

AWS re:Invent

NOV. 28 – DEC. 2, 2022 | LAS VEGAS, NV

SEC326

Establishing a data perimeter on AWS, featuring Goldman Sachs

Tatyana Yatskevich

Principal Solutions Architect
Amazon Web Services

Shubham Shukla

Vice President, Cloud Enablement
Goldman Sachs



© 2022, Amazon Web Services, Inc. or its affiliates. All rights reserved.

Agenda

Data perimeter concept

Data perimeter capabilities and implementation

Automating and enforcing data perimeter controls
in the Goldman Sachs environment

Data perimeter concept



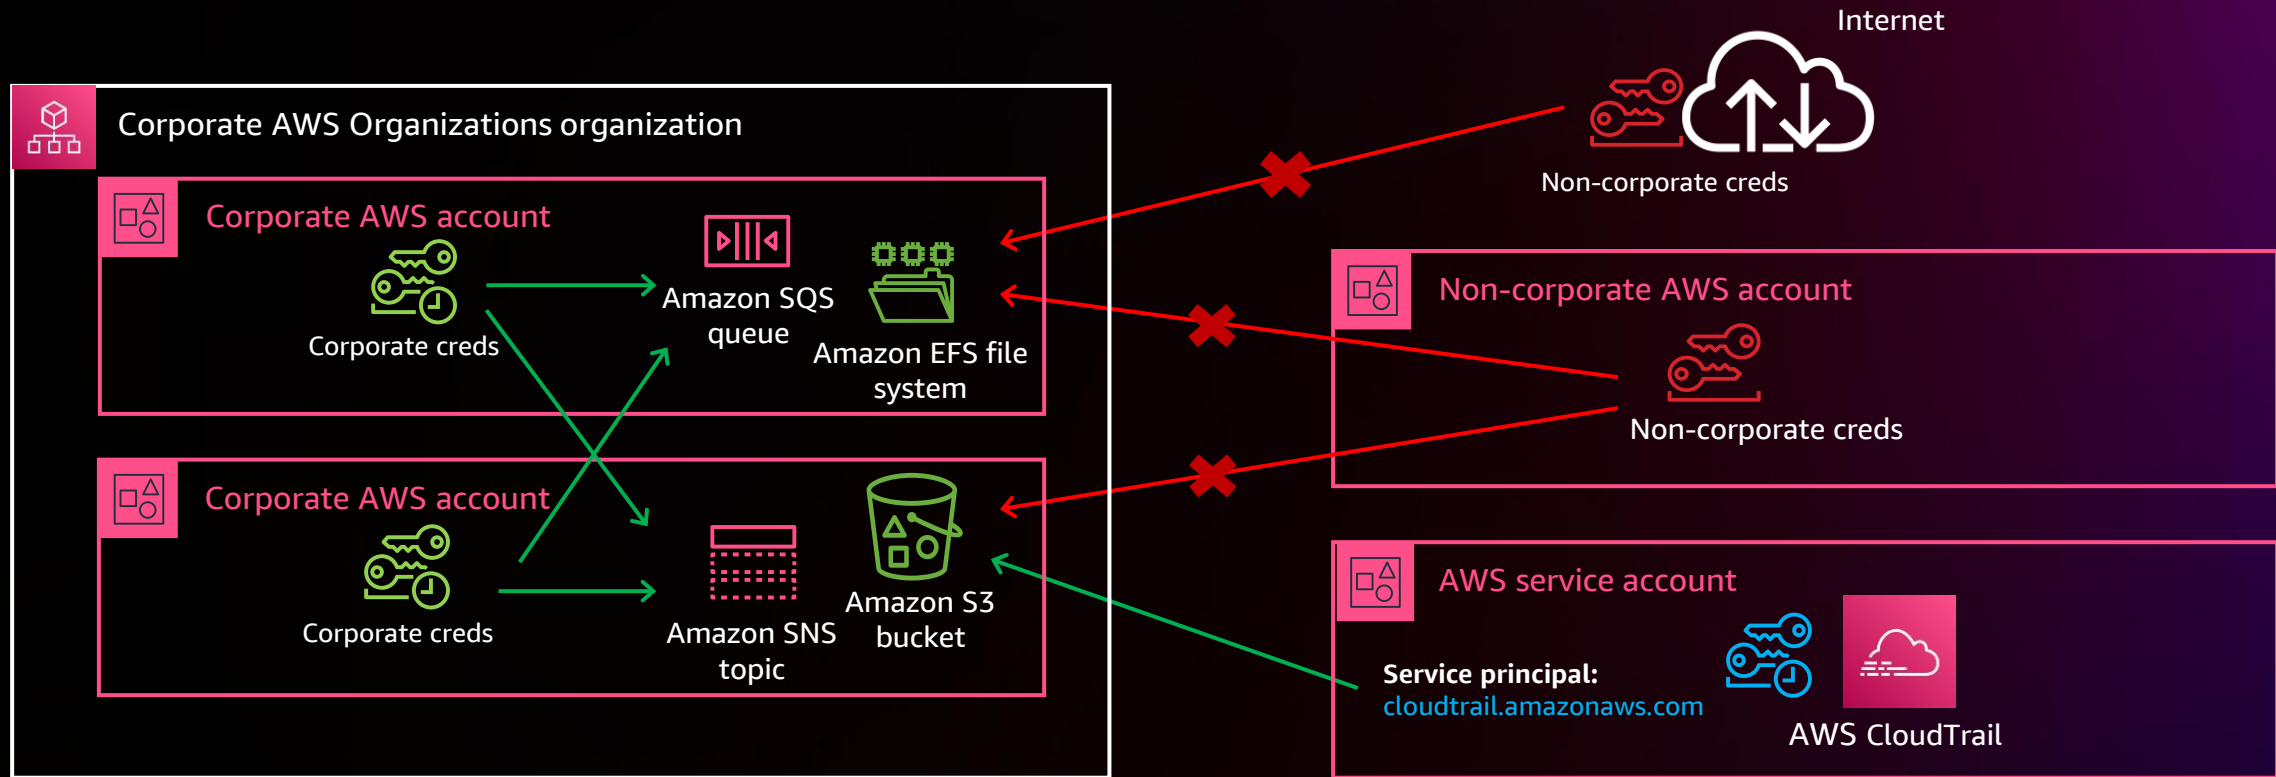
What is a data perimeter?

A set of preventive guardrails in your AWS environment which help ensure that only your **trusted identities** are accessing **trusted resources** from **expected networks**



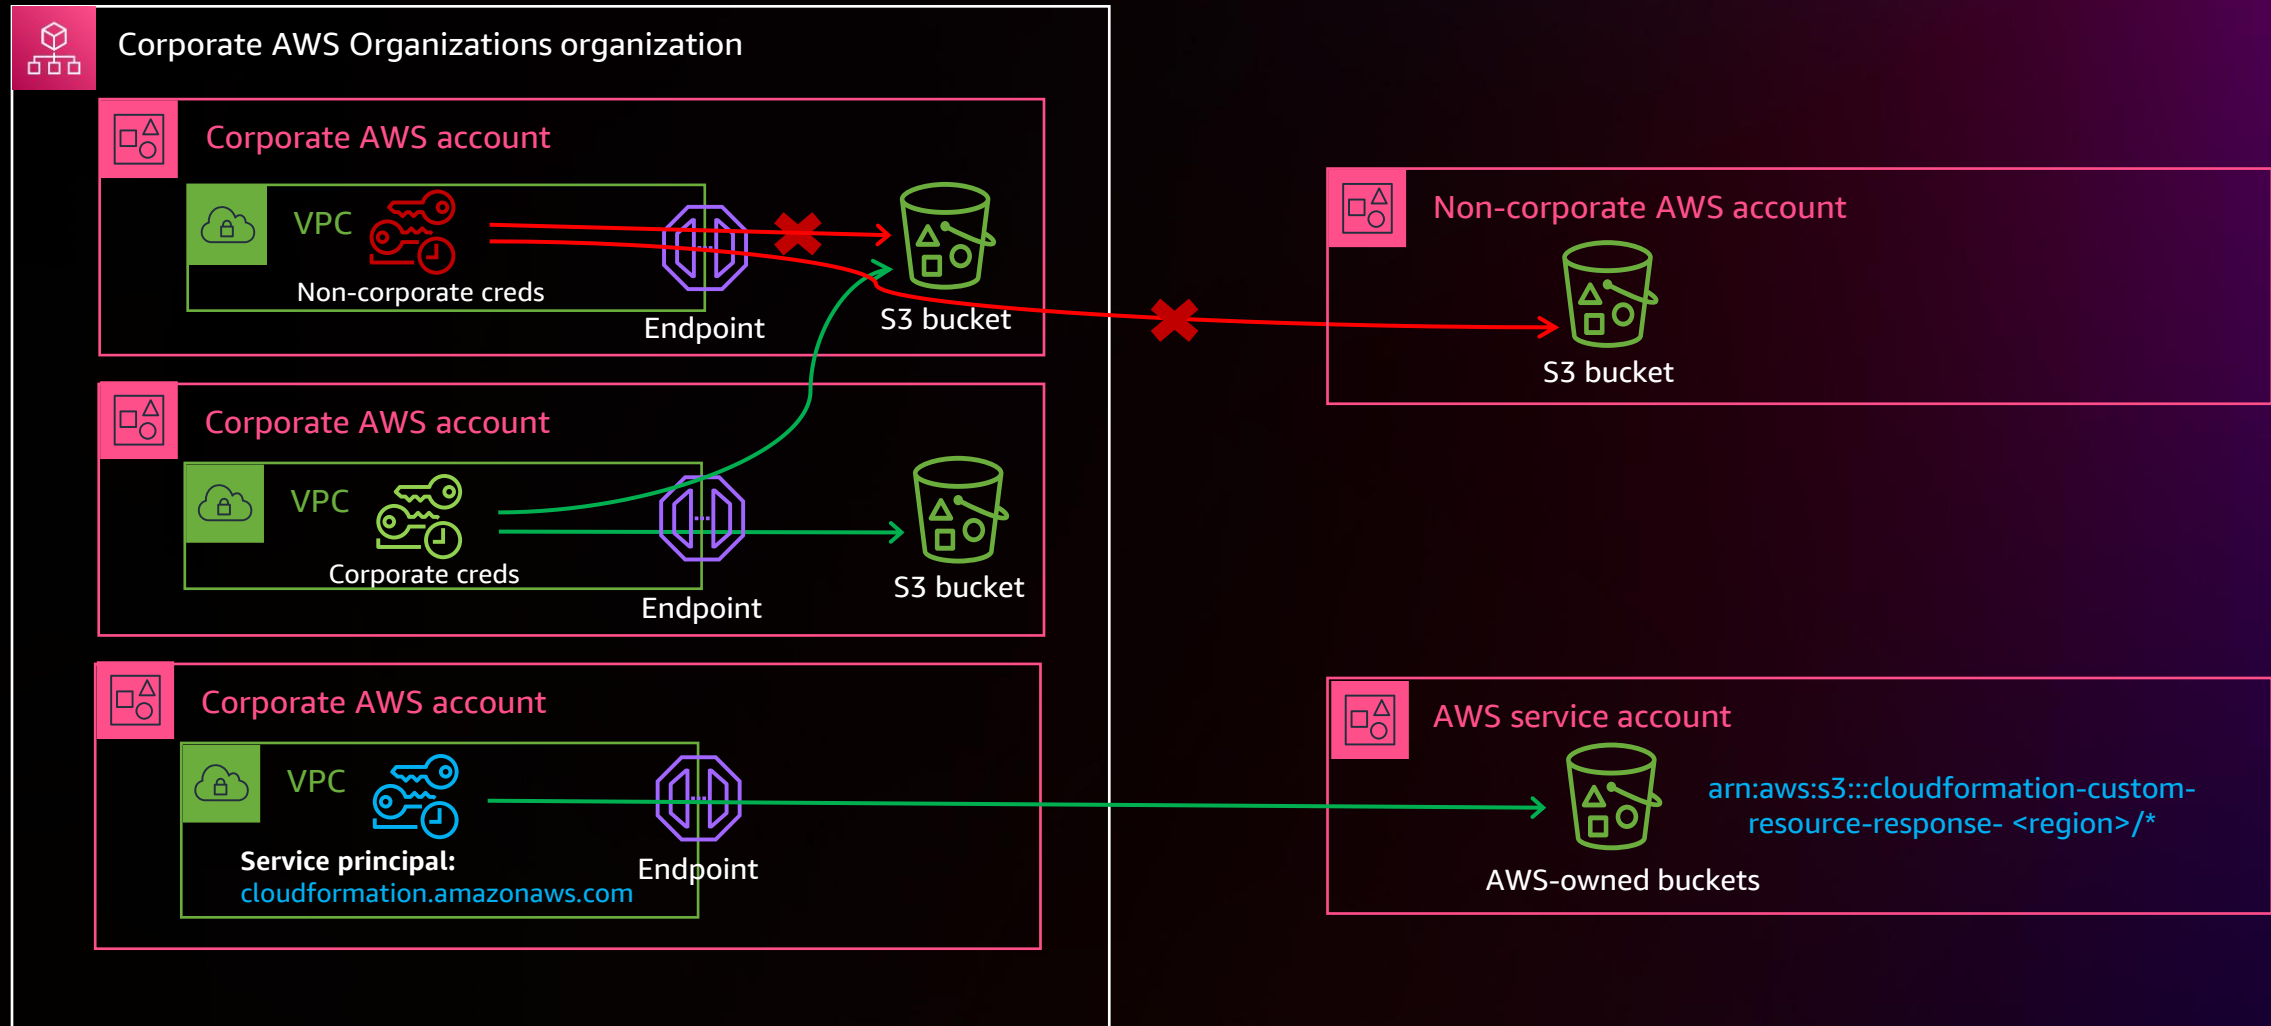
Identity perimeter

ONLY TRUSTED IDENTITIES CAN ACCESS MY RESOURCES



Identity perimeter


ONLY TRUSTED IDENTITIES ARE ALLOWED FROM MY NETWORK



Identity perimeter

ONLY TRUSTED IDENTITIES CAN ACCESS MY RESOURCES

ONLY TRUSTED IDENTITIES ARE ALLOWED FROM MY NETWORK



Identity

My corporate credentials

aws:PrincipalOrgID

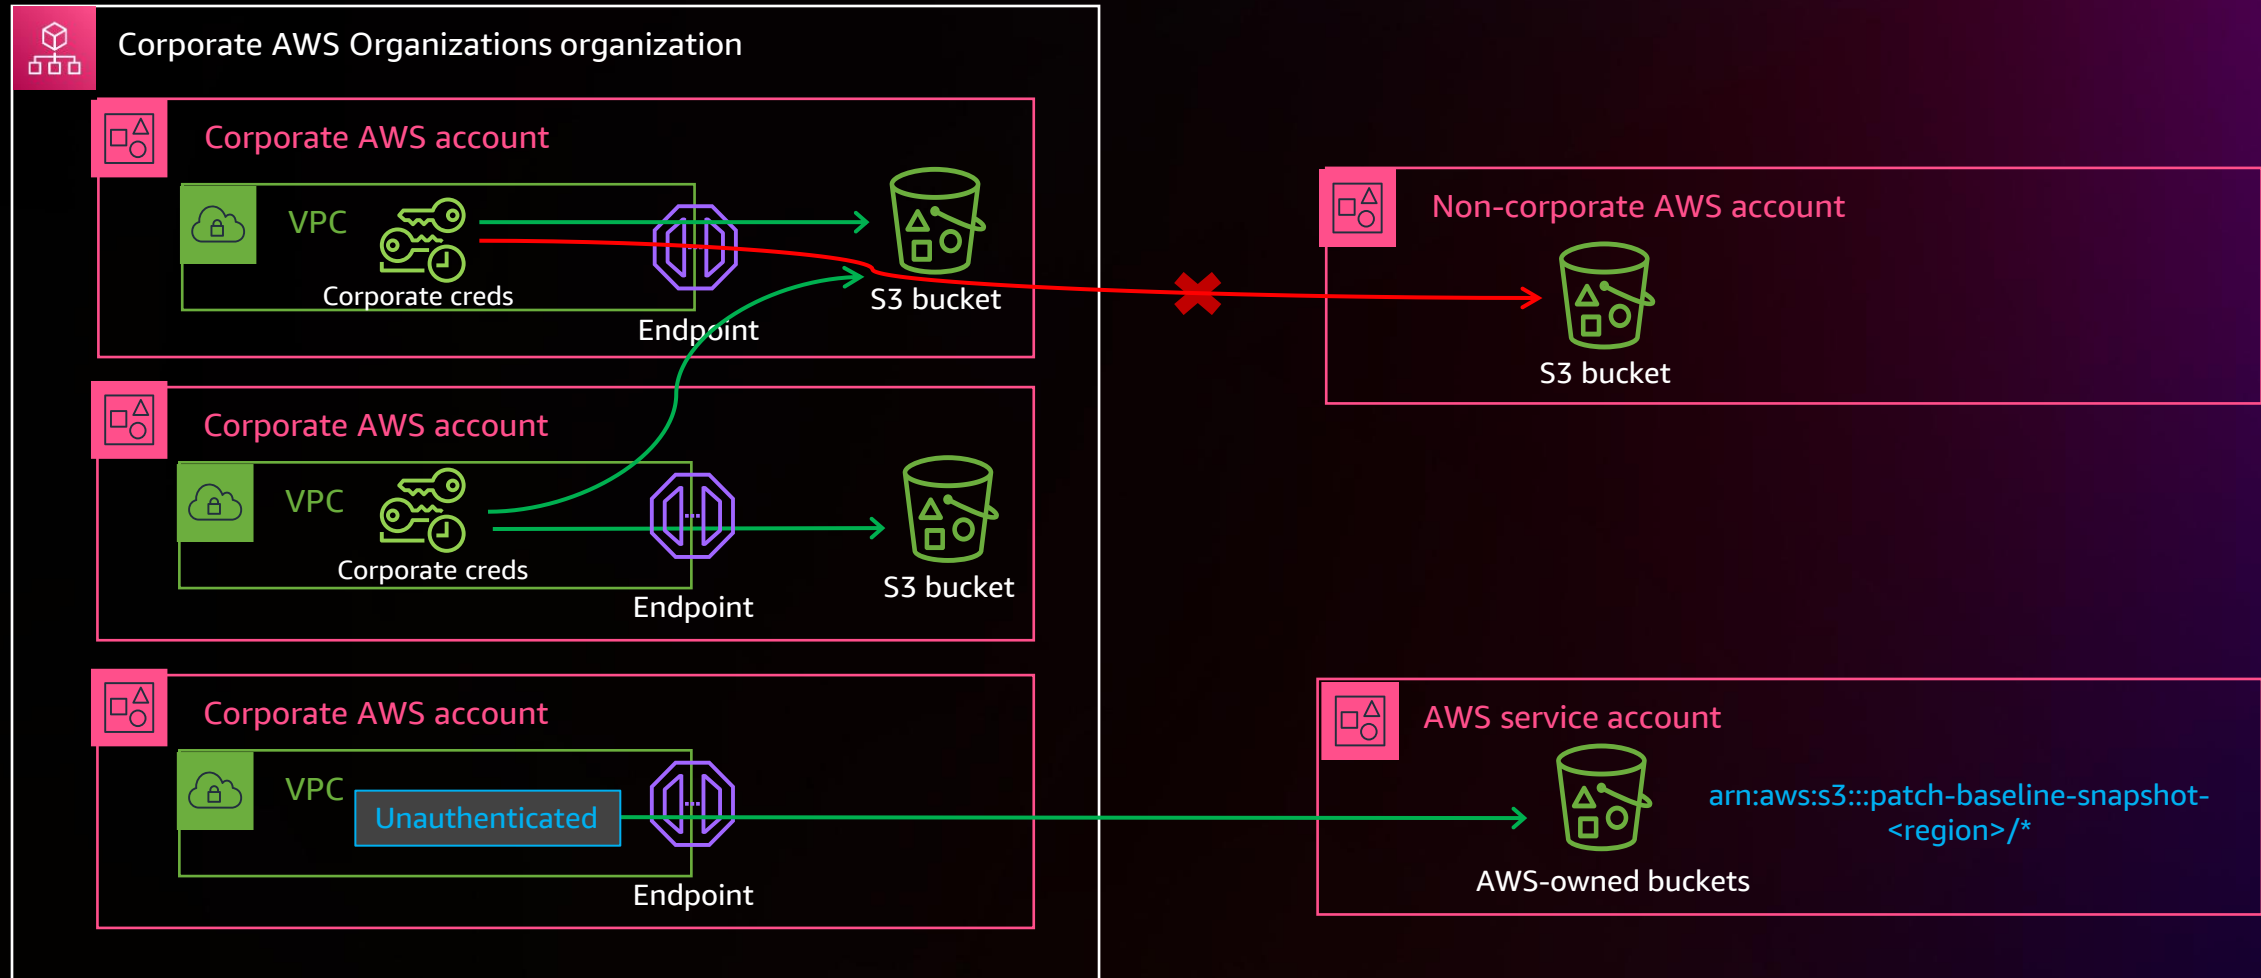
AWS service principal

aws:PrincipalIsAWSService

Resource perimeter

MY IDENTITIES CAN ACCESS ONLY TRUSTED RESOURCES

ONLY TRUSTED RESOURCES CAN BE ACCESSED FROM MY NETWORK



Resource perimeter

MY IDENTITIES CAN ACCESS ONLY TRUSTED RESOURCES

ONLY TRUSTED RESOURCES CAN BE ACCESSED FROM MY NETWORK



Resource

My corporate resources

`aws:ResourceOrgID`

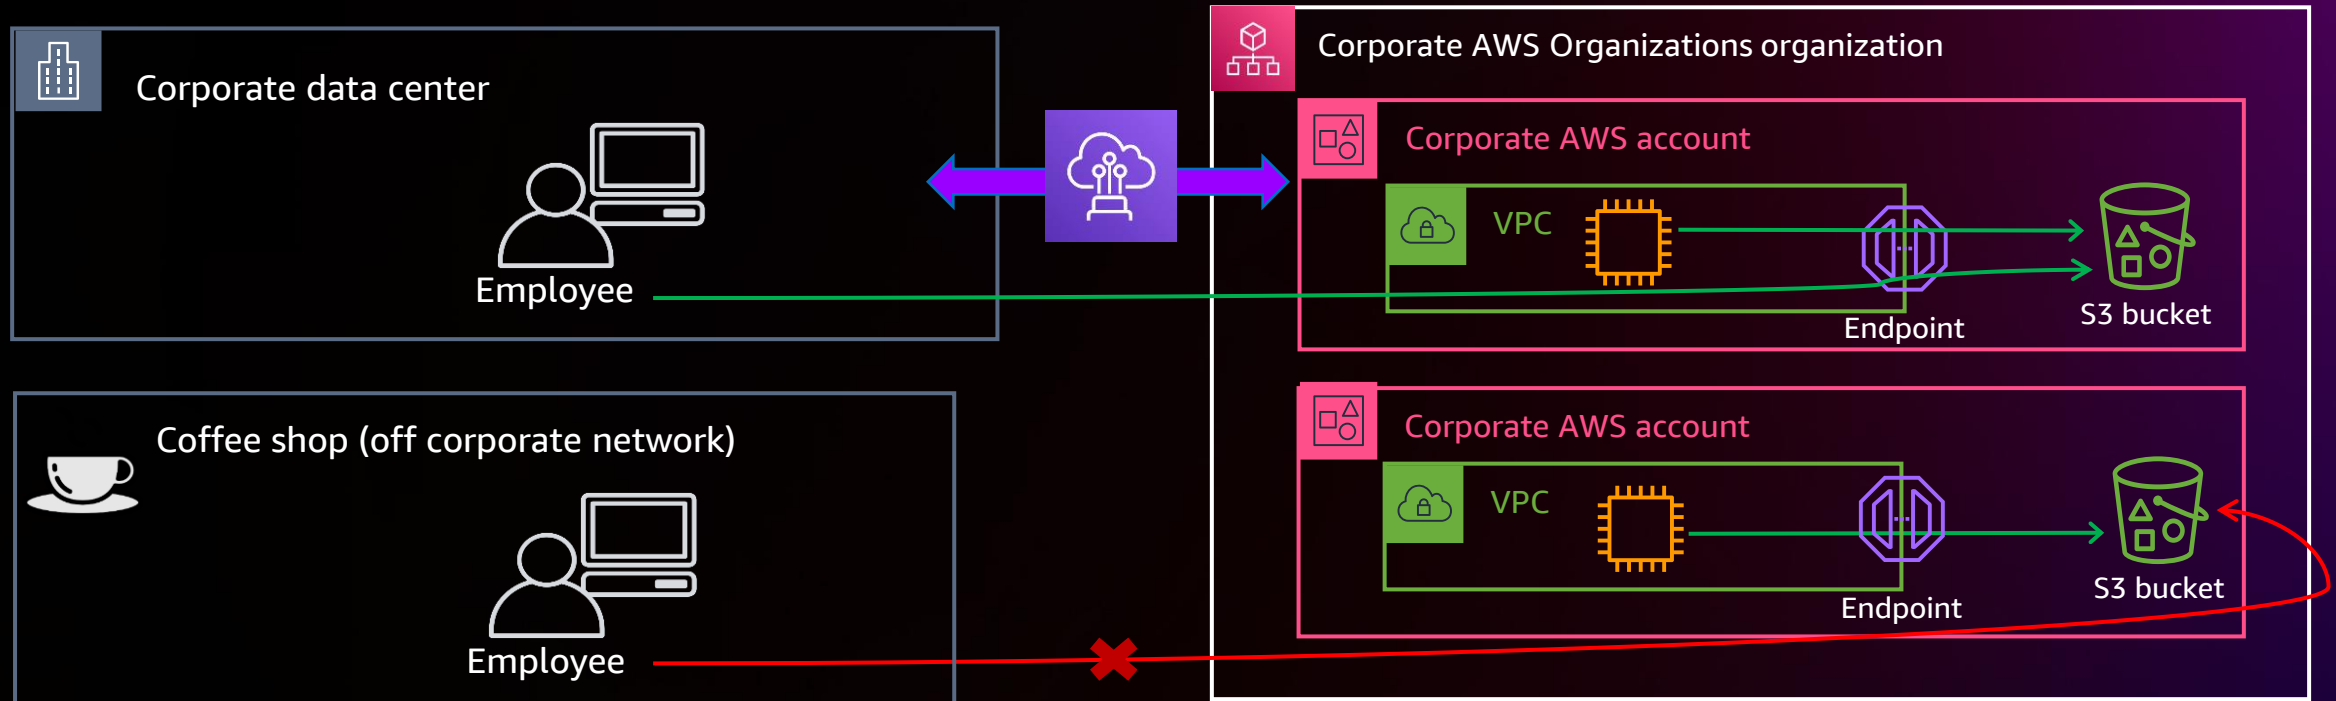
AWS resources

`arn:aws:s3:::patch-baseline-snapshot-
<region>/*`

Network perimeter

MY RESOURCES CAN ONLY BE ACCESSED FROM EXPECTED NETWORKS

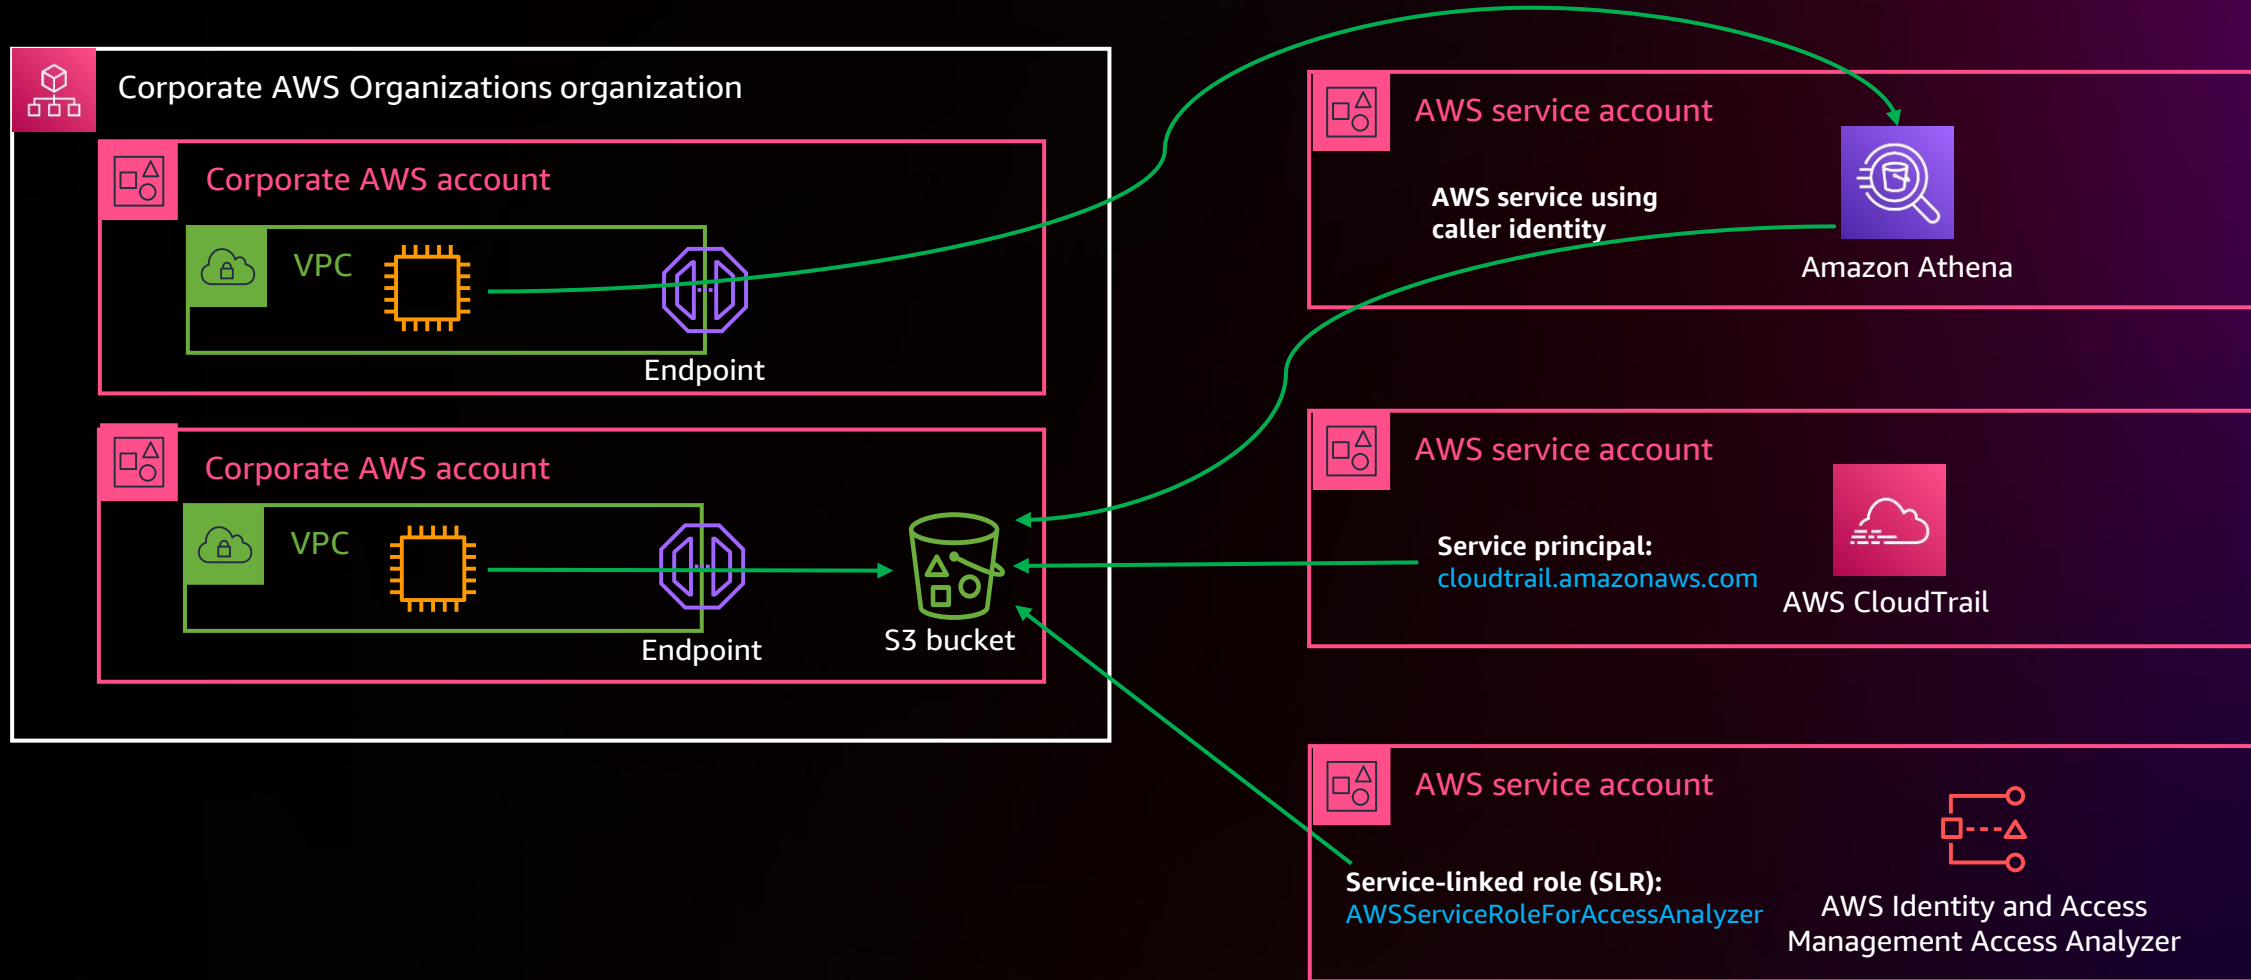
MY IDENTITIES CAN ACCESS RESOURCES ONLY FROM EXPECTED NETWORKS



Network perimeter

MY RESOURCES CAN ONLY BE ACCESSED FROM EXPECTED NETWORKS

MY IDENTITIES CAN ACCESS RESOURCES ONLY FROM EXPECTED NETWORKS



Network perimeter

MY RESOURCES CAN ONLY BE ACCESSED FROM EXPECTED NETWORKS

MY IDENTITIES CAN ACCESS RESOURCES ONLY FROM EXPECTED NETWORKS



Network

My corporate data center/My corporate AWS network

`aws:SourceIp/aws:SourceVpc`

AWS service using caller identity

`aws:viaAWSService`

AWS service principal

`aws:PrincipalIsAWSService`

SLR

`arn:aws:iam::012345678901:role/aws-service-role/*`

Data perimeter controls

Perimeter	Intent/Control objective	Applied on	Using	Primary IAM feature
Identity	Only trusted identities can access my resources	Resources	Resource-based policy	aws:PrincipalOrgID aws:PrincipalIsAWSService
	Only trusted identities are allowed from my network	Network	VPC endpoint policy	aws:PrincipalOrgID aws:PrincipalIsAWSService
Resource	My identities can access only trusted resources	Identities	SCP policy	aws:ResourceOrgID
	Only trusted resources can be accessed from my network	Network	VPC endpoint policy	aws:ResourceOrgID
Network	My identities can access resources only from expected networks	Identities	SCP policy	aws:SourceIp aws:SourceVpc/aws:SourceVpce aws:ViaAWSService
	My resources can only be accessed from expected networks	Resources	Resource-based policy	aws:SourceIp aws:SourceVpc/aws:SourceVpce aws:ViaAWSService aws:PrincipalIsAWSService

Automating and enforcing data perimeter controls in the Goldman Sachs environment

Who is Goldman Sachs?

Goldman Sachs

- Leading investment bank and financial services company
- \$2T+ AUS, 100 Locations
- Technology-first – 12,000+ engineers

Cloud enablement @ Goldman Sachs

- 80+ enterprises
- Advisory, implementation, and security
- Continuous deployment – Cloud Fast Track

Agenda

The Problem: Cloud security at scale

Our answer: Cloud Fast Track

Available tools in the AWS toolbox: Data perimeter

Guardrails: Policy as code for data perimeter

Conclusion: Examples, demo, learnings

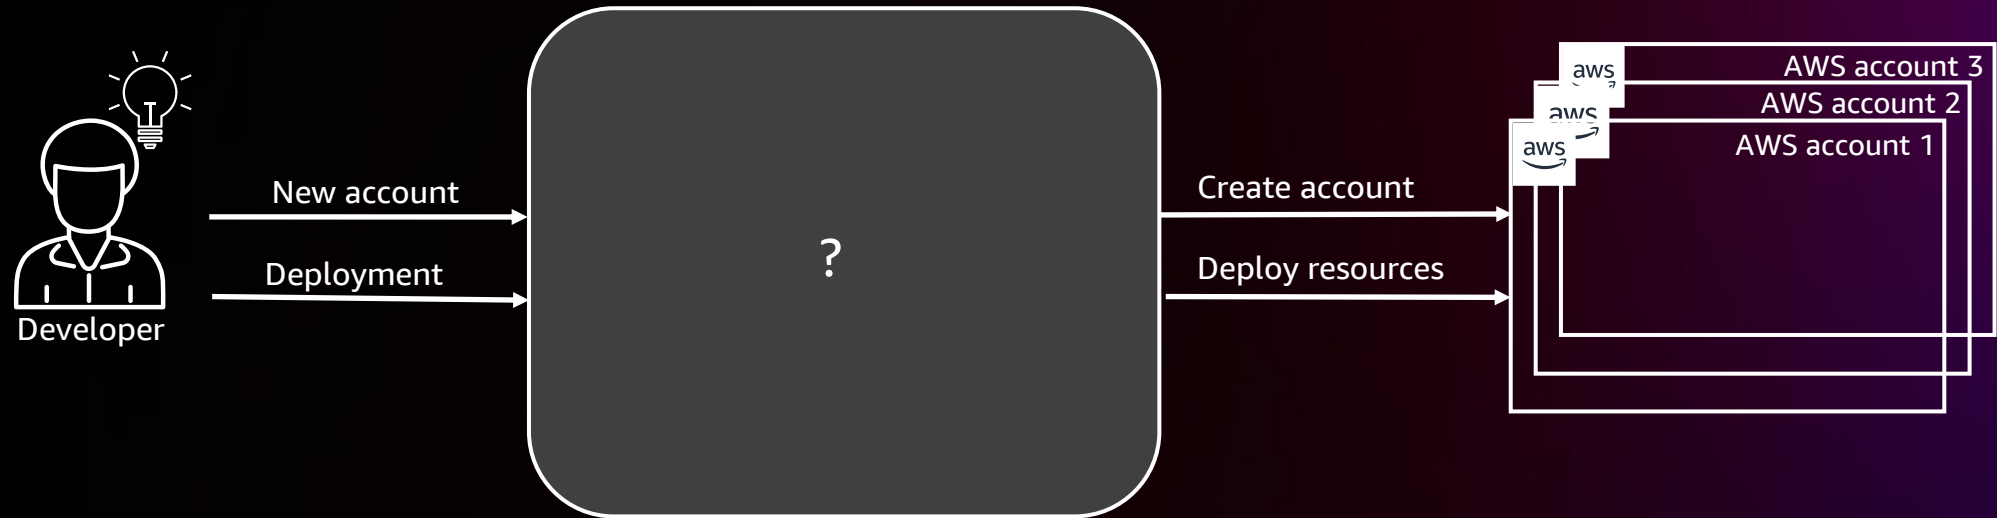
The Problem

CLOUD SECURITY AT SCALE



Desired state

WHAT IF...




Tenets

SECURE CONTINUOUS DEPLOYMENT AT SCALE



Self-service

We succeed if you never
have to talk to us



Shared responsibility

We provide a secure
platform, you build a
secure app



Secure by design

Implement policy as
code to ensure firm-
compliant resources

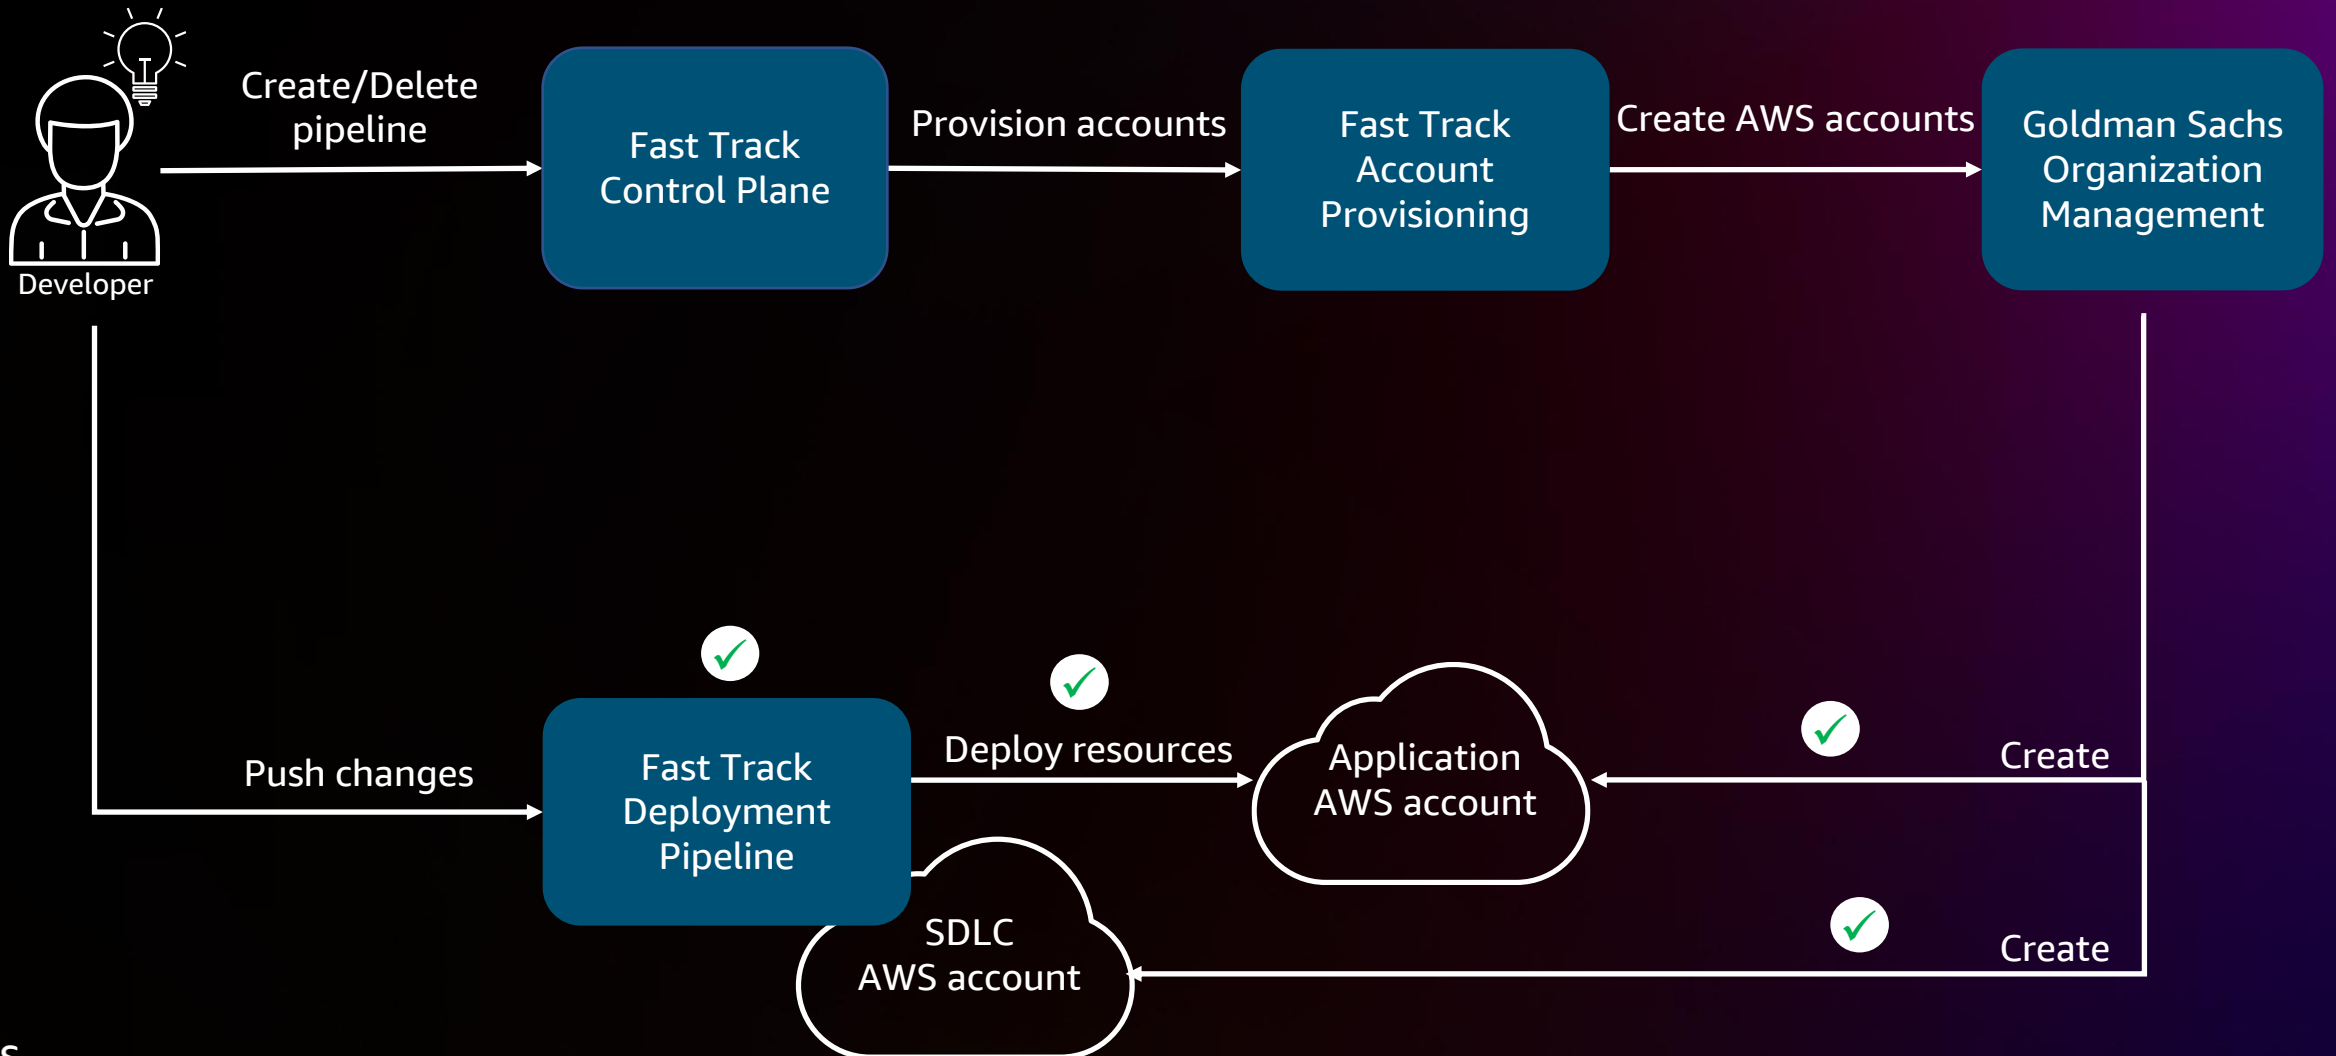


Cloud native

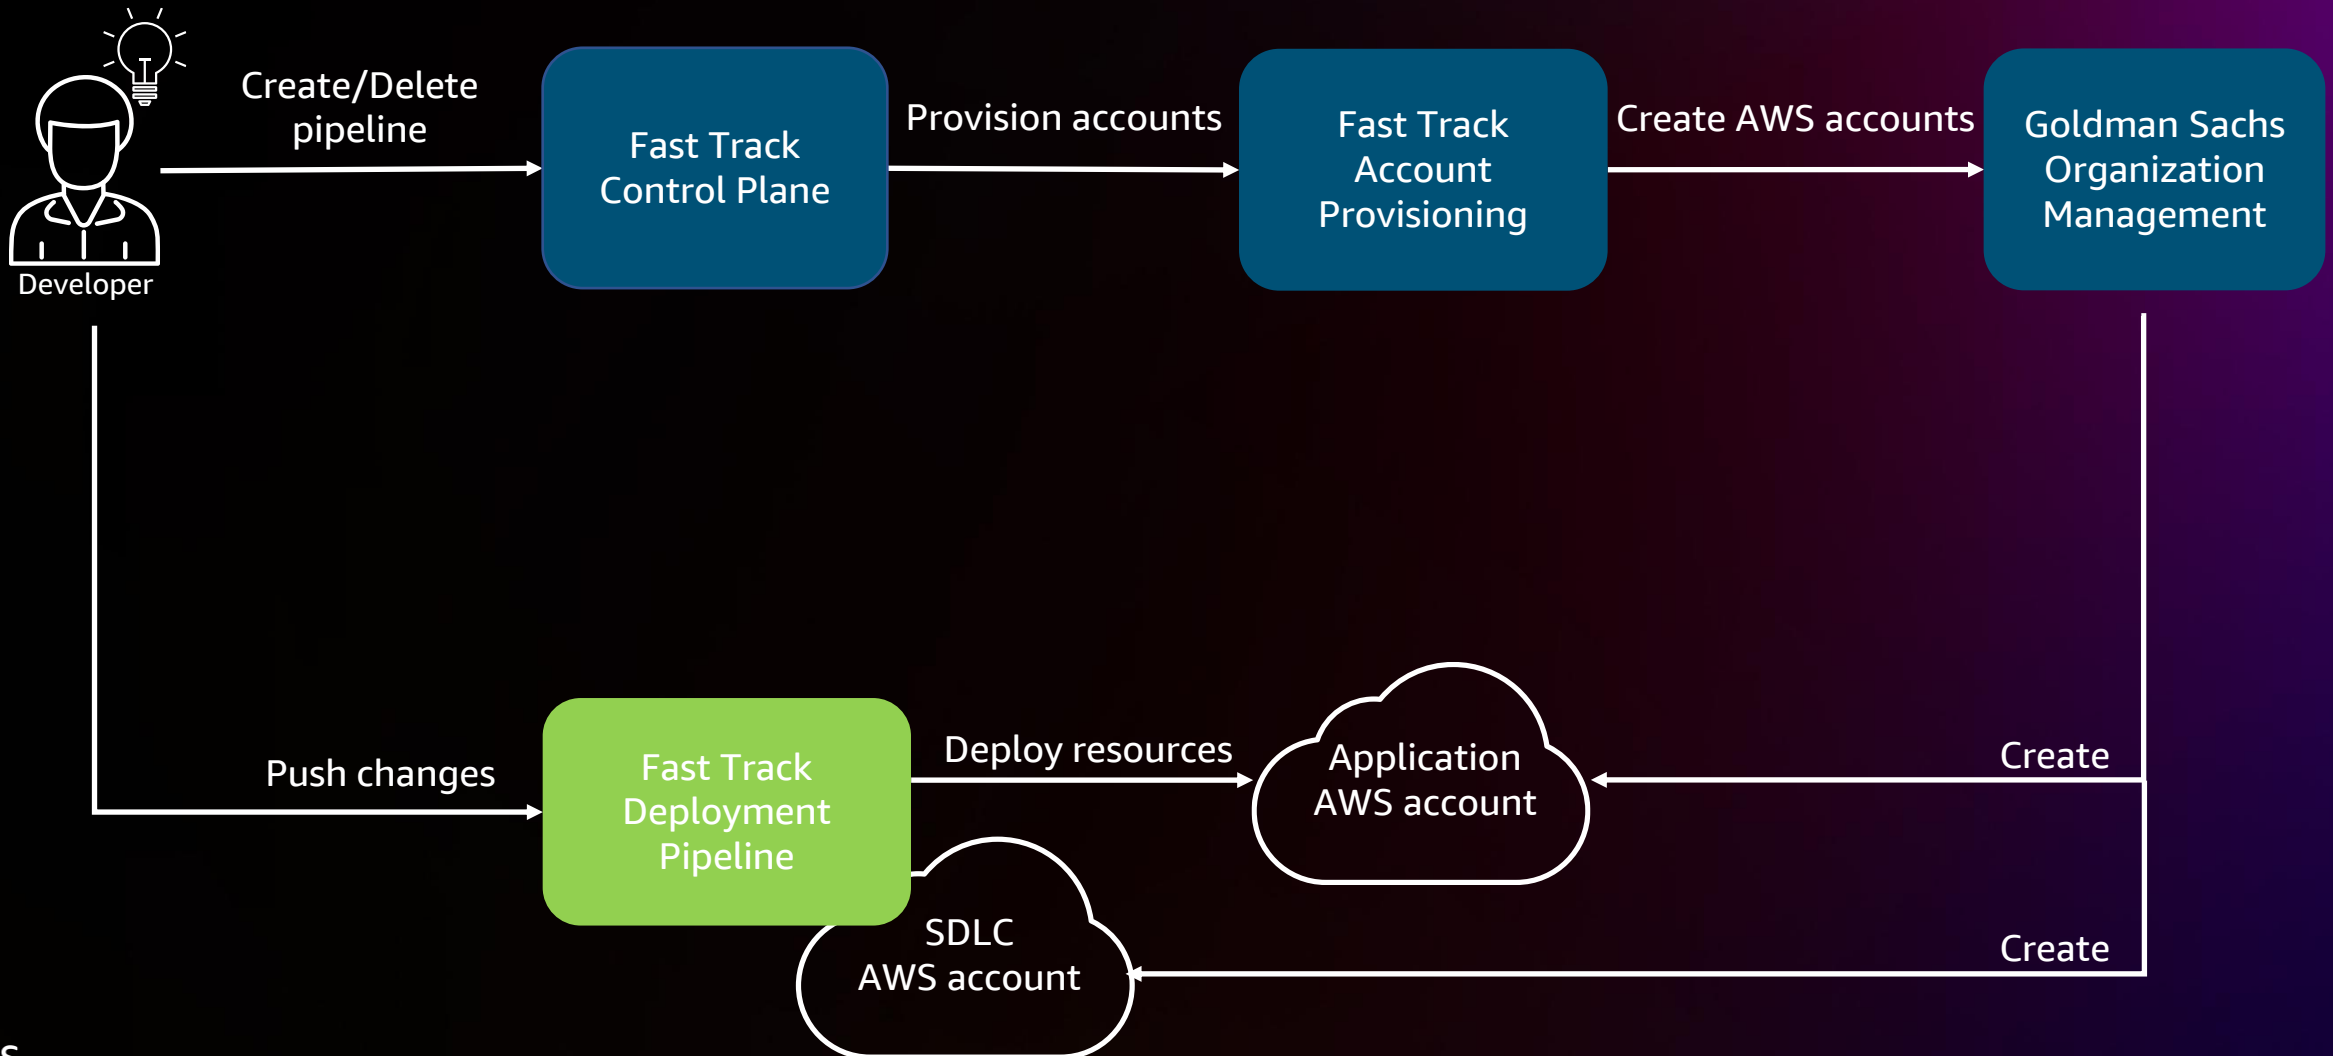
Use public tooling and
documentation

Enter Cloud Fast Track

CUSTOMER JOURNEY

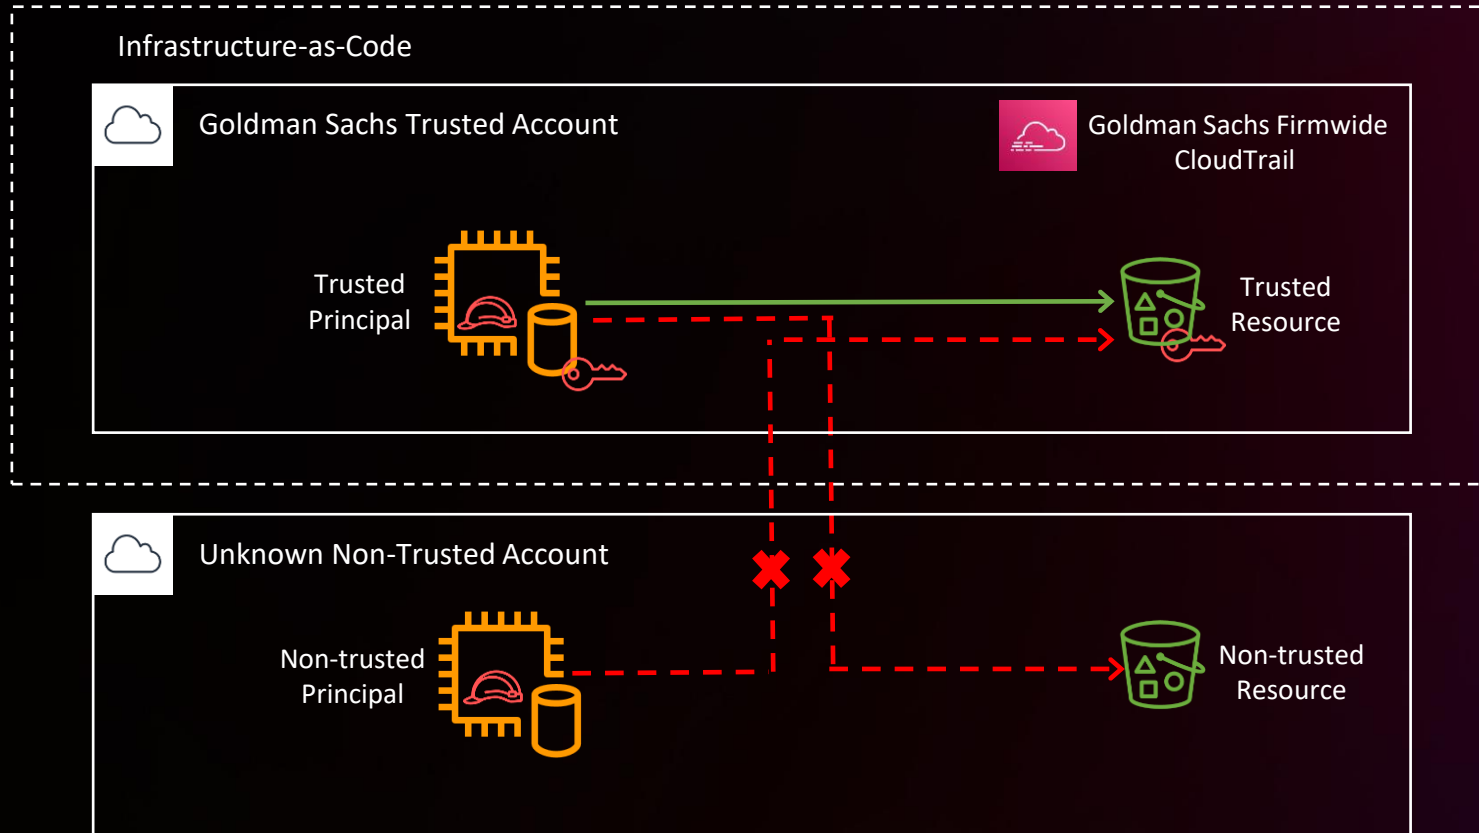


Nuts and bolts: Checkpoint 5



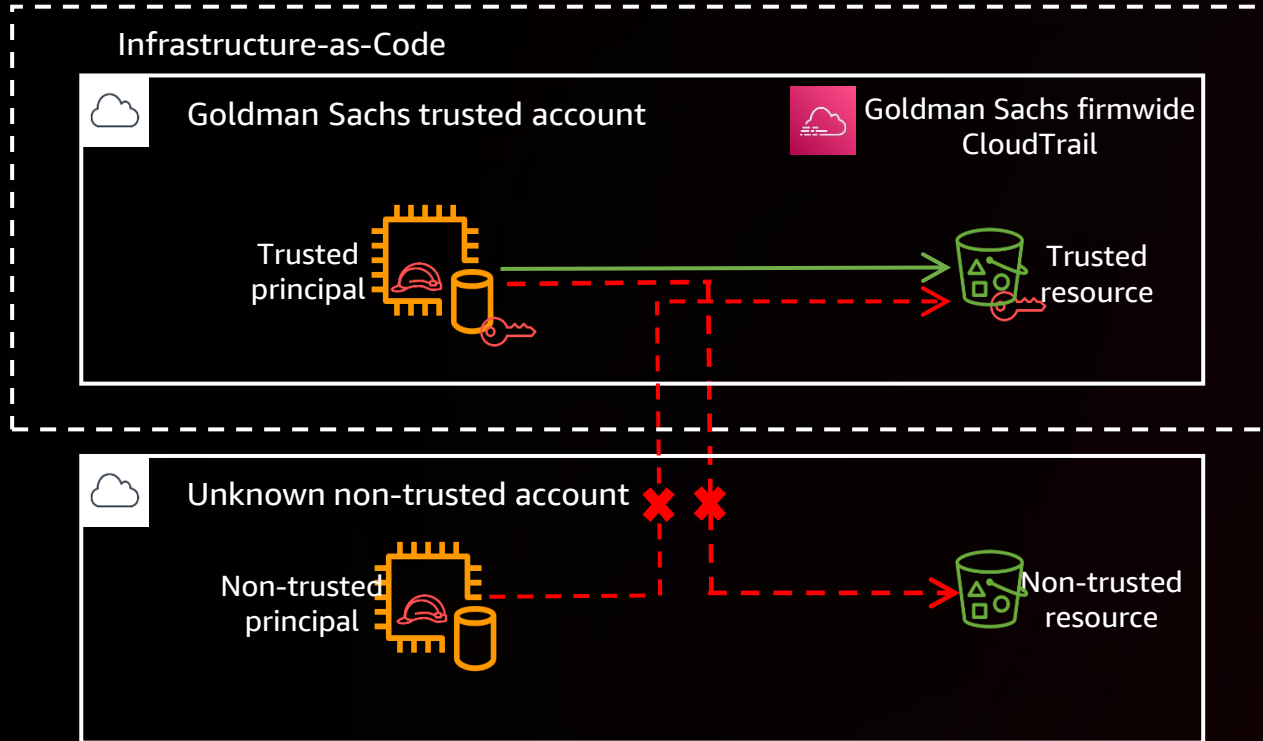
Secure deployment – Threats

WHAT THREAT VECTORS ARE WE SECURING AGAINST?



Secure deployment – Control examples

SO HOW DO YOU SECURE YOUR DEPLOYMENTS?



Thou shall:

- Always attach a vended permission boundary to your `AWS::IAM::Role`
- Only allow actions on `AWS::KMS::Key` from principals in the trusted org path
- Only use vended `AWS::EC2::VPC`
- Always encrypt `AWS::S3::Bucket` with KMS-CMK

AWS tools – Security engineering

AWS OFFERINGS



Trusted identities

Principals within your AWS accounts, or AWS services acting on your behalf

- `aws:PrincipalOrgID`
- `aws:PrincipalOrgPaths`
- `aws:PrincipalAccount`
- `aws:PrincipalIsAWSService`



Trusted resources

Resources owned by your AWS accounts or by AWS services acting on your behalf

- `aws:ResourceOrgID`
- `aws:ResourceOrgPaths`
- `aws:ResourceAccount`



Expected networks

Your on-premises data centers and virtual private clouds (VPCs), or networks of AWS services acting on your behalf

- `aws:SourceIp`
- `aws:SourceVpc`
- `aws:SourceVpce`

How do you enforce controls?

SO YOU WANT TO WRITE GUARDRAILS?

- Guardrails: Policy as code
 - Coarse-grained controls
 - Written in Rego and evaluated using Open Policy Agent (OPA)
- Analogy
 - No matter how bad you drive, you can never drive outside the guardrails
- Enforcement
 - Through managed pipelines

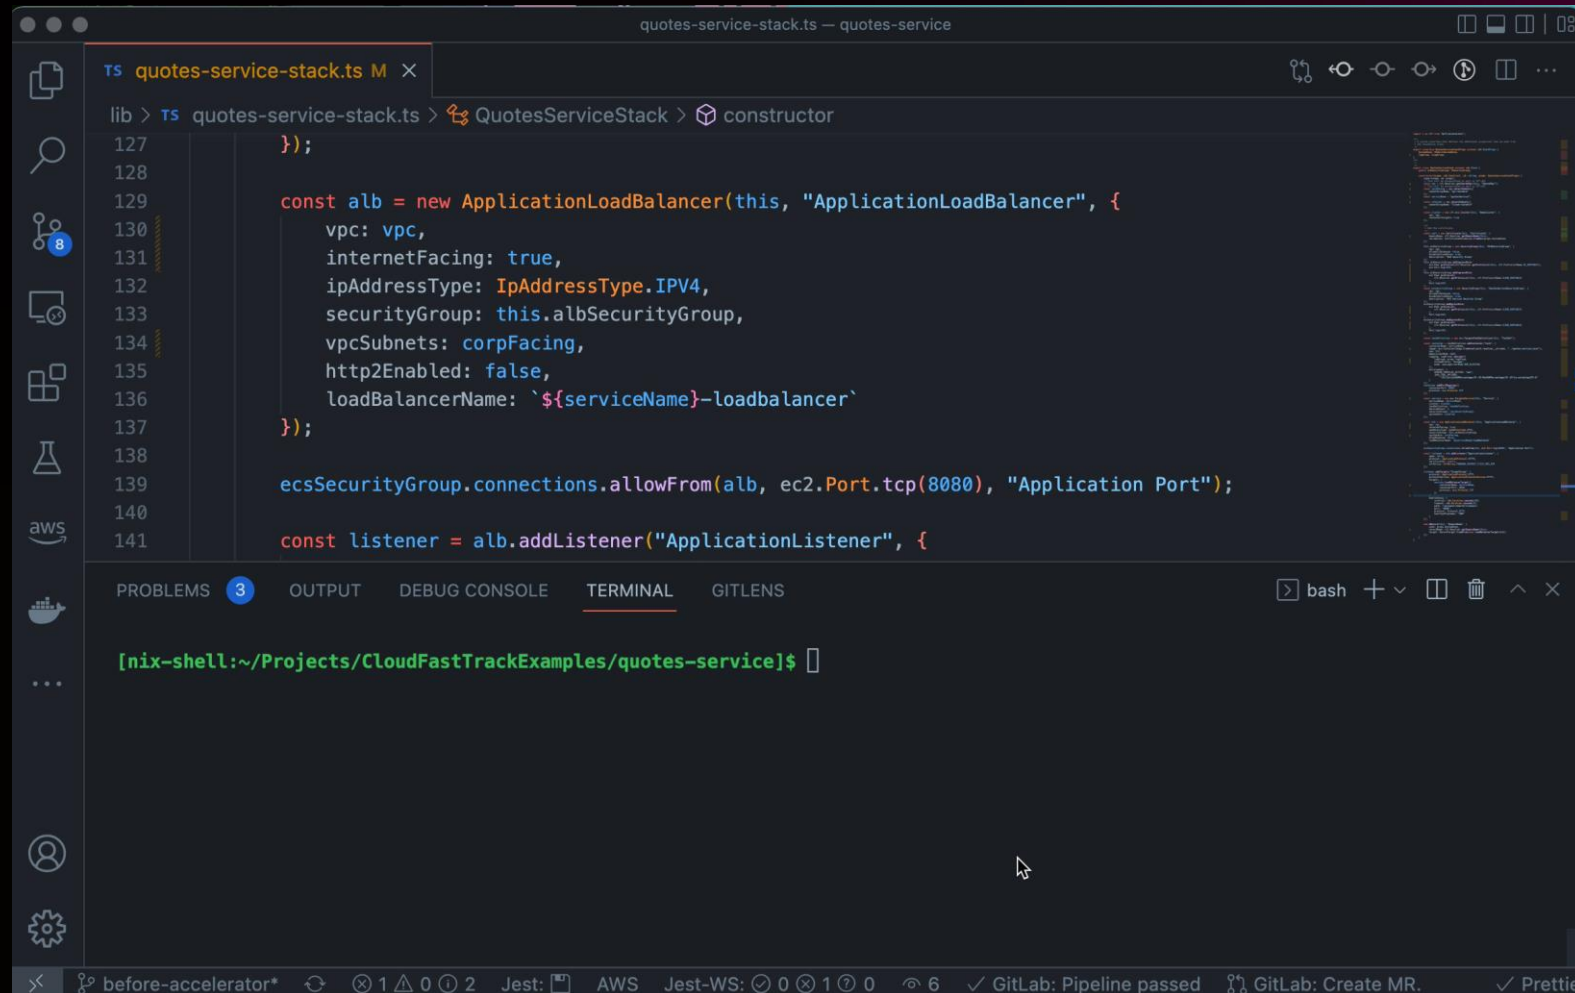
Why guardrails?

ARE YOU OVER-ENGINEERING?

- Highly customizable
 - Security that works for YOUR organization
- Faster iteration cycle
 - Proactive compliance: No need for manual security reviews, commits, deployments
- Faster evaluation
 - 1000+ line CFN template, 200+ Rego policies, <1s evaluation time
- Elaborate feedback
 - No more generic 403 Not Authorized

Guardrails demo

GUARDRAILS, HOW DO THEY WORK?



```
lib > ts quotes-service-stack.ts > QuotesServiceStack > constructor
127     });
128
129     const alb = new ApplicationLoadBalancer(this, "ApplicationLoadBalancer", {
130         vpc: vpc,
131         internetFacing: true,
132         ipAddressType: IpAddressType.IPV4,
133         securityGroup: this.albSecurityGroup,
134         vpcSubnets: corpFacing,
135         http2Enabled: false,
136         loadBalancerName: `${serviceName}-loadbalancer`
137     });
138
139     ecsSecurityGroup.connections.allowFrom(alb, ec2.Port.tcp(8080), "Application Port");
140
141     const listener = alb.addListener("ApplicationListener", {
```

PROBLEMS 3 OUTPUT DEBUG CONSOLE TERMINAL GITLENS

[nix-shell:~/Projects/CloudFastTrackExamples/quotes-service]\$


before-accelerator* 1 0 2 Jest: AWS Jest-WS: 0 1 0 6 ✓ GitLab: Pipeline passed GitLab: Create MR. ✓ Prettie



Recap: Identity perimeter

ONLY TRUSTED IDENTITIES CAN ACCESS MY RESOURCES

ONLY TRUSTED IDENTITIES ARE ALLOWED FROM MY NETWORK



Identity

My corporate credentials

`aws:PrincipalOrgID/aws:PrincipalOrgPaths`

AWS service principal

`aws:PrincipalIsAWSService`

Fast Track: Identity perimeter (1/2)

ONLY TRUSTED IDENTITIES CAN ACCESS MY RESOURCES

ONLY TRUSTED IDENTITIES ARE ALLOWED FROM MY NETWORK



```
cfn_resource_type = "AWS::SecretsManager::ResourcePolicy"
```

For this resource

```
explicitly_allowed_resources[resource] {  
  some resource  
  all_resources[resource]  
  resource_policy := all_resources[resource].ResourcePolicy  
  expected_denies := [  
    s | s := required_deny_statements[_];  
    iam.get_policy_deny_statements(resource_policy, s)  
  ]  
  count(expected_denies) == count(required_deny_statements)  
}
```

Fetch resource policy

Filter all deny policies

Assert expected denies

Fast Track: Identity perimeter (2/2)



ONLY TRUSTED IDENTITIES CAN ACCESS MY RESOURCES

ONLY TRUSTED IDENTITIES ARE ALLOWED FROM MY NETWORK

```
cfn_resource_type = "AWS::SecretsManager::ResourcePolicy"
```

```
required_deny_statements := [
```

```
{
```

```
  "Action": "secretsmanager:*",
```

```
  "Condition": {
```

```
    "BoolIfExists": {
```

```
      "aws:PrincipalIsAWSService": "false"
```

```
    },
```

```
    "ForAllValues:StringNotLikelfExists": {
```

```
      "aws:PrincipalOrgPaths": [
```

```
        path | path := cft_managed_parameters.VENDED_PRINCIPAL_ORG_PATHS[_]
```

```
      ]},
```

```
    },
```

```
    "Effect": "Deny",
```

```
    "Principal": {"AWS": "*"},
```

```
    "Resource": "*",
```

```
  },
```

```
]
```

For this resource

AWS service principal

Whitelisted org paths

Recap: Resource perimeter

MY IDENTITIES CAN ACCESS ONLY TRUSTED RESOURCES

ONLY TRUSTED RESOURCES CAN BE ACCESSED FROM MY NETWORK



Resource

My corporate resources:

aws:ResourceOrgID

Custom perimeter: Lambda bootstrap

Use whitelisted bucket

Fast Track: Resource perimeter (1/2)



MY IDENTITIES CAN ACCESS ONLY TRUSTED RESOURCES

ONLY TRUSTED RESOURCES CAN BE ACCESSED FROM MY NETWORK

```
endpointPolicy(): PolicyStatement[] {  
  const endpointStatements = new Array<PolicyStatement>();  
  endpointStatements.push(  
    new PolicyStatement({  
      effect: Effect.DENY,  
      actions: ["*"],  
      resources: ["*"],  
      principals: [new StarPrincipal()],  
      conditions: {  
        StringNotEquals: {  
          "aws:ResourceOrgID": [ this.props.parentStackProps.orgId, SharedNames.FASTTRACK_ORG_ID ]  
        },  
        BoolIfExists: {  
          "aws:PrincipalIsAWSService": "false"  
        }  
      }  
    })  
  })  
}
```

For whitelisted orgs

AWS principal

Fast Track: Resource perimeter (2/2)

MY IDENTITIES CAN ACCESS ONLY TRUSTED RESOURCES

ONLY TRUSTED RESOURCES CAN BE ACCESSED FROM MY NETWORK



```
resource_type = "AWS::Lambda::Function"
```

For this resource

```
bootstrap_s3_code_bucket_noncompliant_resources(resource_type) = noncompliant_resources {  
  explicitly_allowed_bucket_reference = {resource |  
    some resource  
    all_resources[resource]  
    input.Resources[resource].Properties.Code.S3Bucket == cft.VENDED_CODE_BUCKET_NAME  
  }  
  
  noncompliant_resources := {r | all_resources[r]} - explicitly_allowed_bucket_reference  
  ...  
}
```

Use vended bucket

Non-compliant resource



Network perimeter

MY RESOURCES CAN ONLY BE ACCESSED FROM EXPECTED NETWORKS

MY IDENTITIES CAN ACCESS RESOURCES ONLY FROM EXPECTED NETWORKS



Network

My corporate data center/My corporate AWS network

`aws:SourceIp/aws:SourceVpc`

AWS service using caller identity

`aws:viaAWSService`

AWS service principal

`aws:PrincipalIsAWSService`

SLR

`arn:aws:iam::012345678901:role/aws-service-role/*`

Custom: `Vended VPCs, Vended PLs`



Fast Track: Network perimeter (1/2)

MY RESOURCES CAN ONLY BE ACCESSED FROM EXPECTED NETWORKS

MY IDENTITIES CAN ACCESS RESOURCES ONLY FROM EXPECTED NETWORKS



```
resource = "AWS::ElasticLoadBalancingV2::TargetGroup"
```

For TargetGroup resource

```
explicitly_allowed_resources[resource] {  
    some resource  
    resource_definition := all_resources[resource]  
    resource_definition.TargetType == "ip"  
    cft_managed_parameters.VENDED_VPC_IDS[resource_definition.VpcId]  
}
```

If TargetType is set to IP

```
noncompliant_resources = {r | all_resources[r]} - explicitly_allowed_resources)
```

Assert vended-VPC usage

Fast Track: Network perimeter (2/2)

MY RESOURCES CAN ONLY BE ACCESSED FROM EXPECTED NETWORKS

MY IDENTITIES CAN ACCESS RESOURCES ONLY FROM EXPECTED NETWORKS



```
security_group_resource = "AWS::EC2::SecurityGroup"
```

```
vended_prefix_list_id[resource] {
```

```
...
```

```
some resource, egress_rules
```

```
input.Resources[resource].Type == security_group_resource
```

```
cfthelpers.list_contains(
```

```
    input.Resources[resource].Properties.SecurityGroupEgress[egress_rules].DestinationPrefixListId,
```

```
    cft_managed_parameters.VENDED_PREFIX_LIST_IDS
```

```
)
```

```
}
```

```
...
```

```
...
```

```
compliant_security_groups := same_stack_reference | vended_sg_id | vended_prefix_list_id
```

For SecurityGroup resource

Assert vended PLs are used

Impact statement

WHAT WE'VE ACCOMPLISHED

Policy as code for data perimeter:

- Infrastructure security reviews – No more
- App deployment – Self-service
- Account provisioning from weeks to minutes
- Continuous evaluation – drift detection and continuous evaluation

Learnings

WHAT WE COULD DO BETTER

- Day 1 culture
 - A journey, not a destination – Raise the floor every day
- See the forest for the trees
 - Enable use cases, not services
- Platform operationalization
 - Telemetry, monitoring, alarming

Available resources



Thank you!

Tatyana Yatskevich

[linkedin.com/in/tatyana-yatskevich](https://www.linkedin.com/in/tatyana-yatskevich)

Shubham Shukla

[linkedin.com/in/shu-shukla](https://www.linkedin.com/in/shu-shukla)



Please complete the session survey in the **mobile app**



© 2022, Amazon Web Services, Inc. or its affiliates. All rights reserved.