



Tutorials Dojo Study Guide



TABLE OF CONTENTS

| INTRODUCTION | 7 |
|--|----|
| AWS CERTIFIED DEVOPS ENGINEER PROFESSIONAL EXAM OVERVIEW | 8 |
| Exam Details | 9 |
| Exam Domains | 10 |
| Exam Domain I: SDLC Automation | 11 |
| Exam Domain II: Configuration Management and IaC | 13 |
| Exam Domain III: Resilient Cloud Solutions | 14 |
| Exam Domain IV: Monitoring and Logging | 16 |
| Exam Domain V: Incident and Event Response | 18 |
| Exam Domain VI: Security and Compliance | 19 |
| Old DOP-C01 vs the New DOP-C02 Exam Version | 21 |
| Exam Scoring System | 22 |
| Related Exam Topics | 23 |
| Excluded Exam Topics | 26 |
| Exam Benefits | 26 |
| AWS CERTIFIED DEVOPS ENGINEER PROFESSIONAL EXAM - STUDY GUIDE AND TIPS | 27 |
| Study Materials | 27 |
| AWS Services to Focus On | 29 |
| Common Exam Scenarios | 30 |
| Validate Your Knowledge | 36 |
| Sample Question 1 | 37 |
| Sample Question 2 | 39 |
| Domain 1: Software Development Life Cycle (SDLC) Automation | 43 |
| Overview | 44 |
| What is DevOps? | 45 |
| A Brief History of the DevOps Exam in AWS | 47 |
| Why Automate? | 48 |
| Types of Blue Green Deployment via ELB, Route 53, Elastic Beanstalk | 49 |
| AWS Lambda Function Alias Traffic Shifting | 54 |
| Basic Blue/Green Deployment using Route 53 | 57 |
| Using a Cloned Stack in OpsWorks for Implementing Blue/Green Deployment | 59 |
| AWSCode Commit Full Access, AWSCode Commit Power User, AWSCode Commit Read Only-Permissions and the property of the property | 62 |
| Lifecycle Event Hook Availability (CodeDeploy Concept) | 63 |
| Automatically Run CodeBuild Tests After a Developer Creates a CodeCommit Pull Request | 66 |



| Managing Artifacts in AWS CodeBuild and CodePipeline | 69 |
|---|-----|
| DynamoDB – Fetch vs Projected Attributes | 74 |
| CodeBuild with CloudWatch Logs, Metrics, and Alarms | 76 |
| CodeDeploy with CloudWatch Logs, Metrics, and Alarms | 84 |
| CodePipeline and CloudWatch Events Integration | 93 |
| CodeDeploy - Linear, Canary and All-at-Once (Pre-defined Deployments) | 101 |
| Elastic Beanstalk - Deployment Policies and Settings | 103 |
| Domain 2: Configuration Management and Infrastructure-as-Code | 105 |
| Overview | 106 |
| What is Configuration Management? | 107 |
| What is Infrastructure-as-Code (IaC)? | 109 |
| CloudFormation Cross-Stack Reference | 110 |
| Lambda Function Artifact from S3 or CloudFormation Inline | 112 |
| AutoScalingReplacingUpdate vs AutoScalingRollingUpdate Policy | 115 |
| Time-Based vs Load-Based Instance | 117 |
| Discovery Agent vs Discovery Connector | 120 |
| CloudFormation Template for ECS, Auto Scaling and ALB | 121 |
| Domain 3: Resilient Cloud Solutions | 123 |
| Overview | 124 |
| High Availability vs. Fault Tolerance | 125 |
| Multi-AZ vs Multi-Region Architecture | 127 |
| Disaster Recovery Objectives | 128 |
| Amazon Route 53 Routing Policies | 129 |
| Amazon RDS Disaster Recovery Strategies | 131 |
| Auto Scaling Group with MinSize = 1 and MaxSize = 1 | 134 |
| Auto Scaling Lifecycle Hooks | 137 |
| Amazon EKS Auto Scaling | 139 |
| Amazon EKS Networking | 144 |
| Automated Patching for Hybrid Environments | 149 |
| Amazon EC2 Image Builder | 154 |
| Automated RDS Read Replica Promotion | 162 |
| Amazon Lookout for Metrics | 166 |
| Amazon S3 Object tagging and Access control | 172 |
| Amazon Storage Gateway - S3 File Gateway | 176 |
| Amazon CodeGuru Reviewer | 180 |
| Amazon CodeGuru Profiler | 185 |
| Domain 4: Monitoring and Logging | 189 |



| Overview | 190 |
|---|--------|
| | 190 |
| AWS Config Multi-Account Multi-Region Data Aggregation Consolidating VPC Flow Logs From Multiple Sources | 191 |
| Consolidating CloudTrail Log Files from Multiple Sources | 194 |
| | |
| Ensuring the Integrity of the CloudTrail Log Files | 196 |
| Fetching Application Logs from Amazon EC2, ECS and On-premises Servers | 197 |
| CloudWatch Logs Agent to CloudWatch Logs Subscription | 199 |
| Monitoring Service Limits with Trusted Advisor | 201 |
| Domain 5: Incident and Event Response | 205 |
| Overview | 206 |
| Incident and Event Response Management on AWS | 207 |
| Amazon S3 Event Notifications | 209 |
| Amazon RDS Event Notifications | 211 |
| AWS_RISK_CREDENTIALS_EXPOSED Event | 212 |
| AWS-Scheduled Maintenance Notification to Slack Channel | 216 |
| Using AWS Health API and CloudWatch Events for Monitoring AWS-Scheduled Deployments/Change | es 219 |
| Monitoring Amazon EC2 Auto Scaling Events | 220 |
| Monitoring Amazon S3 Data Events in AWS CloudTrail | 222 |
| AWS CodePipeline Event Patterns | 223 |
| Monitoring Deployments in AWS CodeDeploy | 226 |
| Orchestrating Events in AWS CodePipeline | 227 |
| Monitoring OpsWorks Auto-Healing Events | 229 |
| Domain 6: Security and Compliance | 230 |
| Overview | 231 |
| Management and Governance on AWS | 232 |
| AWS CodeBuild Configuration Best Practices | 236 |
| AWS CodeCommit Managed Policies with Customized Permissions | 238 |
| S3 Bucket Policy to Only Allow HTTPS Requests | 240 |
| Secrets Manager vs. Systems Manager Parameter Store | 242 |
| AWS Managed Policy | 244 |
| Using Systems Manager Automation to create AMIs | 247 |
| AWS SSM Session Manager to Record Sessions on your Instances | 249 |
| AWS Systems Manager Inventory | 253 |
| Systems Manager Patch Manager and Maintenance Windows | 257 |
| RELATED AWS SERVICES TO FOCUS | 262 |
| | 263 |
| Amazon CuickSight | |
| Amazon QuickSight | 263 |



| Amazon AppFlow | 263 |
|---|-----|
| Amazon Eventbridge | 264 |
| AWS App Runner | 264 |
| AWS App2Container | 264 |
| AWS Copilot | 264 |
| Amazon EKS Deployment Options | 265 |
| Amazon EKS on AWS Outposts | 265 |
| Amazon EKS Distro | 265 |
| Red Hat OpenShift Service on AWS (ROSA) | 266 |
| AWS Database Migration Service (AWS DMS) | 266 |
| Amazon DocumentDB | 266 |
| Amazon MemoryDB for Redis | 266 |
| Amazon ElastiCache | 267 |
| AWS Cloud Development Kit (AWS CDK) | 267 |
| AWS CloudShell | 268 |
| AWS CodeArtifact | 268 |
| AWS CodeStar | 269 |
| Amazon CodeGuru | 269 |
| AWS Fault Injection Simulator (AWS FIS) | 269 |
| AWS Control Tower | 270 |
| Customizations for AWS Control Tower (CfCT) | 270 |
| Amazon Lookout for Metrics | 271 |
| Amazon Compute Optimizer | 271 |
| Amazon Managed Grafana | 271 |
| Amazon Managed Service for Prometheus | 272 |
| Amazon FSx | 272 |
| AWS Backup | 273 |
| AWS Elastic Disaster Recovery | 273 |
| AWS Proton | 274 |
| AWS CloudHSM | 274 |
| AWS Network Firewall | 275 |
| Amazon Detective | 275 |
| AWS Virtual Private Network (AWS VPN) | 275 |
| AWS PrivateLink | 276 |
| AWS License Manager | 276 |
| AWS Service Health Dashboard | 276 |
| AWS Health API | 276 |



| AWS Personal Health Dashboard | 277 |
|---|-----|
| AWS Resilient Hub | 277 |
| AWS CHEAT SHEETS | 278 |
| AWS Compute Services | 279 |
| Amazon Elastic Compute Cloud (EC2) | 279 |
| Amazon Elastic Container Registry (ECR) | 283 |
| Amazon Elastic Container Service (ECS) | 284 |
| AWS Elastic Beanstalk | 287 |
| AWS Lambda | 288 |
| AWS Serverless Application Model (SAM) | 289 |
| AWS Storage Services | 290 |
| Amazon EBS | 290 |
| Amazon EFS | 290 |
| Amazon S3 | 290 |
| Amazon S3 Bucket Policies for VPC Endpoints | 293 |
| AWS Database Services | 294 |
| Amazon Aurora | 294 |
| Amazon DynamoDB | 295 |
| Lambda Integration With Amazon DynamoDB Streams | 297 |
| Amazon RDS | 300 |
| AWS Networking & Content Delivery | 302 |
| Amazon API Gateway | 302 |
| Amazon Route 53 | 304 |
| AWS Elastic Load Balancing (ELB) | 306 |
| AWS Transit Gateway | 307 |
| AWS Security & Identity Services | 308 |
| Amazon GuardDuty | 308 |
| Amazon Inspector | 309 |
| Amazon Macie | 311 |
| AWS Identity & Access Management (IAM) | 312 |
| AWS Key Management Service | 315 |
| AWS Secrets Manager | 317 |
| AWS Certificate Manager | 318 |
| AWS Directory Service | 321 |
| AWS Resource Access Manager | 326 |
| AWS Security Hub | 327 |
| AWS Shield | 329 |



| AWS WAF | 331 |
|---|------------|
| AWS Management Tools | 333 |
| AWS Organizations | 333 |
| Amazon CloudWatch | 335 |
| AWS Auto Scaling | 338 |
| AWS CloudFormation | 341 |
| AWS CloudTrail | 344 |
| AWS Config | 344 |
| AWS Health | 343 |
| | 347 |
| AWS OpsWorks | 351 |
| AWS Systems Manager AWS Trusted Advisor | |
| | 355 356 |
| AWS Service Catalog AWS Analytics Services | 359 |
| • | 359 |
| Amazon OpenSearch Service Amazon Kinesis | 361 |
| Amazon Athena | 365 |
| AWS Developer Tools | 367 |
| AWS CodeBuild | 367 |
| AWS CodeCommit | 369 |
| AWS CodeCommit AWS CodeDeploy | 372 |
| AWS CodePipeline | 372 |
| AWS X-Ray | 377 |
| AWS Application Services | 380 |
| Amazon SNS | 380 |
| AWS Step Functions | 382 |
| Comparison of AWS Services | 384 |
| AWS CloudTrail vs Amazon CloudWatch | 384 |
| CloudWatch Agent vs SSM Agent vs Custom Daemon Scripts | 385 |
| EC2 Container Services ECS vs Lambda | 386 |
| EC2 Instance Health Check vs ELB Health Check vs Auto Scaling and Custom Health Check | 387 |
| Elastic Beanstalk vs CloudFormation vs OpsWorks vs CodeDeploy | 388 |
| Service Control Policies vs IAM Policies | 390 |
| FINAL REMARKS AND TIPS | 390 |
| ABOUT THE AUTHORS | 392 |
| | 0,2 |



INTRODUCTION

As more companies build their DevOps practices, there will always be a growing demand for certified IT Professionals that can do agile software development, configuration management, task automation, and continuous integration/continuous delivery (CI/CD). This Study Guide and Cheat Sheets eBook for AWS Certified DevOps Engineer - Professional aims to equip you with the necessary knowledge and practical skill sets needed to pass the latest version of the AWS Certified DevOps Engineer - Professional exam.

This eBook contains the essential concepts, exam domains, exam tips, sample questions, cheat sheets, and other relevant information about the AWS Certified DevOps Engineer – Professional exam. This study guide begins with the presentation of the exam structure, giving you an insight into the question types, exam domains, scoring scheme, and the list of benefits you'll receive once you pass the exam. We used the official AWS exam guide to structure the contents of this guide, where each section discusses a particular exam domain. Various DevOps concepts, related AWS services, and technical implementations are covered to provide you with an idea of what to expect on the actual exam.

DevOps Exam Notes:

Don't forget to read the boxed "**exam tips**" (like this one) scattered throughout the eBook, as these are the key concepts that you will likely encounter on your test. After covering the six domains, we have added a bonus section containing a curated list of AWS Cheat Sheets to fast-track your review. The last part of this guide includes a collection of articles that compares two or more similar AWS services to supplement your knowledge.

The AWS Certified DevOps Engineer - Professional certification exam is a difficult test to pass; therefore, anyone who wants to take it must allocate ample time for review. The exam registration cost is not cheap, which is why we spent considerable time and effort to ensure that this study guide provides you with the essential and relevant knowledge to increase your chances of passing the DevOps exam.

** **Note:** This eBook is meant to be just a supplementary resource when preparing for the exam. We highly recommend working on <u>hands-on sessions</u> and <u>practice exams</u> to further expand your knowledge and improve your test-taking skills.



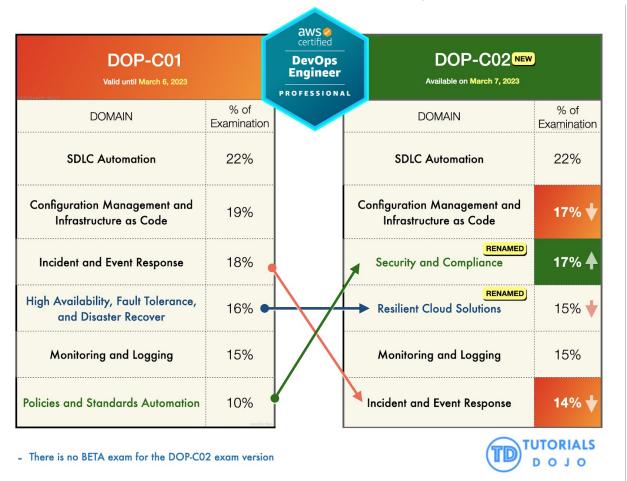
Old DOP-C01 vs the New DOP-C02 Exam Version

In terms of the number of exam domains, the old version of the AWS Certified DevOps Engineer Professional (DOP-C01) has the exact same number of domains as compared with the new DOP-C02 version. However, there are differences in terms of exam coverage and names for some domains.

The biggest exam domain is still the SDLC (Software Development Lifecycle) Automation domain which retains its 22% percent exam coverage. The same goes for the Monitoring and Logging domain, which still has 15 percent. This is followed by the Configuration Management and Infrastructure as Code (IaC) domain which is down to only 17% exam coverage from the previous 19% percent. The Incident and Event Response domain has a huge 4% decline as it only has 14% coverage coming from an 18% high on the previous version.

You can also notice that two exam domains have changed their name:

- The "High Availability, Fault Tolerance, and Disaster Recovery" domain has been renamed and is now called the "Resilient Cloud Solutions" domain.
- The "Policies and Standards Automation" domain is now "Security and Compliance"





The concept of resiliency is related to High Availability, Fault Tolerance, and Disaster Recovery. This is the primary reason why AWS renamed this lengthy domain as "Resilient Cloud Solutions" for brevity. From 16%, this exam domain has a slight decrease of coverage at 15% percent.

Security in AWS can be implemented through IAM Policies, Service Control Policies (SCPs), Bucket Policies, VPC Endpoint Policies, and other types of policies. The term "standards" is synonymous with the word "compliance" in the IT industry. The name of the Policies and Standards Automation exam domain was simplified and is now officially the Security and Compliance domain. It's interesting to note that on the previous DOP-C01 version, this domain has the lowest exam coverage at 10%, but now, it has become the second largest exam domain for DOP-C02 with 17% coverage.

As you can notice, the DevOps Pro exam has significantly included many security-related topics based on its new exam domain content distribution. This means that you have to focus on various security topics and security services offered by AWS.

Exam Scoring System

You can get a score from 100 to 1,000 with a minimum passing score of **750** when you take the DevOps Engineer Professional exam. AWS uses a scaled scoring model to equate scores across multiple exam types that may have different difficulty levels. The complete score report will be sent to you by email after a few days. Right after you complete the actual exam, you'll immediately see a pass or fail notification on the testing screen. A "Congratulations! You have successfully passed..." message will be shown if you pass the exam.

Individuals who unfortunately do not pass the AWS exam must wait 14 days before they are allowed to retake the exam. Fortunately, there is no hard limit on exam attempts until you pass the exam. Take note that on each attempt, the full registration price of the AWS exam must be paid.

Within 5 business days of completing your exam, your AWS Certification Account will have a record of your complete exam results. The score report contains a table of your performance at each section/domain, which indicates whether you met the competency level required for these domains or not. AWS is using a compensatory scoring model, which means that you do not necessarily need to pass each and every individual section, only the overall examination. Each section has a specific score weighting that translates to the number of questions; hence, some sections have more questions than others. The Score Performance table highlights your strengths and weaknesses that you need to improve on.



Related Exam Topics

The new AWS Certified DevOps Engineer – Professional exam (DOP-C02) is focused on the various tools, services, and knowledge areas that revolve around DevOps in AWS. The official exam guide provides a list of AWS services, general tools, and technologies that are grouped according to their primary functions. Keep in mind that even though some of these topics will likely be covered more than others on the exam, the placement or order of these exam topics/ AWS services in this list is not an indication of any relative weight or importance.

The relevant exam topics that you should be familiar with on your upcoming DOP-C02 exam are:

- Application Deployment
- Application Integration
- Application pipelines
- Automation
- Code repository best practices
- Cost optimization
- Deployment requirements
- Hybrid deployments
- IAM policies
- Metrics, monitoring, alarms, and logging
- Network ACL and security group design and implementation
- Operational best practices
- Rollback procedures



Here is the list of relevant AWS services that are covered in the AWS Certified DevOps Engineer – Professional (DOP-C02) exam based on the official exam guide. You must focus on these AWS services and their respective features for your upcoming test:

Analytics:

- Amazon Athena
- Amazon Elastic MapReduce (Amazon EMR)
- Amazon Kinesis Data Firehose
- Amazon Kinesis Data Streams
- Amazon OpenSearch Service
- Amazon QuickSight

Compute:

- AWS App Runner
- Amazon EC2
- Amazon EC2 Auto Scaling
- EC2 Image Builder
- AWS Elastic Beanstalk
- AWS Serverless Application Repository

Containers:

- AWS App2Container
- AWS Copilot
- Amazon Elastic Container Registry (Amazon ECR)
- Amazon Elastic Container Service (Amazon ECS)
- Amazon Elastic Kubernetes Service (Amazon EKS)
- Amazon EKS Distro
- AWS Fargate
- Red Hat OpenShift Service on AWS (ROSA)

Database:

- Amazon Aurora
- Amazon Aurora Serverless v2
- AWS Database Migration Service (AWS DMS)
- Amazon DocumentDB (with MongoDB compatibility)
- Amazon DynamoDB
- Amazon ElastiCache
- Amazon MemoryDB for Redis
- Amazon RDS
- Amazon Redshift

Management and Governance:

- AWS Auto Scaling
- AWS CloudFormation
- AWS CloudTrail
- Amazon CloudWatch
- Amazon CloudWatch Logs
- AWS Compute Optimizer
- AWS Config
- AWS Control Tower
- AWS Health
- AWS License Manager
- Amazon Managed Grafana
- Amazon Managed Service for Prometheus
- AWS OpsWorks
- AWS Organizations
- AWS Personal Health Dashboard
- AWS Proton
- AWS Resilience Hub
- AWS Service Catalog
- AWS Systems Manager
- AWS Trusted Advisor

Networking and Content Delivery:

- Amazon API Gateway
- AWS Client VPN
- Amazon CloudFront
- Elastic Load Balancing (ELB)
- AWS PrivateLink
- AWS Site-to-Site VPN
- Amazon Route 53
- AWS Transit Gateway
- Amazon VPC

Serverless:

- AWS Lambda
- AWS Serverless Application Model (AWS SAM)
- Amazon Simple Notification Service (Amazon SNS)
- Amazon Simple Queue Service (Amazon SQS)
- AWS Step Functions
- Application Integration:
- Amazon AppFlow
- Amazon EventBridge (Amazon CloudWatch Events)

Developer Tools:

- AWS Cloud Development Kit (AWS CDK)
- AWS CloudShell
- AWS CodeArtifact
- AWS CodeBuild
- AWS CodeCommit
- AWS CodeDeploy
- Amazon CodeGuru
- AWS CodePipeline
- AWS CodeStar
- AWS Command Line Interface (AWS CLI)
- AWS Fault Injection Simulator
- AWS SDKs and Tools
- AWS X-Ray

Security, Identity, and Compliance:

- AWS Certificate Manager (ACM)
- AWS CloudHSM
- Amazon Cognito
- Amazon Detective
- AWS Directory Service
- Amazon GuardDuty
- AWS Identity and Access Management (IAM)
- Amazon Inspector
- AWS Key Management Service (AWS KMS)
- Amazon Macie
- AWS Network Firewall
- AWS Resource Access Manager (AWS RAM)
- AWS Secrets Manager
- AWS Security Hub
- AWS Security Token Service (AWS STS)
- AWS Shield
- AWS IAM Identity Center (AWS Single Sign-On)
- AWS WAF

Storage:

- AWS Backup
- Amazon Elastic Block Store (Amazon EBS)
- AWS Elastic Disaster Recovery (CloudEndure Disaster Recovery)
- Amazon Elastic File System (Amazon EFS)
- Amazon FSx for Lustre
- Amazon FSx for NetApp ONTAP
- Amazon FSx for OpenZFS
- Amazon FSx for Windows File Server
- Amazon S3
- Amazon S3 Glacier
- AWS Storage Gateway



Excluded Exam Topics

Usually, the official exam guides provide a list of both the relevant and irrelevant AWS services for the exam. This is not the case for the latest AWS Certified DevOps Engineer Professional DOP-C02 exam. The official DOP-C02 exam guide doesn't come with a list of exam topics that are not in scope for this certification test. However, we can deduce the out-of-scope AWS topics by comparing the exam guide for the AWS Certified DevOps Engineer – Professional DOP-C02 exam.

Just a friendly reminder that the following AWS services and features do not represent each and every AWS offering that is excluded from the DOP-C02 exam content. This list is only a hint of what topics are not covered on the AWS Certified DevOps Engineer — Professional exam, which you should not focus on:

- Machine Learning
- Internet-of-Things (IoT)
- Frontend development for mobile apps
- 12-factor app methodology
- AWS Direct Connect

Exam Benefits

If you successfully passed any AWS exam, you will be eligible for the following benefits:

- **Exam Discount** You'll get a 50% discount voucher that you can apply for your recertification or any other exam you plan to pursue. To access your discount voucher code, go to the "Benefits" section of your AWS Certification Account, and apply the voucher when you register for your next exam.
- AWS Certified Store All AWS-certified professionals will be given access to exclusive AWS Certified merchandise. You can get your store access from the "Benefits" section of your AWS Certification Account.
- Certification Digital Badges You can showcase your achievements to your colleagues and employers
 with digital badges on your email signatures, Linkedin profile, or on your social media accounts. You
 can also show your Digital Badge to gain exclusive access to Certification Lounges at AWS re:Invent,
 regional Appreciation Receptions, and select AWS Summit events. To view your badges, simply go to
 the "Digital Badges" section of your AWS Certification Account.
- Eligibility to join AWS IQ With the AWS IQ program, you can monetize your AWS skills online by
 providing hands-on assistance to customers around the globe. AWS IQ will help you stay sharp and be
 well-versed in various AWS technologies. You can work in the comforts of your home and decide when
 or where you want to work. Interested individuals must have an Associate, Professional, or Specialty
 AWS Certification and be over 18 of age.



AWS CERTIFIED DEVOPS ENGINEER PROFESSIONAL EXAM - STUDY GUIDE AND TIPS

This certification is the pinnacle of your DevOps career in AWS. The AWS Certified DevOps Engineer Professional (or AWS DevOps Pro) is the advanced certification of both <u>AWS SysOps Administrator Associate</u> and <u>AWS Developer Associate</u>. This is similar to how the AWS Solutions Architect Professional role is a more advanced version of the AWS Solutions Architect Associate.

Generally, AWS recommends that you first take (and pass) both AWS SysOps Administrator Associate and AWS Developer Associate certification exams before taking on this certification. Previously, it was a prerequisite that you obtain the associate level certifications before you are allowed to go for the professional level. Last October 2018, AWS removed this ruling to provide customers with a more flexible approach to the certifications.

Study Materials

The <u>FREE AWS Exam Readiness course</u>, <u>official AWS sample questions</u>, Whitepapers, FAQs, AWS Documentation, Re:Invent videos, forums, labs, <u>AWS cheat sheets</u>, <u>practice tests</u>, and personal experiences are what you will need to pass the exam. Since the DevOps Pro is one of the most difficult AWS certification exams out there, you have to prepare yourself with every study material you can get your hands on. If you need a review on the fundamentals of AWS DevOps, then do check out our review guides for the <u>AWS SysOps</u> <u>Administrator Associate</u> and <u>AWS Developer Associate</u> certification exams. Also, visit this <u>AWS exam blueprint</u> to learn more details about your certification exam.

For virtual classes, you can attend the <u>DevOps Engineering on AWS</u> and <u>Systems Operations on AWS</u> classes since they will teach you concepts and practices that are expected to be in your exam.

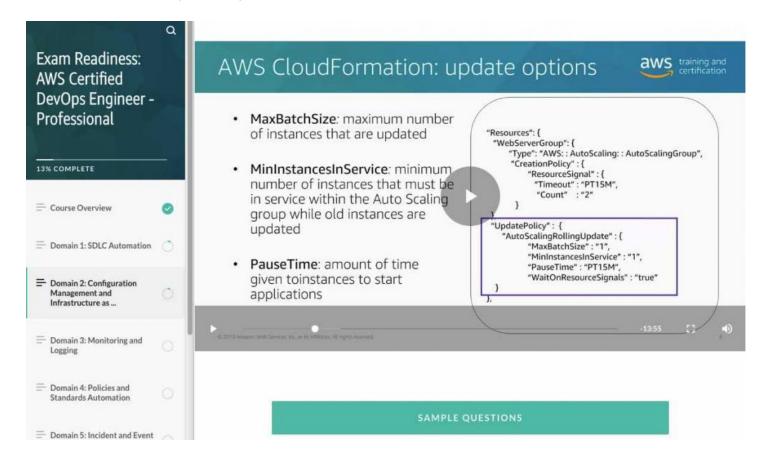
For whitepapers, focus on the following:

- 1. Running Containerized Microservices on AWS
- 2. Microservices on AWS
- 3. Infrastructure as Code
- 4. Introduction to DevOps
- 5. Practicing Continuous Integration and Continuous Delivery on AWS
- 6. Jenkins on AWS
- 7. Blue/Green Deployments on AWS whitepaper
- 8. Development and Test on AWS

Almost all online training you need can be found on the AWS web page. One digital course that you should check out is the Exam Readiness: AWS Certified DevOps Engineer - Professional course. This digital course



contains lectures on the different domains of your exam, and they also provide a short quiz right after each lecture to validate what you have just learned.



Lastly, do not forget to study the AWS CLI, SDKs, and APIs. Since DevOps Pro is also an advanced certification for Developer Associate, you need to have knowledge of programming and scripting in AWS. Go through the AWS documentation to review the syntax of the CloudFormation template, Serverless Application Model template, CodeBuild buildspec, CodeDeploy appspec, and IAM Policy.

Also, check out this article: Top 5 FREE AWS Review Materials.



AWS Services to Focus On

Since this exam is a professional level one, you should already have a deep understanding of the AWS services listed under our SysOps Administrator Associate and Developer Associate review guides. In addition, you should familiarize yourself with the following services since they commonly come up in the DevOps Pro exam:

- 1. AWS CloudFormation
- 2. AWS Lambda
- 3. Amazon CloudWatch
- 4. Amazon EventBridge
- 5. Amazon CloudWatch Alarms
- 6. AWS CodePipeline
- 7. AWS CodeDeploy
- 8. AWS CodeBuild
- 9. AWS CodeCommit
- 10. AWS Config
- 11. AWS Systems Manager
- 12. Amazon ECS
- 13. Amazon Elastic Beanstalk
- 14. AWS CloudTrail
- 15. AWS OpsWorks
- 16. AWS Trusted Advisor

The FAQs provide a good summary for each service, however, the AWS documentation contains more detailed information that you'll need to study. These details will be the deciding factor in determining the correct choice from the incorrect choices in your exam. To supplement your review of the services, we recommend that you take a look at <u>Tutorials Dojo's AWS Cheat Sheets</u>. Their contents are well-written and straight to the point, which will help reduce the time spent going through FAQs and documentation.



Common Exam Scenarios

| Scenario | Solution |
|--|--|
| Software Development and L | ifecycle (SDLC) Automation |
| An Elastic Beanstalk application must not have any downtime during deployment and requires an easy rollback to the previous version if an issue occurs. | Set up Blue/Green deployment, deploy a new version on a separate environment then swap environment URLs on Elastic Beanstalk. |
| A new version of an AWS Lambda application is ready to be deployed, and the deployment should not cause any downtime. A quick rollback to the previous Lambda version must be available. | Publish a new version of the Lambda function. After testing, use the production Lambda Alias to point to this new version. |
| In an AWS Lambda application deployment, only 10% of the incoming traffic should be routed to the new version to verify the changes before eventually allowing all production traffic. | Set up Canary deployment for AWS Lambda. Create a Lambda Alias pointed to the new Version. Set Weighted Alias value for this Alias as 10%. |
| An application hosted in Amazon EC2 instances behind an Application Load Balancer. You must provide a safe way to upgrade the version on Production and allow easy rollback to the previous version. | Launch the application in Amazon EC2 that runs the new version with an Application Load Balancer (ALB) in front. Use Route 53 to change the ALB A-record Alias to the new ALB URL. Rollback by changing the A-record Alias to the old ALB. |
| An AWS OpsWorks application needs to safely deploy its new version on the production environment. You are tasked to prepare a rollback process in case of unexpected behavior. | Clone the OpsWorks Stack. Test it with the new URL of the cloned environment. Update the Route 53 record to point to the new version. |
| A development team needs full access to AWS CodeCommit, but they should not be able to create/delete repositories. | Assign the developers with the AWSCodeCommitPowerUser IAM policy |
| During the deployment, you need to run custom actions before deploying the new version of the application using AWS CodeDeploy. | Add lifecycle hook action BeforeAllowTraffic |

| You need to run custom verification actions after the new version is deployed using AWS CodeDeploy. | Add lifecycle hook action AfterAllowTraffic |
|---|---|
| You need to set up AWS CodeBuild to automatically run after a pull request has been successfully merged using AWS CodeCommit | Create CloudWatch Events rule to detect pull requests and action set to trigger CodeBuild Project. Use AWS Lambda to update the pull request with the result of the project Build |
| You need to use AWS CodeBuild to create artifact and automatically deploy the new application version | Set CodeBuild to save artifact to S3 bucket. Use CodePipeline to deploy using CodeDeploy and set the build artifact from the CodeBuild output. |
| You need to upload the AWS CodeBuild artifact to Amazon S3 | S3 bucket needs to have versioning and encryption enabled. |
| You need to review AWS CodeBuild Logs and have an alarm notification for build results on Slack | Send AWS CodeBuild logs to CloudWatch Log group. Create CloudWatch Events rule to detect the result of your build and target a Lambda function to send results to the Slack channel (or SNS notification) |
| Need to get a Slack notification for the status of the application deployments on AWS CodeDeploy | Create Amazon EventBridge rule to detect the result of CodeDeploy job and target a notification to Amazon SNS or a Lambda function to send results to Slack channel |
| Need to run an AWS CodePipeline every day for updating the development progress status | Create Amazon EventBridge rule to run on schedule every day and set a target to the AWS CodePipeline ARN |
| Automate deployment of a Lambda function and test for only 10% of traffic for 10 minutes before allowing 100% traffic flow. | Use CodeDeploy and select deployment configuration CodeDeployDefault.LambdaCanary10Percent10M inutes |
| Deployment of Elastic Beanstalk application with absolutely no downtime. The solution must maintain full compute capacity during deployment to avoid service degradation. | Choose the "Rolling with additional Batch" deployment policy in Elastic Beanstalk |

| Deployment of Elastic Beanstalk application where the new version must not be mixed with the current version. | Choose the "Immutable deployments" deployment policy in Elastic Beanstalk |
|---|--|
| Configuration Management and | d Infrastructure-as-Code (IaC) |
| The resources on the parent CloudFormation stack needs to be referenced by other nested CloudFormation stacks | Use Export on the Output field of the main CloudFormation stack and use Fn::ImportValue function to import the value on the other stacks |
| On which part of the CloudFormation template should you define the artifact zip file on the S3 bucket? | The artifact file is defined on the AWS::Lambda::Function code resource block |
| Need to define the AWS Lambda function inline in the CloudFormation template | On the AWS::Lambda::Function code resource block, the inline function must be enclosed inside the ZipFile section. |
| Use CloudFormation to update Auto Scaling Group and only terminate the old instances when the newly launched instances become fully operational | Set AutoScalingReplacingUpdate: WillReplace property to TRUE to have CloudFormation retain the old ASG until the instances on the new ASG are healthy. |
| You need to scale down the EC2 instances at night when there is low traffic using OpsWorks. | Create <i>Time-based</i> instances for automatic scaling of predictable workload. |
| Can't install an agent on on-premises servers but need to collect information for migration | Deploy the Agentless Discovery Connector VM on your on-premises data center to collect information. |
| Syntax for CloudFormation with an Amazon ECS cluster with ALB | Use the AWS::ECS::Service element for the ECS Cluster, AWS::ECS::TaskDefinition element for the ECS Task Definitions, and the AWS::ElasticLoadBalancingV2::LoadBalancer element for the ALB. |
| Monitoring a | nd Logging |
| Need to centralize audit and collect configuration settings on all regions of multiple accounts | Setup an Aggregator on AWS Config. |

| Consolidate CloudTrail log files from multiple AWS accounts | Create a central S3 bucket with bucket policy to grant cross-account permission. Set this as destination bucket on the CloudTrail of the other AWS accounts. |
|--|--|
| Ensure that CloudTrail logs on the S3 bucket are protected and cannot be tampered with. | Enable Log File Validation on CloudTrail settings |
| Need to collect/investigate application logs from EC2 or on-premises server | Install CloudWatch Logs Agent to send the logs to CloudWatch Logs for storage and viewing. |
| Need to review logs from running ECS Fargate tasks | Enable awslogs log driver on the Task Definition and add the required logConfiguration parameter. |
| Need to run real-time analysis for collected application logs | Send logs to CloudWatch Logs, create a Lambda subscription filter, Elasticsearch subscription filter, or Kinesis stream filter. |
| Need to be automatically notified if you are reaching the limit of running EC2 instances or limit of Auto Scaling Groups | Track service limits with Trusted Advisor on CloudWatch Alarms using the ServiceLimitUsage metric. |
| Security and Compliance | |
| Need to secure the buildspec.yml file, which contains the AWS keys and database password stored in plaintext. | Store these values as encrypted parameter on SSM Parameter Store |
| Using default IAM policies for AWSCodeCommitPowerUser but must be limited to a specific repository only | Attach additional policy with Deny rule and custom condition if it does not match the specific repository or branch |
| You need to secure an S3 bucket by ensuring that only HTTPS requests are allowed for compliance purposes. | Create an S3 bucket policy that Deny if checks for condition aws:SecureTransport is false |
| Need to store a secret, database password, or variable, in the most cost-effective solution | Store the variable on SSM Parameter Store and enable encryption |
| Need to generate a secret password and have it rotated automatically at regular intervals | Store the secret on AWS Secrets Manager and enable key rotation. |

| Several team members with designated roles need to be granted permission to use AWS resources | Assign AWS-managed policies on the IAM accounts such as, ReadOnlyAccess, AdministratorAccess, PowerUserAccess | |
|---|--|--|
| Apply latest patches on EC2 and automatically create an AMI | Use Systems Manager automation to execute an Automation Document that installs OS patches and creates a new AMI. | |
| Need to have a secure SSH connection to EC2 instances and have a record of all commands executed during the session | Install SSM Agent on EC2 and use SSM Session Manager for the SSH access. Send the session logs to S3 bucket or CloudWatch Logs for auditing and review. | |
| Ensure that the managed EC2 instances have the correct application version and patches installed. | Use SSM Inventory to have visibility of your managed instances and identify their current configurations. | |
| Apply custom patch baseline from a custom repository, and schedule patches to managed instances | Use SSM Patch Manager to define a custom patch baseline and schedule the application patches using SSM Maintenance Windows | |
| Incident and Event Response | | |
| Need to get a notification if somebody deletes files in your S3 bucket | Setup Amazon S3 Event Notifications to get notifications based on specified S3 events on a | |
| | particular bucket. | |
| Need to be notified when an RDS Multi-AZ failover happens | Setup Amazon RDS Event Notifications to detect specific events on RDS. | |
| | Setup Amazon RDS Event Notifications to detect | |

| Get notified of any AWS maintenance or events that may impact your EC2 or RDS instances | Create an Amazon EventBridge rule for detecting any events on AWS Health Service and send a message to an SNS topic or invoke a Lambda function. | |
|---|--|--|
| Monitor scaling events of your Amazon EC2 Auto Scaling Group, such as launching or terminating an EC2 instance. | Use Amazon EventBridge or CloudWatch Events for monitoring the Auto Scaling Service and monitor the EC2 Instance-Launch Successful and EC2 Instance-Terminate Successful events. | |
| View object-level actions of S3 buckets, such as upload or deletion of object in CloudTrail | Set up Data events on your CloudTrail trail to record object-level API activity on your S3 buckets. | |
| Execute a custom action if a specific CodePipeline stage has a FAILED status | Create CloudWatch Event rule to detect failed state on the CodePipeline service, and set a target to SNS topic for notification or invoke a Lambda function to perform custom action. | |
| Automatically rollback a deployment in AWS CodeDeploy when the number of healthy instances is lower than the minimum requirement. | On CodeDeploy, create a deployment alarm that is integrated with Amazon CloudWatch. Track the MinimumHealthyHosts metric for the threshold of EC2 instances and trigger the rollback if the alarm is breached. | |
| Need to complete QA testing before deploying a new version to the production environment | Add a Manual approval step on AWS CodePipeline, and instruct the QA team to approve the step before the pipeline can resume the deployment. | |
| Get notified for OpsWorks auto-healing events | Create an Amazon EventBridge rule for the OpsWorks Service to track the auto-healing events | |
| Resilient Cloud Solutions | | |
| Need to ensure that both the application and the database are running in the event that one Availability Zone becomes unavailable. | Deploy your application on multiple Availability Zones and set up your Amazon RDS database to use Multi-AZ Deployments. | |
| In the event of an AWS Region outage, you have to make sure that both your application and database will still be running to avoid any service outages. | Create a copy of your deployment on the backup AWS region. Set up an RDS Read-Replica on the backup region. | |

| Automatically switch traffic to the backup region when your primary AWS region fails | Set up Route 53 Failover routing policy with health check enabled on your primary region endpoint. |
|--|--|
| Need to ensure the availability of a legacy application running on a single EC2 instance | Set up an Auto Scaling Group with MinSize=1 and MaxSize=1 configuration to set a fixed count and ensure that it will be replaced when the instance becomes unhealthy |
| Ensure that every EC2 instance on an Auto Scaling group downloads the latest code first before being attached to a load balancer | Create an Auto Scaling Lifecycle hook and configure the Pending:Wait hook with the action to download all necessary packages. |
| Ensure that all EC2 instances on an Auto Scaling group upload all log files in the S3 bucket before | Use the Auto Scaling Lifecycle and configure the Terminating:Wait hook with the action to upload |

Validate Your Knowledge

After your review, you should take some practice tests to measure your preparedness for the real exam. AWS offers a sample practice test for free which you can find here.. You can also opt to buy the longer AWS sample practice test at aws.training and use the discount coupon you received from any previously taken certification exams. Be aware though, that the sample practice tests do not mimic the difficulty of the real DevOps Pro exam.

Therefore, we highly encourage using other mock exams such as our very own <u>AWS Certified DevOps Engineer Professional Practice Exam</u> course, which contains high-quality questions with complete explanations on correct and incorrect answers, visual images and diagrams, YouTube videos as needed, and also contains reference links to official AWS documentation as well as our cheat sheets and study guides. You can also pair our practice exams with our <u>AWS Certified DevOps Engineer Professional Exam Study Guide eBook</u> to further help in your exam preparations.



Sample Question 1

An application is hosted in an Auto Scaling group of Amazon EC2 instances with public IP addresses in a public subnet. The instances are configured with a user data script that fetches and installs the required system dependencies of the application from the Internet upon launch. A change was recently introduced to prohibit any Internet access from these instances to improve the security, but after its implementation, the instances could not get the external dependencies anymore. Upon investigation, all instances are properly running, but the hosted application is not starting up completely due to the incomplete installation.

Which of the following is the MOST secure solution to solve this issue and also ensure that the instances do not have public Internet access?

- 1. Download all of the external application dependencies from the public Internet and then store them in an S3 bucket. Set up a VPC endpoint for the S3 bucket and then assign an IAM instance profile to the instances in order to allow them to fetch the required dependencies from the bucket.
- 2. Deploy the Amazon EC2 instances in a private subnet and associate Elastic IP addresses on each of them. Run a custom shell script to disassociate the Elastic IP addresses after the application has been successfully installed and is running properly.
- 3. Use a NAT gateway to disallow any traffic to the VPC which originated from the public Internet. Deploy the Amazon EC2 instances to a private subnet then set the subnet's route table to use the NAT gateway as its default route.
- 4. Set up a brand new security group for the Amazon EC2 instances. Use a whitelist configuration to only allow outbound traffic to the site where all of the application dependencies are hosted. Delete the security group rule once the installation is complete. Use AWS Config to monitor the compliance.

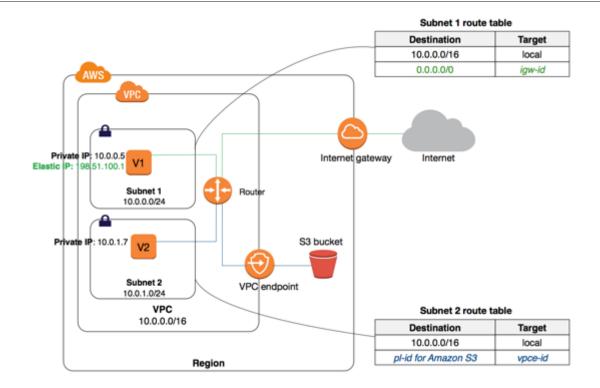
Correct Answer: 1

A VPC endpoint enables you to privately connect your VPC to supported AWS services and VPC endpoint services powered by PrivateLink without requiring an Internet gateway, NAT device, VPN connection, or AWS Direct Connect connection. Instances in your VPC do not require public IP addresses to communicate with resources in the service. Traffic between your VPC and the other service does not leave the Amazon network.

Endpoints are virtual devices. They are horizontally scaled, redundant, and highly available VPC components that allow communication between instances in your VPC and services without imposing availability risks or bandwidth constraints on your network traffic.

There are two types of VPC endpoints: *interface endpoints* and *gateway endpoints*. You can create the type of VPC endpoint required by the supported service. S3 and DynamoDB are using Gateway endpoints, while most of the services are using Interface endpoints.





You can use an S3 bucket to store the required dependencies and then set up a VPC Endpoint to allow your EC2 instances to access the data without having to traverse the public Internet.

Hence, the correct answer is the option that says: **Download all of the external application dependencies from** the public Internet and then store them to an S3 bucket. Set up a VPC endpoint for the S3 bucket and then assign an IAM instance profile to the instances in order to allow them to fetch the required dependencies from the bucket.

The option that says: Deploy the Amazon EC2 instances in a private subnet and associate Elastic IP addresses on each of them. Run a custom shell script to disassociate the Elastic IP addresses after the application has been successfully installed and is running properly is incorrect because it is possible that the custom shell script may fail and the disassociation of the Elastic IP addresses might not be fully implemented, which will allow the EC2 instances to access the Internet.

The option that says: Use a NAT gateway to disallow any traffic to the VPC which originated from the public Internet. Deploy the Amazon EC2 instances to a private subnet then set the subnet's route table to use the NAT gateway as its default route is incorrect because although a NAT Gateway can safeguard the instances from any incoming traffic that were initiated from the Internet, it still permits them to send outgoing requests externally.

The option that says: Set up a brand new security group for the Amazon EC2 instances. Use a whitelist configuration to only allow outbound traffic to the site where all of the application dependencies are hosted.



Delete the security group rule once the installation is complete. Use AWS Config to monitor the compliance is incorrect because this solution has a high operational overhead since the actions are done manually. This is susceptible to human error, such as in the event that the DevOps team forgets to delete the security group. The use of AWS Config will just monitor and inform you about the security violation, but it won't do anything to remediate the issue.

References:

https://docs.aws.amazon.com/vpc/latest/userguide/vpc-endpoints.html https://docs.aws.amazon.com/vpc/latest/userguide/vpce-gateway.html

Check out this Amazon VPC Cheat Sheet:

https://tutorialsdojo.com/amazon-vpc/

Sample Question 2

Due to the growth of its regional e-commerce website, the company has decided to expand its operations globally in the coming months ahead. The REST API web services of the app are currently running in an Auto Scaling group of EC2 instances across multiple Availability Zones behind an Application Load Balancer. For its database tier, the website is using a single Amazon Aurora MySQL database instance in the AWS Region where the company is based. The company wants to consolidate and store the data of its offerings into a single data source for its product catalog across all regions. For data privacy compliance, they need to ensure that the personal information of their users, as well as their purchases and financial data, are kept in their respective regions.

Which of the following options can meet the above requirements and entails the LEAST amount of change to the application?

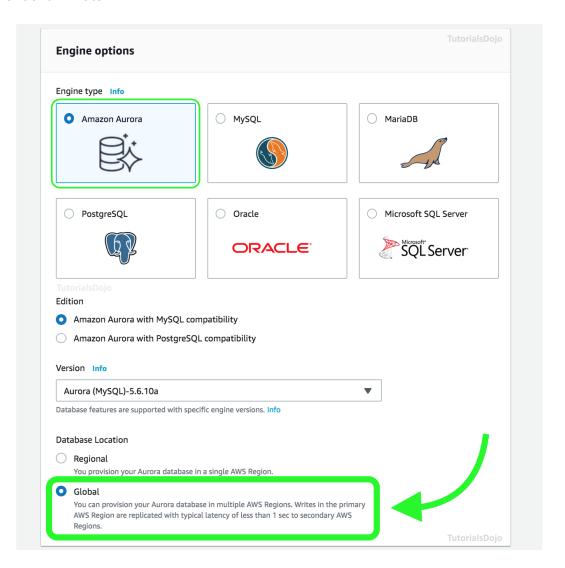
- 1. Set up a new Amazon Redshift database to store the product catalog. Launch a new set of Amazon DynamoDB tables to store the personal information and financial data of their customers.
- 2. Set up a DynamoDB global table to store the product catalog data of the e-commerce website. Use regional DynamoDB tables for storing the personal information and financial data of their customers.
- 3. Set up multiple read replicas in your Amazon Aurora cluster to store the product catalog data. Launch an additional local Amazon Aurora instances in each AWS Region for storing the personal information and financial data of their customers.
- 4. Set up multiple read replicas in your Amazon Aurora cluster to store the product catalog data. Launch a new DynamoDB global table for storing the personal information and financial data of their customers.

Correct Answer: 3

An Aurora global database consists of one primary AWS Region where your data is mastered, and one read-only, secondary AWS Region. Aurora replicates data to the secondary AWS Region with typical latency of



under a second. You issue write operations directly to the primary DB instance in the primary AWS Region. An Aurora global database uses dedicated infrastructure to replicate your data, leaving database resources available entirely to serve application workloads. Applications with a worldwide footprint can use reader instances in the secondary AWS Region for low-latency reads. In the unlikely event, your database becomes degraded or isolated in an AWS region, you can promote the secondary AWS Region to take full read-write workloads in under a minute.



The Aurora cluster in the primary AWS Region where your data is mastered performs both read and write operations. The cluster in the secondary region enables low-latency reads. You can scale up the secondary cluster independently by adding one or more DB instances (Aurora Replicas) to serve read-only workloads. For disaster recovery, you can remove and promote the secondary cluster to allow full read and write operations.



Only the primary cluster performs write operations. Clients that perform write operations connect to the DB cluster endpoint of the primary cluster.

Hence, the correct answer is: Set up multiple read replicas in your Amazon Aurora cluster to store the product catalog data. Launch an additional local Amazon Aurora instances in each AWS Region for storing the personal information and financial data of their customers.

The option that says: Set up a new Amazon Redshift database to store the product catalog. Launch a new set of Amazon DynamoDB tables to store the personal information and financial data of their customers is incorrect because this solution entails a significant overhead of refactoring your application to use Redshift instead of Aurora. Moreover, Redshift is primarily used as a data warehouse solution and is not suitable for OLTP or e-commerce websites.

The option that says: Set up a DynamoDB global table to store the product catalog data of the e-commerce website. Use regional DynamoDB tables for storing the personal information and financial data of their customers is incorrect because although the use of Global and Regional DynamoDB is acceptable, this solution still entails a lot of changes to the application. There is no assurance that the application can work with a NoSQL database, and even so, you have to implement a series of code changes in order for this solution to work.

The option that says: Set up multiple read replicas in your Amazon Aurora cluster to store the product catalog data. Launch a new DynamoDB global table for storing the personal information and financial data of their customers is incorrect because although the use of Read Replicas is appropriate, this solution still requires you to do a lot of code changes since you will use a different database to store your regional data.

References:

https://docs.aws.amazon.com/AmazonRDS/latest/AuroraUserGuide/aurora-global-database.html#aurora-global-database.html#aurora-global-database.advantages

https://docs.aws.amazon.com/AmazonRDS/latest/AuroraUserGuide/AuroraMySQL.Replication.CrossRegion.html

Check out this Amazon Aurora Cheat Sheet:

https://tutorialsdojo.com/amazon-aurora/

Click here for more AWS Certified DevOps Engineer Professional practice exam questions.

More AWS reviewers can be found here



At this point, you should already be very knowledgeable on the following topics:

- 1. Continuous Integration/Continuous Delivery (CI/CD)
- 2. Application Development
- 3. Automation
- 4. Configuration Management and Infrastructure as Code
- 5. Monitoring and Logging
- 6. Incident Mitigation and Event Response
- 7. Implementing Resilient Cloud Solutions
- 8. Security and Compliance

As an AWS DevOps practitioner, you shoulder a lot of roles and responsibilities. Many professionals in the industry have attained proficiency through continuous practice and producing results of value. Therefore, you should properly review all the concepts and details that you need to learn so that you can also achieve what others have achieved.

The day before your exam, be sure to double-check the schedule, location, and items to bring for your exam. During the exam itself, you have 180 minutes to answer all questions and recheck your answers. Be sure to manage your time wisely. It will also be very beneficial for you to review your notes before you go in to refresh your memory. The AWS DevOps Pro certification is very tough to pass, and the choices for each question can be very misleading if you do not read them carefully. Be sure to understand what is being asked in the questions and what options are offered to you. With that, we wish you all the best in your exam!



Domain 1: Software Development Life Cycle (SDLC) Automation



Overview

The first domain of the AWS Certified DevOps Engineer Professional exam checks your preparedness on how well you understand the integration between the AWS services necessary for code development and deployment, such as AWS CodeCommit, AWS CodeBuild, AWS CodeDeploy, and AWS CodePipeline. You will also need a working knowledge of how they complement each other for software development as well as integrations with Amazon EventBridge, CloudWatch Events, Amazon S3, and AWS Lambda. A big part of a normal workday for a DevOps Engineer deals with the software development life cycle.

Roughly 22% of questions in the actual DevOps exam revolve around these topics. The Software Development Life Cycle (SDLC) Automation domain is the biggest exam domain of the DOP-C02 exam, so ensure that you allocate ample time reviewing the topics under this section.

This domain will challenge your know-how in doing the following:

- Configuring code, image, and artifact repositories
- Using version control to integrate pipelines with application environments
- Setting up build processes via AWS CodeBuild and other related services
- Managing build and deployment secrets using AWS Secrets Manager and AWS Systems Manager
 Parameter Store as well as securing other sensitive credentials
- Determining appropriate deployment strategies in AWS CodeDeploy, AWS Elastic Beanstalk and other deployment services
- Running builds or tests when generating pull requests or code merges via AWS CodeCommit and AWS CodeBuild.
- Running load/stress tests, performance benchmarking, and application testing at scale
- Measuring application health based on application exit codes
- Automating unit tests and code coverage
- Invoking AWS services in a pipeline for testing
- Creating and configuring artifact repositories using AWS CodeArtifact, Amazon S3, Amazon Elastic Container Registry (Amazon ECR) and other related services.
- Configuring build tools for generating artifacts using AWS CodeBuild and AWS Lambda
- Automating Amazon EC2 instance and container image build processes via EC2 Image Builder
- Configuring security permissions to allow access to artifact repositories using the AWS Identity and Access Management, CodeArtifact, et cetera.
- Configuring deployment agents via CodeDeploy agent, SSM agent, and others.
- Troubleshooting deployment issues
- Using different deployment methods such as blue/green and canary deployments

In this chapter, we will cover all of the related topics for SDLC automation in AWS that will likely show up in your DevOps Professional exam.



What is DevOps?

Do you ever wonder what DevOps is and why it is so popular in the IT Industry today? Is it a tool, a process, a corporate culture, or a combination of all of these? Why are companies giving competitive salaries for this kind of role?

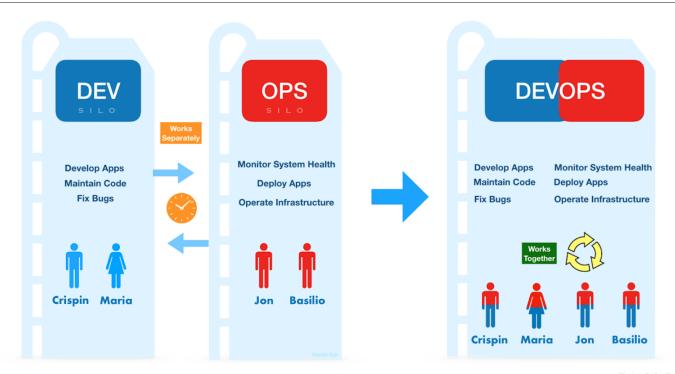
If you typed the word "DevOps" on any job market website today, you would see many available positions that require knowledge of both programming and infrastructure management. You will usually see an advertisement looking for a candidate who knows how to program in Python or any other language. The requirements include being capable of managing servers with configuration management tools such as Ansible, Chef, or Puppet, as well as the provisioning of entire cloud environments using Infrastructure-as-code tools like Terraform or CloudFormation. The salary range offered for these positions is remarkably high too!

Traditional IT companies have a dedicated Development (Dev) team that builds enterprise applications and an Operations (Ops) team that handles the servers and network infrastructure. These two teams are often siloed or isolated from each other. While the Dev team writes the software code, the Ops team prepares the server, database, and other infrastructure needed to run the soon-to-be-released application. In this setup, the developers are entirely oblivious to what the system operators are doing and vice versa. A lot of time is wasted waiting for the Dev Team to fix minor bugs while developing new features and for the Ops team to provision, deploy and scale the needed server resources. When bugs and incompatibility issues are detected in the development cycle, the Ops team waits for the Dev team to address the issue since it is strictly the job of the Developers to fix it. The same is true when there are issues during deployments when the Ops are not familiar with the application and make wrong assumptions which can cause further delays in the deployment targets. Due to this lack of coordination, both the business and its customers are impacted. This is where DevOps comes in!

DevOps is not just the combination of Development (Dev) and Operations (Ops). DevOps is the fusion of practices, processes, tools, and corporate culture that expedite the organization's ability to deliver applications and services at a higher velocity, faster than traditional software development processes. It's not merely a tool or a process that your team adopts, but a synergy of values, corporate structure, and internal processes to attain the digital transformation of the business enterprise. It tears down the traditional and isolated silos of the Development, Operations, IT Security, and other teams, enabling collaboration and improving overall business performance. With DevOps, Developers are empowered to directly influence the deployment life cycle, and the IT Operations folks have the ability to report and fix possible bugs or incompatibilities in the application.

DevOps is not just a framework, rather, it's a cultural approach and a mindset combining operations and development skills and delivering a product (or service) from inception to retirement. Company executives also play a crucial role in allocating budgets and adopting this new status quo within their respective organizations.





Tutorials Dojo

With the advent of Cloud Computing, companies can easily unify their software development and system operation processes. AWS enables organizations to rapidly build, deliver, and manage their products, following DevOps practices with just a click of a button. The efficiency of provisioning new resources, managing infrastructure, deploying application code, automating software release processes, and many other tasks in AWS contributes to overall productivity and business profitability. Because of this massive benefit, companies are willing to pay competitive remuneration for their DevOps Engineers, especially those who are AWS Certified.



A Brief History of the DevOps Exam in AWS

In 2013, Amazon Web Services (AWS) began the Global Certification Program to validate the technical skills and knowledge for building secure and reliable cloud-based applications using the AWS platform. The first-ever certification launched by Amazon is the AWS Certified Solutions Architect – Associate, followed by SysOps Administrator and Developer Associate. A year later, AWS released the first Professional-level certification: AWS Certified Solutions Architect - Professional, and in February 2015, they released the AWS Certified DevOps Engineer Professional.

The AWS Certified DevOps Engineer Professional certification enables technology professionals to showcase their DevOps skills, and it allows companies to identify top candidates to lead their internal DevOps initiatives. It validates your technical expertise in provisioning, managing, and operating distributed application systems on the AWS Cloud platform. It tests your ability to implement and manage Continuous Integration/Continuous Delivery (CI/CD) systems and methodologies on AWS following the industry's best practices, as well as to automate security controls and handle governance processes and meet compliance. The exam also covers core topics such as Software Development Lifecycle (SDLC) automation, security, compliance, monitoring, logging, configuration management, and incident/event response.

As Amazon Web Services continue to evolve, new and updated versions of the AWS certification exams are released regularly to reflect the service changes and to include new knowledge areas. Four years after its initial release, an updated version of the AWS Certified DevOps Engineer - Professional certification was launched in February 2019 with an exam code of DOP-C01.

The latest version of the AWS Certified DevOps Engineer - Professional certification exam was unveiled on March 2023 with an exam code of DOP-C02. AWS is continuously adding more services and features to help organizations and companies to improve their DevOps processes.



Why Automate?

Automation is at the heart of every DevOps engineer. Automation is a key highlight of DevOps practice. There is a saying in the DevOps community to "Automate Everything" and it starts from code inception, releasing to production, application retirement, and everything in between. Eliminating repetitive tasks, reducing toil, and minimizing manual work are the key aspects that you want to solve through automation. Automation in DevOps fosters speed, greater accuracy, consistency, reliability, and rapid delivery.

Here are some major benefits of automation:

- **Speed** to innovate your product faster and adapt to changing market trends. Team members are empowered to make changes quickly as needed, either on the development side or the operational side.
- Rapid delivery increase the pace of your releases by automating your entire deployment pipeline. This is the concept of "fail-fast, iterate faster" in which your companies are incentivized to release minor changes as often as possible, which keeps them on top of competitors.
- **Reliability** continuous integration and continuous delivery processes allow you to reliably and consistently deliver your product to end-users. This also reduces human error as automation rarely makes mistakes as humans do.
- Scale infrastructure as code helps you manage your environments in a repeatable and more efficient manner and scale easily as needed. It gives you a robust system to manage your infrastructure no matter how big or small it is.
- Improved collaboration reduce inefficiencies when collaborating with teams. Automation allows the easier integration of development, testing, and deployment processes. It facilitates faster collaboration between Dev and Ops, which results in an improved turnaround time for bug fixing, deployment, etc.
- Security reduces risk through integrated security testing tools, automating adoption of compliance requirements. It allows you to declare and script your security compliance requirement and make sure they are applied to needed resources in your environments.



Types of Blue Green Deployment via ELB, Route 53, Elastic Beanstalk

AWS Elastic Beanstalk Blue/Green

Elastic Beanstalk, by default, performs an in-place update when you deploy a newer version of your application. This can cause a short downtime since your application will be stopped while Elastic Beanstalk performs the application update.

Blue/Green deployments allow you to deploy without application downtime.

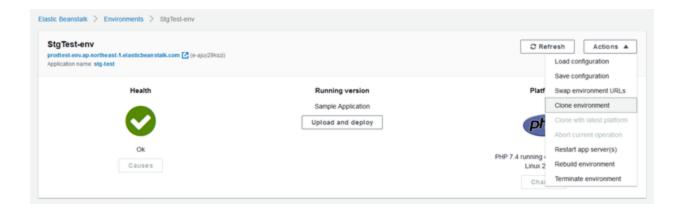
DevOps Exam Notes:

Remember these key points on when to use blue/green deployments:

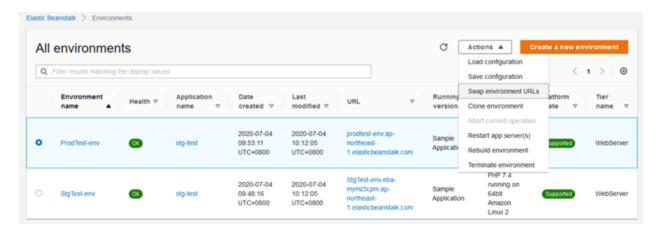
- No downtime during deployment because you are deploying the newer version on a separate environment
- CNAMEs of the environment URLs are swapped to redirect traffic to the newer version.
- Route 53 will swap the CNAMEs of the application endpoints.
- Fast deployment time and quick rollback since both old and new versions are running at the same time, you just have to swap back the URLs if you need to rollback.
- Useful if your newer version is incompatible with the current platform version of your application. (ex. Jumping from major versions of NodeJS, Python, Ruby, PHP, etc.)
- Your RDS Database instance should be on a separate stack because the data will not transfer to your second environment. You should decouple your database from the web server stack.

To implement a Blue/Green deployment for your Elastic Beanstalk application, you can perform the following steps:

1. Create another environment on which you will deploy the newer version of your application. You can clone your current environment for easier creation.

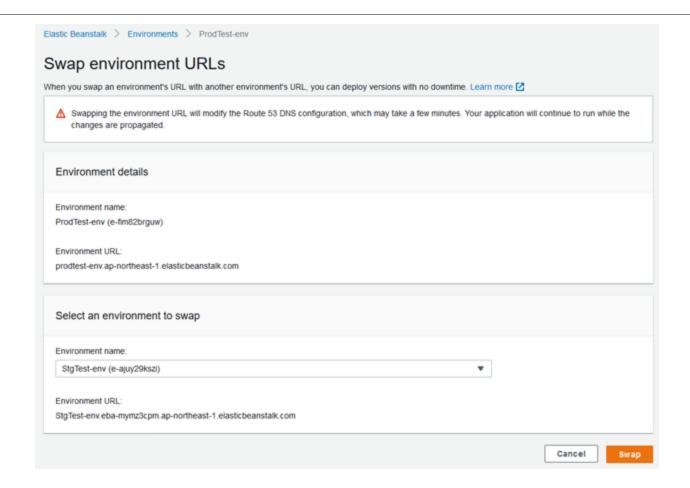


- 2. Once the new environment is ready, deploy a new version of your application. Perform your tests on the URL endpoint of your new environment.
- 3. After testing, select your Production environment, click Actions > Swap environment URLs.



4. On the Swap Environment URLs page, select the newer environment and click Swap to apply the changes.







AWS Lambda Blue/Green

You can also implement Blue/Green deployments on your Lambda functions. The concept is the same as in Elastic beanstalk blue/green deployment i.e. you will need to create two versions of your Lambda function and use function Aliases to swap the traffic flow.

Lambda versions – lets you publish a new version of a function that you can test without affecting the current application accessed by users. You can create multiple versions as needed for your testing environments. The ARN of Lambda version is the same as the ARN of the Lambda function with added version suffix.

arn:aws:lambda:aws-region:acct-id:function:helloworld:\$LATEST

Lambda aliases – Aliases are merely pointers to specific Lambda versions. You can't select a Lambda alias and edit the function. You need to select the LATEST version if you want to edit the function. Aliases are helpful for blue/green deployments because it allows you to use a fixed ARN and point it to a particular Lambda version that you want to deploy.

DevOps Exam Notes:

Remember the difference between Lambda \$LATEST, Lambda Versions, and Lambda Aliases:

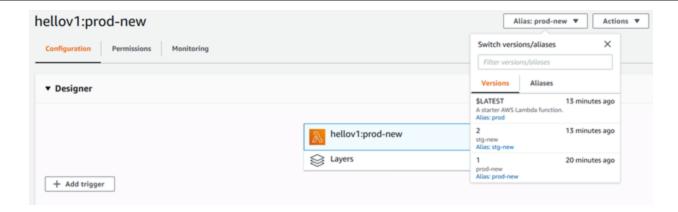
\$LATEST - this is the latest version of your Lambda function. You can freely edit this version.

Lambda Version - fixed version of your function. You can't edit this directly.

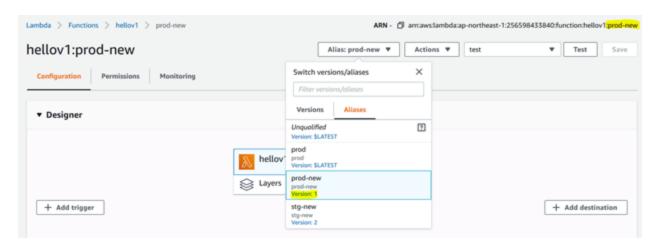
Lambda Alias - a pointer to a specific Lambda version. You can perform blue/green deployment with Aliases by pointing to a newer version.

The following steps will show how blue/green deployment can be done on Lambda functions.

1. The current version of your Lambda function is deployed on Version 1. Create another version and make your changes, this will be Version 2.



2. Create an Alias that will point to the current production version. Use this alias as your fixed production ARN.



3. Create another Alias that you will use for your newer version. Perform your testing and validation on this newer version. Once testing is complete, edit the production alias to point to the newer version. Traffic will now instantly be shifted from the previous version to the newer version.

Sources:

https://docs.aws.amazon.com/elasticbeanstalk/latest/dg/using-features.CNAMESwap.html https://docs.aws.amazon.com/lambda/latest/dg/configuration-aliases.html https://docs.aws.amazon.com/lambda/latest/dg/configuration-versions.html



AWS Lambda Function Alias Traffic Shifting

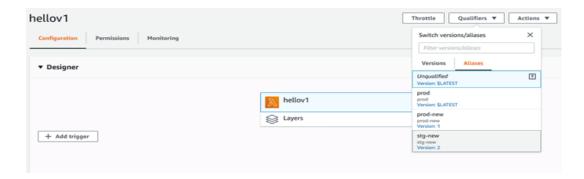
On EC2 instances, you can perform a canary deployment by deploying a newer application version on a single EC2 instance and analyzing the production traffic flowing through it. If you are satisfied with it, you'll proceed to deploy the newer version on all EC2 instances. However, for deployments on your Lambda functions, you can't use a canary deployment since you don't deploy your application directly on EC2 instances.

DevOps Exam Notes:

To provide similar functionality as a canary deployment, AWS Lambda gives you the ability to use Function Aliases to shift the percentage of traffic from one version to another. Essentially, you will create an Alias that points to the current version of the Lambda function, then use a weighted alias to define a newer version of the Lambda function. You can then define the weight (percent of traffic) that you want to forward to this version. After validation, you can completely shift traffic to the newer version.

You can consult the previous section (Types of Blue/Green deployment - AWS Lambda) on how to create AWS Lambda Versions and Aliases. Here's an example of how to control that percentage of traffic flowing to different Lambda functions using function alias. This is similar to the way a canary deployment works.

1. Select the function alias pointing to the current production version.

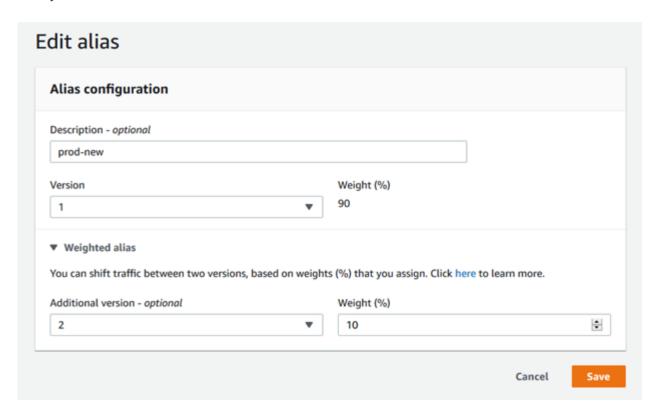


2. Edit the alias configuration.

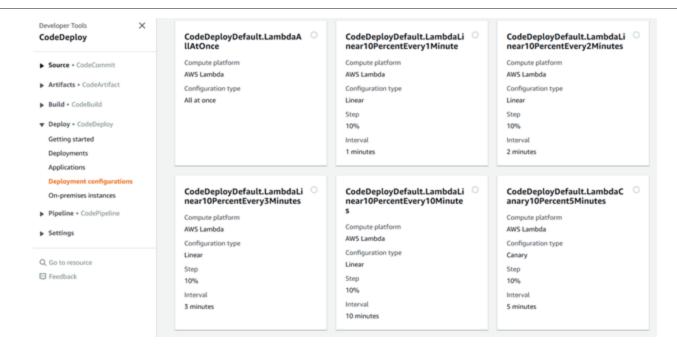




3. On the weighted alias section, select the newer version of your Lambda function and assign the percentage of traffic to shift to the newer version. You can repeat this step multiple times if you want to slowly shift traffic from the older version to the newer version.



If you are using AWS CodeDeploy to deploy your Lambda functions, CodeDeploy uses Aliases to shift traffic to the newer version. As you can see on the options on the deployment configurations, CodeDeploy can automate this gradual traffic shifting for your Lambda Functions.



See the <u>CodeDeploy - Linear, Canary and All-at-Once (Pre-defined Deployments)</u> topic for more discussion on these deployment strategies.

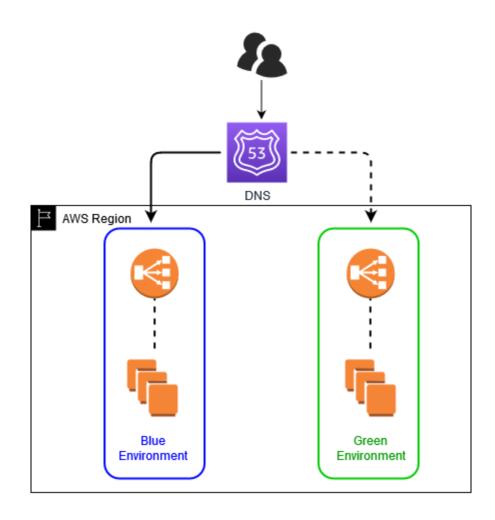
Source:

https://aws.amazon.com/blogs/compute/implementing-canary-deployments-of-aws-lambda-functions-with-alias-traffic-shifting/



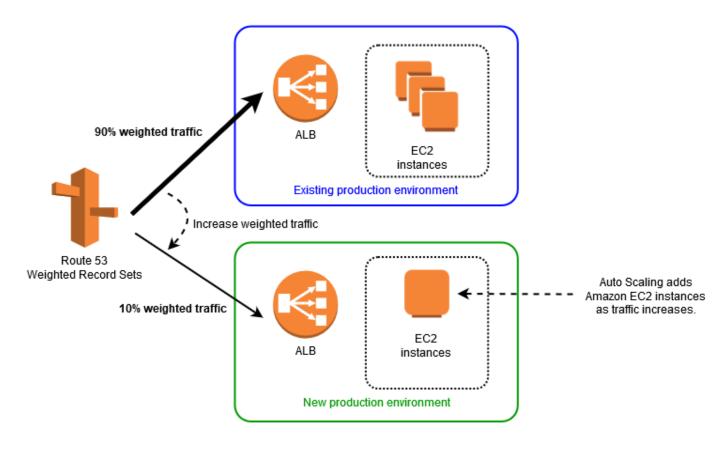
Basic Blue/Green Deployment using Route 53

Blue/green deployment on the AWS platform provides a safer way to upgrade production software. This deployment usually involves two environments, the production environment (blue) and the new updated environment (green).



Once the new version is deployed on the green environment, you can validate the new software before going live. Then, you start shifting traffic away from the blue environment and sending it to the green one. Normally, you'd use Route 53 weighted routing policy because it gives you an easy way to push incremental traffic to the green environment or revert traffic back to the blue environment in case of issues. If you want to, you can switch the traffic immediately by updating the production Route 53 record to point to the green endpoint. Users will not see that you changed the endpoint since from their perspective, the production URL is the same.





You can also shift a small portion (like 10%) of traffic on the Green environment by using a weighted routing policy on Route 53. This way, you can test live traffic on the new environment, analyze the new logs, and then you can easily revert to the original environment if you find any problems. This process is also called a canary deployment.

Source:

https://aws.amazon.com/blogs/startups/upgrades-without-tears-part-2-bluegreen-deployment-step-by-step-on-aws/