

# Test Project

## *Cloud Computing*

Independent Test Project Designers: Shawn Xu, Xiaowei Wang, Qiang Wang, Xu Chen  
Independent Test Project Validator: Chit Boon Lee, SCM

# Contents

Description of project and tasks.....	3
Tasks .....	3
Technical Details.....	3
Module 1: Troubleshooting .....	3
<i>Deploying serverless applications with AWS CodePipeline</i> .....	3
Inventory .....	4
<i>Build a self-service RESTful API to save your time</i> .....	4
Inventory .....	4
<i>CodeDeploy Grand Battle</i> .....	4
Inventory .....	5
<i>Upgrade Peering Connections</i> .....	5
Inventory .....	6
<i>Rebuild Your Application in a Disrupted Network</i> .....	6
Inventory .....	6
<i>DevOps Mobile Development</i> .....	6
Inventory .....	7
<i>EKS Troubleshooting</i> .....	7
Inventory .....	7
<i>ECS Troubleshooting</i> .....	7
Inventory .....	8
<i>Centralize Your Network Egress</i> .....	8
Inventory .....	8
<i>Complete API Functions Using Lambda</i> .....	8
Inventory .....	9
<i>EKS Scaling</i> .....	9
Inventory .....	9
<i>Data Analytics</i> .....	9
Inventory .....	10
Module 2: Quest .....	11
<i>Security Chaos</i> .....	11

# Contents

This Test Project consists of the following documentation/files:

1. WSC2024\_TP53\_main\_document\_actual\_en
2. WSC2024\_TP53\_day1\_actual\_en
3. WSC2024\_TP53\_day2\_actual\_en
4. **WSC2024\_TP53\_day3\_actual\_en**

## Description of project and tasks

This Test Project aims to rigorously evaluate your cloud computing skills through a series of challenges. By completing these modules, you will demonstrate your ability to effectively troubleshoot complex cloud environments and implement innovative solutions.

Today's Test Project comprises two core components:

- **Troubleshooting Module:** This section will test your problem-solving skills through a series of challenges designed to replicate real-world cloud scenarios.
- **Quest Module:** A comprehensive challenge that will assess your overall cloud architecture and implementation abilities.

## Tasks

1. Log into the Cloud Raiser platform
2. Read the documentation thoroughly
3. Continue until the test project day has completed

## Technical Details

### Module 1: Troubleshooting

#### Deploying serverless applications with AWS CodePipeline

##### Summary

In today's fast-paced digital world, businesses are increasingly leveraging cloud technologies to achieve scalability, flexibility, and cost efficiency. Deploying serverless applications has become a popular approach due to its ability to automatically manage the infrastructure, allowing developers to focus solely on writing code. AWS CodePipeline is a powerful continuous integration and continuous delivery (CI/CD) service that automates the build, test, and deployment phases of application development, ensuring rapid and reliable updates.

In this troubleshooting guide, we will provide a comprehensive walkthrough of deploying a serverless application using AWS CodePipeline. Throughout this process, we will encounter and resolve a variety of issues that commonly arise during deployment and access of the application. By addressing these challenges, businesses can streamline their deployment processes, minimize downtime, and enhance the overall reliability and performance of their serverless applications.

## REMINDER

This is a Troubleshooting Challenge.

Please use us-east-1 region.

## Inventory

- CodePipeline

## Build a self-service RESTful API to save your time

### Summary

You are a technical support engineer at ABC Company, a leading firm in the tech industry known for its innovative solutions and efficient operations. Recently, you have received numerous requests from various departments to check and update employee data in the DynamoDB tables you manage. This data is crucial for various HR processes, payroll management, and compliance reporting.

So far, you have had to complete these tasks manually through the system console, which is both time-consuming and prone to errors. As the company continues to grow and the volume of these requests increases, it has become clear that a more efficient solution is needed.

Recognizing the inefficiency and potential for human error, you have decided to build a self-service solution using a serverless architecture to automate the process. This will not only save time but also ensure accuracy and consistency in data management.

Next time you need to update employee data, you will only need to provide the requester with a RESTful API, streamlining the process and allowing for quicker responses to data update requests. This approach aligns with ABC Company's commitment to leveraging cutting-edge technology to enhance productivity and operational efficiency.

## REMINDER

This is a Troubleshooting Challenge.

Please use us-east-1 region.

## Inventory

- API Gateway
- Lambda
- DynamoDB

## CodeDeploy Grand Battle

### Summary

Due to a personnel change, you have recently taken over the operations and maintenance of the company's core systems. These systems are critical for maintaining the company's business operations, including customer management, transaction processing, and data analytics. They ensure that the business runs smoothly and efficiently, providing a seamless experience for customers and stakeholders alike.

Recently, there has been a significant issue: for some reason, someone deleted the contents of CodeDeploy and CodePipeline, which were used to publish the system. This deletion has caused a major disruption, as these services are essential for the deployment and continuous integration/continuous deployment (CI/CD) pipelines that keep the system updated and running smoothly. To make matters worse, there are no past logs to look up the cause of this deletion, leaving you with no clues about what went wrong or how to prevent it from happening again.

Given the urgency of the situation, the system now needs a new release, and as an AWS expert, you are urgently required to rebuild the CodeDeploy and CodePipeline services used for the release. This task is critical not only to restore normal operations but also to ensure the continuity and reliability of the company's core systems.

Your responsibilities include setting up CodeDeploy and CodePipeline from scratch, configuring them to match the previous deployment setup as closely as possible, and ensuring that they are robust enough to handle future releases without similar issues. This will involve working closely with the development team to understand the system's requirements, configuring the necessary AWS resources, and thoroughly testing the new deployment pipeline to ensure it works correctly.

The goal is to quickly restore the deployment capabilities so that the system can receive updates and improvements, minimizing downtime and maintaining the company's high standards of operational efficiency and customer satisfaction.

### REMINDER

This is a Troubleshooting Challenge

Please use us-east-1 region.

### Inventory

A single Web server hosting the application. Its address is in output properties.

- CodeDeploy
- Codepipeline

## Upgrade Peering Connections

### Summary

Your company, Aegis Corp, is a heavily regulated giant in the financial industry that began migrating to the cloud last year. A key issue hindering their strategy has been the lack of internal knowledge and experience in operating AWS workloads. With the rapid development of the AWS partner ecosystem, should the company really be held back by this? Of course not! Aegis Corp thought the same and partnered with the trusted AWS partner, Ninja Operations. After successfully running the initial applications for a year, Aegis Corp has now decided to fully embrace cloud computing. So far, they have been hosting all workloads within a single VPC, but future scalability demands more! You were recently hired as a network expert to help guide this effort.

Ninja Operations' tools need to establish a private network connection with your VPC to manage the workloads within it. Until now, this has been achieved through intra-regional VPC peering. With the planned addition of more VPCs in the coming months, you know that now is the right time to use different AWS network services and move away from VPC peering. To gain stakeholders' trust, you aim to complete this task within the next two days. The biggest challenge is ensuring good security practices - remember, this is a heavily regulated client!

Aegis Corp's security team will provide guidance throughout the process, outlining the requirements that the network setup needs to meet. Time is of the essence. You need to act swiftly!

The current setup includes the following components:

### REMINDER

This is a Troubleshooting Challenge

Please use us-east-1 region.

## Inventory

- VPC
- Transit Gateway

## Rebuild Your Application in a Disrupted Network

### Summary

SuperTech Co., Ltd., a leading technology firm, recently had one of its test accounts compromised by hackers. The attackers managed to delete most of the system resources associated with this account and disrupted network connectivity.

Despite this breach, the application source code and documentation are backed up and intact. As the company's Chief Cloud Architect, your immediate task is to swiftly restore network connectivity for this test account and ensure that the applications continue to function seamlessly.

This incident underscores the critical importance of robust cybersecurity measures in safeguarding sensitive company data and operations. Your expertise and rapid response are crucial in mitigating the impact of this security breach and restoring normal operations to maintain business continuity.

Upon initial investigation, you discovered that only five VPCs from the architecture diagram were preserved in the compromised test account. The remaining resources were deleted by the hackers.

### REMINDER

This is a Troubleshooting Challenge

Please use us-east-1 region.

## Inventory

- RDS
- TGW

## DevOps Mobile Development

### Summary

"If it works on my machine." If you've been in development for a while, you've probably heard this phrase. It often reflects the frustration of encountering an elusive error that seems to arise from unique settings on the developer's personal workstation.

In a business context, relying solely on individual developer environments can lead to inconsistencies and delays in software deployment. To mitigate these challenges, integrating continuous integration into your workflow is crucial. This approach ensures that when code is committed to the repository, it undergoes automated builds using an approved build system. Unit tests are executed, and the resulting compiled code, along with comprehensive logs (essential for troubleshooting and ready-to-deploy configurations), is uploaded to a centralized and accessible location. This practice not only streamlines the development process but also enhances collaboration and accelerates the deployment of reliable software solutions across the organization.

### REMINDER

This is a Troubleshooting Challenge

Please use us-east-1 region.

## Inventory

- IAM
- CodePipeline

## EKS Troubleshooting

### Summary

As a Cloud Platform Engineer at a multinational internet company, you oversee the operations of the cloud infrastructure team. Recently, you received two urgent incident reports. Firstly, an EKS node was flagged as NotReady, impacting the availability of critical business applications. Secondly, the 2048 game application hosted on the EKS cluster became inaccessible, severely affecting user experience and interaction.

Your task is to swiftly address both issues. You need to diagnose the root causes and implement corrective measures to restore the node to normal operation and ensure seamless access to the game application. By meticulously documenting your steps, monitoring outcomes, and taking appropriate actions, you aim to uphold business continuity and enhance user satisfaction.

These incidents highlight the critical importance of maintaining robust cloud infrastructure for uninterrupted service delivery and positive user engagement, crucial for sustaining the company's competitive edge in the digital marketplace.

### REMINDER

This is a Troubleshooting Challenge

Please use us-east-1 region.

## Inventory

- EKS
- EC2

## ECS Troubleshooting

### Summary

As a Cloud Architect for a rapidly growing company, you are tasked with deploying a containerized application across two different VPCs for workload isolation and security. This deployment is crucial for supporting the company's expansion and data protection needs.

The application, part of a new initiative for real-time data processing and analytics, handles sensitive customer data and internal metrics. Using two VPCs ensures component isolation, reducing vulnerabilities and allowing for granular access controls.

A key requirement is a shared file system for configuration files and log data, accessible from containers in different VPCs. You have chosen Amazon ECS for container orchestration and Amazon EFS for the shared file system.

This project is vital for the company's digital transformation strategy, enhancing operational efficiency and delivering superior services to customers. Leveraging AWS services, you aim to create a resilient, secure, and scalable infrastructure supporting the company's growth.

### REMINDER

This is a Troubleshooting Challenge

Please use us-east-1 region.

## Inventory

- ECS
- EC2
- Elastic Load Balancing

## Centralize Your Network Egress

### Summary

When you joined ABC Company's cloud security team, you were assigned a critical project. This project involves implementing security measures for multiple Virtual Private Clouds (VPCs) within ABC Company and ensuring that all traffic from these VPCs undergoes rigorous inspection before leaving the system.

ABC Company operates a distributed infrastructure where each department has its own VPC, hosting various business applications and sensitive data. To protect these assets, ABC Company decided to establish a centralized security architecture for greater visibility and control.

As a Cloud Security Engineer, your primary task is to design a centralized egress VPC, serving as the exit point for all VPC traffic. In this egress VPC, you will use AWS Network Firewall to inspect and filter the traffic. This powerful tool will help you monitor and block potential threats from the various spoke VPCs.

In the centralized egress VPC, you will configure and manage AWS Network Firewall rules to inspect traffic based on predefined security policies. These rules can be customized based on IP addresses, ports, protocols, and other key factors, ensuring only verified traffic can leave the system.

With this security architecture, you will enhance the overall network security of ABC Company and provide unified security controls for each department's VPC. You will regularly monitor and audit the traffic, promptly detecting and responding to potential security threats, ensuring ABC Company's operations run in a trusted environment.

### REMINDER

This is a Troubleshooting Challenge

Please use us-east-1 region.

## Inventory

- VPC
- AWS Network Firewall
- Transit Gateway

## Complete API Functions Using Lambda

### Summary

A city is planning to host the Happy Cup University Football Friendship Tournament, featuring four football teams competing against each other. Your task is to calculate the points based on the match results and determine the final rankings of each university team.

As a Cloud Engineer, your specific task for this project is to use AWS Lambda to fulfill the business requirements. Other team members will handle related services such as API Gateway.

This project is part of a larger initiative to promote sports and camaraderie among university students in the city. The tournament aims to foster healthy competition and teamwork while providing an exciting event for the local community. Efficiently managing and processing the match data using cloud-based solutions is crucial for the success of this initiative.



## REMINDER

This is a Troubleshooting Challenge

Please use us-east-1 region.

## Inventory

- Lambda

## EKS Scaling

### Summary

Your company recently launched a new e-commerce brand on EKS, which has garnered significant acclaim and attracted a large number of users. The DevOps team has been tasked with quickly implementing auto-scaling to handle increased user traffic while operating EKS in a cost-effective manner.

To address this challenge, the team decided to implement Karpenter, a tool that offers faster and more diverse auto-scaling configurations compared to the existing Cluster Autoscaler (CAS). Karpenter is designed to optimize the scaling of Kubernetes clusters by automatically launching just the right compute resources to handle changing application loads.

This implementation is part of a larger strategy to ensure the e-commerce platform remains highly available and performs well under varying traffic conditions. The project aims to enhance user experience by minimizing downtime and latency, which are crucial for maintaining customer satisfaction and loyalty in the competitive e-commerce market.

In this challenge, you will deploy Karpenter to manage capacity-related auto-scaling within the EKS environment. You will also practice using the various settings and options provided by Karpenter to fine-tune the scaling behavior, ensuring that the system can dynamically adjust to traffic demands while optimizing resource usage and costs. This initiative is essential for supporting the company's growth and maintaining its reputation for reliability and efficiency.

## REMINDER

This is a Troubleshooting Challenge

Please use us-east-1 region.

## Inventory

- EKS

## Data Analytics

### Summary

As a data engineer for a global e-commerce company, you are responsible for maintaining and optimizing the company's real-time sales data ETL (Extract-Transform-Load) pipeline. The primary purpose of this pipeline is to collect sales order data from e-commerce websites around the world and dynamically add the corresponding country/region codes based on the country/region information in the orders. The processed data is ultimately stored in an Amazon S3 bucket to serve as a data source for subsequent business analysis.

However, due to a recent infrastructure upgrade within the company, the entire data processing pipeline has experienced a severe failure, rendering it completely inoperative. As the engineer responsible for this pipeline, you need to comprehensively investigate the root causes of the failure and perform necessary repairs and optimizations to ensure the pipeline's stable operation and data accuracy.

This real-time data ETL pipeline is built on AWS Kinesis Data Streams, AWS Glue, and AWS Glue Streaming Jobs. The specific architecture is as follows:

- A Python script utilizes boto3 to push sales order data into a Kinesis Data Stream.
- AWS Glue captures data in real-time from the Kinesis Data Stream.
- An AWS Glue Streaming Job reads country/region code and name mapping data from an S3 bucket and dynamically adds the corresponding country/region code based on the country/region name in the order data.
- The processed data is written into a designated S3 bucket.

In the following tasks, you need to analyze and resolve various potential issues within the pipeline and implement the corresponding optimization objectives.

#### REMINDER

This is a Troubleshooting Challenge

Please use us-east-1 region.

#### Inventory

- Kinesis
- Key Management Service
- Glue
- S3 (Simple Storage Service)

## Module 2: Quest

### Security Chaos

#### Summary

#### Company Background

ABC, a dynamic and forward-thinking company, has established itself as a leader in the tech industry. With a strong commitment to innovation and excellence, ABC provides cutting-edge solutions and services to its diverse clientele.

#### Business Challenge

As ABC expanded its operations, the company faced significant challenges with its existing on-premises infrastructure. The need for greater scalability, enhanced performance, and cost-effective solutions became critical. Managing and scaling physical servers was not only resource-intensive but also hindered the company's ability to quickly adapt to changing market demands.

#### Solution

To overcome these challenges, ABC chose to deploy Amazon EC2 instances on a public cloud platform as a pilot for the initial POC project. If progress is satisfactory and all parameters meet expectations, the migration scope will gradually expand, ultimately achieving the company's overall goal of cloud adoption. This strategic initiative aims to leverage the flexibility, scalability, and cost-efficiency of cloud computing to support their growing business needs.

#### Security

Information security holds a high priority in this project. During this POC phase, it is imperative to thoroughly address security-related requirements to ensure that the company's information security standards are robustly upheld throughout the cloud migration process. Special attention will be given to the security measures for EC2 instances, ensuring comprehensive protection and adherence to the highest security protocols.

#### Logging

When designing your AWS cloud architecture for this POC project, it is essential to enable extensive logging features. Maximizing logging capabilities ensures comprehensive visibility into your system's performance, security, and operational health. This is crucial for effective troubleshooting and optimization. Additionally, you need to set up alarms to promptly address any issues affecting the availability of the web application.

#### REMINDER

This is a Quest Challenge

Please use **us-east-1** region.