AWS re:Inforce

JUNE 10 - 12, 2024 | PHILADELPHIA, PA

APS224

Deep dive into Amazon Bedrock security architecture

Raj Pathak
(he/him)
Principal Solutions Architect
Amazon Web Services

Agenda

Introduction to Amazon Bedrock

How Amazon Bedrock manages security

Amazon Bedrock model environments

What happens during a request to Amazon Bedrock

Securing prompts and responses with Guardrails for Amazon Bedrock

Getting started

Introduction to Amazon Bedrock

Amazon **Bedrock**

The easiest way to build and scale generative AI applications with foundation models (FMs)

Accelerate development of generative AI applications, agents, and RAG using FMs through a single API, without managing infrastructure

Choose FMs from Amazon, AI21 Labs, Anthropic, Cohere, Meta, Mistral AI, and Stability AI to find the right FM for your use case

Privately customize FMs using your organization's data

Amazon Bedrock

Broad choice of models

AI21labs	amazon	ANTHROPIC	cohere	Meta	MISTRAL AI_	stability.ai
Contextual answers, summarization, paraphrasing	Text summarization, generation, Q&A, search, image generation	Summarization, complex reasoning, writing, coding	Text generation, search, classification	Q&A and reading comprehension	Text summarization, Q&A, text classification, text completion, code generation	High-quality images and art
Jurassic-2 Ultra	Amazon Titan Text Premier	Claude 3 Opus	Command	Llama 3 8B	Mistral Large	Stable Diffusion XL 1.0
Jurassic-2 Mid	Amazon Titan Text Lite	Claude 3 Sonnet	Command Light	Llama 3 70B	Mistral 7B	Stable Diffusion XL 0.8
	Amazon Titan Text Express	Claude 3 Haiku	Embed English	Llama 2 13B	Mixtral 8x7B	
	Amazon Titan Text Embeddings	Claude 2.1	Embed Multilingual	Llama 2 70B	Mistral Small	
	Amazon Titan Text Embeddings V2	Claude 2	Command R+			
	Amazon Titan Multimodal Embeddings	Claude Instant	Command R			
	Amazon Titan Image Generator					

aws

How Amazon Bedrock manages security

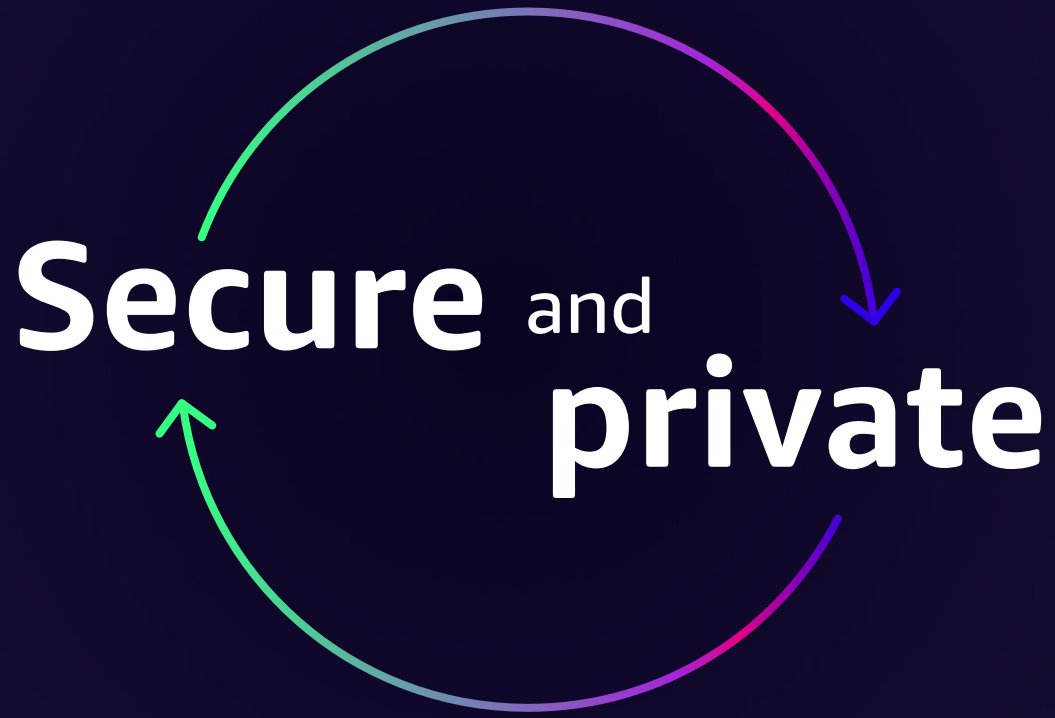

Everything you expect
from an AWS service

IAM | AWS monitoring and logging | AWS PrivateLink

AWS service compliance | AWS data privacy and encryption

Amazon Bedrock keeps data secure and private

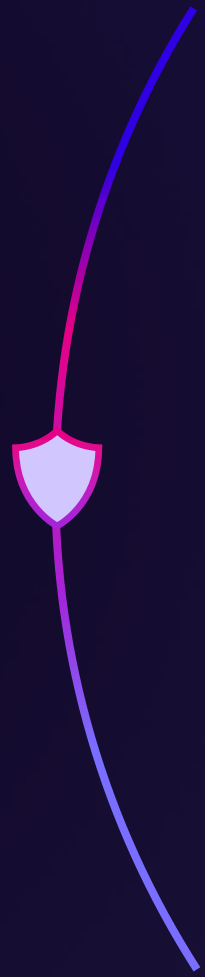

None of the customer's data is used to train the underlying foundation models

All data is encrypted at rest using AWS KMS and encrypted in transit with TLS 1.2 (minimum); also supports Amazon managed keys or customer managed keys (CMKs)

Fine-tuned models are encrypted and stored using customer AWS KMS key

Only you have access to your customized models

Integration with AWS Identity and Access Management (IAM) for fine-grained access controls

Support for data privacy standards, including GDPR, HIPAA, and PCI

Amazon Bedrock

Monitoring and compliance

Amazon CloudWatch integration

Track usage metrics and build customized dashboards

AWS CloudTrail integration

Monitor API activity and troubleshoot issues

Compliance frameworks

Support for FedRAMP Moderate, GDPR, SOC, ISO, CSA compliance, and HIPAA eligibility

IAM permissions with Amazon Bedrock

IAM

Example of an allow policy

```
{
 "Version": "2012-10-17",
 "Statement": {
  "Sid": "AllowInference",
  "Effect": "Allow",
  "Action": [
   "bedrock:InvokeModel",
   "bedrock:InvokeModelWithResponseStream"
  ],
  "Resource": "arn:aws:bedrock:*::foundation-
model/model-id"
 }
}
```

Example of a deny policy

```
{
   "Version": "2012-10-17",
   "Statement": {
      "Sid": "DenyInference",
      "Effect": "Deny",
      "Action": [
         "bedrock:InvokeModel",
         "bedrock:InvokeModelWithResponseStream"
       ],
      "Resource": "arn:aws:bedrock:*::foundation-
model/model-id"
   }
}
```

Model access within Amazon Bedrock

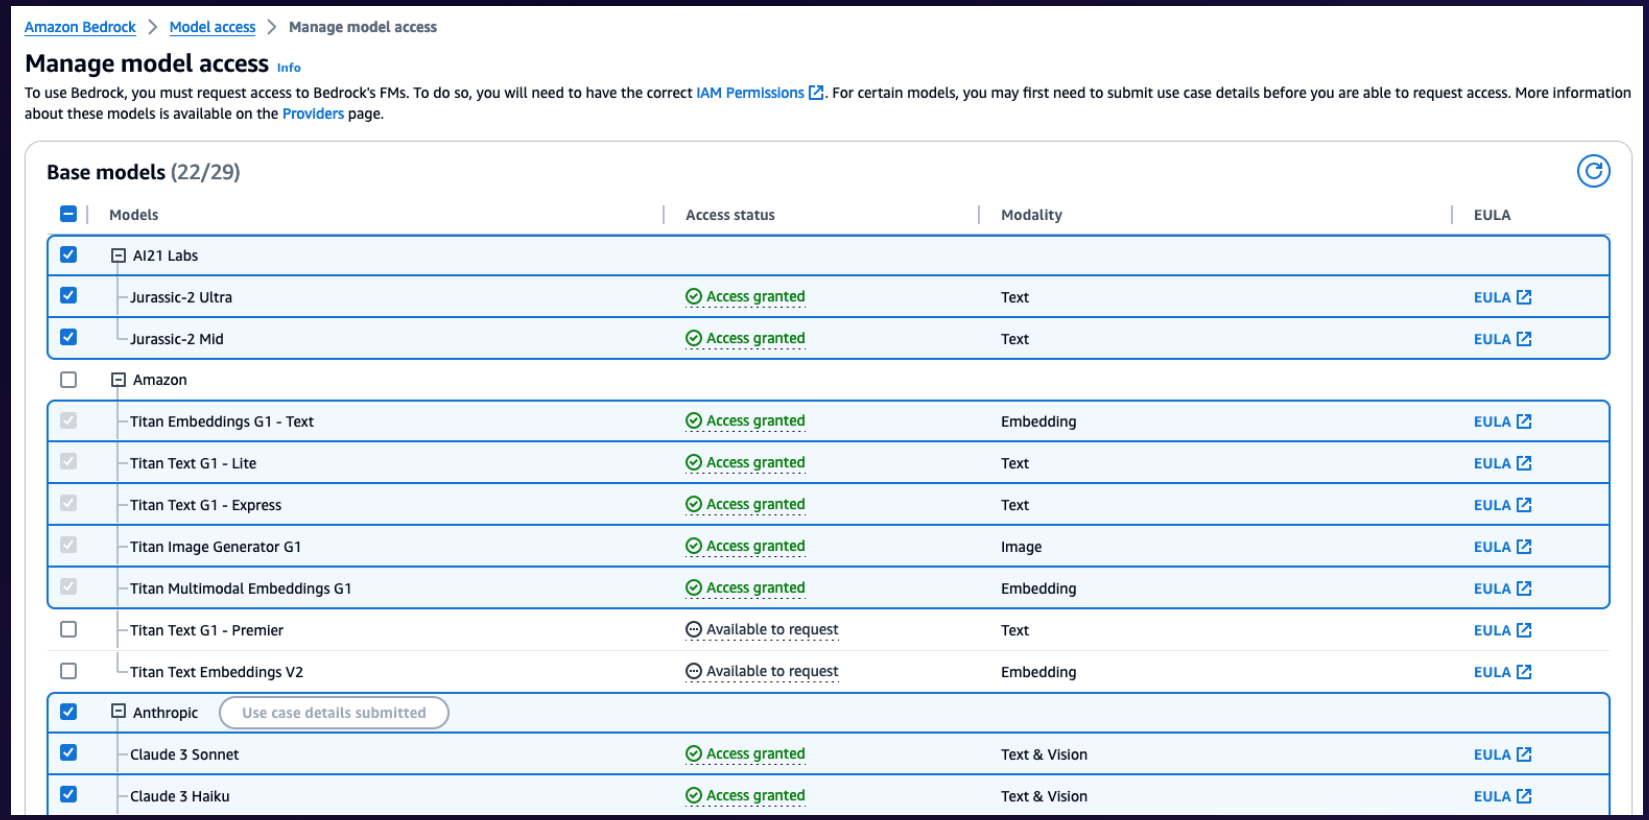

IAM alone does not provide access to models in Amazon Bedrock

- Choose which models to enable in each AWS Region
- Each model comes with a corresponding end-user license agreement (EULA)

Amazon Bedrock model environments

On-demand vs. provisioned compute capacity

On-demand compute

Common compute features

Provisioned capacity compute

Deployment available to all customers

No inference request's input or output text is used to train any model

Deployments are inside an AWS account owned and operated by the Amazon Bedrock service team

Model vendors have no access to any customer data

Deployment available to a single customer

On-demand vs. provisioned compute capacity

On-demand compute

Common compute features

Provisioned capacity compute

Deployment available to all customers

Holds a baseline version of a supported model

No inference request's input or output text is used to train any model

Deployments are inside an AWS account owned and operated by the Amazon Bedrock service team

Model vendors have no access to any customer data

Deployment available to a single customer

Holds a private copy of a baseline model *(potentially fine-tuned by a customer)*

What happens during a request to Amazon Bedrock

Client connectivity

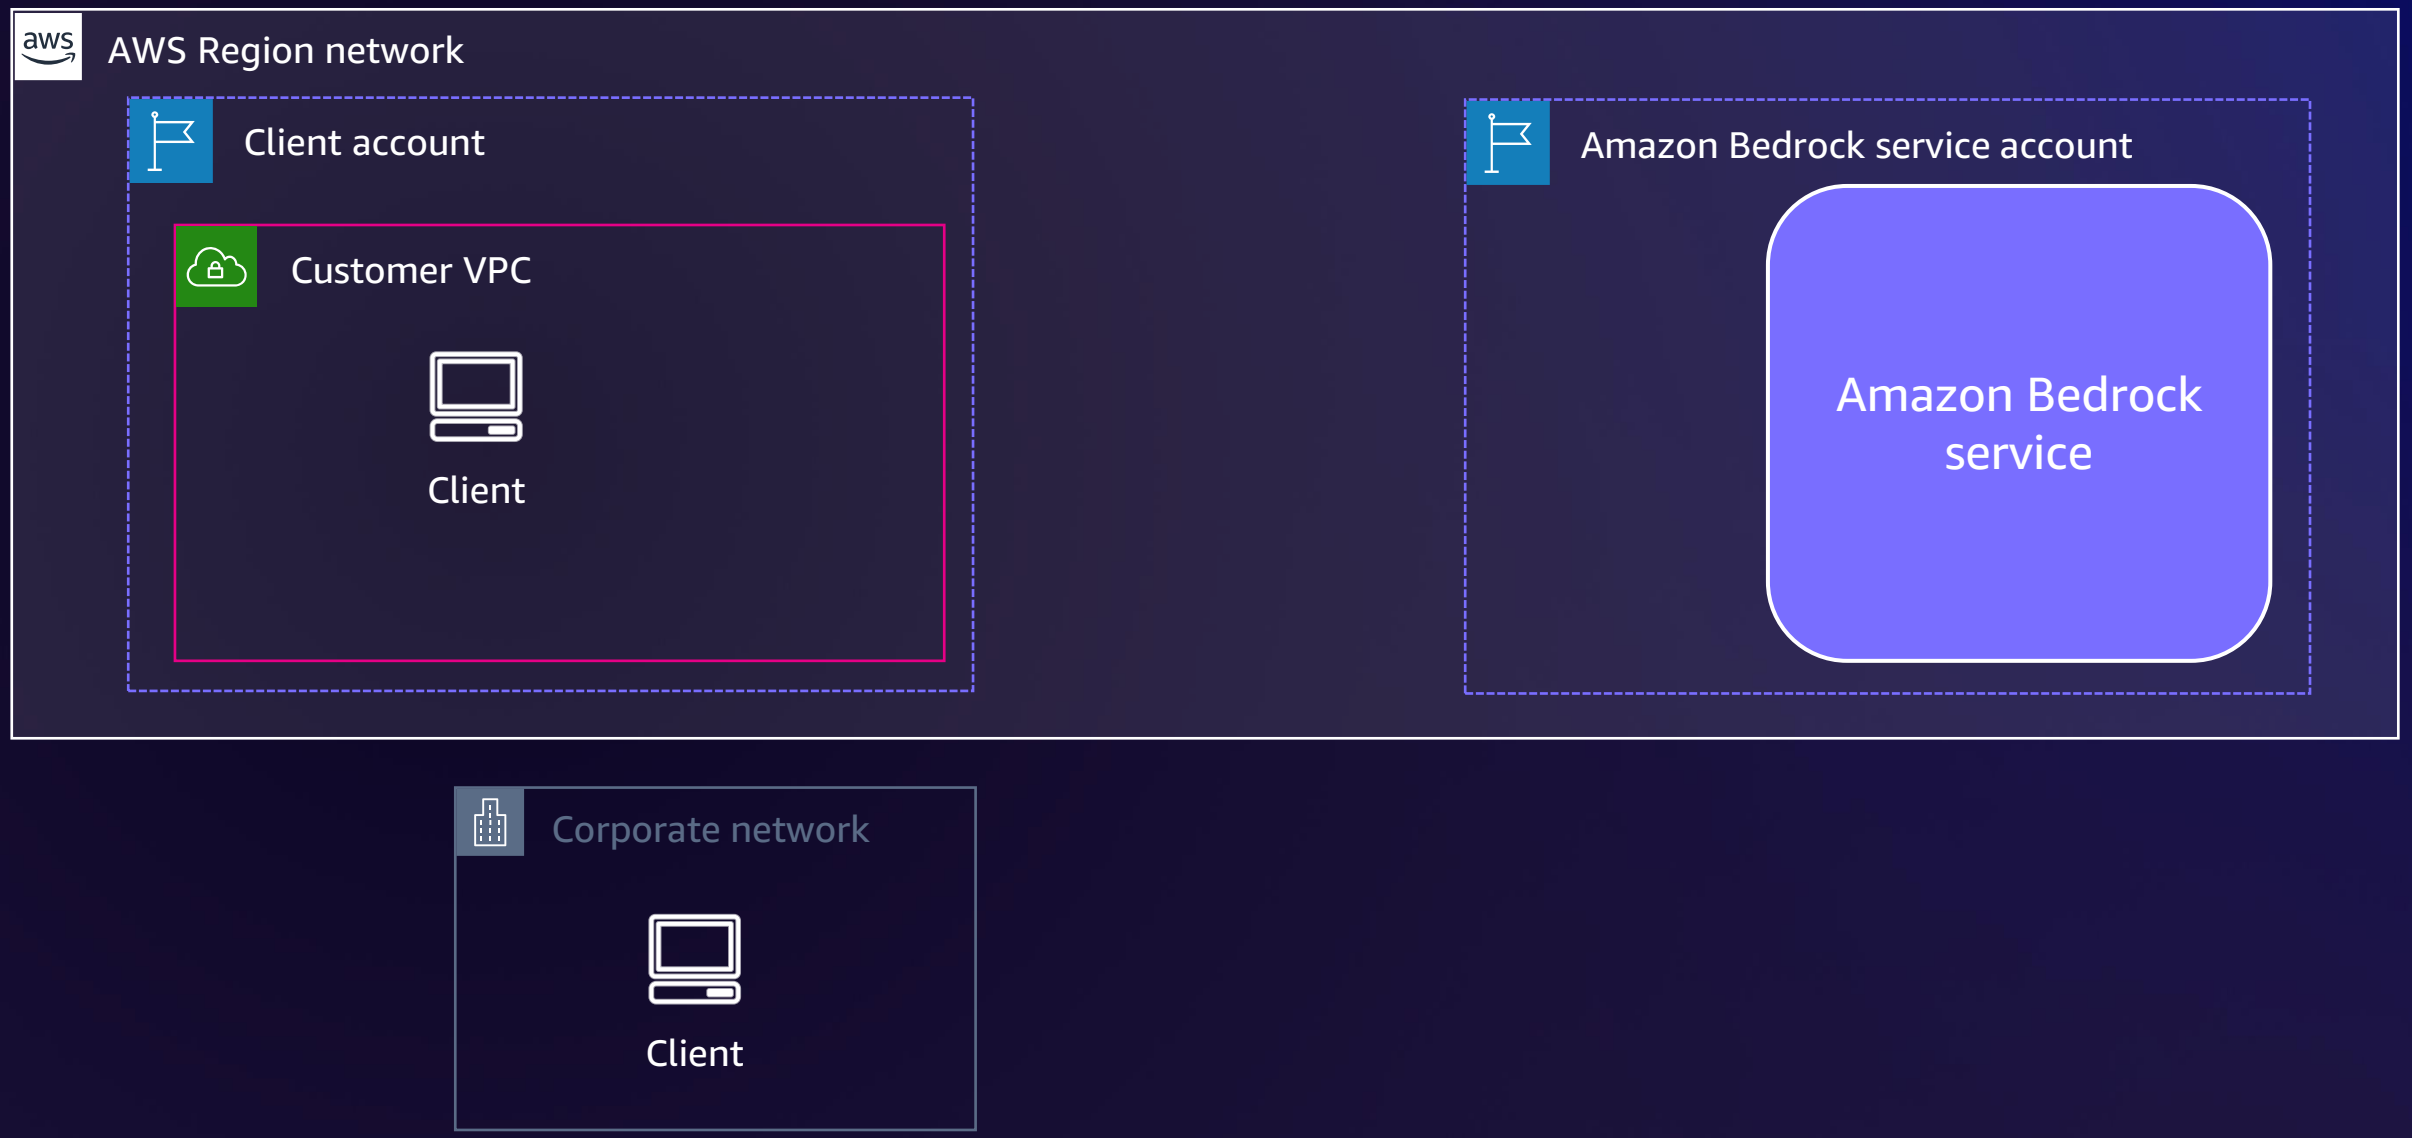

Client connectivity

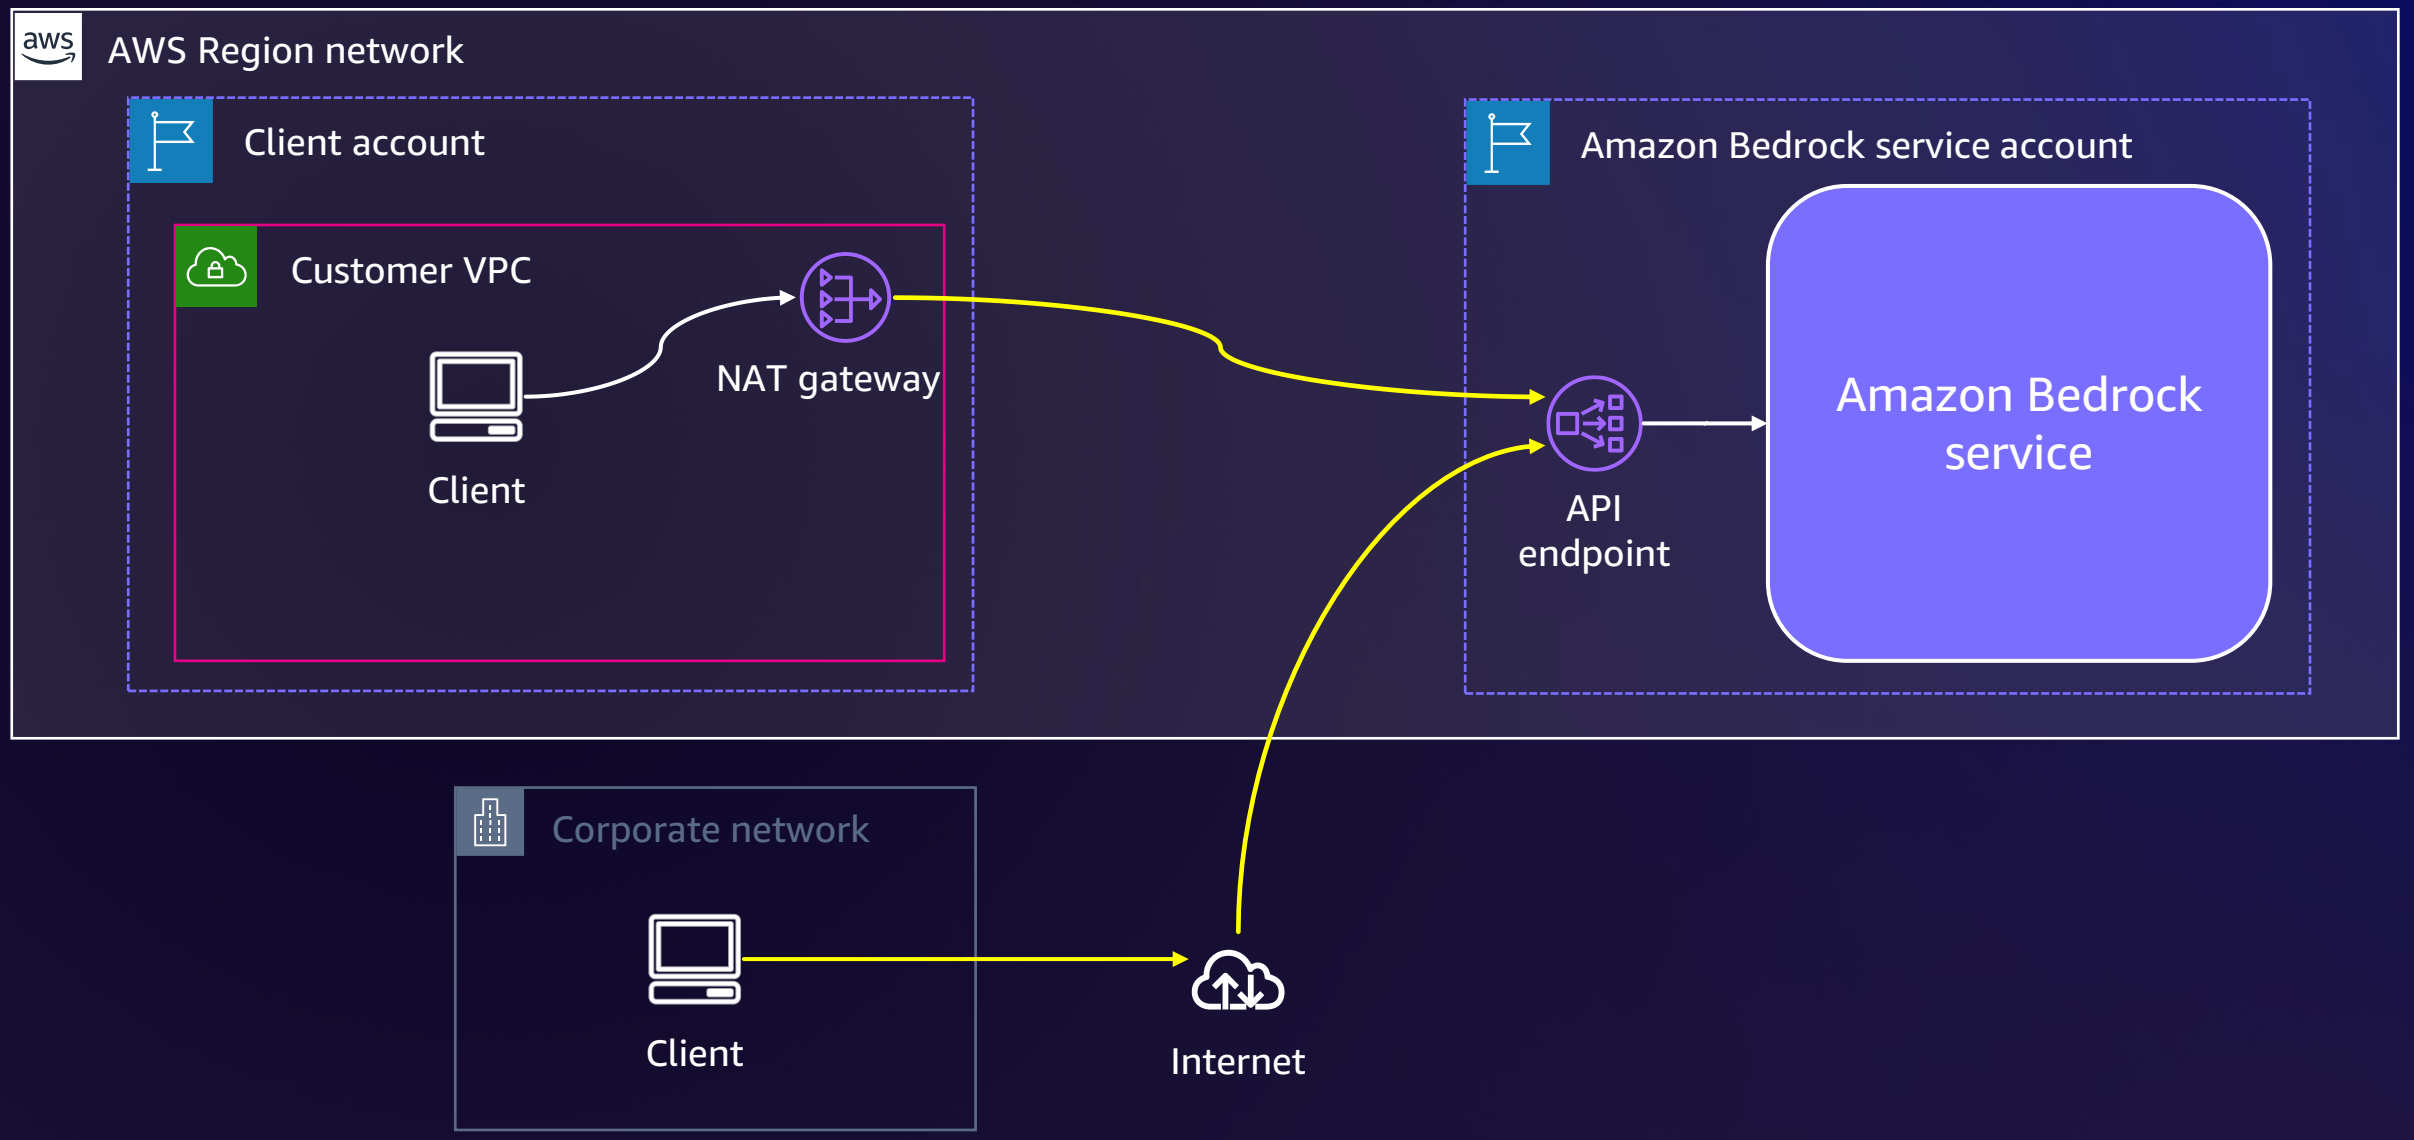

Client connectivity

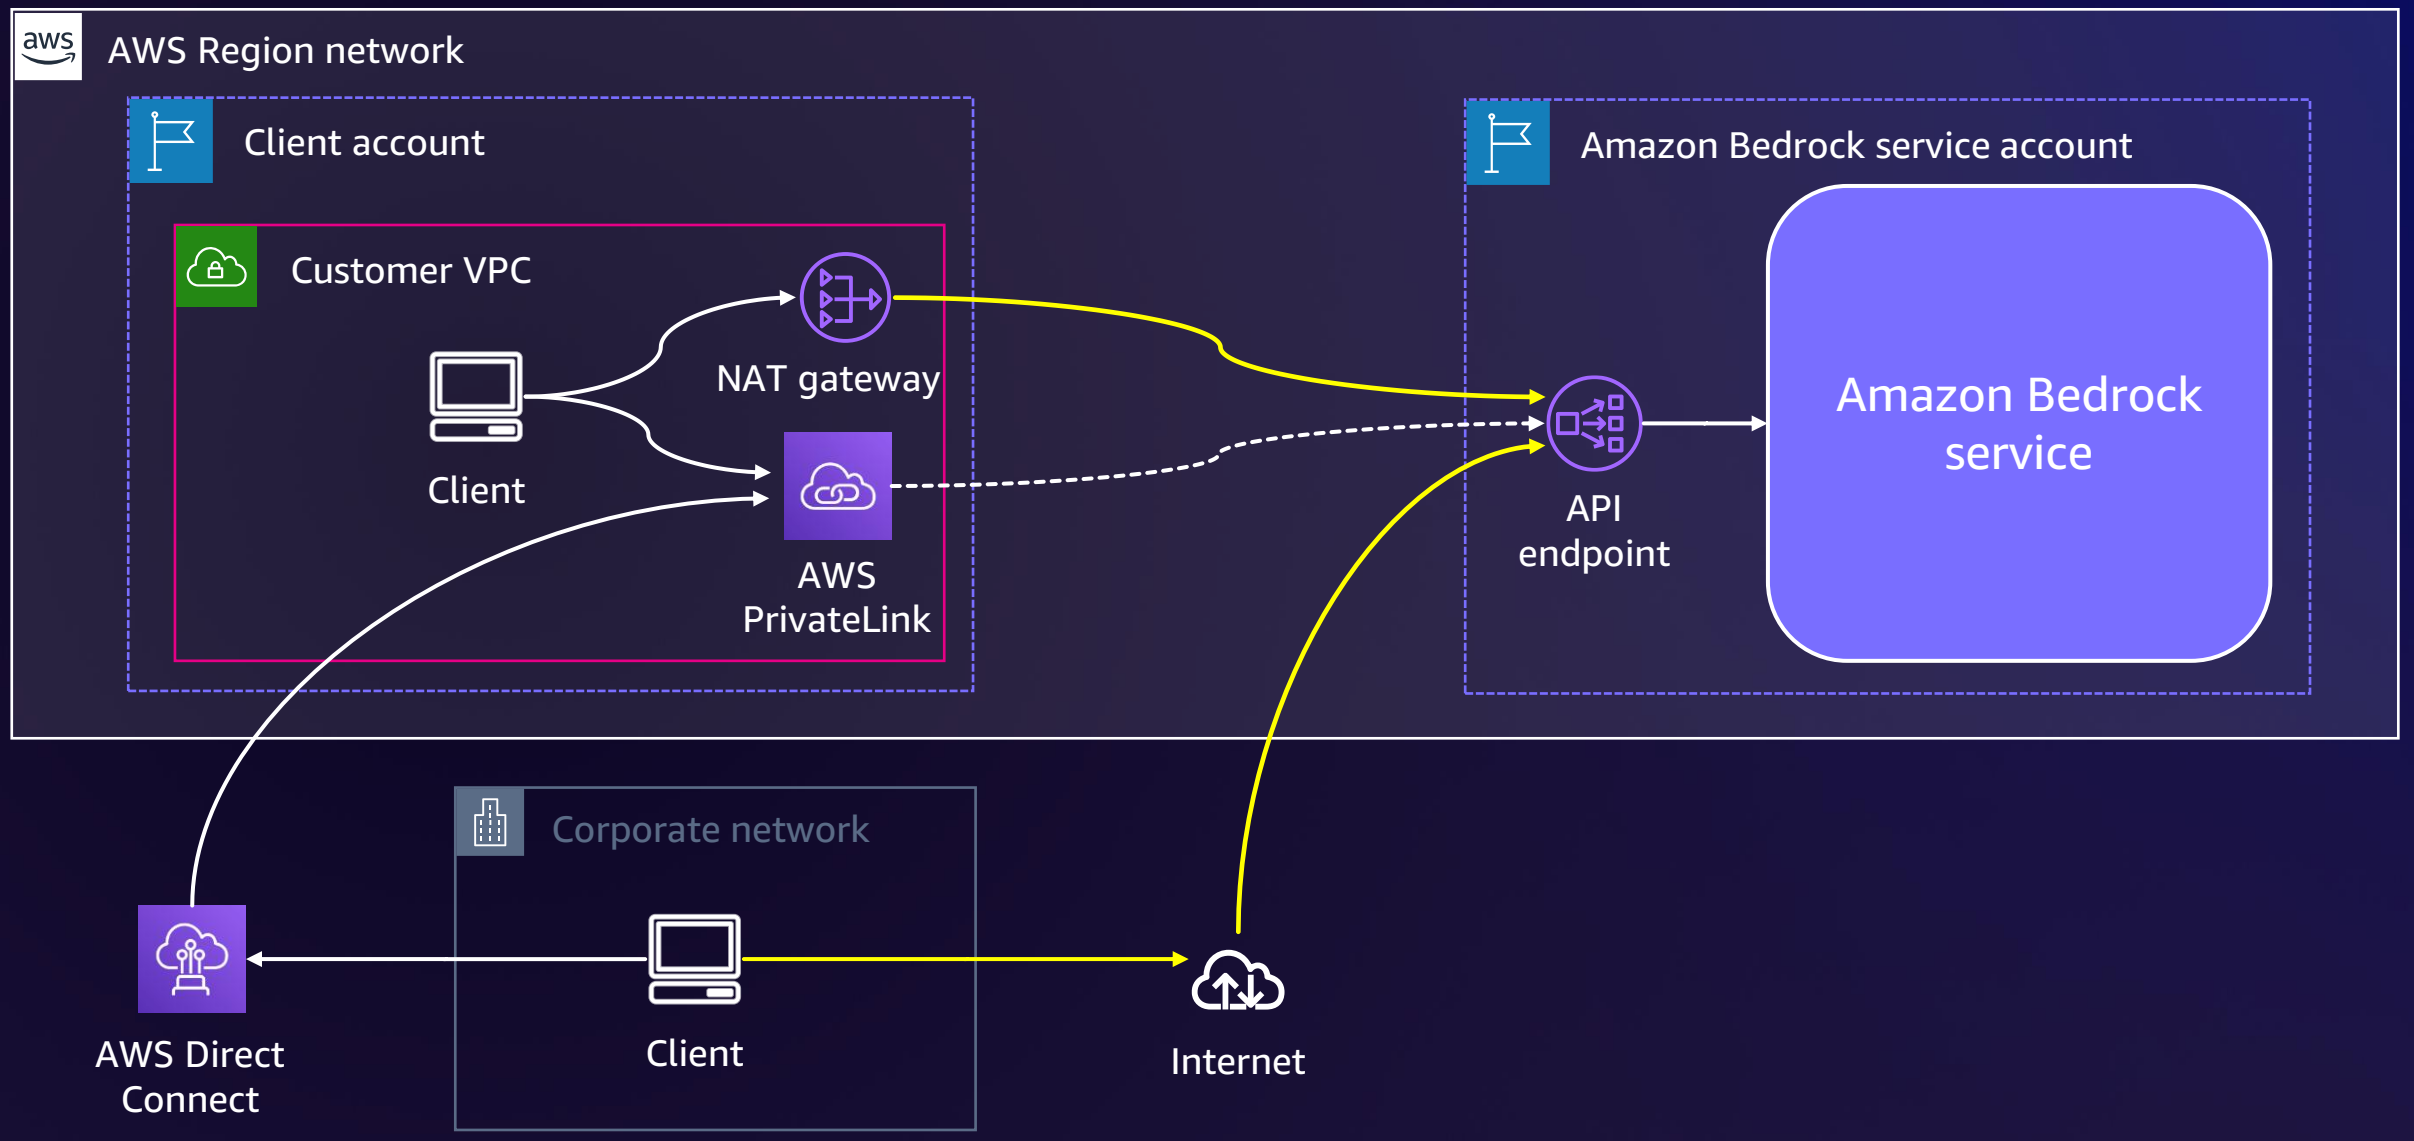

On-demand compute architecture overview

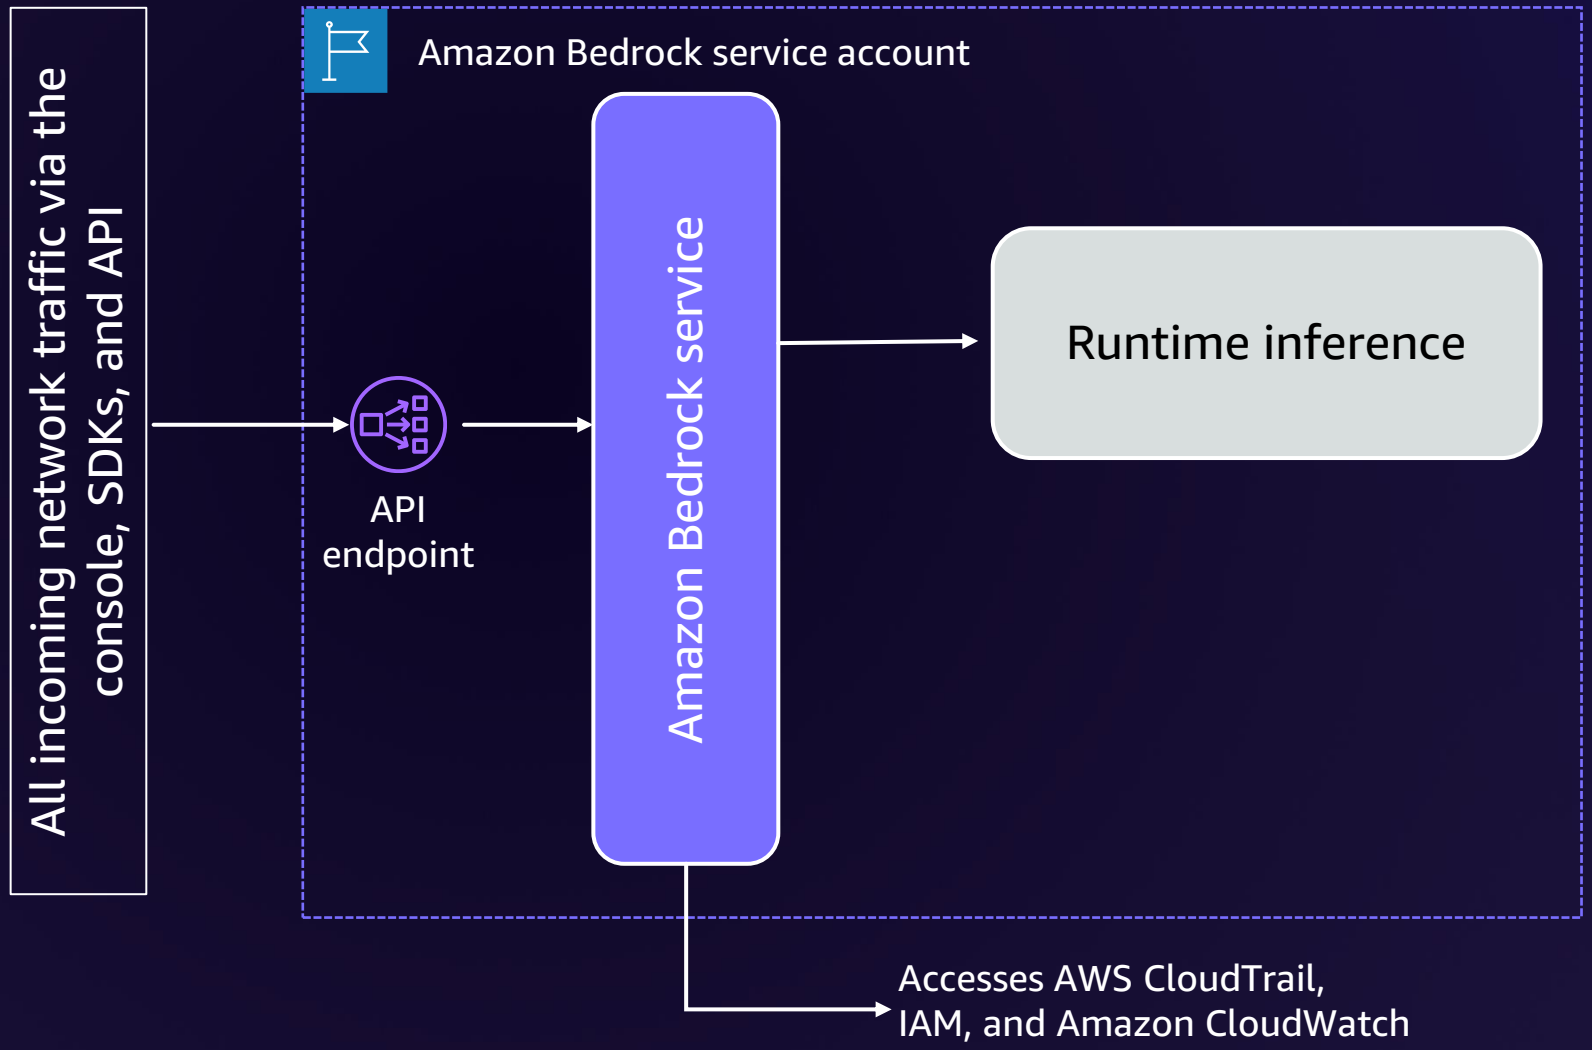

On-demand compute architecture overview

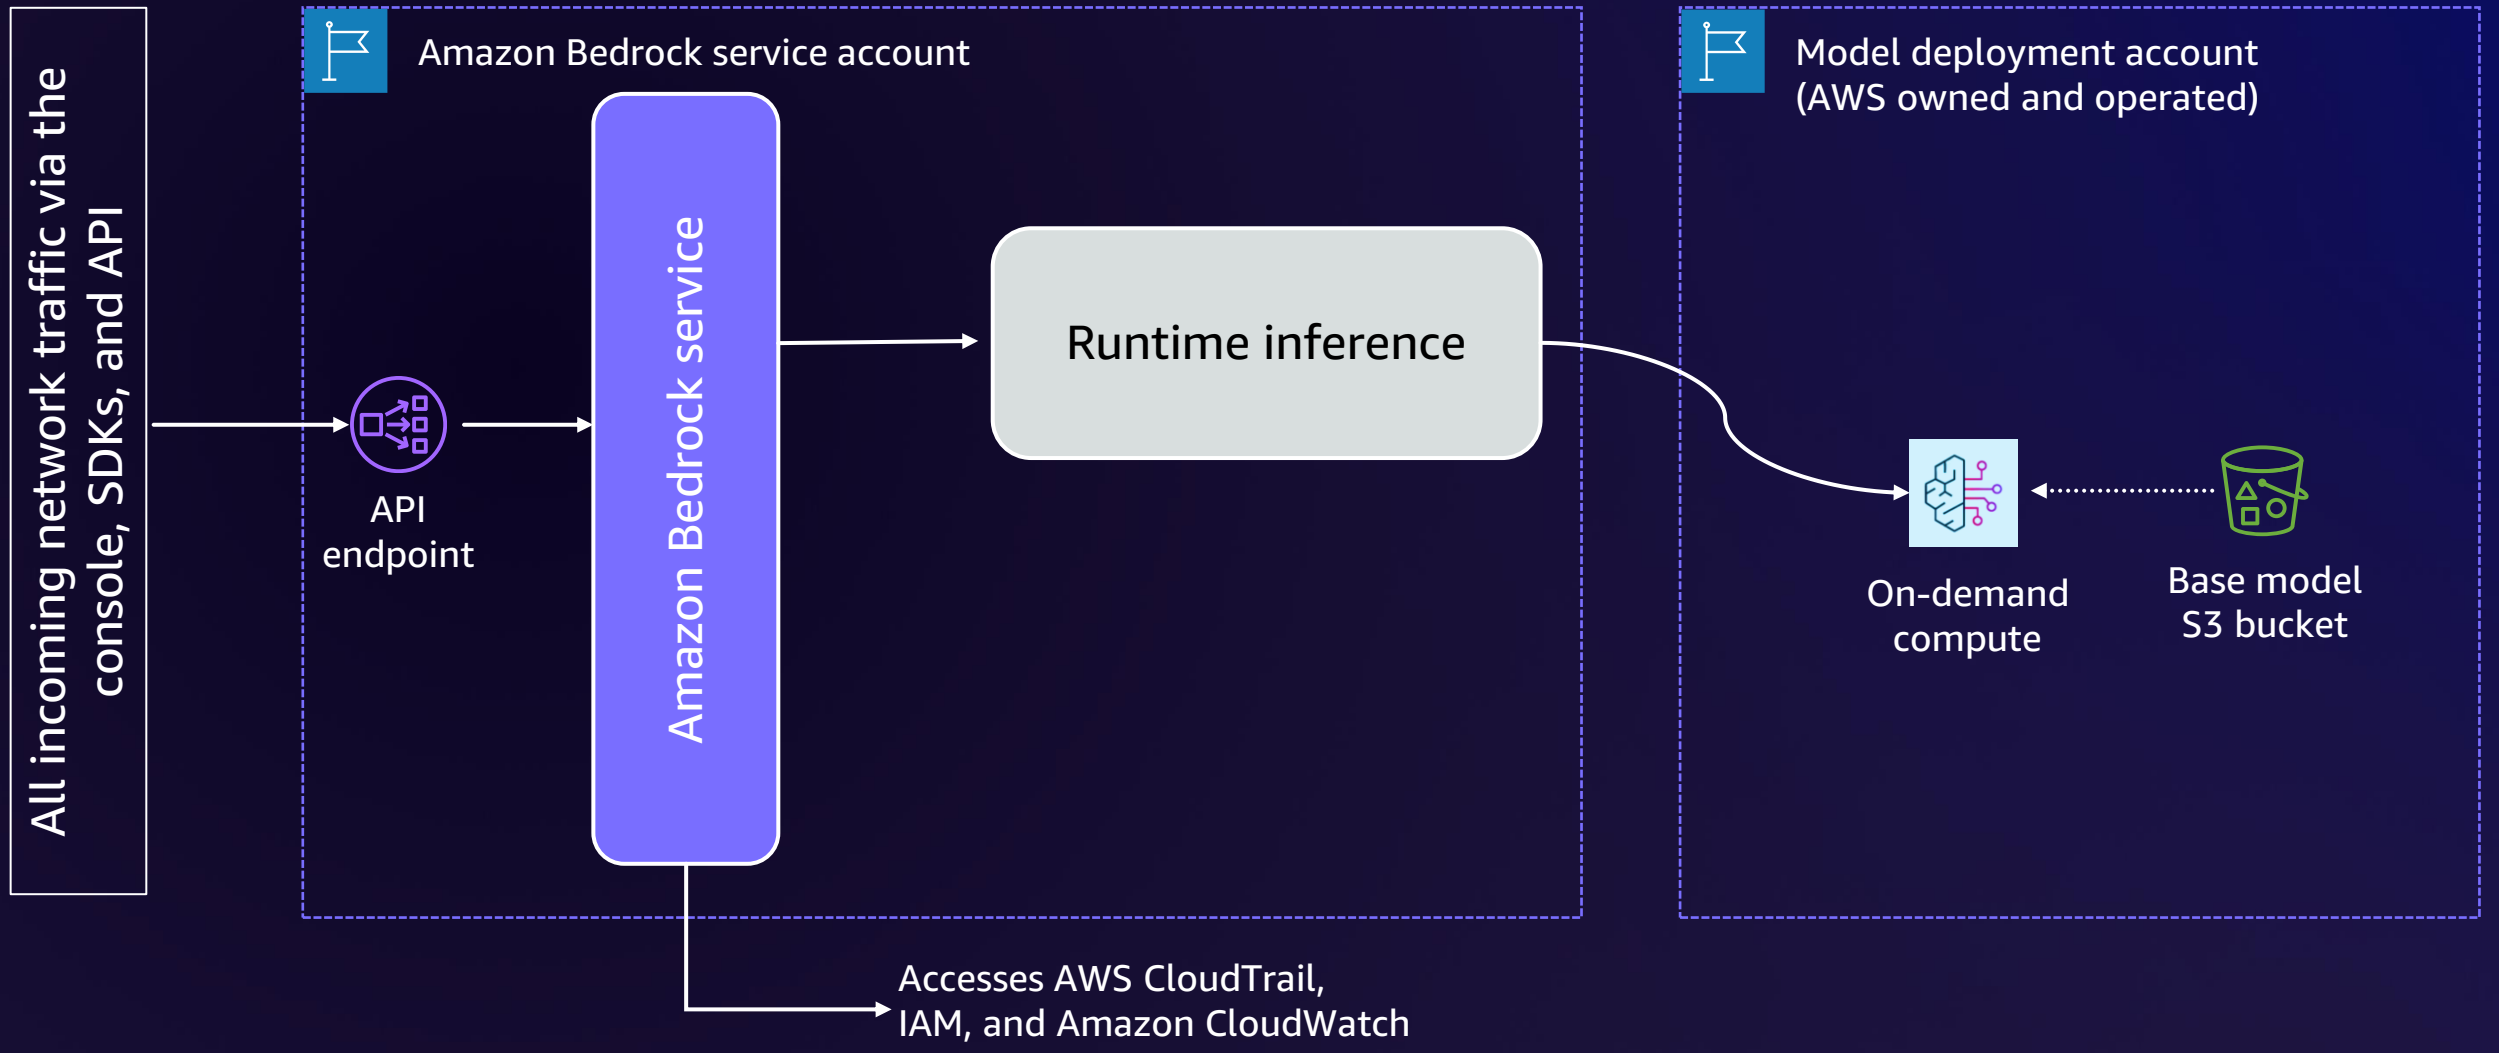

Provisioned capacity architecture overview

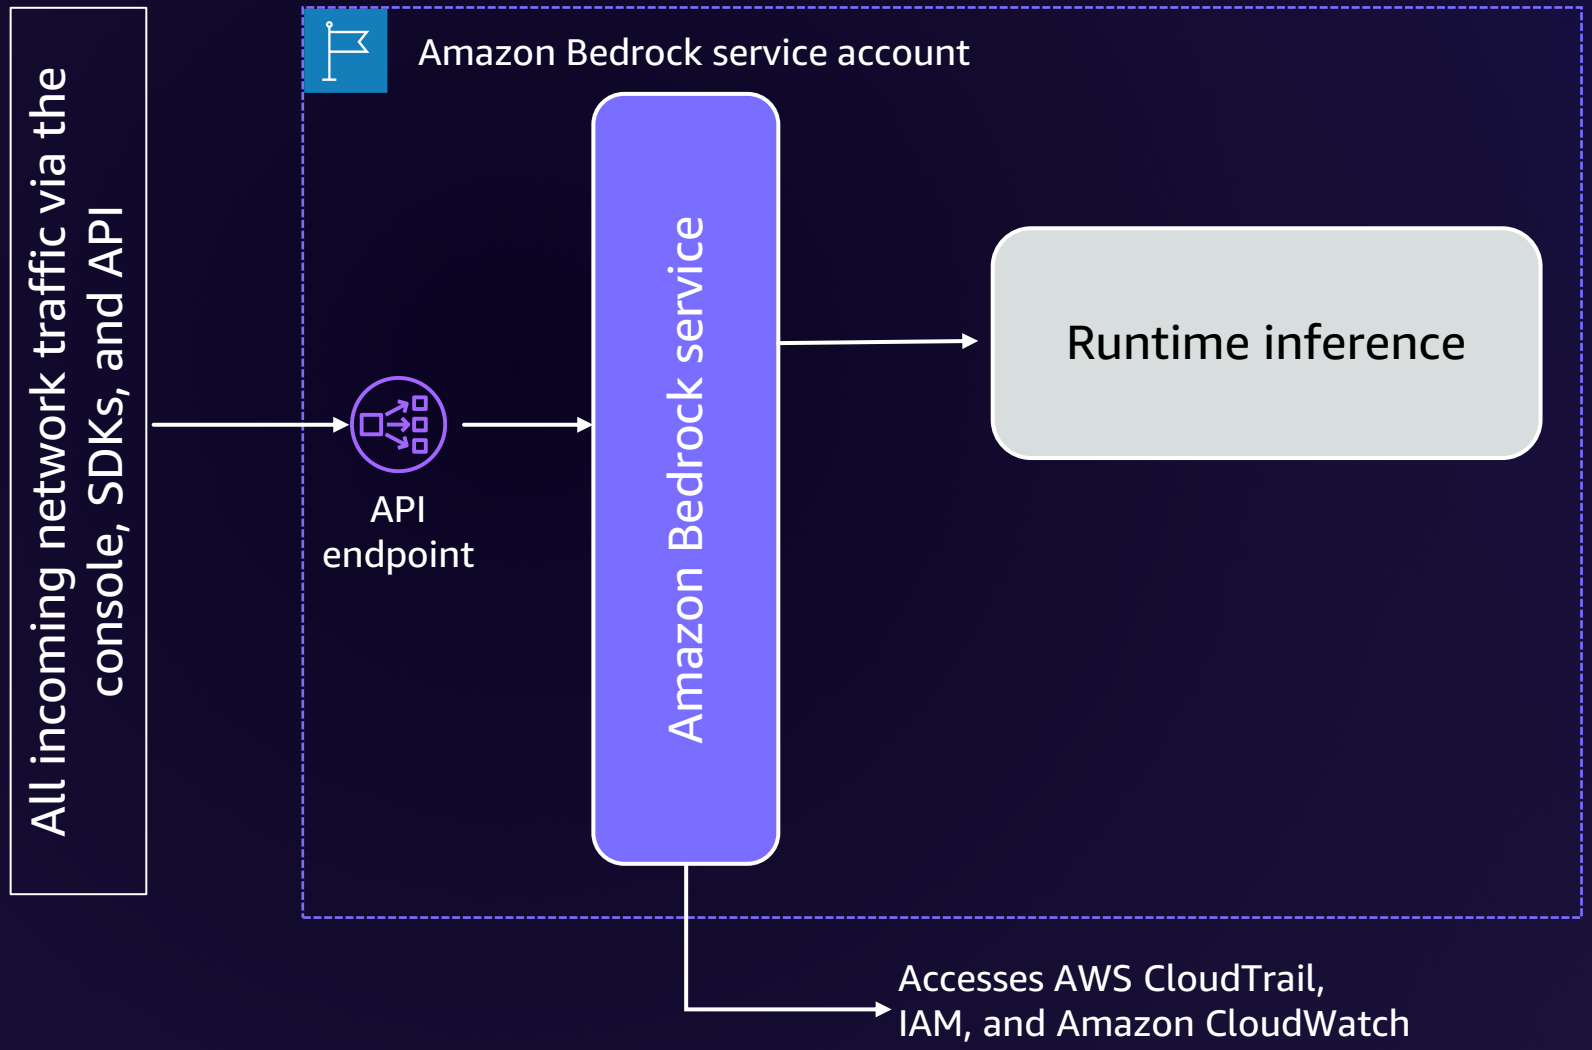

Provisioned capacity architecture overview

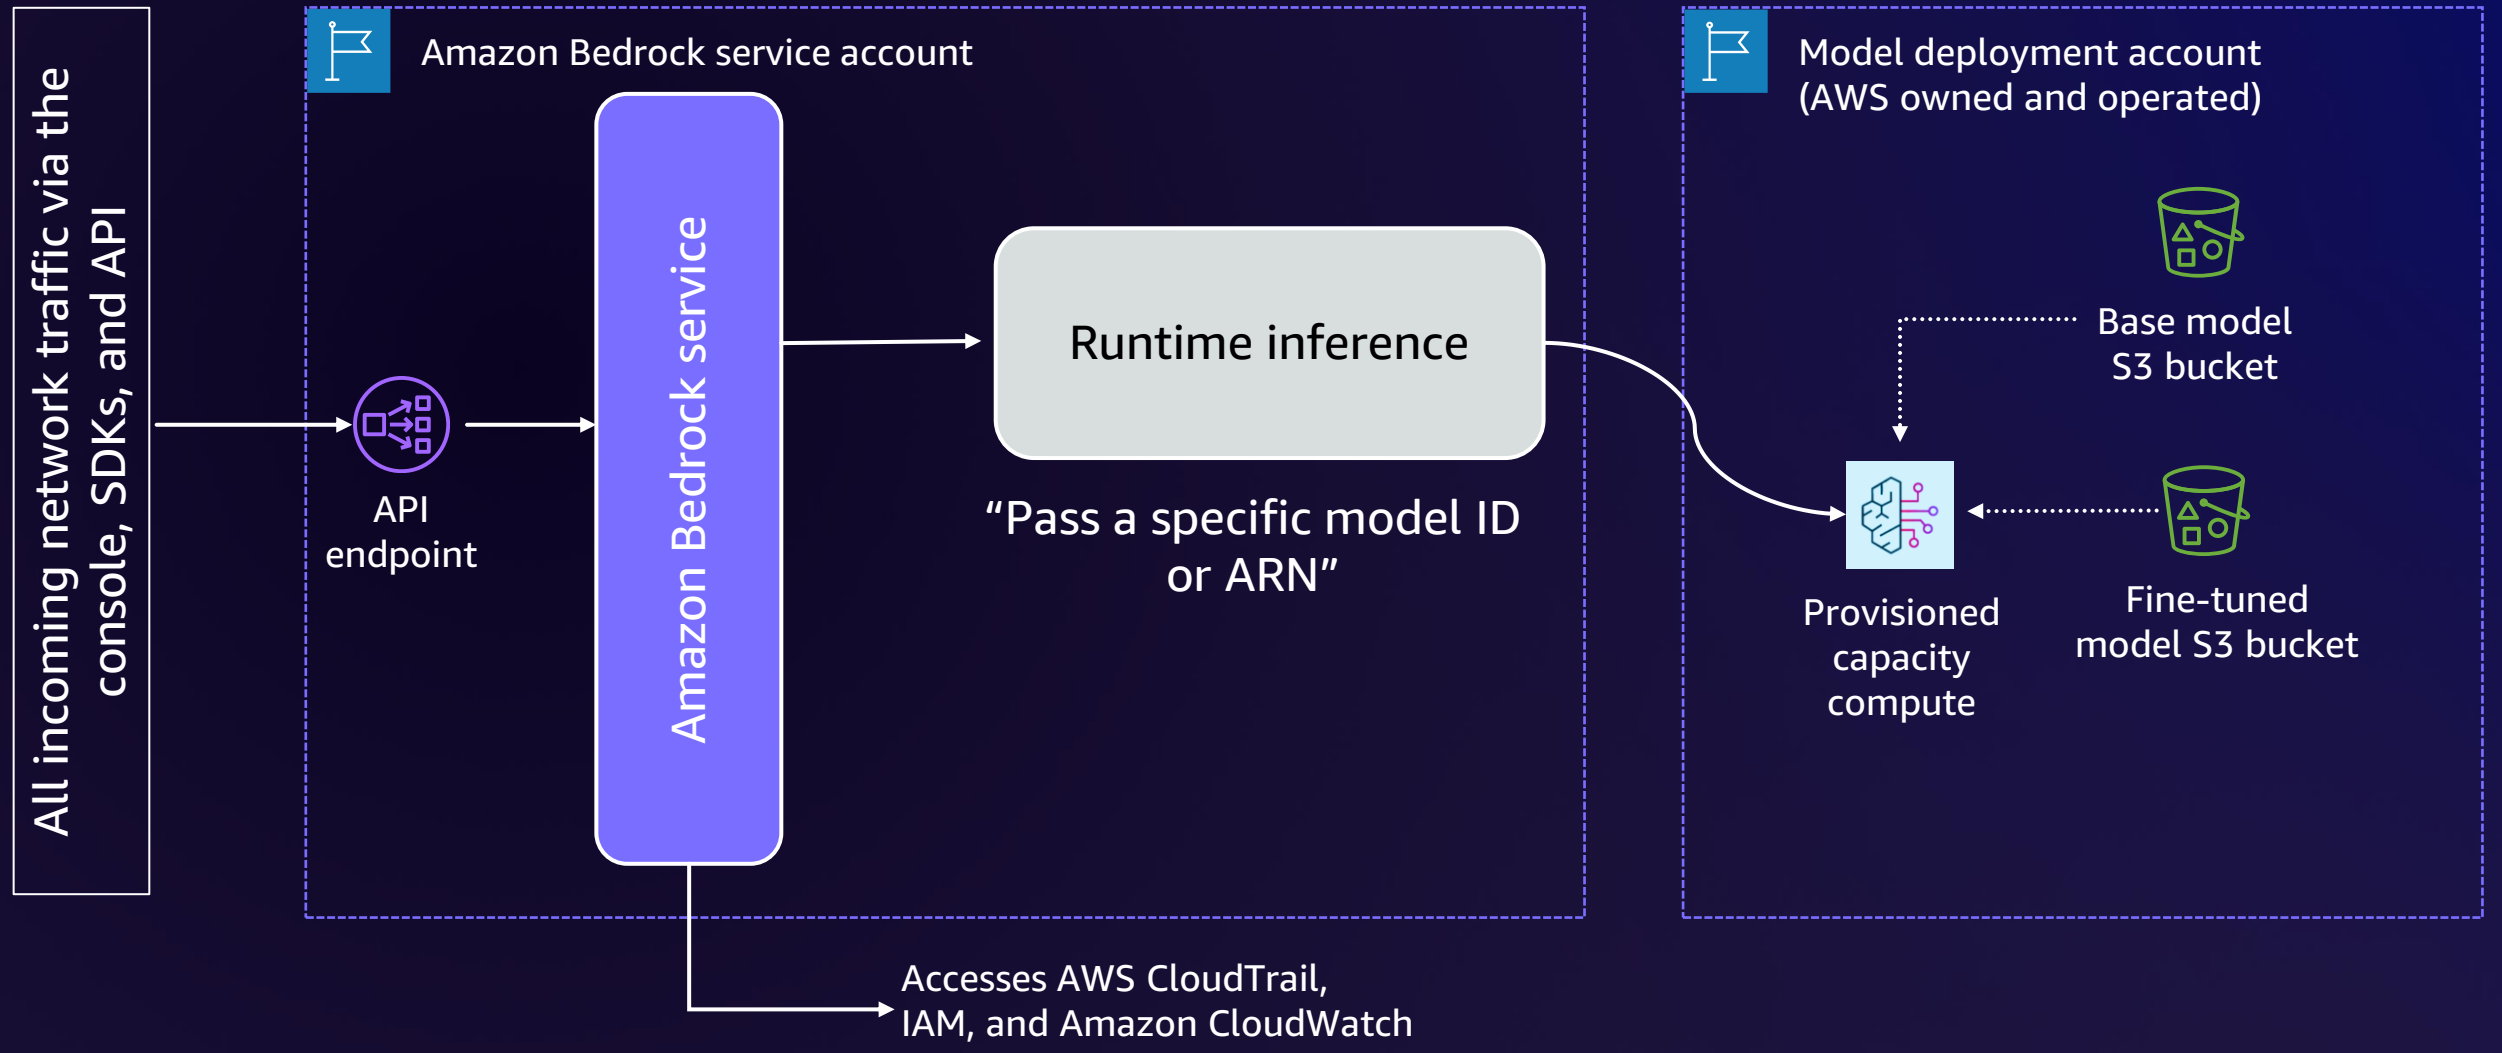

Model fine-tuning architecture overview

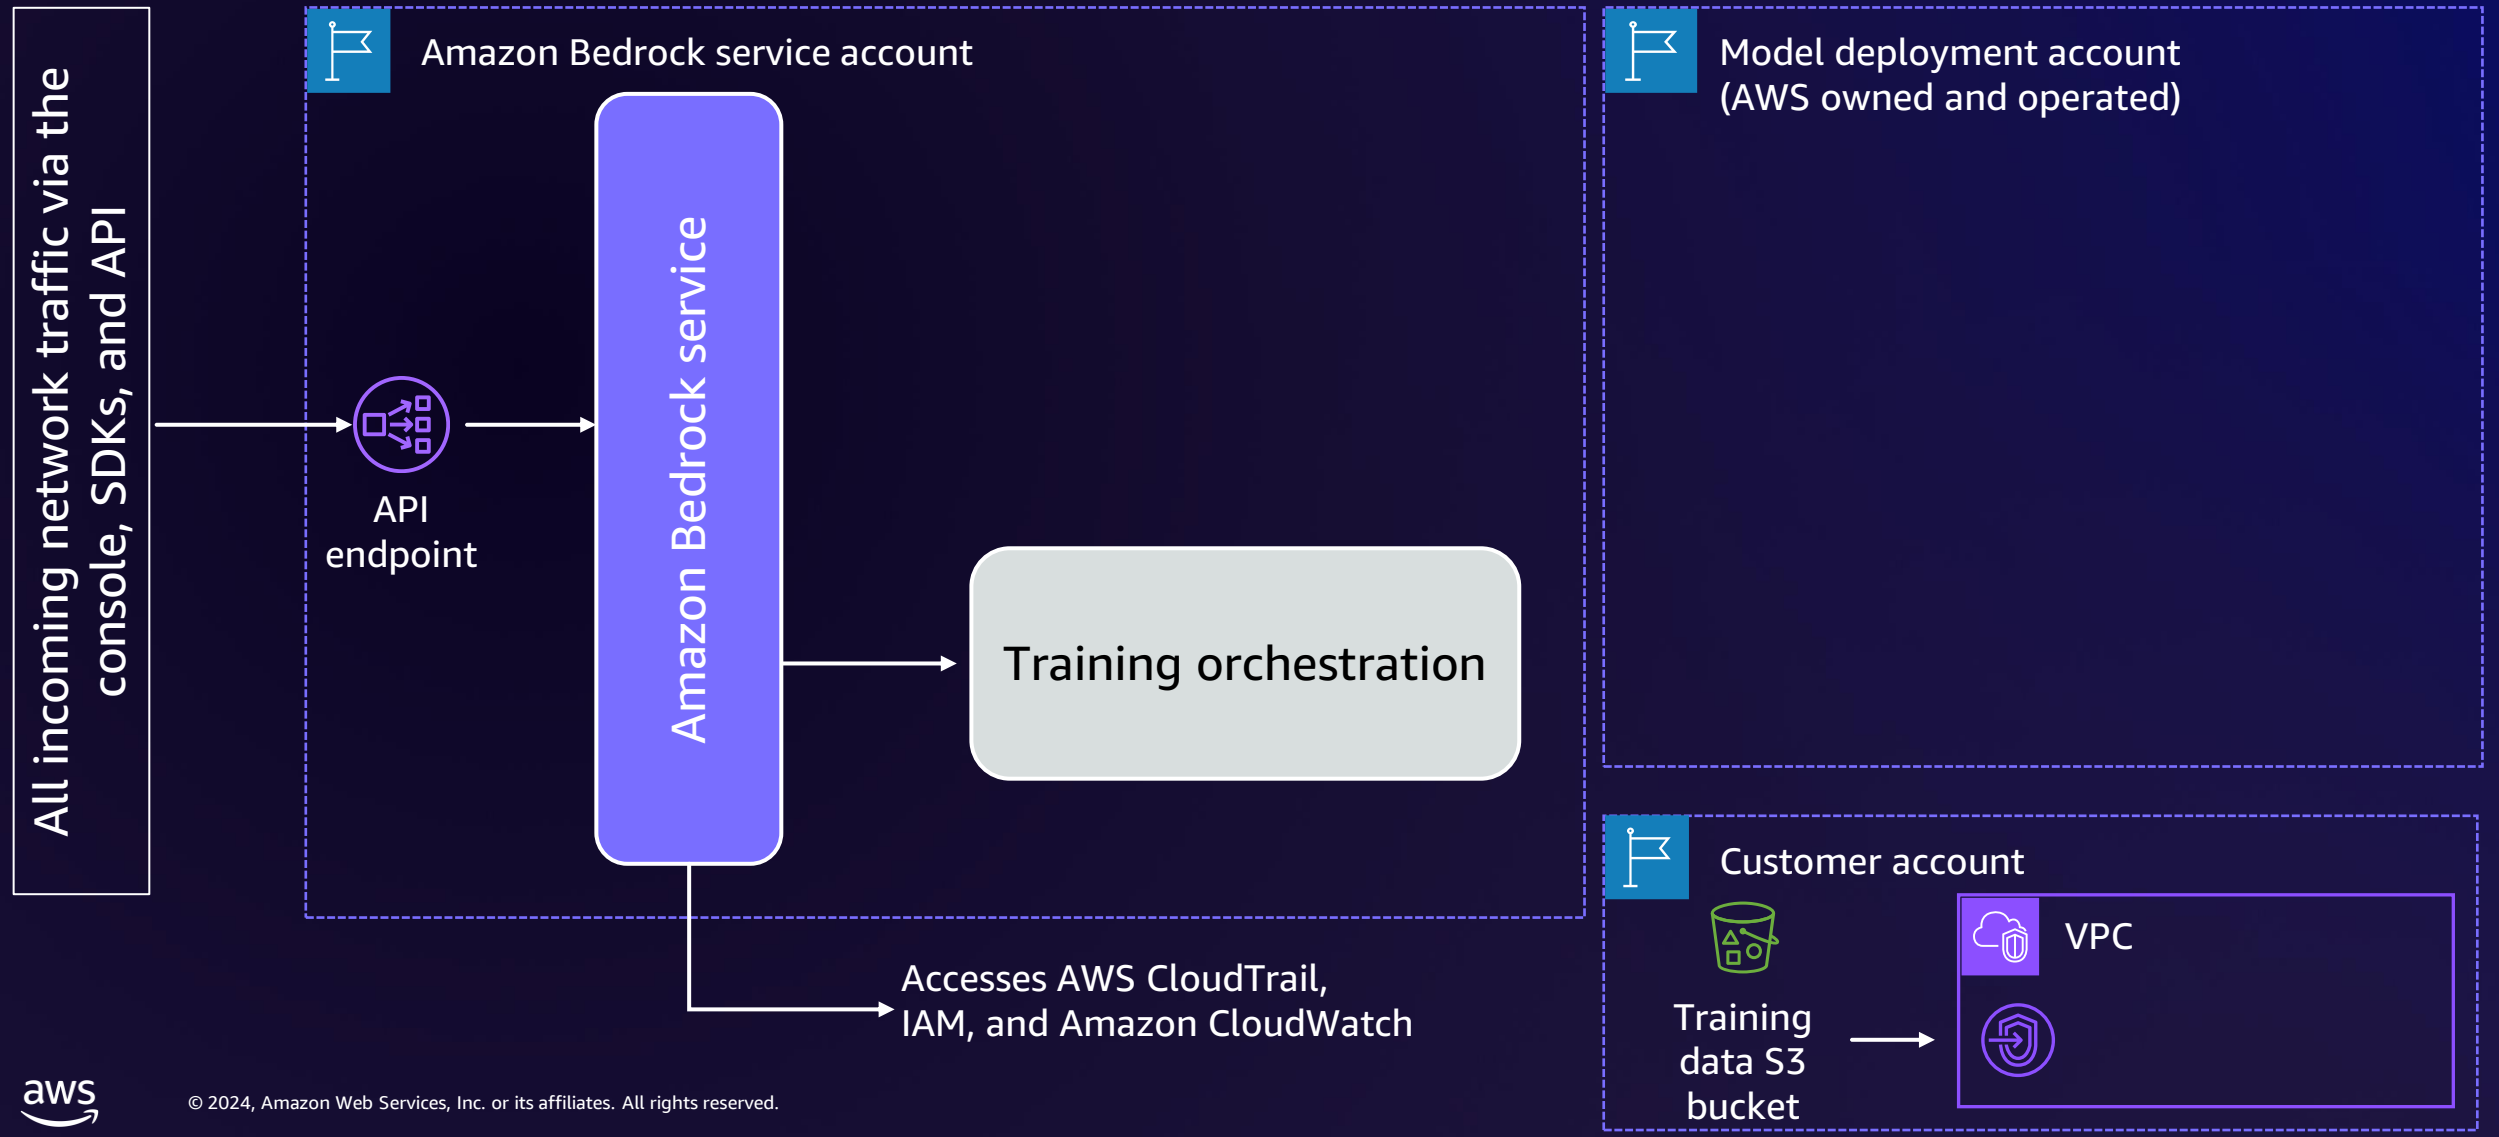

Model fine-tuning architecture overview

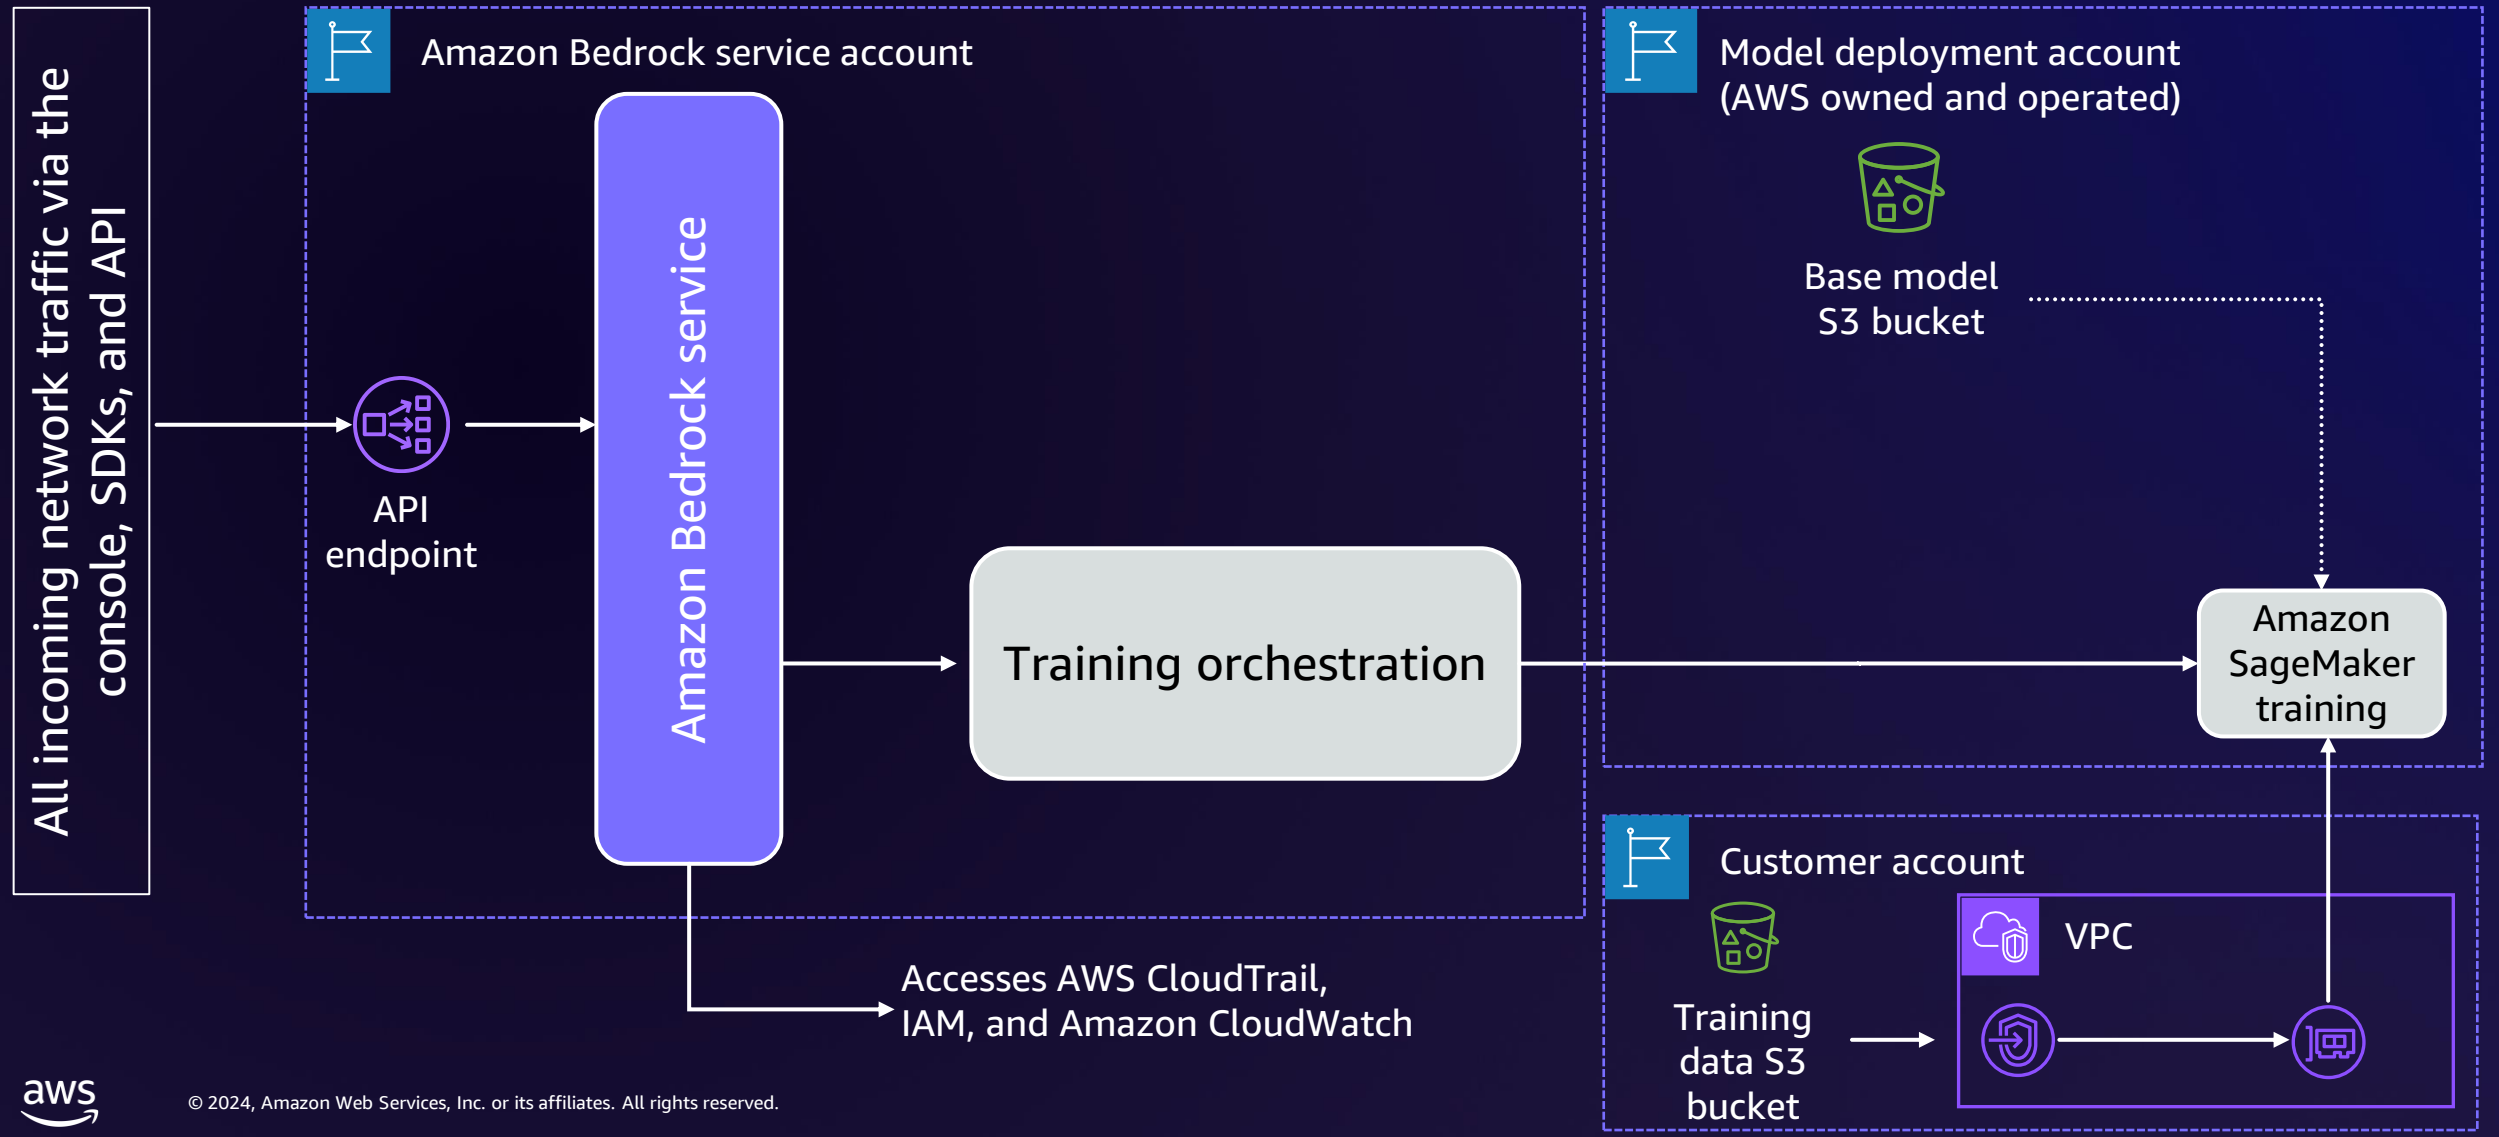

Model fine-tuning architecture overview

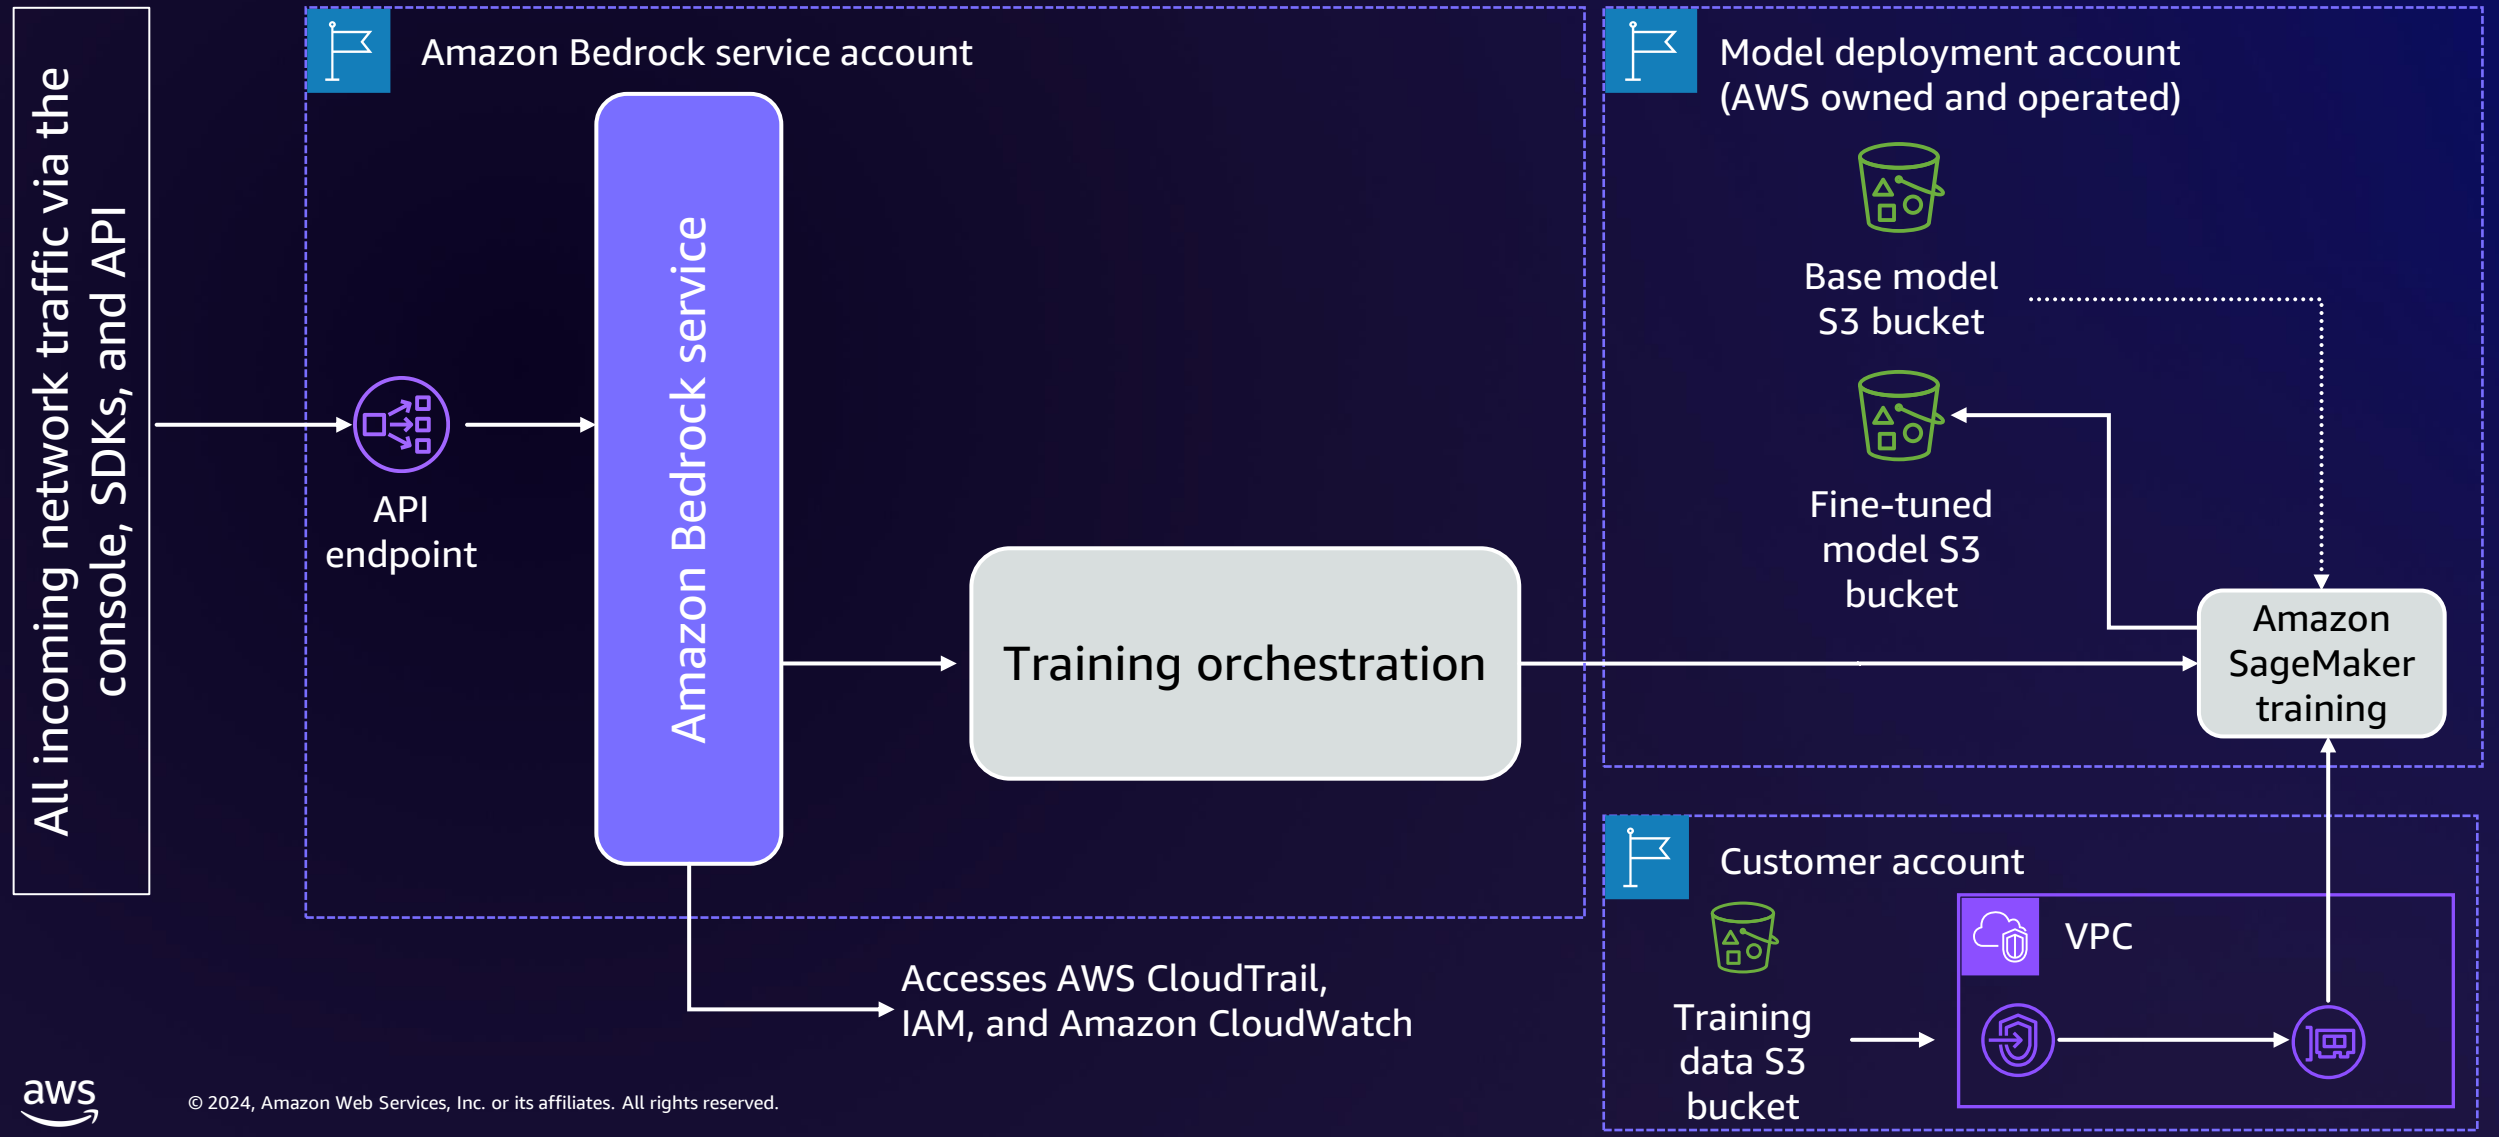

Securing model prompts and responses with Guardrails for Amazon Bedrock

Building generative apps brings new challenges

Undesirable and irrelevant topics

Controversial queries and responses

Toxicity and safety (including brand risk)

Harmful or offensive responses

Privacy protection

Protect user information or sensitive data

Bias/stereotype propagation

Biased results or unfair user outcomes

Organization-specific security and compliance

Consistent prompt and response controls across different models

Guardrails for **Amazon Bedrock**

Implement safeguards customized to your application requirements and responsible AI policies

Apply guardrails to multiple foundation models and agents for Amazon Bedrock

Configure harmful content filtering based on your responsible AI policies

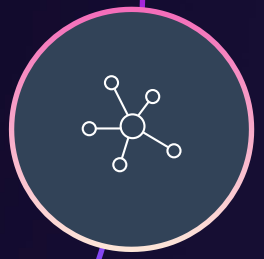

Define and disallow denied topics with short natural language descriptions

Redact sensitive information such as PII in FM responses, along with specific word and regex filters

How it works: Guardrails for Amazon Bedrock

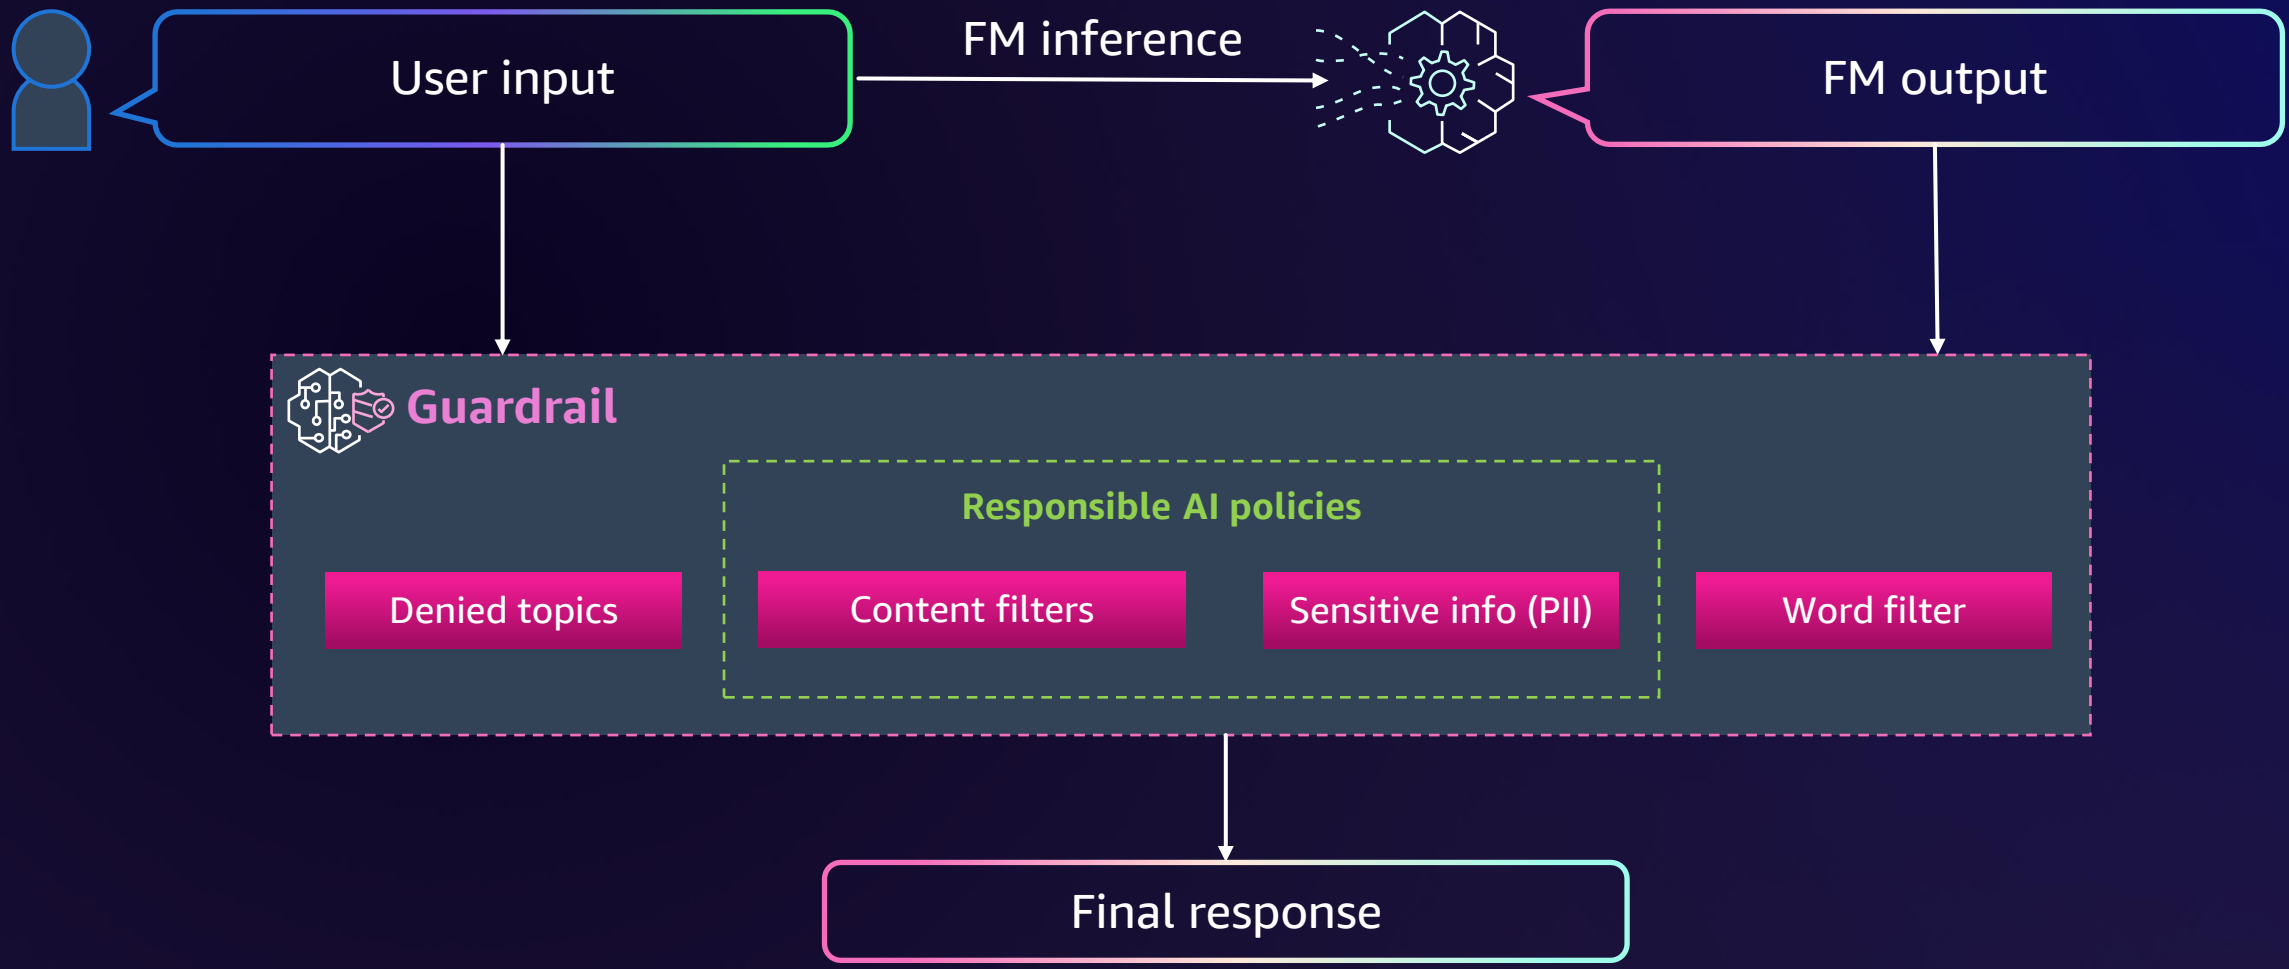

Getting started

Amazon Bedrock code samples

Amazon Bedrock workshop

Amazon Bedrock blog