






2016



Contents

Announced in January			
 Certificate Manager			
Announced in February			
 GameLift			
Announced in March			
 Amazon Lumberyard			
Announced in April			
Announced in May			

 Application Discovery Service			
Announced in June			
Announced in July			
Announced in July			
Announced in September			
			

<u>Deep Learning AMIs</u>			
Announced in October			
 <u>Server Migration Service</u>			
Announced in November			
 <u>AWS Organizations</u>	 <u>Amazon Lex</u>	 <u>Greengrass</u>	 <u>Athena</u>
 <u>Amazon Lightsail</u>	 <u>AWS IoT Button</u>	 <u>Amazon Polly</u>	<u>Amazon Rekognition</u>
<u>Serverless Application Model</u>			
Announced on 1st December			
 <u>Step Functions</u>	 <u>CodeBuild</u>	 <u>AppStream 2.0</u>	 <u>Amazon Pinpoint</u>

 Personal Health Dashboard	 Systems Manager	 Shield	 Batch
 X-ray	 Glue		
Announced on 7th December			
 Artifact			
Announced on 12th December			
 AWS Managed Services			



AWS Certificate Manager

Certificate Manager

15. QUESTION

A company uses an Amazon RDS MySQL database instance to store customer order data. The security team have requested that SSL/TLS encryption in transit must be used for encrypting connections to the database from application servers. The data in the database is currently encrypted at rest using an AWS KMS key.

How can a Solutions Architect enable encryption in transit?

- ☐ Add a self-signed certificate to the RDS DB instance. Use the certificates in all connections to the RDS DB instance.
- ☐ Enable encryption in transit using the RDS Management console and obtain a key using AWS KMS.
- ☐ Download the AWS-provided root certificates. Use the certificates when connecting to the RDS DB instance.
- ☐ Take a snapshot of the RDS instance. Restore the snapshot to a new instance with encryption in transit enabled.

Explanation:

Amazon RDS creates an SSL certificate and installs the certificate on the DB instance when Amazon RDS provisions the instance. These certificates are signed by a certificate authority. The SSL certificate includes the DB instance endpoint as the Common Name (CN) for the SSL certificate to guard against spoofing attacks.

You can download a root certificate from AWS that works for all Regions or you can download Region-specific intermediate certificates.

CORRECT: "Download the AWS-provided root certificates. Use the certificates when connecting to the RDS DB instance" is the correct answer.

INCORRECT: "Take a snapshot of the RDS instance. Restore the snapshot to a new instance with encryption in transit enabled" is incorrect. There is no need to do this as a certificate is created when the DB instance is launched.

INCORRECT: "Enable encryption in transit using the RDS Management console and obtain a key using AWS KMS" is incorrect. You cannot enable/disable encryption in transit using the RDS management console or use a KMS key.

INCORRECT: "Add a self-signed certificate to the RDS DB instance. Use the certificates in all connections to the RDS DB instance" is incorrect. You cannot use self-signed certificates with RDS.

AWS Certificate Manager introduces Enterprise Controls to help govern certificate issuance

Posted On: Aug 23, 2023

Enterprise, network and security admins can now use AWS Identity and Access Management (IAM) [condition context keys](#) with [AWS Certificate Manager](#) (ACM) to help ensure that users are issuing certificates that conform to their organization's public key infrastructure (PKI) guidelines. For example, you can use condition keys to allow only DNS validation. Or, you can authorize which of your users can request certificates for specific domain names such as accounting.example.com and/or wildcard names.

Using these new context keys, you can define how your ACM users customize certificate issuance parameters to authorize 1) a specific certificate validation method, 2) who can request certificates for specific domain names including wildcard names, 3) specific certificate key-algorithm(s), and 4) the request of public or private certificate type. Additionally, you can prevent users from disabling Certificate Transparency (CT) logging or requesting certificates from specific AWS Private Certificate Authorities.

You can distribute and enforce your condition keys across your users and accounts using either [IAM](#) or [Service control policies](#) (SCPs) from [AWS Organizations](#). You can enforce organization-wide policies or have specific policies for organization units. For example, you can authorize your HR unit to issue certificates for the domain name HR.example.com while your IT department can only issue certificates for IT.example.com. You can also enforce these policies at account creation through [AWS CloudFormation](#).

Learn more about this feature [here](#) and [get started with ACM](#). This feature is available in all [AWS Regions](#) where ACM is available, including the AWS GovCloud (US) Regions.

Bibliography

- I. Official
- II. Unofficial
- III. Critical
- IV. General

I. Official

[AWS 2023]

AWS Certificate Manager introduces Enterprise Controls to help govern certificate issuance. Available at: <https://aws.amazon.com/about-aws/whats-new/2023/08/aws-certificate-manager-enterprise-controls-certificate-issuance/?ck_subscriber_id=1560524742>

II. Unofficial


III. Critical

IV. General

GameLift

Amazon Lumberyard

Application Discovery Service



Jeff Barr tells us that:

The new AWS [Application Discovery Service](#) (first announced at the AWS Summit in Chicago) is designed to help you to dig in to your existing environments, identify what's going on, and provide you with the information and visibility that you need to have in order to successfully migrate existing applications to the cloud.

[Barr 2016]

This tells us a few interesting things. First, this product was first announced at the Summit in Chicago. I think I know which presentation it was announced at. Barr tells us that this product provides you with information and visibility – we can “dig into” our existing environments. We are then told that:

This service is an important part of the [AWS Cloud Adoption Framework](#). The framework helps our customers to plan for their journey.

This is helpful, telling us that this Application Discovery Service is in fact considered one component of the CAF (announced in 2015). Barr then gets into the details:

The Discovery Agent

To get started, you simply install the small, lightweight agent on your source hosts. The agent unobtrusively collects the following system information:

- Installed applications and packages.
- Running applications and processes.
- TCP v4 and v6 connections.
- Kernel brand and version.
- Kernel configuration.
- Kernel modules.
- CPU and memory usage.
- Process creation and termination events.
- Disk and network events.
- TCP and UDP listening ports and the associated processes.
- NIC information.
- Use of DNS, DHCP, and Active Directory.

We are told about the types of servers that the agent can be run on:

The agent can be run on Ubuntu 14, Red Hat 6-7, CentOS 6-7, and Windows (Server 2008 R2, Server 2012, Server 2012 R2). We plan to add additional options over time so be sure to let us know what you need.



Figure 1 Nirav Kothari in Chicago, describing AWS Application Discovery Service

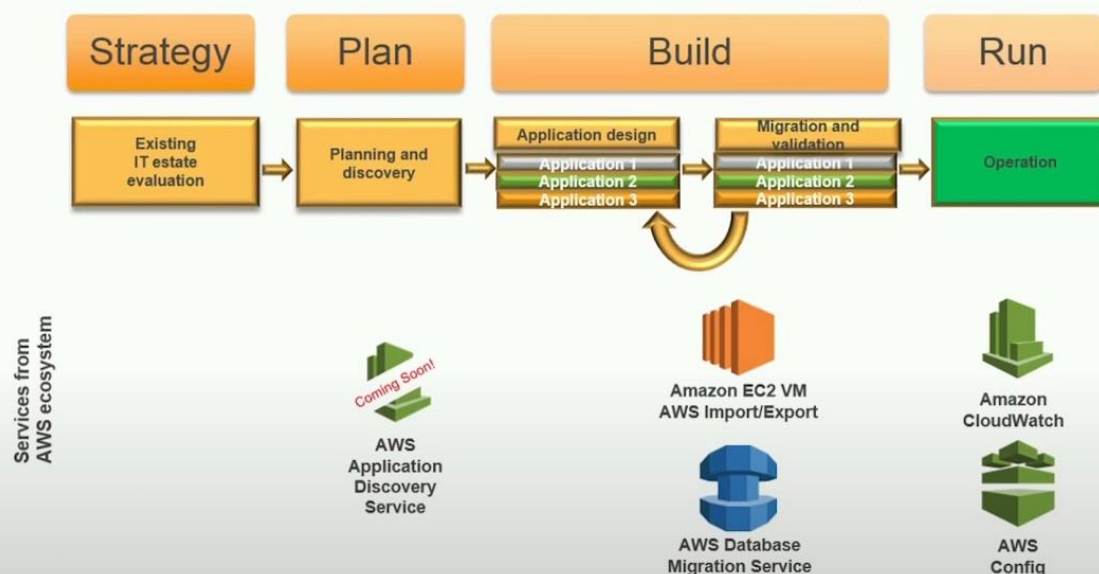
AWS Application Discovery Service

Nirav Kothari

Principal Consultant

AWS Professional Services

Innovations to accelerate cloud migration




AWS Application Discovery Service streamlines planning for your cloud migration




AWS Application Discovery Service **Automate data center application discovery**



Identify application
Inventory



Map application
dependencies




Baseline system and
process performance

AWS Application Discovery Service

Overview

- Agents deployed on source hosts
 - Windows & Linux support
- Capture system inventory, performance, and dependencies
- Capture and store secured data to AWS
- API access to discovered assets
- Output to CSV or XML
 - Can be imported into a third-party migration or visualization tool



AWS Application Discovery Service

Discovery Agent

Binaries deployed using customer's existing software management tools

Support for Windows and Linux hosts

Very small footprint and minimal overhead

Offline mode to inspect captured data

Proxy agnostic, needs outbound 443 connection only

AWS Application Discovery Service

Data capture capabilities

Infrastructure and application level information captured from hosts

Inventory of installed software applications

Information about running applications and processes

- User information, group information, list of kernel modules, operation systems information

System and process performance, and resource utilization

- CPU and memory usage, all create and stop process events, disk and network

Track network activity on a server to infer dependencies

- TCP/UDP listening ports (with processes), TCP v4/v6 connections, NIC information

AWS Application Discovery Service

Features

Public API operations to manage and consume the service

- ListConfigurationItem—search discovered hosts
- GetConfigurationAttributes—describe a discovered host
- ExportConfiguration—ability to export the discovery data as CSV
- Create/Delete Tags—group servers and plan migrations

AWS Command Line Interface (AWS CLI) and SDK utility to access the API operations

Offline processing and visualization

A migration example

Prepare for a LAMP stack for migration

Application Discovery Service views dependencies of a LAMP stack

- Workloads comprised of Apache, MySQL, and PHP servers
- Depends on infrastructure services: LDAP, hardware load balancers, DHCP server, DNS server, Syslog server
- Analyze dependencies and right-size using API operations, import to Microsoft Excel

SI Partner or AWS Professional Services would:

- Create a suitable VPC and subnet on AWS to house application
- Establish a connection by AWS Direct Connect from on-premises to AWS Region
- Plan out the migration
 - Apache and PHP VMs => Amazon EC2
 - Hardware LB => Elastic Load Balancing
 - MySQL => Amazon RDS
 - DNS => Amazon Route 53
 - LDAP and log servers stay on-premises (for other on-premises applications)

Inventory by application type

```
discoverycli> count --hosts --application
Found a total of 6 Web servers(Identified by processes nginx, httpd)
Found a total of 3 DB servers(Identified by processes mysqld)
```

A migration example:

How to use the CLI to prepare for a LAMP stack



```

discoverycli> list --stack --lamp
Client(ClientProcess)
-----
ip-172-31-41-222(python2.7) ,
ip-172-31-37-98(python2.7) ,
ip-172-31-24-0(python2.7) ,

NginxServer(NginxSrvProc) ,
-----
ip-172-31-23-246(nginx / uwsgi) ,
ip-172-31-39-165(nginx / uwsgi) ,
ip-172-31-32-196(nginx / uwsgi) ,

MySQLServer(MySqlSrvProc)
-----
ip-172-31-26-170(mysql) ,
ip-172-31-28-69(mysql) ,
ip-172-31-28-69(mysql)
  
```

Search for dependent hosts

```

discoverycli> list --hosts --dependents ip-172-31-26-170
The following hosts are dependent on host ip-172-31-26-170
  
```

SourceIp	SrcHostName	SourceProcessName	DstIp	DstPort	DstHostName	DstProcessName	TagKey	TagValue
172.31.23.246	ip-172-31-23-246	uwsgi	172.31.26.170	3306	ip-172-31-26-170	mysqld	server_type	wave1

Found a total of 1 connections

```

discoverycli> list --hosts --dependents ip-172-31-23-246
The following hosts are dependent on host ip-172-31-23-246
  
```

SourceIp	SrcHostName	SourceProcessName	DstIp	DstPort	DstHostName	DstProcessName	TagKey	TagValue
172.31.41.222	ip-172-31-41-222	python2.7	172.31.23.246	80	ip-172-31-23-246	nginx	NULL	NULL

Found a total of 1 connections

```

discoverycli> tag_hosts ip-172-31-26-170, ip-172-31-23-246 wave1
  
```

Plan EC2 resource provisioning

```
discoverycli> list --hosts --tag wave1
      agentId      Host name
-----
i-7cd0e1b8      ip-172-31-23-246
i-4ed0e18a      ip-172-31-26-170
```

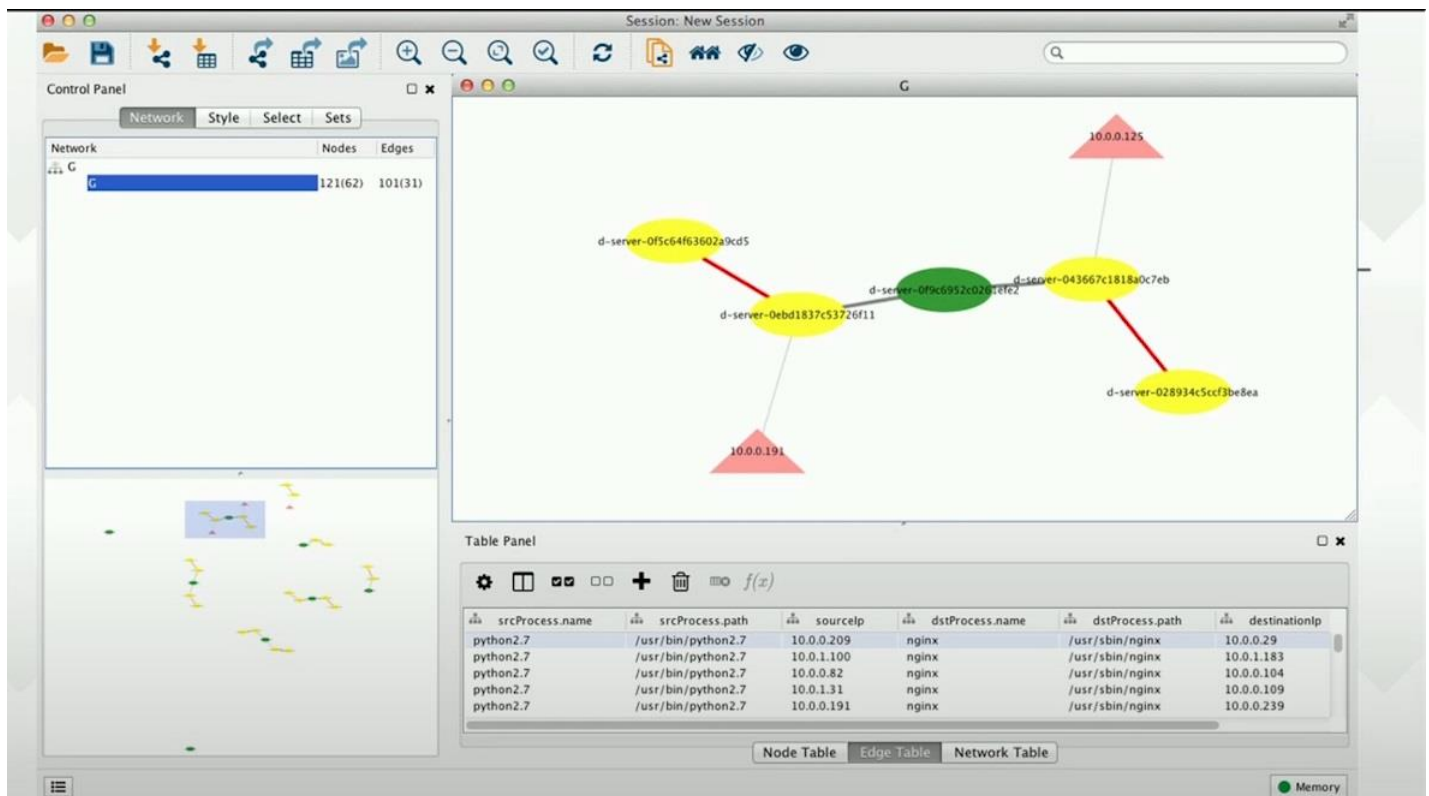
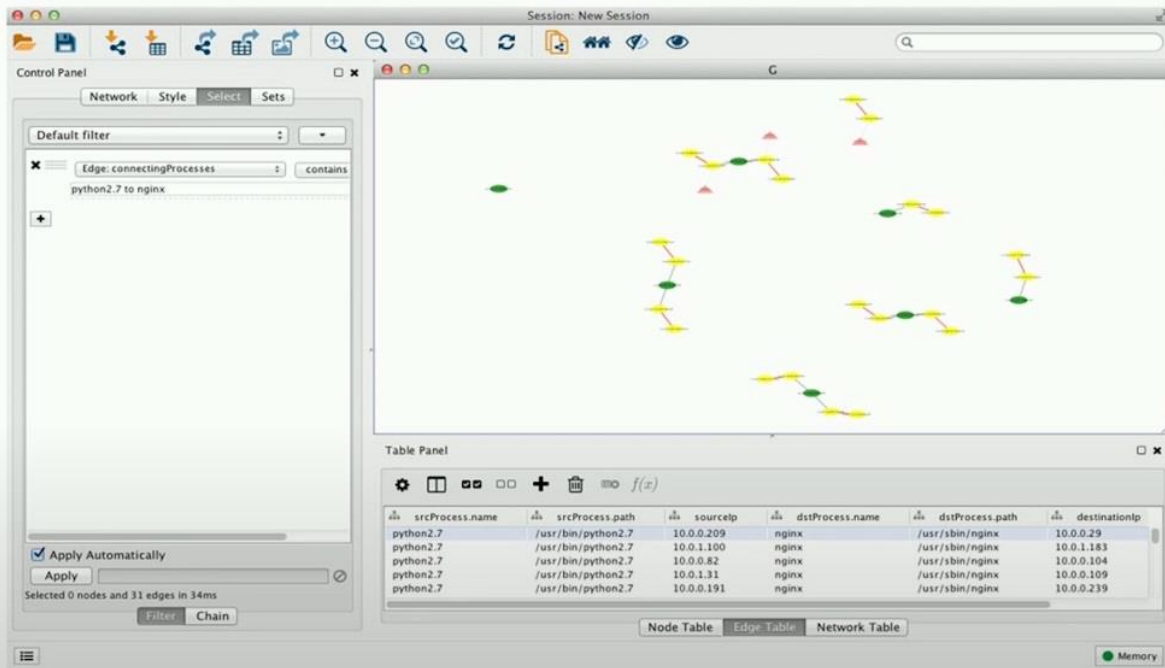
```
discoverycli> describe --hosts --tag wave1
      agentId      Host name      Provisioned CPU      Max CPU      Provisioned memory(in MB)      Max memory(in MB)
-----
i-7cd0e1b8      ip-172-31-23-246      1      90.0      592.30078125      144.03125
i-4ed0e18a      ip-172-31-26-170      1      81.0      592.30078125      249.3828125
```

Availability

Available through AWS Professional Services consultants or SI partners

- Visit service detail page on the AWS marketing website
- Fill the form to sign up for beta program
- We will get you in touch with AWS Professional Services consultant or SI partner

Visualization



AWS Application Discovery Service

Integration with cloud migration tools

Provides a set of public API operations and open data formats, making it easy to integrate or coexist with other tools

Integrates with:

- Discovery tools that can publish information
- Migration tools that leverage the discovered assets
- Existing migration frameworks



A stage presentation slide for AWS re:Invent. On the left, a man in a light blue shirt and dark trousers stands on a stage with blue curtains. To his left is a screen displaying the AWS re:Invent logo. The slide background is dark with a vibrant pink and purple abstract graphic on the right side. The text on the slide includes the session ID 'ENT331', the title 'Save Money and Migrate Faster with Rapid Discovery and Analysis', and the names and titles of the speakers, David Zipkin and Paul Verney. The AWS re:Invent logo is at the bottom left, and the AWS logo is at the bottom right.

ENT331

Save Money and Migrate Faster with Rapid Discovery and Analysis

David Zipkin
Senior Manager
AWS

Paul Verney
Senior Project Manager
Hannover Re SE

AWS re:Invent

© 2018, Amazon Web Services, Inc. or its affiliates. All rights reserved.

aws



Figure 2 Carmen Puccio gives a presentation in 2019

Carmen Puccio works for AWS. In 2019, at the ReInvent conference, Carmen Puccio gives a presentation which talks about strategies for migrating to the cloud. The presentation features the AWS Application Discovery Service.

20. QUESTION

A company runs hundreds of applications across several data centers and office locations. The applications include Windows and Linux operating systems, physical installations as well as virtualized servers, and MySQL and Oracle databases. There is no central configuration management database (CMDB) and existing documentation is incomplete and outdated. A Solutions Architect needs to understand the current environment and estimate the cloud resource costs after the migration.

Which tools or services should the Solutions Architect use to plan the cloud migration (Select THREE.)

- ☐ AWS Migration Hub
- ☐ AWS Cloud Adoption Readiness Tool (CART)
- ☐ AWS Application Discovery Service
- ☐ AWS Server Migration Service
- ☐ AWS Config
- ☐ AWS CloudWatch Logs

Bibliography

- I. Official
- II. Unofficial
- III. Critical
- IV. General

- I. Official

[Barr 2016]

Barr, Jeff (2016). New – AWS Application Discovery Service – Plan Your Cloud Migration. May 12th 2016. Available at:
<<https://aws.amazon.com/blogs/aws/new-aws-application-discovery-service-plan-your-cloud-migration/>>

- II. Unofficial

III. Critical

IV. General

[Puccio 2019]


Strategies and tools to migrate

[Kothari 2016]


Kothari, Nirav (2016). Chicago – AWS Application Discovery Service. Available at:

<https://www.youtube.com/watch?v=Q6EJE3ndOqI&t=1584s&ab_channel=AmazonWebServices>

Deep Learning AMIs



AWS Server Migration Service



Jim Huang explaining the Server Migration Service

AWS Organisations



Announcing AWS Organizations, Now in Preview

Posted On: Nov 29, 2016

AWS Organizations makes it easier for IT teams to manage multiple AWS accounts. You can use Organizations to create groups of AWS accounts, and then apply policies to these groups to help centrally manage security and automation settings.

Organizations helps improve security by helping to control which AWS services can be accessed by users in individual AWS accounts. If Organizations specifies tighter controls than the [AWS Identity and Access Management \(IAM\)](#) policies for the account, then users will be held to the more restrictive Organizations policy. Organizations enables you to automate new account creation by using APIs to create and add new accounts to a group. Organizations also includes the features of [Consolidated Billing](#), automatically linking accounts in your organization so they are paid for using a single payment method.

With Organizations, you can more easily manage your AWS accounts at scale by application, environment, team, or any other grouping that makes sense for your business.

To learn more visit the [AWS Organizations home page](#) and sign up for the AWS Organizations Preview today.

5. QUESTION

An AWS Organization has an OU with multiple member accounts in it. The company needs to restrict the ability to launch only specific Amazon EC2 instance types. How can this policy be applied across the accounts with the least effort?

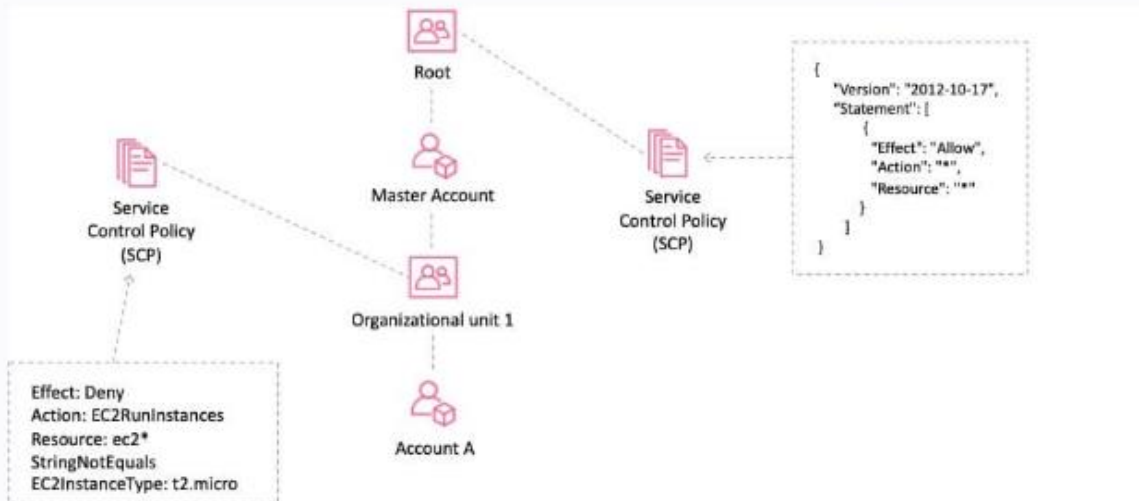
- ☐ Create an SCP with an allow rule that allows launching the specific instance types
- ☒ Create an SCP with a deny rule that denies all but the specific instance types
- ☐ Create an IAM policy to deny launching all but the specific instance types
- ☐ Use AWS Resource Access Manager to control which launch types can be used

Correct

Explanation:

To apply the restrictions across multiple member accounts you must use a Service Control Policy (SCP) in the AWS Organization. The way you would do this is to create a deny rule that applies to anything that does not equal the specific instance type you want to allow.

The following architecture could be used to achieve this goal:



CORRECT: "Create an SCP with a deny rule that denies all but the specific instance types" is the correct answer.

INCORRECT: "Create an SCP with an allow rule that allows launching the specific instance types" is incorrect as a deny rule is required.

INCORRECT: "Create an IAM policy to deny launching all but the specific instance types" is incorrect. With IAM you need to apply the policy within each account rather than centrally so this would require much more effort.

INCORRECT: "Use AWS Resource Access Manager to control which launch types can be used" is incorrect. AWS Resource Access Manager (RAM) is a service that enables you to easily and securely share AWS resources with any AWS account or within your AWS Organization. It is not used for restricting access or permissions.

References:

https://docs.aws.amazon.com/organizations/latest/userguide/orgs_manage_policies_example-scps.html#example-ec2-instances

Save time with our AWS cheat sheets:

<https://digitalcloud.training/aws-organizations/>

4. QUESTION

A company has multiple accounts that are managed using AWS Organizations. A security engineer must setup a shared S3 bucket in a central account and grant read-only access for all users in any account within the AWS Organization. There should be no public access to the S3 bucket data.

Which parameters should the security engineer use to accomplish this goal MOST efficiently?

- ☐ Specify '*' as the principal and aws:PrincipalOrgId as a condition.
- ☒ Specify aws:PrincipalOrgId as the principal with the organization ID value.
- ☐ Specify all account numbers within an array as the principal.
- ☐ Specify the organization's master account as the principal.

Incorrect

Explanation:

You can use a condition key, `aws:PrincipalOrgID`, in policies to require all principals accessing the resource to be from an account (including the master account) in the organization. To set this up for this scenario you must specify '*' as the principal, to allow any user access, and then restrict only to users within the AWS Organization using the condition key. The `aws:PrincipalOrgId` condition key should be used with the organization ID value specified.

The example policy below could be used for this scenario (allows `s3:GetObject` only):

```
{
  "Version": "2012-10-17",
  "Statement": [
    {
      "Sid": "AllowGetObject",
      "Effect": "Allow",
      "Principal": "*",
      "Action": "s3:GetObject",
      "Resource": "arn:aws:s3::my-secured-bucket/*",
      "Condition": {
        "StringEquals": {
          "aws:PrincipalOrgID": [ "o-xxxxxxxxxx" ]
        }
      }
    }
  ]
}
```

CORRECT: "Specify "*" as the principal and `aws:PrincipalOrgId` as a condition" is the correct answer (as explained above.)

INCORRECT: "Specify `aws:PrincipalOrgId` as the principal with the organization ID value" is incorrect.

The value mentioned is used in a condition, not in the principal.

INCORRECT: "Specify all account numbers within an array as the principal" is incorrect.

This is less efficient as you must specify all account numbers and you must come back and add account numbers if new accounts are added to the organization.

INCORRECT: "Specify the organization's master account as the principal" is incorrect.

This is not the correct method and will not grant access for users from other accounts within the organization.

References:

<https://aws.amazon.com/blogs/security/control-access-to-aws-resources-by-using-the-aws-organization-of-iam-principals/>

AWS Security Blog

An easier way to control access to AWS resources by using the AWS organization of IAM principals

by Sulay Shah | on 17 MAY 2018 | in [AWS Identity And Access Management \(IAM\)](#), [AWS Organizations](#), [Security](#), [Identity](#), & [Compliance](#) | [Permalink](#) | [Comments](#) | [Share](#)

AWS Identity and Access Management (IAM) now makes it easier for you to control access to your AWS resources by using the [AWS organization](#) of IAM principals (users and roles). For some services, you grant permissions using resource-based policies to specify the accounts and principals that can access the resource and what actions they can perform on it. Now, you can use a new [condition key](#), `aws:PrincipalOrgID`, in these policies to require all principals accessing the resource to be from an account (including the master account) in the organization. For example, let's say you have an [Amazon S3 bucket policy](#) and you want to restrict access to only principals from AWS accounts inside of your organization. To accomplish this, you can define the `aws:PrincipalOrgID` condition and set the value to your [organization ID](#) in the bucket policy. Your organization ID is what sets the access control on the S3 bucket. Additionally, when you use this condition, policy permissions apply when you add new accounts to this organization without requiring an update to the policy.

In this post, I walk through the details of the new condition and show you how to restrict access to only principals in your organization using S3.

13. QUESTION

A company has created an organization within AWS Organizations. A security engineer created an organizational unit (OU) and moved several AWS accounts into the OU. The Amazon EC2 service is restricted with the following SCP:

```
{
  "Version": "2012-10-17",
  "Statement": [
    {
      "Effect": "Deny",
      "Action": "ec2:*",
      "Resource": "*"
    }
  ]
}
```

One of the AWS accounts in the OU is used for data analytics and the data analysts require access to Amazon EC2 instances for running analytics software.

How can the security engineer provide the data analysts with Amazon EC2 access without affecting the other accounts in the OU?

- ☐ Instruct the data analysts to login to the data analytics account using root credentials to avoid the restrictions.
- ☐ Move the SCP that denies the EC2 service to the root OU of Organizations to limit the accounts it applies to.
- ☐ Create a new OU without the SCP restricting EC2 access. Move the data analytics account to the new OU.
- ☒ Add an allow statement for the EC2 service to the SCP with a condition that limits it to the data analytics account.

2. QUESTION

A company uses AWS Organizations and federation between and on-premises identity provider (IdP). Users authenticate to AWS using credentials in the IdP. A security engineer needs to audit requests to AWS Organizations for creating new AWS accounts.

What should the engineer review to determine who made the request?

- ☐ Federated identity provider logs for the user name.
- ☐ AWS IAM Access Analyzer for the federated user name.
- ☒ AWS CloudTrail for the federated identity user name.
- ☐ AWS X-Ray traces for the federated identity user name.

Correct

Explanation:

AWS Organizations enables you to create new accounts through the console or programmatically using the organizations API. When you create accounts through organizations the API calls are logged in AWS CloudTrail.

In this case AWS CloudTrail can be used to track the activity of the federated users. CloudTrail records the following AWS Security Token Service (AWS STS) API calls: AssumeRoleWithWebIdentity and AssumeRoleWithSAML.

Records of these API calls will be stored in CloudTrail and the user name of the federated user who made the call can be identified.

CORRECT: "AWS CloudTrail for the federated identity user name" is the correct answer (as explained above.)

INCORRECT: "AWS IAM Access Analyzer for the federated user name" is incorrect.

This service helps you identify the resources in your organization and accounts, such as Amazon S3 buckets or IAM roles, that are shared with an external entity.

INCORRECT: "AWS X-Ray traces for the federated identity user name" is incorrect.

This service is used for analyzing and debugging production, distributed applications. It is not used for auditing account activity.

INCORRECT: "Federated identity provider logs for the user name" is incorrect.

The request to AWS Organizations will be recorded by CloudTrail and the federated identity user name will be recorded in the log entry.

How to Easily Identify Your Federated Users by Using AWS CloudTrail

by Akshat Goel | on 28 MAR 2016 | in [AWS CloudTrail](#), [How-To](#), [Identity](#) | [Permalink](#) | [Comments](#) | [Share](#)

Starting today, you can use AWS CloudTrail to track the activity of your federated users ([web identity federation](#) and [Security Assertion Markup Language \[SAML\]](#)). For example, you can now use CloudTrail to identify a SAML federated user who terminated an Amazon EC2 instance in your AWS account, or to identify a mobile application user who has signed in using her Facebook account and has deleted a photo (an Amazon S3 object) from an Amazon S3 bucket. The ability to track federated users can help make it easier for you to conduct audits of their activity, which in turn can help you with your compliance and security efforts.

Akshat Goel's 2016 article

8. QUESTION

A financial services company has an organization in AWS organizations with several member accounts. Amazon S3 buckets are used to store sensitive data backups from common applications within each AWS account. The company needs to restrict users from deleting any S3 buckets or objects across the organization.

What is the MOST scalable solution that meets these requirements?

- ☐ Permissions boundaries in AWS IAM
- ☐ Service Control Policies (SCPs)
- ☐ S3 bucket policies
- ☐ S3 bucket ACLs

Explanation:

The most scalable solution is to use a service control policy as this will automatically apply to any additional accounts that are added to the organization. The following example SCP prevents users or roles in any affected account from deleting any S3 bucket or objects.

```
{
  "Version": "2012-10-17",
  "Statement": [
    {
      "Action": [
        "s3:DeleteBucket",
        "s3:DeleteObject",
        "s3:DeleteObjectVersion"
      ],
      "Resource": "*",
      "Effect": "Deny"
    }
  ]
}
```

CORRECT: "Service Control Policies (SCPs)" is the correct answer (as explained above.)

INCORRECT: "Permissions boundaries in AWS IAM" is incorrect.

Permissions boundaries are implemented within IAM and therefore must be configured in each AWS account within the organization. This is a much less scalable solution.

INCORRECT: "S3 bucket policies" is incorrect.

Bucket policies must be configured on every bucket in every account within the organization. This is also a much less scalable solution.

INCORRECT: "S3 bucket ACLs" is incorrect.

Bucket ACLs offer only limited permission options and would also not represent a scalable solution.

References:

https://docs.aws.amazon.com/organizations/latest/userguide/orgs_manage_policies_scps.html


ggg

`https://betterprogramming.pub/how-to-migrate-your-aws-account-into-an-organization-b438d6301dee`


Image credits

The building is Willis Tower, in Chicago. From <https://theconstructor.org/architecture/tube-structural-system-types-features/560166/>


Amazon Lex



Greengrass



Athena




It is very common for people to get confused between Athena and Aurora. Very commonly, when speaking, people say one when they mean the other. Recall that Aurora was the relational database engine released by AWS way back in 2014. We're about to look at Amazon Athena, announced in 2016.

Aurora was a Roman goddess, used to personify dawn, and the coming of the **light**. The “northern lights”, with the blue and green colours in the sky, are named after the goddess. AWS created Aurora, a database that really **roared**, in terms of performance. In their 2017 paper, Alexandre Verbitski and Anurag Gupta—along with eight others—set out how they essentially made the database engine **lighter**. Think about two layers: database and storage. They talk about various functions being “offloaded” to the storage layer. Database engines involve something called a redo log. At one point they even exclaim: the database **IS** the log. They rethink failure. The cloud is raucous with failures, the Aurora authors note (background failures happen all the time). So we should demand 4 of 6 quorum for writes and 3 of 6 quorum for reads—they argue. Aurora is a light engine, that roars. The Roman goddess, by the way, was not particularly intelligent. Her lover was a mortal man, whom she granted eternal life. However, she forget to ask for eternal youth. The man eventually became so decrepit that she

decided to turn him into a cicada. If you want someone intelligent, look to Athena.

Athena is the Greek goddess of wisdom and craft. In 2016, AWS *thankfully* launched Athena. This is for people whose S3 buckets have become impenetrable **th**ickets. The sides of the bucket are **th**warted with the thick accretion of objects within it. Having **th**robbing storage is not everything. Seriously, being think of inaccessible things as a *th*reat. I say “thumbs up AWS”, for introducing ATHENA.



13. QUESTION

A company manages an application that runs on Amazon EC2 instances behind a Network Load Balancer (NLB). The NLB has access logs enabled which are being stored in an Amazon S3 bucket. A security engineer requires a solution to run ad hoc queries against the access logs to identify application access patterns.

How should the security engineer accomplish this task with the least amount of administrative overhead?

- ☐ Create an Amazon Athena table that uses the S3 bucket containing the access logs. Run SQL queries using Athena.
- ☐ Use the S3 copy command to copy logs to a separate bucket. Enable S3 analytics to analyze access patterns.
- ☐ Write an AWS Lambda function to query the access logs. Use event notifications to trigger the Lambda functions when log entries are added.
- ☐ Import the access logs into Amazon CloudWatch Logs. Use CloudWatch Logs Insights to analyze the log data.

Incorrect

Explanation:

Amazon Athena is a serverless service you can use to run SQL queries against data in Amazon S3. You just need to point Athena to your data in Amazon S3, define the schema, and start querying using the built-in query editor. This is ideal for running ad-hoc queries on access logs stored in an S3 bucket.

CORRECT: "Create an Amazon Athena table that uses the S3 bucket containing the access logs. Run SQL queries using Athena" is the correct answer (as explained above.)

INCORRECT: "Use the S3 copy command to copy logs to a separate bucket. Enable S3 analytics to analyze access patterns" is incorrect.

There's no need to copy the data and S3 analytics is used to identify object access patterns for requests to S3 objects. It is used for storage class analytics. It does not help with identifying access patterns for your application by reading the file and looking at source IP addresses (for example).

INCORRECT: "Write an AWS Lambda function to query the access logs. Use event notifications to trigger the Lambda functions when log entries are added" is incorrect.

This will be more complex and is less useful for running ad hoc queries as it is something that will run every time a file is added.

INCORRECT: "Import the access logs into Amazon CloudWatch Logs. Use CloudWatch Logs Insights to analyze the log data" is incorrect.

You cannot natively import logs into CloudWatch Logs from Amazon S3. You may be able to achieve this with a custom Lambda function, but it will be more work.

References:

<https://docs.aws.amazon.com/elasticloadbalancing/latest/application/load-balancer-access-logs.html>

<https://aws.amazon.com/athena/features/>

Bibliography

- I. Official
- II. Unofficial
- III. Critical
- IV. General

I. Official

II. Unofficial



Amazon Fills Big Data Hole. Available at:

<<https://www.datanami.com/2017/10/17/amazon-fills-big-data-hole-athena/>>

III. Critical


IV. General

Lightsail



AWS IoT Button

Amazon Polly



Amazon Rekognition

Bibliography

- I. Official
- II. Unofficial
- III. Critical
- IV. General

I. Official

II. Unofficial


III. Critical

[Burgess 2024]

Burgess, Matt (2024). Amazon-Powered AI Cameras Used to Detect Emotions of Unwitting UK Train Passengers. *Wired*. Accessed via *Last Week in AWS*. Available at: <<https://www.wired.com/story/amazon-ai-cameras-emotions-uk-train-passengers/>>

IV. General

Serverless Application Model (SAM)



PROJECT

flourish

Introducing the AWS Serverless Application Model

Posted On: Nov 18, 2016

The AWS Serverless Application Model (AWS SAM, previously known as Project Flourish) extends AWS CloudFormation to provide a simplified way of defining the Amazon API Gateway APIs, AWS Lambda functions, and Amazon DynamoDB tables needed by your serverless application.

Previously, CloudFormation lacked the specialized resource types optimized for defining serverless applications. Now, you can use new resource types to write CloudFormation templates that define serverless resources with only a few lines of text. Additionally, two new commands have been added to the AWS CloudFormation CLI to simplify the process of packaging a serverless application and deploying it with CloudFormation.

As part of this release, the AWS Serverless Application Model is made available under the Apache 2.0 license, enabling others in the ecosystem to adopt and incorporate it into build, deployment, monitoring and management tools with a commercial-friendly license.

To learn more about using AWS SAM to build serverless applications, visit our [documentation](#) and check out the model on [GitHub](#).

SPEED UP NEW SERVERLESS APPLICATION DEVELOPMENT WITH CUSTOMIZED SAM TEMPLATES

JULY 1, 2023
3-MINUTE READ

Learn how you can use custom templates with SAM to speed up initial application setup.



As you start setting patterns and best practices within your projects at some point you will want to share these to simplify the lives of other developers.

SAM allows you to initialize projects by using the `init` command and selecting a template for your project. SAM has built-in templates to generate simple Hello-World starter projects (you can find the list of templates [here](#)). These will help you get

AWS SAM CLI announces remote invoke feature for AWS Lambda

Posted On: Jun 26, 2023

The [AWS Serverless Application Model \(SAM\)](#) Command Line Interface (CLI) announces the launch of remote invoke command which enables developers to quickly invoke their AWS Lambda functions deployed to the AWS cloud. The AWS SAM CLI is a developer tool that makes it easier to build, test, package, and deploy serverless applications.

Customers can now use the SAM CLI to test a Lambda function with a simple command `sam remote invoke`. This command takes an event payload and Lambda logical id as input to invoke the Lambda function, then outputs the response. It supports multiple invocation modes such as response streaming, asynchronous invoke, dry run and request-response. Previously, customers had to use AWS Console or AWS CLI to test their Lambda functions. With this launch, SAM CLI users can use this feature along with `sam sync` command to speed up the build-deploy-test iteration loop.

This feature is available with SAM CLI version 1.88.0+. To learn more about this feature, please see the [blogpost](#) and [documentation](#). You can install the latest version of the SAM CLI by following the instructions in the [documentation](#).

Bibliography

- I. Official
- II. Unofficial
- III. Critical
- IV. General

- I. Official
- II. Unofficial

[Moreno 2023]


Moreno, Andres (2023). Speed up new serverless application development with customized SAM templates. July 1st 2023. Available at: <https://www.andmore.dev/blog/how-to-build-a-custom-sam-template/?utm_source=substack&utm_medium=email>

[Pyle 2023]

https://www.binaryheap.com/building-serverless-applications-with-aws-data/?utm_source=substack&utm_medium=email

- III. Critical
- IV. General

Step Functions



“Asynchronously
invoke”

What does that even *mean*?

18. QUESTION

An application running on Amazon EC2 needs to asynchronously invoke an AWS Lambda function to perform data processing. The services should be decoupled.

Which service can be used to decouple the compute services?

- ☐ AWS Config
- ☐ Amazon SNS
- ☐ Amazon MQ
- ☐ Amazon Step Functions

Correct

Explanation:

You can use a Lambda function to process Amazon Simple Notification Service notifications. Amazon SNS supports Lambda functions as a target for messages sent to a topic. This solution decouples the Amazon EC2 application from Lambda and ensures the Lambda function is invoked.

CORRECT: "Amazon SNS" is the correct answer.

INCORRECT: "AWS Config" is incorrect. AWS Config is a service that is used for continuous compliance, not application decoupling.

INCORRECT: "Amazon MQ" is incorrect. Amazon MQ is similar to SQS but is used for existing applications that are being migrated into AWS. SQS should be used for new applications being created in the cloud.

INCORRECT: "AWS Step Functions" is incorrect. AWS Step Functions is a workflow service. It is not the best solution for this scenario.

States-language.net

AWS Step Functions launches Versions and Aliases

Posted On: Jun 22, 2023

AWS Step Functions announces the availability of Versions and Aliases, improving resiliency for deployments of serverless workflows. [AWS Step Functions](#) is a visual workflow service capable of orchestrating over 11,000+ API actions from over 250 AWS services to automate business processes and data processing workloads.

Now, AWS Step Functions supports more resilient deployments with Versions and Aliases for workflows, a new set of capabilities that makes it easier for you to set up continuous deployment, to help you iterate faster and release safely into production. Using Step Functions Versions and Aliases you can maintain multiple versions of your workflows, track which version was used for each execution, and create aliases that route traffic between workflow versions. You can deploy your workflows gradually using industry standard techniques such as blue-green and canary style deployments with fast rollbacks to your Step Functions workflows, increasing deployment safety and reducing downtime and risk.

There is no additional fee for Versions and Aliases, so you only pay for what you use as per existing AWS Step Functions pricing. Please visit [Step Functions pricing](#) to learn more.

You can get started using Versions and Aliases in the [AWS console](#), AWS CloudFormation, the AWS Command Line Interface (CLI), or the AWS Cloud Development Kit (CDK). To learn more, please see the AWS Step Functions [Developer Guide](#) and the [Launch Blog](#) to get started.

AWS Step Functions Versions and Aliases is available in the [regions listed here](#). For a complete list of regions and service offerings, see [AWS Regions](#).

[AWS Step Functions launches Versions and Aliases](#) - A good enhancement for all of you chumps who somehow don't write perfect code the first time like I do.

Corey Quinn on June 26th 2023

AWS Step Functions adds integration for 7 services including Amazon VPC Lattice

Posted On: Jun 15, 2023

AWS Step Functions expands its AWS SDK integrations with support for 7 additional AWS services including Amazon VPC Lattice, Amazon CloudWatch Internet Monitor, AWS IoT TwinMaker, and Amazon OpenSearch Ingestion.

[AWS Step Functions](#) is a visual [workflow](#) service capable of orchestrating over 12,000+ API actions from over 320 AWS services to help customers build distributed applications at scale. By directly invoking AWS services or their API actions from AWS Step Functions, customers can write less code, simplify their architecture and save costs. In addition to newly added services, Step Functions also added support for over 460 new API actions from new and existing AWS services such as Amazon ECS, Amazon EC2, Amazon Athena, and Amazon Quicksight. For the full list of added services, visit [AWS SDK service integrations](#).

These enhancements are now generally available in all regions where AWS Step Functions is available. Specific services and API actions are subject to the availability of the target services in the AWS Region. For a complete list of regions and service offerings, see [AWS Regions](#). To learn more about the enhancements including AWS SDK integration, read the [Developer Guide](#), and try building a state machine using [our AWS SDK integration tutorial](#).

Bibliography

- I. Official
- II. Unofficial
- III. Critical
- IV. General

Official

[AWS 2023]

AWS Step Functions Launches Versions and Aliases.
[Announcement]. June 22nd 2023. Available at:
<https://aws.amazon.com/about-aws/whats-new/2023/06/aws-step-functions-versions-aliases/?ck_subscriber_id=1560524742>

I.

II. Unofficial

[Pyle 2023]

Pyle, Benjamin (2023). AWS Step Functions Versions and Aliases. June 27th 2023. Available at: <https://www.binaryheap.com/aws-step-function-versions-and-aliases/?utm_source=substack&utm_medium=email>

III. Critical

IV. General

[Stack 2011]

Can an object oriented program be seen as a Finite State Machine? *Software Engineering Stack Exchange* [Forum]. Available at: <<https://softwareengineering.stackexchange.com/questions/95037/can-an-object-oriented-program-be-seen-as-a-finite-state-machine/95039#95039>>

[Cui 2023]


Cui, Yan (2023). Testing Step Functions: how to skip time when testing timeout and Wait states. Available at: <<https://theburningmonk.com/2023/06/testing-step-functions-how-to-skip-time-when-testing-timeout-and-wait-states>>

Image credits

A set of human footprints uncovered by researchers in White Sands National Park. According to a new study, these tracks date to between 21,000 to 23,000 years ago—a time when massive ice sheets are believed to have blocked human migration into the Americas.

PHOTOGRAPH BY DAN ODESS

CodeBuild



What is CodeBuild?

This essay was written in response to a specific problem. I'm learning about AWS. Sooner or later, you are introduced to a variety of AWS services whose names all begin with “code”. There is CodeBuild, CodeDeploy, CodeCommit and CodePipeline. These are thrown onto your lap, and you are left wondering what the difference is between them. We do not get the differences made explicit. Trying to mark out each service vividly is my aim here.

Helpful background

Chris Munns provides a nice introduction to how AWS think about software development. He said:

At AWS, we see that software development typically has four major phases:

- (i) Source
- (ii) Build
- (iii) Test
- (iv) Production

[Munns 2017]

Munns [explains](#) that the SOURCE phase involves writing code. People are “checking in” code, to a source control service.

In the BUILD stage, we’re compiling code. We might be creating container images.

In the TEST stage, we’re integrating the code with other systems. We are performing load testing. We might be performing penetration testing.

At the PRODUCTION stage, we are deploying the software into a production environment. PRODUCTION can mean different things. If you are building a mobile application, PRODUCTION involves pushing the application to the app store; if you are building a web application, PRODUCTION involves pushing the application to a web server.

Sidebar: What *is* meant by “production environment”?

According to [Suse], a production environment is:

The set of computers where finished, user-ready software is deployed and executed.

When software code is moved to the production environment, it is the final step in a 4-tier architecture that includes development, testing, staging and production.

Before it reaches the production environment, the software is still under development, and bugs and flaws are actively found and fixed.

They go on to say that this set of computers is

Often referred to as *live*, especially for servers.

There is some flexibility in where the production environment is housed:

The production environment can be housed in a single data centre, a network of geographically dispersed machines in multiple data centres, or virtual machines (cloud computing).

[Suse]

A production environment is sometimes called a deployment environment. There is currently no Wikipedia article entitled “Production environment”. There is, though, “Deployment environment”. This article is not devoted to production environments. It is

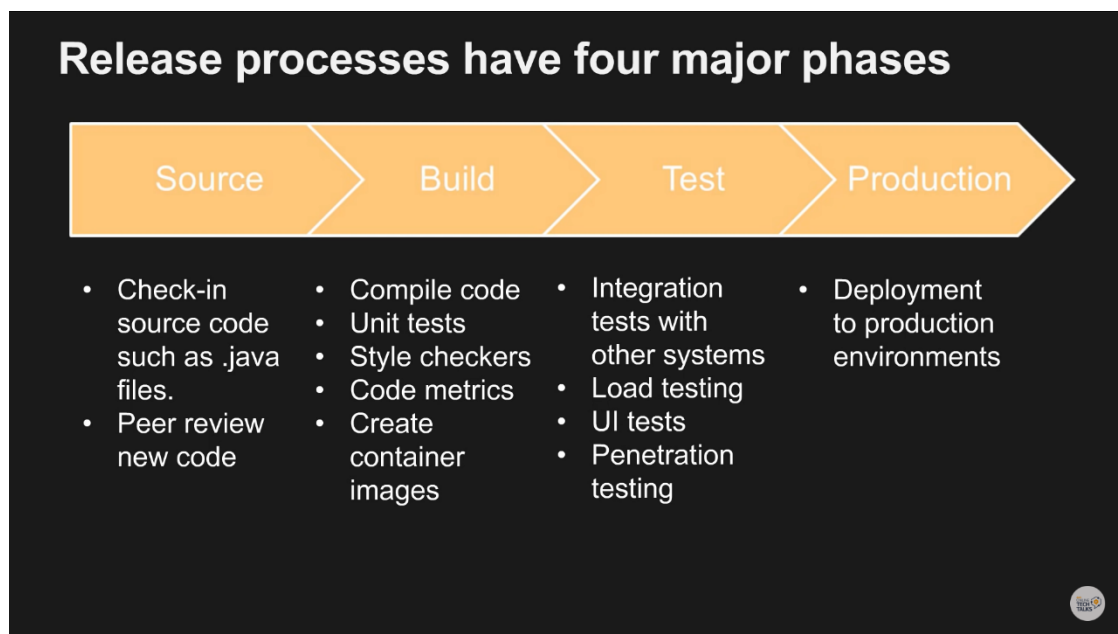
devoted to the general concept of a deployment environment; it has a sub-section on the development environment, a sub-section on the testing environment, and so on.

This is confusing. It seems that the expression “**deployment environment**” can sometimes refer to

- the general concept
possibly instantiated as development, staging, production

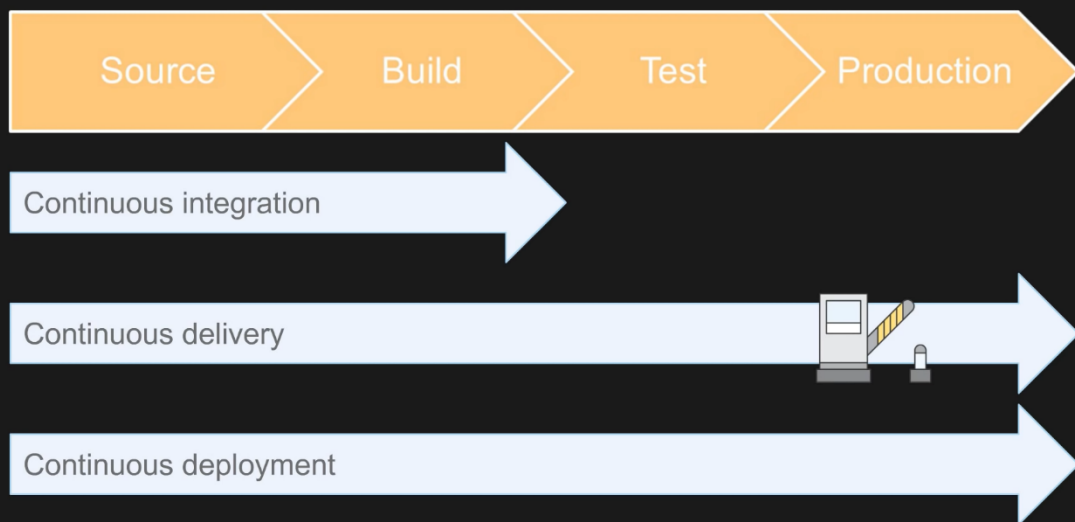
and sometimes to

- the *specific* environment, (which comes after development, staging etc), known as production.



Once Munns has introduced the four major phases (SBTP), he introduces three “process levels”. The three *process levels* are:

Release processes levels



Importantly, all the services which begin with the word “code”, such as `CodeDeploy` and `CodeCommit`, are sometimes referred to as “CodeStar”. CodeStar is not a distinct service, which you are yet to read about. CodeStar is, really “Code asterisk” or `Code*`, where the asterisk is similar to a wildcard.

CodeDeploy is deployed

The oldest service is CodeDeploy. This was released on 12th November 2014. This is the same year AWS Config, KMS, and AWS Directory Service was launched. Andy Jassy [announced](#) CodeDeploy at Re:Invent. Jassy said:

The capabilities you have with CodeDeploy are that you can choose to deploy to all of your instances at once.

But most who have large amounts of instances—what they *really* want to do it deploy to groups of instances, to make sure the deployment is working out.

So, we allow you to do rolling upgrades, rolling deployments.

So, that was CodeDeploy. It is the oldest tool and was announced by Andy Jassy.

It was 2014. Nations had deployed athletes at the Sochi 2014 Winter Olympics; vaccines had been deployed to help with the [Ebola epidemic](#) in Africa; investigators were deployed to search for [missing plane](#) MH370; troops were [deployed by Russia](#), to annex Crimea; and in Scotland, worker were deployed to [count ballot papers](#), to find out that Scotland decided to remain in the United Kingdom. **In 2014 CodeDeploy was deployed.**

[Troutman’s talk]

In 2015, Rob Brigham delivered a presentation with Clare Liguori¹ [\[Brigham 2015\]](#).

Brigham’s talk has X parts:

1. Monolith development lifecycle

¹ Liguori is pronounced “lih-GOR—ee”.

- a. EVEN when you have that large, new version

2. Service-oriented ARCHITECTURE

- a. One whose sole job it was to render the BUY button correctly.
- b. We had a single-purpose service whose sole job was to calculate the tax correctly on the checkout process.
- c. We also had a rule that:
 - i. They could only talk to one another through their web service APIs; there was no backend, shared data access allowed.
- d. What this enabled us to do was
 - i. to create a very highly decoupled architecture,
 - ii. where these services could **operate independently** of each other,
 - iii. without any coordination between these services.
 - iv. As long as they adhere to that standard web service interface.
- e. To give you an idea of what this ARCHITECTURE looked like, I've included this **graphic**.
 - i. What this represents is the Amazon.com retail website, circa 2009
 - ii. And all of these individual services that made up that experience.
 - iii. Back then when we made this ARCHITECTURAL SHIFT, we didn't have this term—but we'd call

this a **MICROSERVICES
ARCHITECTURE.**

- f. In addition to that **architectural change**, we also made an *organisational change*.
- 3.

4. Tools group

To fix that, we started a new tools group, that built a new breed of developer tools.

These tools had some unique characteristics.

(a) self-service

The tools had to be self-service. There's no way one tools group would be able to on-board thousands of "two pizza teams" if they required hand-holding.

This was very AWS-like, even before AWS had started.

(b) Technology-agnostic

We'd given these "two pizza" teams full autonomy to make whatever decision they wanted, and they took full advantage of that.

They chose different operating systems, different programming languages, architectures, app frameworks.

(c) Encourage best practices

We still wanted the two-pizza teams to **share** their learnings. If a team learned a best practice, we would take that and bake it into the toolset. That made it easy for other teams to both discover this best practice as well as adopt it themselves.

(d) Single-purpose services

You could say we drank the microservices koolaid. Just as we were teasing apart the website architecture, we didn't want to deliver an end-to-end tool chain, that was tightly coupled together.

We wanted to deliver it as these functional building block units. Teams could pick the units that work best for them, and then **tie them together** as needed.

So I want to talk about a couple of these building-block tool services, that we use internally at Amazon.

5. Apollo

The first is Apollo, which is our deployment engine. It's job is essentially to get bits onto a box. We've been using Apollo to deploy the retail website for over a dozen years now, and we also use it to deploy our Amazon Web Services. Over that time, we've learned a lot about how to **do deployments well**.

We've taken those learnings, and baked them into the tool.

a) Deployment service:

b) No downtime deployments

- As you can imagine, we're not allowed to take down the retail website any time we want to push a code change
- So we use **rolling updates**
- When Apollo is updating a fleet of application servers, it's only going to update a small fraction of those at a time
- It will incrementally work its way across the fleet

c) Health tracking:

- Occasionally, a bad code change can make its way through testing, and roll out into production
- What we want to make sure is that that bad code change is not going to take down the whole fleet
- It will automatically cut off that deployment

d) Versioned artifacts and rollbacks

6. The next service I want to talk about is PIPELINES

- a. This is our internal **continuous delivery engine**.
- b. Even **AFTER** we built Apollo and had automated deployments,
- c. We **STILL** noticed that it took a long time for a code change to go from a developer check-in to be running in production
- d. Being a data-driven company, we did a study on that.
- e. When we added up that data, and looked at the results, we were embarrassed
- f. It was on the order of weeks
- g. What we saw was that it wasn't the duration of one of those actions
 - i. It wasn't so much the *duration* of any one of those actions.
 - ii. It wasn't the duration of a build, or the duration of a test run, or the duration of a deployment
 - iii. It was all of this dead time in between.
- h. Bunch of **inefficient, manual** hand-offs
 - i. After one task was done, the person would take that and notify the next person.
 - ii. That normally happened in the form of an email.

A company like Amazon “**prides itself on efficiency**”

- i. A company that uses robots in its fulfilment centres
 - j. You can imagine how crazy it was that we were using humans to pass around these virtual bits in our software delivery process
 - k. So we had to fix that. We did that using Pipelines.
7. Pipelines allow these teams to model out their complete end-to-end release process
- a. They can specify how they want their source code changes to be automatically built and unit tested
 - b. How they wanted those to then be deployed to their test environments.
 - c. What tests they want to run on those
 - d. And then how they wanted those changes to move out into a production deployment.
 - e. After they modelled out that release process, Pipelines would automatically handle all those code changes.
 - f. It automatically trigger off the builds, automatically check the results.
8. We also saw another **improvement that we didn't expect.**
- a. The teams that had fully automated pipelines had more reliable releases than those that had manual steps involved.
 - b. That was a little unintuitive at first
 - c. Teams that fully dedicated themselves to making sure that every validation step that they wanted, was baked into an automated test—those teams had more reliable releases
 - d. ...that required less rollbacks

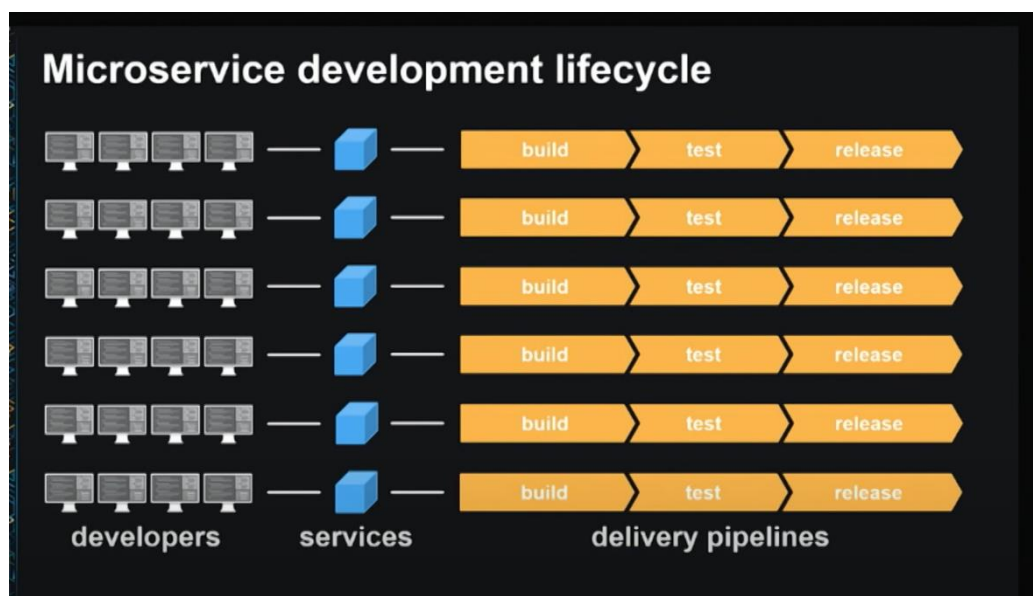
9. So, with these two advantages, of having FASTER and MORE RELIABLE releases, it's been incredibly successful inside of Amazon.

10.

In the first half of the talk, it appears that Rob Brigham is “telling a story”, or providing background, as is conventional for presentations of these sorts. However, this presentation is essential viewing. You will see concepts and terms explained here which you’ve seen others mentioning for a while. This is likely where they go it from. Brigham explains what “two pizza teams” are, how they structure the developers within Amazon itself, and microservice architectures.

He explains that on the retail website





We fully unblock these two-pizza teams from working independently.

Brigham's answer to how Amazon is able to move so fast is that:

it might look like large Amazon is a large organisation that has some internal overhead.

We're really structured like a bunch of small, start-up teams that are all operating and organised very efficiently, and moving as fast as they can.

Brigham tells us that CodeDeploy “has rolling updates”. This means that you can deploy without downtime. Brigham’s slide says:

- Easy and reliable deployments
- Scale with ease
- Deploy to any server

So, unsurprisingly, CodeDeploy helps with deployments. Two more features are of note: it scales well and has health checks (to cut off bad deployments). Brigham was keen to emphasise that when they launched CodeDeploy (Nov 2014), they

only supported deploying to Amazon EC2 instances.

And they had just added support for on-premises deployments [Brigham 2015]. Think of your own private data centre and VMs in *other* clouds. The idea was that you could now use CodeDeploy as your “**central tool**, to manage all your deployments, to all your different applications, in all of your different environments” [Brigham 2015].

Super Thursday

In 2015, two new services were launched on the very same day. Specifically, on Thursday 9th July 2015², CodeCommit and CodePipeline were released. This was *Super Thursday*.

What is CodeCommit?

In 2016, a [short promotional video](#) describing CodeCommit was released. The video starts by talking about the **problem** of storing and versioning code. It talks about the **problem** of managing and scaling your source control system. We will explain what a source control system is shortly. CodeCommit had actually been [announced](#) by Jassy at Re:Invent 2014. Jassy describes CodeCommit as a “managed code repository in the cloud”.

In the words of Brigham, CodeCommit *is* Git. It is

Git source control, re-implemented on top of S3 storage.

[Brigham 2015]

Brigham distinguishes between the front end and back end. On the front end, CodeCommit

works like any other Git **source control system** out there.

You use the same Git tools, issue the same Git commands. However, it’s on the *back* end CodeCommit uses some unique AWS features. Brigham tells us:

We’ve implemented Git on top of S2 and DynamoDB.

So, CodeCommit relies on the simple storage service and the non-relational database service DynamoDB. Finally,

² Boris Johnson—the mayor of London—could not get London Underground workers to commit to a pay deal. As a result, this Thursday saw the [largest strike](#) in 13 years. It was chaos.

In Russia, the [BRICS conference](#) saw several emerging economies **commit** to a [new bank](#).

encryption-at-rest is achieved. CodeCommit automatically encrypts your repositories using customer-specific keys.

This raises the question *What is a source control system?* This is sometimes called a **Source Code Control System (SCCS)**. An SCCS is a system designed to track changes in **source code**.

What is source code?

Source code is generally understood to mean programming statements that are created by a programmer with a text editor or a visual programming tool and then saved in a file [Washington].

Source code is different from object code. **Object code** generally refers to

the output, a compiled file, which is produced when the Source Code is compiled with a C compiler.

[Washington]

The idea is that SOURCE code is written using a human-readable programming language, usually as plain text [Wikipedia 2022]. In contrast, object code consists of machine-readable instructions. These instructions are processed by the CPU in a computer. Operating system or application software is usually in the form of compiled object code. Experts from the Linux Information Project put it this way:

Source code (also referred to as source or code) is the version of software as it is originally written (i.e. types into a computer) by a human in plain text (i.e. human readable alphanumeric characters).

[Linfo]

Let's bring this back to CodeCommit. This is a Source Code Control System, released on that Thursday in 2015. CodeCommit is a version control system. In other words, it's a system for identifying changes to the code, giving them a number or letter code. We are organising and controlling revisions to the code.

What is CodePipeline?



As we have seen, CodePipeline was released by AWS on Thursday in 2015 (although [here](#) is Jassy announcing it in 2014).

Rob Brigham outlines CodePipeline [here](#), saying:

After you define your own custom release process, CodePipeline is going to manage all of your

In the textbook by Piper and Clinton (2021) their discussion of Pipeline involves introducing these concepts:

1. Continuous integration
2. Continuous delivery
3. Source, built, test, [approval](#), [deploy](#), [invoke](#)
4. Artifact

The list of six items are action types. *Actions* are found in *stages*. And we are told:

Every CodeDeploy pipeline must include at least two stages and can have up to 10.

So, there is a little bit of terminology to get on top of. For example, ARTIFACT³ refers to

The collection of data, such as application source code, built applications, dependencies, definitions files, templates and so on,

That is worked on by pipeline actions.

³ Artefact is a British spelling (artifact is *also* a British spelling). British traditionalists have allowed the Americanism for a long time now.

Artifacts are produced by some actions and consumed by others.

In a pipeline, artifacts can be the set of files worked on by an action (input artifacts) or the updated output of a completed action (output artifacts).

AWS Documentation

[Rimple] [says](#) an artifact is:

A built object or objects that can be used by down-stage build tools like CodePipeline.

There are some ways to remember those six items by the way⁴. There are two things to note about the 6 action types.

First, the penultimate item is DEPLOY. Given that CodeDeploy is concerned with deployment of code, it is immediately clear that CodePipeline is a broad, overarching service. Just *one* of its action types is DEPLOY. The second thing to note is that this list is similar to a list introduced by Chris Munns (2017). He said that software development (in AWS at least), can be thought of as occurring in four stages:

1. Source
2. Build
3. Test
4. Production

You can think of his PRODUCTION category as containing the final three action types. That is, PRODUCTION covers *Approve*, *Deploy*, and *Invoke*.

⁴ You're sat having a cuppa with a developer who is explaining Pipeline. Suddenly, to your horror, he dips his biscuit in his drink. He explains "some biscuits taste amazing dipped in!" (SBT ADI).

- Snarky British tourists *always* demand ice.
-

In his 2015 presentation, Brigham tells a long background story. Then, the services are [introduced](#):

We have **CodeDeploy**, which you can use for automated deployments;

CodePipeline that you can use for end-to-end release automation;

And then, if you want to move your source code to the cloud—if you want your entire pipeline hosted on AWS—you can use CodeCommit to store your source code also.

[Brigham 2015]

Clare Liguori's presentation

The Release of CodeBuild

CodeBuild was released in late 2016. Specifically, Thursday [1st December 2016](#). Liguori reminds us that it is a **fully managed build service**:

It's a fully managed build service. What that means is that you don't have to worry about provisioning servers, patching servers, scaling.

It's continuously scaling.

[Liguori 2017].

Chris Munns did a talk on CodeBuild in January 2017 [Munns 2017]. Also see [Rimple 2020] for an extended presentation. Munns describes CodeBuild as a:

Fully managed build service that compiles source code, runs tests, and produces software packages.

Learning that CodeBuild is a build service doesn't seem to be learning much. I want to drill down, and see if we can appreciate some more concrete features of CodeBuild. Fortunately, a demonstration of the CodeBuild console is provided by [Maarek 2020b].

Choice Cuts

[AWS CodeBuild now supports GitHub Actions](#) - The documentation here is unclear; what I think is happening is that you can now run the giant pile of Actions that folks publish over on GitHub inside of your CodeBuild jobs. If so this is frankly genius; it really expands the scope of possibilities that CodeBuild has. A reminder: CodeBuild is my favorite service to run large jobs on a scheduled basis; it's still arguably AWS's best serverless container offering.

Email from Corey Quinn on July 17th 2023

AWS CodeBuild now supports GitHub Actions

Posted On: Jul 7, 2023

AWS CodeBuild customers can now use GitHub Actions during the building and testing of software packages. AWS CodeBuild is a fully managed continuous integration service that compiles source code, runs tests, and produces ready-to-deploy software packages. Customers' CodeBuild projects are now able to leverage many of the pre-built actions available in GitHub's marketplace. GitHub Actions are open source applications for the GitHub Actions platform that perform a complex but frequently repeated task.

With CodeBuild's integration with GitHub Actions, you can now extend your buildspec definition to invoke third-party solutions. There is no need to author and maintain custom integrations, or learn how integrate others' solutions into your build process.

CodeBuild's integration with GitHub actions is available in US East (N. Virginia), US East (Ohio), US West (N. California), see the US West (Oregon), South America (Sao Paulo), Asia Pacific (Singapore), Asia Pacific (Tokyo), Asia Pacific (Sydney), Asia Pacific (Hyderabad), Asia Pacific (Melbourne), Asia Pacific (Seoul), Asia Pacific (Mumbai), Asia Pacific (Osaka), Asia Pacific (Hong Kong), Asia Pacific (Jakarta), Europe (Ireland), Europe (Frankfurt), Europe (Zurich), Europe (Paris), Europe (Stockholm), Europe (Milan), Europe (London), Europe (Spain), Canada (Central), Middle East (Bahrain), Middle East (UAE), and Africa (Cape Town). For more information about the AWS Regions where CodeBuild is available, see the [AWS Regions page](#).

To learn more about how to get started with CodeBuild, visit the [AWS CodeBuild](#) product page. To learn more about CodeBuild's integration with GitHub Actions, see CodeBuild's documentation for [Running GitHub Actions](#).

AWS CodeBuild now supports custom images for AWS Lambda compute - I'm starting to wonder what the value is at all of the CodeBuild managed environments, given how effective it's become at managing other, less tetchy compute options.

Corey Quinn writing in an email on 25h March 2024

References

[AWS 2015]

“Deploying AWS CodeDeploy”. Available at:
https://www.youtube.com/watch?v=jcR9iIWdU7E&ab_channel=AmazonWebServices

[AWS 2016]

Introduction to AWS CodeCommit. Available at:
https://www.youtube.com/watch?v=46PRLMW8otg&ab_channel=AmazonWebServices

[AWS 2020]

‘Streamline your Software Release Process Using AWS CodePipeline’. Available at:
https://www.youtube.com/watch?v=zMa5gTLrz_mQ&ab_channel=AmazonWebServices

[AWS 2023]

AWS CodeBuild now supports GitHub Actions. Available at:
<https://aws.amazon.com/about-aws/whats-new/2023/07/aws-codebuild-github-actions/?ck_subscriber_id=1560524742>

[Brigham 2015]

“DevOps at Amazon: A Look at Our Tools and Processes”. October 2015. Available at:
https://www.youtube.com/watch?v=esEFaY0FDKc&ab_channel=AmazonWebServices

[Davis 2022]

Davis, Neal. What is AWS CodeDeploy. Available at:
https://www.youtube.com/watch?v=lbBg9FMnnPM&ab_channel=DigitalCloudTraining

[Jassy 2014a]

Jassy, Andy. (2014). ‘Announcing AWS CodeDeploy’. Available at:
https://www.youtube.com/watch?v=0fNWTYbwAfg&ab_channel=AmazonWebServices

[Jassy 2014b]

Jassy, Andy. (2014). Announcing AWS CodePipeline. Available at:
https://www.youtube.com/watch?v=d01AKo7HAng&ab_channel=AmazonWebServices

[Liguori 2017]

Liguori, Clare (2017). Continuous Integration Best Practices for Software Development. Available at:
https://www.youtube.com/watch?v=GEPJ7Lo346A&ab_channel=AmazonWebServices

[Linfo]

“Source Code Definition”. The Linux Information Project. Available at:
http://www.linfo.org/source_code.html

[Maarek 2020a]

Maarek, Stephanie (2020). “AWS CodePipeline – Artifacts, Encryption, S3 Introduction”. Available at: https://www.youtube.com/watch?v=u5UV-wr-IIA&ab_channel=StephaneMaarek

[Maarek 2020b]

Maarek, Stephanie (2020). “AWS CodeBuild Tutorial”. Available at: https://www.youtube.com/watch?v=qGgNyOkZEb0&ab_channel=StephaneMaarek

[Munns 2017]

Munns, Chris (2017). Announcing AWS CodeBuild. Available at: https://www.youtube.com/watch?v=E6EBg46vvu0&ab_channel=AWSOnlineTechTalks

[Suse]

Available at: <https://www.suse.com/suse-defines/definition/production-environment/#:~:text=A%20production%20environment%2C%20sometimes%20called,%2C%20testing%2C%20staging%20and%20production.>

[Troutman 2015]

Troutman, Andy (2015). “Getting started with AWS CodeDeploy”. 29th April 2015. Available at: https://www.youtube.com/watch?v=zHOQnTNxIKM&ab_channel=AWSOnlineTechTalks

[Rimple 2020]

Rimple, Ken (2020). “All the AWS CodeBuild you can stomach in 45 minutes” Available at: https://www.youtube.com/watch?v=yCVR-uqc4qk&ab_channel=ChariotSolutions

[Washington]

University of Washington Research. Available at: <https://www.washington.edu/research/glossary/source-code-and-object->

[code/#:~:text=Source%20code%20is%20generally%20understood,compiled%20with%20a%20C%20compiler.](#)

[Wikipedia 2022a]

Wikipedia. “Source Code”. Available at:
https://en.wikipedia.org/wiki/Source_code [Accessed
11th August 22]


[Wikipedia 2022b]

Deployment environment. Available at:
[https://en.wikipedia.org/wiki/Deployment_environment
#Production](https://en.wikipedia.org/wiki/Deployment_environment#Production)

Image credits

HDPE pipeline in Australia. See
[https://commons.wikimedia.org/wiki/File:HDPE_Pipeline_in_a
_harsh_Australian_environment.jpg](https://commons.wikimedia.org/wiki/File:HDPE_Pipeline_in_a_harsh_Australian_environment.jpg)

AppStream 2.0



Bibliography

- I. Official
- II. Unofficial
- III. Critical
- IV. General

I. Official


[Woodard 2023]

Woodard, Sam and Irshad Buchh and Anil Konda (2023). Drug Analyzer on AWS Provides Analytics That Inform Decisions and Support New Therapies. *AWS Partner Network (APN) Blog*. June 5th 2023. Available at: https://aws.amazon.com/blogs/apn/drug-analyzer-on-aws-provides-analytics-that-inform-treatment-decisions-and-support-new-therapies/?ck_subscriber_id=1560524742

II. Unofficial

III. Critical

IV. General



Bibliography

- I. Official
- II. Unofficial
- III. Critical
- IV. General

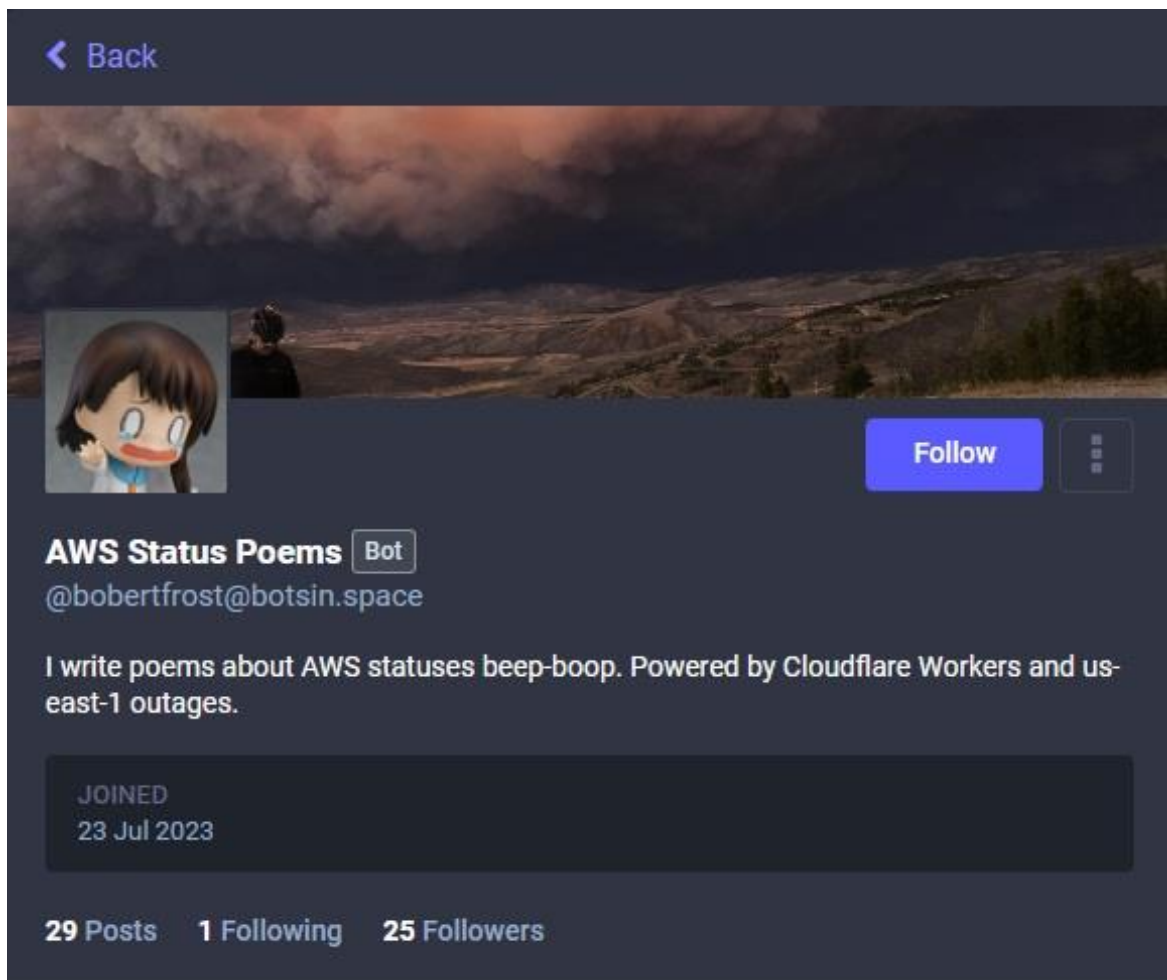

I. Official

II. Unofficial

III. Critical

IV. General

Personal Health Dashboard





AWS Status Poems

@bobertfrost

📅 21 Aug

Service is operating normally: [RESOLVED] Increased SAML Error Rates

Roses are red,
Violets are blue,
AWS had an incident,
With SAML error rates it went askew.
But fear not, it's resolved,
Service is operating normally, too.
A configuration change was to blame,
But now it's fixed, the issue tamed.
So logins can continue without any fuss,
Thankful for AWS and their quick adjust.

Source: status.aws.amazon.com/#signin-...



Bibliography

- I. Official
- II. Unofficial
- III. Critical
- IV. General

I. Official

II. Unofficial


[Mitchell]

Mitchell, Colin. AWS Status poems. Available at:
https://botsin.space/@bobbertfrost/with_replies

III. Critical

IV. General

Systems Manager






Werner Vogels announcing "EC2 Systems Manager" in 2016

Spid Spar

(?)



S - STATE MANAGER
P - PARAMETER STORE
I - INVENTORY
D - DISTRIBUTOR

S - SESSION MANAGER
P - PATCH MANAGER
A - AUTOMATION
R - RUN COMMAND



Two spider man people sparring with one another



Maitreya Ranganath explaining Systems Manager in 2017

Contents

2. State manager
3. Parameter store
4. Inventory
5. Distributor
6. Session manager
7. Patch manager
8. Automation
9. Run Command

State Manager

QUESTION 20 OF 33

20. QUESTION

A security engineer requires a solution for allowing employees to connect to a command line interface on Amazon EC2 Linux instances without using SSH keys or ports.

Which solutions meets these requirements?

- ☐ Use AWS Secrets Manager to store SSH keys. Instruct the employees to use the AWS CLI to retrieve the SSH key and connect to the EC2 Linux instances.
- ☐ Use AWS Systems Manager Run Command to open an SSH connection to the EC2 Linux instances. Grant the IAM user accounts permissions to use Run Command.
- ☒ Use AWS Systems Manager Session Manager. Grant the IAM user accounts permissions to use Systems Manager Session Manager.
- ☐ Use a bastion host EC2 instance in a public subnet. Use the bastion instance to connect to the EC2 Linux instances using an X.509 certificate.

Correct

Explanation:

Session Manager is a fully managed AWS Systems Manager capability. With Session Manager, you can manage your Amazon Elastic Compute Cloud (Amazon EC2) instances, edge devices, and on-premises servers and virtual machines (VMs).

You can use either an interactive one-click browser-based shell or the AWS Command Line Interface (AWS CLI). Session Manager provides secure and auditable node management without the need to open inbound ports, maintain bastion hosts, or manage SSH keys.

Session Manager helps you improve your security posture by letting you close SSH ports, freeing you from managing SSH keys and certificates, bastion hosts, and jump boxes.

CORRECT: "Use AWS Systems Manager Session Manager. Grant the IAM user accounts permissions to use Systems Manager Session Manager" is the correct answer (as explained above.)

INCORRECT: "Use AWS Systems Manager Run Command to open an SSH connection to the EC2 Linux instances. Grant the IAM user accounts permissions to

INCORRECT: "Use AWS Secrets Manager to store SSH keys. Instruct the employees to use the AWS CLI to retrieve the SSH key and connect to the EC2 Linux instances" is incorrect.

Secrets Manager can be used for storing secrets but storing SSH keys does not provide a solution as once retrieved the users would still need to connect via the SSH protocol. The public keys must also be stored on the server.

References:

<https://docs.aws.amazon.com/systems-manager/latest/userguide/session-manager.html>

8. QUESTION

A security team has requested that all existing and new Amazon RDS databases are encrypted at rest using AWS Key Management Service (KMS) encryption keys. A security engineer must identify which RDS databases are currently unencrypted and devise a plan for enabling encryption.

Which combination of steps should the security engineer take to accomplish this? (Select TWO.)

- ☐ Create a snapshot of unencrypted databases. Copy the unencrypted snapshots to created encrypted snapshots. Restore the databases from the encrypted snapshots.
- ☐ Use AWS Config to detect any existing and new unencrypted databases. Configure an Amazon SNS notification to alert the security team.
- ☐ Create an encrypted read replica of unencrypted database instances.
- ☐ Use AWS System Manager State Manager to detect the RDS database encryption status. Create an Amazon SNS notification to alert the security team.
- ☐ Enable encryption for the Amazon RDS database instances.

INCORRECT: "Create an encrypted read replica of unencrypted RDS database instances. Promote the replicas to be the primary instances and delete the unencrypted database instances" is incorrect.

You cannot enable encryption after creation of the database, and this includes for any instances created from the o standby instances.

INCORRECT: "Use AWS System Manager State Manager to detect the RDS database encryption status. Create an Amazon SNS notification to alert the security team" is incorrect.

AWS Systems Manager State Manager is used to manage the configuration on EC2 instances and on-premises servers. It is not used for RDS database encryption.

References:

<https://docs.aws.amazon.com/AmazonRDS/latest/UserGuide/Overview.Encryption.html>

<https://docs.aws.amazon.com/config/latest/developerguide/notifications-for-AWS-Config.html>

<https://docs.aws.amazon.com/config/latest/developerguide/rds-storage-encrypted.html>

Incorrect

Explanation:

With AWS Config, you can continuously monitor and record configuration changes of your AWS resources. There are several managed rules that you can use. The `rds-storage-encrypted` managed rule checks whether storage encryption is enabled for your RDS DB instances. You can configure an Amazon SNS notification based on the result using Amazon SNS.

Once the unencrypted databases have been discovered the next task is to enable encryption. The key fact to remember here is that you cannot enable encryption on an existing RDS database instance. You also cannot create encrypted replicas from unencrypted database instances.

The only solution is to create a snapshot (which will be unencrypted) and subsequently create an encrypted copy of the snapshot. Then, create a new database instance from the encrypted snapshot. The new database will be encrypted and will have a new endpoint address.

CORRECT: "Create a snapshot of unencrypted databases. Copy the unencrypted snapshots to created encrypted snapshots. Restore the databases from the encrypted snapshots" is a correct answer (as explained above.)

CORRECT: "Use AWS Config to detect any existing and new unencrypted databases. Configure an Amazon SNS notification to alert the security team" is also a correct answer (as explained above.)

Bibliography

- I. Official
- II. Unofficial
- III. Critical
- IV. General

I. Official

II. Unofficial

III. Critical

IV. General

Parameter Store

24. QUESTION

A company has several AWS Lambda functions. While reviewing the Lambda functions a security engineer discovers that sensitive information is being stored in environment variables and is viewable as plaintext in the Lambda console. The values of the sensitive information are less than 8 KB and there are over 10,000 values stored across the functions.

What is the MOST cost-effective way to address this security issue?

- ☐ Store the environment variables in an encrypted Amazon EFS file system and access them at runtime. Use POSIX permissions to restrict access to only the Lambda functions that require access.
- ☐ Store the environment variables in AWS Secrets Manager and access them at runtime. Use IAM permissions to restrict access to the secrets to only the Lambda functions that require access.
- ☐ Use AWS Config to store the environment variables. Access the environment variables at runtime. Use IAM permissions to restrict access to the environment variables to only the Lambda functions that require access.
- ☒ Store the environment variables in AWS Systems Manager Parameter Store as secure string parameters and access them at runtime. Use IAM permissions to restrict access to the parameters to only the Lambda functions that require access.

Incorrect

Explanation:

AWS Secrets Manager is well suited to this use case and is the most cost-effective option. For small amounts of data under 4 KB, Systems Manager Parameter store would be cheaper when using standard parameters. However, if you have more than 10,000 parameters or the parameters are over 4 KB (and up to 8 KB) in size, advanced parameters must be used which would be more expensive.

Both Systems Manager Parameter Store and Secrets Manager support controlling access to the information stored using IAM policies.

CORRECT: "Store the environment variables in AWS Secrets Manager and access them at runtime. Use IAM permissions to restrict access to the secrets to only the Lambda functions that require access" is the correct answer (as explained above.)

INCORRECT: "Store the environment variables in AWS Systems Manager Parameter Store as secure string parameters and access them at runtime. Use IAM permissions to restrict access to the parameters to only the Lambda functions that require access" is incorrect.

This would be more expensive for this specific use case as explained above.

INCORRECT: "Store the environment variables in an encrypted Amazon EFS file system and access them at runtime. Use POSIX permissions to restrict access to only the Lambda functions that require access" is incorrect.

EFS would be an expensive solution and Lambda cannot mount an EFS file system.

INCORRECT: "Use AWS Config to store the environment variables. Access the environment variables at runtime. Use IAM permissions to restrict access to the environment variables to only the Lambda functions that require access" is incorrect.

Config cannot be used to store information such as environment variables.

“For small amounts of data under **4 KB**, Systems Manager Parameter store would be cheaper when using standard parameters”

11. QUESTION

A Solutions Architect is developing a mechanism to gain security approval for Amazon EC2 images (AMIs) so that they can be used by developers. The AMIs must go through an automated assessment process (CVE assessment) and be marked as approved before developers can use them. The approved images must be scanned every 30 days to ensure compliance. Which combination of steps should the Solutions Architect take to meet these requirements while following best practices? (Select TWO.)

- ☐ Use AWS GuardDuty to run the CVE assessment package on the EC2 instances launched from the approved AMIs.
- ☐ Use AWS Lambda to write automatic approval rules. Store the approved AMI list in AWS Systems Manager Parameter Store. Use a managed AWS Config rule for continuous scanning on all EC2 instances and use AWS Systems Manager Automation documents for remediation.
- ☐ Use Amazon Inspector to run the CVE assessment package on the EC2 instances launched from the approved AMIs.
- ☐ Use AWS Lambda to write automatic approval rules. Store the approved AMI list in AWS Systems Manager Parameter Store. Use Amazon EventBridge to trigger an AWS Systems Manager Automation document on all EC2 instances every 30 days.
- ☐ Use the AWS Systems Manager EC2 agent to run the CVE assessment on the EC2 instances launched from the approved AMIs.

11. QUESTION

A Solutions Architect is developing a mechanism to gain security approval for Amazon EC2 images (AMIs) so that they can be used by developers. The AMIs must go through an automated assessment process (CVE assessment) and be marked as approved before developers can use them. The approved images must be scanned every 30 days to ensure compliance. Which combination of steps should the Solutions Architect take to meet these requirements while following best practices? (Select TWO.)

- ☐ Use AWS GuardDuty to run the CVE assessment package on the EC2 instances launched from the approved AMIs.
- ☐ Use AWS Lambda to write automatic approval rules. Store the approved AMI list in AWS Systems Manager Parameter Store. Use a managed AWS Config rule for continuous scanning on all EC2 instances and use AWS Systems Manager Automation documents for remediation.
- ☒ Use Amazon Inspector to run the CVE assessment package on the EC2 instances launched from the approved AMIs.
- ☒ Use AWS Lambda to write automatic approval rules. Store the approved AMI list in AWS Systems Manager Parameter Store. Use Amazon EventBridge to trigger an AWS Systems Manager Automation document on all EC2 instances every 30 days.
- ☐ Use the AWS Systems Manager EC2 agent to run the CVE assessment on the EC2 instances launched from the approved AMIs.

Correct

Explanation:

AWS Lambda can be used to run the approval process for the AMIs and then automatically store the results in AWS Systems Manager Parameter Store.

For the CVE assessment, Amazon Inspector can be used to perform security assessments of Amazon EC2 instances by using AWS managed rules packages such as the Common Vulnerabilities and Exposures (CVEs) package.

Amazon EventBridge (CloudWatch Events) can then be used to create scheduled triggers that run AWS Systems Manager Automation documents on a recurring schedule (30 days). AWS Systems Manager will update the running instances to ensure they are up to date with any security updates that need to be applied.

CORRECT: "Use AWS Lambda to write automatic approval rules. Store the approved AMI list in AWS Systems Manager Parameter Store. Use Amazon EventBridge to trigger an AWS Systems Manager Automation document on all EC2 instances every 30 days" is a correct answer.

CORRECT: "Use Amazon Inspector to run the CVE assessment package on the EC2 instances launched from the approved AMIs" is also a correct answer.

INCORRECT: "Use the AWS Systems Manager EC2 agent to run the CVE assessment on the EC2 instances launched from the approved AMIs" is incorrect. Systems Manager does not have a CVE assessment, use Amazon Inspector which is designed for this purpose and has a package preconfigured.

INCORRECT: "Use AWS Lambda to write automatic approval rules. Store the approved AMI list in AWS Systems Manager Parameter Store. Use a managed AWS Config rule for continuous scanning on all EC2 instances and use AWS Systems Manager Automation documents for remediation" is incorrect. Amazon Inspector is a better fit for a CVE assessment.

INCORRECT: "Use AWS GuardDuty to run the CVE assessment package on the EC2 instances launched from the approved AMIs" is incorrect. GuardDuty is an intelligent threat detection service. It is not suitable for a CVE assessment.

14. QUESTION

A security engineer is attempts to encrypt a secure string parameter value in AWS Systems Manager Parameter Store with an AWS KMS key and receives an *InvalidKeyId* error message.

Why was this error message generated?

- ☐ The KMS key specified is currently in use.
- ☐ The KMS key specified does not exist.
- ☒ The KMS key specified is not enabled.
- ☐ The KMS key specified is not compliant.

Correct

Explanation:

To perform any operation on a secure string parameter, Parameter Store must be able to use the Amazon KMS key that you specify for your intended operation. Most of the Parameter Store failures related to KMS keys are caused by the following problems:

- The credentials that an application is using do not have permission to perform the specified action on the KMS key.
- The KMS key is not found. This typically happens when you use an incorrect identifier for the KMS key.
- The KMS key is not enabled. When this occurs, Parameter Store returns an **InvalidKeyId** exception with a detailed error message from Amazon KMS.

The specific error message received indicates that the issue is due to the KMS key being disabled.

CORRECT: "The KMS key specified is not enabled" is the correct answer (as explained above.)

INCORRECT: "The KMS key specified is not compliant" is incorrect.

There is no compliance requirement for a key to work with Parameter Store.

INCORRECT: "The KMS key specified does not exist" is incorrect.

The specific error generated indicates that the key is not enabled.

INCORRECT: "The KMS key specified is currently in use" is incorrect.

KMS keys do not get locked to a single process and can be used by multiple processes at the same time.

References:

https://docs.amazonaws.cn/en_us/kms/latest/developerguide/services-parameter-store.html

12. QUESTION

An application running on Amazon EC2 instances reads secrets stored in AWS Systems Manager Parameter Store. The application issued GetParameter API calls for secure string parameters and the calls failed.

Which factors could be the cause of this failure? (Select TWO.)

- ☐ The IAM role assigned to the EC2 instance profile does not have permissions to retrieve parameters in Systems Manager Parameter Store.
- ☒ Systems Manager Parameter Store does not have decrypt permissions on the AWS KMS key used to encrypt the parameter.
- ☐ The IAM role assigned to the EC2 instance profile does not have encrypt permissions on the AWS KMS key used to encrypt the parameter.

Incorrect

Explanation:

To perform a GetParameter API call to a secure string parameter there are two permissions required. Firstly, the IAM role assigned to the EC2 instance must have the permission to execute the `ssm:GetParameter` API action which is an AWS Systems Manager Parameter Store permission.

The parameters are stored as secure strings which means they will be encrypted with an AWS KMS key. To read the values they must be decrypted so the EC2 instance will also need the `Decrypt` API permission which is associated with AWS KMS.

CORRECT: "The IAM role assigned to the EC2 instance profile does not have decrypt permissions on the AWS KMS key used to encrypt the parameter" is a correct answer (as explained above.)

CORRECT: "The IAM role assigned to the EC2 instance profile does not have permissions to retrieve parameters in Systems Manager Parameter Store" is also a correct answer (as explained above.)

INCORRECT: "The IAM role assigned to the EC2 instance profile does not have encrypt permissions on the AWS KMS key used to encrypt the parameter" is incorrect.

In this case the instance is attempting to retrieve and therefore decrypt the secure string, not encrypt it.

INCORRECT: "Systems Manager Parameter Store does not have decrypt permissions on the AWS KMS key used to encrypt the parameter" is incorrect.

Parameter store is not executing the API action to decrypt the secure string, the EC2 instance is executing the API action.

INCORRECT: "Systems Manager Parameter Store does not have encrypt permissions on the AWS KMS key used to encrypt the parameter" is incorrect.

Parameter store is not executing the API action to encrypt or decrypt the secure string, the EC2 instance is executing the API action.

References:

18. QUESTION

A security engineer is building an application that is running on Amazon EC2. The application communicates with an Amazon RDS MySQL instance and authenticates with a user name and password. The credentials should be encrypted and rotated every 60 days.

Which steps should the engineer take to protect the credentials and ensure they can be automatically rotated?

- ☐ Store the credentials as an encrypted string parameter in AWS Systems Manager Parameter Store. Enable automatic rotation every 60 days and grant permission to the EC2 instance role to retrieve the parameter programmatically.
- ☐ Store the credentials in AWS Secrets Manager and choose an AWS KMS key. Enable automatic rotation every 60 days and configure the application to retrieve the secret programmatically.
- ☐ Store the credentials in an Amazon S3 bucket configured with SSE-KMS encryption. Grant permission to the EC2 instance role to retrieve the credentials from S3 programmatically. Update the credential files every 60 days.
- ☐ Store the credentials on an encrypted Amazon EFS volume. Configure the application instances to mount the volume. Use an AWS Lambda function to update the credentials on the EFS volume every 60 days.

Explanation:

AWS Secrets Manager can automatically rotate secrets for Amazon RDS Database Service (Amazon RDS). Secrets Manager can be configured to retrieve the secrets from Secrets Manager using the API.

This is an expensive and unnecessary solution. Secrets Manager is built for this purpose and will be lower cost and more efficient.

References:

<https://aws.amazon.com/blogs/security/how-to-use-aws-secrets-manager-rotate-credentials-amazon-rds-database-types/>

CORRECT: "Store the credentials in AWS Secrets Manager and choose an AWS KMS key. Enable automatic rotation every 60 days and configure the application to retrieve the secret programmatically" is the correct answer (as explained above.)

INCORRECT: "Store the credentials as an encrypted string parameter in AWS Systems Manager Parameter Store. Enable automatic rotation every 60 days and grant permission to the EC2 instance role to retrieve the parameter programmatically" is incorrect.

Parameter Store does not automatically rotate credentials.

INCORRECT: "Store the credentials in an Amazon S3 bucket configured with SSE-KMS encryption. Grant permission to the EC2 instance role to retrieve the credentials from S3 programmatically. Update the credential files every 60 days" is incorrect.

This is not an example of automatic rotation; Secrets Manager is a better solution.

INCORRECT: "Store the credentials as an encrypted string parameter in AWS Systems Manager Parameter Store. Enable automatic rotation every 60 days and grant permission to the EC2 instance role to retrieve the parameter programmatically" is incorrect.

Parameter Store does not automatically rotate credentials.

INCORRECT: "Store the credentials in an Amazon S3 bucket configured with SSE-KMS encryption. Grant permission to the EC2 instance role to retrieve the credentials from S3 programmatically. Update the credential files every 60 days" is incorrect.

This is not an example of automatic rotation; Secrets Manager is a better solution.

INCORRECT: "Store the credentials on an encrypted Amazon EFS volume. Configure the application instances to mount the volume. Use an AWS Lambda function to update the credentials on the EFS volume every 60 days" is incorrect.

This is an expensive and unnecessary solution. Secrets Manager is built for this purpose and will be lower cost and more efficient.

References:

<https://aws.amazon.com/blogs/security/how-to-use-aws-secrets-manager-rotate-credentials-amazon-rds-database-types-oracle/>

8. QUESTION

A company must ensure that AWS CloudTrail is recording API activity across all AWS Regions within their account. An automated solution is required to check that CloudTrail is enabled and to turn it back on if it has been turned off.

What is the MOST efficient way to implement this solution?

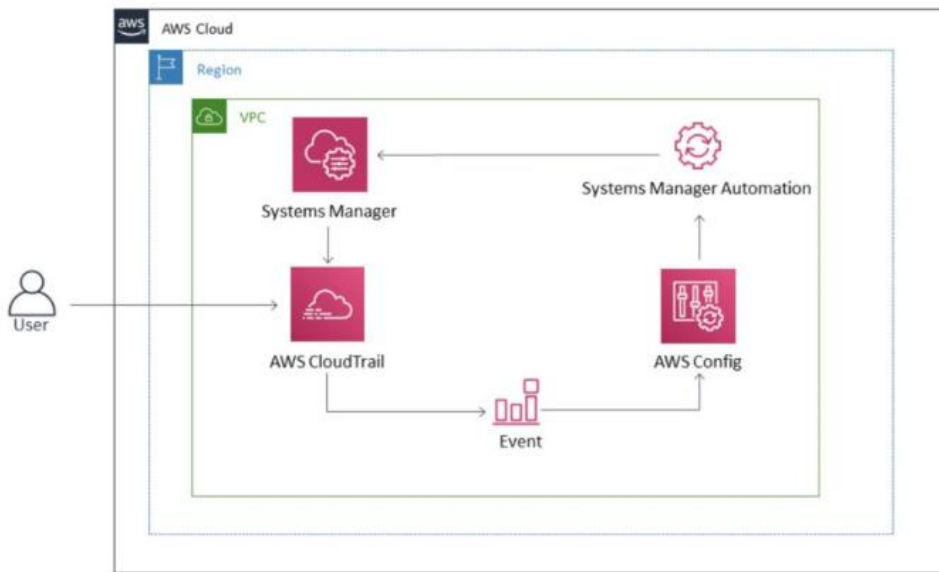
- ☐ Use Amazon Athena to monitor the Amazon S3 buckets where CloudTrail logging occurs. If logging ceases trigger an automated action that executes Systems Manager Automation to remediate the issue.
- ☐ Create an Amazon CloudWatch alarm for the AWS CloudTrail StopLogging API action. Configure remediation using an AWS Step Functions State Machine with an AWS Lambda function that turns CloudTrail back on.
- ☐ Create an Amazon EventBridge event with the event type "AWS API Call via CloudTrail" and configure AWS Lambda as a target. Configure the Lambda function to turn CloudTrail back on.
- ☒ Use AWS Config with the managed rule cloudtrail-enabled to check that CloudTrail is enabled. If the rule is NON_COMPLIANT use Systems Manager Automation to automatically remediate the issue.

Correct

Explanation:

To ensure that CloudTrail remains enabled in your account, AWS Config provides the cloudtrail-enabled managed rule. If CloudTrail is turned off, the cloudtrail-enabled rule automatically re-enables it by using automatic remediation.

This solution uses AWS Config to identify if CloudTrail logging is turned off and then an AWS Systems Manager Automation runbook to remediate the issue. The following diagram depicts this solution:



CORRECT: "Use AWS Config with the managed rule cloudtrail-enabled to check that CloudTrail is enabled. If the rule is NON_COMPLIANT use Systems Manager Automation to automatically remediate the issue" is the correct answer (as explained above.)

INCORRECT: "Create an Amazon EventBridge event with the event type "AWS API Call via CloudTrail" and configure AWS Lambda as a target. Configure the Lambda function to turn CloudTrail back on" is incorrect.

This event will trigger when certain API calls are made via AWS CloudTrail. It does not check if CloudTrail itself is being modified unless specifically targeting those API calls. This would require more work to implement properly compared to the correct answer.

INCORRECT: "Use Amazon Athena to monitor the Amazon S3 buckets where CloudTrail logging occurs. If logging ceases trigger an automated action that executes Systems Manager Automation to remediate the issue" is incorrect.

Athena is used for running SQL queries on S3 data. This is not a good way to check for changes to AWS CloudTrail.

INCORRECT: "Create an Amazon CloudWatch alarm for the AWS CloudTrail StopLogging API action. Configure remediation using an AWS Step Functions State Machine with an AWS Lambda function that turns CloudTrail back on" is incorrect.

You can alert based on metric filters that check for CloudTrail API actions. However, this would be more complex, and Step Functions would not be used as the target, Lambda would be.

References:

<https://docs.aws.amazon.com/prescriptive-guidance/latest/patterns/automatically-re-enable-aws-cloudtrail-by-using-a-custom-remediation-rule-in-aws-config.html>

AWS Systems Manager Parameter Store increases API throughput limit - This is kind of amazing. I remember when Parameter Store first launched and I smacked face-first into the original rate limit (which as I recall was something like 30 requests per second?) when spinning up a pile of ECS containers all at once at a client. With a new default limit of 10,000 transactions per second, I can still cause this behavior but now I have to work really really hard to do it.

Corey Quinn in an email on 10th July 2023

AWS Systems Manager Parameter Store increases API throughput limit

Posted On: Jul 3, 2023

Parameter Store, a capability of AWS Systems Manager, now supports up to 10,000 transactions per second (TPS) for [GetParameter](#) and [GetParameters](#) APIs, increased from the previous 3,000 TPS limit. Parameter Store allows you to securely store configuration data and secrets as hierarchical key-value pairs. You can flexibly store parameters such as API keys, subnet IDs, and passwords, and you can reference those parameters in your code and through AWS services such as AWS Lambda, Amazon EC2, and AWS CloudFormation. This increased API limit

makes it easier to use Parameter Store to support high-traffic applications and demanding workloads without sacrificing performance due to API throttling.

The 10,000 TPS limit is available when you enable the higher throughput setting in your Parameter Store account. If you already have the higher throughput setting enabled, you can start using 10,000 TPS today with no additional steps required.

To turn on higher throughput for Parameter Store, visit the [documentation](#). For information about pricing, visit the [Systems Manager Pricing page](#). This feature is available in all [AWS Regions](#) where AWS Systems Manager is available.

[AWS Cloud Operations & Migrations Blog](#)

Downgrade SQL Server Enterprise edition using AWS Systems Manager Document to reduce cost

by Yogi Barot, Sabarinath Nair, and Vikas Babu Gali | on 05 JUL 2023 | in [Amazon EC2](#), [AWS Systems Manager](#), [Integration & Automation](#), [Management Tools](#), [Technical How-To](#) | [Permalink](#) | [Share](#)

In this post, we will show how to downgrade SQL Server from Enterprise edition to Standard edition on [Amazon Elastic Compute Cloud\(EC2\)](#) instances to help you reduce cost. If you are not using any of the features of [Enterprise edition](#), you can downgrade to Standard edition.

Here is the flowchart that can help you identify the most used Enterprise edition features.

Bibliography

- I. Official
- II. Unofficial
- III. Critical
- IV. General

I. Official

[Barot 2023]


Barot, Yogi and Sabarinath Nair and Vikas Gali (2021). Downgrade SQL Server Enterprise Edition using AWS Systems Manager Document to reduce cost. Courtesy of *Last Week in AWS*. Available at: <<https://aws.amazon.com/blogs/mt/downgrade-sql-server-enterprise-edition-using-aws-systems-manager-document-to-reduce-cost/>>

II. Unofficial

III. Critical

IV. General

AWS Shield



Bibliography

- I. Official
- II. Unofficial
- III. Critical
- IV. General


I. Official

II. Unofficial

III. Critical

IV. General

Batch



Bibliography

- I. Official
- II. Unofficial
- III. Critical
- IV. General


I. Official

II. Unofficial

III. Critical

IV. General

X-ray



Bibliography

- I. Official
- II. Unofficial
- III. Critical
- IV. General


I. Official

II. Unofficial

III. Critical

IV. General

Glue



Bibliography

- I. Official
- II. Unofficial
- III. Critical
- IV. General

I. Official

II. Unofficial

III. Critical

IV. General

Artifact



Introducing AWS Artifact: Speeding Access to Compliance Reports

by Sara Duffer | on 07 DEC 2016 | in [AWS](#)
[Artifact](#), [Compliance](#) | [Permalink](#) | [Comments](#) | [Share](#)



I am pleased to announce AWS Artifact, a no-cost, self-service audit report and certification retrieval portal in the AWS Management Console that gives AWS customers on-demand access to [AWS compliance reports](#).

To document the current and historical compliance of the AWS infrastructure and services, many AWS customers provide compliance reports—including those for [ISO](#), [SOC](#), and [PCI](#)—to their auditors or regulators. You can now sign in to the AWS Management Console on your computer or mobile phone, and pull relevant reports in minutes. You can also give auditors and regulators direct access to one or more AWS compliance reports [using AWS Identity and Access Management \(IAM\) permissions](#).

AWS Director of Risk and Compliance Chad Woolf spoke about his vision of Artifact: “Naturally, we’re excited to provide customers and their auditors with selection and convenience when assessing the security that AWS provides,” Woolf said. “The release of AWS Artifact sets the stage for AWS to transform the auditing industry, moving auditing from being time-intensive and manual to highly automated and continuous in the cloud.”

You can start downloading the audit reports in the AWS Management Console today. Many of the documents are confidential and require you to accept Amazon’s confidentiality

terms and conditions, but after you review and agree to those terms, you will be granted instant access to review documents. You can also see [Getting Started with AWS Artifact](#) for more details.

To learn more about Artifact, see the [Artifact home page](#). See the [AWS Cloud Compliance home page](#) for more about AWS Cloud compliance and certifications.

– Sara

TAGS: [AWS Artifact](#), [Compliance reports](#), [Security Blog](#)

Above is the AWS blogpost introducing AWS Artifact. Have a read of the article. There are a few acronyms which you will need to master: ISO is the International Organization for Standardisation, SOC stands for System and Organization Controls, PCI stands for Payment Card Industry.

What is SOC?

SOC stands for Service Organization Controls. This is a standardized audit report, and evolved from an older audit report called SAS-70. SAS stands for Statement on Auditing Standards.

The American Institute of Certified Public Accountants (AICPA) was responsible for SAS-70.

The AICPA has its headquarters in Durham, North Carolina.



Durham, North Carolina

City of Medicine

The City of Durham is famously known as the “City of Medicine,” with healthcare as a major industry including more than 300 medical and health-related companies and medical practices. You are now one of our 2,400 valued employees working in one of the City's 24 departments serving 245,475 city residents.



DurhamNC.gov

<https://www.durhamnc.gov> › Welcome-to-The-City-of-D...

Welcome to The City of Durham




About featured snippets



Feedback

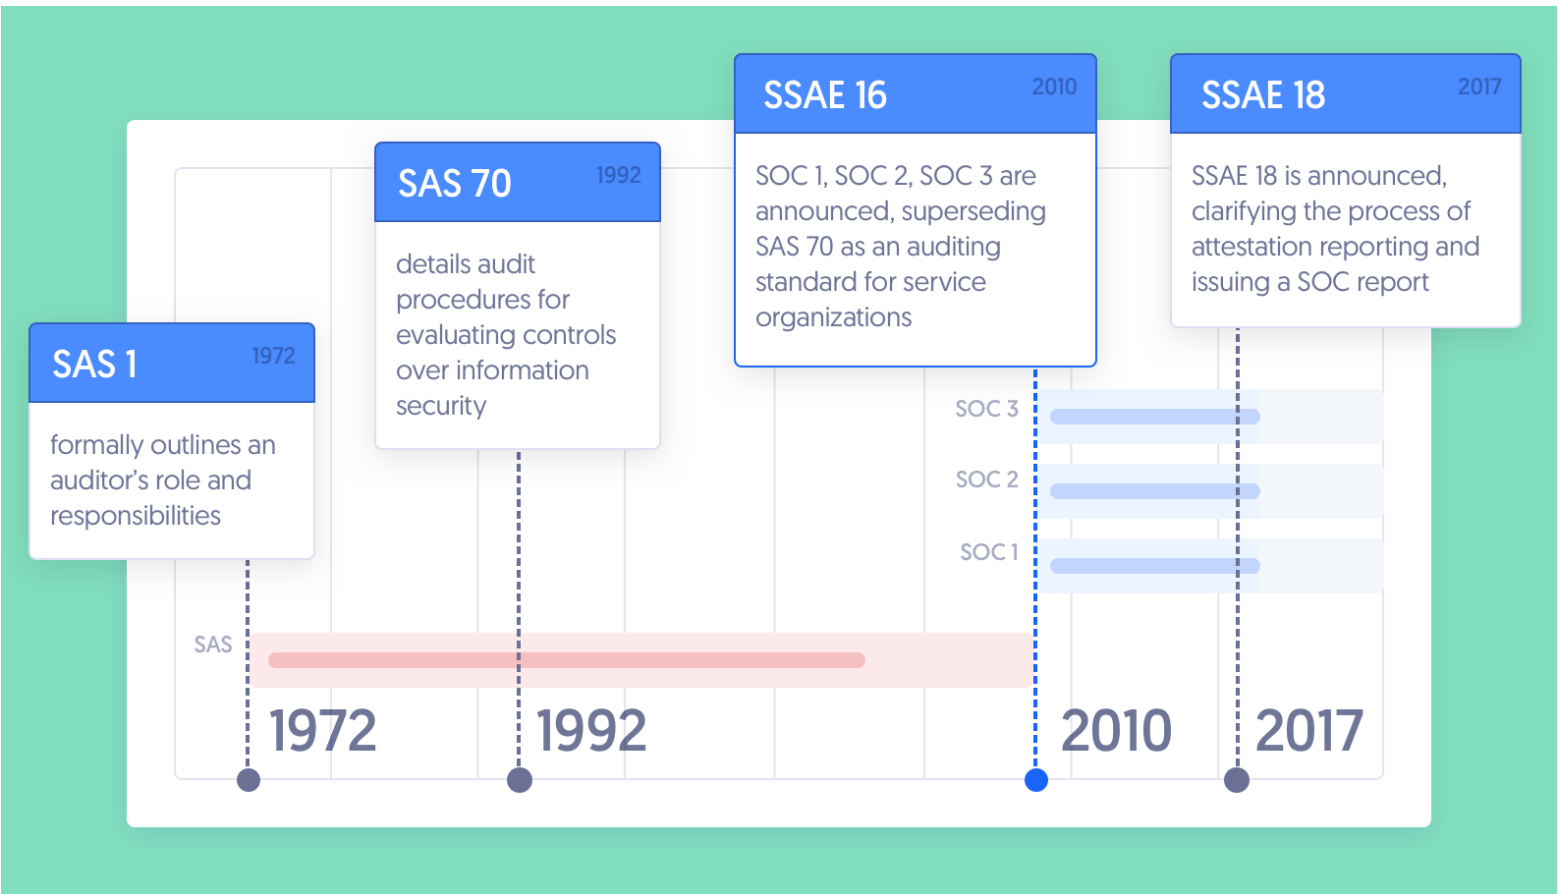
Durham is home to several recognized institutions of higher education, most notably [Duke University](#) and [North Carolina Central University](#). Durham is also a national leader in health-related activities, which are focused on the [Duke University Hospital](#) and many private companies. Duke and its [Duke University Health System](#) are the largest employers in the city. North Carolina Central University is a [historically black university](#) that is part of the [University of North Carolina](#) system. Together, the two universities make Durham one of the vertices of the [Research Triangle](#) area; central to this is the [Research Triangle Park](#)^[16] south of Durham, which encompasses an area of 11 square miles and is devoted to research facilities.

[Wikipedia 2023]



American Tobacco Campus, Durham

Below, we can see that there has been a long history of different frameworks:




Notice how the author is Sara Duffer. Below, we can see that Sara Duffer delivered a presentation with Chad Woolf:



This was a presentation at the Re:Invent conference, in 2016. We can see that the title of the presentation is

“How AWS Automates Compliance at Massive Scale using AWS Services”.





Screenshot from a YouTube video explaining how AWS Artifact works [AWS 2016]

From the Mouth of AWS Horse

Customer Compliance Guides now available on AWS Artifact - This serves as a handy excuse to remind you all that AWS Artifact exists, and you should absolutely use it rather than trying to white-knuckle responses to auditor questions about the AWS environment. I did that once and was asked to set up a tour of us-east-1.



Email from Corey Quinn on July 6th 2023

AWS Security Blog

Customer Compliance Guides now available on AWS Artifact

by Kevin Donohue and Travis Goldbach, | on 23 JUN 2023 | in [Announcements](#), [Foundational \(100\)](#), [Security, Identity, & Compliance](#) | [Permalink](#) | [Comments](#) | [Share](#)

[Amazon Web Services \(AWS\)](#) has released Customer Compliance Guides (CCGs) to support customers, partners, and auditors in their understanding of how compliance requirements from leading frameworks map to AWS service security recommendations. CCGs cover 100+ services and features offering security guidance mapped to 10 different compliance frameworks. Customers can select any of the available frameworks and services to see a consolidated summary of recommendations that are mapped to security control requirements.

CCGs summarize key details from [public AWS user guides](#) and map them to related security topics and control requirements. CCGs don't cover compliance topics such as physical and maintenance controls, or organization-specific requirements such as policies and human resources controls. This makes the guides lightweight and focused only on the unique security considerations for AWS services.

AWS Artifact launches email notifications - If you care about compliance reports, turn this on and save yourself some work. But you don't care about compliance reports; you just have to pretend to so you don't get fired. So turn this on and then ignore it.

Email from Corey Quinn on 14th August 2023

Navigating GDPR Compliance on AWS

AWS Whitepaper



The above whitepaper was first published in
November 2017

AWS Compliance Program

AWS continually maintains a high bar for security and compliance across all of our global operations. Security has always been our highest priority – truly "job zero." AWS regularly undergoes independent third-party attestation audits to provide assurance that control activities are operating as intended. More specifically, AWS is audited against a variety of global and regional security frameworks dependent on region and industry. Currently, AWS participates in over 50 different audit programs.

The results of these audits are documented by the assessing body and made available for all AWS customers through [AWS Artifact](#). AWS Artifact is a no-cost, self-service portal for on-demand access to AWS compliance reports. When new reports are released, they are made available in AWS Artifact, allowing customers to continuously monitor the security and compliance of AWS with immediate access to new reports.

Customers can take advantage of internationally recognized certifications and accreditations, demonstrating compliance with rigorous international standards, such as ISO 27017 for cloud security, ISO 27018 for cloud privacy, SOC 1, SOC 2 and SOC 3, PCI DSS Level 1 and others. AWS also helps customers meet local security standards such as BSI's Common Cloud Computing Controls Catalogue (C5), a German government-backed attestation.

For more detailed information about the AWS certification programs, reports, and third-party attestations, see [AWS Compliance Programs](#). For service-specific information, see [AWS Services in Scope](#).

US vendor accused of violating GDPR by reputation-scoring EU citizens

TeleSign and Belgian parent did almost everything wrong, alleges Max Schrems

Brandon Vigliarolo

Fri 23 Jun 2023 // 18:00 UTC

A US-based fraud prevention company is in hot water over allegations it not only collected data from millions of EU citizens and processed it using automated tools without their knowledge, but that it did so in the United States, all in violation of the EU's data protection rules.

The [complaint](#) was filed by Austrian privacy advocacy group noyb, helmed by lawyer [Max Schrems](#), and it doesn't pull any punches in its claims that TeleSign, through its former Belgian parent company BICS, secretly collected data on cellphone users around the world.



Monday, June 26, 2023

Good Morning!

Amazon did not have a great week, regulatorily speaking. The FTC comment period about [the business of cloud computing](#) ended, their warehouse practices are now the [focus of a senate probe](#), the FTC is suing Amazon for its [Prime enrollment dark patterns](#), their iRobot acquisition is now [the subject of an EU investigation](#), and the launch of Amazon Clinic is [being delayed after the senate asked some hard questions](#). Ever think that maybe, just maybe, your company should have a strategy that's slightly more sophisticated than "yes?"

Glossary

ISO –

CCG – Customer Compliance Guides

Bibliography

- I. Official
- II. Unofficial
- III. Critical
- IV. General

I. Official

[Amazon 2023]

The INFORM Consumers Act takes effect on June 27. Here's how Amazon is protecting our customers and sellers from bad actors. Courtesy of *Last Week in AWS*. Available at:
<<https://www.aboutamazon.com/news/policy-news-views/inform-consumers-act-takes-effect-on-june-27>>

[Duffer 2016]

Duffer, Sara (2016). Introducing AWS Artifact: Speeding Access to Compliance Reports. *AWS Security Blog*. Dec 7th 2016. Available at:
<<https://aws.amazon.com/blogs/security/introducing-aws-artifact-speeding-access-to-compliance-reports/>>

[AWS 2016]

Learn How to Use AWS Artifact to Retrieve Your Compliance Reports. YouTube Channel: Amazon Web Services. Available at:
<https://www.youtube.com/watch?v=O307zdIOMWo&ab_channel=AmazonWebServices>

[Woolf 2016]

EU Compliance Update. Available at:
<<https://aws.amazon.com/blogs/security/eu-compliance-update/>>

[Donohue 2023]

Donohue, Kevin and Travis Goldbach (2023). Customer Compliance Guides now available on AWS Artifact. Available at:
<https://aws.amazon.com/blogs/security/customer-compliance-guides-now-available-on-aws-artifact/?ck_subscriber_id=1560524742>

II. Unofficial

[Shah 2020]

AWS Tutorial – AWS Artifact – Reports and Agreements. YouTube Channel: NamrataHShah. Uploaded 12th Feb 2020. Available at: https://www.youtube.com/watch?v=lCUATy4emgY&ab_channel=NamrataHShah

III. Critical

IV. General

[Sherinksy 2010]

Sherinksy, Judith M. (2010). Replacing SAS-70. *Journal of Accountancy*, Aug 1st 2010. Available at: <https://www.journalofaccountancy.com/issues/2010/aug/20103009.html>

[Wikipedia 2023]

Durham, North Carolina. Available at: https://en.wikipedia.org/wiki/Durham,_North_Carolina

AWS Managed Services

Bibliography

- I. Official
- II. Unofficial
- III. Critical
- IV. General

- I. Official
- II. Unofficial
- III. Critical
- IV. General