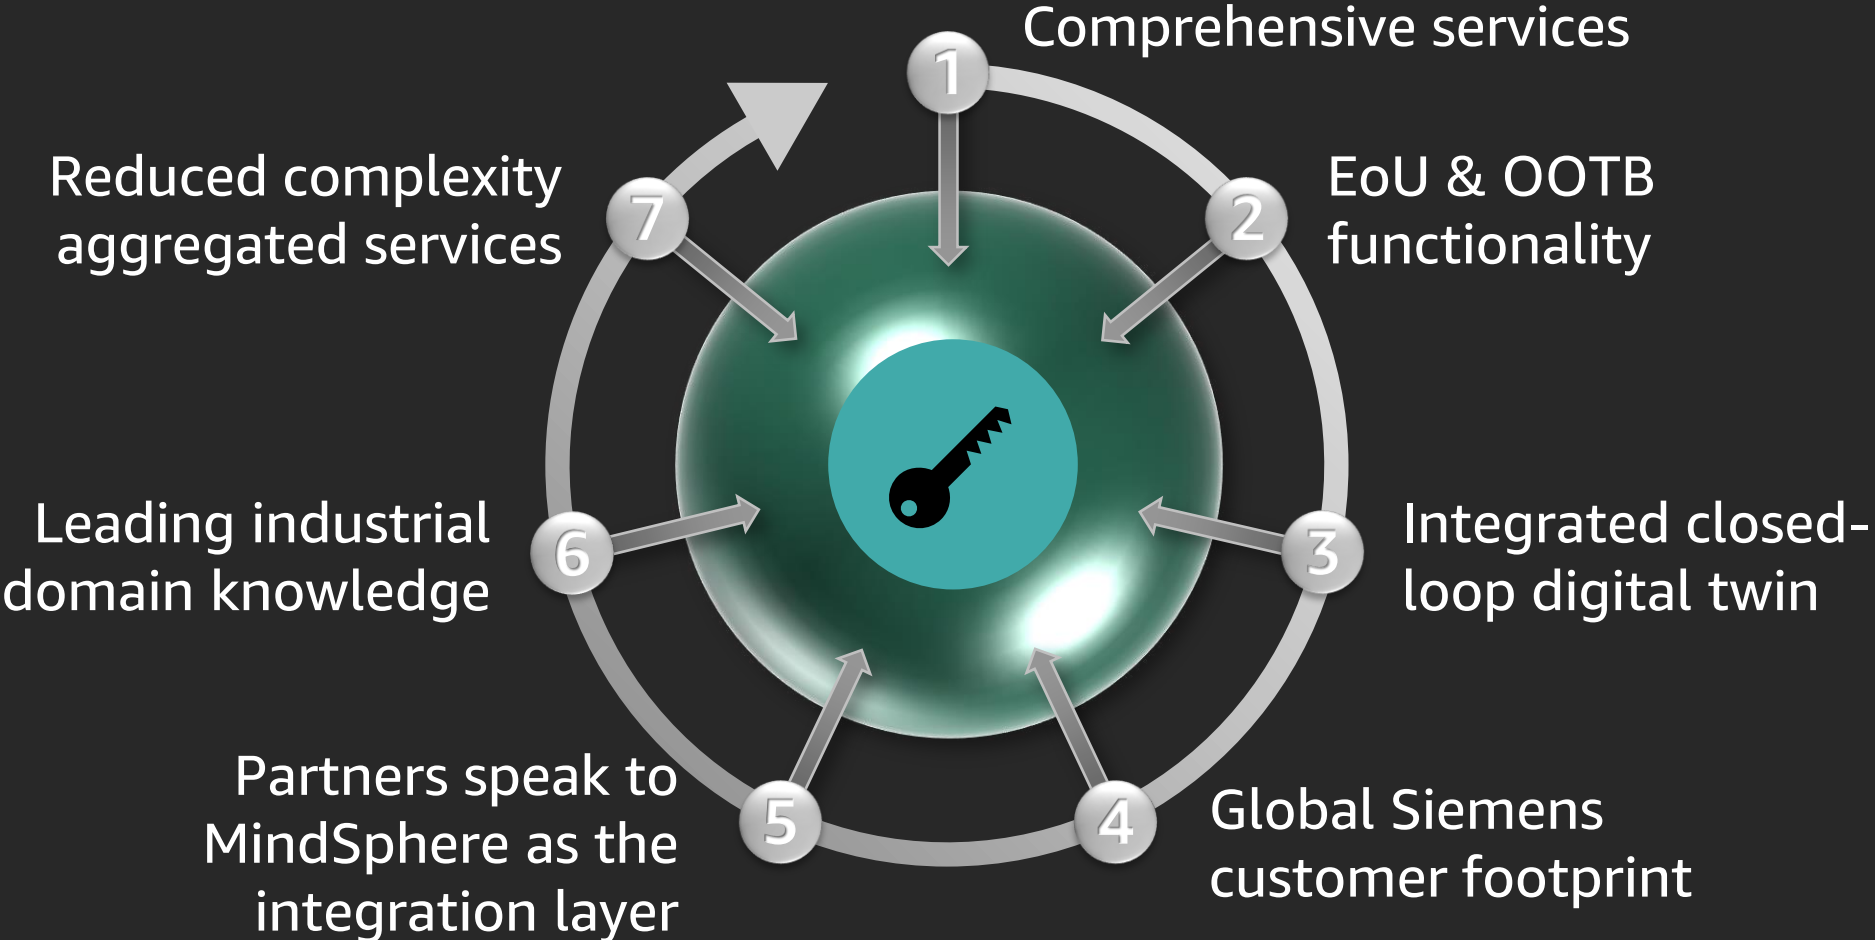


Mendix ecosystem

>90,000 developers



Know your differentiators



Approach: Cloud-optimized

“We don’t want to manage anything we don’t have to manage.”

- Consensus of MindSphere architecture team

“We want to ride AWS’s innovation curve, not fight it.”

- David Mitchell, VP MindSphere Products

Managed services

Fully managed
Automatic data replication to 3 AZs
Easy split/merge shards

Self-managed
Monitor & scale
Manage backups

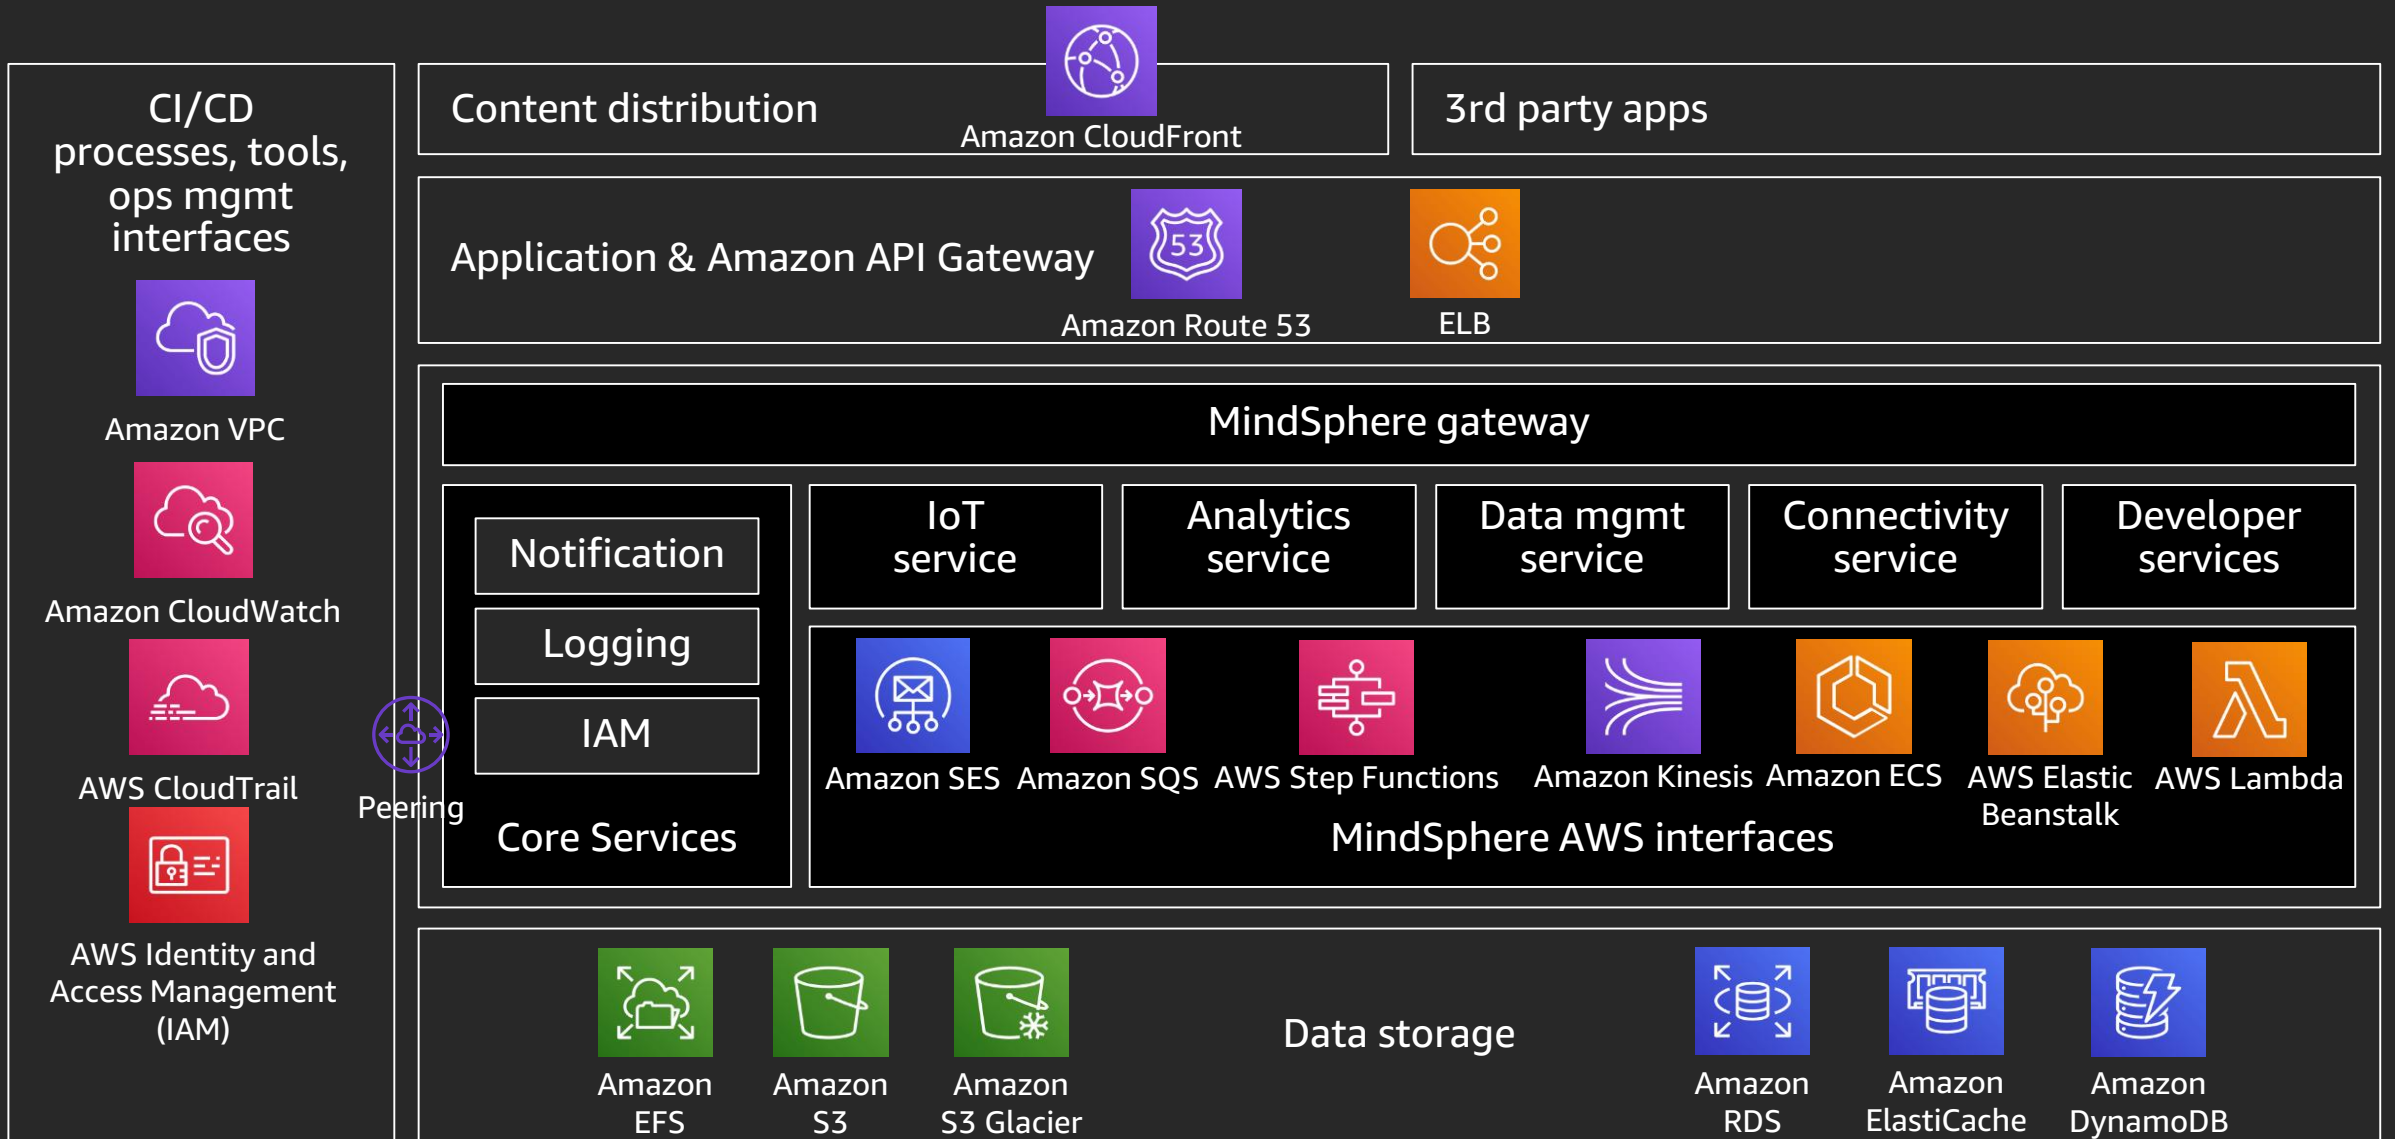


**Apache
Kafka**



**Amazon
Kinesis**

MindSphere on AWS overview



Services in use in 2018

Amazon ECS	AWS Elastic Beanstalk	Amazon CloudWatch
Amazon CloudFront	Application Load Balancer	AWS CloudTrail
Amazon DynamoDB	Amazon API Gateway	
Amazon EBS	Amazon S3	AWS Direct Connect
Amazon EC2		AWS KMS
Amazon Kinesis		Amazon Registrar
AWS Lambda	Amazon ElastiCache	Amazon Route 53
Amazon RDS	Amazon EMR	AWS Secrets Manager
Amazon SES		Amazon Shield
Amazon SNS	Amazon VPC	Amazon Shield Advanced
Amazon SQS	AWS WAF	
AWS Step Functions		

Services in use now

Amazon ECR

Amazon ECS

Amazon Kinesis Data Firehose

Amazon CloudFront

Amazon DynamoDB

Amazon EBS

Amazon EC2

Amazon EKS

Amazon Kinesis

AWS Lambda

Amazon Neptune

Amazon RDS

Amazon SES

Amazon SNS

Amazon SQS

AWS Step Functions

AWS Elastic Beanstalk

Amazon EFS

Amazon WorkSpaces

Application Load Balancer

Amazon API Gateway

Amazon S3

Amazon Athena

AWS Batch

Amazon DocumentDB

Amazon ElastiCache

Amazon ES

Amazon EMR

AWS Glue

Amazon VPC

AWS WAF

AWS Budgets

Amazon CloudWatch

AWS CloudTrail

AWS Config

Amazon DMS

AWS Direct Connect

Amazon GuardDuty

AWS KMS

Amazon Registrar

Amazon Route 53

AWS Secrets Manager

Amazon Shield

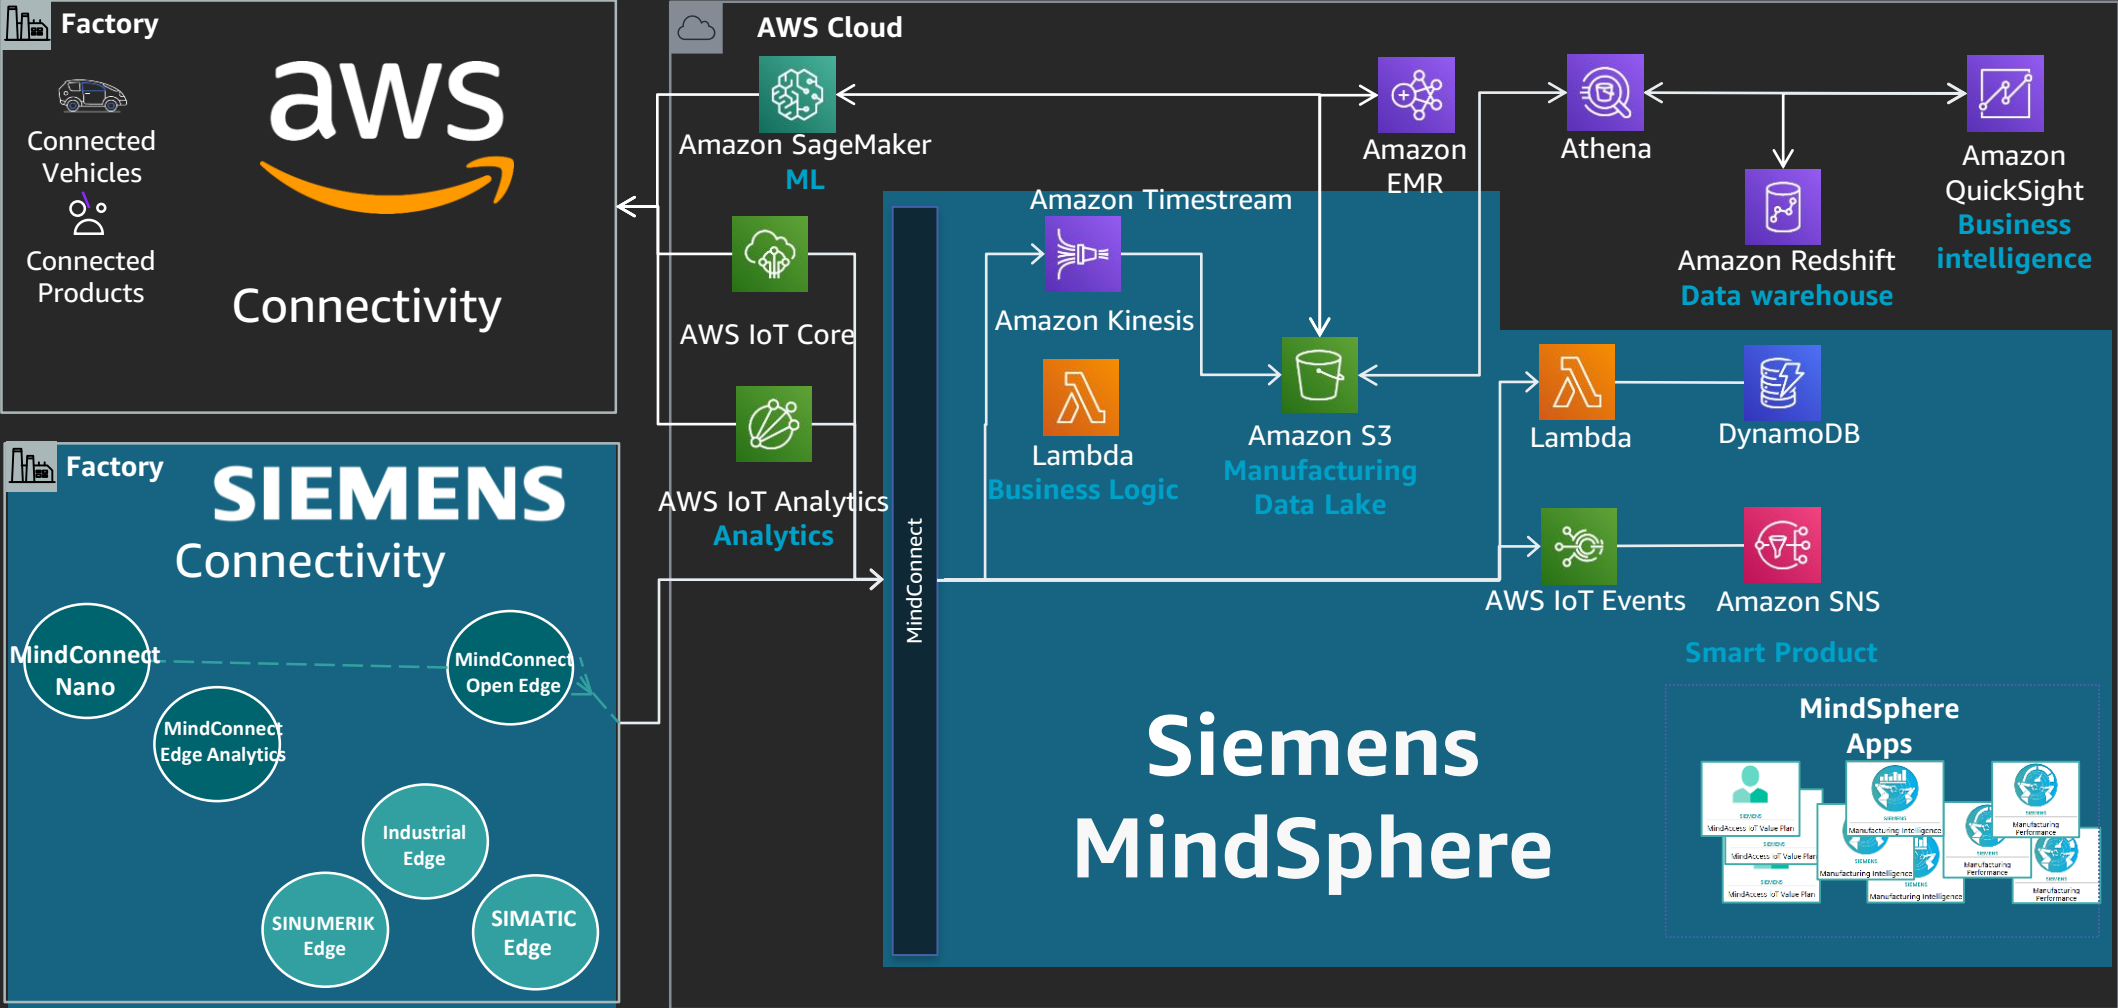
Amazon Shield Advanced

AWS Systems Manager

AWS Trusted Advisor

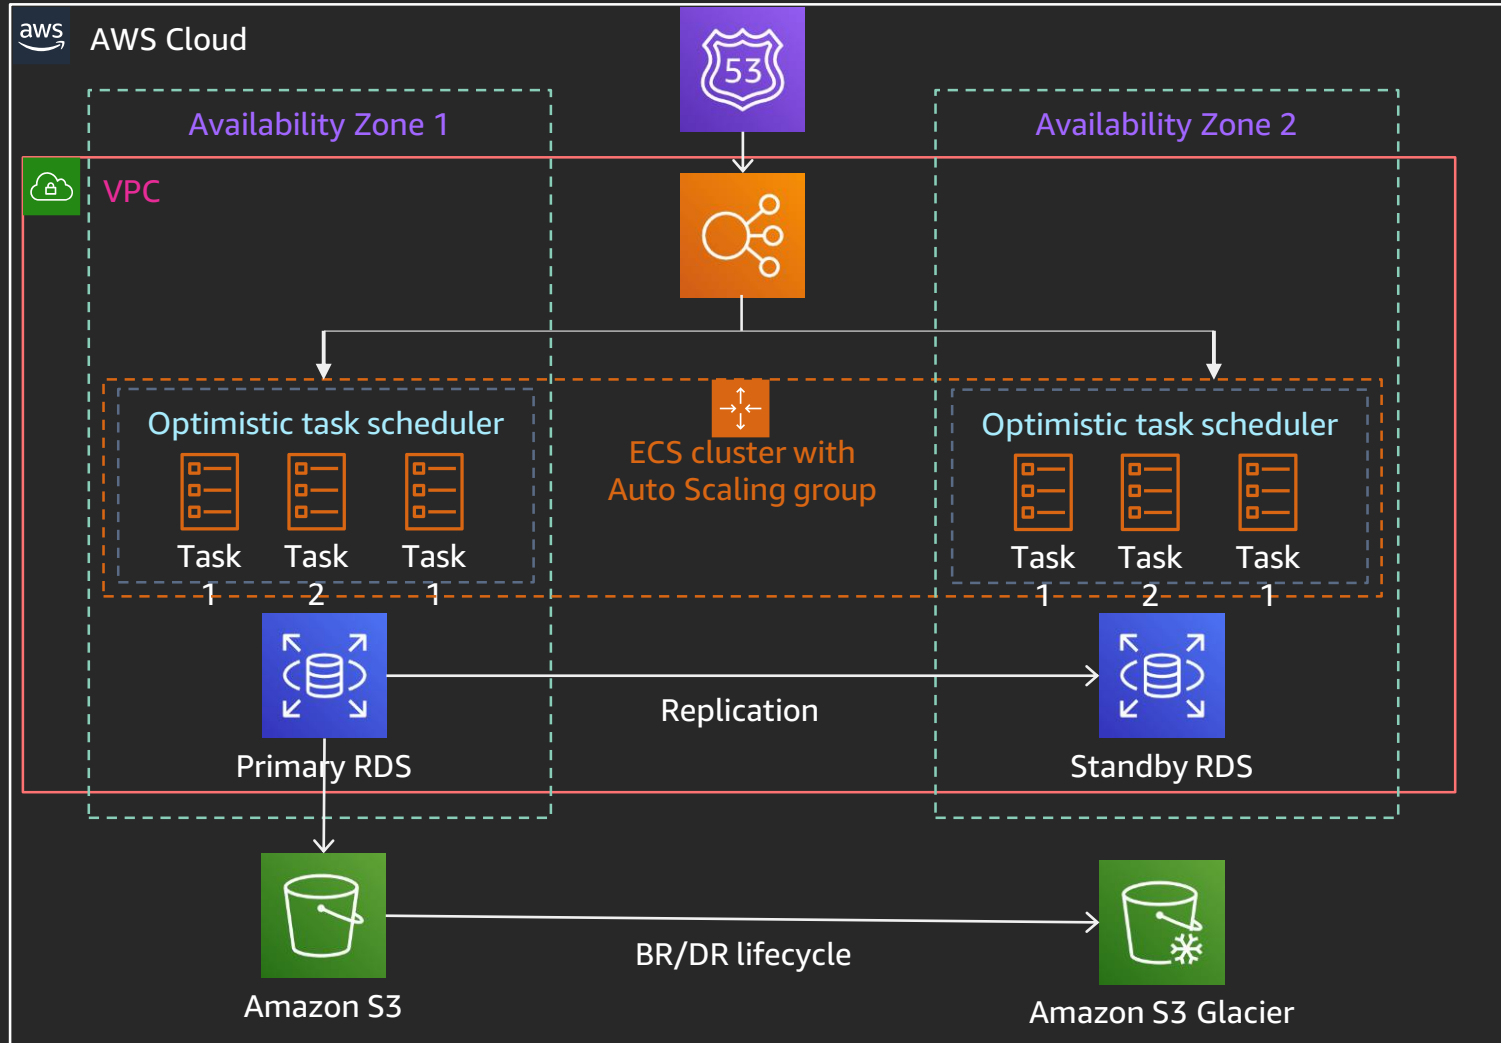
AWS X-Ray

MindSphere + AWS accelerates solutions



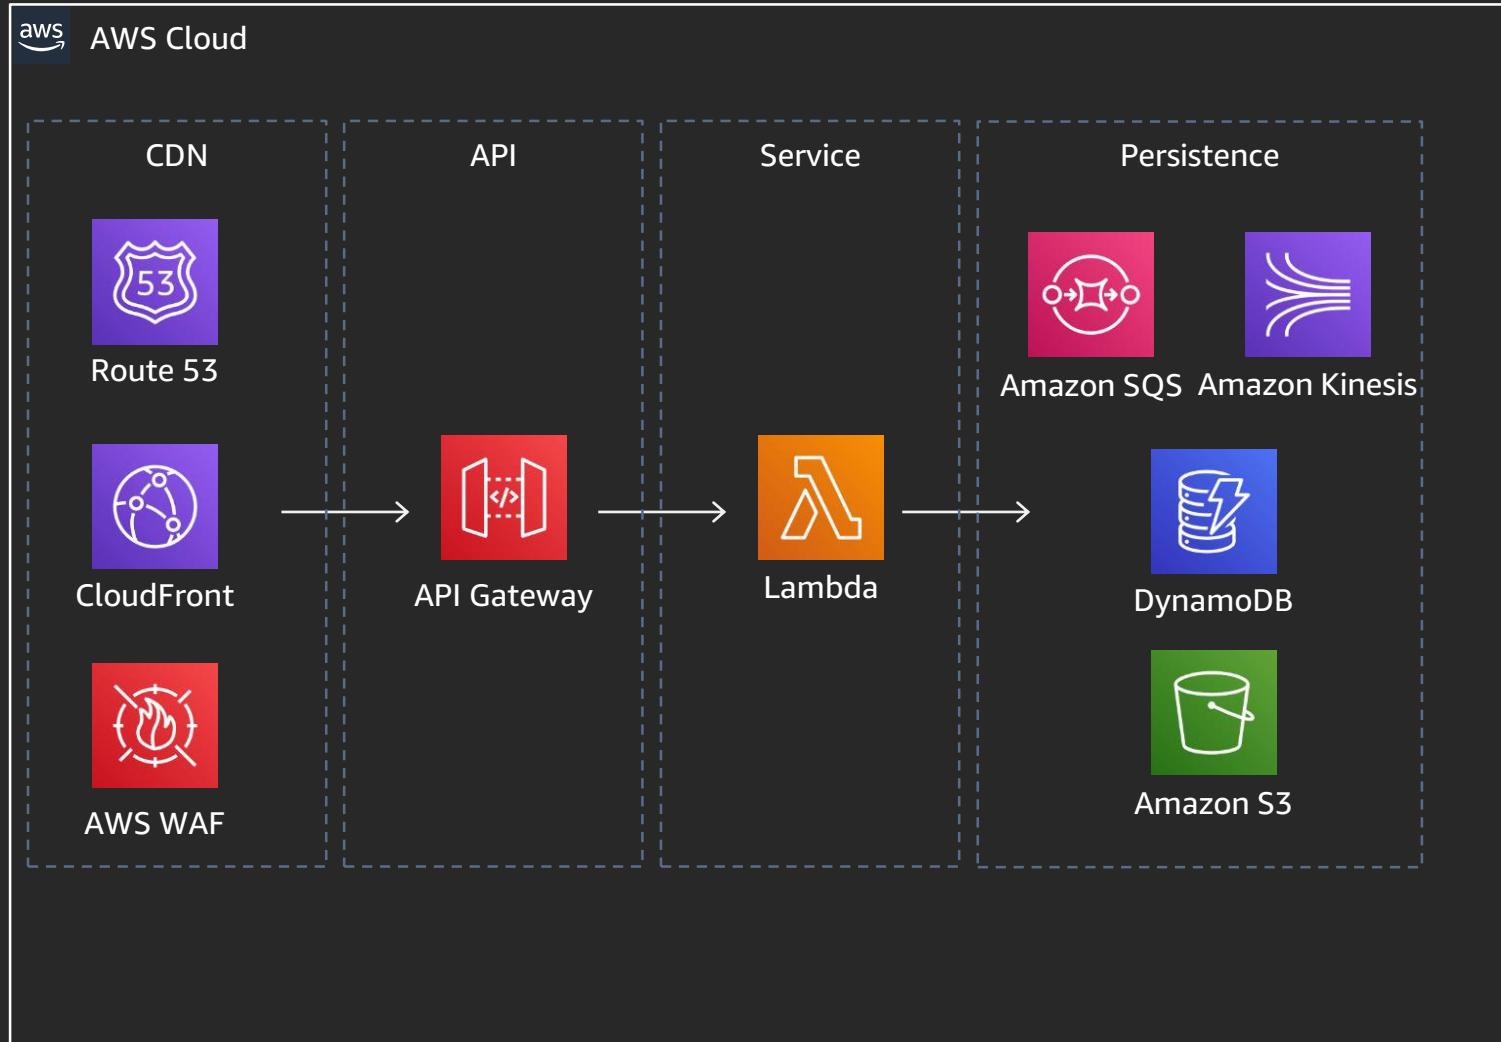
Design patterns

Design patterns: Containerized microservices



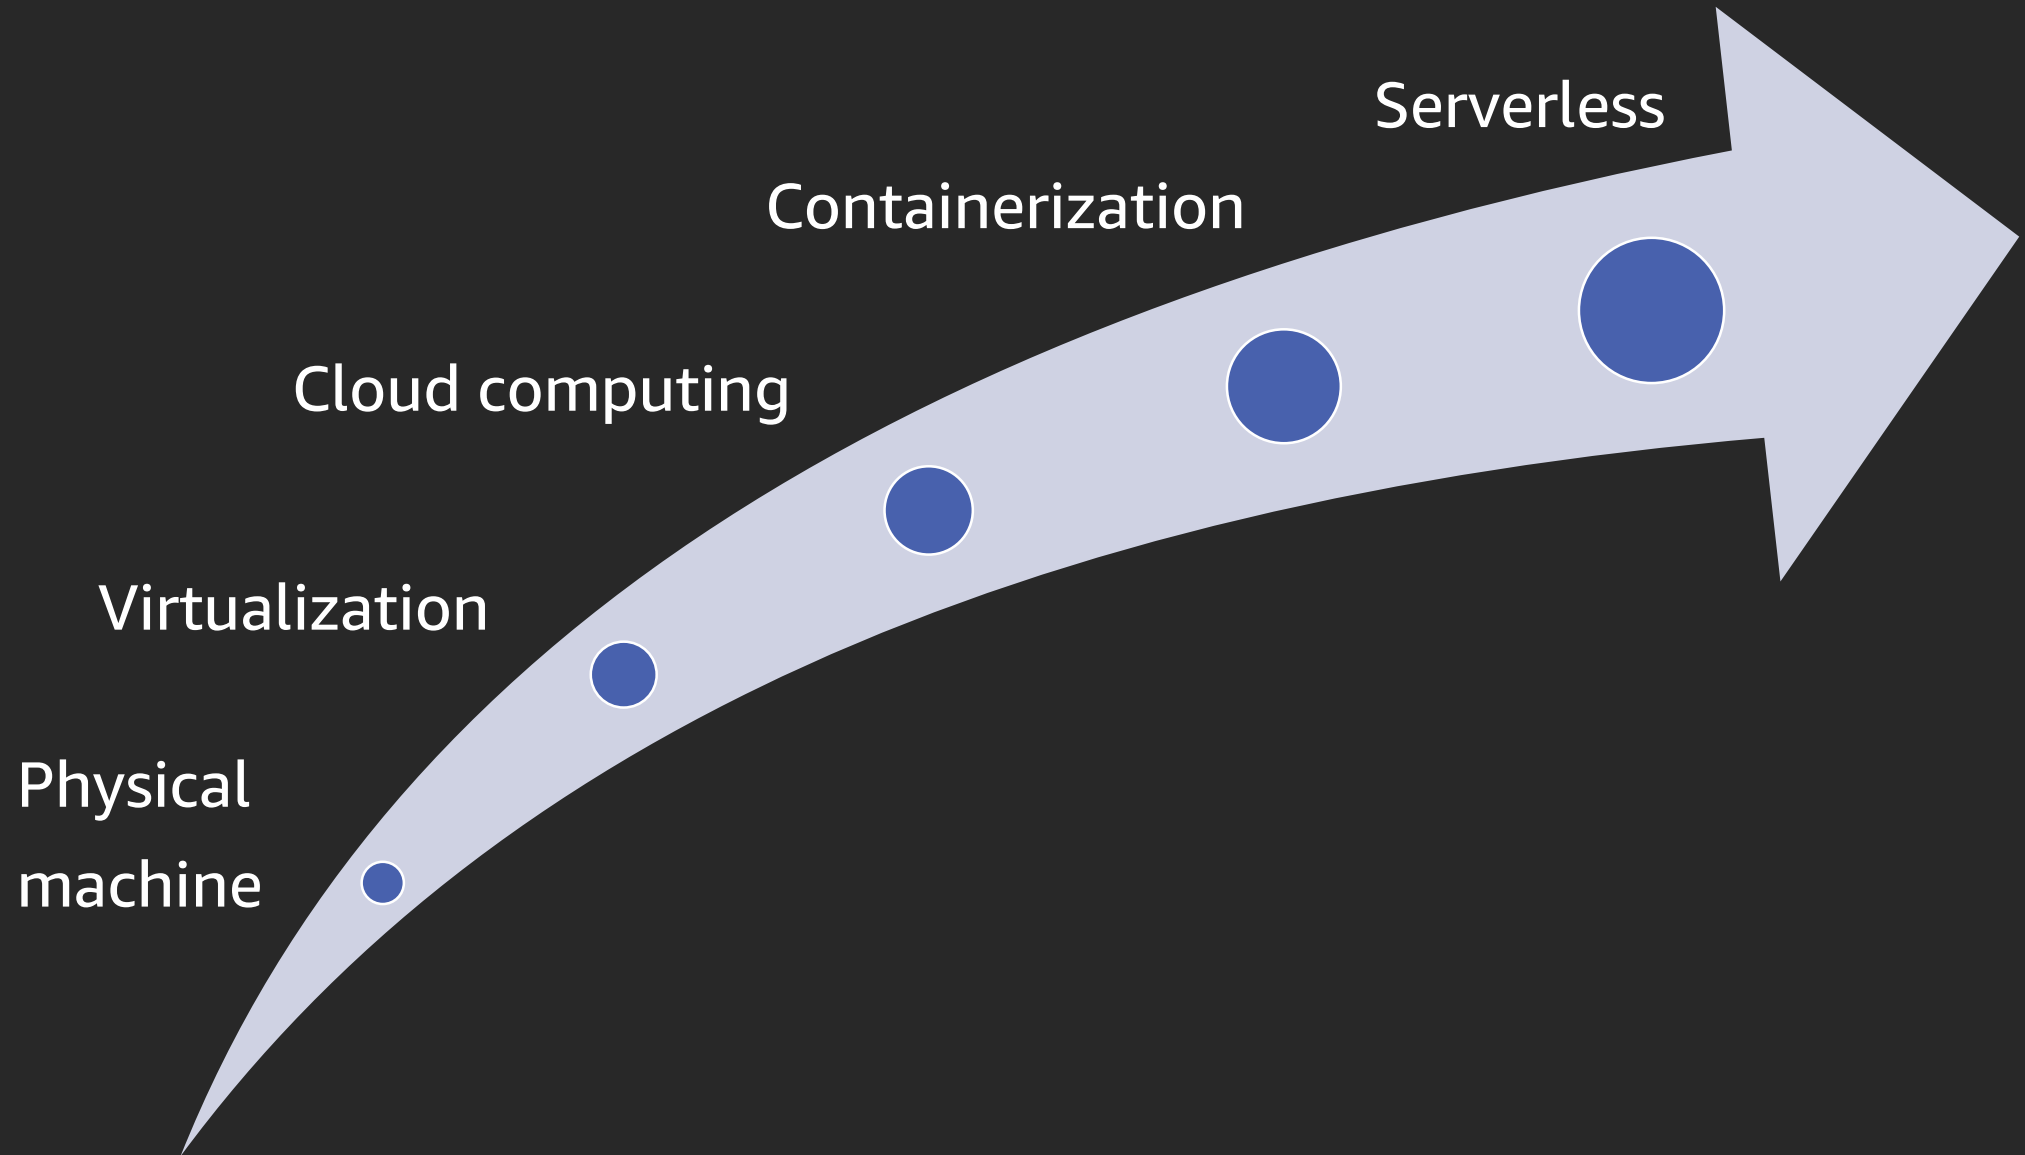
1	Resilient
2	Auto Scaling
3	Secure
4	Durable

Design patterns: Serverless



1	Asynchronous
2	Latency tolerant
3	Lower TCO - PAYGO
4	SLA + lower MGMT

Evolution of architectures resulting in increased productivity

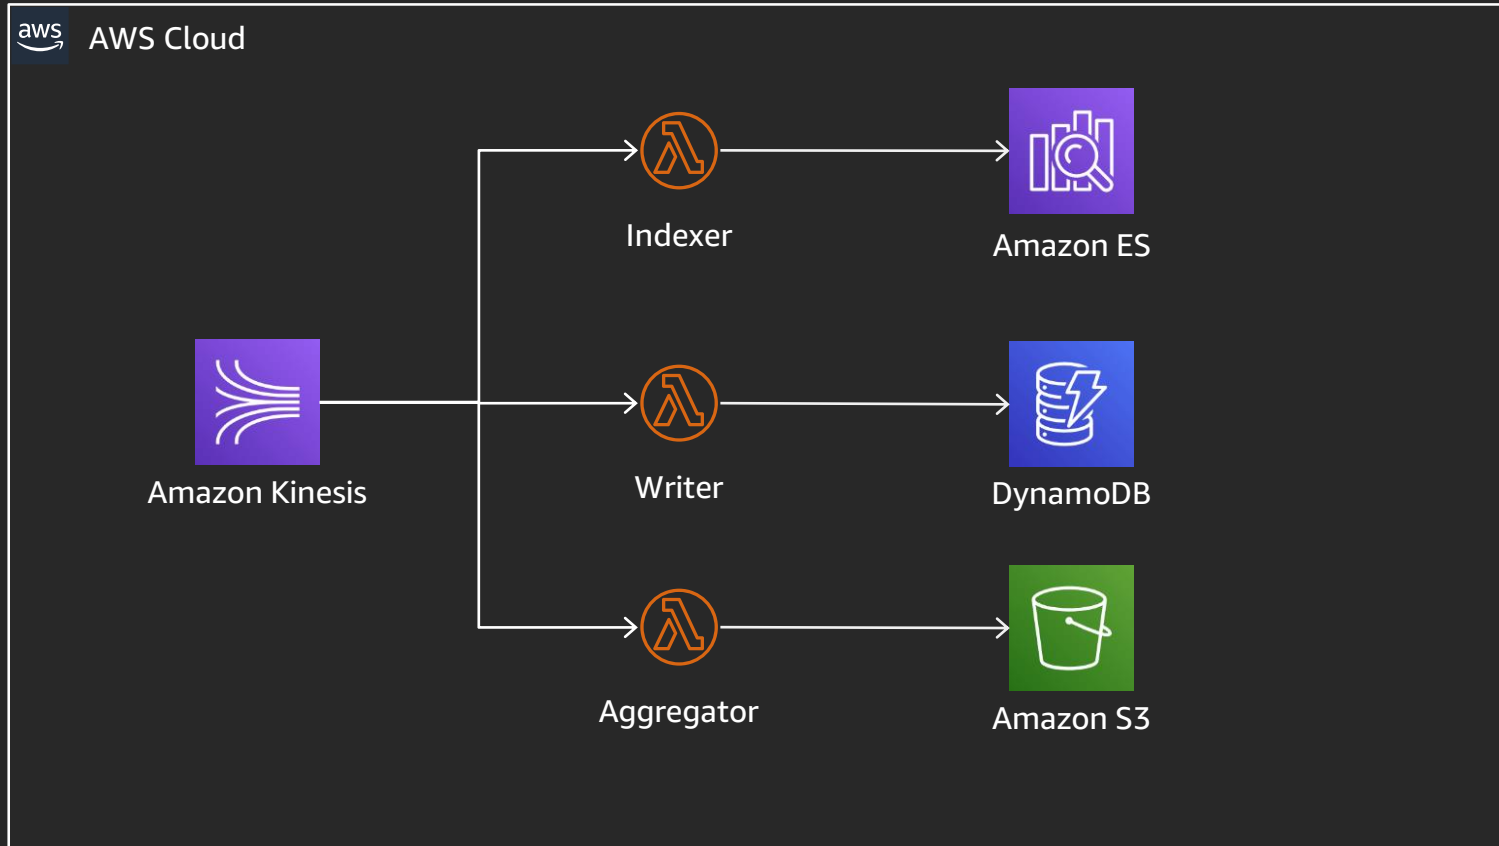


Data lake versus data swamp



Challenges running @ scale

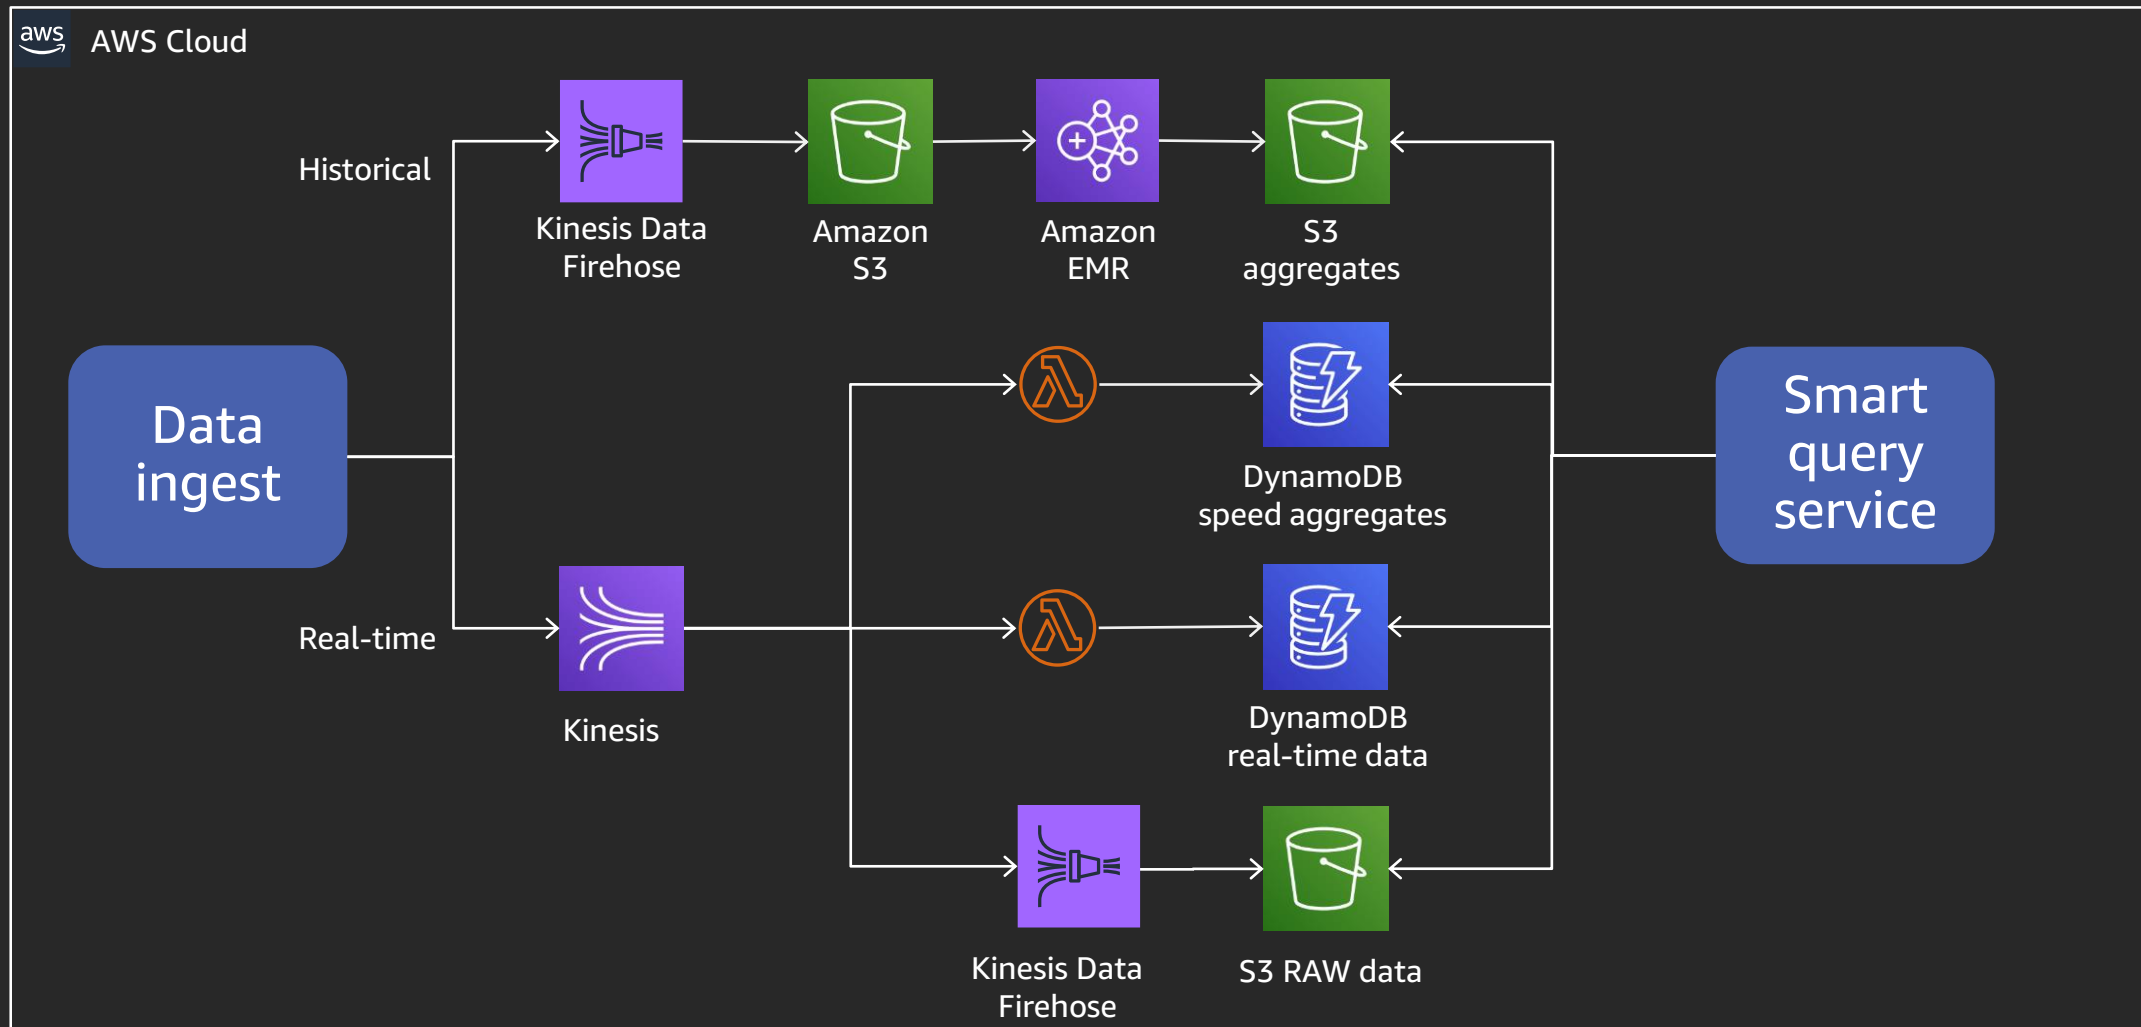
Unpredictable workloads



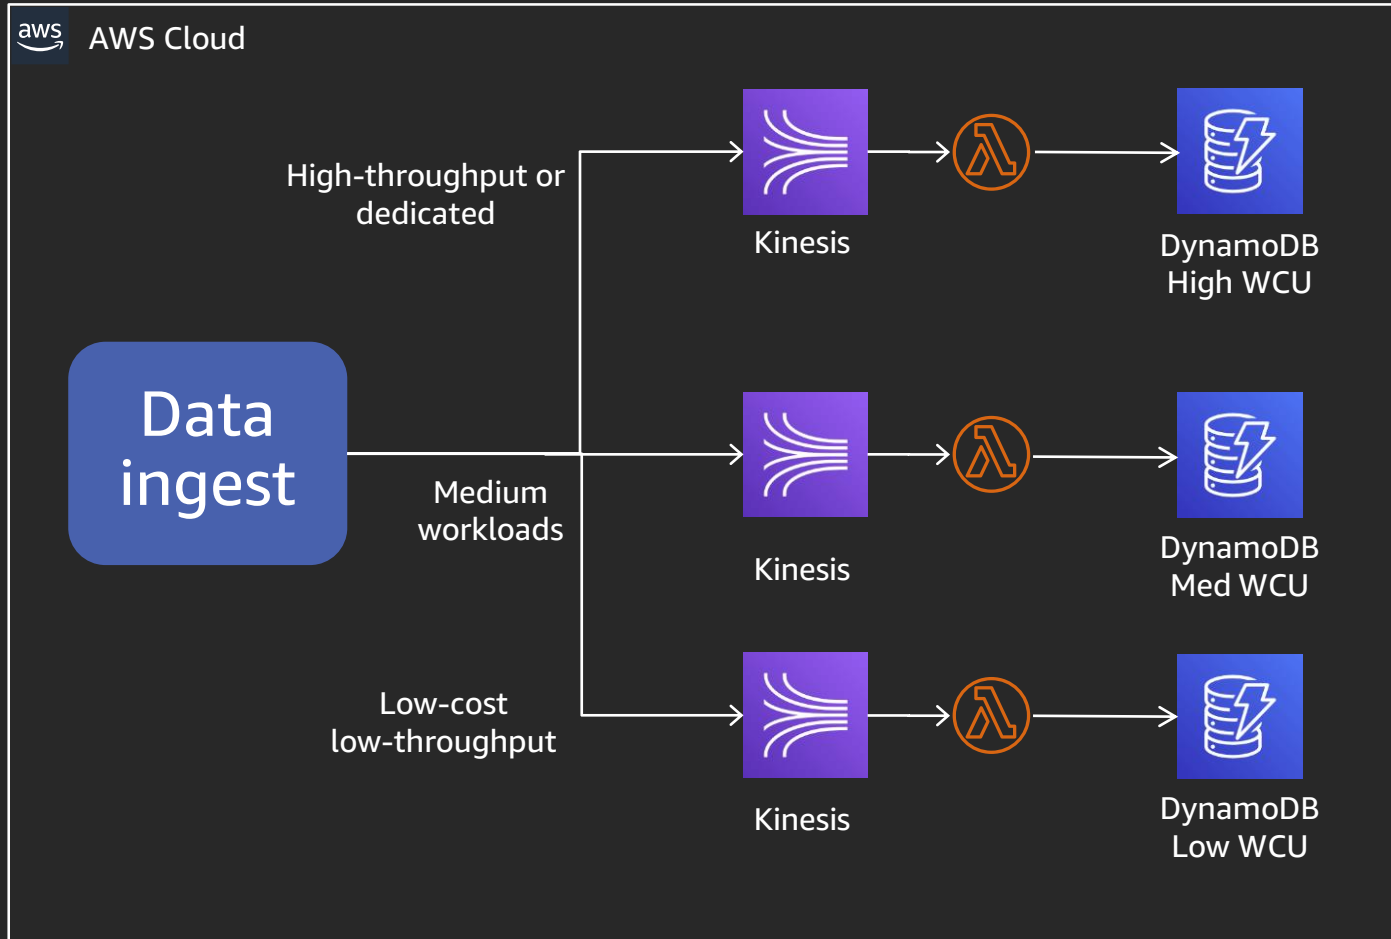
1	Over-provisioning high cost
2	Buffer for peak workloads
3	Separate real time from historical
4	Separate mixed workloads

... the real world is never what you expect

Separate real-time from non-real-time data



Split mixed workloads



Not all devices are the same

Over-provisioning for worst case is expensive

Key learnings for scalability

- Understand service limits & design with them in mind
- Scaling takes time
- The real world will not be what you expect
- Monitor access patterns—learn & adapt
- Serverless makes it all easier—but still has limits

Cost optimization

Managing costs: Understanding cost factors

- Use cost calculators
- Storage costs
- WCU & RCU
- Reading costs
- Provisioned vs on-demand

Indexed Data Storage:

Dataset Size: GB

On-demand Capacity *:

Item Size (All attributes): KB

Read Consistency: Strongly Consistent Eventually Consistent

Total number of items read per month: in millions

Number of transactional items read per month (subset of above): in millions

Number of items read with eventual consistency: in millions

Number of items read with strong consistency: in millions

Number of items written per month: in millions

Number of transactional items written per month (subset of above): in millions

Provisioned Capacity *:

Item Size (All attributes): KB

Read Consistency: Strongly Consistent Eventually Consistent

Total number of items read per second: Reads/Second

Number of transactional items read per second (subset of above): Reads/Second

Number of items read with eventual consistency: Reads/Second

Number of items read with strong consistency: Reads/Second

Number of items written per second: Writes/Second

Number of transactional items written per second (subset of above): Writes/Second

DynamoDB Streams:

Read Request Units: Units/Month

Data Transfer (optional):

Data Transfer Out of Region: GB/Month

Data Transfer In: GB/Month

General practices for cost optimization

- Automate everything
- Turn it off!
- Low-cost development environments
- Rightsize your instances
- Manage caches
- Polling on Amazon SQS vs. push on Amazon SNS
- Plan infrastructure based on usage patterns
- Use the billing console & Trusted Advisor

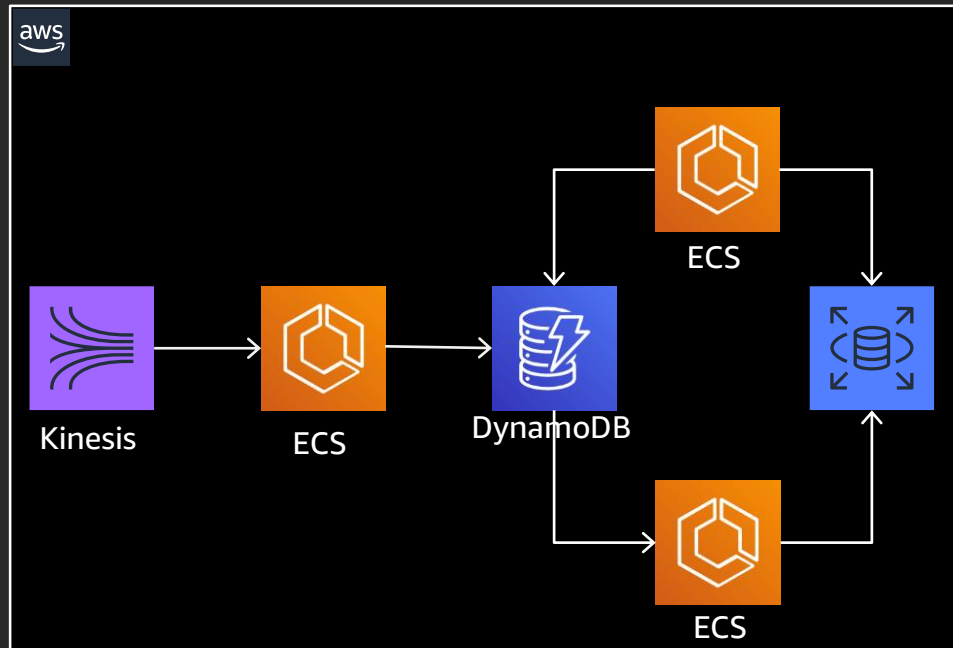
Cost optimize with serverless

Stream Kinesis

ECS reader to **DynamoDB**

Aggregator to **Amazon RDS**

Query **DynamoDB** & **Amazon RDS**



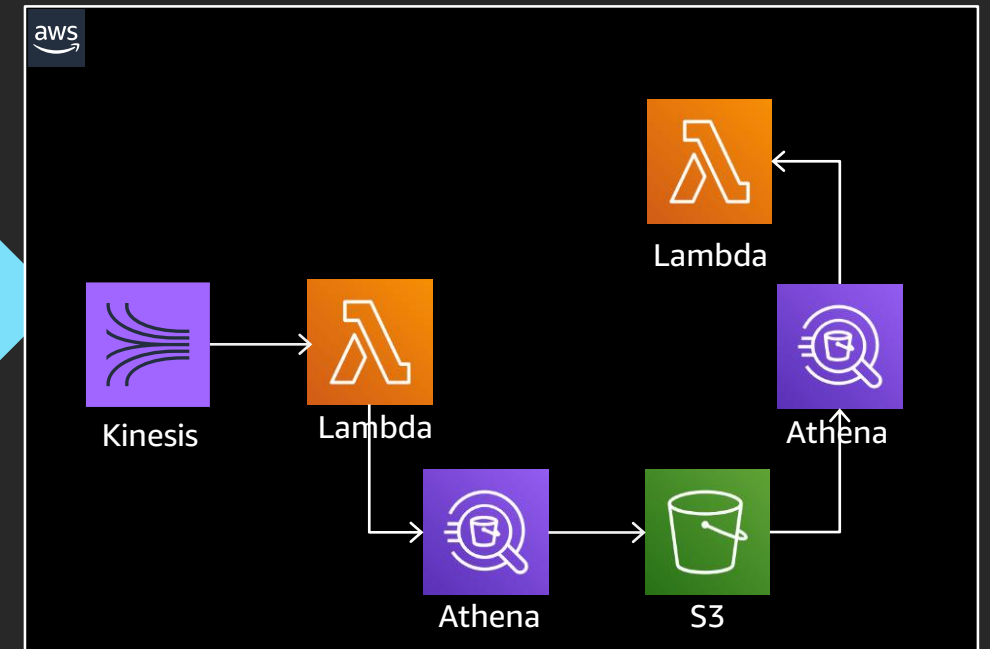
90%
cost
reduction

Stream Kinesis

Lambda reader to **Amazon S3**

Athena aggregator

Query **Athena** on **Amazon S3**



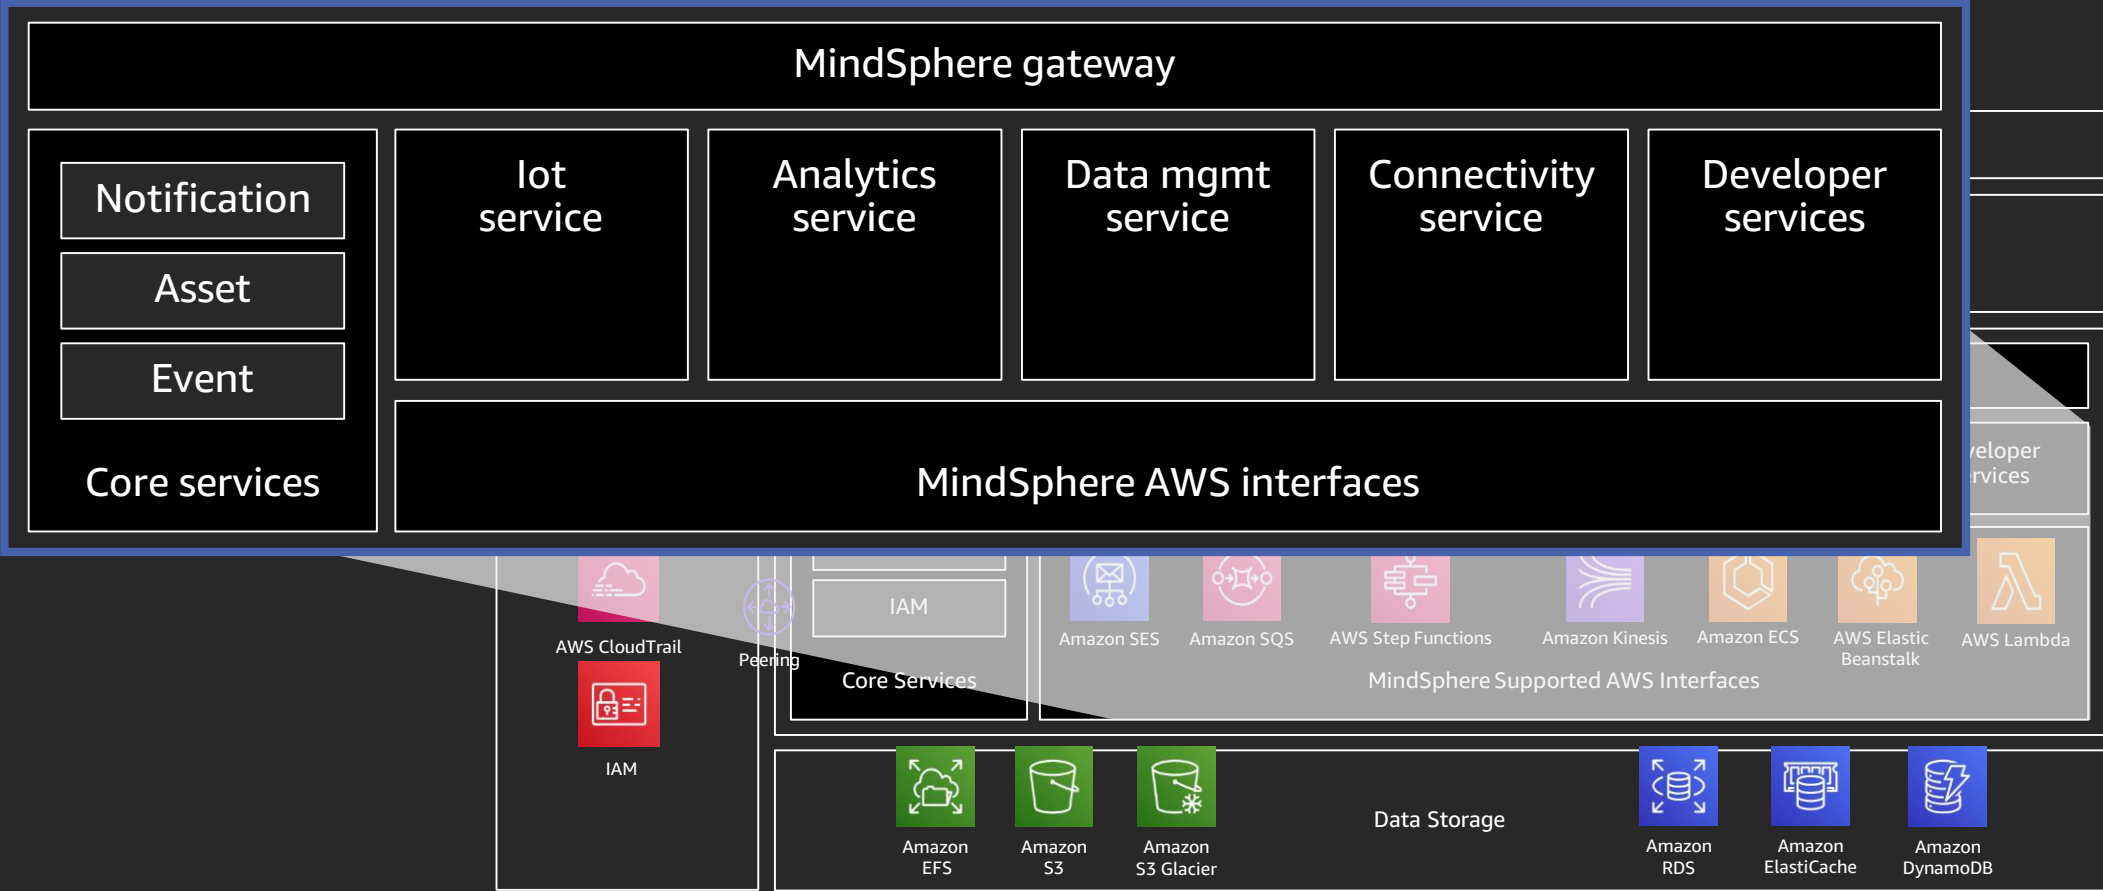
“You have to understand that the system you’re building now is not the system you will be running in 6 months or 2 years.”

Werner Vogels

CTO, Amazon.com


Delivery & operations

Microservices—independent delivery



Operational readiness

Security 

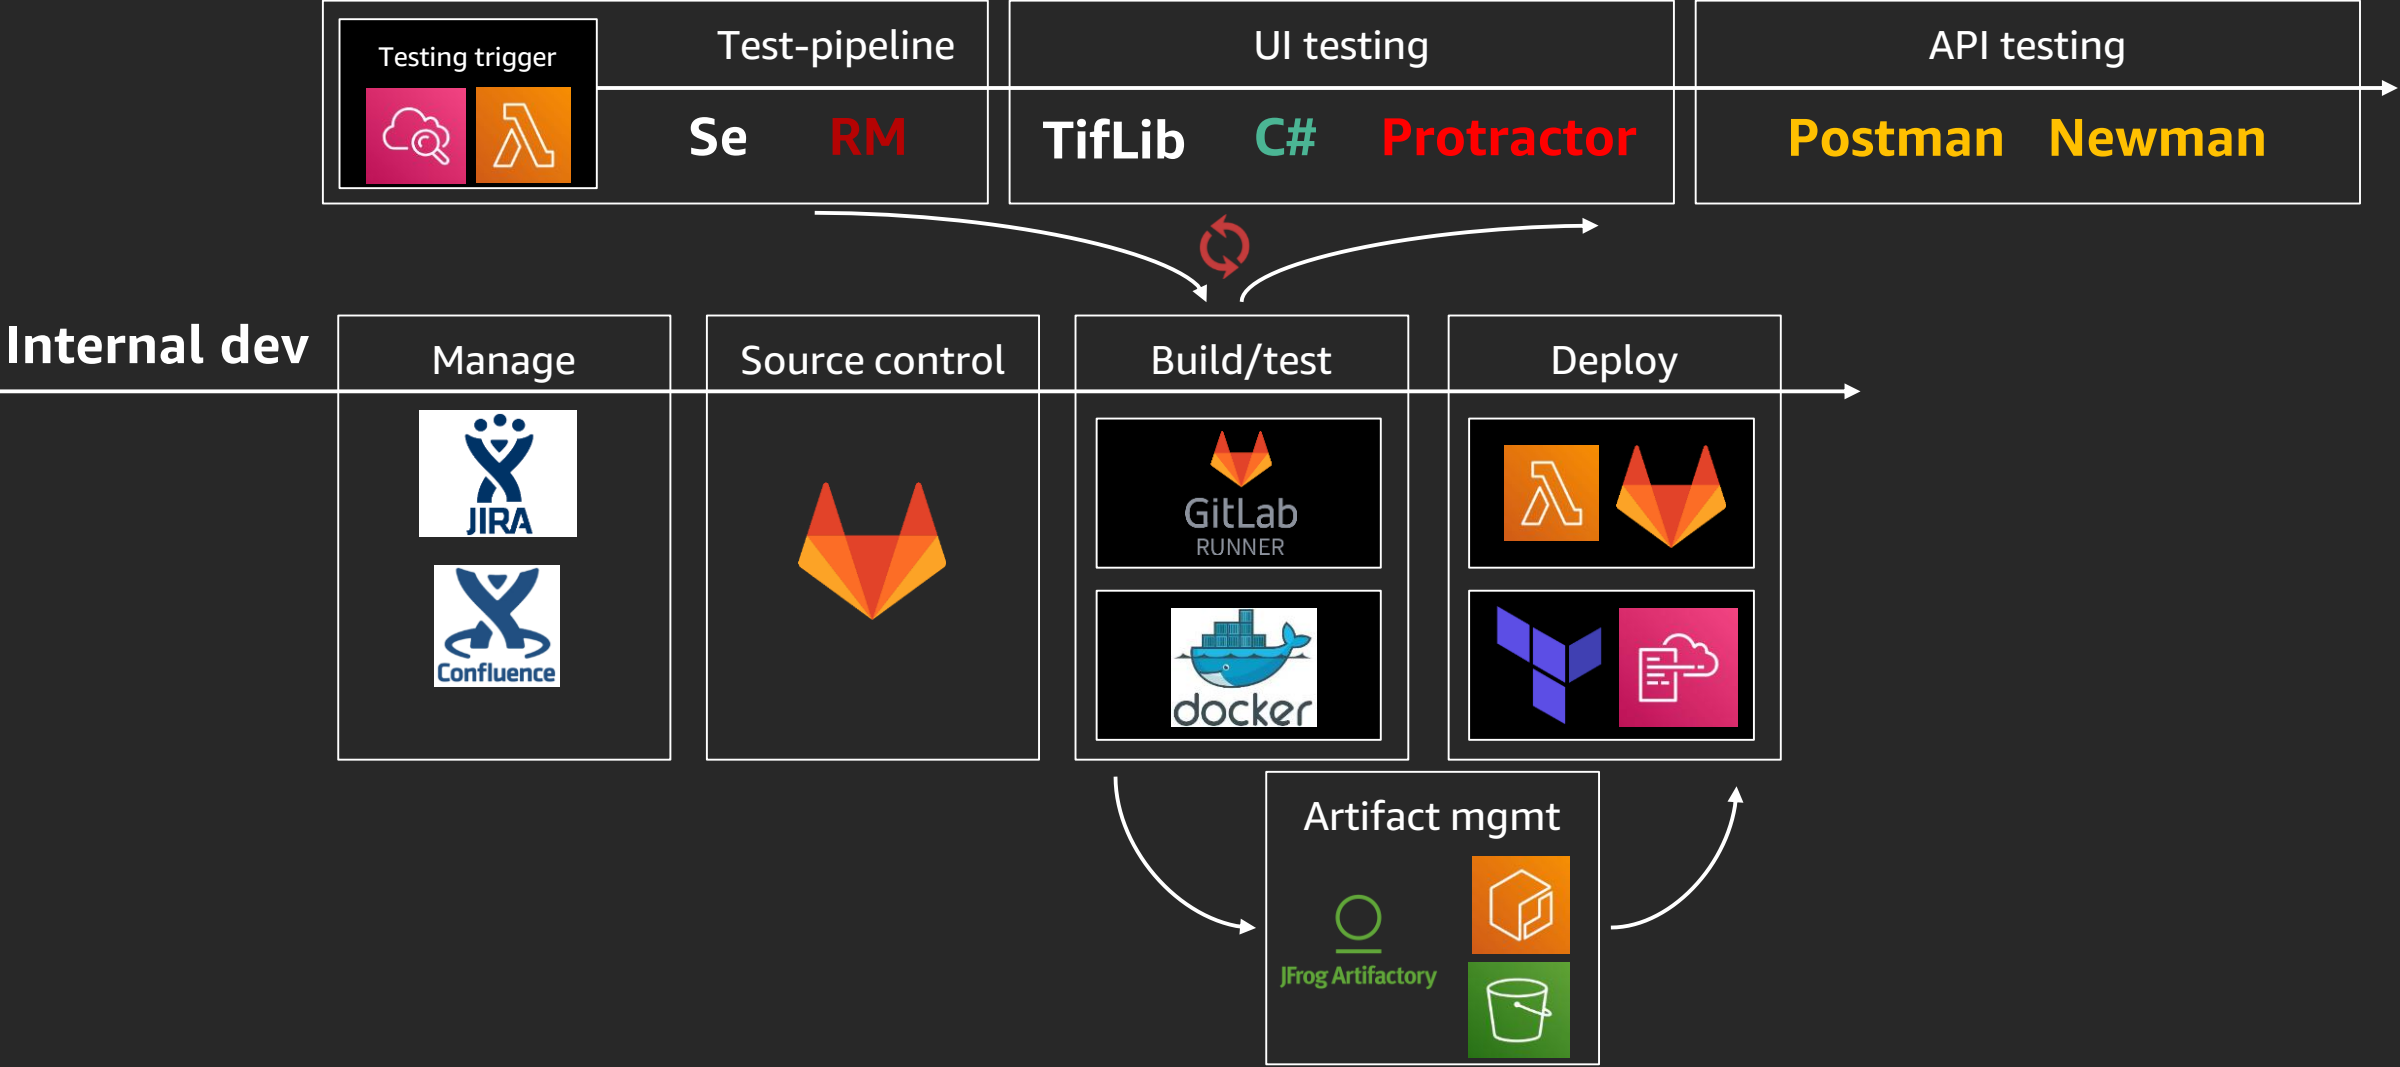
 Quality of service:
availability & operations

Quality concerns:
DoD, E2E & defect culture 

 Deployment


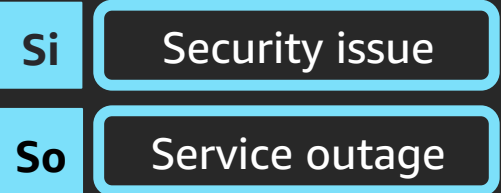
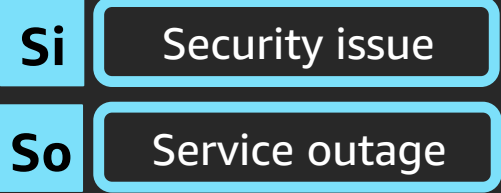
Speed & reproducibility 

CI/CD automate for speed of delivery




Continuous monitoring

Customer support team



System generated



Key takeaways



Align your architecture & strategy



Security, reliability & scalability



Plan for change



Monitor, analyze & improve



Practice & experiment

Thank you!

Kathleen DeValk

Kathleen.devalk@siemens.com

Frank Kovacs

kovfrank@amazon.com



Please complete the session survey in the mobile app.