EC-Council

C|CSE™

Certified Cloud Security Engineer

hide01.ir

Business Continuity and Disaster Recovery in Cloud

Module 09

This page is intentionally left blank.

hide01.ir

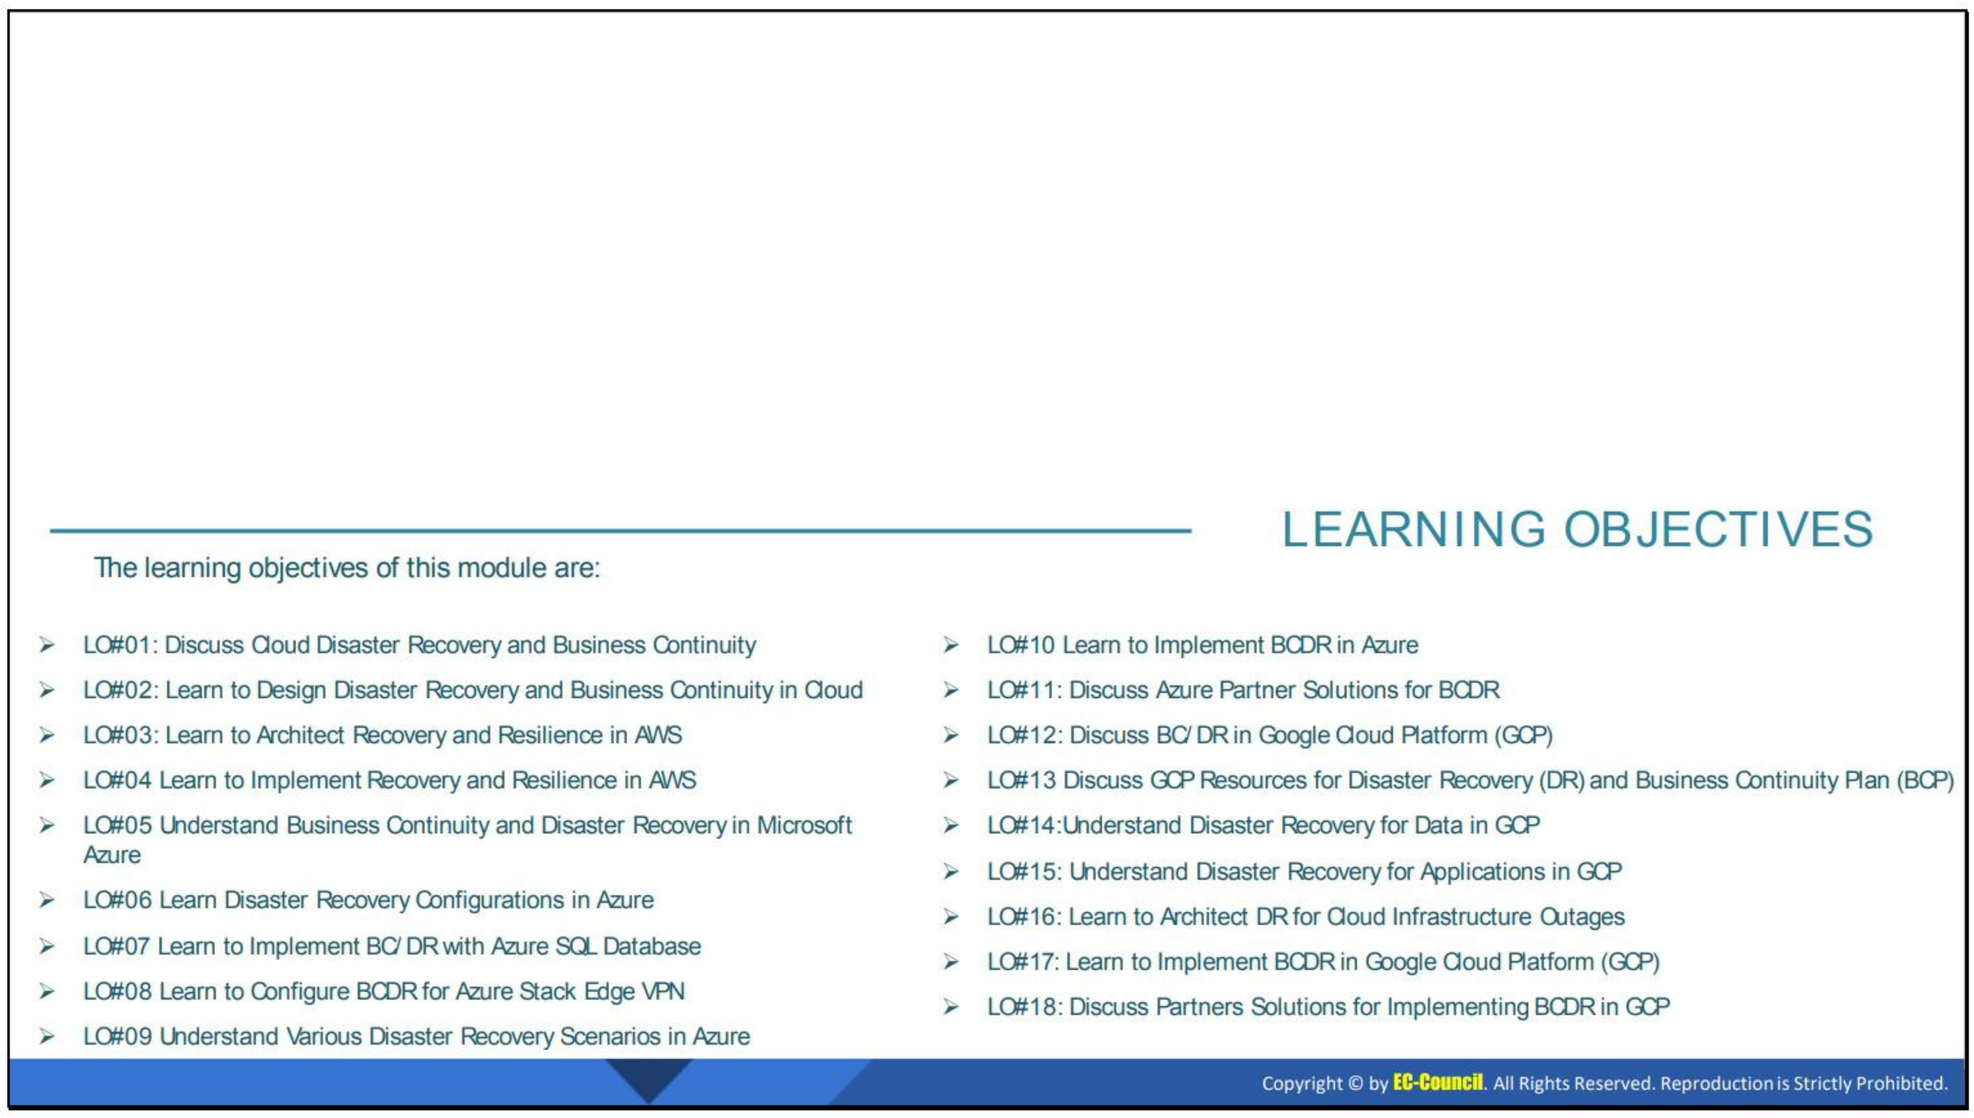

Learning Objectives

After adopting the cloud environment, the organization should design and implement a robust business continuity and disaster recovery plan to safeguard the data and the production environment. A disruptive event can strike an organization at any point in time; therefore, every organization that adopts the cloud environment should prepare itself in advance. The cloud security engineer plays an important role in safeguarding the organization against any disaster event.

The cloud security engineers should be well-acquainted with business continuity and disaster recovery plans to restore the business function to normal within a short stint of time following a disruptive event. They should be capable of preventing the disaster event or at least ensuring that the disaster event causes the least damage to core functionality. This module provides robust information to the cloud security engineer regarding business continuity and disaster recovery in the cloud ecosystem. Through this module, the cloud security engineer will learn how to protect and secure their organization from any type of disruptive event.

The learning objectives of this module are as follows:

- Introduction to cloud disaster recovery and business continuity
- Design a disaster recovery and business continuity in cloud
- Learn architecting for recovery and resilience in AWS
- Implement recovery and resilience in AWS
- Understand business continuity and disaster recovery in Microsoft Azure
- Understand disaster recovery configuration in Azure

Certified Cloud Security Engineer Copyright © by **EC-Council**
All Rights Reserved. Reproduction is Strictly Prohibited.

- Learn to implement BCDR with Azure SQL Database
- Learn to configure BCDR for Azure Stack Edge VPN
- Understand various recovery scenarios in Azure
- Learn to implement BCDR in Azure
- Understand business continuity and disaster recovery in GCP
- Learn data disaster recovery in GCP
- Learn applications disaster recovery in GCP
- Implement BCDR in GCP

hide01.ir

Certified Cloud Security Engineer Copyright © by **EC-Council**
All Rights Reserved. Reproduction is Strictly Prohibited.

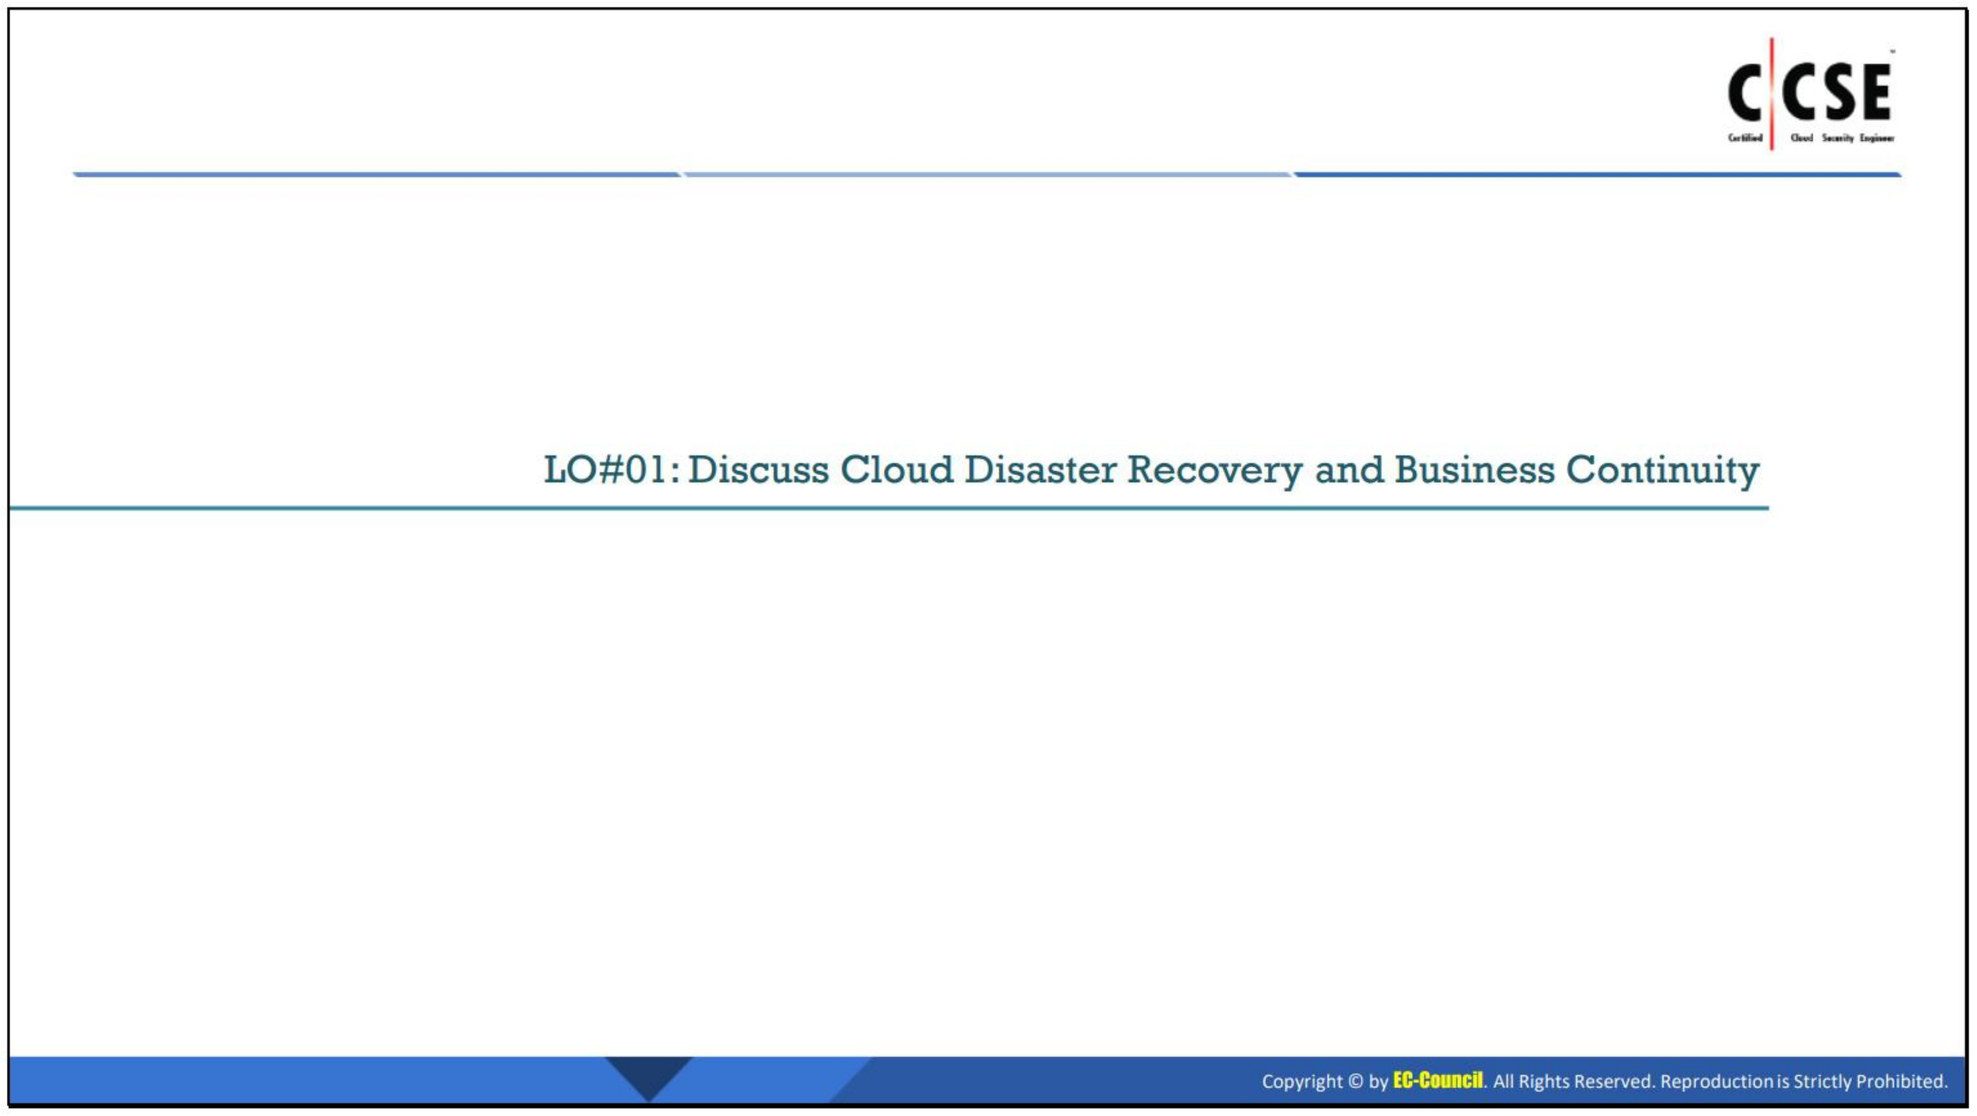

LO#01: Introduction to Cloud Disaster Recovery and Business Continuity

Gaining knowledge on important topics such as cloud disaster recovery, DR objectives, business continuity (BC) in the cloud, etc., allows cloud security engineers to understand how to effectively prepare for business continuity and disaster recovery (BCDR) in their organizations. The objective of this section is to understand important concepts of cloud BCDR that are essential for the business continuity of an organization following a disaster.

Certified Cloud Security Engineer Copyright © by **EC-Council**
All Rights Reserved. Reproduction is Strictly Prohibited.

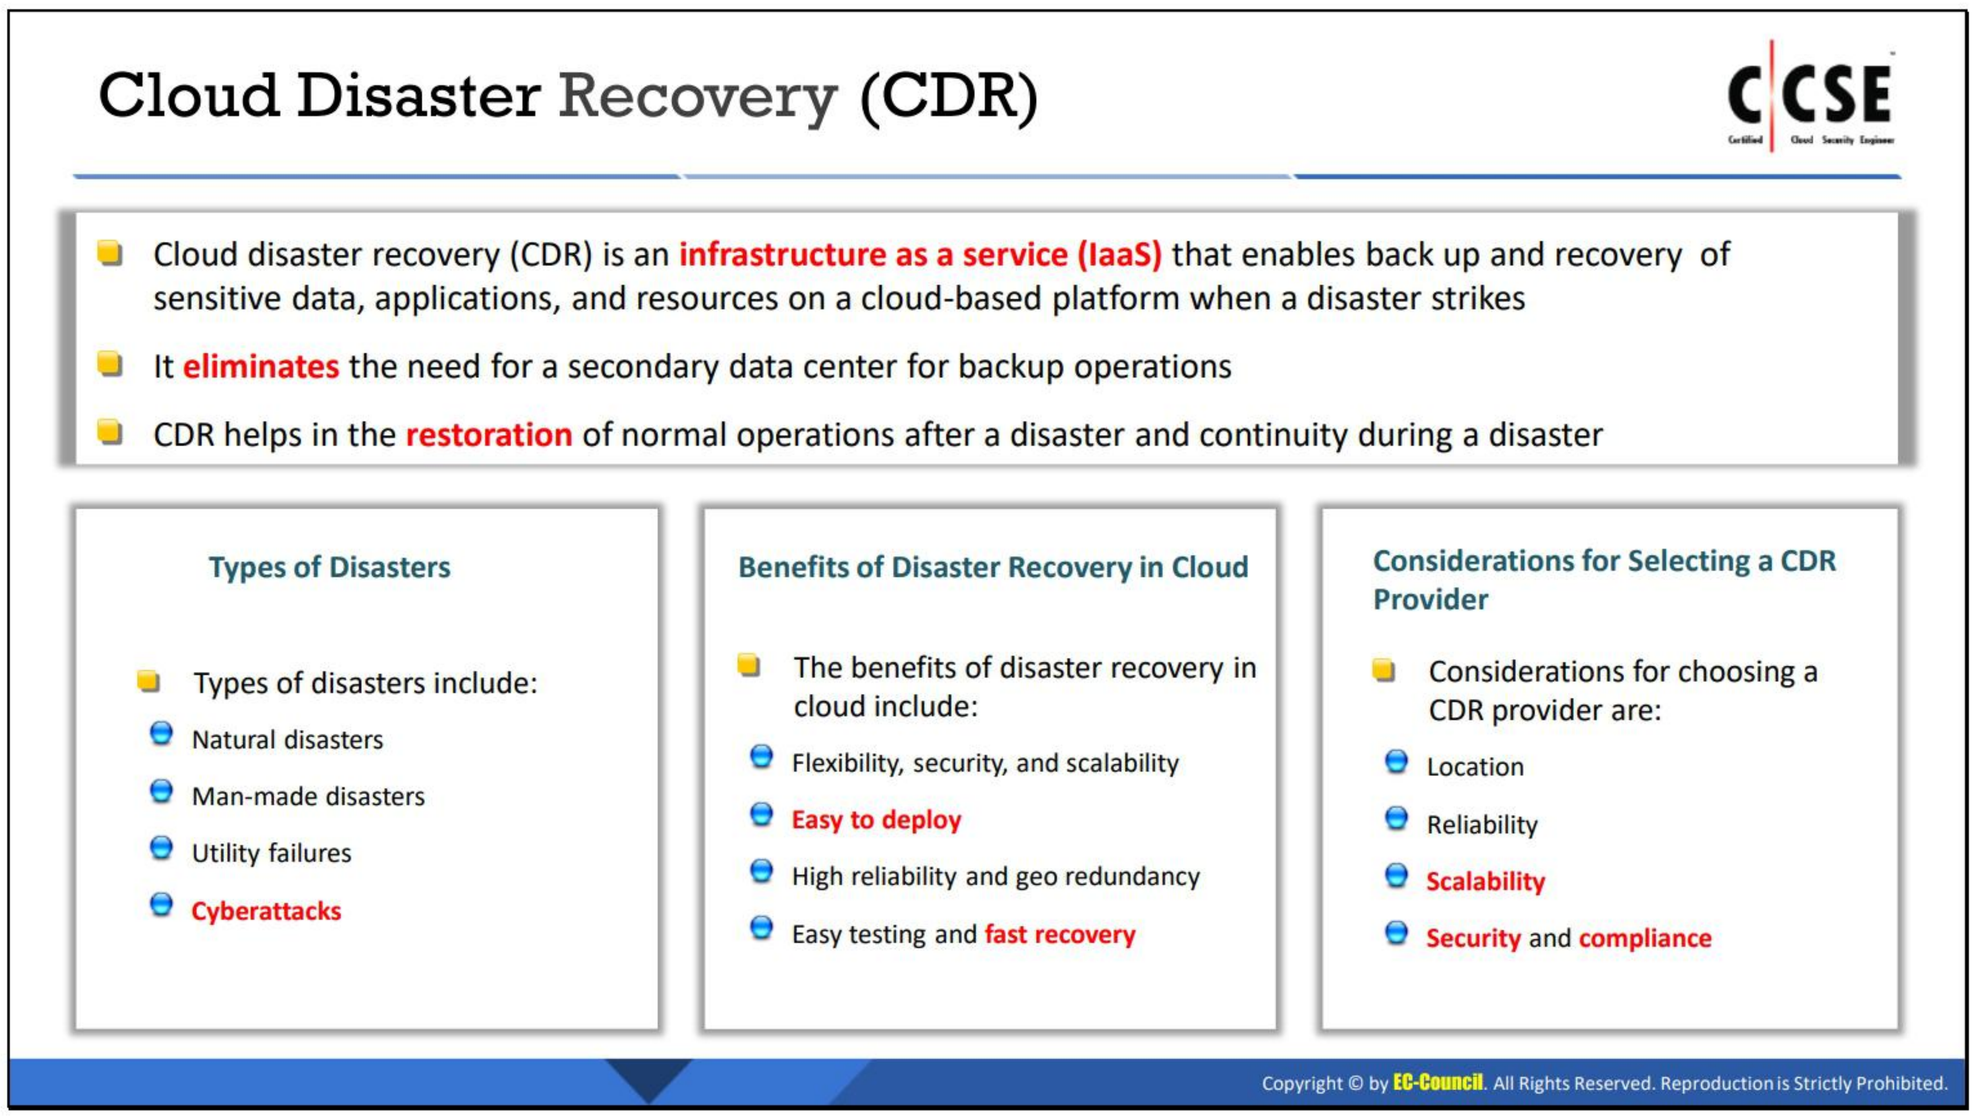

Cloud Disaster Recovery (CDR)

What is Disaster Recovery in Cloud?

CDR is the method that is employed to backup the data, applications, and other resources to the cloud when a disaster strikes the infrastructure. CDR is an IaaS that enables backup and recovery of sensitive data, applications, and resources on a cloud-based platform when a disaster strikes. CDR is similar to traditional disaster recovery, where the valuable assets and sensitive information are protected, and can restore normal operations when a disaster occurs.

CDR automates the process of transferring the data to the recovery sites, thereby providing high-availability infrastructure. CDR offers an automated, scalable, high performance disaster recovery services without the need of a datacenter. CDR helps in the restoration of normal operations after a disaster and continuity during a disaster.

Types of Disasters

- **Natural disasters:** Natural disasters include forces of nature that pose a risk to the critical infrastructure. These include wildfires, tornadoes, hurricanes, floods, and earthquakes.
- **Man-made disasters:** These include that cause disruption to the critical infrastructure due to human negligence or mistakes.
- **Utility failures:** These include utility failures that cause disruption in normal operations such as electricity or power failure, loss of communication, etc.
- **Cyberattacks:** Cyberattacks include attack on the technical assets such as information leak, ransomware, SQL injection, DOS attacks, etc.

Certified Cloud Security Engineer Copyright © by **EC-Council**
All Rights Reserved. Reproduction is Strictly Prohibited.

Benefits of Disaster Recovery in Cloud

- **Flexibility and scalability**: CDR services offer a large number of resources on demand for disaster recovery, and as the infrastructure grows the disaster recovery solutions can grow with it.
- **Easy to deploy**: CDR services are easy to set up. For backup services, the provider can transfer the data to the cloud quickly.
- **High reliability and geo-redundancy**: Cloud providers have global coverage, that is, multiple data centers can support users in different regions. This ensures service availability and high reliability.
- **Easy testing and fast recovery**: For testing the availability of the workloads, cloud workloads operate with virtual machines (VM); thus, VM image files can be copied to test servers without any issues on the production workloads. Data transfer speeds can be optimized through high bandwidth and fast disk input/output options.

How to Select a Cloud Disaster Recovery Provider?

The CDR provider can be selected based on six factors.

- **Location**: The location of the DR provider should be selected by considering the physical distance and latency. If the location is nearby, there is a risk of shared physical disaster risk; however, if it is far away, the latency increases.
- **Reliability**: Service downtime during disaster and during recovery should be minimal.
- **Scalability**: The cloud provider must be able to protect selected data and resources; they should also accommodate additional resource to deliver best performance.
- **Security and compliance**: The cloud provider must offer authentication, virtual private network (VPN), and encryption to protect the business; they should also meet compliance standards such as ISO 27001, SOC 2, SOC 3, and PCI-DSS.

Certified Cloud Security Engineer Copyright © by **EC-Council**
All Rights Reserved. Reproduction is Strictly Prohibited.

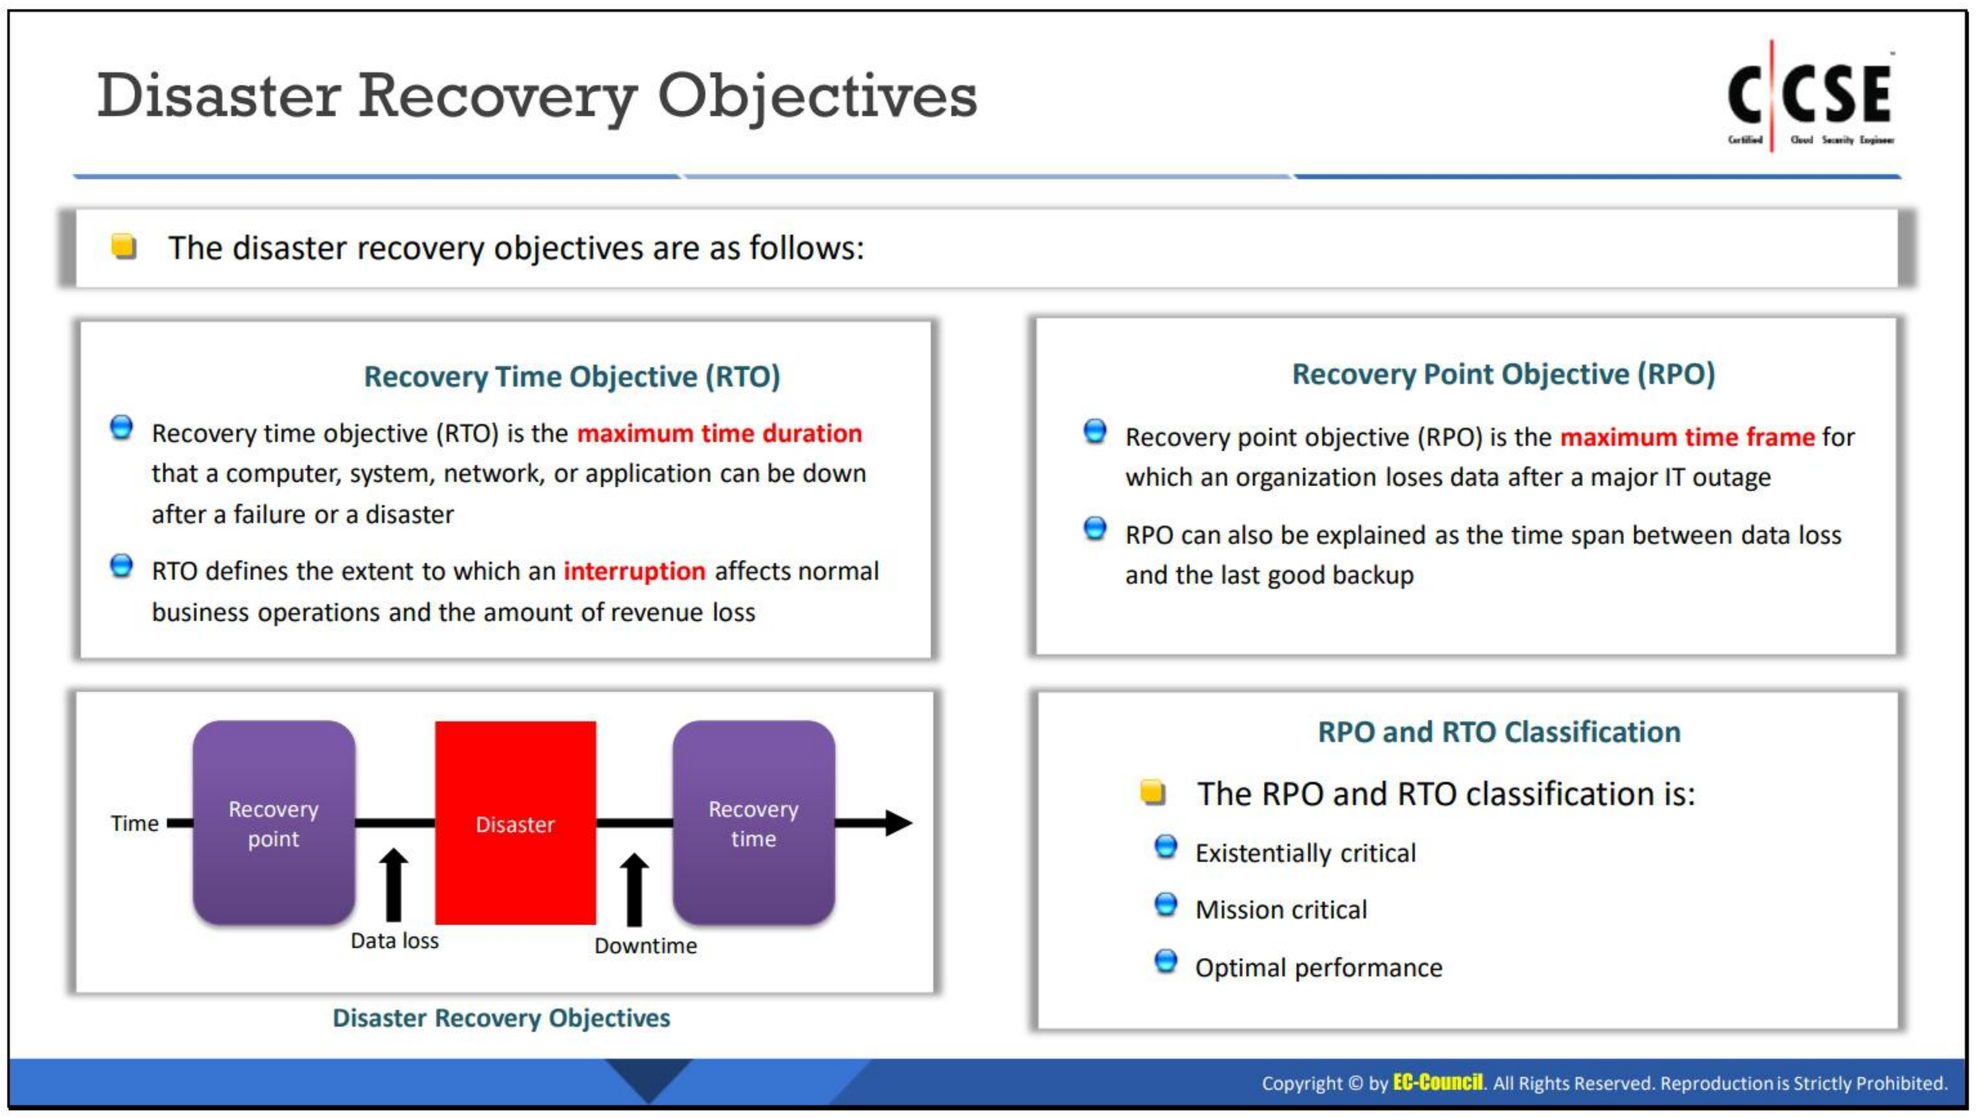

Disaster Recovery Objectives

The objectives of disaster recovery are described below.

Recovery Point Objective (RPO)

RPO is the maximum time frame for which an organization loses data after a major IT outage. RPO can also be explained as the time span between data loss and the last good backup. Every organization must calculate how long it can operate without the required data before a failure.

The factors that affect RPO are as follows:

- The maximum data loss that the organization can handle.
- Industry-specific factors such as the businesses handling sensitive information.
- Cloud data storage helps in faster recovery of the data.
- Financial loss due to data loss and disruption in operations.
- Compliance schemes include provisions for disaster recovery and data loss.

Recovery Time Objective (RTO)

RTO is the maximum tolerable length of time that a computer, system, network, or application can be down after a failure or a disaster. RTO defines the extent to which an interruption affects normal business operations and the amount of revenue loss due to such an interruption. RTO is expressed in minutes. For example, an RTO of 45 minutes implies that the IT operations must be restarted within 45 minutes.

Certified Cloud Security Engineer Copyright © by **EC-Council**
All Rights Reserved. Reproduction is Strictly Prohibited.

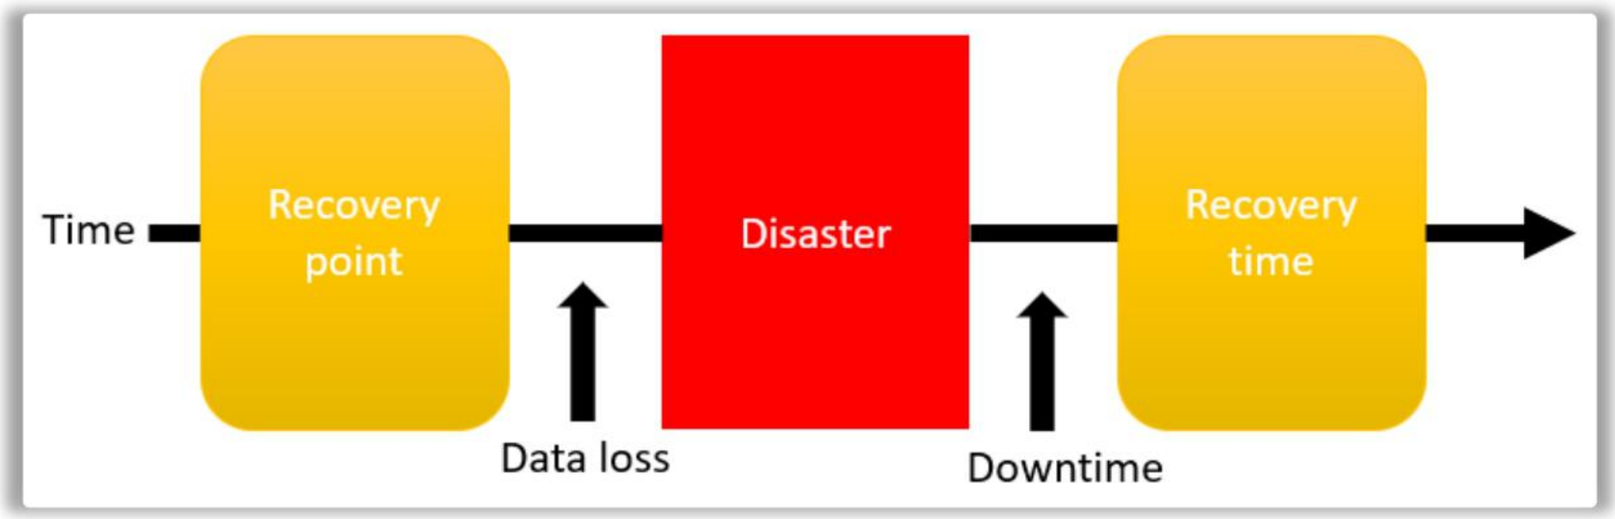

Figure 9.1: Disaster Recovery Objectives

RTO and RPO in the Scope of Cloud

To review RPO and RTO, it is important to classify and determine the applications based on criticality.

- **Existentially critical:** Organizations cannot exist if such first-tier systems are compromised.
- **Mission-critical:** If such second-tier systems are compromised, it will affect employee productivity, brand reputation, and financial loss.
- **Optimal performance:** If these third-tier systems are compromised, the overall organizational efficiency will be reduced.

SaaS platforms have certain protections for critical operations; however, RPO and RTO should be established to recover from errors or malwares that may corrupt the data.

ride1.ir

Certified Cloud Security Engineer Copyright © by **EC-Council**
All Rights Reserved. Reproduction is Strictly Prohibited.

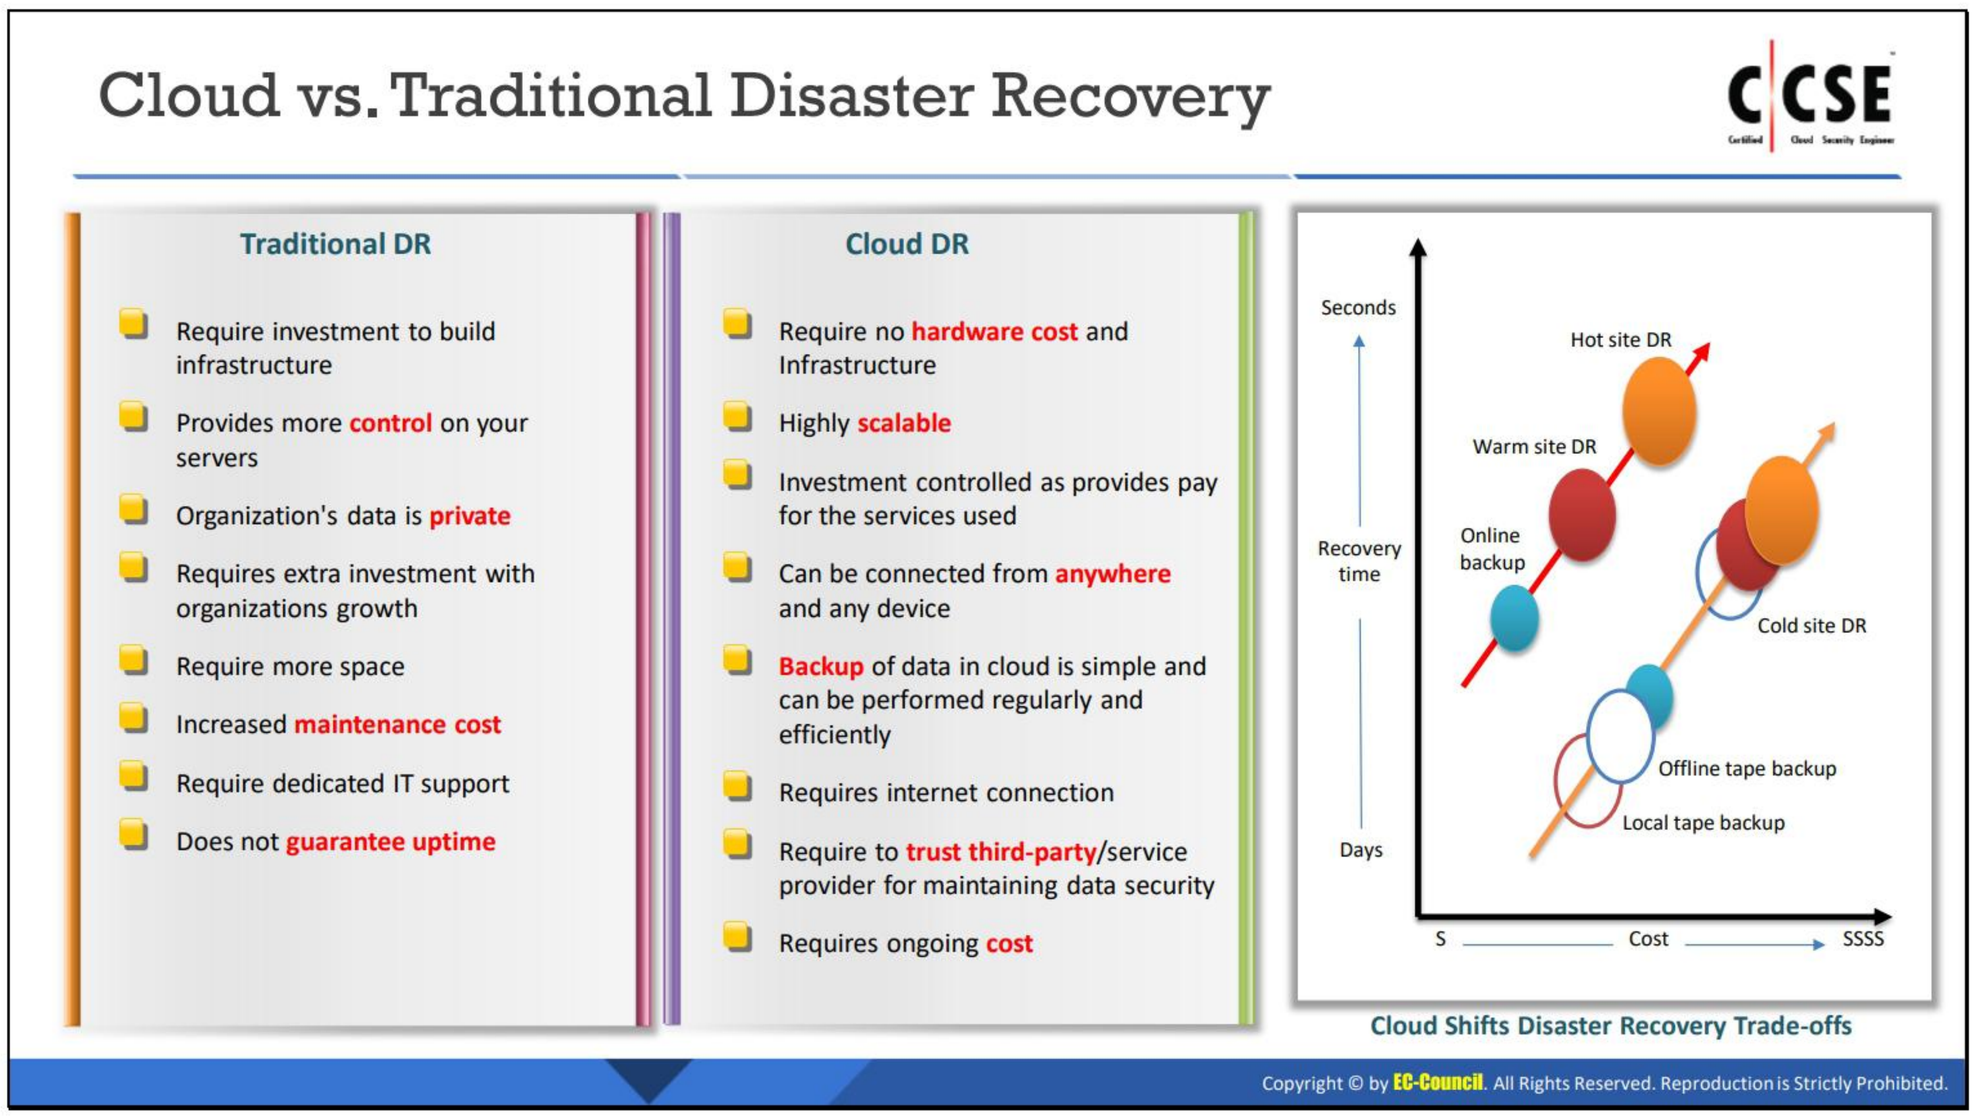

Cloud vs Traditional Disaster Recovery

Traditional disaster recovery consists of a secondary data center for data backup. Traditional disaster recovery consists of a large number of servers, better internet connectivity, and firewalls, switches, and routers for a reliable network connection. The benefits of CDR over traditional disaster recovery includes the following:

- It does not require the user to build any kind of physical infrastructure. Cloud storage access is sufficient for a disaster recovery site.
- CDR is scalable as resources can be added based upon the business requirements.
- Cloud provides' pay as you go pricing model reduces costs.
- CDR can be performed from any location.
- Single points of failures are eliminated by storing the backup data across different geographical locations.
- CDR includes state of the art infrastructure where issues are resolved quickly, and it ensures 24×7 support.

Traditional disaster recovery consists of the user setting up a physical secondary data center for disaster recovery, which needs to be provided with backup power, cooling systems, and on-site maintenance. With CDR the cloud service provider (CSP) offers all those functionalities, and the user only needs to configure the services for performing disaster recovery.

Certified Cloud Security Engineer Copyright © by **EC-Council**
All Rights Reserved. Reproduction is Strictly Prohibited.

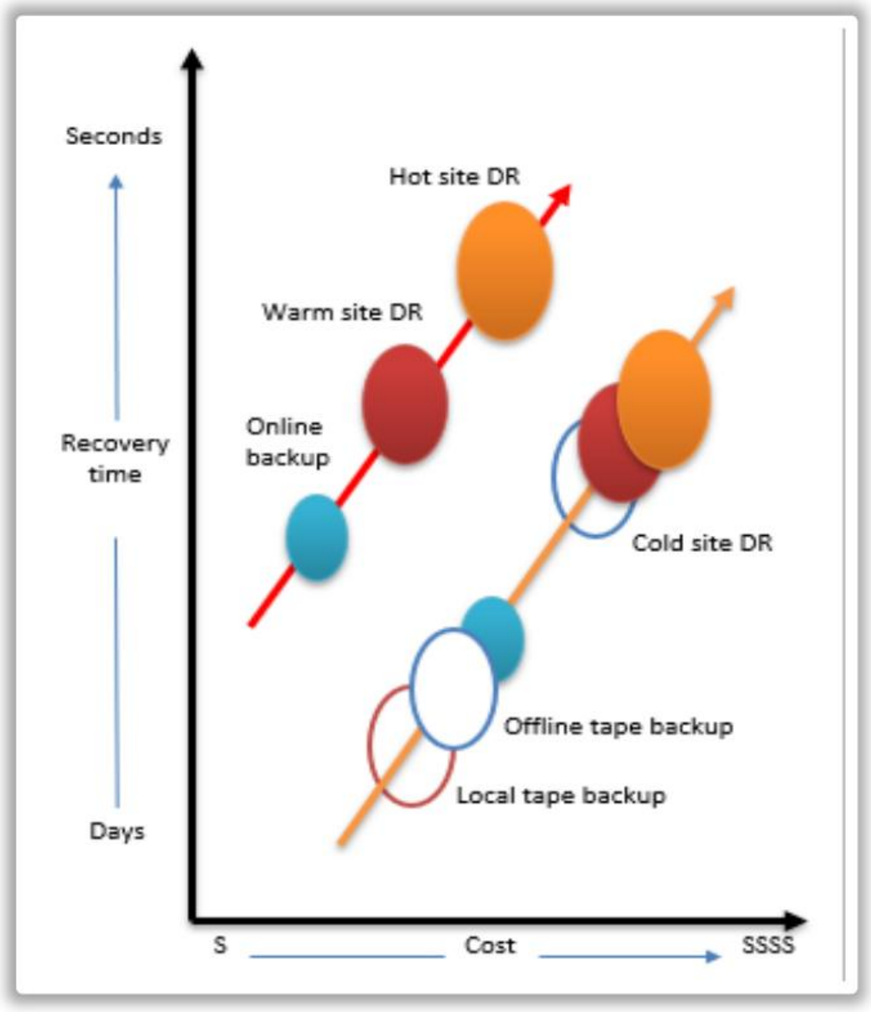

Figure 9.2: Cloud Shifts Disaster Recovery Tradeoffs

Traditional Disaster Recovery

- Requires investment to build infrastructure.
- Provides more control on your servers.
- Organization's data is private.
- Requires extra investment with organizations growth.
- Requires more space.
- Increased maintenance cost.
- Require dedicated IT support.
- Does not guarantee uptime.

hide01.ir

Cloud Disaster Recovery

- Requires no hardware cost and infrastructure.
- Highly scalable.
- Investment controlled as provides pay for the services used.
- Can be connected from anywhere and any device.
- Backup of data in cloud is simple and can be performed regularly and efficiently.
- Requires internet connection.
- Require trusting third-party/service provider for maintaining data security.
- Requires ongoing cost.

Certified Cloud Security Engineer Copyright © by **EC-Council**
All Rights Reserved. Reproduction is Strictly Prohibited.

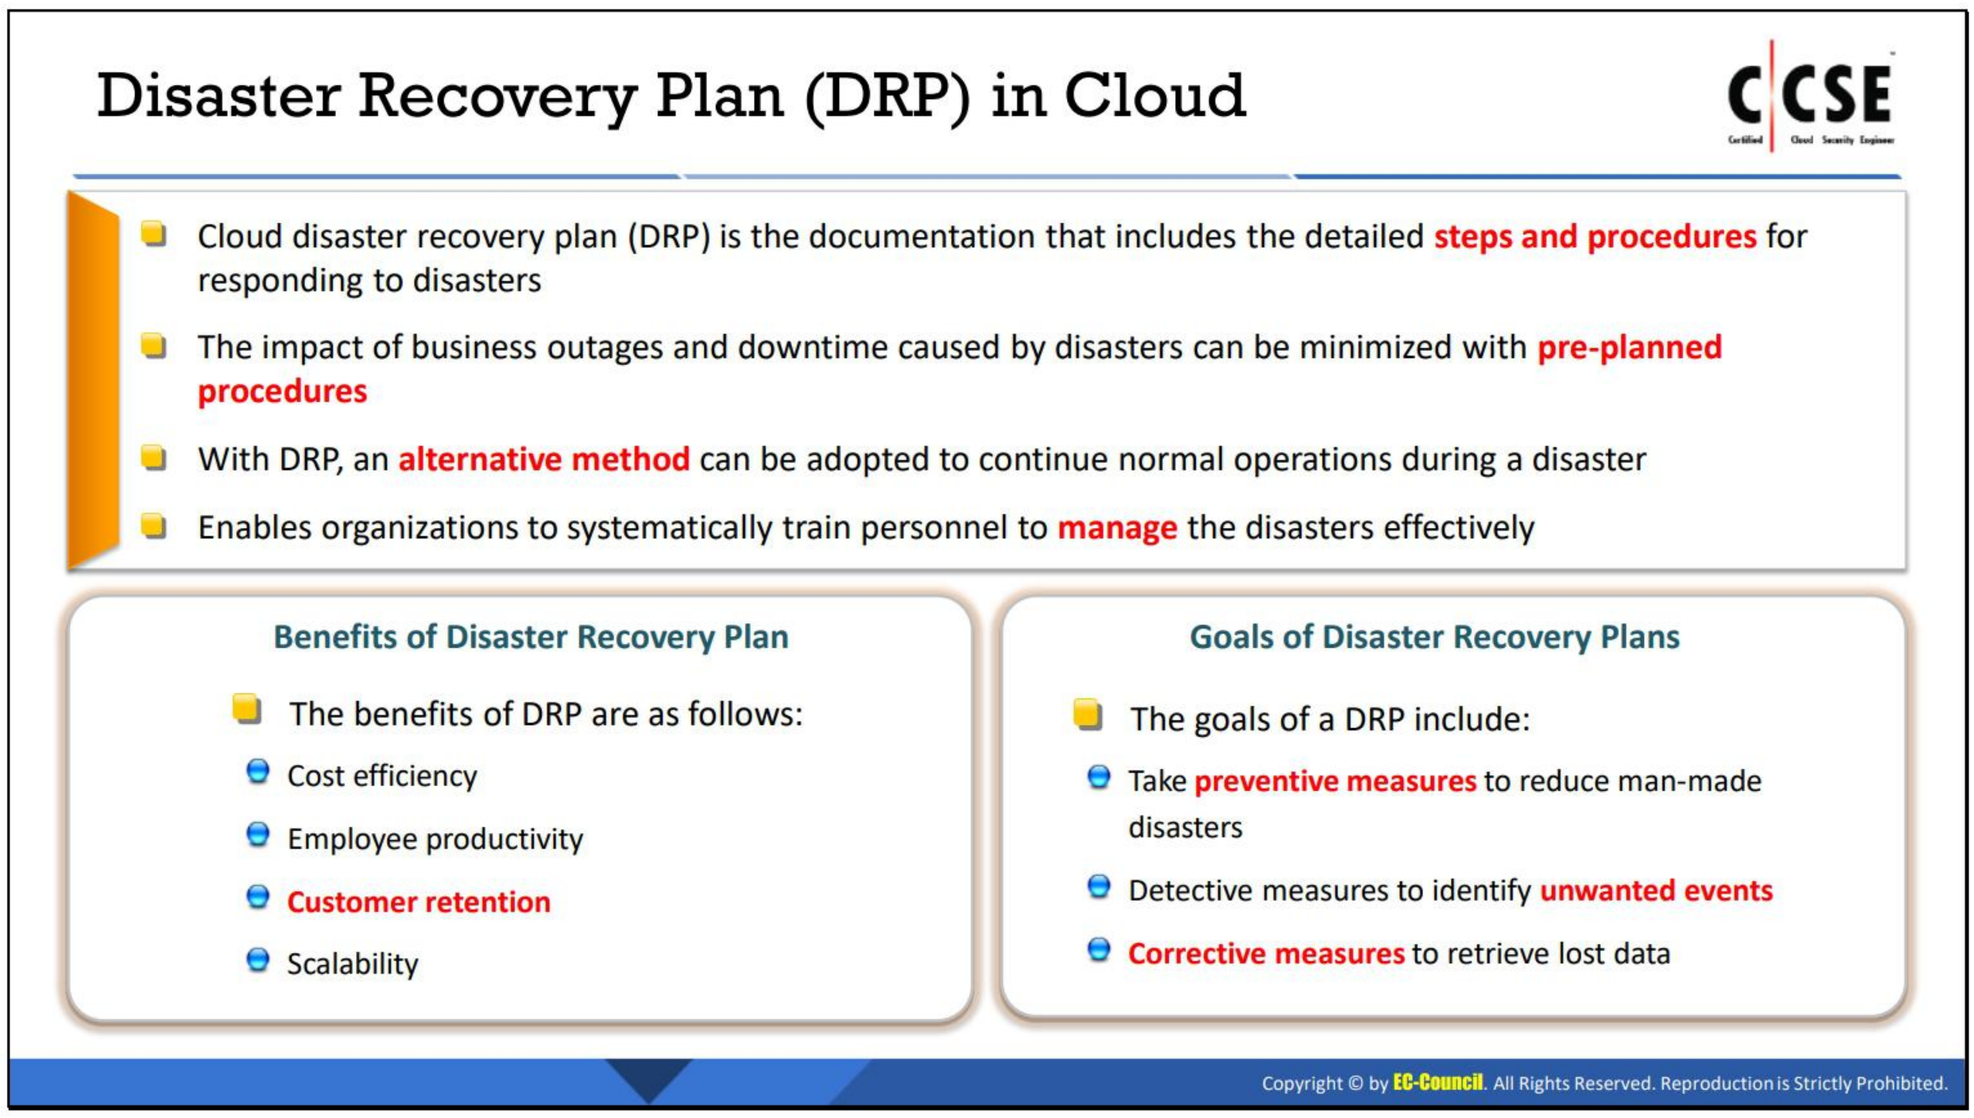

Disaster Recovery Plan (DRP) in the Cloud

What is a Cloud Disaster Recovery Plan?

A CDR plan is a detailed document created by the organization that contains the steps and procedures to respond to disasters, outages, and cyberattacks. The DRP contains all the strategies to mitigate a disaster without affecting normal business operations. The CDR plan also includes the recovery plans for any type of natural outage, apart from technical outages.

Importance of a Disaster Recovery Plan in Cloud

The impact of outages and downtime is extremely difficult for businesses to handle. A DRP in a hybrid multi-cloud environment includes skills, strategies, and technologies for the protection of data and continuity of normal operations. Interruptions to normal operations can be minimized. DRP is important because

- The disruptions caused by a disaster can be effectively reduced.
- The financial loss caused by a disaster can be reduced.
- When disaster strikes, an alternative method can be adopted to continue normal operations.
- The staff can be trained to manage disasters.
- The services can be restored rapidly.

Benefits of a Disaster Recovery Plan

- **Cost efficiency:** The DRP includes different aspects such as
    - Prevention of man-made disasters.

Certified Cloud Security Engineer Copyright © by **EC-Council**
All Rights Reserved. Reproduction is Strictly Prohibited.

- Detective measures of unwanted events.
- Corrective measures to retrieve the lost data and restoration of normal operations.

The DRP includes goals for analysis of threats, better performance of the assets, and solutions that focus on BC plans. Thus, with CDR planning, the cost of archive maintenance can be reduced.

- **Employee productivity:** When employees are assigned roles and responsibilities to be followed during a disaster, it helps in increasing overall employee productivity.
- **Customer retention:** With a CDR plan, availability can be ensured and downtime can be minimized for the customers because CDR has automatic failover capabilities.
- **Better scalability:** With cloud-based disaster recovery, the effectiveness of backups is better and the cost of disaster recovery is low compared to traditional DRPs.

Goals of Disaster Recovery Plan

The DRP should include the following goals for continuing normal operation after a disaster.

- **Preventive measures to reduce man-made disasters:** Man-made disasters in the cloud environment can be reduced by implementing access for individuals based on principle of least privilege, training for the employees, etc. in the DR plan.
- **Detective measures to identify unwanted events:** The detective measures for DR plan include implementing intrusion prevention systems (IPSs), IDSs, anomaly detectors, etc. to detect security disasters.
- **Corrective measures to retrieve lost data:** DR plan should focus on disaster recovery measures such as backup, failover, etc.

hide01.ir

Certified Cloud Security Engineer Copyright © by **EC-Council**
All Rights Reserved. Reproduction is Strictly Prohibited.

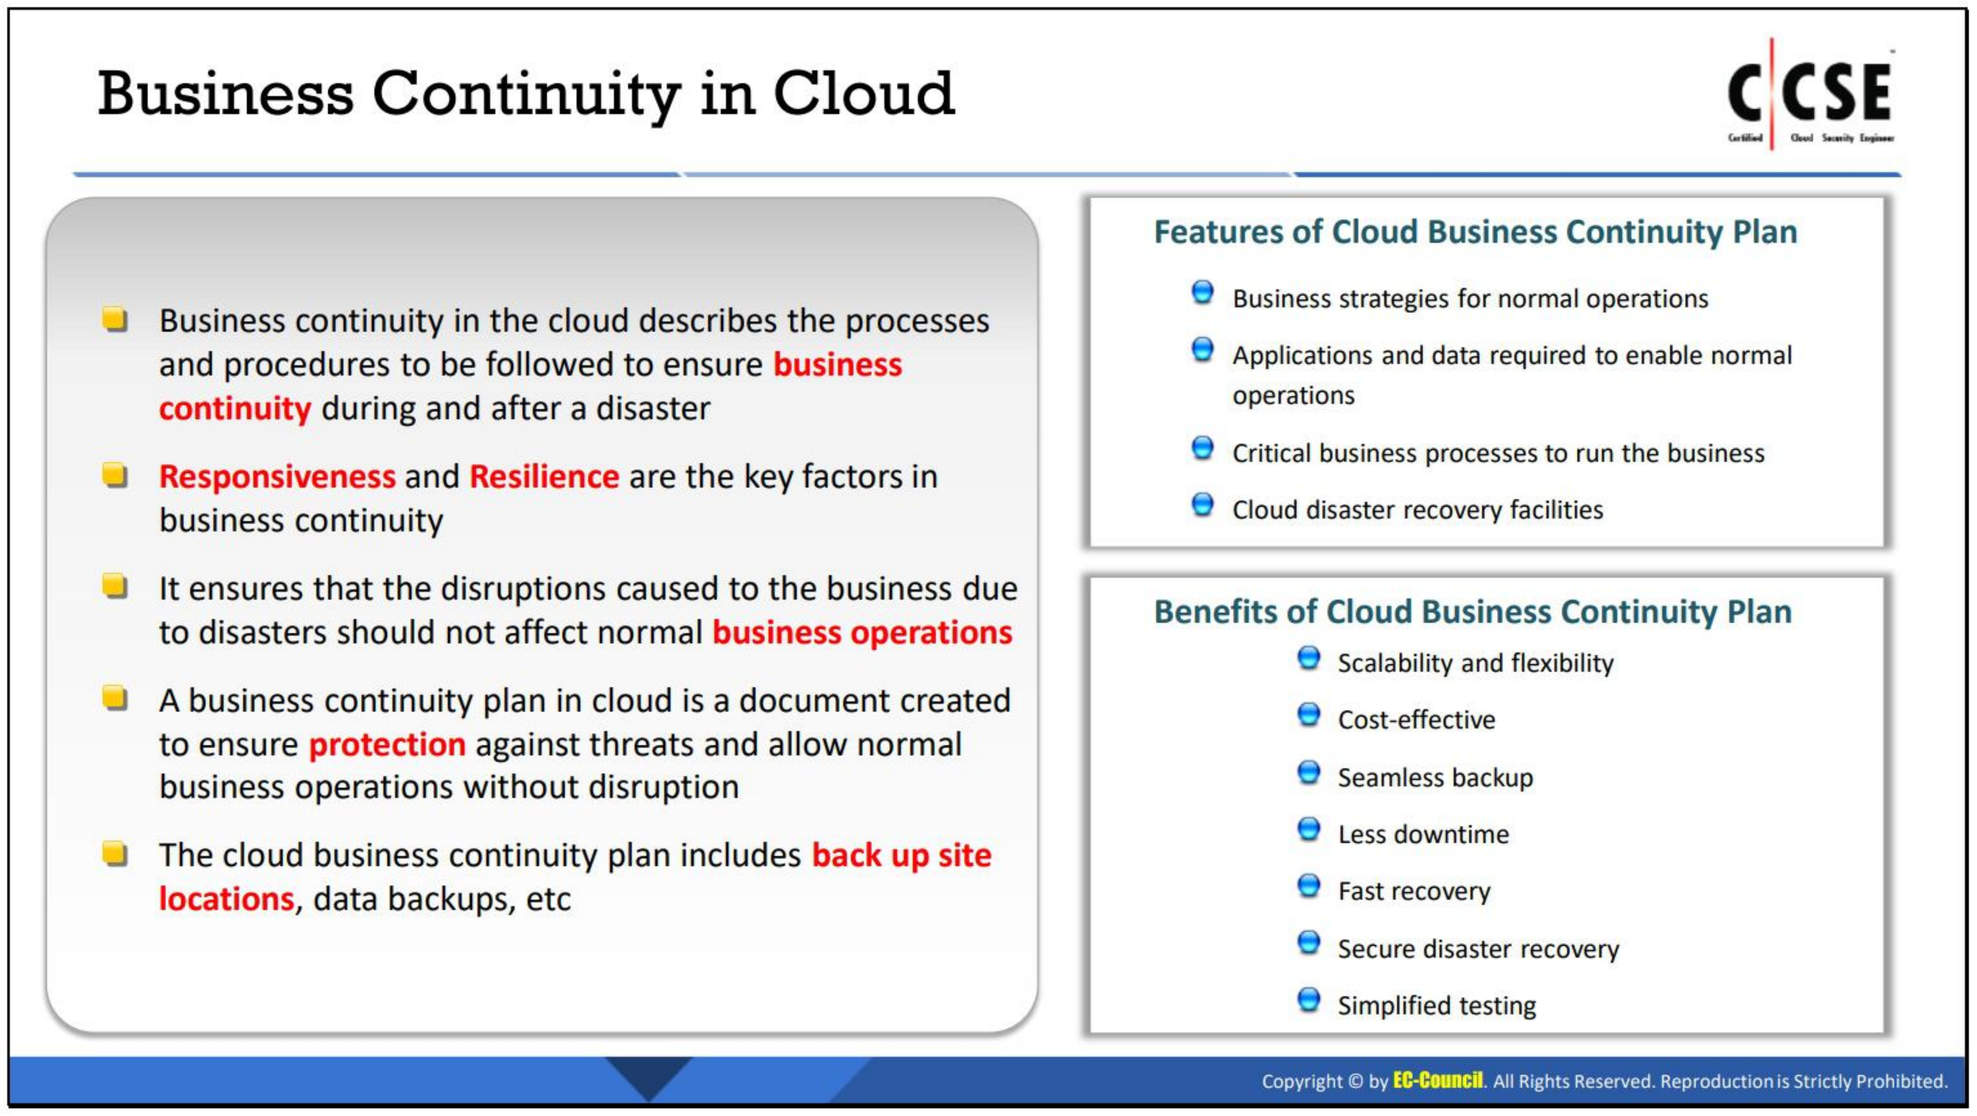

Business Continuity in Cloud

What is BC in Cloud?

BC in cloud describes the processes and procedures that should be followed to ensure the continuity of an organization's critical business functions during and after a disaster. For ensuring BC in cloud, two factors are considered: responsiveness and resiliency.

- **Responsiveness:** Responsiveness is the ability to reduce business impact when a disaster strikes.
- **Resiliency:** Resiliency refers to the ability of the business to recover from the disaster quickly.

Importance of Business Continuity in Cloud

The business operations must withstand the environmental changes and pandemic situations to deliver the best end-user experience. Malicious actors, phishing attacks, ransomware, and advanced persistent threats (APTs) exploit the organization. The disruptions caused due to such issues affects financial growth, customer satisfaction, and brand reputation. An increasing number of systems are now integrated in cloud infrastructure, which increases the number of potential vulnerabilities. If one chain breaks, the organization will be affected. Thus, it is important to develop a secure, robust BC system within the organization.

What is Business Continuity Plan in Cloud?

A BC plan in cloud is a comprehensive document that is formulated to ensure resilience against potential threats and to allow operations to continue under adverse or abnormal conditions. Such a plan includes the backup site locations, data backups, supplies, and equipment. The plan

Certified Cloud Security Engineer Copyright © by **EC-Council**
All Rights Reserved. Reproduction is Strictly Prohibited.

also includes the contact information of the backup site providers, emergency security personnel, etc. The features of a cloud BC plan are described below.

- **Strategy**: The entities that are integrated with business strategies for normal operations.
- **Applications and Data**: The entities related to the software to be used to enable daily business operations.
- **Processes**: The entities that are related to critical business process to run the business.
- **Facilities**: The entities related to the CDR provider if the primary site is destroyed.

Benefits of a Cloud Business Continuity Plan

- **Scalability and flexibility**: The BC plan in cloud can be designed based on the required services and can be scaled up and down depending on changing business needs. It also offers scalability where the most critical services can be stored in the cloud and the less critical ones on in-house servers.
- **Cost-effective**: Cloud-based BC plan is cost-effective as it does not need any additional infrastructure such as hardware, servers, periodic maintenance, etc.
- **Seamless backup**: Cloud-based data backup is reliable and it eliminates any single point of failure as the data is stored at multiple places in the cloud. Data is automatically backed-up to the data centers distributed globally without any harm to the business.
- **Less downtime**: Cloud-based disaster recovery allows configuration of faster recovery times, thereby offering minimal down time.
- **Fast recovery**: As the data can be replicated to multiple locations, the recovery of the data is easy when a disaster strikes.
- **Secure disaster recovery**: Access is limited through authentication procedures, and data is encrypted at storage or in transit.
- **Simplified testing**: The complexity in testing the DR resources is eliminated as the cloud completely automates the process by using the cloud providers network, computer, and storage resources.

hidoc1.ir

Certified Cloud Security Engineer Copyright © by **EC-Council**
All Rights Reserved. Reproduction is Strictly Prohibited.

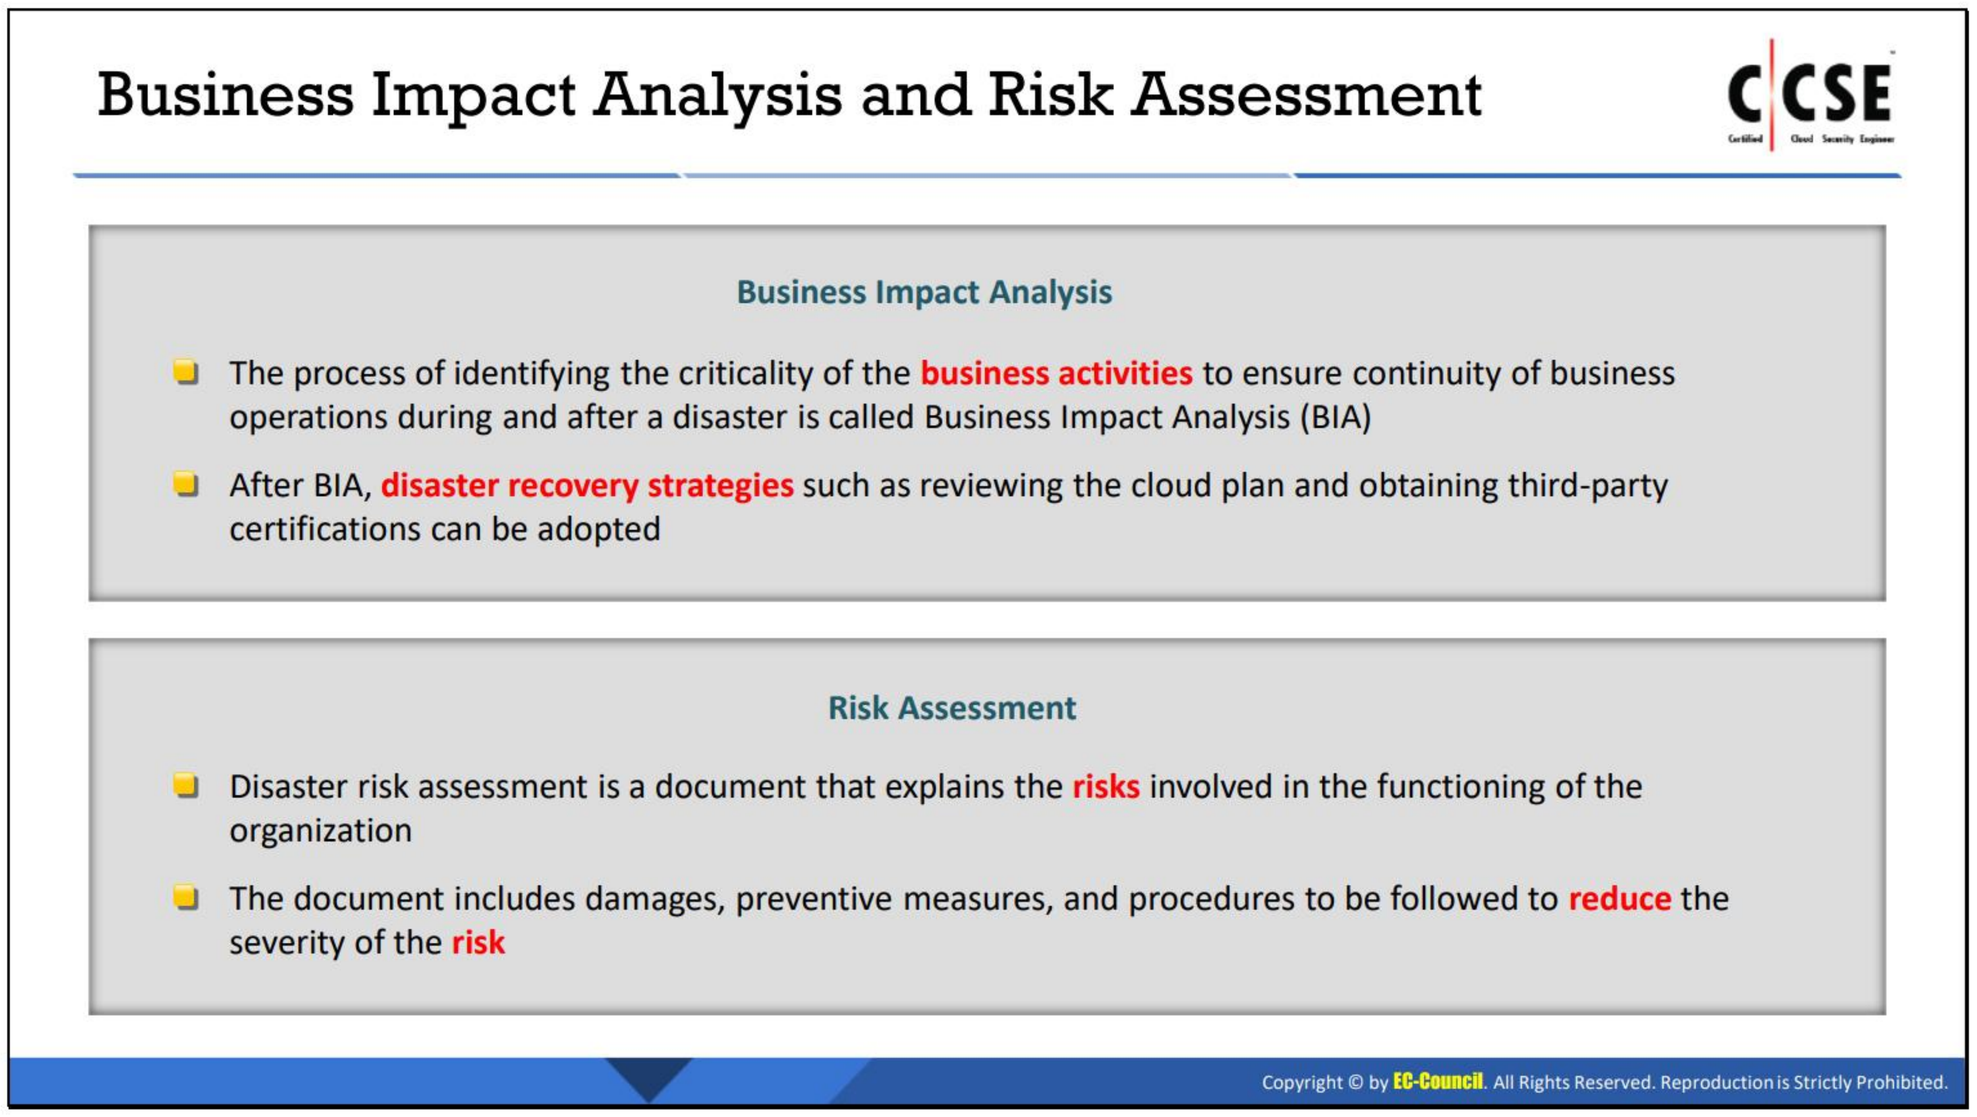

Business Impact Analysis and Risk Assessment

Business Impact Analysis

The process of identifying the criticality of the business activities to ensure continuity of business operations during and after a disaster is called business impact analysis (BIA). In the cloud, applications are purchased from third-party providers and not on application listings. These services should be identified, and the criticality of these services should be determined based on the requested RTO.

Recovery Strategies after Business Impact Analysis in Cloud

After determining the critical services during BIA, suitable recovery strategies can be developed. The suitable way for conforming cloud recovery capabilities includes

- Reviewing the cloud plans.
- Ensuring third-party certifications such as ISO 22301.

Plan Development after Business Impact Analysis

- When the application is not accessible, all manual workarounds should be documented.
- Emergency contacts, such as support desk, account manager, website address displaying the service status, should be documented.

Risk Assessment in Cloud Disaster Recovery

The disaster risk assessment is a document that explains about the risks involved in the functioning of the organization. The risk assessment document includes the following:

- The damages the incident may cause.

Certified Cloud Security Engineer Copyright © by **EC-Council**
All Rights Reserved. Reproduction is Strictly Prohibited.

- The time and effort required to mitigate the incident.
- Preventive Measures.
- The steps and procedures to reduce the severity of the incident.

hide01.ir

Certified Cloud Security Engineer Copyright © by **EC-Council**
All Rights Reserved. Reproduction is Strictly Prohibited.

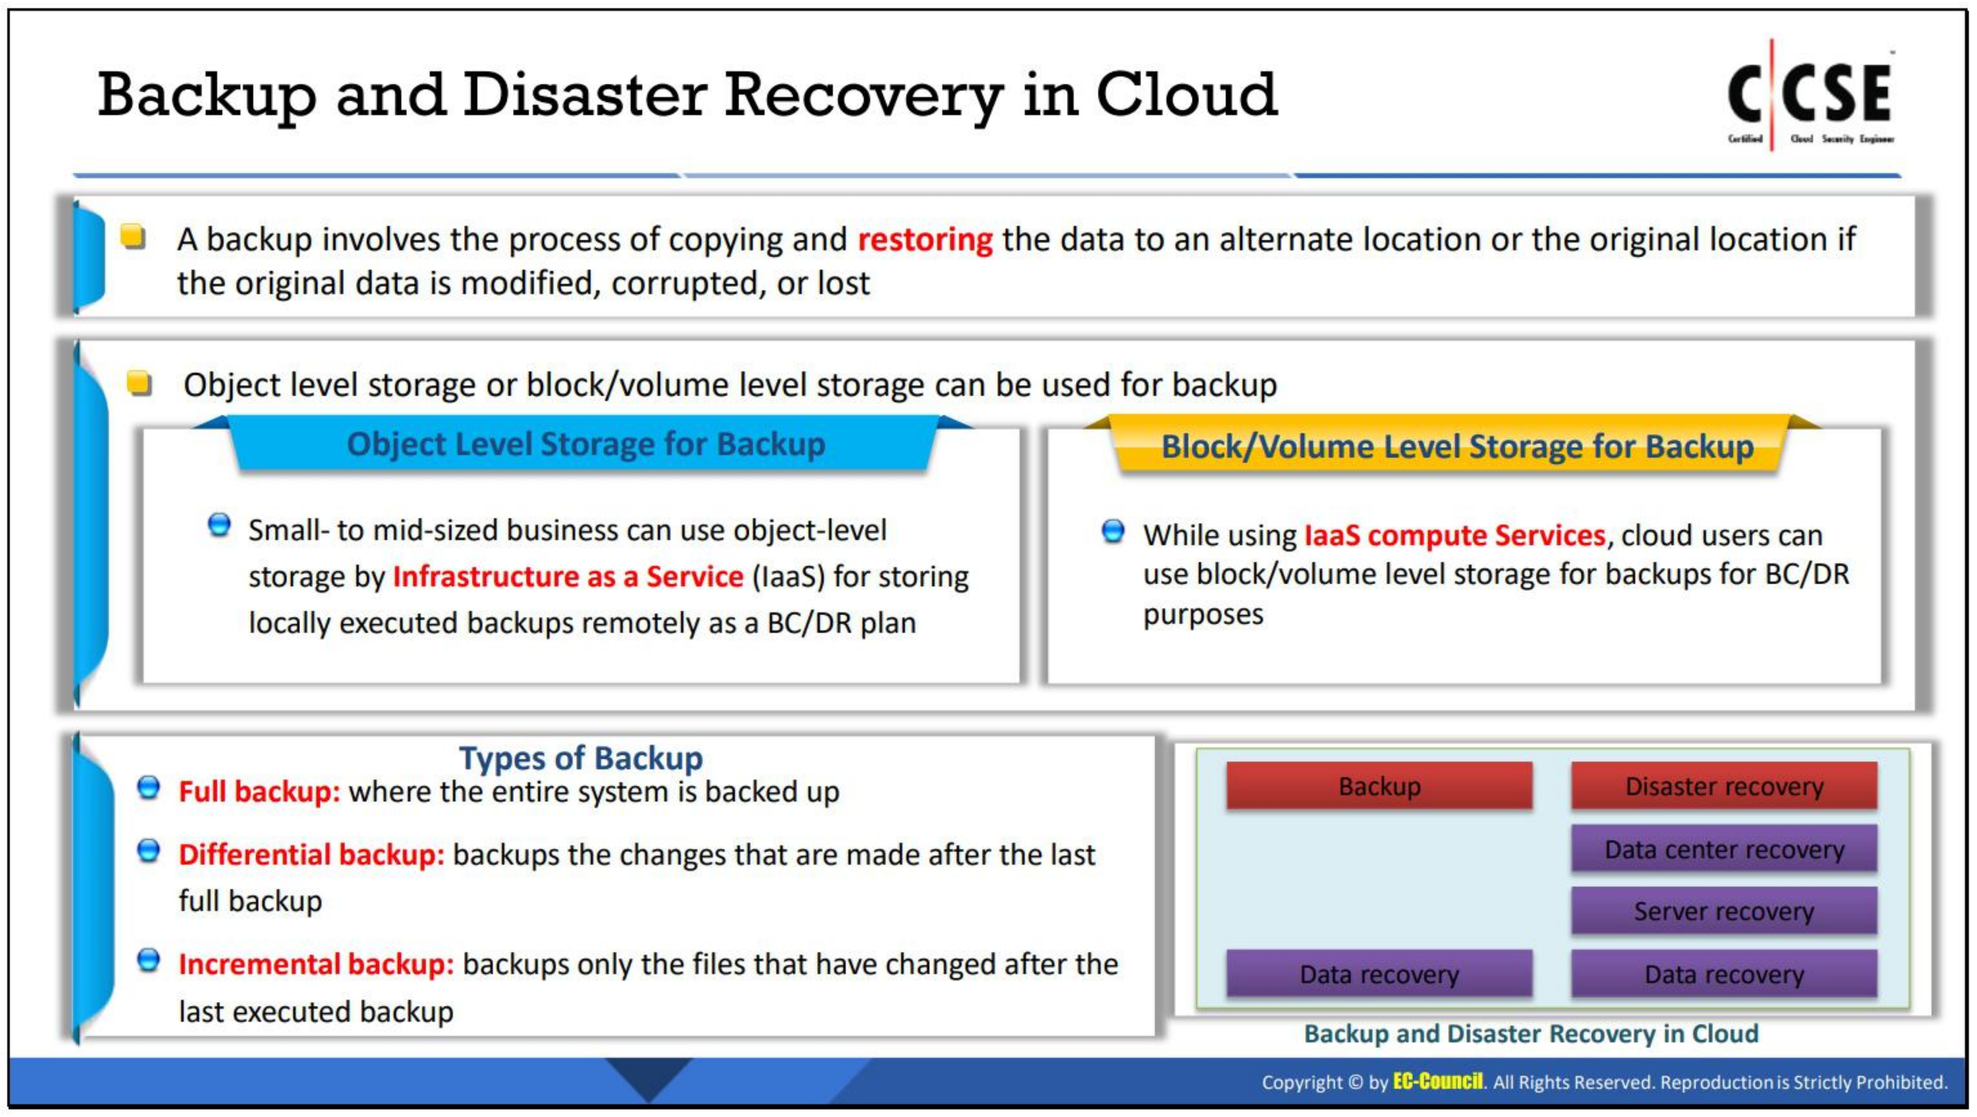

Backup and Disaster Recovery in the Cloud

Backup and Recovery

This includes the process of copying and restoring the data to an alternate location or to the original location if the original data is modified corrupted or lost. The different types of backups in cloud are described below:

- **Full backup:** Here, the entire system is backed-up in a single process so that organizations can take the executed backup and restore it from a single media.
- **Differential backup:** In differential backups, only the changes that are made after the last full backup are executed.
- **Incremental backup:** Incremental backups are faster and include only those files that have changed since the last executed backup.

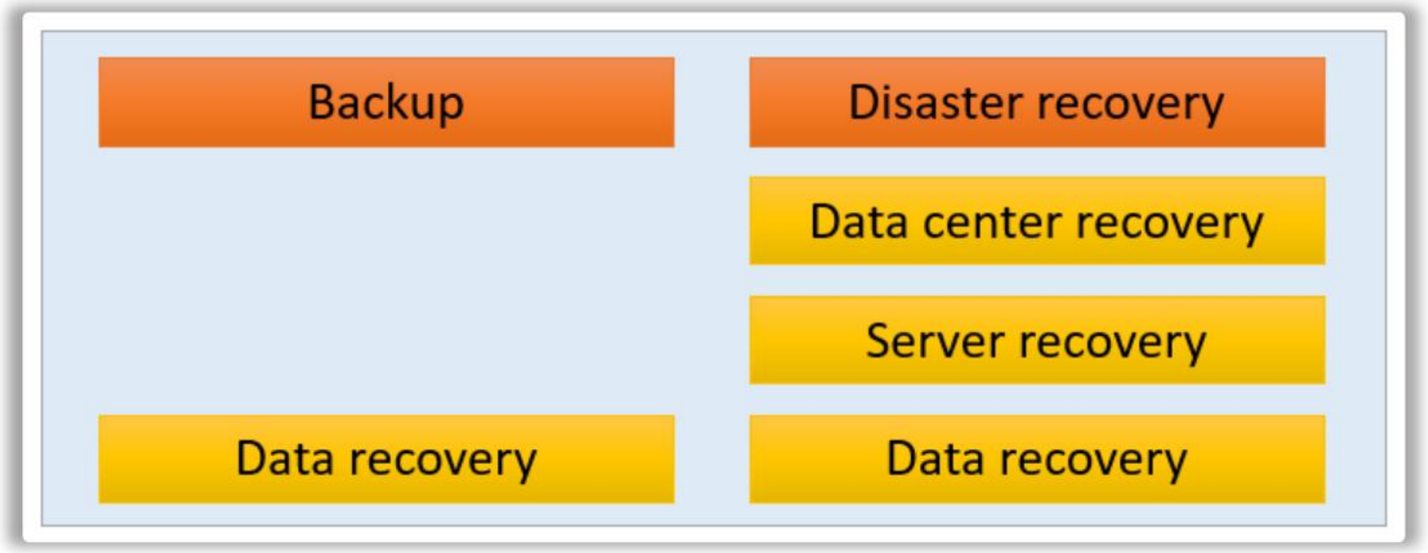

Figure 9.3: Backup and DR in Cloud

Certified Cloud Security Engineer Copyright © by **EC-Council**
All Rights Reserved. Reproduction is Strictly Prohibited.

Types of Services

- **Object-level storage for backups**: Small- to mid-sized business can use object-level storage via IaaS for storing locally executed backups remotely as a BCDR plan. This method provides safe off-site data storage. The concerns with this method are
    - **Replication management:** Backup is static; therefore, it is easier to recover from errors by reverting to the backup. However, the replication snapshots should be taken at specified intervals for reversion.
    - **Legal and regulatory requirements:** The legal and regulatory requirements for data retention should be followed.
- **Block/Volume-level storage for backup:** While using IaaS compute services, cloud users can use block/volume-level storage for backup for BCDR purposes. Block-level replication requires only a small-bandwidth link to the CSP because it is bandwidth efficient.

hide01.ir

Certified Cloud Security Engineer Copyright © by **EC-Council**
All Rights Reserved. Reproduction is Strictly Prohibited.

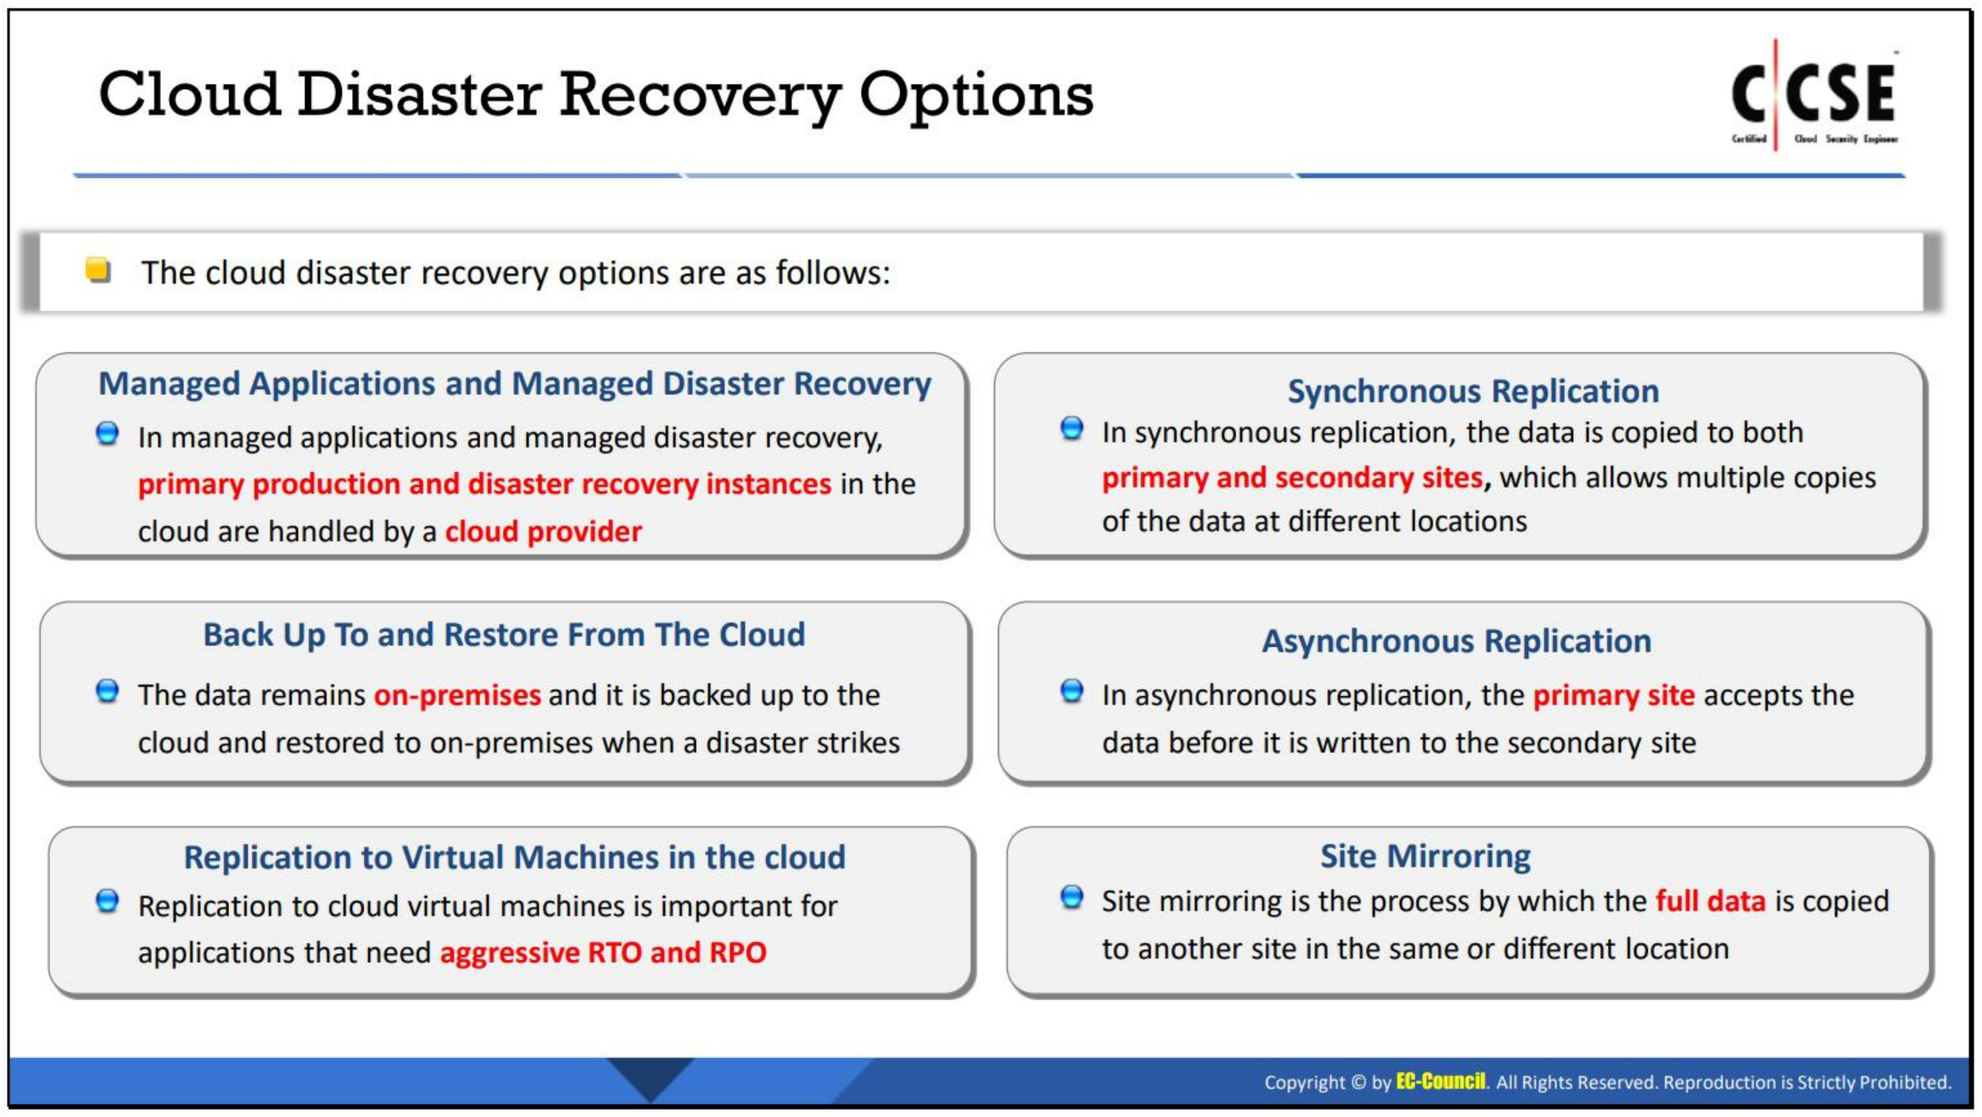

Cloud Disaster Recovery Options

The CDR options are as follows:

- **Managed applications and managed disaster Recovery**: In managed applications and managed DR, primary production and disaster recovery instances in the cloud are handled by a cloud provider. When the control is hand overed to the cloud provider, it should be ensured that the provider delivers services without any availability problems within the defined service level agreements (SLAs).
- **Backup to and restore from the cloud**: In this approach, the data remains on the premises and it is backed-up to the cloud and restored to the premises when a disaster strikes.
- **Backup to and restore to the cloud**: In this option, the data is restored to VMs in the cloud. The restoration of data occurs when a disaster strikes. For continuous restorations, the VMs should be up to date; this helps in meeting the determined RTO.

Replication to VMs in the Cloud

Replication to cloud VMs is important for applications that need aggressive RTO and RPO. Replication is supported for both cloud VM-to-cloud VM and also on-premises-to-cloud VM.

Replication

The fault tolerance and reliability of system data can be improved via replication. The two types of replication are synchronous and asynchronous.

- **Synchronous replication:** Here, the data is copied to both primary and secondary sites; this allows multiple copies of the data at different locations. Synchronous replication can affect the performance of the application that is replicated.

Certified Cloud Security Engineer Copyright © by **EC-Council**
All Rights Reserved. Reproduction is Strictly Prohibited.

- **Asynchronous replication:** Here, the primary site accepts the data before it is written to the secondary site. In asynchronous replication, if the primary site fails before the data gets written to the secondary site, the data may be lost.

Site Mirroring

Site mirroring is the process through which the full data is copied to another site in the same, or different, location. In site mirroring, reliability depends on the type of process implemented in the organization.

hide01.ir

Certified Cloud Security Engineer Copyright © by **EC-Council**
All Rights Reserved. Reproduction is Strictly Prohibited.

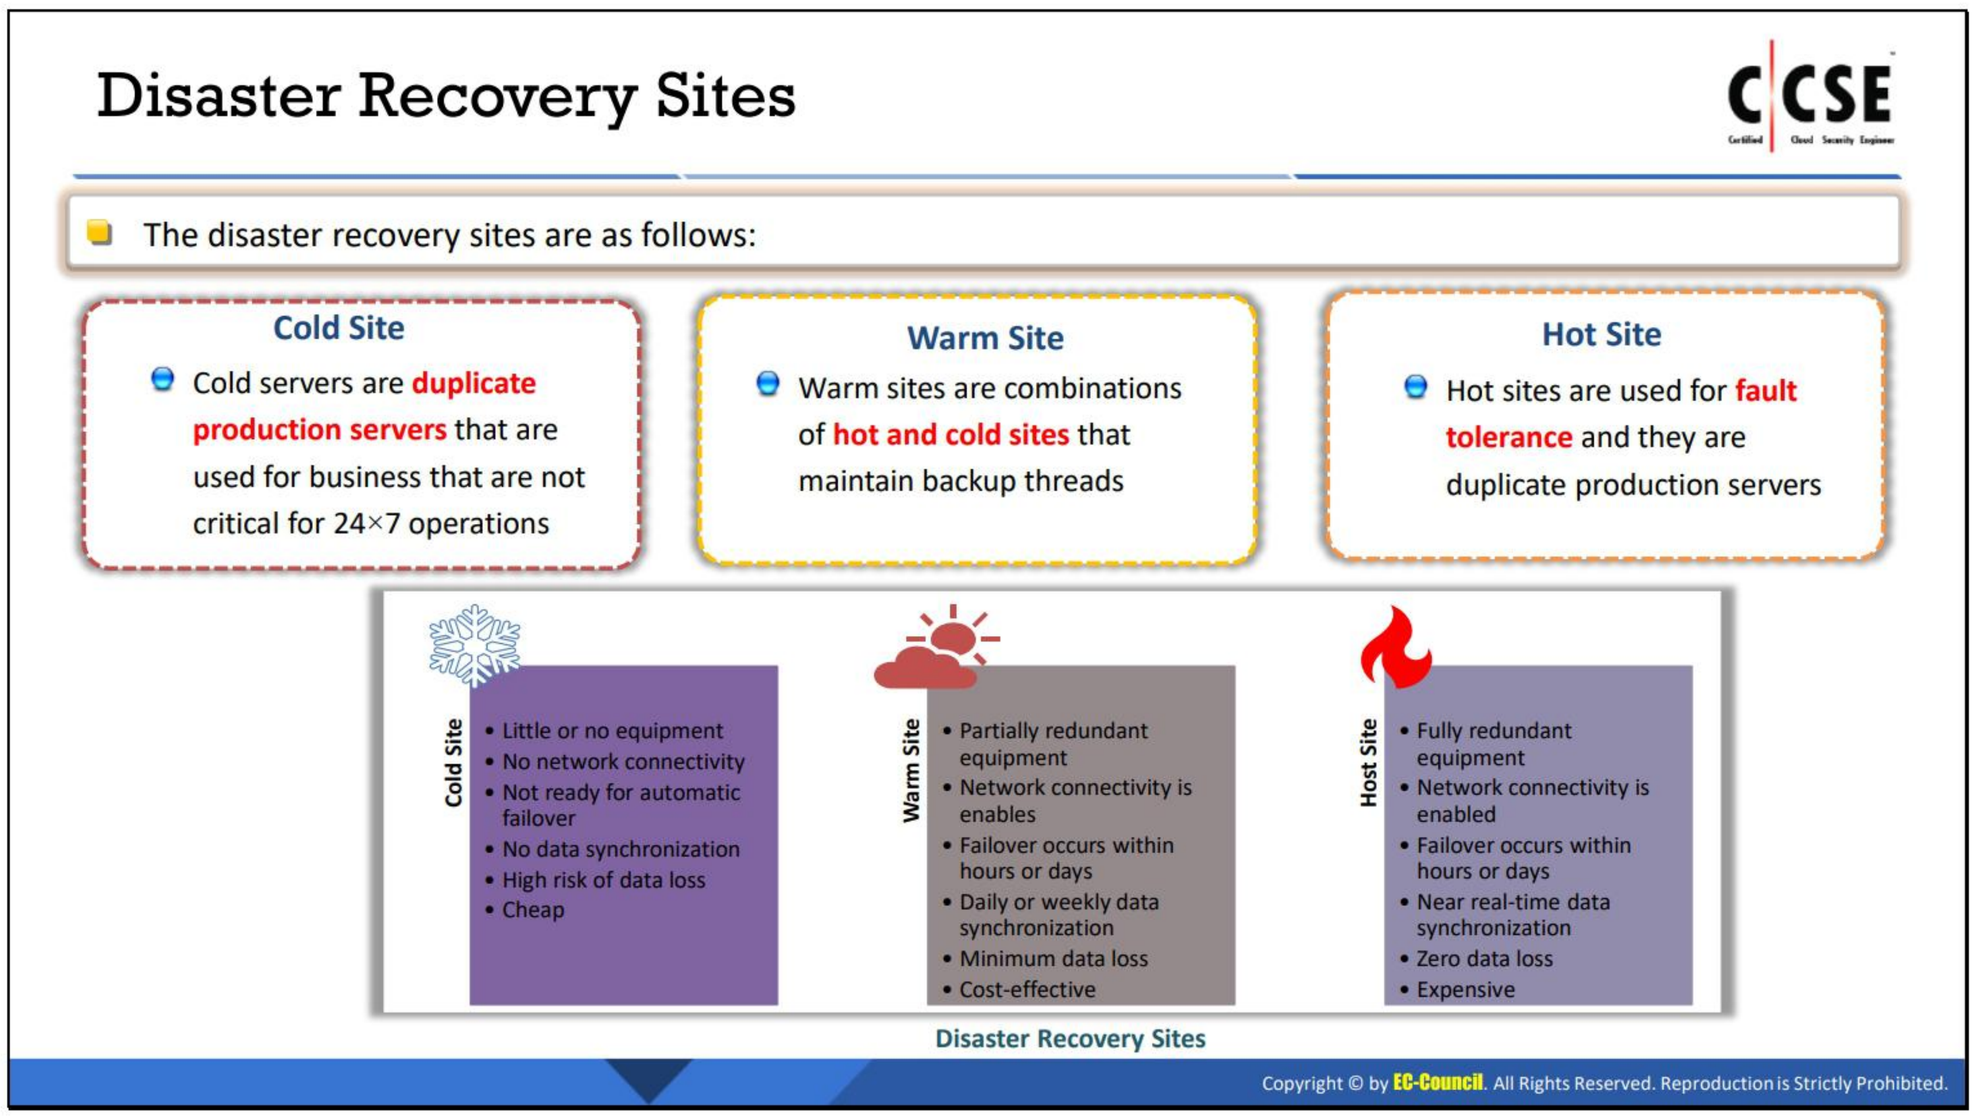

Disaster Recovery Sites

The various disaster sites can be described as follows:

- **Hot sites:** Hot sites are used for fault tolerance. They are duplicate servers of the actual production server. Hot sites are backup services that consist of a mirrored copy of the primary data center. Hot sites are always live and are equipped with all the hardware and software to perform critical operations. When a disaster strikes, failover can be performed to the DR site with minimal downtime and without any data loss.
- **Cold sites**: A cold site is a backup service that does not contain any preinstalled hardware. Cold sites take considerable time during disaster recovery as they require the support of the IT team to make the service functional. Cold sites are backup services with no network connectivity, and they do not perform automatic failover and data synchronization. They are duplicate servers that is not synchronized with the production server 24×7, instead the updates are made at the end of the day. They are useful for businesses that do not require live updating.
- **Warm site**: Warm sites share the features of hot and cold sites and maintain backup threads rather than an updated entire backup. A warm site is a backup service that is installed with all the necessary hardware; however, it cannot function as efficiently as primary data centers. In warm sites, the data is synchronized between the primary and disaster recovery sites daily or weekly, which results in data loss.

Certified Cloud Security Engineer Copyright © by **EC-Council**
All Rights Reserved. Reproduction is Strictly Prohibited.

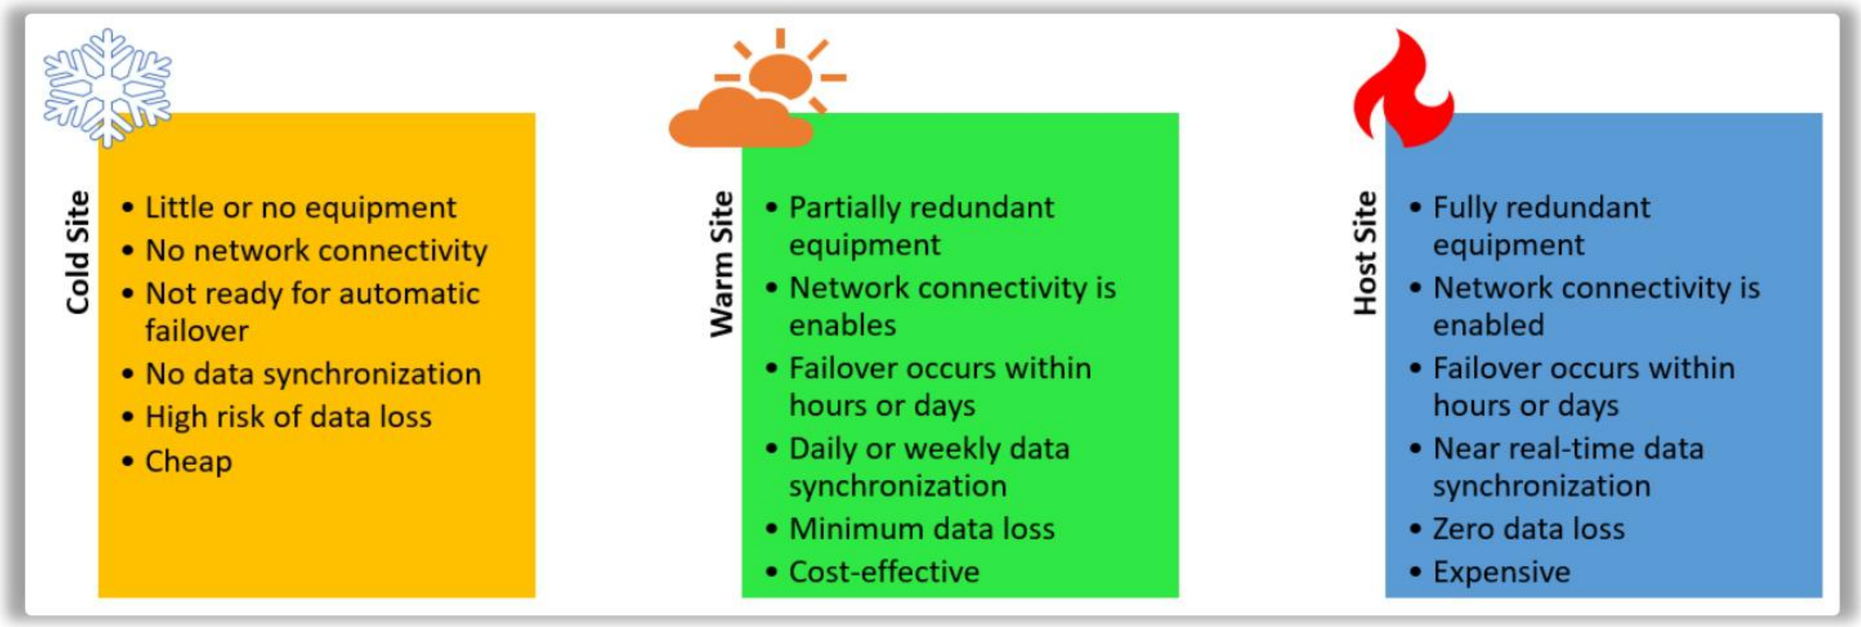

Figure 9.4: Different Disaster Recovery Sites

hide01.ir

Certified Cloud Security Engineer Copyright © by **EC-Council**
All Rights Reserved. Reproduction is Strictly Prohibited.

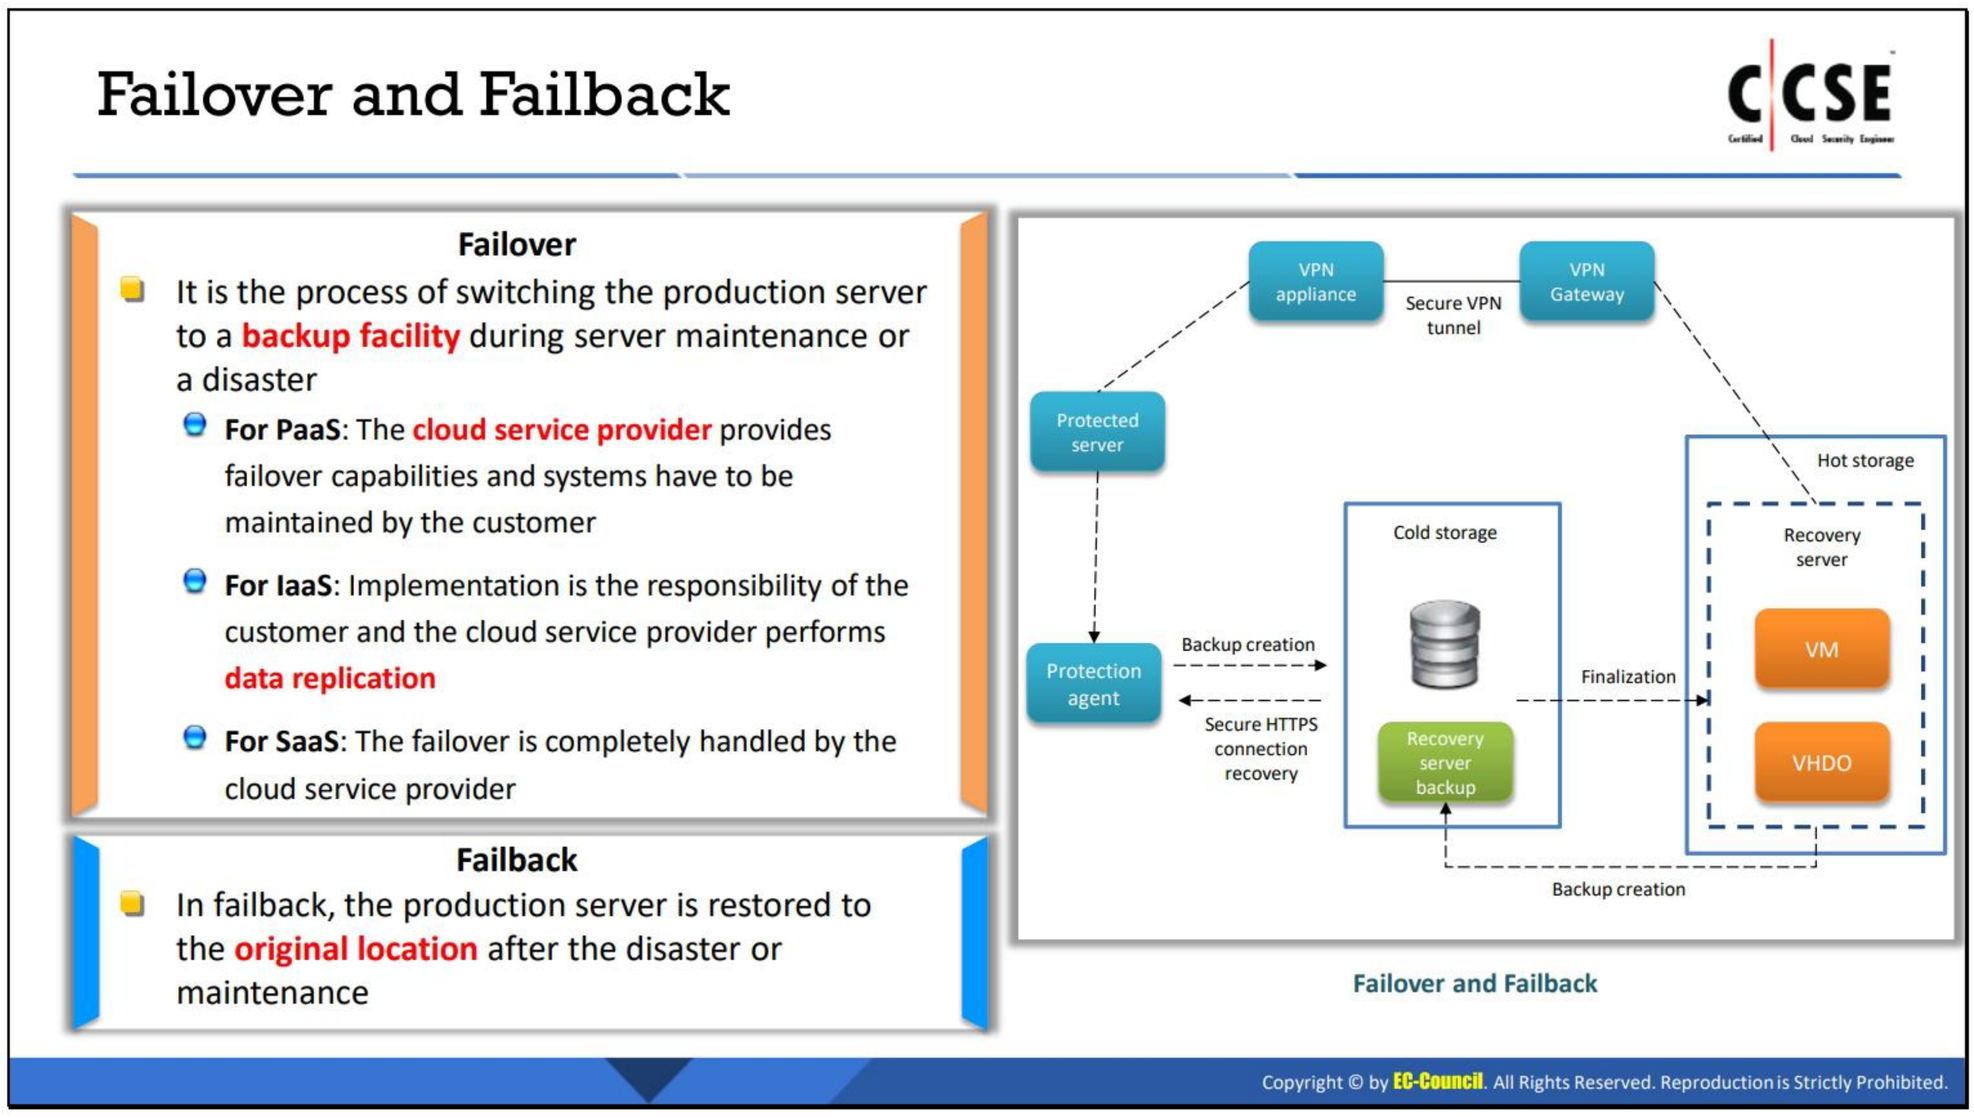

Failover and Failback

Failover

The process of switching the production server to a backup facility, that is, a recovery site for production server maintenance or during a disaster, is called failover.

- **For PaaS**: The CSP will provide failover capabilities and certain systems must be maintained by the customer.
- **For IaaS**: The implementation process is the responsibility of the customer and the CSP performs data replication**.**
- **For SaaS**: Systems in which the failover is completely handled by the CSP.

Failback

Failback is the reverse of failover; here, the production server is restored back to the original location after a disaster or maintenance.

Certified Cloud Security Engineer Copyright © by **EC-Council**
All Rights Reserved. Reproduction is Strictly Prohibited.

How Failover and Failback Work in Cloud Computing

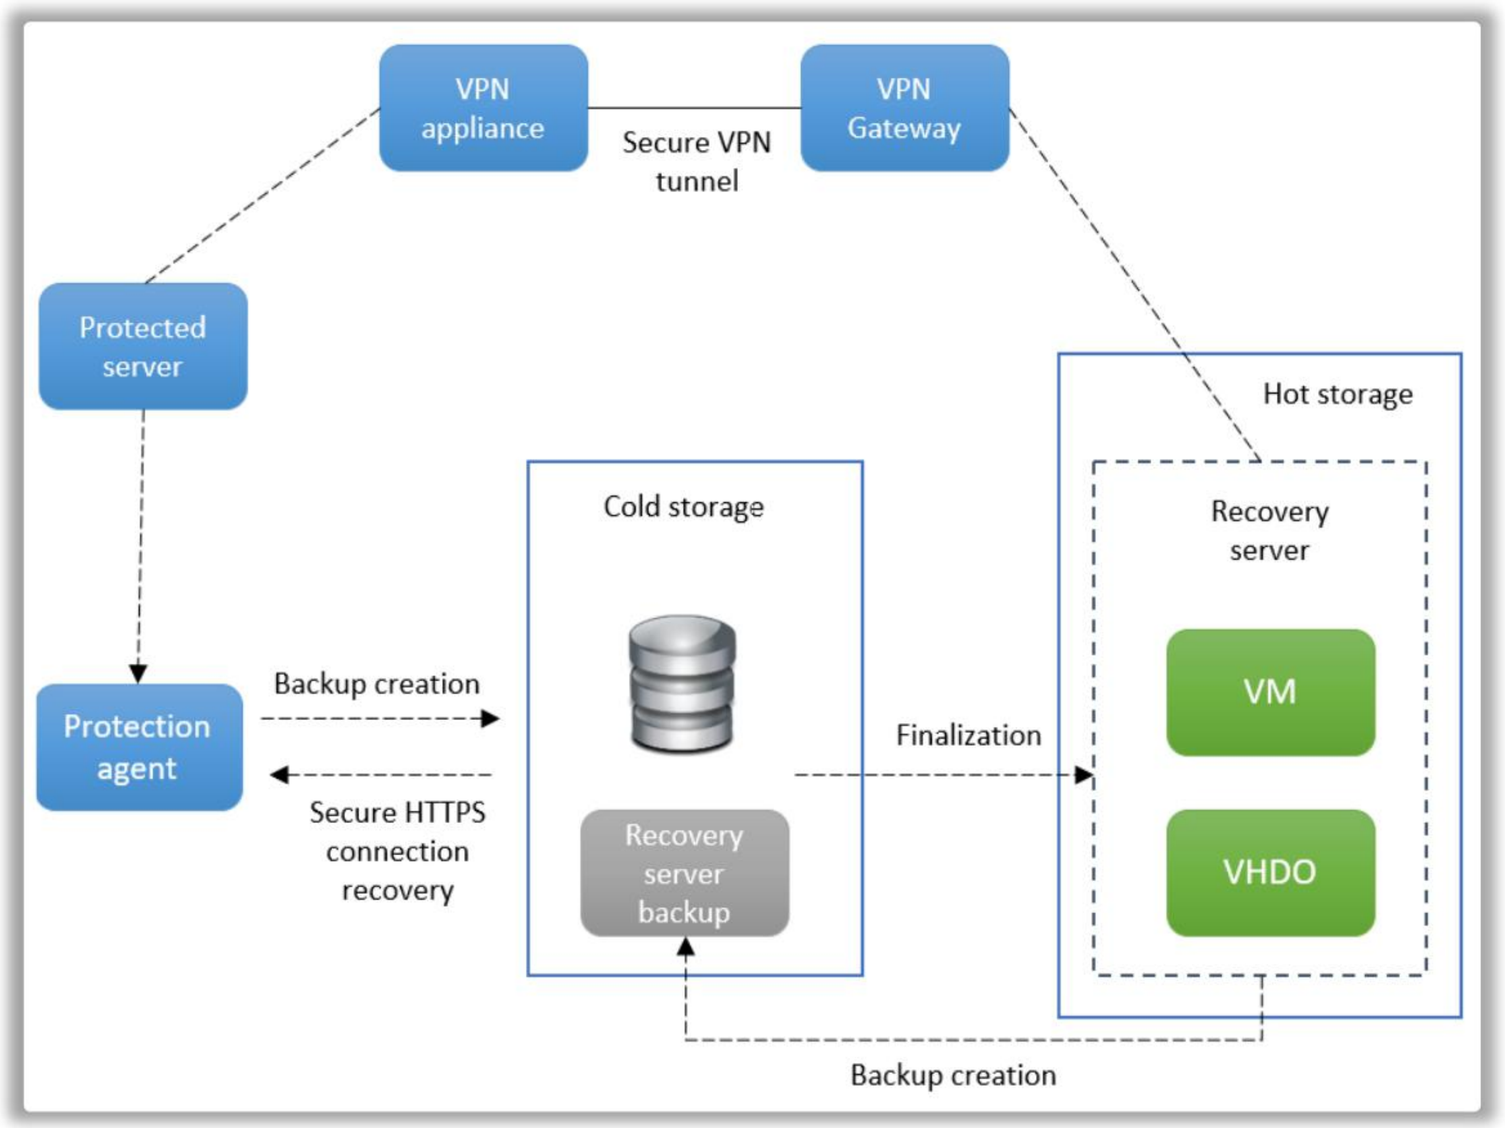

Figure 9.5: Fail Over and Fail Back

When the recovery server is created, it is initially in the standby state. A disk image backup of the original machine is to be created before the failover process initiates.

First, a recovery point of the original machine is to be selected, from which a VM is created. The VM is then used by failover operation. The recovery server moves to the finalization state, where the server's virtual disks from backup storage are transferred to the disaster recovery storage. The server can be operated during the finalization state but with minimal performance. After the finalization is complete, the server performance normalizes and the server this fails over from the original server to the cloud recovery server.

Certified Cloud Security Engineer Copyright © by **EC-Council**
All Rights Reserved. Reproduction is Strictly Prohibited.

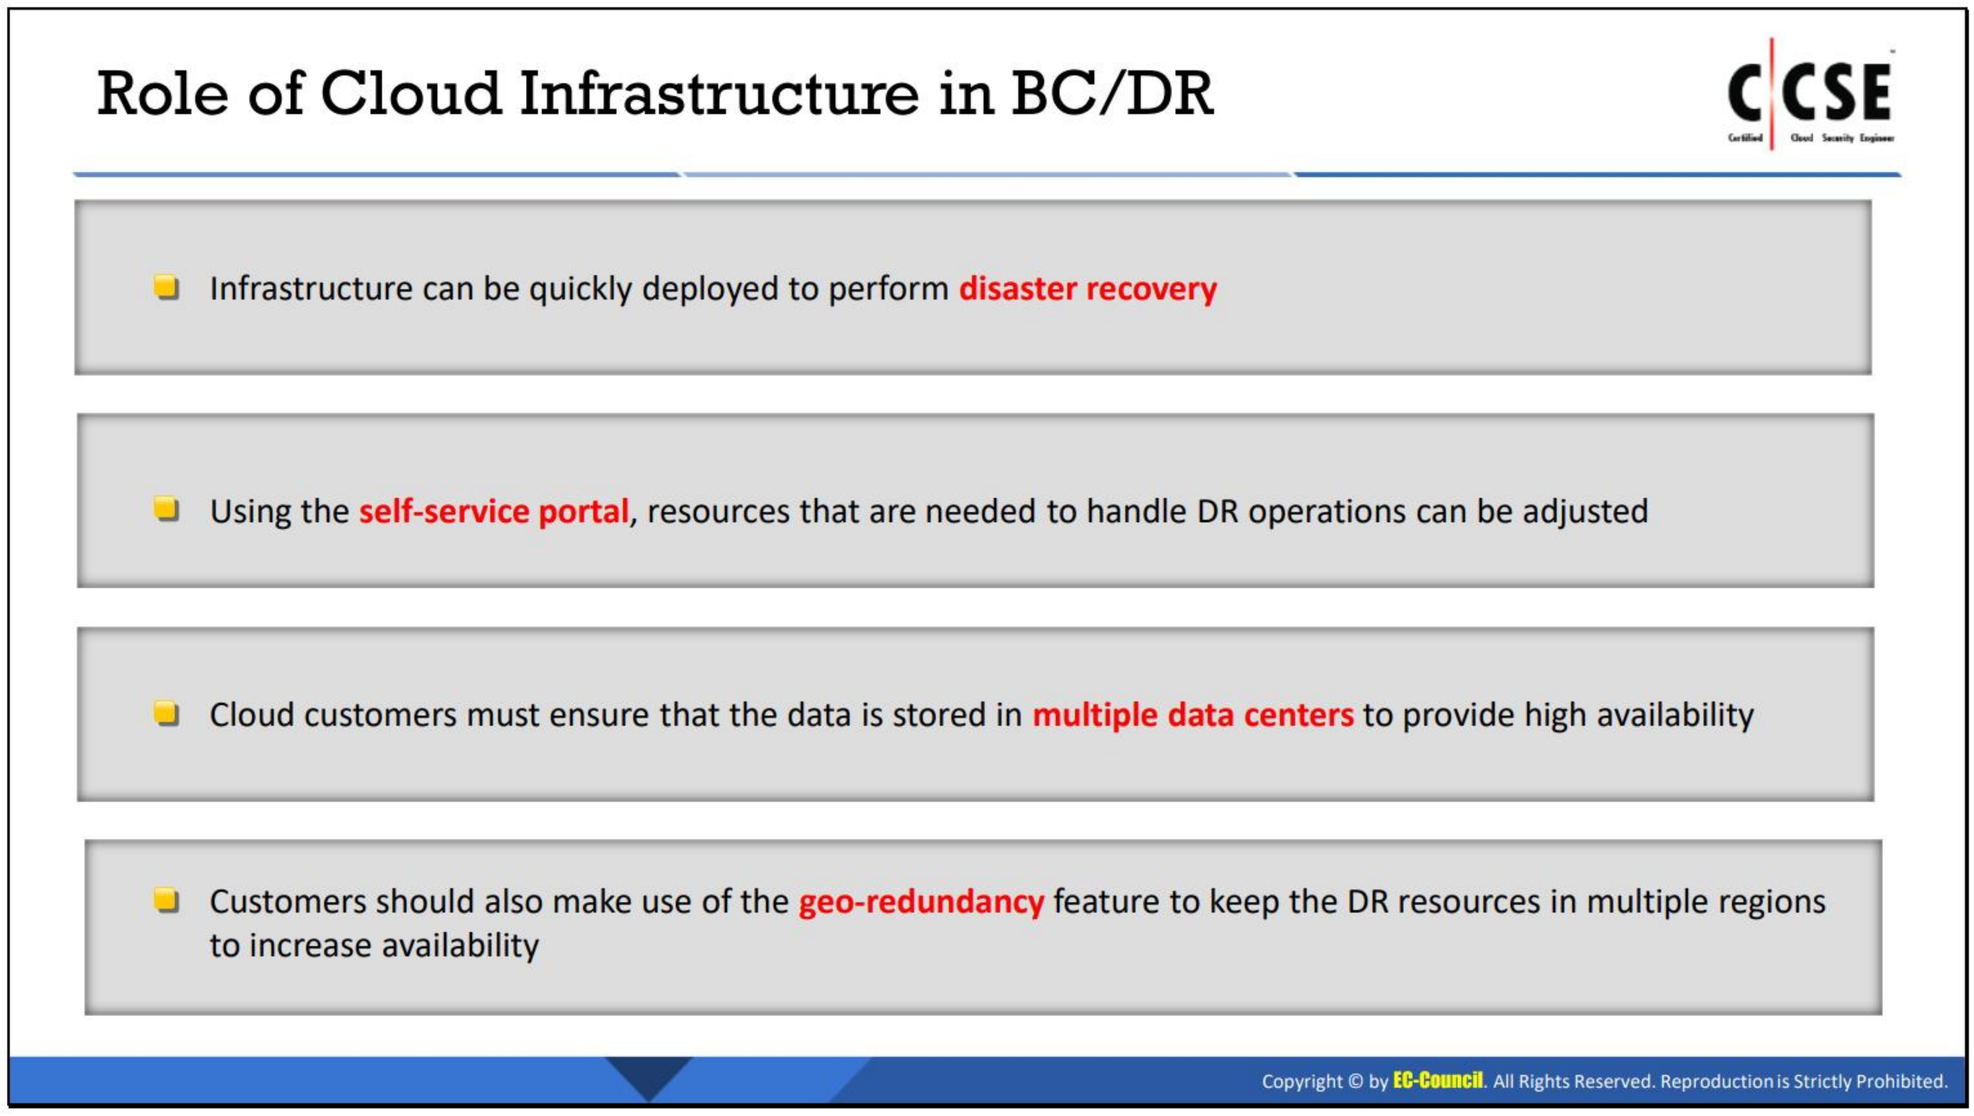

Role of Cloud Infrastructure in BC/DR

Any CSP provides features to support high availability. All the features work only when the architecture is adjusted. When the applications are moved to the cloud without redesigning the architecture, it could lead to lower availability.

The DR plan should consider infrastructure, threats, vulnerabilities, and DR strategies that work. When the DR plan is integrated with cloud computing services, each step of the recovery process can be automated.

Cloud infrastructure offers rapid elasticity. The on-demand self-service of the cloud provides flexible infrastructure. The infrastructure can be quickly deployed to perform disaster recovery. Through the self-service portal, resources can be adjusted that are needed to handle DR operations.

Cloud customers must ensure that the data is stored in multiple data centers to provide high availability. This also helps improve service reliability and provides redundancy. The customers should also make use of the geo-redundancy feature to keep the DR resources in multiple regions to increase availability.

It is essential to understand the cost model before storing data in multiple provider's physical locations. Some of the data and assets must be converted to work in different cloud locations. All the required assets must be included in the plan.

Certified Cloud Security Engineer Copyright © by **EC-Council**
All Rights Reserved. Reproduction is Strictly Prohibited.

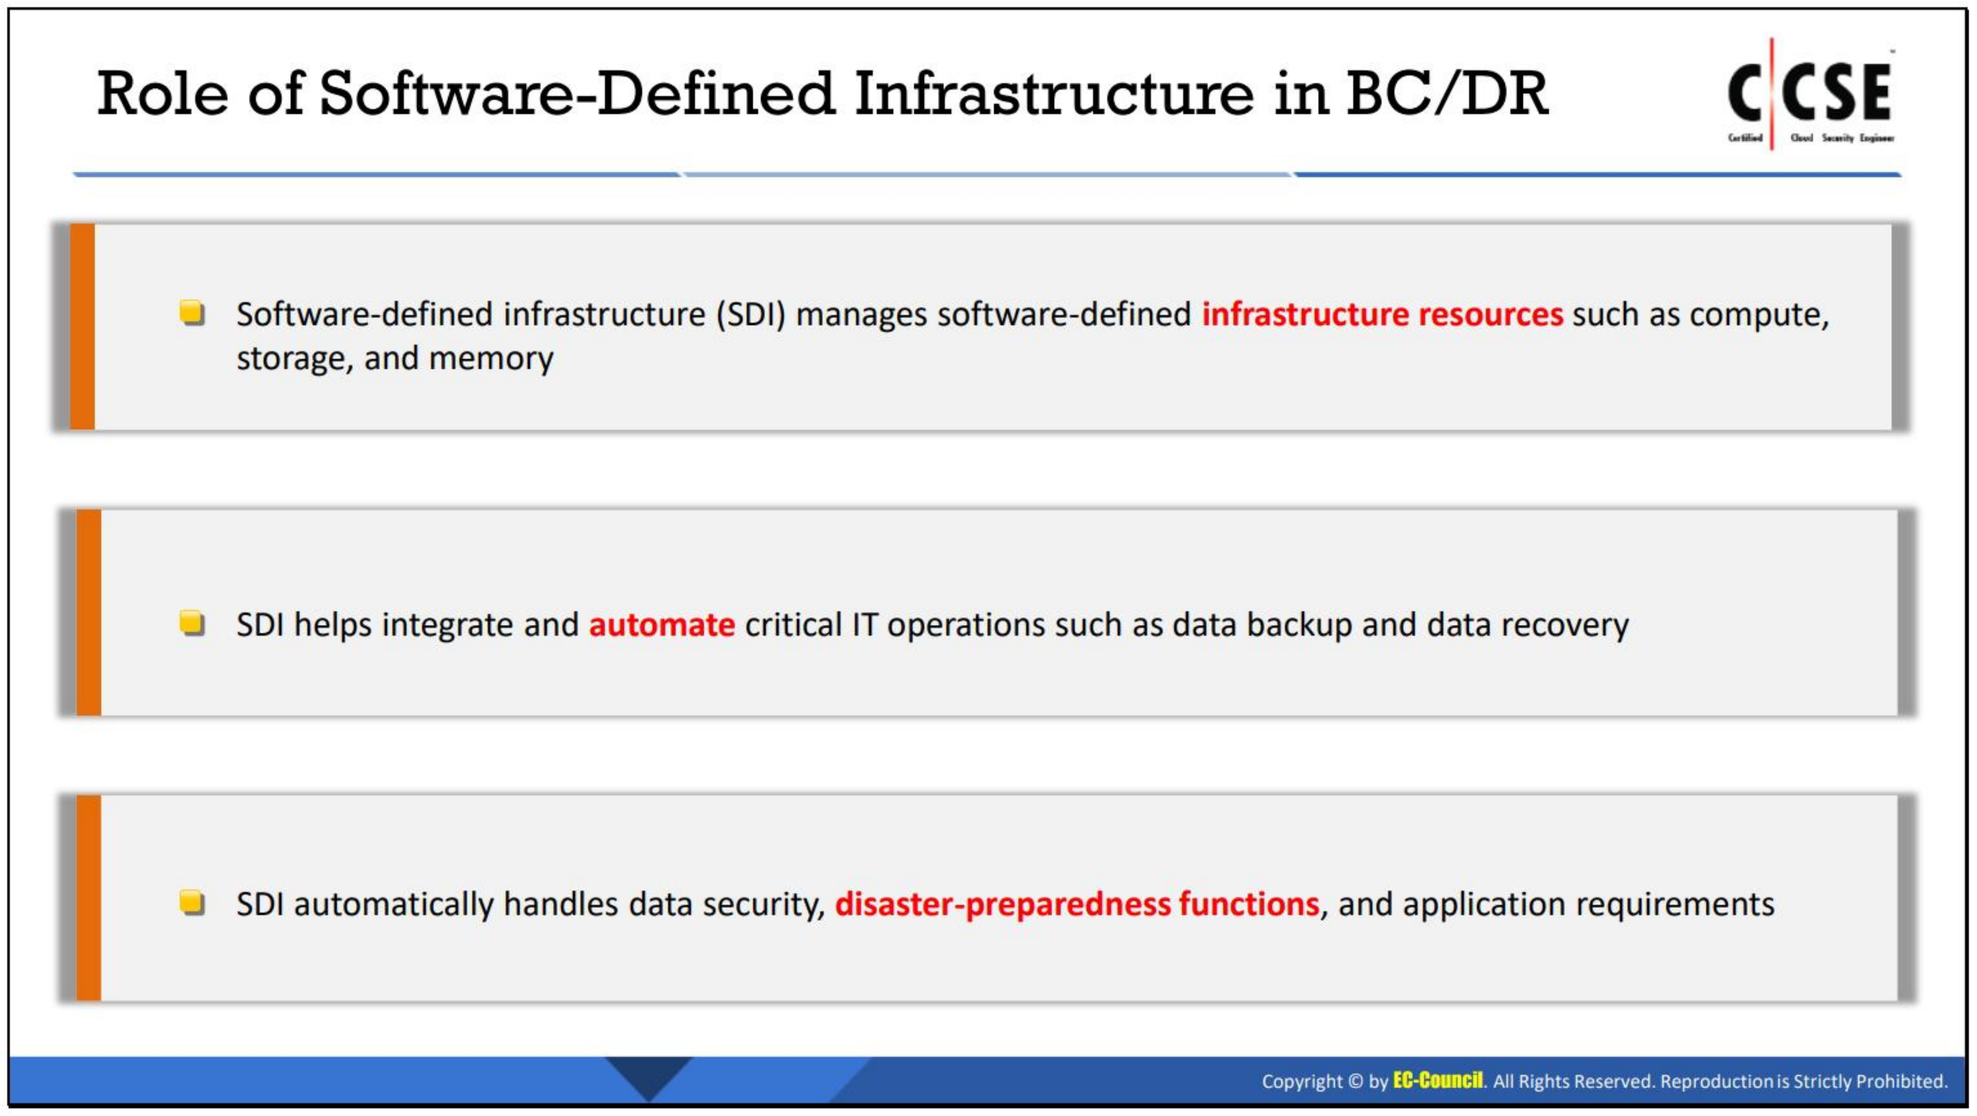

Role of Software-Defined Infrastructure in BCDR

Software-defined infrastructure (SDI) has become a popular approach to meet the demands on maximizing the value of the infrastructure deployment. SDI manages infrastructure resources such as compute, storage, and memory using software technologies. It does not require any human intervention. SDI helps build a robust and fault-tolerant IT environment.

SDI includes software-defined network (SDN), software-defined storage (SDS), and software-defined compute (SDC). SDI helps integrate and automate critical IT operations such as data backup and data recovery. It automatically handles data security, disaster-preparedness functions, and application requirements.

The storage of SDS is separated from hardware to make it more flexible to use. The resources of SDS can be used with any software, which is vital for a company that operates remotely. SDS is helpful while planning for an efficient BC plan. SDS plays an important role in preventing damage caused by cyberattacks. Any abnormal data access attempt can be detected and prevented with encryption, auditing, and virus detection. It also helps with data recovery by protecting the copies of data through air-gap measures. Air-gap measures physically separate the data and export files to off-site locations that cannot be accessed online.

The SDN can be used as a part of cloud platforms that support data backup and replication requirements. SDN can be optimized to disaster recovery systems where the requirements are defined by the customer and are provided by the SDN service provider.

SDN can be linked with disaster recovery-as-a-service (DRaaS). DraaS when linked with SDN can launch DR procedures by using the network as a transport medium. The cloud customers need to specify their requirements for the DR and should also consider RTO and RPO values.

Certified Cloud Security Engineer Copyright © by **EC-Council**
All Rights Reserved. Reproduction is Strictly Prohibited.

Then, the customers should determine how the SDN can be used and how it can be integrated with the existing network infrastructure.

hide01.ir

Certified Cloud Security Engineer Copyright © by **EC-Council**
All Rights Reserved. Reproduction is Strictly Prohibited.

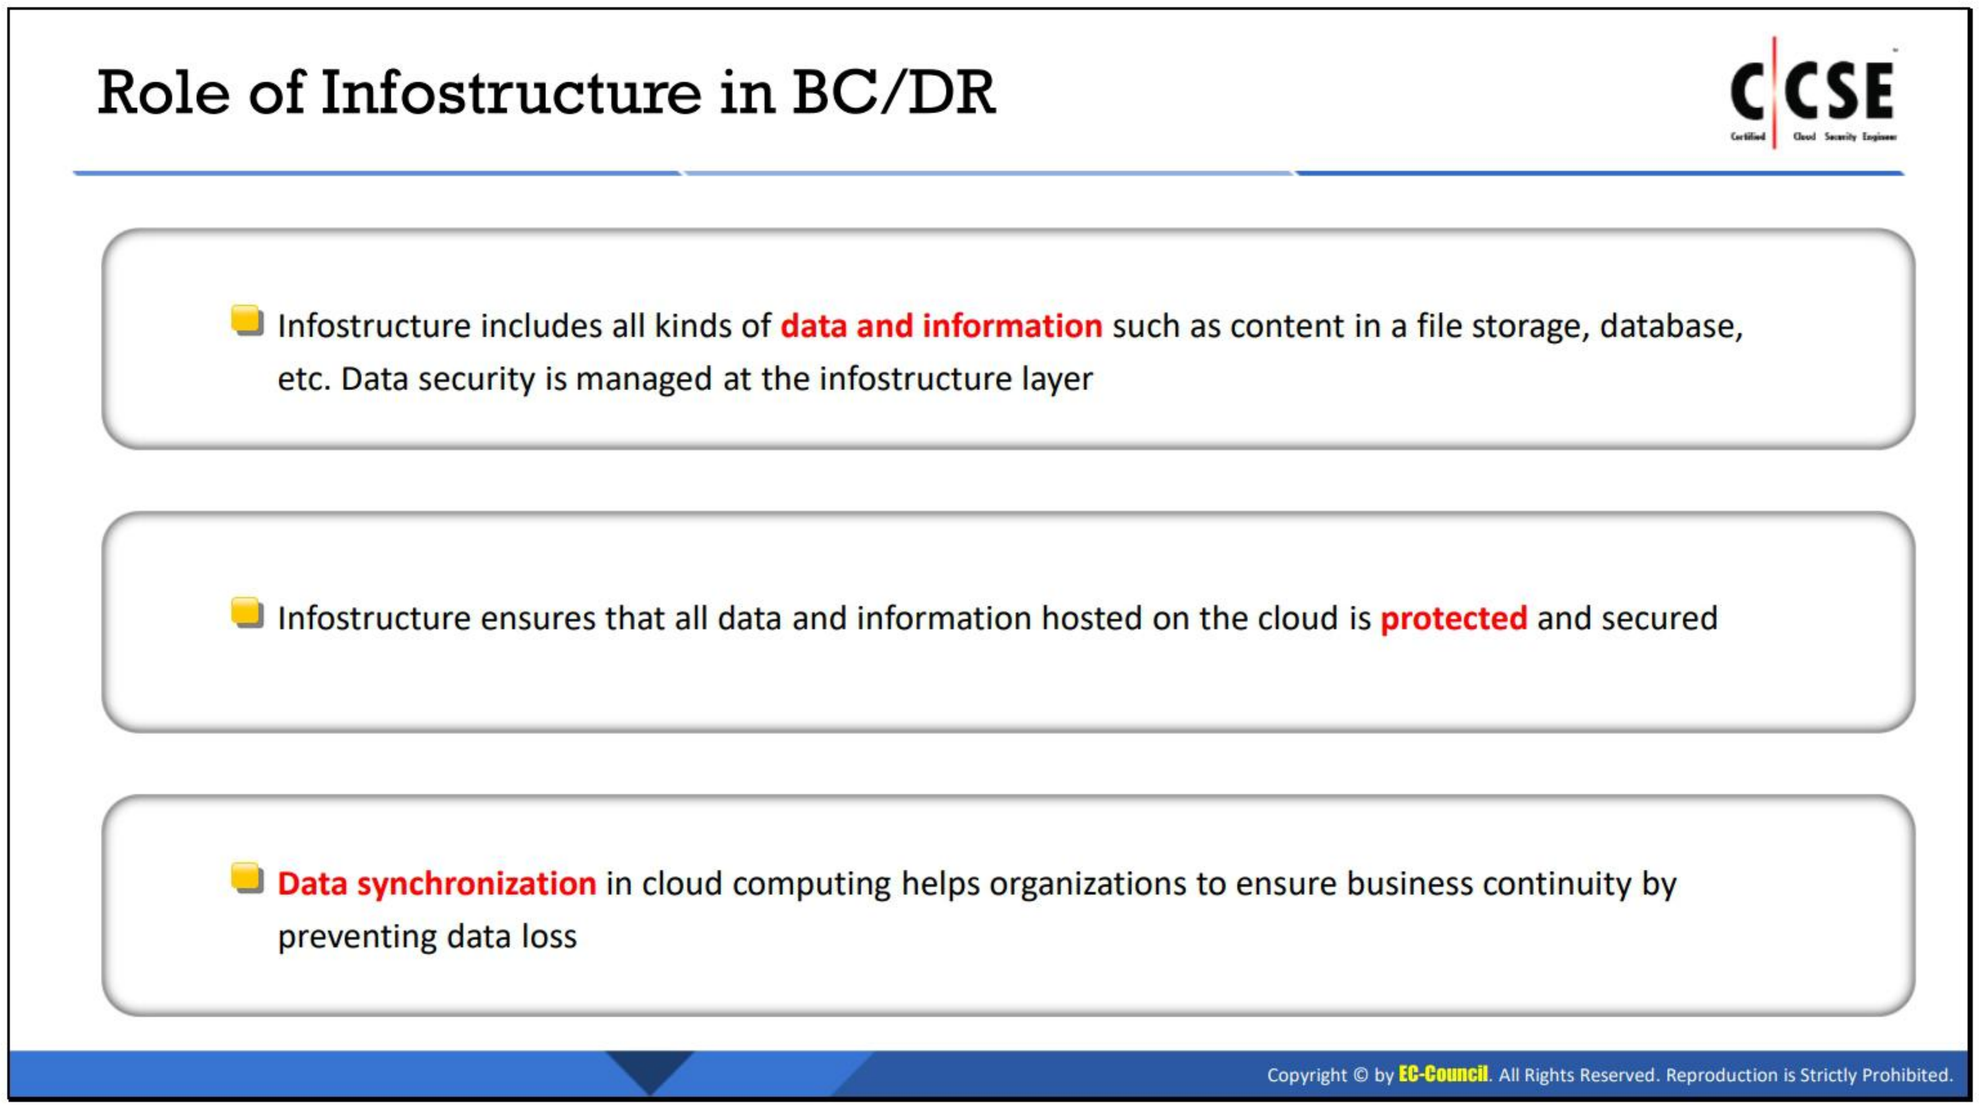

Role of Infostructure in BCDR

A layout of information that allows navigation of information is known as the infostructure. It is the top-most layer in the cloud logical model and includes data and information. It covers all kinds of data and information such as content in a file storage, database, etc. Data security is managed at the infostructure layer.

Infostructure ensures that all data and information hosted on the cloud is protected and secured. To achieve BC, all data must be synchronized to the backup site so that it can be made available when the primary site becomes unavailable.

Though storage costs are manageable, data synchronization is often a challenge that is difficult to manage across locations. In cloud environments, data synchronization is difficult because the size of data sets is large, and data must be in sync across all locations and services. Synchronizing data is difficult even in a single storage location.

Data Backups

Backing up data is the traditional and most common way to ensure BC. Most organizations use data backups to protect their digital assets when they are hit by disasters. Backing up of data involves transferring the changes that take place in local storage to either a cloud backup service or a local backup service. Typically, changes that are transferred to the backup site are incremental and are stored in backup archives. A good backup service stores ongoing backups and allows users to recover an old document.

Though this method is commonly used to secure data and information, it is quite cumbersome and involves considerable effort. Data recovery involves various steps such as launching a backup program, going through multiple backup instances, and initiating a restore. Sometimes, files and content in backups cannot be viewed until they are restored.

Certified Cloud Security Engineer Copyright © by **EC-Council**
All Rights Reserved. Reproduction is Strictly Prohibited.

Cloud Sync

Cloud synchronization is another method that can be used to protect data from unexpected events. Cloud sync involves moving of files that exist on your local computer to the cloud. It is often recommended to work on files that are stored in the cloud. This ensures that data is not lost when your primary storage is hit by disasters and the data stored on the cloud can be used from anywhere and at any time.

Data Synchronization

Data synchronization is the process of syncing data over two endpoints. Remote data synchronization involves real-time transfer of data from one endpoint to another over a network. Data synchronization in cloud computing helps organizations ensure BC by preventing data loss. With data sync, data remains always consistent on the primary site and backup site. Any data modification done at any site is synchronized on the other site to prevent mistakes, privacy breaches, and ensure that the information is up to date.

hide01.ir

Certified Cloud Security Engineer Copyright © by **EC-Council**
All Rights Reserved. Reproduction is Strictly Prohibited.

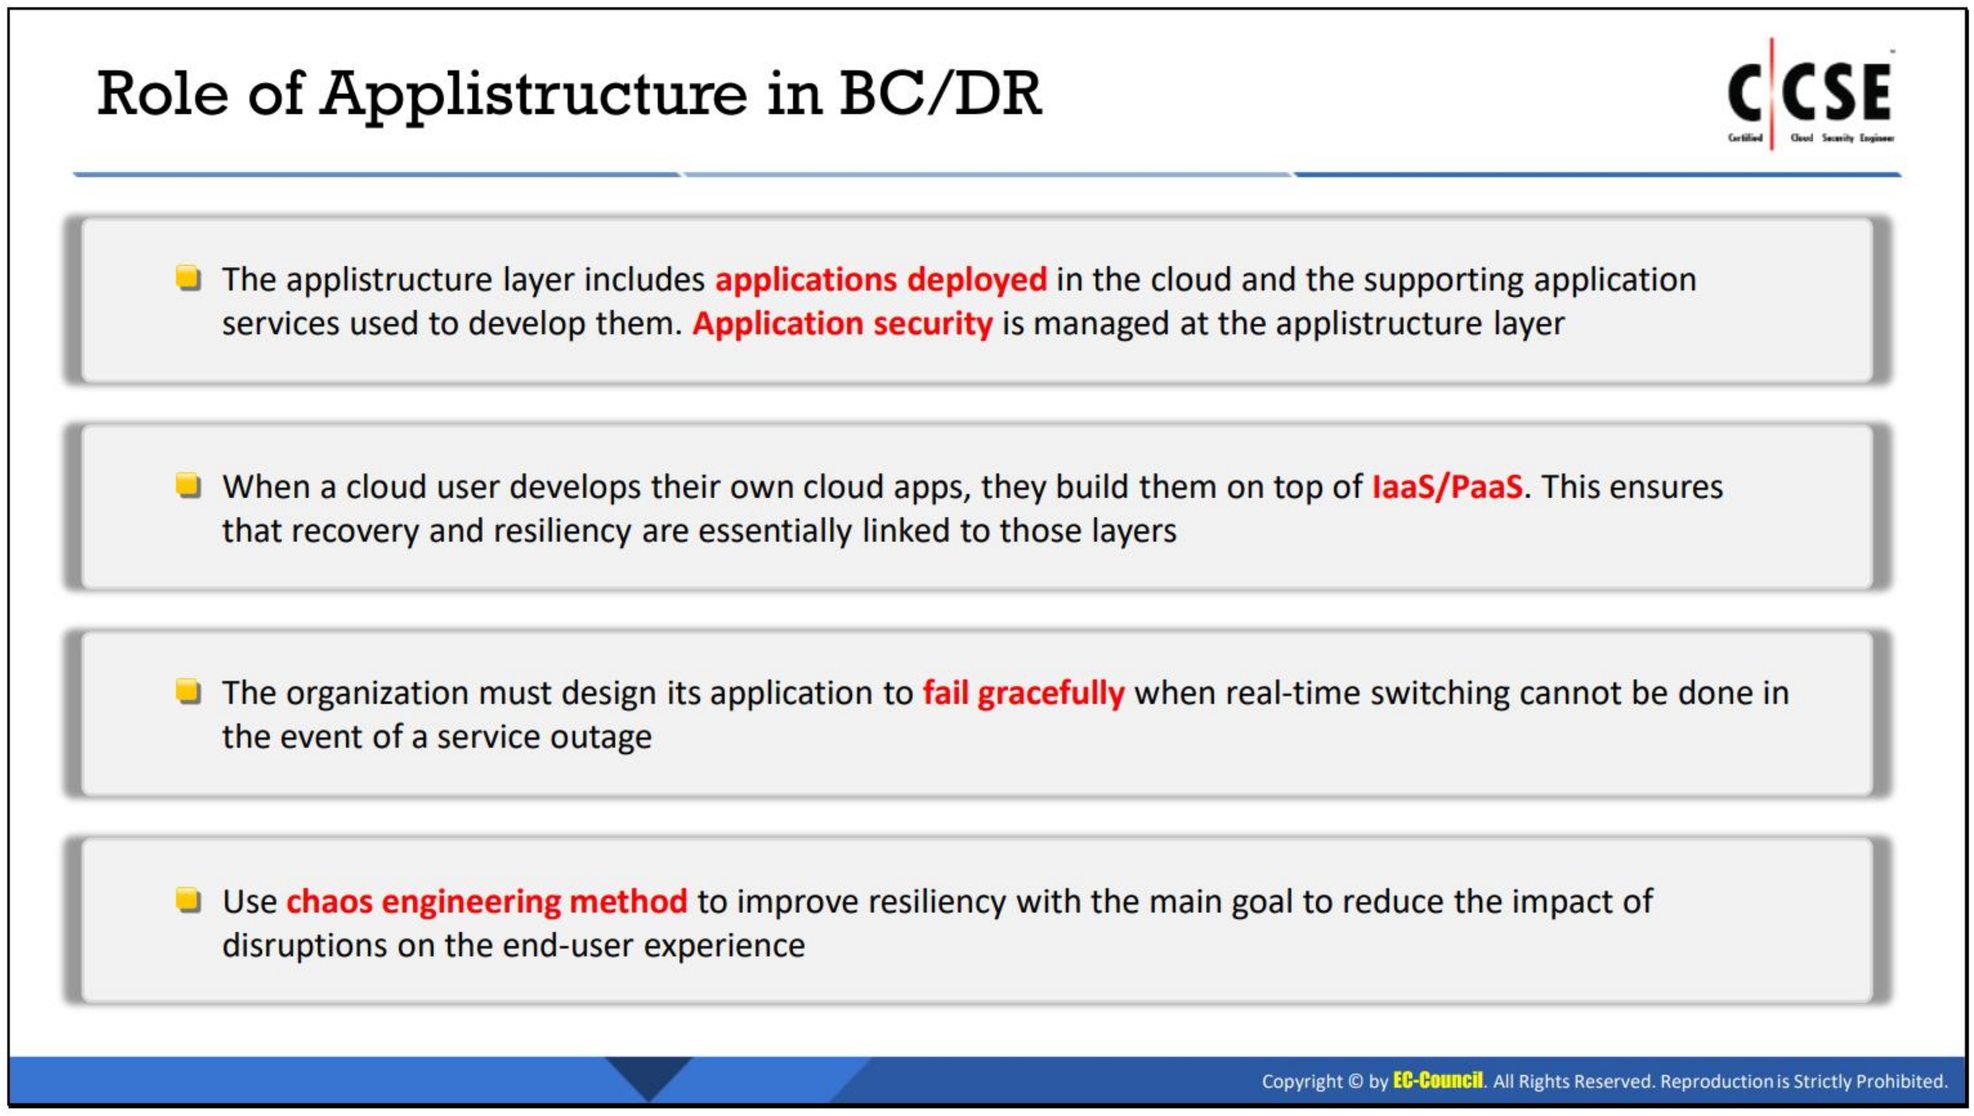

Role of Applistructure in BCDR

Applistructure refers to an amalgamation of technological infrastructure and applications. The applistructure layer is that part of the cloud computing logical model that includes applications deployed in the cloud and the supporting application services used to develop them; for example, PaaS features such as artificial intelligence analysis, message queues, or notification services. Application security is managed at the applistructure layer.

When a cloud user develops their own cloud apps, they build them on top of IaaS/PaaS. This ensures that recovery and resiliency are essentially linked to those layers. However, applistructure contains everything in an application. It is essential for an organization to be aware of PaaS limitations and lock-ins and prepare for PaaS component outages. A wide range of functions is part of platform services used in applications, ranging from authentication systems to notifications. It is normal for modern apps to integrate these types of services from various cloud providers.

Discussing with the Provider

The organization must have discussions with its provider on the availability of its service/component. It is necessary because the provider may not provide the same level of performance for two different services. For example, an infrastructure provider may not provide the same level of availability and performance for the database service as it provides for VM hosting service.

Design the Application to Fail Gracefully

The organization must design its application to fail gracefully when real-time switching cannot be performed in the event of a service outage. If something is designed to fail gracefully, it causes the least damage to its core functionality. If the organization designs the application to fail

Certified Cloud Security Engineer Copyright © by **EC-Council**
All Rights Reserved. Reproduction is Strictly Prohibited.

gracefully, it ensures that the application's core functionality remains usable even if the app's design fails. The organization can use various automation techniques to support the application for graceful failure. For example, if the organization's message queue service fails, it should stop the front end so that no message is lost.

Accept the Downtime

Downtime is another option that the organization can choose in case of a disaster or outage. If the organization does not need perfect availability, it can prepare to accept an outage. However, if it plans to accept an outage, it must ensure that it fails gracefully with emergency downtime notifications and responses. It can be achieved with the help of a static standby through DNS redirection.

Chaos Engineering

To help build resilient cloud deployments, chaos engineering is frequently used. Chaos engineering refers to the process of experimenting on a software system that is deployed in production to check the system's capability to withstand sudden and unexpected conditions. It has the potential to bring significant changes to BC. This method leverages tools to selectively degrade parts of the cloud to test BC continuously.

The main idea behind chaos engineering is to deliberately cause problems in an application/software system deployed in the production environment as part of the testing strategy and see how the application/software system deals with the disruption, detects vulnerabilities, and fixes them. This is the best method to improve resiliency with the main goal to reduce the impact of disruptions on the end-user experience. Additionally, this method helps engineers design systems that can withstand failures and disruptions.

hiteo1.ir

Certified Cloud Security Engineer Copyright © by **EC-Council**
All Rights Reserved. Reproduction is Strictly Prohibited.

LO#02: Learn Designing a Disaster Recovery and Business Continuity in Cloud

The objective of this section is to understand how to design a BCDR plan for a cloud environment to ensure that business operations and services remain unaffected or least affected during a disaster. This section will help cloud security engineers carefully design a BCDR plan by considering important factors such as BCDR architectures for PaaS, IaaS, and SaaS; BCDR implementation guidelines; DR checklists; DR tests, etc.

hid01.ir

Certified Cloud Security Engineer Copyright © by **EC-Council**
All Rights Reserved. Reproduction is Strictly Prohibited.

Designing a Disaster Recovery Plan in Cloud

- Designing a disaster recovery plan in the cloud involves the following:
  - **Audit the infrastructure and assess the risks**: All the assets should be **audited** and perform a risk assessment on the assets
  - **Perform a business impact analysis**: A business impact analysis (BIA) can be performed by considering the recovery **point objective (RPO)** and **recovery time objective (RTO)**

- Create a disaster recovery plan based on a replication approach and consider:
  - Backup and restore
  - Pilot light approach
  - Warm standby
  - Full replication in cloud
  - Multi-cloud option

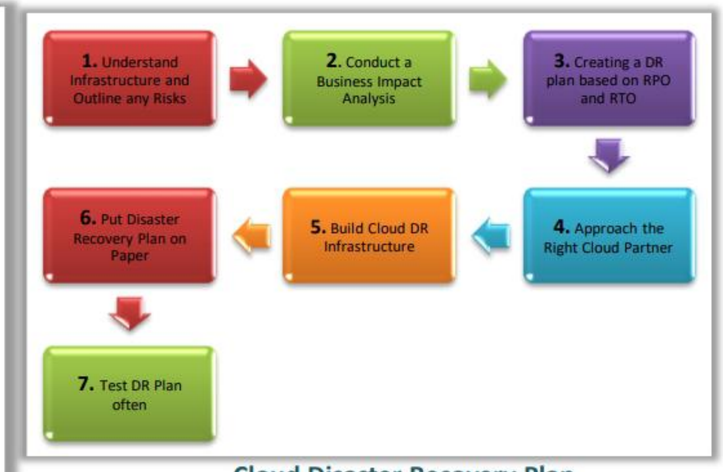

Cloud Disaster Recovery Plan

Copyright © by EC-Council. All Rights Reserved. Reproduction is Strictly Prohibited.

Designing a Disaster Recovery Plan in Cloud (Cont'd)

hide01.ir

- **Select the cloud provider:** The cloud provider should be selected after determining the **replication approach** to meet the DR goals
- **Set up the cloud disaster recovery infrastructure:** Cloud disaster recovery infrastructure can be setup based upon the **services provided** by the selected cloud provider
- **Documentation:** The disaster recovery plans should be documented which contains the instructions regarding the **roles and responsibilities** of every individual when a disaster strikes
- **Regular testing of the disaster recovery plan:** A full DR test is to be done and daily **back up snapshots** should be taken to ensure that everything is working fine

Copyright © by EC-Council. All Rights Reserved. Reproduction is Strictly Prohibited.

Designing a Disaster Recovery Plan in Cloud

Designing a DRP in the cloud involves the following:

- **Audit the infrastructure and assess the risks:** A DRP in cloud can be created by conducting an audit of the assets and how worthy are the assets. A risk assessment should be performed on all the assets. The disasters can be in different forms such as
  - **Site disaster:** Fire/power outages that make the datacenter unusable.

Certified Cloud Security Engineer Copyright © by **EC-Council**
All Rights Reserved. Reproduction is Strictly Prohibited.

- **Area disaster:** Floods, storms, and pandemics that make the datacenter unusable.
- **Regional disaster:** Terror attacks, financial failures, etc. that make the data center unusable.

Once the assets and the potential risks are assessed, proper disaster recovery plans can be developed to mitigate those risks.

- **Perform a business impact analysis** A BIA can be performed by considering the RPO and RTO.
  - **RTO:** Consider the case of critical business operations that runs 24×7 (e.g., a financial operations platform). In this case, an RTO can be set for 15 min or less and a DRP needs to be developed to perform full recovery in that time.
  - **RPO:** Consider the case of a co-working space company; here, the data may not be time sensitive. The RPO can be set for a couple of days here and the DRP can be set accordingly.

Figure 9.6: Cloud DR Plan

- **Creating a DRP based on RPO and RTO:** After determining the RPO and RTO, a DRP can be set up that meets the disaster recovery goals. The different disaster recovery approaches to implement the DR plan are
  - **Backup and restore:** The backup-and-restore plan can be easily created if the RPO and RTO is long. The on-premises data will get replicated to the cloud storage service such as Amazon S3 or Azure Storage. To avoid a larger downtime, faster restoration of the replicated data should be preferred.
  - **Pilot light approach:** This option helps in faster recovery of systems compared to the backup-and-restore approach. The cloud can be used to host critical elements

Certified Cloud Security Engineer Copyright © by **EC-Council**
All Rights Reserved. Reproduction is Strictly Prohibited.

of the application. If a disaster strikes, the application can be fully restored. Pilot light infrastructure depends upon the RTO.

The elements in this approach include replication of data in cloud, preconfigured servers ready to be launched, and installation packages or configuration information. The preconfigured servers can be fired up quickly for recovery.

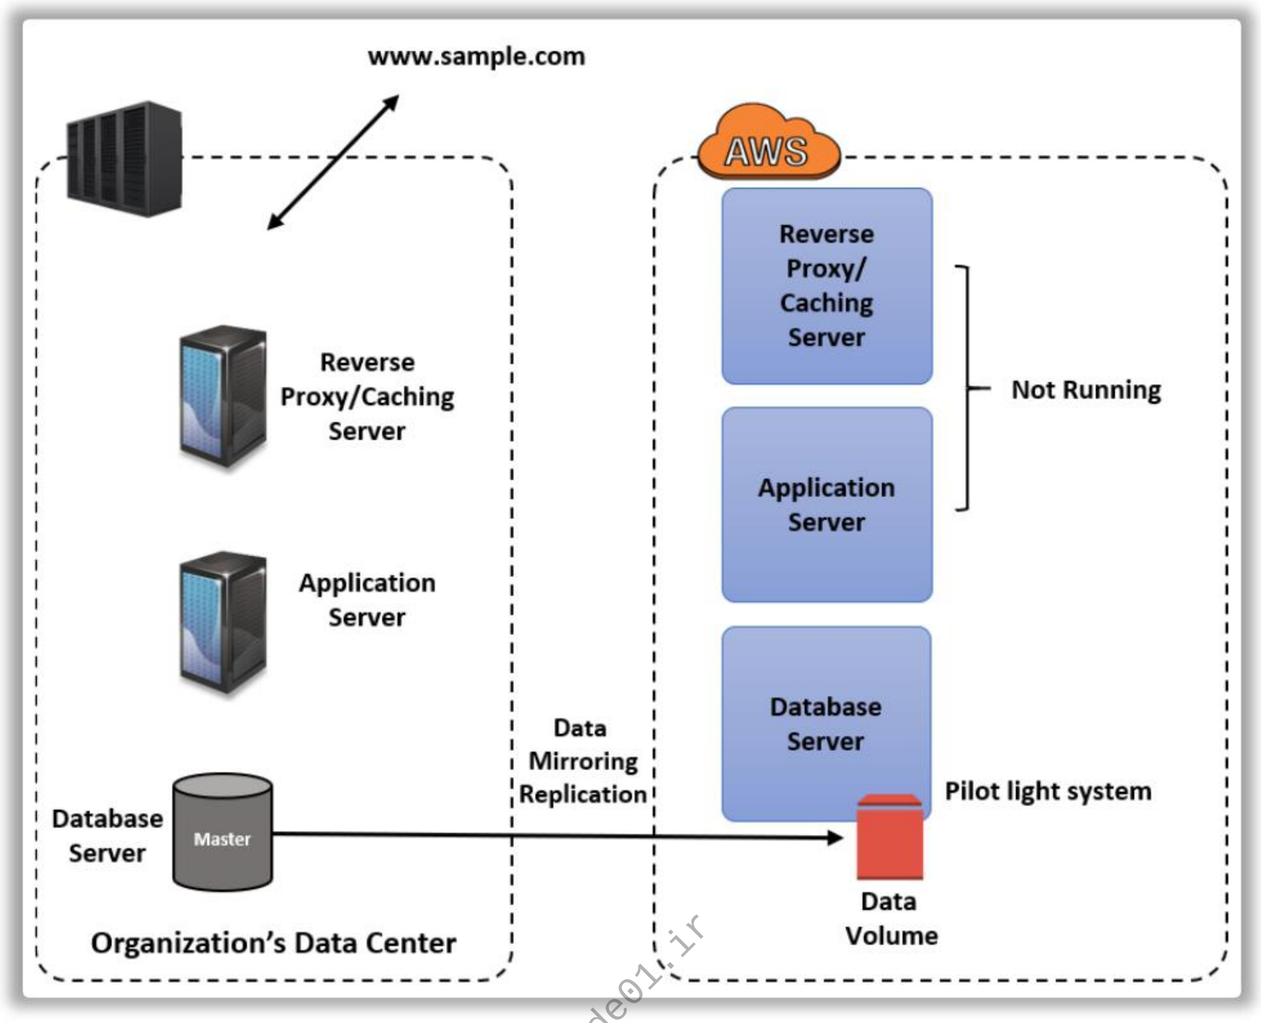

Figure 9.7: Pilot Light Approach

- **Warm standby:** Here, a scaled-down copy of the fully functioning environment in the cloud is created. In warm standby, the speed of recovery is faster as a full replica is always running in the cloud. When disaster strikes, the capacity of the VMs can be increased and the traffic can be rerouted across them.

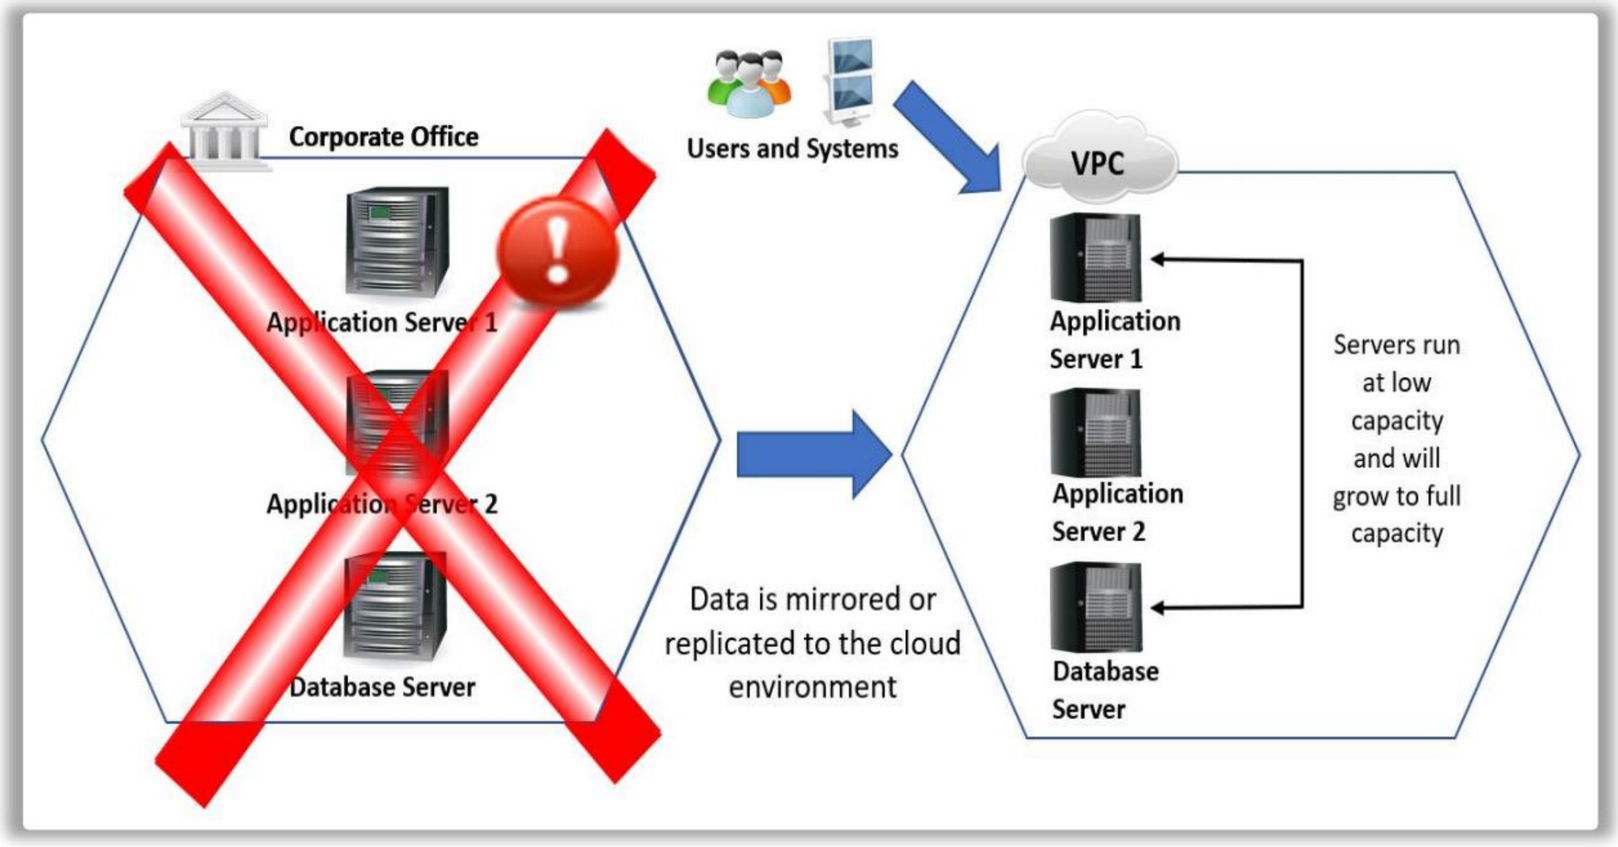

Figure 9.8: Warm Standby

Certified Cloud Security Engineer Copyright © by **EC-Council**
All Rights Reserved. Reproduction is Strictly Prohibited.

- **Full replication in cloud:** Through this approach, data can be quickly recovered as multiple replicas are running on-site and in the cloud. In this case, the on-premises environment is fully replicated in the cloud and the traffic get distributed between the environments. The disadvantage of full replication in cloud approach is that it is expensive compared to the other approaches because of the fully operational cloud infrastructure and the traffic routed to it.

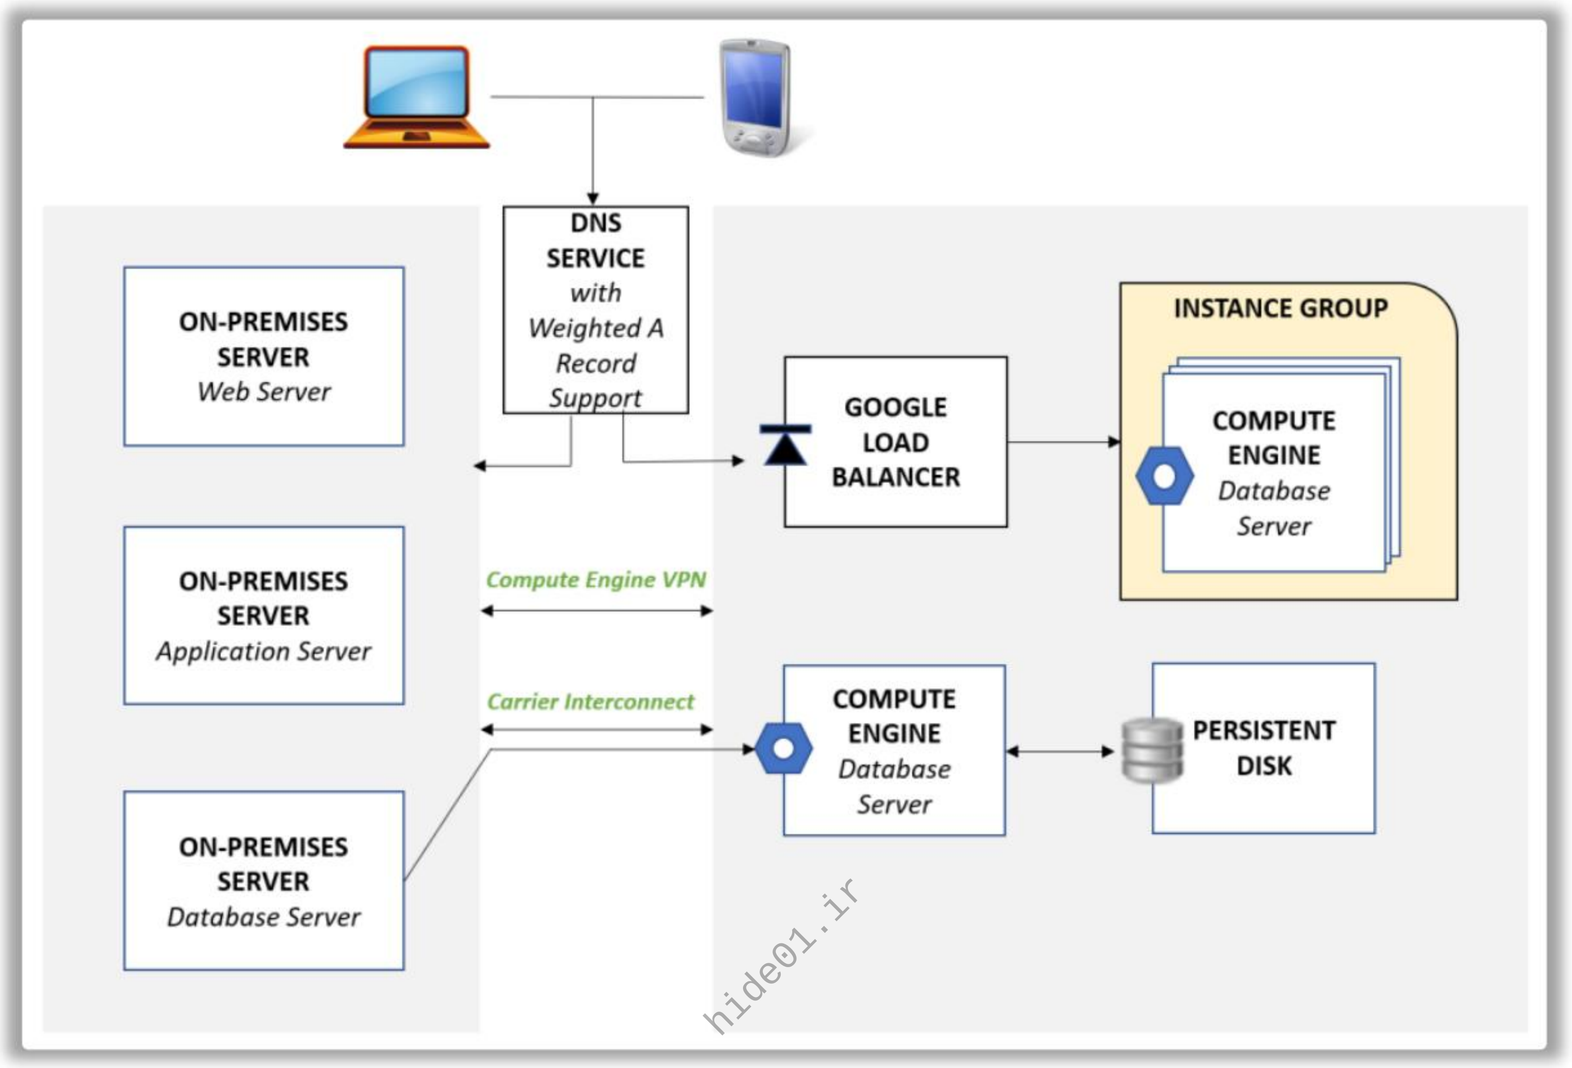

Figure 9.9: Full Replication in Cloud

- **Multi-cloud option**: If the complete infrastructure is deployed in the cloud, another cloud provider can be used as an option for disaster recovery. Within the cloud, backup can be performed by splitting the infrastructure into different availability zones. To ensure availability, it is always a better option to use two different cloud providers.

- **Select the cloud provider:** The cloud provider should be selected after determining the approach to meet the DR goals. If environment replication is also required, other factors such as reliability, speed of recovery, usability, simplicity in setup and recovery, scalability, security, and compliance also need to be considered. Some of the big cloud providers are AWS, Azure, Google Cloud, etc.
- **Set up the CDR infrastructure:** After the cloud provider is selected, the disaster recovery infrastructure can be set up. The following factors to be considered while designing the disaster recovery infrastructure.
  - Quantity of infrastructure components required.
  - How the data can be copied to the cloud.
  - The methods to be adopted for user authentication and access management.

Certified Cloud Security Engineer Copyright © by **EC-Council**
All Rights Reserved. Reproduction is Strictly Prohibited.

  - The security and compliance best practices.
  - The security measures to be implemented to face a disaster.

- **Documentation:** The DRPs should include clear instructions regarding the roles and responsibilities of every individual when a disaster strikes. A proper documentation ensures that the DRP is executed effectively when a disaster strikes.
- **Regular testing of the DRP:** After the disaster recovery plan is created, it is important to test and validate it regularly to ensure that it is effective. Testing of the plans will help in identifying the security loopholes when a disaster strikes. A full DR test is to be done and daily backup snapshots should be taken to ensure that everything is working fine. Testing the plan will help in ensuring that the organization is ready to face any emergencies.

hide01.ir

Certified Cloud Security Engineer Copyright © by **EC-Council**
All Rights Reserved. Reproduction is Strictly Prohibited.

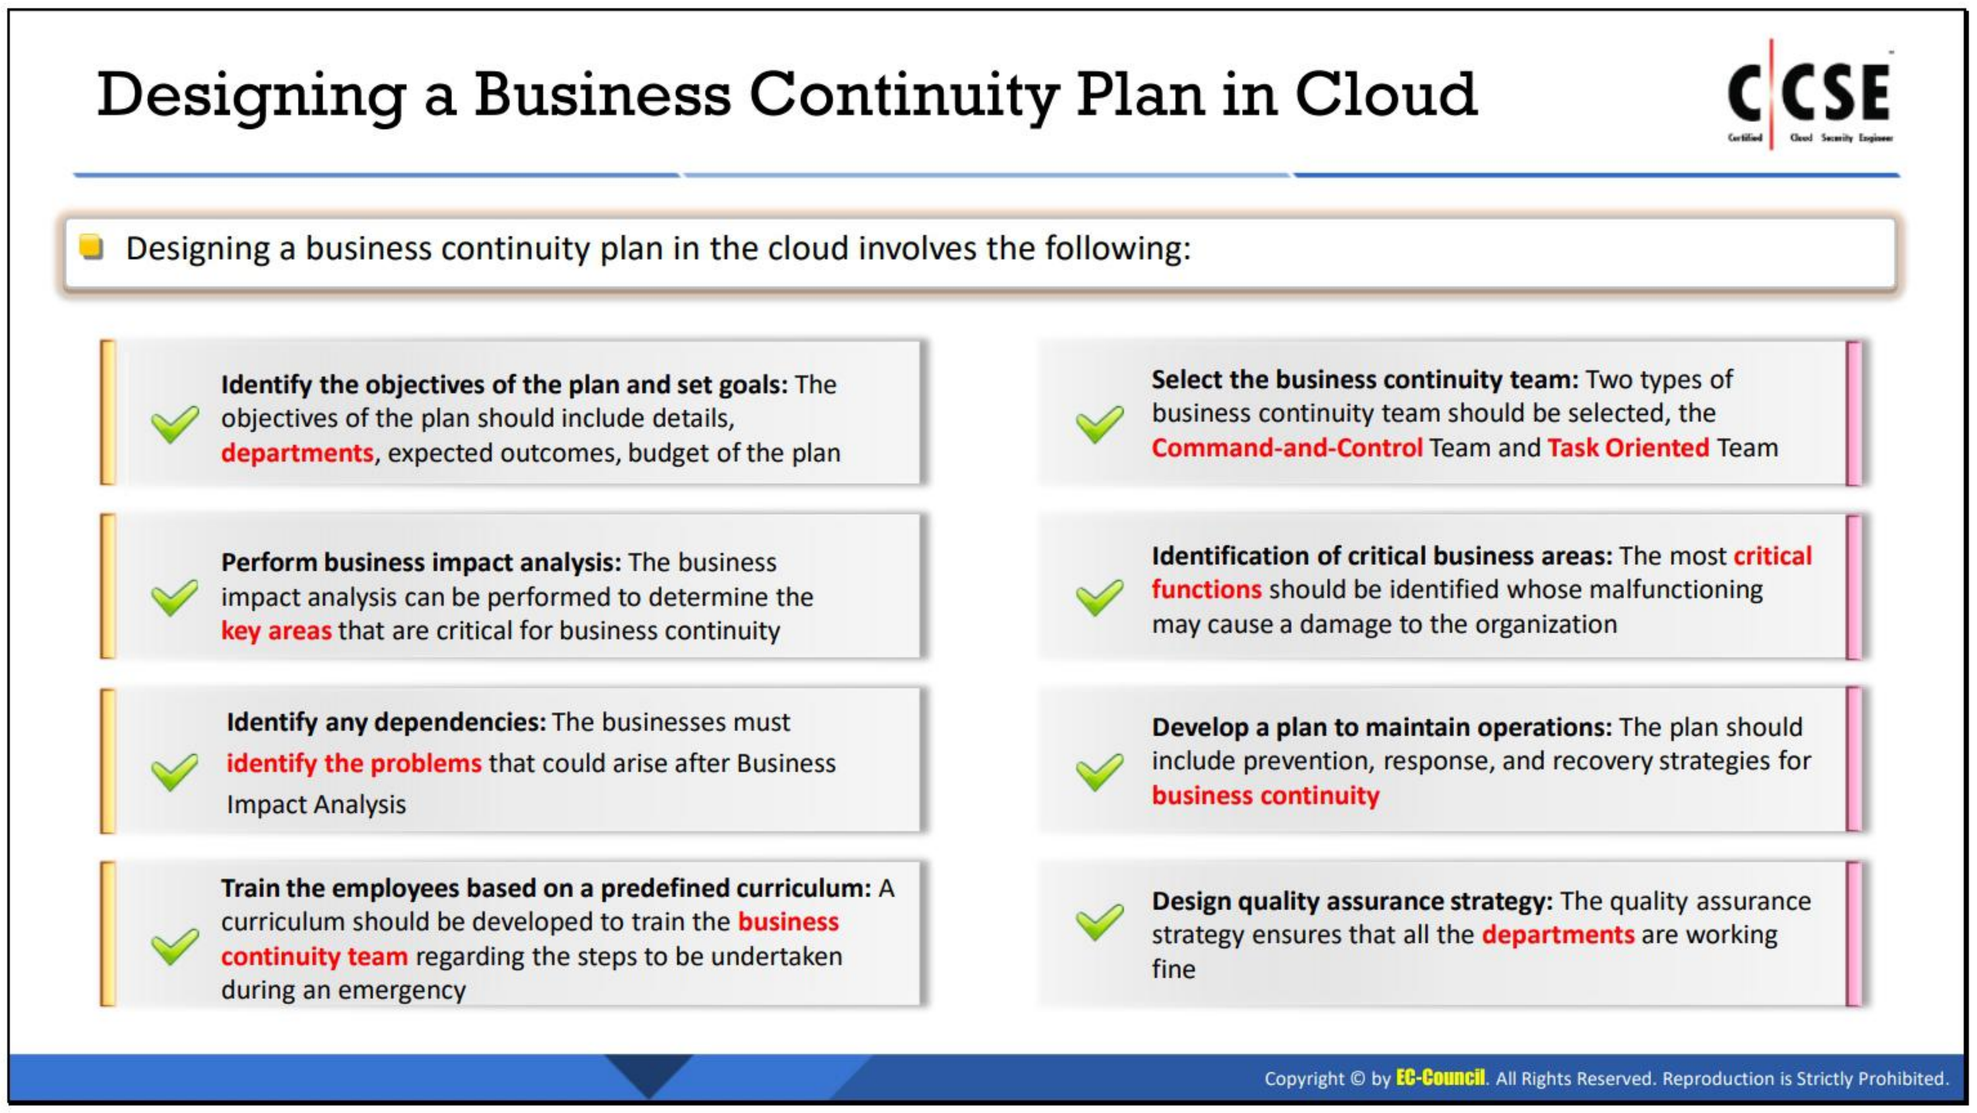

Designing a Business Continuity Plan in Cloud

Designing a BC plan in the cloud involves the following:

- **Identify the objectives of the plan and set goals:** While designing a BC plan in cloud, the objectives of the plan should be considered first. These include
  - The details of the plan
  - The departments covered by the plan
  - The expected outcomes of the plan
  - The milestones to be tracked
  - The budget of the plan
- **Select the BC team:** A BC team should be selected with the set of responsibilities assigned to them. It should include the contact information and the backup contact details of the team. The two types of sub-teams to consider are
- **Command-and-control teams:** The command-and-control team is the crisis management team that makes sure that everything is perfect, and the resources are ready to function.
- **Task-oriented teams:** The task-oriented team consists of personnel deployed to undertake specific tasks such as
  - Internal and external communication
  - Disaster recovery
  - Legal
  - IT team

Certified Cloud Security Engineer Copyright © by **EC-Council**
All Rights Reserved. Reproduction is Strictly Prohibited.

- **Perform business impact analysis:** BIA is the assessment of the threats that could affect a business. BIA can be performed to determine the key areas that are critical for BC.
- **Identification of the critical business areas:** In this phase, the most critical functions whose malfunctioning may cause a damage to the organization should be analyzed. While identifying the critical business functions, the factors to be considered are
  - The departments associated with this critical operation.
  - How often the critical function occurs.
  - The financial loss if this critical function fails.
  - The legal issues if this function is impacted.
- **Identify any dependencies:** Businesses must identify the problems that could arise after BIA. The maximum tolerable downtime of time sensitive applications could be analyzed. The resources can be allocated based on critical rating of the important business functions.
- **Develop a plan to maintain operations:** In this phase, an analysis of the current disaster recovery capabilities and improvements is to be performed based on those capabilities.
  - **Prevention strategies:** The BC plan should include the prevention strategies to be followed before a disaster strikes. The BIA also helps in the identification of areas that require mitigation measures to be adopted.
  - **Response strategies:** The organizations could prepare an emergency response plan and assign the staff based on their roles and responsibilities to respond to an incident.
  - **Recovery strategies:** This recovery strategies indicate who is responsible for implementation of the recovery measures. The RTO should be considered for implementing all the recovery plans.
- **Train employees based on a predefined curriculum:** A curriculum should be developed to train the BC team regarding the steps to be undertaken during an emergency. A mock emergency scenario can be created to test the overall efficiency of the BC team.

  The employees should be directed not to publish any kind of reports via social media to implement better business operations.
- **Quality assurance strategy:** The quality assurance strategy ensures that all the departments are working fine and the organization is prepared for facing a disaster.
  - **Internal reviews:** Internal reviews of the business should be conducted in situations such as
    - Threats
    - Changes to the company structure
    - Employee distribution geographically

Certified Cloud Security Engineer Copyright © by **EC-Council**
All Rights Reserved. Reproduction is Strictly Prohibited.

- **External reviews:** The external review should mention the external person who performs an audit of the BC plan and suggests improvements.

hide01.ir

Certified Cloud Security Engineer Copyright © by **EC-Council**
All Rights Reserved. Reproduction is Strictly Prohibited.

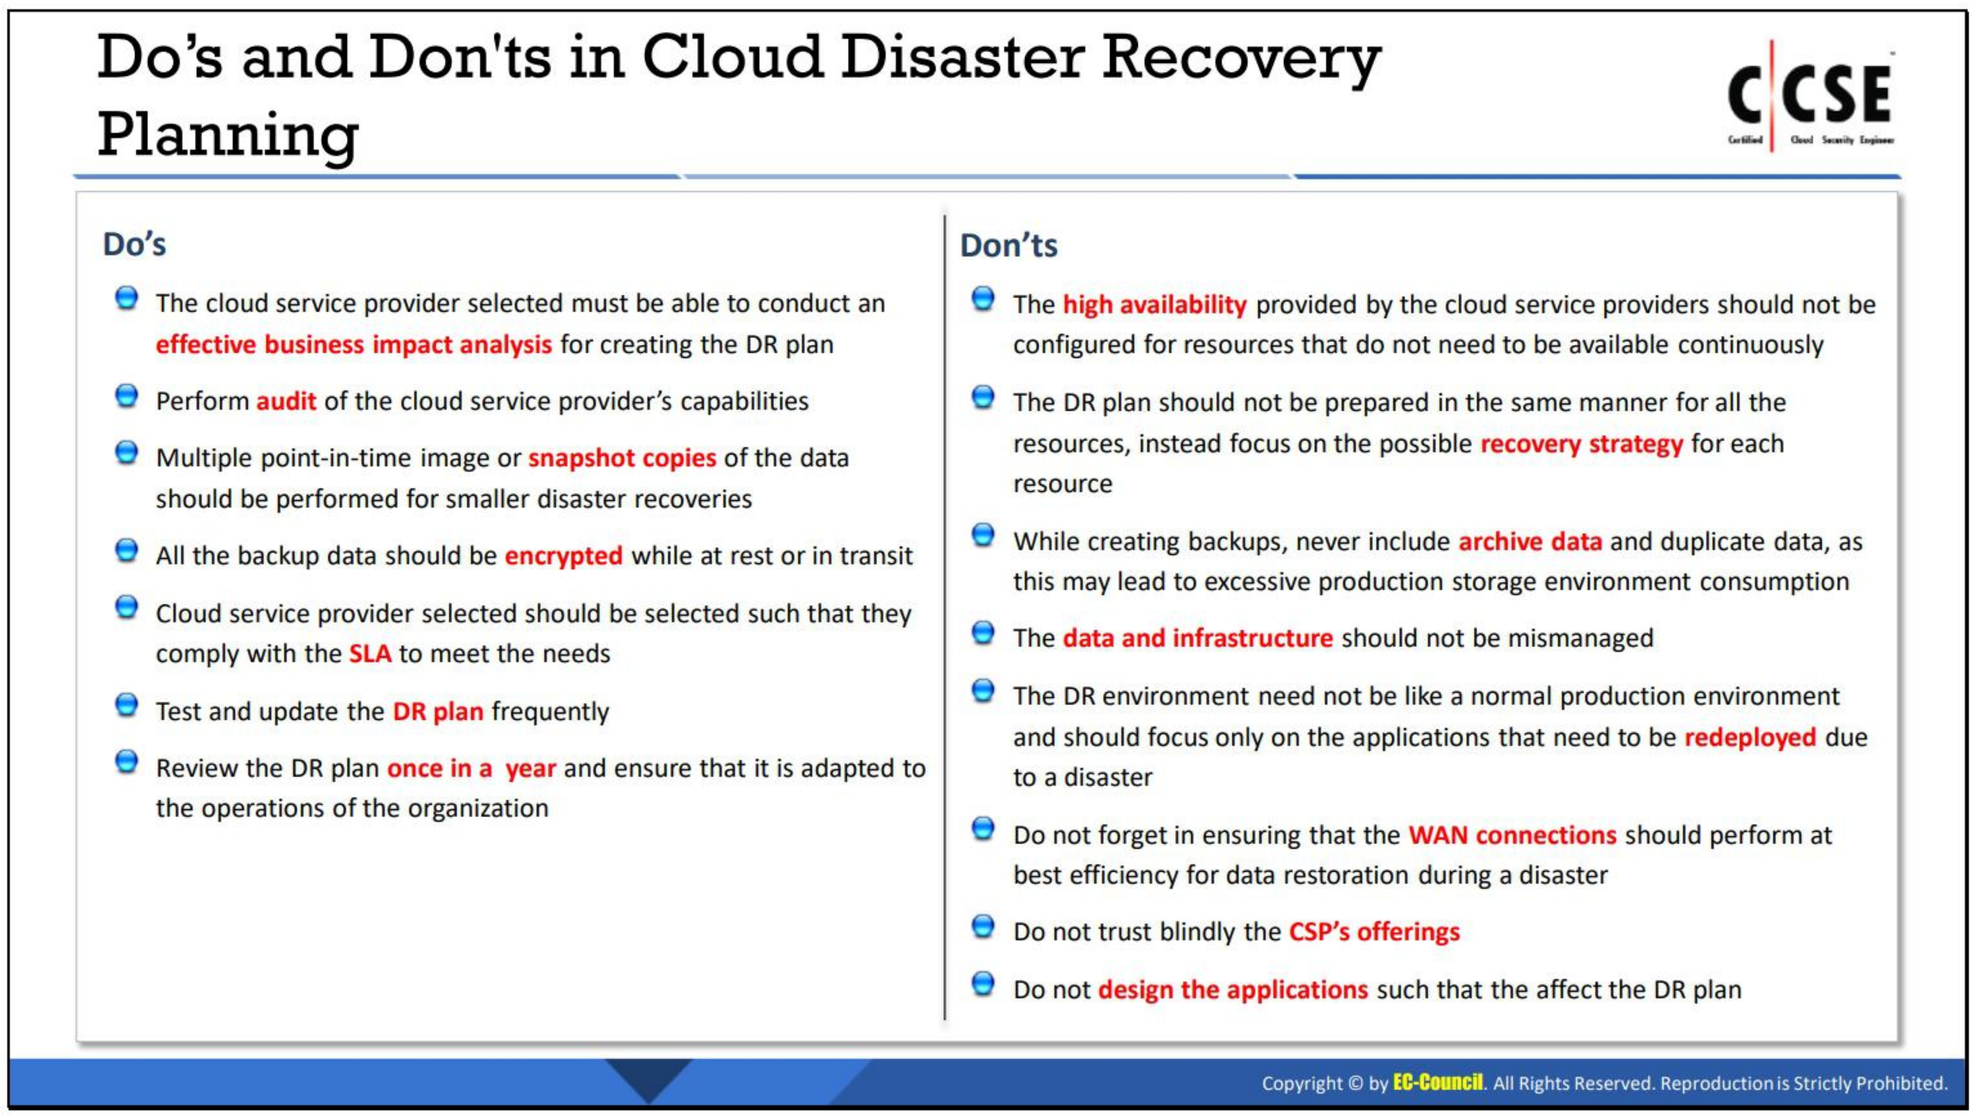

Do's and Don'ts in Cloud Disaster Recovery Planning

Do's

- **CSP must be able to conduct an effective BIA for creating the DR plan:** The CSP selected for implementing the services should have capabilities to conduct a BIA on the environment and develop the DR plan accordingly. The CSP should be able to provide answers such as whether the infrastructure be able to survive with a replication to failover systems or whether a full backup is essential.
- **Audit the CSP capabilities:** The ideal CSP will verify the risks associated with the organization and will help in figuring out the RTOs and RPOs of various services. The CSP should also comply with some regulatory compliance requirements such as HIPAA, GLBA, PCI-DSS etc. The user must also confirm the security measures adopted by the CSP in protecting the cloud infrastructure.
- **Multiple point-in-time image or snapshot copies of the data should be performed for smaller disaster recoveries:** Most often, disasters are having less impact and will be localized. To perform faster recovery in such cases, create multiple point-in-time images or snapshot copies of the data and also consider using replication services.
- **Backup data should be encrypted while at rest or in transit:** The backup solution should be such that the data should be encrypted while in transit or at rest.
- **Ensure that the CSP complies with the SLA:** The DR plan should be prepared after reviewing the SLAs with the CSP. For DR planning, the SLA must include
    - Data and applications included with the SLA.
    - The time period of performing testing and upgrades.

Certified Cloud Security Engineer Copyright © by **EC-Council**
All Rights Reserved. Reproduction is Strictly Prohibited.

- How CSP guarantees confidentiality, integrity, and availability of data.

- **Test and update the DR plan frequently:** The DR plan should be tested to ensure that the implemented practices are sufficient to protect the infrastructure from disasters. DR test should be performed in a simulation environment to test whether the plan works effectively. The plan should be updated after every deployment and other major changes implemented in the infrastructure. If the CSP intends to update and test the DR plan, ensure that it is included in the SLA.
- **Review the plan regularly:** The implemented DR plan should be reviewed periodically across the whole infrastructure and it should be ensured that major changes in the organization can adapt to the implemented DR plan.

Don'ts

- **High-availability configuration must not be configured for all resources:** High availability for performing failover is not necessary for resources that do not need to be available continuously.
- **Do not have the same DR approach for all workloads:** All the workloads do not require the same strategy for its operations. The best fit can be applied to the workloads based on its requirements. For data backup, different methods such as disk-to-disk replication, disk-to-disk mirroring, and replication by snapshots can be deployed.
- **Do not backup everything as a single backup:** All the data should not be backed-up as a single backup as the data may contain a combination of data, including archive data, duplicate data, etc. If everything is backed-up together, it may consume excess storage and the recovery process may take a considerably large amount of time. Perform backup or replication only for the day-to-day changes.
- **Do not mismanage the data and infrastructure:** Mismanagement of data and infrastructure is often the root cause of a disaster. Data should be prioritized based on its priority in the DR plan. Infrastructure monitoring should be performed to avoid failures in the different types of equipment. Data classification tools can be used to prioritize the data based on the restore priority and infrastructure can be managed by using resource management tools.
- **Do not design the DR environment as similar to production environment:** The DR environment should not be an exact copy of the production environment. When a disaster strikes, only certain components are affected and only these components need to be redeployed. Additionally, server virtualization can be used for applications to run on the virtual environment.
- **Do not forget to ensure the efficiency of WAN connection:** The WAN connection must perform at full efficiency to perform data recovery or access to applications. A secondary WAN connection should be provided in addition to the primary. This is to ensure connectivity that, in case the cause of the disaster in the production environment was the issue with primary WAN, disaster recovery can be performed only with a secondary WAN

Certified Cloud Security Engineer Copyright © by **EC-Council**
All Rights Reserved. Reproduction is Strictly Prohibited.

connectivity. The remote disaster recovery site should be at least 80 km away from the production environment to recover in the case of regional disasters.

- **Do not trust the CSP blindly:** While using cloud services for DR, do not trust the CSP blindly with the offerings they provide. Confirm the offerings regarding the security infrastructure and ensure them by including them in the SLAs.
- **Do not design the applications such that it does affect DR plan:** Insecure design of applications, such as deploying applications in proprietary server hypervisors, insecure functions for building applications, or high caching that leads to data loss, will cause the disaster recovery to be a failure. Therefore, for better disaster recovery, the DR plan should be included in the design of the application itself.

hide01.ir

Certified Cloud Security Engineer Copyright © by **EC-Council**
All Rights Reserved. Reproduction is Strictly Prohibited.

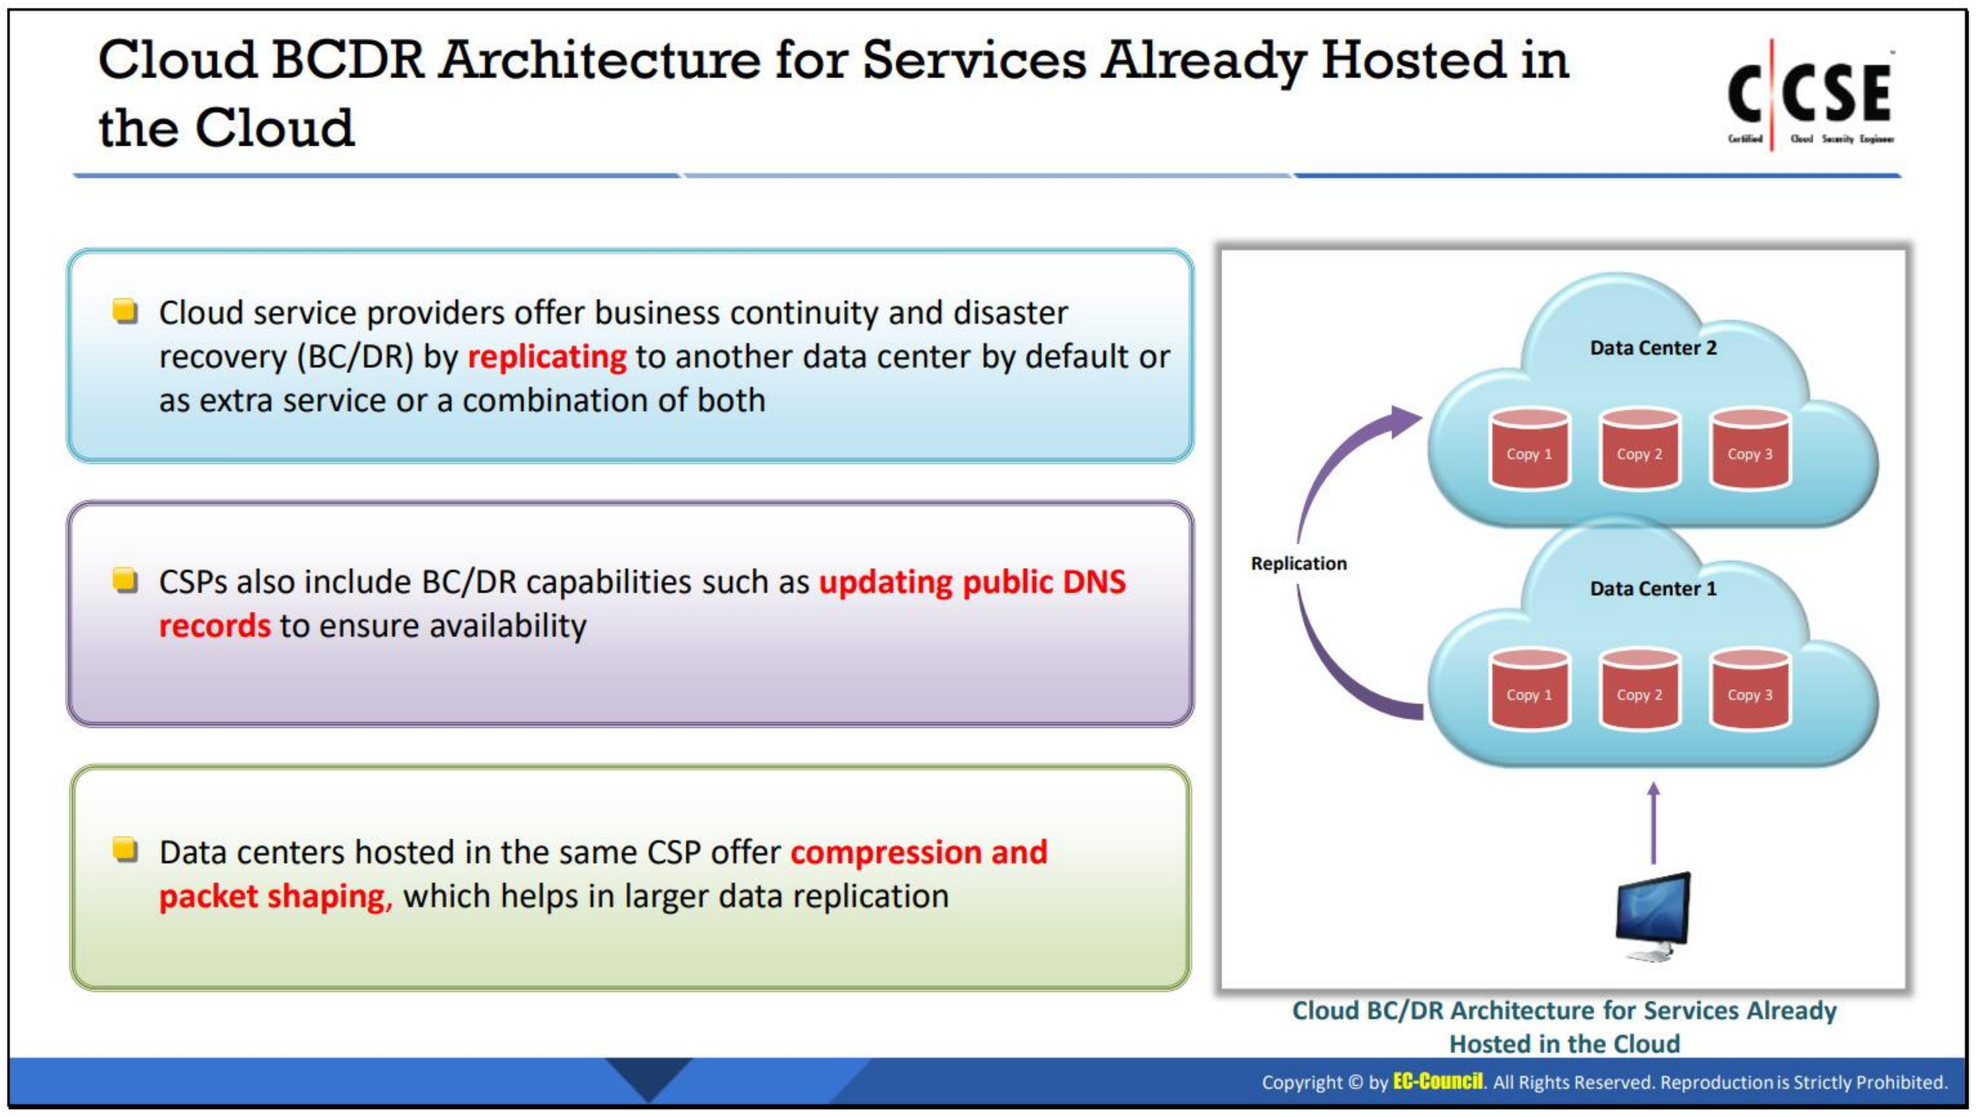

Cloud BCDR Architecture for Services Already Hosted in the Cloud

Cloud-to-Cloud Disaster Recovery

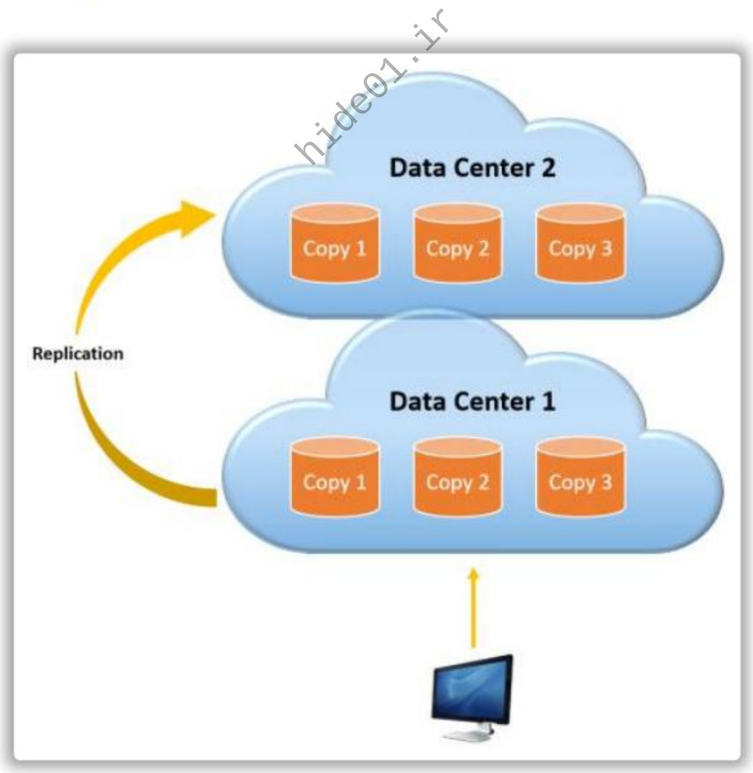

Figure 9.10: Cloud BCDR Architecture for Services Already Hosted in the Cloud

If disaster recovery is to be used for production environments, it should be ensured that the services are replicated to another site. CSPs offer BCDR by replicating to another data center by default or as extra service or a combination of both. Disaster recovery when deployed in the cloud is simpler than on-premises deployment as the CSP manages replication, bandwidth, networking, etc. CSPs also includes BCDR capabilities such as updating public DNS records to ensure availability.

Certified Cloud Security Engineer Copyright © by **EC-Council**
All Rights Reserved. Reproduction is Strictly Prohibited.

In the cloud-to-cloud disaster recovery shown in the figure, the data is replicated across three devices at Data Center 1 and also to the disaster recovery site. If the disaster recovery gets activated, all external DNS records would be pointed to Data Center 2. The data centers hosted in the same CSP have compression and packet shaping capabilities, which helps in larger data replication.

hide01.ir

Certified Cloud Security Engineer Copyright © by **EC-Council**
All Rights Reserved. Reproduction is Strictly Prohibited.

Cloud BCDR Architecture for In-House/Traditional Non-Cloud Services

- In this model, the system will be hosted **on-premises** and, depending upon the service, the data or systems get replicated to the cloud service provider
- CSP provides disaster recovery capabilities for traditional on-premises systems, which depends upon the **RPO and RTO** of the cloud customer

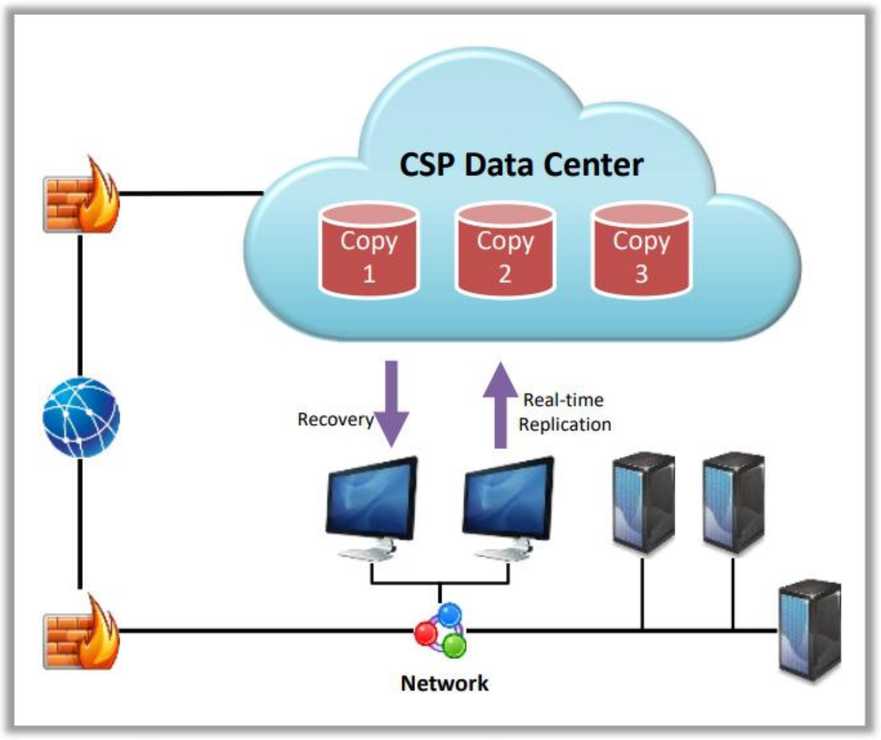

Cloud BC/DR Architecture for In-House/Traditional Non-Cloud Services

Copyright © by EC-Council. All Rights Reserved. Reproduction is Strictly Prohibited.

Cloud BCDR Architecture for In-House/Traditional Non-Cloud Services

Traditional Data Center to Cloud Service Provider Replication

hide01.ir

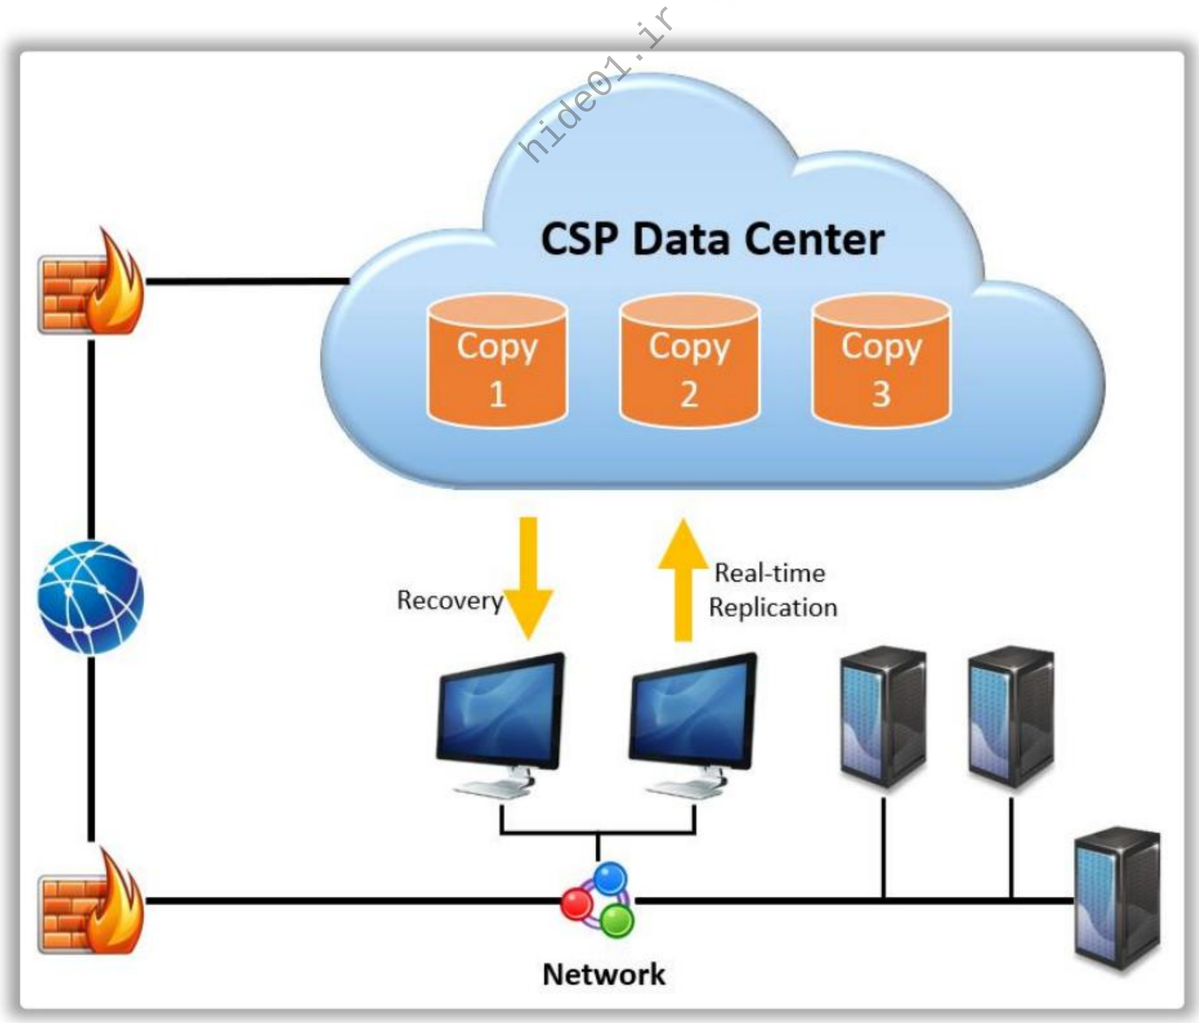

Figure 9.11: Cloud BCDR Architecture for In-House/Traditional Noncloud Services

In this model, the system is hosted on-premises and, depending upon the service, the data or systems get replicated to the CSP. In hybrid-cloud models, the CSP hosts the infrastructure on-site to manage monitoring, replication, and failover.

Certified Cloud Security Engineer Copyright © by **EC-Council**
All Rights Reserved. Reproduction is Strictly Prohibited.

The CSP provides disaster recovery capabilities for traditional on-premises systems depending on the RPO and RTO of the cloud customer. The configuration principles should consist of the following:

- Configuration at the local systems to replicate data.
- Configuration of the storage at CSP to get replicated.
- The required specifications at CSP to be defined to invoke disaster recovery.
- Configurations to be performed for third parties to provide services for the cloud hosted systems as on-premises systems when DR is invoked.
- Access to the CSP should be tested from relevant sites.
- Performance tests should be performed to ensure that the cloud service behaves efficiently.
- All the applications and systems should be supported in the cloud environment.
- Alternative solutions should be developed for systems that will not function in cloud.

hide01.ir

Certified Cloud Security Engineer Copyright © by **EC-Council**
All Rights Reserved. Reproduction is Strictly Prohibited.

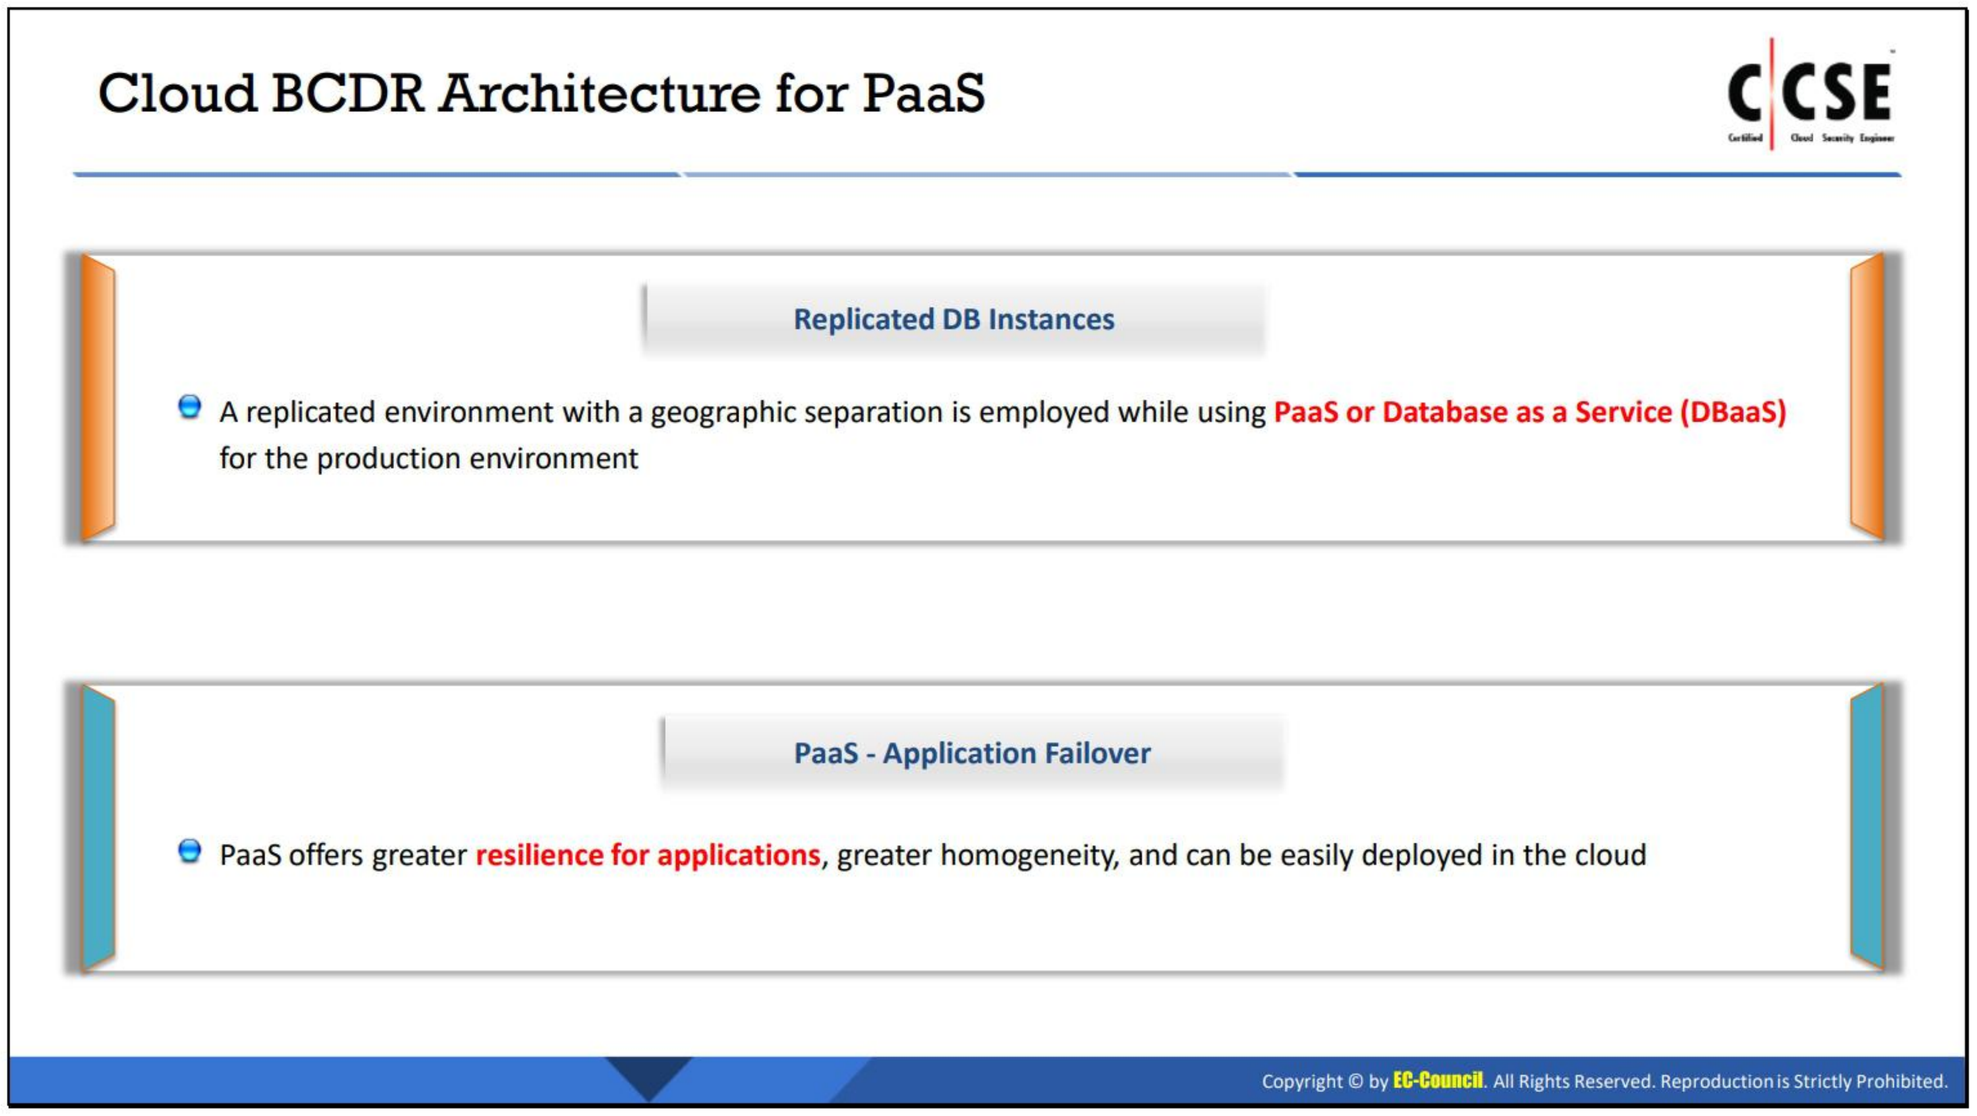

Cloud BCDR Architecture for PaaS

PaaS Replicated Database Instances

A replicated environment with a geographic separation is employed while using PaaS or Data base-as-a-Service (DBaaS) for production environment. A replicated system provides load balancing and is more cost-effective.

h1ee0.ir

PaaS – Application Failover

PaaS offer greater resilience for applications, greater homogeneity, and can be easily deployed in cloud. PaaS allows a single application image to operate on IaaS. For BCDR, it provides the following features:

- Cloud bursting
- On-demand scaling
- Service provider redundancy
- Geographic independence

Certified Cloud Security Engineer Copyright © by **EC-Council**
All Rights Reserved. Reproduction is Strictly Prohibited.

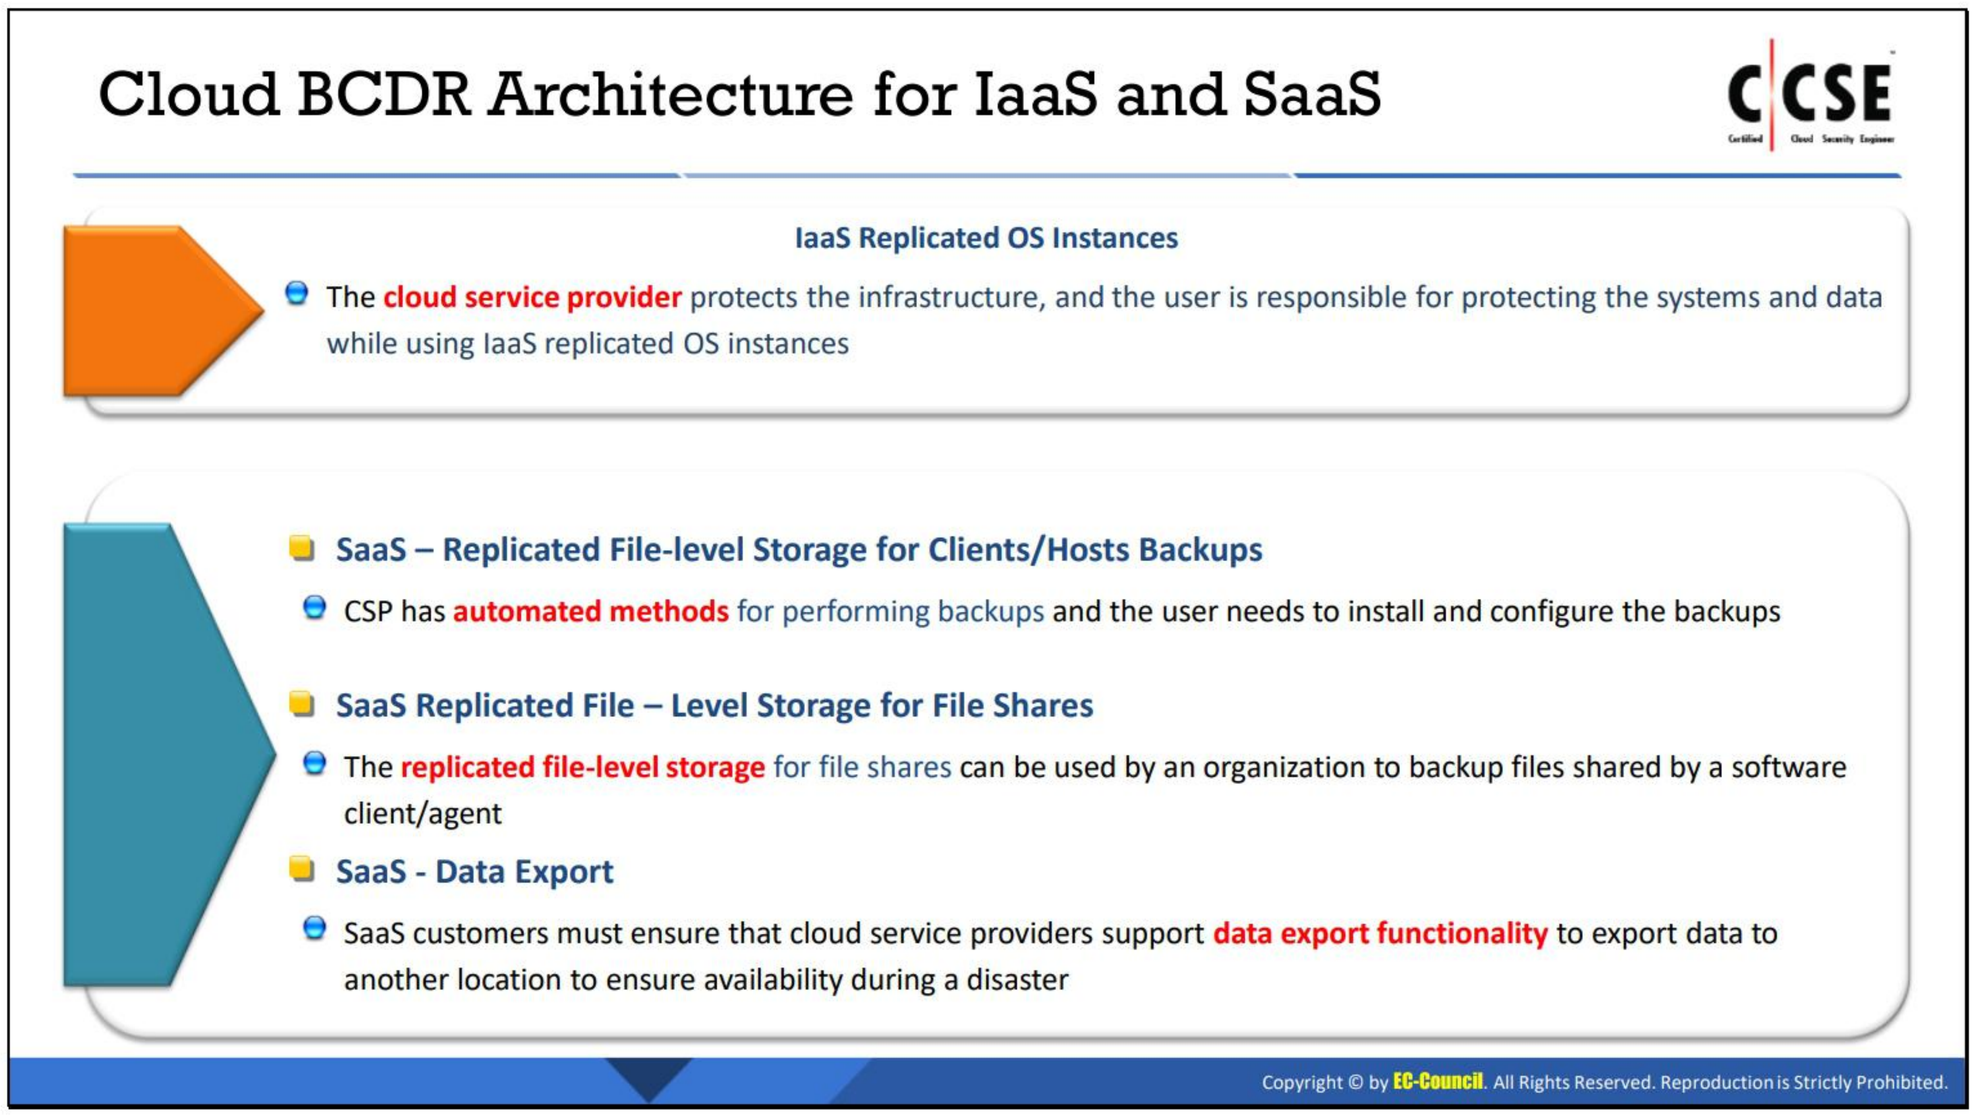

Cloud BCDR Architecture for IaaS and SaaS

IaaS Replicated OS Instances

For replicated OS instances, the shared responsibility model is followed. The CSP protects the infrastructure, and the user is responsible for protecting the systems and data. If IaaS is used for disaster recovery, the customer is responsible for all the configurations of the disaster recovery system. IaaS offers greater flexibility for disaster recovery, and the limitations are only from CSPs regarding the supported operating systems and versions. IaaS can be deployed for business that are less compatible with PaaS or SaaS.

SaaS – Replicated File-level Storage for Clients/Hosts Backups

CSPs have automated methods for performing backups, and the user needs to install and configure it. The user should test the backup every six months.

SaaS Replicated File-level Storage for File Shares

The replicated file-level storage for file shares can be used by an organization to backup files shared via a software client/agent. Thus, such a solution is implemented by small to mid-sized businesses.

SaaS - Data Export

SaaS customers must ensure that the CSP supports data export functionality to export data to another location to ensure availability during a disaster. SaaS data export can also be used to ensure that proper RPO and RTO are met. Customers should ensure that the CSP follows a standard for data export. Standards such as ASCII/a common data model and data dictionary/master data management (MDM), help in automated export of information.

Certified Cloud Security Engineer Copyright © by **EC-Council**
All Rights Reserved. Reproduction is Strictly Prohibited.

Disaster Recovery as a Service (DRaaS)

- Disaster recovery-as-a-service (DRaaS) is a **cloud backup model** for organizations where the applications and data are protected from disruptions by using cloud resources
- DRaaS is also called as **business continuity-as-a-service (BCaaS)**

DRaaS Models

- DRaaS models are as follows:
  - **Self-service DRaaS**: Self-service DRaaS is for organizations with **extensive internal expertise** to plan deploy and execute the DRaaS strategy
  - **Assisted DRaaS**: In assisted DRaaS, the **cloud provider** supplies all the tools required for performing backup operations and it assists in implementing the recovery plan
  - **Managed DRaaS**: In managed DRaaS, all the aspects of **disaster recovery** such as protection of the infrastructure resources, DR testing, and validation of the resources lie with the cloud provider

Copyright © by EC-Council. All Rights Reserved. Reproduction is Strictly Prohibited.

Disaster Recovery as a Service (DRaaS)(Cont'd)

hide01.ir

The benefits of DRaaS include:

Multisite Backup	The data is replicated to **multiple sites** and failover occurs during disasters to ensure continuous availability
Speed of Recovery	DRaaS solutions are equipped with **automated disaster recovery** capabilities for faster disaster recovery
Cost Savings	DRaaS is **cost-effective** considering additional expenses that may occur with traditional disaster recovery
Greater Access and Administration	DRaaS provides **greater access** as it is cloud-based, and it reduces the operational overhead of the IT team of the organization
Scalability	DRaaS is **scalable** according to customer needs as it is deployed on a virtual platform
Redundancy	DRaaS provides **cost-effective redundancy** and performs periodic testing to check whether the recovery plan is working as expected

Copyright © by EC-Council. All Rights Reserved. Reproduction is Strictly Prohibited.

Disaster Recovery-as-a-Service (DRaaS)

DRaaS is a cloud backup model where the applications and data are protected from disruptions by using cloud resources. DRaaS provides an organization with a total backup to allow BC in the event of a disruption caused by disaster. DRaaS allows the applications to run on VMs during a disaster and also without a disaster. DRaaS is also called as business continuity-as-a-service (BCaaS). The three DRaaS models are as follows:

Certified Cloud Security Engineer Copyright © by **EC-Council**
All Rights Reserved. Reproduction is Strictly Prohibited.

- **Self-service DRaaS**: Self-service DRaaS offers greater flexibility wherein the user itself is responsible for the disaster recovery solution. Self-service DRaaS is perfect for organizations with extensive internal expertise to plan deploy and execute the DRaaS strategy. DRaaS strategy management includes planning, testing, ongoing management, and management of failover. Thus, DRaaS service is suitable for users who have a better understanding of the storage, compute, and networking of the infrastructure as well as the relation between them. DRaaS helps in hardening the RTO and RPO goals by implementing automation at the IaaS level, and the critical application services can be prioritized for BC.
- **Assisted DRaaS**: In assisted DRaaS, the cloud provider supplies all the tools required for performing backup operations and it also assists in implementing the recovery plan. Thus, the management burden of the users gets reduced here. Assisted DRaaS provides additional resources for proper management of disaster recovery during the DR testing phase, which helps users perform their tests efficiently. The providers also offer consultation, advice, and a dedicated team of engineers for assistance, which helps users optimize the disaster recovery procedures.
- **Managed DRaaS**: In managed DRaaS, all the aspects of disaster recovery lie with the cloud provider. The service provider manages the protection of the infrastructure resources, DR testing, and validation of the resources. Users can achieve the best results for disaster recovery from the provider based on SLAs. Managed DRaaS also comes under a shared responsibility model, where the users need to communicate with the service providers regarding the changes made to the application and services to effectively manage disaster recovery. Managed DRaaS also provides a cross-platform technology where it provides failover with on-premises data center or across any other cloud provider.

Benefits of DRaaS

- **Multi-site backup**: As DRaaS is fully cloud-based, the data is replicated to multiple sites to ensure continuous availability if one or more sites are affected by a disaster.
- **Speed**: DRaaS solutions are equipped with automated disaster recovery capabilities, which helps harden the RTO and RPO objectives. DRaaS provides faster and reliable disaster recovery solutions.
- **Cost savings**: In-house disaster recovery programs include additional costs due to infrastructure, maintenance, and capital expenses. DRaaS is cost-effective considering the additional expenses that may occur with traditional disaster recovery.
- **Greater access and administration**: As DRaaS is a cloud-based solution, it can be accessed from anywhere. Similarly, because a professional service provider administers the disaster recovery process, the operational overhead of the IT team of the organization is reduced.
- **Redundancy and scalability**: As user data is stored on multiple places, the DRP is such that it eliminates single points of failures. DRaaS is scalable as it is stored in a virtual

Certified Cloud Security Engineer Copyright © by **EC-Council**
All Rights Reserved. Reproduction is Strictly Prohibited.

environment and can adapt according to the user needs. The service provider performs periodic testing of the DRP to check whether the DRP is working fine.

hide01.ir

Certified Cloud Security Engineer Copyright © by **EC-Council**
All Rights Reserved. Reproduction is Strictly Prohibited.

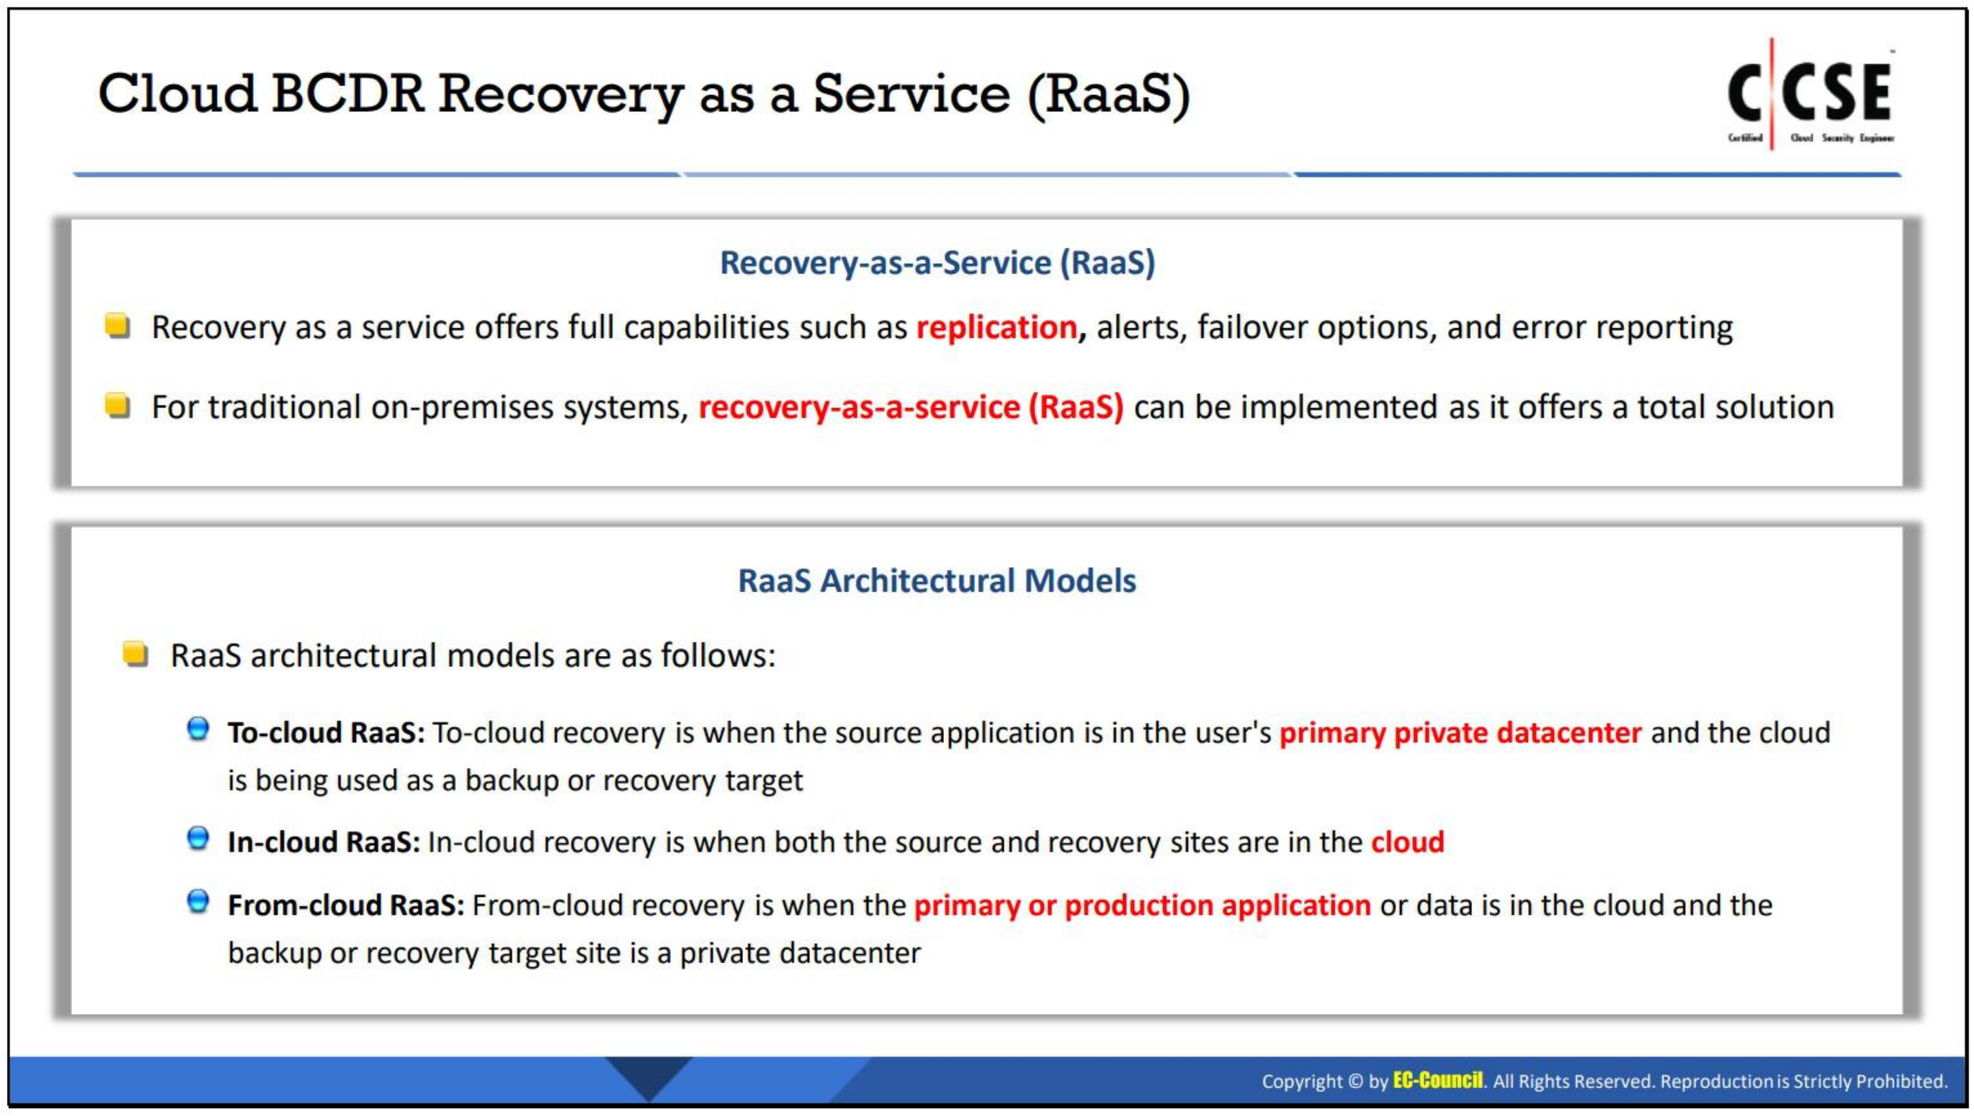

Cloud BCDR Recovery-as-a-Service (RaaS)

Recovery-as-a-Service (RaaS)

RaaS is a complete service, which includes the following:

hide01.ir

- Replication to the CSP.
- Alerts regarding issues.
- Automated failover options.
- Error reporting.
- Compression and replication.

Implementation Considerations for Different Models

- **Traditional on-premises systems:** For traditional on-premises systems, RaaS offers a total solution. The CSP helps in hosting the hardware at the local site or provides partially automated failover services. The CSP also helps in the failback process to return the business to normal operations.
- **Cloud-based systems recovery:** A business fully implemented in cloud can use the in-built BCDR capabilities used by the cloud. If an IaaS solution is used, CSPs offer automatic failover and automatic replication capabilities. RaaS offers all available capabilities for replication.

Certified Cloud Security Engineer Copyright © by **EC-Council**
All Rights Reserved. Reproduction is Strictly Prohibited.

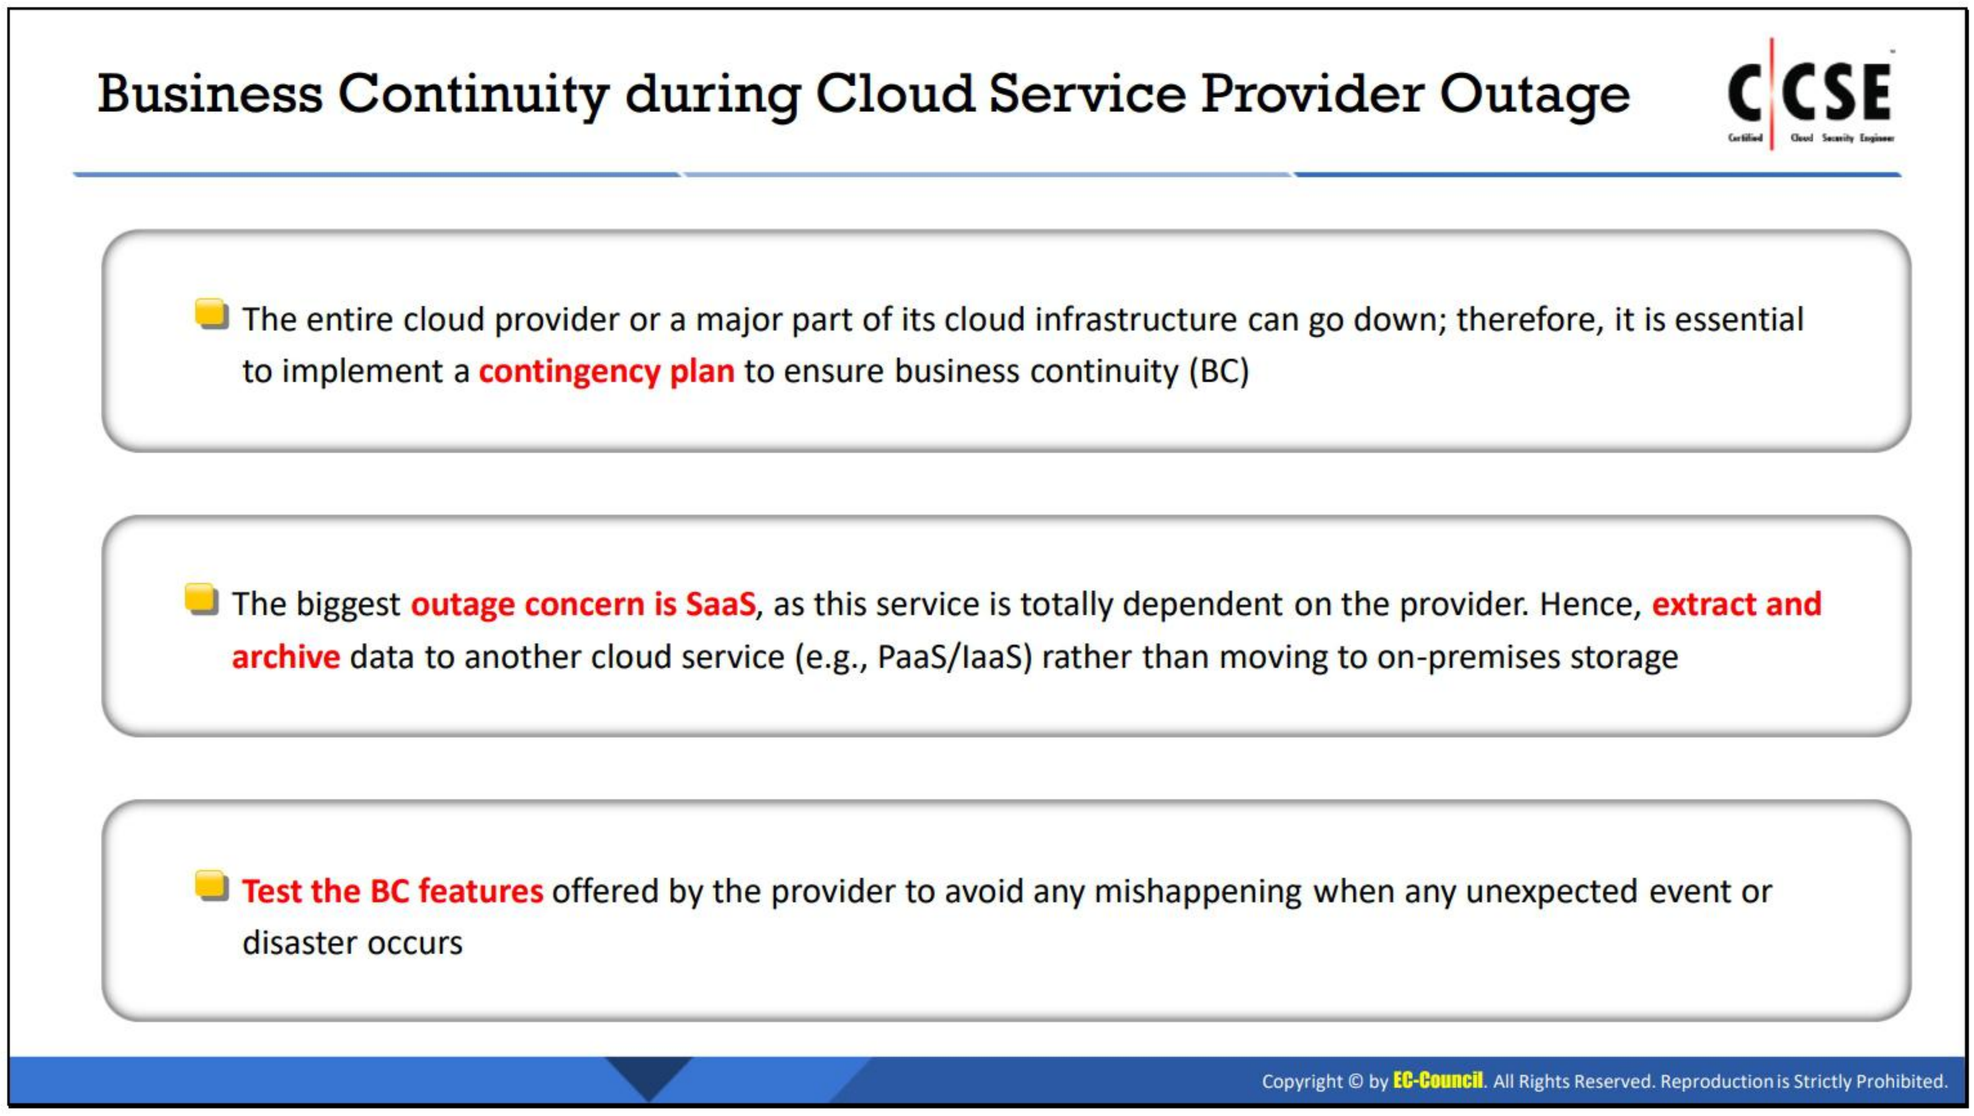

Business Continuity during Cloud Service Provider Outage

In certain scenarios, an entire CSP or a major part of its cloud infrastructure (located at one geography) can fail. In such cases, it is essential to implement a contingency plan to ensure BC. It is not easy to plan for CSP outages because of the lock-in agreement with the provider. In certain cases, migration to another part of the provider's service is possible; however, in other cases, there is no chance of internal migration. An organization must accept this risk based on the history of the provider and its internal availability capabilities.

Another option that the organization should consider is downtime. Downtime mainly depends on the RTO of the organization. Through DNS redirection, the organization should ensure that some sort of static standby is made available to serve its users. It must include failure responses to API calls (if they offer APIs) as part of graceful failure or fault tolerance.

Have a Contingency Plan

The organization should be careful while considering a CSP as its provider. It should consider the benefits and risks of moving ahead with the provider in terms of BC. It should have a contingency plan to handle possible disaster situations at the provider's end. Various scenarios must be defined in the plan that define the course of action when disasters hit the provider. The organization must ensure that the contract signed with the provider contains provisions that define what happens with data in case of provider bankruptcy, service changes, etc.

While choosing the provider, the organization must ensure that the provider takes data backups regularly. It should also know how the provider manages these backups and how they can be used for BC. CSPs usually provide backup services for an additional cost.

Certified Cloud Security Engineer Copyright © by **EC-Council**
All Rights Reserved. Reproduction is Strictly Prohibited.

Consider Multiple Cloud Providers

The organization must choose a secondary service or provider if its service is reliant on the same provider. Most large organizations prefer to have two or more cloud providers to minimize risk. If one provider goes out of business, they can use the other provider to ensure BC. Moving data between different providers may be difficult. However, it is relatively easier than moving security controls, meta-structure, logging, etc., that might be incompatible between cloud platforms.

Organizations that have multiple CSPs need to document the process for their data and application migration from one provider to another when any service interruption happens. Additionally, they need to test this process regularly to ensure it is working as intended. The BC plan must include all important aspects that affect an organization's business such as RTO, RPO, and MTD.

The biggest outage concern is SaaS because this service is totally dependent on the provider. With SaaS, BC options include scheduled data extraction and data archiving, apart from accepting downtime. It is better to extract and archive data to another cloud service (e.g., PaaS/IaaS) rather than move it to on-premises storage. The organization should consider the history of the provider as part of its risk-based approach toward BC.

In addition, the organization must have an alternate application where it can migrate its data even when it has its data because if it cannot use the data available at its end, it will not have a viable recovery strategy.

Finally, test the BC features offered by the provider to avoid any mishappening when any unexpected event or disaster occurs. This will help evaluate the effectiveness of the BC features and fix issues, if any.

h1te01.ir

Certified Cloud Security Engineer Copyright © by **EC-Council**
All Rights Reserved. Reproduction is Strictly Prohibited.

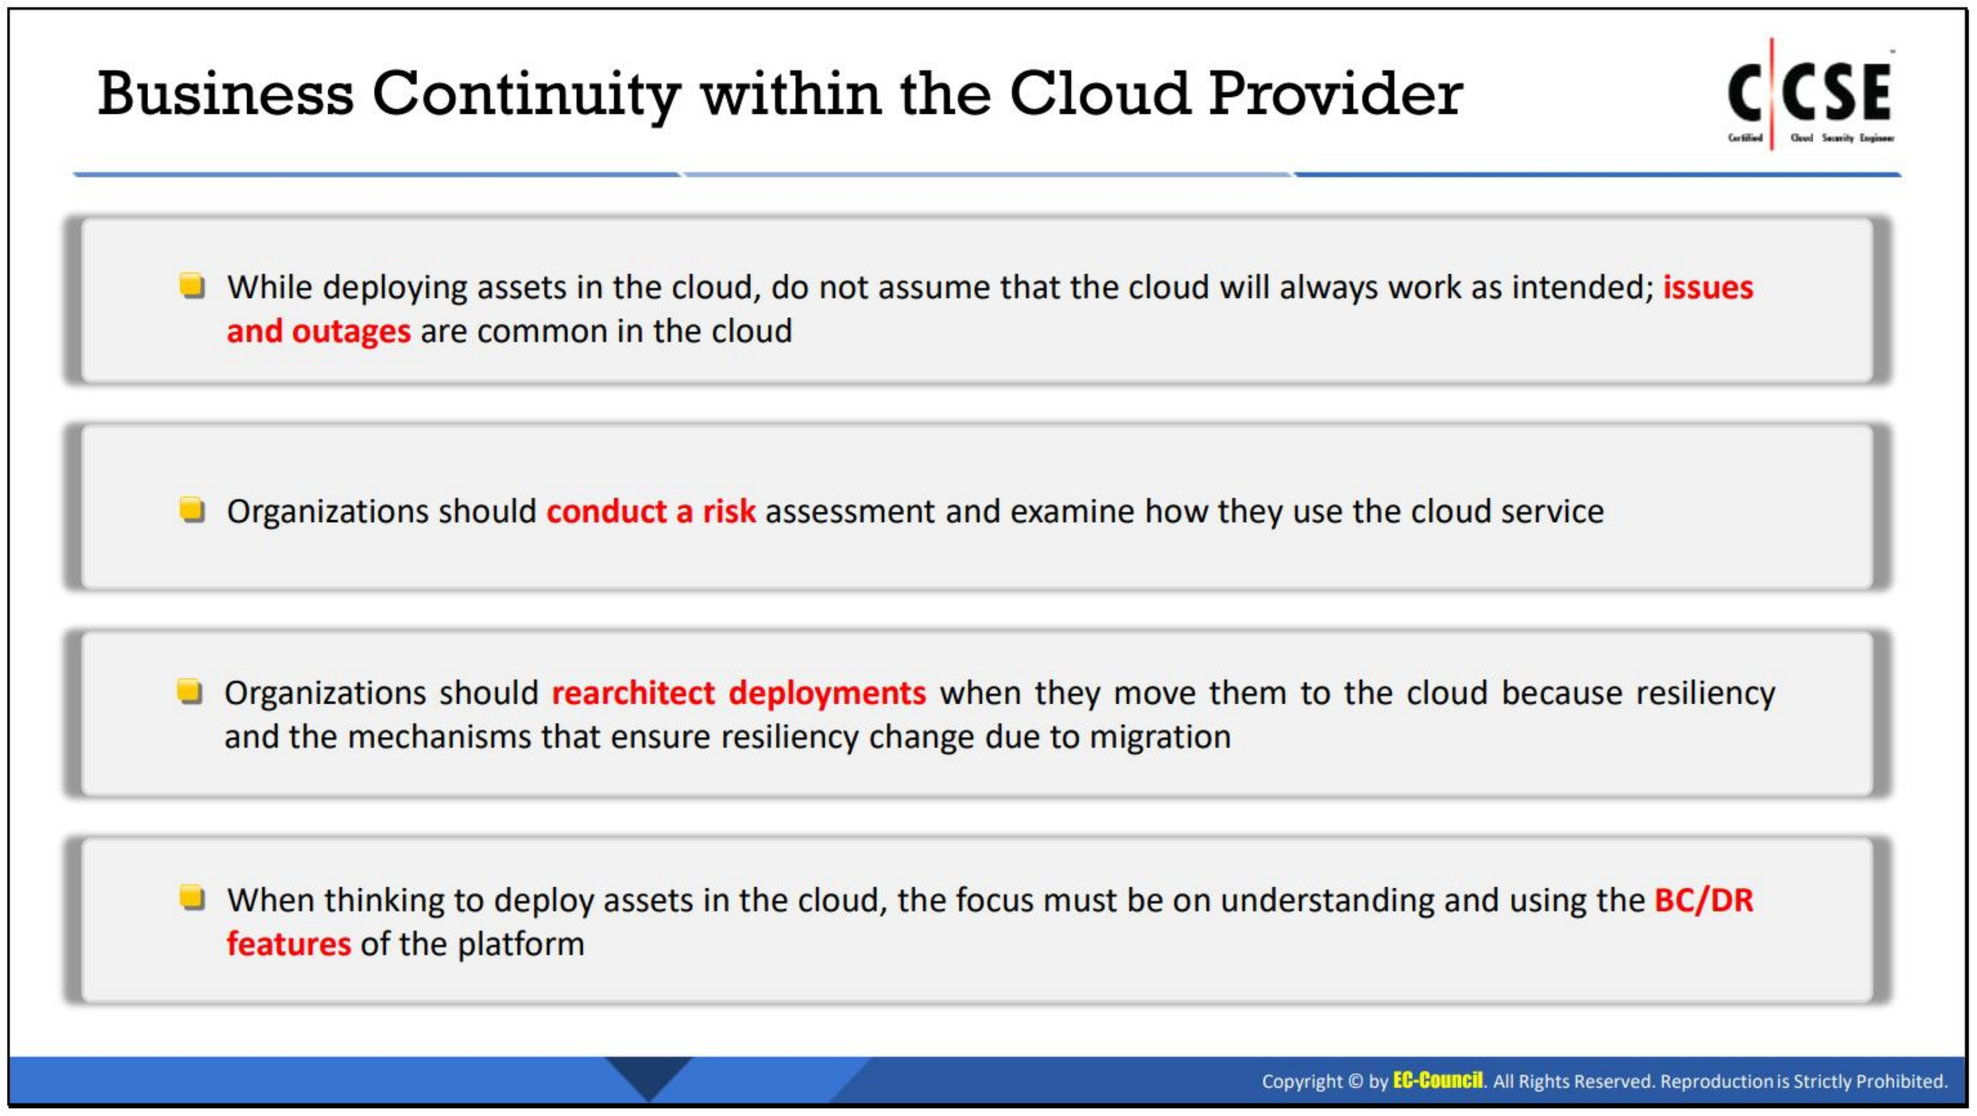

Business Continuity within the Cloud Provider

While deploying assets in the cloud, it cannot be assumed that the cloud will always work as intended. Like other technologies, issues and outages are common in the cloud. However, the CSP can make the cloud more resilient by leveraging mechanisms that help build resilient applications.

hide01.ir

For single assets (e.g., a VM), virtualizing resources into pools creates less resiliency. Using software to manage things and abstracting resources increase flexibility to enable resiliency features more easily. Resiliency features include cross-geographic load balancing and durable storage.

Various cloud providers are available in the market; however, all are not the same in terms of their offerings. The cloud is not always better than traditional IT infrastructure and a decision in this regard needs to be made on a case-by-case basis. Using a risk assessment and examining how they use the cloud service will help organizations make better decisions in this regard.

Therefore, it is best for organizations to re-architect deployments when they move them to the cloud. This is because both resiliency as well as mechanisms that ensure resiliency change due to migration. Re-architecting deployments will ensure that resiliency is maintained even in the cloud. Direct lift-and-shift migrations can neither resolve failures nor benefit from potential improvements by using the platform or service-specific capabilities.

When thinking to deploy assets in the cloud, the focus must be on understanding and using the BCDR features of the platform. After deciding to deploy in the cloud, optimize the use of BCDR features provided by the platform before adding extra capabilities via third-party tools.

Certified Cloud Security Engineer Copyright © by **EC-Council**
All Rights Reserved. Reproduction is Strictly Prohibited.

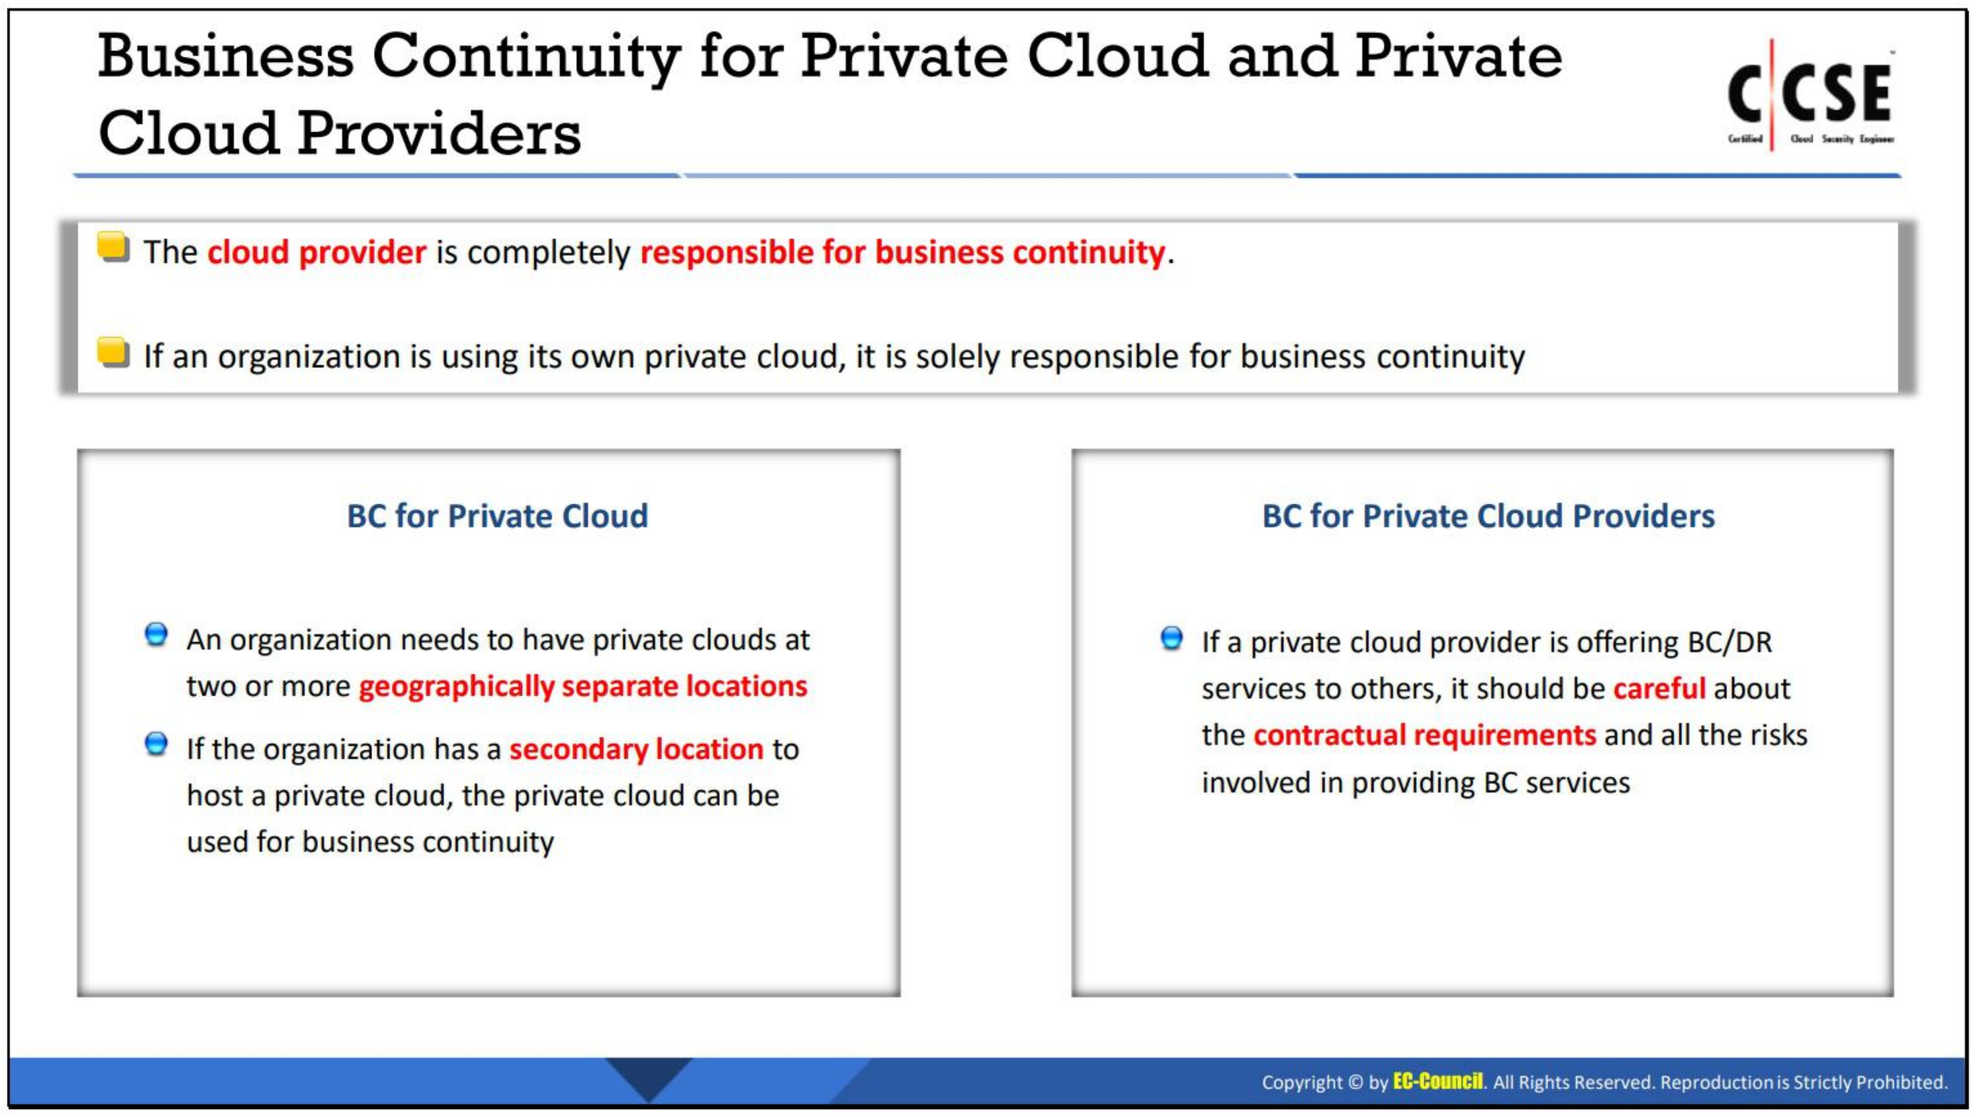

Business Continuity for Private Cloud and Private Cloud Providers

The cloud provider is completely responsible for BC. If an organization is using a private cloud, it is solely responsible for BC. BCDR involves everything, ranging from cloud services to physical facilities. For BC, RTOs and RPOs are tight because business processes will halt if the cloud goes down.

BC for Private Cloud

Private clouds can also be used for BCDR functions. For this purpose, an organization needs to have private clouds in at least two or more geographically separate locations. As the underlying hardware for these clouds is controlled by the organization's IT team, it can easily customize them to be used for BC as per its business requirements. The organization must employ identical hardware for private clouds on both sites.

If the organization has a secondary location to host a private cloud, the private cloud can be used for BC. In such a case, the organization can use the cloud infrastructure hosted in the secondary site for running its test and development environments. If any unexpected event occurs that requires moving of the production environment to the secondary site, the test and development environments can be down prioritized and cloud resources must be allocated to the production environment. To ensure that the secondary site can be used for BCDR, data must be replicated from the primary site to the secondary one. As the organization has full control over both sites, it can choose the technology it wants to use. It must perform frequent failover testing to verify whether the solution is operational.

BC for Private Cloud Providers

If a private cloud provider is offering BCDR services to others, it should be careful about contractual requirements such as data storage while creating its BC plans. It must carefully

Certified Cloud Security Engineer Copyright © by **EC-Council**
All Rights Reserved. Reproduction is Strictly Prohibited.

evaluate all the risks involved in providing BC services. For example, failing over of services to another geography that has a different legal jurisdiction may violate local laws or contracts.

Multiple private cloud providers are now offering DRaaS as part of their offerings. DRaaS enables organizations to use their own data centers for hosting and managing their production environments and leverages a CSP to provide BCDR as a service to a cloud site. DRaaS is offered by providers at reasonable rates; therefore, organizations can eliminate the need for a managed platform and capital expenditure for achieving BC. With DRaaS, CSPs are responsible for managing the BCDR solution and offering organizations SLAs around RPOs and RTOs. This eliminates the overhead of managing RPOs and RTOs because CSPs provide contracted guarantees.

As a strong working relationship is necessary between the organization and CSP, most organizations prefer to choose a local CSP to ensure smooth BCDR functions.

hide01.ir

Certified Cloud Security Engineer Copyright © by **EC-Council**
All Rights Reserved. Reproduction is Strictly Prohibited.

BCDR Implementation Guidelines

The cloud BCDR implementation guidelines are as follows:

- **Implement IAM (Federated Identities):** With federation, **identities repeated provisioning**, security loopholes, and user inconvenience can be reduced
- **Replication of Files and Data:** Different **types of replication** can be done for data recovery
  - **File-level replication** can be performed to copy only the modifications in files
  - **File copy** can be performed for copying the entire file
  - **Block-level replication** can be performed for replicating block-level changes in the storage device
- **Clustering:** Clustering allows ensuring **continuous services** by failover of hardware/software within a group of hardware and software resources
- **Server-Side Encryption:** Server-side encryption makes sure that the **user data** after getting uploaded to the cloud provider are safe
- **Client-Side Encryption:** Client-side encryption helps in ensuring the **confidentiality** and integrity of the data

Copyright © by EC-Council. All Rights Reserved. Reproduction is Strictly Prohibited.

BCDR Implementation Guidelines (Cont'd)

hide01.ir

- **Transport/DIM Encryption:** Transport layer security (TLS) and SSL protocols should be supported by cloud service providers for **security in transit**
- **Active Directory Federation Services (ADFS):** Active Directory Federation Services (ADFS) can be used for **single sign-on (SSO)** mechanisms if client and server use Microsoft windows
- **Security Assertion Mark-up Language (SAML):** SAML is an XML-based **authentication standard** for single sign-on
- **Transaction Journaling:** The systems can be **restored to a particular point** by journaling, which is an audit feature
- **Realm–Based Access Control:** Realm-based access control is **domain-based access control** that can be used for large business

Copyright © by EC-Council. All Rights Reserved. Reproduction is Strictly Prohibited.

BCDR Implementation Guidelines

The cloud BCDR implementation guidelines are as follows:

- **Identity and access management (IAM) – federated identities for BCDR:** Federated identity management can be used for the secure exchange of information. With federated identities, repeated provisioning, security loopholes, and user inconvenience can be

Certified Cloud Security Engineer Copyright © by **EC-Council**
All Rights Reserved. Reproduction is Strictly Prohibited.

reduced. The federated identity management also helps in the login barriers, better customer service, and collaboration with partners.

- **Replication of files and data:** Replication can be explained as the complete copying of the entire system to DR site for recovery from a disaster. The types of replication that can be performed are
    - **File-level synchronization**: In file synchronization, only the modifications are copied over, which is bandwidth efficient. Therefore, the files are synchronized across different locations for ensuring availability.
    - **Block-level replication**: In block-level replication, block-level changes are replicated. Block-level is not compatible for immediate failover operations, therefore, it cannot be used for systems that are required to be live.

The following factors are to be considered for data replication in the cloud.

- Systems for which data is replicated.
- Mode of disaster recovery (such as data replication, live system replication, etc.).
- Bandwidth requirements.
- Type of replication supported by the CSP.

Clustering

A group of hardware and software resources that are joined as one endpoint are clustered. Clustering ensures continuous services through failover of hardware/software within this group. As clustering is specific to one location, location failures affect the clustering. For efficient clustering, services should be hosted on multiple IaaS instances through horizontal scaling and high availability.

Server-side encryption

Server-side encryption makes sure that the user data after getting uploaded to the cloud provider is safe. Server-side encryption ensures the security of the data at rest.

Client-side encryption

Client-side encryption ensures confidentiality and integrity of the data. Data at rest (DAR) encryption solutions can be used for client-side encryption.

Transport/DIM encryption

Transport layer security (TLS) and secure sockets layer (SSL) protocols should be supported by CSPs. Secure shell (SSH) helps in ensuring the security of the data in transit.

ADFS

Active directory federation services (ADFS) can be used for single sign-on (SSO) mechanisms if client and server use Microsoft Windows OS. ADFS use the lightweight directory access protocol (LDAP). The authentication and authorization can be controlled by AD group policy objects by enabling role-based access controls (RBAC).

Certified Cloud Security Engineer Copyright © by **EC-Council**
All Rights Reserved. Reproduction is Strictly Prohibited.

Security assertion mark-up language (SAML)

SAML can be used for SSO process. It is an XML-based authentication and authorization standard. SAML is a web-based platform and is offered by an identity provider on behalf of a cloud user.

Transaction journaling

Journaling provides an audit feature. The systems can be restored to a particular point by journaling. Transaction journaling also provides a detective safeguard for auditors.

Realm-based access control

Realm-based access control refers to domain-based access control. Realm-based access control is used by large businesses that have multiple different domains.

hide01.ir

Certified Cloud Security Engineer Copyright © by **EC-Council**
All Rights Reserved. Reproduction is Strictly Prohibited.

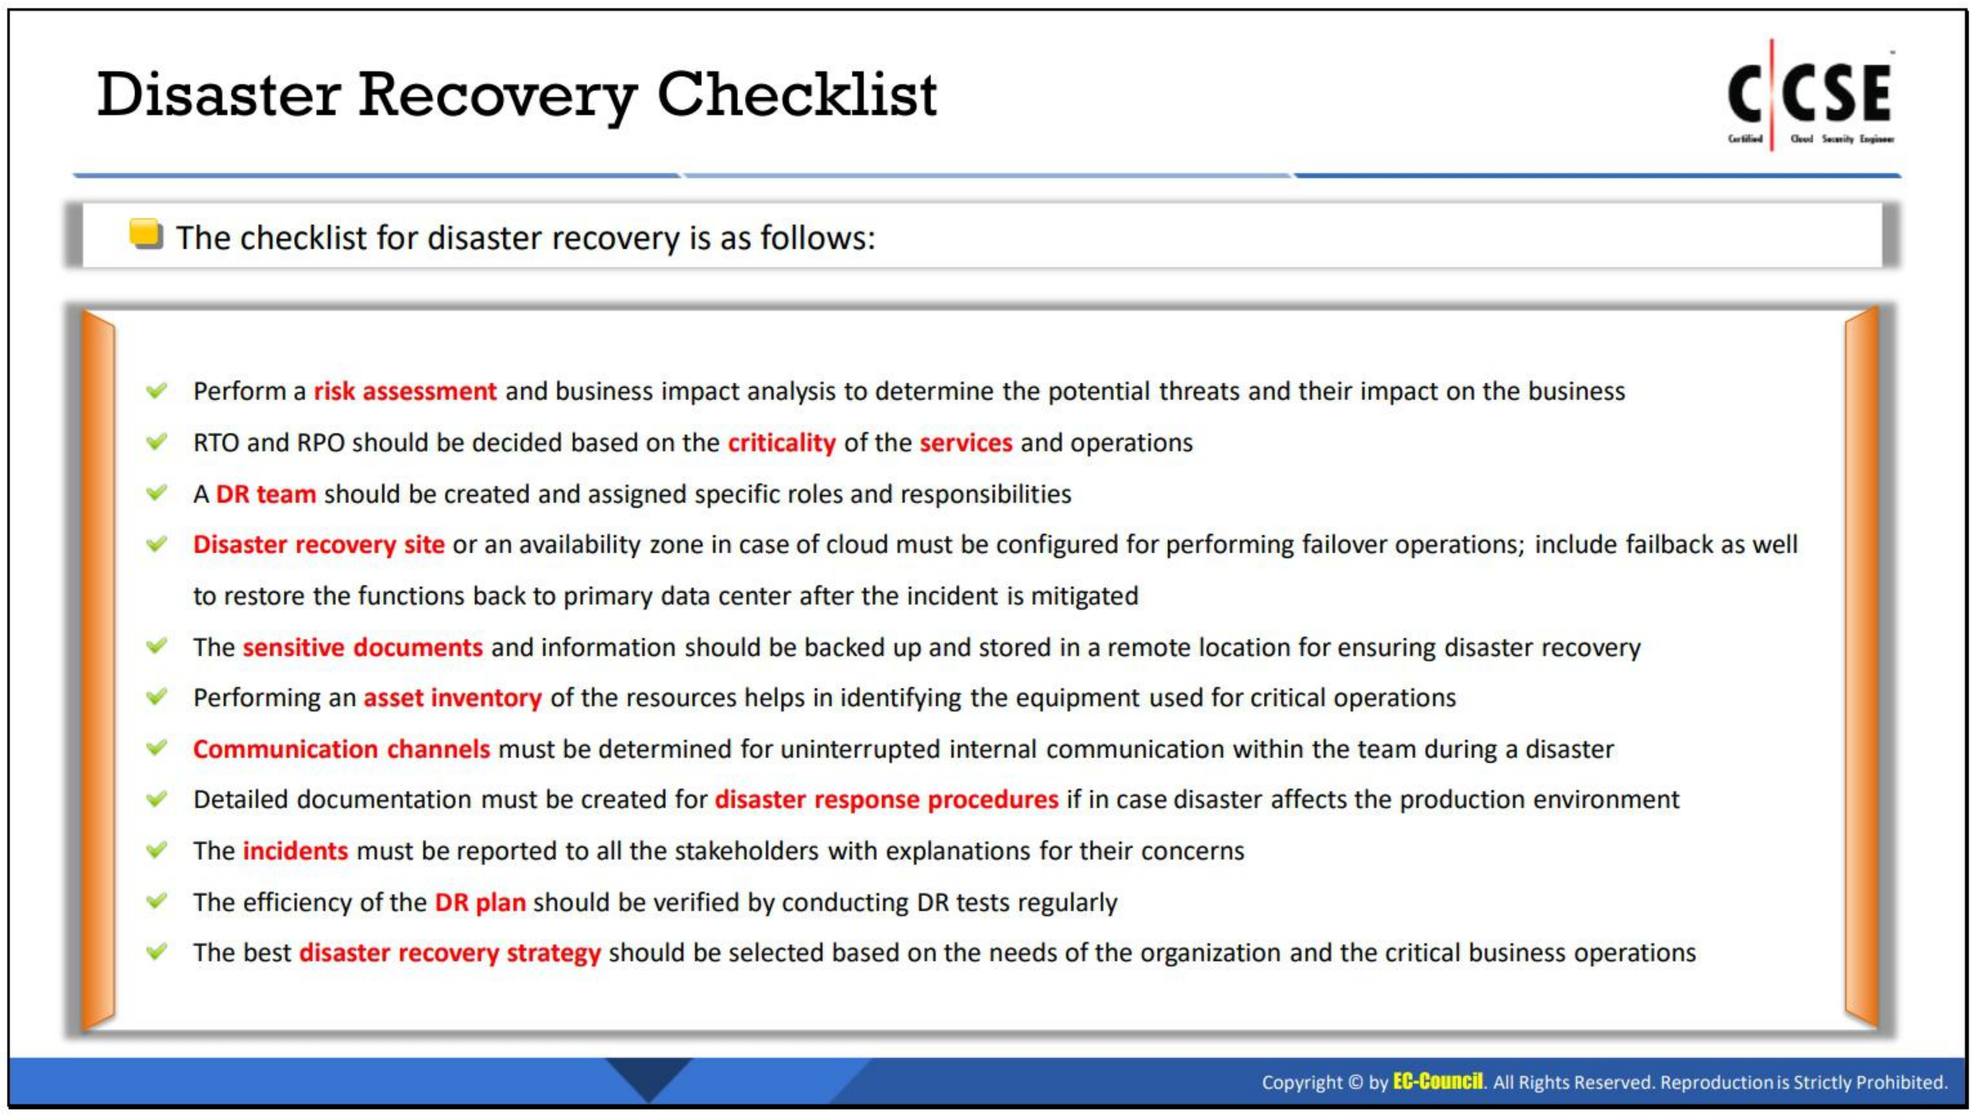

Disaster Recovery Checklist

The checklist for disaster recovery is as follows:

- **Perform a risk assessment and BIA to determine the potential threats and its impact on the business:** A risk assessment on the resources in the organization helps in identifying the potential threats and vulnerabilities affecting the business. A BIA can be performed to determine the impact of the particular vulnerability or threat. In the DR plan prevention and mitigation measures can be included to address this issue.
- **Determine the RTO and RPO:** The critical functions and operations of the business needs to be analyzed. RPO and RTO helps in determining the recovery strategies to be adopted for these critical operations.
- **Create a DR team and assign them roles and responsibilities:** For the effective working of a DR plan, a team should be created and assigned specific roles and responsibilities to respond to an incident. This will help in eliminating confusion among the personnel regarding what to perform when a disaster strikes.
- **Create disaster recovery site to perform failover:** To resume normal operations of a business and to ensure continuous availability, a DR site (availability zones (AZs) in cloud) should be created to perform failover. Failback should also be included in the DR plan as the DR site cannot function as a primary data center.
- **Backup sensitive information:** Sensitive information is an important asset for an organization. Loss of sensitive information will cause failure of the DRP. Backup the sensitive information in remote locations to continue with normal business operations. Additionally, infrastructure documentation should be performed, which contains information about the original assets of the organization.

Certified Cloud Security Engineer Copyright © by **EC-Council**
All Rights Reserved. Reproduction is Strictly Prohibited.

- **Asset inventory:** A complete asset inventory of the organization helps in determining the equipment required for restoring the normal state and the assets that are important for critical operations.
- **Communication channels:** During a disaster, all the team members must be able to communicate without interruption to successfully perform disaster recovery. The communication channels must be determined for such a situation when the network gets affected.
- **Detailed documentation of disaster recovery procedures:** Even with all the security measures in place, a disaster affects the organization. A detailed documentation explaining the disaster response procedures should also be part of the DR plan. The documentation must include execution of DR activities, the personnel roles, and responsibilities for DR process, failover procedure, and verification that the system has recovered.
- **Report incidents to stake holders:** The incidents must be reported to all the customers, third-party suppliers, and vendors, and appropriately explain them their concerns. Creating a press release should also be part of a DR plan to ensure that there is no gap in communication when a disaster strikes.
- **Verify and update the DR plan:** As new resources are regularly added to organizations, the current DRPs might not be sufficient to address new security issues. Therefore, the DR plan should be updated based on changing needs of the organization. A DR test can be performed periodically to ensure that the DR plan is efficient in handling a disaster.
- **Select the best disaster recovery strategy:** DR strategies selected must be based on the requirements. It should be such that the selected strategy complies with the business needs. The disaster recovery strategy can be either a do-it-yourself strategy or assigned to a third-party vendor based on every security aspect.

Certified Cloud Security Engineer Copyright © by **EC-Council**
All Rights Reserved. Reproduction is Strictly Prohibited.

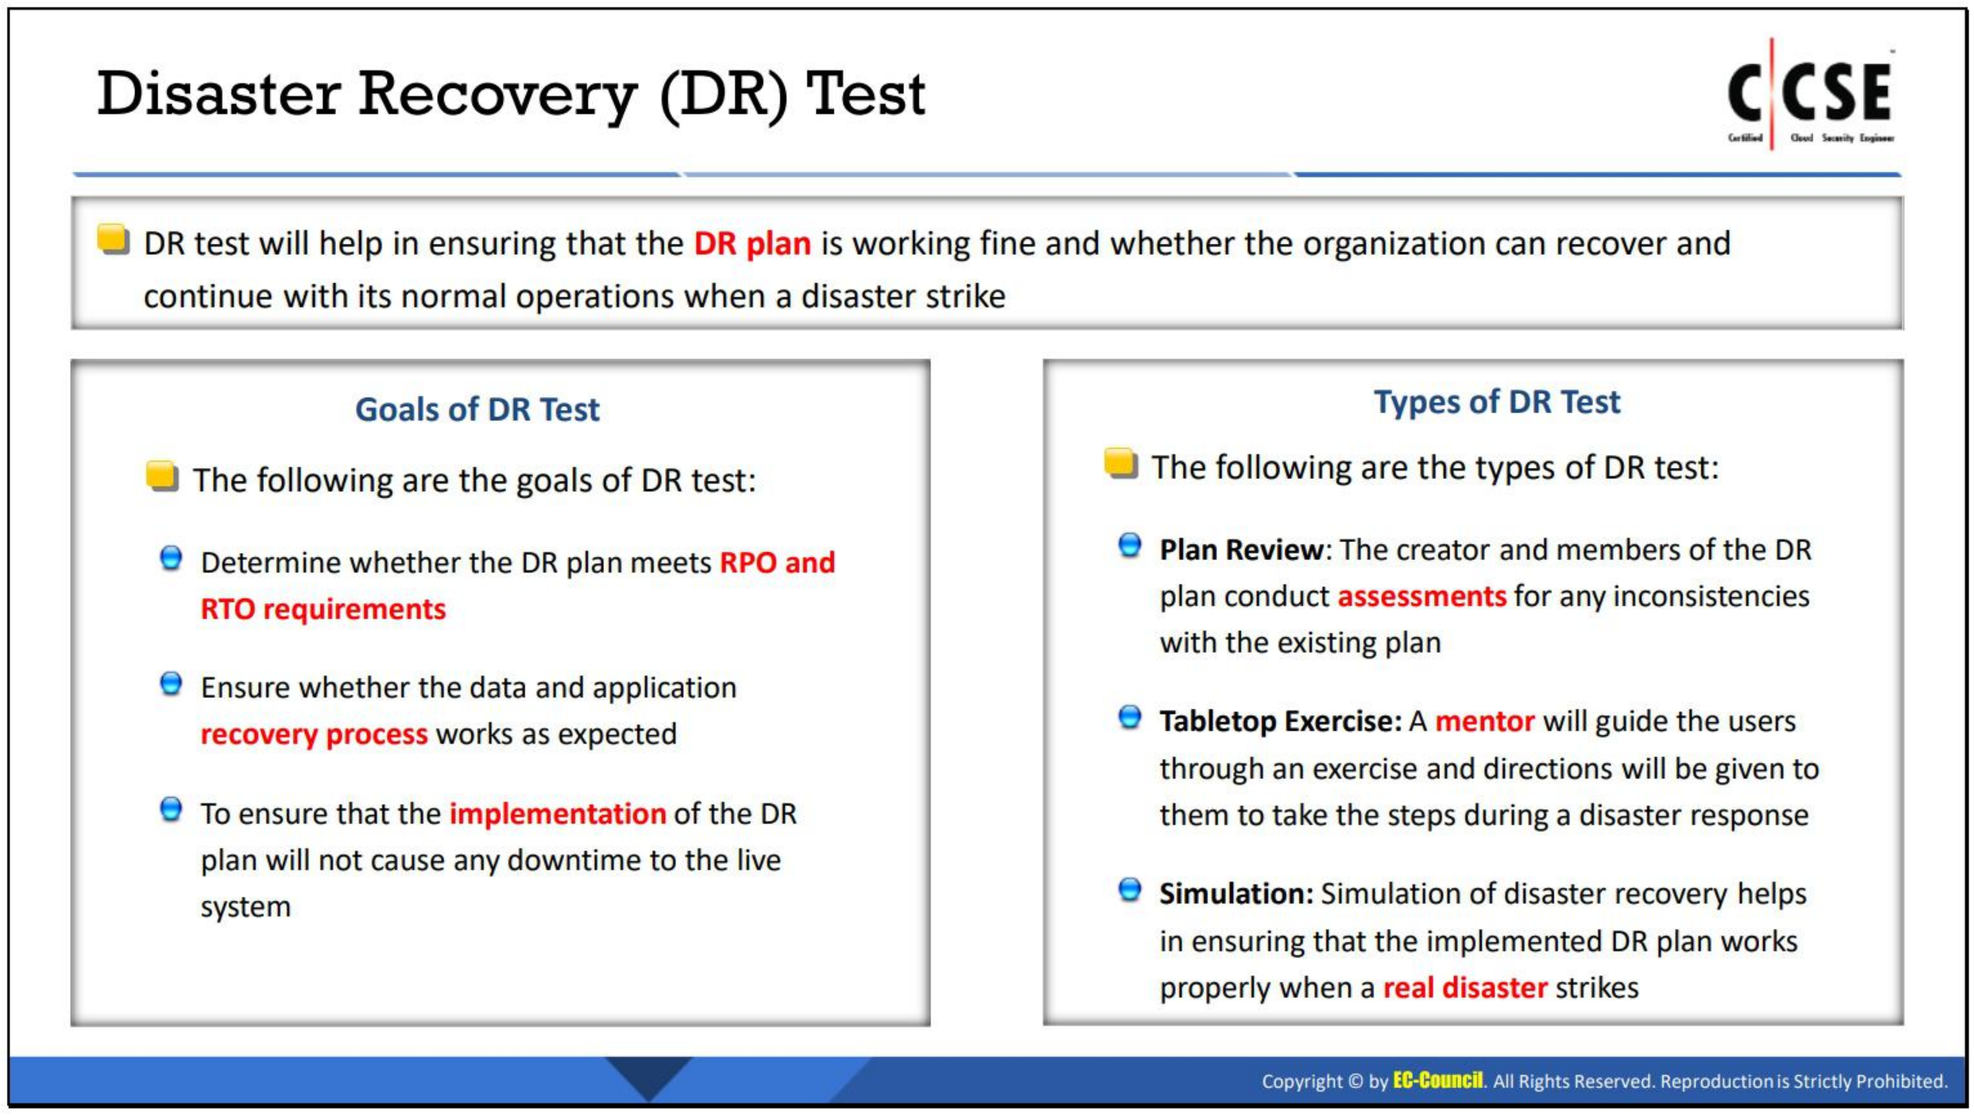

Disaster Recovery (DR) Test

Verification of each step in the DRP is essential to ensure that the DR plan works as expected; this is referred to as a disaster recovery test. A DR test helps in ensuring that the DR plan is working fine and that the organization can recover and continue with its normal operation when a disaster strike. If the organization fails to conduct a DR test, the DR plan may not function as expected when disaster strikes.

After the DR test is conducted, the logs should be analyzed to determine

- What all worked efficiently?
- What did not function properly?
- Changes to be made in DR plan for proper operations.

Goals of DR Test

- Determine whether the DR plan meets RPO & RTO requirements.
- Ensure whether the data and application recovery process work as expected.
- DR test needs to be conducted when new products or new applications are added into the infrastructure to ensure its compliance with disaster recovery.

Types of Disaster Recovery Tests

The three types of disaster recovery testing are

- **Plan review**: In plan review, the creator and other team members of the DRP conduct an assessment to verify whether there are any inconsistencies and anything missing with the existing plan.

Certified Cloud Security Engineer Copyright © by **EC-Council**
All Rights Reserved. Reproduction is Strictly Prohibited.

- **Table-top exercise:** Table-top exercise deals with simulated disaster scenarios, which helps the users of the organizations get familiarized with the disaster response implementation. The administrators can also ensure the effectiveness of the organization in handling the disaster via this exercise. In table-top exercises, a mentor guides the users through an exercise where directions will be given to them regarding the steps to be taken during a disaster response. This helps users ensure that they are clear regarding the guidelines for their responsibilities in disaster response and any missing information can be added.
- **Simulation**: Simulation of disaster recovery helps in ensuring that the implemented procedures for backup and recovery during a disaster work properly when it actually strikes.

hide01.ir

Certified Cloud Security Engineer Copyright © by **EC-Council**
All Rights Reserved. Reproduction is Strictly Prohibited.

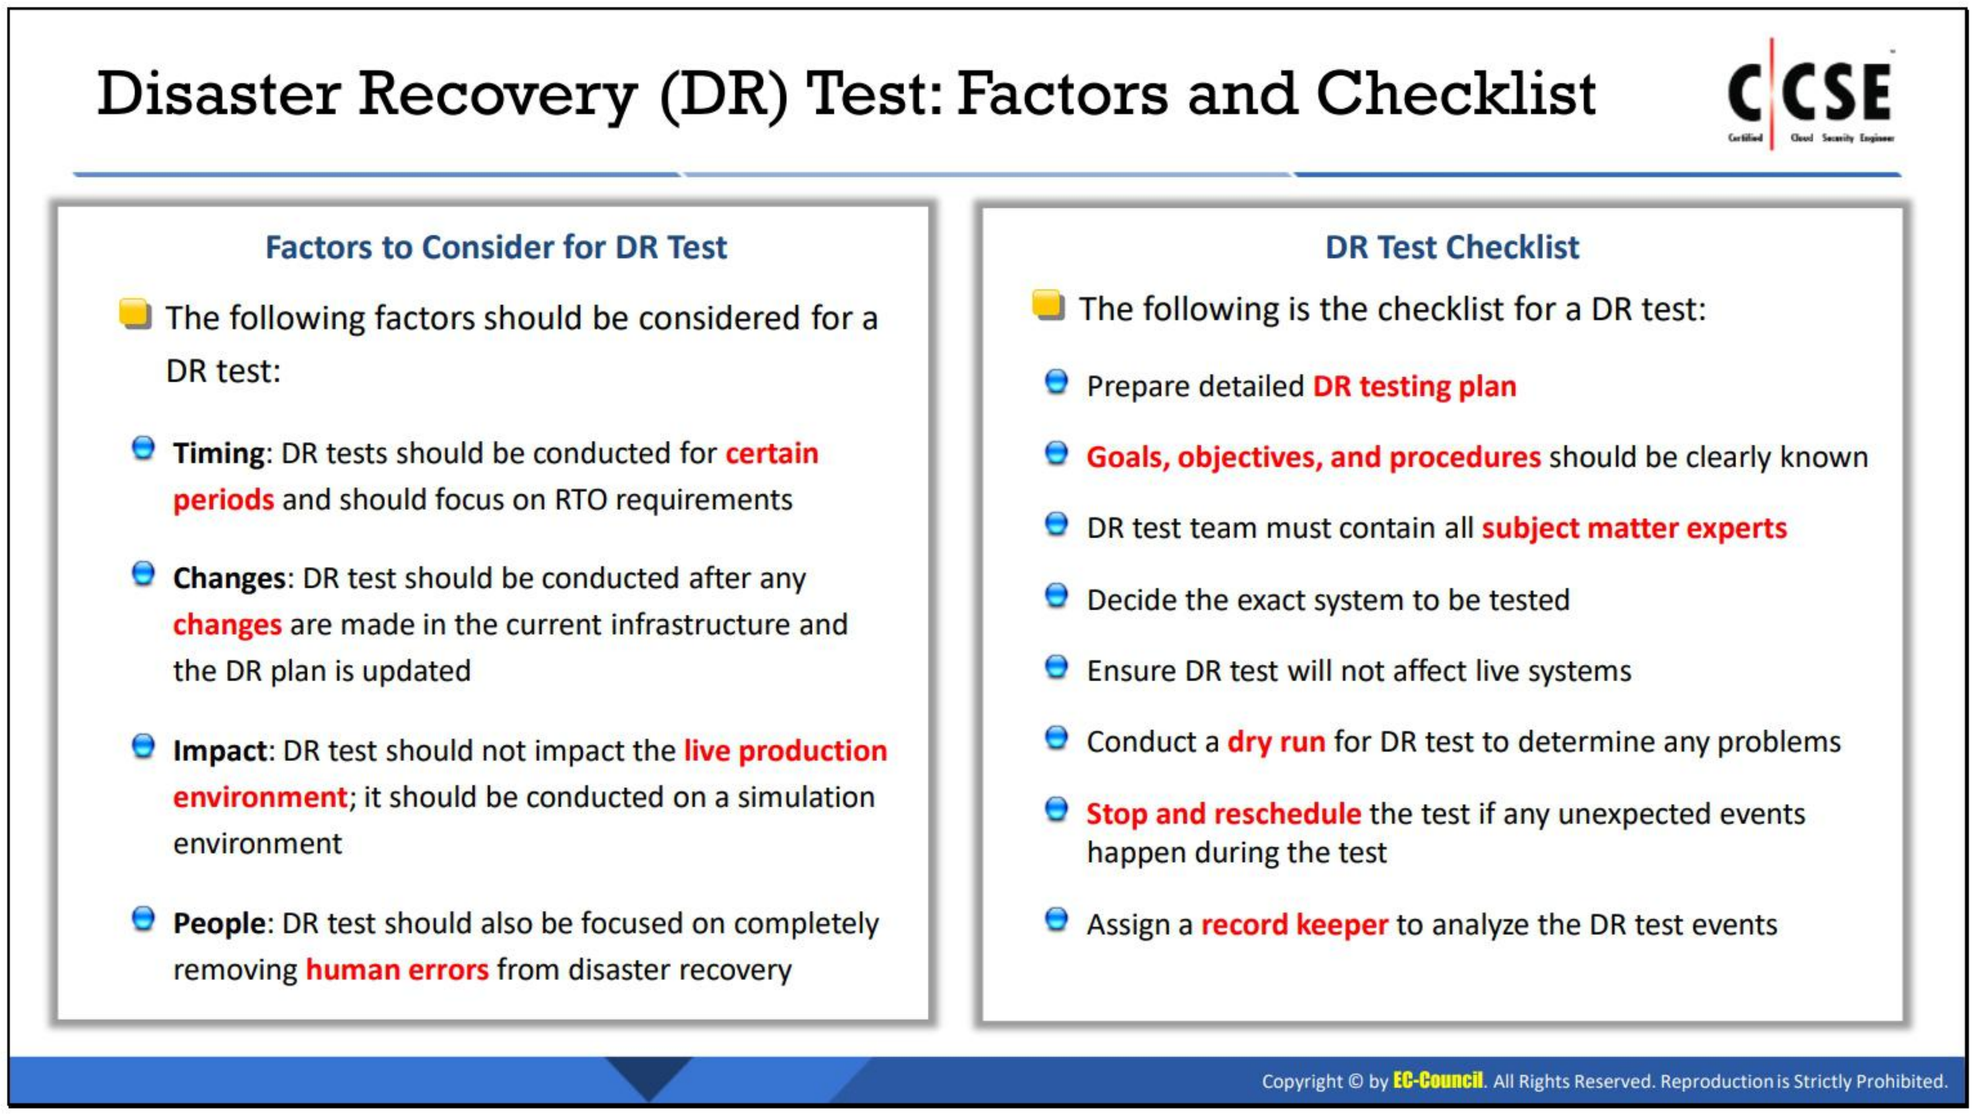

Disaster Recovery (DR) Test: Factors & Checklist

Factors to be considered for Disaster Recovery (DR) Test

The factors to be considered for an effective disaster recovery testing are

- **Timing:** As the time duration between the disaster recovery tests increases, it may cause failure of the DRP. This happens because as resources get added to the infrastructure, a DR test should be done to ensure that the added resources comply with the implemented DR plan. The recovery time from the RTO point of view is also to be considered.
- **Changes**: The DR test also confirms whether the implemented backup-and-restore plan is unaffected by changes. The test should be conducted after any changes in the current infrastructure and DR plan should be rewritten based on the test results.
- **Impact**: It should also be ensured that the disaster recovery test does not cause any downtime in the production environment. The DR test should be conducted in a disaster simulation environment; however, it should also be ensured that the DR plan works similarly when a real disaster strikes.
- **People**: The DR test can also be focused on completely removing human errors from disaster recovery. Apart from employees, a variety of other people/experts can be deployed to validate the effectiveness of the DR plan.

DR Testing Checklist

A DR testing checklist must contain the following:

- A detailed DR testing plan should be created to run DR tests.

Certified Cloud Security Engineer Copyright © by **EC-Council**
All Rights Reserved. Reproduction is Strictly Prohibited.

- Before conducting the DR test, the goals, objectives, and procedures should be clearly known.
- The DR test team should contain all the subject matter experts.
- Decide the exact system to be tested and make changes to the DR plan after the test.
- Make sure that the DR test does not affect any of the live production systems.
- A dry run before the actual test helps in determining any potential issues.
- The DR test should be stopped if any issue arises; verify the issue and then either continue or reschedule.
- A record keeper can be designated to record what occurred during the test, what did not work as expected, and what improvements can be made.
- The DR plans should be updated based on the DR test.

hide01.ir

Certified Cloud Security Engineer Copyright © by **EC-Council**
All Rights Reserved. Reproduction is Strictly Prohibited.

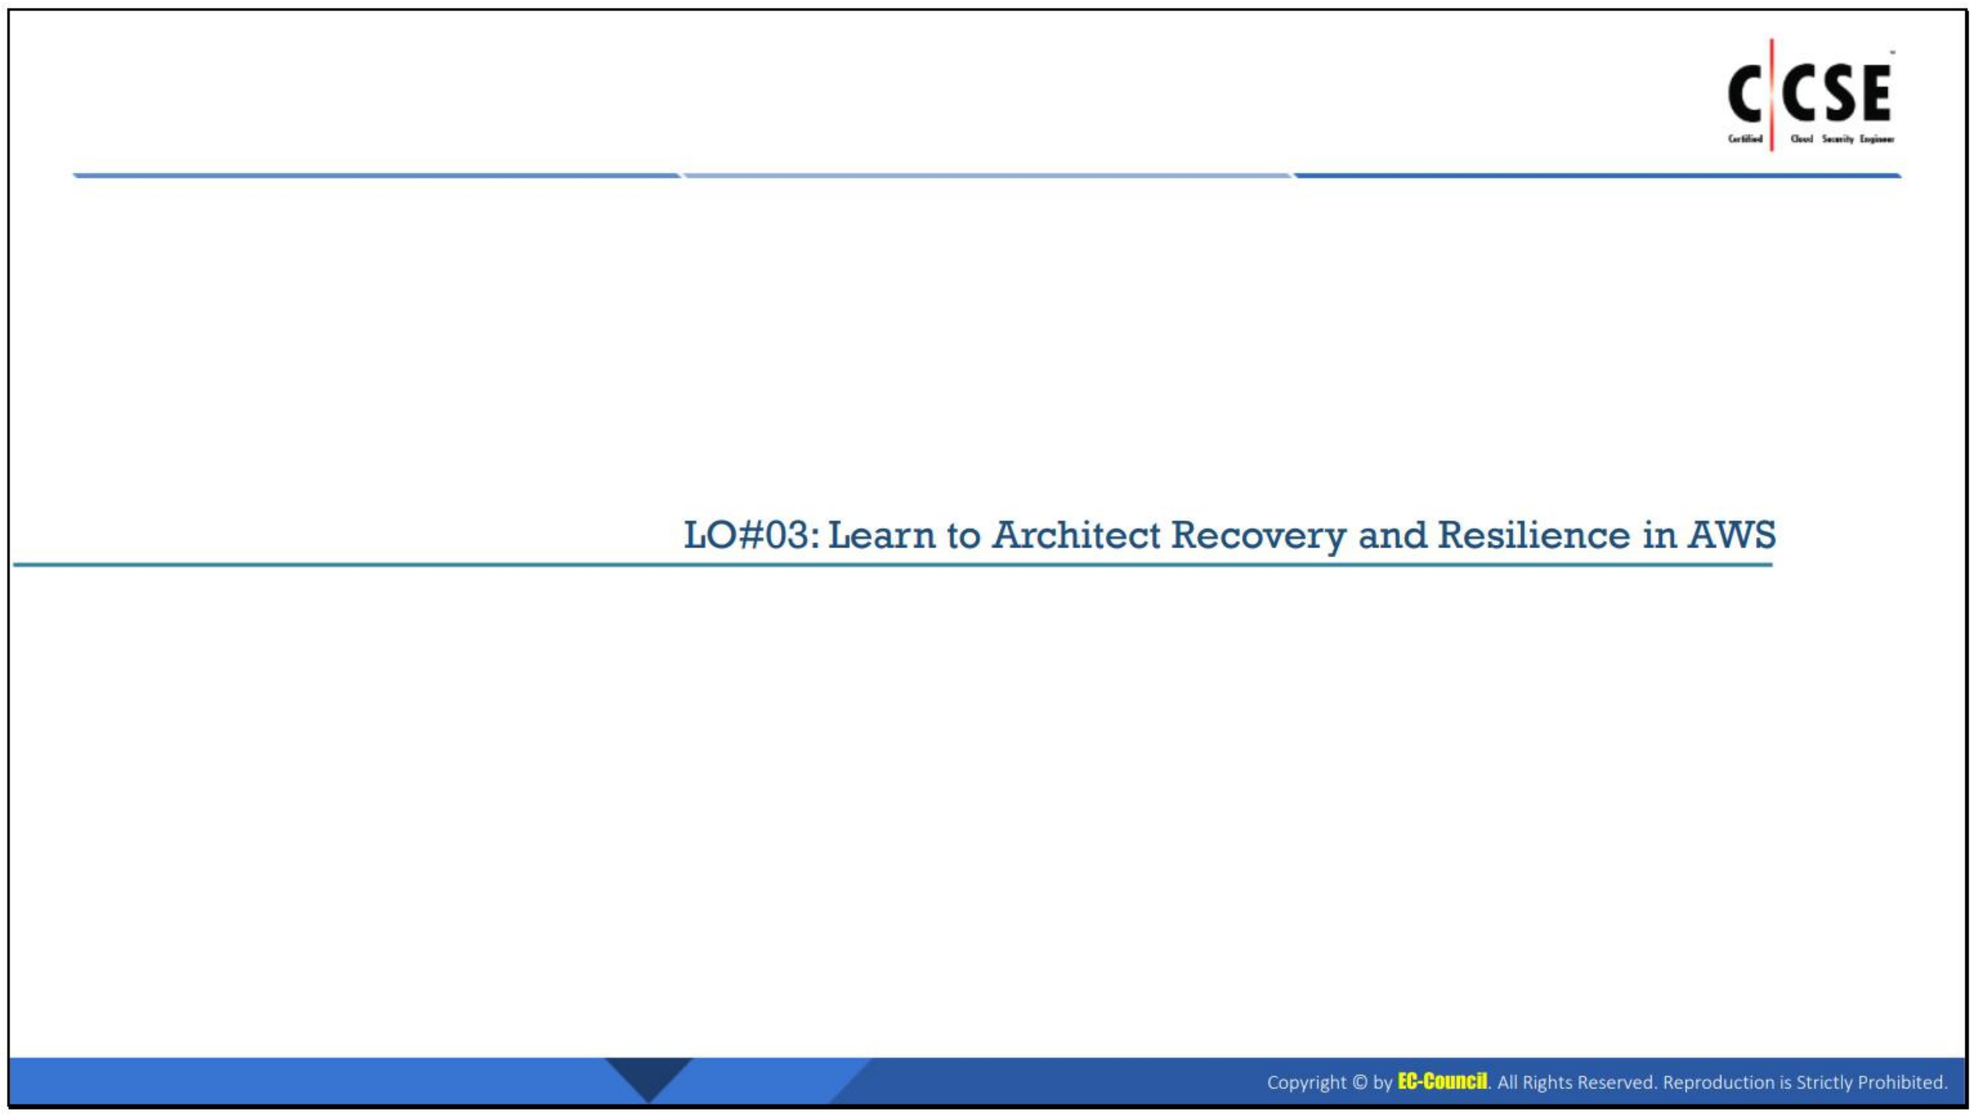

LO#3: Learn Architecting for Recovery and Resilience in AWS

The objective of this section is to understand the various DR options and architectures that are available on Amazon Web Services (AWS). Cloud security engineers can make themselves familiar with the AWS services that can be used for BCDR on AWS so that they can use them effectively to recover data and services during disasters.

hidoo1.ir

Certified Cloud Security Engineer Copyright © by **EC-Council**
All Rights Reserved. Reproduction is Strictly Prohibited.

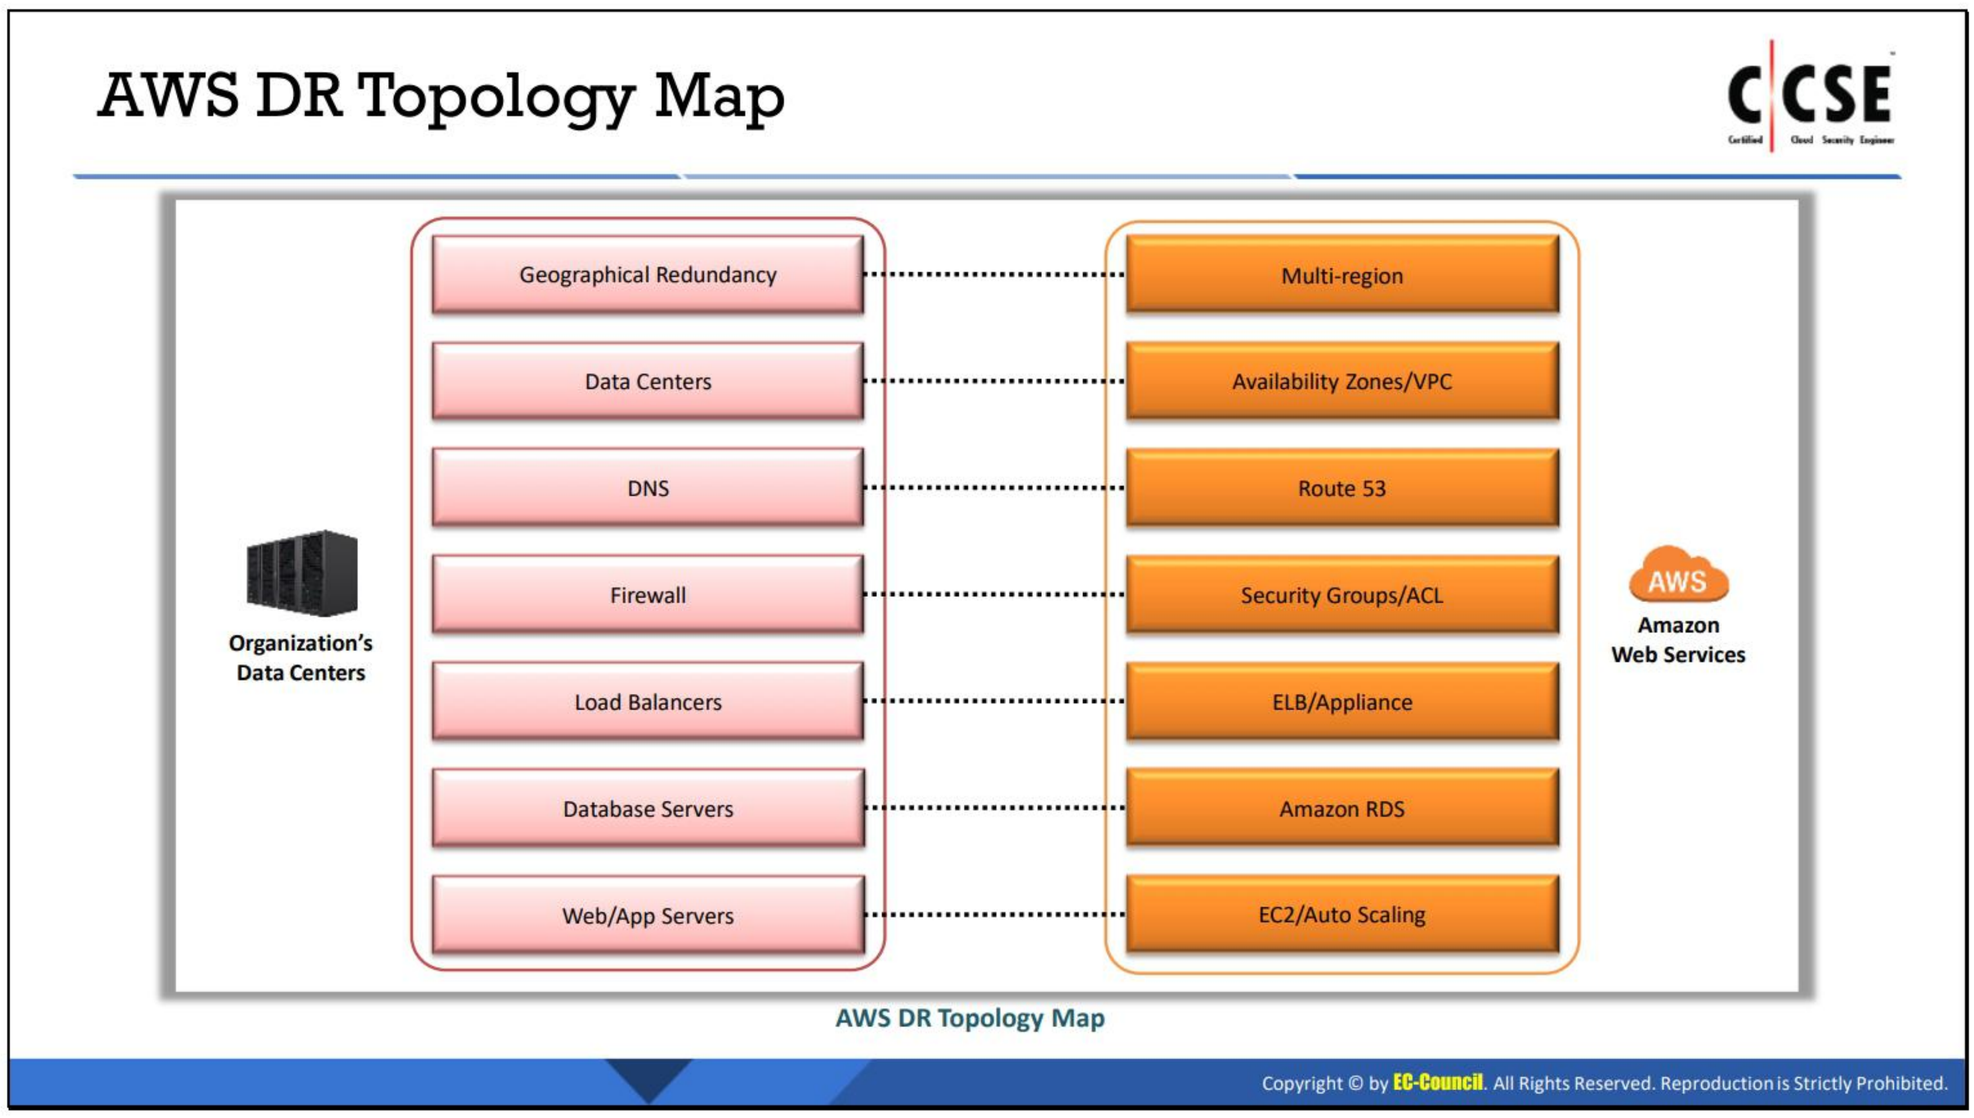

AWS DR Topology Map

Implementing a DR plan on AWS has many advantages. Customers do not have to worry about the expenses that they would bear if they used a traditional data center for disaster recovery. AWS offers disaster recovery to its customers at a nominal cost. Customers can pay a small amount of money for DR in advance for future use of DR service (CAPEX) or pay as they use (OPEX). Customers do not have to pay in advance if they use AWS's "pay as you use" (OPEX) model.

When one part of an application fails to work properly, the customer does not have to perform a failover of the complete site. It is possible just because of the high flexibility provided by AWS. AWS provides fast and easy scaling of applications during a disaster event and allows customers to completely automate their DR plans. This makes testing, documentation, and maintenance of the DR plan much easier.

Certified Cloud Security Engineer Copyright © by **EC-Council**
All Rights Reserved. Reproduction is Strictly Prohibited.

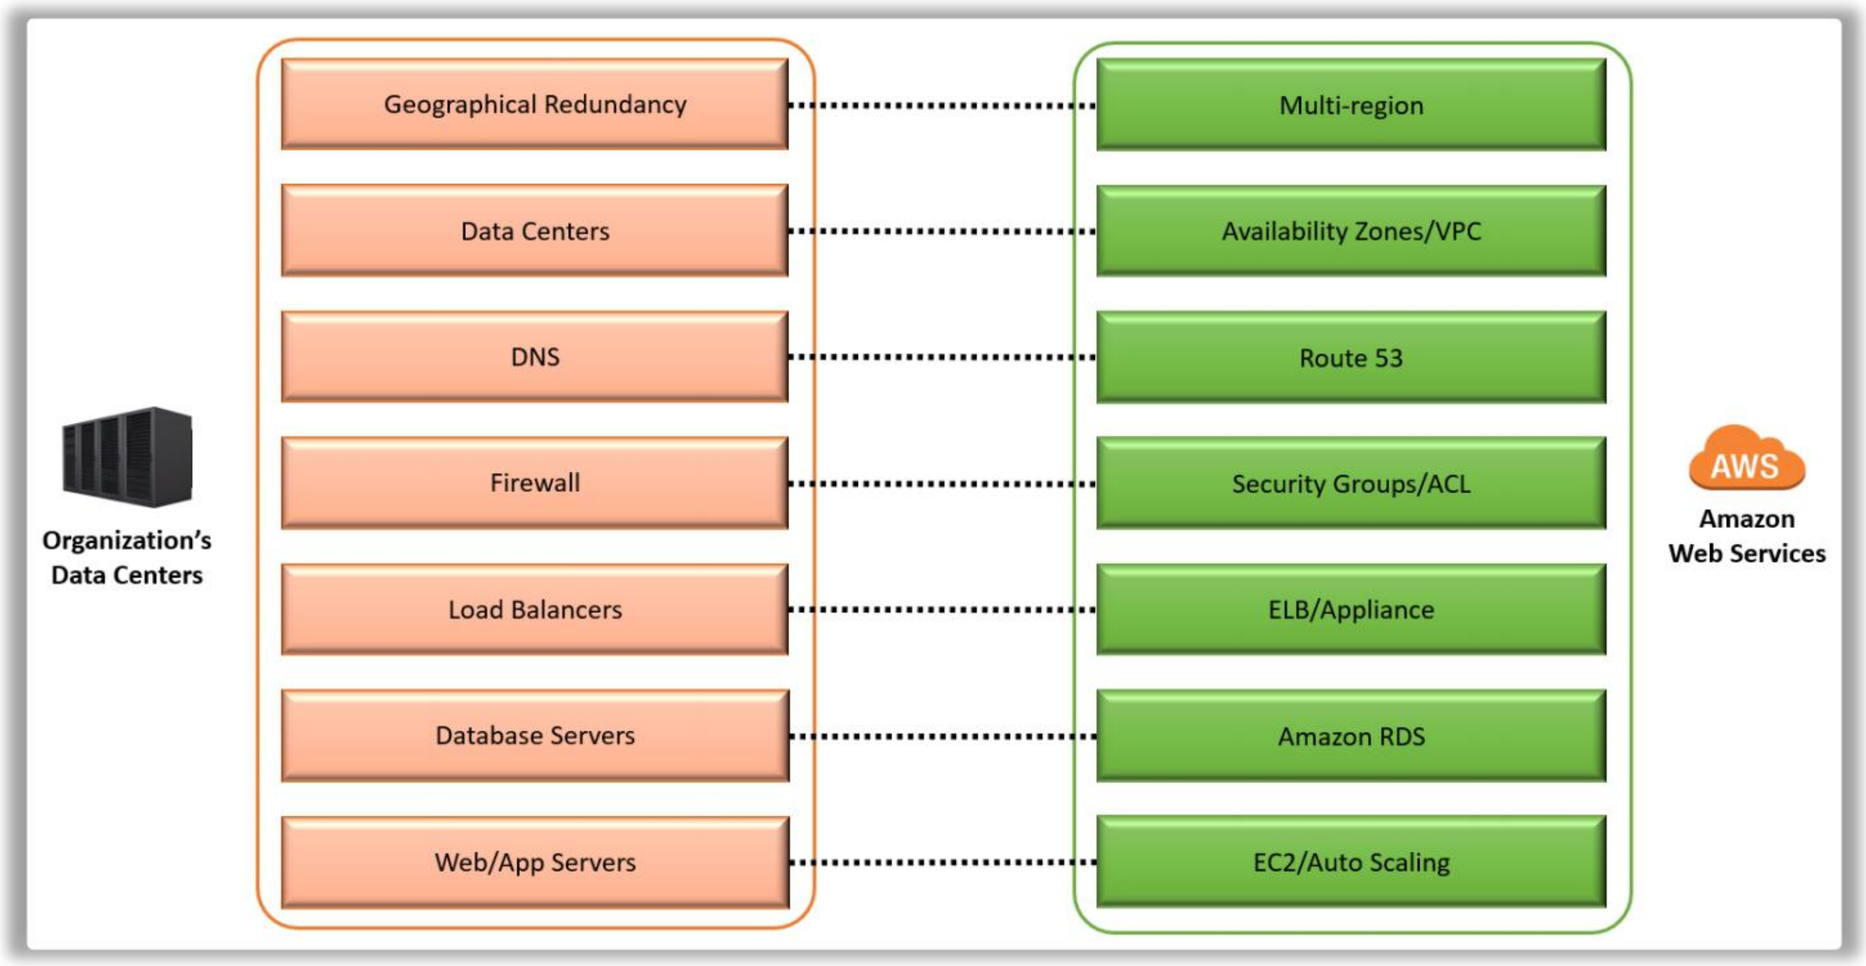

Figure 9.12: AWS DR Topology Map

The above image shows the AWS service equivalents to the infrastructure present in an on-premises data center.

hide01.ir

Certified Cloud Security Engineer Copyright © by **EC-Council**
All Rights Reserved. Reproduction is Strictly Prohibited.

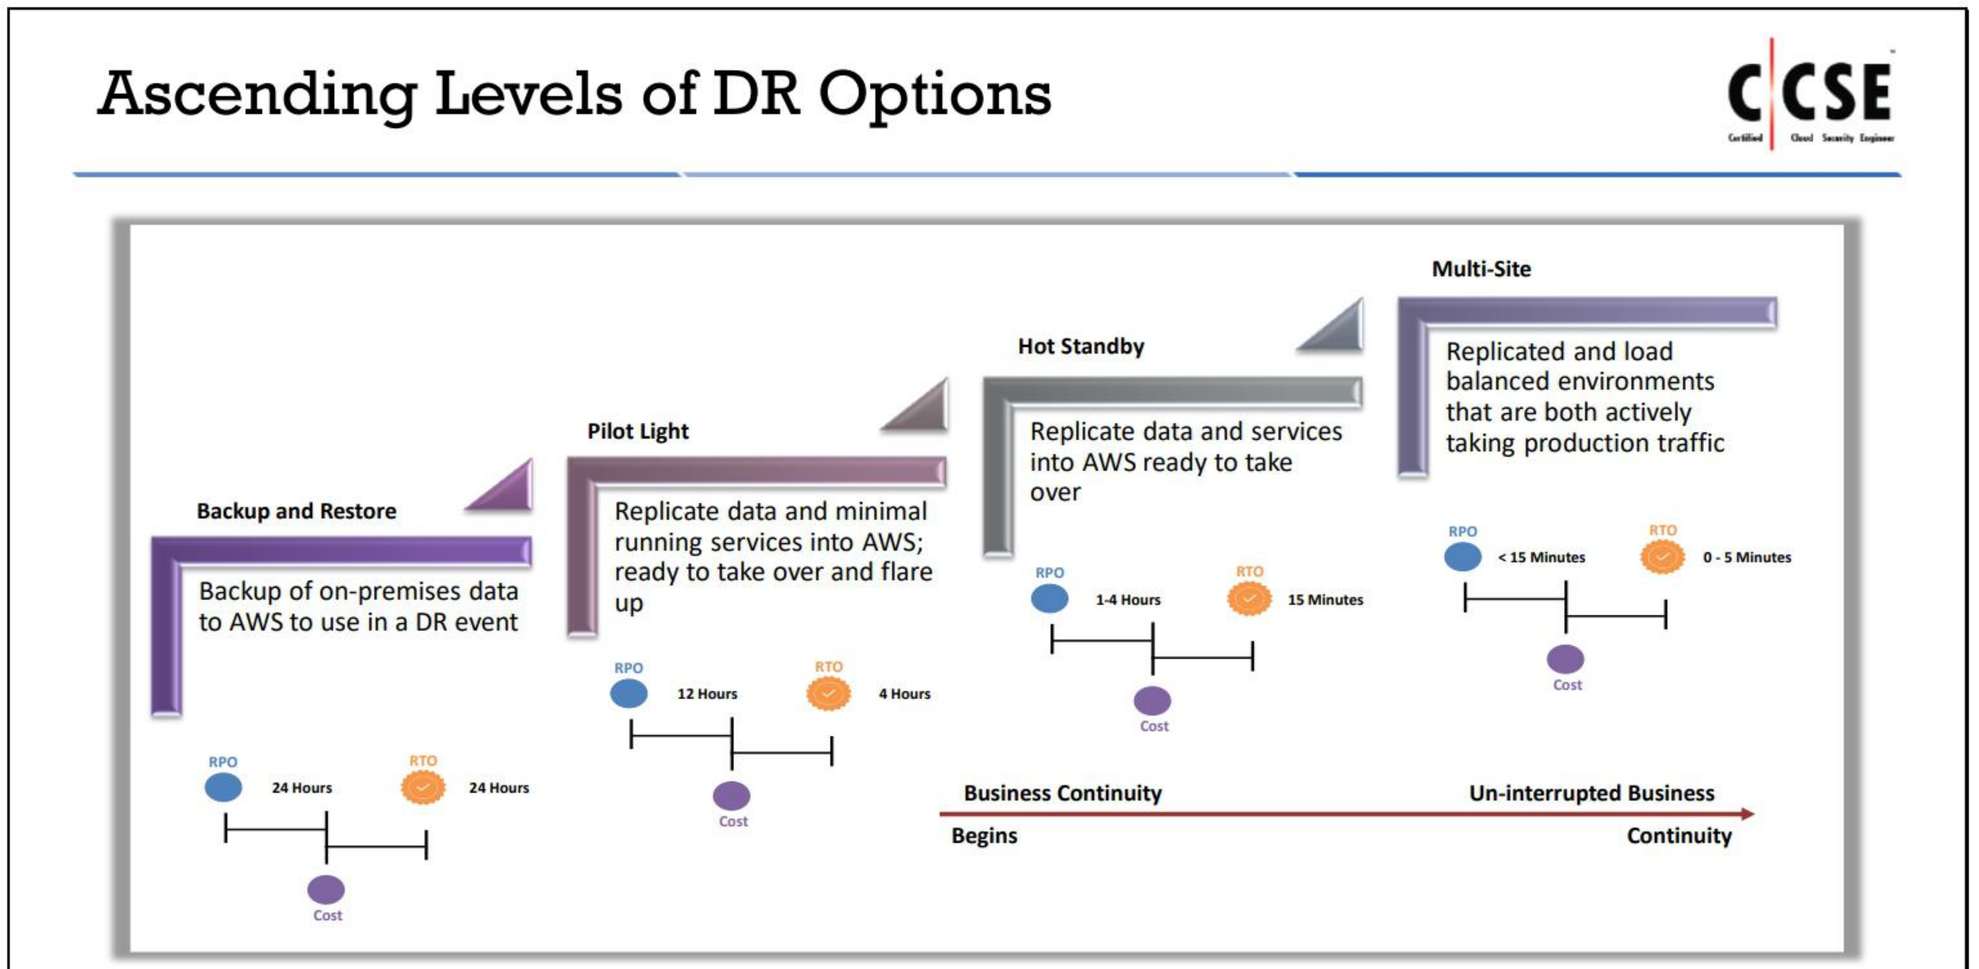

Ascending Levels of DR Options

The disaster recovery strategies offered by AWS can be classified into four types, which start from low cost and low complex strategy, and then move toward high cost and highly complex strategy. The DR plan should be tested on a regular basis to ensure that it works efficiently in the event of a disaster.

hide01.ir

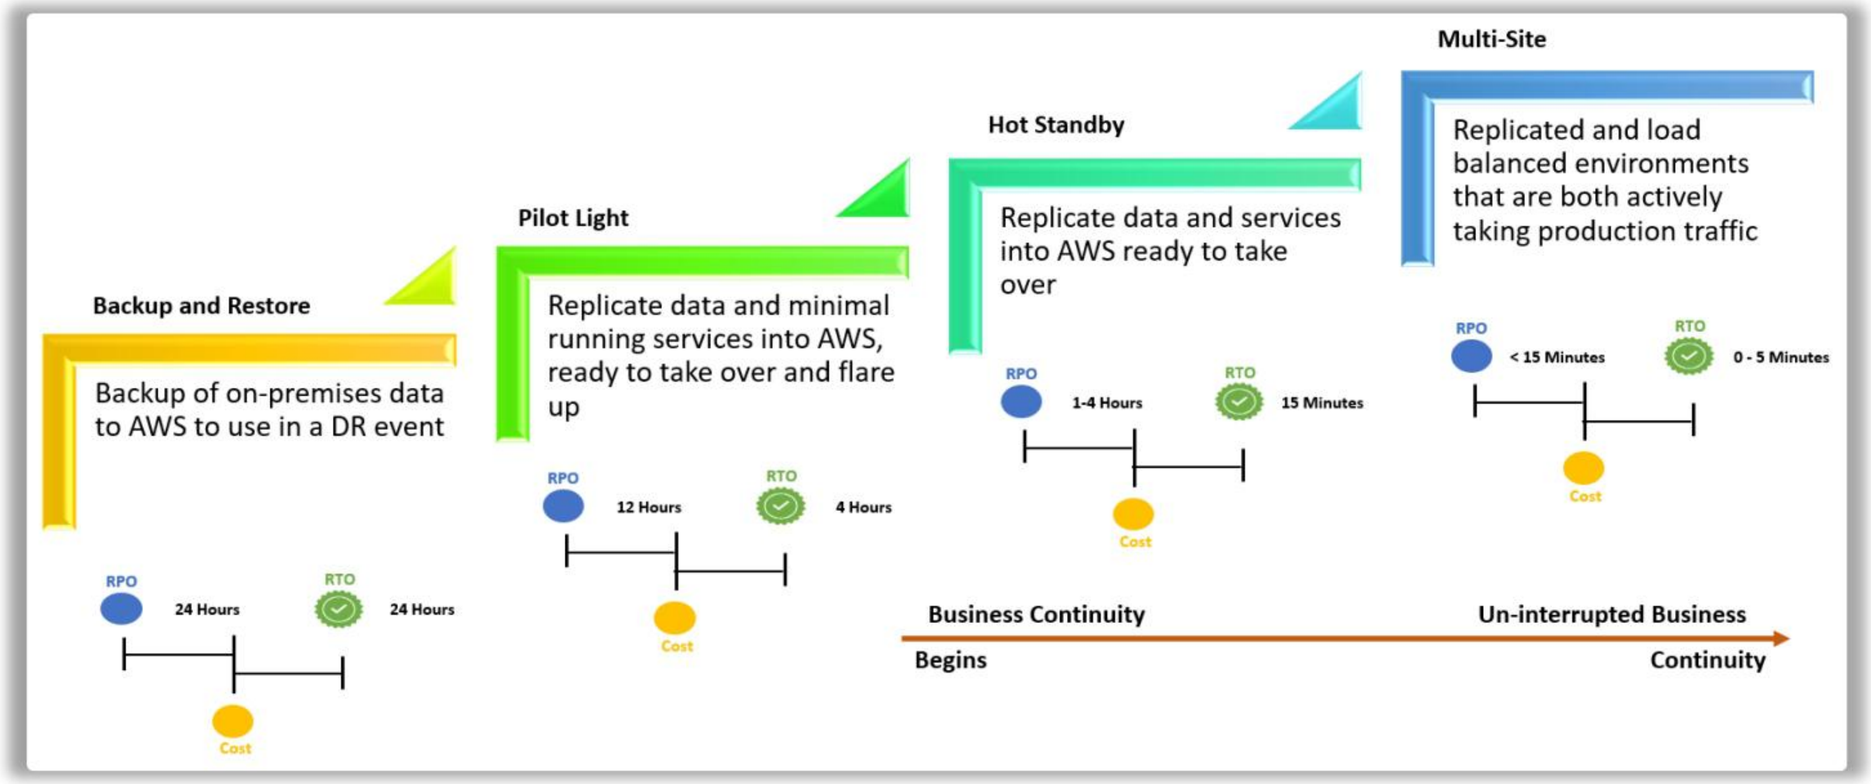

Figure 9.13: Ascending Levels of DR Plan

The ascending level of DR options are as follows:

- **Backup and restore**: The backup-and-restore disaster recovery strategy is cost-effective and is useful for lower-level business-critical processes. It will help in backing up on-premises data to AWS. The AWS customer can restore the data and can deploy resources

Certified Cloud Security Engineer Copyright © by **EC-Council**
All Rights Reserved. Reproduction is Strictly Prohibited.

after the disaster event. RTO and RPO will be high when compared to other disaster recovery strategies provided by AWS.

- **Pilot Light**: The pilot light disaster recovery strategy helps organizations configure and run the core services of their systems. It will help in replicating data and minimal running services into AWS. This DR strategy costs more than the backup-and-restore strategy and is most suitable where lower RTO and RPO are required.
- **Warm standby**: A warm standby disaster recovery strategy helps the organization run duplicate business-critical systems on the cloud; the infrastructure is up and running in the cloud. It will replicate services and data into AWS. It is a more expensive than the previous two strategies. This method is most suitable for core business-critical functions and solutions that need RTO and RPO to be within minutes.
- **Multi-site active/active**: A multi-site active/active disaster recovery strategy is the most expensive, but it offers the least RTO and RPO because it involves running a production environment on both sites simultaneously. It will replicate and load-balance the environment. This strategy is suitable for mission-critical processes and provides zero downtime.

If the disaster event disrupts or causes loss to one physical data center or important workloads, then the AWS customer requires backup-and-restore disaster recovery strategy. If the disruption caused by the disaster event is severe or it causes loss to the data center, or the AWS user needs to meet regulatory requirements, then the AWS user can opt for warm standby, pilot light, or multi-site active/active.

hide01.ir

Certified Cloud Security Engineer Copyright © by **EC-Council**
All Rights Reserved. Reproduction is Strictly Prohibited.

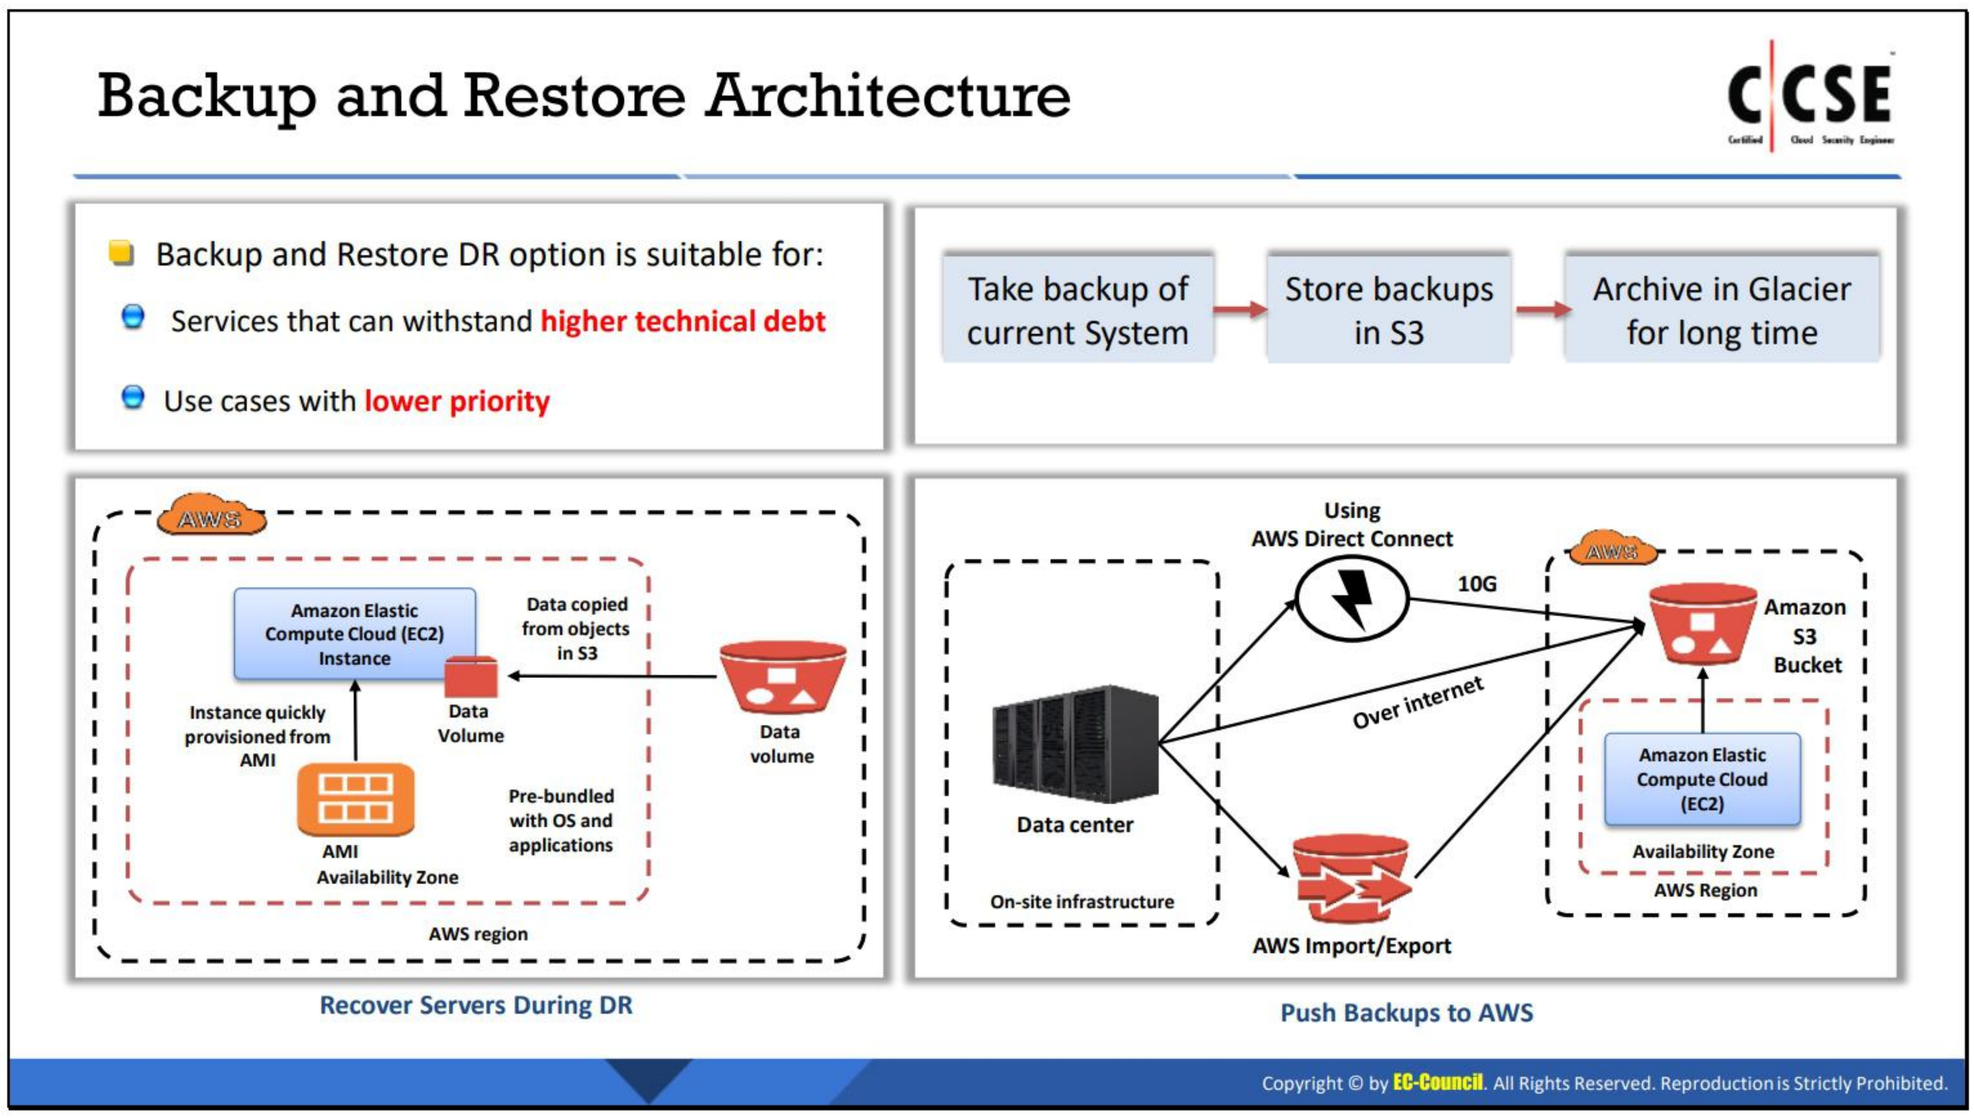

Backup-and-Restore Architecture

Backup and Restore

Traditional DR methods include backing up data on tapes and storing them off-site. Restoring of the system quickly in case of a disaster becomes difficult in this way. AWS provides services that can help organizations perform backup-and-restore operations very easily and quickly. Proper planning should be done to use appropriate service perform backup-and-restore operations as part of the DRP.

Below are some AWS services that help organizations in faster and easier disaster recovery.

- **Amazon S3** is best suited for backing up data that requires quick retrieval. Backup data stored in Amazon S3 can be quickly restored when any disaster strikes. This service is accessible from any location, and a network is required to transfer data from and to Amazon S3. Many open-source and commercial backup solutions support integration with Amazon S3 these days.
- **AWS Import/Export** service is a data transfer service that can be utilized to transfer large volumes of data to and from AWS by connecting storage devices.
- **Amazon Glacier** can be used to backup data that does not need quick retrieval; therefore, it can be used where data is to be stored for long-term and a retrieval time of many hours is adequate. It offers the same durability as Amazon S3, thus, both can be used in conjunction to create a tiered backup solution.
- **Amazon Storage Gateway** can be used for creating backups by copying snapshots of on-premises data volumes transparently to Amazon S3. These snapshots can be used to create Amazon Elastic Block Store (EBS) volumes or local volumes. By using gateway-cached volumes service, primary data can be stored in Amazon S3 and frequently

Certified Cloud Security Engineer Copyright © by **EC-Council**
All Rights Reserved. Reproduction is Strictly Prohibited.

accessed data can be stored locally to provide low-latency access. When any disaster happens, the cache volumes can be restored either to Amazon EC2 instances or a second site running a storage cache gateway. The gateway virtual tape library (VTL) configuration of Storage Gateway can be used as a replacement for traditional tape backup and as a backup target for existing backup management software.

The application is preconfigured on an EC2 instance, and it is converted to an Amazon Machine Instance (AMI) stored in Amazon S3 to enable quick recovery when a disaster hits.

The below image shows various data backup options to Amazon S3, either from AWS or on-premises infrastructure.

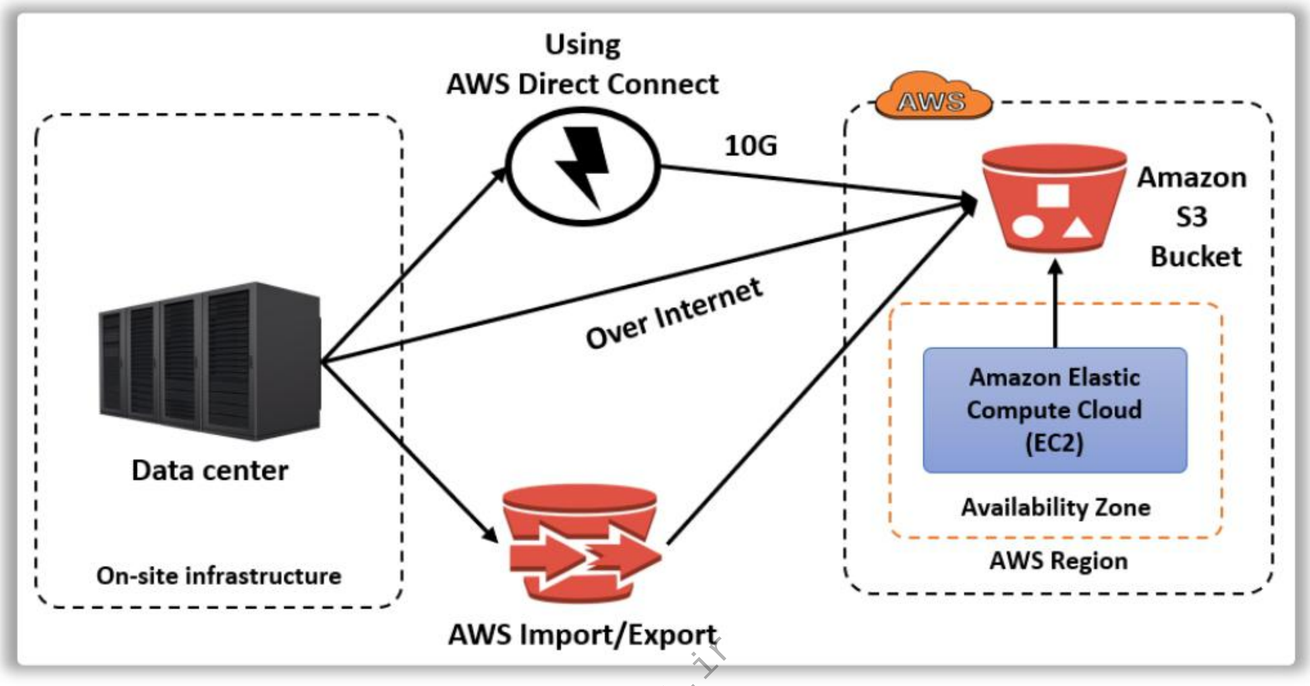

Figure 9.14: Data Backup from On-premises Infrastructure

The next step after the data backup is to restore the data quickly and reliably after a disaster; thus, systems should be configured properly to securely retain the data. It is recommended to test the recovery processes before any disaster happens to avoid any data loss during and after a disaster. When a disaster strikes, the data is recovered from Amazon S3 and a preconfigured AMI is utilized to recreate a production environment on AWS.

The below image shows how Amazon S3 backups can be used to restore a system to Amazon EC2 as part of the recovery process.

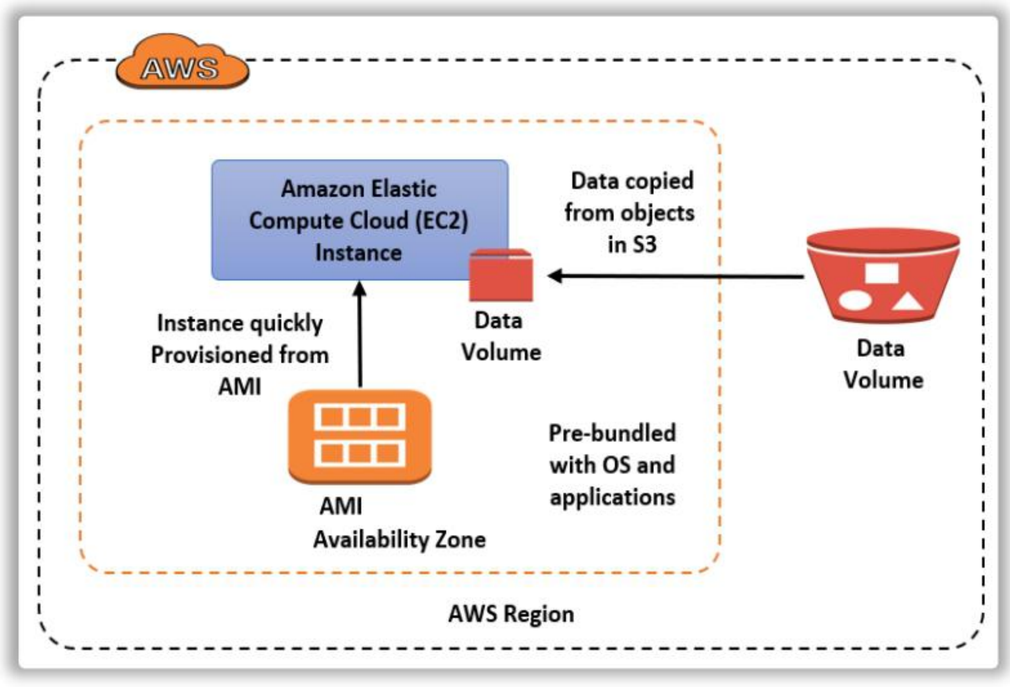

Figure 9.15: Restoring a System to Amazon EC2 from S3 Backups

Certified Cloud Security Engineer Copyright © by **EC-Council**
All Rights Reserved. Reproduction is Strictly Prohibited.

Following are the key steps for backup and restore:

- Select the desired method or service to backup your data to AWS.
- Make sure there is an appropriate data retention policy for this data.
- Proper security measures, such as access control policies and encryption, must be implemented for the data.
- The data recovery and restoration process of the system should be tested on a regular basis.

hide01.ir

Certified Cloud Security Engineer Copyright © by **EC-Council**
All Rights Reserved. Reproduction is Strictly Prohibited.

Pilot Light Pattern

Build resources around replicated dataset → Scale resources in AWS in response to a DR event

- Replicate core databases by keeping "**pilot light**" on
- Build **AWS resources** around the dataset and leave them in the stopped state

- Start up pool of resources in AWS when **events dictate**
- Match current production capacity through **auto-scaling policies**

- Pilot light DR option is suitable for
  - **Business-critical services**
  - Meeting **lower RTO and RPO** requirements

Copyright © by EC-Council. All Rights Reserved. Reproduction is Strictly Prohibited.

Pilot Light Architecture (Cont'd)

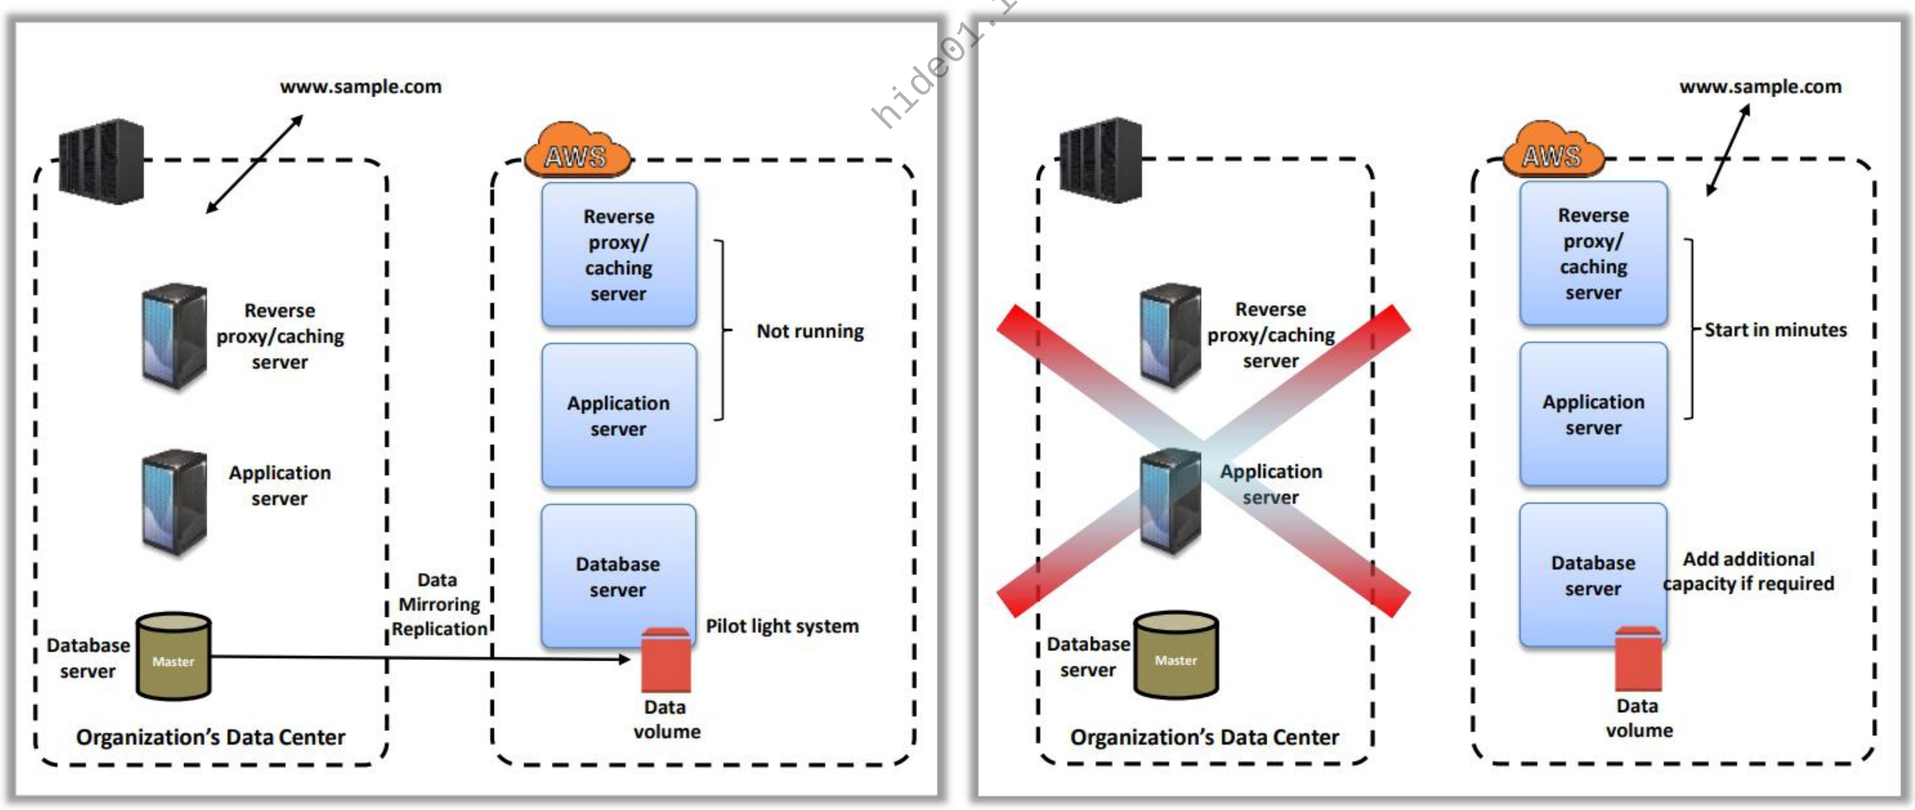

hide01.ir

Copyright © by EC-Council. All Rights Reserved. Reproduction is Strictly Prohibited.

Pilot Light Pattern

Pilot light disaster recovery strategy is used for quick recovery from any disaster, and it works faster than the backup-and-restore method. This strategy works by running a minimal version of an environment on the cloud so that it can be restored quickly when a disaster strikes. AWS helps organizations configure and run the most critical core elements of their systems (known as the pilot light) in AWS. At the time of recovery, they provide a full-scale production environment around critical core elements instantly.

Certified Cloud Security Engineer Copyright © by **EC-Council**
All Rights Reserved. Reproduction is Strictly Prohibited.

The critical core elements include the elements around which infrastructure components can be provisioned rapidly to restore the entire system in AWS. The core infrastructure elements include database servers (utilized for duplicating data to AWS services such as EC2 and Relational Database Service (RDS)) and other important data outside the database that needs to be replicated based on the system to AWS.

Certain preconfigured servers bundled as AMIs are used to provision the remaining infrastructure that is needed to restore business-critical services. In the event of disaster recovery, the instances from these AMIs start working with their predefined roles, such as web or app servers, within their deployment around the critical infrastructure.

There are two main AWS networking services that can be used as options for provisioning.

- **Elastic IP addresses**: Use elastic IP addresses that can be pre-allocated and known or identified in the DR preparation phase and associate them with your instances. Elastic network interfaces (ENI) can be utilized for MAC address-based software licensing. It has a pre-allocated MAC address that can be used to allot license. These can be associated with the instances like elastic IP addresses. Like elastic IP addresses, ENI can also be associated with the instances.
- **Elastic Load Balancing (ELB)**: Utilize ELB to distribute traffic to multiple instances and then update the DNS records to the load balancer by using a CNAME or point at Amazon EC2 instances. This option is recommended for traditional web-based applications.

For less critical systems, AWS customers should make sure that the information related to configuration and the installation packages is available in AWS as EBS snapshots. It will quicken the setup process for the application server because multiple volumes can be created in multiple AZs to attach to Amazon EC2 instances. Then, it will be easy to install and configure by using the backup-and-restore method.

AWS now allows automation of the provisioning and configuration of infrastructure resources, which is very helpful in saving time and preventing human errors. However, certain installation and configuration tasks need to be performed to fully recover the applications.

Preparation Phase

In the preparation phase of the disaster recovery process using the pilot light method, the data that continuously keeps changing needs to be replicated into the pilot light system. The data that is updated less frequently, such as the OS and applications, can be stored and updated as AMIs.

Key steps needed for the preparation:

- Setup Amazon EC2 instances in AWS so that the data can be replicated.
- Make sure all supporting custom software packages are present in AWS.
- Where fast recovery is needed, create and maintain AMIs of key servers.
- Test these servers by running them regularly and apply any configuration changes and software updates.
- Automate the provisioning of AWS resources.

Certified Cloud Security Engineer Copyright © by **EC-Council**
All Rights Reserved. Reproduction is Strictly Prohibited.

The preparation phase of the pilot light scenario is shown in the image given below.

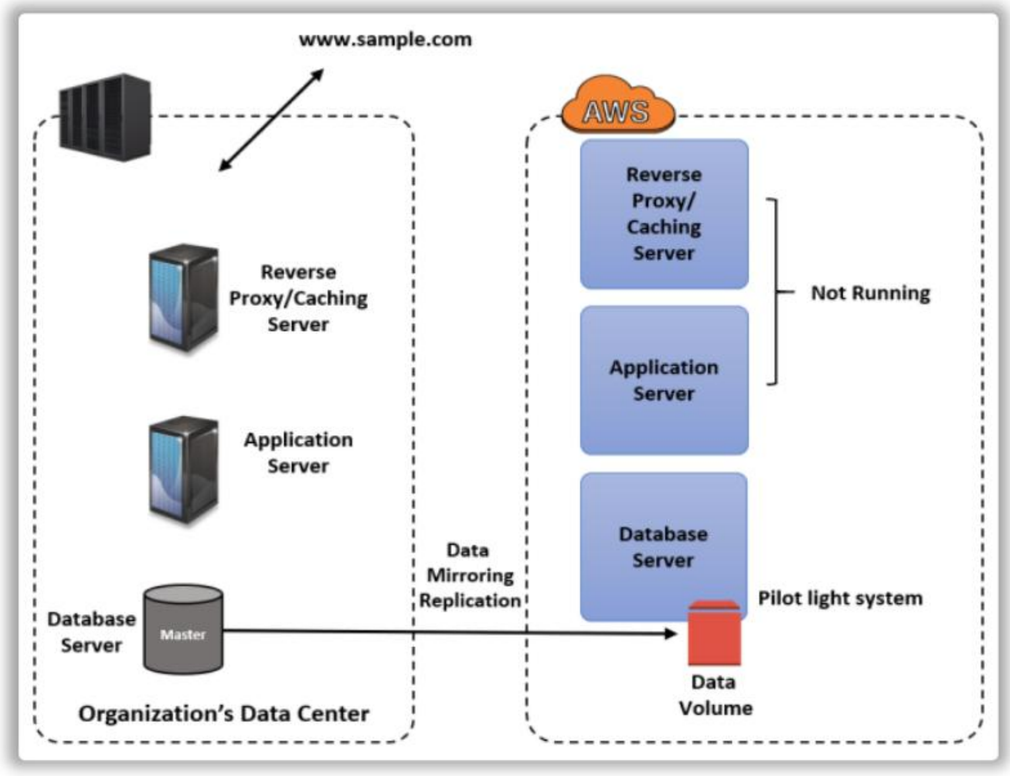

Figure 9.16: Recovery Phase of Pilot Light Scenario

This will enable the minimum version of the environment to run in AWS with the core infrastructure (pilot light).

Recovery Phase

The rest of the environment around the core infrastructure can be recovered by starting the system using AMI on appropriate instance types. Resize dynamic servers as needed or increase their capacity to handle production volumes. The scalable and cost-effective approach to add capacity to a system is horizontal scaling. For example, more web servers can be added to the system at peak times. But scaling vertically is also possible by opting for larger EC2 instance types. From a networking point of view, required DNS updates can be performed in parallel.

Ensure that redundancy is restored soon after recovery. Maintain additional redundancy at the data layer and take regular backups of your systems.

The recovery phase of the pilot light scenario is shown in the below image.

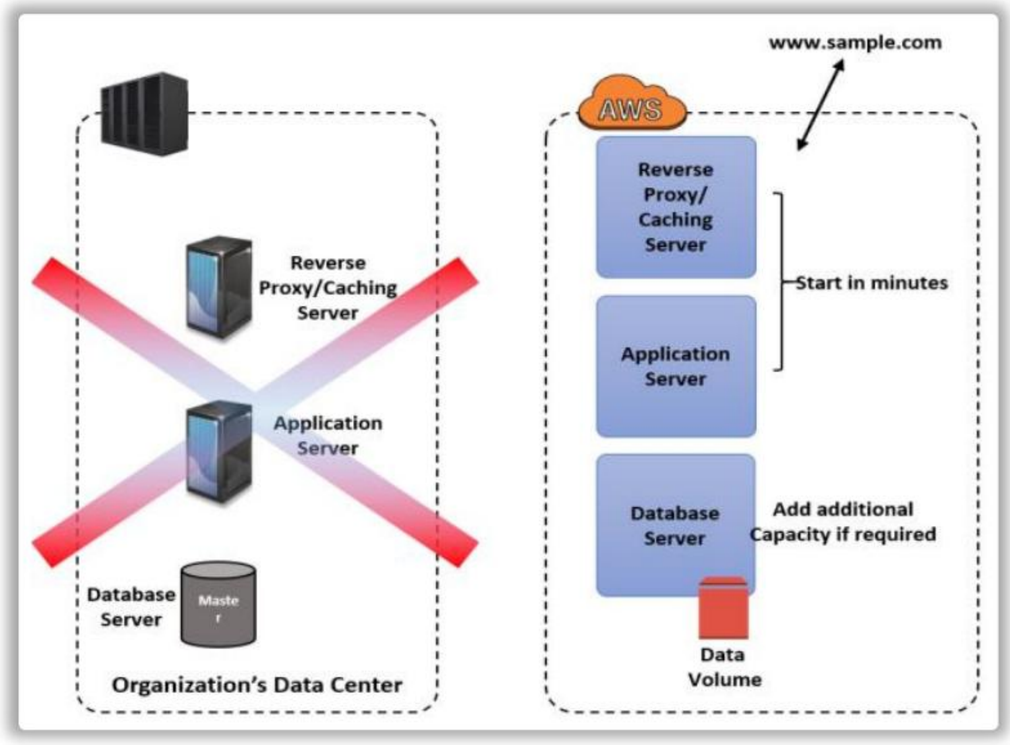

Figure 9.17: Recovery Phase of Pilot Light Scenario

Certified Cloud Security Engineer Copyright © by **EC-Council**
All Rights Reserved. Reproduction is Strictly Prohibited.

The key steps needed for recovery are listed below:

- From your custom AMIs, start your application Amazon EC2 instances.
- Resize your data store instances/existing database to make them ready to process increased traffic.
- Change DNS in a way that it points to Amazon EC2 servers.
- Utilize an automated way for installing and configuring non-AMI based systems.

hide01.ir

Certified Cloud Security Engineer Copyright © by **EC-Council**
All Rights Reserved. Reproduction is Strictly Prohibited.

Warm Standby Architecture

Build resources around replicated environment → Scale resources in AWS in response to a DR event

- Operate a warm standby by replicating **app server** and core databases
- Build AWS resources around dataset and **run** in a **limited capacity**
- The **scale-up** pool of resources in AWS when events dictate
- Match current production capacity through **auto-scaling policies**

- Warm standby DR option is suitable for:
  - Solutions that require **RTO and RPO in minutes**
  - **Core business-critical functions**

Copyright © by EC-Council. All Rights Reserved. Reproduction is Strictly Prohibited.

Warm Standby Architecture (Cont'd)

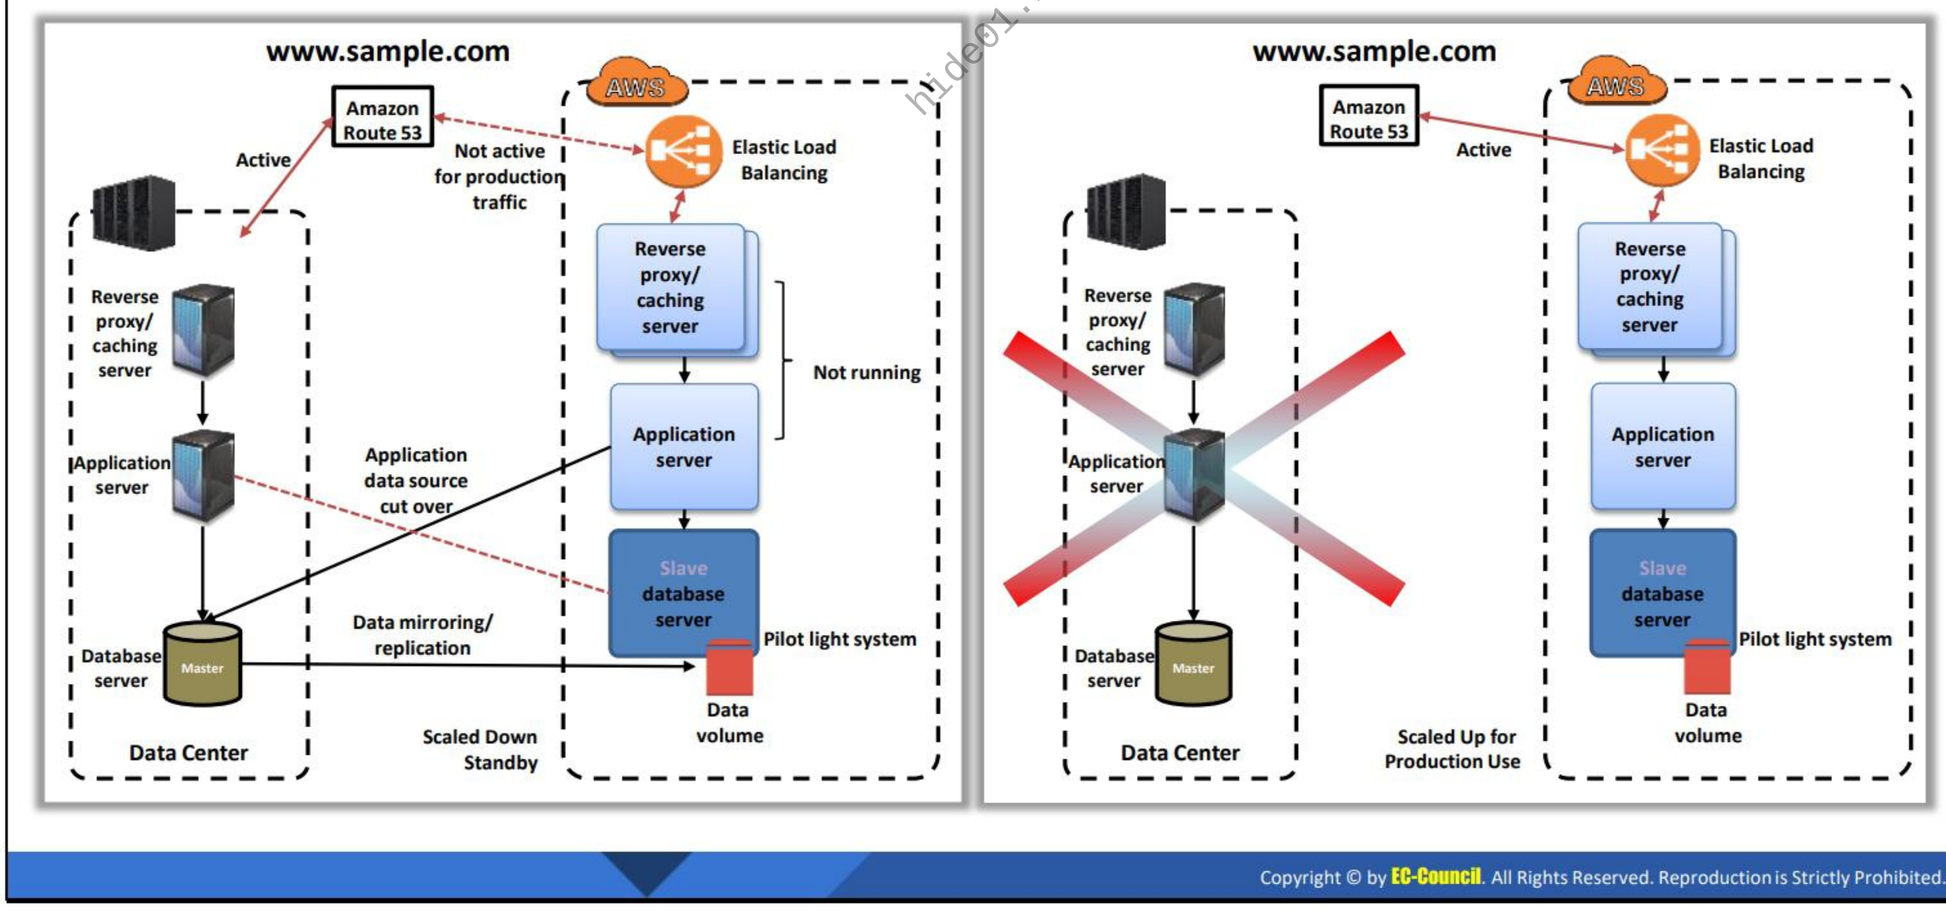

hide01.ir

Copyright © by EC-Council. All Rights Reserved. Reproduction is Strictly Prohibited.

Warm Standby Architecture

In a warm standby strategy, a scaled-down version of the production environment runs in the cloud. It reduces the recovery time in the event of disasters because there are always some services running in the cloud. In this scenario, it is possible to always run fully duplicate systems of your business-critical systems on the cloud after identifying them. The duplicate systems run on a small group of EC2 instances. This solution is fully functional but is not scaled to take full

Certified Cloud Security Engineer Copyright © by **EC-Council**
All Rights Reserved. Reproduction is Strictly Prohibited.

production load. The warm standby scenario is suitable for testing, quality assurance, and internal use.

At the time of a disruptive event, the standby system will be instantly scaled up to manage the incoming production load by adding new instances to the load balancer and by running servers that have small size on the larger EC2 instance. In this situation, horizontal scaling is preferable over vertical scaling for adding capacity.

Preparation Phase

In the preparation phase of warm standby, an AWS solution and an on-site solution are run simultaneously, as shown below.

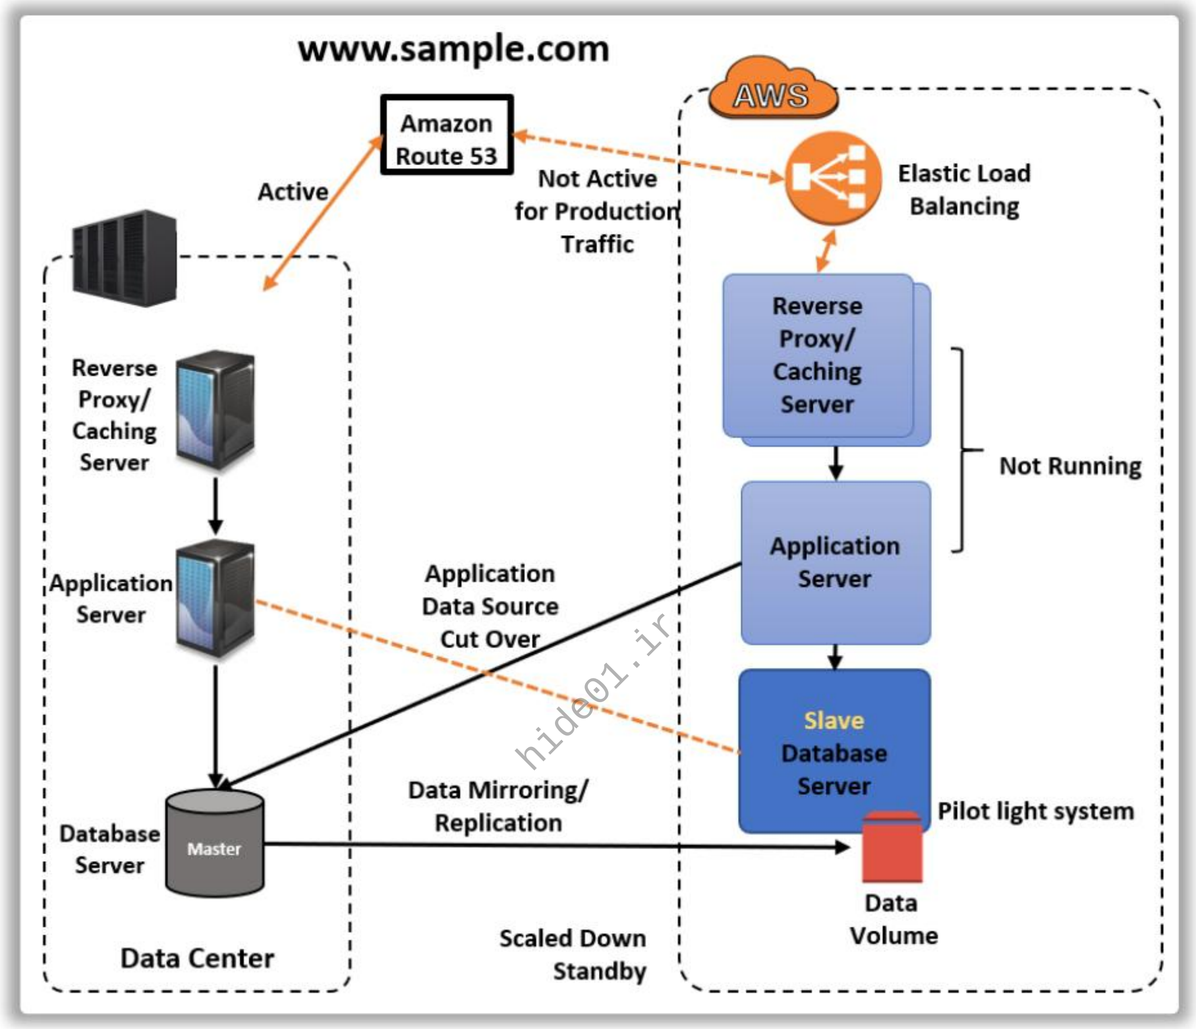

Figure 9.18: Preparation Phase of Warm Standby

The key steps for preparation are listed below:

- Set up Amazon EC2 instances in AWS so that the data can be replicated.
- Create AMIs and maintain them.
- Use a minimum footprint of AWS infrastructure or Amazon EC2 instances to run your application.
- Software should be patched and updated and the configuration files should also be updated according to the live environment.

Recovery Phase

When the production system fails, the standby system will be scaled up to handle production load and DNS records will be modified to route all traffic to AWS.

The key steps for recovery are as follows:

Certified Cloud Security Engineer Copyright © by **EC-Council**
All Rights Reserved. Reproduction is Strictly Prohibited.

- Use the load balancer to increase the size of Amazon EC2 fleets in service (horizontal scaling).
- Depending on the requirement, run the applications on larger EC2 (vertical scaling).
- To route all traffic to AWS, change the DNS records manually or use Amazon Route 53 automated health checks.
- To right-size the Amazon EC2 fleet or accommodate the increased load, consider using Auto Scaling.

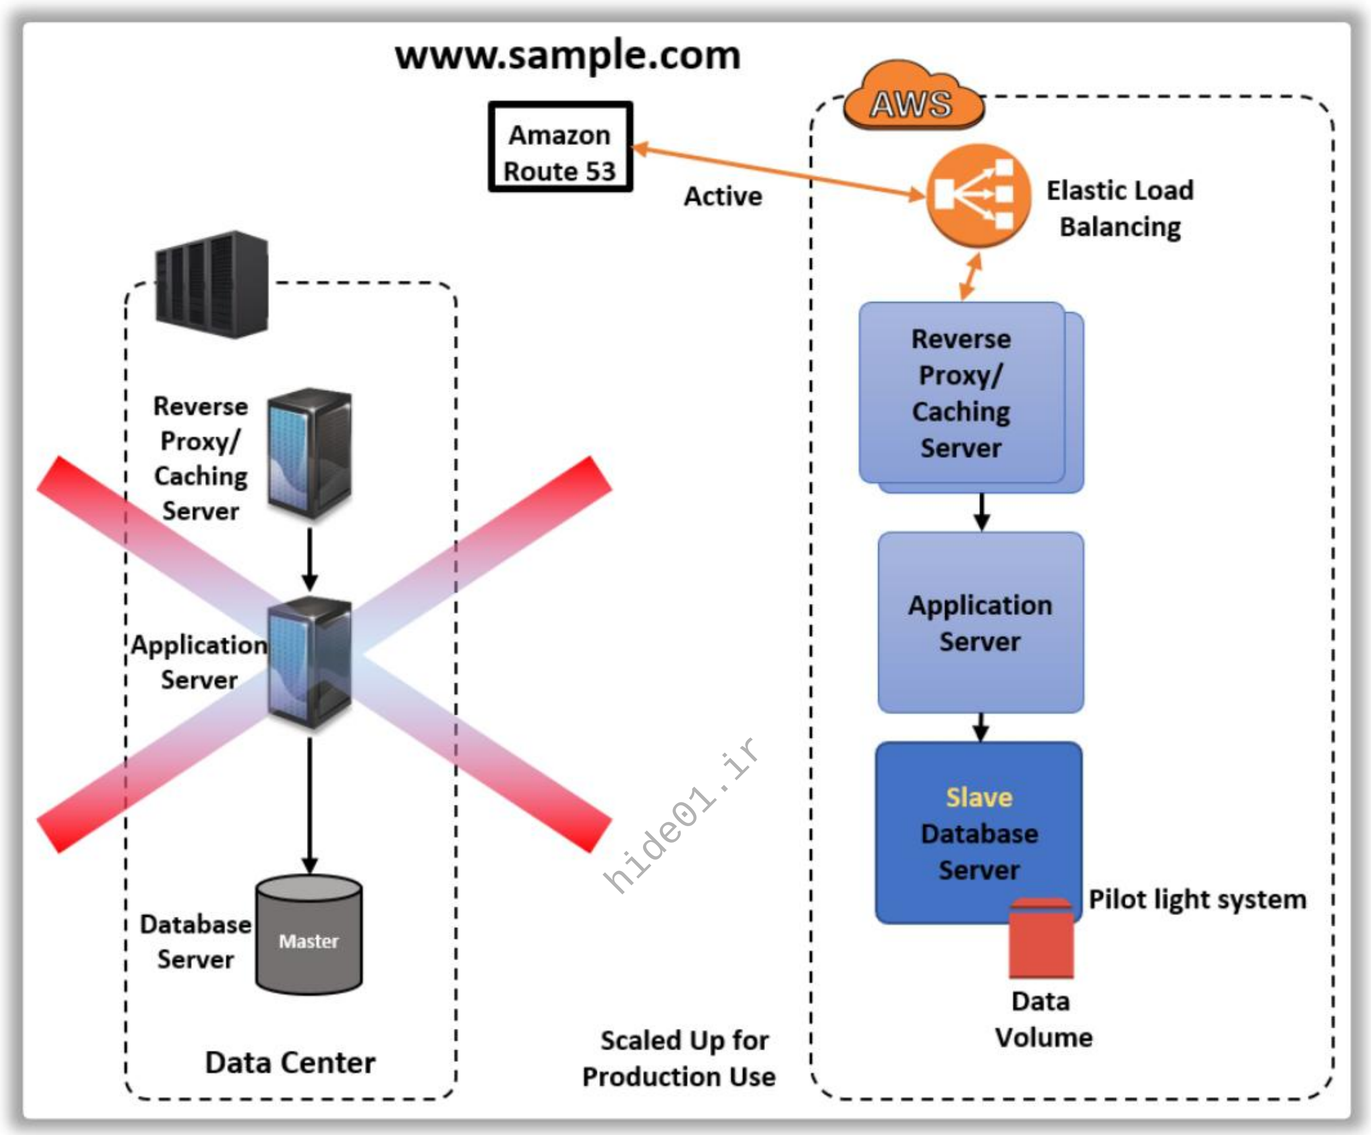

Figure 9.19: Recovery Phase of Warm Standby

Certified Cloud Security Engineer Copyright © by **EC-Council**
All Rights Reserved. Reproduction is Strictly Prohibited.

Hot Multi-Site Architecture

Deploy resources necessary to operate full production → Fail over to AWS in response to a DR event

- Operate a full stack by replicating **app servers** and **core databases**
- Sufficient resources in AWS to handle **full peak load**

- Hot DR option is suitable for:
  - Solutions that require **zero RTOs and RPOs**
  - **Core business-critical functions**
  - Multi-site architecture with automatic load balancing of traffic to **multiple availability zones (AZs)** and even AWS regions

Copyright © by EC-Council. All Rights Reserved. Reproduction is Strictly Prohibited.

Hot Multi-Site Architecture (Cont'd)

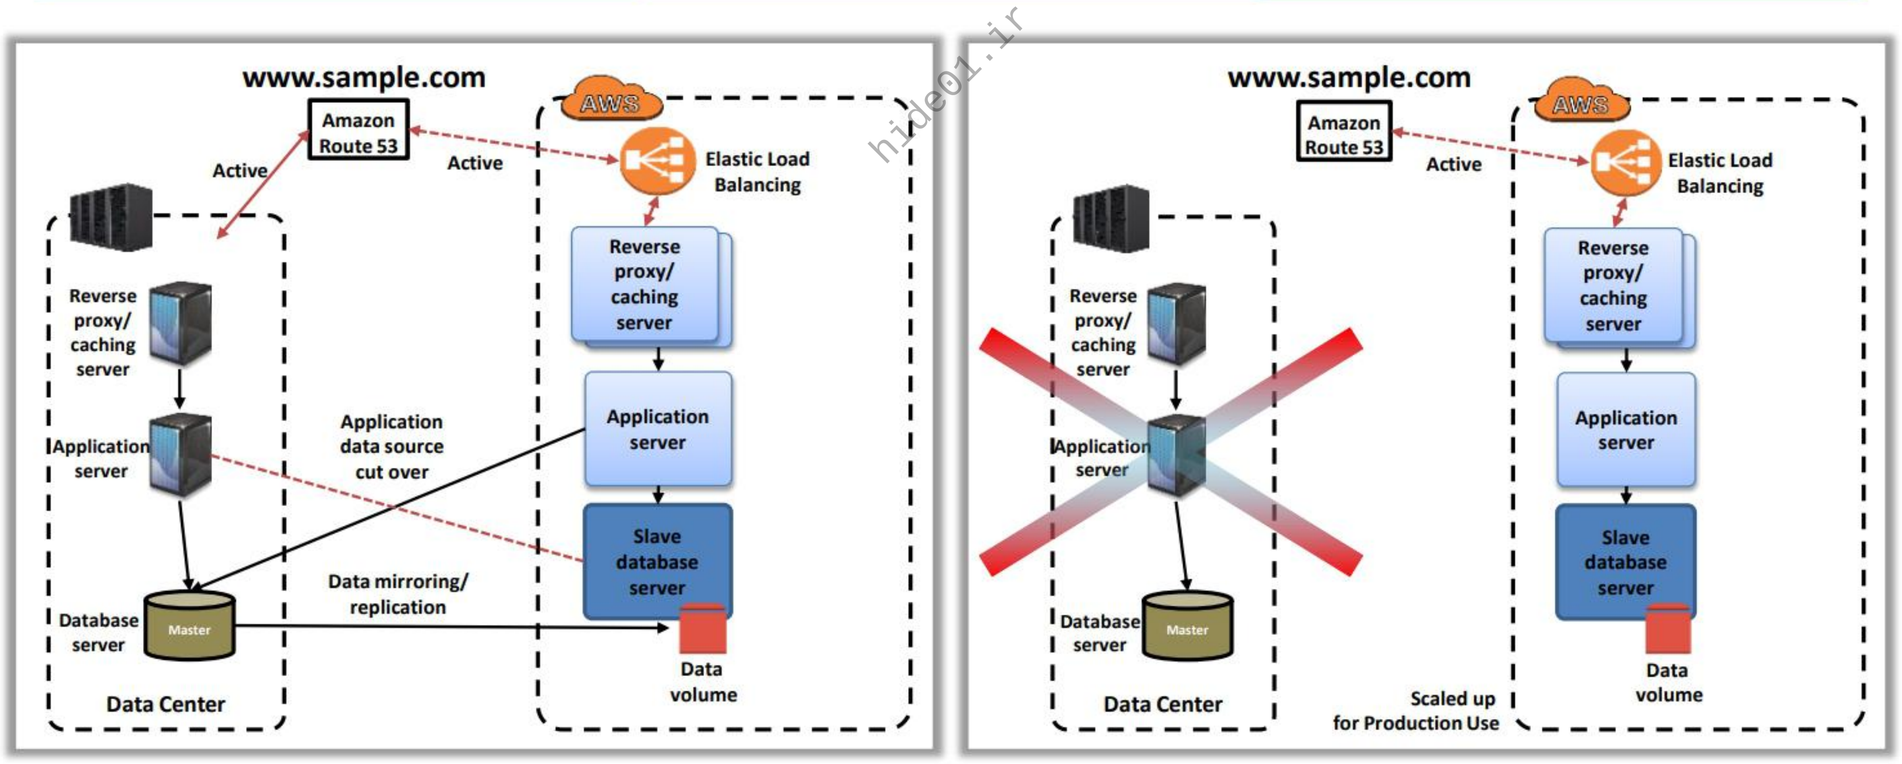

hide01.ir

Copyright © by EC-Council. All Rights Reserved. Reproduction is Strictly Prohibited.

Hot Multi-Site Architecture

The fourth disaster recovery scenario is the multi-site solution, which is deployed on both AWS and on-premises. In this scenario, a multi-site solution works on the on-site infrastructure and on AWS in an active-active configuration (both are active and running). The implementation of this scenario is most expensive; however, it offers the least RTO and RPO because it involves running of a production environment on both sites (on-site infrastructure and AWS) simultaneously.

Certified Cloud Security Engineer Copyright © by **EC-Council**
All Rights Reserved. Reproduction is Strictly Prohibited.

In this scenario, a DNS service such as Amazon Route 53 that supports weighted routing is used to route a part of the production traffic to the on-site infrastructure and the remaining part of the traffic to the AWS infrastructure. Data is replicated to AWS through replication methods such as synchronous and asynchronous methods.

In a disaster situation when the on-site production infrastructure fails to handle the production load, the DNS weighting can be adjusted or updated to redirect all the production traffic to the AWS infrastructure and the capacity of the AWS servers can be increased or scaled up rapidly to handle full production load. This process can be automated by using **EC2 Auto Scaling**. Some application logic is needed to detect the primary database failure and upgrade the slave database on AWS to the master database.

Preparation Phase

In the preparation phase, a DNS service (Amazon Route 53) is used to route some part of production traffic to AWS. The data sources in the on-site production system may be accessed by the application on AWS. Data is replicated to the AWS infrastructure. The below image shows the preparation phase.

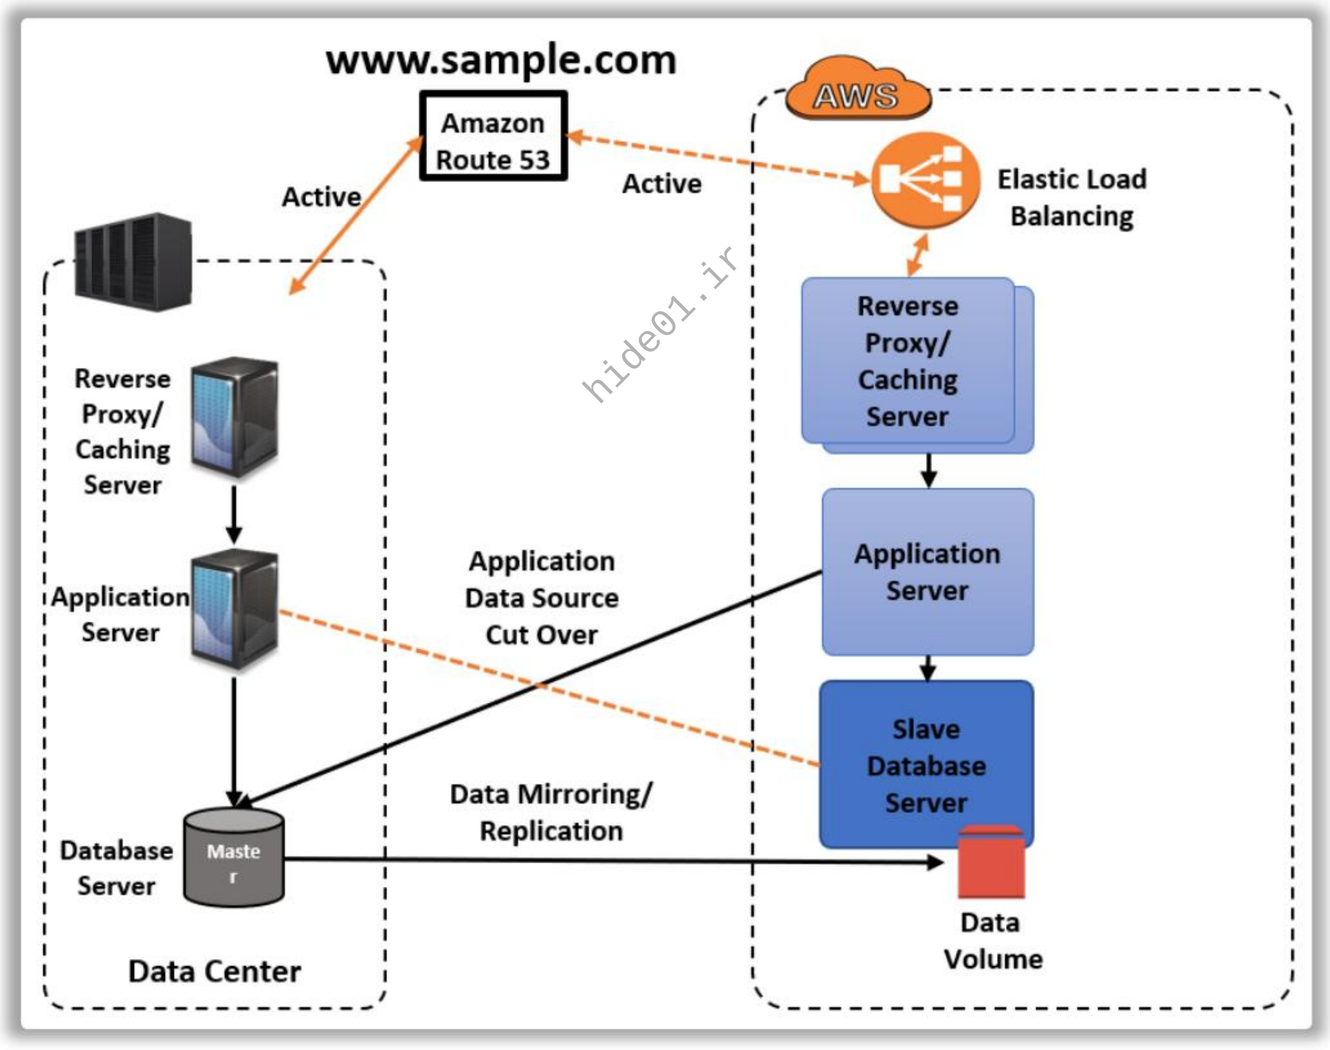

Figure 9.20: Multi-Site Preparation

The key points for preparation are as follows:

- Set up an environment on AWS to duplicate the production environment.
- Distribute the incoming requests to both sites (on-site and AWS) by setting up DNS weighting or any traffic routing technology. Configure automated failover to redirect the production traffic away from the affected site.

Certified Cloud Security Engineer Copyright © by **EC-Council**
All Rights Reserved. Reproduction is Strictly Prohibited.

Recovery Phase

When a disaster strikes and the production environment fails to handle the production load, all the traffic is redirected to the AWS infrastructure by updating the DNS.

The key points for recovery include:

- Change the DNS weighting to redirect all the production traffic to AWS.
- Have some application logic for failover to use the AWS database servers for all production queries.
- To automatically right-size the AWS fleet, use Auto Scaling.

The below image shows how the recovery is performed.

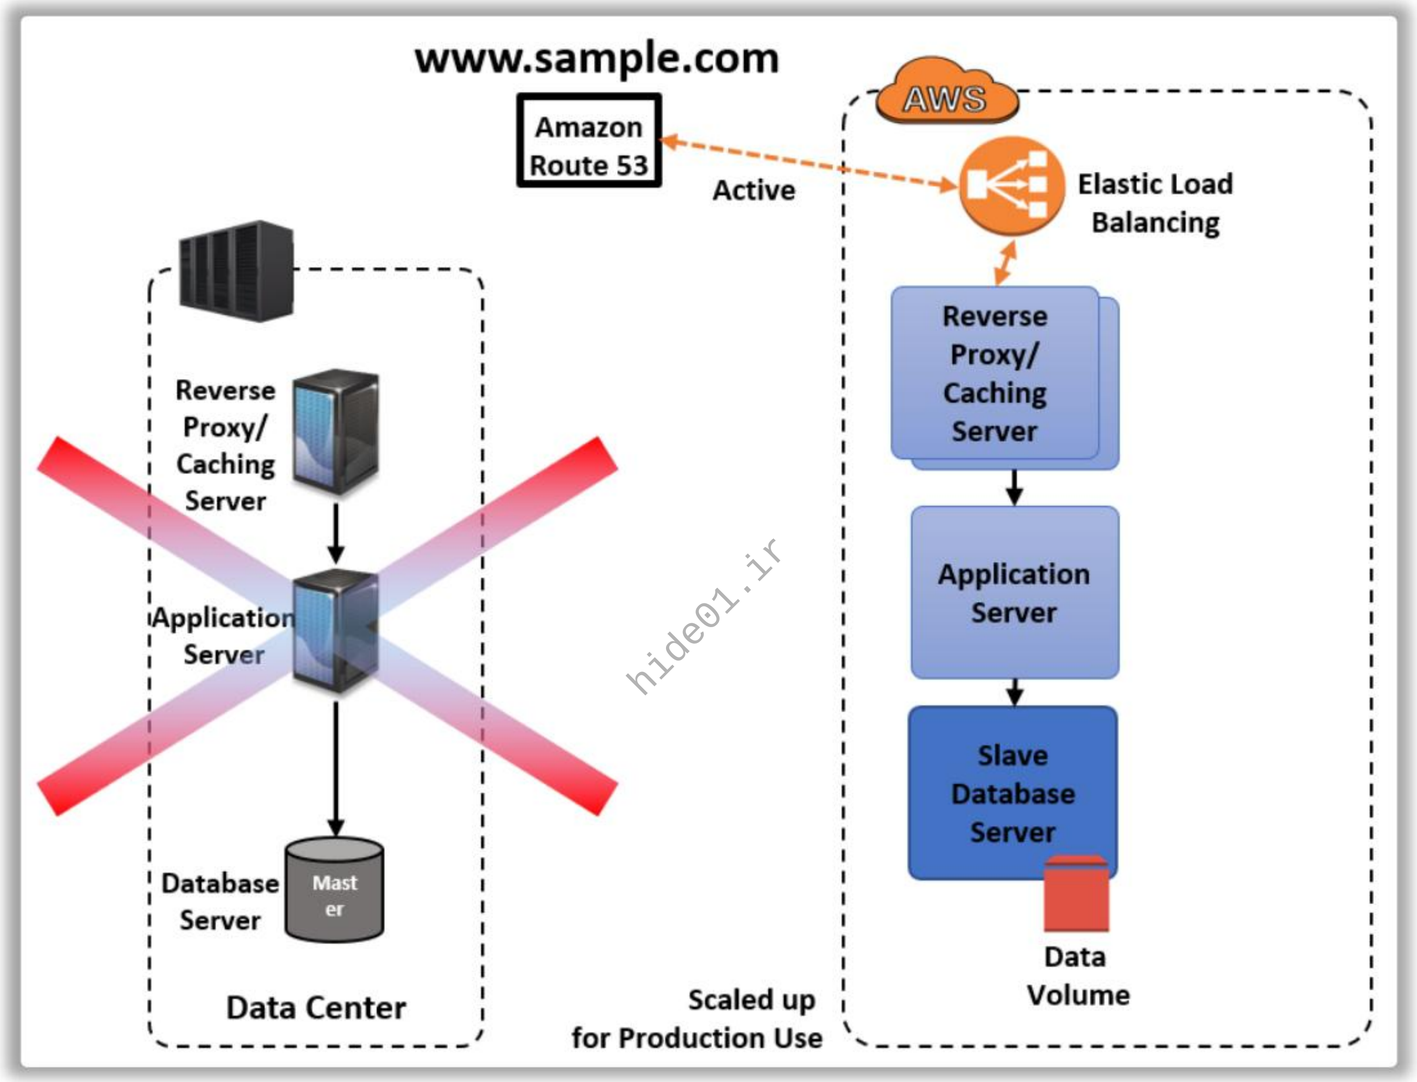

Figure 9.21: Recovery Phase of Multi-site

Certified Cloud Security Engineer Copyright © by **EC-Council**
All Rights Reserved. Reproduction is Strictly Prohibited.

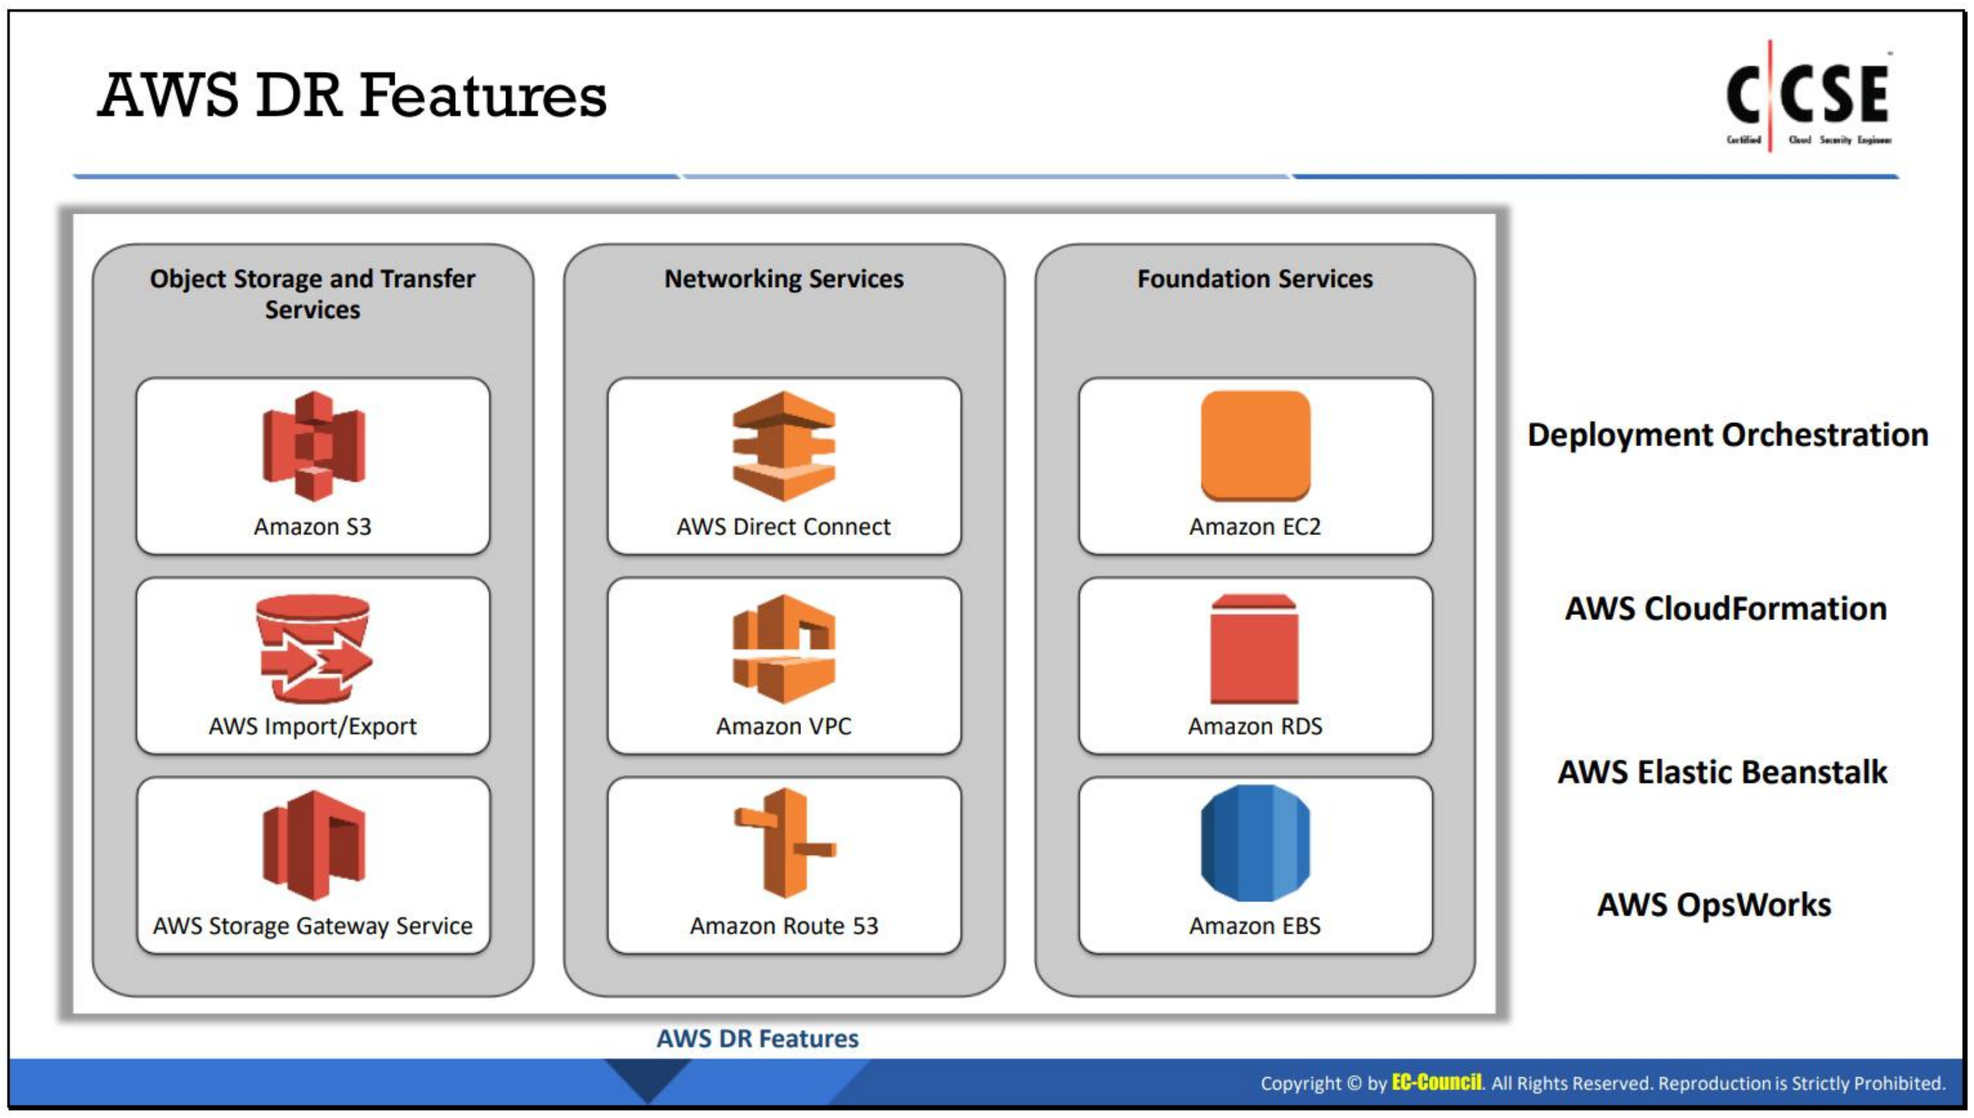

AWS DR Features

At the time of preparing and planning for a DRP using AWS, it is important to have a thorough understanding of the services that you can use. Additionally, the AWS customer should have clarity whether the selected services will support durable storage and data migration. Some features and services that the AWS customer should incorporate in their DRP are object storage and transfer services, networking services, foundation services, deployment orchestration, etc.

Object Storage and Transfer Services

AWS object storage and transfer services that are used for disaster recovery are as follows:

- **Amazon S3**: Amazon S3 provides a durable storage infrastructure that can be used for primary and mission-critical data storage. It is used for disaster recovery because it can store objects redundantly on multiple devices in multiple places within a region and can provide high durability. Versioning and replicating data within an AWS region across several AZs offer extra protection for data retention and archiving. It enables applications to periodically backup files, DB snapshots, and volumes; thus, it is useful for disaster recovery. Backups in the form of ETL jobs are sent to S3 periodically. Through cross-region replication, replication of content is possible across different regions based on the disaster recovery requirements. Amazon S3, along with other AWS services such as compute, networking, and database, helps recover from outages that happen due to natural disasters, human errors, and system failures.
- **Amazon Glacier**: Amazon Glacier is a cloud storage service, like Amazon S2, offered by AWS. It is used as a long-term backup for data that is used infrequently. It supports disaster recovery of inactive data that is seldom used, and it takes a little longer (several

Certified Cloud Security Engineer Copyright © by **EC-Council**
All Rights Reserved. Reproduction is Strictly Prohibited.

hours) for the backup data to be retrieved; however, it is as durable as Amazon S3. This service should not be used for disaster recovery of data that needs quick retrieval.

- **Amazon EBS**: Amazon EBS service is a block storage service that is designed to use with Amazon Elastic Compute Cloud (EC2) and it can be used to create point-in-time snapshots (backups) of data volumes. These backups are stored in Amazon S3 redundantly in multiple AZs. The first EBS snapshot contains a complete copy of the volume and other ongoing snapshots contain only incremental-level changes. This way of disaster recovery is fast and reliable. Data can be securely stored and retrieved easily using Amazon EBS.
- **AWS Import/Export**: This service uses portable storage devices to accelerate the movement of huge amounts of data into and out of AWS. This service provides fast and direct data transfer from and to data storage devices with the help of a high-speed internal network provided by AWS. It can be utilized for transferring data from and to AWS services such as Amazon S3 buckets and EBS snapshots. It offers faster data transfer than the internet; thus, it helps in faster backup-and-restore operations for disaster recovery.
- **AWS Storage Gateway**: It is utilized to connect on-premises software appliance and storage devices with cloud-based storage to enable secure integration between the on-premises infrastructure and AWS storage devices. This service can be used for cloud backup and data recovery when the data center is on-premise. The data stored in the AWS cloud will be available for failover in case of a disaster. AWS Storage Gateway provides disaster recovery with three configurations.
  - **Gateway-cached volumes**: It allows storing of primary data into Amazon S3 and storing of frequently accessed data locally (on-premise). The frequently accessed data is cached in the local storage. This configuration is helpful when fast access to frequently used data is needed.
  - **Gateway-stored volumes**: In this configuration, the primary data is stored in the local data center; however, it is asynchronously backed-up to Amazon S3. This configuration is used when users need low-latency access to the full data set.
  - **Gateway virtual tape library**: This configuration allows the storage of the virtual tape in an Amazon S3 backed VTL or an Amazon Glacier backed virtual tape shelf (VTS).
- **AWS Elastic File System (EFS)**: It is a simple and scalable file storage service that can be used to mount multiple EC2 instances. Through this service, the data present in Amazon EFS can be replicated across several AZs automatically. It provides better durability than Amazon EBS (as it replicates data within an AZ). It is highly available and sturdy for several EC2 instances connected to the service. It stores files in multiple AZs; thus, an IT professional can access them from different AZs. It supports periodic backups to Amazon EFS from on-premises storage services for disaster recovery.

Certified Cloud Security Engineer Copyright © by **EC-Council**
All Rights Reserved. Reproduction is Strictly Prohibited.

Networking Services

During a disaster, when the existing system is failing over to another site, network settings must be modified. The features and services offered by AWS to manage and modify network settings are described below.

- **Amazon Direct Connect**: This service of AWS connects a premise to AWS by setting up a dedicated network connection between them. This service is very useful in performing disaster recovery as it provides better connectivity than internet-based connections.

  For achieving the SLA objective, Amazon Direct Connect Resiliency Toolkit offers a connection wizard that has various resiliency models, which can be used to establish a dedicated connection. The resiliency models are created to make sure that the AWS customer can have a suitable number of dedicated connections in various locations.

- **Amazon Virtual Private Cloud**: Amazon VPC enables the user to use an isolated and private section of the cloud to launch AWS resources in a virtual network defined by the user. A VPN connection can be established between the user's VPC and the on-premises data center. AWS cloud acts as an extension to their organization's data center.

  Amazon VPC can be used to extend the existing network topology to the cloud for disaster recovery. This can help in recovering enterprise applications present on the internal network.

  The AWS customer utilizing Amazon Virtual Private Cloud for disaster recovery can draw all the advantages of a disaster recovery site at a less price. With EBS volume, the user can periodically backup the data from the data center to Amazon EC2 instances or can import the VM images to EC2. For BC, the customer can immediately launch replacement compute capacity in AWS. After the disaster, the user can send the data back to the data center and can terminate the EC2 instance that is not required.

- **Amazon Route 53**: Amazon Route 53 is a scalable and highly available DNS service offered by AWS that is used to route users to internet applications. It can failover between various endpoints and static websites that are hosted in Amazon S3 and has various load-balancing capabilities that are useful in DR scenarios like DNS endpoint health checks.

  Elastic IP addresses are different from traditional IP addresses as they are static IP addresses specifically designed for dynamic cloud computing. They facilitate masking of AZs or instance failures by programmatically remapping the public IP addresses to instances in the user account in a particular region. These elastic IP addresses can be used for disaster recovery by pre-allocating certain IP addresses for most critical systems before a disaster strikes. This will make the execution of the DRP easy as the user already knows the IP addresses.

- **Amazon ELB**: Amazon ELB distributes the incoming application traffic across several Amazon EC2 instances. It offers load-balancing capacity needed in response to incoming application traffic, which helps in achieving great fault tolerance in applications. ELB helps in data resiliency by distributing the incoming traffic to the various instances in a single AZ or multiple AZs. The AWS customer can utilize AWS Global Accelerator with Application

Certified Cloud Security Engineer Copyright © by **EC-Council**
All Rights Reserved. Reproduction is Strictly Prohibited.

Load Balancers for distributing the incoming traffic to various load balancers in one or more AWS regions.

Foundation Services

AWS foundation services that are used for disaster recovery are as follows:

- **Amazon Elastic Compute Cloud (Amazon EC2)**: This service provides the elastic capacity to AWS. Utilizing Amazon EC2, the customer can create EC2 instances within a short time and retain control before, after, and during the DR event. It offers resizable compute capacity in the cloud; thus, it is very helpful in disaster recovery when creating VMs that can be controlled by the user (IT professional).

  AWS stores system images in AMI and these are preconfigured with OS. Certain preconfigured AMIs may have application stacks; however, the user can configure their own AMIs. For DR, it is better to configure and identify own AMIs to launch them as part of the recovery process. Preconfiguring such AMIs with the desired OS and application stacks is recommended.

  Applications running in instances can be protected from the failure of a single location by launching instances in separate AZs. Each region has multiple AZs and one AZ is isolated from another AZ; thus, the failure of one AZ does not cause any effect on the other one.

- **Amazon EC2 VM Import Connector**: This feature helps import virtual machine images to Amazon EC2 instances from an existing environment. The AWS customer can make a DR repository for VM images. For disaster recovery and backup purposes, the customer can import their local VM images into Amazon EC2. The VMs can be imported and stored as AMIs. Based on the requirement of the customer, the created AMIs will be launched quickly in Amazon EC2. If the disaster event strikes the local environment, then the customer can rapidly launch instances to safeguard BC, and concurrently export them to restore the local infrastructure.

  - **Instance Import**: Initially, the AWS user needs to prepare their VM for export, and then export the VM utilizing the supported format. The image of the VM should be uploaded to the Amazon S3 and the instance importing task should be initiated. After completing the importing task, the customer can create an AMI from the stopped instance. The AWS user can copy AMI to another region to launch the instance. The user also has the provision of exporting the previously imported instance. The below diagram illustrates the process of exporting the VMs to Amazon EC2 as an instance.

Certified Cloud Security Engineer Copyright © by **EC-Council**
All Rights Reserved. Reproduction is Strictly Prohibited.

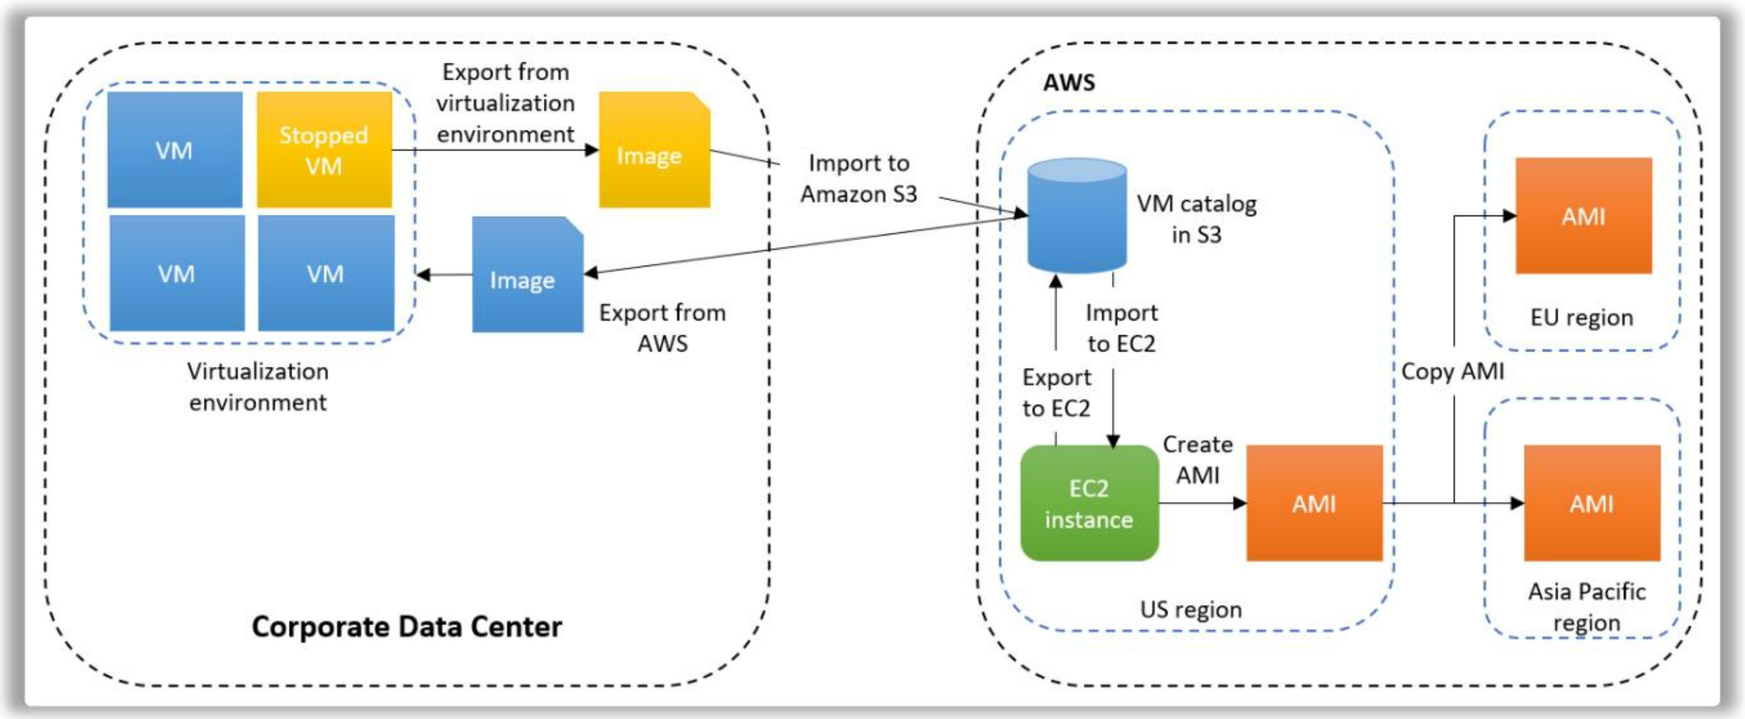

Figure 9.22: Amazon EC2 VM Import Connector

- **Amazon EBS**: Amazon EBS service is a block storage service that is designed to use with Amazon Elastic Compute Cloud (EC2) and it can be used to create point-in-time snapshots (backups) of data volumes. These backups are stored in Amazon S3 redundantly in multiple AZs. The first EBS snapshot contains a complete copy of the volume and other ongoing snapshots contain only incremental-level changes. This way of disaster recovery is fast and reliable. Data can be securely stored and retrieved easily using Amazon EBS.
- **Amazon RDS**: This service can be utilized to setup, operate, and scale the relational database in the cloud. It can also be used to store critical data in an existing database that is running in the DR preparation phase.

  It can also be used to run the production database in the recovery phase. It also helps organizations take backup or snapshots of data from one location to another location. In case the organization operates in multiple regions, it will allow the organization to have a read replica in another region.

  By default, RDS automated backup features are turned on. RDS makes the storage volume snapshot of the entire database instance; it will backup the entire database instance rather than a single database. The initial backup contains the entire database instance backup. The other backups are all incremental ones that have snapshots containing blocks that have changed since the last backup.

  RDS creates the storage volume snapshot of the DB instance, which will backup the entire DB instance rather than individual databases. The first backup consists of a full instance backup. Subsequent backups are incremental with snapshots containing only the blocks that changed since the previous backup. Removal of earlier snapshots would not lead to data loss if a single snapshot referenced to the data is available.
- **Amazon DynamoBD**: This is a fast, simple, and cost-effective database service provided by AWS. It is a fully managed NoSQL database service that can be used to retrieve any amount of data. This service can be used to copy data to Amazon S3 or DynamoBD in another region in the DR preparation phase. Its backup-and-restore operation can help scale up seamlessly within minutes with a single API call in the recovery phase. Backup

hide01.ir

Certified Cloud Security Engineer Copyright © by **EC-Council**
All Rights Reserved. Reproduction is Strictly Prohibited.

creation, restoration, retention, and deletion can be automated using this service through APIs.

Using Amazon DynamoBD**,** the AWS customer can create an on-demand backup of Amazon DynamoBD tables. For enabling continuous backup of data, the customer can use point-in-time restore (PITR), which will safeguard DynamoBD from incidental deletion or writing. The AWS customer using PITR would not have to worry about creation, scheduling, or maintaining on-demand backups. For instance, a test script was written accidentally on production DynamoBD. The AWS customer can use PITR to restore the table to the desired point in time within the last 35 days. The AWS customer can also bring back the table to the desired point in time utilizing DynamoBD API, AWS Management Console, or AWS Command Line Interface (CLI). The incremental backups of the AWS customer table are maintained by DynamoDB.

- **Amazon Redshift**: This is a fast, simple, and cost-effective data warehouse service provided by AWS that can help analyze data by using current business intelligence tools. It is a fully managed data warehouse service that can be used for disaster recovery. It can be used to snapshot a data warehouse to be stored in Amazon S3 durably in the DR preparation phase. It can be used to rapidly restore the data warehouse into the same region or another AWS region in the recovery phase of DR.

  In Amazon Redshift, the snapshots are backups of the cluster. Using the SSL connection, Amazon Redshift internally stores the snapshots in Amazon S3. Redshift takes the incremental snapshots automatically. These automated snapshots consist of all the data that is required to create a cluster. If AWS customers enable automated snapshots for the cluster, then Redshift will take the snapshots periodically—the snapshots are taken every 8 hours or data change amounting to 5 GB per node. The AWS customer can also make a snapshot schedule for taking automated snapshots. The automated snapshots are enabled by default in Amazon Redshift at the time of cluster creation. After the retention period, the snapshots are automatically deleted.

  The AWS customer can also configure Amazon Redshift to automatically copy the snapshot from a cluster to another region of AWS. Whenever the AWS customer creates a snapshot in the primary region of the cluster, then the snapshot will be automatically copied to the secondary region of the AWS. These two regions are considered as source AWS region and destination AWS region. If disaster strikes and affects the source AWS region, then the AWS customer can easily create the cluster using snapshots from the destination AWS region.

Deployment Orchestration

AWS customers can enhance their recovery capability by utilizing automation processes and post-startup software installation or configuration. The deployment orchestration services provided by AWS for BCDR are as follows:

- **AWS CloudFormation**: This AWS service helps system administrators and developers make a set of related AWS resources easily and use or provision them in an orderly

Certified Cloud Security Engineer Copyright © by **EC-Council**
All Rights Reserved. Reproduction is Strictly Prohibited.

manner. It can be used to deploy associated resource collections as required and create templates for environments.

AWS CloudFormation is a robust tool provided by AWS that helps in implementing the deployed infrastructure across various AWS regions. Using a single operation, it will allow the AWS customer to create, delete, or update CloudFormation stacks among multiple accounts and regions. It utilizes JSON or YAML for defining Infrastructure as a Code (IaC); however, using AWS Cloud Development Kit (CDK), the AWS customer can define IaC utilizing a well-acquainted programming language. After converting the code to CloudFormation, it is used for deploying resources in AWS.

The AWS customer needs to deploy the scaled-down version of their core workloads with minimum resources in the disaster recovery region. With the help of CloudFormation, the AWS customer can define their infrastructure and consistently deploy it among AWS accounts across various regions. CloudFormation utilizes predefined parameters for finding the AWS accounts or regions of its deployment. Hence, the AWS customer should implement conditional logic in their CloudFormation templates for deploying only the scaled-down version in the disaster recovery region.

- **AWS Elastic Beanstalk**: This service helps in the deployment and scaling of web applications and services that are developed with PHP, Java, .NET, Ruby, Python, and Docker. This will help in disaster recovery by making the recovery process fast. When the user deploys an application code, it will provide the operating environment for the application.

  The AWS global infrastructure is made across AWS AZs and regions. There are many AZs that are physically separated in a region. Using AZs, the AWS customer can operate the application and failover to another AZ without any hindrance. On behalf of AWS customer, the AWS Elastic Beanstalk automates and manage the usage of AWS global infrastructure. With AWS Elastic Beanstalk, the customer can draw benefits of fault tolerance mechanisms and availability offered by AWS.

- **AWS OpsWorks**: It is an application management service offered by AWS that makes the deployment and operation of applications (of all types and sizes) very easy. It can replace its host automatically in the event of an instance failure. It can be used to template a user environment in the preparation phase of DR. During the recovery phase, it can also be used along with AWS CloudFormation.

  A user can provision a new stack instantly by utilizing the stored configuration that supports defined RTO. It makes the deployment process very easy in the disaster recovery phase because it provides essential tools to users for creating an environment as a series of layers and each layer can be configured as a tier of their application.

Certified Cloud Security Engineer Copyright © by **EC-Council**
All Rights Reserved. Reproduction is Strictly Prohibited.

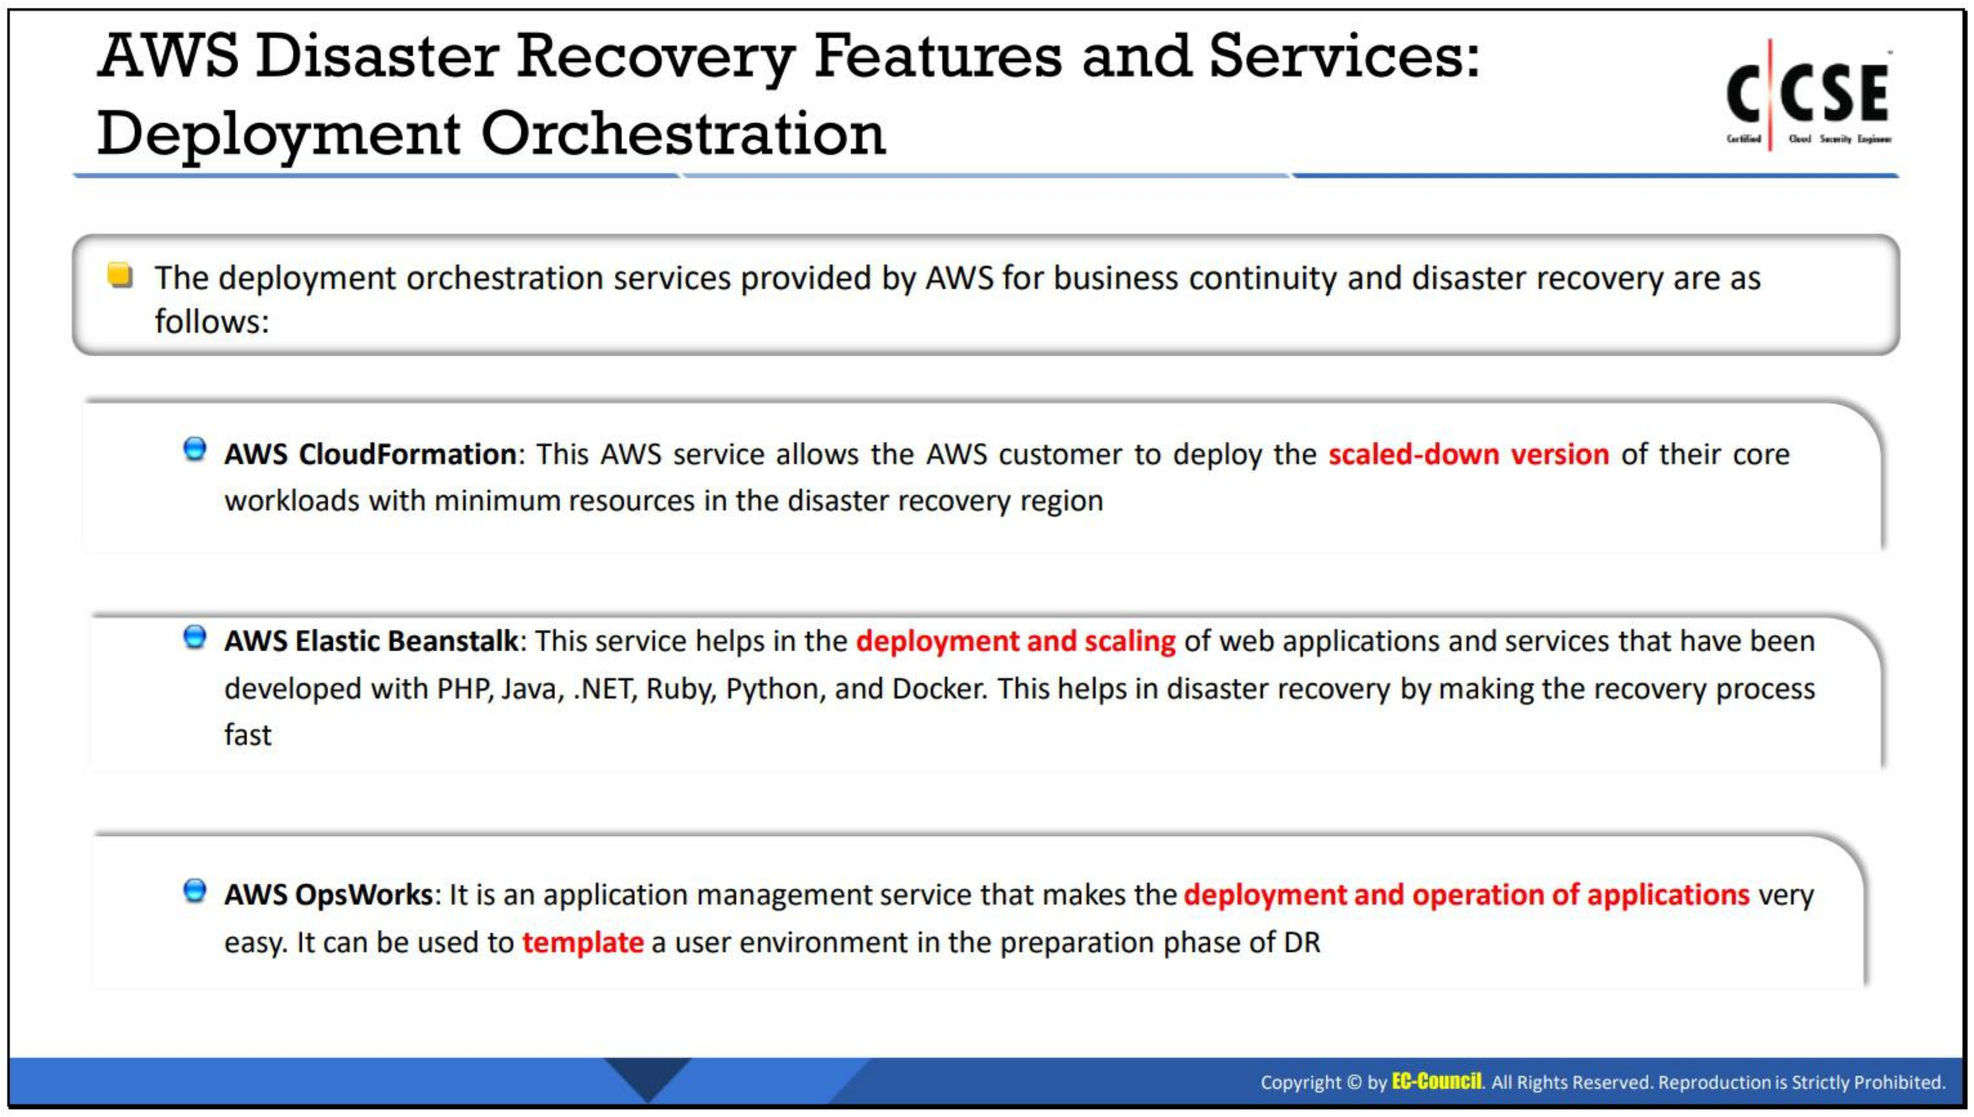

AWS Disaster Recovery Features and Services: Deployment Orchestration

Amazon EC2 supports the use of deployment automation, post-start up installation or configuration processes, and tools that can help in the recovery phase of the DR. This will help create the necessary resources in an automated manner.

The deployment orchestration services provided by AWS for BCDR are as follows:

- **AWS CloudFormation**: This AWS service helps system administrators and developers make a set of related AWS resources easily and use or provision them in an orderly manner. It can be used to deploy associated resource collections as required and create templates for environments.
- **AWS Elastic Beanstalk:** This service helps in the deployment and scaling of web applications and services that are developed with PHP, Java, .NET, Ruby, Python, and Docker. This will help in disaster recovery by making the recovery process fast. When the user deploys an application code, it will provision the operating environment for the application.
- **AWS OpsWorks**: It is an application management service offered by AWS that makes the deployment and operation of applications (of all types and sizes) very easy. It has the ability to replace its host automatically in the event of an instance failure. It can be used to template a user environment in the preparation phase of DR and can be used together with AWS CloudFormation in the recovery phase. The user can use the stored configuration that supports the defined RTO to provision a new stack quickly. It makes the deployment process very easy in the disaster recovery phase because it provides essential tools to users for creating an environment as a series of layers and each layer can be configured as a tier of their application.

Certified Cloud Security Engineer Copyright © by **EC-Council**
All Rights Reserved. Reproduction is Strictly Prohibited.

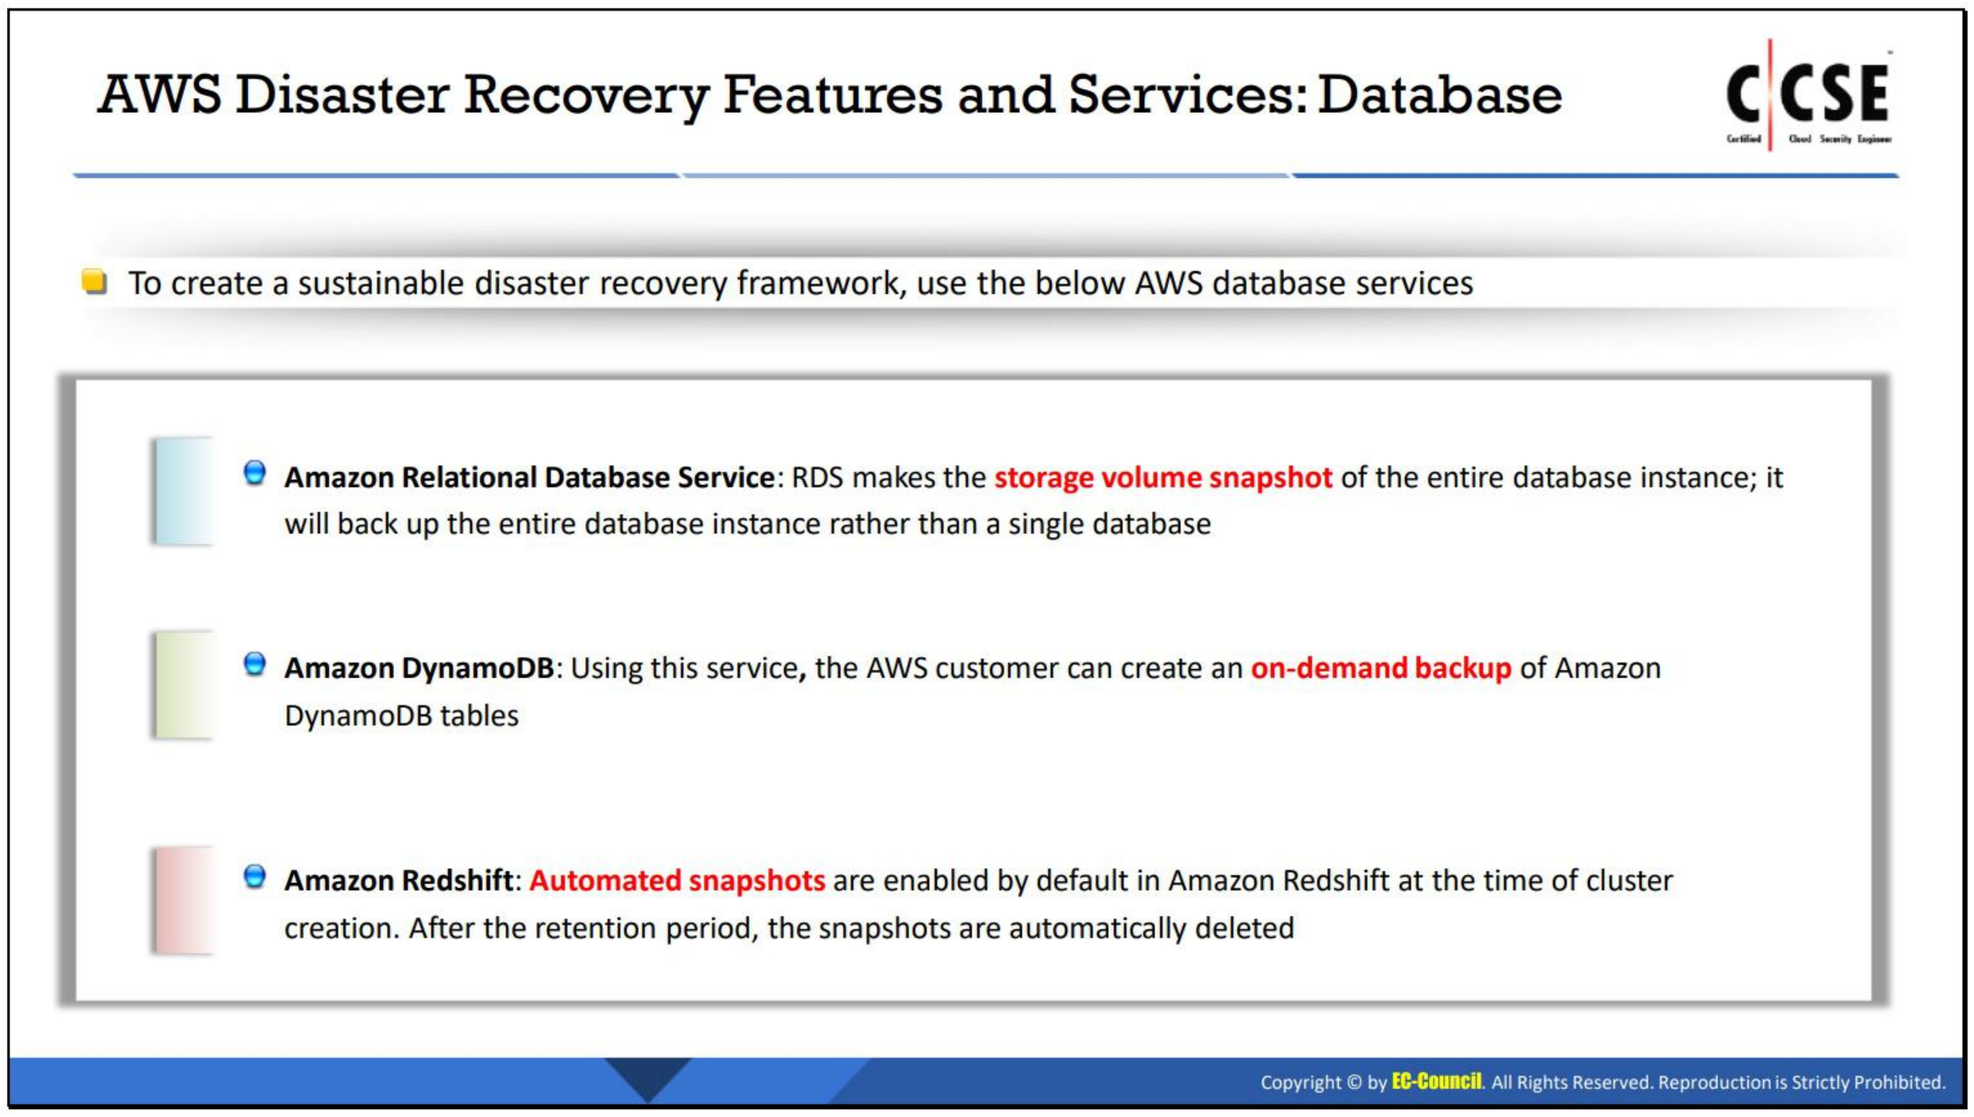

AWS Disaster Recovery Features and Services: Database

AWS provides certain important features to fulfill database requirements for disaster recovery. To create a sustainable disaster recovery framework, use the following AWS database services:

- **Amazon RDS**: This service helps in setting up, operating, and scaling a relational database in the cloud and can be used to store critical data in an existing database that is running in the DR preparation phase. It can also be used to run the production database in the recovery phase. It also enables organizations to snapshot data from one region to another and allows them to have a read replica running in another region if they operate in multiple regions.
- **Amazon DynamoBD**: This is a fast, simple, and cost-effective database service provided by AWS. It is a fully managed NoSQL database service that store and retrieve data without any limits and can handle any amount of network traffic. This service can be used to copy data to Amazon S3 or DynamoBD in another region in the DR preparation phase. Its backup-and-restore operation can help to scale up seamlessly within minutes with a single API call in the recovery phase. Its PITR feature offers continuous backups of DynamoBD table data. Backup creation, restoration, retention, and deletion can be automated by using this service through APIs.
- **Amazon Redshift**: This is a fast, simple, and cost-effective data warehouse service provided by AWS that can analyze data by using current business intelligence tools. It is a fully managed data warehouse service that can be used for disaster recovery. It can be used to snapshot a data warehouse to be stored in Amazon S3 durably in the DR preparation phase. It can be used to rapidly restore the data warehouse into the same region or another AWS region in the recovery phase of DR.

Certified Cloud Security Engineer Copyright © by **EC-Council**
All Rights Reserved. Reproduction is Strictly Prohibited.

AWS Disaster Recovery Features and Services: Storage

- AWS storage services that are used for disaster recovery are as follows:
  - **Amazon S3**: It is used for disaster recovery because it can **store objects redundantly** on multiple devices in multiple places within a region and can provide high durability
  - **Amazon Glacier**: It supports disaster recovery of **inactive data that is seldom used** and it takes a little longer (several hours) for the backup data to be retrieved but it is as durable as Amazon S3
  - **Amazon Elastic Block Store**: It is used to **create point-in-time snapshots** (backups) of data volumes

Copyright © by EC-Council. All Rights Reserved. Reproduction is Strictly Prohibited.

AWS Disaster Recovery Features and Services: Storage (Cont'd)

hide01.ir

- **AWS Import/Export**: It offers faster data transfer than the internet; thus, it helps in **faster backup-and-restore operations** for disaster recovery
- **AWS Storage Gateway**: This service can be used for cloud backup and data recovery when the **data center** is on-premises. The **data stored** in the AWS cloud will be available for **failover** in case of a disaster
- **AWS Elastic File System**: It supports **periodic backups** to Amazon EFS from **on-premises storage services** for disaster recovery

Copyright © by EC-Council. All Rights Reserved. Reproduction is Strictly Prohibited.

AWS Disaster Recovery Features and Services: Storage

AWS storage services that are used for disaster recovery are as follows:

- **Amazon S3**: Amazon S3 provides a durable storage infrastructure that can be used for primary and mission-critical data storage. It is used for disaster recovery because it can store objects redundantly on multiple devices in multiple places within a region and can provide high durability. Versioning and replicating data within an AWS region across

Certified Cloud Security Engineer Copyright © by **EC-Council**
All Rights Reserved. Reproduction is Strictly Prohibited.

several AZs offer extra protection for data retention and archiving. It enables applications to periodically backup files, DB snapshots, and volumes; thus, it is useful for disaster recovery. Backups in the form of ETL jobs are sent to S3 periodically. Through cross-region replication, replication of content is possible across different regions based on the disaster recovery requirements. Amazon S3, along with other AWS services like compute, networking, and database, helps recover from outages that happen due to natural disasters, human errors, and system failures.

- **Amazon Glacier**: Amazon Glacier is a cloud storage service, like Amazon S2, offered by AWS. It is used as a long-term backup for data that is used infrequently. It supports disaster recovery of inactive data that is seldom used, and it takes a little longer (several hours) for the backup data to be retrieved; however, it is as durable as Amazon S3. This service should not be used for disaster recovery of data that needs quick retrieval.
- **Amazon EBS**: Amazon EBS service is a block storage service that is designed to use with Amazon Elastic Compute Cloud (EC2) and it can be used to create point-in-time snapshots (backups) of data volumes. These backups are stored in Amazon S3 redundantly in multiple AZs. The first EBS snapshot contains a complete copy of the volume and other ongoing snapshots contain only incremental-level changes. This way of disaster recovery is fast and reliable. Data can be securely stored and retrieved easily using Amazon EBS.
- **AWS Import/Export**: This service uses portable storage devices to accelerate the movement of huge amounts of data into and out of AWS. This service provides fast and direct data transfer from and to data storage devices with the help of a high-speed internal network provided by AWS. It can be used to transfer data into and out of Amazon S3 buckets and Amazon EBS snapshots. It offers faster data transfer than the internet; thus, it helps in faster backup-and-restore operations for disaster recovery.
- **AWS Storage Gateway**: This AWS service helps customers to connect their on-premises storage devices with cloud-based storage to enable secure integration between their on-premises IT infrastructure and AWS storage infrastructure. This service can be used for cloud backup and data recovery when the data center is on-premise. The data stored in the AWS cloud will be available for failover in case of a disaster. AWS Storage Gateway provides disaster recovery with three configurations.
  - **Gateway-cached volumes**: It allows storing of primary data into Amazon S3 and storing of frequently accessed data locally (on-premise). The frequently accessed data is cached in the local storage. This configuration is helpful when fast access to frequently used data is needed.
  - **Gateway-stored volumes**: In this configuration, the primary data is stored in the local data center; however, it is asynchronously backed-up to Amazon S3. This configuration is used when users need low-latency access to the full data set.
  - **Gateway virtual tape library**: This configuration allows the storage of the virtual tape in an Amazon S3 backed VTL or an Amazon Glacier backed VTS.

Certified Cloud Security Engineer Copyright © by **EC-Council**
All Rights Reserved. Reproduction is Strictly Prohibited.

- **AWS EFS**: It is a simple and scalable file storage service that can be used to mount multiple EC2 instances. Through this service, the data present in Amazon EFS can be replicated across several AZs automatically. It provides better durability than Amazon EBS (as it replicates data within an AZ). It is highly available and sturdy for several EC2 instances connected to the service. It stores files in multiple AZs; thus, an IT professional can access them from different AZs. It supports periodic backups to Amazon EFS from on-premises storage services for disaster recovery.

hide01.ir

Certified Cloud Security Engineer Copyright © by **EC-Council**
All Rights Reserved. Reproduction is Strictly Prohibited.

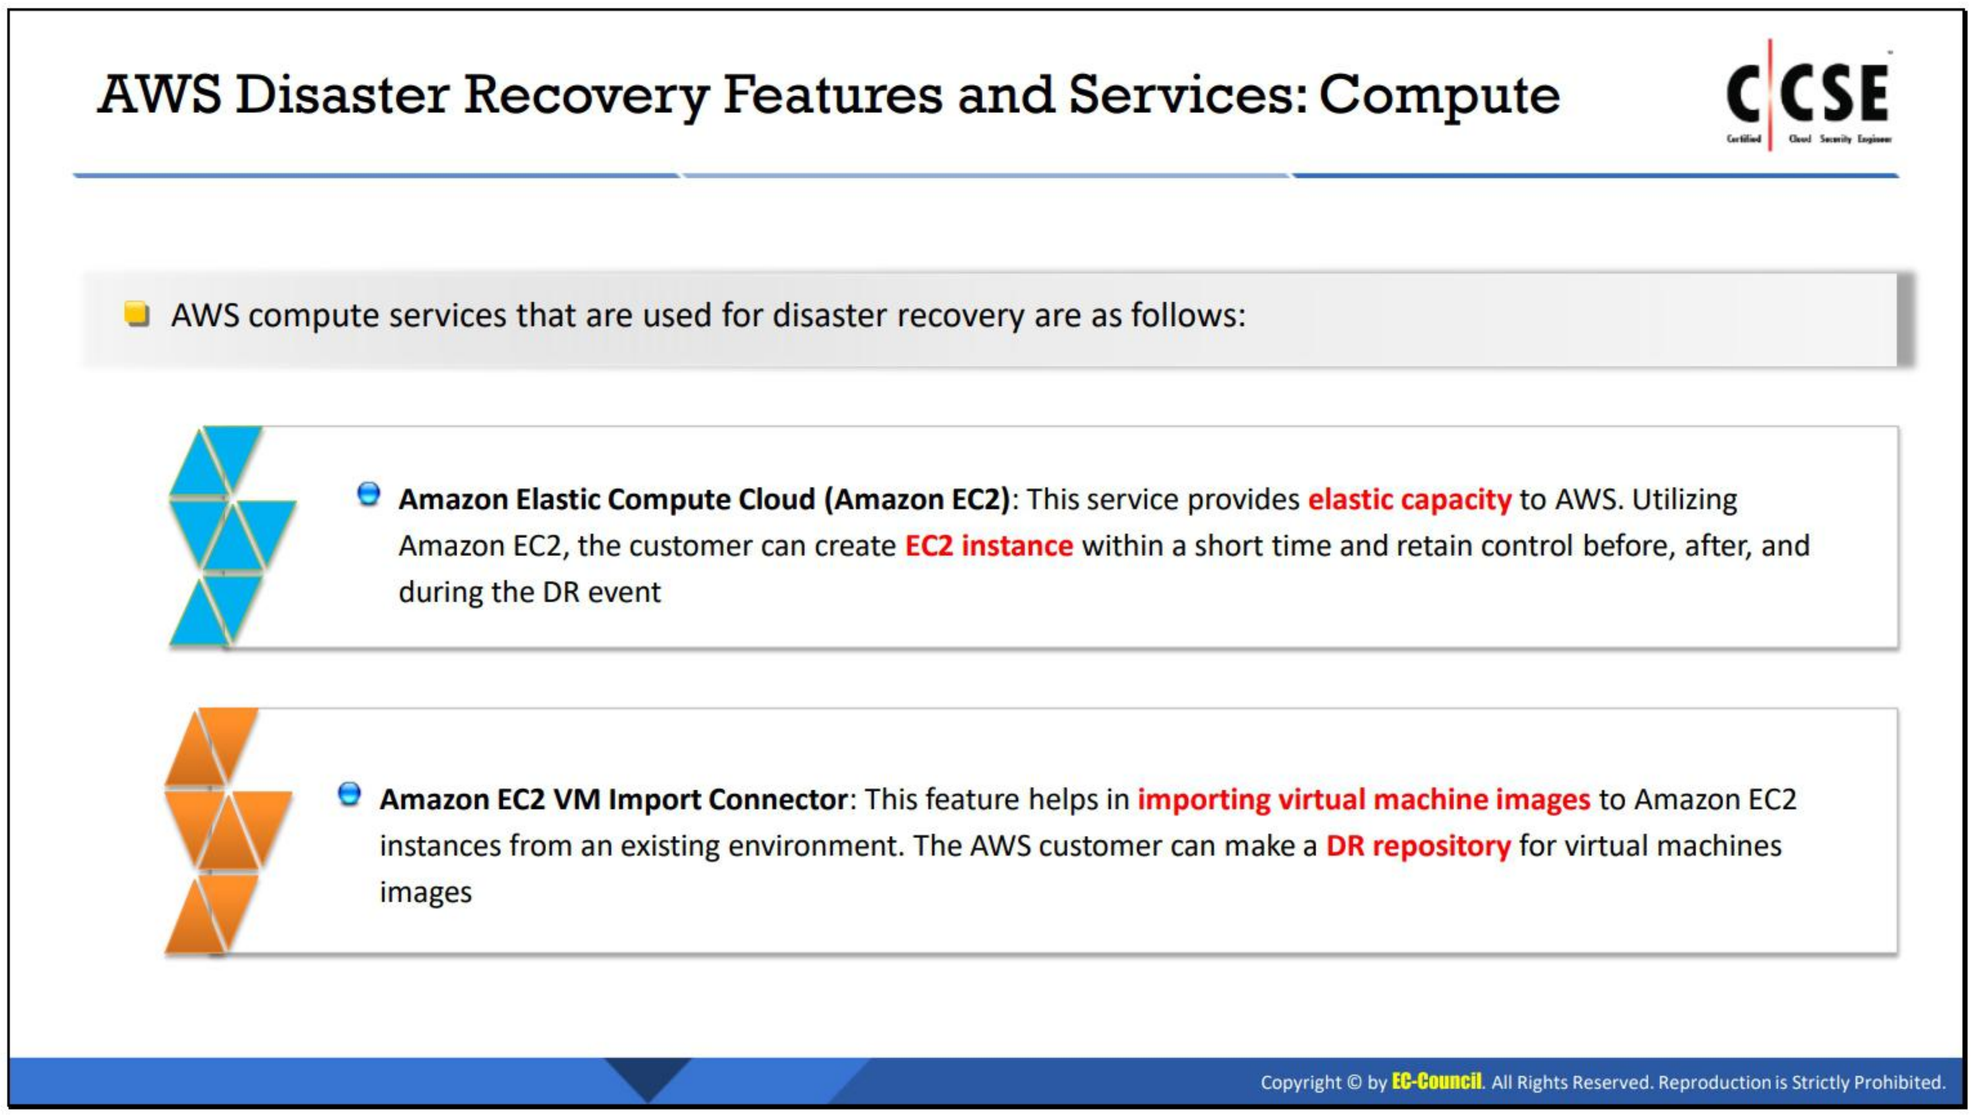

AWS Disaster Recovery Features and Services: Compute

AWS compute services that are used for disaster recovery are as follows:

- **Amazon Elastic Compute Cloud (Amazon EC2)**: This service provides the elastic capacity to AWS. Utilizing Amazon EC2, the customer can create EC2 instances within a short time and retain control before, after, and during the DR event. It offers resizable compute capacity in the cloud; thus, it is very helpful in disaster recovery when creating VMs that can be controlled by the user (IT professional).

  AWS stores system images in AMI and these are preconfigured with OS. Certain preconfigured AMIs may have application stacks; however, the user can configure their own AMIs. For DR, it is better to configure and identify own AMIs to launch them as part of the recovery process. Preconfiguring such AMIs with the desired OS and application stacks is recommended.

  Applications running in instances can be protected from the failure of a single location by launching instances in separate AZs. Each region has multiple AZs and one AZ is isolated from another AZ; thus, the failure of one AZ does not cause any effect on the other one.

- **Amazon EC2 VM Import Connector**: This feature helps import virtual machine images to Amazon EC2 instances from an existing environment. The AWS customer can make a DR repository for VM images. For disaster recovery and backup purposes, the customer can import their local VM images into Amazon EC2. The VMs can be imported and stored as AMIs. Based on the requirement of the customer, the created AMIs will be launched quickly in Amazon EC2. If the disaster event strikes the local environment, then the customer can rapidly launch instances to safeguard BC, and concurrently export them to restore the local infrastructure.

Certified Cloud Security Engineer Copyright © by **EC-Council**
All Rights Reserved. Reproduction is Strictly Prohibited.

- **Instance Import**: Initially, the AWS user needs to prepare their VM for export, and then export the VM utilizing the supported format. The image of the VM should be uploaded to the Amazon S3 and the instance importing task should be initiated. After completing the importing task, the customer can create an AMI from the stopped instance. The AWS user can copy AMI to another region to launch the instance. The user also has the provision of exporting the previously imported instance. The below diagram illustrates the process of exporting the VMs to Amazon EC2 as an instance.

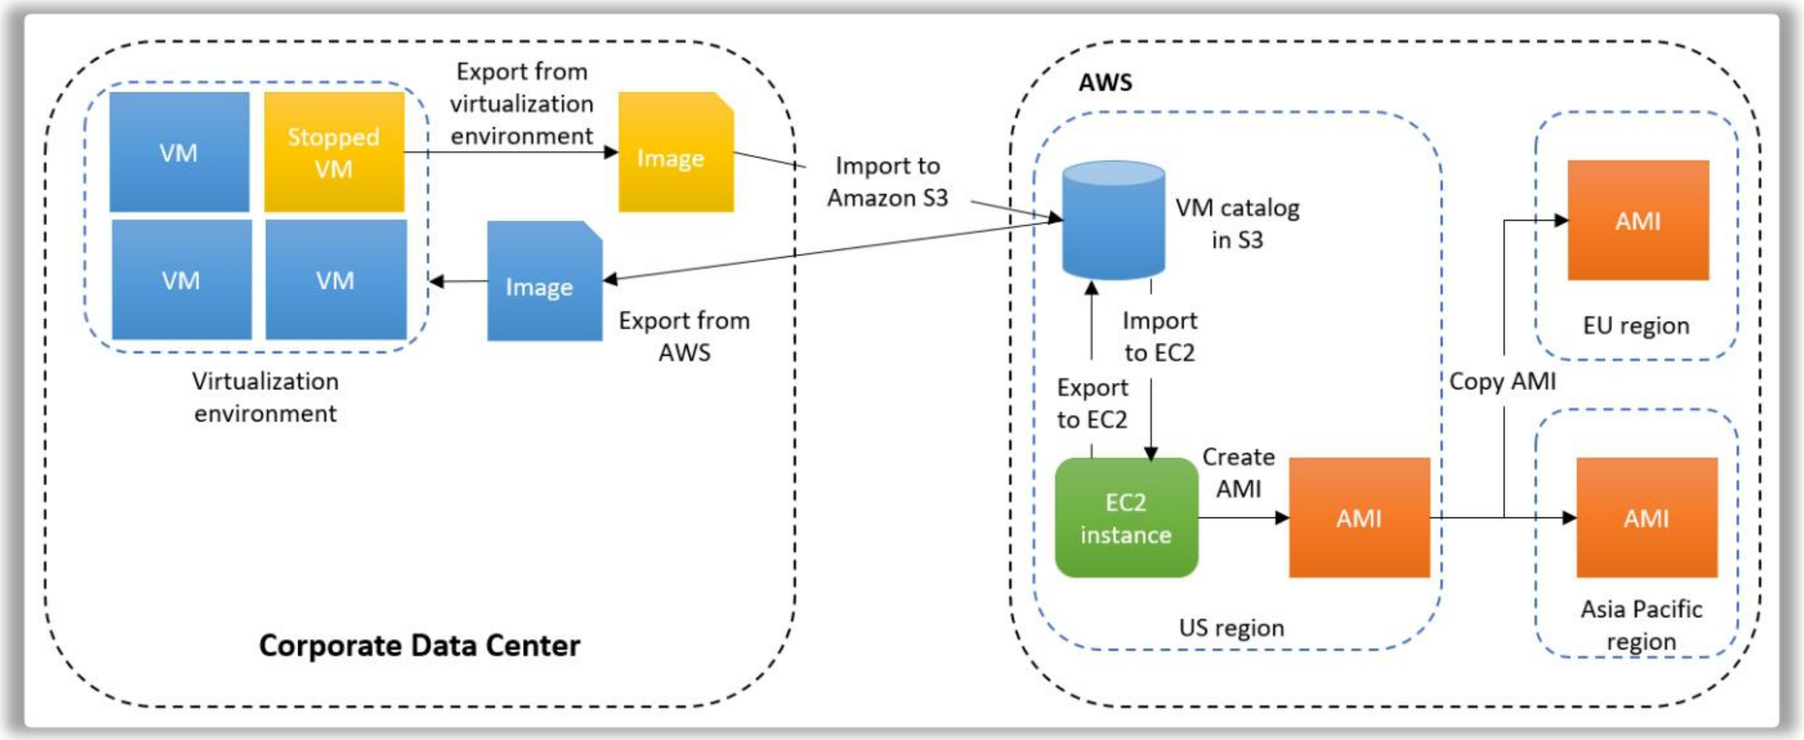

Figure 9.23: Exporting VM to Amazon EC2 VM

hide01.ir

Certified Cloud Security Engineer Copyright © by **EC-Council**
All Rights Reserved. Reproduction is Strictly Prohibited.

AWS Disaster Recovery Features and Services: Networking

- During a disaster, when the existing system is failing over to another site, **network settings** must be modified
- Some **features and services** offered by AWS to manage and modify network settings are as follows:

Amazon Route53	It can **failover** between various **endpoints and static websites** that are hosted in **Amazon S3** and has various load-balancing capabilities that are useful in DR scenarios such as DNS endpoint health checks
Elastic Load Balancing	Elastic load balancing helps in data resiliency by **distributing the incoming traffic** to various instance in single AZ or multiple AZs
Amazon Direct Connect	For achieving the **SLA objective**, Amazon Direct Connect Resiliency Toolkit offers a connection wizard that has various **resiliency models**, which can be used to **establish a dedicated connection**
Amazon Virtual Private Cloud	Amazon VPC can be used to extend the **existing network topology** to the cloud for disaster recovery, which can help in recovering enterprise applications present on the internal network

Copyright © by EC-Council. All Rights Reserved. Reproduction is Strictly Prohibited.

AWS Disaster Recovery Features and Services: Networking

During a disaster, when the existing system is failing over to another site, network settings must be modified. Some features and services offered by AWS to manage and modify network settings are described below.

hide01.ir

- **Amazon Route 53**: Amazon Route 53 is a scalable and highly available DNS service offered by AWS that is used to route users to internet applications. It can failover between various endpoints and static websites that are hosted in Amazon S3 and has various load-balancing capabilities that are useful in DR scenarios like DNS endpoint health checks.

  Elastic IP addresses are different from traditional IP addresses as they are static IP addresses specifically designed for dynamic cloud computing. They facilitate masking of AZs or instance failures by programmatically remapping the public IP addresses to instances in the user account in a particular region. These elastic IP addresses can be used for disaster recovery by pre-allocating certain IP addresses for most critical systems before a disaster strikes. This will make the execution of the DRP easy as the user already knows the IP addresses.

- **Amazon ELB**: Amazon ELB distributes the incoming application traffic across several Amazon EC2 instances. It offers load-balancing capacity needed in response to incoming application traffic, which helps in achieving great fault tolerance in applications. ELB helps in data resiliency by distributing the incoming traffic to the various instances in a single AZ or multiple AZs. The AWS customer can utilize AWS Global Accelerator with Application Load Balancers for distributing the incoming traffic to various load balancers in one or more AWS regions.

Certified Cloud Security Engineer Copyright © by **EC-Council**
All Rights Reserved. Reproduction is Strictly Prohibited.

- **Amazon Direct Connect**: This service of AWS connects a premise to AWS by setting up a dedicated network connection between them. This service is very useful in performing disaster recovery as it provides better connectivity than internet-based connections.

  For achieving the SLA objective, Amazon Direct Connect Resiliency Toolkit offers a connection wizard that has various resiliency models, which can be used to establish a dedicated connection. The resiliency models are created to make sure that the AWS customer can have suitable number of dedicated connections in various locations.

- **Amazon Virtual Private Cloud**: Amazon VPC enables the user to use an isolated and private section of the cloud to launch AWS resources in a virtual network defined by the user. The user can create a VPN connection between their VPC and data center and use the AWS cloud as an extension of their corporate data center. Amazon VPC can be used to extend the existing network topology to the cloud for disaster recovery. This can help in recovering enterprise applications present on the internal network.

  The AWS customer utilizing Amazon Virtual Private Cloud for disaster recovery can draw all the advantages of disaster recovery site at a less price. With EBS volume, the user can periodically backup the data from the data center to Amazon EC2 instances or can import the VM images to EC2. For BC, the customer can immediately launch replacement compute capacity in AWS. After the disaster, the user can send the data back to the data center and can terminate the EC2 instance that are not required.

hide01.ir

Certified Cloud Security Engineer Copyright © by **EC-Council**
All Rights Reserved. Reproduction is Strictly Prohibited.

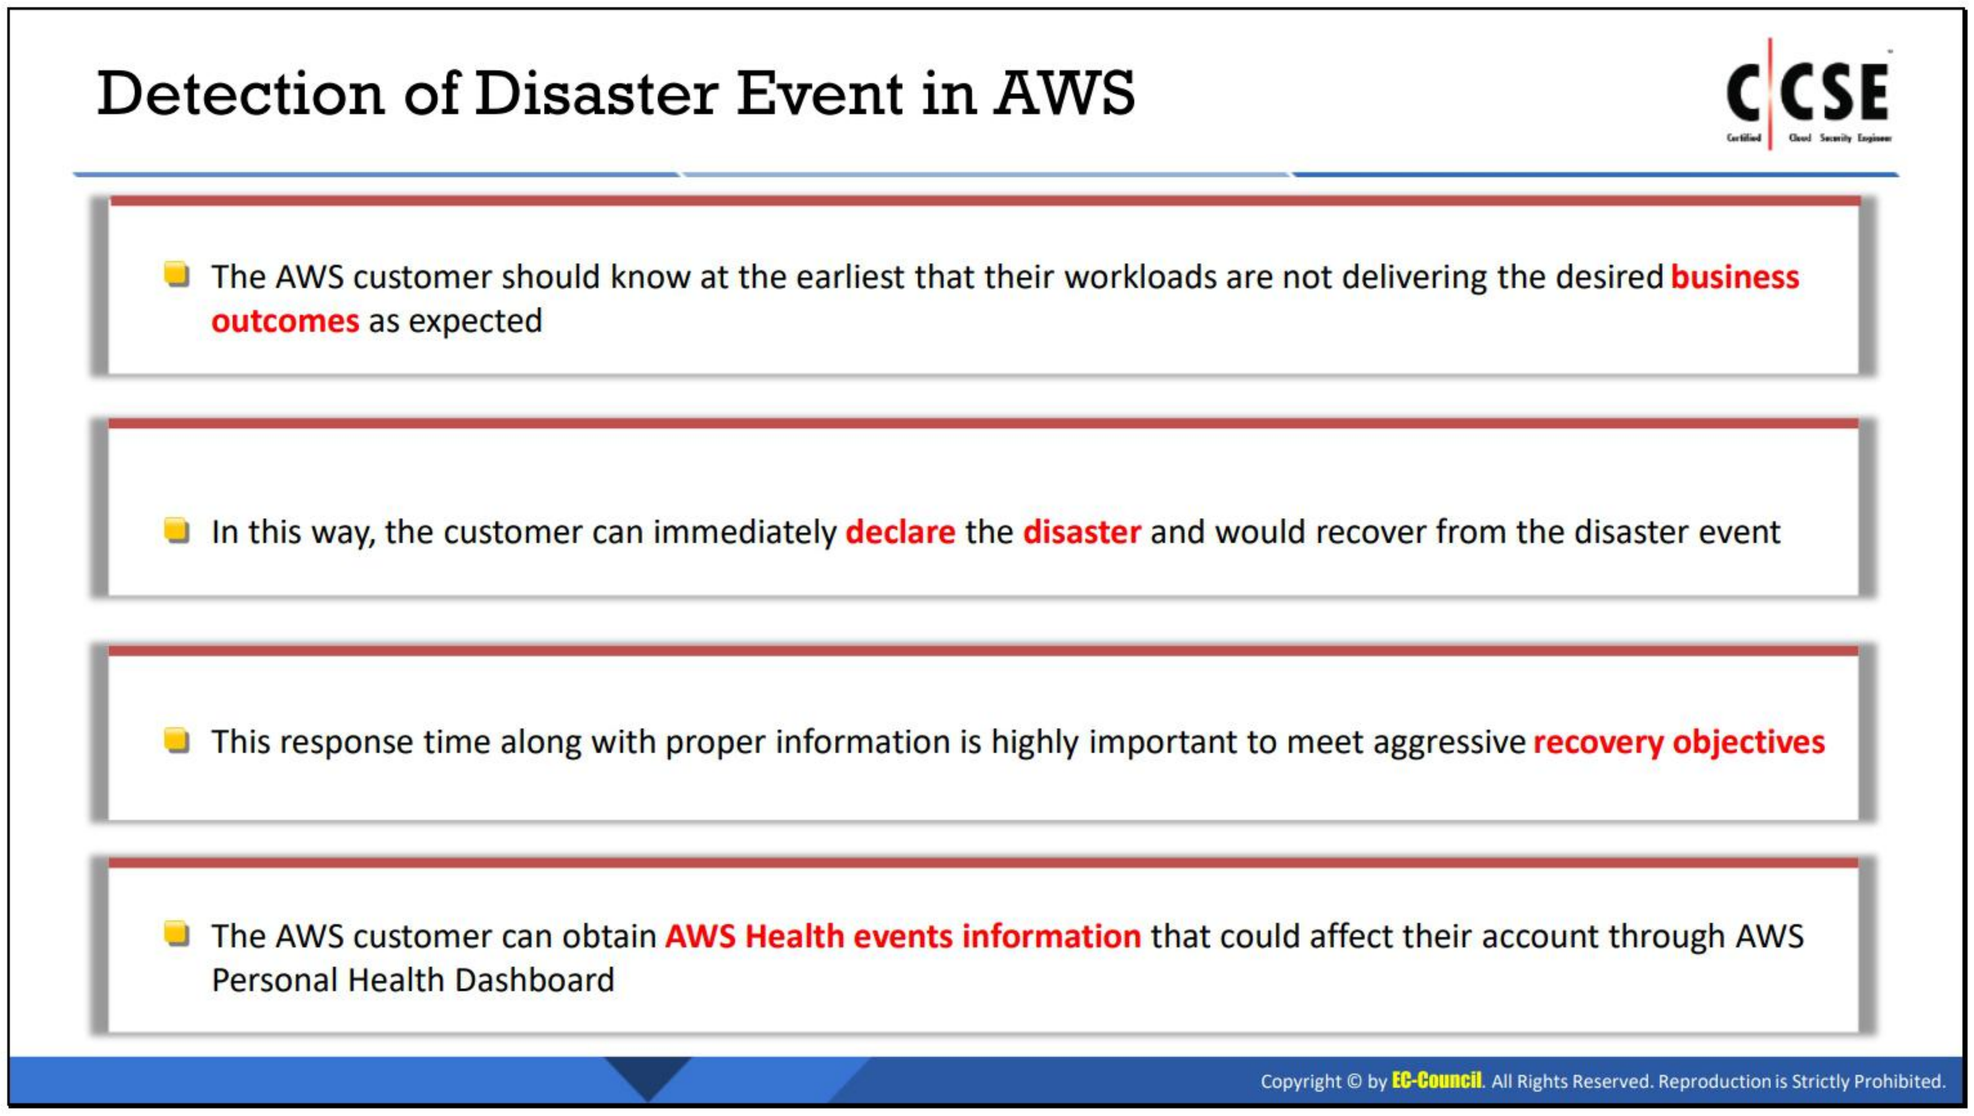

Detection of Disaster Event in AWS

The AWS customer should know at the earliest that their workloads are not delivering the desired business outcomes as expected. In this way, the customer can immediately declare the disaster and would recover from the disaster event. This response time along with proper information is highly important to meet aggressive recovery objectives. If the AWS customer RPO is 1 hour, then the AWS customer must detect an incident; inform the respective personnel regarding the incident; get involved in the escalation process, without starting the DRP if there is any information on RTO, then start evaluating it; declare the disaster and recover in an hour.

If the stakeholder decides not to invoke disaster recovery despite the RTO being at risk, then the DRPs and objectives should be re-evaluated. The decision of not invoking the DRP suggests that the plan is not adequate or there is a lack of confidence in implementing the DRP.

To achieve objectives that deliver business value, it is crucial to factor-in incident detection, notification, escalation, discovery, and declaration in the planning and objectives.

The AWS customer can find the updated service availability information on the Service Health Dashboard. For obtaining the status information, the user can check the dashboard at any point in time. The AWS customer can also subscribe to the RSS feed to obtain an alert on interruption to each service. The customer can also create a support request if they encounter any real-time operation issue with any AWS service that is not shown on the Service Health Dashboard.

The AWS customer can obtain AWS Health events information that could affect their account through AWS Personal Health Dashboard. This information can be found in two ways. One way is through the dashboard, which will show the recent and upcoming events based on the selected category, and the other way is through the full event log that will show all events over the past 90 days.

Certified Cloud Security Engineer Copyright © by **EC-Council**
All Rights Reserved. Reproduction is Strictly Prohibited.

Depending on the health checks, if the RTO requirements are stringent, then enforce automated failover. The health checks should be designed in such a way that it represents the user experience; moreover, it should be based on the desired key performance indicators. Depending on multiple signals, utilize the deep health checks. It is recommended to use caution with this approach to prevent triggering of false alarms as failing over without any necessity could bring in availability risks.

hide01.ir

Certified Cloud Security Engineer Copyright © by **EC-Council**
All Rights Reserved. Reproduction is Strictly Prohibited.

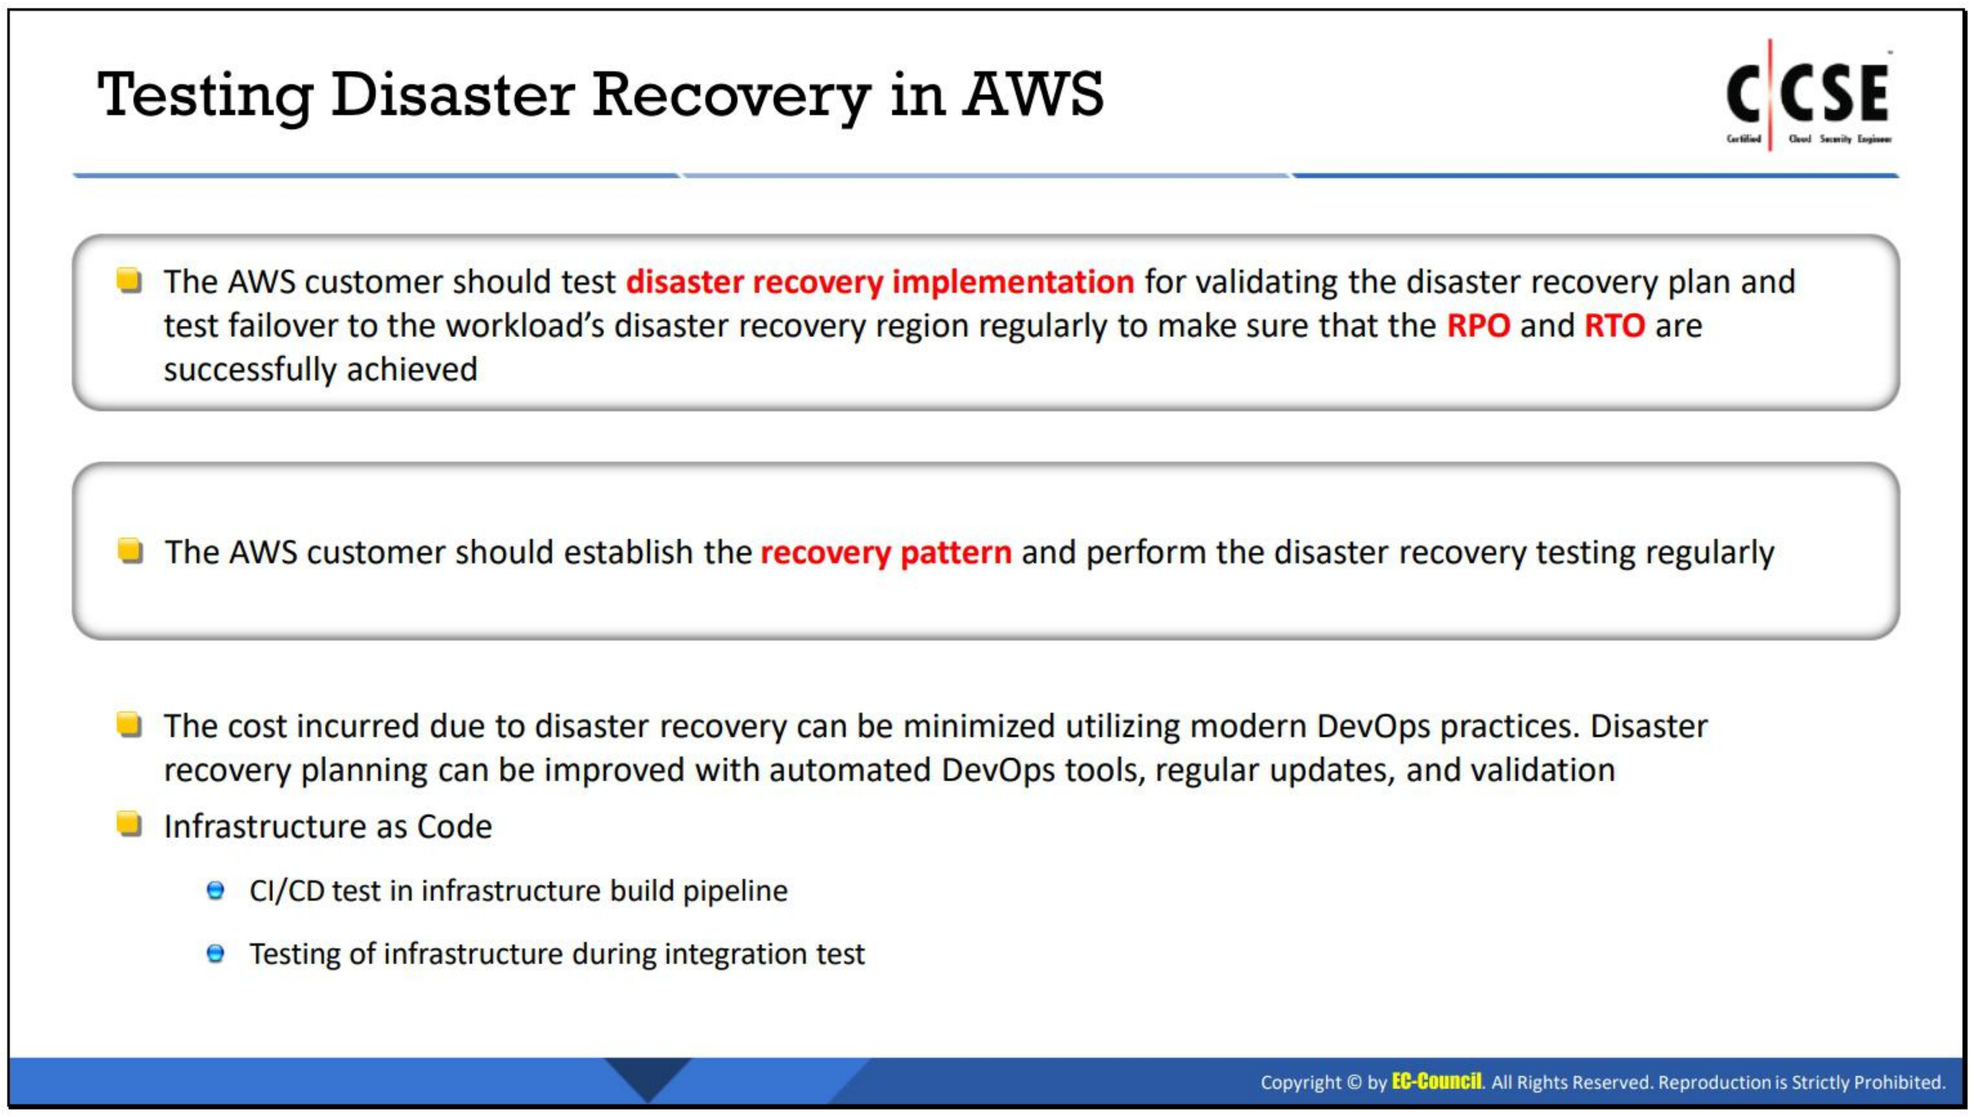

Testing Disaster Recovery in AWS

The AWS customer should test disaster recovery implementation for validating the DRP and test failover to the workload's disaster recovery region regularly to make sure that the RPO and RTO are successfully achieved.

The recovery paths that are rarely executed should be avoided. For example, the AWS customer has a secondary data store, which is utilized for read-only queries. If the customer writes to the primary data store and it fails, then the customer would like to failover to the secondary data store. If the customer did not regularly test the failover, then they may find that the expectations regarding the capabilities of the secondary data store are not up to the mark. The secondary data store capacity might be sufficient when it was tested last time; however, under the influence of the disaster event, it would not be able to handle the load.

Hence, the path that was tested regularly will have a high probability of recovering from a disaster event. It is recommended to have few recovery paths.

The AWS customer should establish the recovery pattern and perform the disaster recovery testing regularly. If there are recovery paths that are complex or critical, the AWS customer still needs to perform the testing regularly to validate that the path is working properly.

At the disaster recovery region, the AWS customer should manage the configuration drift. The customer should ensure that the infrastructure, configuration, and data as per requirement without any change or modification in the disaster recovery region. For example, verify that the service quotas and AMIs are updated.

The AWS customer can use AWS Config for monitoring and recording the AWS resource configuration continuously. The drift can be easily detected with AWS Config. After detecting the

Certified Cloud Security Engineer Copyright © by **EC-Council**
All Rights Reserved. Reproduction is Strictly Prohibited.

drift, AWS Config prompts the AWS system Manager Automation to repair drift and trigger alarm. The AWS customer can also use AWS CloudFormation for detecting drift in the deployed stacks.

hide01.ir

Certified Cloud Security Engineer Copyright © by **EC-Council**
All Rights Reserved. Reproduction is Strictly Prohibited.

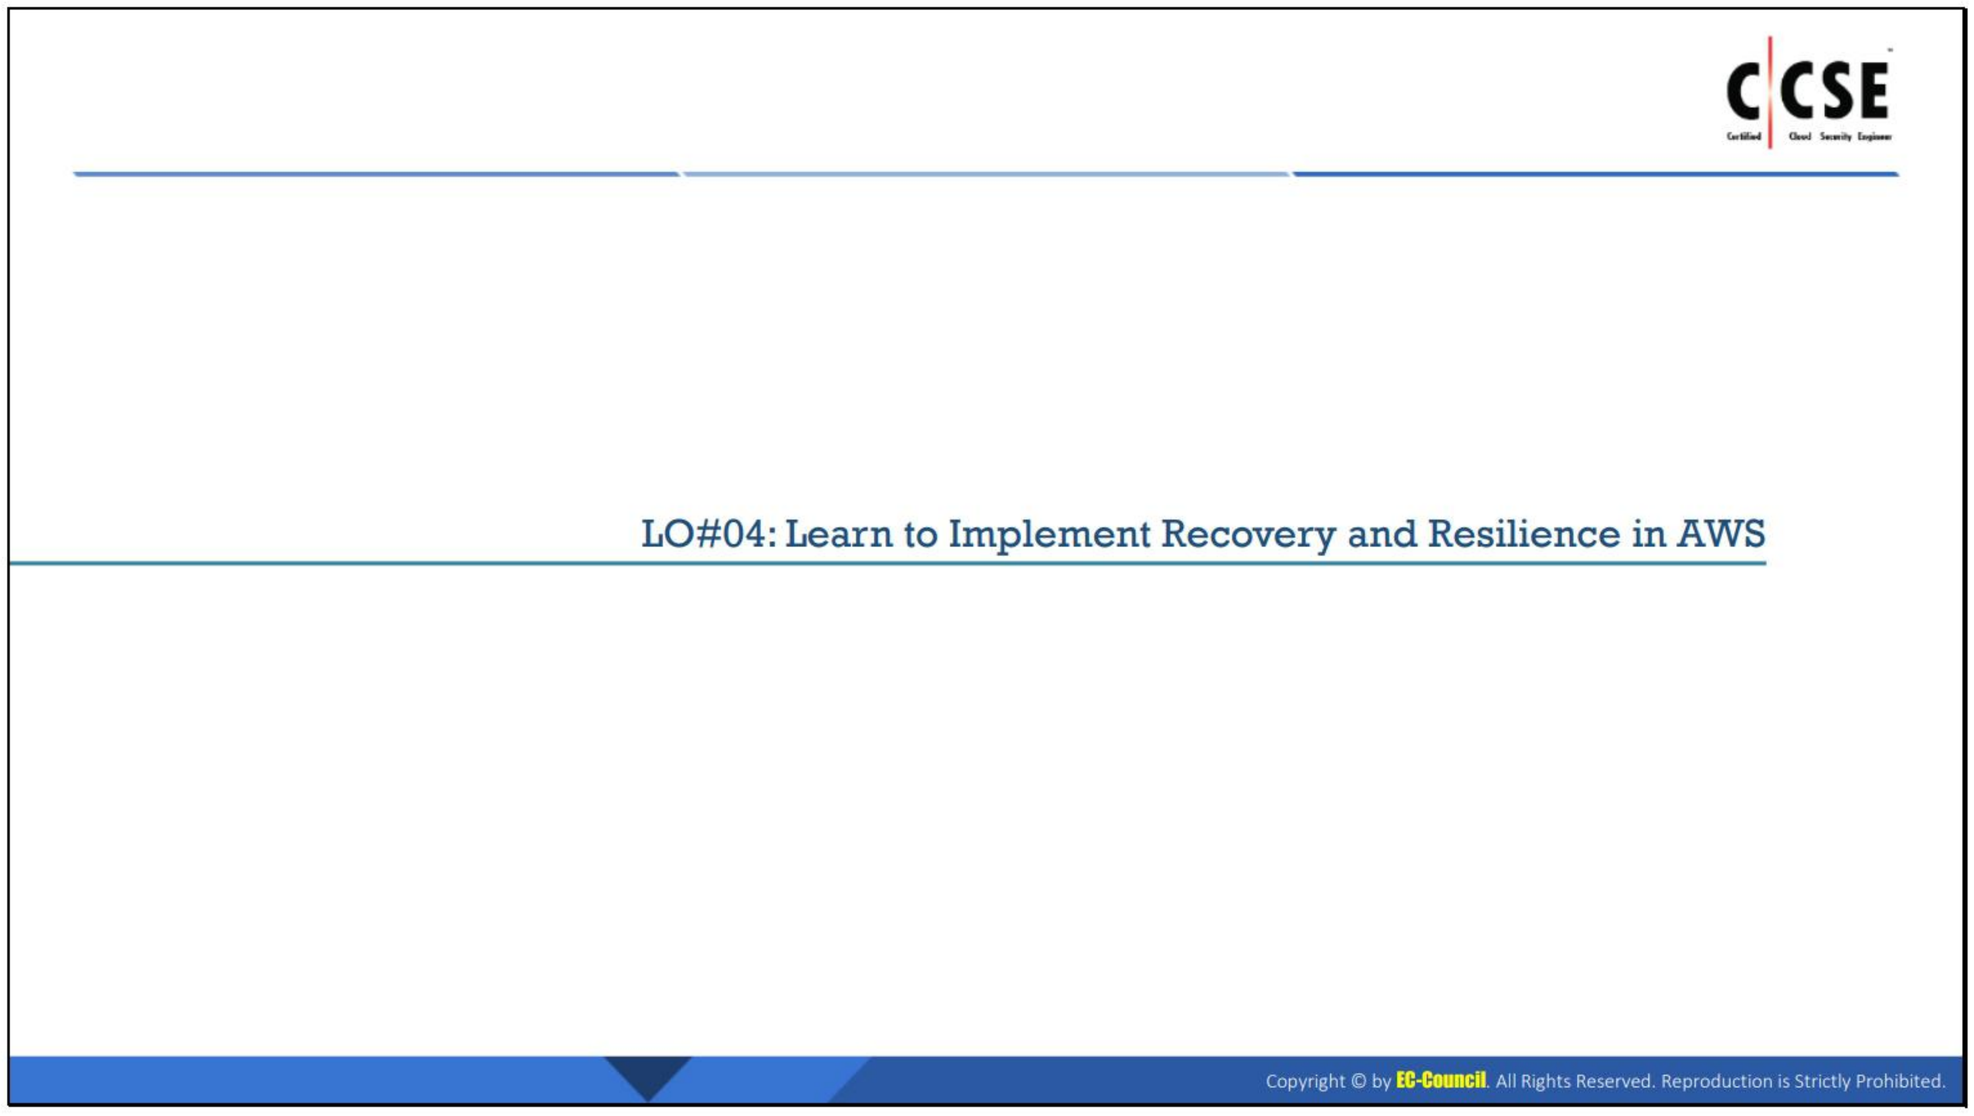

LO#04: Learn Implementing Recovery and Resilience in AWS

The objective of this section is to understand how a DR plan can be implemented on AWS. It helps cloud security engineers to learn the various ways they can use to backup and restore data and services using AWS options. In addition, the cloud security engineers will become well-versed with AWS partner network (APN) and BCDR tools such as Zerto, Commvault, etc.

Certified Cloud Security Engineer Copyright © by **EC-Council**
All Rights Reserved. Reproduction is Strictly Prohibited.

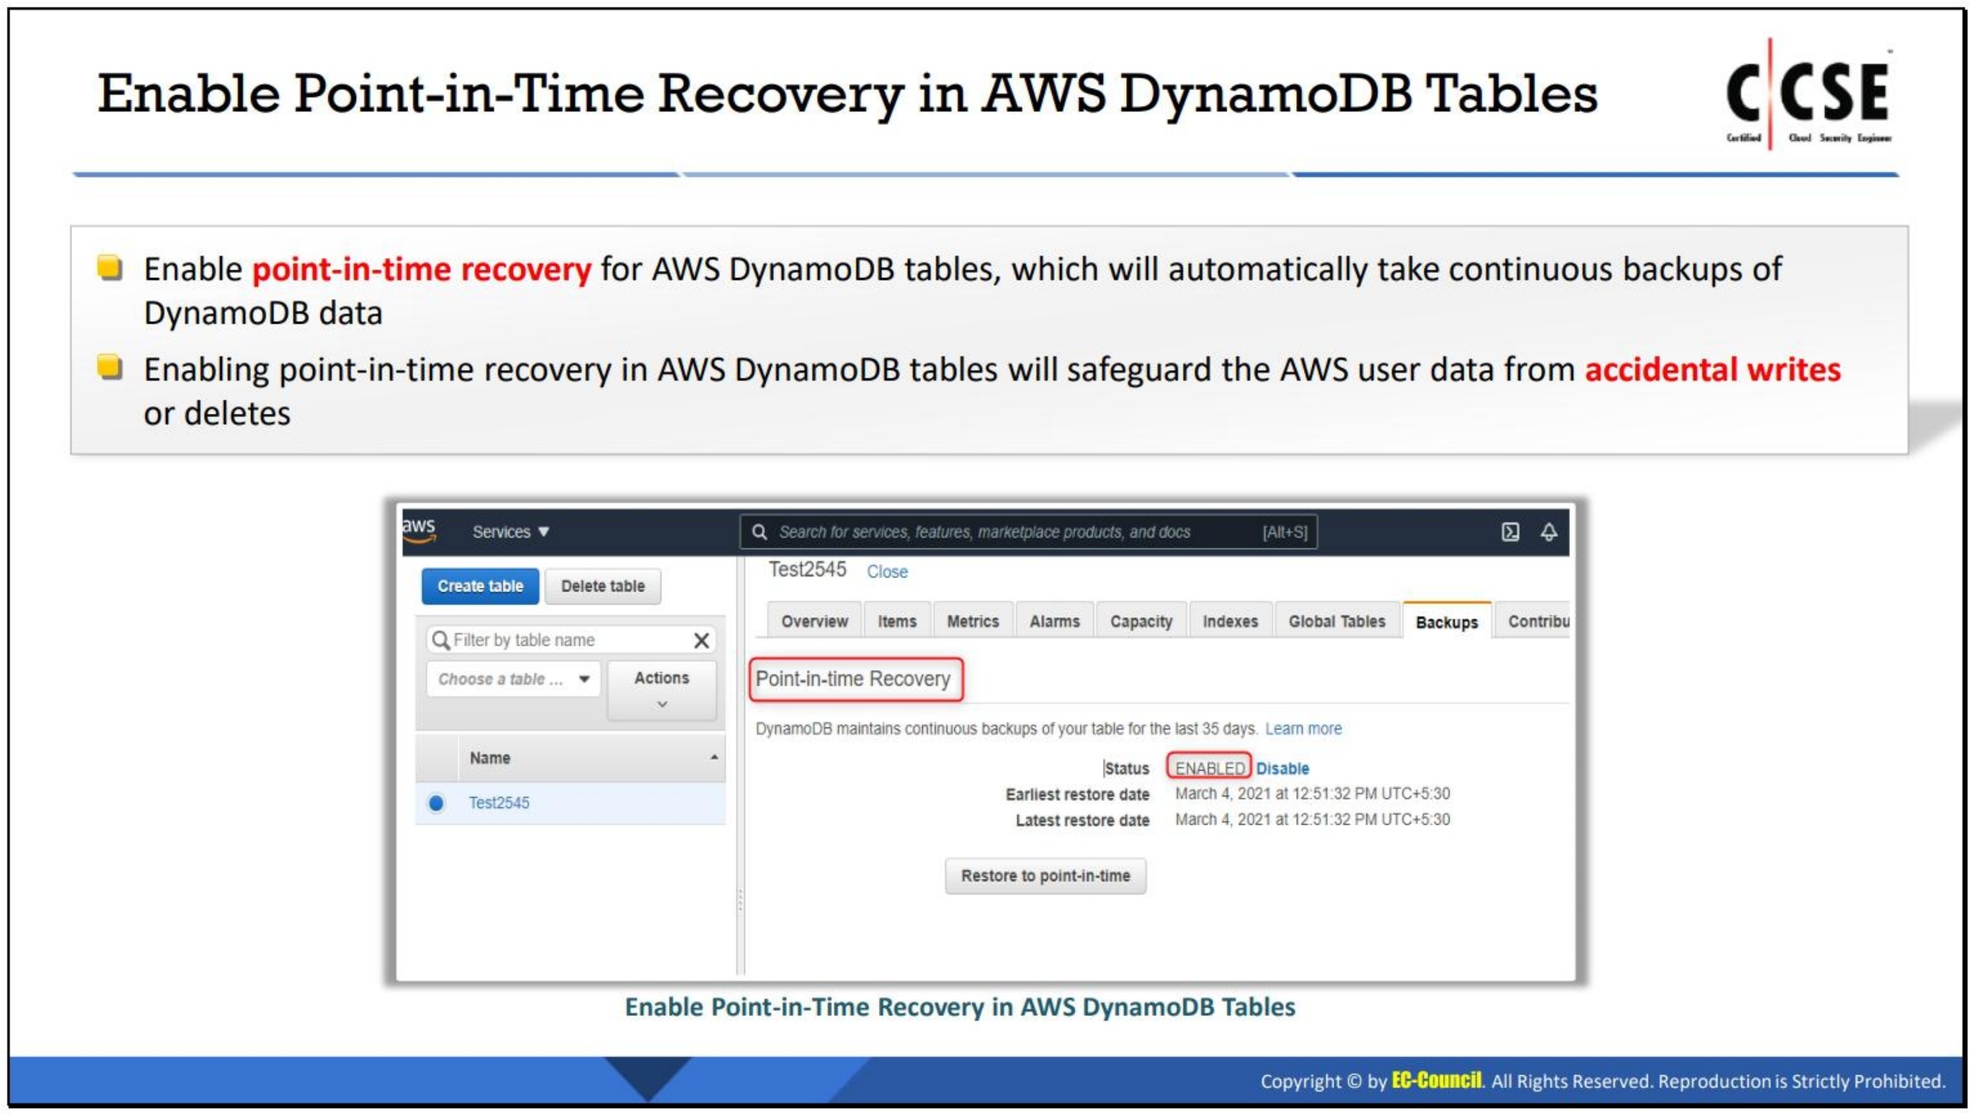

Enable Point-in-Time Recovery in AWS DynamoDB Tables

Enable Point-in-Time Recovery in AWS DynamoDB Tables

The AWS cloud security professional should enable PITR for AWS DynamoDB tables, which will take continuous backups of DynamoDB data automatically. AWS DynamoDB service will be used to take the data backup with per-secondary granularity and allow the administrator to use PITR to restore the table to the desired point in time within the last 35 days. Enabling PITR in AWS DynamoDB tables will safeguard the AWS user data from accidental writes or deletes.

To Enable Point-in-time Recovery in AWS DynamoDB Tables

- Sign in to AWS Management Console.

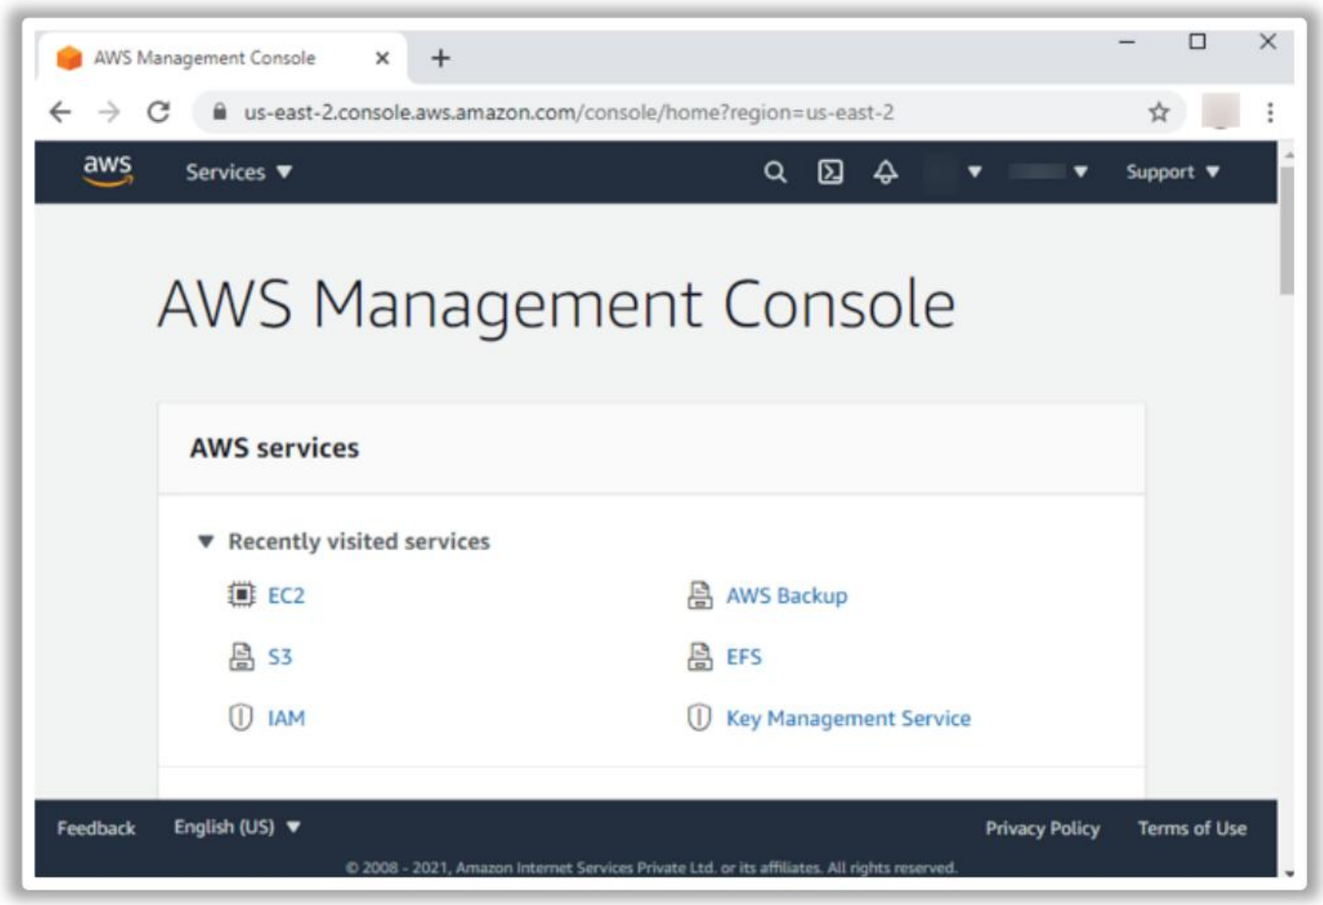

Figure 9.24: AWS Management Console

Certified Cloud Security Engineer Copyright © by **EC-Council**
All Rights Reserved. Reproduction is Strictly Prohibited.

- Click on the following link to open DynamoDB: https://aws.amazon.com/dynamodb/getting-started/

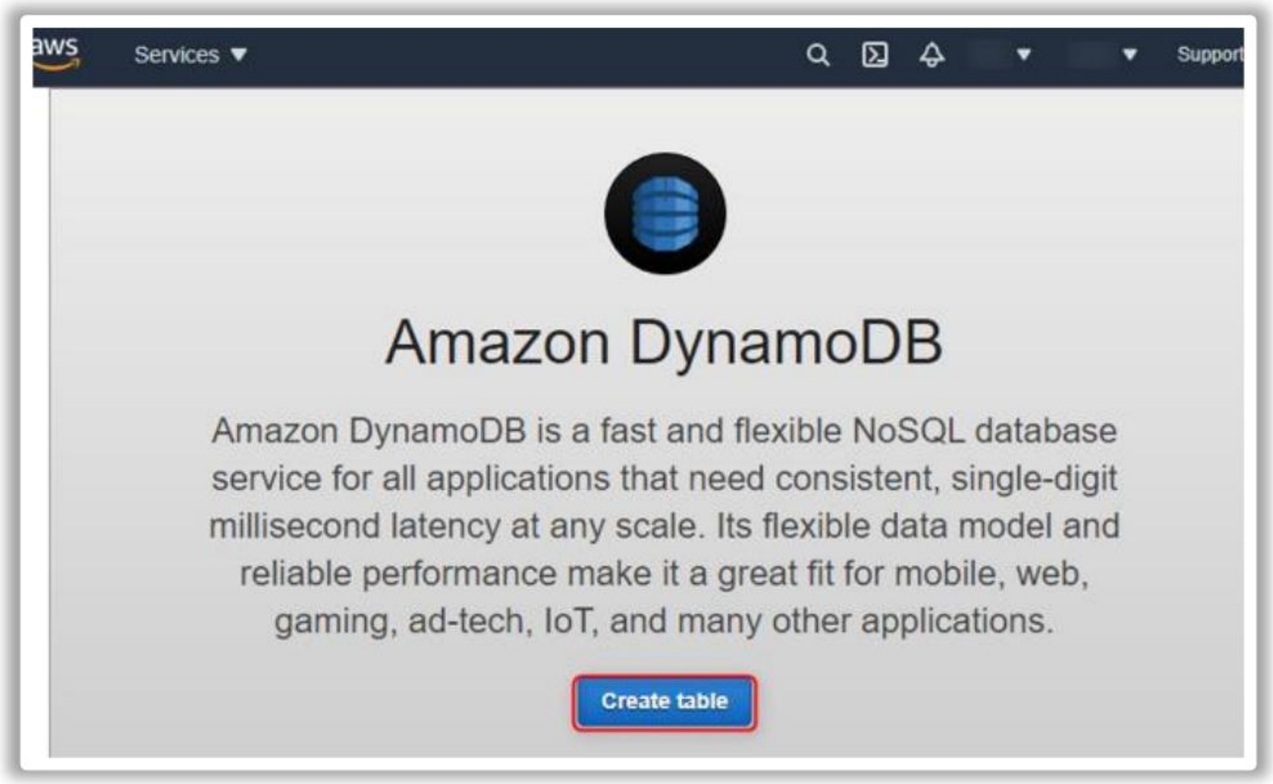

Figure 9.25: Amazon DynamoDB

- Click on Create table button, enter the Table name, Primary key, and choose the Primary key from the drop-down; then click on the Create button.

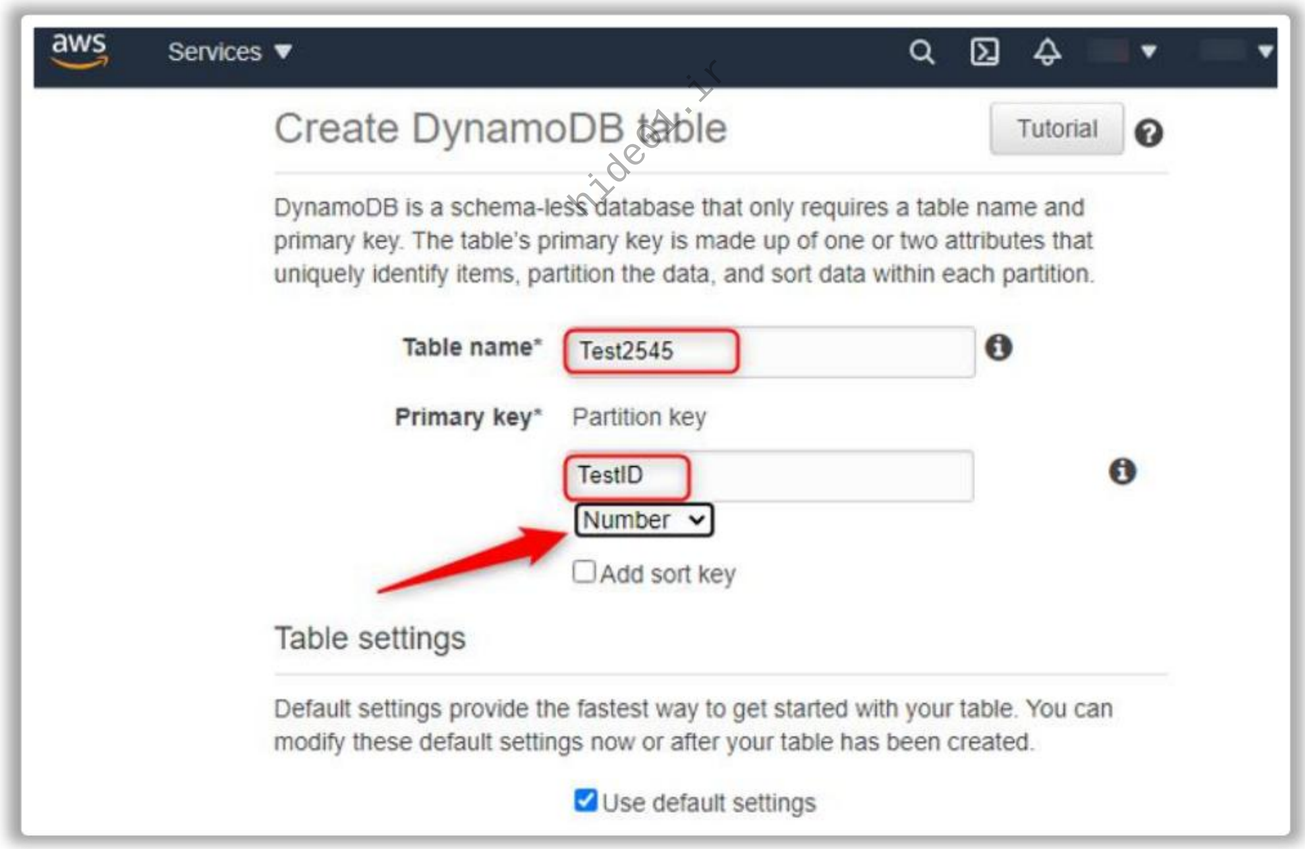

Figure 9.26: Create DynamoDB table Console

Certified Cloud Security Engineer Copyright © by **EC-Council**
All Rights Reserved. Reproduction is Strictly Prohibited.

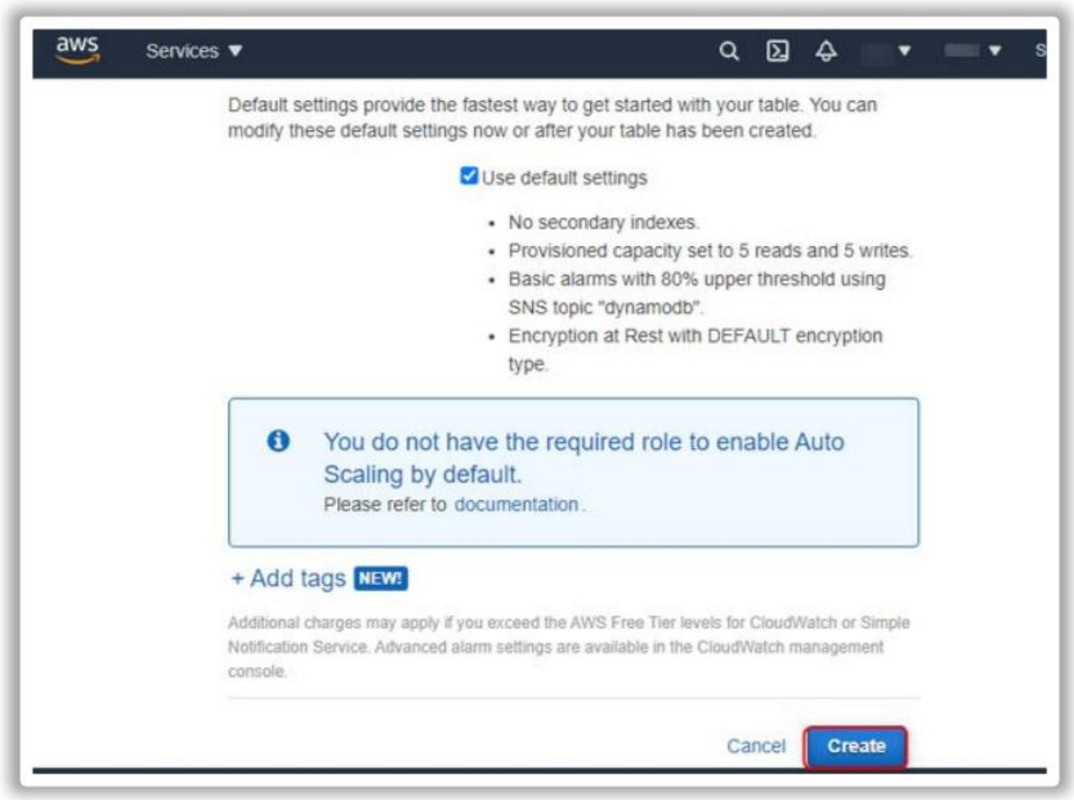

Figure 9.27: Creating DynamoDB Table

- Click on the **items** tab and then click on **Create item** button.

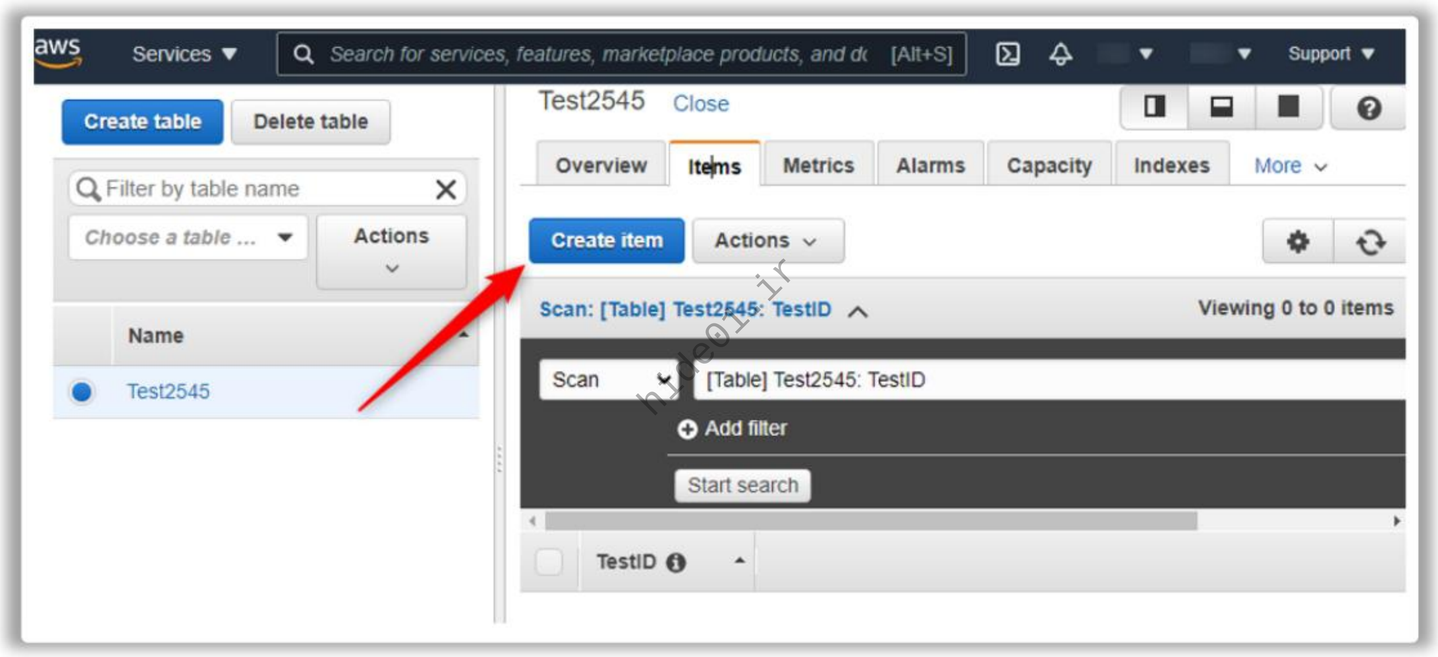

Figure 9.28: Create Table Console

- Fill in the data as per requirement in the table and click on the **Save** button.

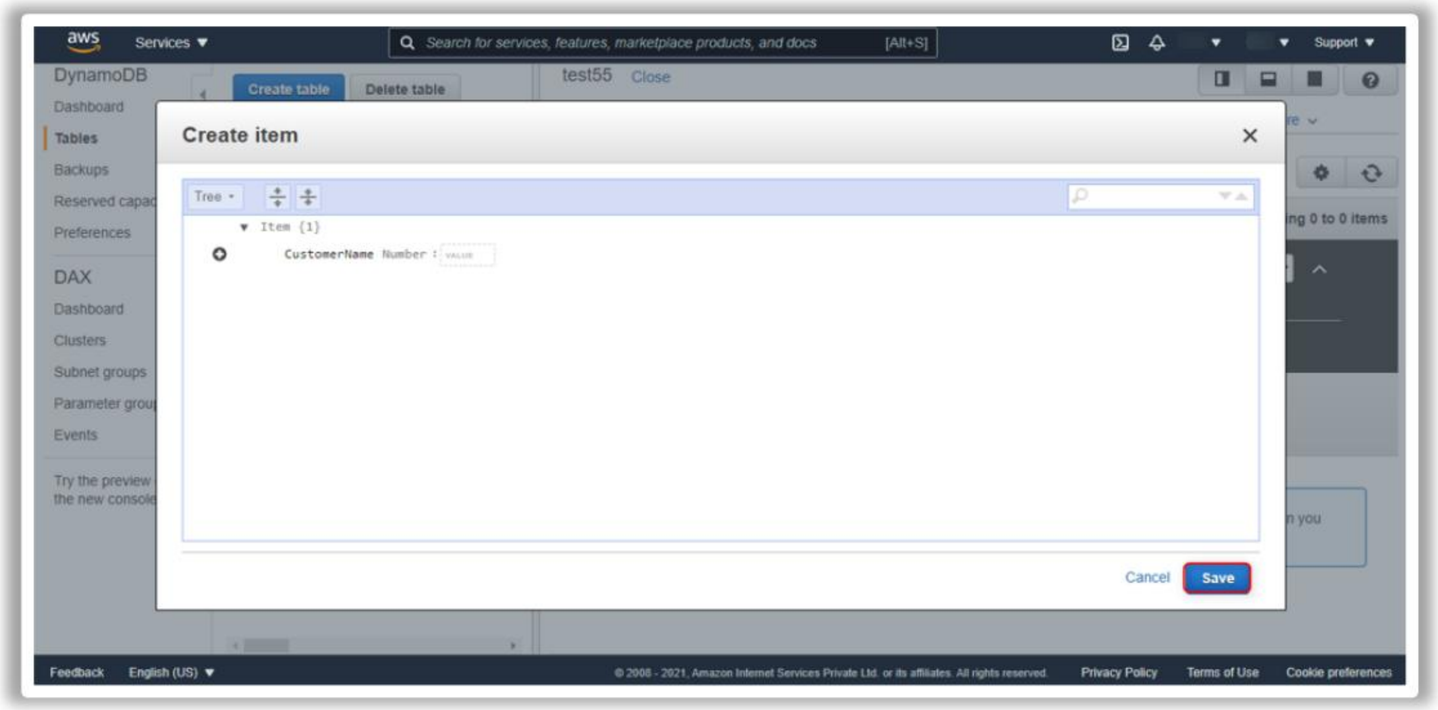

Figure 9.29: Creating item Console

Certified Cloud Security Engineer Copyright © by **EC-Council**
All Rights Reserved. Reproduction is Strictly Prohibited.

- Similarly, you can add as many rows as you can, then click on the Backups tab.

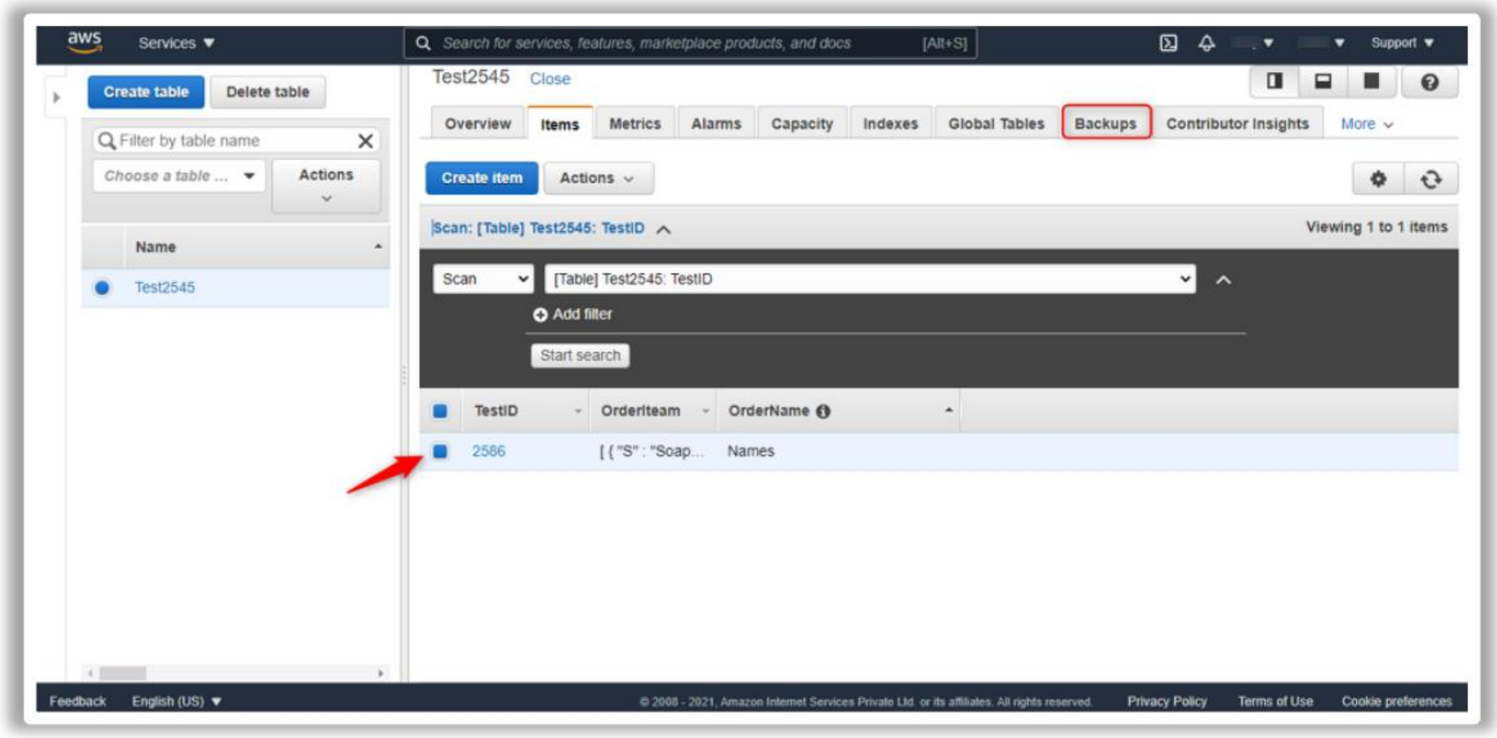

Figure 9.30: Creating table Console

- Under the Point-in-Time Recovery section, click on **Enable** button. Then, **Enable Point-in-Time Recovery** pop will be displayed, click again on **Enable** button.

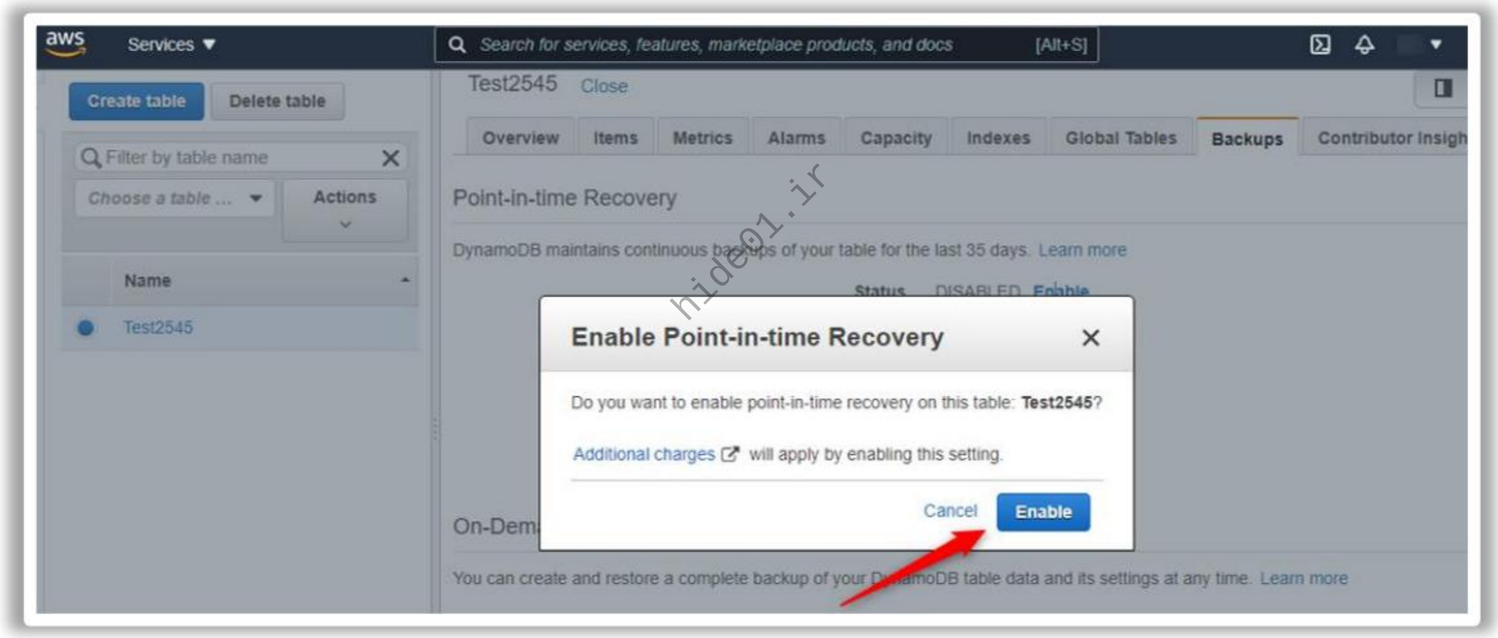

Figure 9.31: Enable Point-in-time Recovery pop-up

- Point-in-Time Recovery for continuous data backup for Amazon DynamoDB is successfully enabled.

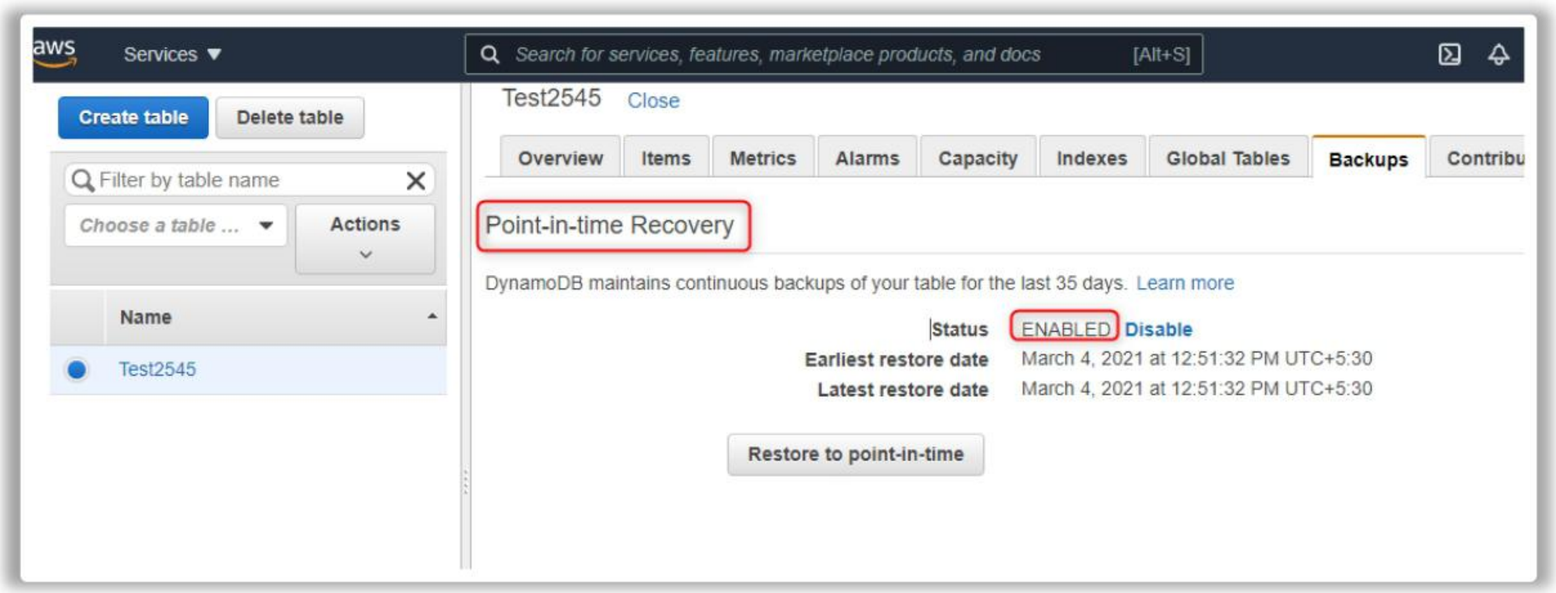

Figure 9.32: Point-in-time Console

Certified Cloud Security Engineer Copyright © by **EC-Council**
All Rights Reserved. Reproduction is Strictly Prohibited.

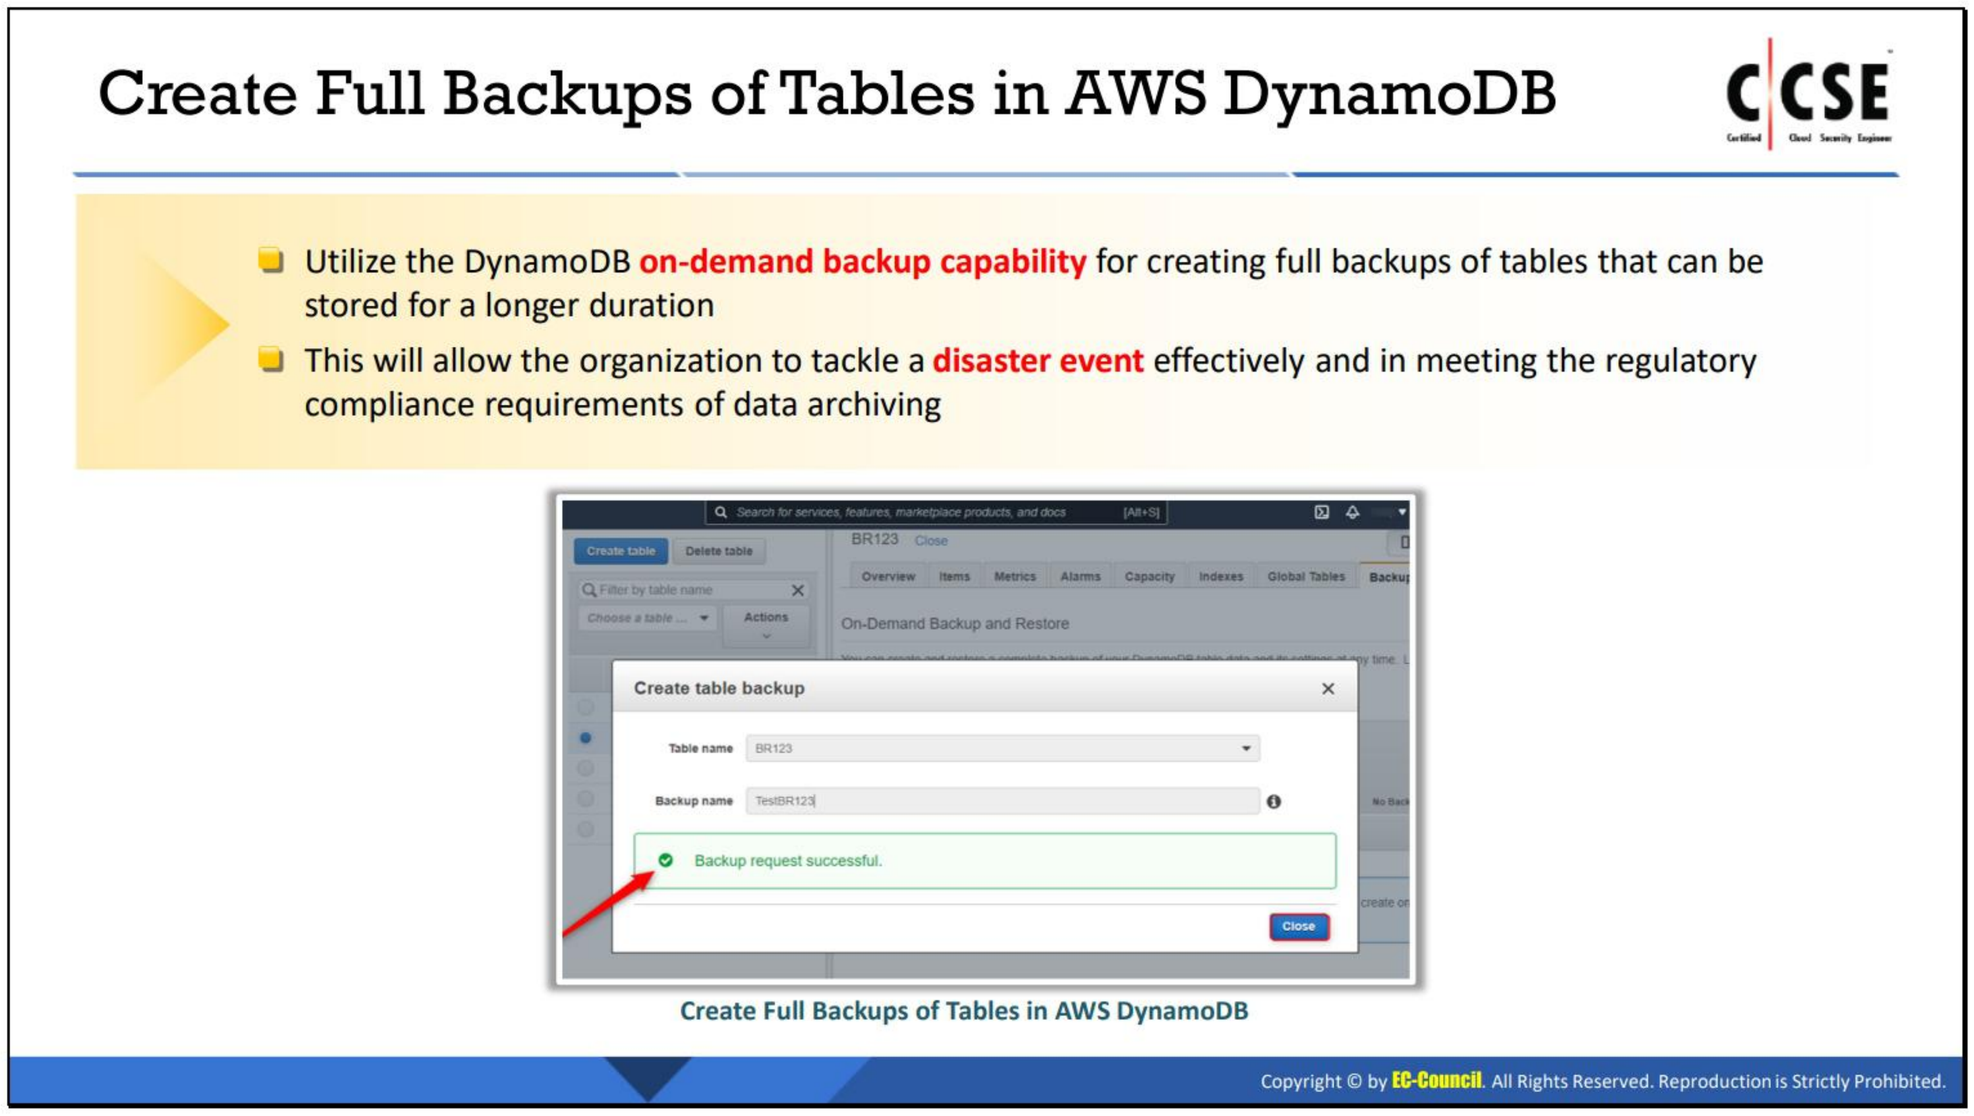

Create Full Backups of Tables in AWS DynamoDB

The cloud security professional can utilize DynamoDB's on-demand backup capability for creating full backups of tables that can be stored for a longer duration. This will allow the organization to tackle a disaster event effectively and in meeting the regulatory compliance requirements of data archiving. The cloud administrator can backup and store the table data at any point in time. The on-demand backup process can scale without affecting the availability or performance of an application. It utilizes a distributed technology that is new and unique, which allows the administrator to complete the backup process within a short time. An on-demand backup will be preserved even though the respective table is deleted.

To Create Full Backups of Tables in AWS DynamoDB

- Sign in to AWS Management Console.

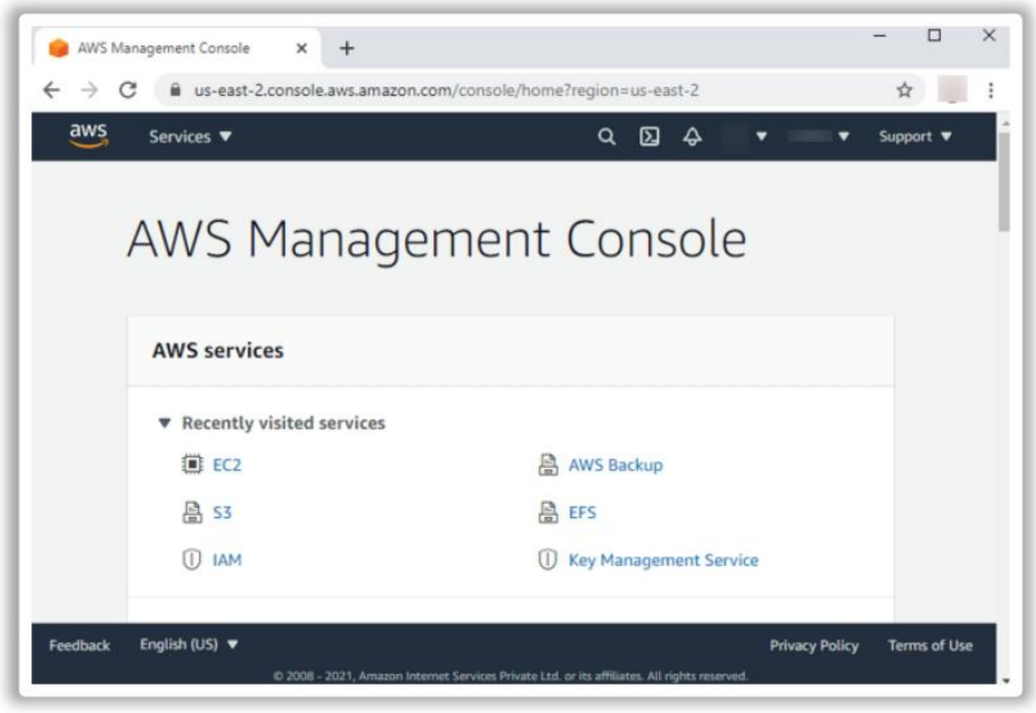

Figure 9.33: AWS Management Console

Certified Cloud Security Engineer Copyright © by **EC-Council**
All Rights Reserved. Reproduction is Strictly Prohibited.

- Click on **DynamoDB.**

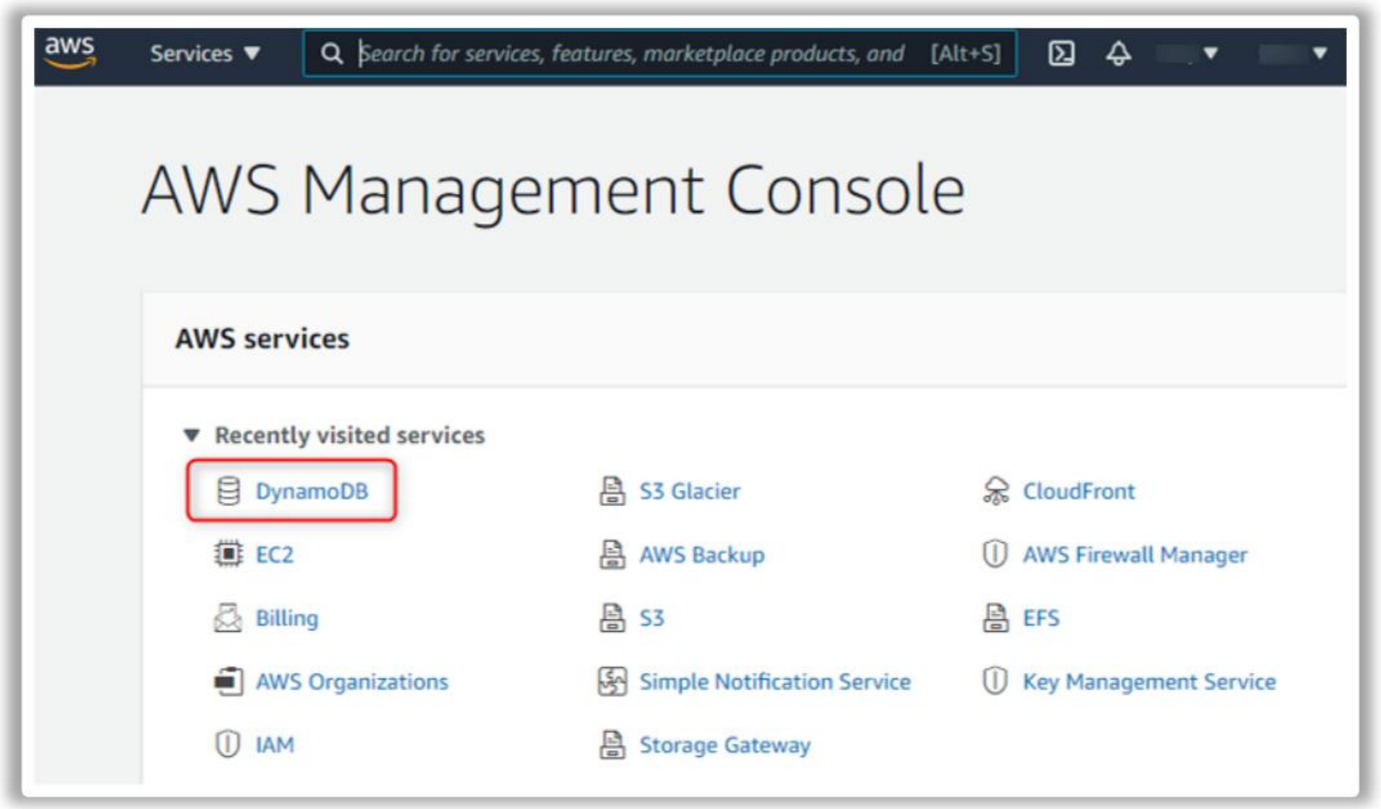

Figure 9.34: AWS Management Console

- You will be redirected to the DynamoDB dashboard, click **Tables** on the left pane. You will find the existing tables. Select the table to create full backups.

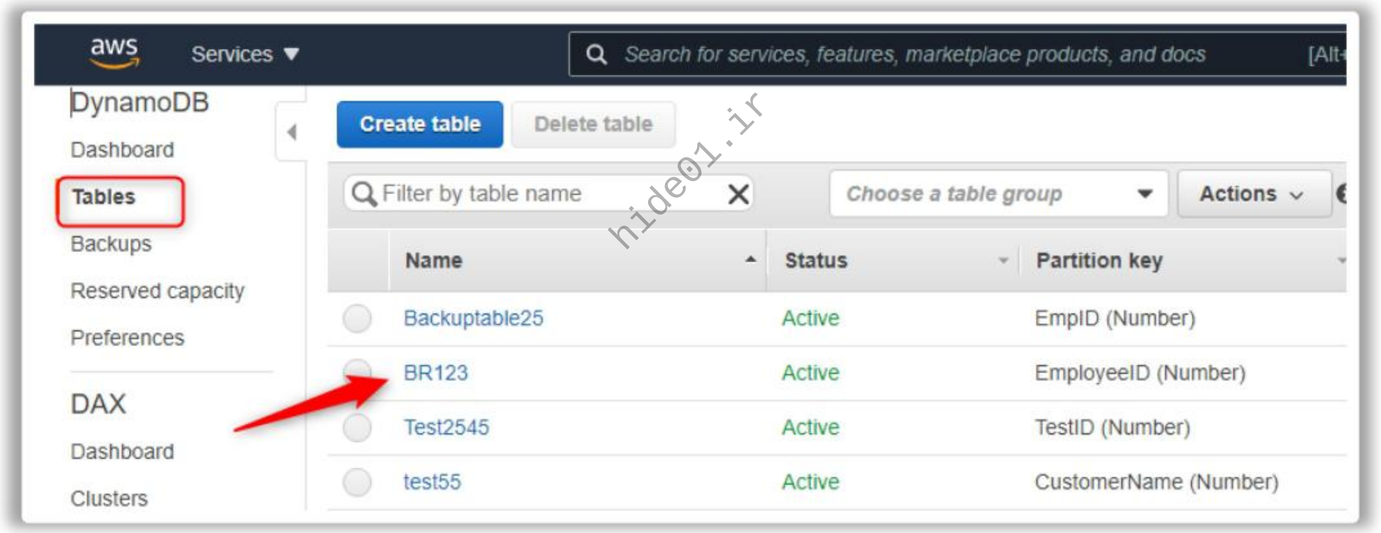

Figure 9.35: Create Table Console

- Click on **Backups** tab, then under **On-Demand Backup and Restore**, click on **Create backup** button.

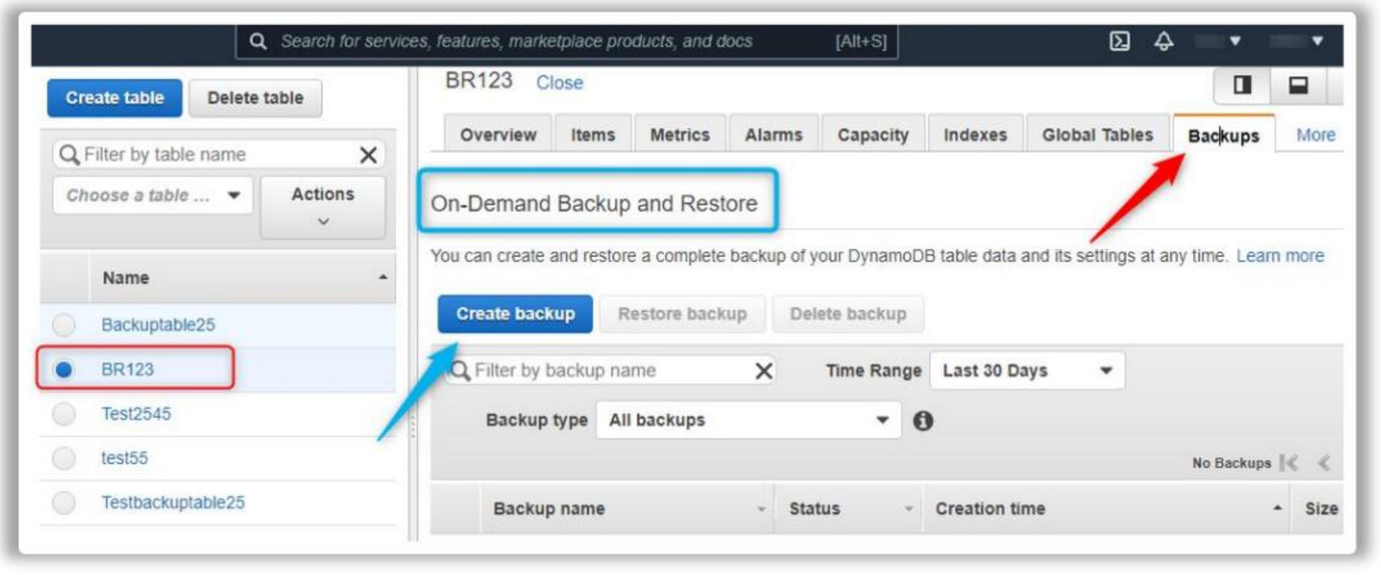

Figure 9.36: Create backup Console

Certified Cloud Security Engineer Copyright © by **EC-Council**
All Rights Reserved. Reproduction is Strictly Prohibited.

- In the **Backup name** field, enter the name of the backup and click the **Create** button.

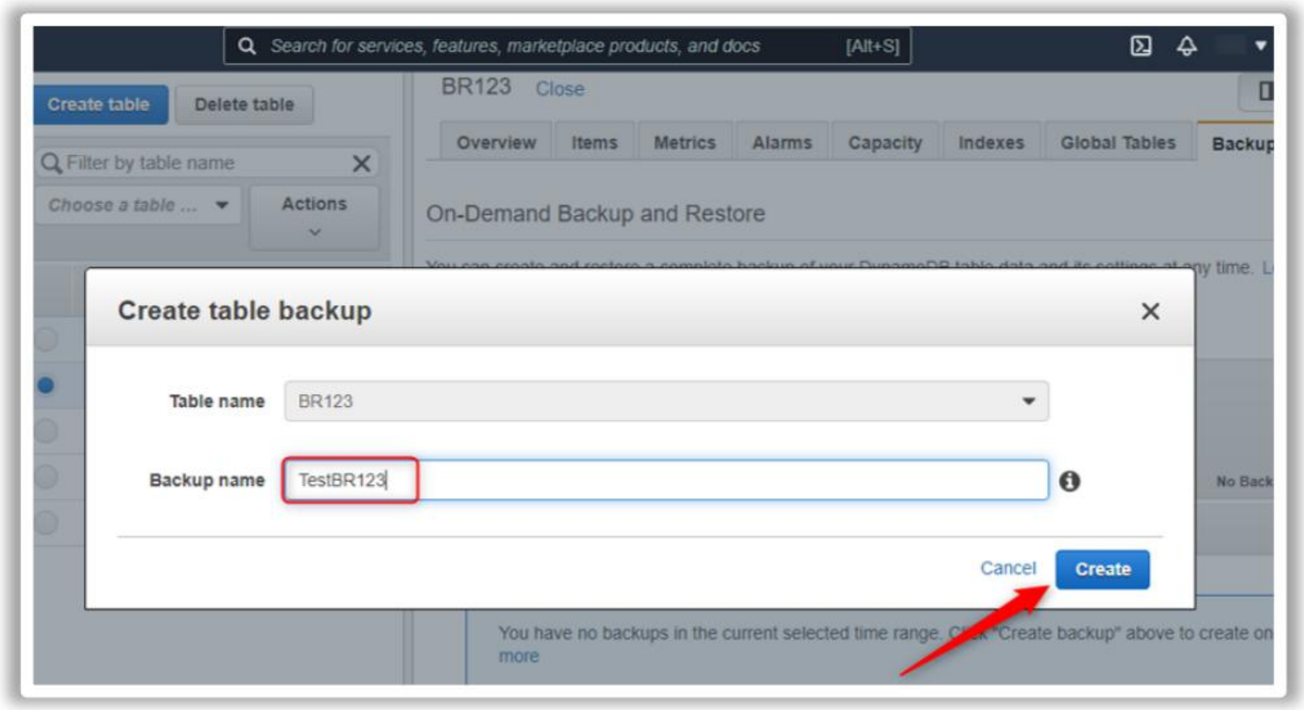

Figure 9.37: Create Table Backup Console

- Once on-demand backup is successfully created, then click the **Close** button.

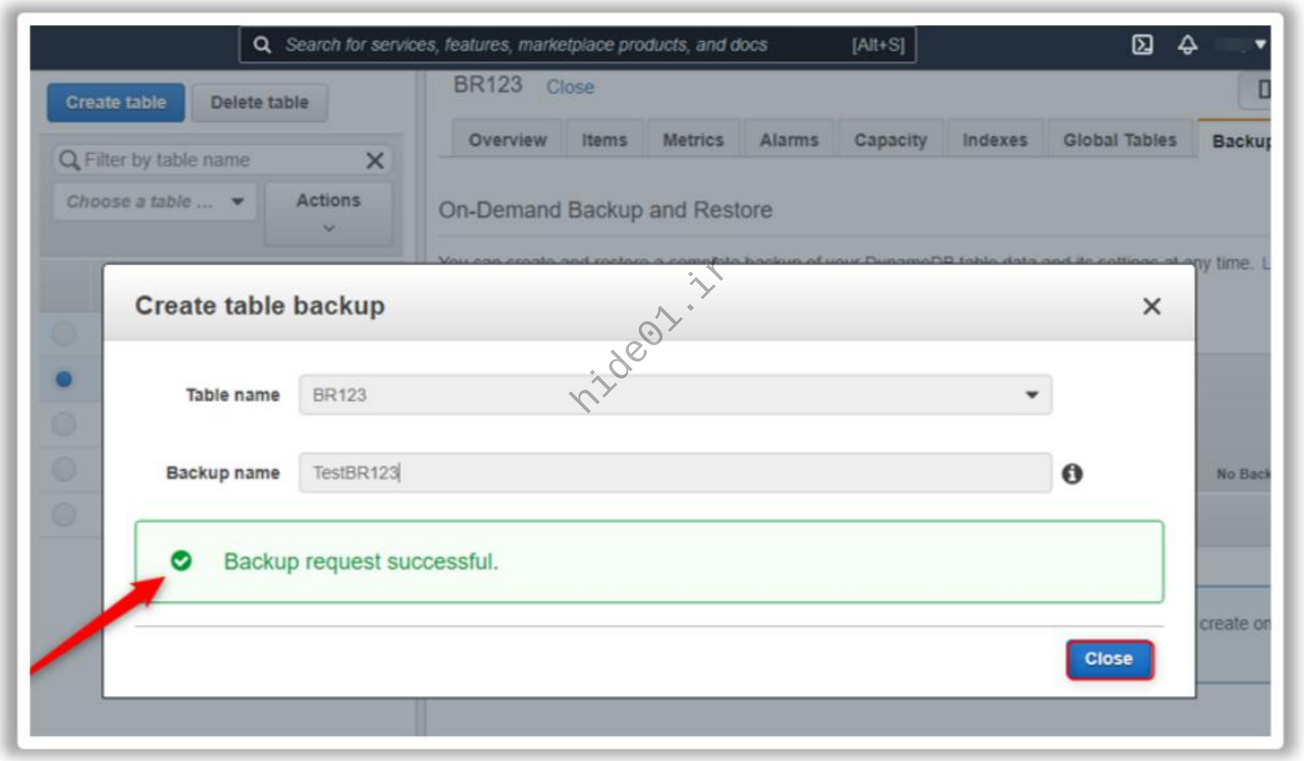

Figure 9.38: Create Table backup Console

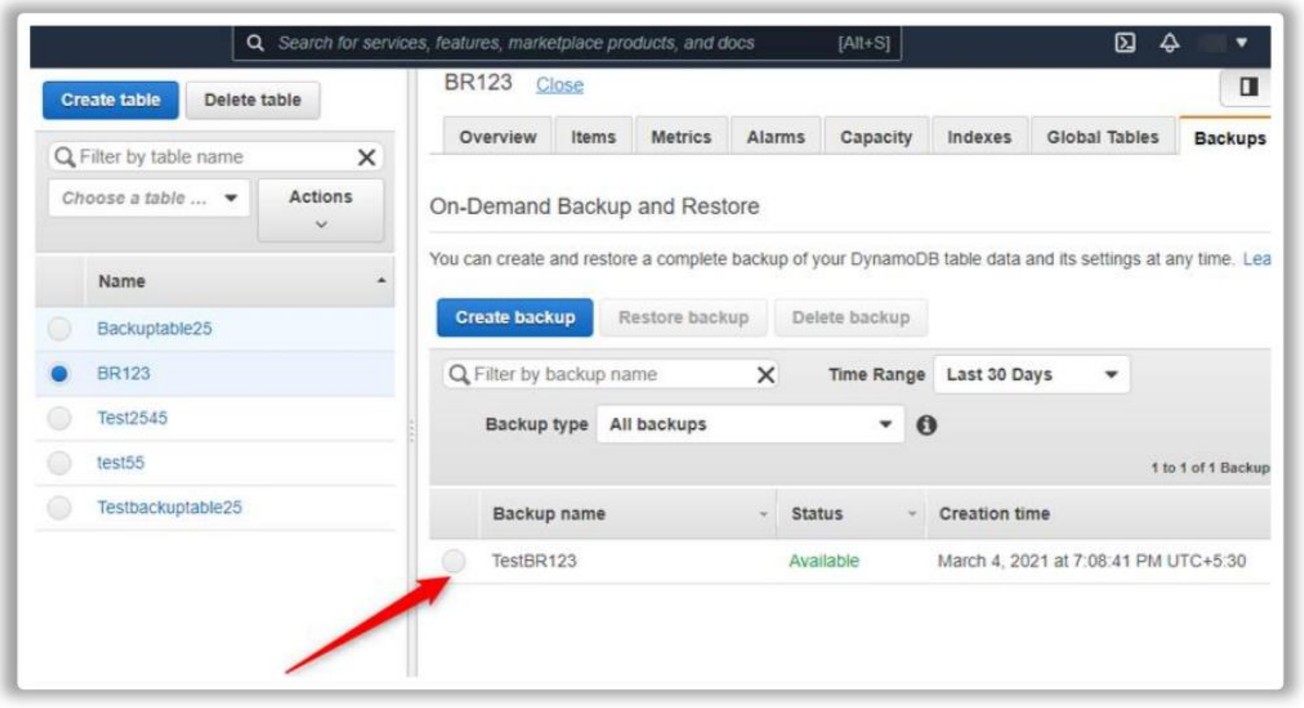

Figure 9.39: Table backup Console

Certified Cloud Security Engineer Copyright © by **EC-Council**
All Rights Reserved. Reproduction is Strictly Prohibited.

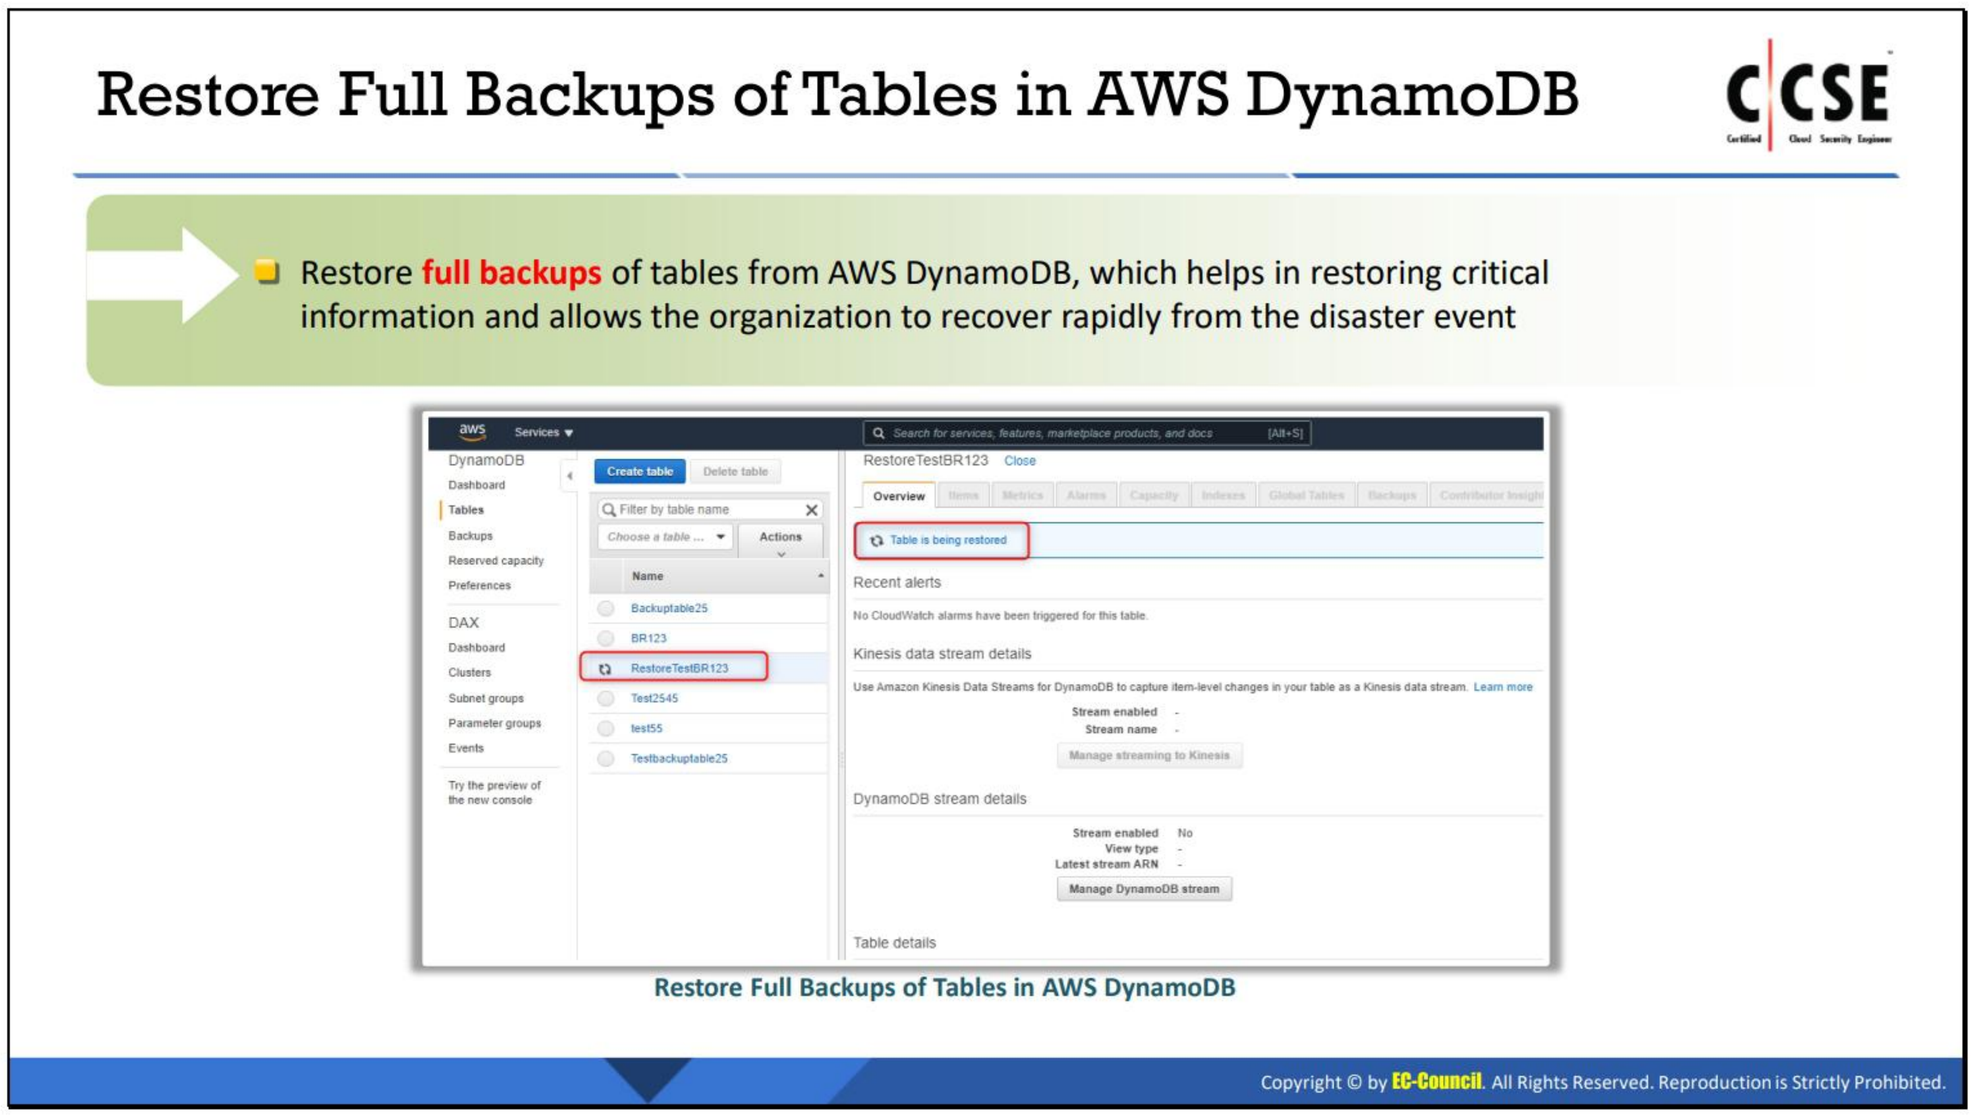

Restore Full Backups of Tables in AWS DynamoDB

Restore Full Backups of Tables in AWS DynamoDB

The cloud security professional can easily restore full backups of tables from AWS DynamoDB. Restoring full backups of tables in AWS DynamoDB would not affect the availability or performance of an application. The administrator can restore the data at any point in time. It helps in restoring critical information and allows the organization to recover rapidly from the disaster event.

Restore Full Backups of Tables in AWS DynamoDB

- Sign in to AWS Management Console.

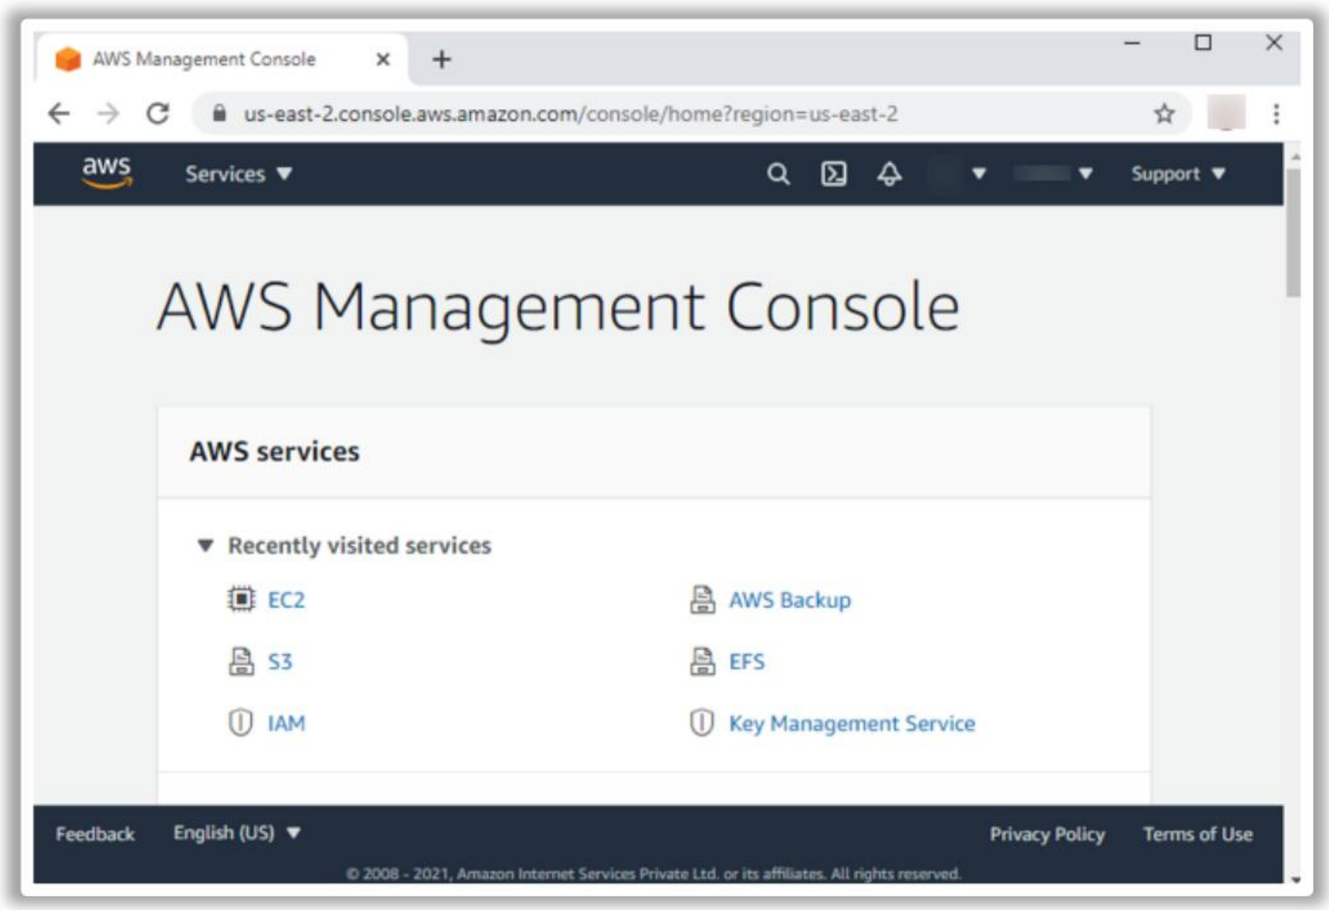

Figure 9.40: AWS Management Console

Certified Cloud Security Engineer Copyright © by **EC-Council**
All Rights Reserved. Reproduction is Strictly Prohibited.

- Click on **DynamoDB.**

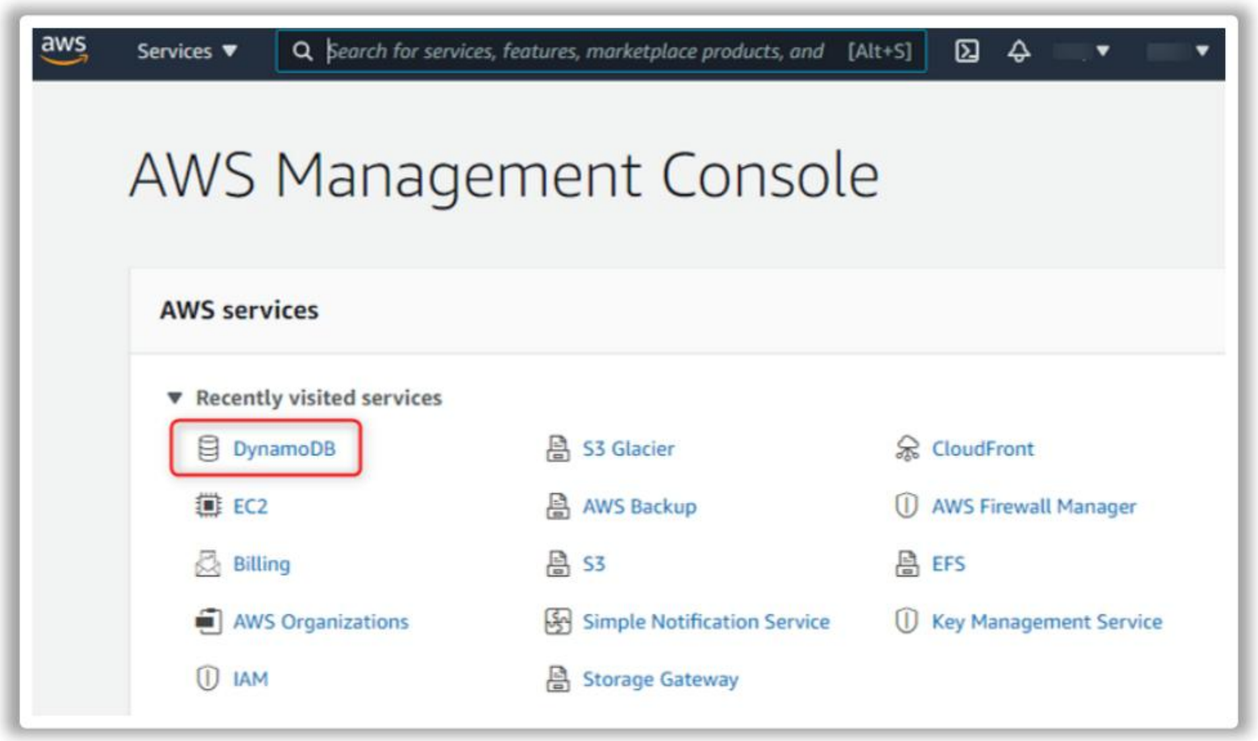

Figure 9.41: AWS Management Console

- You will redirect to the DynamoDB dashboard, click on **Tables** on the left pane. You will find the existing tables. Select the table to restore the full backup.

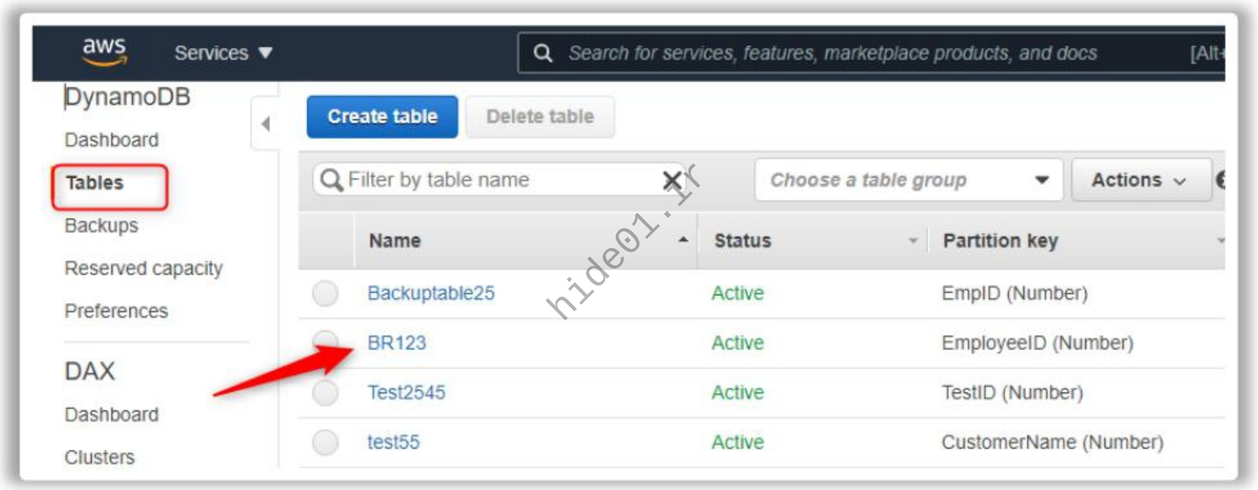

Figure 9.42: Create table Console

- Click on tab **Backups**, then under **On-Demand Backup and Restore**, select the table and click on the **Restore backup** button**.**

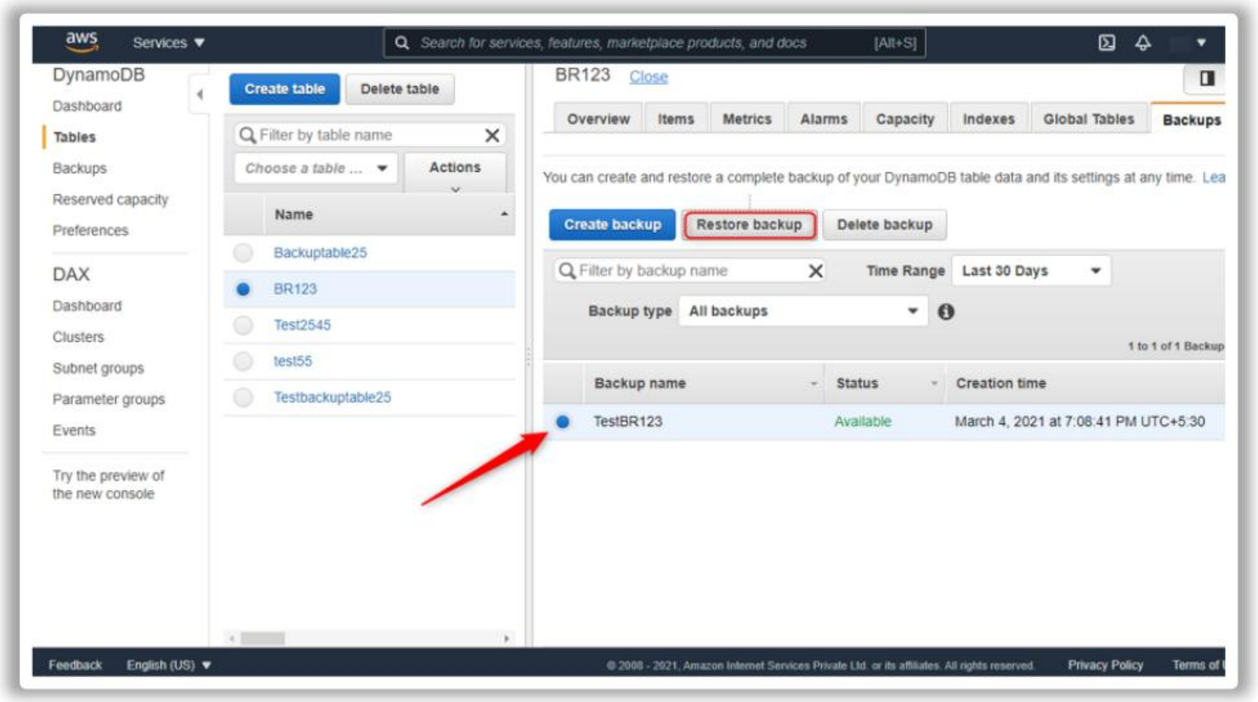

Figure 9.43: Restore Backup Console

Certified Cloud Security Engineer Copyright © by **EC-Council**
All Rights Reserved. Reproduction is Strictly Prohibited.

- In the **new table name field**, enter the name of the table. Navigate and click on the **Restore table** button.

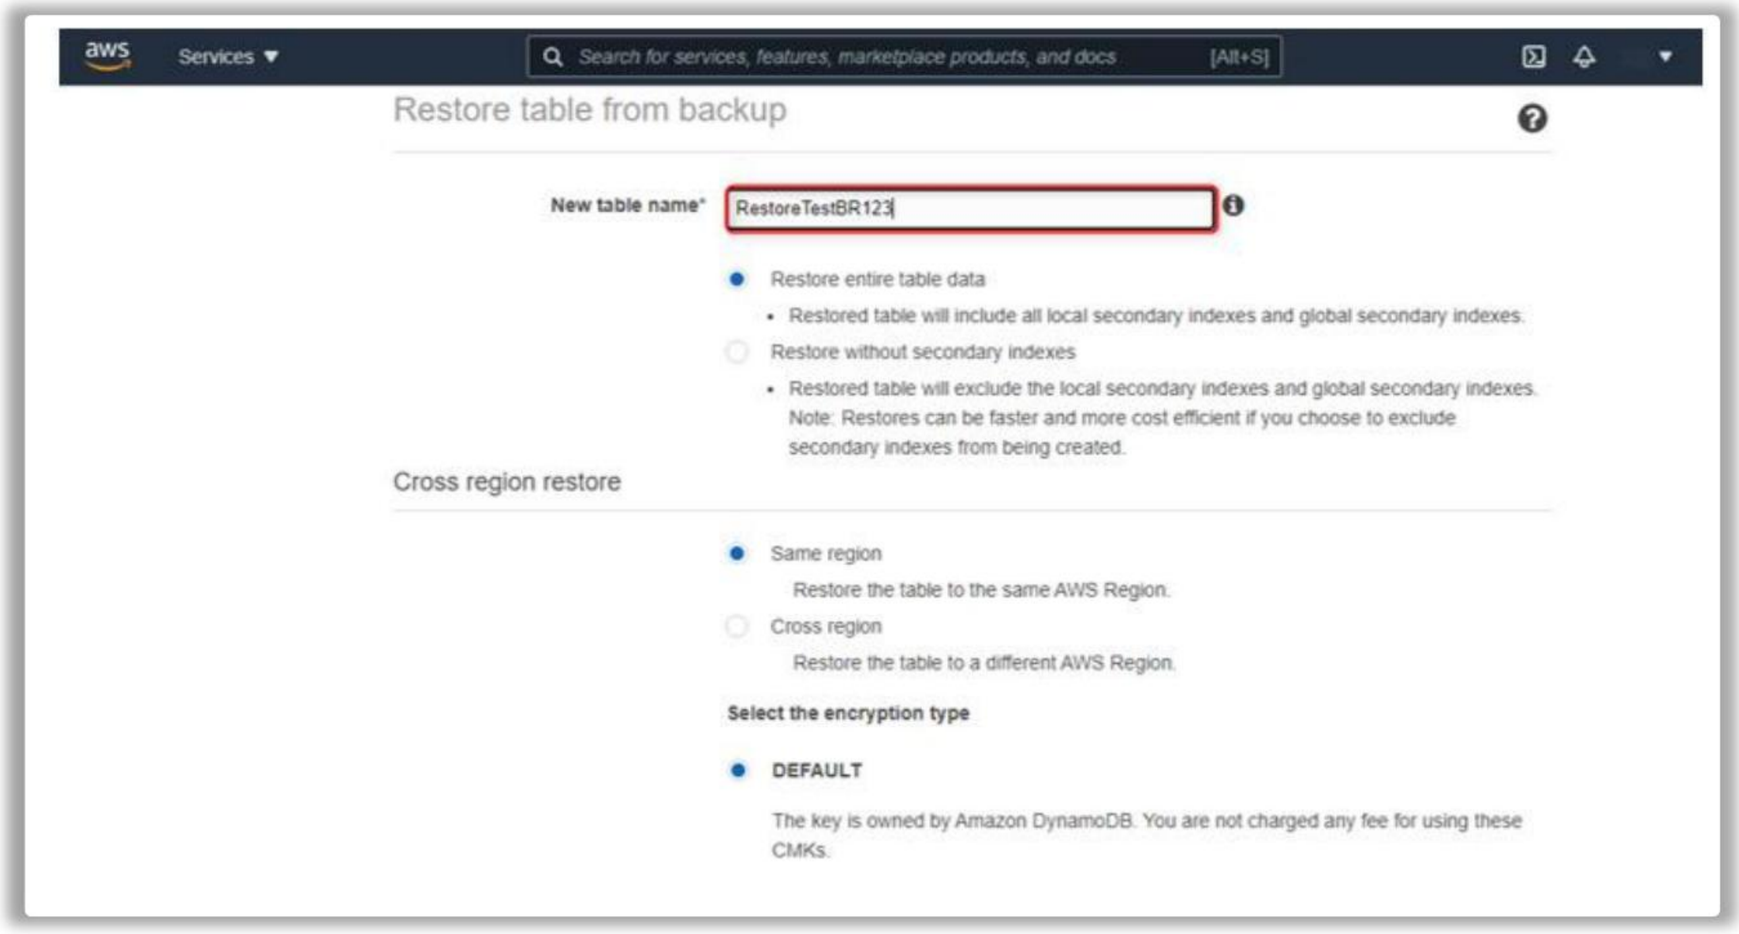

Figure 9.44: Restore table from Backup Console

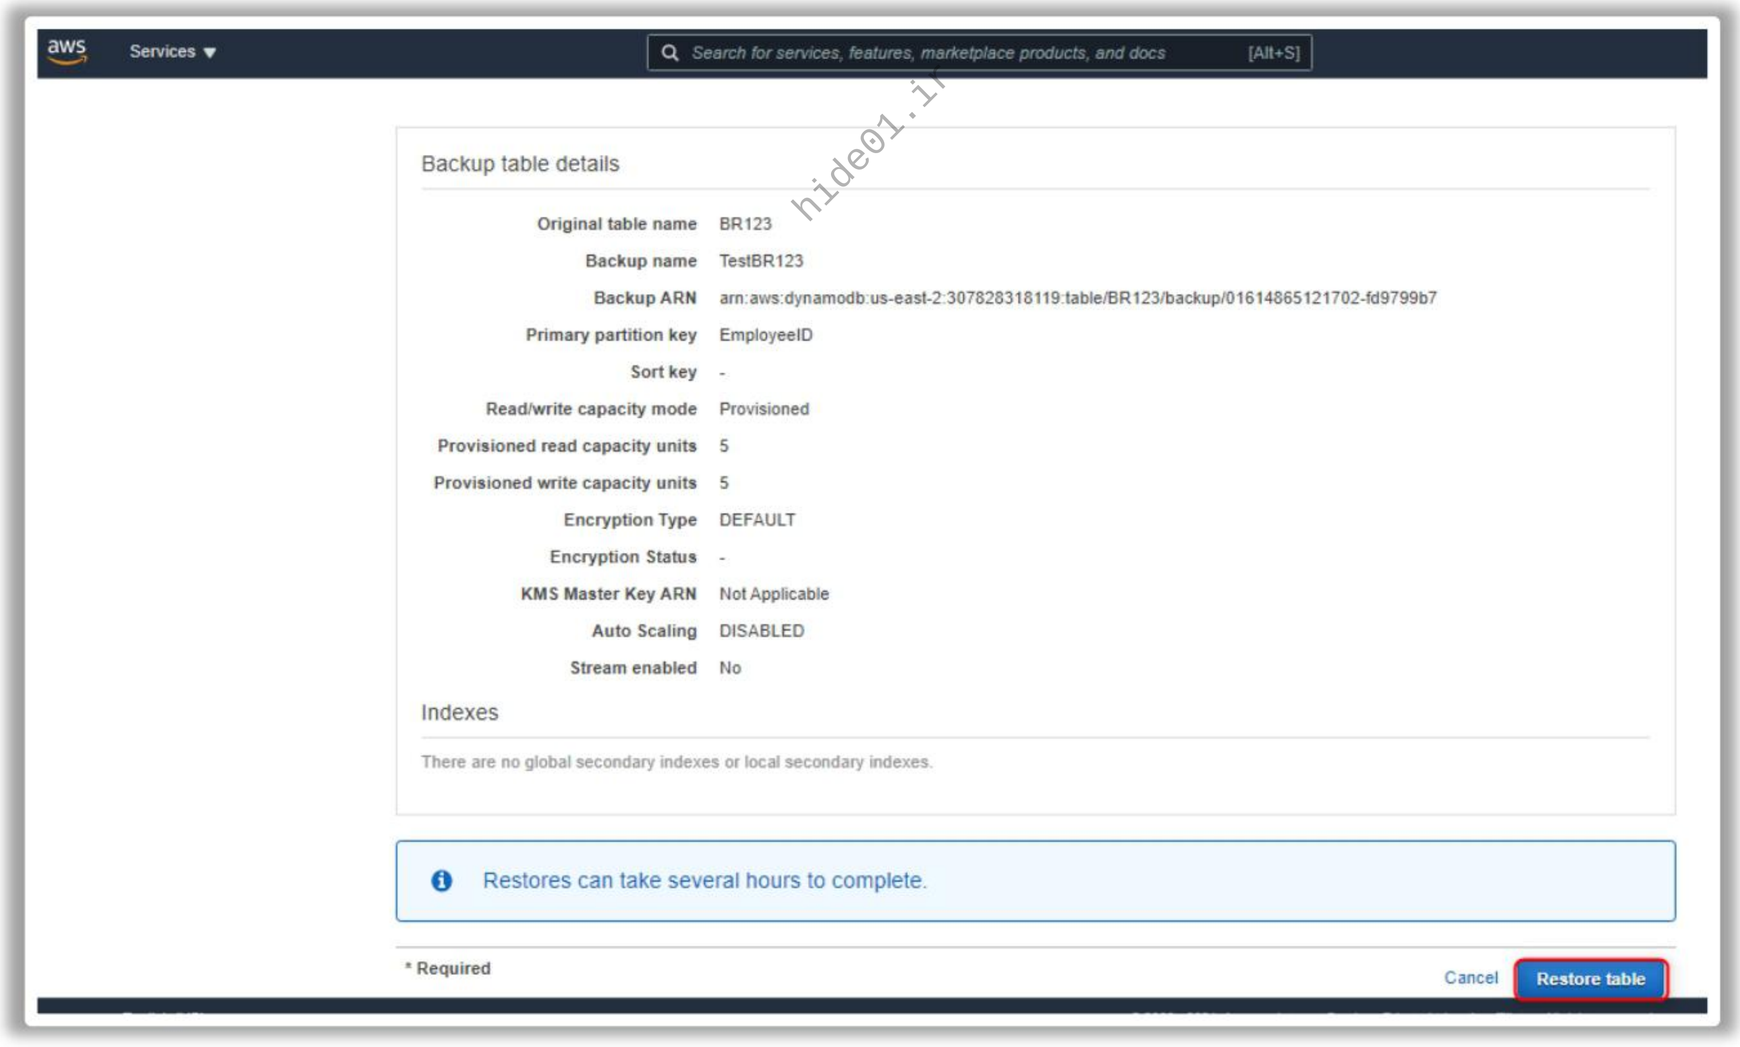

Figure 9.45: Restore table from Backup Console

Certified Cloud Security Engineer Copyright © by **EC-Council**
All Rights Reserved. Reproduction is Strictly Prohibited.

- Within a short time; the full backup table will be successfully restored.

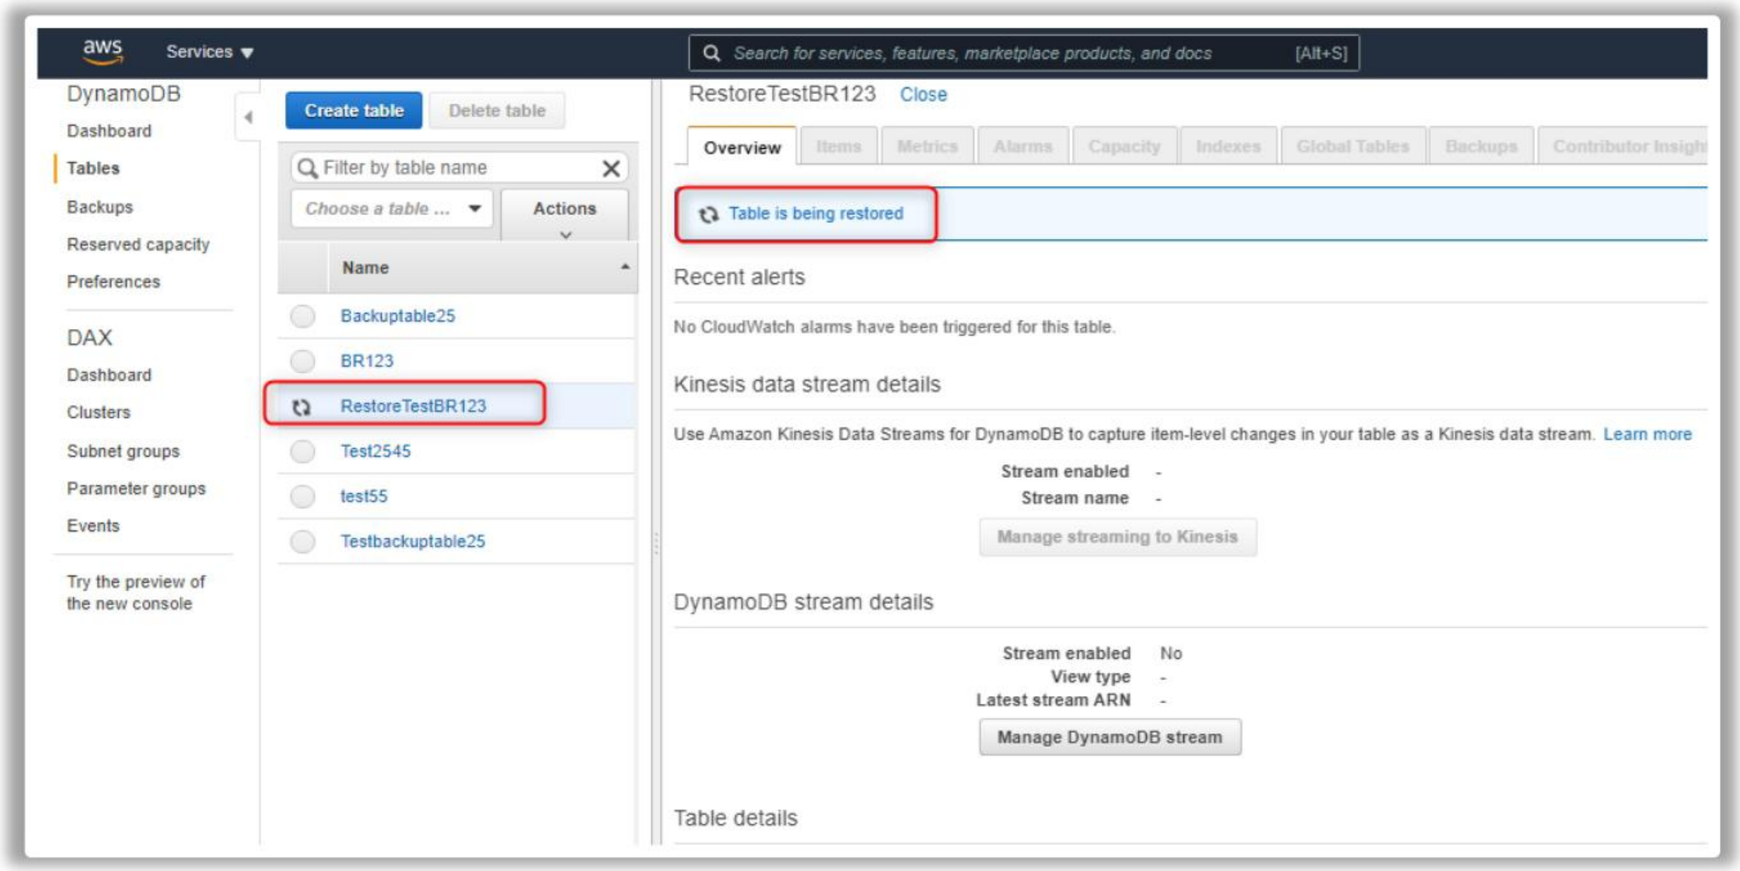

Figure 9.46: Restore table Backup Console

- After successful restoration of the table, click on the **Items** tab; you will observe that the table is restored successfully.

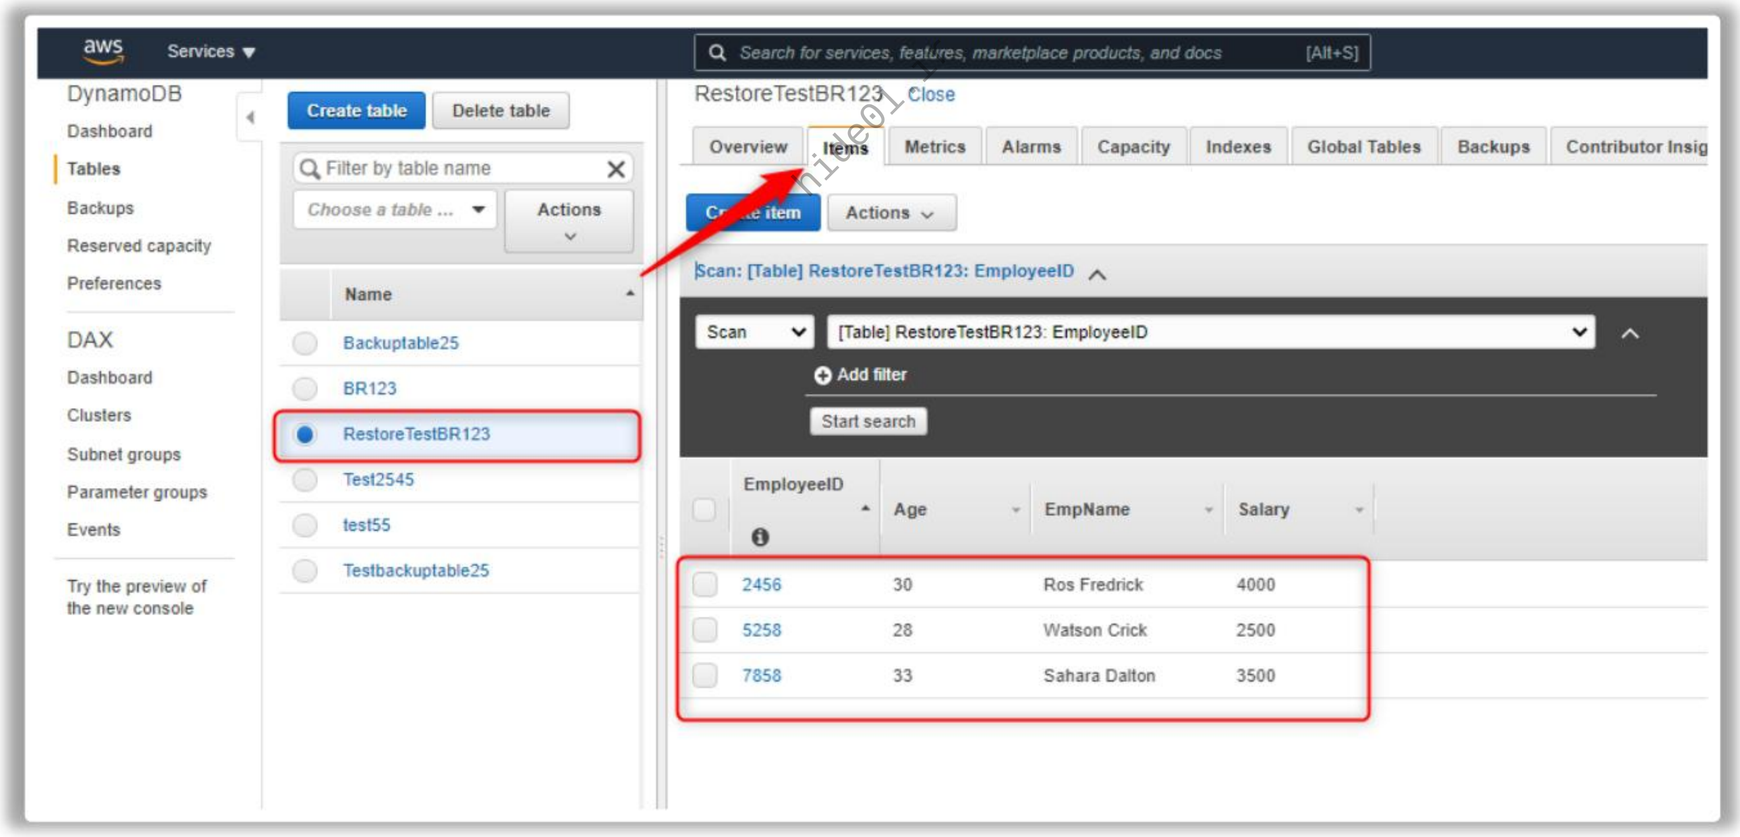

Figure 9.47: Restore table from Backup Console

Certified Cloud Security Engineer Copyright © by **EC-Council**
All Rights Reserved. Reproduction is Strictly Prohibited.

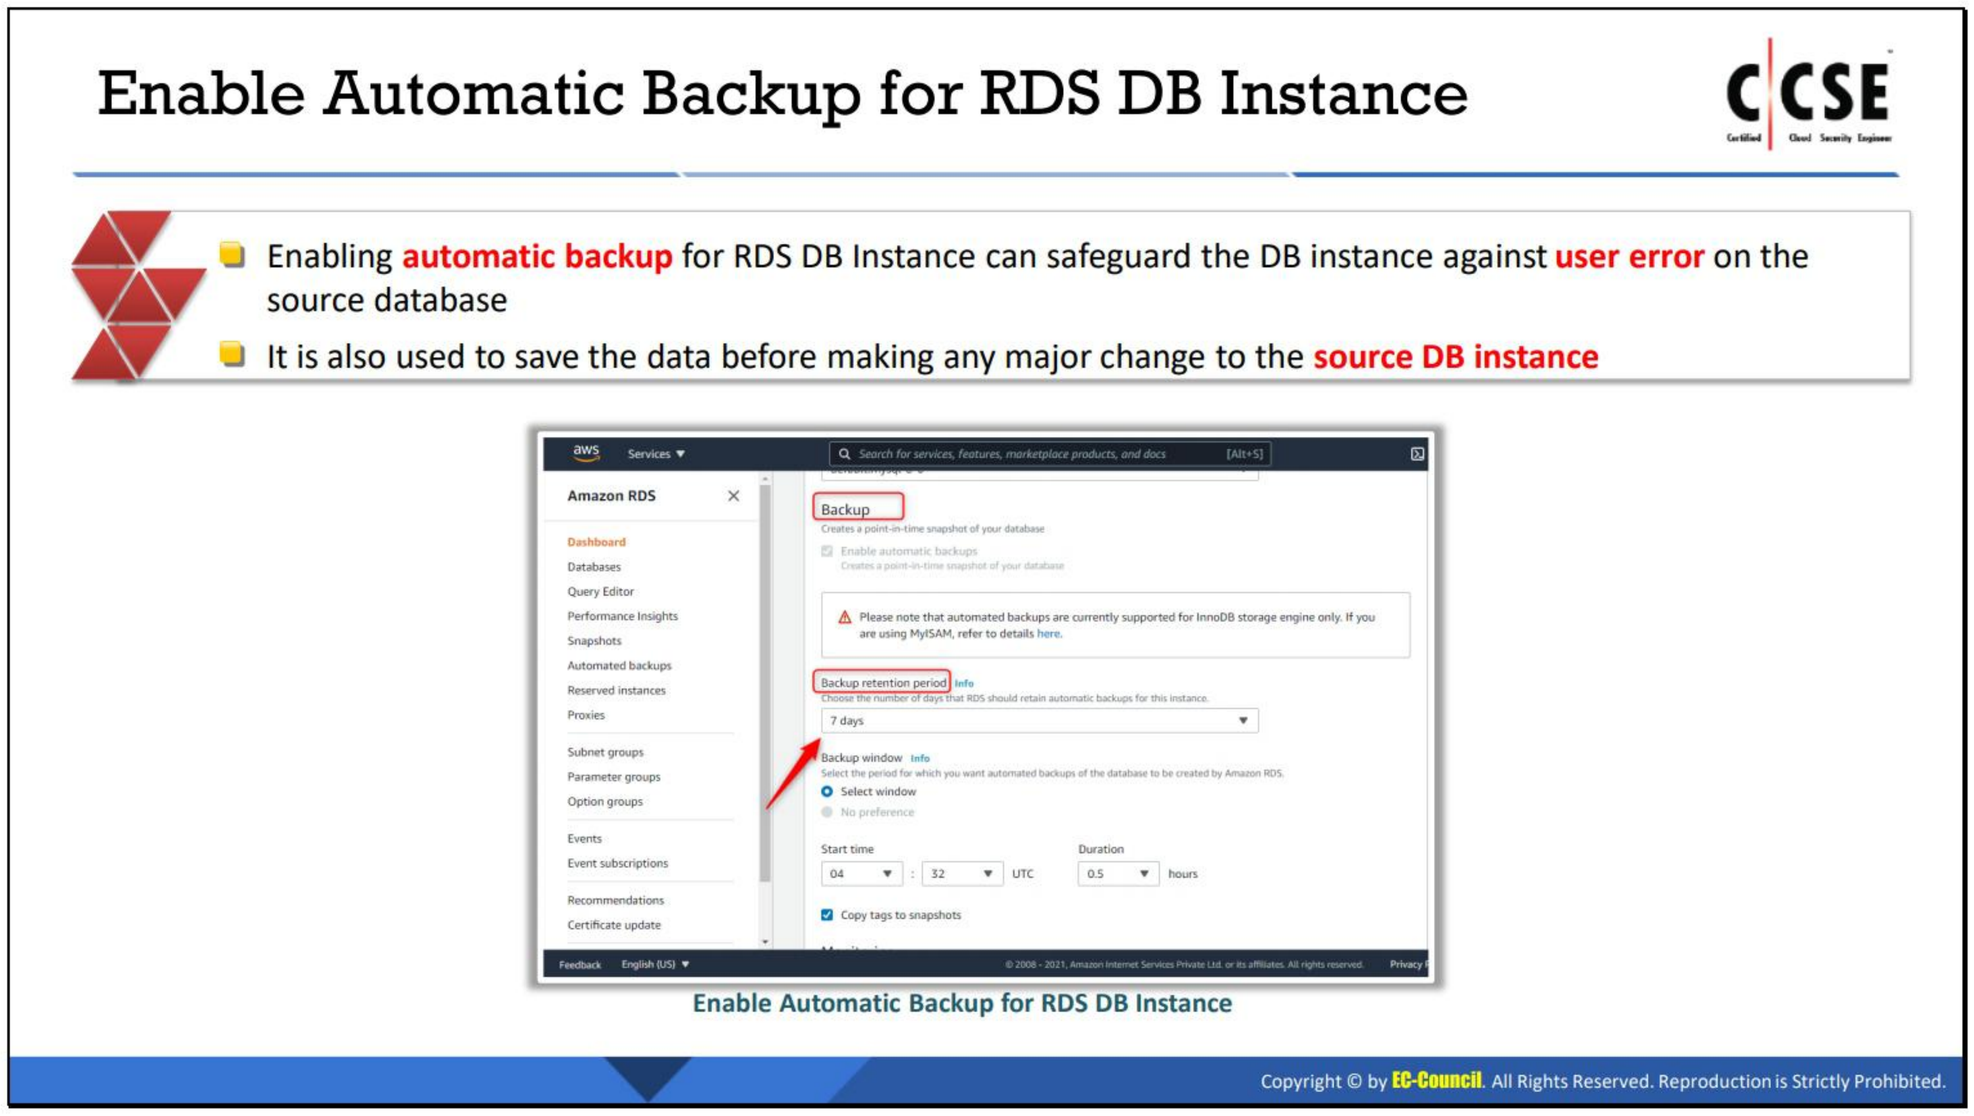

Enable Automatic Backup for RDS DB Instance

Enable Automatic Backup for RDS DB Instance

The cloud security professional should ensure that the automatic backup for the RDS DB instance is enabled. By enabling this feature, the cloud security administrator can allow the RDS DB to take the snapshot of the data automatically and keep the backup data for the selected time, which was defined at the time of enabling automatic backup for RDS DB instance. This will allow the security administrator to efficiently handle the data restoration process. Enabling automatic backup for RDS DB instance can safeguard the DB instance against user error on the source database and it is used to save the data before making any major change to the source DB instance.

To Enable Automatic Backup for RDS DB Instance

- Sign in to AWS Management Console.

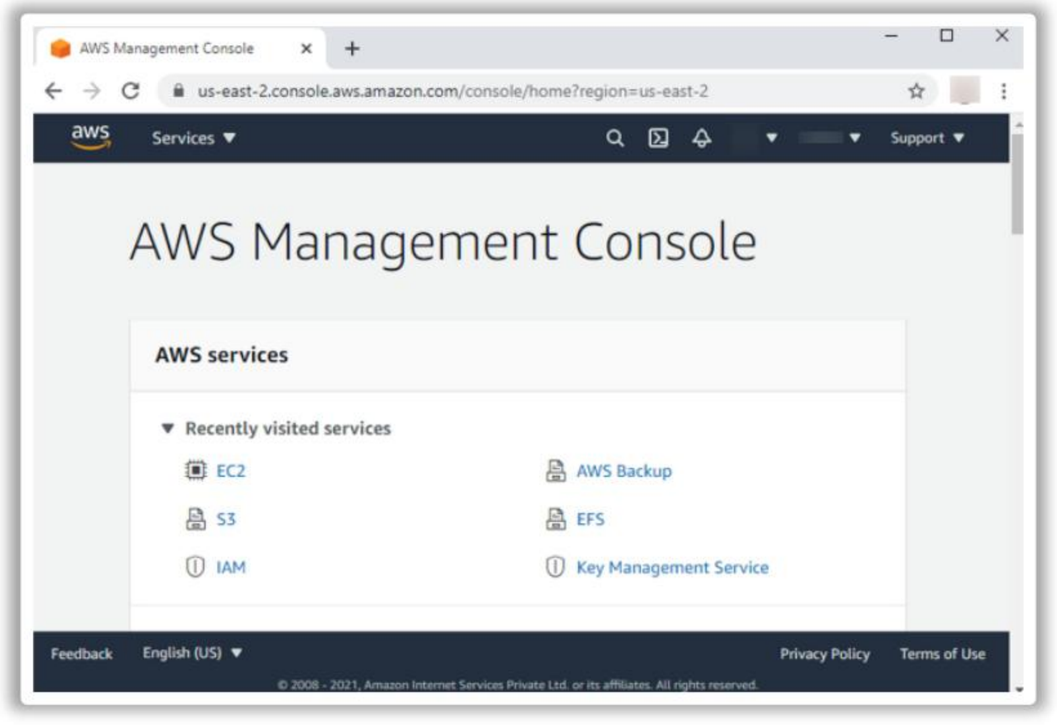

Figure 9.48: AWS Management Console

Certified Cloud Security Engineer Copyright © by **EC-Council**
All Rights Reserved. Reproduction is Strictly Prohibited.

- Click on the following link to open the RDS database: https://console.aws.amazon.com/rds/home. Click on **Databases** in the navigation pane. Then select the database instance that you want to modify and click on it.

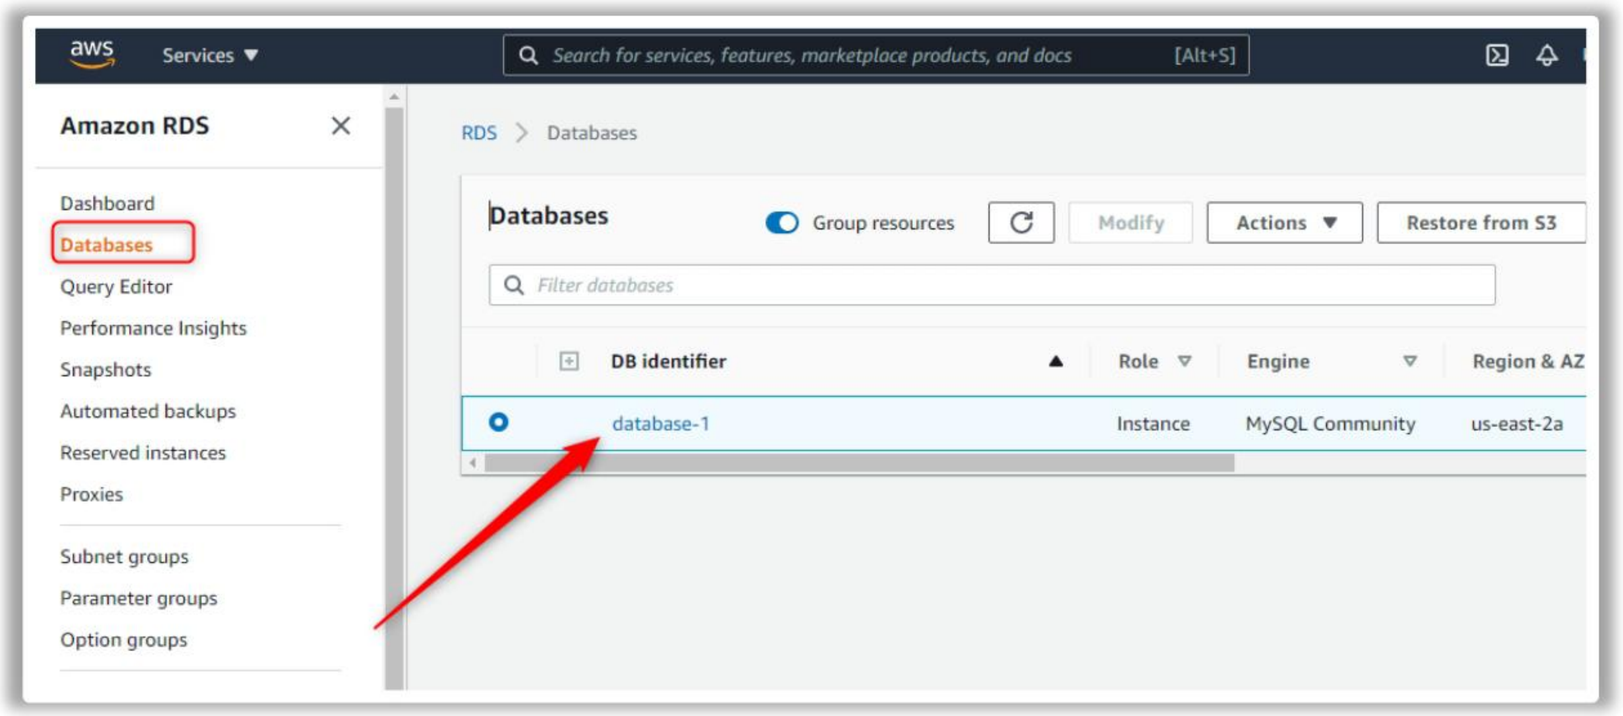

Figure 9.49: Database Console

- Click on **Modify** button.

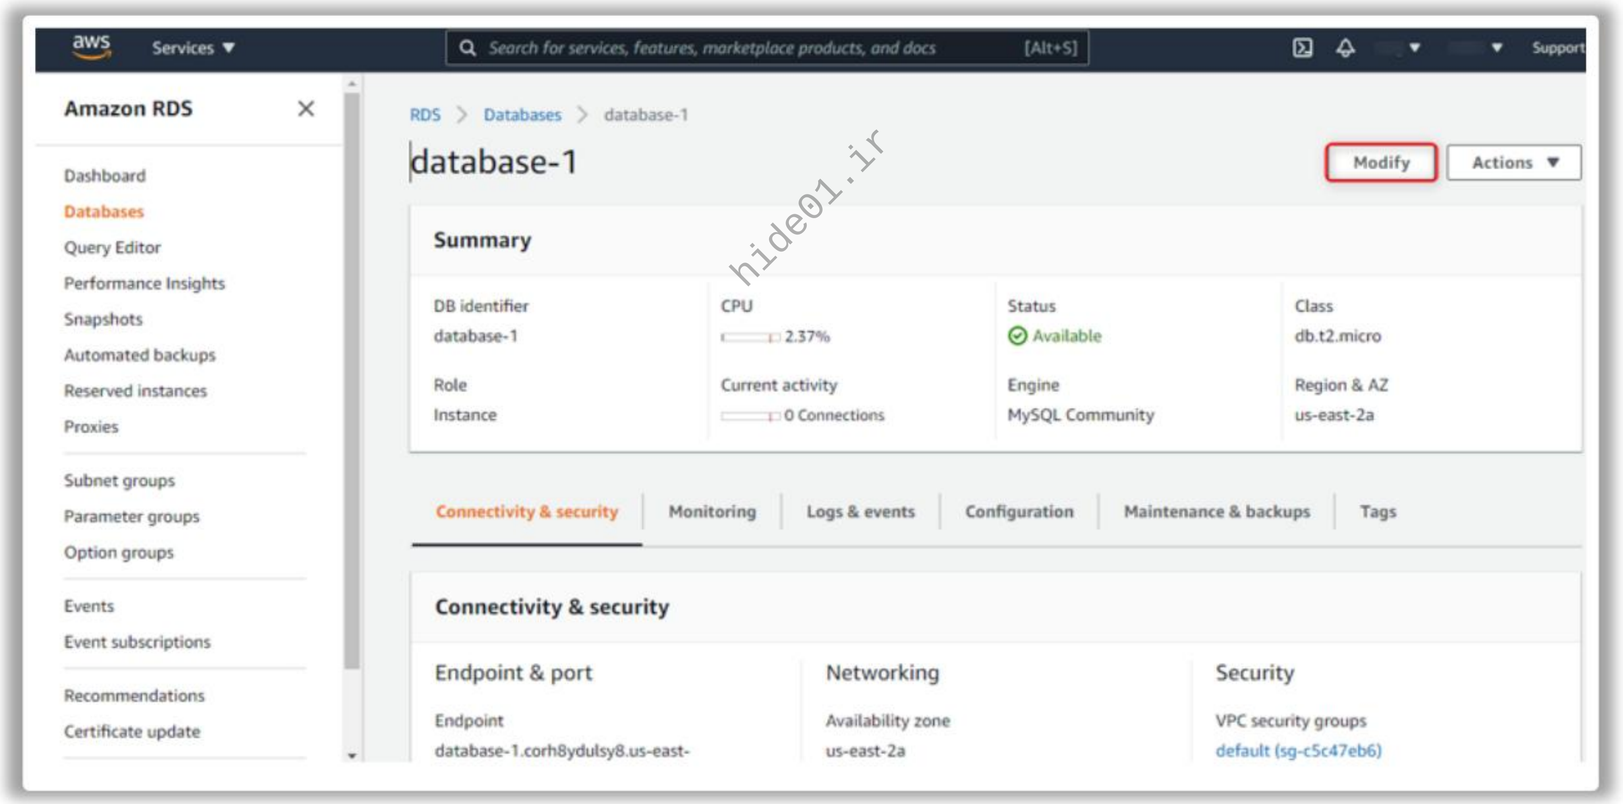

hide01.ir

Figure 9.50: Database Console

Certified Cloud Security Engineer Copyright © by **EC-Council**
All Rights Reserved. Reproduction is Strictly Prohibited.

- **Modify DB Instance** page will page. Navigate to the Backup section and, under the Backup retention period, select the number of days that the RDS would retain the automatic backup.

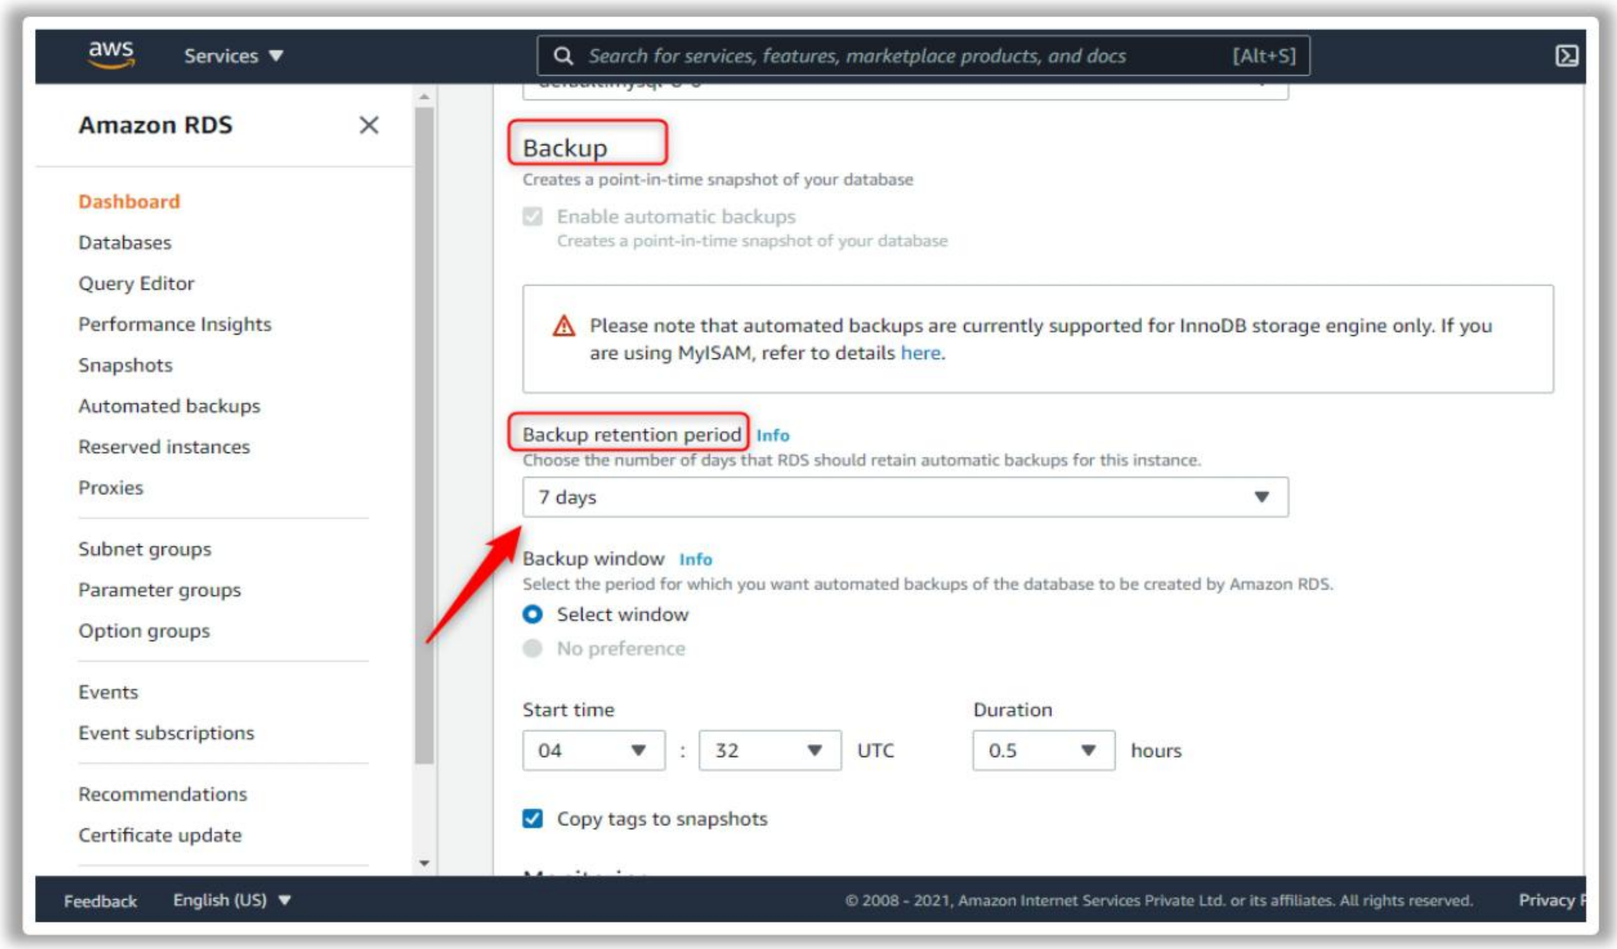

Figure 9.51: Backup Console

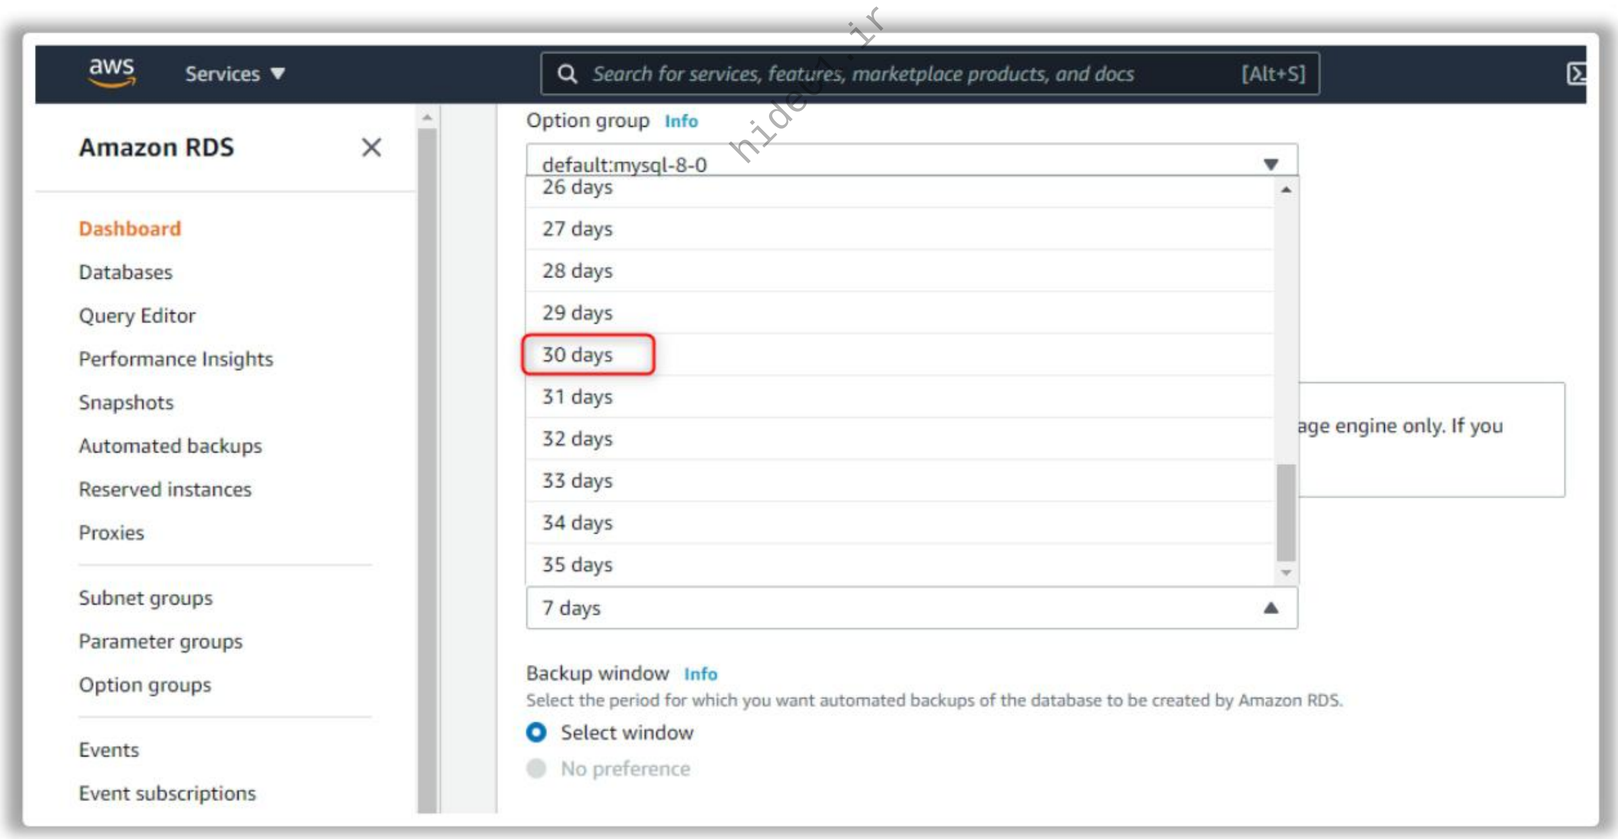

Figure 9.52: Dashboard Console

Certified Cloud Security Engineer Copyright © by **EC-Council**
All Rights Reserved. Reproduction is Strictly Prohibited.

- Then, click on the **Continue** button.

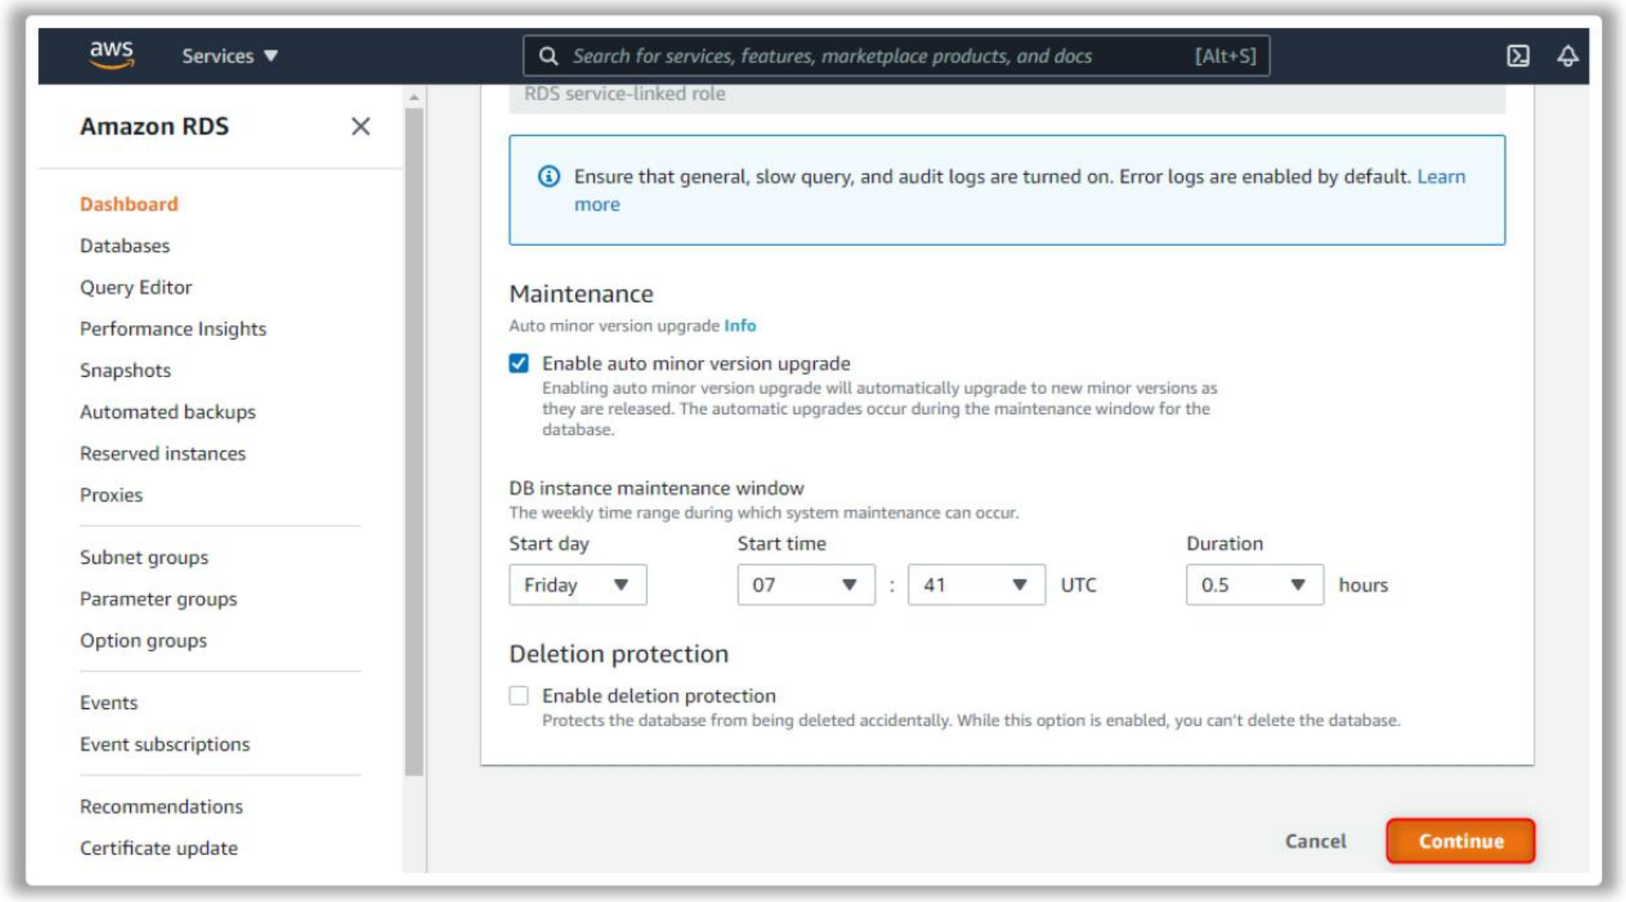

Figure 9.53: Database Console

- Then, click on **Modify DB Instance** button.

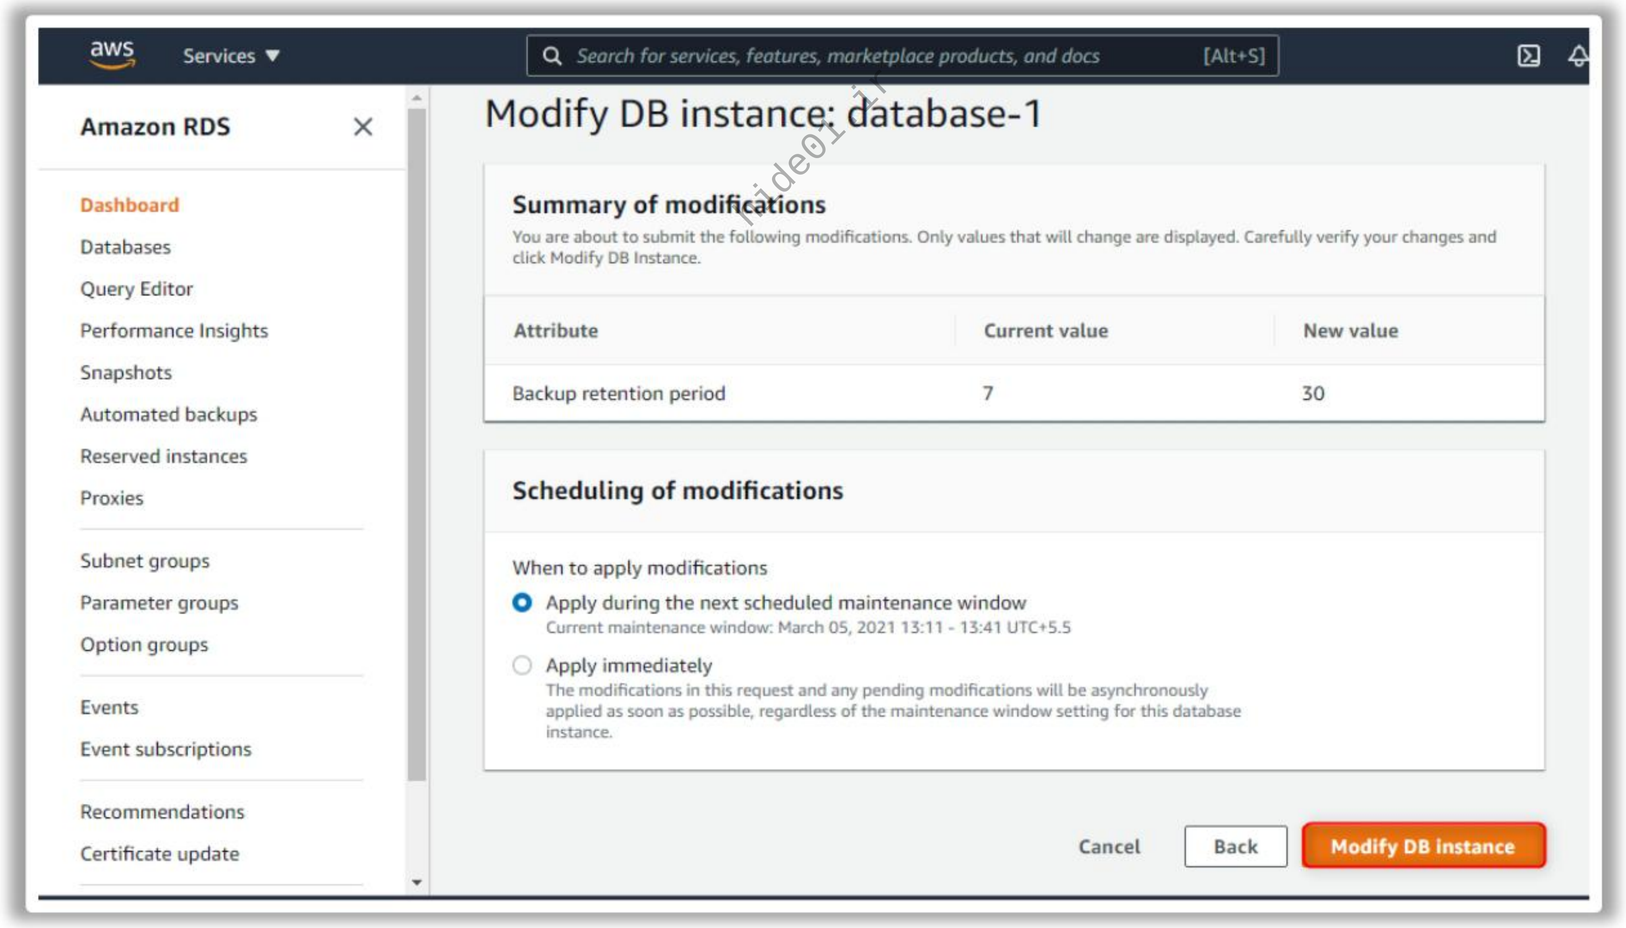

Figure 9.54: Modify Database Console

Certified Cloud Security Engineer Copyright © by **EC-Council**
All Rights Reserved. Reproduction is Strictly Prohibited.

- An automatic backup will be successfully enabled.

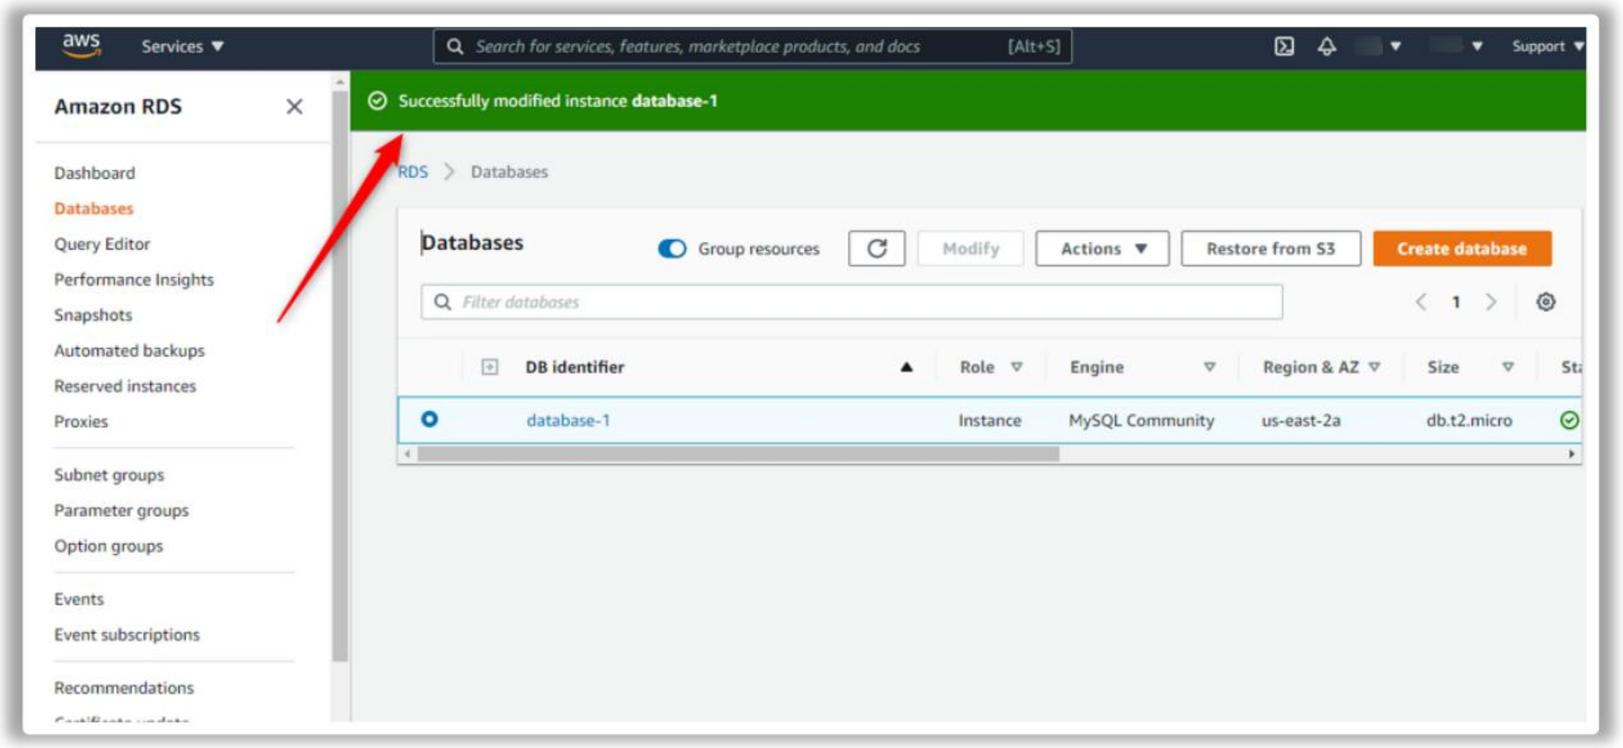

Figure 9.55: Modify Database Console

hide01.ir

Certified Cloud Security Engineer Copyright © by **EC-Council**
All Rights Reserved. Reproduction is Strictly Prohibited.

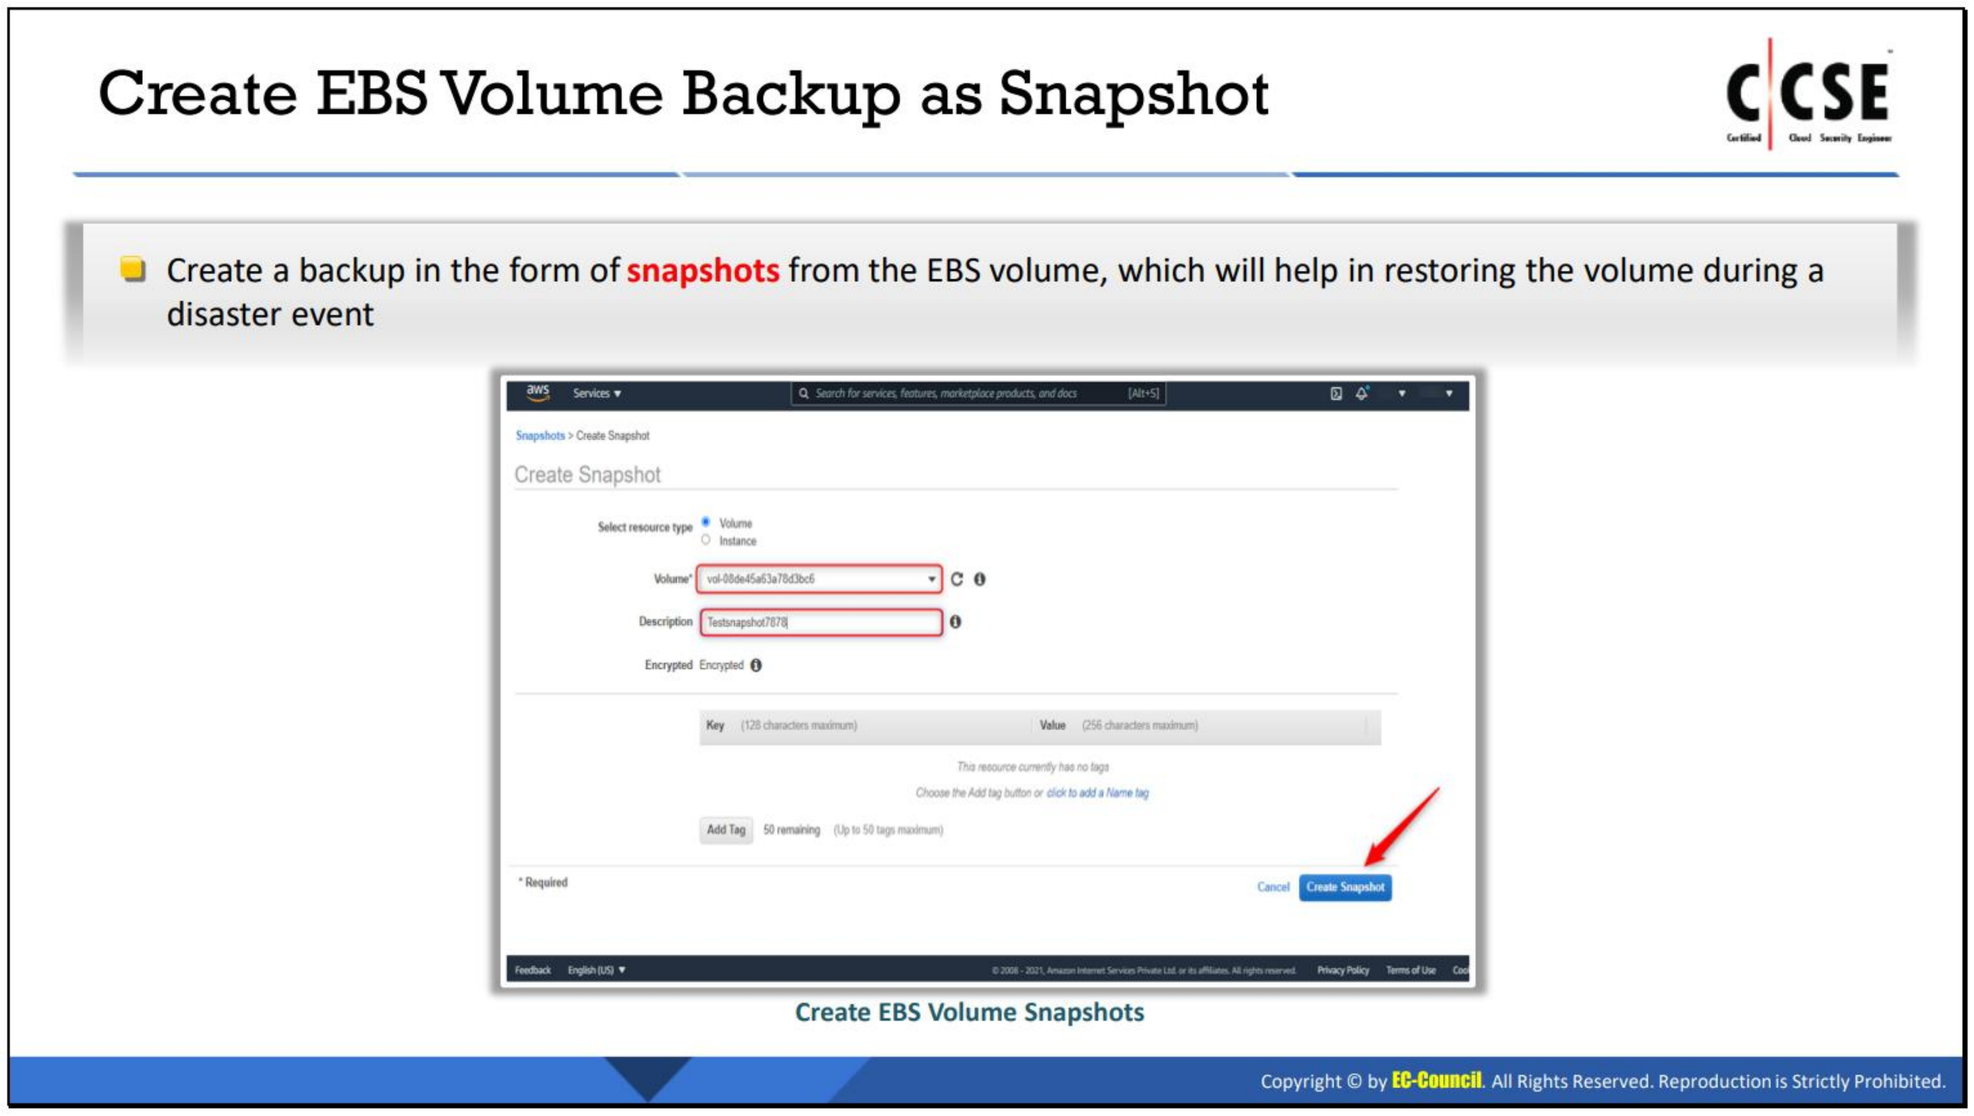

Create EBS Volume Backup as Snapshot

The services that are running in EC2 instance can utilize the EBS volume for accessing and storing the primary data. In Amazon EBS, the cloud security professional can create a backup in the form of snapshots from the EBS volume. The copy of the volume will be placed in Amazon S3 and is stored redundantly in various geographical regions of AWS. The first snapshot will store the entire copy of the volume and the next snapshots will store the incremental block-level changes. It is considered a rapid and trustworthy approach to store the full backup.

To Create EBS Volume Snapshots

- Sign in to AWS Management Console.

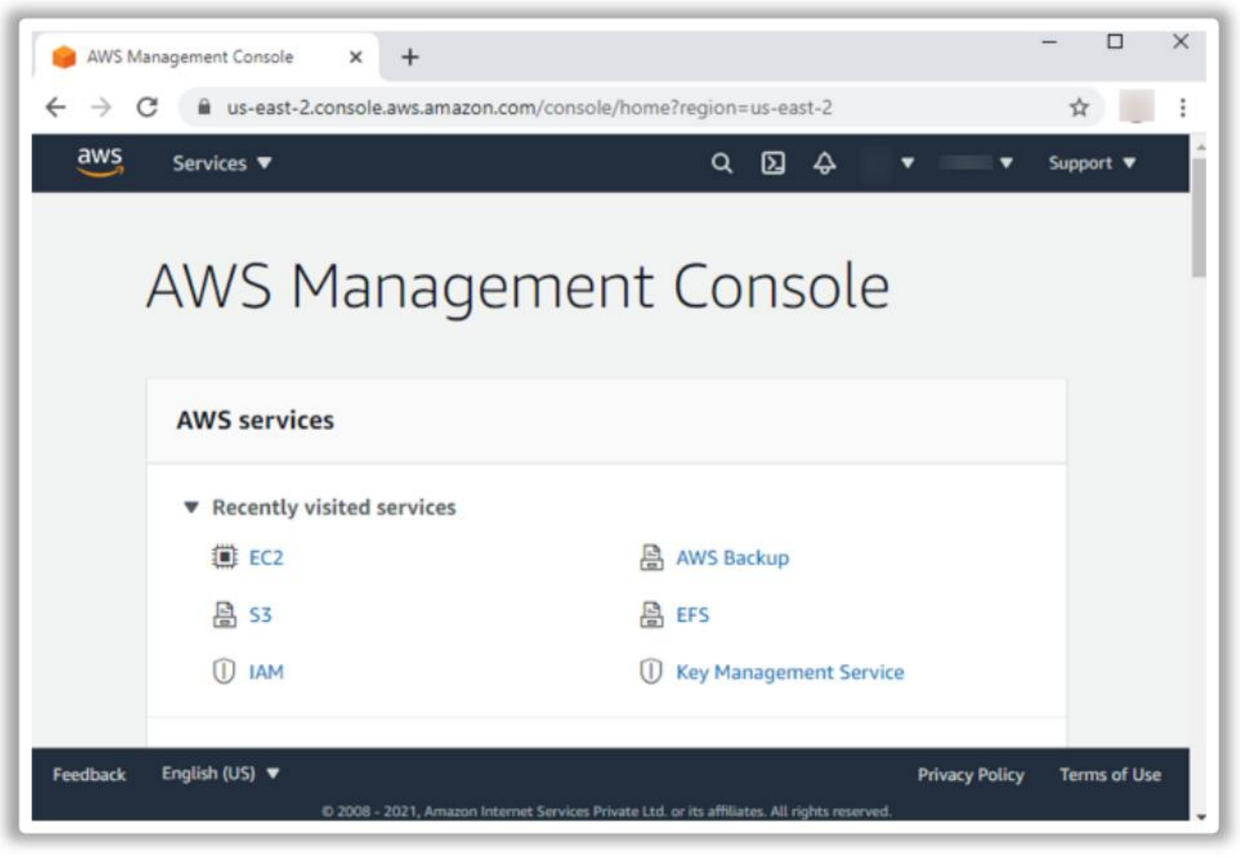

Figure 9.56: AWS Management Console

Certified Cloud Security Engineer Copyright © by **EC-Council**
All Rights Reserved. Reproduction is Strictly Prohibited.

- Click on the following link to open the EC2 dashboard: https://console.aws.amazon.com/ec2/.

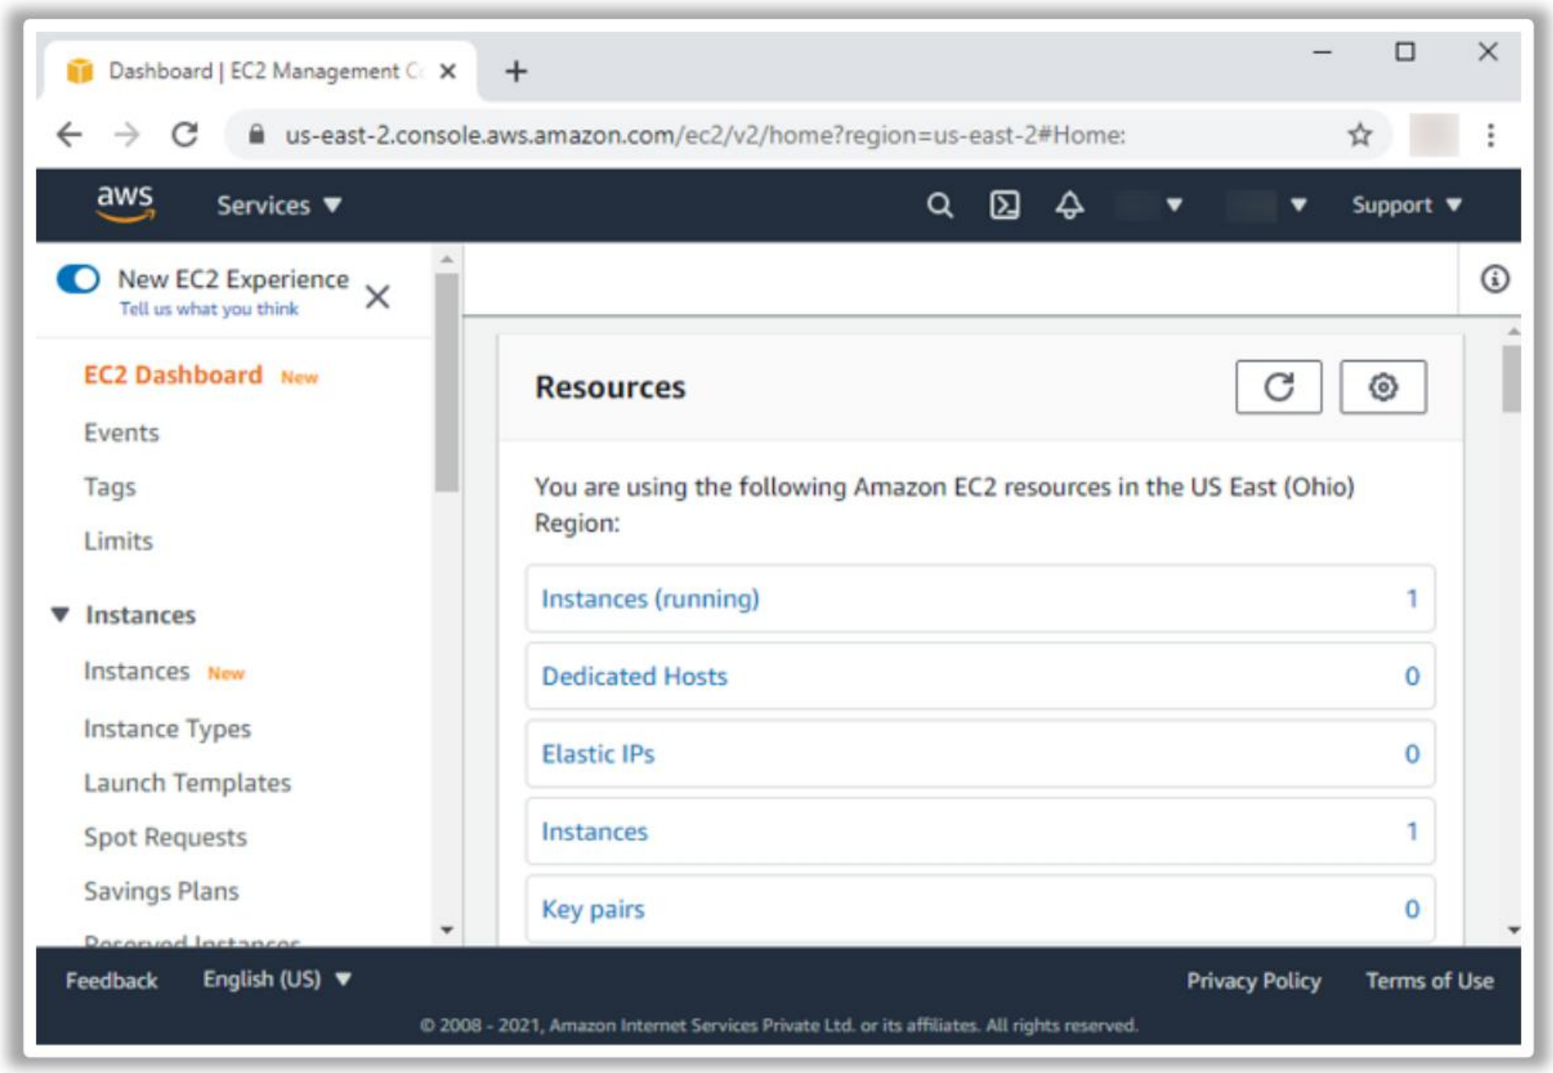

Figure 9.57: EC2 Dashboard Console

- Under **Elastic Block Store**, you will find various volumes; remember the Volume ID for which you want to create a snapshot.

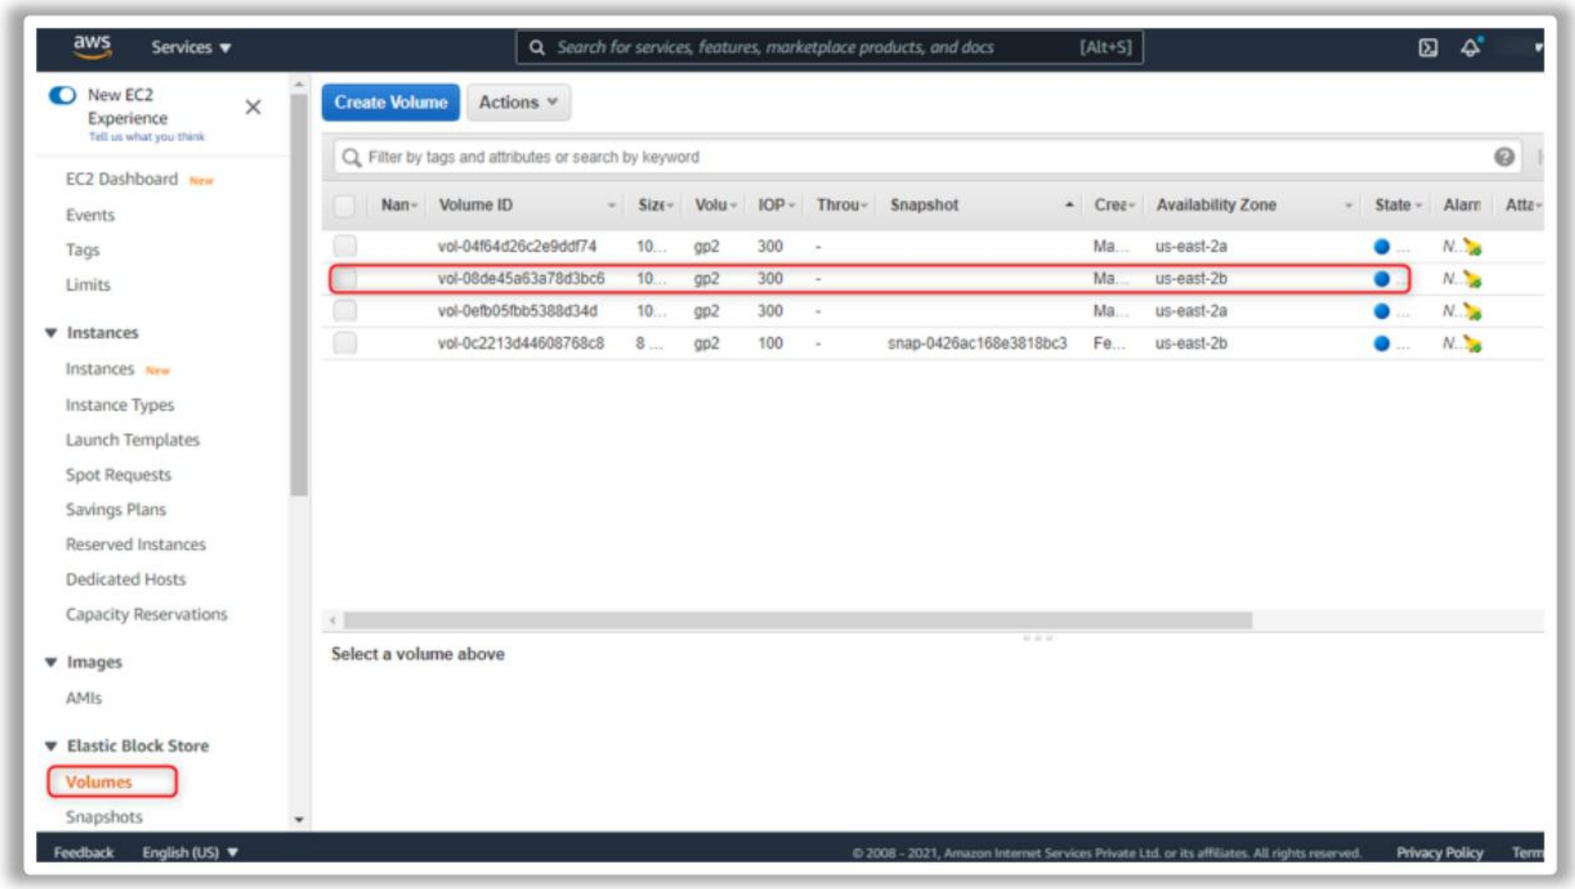

Figure 9.58: Create Volume Console

Certified Cloud Security Engineer Copyright © by **EC-Council**
All Rights Reserved. Reproduction is Strictly Prohibited.

- Go to the **Elastic Block store** in the navigation panel and click on **Snapshots**; then click on **Create Snapshot button.**

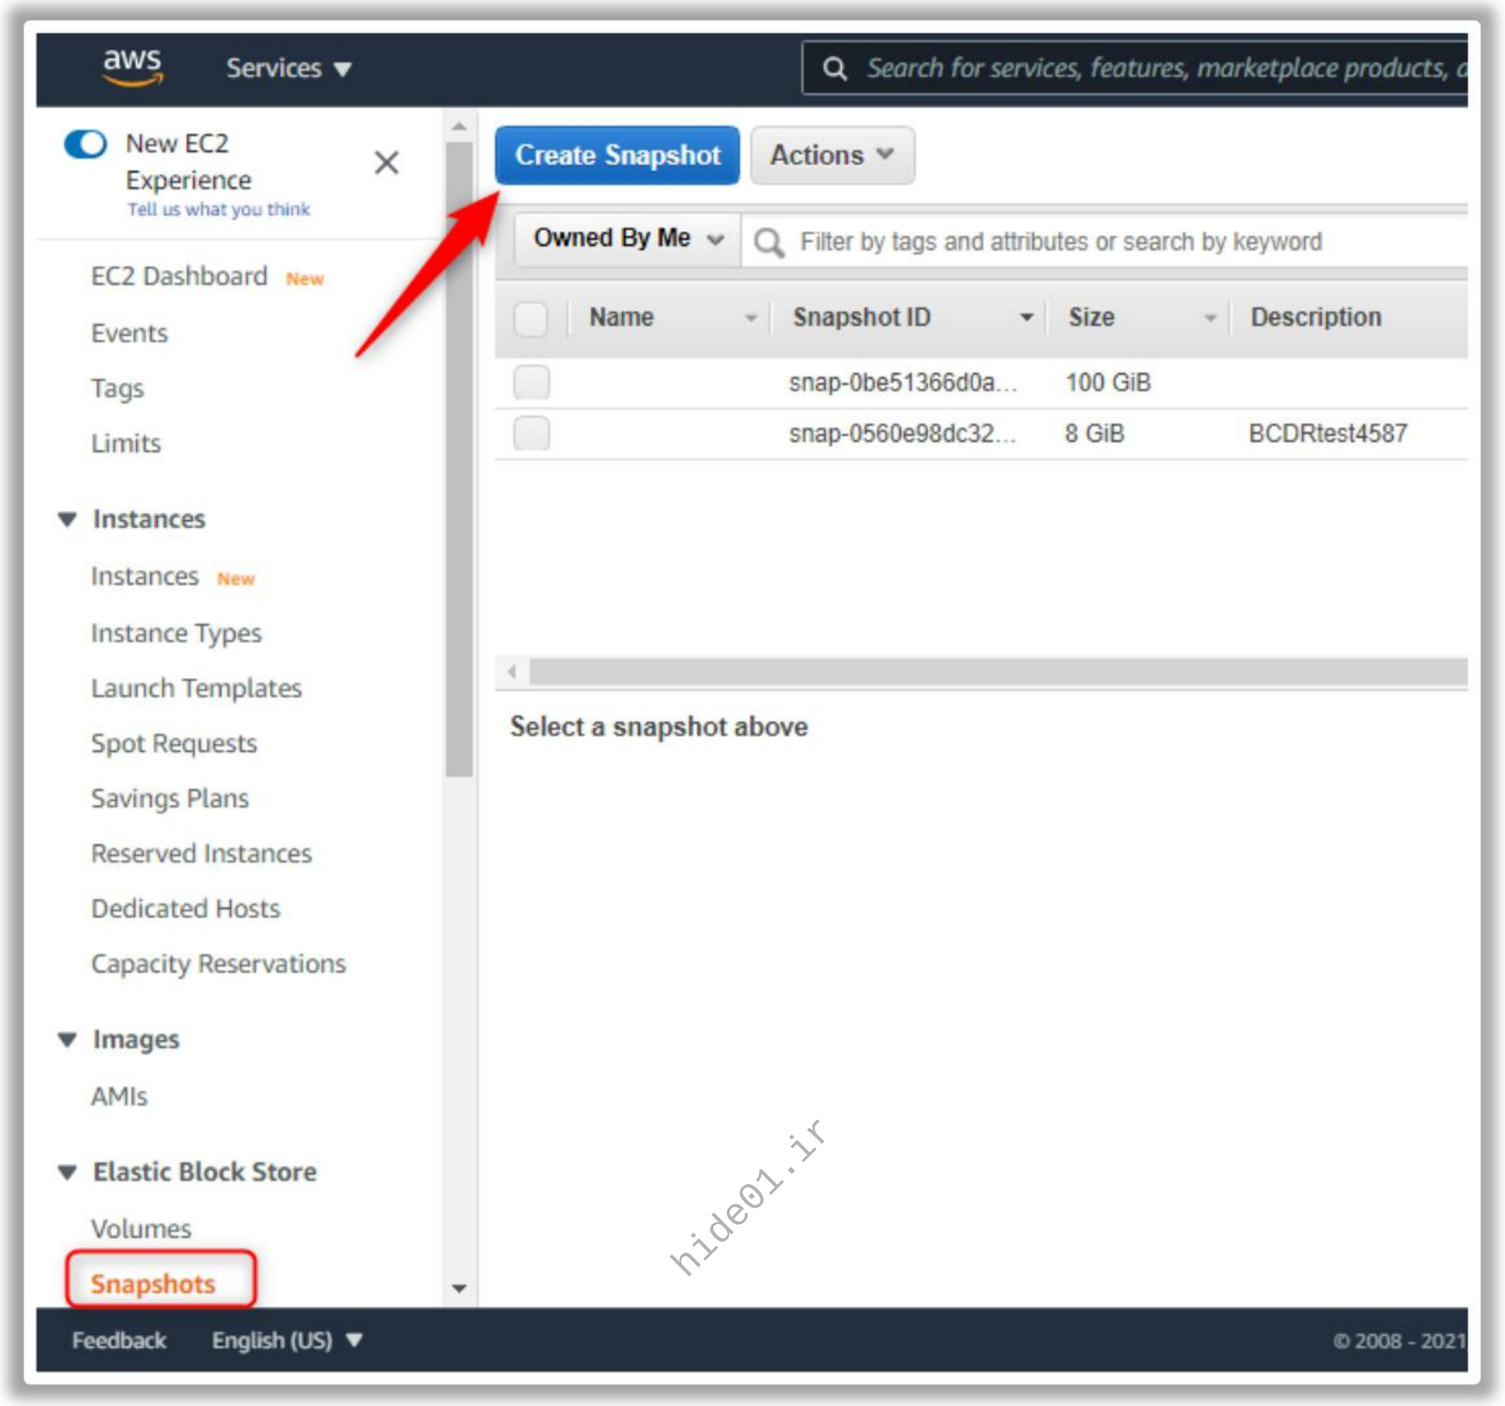

Figure 9.59: Create Snapshot Console

- Enter the Volume ID in the **Volume** field and then describe the Snapshot in the **Description** field; click on **Create Snapshot** button.

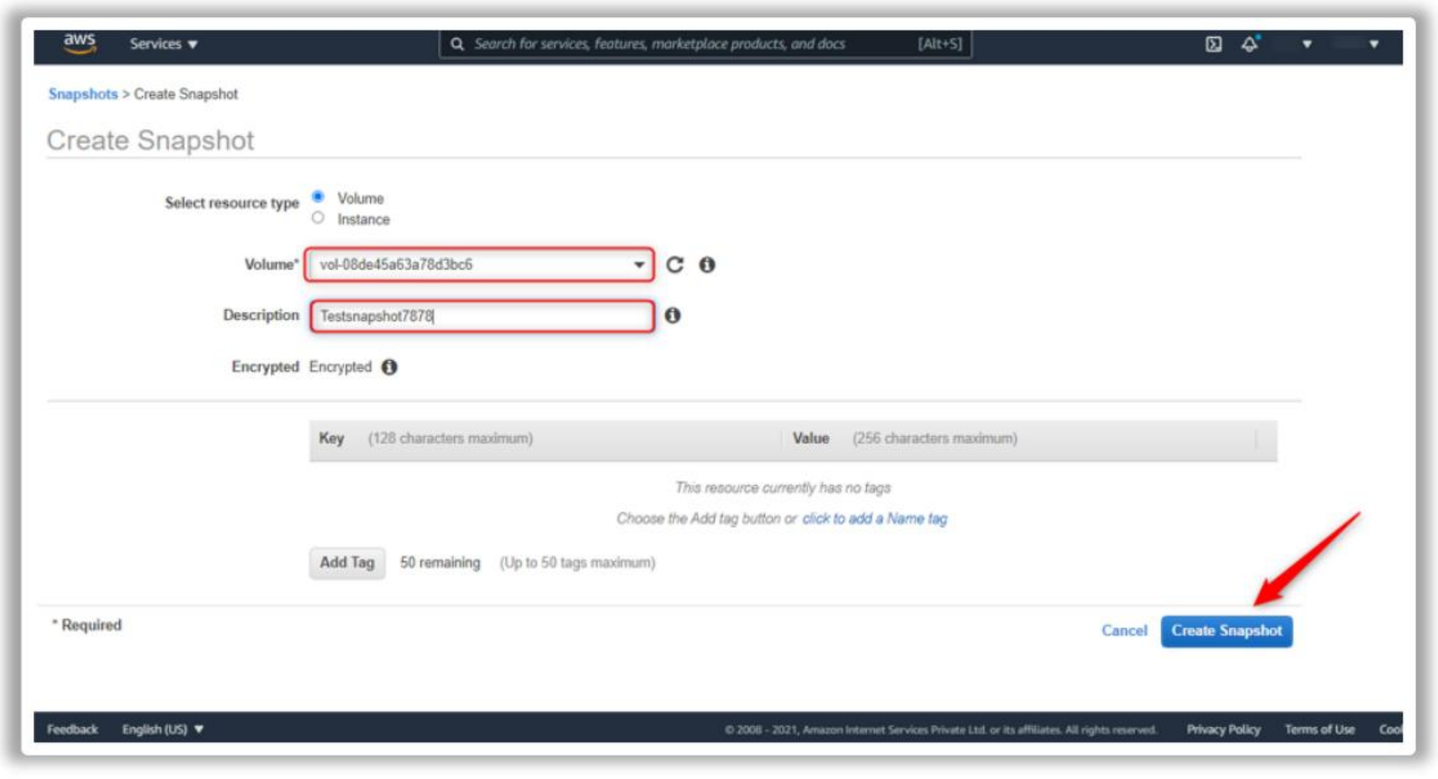

Figure 9.60: Create Snapshot Console

Certified Cloud Security Engineer Copyright © by **EC-Council**
All Rights Reserved. Reproduction is Strictly Prohibited.

- The snapshot is successfully created, click the **Close** button.

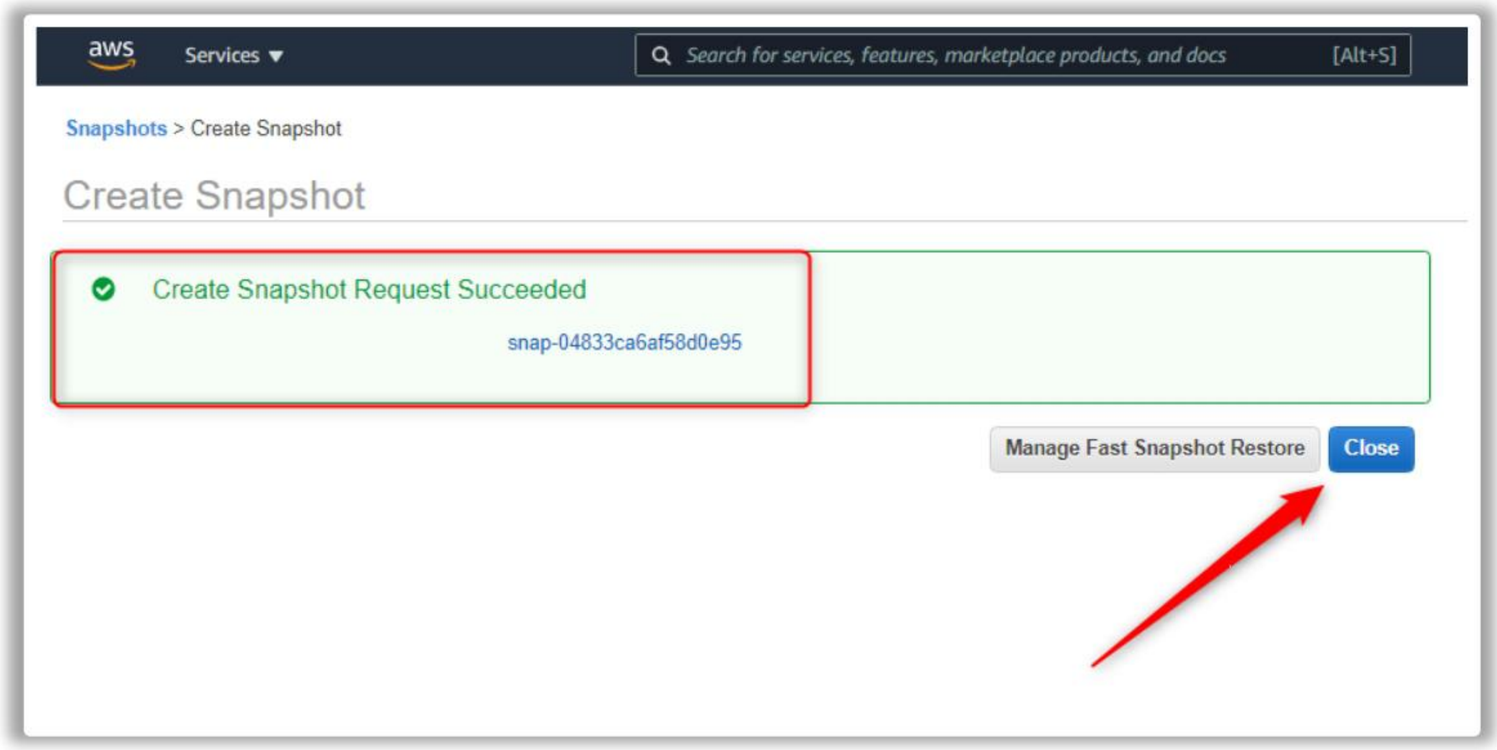

Figure 9.61: Snapshot Created

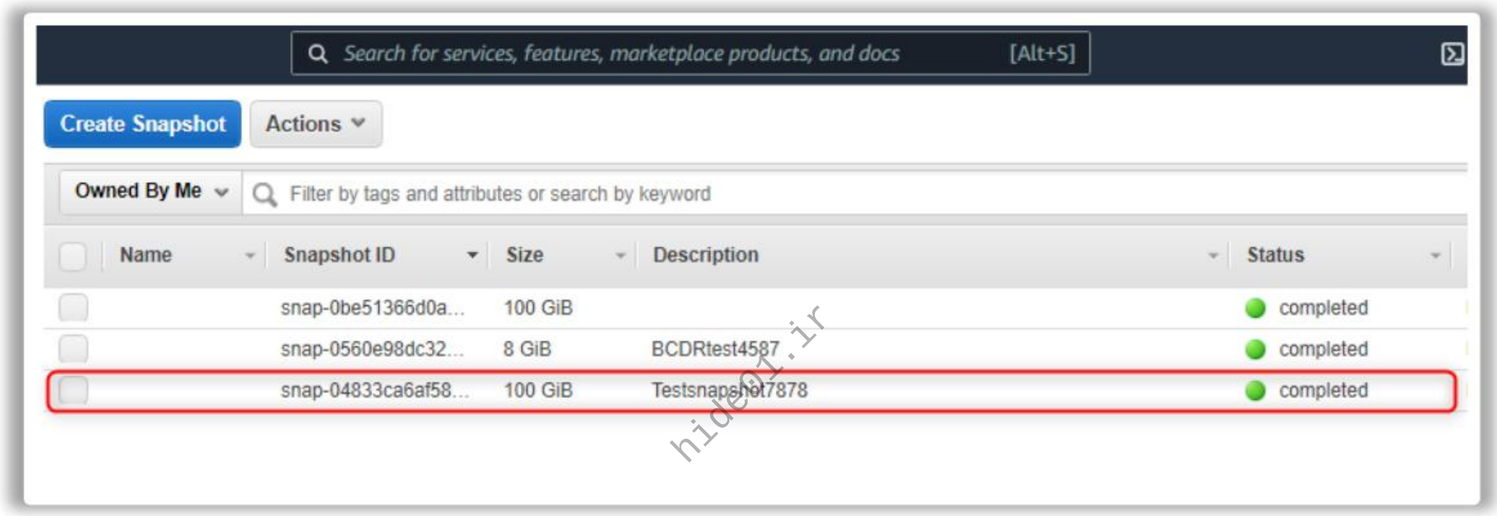

Figure 9.62: Snapshot Console

Certified Cloud Security Engineer Copyright © by **EC-Council**
All Rights Reserved. Reproduction is Strictly Prohibited.

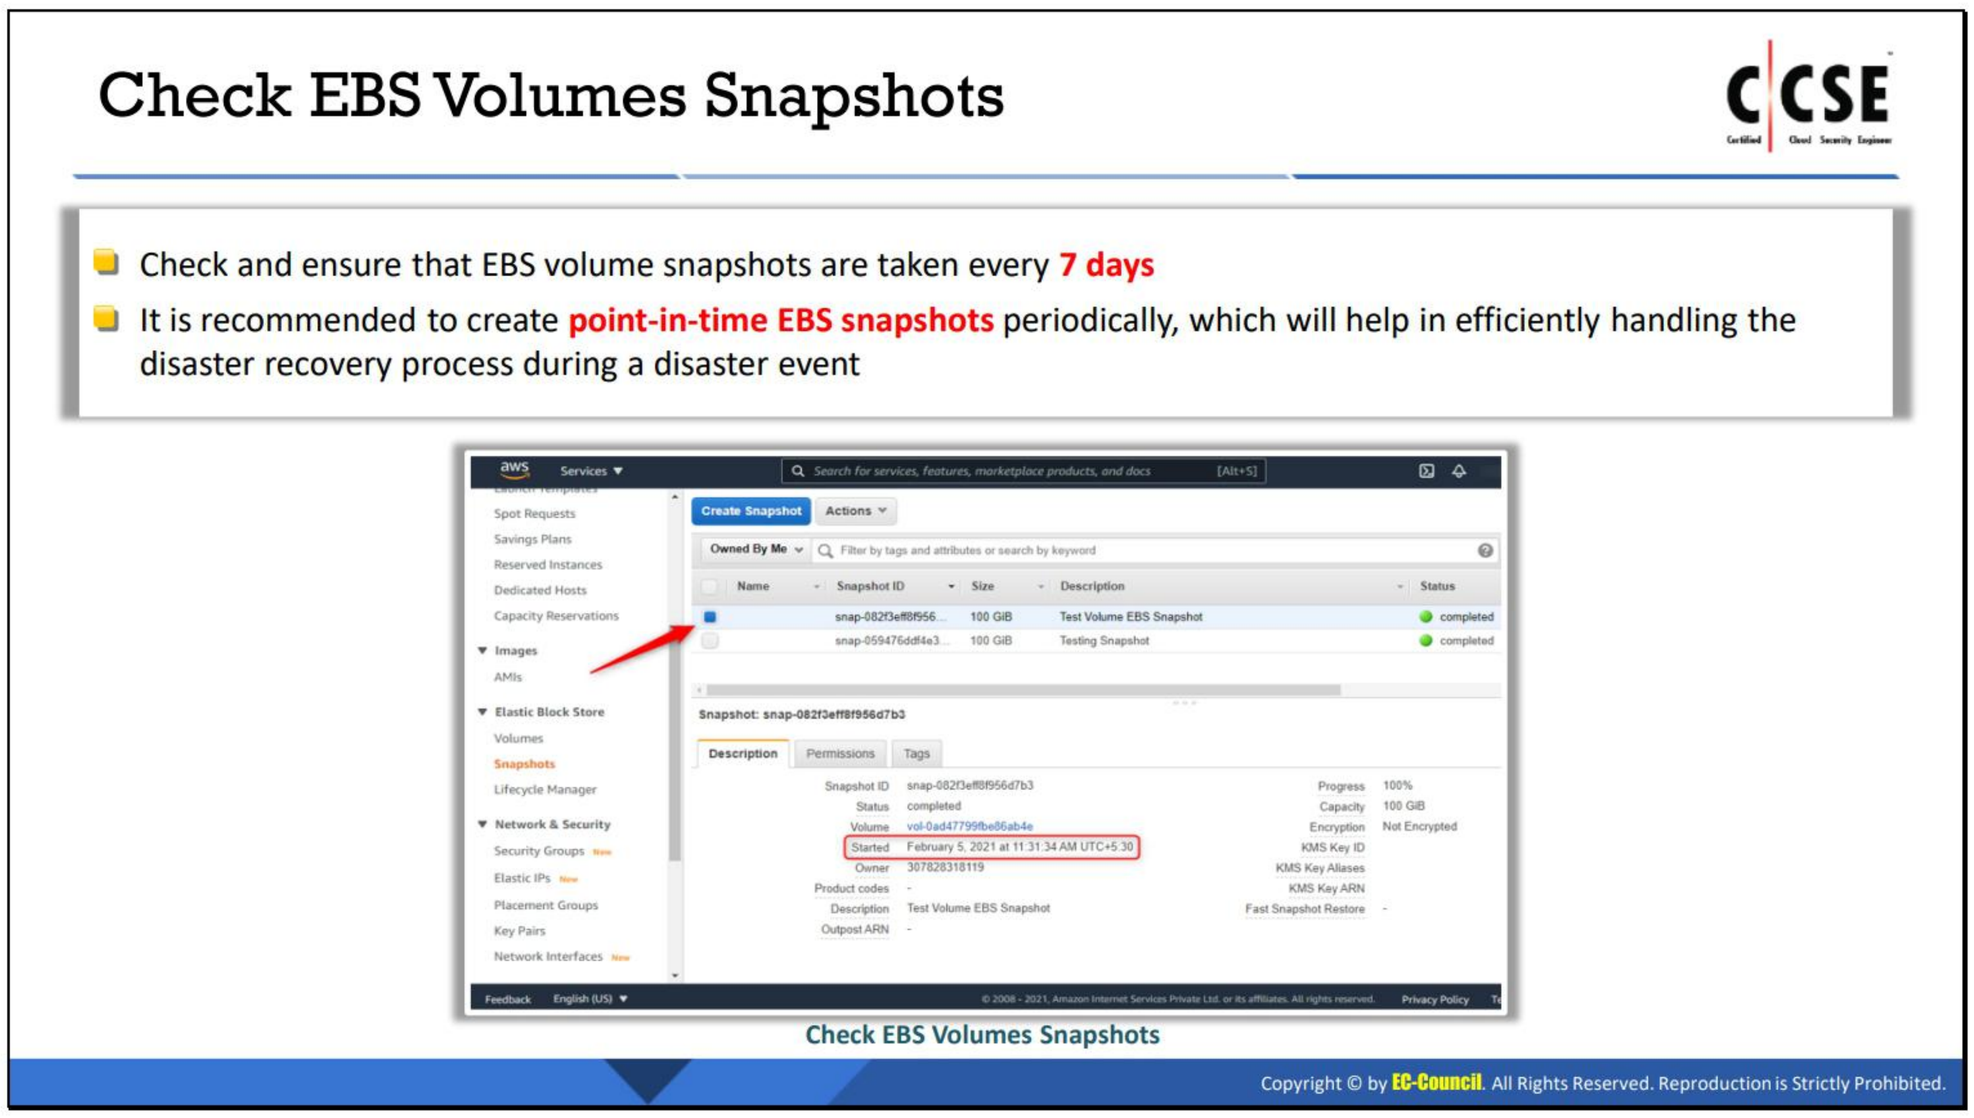

Check EBS Volumes Snapshots

Copyright © by EC-Council. All Rights Reserved. Reproduction is Strictly Prohibited.

Check EBS Volumes Snapshots

The cloud security professional should check the EBS volume snapshots to ensure that they are recently taken. The snapshots should be taken every 7 days. It is recommended to create point-in-time EBS snapshots periodically, which will help in efficiently handling the disaster recovery process during a disaster event. It will save the data before the EC2 instance shutdown. It also allows that the organization to meet the compliance standards of NIST.

To Check EBS Volumes Snapshots

- Sign in to AWS Management Console.

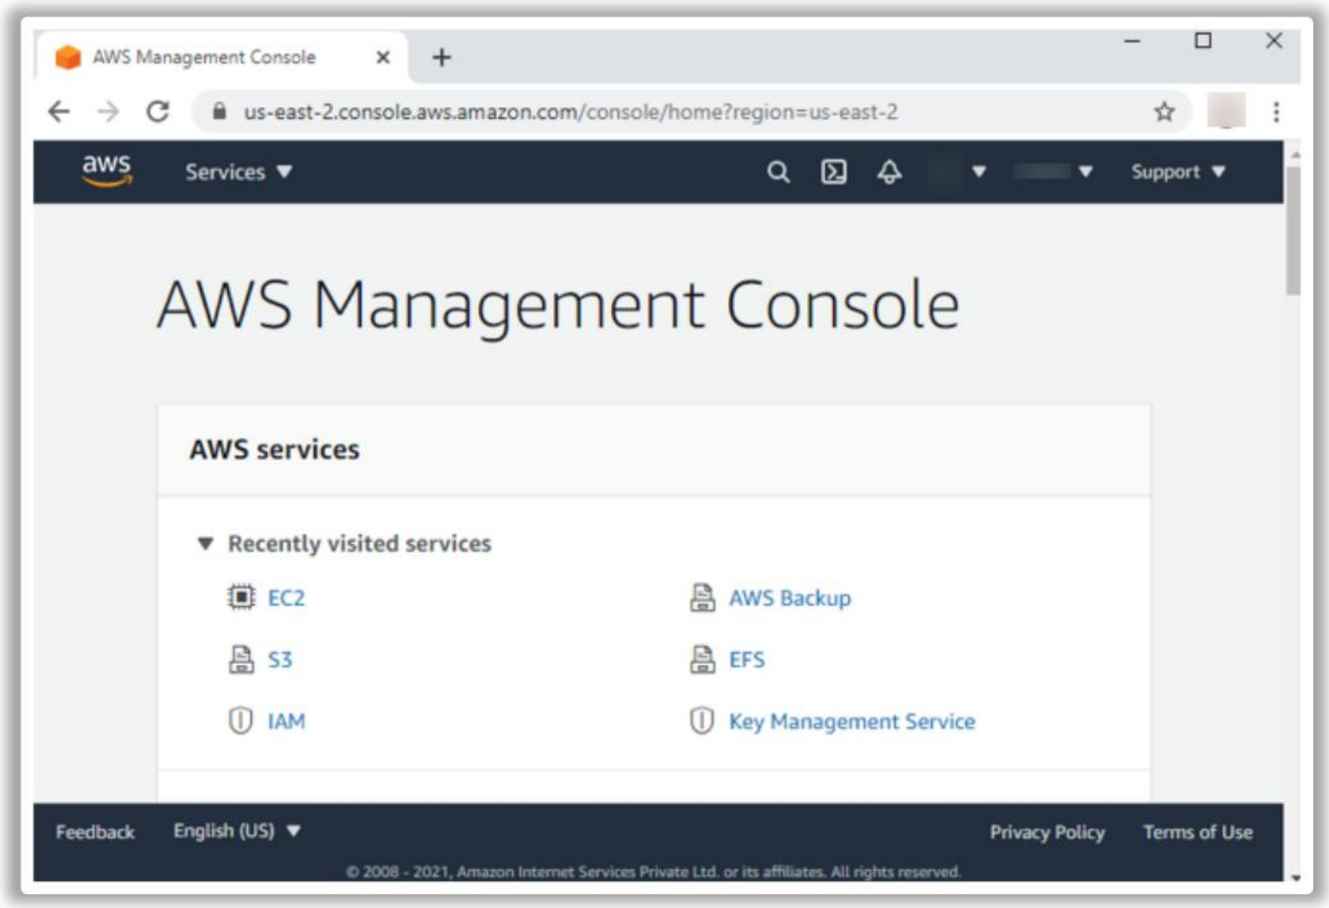

Figure 9.63: AWS Management Console

Certified Cloud Security Engineer Copyright © by **EC-Council**
All Rights Reserved. Reproduction is Strictly Prohibited.

- Click on the following link to open EC2 dashboard: https://console.aws.amazon.com/ec2/

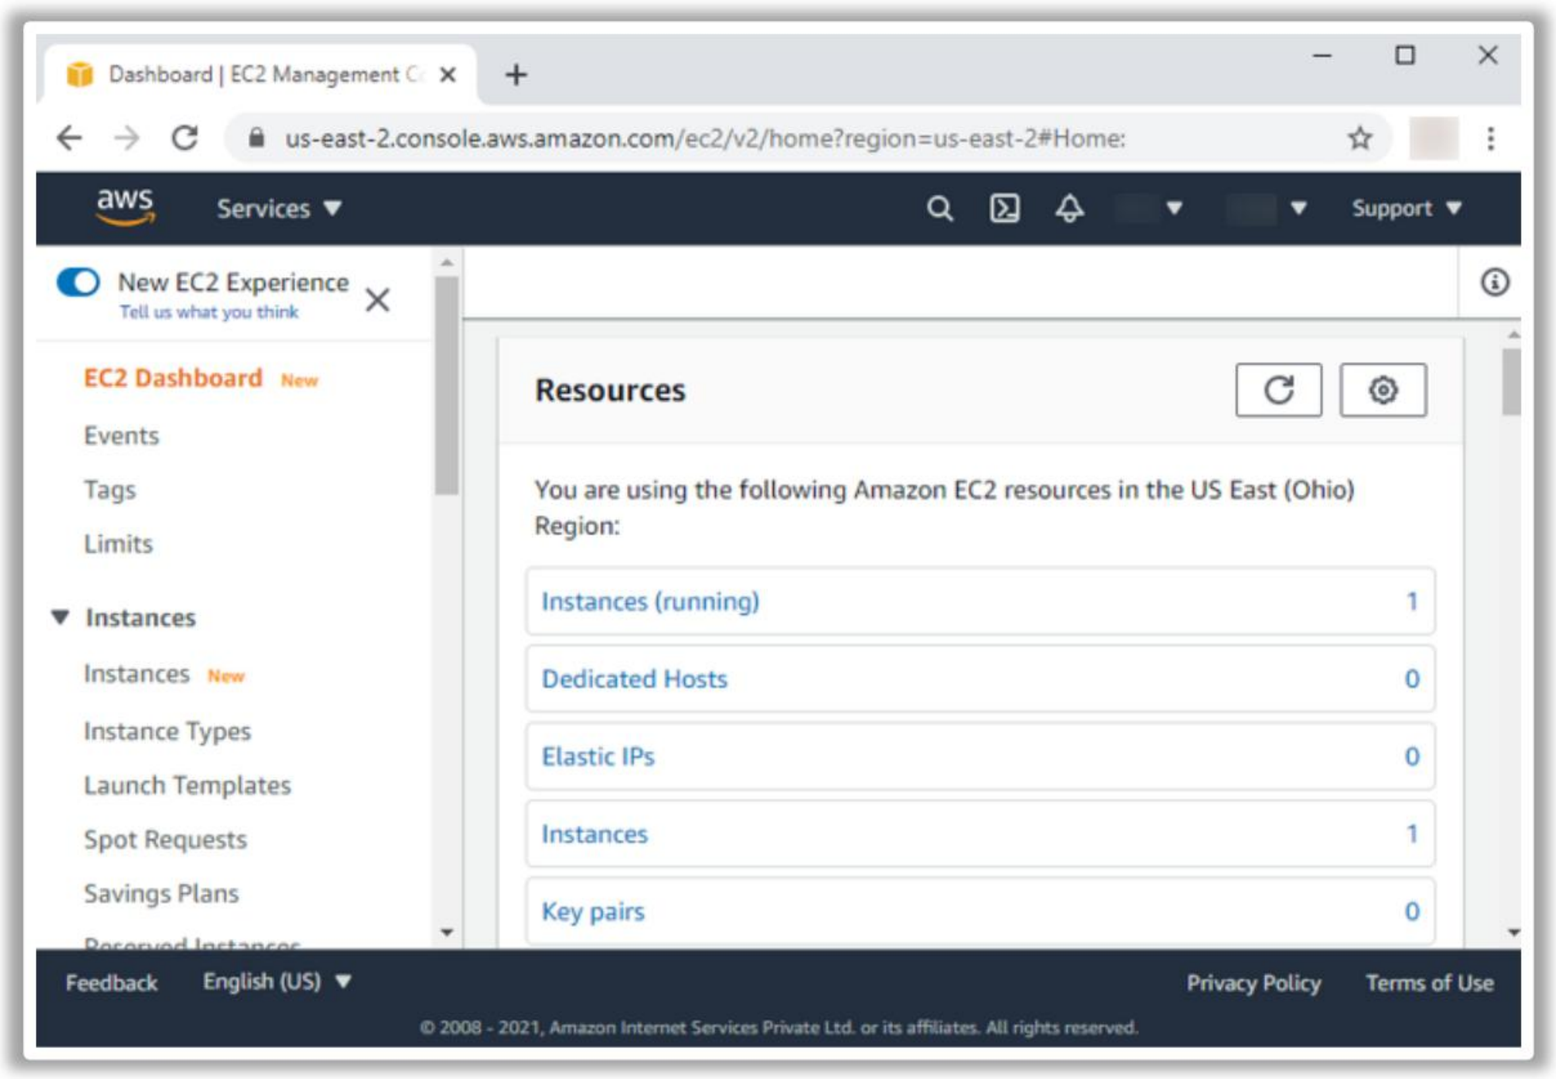

Figure 9.64: EC2 Dashboard Console

- In the navigation pane, under **Elastic Block Store**, click on **Snapshots**.

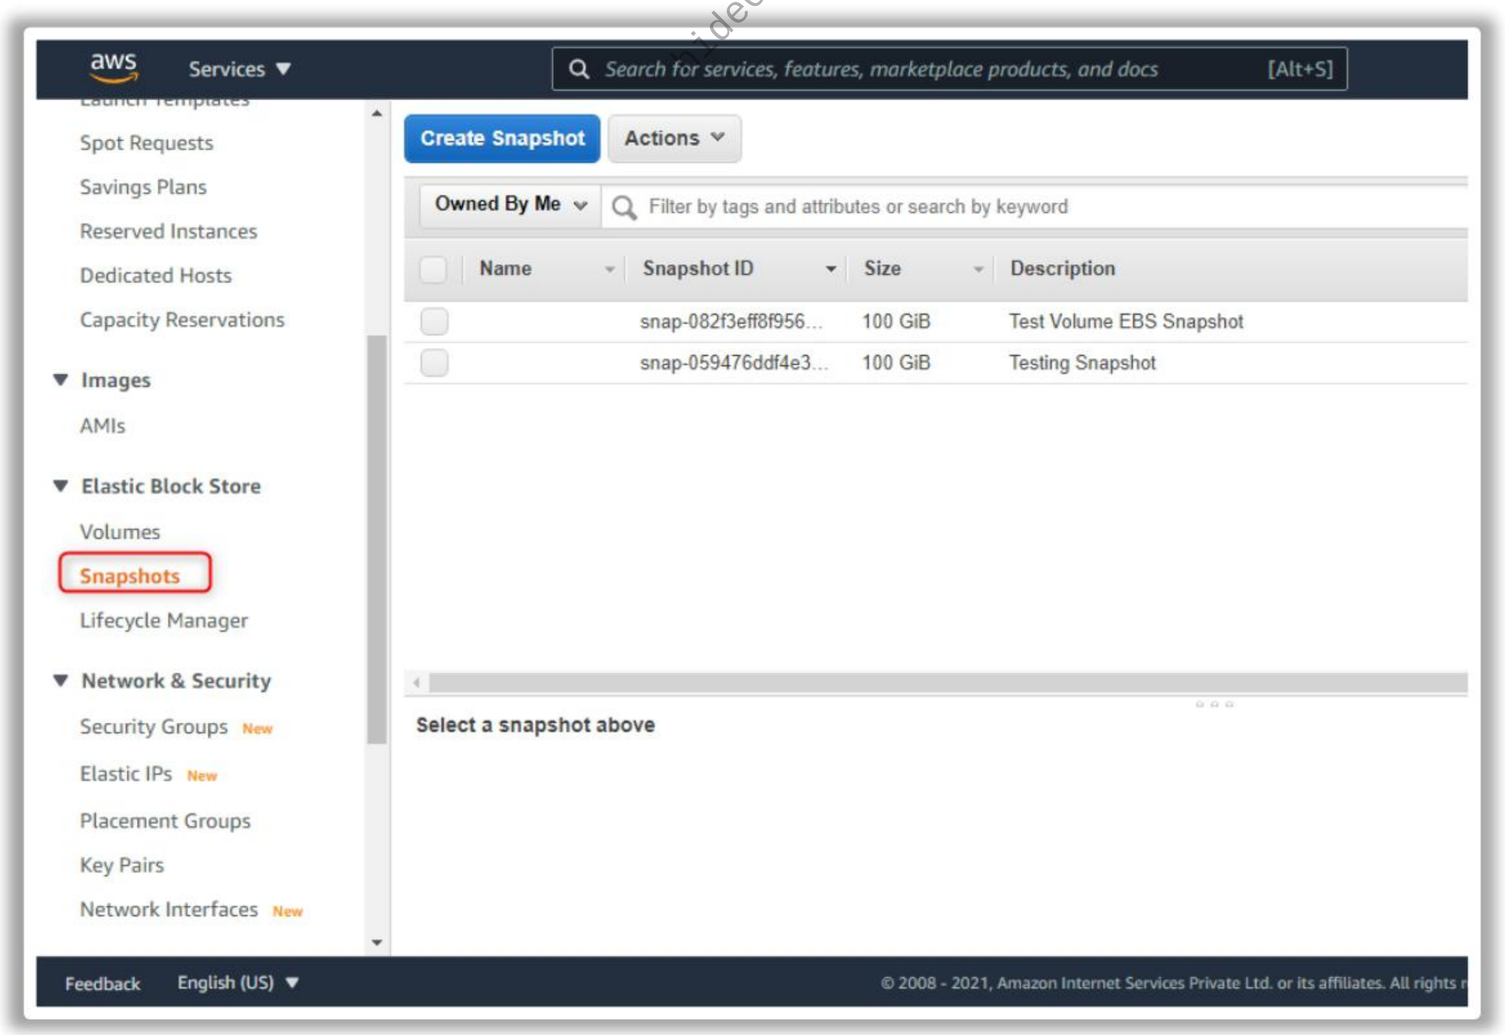

Figure 9.65: Create Snapshot Console

Certified Cloud Security Engineer Copyright © by **EC-Council**
All Rights Reserved. Reproduction is Strictly Prohibited.

- Choose the EBS volume snapshot that you want to examine. In the bottom panel, click on the **Description** tab. Near **Started** parameter, check the time and date of the last snapshot. If the snapshot of the volume was created 7 days ago, then the user should take a new snapshot.

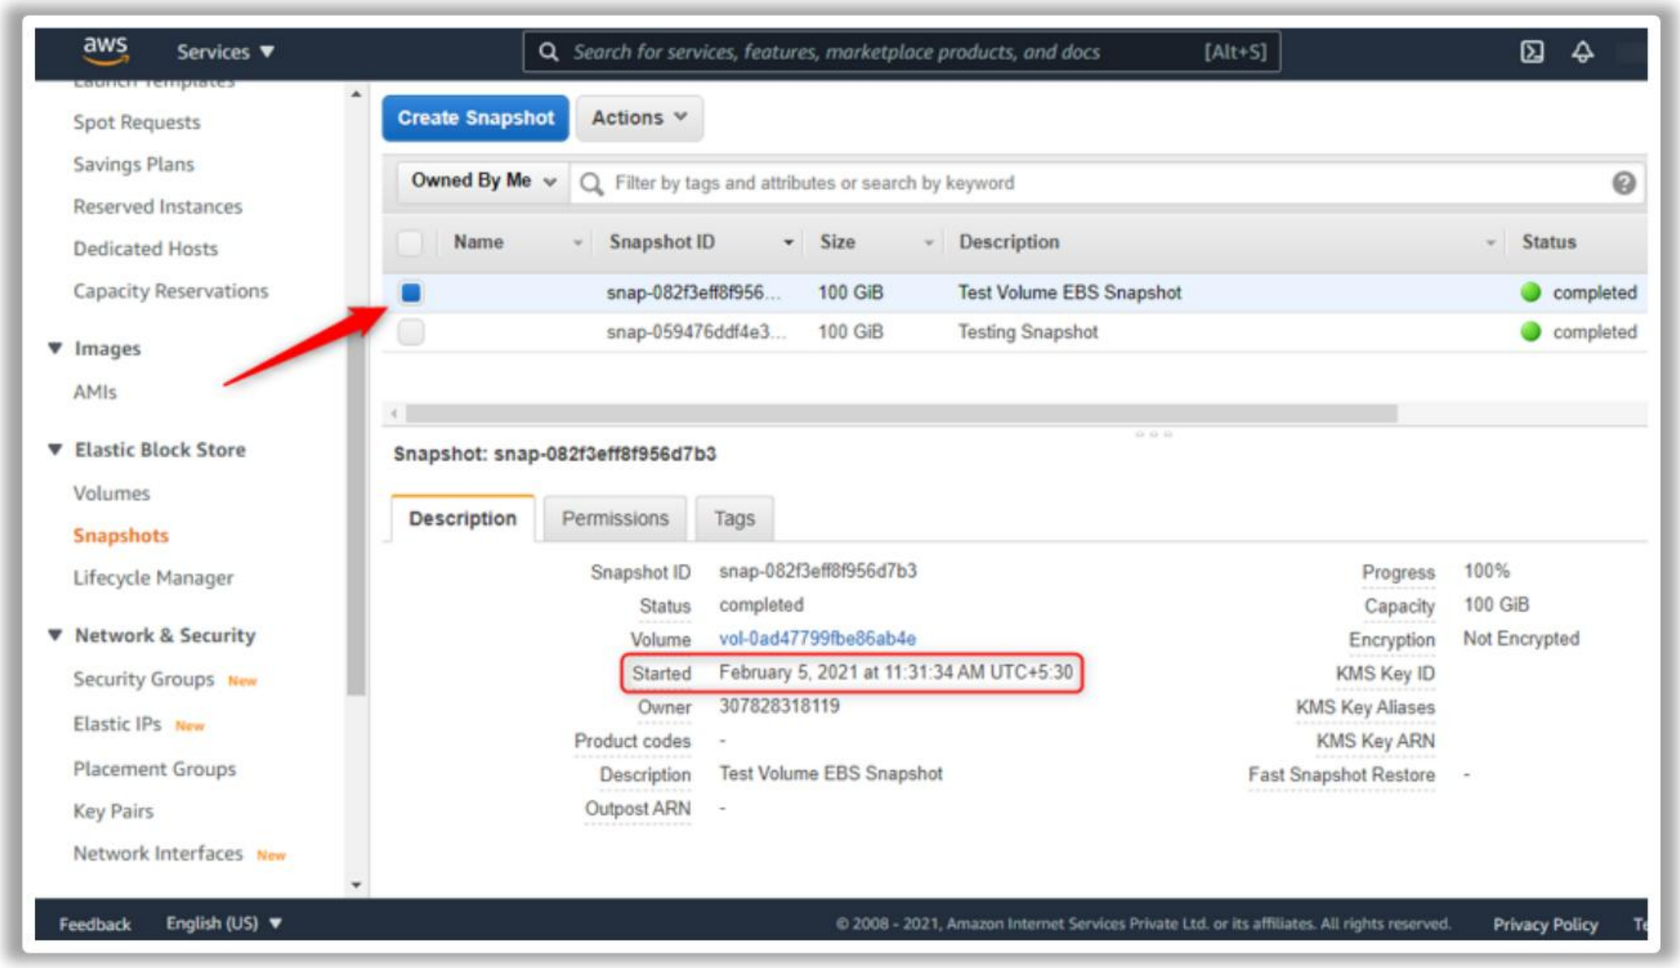

Figure 9.66: Verifying time and date of last snapshot

hide01.ir

Certified Cloud Security Engineer Copyright © by **EC-Council**
All Rights Reserved. Reproduction is Strictly Prohibited.

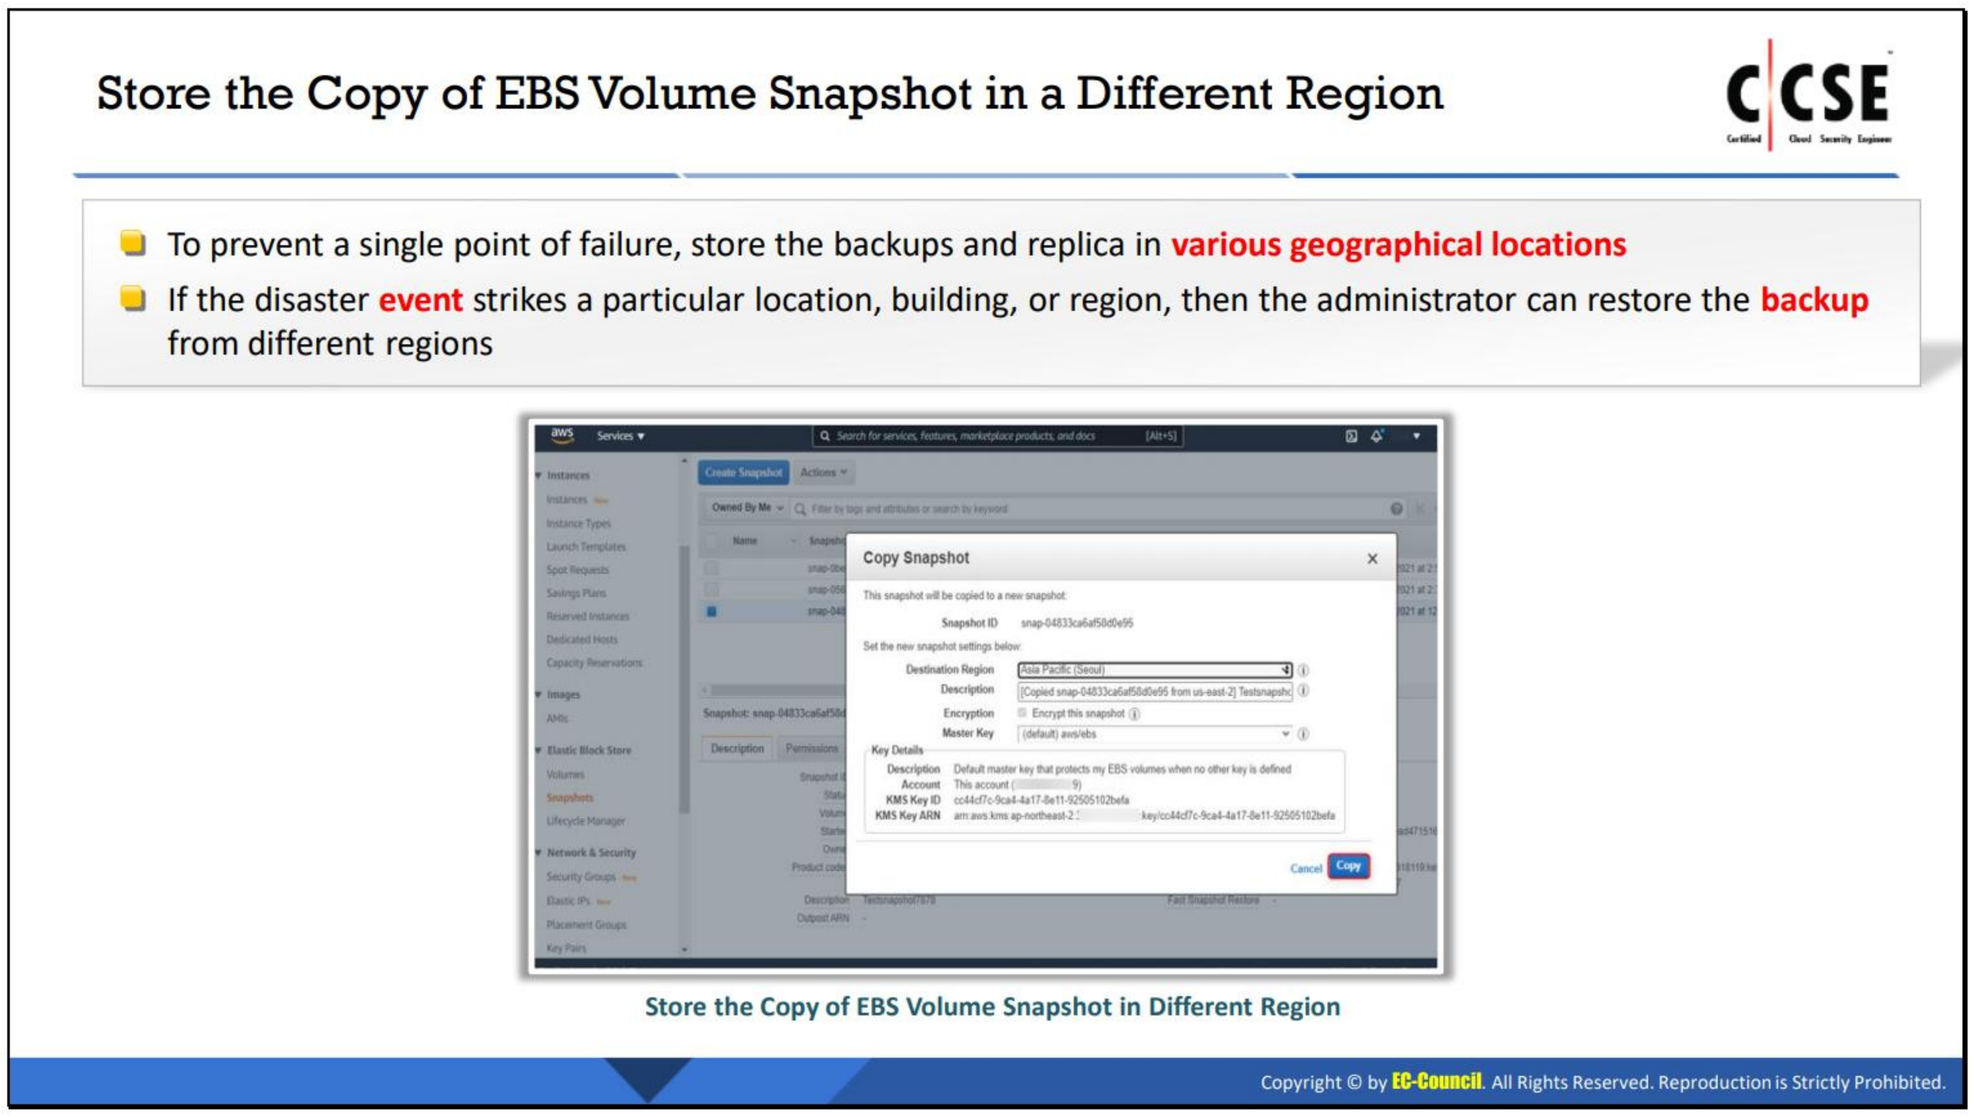

Store the Copy of EBS Volume Snapshot in a Different Region

To prevent a single point of failure, the cloud security professional should store the backups and replica in various geographical locations. If the disaster event strikes a particular location, building, or region, then the administrator can restore the backup from the different regions.

To Store the Copy of EBS volume Snapshot in a Different Region

- Sign in to AWS Management Console.

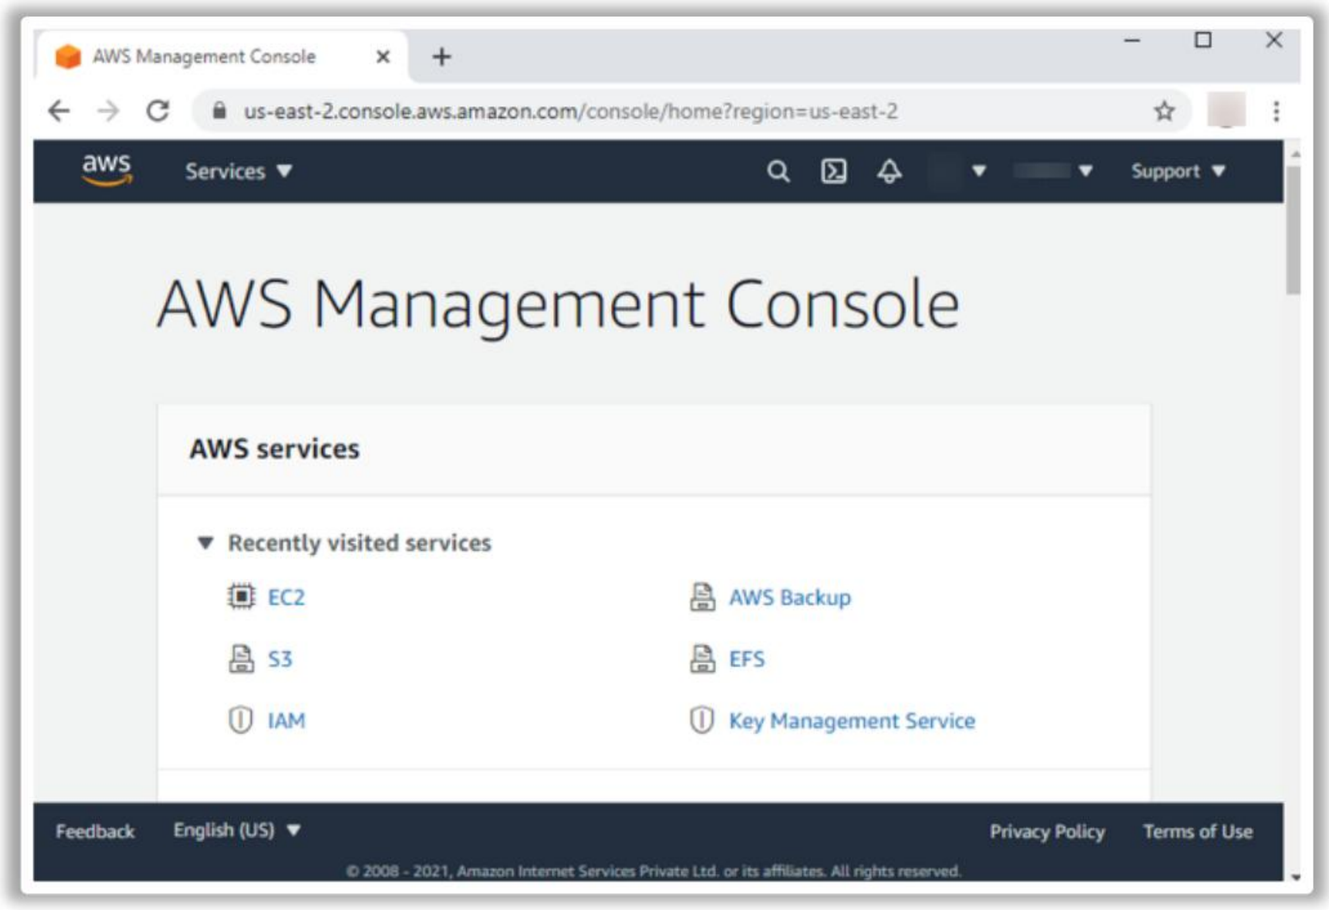

Figure 9.67: AWS Management Console

Certified Cloud Security Engineer Copyright © by **EC-Council**
All Rights Reserved. Reproduction is Strictly Prohibited.

- Click on the following link to open EC2 dashboard: https://console.aws.amazon.com/ec2/

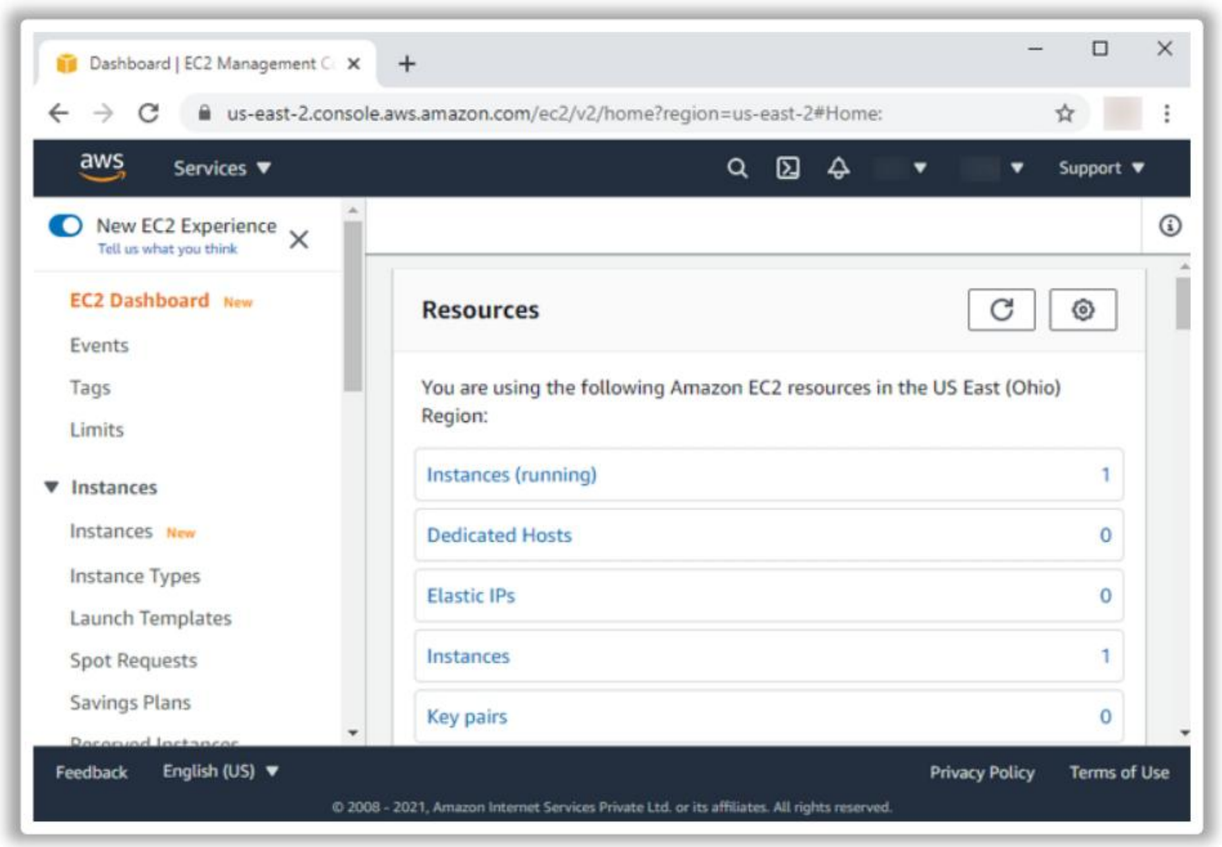

Figure 9.68: EC2 Dashboard Console

- Under **Elastic Block Store**, click on Snapshots; select the snapshot.

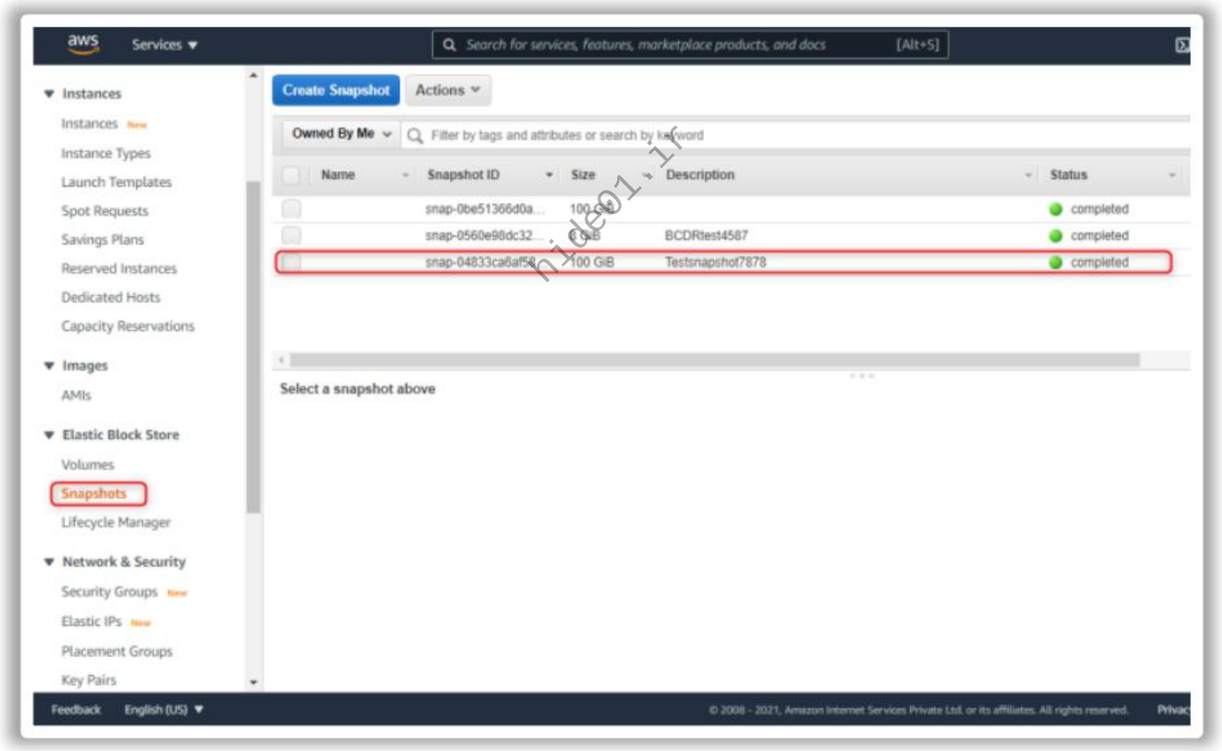

Figure 9.69: Create Snapshot Console

- Click on the **Actions** drop-down; then click on **Copy**.

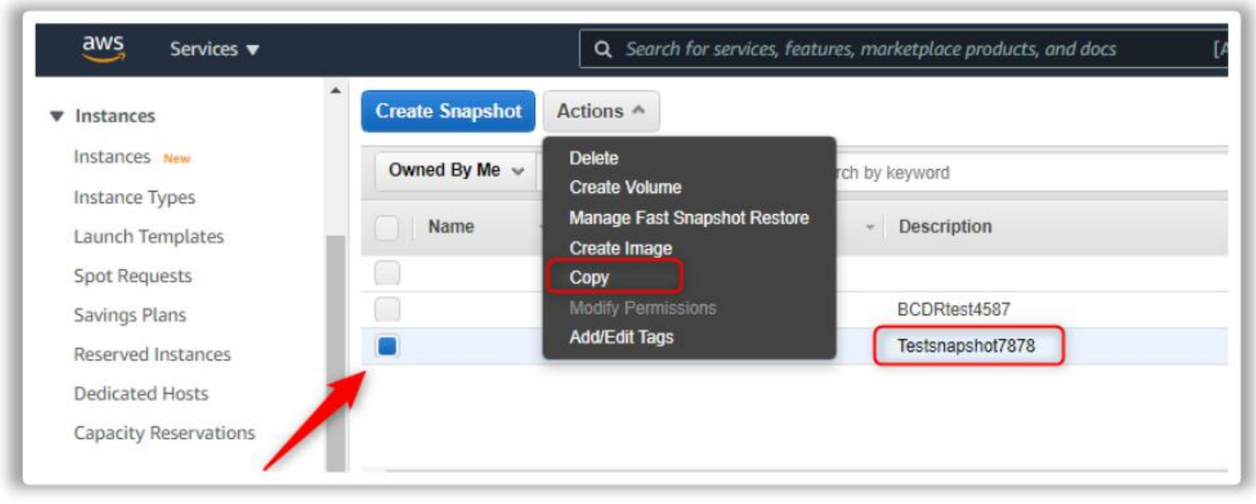

Figure 9.70: Click Action Drop-down

Certified Cloud Security Engineer Copyright © by **EC-Council**
All Rights Reserved. Reproduction is Strictly Prohibited.

- Copy Snapshot window will appear; select the destination region from the drop-down.

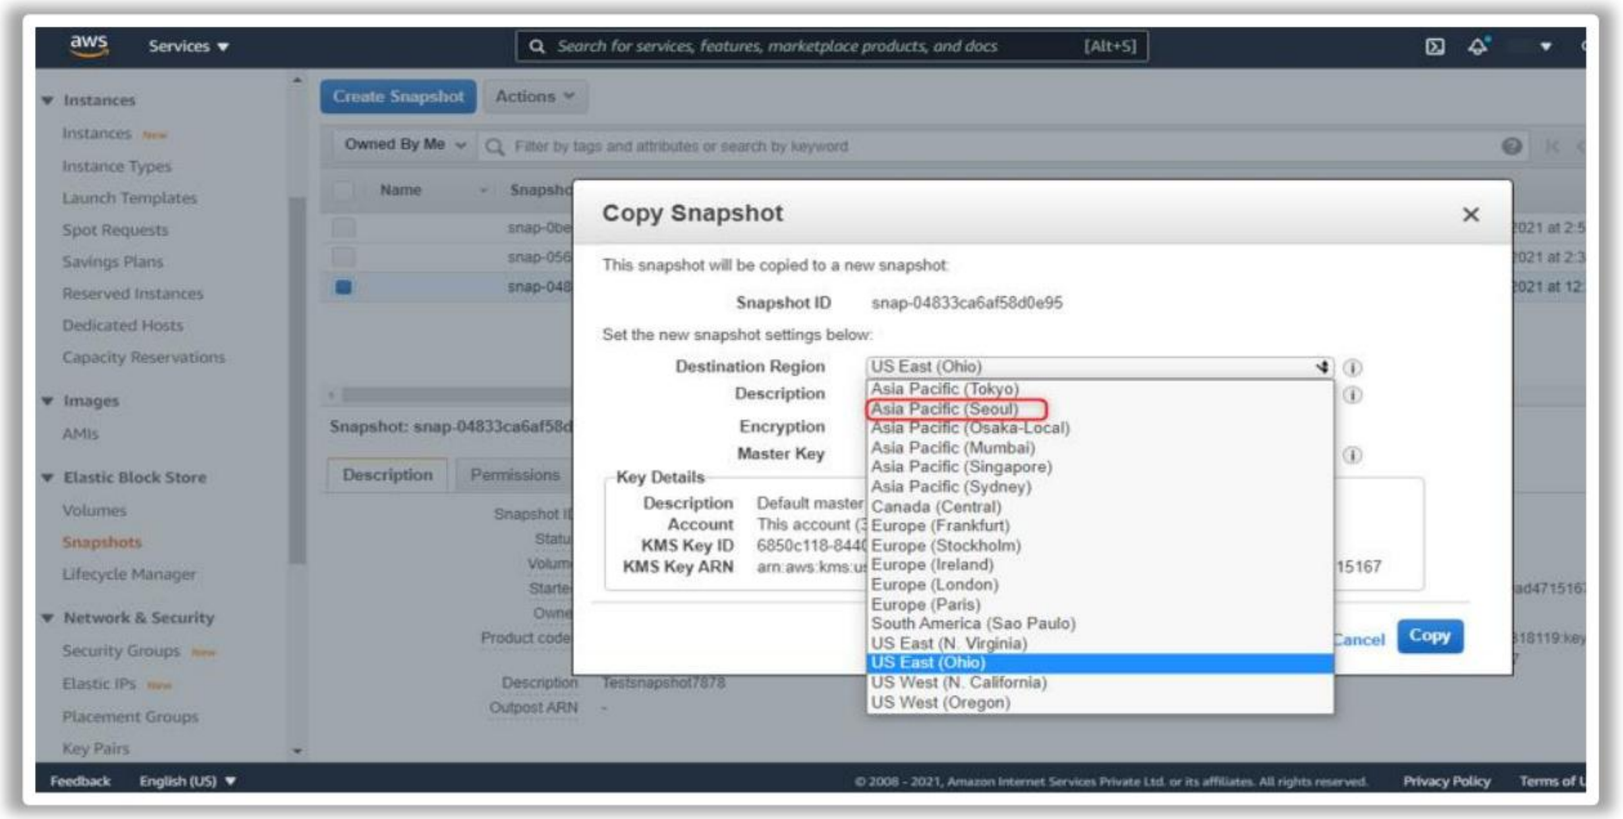

Figure 9.71: Select the Destination Region

- Then, click on the **Copy** button.

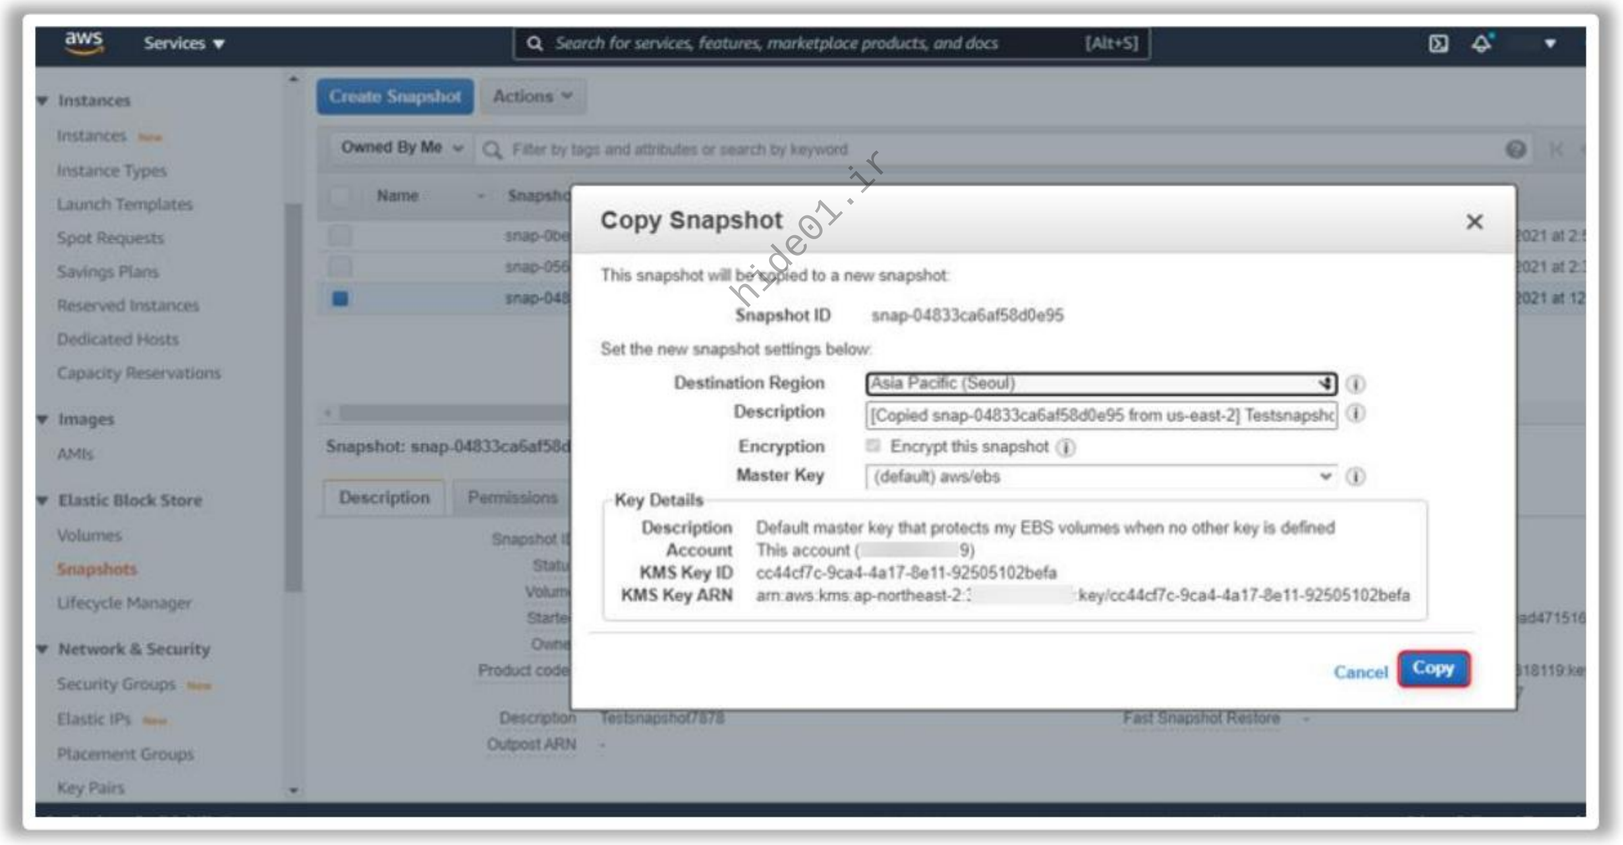

Figure 9.72: Click Copy

Certified Cloud Security Engineer Copyright © by **EC-Council**
All Rights Reserved. Reproduction is Strictly Prohibited.

- You will get a message indicating that the **Snapshot copy operation has been initiated**. Wait for few seconds and then click on the hyperlink **Snapshots** mentioned in the next line.

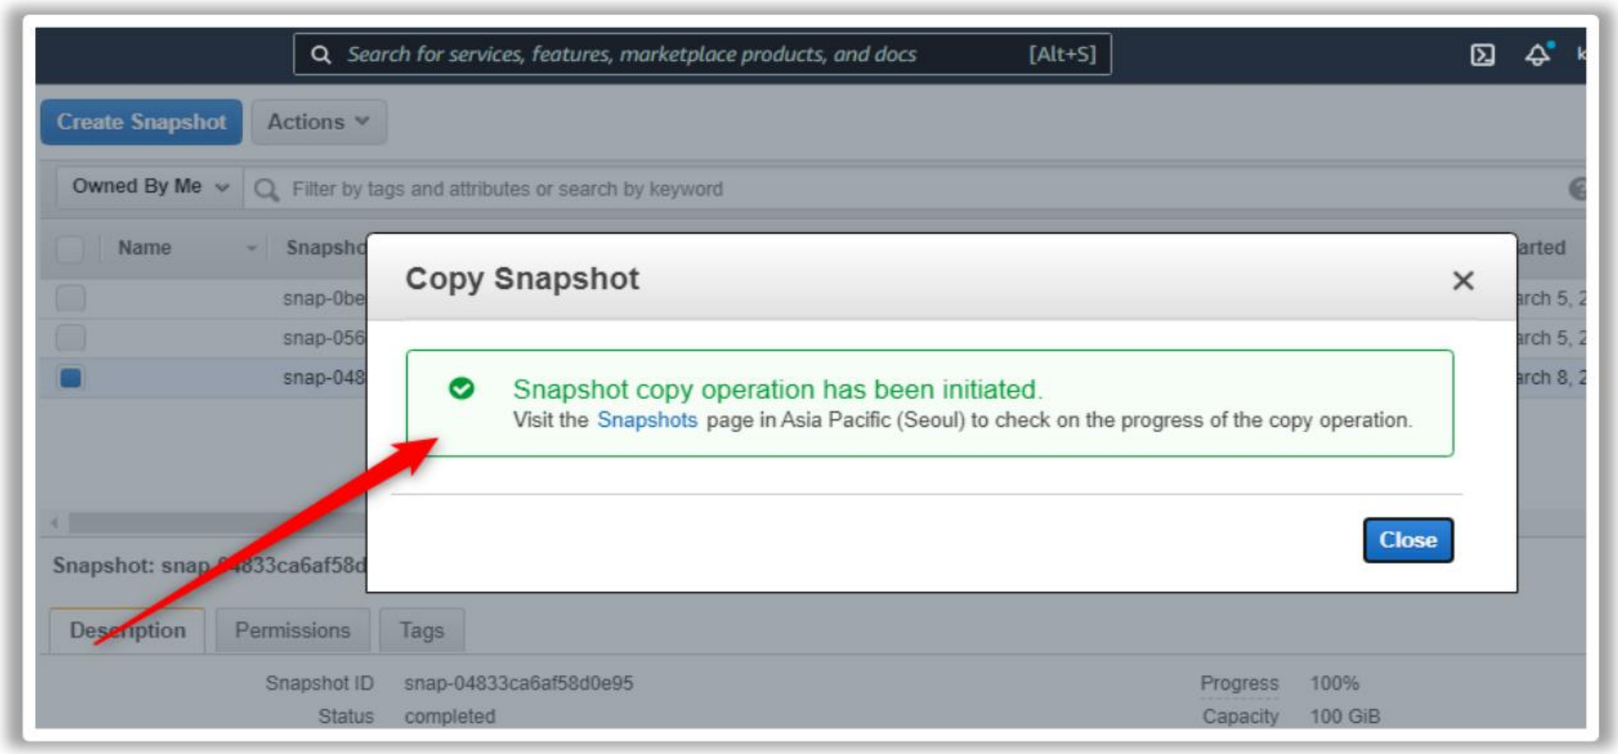

Figure 9.73: Select Close

- You will be redirected to a different region where your snapshot is successfully stored.

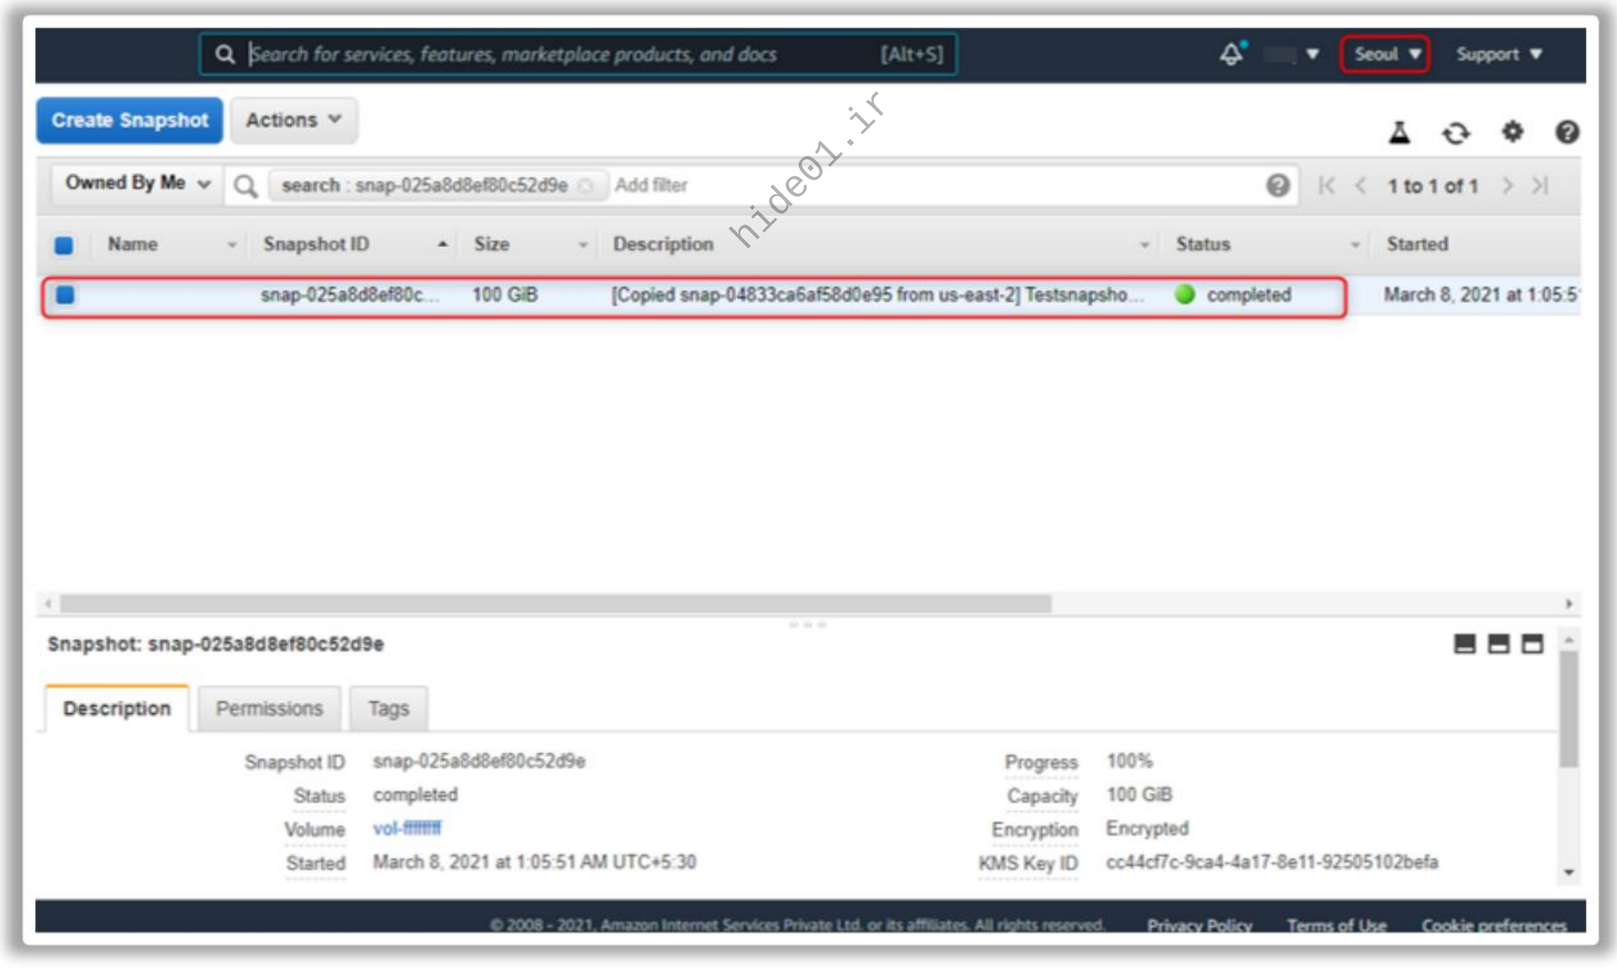

Figure 9.74: Snapshot Console

Certified Cloud Security Engineer Copyright © by **EC-Council**
All Rights Reserved. Reproduction is Strictly Prohibited.

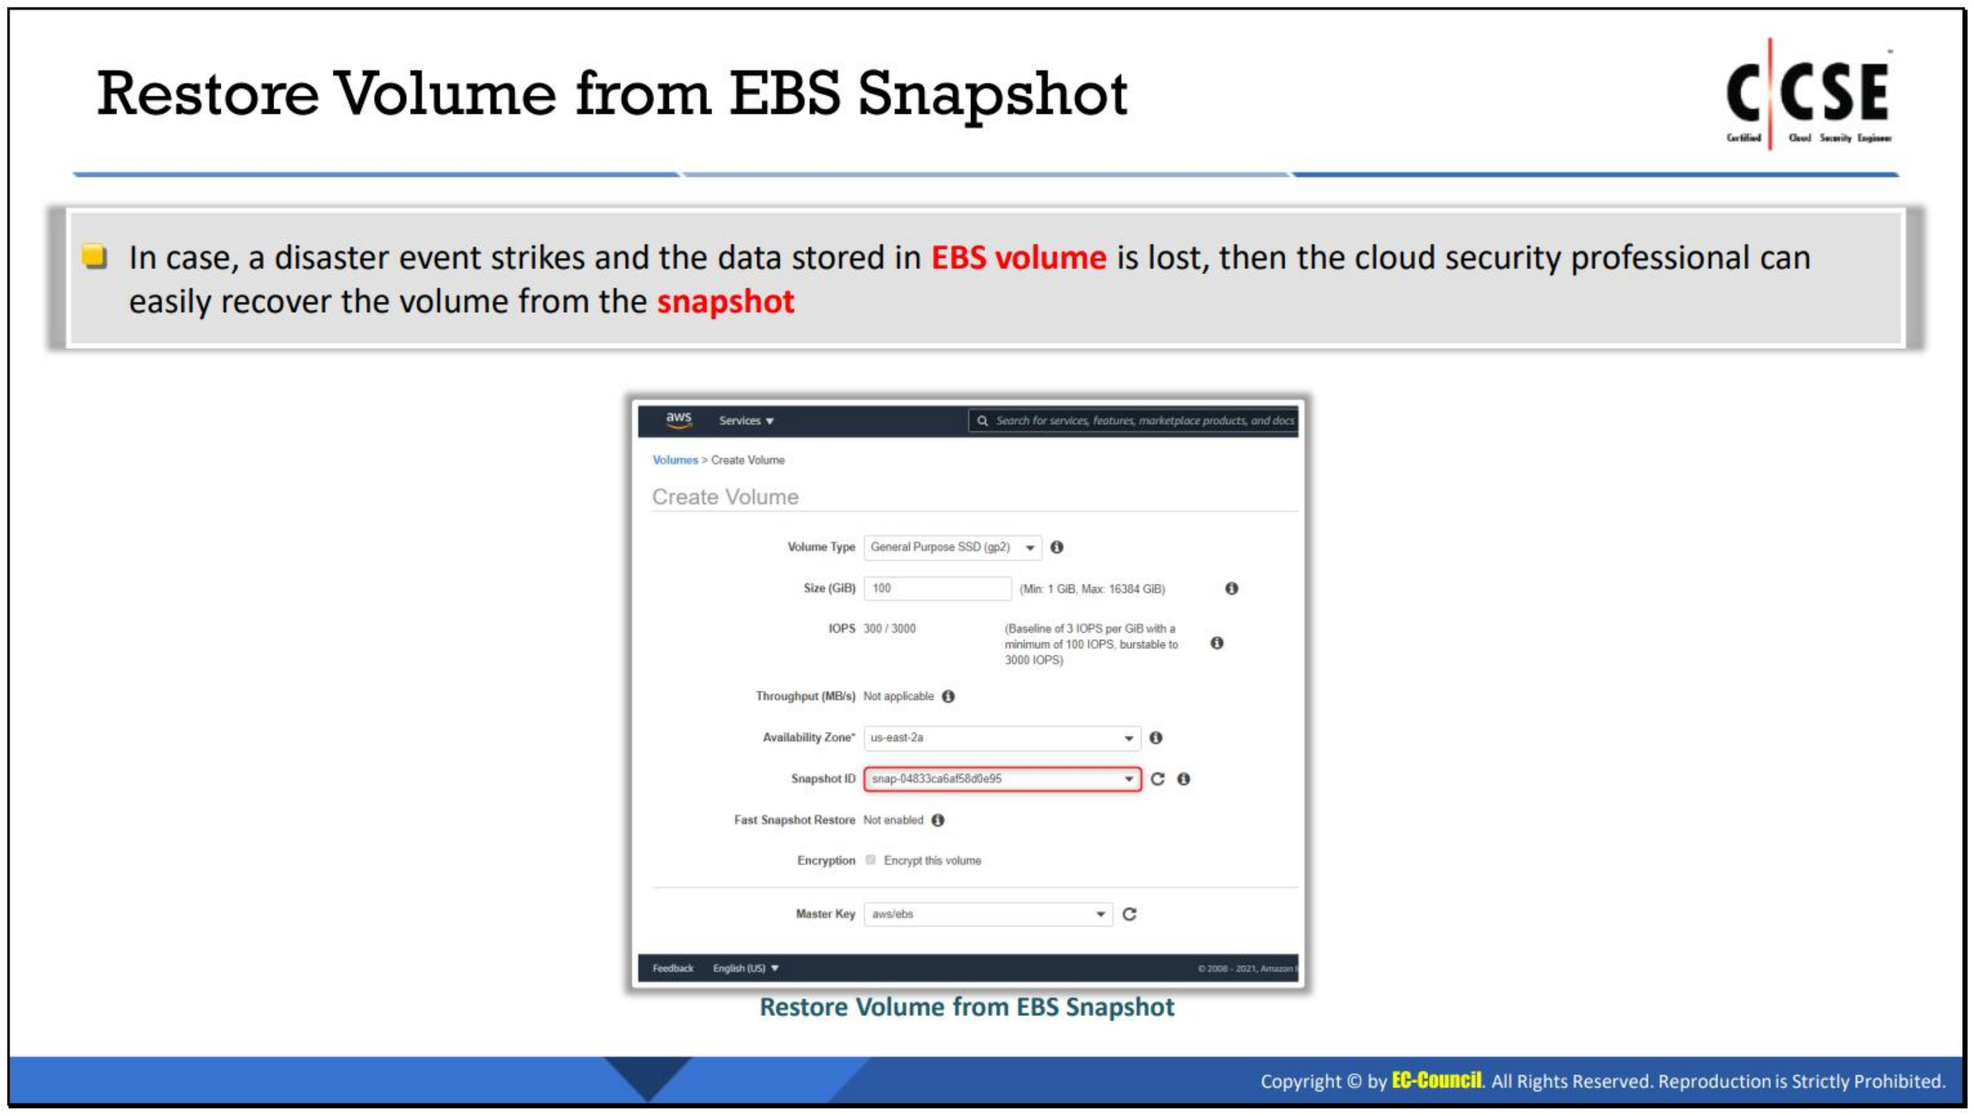

Restore Volume from EBS Snapshot

Restore Volume from EBS Snapshot

In case a disaster event strikes and the block data stored in EBS volume is lost, then the cloud security professional can easily recover the data from the snapshot. Using the snapshots stored in the same or different AZ, the cloud security professional can easily restore the entire volume from EBS Snapshot.

To Restore Volume from EBS Snapshot

- Sign in to AWS Management Console.

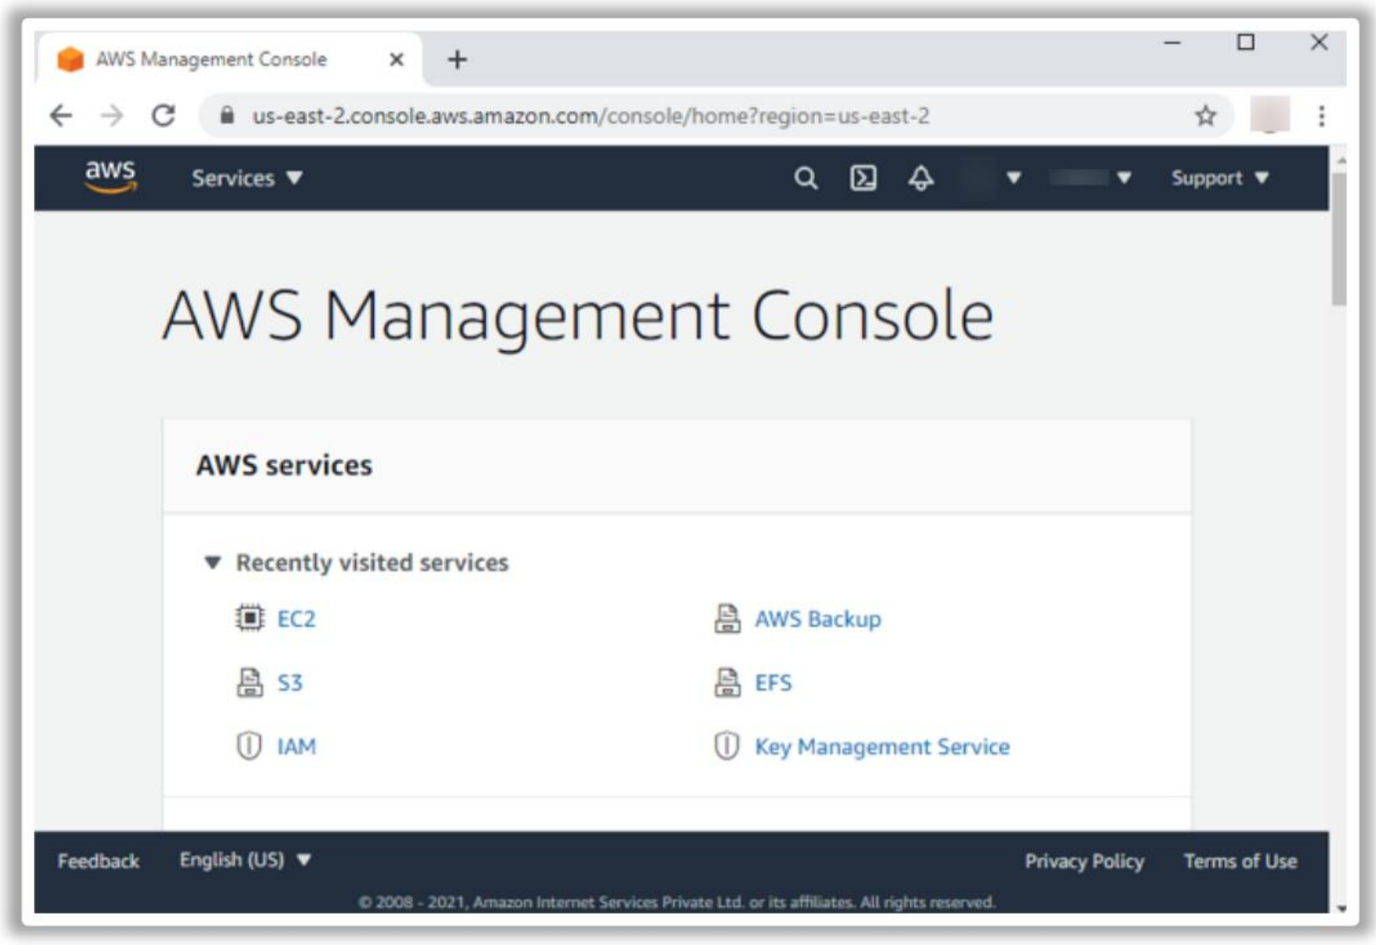

Figure 9.75: AWS Management Console

Certified Cloud Security Engineer Copyright © by **EC-Council**
All Rights Reserved. Reproduction is Strictly Prohibited.

- Click on the following link to open EC2 dashboard: https://console.aws.amazon.com/ec2/

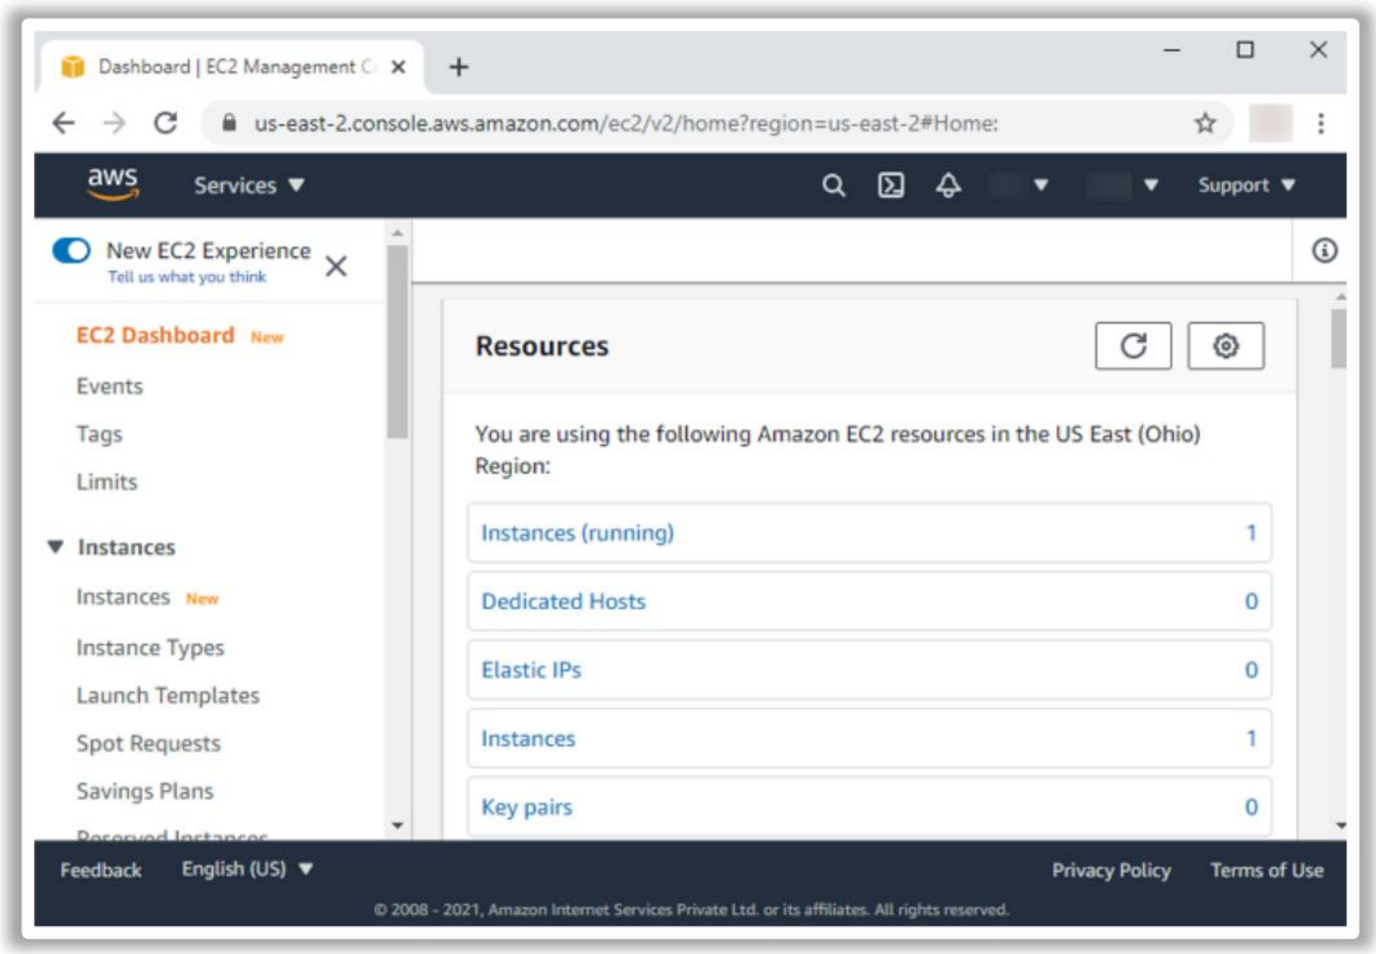

Figure 9.76: EC2 Dashboard Console

- Under **Elastic Block Store**, click on **Snapshots** and copy the **Snapshot ID** to restore the volume related to that snapshot.

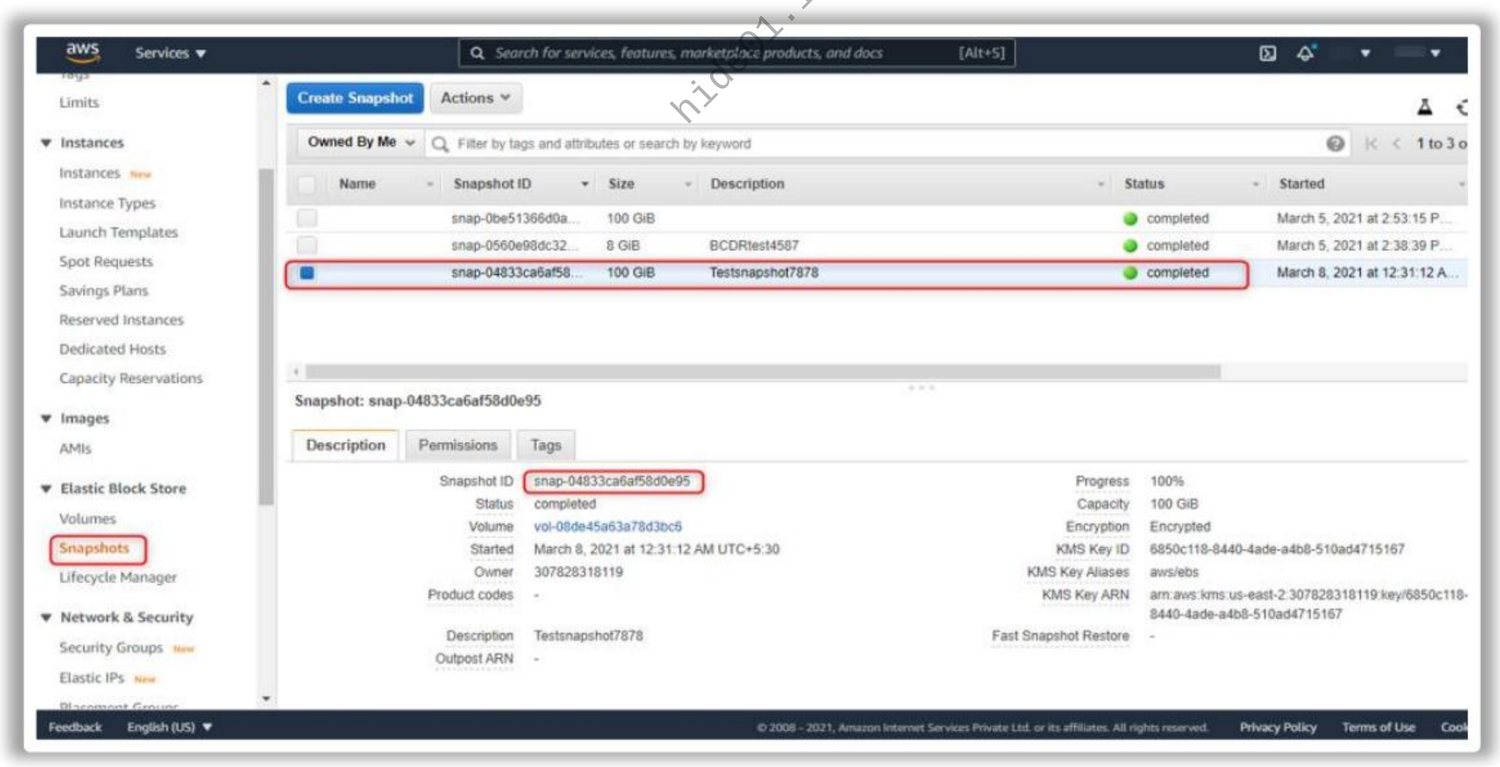

Figure 9.77: Snapshot Console

Certified Cloud Security Engineer Copyright © by **EC-Council**
All Rights Reserved. Reproduction is Strictly Prohibited.

- Under **Elastic Block Store**, click on **Volume** and then click on **Create Volume** button.

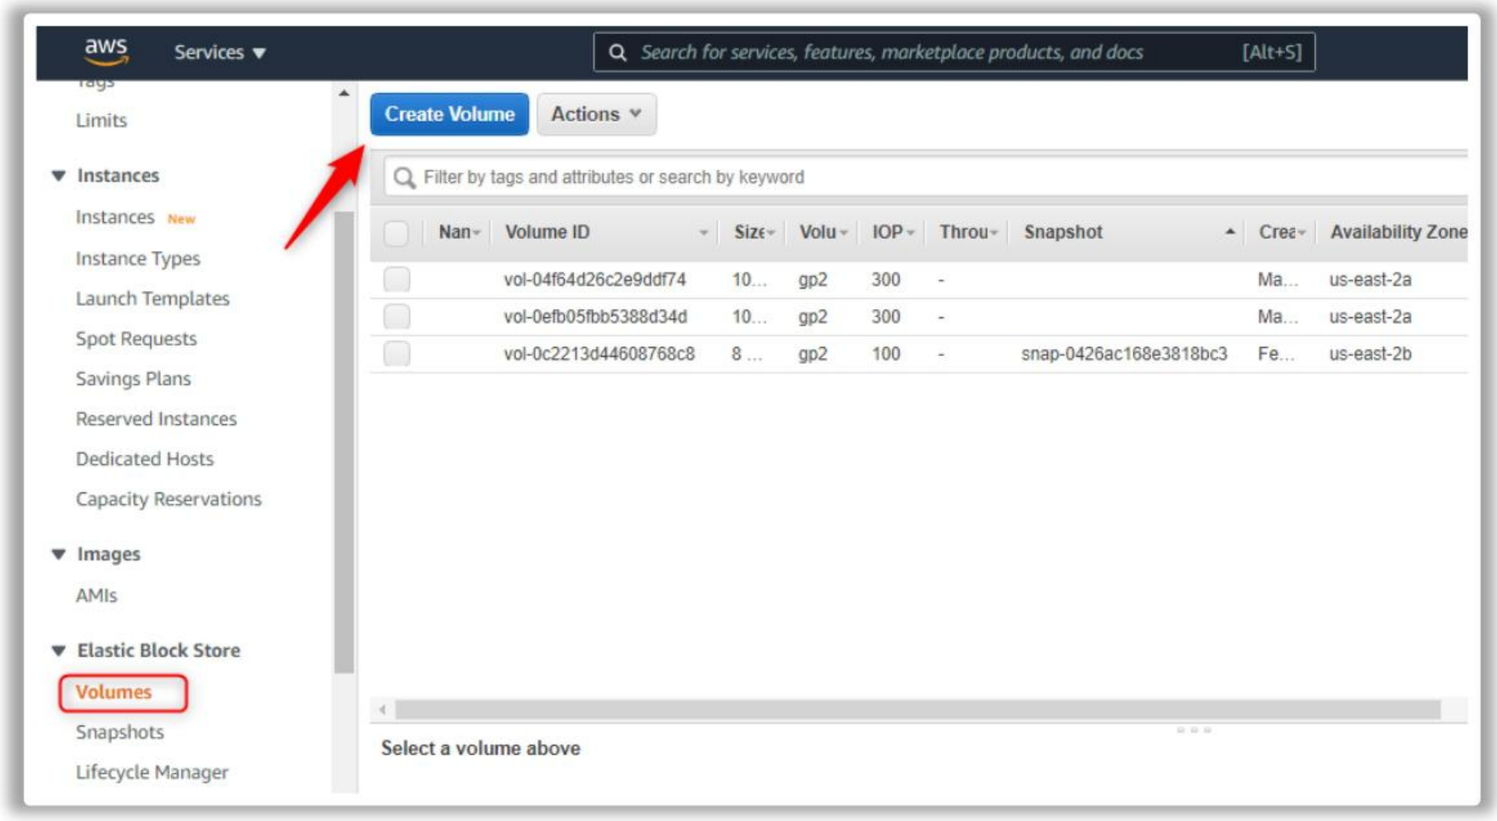

Figure 9.78: Volume Console

- In the **Snapshot ID** field, enter Snapshot ID. Scroll down and click on **Create Volume** button.

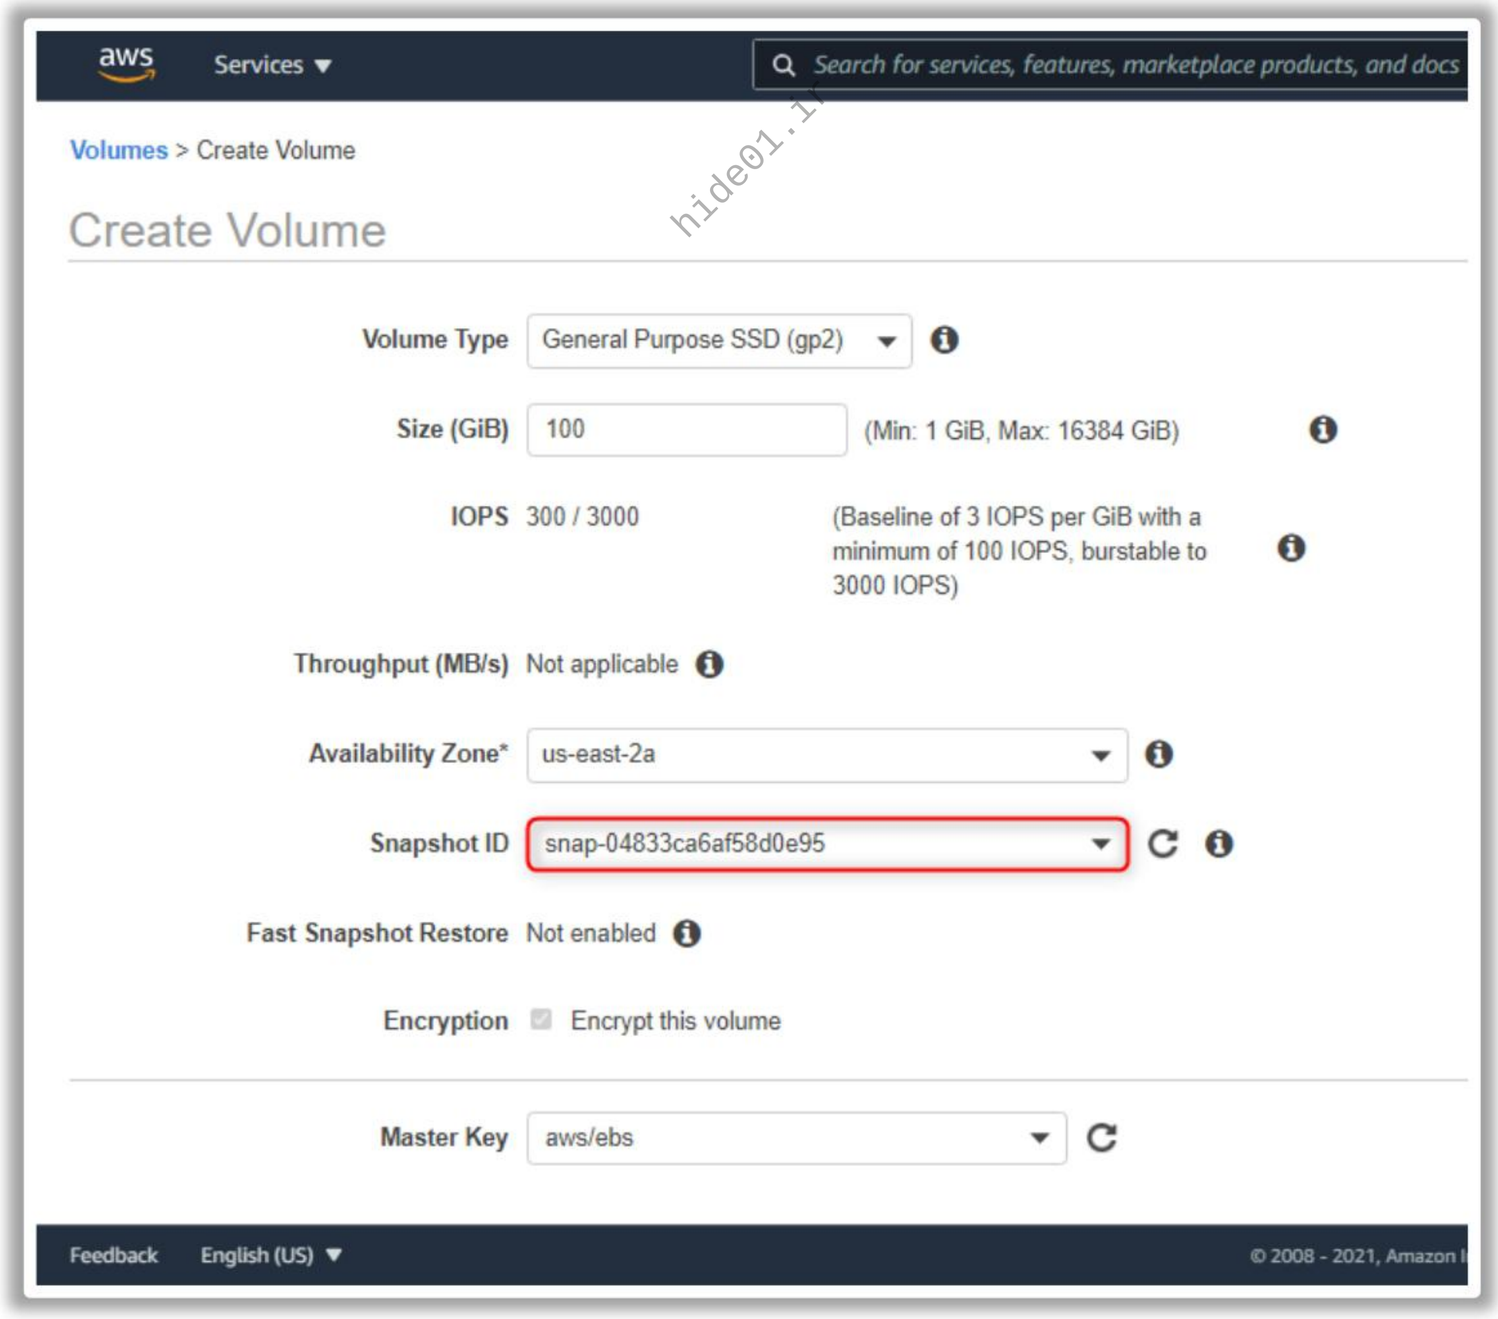

Figure 9.79: Create Volume Console

Certified Cloud Security Engineer Copyright © by **EC-Council**
All Rights Reserved. Reproduction is Strictly Prohibited.

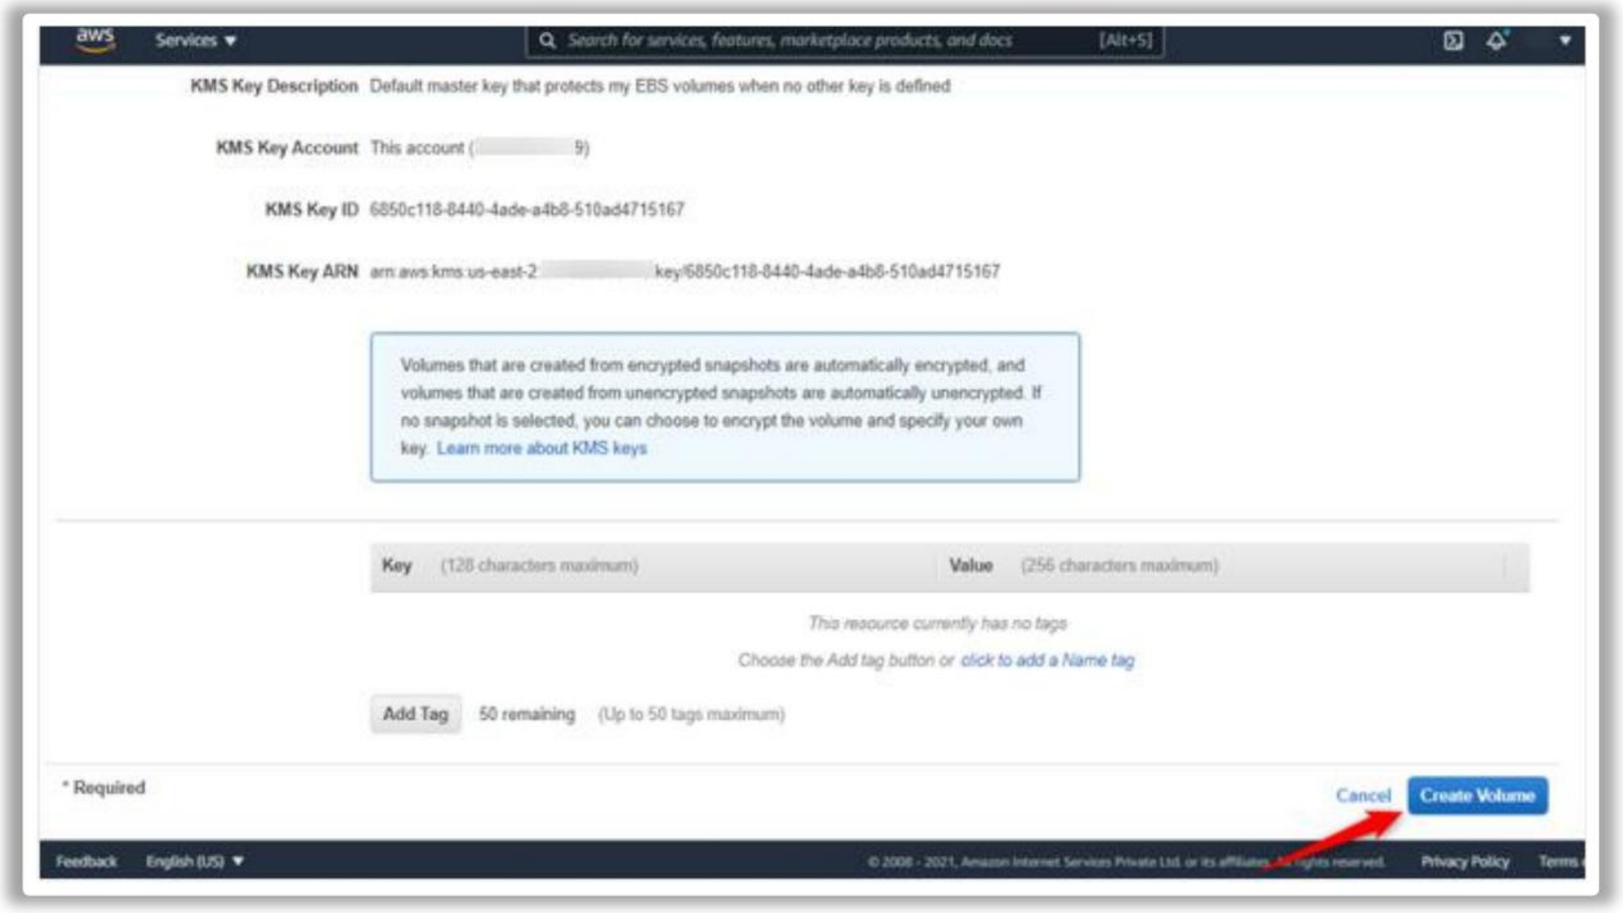

Figure 9.80: Creating Volume

- You will see a message indicating that the **volume is successfully created**. Click on the **Close** button.

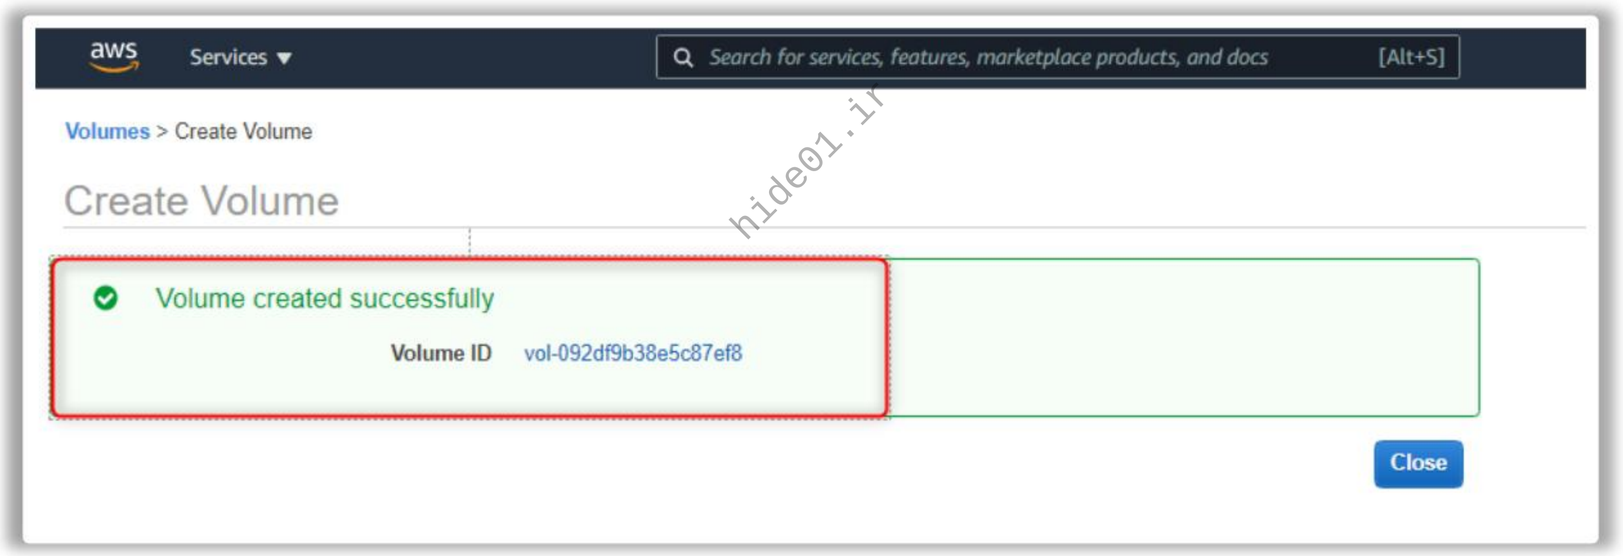

Figure 9.81: Volume Created Successfully

- Crosscheck the **snapshot ID.** Thus, the Amazon EBS volume is successfully created from the EBS snapshot.

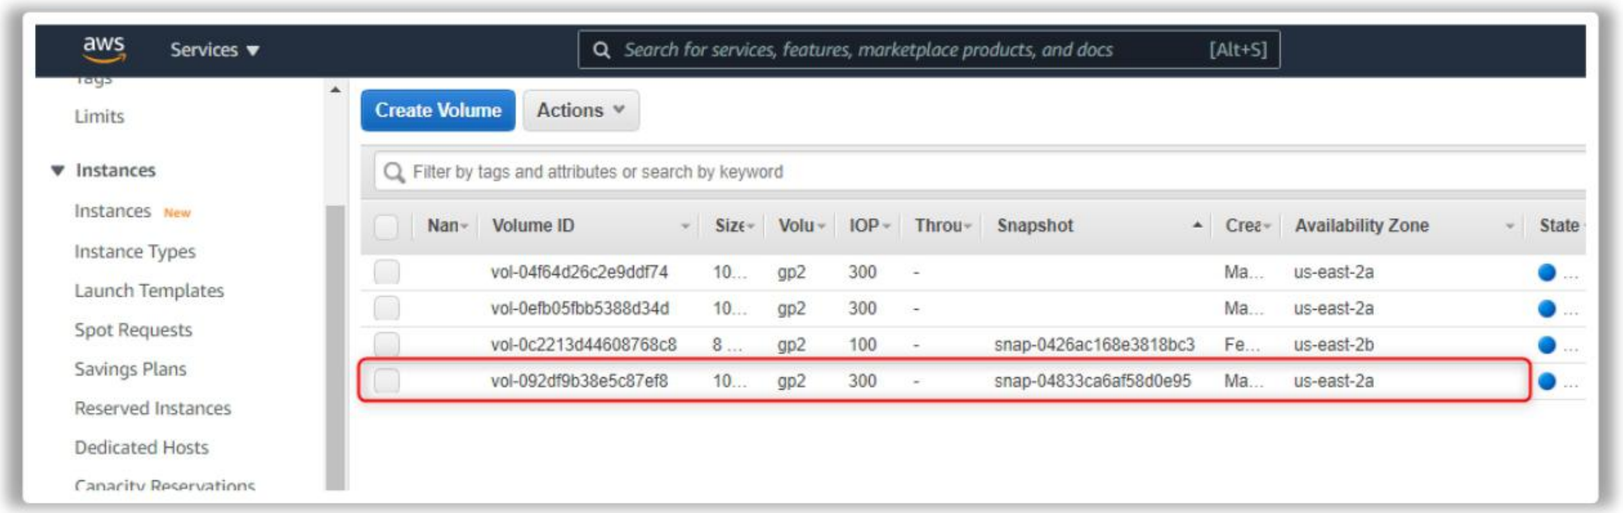

Figure 9.82: Creating Snapshot ID

Certified Cloud Security Engineer Copyright © by **EC-Council**
All Rights Reserved. Reproduction is Strictly Prohibited.

Enable Fast Snapshot Restore

- To implement disaster recovery and enhance performance:
  - Enable **Fast Snapshot Restore** on the existing or new snapshot in EBS on a per availability zone
  - Create a volume that can provide **robust performance** and does not require initialization

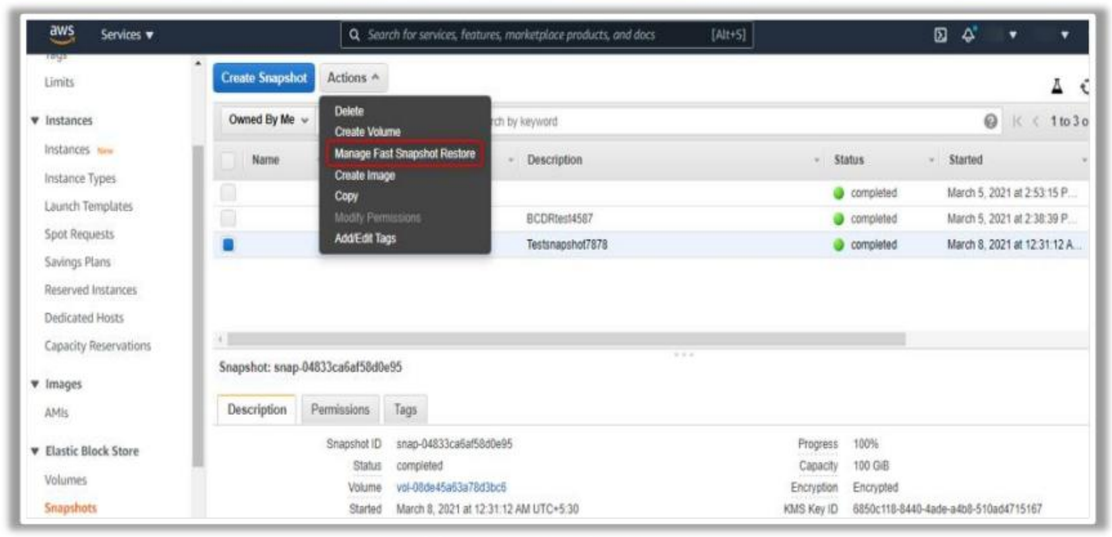

Enable Fast Snapshot Restore

Copyright © by EC-Council. All Rights Reserved. Reproduction is Strictly Prohibited.

Enable Fast Snapshot Restore

The cloud security professional can enable Fast Snapshot Restore on the existing or new snapshot in EBS on a per AZ basis, and then create a volume that can provide robust performance and does not require initialization. This will not only help in disaster recovery but also enhances performance, which in turn helps build a faster and responsive AWS system.

To Enable Fast Snapshot Restore

- Sign in to AWS Management Console.

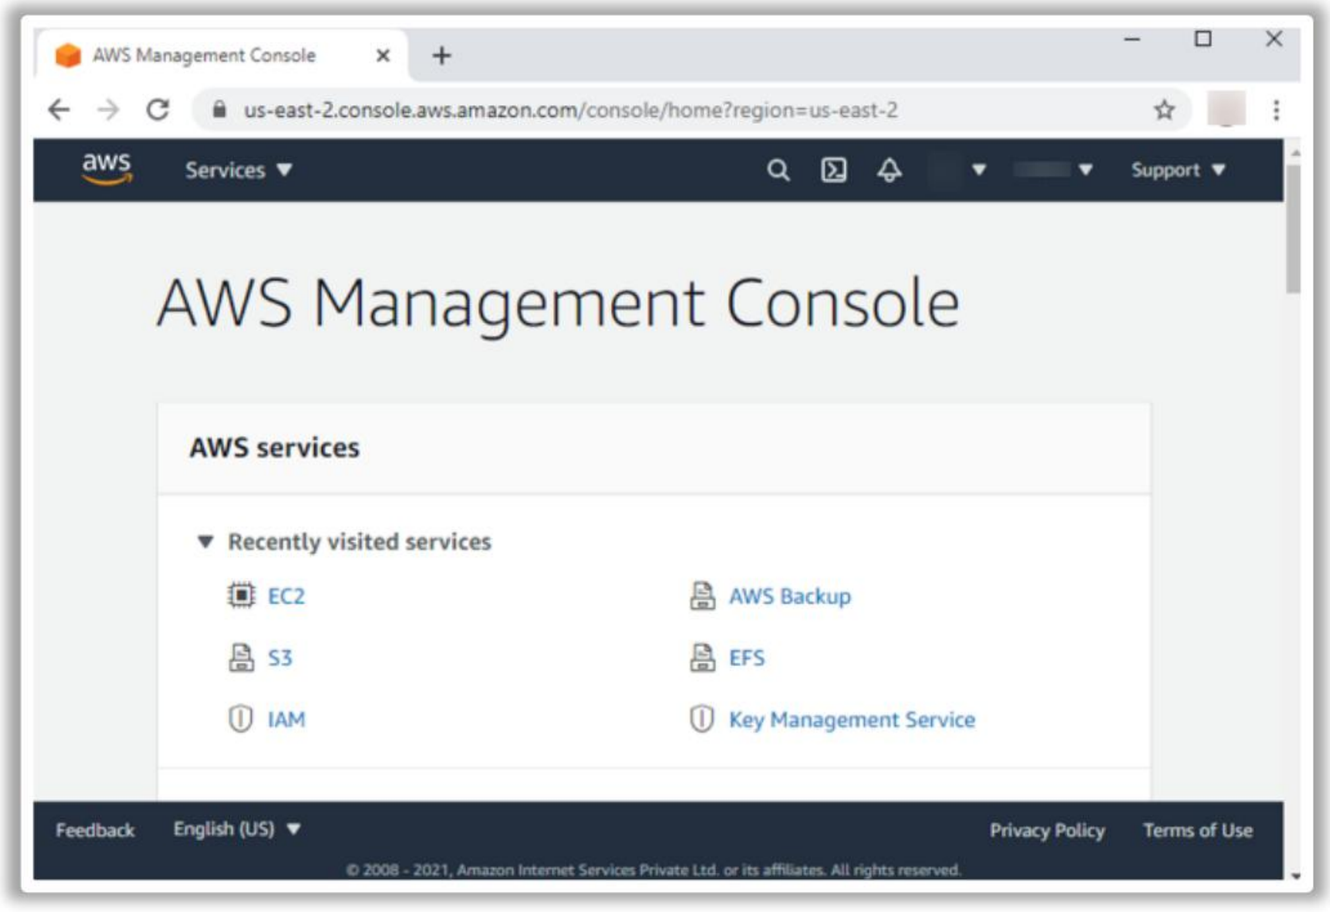

Figure 9.83: AWS Management Console

Certified Cloud Security Engineer Copyright © by **EC-Council**
All Rights Reserved. Reproduction is Strictly Prohibited.

- Click on the following link to open EC2 dashboard: https://console.aws.amazon.com/ec2/

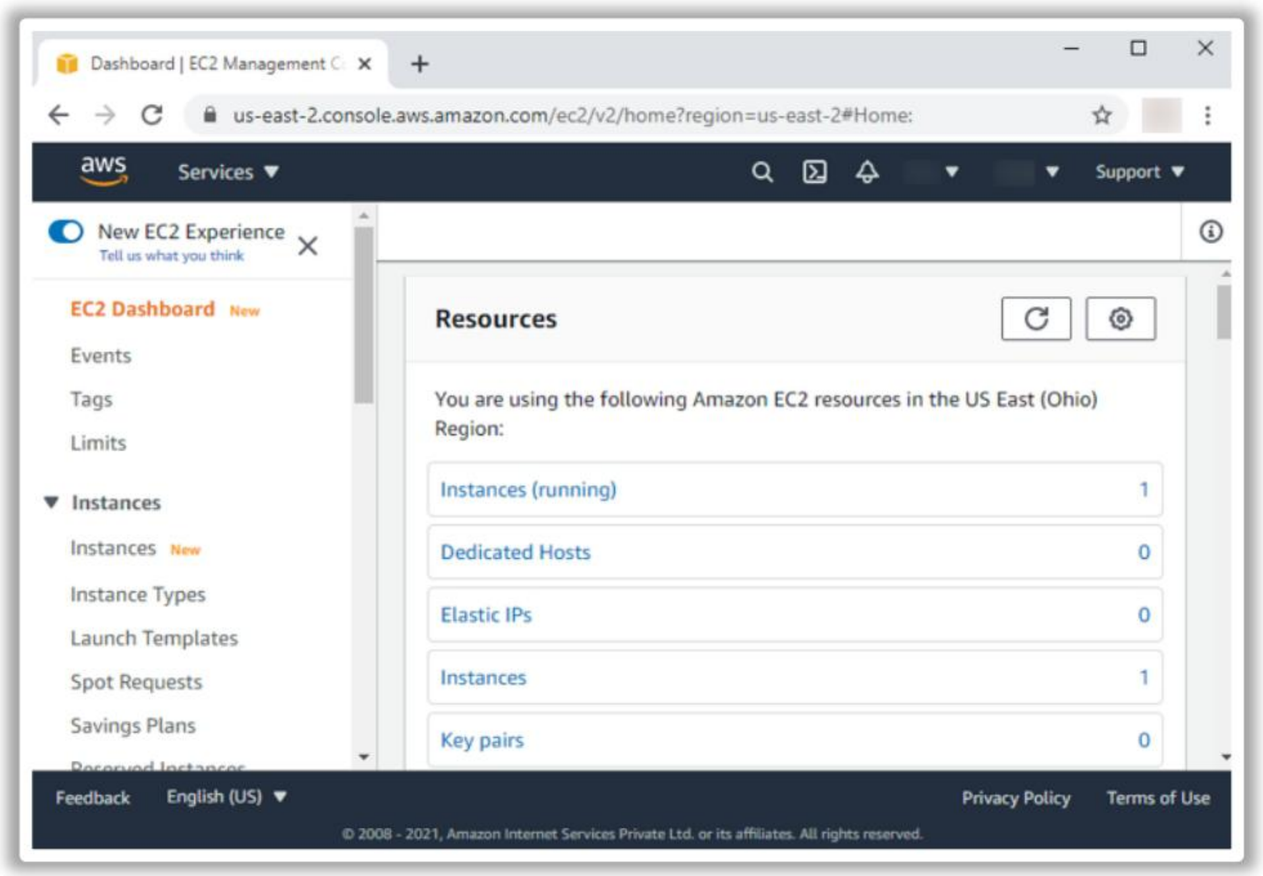

Figure 9.84: EC2 Dashboard Console

- Under **Elastic Block Store**, click on **Snapshots** and chose the snapshot that you want to set up for fast restore.

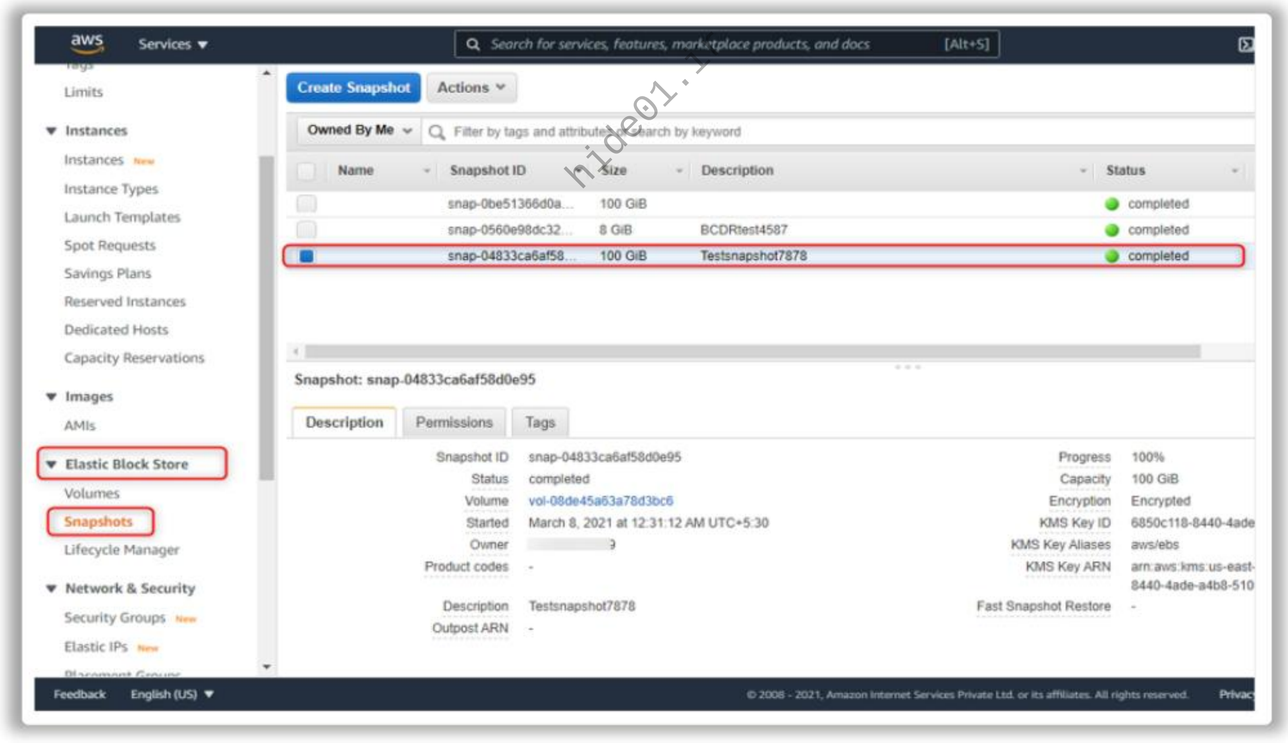

Figure 9.85: EC2 Snapshot Console

Certified Cloud Security Engineer Copyright © by **EC-Council**
All Rights Reserved. Reproduction is Strictly Prohibited.

- Click on the **Actions** drop and then click on **Manage Fast Snapshot Restore**.

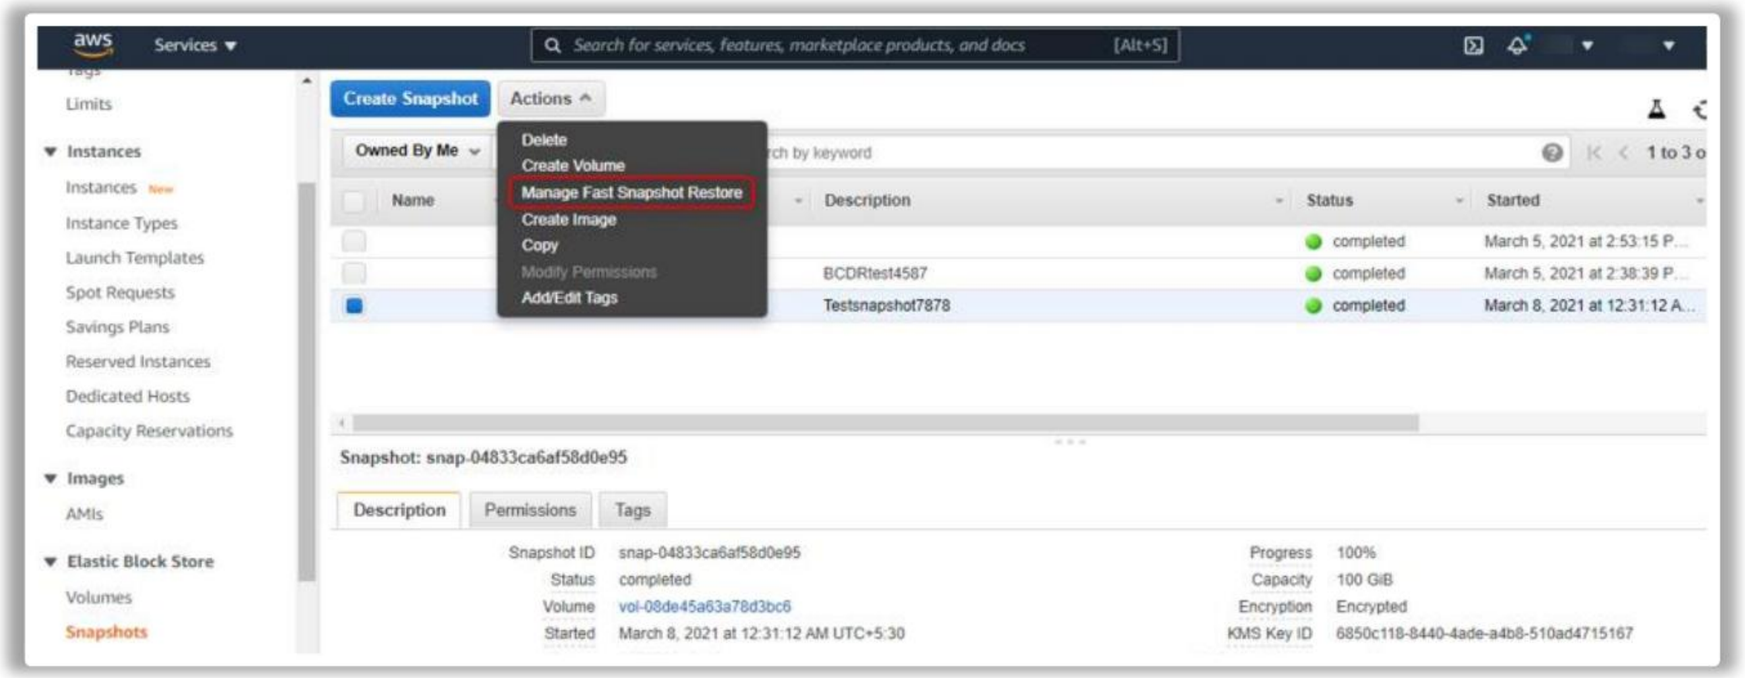

Figure 9.86: Click on Action Drop-box

- Select the region and then click on the **Save** button.

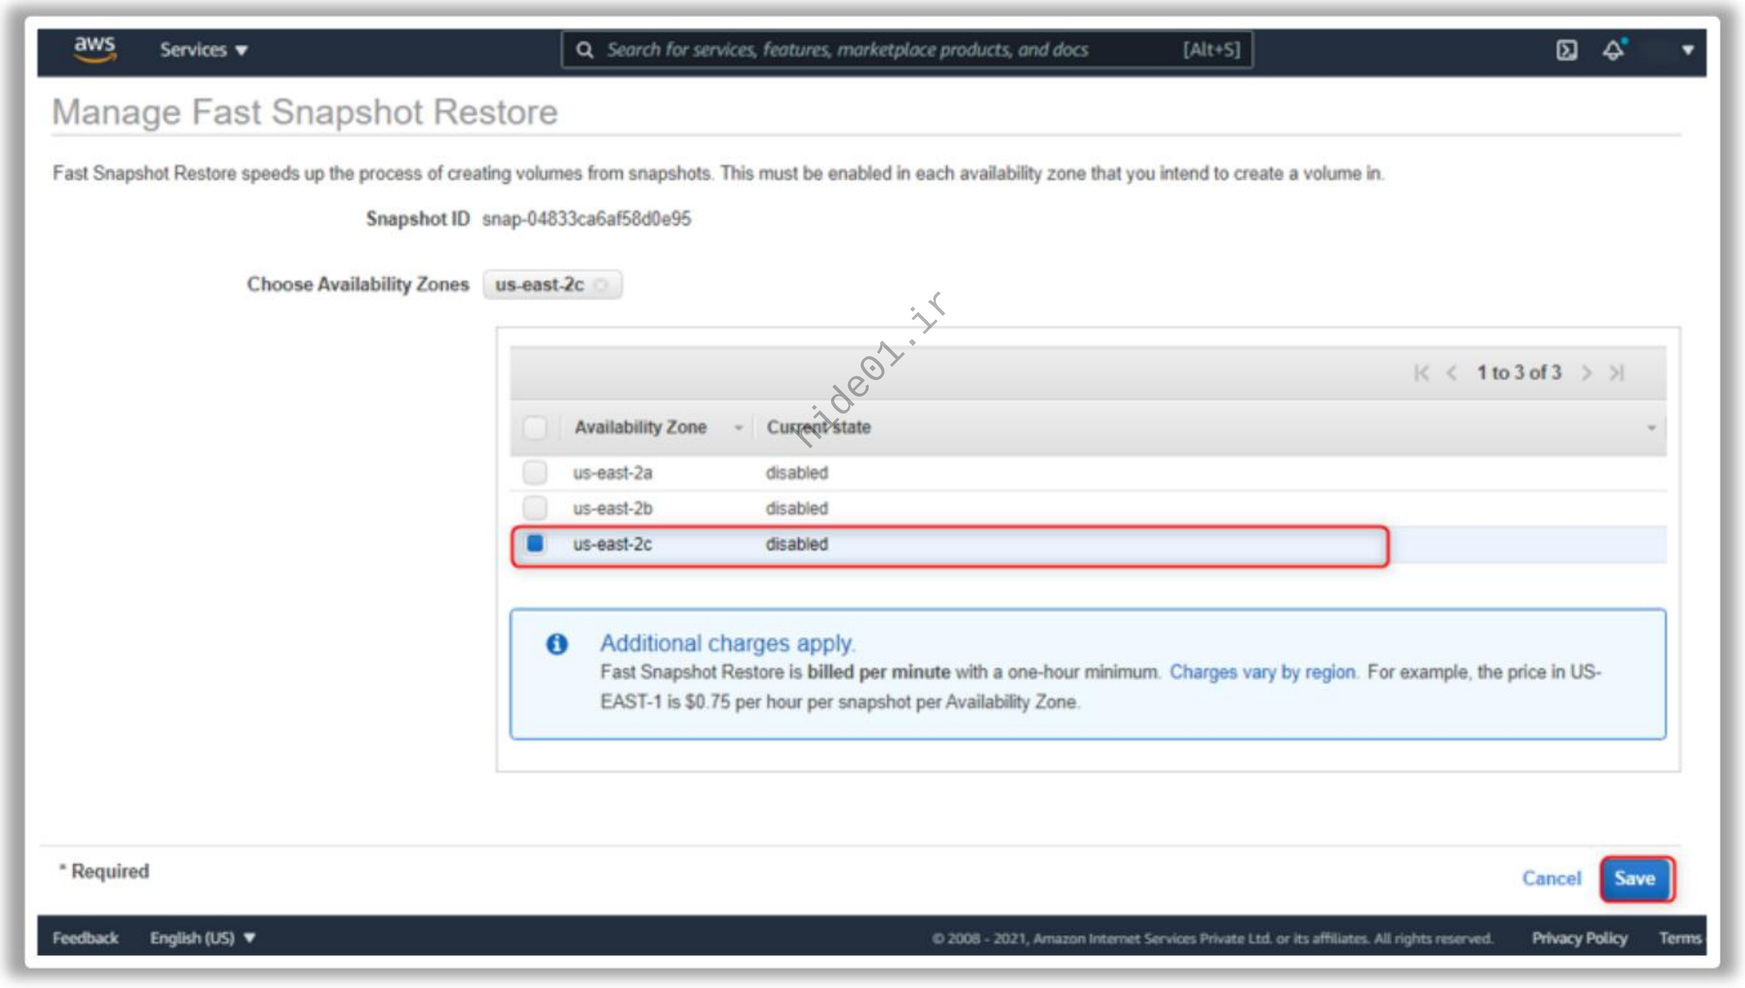

Figure 9.87: Manage Fast Snapshot Restore Console

- After clicking on the Save button, you will see the message that the **Fast Snapshot Restore settings are saved**.

Note: EBS Fast Snapshot Restore should be enabled if your organization would suffer a high impact even with slight downtime.

Certified Cloud Security Engineer Copyright © by **EC-Council**
All Rights Reserved. Reproduction is Strictly Prohibited.

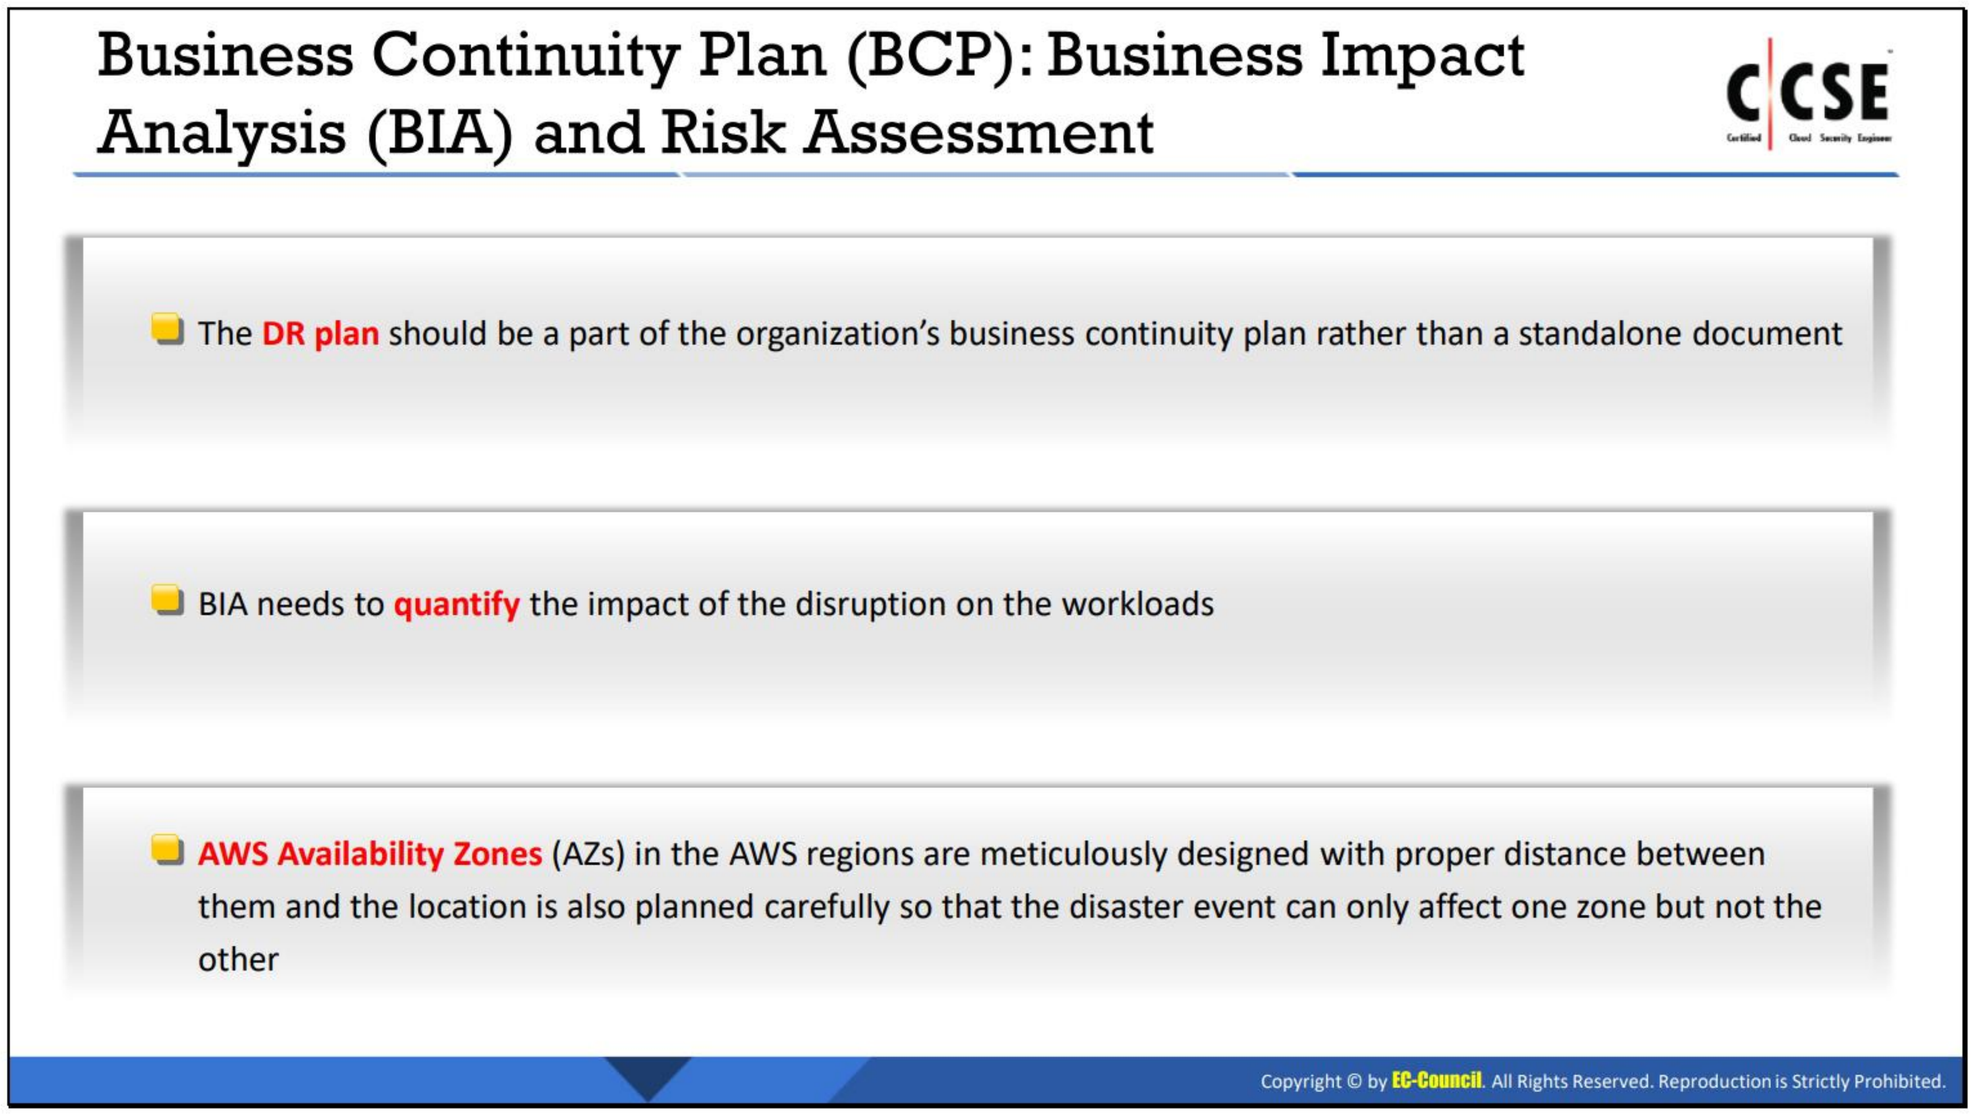

Business Continuity Plan (BCP): Business Impact Analysis (BIA) and Risk Assessment

The DR plan should be a part of the organization's BCP rather than a standalone document. The AWS customer should not maintain an aggressive disaster recovery target for a workload that does not meet the BC objectives due to the impact of the disaster event on the business. For example, a natural disaster such as an earthquake or flood can prevent the transportation of products that are purchased from an online portal. The disaster recovery strategy might keep the workload functioning, but the BCP will not be able to support transportation. Hence, the disaster recovery strategy must focus on business context, priorities, and requirements.

BIA needs to quantify the impact of the disruption on the workloads. The impact of the disruptive event that does not allow the utilization of your workloads to external as well as internal customers and the adverse effect of the disruption on the business should be determined through BIA. This analysis will be used to decide how much data loss is tolerable and how fast the workloads can be made available. Avoid making the recovery objectives in isolation because the likelihood of disruption and recovery costs are important factors in disaster recovery of your workloads. The impact on the business due to a disaster event might be time-dependent. For instance, the disruption in the payroll system has a high impact on the business before everyone gets the payment; however, it would have a low impact after everyone has received the payment.

To determine the likelihood of disruption for each disruptive event, conduct risk assessment on the type of disaster and geographical impact of the disruptive event, including an overview of technical implementation of the workload. Therefore, for highly critical workloads, consider high availability in many regions along with continuous backups to reduce the impact of disaster events on the business, whereas, for less critical workloads, there is no need of placing a disaster recovery strategy.

Certified Cloud Security Engineer Copyright © by **EC-Council**
All Rights Reserved. Reproduction is Strictly Prohibited.

AWS AZs in AWS regions are meticulously designed with proper distance between them and the location is also planned carefully so that the disaster event can only affect one zone but not the other. Hence, the multiple AZ architecture of AWS within a region will meet the AWS customer's risk mitigation requirements.

To make sure that the disaster recovery strategy delivers the appropriate business value by considering the risk and the impact of the disaster event on the business, evaluate the cost of disaster recovery options. Using this information, the AWS customer can document the risk, threat, impact, and cost of various disaster events and the related disaster recovery options. Using this information, every workload recovery objective can be determined.

hide01.ir

Certified Cloud Security Engineer Copyright © by **EC-Council**
All Rights Reserved. Reproduction is Strictly Prohibited.

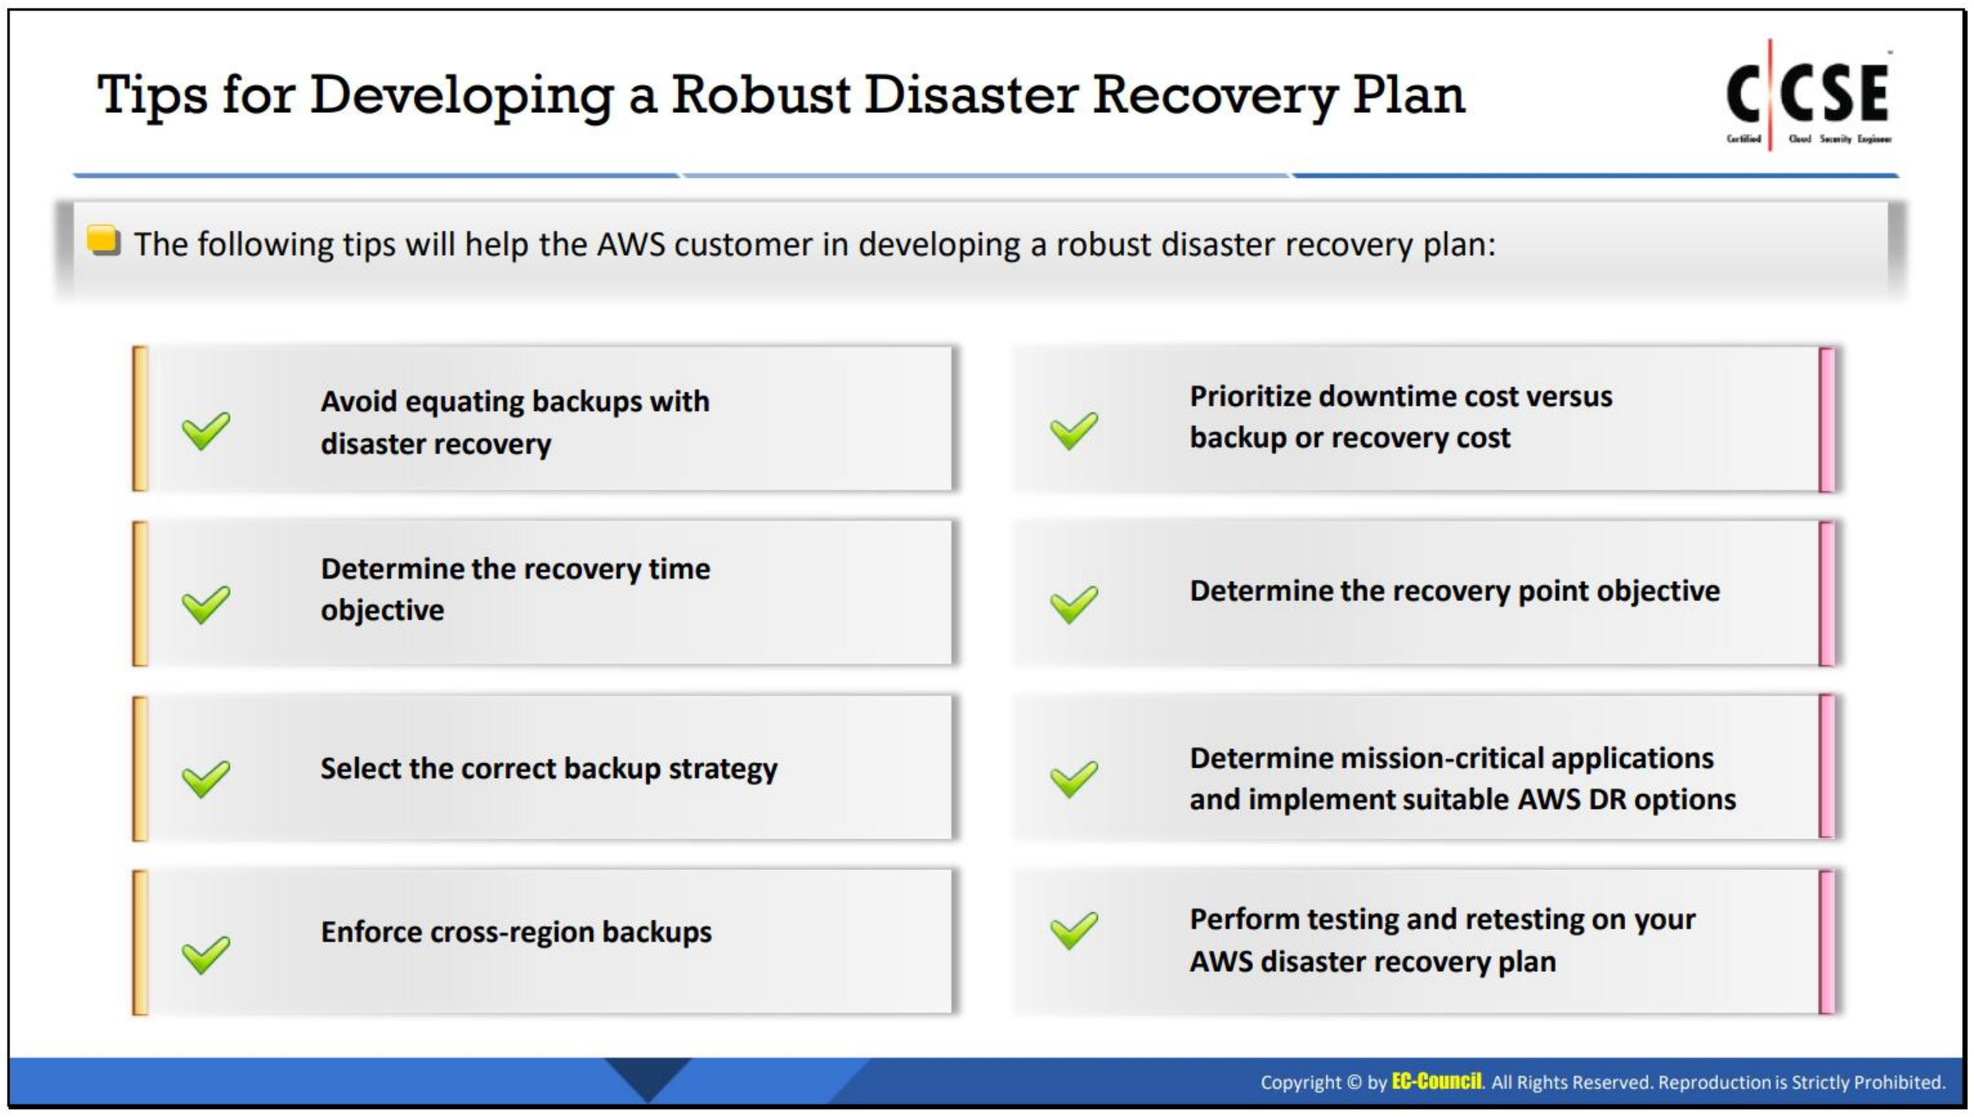

Tips for Developing a Robust Disaster Recovery Plan (DRP)

AWS offers customers easily scalable and flexible solutions to meet their cloud computing requirements. AWS customers do not have to invest money in making, using, and maintaining costly physical infrastructure as all these facilities are provided by AWS at a very low price. In case of DR, the AWS decentralized system is one of the best business solutions. However, while the essential features and tools that help in disaster recovery are provided by AWS, they will be useless if the customer does not have a comprehensive disaster recovery strategy. The following tips will help the AWS customer in developing a robust DRP:

Avoid Equating Backups with Disaster Recovery

It is important to make a backup of your data regularly, but do not equate this process with disaster recovery. Disaster recovery is a policy, process, and procedure, which is associated with preparation for recovery or continuation of a business process during a time of crisis. Hence, backing up of data would not be helpful unless you have a proper process in place that can retrieve the data quickly and also keep the data for usage effectively.

Prioritize: Downtime Cost Versus Backup or Recovery Cost

AWS DRP should meet all the requirements of the organization. The organization should choose between the amount of money that the organization will lose due to downtime versus the amount of money that the organization will be spending on disaster recovery of data. If the organization would not have a high impact with a lengthy disaster event, then the organization can opt for a less expensive backup and recovery strategy, whereas if the organization would have a high impact with slight downtime, then the organization should opt for duplicate production environment provided by AWS.

Certified Cloud Security Engineer Copyright © by **EC-Council**
All Rights Reserved. Reproduction is Strictly Prohibited.

Determine the Recovery Time Objective

As RTO is the maximum duration of time that an organization's recovery process takes to bring everything back online without causing any undesirable loss to the organization. The organization should determine the RTO before evaluating the cost of backup and recovery against downtime cost.

Determine the Recovery Point Objective

RPO is the maximum time frame an organization is willing to lose data during a disaster event. For instance, if a disaster event occurred in an organization 2 hours after the last backup, the changes that have occurred in the data during these 2 hours will be lost. If the organization can tolerate this amount of data loss, then 2 hours RPO is considered a good fit. In contrast, if 2 hours of data loss is considered as a major loss to the organization, then the RPO time should be minimized by providing frequent backups.

Select the Correct Backup Strategy

As taking backup is considered a part of effective disaster recovery planning, an organization needs to select the right backup recovery plan. When customers opt for cloud-based solution, they must select backup options such as EBS snapshots or AMI.

Determine Mission-Critical Applications and Implement Suitable AWS DR Options

After ascertaining the RPO, RTO, and suitable backup strategy, it is important to determine which AWS disaster recovery is suitable for their organization. Based on the selected option, the organization should find and give priority to the mission-critical applications.

hidoi.ir

Enforce Cross-Region Backups

If a disaster event strikes your primary data storage, then having the backup in the same building or same region would hardly help you in disaster recovery. The AWS backup data storage in multiple AZs and various regions allows the customer to recover the data even if the primary data storage is knocked offline.

Perform Testing and Retesting on your AWS Disaster Recovery Plan (DRP)

Sometimes the best AWS DRP would fail when it is put into actual practice. Therefore, it is important to perform tests and retests on your AWS DRP periodically for flaws. AWS duplicate environment allows the AWS customer to test their plan utilizing real-time scenarios without affecting the performance of the actual production environment.

Certified Cloud Security Engineer Copyright © by **EC-Council**
All Rights Reserved. Reproduction is Strictly Prohibited.

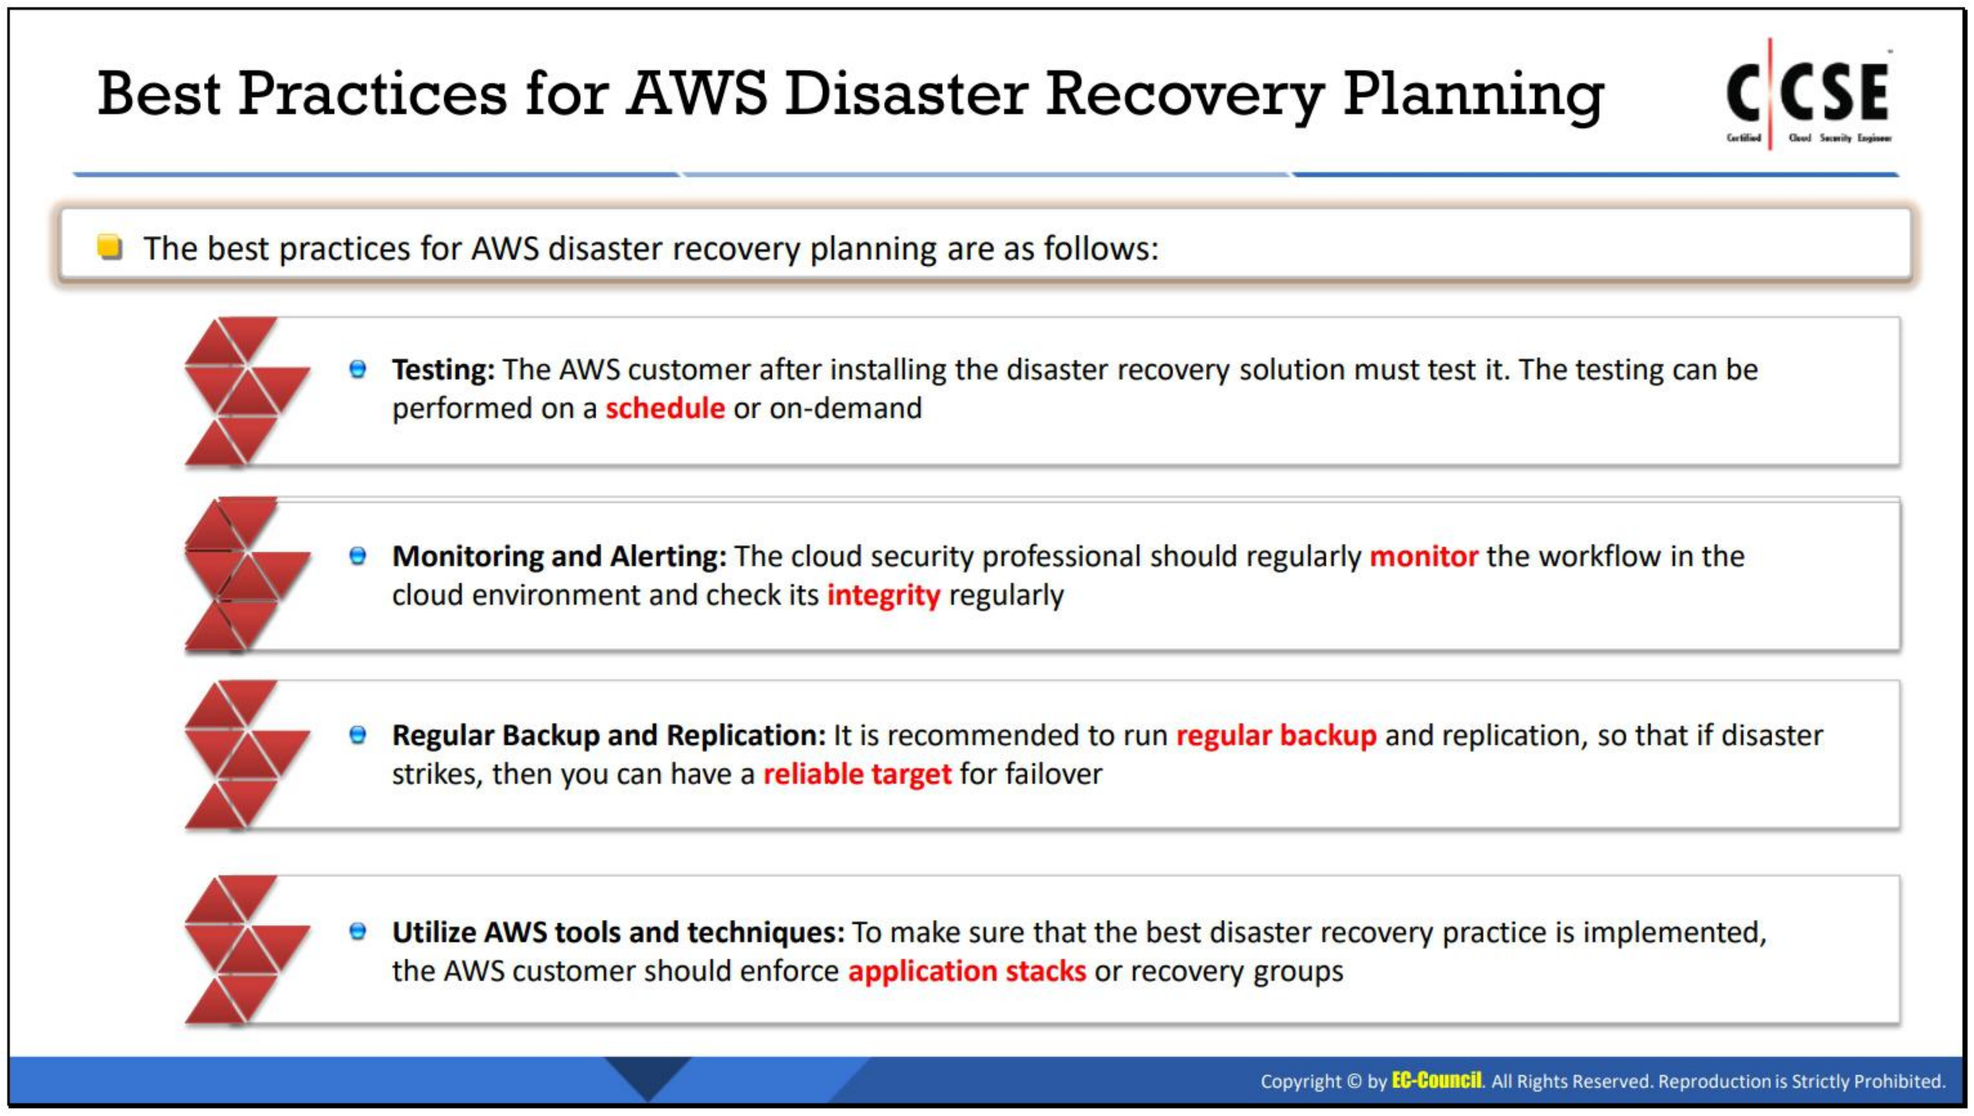

Best Practices for AWS Disaster Recovery Planning

- **Testing**: The AWS customer after installing the disaster recovery solution must test it. The testing can be performed on a schedule or on-demand. Perform game-day testing to verify that the DRP is working as expected and RTOs is within the expected range. To achieve this, utilize AWS CloudFormation to deploy the entire environment on Amazon EC2. The AWS customer can make a resource template, which can be used to model and manage the components of infrastructure in the cloud environment. Perform periodic testing to check whether all the infrastructure components are working properly, and the RPO and RTO can be successfully met when the disaster strikes.
- **Monitoring and alerting:** The organization should detect the issues at the earliest before they strike the organization in the form of a disaster event. The cloud security professional should regularly monitor the workflow in the cloud environment and check its integrity regularly. Due to this, the administrator can detect the threat like server failure, application shutdown, or connectivity issues. AWS services such as Amazon CloudWatch will help the AWS user in evaluating the AWS resource performance. It will send alerts and notifications to the administrator if certain metrics reach or crosses the critical level.
- **Regular backup and replication:** Backup and replication should be regularly performed so that if the organization is hit by a disaster event, then the organization would have a reliable target for failover. After recovering everything and switching into the cloud environment, continue taking regular backup and replication. Store the backups and replicas in the various geographical location provided by AWS to prevent the risk of a single point of failure.

Certified Cloud Security Engineer Copyright © by **EC-Council**
All Rights Reserved. Reproduction is Strictly Prohibited.

- **Utilize AWS tools and techniques:** To make sure that the best disaster recovery practice is implemented, the AWS customer should enforce application stacks or recovery groups. This will allow the AWS customer to arrange the recovery of infrastructure properly. For example, give priority to the restoration of business-critical applications, etc.

hide01.ir

Certified Cloud Security Engineer Copyright © by **EC-Council**
All Rights Reserved. Reproduction is Strictly Prohibited.

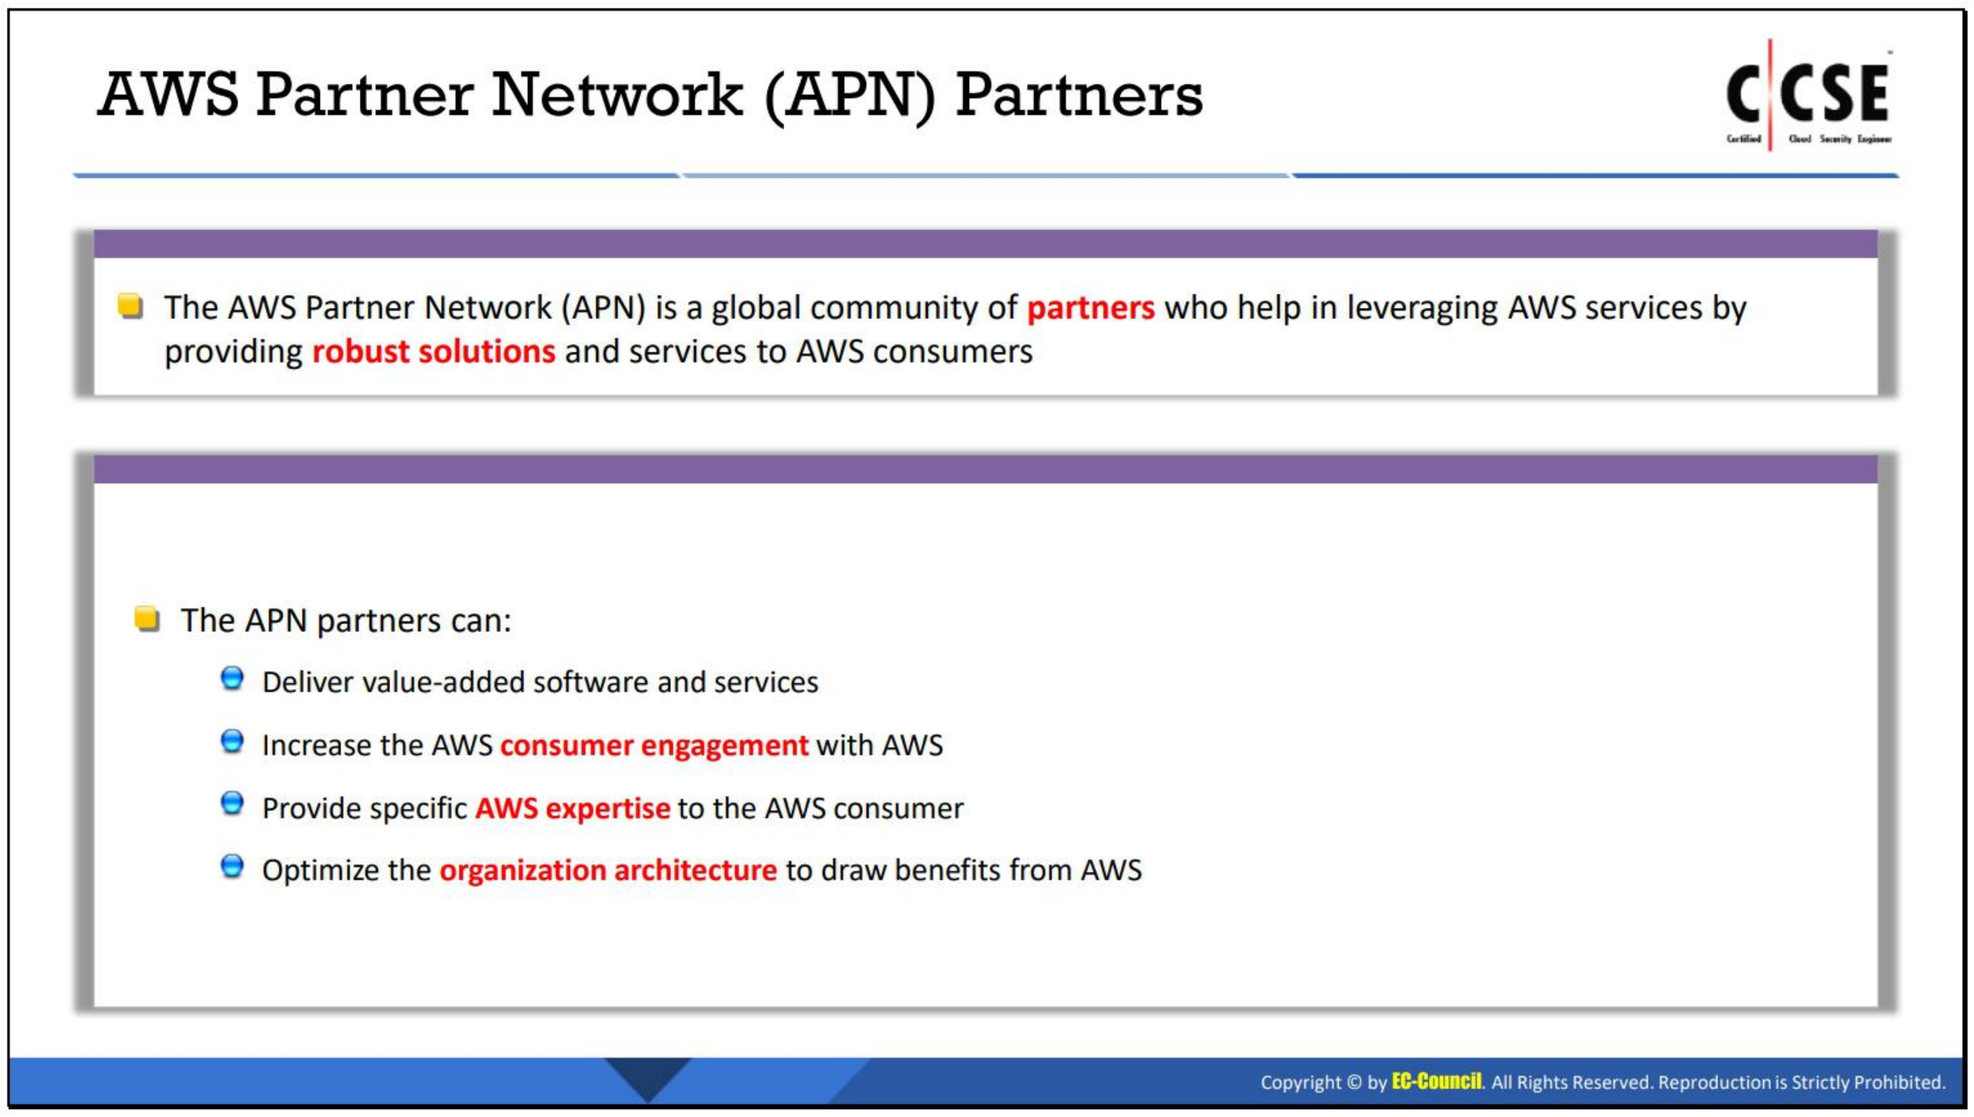

AWS Partner Network (APN) Partners

The APN is a global community of partners who help in leveraging AWS services by providing robust solutions and services to AWS consumers. AWS supports partners in building, selling, and marketing their AWS products.

AWS helps their customer in finding and connecting partners with high AWS expertise across the globe. These partners will help the cloud consumer organization at any point during their cloud adoption journey.

Whether an organization is starting to adopt the AWS or stretching their business using AWS, the APN is the best place to start.

APN partners can

- Deliver value-added software and services.
- Increase the AWS consumer engagement with AWS.
- Provide specific AWS expertise to the AWS consumer.
- Optimize the organization architecture to draw benefits from AWS.

Certified Cloud Security Engineer Copyright © by **EC-Council**
All Rights Reserved. Reproduction is Strictly Prohibited.

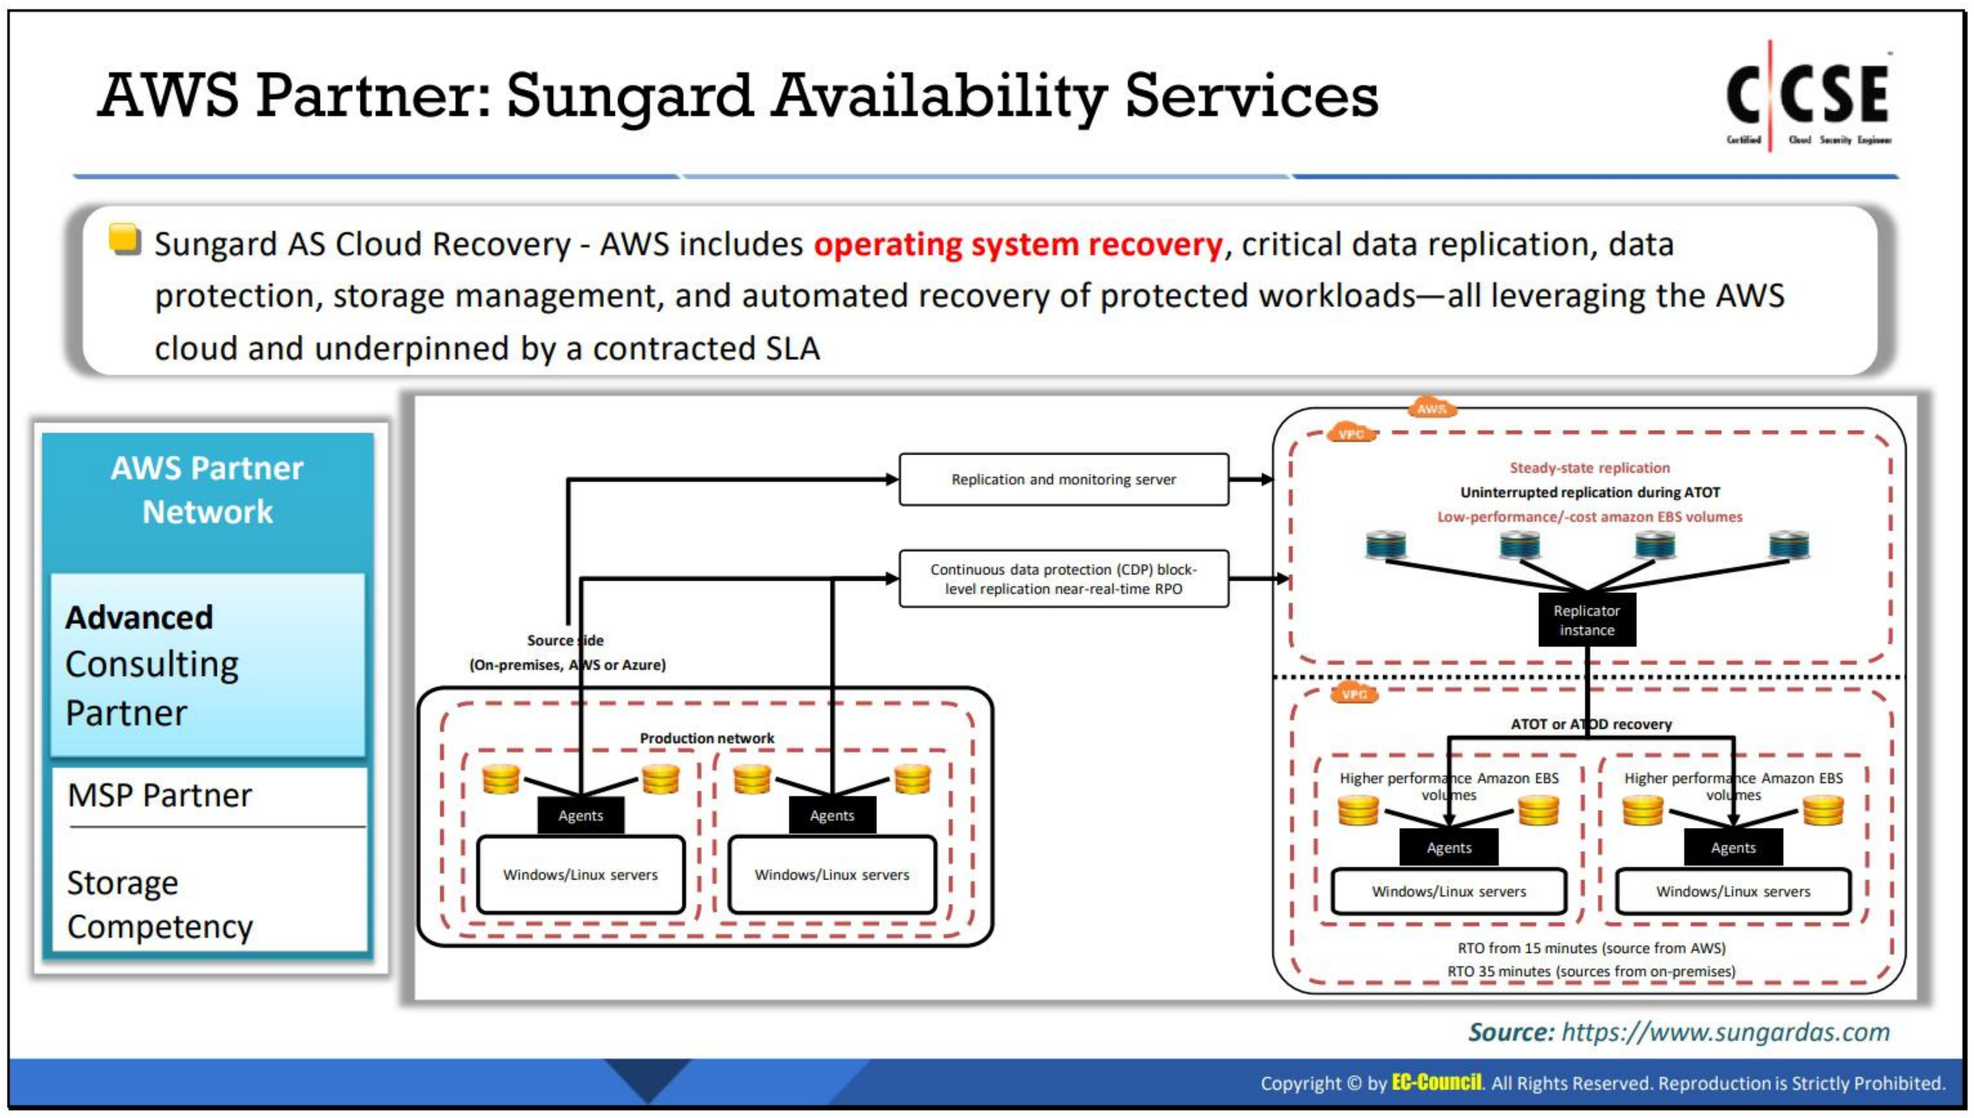

AWS Partner: Sungard Availability Services

Source: https://www.sungardas.com

More than a million customers rely on AWS cloud services for faster innovation, lower costs, and global application scalability. As an Advanced Consulting Partner in APN, an AWS Managed Service Provider, and an AWS Storage Competency Partner with 40 years of experience, Sungard AS is uniquely positioned to help AWS customers leverage AWS cloud resources across the production and recovery environments.

Sungard AS Cloud Recovery - AWS includes operating system recovery, critical data replication, data protection, storage management, and automated recovery of protected workloads, all leveraging the AWS cloud and underpinned by a contracted SLA.

All the AWS cloud services offered by Sungard are backed by consulting services and a tailored approach to ease cloud adoption—from planning, design, and migration to running, managing, and optimizing the environment for you.

With Sungard AS, the customer can keep your data protected, your recovery seamless, and your business operations running smoothly while you take advantage of AWS cloud services. Some AWS related key services offered by Sungard AS are as follows:

- **Cloud Recovery – AWS:** A fully managed recovery service that makes the AWS cloud your recovery destination.
- **Managed Backup – AWS:** Stores data on local appliances and in the AWS cloud. Unifies data protection for compliance and cost control; simplifies and speeds restoration and recovery.

Certified Cloud Security Engineer Copyright © by **EC-Council**
All Rights Reserved. Reproduction is Strictly Prohibited.

- **Managed Cloud – AWS:** Infrastructure management for workloads deployed in dedicated virtual private clouds within an AWS recovery environment.
- **Managed Cloud - AWS, UK Sovereign**: The agility of AWS cloud with the security/compliance needed for sensitive data, including "Official" and "Official Sensitive" workloads (the UK only)

hide01.ir

Certified Cloud Security Engineer Copyright © by **EC-Council**
All Rights Reserved. Reproduction is Strictly Prohibited.

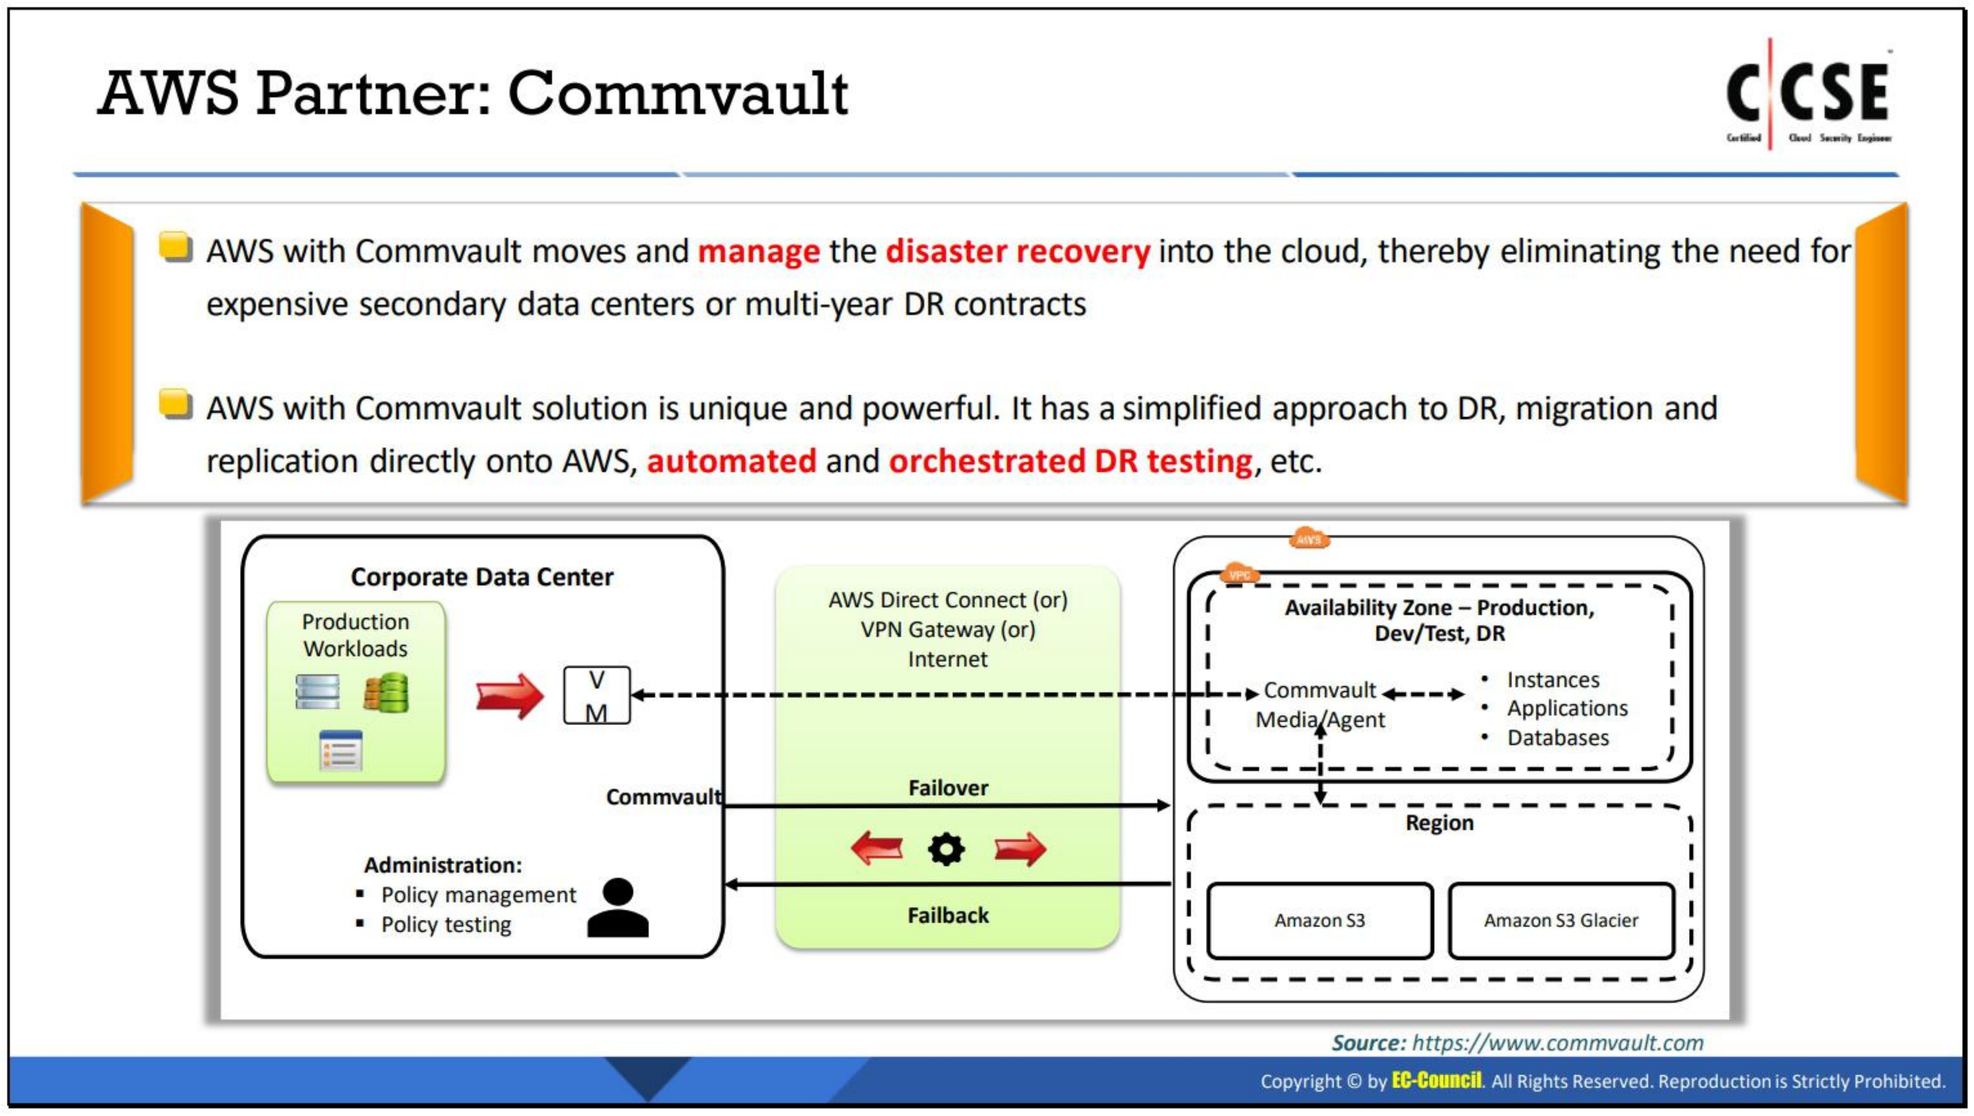

AWS Partner: Commvault

Source: https://www.commvault.com

AWS and Commvault move and manage disaster recovery into the cloud, thereby eliminating the need for expensive secondary data centers or multi-year DR contracts. Commvault's policy-driven data platform manages the RTO and RPO requirements across different tiers of workloads. Select the policies, and Commvault software does the rest.

Commvault DR orchestration keeps data on low-cost AWS storage until it is time for monthly testing — or when there is a need to recover from an incident.

AI and in-depth analytics monitor risk profile and data pool. Alerts warn of potential ransomware/malware attacks so it is possible to quickly isolate the attack and minimize any impact. Automation reduces the risk of human error and foundation and application hardening halts attacks in their tracks.

AWS + Commvault solutions are unique and powerful with the following features:

- **Simplified approach to DR**: Lower cost and complexity — gain a single view into all data.
- **Migration and replication directly onto AWS**: No need for cloud connectors, gateways, or costly proprietary hardware.
- **Power management:** Only power up resources as needed — power down when finished. Significantly lower cloud compute costs.
- **Automated and orchestrated DR testing**: Reduce the risk of human error by eliminating scripting.

Certified Cloud Security Engineer Copyright © by **EC-Council**
All Rights Reserved. Reproduction is Strictly Prohibited.

- **Built-in comprehensive ransomware protection**: Limit disruption and damage to the business.
- **Commvault deduplication reduces data sent to AWS**: Speeds up transmission and recovery times; lowers storage costs.

hide01.ir

Certified Cloud Security Engineer Copyright © by **EC-Council**
All Rights Reserved. Reproduction is Strictly Prohibited.

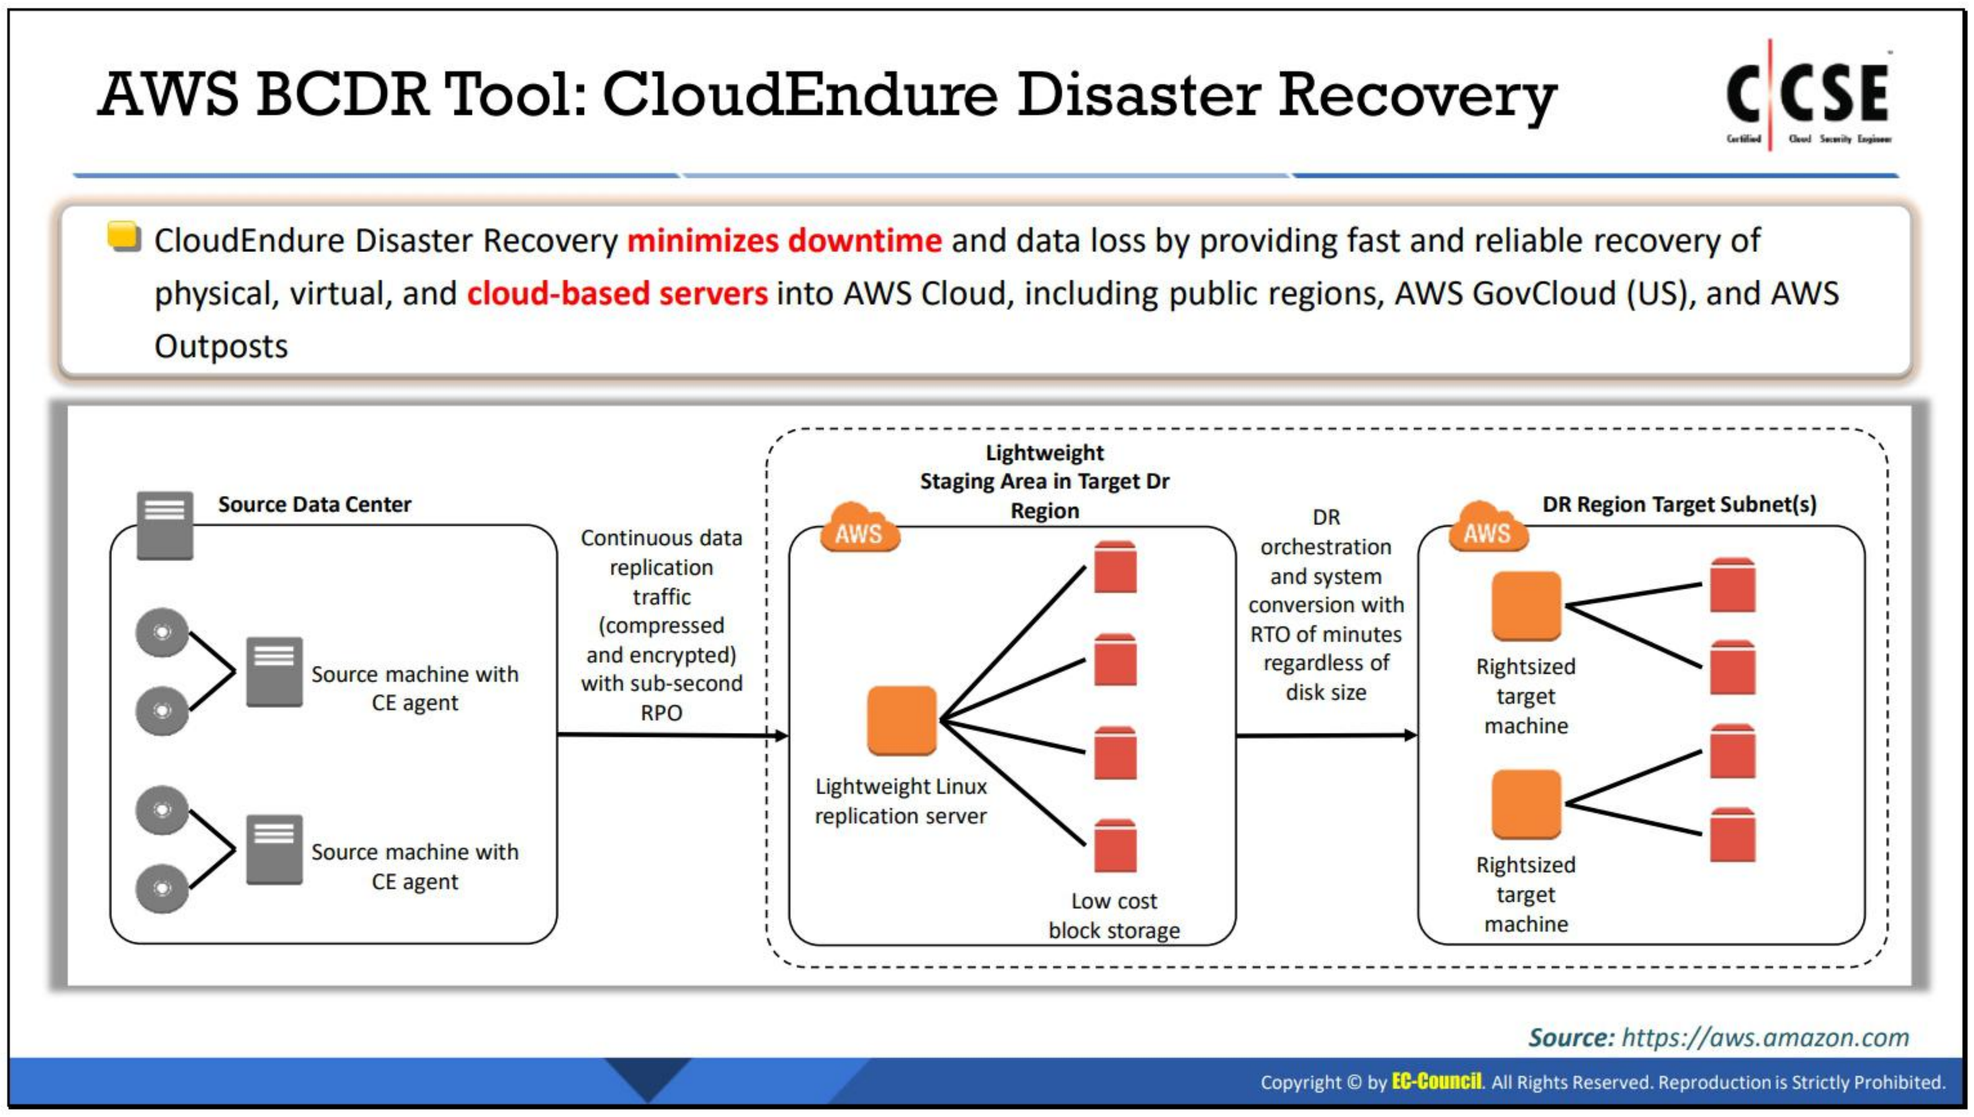

AWS BCDR Tool: CloudEndure Disaster Recovery

Source: https://aws.amazon.com

IT disasters such as data center failures, server corruptions, or cyberattacks can not only disrupt business but also cause data loss, impact your revenue, and damage your reputation. CloudEndure Disaster Recovery minimizes downtime and data loss by providing fast, reliable recovery of physical, virtual, and cloud-based servers into AWS Cloud, including public regions, AWS GovCloud (US), and AWS Outposts.

CloudEndure Disaster Recovery can be used to protect the most critical databases, including Oracle, MySQL, and SQL server, as well as enterprise applications such as SAP.

CloudEndure Disaster Recovery continuously replicates machines (including operating system, system state configuration, databases, applications, and files) into a low-cost staging area in the target AWS account and preferred region.

Benefits

- Reduce downtime and protect against data loss.
- Simplify implementation and increase reliability.
- Protect enterprise applications and databases with a single tool.
- Decrease your total cost of ownership.
- Accelerate your migration to the cloud.
- Protect your data from ransomware attacks.

Certified Cloud Security Engineer Copyright © by **EC-Council**
All Rights Reserved. Reproduction is Strictly Prohibited.

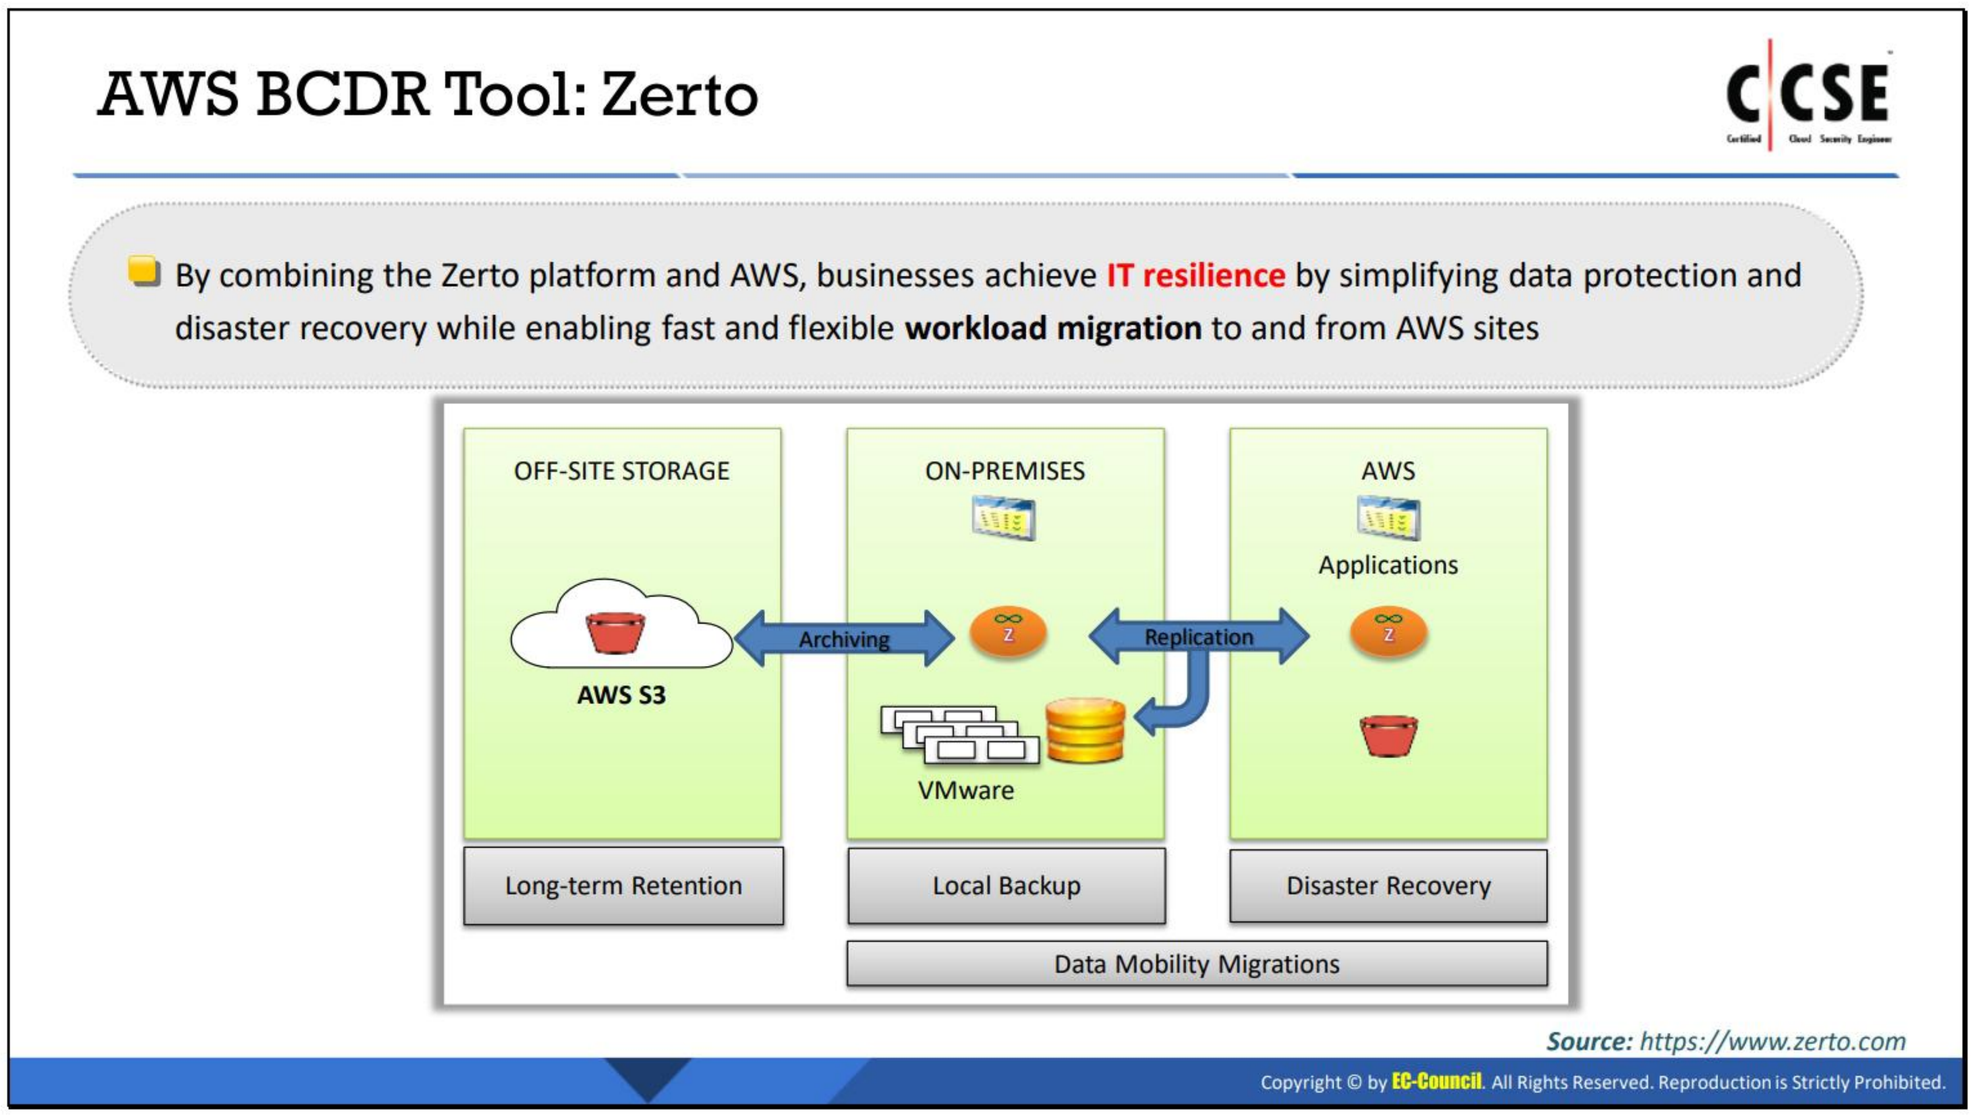

AWS BCDR Tool: Zerto

Source: https://www.zerto.com

Zerto's platform empowers businesses to truly leverage cloud efficiencies with the flexibility to move from, to, or between clouds, such as AWS. Zerto's industry-leading cloud data management and protection platform is the premiere, single solution for protection, recovery, backup, and migration of data in the public cloud.

By combining the Zerto platform and AWS, businesses achieve IT resilience, by simplifying data protection and DR, while enabling fast and flexible workload migration to and from AWS sites. AWS can also be used as a cloud-based disaster recovery site, eliminating the capital costs associated with owning and operating your data center.

Zerto gives businesses the freedom to choose a backup and recovery option that fits their needs. Having disaster recovery or backup target site on standby in the cloud, without having to manage the hardware, power, or cooling, makes AWS the perfect solution for cloud-based data protection.

Key Benefits

- Reduce DR costs
- Simple and easy to deploy and use
- TenCate migrates to AWS with Zerto
- Replication to and from AWS Sites
- Zerto improves DR strategy and migration to AWS

Certified Cloud Security Engineer Copyright © by **EC-Council**
All Rights Reserved. Reproduction is Strictly Prohibited.

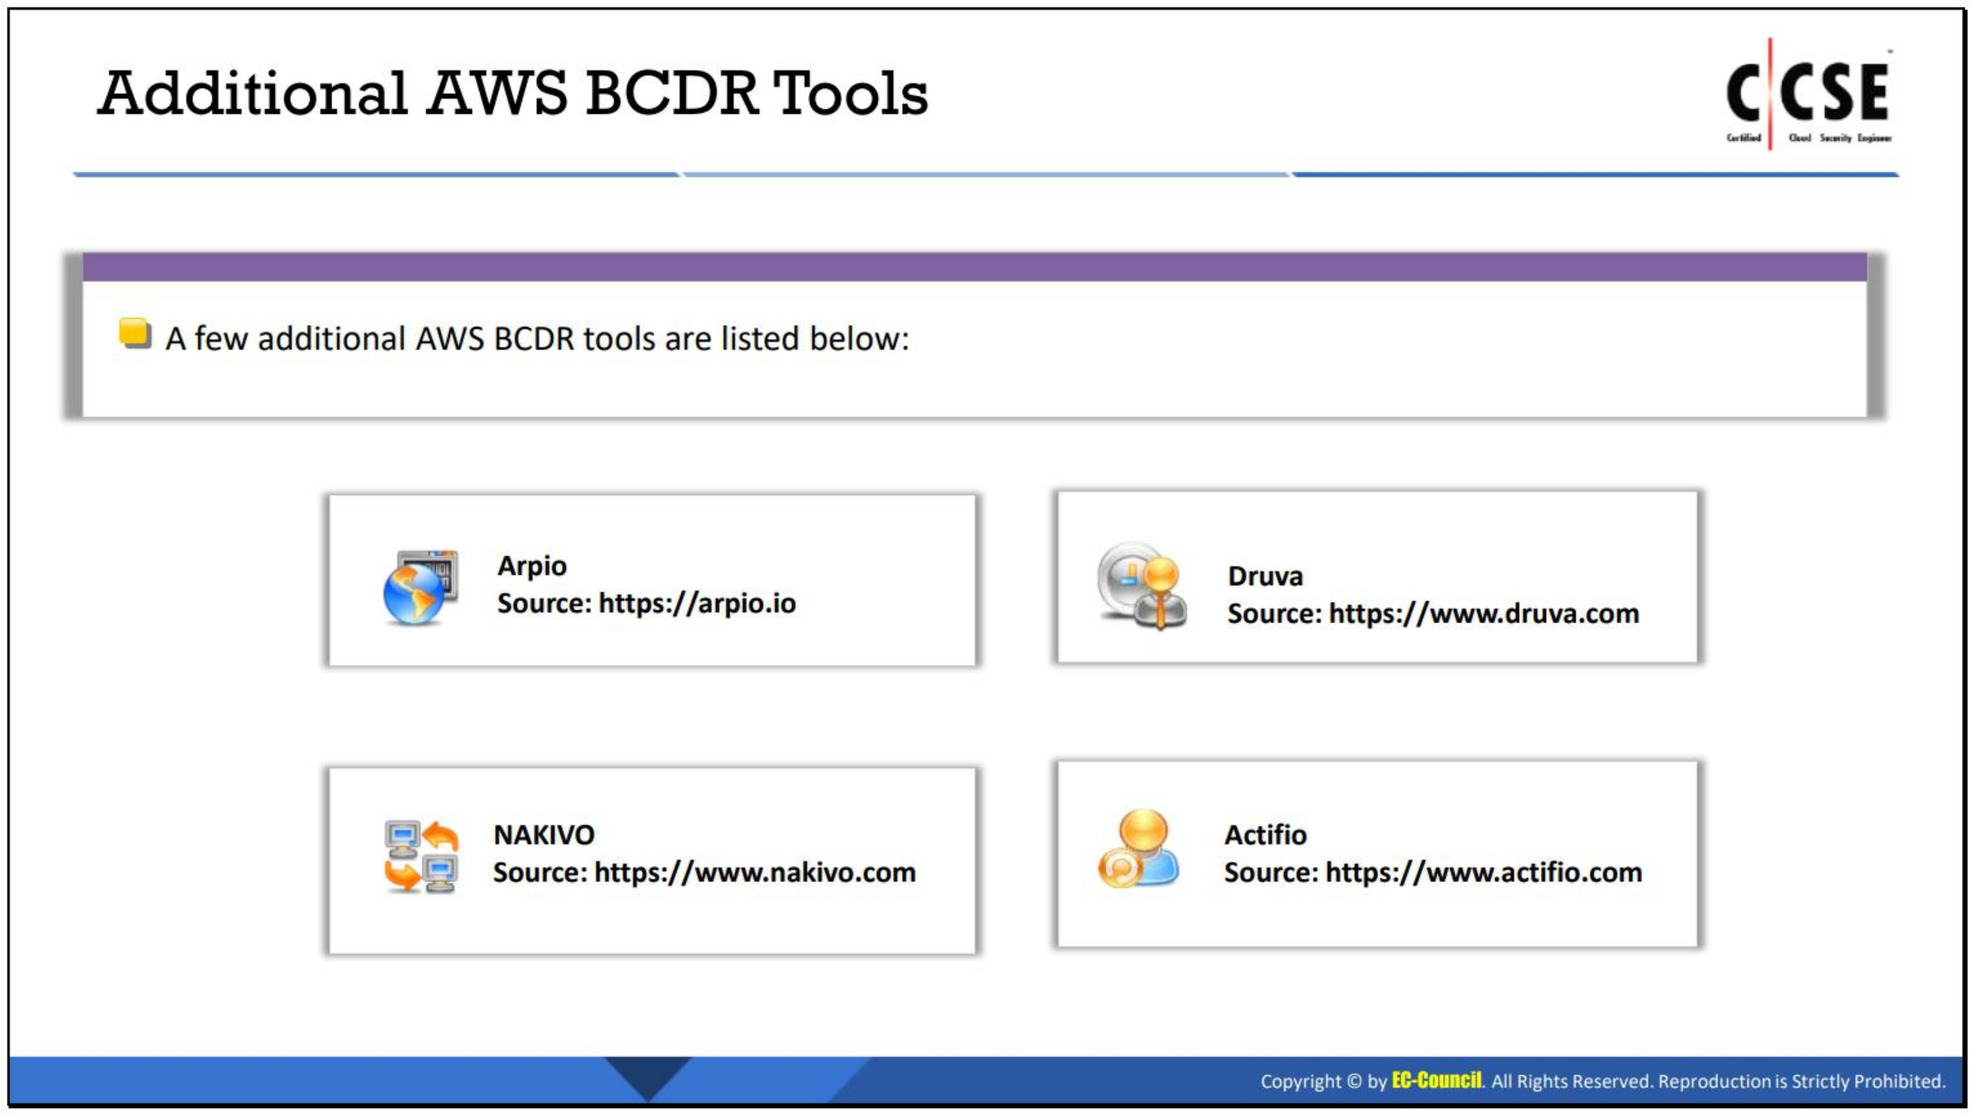

Additional AWS BCDR Tools

Arpio

Source: https://arpio.io

ride01.ir

Arpio protects AWS environments from catastrophic downtime. Arpio continuously replicates the entire environment – data and infrastructure – to an alternate region, so that the customer can opt-out of region-wide downtime. Arpio secures data backups in an immutable vault so customers will always have a path to recovery. If the customer AWS environment breaks, Arpio can rewind the clock on those breaking changes, thereby getting the customer back online fast.

Druva

Source: https://www.druva.com

Unlike server-based solutions that do not scale, Druva provides AWS cloud backup for enterprises to simplify the protection of resources, delivered as-a-service. Built for multi-tenancy to protect multiple accounts, from a single console. Apply policies across multiple accounts using include/exclude rule options. Druva AWS cloud backup is easily scaled across multiple accounts, configure GFS retention policy settings, copy data between regions, and accounts, perform application-consistent backups, and transition EBS snapshots to S3 storage.

NAKIVO Backup and Replication

Source: https://www.nakivo.com

NAKIVO Backup and Replication was designed with a focus on fast and reliable AWS disaster recovery. With the product in place, customers can be sure that they are well-equipped for dealing with disruptive events of any nature. Some key features of the NAKIVO Backup and Replication tool include advanced AWS backup and disaster recovery, superior AWS backup

Certified Cloud Security Engineer Copyright © by **EC-Council**
All Rights Reserved. Reproduction is Strictly Prohibited.

performance, smart storage space utilization, automated AWS data protection tasks, nondisruptive AWS DRP testing, recovery of instances, individual files, or application objects, and an easy-to-use and intuitive web interface.

Actifio

Source: https://www.actifio.com

Enterprises use Actifio for cloud backup to AWS to reduce their data center footprint and costs. Backup, recover, and clone, in just minutes, multi-tier applications including databases, VMs, physical servers to AWS. Reuse backups in AWS S3 to provision rapid database clones to AWS EC2 instances for test/dev/analytics, without consuming any extra AWS EBS storage. Protect AWS cloud resident workloads using AWS EBS snapshots. Catalog and recover AWS EC2 instances from AWS EBS snapshots at scale in any region.

hide01.ir

Certified Cloud Security Engineer Copyright © by **EC-Council**
All Rights Reserved. Reproduction is Strictly Prohibited.

LO#05: Understand Business Continuity and Disaster Recovery in Microsoft Azure

The objective of this section is to understand what BCDR services are provided in Azure and how they work. This section will help cloud security engineers know understand the various Azure services that support BCDR such as Azure Site Recovery, Azure Backup, Recovery Service Vault, etc.

Certified Cloud Security Engineer Copyright © by **EC-Council**
All Rights Reserved. Reproduction is Strictly Prohibited.

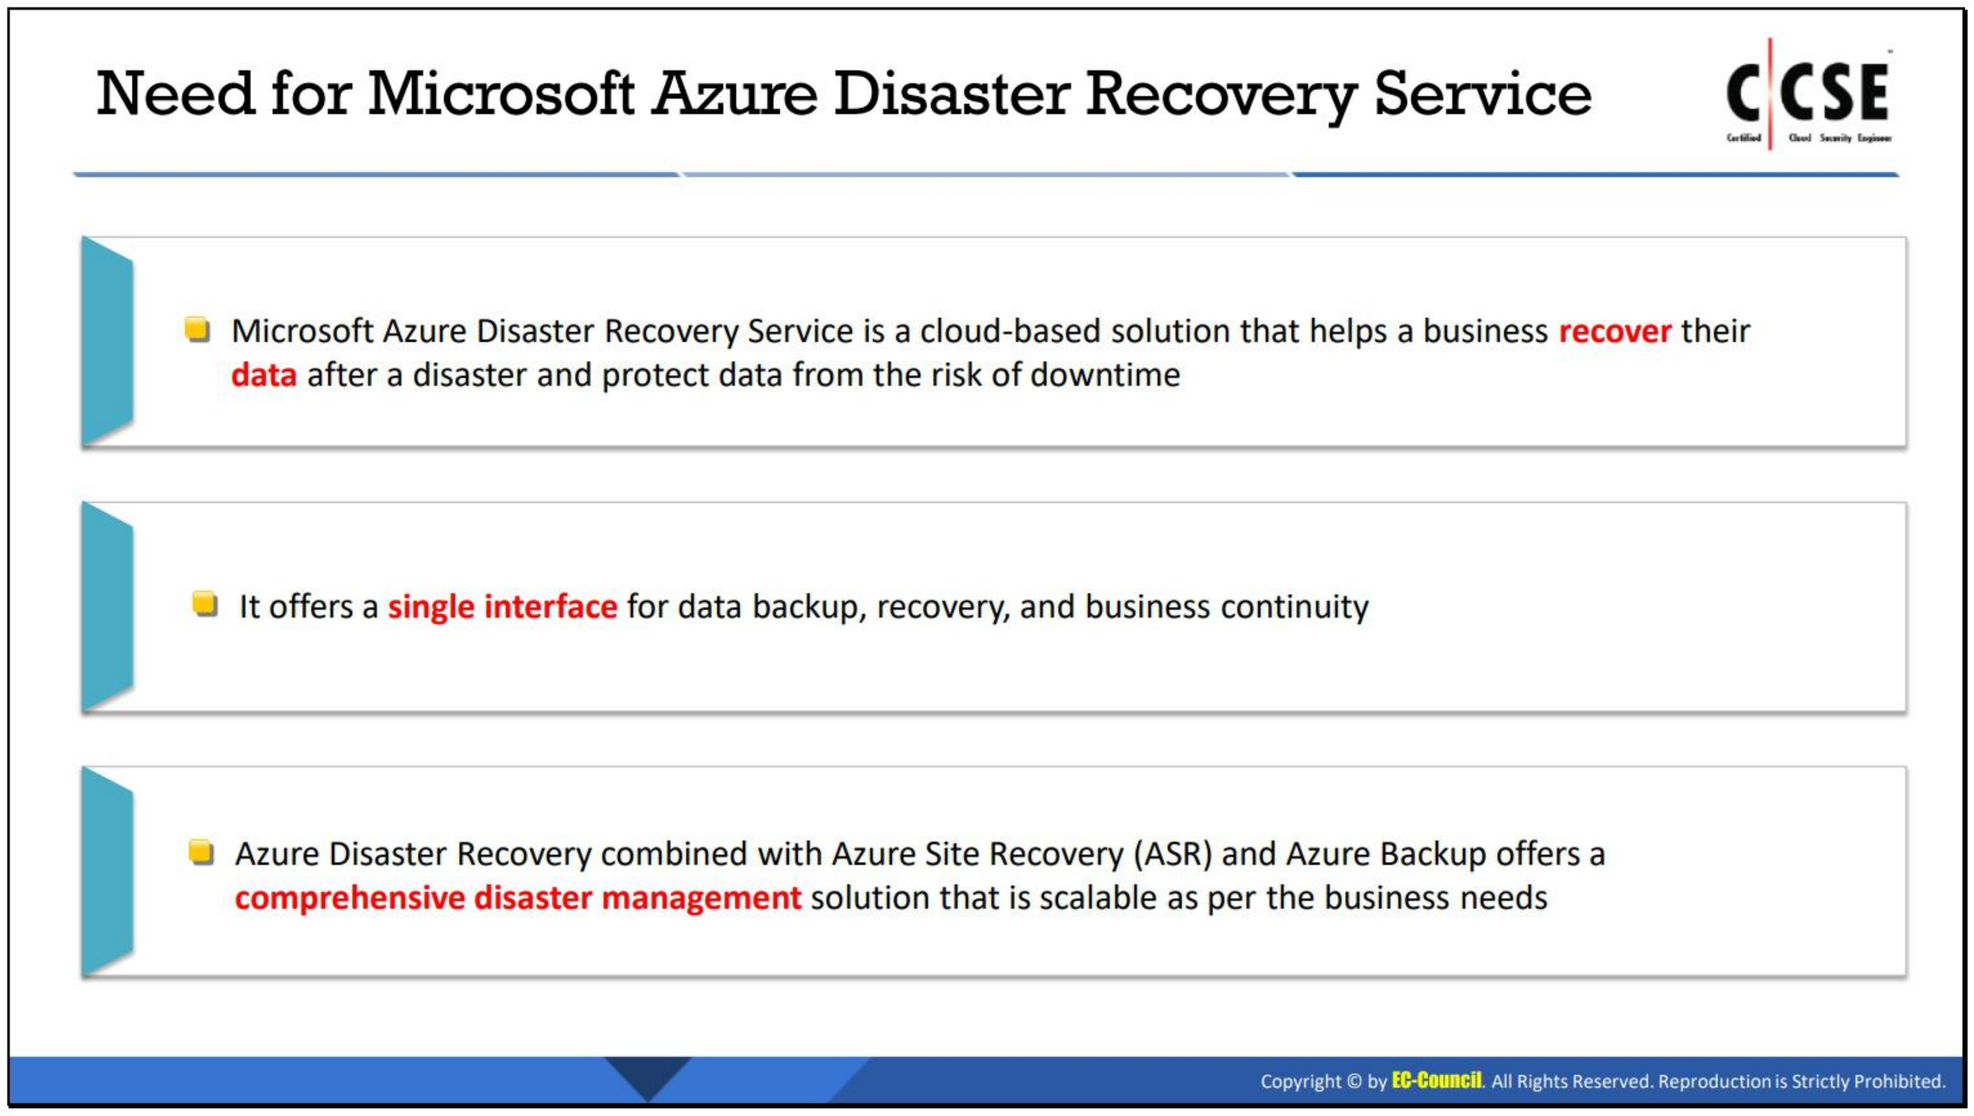

Need for Microsoft Azure Disaster Recovery Service

In this era, every organization relies on technology to function properly. All the staff and customers would also expect high availability and accessibility. The business is dependent on the IT system and, if the connection is lost, it affects the productivity and efficiency of the business. It may lead to data loss if there is no backup to recover. This will have a drastic effect on the business. Hence, it important for every business to implement fault tolerance and BC. The businesses should be able to recover from any type of disaster with minimal downtime. Every business should implement a BCDR plan using a DRaaS solution.

Traditional BCDR solutions require the maintenance of infrastructure, which is costly. Microsoft Azure Disaster Recovery Service is a cloud-based solution that helps a business recover their data after a disaster and protect data from the risk of downtime. A business of any size can make use of the Azure Disaster Recovery service, which offers a single interface for advanced backup, disaster recovery, and BC solutions.

Azure Disaster Recovery offers the following services to contribute to BCDR:

- Azure Site Recovery (ASR) – DRaaS
- Azure Backup – Backup-as-a-Service (BaaS)

Azure Disaster Recovery combined with these two services offers a comprehensive disaster management solution that is scalable as per the business' needs. Most companies have a misconception that backup is enough to recover from a disaster. When an outage takes place in the business, they realize that backup is not sufficient to keep the business up. Hence, ASR and Azure Backup go hand in hand.

Certified Cloud Security Engineer Copyright © by **EC-Council**
All Rights Reserved. Reproduction is Strictly Prohibited.

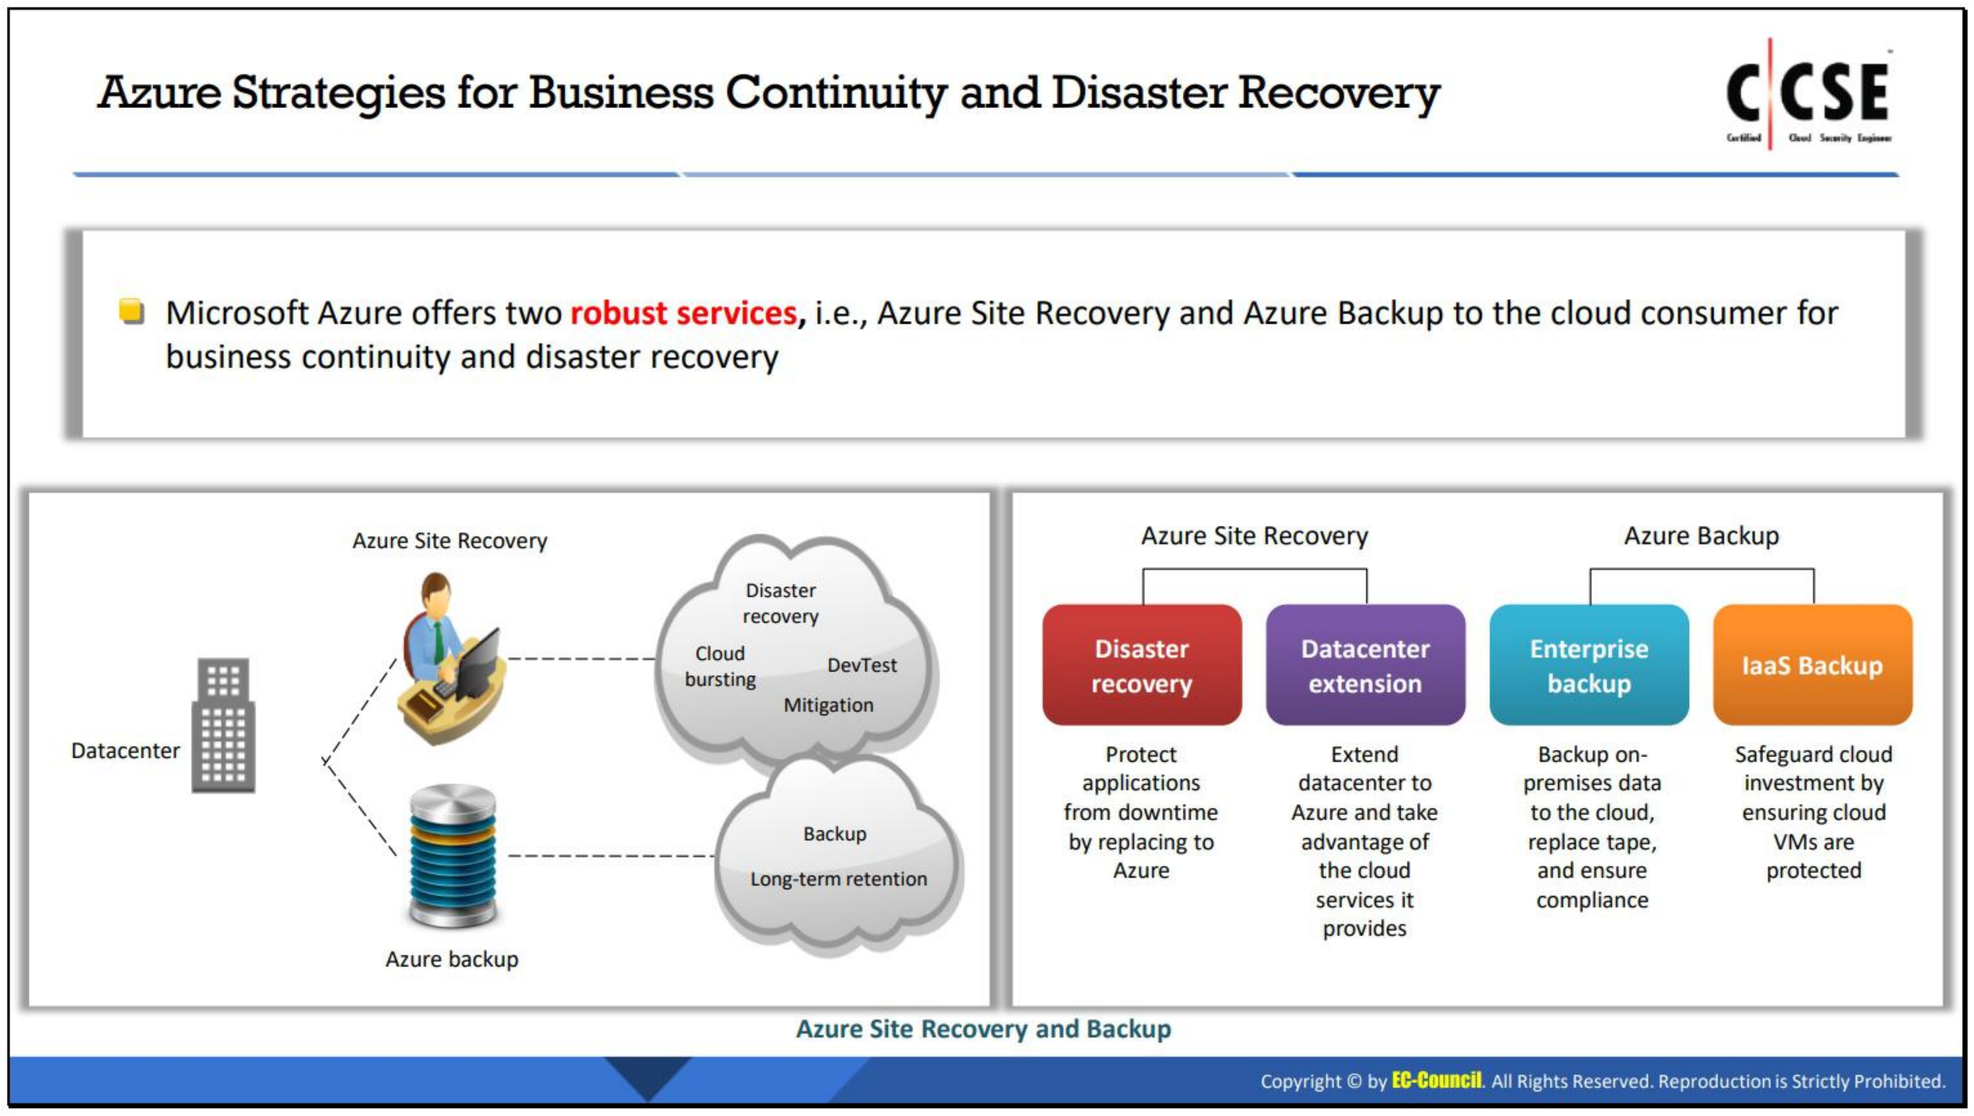

Azure Strategies for Business Continuity and Disaster Recovery

Every organization that has adopted a cloud ecosystem stores data on cloud servers. These organization completely relies on the cloud for their data. If there is any loss in the data, it would lead to tremendous loss to the business of the organization. The loss could be due to a disaster event, natural calamities, human error, or power outage. To prevent this, Microsoft Azure offers two robust services, Azure Backup and ASR, to the cloud consumer for BCDR. Using these services, the organization can backup and restore its data.

Azure Backup

Azure Backup enhances the backup efficiency and reduces storage consumption. It will be utilized in case of accidental deletion, security, alternate location recovery, patch testing, etc. It will be a part of Azure Operations Management Suite or can be sold as a standalone. The long-term retention policy of Azure Backup delivers granular recovery. There will be variation in Azure Backup in terms of acceptable RPO. For VM backup, it will take 1 day, and for database backup, it will approximately take 15 days. Azure Backup provides high reliability and security. Generally, the Azure Backup data is retained for 30 days or less.

Azure enterprise backup can backup on-premises data to the cloud, replace tape, and allow the organization to meet compliance. Azure IaaS backup safeguards cloud investments by protecting the cloud VMs.

Azure Site Recovery

ASR performs failover, failback, and replication. The automated protection and recovery for the workstations and the local servers are provided by ASR. It coordinates various scenarios of DR, migration, testing, and developing environment. The RTO and RPO are

Certified Cloud Security Engineer Copyright © by **EC-Council**
All Rights Reserved. Reproduction is Strictly Prohibited.

very low for ASR. The RTO is in the range of 5-30 minutes and RPO is approximately 30 seconds. ASR requires only operational recovery data. ASR can help create customizable recovery plans and on-demand test copies.

Both Azure Backup and ASR are efficient disaster solutions that help in data backup and restoration. They work along with each other to deliver robust BC; however, there is a slight difference in their purposes. With Azure Backup, you can restore a file that is corrupted, whereas ASR performs replication of the data and system configuration to another datacenter. Apart from this, Azure Backup safeguards from data crisis regularly, whereas ASR is a DR process meant for an apocalyptic scenario. Therefore, based on your business requirements, select either any one of them or work with both.

ASR safeguards applications from downtime by replicating them to Azure. It expands the client datacenter to Azure and draws benefits of the cloud services provided by Azure.

hide01.ir

Certified Cloud Security Engineer Copyright © by **EC-Council**
All Rights Reserved. Reproduction is Strictly Prohibited.

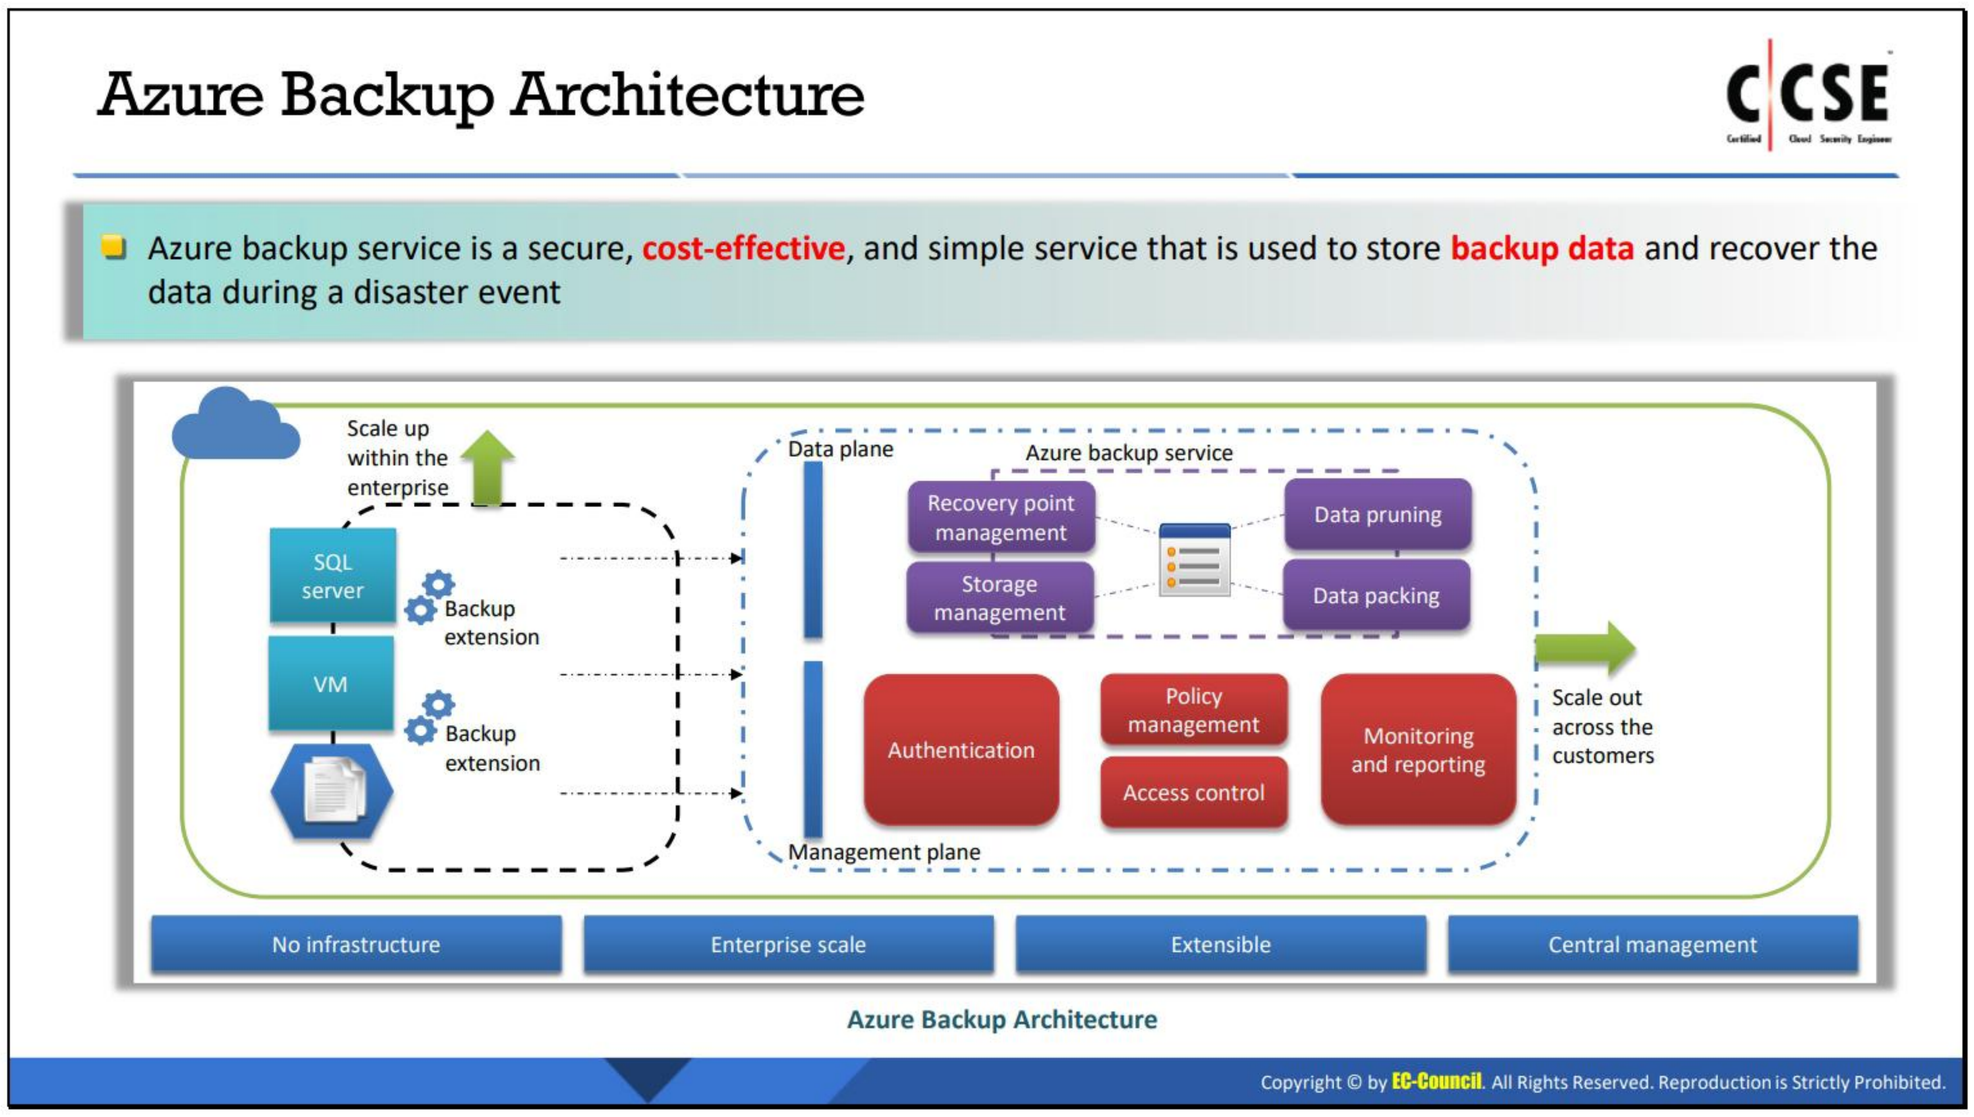

Azure Backup Architecture

Azure Backup service is a secure, cost-effective, and simple service that is used to store backup data and recover it during a disaster event. Azure Backup is used to backup Azure VMs, on-premises files and folders, Azure blobs, Azure managed disks, SQL server in Azure VMs, etc.

The Azure Backup provides a robust solution to backup the on-premises resources to the cloud. There is no need to deploy complex on-prem backup solutions in Azure for short and long-term backup. The isolated and independent backups of Azure will safeguard against accidental data deletion. The backup in Azure will be stored in a Recovery Services vault, which has in-built management of recovery points. Azure Backup provides simple configuration and scalability, and you can restore the backup easily.

Azure Backup utilizes the limitless scale of Azure cloud to provide high availability with zero maintenance or monitoring overhead. Azure Backup neither restricts the Azure user on the amount of data transfer nor levies any charge for the transferred data. Azure Backup offers a robust solution to secure data both at rest and in transit.

The organization can store its backup data in Azure Backup and can manage the data easily. Azure Backup service helps in recovery point management, data packing, data pruning, storage management, and policy management. It secures the data by implementing authentication and access control.

Azure Backup provides centralized monitoring and management. In the Recovery Services vault of the Azure Backup, there is a provision for in-built alerting and monitoring capabilities. By utilizing Azure monitor, Azure users can scale up their monitoring and reporting.

Certified Cloud Security Engineer Copyright © by **EC-Council**
All Rights Reserved. Reproduction is Strictly Prohibited.

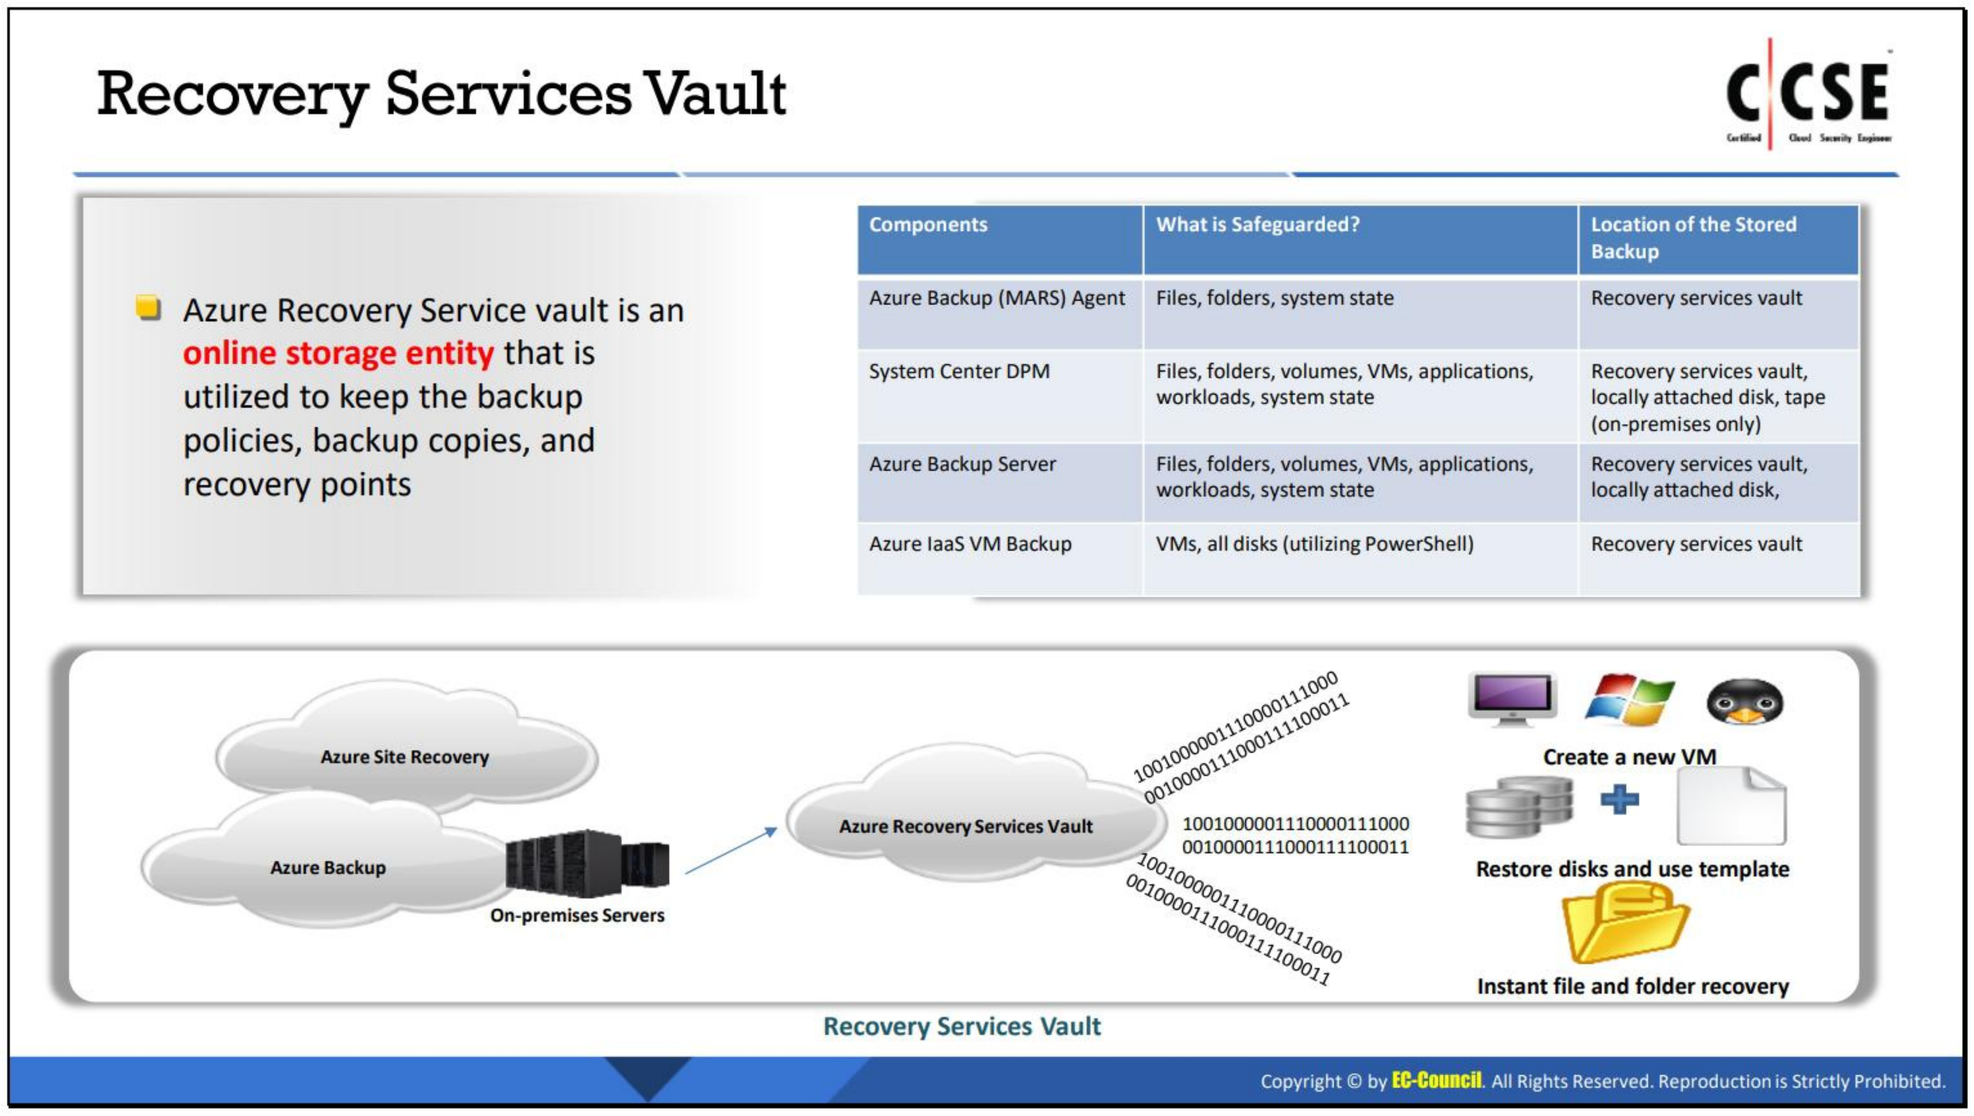

Recovery Services Vault (RSV)

Azure RSV is an online storage entity that is utilized to keep backup policies, backup copies, and recovery points. It can hold the backup data for Azure services, workstations, and on-site servers. RSVs help organize backup data easily and minimize management overheads. Cloud administrators can use these vaults to store backup data of various Azure services such as Azure SQL databases, Windows Server, IaaS VMs (Linux and Windows), Azure Backup Server, and so on. RSVs are based on the Azure Resource Manager model.

Each Azure Subscription supports the creation of up to 500 RSVs. Cloud administrators can choose locally redundant storage (LRS) or geo-redundant storage (GRS) while storing data on RSV. They can decide which storage (LRS or GRS) to use based on the following information:

- **LRS**: With LRS, data gets copied three times in a data center's storage scale unit. As all the three data copies reside in the same region, LRS can protect the data from local hardware failures.
- **GRS**: With GRS, data is copied to a secondary region that is far away from where the source data resides (primary region). As the backup data resides in another region, data is protected even in the event of outages in the primary region. This makes GRS costlier than LRS.

Certified Cloud Security Engineer Copyright © by **EC-Council**
All Rights Reserved. Reproduction is Strictly Prohibited.

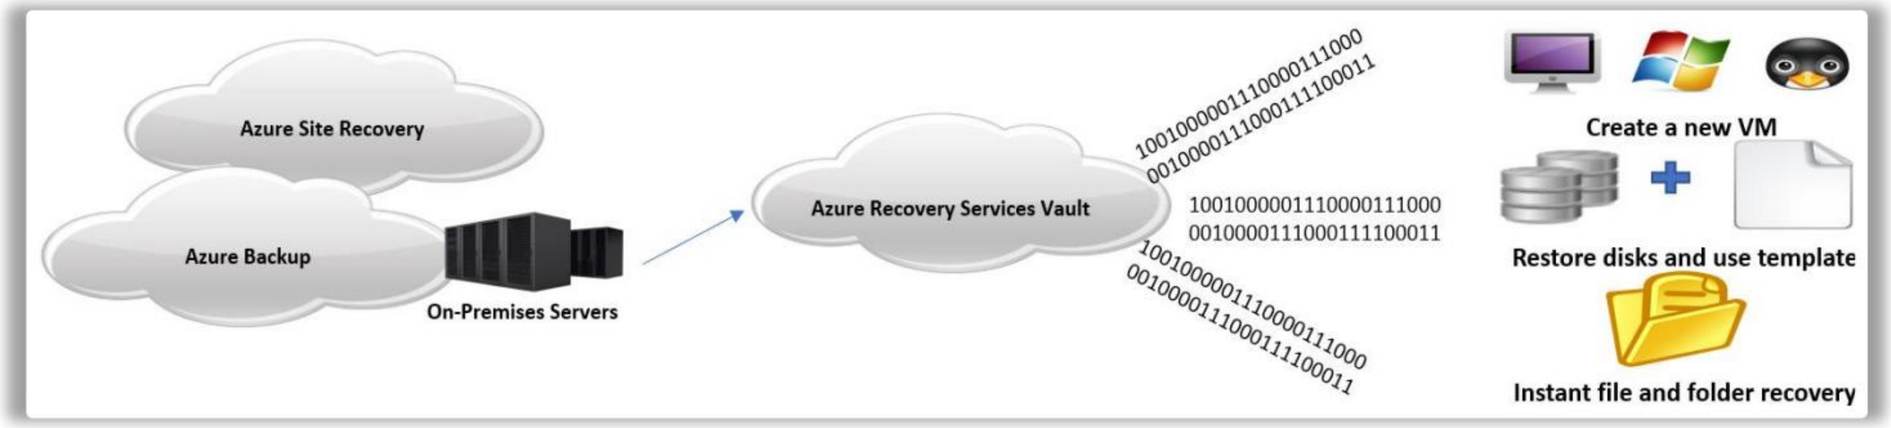

Figure 9.88: Recovery Services Vault

Azure RSV

Azure RSV can help cloud administrators restore a VM from its saved recovery points when they want to repair or rebuild the VM. While restoring a recovery point, they can create a new VM (which is a point-in-time representation of the backed-up VM), perform individual file recovery, or restore drives and use templates to customize a restored VM.

After restoring the VM, they can also do the following:

- Utilize a template to customize the restored VM.
- Attach the restored drives to an existing VM.
- Create a new VM from the restored drives.
- Recover folders and files instantly from an Azure VM backup.

hide01.ir

Azure Backup Components

Various Azure Backup components protect various forms of data such as workloads, files, apps, etc. The following table shows which component can be used to protect which form of data:

Component	What is protected?	Location of the stored backup
Azure Backup (MARS) Agent	Files, folders, and system state	Recovery services vault (RSV)
System Center Data Protection Manager (DPM)	Files, folders, volumes, VMs, applications, workloads, and system state	RSVs, locally attached disk, and tape (on-premises only)
Azure Backup Server	Files, folders, volumes, VMs, applications, workloads, and system state	RSV and locally attached disk
Azure IaaS VM Backup	VMs and all disks (utilizing PowerShell)	RSVs

Table 9.1: Form of Data

Certified Cloud Security Engineer Copyright © by **EC-Council**
All Rights Reserved. Reproduction is Strictly Prohibited.

Security of Azure Backup

- **Network Security:** Encryption is performed on the **entire backup data** from the customer servers to the recovery services vault utilizing AES256
- Backup data is transferred through **HTTP**
- The Azure customer can only decrypt the stored data in Recovery Services vault using **passphrase**
- **Data Security**: The backup data is encrypted using **Azure Disk Encryption** that utilizes dm-crypt on Linux VMs and BitLocker on Windows VMs

Feature	Azure Backup Agent	System Center DPM	Azure Backup Server	Azure IaaS VM Backup
Recovery Services Vault	●	●	●	●
Disk Storage		●	●	
Tape Storage		●		
Compression (in Recovery Services Vault)	●	●	●	
Incremental Backup	●	●	●	●
Disk deduplication		◗	◗	
Key	● = Supported	◗ = Partly Supported	<blank> = Not Supported	

Features of Azure Backup Components

Copyright © by EC-Council. All Rights Reserved. Reproduction is Strictly Prohibited.

Security of Azure Backup

Azure Backup offers robust security for both data at rest and in transit. The following security capabilities of Azure Backup help in safeguarding backup data and meeting the business requirements.

Various security controls are in-built in the Azure Backup service that can detect, respond, and prevent vulnerabilities. Encryption is performed on the entire backup data from the customer servers to the RSV utilizing AES256. Backup data is transferred through HTTPS. The Azure customer can only decrypt the stored data in the RSV using a passphrase.

The Azure Backup data is automatically encrypted utilizing platform-managed keys and the customer would not have to take any action for enabling it. However, the customer can encrypt the backed-up data utilizing customer-managed keys that are stored in Azure Key Vault. This will be applied to all the workloads that are backed-up in the RSV. The backup data is encrypted using Azure Disk Encryption, which utilizes dm-crypt on Linux VMs and BitLocker on Windows VMs

The storage accounts that are utilized by RSVs are completely separated and the unauthorized users cannot not have access to it. Azure Backup management operations can allow access. Azure RBAC provides fine-grained access to Azure Backup. Using Azure RBAC, the administrator can segregate the duties in a team and grant limited access to users to accomplish their work.

In Azure Backup, the backup data will be stored in Azure Storage and the guest would not have any direct access to the backup storage. In VM backup, Azure creates the backup snapshot and stores it; the guest does not have any involvement except quiescing the

Certified Cloud Security Engineer Copyright © by **EC-Council**
All Rights Reserved. Reproduction is Strictly Prohibited.

workload for application-consistent backups. The guest cannot delete or tamper with the existing backup data.

Feature	Azure Backup Agent	System Center DPM	Azure Backup Server	Azure IaaS VM Backup
Recovery Services Vault	●	●	●	●
Disk Storage		●	●	
Tape Storage		●		
Compression (in Recovery Services Vault)	●	●	●	
Incremental Backup	●	●	●	●
Disk deduplication		◗	◗	

Key ● = Supported ◗ = Partly Supported <blank> = Not Supported

Figure 9.89: Features of Azure Backup Components

Features of Azure Backup Components

The robust features of Azure Backup allow the customer to retrieve their data even after deletion. If the customer deletes the backup of a VM accidentally, then the backup data is stored in Azure Backup for 14 more days, which can be recovered without any loss.

Azure Backup has in-built monitoring and alerting capabilities that allow the user to view and configure all the events that happen in Azure Backup. The monitoring and reporting tools of Azure Backup will send an alert to the user as soon as any unauthorized access, malicious or suspicious activity is performed in the backup environment.

Azure Backup provides a complete set of certifications and attestations that allows an enterprise to comply with regional, national, and industry-specific requirements.

Certified Cloud Security Engineer Copyright © by **EC-Council**
All Rights Reserved. Reproduction is Strictly Prohibited.

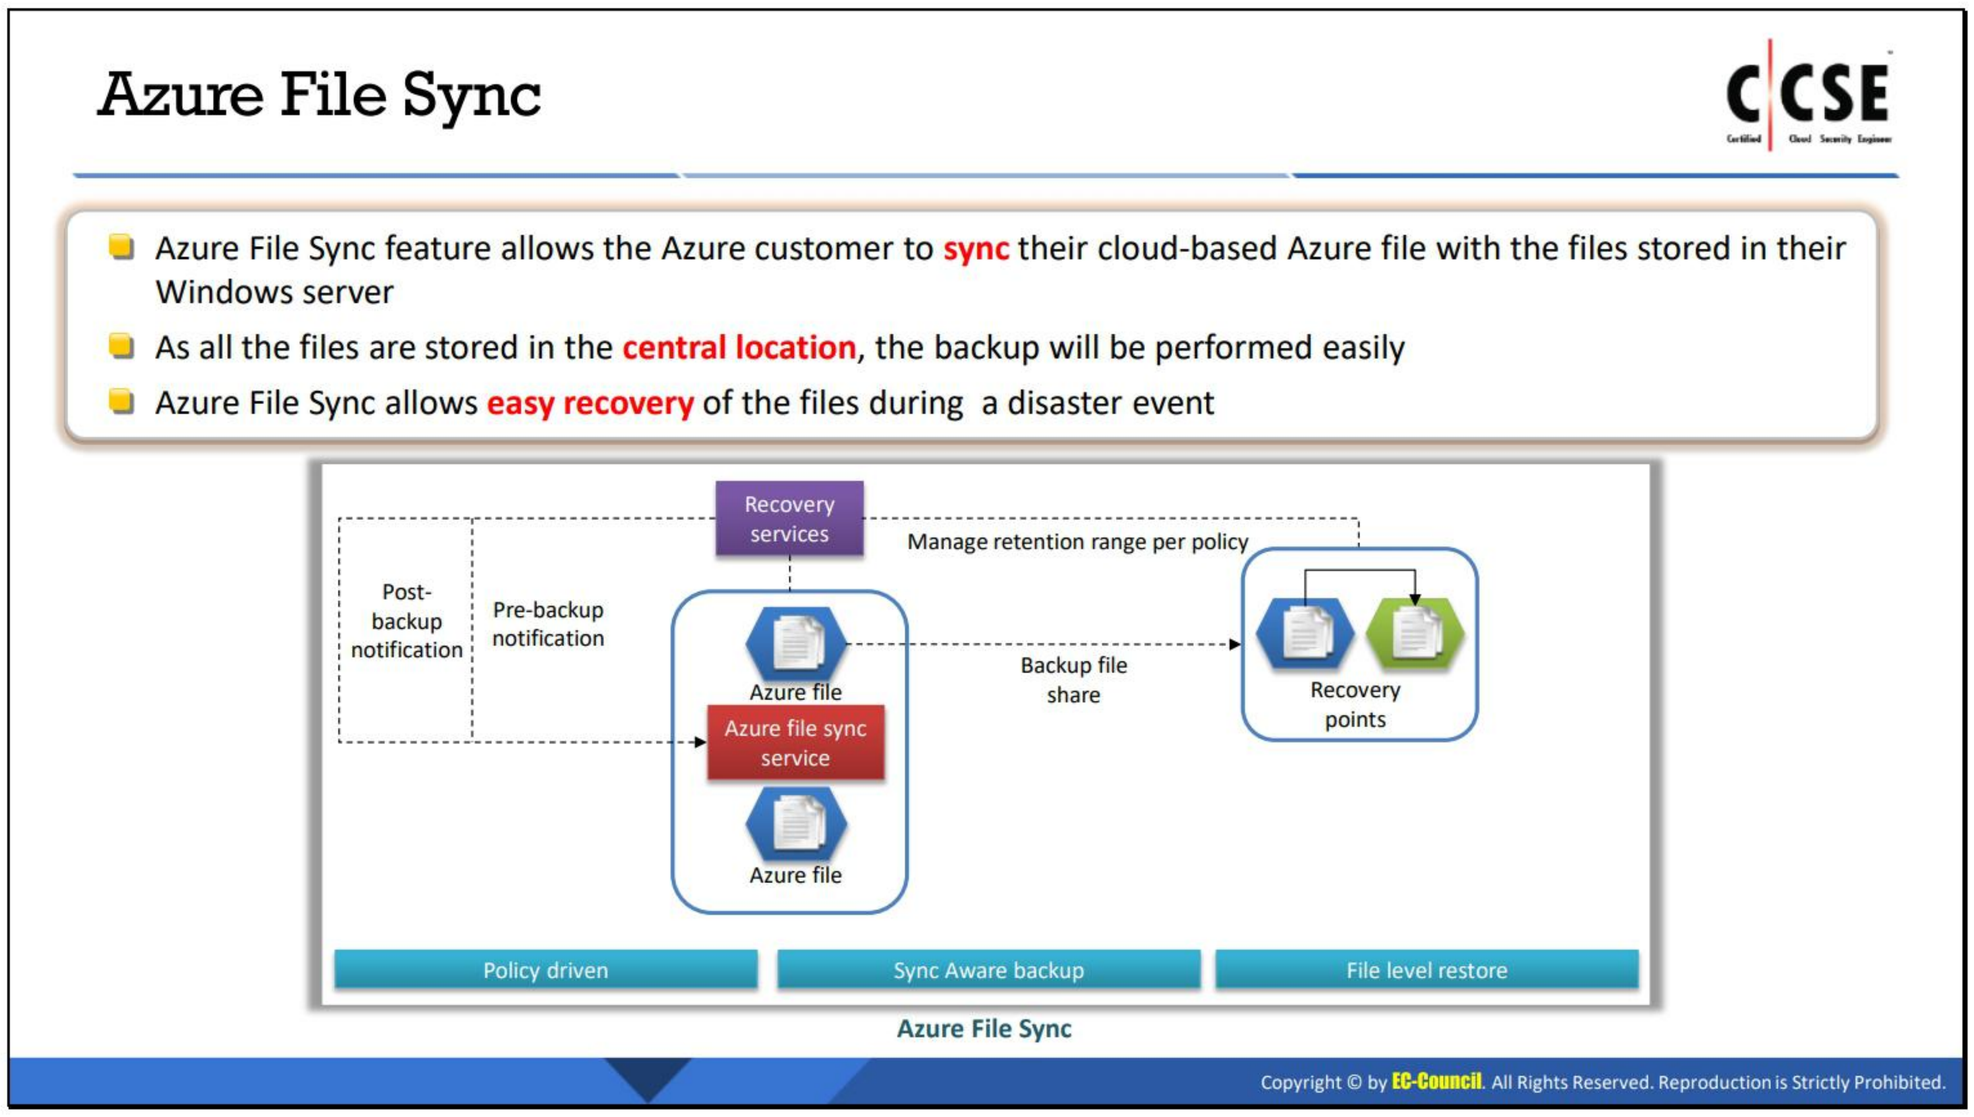

Azure File Sync

File Sync is a unique feature provided by Azure in Azure Files that eliminates the need for choosing between the server and the cloud. The Azure File Sync feature allows Azure customers to sync their cloud-based Azure Files with the files stored in their Windows servers. The File Sync feature allows them to save their recently accessed or newer files in the on-premises server, while leveraging the cloud for accessing files they may require.

With Azure File Sync (AFS), cloud administrators can integrate their Windows servers into Azure Files, eliminating the need for transferring files. This makes sharing of files with clients and remote employees easier. File Sync will move older files that are no longer accessed or rarely accessed to the cloud automatically, thereby freeing up storage space in the on-premises server.

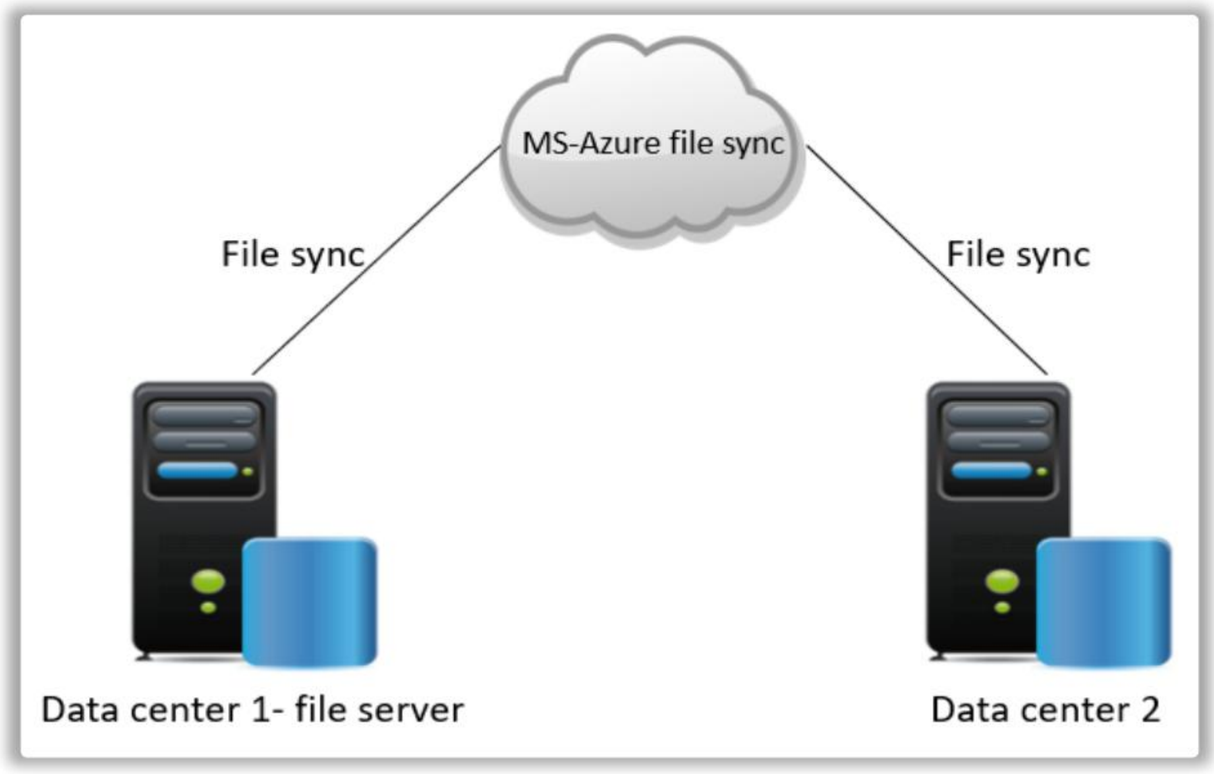

Figure 9.90: Azure File Sync

Certified Cloud Security Engineer Copyright © by **EC-Council**
All Rights Reserved. Reproduction is Strictly Prohibited.

AFS can help organizations do the following:

- Eliminate local backup and disaster recovery.
- Centralize file services inside Azure Storage.
- Cache data in several locations for local, fast performance.

Features of AFS

AFS provides the following features:

- **File synchronization**: AFS synchronizes file shares using Azure Files. It synchronizes everything to a general storage account. Azure File shares are created in the general storage account, keeping the Azure copy as the master copy.
- **File sharing inside the office**: AFS allows file shares across multiple offices by synchronizing the Azure master shares across various file servers present in each office.
- **File sharing outside the office**: AFS allows sharing of files across continents. Files can be replicated in any Azure region and servers can be configured to connect to the closest version of file shares.
- **File backups**: AFS uses a central location to save everything. As all the files are stored in the central location, the backups are performed easily. Backups of Azure Files can be saved to the cloud using Azure Backup. This method would create backups of the storage file account in an RSV.
- **File recovery**: AFS allows easy recovery of the files during a disaster event. If an on-premises file server is down or lost, the cloud administrator can connect a new server to Azure and download files again.
- **Backup with tiering**: Azure Backup takes care of backing up files to provide safe file storage. Ensure to backup files using Azure Backup only after enabling cloud tiering. With cloud tiering, AFS stores recently accessed files only on local servers.

hiile01.ir

Certified Cloud Security Engineer Copyright © by **EC-Council**
All Rights Reserved. Reproduction is Strictly Prohibited.

Linux Application Consistent Backup

- Linux application consistent backup is a powerful framework that allows the customer to obtain **application native backup** by invoking native backup APIs during virtual machine snapshot
- Azure backup service will take this backup to offsite recovery services vault and **secure** it against VM corruption, disaster recovery scenarios, or VM deletion

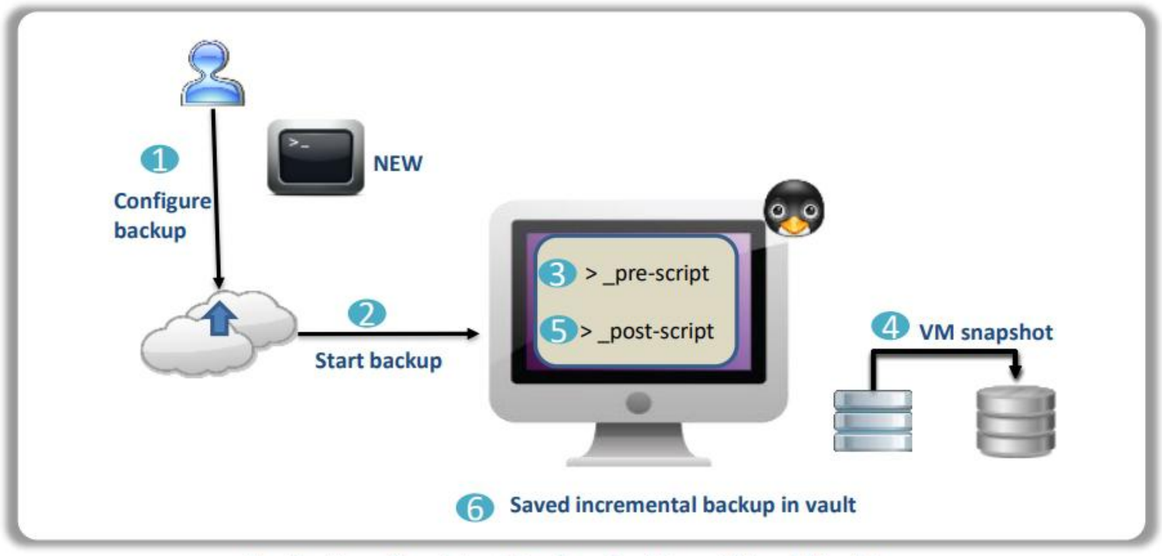

Application Consistent Backup for Linux Virtual Machines

Copyright © by EC-Council. All Rights Reserved. Reproduction is Strictly Prohibited.

Linux Application-Consistent Backup

Application consistency refers to starting the application at the time of VMs boot after being restored. Application-consistent backups ensure that the backup data is transactionally consistent. To ensure transactional consistency, the application should be quiesced, and while taking a VM snapshot, there should not be any unfinished transactions.

To make sure that there is application-consistent VM backup, Windows provides Volume Snapshot Service (VSS) framework; however, in Linux, there is no such framework. Due to the high usage of Azure Linux VMs, there is a huge requirement for backing up critical applications that are running in Azure VMs. To make sure that there is application-consistent VM backup for Linux application, a framework similar to the VSS-like generic framework has been introduced. To ensure that the Linux VMs are application-consistent, the user can utilize the Linux pre-script and post-script framework for taking application-consistent backups. During the VM backup process, this framework provides the flexibility to run custom pre- and postscripts. Using these scripts, application IOs can be quiesced at the time of VM snapshot, which guarantees application consistency.

Linux application-consistent backup is a powerful framework that allows the customer to obtain application-native backups by invoking native backup APIs during VM snapshots. Azure Backup service will take this backup to the off-site RSV and secure it against VM corruption, disaster recovery scenarios, or VM deletion.

This framework gives a lot of power to the users; thus, it is very important to secure and provide permission to only root users for critical JSON and script files. Ensure that the root user is the owner of these files and also ensure that the root user has read, write, and execute permissions for the script file.

Certified Cloud Security Engineer Copyright © by **EC-Council**
All Rights Reserved. Reproduction is Strictly Prohibited.

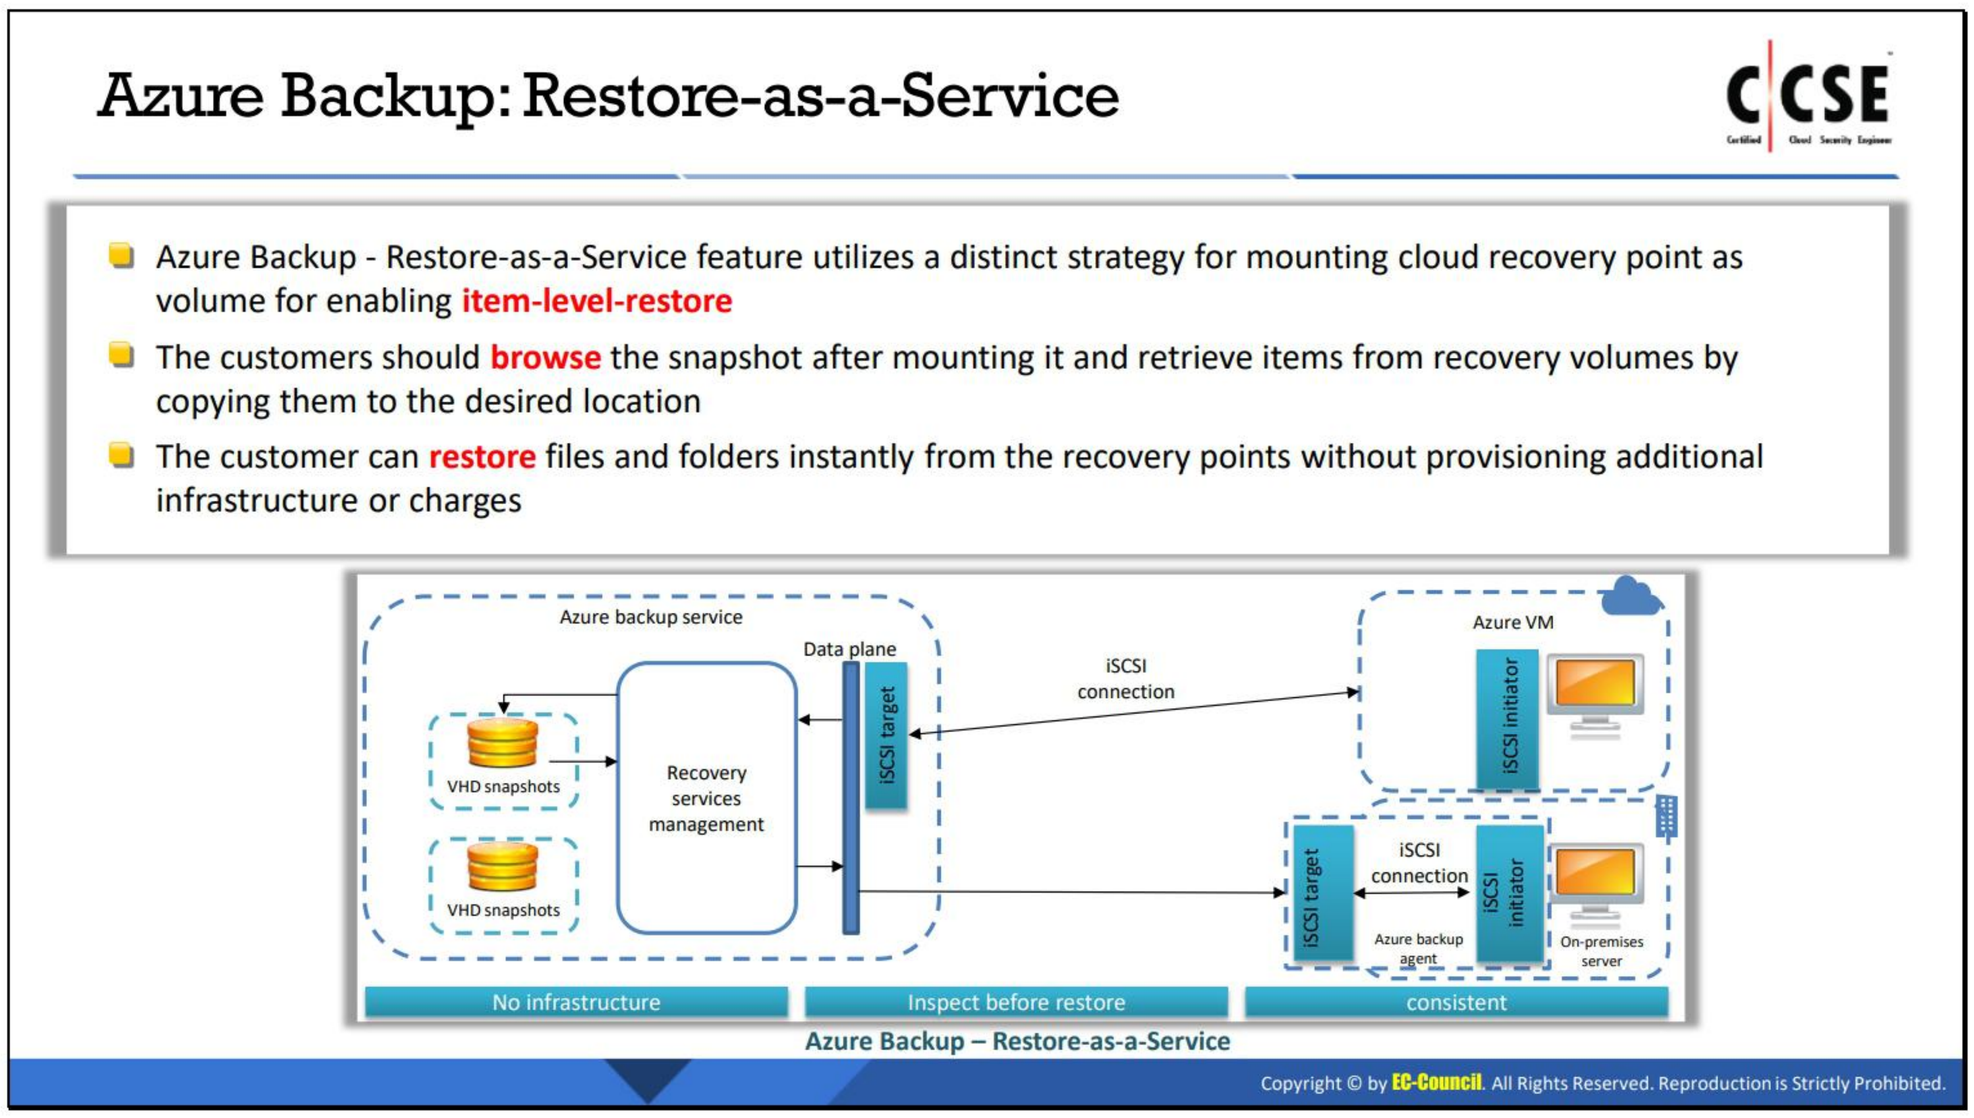

Azure Backup: Restore-as-a-Service

Restore-as-a-Service is a part of Azure Backup that provides instant backup restore. The Azure Backup-Restore-as-a-Service feature utilizes a distinct strategy for mounting cloud recovery point as volume for enabling item-level restore. This feature allows instant restoration of files and folders to a Windows OS from cloud-based recovery points. To benefit from this feature, a customer simply needs to backup their VMs using the Azure Backup service.

After mounting the snapshot, the customer needs to browse through it and recover items by copying them from recovery volumes to the desired location. The customer can restore files and folders instantly from the recovery points without provisioning additional infrastructure or charges. This is because the restore-as-a-service feature does not require the provision of any additional infrastructure and provides free egress from Azure.

Figure 9.91: Azure Backup- Restore-as-a-Service

Certified Cloud Security Engineer Copyright © by **EC-Council**
All Rights Reserved. Reproduction is Strictly Prohibited.

Key Benefits of Azure Backup Restore-as-a-Service

- **Single restore mechanism:** The restore-as-a-service model provides a single restore mechanism for all backup sources. It uses a unified approach to recover files and folders that are backed-up from various sources, including cloud and on-premises. Cloud administrators can use instant restore when they backup on-premises data to the cloud as well as protect VMs by using Azure VM backup.
- **Instant file recovery:** Recover files instantly from cloud backups of on-premises file servers or Azure VMs. Instant restore significantly reduces the time required for file recovery.
- **Option to review files before restoring:** Restore-as-a-service allows cloud administrators to directly open and see application files (e.g., SQL) from the cloud recovery point snapshots before restoring. This eliminates the need for restoring and attaching them to live app instances.
- **Recover any combination of files:** The restore-as-a-service feature allows the restoration of multiple files from several folders to any target (e.g., local server).
- **Flexibility in restoring data:** With restore-as-a-service, system state files can be restored from Azure without multiple changes.

hide01.ir

Certified Cloud Security Engineer Copyright © by **EC-Council**
All Rights Reserved. Reproduction is Strictly Prohibited.

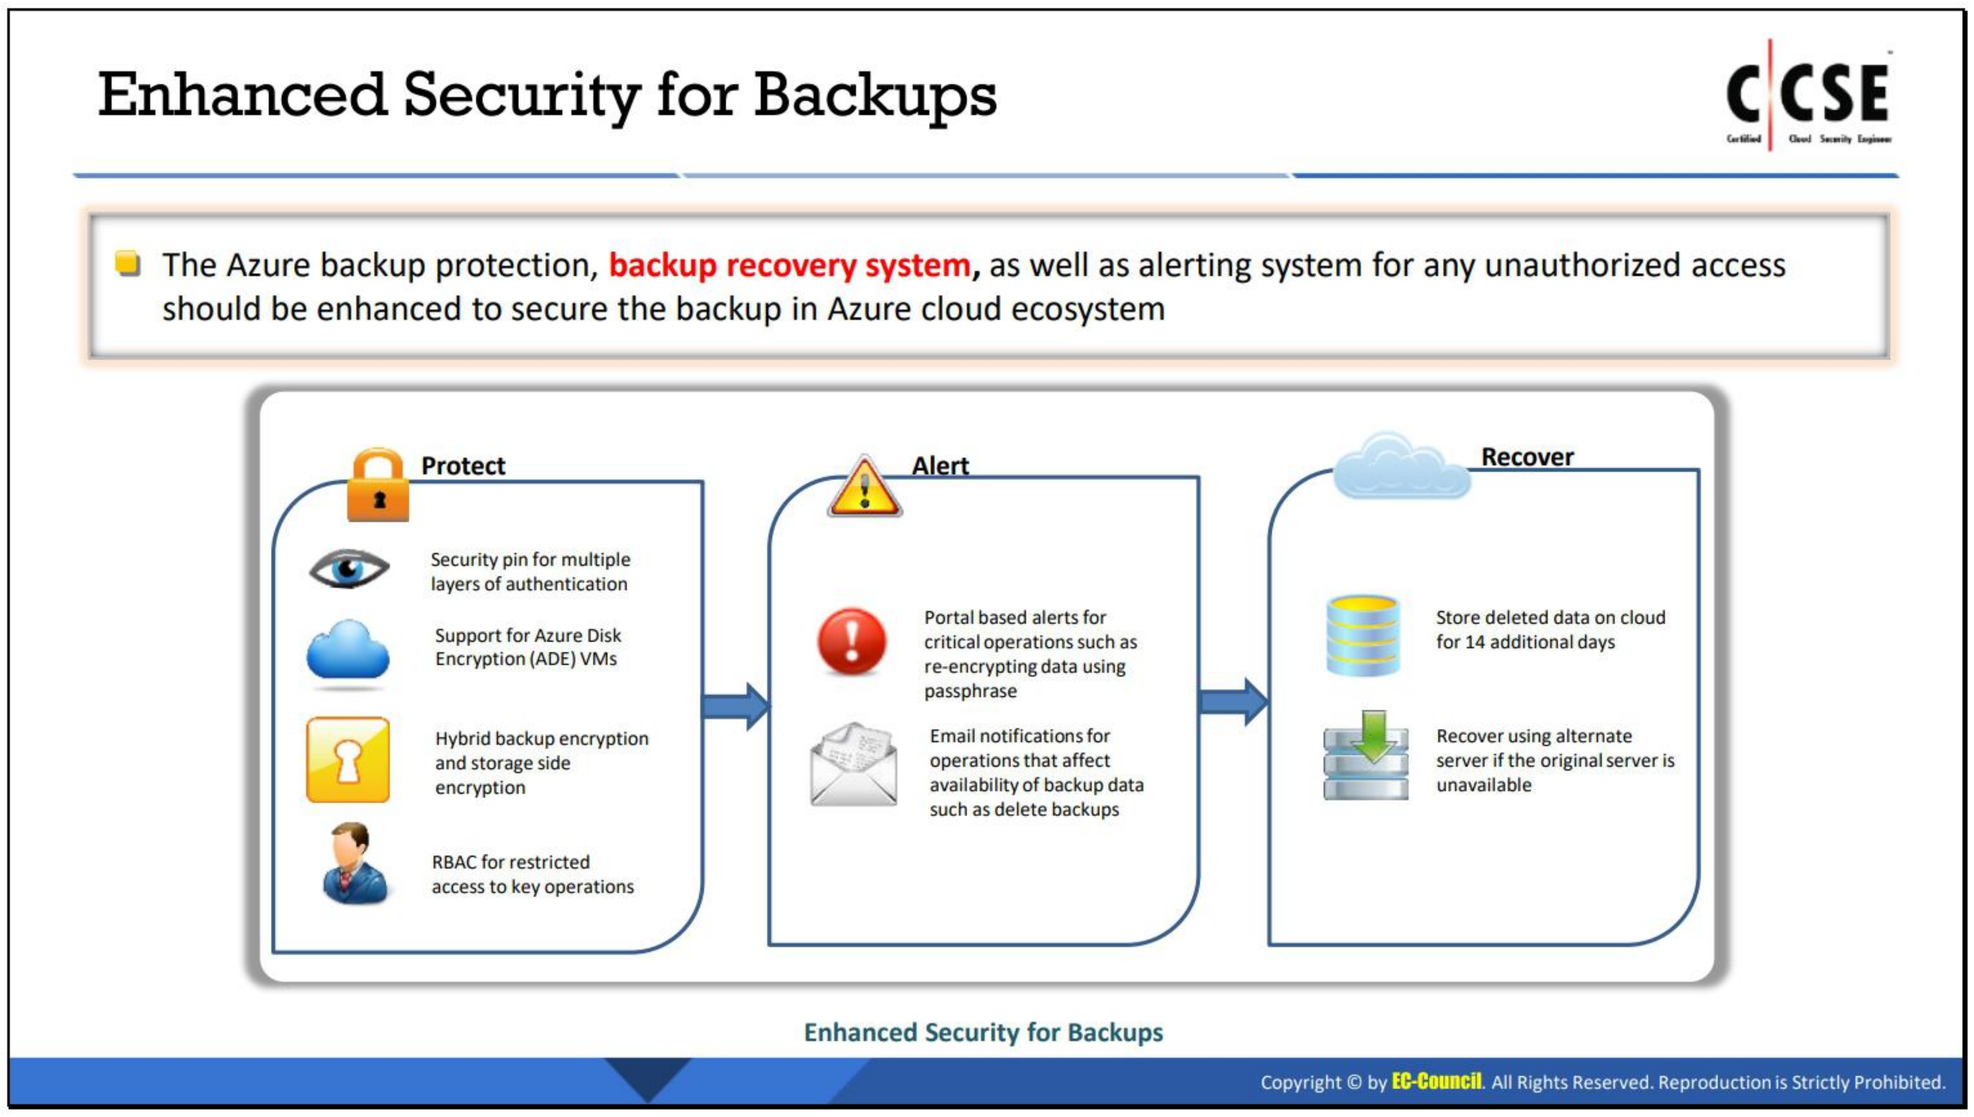

Enhanced Security for Backups

Production servers of organizations are mostly affected by malware attacks that try to re-encrypt data or delete it permanently. Backups also get affected if production data is compromised. Azure Backup offers certain security capabilities that help protect cloud backups. Azure customers can use these security features to protect their backups and recover data by using cloud backups when production and backup servers are attacked or compromised.

Figure 9.92: Enhanced Security for Backups

To equip organizations with a strong backup solution and enable them to be prepared against attacks, Microsoft has introduced backup security features that are built on the following three principles:

- **Prevention**: Azure Backup provides the below security features that can be utilized to prevent backups from attacks.
    - **Security PIN for several layers of protection**: For critical operations (e.g., Change Passphrase and Delete Backup Data), Azure Backup has added a new authentication layer. These operations require a security PIN, which is provided only to users that have valid credentials.

Certified Cloud Security Engineer Copyright © by **EC-Council**
All Rights Reserved. Reproduction is Strictly Prohibited.

- **Support for Azure Disk Encryption VMs**: It supports backup and restore of Azure Disk Encrypted (ADE) VMs (VMs having their data disks or OS encrypted with Azure Disk Encryption) as well as customer master key (CMK)-encrypted VMs.
- **Hybrid backup encryption and storage-side encryption**: Azure Backup supports storage-side encryption and hybrid backup encryption (data is encrypted prior to leaving the user site) to secure backups on Azure.
- **RBAC for controlled access**: Azure Backup allows use of RBAC to control managed operations via fine-grained access. With RBAC, assign users only the amount of access required to perform their duties.

- **Alerting**: Azure Backup provides the below security features that alert regarding attacks or suspicious activity.
  - **Alerts for critical operations**: Azure Backup allows cloud administrators to view and configure actions for events occurring in Azure Backup. It offers portal-based alerts for critical operations such as data re-encryption using a passphrase. Additionally, its monitoring and alerting tools can be used to get alerts for any suspicious, malicious, or unauthorized activity. These alerts help users detect attacks soon after they occur.
  - **Email notifications for operations affecting backup data availability**: It also offers email notifications for all operations that impact the availability of backup data (e.g., delete backup).

hide01.ir

- **Recovery**: Azure Backup provides the below security features that address issues related to backup recovery.
  - **Deleted data is stored for 14 days after deletion**: It retains deleted data for 14 days to protect backup data from unintentional deletes. If a user deletes the backup data of a VM by mistake, the backup data is not lost. It can be recovered within 14 days after deletion.
  - **Alternate servers for backup recovery**: Azure Backup allows recovering backup data from an alternate server if the primary backup server is unavailable. It always maintains a minimum number of recovery points so that there are always enough recovery points that can be used to recover data.

Certified Cloud Security Engineer Copyright © by **EC-Council**
All Rights Reserved. Reproduction is Strictly Prohibited.

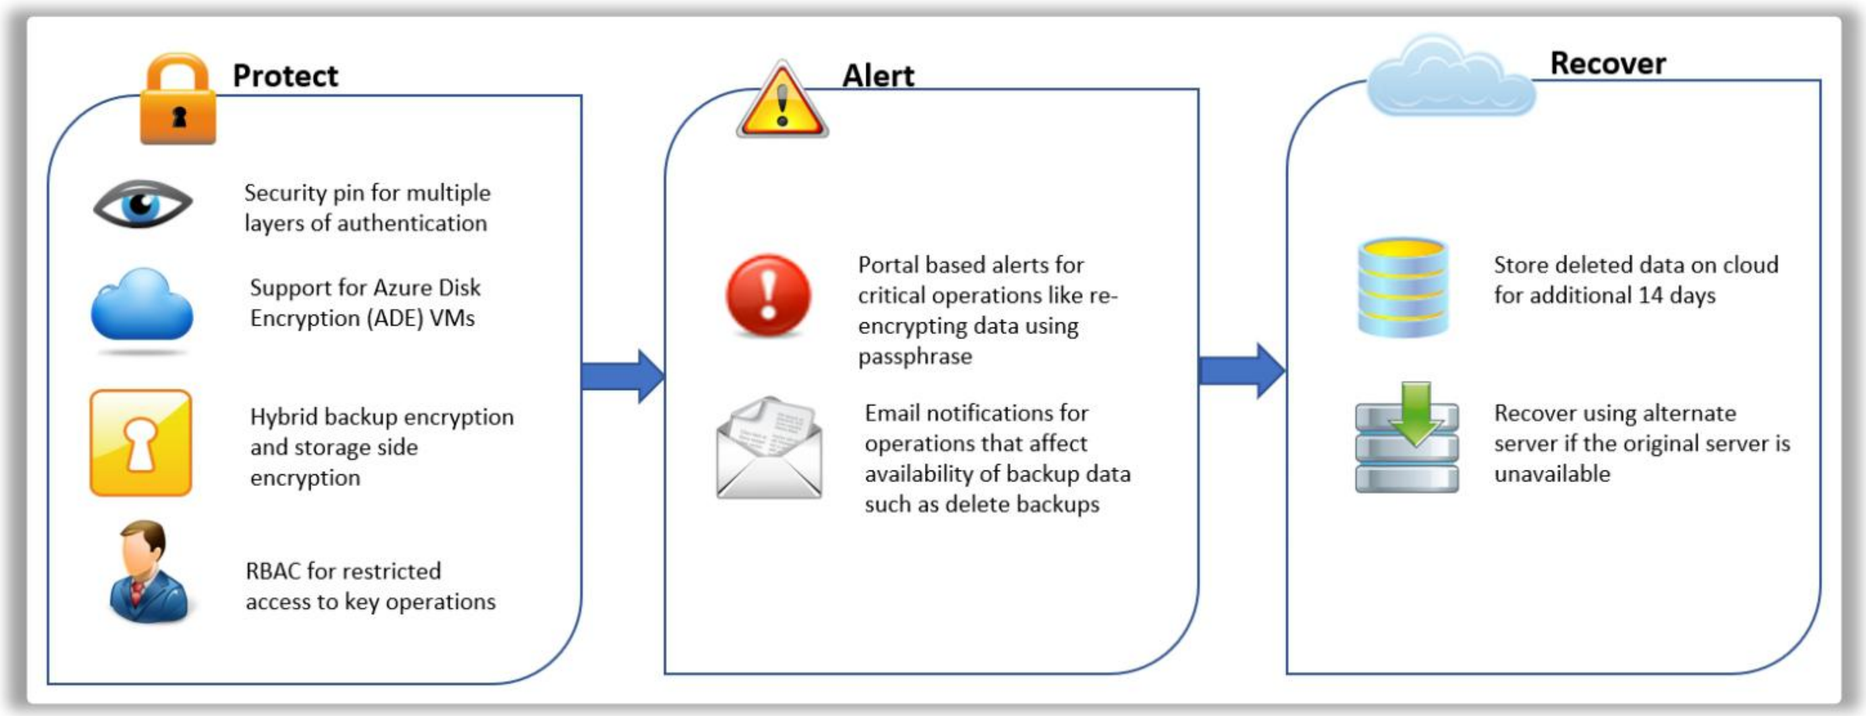

Figure 9.93: Enhanced Security for Backup

hide01.ir

Certified Cloud Security Engineer Copyright © by **EC-Council**
All Rights Reserved. Reproduction is Strictly Prohibited.

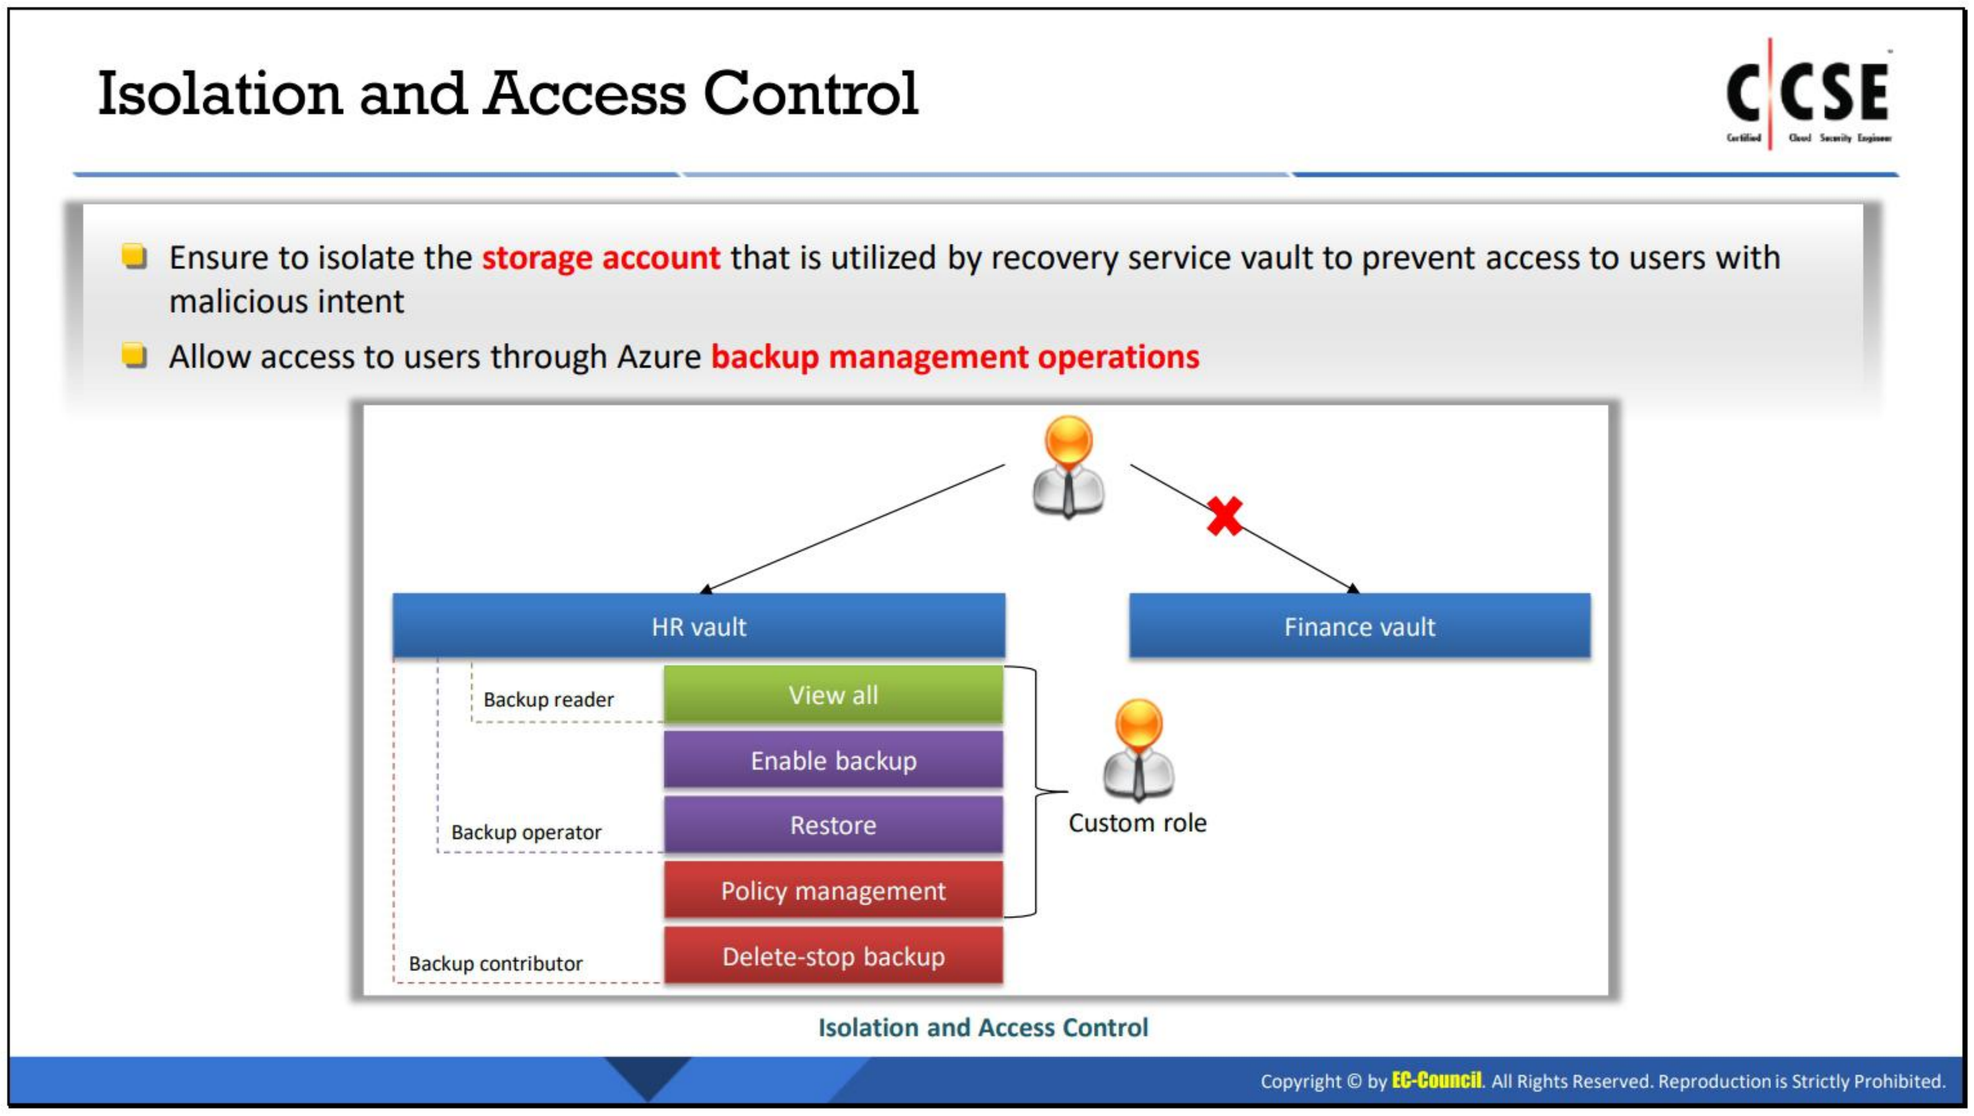

Isolation and Access Control

Azure implements sufficient measures to manage and control the identity and user access to backups. It isolates storage accounts utilized by RSV and restricts user access to it. Users cannot perform any malicious activity because all access to backup data is granted via Azure Backup management operations.

Cloud administrators can control managed operations via fine-grained access by leveraging Azure RBAC. RBAC enables them to divide responsibilities within their team and permit only the required amount of access to users necessary for their jobs.

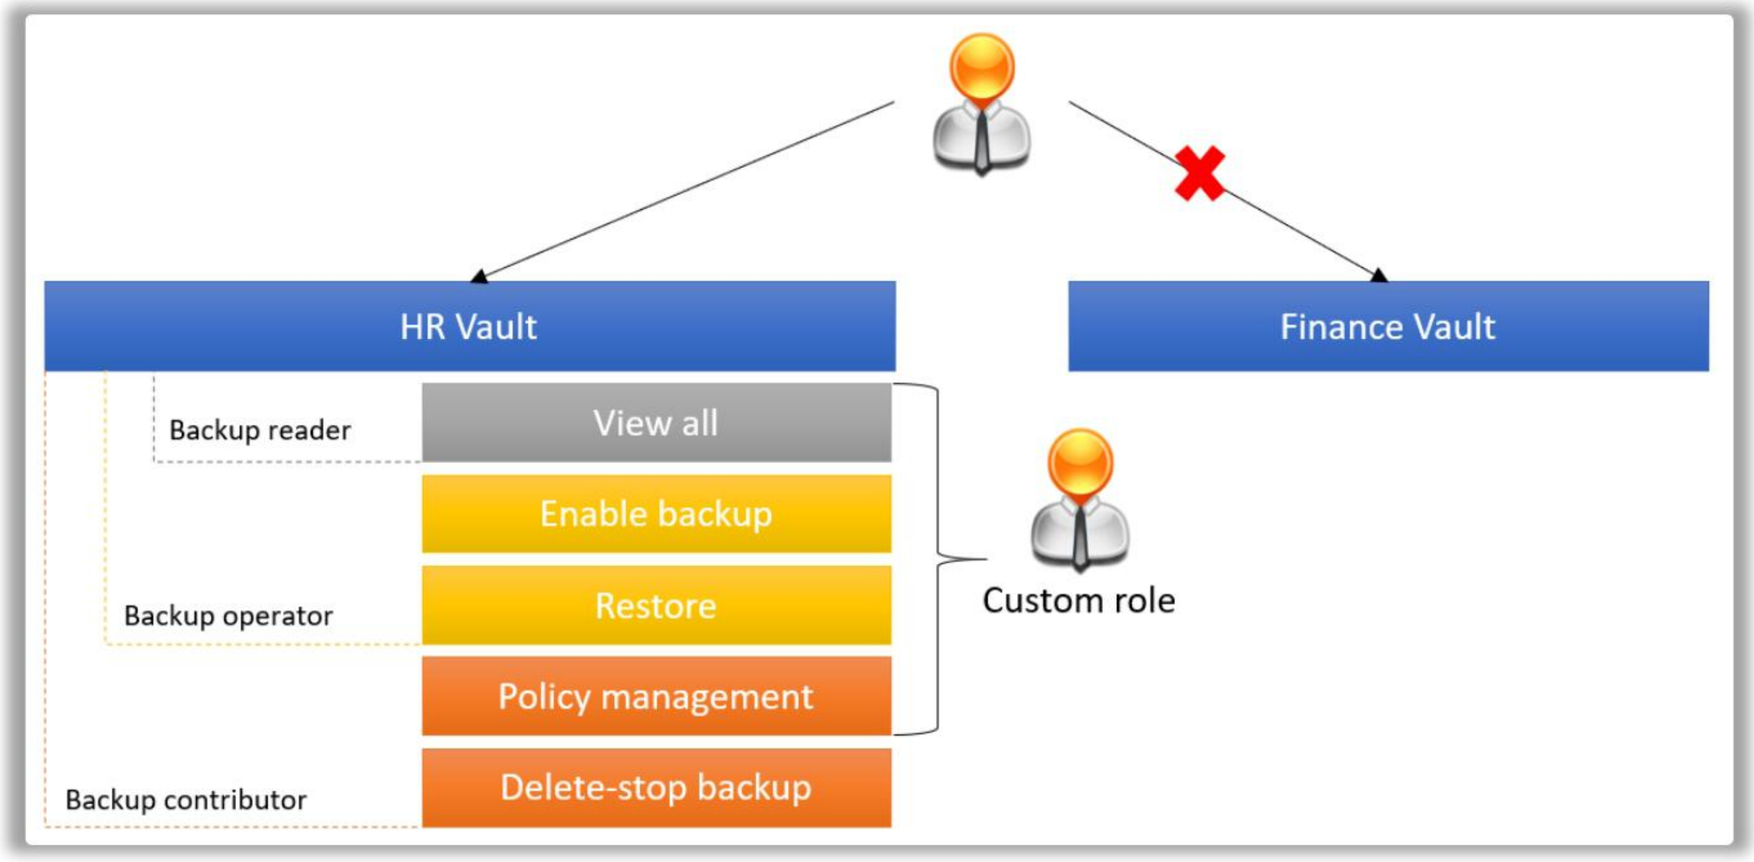

Figure 9.94: Job Responsibilities

Certified Cloud Security Engineer Copyright © by **EC-Council**
All Rights Reserved. Reproduction is Strictly Prohibited.

To control backup management operations, Azure Backup offers the following built-in roles:

- **Backup contributor**: A backup contributor is the administrator of backup management. This role can perform all backup management operations such as creation and management of backups, but cannot delete RSV and provide access to others. Users with this role have the highest privileges as they can perform policy management and destructive operations such as deleting a backup and stopping a backup.
- **Backup operator**: A backup operator can perform all operations that a backup contributor can do, except performing the backup deletion and managing backup policies. They cannot perform destructive operations (e.g., **remove the registration of on-premises resources** and **stop a backup with delete data)**. Users with this role have limited privileges as they can create and restore backups and view operations but cannot do other operations.
- **Backup reader**: A backup reader has permissions to only see the backup management operations. It is a monitoring role, which cannot be used to make any changes to backup services. Users with this role have the least privilege as they can only view operations happening in Azure Backup.

hide01.ir

Certified Cloud Security Engineer Copyright © by **EC-Council**
All Rights Reserved. Reproduction is Strictly Prohibited.

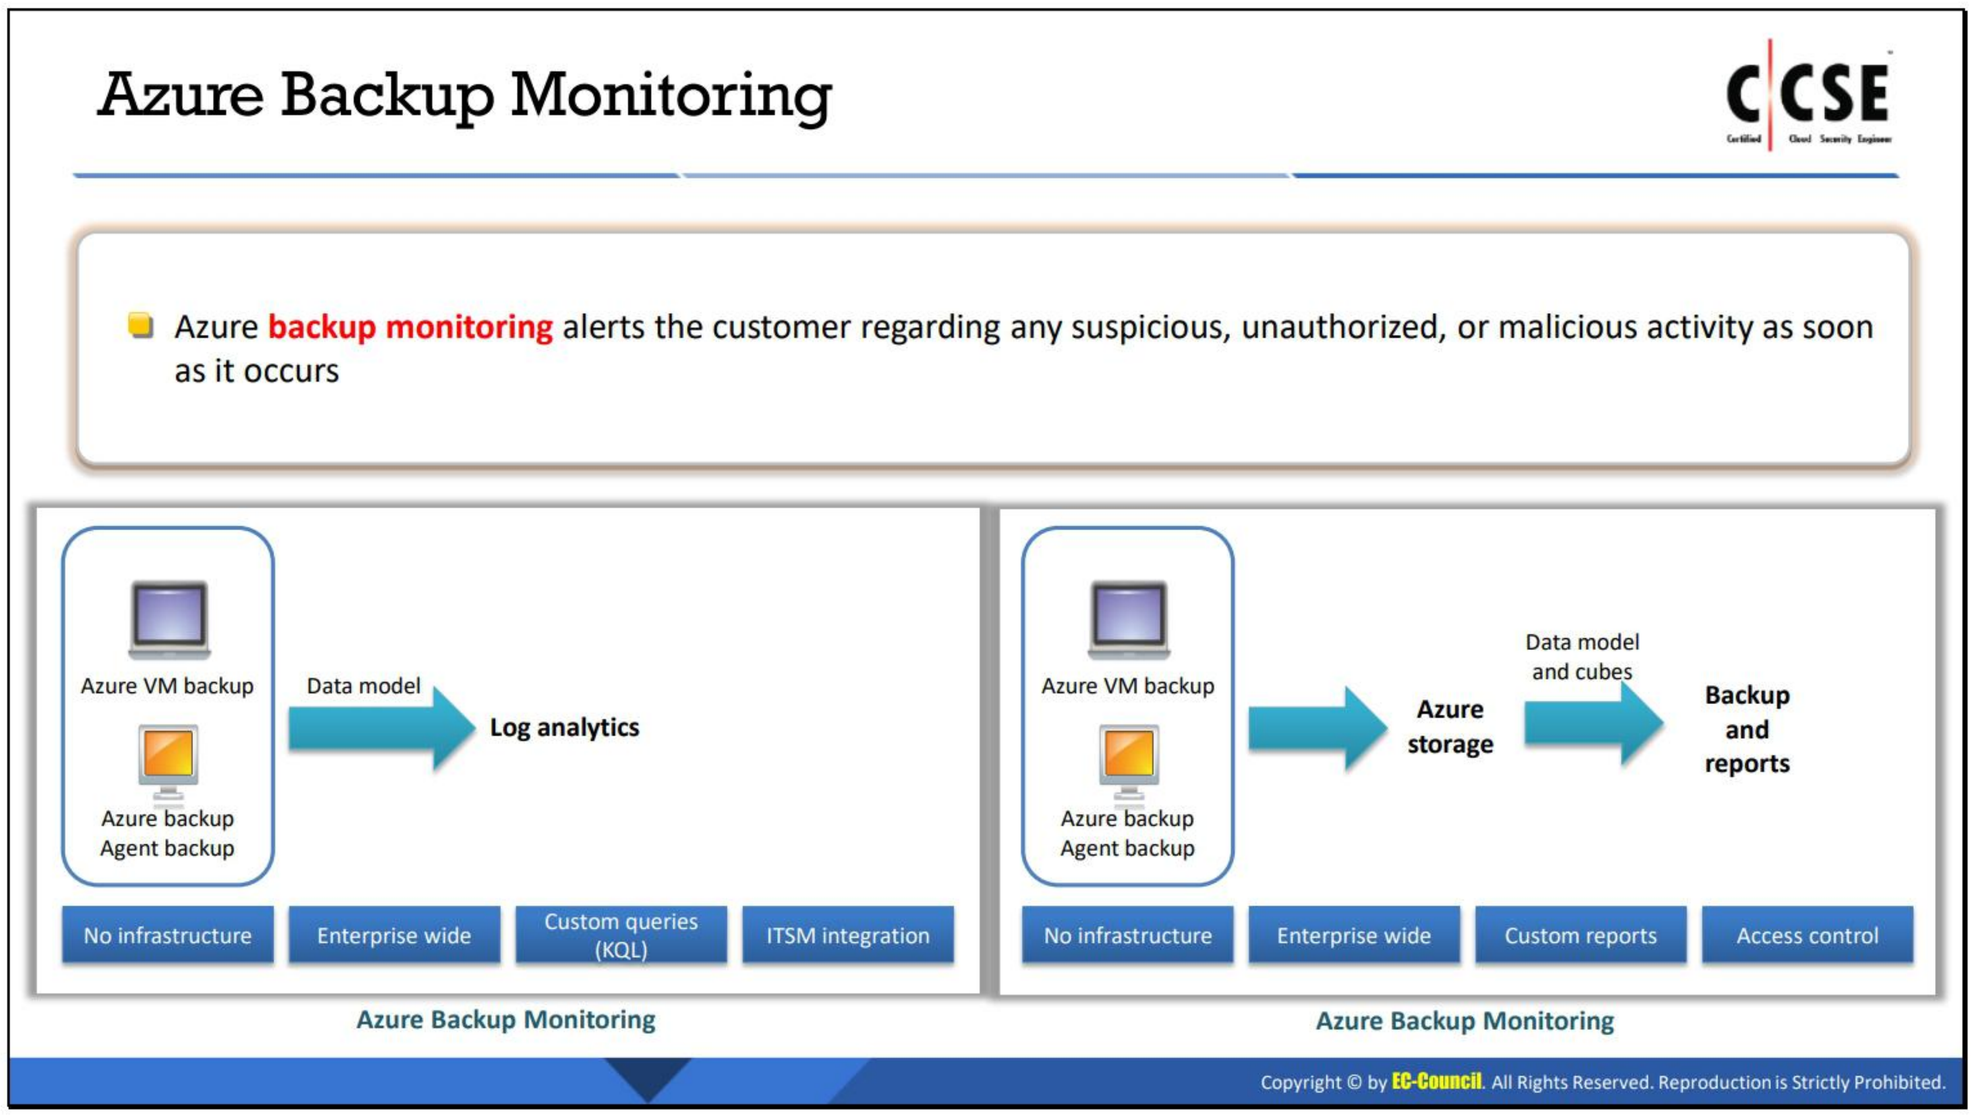

Azure Backup Monitoring

Various built-in monitoring and alerting capabilities are offered by Azure Backup to see and configure actions for events that are related to Azure Backup. It provides backup reports that help cloud administrators track usage, identify key trends at various levels of granularity, and audit backups and restores. They can use the monitoring and alerting tools of Azure Backup to get alerts when something suspicious is detected. Azure Backup monitoring alerts the customer regarding any suspicious, unauthorized, or malicious activity as soon as it occurs**.**

Monitor Azure Backup Protected Workloads Through Log Analytics

Source: https://www.microsoft.com

Azure Backup allows organizations to monitor important backup parameters across RSVs and subscriptions by using log analytics. For custom monitoring requirements for all workloads, configure alerts and actions. This LA-based solution will cover all workloads saved or protected by Azure Backups such as System Center DPM-A, Azure VMs, and file-folder backup from the backup agent.

Configure Diagnostic Settings

Follow these steps to configure diagnostic settings:

- Go to **Diagnostic setting** from **Azure Monitor** or **Azure Recovery Services Vault**.
- Choose **Send data to log analytics** in the **Diagnostic setting** window. Select the relevant LA workspace and log accordingly, choose **AzureBackupReport** and then choose **Save**.

  Ensure to select the same workspace for all vaults because it will help get a centralized view. Wait for 24 hours after completing the configuration for initial data push.

Certified Cloud Security Engineer Copyright © by **EC-Council**
All Rights Reserved. Reproduction is Strictly Prohibited.

Deploy Solution Template

Deploy Azure quick-start template after the data is in the workspace to visualize the monitoring data. Configure it and install the template. Ensure that the workspace name, resource group, and workspace location remain the same.

An overview tile appears for Azure Backup in the workspace dashboard after the template is deployed. To go to the solution dashboard, choose the overview tile.

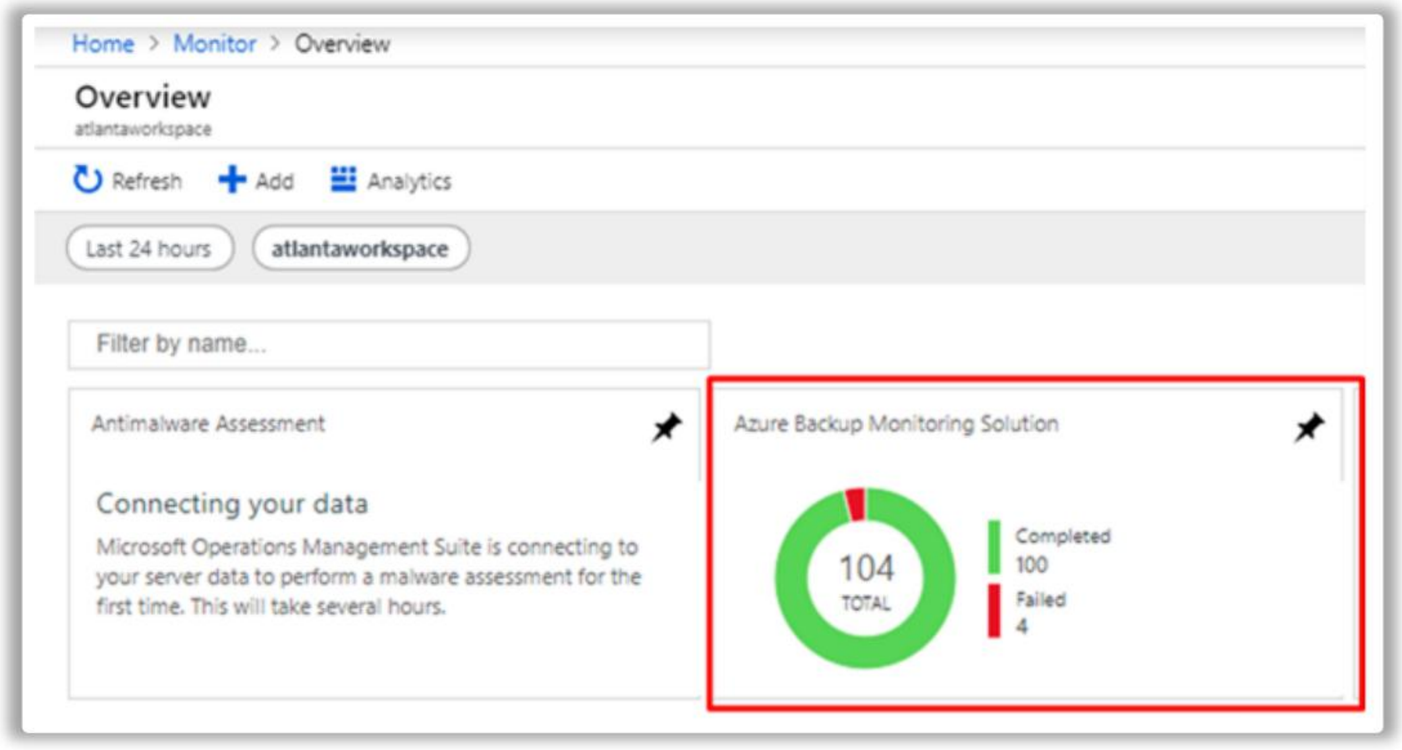

Figure 9.95: Overview Console

Monitor Azure Backup Data

- **Monitor backups and restores:** Monitor daily backups regularly for all workloads protected by Azure Backup. With this, even log backups for SQL databases can be monitored.

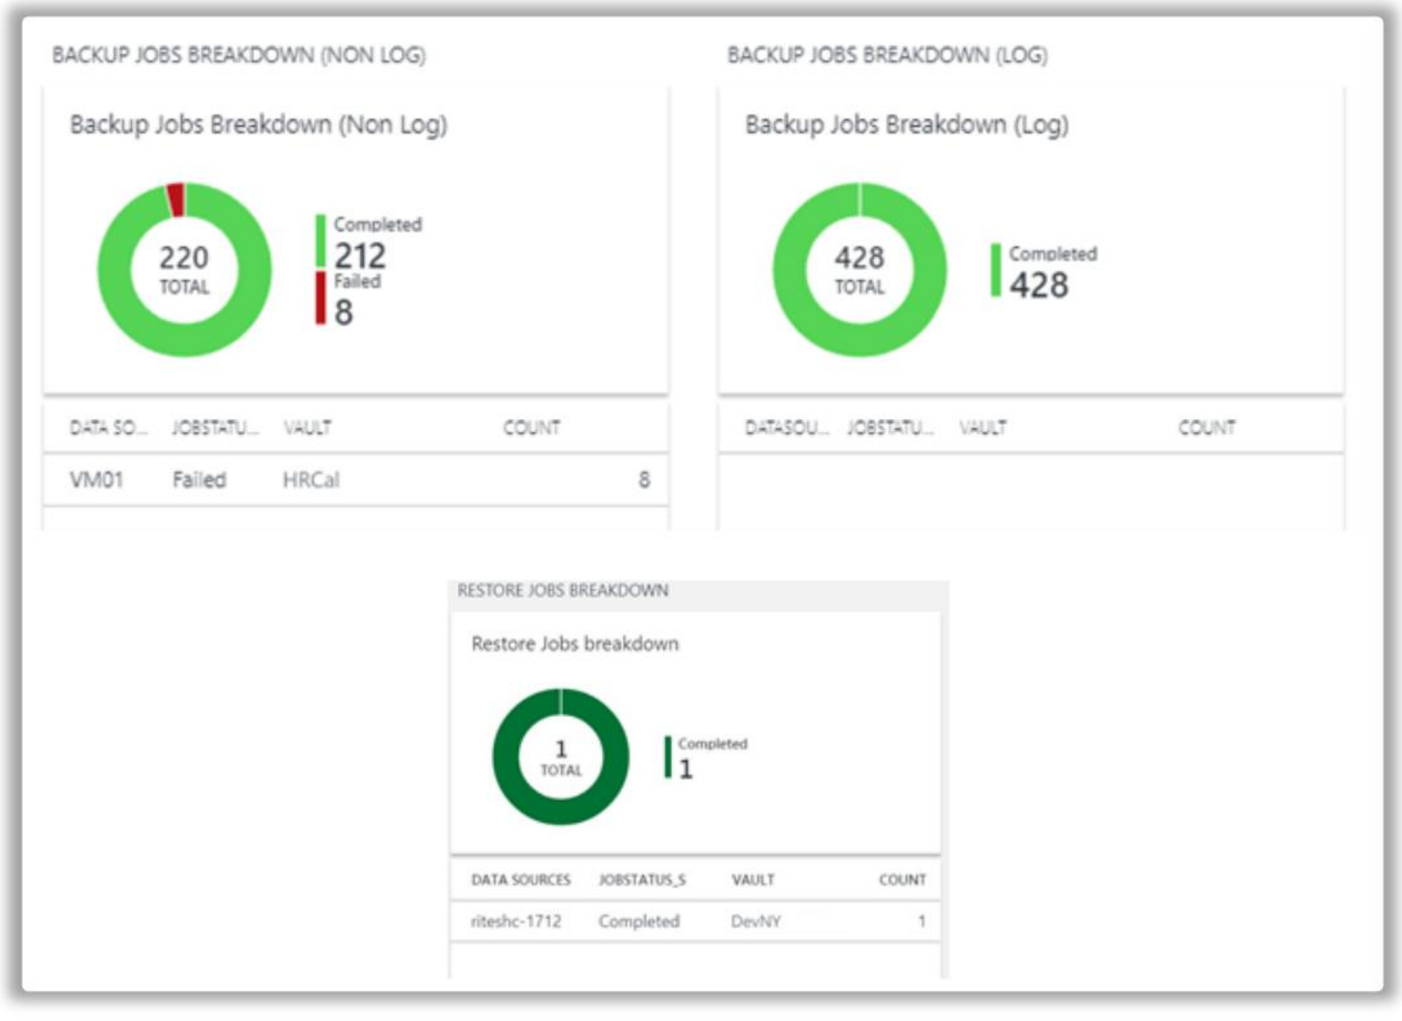

Figure 9.96: Monitoring Azure Backup and Restore

Certified Cloud Security Engineer Copyright © by **EC-Council**
All Rights Reserved. Reproduction is Strictly Prohibited.

- **Monitor all data sources:** By using the active data sources graph, monitor a reduction or spike in the number of backed-up data sources.

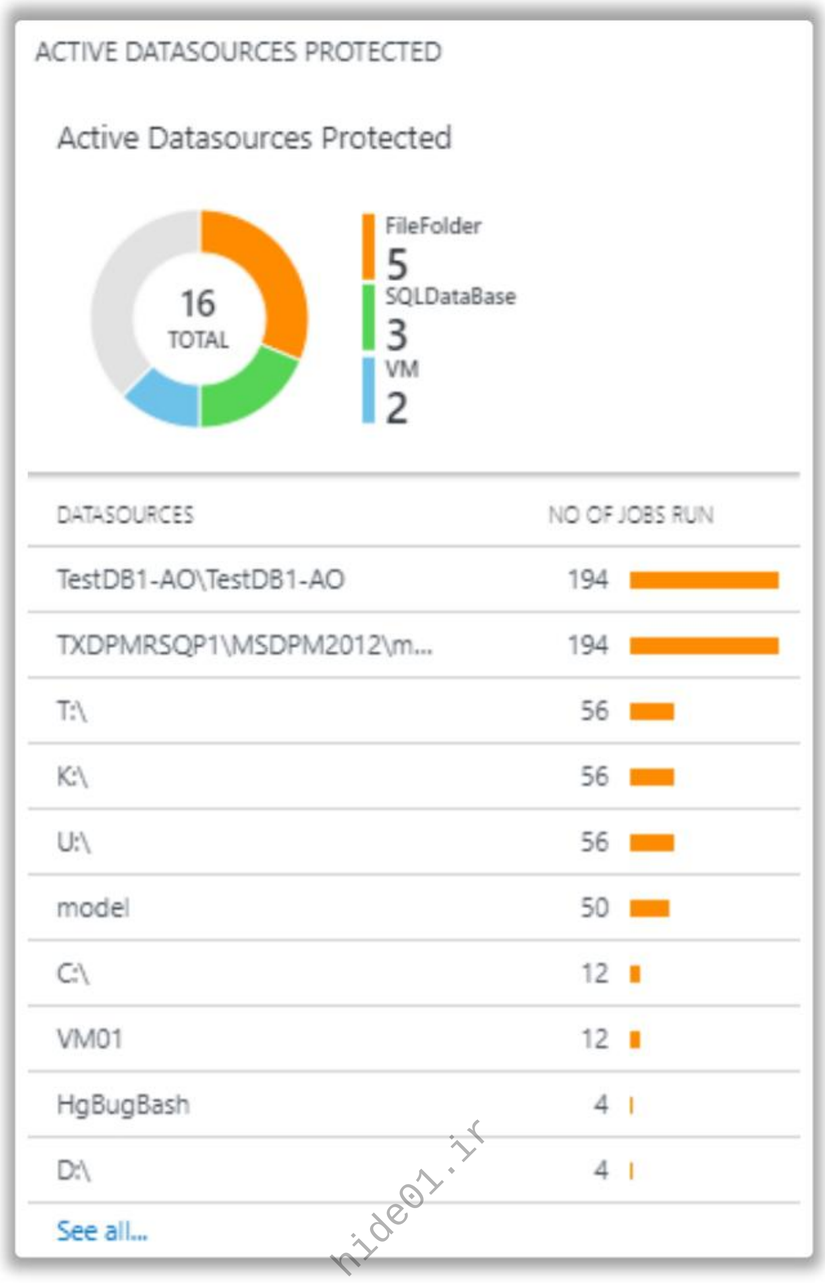

hide01.ir

Figure 9.97: Monitoring All Data Sources

- **Monitor Azure Backup alerts:** If a backup/restore job fails, Azure Backup generates alerts that can be viewed in a single place (as shown below).

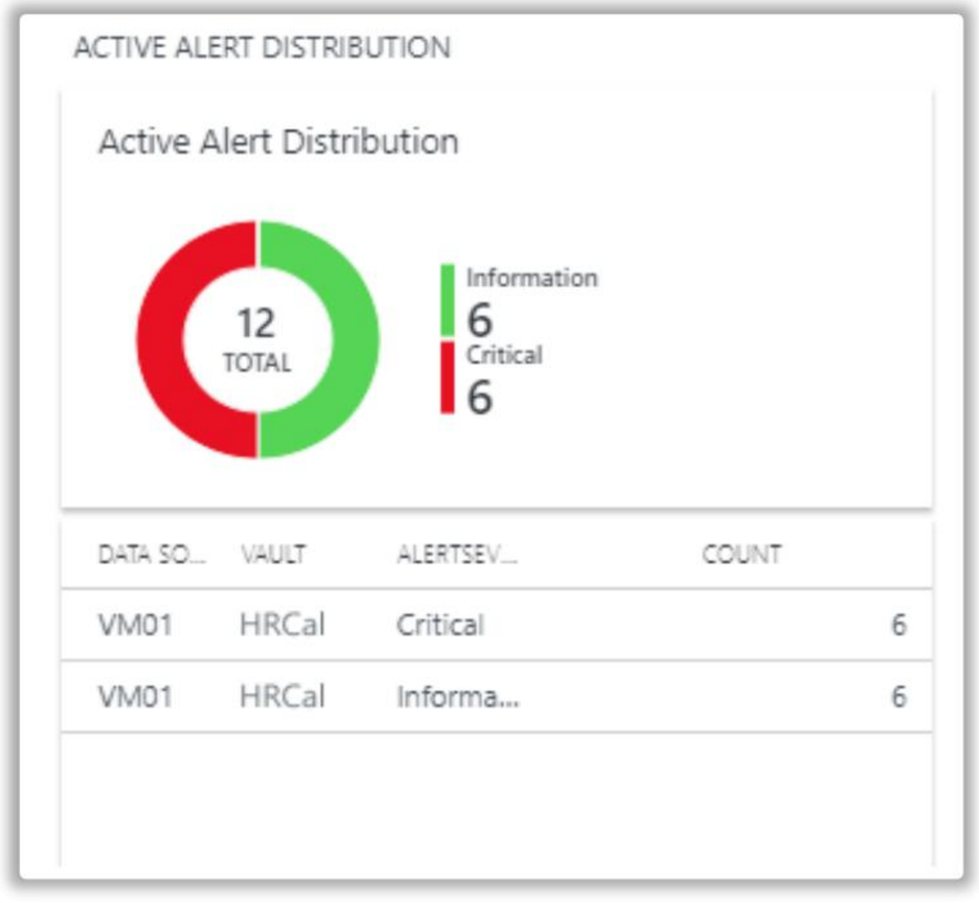

Figure 9.98: Monitoring Azure Backup Alerts

Certified Cloud Security Engineer Copyright © by **EC-Council**
All Rights Reserved. Reproduction is Strictly Prohibited.

- Choose the relevant time range (e.g., start and end dates) to monitor alerts.

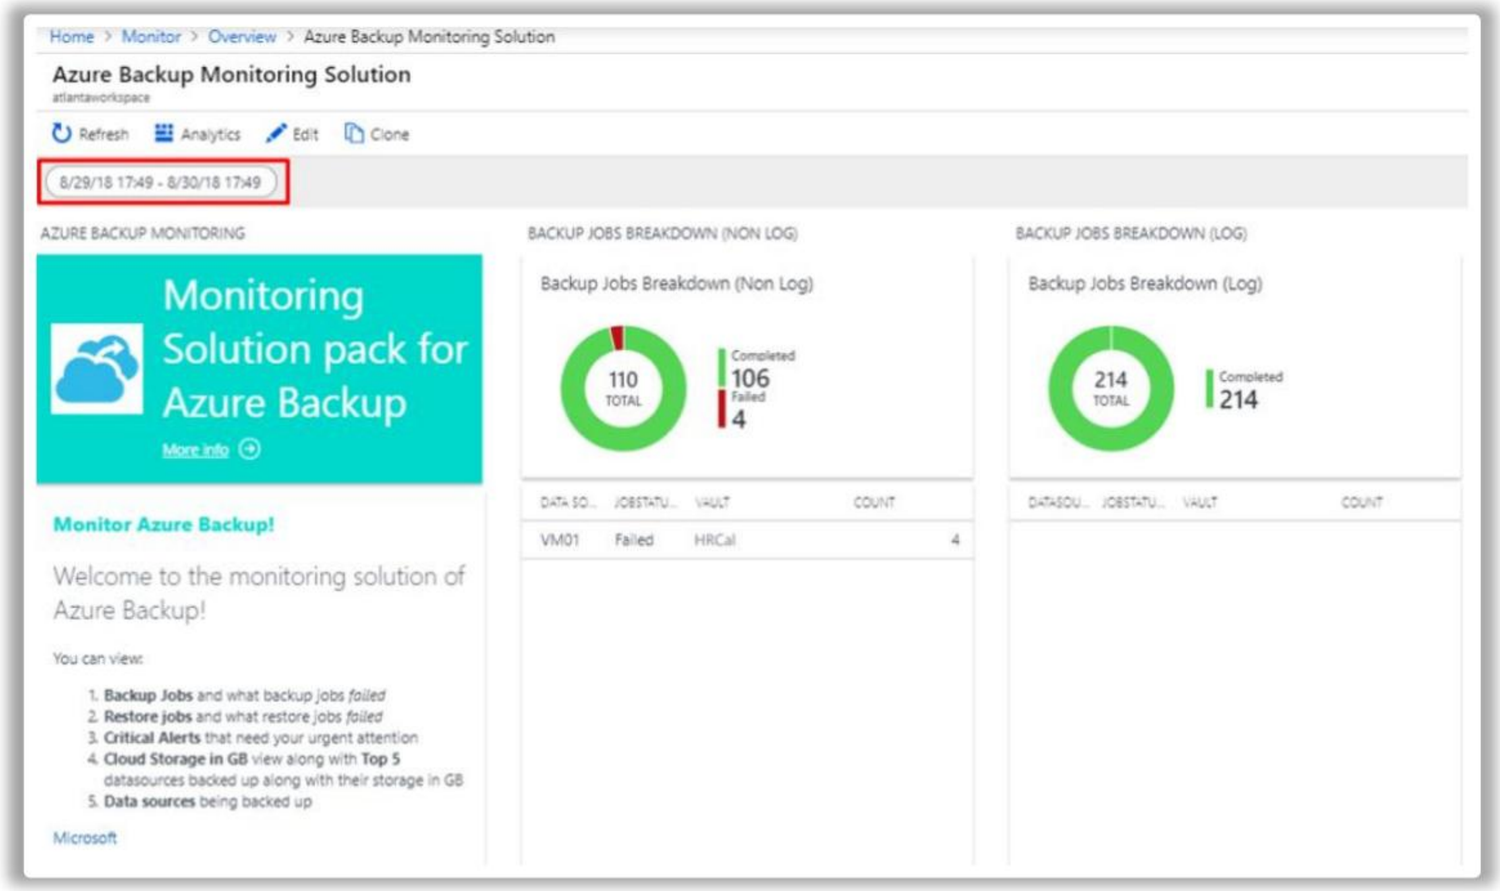

Figure 9.99: Azure Backup Monitoring Solution Console

hide01.ir

Certified Cloud Security Engineer Copyright © by **EC-Council**
All Rights Reserved. Reproduction is Strictly Prohibited.

- **Generate custom alerts:** In the above graphs, click on any row to see a more detailed view in the Log Search window where custom alerts can be generated for that scenario.

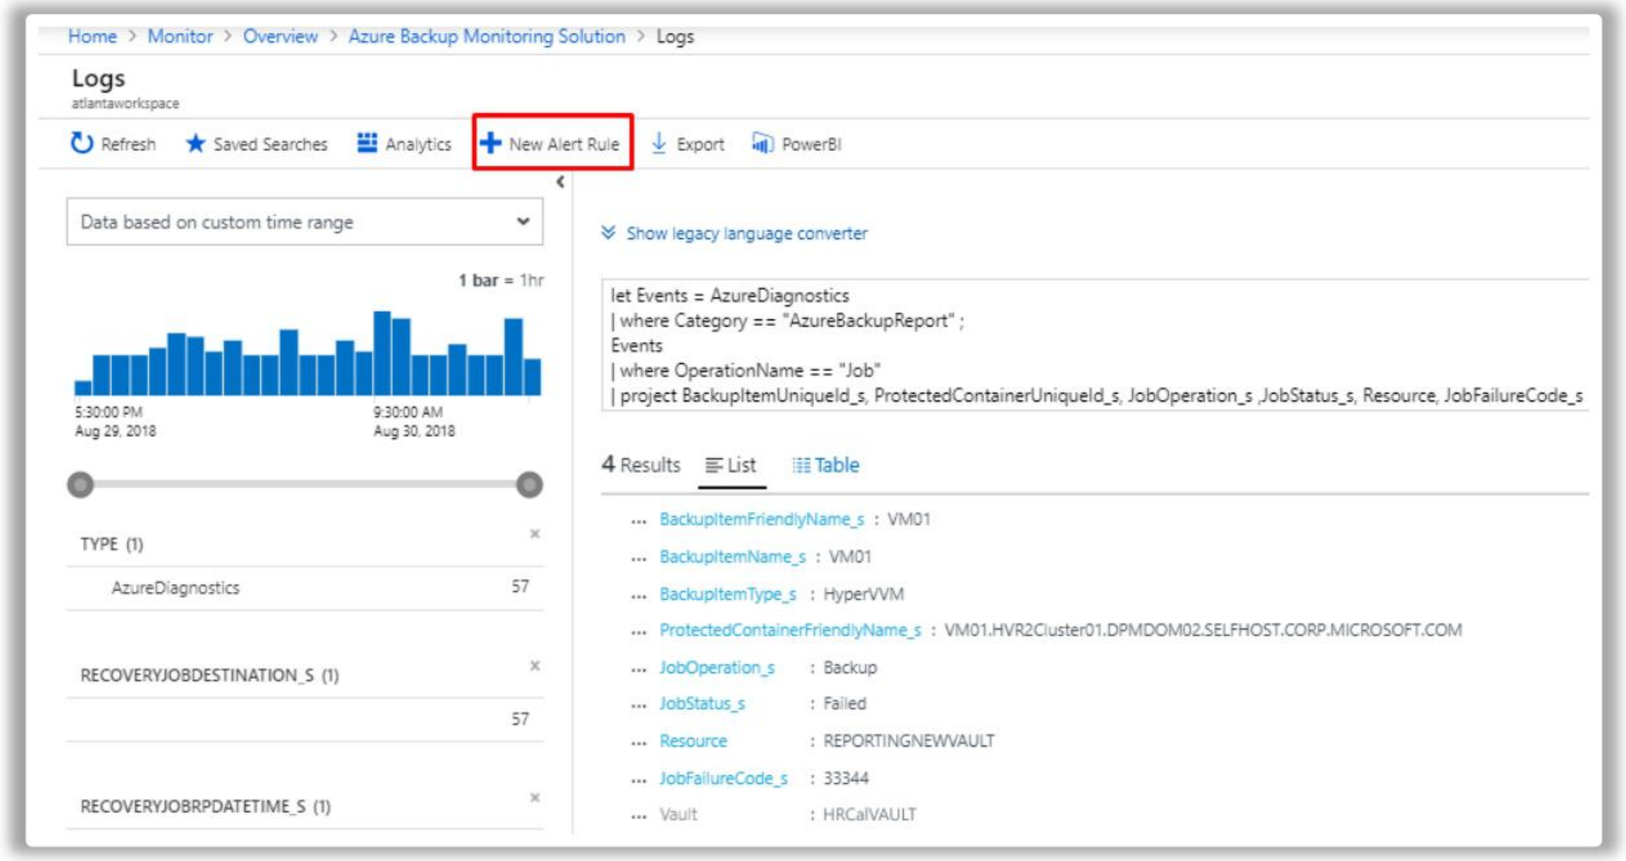

Figure 9.100: Generating Custom Alerts

Azure Backup Reports (Power BI)

The Azure Backup service can be connected to Power BI for reporting. Reports can be viewed, created, and customized directly by going to Power BI after it is integrated with Azure Backup.

hide01.ir

Support Scenarios

- Azure Backup supports Azure file/folder backups and VM backups to the cloud via Azure RSV agent.
- At this time, reports for Azure Backup Server, Azure SQL, and DPM are not supported.
- If the same storage is configured for all vaults, reports across vaults and subscriptions can be viewed. However, the selected storage account must be in the same region as RSV.
- In Power BI, the scheduled refresh occurs at a frequency of 24 hours. Ad-hoc refresh can also be performed in Power BI. In that case, reports are rendered using the latest data in the customer storage account.

Configuring Reporting

Follow the below-mentioned steps to configure reporting:

- Under **Recovery Services Vault**, backup reports are present.

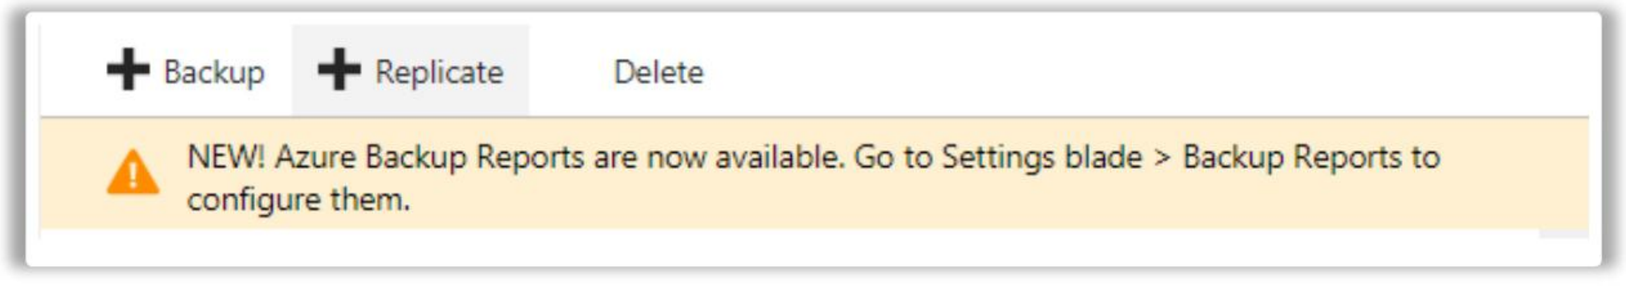

Figure 9.101: Recovery Service Vault Console

Certified Cloud Security Engineer Copyright © by **EC-Council**
All Rights Reserved. Reproduction is Strictly Prohibited.

- Go to **Backup reports** and select **Configure** to enable it.

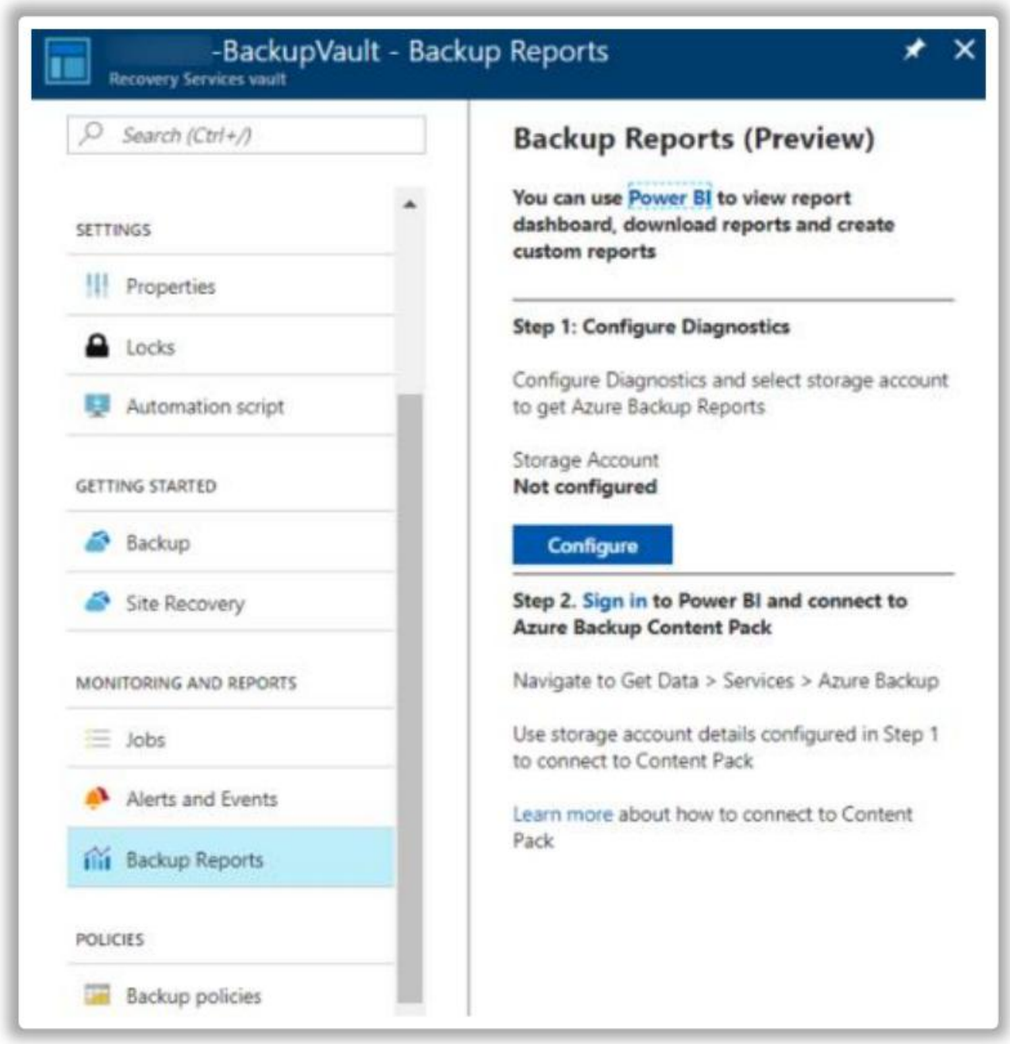

Figure 9.102: Backup Report Console

- In the **Diagnostic setting** page, select the checkbox of **Archive to a storage account** and then give the storage account. Choose the retention period (in days) for reports and click on **OK**.

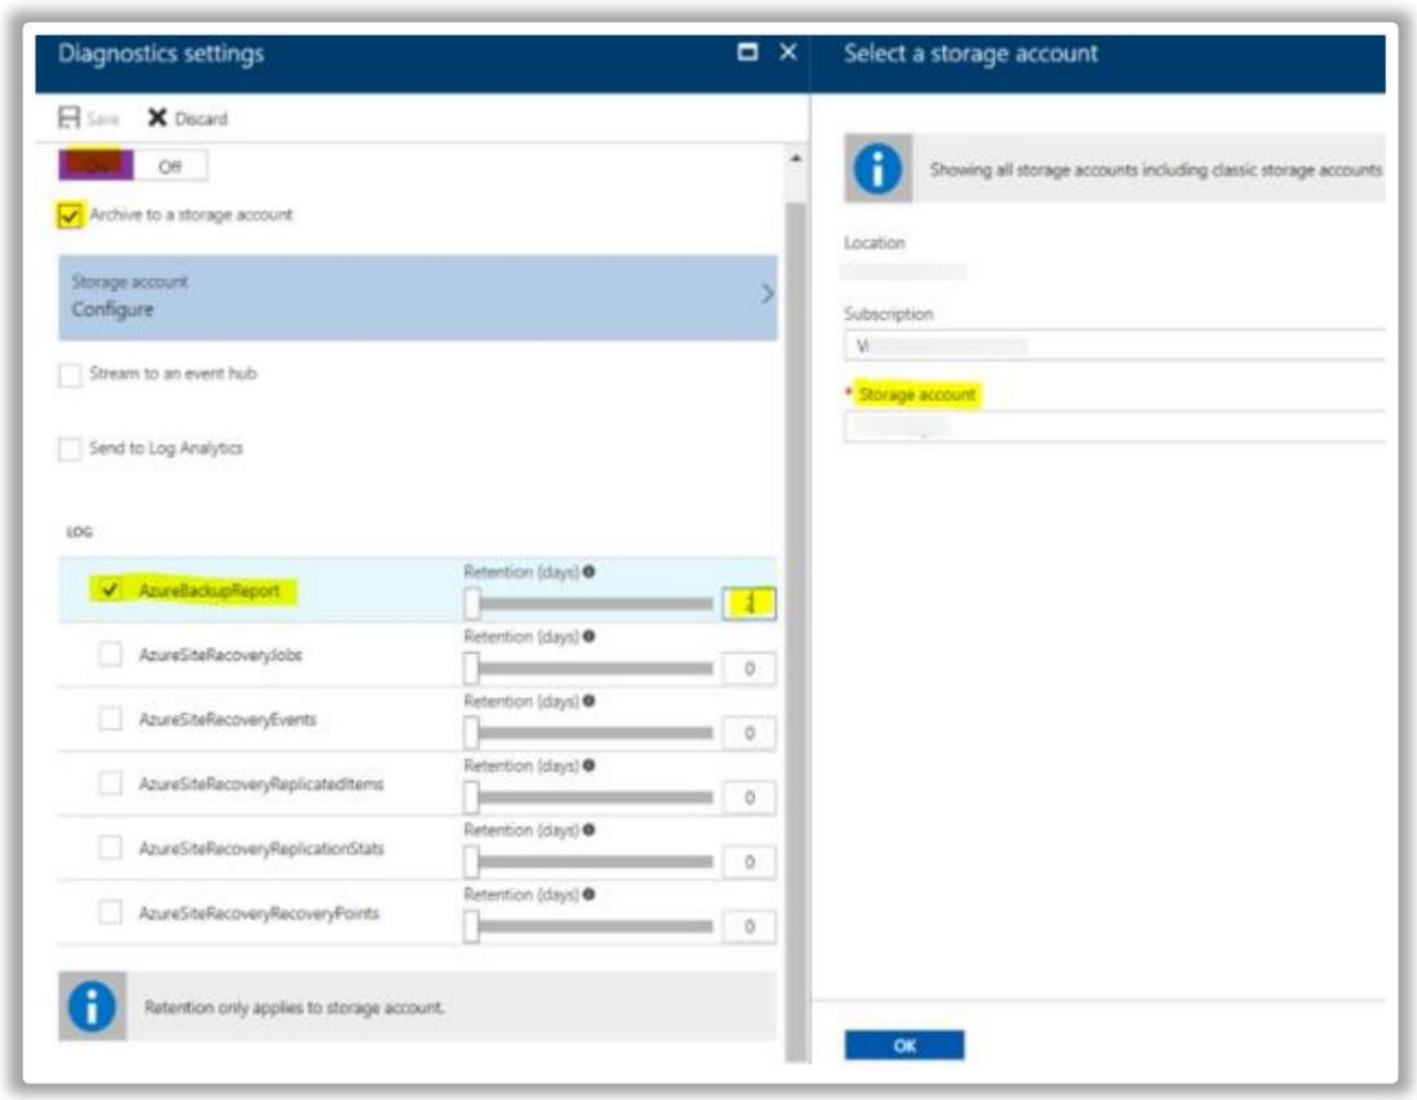

Figure 9.103: Diagnostic Setting Console

Certified Cloud Security Engineer Copyright © by **EC-Council**
All Rights Reserved. Reproduction is Strictly Prohibited.

- Use the organizational account to sign in to Power BI.

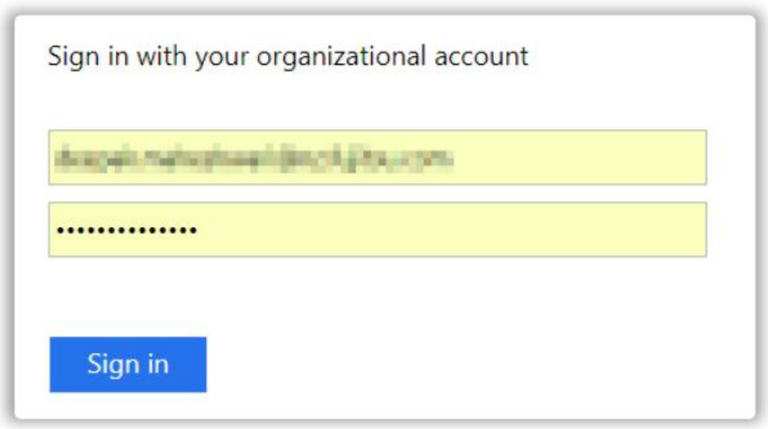

Figure 9.104: Sign in page

- In Power BI, search for "Get Data". Choose Services and then click on Get.

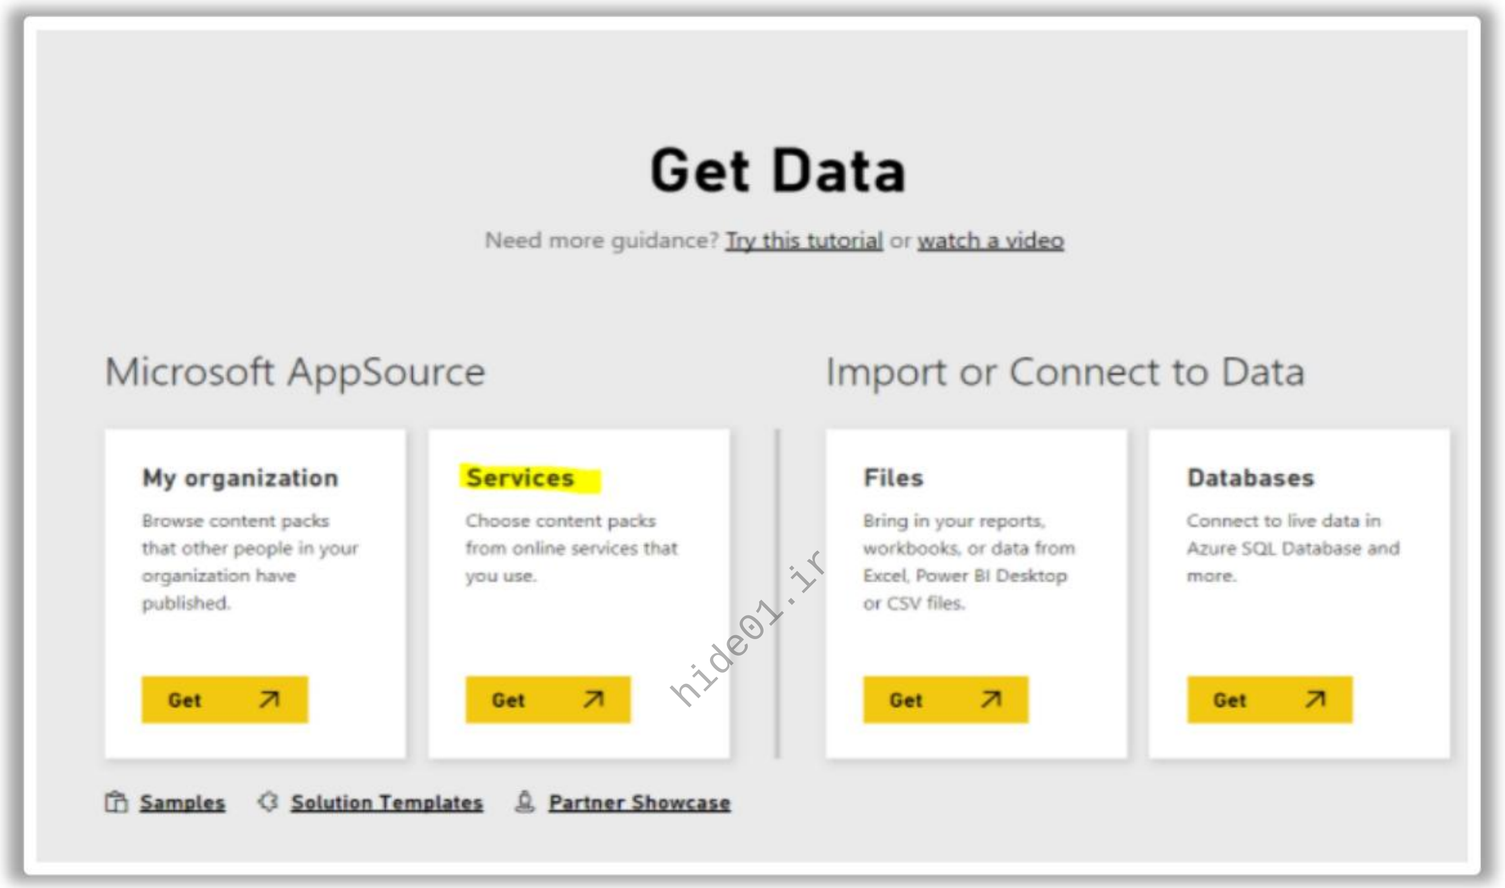

Figure 9.105: Power BI Console

- Now, look for "Azure Backup" and choose **Get it now**.

Figure 9.106: AppSource Console

Certified Cloud Security Engineer Copyright © by **EC-Council**
All Rights Reserved. Reproduction is Strictly Prohibited.

- Enter the name of the storage account.

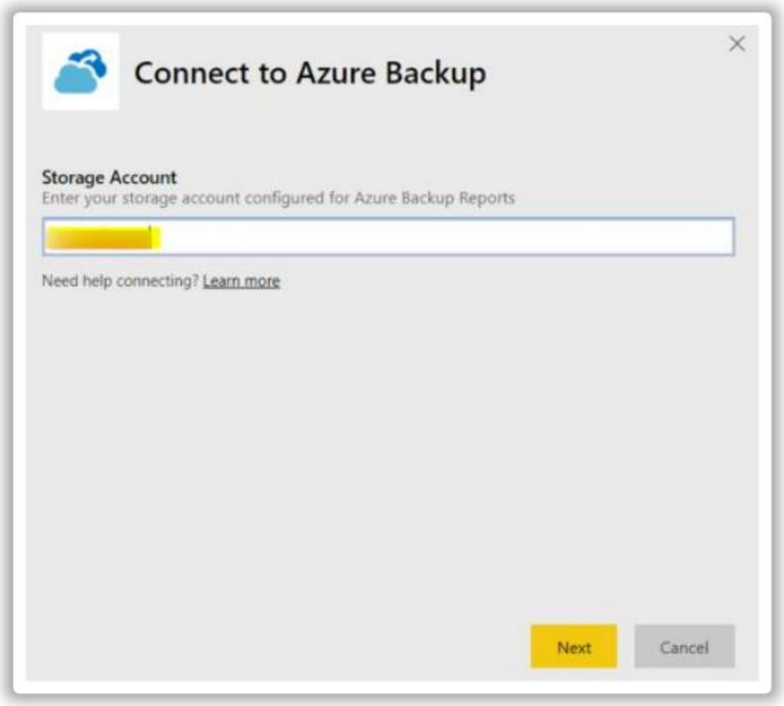

Figure 9.107: Storage Account Console

- From the Azure portal, copy one **Storage Access key**.

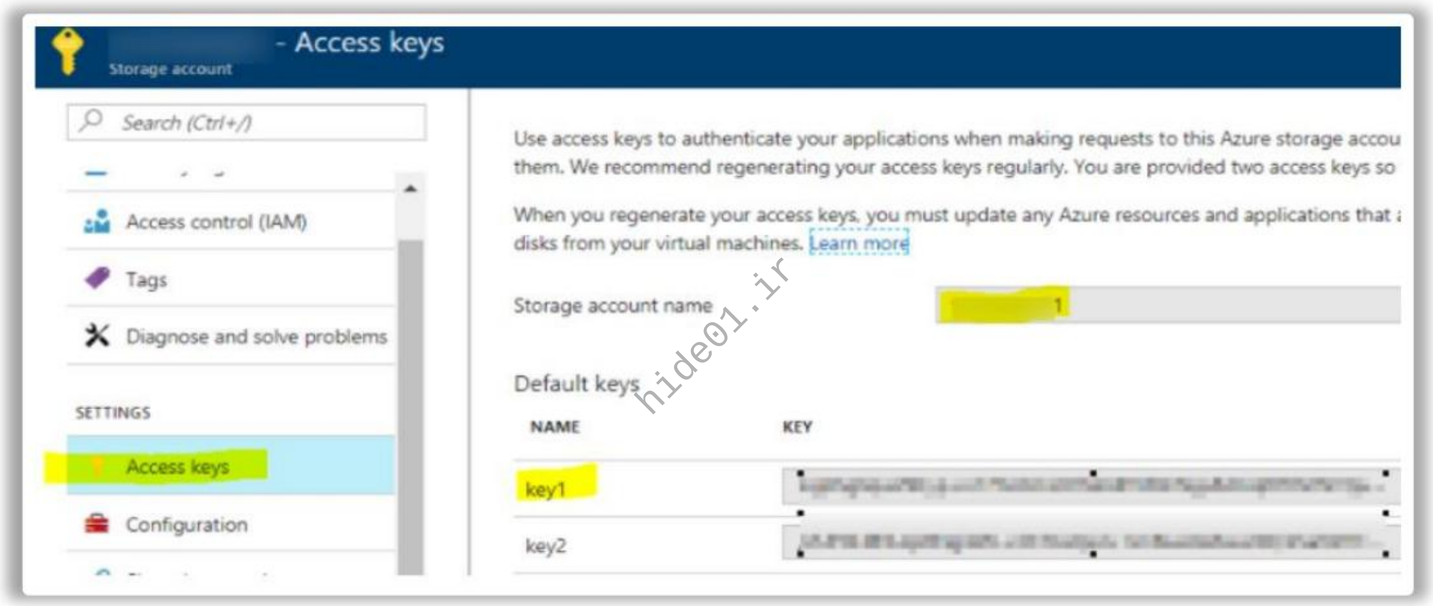

Figure 9.108: Access Keys Console

- To connect to storage, paste the access key.

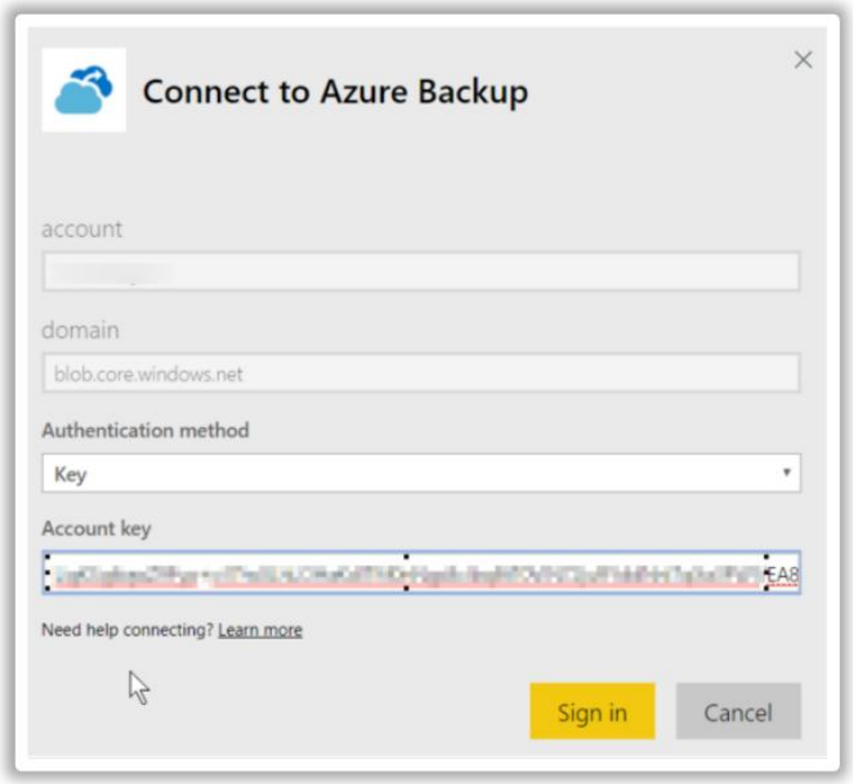

Figure 9.109: Azure Backup Sign in Console

Certified Cloud Security Engineer Copyright © by **EC-Council**
All Rights Reserved. Reproduction is Strictly Prohibited.

- Once the connection is successful, Power BI displays Azure Backup.

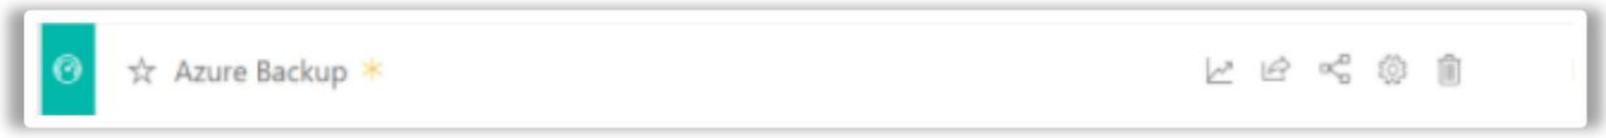

Figure 9.110: Azure Backup Panel

- Now, Power BI shows all backup jobs and other reports.

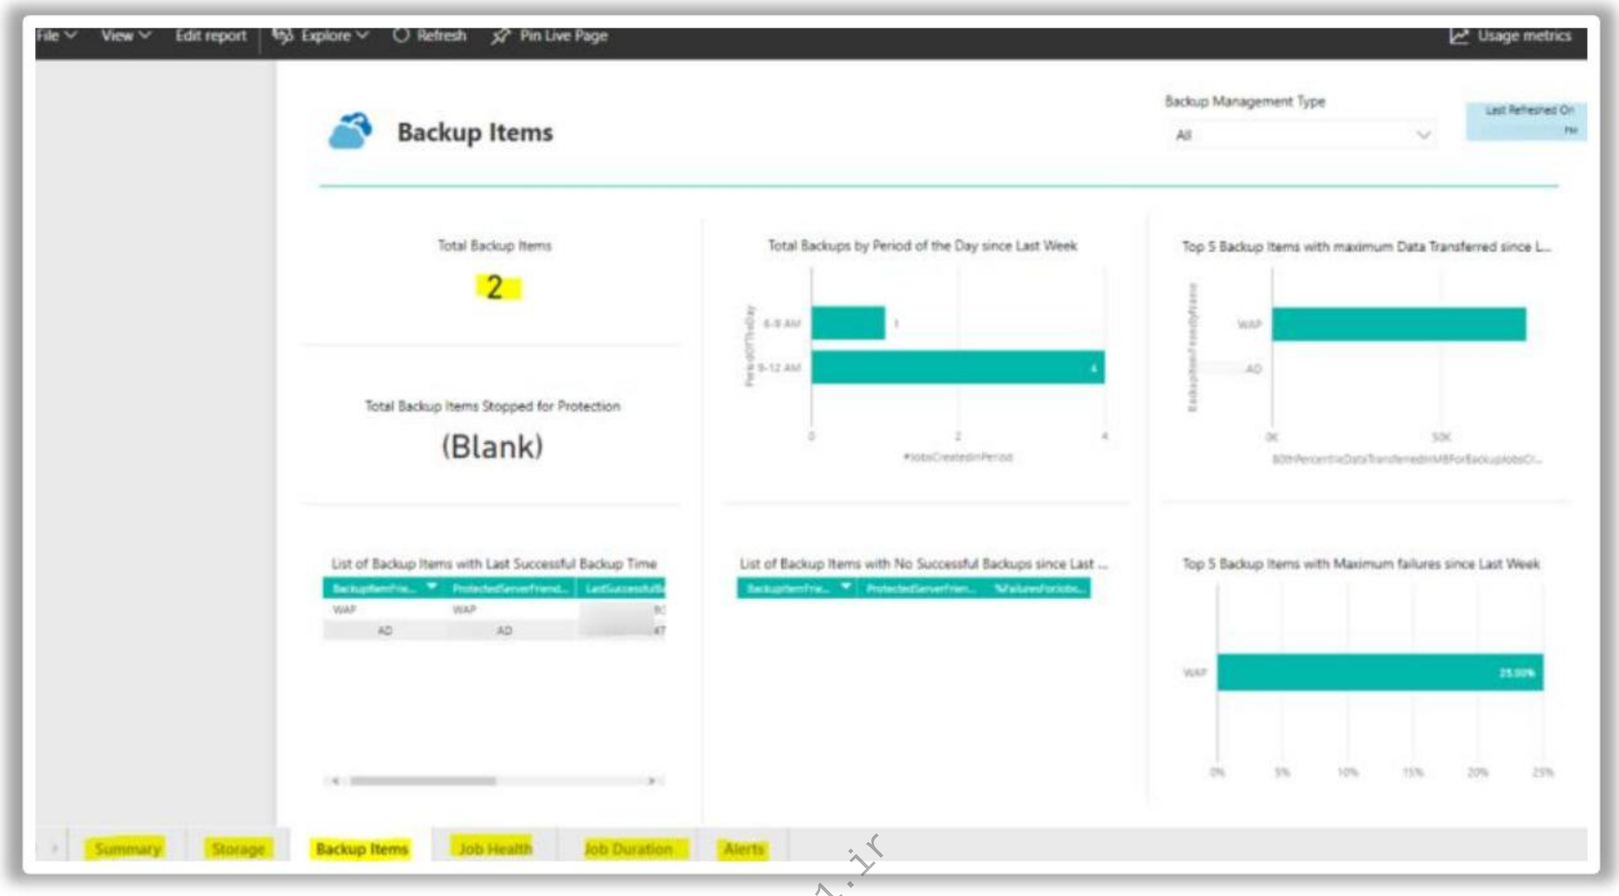

Figure 9.111: Backup Items Console

hade01.ir

Certified Cloud Security Engineer Copyright © by **EC-Council**
All Rights Reserved. Reproduction is Strictly Prohibited.

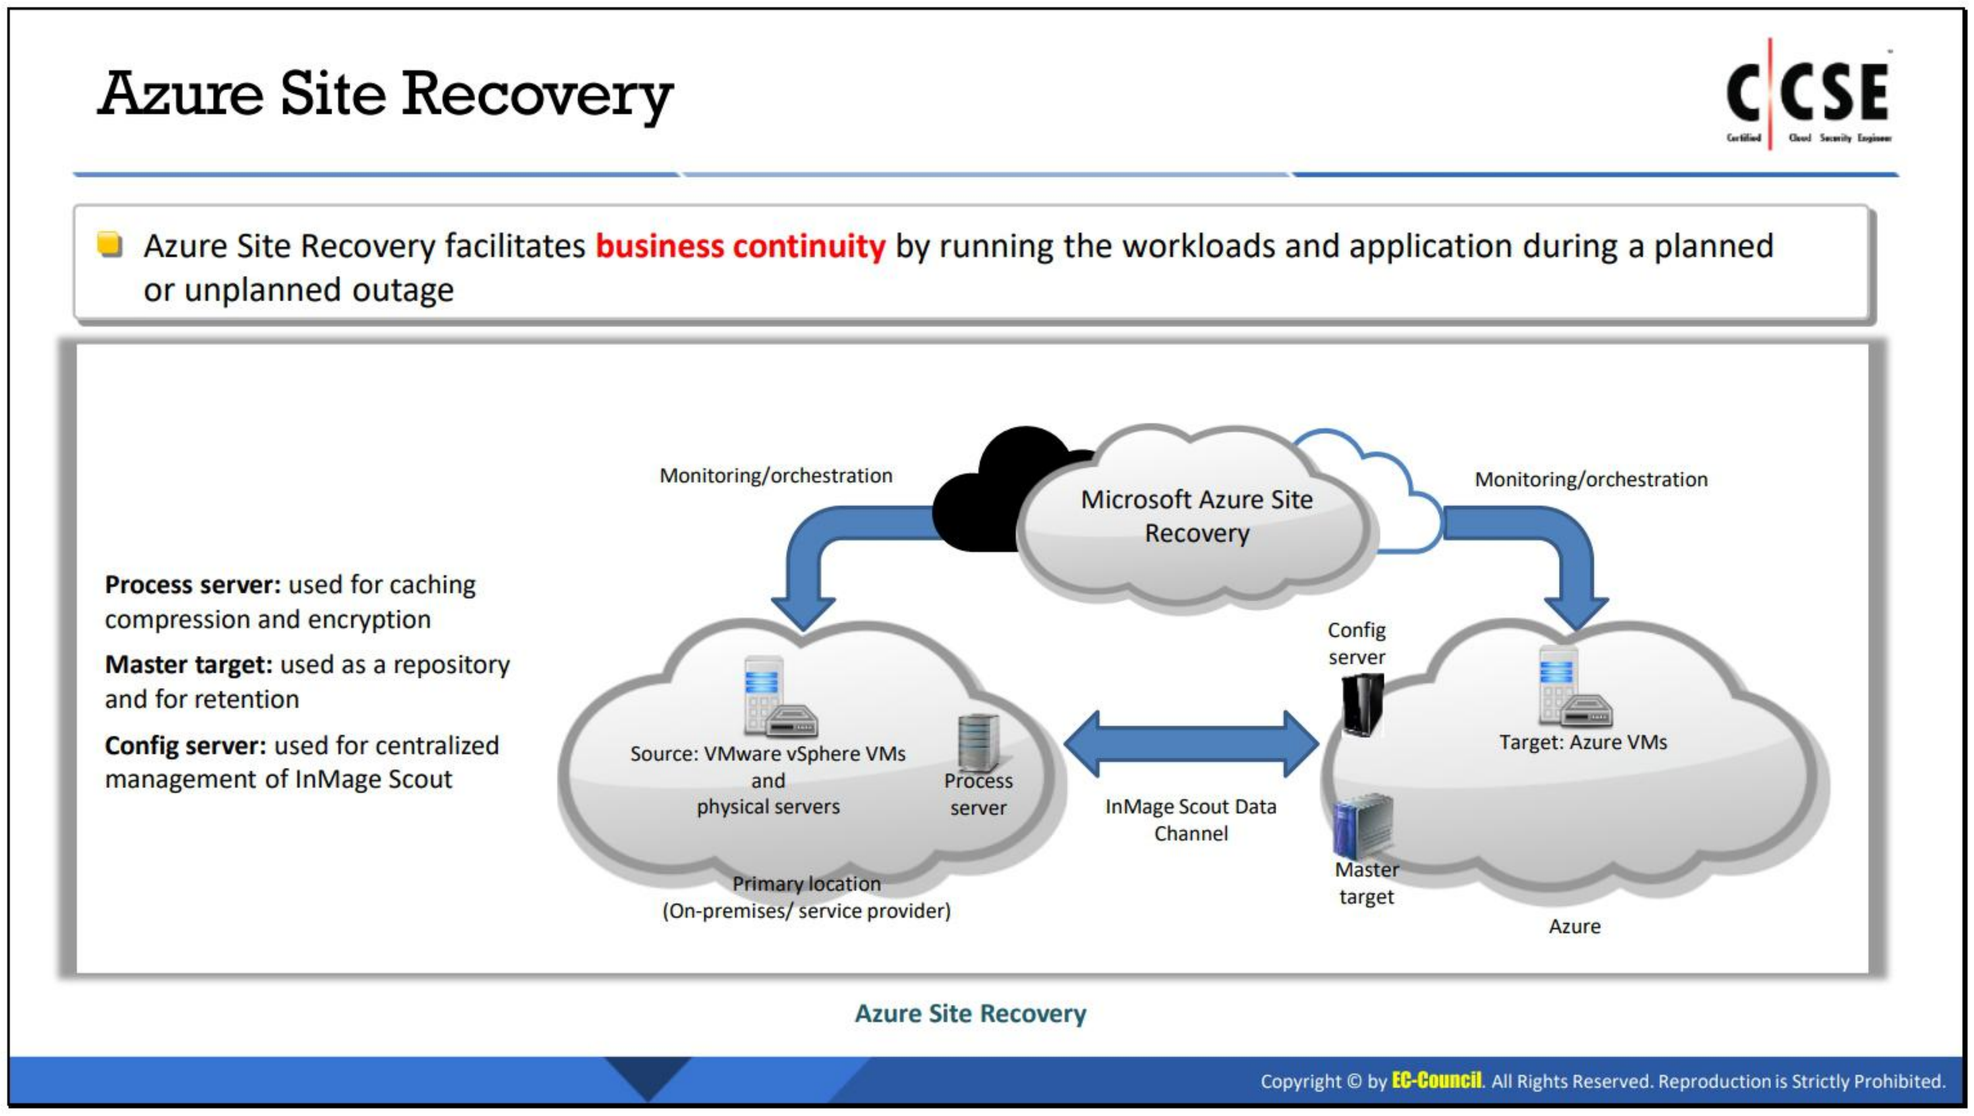

Azure Site Recovery

For the safety of data, any organization must adopt a BCDR plan. This helps in proving your applications and services during any planned or unplanned outages.

Azure Recovery Services also includes the following in your BCDR plan:

- **Site Recovery service:** The ASR service ensures that all the business applications and data can be used even during disasters. This helps in BC. The main function of Site Recovery is to replicate physical and VMs workloads from primary site to secondary site. Failover to secondary location helps applications access if there is an outage in the primary site. Once the primary location is ready to go live, then you can again fail back to it.
- **Backup service:** The Azure Backup service is used to maintain data safe and recover when required.

Site Recovery allows replication management of the following:

- Replication of Azure VMs between Azure regions.
- Azure Stack VMs, On-premises VMs, and physical servers.

Certified Cloud Security Engineer Copyright © by **EC-Council**
All Rights Reserved. Reproduction is Strictly Prohibited.

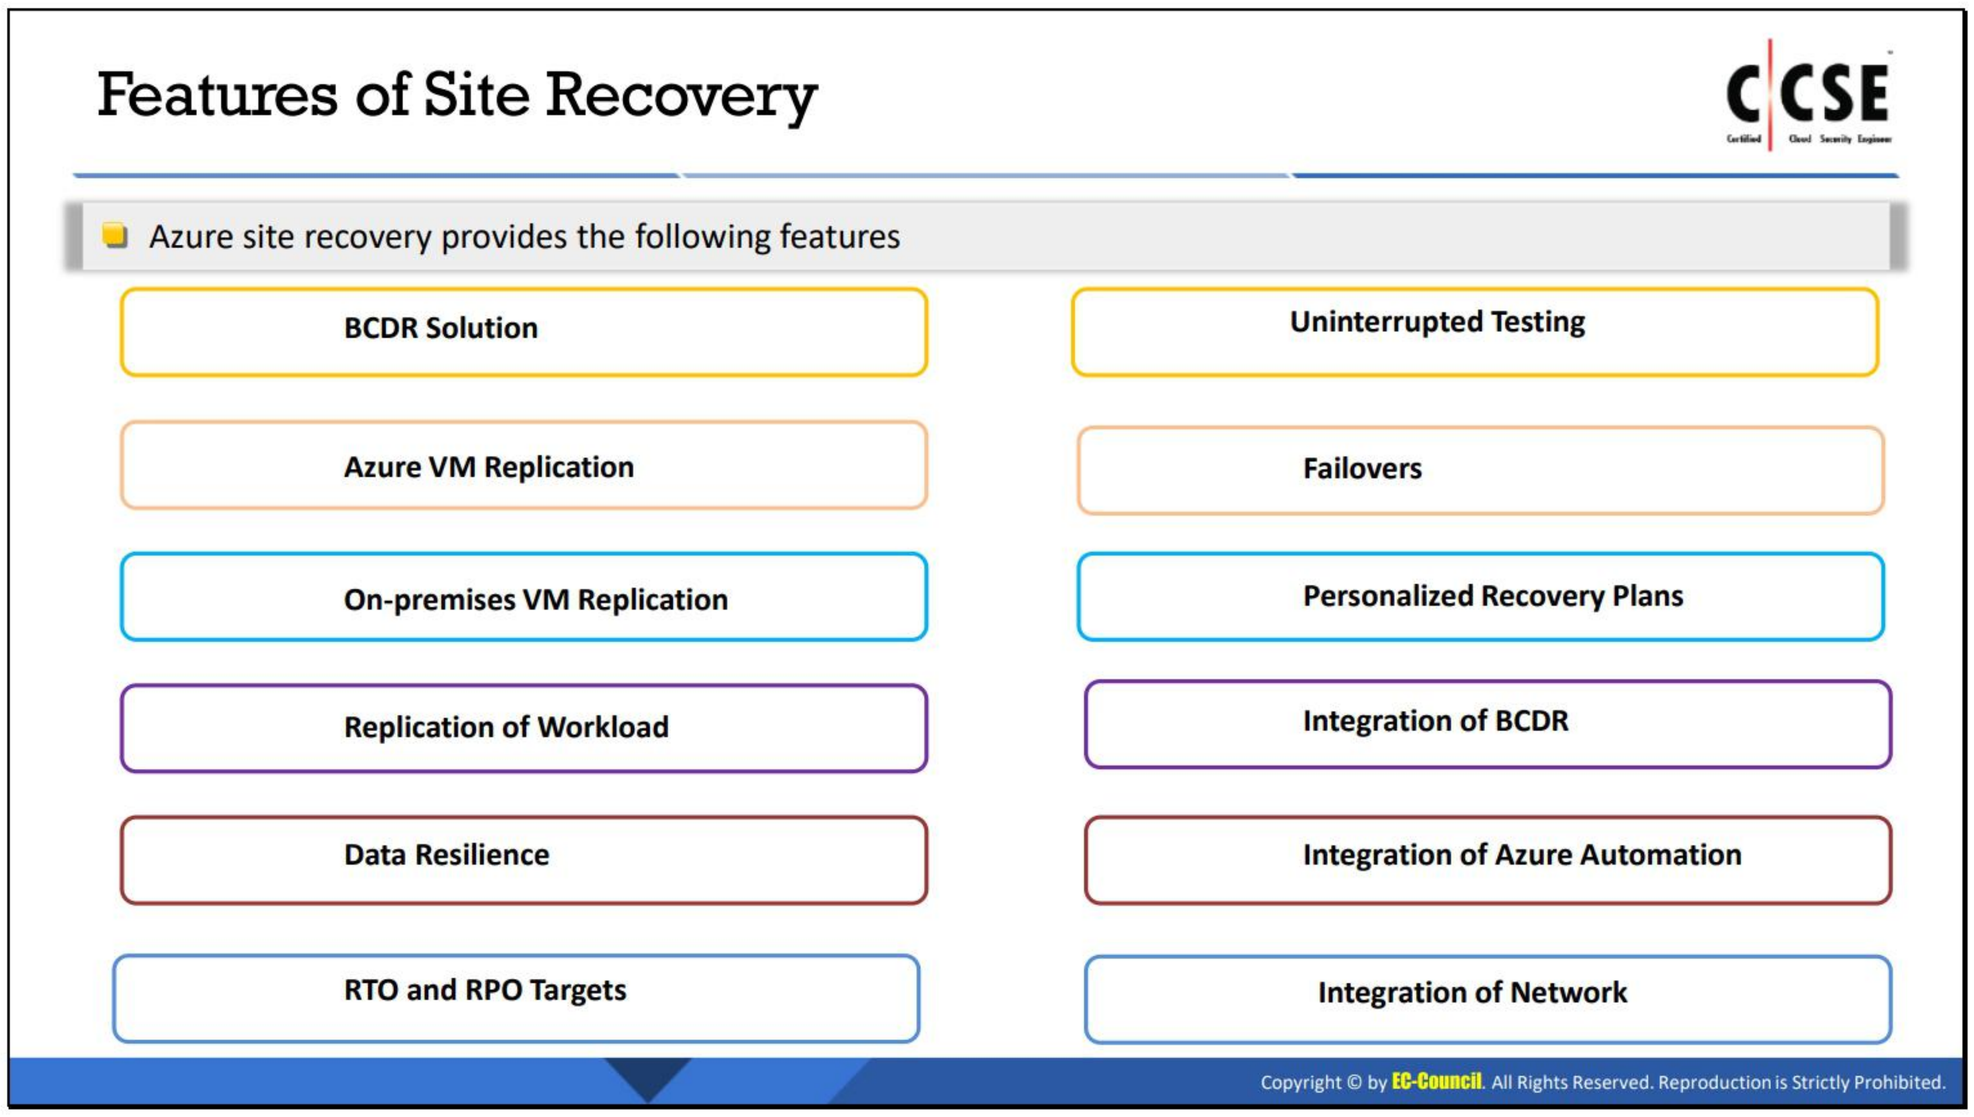

Features of Site Recovery

ASR provides the following features:

- **BCDR solution:** Utilizing ASR, the Azure customer can setup and manage failback, replication, failover from a single location in the Azure portal.
- **Azure VM replication**: Disaster recovery of Azure VMs can be set up from a primary region to a secondary region.
- **On-premises VM replication**: The VMs and physical servers can be replicated to Azure or any secondary on-premises datacenter. To avoid the cost and maintenance complexity, it is better to replicate to Azure than opting secondary datacenter.
- **Replication of workload**: Any workload that is running on the supported Azure VMs, Hyper-V, on-premises, VMware VMs, and Linux/Windows servers can be replicated.
- **Data resilience**: Replication without intercepting application data is possible using Site Recovery. If the data is replicated to Azure, the data is stored in Azure Storage with the same security provisions. In case of failover, Azure VMs are created based on replicated data.
- **RTO and RPO targets**: RTO and RPO should be set within the limits of the organization. Continuous replication for VMware VMs and Azure VMs is provided by Site Recovery. The Hyper-V replication frequency is as low as 30 seconds. RTO can be further reduced with the integration of the Azure Traffic Manager.
- **Uninterrupted testing**: Disaster drills can be performed easily where it does not affect ongoing replication.

Certified Cloud Security Engineer Copyright © by **EC-Council**
All Rights Reserved. Reproduction is Strictly Prohibited.

- **Failovers**: With zero-data loss planned failovers can be run for expected outages. At the same time, expect minimal data loss for unplanned failovers in unexpected disasters, this further depends on the frequency of replication. Failback is possible immediately after the availability of the primary location.
- **Personalized recovery plans**: Recovery plans allow you to customize and sequence the failover. It also allows you to recover multi-tier applications that are running on multiple VMs. In a recovery plan, you can group machines and optionally add scripts and manual actions. Azure Automation runbooks and recovery plans integration could be done.
- **Integration of BCDR**: Other BCDR technologies integrations are done with Site Recovery. For example, Site Recovery can be used for the protection of corporate workloads SQL server backend, with always-on native support for SQL server, to manage availability groups failover.
- **Integration of Azure automation**: For the facility of downloading and integration with Site Recovery, the Azure Automation library provides application-specific scripts and production-ready phase.
- **Integration of network**: For application network management Site Recovery integrates with Azure. For example, Use Azure Traffic Manager for effective network switchovers of IP address reservation and load-balancer configuration.

hide01.ir

Certified Cloud Security Engineer Copyright © by **EC-Council**
All Rights Reserved. Reproduction is Strictly Prohibited.

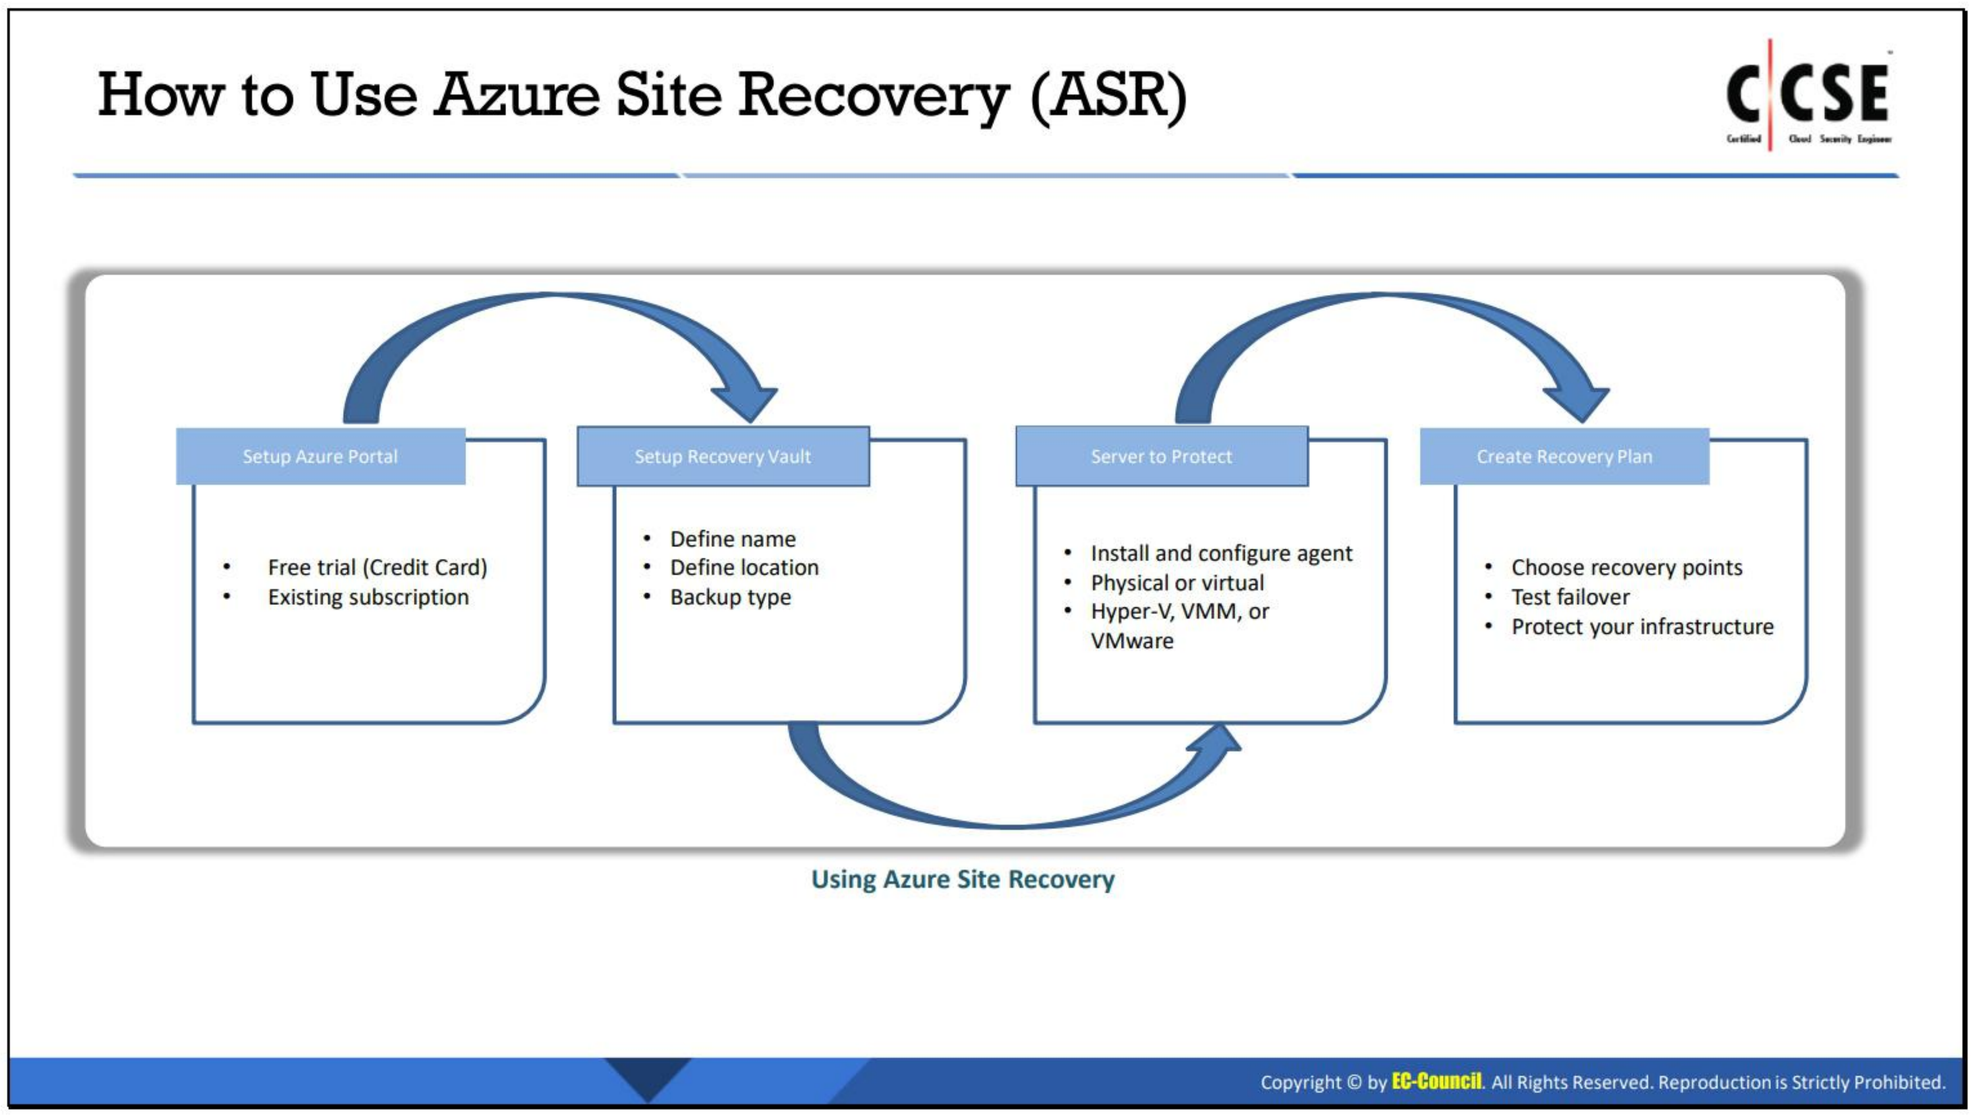

How to Use Azure Site Recovery (ASR)

ASR is a part of Azure Subscription as well the Operations Management Suite (OMS). Both provide free trials, which the cloud consumer can sign up to use.

The setup of Azure is simple and easily understandable. If you do not have a subscription, then you need to set up an Azure Subscription. Next, sign in to the Azure portal. The next step is provisioning RSVs. RSVs can hold the backup data for Azure services, workstations, and on-site servers. They help organize backup data easily and minimize management overheads. In the recovery vault, you need to define the name, location, type of backup, etc.

In the next step, install and configure the agent on the server, then identify the servers that you want to protect and configure them. Finally, create a recovery plan that includes testing failover, selecting recovery points, protecting your infrastructure, etc.

The user can utilize the recovery plan for failback from Azure and failover to Azure. Ensure that the source location has systems that are enabled for failover and recovery. The machines that are used in the recovery plan should have similar source and target. In the same plan, the user can incorporate Hyper-V VMs and VMware VMs managed by VMM. The physical servers and VMware VMs can be included in the same plan.

Certified Cloud Security Engineer Copyright © by **EC-Council**
All Rights Reserved. Reproduction is Strictly Prohibited.

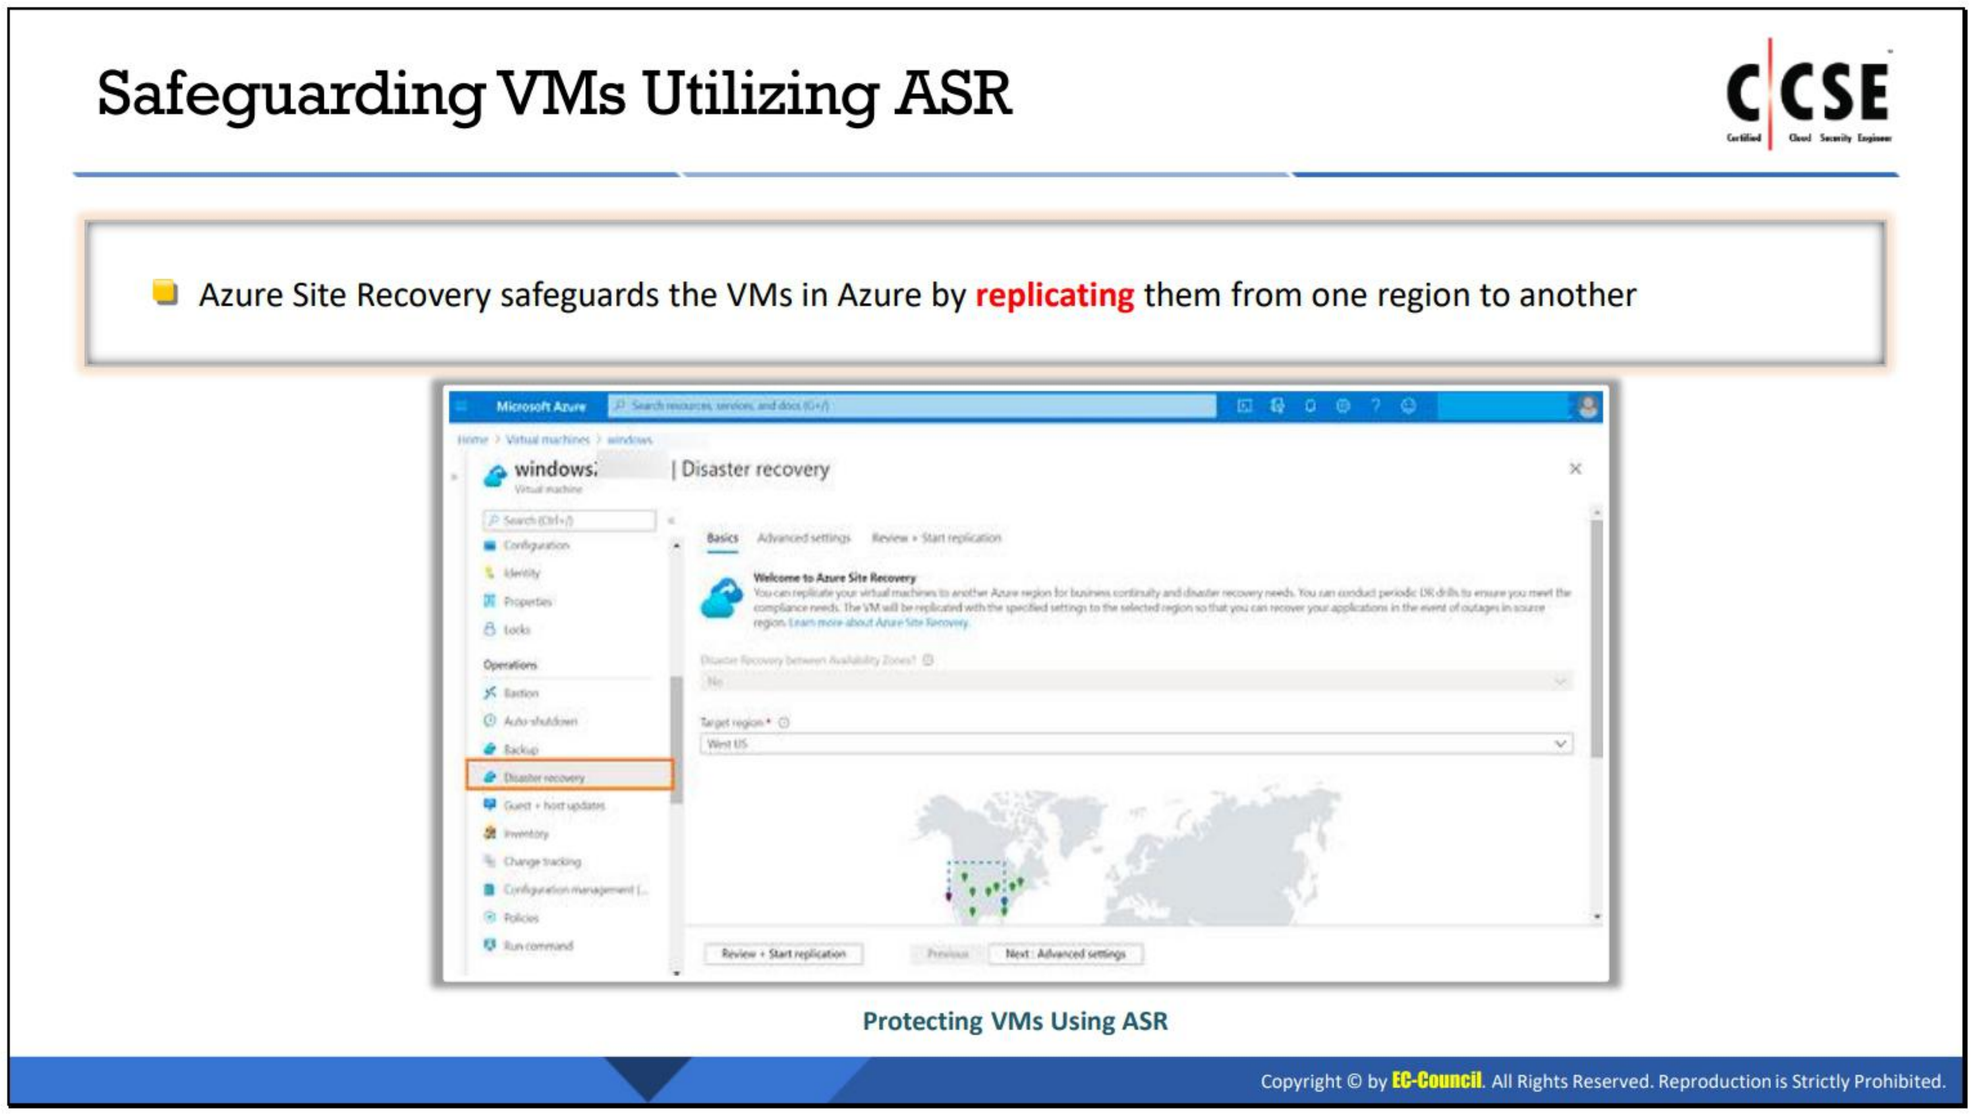

Safeguarding VMs Utilizing ASR

In Azure, the ASR is utilized to protect the VMs by replicating them from one region to another.

To safeguard VMs using ASR, perform the following steps:

- Log in to the **Azure portal**, navigate to **Virtual machines**, and under **Operations**, click on **Disaster recovery**.

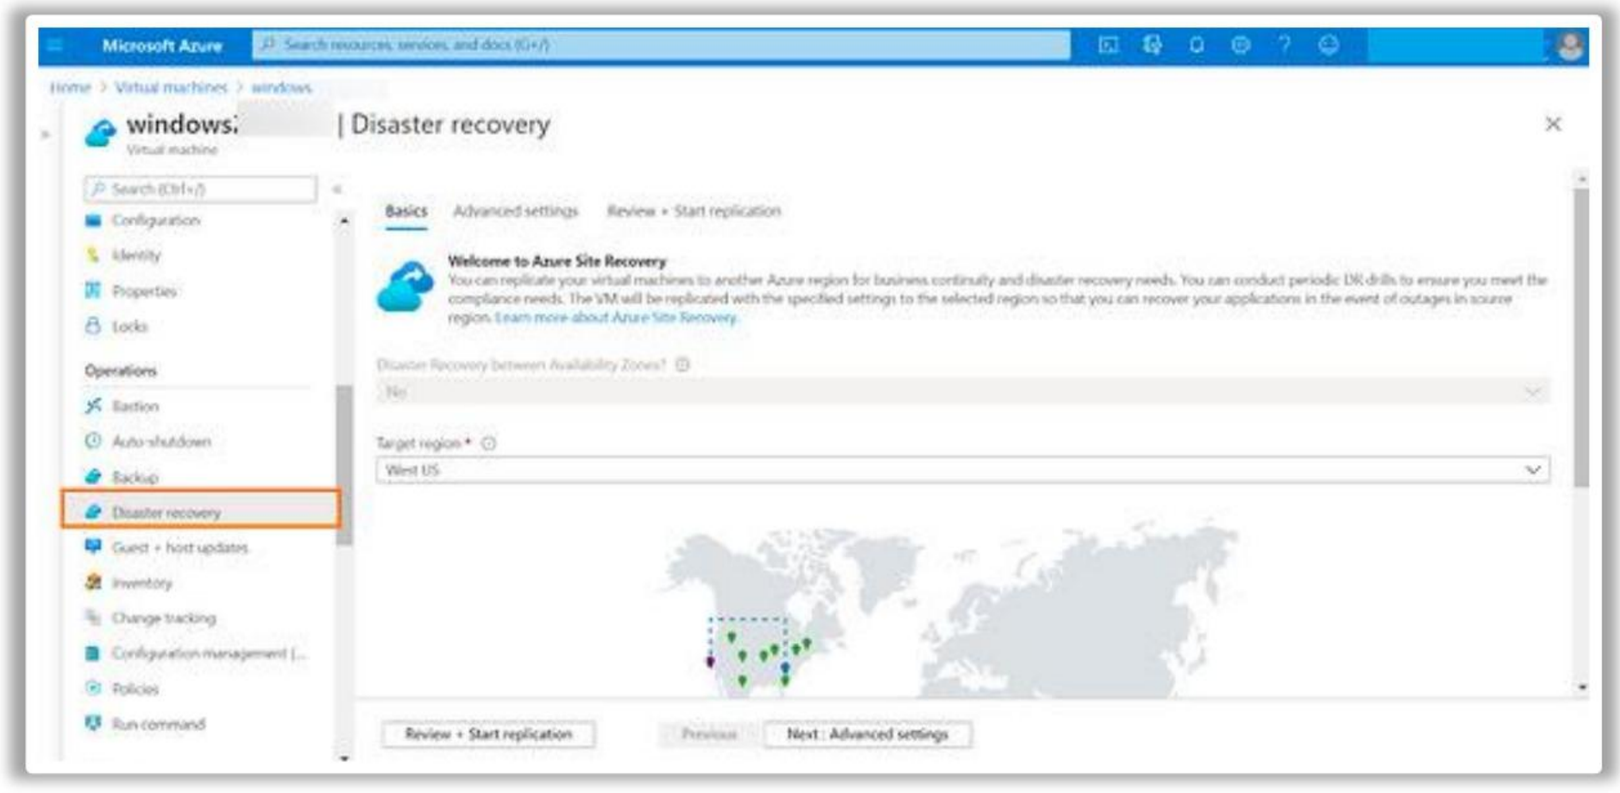

Figure 9.112: Disaster Recovery Console

Certified Cloud Security Engineer Copyright © by **EC-Council**
All Rights Reserved. Reproduction is Strictly Prohibited.

- From the geographical cluster, choose the target region, then click on **Next: Advanced Settings**.

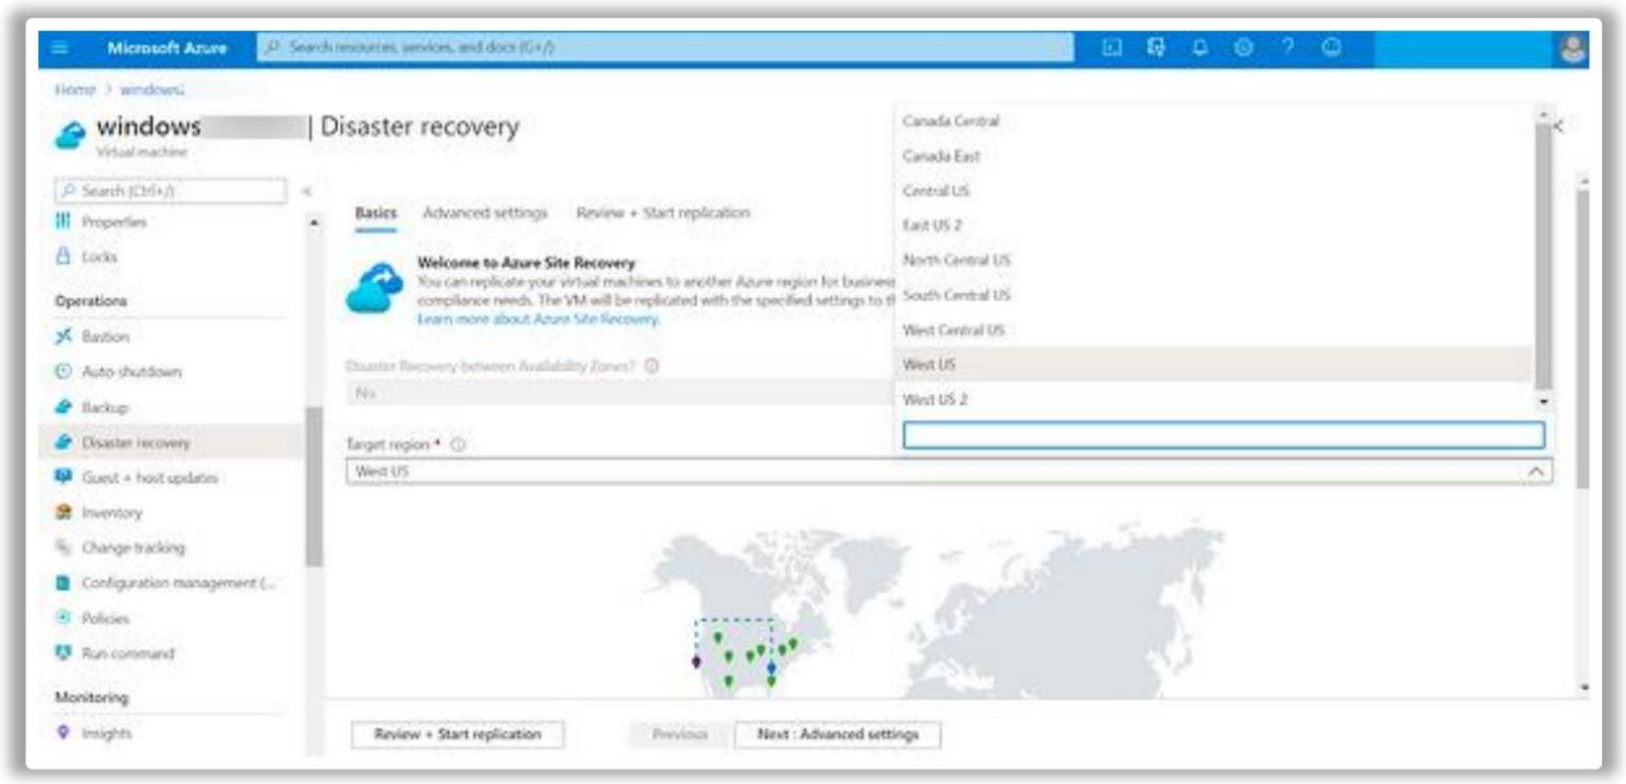

Figure 9.113: Creating Advanced Settings

- Set up the target environment by select the target for Subscription, VM resource group, virtual network, Availability. You can select the existing resource groups, virtual network, or create a new one.

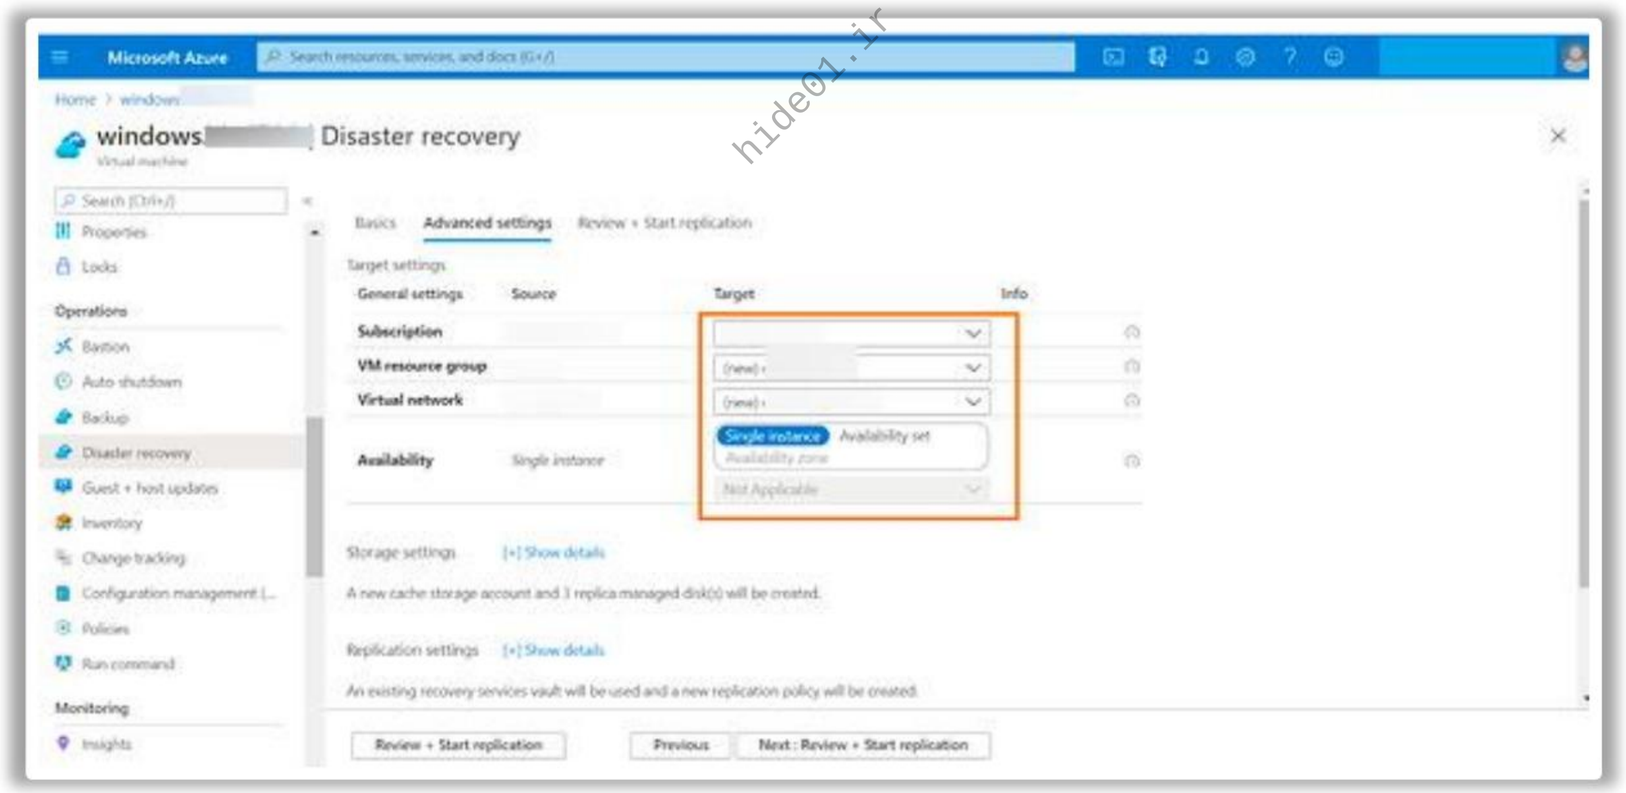

Figure 9.114: Setting Target Environment

Certified Cloud Security Engineer Copyright © by **EC-Council**
All Rights Reserved. Reproduction is Strictly Prohibited.

- Before replicating the target, choose the cache storage that can be utilized for temporary storing data in the source region. You may also choose the disks that need to be replicated and the target disk type (Premium SSD or Standard SSD/HDD). You can choose the name of the vault and the replication policy that can be utilized for replication. Then, click on the **Review + Start replication** button.

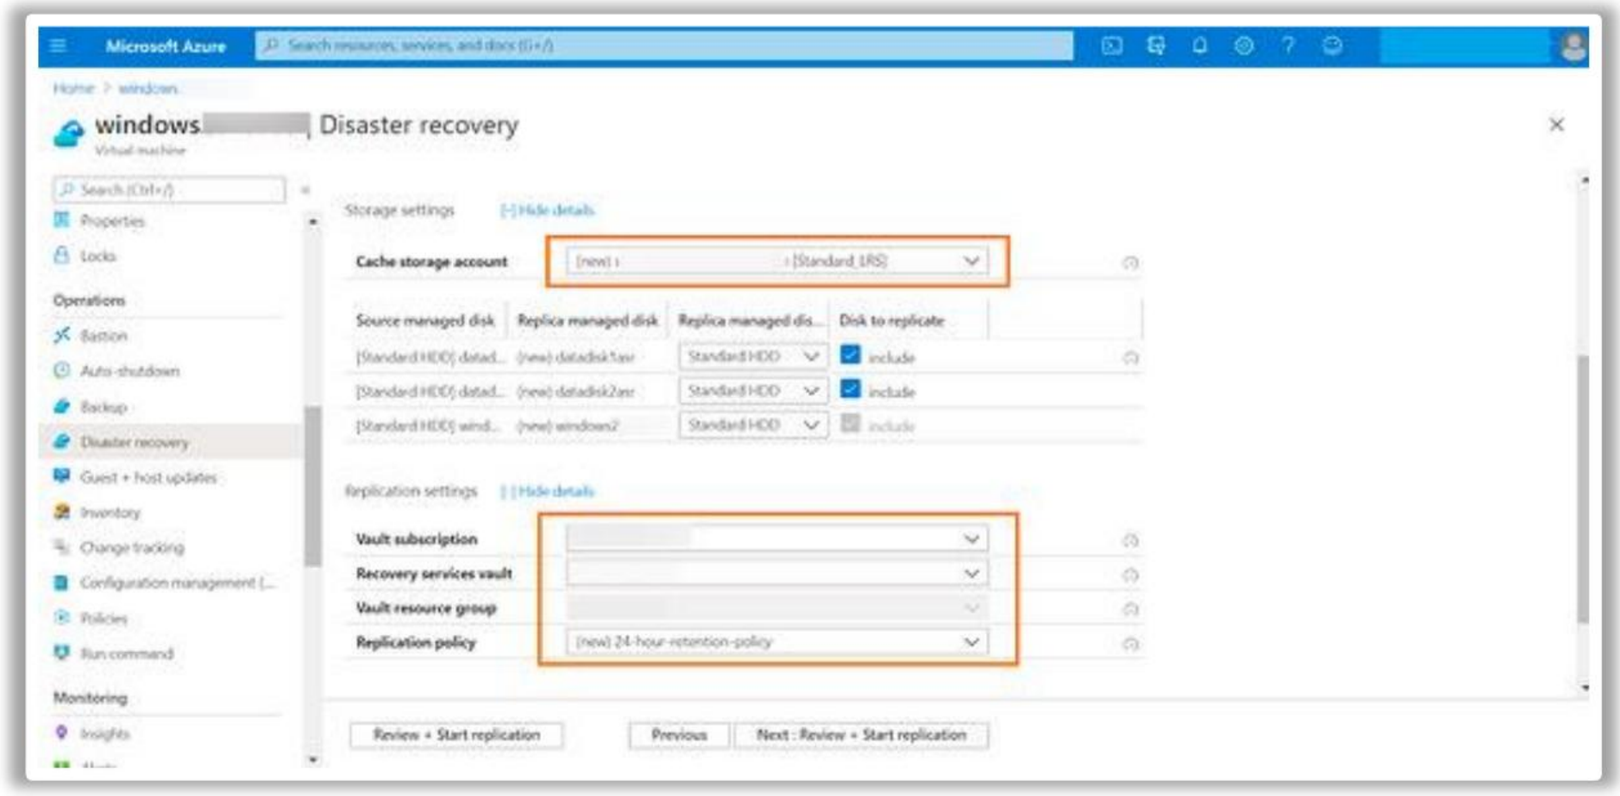

Figure 9.115: Review Settings

- On the next page, click on **Start replication**

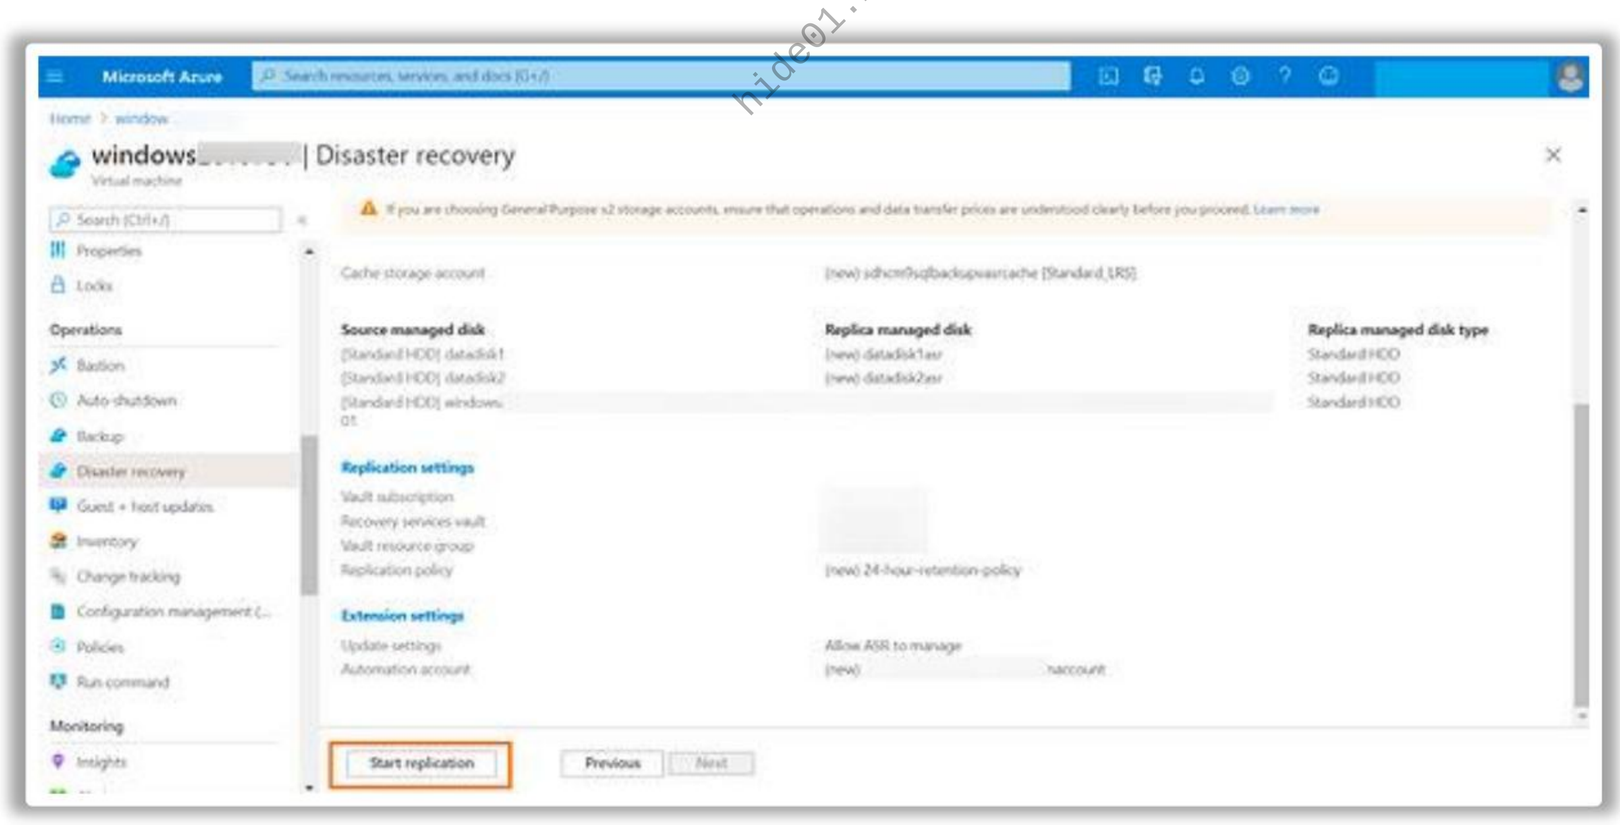

Figure 9.116: Start Replication

Certified Cloud Security Engineer Copyright © by **EC-Council**
All Rights Reserved. Reproduction is Strictly Prohibited.

- After completion of initial replication, you can check the protection status by navigating to **Virtual machine**, **Operations**, and **Disaster Recovery**.

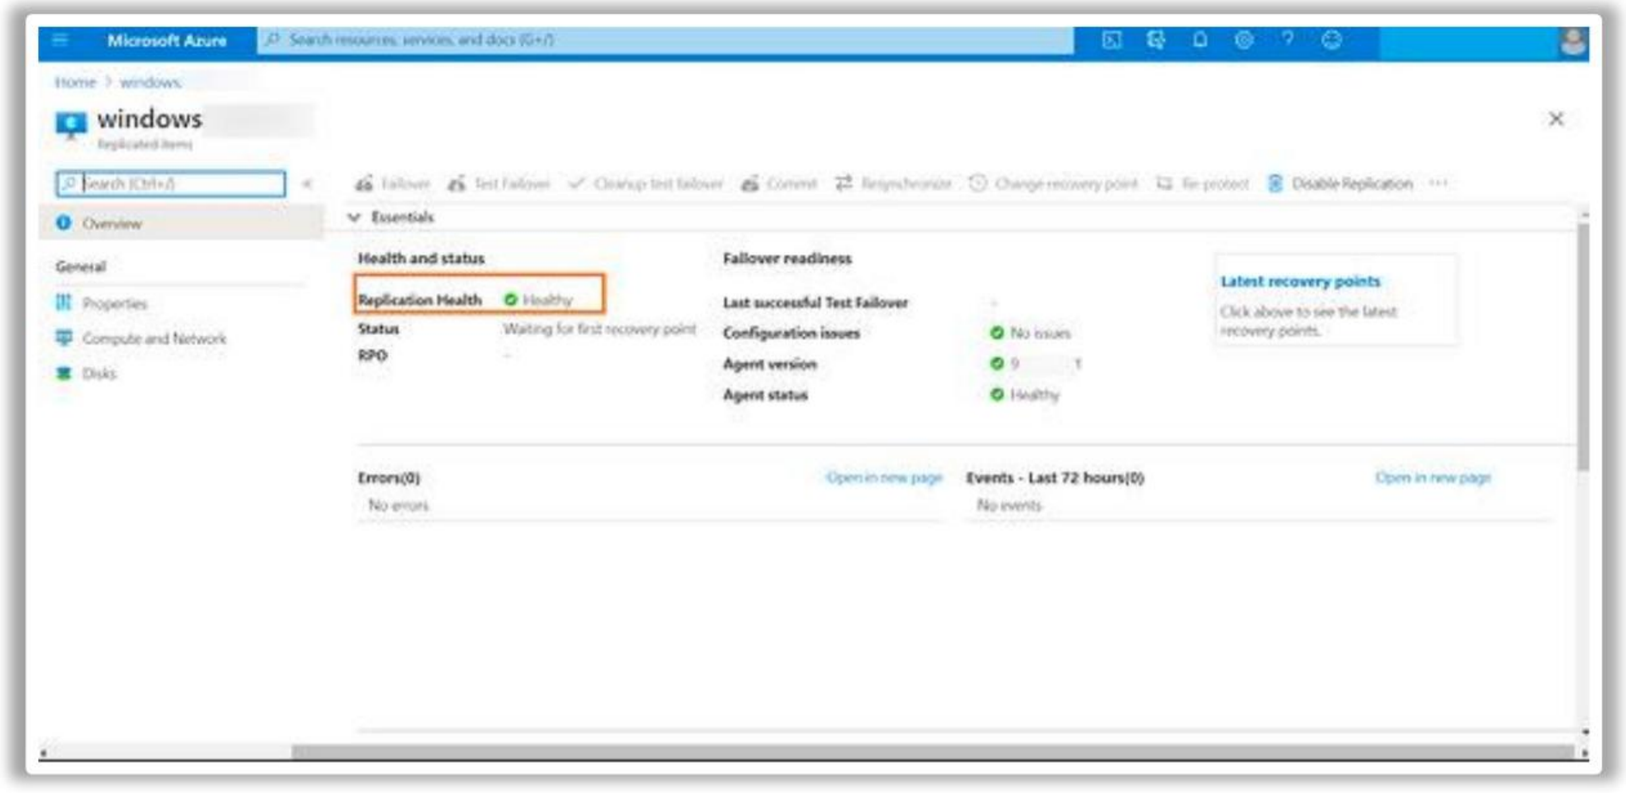

Figure 9.117: Checking Replication Health

hide01.ir

Certified Cloud Security Engineer Copyright © by **EC-Council**
All Rights Reserved. Reproduction is Strictly Prohibited.

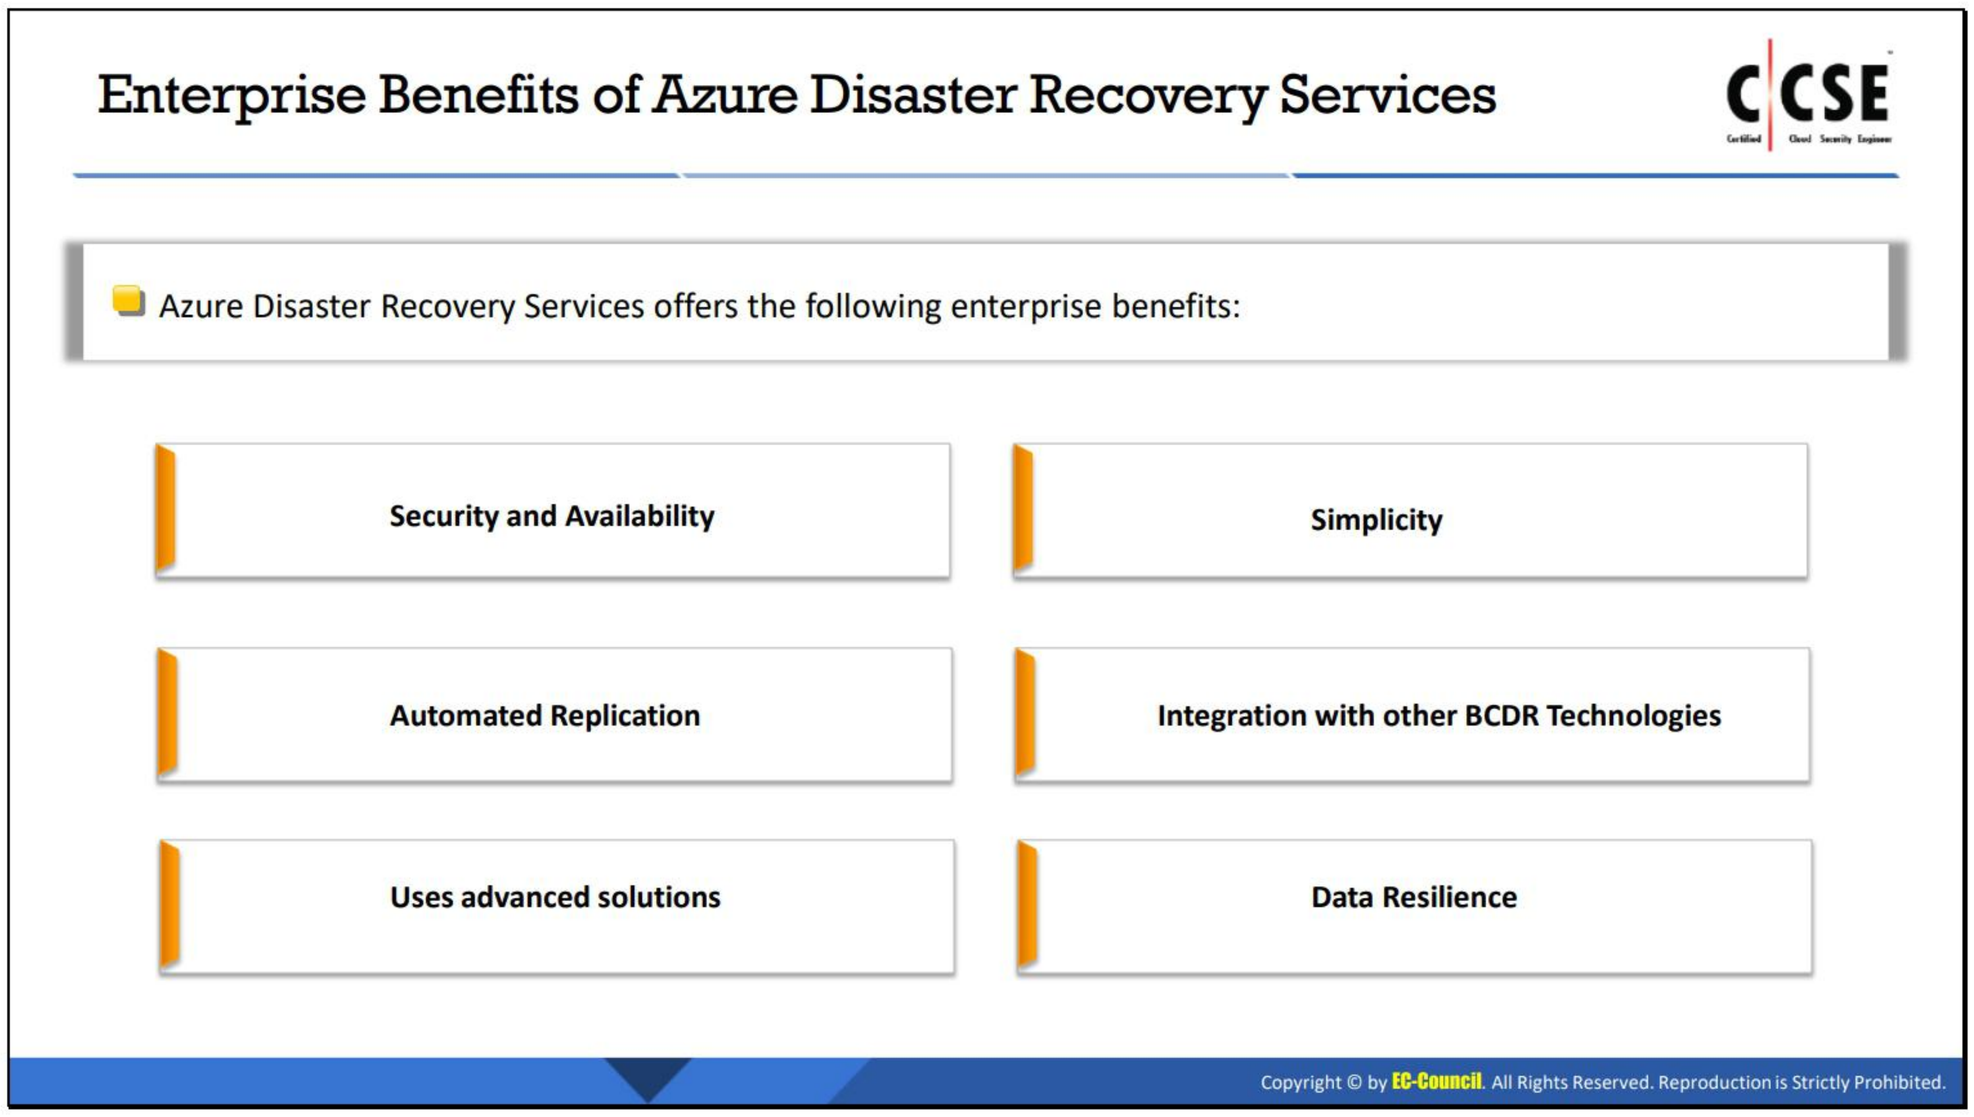

Enterprise Benefits of Azure Disaster Recovery Services

Azure Disaster Recovery Services offers the following enterprise benefits:

- **Security and availability:** Azure uses advanced security features such as multi-factor authentication (MFA) and encryption. These security features help a business fulfill compliance requirements and protect sensitive data. ASR offers 99.9% service availability with 24/7 support. Azure Backup stores the copies of data in three different locations within a primary Azure data center. It encrypts the data to ensure security and BC.
- **Simplicity:** ASR reduces the complexity of deploying and maintaining the infrastructure for DR. DR plans can be managed from a single location from the Azure portal.
- **Automated replication:** Downtime could cause a huge loss to the business and ASR reduces the downtime. ASR offers low RPO values. The RTO value can be minimized using automation. Azure Traffic Manager can also be used to make the DR process efficient and repeatable.
- **Integration with other BCDR technologies:** ASR can be integrated with other BCDR technologies to manage failover, etc.
- **Uses advanced solutions:** ASR uses advanced backup, disaster recovery, and backup solutions. A business of any size can make use of the Azure Disaster Recovery service to implement an effective BCDR plan.
- **Data resilience:** ASR orchestrates replication without affecting the application data. The data replicated to Azure is stored in Azure Storage with resilience. At the time of failover, VMs are created based on the replicated data.

Certified Cloud Security Engineer Copyright © by **EC-Council**
All Rights Reserved. Reproduction is Strictly Prohibited.

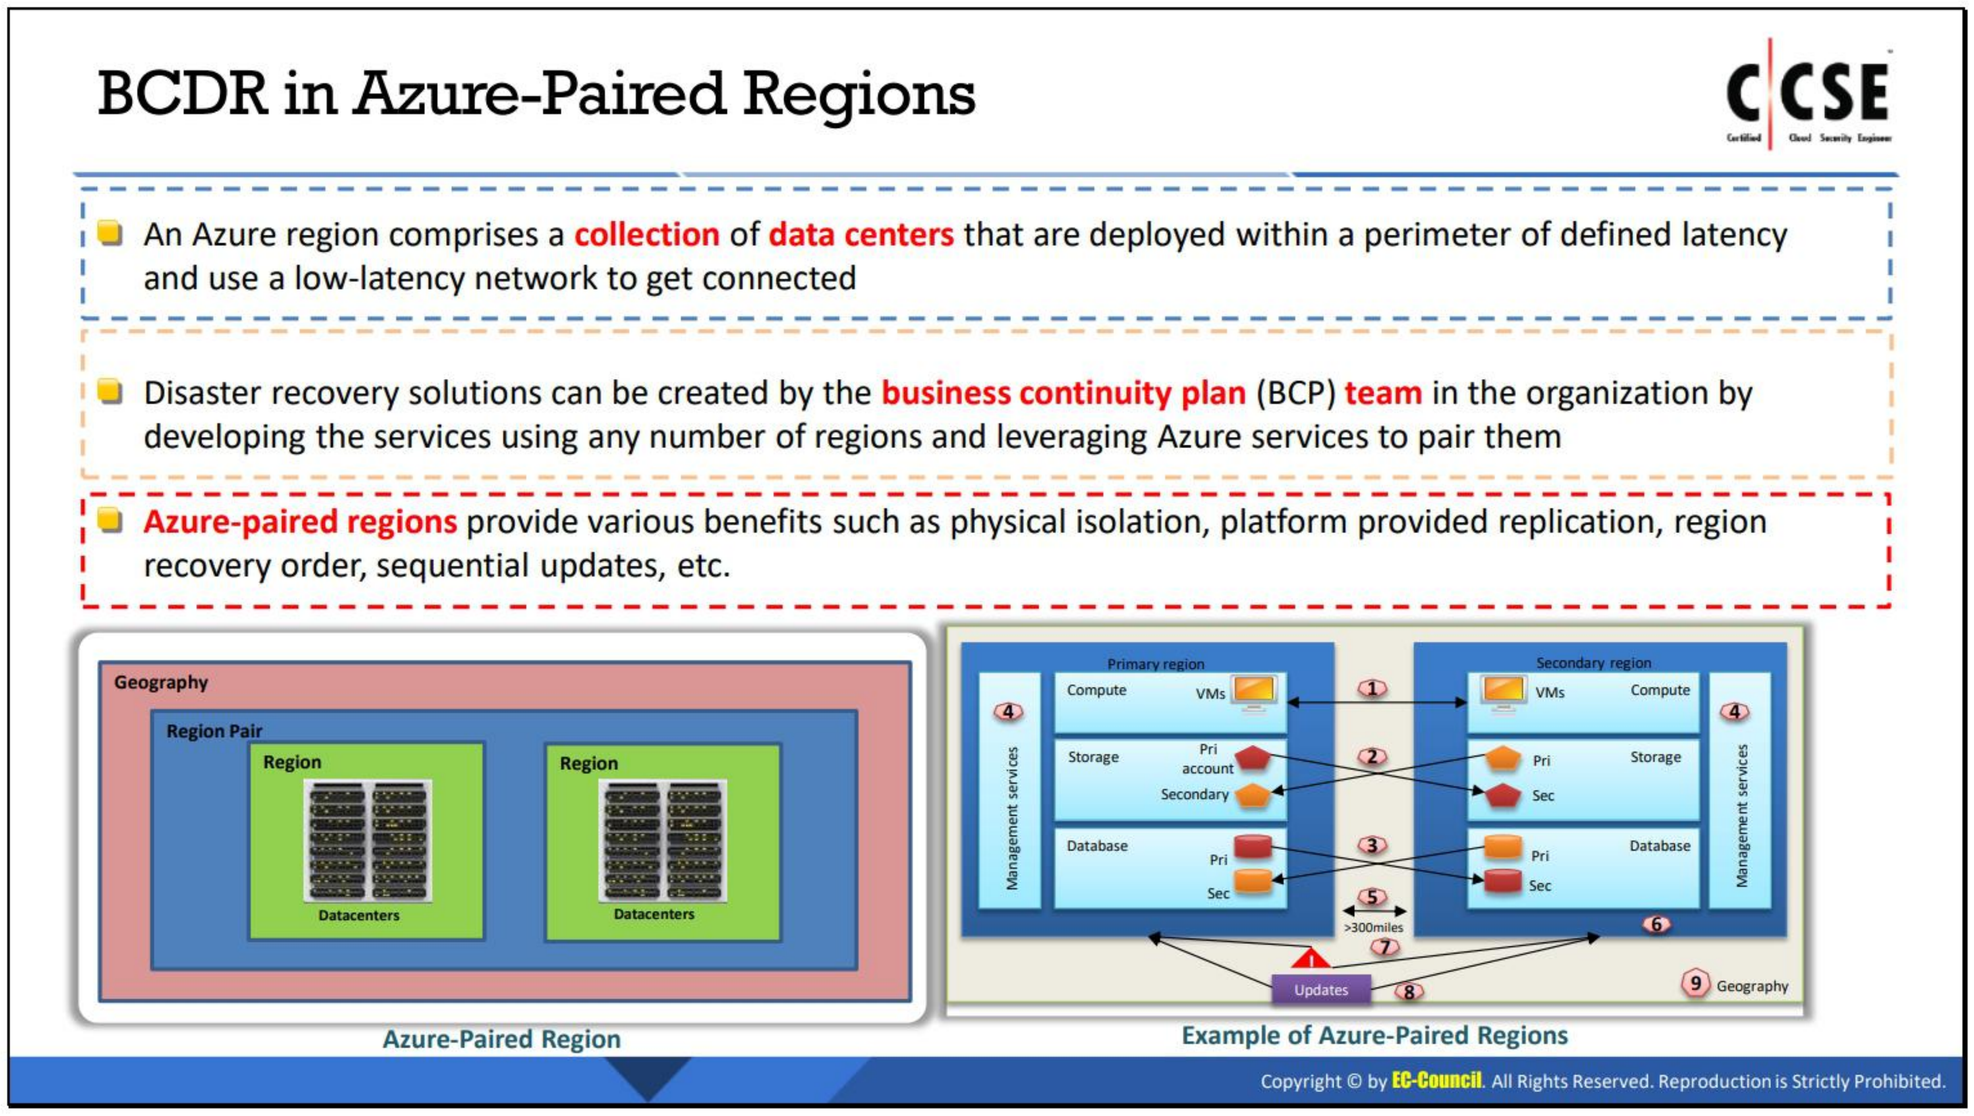

BCDR in Azure-Paired Regions

An Azure region comprises a collection of data centers that are deployed within a perimeter of defined latency and use a low-latency network to get connected. The use of low-latency networks makes sure that the Azure services provide the best possible performance and security. The geography defined by Azure is the geographical location (in any part of the world) containing at least one Azure region. Geographies are isolated areas in the world and they contain two or more regions that help to handle data redundancy issues.. Two regions located in the same geography comprise a regional pair. All the Azure regional pairs are updated with planned maintenance, and each pair should be updated at a time. At least one region in each pair will be prioritized for recovery when an outage affects multiple regions.

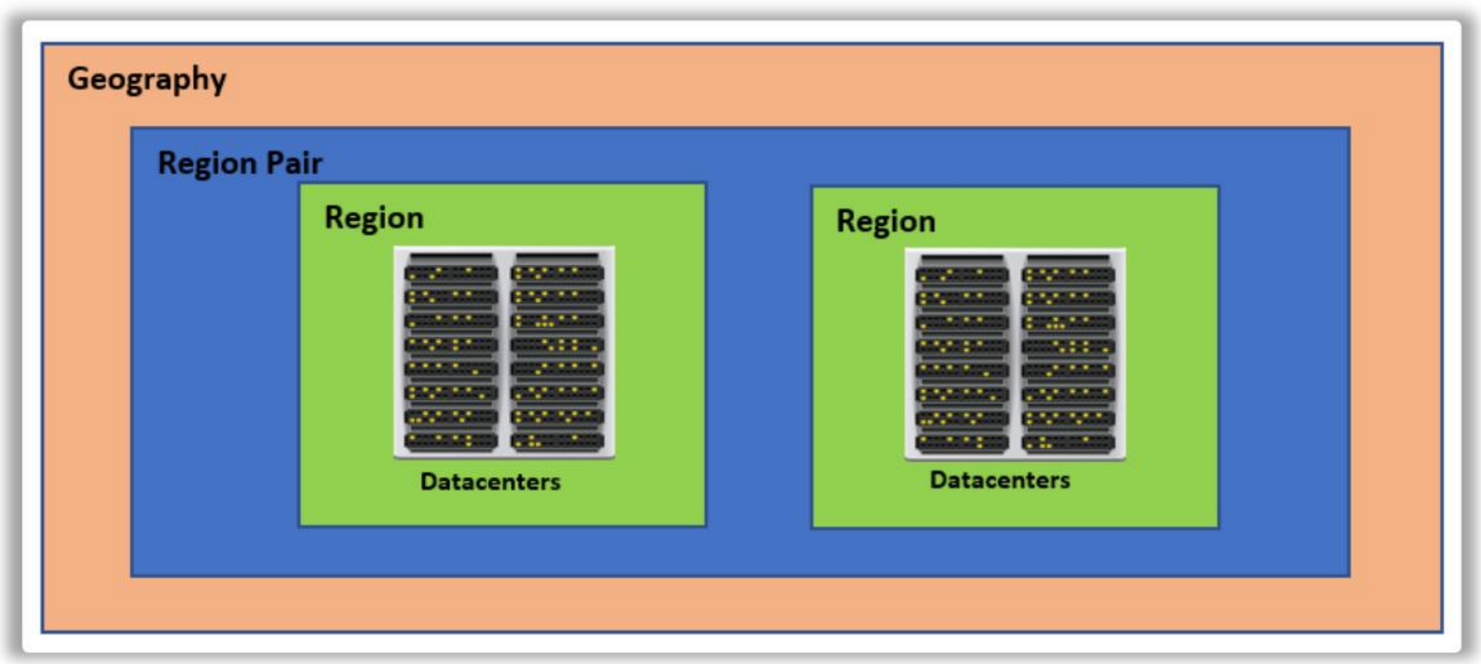

Figure 9.118: Azure Paired Regions

To ensure BC and protection against data loss, some of the Azure services take advantage of paired regions to ensure data availability. The advantage is nothing but the usage of secondary

Certified Cloud Security Engineer Copyright © by **EC-Council**
All Rights Reserved. Reproduction is Strictly Prohibited.

geo location in case of primary unavailability where the data is replicated to a secondary region by the Azure. All the Azure services might not be automatically replicated and changed to their secondary region. Recovery and replication can be configured in such cases by the customer.

Can Regional Pairs Be Selected?

There is no facility to select the regional pairs. Certain Azure services such as Azure redundant storage depend on regional pairs. They will not allow you to create or select any new regional pairings. Due to the planned maintenance and recovery prioritization pairings for regional pairs, there is no chance to create or take advantage of own regional pairs. However, customized disaster recovery solutions can be created by developing the services using any number of regions and leveraging Azure services to pair them. To survive the loss of the primary region services such as Azcopy, Azure DNS, and Azure Traffic Manager can be used to schedule backup and design a resilient architecture, respectively.

The below image shows the regional pair of disaster recovery. Cross-region activities are denoted with green numbers.

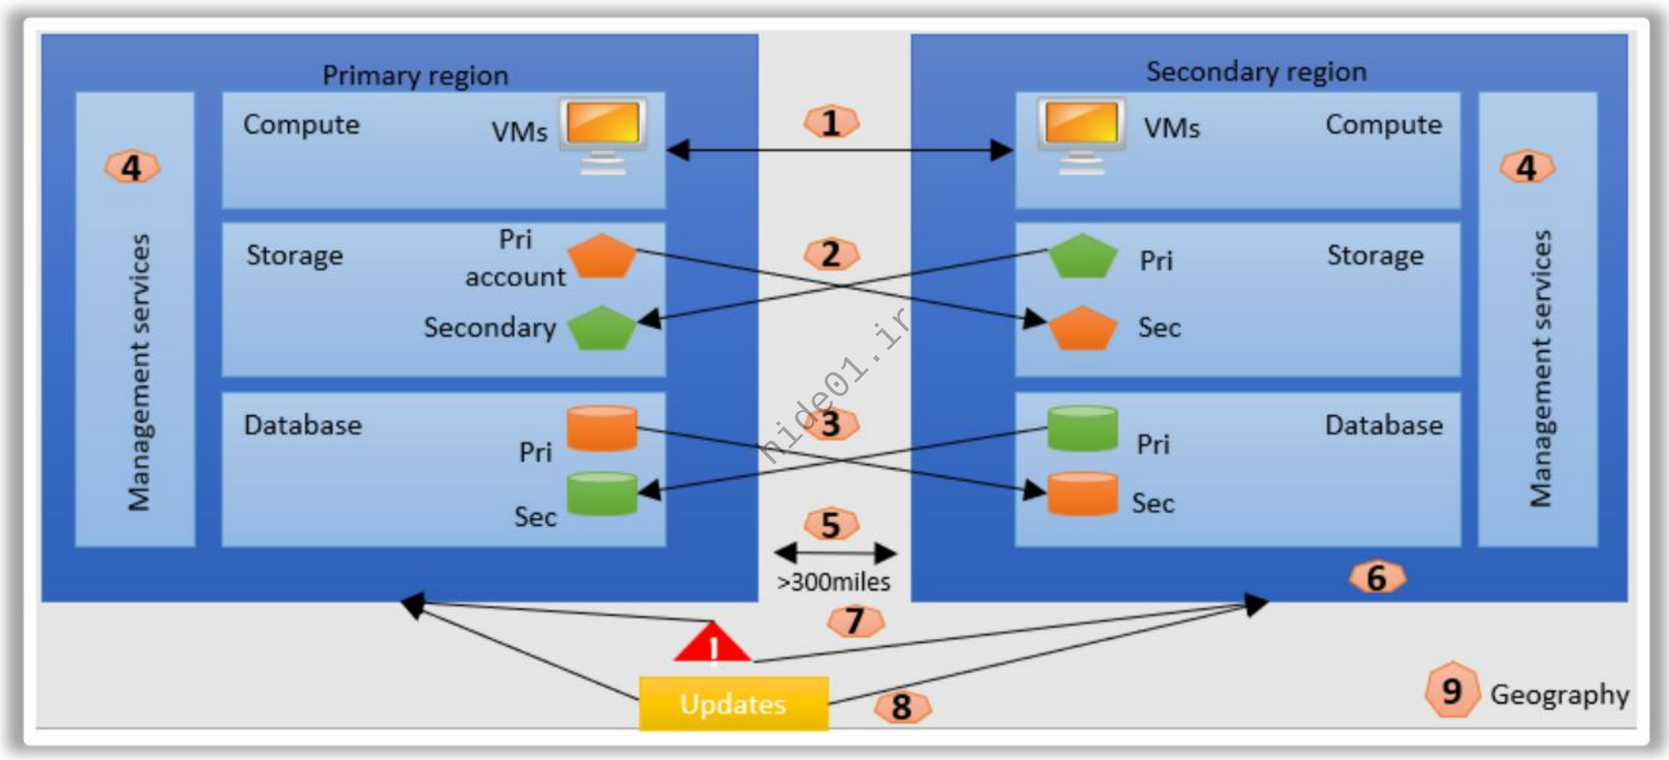

Figure 9.119: Hypothetical Azure Regional Pair

Cross-region Activities

- **Azure compute (IaaS):** During a disaster, make sure that the resources are available in another region; provision additional compute resources for this purpose.
- **Azure Storage:** The configuration of GRS is done by default at the time of creating a storage account. This will replicate the data automatically. The replication combination includes three times-primary region and three times-paired region.
- **Azure SQL Database:** To configure asynchronous replication to any region in the world, Azure SQL Database Geo-Replication can be used. This replicates the transactions. For disaster recovery scenarios, it is recommended to deploy these resources in the paired region.

Certified Cloud Security Engineer Copyright © by **EC-Council**
All Rights Reserved. Reproduction is Strictly Prohibited.

- **Azure Resource Manager:** The isolation of components across regions is provided by Azure Resource Manager. It reduces the impact of one region to another due to logical failures.

Benefits of Paired Regions

- **Physical isolation:** Azures tries to separate data centers with a physical distance of 300 miles wherever possible. This is subject to power, network infrastructure availability, geography size, and regulations, etc. It helps in avoiding natural disasters, physical network outages, and natural disasters as there are very low chances of them occurring together at both places.
- **The platform provided replication:** Automatic replication to the paired region is provided by GRS.
- **Region recovery order:** From every pair, recovery of one region is prioritized. It ensures the applications deployed across the paired regions can have one recovered priority region. There are chances of delay in the recovery for the applications that are deployed across unpaired regions. The selected regions may be the last two available the recovery in the worst case.
- **Sequential updates:** To minimize downtime, logical failures, and effect of bugs, the Azure system updates are sequentially rolled out to paired regions.
- **Data residency:** A region is a location within the same geography as its pair, except Brazil South. This meets the requirements of data residency for tax and law enforcement jurisdiction.

hide01.ir

Certified Cloud Security Engineer Copyright © by **EC-Council**
All Rights Reserved. Reproduction is Strictly Prohibited.

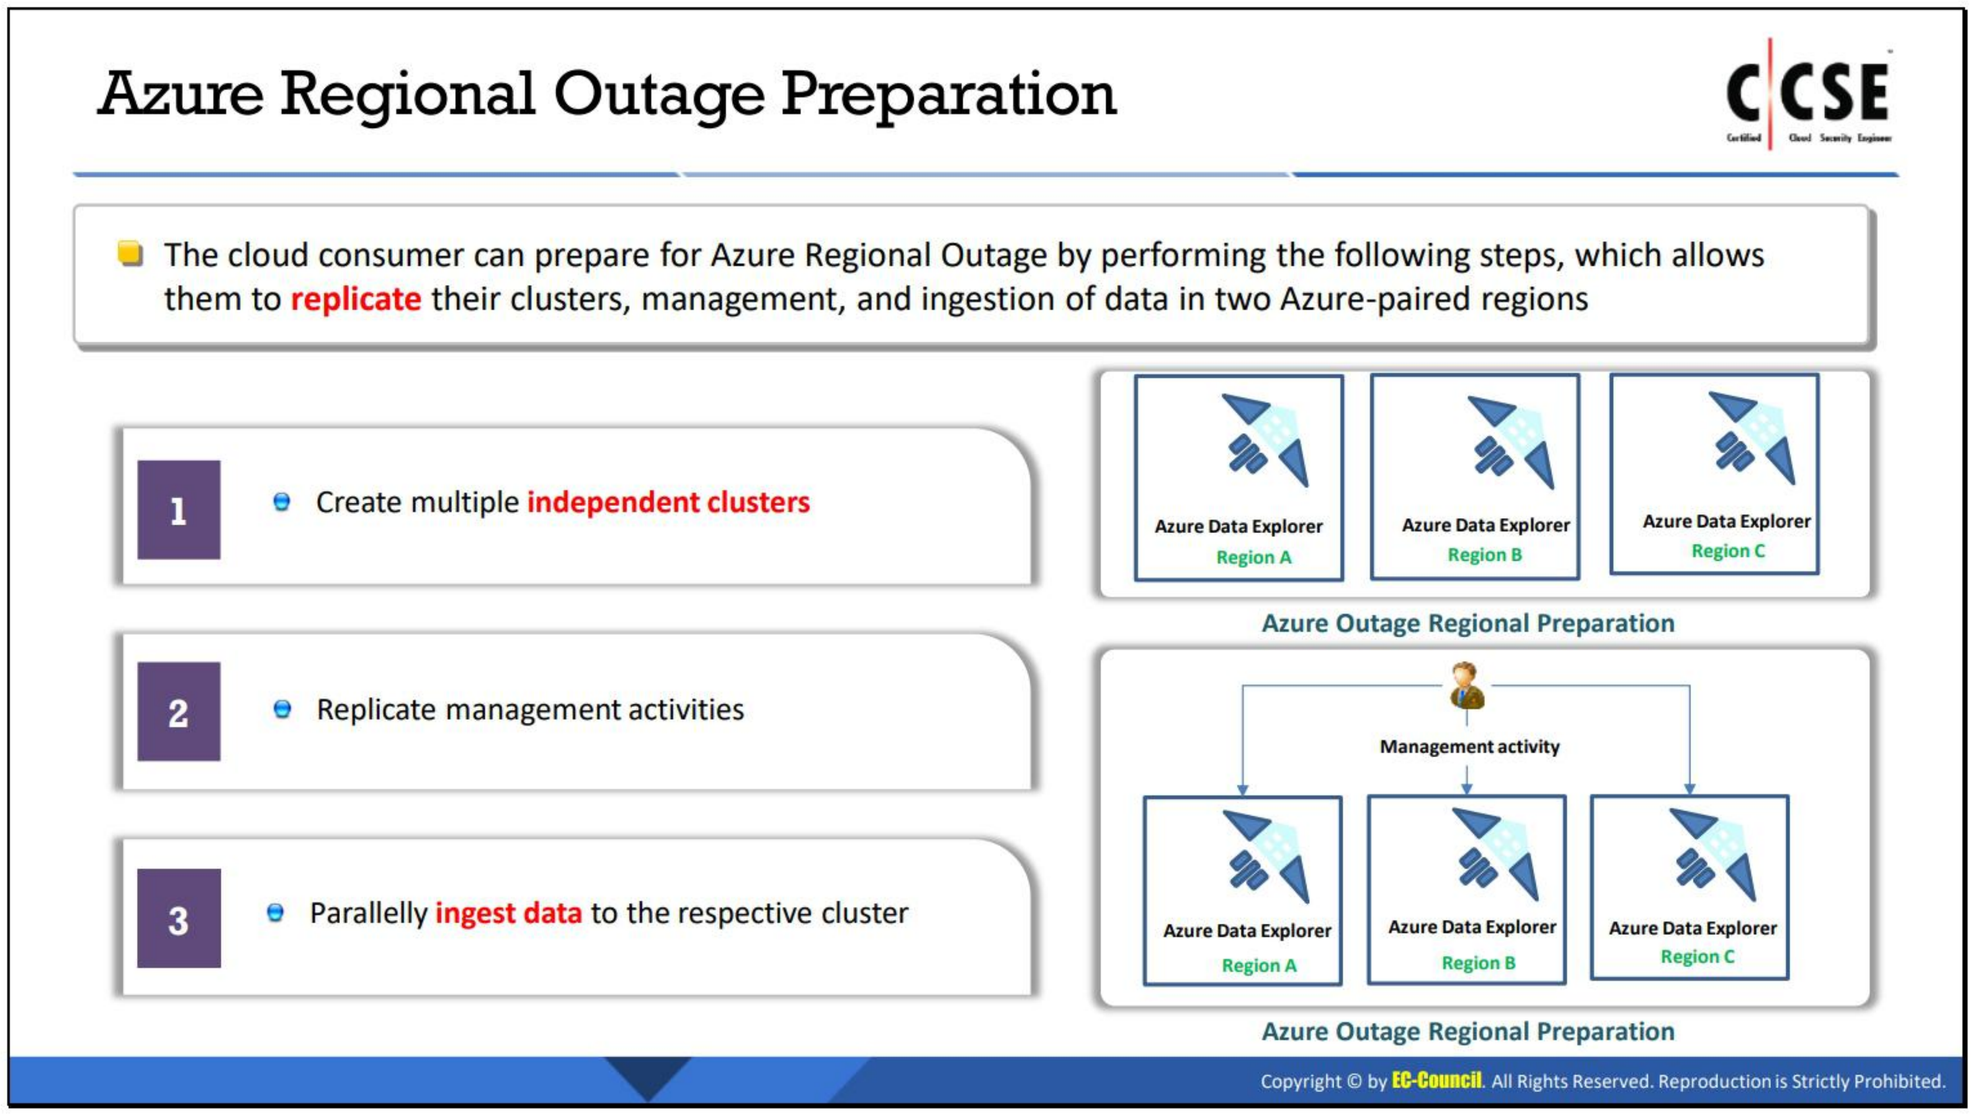

Azure Regional Outage Preparation

When an outage of entire Azure regions takes place, Azure Data Explorer does not support automatic protection. Such a disruption most likely occurs during a natural disaster. To ensure continuity of business, the following steps should be performed if a solution for a disaster recovery situation is required. Clusters, data ingestion, and management are replicated in two Azure-paired regions.

- In Azure-paired regions, create two or more independent clusters.
- Management activities replication like managing user roles or creating new tables on each cluster.
- Ingest data to each cluster in parallel.

Create Multiple Independent Clusters

Azure Data Explorer cluster should be created in more than one region. Two of the clusters among them should be created in the Azure-paired region.

Figure 9.120: The Three Clusters in Three Different Regions

Certified Cloud Security Engineer Copyright © by **EC-Council**
All Rights Reserved. Reproduction is Strictly Prohibited.

Replicate Management Activities

For every replica to have the same cluster configuration, the management activities have to be replicated.

- The same has to be created on each replica:
  - Databases: Either Azure portal or one of the Microsoft SDKs can be used for the creation of the database.
  - Tables
  - Mappings
  - Policies
- Authentication and authorization should be managed on each replica.

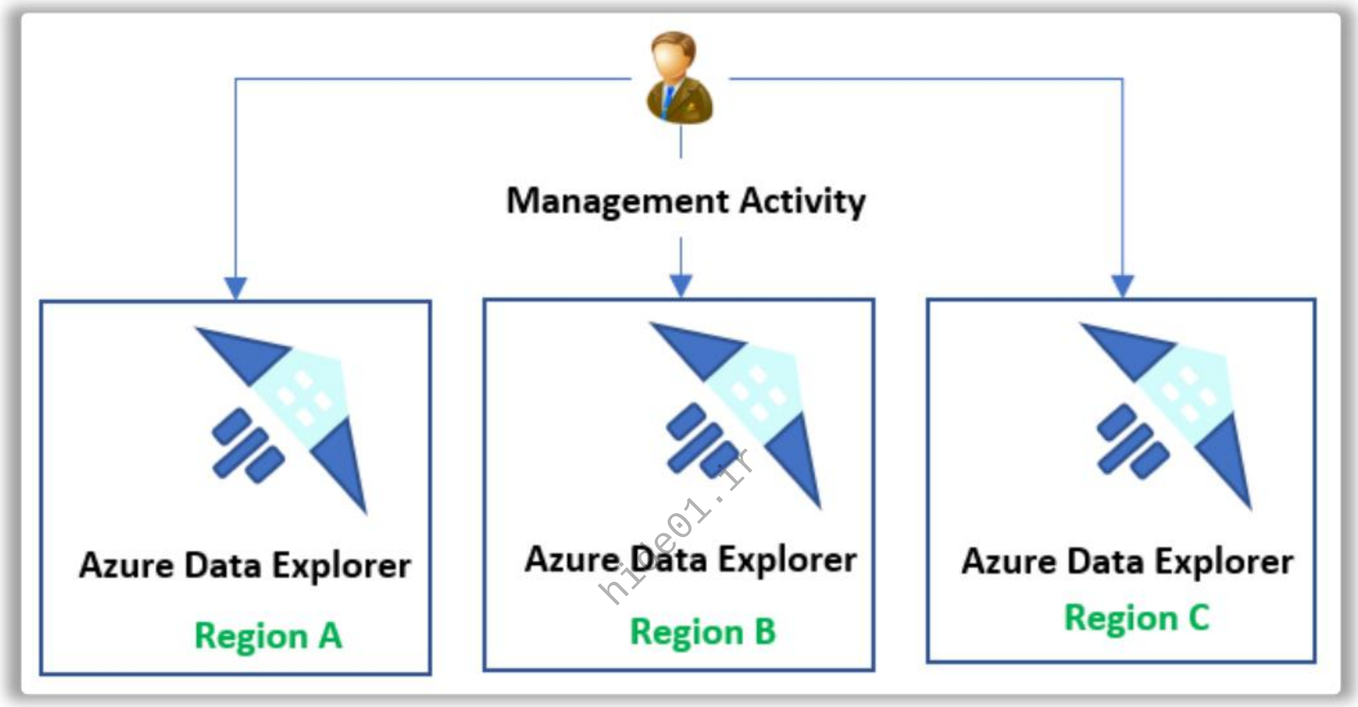

Figure 9.121: Image Showing the Replication of Management Activities

Certified Cloud Security Engineer Copyright © by **EC-Council**
All Rights Reserved. Reproduction is Strictly Prohibited.

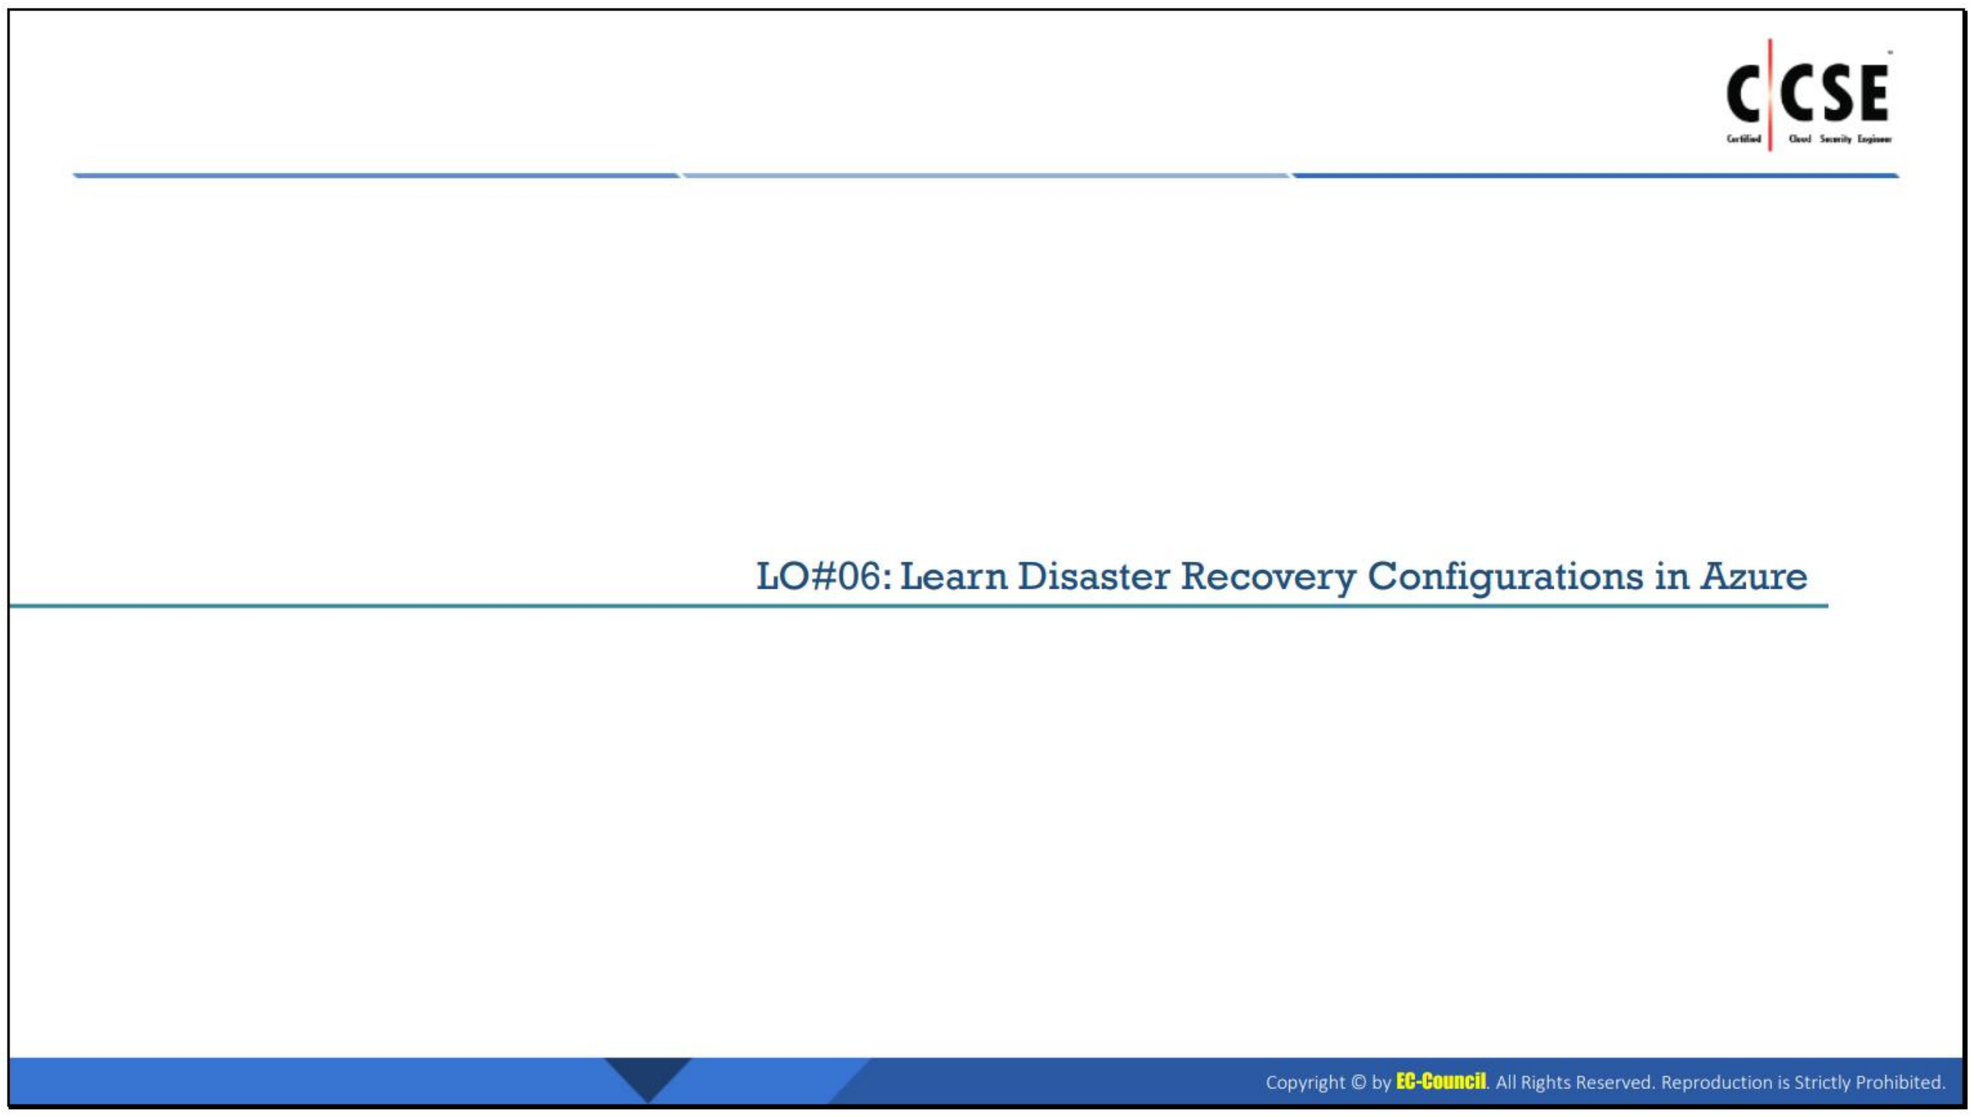

LO#06: Learn Disaster Recovery Configuration in Azure

The objective of this section is to understand different DR configurations that are available in Azure. This section will help cloud security engineers understand and decide which DR configuration in Azure is suitable for their organization.

hidari.jp

Certified Cloud Security Engineer Copyright © by **EC-Council**
All Rights Reserved. Reproduction is Strictly Prohibited.

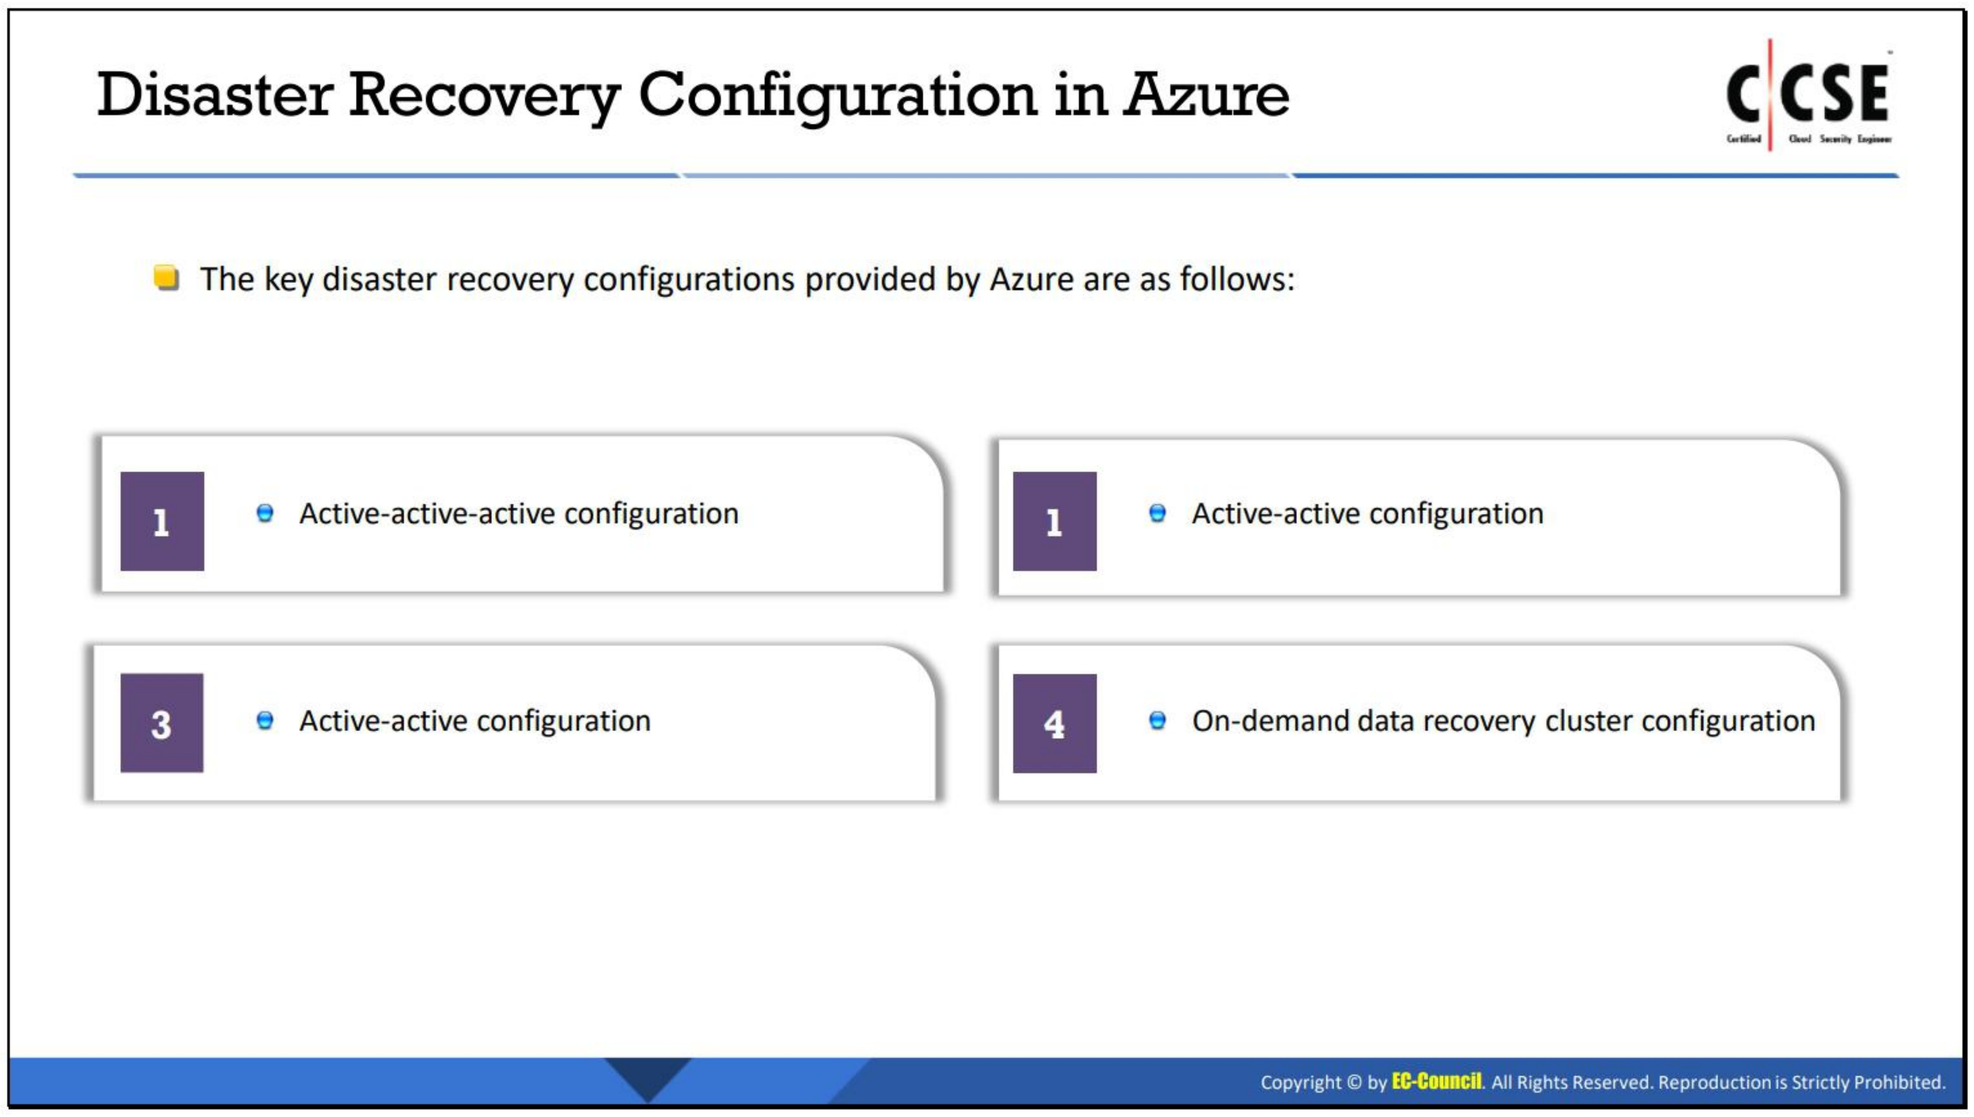

Disaster Recovery Configuration in Azure

As mentioned earlier, the time to recover from disruption is known as RTO. For example, if the RTO is 2 hours, it means that the application should be brought back to the normal state and it should be running within 2 hours of the disruption. The time frame during the disruption before the quantity of data lost during that period greater than the allowable threshold is RPO.

Consider a scenario where the RPO is 24 hours and the application consists of data from initial 15 years; this is still within the limits of agreed-upon RPO parameters. For disaster recovery, the processing, ingestion, and curation process should be planned with utmost care. Integration of various sources data into Azure Data Explorer is “Integration.” Transformation and other similar activities come into “processing.” Materialized views, exports to the data lake come into “curation.” Popular disaster recovery configurations are listed below:

- Active-active-active configuration
- Active-active configuration
- Active-hot standby configuration
- On-demand data recovery cluster configuration

Certified Cloud Security Engineer Copyright © by **EC-Council**
All Rights Reserved. Reproduction is Strictly Prohibited.

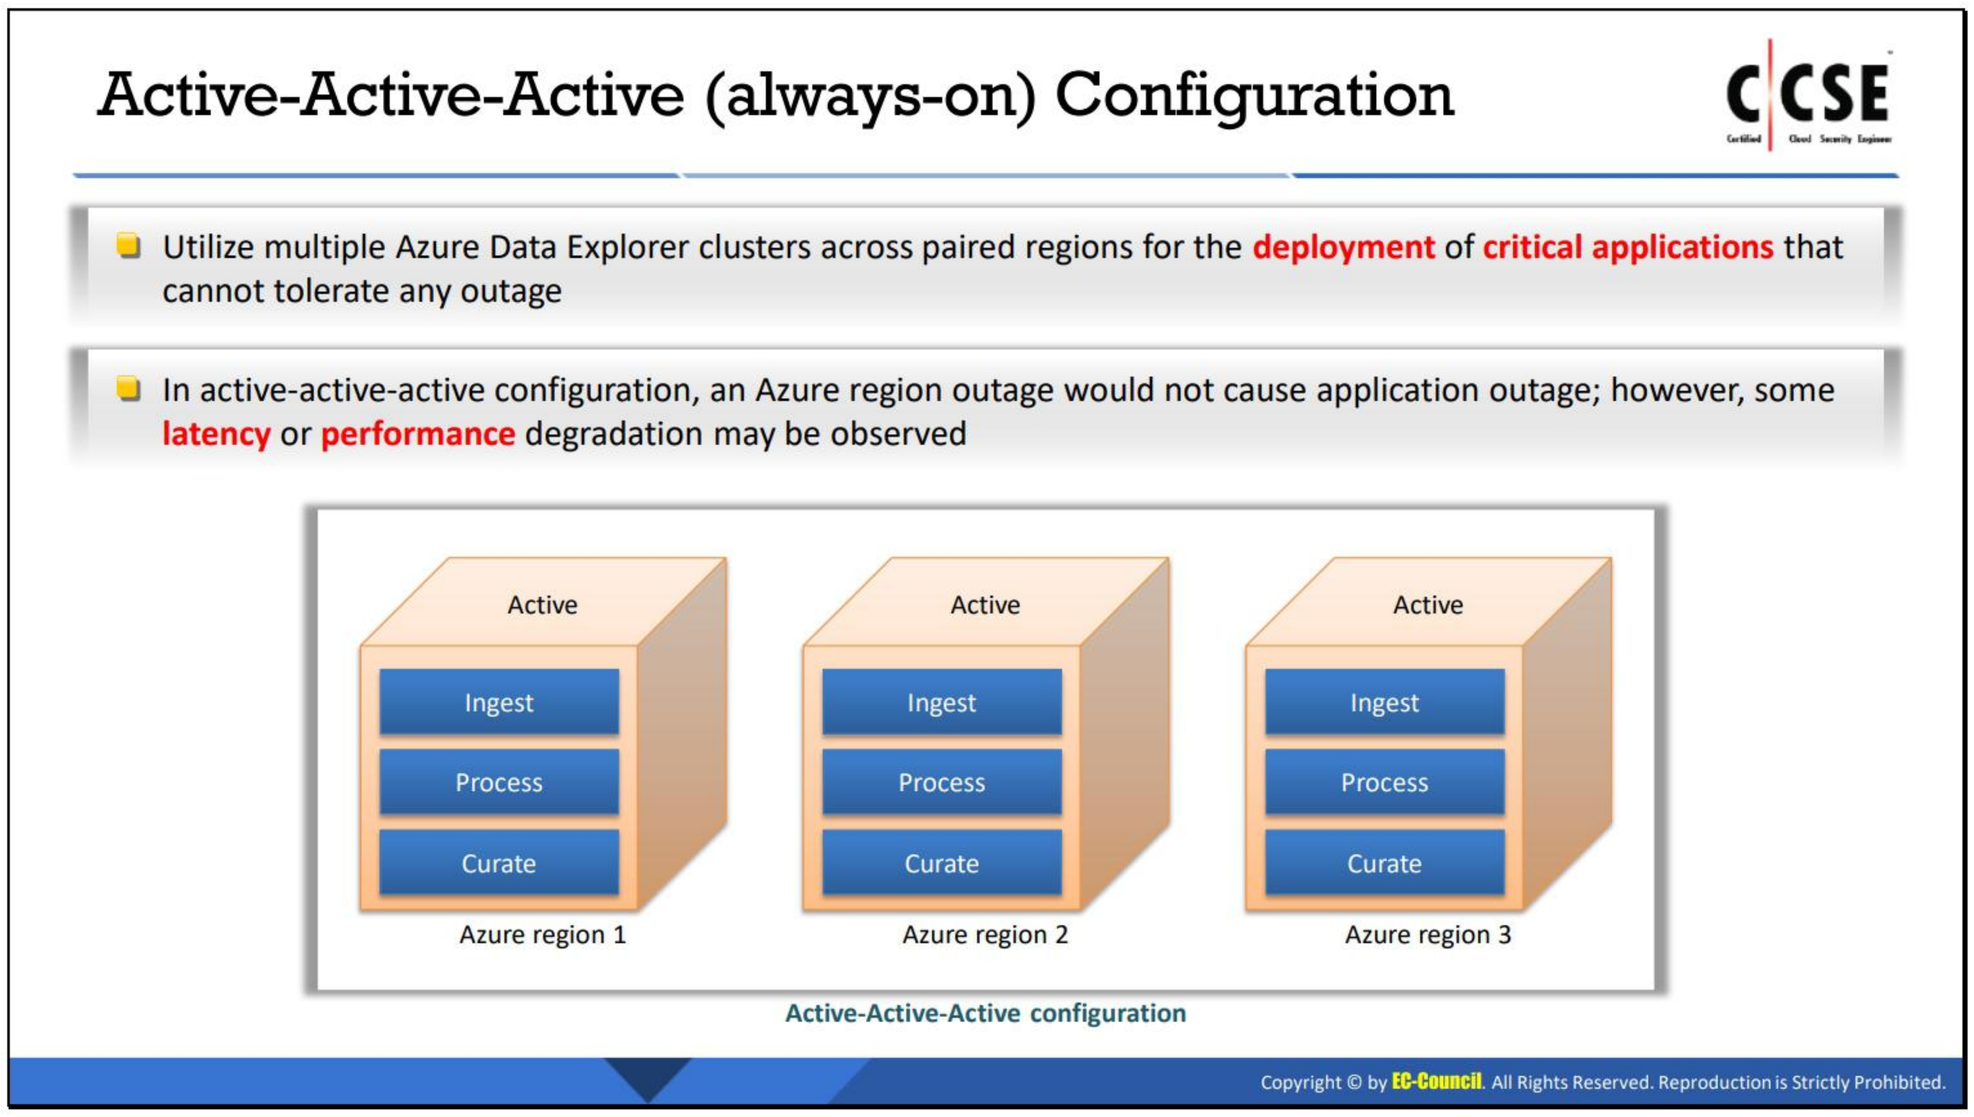

Active-Active-Active (always-on) Configuration

This is also called an "always-on" configuration. Multiple Azure Data Explorer clusters across Azure-paired regions should be used in no tolerance for outages and for the deployment of critical applications. For all the clusters, processing, ingestion, and curation should be set up. Stock-keeping unit (SKU) of the cluster must be the same across regions. Make sure that the updates are rolled out stumbled across Azure-paired regions. Some latency may be noticed instead of application outage due to Azure region outage.

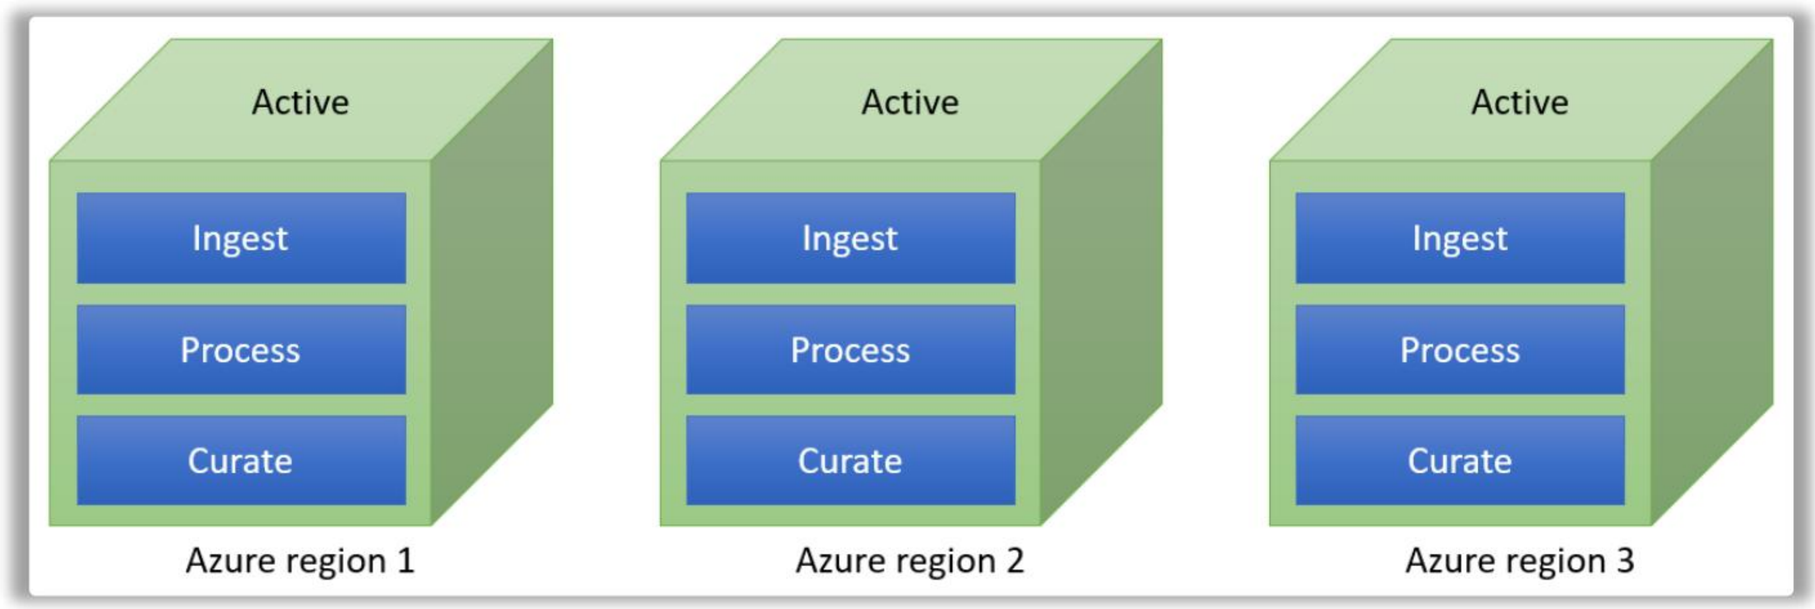

Figure 9.122: Active-active-active Configuration

Certified Cloud Security Engineer Copyright © by **EC-Council**
All Rights Reserved. Reproduction is Strictly Prohibited.

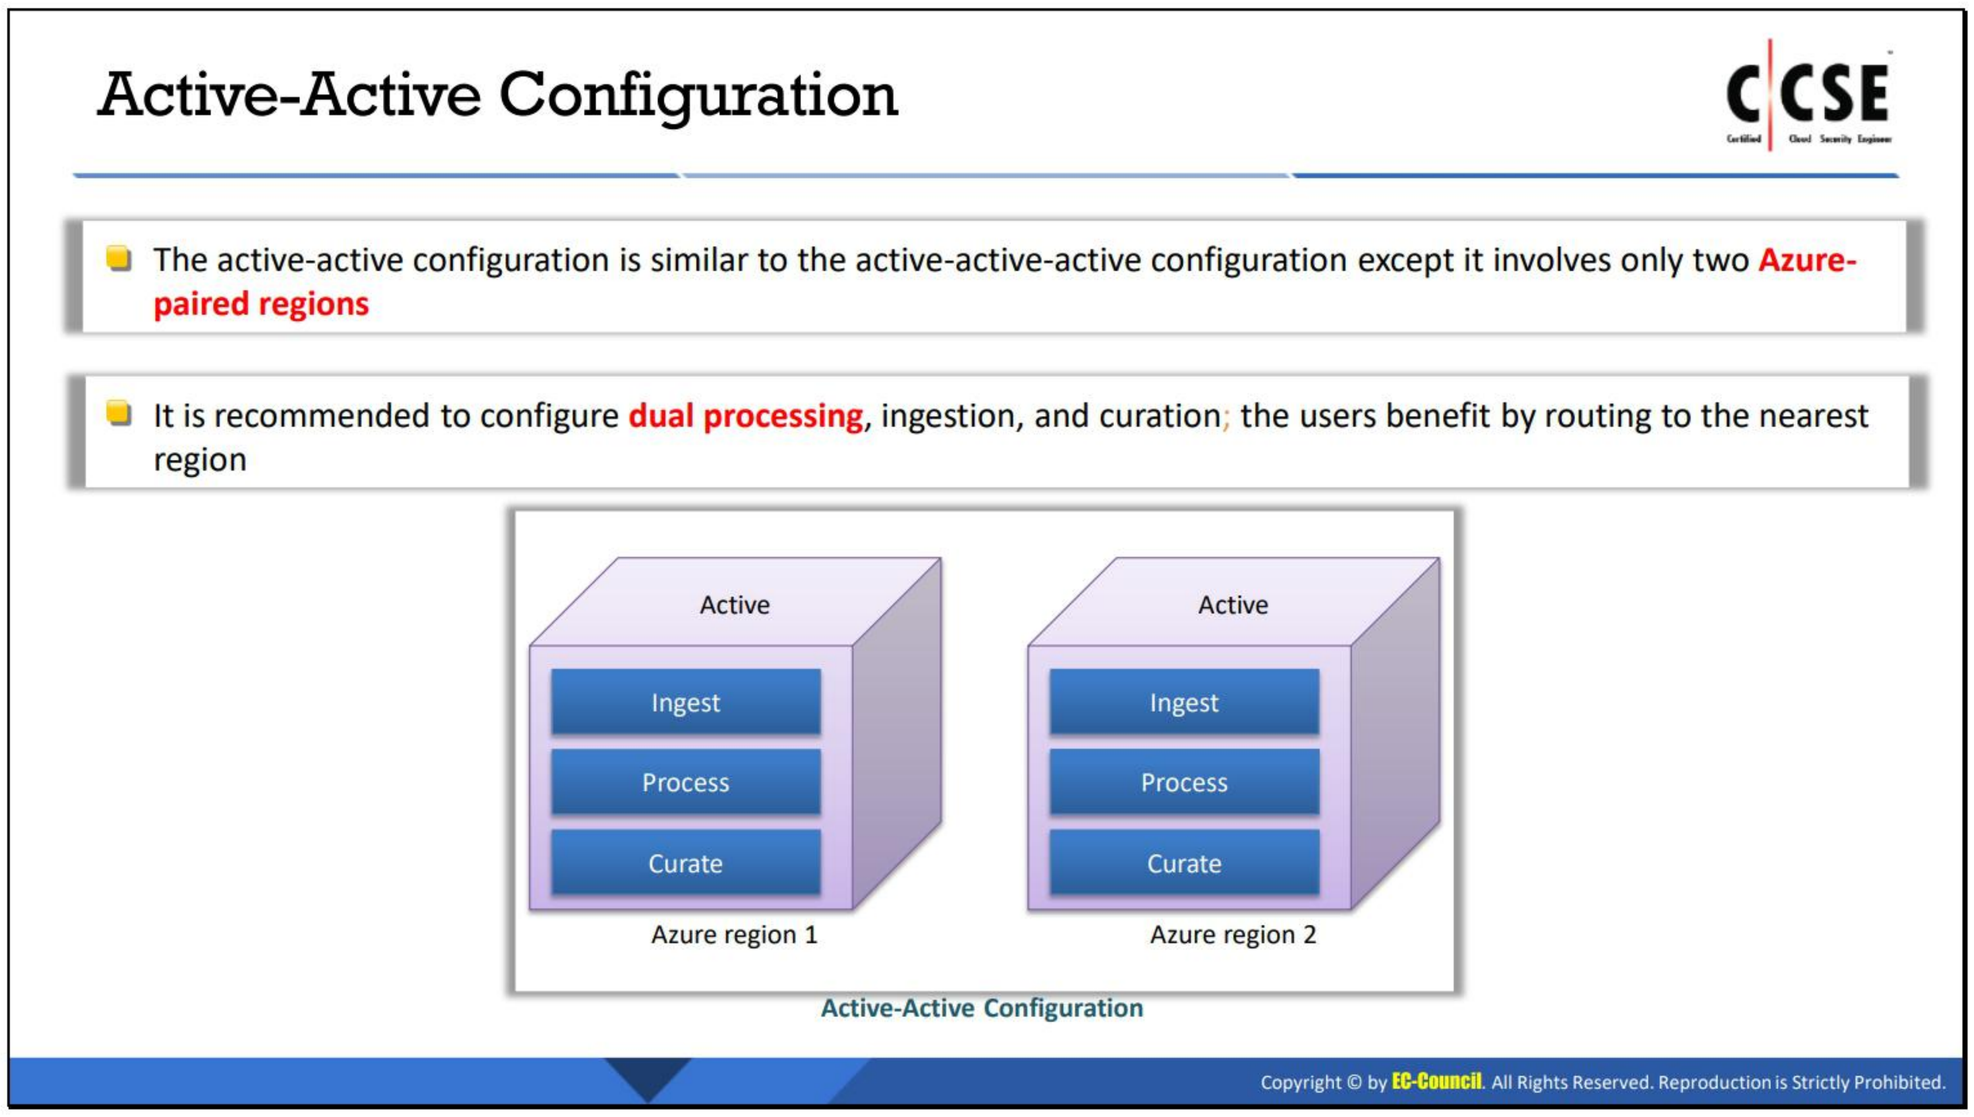

Active-Active Configuration

The active-active configuration involves only two Azure-paired regions; apart from this, it is similar to the active-active-active configuration. SKU of the cluster should be the same across regions. It is recommended to configure dual processing, ingestion, and curation. Here, the users are benefited by routing to the nearest region.

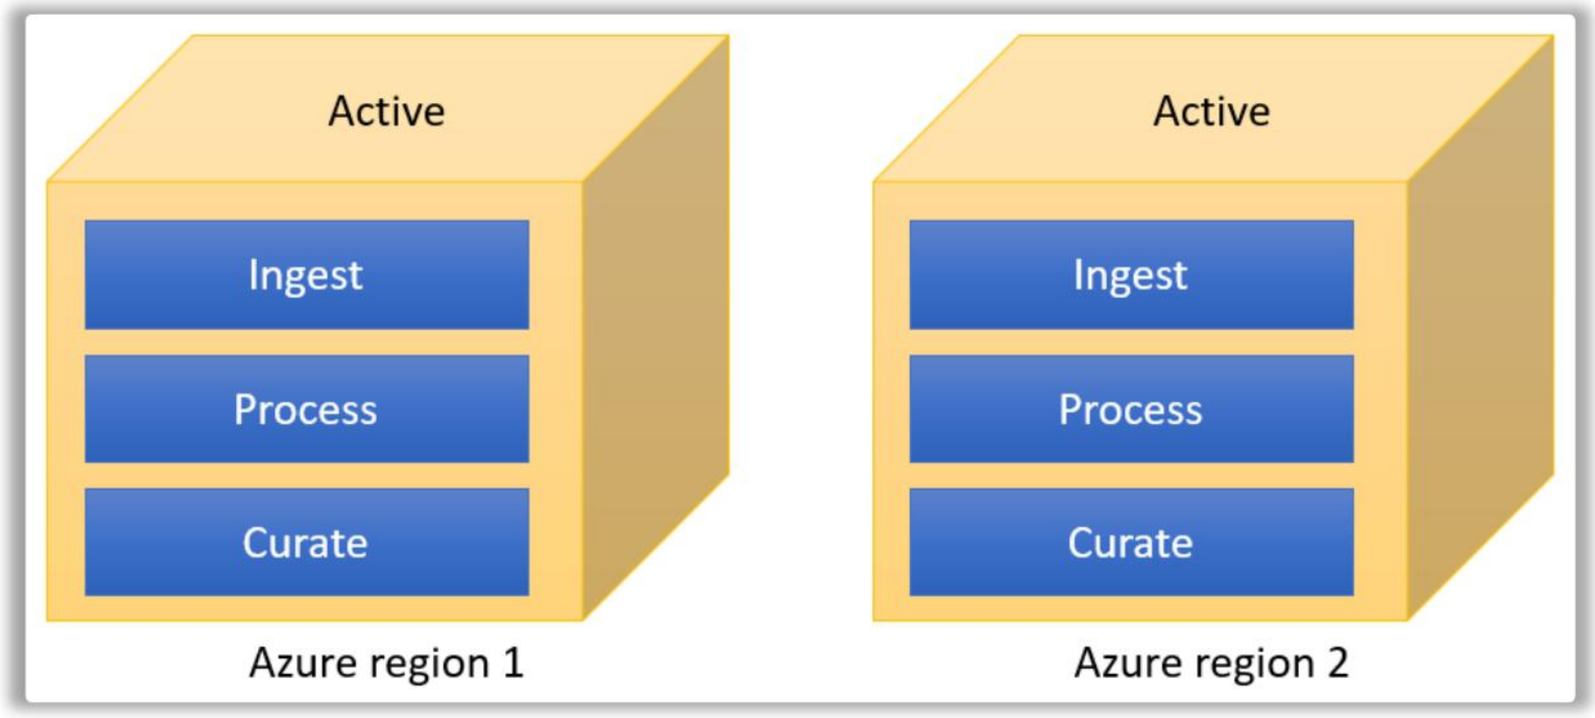

Figure 9.123: Active-active Configuration

Certified Cloud Security Engineer Copyright © by **EC-Council**
All Rights Reserved. Reproduction is Strictly Prohibited.

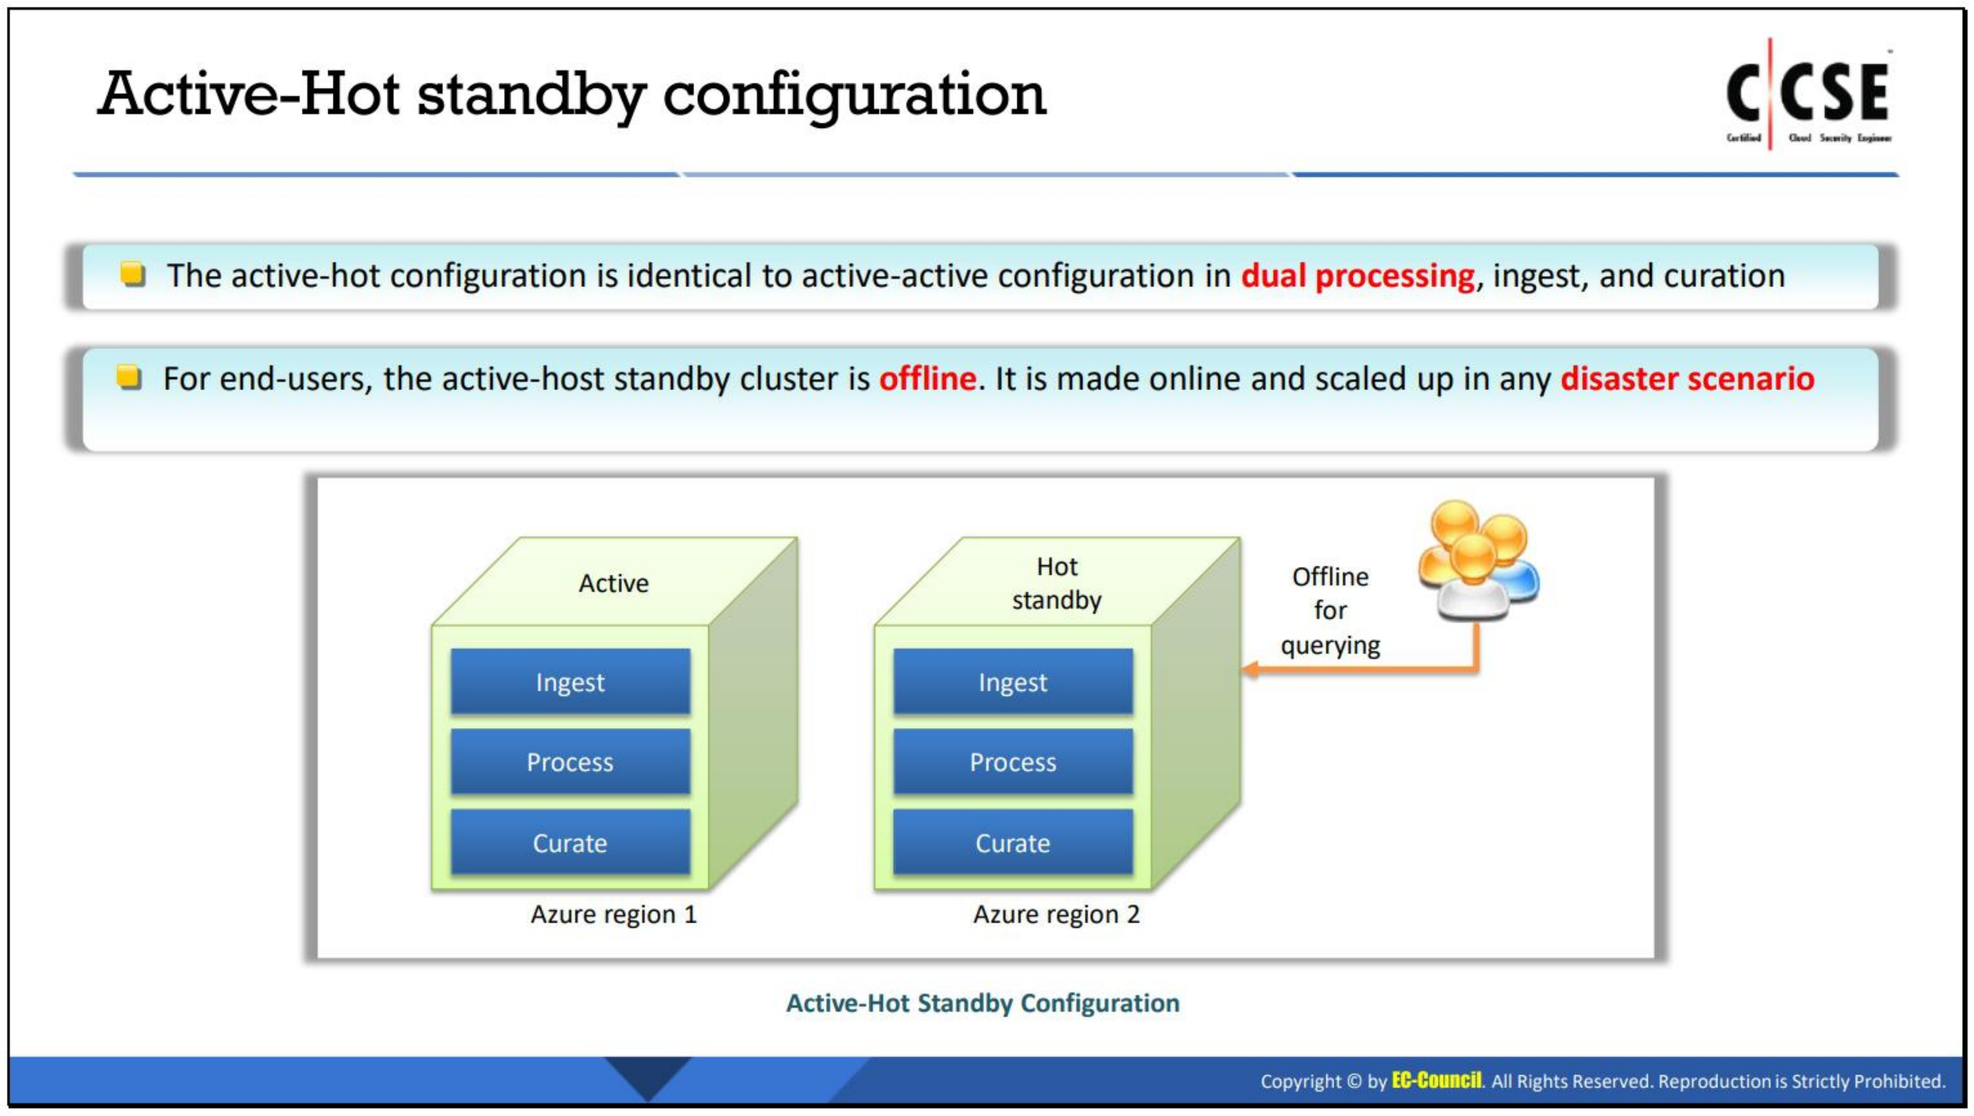

Active-Hot Standby Configuration

In dual processing, ingest, and curation point of view, the active-hot configuration and active-active configuration are identical. Here, the cluster need not be in the same SKU as the primary as the standby cluster is offline to end-users. SKU and scale can be smaller in hot standby and as such are less performant. This standby cluster will be made online and is scaled up in any disaster scenario.

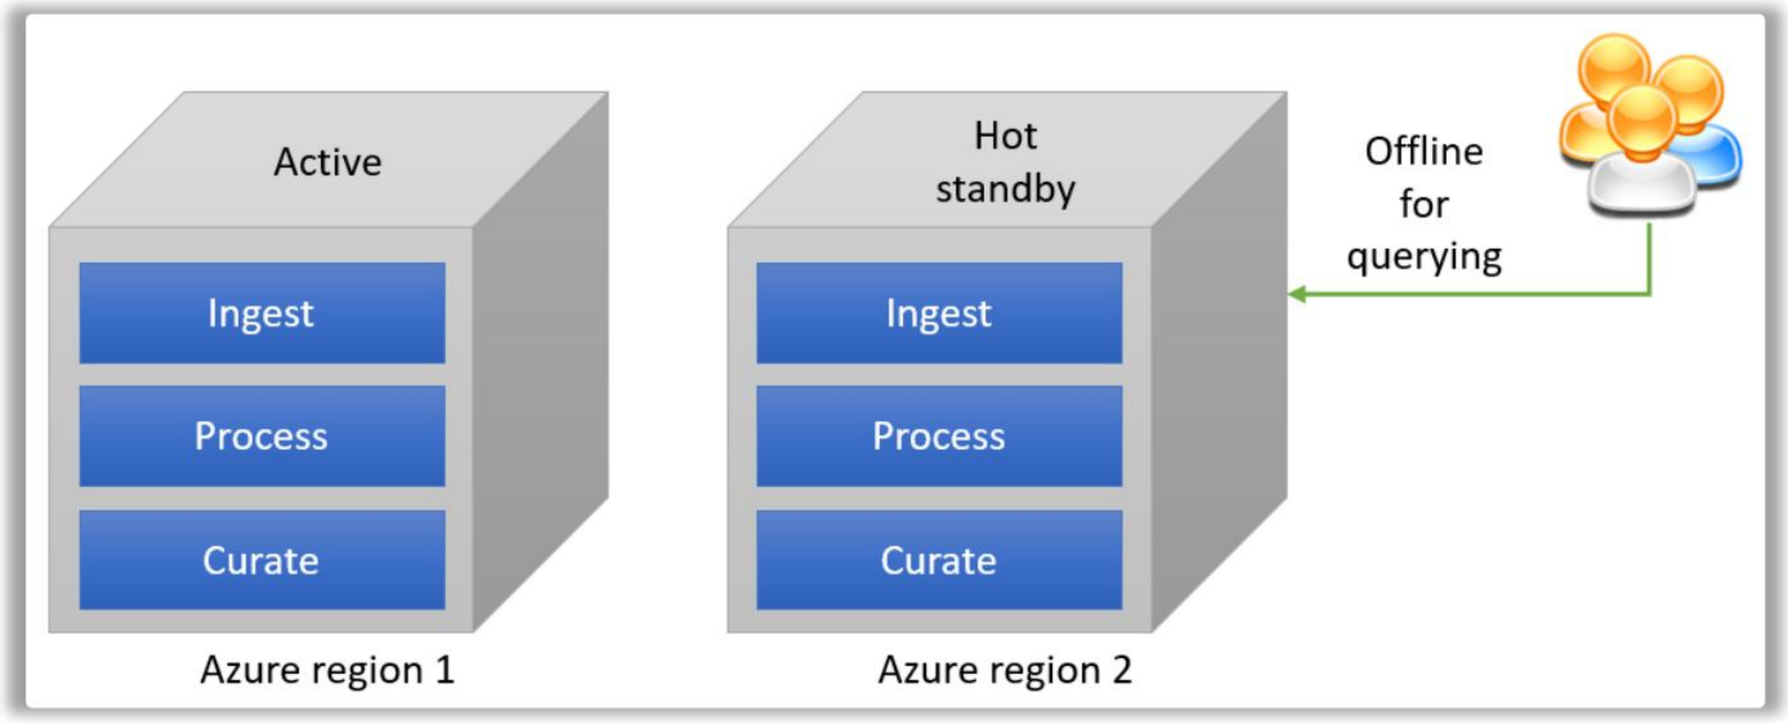

Figure 9.124: Active-Hot Standby Configuration

Certified Cloud Security Engineer Copyright © by **EC-Council**
All Rights Reserved. Reproduction is Strictly Prohibited.

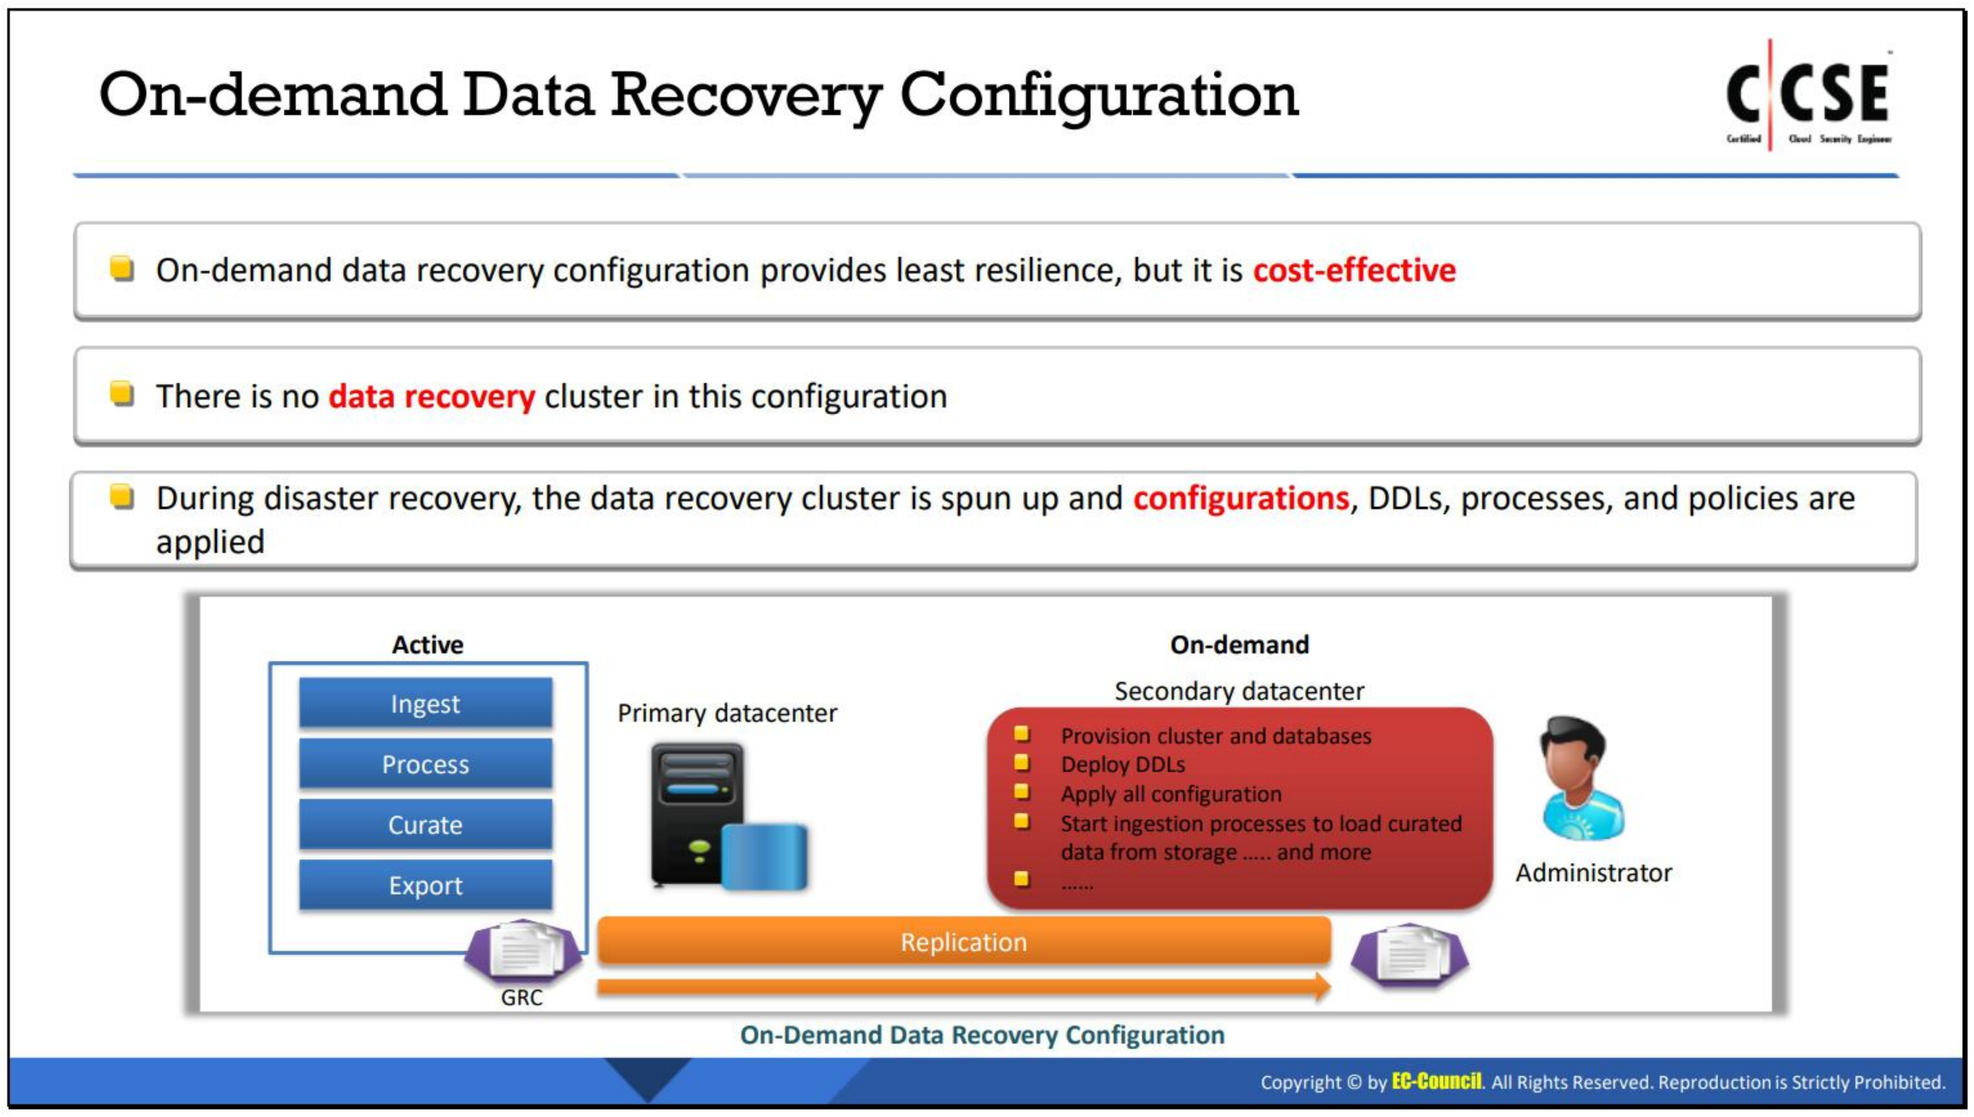

On-Demand Data Recovery Configuration

The least resilience is offered by on-demand data recovery configuration, that is, it offers the highest RPO, RTO, and efforts toward them are high. Fortunately, the cost required for the solution is too low. There is no data recovery cluster in this configuration. GRS storage account receives the continuous configured exported curated data. If there is a DR, then the DR cluster is turned up. During the data recovery cluster usage, configurations, DDLs, processes and, policies are applied. To over-ride the ingestion time that defaults to system time, data is ingested from storage with the ingestion property “kustoCreationTime”.

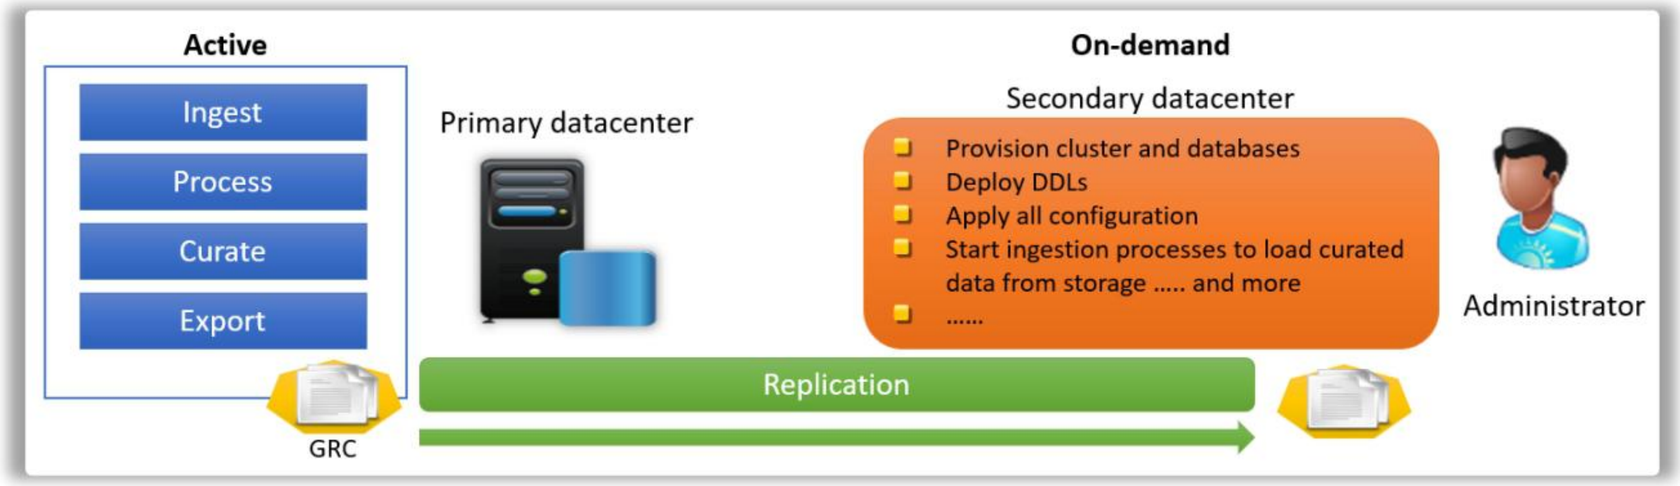

Figure 9.125: Image Showing On-demand Data Recovery Configuration

Certified Cloud Security Engineer Copyright © by **EC-Council**
All Rights Reserved. Reproduction is Strictly Prohibited.

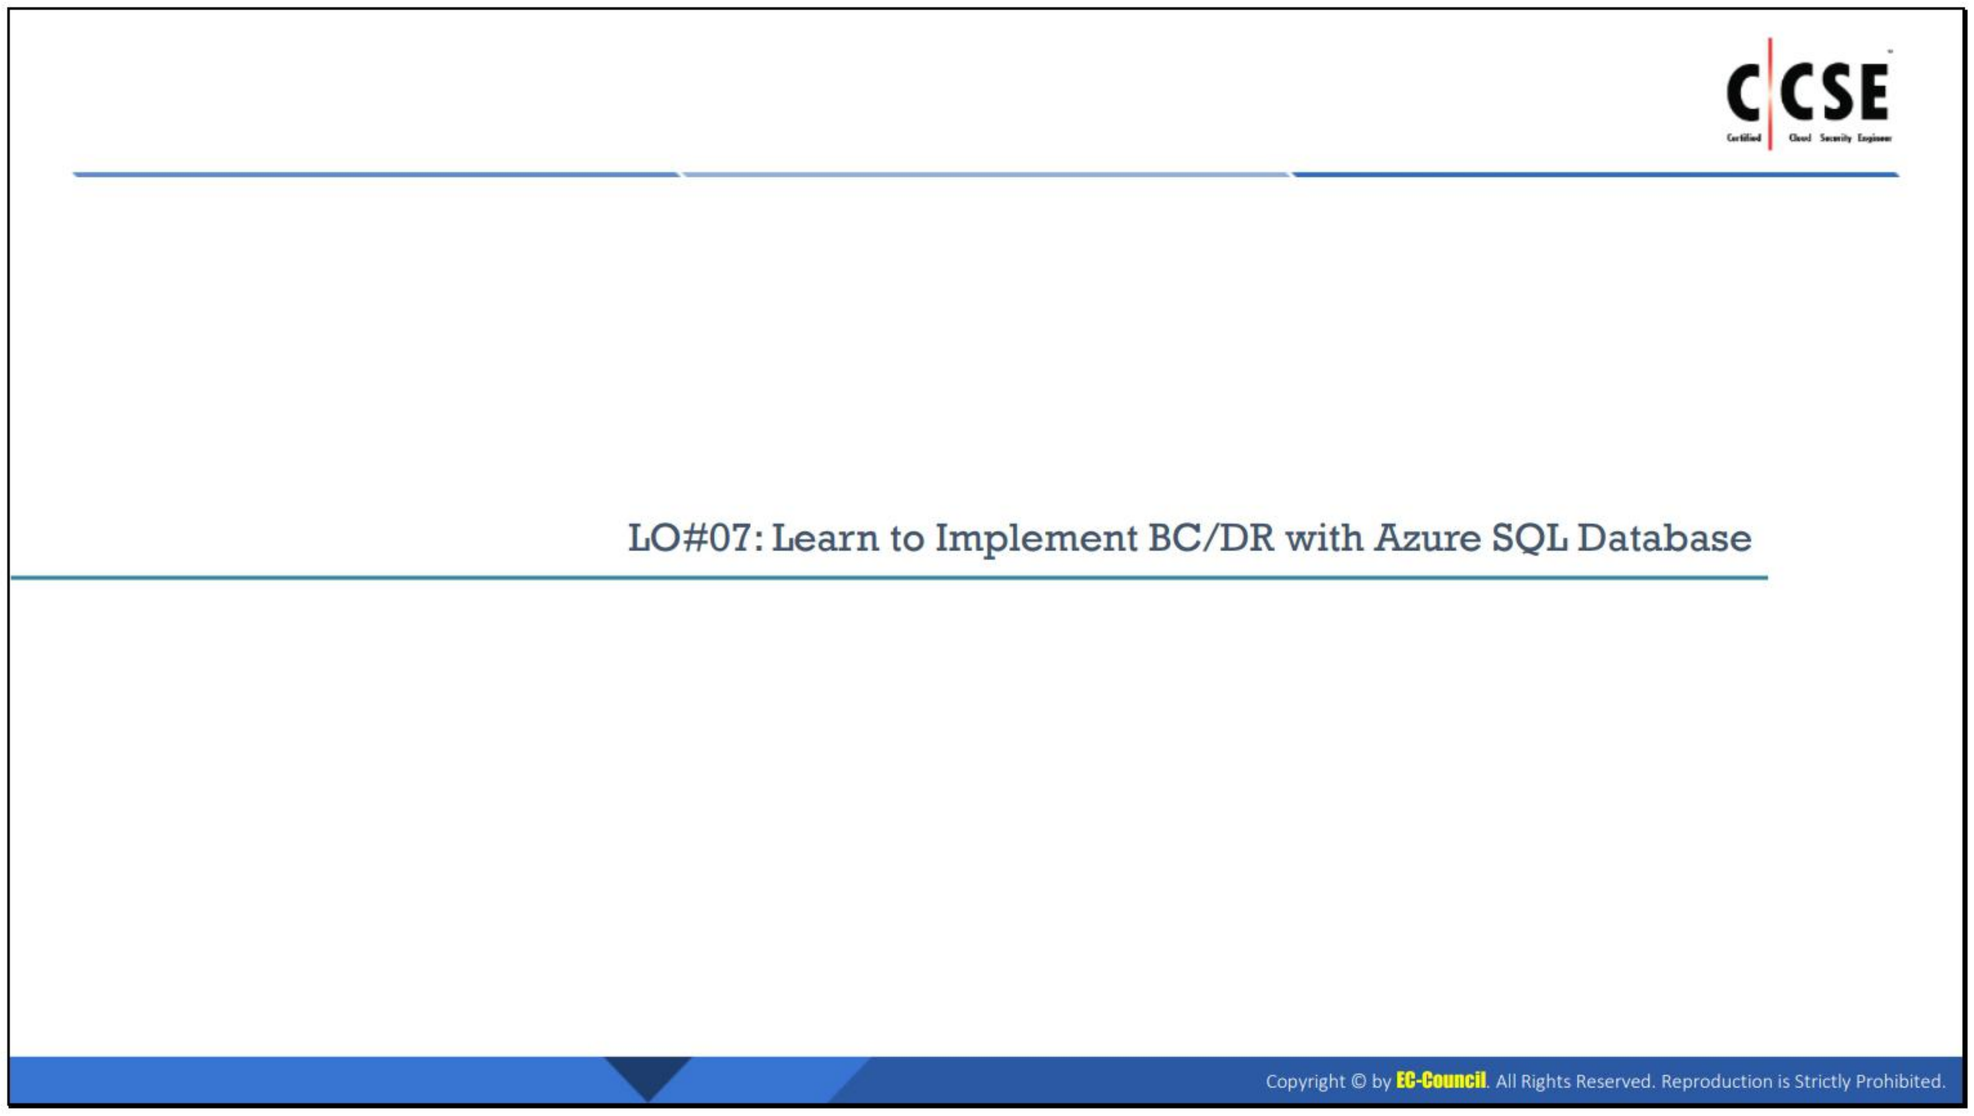

LO#07: Learn to Implement BCDR with Azure SQL Database

The objective of this section is to understand the BCDR features and capabilities of Azure SQL databases. Understanding these features and capabilities will help cloud security engineers perform backup and recovery using SQL databases.

hidden

Certified Cloud Security Engineer Copyright © by **EC-Council**
All Rights Reserved. Reproduction is Strictly Prohibited.

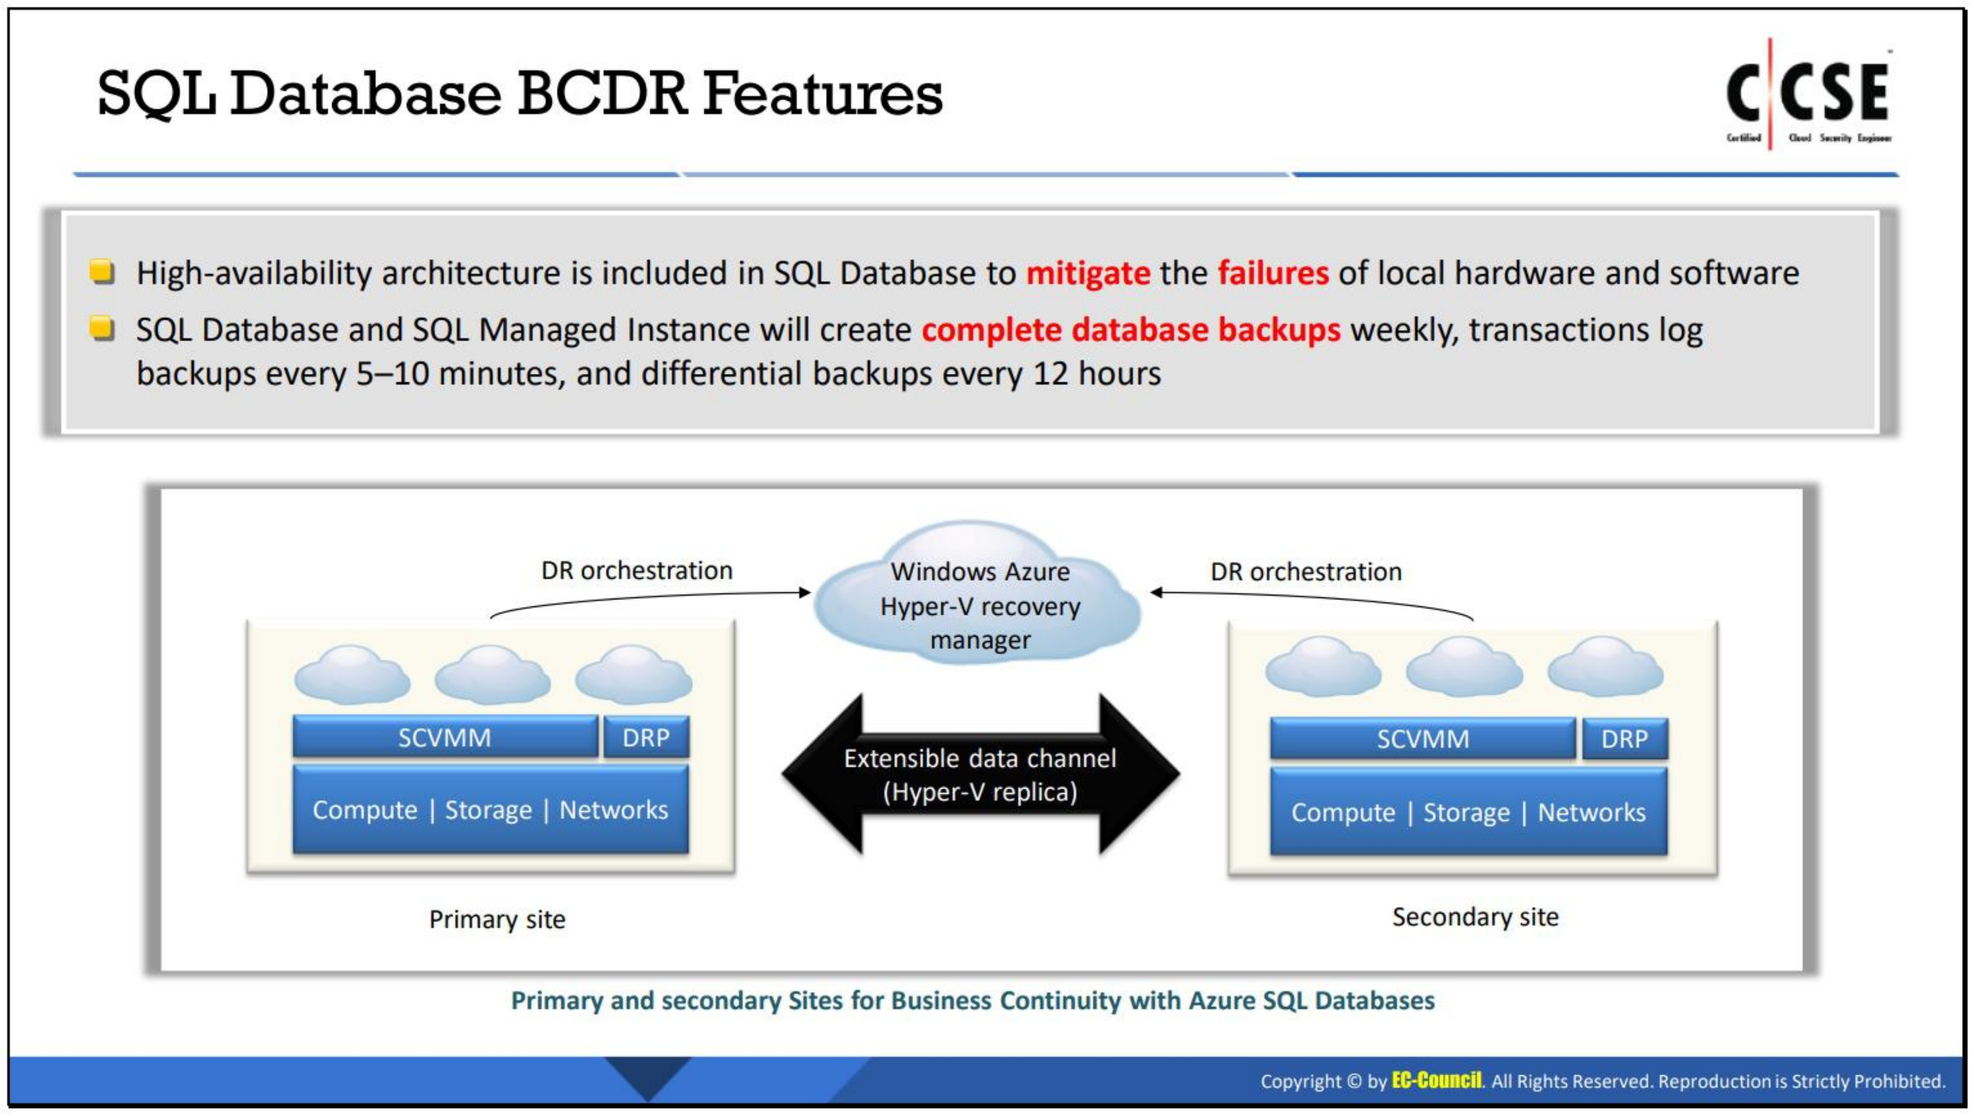

SQL Database BCDR Features

There are four major potential disruptions from the perspective of a database:

- Local failures including both hardware and software that affect database node such as disk-drive failure.
- Deletion or corruption of data usually occurred due to human error or application bug. These application-specific failures are generally not detected by the database service.
- Datacenter outage due to a natural disaster; redirecting to an alternate datacenter requires some level of geo-redundancy with application failover.
- Maintenance or upgrade errors; the occurrence of unanticipated issues during planned infrastructure maintenance or any upgrades that require rollback to a prior database state.

hide01.ir

HA architecture is included in SQL Database to mitigate the failures of local hardware and software. To overcome this failure, it guarantees automatic recovery with SLA of up to 99.995% availability. SQL Database and SQL Managed Instance create complete database backups weekly, transactions log backups every 5–10 minutes, and differential backups every 12 hours. For all service tiers, the backups are stored in RA-GRS for at least 7 days. The retention period for PITR for all service tiers, except basic configurable backup, is up to 35 days. To mitigate various unplanned scenarios, SQL Database and SQL Managed Instance will provide multiple BC features.

- Restoration of row versions from any point in time is possible using temporal tables.
- To restore the complete database to some point in time within 35 days of the configured retention period, “Built-in automated backups” and “Point in Time Restore” can be used.
- A deleted database can be restored to its last deleted stage if the server is not deleted.

Certified Cloud Security Engineer Copyright © by **EC-Council**
All Rights Reserved. Reproduction is Strictly Prohibited.

- A limited public preview for SQL Managed Instance provides long-term backup retention to store backups up to 10 years.
- In case of a data center outage, to create readable replicas and manually failover to any replica, active geo-replication provides this facility.
- For automatic recovery in case of a data center outage, the auto-failover group provides auto-recovery to the application.

hide01.ir

Certified Cloud Security Engineer Copyright © by **EC-Council**
All Rights Reserved. Reproduction is Strictly Prohibited.

Capabilities of Azure SQL Database to Recover from an Outage

Azure SQL Database offers various features to recover from an outage such as:

Azure Active Geo-Replication

- It helps create **readable secondary databases** on a server either in the same region or a different region

Azure Auto-Failover Groups

- Azure auto-failover groups support **automatic failover** of multiple databases after a disaster or any unplanned event

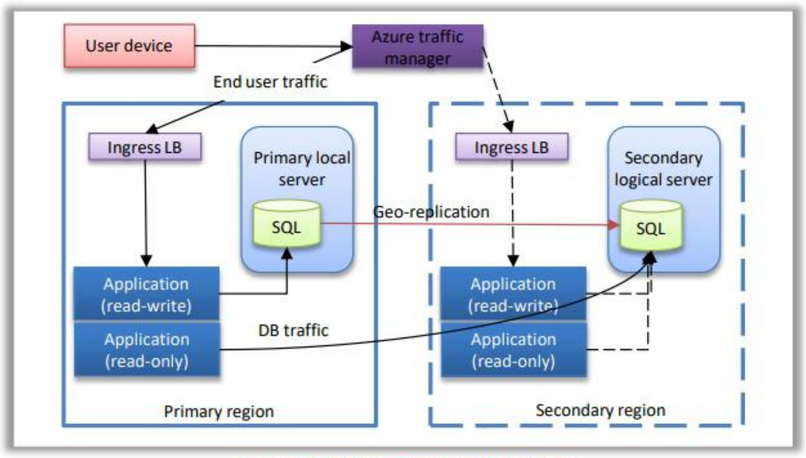

Azure Active Geo-replication

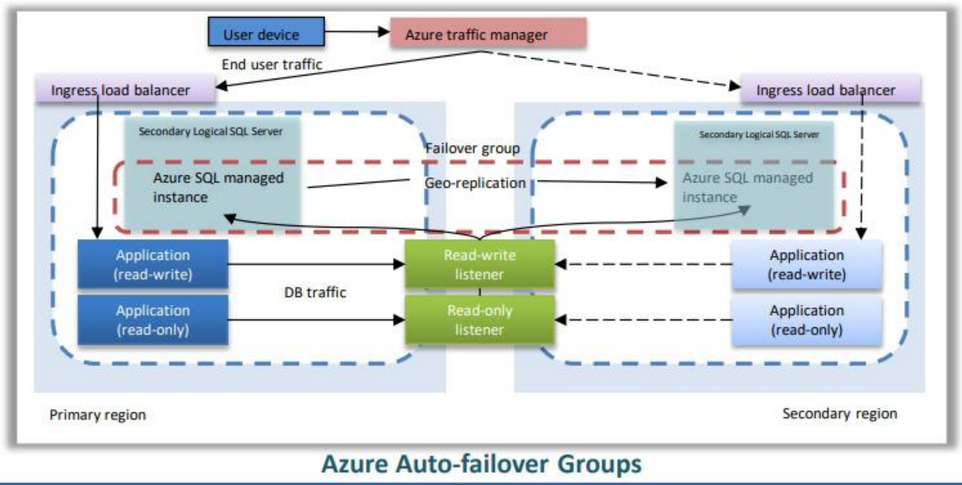

Azure Auto-failover Groups

Copyright © by EC-Council. All Rights Reserved. Reproduction is Strictly Prohibited.

Capabilities of Azure SQL Database to Recover from an Outage (Cont'd)

hide01.ir

Azure Geo-Restore

- The geo-restore feature allows restoring a database from a geo-redundant backup for **creating** a **new database**. It is automatically enabled for all the service tiers

Zone Redundant Configuration

- The zone redundant configuration makes use of the Azure availability zones to **replicate** the **database** in various physical locations in an Azure region

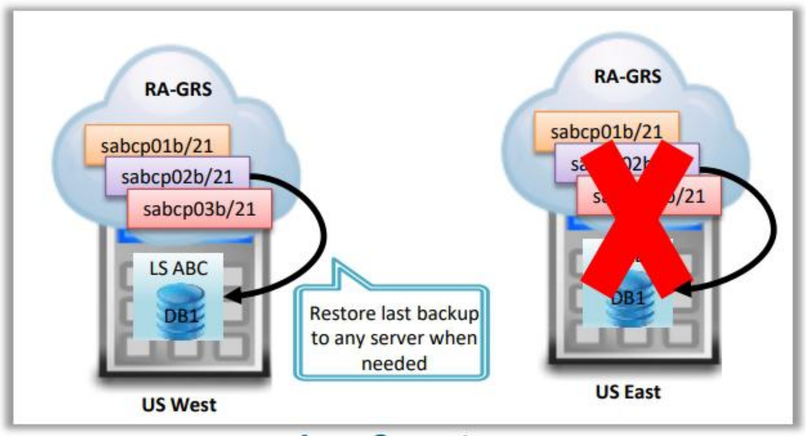

Azure Geo-restore

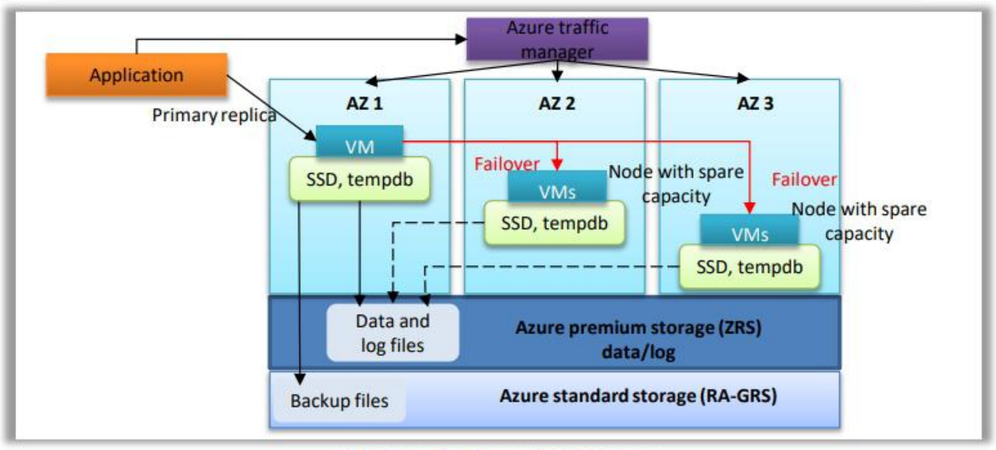

Zone-redundant Databases

Copyright © by EC-Council. All Rights Reserved. Reproduction is Strictly Prohibited.

Capabilities of Azure SQL Database to Recover from an Outage

Azure SQL Database offers various features to recover from an outage such as:

- Active geo-replication
- Auto-failover groups
- Geo-restore

Certified Cloud Security Engineer Copyright © by **EC-Council**
All Rights Reserved. Reproduction is Strictly Prohibited.

- Zone-redundant databases

It is essential to understand each feature to make use of them in the event of an outage.

Active Geo-Replication

The active geo-replication feature helps create readable secondary databases on a server either in the same region or a different region. It is mainly designed to provide a solution for BC. Active geo-replication allows an application to recover from a disaster of individual databases due to a regional disaster or an outage.

By activating geo-replication, the application initiates a failover to a secondary database during an outage of the primary database for a long time. A maximum of four secondary databases are supported in a region and they can be used for read-only access queries as well. However, the failover needs to be manually initiated by the user or the application. Once the failover is successful, the new primary database will have a new endpoint connection.

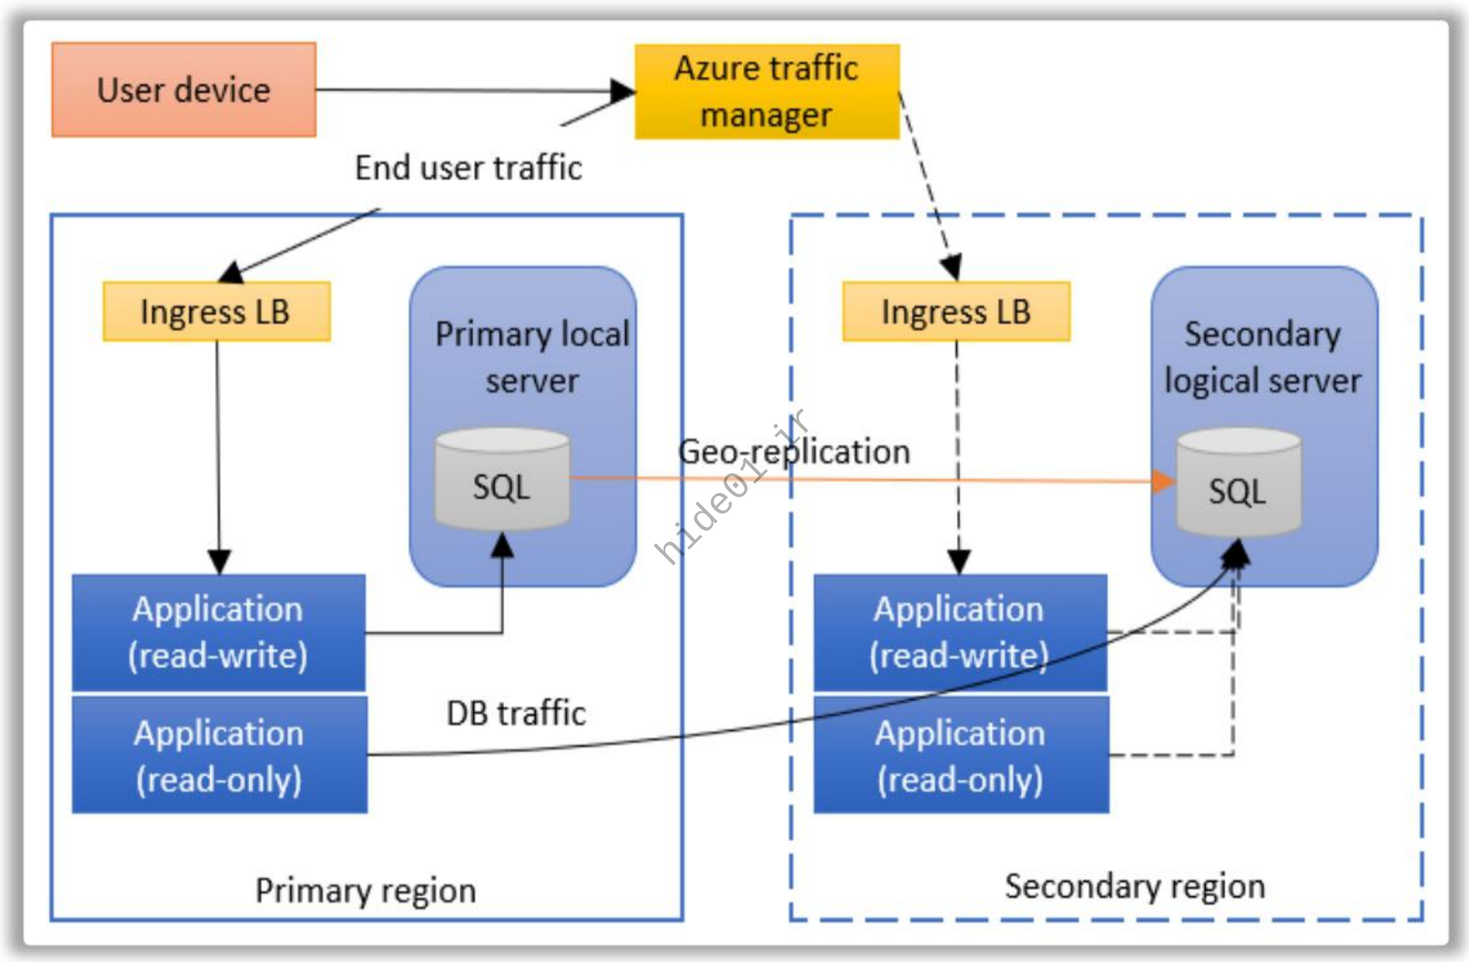

Figure 9.126: Active Geo-Replication

If the primary database fails due to any reason or it needs to be made offline, the failover can be initiated to any of the secondary databases. After activating the failover to a secondary database, all the other secondary databases automatically get connected to the new primary database.

Active geo-replication can be used to manage replication and failover of individual or a set of databases through

- The Azure portal
- REST API: Single database
- Transact-SQL: Single database or elastic pool
- PowerShell: Elastic pool
- PowerShell: Single database

Certified Cloud Security Engineer Copyright © by **EC-Council**
All Rights Reserved. Reproduction is Strictly Prohibited.

Geo-replication asynchronously replicates committed transactions on the primary database to a secondary database using snapshot isolation. The secondary databases can also be used to offload read-only workloads.

Apart from DR, active geo-replication is also useful for the following:

- A database can be migrated from one server to another online server with minimal downtime.
- Extra secondary databases can be created as a failure back copy during application upgrades.

If more than four secondary databases are required for a primary database, geo-replication can be configured on a secondary database. This lets the secondary database itself to have four more secondary databases.

Auto-Failover Groups

Geo-replication does not allow simultaneous replication and failover of multiple databases. The failover needs to be performed manually, which could lead to data loss. These drawbacks can be addressed with auto-failover groups.

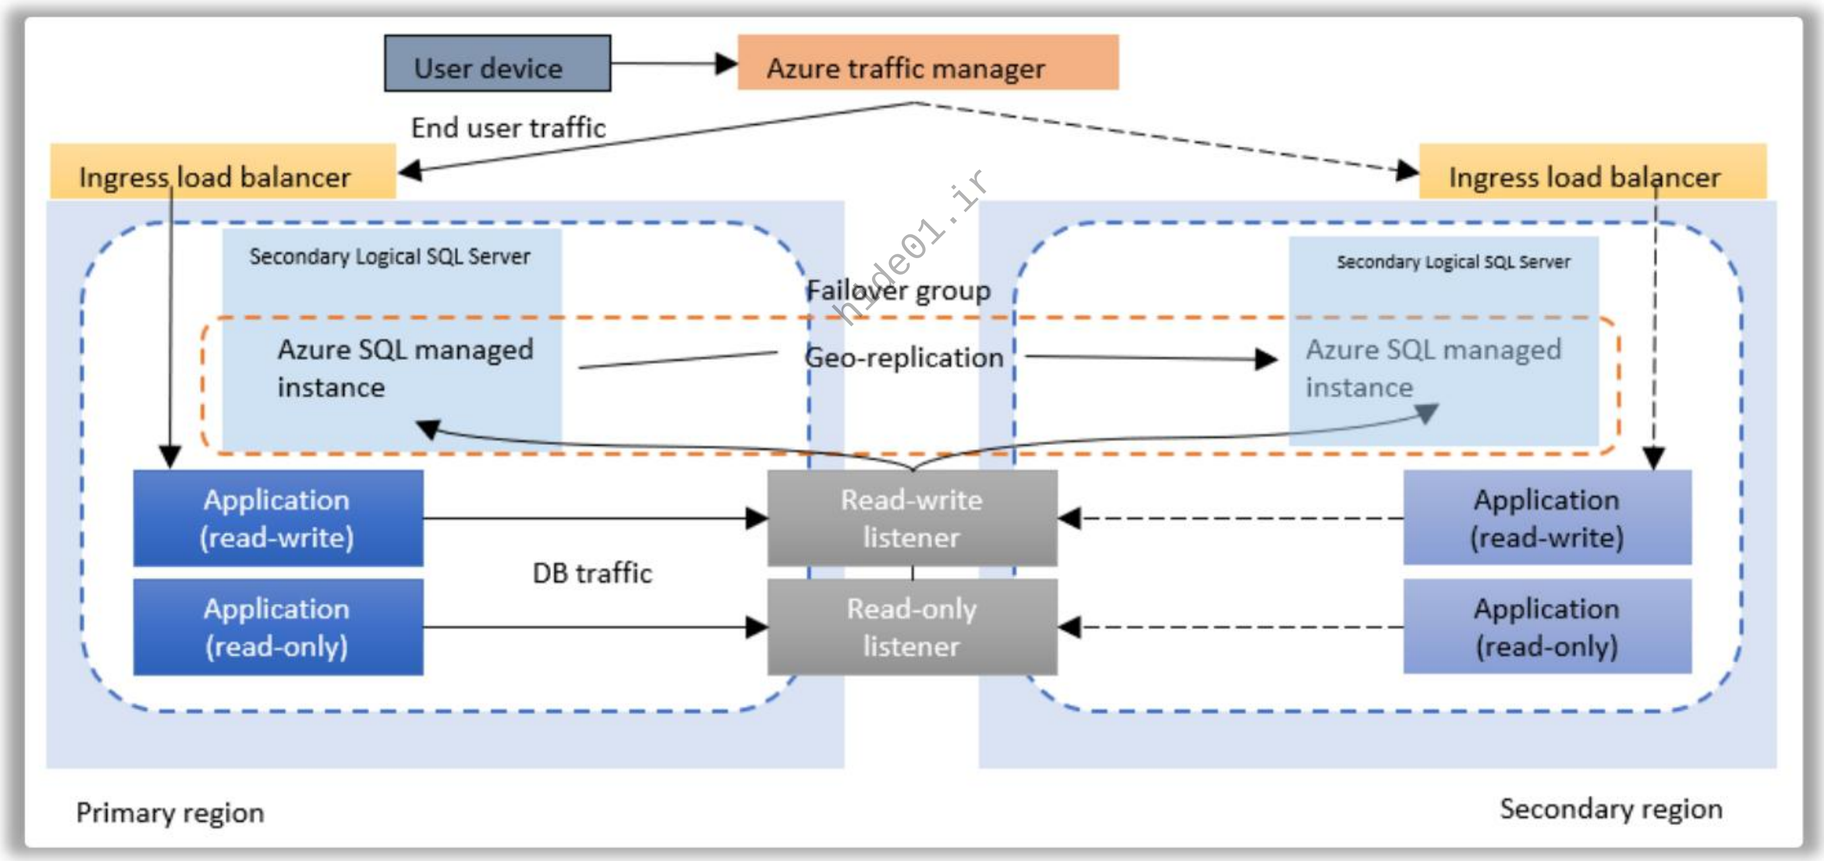

Figure 9.127: Auto-Failover Groups

Auto-failover groups support automatic failover of multiple databases after a disaster or any unplanned event. An auto-failover group can have one or more databases. All the databases in the failover group must be configured on a primary server. The databases in an auto-failover group can be replicated to only one secondary server or an instance in another region.

When auto-failover groups are used along with an automatic failover policy, an outage that affects the databases in the group leads to automatic failover. The auto-failure groups support read-write and read-only listener endpoints that remain the same during failovers. The read-write listener is configured to automatically point to the primary server in the group. The read-only listener is configured to point to the secondary server in the group. The automatic failover

Certified Cloud Security Engineer Copyright © by **EC-Council**
All Rights Reserved. Reproduction is Strictly Prohibited.

policy controls the automatic failover of the read-write listener after a primary server failure. The read-only policy is controls the automatic failover of the read-only listener following a failure of the secondary server.

Whether failover activation is performed manually or automated, failover switches all the secondary databases in the group. After the failover is completed, the DNS record is updated to redirect the endpoints to the new region.

Geo-restore

Geo-store is a cost-effective disaster recovery solution for the Azure SQL Database. The geo-restore feature allows restoring a database from a geo-redundant backup for creating a new database. It is automatically enabled for all the service tiers. It can be created on any server and in any Azure region. It relies on the geo-replicated backup as its source. Geo-restore is suitable for applications that use small databases.

Geo-restore can be used even when the database is inaccessible due to an outage. It uses the same technology as point-in-time restore but with a major difference. It restores the database using a copy of the daily backup in geo-replicated blob storage. The service maintains a backup chain for every active database. The backup chain includes weekly full backup and transaction logs that are saved for every 5 minutes. The blobs are geo-replicated, which makes the backup available even after a failure in the primary region.

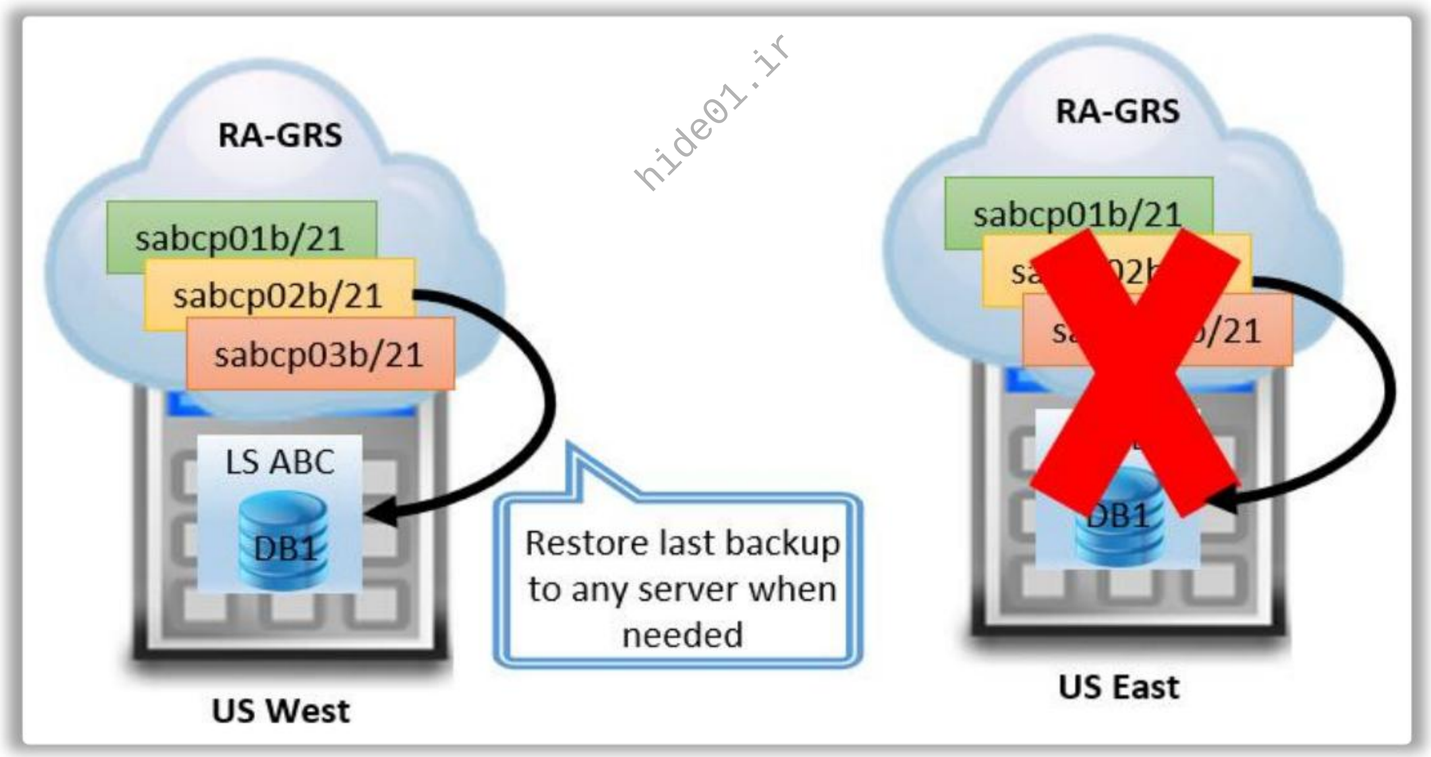

Figure 9.128: Geo-restore

If an incident in a region makes the database application unavailable, geo-restore can be used to restore a database from the recent backup in another region. Though the backups are geo-replicated, there could be some delay between the time when the backup was taken and geo-replication to Azure Blob in the different regions. For this reason, the restored database can be an hour behind the original database.

The recovery time is based on the size of the database, the number of parallel restore requests processed in that region, and the performance of the database. If the outage in a region is for a long time, many geo-restore requests will be processed by other regions. The service will limit the resources used by the restore operation to reduce the effect on workloads in that region. The

Certified Cloud Security Engineer Copyright © by **EC-Council**
All Rights Reserved. Reproduction is Strictly Prohibited.

recovery time for the database in that region may increase if there are a large number of requests.

Zone-redundant Databases

The zone-redundant configuration makes use of the Azure AZs to replicate the database in various physical locations in an Azure region. It can be applied to both the existing and new general-purpose databases and elastic pools. Zone-redundant configuration makes the elastic pool and general-purpose database resilient to a large number of failures.

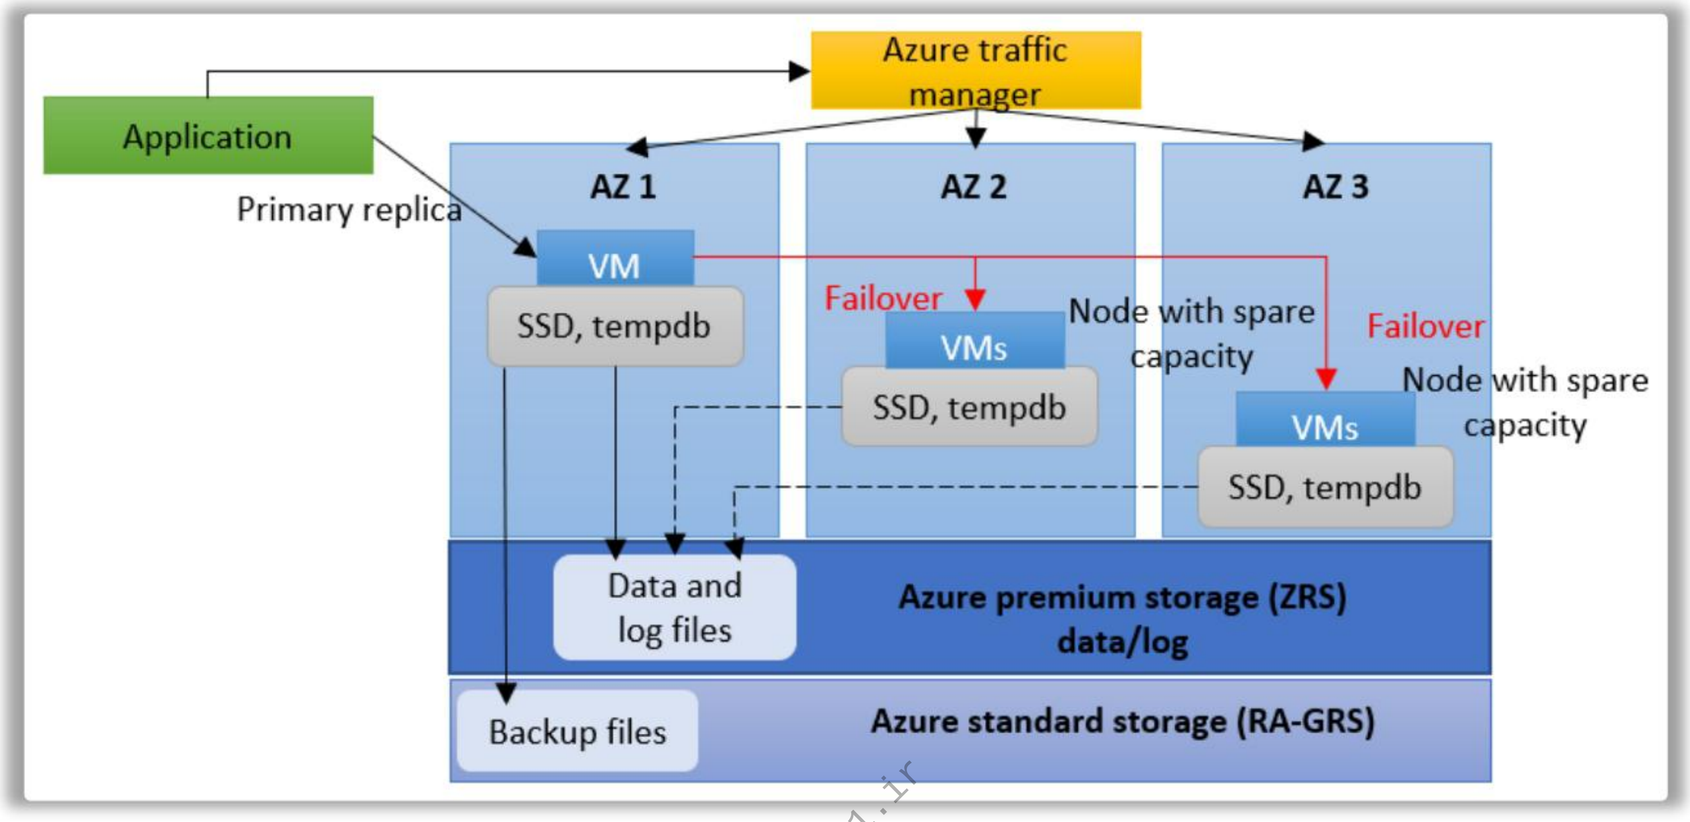

Figure 9.129: Zone-redundant Database

When the zone-redundant option is enabled, the Azure SQL Database automatically reconfigures the database, without any downtime. This configuration can be enabled from Azure portal, PowerShell, CLI, and Azure Resource Manager API.

The zone-redundant general-purpose database and elastic pools are available in East US, East US 2, West US 2, North Europe, West Europe, Southeast Asia, and Australia East.

Certified Cloud Security Engineer Copyright © by **EC-Council**
All Rights Reserved. Reproduction is Strictly Prohibited.

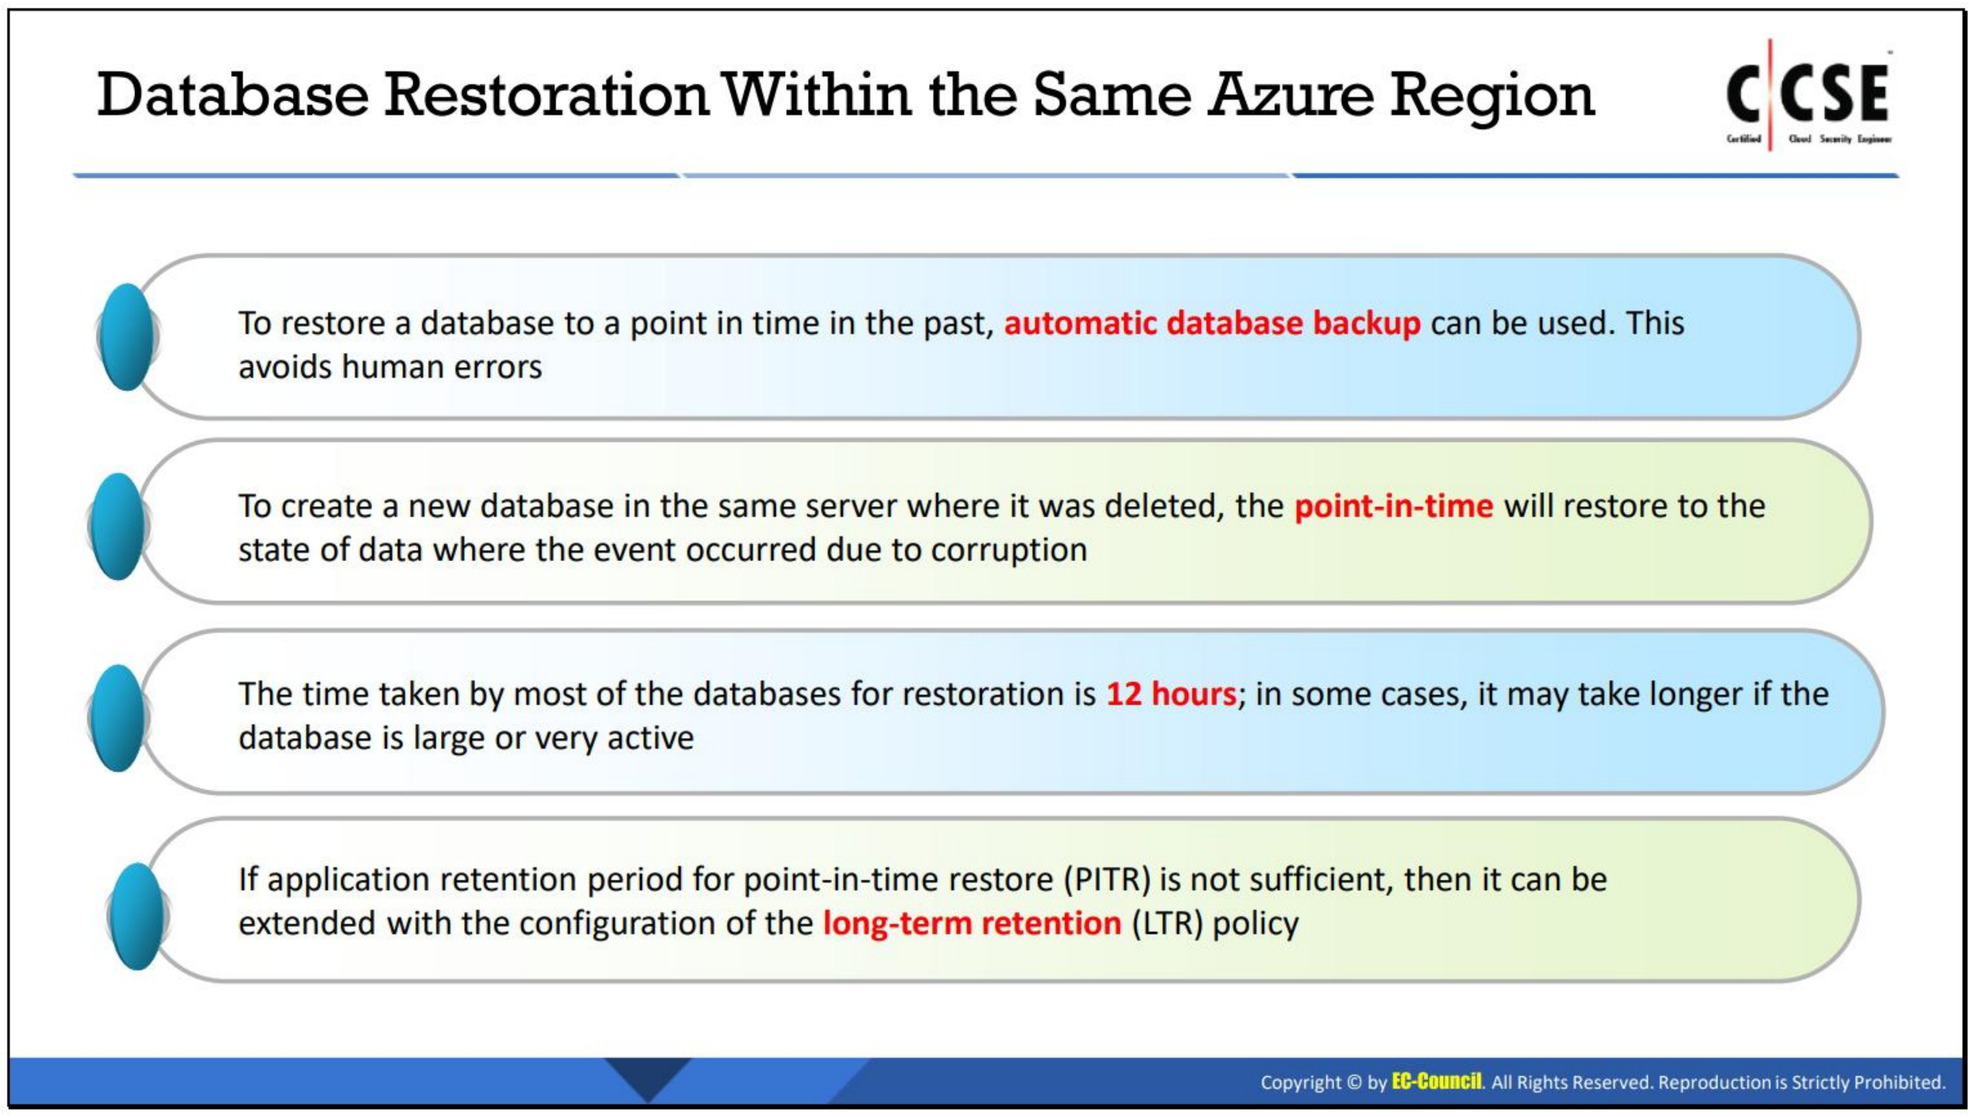

Database Restoration Within the Same Azure Region

To restore a database to a point in time in the past, automatic database backup can be used. Data corruption due to human error can be avoided using this way. To create a new database in the same server where it was deleted, the point-in-time will restore to the state of data where the event occurred due to corruption.

The time taken by most of the databases for restoration is 12 hours; however, in some cases, it may take a longer time if the database is large or very active. If your application retention period for PITR is not sufficient, then it can be extended with the configuration of the long-term retention (LTR) policy.

Certified Cloud Security Engineer Copyright © by **EC-Council**
All Rights Reserved. Reproduction is Strictly Prohibited.

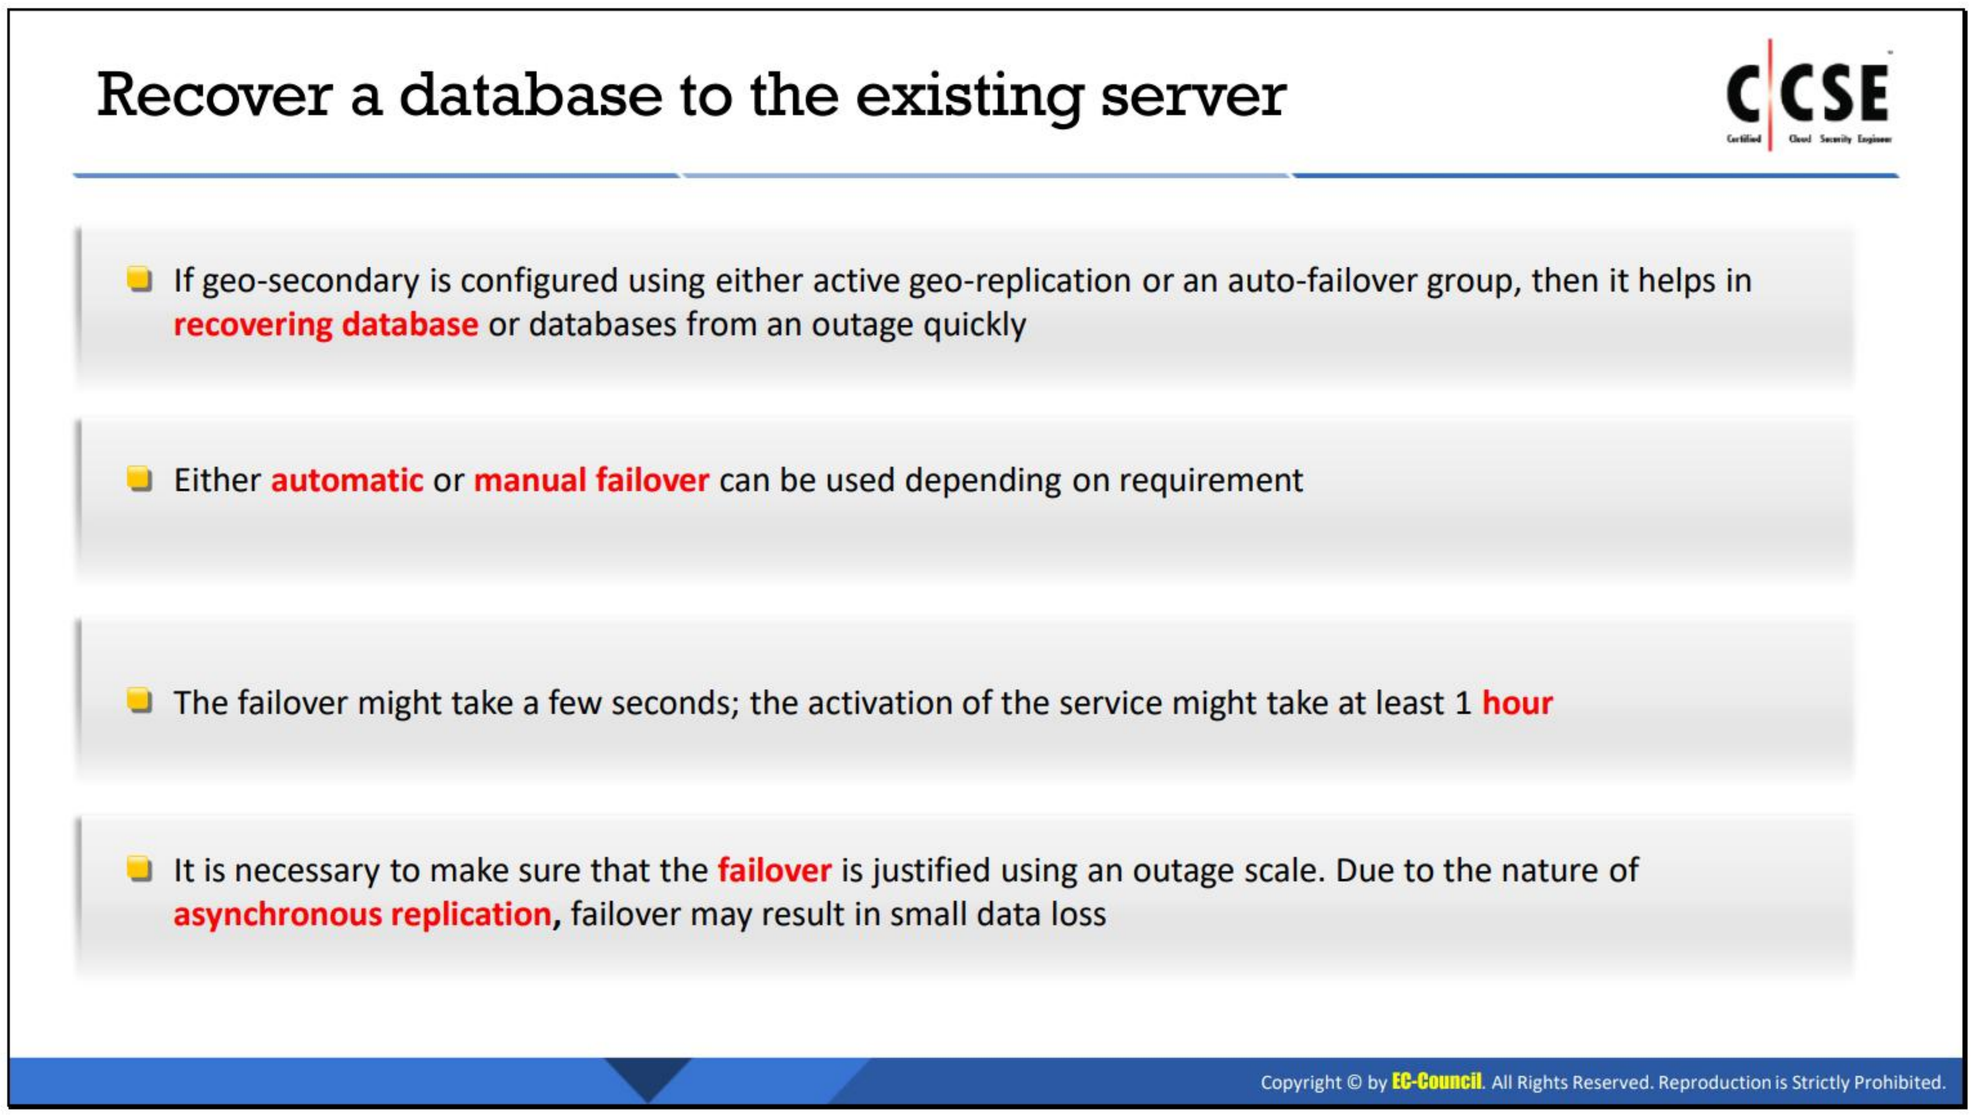

Recover a Database to the Existing Server

Although it is a rare scenario, even an Azure datacenter may face an outage. Such an outage disrupts the business, which might last few minutes or hours depending on the situation.

- The foremost simple way is to wait for the database to get online immediately after the outage is completed. It works for applications that can afford an offline database. For example, a development project is not required to work constantly. One may not know the time taken for outage clearance; thus, it works only if the database is not required for a while.
- Alternatively, use Azure geo-redundant database backup for restoration (geo-restore). Geo-restore allows restoring from geo-redundant backups in any database or datacenter outage scenario.
- Configuring geo-secondary using either active geo-replication or an auto-failover group helps in recovering database or databases from an outage quickly. Either automatic or manual failover can be used depending on your choice. While the failover might take a few seconds, the activation of the service might take at least 1 hour. It is necessary to make sure that the failover is justified using an outage scale. Due to the nature of asynchronous replication, the failover may result in small data loss.

Prepare for an Outage

The following items are necessary irrespective of the BC feature being used:

- Create a target server after identification that includes logins, server-level IP firewall rules, and master database-level permissions.
- Establish a plan and rules to redirect the clients and their applications to the new server.

Certified Cloud Security Engineer Copyright © by **EC-Council**
All Rights Reserved. Reproduction is Strictly Prohibited.

- Other dependencies such as audit settings and alerts should be documented.

If these are not tested at an earlier stage, they may take more time for configuration and troubleshooting at the time of failover, which can lead to a stressful situation for the administrators.

hide01.ir

Certified Cloud Security Engineer Copyright © by **EC-Council**
All Rights Reserved. Reproduction is Strictly Prohibited.

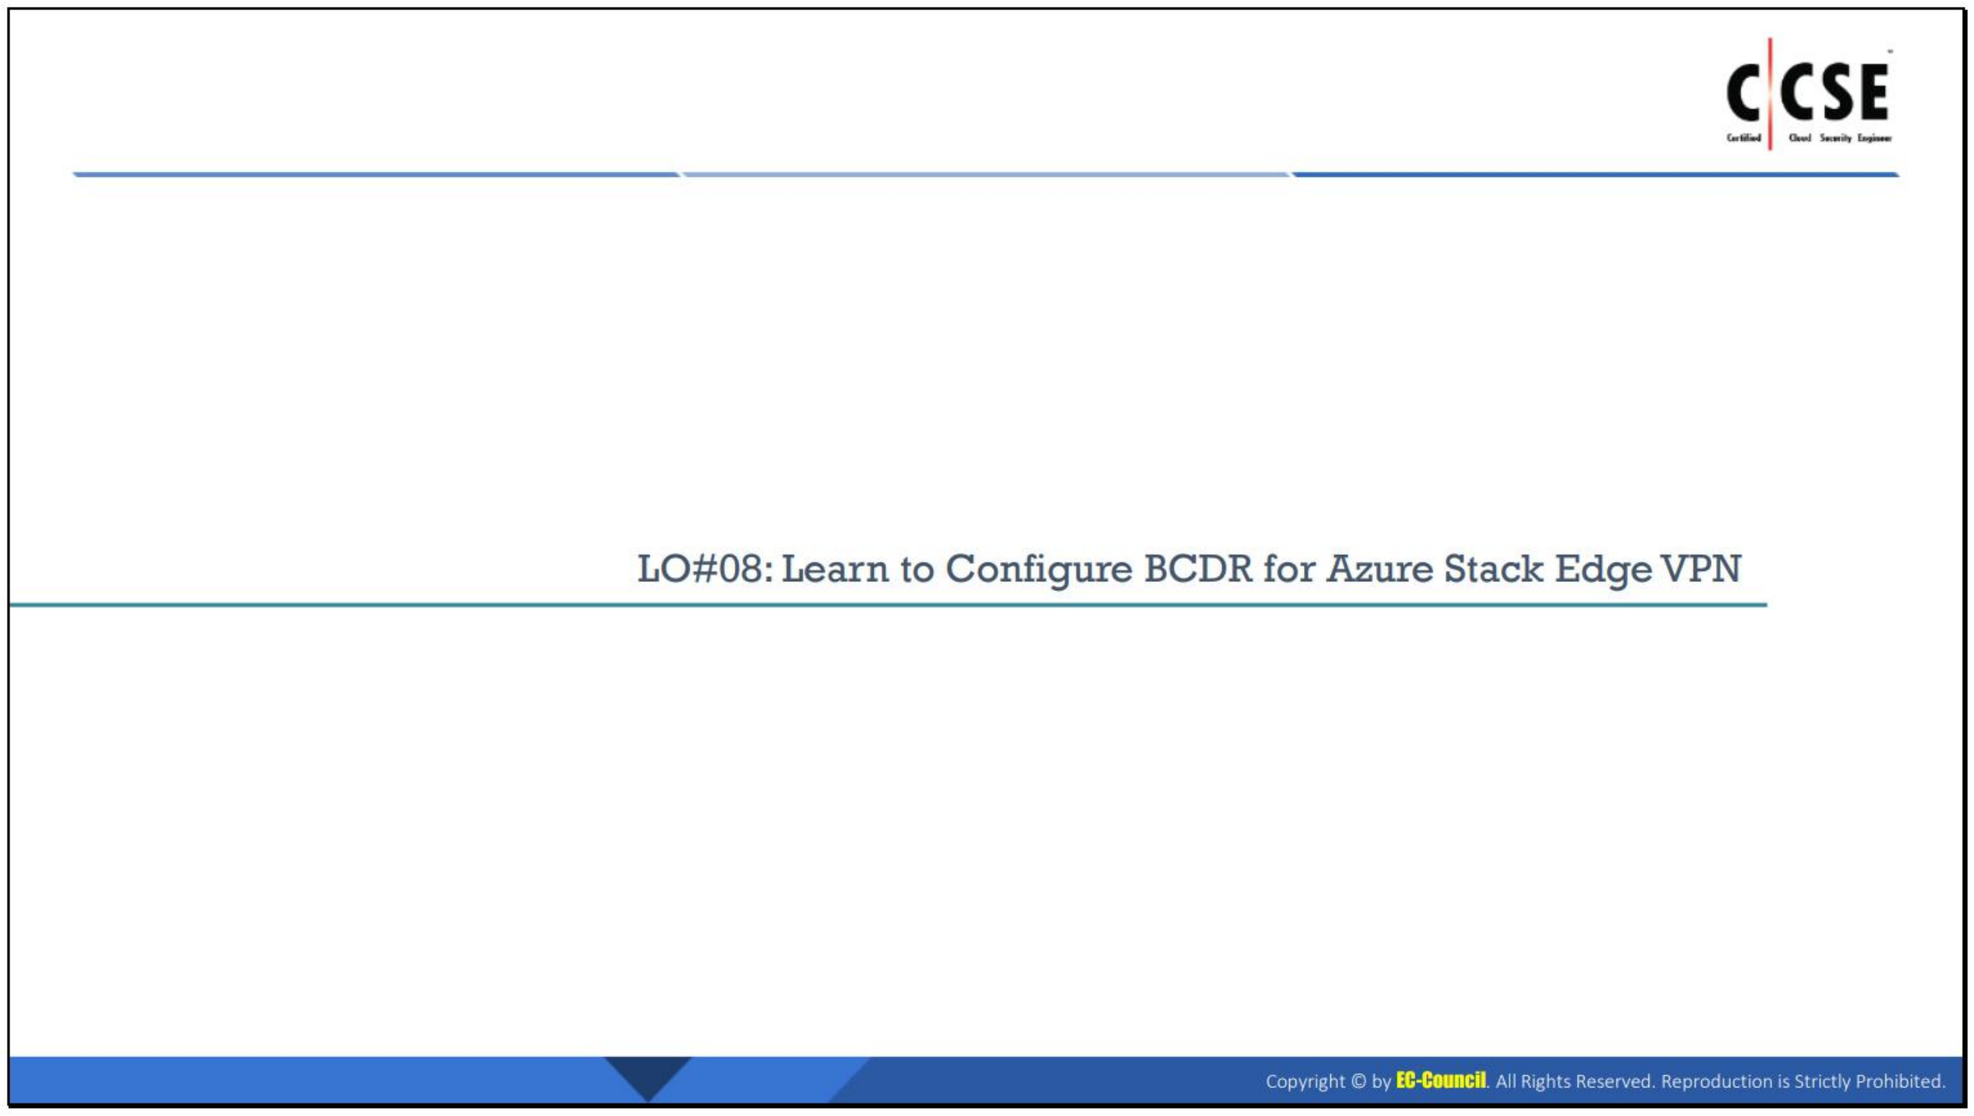

LO#08: Learn to Configure BCDR for Azure Stack Edge VPN

The objective of this section is to understand how to configure BCDR for Azure Stack Edge VPN. By understanding this section, cloud security engineers can know how to configure failover to a paired region as well as how to recover from a failed region.

hide1.ir

Certified Cloud Security Engineer Copyright © by **EC-Council**
All Rights Reserved. Reproduction is Strictly Prohibited.

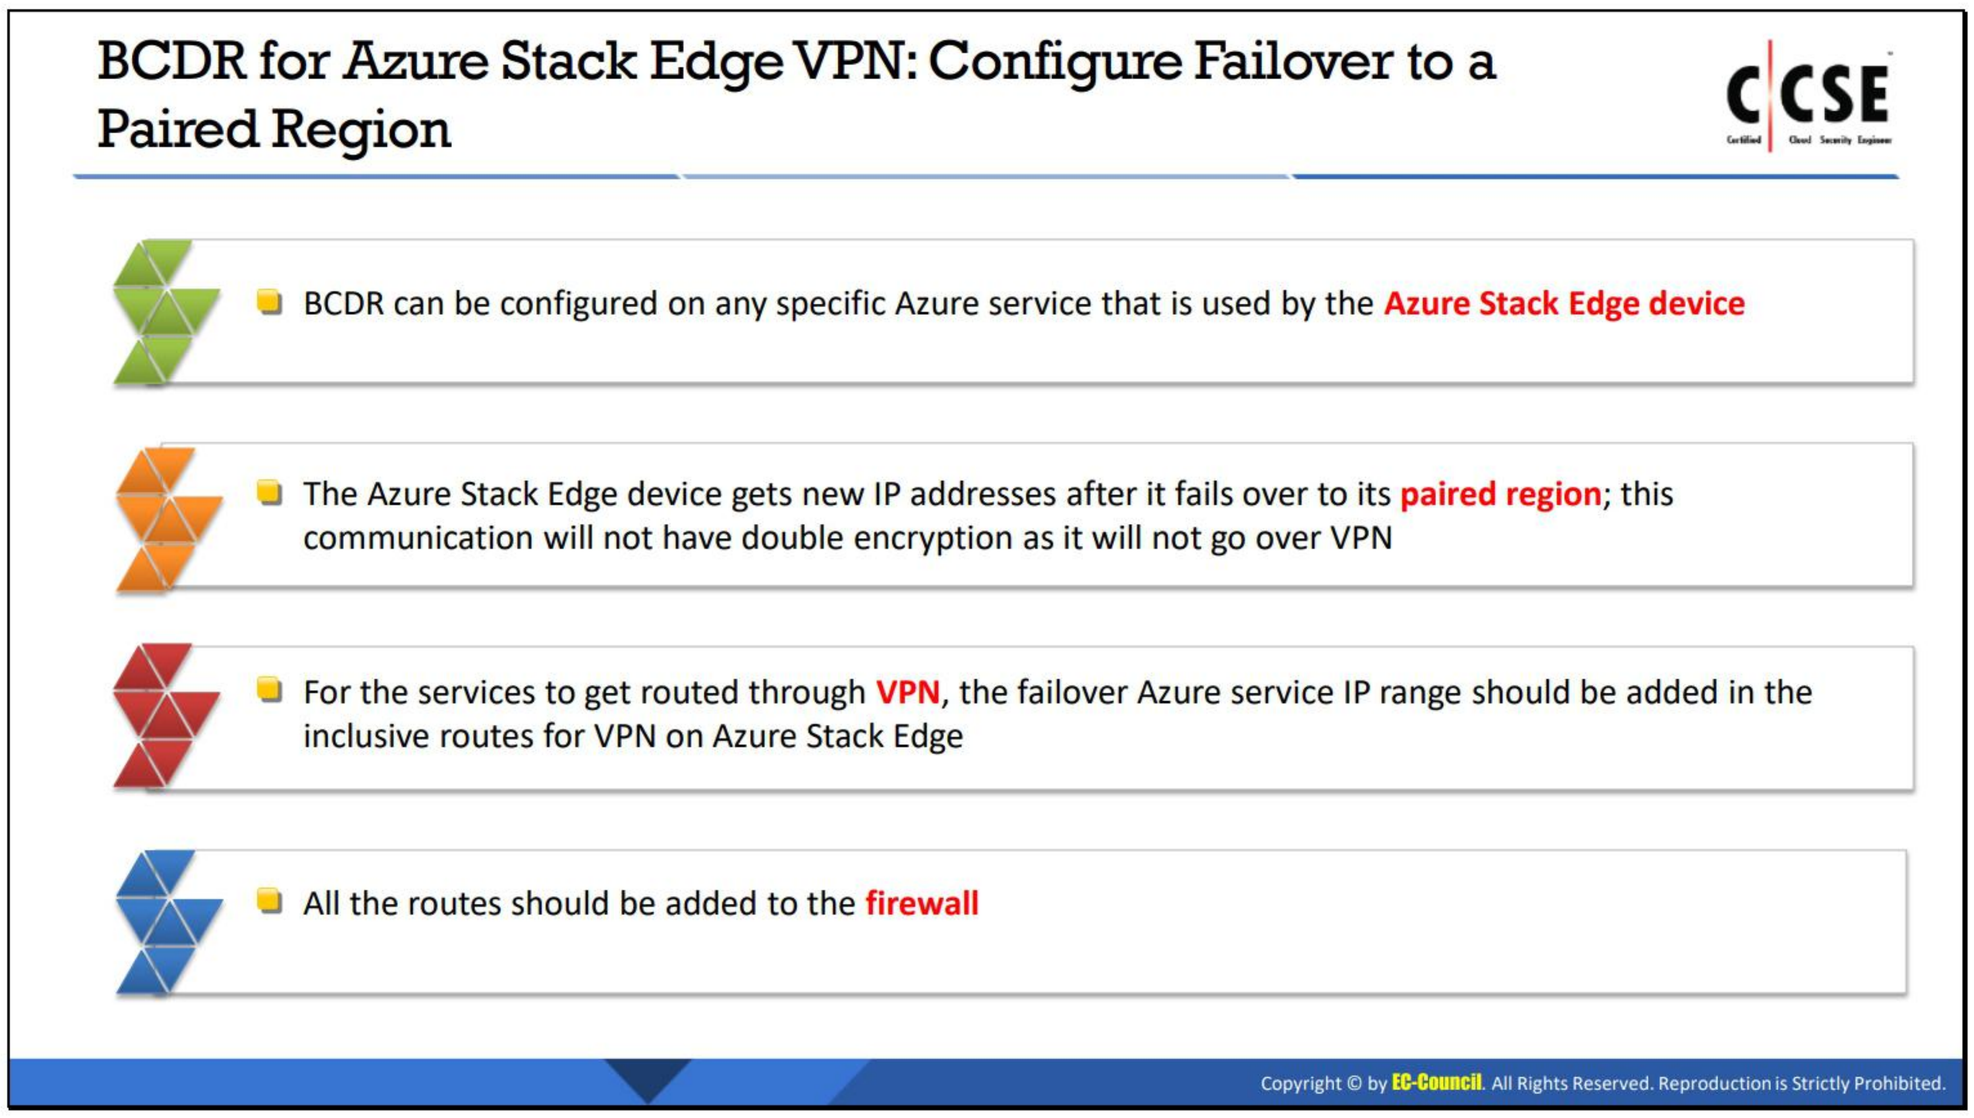

BCDR for Azure Stack Edge VPN: Configure Failover to a Paired Region

Azure Stack Edge devices use other Azure services such as Azure Storage. BCDR can be configured on any specific Azure service that is used by the Azure Stack Edge device. The Azure Stack Edge device gets new IP addresses after it fails over to its paired region; thus, this communication will not have double encryption. The data and services configured in the home region go over the VPN tunnel as the Azure Stack Edge device uses split tunneling; here, the home region is the region that is associated with the Azure Stack Edge device. The data will not go over VPN if the Azure services fail over to outside of the home region. Thus, this is not doubly encrypted.

The impact of this issue reflects on only a few Azure services. The below changes should be made in Azure Stack Edge VPN configuration to resolve the encryption issue:

- For the services to get routed through VPN, the failover Azure service IP range should be added in the inclusive routes for VPN on Azure Stack Edge.

  A JSON file with service-specific routes should be downloaded for the addition of inclusive routes.
- In the Azure Route table, the corresponding Azure service IP ranges should be added.
- All the routes should be added to the firewall.

Certified Cloud Security Engineer Copyright © by **EC-Council**
All Rights Reserved. Reproduction is Strictly Prohibited.

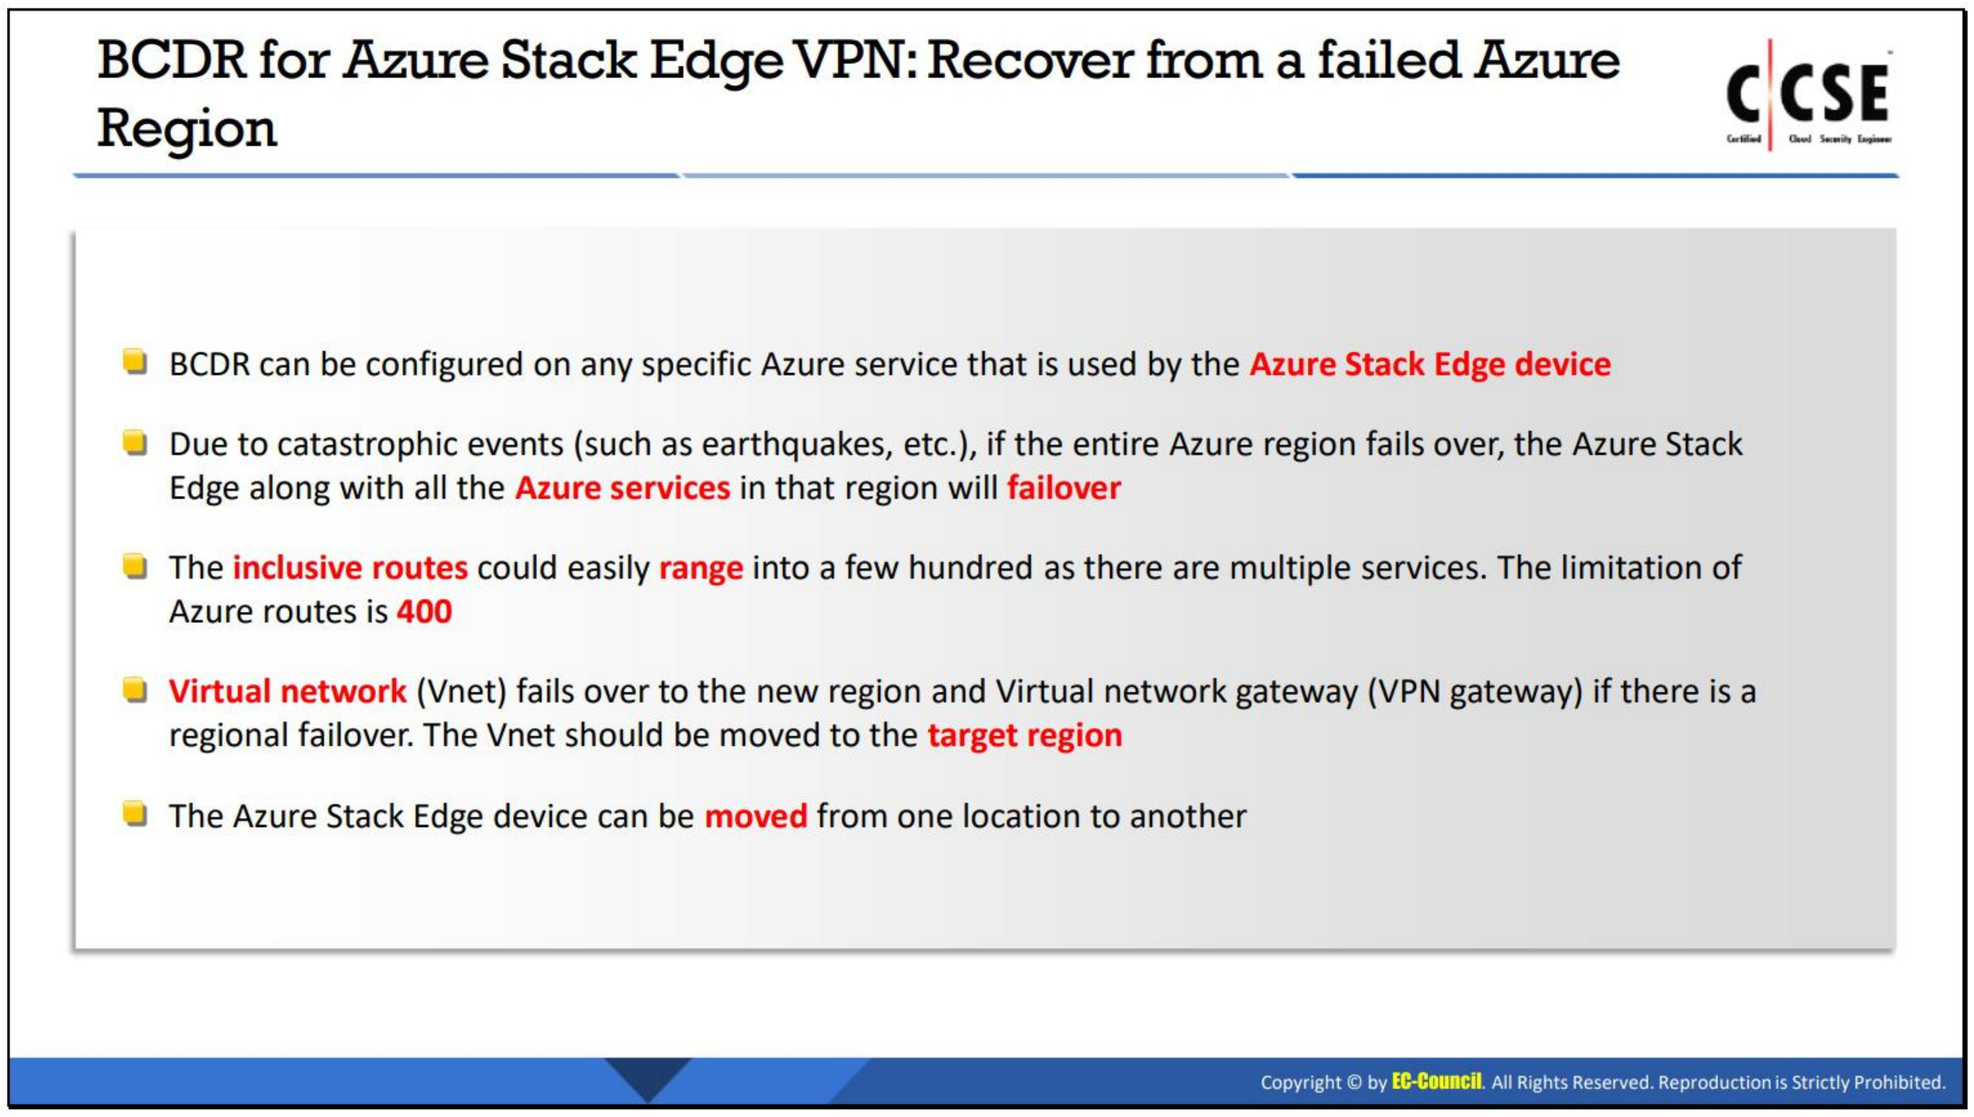

BCDR for Azure Stack Edge VPN: Recover from a Failed Azure Region

In case of catastrophic events (such as earthquakes, etc.), if an entire Azure region fails over, the Azure Stack Edge along with all the Azure services in that region will failover. The inclusive routes could easily range into a few hundred as there are multiple services. The limitation of Azure routes is 400. Virtual network (Vnet) fails over to the new region and virtual network gateway (VPN gateway) if there is a regional failover. Make the below changes in the Azure Stack Edge VPN configuration to resolve this issue.

- The Vnet should be moved to the target region.
- New Azure VPN gateway should be deployed in the target region (the region where Vnet is moved).
- To use the above VPN gateway in VPN connection, the Azure Stack Edge VPN configuration should be updated. To add routes that use a VPN gateway, the target region should be selected.
- If there is any change in the client address pool, the incoming Azure route table needs to be updated.

To Move the Azure Region from One to Another

The Azure Stack Edge device can be moved from one location to another. You need to configure the device for a new home region to use the nearest region to the deployed device. The following changes need to be changed:

- To use the new region's VPN gateway, update the configuration of the Azure Stack Edge VPN; for the addition of routes that use VPN gateways, new region should be selected.

Certified Cloud Security Engineer Copyright © by **EC-Council**
All Rights Reserved. Reproduction is Strictly Prohibited.

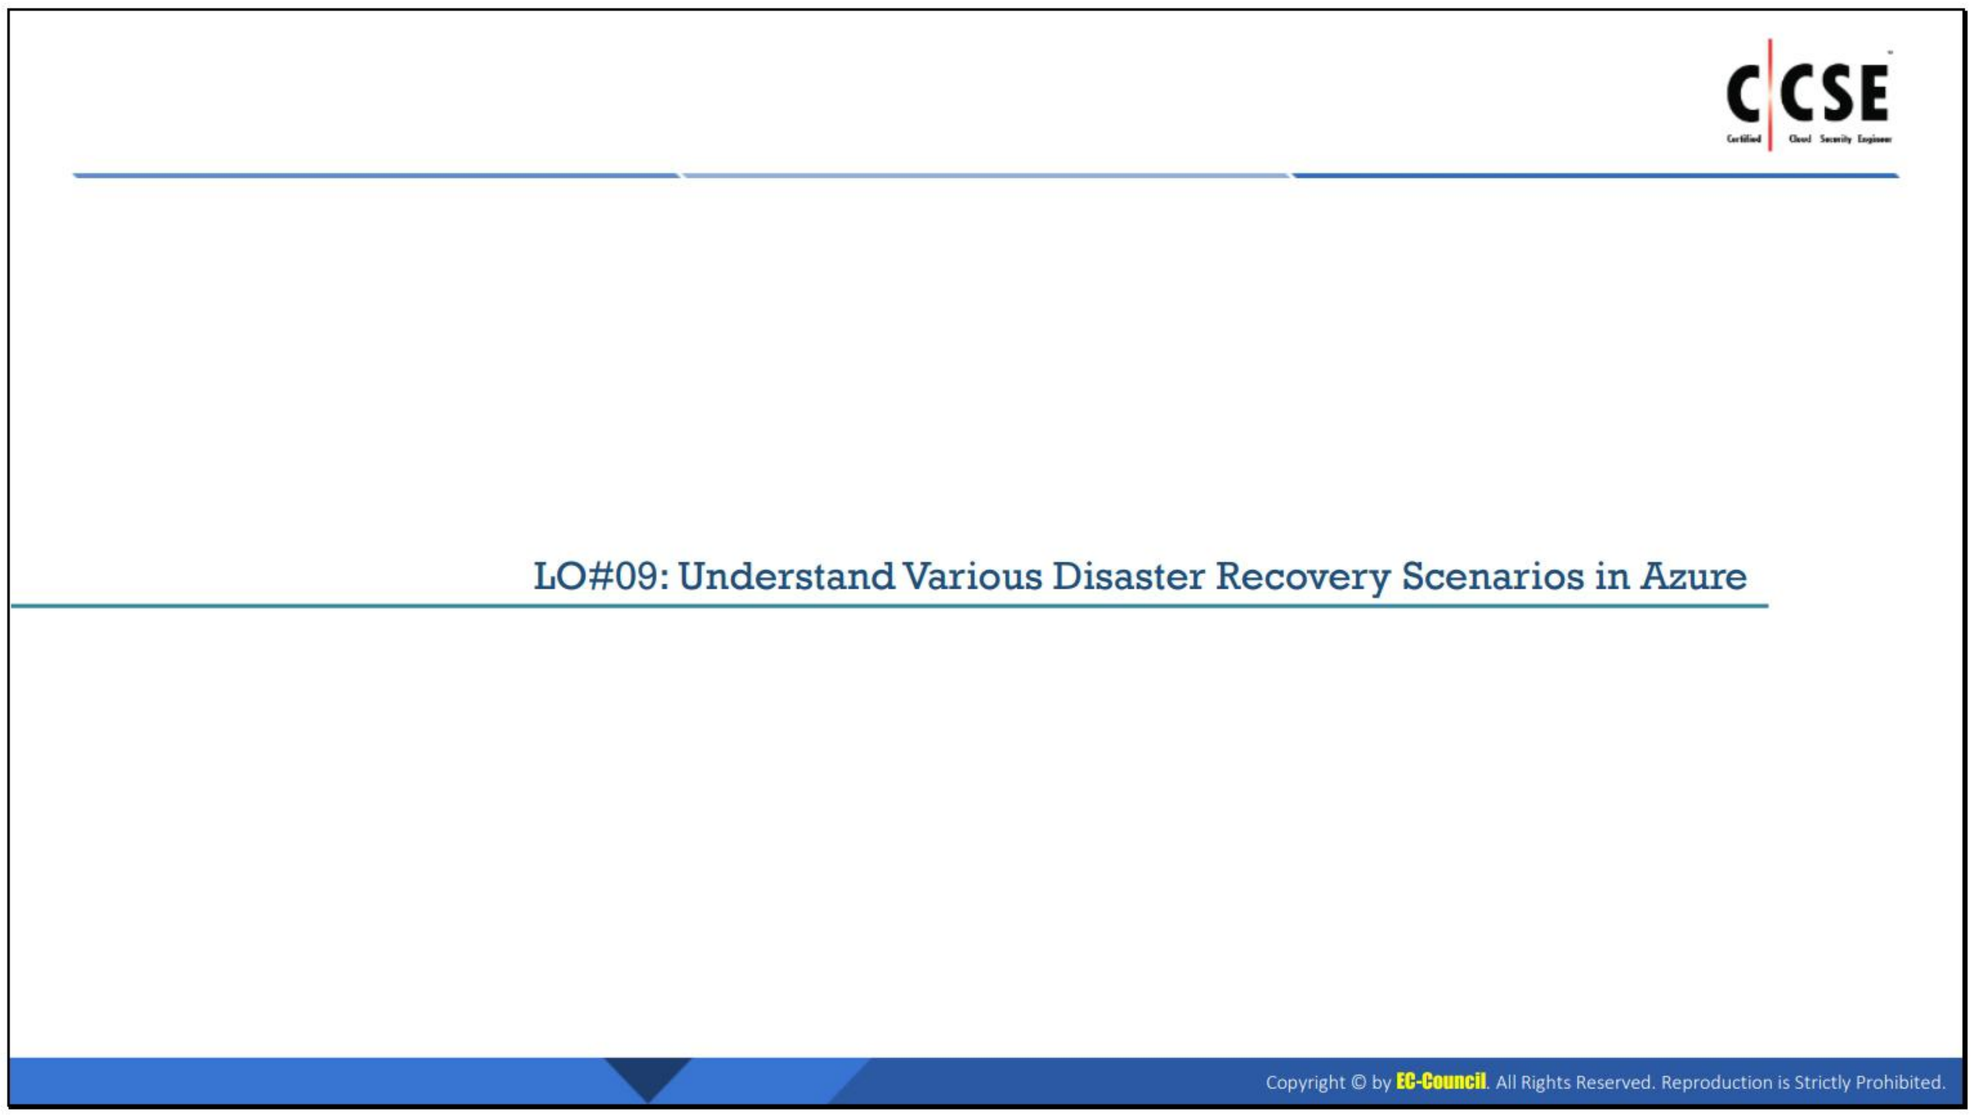

LO#09: Understand Various Disaster Recovery Scenarios in Azure

The objective of this section is to understand the different DR scenarios supported by Azure. Cloud security engineers should gain complete knowledge on the various disaster recovery scenarios in Azure to prepare their organization before a disaster event. This section includes important topics such as various DR scenarios, enterprise-scale BCDR, best practices, etc.

Certified Cloud Security Engineer Copyright © by **EC-Council**
All Rights Reserved. Reproduction is Strictly Prohibited.

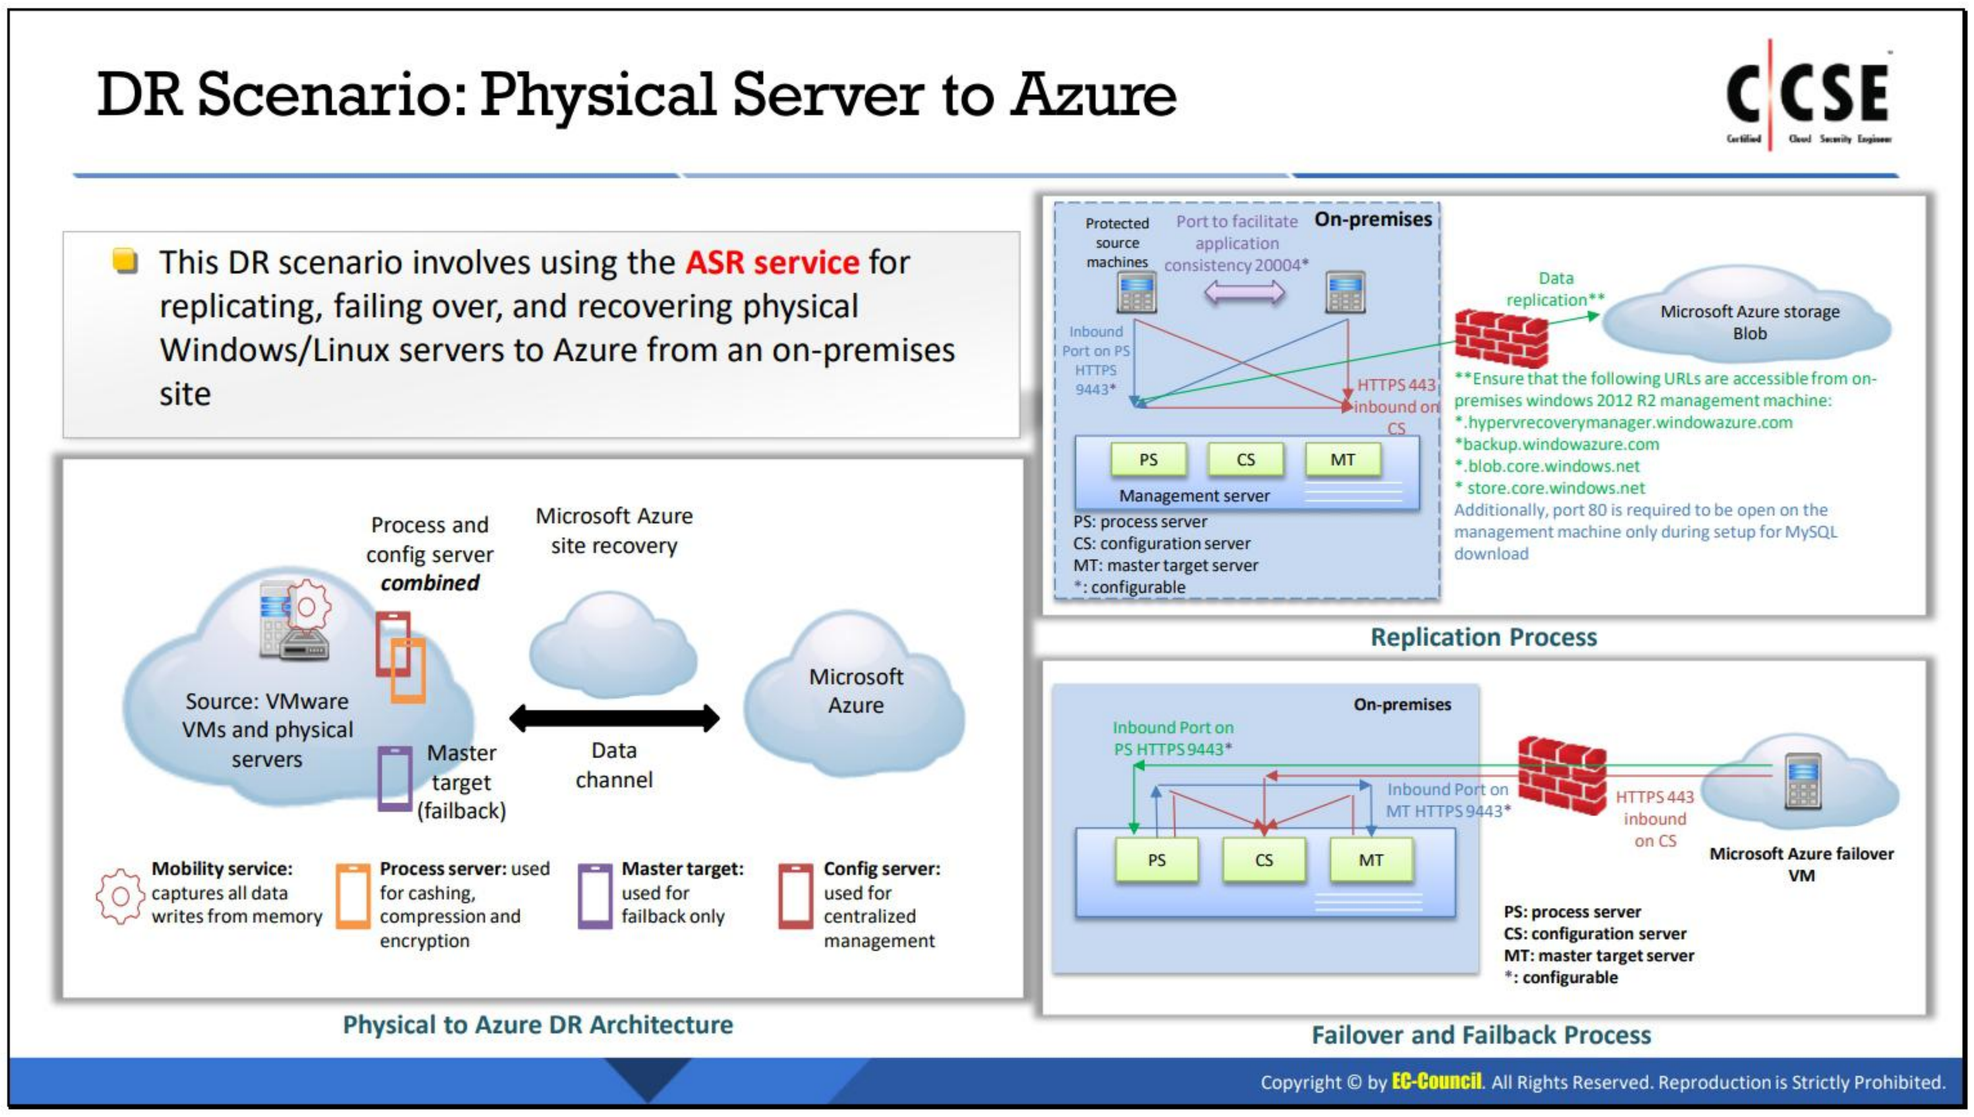

DR Scenario: Physical Server to Azure

This DR scenario involves using the ASR service for replicating, failing over, and recovering physical Windows/Linux servers to Azure from an on-premises site.

hide01.ir

Architectural Components

- **Azure**: Replicated data from on-premises physical servers is stored in Azure managed disks. This replicated data is used to create VMs when the customer performs a failover from on-premises to Azure.
- **Process server**: It installs the mobility service on the servers that are to be replicated and acts as a replication gateway. It receives replication data from servers, optimizes it with caching, encryption and compression, and sends it to Azure.
- **Master target server**: During failback from Azure, the master target server handles the replication data.
- **Replicated servers**: On every server that is to be replicated, the mobility service is installed. Customers can install the service via an automated deployment method or manually.

Outbound Network Connectivity

The customer needs to modify outbound network connectivity to enable their environment to replicate data to Azure. Certain URLs need to be allowed if a URL-based firewall proxy is used to control outbound network connectivity to ensure Site Recovery works as expected.

Certified Cloud Security Engineer Copyright © by **EC-Council**
All Rights Reserved. Reproduction is Strictly Prohibited.

The following image shows the physical-to-Azure DR architecture.

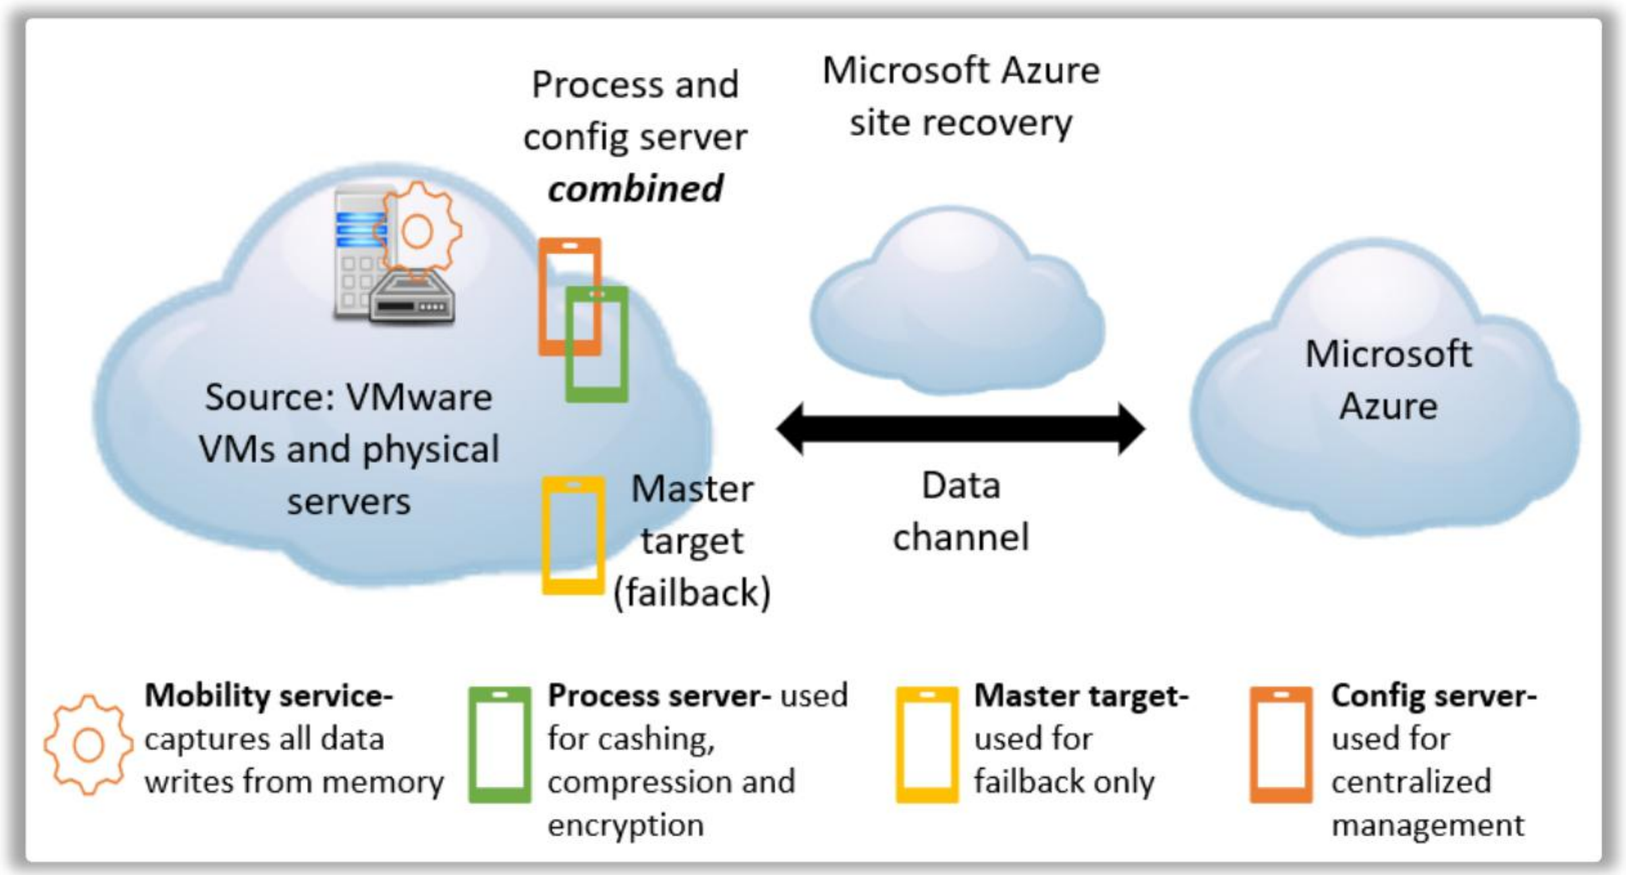

Figure 9.130: Physical to Azure DR Architecture

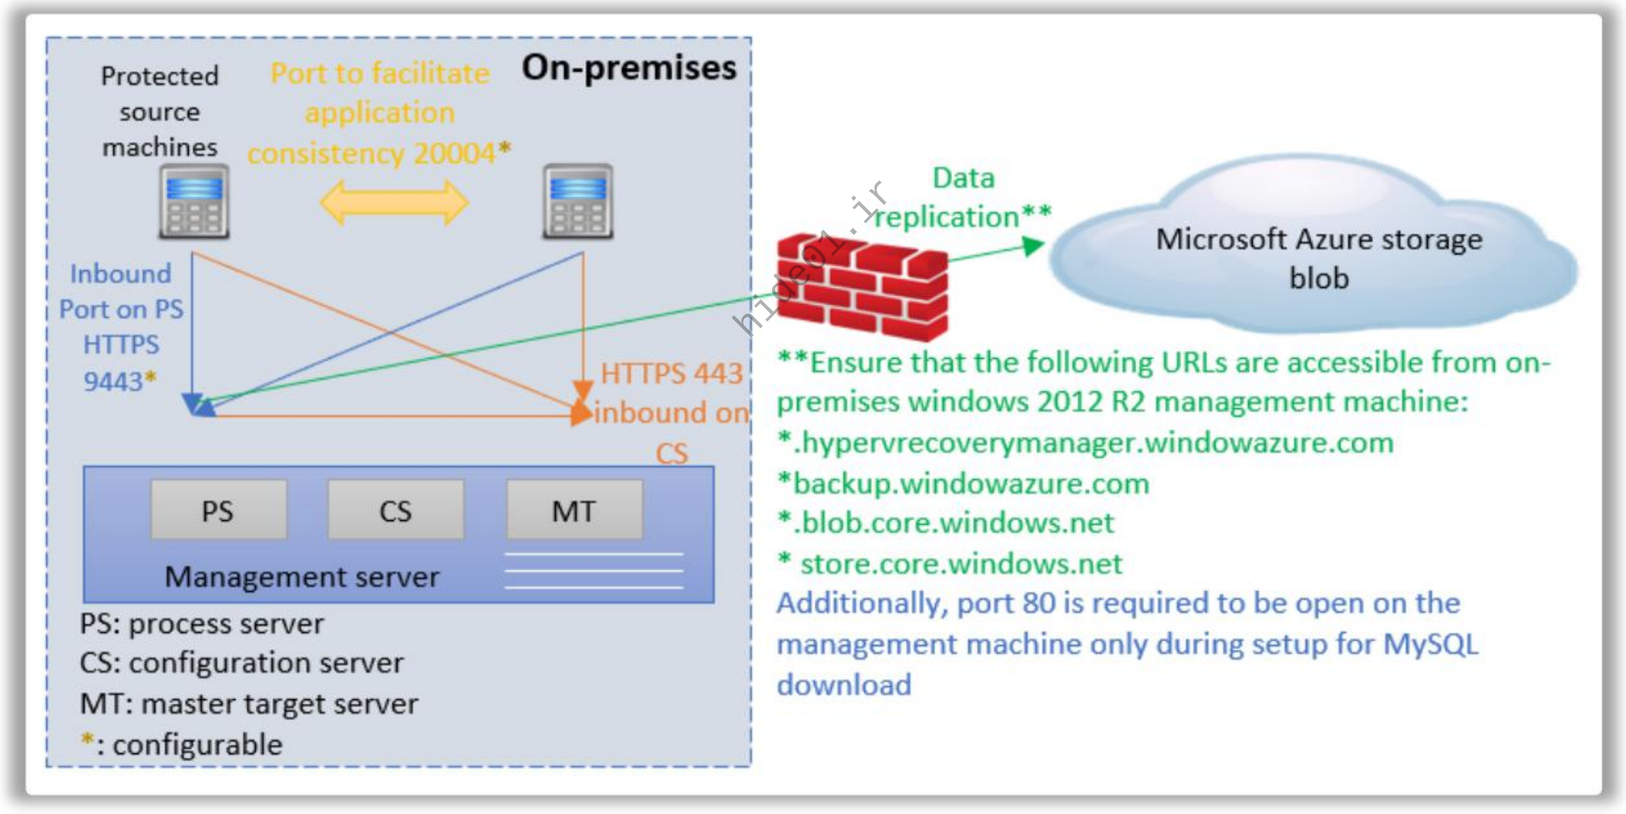

Figure 9.131: Replication Process

- The customer needs to set up the deployment which includes Azure and on-premises components. Then, the customer must specify the replication source and target, create a replication policy, set up the configuration server, and enable replication in the Recovery Site Vault.
- The replication policy helps machines to replicate, and the copy of server data is replicated by machines to Azure Storage.
- Replication of delta changes to Azure starts after the initial replication is completed. Tracked changes for machines is stored in a file (with .hrl extension)
    - For replication management, machines communicate with the configuration server over HTTP port 443 inbound.

Certified Cloud Security Engineer Copyright © by **EC-Council**
All Rights Reserved. Reproduction is Strictly Prohibited.

- The process server receives replication data from machines over HTTP port 9443 inbound.
- The configuration server uses HTTP port 443 outbound to orchestrate replication management with Azure.
- Data from source machines are received by the process server that optimizes, and encrypts it and then, transfers it to Azure Storage by using HTTP port 443 outbound.
    - Traffic is replicated over the internet to Azure Storage public endpoints. Customers can use Azure ExpressRoute public peering as an alternate option.

Failover and Failback Process

After setting up replication, it is good to run a DR drill (test failover) to confirm everything is working as expected. The following things must be considered.

- Azure does not support planned failover.
- The customer should have an on-premises VMware infrastructure because failback to the on-premises VMware VM is required.
- When a failover is run, Azure VMs are created in Azure Storage from replicated data.
- After successful failover, the customer is committed to access their workloads from the Azure VM.
- The customer can failback once their primary site is operational again.
- The customer needs to set up a failback infrastructure that contains a VPN connection, a temporary process server in Azure, a failback policy, VMware infrastructure, and a separate master target server.
- Once the components are set up, failback happens in three stages.
    - Re-protect Azure VMs to allow them to replicate data from Azure to the on-premises VMware VM.
    - Perform a failover to the on-premises site.
    - Re-enable replication after the workloads have failed back.

hide01.ir

Certified Cloud Security Engineer Copyright © by **EC-Council**
All Rights Reserved. Reproduction is Strictly Prohibited.

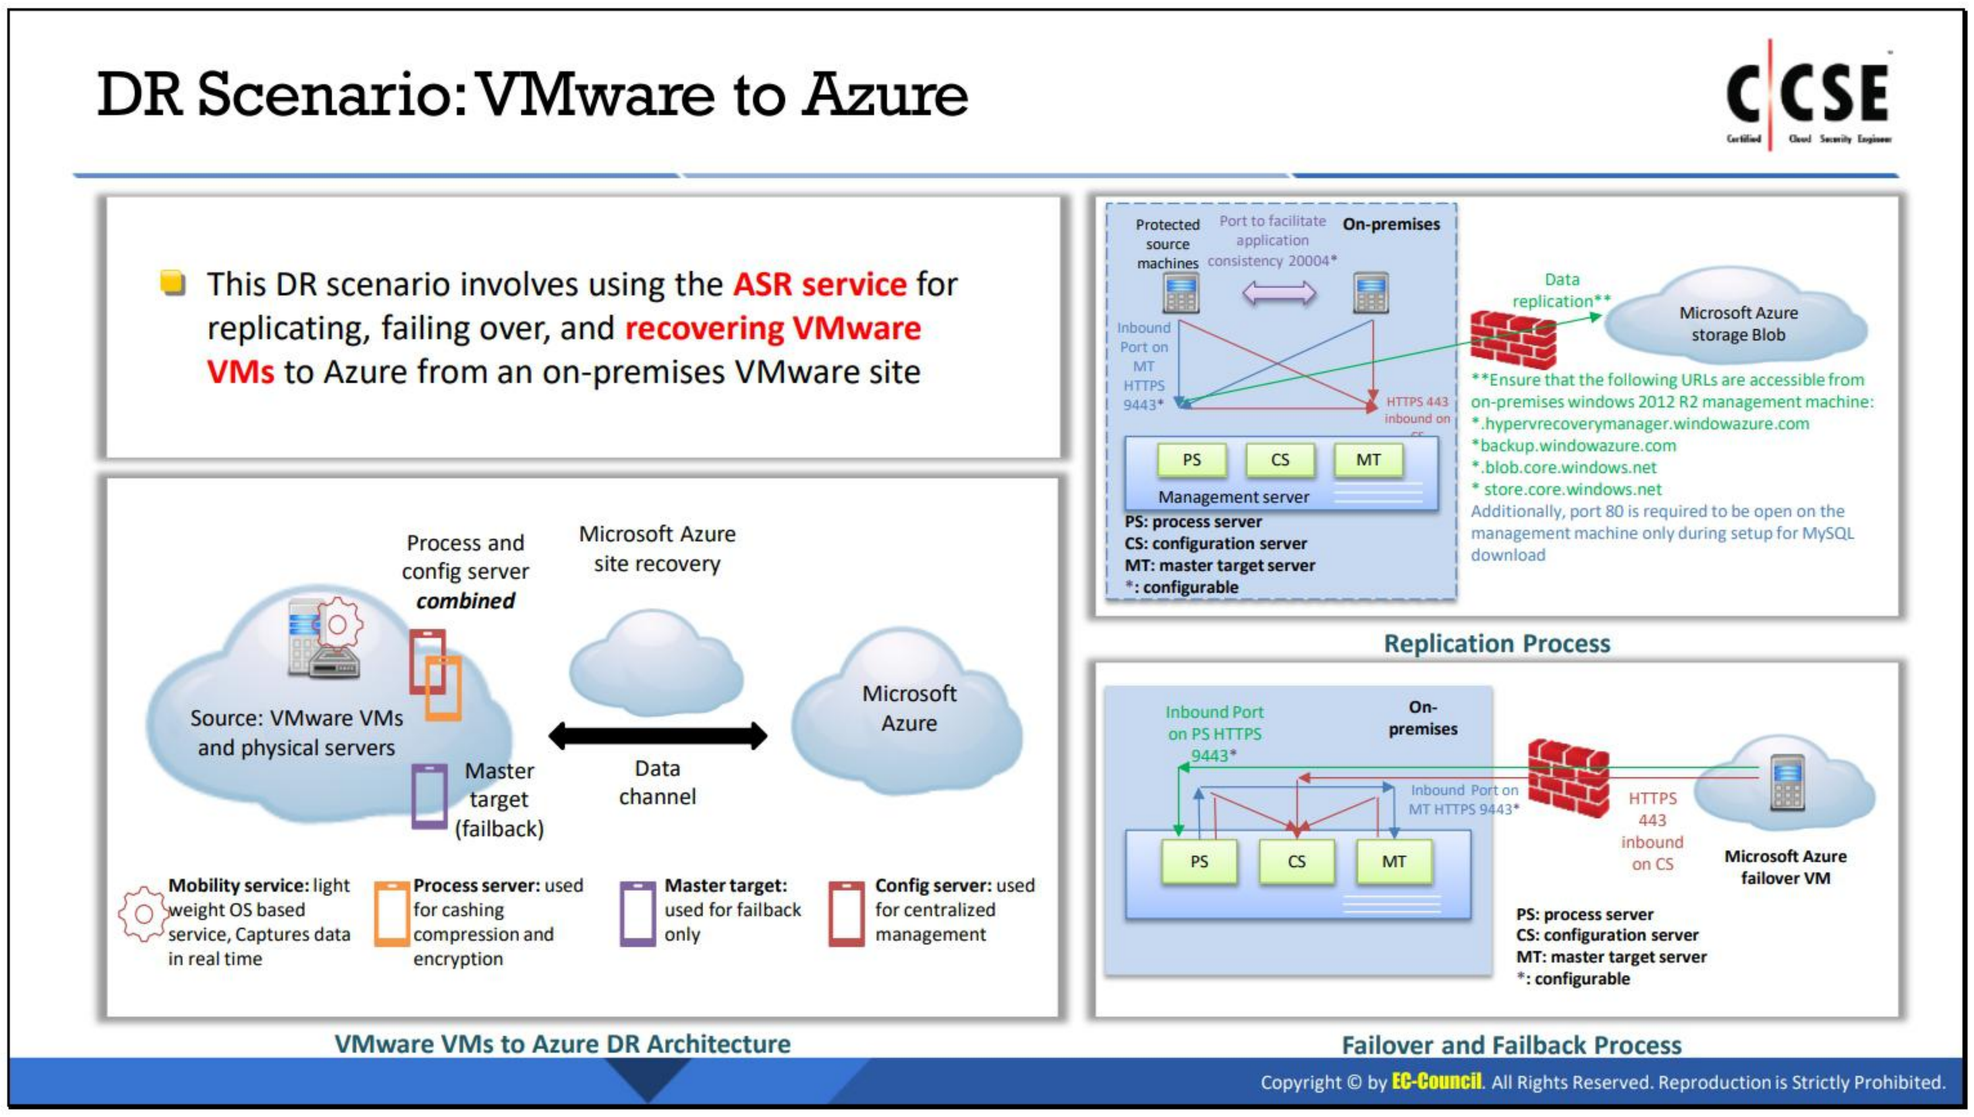

DR Scenario: VMware to Azure

This DR scenario involves using the ASR service for replicating, failing over, and recovering VMware VMs to Azure from an on-premises VMware site.

hide01.ir

Architectural Components

- **Azure**: Replicated data from on-premises physical servers are stored in Azure managed disks. This replicated data is used to create VMs when the customer performs a failover from on-premises to Azure.
- **Configuration server machine**: All on-premises Site Recovery components like the master target server, process server, and configuration server are run by this machine.
- **VMware servers**: The customer must add VMware servers to the RSV during Site Recovery deployment.
- **Replicated servers**: On every server that is to be replicated, the mobility service is installed. Customers can install the service via an automated deployment method or manually.

Certified Cloud Security Engineer Copyright © by **EC-Council**
All Rights Reserved. Reproduction is Strictly Prohibited.

The following image shows the VMware VMs-to-Azure DR architecture.

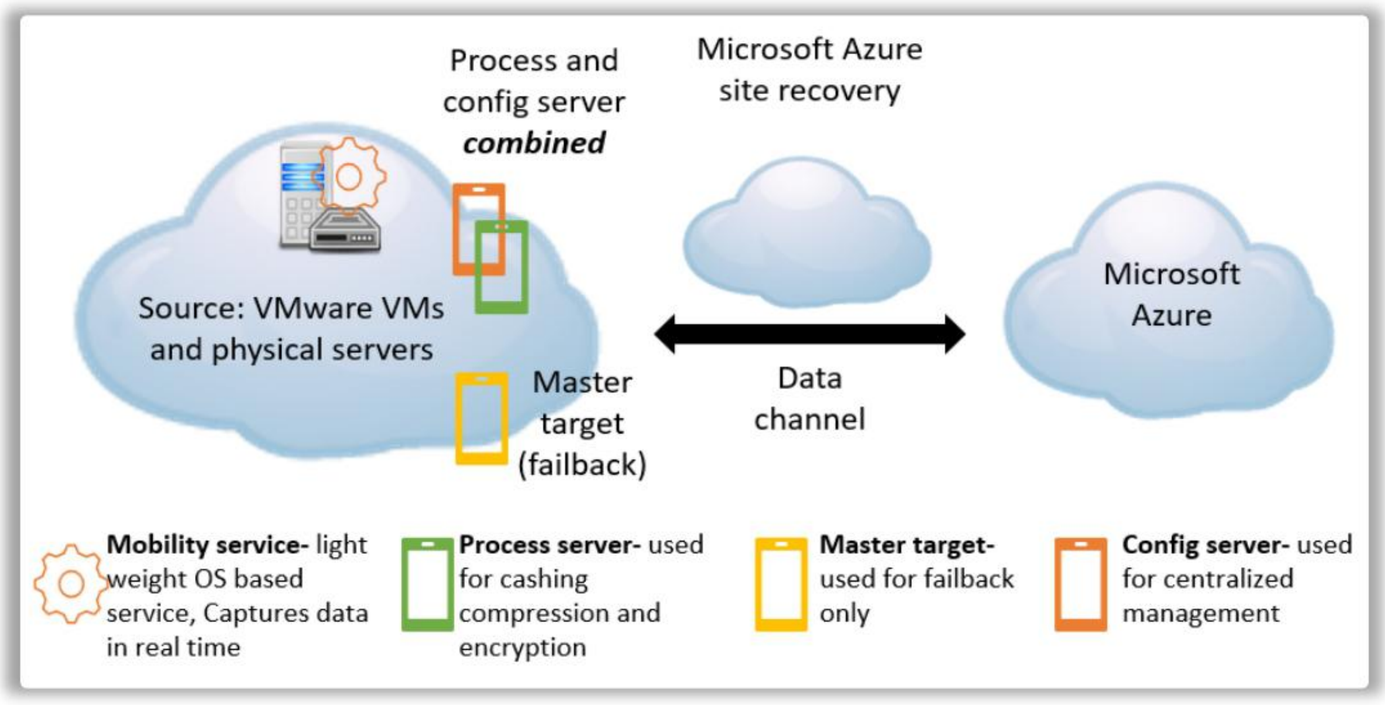

Figure 9.132: VMware VMs-to-Azure DR architecture

Outbound Network Connectivity

The customer needs to modify outbound network connectivity to enable their environment to replicate data to Azure. Certain URLs need to be allowed if a URL-based firewall proxy is used to control outbound network connectivity to ensure Site Recovery works as expected.

Replication Process

- Initial replication to Azure Storage starts once the customer enables replication from a VM by using the specified replication policy.
- Traffic is replicated over the internet to Azure Storage public endpoints. Customers can use Azure ExpressRoute public peering as an alternate option.
- The initial replication operation sends complete data on the machine to Azure at the time of enabling the replication. Replication of delta changes to Azure starts after the initial replication is completed. Changes that are tracked for machines are sent to the process server.
- The flow of communication is as follows:
    - For replication management, VMs communicate with the configuration server over HTTP port 443 inbound.
    - The process server receives replication data from VMs over HTTP port 9443 inbound. This port can be changed.
    - The configuration server uses HTTP port 443 outbound to orchestrate replication management with Azure.
    - Data from VMs is received by the process server that optimizes it, encrypts it, and then transfers it to Azure Storage by using HTTP port 443 outbound.

Certified Cloud Security Engineer Copyright © by **EC-Council**
All Rights Reserved. Reproduction is Strictly Prohibited.

- Replication data logs are stored first in a cache storage account. After processing these logs, an Azure Managed Disk stores the data. On the disk, recovery points are created.

Re-synchronization Process

- Sometimes, network connectivity issues occur between the process server to Azure or between source machines to the process server while transferring delta changes or during initial data replication. Any of them can cause failures in data transfer for some time.
- A machine is marked for re-synchronization by Site Recovery to prevent data integrity issues and reduce data transfer costs.
- Site Recovery can also mark a machine for re-synchronization.
  - If a machine experiences force shut down.
  - If it experiences configurational changes such as disk re-sizing.
- Only delta data is sent to Azure via re-synchronization.
- Re-synchronization is by default scheduled to run automatically in odd hours. If the customer does not want to wait until default re-synchronization occurs, they can manually resynchronize a VM.
- An error will be shown in the Azure portal on the specific machine where the default re-synchronization fails and manual intervention is needed for re-synchronization. This helps the customer to resolve the issue and initiate a manual re-synchronization.
- Once the re-synchronization is completed, the replication of delta changes will start again.

Replication Policy

Once the customer enables Azure VM replication, Site Recovery, by default, creates a new replication policy with the (below) default policy settings.

- **point retention**: It defines how long Site Recovery can keep recovery points. The default time is 24 hours.
- **App-consistent snapshot frequency**: It shows how often app-consistent snapshots are taken by Site Recovery. The default frequency for taking snapshots is every four hours.

Managing Replication Policies

The customer can modify and manage the default replication policy as follows:

- When they enable replication, they can change the default policy settings.
- They can create a replication policy first and then apply it while they enable replication.

Snapshots and Recovery Points

Recovery points are created from snapshots (backups) of VM disks that are taken at a particular point-in-time. The customer can use a recovery point to restore their VM in the specified target location when they fail over a VM. By default, Site Recovery takes crash-consistent snapshots of

Certified Cloud Security Engineer Copyright © by **EC-Council**
All Rights Reserved. Reproduction is Strictly Prohibited.

data. If the customer specifies a frequency for taking snapshots, Site Recovery takes app-consistent snapshots.

- **Crash-consistent snapshot**: These snapshots do not include anything in memory, and they capture data that was present on the disk when the snapshot was taken. By default, crash-consistent recovery points are created every 5 minutes by Site Recovery and this time cannot be changed.
- **App-consistent Snapshot**: It contains complete information that is present in a crash-consistent snapshot as well as data in memory and transactions that are in progress. Crash-consistent snapshots are used to create app-consistent recovery points.

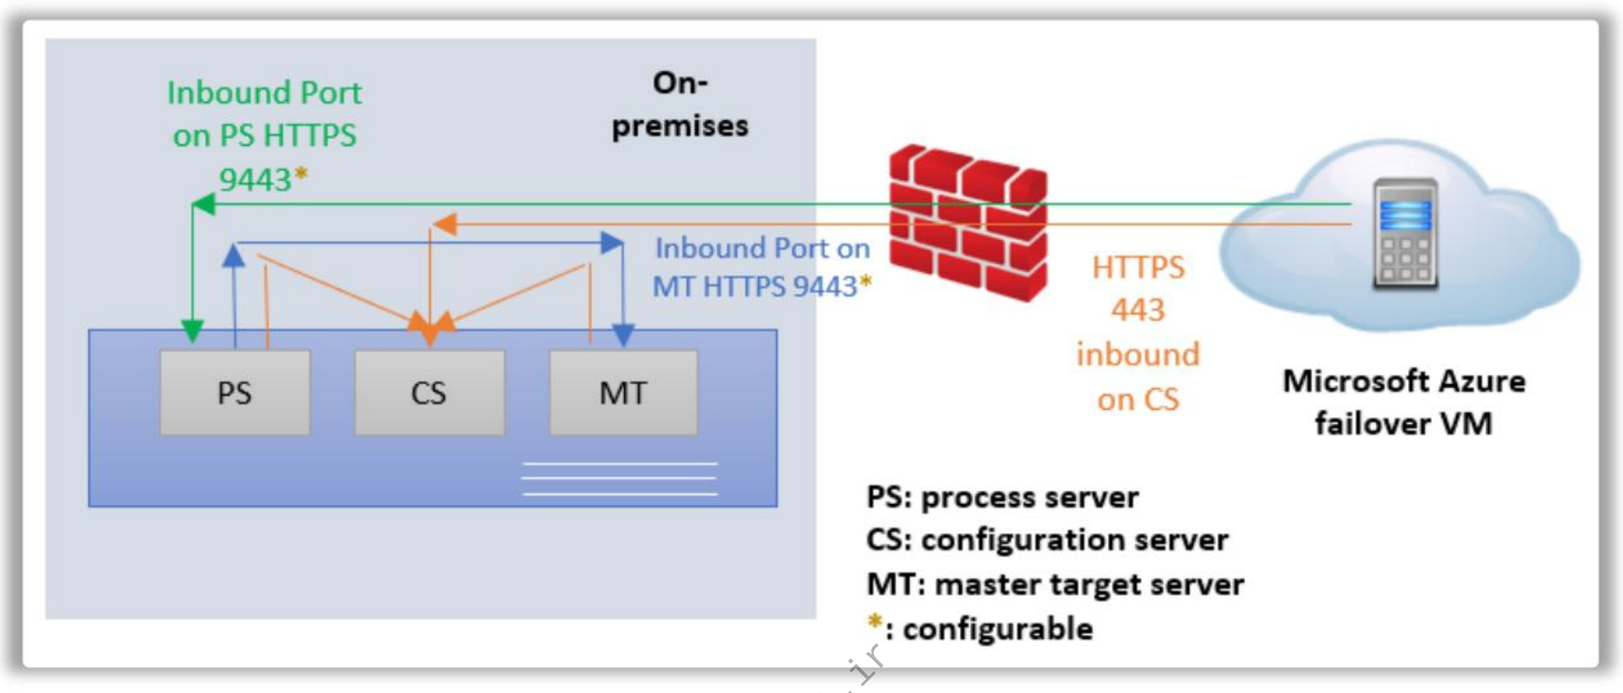

Figure 9.133: Failover and Failback Process

After setting up replication, it is good to run a DR drill (test failover) to confirm everything is working as expected. The following things must be considered:

- The customer can fail over one VM or create recovery plans to failover many VMs.
- After starting the initial failover, the customer is committed to access their workloads from the Azure VM.
- The customer can failback once their primary site is operational again.
- The customer needs to set up a failback infrastructure that contains a VPN connection, a temporary process server in Azure, a failback policy, and a separate master target server.
- Once the components are set up, failback happens in three stages.
    - Re-protect Azure VMs to allow them to replicate data from Azure to the on-premises VMware VM.
    - Perform a failover to the on-premises site.
    - Re-enable replication for on-premises VMs after the workloads have failed back.

Certified Cloud Security Engineer Copyright © by **EC-Council**
All Rights Reserved. Reproduction is Strictly Prohibited.

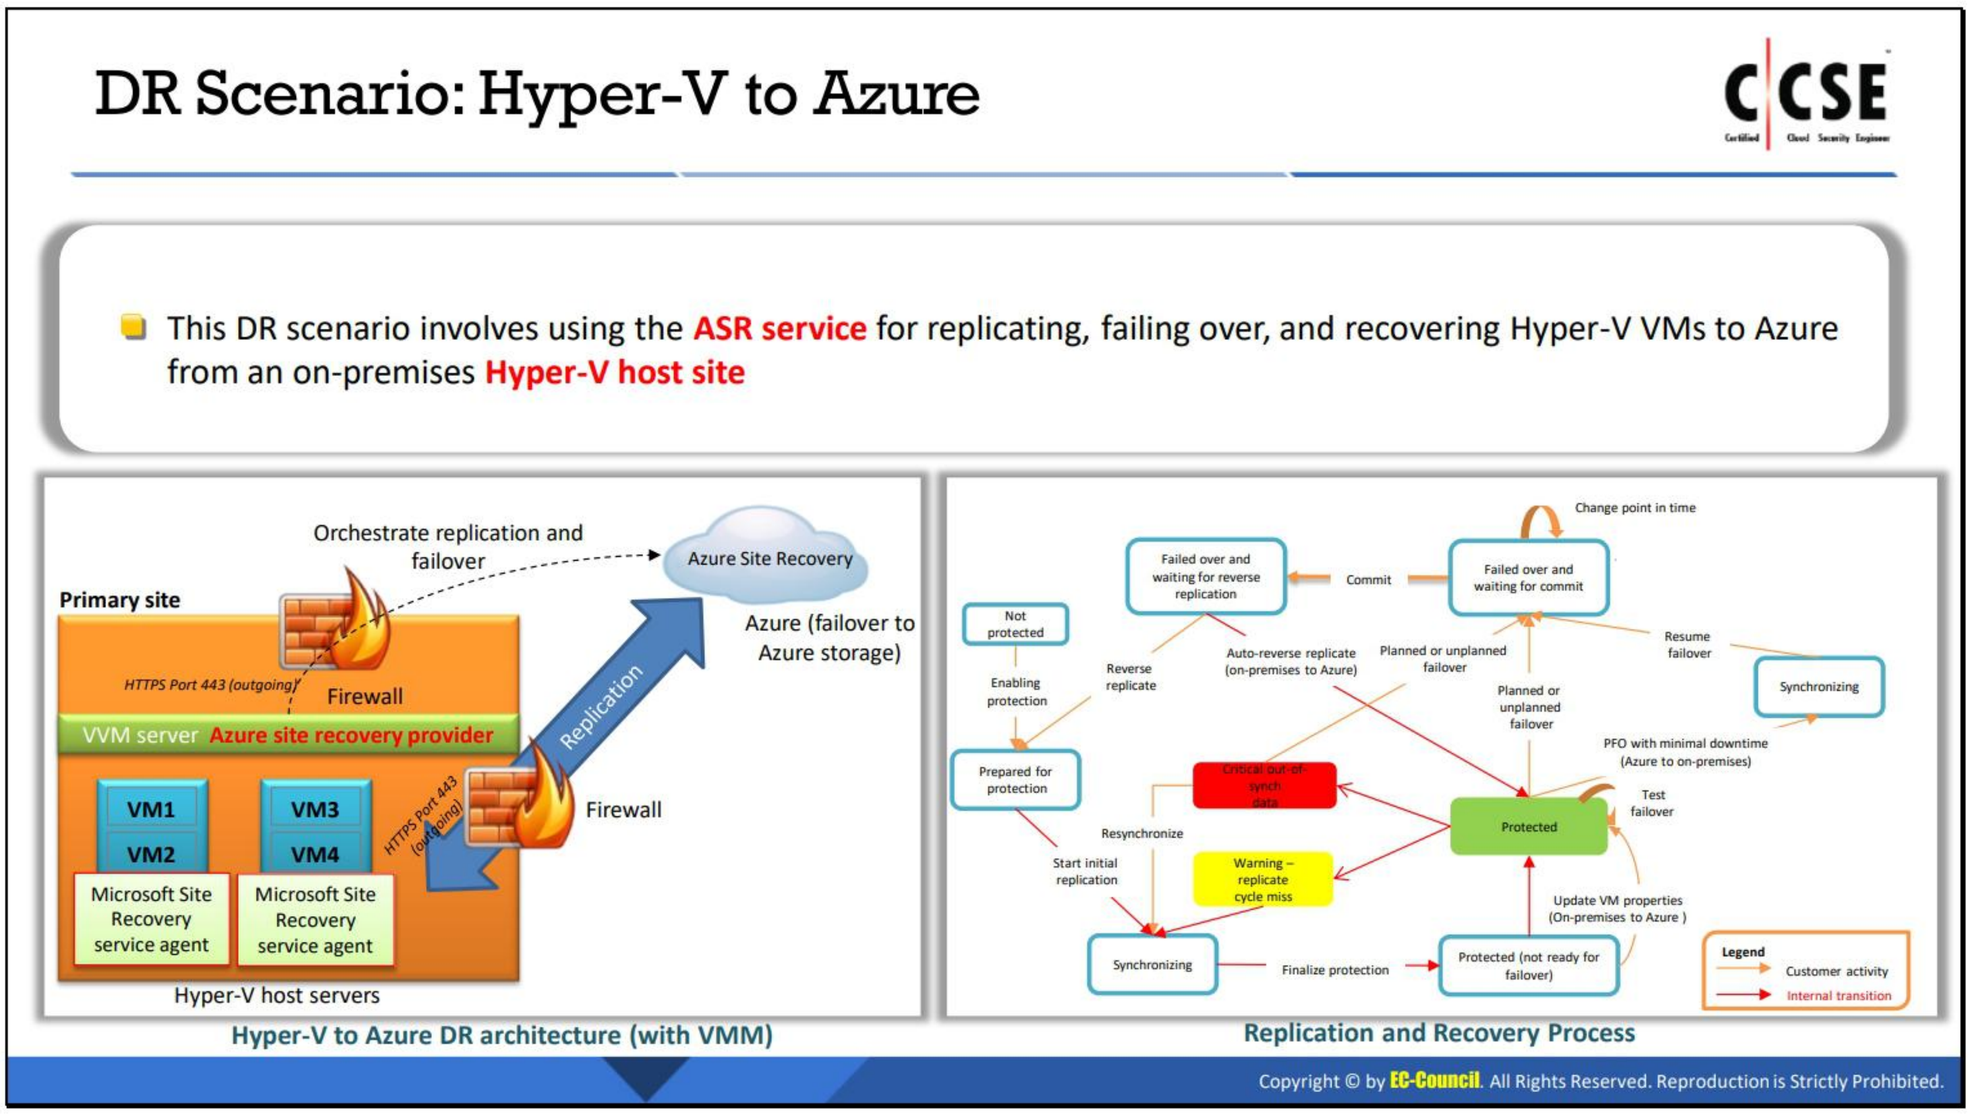

DR Scenario: Hyper-V to Azure

This DR scenario involves using the ASR service for replicating, failing over, and recovering Hyper-V VMs to Azure from an on-premises Hyper-V host site.

h1e01.ir

Architectural Components for Hyper-V without VMM

- **Azure**: Replicated data from on-premises physical servers are stored in Azure managed disks. This replicated data is used to create VMs when the customer performs a failover from on-premises to Azure.
- **Hyper-V**: Hyper-V hosts and clusters are gathered by the customer into Hyper-V sites during Site Recovery deployment. The customer has to install the Recovery Services agent and Site Recovery Provider on every standalone Hyper-V host or cluster node. The Site Recovery Provider helps in orchestrating replication with Site Recovery on the internet. Data replication is handled by the Recovery Services agent.
- **Hyper-V VMs**: It is required to have one or more VMs running on Hyper-V.

Architectural Components for Hyper-V with VMM

- **Azure**: Replicated data from on-premises physical servers are stored in Azure managed disks. This replicated data is used to create VMs when the customer performs a failover from on-premises to Azure.
- **Hyper-V VMs**: It is required to have one or more VMs running on Hyper-V.
- **VMM server**: It has one or more clouds that contain Hyper-V hosts. The customer should install the Site Recovery Provider on the VMM server.

Certified Cloud Security Engineer Copyright © by **EC-Council**
All Rights Reserved. Reproduction is Strictly Prohibited.

- **Hyper-V host**: The customer should install the Recovery Services agent on every Hyper-V host or cluster node.
- **Networking**: VM networks are mapped to Azure virtual networks as part of networking. Azure VMs that are created after failover are added to the Azure network, which is mapped to the VM network.

The following image shows the Hyper-V-to-Azure DR architecture (with VMM).

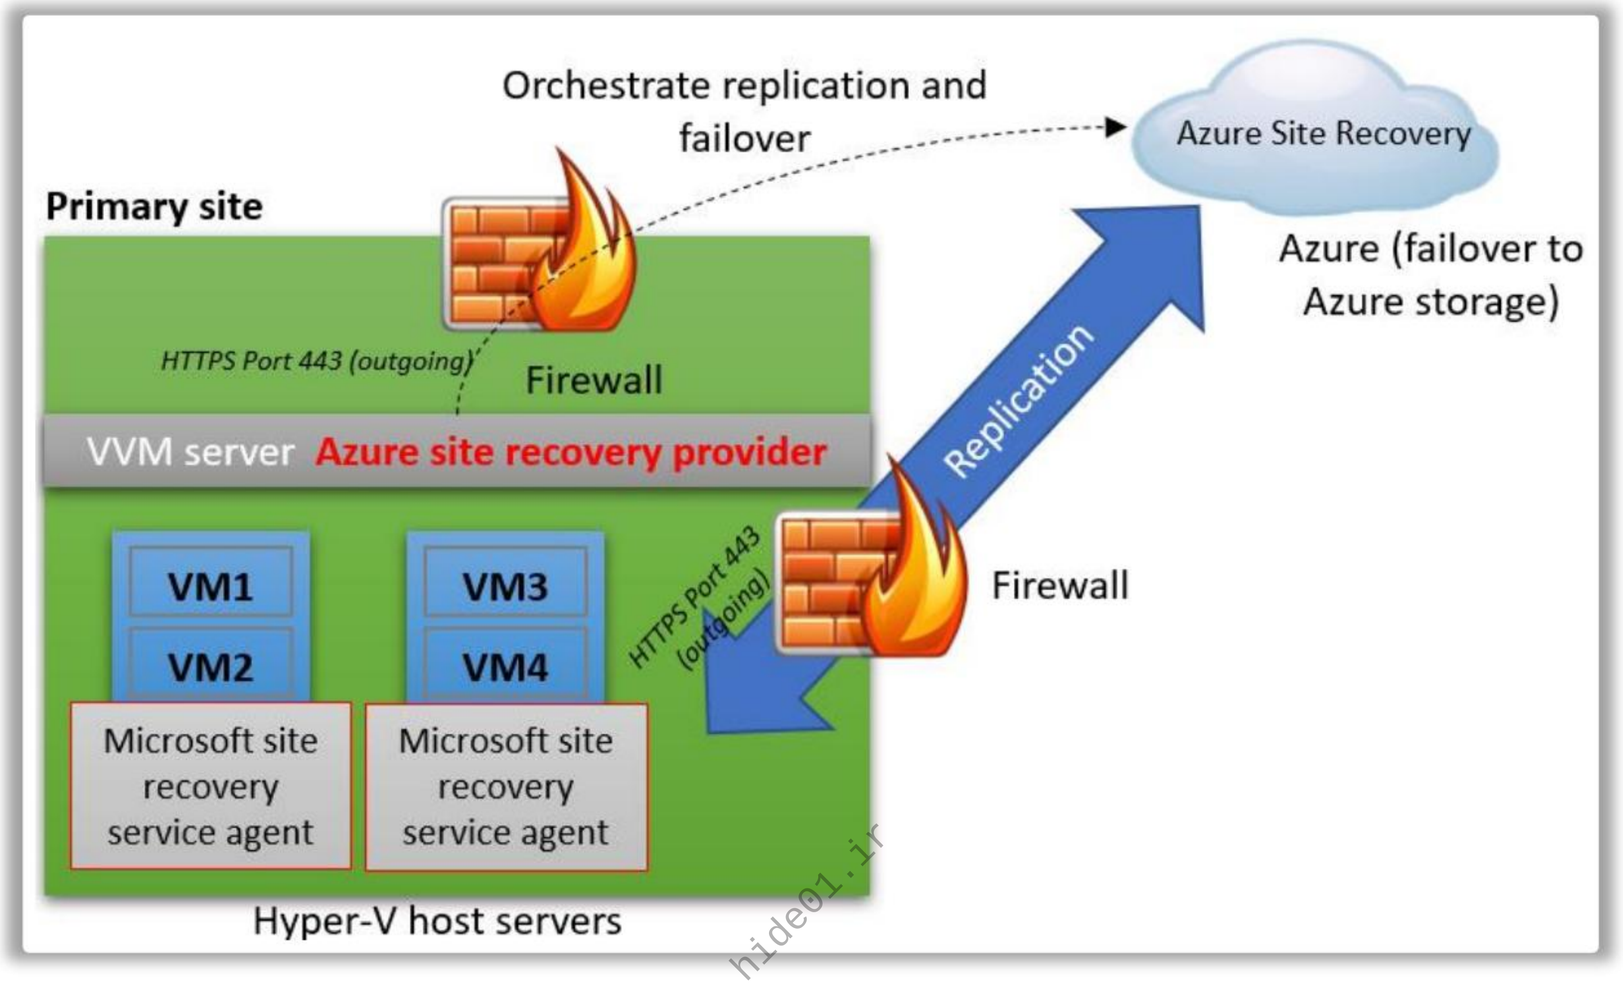

Figure 9.134: Hyper-V to Azure DR architecture (with VMM)

Source: https://www.microsoft.com

Outbound Network Connectivity

The customer needs to modify outbound network connectivity to enable their environment to replicate data to Azure. Certain URLs need to be allowed if a URL-based firewall proxy is used to control outbound network connectivity to ensure Site Recovery works as expected.

Certified Cloud Security Engineer Copyright © by **EC-Council**
All Rights Reserved. Reproduction is Strictly Prohibited.

Replication Process

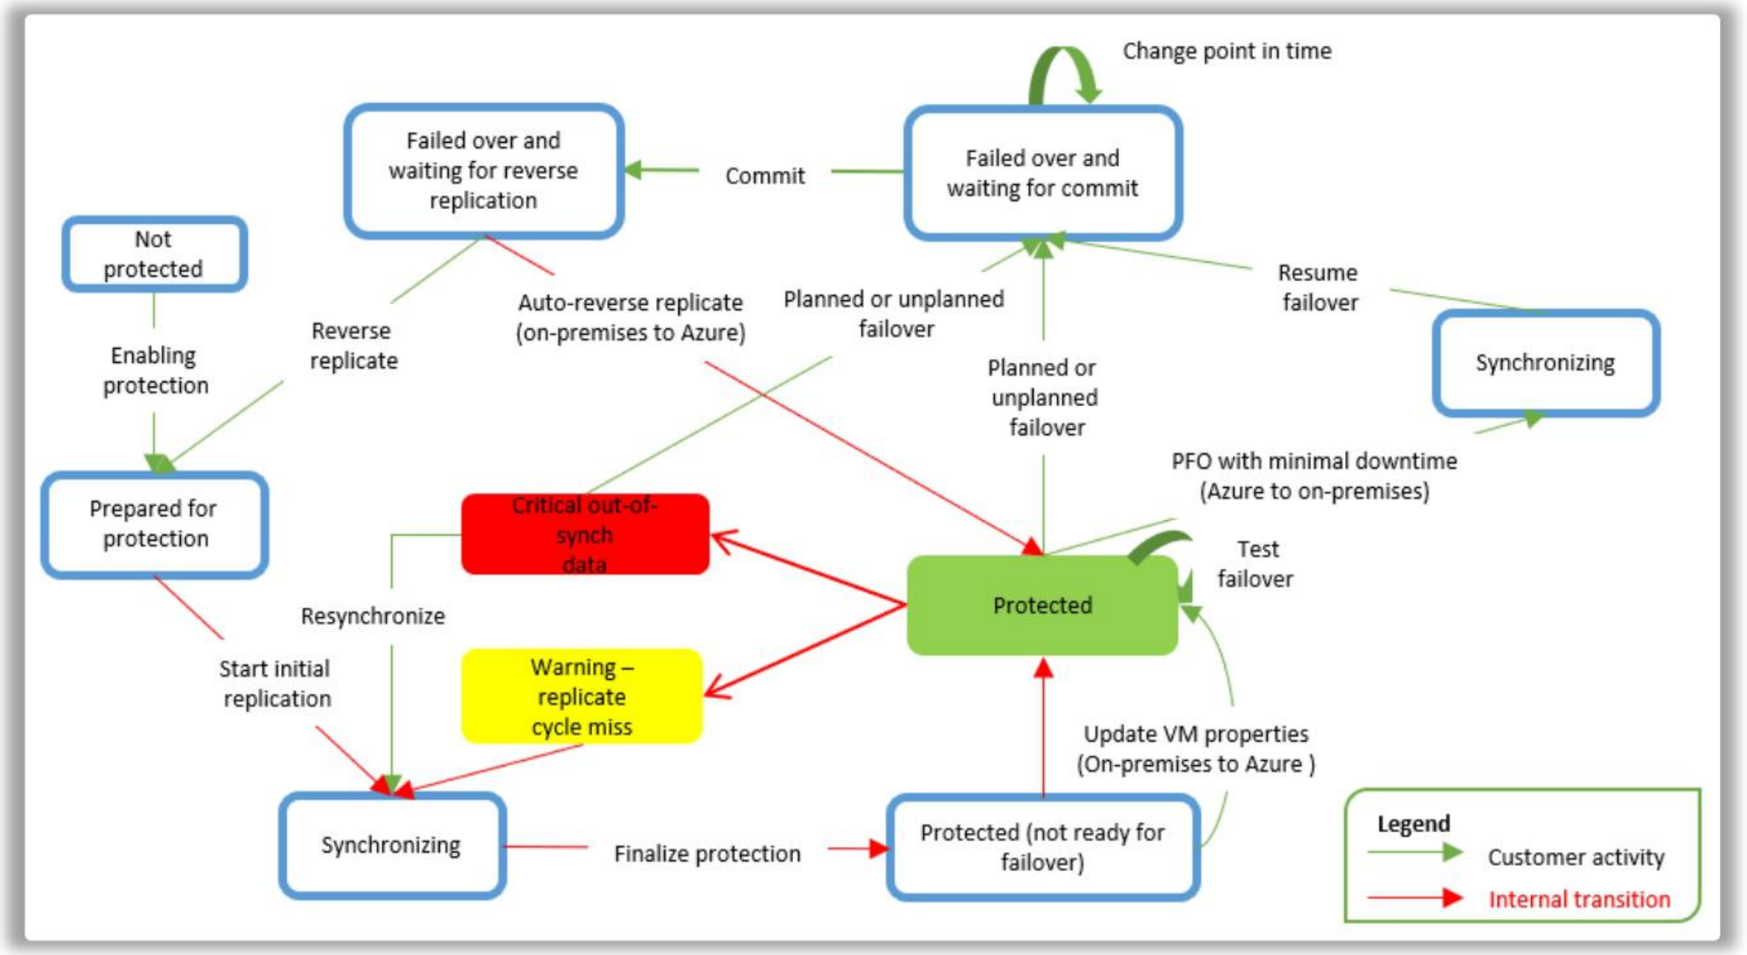

Figure 9.135: Replication and Recovery Process

The above image shows the replication and recovery process.

Enable Protection

- The **Enable protection** begins when the customer enables protection for a Hyper-V VM on-premises or in the Azure portal.
- Before invoking the **CreateReplicationRelationship,** this job verifies that the machine complies with prerequisites to set up replication with the configured settings.
- To begin a full VM replication and send virtual disks of the VM to Azure, the job starts the initial replication by using the **StartReplication** method.
- The customer can use the **Jobs** tab to monitor the job.

hidoo1.ir

Certified Cloud Security Engineer Copyright © by **EC-Council**
All Rights Reserved. Reproduction is Strictly Prohibited.

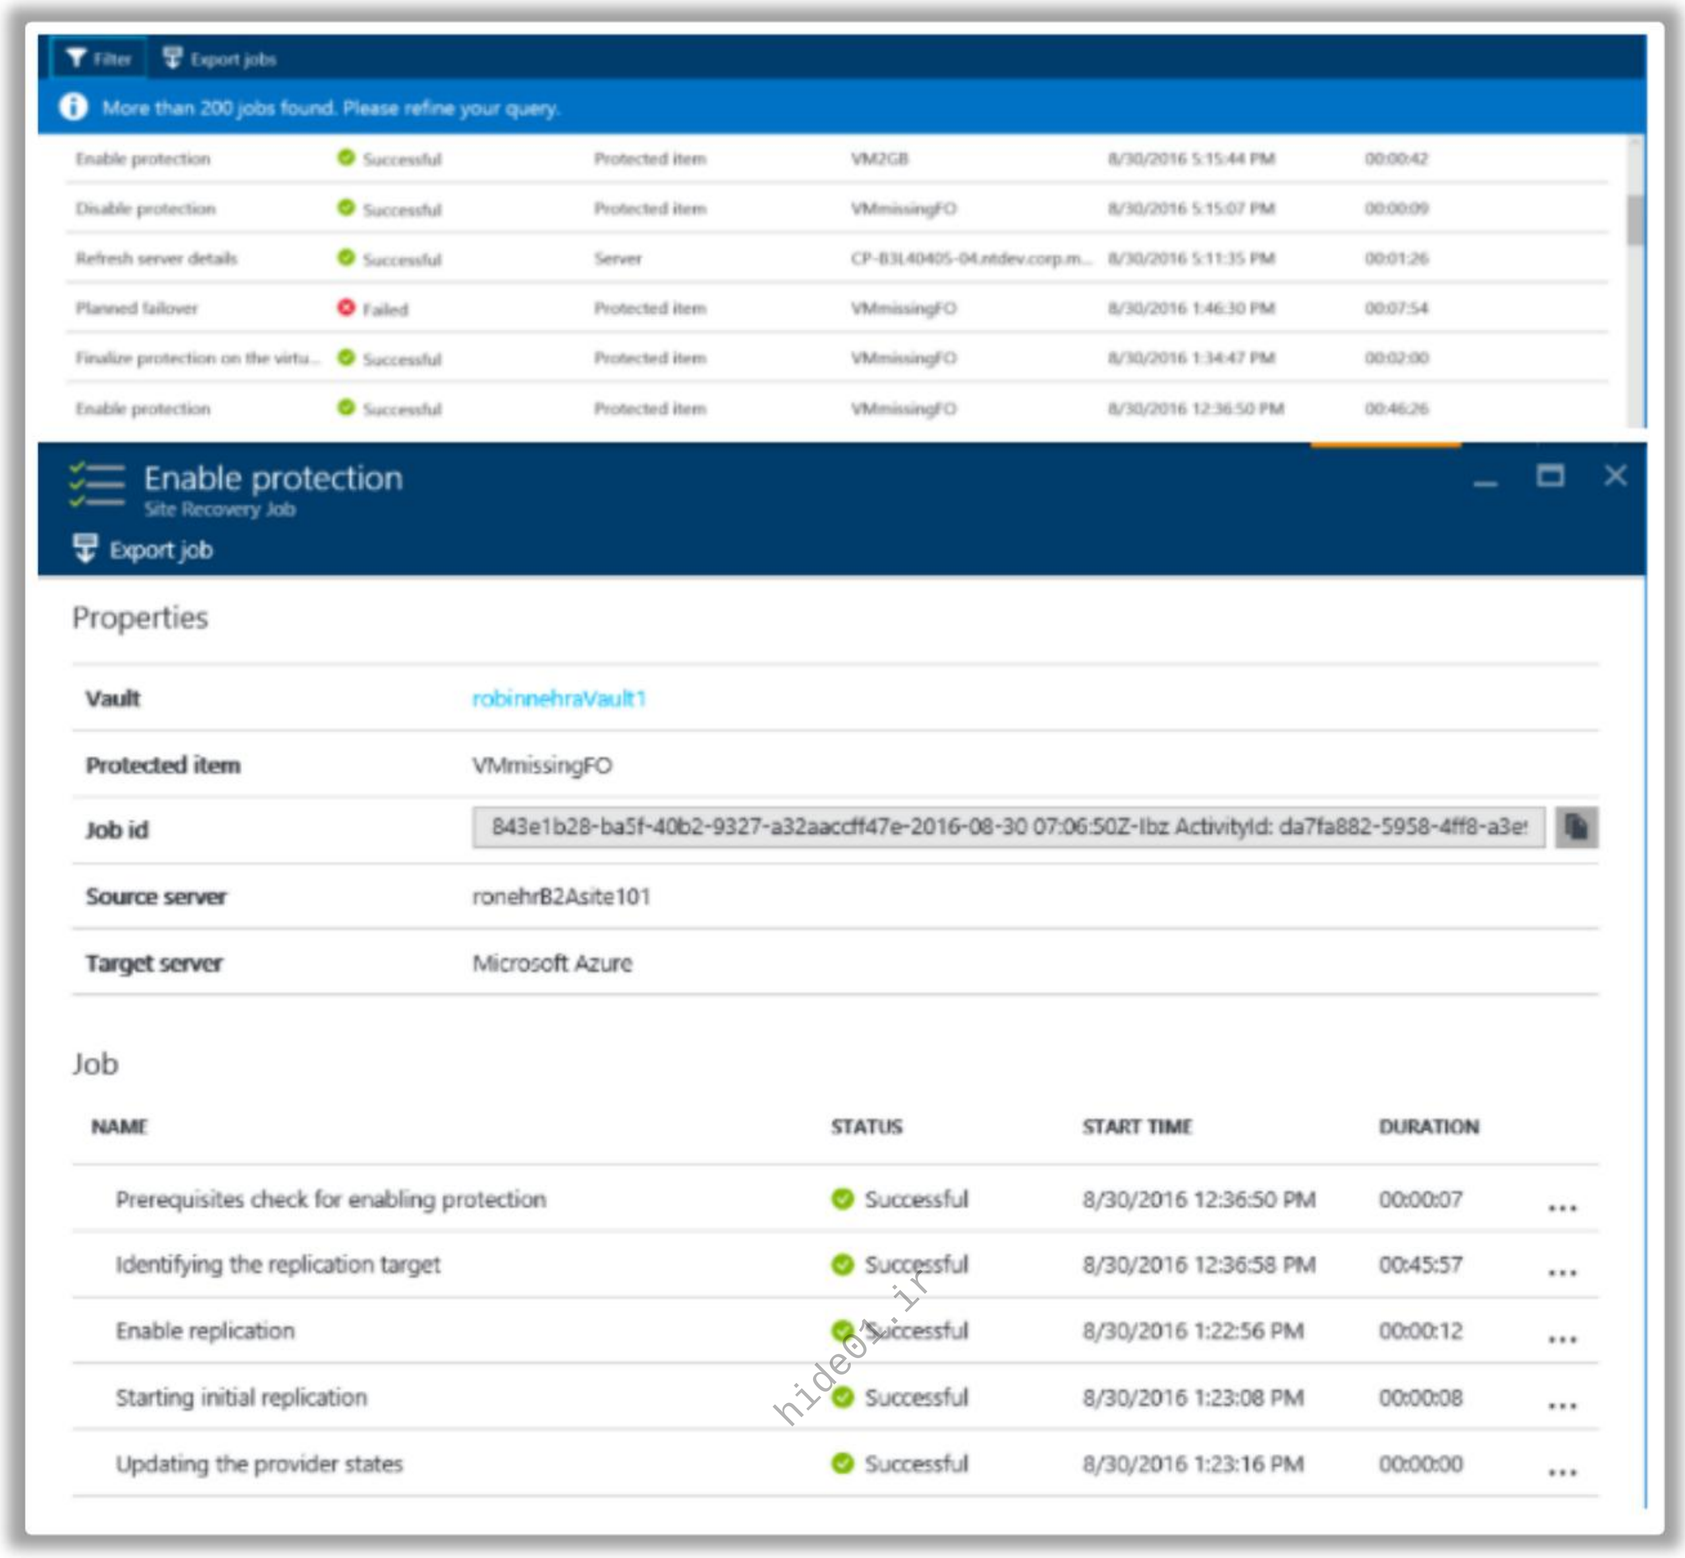

Figure 9.136: Job Tab Console

Initial Replication

- A Hyper-V VM snapshot is taken when initial replication is set or triggered.
- Virtual disks present on the VM get replicated one by one until they all get copied to Azure.
- The Hyper-V Replica Replication Tracker tracks the disk changes that occur while initial replication is going on as Hyper-V replication logs (.hrl). All these logs are stored in the same folder where the disks are present. There is an associated .hrl file with each disk that is sent to the secondary storage. The disk resources are consumed by log files and the snapshot when initial replication takes place.
- The snapshot is deleted when the initial replication is completed.
- Delta disk changes are synchronized and added to the persistent disk.

Finalize Protection Process

- When the initial replication is completed, the **Finalize protection on the VM** job starts, which configures network and other post-replication settings to protect the VM.

Certified Cloud Security Engineer Copyright © by **EC-Council**
All Rights Reserved. Reproduction is Strictly Prohibited.

- Check the VM settings to confirm that it is ready for failover. Test if the failover occurs as expected by running a DR drill for the VM.

Delta Replication

- Delta replication is done after the initial replication is over as per the replication policy.
- The Hyper-V Replica Replication Tracker tracks the disk changes that occur while initial replication is going on as Hyper-V replication logs in an associated file with (.hrl) extension.
- There is an associated .hrl file with each disk that is sent to the secondary storage. When a log is on the way to Azure, the changes to the primary disks are tracked in another log file stored in the same folder.
- The customer can use the Azure portal to monitor the VM during initial replication and delta replication.

Re-synchronization Process

- When failure occurs while performing delta replication or a full replication is be feasible, then a VM is usually marked for re-synchronization.
- Only delta data is sent to Azure via re-synchronization.
- Once the re-synchronization is complete, the replication of delta changes will start again.
- If the customer does not want to wait until default re-synchronization occurs, they can manually resynchronize a VM as shown below.

h1901.ir

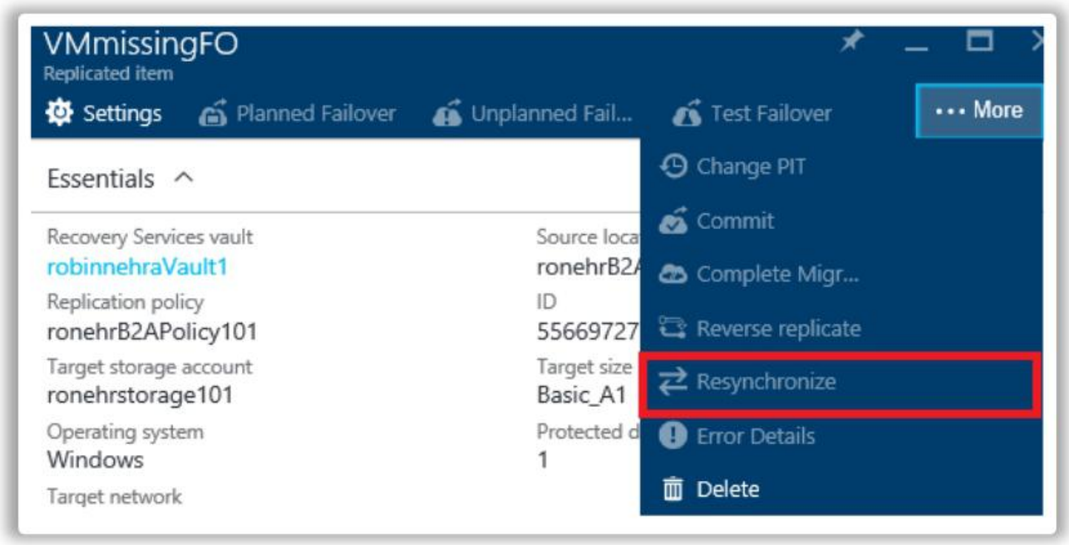

Figure 9.137: Click Resynchronize

Failover and Failback Process

- The customer can perform a planned or unplanned failover from on-premises Hyper-V VMs to Azure. Source VMs are shut down to make sure no data is lost when a planned failover is performed. If the primary site becomes unavailable, an unplanned failover is run.
- The customer can fail over one machine or create recovery plans to failover many machines.

Certified Cloud Security Engineer Copyright © by **EC-Council**
All Rights Reserved. Reproduction is Strictly Prohibited.

- The customer can see the created replica VMs in Azure after the first stage of the failover is completed.
- Then, they can access the workloads from the replica Azure VM.

Failback can be performed once the primary site is working and can handle production load. Failback can be performed by following these steps.

- Start a planned failover from Azure to the on-premises site.
- After initial synchronization is complete, choose to complete the failover. Log in to the on-premises VM to verify things are working as expected.
- Complete the failback process and start using the workloads from the on-premises VM.
- Re-enable replication for on-premises VMs after the workloads have failed back.

hide01.ir

Certified Cloud Security Engineer Copyright © by **EC-Council**
All Rights Reserved. Reproduction is Strictly Prohibited.

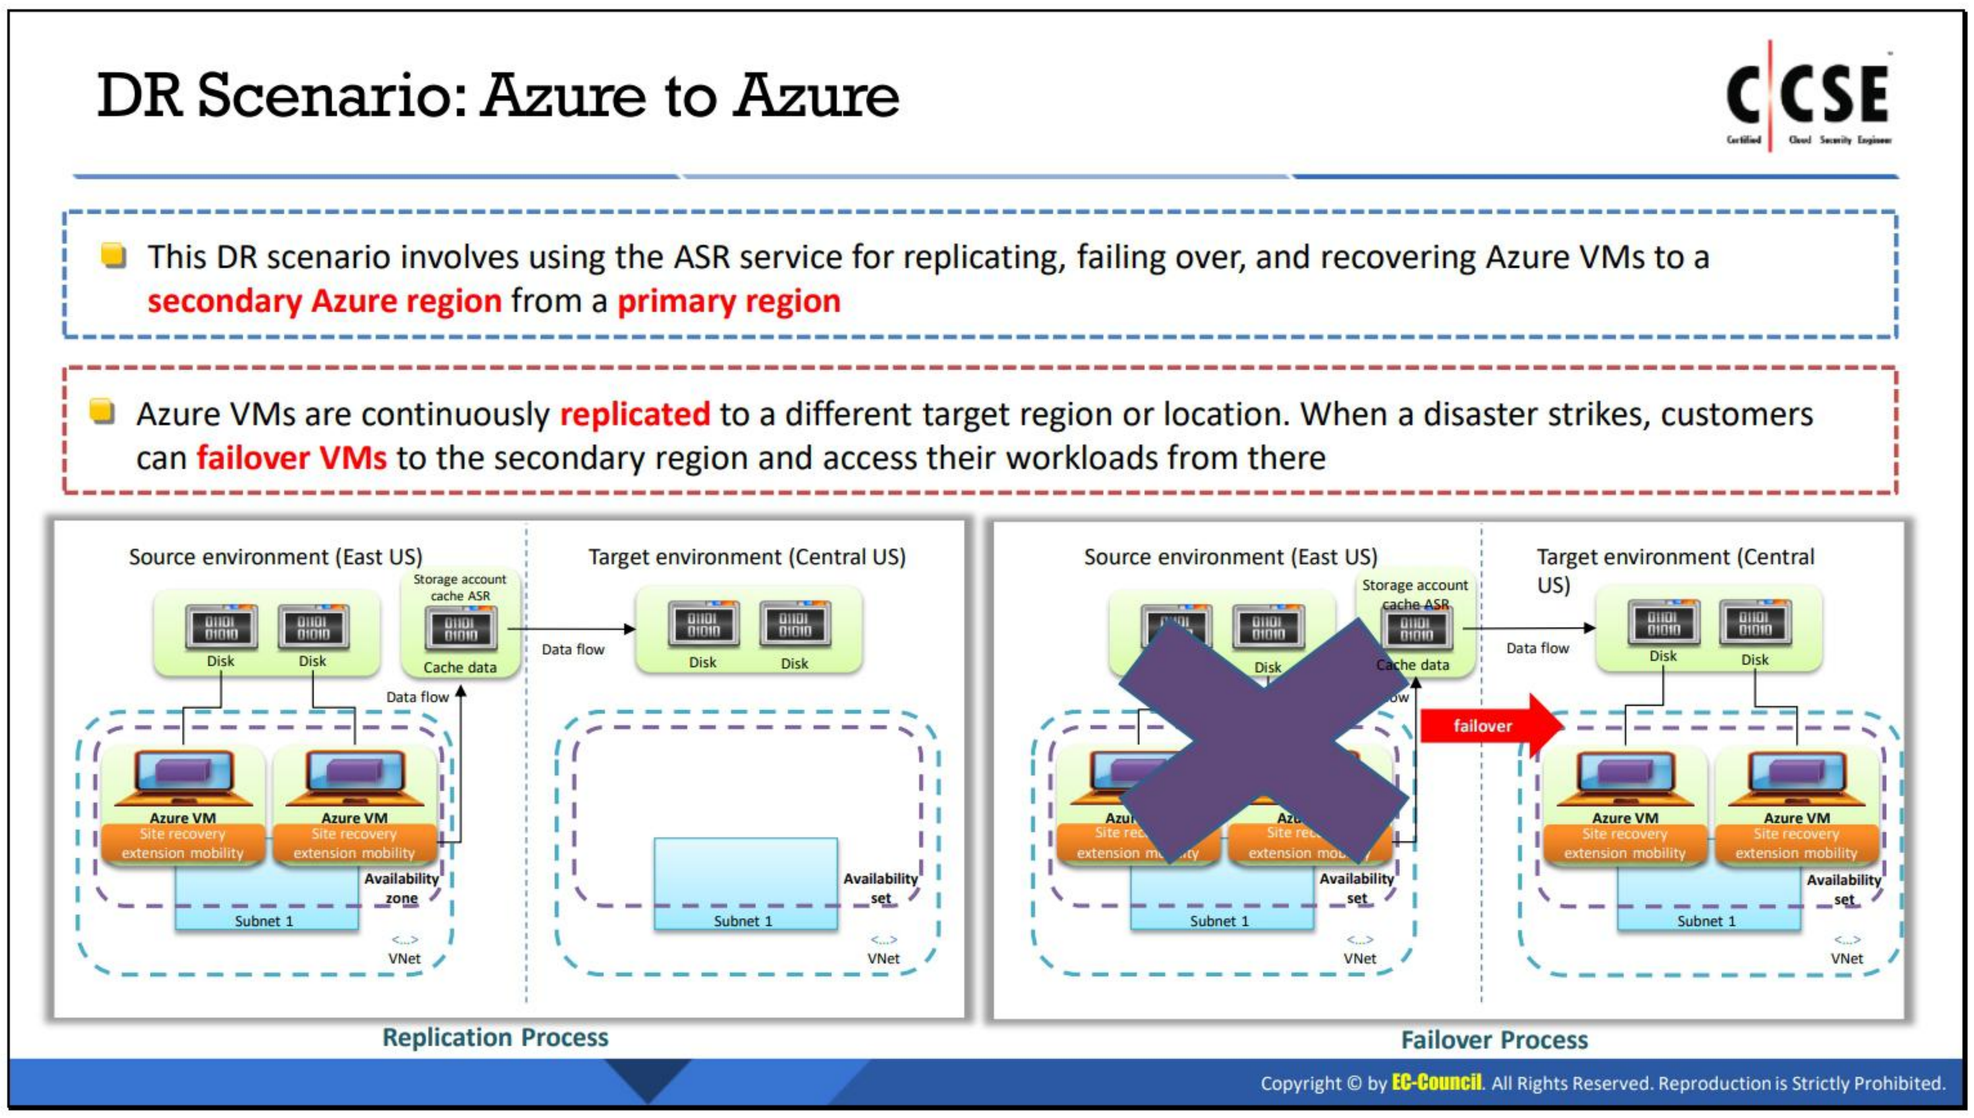

DR Scenario: Azure to Azure

This DR scenario involves using the ASR service for replicating, failing over, and recovering Azure VMs to a secondary Azure region from a primary region. Azure VMs are continuously replicated to a different target region or location. When a disaster strikes, customers can fail over VMs to the secondary region and access their workloads from there. They can failback once their primary region is working well.

Architectural Components

- **VMs in the source region:** One or more VM should be there in the source region.
- **Source VM storage**
- **Source VM networks:** VMs run in one or more subnets in a virtual network of the source region.
- **Cache storage account:** In the source region, a cache storage account is required because VM changes are stored in the cache storage before they are sent to the target storage during replication.
- **Target resources:** Target resources can be set up by default or manually. They are used during replication and while a failover occurs.

Certified Cloud Security Engineer Copyright © by **EC-Council**
All Rights Reserved. Reproduction is Strictly Prohibited.

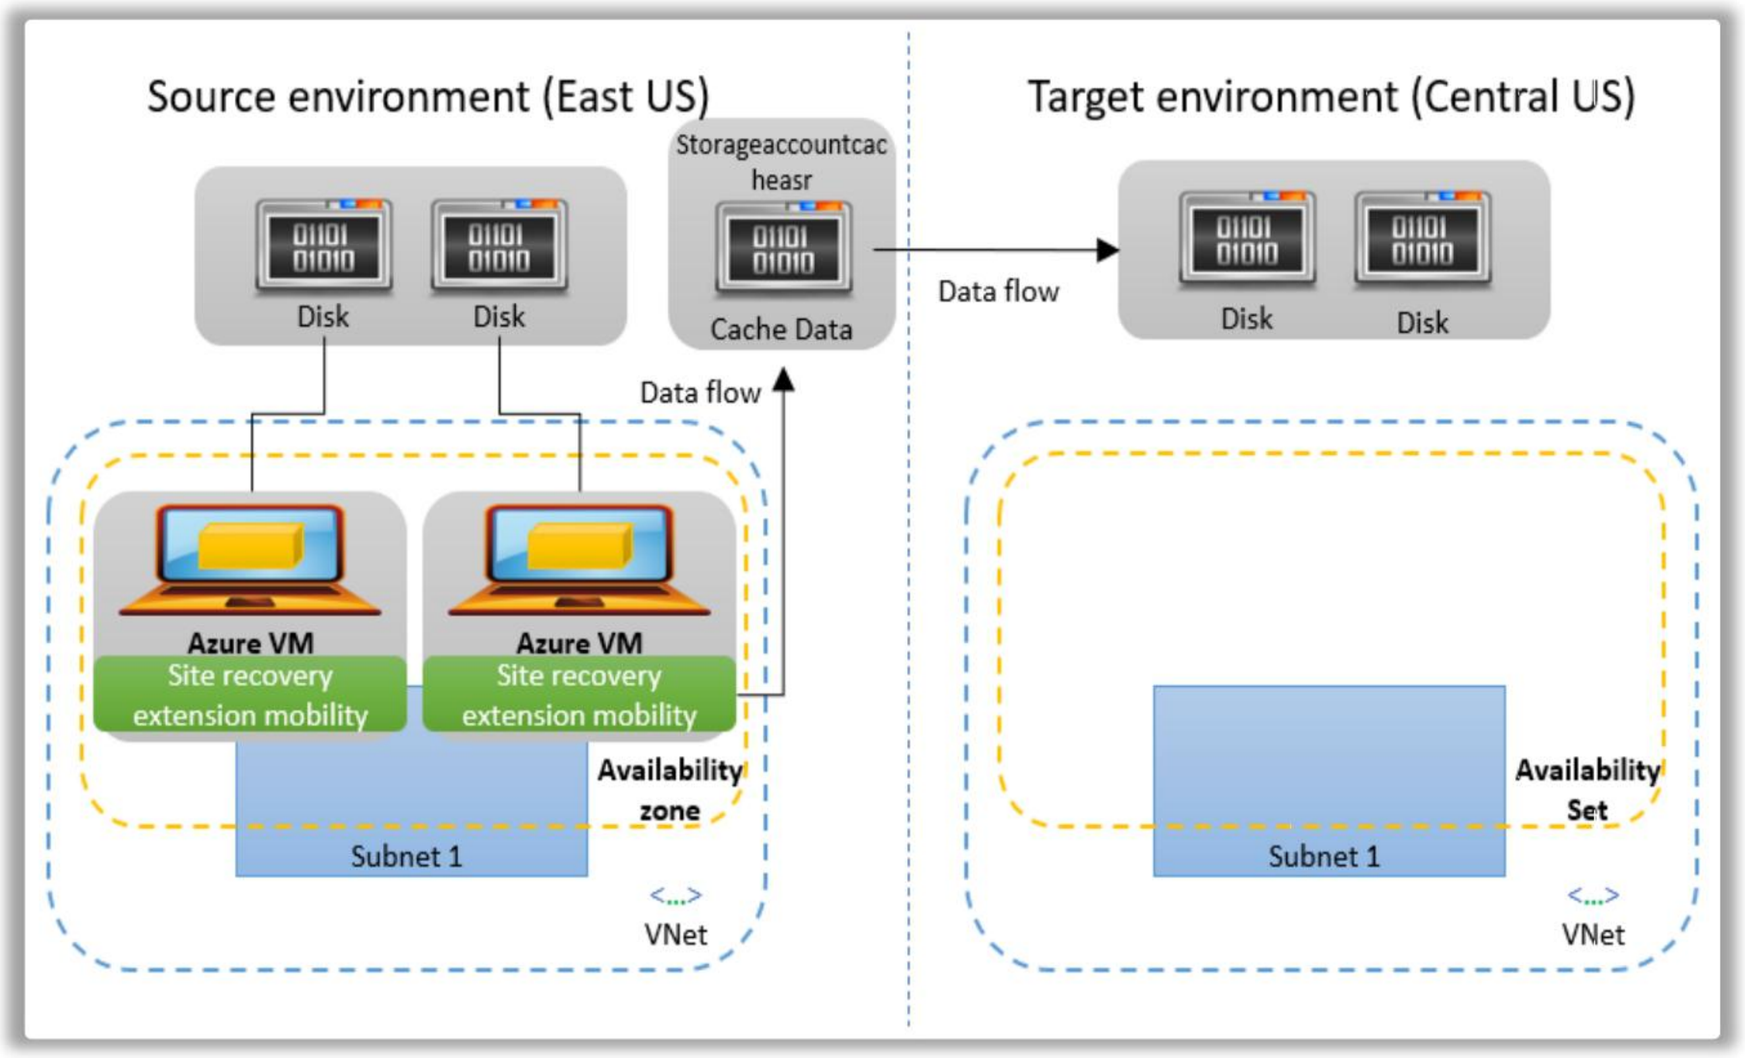

Figure 9.138: Replication Process

Target Resources and their Management

Site Recovery provides the option to create target resources automatically when replication of a VM is enabled. Customers can modify the default target settings when they enable replication or while replication is in progress.

Replication Policy

Once Azure VM replication is enabled, Site Recovery, by default, creates a new replication policy with the (below-mentioned) default policy settings.

- **Recovery point retention**: It defines how long Site Recovery can keep recovery points. The default time is 24 hours.
- **App-consistent snapshot frequency**: It shows how often app-consistent snapshots are taken by Site Recovery. The default frequency for taking snapshots is every four hours.

Managing Replication Policies

The customer can modify and manage the default replication policy as follows:

- When they enable replication, they can change the default policy settings.
- They can create a replication policy first and then apply it while they enable replication.

Snapshots and Recovery Points

Recovery points are created from snapshots (backups) of VM disks that are taken at a particular point in time. The customer can use a recovery point to restore their VM in the specified target location when they fail over a VM. By default, Site Recovery takes crash-consistent snapshots of

Certified Cloud Security Engineer Copyright © by **EC-Council**
All Rights Reserved. Reproduction is Strictly Prohibited.

data. If the customer specifies a frequency for taking snapshots, Site Recovery takes app-consistent snapshots.

- **Crash-consistent snapshot**: These snapshots do not include anything in memory, and they capture data that was present on the disk when the snapshot was taken. By default, crash-consistent recovery points are created every 5 minutes by Site Recovery and this time cannot be changed.
- **App-consistent snapshot**: It contains complete information of a crash-consistent snapshot as well as data in memory and transactions that are in progress. Crash-consistent snapshots are used to create app-consistent recovery points.

Replication Process

When replication for an Azure VM is enabled, the following happens:

- The Site Recovery Mobility service extension is installed on the VM automatically.
- The VM is registered with Site Recovery by extension.
- Continuous replication for the VM starts. In the source region, disk writes are sent to the cache storage account immediately.
- The data in the cache storage account is processed by Site Recovery and sent to the target storage account or replica managed disks.
- Crash-consist recovery points are generated with an interval of five minutes after the data is processed. As per the setting defined in the replication policy, app-consistent recovery points are generated.

Connectivity Requirements

Outbound connectivity is required for Azure VMs that are to be replicated and inbound connectivity is required by Site Recovery to connect with the VM. Some URLs and IP addresses should be allowed by the customers to enable proper functioning. For example, allow HTTPS outbound: port 443.

- **Control access with network security group (NSG) rules:** If VM connectivity is controlled by filtering network traffic that comes in and moves out of Azure networks using NSG rules, the following requirements should be considered:
  - For traffic replication, outbound access should be allowed by NSG rules at the source Azure region.
  - It is recommended to create rules in a test environment before they are applied to production.
  - Instead of allowing IP addresses, use service tags because they represent a collection of IP address prefixes that are gathered together to avoid complexity.

Certified Cloud Security Engineer Copyright © by **EC-Council**
All Rights Reserved. Reproduction is Strictly Prohibited.

Failover Process

When a failover is performed, VMs are created in the target resource group, target availability set, target virtual network, and target subnet. The customer can use any recovery point during the failover.

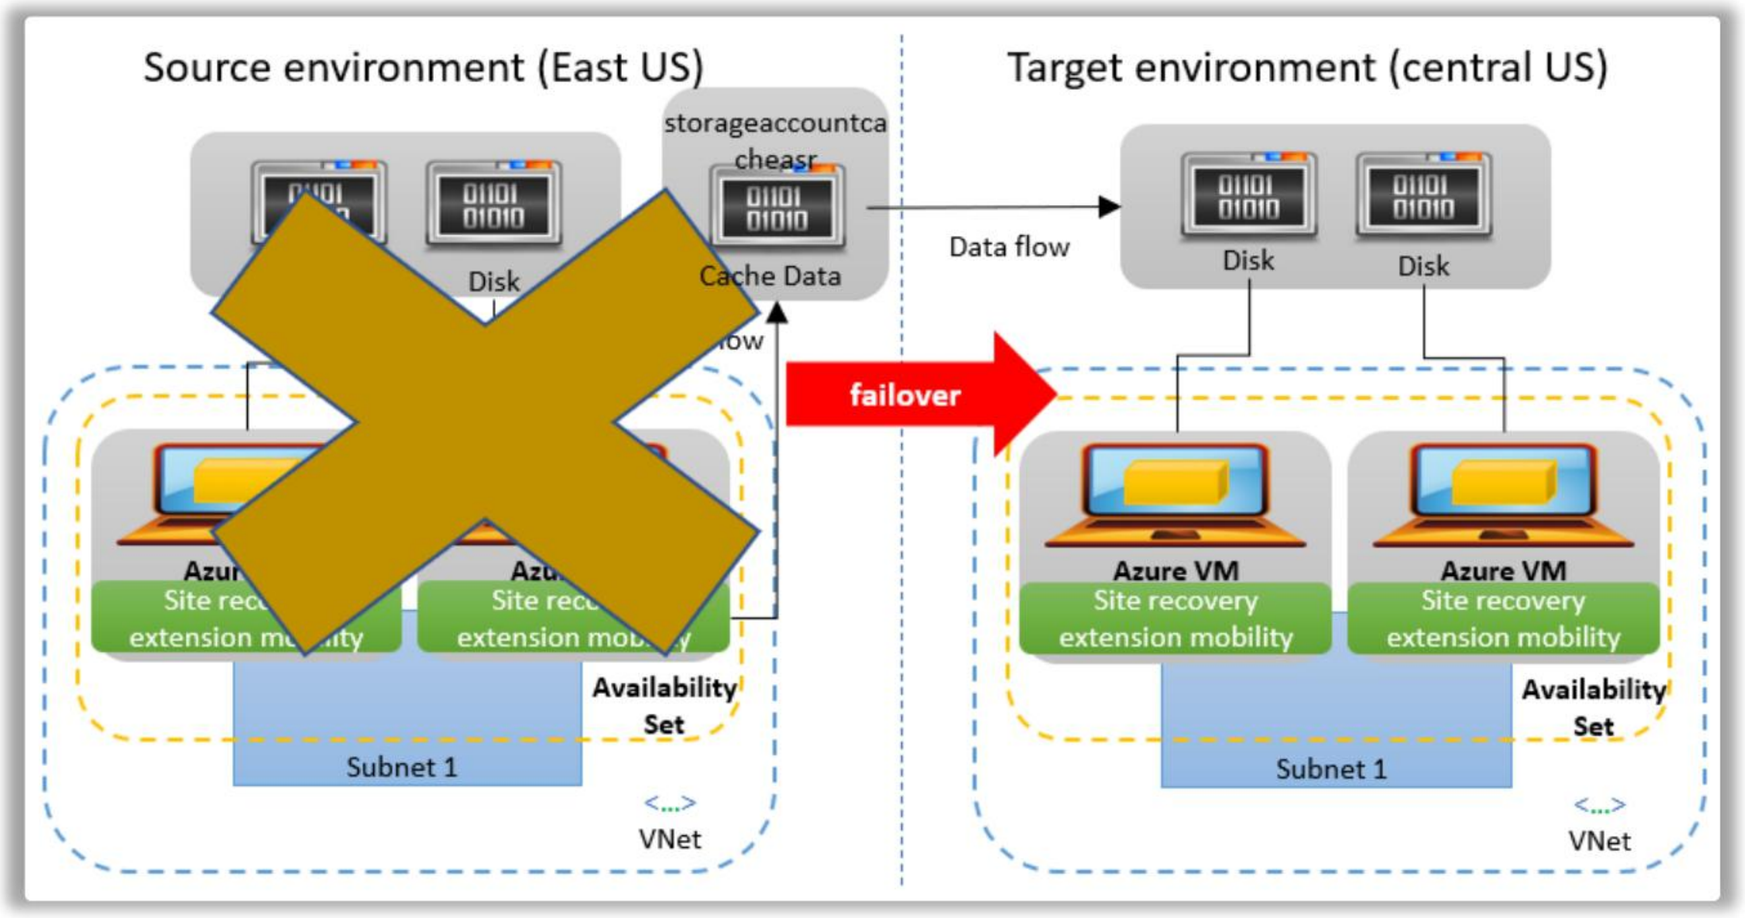

Figure 9.139: Failover Process

hide01.ir

Certified Cloud Security Engineer Copyright © by **EC-Council**
All Rights Reserved. Reproduction is Strictly Prohibited.

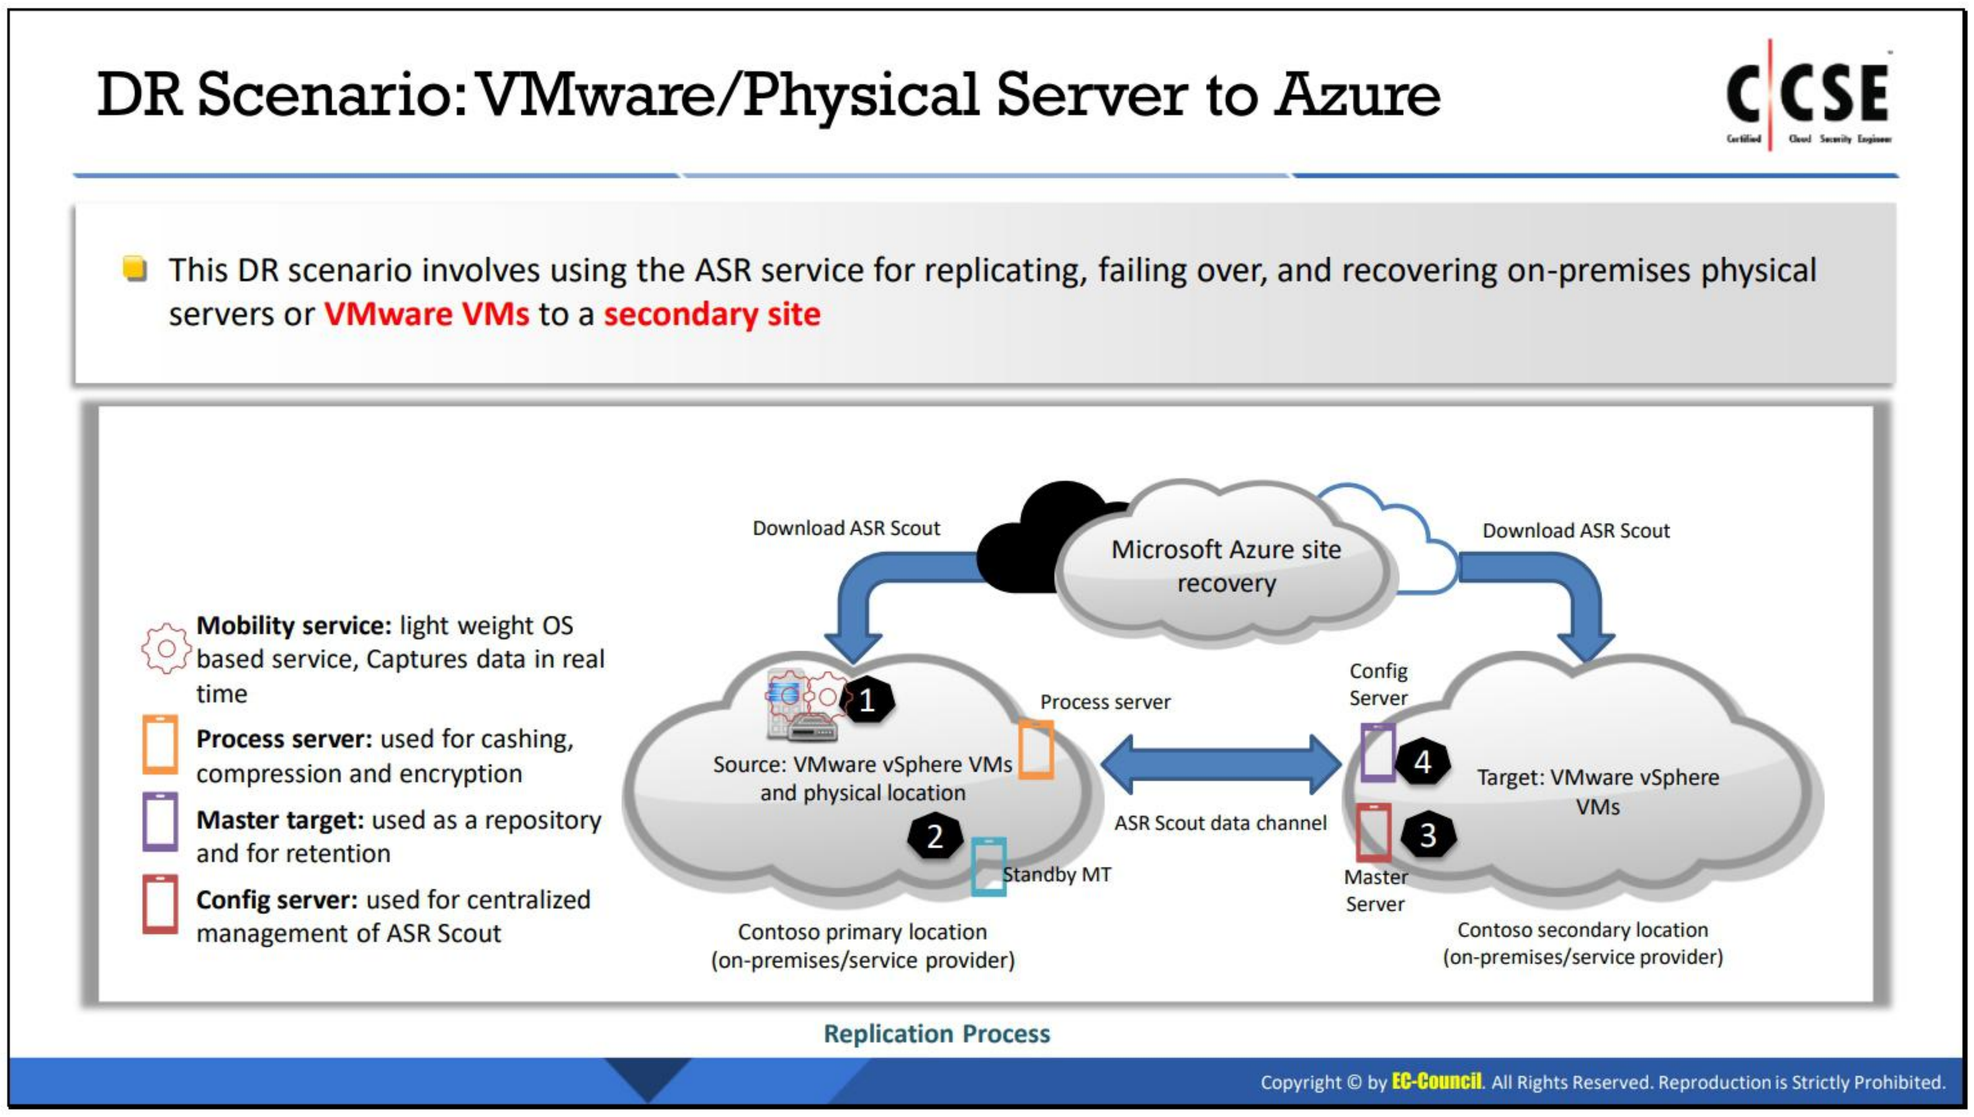

DR Scenario: VMware/Physical Server to Azure

This DR scenario involves using the ASR service for replicating, failing over, and recovering on-premises physical servers or VMware VMs to a secondary site.

hide01.ir

Architectural Components

- **Azure**: Download InMage Scout and install its latest updates for setting up deployment after creating an RSV.
- **Process server**: It is located at the primary site in this scenario and used to optimize, cache, and compress data.
- **Configuration server**: It is located at the secondary site. It manages, configures, and monitors the deployment via the vContinuum Console or the management website.
- **vContinuum server**: It is also located at the secondary site along with the configuration server. It helps in managing and monitoring the target environment by providing a console.
- **Master target server**: It helps in holding replicated data at the secondary site after receiving data from the process server. It creates a replica machine and holds the data retention points.
- **VMware ESX/ESXi and vCenter Server**: Hosting of VMs is done on ESX/ESXi hosts and a vCenter server manages these hosts.

Certified Cloud Security Engineer Copyright © by **EC-Council**
All Rights Reserved. Reproduction is Strictly Prohibited.

Outbound Network Connectivity

The customer needs to modify outbound network connectivity to enable their environment to replicate data to Azure. Certain URLs need to be allowed if a URL-based firewall proxy is used to control outbound network connectivity to ensure Site Recovery works as expected.

Replication Process

- On each site, set up the component servers (configuration server, process server, and master target server). On the machines that are to be replicated, install the Unified Agent.
- After initial replication gets completed, delta replication changes are sent by the agent present on each machine to the process server.
- The process server is responsible for optimizing the data received and transferring it to the master target server present on the secondary site. The replication process is managed by the configuration server.

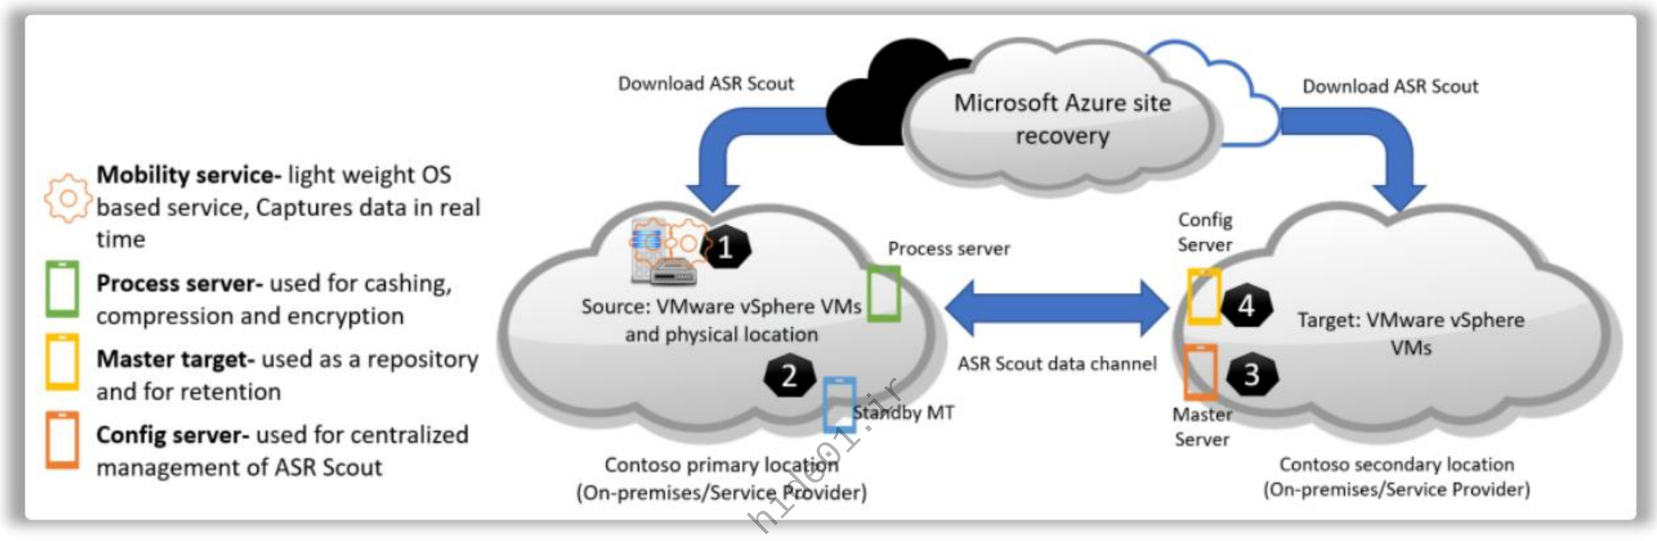

Figure 9.140: Replication Process

Certified Cloud Security Engineer Copyright © by **EC-Council**
All Rights Reserved. Reproduction is Strictly Prohibited.

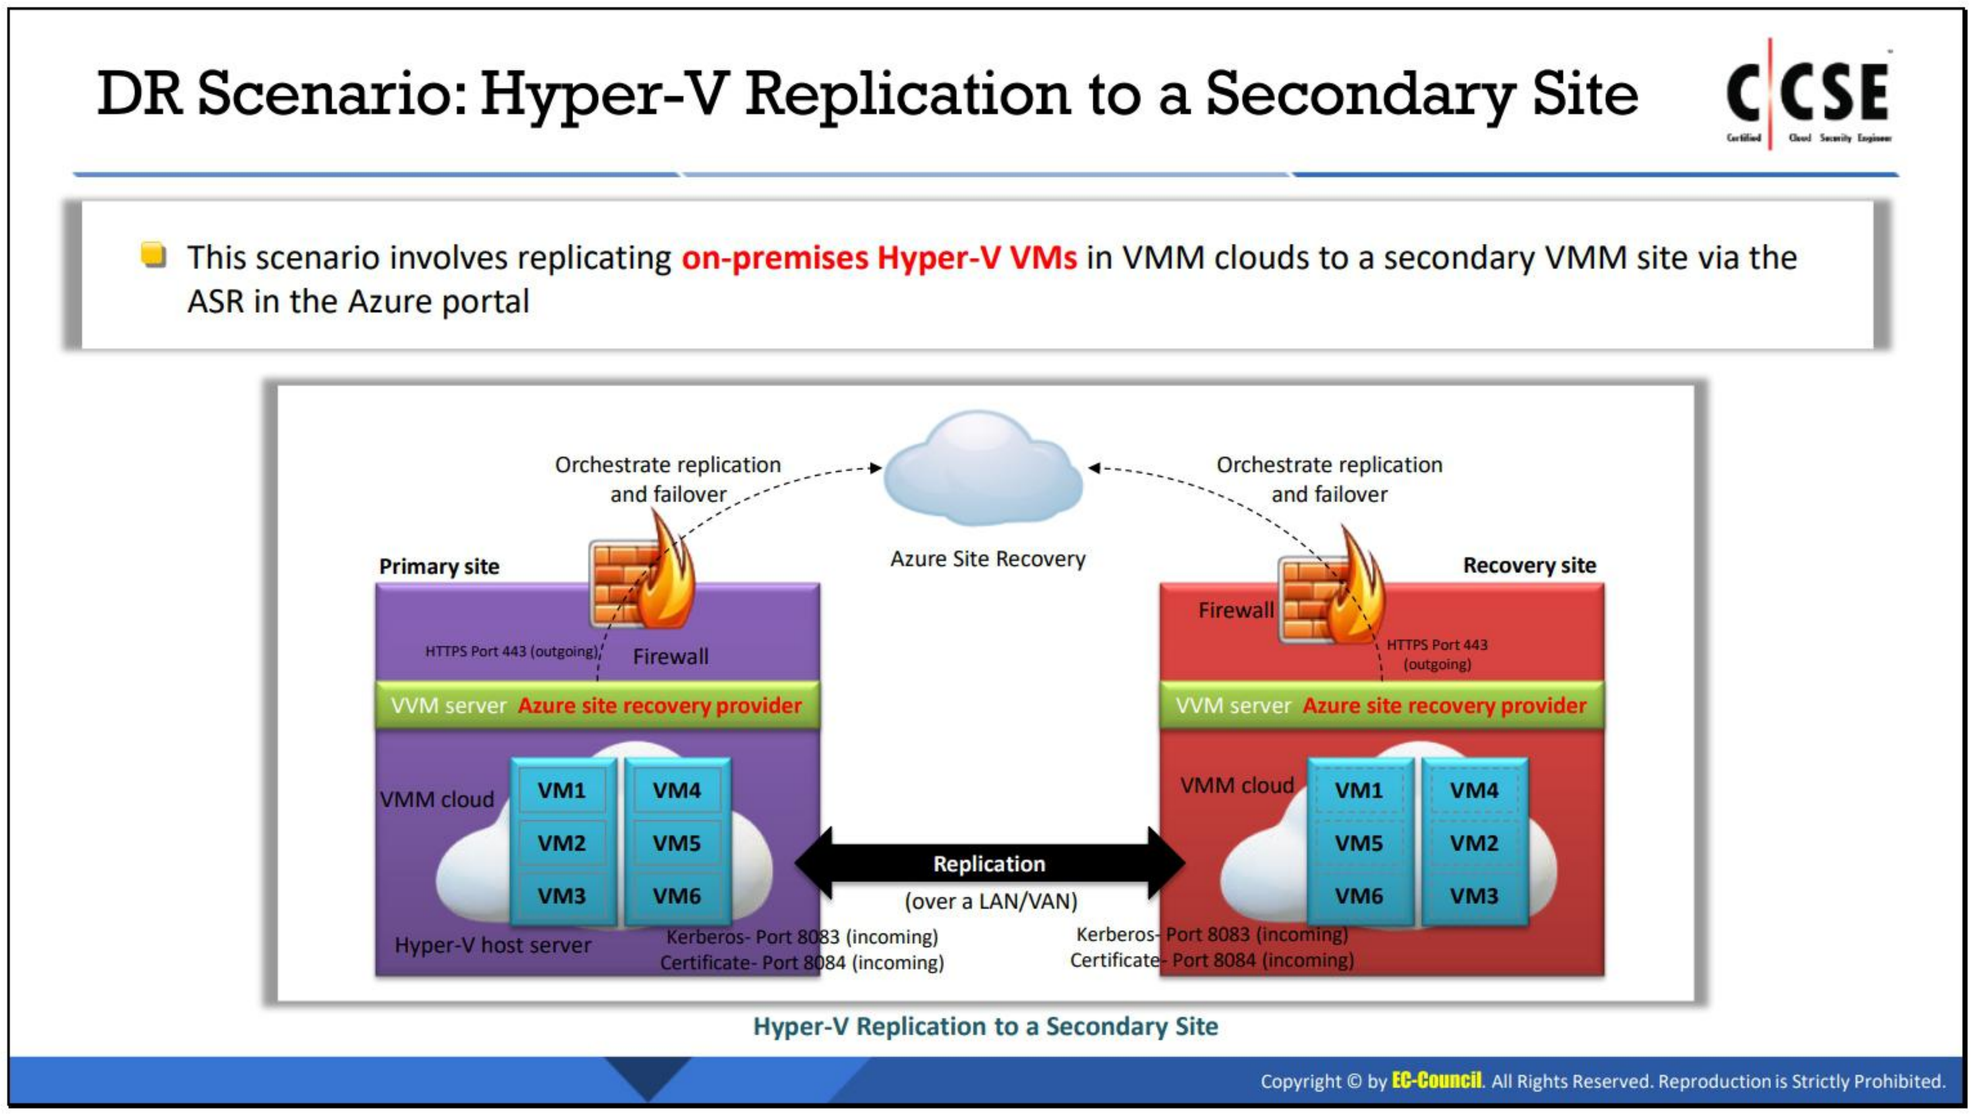

DR Scenario: Hyper-V Replication to a Secondary Site

This scenario involves replicating on-premises Hyper-V VMs in VMM clouds to a secondary VMM site via the ASR in the Azure portal.

hide01.ir

Architectural Components

- **Azure:** To organize and manage replication between two VMM locations, create a RSV in Azure.
- **VMM server:** One VMM server is required at both the primary location and the secondary location.
- **Hyper-V server:** In both primary and secondary VMM clouds, one or more Hyper-V host servers are needed because they replicate data between the two locations over the VPN or LAN via certificate authentication or Kerberos.
- **Hyper-V VMs:** At least one VM should be present at the source host server that has to be replicated.

Certified Cloud Security Engineer Copyright © by **EC-Council**
All Rights Reserved. Reproduction is Strictly Prohibited.

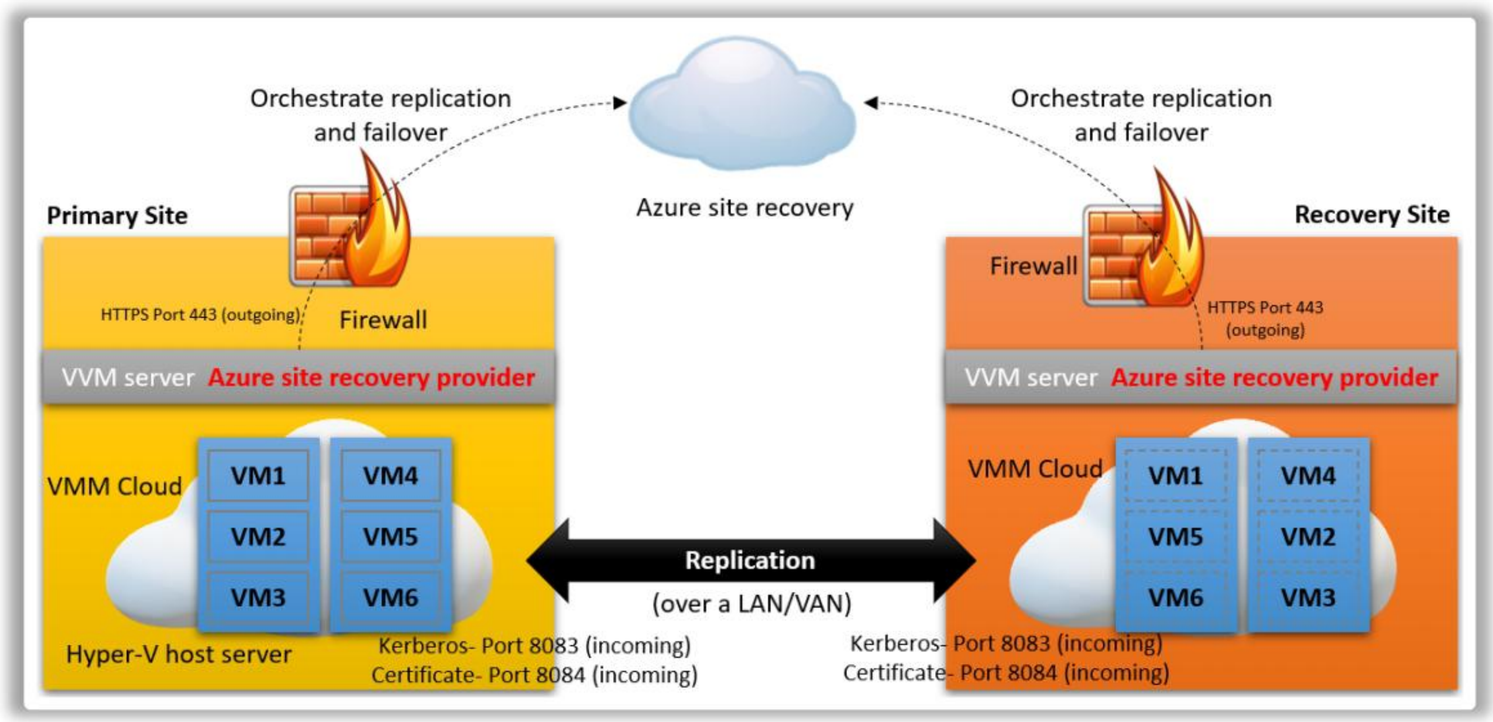

Figure 9.141: Hyper-V Replication to a Secondary Site

Replication Process

- A Hyper-V VM snapshot is taken when initial replication is set or triggered.
- Virtual disks present on the VM get replicated one by one until they all get copied to Azure.
- The Hyper-V Replica Replication Tracker tracks the disk changes that occur while initial replication is going on as Hyper-V replication logs (.hrl). All these logs are stored in the same folder where the disks are present. There is an associated .hrl file with each disk that is sent to the secondary storage. The disk resources are consumed by log files and the snapshot when initial replication takes place.
- The snapshot is deleted when the initial replication is completed.
- Delta disk changes are synchronized and added to the persistent disk.

Failover and Failback Process

- The customer can fail over one machine or create recovery plans to failover many machines.
- Between on-premises sites, a planned or unplanned failover can be performed. Source VMs are shut down to make sure no data is lost when a planned failover is performed.
    - If an unplanned failover is performed, machines in the secondary site are not protected after the failover.
    - If a planned failover is performed, machines in the secondary sire are protected after the failover.
- The customer must access the workloads from the replica VM once the initial failover runs.
- The customer can perform failback once the primary site is available again.

Certified Cloud Security Engineer Copyright © by **EC-Council**
All Rights Reserved. Reproduction is Strictly Prohibited.

- To begin replicating from the secondary site to the primary one, the customer must initiate reverse replication because it brings the VMs into a protected state. However, the secondary site remains to be the active location.
- The customer needs to start a planned failover from the secondary site to the primary site after another reverse replication to convert the primary site to the active location again.

hide01.ir

Certified Cloud Security Engineer Copyright © by **EC-Council**
All Rights Reserved. Reproduction is Strictly Prohibited.

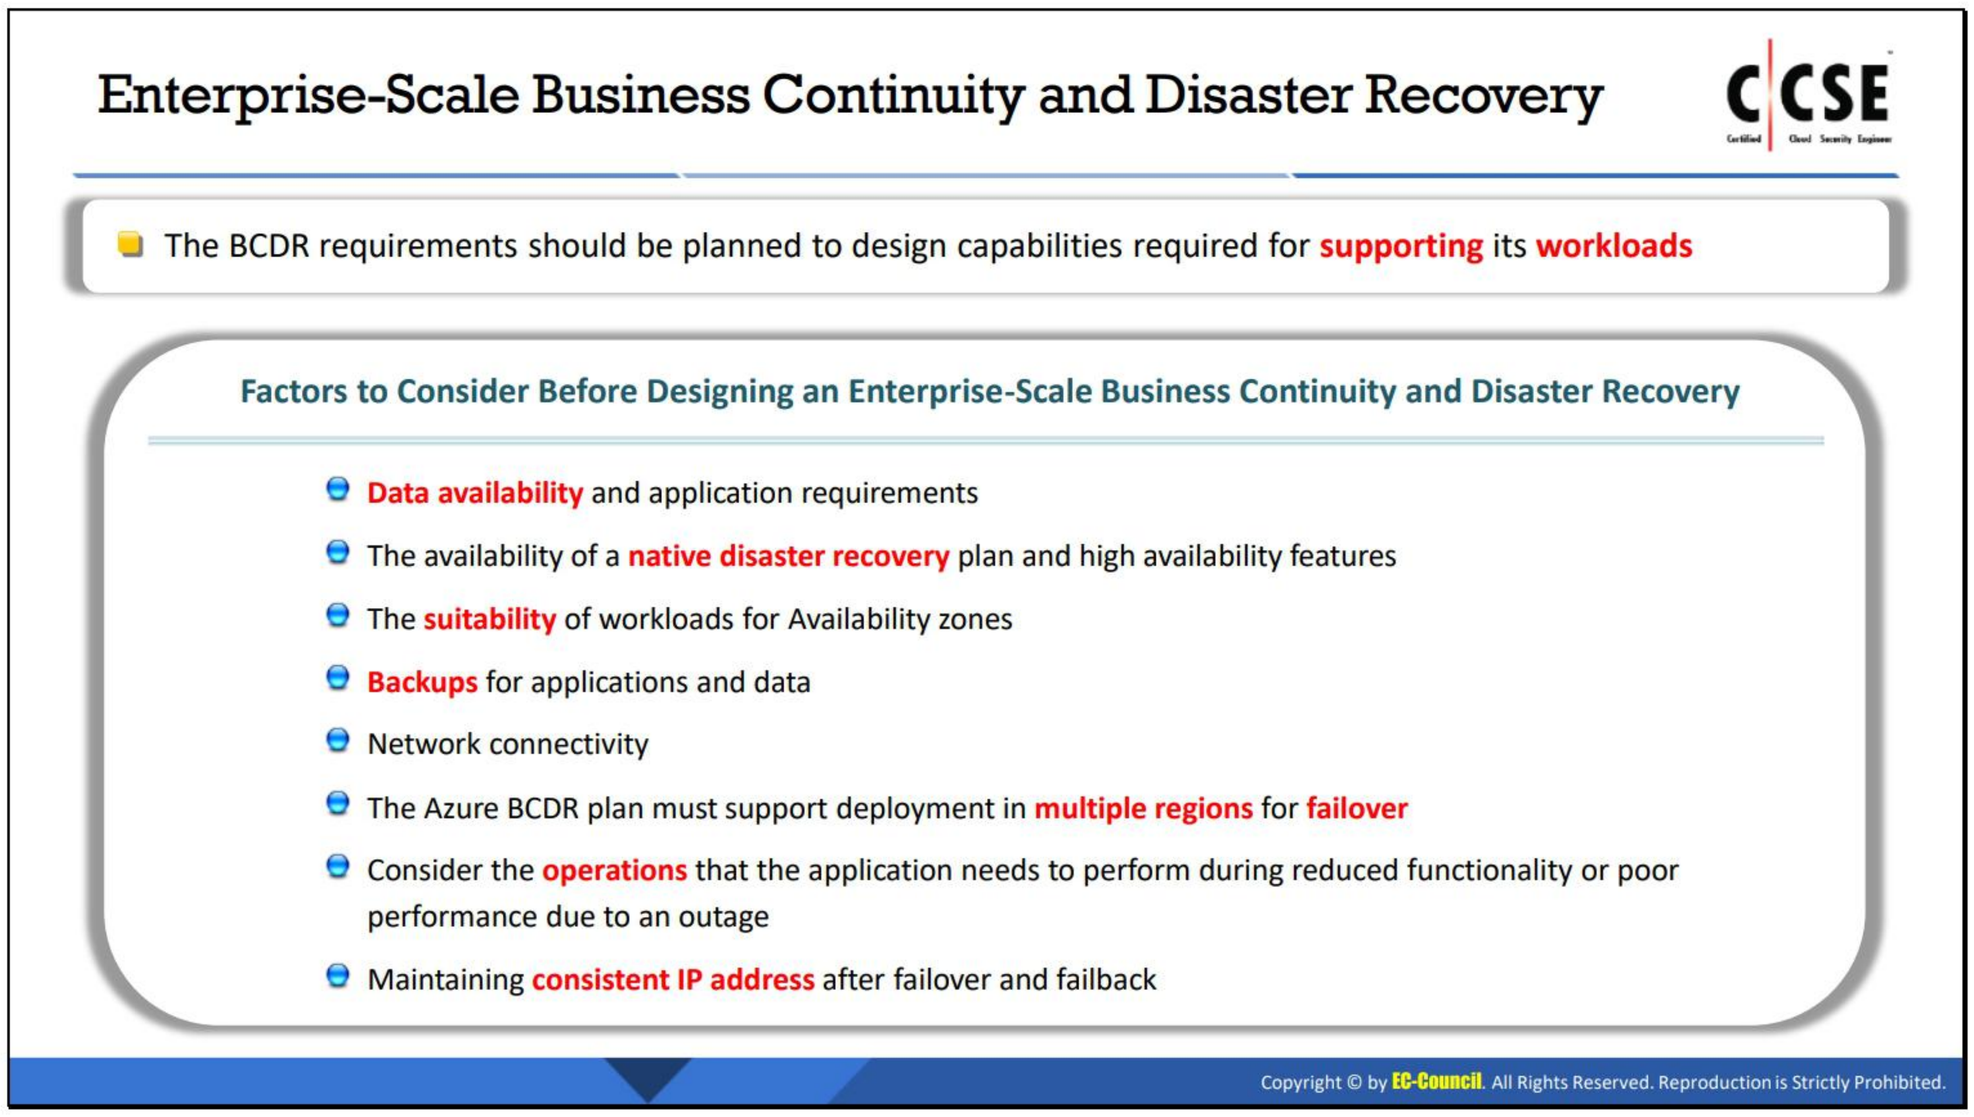

Enterprise-Scale Business Continuity and Disaster Recovery

Every organization considers designing efficient platform-level capabilities that are used by application workloads to meet specific requirements. These applications workloads have certain requirements in terms of RTO and RPO. The organization also needs a BCDR plan to protect the data and production environment. The BCDR requirements should be planned to design capabilities required for supporting its workloads.

Factors to be Considered

The following factors should be considered before designing an enterprise-scale BCDR:

- Data availability and application requirements should be considered for designing the Azure BCDR. Additionally, the usage of active-passive and active-active patterns such as RTO and RPO requirements for workloads should be considered.
- The availability of a native DRP and high-availability features should also be considered. The organization should also consider BCDR for PaaS.
- The Azure BCDR plan must support deployment in multiple regions for failover.
- Consider the operations that the application needs to perform in case of reduced functionality or poor performance due to an outage.
- The suitability of workloads for AZs should be considered such as data sharing between different zones, Azure ultra-disk storage, etc.
- Backups for applications and data should also be considered. This includes instance snapshots, using Azure Backup and service vaults, and geo-replication.

Certified Cloud Security Engineer Copyright © by **EC-Council**
All Rights Reserved. Reproduction is Strictly Prohibited.

- Network connectivity is an important factor to be considered. Network connectivity when a failover occurs such as planning bandwidth capacity for Express Route, traffic routing during an outage needs to be considered.
- Maintaining consistent IP address after failover and failback, engineering DevOps capabilities related to planned and unplanned failover should be considered.

hide01.ir

Certified Cloud Security Engineer Copyright © by **EC-Council**
All Rights Reserved. Reproduction is Strictly Prohibited.

Best Practices for Enterprise-Scale Business Continuity and Disaster Recovery (BCDR)

The following best practices should be considered for designing an enterprise-scale BCDR architecture:

- **Use ASR:** ASR should be used for Azure-to-Azure CDR cases. This helps replicate workloads across different regions. ASR provides capabilities for VM workloads to achieve lower values of RPO and RTO. This can be achieved through recovery automation and real-time replication. ASR also allows conducting recovery drills without disturbing the production workloads. The Azure policy can be used to activate replication and to audit the protection of VMs.
- **Use Native PaaS disaster recovery capabilities:** The native capabilities provide an easy solution for the design and deployment of automation. It is easy to build replication and failover into the workload environment using the built-in-features.
- **Use Azure-native backup capabilities:** Using Azure-native backup capabilities eliminates the need for configuring and maintaining third-party backup software tools and infrastructure. Using other native features, a backup configuration can be set up, audited, and enforced. This helps ensure that the services will meet business requirements.
- **Use multiple regions:** Multiple regions and peering locations for ExpressRoute connectivity should be used. This ensures cross-premise connectivity when an outage affects an Azure region or a peering location.
- **Avoid usage of overlapping IP addresses:** The BCDR network architecture should provide parallel connectivity to all websites. A DR network that uses classless inter-domain routing

Certified Cloud Security Engineer Copyright © by **EC-Council**
All Rights Reserved. Reproduction is Strictly Prohibited.

(CIDR) blocks needs a network failover process to delay application failover during an outage.

hide01.ir

Certified Cloud Security Engineer Copyright © by **EC-Council**
All Rights Reserved. Reproduction is Strictly Prohibited.

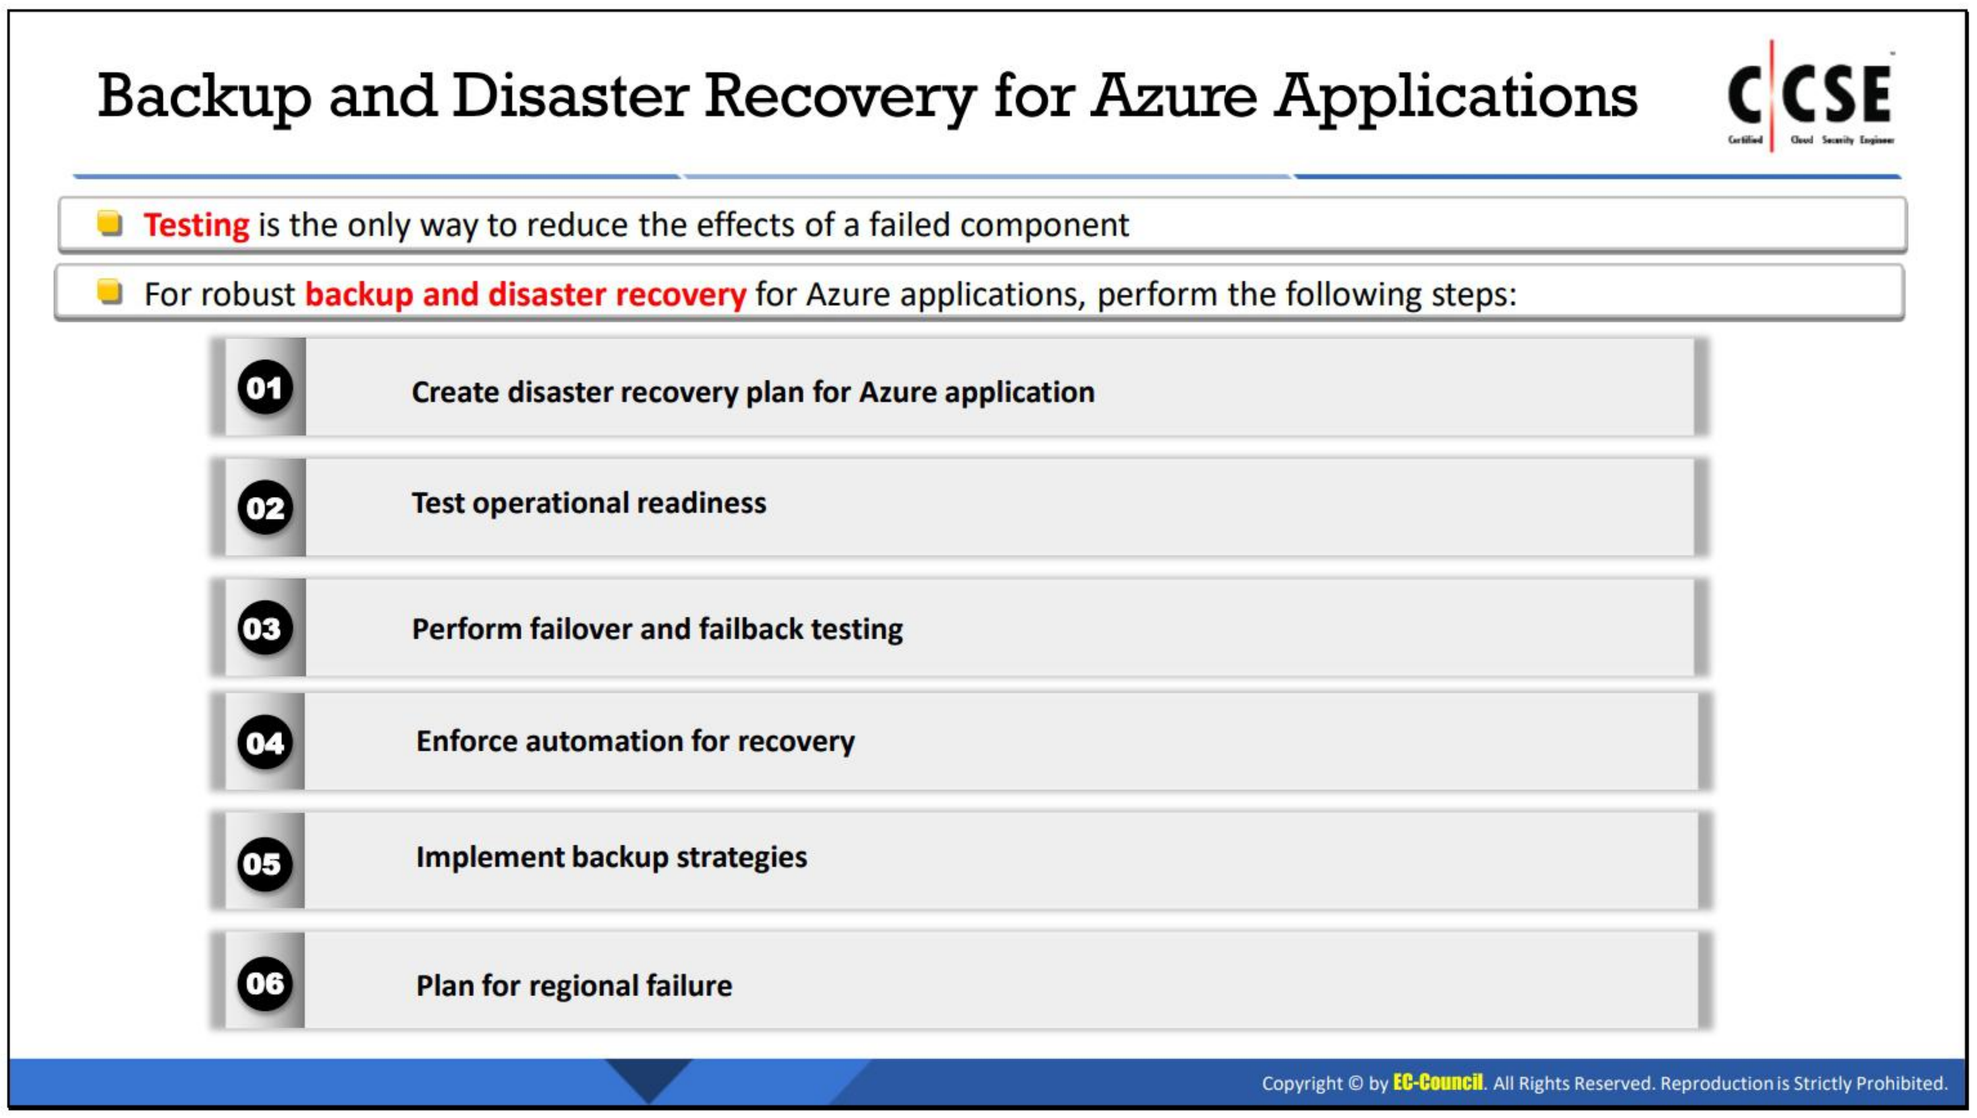

Backup and Disaster Recovery for Azure Applications

Disaster recovery for Azure applications involves the process of restoring the functionality of the applications after a disaster or loss. Though Azure offers reliable services, failures may happen. Instead of focusing on preventing the failures, the focus should be on reducing the effects of a failed component.

hide01.ir

Testing is the only way to reduce the effects of a failed component. The testing can be automated; however, it is important to be prepared when the applications fail. Backup and recovery methods are essential and important during application failures.

For a few applications, reduced functionality or partial unavailability of the application may be acceptable. For some applications, even a minute reduction in the functionality might not be acceptable.

Disaster Recovery Plan (DRP) for Azure Applications

In the first step, a recovery plan should be created. After testing the plan from all the corners, it should be considered as complete. People, processes, and applications should be included in the plan to restore the functionality of the application as per the defined SLA.

The following suggestions should be considered for creating and testing the DRP:

- The process for contacting the support team and escalation issues should be included. This helps reduce the chances of a prolonged downtime.
- The impact of application failures on the business should be evaluated.
- For mission-critical applications, a cross-region recovery architecture should be chosen.
- An owner should be assigned for a DRP, automation, and testing.

Certified Cloud Security Engineer Copyright © by **EC-Council**
All Rights Reserved. Reproduction is Strictly Prohibited.

- The process should be documented including the manual steps.
- A backup strategy should be established for transaction data, and backup restoration should be tested on regular basis.
- Alerts should be set for Azure services used by the application.
- The operation staff should be trained to execute the DRP.
- Regular disaster simulations should be performed to improve and validate the plan.

If ASR is used to replicate VMs, a fully automated recovery plan should be created to failover the complete application.

Operational Readiness Testing

An operational readiness test should be performed for failback to the primary region and failover to the secondary region. Most of the Azure services support test failover and manual failover to conduct drills for disaster recovery. Apart from that, an outage can be simulated by removing the Azure services or by shutting down. As a part of the application lifecycle, automated operational responses should be tested regularly.

Failover and Failback Testing

Failover and failback should be tested to ensure that the dependent services are back in a synchronized way during disaster recovery. Though the changes to system and operations might affect the failover and failback functions, they might not be identified till the main system gets overloaded or fails. The users should ensure that the dependent services should failover and failback in the correct manner.

If ASR is used for replicating the VMs, disaster recovery drills should be performed periodically to validate the replication strategy. The ongoing VM replication and production environment should not be affected by a test failover.

Understand the Implication of Outages

Users must understand the result of service disruption for all the dependent services. They must also be aware of how the applications will respond to those service disruptions. Most Azure services offer features for resilience and evaluating each service helps improve the DRP. Similarly, many Azure services provide resilience features and evaluation of each service would enhance the DRP. When a few components of the Azure network become inaccessible, the users may not be able to access the application or data. In those situations, the DRP should be designed to run the applications with less functionality. If reducing the functionality does not help, other options such as application downtime or failover to other regions can be used.

Recovery Automation

The steps for recovering the application or failover to a secondary region should be codified. It can be codified using automation to ensure that it can respond to an outage to reduce the impact. Codified steps should also exist to capture the failback application process to the primary region after addressing the failover.

Certified Cloud Security Engineer Copyright © by **EC-Council**
All Rights Reserved. Reproduction is Strictly Prohibited.

Tools used for failover orchestration should be considered in the failover strategy while automating the failover procedures.

Backup Strategies

There are various alternative strategies for implementing distributed compute across different regions. These strategies must adapt to the business requirements.

All such strategies can be divided into the following categories at a high level:

- **Redeploy on disaster:** In this method, there is redeployment of the application from the beginning in the event of a disaster. This approach is suitable for noncritical applications and applications that do not require a guaranteed recovery time.
- **Warm spare:** In this method, a secondary hosted service is created in another region. Roles are assigned to that region to handle minimal capacity; however, they do not receive any network traffic related to production. This approach is suitable for applications that are not designed to distribute traffic in various regions.
- **Hot spare:** In this approach, the application is designed in such a way that it receives production traffic in multiple regions. The services in each region are configured with a higher capacity for disaster recovery. The cloud services also scale-out when required during disaster and failover. The hot spare approach requires investment for designing the application. It offers huge benefits such as guaranteed and low recovery time, frequent testing of the recovery locations, and proper usage of capacity.

hide01.ir

Handling Regional Failures

Azure is divided into various units known as regions. Each region has data centers and some regions have AZs. The AZs can be used for resilience during an outage in a single data center. Hence, regions with AZs should be used to provide high availability.

In rare cases, the services and facilities in an AZ or a region may become inaccessible due to network failure. The facilities may also be permanently lost due to a natural disaster.

Microsoft Azure can create applications that can be distributed across different regions and zones. This distribution helps reduce the effect of a region or zone failure on other zones or regions.

Certified Cloud Security Engineer Copyright © by **EC-Council**
All Rights Reserved. Reproduction is Strictly Prohibited.

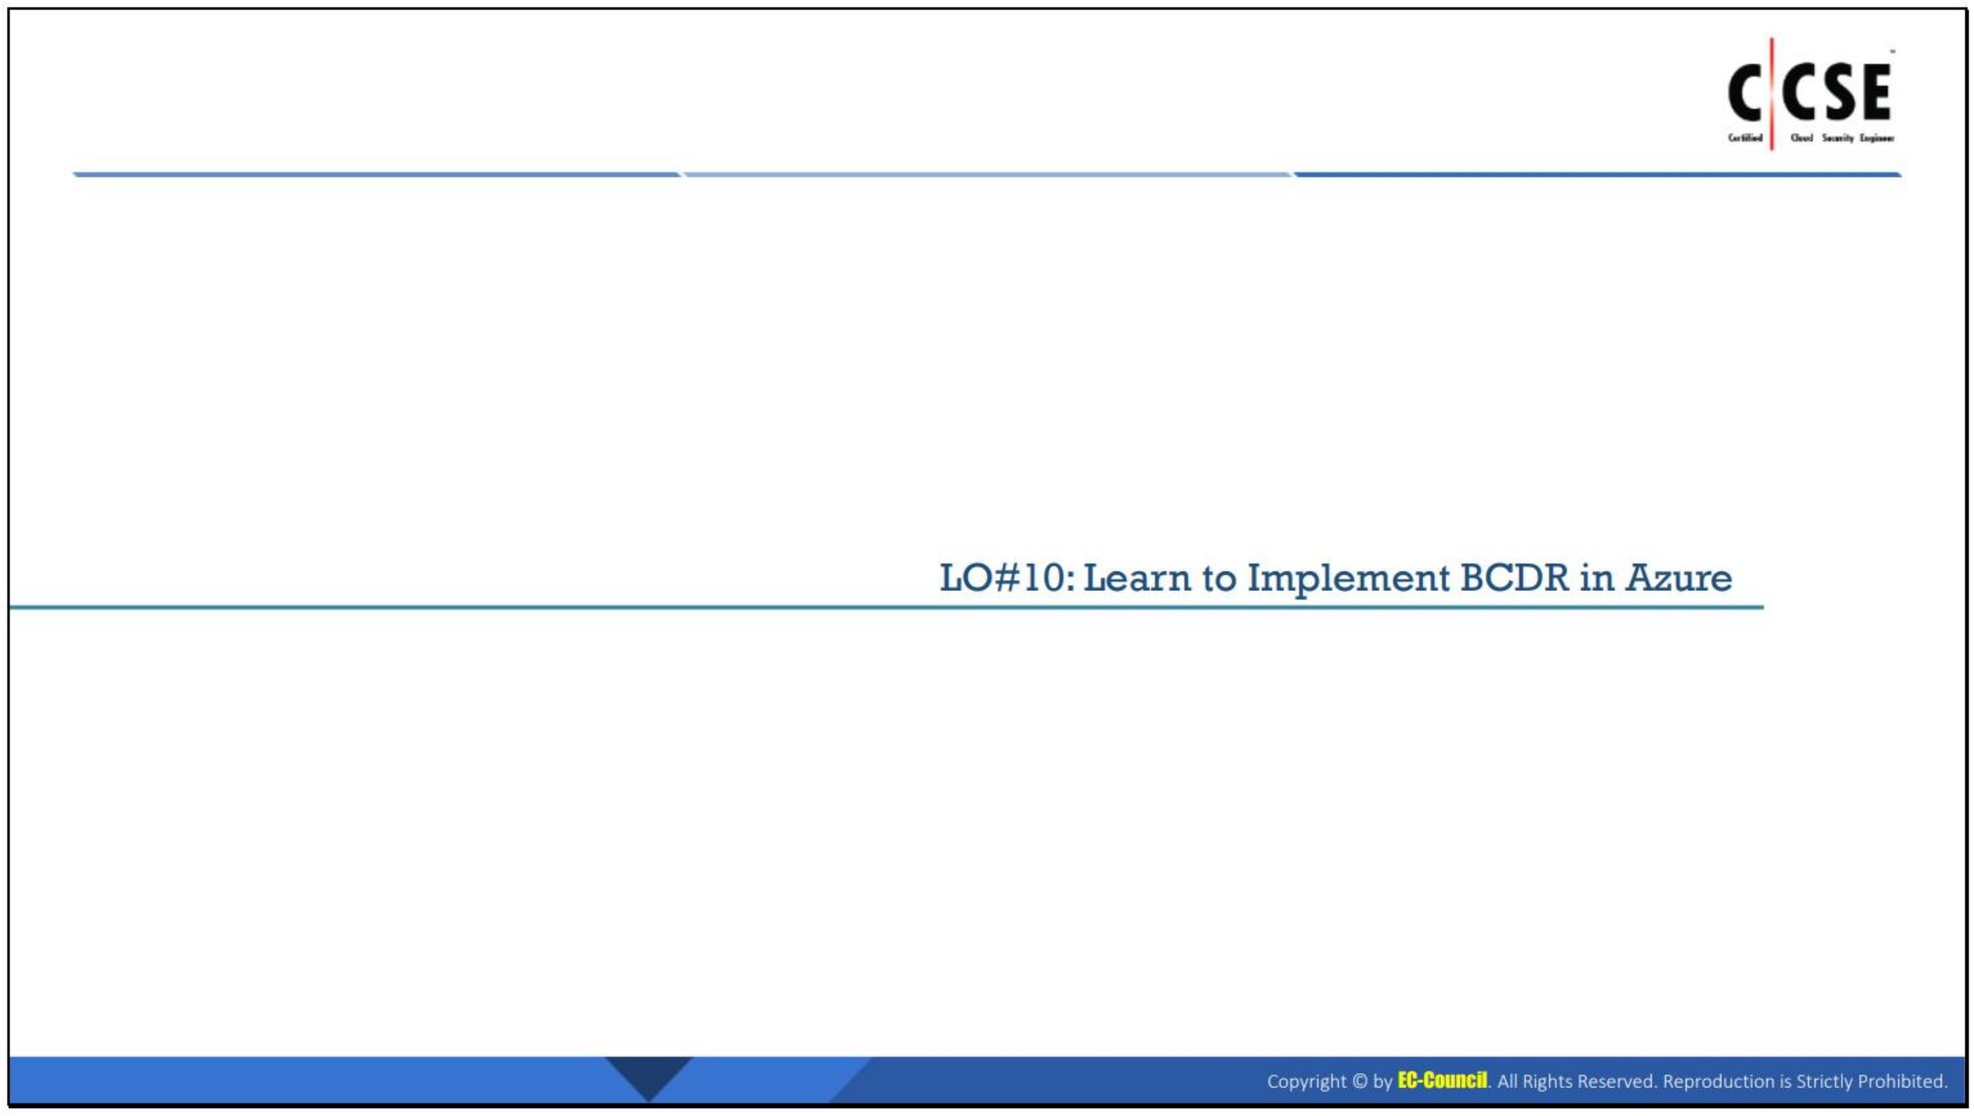

LO#10: Learn to Implement BCDR in Azure

The objective of this section is to understand the implementation of BCDR in Azure. This section provides information about performing BCDR activities using Azure features and other partner solutions. In this section, cloud security engineers will learn how to practically implement BCDR in Azure.

hide01.ir

Certified Cloud Security Engineer Copyright © by **EC-Council**
All Rights Reserved. Reproduction is Strictly Prohibited.

Enable Automated Backups

- To follow security and regulatory requirements, the cloud security professional should ensure that **App Services** in the Microsoft Azure account should be **configured** with a sufficient daily backup retention period for scheduled/automated backups

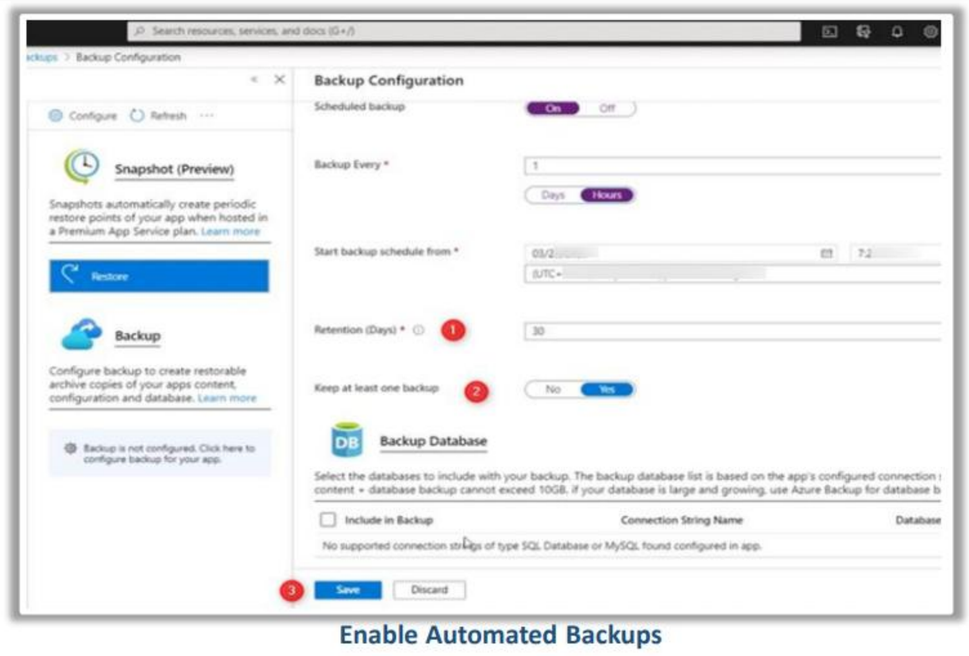

Enable Automated Backups

Copyright © by EC-Council. All Rights Reserved. Reproduction is Strictly Prohibited.

Enable Automated Backups

To follow security and regulatory requirements, ensure that App Services in the Microsoft Azure account are configured with a sufficient daily backup retention period for scheduled/automated backups. As a best practice, configuring optimal backup retention period for Azure App Services applications will enforce backup strategy. This also proves that the organization complies with its regulations. For effective handling of app data, the retention of application backups can be configured for a longer period.

This backup helps against accidental deletion or corruption by providing a facility to restore to the previous point. You can configure backups in such a way that it creates restorable copies of configuration files, app content, and database connected to that application. Once the Backup and Restore feature is enabled, then the selected applications are backed-up according to the schedules. Whenever required, the applications can be restored to previous points. Application backups can also be configured to retain up to unlimited time.

To Enable Automated Backups

- Sign in to the Azure portal.
- Click on **More services**.

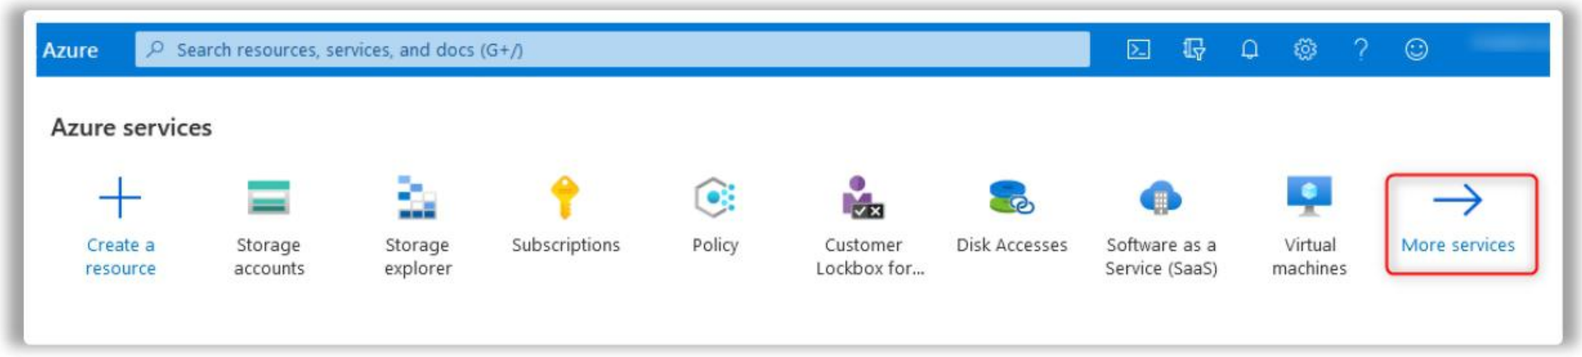

Figure 9.142: Microsoft Azure Console

Certified Cloud Security Engineer Copyright © by **EC-Council**
All Rights Reserved. Reproduction is Strictly Prohibited.

- In the left navigation panel select **Web** and then click on **App services** on the right-side panel.

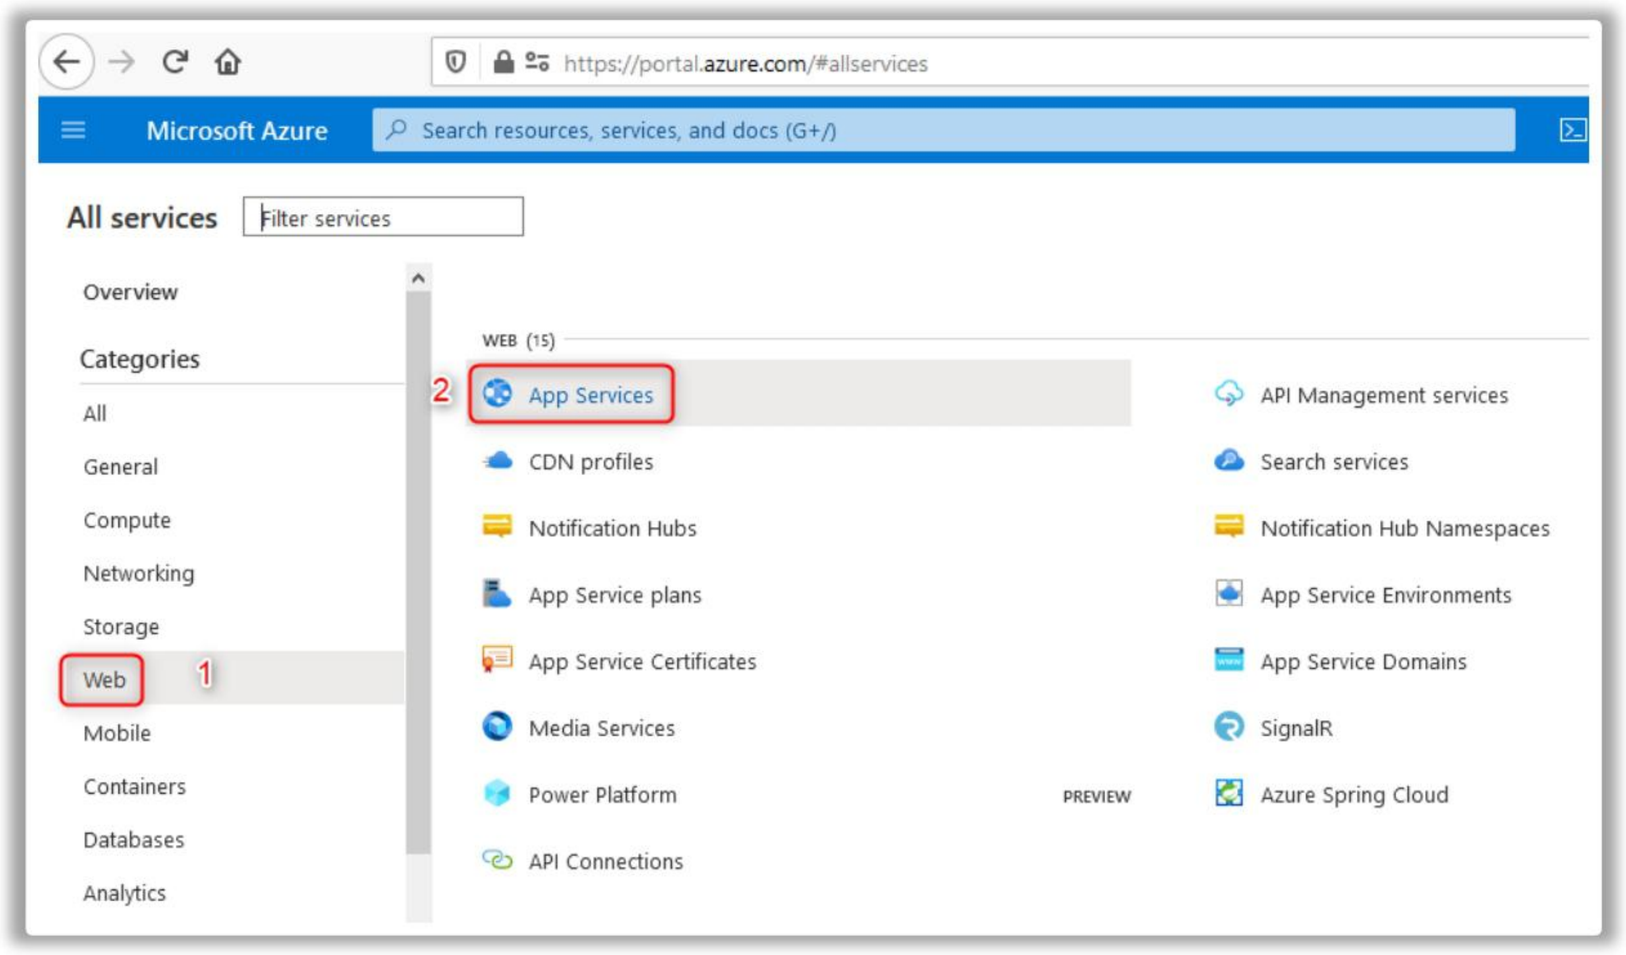

Figure 9.143: Click App Service

- A list of available App services will be displayed; select those for which you want to enable automated backups.
- In the left navigation panel, locate **Settings** and then go to **Backups** and select it.
- Click on **Configure** to set up the backup retention period.

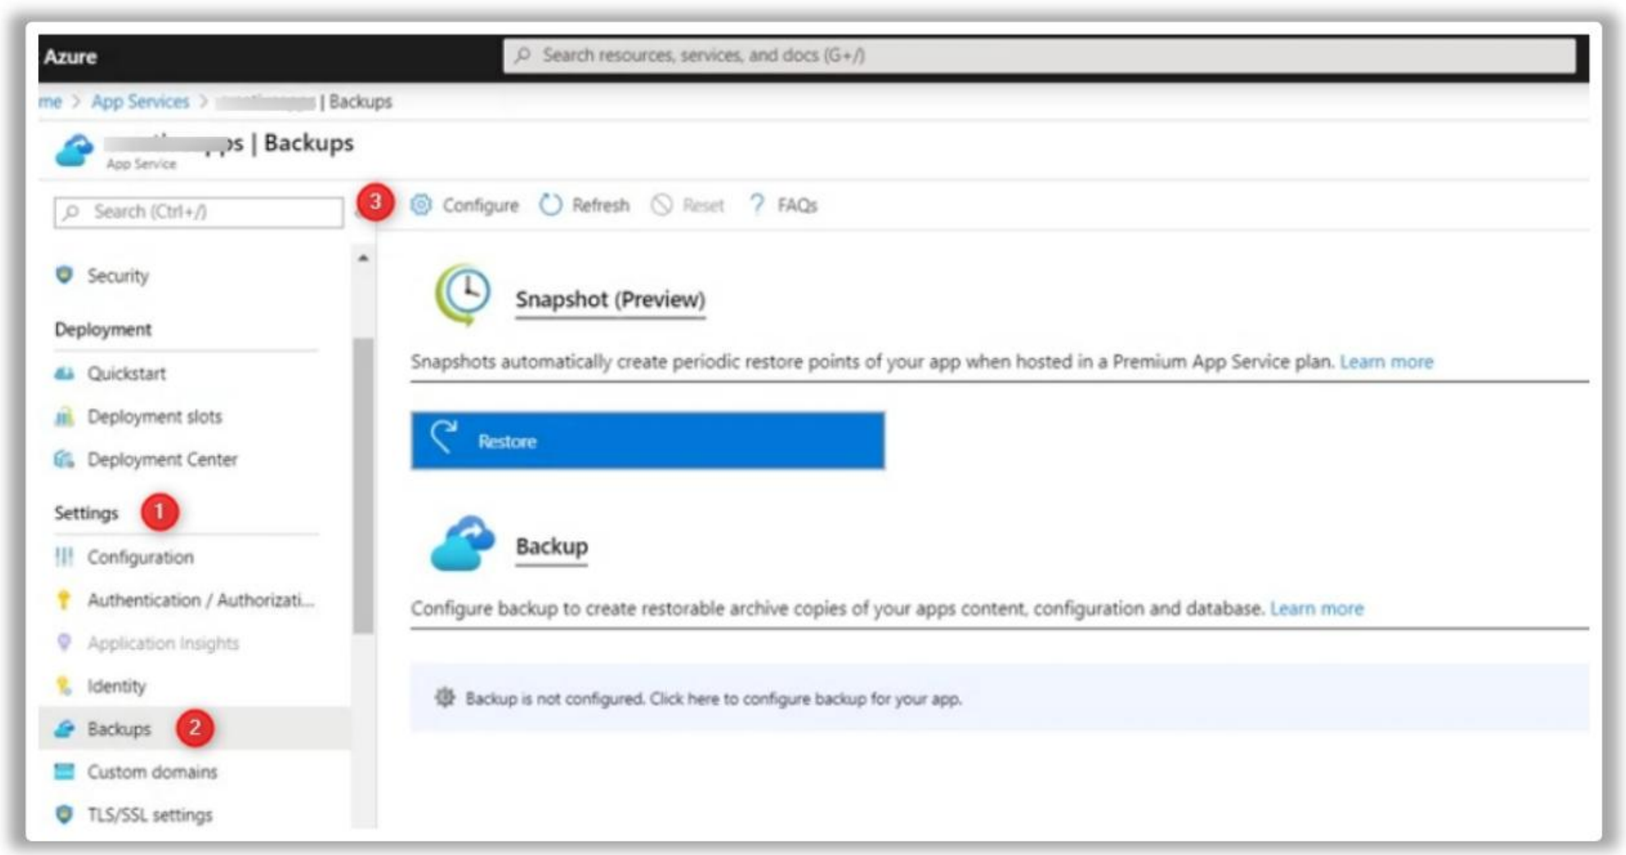

Figure 9.144: Microsoft Azure Backup Console

Certified Cloud Security Engineer Copyright © by **EC-Council**
All Rights Reserved. Reproduction is Strictly Prohibited.

- Click on **Storage Settings** to select storage account.

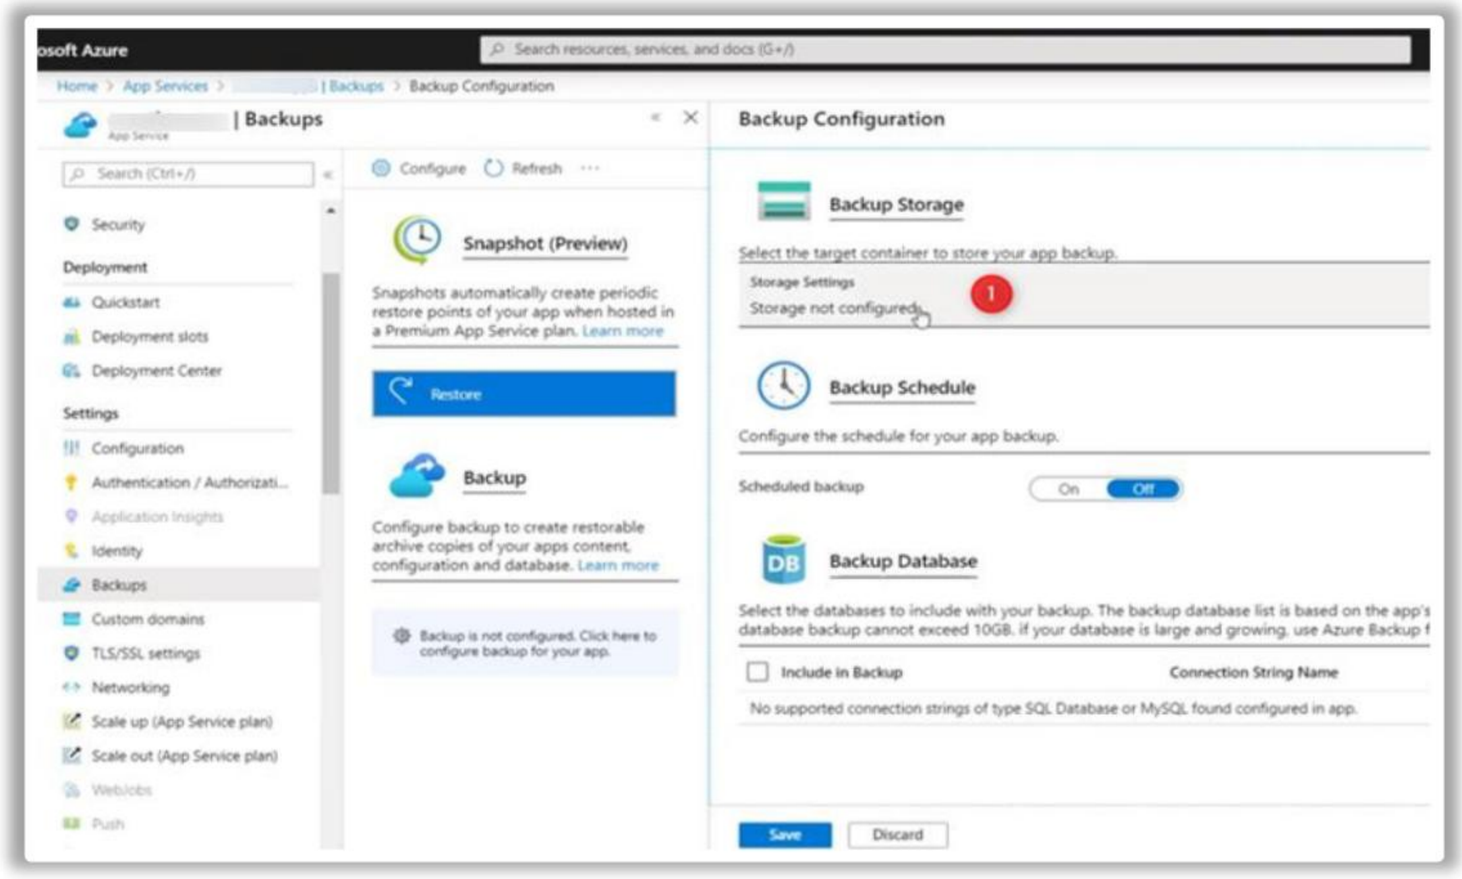

Figure 9.145: Backup Setting Console

- A list of available storage accounts is displayed, select the desired account.

Figure 9.146: Microsoft Azure Console

- Now select the containers from the available list.
- Then click on **Select**.

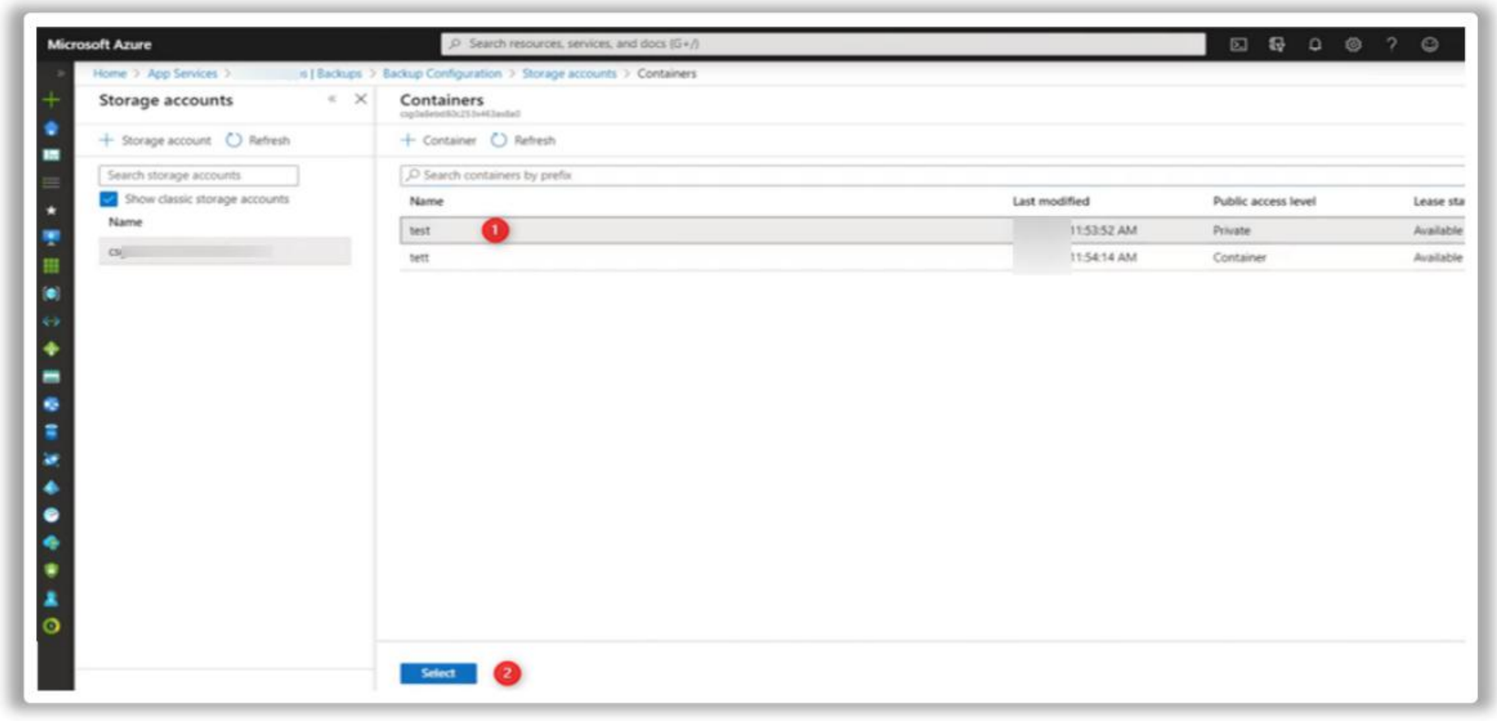

Figure 9.147: Container Console

- Under the **Backup schedule**, turn ON the **Scheduled backup.**

Certified Cloud Security Engineer Copyright © by **EC-Council**
All Rights Reserved. Reproduction is Strictly Prohibited.

- Select the **Backup Every*** to either days or hours and configure the number as per the requirement of the organizational policy.
- Verify date and time in the field of **Start backup schedule from**.

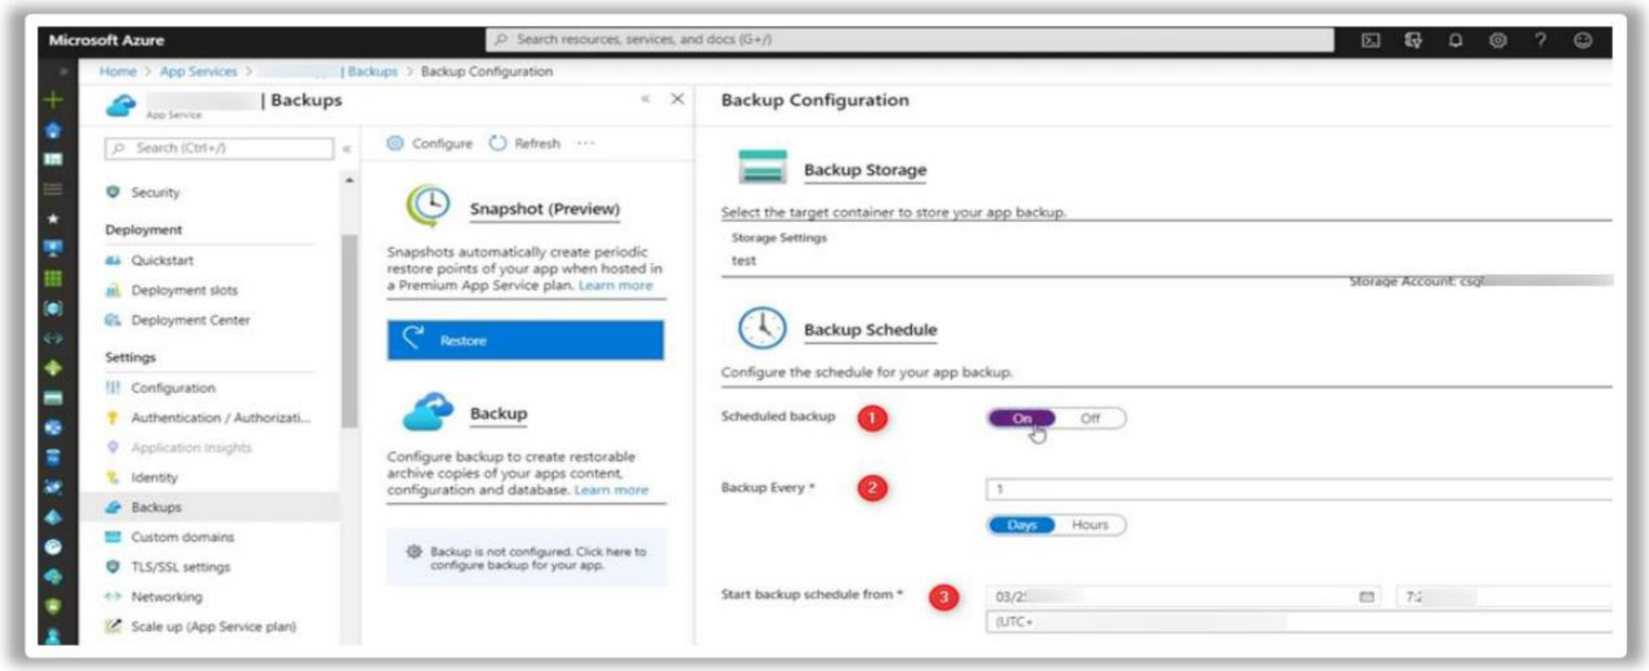

Figure 9.148: Backup Configuration

- Select the **Retention (Days)** to the required value, by default it is 30.
- Select the **Keep at least one backup** as per requirement.
- Under the **Backup Database** select databases if you want to include them in the app backups.
- Click on **Save.**

hide01.ir

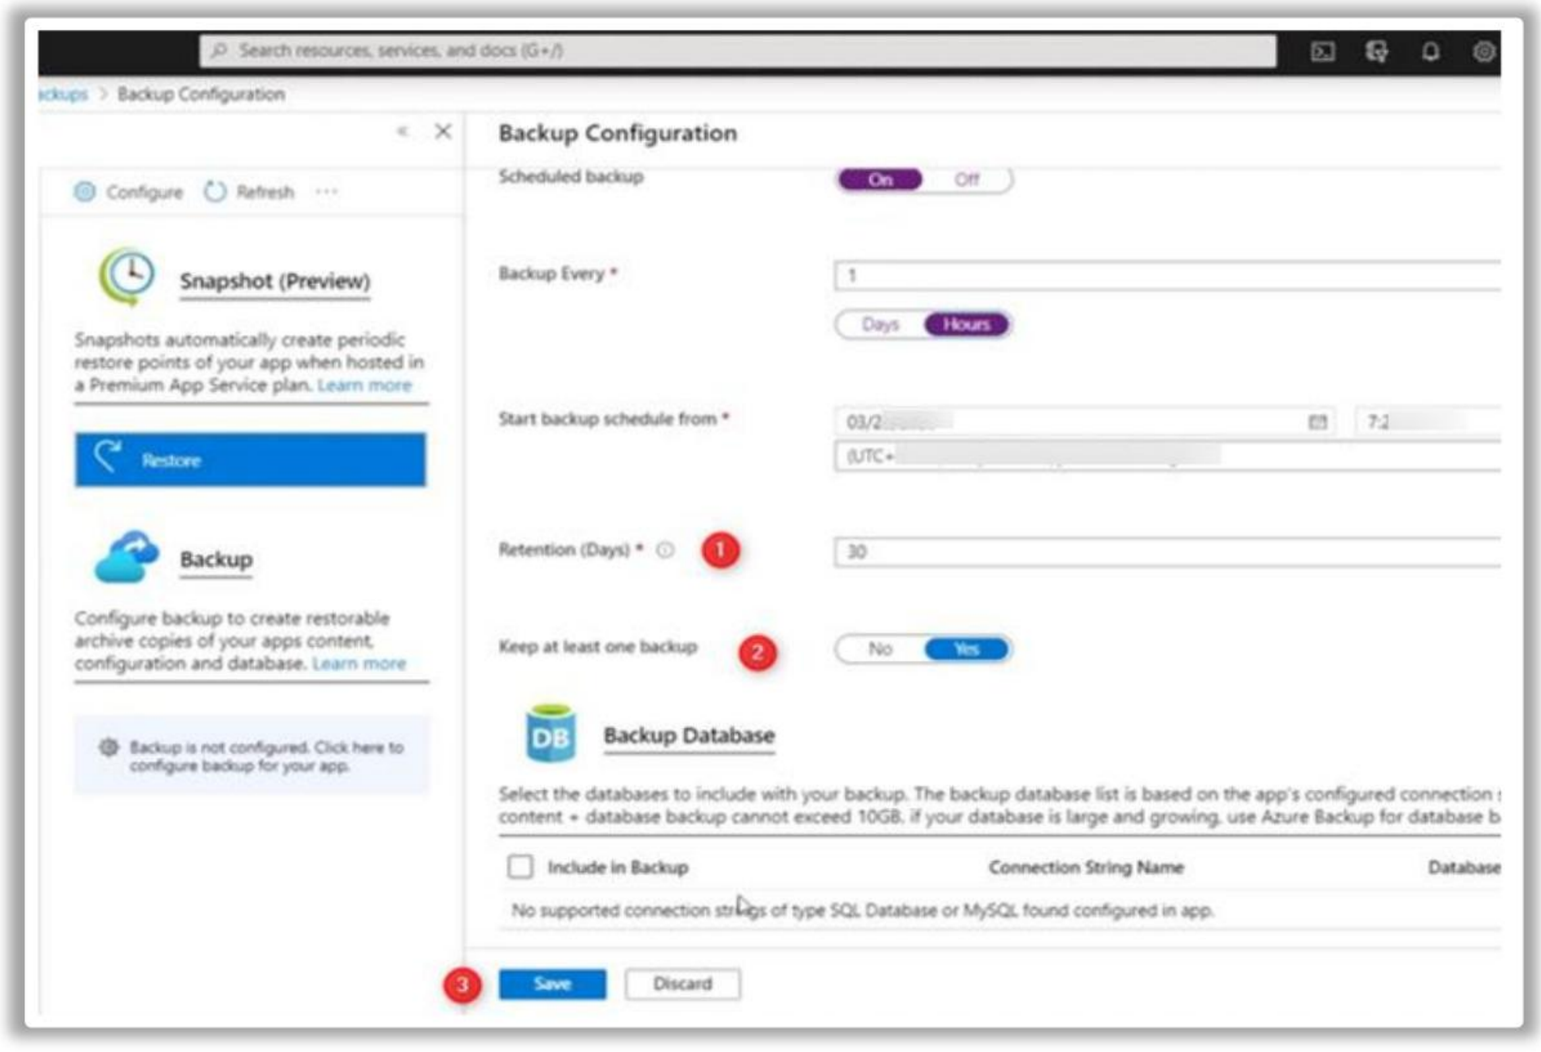

Figure 9.149: Successful Backup Configuration

- The backup retention period has been successfully configured.

Certified Cloud Security Engineer Copyright © by **EC-Council**
All Rights Reserved. Reproduction is Strictly Prohibited.

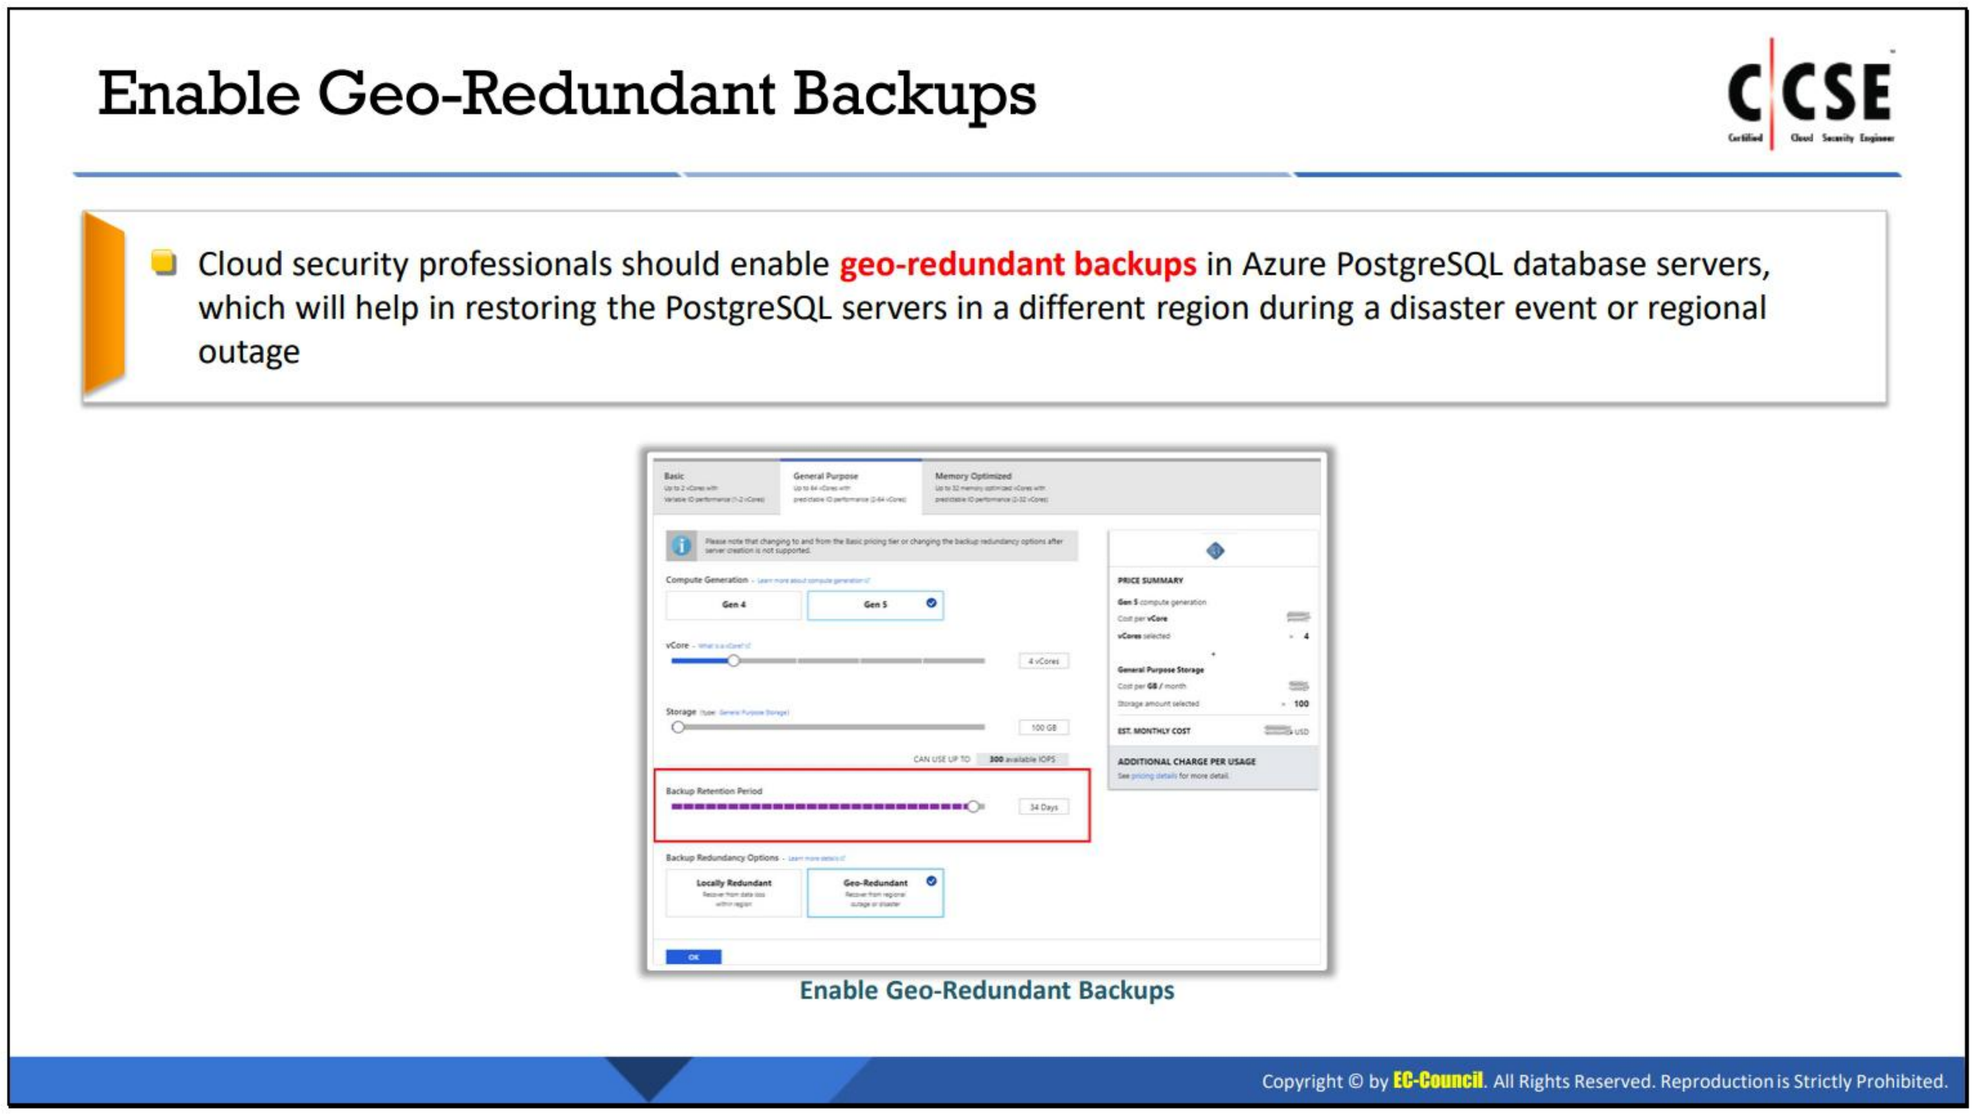

Enable Geo-Redundant Backups

Enable Geo-Redundant Backups

Cloud security professionals should enable geo-redundant backups in Azure PostgreSQL database servers, which will help in restoring the PostgreSQL servers in a different region during a disaster event or regional outage.

In event of PostgreSQL database server unavailability due to natural disaster, the default recovery option is geo-restore. Up to 4 TB of storage can be restored to its geo-paired region and up to 16 TB of storage to any other regions that support it. The PostgreSQL server configuration can be changed during geo-restore. The changes include vCore, compute generation, backup redundancy options, and backup retention period.

To Enable Geo-Redundant Backups

- Sign in to the Azure portal.
- Click on **More services**.

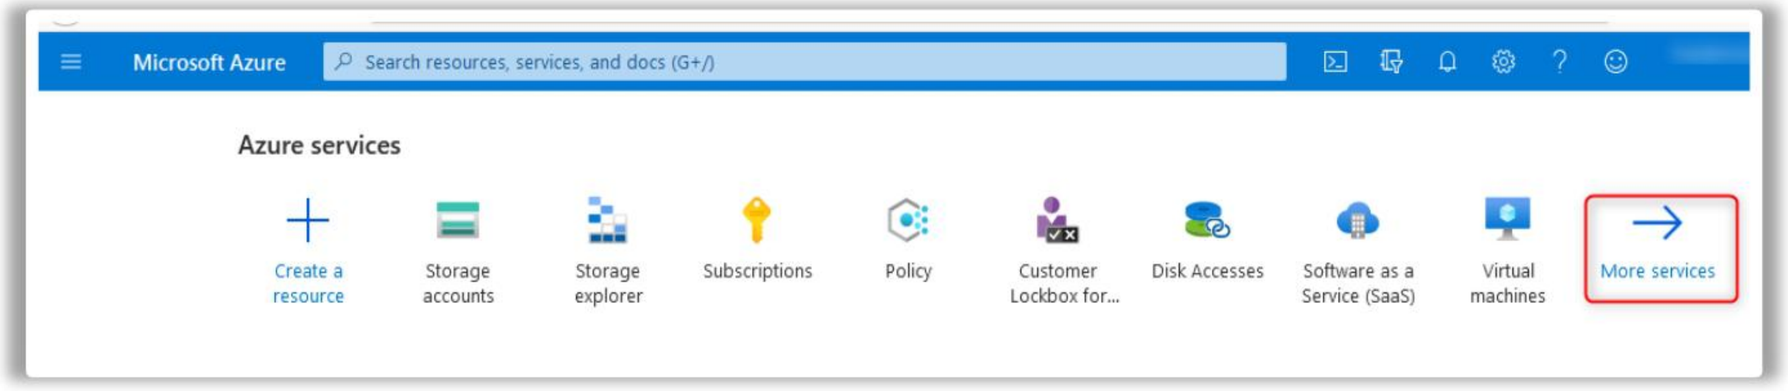

Figure 9.150: Microsoft Azure Console

Certified Cloud Security Engineer Copyright © by **EC-Council**
All Rights Reserved. Reproduction is Strictly Prohibited.

- Search for **Azure Database for PostgreSQL server** and select the same in the right-side panel.

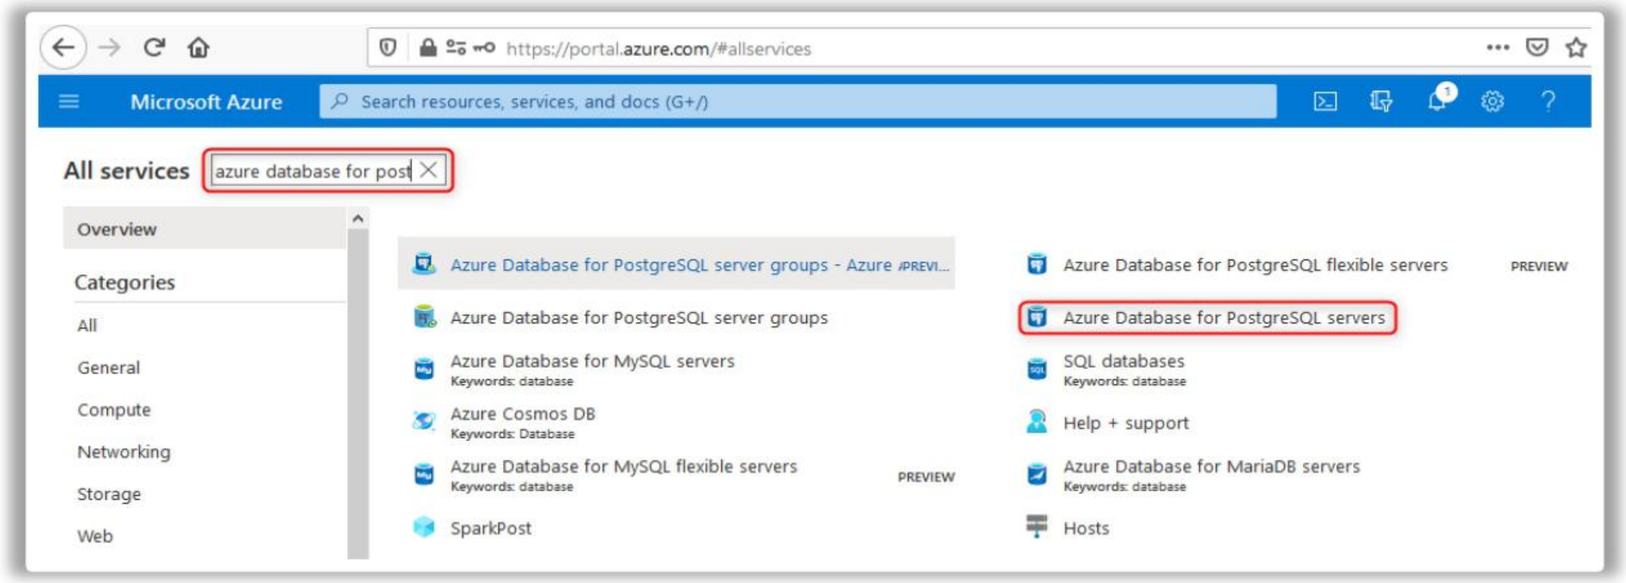

Figure 9.151: Microsoft Azure All Service Console

- Click on **+ Add** to create a new server.

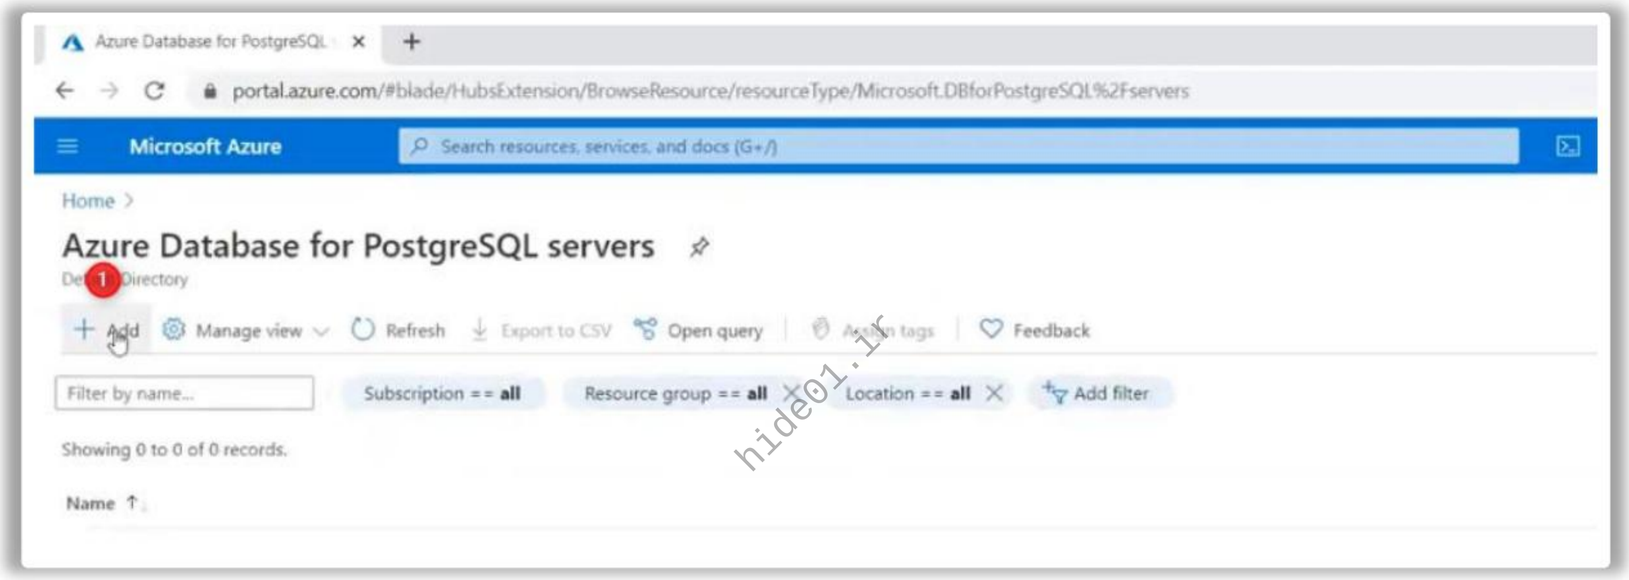

Figure 9.152: Click Add

- Under **Single server** click on **Create**.

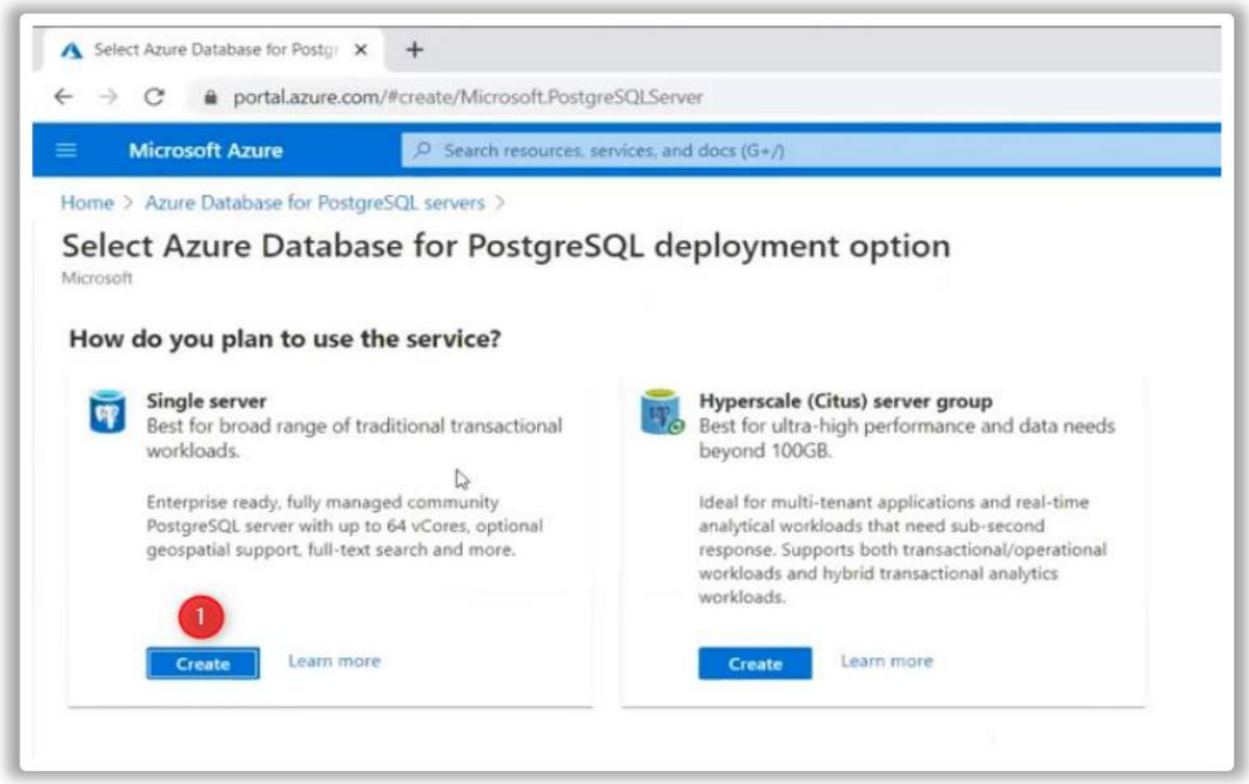

Figure 9.153: Click Create

Certified Cloud Security Engineer Copyright © by **EC-Council**
All Rights Reserved. Reproduction is Strictly Prohibited.

- Select a **Resource group** from the drop-down list or create a new resource and click on **OK**.

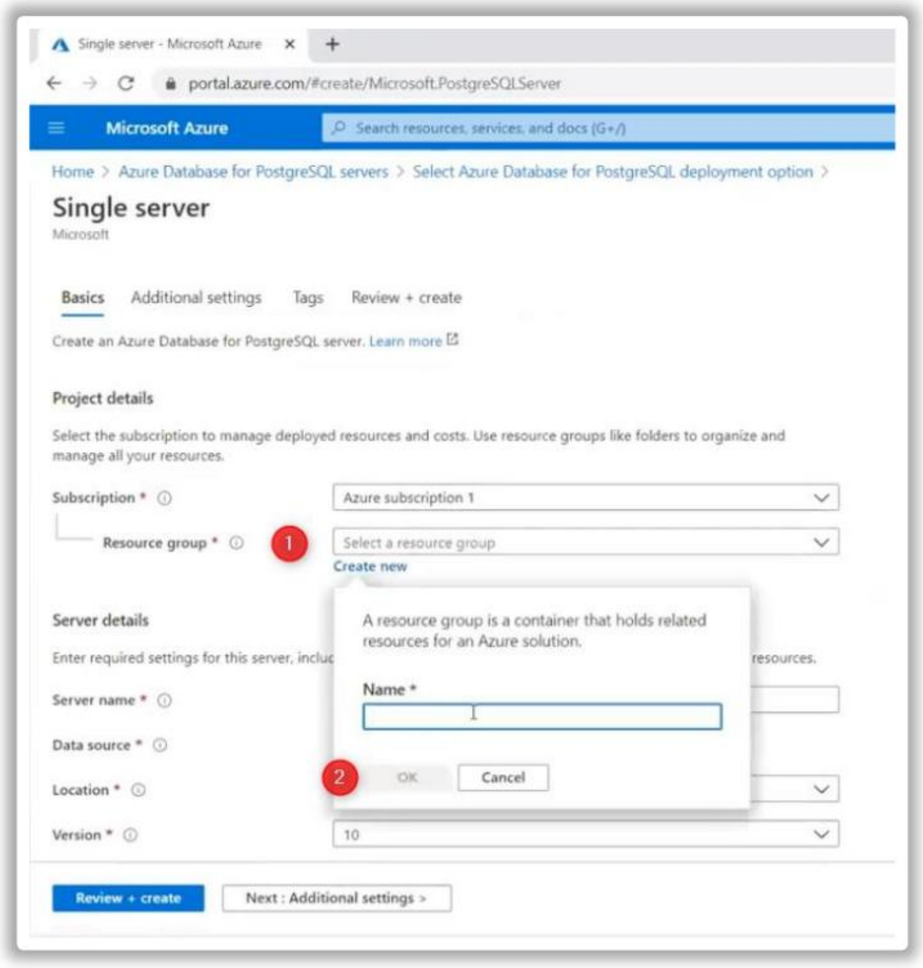

Figure 9.154: Select Resource Group

- Enter required details such as **Server name**, **Location**, and **Version**.
- Click on **Configure server** beside **Compute + storage**.

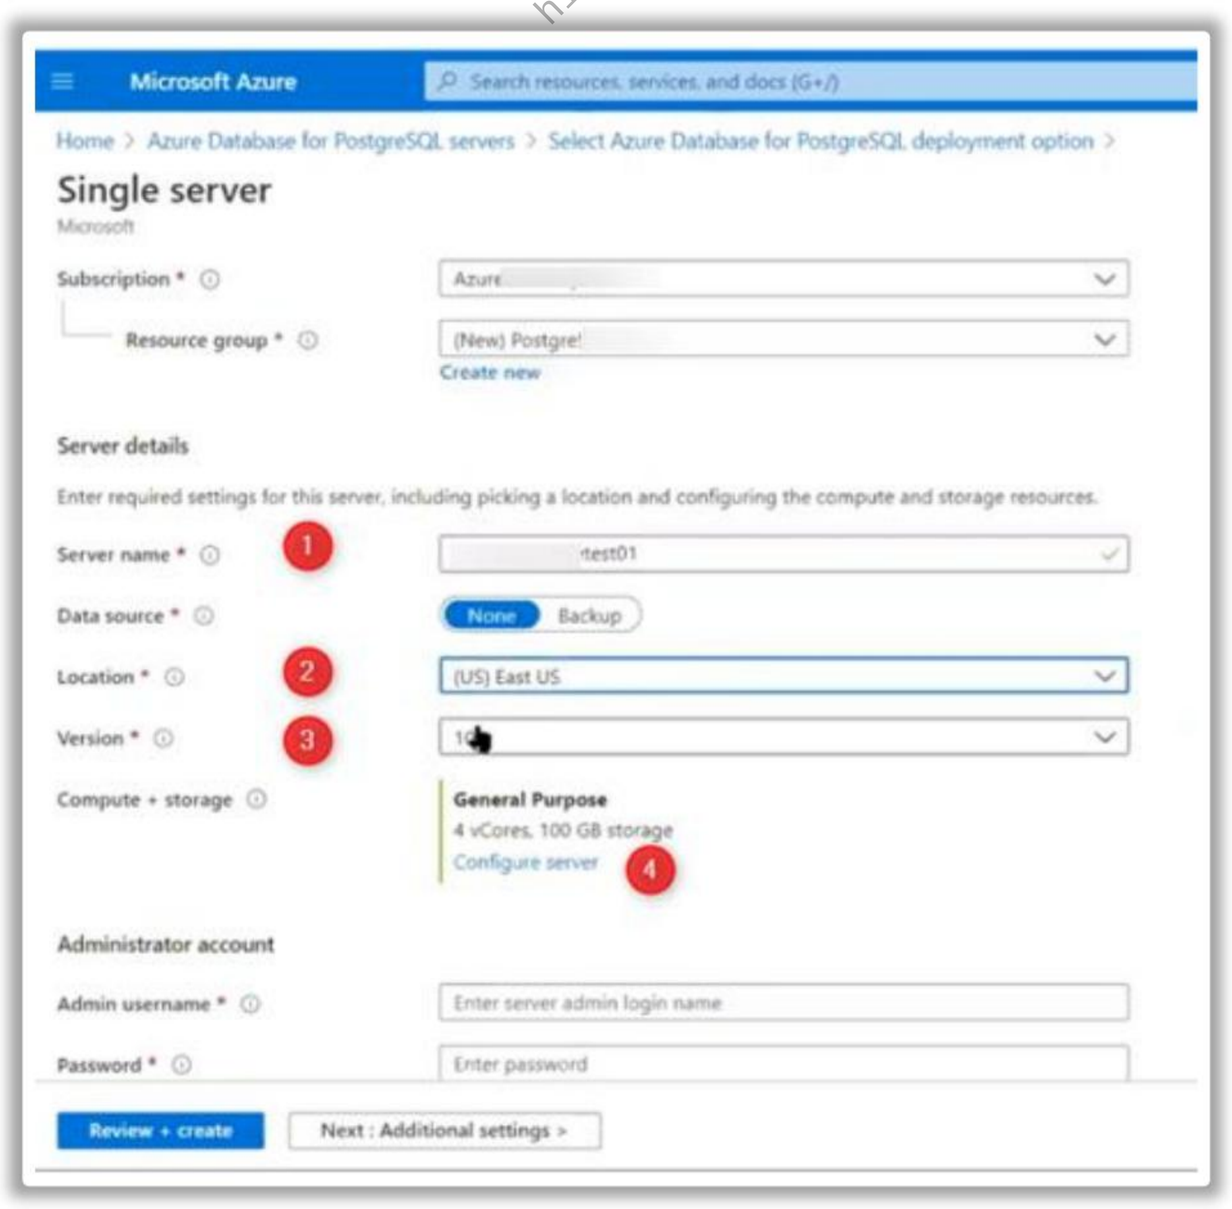

Figure 9.155: Single Server Console

Certified Cloud Security Engineer Copyright © by **EC-Council**
All Rights Reserved. Reproduction is Strictly Prohibited.

- Select **Geo-Redundant** under **Backup Redundancy Options**.
- Select the **Backup Retention Period** by dragging on the scale provided; the maximum retention period is 35 days.
- Click on **OK**.

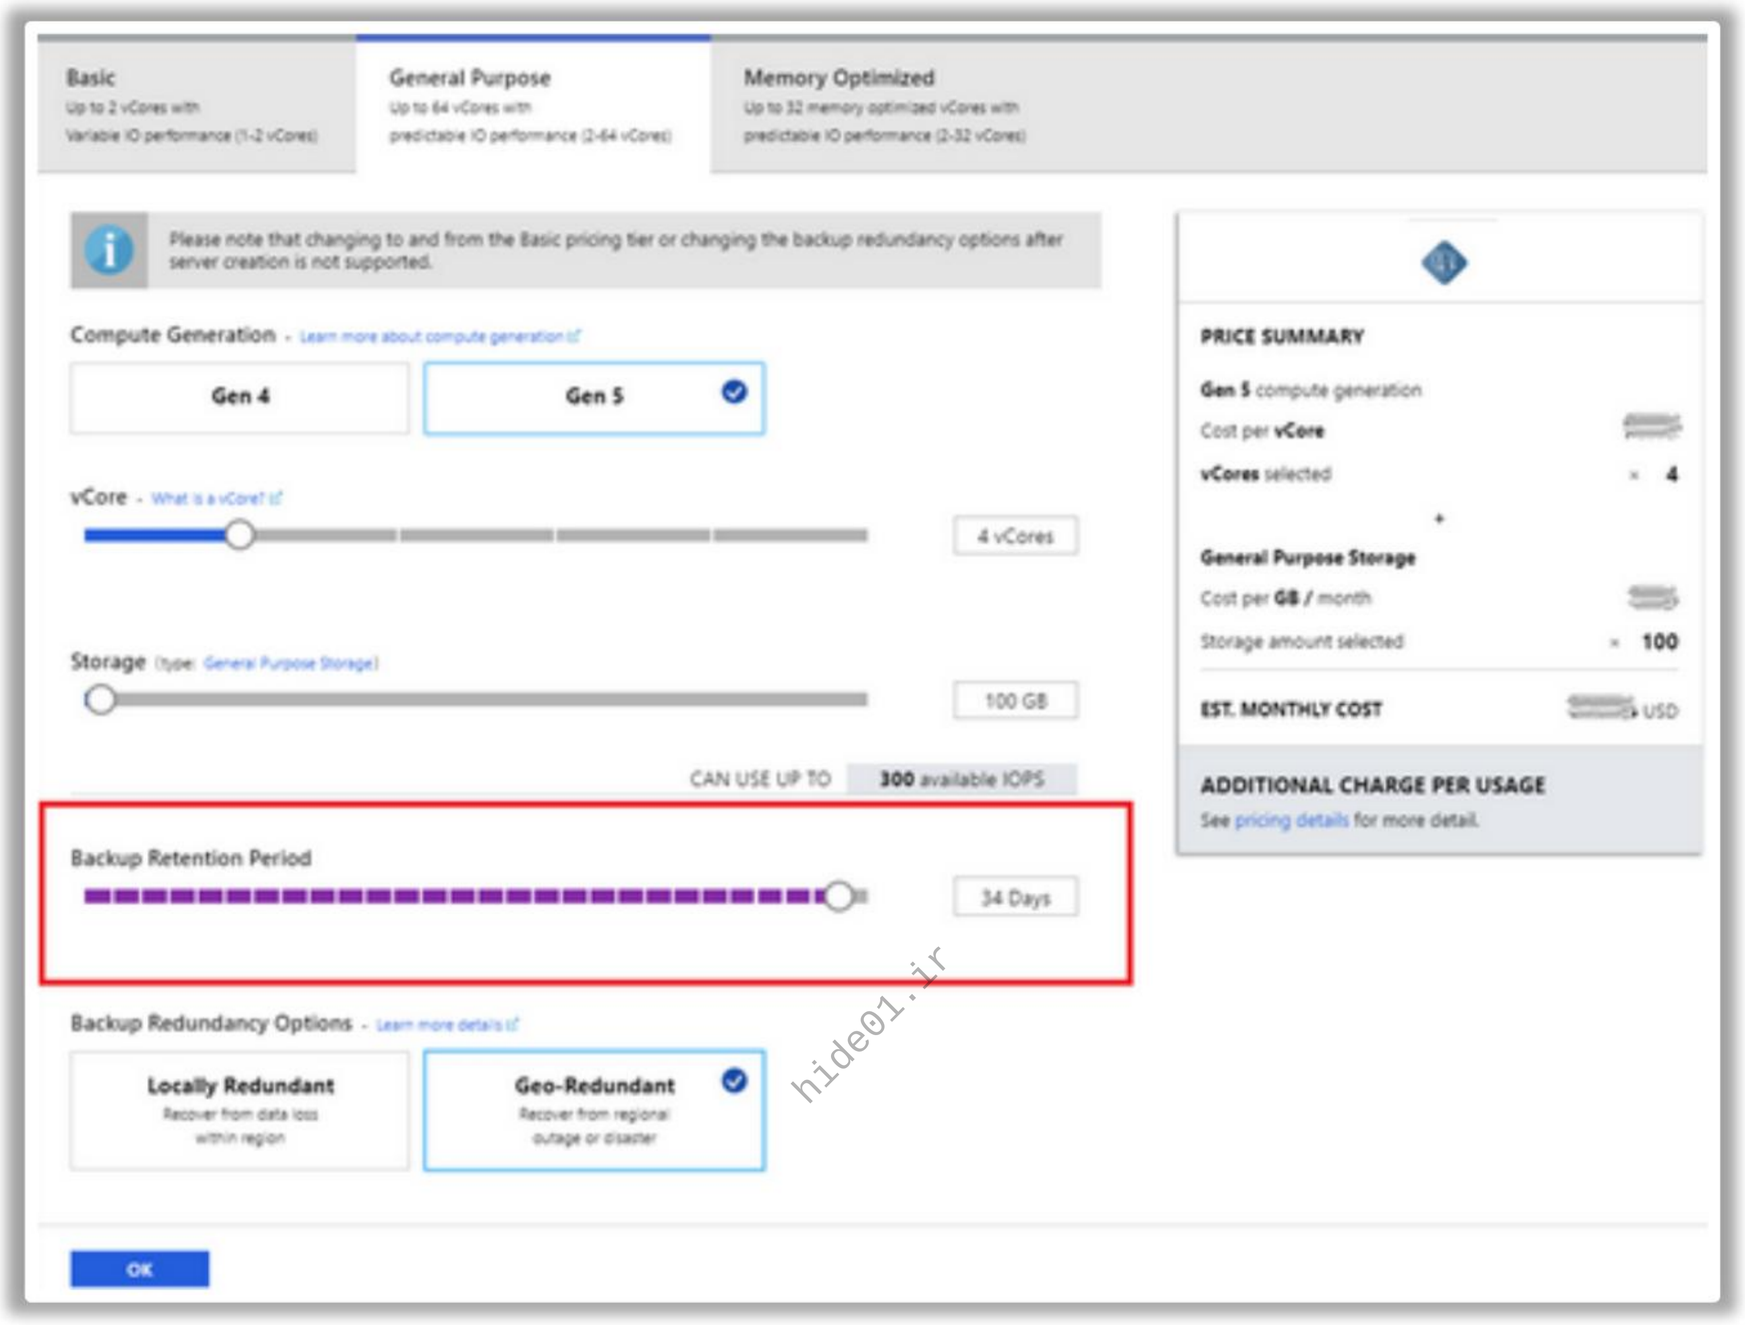

Figure 9.156: Click OK

- Click **Next** to proceed with the setup.

Certified Cloud Security Engineer Copyright © by **EC-Council**
All Rights Reserved. Reproduction is Strictly Prohibited.

- On the **Tags** panel any required tags can be set up; then click on **Review + create.**

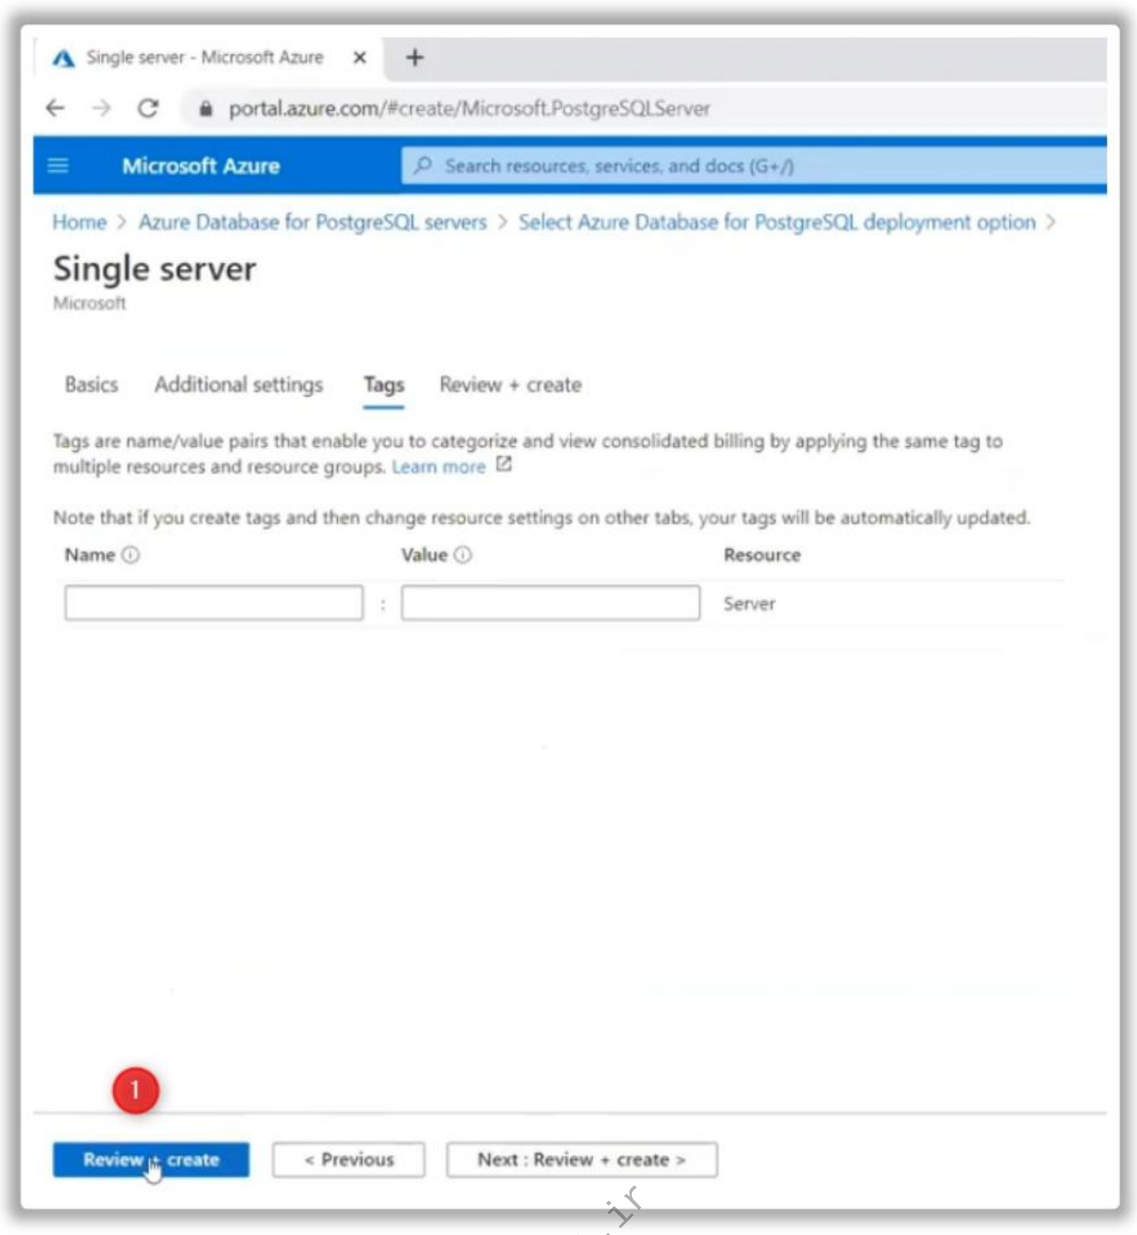

Figure 9.157: Click Review+ create

- Review the server configuration and click on **Create**.

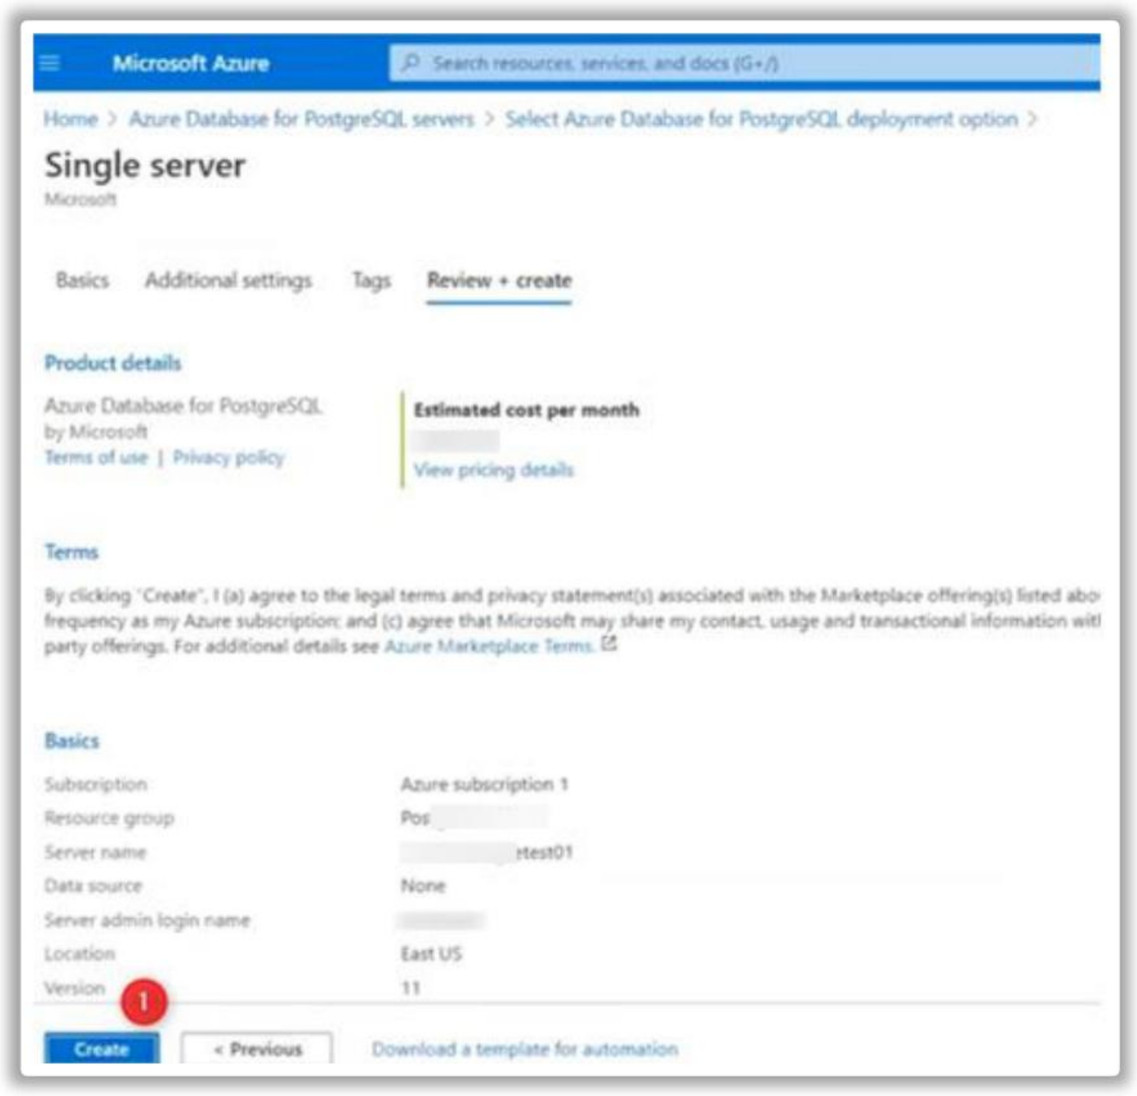

Figure 9.158: Click Create

Certified Cloud Security Engineer Copyright © by **EC-Council**
All Rights Reserved. Reproduction is Strictly Prohibited.

- The PostgreSQL server is successfully created with Geo-redundant Backup.

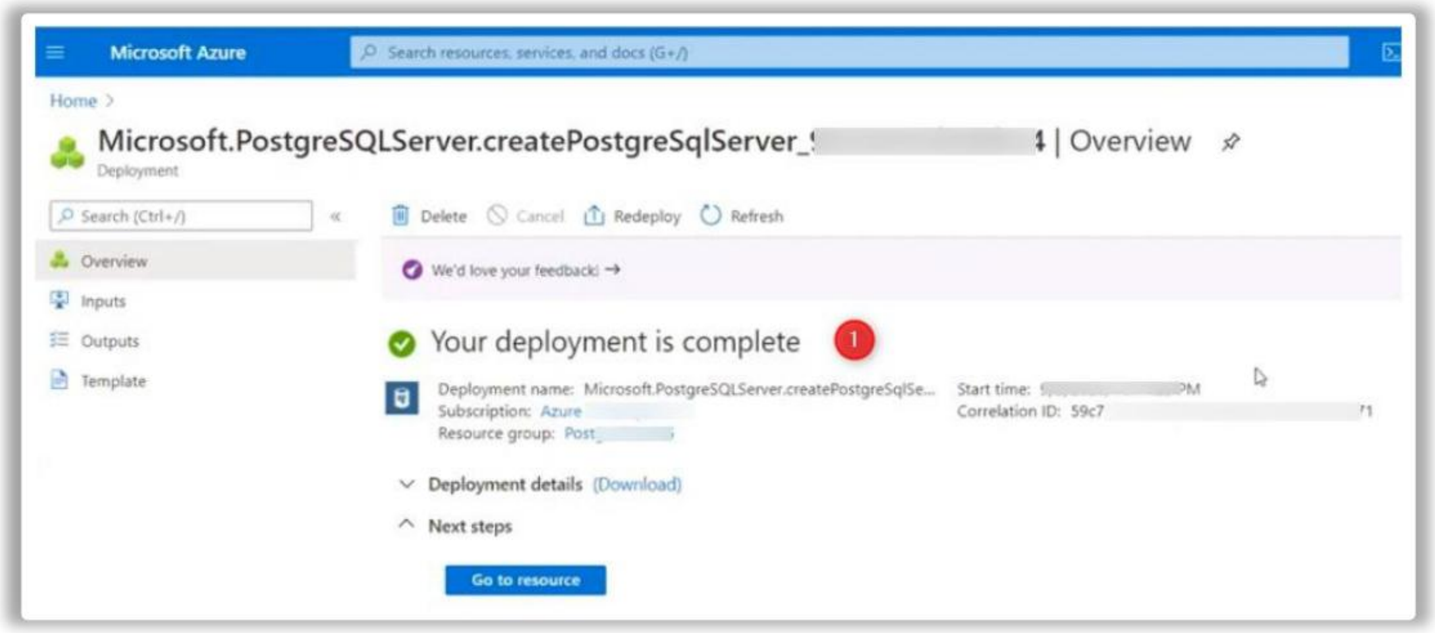

Figure 9.159: PostgreSQL server is successfully created

hide01.ir

Certified Cloud Security Engineer Copyright © by **EC-Council**
All Rights Reserved. Reproduction is Strictly Prohibited.

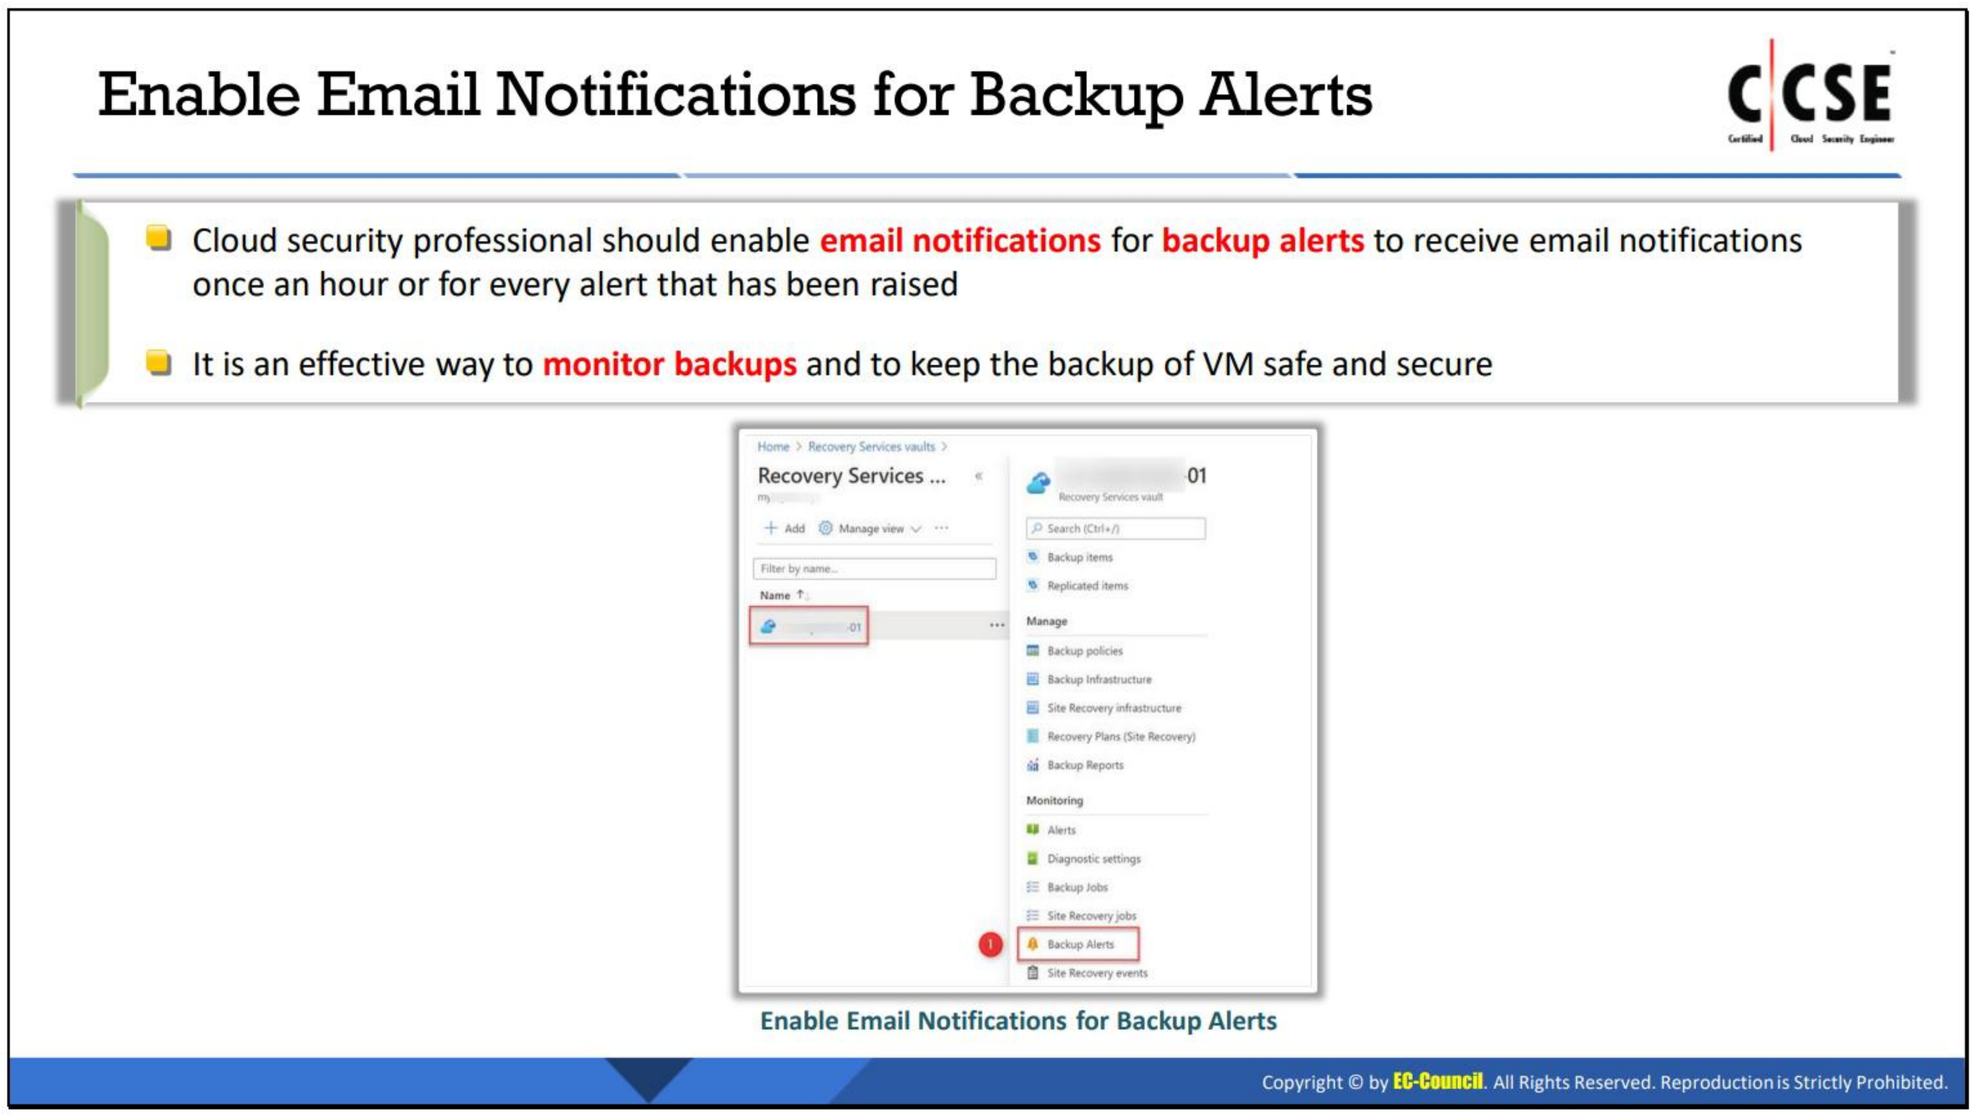

Enable Email Notifications for Backup Alerts

Enable Email Notification for Backup Alerts

It is recommended to configure notifications to trigger emails in the event of any warning or critical backup alerts in the Azure cloud account. This feature allows you to receive email notifications once an hour or for every alert that has been raised. Email notifications for alerts represent an effective way to monitor backups and also keep the backup of VM safe and secure. The notification contains details such as what backup has been deleted, how long is it available to recover, and allows you to take action.

To Enable Email Notifications for Backup Alerts

- Sign in to the Azure portal.
- Click on **More services**.

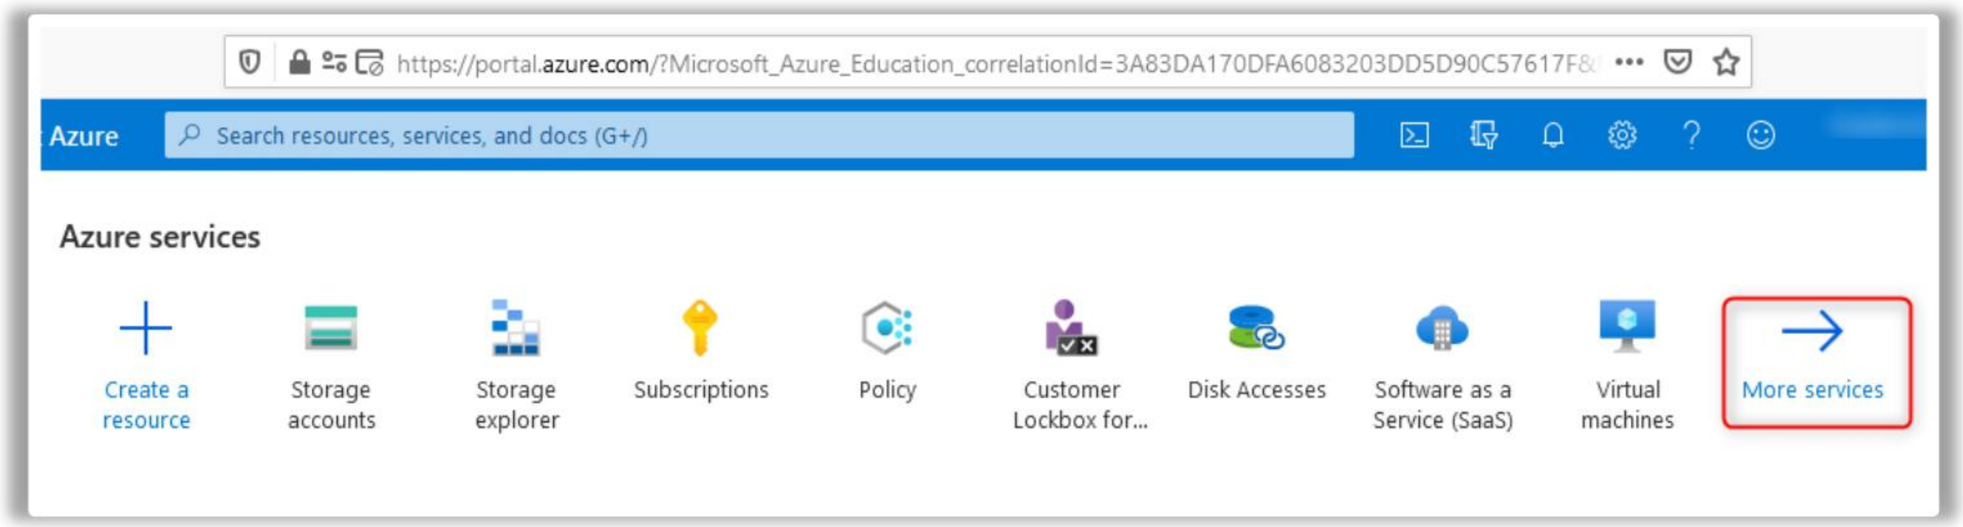

Figure 9.160: Microsoft Azure Console

- Search **Recovery Services vault** and open it.
- Select the vault for which you want to enable notifications.

Certified Cloud Security Engineer Copyright © by **EC-Council**
All Rights Reserved. Reproduction is Strictly Prohibited.

- Navigate to **Monitoring** and select **Backup Alerts**.

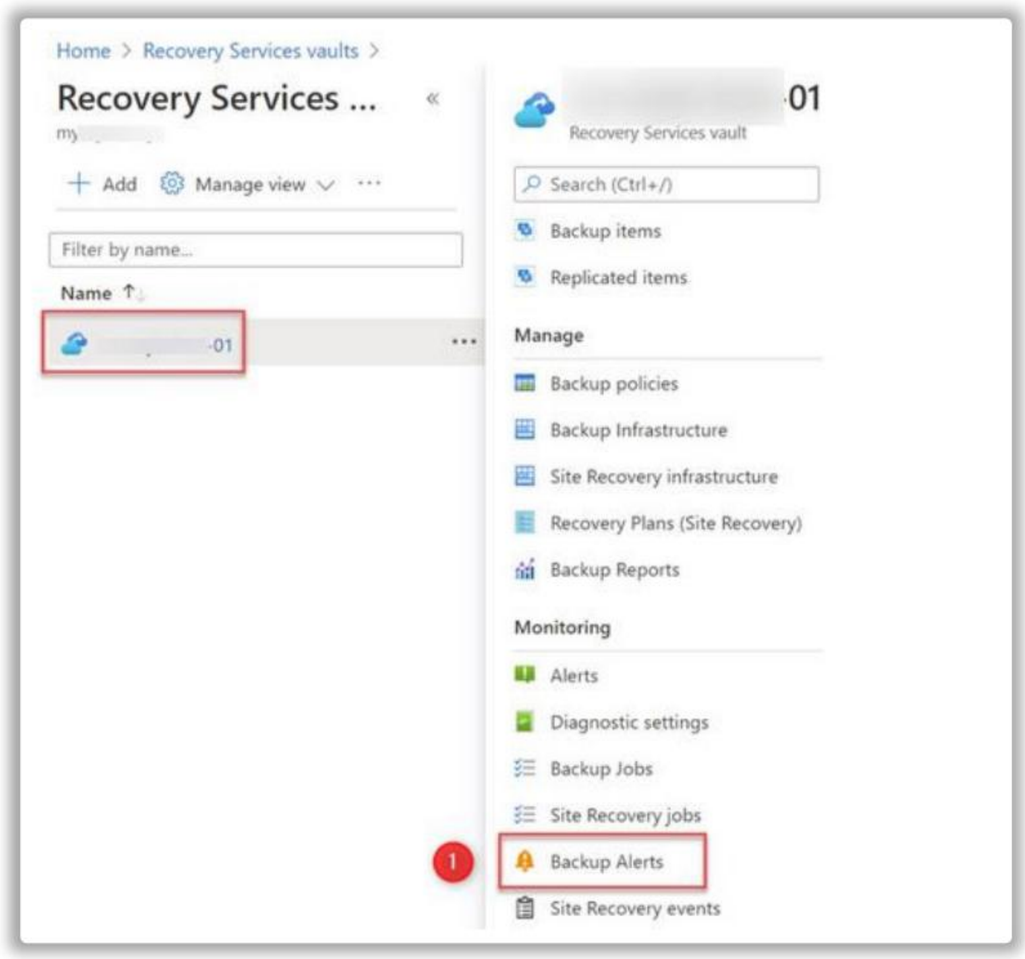

Figure 9.161: Click Backup Alerts

- Now click **Configure notifications** tab.

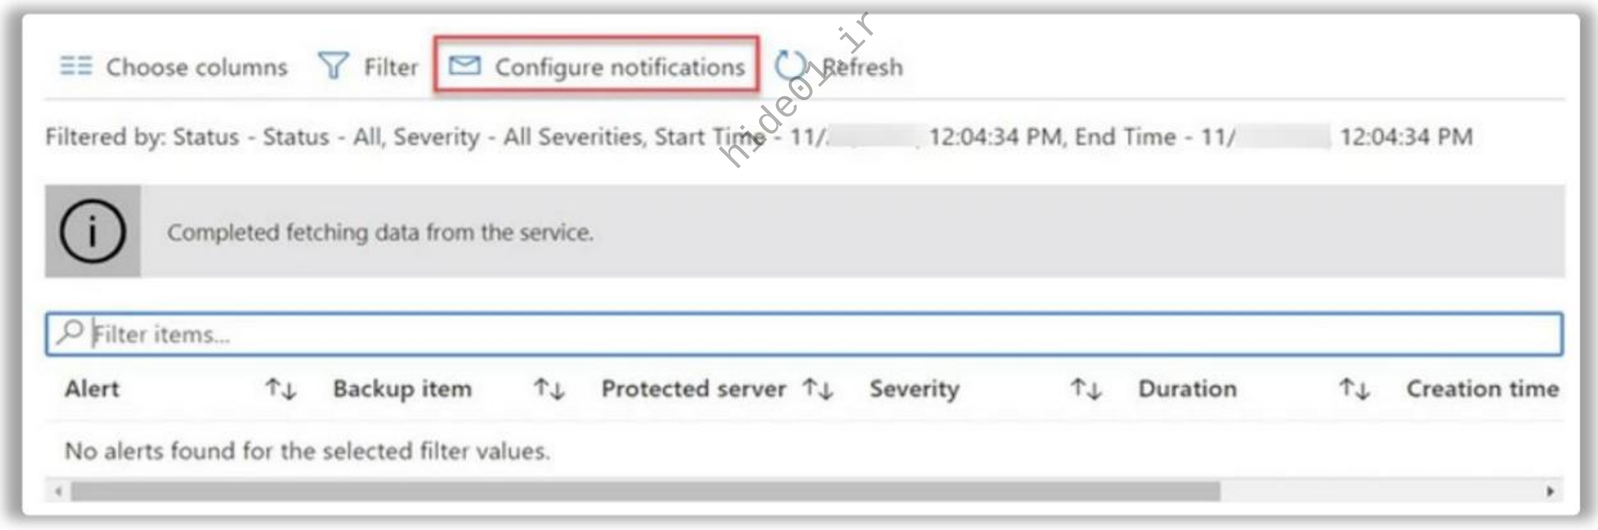

Figure 9.162: Click Configure notification

- Go to **Configure notifications,** under **Email notifications,** turn the toggle bar to **On**.
- Select either **Per Alert** or **Hourly Digest** under **Notify** as per requirement.

Certified Cloud Security Engineer Copyright © by **EC-Council**
All Rights Reserved. Reproduction is Strictly Prohibited.

- Click on **Save**.

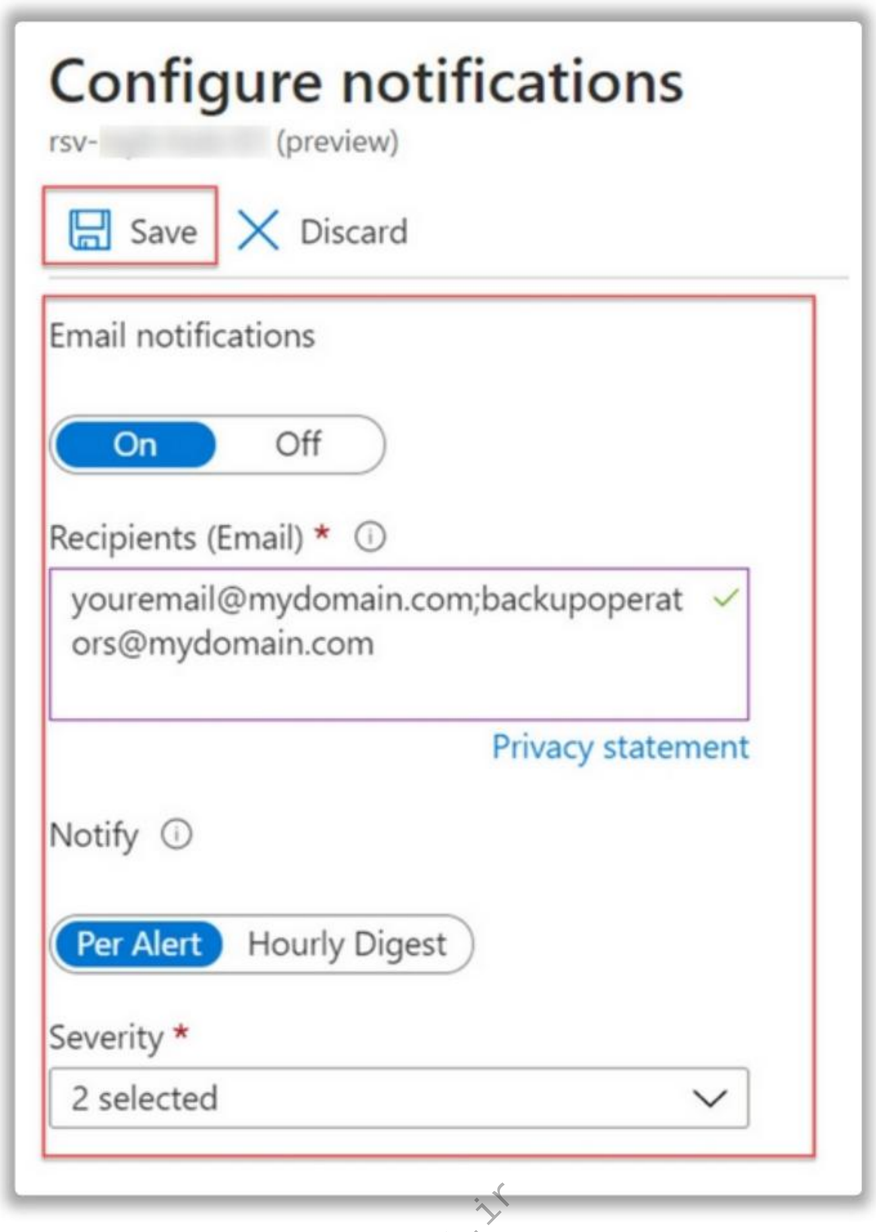

Figure 9.163: Click Save

- Email notification configuration is successful.

Figure 9.164: Configure alert

Certified Cloud Security Engineer Copyright © by **EC-Council**
All Rights Reserved. Reproduction is Strictly Prohibited.

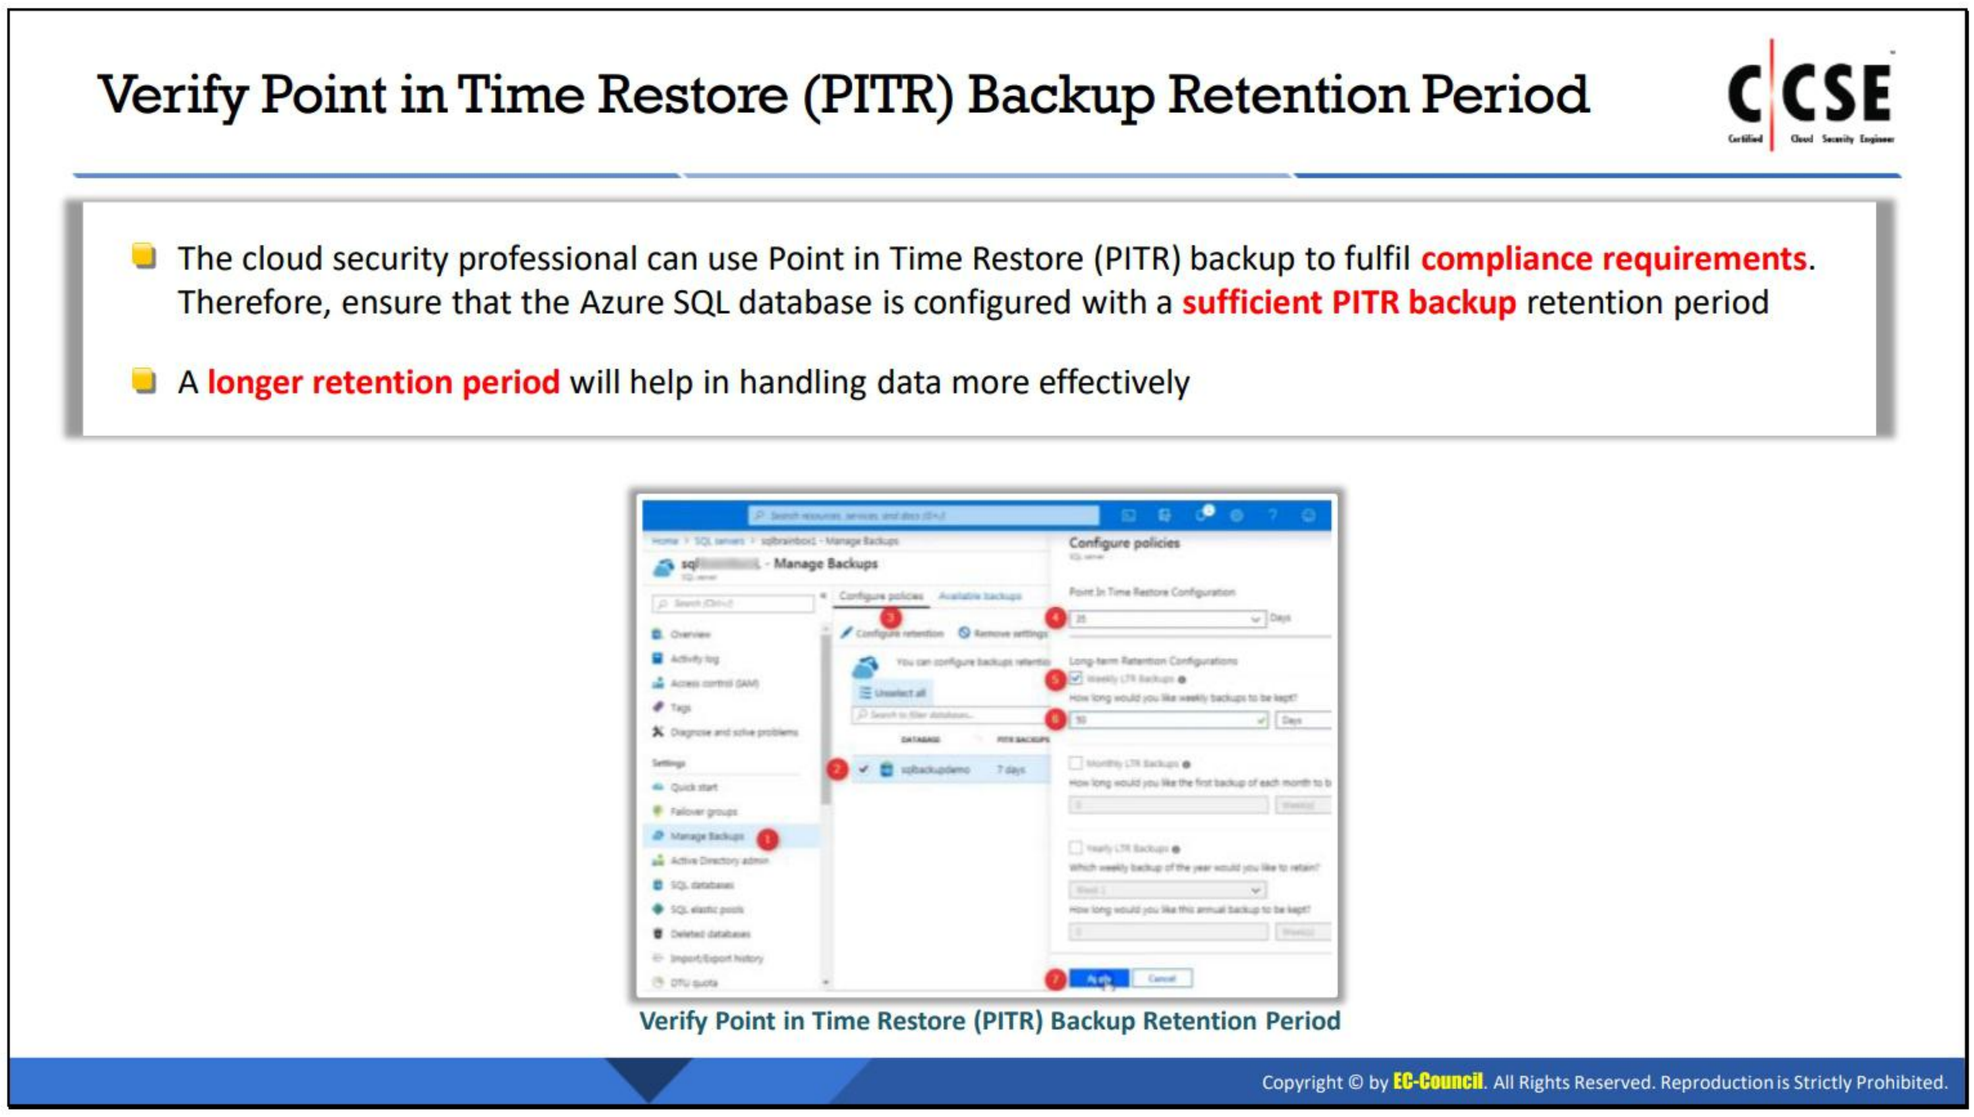

Verify Point in Time Restore (PITR) Backup Retention Period

PITR backup allows organizations to fulfil compliance requirements. Ensure that the Azure SQL Database is configured with a sufficient PITR backup retention period. The backups can be stored between 7 to 35 days. To make the backups available even if the primary data center is not available, SQL service uses read-access geo-redundant storage (RA-GRS). For Azure SQL databases, you need to have an optimal backup retention period as a best practice. A longer retention period will help in handling data more effectively.

To Check for Sufficient Point in Time Restore (PITR) Backup Retention Period

- Sign in to the Azure portal.
- Go to **More services**.

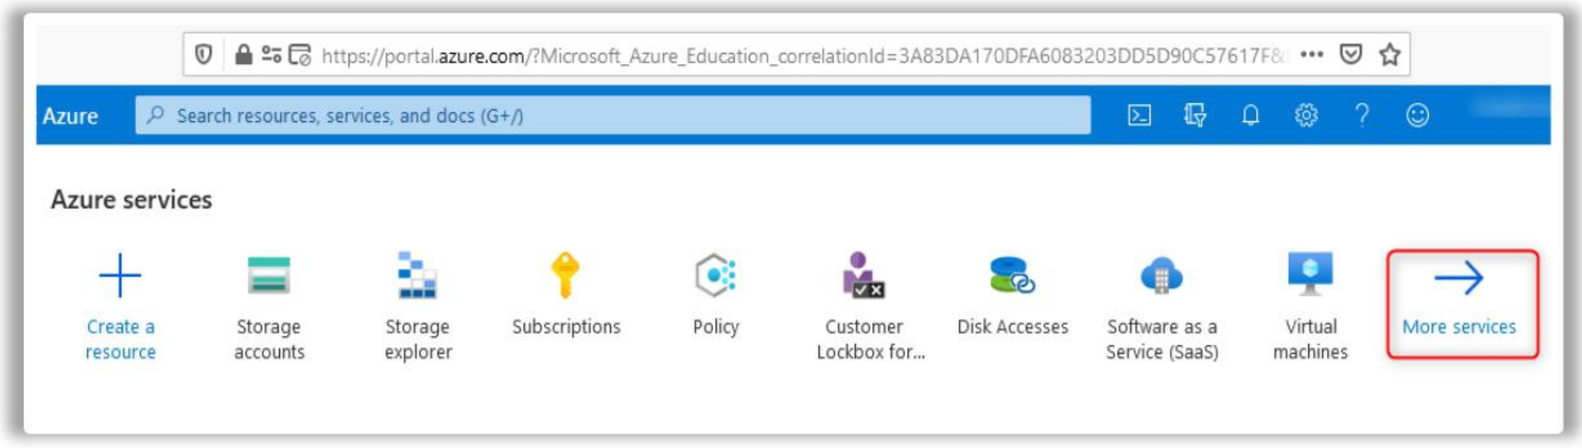

Figure 9.165: Microsoft Azure Console

- In the left navigation panel, search for **SQL Server** and click on it.

Certified Cloud Security Engineer Copyright © by **EC-Council**
All Rights Reserved. Reproduction is Strictly Prohibited.

- A list of available SQL servers will be displayed; select the SQL server for which you need to verify the PITR backup retention period.

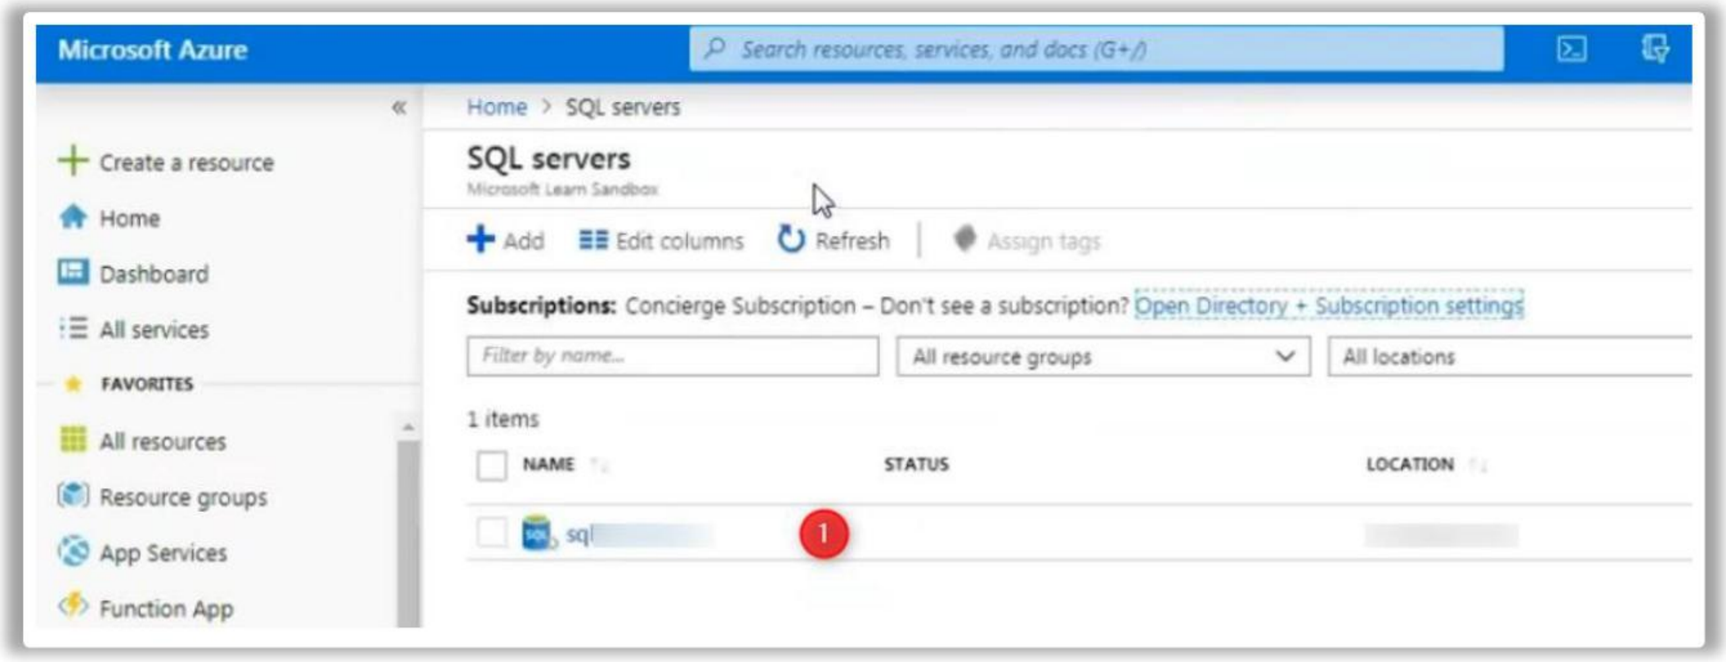

Figure 9.166: SQL servers Console

- In the Navigation panel locate **Settings** and go to **Manage Backups**.
- Select the check box of the database; then click on **configure retention.**
- Select the no of days in **Point In Time Restore Configuration** as per requirement.
- Under "**Long-term Retention Configurations**", select the check box of **Weekly LTR Backups** and enter the required number of days below.
- Finally, click on **Apply**.

hide01.ir

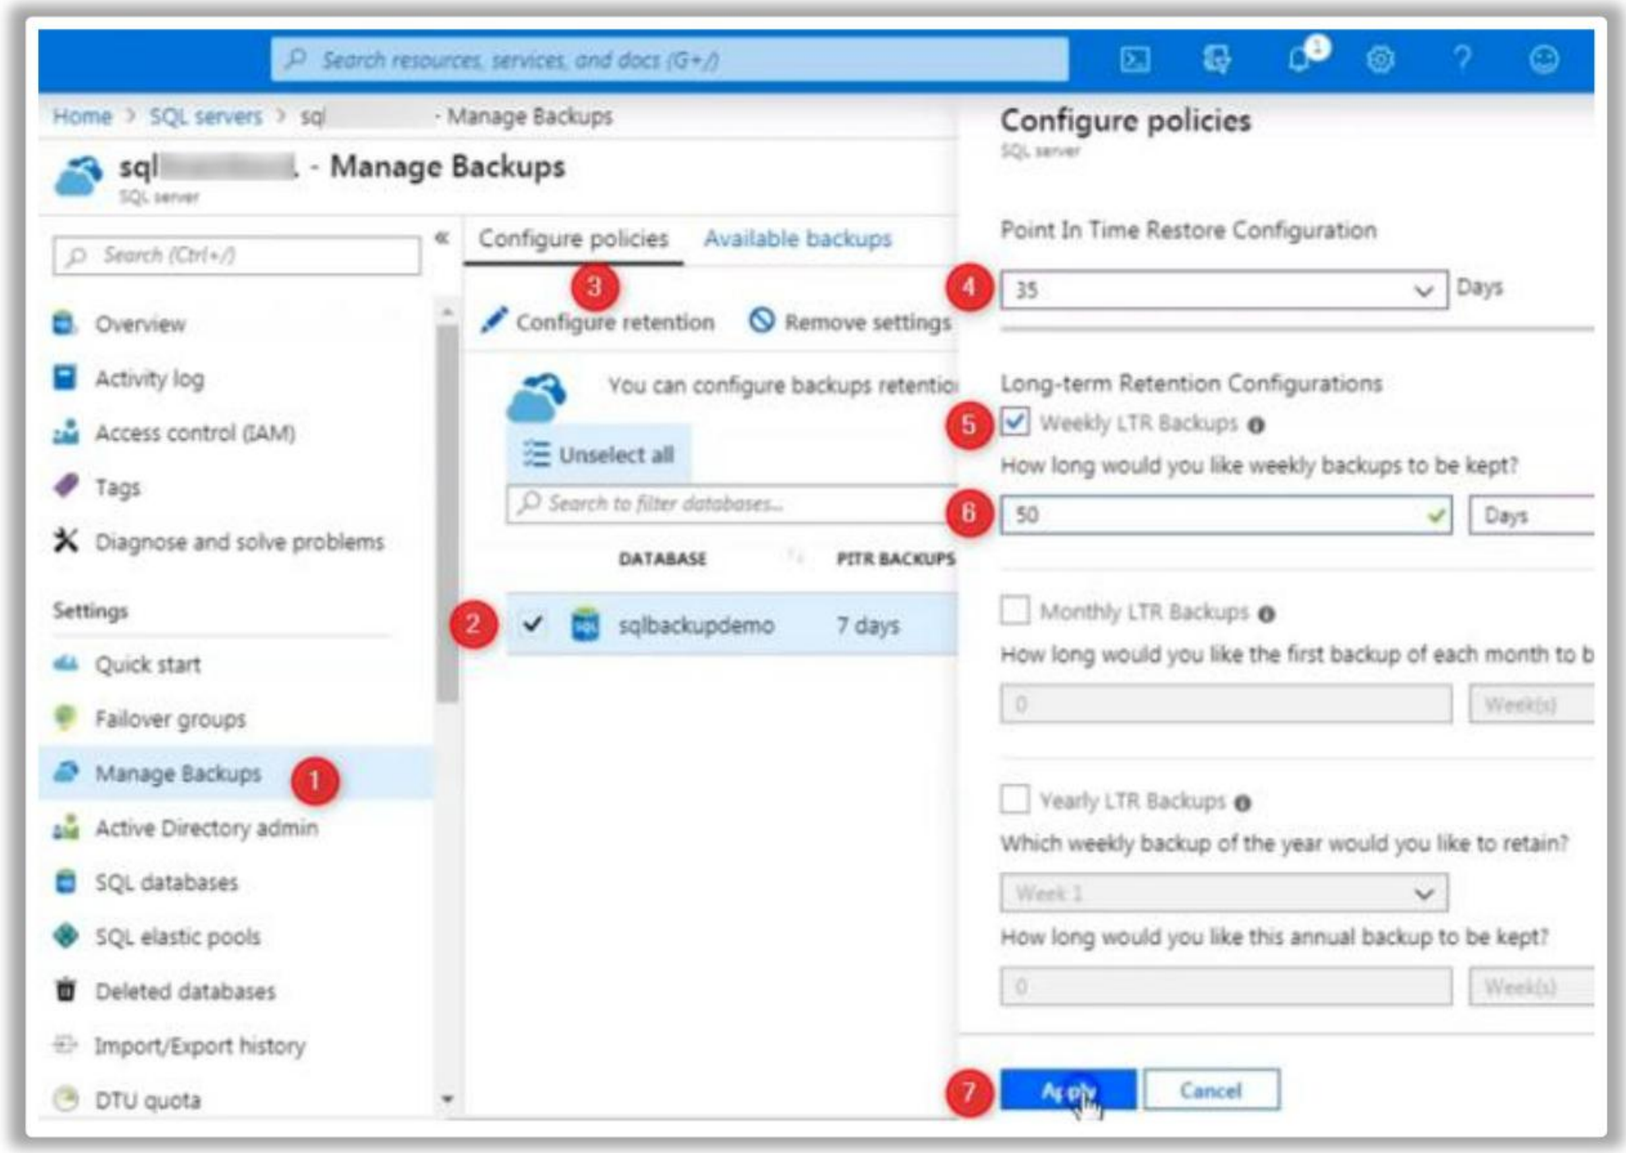

Figure 9.167: Click Apply

Certified Cloud Security Engineer Copyright © by **EC-Council**
All Rights Reserved. Reproduction is Strictly Prohibited.

- Click on **Yes.**

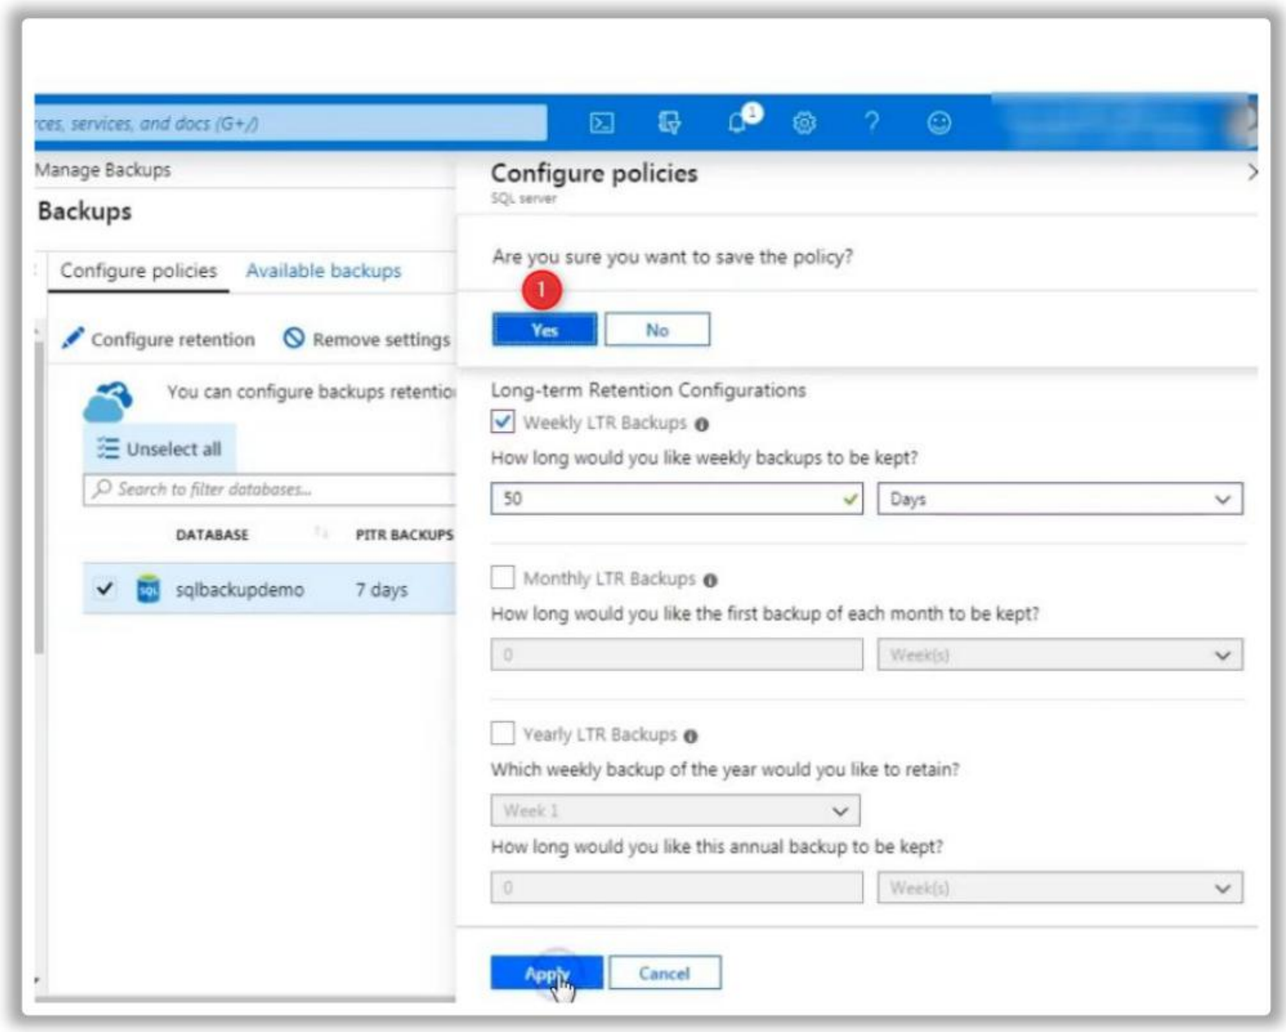

Figure 9.168: Click Yes

- Click on **Refresh** to reflect changes; now the PITR backup retention period is successfully updated.

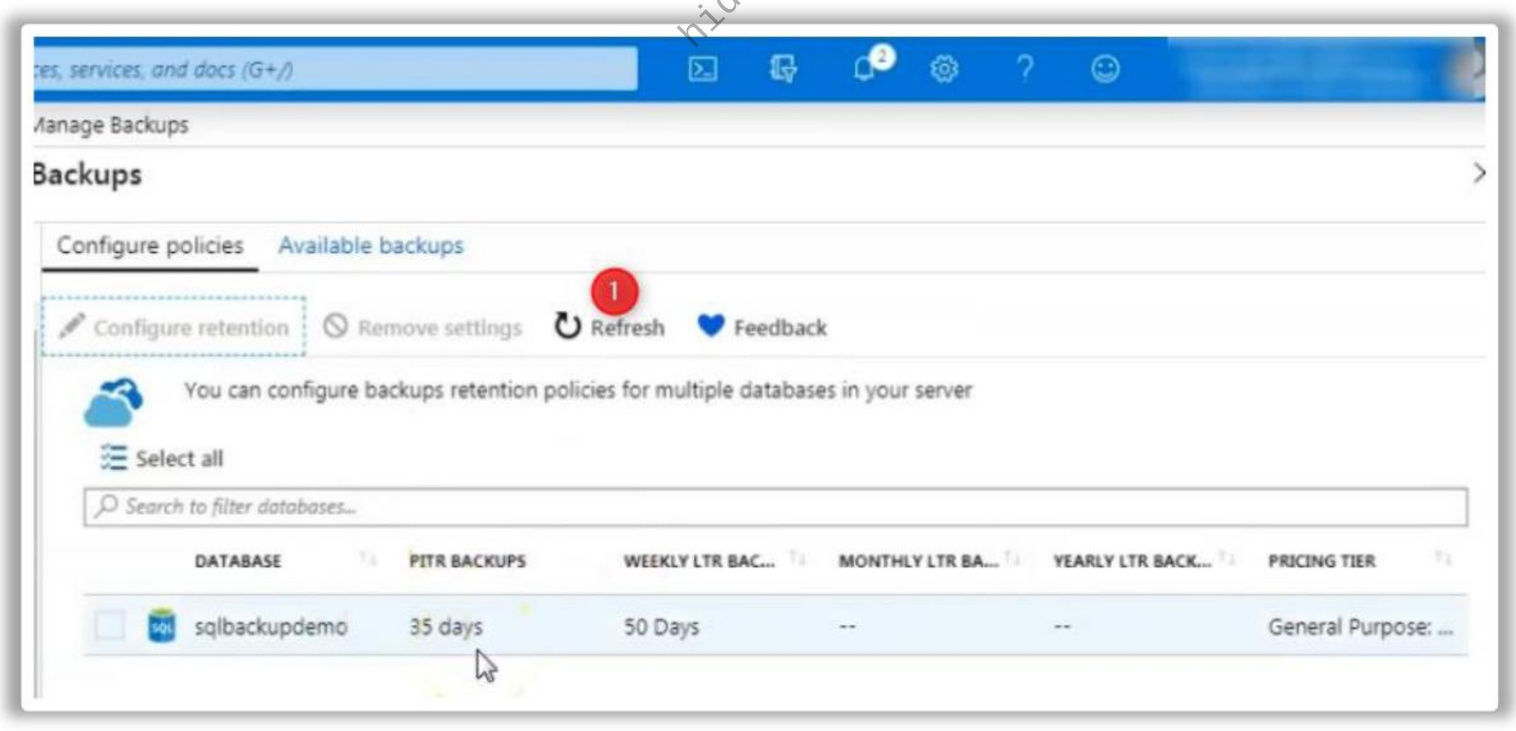

Figure 9.169: Microsoft Azure Console

Certified Cloud Security Engineer Copyright © by **EC-Council**
All Rights Reserved. Reproduction is Strictly Prohibited.

Configure and Enable Azure Backup Service for VM

- To meet **compliance requirements** and to provide security, configure and enable Azure Backup service for Azure virtual machines
- It is a **cost-effective** and simple backup solution that helps in the data recovery of the virtual machine

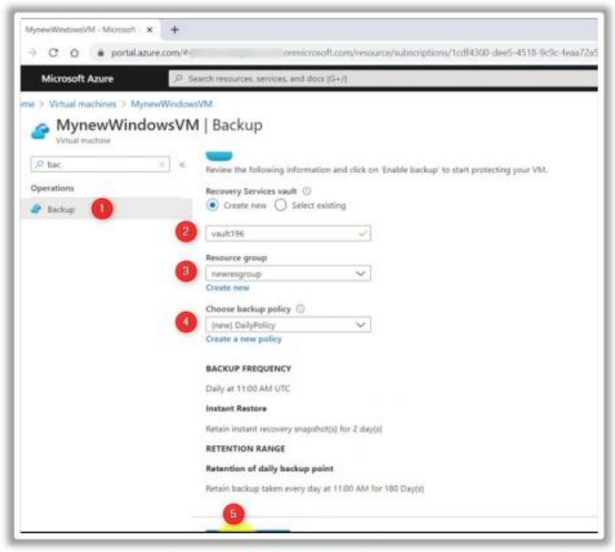

Configure and Enable Azure Backup Service for VM

Copyright © by EC-Council. All Rights Reserved. Reproduction is Strictly Prohibited.

Configure and Enable Azure Backup Service for VM

To meet compliance requirements and to provide security, make sure that the Azure Backup service is configured and enabled for your Azure VMs. It is a cost-effective and simple backup solution that helps in the data recovery of the VM. To avoid data loss due to accidental deletion or corruption, it is recommended to configure and enable VM backups. Once the configuration of the Azure Backup service is done, then the backup service will run at the scheduled time. Using these backups, the recovery points are created, and these recovery points are stored in the Azure RSV.

To configure and enable Azure Backup service for VM

- Sign in to the Azure portal.
- Click on **More services**.

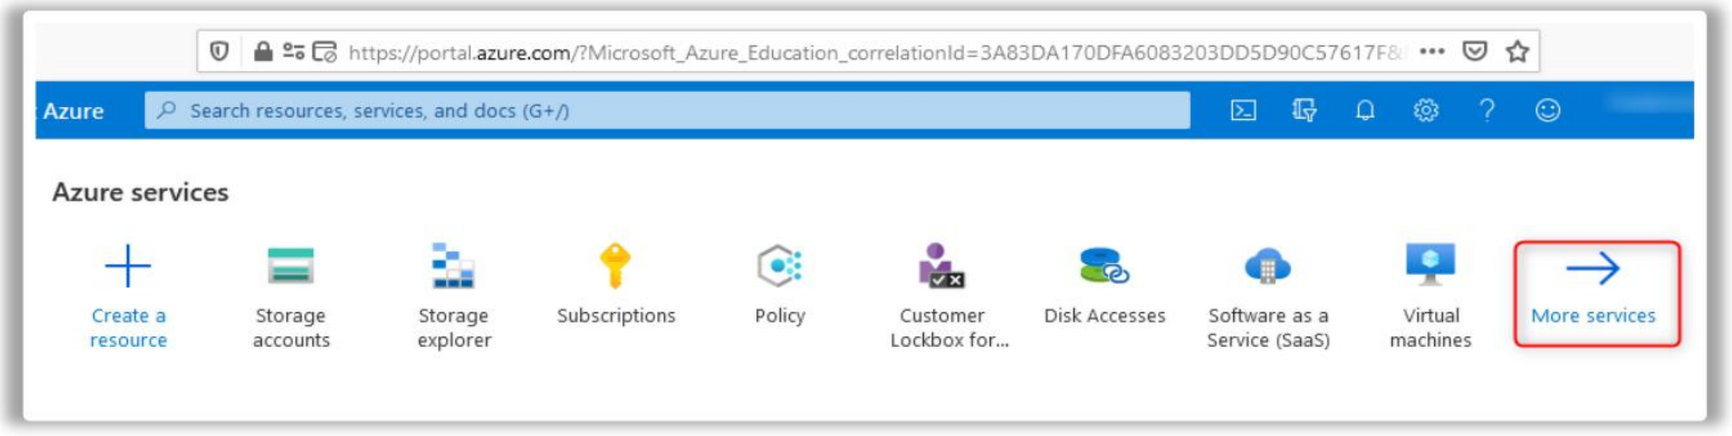

Figure 9.170: Microsoft Azure Console

Certified Cloud Security Engineer Copyright © by **EC-Council**
All Rights Reserved. Reproduction is Strictly Prohibited.

- Search for **Virtual machines** and click on the same.

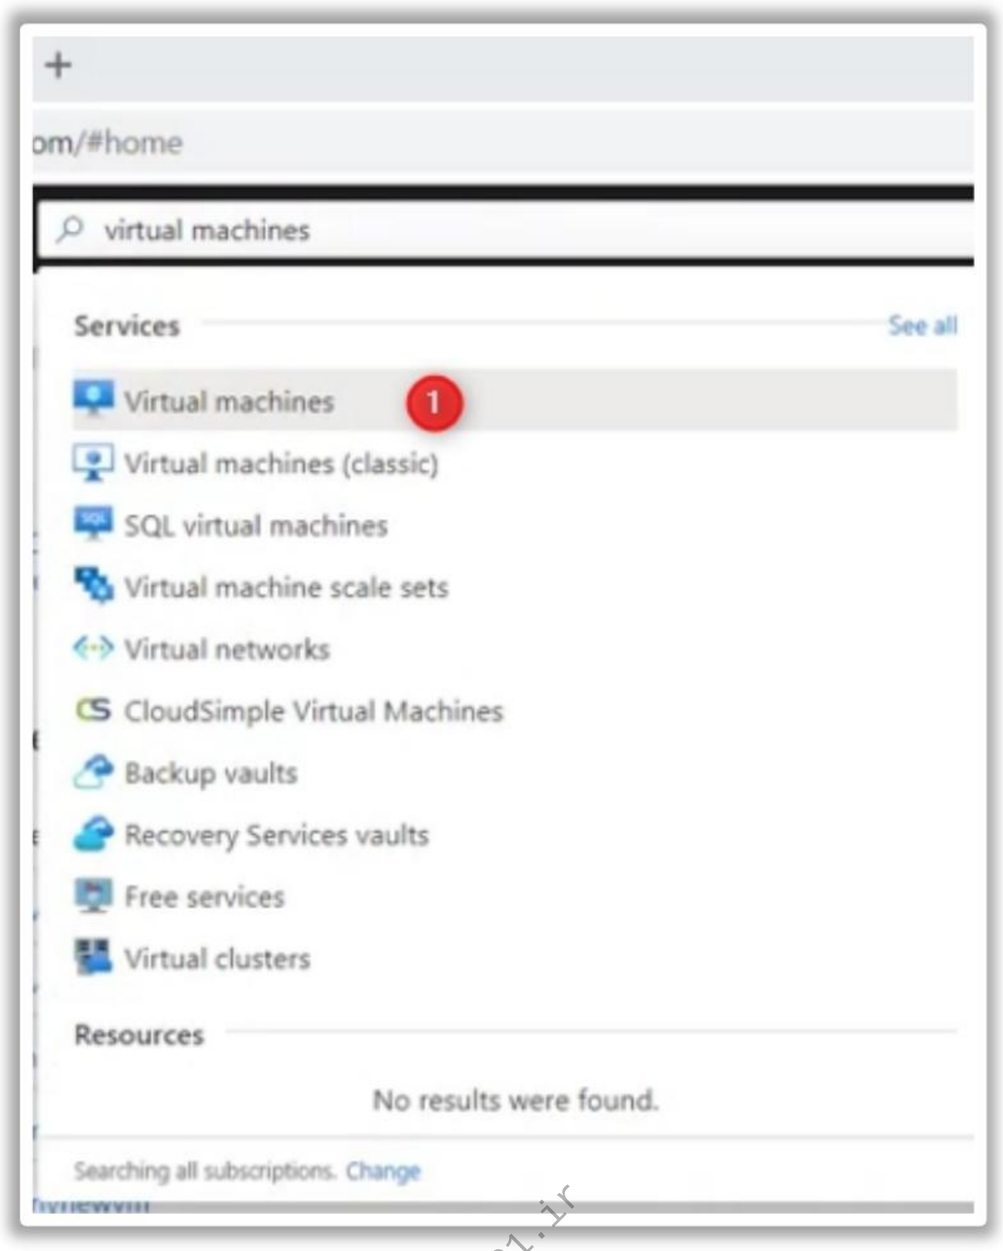

Figure 9.171: Searching VM

- A list of available VMs is displayed.

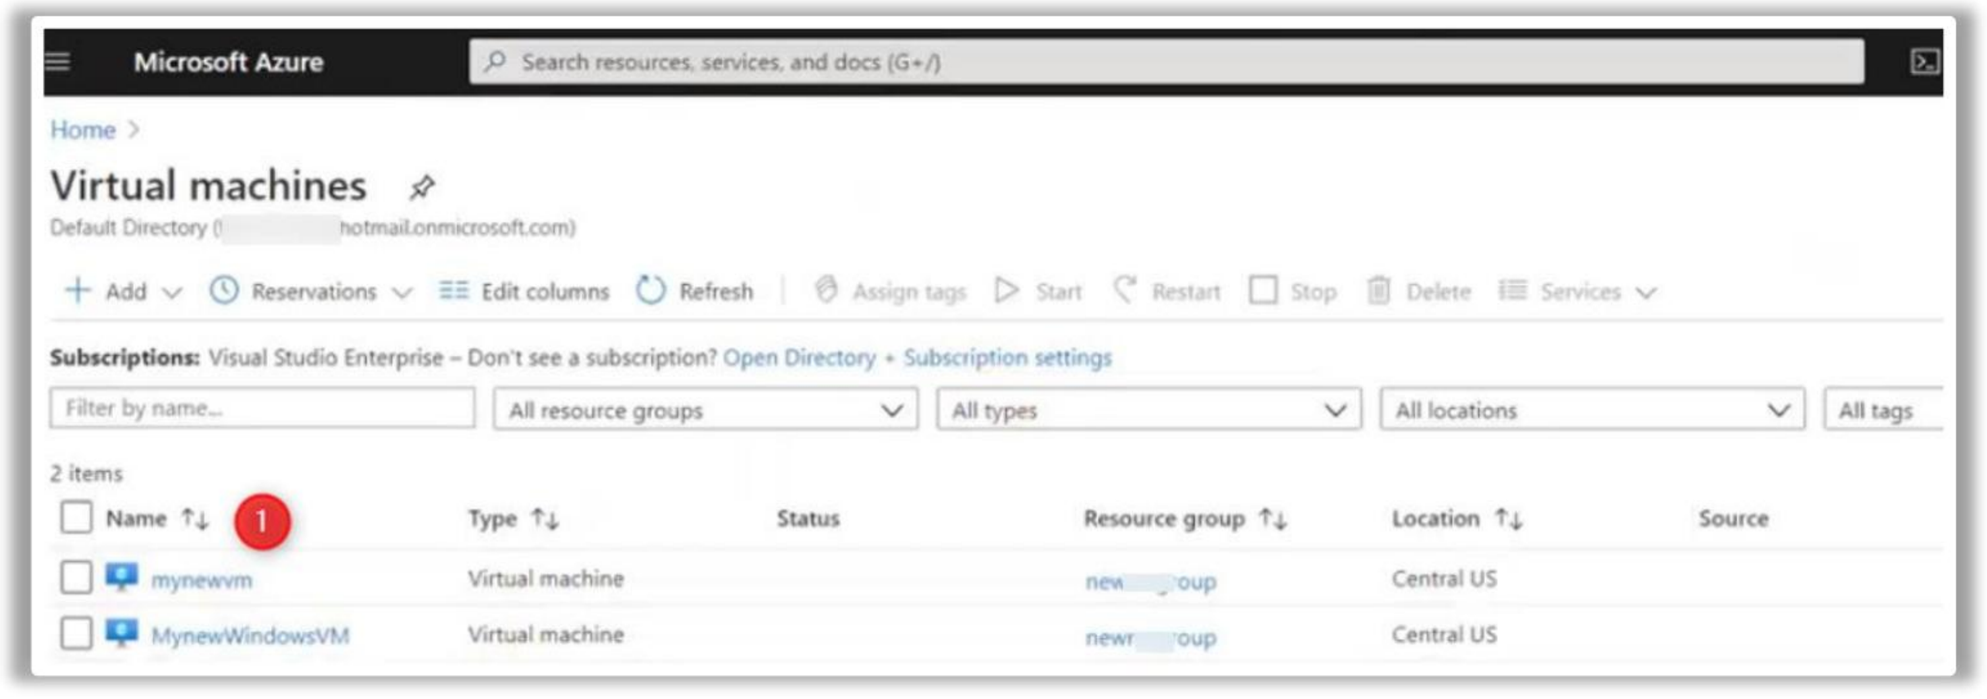

Figure 9.172: List of VM

- Select the VM that you want to configure.
- From the left navigation panel locate **Operations** and click on **Backup**.
- Under the **Recovery Services vault**, you can either create a new one or select an existing one.

Certified Cloud Security Engineer Copyright © by **EC-Council**
All Rights Reserved. Reproduction is Strictly Prohibited.

- Select a resource from the drop-down under **Resource group**.
- Choose the backup policy from the drop-down under **Choose backup policy**.
- Now click on **Enable Backup**.

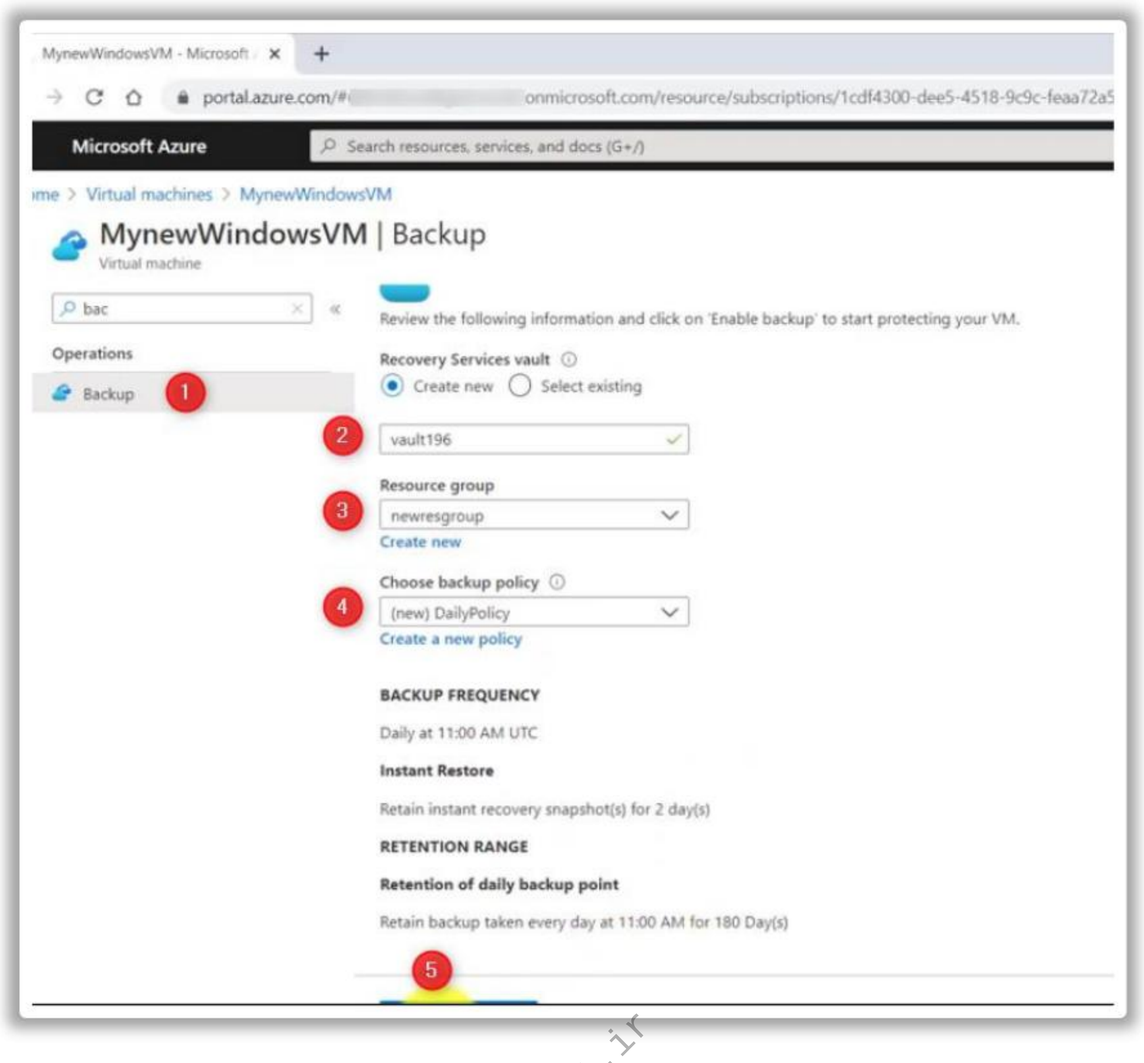

Figure 9.173: Enabling Backup

- You can wait for the scheduled backup or click on **Backup now**.

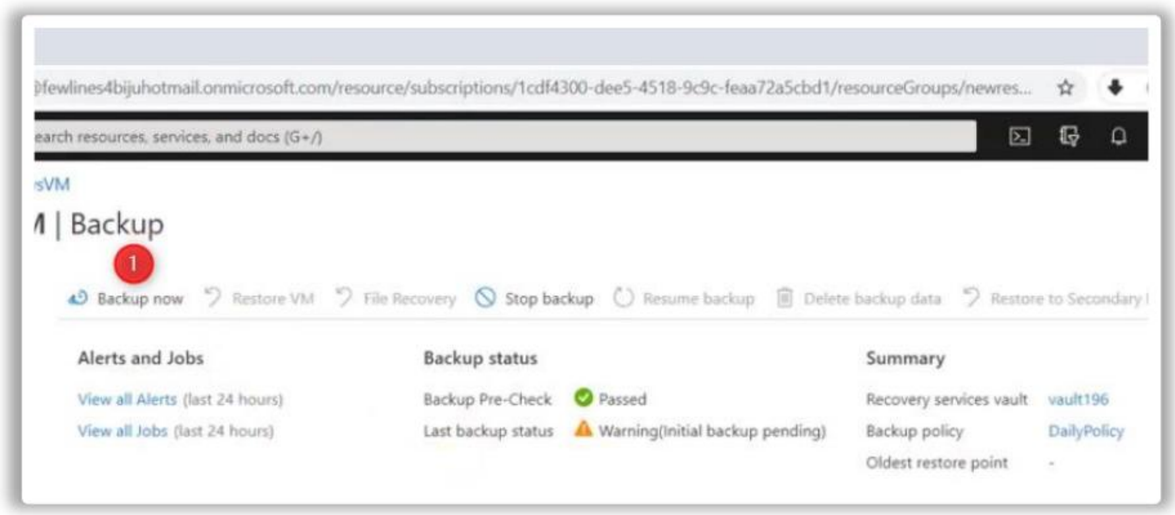

Figure 9.174: Schedule Backup

- Backup has been enabled successfully in the VM.

Certified Cloud Security Engineer Copyright © by **EC-Council**
All Rights Reserved. Reproduction is Strictly Prohibited.

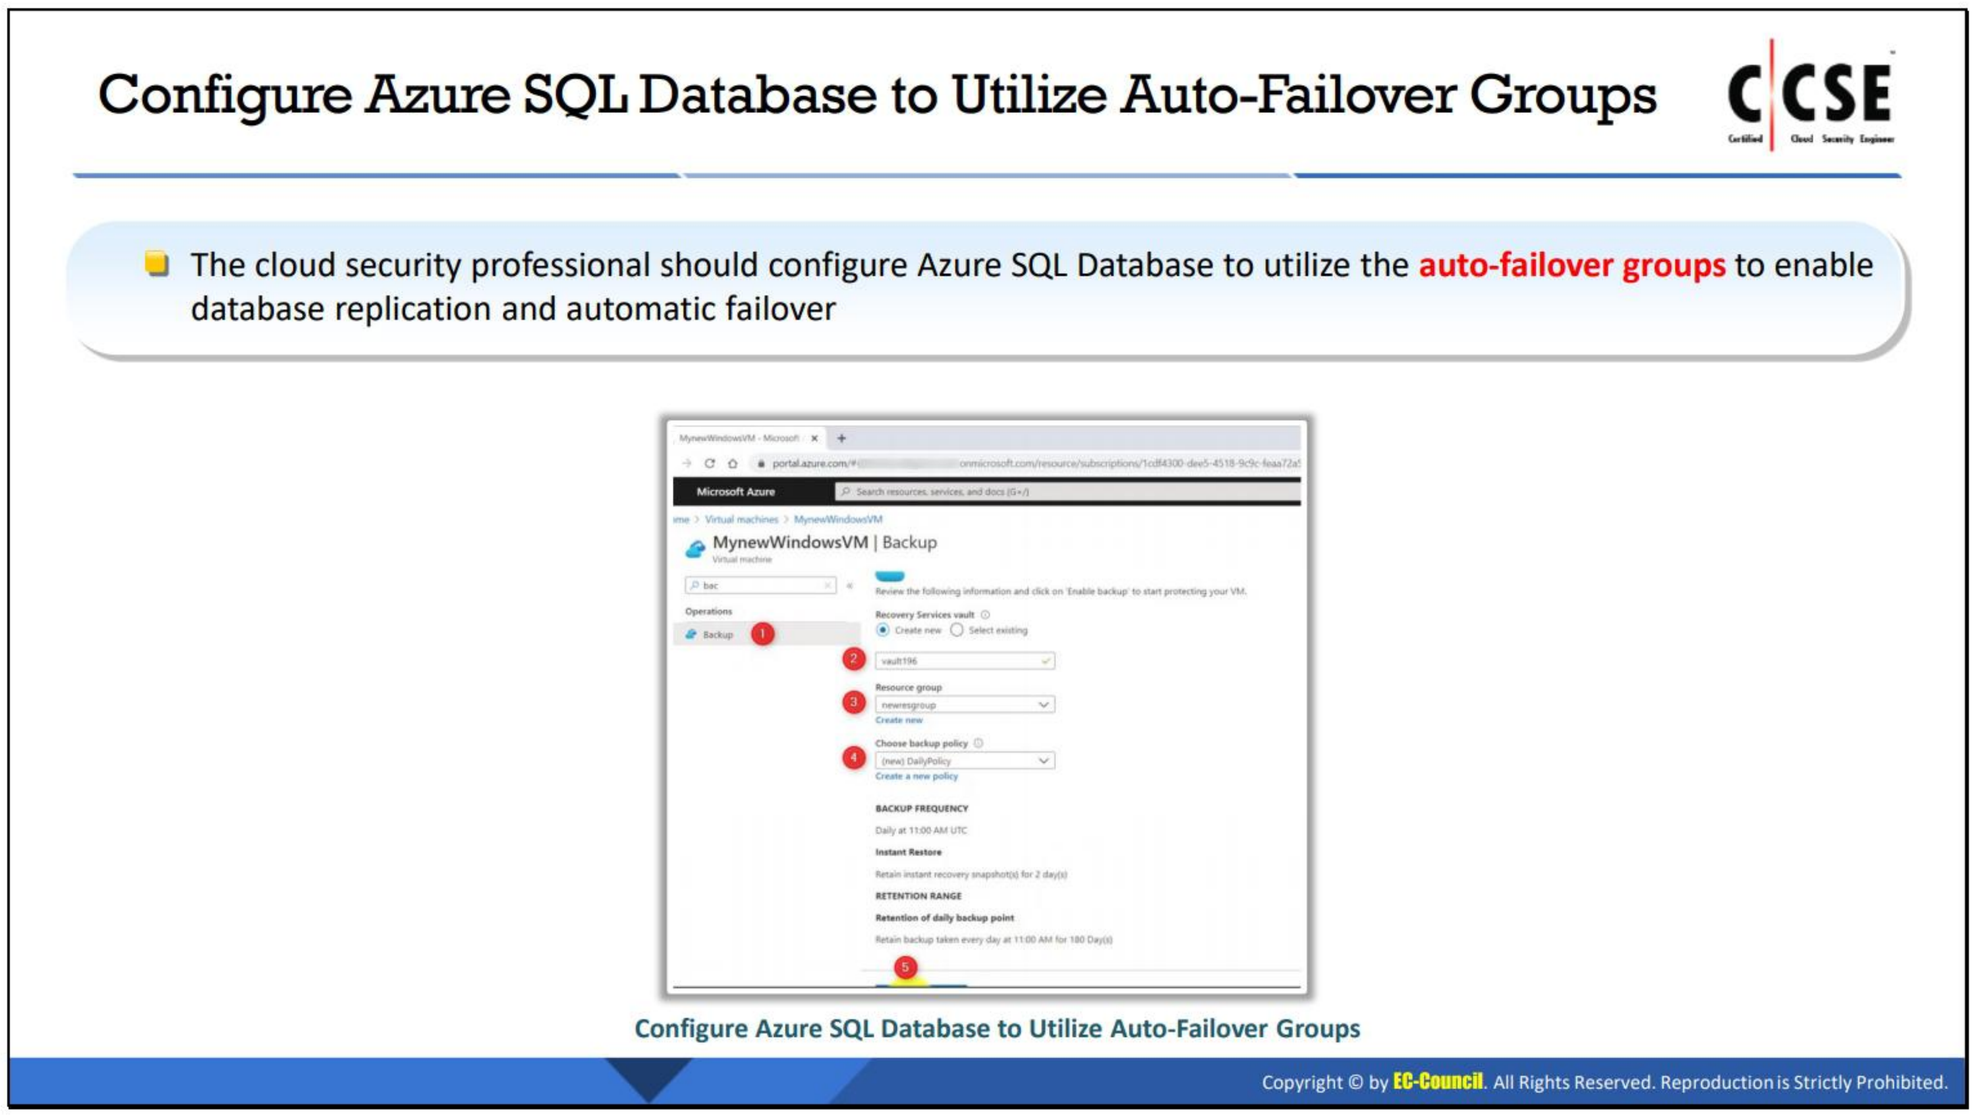

Configure Azure SQL Database to Utilize Auto-Failover Groups

Configure Azure SQL Database to Utilize Auto-Failover Groups

The Azure SQL failover group helps in failover, automatically managing replication, connectivity, etc. The cloud security professional should configure the Azure SQL Database to utilize the auto-failover groups to enable database replication and automatic failover.

When using auto-failover groups, in case an outage affects single or multiple SQL databases that are present in the group, it results in automatic failover. In the failover process, all the secondary databases that are present in the group are switched to primary databases. After completion of failover, the DNS records are updated automatically, which helps in redirecting the Azure endpoints to the new Azure region.

To Configure Azure SQL Database to Utilize Auto-Failover Groups

- Sign in to the Azure portal.
- Click on **More services**.

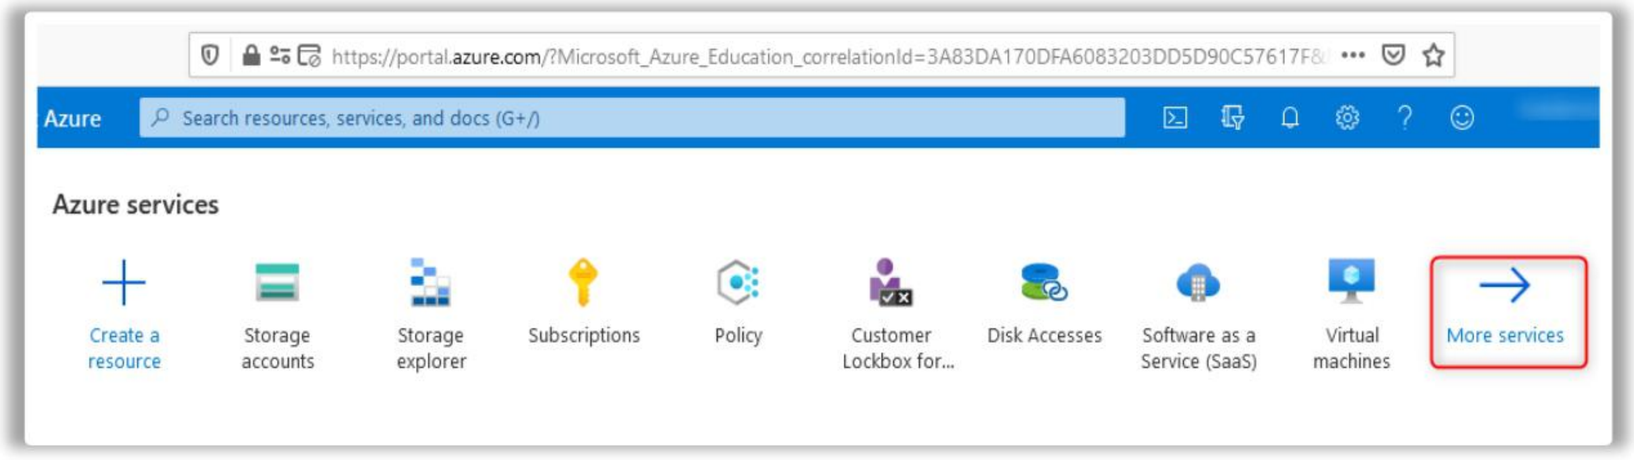

Figure 9.175: Microsoft Azure Console

- Search for SQL Server and click on the same.

Certified Cloud Security Engineer Copyright © by **EC-Council**
All Rights Reserved. Reproduction is Strictly Prohibited.

- A list of available SQL servers is displayed.

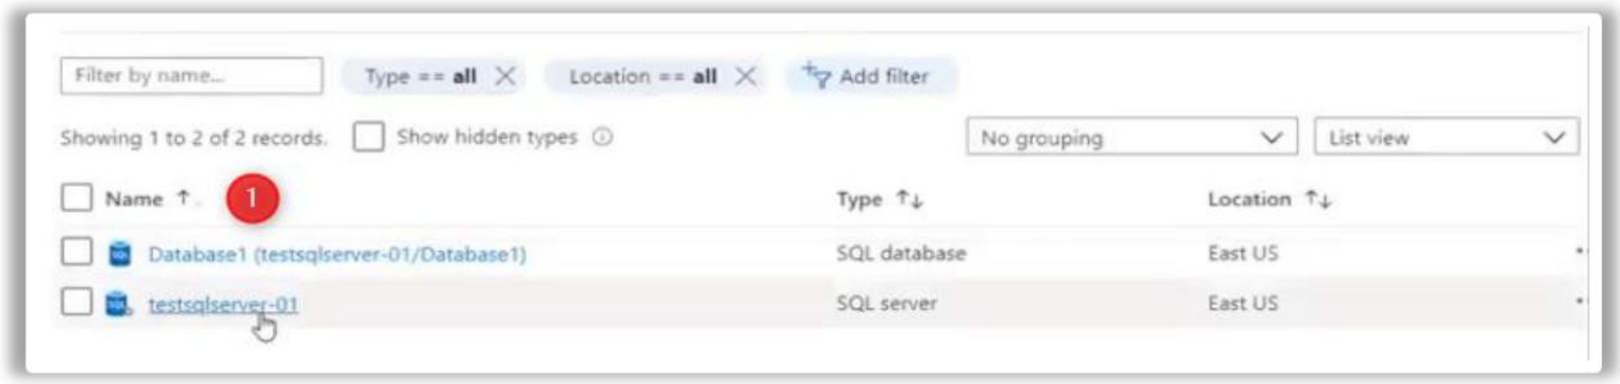

Figure 9.176: List of VM

- Select the server in which you want to configure auto-failover.
- In the left navigation panel under **Settings** go to **Failover groups**.
- Click on **+ Add group**.

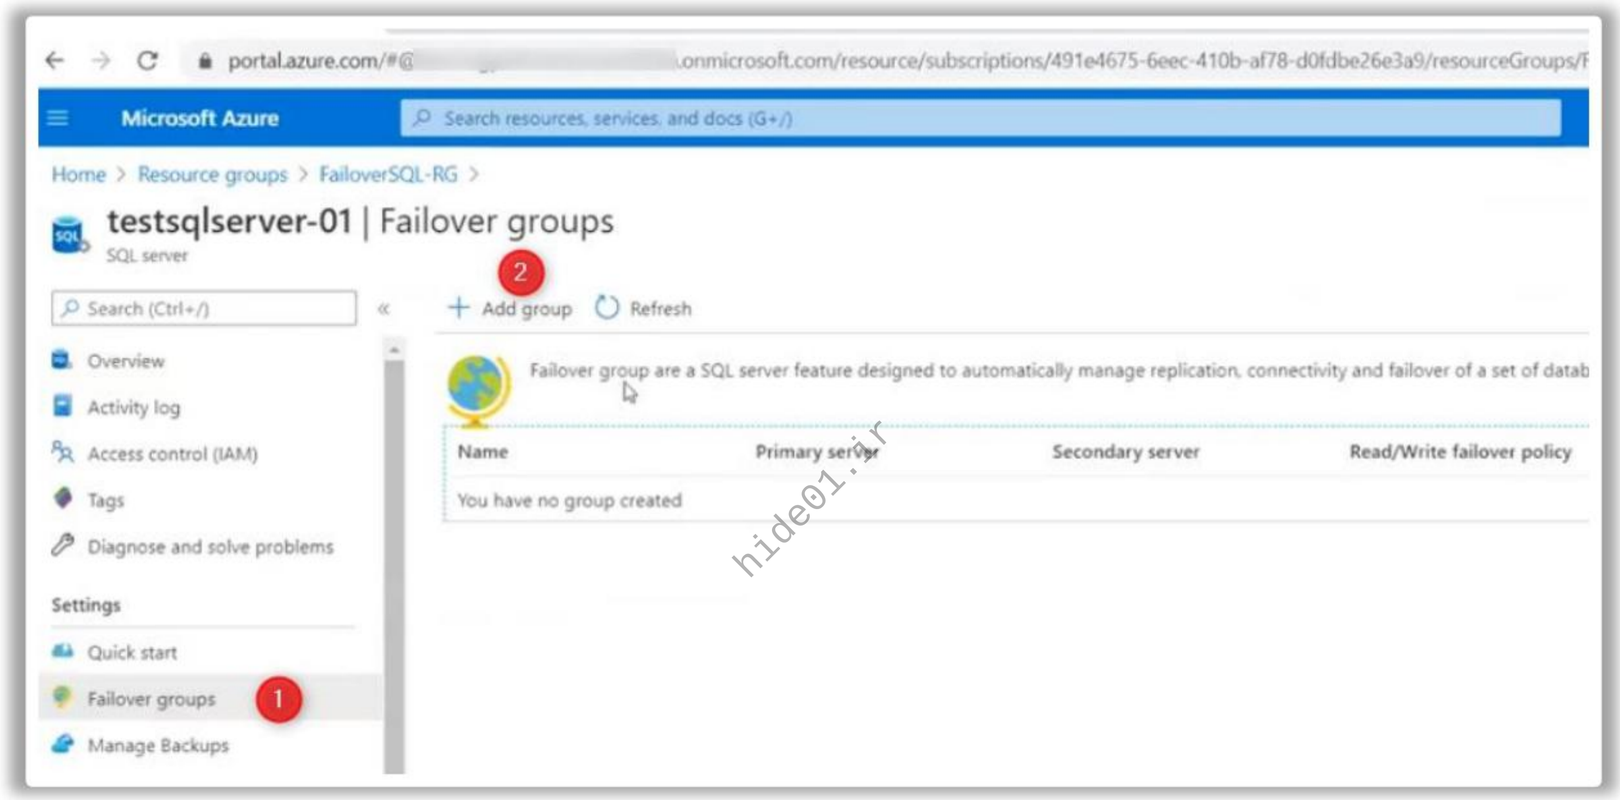

Figure 9.177: Adding group

- Enter the **Failover group name**.
- Select **Secondary server**, then click on **+ Create a new server**.

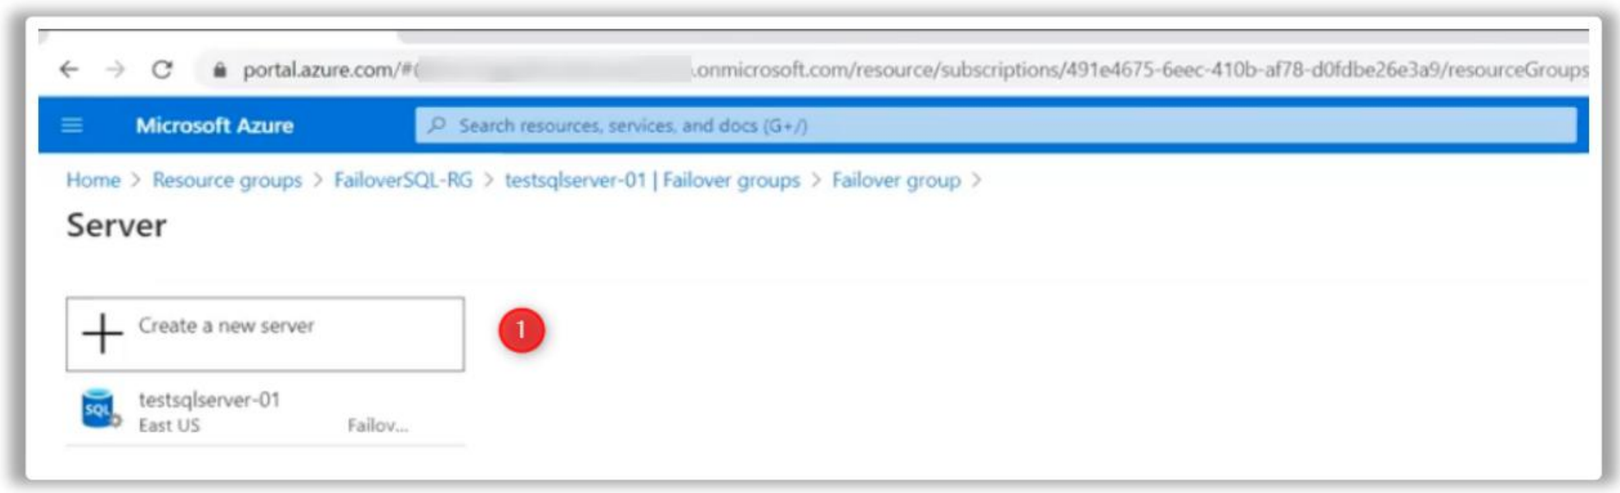

Figure 9.178: Selecting Secondary server

- Enter details for **New server**, including **Server name**, **Server admin login**, **Password**, **Confirm Password**, and **location**.

Certified Cloud Security Engineer Copyright © by **EC-Council**
All Rights Reserved. Reproduction is Strictly Prohibited.

- Now click on **Select**.

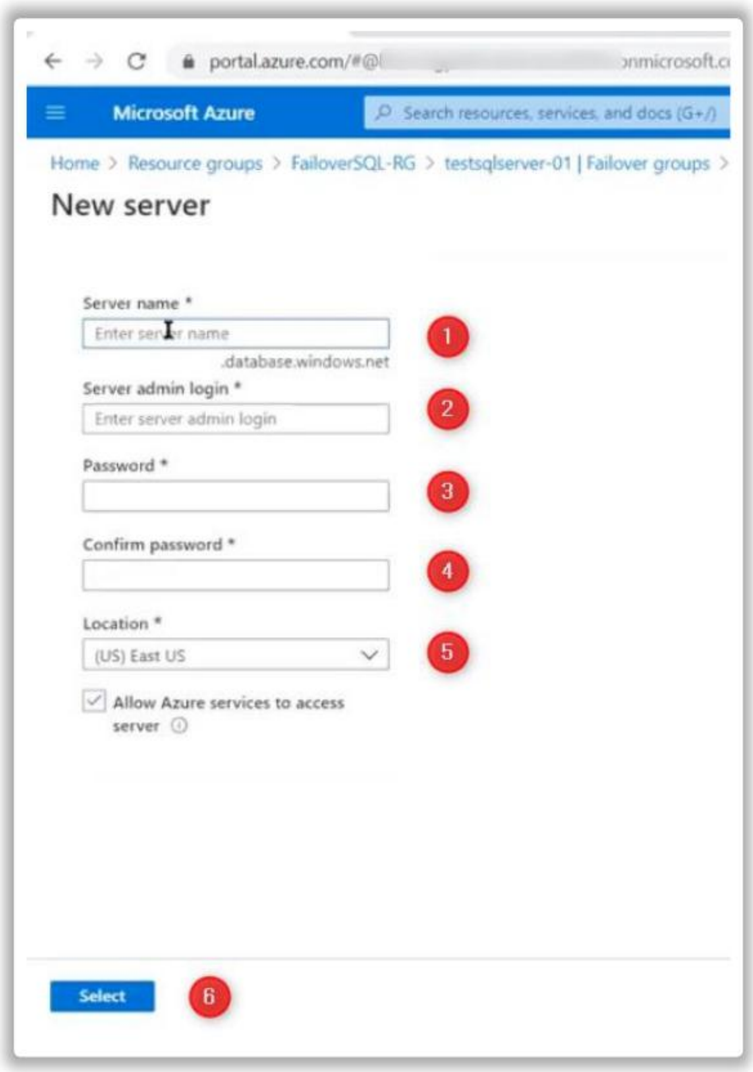

Figure 9.179: Click Select

- Select **Read/Write policy** to **Automatic** from the drop-down.
- Select the required time from the drop-down under "**Read/Write grace period (hours)**."
- Go to **Database within the group** to add a database.

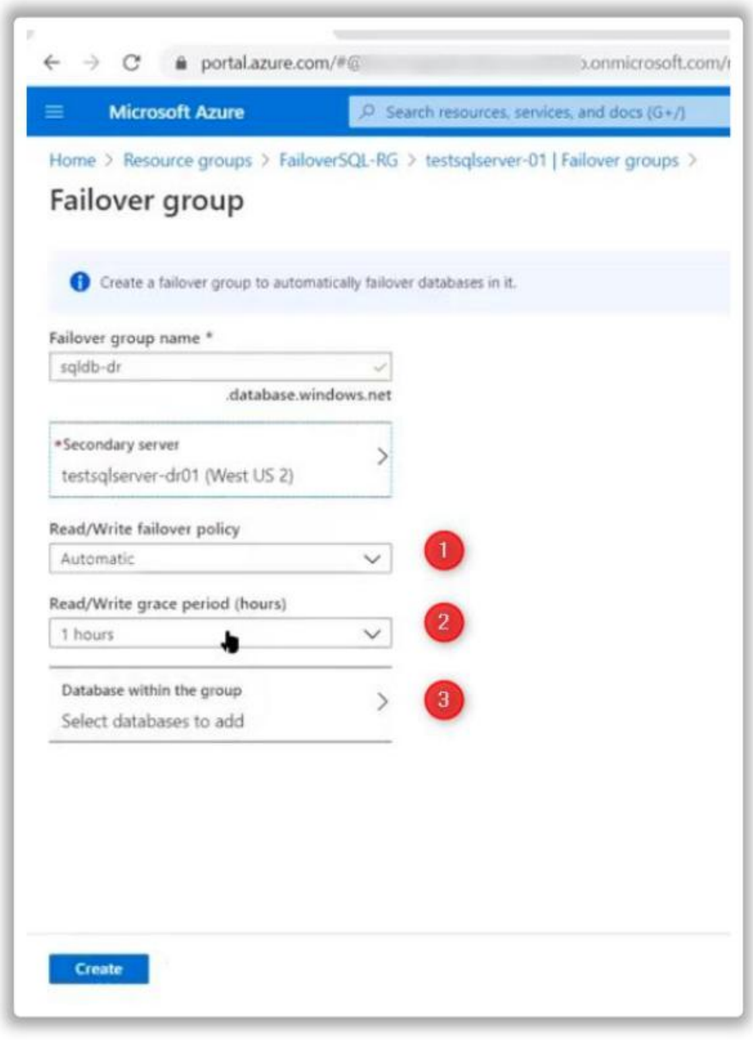

Figure 9.180: Failover group Console

Certified Cloud Security Engineer Copyright © by **EC-Council**
All Rights Reserved. Reproduction is Strictly Prohibited.

- A list of available databases is shown; select the preferred database.
- Then click on **Select**.

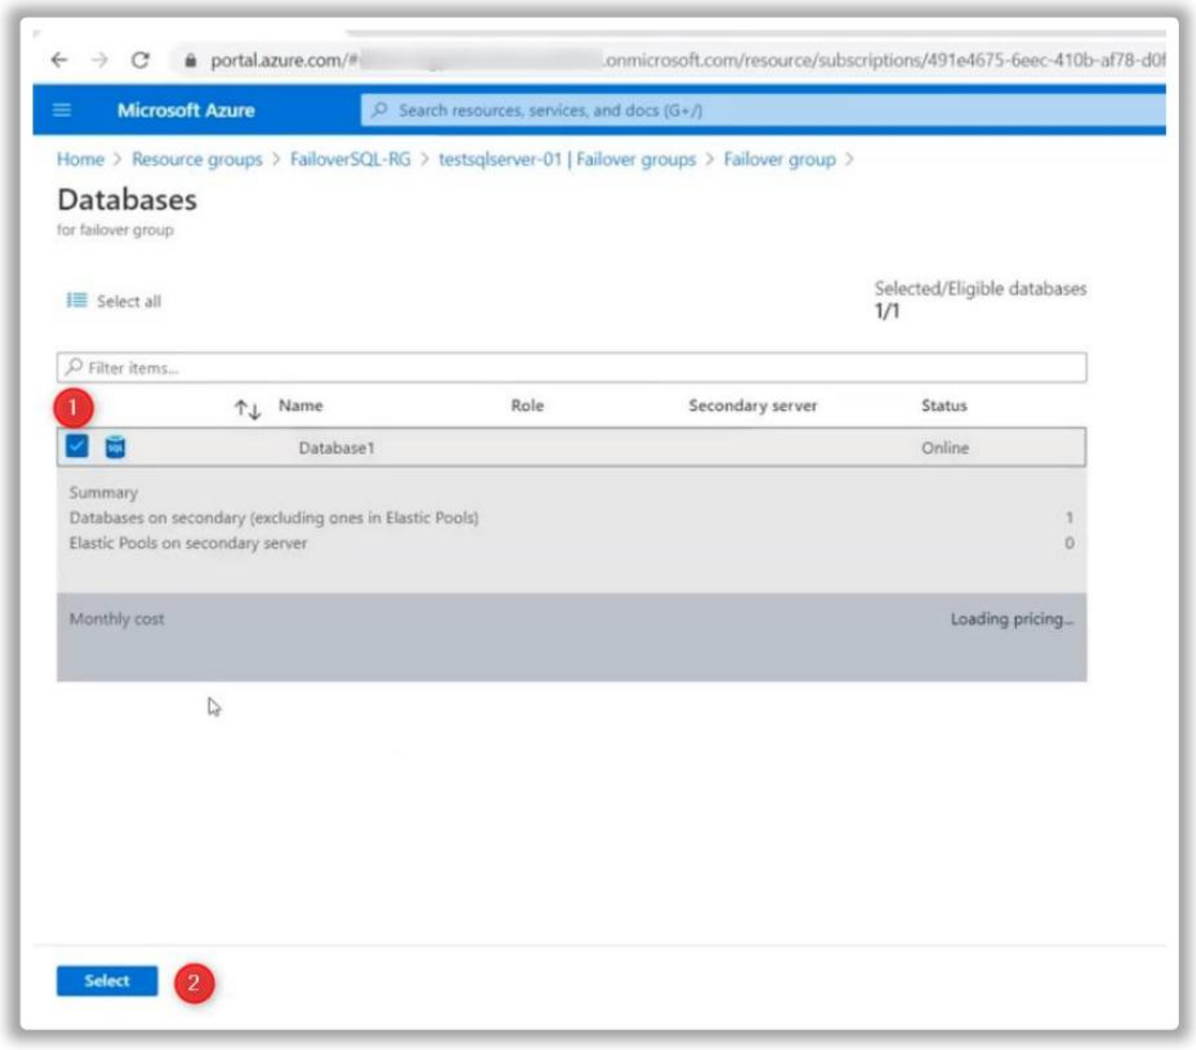

Figure 9.181: Click Select

- Now click on **Create**.

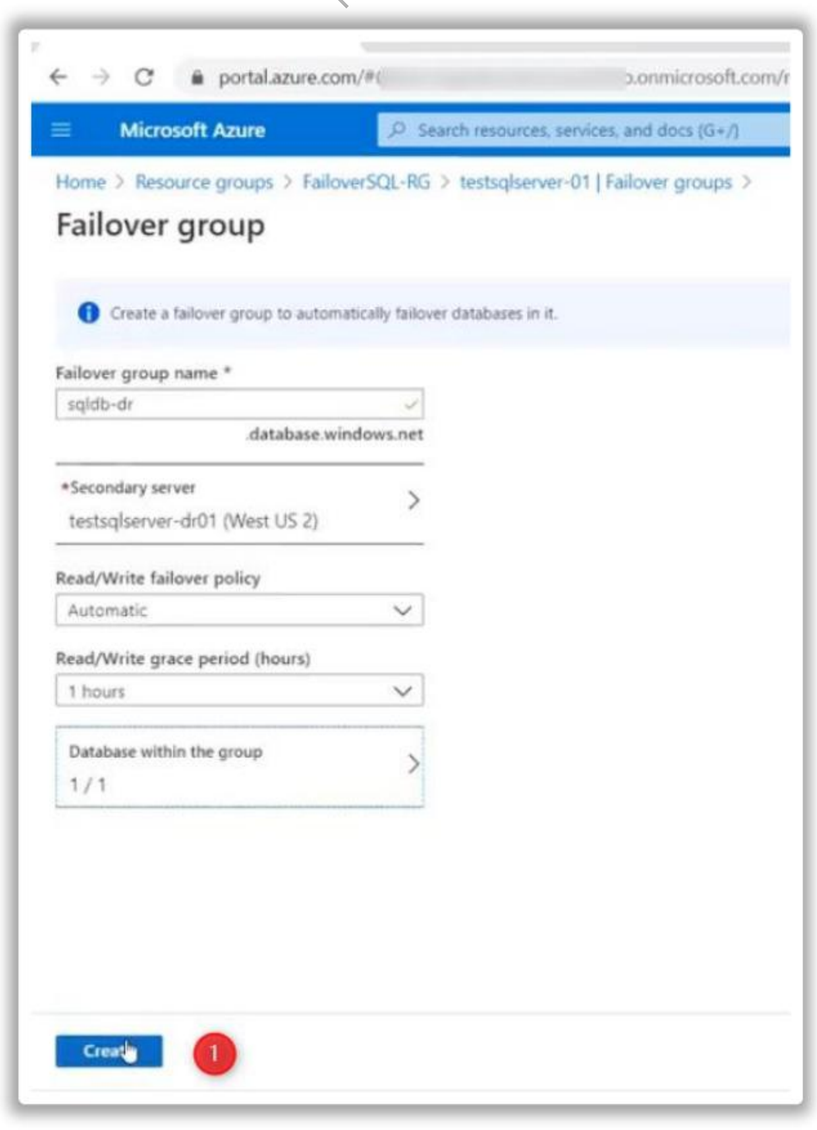

Figure 9.182: Click Create

Certified Cloud Security Engineer Copyright © by **EC-Council**
All Rights Reserved. Reproduction is Strictly Prohibited.

- The deployment is successful.

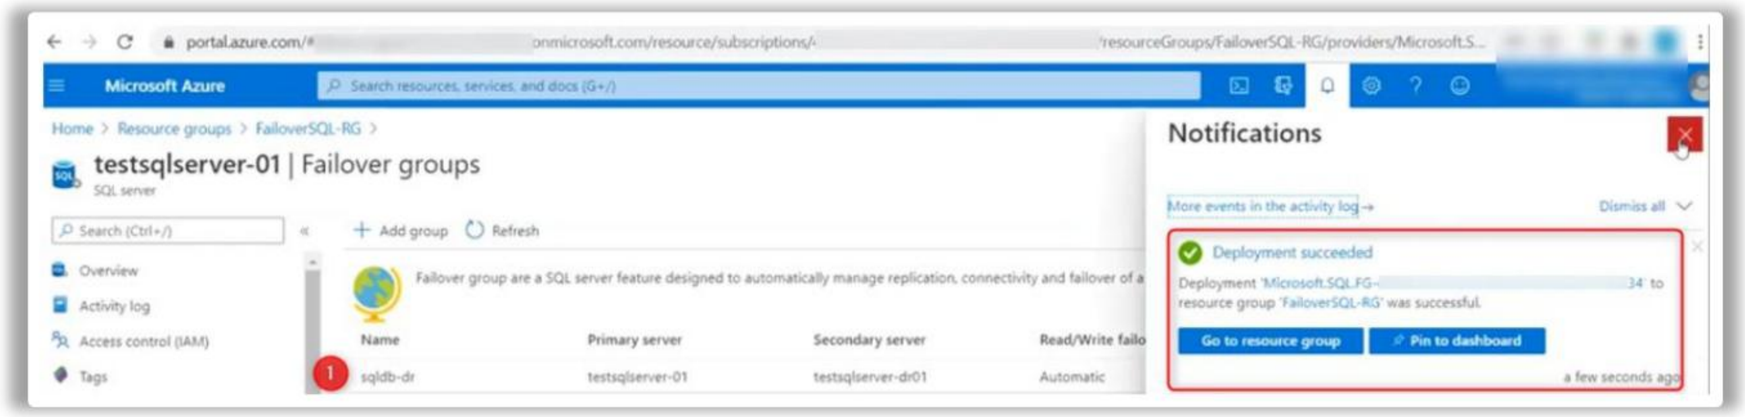

Figure 9.183: Successful Deployment

hide01.ir

Certified Cloud Security Engineer Copyright © by **EC-Council**
All Rights Reserved. Reproduction is Strictly Prohibited.

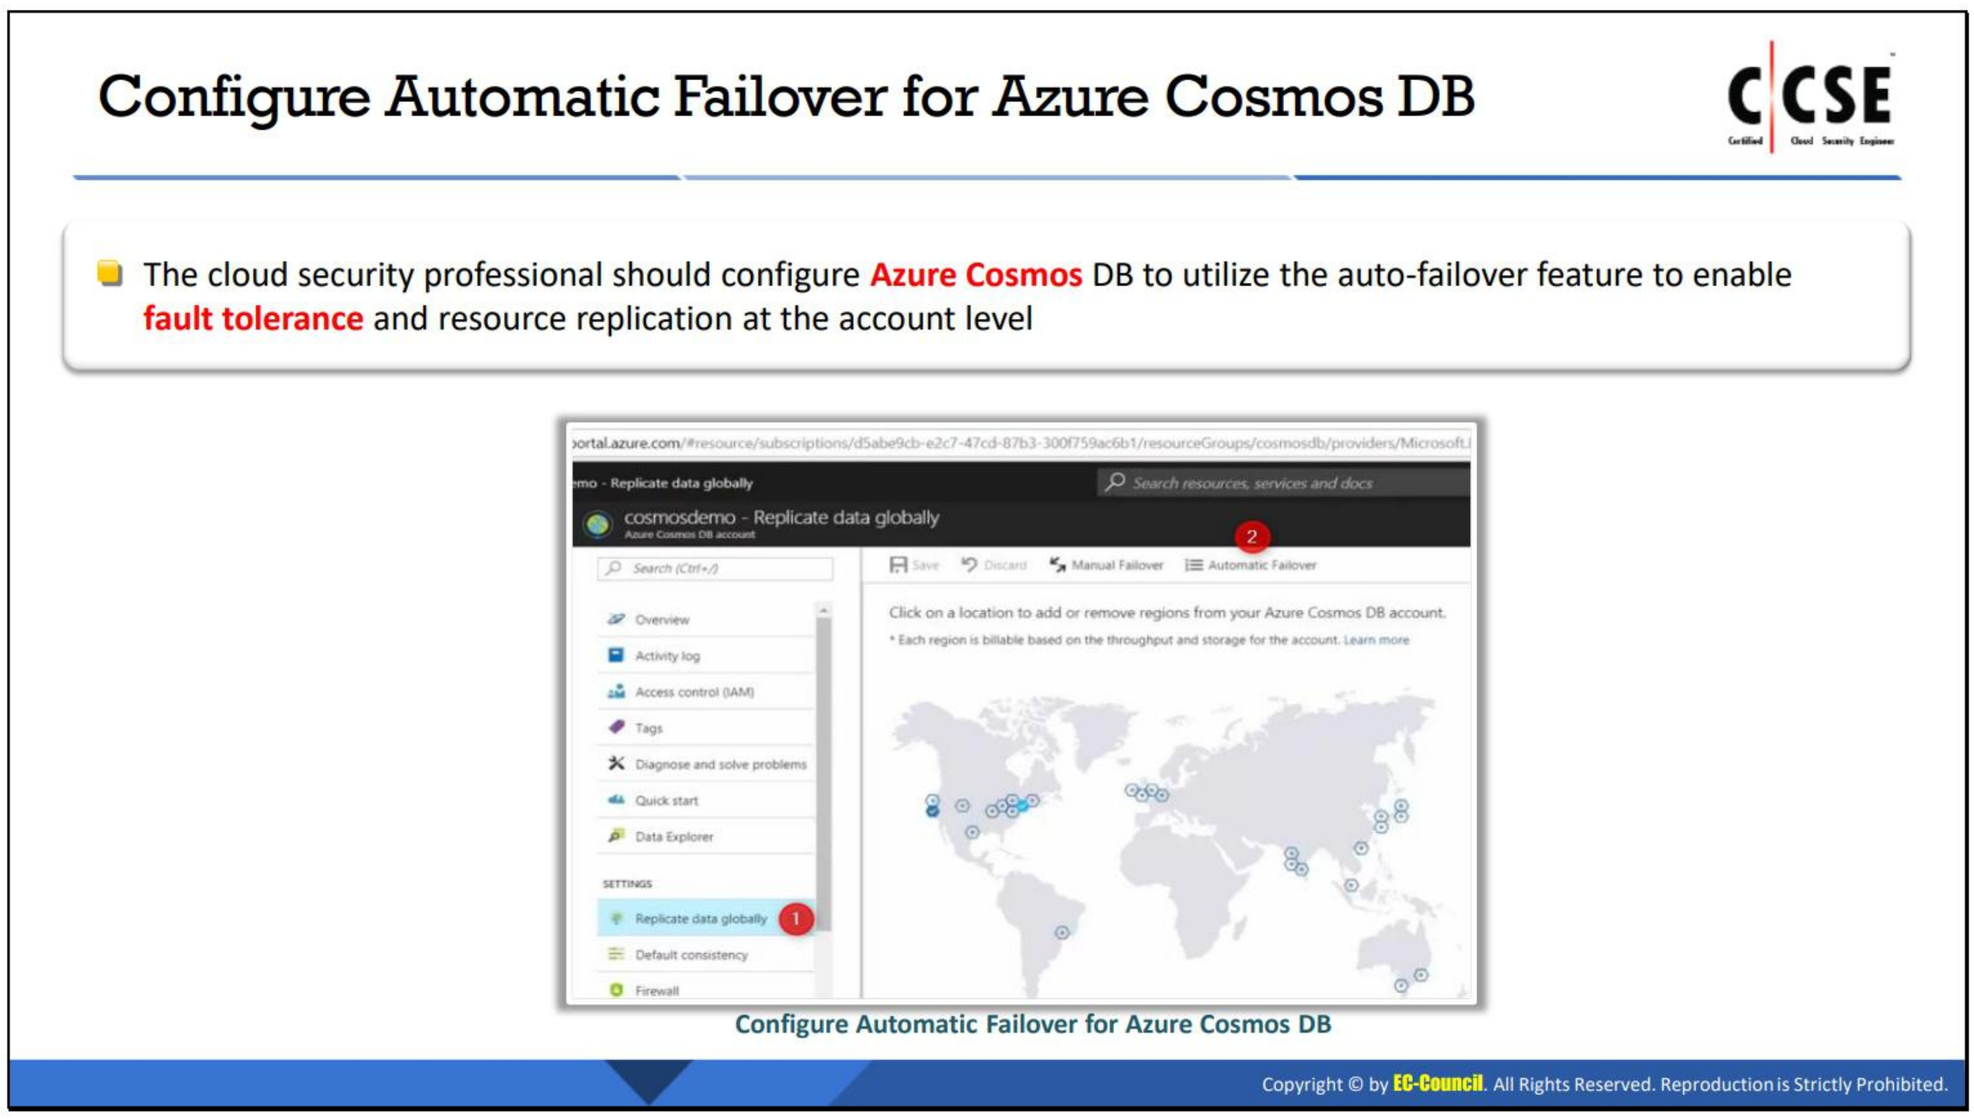

Configure Automatic Failover for Azure Cosmos DB

Configure Automatic Failover for Azure Cosmos DB

The cloud security professional should configure the Azure Cosmos DB to utilize the auto-failover feature to enable fault tolerance and resource replication at the account level. Automatic failover allows Azure Cosmos DB to failover automatically to the Azure cloud region; this is done with the highest failover priority when the source region is not available. To enable this feature, the Cosmos DB account needs to have two or more regions configured. When the Azure Cosmos DB account is configured for automatic failover, if a write region outage occurs, then the Cosmos DB account will automatically switch the secondary region as a new primary write region.

To Enable Automatic Failover

- Sign into the Azure portal.
- Click on **More services**.

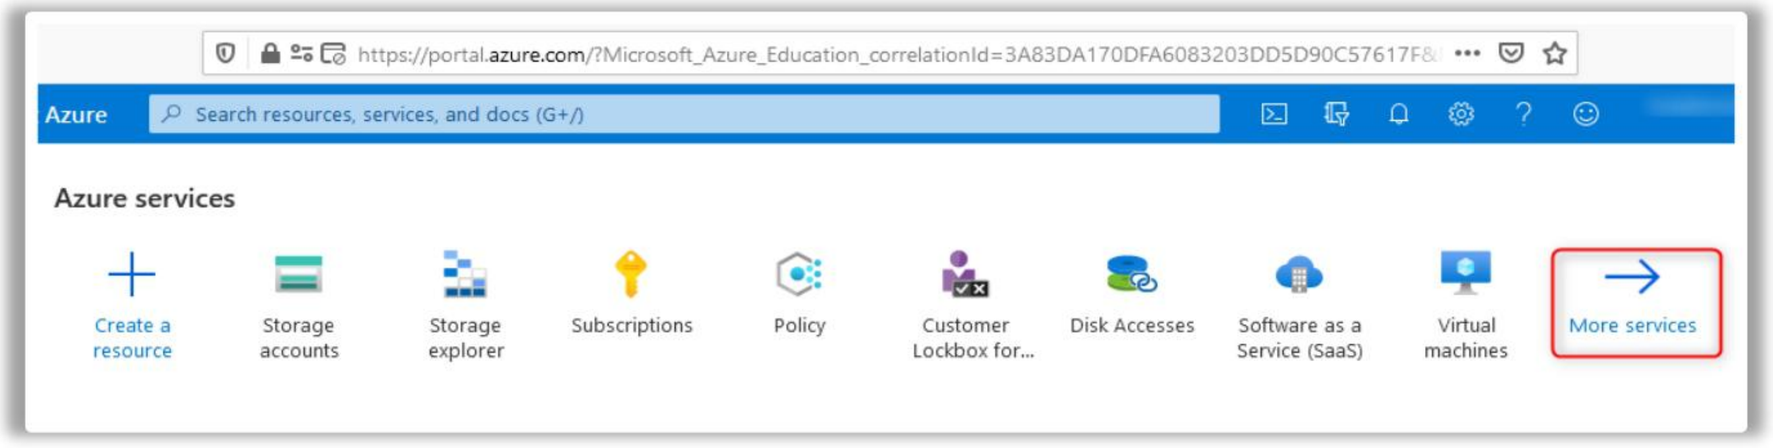

Figure 9.184: Microsoft Azure Console

Certified Cloud Security Engineer Copyright © by **EC-Council**
All Rights Reserved. Reproduction is Strictly Prohibited.

- Search for Azure Cosmos DB and select the same.

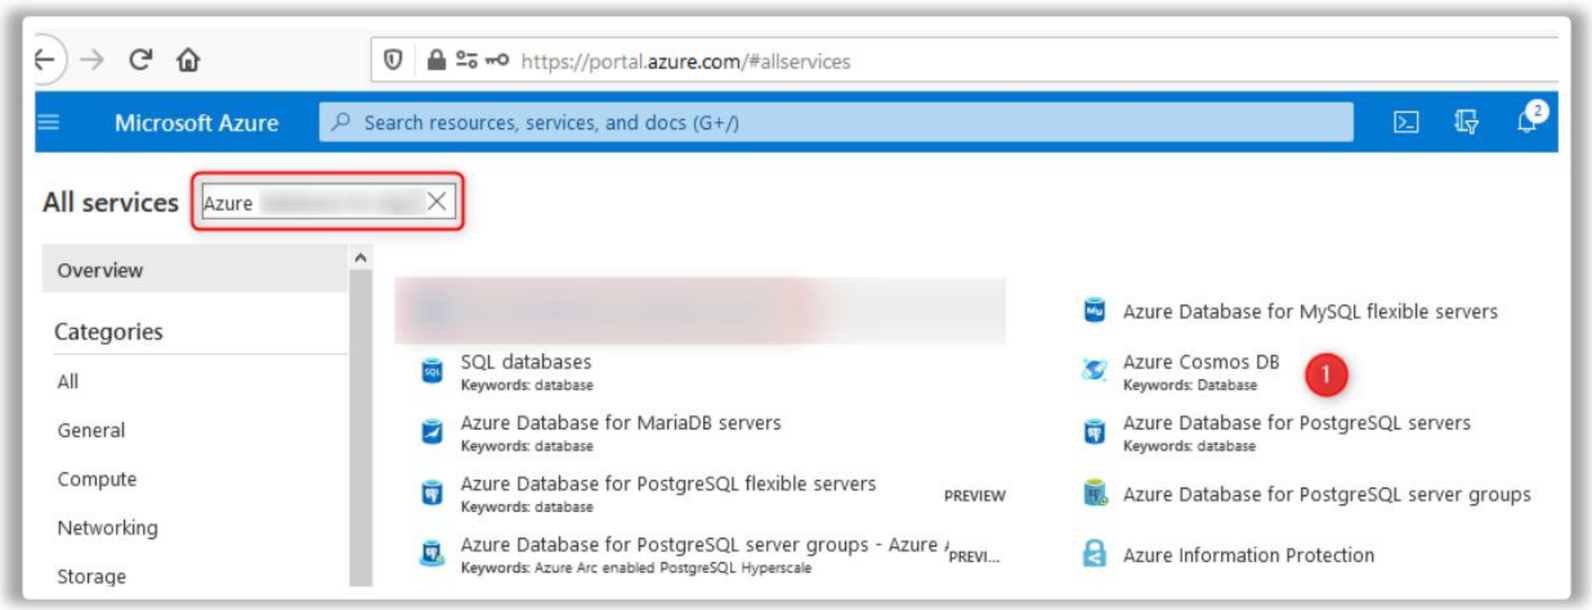

Figure 9.185: All Service Console

- A list of available Azure Cosmos DBs is listed; select the one in which you need to configure automatic failover.
- In the left navigation panel, under **SETTINGS**, go to **Replicate data globally.**
- At least two or more regions should be configured to use the automatic failover feature.
- Then click on **Automatic Failover**.

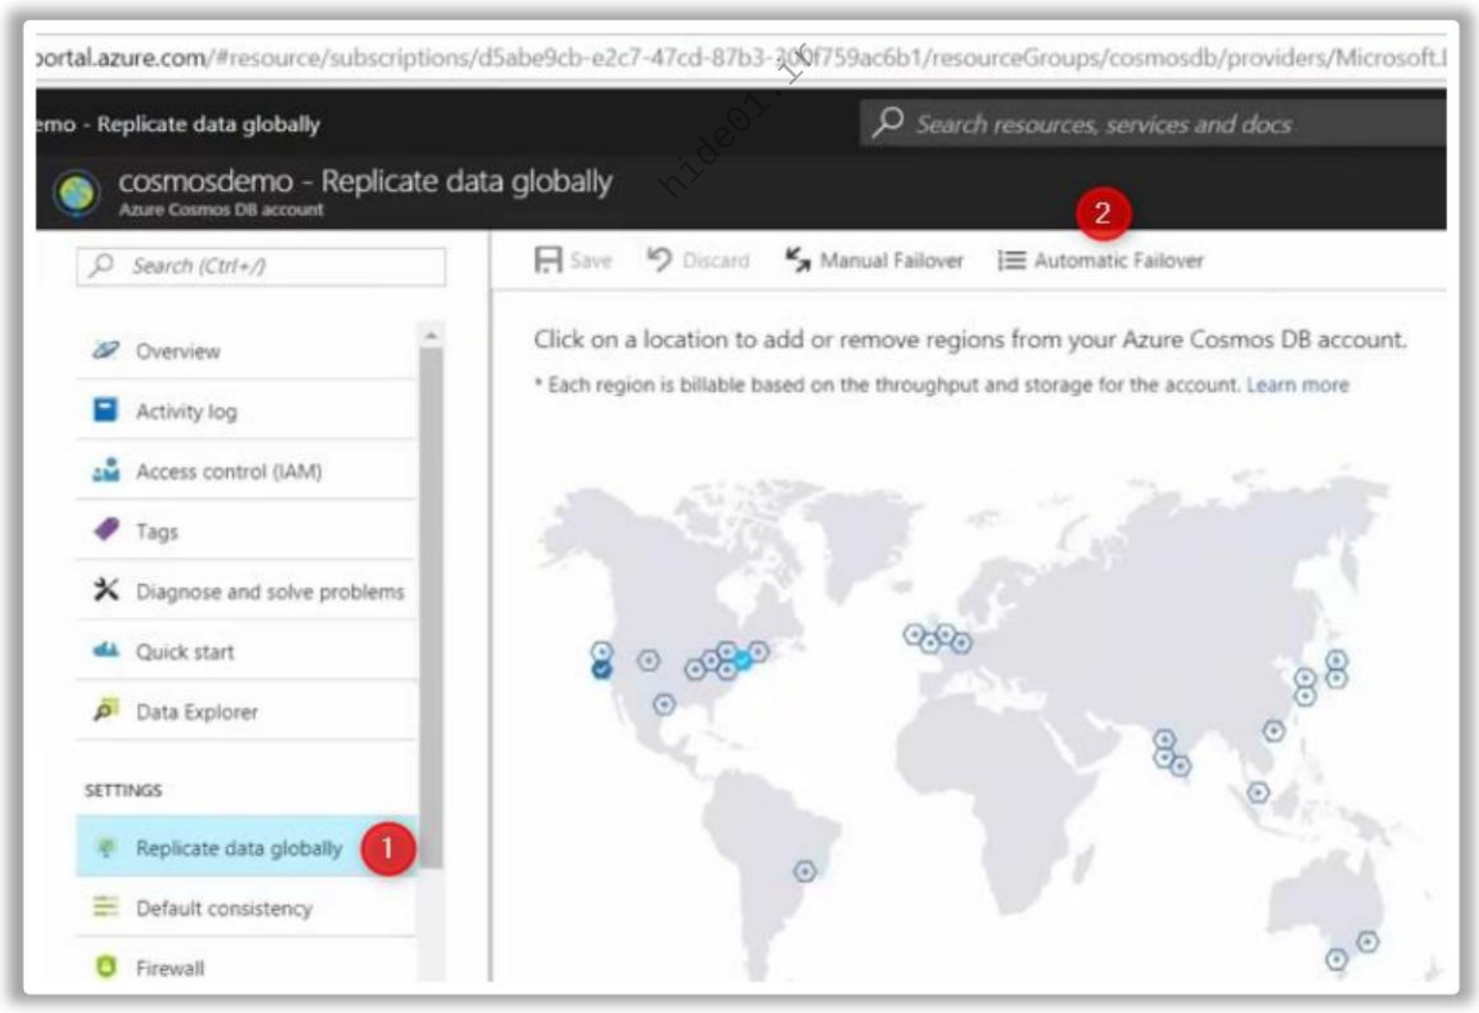

Figure 9.186: Click Automatic Failover

- Select **ON** under **Enable Automatic Failover**.

Certified Cloud Security Engineer Copyright © by **EC-Council**
All Rights Reserved. Reproduction is Strictly Prohibited.

- Then click on **OK** to apply changes.

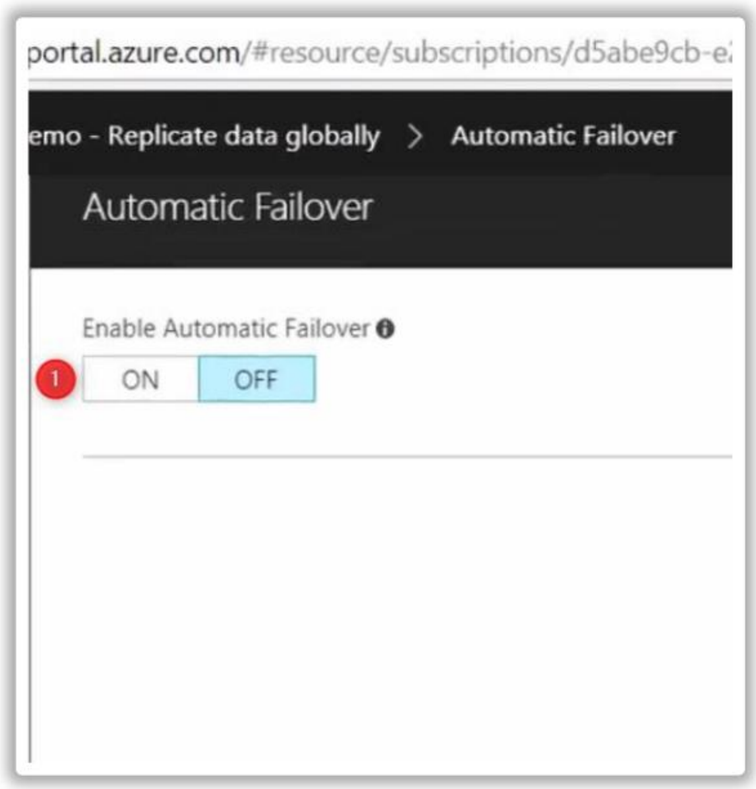

Figure 9.187: Click ON

- The automatic failover to Azure Cosmos DB has been successfully enabled.

hide01.ir

Certified Cloud Security Engineer Copyright © by **EC-Council**
All Rights Reserved. Reproduction is Strictly Prohibited.

LO#11: Discuss Azure Partner Solutions for BCDR

The objective of this section is to understand the various partner solutions for implementing BCDR in Azure.

hide01.ir

Certified Cloud Security Engineer Copyright © by **EC-Council**
All Rights Reserved. Reproduction is Strictly Prohibited.

Veeam

- Veeam provides enterprise-ready Azure **backup** and **recovery capabilities**, fully automated for protecting and managing all Azure data at scale with ease

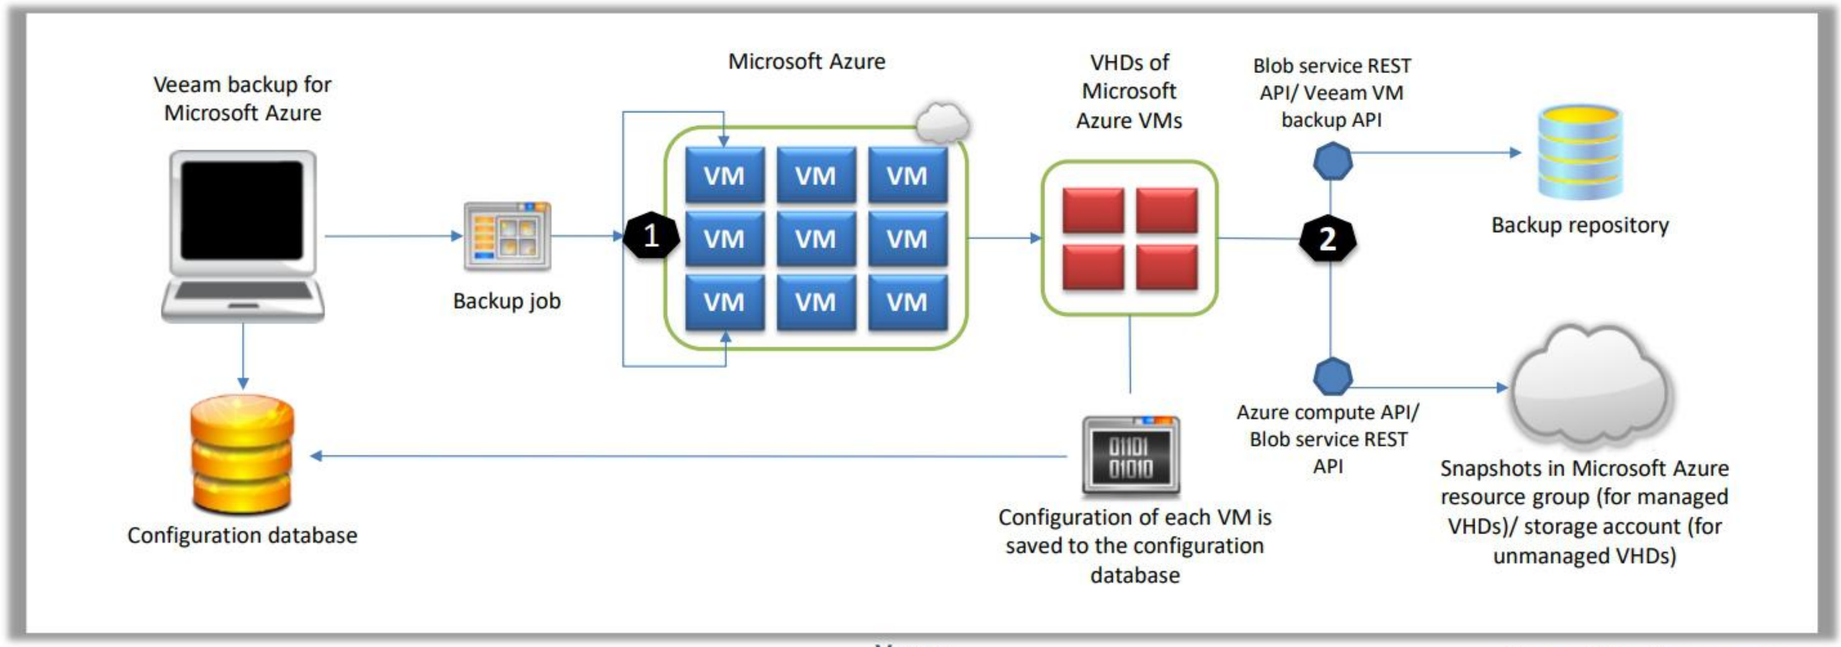

Veeam

Source: https://www.veeam.com

Copyright © by EC-Council. All Rights Reserved. Reproduction is Strictly Prohibited.

Veeam

Source: https://www.veeam.com

Veeam provides enterprise-ready Azure Backup and recovery capabilities, fully automated for protecting and managing all your Azure data at scale with ease.

Veeam automate Azure snapshots and backups with fully customizable policies and scheduling. Automatically tier backup to hot and cool Azure Blob tiers for cost-effective LTR and compliance. Zero-data loss backup and recovery without shutting down VMs or disconnecting volumes. The key features of Veeam are as follows:

- **Reliable backup and recovery**: Azure-native backup and recovery; fully automated for simple yet powerful data protection that will not let customer down.
- **Actualize cloud savings**: Industry-first cost management tools and low-cost tiering options eliminate the potential of cloud overspend.
- **Secure and compliant protection**: Fight back against ransomware and cybersecurity threats while ensuring compliant data protection and archive.
- **Hybrid-cloud ready**: Unify Azure data protection and management with all hybrid- and multi-cloud data from a common control plane.

Certified Cloud Security Engineer Copyright © by **EC-Council**
All Rights Reserved. Reproduction is Strictly Prohibited.

Commvault

- Just as Azure Stack and Azure provide a single platform for **hybrid cloud**, Commvault mirrors this approach by providing a **single platform** to manage, migrate, and protect data and applications across Azure, Azure Stack, and on the edge—creating the perfect hybrid cloud

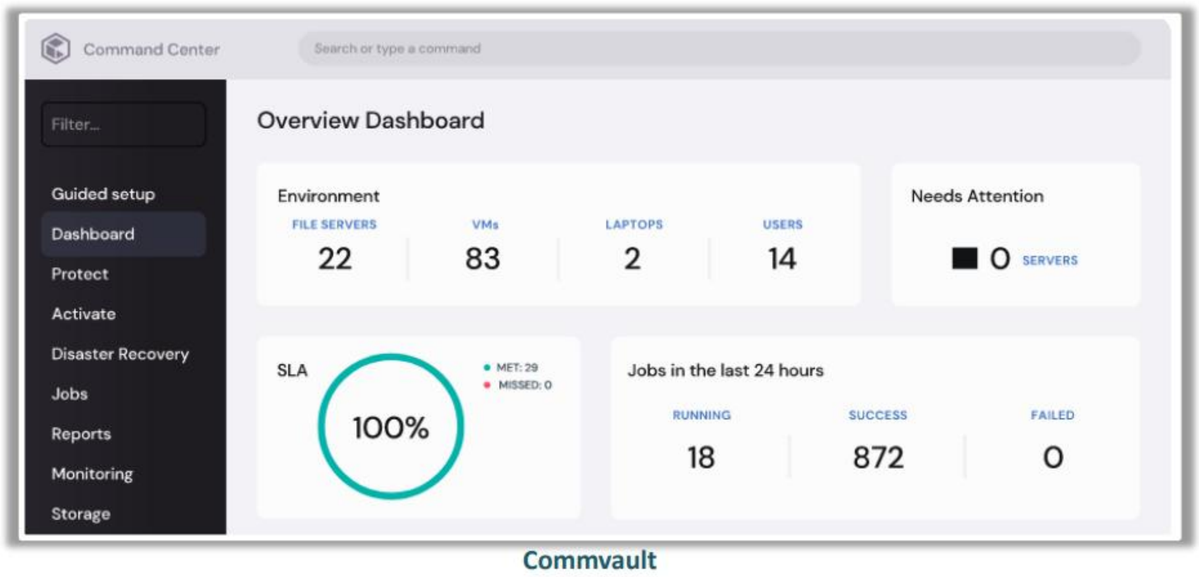

Commvault

***Source**: https://www.commvault.com*

Copyright © by **EC-Council**. All Rights Reserved. Reproduction is Strictly Prohibited.

Commvault

Source: https://www.commvault.com

Just as Azure Stack and Azure provide a single platform for hybrid cloud, Commvault mirrors this approach by providing a single platform to manage, migrate, and protect data and applications across Azure, Azure Stack, and on the edge—creating the perfect hybrid cloud.

With the advantage of running fully native in Azure Stack, Commvault enables immediate migration from wherever your data is to and from Azure, Azure Stack, or anywhere else data resides. Some key features of Commvault are as follows:

- Fast replication of VMs, applications, and storage snapshots
- Automation for DR scripts, compliance reporting and outage discovery
- Flexible copy data management for verifiable recovery and repurposing for testing and DevOps
- Flexible sub-minute RPOs and RTOs to balance mission criticality of workloads
- Security scans of replicas at the DR site with resilient ransomware protection
- One-click failover, failback, and recovery

Certified Cloud Security Engineer Copyright © by **EC-Council**
All Rights Reserved. Reproduction is Strictly Prohibited.

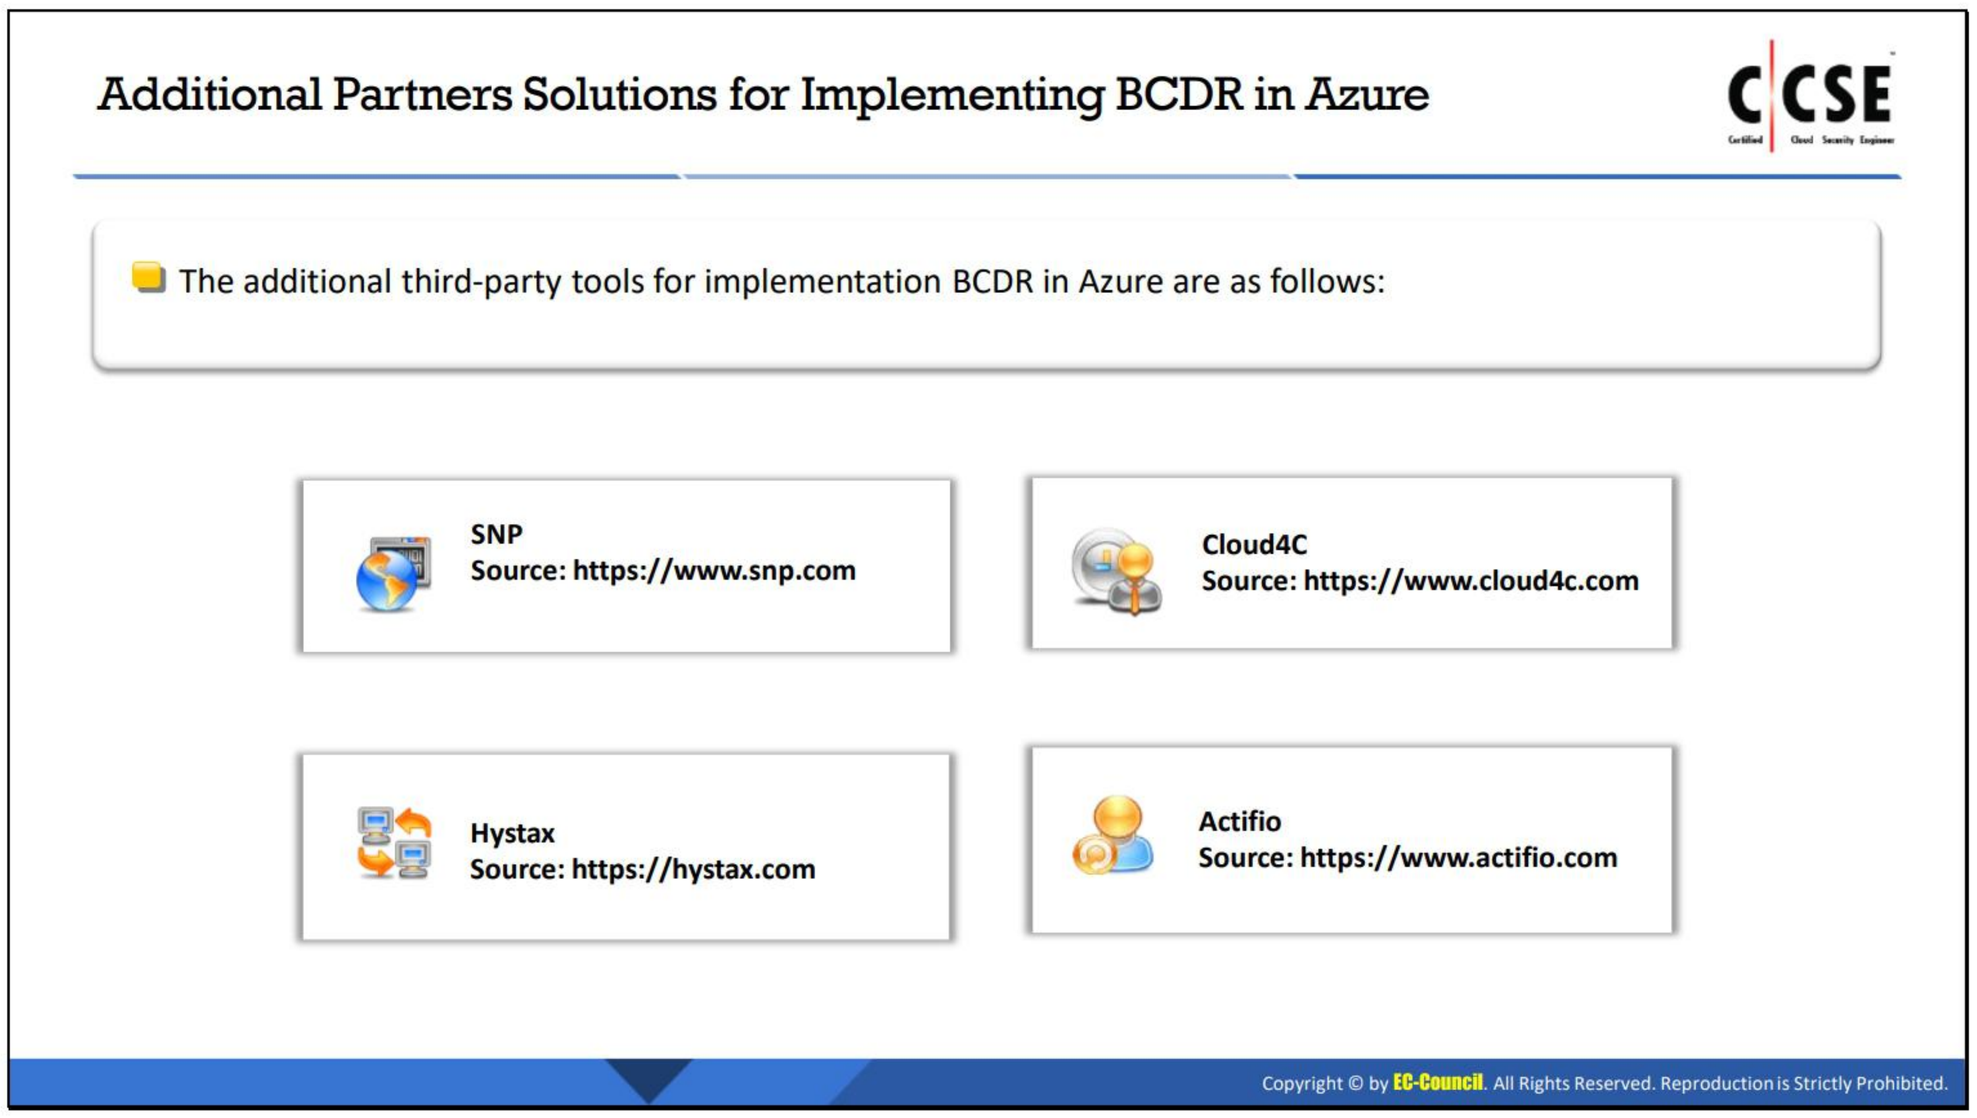

Additional Partners Solutions for Implementing BCDR in Azure

SNP

Source: https://www.snp.com

pride01.ir

SNP's backup and disaster recovery service provides simple, secure, and cost-effective solutions to backup and recover data and applications from the Microsoft Azure cloud. SNP's disaster recovery service can reliably recover single-tier or multi-tier on-premises applications, without the need for the organization to invest in duplicate systems or additional staff. Leveraging ASR, SNP experts help implement a seamless DR and backup solution by managing and orchestrating replication, failover, and failback of on-premises machines, and Azure VMs and integrate them to Azure, exceeding the capabilities of most on-premises disaster recovery solutions, at a lower total cost of ownership (TCO) than competing solutions.

Cloud4C

Source: https://www.cloud4c.com

Cloud4C's DRaaS framework is designed to meet the demands of enterprises such as delivering uninterrupted services and brinCloud4C keeps RPOs and RTOs in sync and under budget for an enterprise. Moreover, it offers a compliance-ready and single SLA on the Azure platform. Cloud4C enables orchestrated DR functions, including failover, failback, testing, and reverse protection.

Hystax

Source: https://hystax.com

Hystax Acura is a cloud-native disaster recovery solution with robust RPO and RTO, consistent replication, and storage agnostic snapshots. Hystax Acura Disaster Recovery is available for end-customers and partners in the forms of DRaaS platform, on-premises disaster recovery and a

Certified Cloud Security Engineer Copyright © by **EC-Council**
All Rights Reserved. Reproduction is Strictly Prohibited.

cloud backup solution. By using Hystax Acura for failover between AZs in Microsoft Azure cloud, the customer can improve the overall availability of workloads and ensure resiliency on Azure.

Actifio

Source: https://www.actifio.com

Enterprises use Actifio for cloud backup to Microsoft Azure (Azure) and to reduce their data center footprint and costs. Using a single pane of glass for all workloads, they can also protect Azure resident workloads, perform one-click DR orchestration, and lower RTO and operational burden. Lastly, they can reuse the backups to provision rapid database clones for test/dev/analytics in Azure.

hide01.ir

Certified Cloud Security Engineer Copyright © by **EC-Council**
All Rights Reserved. Reproduction is Strictly Prohibited.

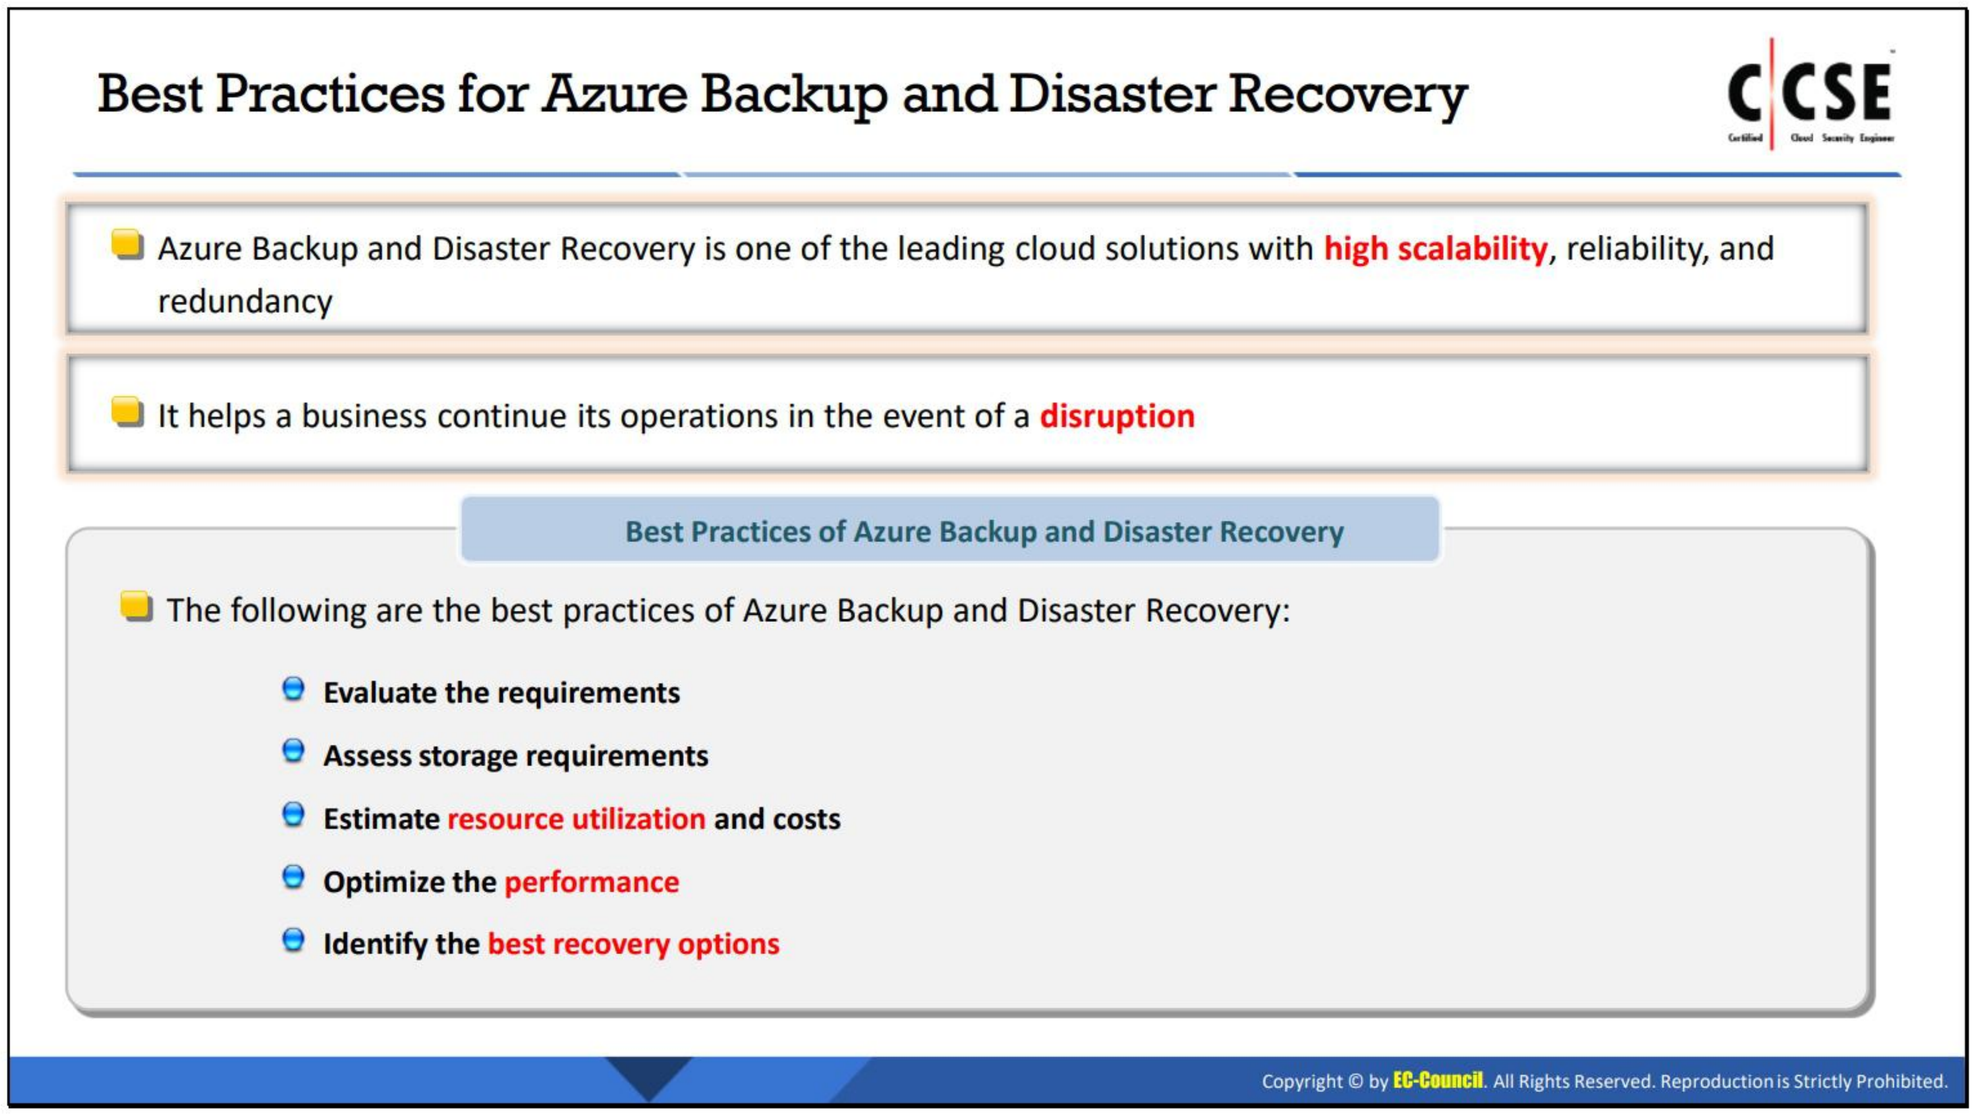

Best Practices of Azure Backup and Disaster Recovery

As the business grows and generates large amounts of data for running the business operations, on-premises storage may seem expensive, inefficient, and unsafe. This can have a drastic effect on business operations. The business must have a business backup and recovery solution to implement a BCDR. An efficient way to handle business-critical data is to replicate the data to an environment other than the source.

Azure Backup and Disaster Recovery is one of the leading cloud solutions with high scalability, reliability, and redundancy. It helps a business continue its operations in the event of a disruption.

The following are the best practices of Azure Backup and Disaster Recovery:

- **Evaluate the requirements:** There are various recover solutions to serve different requirements. Azure itself offers different solutions for recovery, backup, and archiving. The ASR helps keep business applications and workloads functional during an outage or disaster. To archive on-premises data, Azure Blob storage can be used. In contrast, Azure Backup provides a solution to backup on-premises or cloud data to Azure cloud. Hence, it is important to evaluate requirements before deciding on a disaster recovery solution.
- **Assess storage requirements:** Azure Backup can be used to backup files, folders, SQL servers, system states, and VMs. Azure Backup is used to automate storage, reduce the dependency on on-premises backups, and secure data at rest and during transit. There are certain limitations to using Azure Backup with Microsoft Azure Recovering Services for backing up files. Azure Backup cannot accommodate read-only volumes, Bitlocker protected volumes, network shares, offline volumes, and removable media. To backup files using Azure Backup service, NTFS is the only file system.

Certified Cloud Security Engineer Copyright © by **EC-Council**
All Rights Reserved. Reproduction is Strictly Prohibited.

- **Estimate resource utilization and costs:** Azure Backup offers flexible pricing plans for customers. The cost is based on the data stored, the number of instances that need to be protected, and the storage model. Hence, users should estimate factors such as required bandwidth, backup schedules, policies to be created, and items to be protected.
- **Optimize the performance:** The performance of Azure Backup can be optimized by using appropriate policy configuration. As the VM disks are backed-up in parallel, users must ensure to backup VMs at different time intervals. The delay between creating snapshots and copying them to the vault should also be considered.
- **Identify the best recovery options:** Azure Backup provides the ability to restore the workloads that are backed-up. It is not only possible to select a restore point, but it also offers various restore options. From the restore point, users can make a new VM or they can restore a VM disk to create a new VM. It is also possible to restore and replace a disk. This option is suitable for customizing new VMs. The purpose of the recovery is the basis for choosing the recovery option. To improve the performance of backup and recovery, the Instance Restore option can be used.

hide01.ir

Certified Cloud Security Engineer Copyright © by **EC-Council**
All Rights Reserved. Reproduction is Strictly Prohibited.

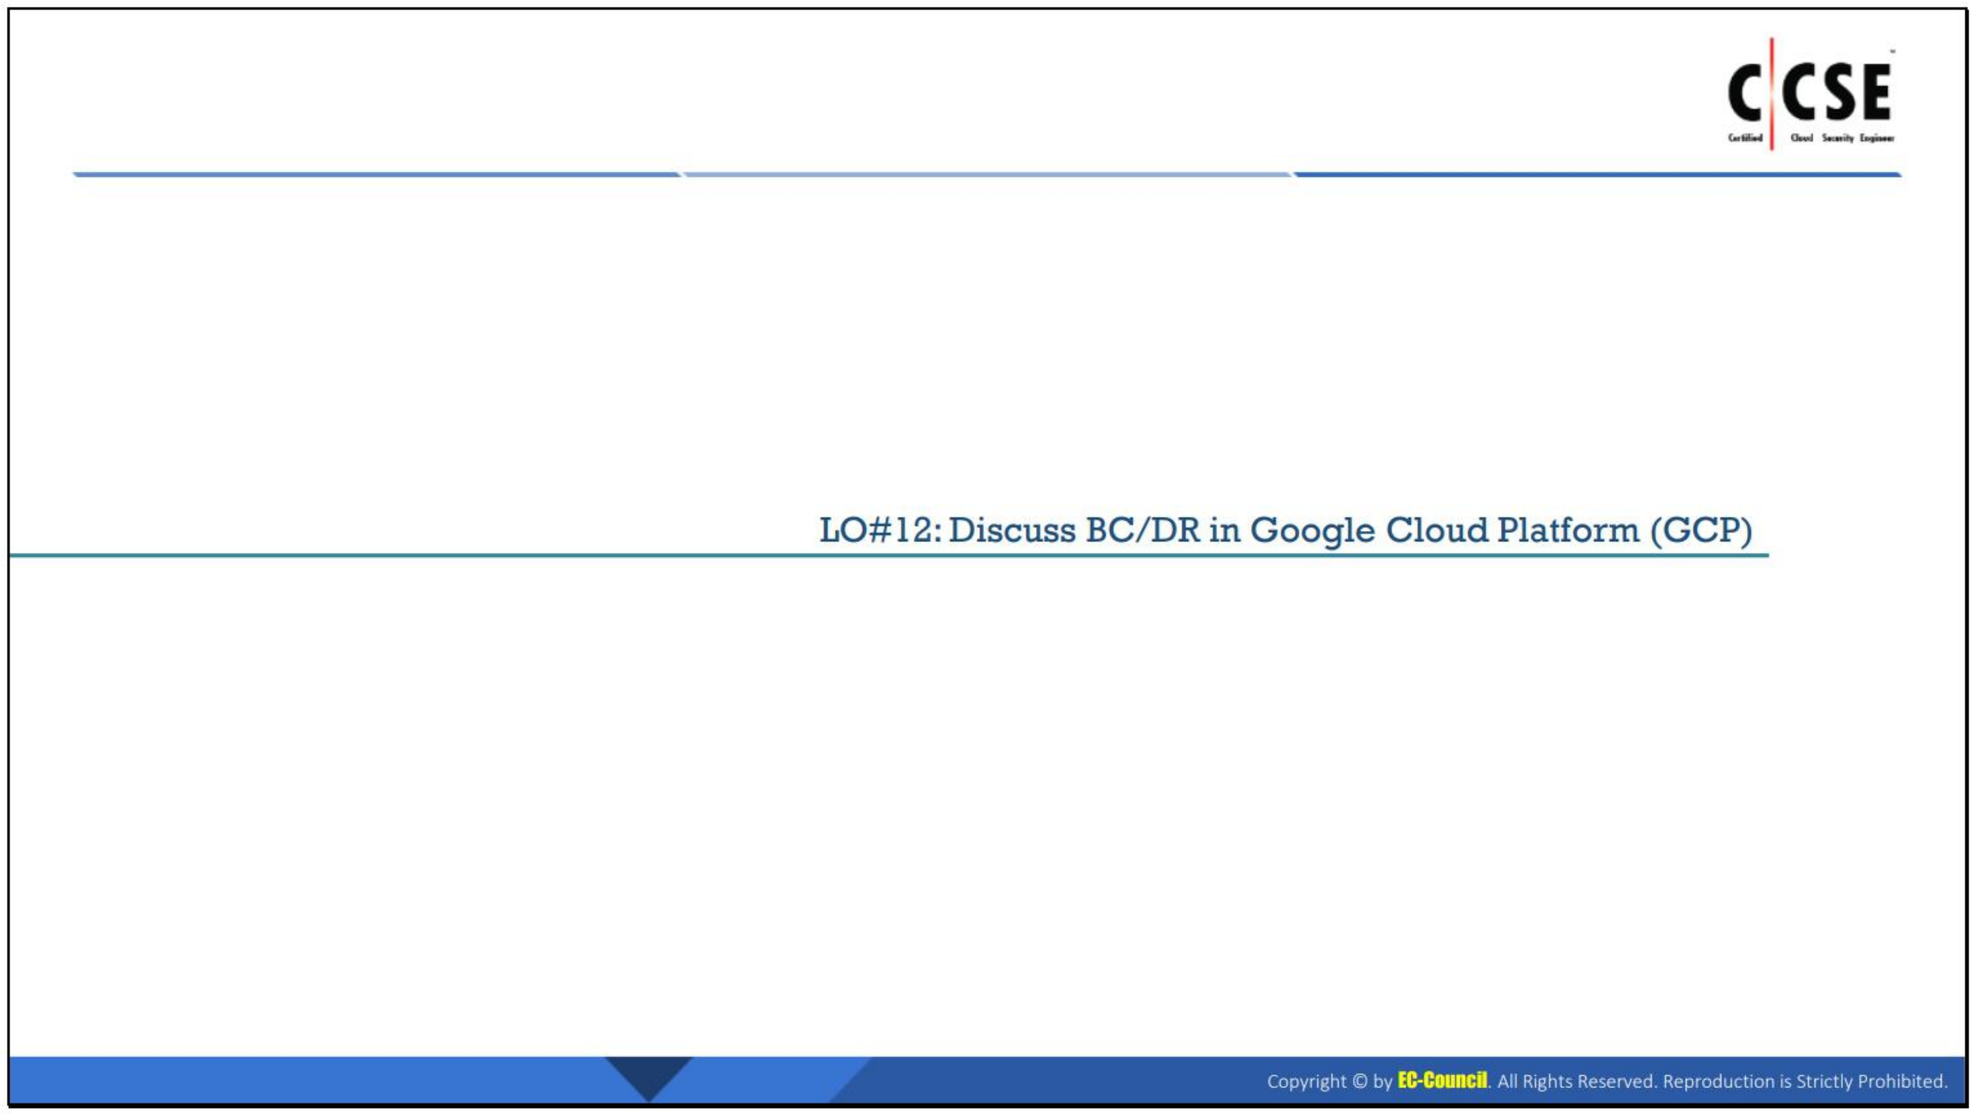

LO#12: Understand Business Continuity and Disaster Recovery in Google Cloud Platform

The objective of this section is to understand BCDR in Google Cloud Platform (GCP). The section illustrates the tools and features that GCP provides for BCDR. By understanding these tools and features, cloud security engineers can decide how to use them effectively.

Certified Cloud Security Engineer Copyright © by **EC-Council**
All Rights Reserved. Reproduction is Strictly Prohibited.

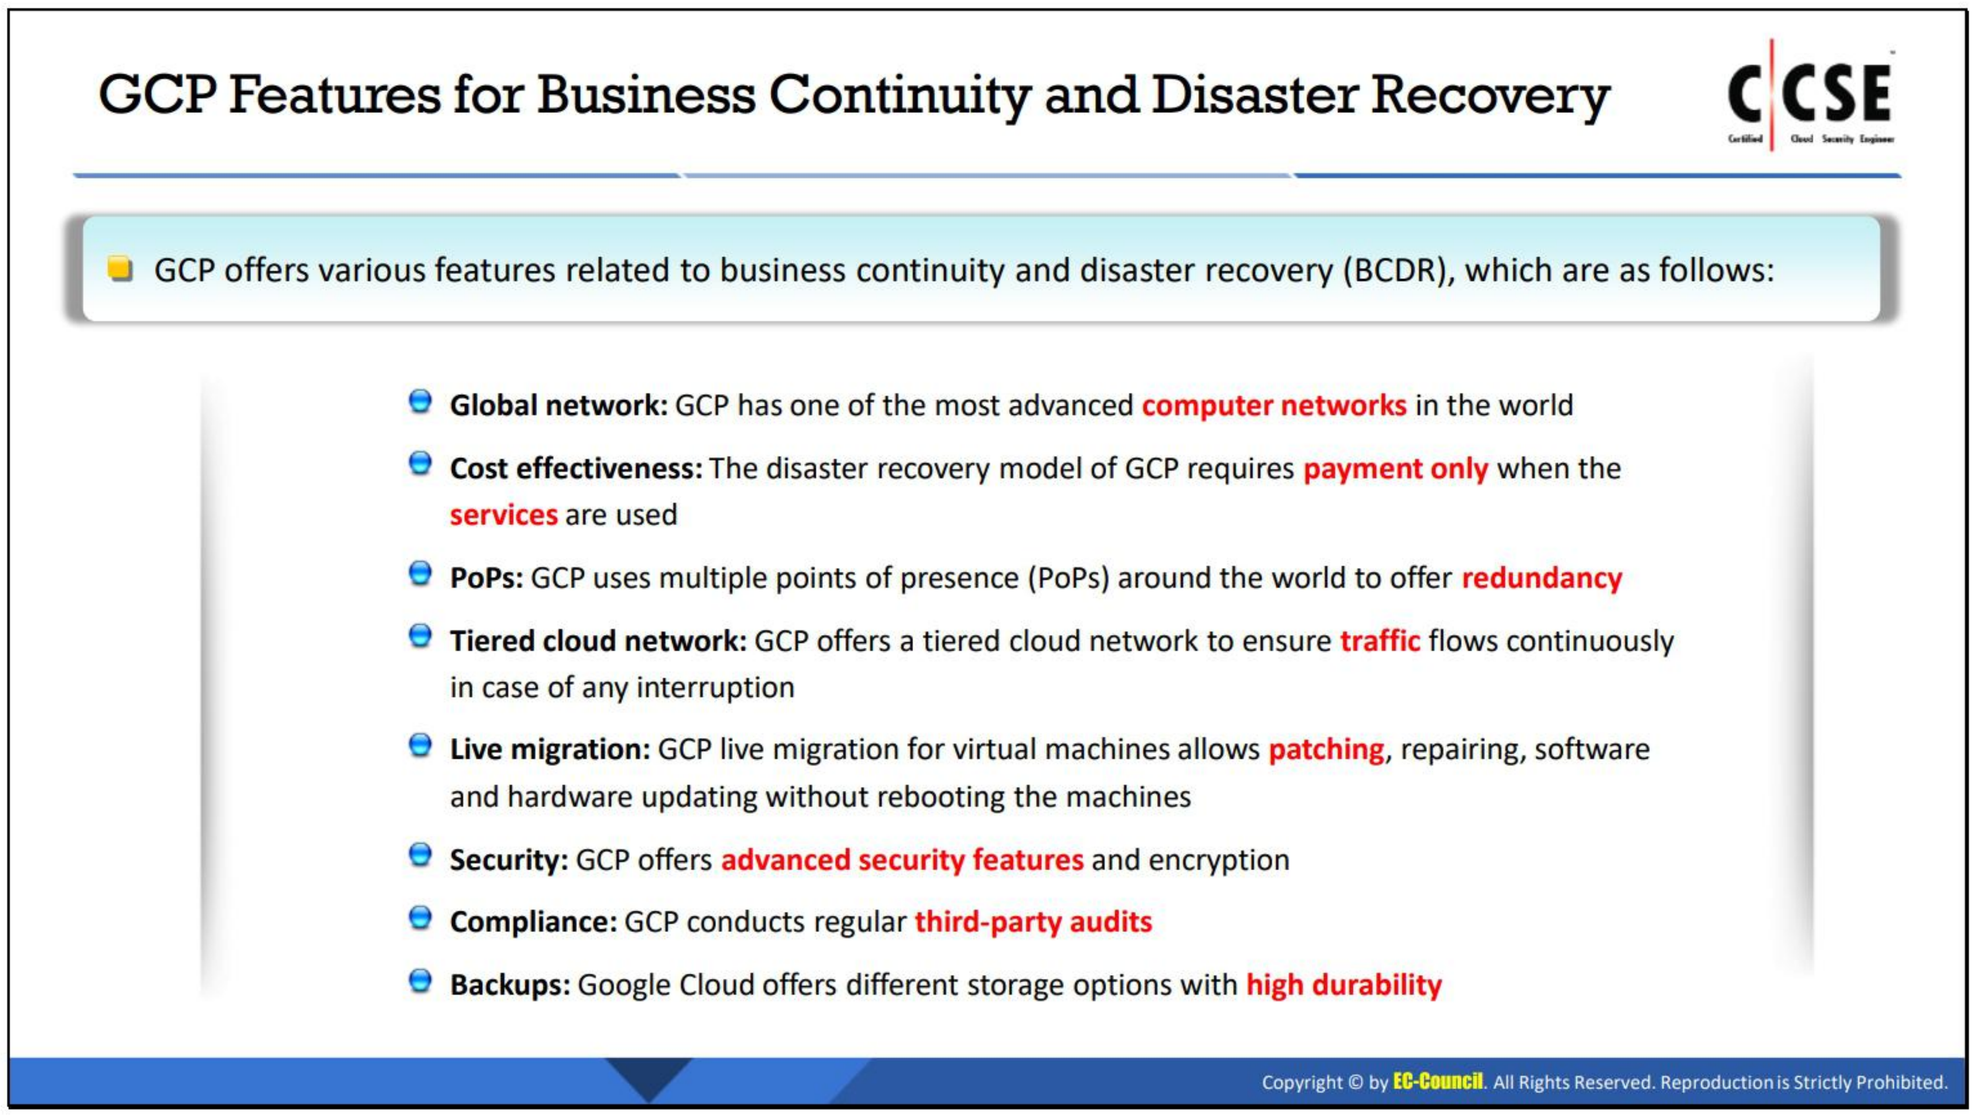

GCP Features for Business Continuity and Disaster Recovery

Many organizations have been adopting Google Cloud due to its cost-effectiveness, flexibility, and security features. GCP also dominates the traditional methods of BCDR by handling complicated planning requirements.

hide01.ir

Google Cloud eliminates most complicated factors by using a highly managed solution. GCP offers various features related to DR plan such as:

- **Global network:** GCP has one of the most advanced computer networks in the world. It uses cutting-edge software-defined networking and edge-caching services for its backbone network to offer fast and scalable performance.
- **Cost-effectiveness:** The disaster recovery model of GCP requires payment only when the services are used.
- **PoPs:** GCP uses multiple points of presence (PoPs) around the world to offer redundancy. The copies of the data are automatically stored in multiple storage devices in various locations.
- **Tiered cloud network:** GCP offers a tiered cloud network to ensure traffic flows continuously in case of any interruption.
- **Live migration:** GCP hosting offers live migration for VMs. This ensures that the patching, repairing, and software and hardware updates are performed without rebooting the machines.
- **Security:** GCP offers advanced security features and encryption. This ensures that the customer data is safe in Google Cloud. The site reliability engineering team also ensures to provide high availability for the customers.

Certified Cloud Security Engineer Copyright © by **EC-Council**
All Rights Reserved. Reproduction is Strictly Prohibited.

- **Compliance:** The Google Cloud conducts regular third-party audits to ensure that Google Cloud meets the required security, compliance regulations, privacy, and best practices.
- **Backups:** The Google Cloud offers different storage options such as Nearline, Coldline, Regional, and Multi-regional storage with high durability in case of a disaster.

hide01.ir

Certified Cloud Security Engineer Copyright © by **EC-Council**
All Rights Reserved. Reproduction is Strictly Prohibited.

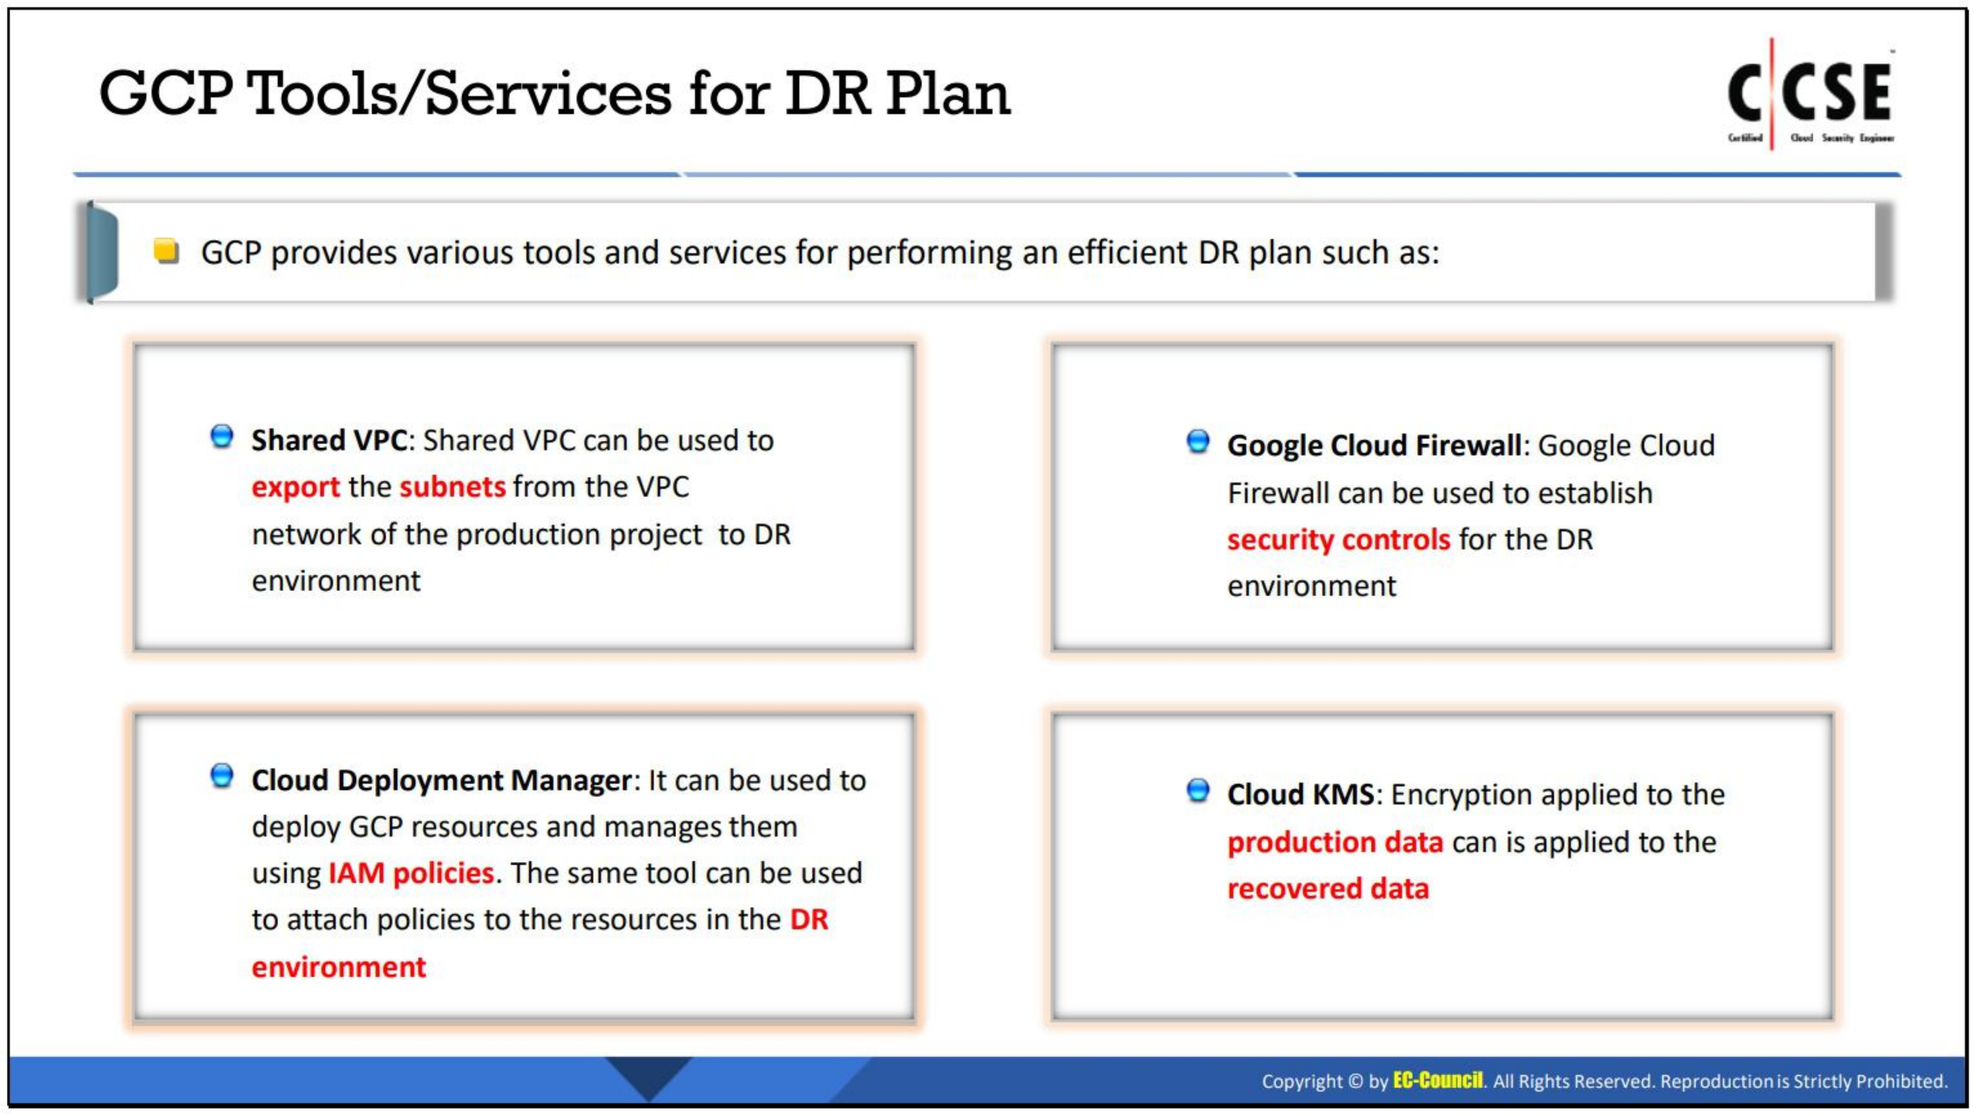

GCP Tools/Services for DR Plan

GCP provides various tools and services for performing an efficient DR plan.

Shared VPC

ride01.ir

Security is an important aspect while designing a DR plan. The controls that are applied to the production environment must be applied to the recovered data as well. Cloud customers must also ensure that the network controls applied to the recovered data provide the same separation and blocking as the production environment.

Shared VPC can be used to export the subnets from the VPC network of the production project to other service projects such as DR in the same organization. It can also help establish centralized security and network controls to the deployed environment. With shared VPC, the instances in the service project or DR can have a network connection with the shared subnet of the production environment.

It also provides an additional feature for preventing the host project from accidental deletion called “lien.” When the host project is configured as a shared VPC host project, the lien is placed on it and this prevents the project from accidental deletion. When the host project is no longer a shared VPC, the lien automatically gets removed.

Google Cloud Firewalls

Google Cloud Firewall can be used to establish security controls along with the inbound and outbound traffic for the disaster recovery environment. It allows hierarchical firewall policies where users can create firewalls at different levels such as organization and folder level.

GC Firewall also supports VPC Firewall rules where the traffic to and from the VM instances can be controlled. Though firewall rules are applied at a network level, connections are established

Certified Cloud Security Engineer Copyright © by **EC-Council**
All Rights Reserved. Reproduction is Strictly Prohibited.

for a per-instance basis. Using VPC rules logging, the cloud customers can audit, verify, and analyze the effects of firewall rules on a cloud environment.

Cloud Deployment Manager

The cloud administrator must ensure that the users are provided with the same access to the DR environment that they provided in the production environment. If the source production environment is Google Cloud, applying IAM policies to the DR environment is easy.

Cloud Deployment Manager can be used to deploy GCP resources and manage them using IAM policies. Users need to define the resources in the configuration file so that the process repeats whenever required. The same tool can be used to attach policies to the resources in the DR environment.

Cloud Deployment Manager follows a declarative approach where the users only need to specify the configuration and the system performs the required steps.

Cloud KMS

Cloud customers must also ensure that the encryption applied to the production data is applied to the recovered data. The keys should be secretly managed so that only authorized users can access DR resources.

Cloud Key Manager can be used to manage symmetric and asymmetric keys of the GCP cloud resources. Cloud customers can create, rotate, and destroy the cryptographic keys. They can use either software-backed encryption keys or customer-manager encryption keys (CMEK) to encrypt the data in GCP resources.

hide01.ir

Certified Cloud Security Engineer Copyright © by **EC-Council**
All Rights Reserved. Reproduction is Strictly Prohibited.

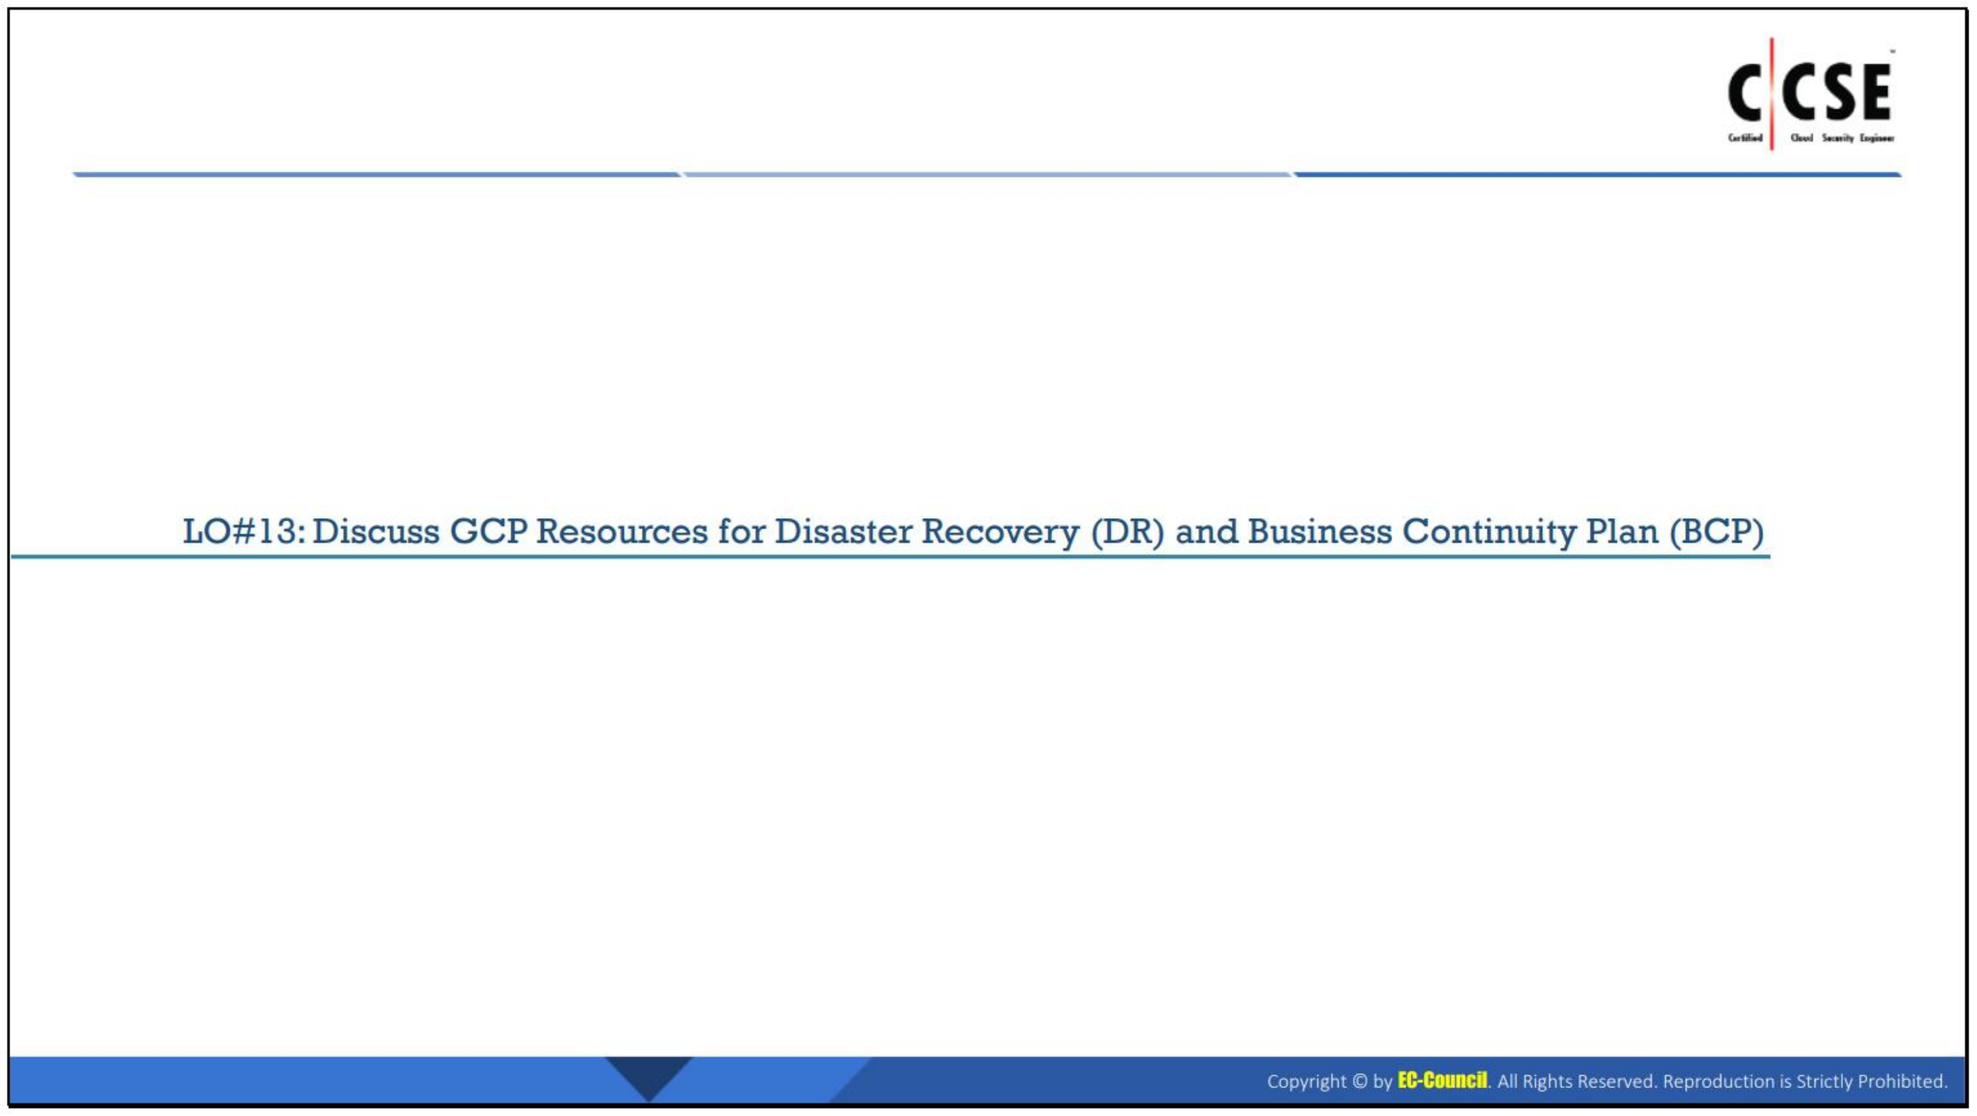

LO#13: Discuss GCP Resources for Disaster Recovery (DR) and Business Continuity Plan (BCP)

The objective of this section is to understand how GCP helps with backing up and recovering data and databases following a disaster event. By the end of this section, Cloud Security.

hidden1.ir

Certified Cloud Security Engineer Copyright © by **EC-Council**
All Rights Reserved. Reproduction is Strictly Prohibited.

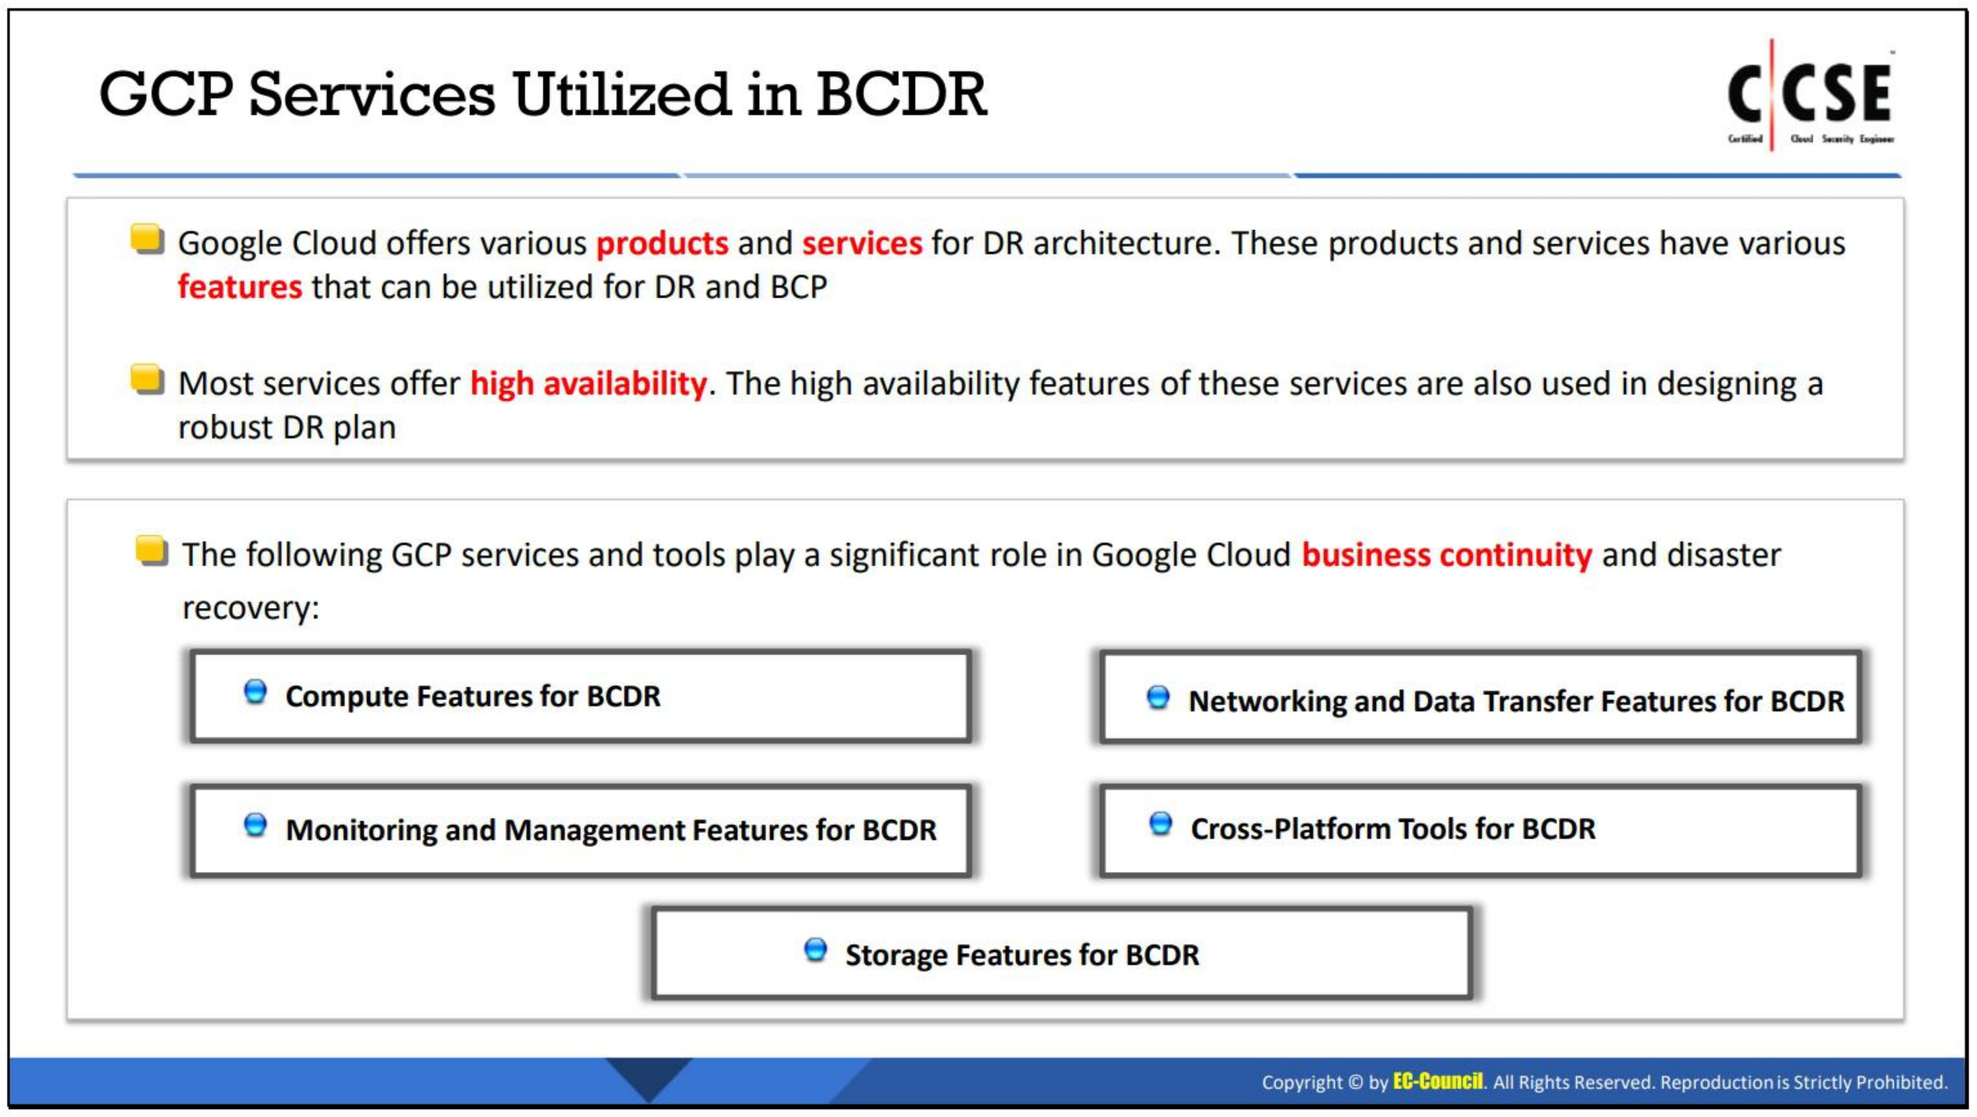

GCP Services Utilized in BCDR

Google Cloud offers various products and services for DR architecture. These products and services have numerous features that can be utilized for DR and BCP. Most services offer high availability. The high-availability features of these services are also used in designing a robust DR plan. For instance, by using the benefits of high-availability features, the consumer can design an architecture that can improve uptime and alleviate small-scale failures such as single VM failure.

The following GCP services play a significant role in Google Cloud BCDR:

- Compute features for BCDR
- Networking and data transfer features for BCDR
- Monitoring and management features for BCDR
- Cross-platform tools for BCDR
- Storage features for BCDR

Certified Cloud Security Engineer Copyright © by **EC-Council**
All Rights Reserved. Reproduction is Strictly Prohibited.

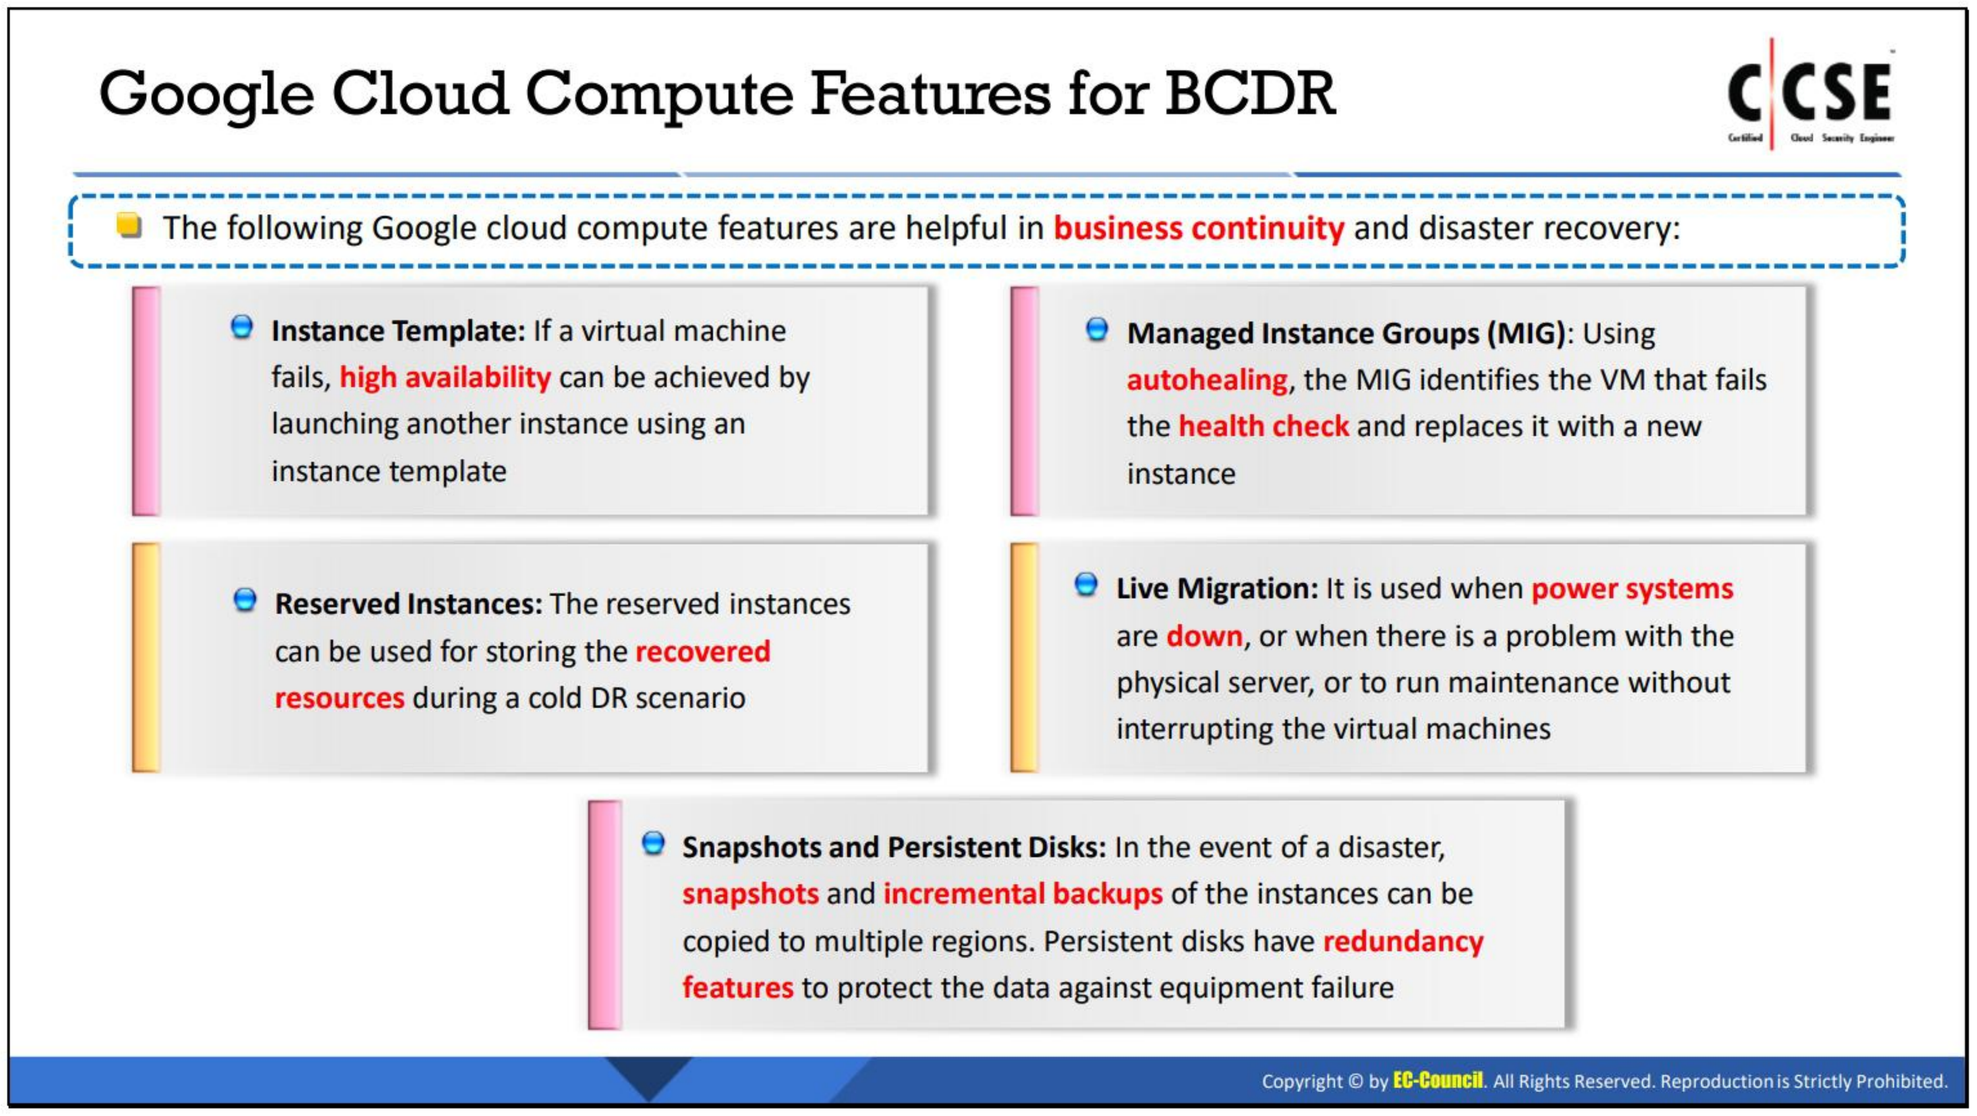

Google Cloud Compute Engine Features for BCDR

Compute Engine offers VM instances. To implement a DR plan, several features other than launching, configuring, and monitoring instances need to be used. For DR, "deletion protection" can be applied to instances to prevent accidental deletion. This feature is very useful while hosting databases. HA can be achieved in GCP using various features. The following features are helpful in BCDR:

- **Instance template:** An instance template can be used to save the configuration of an instance. Using the template, any number of instances can be created with the same configurations such as machine type, boot disk image, etc. This feature is also helpful for launching multiple instances with the same specifications to achieve HA. If a VM fails, high availability can be achieved by launching another instance using an instance template.
- **Managed instance groups (MIGs)**: A managed instance group can be used to create a group of VMs that run the same services in the same configuration. MIGs that work with load balancer distributes the traffic to a group of instances that are available in multiple zones can also be used in case of disaster recovery. The traffic can be sent to other instances in case of a failure at an AZ. It also has a feature to improve availability called autohealing. Using autohealing, the MIG identifies the VM that fails the health check and replaces it with a new instance.
- **Reserved instances:** The VM instances can be reserved in a zone with predefined machine types. The Reserved Instances can be used for storing the recovered resources during a cold DR scenario. The recovered resources can be used in a failover condition without the need for configuring and deploying them in advance.

Certified Cloud Security Engineer Copyright © by **EC-Council**
All Rights Reserved. Reproduction is Strictly Prohibited.

- **Snapshots and persistent disks:** In the event of a disaster, snapshots and incremental backups of the instances can be copied to multiple regions and can be used to create persistent disks. Persistent disks have redundancy features to protect the data against equipment failure. Snapshots of the persistent disks can also be created to protect the data due to a user error.
- **Live migration:** This feature enables the instances to be migrated to other hosts in the same zone instead of rebooting the instances. Live migration is used when power systems are down, or when there is a problem with the physical server, or to run maintenance without interrupting the VMs.

hide01.ir

Certified Cloud Security Engineer Copyright © by **EC-Council**
All Rights Reserved. Reproduction is Strictly Prohibited.

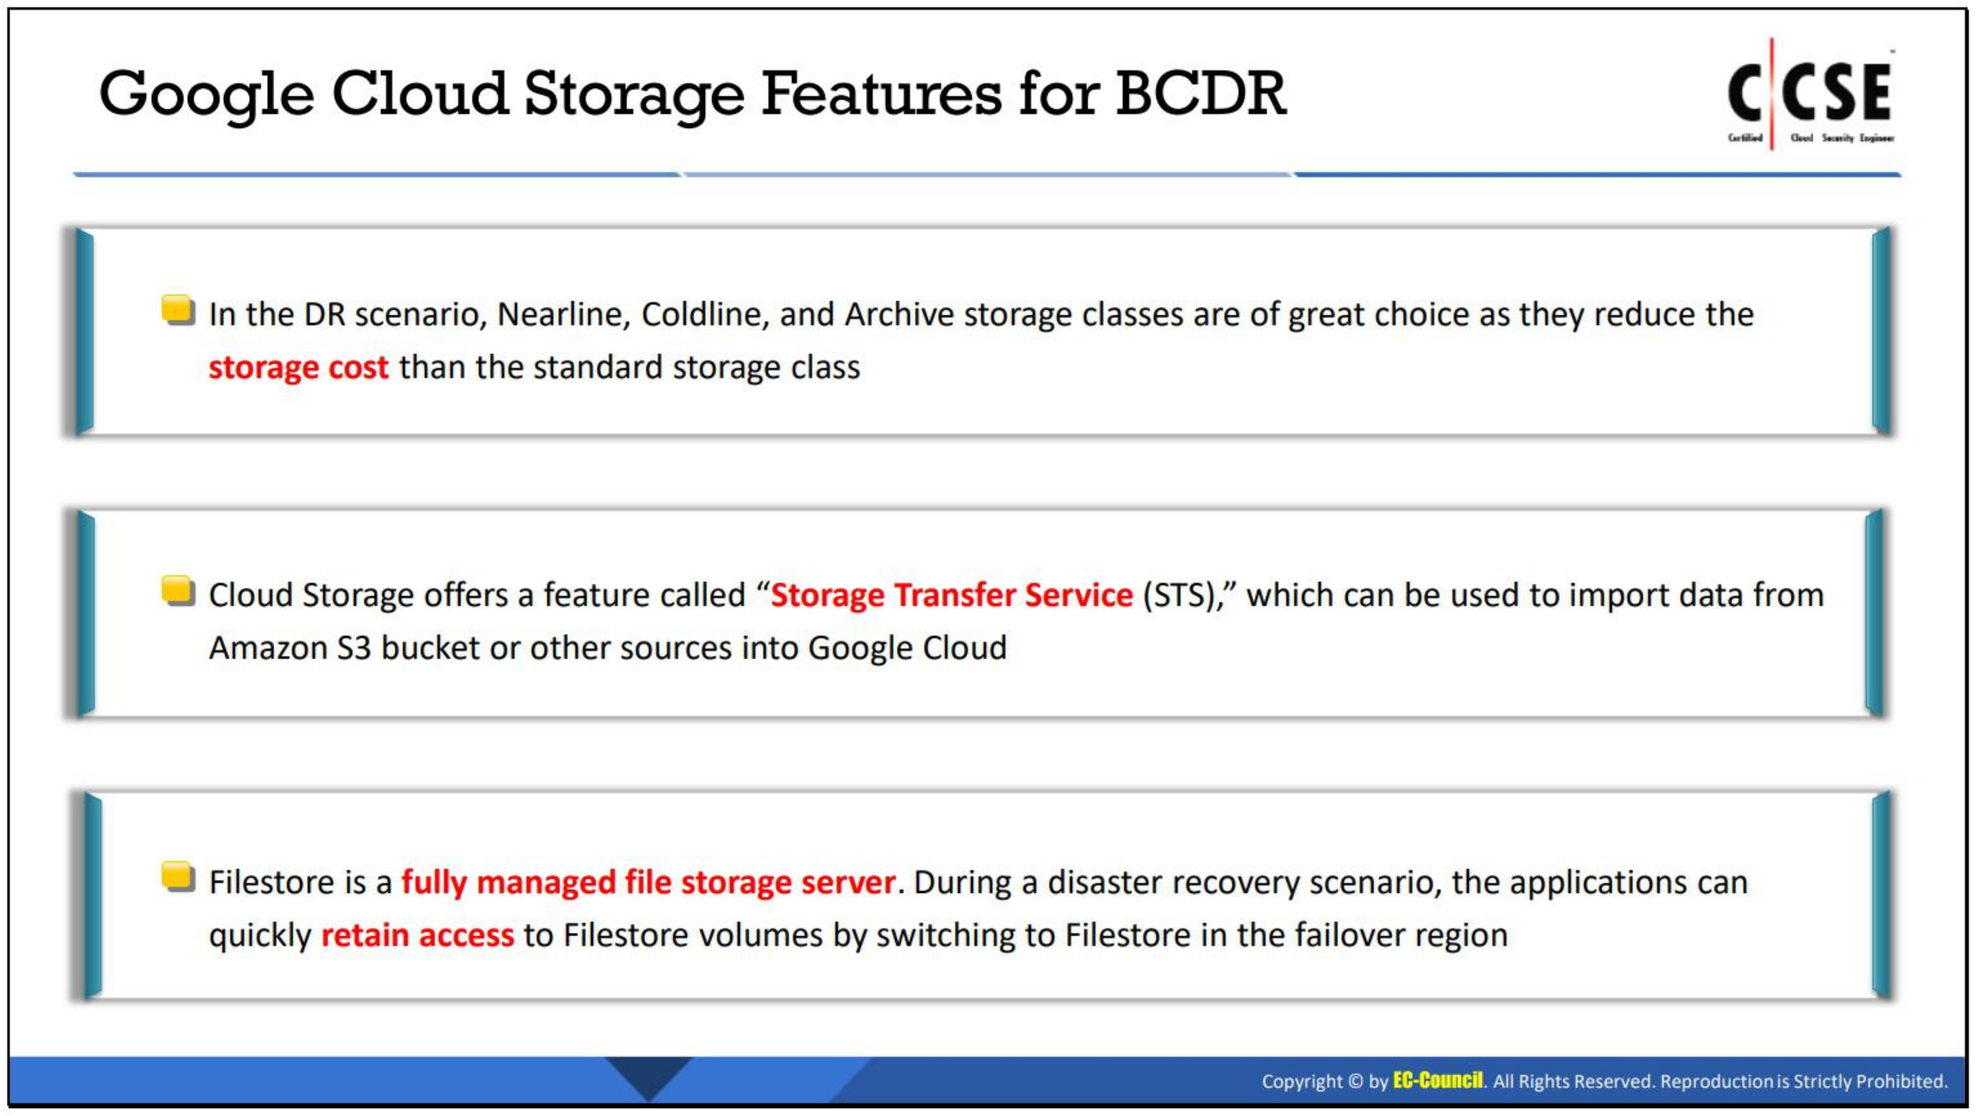

Google Cloud Storage Features for BCDR

Cloud Storage is an object storage service that is suitable for storing backup data. Google Cloud offers different storage classes such as Standard, Nearline, Coldline, and Archive to meet specific use cases.

In the DR scenario, Nearline, Coldline, and Archive storage classes are of great choice as they reduce the storage cost compared to the Standard storage class. Nearline, Coldline, and Archive storage classes are designed for infrequently accessed data.

Cloud Storage offers a feature called storage transfer service (STS), which can be used to import data from Amazon S3 bucket or other sources into Google Cloud. In DR scenarios, STS can be used to backup data from other storage providers to the cloud storage bucket. It can also be used to move data from a bucket in a multi-region to a bucket in a single region to reduce storage costs.

Filestore is a fully managed file storage server. It is fully integrated with Compute Engine and Google Kubernetes Engine (GKE) so that the instances and containers can refer to the same shared data. Filestore instances are zonal, and the instance becomes unavailable if the zone is down. To mitigate this, cloud customers should regularly backup their data by synching Filestore volume to a Filestore instance in another region. A job must be scheduled to run on GKE clusters or Compute Engine instances to backup the data.

During a disaster recovery scenario, the applications can quickly retain access to Filestore volumes by switching to Filestore in the failover region.

Certified Cloud Security Engineer Copyright © by **EC-Council**
All Rights Reserved. Reproduction is Strictly Prohibited.

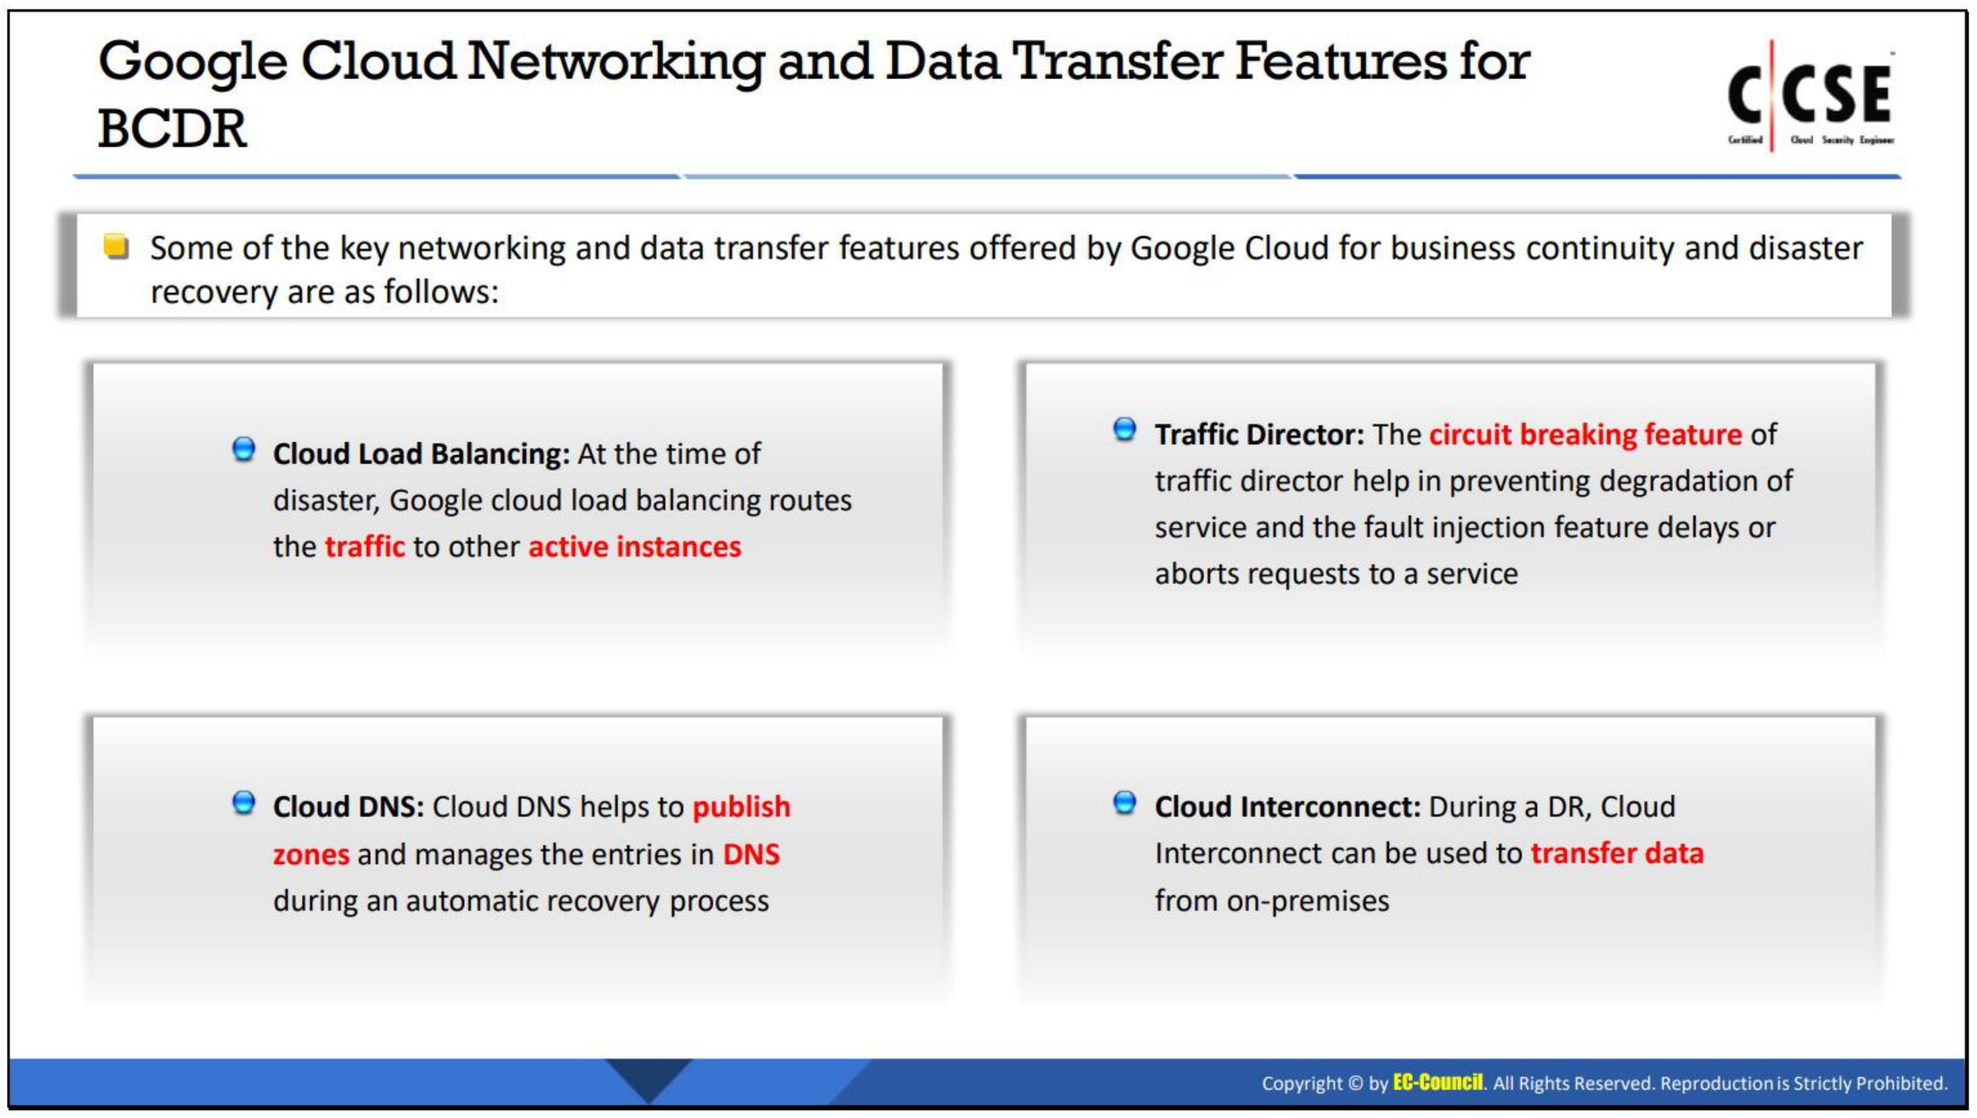

Google Cloud Networking and Data Transfer Features for BCDR

For a DR plan, it is important to consider networking and data transfer factors as well. GCP provides the below services to handle networking and data transfer for a DR scenario.

hide01.ir

Cloud Load Balancing

Google Cloud Load Balancing feature provides high availability for Compute Engine instances by distributing the traffic to a set of multiple instances. This feature is helpful in case of a disaster as the traffic can be routed to other active instances. Load balancing has a health check feature, where it checks if an instance is working well. If the instance is not working, the load balancing will not direct the traffic to that instance and routes the traffic to other healthy instances.

Cloud Load Balancing provides a global IP address that can be routed to Compute Engine instances. With load balancing, the user requests will be routed to the instances in their nearest location. This reduces the response time and improves efficiency.

Traffic Director

Traffic Director is a fully managed network platform and service mesh. Using this, cloud customers can create a fully managed traffic control to their service mesh. Traffic Director is responsible for managing the configuration of service proxies on Compute Engine and GKE.

It can offload health checks from service proxies and cloud customers can configure their traffic control policies. It has additional features such as circuit breaking and fault injection that are helpful for a DR scenario. Using the circuit breaking feature, cloud customers can limit the requests to a service, and further requests will be prevented from reaching the service. This helps prevent further degradation of service.

Certified Cloud Security Engineer Copyright © by **EC-Council**
All Rights Reserved. Reproduction is Strictly Prohibited.

Using fault detection feature, Traffic Director delays or aborts requests to a service so that cloud customers can test their service for handling request delays or aborted requests.

Cloud DNS

Cloud DNS offers both public zone and privately managed DNS zones. The public zone is visible to everyone, whereas a private zone is visible only to the specified VPN. Cloud DNS helps publish zones and manages the entries in DNS during an automatic recovery process.

It uses a global network of anycast name servers to handle the cloud customer's DNS zones from locations around the world that are redundant. This ensures high availability and low latency for the service users.

Cloud Interconnect

Cloud Interconnect can be used for high availability and low-latency connection to transfer data from on-premises to Google Cloud VPC networks. During a DR, Cloud Interconnect can be used to transfer data from on-premises.

hide01.ir

Certified Cloud Security Engineer Copyright © by **EC-Council**
All Rights Reserved. Reproduction is Strictly Prohibited.

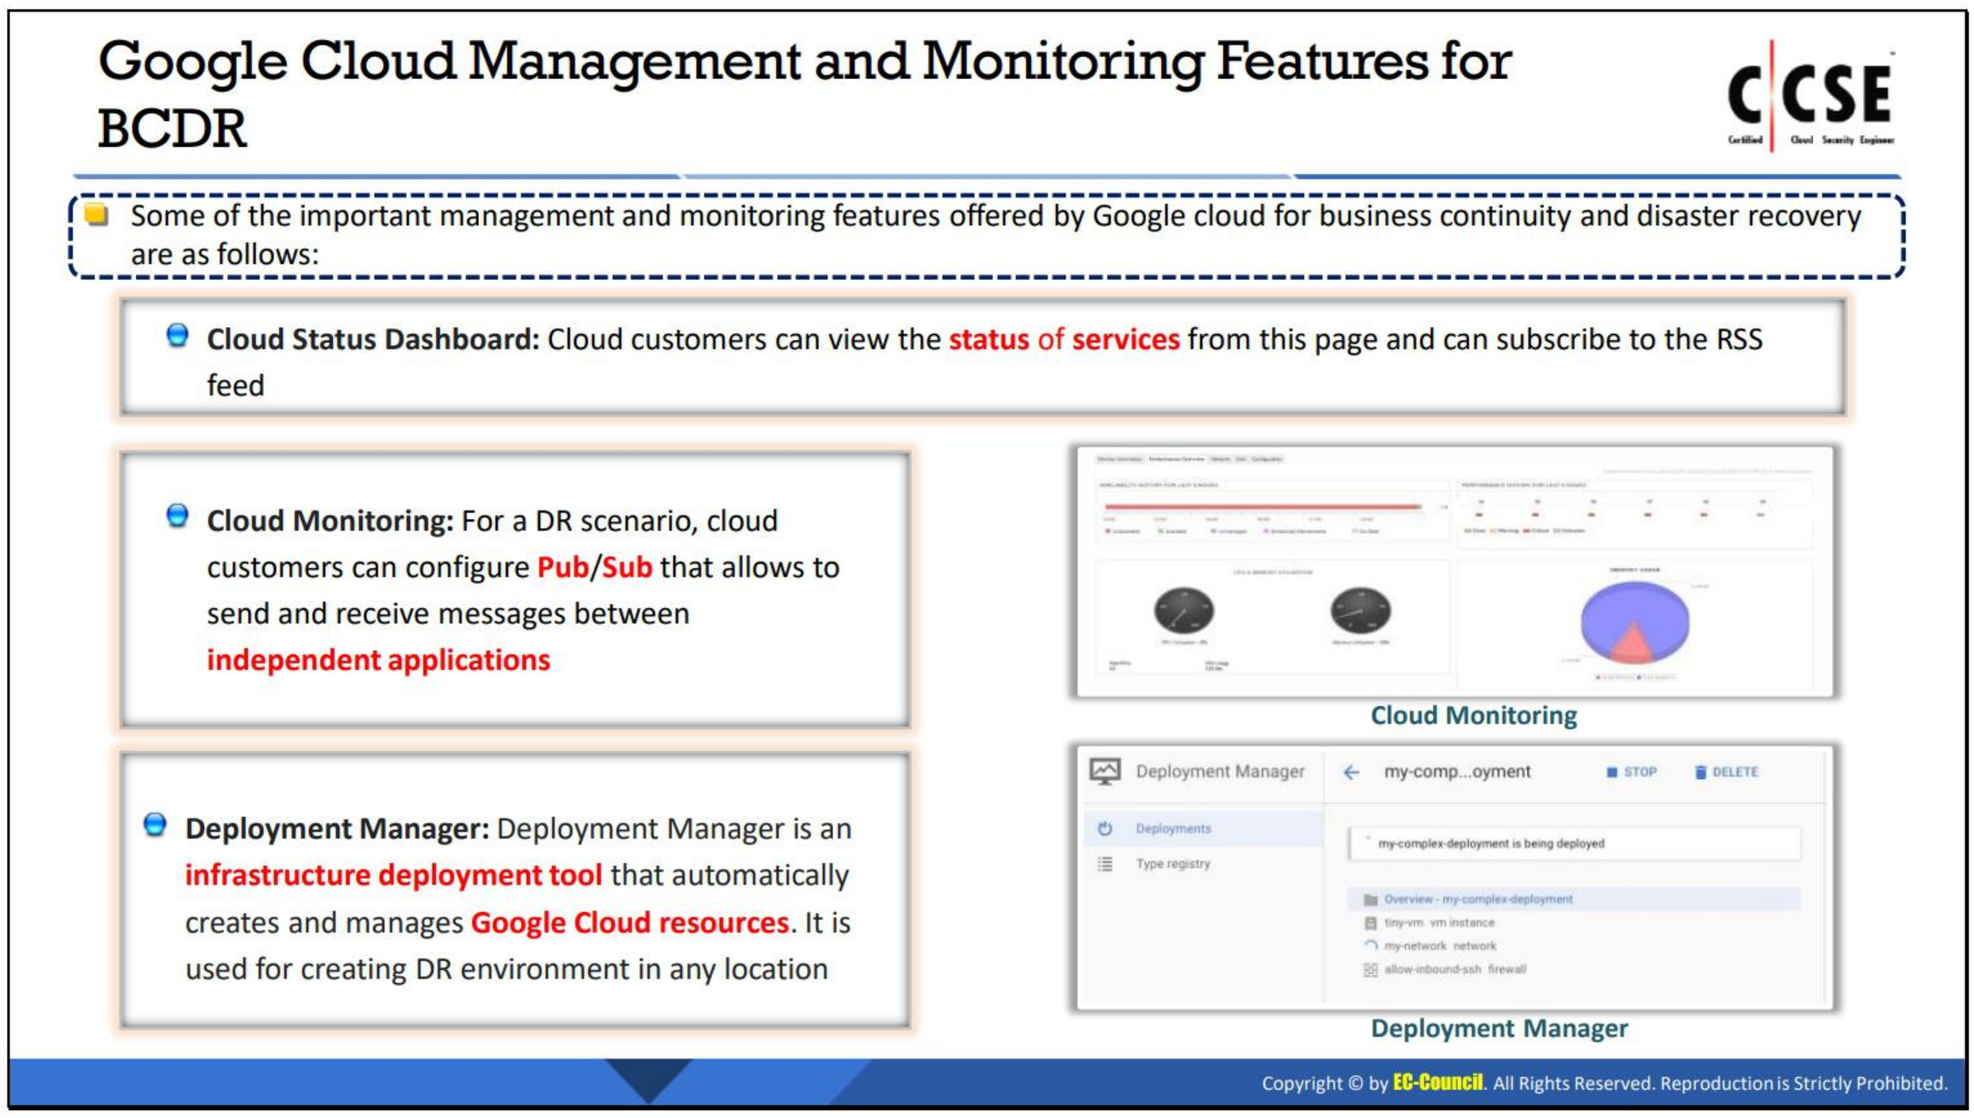

Google Cloud Management and Monitoring Features for BCDR

There are various services provided by GCP for monitoring and managing BCDR. Some of the important services and their features are described below.

hide01.ir

Cloud Status Dashboard

Cloud Status Dashboard provides information about the availability of Google Cloud services. Customers can view the status of services from this page and can subscribe to the RSS feed. An RSS feed is a file that contains information about the new updates of services.

Cloud Monitoring

Google Cloud Monitoring can be used to view the entire stack of a project. It provides metrics, metadata, and events from Google Cloud, AWS, application instrumentation, and hosted uptime probes.

Cloud customers can configure alerts to send notifications to other tools such as slack and other incident management tools to provide regular updates to cloud administrators. For a DR scenario, cloud customers can configure Pub/Sub that allows them to send and receive messages between independent applications. Then, they can use Cloud Functions to initiate an automated response to the Cloud Monitoring alert.

Deployment Manager

Deployment Manager is an infrastructure deployment tool that automatically creates and manages Google Cloud resources. Cloud customers need to write template and configuration files. Using these templates, cloud customers can create an environment with different Cloud services with just a single command. Customers can also break down the environment with a single command.

Certified Cloud Security Engineer Copyright © by **EC-Council**
All Rights Reserved. Reproduction is Strictly Prohibited.

This makes Google Deployment Manager suitable for a DR architecture and customers can create a DR environment in any location they want.

hide01.ir

Certified Cloud Security Engineer Copyright © by **EC-Council**
All Rights Reserved. Reproduction is Strictly Prohibited.

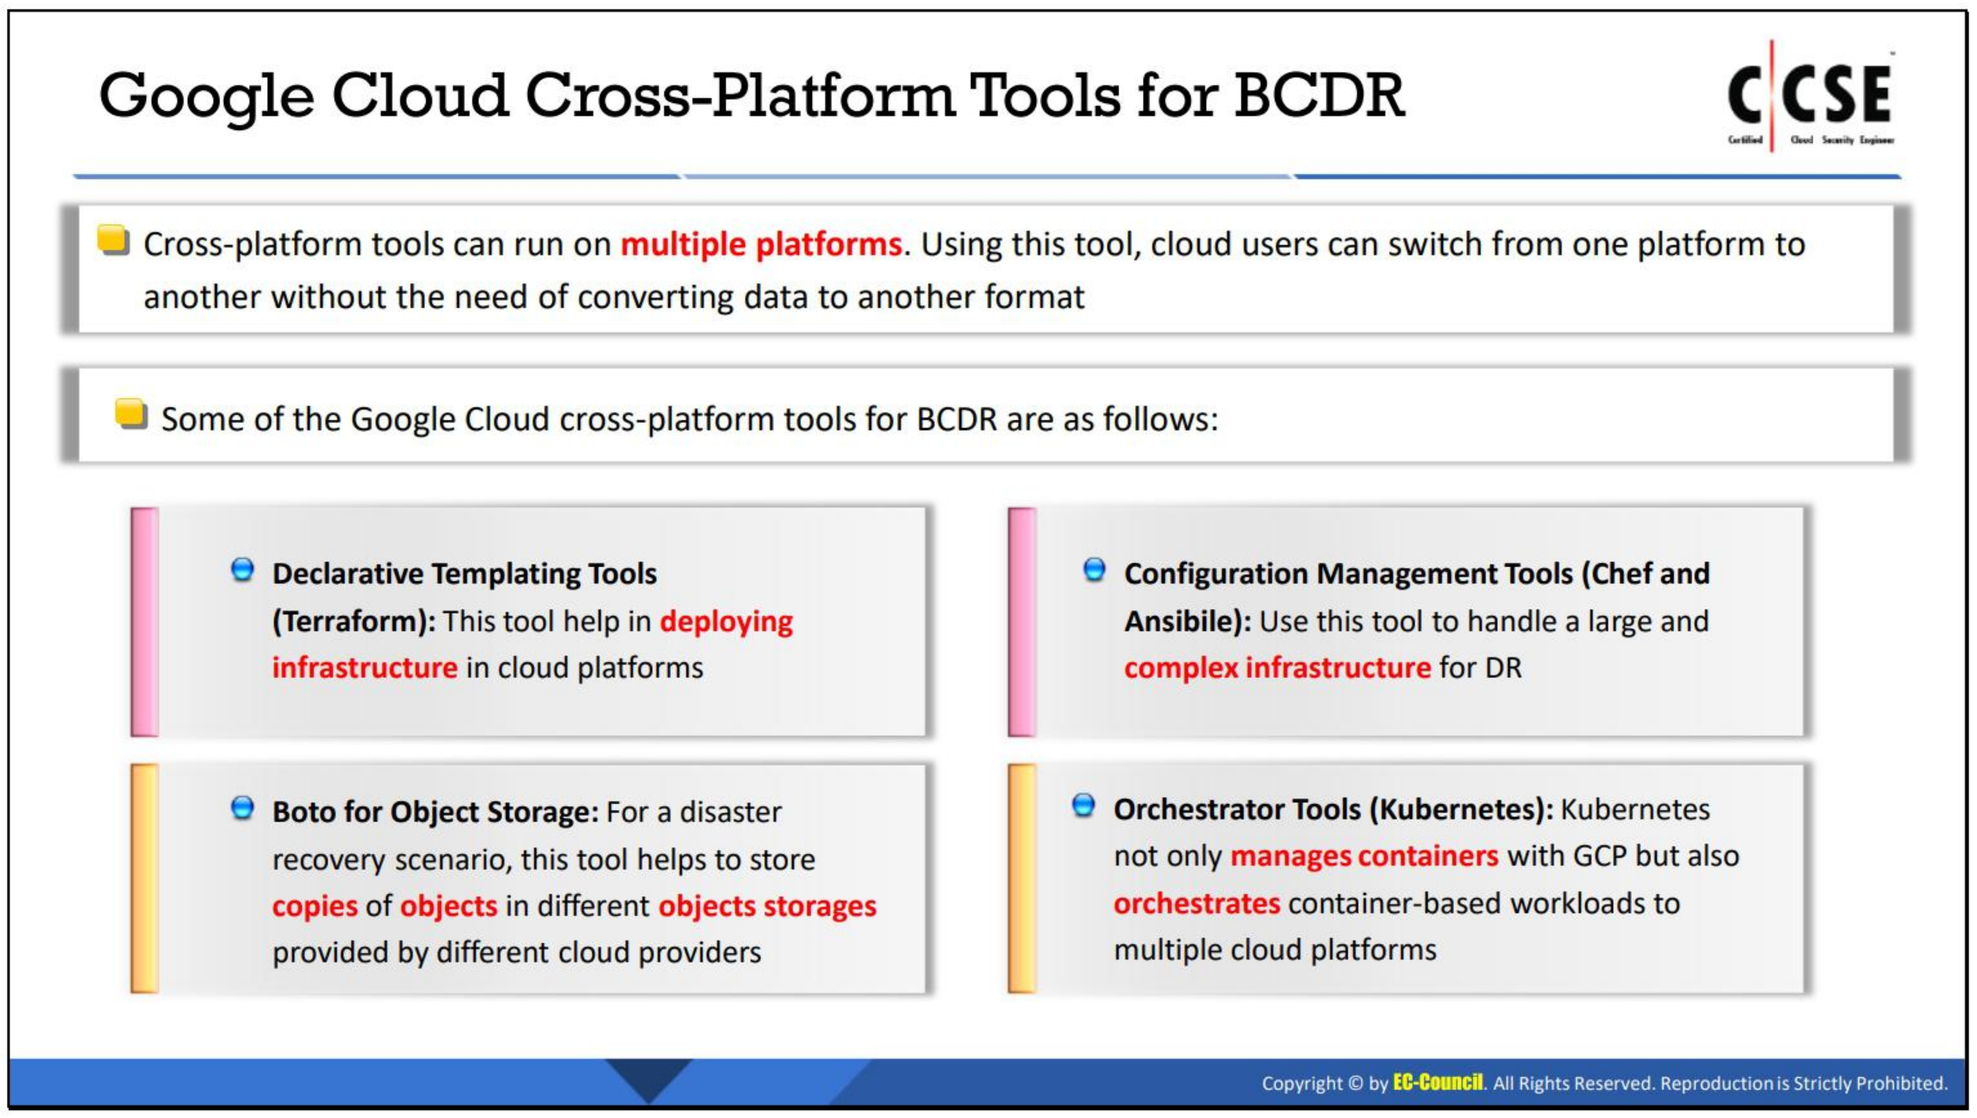

Google Cloud Cross-Platform Tools for BCDR

During a DR, when the workloads run on different platforms, it is important to choose a tool that works with all the platforms to reduce operational overhead. Cross-platform tools can run on multiple platforms. Cloud users can switch from one platform to another without the need of converting data to another format. Google Cloud offers various tools that support cross-platform during a DR scenario.

Declarative Templating Tools (Terraform)

Declarative templating tools help deploy infrastructure in cloud platforms. Deployment Manager can be used for deployment only in GCP. Terraform is one of the most popular cross-platform deployment tools.

Terraform is an open-source IaC tool. Cloud users can manage hundreds of cloud services through its CLI workflow. They can write IaC using declarative configuration files. By provisioning infrastructure as code, Terraform reduces human error and increases automation.

Configuration Management Tools (Chef and Ansible)

To handle a large and complex infrastructure for DR, it is suggested to use platform-agnostic configuration management tools such as Chef and Ansible. These tools help apply reproducible configuration irrespective of the location of compute workload.

Boto for Object Storage

For a disaster recovery scenario, it is important to store copies of objects in different objects storages provided by different cloud providers. Boto is one of the cross-platform tools for object storage. Boto works with both Amazon S3 and Cloud Storage. It is a Python package that provides support for both public cloud and private cloud.

Certified Cloud Security Engineer Copyright © by **EC-Council**
All Rights Reserved. Reproduction is Strictly Prohibited.

Orchestrator Tools (Kubernetes)

Containers are also considered as building blocks for disaster recovery. A container helps package services and maintain consistency across cloud platforms. Containers typically perform orchestration.

Kubernetes not only manages containers with GCP but also orchestrates container-based workloads to multiple cloud platforms. Using DNS service, traffic to Kubernetes clusters on various cloud platforms can be distributed.

Cloud customers ensure that they can access the image registry and pull the image to the target environment during a disaster. For this, Container Registry, which is platform agnostic, can be used. Container Registry provides secure and private Docker image storage on GCP.

hide01.ir

Certified Cloud Security Engineer Copyright © by **EC-Council**
All Rights Reserved. Reproduction is Strictly Prohibited.

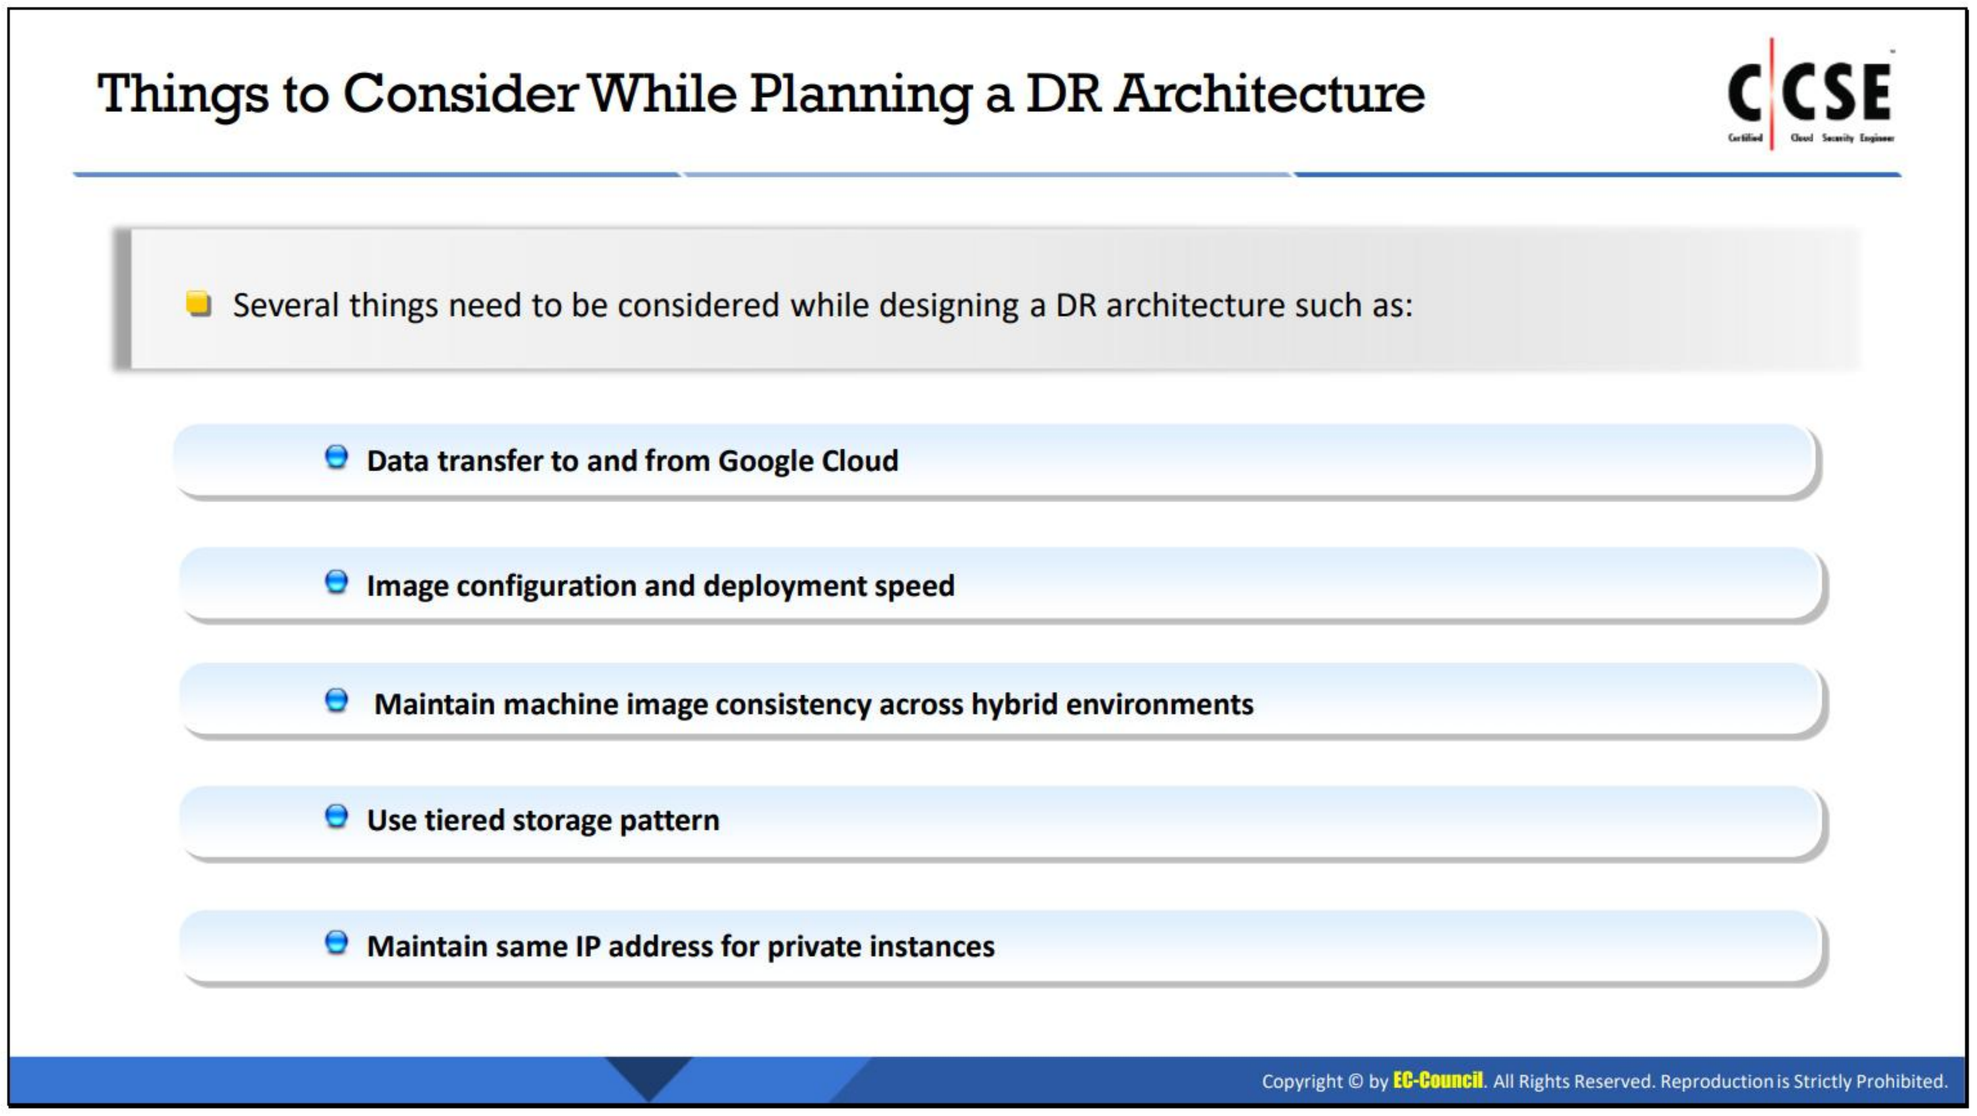

Things to Consider While Planning a DR Architecture

Several things need to be considered while designing a DR architecture. They are described below.

hide01.ir

Data Transfer to and from Google Cloud

Data transfer to and from Google Cloud is an important aspect for planning a DR. Data transfer is a complex task if the DR plan is based on moving data from on-premises or other cloud providers to GCP.

If the data transfer is performed using a public internet connection, the network throughput is unpredictable. Cloud Interconnect ensures a good throughput for data transfer. Cloud Interconnect offers various options to connect with Google Cloud such as:

- **Cloud VPN:** Cloud VPN helps create IPsec VPC tunnels between the target network and a GCP VPC network. VPC tunnel connects two VPN gateways and acts as a medium for encrypted traffic. With Cloud VPN, the traffic between the networks is encrypted by one gateway and is decrypted by another gateway. This helps protect the data passing through the internet. HA VPN can also be used to create high-availability VPN connection with 99.99% service availability.
- **Direct peering:** Direct peering can be used to create and exchange traffic between customer's networks and Google's edge network. It provides a direct connection from an on-premises network to Google public IP addresses.
- **Dedicated interconnect:** It provides direct physical connection between Google's network and on-premises network. The circuit speed can be 10 Gbps or 100 Gbps. It provides consistent throughput for large data transfers over large bandwidth, which is cost-effective.

Certified Cloud Security Engineer Copyright © by **EC-Council**
All Rights Reserved. Reproduction is Strictly Prohibited.

- **Partner interconnect:** Partner interconnect provides a connection between on-premises and GCP through a service provider. It offers the same capabilities as dedicated interconnect but the circuit speed is less than 10 Gbps.

Image Configuration and Deployment Speed

When a machine image is configured for deploying new instances, it is important to consider the effect of configuration on the speed of deployment. If a machine image is minimally configured, the time for launching the instance is high as it needs to be downloaded and installed individually. If the machine is highly configured, the instance will be quickly launched but the image needs to be updated frequently. The time for launching an instance has a direct relation with the RTO.

Maintain Machine Image Consistency across Hybrid Environments

While implementing a hybrid environment, cloud customers need to find a way to maintain VM consistency in the production environment. For this, a cloud security professional can make use of "Packer," which creates similar machine images for different platforms. The same scripts can be used with platform-specific configuration files. With Packer, customers can put the configuration file in version control to know the deployed version in production.

Cloud customers can also use configuration management tools such as Chef, Ansible, and Puppet to create instances with fine granularities, and fully or minimally configured images. They can also manually convert or make use of existing machine images such as VirtualBox images, Amazon Machine Image, etc.

hide01.ir

Use Tiered Storage Pattern

A tiered storage pattern is used for backups where recent data is stored in faster storage and the older backups are moved to cheaper storage classes. The implementation is based on the source of the data such as on-premises or Google Cloud. In both cases, the objects between storage buckets in different classes are migrated.

The cloud security professional can also change the storage classes of objects using object lifecycle rules.

Maintain Same IP Address for Private Instances

If a VM needs to be replaced due to any reason, the replaced VM should appear as the original VM. For this, the IP address through which the clients connect should remain the same.

The simplest way to do this is to use a managed instance group. The managed instance group can be configured with a load balancer to ensure that the same IP address is used to connect to the instance.

Certified Cloud Security Engineer Copyright © by **EC-Council**
All Rights Reserved. Reproduction is Strictly Prohibited.

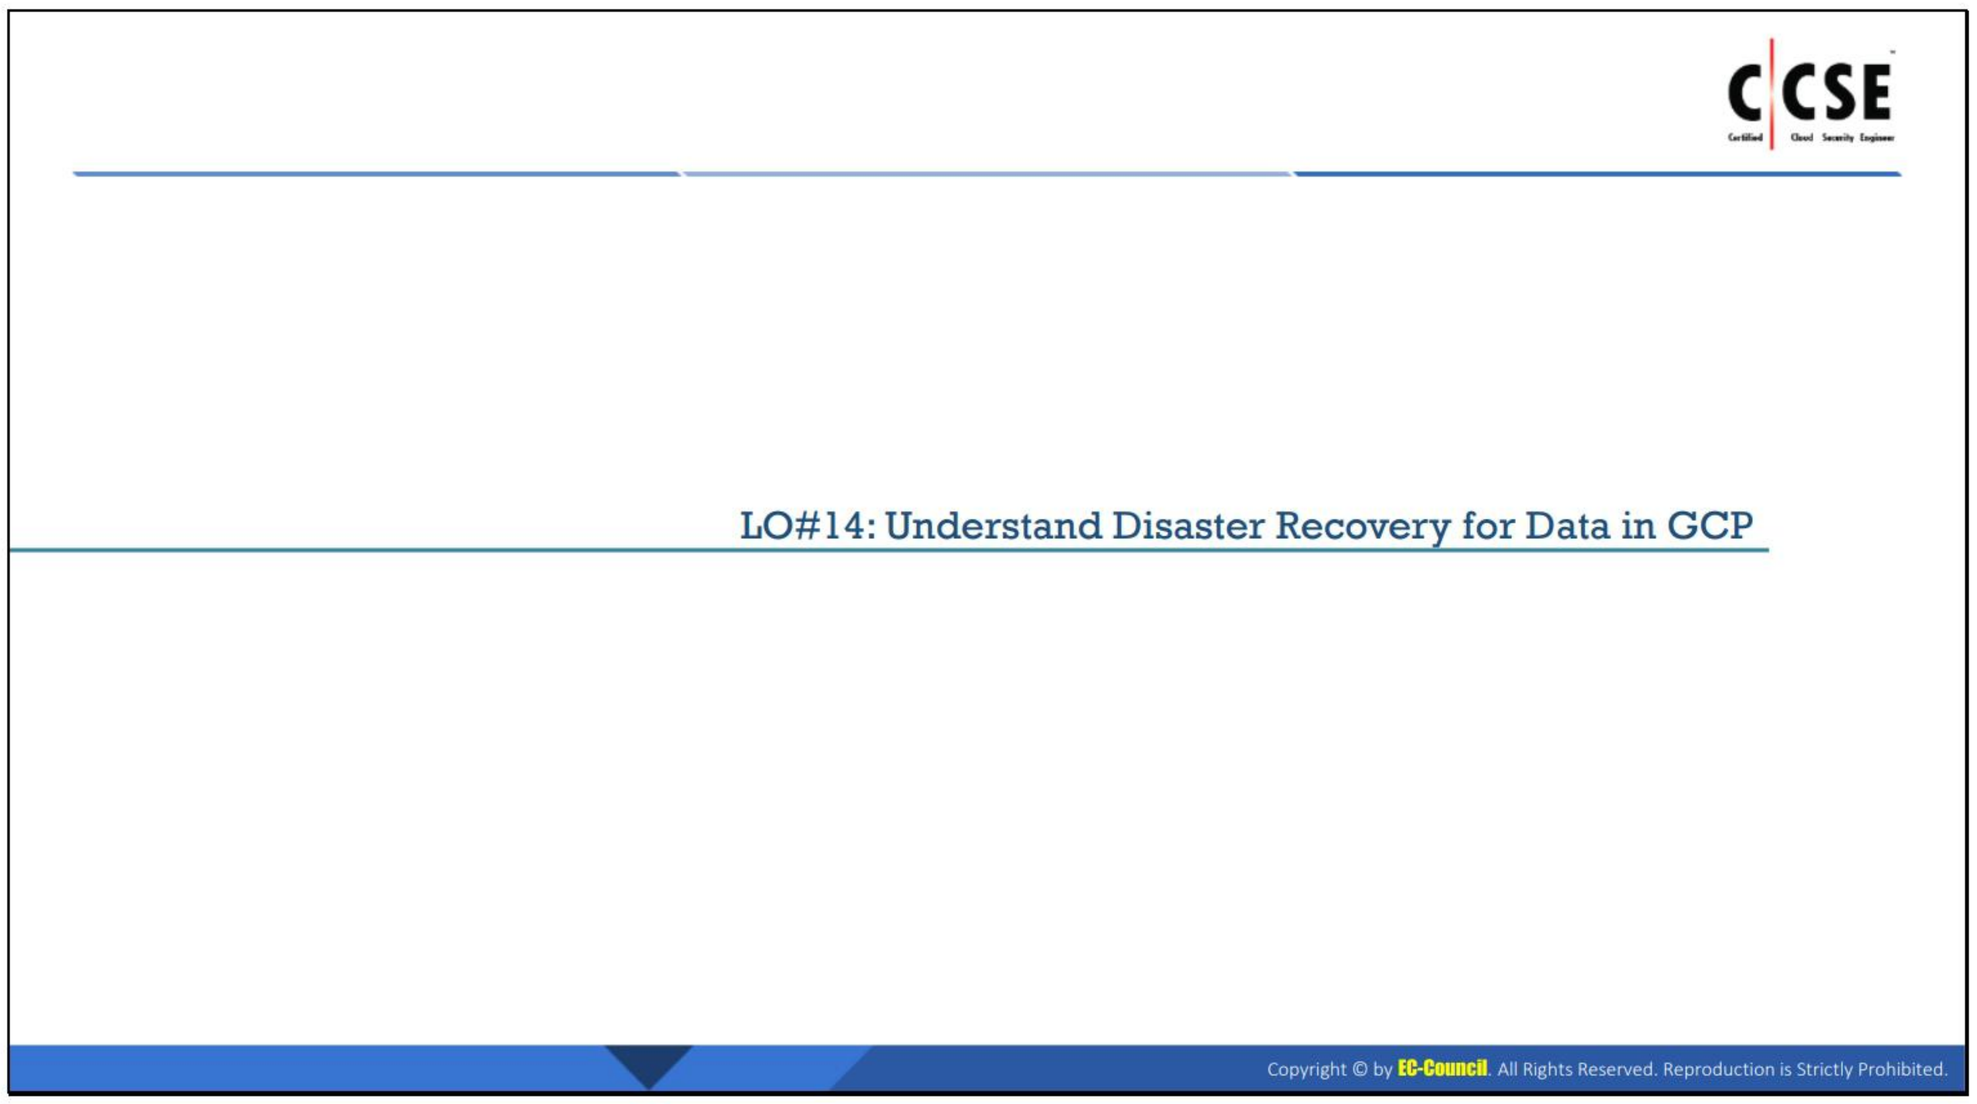

LO#14: Learn Disaster Recovery for Data in GCP

The objective of this section is to understand how GCP helps with backing up and recovering data and databases following a disaster event. By the end of this section, Cloud Security

hide01.ir

Certified Cloud Security Engineer Copyright © by **EC-Council**
All Rights Reserved. Reproduction is Strictly Prohibited.

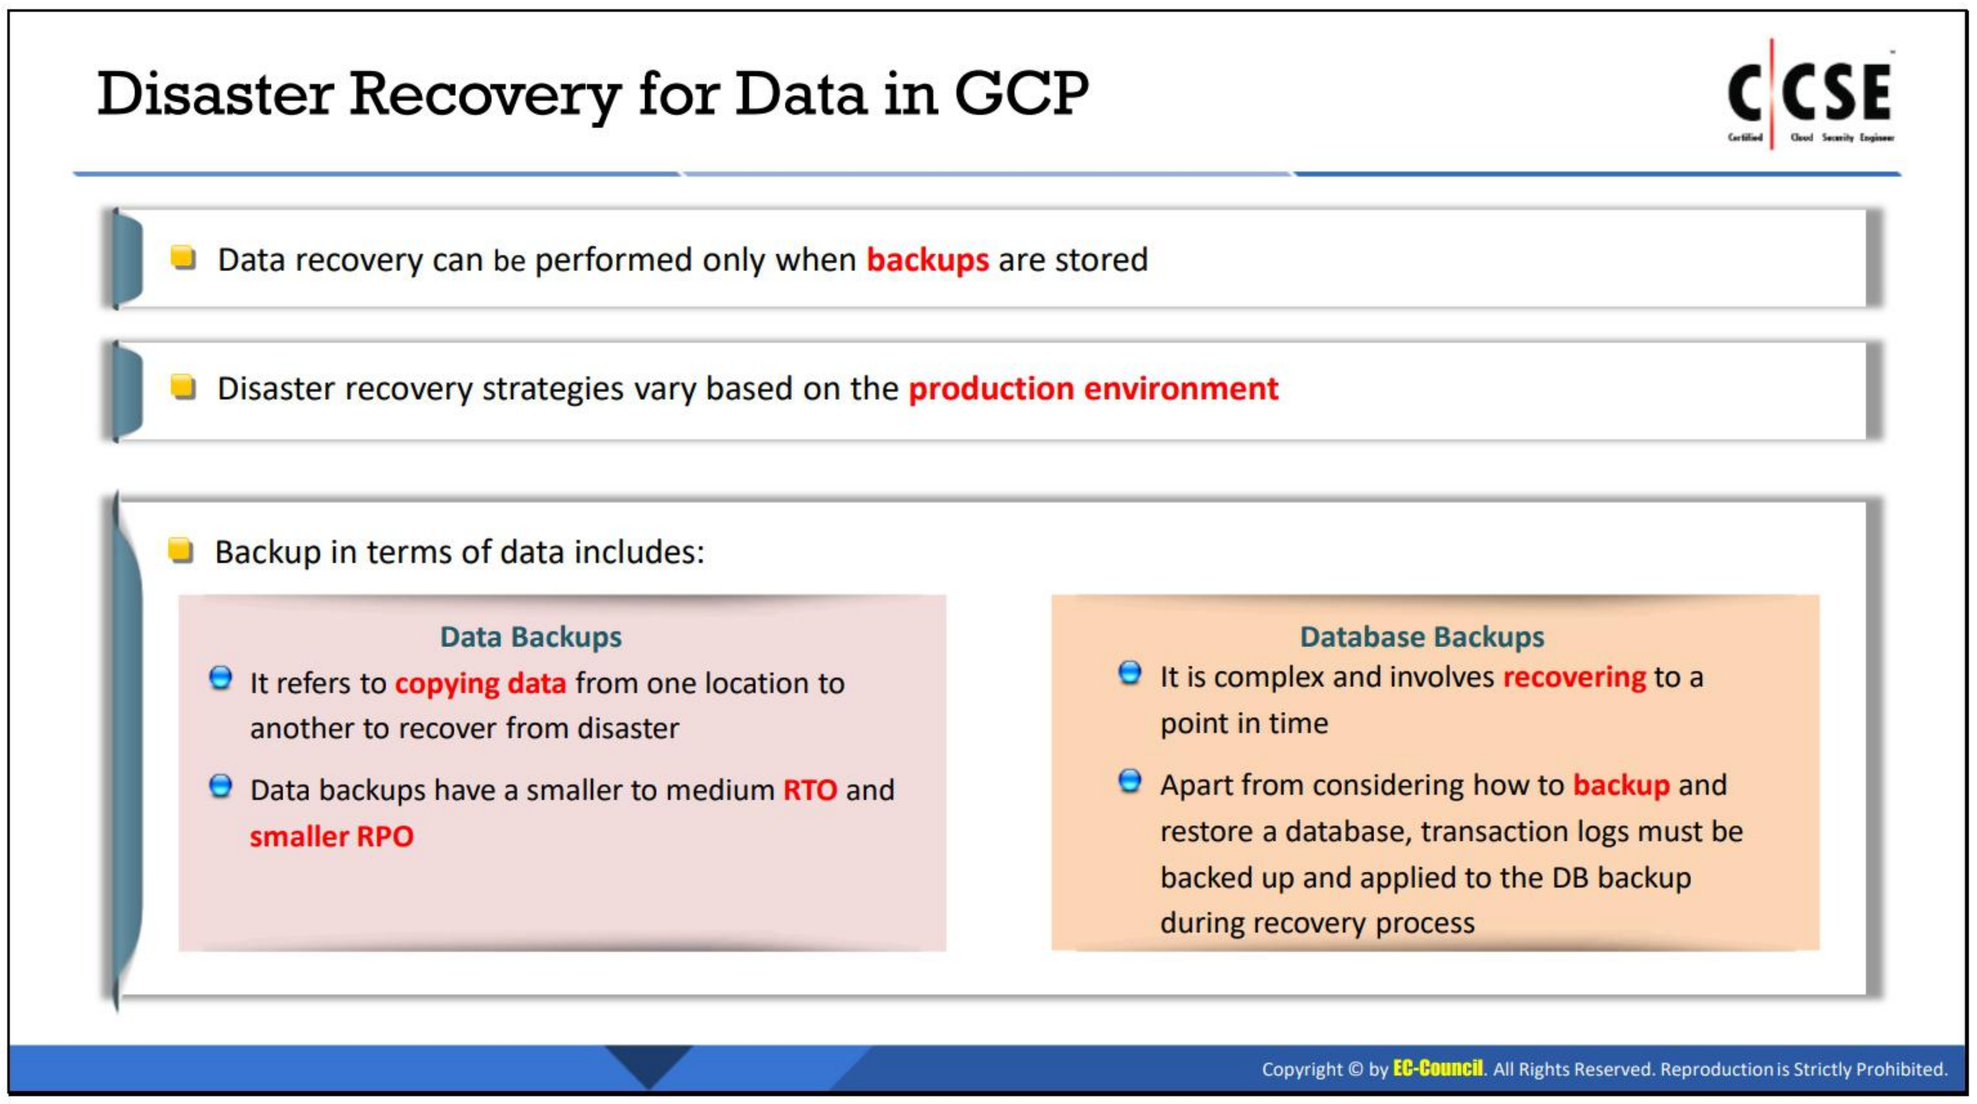

Disaster Recovery for Data in GCP

A DRP must mention how to avoid data loss during a disaster. Here, data refers to two terms, namely, data and database. The backup of data, log data, and recovering data belongs to one of the following categories:

Data backups: Data backup refers to copying a certain amount of data from one location to another. Data backups have a smaller to medium RTO and smaller RPO. Backups are a part of the recovery plan. Data backup is made either to recover from corrupted data so that it can be restored to a good state, or to restore data in DR if the production environment is down.

Database backups: Database backups are complex as they need to be recovered to a point in time. Apart from considering how to backup and restore a database, it is equally important to consider how to backup transaction logs. During database recovery, database functionality needs to be restored and then the latest database backup needs to be updated. Now, the recovered transaction logs that were backed-up after the last backup need to be applied. Due to these complications, it is important to adopt a HA method to reduce the time to recover from a disaster.

Certified Cloud Security Engineer Copyright © by **EC-Council**
All Rights Reserved. Reproduction is Strictly Prohibited.

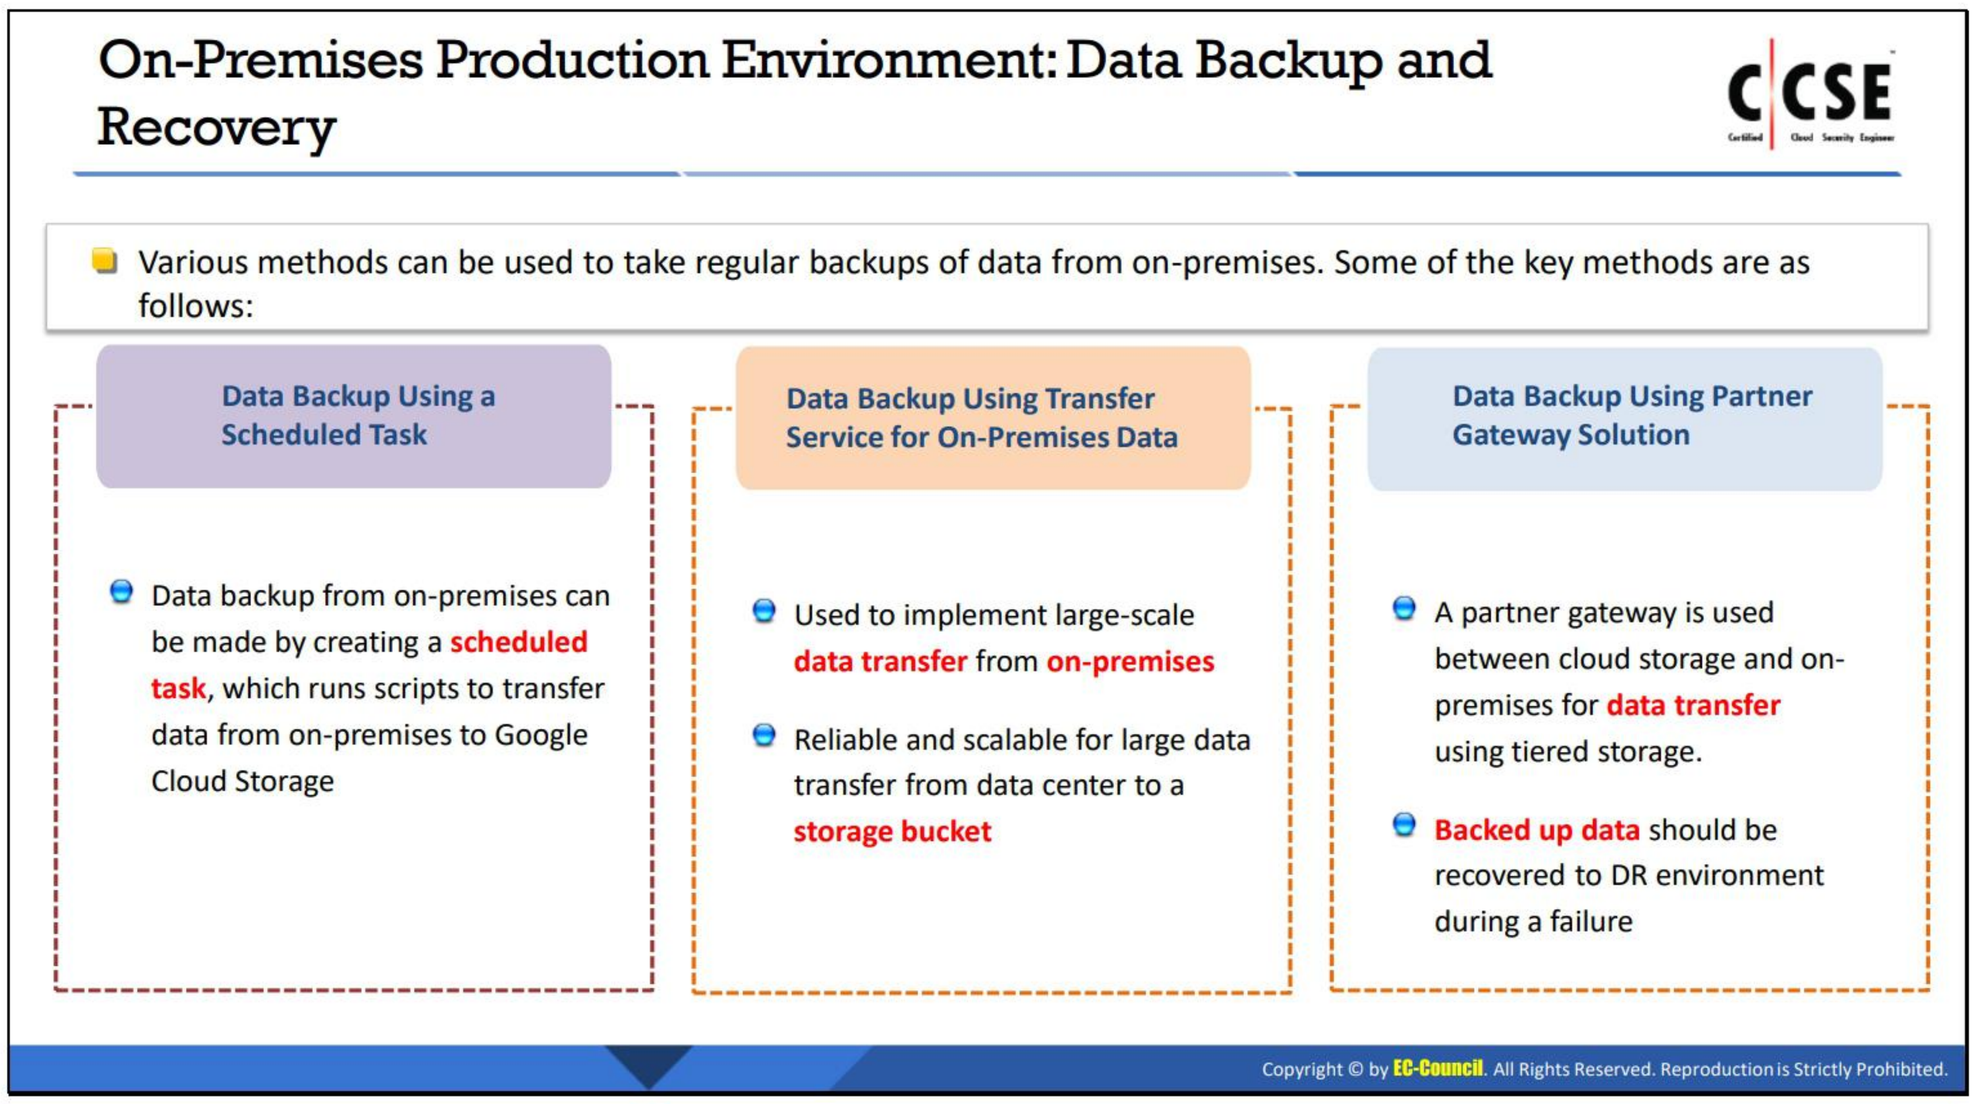

On-Premises Production Environment: Data Backup and Recovery

There are different methods through which regular data backups can be taken from on-premises to Google Cloud. Common methods include the following:

- **Data backup using a scheduled task:** Data backup from on-premises can be made by creating a scheduled task. A scheduled task runs scripts to transfer data from on-premises to Google Cloud Storage. The backup to cloud storage can be automated using client libraries or by using "gsutil."
- **Data backup using transfer service for on-premises data:** Proper planning and robust execution strategies are required for transferring large volumes of data through a network. It is difficult to develop custom scripts that are reliable and scalable. The use of custom scripts can also increase the risk of data loss and can lead to lower RPO values.

  Transfer service for on-premises data can be used to implement large-scale data transfer from on-premises. It is a reliable, scalable, and managed software service for transferring large amounts of data from a data center to a cloud storage bucket. It reduces the cost required for investing in engineering teams or for buying other transfer solutions.

  Cloud customers need to install a Docker Container with an on-premises agent for Linux on their data center. The transfer service for on-premises data transfer coordinates with the agents to transfer data to the cloud storage.
- **Data backup using partner gateway solution:** On-premises applications integrate with third-party solutions for data backups and recovery. They use a tiered storage structure where recent backups are on the faster storage and old backups are moved to slower and cheaper storage classes.

Certified Cloud Security Engineer Copyright © by **EC-Council**
All Rights Reserved. Reproduction is Strictly Prohibited.

This structure can be implemented by using a partner gateway between cloud storage and on-premises storage for performing data transfer. The following diagram shows the data transfer from on-premises to Google Cloud using a partner solution.

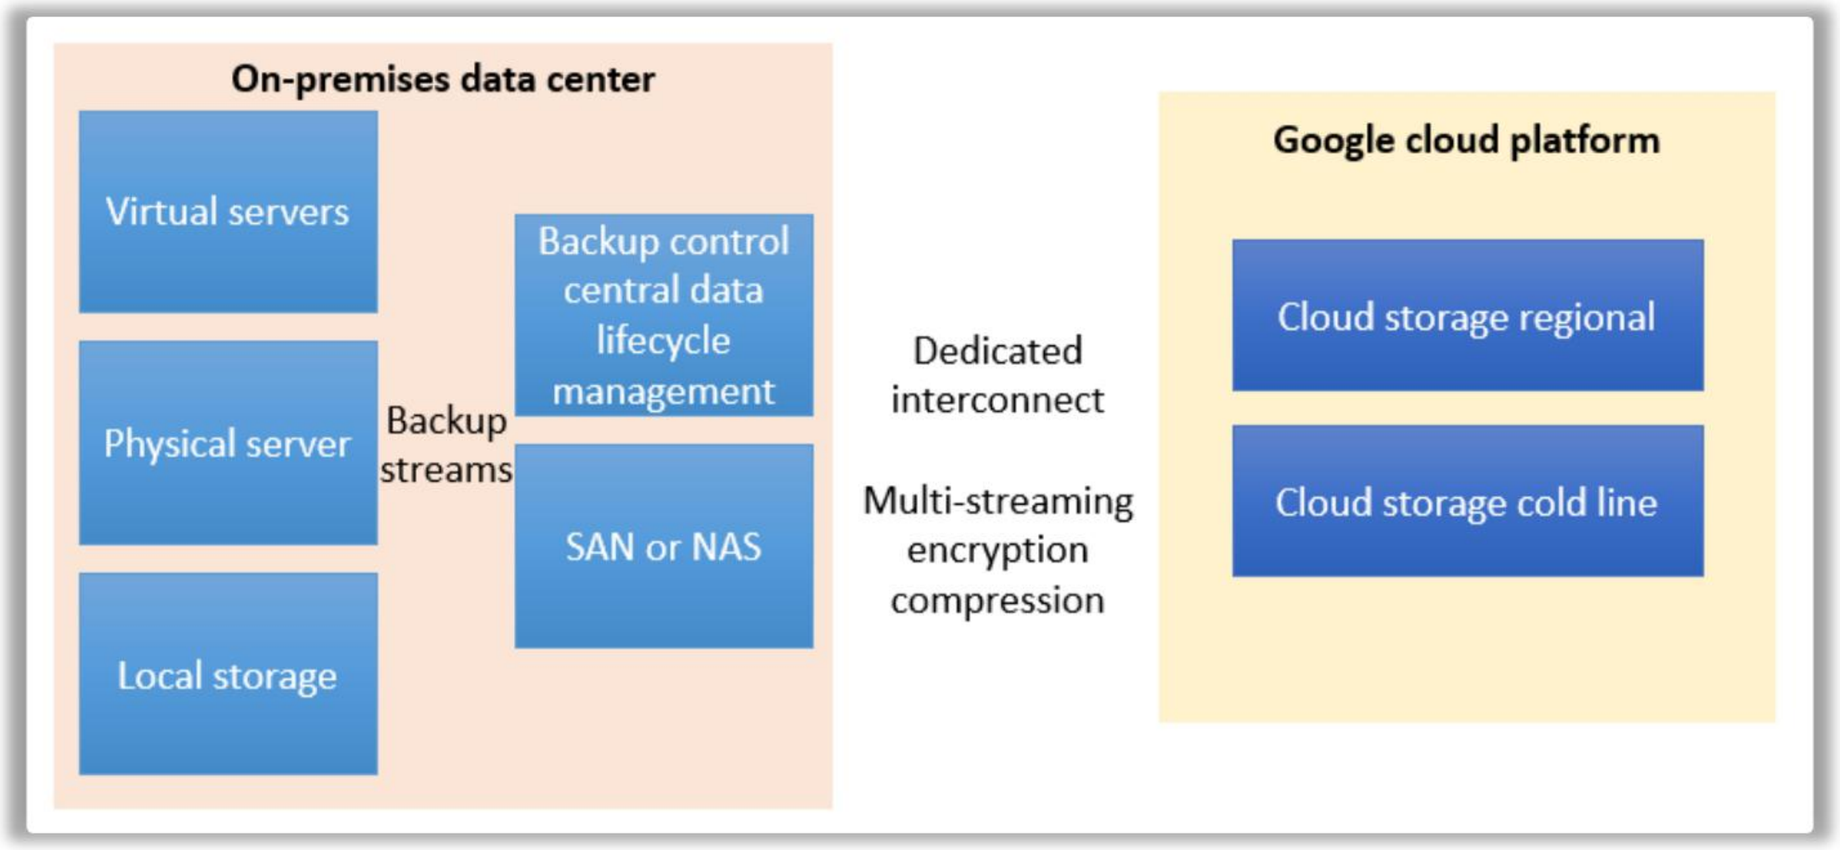

Figure 9.188: Data transfer from On-premises to Google Cloud

Whenever a failure occurs, the backed-up data must be recovered to the DR environment. The DR environment will serve the traffic till the production environment is reverted.

hide01.ir

Certified Cloud Security Engineer Copyright © by **EC-Council**
All Rights Reserved. Reproduction is Strictly Prohibited.

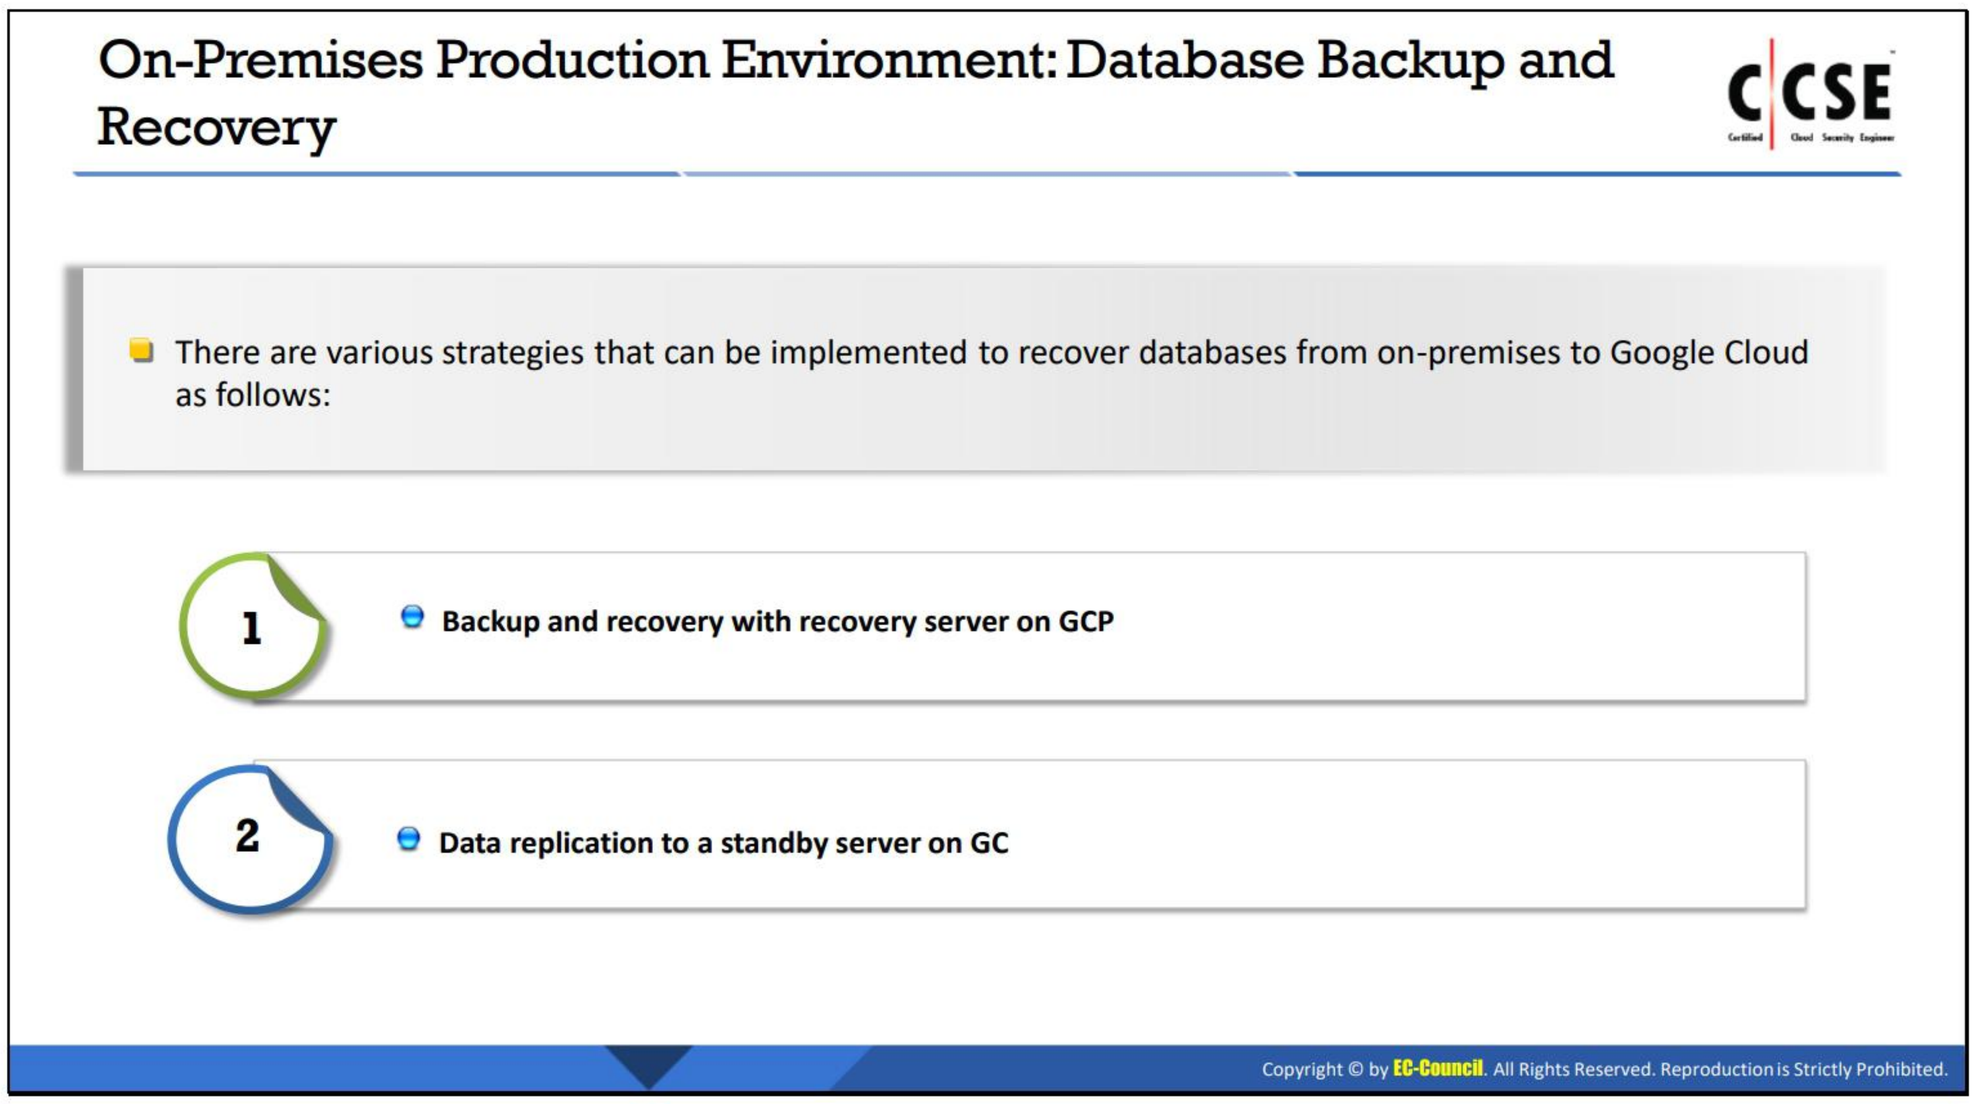

On-Premises Production Environment: Database Backup and Recovery

There are various strategies that can be implemented to recover databases from on-premises to Google Cloud.

hidel01.ir

Backup and Recovery with Recovery Server on GCP

The following steps must be taken before a disaster:

- Create a database backup of the database management system using built-in mechanisms.
- Connect on-premises network to Google Cloud network.
- Create a storage bucket for data backup in GCP.
- Now, copy the backup files to cloud storage using gsutil or using a partner gateway.
- Copy the transaction logs to the recovery location on Google Cloud.

Once the backup topology is configured, cloud customers must ensure that they can recover to the system on Google Cloud. For this, apart from restoring the backup files to a target database, transaction logs must be replayed to get the smallest RTO. The following are the recovery steps:

- Create a custom image for the database server on GCP. It should have the same configuration as the on-premises database server.
- Implement a process for copying the on-premises backup files and transaction logs to cloud storage.
- Create a minimal instance using a custom image and attach persistent disks.
- Make “auto delete” flag as “false” so that persistent disks will not be deleted accidentally.

Certified Cloud Security Engineer Copyright © by **EC-Council**
All Rights Reserved. Reproduction is Strictly Prohibited.

- During a disaster, apply the latest backup files and transaction logs that were copied to Google Cloud.
- Now, replace the minimal instance with a larger instance that can handle the production traffic.
- Now, switch the clients to point of the recovered database in GCP.

When this production environment can run and handle production workloads, the cloud security professional must reverse the steps followed for failover to the GCP recovery environment. The steps include the following:

- Collect a backup of the database running on GCP.
- Copy the backup files to the production environment.
- Now, apply the backup file to a production database system.
- Stop clients from connecting to the database system on Google Cloud. From here, the application will not be available until the production environment is restored.
- Finally, redirect the client connections to the production environment.

Data Replication to a Standby Server on GCP

To achieve smaller RTO and RPO values, a standby server can be used on Google Cloud for data replication. The following are the steps:

- Connect the on-premises network to the Google Cloud network
- Create a custom image of the database on GCP with configurations to the on-premises database server.
- Launch an instance from a custom image and attach persistent disks. Set the auto-delete flag to false.
- Now, configure replication between the target database server on GCP and the on-premises database server.
- Clients are configured to direct to database server on-premises.

After configuring this topology, switch the clients to the standby server running on the GCP network.

When the production environment on-premises is up and running, resynchronize the production database server with the GCP database server. Now, redirect the client requests to the on-premises production environment.

Certified Cloud Security Engineer Copyright © by **EC-Council**
All Rights Reserved. Reproduction is Strictly Prohibited.

Google Cloud Production Environment: Data Backup and Recovery

- The production environment and the disaster recovery environment are the same
- Data Backup and Recovery
  - Uses **tiered storage** for data backup
  - The backed-up data is **migrated** to a **lower-cost storage** tier

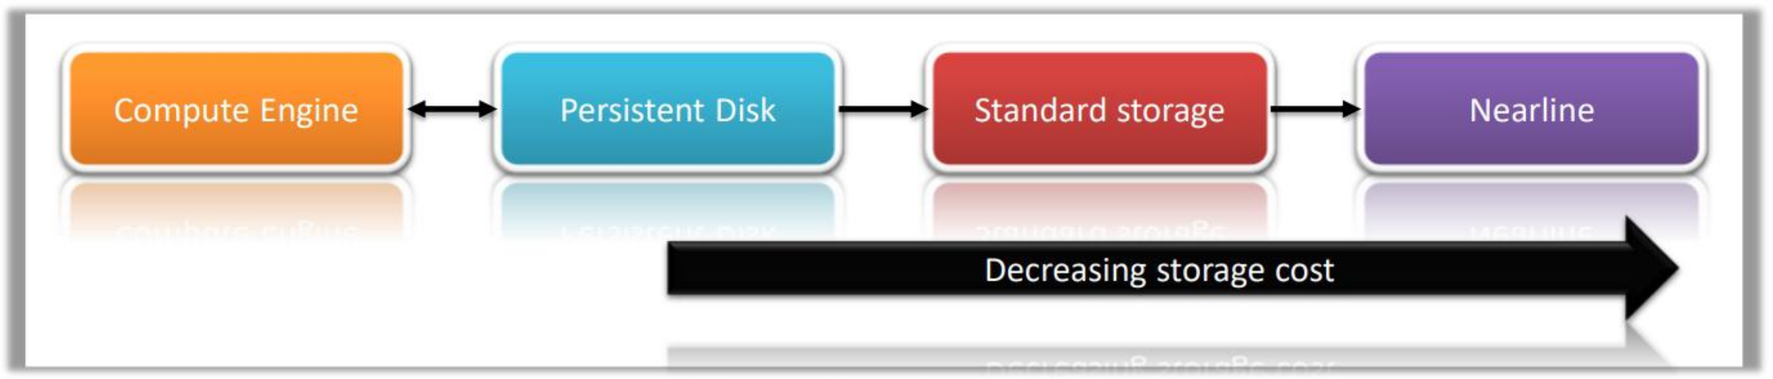

Decreasing Storage Cost Using Tiered Storage

Copyright © by EC-Council. All Rights Reserved. Reproduction is Strictly Prohibited.

Google Cloud Production Environment: Database Backup and Recovery

hide01.ir

- GCP provides different options to **design** the **database disaster recovery** architecture for self-managed databases on Google Cloud Platform such as
  - **Database server recovery without synchronization**
  - **Data recovery from partial corruption of a large database**
  - **Backup and recovery for managed database services on Google Cloud**

Copyright © by EC-Council. All Rights Reserved. Reproduction is Strictly Prohibited.

Certified Cloud Security Engineer Copyright © by **EC-Council**
All Rights Reserved. Reproduction is Strictly Prohibited.

Google Cloud Production Environment: Database Backup and Recovery (Cont'd)

DR for Managed Database Services on GCP: Certain methods can be used to implement backup and recovery strategies in each managed databases, which are as follows:

Cloud Bigtable

- Provides **Bigtable replication** to increase the availability
- Allows to **save** a **copy** of the data and table's schema

BigQuery

- Stores data in **different tables** for a specific time to handle DR
- Make **copies** of **datasets** at certain time periods for data corruption

Firestore

- Offers export and import service to recover from **accidental deletion** of data

Cloud SQL

- Enables **automated backups** and also binary logging for point-in-recovery

Cloud Spanner

- Uses **dataflow connectors** for controlled backups
- Uses "**commit timestamp**" feature for incremental backups

Cloud Composer

- Used for **scheduling regular backups** of different Google Cloud databases

Copyright © by EC-Council. All Rights Reserved. Reproduction is Strictly Prohibited.

Google Cloud Production Environment

In this case, both the production environment and the disaster recovery environment are the same.

hide01.ir

Data Backup and Recovery

A common method for data backup in this scenario is to use a tiered storage pattern. In this method, the backed-up data is migrated to a lower cost storage tier. As mentioned earlier, Google Cloud offers Nearline, Coldline, and Archive storage classes. These storage classes are intended for storing less frequently accessed data.

There will be additional costs for retrieving the data from these storage classes and there are minimal durations for which the customers are charged for.

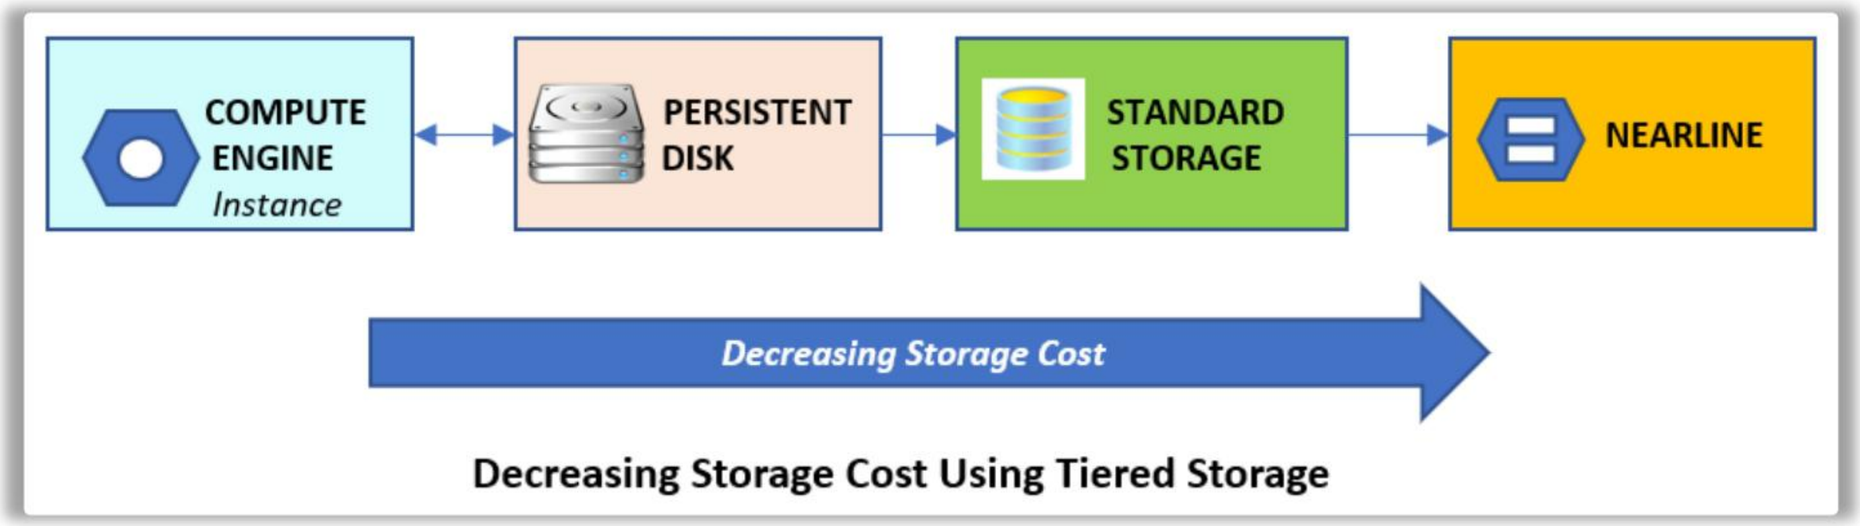

Decreasing Storage Cost Using Tiered Storage

Figure 9.189: Decreasing Storage Cost Using Tiered Storage

Certified Cloud Security Engineer Copyright © by **EC-Council**
All Rights Reserved. Reproduction is Strictly Prohibited.

Database Recovery and Backup

When a self-managed database is used, the same operational concerns apply as the managing production database on-premises, but customers no longer need to maintain or manage the underlying infrastructure.

By using appropriate features of DR building blocks, cloud customers can set up HA configuration to keep smaller RO values. The database configuration must be designed in such a way that it is easy to recover to a state that is closer to the pre-disaster state. This helps keep smaller RPO values. There are various options offered by GCP for designing the database disaster recovery architecture for self-managed databases on GCP.

Recovering a Database Server Without Synchronizing State

The most common pattern is to recover the database server without synchronizing the system state. The building blocks used for this scenario are Compute Engine, MIGs, and Cloud load balancing.

The following are the steps to configure the above scenario:

- Create a VPC network in Google Cloud.
- Create a custom image with the configuration of the database server with the following steps:
    - Configure the database server such that database files and log files can be moved to persistent disks.
    - Create a snapshot from the persistent disk.
    - Now, configure a startup script for creating a persistent disk from the snapshot and for mounting the disk.
    - Then, create a custom image of the boot disk.
- Create an instance template using the custom image.
- Configure a regional managed instance group using the instance template with the target size as 1.
- Using Cloud Monitoring metrics, configure health checks.
- Utilizing the existing regional managed group, set an internal load balancing.
- Now, take regular snapshots of persistent disk by setting up a scheduled task.

When the database instance needs to be replaced, the above configuration creates a new database server with the correct version. It then attaches the persistent disks with the latest backup and transaction log files to the new database instance. It also eliminates the need for reconfiguring the clients who communicate with the database server in response to an event. As the replacement instance is created from the instance template, the security controls applied to the database server will be applied to the recovered database server.

Certified Cloud Security Engineer Copyright © by **EC-Council**
All Rights Reserved. Reproduction is Strictly Prohibited.

The cloud security professional does not need to initiate steps for failover as it automatically takes place in case of a disaster. The internal load balancer ensures that the replacement database instance has the same IP address as the original database server. The instance template and custom image are responsible for ensuring that the configuration of the replacement instance is similar to the instance it is replacing. By taking regular snapshots, cloud customers ensure that the disks are recreated using snapshots and are attached to the new replacement instance. The transaction log files that were written to the persistent disks are also automatically restored.

The regional instance group provides HA, and it reacts to the failover at the application, instance, or zonal level. The target size 1 ensures that there will be only one running instance.

Data Recovery from Partial Corruption of a Large Database

When cloud customers have a large database that can store petabytes of data, there can be an outage that could affect a portion of the data. In that case, cloud customers need to minimize the amount of data that needs to be restored. It is not required to restore the entire database for restoring a portion of the data.

The following strategies would help cloud customers in this scenario:

- The data can be stored in different tables for certain time periods. This helps ensure that they need to restore only a subset of data to a new table but not the entire data.
- The original data can be stored. This lets the cloud security professional create a new table and reload only the uncorrupted data. From here, the applications can be pointed to the new table.

hidden.it

If the organization's RTO allows, cloud customers can also prevent access to the table that contains the corrupted data. This can be achieved by making the application offline till the corrupted data is restored to a new table.

Backup and Recovery for Managed Database Services on Google Cloud

There are certain methods that can be used to implement backup and recovery strategies for the managed database services on Google Cloud. Managed databases are designed to provide scaling and hence the backup-and-restore mechanisms of traditional RDBMS are not applicable. The amount of data that need to be restored during a DR scenario needs to be reduced.

There are several strategies for each type of managed database for minimizing the amount of data that need to be restored.

Cloud Bigtable

Cloud Bigtable is a managed and scalable NoSQL database service for large-scale operational workloads. It provides Bigtable replication that increases availability by copying data to multiple regions or zones. It is also useful for resilience during an event of zonal or regional failures.

Cloud Bigtable is a fully managed service. It allows organizations to save a copy of the data and table's schema. From the backup, data can be restored to a new table when required.

Certified Cloud Security Engineer Copyright © by **EC-Council**
All Rights Reserved. Reproduction is Strictly Prohibited.

The tables from Cloud Bigtable can also be exported as a series of Hadoop sequence files. The sequence files can be stored in cloud storage or they can be used to import the data into another instance in Bigtable. Datasets can also be replicated asynchronously in various zones within a region.

BigQuery

Long-term storage of BigQuery can be used for archive data. When a table is not edited for more than 90 consecutive days, the storage cost for that table drops by 50% without degrading the availability, performance, and durability.

BigQuery is replicated automatically; however, it does not help with data corruption. Cloud customers should have a plan to recover from the DR scenarios. The cloud security professional can perform the following to handle DR:

- If the data corruption is noticed within 7 days, the table can be queried to a point in time in the past for recovering the table using snapshot decorators.
- Data can be exported from BigQuery. A new table should be created with the exported data and the corrupted data should be excluded.
- The data can be stored in different tables for a specific time. This helps ensure that only a subset of data needs to be restored to a new table and not the entire dataset.
- Make copies of datasets at certain time periods. The copies can be used in case of a data corruption event that occurred beyond a point in time the query can capture. The dataset can also be copied from one region to another to provide high availability in case of region failures.
- The original data can be stored on Google Cloud Storage. This allows the cloud security professional to create a new table and reload only the uncorrupted data. From then, the applications can be pointed to the new table.

hi0n01.ir

Firestore

Firestore offers export and import service to recover from accidental deletion of data. Using this service, Firestore entities can be imported and exported using a cloud storage bucket. The data exported from Firestore can also be loaded into BigQuery.

Cloud SQL

While Cloud SQL is used, the cloud security professional should enable automated backups as well as binary logging for the Cloud SQL instances. This helps perform a PITR that restores the database from the backup. The restored database can be recovered to a new Cloud SQL instance.

Cloud SQL can be configured with HA configuration and cross-region replicas to increase the uptime in case of a zonal or regional failure.

Cloud Spanner

Cloud Spanner is a fully managed relational database with unlimited scaling and 99.999% availability. Using a dataflow template, the database can be exported to a set of Avro files in the

Certified Cloud Security Engineer Copyright © by **EC-Council**
All Rights Reserved. Reproduction is Strictly Prohibited.

cloud storage bucket. Now, by using another dataflow template, the exported files can be reimported to a new Cloud Spanner database.

Dataflow connectors can also be used for more controlled backups. Dataflow connector lets the cloud security professional read and write the data to Cloud Spanner in a dataflow pipeline. The speed for reading the data from Spanner is based on the number of nodes configured.

Cloud Spanner has a "commit timestamp" feature for incremental backups. It allows selection of only those rows that were added or modified after the last backup. Cloud Spanner Backup and Restore can be used for managed backups. Cloud Spanner Backup and Restore allows creation of backups of the Spanner database whenever required, which can be later restored when data corruption occurs.

To achieve smaller RTO values, a warm standby Cloud Spanner instance can be configured with a lesser number of nodes to meet storage requirements.

Using Spanner PITR, data can be recovered from a specific point of time. If an operator accidentally writes data or an application rollout corrupts the database, PITR can be used to recover data from a point in time in the past.

Cloud Composer

Cloud Composer can be used for scheduling regular backups of different Google Cloud databases. A directed acyclic graph (DAG) can be created to schedule copying data to another project, or table, or to export data to cloud storage. Various cloud platform operators can be used for copying and exporting data.

hide01.ir

Certified Cloud Security Engineer Copyright © by **EC-Council**
All Rights Reserved. Reproduction is Strictly Prohibited.

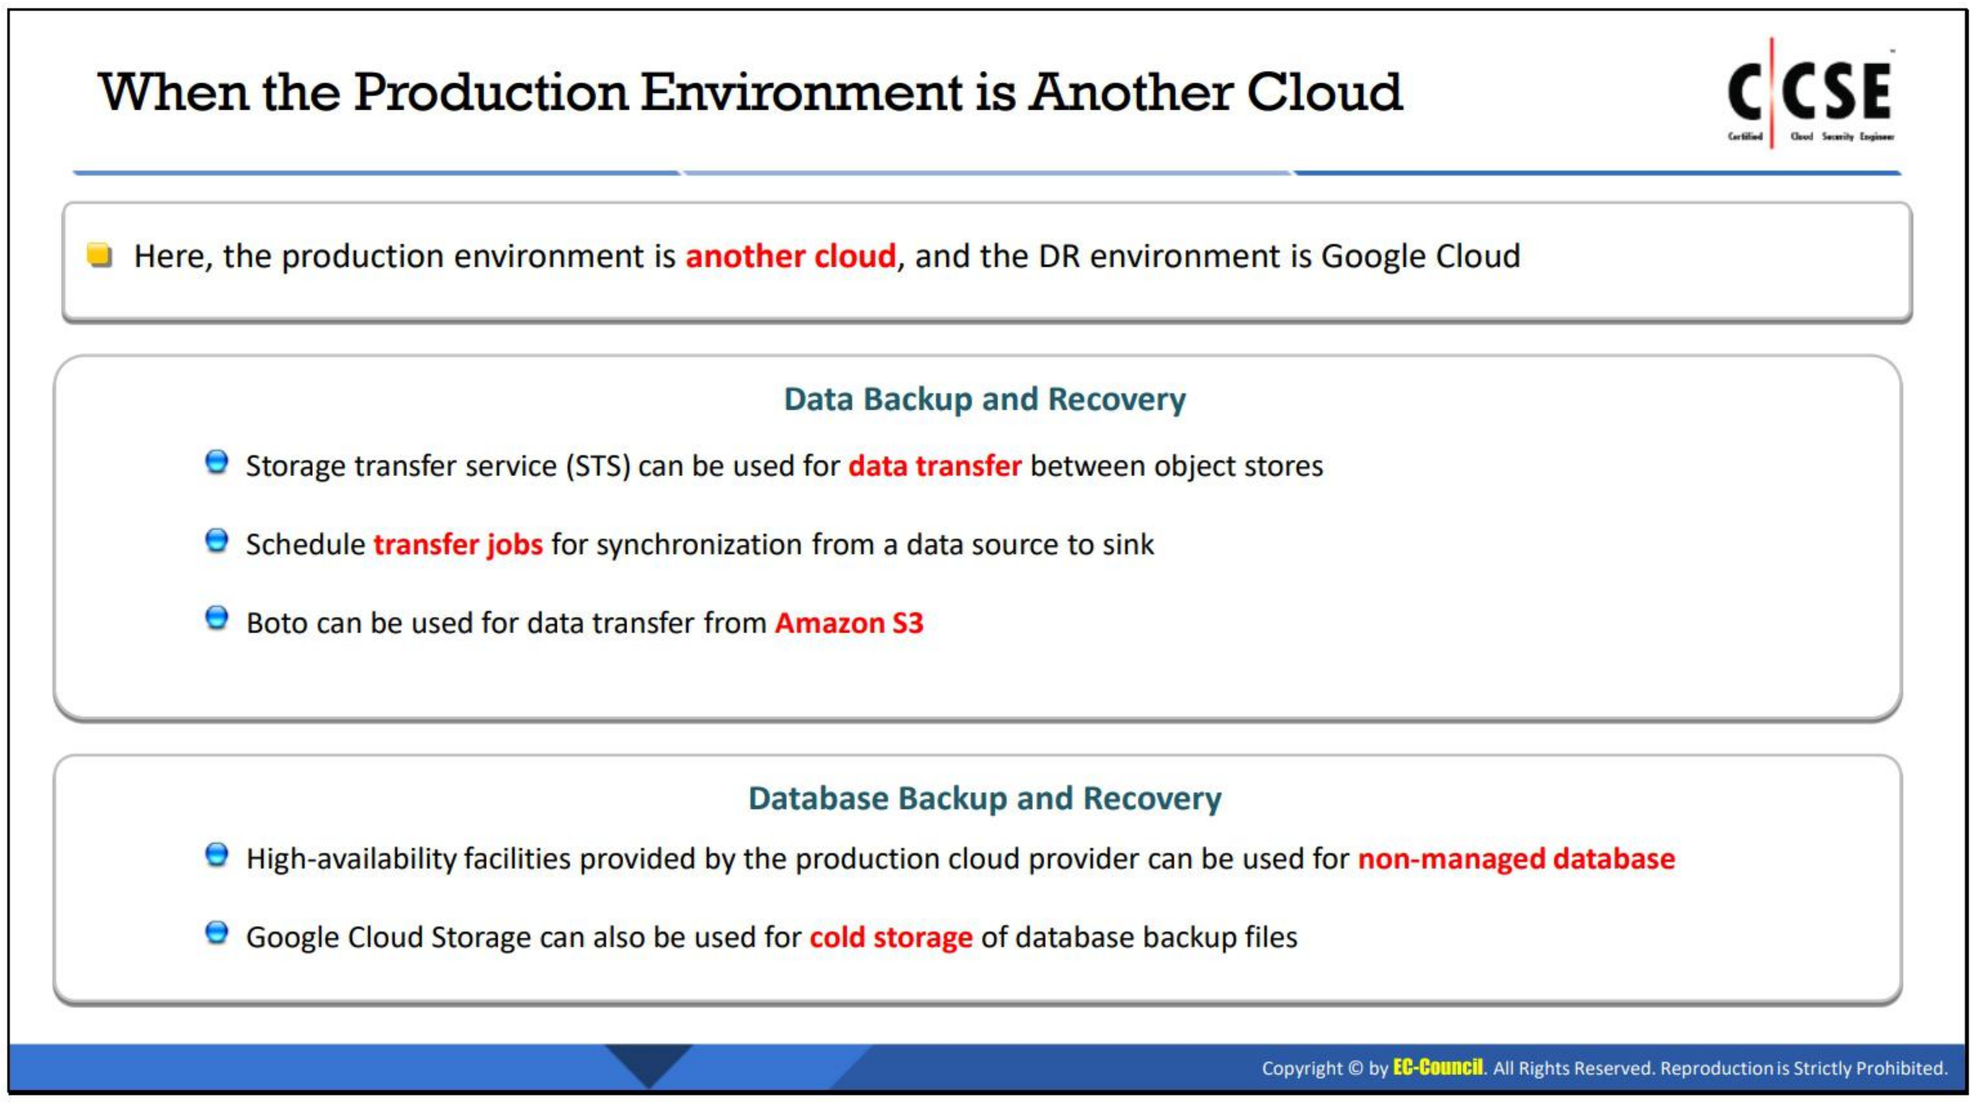

When the Production Environment is Another Cloud

In this case, the production environment is another cloud, and the DR environment is Google Cloud.

hide01.ir

Data Backup and Recovery

It is a common DR use case for transferring data from one object store to another. STS can be used for data transfer between object stores. STS is compatible with the Amazon S3 bucket and it helps in transferring data from the S3 bucket to cloud storage.

The cloud security professional can schedule transfer jobs for synchronization from a data source to sink. Advanced filters can be applied during data transfer.

The following factors must be considered to achieve the desired RPO:

- **Rate of change:** Rate of change is the amount of data that is generated or updated during a certain period. The higher the rate of change, the greater number of resources will be required for transferring the changes to the destination for every incremental transfer period.
- **Transfer performance:** It is the time required for transferring the files. In the case of large file transfers, transfer performance is determined with the available bandwidth between source and destination.

  If the transfer job has a large number of small files, queries per second (QPS) can be a factor. In this case, the cloud security professional can schedule multiple concurrent jobs to increase the performance until the bandwidth is available.
- **Frequency:** Frequency is the interval between backup jobs. The intervals between two transfer jobs should not be more than the RPO objective.

Certified Cloud Security Engineer Copyright © by **EC-Council**
All Rights Reserved. Reproduction is Strictly Prohibited.

- **Monitoring and alerts:** STS provides Pub/Sub notifications for various types of events. Cloud customers must subscribe to these notifications so that they can handle any changes in the job completion time or any unexpected failures.

Data transfer from Amazon S3 can also be performed using Boto, which is a Python tool compatible with Amazon S3 and Cloud Storage.

Database Backup and Recovery

If the cloud customers are using non-managed databases on the compute service, they can make use of high-availability facilities provided by the production cloud provider. It can be extended to incorporate HA deployment to Google Cloud. Google Cloud Storage can also be used as the destination for the cold storage of the production database backup files.

hide01.ir

Certified Cloud Security Engineer Copyright © by **EC-Council**
All Rights Reserved. Reproduction is Strictly Prohibited.

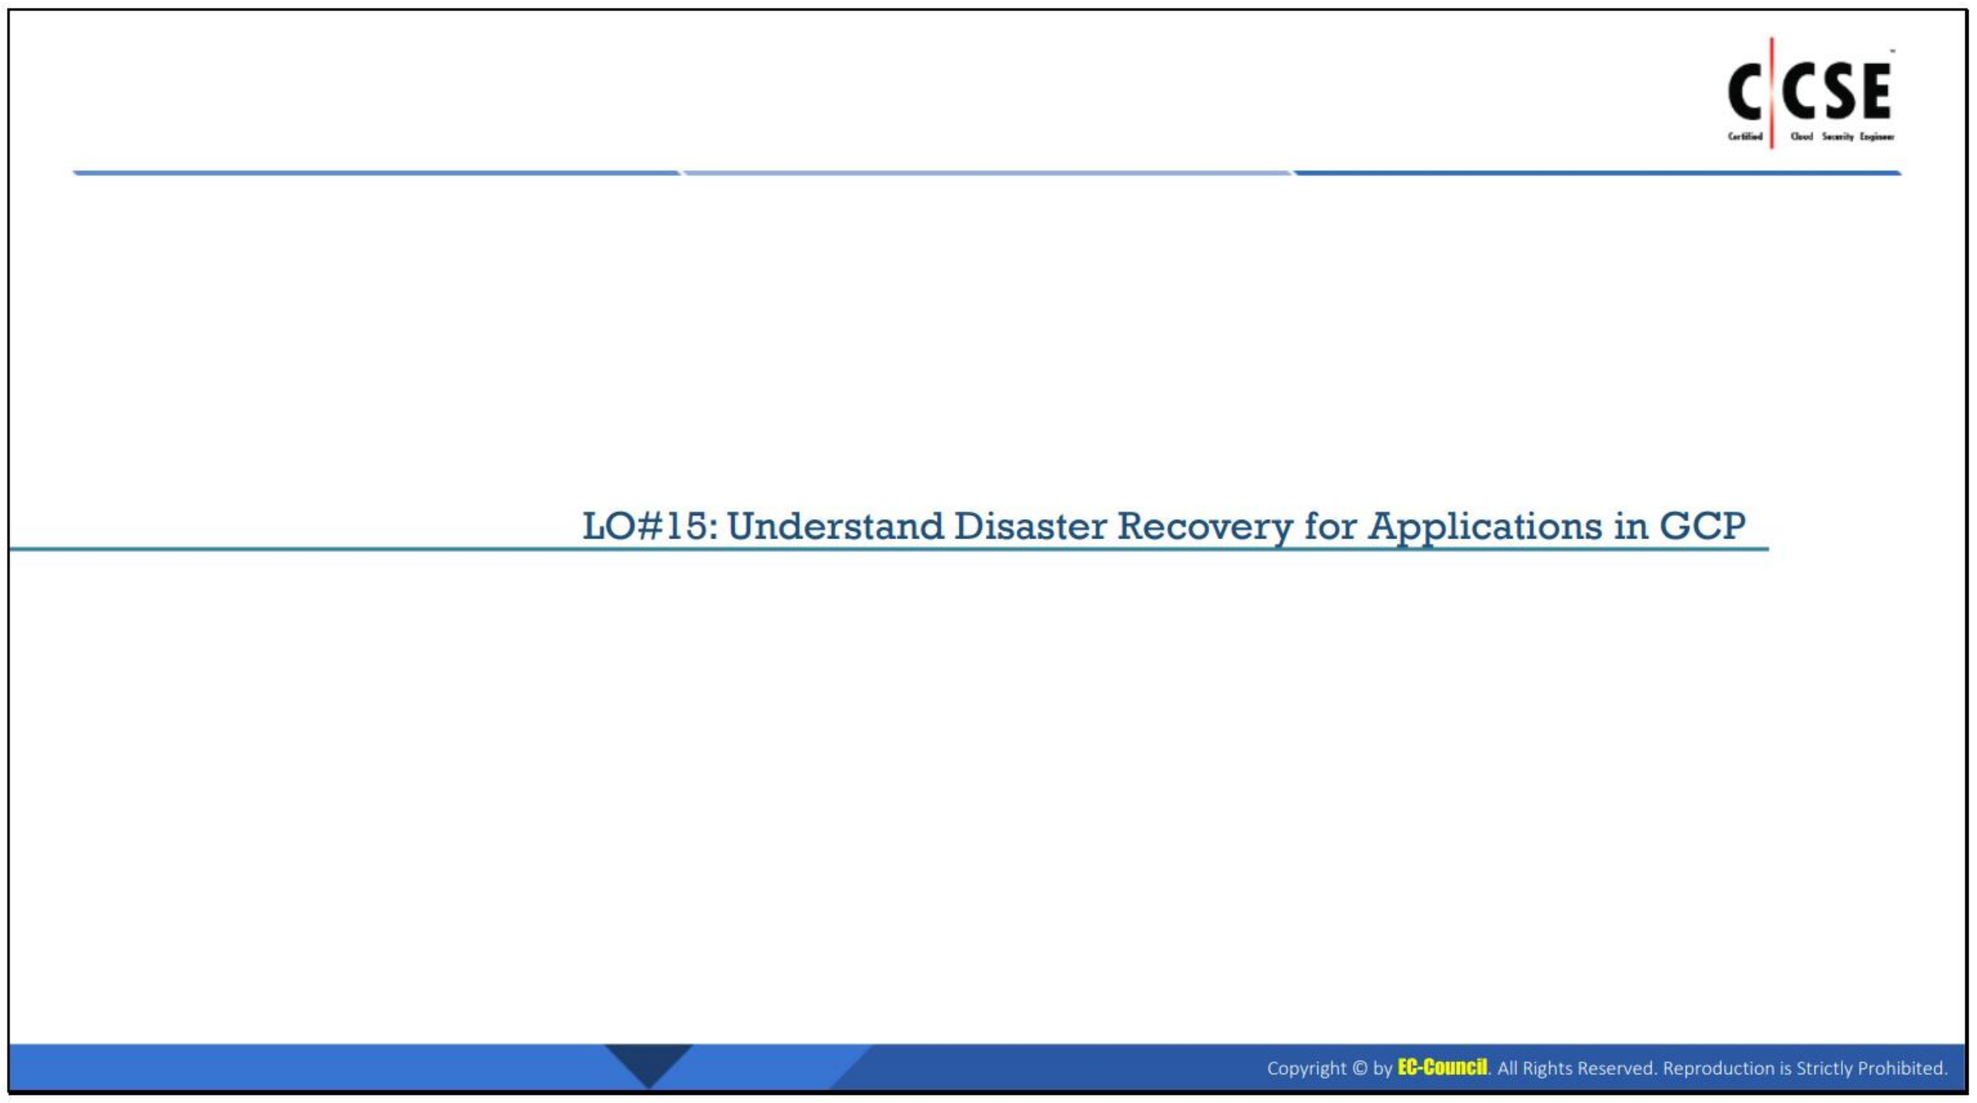

LO#15: Learn Application Disaster Recovery in GCP

The objective of this section is to understand disaster recovery for applications in GCP. This section introduces various disaster recovery and HA architectures for on-premises and Google Cloud applications. In addition, cloud security engineers can learn how to plan DR for cloud infrastructure outages and locality-restricted workloads.

Certified Cloud Security Engineer Copyright © by **EC-Council**
All Rights Reserved. Reproduction is Strictly Prohibited.

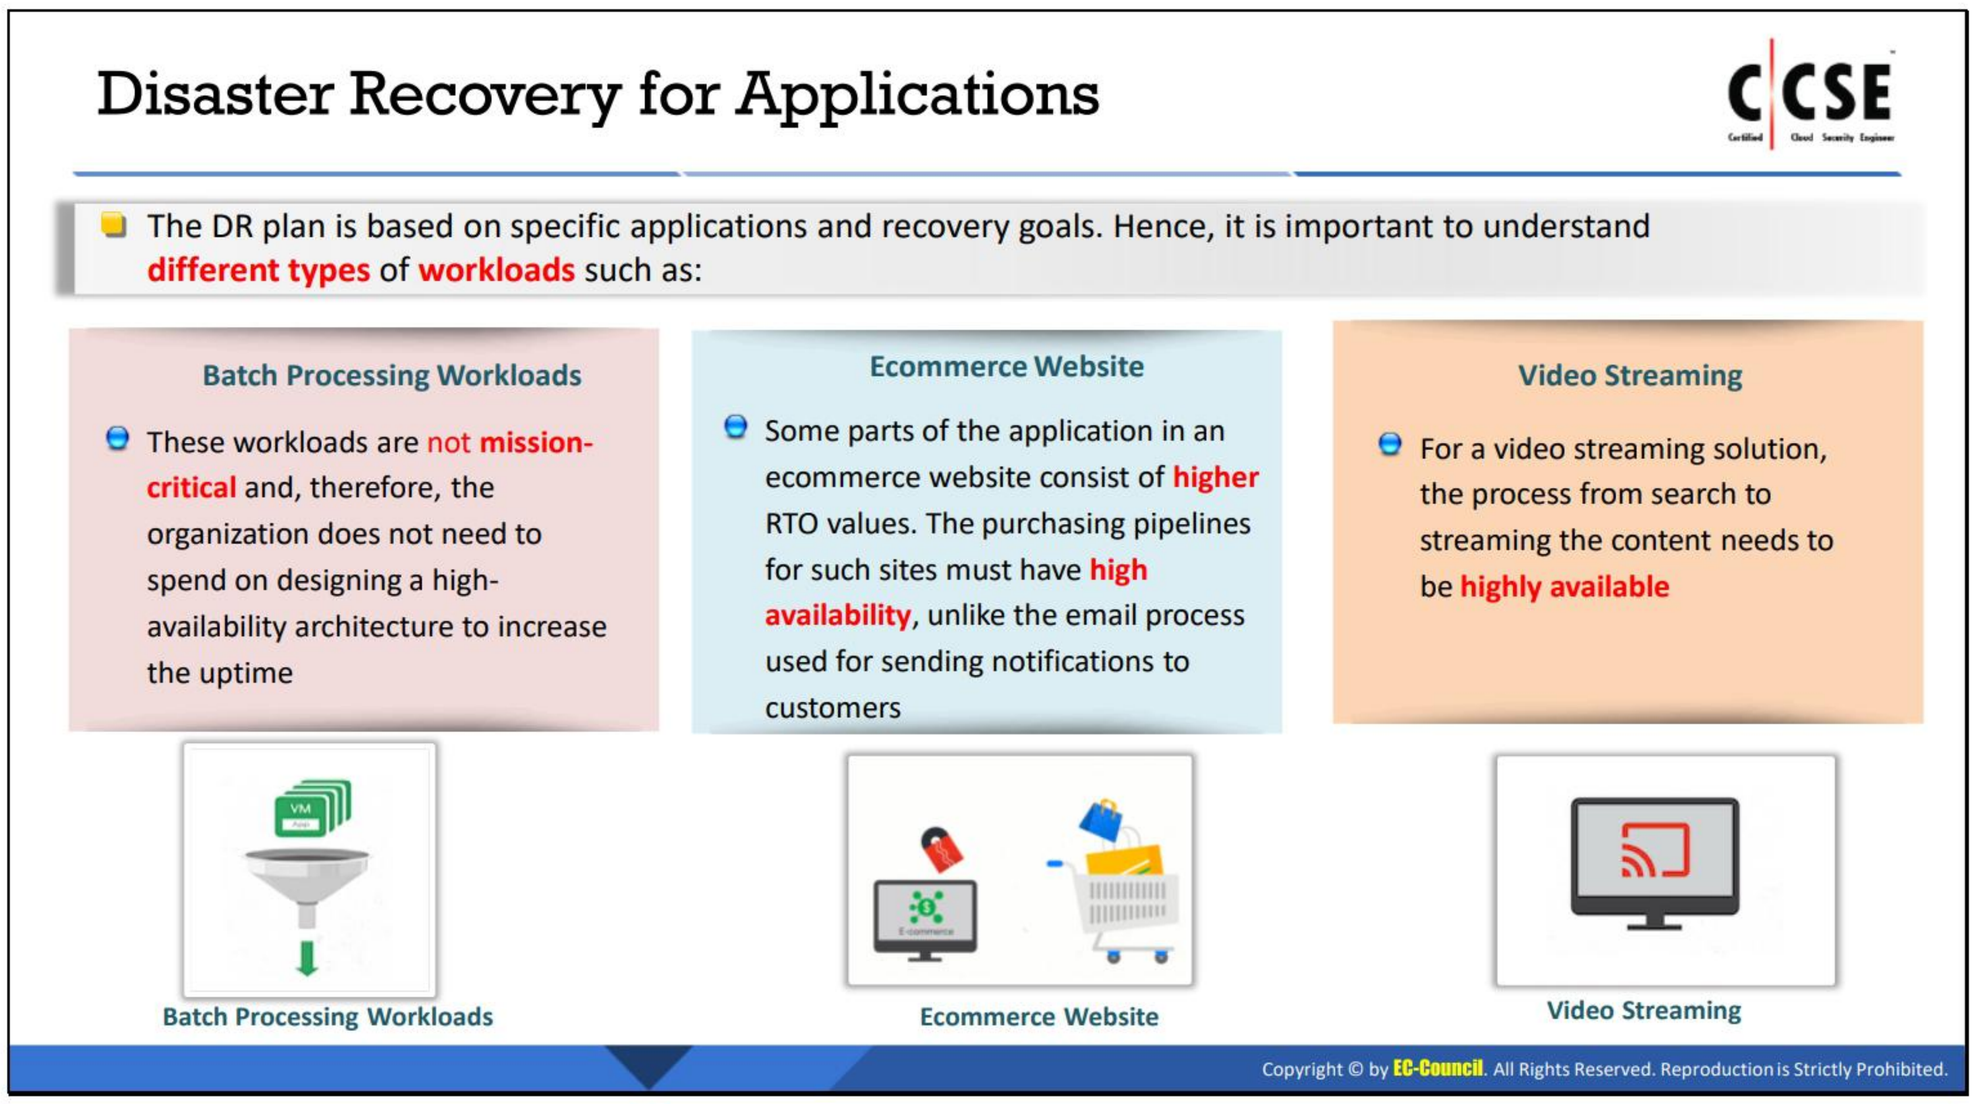

Disaster Recovery for Applications

Applications can be recovered from a disaster using the DR building blocks and their features. The DR plan is based on the specific applications and the recovery goals. Before proceeding, it is important to understand different types of workloads.

Batch Processing Workloads

Batch processing workloads are not mission-critical. Thus, the organization does not need to spend on designing a high-availability architecture to increase the uptime. In such cases, cost-effective products such as “preemptible VM” can be used. A preemptible VM is an instance that can be created at a lower cost than normal instances. By enforcing checkpoints along with processing, the processing job will recover from the point of failure when launching a new VM. If Dataproc is utilized, the MIGs will launch preemptible nodes. It is considered as warm pattern as there is a waiting time for replacement instances.

Ecommerce Sites

Certain parts of an ecommerce website application consist of higher RTO values. The purchasing pipelines for such sites must have HA; however, the email process used for sending notifications to customers can tolerate delays. It is a blend of hot and cold or warm patterns. The transaction part of the application must have minimal RTO values and high uptime. Therefore, high-availability architecture that maximizes this part of the application should be utilized and it is deemed as hot pattern.

Video Streaming

For a video streaming solution, the process from search to streaming the content needs to be highly available. It should also have a low latency to provide a satisfactory experience to the

Certified Cloud Security Engineer Copyright © by **EC-Council**
All Rights Reserved. Reproduction is Strictly Prohibited.

users. If the solution fails to provide a great experience to users, users can easily switch to other competitors.

Hence, high-availability architecture is mandatory for video streaming solutions with smaller RTO values. This case requires a hot pattern throughout the application to reduce the effect of disaster.

hide01.ir

Certified Cloud Security Engineer Copyright © by **EC-Council**
All Rights Reserved. Reproduction is Strictly Prohibited.

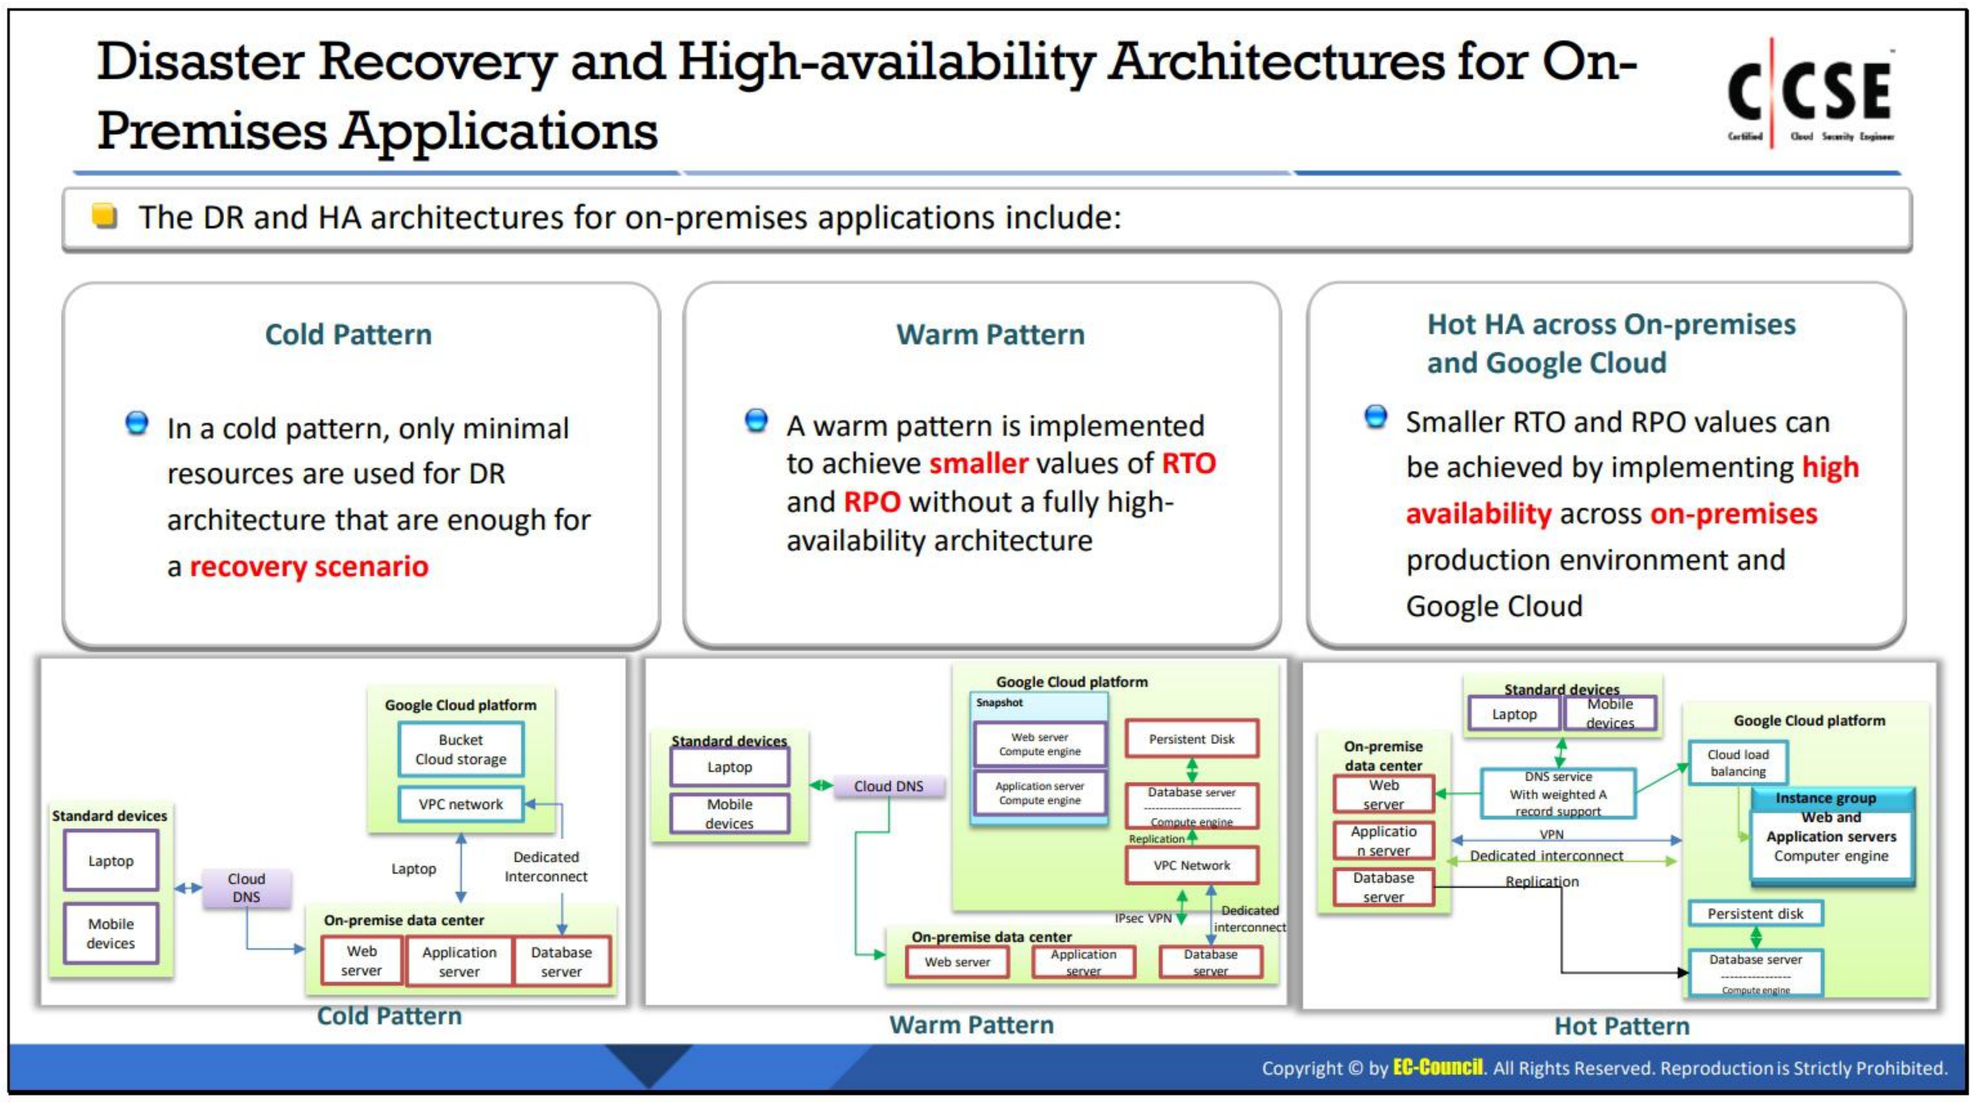

Disaster Recovery and High-availability Architectures for On-premises Applications

The DR and high-availability architectures for on-premises applications include:

Cold Pattern for Recovering Applications to Google Cloud

In a cold pattern, only minimal resources are used for DR architecture that are enough for a recovery scenario. When the production environment stops running production traffic, the failover method should have a mirror of the production environment to get started in GCP. From then, clients can start using the services from the DR environment.

The following are the steps for implementing the architecture:

- Create a VPC network in GCP.
- Establish a connection between the on-premises network and Google Cloud network
- For data backup, create a Storage bucket.
- Create a service account key to upload database backups of on-premises to Google Cloud.
- Copy the service account key to the on-premises location where the scripts are run to upload database backups.
- Create an IAM policy to control the access to the bucket and its objects. Make sure that the service account has required permissions.
- Test that the upload and download are working in the target bucket.
- Create a script for data transfer.
- Create a scheduled task to run the script.

Certified Cloud Security Engineer Copyright © by **EC-Council**
All Rights Reserved. Reproduction is Strictly Prohibited.

- For each type of server in the production environment, generate custom images. The configuration of the custom images should match the configurations of its on-premises servers.
- Now, configure DNS to point to the internet-facing web services.
- Create a Deployment Manager to create application servers in the Google Cloud network using the configured custom images.

In the event of a failure, the cloud security professional can recover to the system that is running on GCP. Using the Deployment Manager template, the recovery environment can be created to initiate the recovery process. When the recovery environment instance is running and are ready to accept traffic, the DNS can be adjusted to point to the webserver in GCP.

The following are the steps for recovery:

- Create a deployment in Google Cloud using Deployment Manager Template.
- Now, configure the database server running in GCP with the latest database backup available in the cloud storage.
- Now, apply the most recent transaction logs in cloud storage.
- To ensure that the application works as per expectation, user scenarios should be simulated on the recovered cloud environment.
- When the test is successful, point the cloud DNS to the web server on Google Cloud.

As soon as the production environment starts working and can support production traffic and workloads, the failover steps used for Google Cloud recovery environment should be reversed. The steps include the following:

- Take the database backup running on GCP.
- Make a copy of the backup file to the production environment.
- Now, apply the backup file to the production database system.
- Stop the connection to the application on Google Cloud. From now on, the application will be unavailable until the production environment is restored.
- Copy the transaction log files to the production environment and apply those log files to the database server.
- Point the DNS to on-premises web service.
- Test the process that copies the data to cloud storage is working as expected.
- Now, delete the deployment.

Warm Pattern for Recovering Applications to Google Cloud

A warm pattern is implemented to achieve smaller values of RTO and RPO without a fully high-availability architecture. If the RTO and RPO values are smaller, then the implementation cost for fully redundant DR architecture is high. Hence, a warm pattern for the DR scenario is a good choice to meet budget and availability.

Certified Cloud Security Engineer Copyright © by **EC-Council**
All Rights Reserved. Reproduction is Strictly Prohibited.

The following are the steps for configuring this scenario:

- Create a VPC network in Google Cloud.
- Establish a connection between the on-premises network and the GCP network.
- Replicate the on-premises server to Google Cloud instances.
- Now, create a custom image of the database server running on GCP that should have configurations of the database server running on-premises.
- Create snapshots of application and web server instances.
- Using the custom image created earlier, start the database instance in Google Cloud. Use the smallest machine type that accepts replicated data from the on-premises production database.
- Attach the persistent disks to the database instance in GCP.
- Initiate the replication process between the on-premises database server and the database server in GCP.
- Adjust the auto-delete flag of the persistent disks to "no auto-delete."
- To take regular snapshots of persistent disks of the database instance running on GCP, set up a scheduled task.
- Now, test the process used for launching instances from snapshots and test the process used for taking persistent disks snapshots.
- Web and application server instances should be created utilizing the snapshots.
- Write a script that can copy the updates of the associated on-premises servers to the web and application servers. Additionally, build a script for creating snapshots of the updated servers.
- Now, point the cloud DNS to the internet-facing web services that are on-premises.

If a failover occurs, an automated failover process can be invoked using a monitoring and alerting system. Database system on the Google Cloud is configured when the on-premises application needs to fail over so that it can accept the production traffic. The recovery steps are as follows:

- Resize the database instance such that it can handle the production workloads.
- Create web server and application server instances using web server and application snapshots, respectively.
- Test the application by simulating the user scenarios on the recovered environment.
- Once the tests are successful, point the cloud DNS to the web service on GCP.
- Once the production environment is up and accepts the production workloads, reverse the steps used for failover to the Google Cloud recovery environment. The steps are as follows:
    - Take a backup of database on GCP.

Certified Cloud Security Engineer Copyright © by **EC-Council**
All Rights Reserved. Reproduction is Strictly Prohibited.

- The backup file should be copied to the production environment.
- Now, the copied backup file should be applied to the production database system.
- Stop all the connections to the applications on the GCP. This can be achieved by modifying the firewall rules to prevent connections to the web server. From now on, the application will not be available until the production restoration is complete.
- Copy the transaction log files to the production environment and apply those transaction logs to the database server.
- Test the application if it is working as expected.
- Set up cloud DNS to point to on-premises web service.
- Run only the reference servers and delete both the web and application server on Google Cloud.
- Reduce the database server to minimal size, which can allow replicated data from the on-premises database.
- Utilizing the instructions provisioned by the database software, initiate the replication process between the on-premises database server and the Google Cloud database server.

Hot High Availability across On-Premises and Google Cloud for Applications On-premises

If the applications need smaller RTO and RPO values, it can be achieved by implementing high availability across on-premises production environment and Google Cloud. This strategy offers a hot pattern as the on-premises environment and Google Cloud are serving the production traffic. The major difference between the warm pattern and the hot high-availability pattern is that the resources in the on-premises production environment and Google Cloud are in production mode and they both regulate the traffic.

The following are the steps for implementing a hot DR that requires minimal intervention during a disaster:

- Create a VPC network.
- Establish a connection between the on-premises network and the Google Cloud network.
- Create custom images for all the servers on on-premises production environment and they should have similar configurations to the on-premises.
- Configure replication between the Google Cloud database server and the on-premises database server by following the steps given by the database software.
- Now, create different instance templates that utilize images for creating the application servers and the web servers
- Create regional MIGs for the web servers and the application servers.
- Enable health checks with Cloud Monitoring.

Certified Cloud Security Engineer Copyright © by **EC-Council**
All Rights Reserved. Reproduction is Strictly Prohibited.

- Set up load balancing utilizing the configured MIGs.
- Create a scheduled task for capturing regular snapshots of the persistent disks.
- Now, distribute the user requests to Google Cloud environment and the on-premises environment by configuring a DNS service.

This architecture will need a DNS service that allows weighted routing to Google Cloud environment and on-premises environment. The system must also be designed to handle partial failures. In this case, the system must reroute the traffic to the same service in other backup environments. Processes should be enforced to check if the custom images on GCP have similar configurations as those on the on-premises images. As part of the standard lifecycle upgrade, custom images should be upgraded. The cloud administrator must also ensure that the Deployment Manager template is using the latest version of custom images.

In this hot scenario, disaster indicates that one of the environments is not available. The following changes must be performed to handle the disaster:

- If the DNS service fails to reroute traffic based on the health check failures, the cloud security professional should manually reroute the traffic to the system that is running.
- If the database system does not promote read-only replica automatically to the writeable primary on failure, the cloud professional should interfere and ensure that the replica is promoted.

When the other production environment starts working and can handle traffic, the databases must be resynchronized. Once the databases are synchronized, adjust the DNS service such that the production traffic is distributed to both the production environments.

hiu01.ir

Certified Cloud Security Engineer Copyright © by **EC-Council**
All Rights Reserved. Reproduction is Strictly Prohibited.

Disaster Recovery and High-availability Architectures for Applications on Google Cloud

When the application environment is designed on Google Cloud or production workload, the **high-availability features** of the architecture will have a **direct influence** on the DR architecture. The DR and high-availability architectures for applications on Google Cloud include:

Cold Pattern with Recoverable Application Server

- Only **one instance** that writes to disk is required
- Makes use of **high-availability features** in GCP
- **Failover** happens **automatically** in case of a disaster

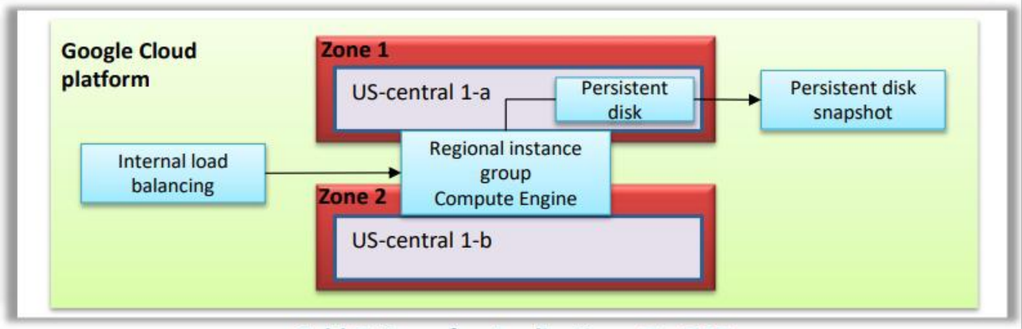

Cold Pattern for Applications On GCP

Warm Pattern Using Static Site Failover

- Handles the service interruption by launching Cloud **Storage-based static site**
- In the production configuration, the cloud **DNS points** to only the primary application
- The DNS points to **static site** when Compute Engine application fails

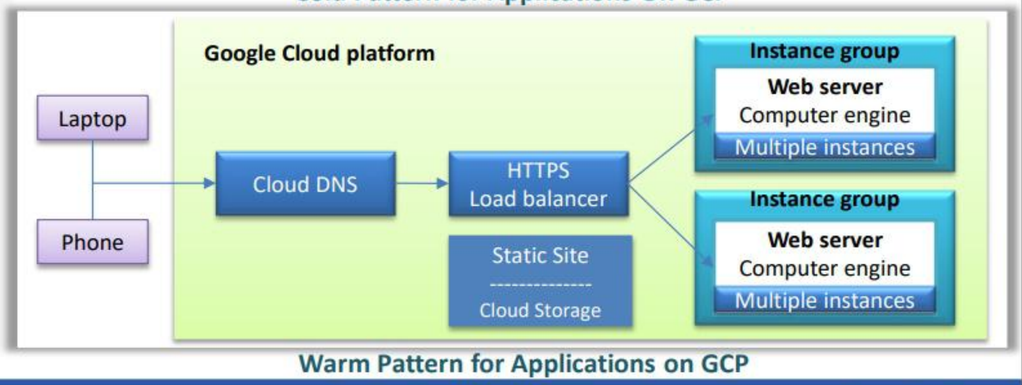

Warm Pattern for Applications on GCP

Copyright © by EC-Council. All Rights Reserved. Reproduction is Strictly Prohibited.

Disaster Recovery and High-availability Architectures for Applications on Google Cloud (Cont'd)

Hot Pattern Using High-availability Web Application

- Implemented by establishing a well-architected **high availability deployment**
- The **failover steps** occur **automatically** during a disaster

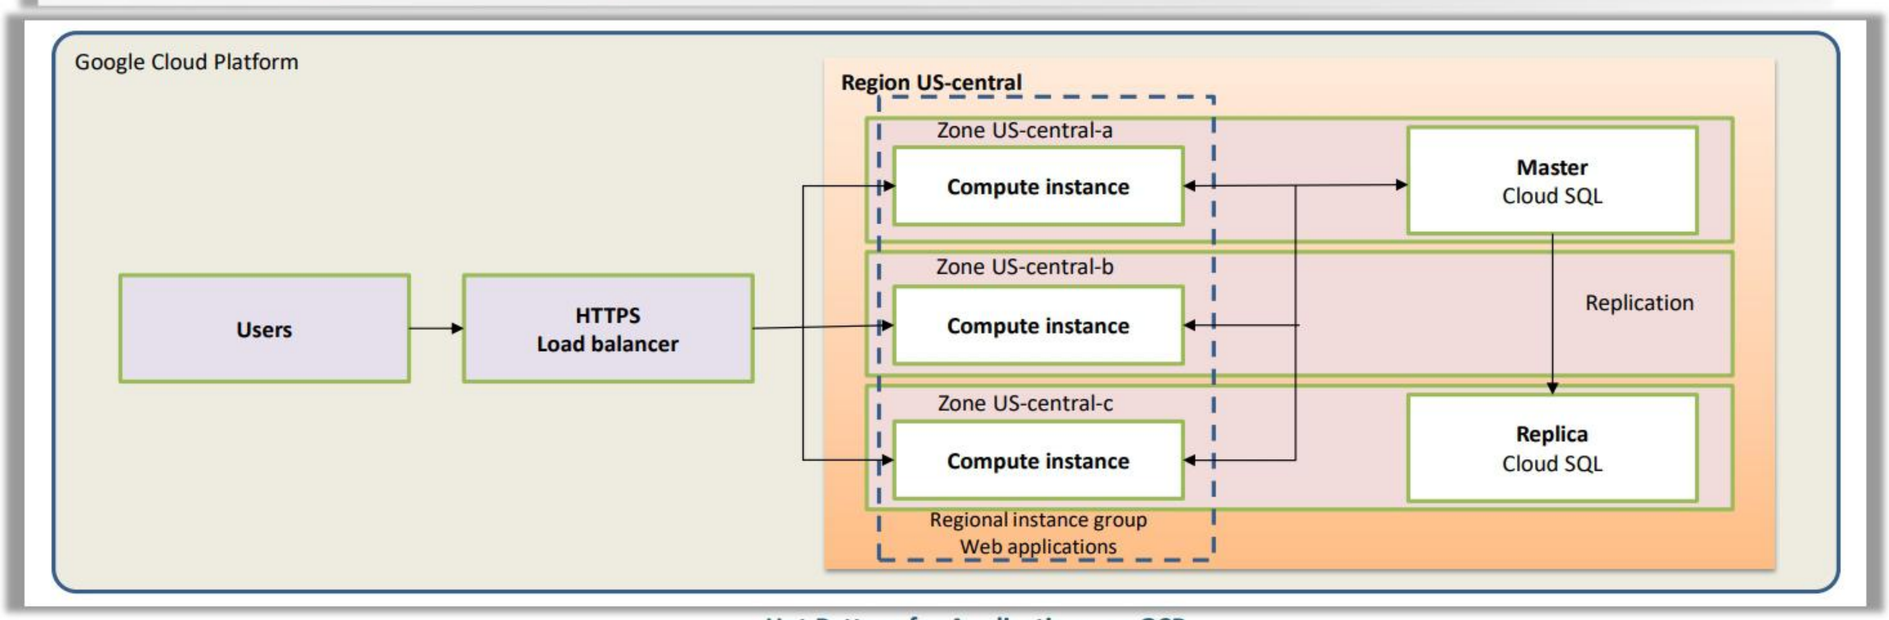

Hot Pattern for Applications on GCP

Copyright © by EC-Council. All Rights Reserved. Reproduction is Strictly Prohibited.

Disaster Recovery and High-availability Architectures for Applications on Google Cloud

When the application environment is designed on Google Cloud or production workload, the high-availability features of the architecture will have a direct influence on the DR architecture.

Certified Cloud Security Engineer Copyright © by **EC-Council**
All Rights Reserved. Reproduction is Strictly Prohibited.

Cold Pattern for Recoverable Application Server on Google Cloud

In a cold scenario where only a single server instance is enough, only one instance that writes to disk is required. For on-premises, this can be achieved with the help of an active/passive cluster. When Google Cloud is used as the production environment, the server instance is also a part of the managed instance group. This managed instance group is used as the backend service for internal load balancing. The Compute Engine and internal load balancing can be used as building blocks for this architecture.

The above architecture can be implemented using the following steps:

- Create a VPC network in Google Cloud.
- Create a custom image that should have the same configuration as the application server. While creating the image, configure the following:
    - The server should be configured so that the data processed by the application can be moved to a persistent disk.
    - Create a snapshot of the persistent disk.
    - Utilizing the updated snapshot, generate a persistent disk and mount the disk by running the startup script.
- Create an instance template that utilizes the image.
- Using the above instance template, configure a managed instance group with the target size as 1.
- Setup a managed instance group with target size 1 by utilizing the instance template.
- Enable health checks utilizing Cloud Monitoring.
- Using the regional MIGs, configure internal load balancing.
- Take regular snapshots of the persistent disks by configuring a scheduled task.

This architecture makes use of high-availability features of Google Cloud. When a disaster strikes, the failover steps are automatically performed, and the cloud security engineer does not have to initiate any action. The same IP address will be utilized for application server, and it is ensured by the internal load balancer. The RPO value is determined by the last captured snapshot. The smaller RPO value can be achieved by taking snapshots more frequently.

In case of a failure, the regional instance group is responsible for launching a replacement instance in another zone within the same region. Using the latest snapshot, a new persistent disk is created and is attached to the new instance.

Warm Pattern for Applications on Google Cloud Using Static Site Failover

In case of instance failure, a cloud-based static site can be launched to handle the service interruption. If web application is static, then a static site serves as an option. In this situation, the primary application runs on instances. These instances are grouped as MIGs, which act as backend services for the HTTP load balancer. The HTTP load balancer is responsible for directing the incoming requests to the instances.

Certified Cloud Security Engineer Copyright © by **EC-Council**
All Rights Reserved. Reproduction is Strictly Prohibited.

The following are the steps for configuring the warm pattern:

- Create a VPC network.
- Create a custom image using the configuration of the application web service.
- Create an instance template for the web servers that use the image.
- Set up a managed instance group for the web servers.
- Using Cloud Monitoring, configure health checks.
- Using the configured managed instance group, configure load balancing.
- Now, create cloud storage based static website.

In the production configuration, the cloud DNS points to the primary application and the static website remain inactive. If the Compute Engine application fails or becomes down, the DNS should be configured to point to this static website.

When the application Compute Engine instances run and can handle production traffic, the recovery steps can be reversed. The cloud DNS should be pointed to the load balancer.

Hot Pattern for Applications on Google Cloud Using High-availability Web Application

The hot pattern for applications running on Google Cloud can be implemented by establishing a well-architected HA deployment. The high-availability deployment can be achieved using Compute Engine load balancing and Cloud SQL.

hide01.ir

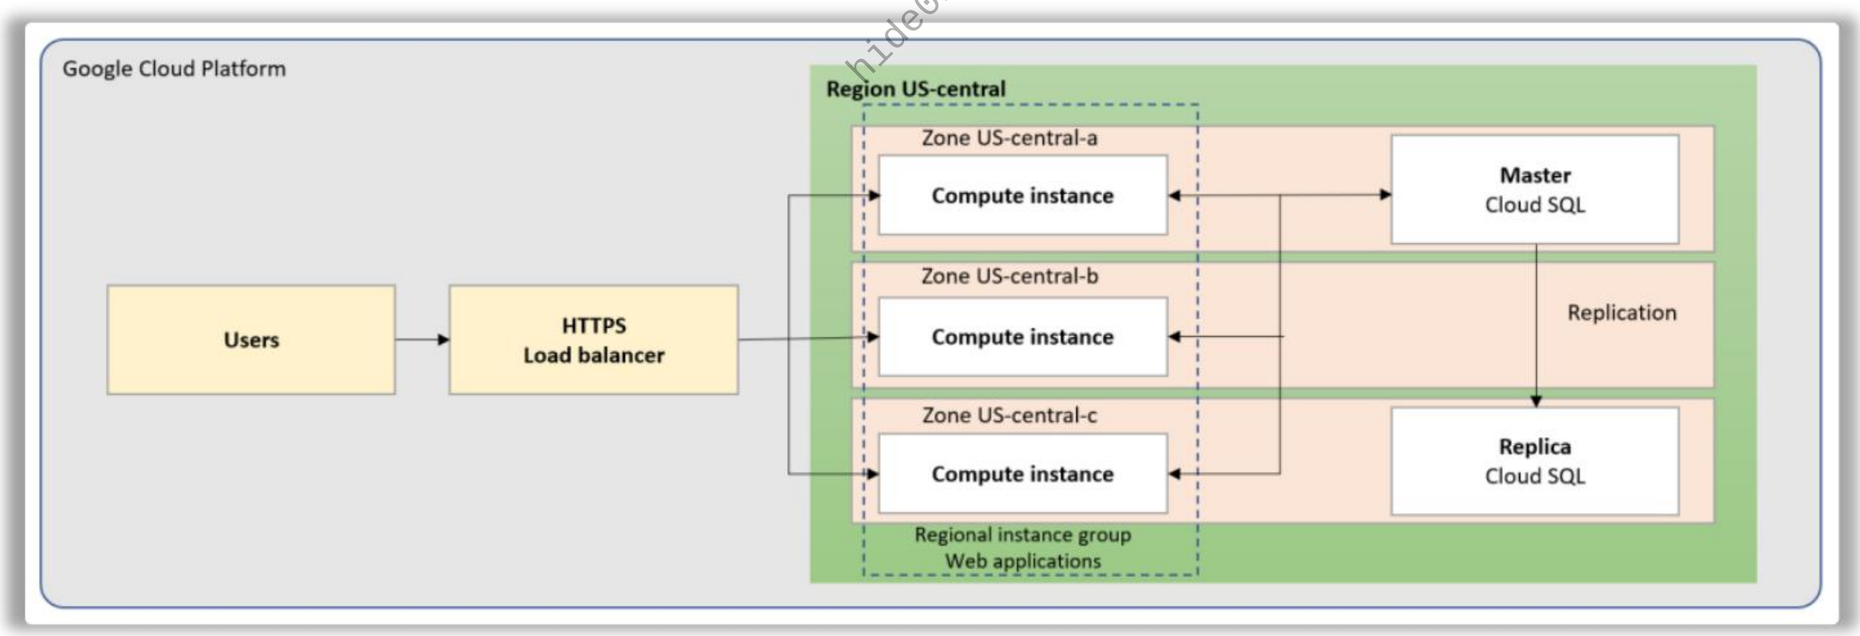

Figure 9.190: Hot Pattern for Application on GCP

The above architecture makes use of the high-availability features of Google Cloud. The failover steps occur automatically in case of a disaster. The instances are distributed across zones by using a regional managed instance group along with the global load balancing and Cloud SQL. The regional managed instance group handles failures at application and zonal levels.

To handle recovery at the application level, while configuring MIGs, HTTP health checks should be configured to check if the services are working on the instances. If the health check determines that the service on an instance has failed, managed instance group recreates that instance automatically.

Certified Cloud Security Engineer Copyright © by **EC-Council**
All Rights Reserved. Reproduction is Strictly Prohibited.

Architecting DR for Locality Restricted Workloads

- Some industries need to adhere to locality-specific requirements such as:
  - Restricting data at rest to a specific country
  - **Processing** of data should take place in the country where it resides
  - Workloads should be accessible only from **predefined locations**
  - Data should be **encrypted** only with **customer-managed keys**

DR Plan for Locality-Restricted Workloads

- The method used for designing the DR is based on the **type** of **applications** and the locality restrictions that need to be met. Planning a DR involves 3 steps, namely:
  - Defining the **locality requirements**
  - Identifying the **useful services**
  - Identifying the **regional granularity** for data and applications

Copyright © by EC-Council. All Rights Reserved. Reproduction is Strictly Prohibited.

Architecting DR for Locality Restricted Workloads (Cont'd)

hive01.ir

Google Cloud Products to Meet Locality Restrictions

- GCP offers products with features suitable for locality-based restrictions such as:
  - **Organization polices**
  - Encryption
  - Storage
  - Compute Engine

Using On-premises or Other Cloud as Production Site

- On-premises and other cloud can be used to meet locality restrictions, such as
  - Anthos
  - Google Cloud **VMware Engine**
  - VPC Service Controls

Copyright © by EC-Council. All Rights Reserved. Reproduction is Strictly Prohibited.

Architecting DR for Locality-Restricted Workloads

For certain industries, it is important to adhere to locality-specific requirements such as:

- Restricting DAR to a specific country.
- Processing of data should take place in the country where it resides.
- Workloads should be accessible only from predefined locations.

Certified Cloud Security Engineer Copyright © by **EC-Council**
All Rights Reserved. Reproduction is Strictly Prohibited.

- Data should be encrypted only with customer-managed keys.

Before designing a DR for locality-restricted workloads, it is important to understand the locality terminology in GCP.

GCP offers services across various regions throughout America, Europe, Asia Pacific, and the Middle East. These regions are further divided into zones where the cloud customers can deploy their resources.

There are three different types of resources, as described below.

- **Multi-regional resources:** Multi-regional resources are designed to be redundant and distributed across various regions. These resources are useful for resilience to a region failure.
- **Regional resources:** Regional resources are redundant across multiple zones in a region. These resources can be used for resilience to a failure of a zone within a region.
- **Zonal resources:** Zonal resources are operated in a single zone. The zonal resources become unavailable when a zone is not reachable or unavailable.

DR Plan for Locality-Restricted Workloads

The method used for designing DR is based on the type of applications and the locality restrictions that need to be met. All these factors have a direct influence on the DR architecture. Planning a DR involves three steps, namely:

hide01.ir

- Defining the locality requirements.
- Identifying the useful services.
- Identifying the regional granularity for data and applications.

Defining the Locality Requirements

Before initiating the design process, define the locality requirements based on the following.

- **DAR:** The location of the DAR decides what services, HA, and DR methods can be used to achieve the desired RTO and RPO values.
- **Encryption techniques:** If the locality requirements can be met by employing encryption methods such as Cloud key management service, multi-regional and dual regional services can be used.
- **Data processing:** If the data can be processed outside the location where it exists, Anthos can be used. Anthos provides a hybrid environment to implement product-specific controls in multiple zones within a region. VPC security controls can be configured to restrict the data access and to restrict the resources that process the data.
- **Data access controls:** If it is required to restrict who can access the applications, various products and services can be used for restricting the access. Identity aware proxy (IAP) verifies a user id and determines if the user can be allowed to access the application.

Certified Cloud Security Engineer Copyright © by **EC-Council**
All Rights Reserved. Reproduction is Strictly Prohibited.

Identify the Useful Services

- Based on the locality and regional granularity requirements, identify the services that can be used for the application. Designing the applications as per the locality requirements needs a proper understanding of which products can be restricted to what regions. It is also important to consider what controls can be applied to meet location restriction requirements.

Identify the Regional Granularity for Data and Applications

Identify the regional granularity for the data and applications by considering the following:

- **Use of multi-regional services:** By using multi-regional services in the design, high-availability resilient architectures can be created.
- **Location restriction for accessing the application:** If access to the application has location restrictions, Google Cloud products such as VPC Services Controls and Google Cloud Armor can be used.
- **Restricting DAR to a specific location:** While using managed services, ensure that the services used can be configured such that the data can be restricted to a specific region.
- **Regional restrictions to applications:** Certain Google products do not offer regional restrictions. Identify the regions where the application can be used and any features that can be used for restricting to a specific region.

Google Cloud Products to Meet Locality Restrictions

Google Cloud offers several products with various features and techniques for DR strategy to meet locality-based restrictions.

- **Organizational Policies:** Organizational policies help to have centralized control over Google Cloud resources. The following constraints can be used while architecting for locality-restricted workloads:
    - **Domain restricted sharing:** This constraint allows only the identities in the allowed list to be added to IAM policies.
    - **Location restricted resources:** This constraint indicates a set of locations where location-based GCP resources can be created.
    - **Encryption:** If the locality-based restrictions need to restrict access to the data, encryption methods need to be implemented. GCP offers products to use the managed keys such as:
        - **Cloud External Key Manager (Cloud EKM):** Cloud EKM can be used for encrypting the data in Compute Engine and BigQuery. Encryption is performed using encryption keys stored in third-party key management system which is deployed outside the GCP.
        - **Customer-supplied Encryption Keys (CSEK)**: CSEK can be used with Compute Engine and cloud storage. These keys are used for encrypting the Google-generated keys for data encryption and decryption.

Certified Cloud Security Engineer Copyright © by **EC-Council**
All Rights Reserved. Reproduction is Strictly Prohibited.

The cloud security professional must ensure to implement proper high availability and recovery processes for the external key manager.

- **Storage:** When designing the DR for locality-restricted workloads, the cloud security professional must ensure that the DAR is located at the required region. Google Cloud Storage and Filestore can be configured to meet the requirements.

  Cloud Storage offers geo-redundancy for locality-restricted workloads. Geo-redundancy allows to store data redundantly in at least two geographical locations that are at least 100 miles away. It ensures maximum availability of data even in case of large disruptions. Objects in the multi-region and dual region are geo-redundant.

  Filestore offers managed file storage services that can be deployed in various regions and zones to meet locality-based restrictions.

- **Compute Engine:** The Compute Engine resources can be global, regional, or zonal. Regional resources can be used by all the resources in that region irrespective of the zone. Whereas zonal resources can be used only by the resources in that zone.

  When resources are placed in multiple zones in a region, resources are isolated from physical infrastructure failures. Keeping resources in different regions also offers high availability in case of a disaster. Compute Engine helps design a robust system with resources located across different failure domains.

Using On-premises or Other Cloud as Production Site

Cloud customers may use a Google location that does not allow dual or multi-regional combination for DR. For this, on-premises or other cloud can be used as production or failover sites to meet locality restrictions.

There are various Google Cloud products that are optimized for hybrid workloads such as:

- **Anthos:** It is a multi-cloud application platform that allows running container-based workloads anywhere in the world. For DR strategy, Anthos allows to run applications on-premises. Whenever a failure or disaster occurs, a cloud security professional can failover to run the workloads on the Anthos cluster in Google Cloud.
- **Google Cloud VMware Engine:** It helps to run workloads in the cloud. If the on-premises workloads are VMware, the DR architecture can be designed to run on the same virtualization solution as the on-premises. The region can be selected as per the locality requirements.
- **VPC Service Controls:** The cloud professional must ensure that the security controls applied to the production environment also extend to the failover environment. VPC Service Controls can be used to define security perimeter from on-premises networks to the projects on Google Cloud. Granular access control policies can be created in GCP based on the attributes. These polices help ensure that proper security controls are applied to on-premises and GCP environment.

Certified Cloud Security Engineer Copyright © by **EC-Council**
All Rights Reserved. Reproduction is Strictly Prohibited.

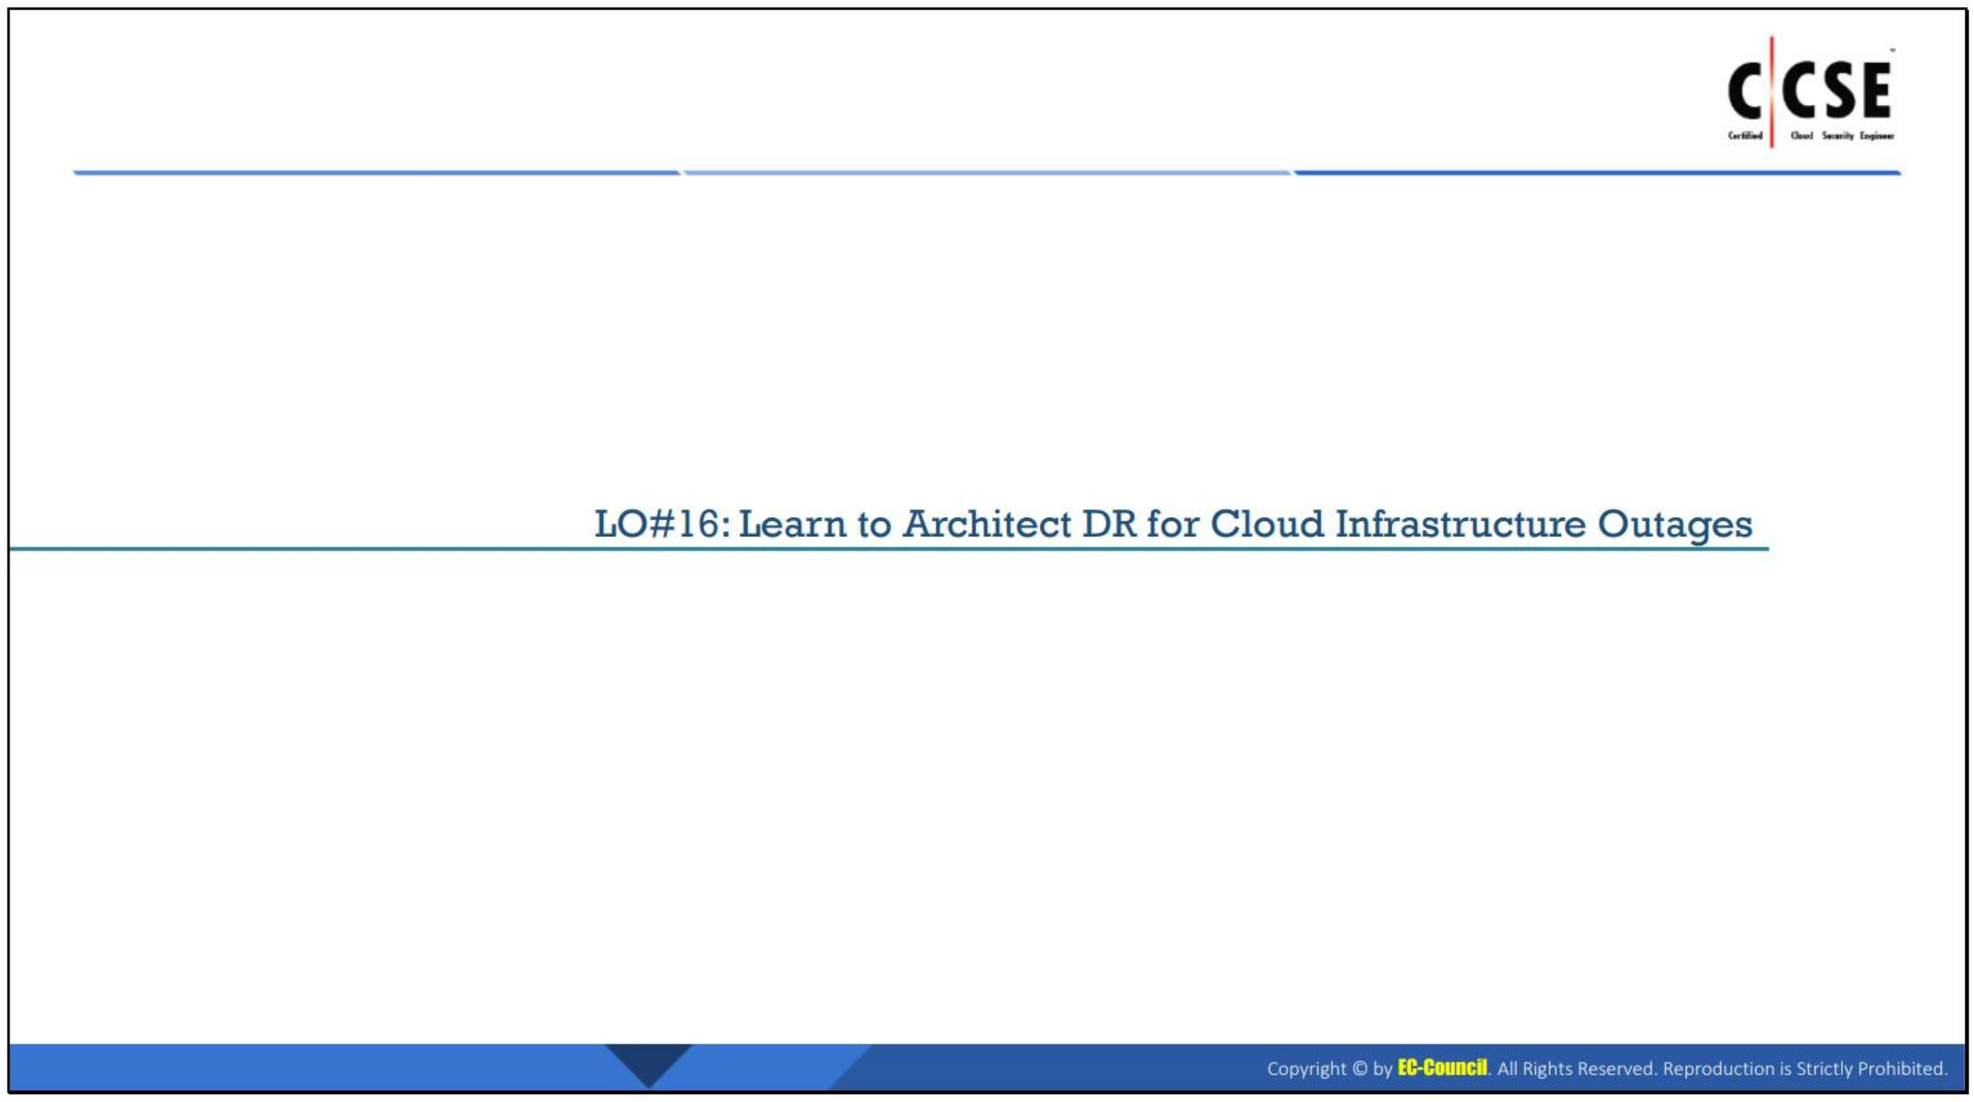

LO#16: Learn to Architect DR for Cloud Infrastructure Outages

The objective of this section is to understand how to design applications to perform well during outages.

hide01.ir

Certified Cloud Security Engineer Copyright © by **EC-Council**
All Rights Reserved. Reproduction is Strictly Prohibited.

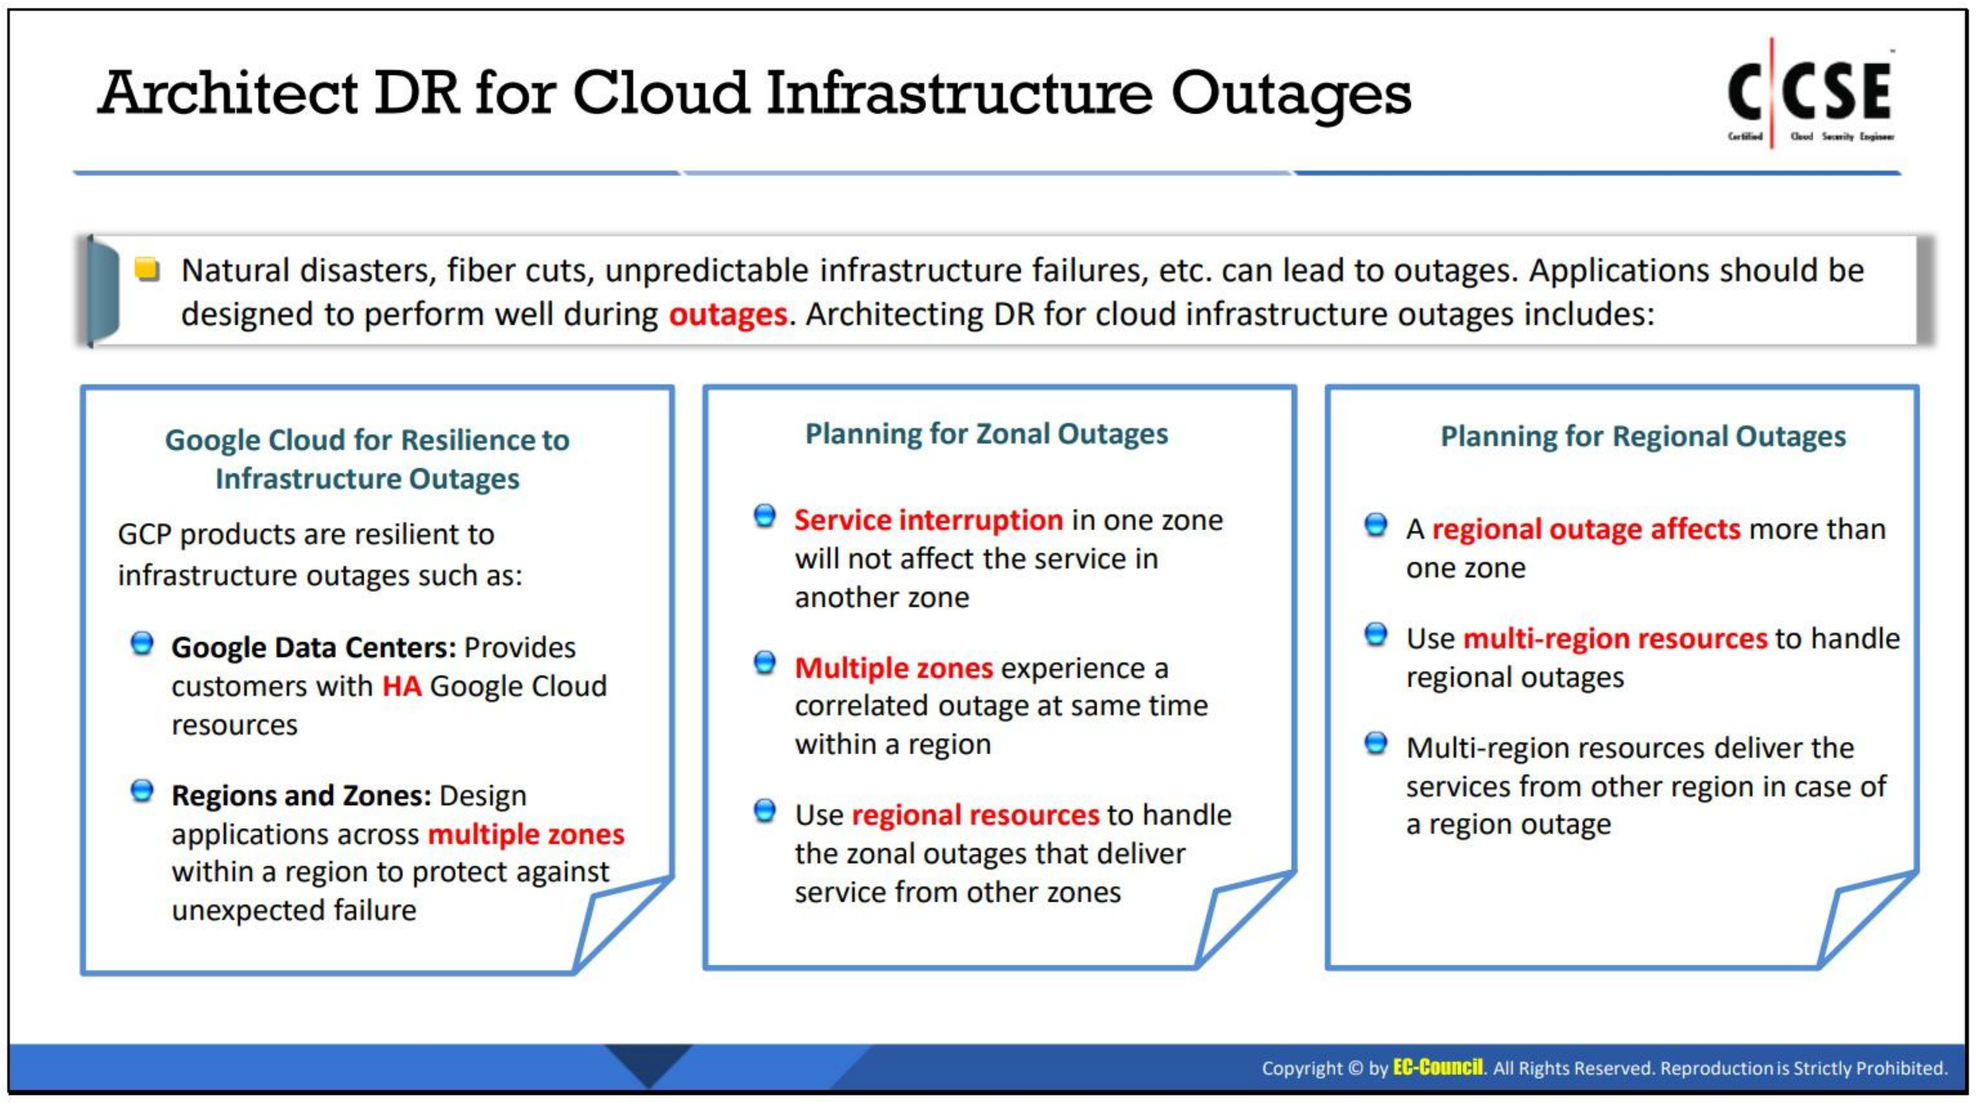

Architecting DR for Cloud Infrastructure Outages

Although Google Cloud offers reliable services, disasters such as natural disasters, fiber cuts, and unpredictable infrastructure failures may occur. These disasters could lead to outages. It is important to be planned for Google Cloud outages. Cloud customers need to design applications such that they perform well during outages using DR features of Google Cloud products.

It is essential to understand how Google Cloud is designed to reduce the frequency of outages and how to design applications that are resilient for Google Cloud outages.

Google Cloud for Resilience

- **Google Data Centers**: Google infrastructure is designed to meet availability targets. It is the leader in designing data centers. Each data center technology in GCP has its own redundancies. Data centers of Google Cloud are designed to provide customers with high-availability Google Cloud resources. Google uses its experience to design the availability of physical and logical system architecture. Google is responsible for handling the risks associated with data centers. This makes cloud customers free from design and operational responsibilities.
- **Regions and Zones**: GCP resources are available across multiple regions. Regions are the geographic locations that contain zones. A zone refers to the deployment area for GCP resources that consists of computing resources.

  Applications can be designed across multiple zones within a region to protect against unexpected failure and to provide HA.

Certified Cloud Security Engineer Copyright © by **EC-Council**
All Rights Reserved. Reproduction is Strictly Prohibited.

Planning for Zonal Outages

Due to the isolation of physical and software in zones, an infrastructure failure may cause outages only in a single zone only. Zones within a region are independent and service interruption in one zone will not affect the service in another zone within the region. An outage at a zone does not necessarily lead to an entire zone outage. A zonal outage may not affect some of the resources in that zone.

In rare cases, multiple zones experience a correlated outage at the same time within a region. In this case, regional outage scope can be applied.

Regional resources are designed to resist the zonal outages by delivering service from other zones. If a zone of the regional resource is interrupted, it makes itself available from other zones.

While using zonal resources such as persistent disks and Compute Engine instances, cloud customers need to perform their own resource composition by building, testing failover, and recovery between zonal resources in multiple zones. Some of the strategies include:

- When a health check determines issues with a zone, traffic can be routed to VMs in another zone using load balancing.
- Instance template and MIGs can be used to create scale similar VM in multiple zones.
- Regional persistent disks can be used to synchronously replicate data to another zone in a region.

ide01.ir

Planning for Regional Outages

A regional outage affects more than one zone in that region. Regional outages are less frequent and are caused by large-scale infrastructure failures and natural disasters. An outage can cause an hour of downtime for a particular product. Thus, critical applications must have a multi-region DR plan.

Multi-region resources can be used to handle region outages. These resources are resistant to region outages by delivering the services from other regions.

If the cloud customers use regional resources that are resilient to regional outages, they need to perform their own resource composition by designing, testing the failover, and recovery between regional resources in multiple regions. The following regional strategies can be applied:

- Regional resources can replicate data to multi-region storage such as cloud storage or using a hybrid-cloud solution such as Anthos.
- Once the regional outage mitigation is in place, it needs to be tested regularly.

Certified Cloud Security Engineer Copyright © by **EC-Council**
All Rights Reserved. Reproduction is Strictly Prohibited.

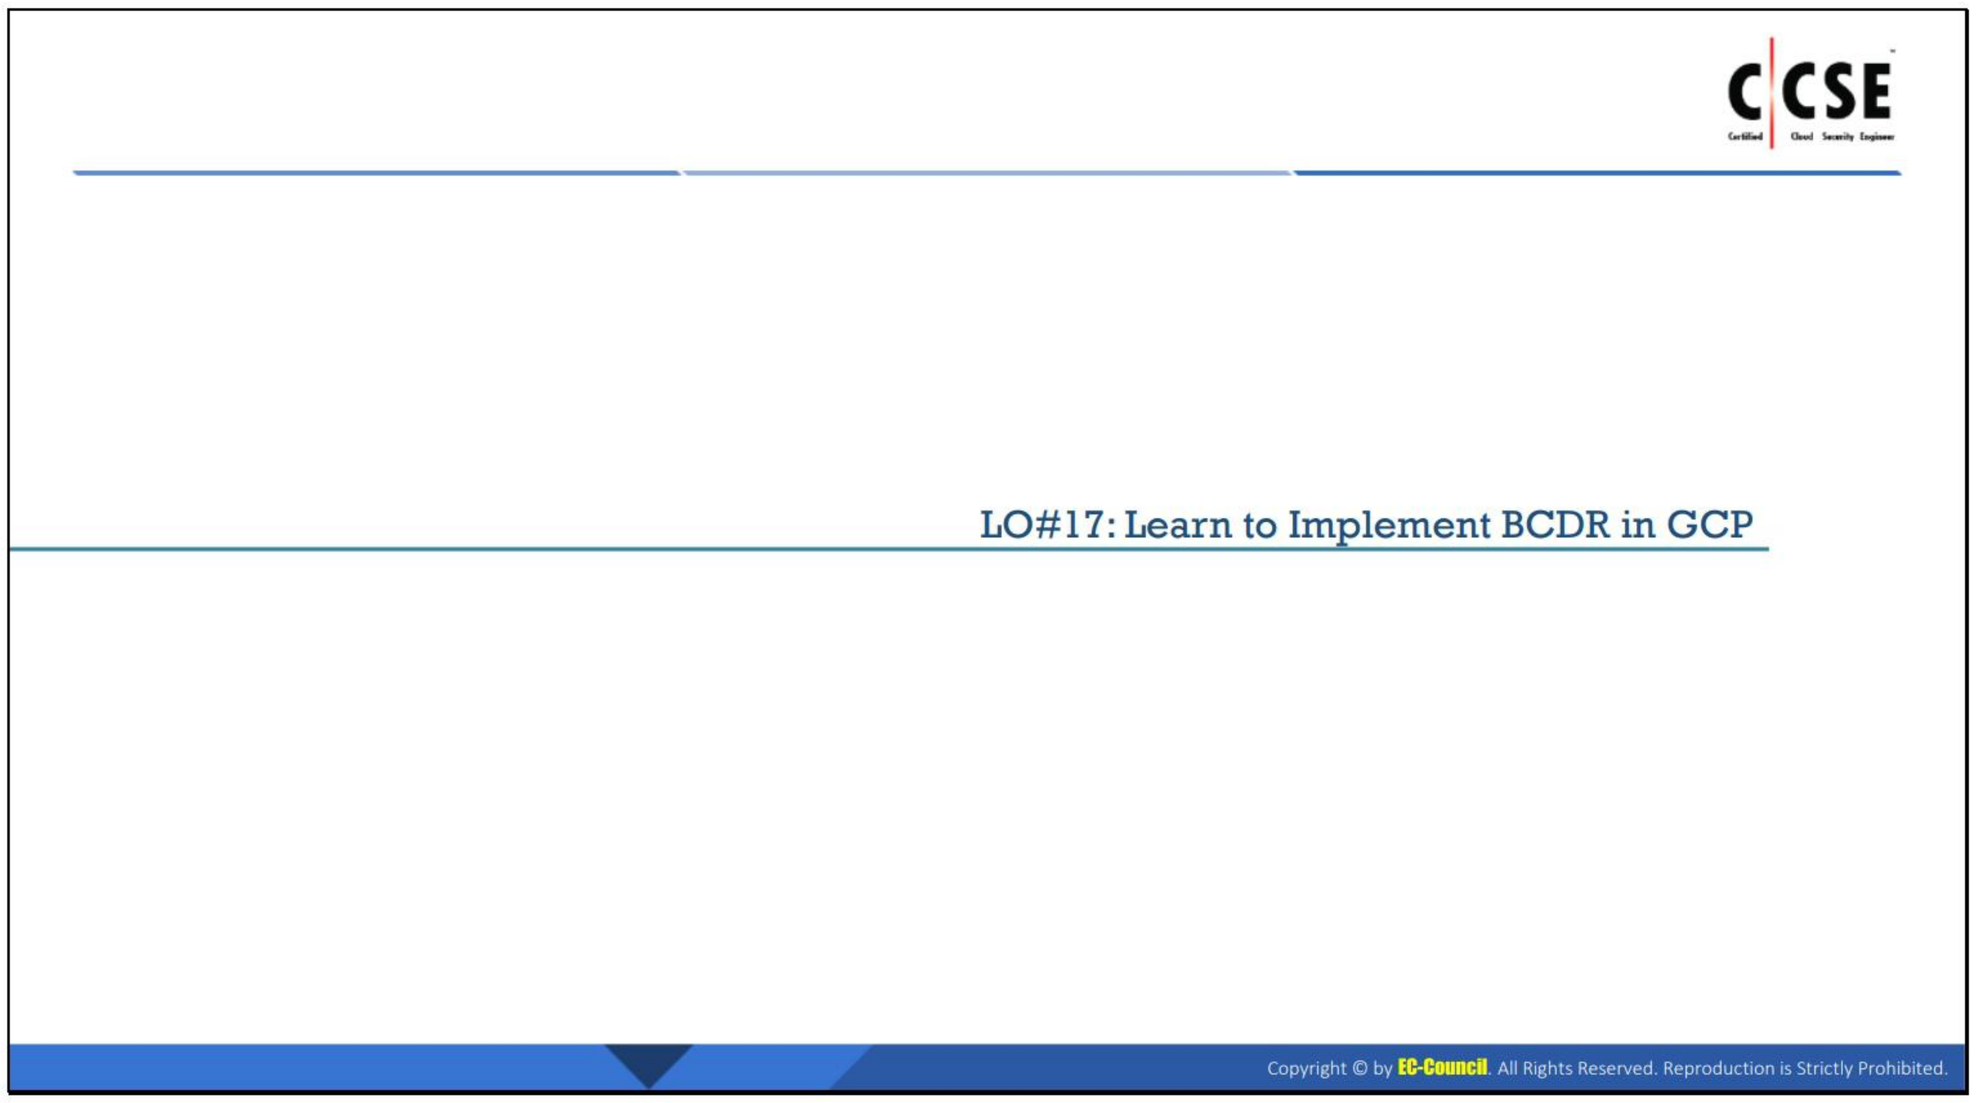

LO#17: Implementing BCDR in GCP

The objective of this section is to understand how to implement BCDR practically in GCP. This section explains the use of various GCP services and tools to implement BCDR. In addition, certain partner tools are discussed.

hide01.ir

Certified Cloud Security Engineer Copyright © by **EC-Council**
All Rights Reserved. Reproduction is Strictly Prohibited.

Enable Point-in-Time Recovery

- Using the point-in-time recovery feature provided by the Google cloud platform, data can be **recovered** to the state before the loss or corruption
- If point-in-time recovery is enabled, then it creates a **new instance**. The newly created instance settings are the replica of the **source instance**

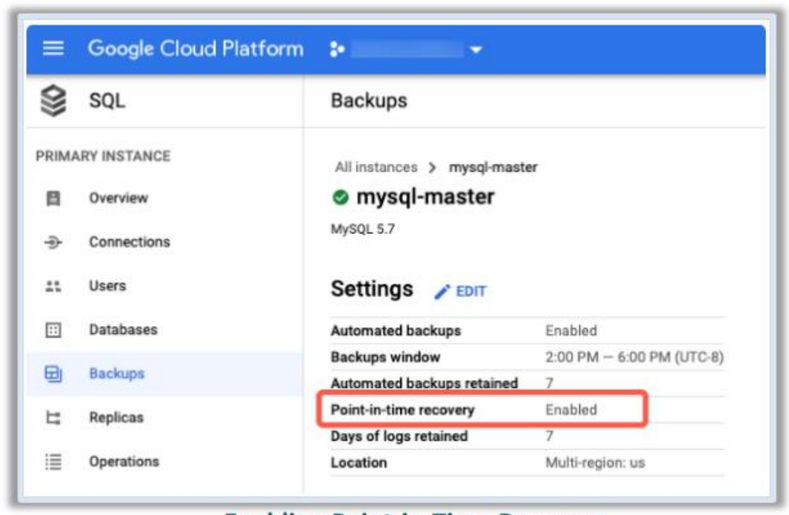

Enabling Point-in-Time Recovery

Copyright © by EC-Council. All Rights Reserved. Reproduction is Strictly Prohibited.

Enable Point-in-Time Recovery in GCP

If there is a data loss due to some fault, then the data can be recovered to the state before the loss or corruption. This is done using the PITR feature provided by the GCP. If the feature is enabled, then it creates a new instance, PITR cannot be performed on an existing instance. The newly created instance settings are the replica of the source instance, which is similar to clone creation works.

Point-in-time Recovery Implementation

By default, when a new Cloud SQL is created, then the PITR is enabled. If this feature is enabled or disabled at a later stage, then the Cloud SQL instance will restart.

Source: https://console.cloud.google.com

Steps to Enable Point-in-time Recovery Using Console

- Sign in to Google Cloud Console and go to the **Cloud SQL Instances** page.
- From the available list, choose the instance that you want to enable point-in-time recovery.
- Click on **Edit**.
- Locate and go to "**Backups, recovery, and high availability"**; then select **Automatic backups**.

Certified Cloud Security Engineer Copyright © by **EC-Council**
All Rights Reserved. Reproduction is Strictly Prohibited.

- Now click on **Enable point-in-time recovery.**

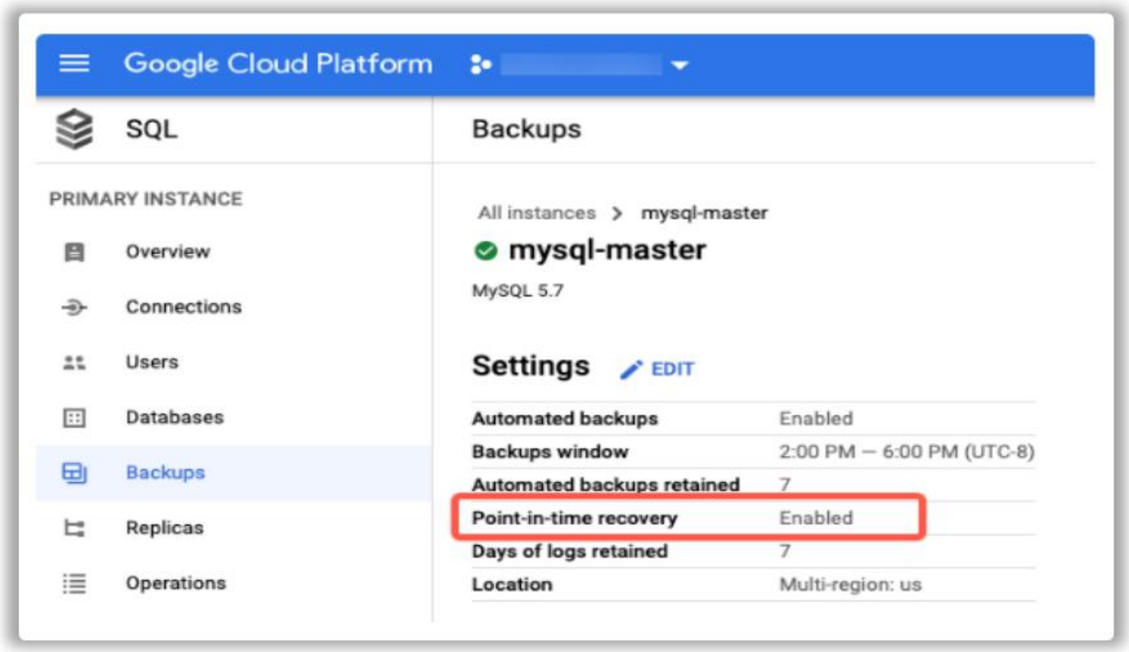

Figure 9.191: SQL Backup Console

- Then click on **Save**.
- Verify that **pointInTimeRecoveryEnabled** is shown as **true** on the **Instance details** page.

Steps to Enable Point-in-time Recovery Using Gcloud

- Use the below command for the instance details to be displayed:

```
gcloud sql instances describe [INSTANCE_NAME]
```

- Under "**backupConfiguration,**" if **enabled: false** is found then enable scheduled backups:

```
gcloud sql instances patch [INSTANCE_NAME] --backup-start-time
[HH:MM]
```

"backup-start-time" needs to be specified using the 24-hour format in UTC±00 time zone.

- Now enable PITR:

```
gcloud sql instances patch [INSTANCE_NAME] --enable-point-in-time-
recovery
```

- Finally, confirm the changes made:

```
gcloud sql instances describe [INSTANCE_NAME]
```

- Under "**backupConfiguration,**" verify **pointInTimeRecoveryEnabled: true**.

Steps to Perform Point-in-recovery Using Console

- Sign in to Google Cloud Console and go to the Cloud SQL Instance page.
- From the available instances, select the instance for which you want to perform recovery; click on the vertical three dots icon and select **CLONE**.
- If required, the new instance name can be updated in **Create a clone** window.
- Select **Clone from the earlier time** option.
- Now enter time for point-in-time recovery.
- Finally, click **Create clone**.

Certified Cloud Security Engineer Copyright © by **EC-Council**
All Rights Reserved. Reproduction is Strictly Prohibited.

Steps to Perform Point-in-recovery Using Gcloud

- Creating PostgreSQL instance clone using PITR:

```
gcloud sql instances clone [SOURCE_INSTANCE_NAME]
[NEW_INSTANCE_NAME] \ --point-in-time [TIMESTAMP]
```

- **TIMESTAMP UTC** is in FRC 3339 format, that is, it looks like **2020-03-21T06:20:00.094z;** this indicates the predefined time for recovery configured by the administrator.

How to Disable Point-in-time Recovery Using Console

- Sign in to Google Cloud Console and go to the Cloud SQL Instances page.
- From the available instances, select the instance that you want to disable point-in-time recovery.
- Click on **Edit**.
- Locate and go to "**Backups, recovery, and high availability"**; then deselect **Enable point-in-time recovery.**
- Click on **Save.**
- Verify that **pointInTimeRecoveryEnabled** is shown as **false** on the **Instance details** page.

How to Disable Point-in-time Recovery Using Gcloud

- To disable point-in-time recovery:

```
gcloud sql instances patch [INSTANCE_NAME] –no-enable-point-in-
                                time-recovery
```

ig01.ir

- Now confirm the changes made:

```
gcloud sql instances describe [INSTANCE_NAME]
```

- Under "**backupConfiguration,"** verify **pointInTimeRecoveryEnabled:**

Certified Cloud Security Engineer Copyright © by **EC-Council**
All Rights Reserved. Reproduction is Strictly Prohibited.

Create Read Replicas in GCP

- A primary instance is the result of a read replica copy; when the original instance becomes unavailable or corrupted, the replica could avoid this **unavailability**. It is recommended to create read replicas in GCP
- The following can be done using read replica:
  - Read request or analytics traffic can be **offloaded** from the primary
  - **Failover** to another region or regional migration can be performed for the reason of disaster recovery

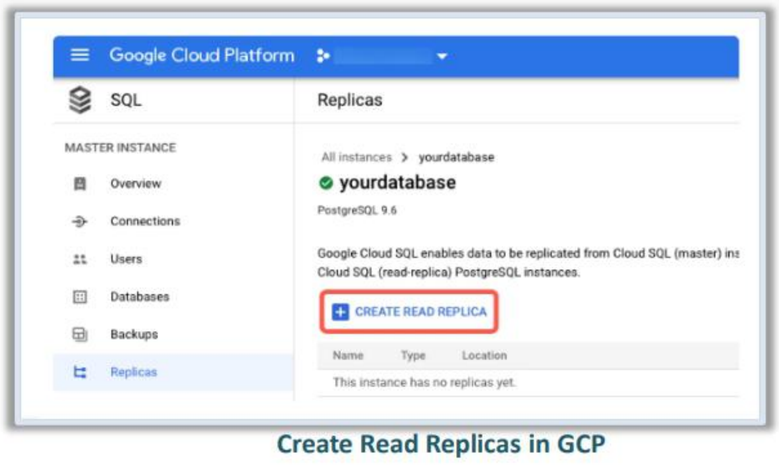

Create Read Replicas in GCP

Copyright © by EC-Council. All Rights Reserved. Reproduction is Strictly Prohibited.

Create Read Replicas

A primary instance is the result of a read replica copy, and primary changes are done in almost real-time. The following can be done using read replica:

- Read request or analytics traffic can be offloaded from the primary.
- Failover to another region or regional migration can be performed for the reason of disaster recovery.

When the original instance becomes unavailable or corrupted, the replica could avoid this unavailability. Taking Cloud SQL instance as a reference, the primary instance is nothing but the replicated instance and the copies are referred to as “read replicas.” All these reside in Cloud SQL. If an on-premises database is taken as reference, the replication situation is defined as replication from an external server. Here, the source database server is the replicated database. The Cloud SQL replicas are the copies of it residing in the Cloud SQL.

Source: https://console.cloud.google.com

To Create a Read Replica Using Console

- Sign in to Google Cloud Console.
- Navigate to the Cloud SQL Instance page.
- Select the instance for which you want to create a replica from the available list, open the icon from the right-side corner for more actions.
- Now select **Create read replica**.
- Then click on **Create**.

Certified Cloud Security Engineer Copyright © by **EC-Council**
All Rights Reserved. Reproduction is Strictly Prohibited.

Read Replica Creation Using Gcloud

- To create the replica, enter the below line.

```
gcloud sql instances create [REPLICA_NAME] --master-instance-
                    name=[MASTER_INSTANCE_NAME]
```

IAM Database Authentication for "Read Replica Configuration"

If the "cloudsql.iam_authentication" flag is enabled in the primary instance, then it is not automatically enabled in the read replica.

The following steps must be followed to configure IAM database authentication in a read replica:

- Log in to the Google Cloud Console.
- Navigate to the Cloud SQL Instance page.
- Click on the required primary instance.
- Search for the "cloudsql.iam_authentication" flag in the Configuration tab. It should be enabled only if the flag is on the list; otherwise it is not necessary.
- Next from the left navigation panel, select **Replicas**.
- Select the replica that you want to edit.
- Under **Configuration options**, click on **Flags** to expand.
- Now click on **+ Add**.
- In the flag name enter "cloudsql.iam_authentication" and ensure that the flag is **On**.
- Finally, click the **Save** button.

hide01.ir

Certified Cloud Security Engineer Copyright © by **EC-Council**
All Rights Reserved. Reproduction is Strictly Prohibited.

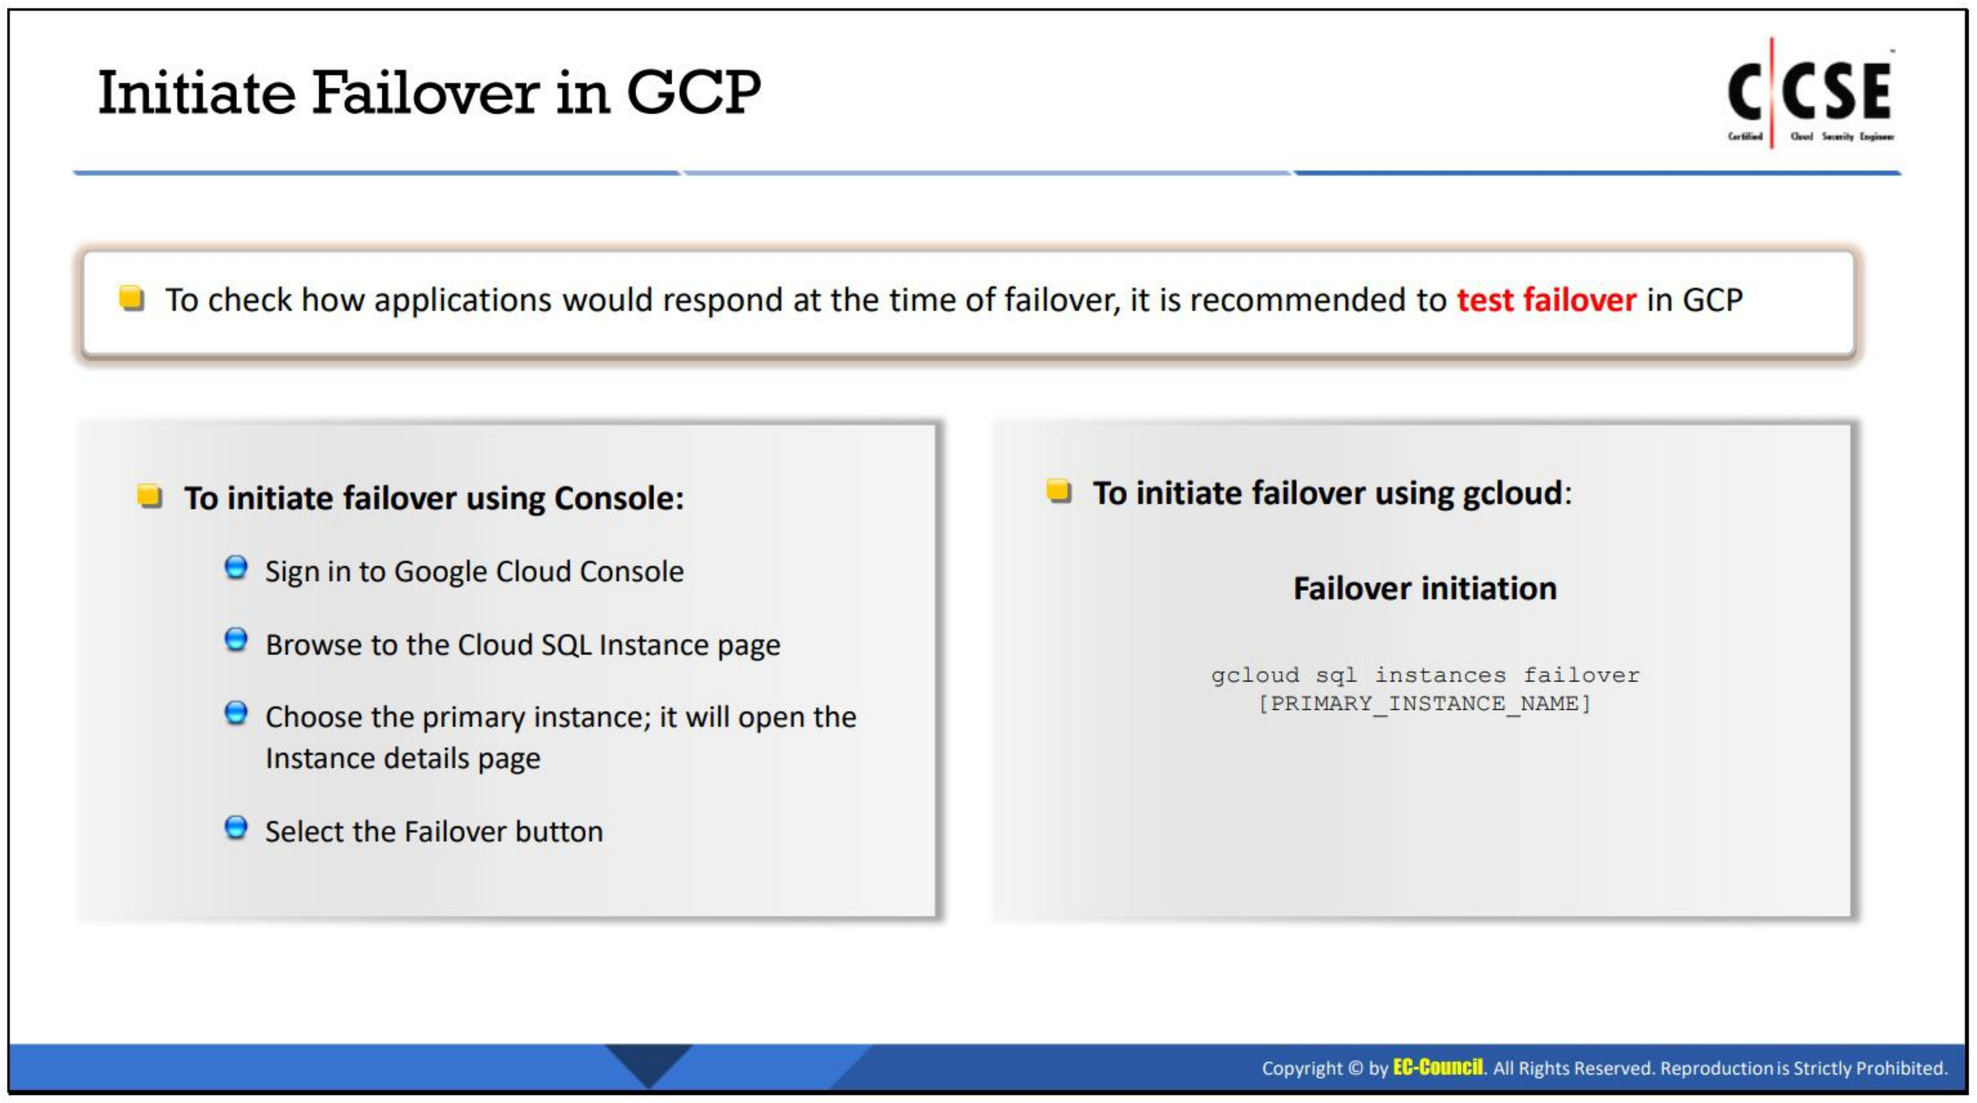

Initiate Failover in GCP

To check how the applications would respond at the time of failover, it is recommended to rest failover in GCP. This is an optional test but it provides a better overview when performed. This feature can be used for testing purposes, which reveals the application response in the event of automatic failover. It is important to ensure that the previous failover on the instance should be completed before performing another failover.

To initiate failover using Console, follow these steps:

- Sign in to Google Cloud Console.
- Navigate to the Cloud SQL Instance page.
- Select the primary instance, then it will open the **Instance details** page.
- Select the **Failover** button.

There might be unavailability of server data for some time.

To initiate failover using gcloud, use the following code:

```
gcloud sql instances failover [PRIMARY_INSTANCE_NAME]
```


Certified Cloud Security Engineer Copyright © by **EC-Council**
All Rights Reserved. Reproduction is Strictly Prohibited.

Enable Replication in Cloud SQL Instances

- The cloud security professional should enable replication of the instances to **migrate** the **data** between regions, platforms, and from on-premises to Cloud SQL for ensuring availability of services when a disaster strikes the primary environment

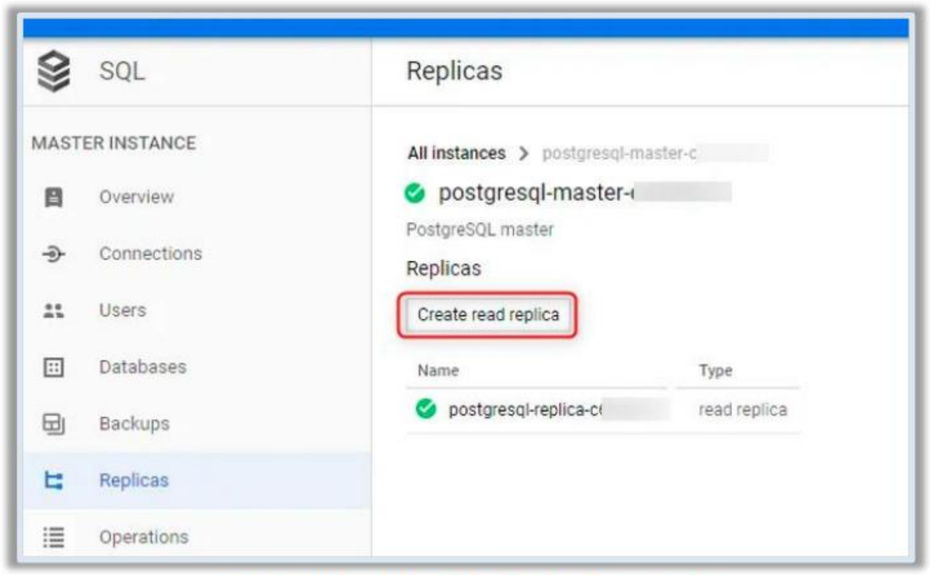

Enabling Replication in Cloud SQL Instances

Copyright © by EC-Council. All Rights Reserved. Reproduction is Strictly Prohibited.

Enable Replication in Cloud SQL instances

Replication helps in migrating the data between regions, between platforms, and from on-premises to Cloud SQL so that the replicated copies can be used for ensuring the availability of services when a disaster strikes the primary environment.

When a Cloud SQL instance is created, it can be configured for HA, which is also called a regional instance, which consists of primary and secondary zones inside that particular region. A failover replica of the primary instance is also created simultaneously during the configuration of high availability so that there will be a primary instance and a standby instance. Synchronous replication makes it possible to replicate all writes made to the primary instance to disks in both primary and secondary zones. During a disaster, failover occurs where the standby instance becomes the primary instance. The original primary instance will be recreated once it becomes available after the disaster and then it becomes the standby instance. When a disaster strikes again, the failover process occurs with the new setup of primary and standby instances.

Cloud SQL supports read replica, cross-region read replica, and external read replica.

- **Read replica**: Read replica is a similar copy of the primary instance and it is read-only. Read replicas can reduce the traffic load on the primary instance by processing the number of queries and requests. Read replicas can be used for regional failover for performing disaster recovery.
- **Cross-region read replica**: Cross-region replicas are read replicas hosted on a different region. They can improve read performance, protect against regional failures, and enable data migration between regions with reduced downtime.

Certified Cloud Security Engineer Copyright © by **EC-Council**
All Rights Reserved. Reproduction is Strictly Prohibited.

- **External read replicas**: The external SQL instances that replicate from a primary instance are external read replicas. The connection latency in external communications can be reduced by using external replicas.

Source: https://console.cloud.google.com

To Enable Replication in Cloud SQL Instances

Step 1: Create read replicas.

With Console

- Navigate to the **Cloud SQL Instances** page in console.
- Open **more actions** menu for the instance to which replica is to be created.
- Click on **Create Read Replica.**

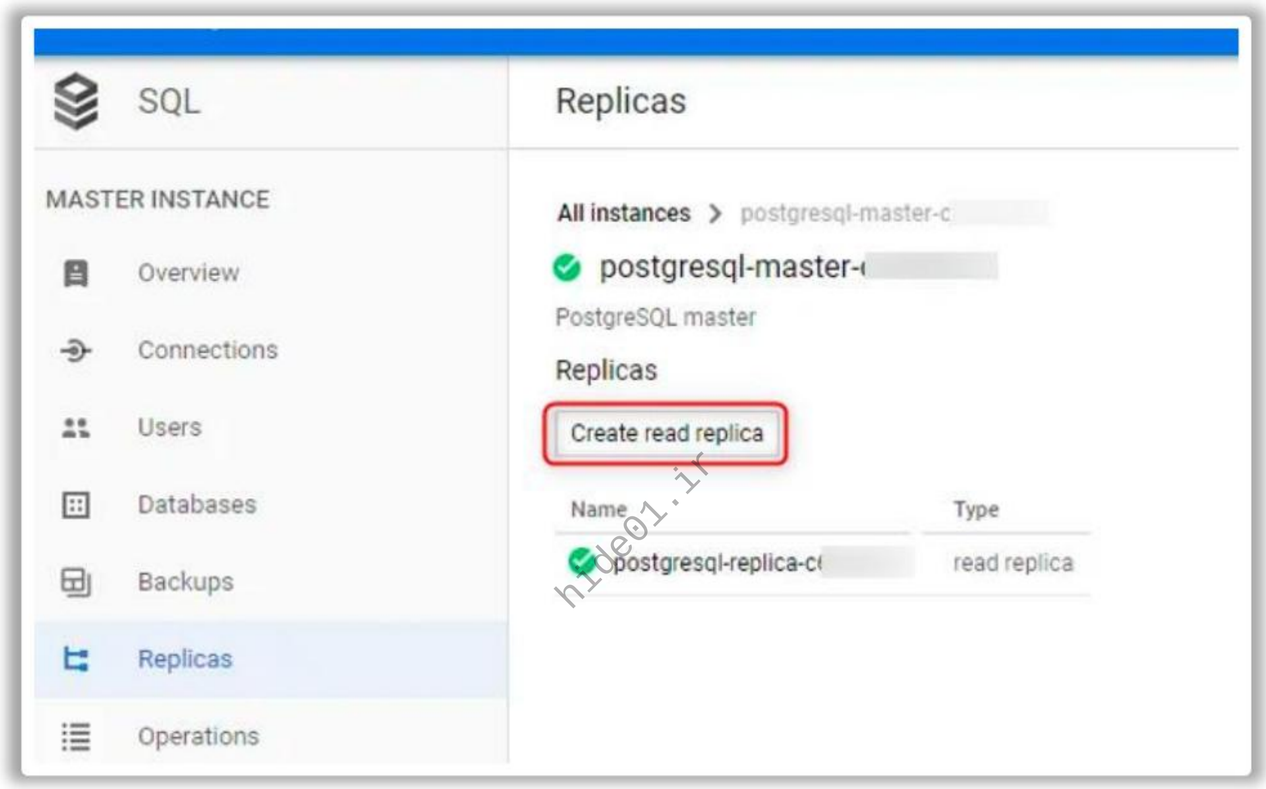

Figure 9.192: Create read replica

- To enable backup and binary logging, click on **Automate Backup** and then click on **Enable Binary Logging.**
- Click **Continue.**
- Click **Save & Restart.**
- Enter the instance id, name, region, and zone in the **Create read replica** page.
- Click on **Create**; this creates the backup and replica.

With gcloud Script

- The status of the primary instance can be checked to confirm that it is not a replica by issuing the following command and if databaseReplicationEnabled is false, it is not a replica.

```
gcloud sql instances describe PRIMARY_INSTANCE_NAME
```

- To enable backups, execute the following command

Certified Cloud Security Engineer Copyright © by **EC-Council**
All Rights Reserved. Reproduction is Strictly Prohibited.

```
gcloud sql instances patch PRIMARY_INSTANCE_NAME --backup-start-time
>HH:MM
```

- To enable binary logs on the primary instance

```
gcloud sql instances patch --enable-bin-log PRIMARY_INSTANCE_NAME
```

- The replica can be created using the following command

```
gcloud sql instances create REPLICA_NAME --master-instance-
name=PRIMARY_INSTANCE_NAME
```

Step 2: Enable high availability for an instance.

With Console

- Navigate to **Cloud SQL Instances** page in Console.
- Click on the instance name to enable HA, which opens the **Overview** page.
- Click on **Edit** and open **Backups, recovery, and high availability**.
- Mark the checkboxes for **Automate backups** and **Enable point in time recovery**.
- Select **High Availability (regional)** and click on **Save.**

With gcloud Script

- HA for the instances can be configured by issuing the following command.

```
gcloud sql instances patch [INSTANCE_NAME] --availability-type
    REGIONAL --enable-bin-log --backup-start-time=[HH:MM]
```

hide01.ir

Step 3: Enable failover for primary instances during disaster.

With Console

- Navigate to the **Cloud SQL Instances page** in Console.
- Click on the name of the primary instance page to which failover is to be initiated to open the **Instance Details** page.
- Click on **Failover.**

With gcloud Script

- To configure failover, issue the following command.

```
gcloud sql instances failover [PRIMARY_INSTANCE_NAME]
```


Certified Cloud Security Engineer Copyright © by **EC-Council**
All Rights Reserved. Reproduction is Strictly Prohibited.

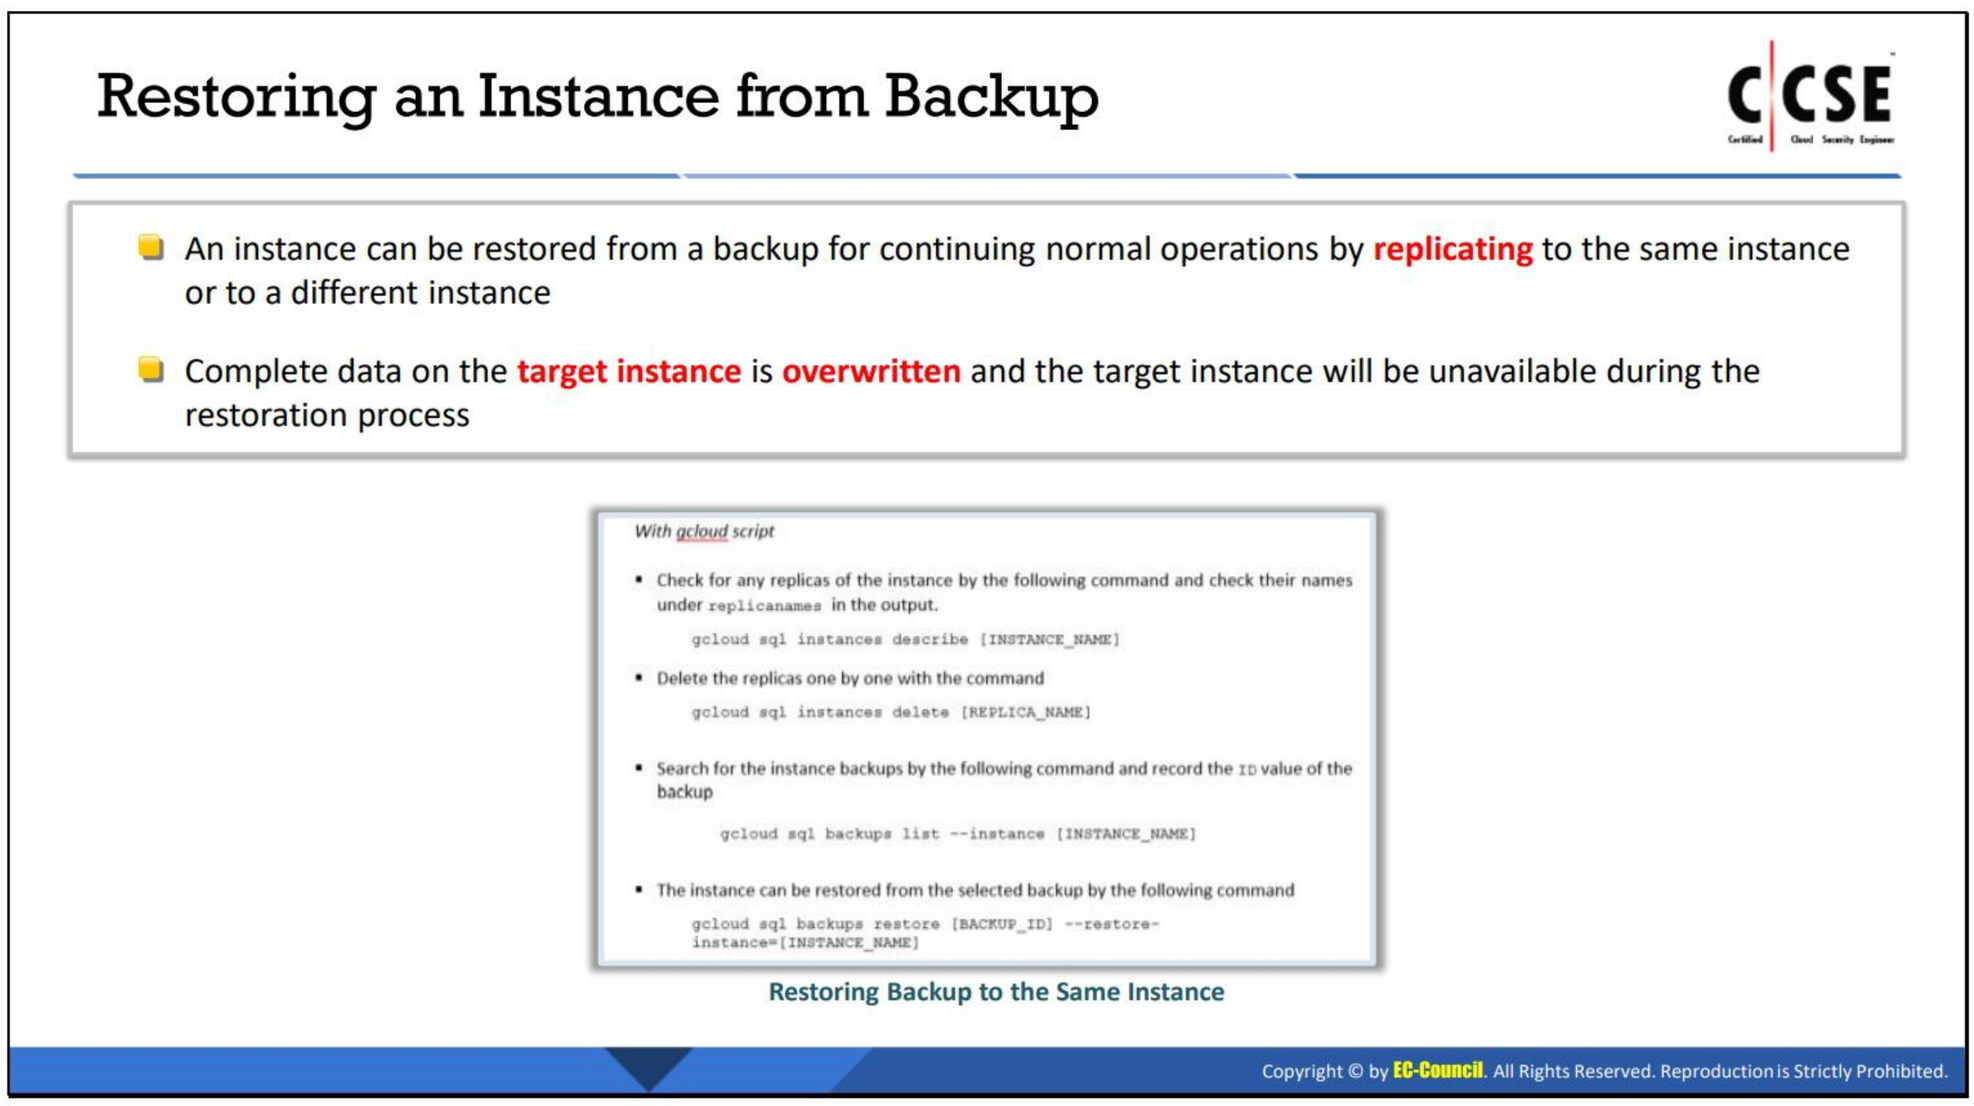

Restore an Instance from Backup

An instance can be restored from a backup for continuing normal operations by replicating to the same instance or a different instance. While performing the restore process from the backup to the same instance, all the data on that instance is returned to the state when the backup was performed. At the time of restoring the process from backup to a different instance, the data on the target instance is updated to the state of the source instance from which the backup was performed. While performing the restore process from a backup, complete data on the target instance is overwritten and the target instance becomes unavailable.

Source: https://console.cloud.google.com

To Restore Backup to the Same Instance

With Console

- Navigate to the **Cloud SQL Instances** page in console.
- Delete any replicas of the target instance by clicking on the **more actions** menu.
- Click on the instance to be restored which opens the **Instance Details** window.
- Click on **Backups**.
- Search for the backup and click **Restore.**
- From the drop-down of the Instance menu in the **Restore an Instance from a backup** window, select the instance to be used for the restore.
- Under **Instance Name**, enter the name of the instance.
- Click **Restore**.

Certified Cloud Security Engineer Copyright © by **EC-Council**
All Rights Reserved. Reproduction is Strictly Prohibited.

- Navigate to the **Operations** page of the instance to check whether the restoration process is complete.
- Restore the replicas deleted initially.

With gcloud Script

- Check for any replicas of the instance using the following command and check their names under replicanames in the output.

```
gcloud sql instances describe [INSTANCE_NAME]
```

- Delete the replicas one by one with the command

```
gcloud sql instances delete [REPLICA_NAME]
```

- Search for the instance backups using the following command and record the ID value of the backup

```
gcloud sql backups list --instance [INSTANCE_NAME]
```

- The instance can be restored from the selected backup using the following command

```
gcloud sql backups restore [BACKUP_ID] --restore-
instance=[INSTANCE_NAME]
```

To Restore from a Backup to Different Instance

With Console

hide01.ir

- Navigate to the **Cloud SQL Instances** page in console.
- Delete any replicas of the target instance by clicking on **more actions** menu.
- Click on the instance to be restored, which opens **Instance Details**.
- Click on **Backups**.
- Search for the backup and click **Restore.**
- From the Instance menu drop-down in the **Restore an Instance from a backup** window, select the instance to which restoration is to be performed.
- From the Instance menu drop-down in the **Restore an Instance from a backup** window, select the instance to be used for the restore process.
- Under **Instance Name**, enter the name of instance.
- Click **Restore**.
- Navigate to **Operations** page of the instance to check whether the restoration process is complete.
- Restore the replicas deleted initially.

With gcloud Script

- Check for any replicas of the target instance using the following command and check their names under replicanames in the output.

Certified Cloud Security Engineer Copyright © by **EC-Council**
All Rights Reserved. Reproduction is Strictly Prohibited.

```
gcloud sql instances describe [TARGET_INSTANCE_NAME]
```

- Delete the replicas one by one with the command

```
gcloud sql instances delete [REPLICA_NAME]
```

- Search for the source instance backups using the following command and record the ID value of the backup

```
gcloud sql backups list --instance [SOURCE_INSTANCE_NAME]
```

- The instance can be restored from the selected backup to the target instance using the following command

```
gcloud sql backups restore [BACKUP_ID] -restore
      instance=[TARGET_INSTANCE_NAME] \
   --backup-instance=[SOURCE_INSTANCE_NAME]
```

hide01.ir

Certified Cloud Security Engineer Copyright © by **EC-Council**
All Rights Reserved. Reproduction is Strictly Prohibited.

Checking the Replication Status using Custom Dashboard

- The replication status and metrics of the replica instances such as the **replication state**, replication lag, slave input/output thread running state and slave SQL thread running state can be viewed by creating a **custom dashboard**

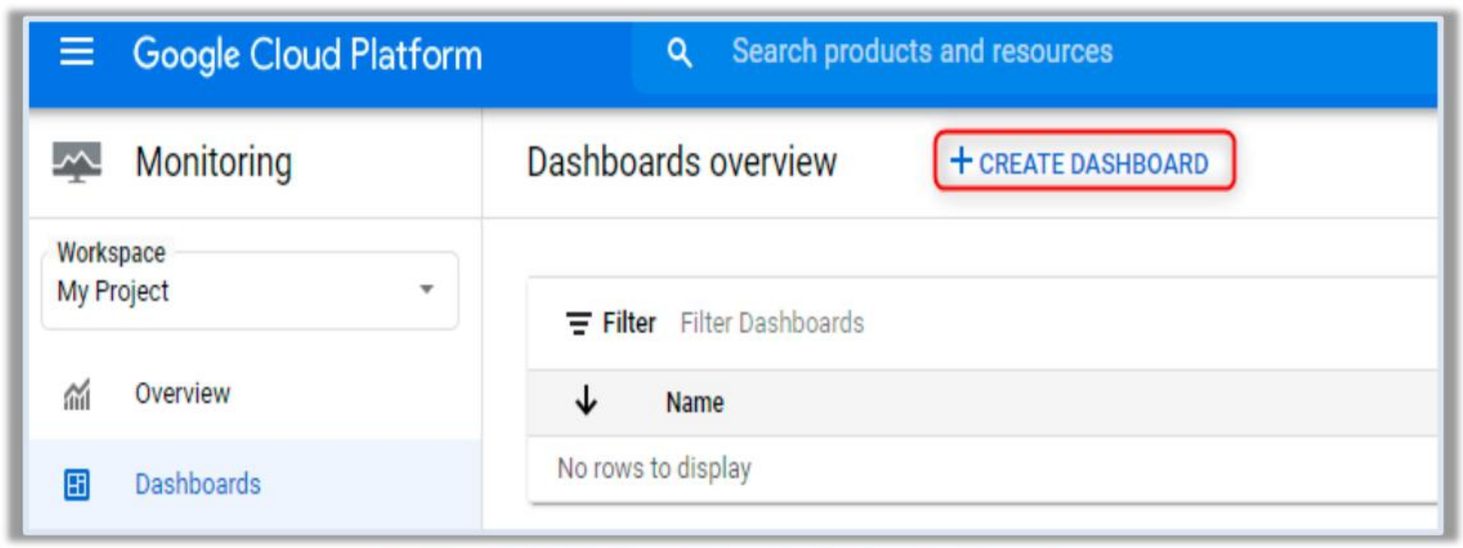

Checking the Replication Status using Custom Dashboard

Copyright © by EC-Council. All Rights Reserved. Reproduction is Strictly Prohibited.

Checking the Replication Status using Custom Dashboard

The replication status and metrics of the replica instances can be viewed from the GCP console. Gcloud command line tool provides a brief insight into the replication status and metrics. Replica instances provide four types of metrics:

- **Replication state**: Replication state indicates the status of replication, that is, whether the primary instance is providing logs to the replicated instance. It has the following values: running, stopped, and error.
- **Replication lag**: Replication lag is the time lag between the replica's state and state of the primary instance. The writes will be counted as lagging even if the replica has received it but not yet applied to the database.
- **Slave I/O thread running state**: Slave I/O thread running state indicates whether the I/O thread used for scanning the binary log of the primary instance runs on the replica. It has the following values: Yes, No, and Connecting.
- **Slave SQL thread running state**: Slave SQL thread running state indicates whether the SQL threads for performing events in relay log runs on the replica. It has the following values: Yes, No, and Connecting.

To Check Replication Status

Step 1: Monitoring replication state and status on Cloud SQL monitoring dashboard.

- Navigate to **Monitoring** in Console.
- Click on **Metrics Explorer**.
- In the **Find resource type and metric** window, enter the name of the monitored resource.

Certified Cloud Security Engineer Copyright © by **EC-Council**
All Rights Reserved. Reproduction is Strictly Prohibited.

- For **Resource Type**, select **Cloud SQL Database.**

Step 2: Creating custom dashboard to view metrics of replicas.

With Console

- Navigate to **Monitoring** in console.

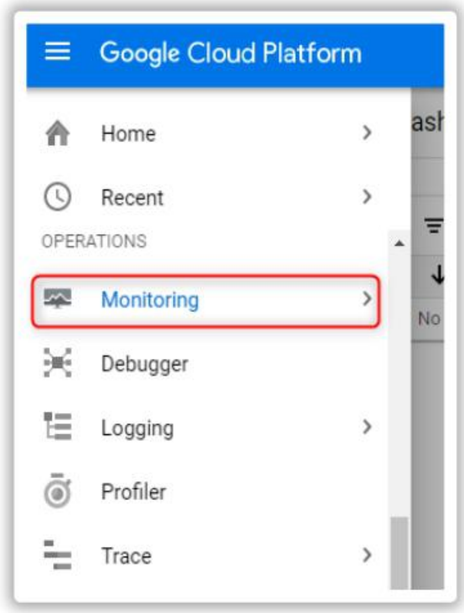

Figure 9.193: Click Monitoring

- Click on **Dashboards.**

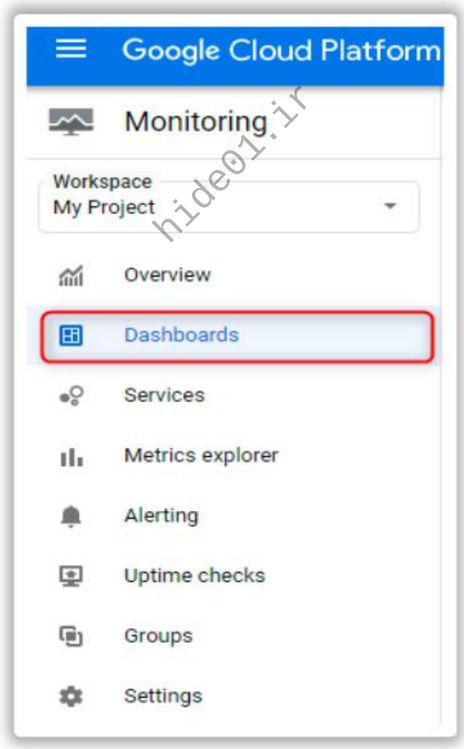

Figure 9.194: Click Dashboard

- Click on **Create Dashboard.**

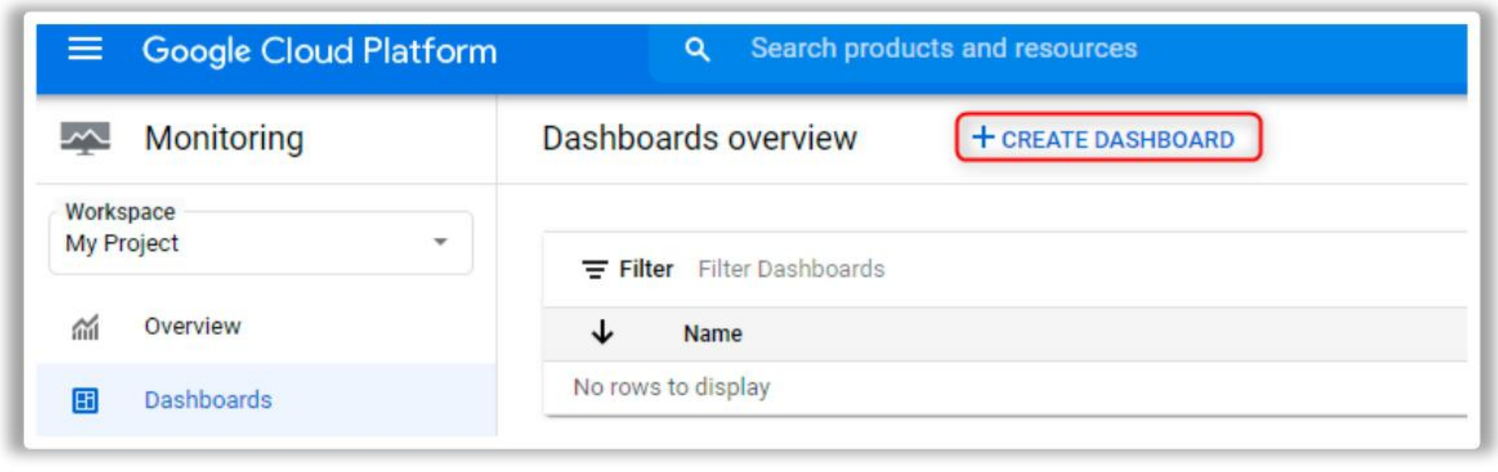

Figure 9.195: Click Create Dashboard

Certified Cloud Security Engineer Copyright © by **EC-Council**
All Rights Reserved. Reproduction is Strictly Prohibited.

- Click on **Add Chart.**

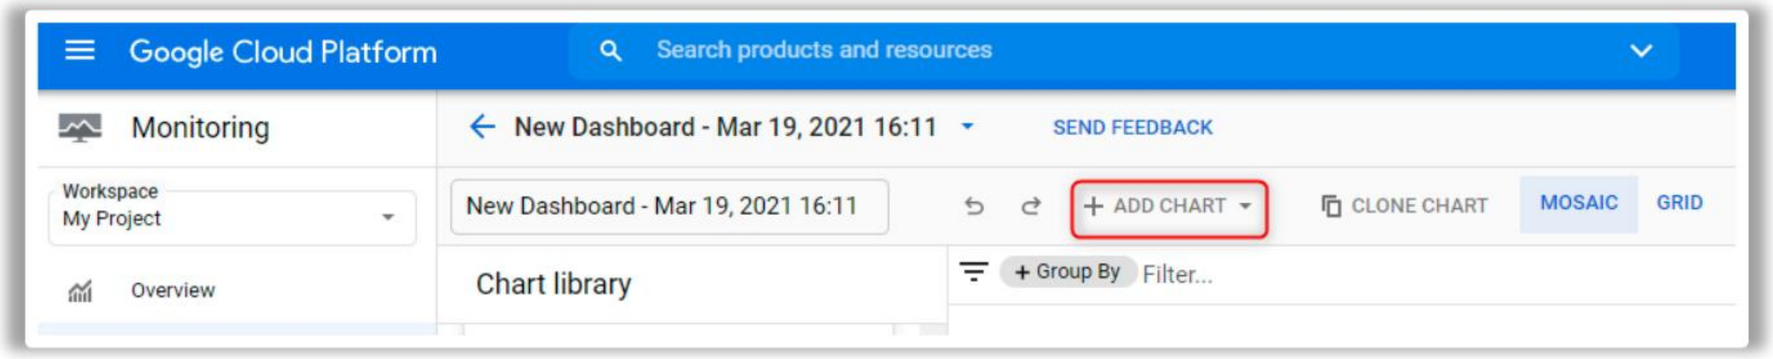

Figure 9.196: Click Add chart

- Select the **Resource type** as **Cloud SQL Database.**
- The custom chart for different metrics can be created by
    - **Monitoring the replication state:** Type **`Replication State`** under **Select a Metric** and add the filter state=" running". If replication is running, 1 will be shown; otherwise, 0 will be shown.
    - **Monitoring the replication lag:** Type **`replica_lag`** under **Select a Metric.**
    - **Monitoring the replica's I/O thread metric:** Type **`Slave I/O thread running state`** under **Select a Metric** and add the filter state="Yes". If the thread is running, 1 will be shown; otherwise, 0 will be shown.
    - **Monitoring the replica's SQL thread metric:** Type **`Slave SQL thread running state`** under **Select a Metric** and add the filter state= "Yes". If the thread is running, 1 will be shown; otherwise, 0 will be shown.

hidedo.ir

With gcloud Script

- The replication status of replicas can be verified with the command below and the properties such as databaseReplicationEnabled and masterInstanceName should be verified in the output

```
gcloud sql instances describe [REPLICA_NAME]
```

- The replicas of primary instances can be checked by issuing the following command and check for replicaNames in output.

```
gcloud sql instances describe [PRIMARY_INSTANCE_NAME]
```


Certified Cloud Security Engineer Copyright © by **EC-Council**
All Rights Reserved. Reproduction is Strictly Prohibited.

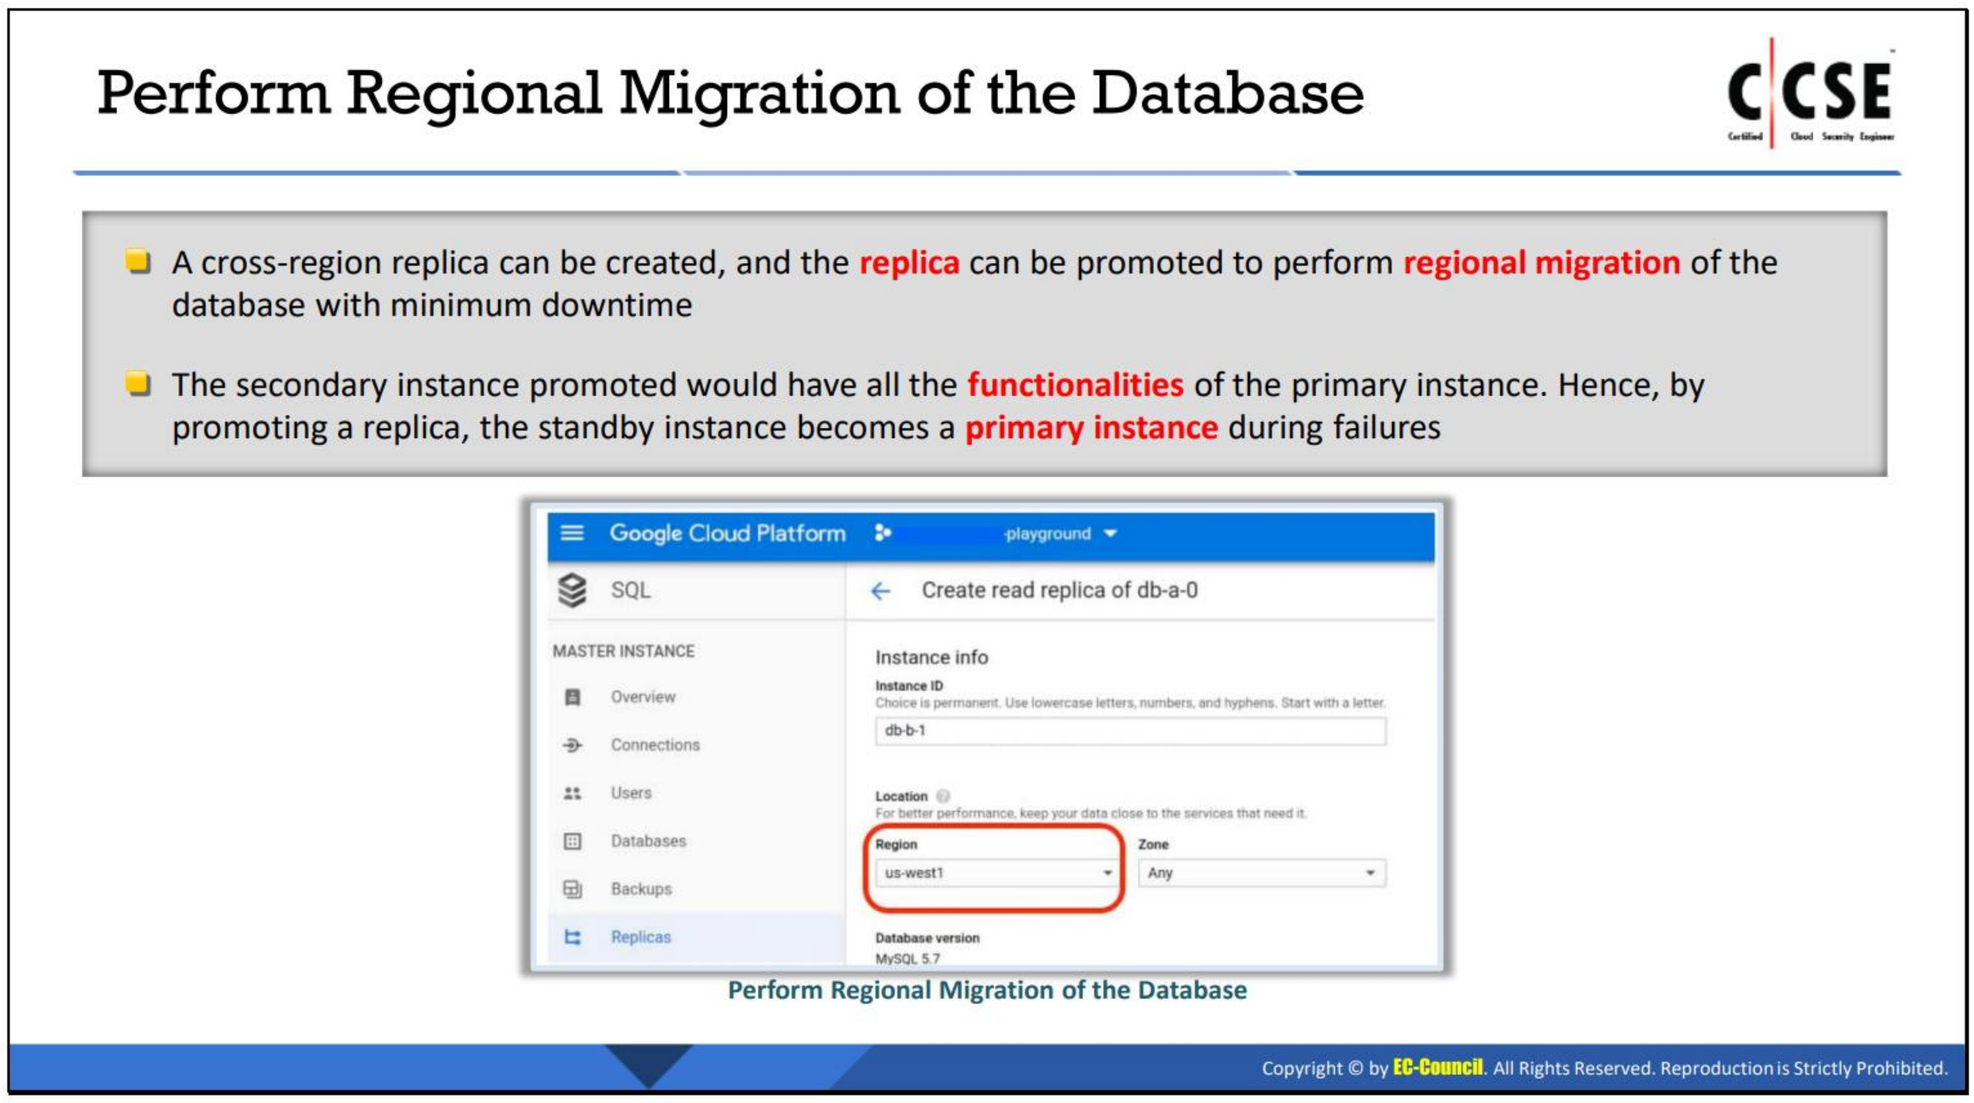

Perform Regional Migration of the Database

A cross-region replica can be used to perform regional migration of the database with reduced downtime. In regional migration the process that occurs is, a replica of the primary instance is created in another region, wait till the replication come up, promote the replica, and then redirecting the users over the new instance. The secondary instance promoted could have all the functionalities of the primary instance. Thus, in general by promoting a replica the standby instance becomes a primary instance during failures.

How to Perform Regional Migration?

Step 1: Create cross-region replicas.

With Console

- Navigate to **Cloud SQL Instances** page in console.
- Open **more actions** menu for the instance to which replica is to be created.
- Click on **Create Read Replica.**

Certified Cloud Security Engineer Copyright © by **EC-Council**
All Rights Reserved. Reproduction is Strictly Prohibited.

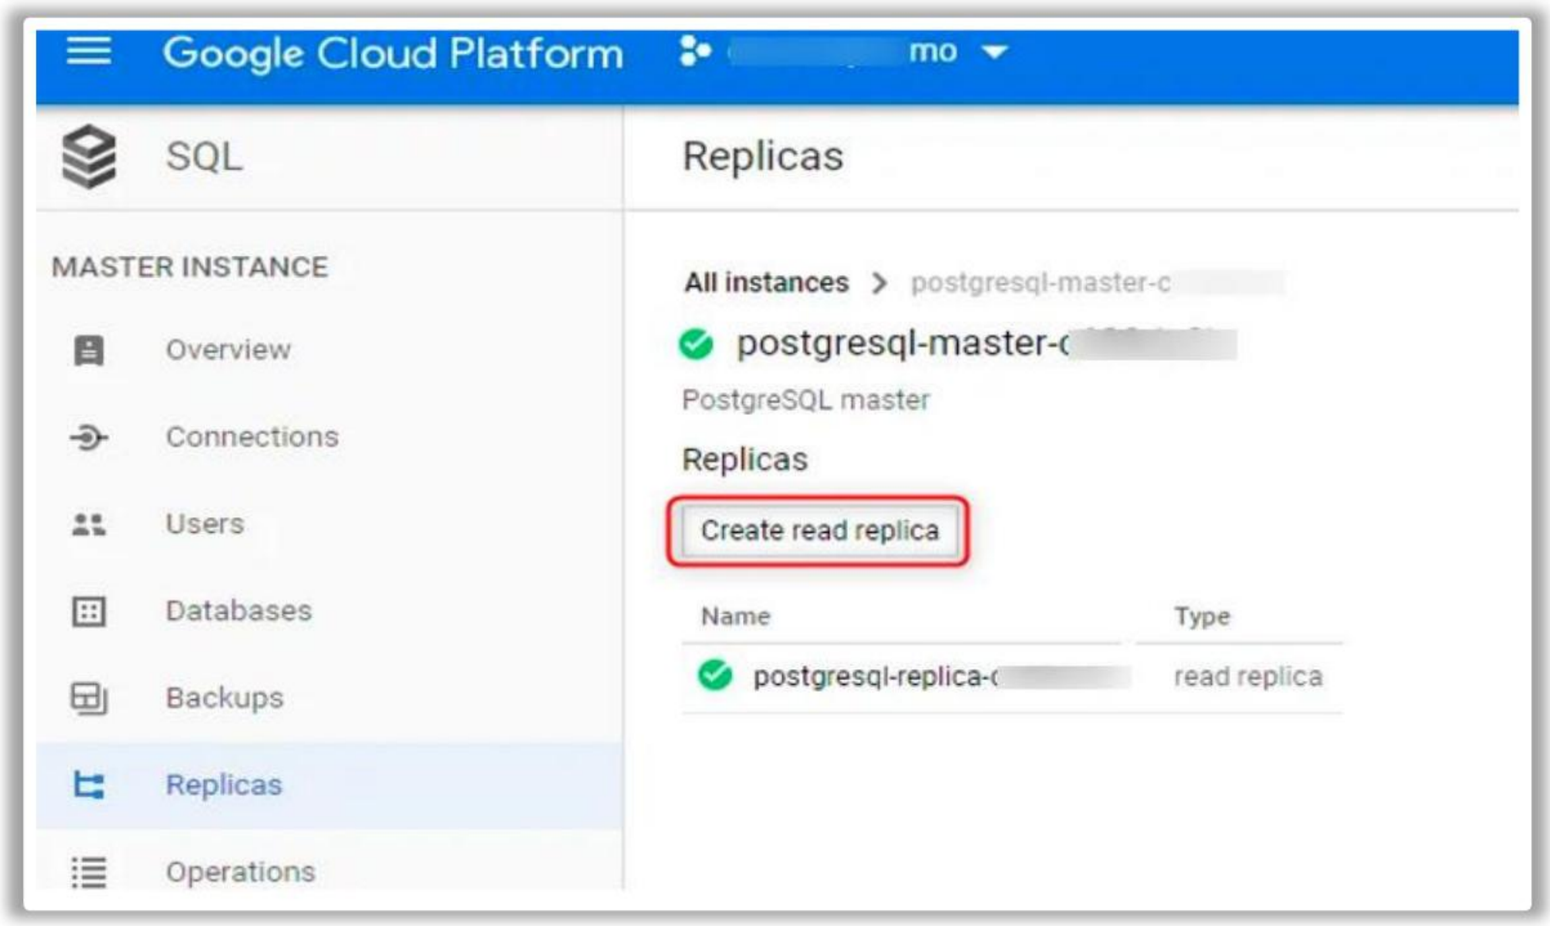

Figure 9.197: Create read replica

- To enable backup and binary logging click on **Automate Backup** and then click on **Enable Binary Logging.**
- Click **Continue**
- Click **Save & Restart.**
- Enter the instance id, name, region (cross-region), and zone in the **Create read replica** page.

hide01.ir

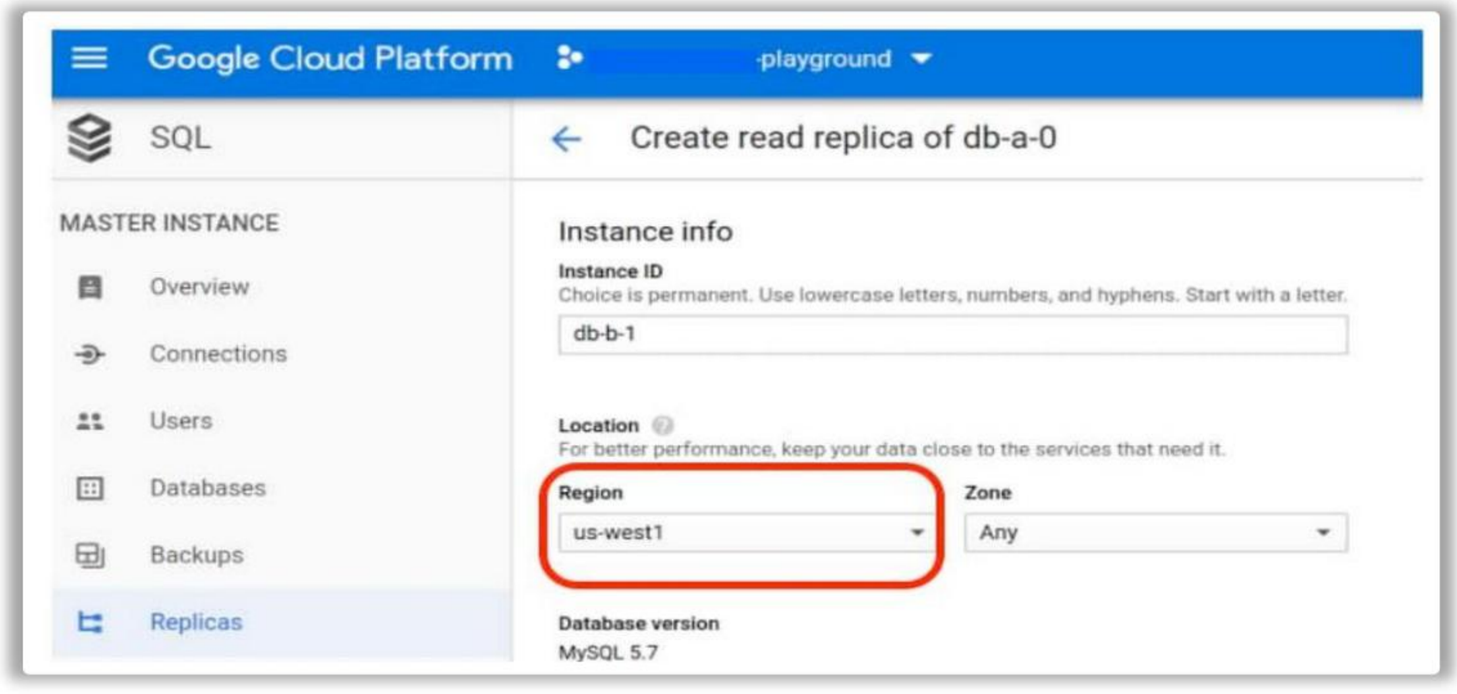

Figure 9.198: Check Region

- Click on **Create** which creates the backup and replica.

With gcloud Script

Certified Cloud Security Engineer Copyright © by **EC-Council**
All Rights Reserved. Reproduction is Strictly Prohibited.

- The status of the primary instance can be checked to confirm that it is not a replica by issuing the following command and if databaseReplicationEnabled is false, it is not a replica

```
gcloud sql instances describe PRIMARY_INSTANCE_NAME
```

- To enable backups, give the following command

```
gcloud sql instances patch PRIMARY_INSTANCE_NAME --backup-start-time
>HH:MM
```

- To enable binary logs on the primary instance

```
gcloud sql instances patch --enable-bin-log PRIMARY_INSTANCE_NAME
```

- The replica can be created using the following command.

```
gcloud sql instances create REPLICA_NAME --master-instance-
name=PRIMARY_INSTANCE_NAME
```

Step 2: Promoting a replica.

With Console

- Navigate to **Cloud SQL Instances** page in Console.
- Click on the name of a replica instance.
- Click on **Promote Replica.**

hide01.ir

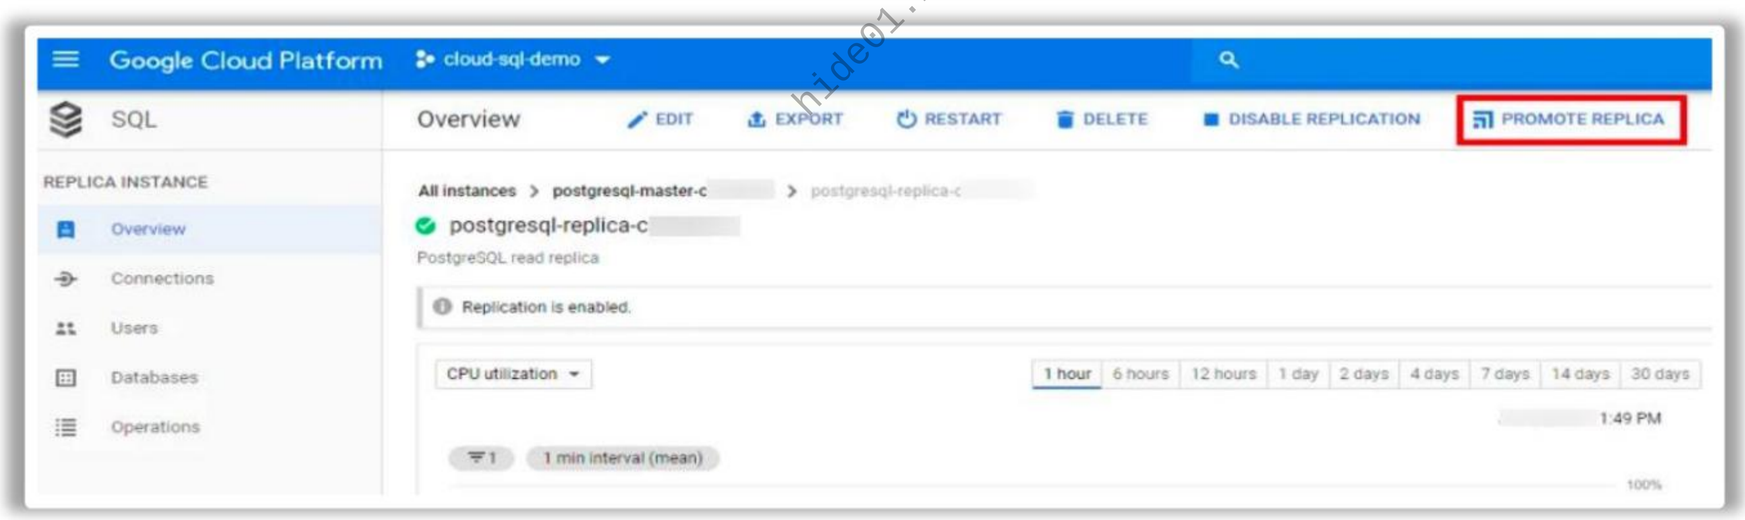

Figure 9.199: Click Promote Replica

- Click Ok.

With gcloud Script

- To promote the replica to a standalone instance, execute the following command

```
gcloud sql instances promote-replica [REPLICA_NAME]
```


Certified Cloud Security Engineer Copyright © by **EC-Council**
All Rights Reserved. Reproduction is Strictly Prohibited.

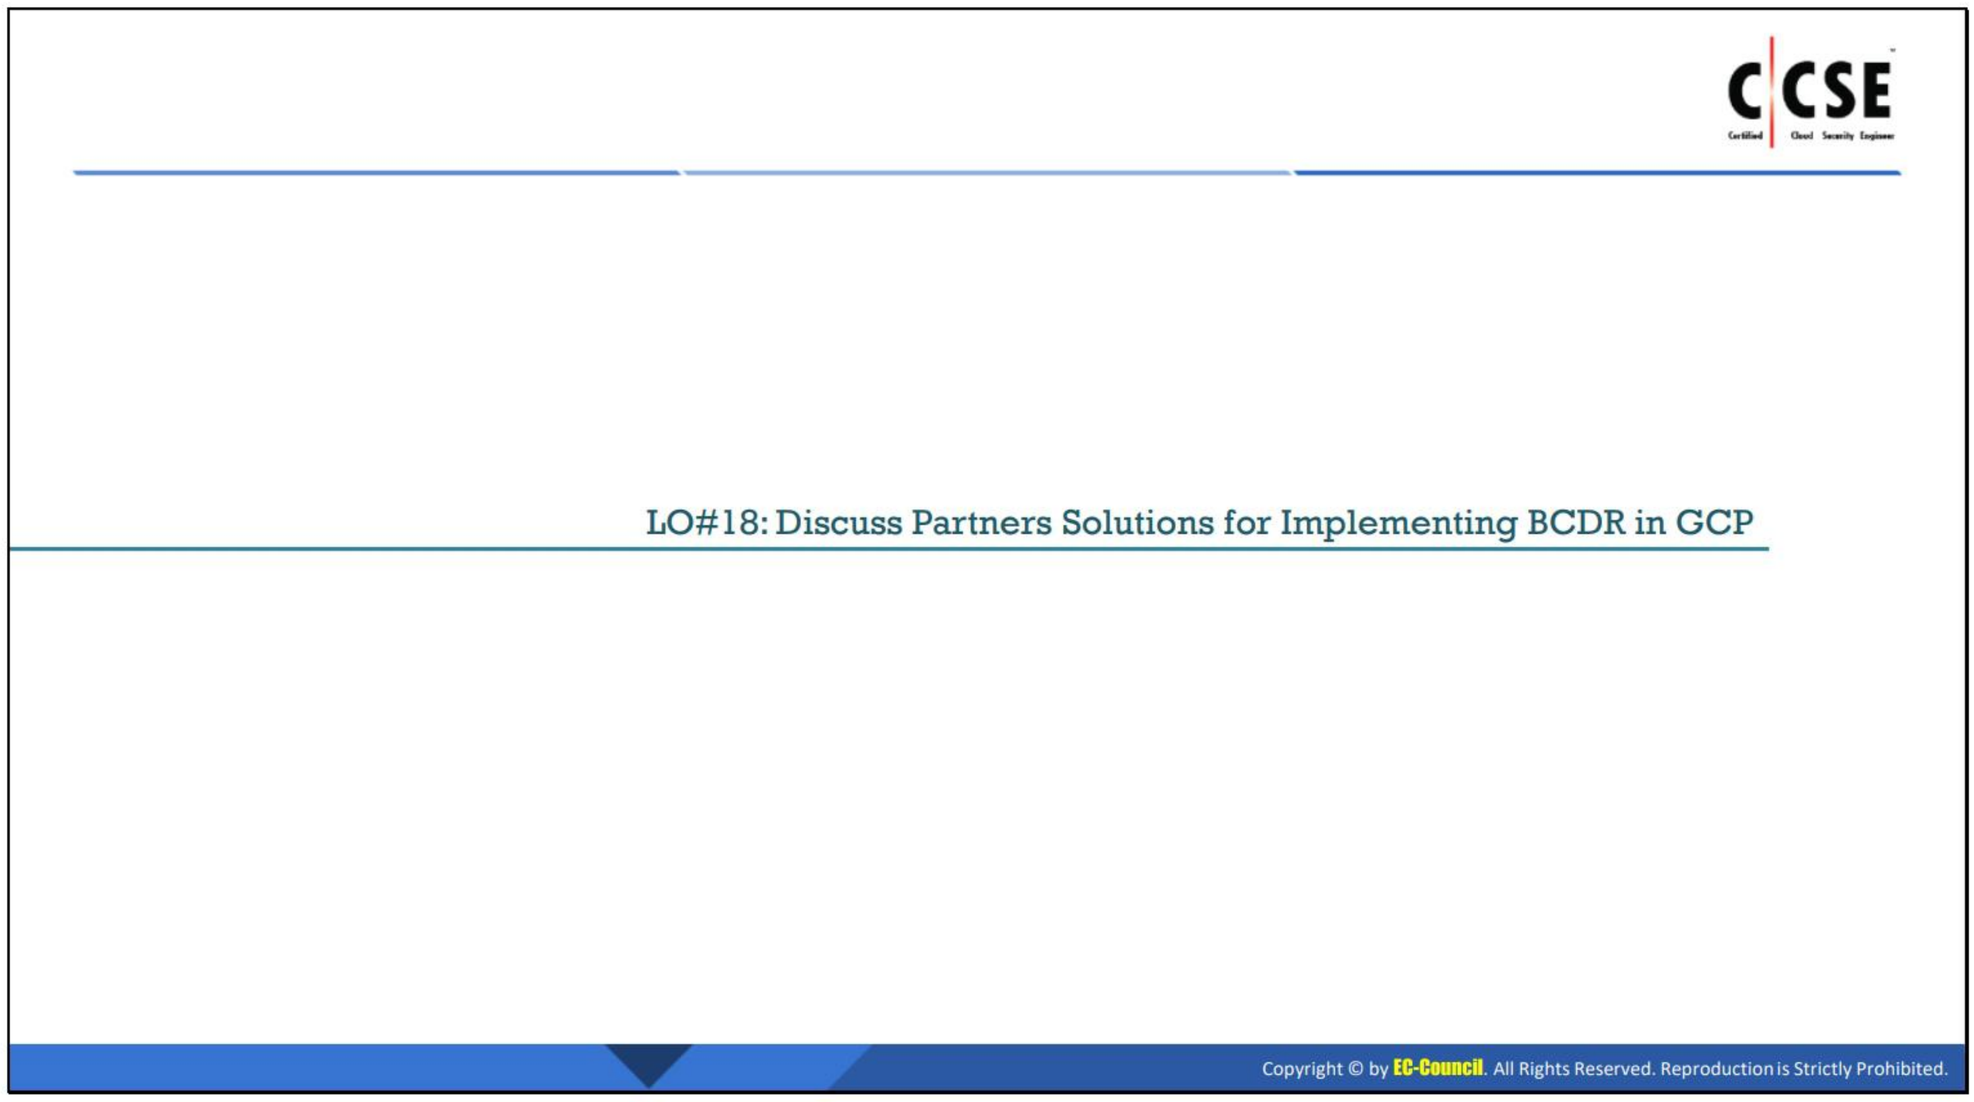

LO#18: Discuss Partners Solutions for Implementing BCDR in GCP

The objective of this section is to understand the various partner solutions for implementing BCDR in GCP.

hide01.ir

Certified Cloud Security Engineer Copyright © by **EC-Council**
All Rights Reserved. Reproduction is Strictly Prohibited.

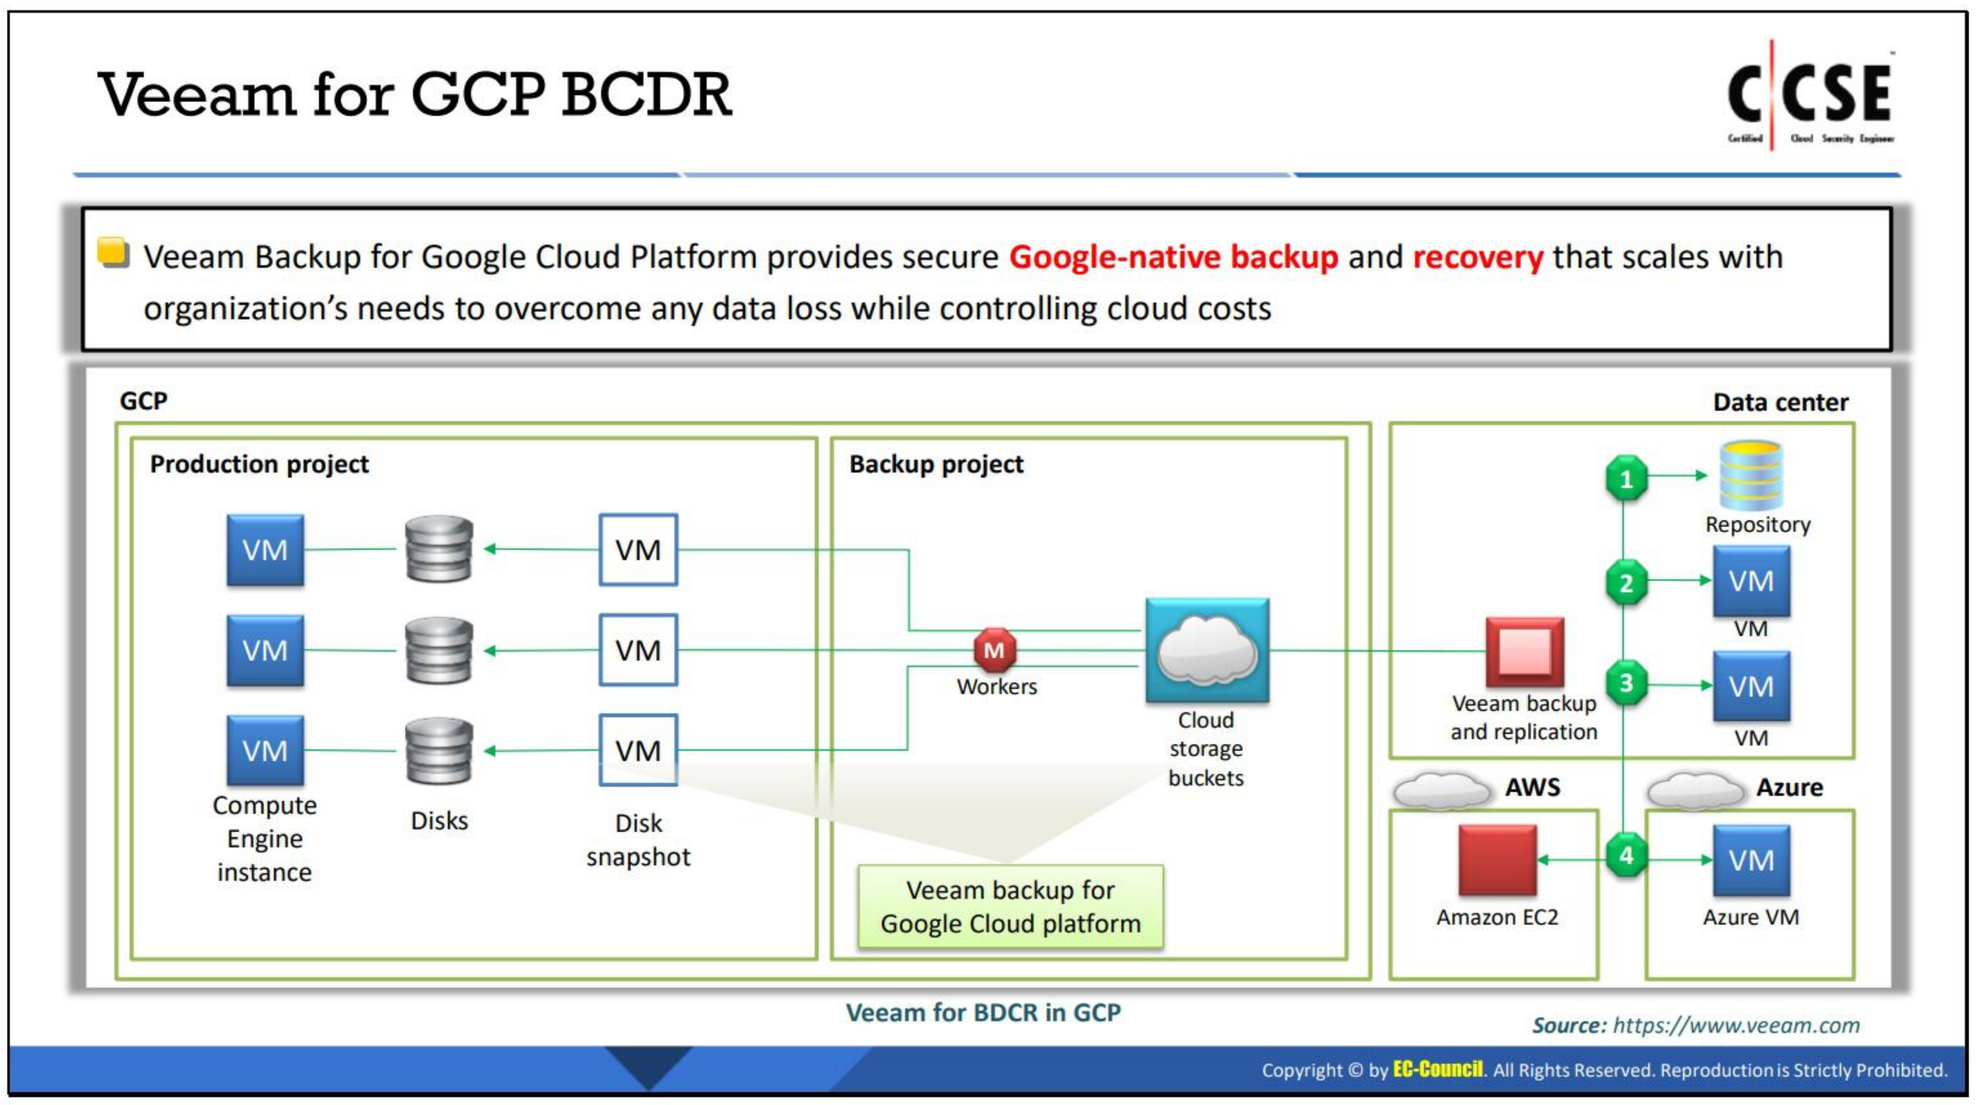

Veeam for GCP BCDR

Source: https://www.veeam.com

Veeam provides a single platform for modernizing backup, accelerating hybrid cloud, and securing data. It provides a simple, flexible, and reliable backup service so that customers can go beyond and focus on their business. GCP and Veeam continue to deliver a tightly integrated portfolio for protecting and managing hybrid-cloud data, inclusive of Google Compute Engine. Veeam Backup for GCP provides secure Google-native backup and recovery that scales with your needs to overcome any data loss while controlling cloud costs. Easily deployable from the GCP Marketplace, you can protect your applications and data in minutes with purpose-built data protection from industry leaders.

Quickly overcome any data loss with flexible instance, volume, and file-level recovery options. Isolate backups from production to secure data against internal and external threats. Effortlessly protect all data—even at scale—with policy-based automation and labels. Automatically tier snapshots to object storage for low-cost retention and compliance.

- **Reliable Google Cloud backup**: Google-native backup and recovery powered by policy-based automation to eliminate the risk of any data loss.
- **Data security and compliance**: Overcome internal and external security threats while ensuring compliant backup and archive.
- **Cost-effective protection**: Industry-first cost calculation and management tools to control costs and avoid cloud overspend.
- **Hybrid-cloud ready**: Backup, recover, and migrate Google workloads within GCP or to private and other public clouds.

Certified Cloud Security Engineer Copyright © by **EC-Council**
All Rights Reserved. Reproduction is Strictly Prohibited.

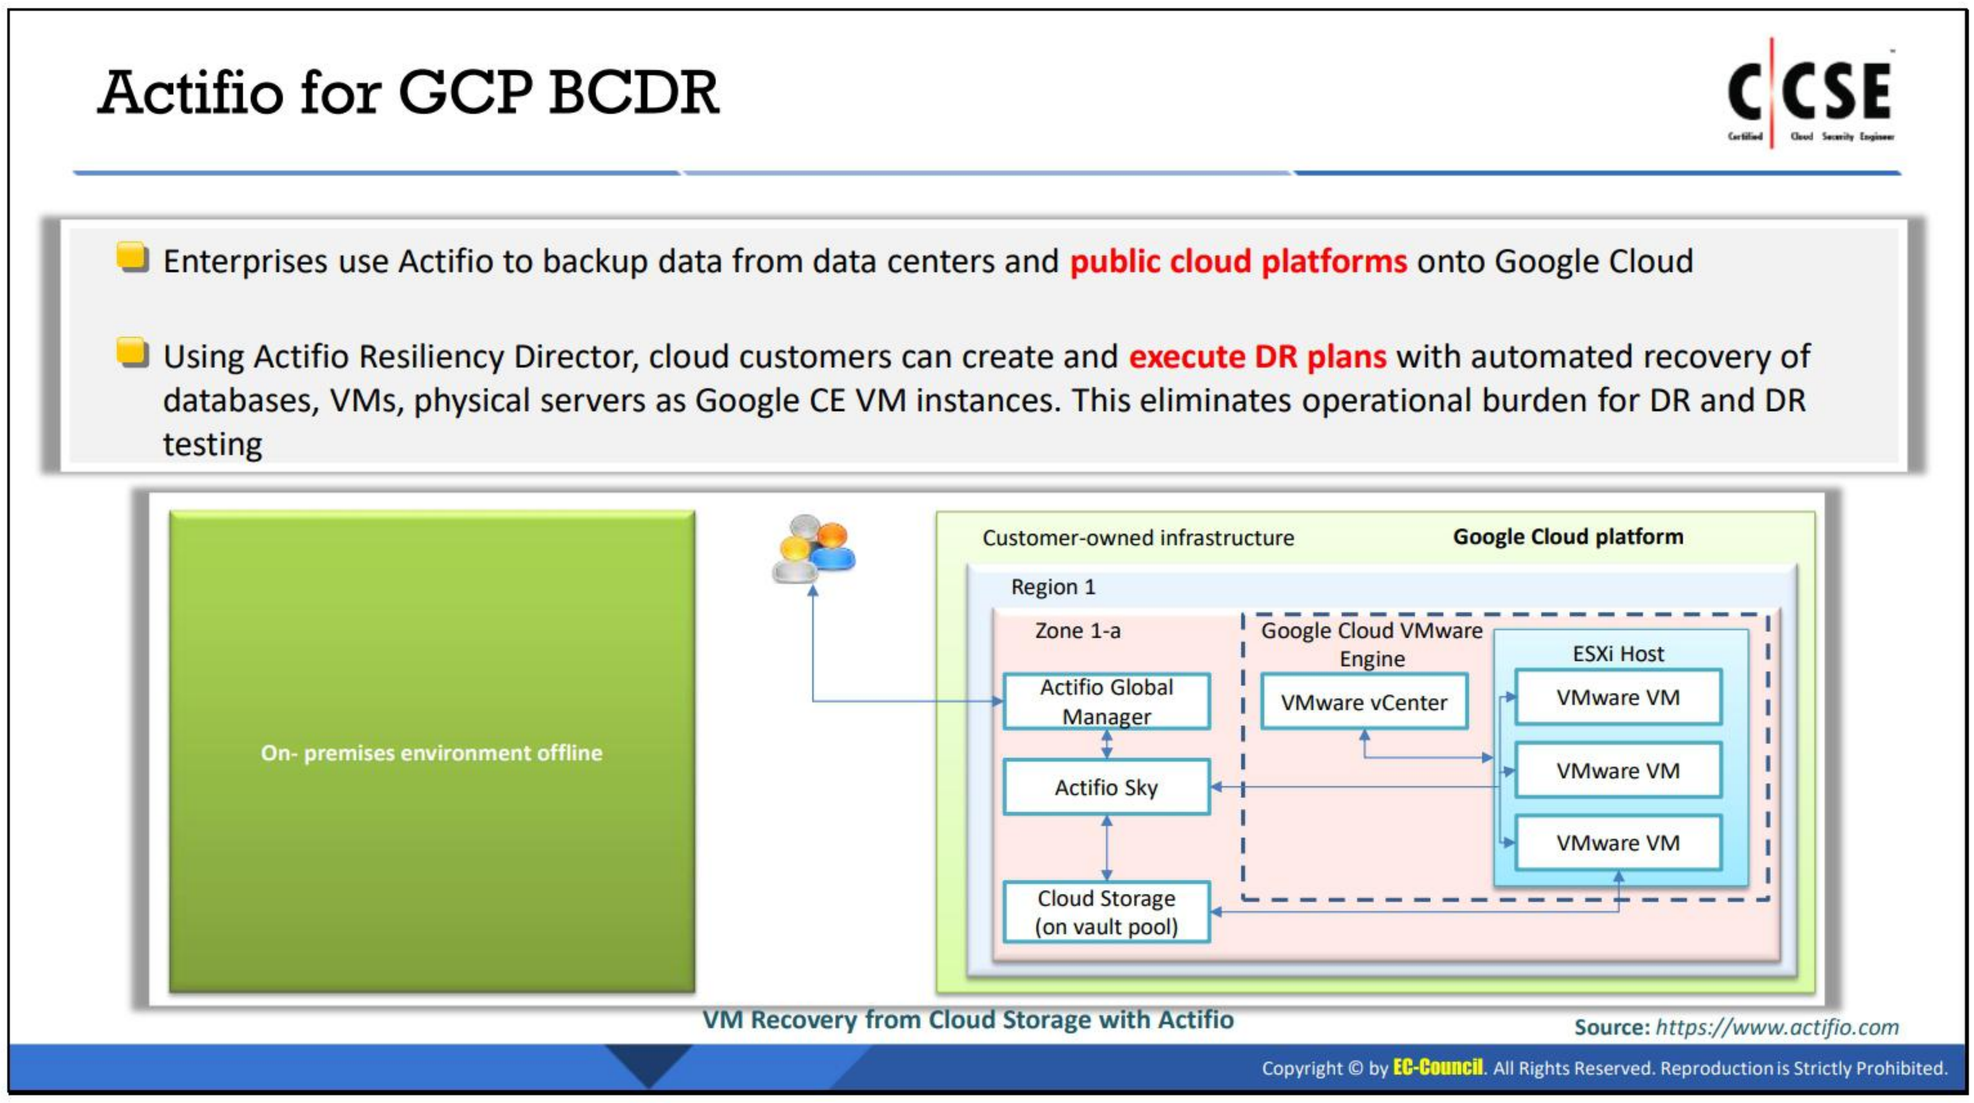

Actifio for GCP BCDR

Source: https://www.actifio.com

Enterprises use Actifio to backup data from data centers and public cloud platforms onto Google Cloud. Using Actifio Resiliency Director, cloud customers can create and execute DR plans with automated recovery of databases, VMs, physical servers as Google CE VM instances. This eliminates the operational burden for DR and DR testing. The backups can be reused to provision rapid database clones for test/dev/analytics in Google Cloud.

Actifio can be combined with GCP VMware Engine to build a cost-effective DR solution for VMware workloads. It reduces the cost by dynamically creating VMware software designed data center (SDDC). It also enables direct and instant data access from Google Cloud Storage through Actifio.

Certified Cloud Security Engineer Copyright © by **EC-Council**
All Rights Reserved. Reproduction is Strictly Prohibited.

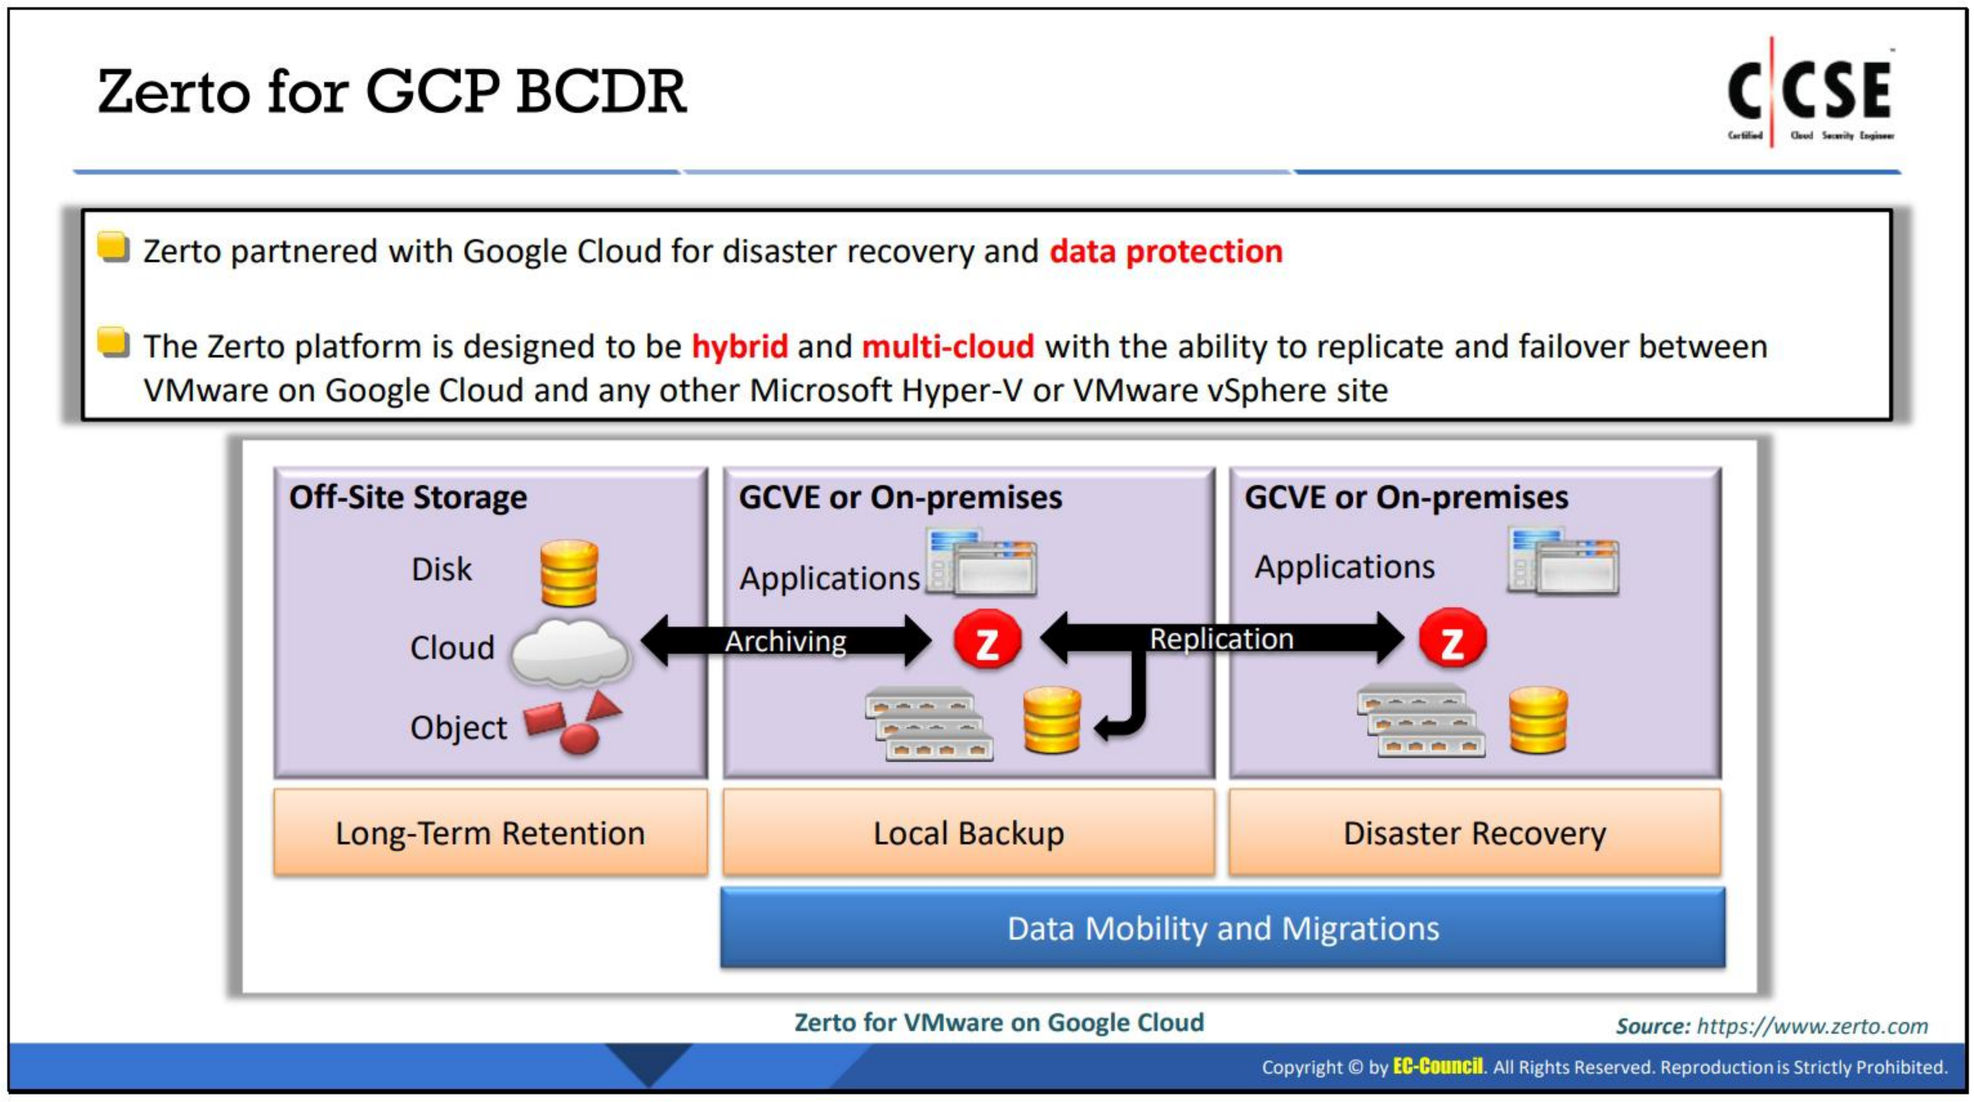

Zerto for GCP BCDR

Source: https://www.zerto.com

Zerto partnered with Google Cloud for disaster recovery and data protection. The Zerto platform is designed to be hybrid and multi-cloud with the ability to replicate and failover between VMware on Google Cloud and any other Microsoft Hyper-V or VMware vSphere site. With always-on replication and journal-based recovery, Zerto helps to seamlessly rewind and recover entire sites, applications, VMs, and individual files. Applications will recover using predefined settings such as boot order, network configuration, and optional re-IP.

Innovative consistency groupings allow for DR replication, migration, and recovery of all VMs associated with an application to or from VMware on Google Cloud in the event of an application failure, site-wide outage, or planned migration.

Certified Cloud Security Engineer Copyright © by **EC-Council**
All Rights Reserved. Reproduction is Strictly Prohibited.

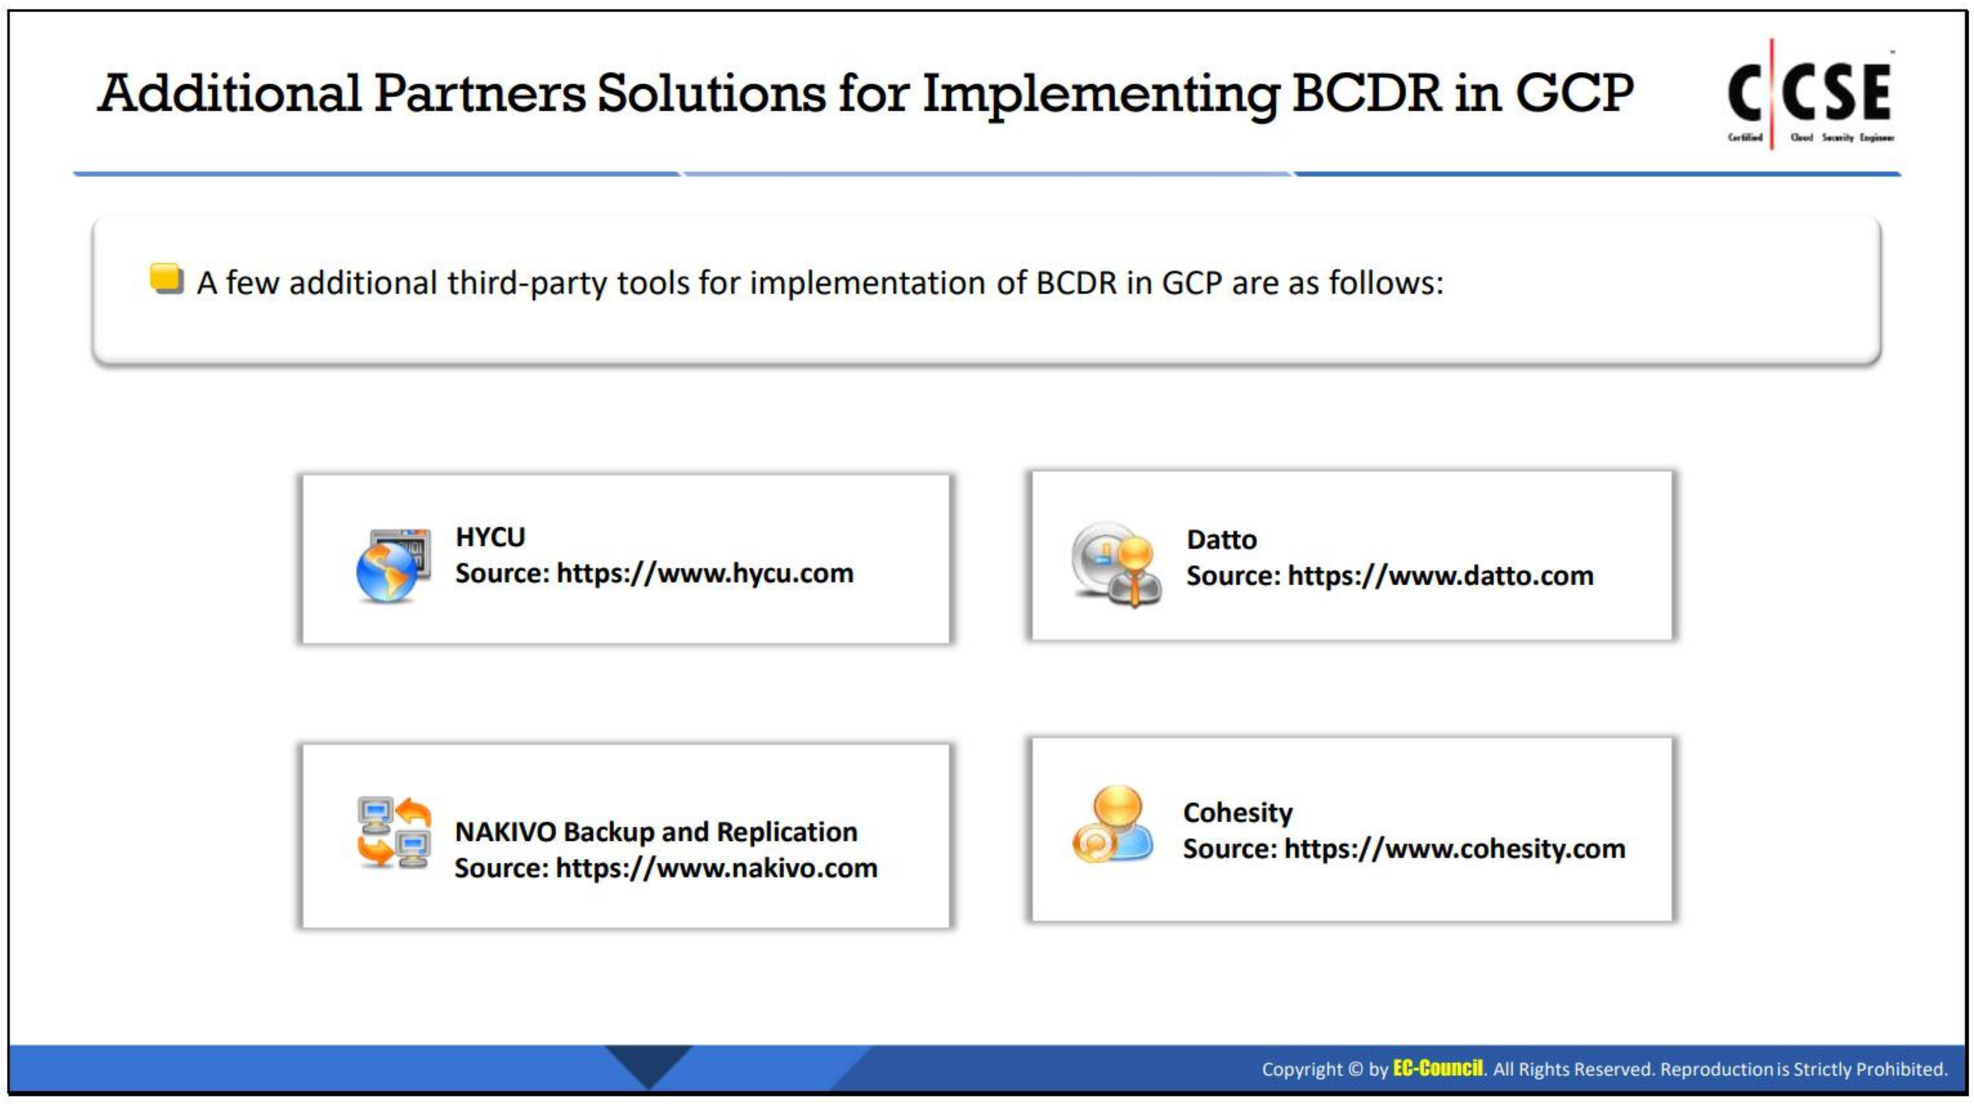

Additional Partners Solutions for Implementing BCDR in GCP

HYCU

Source: https://www.hycu.com

hade01.ir

HYCU delivers the first purpose-built backup-as-a-service for Google Cloud. With HYCU, cloud customers do not need to deploy or manage their backups. They can just choose the virtual machine (VM) that you want to protect and assign the desired policy and HYCU for Google Cloud will take care of the rest. HYCU completely takes care of backup scheduling, automates retention, and storage tiering. Cloud customers can recover individual folders or files from VMs in minutes.

HYCU for Google Cloud also enables customers to protect instances, applications running on instances, and data in buckets. After establishing the data protection environment, data protection can be enabled. After the first backup is successfully completed, data can be restored if it becomes damaged or corrupted.

Datto

Source: https://www.datto.com

Datto SaaS Protection is the leading cloud-to-cloud backup solution built exclusively for minimum support prices (MSPs), protecting thousands of businesses today. Datto SaaS Protection offers comprehensive backup and recovery for critical cloud data that lives in Microsoft 365 and Google Workspace applications. It protects against permanent data loss and recovers from ransomware or user error quickly with a one-click restore.

NAKIVO Backup and Replication

Source: https://www.nakivo.com

Certified Cloud Security Engineer Copyright © by **EC-Council**
All Rights Reserved. Reproduction is Strictly Prohibited.

NAKIVO Backup and Replication provides image-based, application-aware, and incremental backup and replication. It comes with built-in features that increase backup performance, improve reliability, speed up recovery and help save time as well as money. It allows to send VM backups or their copies to private/public clouds such as AWS, GCP, etc.

Cohesity

Source: https://www.cohesity.com

Cohesitfy's native integration with GCP helps consolidate multiple data workloads in a hybrid architecture. Google Cloud and Cohesity give a joint solution to take back control of the data over the hybrid cloud. Cohesity Helios consolidates all data infrastructure into one unified and efficient multi-cloud. The joint solution enables customers to use Google Cloud for data protection, long-term archival, dev/test, and disaster recovery. Cohesity can archive data to all Google Cloud Storage classes for long-term retention data platform. It can protect Google Cloud Compute Engine VMs and applications running on Google Cloud with efficient snapshots APIs. Snapshot backups can be stored on Cohesity to maximize data reduction efficiency or they can be stored directly on Google Cloud Storage and managed remotely from a Cohesity cluster.

When it is deployed on-premises, it provides simple connectivity to Google Cloud to extend the data center to the cloud for LTR, DR, test/dev, and tiering as well providing protection for Google Cloud VMware Engine.

hide01.ir

Certified Cloud Security Engineer Copyright © by **EC-Council**
All Rights Reserved. Reproduction is Strictly Prohibited.

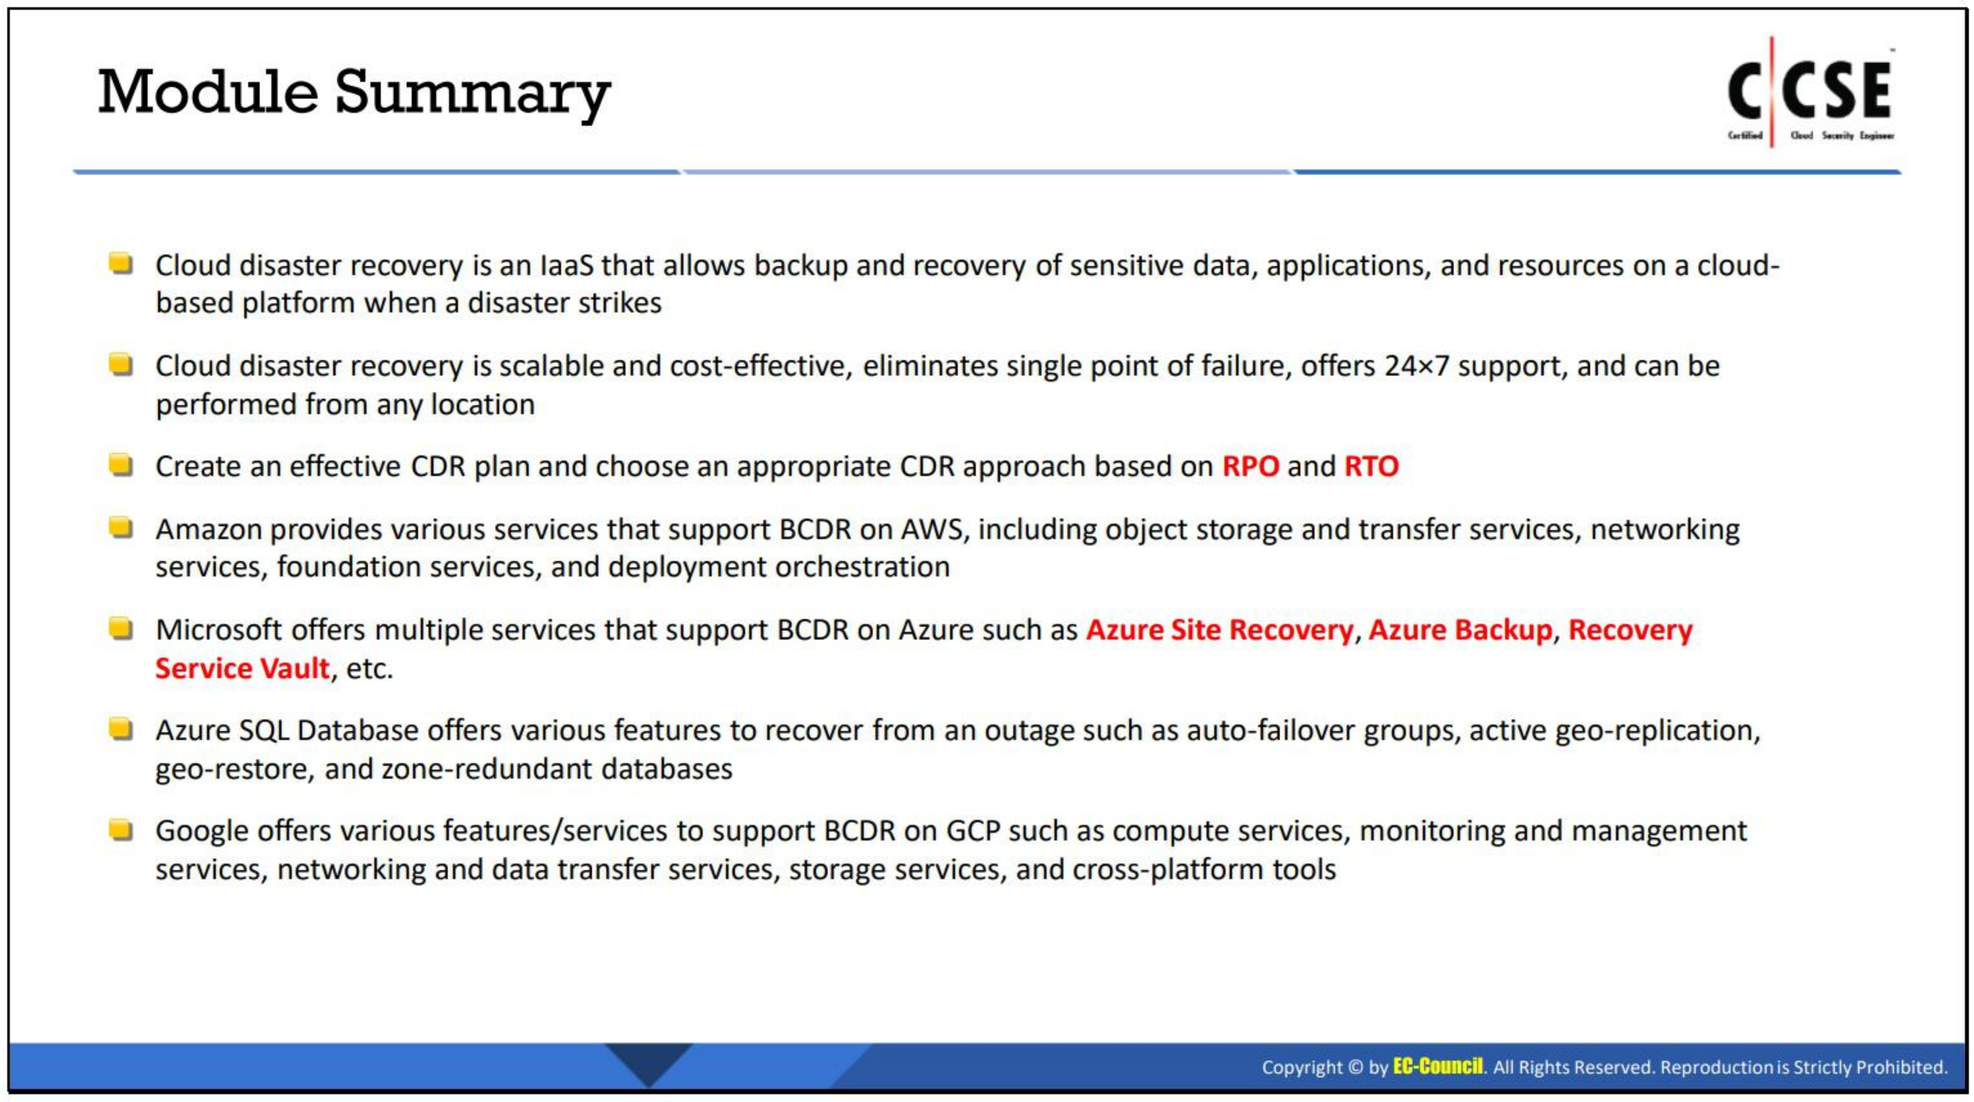

Module Summary

This module described important concepts of cloud BCDR, designing a DR plan in the cloud environment, and various features/services offered by different cloud providers (AWS, Azure, and GCP). By the end of this module, the cloud security engineer will be able to prepare and implement an effective DR plan to safeguard the organization against a disaster event.

Following are the key points discussed in this module:

- CDR is an IaaS that allows backup and recovery of sensitive data, applications, and resources on a cloud-based platform when a disaster strikes.
- CDR is scalable and cost-effective, eliminates single point of failure, offers 24 × 7 support, and can be performed from any location.
- Create an effective CDR plan and choose an appropriate CDR approach based on **RPO** and **RTO.**
- Amazon provides various services that support BCDR on AWS, including object storage and transfer services, networking services, foundation services, and deployment orchestration.
- Microsoft offers multiple services that support BCDR on Azure such as **Azure Site Recovery**, **Azure Backup**, **Recovery Service Vault**, etc.
- Azure SQL Database offers various features to recover from an outage such as auto-failover groups, active geo-replication, geo-restore, and zone-redundant databases.
- Google offers various features/services to support BCDR on GCP such as compute services, monitoring and management services, networking and data transfer services, storage services, and cross-platform tools.cc

Certified Cloud Security Engineer Copyright © by **EC-Council**
All Rights Reserved. Reproduction is Strictly Prohibited.

This page is intentionally left blank.

hide01.ir