EC-Council

C|CSE™

Certified Cloud Security Engineer



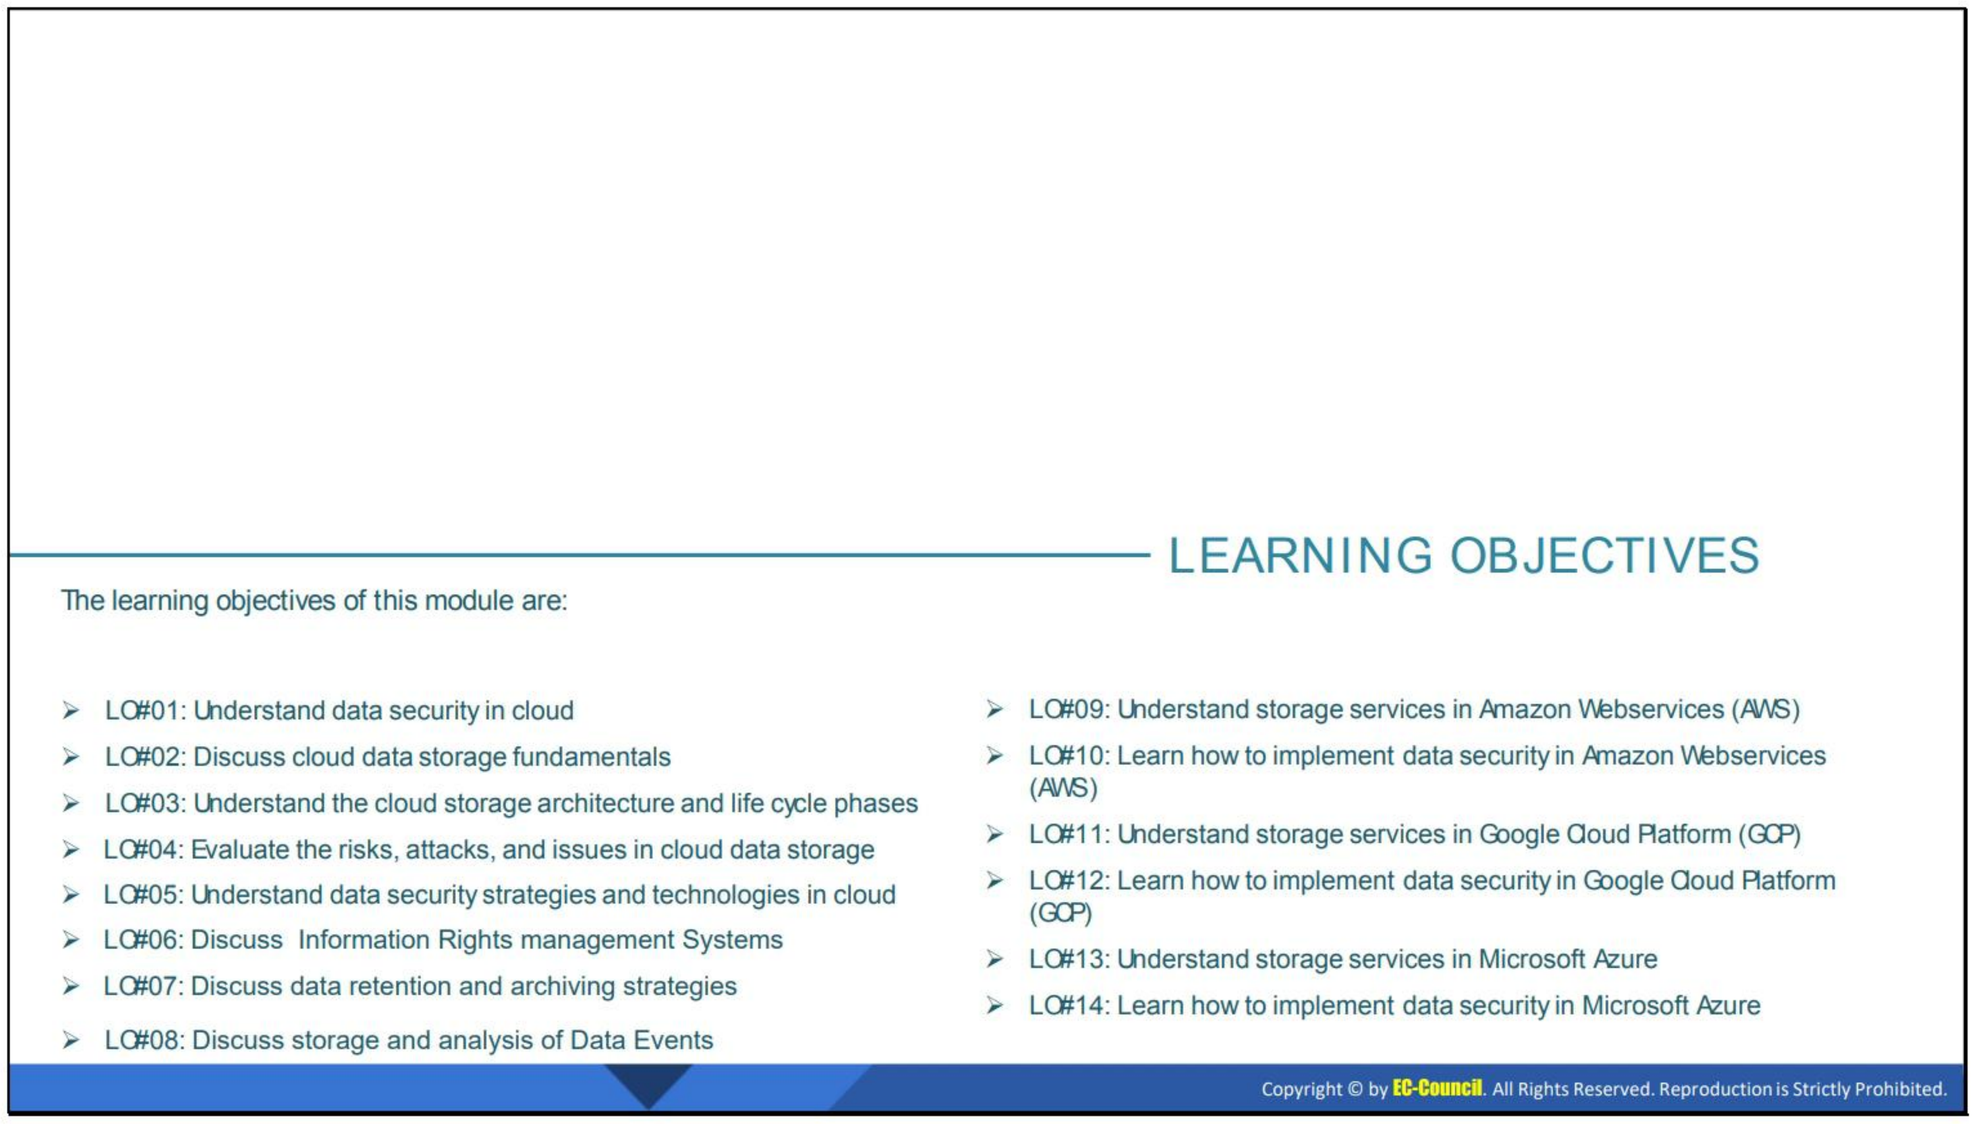
hide01.ir

Data Security in Cloud

Module 04

This page is intentionally left blank.

hide01.ir



Learning Objectives

Migration to cloud environment poses a risk to the organizations data as they are required to transfer their sensitive information in a shared pool. Data security is the key to achieving information governance in the cloud. Data security is dependent on risk-based assessment as every information is not equally important to secure. It is very important to secure sensitive data. Hence, it is essential to understand the various security controls that can secure cloud data. The module discusses the various technical aspects involved in securing cloud data. The module discusses cloud storage types, cloud storage architecture, encryption, key management, data security strategies, information rights management (IRM), data deletion, retention, and archiving policies, data events, and risks in cloud data storage. It also briefly discusses the various storage services and their implementation in Amazon Web Services (AWS), Azure, and Google Cloud Platform (GCP).

The learning objectives of this module are as follows:

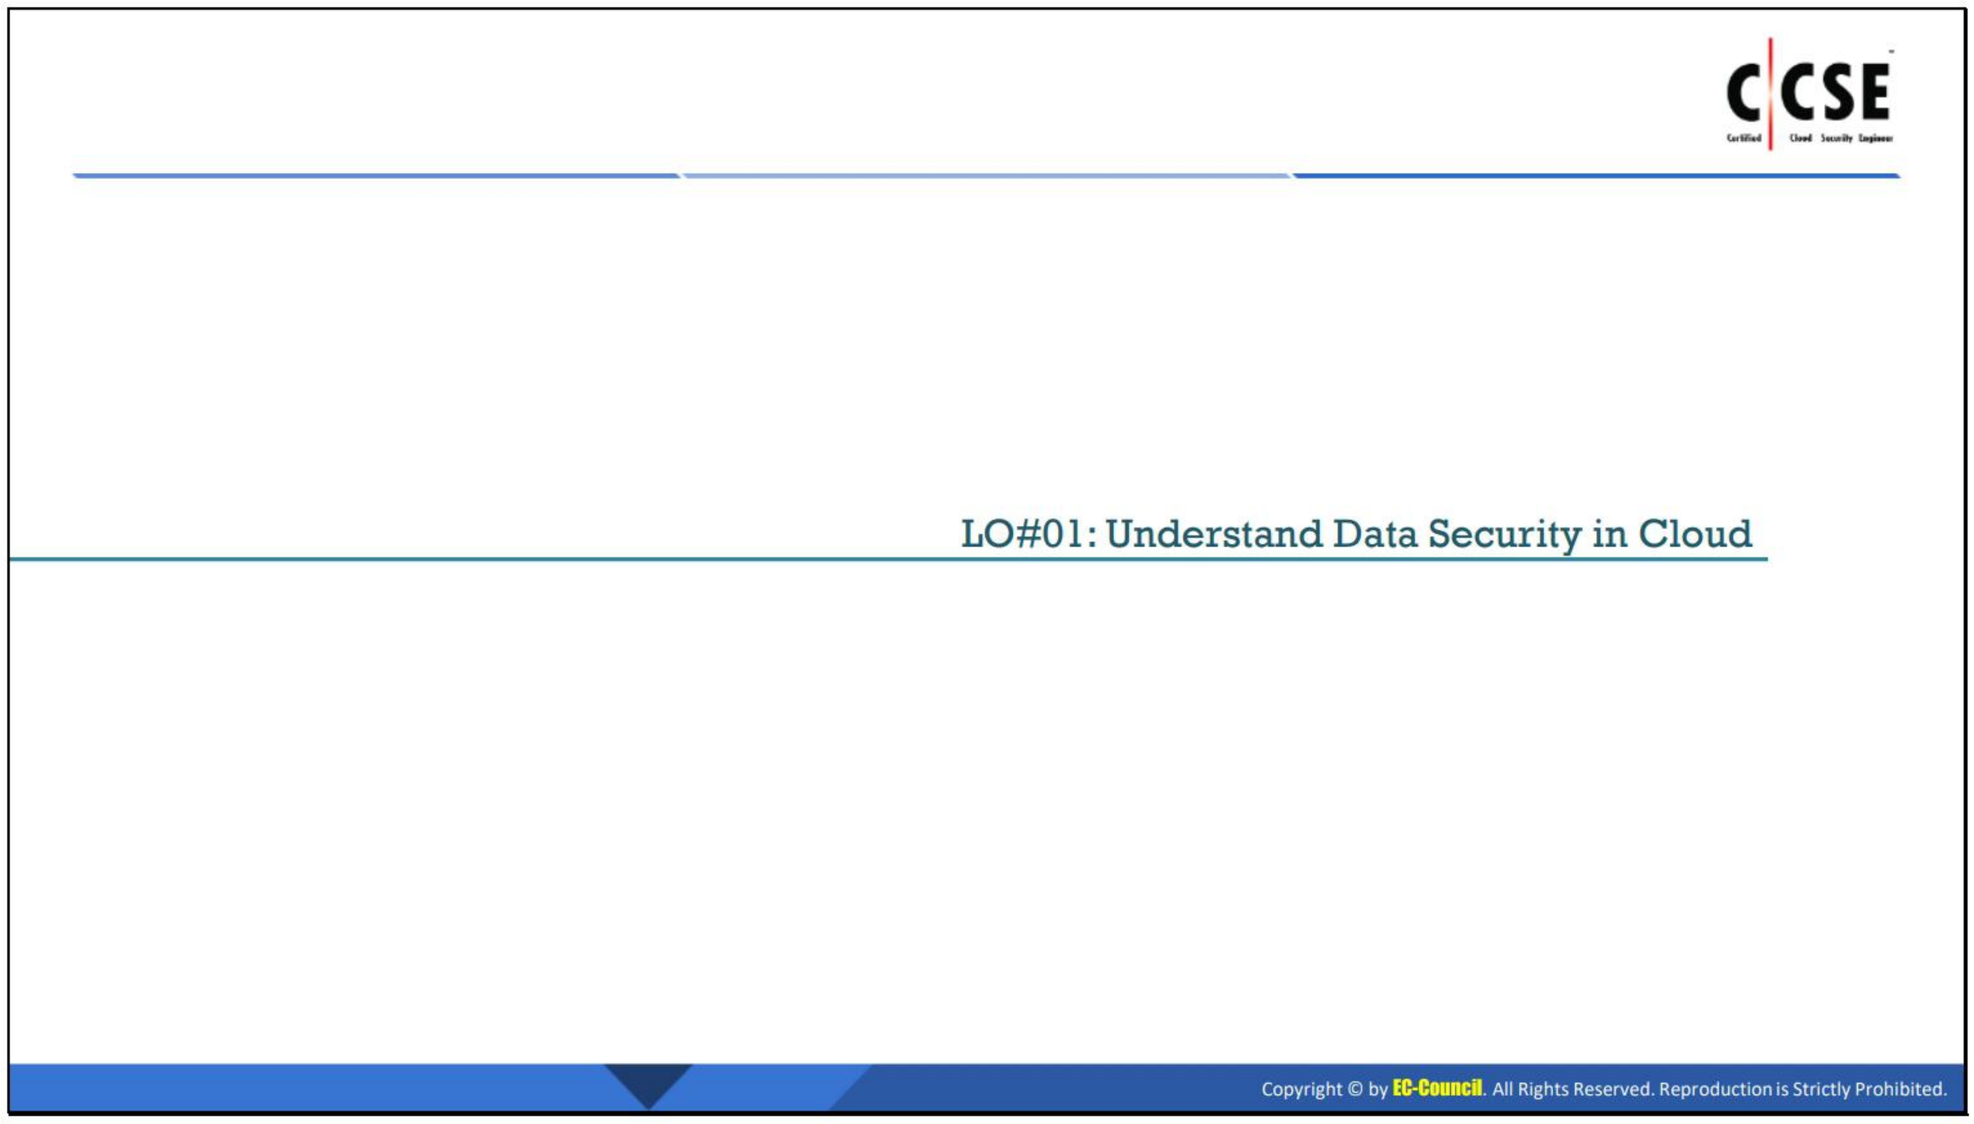
- Understanding Data Security in Cloud
- Understanding Cloud Data Storage Fundamentals
- Understand Cloud Storage Architecture and Cloud Data Life Cycle
- Evaluate Risks, Attacks and Issues in Cloud Data Storage
- Understand Data Security Strategies and Technologies
- Discuss Information Rights Management Systems
- Discuss Data Retention and Archiving Strategies
- Discuss Storage and Analysis of Data Events
- Understand Storage Services in Amazon Webservices (AWS)

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

- Learn How to Implement Data Security in Amazon Webservices (AWS)
- Understand Storage Services in Google Cloud Platform (GCP)
- Learn How to Implement Data Security in Google Cloud Platform (GCP)
- Understand Storage Services in Microsoft Azure
- Learn How to Implement Data Security in Microsoft Azure

hide01.ir

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

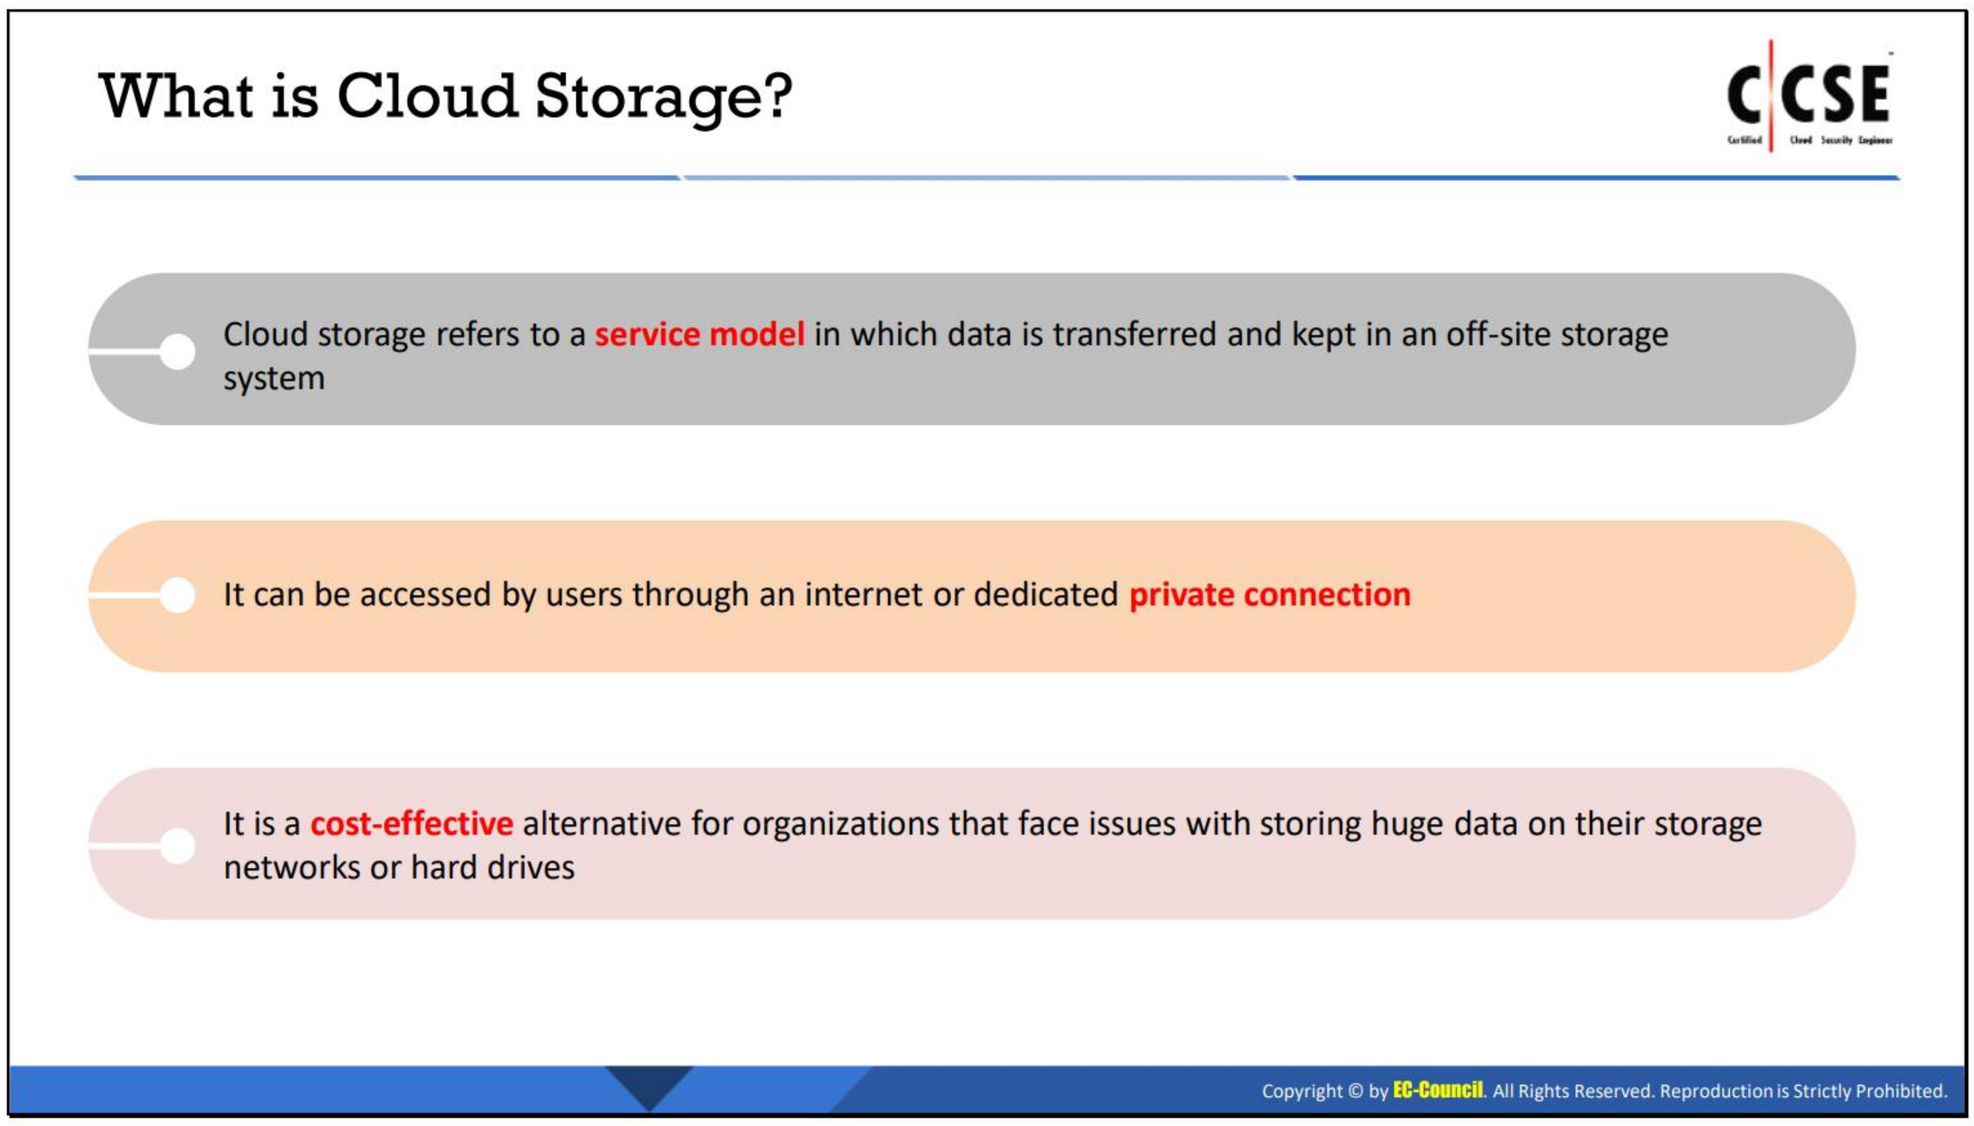


LO#01: Understand Data Security in Cloud

Understanding how the company's data is stored in a cloud environment and the benefits of securing the data is very important. The objective of this section is to explain cloud storage, data security, and the importance of securing cloud data.

hidn91.ir

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

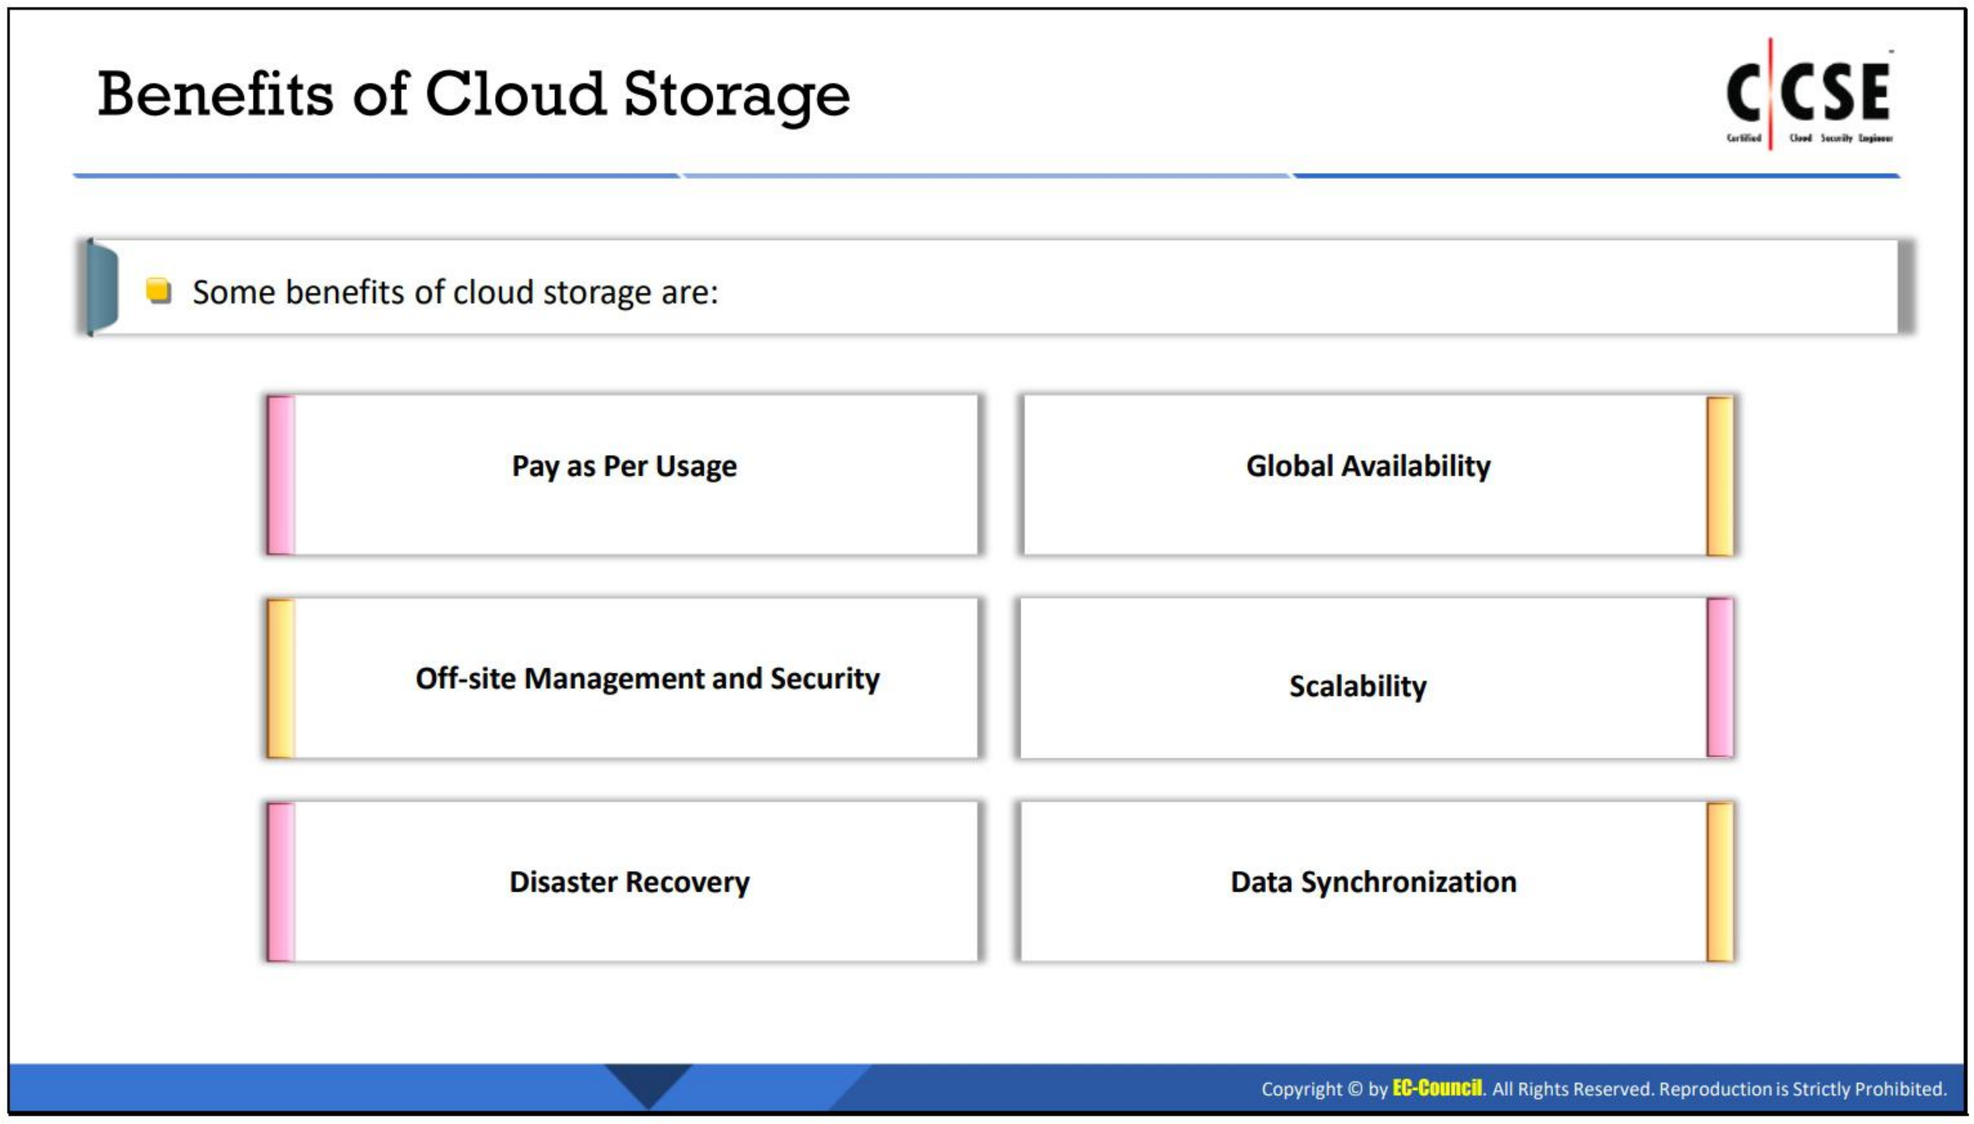


What is Cloud Storage?

Cloud storage refers to a service model where data is transmitted and stored in an offsite location and can be accessed by users, when required, over a network such as the internet or a dedicated private connection. The data is stored on remote systems where the storage provider ensures that the data is properly managed, maintained, and backed up so that users can access it anytime. Cloud storage providers generally charge customers for storing and maintaining their data.

Cloud storage provides a better and cost-effective alternative for organizations that face issues with storing large amounts of data on their storage networks or hard drives. Typically, organizations use storage area networks for storing files and data. The cost of maintaining storage area networks becomes high when they add more data and files as they grow. Instead of adding new servers and investing money to maintain them, organizations can use cloud storage to store and maintain their data securely at a lower cost.

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



Benefits of Cloud Storage

Storing data on the cloud has numerous benefits. Some of them are listed below.

- **Pay as per usage**: Customers who use cloud storage must pay only for the storage they utilize to store their data. They do not have to maintain the entire in-house storage system if they use cloud storage. This significantly reduces the overall costs incurred on storage due to its cost-effective feature.
- **Global availability**: Cloud storage can be accessed anytime and from anywhere; thus, customers do not have to worry about data accessibility, OS capability, and other complex allocation processes. On-premises storage systems offer local accessibility whereas cloud storage provides global accessibility.
- **Offsite management and security**: Once organizations transmit data on the cloud for storage, the provider is responsible for securely maintaining it. This removes the need for the organization to get involved in storage-related tasks, including procurement, administration, installation, and maintenance, which helps increase the staff's focus on other important tasks.
- **Scalability**: With on-premises storage, organizations struggle with data storage constraints. The scalability of storage networks is a major challenge for on-premises systems. Cloud storage allows them to scale (up or down) the storage capacity as per their needs as cloud storage capacity is unlimited. Cloud providers offer flexible plans to customers, which allow them to choose their desired plan as per their business needs.
- **Disaster recovery**: Organizations can use cloud storage for backing up data, which they can use for recovery when they face any disaster or data loss. Cloud platforms provide better disaster recovery services with affordable and flexible plans. Cloud storage also

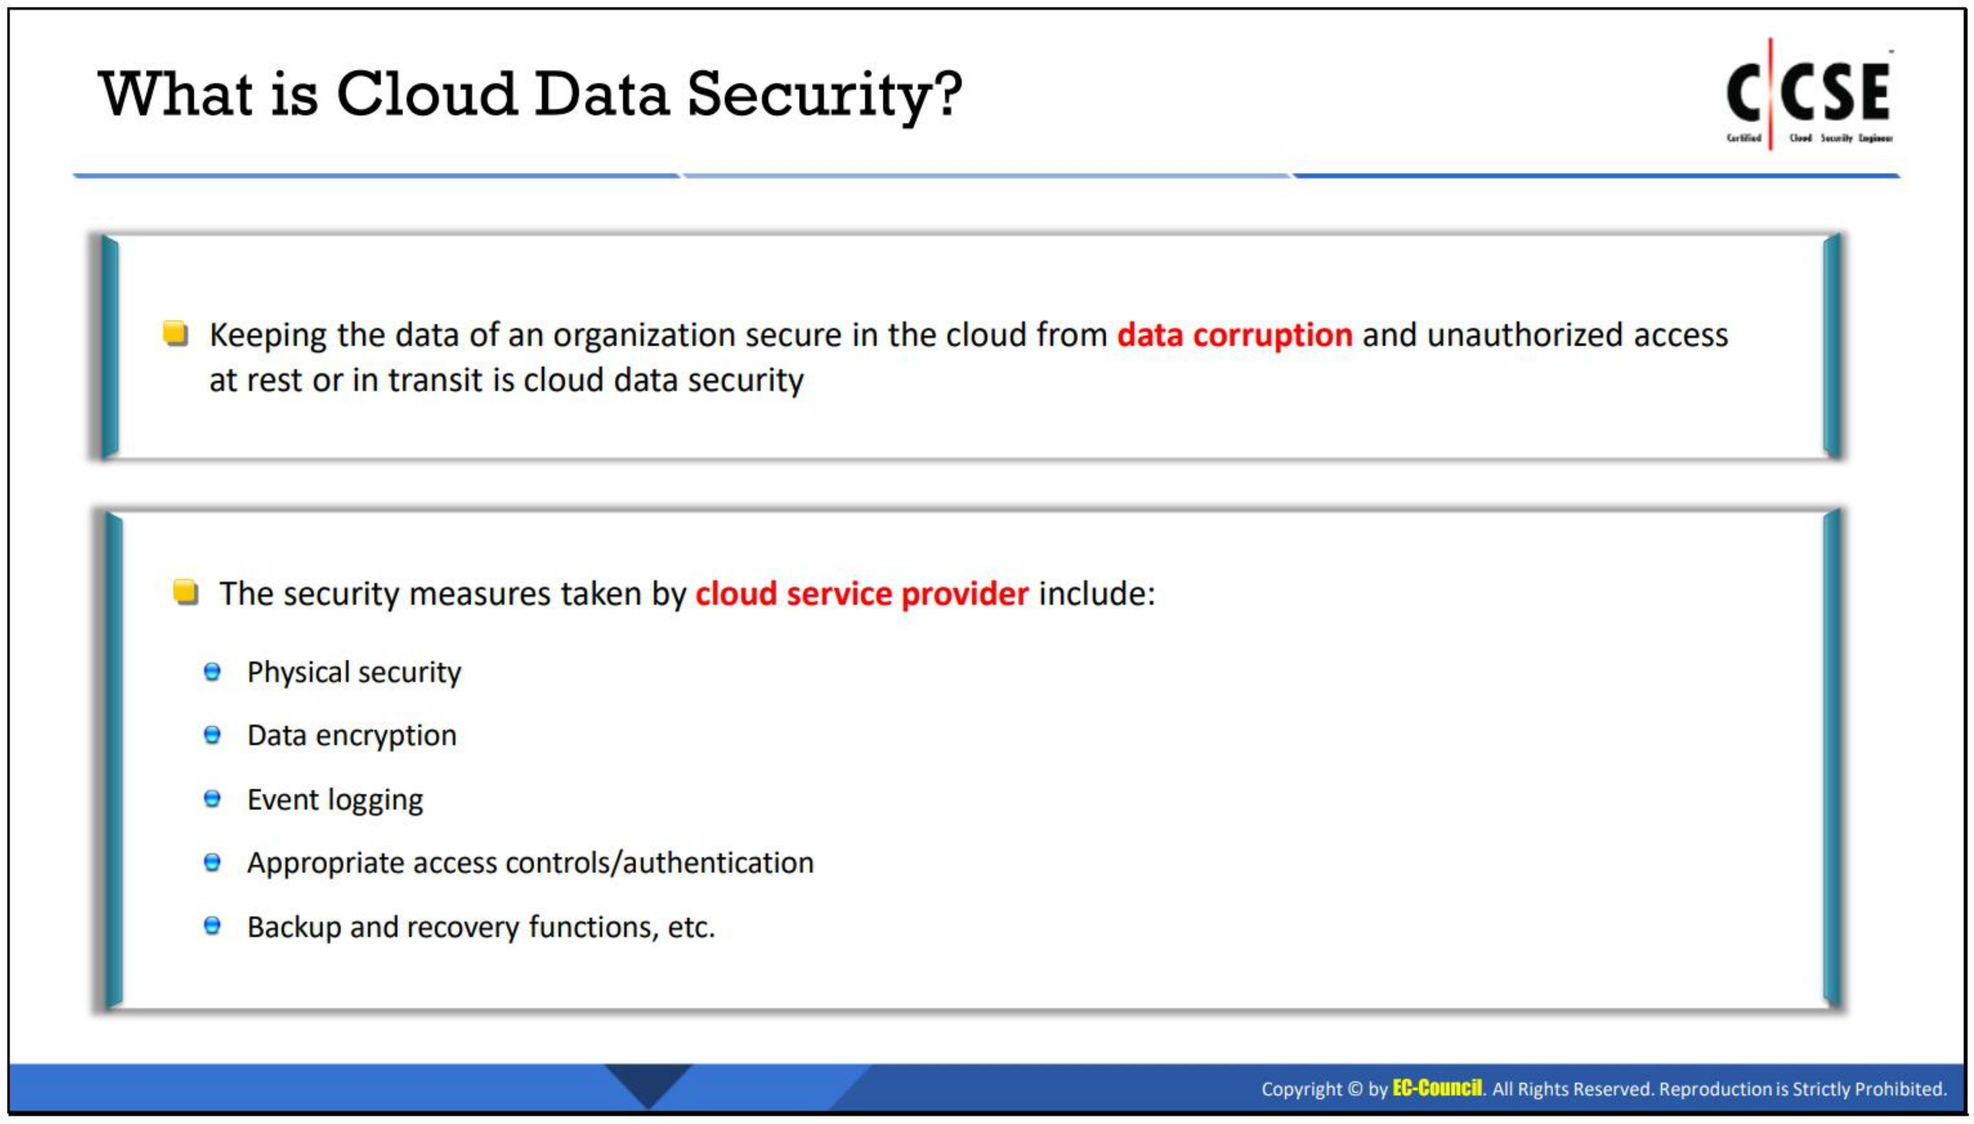
Certified Cloud Security Engineer Copyright © by **EC-Council**
All Rights Reserved. Reproduction is Strictly Prohibited.

makes data retrieval in disaster recovery scenarios much simpler and easier, which helps in business continuity.

- **Data synchronization**: Synchronization of data is crucial. It helps in accessing cloud storage from any device, and cloud storage provides quick synchronization of data. Data stored in the cloud can be accessed from any device in any part of the world. Any changes made to data or files stored on the cloud via any device will be reflected by cloud storage instantly on all affiliated devices.

hide01.ir

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



What is Cloud Data Security?

Cloud data security refers to keeping the data of an organization secure in the cloud. This includes protecting data from data corruption and unauthorized access at rest or in transit. As the data stored on the cloud can be accessed from anywhere and anytime, mostly through the internet, cloud data is subjected to multiple threats and vulnerabilities. Cloud storage providers are responsible for storing, managing, and backing up the data of customers securely; therefore, they consider cloud data security as their top priority.

To ensure cloud data security, these providers use multiple approaches, technologies, and tools. Many security tools required to protect cloud data are already built into the cloud infrastructure. Cloud providers make use of these in-built security features along with other security measures to store data securely.

Some of the security measures taken by cloud service providers are listed below.

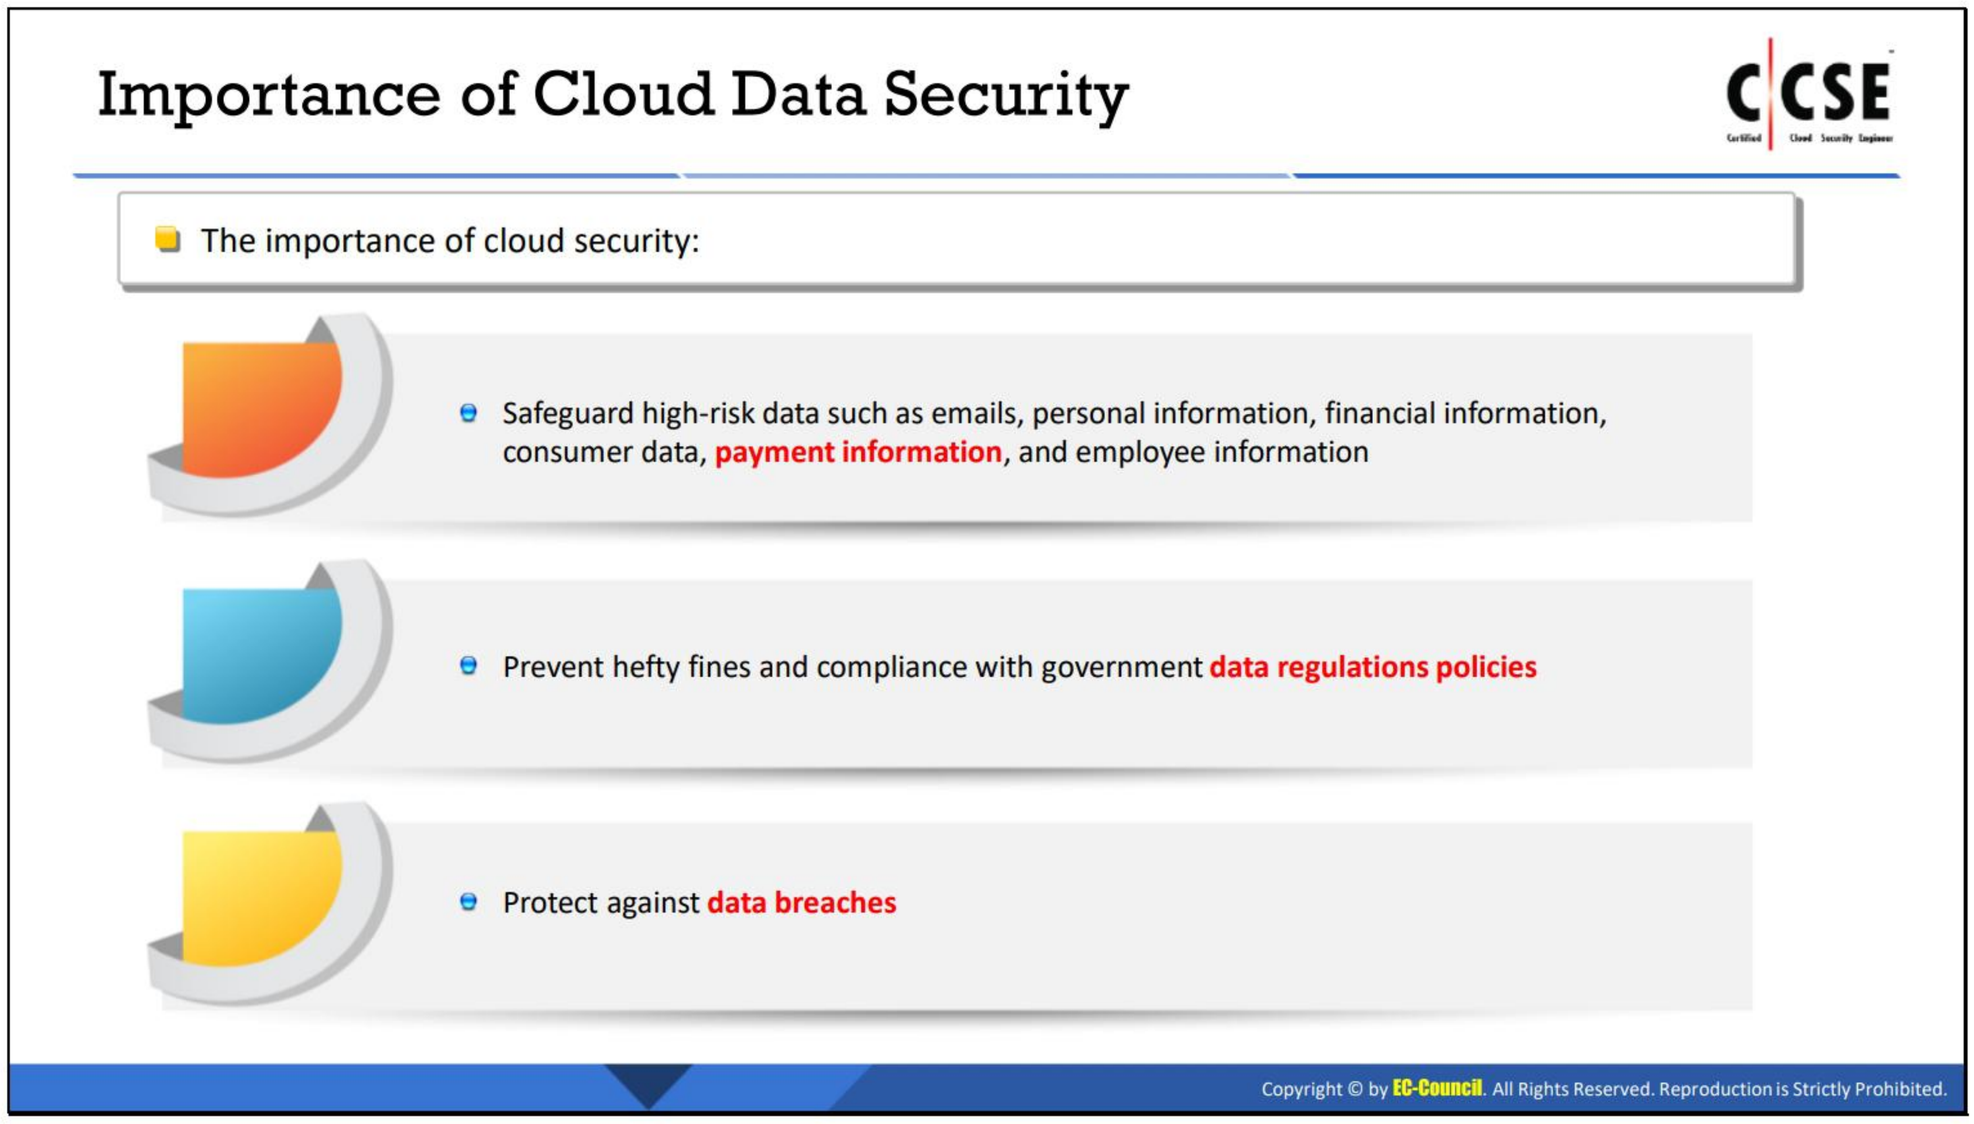
- **Physical security**: Cloud service providers (CSPs) take proper security measures and ensure that their data centers are well equipped with the latest security technologies, 24×7 video surveillance, security fencing, security guards, etc., to protect the data storage facilities. They may also implement geofencing, which involves creating a geographic boundary and detecting suspicious activity by using IP addresses and geolocation data.
- **Advanced firewalls**: CSPs use advanced firewalls that help in securing cloud data by filtering malicious network traffic. Unlike traditional firewalls, advanced firewalls check for each packet's content and then map it with known threats to detect malicious content.
- **Data encryption**: This security measure is critical to secure cloud data from unauthorized access. It is implemented to make data useless and worthless for attackers and people who do not have the secret keys to decrypt it.

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

- **Internal firewalls**: These firewalls are used to limit access for internal users because every user should not have complete access to resources. Internal firewalls even help in protecting resources when an account gets compromised as that account would only have limited access.
- **Event logging**: Event logging is implemented by providers to help security analysts detect suspicious activities and threats while auditing. All network actions are logged in event logs, which helps analysts in detecting and prevent security breaches.
- **Data-aware filtering**: This measure helps organizations to look for particular events and conditions that the users accessed, what kind of information, and when they did it. Data-aware filtering can be attached to role-based authorizations and privileges.
- **Backup and recovery functions**: CSPs offer backup and recovery capabilities to organizations to help them sail through tough times such as when they suffer from disasters, ransomware attacks, data corruption, and deletion.
- **Policy-based lifecycle retention**: CSPs employ data classification policies in their storage systems to effectively automate and manage data storage, retention, archival, and deletion.
- **Appropriate access controls/authentication**: To protect cloud data, CSPs enforce sufficient access controls and multi-factor authentication (MFA). This security measure prevents any data misuse and resource exploitation by internal or external threat actors.

hide01.ir

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



Importance of Cloud Data Security

Securing data in the cloud is very important because it contains important and sensitive information of cloud consumers. The CSP is responsible for providing proper data protection so that customers can store it securely. As companies are increasing their user-base, they are moving their data to the cloud. This data includes personal, financial, and highly sensitive data of their customers, which should be secured at any cost. It is important for organizations and CSPs to secure their data on the cloud from security breaches, data theft, etc., to avoid regulatory fines and legal issues.

Described below are some reasons that show the importance of cloud data security.

- **High-risk data**: High-risk data such as emails, personal information, financial information, consumer data, payment information, and employee information are safeguarded with cloud data security.
- **Government regulations**: Government has various data-regulation policies for companies that deal with data. Companies are liable and accountable for protecting their customers' personal information, financial information, etc., as per various regulations. Cloud data security helps prevent hefty fines and enables compliance with government data-regulation policies.
- **Data breaches**: Organizations must implement cloud data security to avoid security loopholes that may result in data breaches. Data breaches not only defame organizations but also hurt them financially due to regulatory fines. Cloud data security helps protect against data breaches.

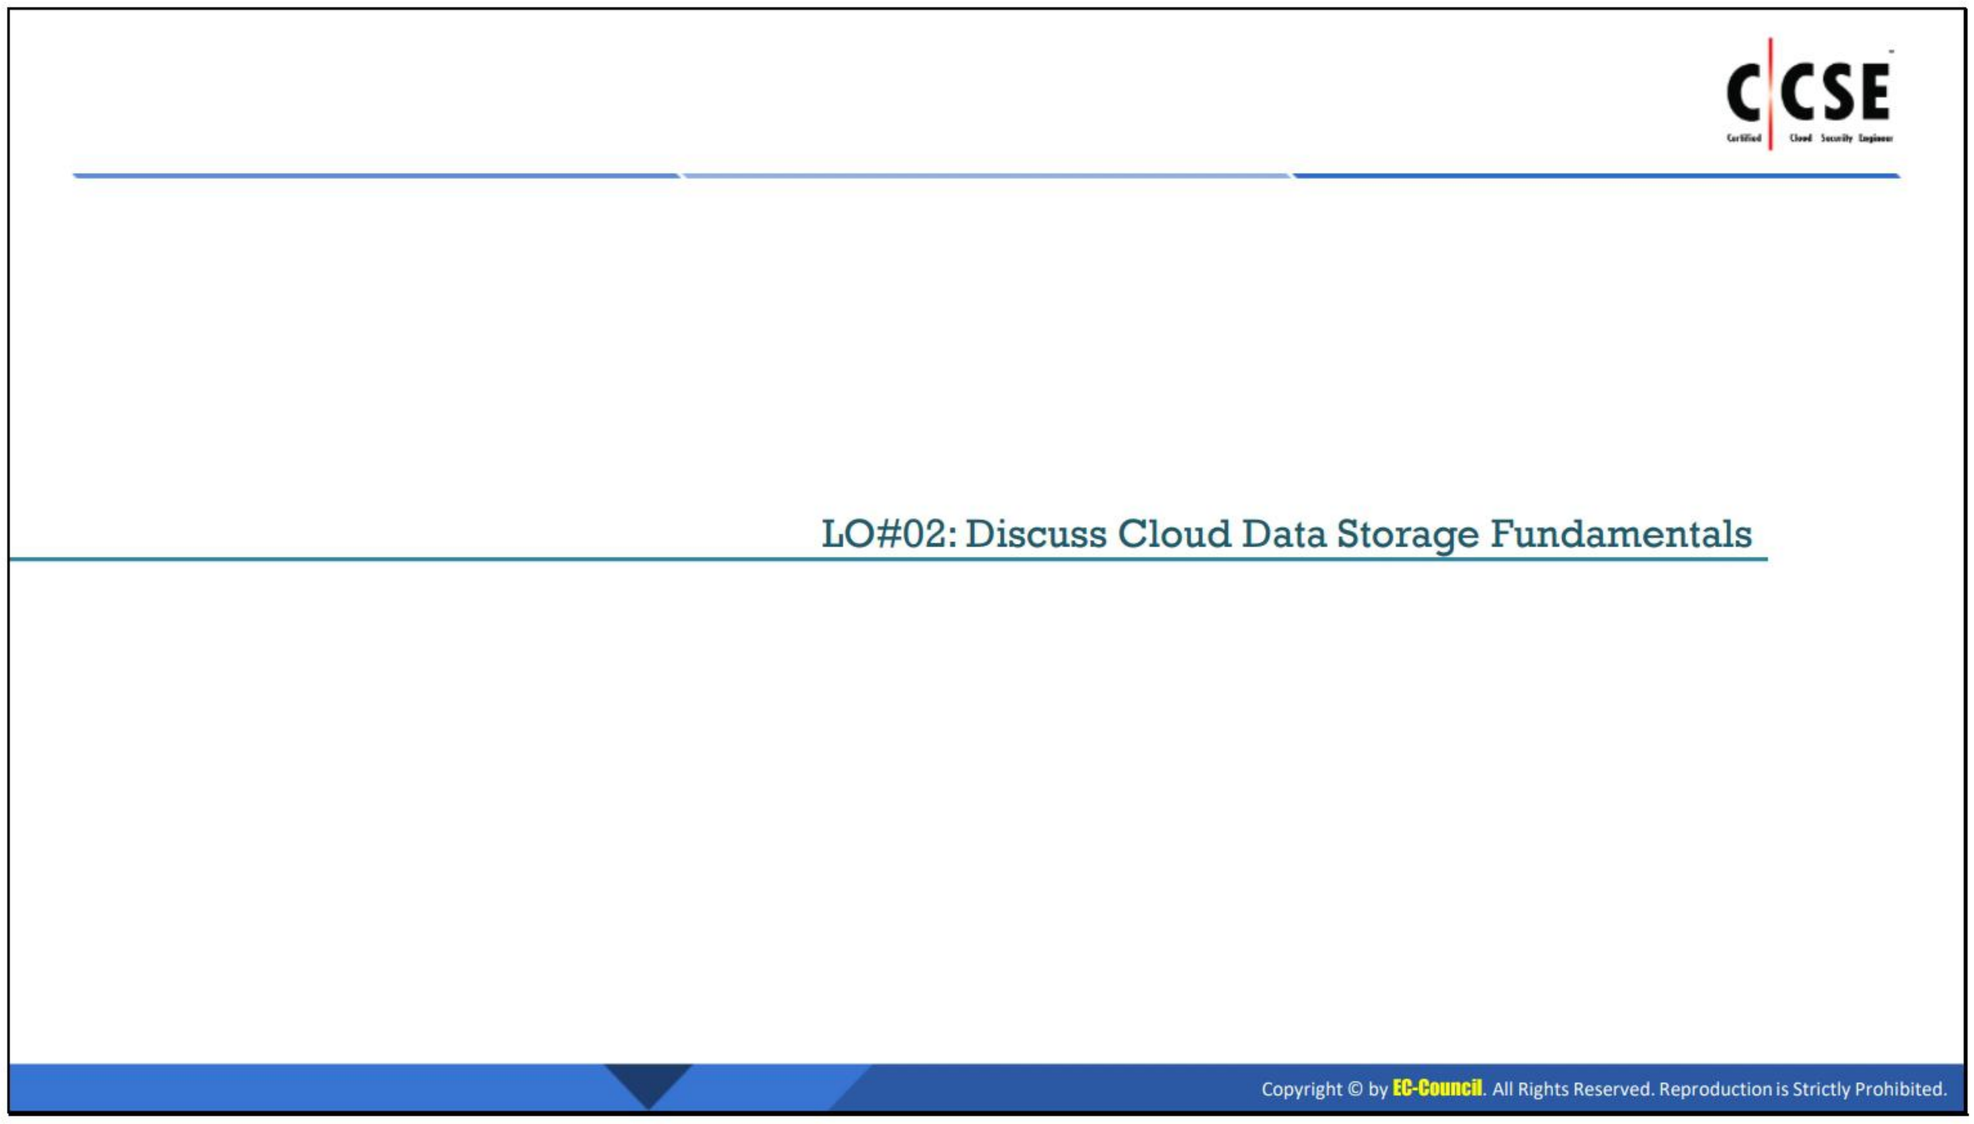
Certified Cloud Security Engineer Copyright © by **EC-Council**
All Rights Reserved. Reproduction is Strictly Prohibited.

Cloud data security allows organizations to do the following:

- Detect and prevent data disruption and loss
- Define security policies
- Detect and mitigate risks
- Better access control
- Secure data and applications across many environments with complete visibility into all activities

hide01.ir

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

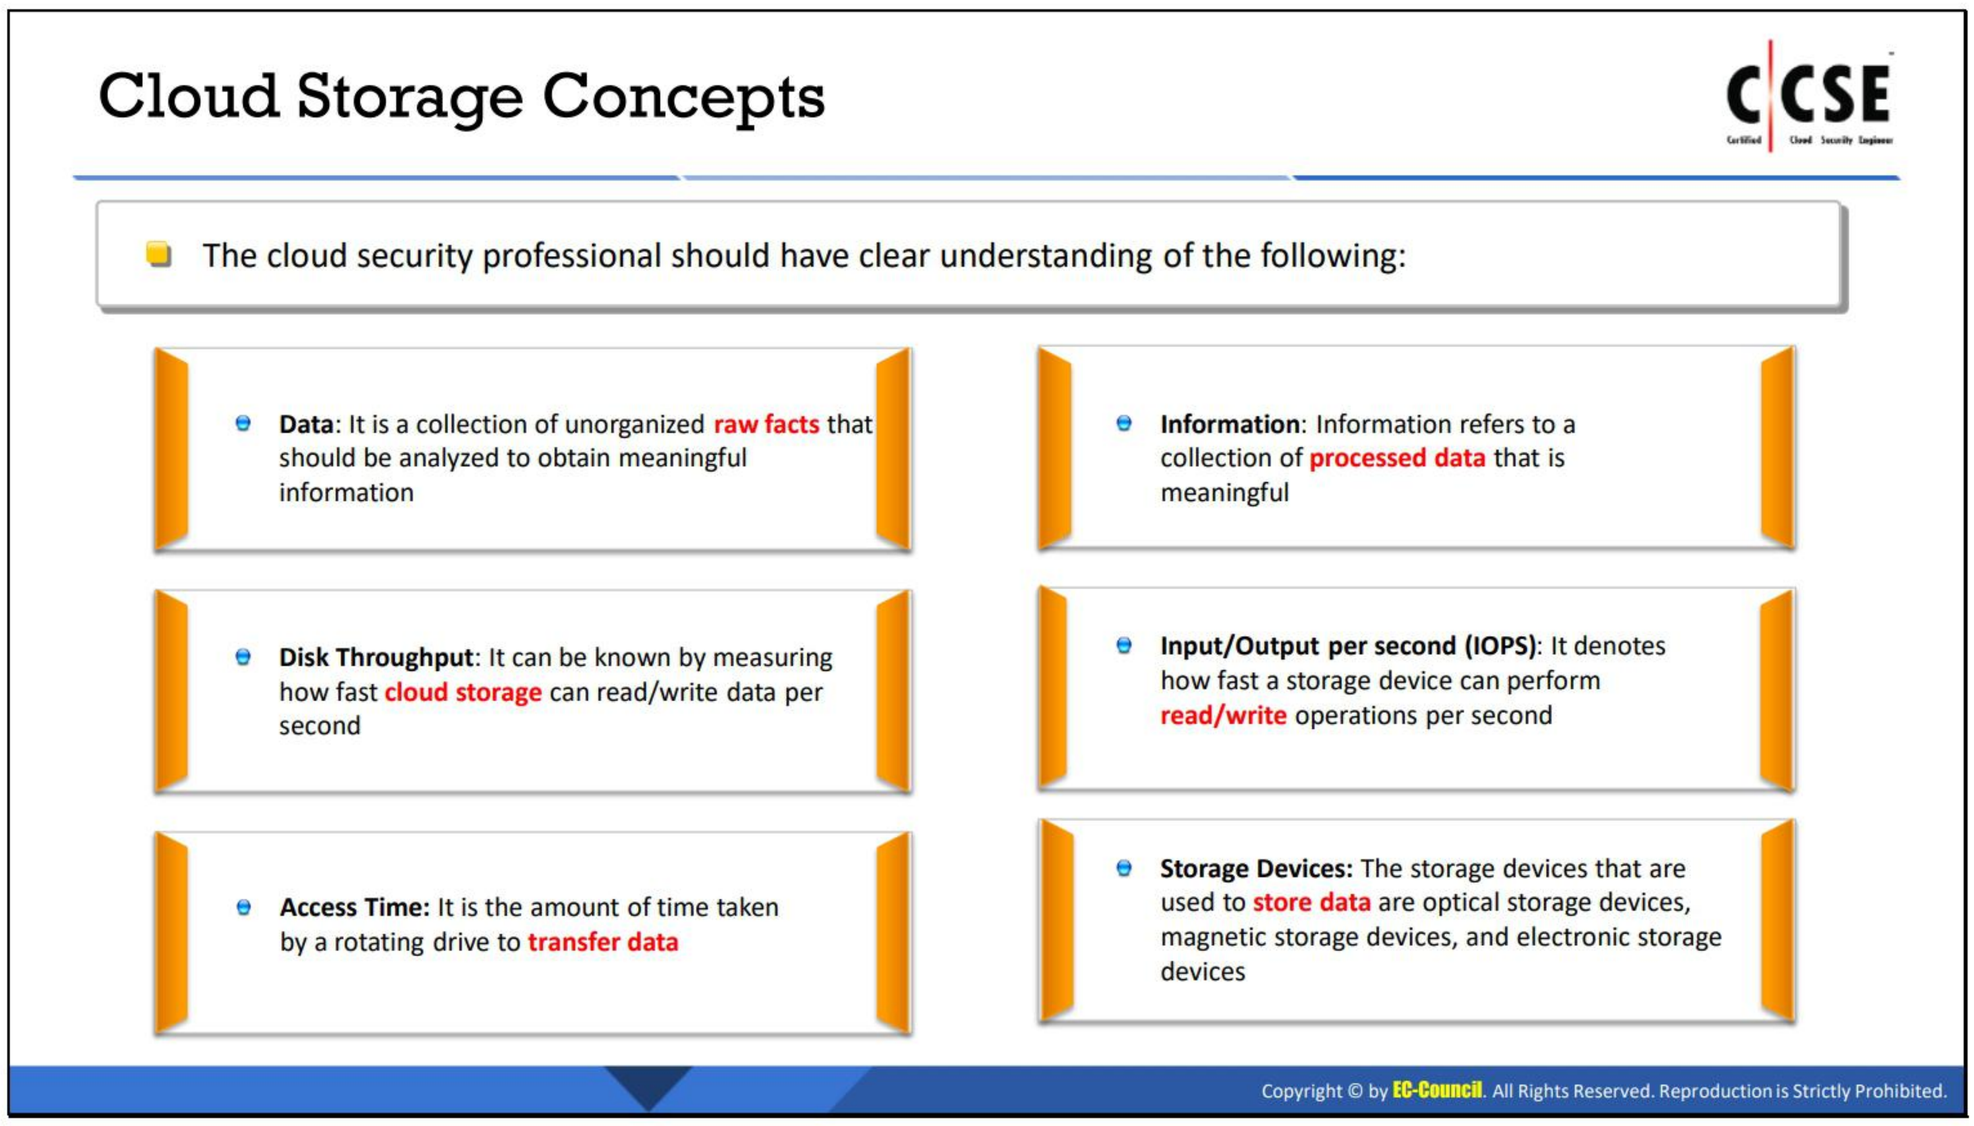


LO#2: Discuss Cloud Data Storage Fundamentals

To implement data security controls in the cloud, it is important to understand the various storage techniques that are followed in cloud. This section focuses on the basics of cloud data storage, including cloud data storage types, cloud storage techniques, cloud storage mechanisms, and cloud storage design patterns.

hidoo1.ir

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



Cloud Storage Concepts

The cloud security professional should have a clear understanding of the below concepts.

Data

ide01.ir

It is a collection of unorganized raw facts that can be analyzed to obtain meaningful information. Data does not have any meaning, unless it is organized. Data consists of numbers, observations, symbols, facts, etc., in a raw format. To derive meaning from data, a machine or human needs to interpret it. Data cannot help in making decisions directly as it is low-level knowledge that is measured in bits and bytes.

There are three types of data structures used for the analysis of data.

- **Structured data:** Structured data is well organized and can be easily interpreted by machine language. It is often referred to as quantitative data and its examples include addresses, names, dates, credit card numbers, etc. This type of data can be organized in relational databases and can be managed by using a programming language called SQL.
- **Unstructured data:** Unstructured data is not easy to analyze as it does not follow a predefined data model. Unlike structured data, unstructured data cannot be understood using traditional programs because it is not organized. Video files, audio files, etc., are examples of unstructured data. Most data generated today is unstructured and cannot be grouped or organized in relational databases. It can be managed by using NoSQL databases and non-relational databases.
- **Semi-structured data:** Semi-structured data can be analyzed more easily than unstructured data as it contains tags or markers that segregate semantic elements and apply hierarchies of fields and records within the data. However, it does not follow the

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

formal data structure models associated with relational databases. XML and JSON are examples of semi-structured data.

- **Metadata**: This is not a type of data structure but an important element for Big Data analysis. It offers extra information about a particular data set. For example, in a photograph, metadata describes when and where the photograph was taken. Metadata essentially provides information about data, so it is widely used by Big Data solutions.

Information

Information refers to a collection of processed data that is meaningful. Data should be processed, organized, or interpreted to consider as information. Information is processed data that has some context, purpose, and relevance.

It is a refined or processed form of data that can be structured as ideas, languages, and thoughts. As information conveys some meaning, it can be directly used to make decisions. It is measured in units such as quantity and time. Information can be used to make meaningful decisions when it is accurate, complete, and timely available.

Disk Throughput

Measurement of the amount of data that travels through a network path is called throughput. Disk throughput can be known by measuring how fast cloud storage can read/write data per second. It measures the speed of reading/writing data by cloud storage. Disk throughput has a direct impact on storage performance.

hide01.ir

Input/Output per Second

Input/output per second (IOPS) is a measurement that determines the performance of storage area networks (SANs) and hard drives (SSD/HDD). It denotes how fast a storage device can perform read/write operations per second. It is a less popular measurement of hard drives, but it is a critical indicator that determines their performance under heavy use.

Access Time

The amount of time taken by a rotating drive to transfer data is called the access time of the rotating drive. Factors related to the rotating disk and its moving head control the access time.

- **Average seek time:** Seek time is the amount of time required by the hard-disk controller for locating a piece of data stored on the drive. Seek time is affected by delays such as transfer time and rotational delay or latency. Whenever data is read or written to the disk drive, the read/write head is required to move to the correct position. Seeking refers to the actual physical position of the disk and the time taken by the head assembly to move from one part of the disk to another. Seek time for a particular disk depends on the distance between the start point and the destination point on the disk.

  Seek time is normally calculated as an average seek time due to these variations. A PC's average seek time has been around 9 ms since 2004. Optical drives such as CDs, floppy drives, and DVDs have much slower seek time due to their larger head construction when compared to a hard disc drive. The average seek time of a rotating drive is calculated by

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

taking the ratio of the time required to do all possible seeks to the number of possible seeks. An acceptable seek time for a hard disk is less than 10 ms.

- **Full-stroke seek time**: It is the amount of time taken by the read/write head to seek or search the entire disc. It is measured in milliseconds and can be between 15 ms to 20 ms. This time is the longest as the head covers the entire width of the disk. It is also called as the maximum seek time.
- **Track-to-track seek time**: It is the amount of time taken by the read/write head to seek between two adjacent tracks. It is measured in milliseconds and can be between 2 ms to 4 ms or 1 ms at the lowest. It is also known as the minimum seek time as the head just moves from one track to another.

Storage Devices

The various storage devices used for storing data are described below.

- **Optical storage devices:** In optical storage devices, light is used to record data. Recording of data takes place by using a system that can write/read to a removable disk and a drive contains that disk. Popular examples of optical storage devices include DVDs, CDs, and Blu-ray discs. These storage devices are lightweight, cheap, and removable. It uses a laser beam to read/write data on devices digitally.

  Storage devices that are marked with "R" can be written once and those with "WR" can be written many times. For example, a DVD-R can be written only once whereas DVD-WR can be written multiple times. In this type of storage, data is sequentially accessed and the storage media are suitable for mass replication. They provide an excellent S/N ratio and a high sample rate.

- **Magnetic storage devices**: Magnetic storage involves storing data on a magnetic medium. In magnetic storage devices, data is read, written, or accessed on the devices through a magnetization process. Popular examples of magnetic storage devices include disk drives, magnetic tapes, and floppy disks. Hard-disk drives are coated with magnetizable medium and run at a very high speed. They are an example of disk drives. Floppy disks are portable storage devices that store limited data. Magnetic tapes are coated with magnetizable elements and have high storage capacity.

  Magnetic storage devices provide a low sample rate and intermediate S/N ratio. In magnetic disks, data can be accessed randomly and only one disk can be used/ accessed at a time.

- **Electronic storage devices**: Electronic storage involves storing data electronically on any digital medium. In electronic storage devices, data is stored and retrieved by using electrical power. Popular examples of electronic storage devices include hard disks, compact disks, audiotapes, magnetic tapes, and videotapes.

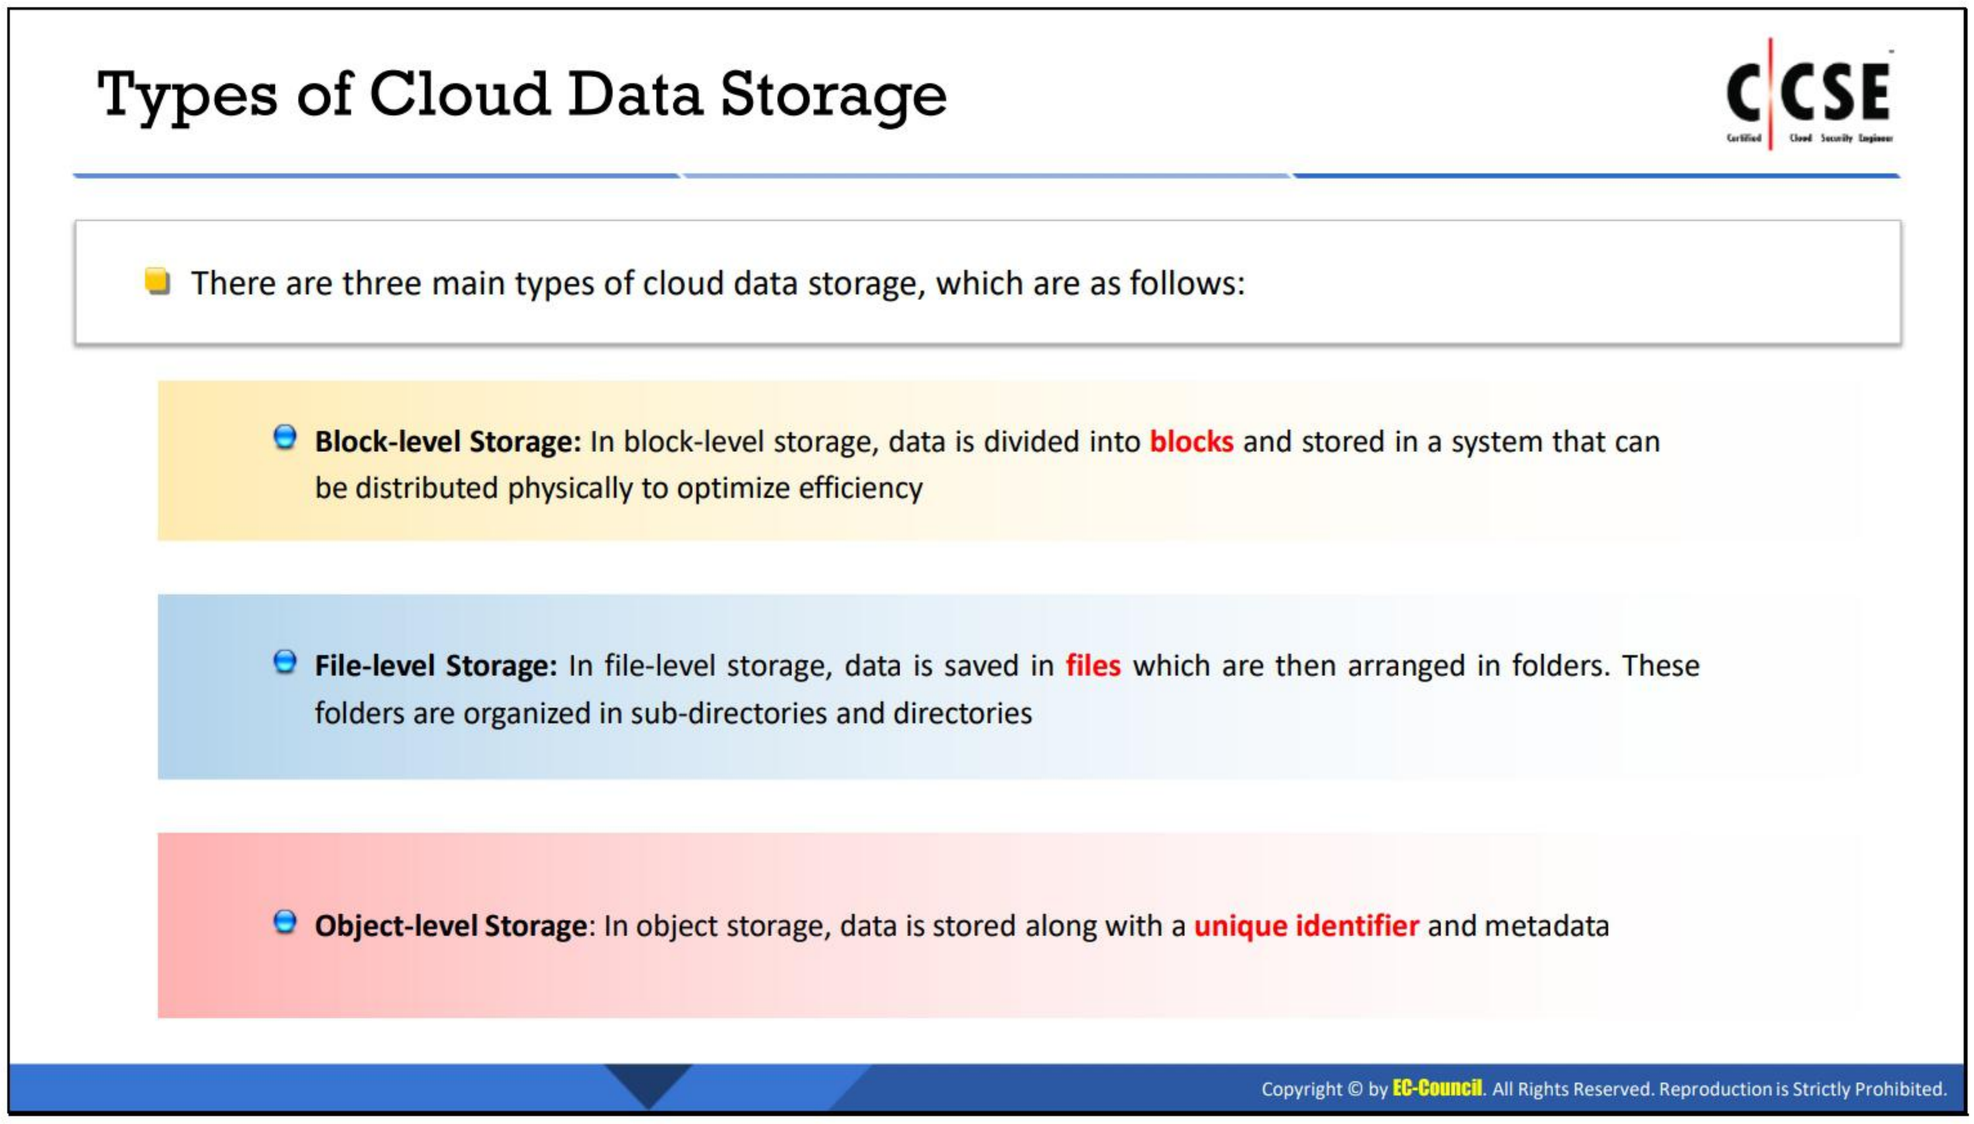
  In flash memory, content can be written and removed by the application of a flash of high voltage pulse at the byte level. Flash memory is electronically erasable programmable read-only memory (EEPROM) that can retain data when there is no power. It uses floating-

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

gate transistors to store a bit of information. It finds its application in multiple devices such as pen drives, SD cards, etc.

hide01.ir

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



Types of Cloud Data Storage

There are three main types of cloud data storage, as described below.

- **Block-level storage**: In block-level storage, data is divided into blocks and stored in a system that can be distributed physically to optimize efficiency. This storage type stores structured data or workloads. A unique identifier is assigned to each block, which helps in data retrieval from the block. It is normally used in SANs. Block storage works well in situations where quick data retrieval is necessary. Block storage is expensive as it caters to larger workloads by providing better scalability and improves IOPS.

  Due to its fast and reliable performance, using block storage is preferred for databases and applications. Block storage primarily finds its use in

  - Database storage
  - Storage for Redundant Array of Independent Disks (RAID) volumes
  - Storing data for critical systems that affect business operations
  - Email servers that run applications such as Microsoft Exchange

- **File-level storage**: In file-level storage, data is generally stored in files that are then organized in folders and these folders are arranged in sub-directories and directories. Users can access a file using a path as the files are arranged hierarchically. In this storage type, naming, modifying, or deleting files is easy because all files are in an organized manner. They are normally used for storing unstructured data and in network-attached storage systems. Devices using this storage type can scale storage up and down as per requirements as file-level storage is dually scalable. They utilize server message block

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

(SMB) and common internet file system (CIFS) protocols for Windows and network file storage (NFS) for Linux.

The main advantage of this storage type is its approachability. In a cloud environment, the administrator can easily manage access and editing permissions for files using file storage. This enables easy access to sharing and collaboration. It is less costly when compared to block storage. File storage is commonly used to

- Store files in a content repository for an office
- Store files in a small data center environment
- Store data that needs the easy deployment and data protection

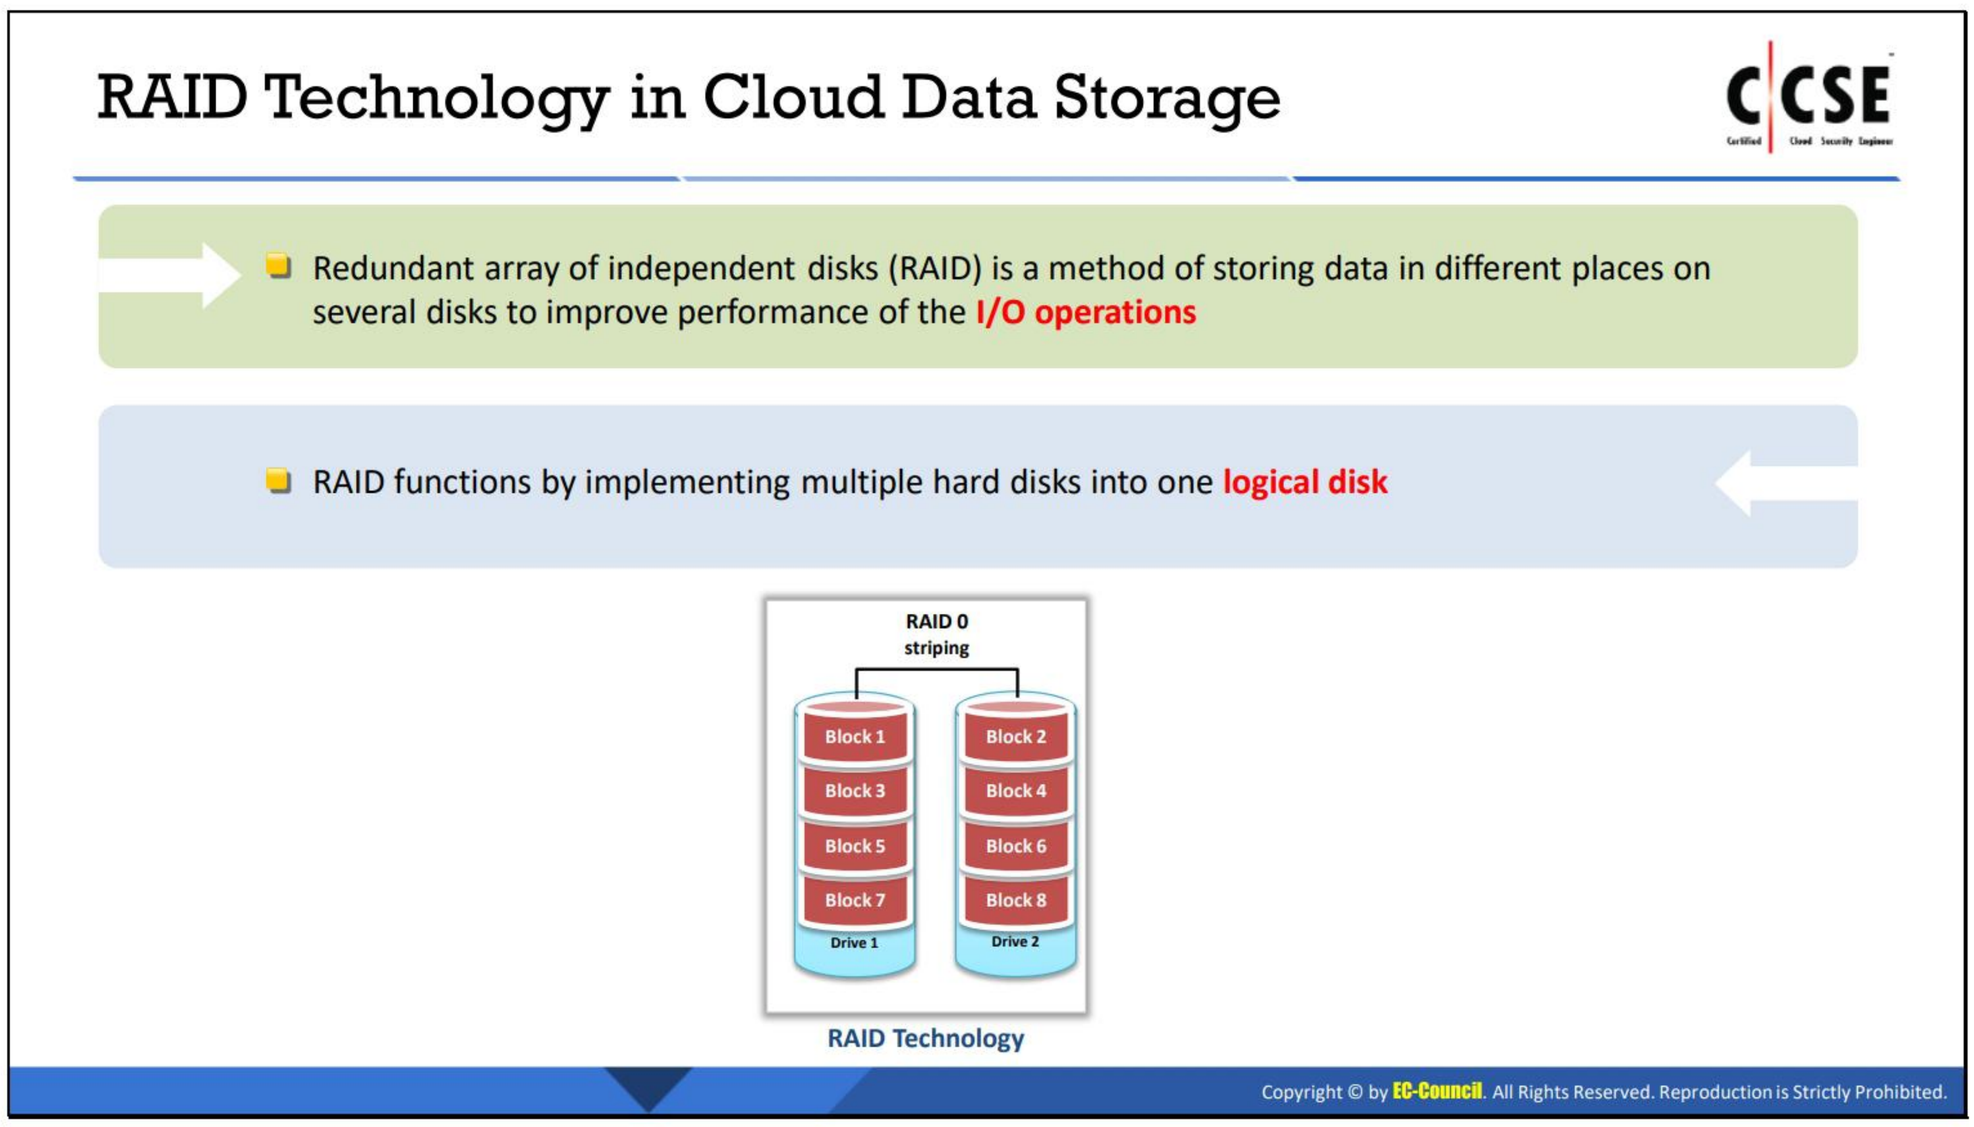
- **Object-level storage**: In object storage, data is stored along with a unique identifier and metadata. Data can be anything, for example, a file, a photo, audio, etc. When it is stored in object storage, an identifier is added to it which acts as its name. This identifier is a 128-bit integer and is also called a globally unique identifier (GUID). Metadata is a component that is attached to an object to classify or characterize it. Metadata can be any type of information that helps in classifying the object.

  Object storage is used commonly by many cloud storage providers because of its ability to store huge amounts of unstructured data along with easy data accessibility. It provides easy accessibility to customers because it uses metadata to classify data. Metadata also helps storage administrators implement their customized policies for data retention, preservation, and deletion because of its unrestricted nature.

  Object storage is mainly used for storing large amounts of data that is occasionally accessed. It is commonly used to

  - Storing large data sets
  - Storing unstructured data
  - Storing large quantities of media assets as an archive

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



RAID Technology in Cloud Data Storage

RAID is a method of storing data in different places on several disks. Storing the data on multiple disks improves performance in terms of I/O operations. RAID functions by implementing multiple hard disks into one logical disk. It allows storing the same data in a balanced way across an array of disks. The effective implementation of this technology helps address the complex issues of fault tolerance. The data organized in RAID levels depends on the RAID storage techniques and installation methods used. Usually, the implementation of RAID is done on a server. Though personal computers do not necessarily need this technology, they can still set up and utilize it in a smaller environment than an enterprise.

RAID technology has the below features.

- **Fault tolerance:** RAID exhibits fault tolerance. If a disk fails to work, other disks will continue to function normally.
- **Performance:** RAID achieves high performance during read and write processes across multiple disks.
- **Competence:** This is defined by the amount of data stored. The storage capacity of the disks depends on the RAID level chosen. The storage capacity may not necessarily be equal to the size of the individual RAID disks.

RAIDs use the below storage techniques.

- **Striping:** Data striping divides the data into multiple blocks. These blocks are further written across the RAID system. Striping improves data storage performance.
- **Mirroring:** Data mirroring makes image copies of the data and simultaneously stores this data across the RAID. This affects fault tolerance and data performance.

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

- **Parity:** Parity uses a striping method to calculate a parity function of a data block. During drive failure, the parity recalculates the function using the checksum method.

The advantages of RAID systems are described below.

- **Performance and reliability:** RAID increases the performance of reading and writing data on disks. The speed of the process is much faster than using a single drive for storage. It improves performance by distributing the I/O. The RAID controller distributes data over several physical drives making sure not to overburden a single drive in the RAID system. RAID ensures reliability of data even if a disk fails. The failed components can be replaced in a RAID system without shutting the system down. This feature is called hot swapping. The replacement process does not affect how the other disks or the network function.
- **Parity check:** Parity check is a process where the RAID system compares the data stored in the crashed system with the data stored in the other disks. This checking process is accomplished on all the drives. The parity check is performed after mirroring the data first. Regularly performing parity checks helps detect the probability of a system crash, thereby preventing loss of data.
- **Data redundancy:** Failure of a disk can occur at any time. Thus, data redundancy is important for the organization. RAID provides enhanced data redundancy in case of a hardware failure.
- **Disk striping:** Disk striping improves the read/write performance of the data. The data is divided into small chunks and spread across multiple disks. Depending on the RAID level implementation, the data is divided into bytes, bits, or blocks. Data reading and writing can be done simultaneously in a RAID system.
- **System uptime:** This is a metric that detects the reliability and stability of a computer. System uptime defines the time the system can be left unattended without any assistance. Configuring RAID on a system helps in enhancing system uptime. High system uptime in an organization signifies that their productivity is high.

h1te01.ir

Types of RAID

Hardware RAID

Hardware RAID uses a physical controller to work and the processing is done on hardware such as on a motherboard or a RAID expansion card. In a hardware-based RAID, logical disks are configured and mirrored on the hardware. A physical controller is located on the PCI bus and it manages the application data and operating system(s). The controller prevents the drives from data loss and enhances the read-write operations.

RAID levels 0, 1, 3, and 5 are compatible with hardware RAID. A hardware RAID provides efficient and non-stop recovery from media failure. Performance-based advantages are much higher with a hardware RAID. For example, the implementation of RAID Level 5 will enhance the data throughput as compared to a software-based RAID. Multiple controllers can be added to improve the read-write performance and total storage capacity.

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

A hardware RAID can be implemented when there is a complex and critical setup or for large databases.

Some advantages of hardware RAID are listed below.

- **Write-back mode**: A typical hardware controller has a battery backup unit (BBU). The hardware RAID can work in the write-back mode because of the BBU feature. If there is power failure while writing data to a drive, the data will not be lost or deleted. The BBU plays a very important role in write-back mode.
- **Hot-swapping**: Many controllers in a hardware RAID support hot-swapping. The disk can be replaced while the server is still running; this avoid any effect on the productivity of the organization.
- **Higher throughput**: With the availability of a BBU, a hardware RAID offers higher read and write throughput, increasing the overall performance of the RAID level.
- **Boot loader**: A hardware RAID can recover from a boot loader failure.

Software RAID

As the name suggests, software RAID uses software instead of hardware for its implementation. Unlike a hardware RAID that uses a controller, software RAID uses system processors and other applications to work. Software RAID is implemented in the operating system or at the kernel level. The performance of a software RAID depends on CPU performance. Software RAID relies on a standard host adapter and executes all I/O commands using mathematical calculations. RAID levels 0, 1, and 5 are compatible with software RAID.

Listed below are the advantages of software RAID.

- **Cost-effective**: Software RAID is part of the operating system. There are no additional items needed. It is more cost-effective than a hardware RAID implementation.
- **Simplicity**: A software RAID does not need a hardware controller. There are no complexities for its implementation.
- **Duplexing**: Duplexing in a software RAID requires only a standard controller for the process.

Levels in RAID

Different RAID levels are used for different situations and requirements. Each level can be used for a specific situation where it enhances the reliability and performance of data storage.

- **RAID 0: Disk striping**

  RAID 0 splits data into sections or blocks and writes it across multiple drives. The storage capacity of RAID 0 is equal to the sum of the disk's capacity in the set. The data distribution in RAID Level 0 is equal among all the disk sets, resulting in high performance. With concurrent high performance, the throughput of the read and write operation on multiple disks is equal to the throughput of the array of disks.

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

RAID 0 does not provide fault tolerance. Failure of one disk can lead to the failure of all the disk in a Level 0 volume. The probability of recovering data from a RAID Level 0 system is minimal. Both software and hardware RAID controllers support RAID 0, which can be used to help boost server performance.

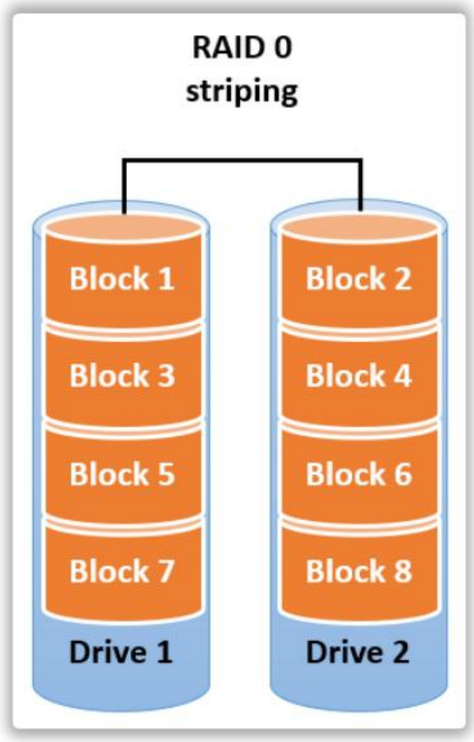


Figure 4.1: RAID 0 Disk Striping

RAID 0 provides the below advantages.

- **Read/write performance:** RAID Level 0 has a very good read and write performance. The performance is even greater if the controller supports independent reads and writes to different disks in the array.
- **Implementation:** RAID 0 is easy to implement as the data is divided into a sequential set of blocks. There is no storage loss as the maximum capacity is used.
- **Cost:** RAID Level 0 is cost-effective compared to the other RAID levels.

RAID 0 systems are ideal for non-critical cloud data storage that needs fast reads/writes. They can be used on video editing stations and the public cloud.

- **RAID 1**

A typical RAID 1 system contains an exact copy of the data on two or more disks. RAID 1 writes data on multiple drives and multiple mirrors drives at the same time. Failure of one drive does not affect the data on the other drives. This allows data retrieval from the mirror drive and ensures reliability.

In a RAID 1 hardware implementation, a minimum of two disks is required. In a software RAID 1, data can be copied to a volume of the disk.

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

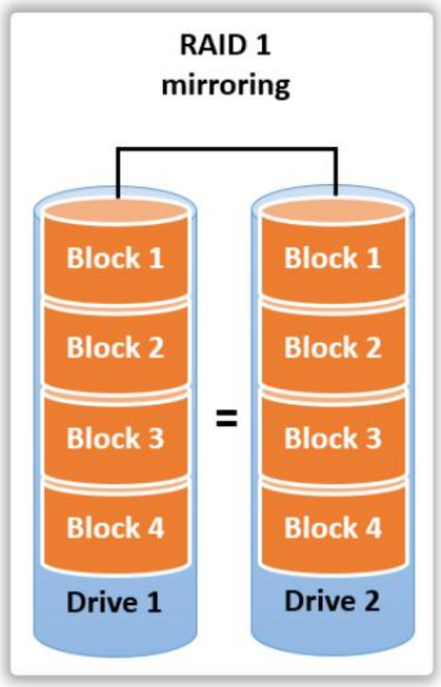


Figure 4.2: RAID 1 Mirroring

RAID 1 provides the below advantages.

- **High read performance:** Because there are two disks, the read performance is higher in a RAID Level 1 system. Data can be read simultaneously while it is being written on the other disk.
- **Reliability:** The mirroring feature in RAID 1 ensures that data will be available even in case of a failure, making it more reliable than RAID 0.
- **Compatibility:** RAID 1 is compatible with hardware and software RAID systems, including controllers.
- **Simplicity:** RAID 1 systems utilize simple technology.

RAID 1 is suitable for environments where read performance matters more than write performance. RAID 1 can be used for mission-critical cloud storage such as accounting, payroll, and other financial applications.

- **RAID 3: Disk striping with parity**

RAID 3 uses striping and parity as its main feature to store the data. To implement a RAID Level 3 system, a minimum of three disks is required. The data is stored on multiple drives at the byte level. This RAID level dedicates one drive to store the parity information. The byte-level division allows the drives to work simultaneously. At any time either a read operation or a write operation can take place.

RAID Level 3 shows a high transfer data rate along with data security. It can perform data recovery and error correction by calculating an exclusive OR (XOR) of the information recorded on the parity drive. RAID 3 is a good choice for specialized cloud databases or single-user systems.

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

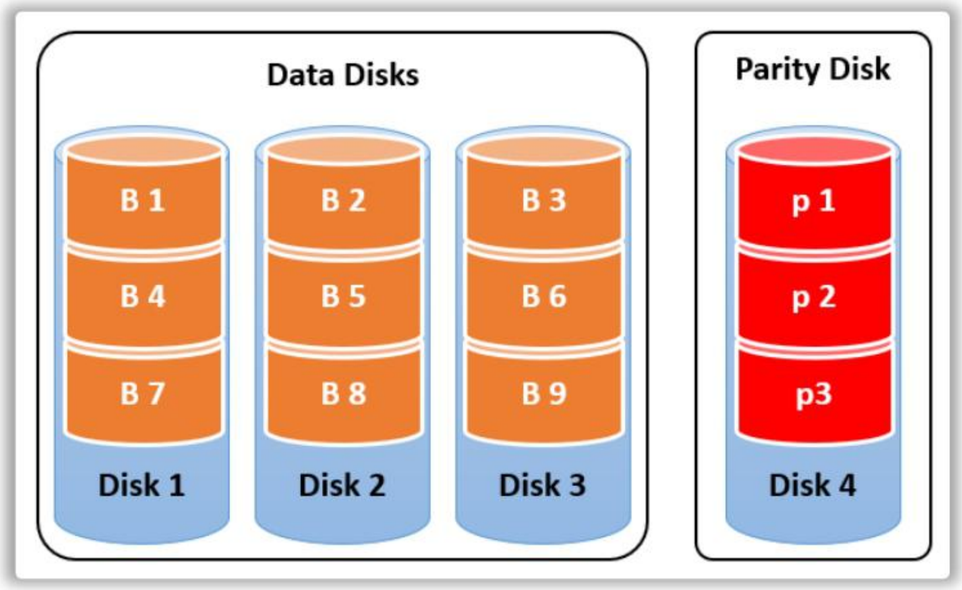


Figure 4.3: RAID 3 Disk Striping with Parity

RAID 3 offers the following benefits:

- **High throughput:** RAID Level 3 provides high throughput for reading and writing operations of large data transfers.
- **Resistant:** This RAID level is resistant to disk failure and breakdown.

- **RAID 5: Disk striping with distributed parity**

  RAID Level 5 includes block-level striping with distributed parity—the parity information is distributed among all the drives, except one drive. The data chunks in a RAID Level 5 system are larger than the regular I/O size, but they can be re-sized. To prevent data loss after drive failure, data can be calculated from the distributed parity.

  RAID 5 needs at least three disks, but for better performance, more than three disks (up to 16 disks) can be used. If a disk fails, it takes a long time to rebuild a RAID 5 array. When the array is being built again, performance can degrade, making it vulnerable to additional disk failure. This level offers significant read performance as the disks satisfy the data requests independently. RAID 5 is not a good choice to write operations on the system.

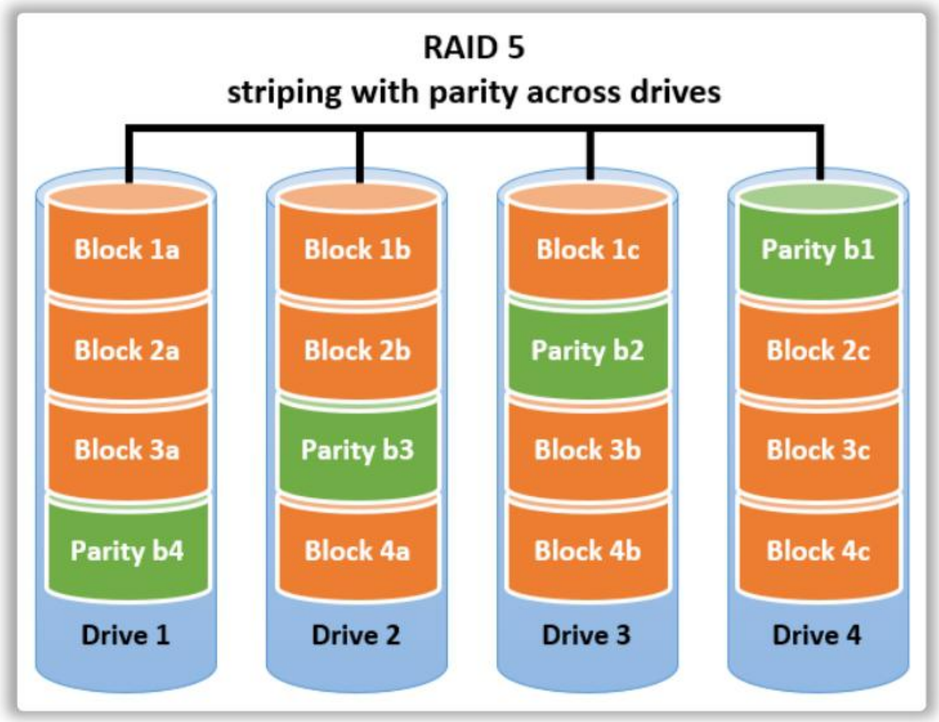


Figure 4.4: RAID 5 Disk Striping with Distributed Parity

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

RAID 5 offers the following advantages:

- **Fast read speed:** Among all the RAID levels, level 5 offers the highest read data transaction rates.
- **Resilience:** RAID 5 can withstand the failure of a single drive without loss of data.
- **Hot swapping:** In case of a disk failure, the failed disk can be replaced with a new one without a server shutdown.

RAID 5 is found most often in cloud file and application servers, database servers, web, email, and news servers. RAID Level 5 systems provide excellent security and good performance as well as efficient storage.

- **RAID 6: Disk striping with double parity**

RAID 6 works like RAID 5 but uses dual parity along with block striping. It stores parity data in two different drives. The RAID Level 6 system requires a minimum of 4 disks to operate and stores an extra parity on every disk. RAID 6 systems can run even if two disks fail to work.

It provides higher redundancy and better read performance than RAID 5. Unlike RAID 5, it can withstand the failure of two disks at the same time. Data will not be lost in this situation. RAID 6 systems are suitable to be used in large storage arrays and cloud application servers. The configuration of this RAID level is difficult to implement in a RAID controller because it needs to calculate two parities for every data block.

hide01.ir

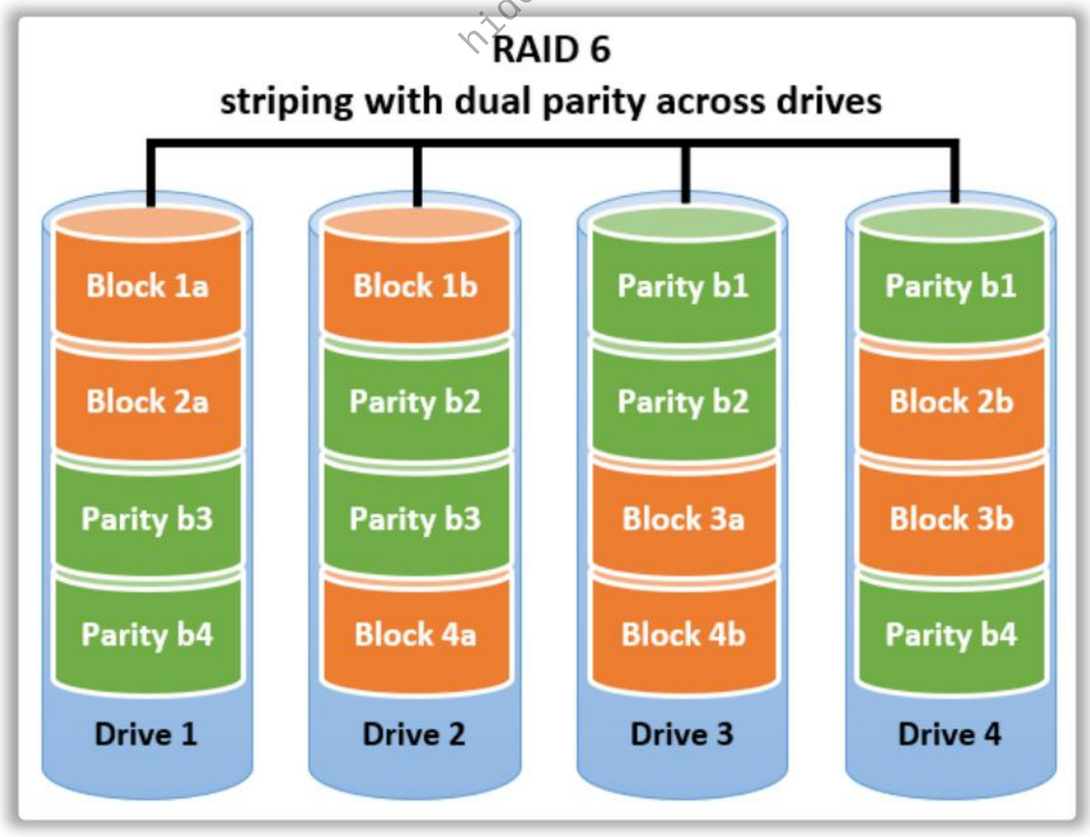


Figure 4.5: RAID 6 Disk Striping with Double Parity

The benefits offered by RAID 6 are listed below.

- **Fast read speed:** It offers a fast read speed like RAID 5.
- **Data security**: It offers better security as data is not lost while replacing damaged disks in a case where two disks fail at the same time.

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

- **RAID 10: Blocks striped and mirrored**

  RAID Level 10 includes disk striping and mirroring in a nested hybrid RAID level. It is a combination of RAID 1 and RAID 0. The level can symbolically be represented as RAID 1+0 or RAID 10. RAID 10 includes the mirroring of RAID 1 without the parity and striping of RAID 0. It requires a minimum of four drives for its operation.

  The performance of RAID 10 is higher than a RAID 1. RAID Level 10 has the same fault tolerance as RAID Level 1. RAID 10 is a great choice for cloud database servers, web servers, email servers, etc., and can be used on hardware or software RAID implementations.

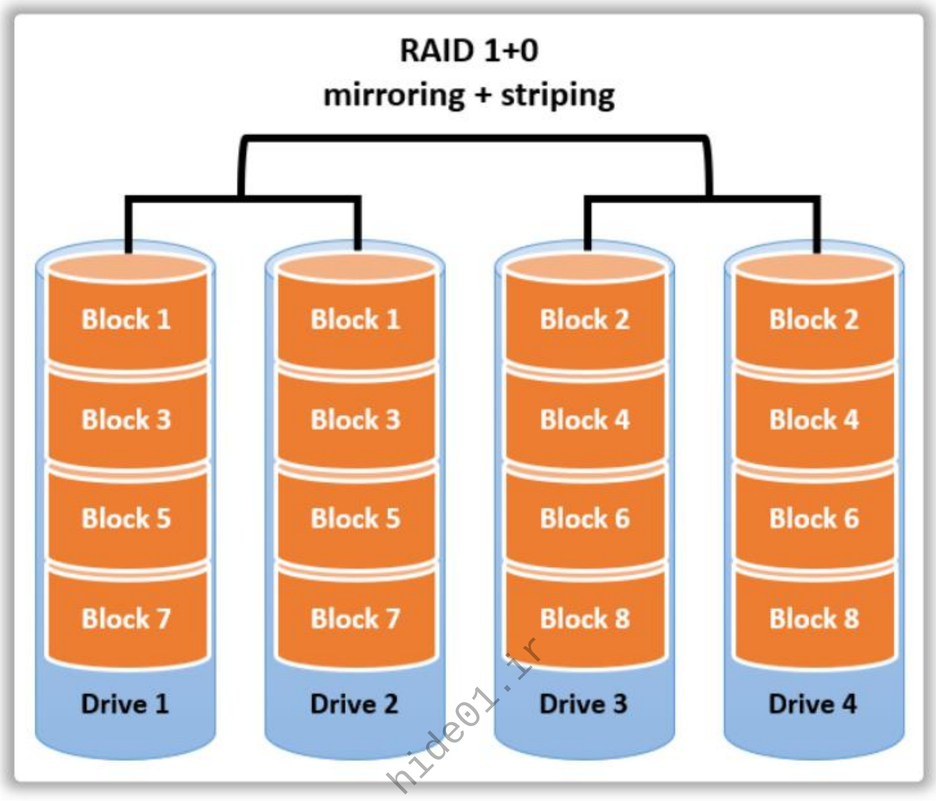


Figure 4.6: RAID 10 Blocks Striped and Mirrored

RAID 10 offers the below benefits.

- **High operations:** Owing to the combination of RAID Levels 1 and 0, it provides high I/O operations.
- **Better throughput:** Compared with other RAID levels, RAID 10 provides better throughput and higher latency.
- **Efficient write operations:** The write operations of this level are efficient; thus, this RAID level is often implemented on cloud database servers and other cloud servers performing write operations.

- **RAID 50: Striping and mirroring across multiple RAID levels**

  RAID Level 50 includes mirroring and striping across multiple RAID levels. This level is a combination of the block-level striping of level 0 and the distributed parity of level 5. The configuration of RAID Level 50 requires a minimum of six drives. This level undergoes a hot-swapping process when a disk fails.

  RAID 50 is an improvement over RAID 5, specifically for its write operation and fault tolerance. RAID Level 50 can be implemented on cloud servers that run applications

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

requiring high fault tolerance, capacity, and random-access performance. This level offers data protection and faster rebuilds compared to a RAID 5 system. When one disk fails in a segment, it only affects that segment and not the entire array. Only that segment is rebuilt. The rest of the array functions normally.

The benefits offered by RAID 50 are listed below.

- **Security:** The data stored in a RAID 50 is more secure than in a RAID 5. With a larger storage capacity, this level offers more data storage than RAID 5.
- **Non-degradable:** With the use of a minimum of six drives in the configuration environment, failure of one disk does not impact the server function configured on this level.
- **Read and write performance:** The read and write performance of RAID Level 50 is far better than RAID Level 5.

hide01.ir

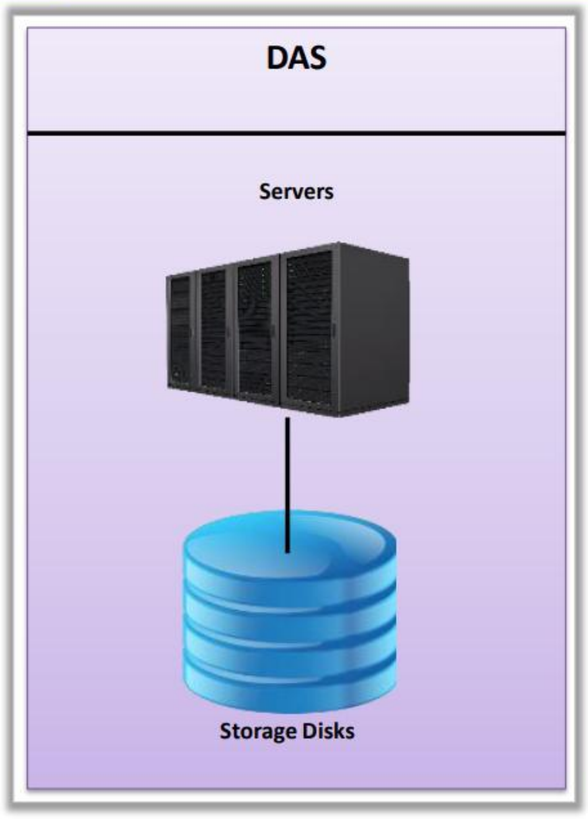
Certified Cloud Security Engineer Copyright © by **EC-Council**
All Rights Reserved. Reproduction is Strictly Prohibited.

Storage Techniques in Cloud

- The storage techniques that are widely used in cloud computing to store business are
  - Data DAS (Direct-attached storage)
  - NAS (Network-attached storage)
  - SAN (Storage area network)

Direct-attached Storage (DAS)

- It consists of **storage device** that directly connects to the computing device through a physical cable



Direct-attached Storage

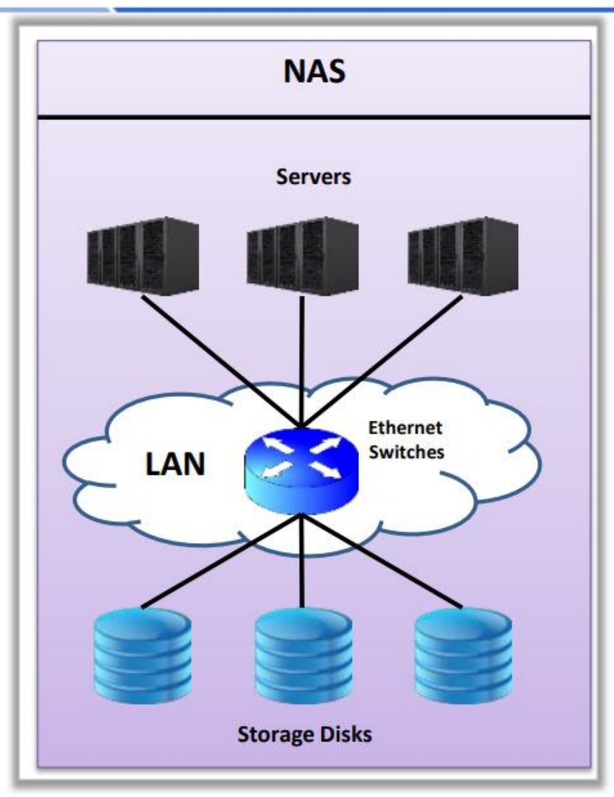
Copyright © by EC-Council. All Rights Reserved. Reproduction is Strictly Prohibited.

Storage Techniques in Cloud (Cont'd)

hide01.ir

Network-attached Storage (NAS)

- It is a storage device that can be attached directly to a network and allows data storage and retrieval from a **central location**
- The protocols used in NAS are:
  - **Common Internet File Services/Server Message Block (CIFS/SMB):** This protocol is utilized in **Windows**
  - **Network File System (NFS)**: This protocol is utilized in **UNIX servers**, and it is the common Linux protocol



Direct-attached Storage

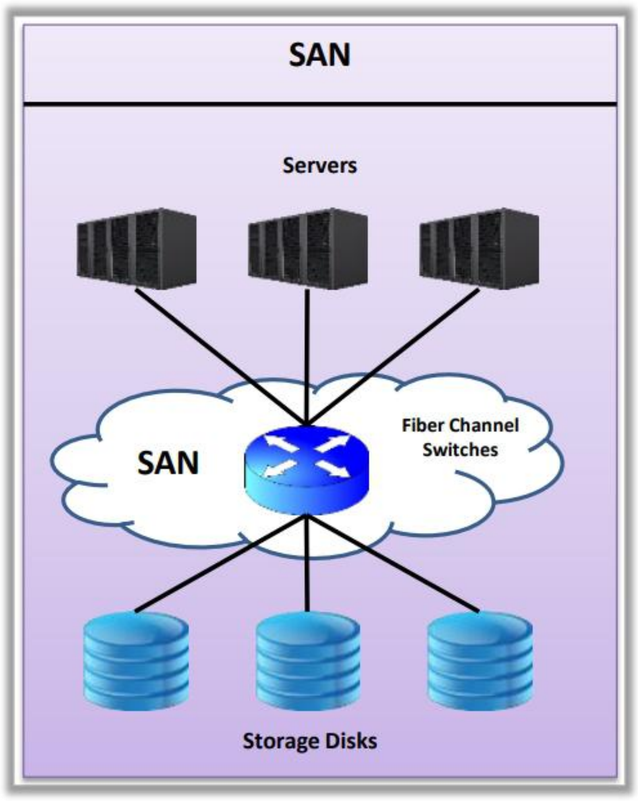
Copyright © by EC-Council. All Rights Reserved. Reproduction is Strictly Prohibited.

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

Storage Techniques in Cloud (Cont'd)

Storage Area Network (SAN)

- **SAN** is an independent **high-speed network** that offloads data from servers/desktops to the storage devices
- The important **protocols** utilized in SAN are FCP, iSCSI, FCoE, etc.
- The following two types of SAN have been used extensively for storing data in the data center:
  - **IP SAN:** IP SAN is a storage area network that uses **internet protocol** (IP) to send the storage traffic instead of fiber channel cables
  - **FC SAN**: FC SAN allows **multiple servers** to connect to or access network storage devices through a fiber channel



Storage Area Network (SAN)

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

Storage Techniques in Cloud

The storage techniques that are widely used in cloud computing to store business data are described below.

hide01.ir

Direct-attached Storage (DAS)

Direct-attached storage consists of storage device that directly connects to the computing device through a physical cable. The stored data in DAS can be accessed through block-level protocol.

This type of storage is directly connected to the computing device (a computer, laptop, etc.) and can be accessed by that device only. Common examples of DAS are flash drives, hard drives, solid-state drives, etc.

DAS in an enterprise can be external drives that are connected directly to the server via serial advanced technology attachment (SATA), small computer system interface (SCSI), etc. Optical storage devices and tapes, HDDs, SSDs are also DAS because they are connected to the system directly.

DAS devices provide better performance than networked storage devices because data does not have to travel when a read/write operation is performed. They are cheap, easy to use, and can be used to create local backups. It is difficult to share data between machines by using DAS devices as the DAS of one computer is not accessible to other computers. DAS works well for small organizations that only require sharing of data locally and have little or no IT support.

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

Network-attached Storage (NAS)

NAS is a storage device that is connected to a network and allows data storage and retrieval from a central location. It allows only authorized users and clients to store and retrieve data. NAS devices are flexible and can be scaled (additional storage can be added when needed). It is cheaper, faster, and offers all the benefits of an onsite public cloud.

It works like a private cloud and lets users access data remotely via a network connection from anywhere and at any time. This makes it easy for employees to access data continually, collaborate easily, resolve customer issues, etc., as all the information is available in a single place.

NAS mainly handles unstructured data such as audio, video, text files, websites, and documents. Organizations use it for better collaboration and sharing data more effectively between teams that are located remotely or work in different time zones. It is connected to a wireless router so that all devices connected to that network can access files and data easily.

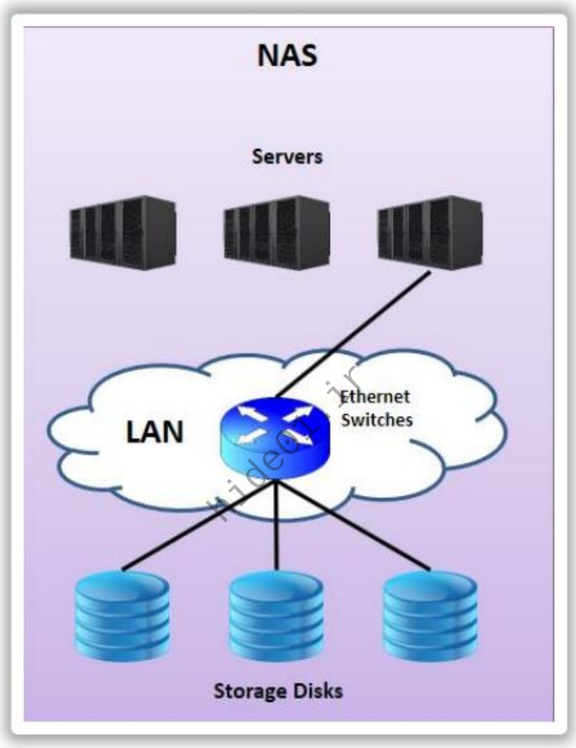


Figure 4.7: Network-attached Storage

NAS is commonly used to

- Share documents and spreadsheets between coworkers
- Store and serve multimedia files at home
- Take a backup for DR and archiving data in an enterprise
- Centralize data storage easily

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

Storage Area Network

SAN is an independent high-speed network that offloads data from servers/desktops to the storage devices. The important protocols utilized in SAN are FCP, iSCSI, FCoE, etc.

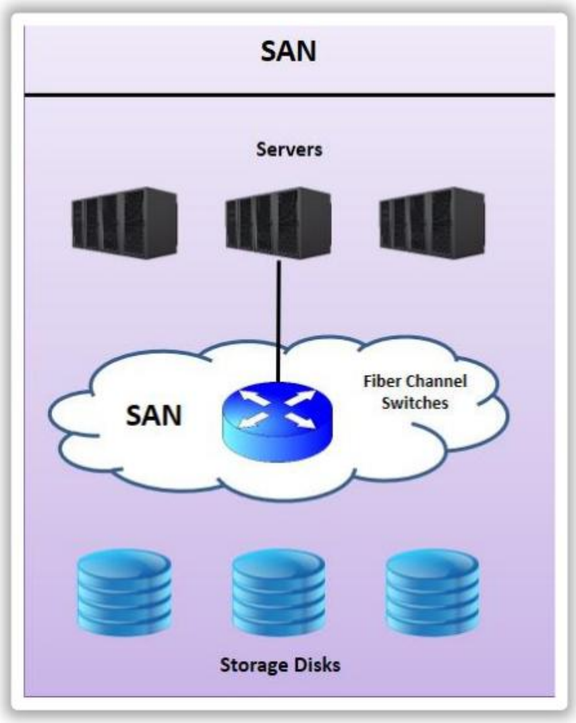


Figure 4.8: Storage Area Network

The following two types of SAN have been used extensively for storing data in the data center:

- **Internet Protocol Storage Area Network (IP SAN):** IP SAN is a storage area network that uses IP to send the storage traffic instead of fiber channel cables. It allows many servers to access block storage devices by using certain storage protocols such as fibre channel (FC) IP and iSCSI. These storage protocols exchange block-based data between a server and storage arrays.

  IP SAN is much simpler to manage when compared to FC SAN as it does not require any special hardware such as FC switches or host bus adapters. Instead, it uses commodity Ethernet networking hardware. The iSCSI protocol is used extensively by IP SAN to encapsulate SCSI commands and collect data into packets to facilitate the transfer between storage devices and host servers. This protocol helps in sharing storage resources over an IP network. The FC IP protocol connects distributed FC SAN islands over an IP network. FC frames are wrapped with the IP payload before they are transferred over the IP network. IP SAN provides easy access to remote storage, competitive speed, and cost-saving by eliminating the need to buy any extra equipment for establishing connections.

- **Internet fibre channel protocol (iFCP):** This protocol helps extend FC SAN across the internet. It uses TCP/IP to transfer data to and from FC storage devices on the internet or in the local SAN. This protocol integrates FC and SCSI networks into the internet. It can also be used along with current FC protocols such as FCoE. iFCP can be used to directly attach FC storage components to an IP network.

  This protocol can fix some problems that FC IP cannot. For example, FCIP can only encapsulate FC data and forward it as an extension of the FC network over a TCP/IP network. FCIP works only inside the FC environment. However, iFCP helps in migrating to

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

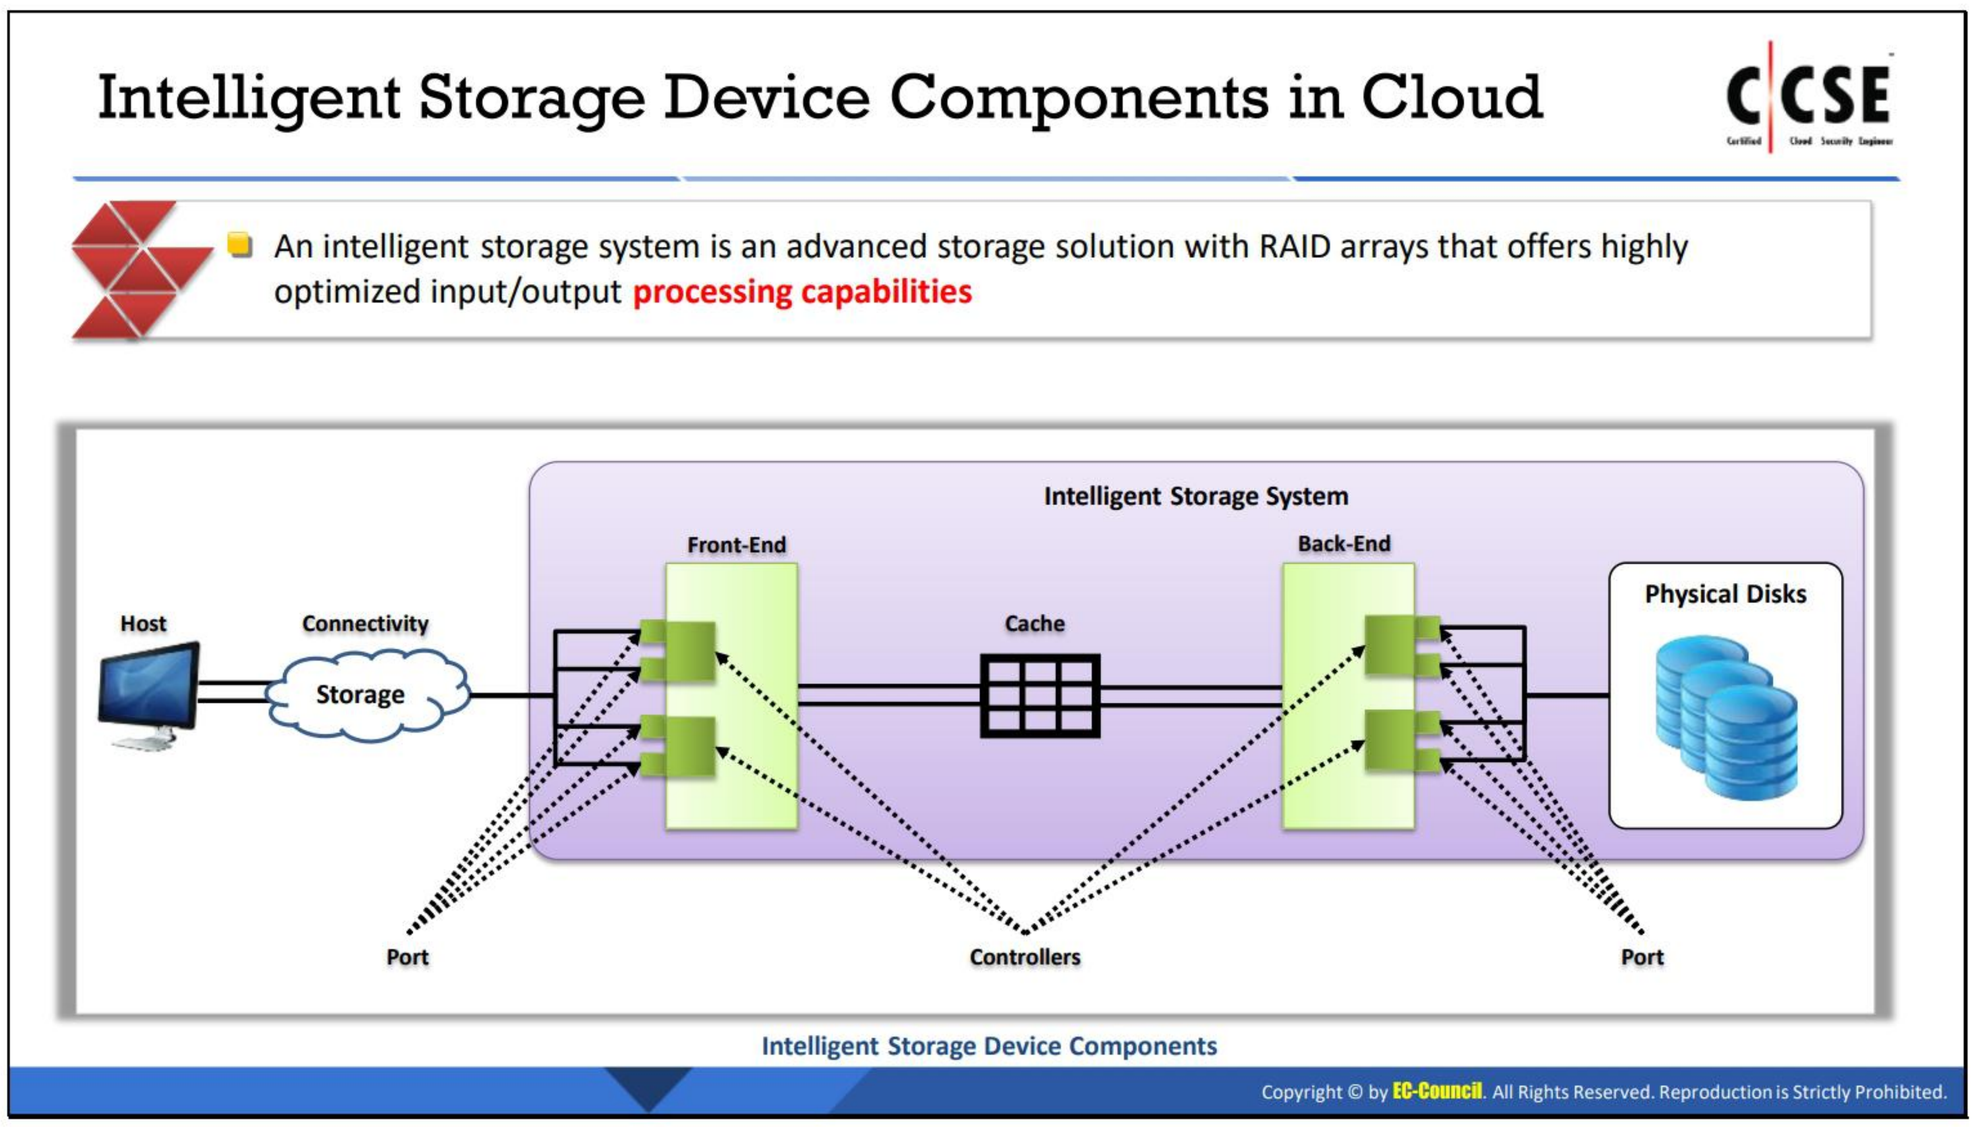
an IP SAN (or a hybrid network) from an FC SAN as iFCP gateways can either complement or replace the current FC fabrics.

- **Internet small computer system interface (iSCSI):** It is a transport layer protocol that transports block-level data over TCP/IP networks between the iSCSI initiator (a piece of hardware or software installed on a server) and the storage target (a server that contains storage resources). It encrypts network packets while sending and decrypts them when they reach the target.

  With the help of iSCSI, it is easy to create a shared-storage network in which several clients and servers are allowed to access central storage resources as if they are connected locally. It sends SCSI commands using initiators to storage targets on remote servers. These targets may be tape, NAS, SAN, SSD, HDD, etc. The administrator can use this protocol to allow hosts to better utilize shared storage by letting them store data on remote-networked storage. With the help of this protocol, the administrator can virtualize remote storage for those applications that need DAS. This protocol does not require costly and complex cards and switches such as FC networks because iSCSI uses standard Ethernet. This makes iSCSI cheaper and easier to implement.

hide01.ir

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



Intelligent Storage Device Components in Cloud

An intelligent storage system is an advanced storage solution with RAID arrays that offers highly optimized input/output processing capabilities. A functional environment controls the allocation, management, and utilization of storage resources. Intelligent storage systems are configured with huge amounts of memory known as cache and they use advanced algorithms to satisfy the input/output requirements of applications that are performance sensitive.

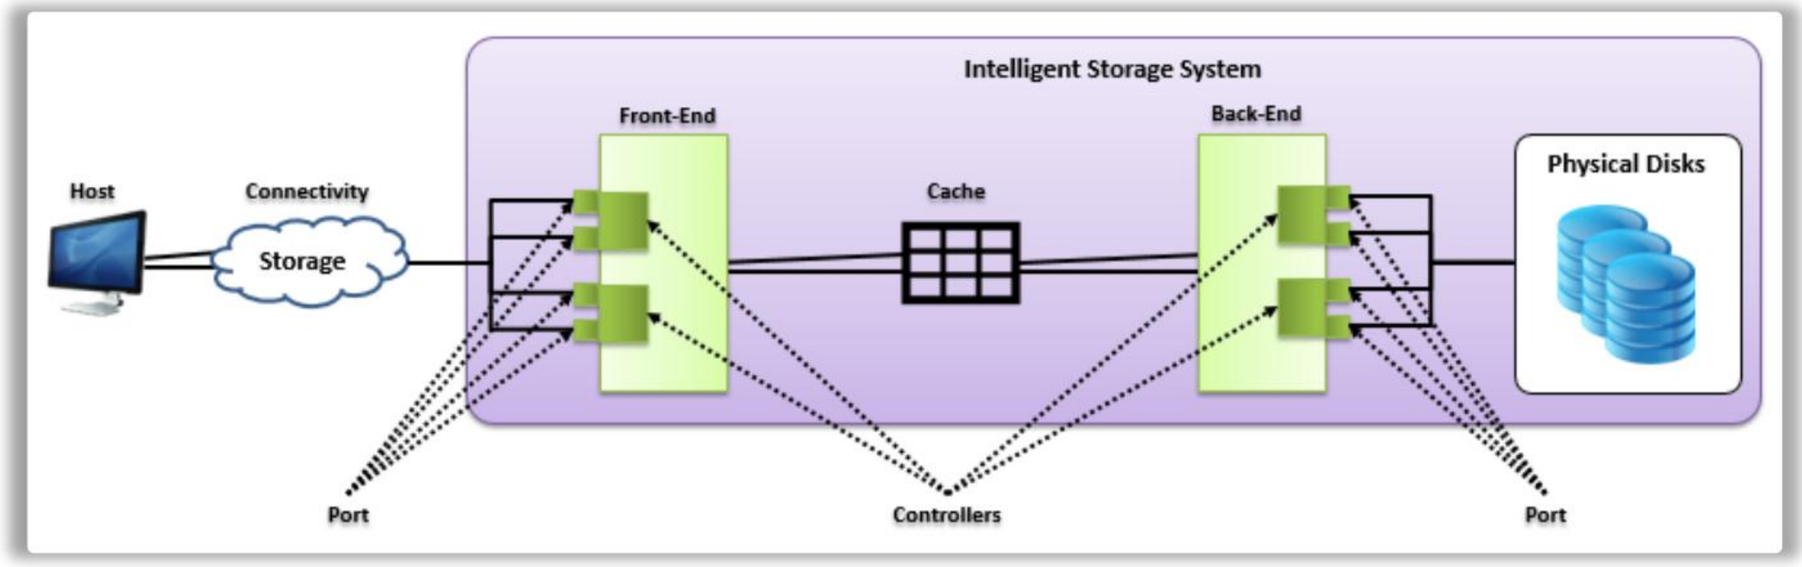


Figure 4.9: Intelligent Storage Device Components

An intelligent storage system has the below components.

- **Storage device front-end**: The front-end serves as an interface between the host and the storage system. There are two components of the front-end.
  - **Front-end ports**: These ports allow hosts and servers to connect to the storage system. For storage connections, appropriate protocols like FC, SCSI, or iSCSI are executed by every front-end port's processing logic.

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

  - **Front-end controllers**: They use the internal data bus for routing data from and to cache. The controller acknowledges to the host when the data is received by cache and uses some command queuing algorithms such as FIFO and seek time optimization to optimize I/O processing.

- **Storage controller**: Storage controllers process I/O requests in the intelligent storage system. They help perform several important functions such as storage management, serving I/O requests, provisioning storage, RAID protection, automated tiering, data encryption, data compression, intelligent cache management, and local and remote data replication. Every controller has some amount of cache memory and one or more processors to attend I/O requests. More than one storage controllers are configured in a cloud intelligent storage system to provide redundancy when a controller fails to work. Load balancing is easily achieved with the help of multiple controllers.
- **Storage processor ports**: Storage processor ports are front-end ports that have some processing logic to choose appropriate transport protocols like SCSI or FC to establish storage connections. Intelligent storage systems for clouds use redundant ports on the front-end to ensure high availability.
- **Dual storage ports**: The intelligent storage systems have dual storage ports configured. If each disk has dual ports, it can connect to two separate controllers using them. This enhances the reliability of the storage system.
- **Storage cache**: Cache plays an important role in enhancing the input/output (I/O) performance in the storage system. It is memory where data is temporarily kept, thereby minimizing the time needed to attend or service I/O requests from the host or server. It reduces the overall delay in I/O request processing. It generally takes some milliseconds for the host to access data from the physical disks directly due to rotational latency and seek times. However, cache provides the same data within a millisecond.
- **Storage back-end**: The back-end serves as an interface between the physical disks and cache. It regulates the transfer of data between the physical disks and cache. It receives data from cache and forwards it to the physical disks.
- **Storage back-end controller**: Storage back-end controller interacts with the physical disks while performing read/write operations. It also offers temporary, limited storage. Back-end controllers implement algorithms that offer error detection and correction as well as RAID functionality.
- **Storage back-end port**: The ports on the storage back-end are connected to the physical disks.
- **Redundant storage back-end ports**: Redundant back-end ports are configured to provide high availability and data protection. Configuring extra back-end ports helps in providing an alternate path to the disks when a port fails to work.
- **Physical disks**: Data is stored in physical disks of the intelligent storage system. A FC or SCSI interface is used to connect the back-end and the physical disks. In the storage

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

system, a combination of FC or SCSI drives and IDE/ATA drives can be used. Physical disks are the slowest components in the whole intelligent storage system.

hide01.ir

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

Cloud Storage Mechanisms

The storage mechanisms in cloud are as follows:

Cloud Storage Mechanisms	
Caching	Caching refers to storing data in a **temporary storage** area known as cache to facilitate faster access to information or data
Data Aging Management	The data aging management mechanism in cloud storage automates the **transfer** of datasets between various cloud storage device types in a policy-driven, automated manner as per each dataset's lifecycle
Data Management	Cloud data management refers to storing an organization's data on a **remote server** which is owned and managed by a third-party vendor who has expertise in cloud data hosting
Device Performance Monitor	The device performance monitor mechanism in cloud storage ensures the specified levels of **performance** are met
Data Placement Auditor	The cloud data placement auditor mechanism helps to control where **datasets** are stored on the cloud
Gateway	Cloud storage gateway refers to a hardware or software appliance that **connects** local applications with cloud storage
Import and Export	The import and export mechanism provide a way of **transporting data** to and from a cloud storage device
LUN Masking	The LUN masking mechanism helps in **configuring** security policies that are needed to make storage LUNs available to systems and storage devices that need access
Raid-level Identifier	The RAID-level identifier mechanism offers **RAID-level** information on cloud storage devices
Storage Path Masking	The storage path masking mechanism helps to **explore** the paths available to reach a physical storage device or cloud storage device

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

Cloud Storage Mechanisms

The storage mechanisms in cloud are explained below.

- **Caching**: Caching refers to storing data in a temporary storage area known as cache to facilitate faster access to information or data. It allows re-accessing or re-using the previously computed or retrieved data efficiently. Caching helps in serving requests faster by storing the previously accessed data in cache; thus, whenever that data is requested again, it can be retrieved faster without server involvement. It reduces the overall delay in I/O request processing.

  Caching provides the following benefits in a cloud environment:

  - Increases throughput
  - Improves application performance
  - Reduces database cost
  - Minimizes the load on the back-end
  - Better accessibility

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

- **Data aging management**: The data aging management mechanism in cloud storage automates the transfer of datasets between various cloud storage device types in a policy-driven and automated manner as per each dataset's lifecycle. Different data management lifecycles can be defined for different datasets. Based on their lifecycle, they may need to be hosted on different cloud storage devices. A data lifecycle can be defined by the cloud consumer through dedicated interfaces. Cloud service consumers can also define policies based on the lifecycle.

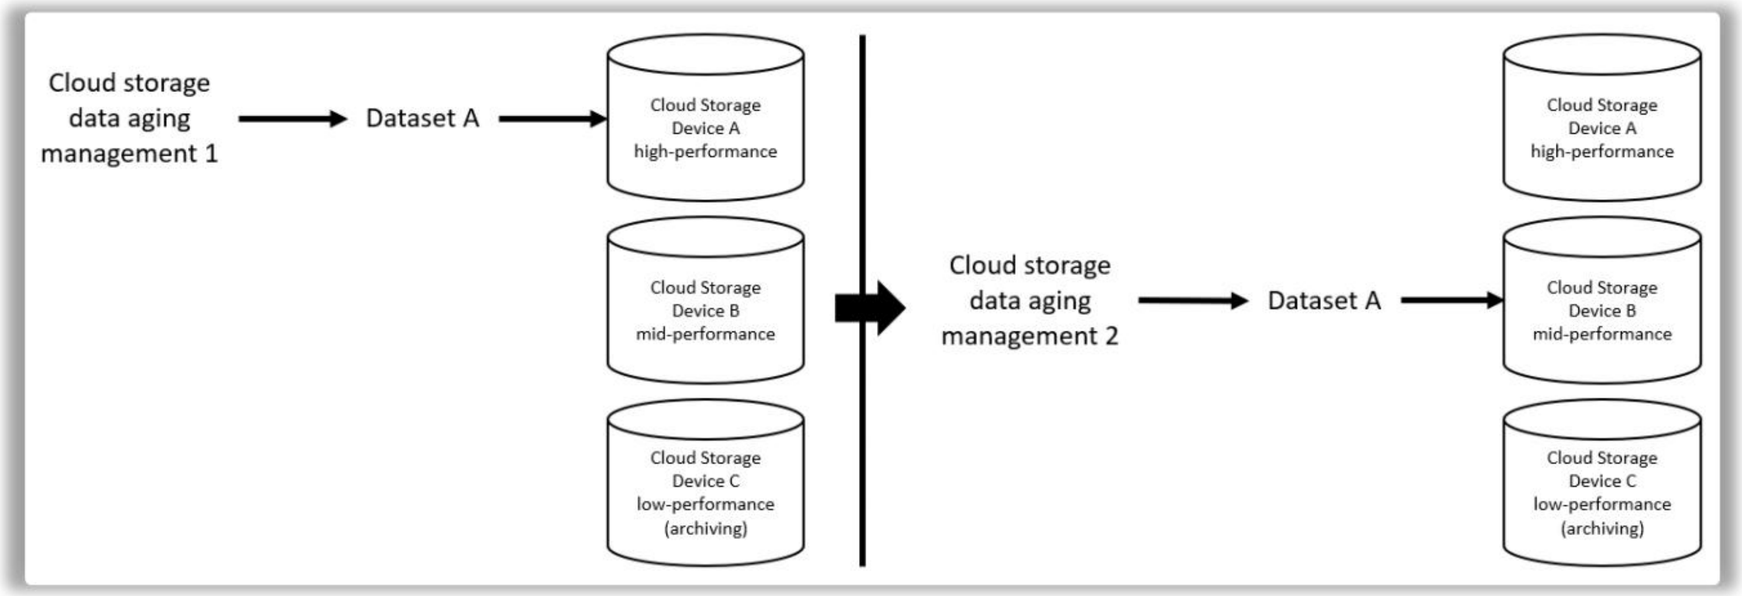


Figure 4.10: Data Aging Management

Different cloud storage devices have different pricing and performance baselines (as shown above). A customer's data may undergo changes in its availability, accessibility, or performance requirements over time as part of their data lifecycle. With the help of other cloud storage mechanisms, the data aging management mechanism can move the dataset (Dataset A) to another storage device without affecting its availability or service.

- **Data management**: Cloud data management refers to storing an organization's data on a remote server, which is owned and managed by a third-party vendor who has expertise in cloud data hosting. It also includes identifying the right or suitable data from the on-premises storage and transferring it to the cloud via data replication, data archiving, or data migration. With cloud data management, organizations can get professional support, an automated backup strategy, and easy access to resources from any location. Cloud data management offers the following benefits:
  - Scalability
  - Governed access
  - Enhanced security
  - Automated backups and disaster recovery
  - Less maintenance cost

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

- **Device performance monitor**: The device performance monitor mechanism in cloud storage ensures that specified levels of performance are met using a policy-based model of storage allocation. It also performs other important functions such as checking the existing location of datasets automatically against the specified required performance metrics. It checks this to ensure that datasets are always in a cloud storage device that satisfies its requirements.

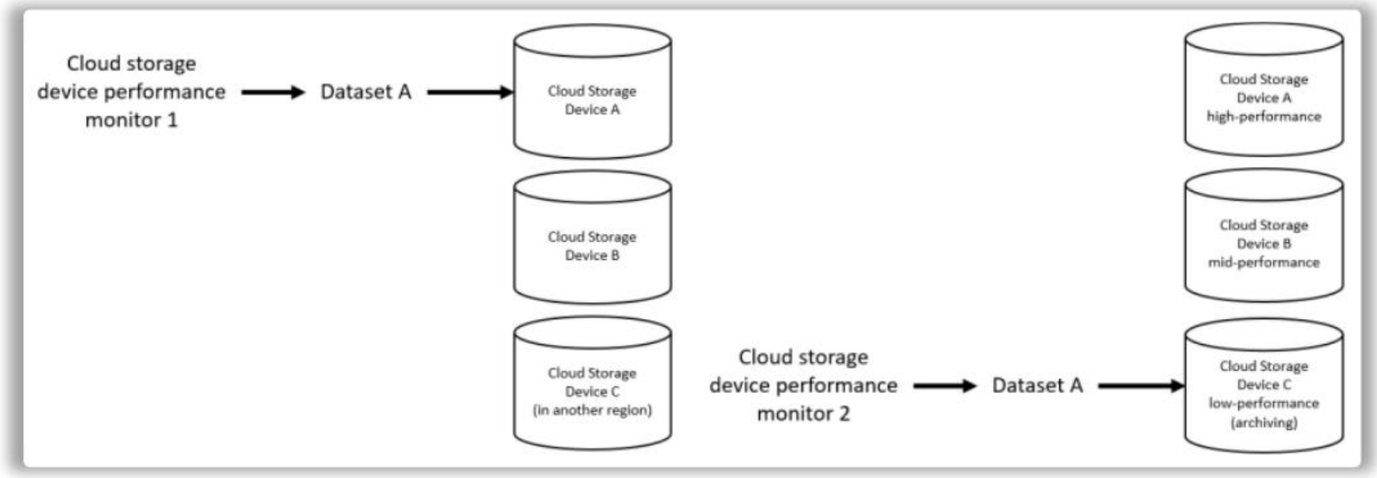


Figure 4.11: Device Performance Monitor

If a dataset resides in a cloud storage device that does not satisfy the consumer's requirements, the device performance monitor raises an alert. It raises an alert even when a particular dataset is moved to a cloud storage device that does not meet the requirements of that dataset. It helps prevent a dataset from being stored on a storage device that cannot offer the level of performance required for that dataset.

- **Data placement auditor**: The cloud data placement auditor mechanism helps control where datasets on the cloud are stored. A cloud administrator can use this mechanism to enforce policies about locations where each dataset can be stored or cannot be stored on the cloud. The administrator can also use data placement auditor for checking and auditing storage locations of datasets to ensure that the datasets are stored in appropriate cloud storage devices based on the requirements specified in the contract. When a dataset is stored in a storage device where it should not be hosted, the auditor mechanism checks regulatory and compliance policies against the location where the data is hosted. The administrator can use SDKs, APIs, a policy engine, and service-level agreement (SLA) monitor mechanism to establish the cloud data placement auditor mechanism.

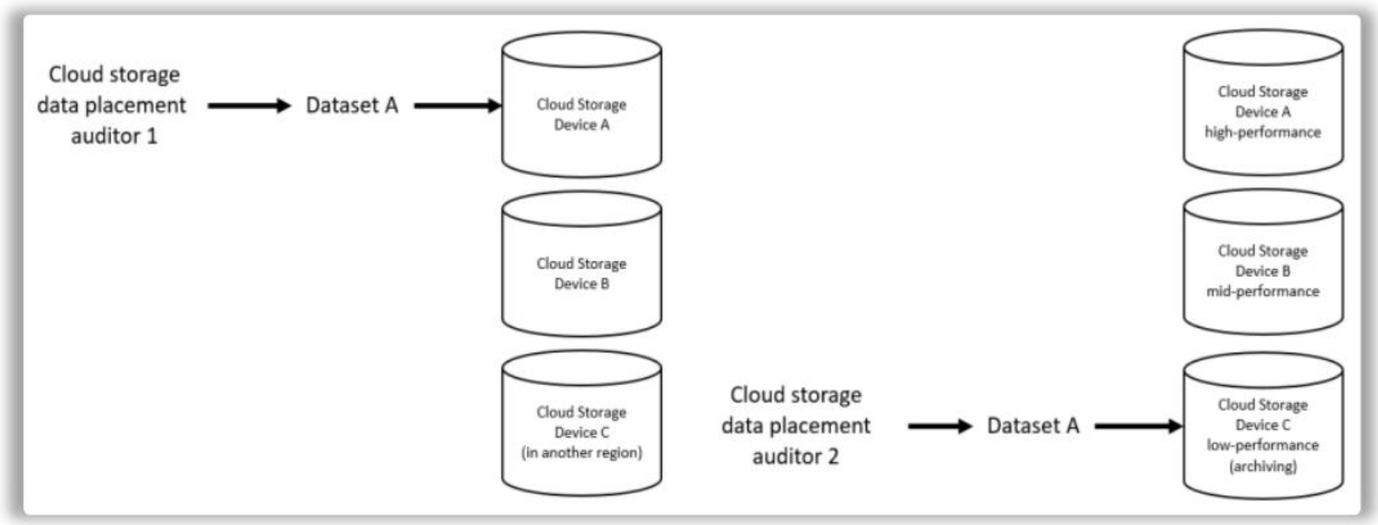


Figure 4.12: Data Placement Auditor

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

- **Gateway**: Cloud storage gateway refers to a hardware or software appliance that connects local applications with cloud storage. It can be a hardware or virtual machine (VM) image that offers basic protocol translation and easy connectivity to help technologies that are not compatible to communicate with each other. Cloud storage gateways can also be used to archive data in the cloud.
- **Import and export**: The import and export mechanism provides a way of transporting data to and from a cloud storage device. By using temporary or permanent links, data can be imported or exported from a cloud environment with this mechanism. This mechanism also helps export data for storage in another cloud storage device for disaster recovery purposes.
- **Logical unit number (LUN) masking**: The LUN masking mechanism helps in configuring security policies that are needed to make storage LUNs available to storage devices and systems that need access. These devices can access the storage LUNs by using configuration options and interfaces offered by physical storage vendors.

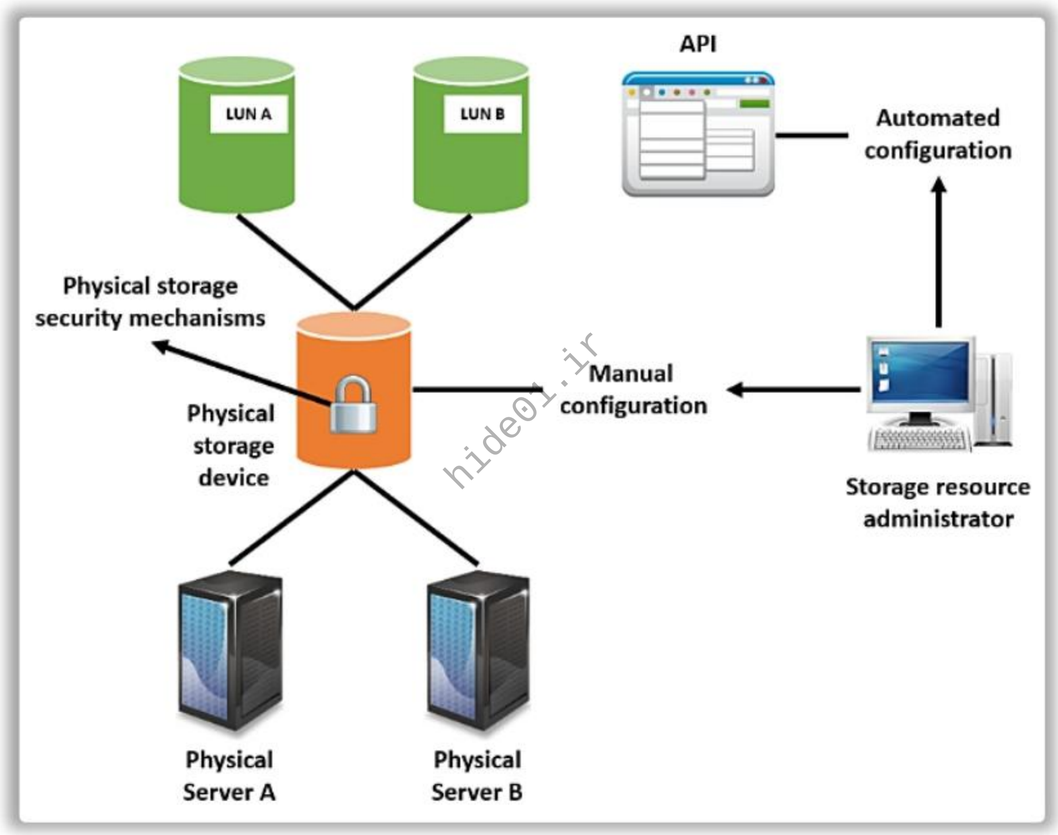


Figure 4.13: LUN Masking

The cloud administrator can use the managed security features provided by a physical storage device to establish the LUN masking mechanism. The administrator can use a management interface or the provided API to define security policies for the servers (shown in the above image) to communicate with the presented LUNs.

- **RAID-level identifier**: The RAID-level identifier mechanism offers RAID-level information on cloud storage devices. The cloud storage administrator can establish this mechanism automatically if cloud storage device vendors offer SDKs or APIs; otherwise, the administrator needs to use the available features such as inventory tags or defined attributes to populate the information manually.
- **Storage path masking**: The storage path masking mechanism helps to explore the paths available to reach a physical storage device or cloud storage device. The cloud

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

administrator can use this mechanism to allow concurrent communication over many pathways and to mask some or all paths to a physical or cloud storage device.

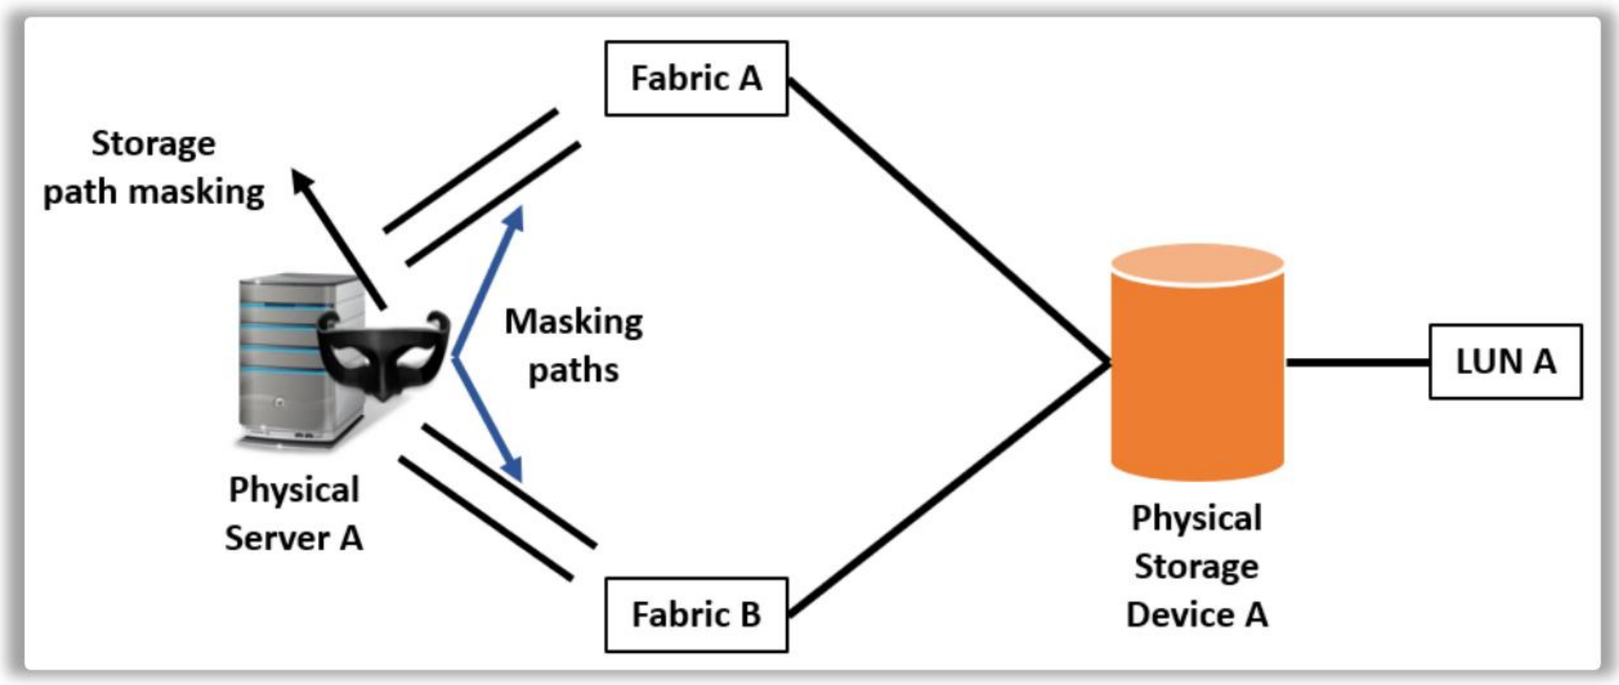


Figure 4.14: Storage Path Masking

As shown in the above image, this mechanism can be utilized to mask two paths for the server to reach the storage device (via Fabric A or Fabric B) while accessing data on LUN A.

hide01.ir

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

Cloud Storage Design Patterns

- Cloud storage design patterns address common **software design problems** by providing a template for solutions to the problems
- Cloud storage design patterns are successfully implemented in the various **cloud storage features**

Performance and Scaling

- The common designs problems related to performance and scaling with field-tested solutions are as follows:
  - Storage workload management
  - Cloud storage device performance enforcement
  - Direct I/O access
  - Direct LUN access
  - Cross-storage device vertical tiering
  - Intra-storage device vertical tiering

Data Migration and Management

- The common designs problems related to data migration and management with field-tested solutions include:
  - Cloud data placement compliance check
  - Single Root I/O virtualization
  - Sub-LUN tiering
  - Cloud data lifecycle management
  - Cloud storage maintenance window
  - Cloud storage data management
  - Cloud dynamic data normalization

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

Cloud Storage Design Patterns (Cont'd)

hide01.ir

Data Integrity and Security

- The common designs problems related to data integrity and security with field-tested solutions are as follows:
  - Cloud storage device masking
  - Cloud storage device path masking
  - IP storage isolation
  - Cloud data-at-rest encryption
  - In-transit cloud data encryption

Reliability and Fail-over

- The common designs problems related to reliability and fail-over with field-tested solutions are as follows:
  - Redundant storage
  - RAID-based data placement
  - Multi-path resource access

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

Cloud Storage Design Patterns

Cloud storage design patterns address common software design problems by providing a template for solutions to the problems. They can be used to build scalable, reliable, secure applications in the cloud by addressing problems that arise while building and designing cloud applications. Common problems related to software designs can be fixed by using design patterns as they are templates (not finalized designs) that can be used to solve problems in various situations. The various cloud storage design patterns are described below.

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

Data Migration and Management

- **Cloud Data Placement Compliance Check**

  Problem: Organizations that use the cloud need to follow regulatory data hosting policies that impose certain location requirements on the storage devices. However, customer organizations may not know the location of the storage devices when they host data in a third-party cloud. How can customers ensure that the storage device which stores their data meets their compliance requirements?

  Solution and application: To monitor the exact location of a cloud storage device, a solution is implemented that can monitor and send notifications if the storage conditions do not satisfy the compliance policies. The cloud storage data placement auditor mechanism helps in enforcing customer or provider-defined policies on a particular storage device or dataset. The below image shows the steps for implementing the data placement compliance check.

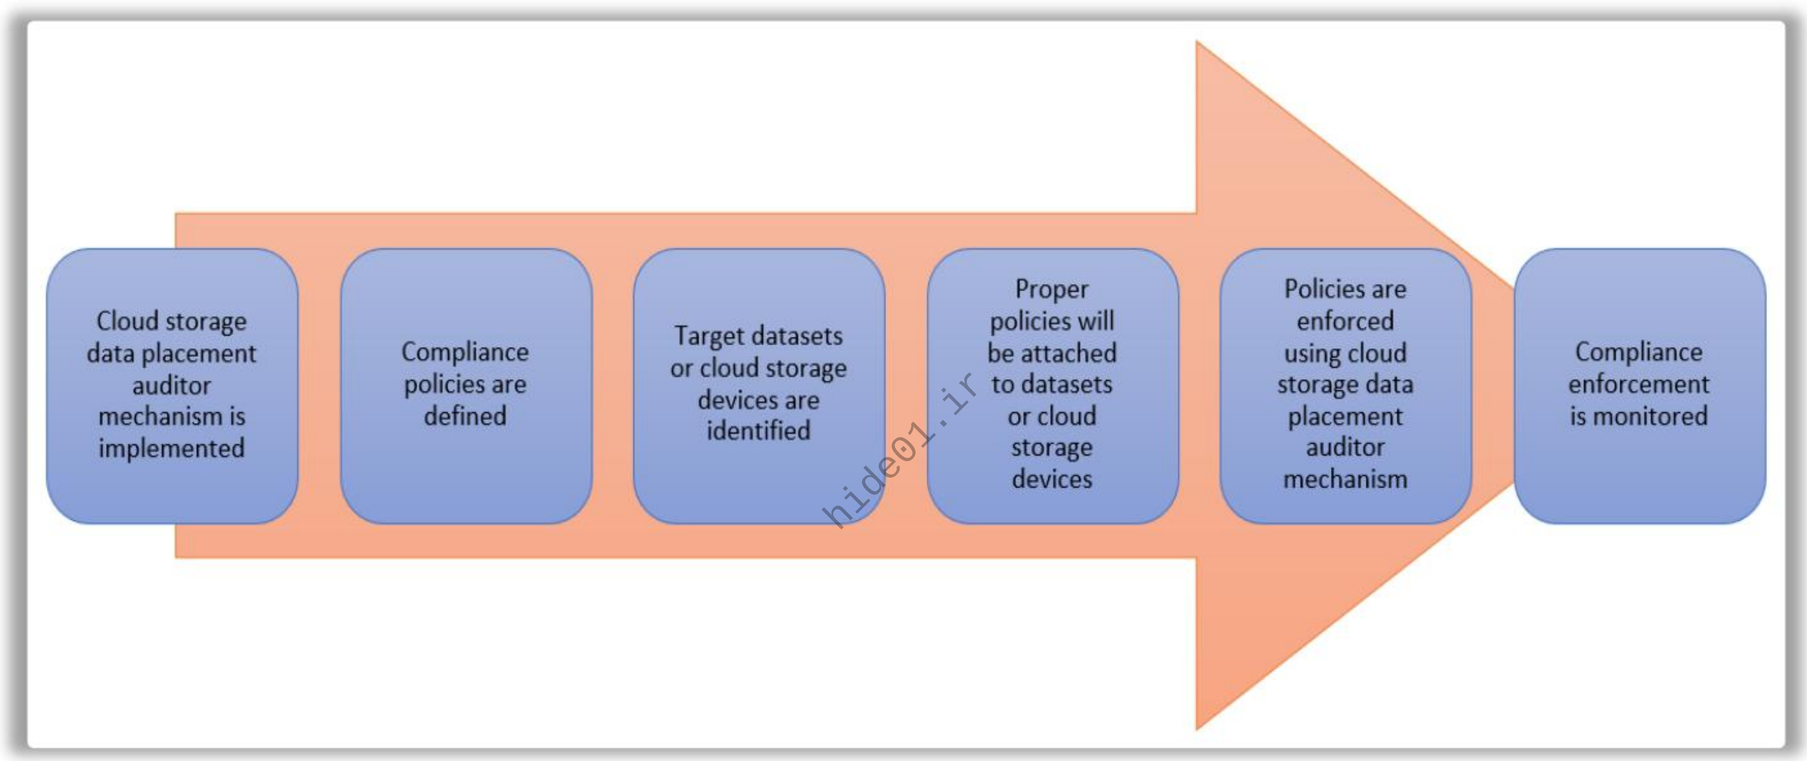


Figure 4.15: Implementing Data Placement Compliance Check

- **Single Root I/O Virtualization**

  Problem: Multiple cloud consumers cannot use a physical single I/O device due to its limitation of being serviceable to one consumer at a time. How can this physical I/O device be virtualized and shared among several consumers to be used by them?

  Solution and application: The physical I/O device's functionality is abstracted and multiple cloud users are presented with its features. To virtualize the device's functionality, PNIC software is utilized, and virtualization techniques are used to present the abstracted functions to consumers. The single root I/O virtualization pattern is implemented using a PNIC software that is attached to a physical server as shown below.

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

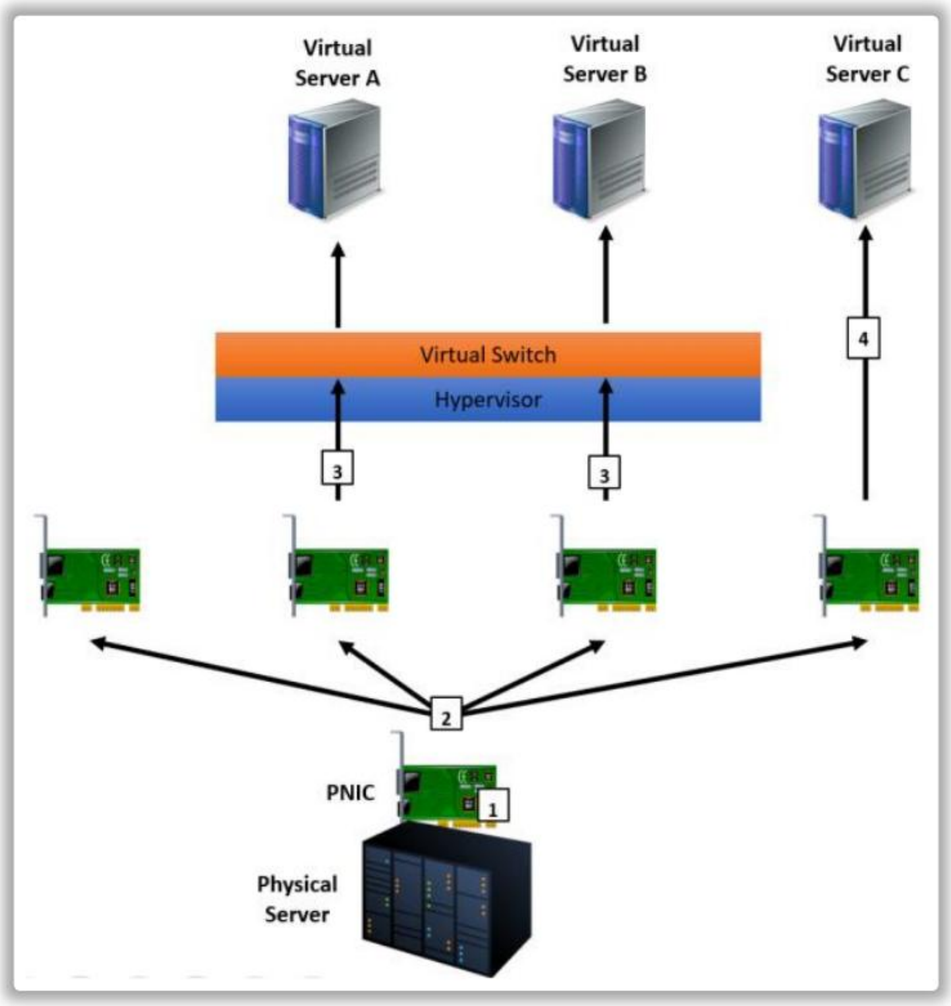


Figure 4.16: Implementing Single Root I/O Virtualization Pattern

- **Sub-LUN Tiering**

  Problem: A subset of data hosted within a LUN has increased performance needs that are not fulfilled in its existing location. How can the data subset stored in that LUN be migrated without moving the entire LUN?

  Solution and application: This problem can be solved by using a solution that has the ability to monitor actual performance levels of the data subset against the required performance levels and by coordinating the dataset migration when required. The automatic scaling listener mechanism can be used to monitor actual performance levels against the required levels and can use a sub-LUN migration mechanism to initiate required migration.

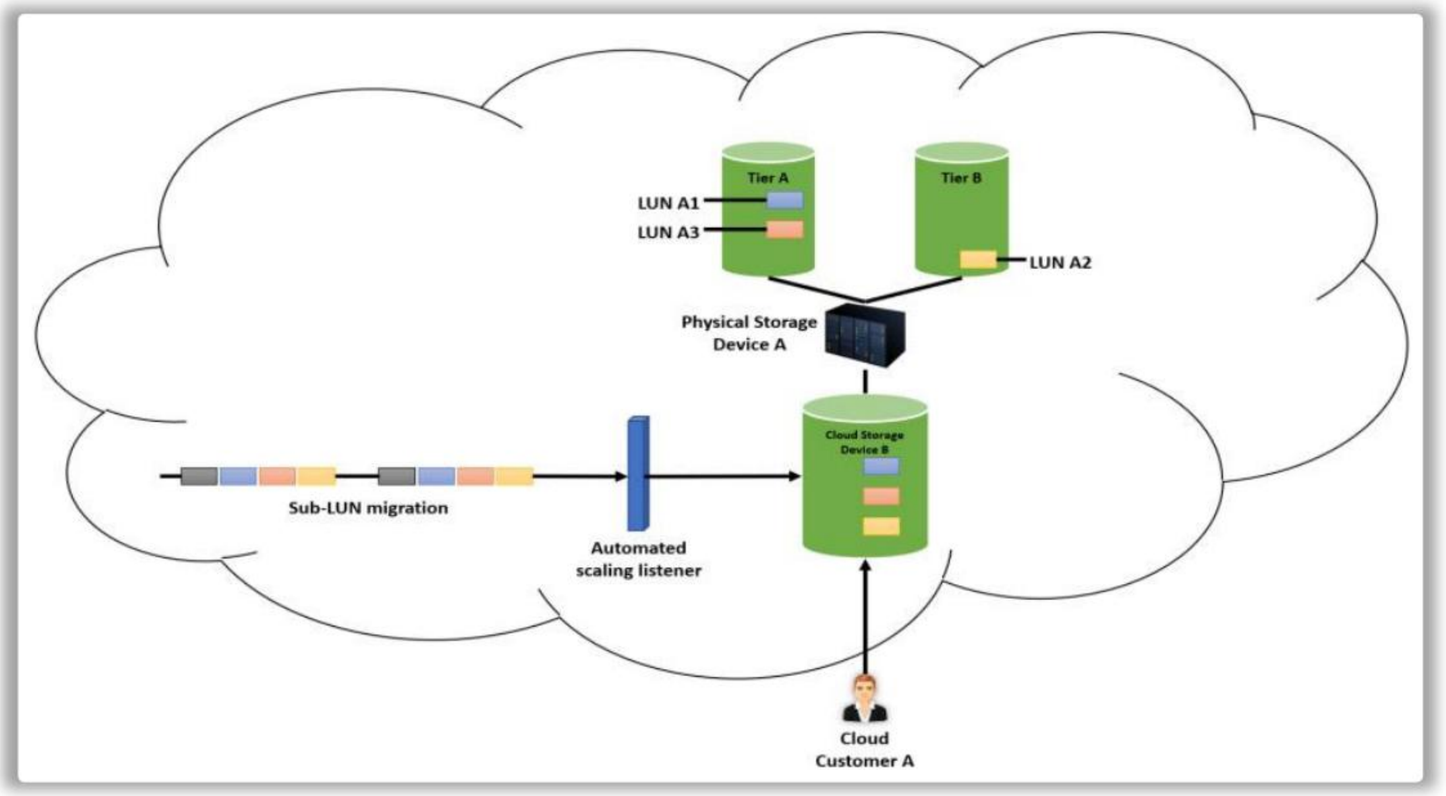


Figure 4.17: Sub-LUN Tiering

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

The Dataset A2 is moved to the higher performance Tier B by the sub-LUN migration mechanism.

- **Cloud Data Lifecycle Management**

  Problem: Organizations have datasets that are no longer required but they occupy storage space in the databases causing performance issues and increasing maintenance and administration burdens. How can be data managed and stored in the cloud as per a defined lifecycle?

  Solution and application: This problem can be solved by using a solution that can manage and move data automatically to a different kind of CS device or delete it based on its state in a specified data lifecycle. To move data to some other storage device or delete it as per the defined lifecycle, the cloud storage data aging management mechanism is used. This mechanism monitors or tracks the state of data against the specified lifecycle to decide whether to move the data or delete it. The steps of applying this pattern are shown below.

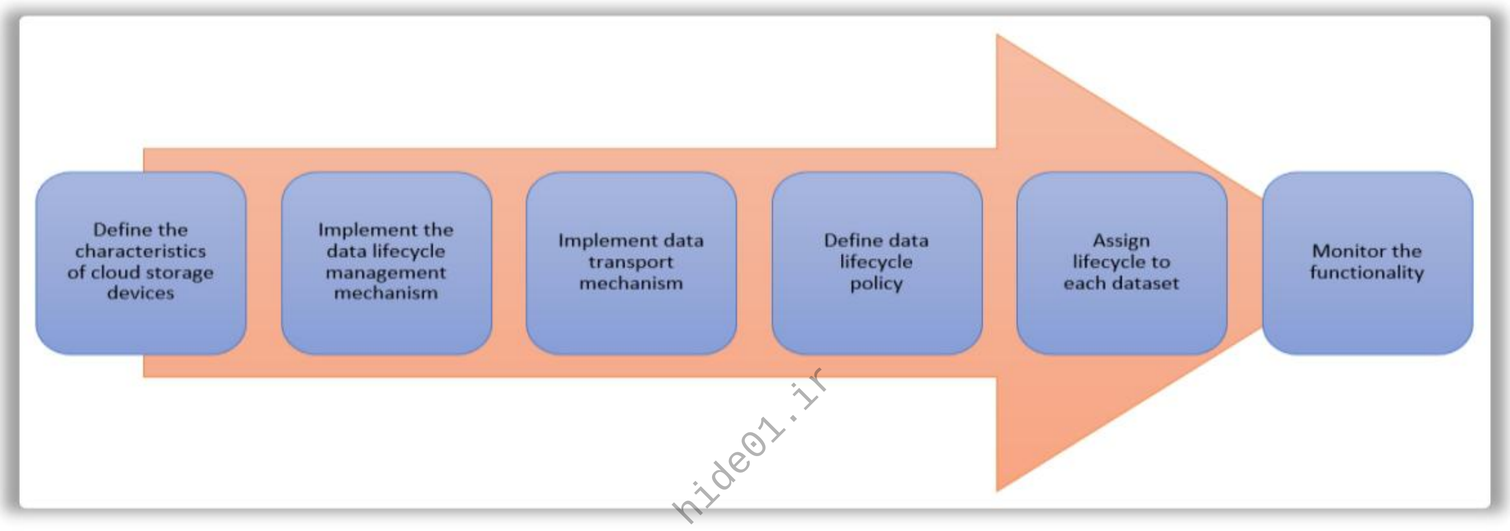


Figure 4.18: Applying Data Lifecycle Management

- **Cloud Storage Data Management**

  Problem: In a cloud environment, not all cloud consumers may have permission to cloud-based storage devices containing their data. This restricts their ability to access as well as manage their important datasets. How can cloud customers access data stored on CS devices directly or provide other cloud customers with appropriate access permission levels?

  Solution and application: This problem can be solved by using a solution that enables cloud customers to access data that is stored in storage devices. The cloud storage management portal mechanism can help them by providing APIs or an interface to interact with the data stored in storage devices. With the help of this portal mechanism, the cloud consumers can manage the data stored on cloud storage devices.

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

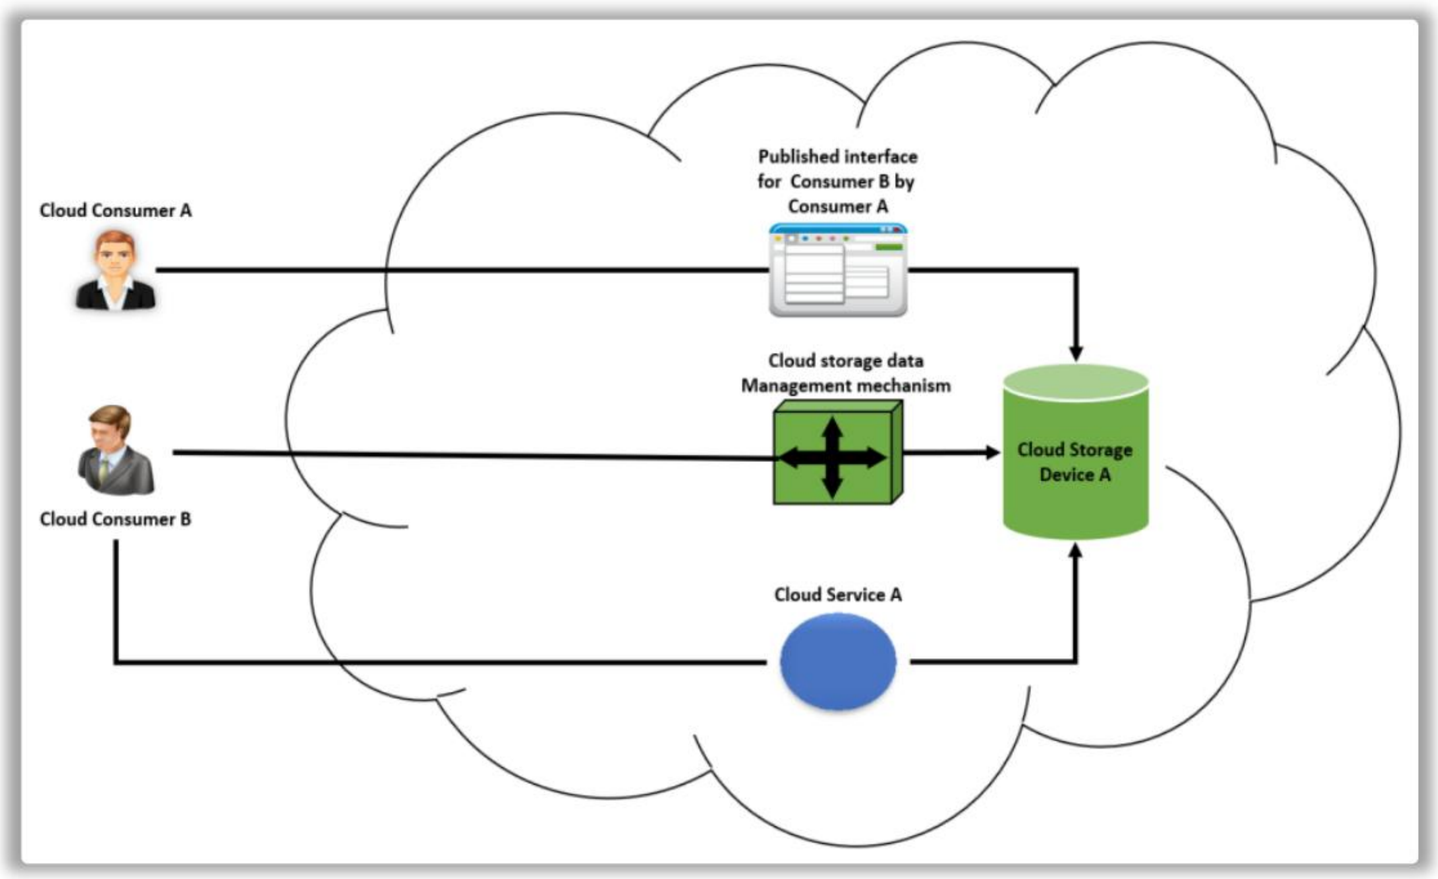


Figure 4.19: Cloud Storage Data Management

- **Storage Maintenance Window**

  Problem: In a cloud environment, shutting down storage devices is required during hardware maintenance, which causes loss of data access and service disruption. What can be done to preserve access to data in a CS device during a maintenance window or outage?

  Solution and application: During maintenance and other outages, an outage prevention system should be used to move data without interruption. LUN migration can be applied to move data temporarily to a different CS device during maintenance.

  When maintenance is scheduled, live storage migration transfers the LUNs to a secondary storage device from the primary storage device. Upon completion of data migration, data requests are directed to the duplicate LUNs in the secondary storage device.

hi-e01.ir

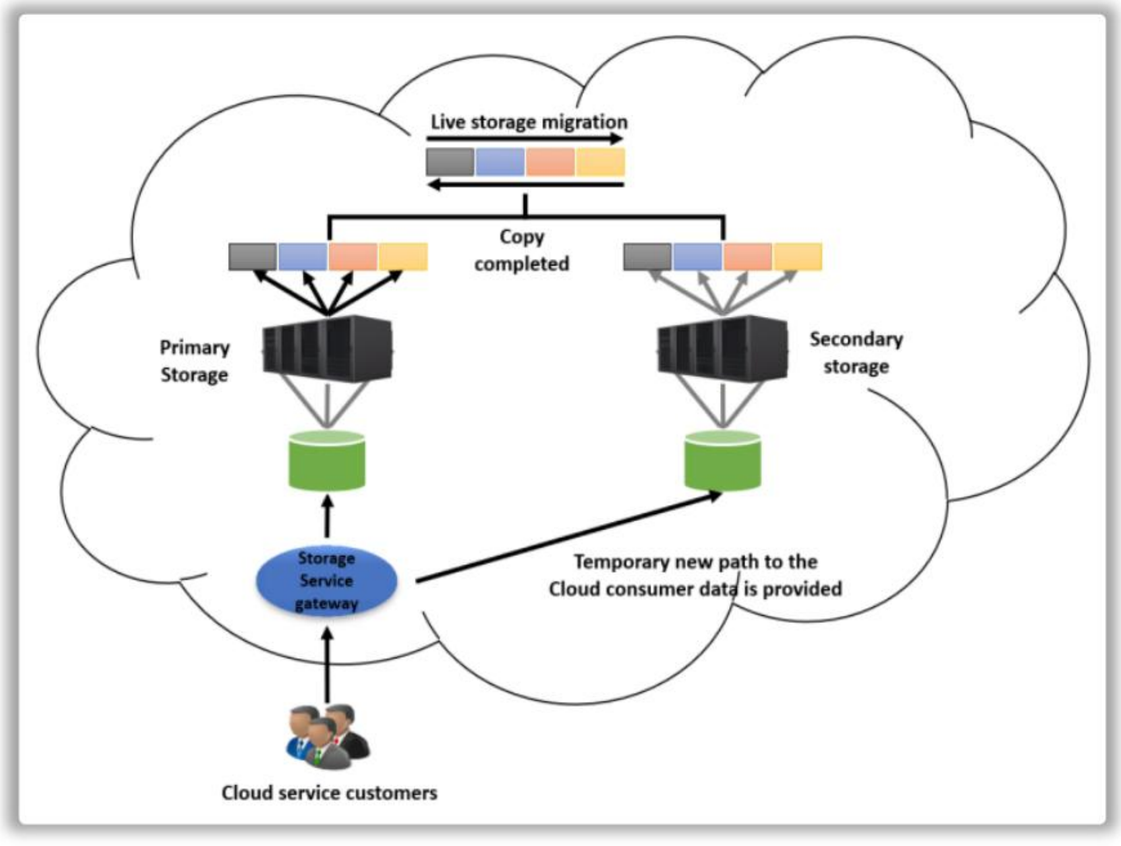


Figure 4.20: Transfer of LUN to Secondary Storage Device

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

Then, the primary storage is shut down for maintenance and the secondary storage device services the data requests until the primary storage device is live again.

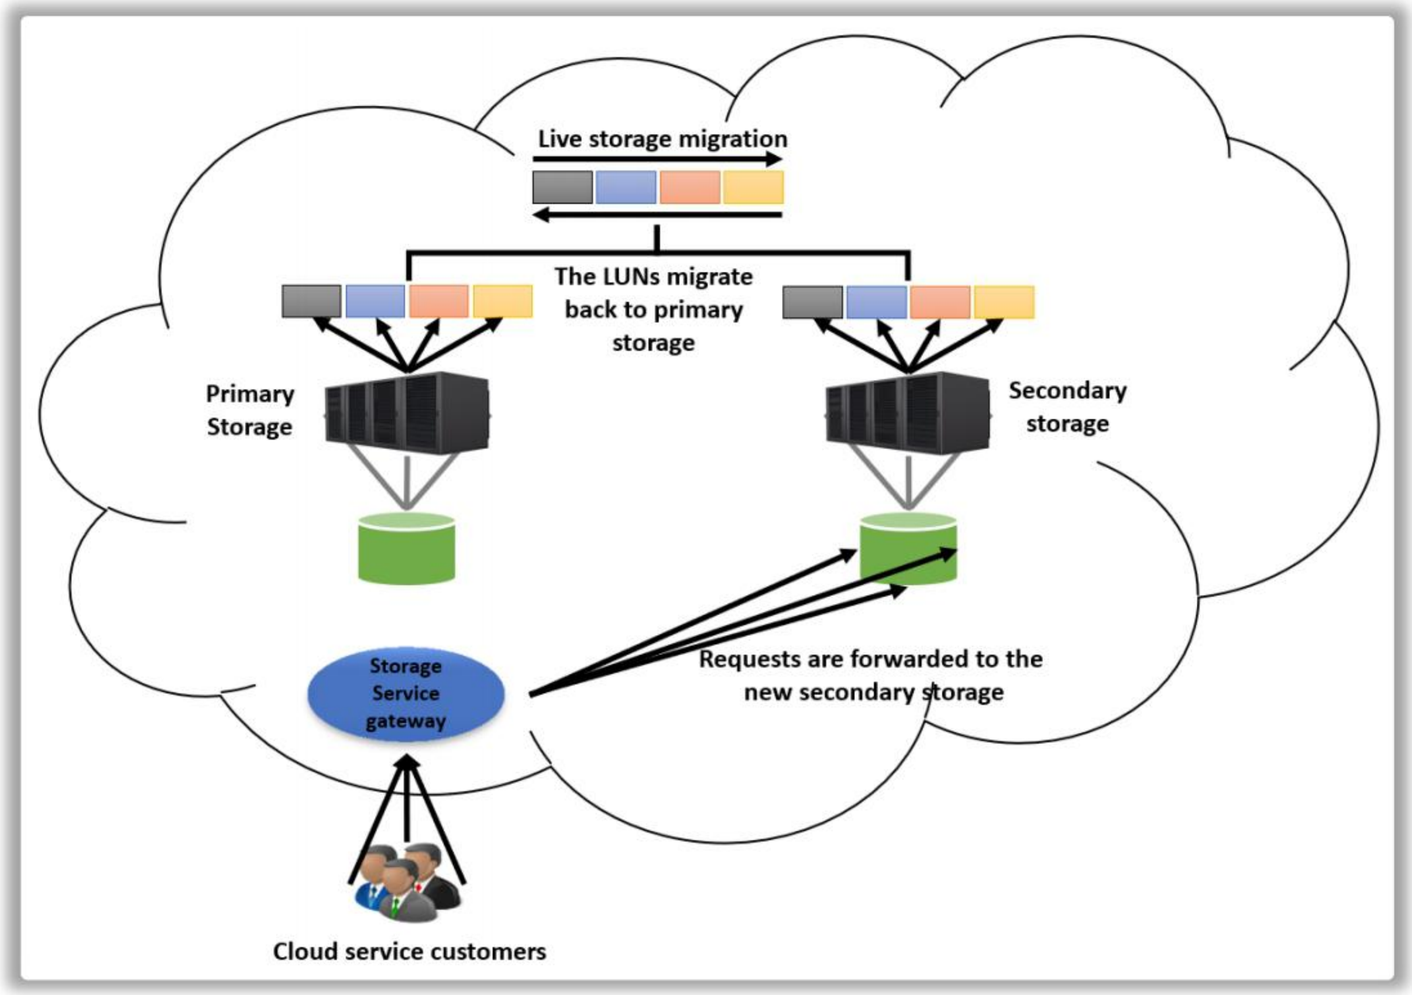


Figure 4.21: Secondary Storage Device Handling the Data Requests

Once the maintenance is completed at the primary storage device, requests for data access are redirected back to it.

hide01.ir

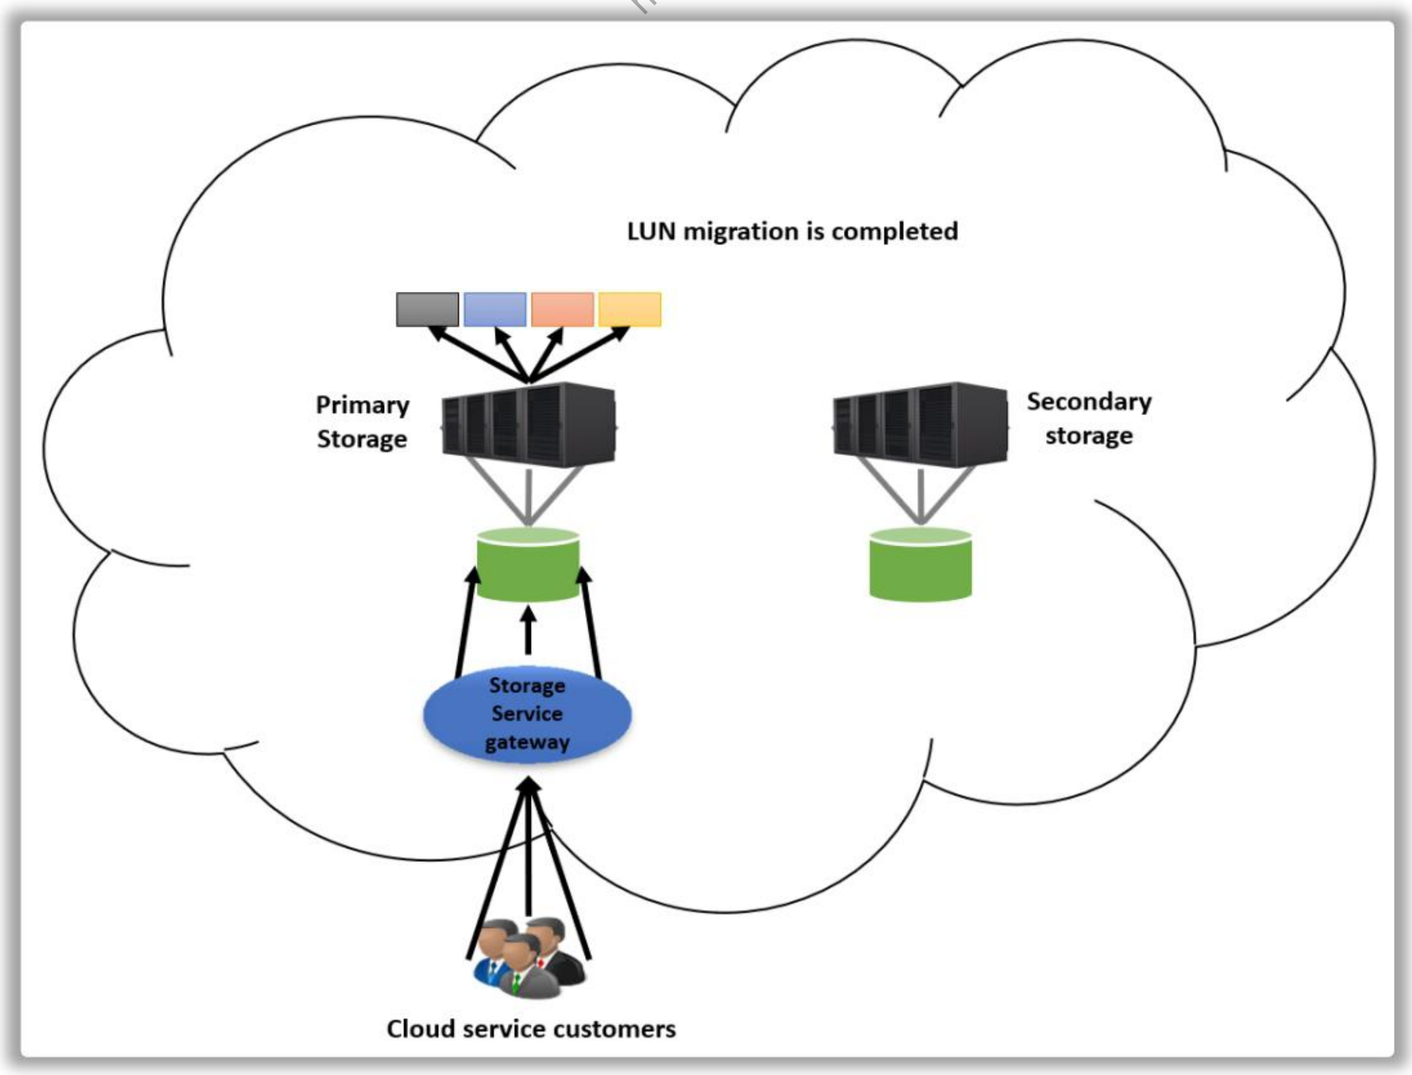


Figure 4.22: Redirecting the Data Requests to Primary Storage Device

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

- **Dynamic Data Normalization**

  Problem: CS devices have huge volumes of redundant data stored by cloud consumers; this bloats the storage architecture and causes data access performance issues. How can this extra or redundant data be automatically eliminated inside CS devices?

  Solution and application: To optimize the cloud storage, device performance, and capacity, consumers' data should be normalized. In order to detect and remove redundant data at file or block-based levels, data de-duplication technology can be used.

Performance and Scaling

- **Storage Workload Management**

  Problem: Over utilization of one cloud storage device occurs when storage-related processing is restricted to that device only. At the same time, other storage devices remain under-utilized or unutilized. This results in a non-optimized CS architecture. What needs to be done to dynamically distribute storage processing workloads across several storage devices?

  Solution and application: A storage capacity system is used to divide and manage LUNs, and to distribute workloads among different storage devices in the network. All storage devices in the cloud are combined in a resource pool. From this pool, they are scaled horizontally along with using a storage capacity monitor as well as LUN migration.



Figure 4.23: Distribution of LUNs Across Cloud Storage Devices

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

The above image shows that LUNs are distributed evenly across cloud storage devices.

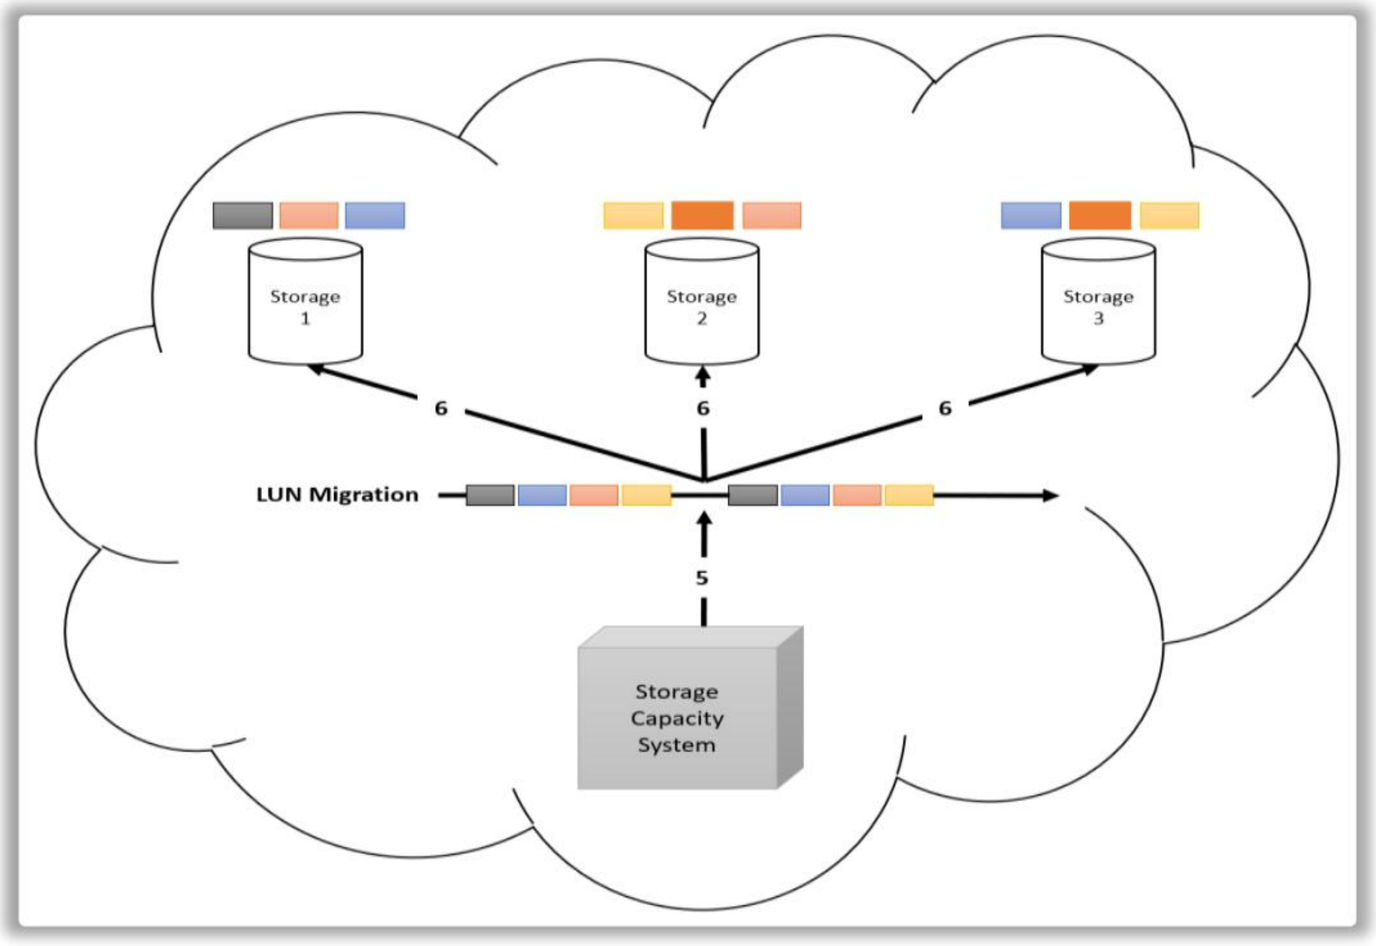


Figure 4.24: Applying Storage Workload Management

By applying the storage workload management pattern, workloads are managed and processed evenly across multiple storage devices as shown above.

- **Cloud Storage Device Performance Enforcement**

  Problem: There can be different performance requirements specified by a cloud consumer for different datasets. If a cloud storage device having fixed performance capabilities host datasets that have varying performance requirements, it cannot meet their performance requirements. How can the data with varying requirements be stored in a storage device that complies with the performance requirements of every dataset?

  Solution and application: This problem can be solved by using a solution that can compare and match the performance requirements or characteristics of datasets against a destination CS device's capabilities. The CS device performance monitor mechanism is useful for managing data stored on a storage device according to its performance characteristics. It can also help enforce policies to avoid data from being moved or copied by sending an alert response.

- **Direct I/O Access**

  Problem: Virtualized networks and their associated virtualized resources carry certain capacity limitations that affect data transfer performance and virtual server communication. How can data transfer capacity thresholds be eliminated, which are imposed on a virtual server by its virtualized environment?

  Solution and application: The problem can be solved by allowing the virtual server to bypass the hypervisor and access the I/O card of the physical server directly. With this solution, total control of the physical server's I/O card is transferred to the virtual server (VS). Then, the VS can accept this I/O card as a recognized hardware device.

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

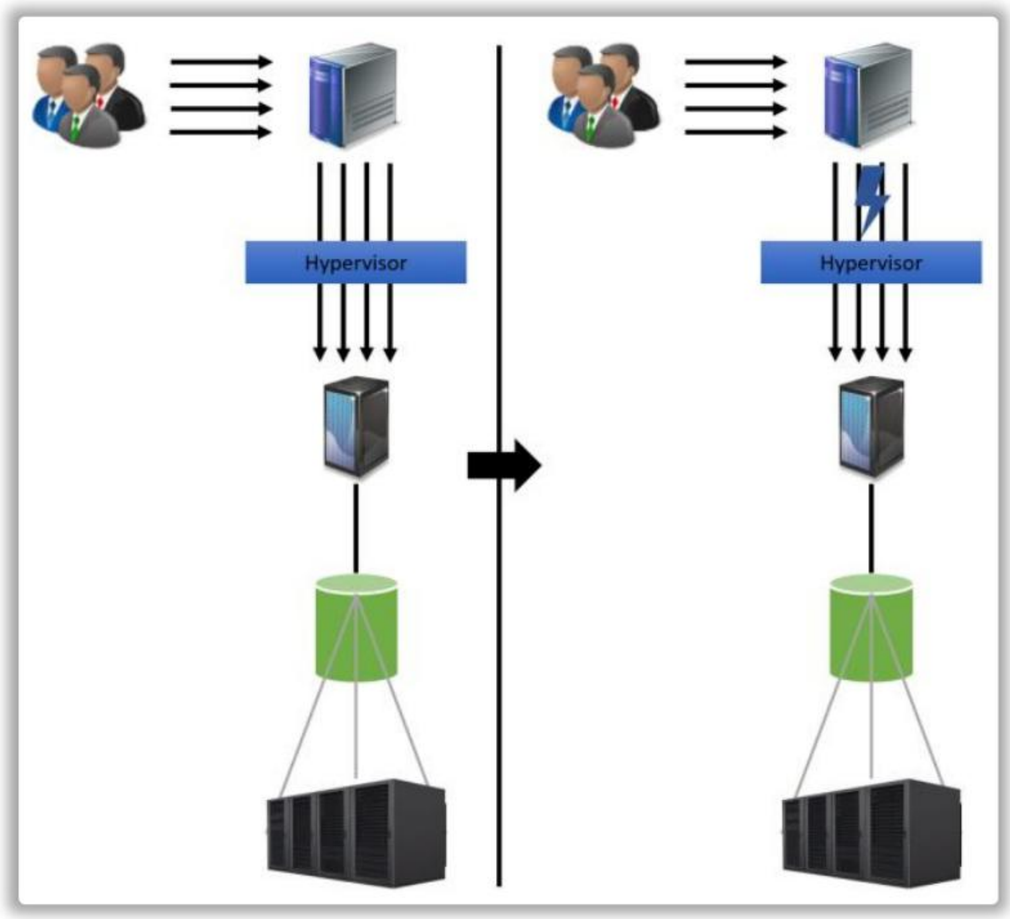


Figure 4.25: Direct I/O Access

A VS accesses the database stored on a SAN storage LUN via a virtual switch. When the amount of requests increases, it causes bandwidth and performance issues (as virtual NIC's bandwidth and performance are not sufficient).

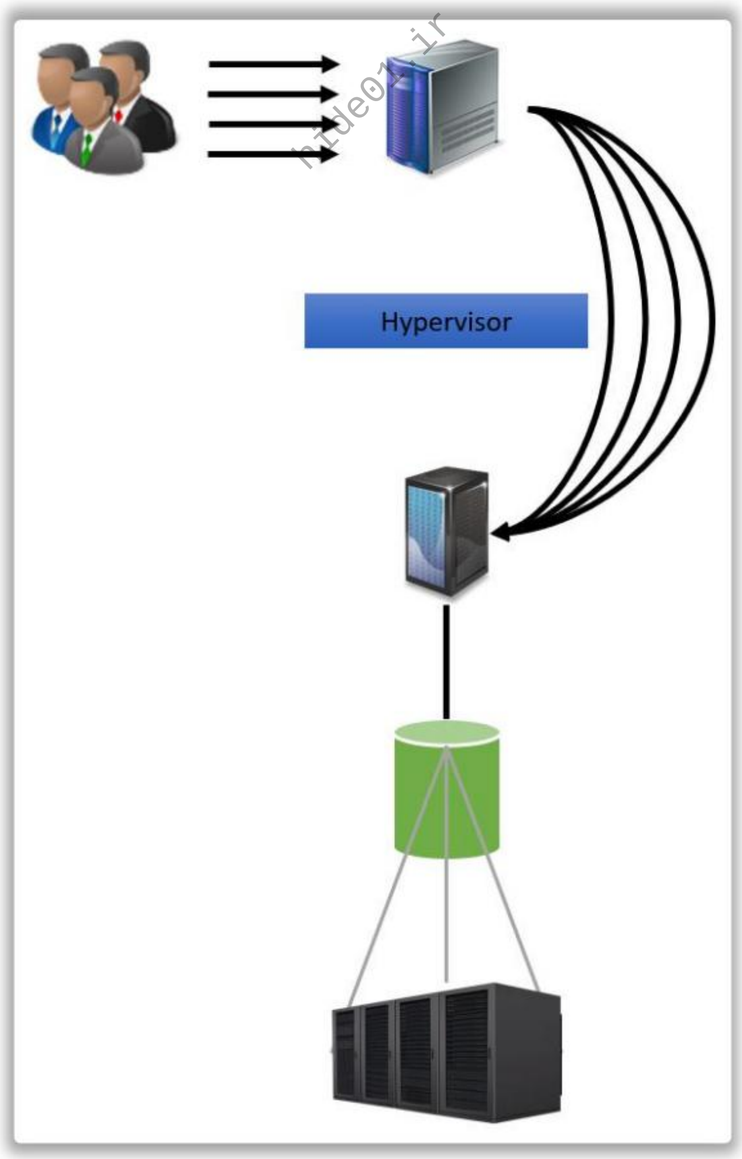


Figure 4.26: Virtual Server Bypassing the Hypervisor

The increased amount of workload can be handled effectively by bypassing the hypervisor via a direct physical link to the physical server in order to access the database.

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

- **Direct LUN Access**

  Problem: LUNs that are mapped on the hypervisor by using a host bus adapter (HBA) may restrict or prevent data access to emulated file-based storage (EFS). This can result in performance limitations. How can a VS control the performance limitations that are imposed by EFS?

  Solution and application: To overcome the performance limitations of the VS, a physical HBA must be used to provide the VS with direct access to block-based storage. Raw-device mapping technology can be leveraged to configure the hypervisor so that it allows the VS to access raw, block-based LUNs. The below image shows how the virtual servers are given direct access to block-based LUNs.

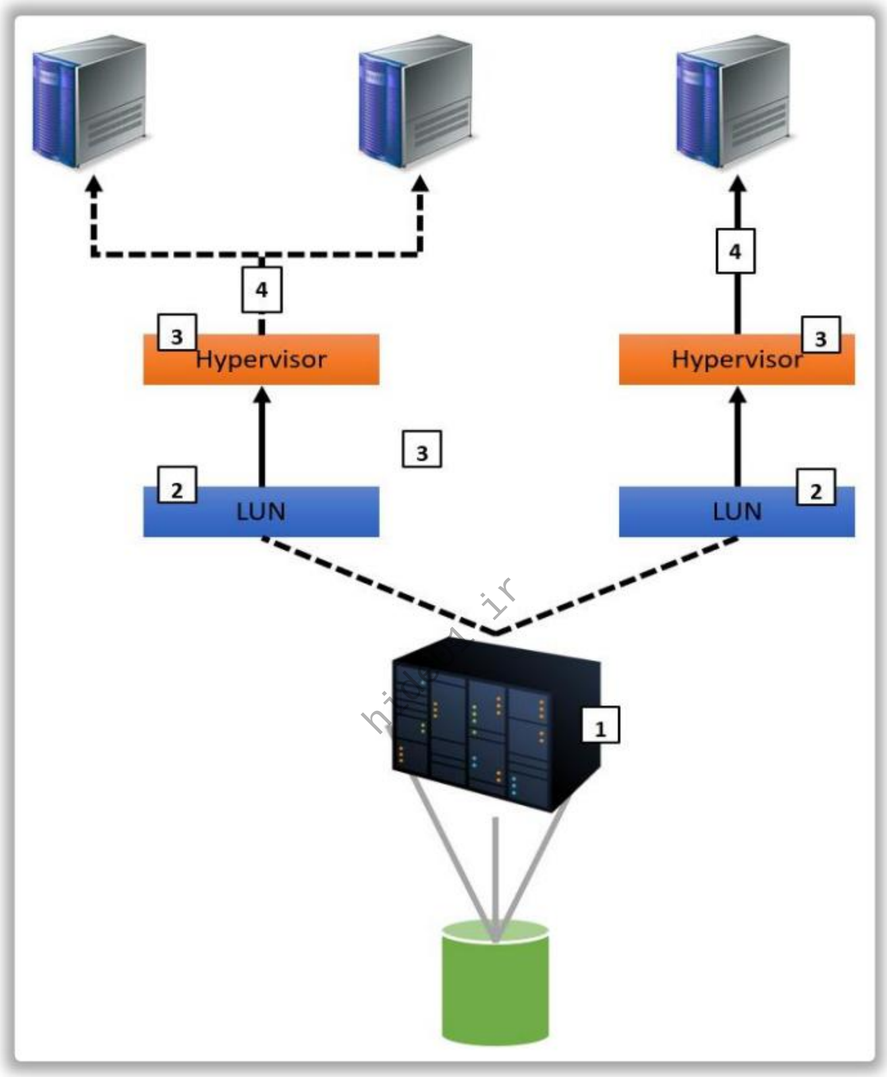


Figure 4.27: Direct LUN Access

- **Cross-Storage Device Vertical Tiering**

  Problem: Manual vertical scaling is needed to increase the processing capacity of data hosted in CS devices, but it is potentially unusable and inefficient. How can vertical scaling of data be done dynamically?

  Solution and application: This problem can be solved by establishing a system in which the vertical scaling of data can be done across several storage devices dynamically. LUN migration can be used to move LUN disks dynamically between cloud storage devices that have different capabilities with the help of predefined capacity thresholds.

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

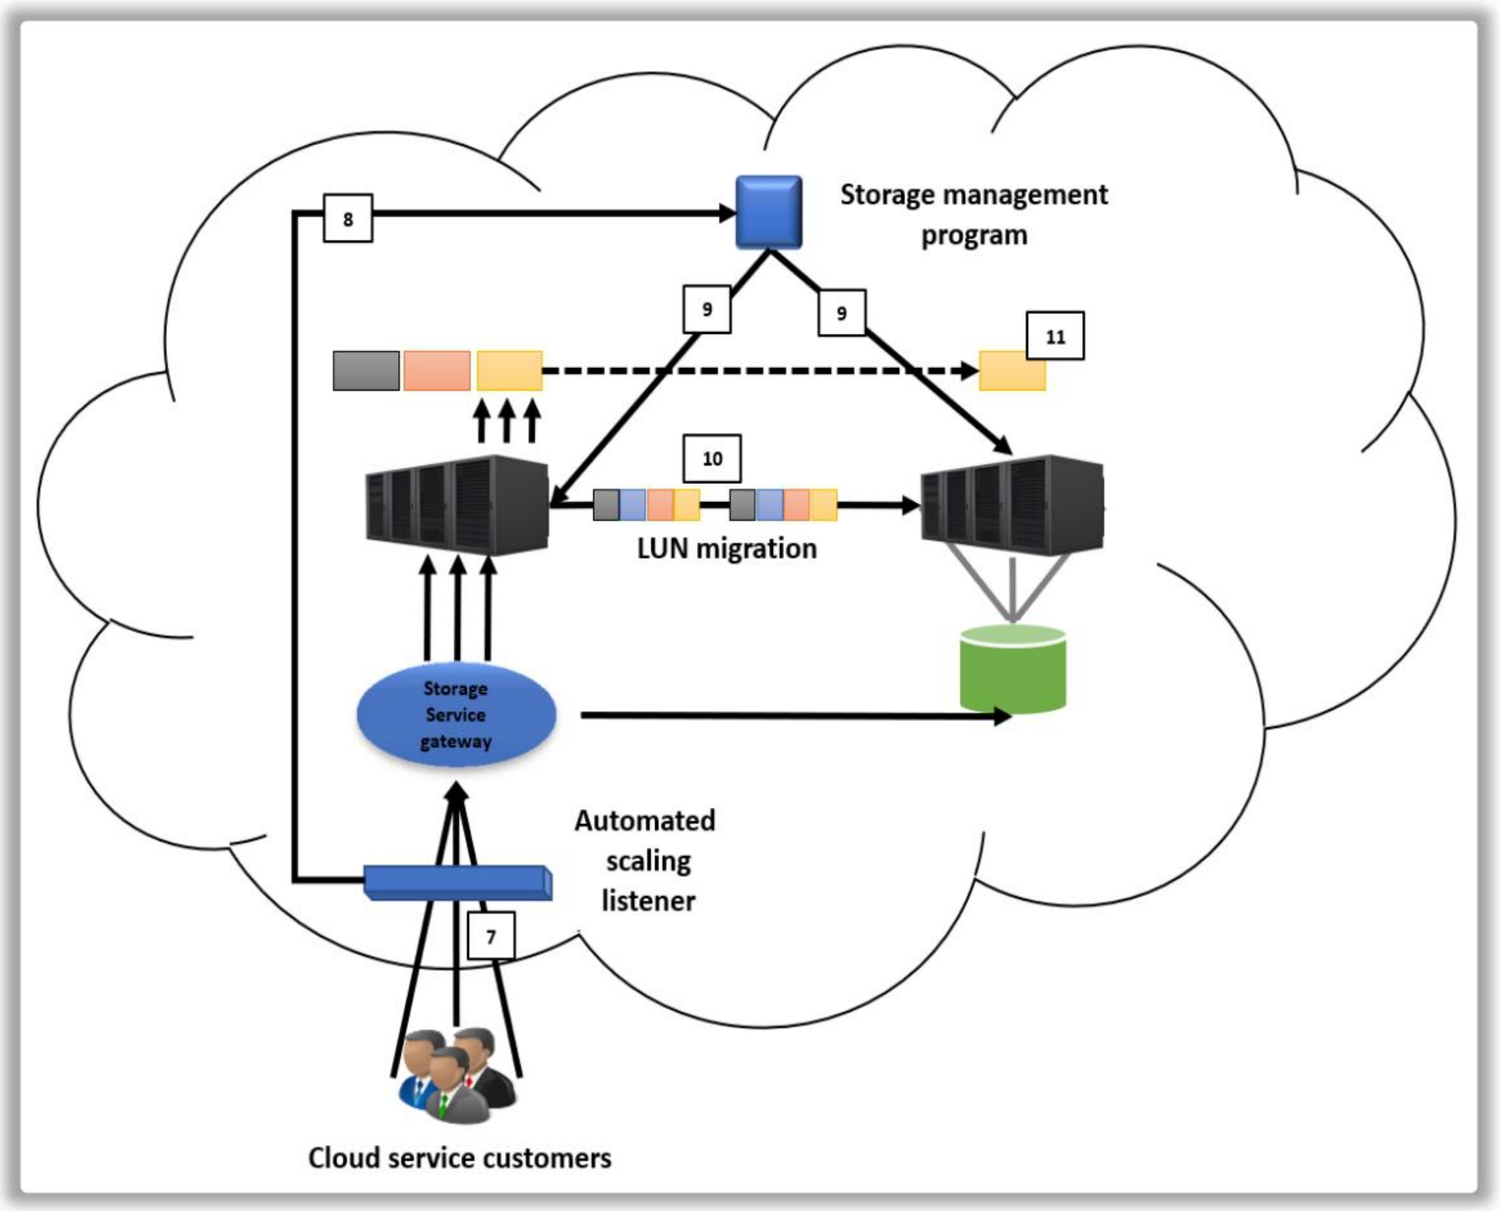


Figure 4.28: Cross-storage Device Vertical Tiering

- **Intra-Storage Device Vertical Tiering**

  Problem: The processing and storage capacity of data is limited to the storage and processing capacity of the CS device in which it is stored. How can the vertical scaling of the data be done inside a storage device?

  Solution and application: A storage device that can support multiple disk types must be used to allow dynamic vertical scaling. Complex cloud storage technology can be leveraged to create storage tiers. Through these tiers, data can be scaled (up or down as required) by using LUN migration.

Data Integrity and Security

- **Cloud Storage Device Masking**

  Problem: In a shared cloud environment, data is vulnerable to unauthorized access attempts by other cloud consumers. How can data residing on a storage device be restricted or isolated to certain specific consumers only?

  Solution and application: This problem can be fixed by using a solution that isolates every storage device from being accessible to unauthorized cloud customers. The LUN masking mechanism can be useful in implementing defined policies at the physical storage array to secure a particular storage device from being accessed by attackers or unauthorized consumers in the shared cloud environment.

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

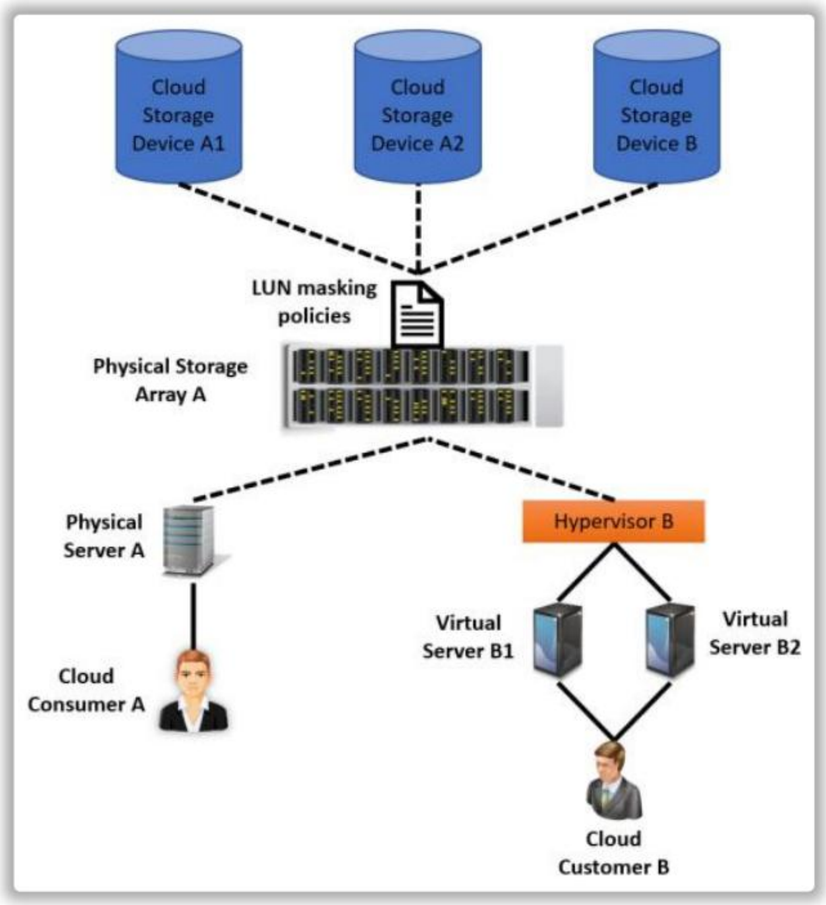


Figure 4.29: Cloud Storage Device Masking

- **Cloud Storage Device Path Masking**

  Problem: In a cloud environment, consumers can use alternative pathways to access cloud storage devices, which may compromise data integrity and security. How can data hosted on a storage device be isolated to particular customers through some pathways?

  Solution and application: To fix this problem, a solution must be implemented at the physical server layer or hypervisor layer to restrict access and usage of certain paths that are available. The storage path masking can be deployed within the hypervisor or OS installed on the physical server. This masking mechanism hides or masks two paths that were previously visible to Hypervisor B, as shown below.

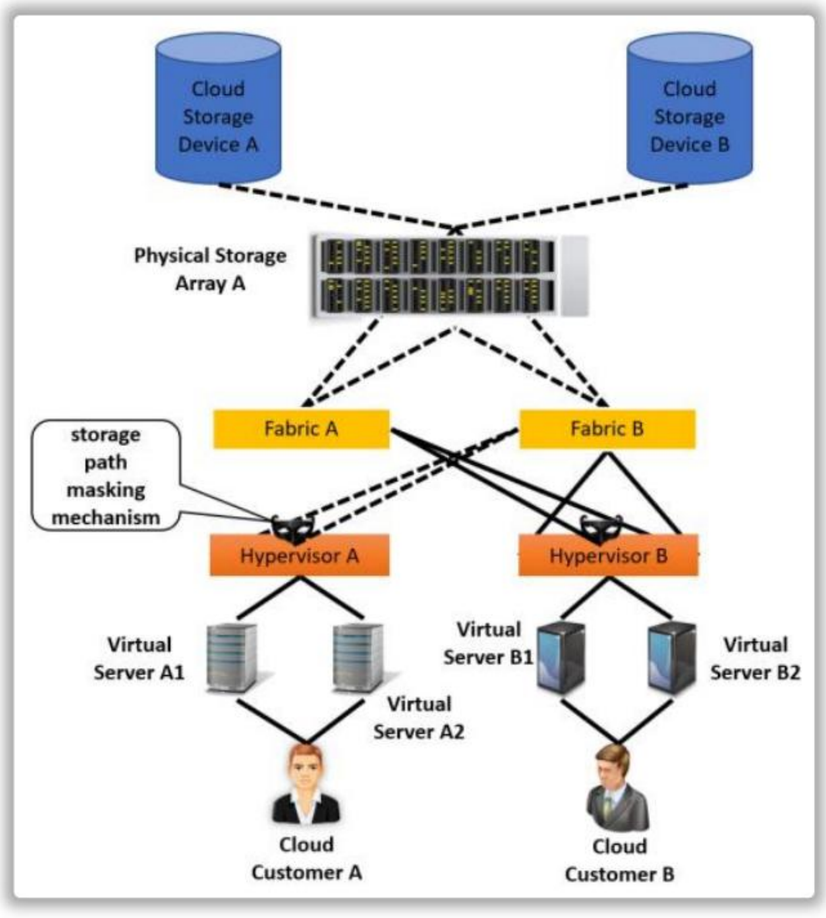


Figure 4.30: Cloud Storage Device Path Masking

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

- **IP Storage Isolation**

  Problem: The same underlying infrastructure is shared by both storage traffic and network traffic. How to secure IP storage when utilizing the same communication layer and channel used by other networking traffic?

  Solution and application: This problem can be fixed by implementing a security mechanism that can manage and provide appropriate isolation for storage devices. A networking mechanism can be used to secure IP storage data from unauthorized users. This mechanism can be utilized to assign a virtually dedicated traffic tunnel to IP storage.

- **Cloud Data-at-Rest Encryption**

  Problem: In a cloud environment, data must be stored securely in the physical disks (referred to as cloud storage devices) and these disks should be protected against unauthorized access. What can cloud providers do to securely store consumers' data on cloud storage devices?

  Solution and application: Data stored on the disks must be secured to prevent unauthorized access. An encryption mechanism that suits the physical storage arrays must be used to encrypt data that is stored on the disks and decrypt it when it leaves the disks automatically. This can secure data stored on the disks.

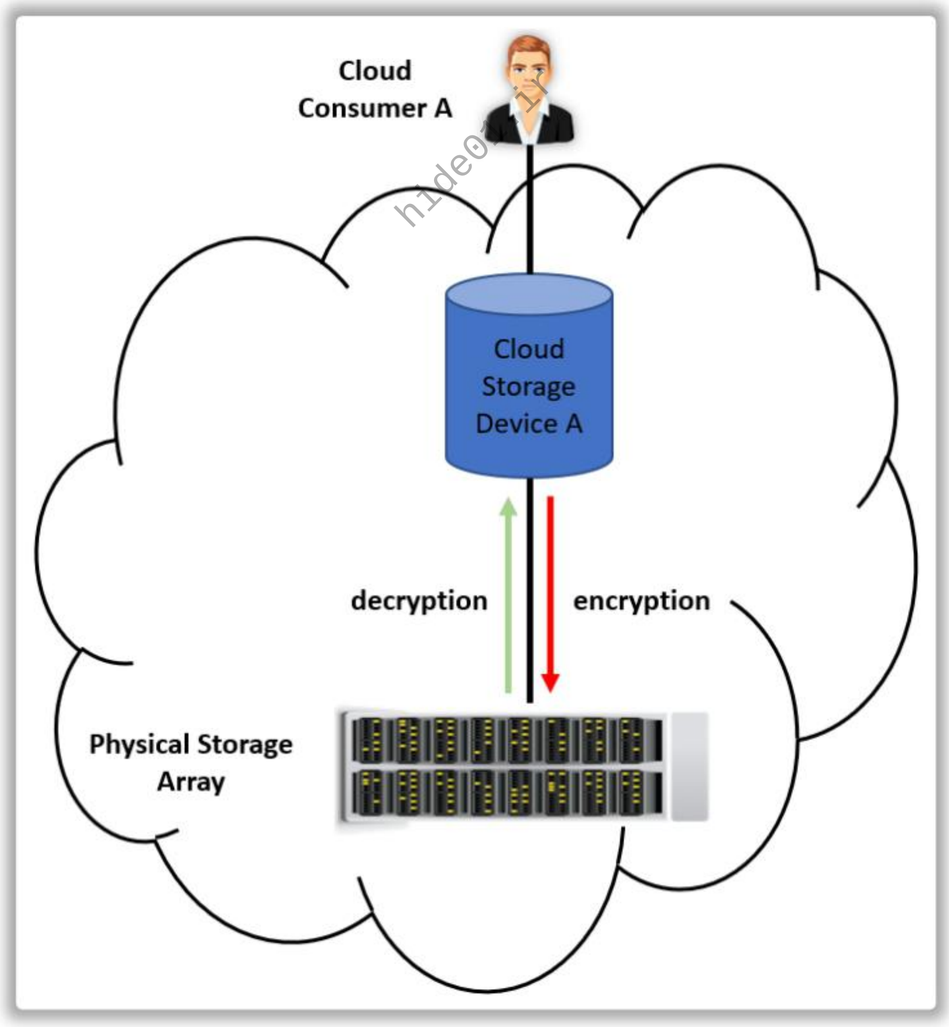


Figure 4.31: Data-at-Rest Encryption

- **In-Transit Cloud Data Encryption**

  Problem: Data that is copied from and to a cloud environment travels via a network and servers where organizations cannot secure it, which may allow attackers to intercept it. How can data be transferred from, to, and within a cloud environment securely?

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

Solution and application: This problem can be fixed by using a solution that can secure data when it travels from the sender to the receiver and can also make sure that the receiver does not accept the data if it is modified. An encryption mechanism can be used to encrypt data for confidentiality when it travels to the receiver from the sender and a digital signature mechanism must be implemented to ensure data integrity.

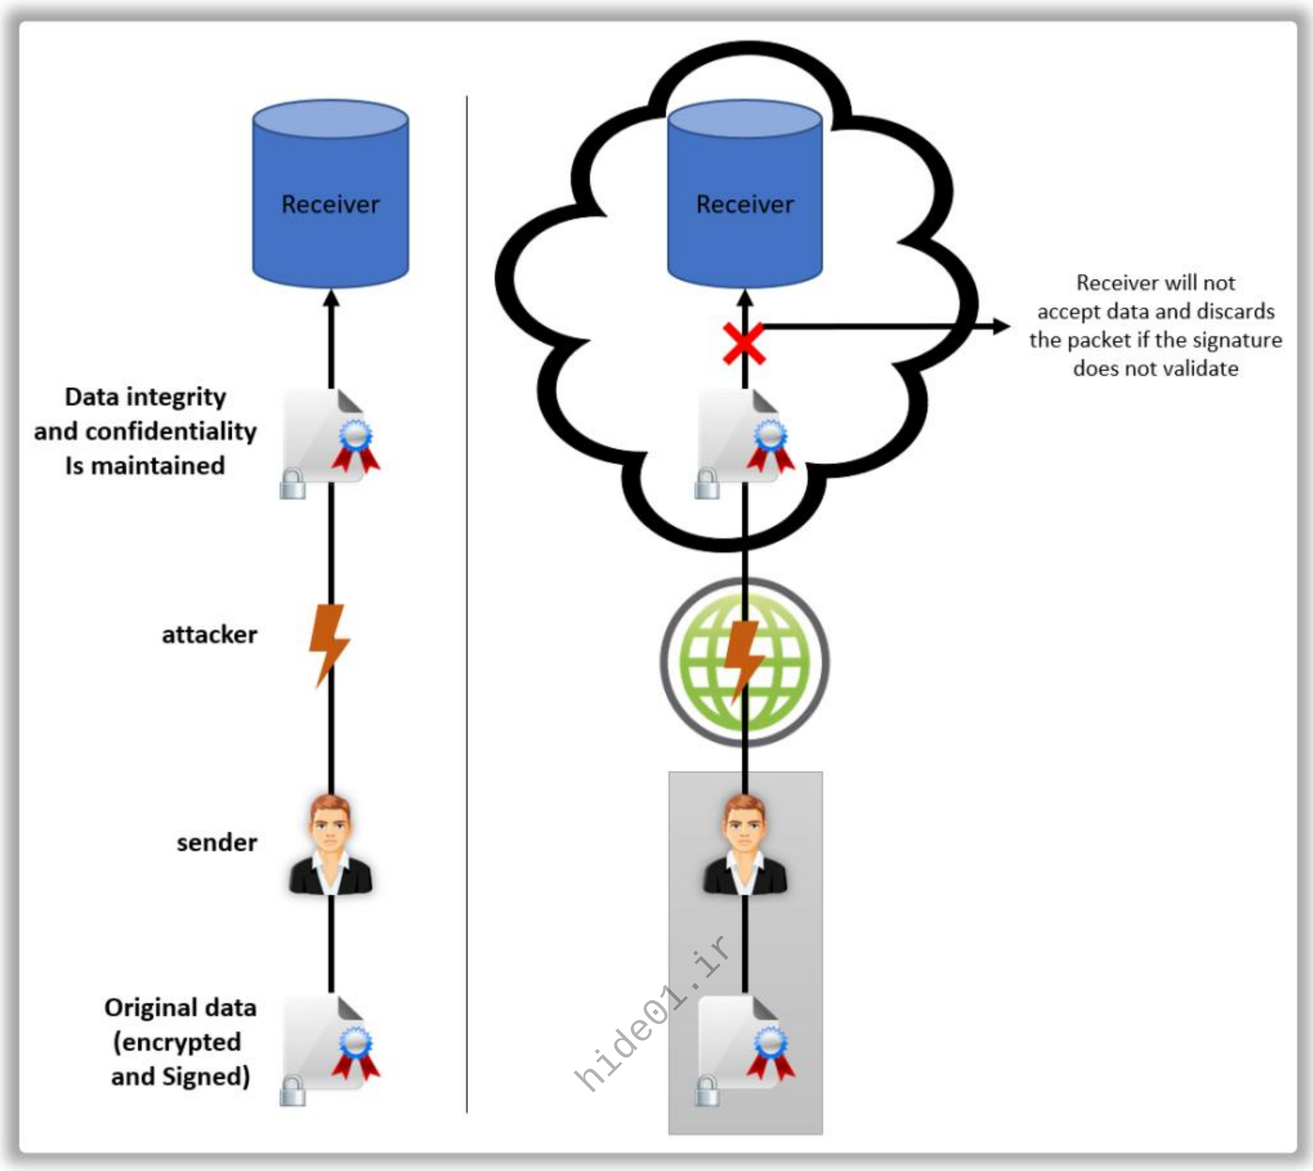


Figure 4.32: In-transit Data Encryption

The data is encrypted as well as signed before it is sent to the receiver. Attackers may try to intercept data while it is in transit. The receiver discards the data packet if it is intercepted by the attacker to ensure data integrity and confidentiality.

Reliability and Failover

- **Redundant Storage**

  Problem: Cloud consumers cannot access data if cloud storage devices become inaccessible or fail. How can the availability and reliability of storage devices be ensured under such failure conditions?

  Solution and application: A fail-safe system consisting of extra storage devices must be used. This ensures that the secondary storage devices will handle all access requests when the primary storage devices become inaccessible or fail. Data gets replicated to the secondary storage devices from the primary storage devices. To redirect data access requests to the secondary storage (when required), a storage device gateway is used.

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

- **RAID-based Data Placement**

  Problem: Cloud storage devices have different RAID levels. Data stored on a particular RAID level that is not compatible with its protection or performance requirements may make that RAID underperform or fail. How can data be hosted on a RAID level that is compatible with its required or specified protection and performance levels?

  Solution and application: An interface must be provided to cloud consumers, which informs them about the existing RAID levels of cloud storage devices. A RAID identifier mechanism can help cloud consumers understand the available RAID levels for a cloud storage device they select by directly communicating with the storage management portal mechanism.

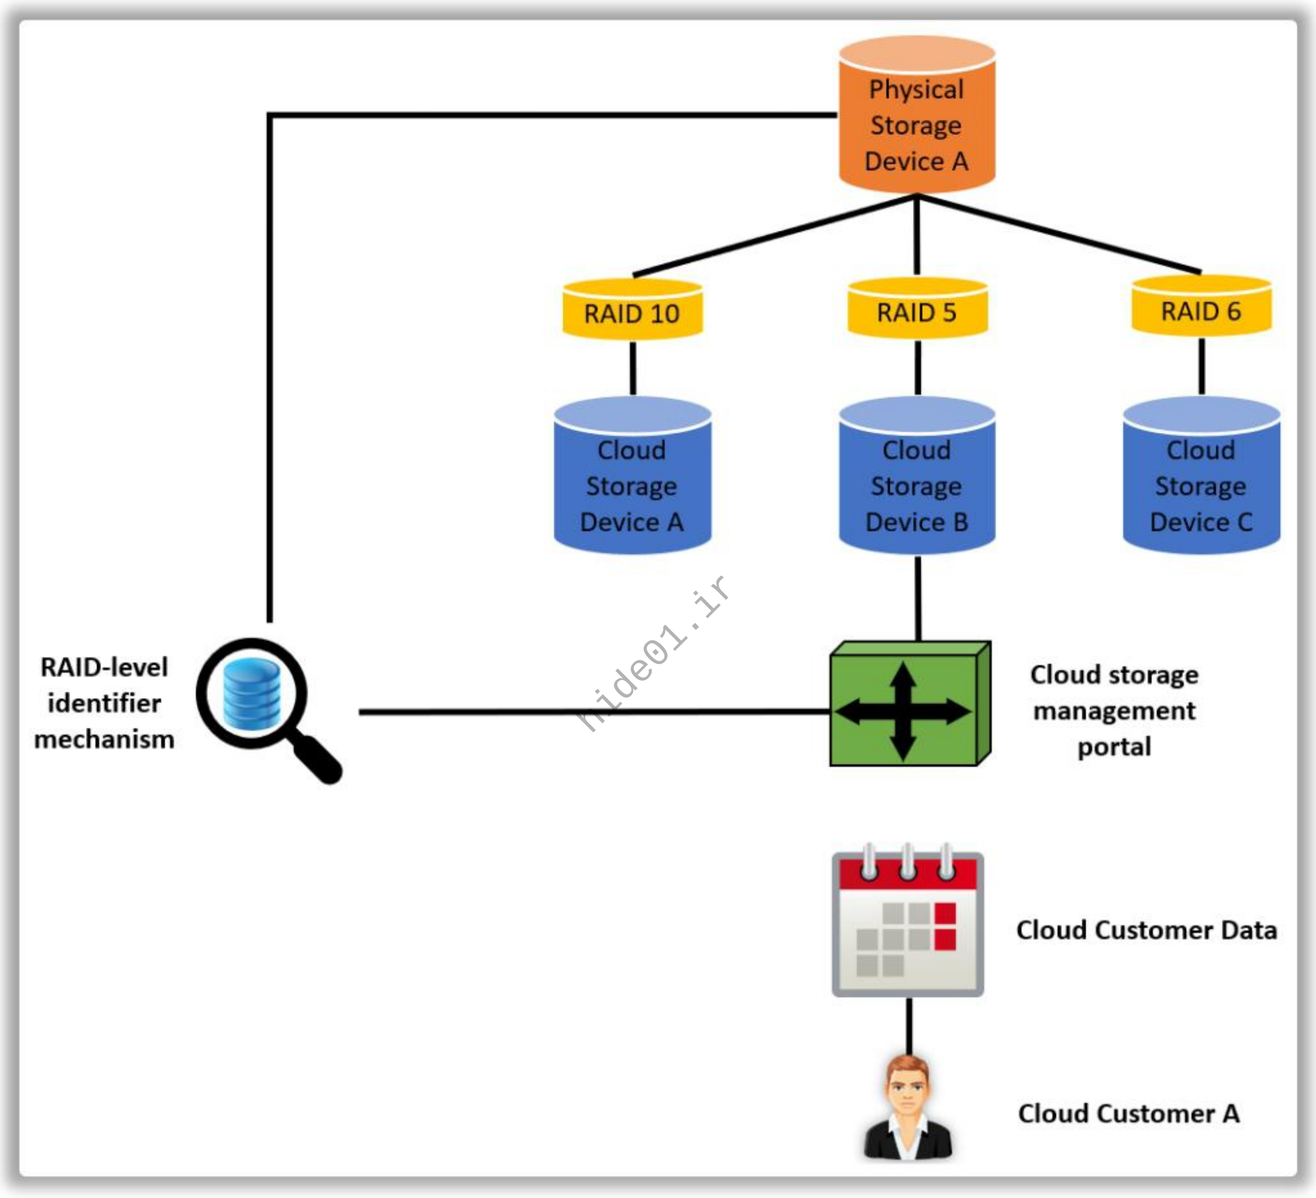


Figure 4.33: Data Placement Mechanism

The RAID identifier mechanism provides information on the particulars of cloud storage devices by using the cloud storage management portal mechanism.

- **Multi-path Resource Access**

  Problem: An IT resource is lost when the path to that resource becomes unavailable or is lost, which will affect the stability of the entire cloud environment. What can be done to access an IT resource when its path becomes unavailable or is lost?

  Solution and application: Cloud consumers are provided with alternative paths to access IT resources to overcome path failures. A multi-pathing system running on the hypervisor or the server is utilized to offer multiple alternative paths to access the same IT resources.

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

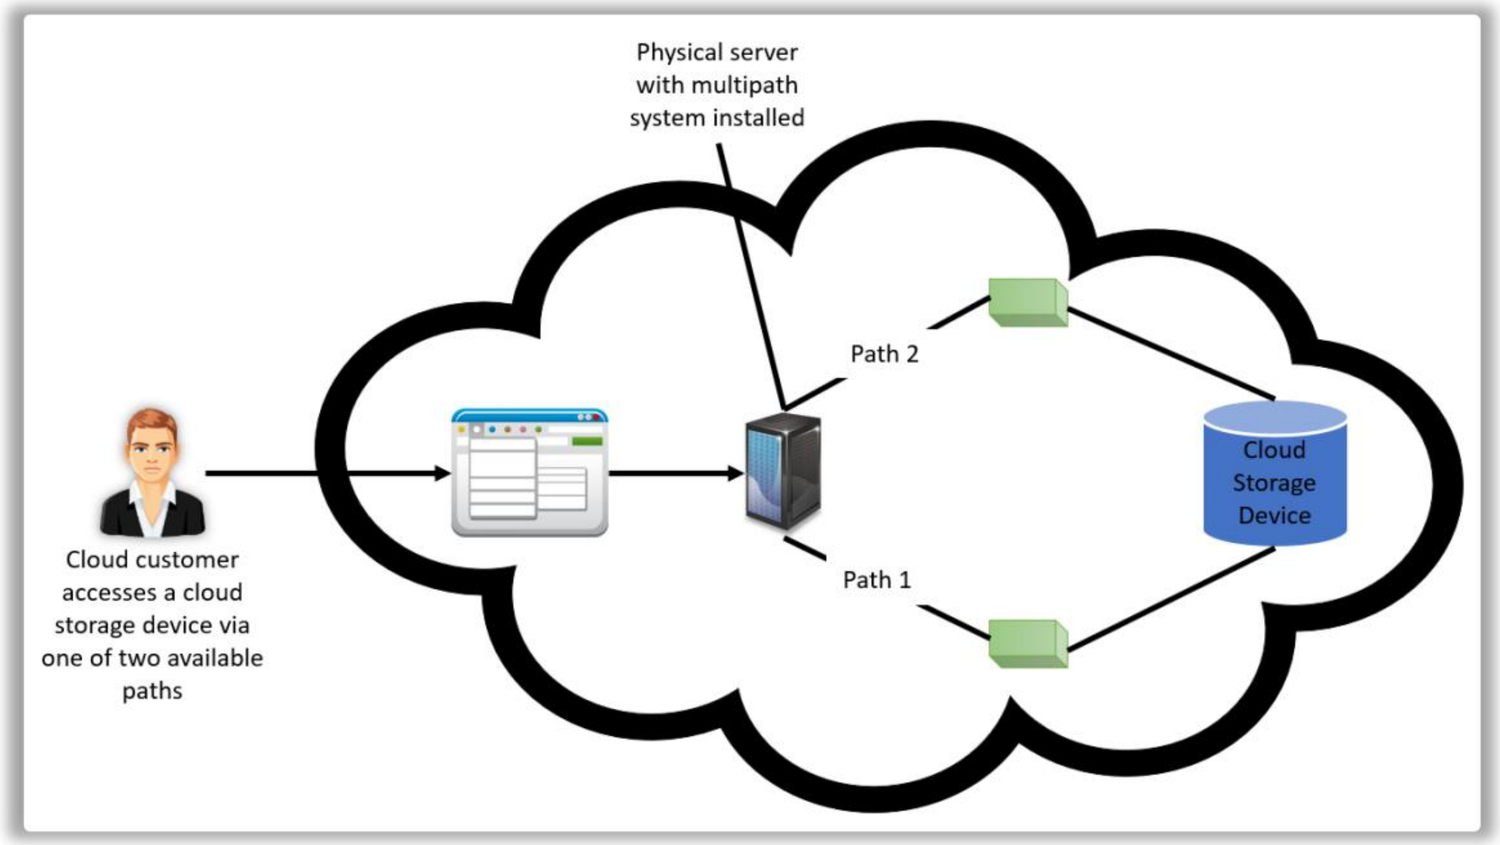
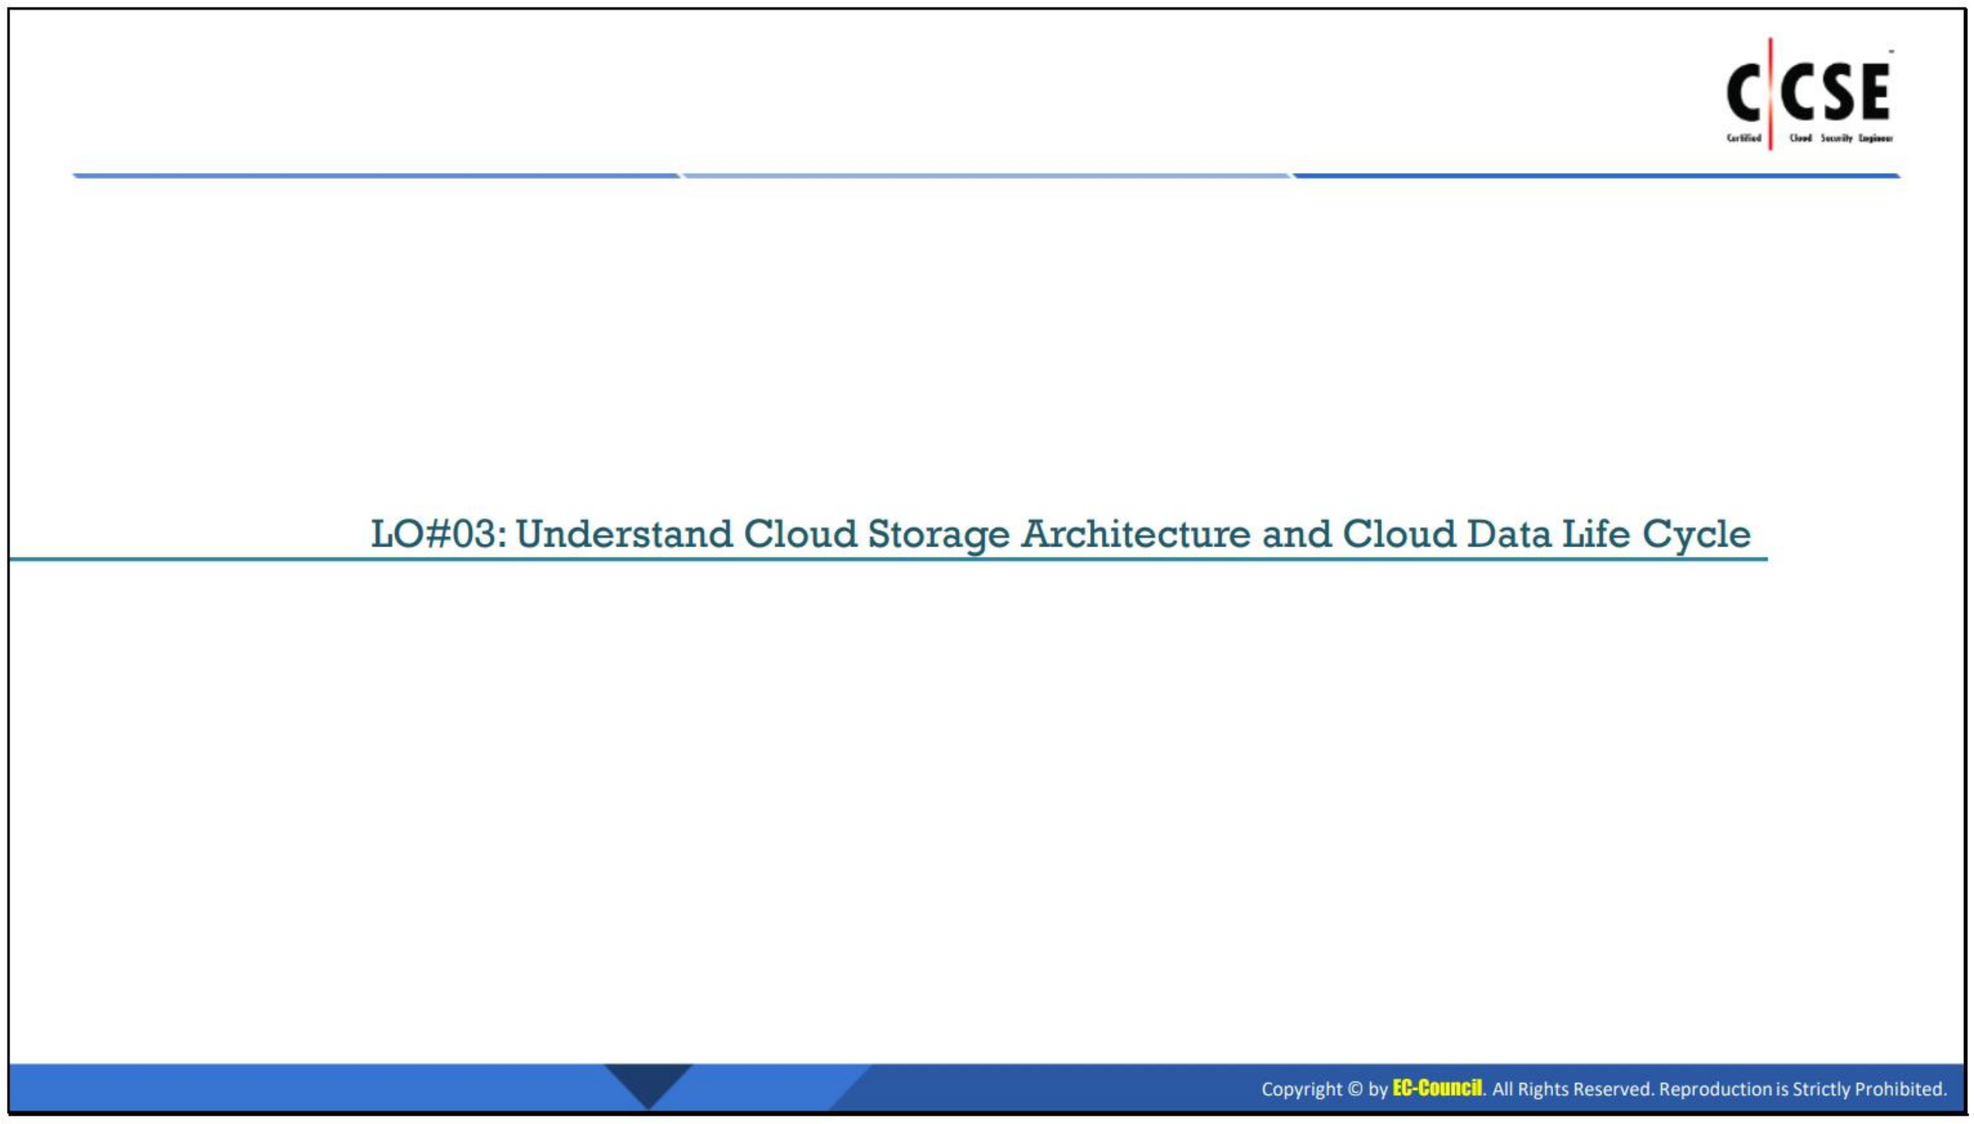


Figure 4.34: Multi-pathing System

hide01.ir

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

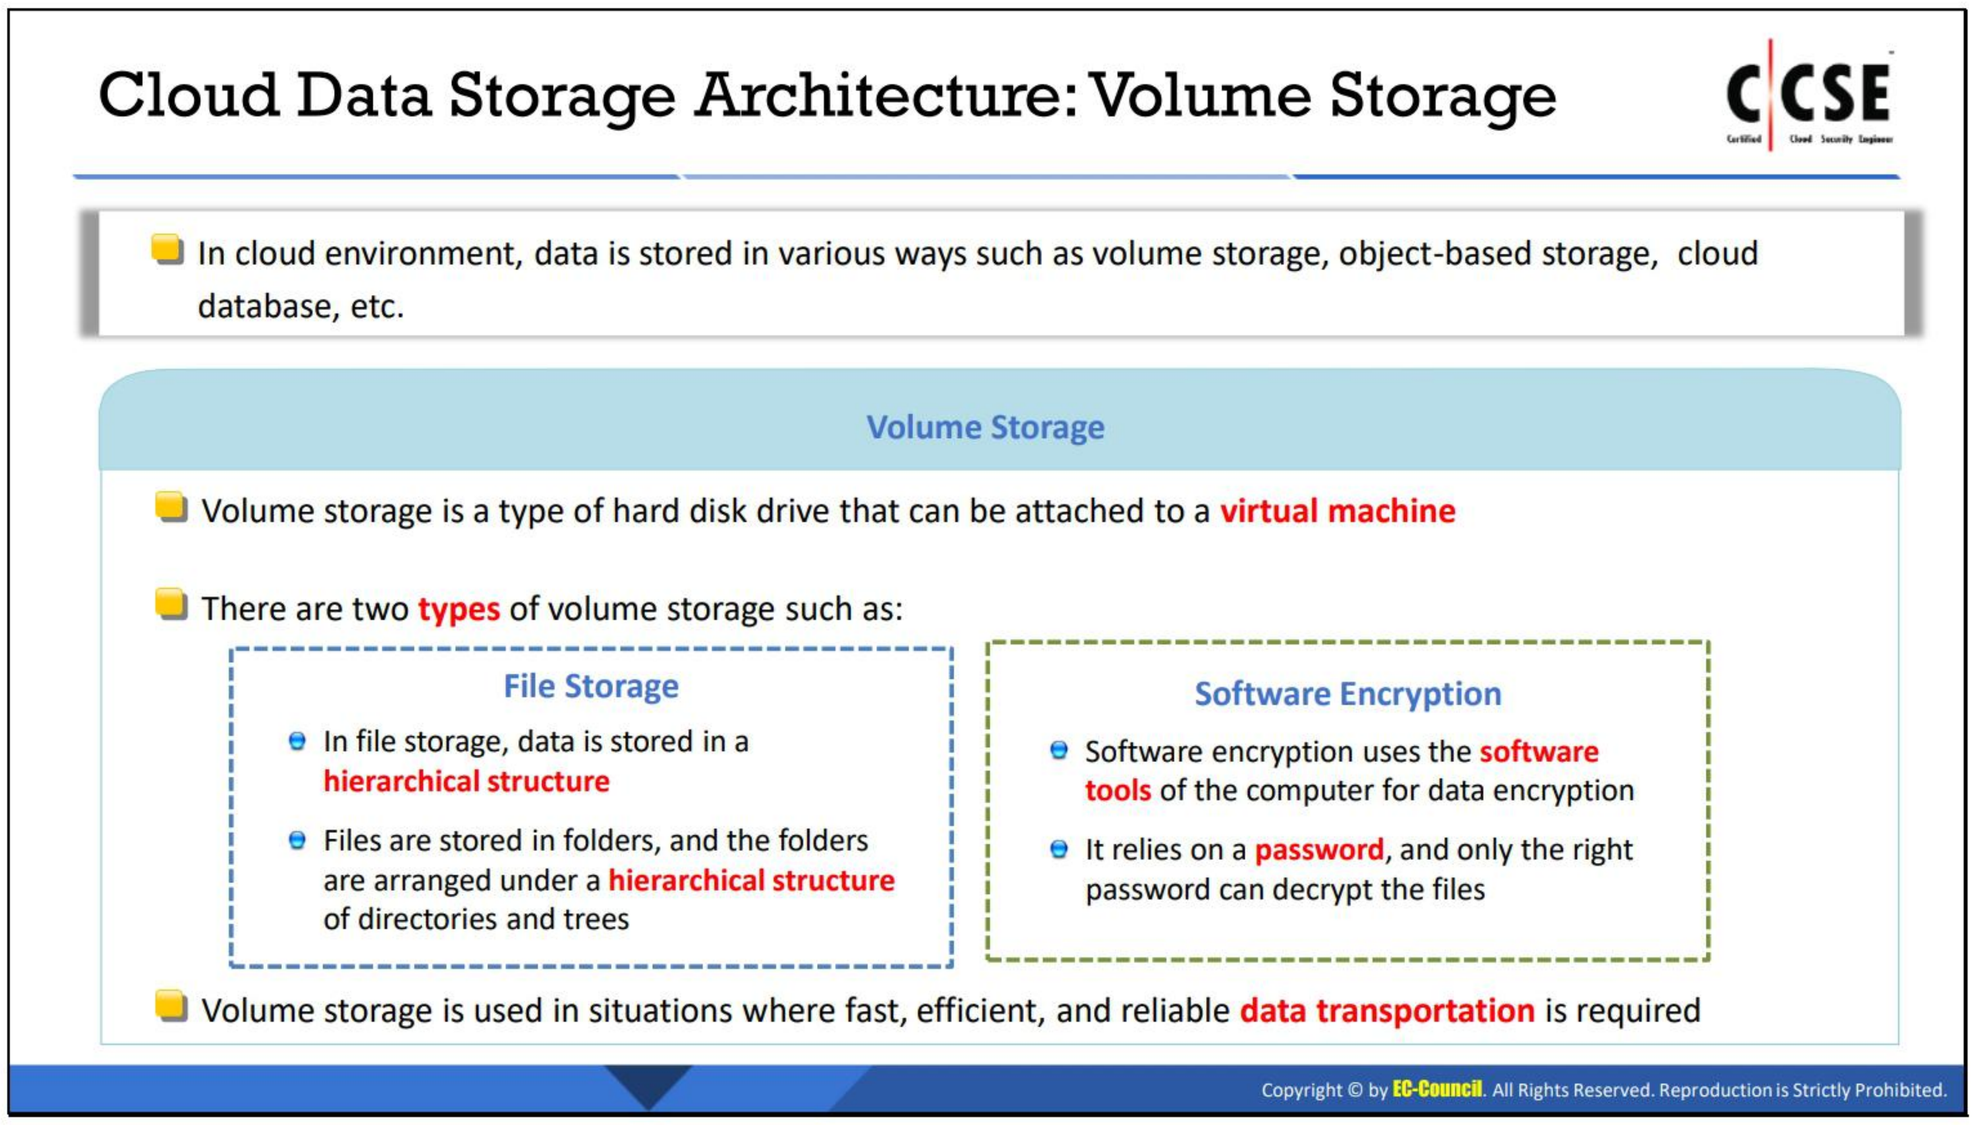


LO#3: Understand Cloud Storage Architecture and Cloud Data Life Cycle

The objective of this section is to describe the different ways of storing cloud data such as object storage, volume storage, cloud database, raw storage, long-term storage, and ephemeral storage. It also explains the different phases of cloud data lifecycle.

hide01.ir

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



Cloud Storage Architecture: Volume Storage

In a cloud environment, data is stored in various ways such as volume storage, object-based storage, cloud database, etc.

hide01.ir

Volume storage is a type of hard-disk drive that can be attached to a VM. The storage space is allotted within the cloud. This virtual drive is like a physical drive attached to a device. Data can be hosted through this drive within the same file system. Volume is the main component of volume storage rather than the VMs. The volumes can be detached and attached from one instance to another. The volumes can work with any operating system.

Volume storage supports various operations such as read/write, and keeps the files of running VMs or instances. The underlying hardware of volume storage includes independent storage nodes that offer standard or SSD disks.

There are two types of volume storage.

- **File storage:** In file storage, data is stored in a hierarchical structure. Files are stored in folders, and the folders are arranged under a hierarchical structure of directories and trees. Users can create, delete, read, and write files and can arrange them under directories. It offers file-sharing feature and the users do not need to write code to have shared file access. It can be used for the direct migration of applications to the cloud that need file-sharing protocols. File storage enables an organization to reuse the technical resources for other projects, which reduces costs.
- **Block storage:** Block storage chunks data into pieces called “blocks.” It is an efficient way to store, use, and manage data. Block storage is designed to dissociate data from the storage environments, and the data is dispersed to multiple environments. This allows the underlying storage system to effectively retrieve and reassemble the data irrespective

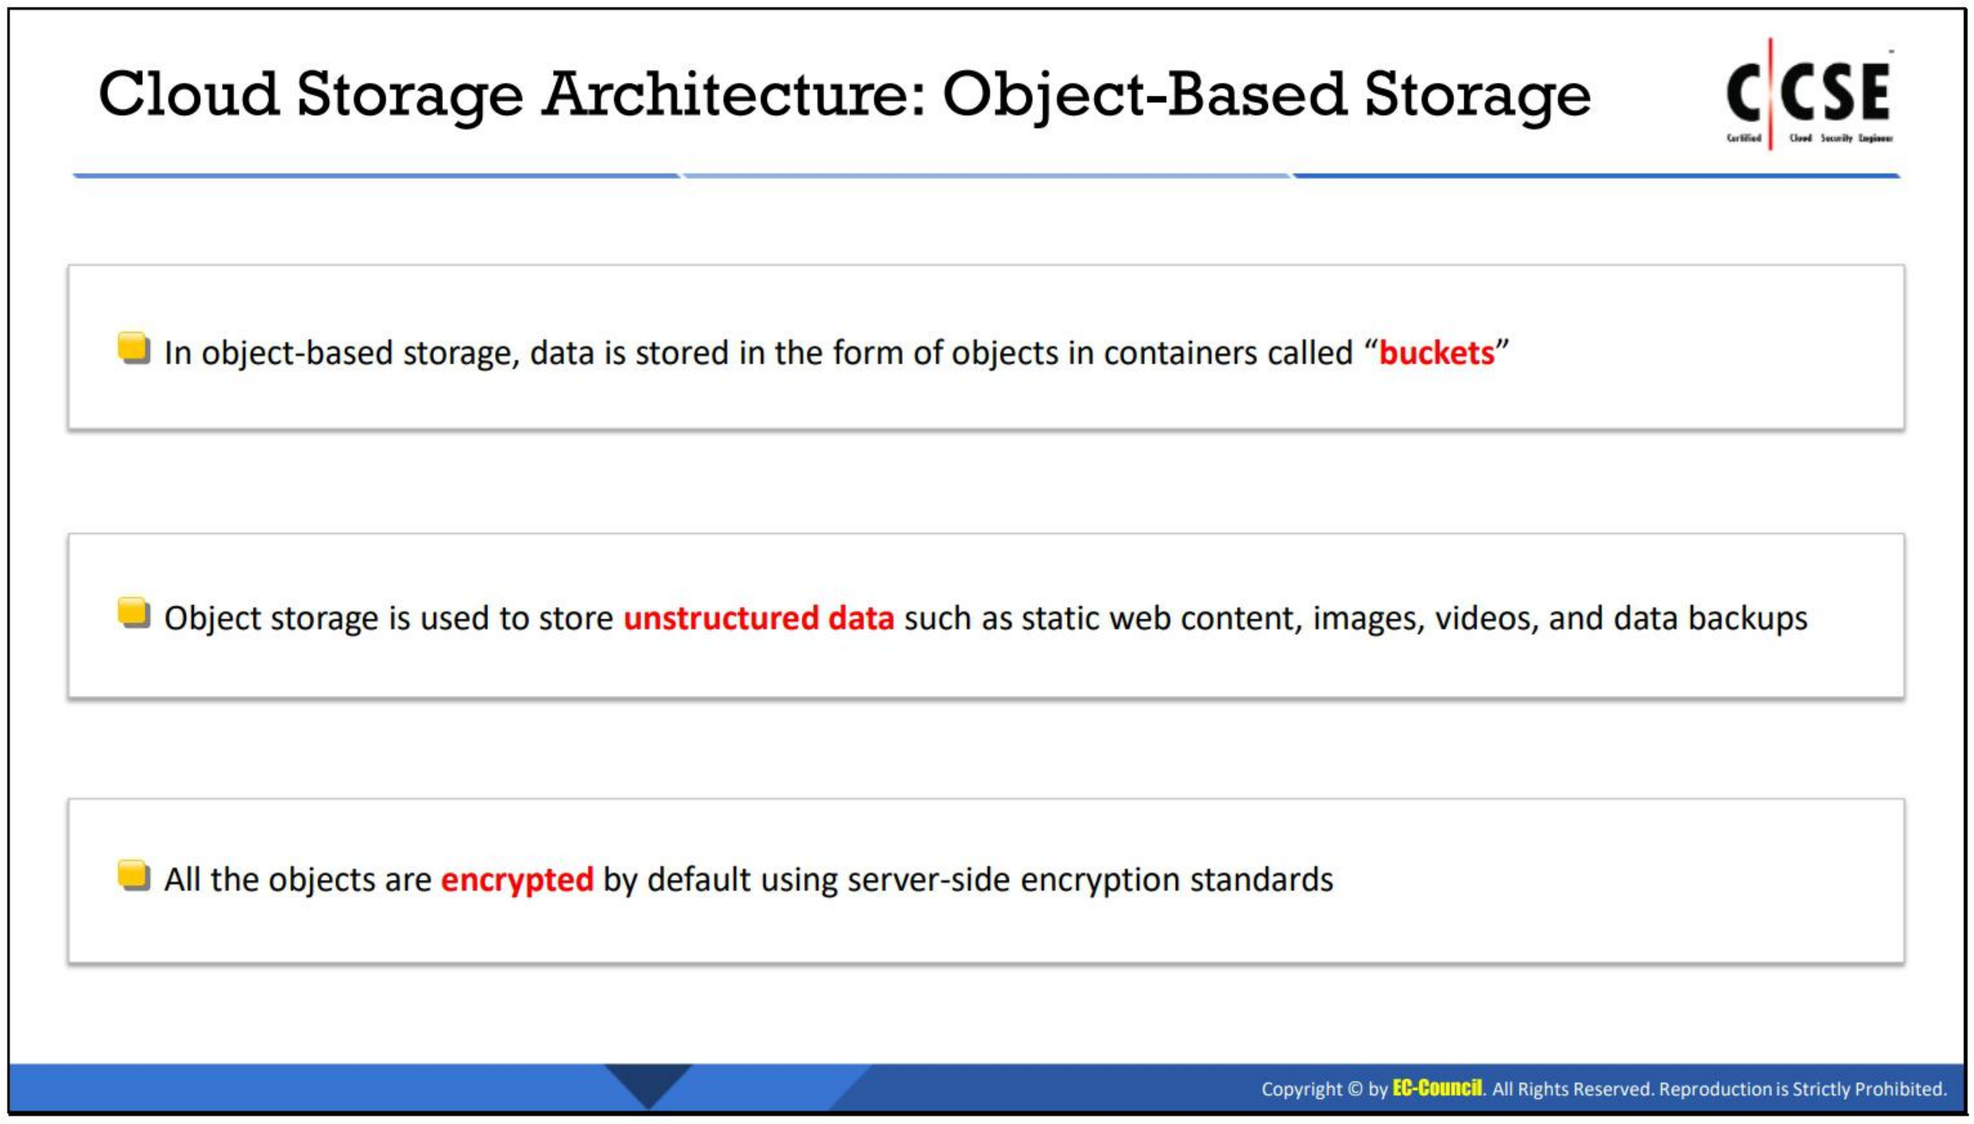
Certified Cloud Security Engineer Copyright © by **EC-Council**
All Rights Reserved. Reproduction is Strictly Prohibited.

of where it gets stored. Hence, it makes the users retrieve data quickly. It is used by organizations that perform large transactions. Examples of block storage include SAN, RAID, virtual machine file system (VMFS), and iSCSI.

Volume storage has a snapshot feature through which an organization can back up data and use it in case of disaster recovery. It is highly suitable for high-performance applications that should have low latency and consistent I/O performance. It is also used in situations where fast, efficient, and reliable data transportation is required.

hide01.ir

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



Cloud Storage Architecture: Object-Based Storage

Object-based storage is different from volume storage; here, the data is stored in the form of objects. Objects are stored in containers called "buckets." Object storage can be enabled at various levels such as device level, interface level, and system level.

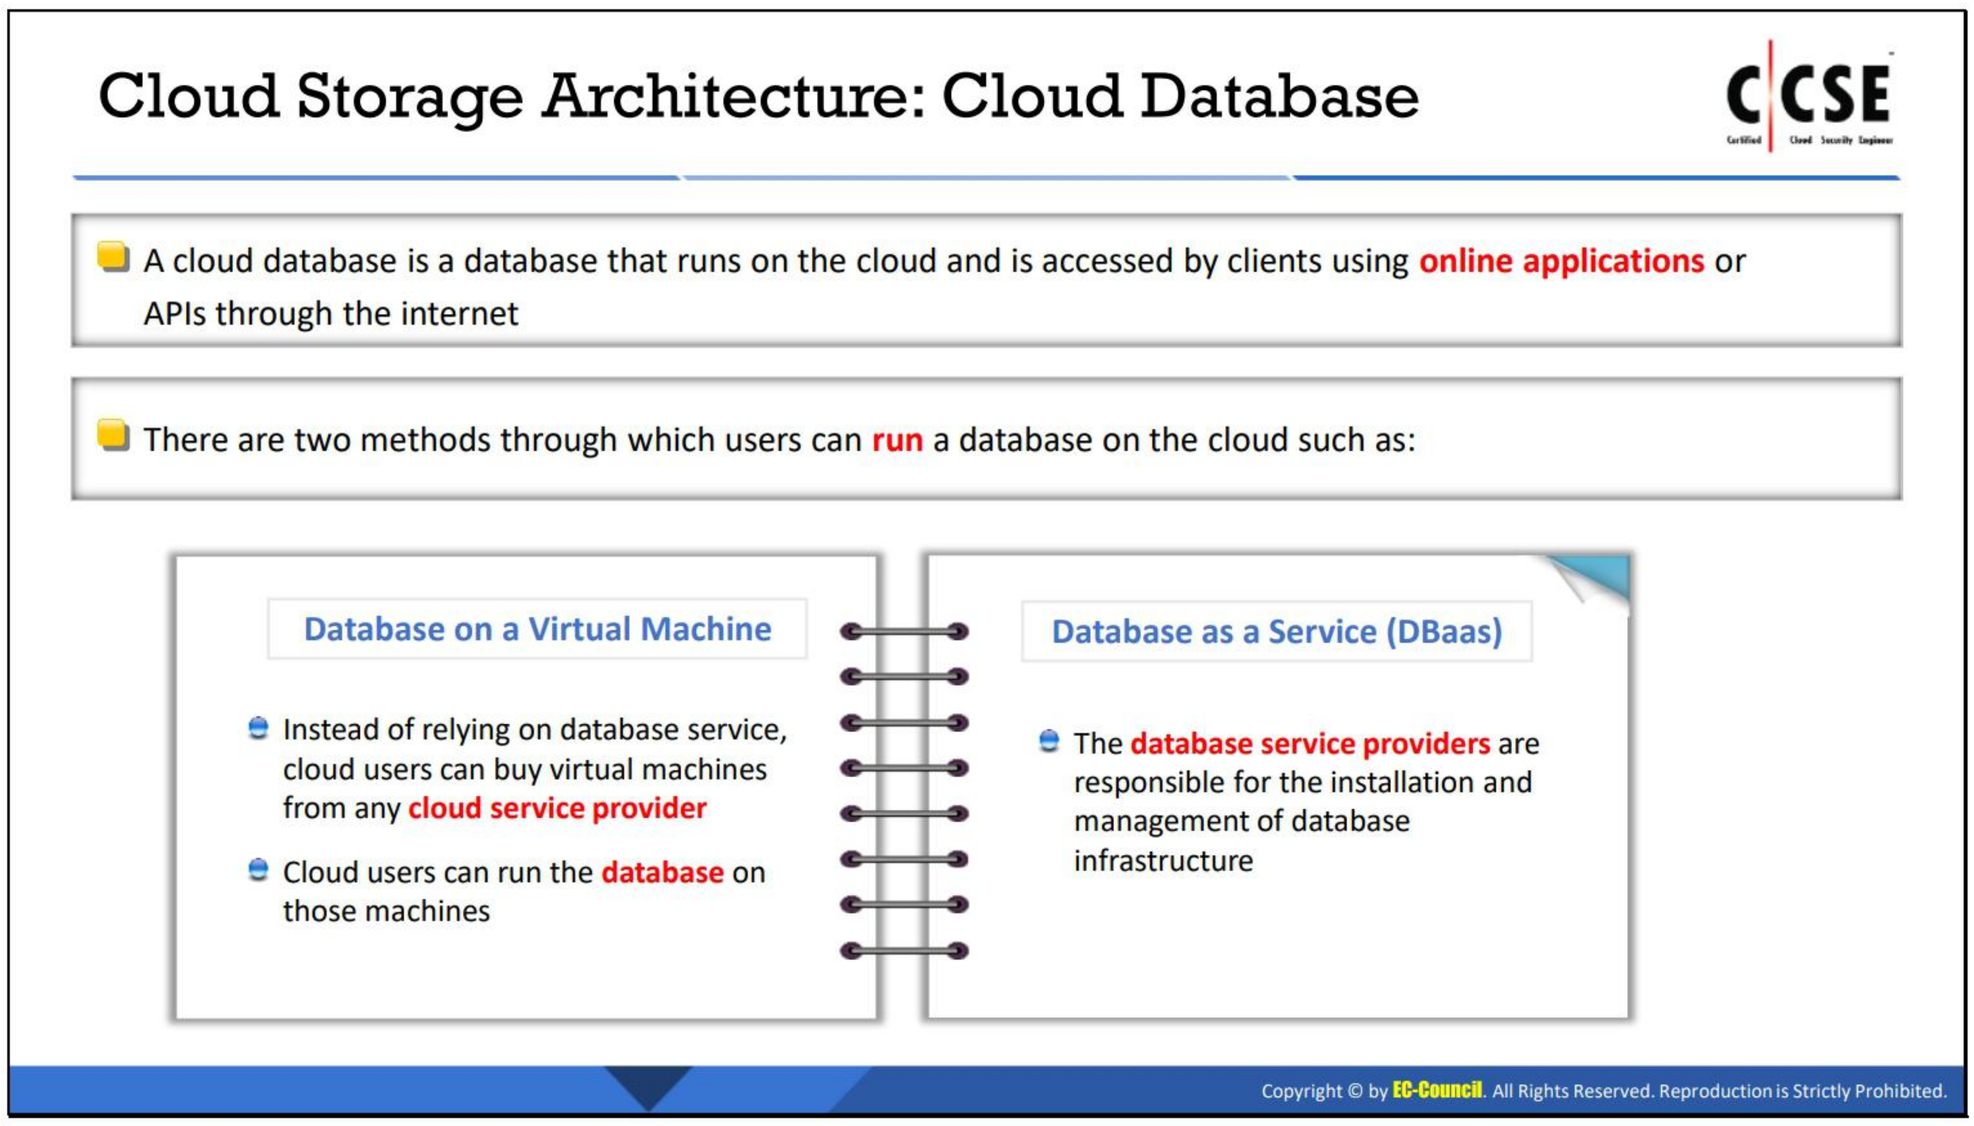
Objects itself consists of data, and it can be of any format. Each object has metadata and a unique identifier. Metadata plays an important role in the success of object storage as it contains complete information about the usage and functions of the data. Metadata is described by the person who created the object storage, and it contains information such as the type of data, how it can be used, etc. The unique identifier is a 128-bit value, and it is used to identify the object over a globally distributed system.

Object storage is used to store unstructured data, which is written once and read many times. The examples include static web content, images, videos, and data backups.

Object storage offers high scalability, and the users can keep adding the data. Cloud users can start with a small piece of data and then scale up without affecting performance. As there is no folder hierarchy in object storage, the data-retrieval process is also fast. Object storage enables the cloud users to optimize resources as it does not have a file hierarchy.

Object storage enables cloud users to securely store and retrieve the data from the cloud platform or the internet. It can be accessed from any location in the world. All the objects are encrypted by default using server-side encryption standards. Each object has a separate data encryption key, and the data encryption key is encrypted with the bucket's encryption key. Cloud users can access object storage through console, command-line interface, and REST API.

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



Cloud Storage Architecture: Cloud Databases

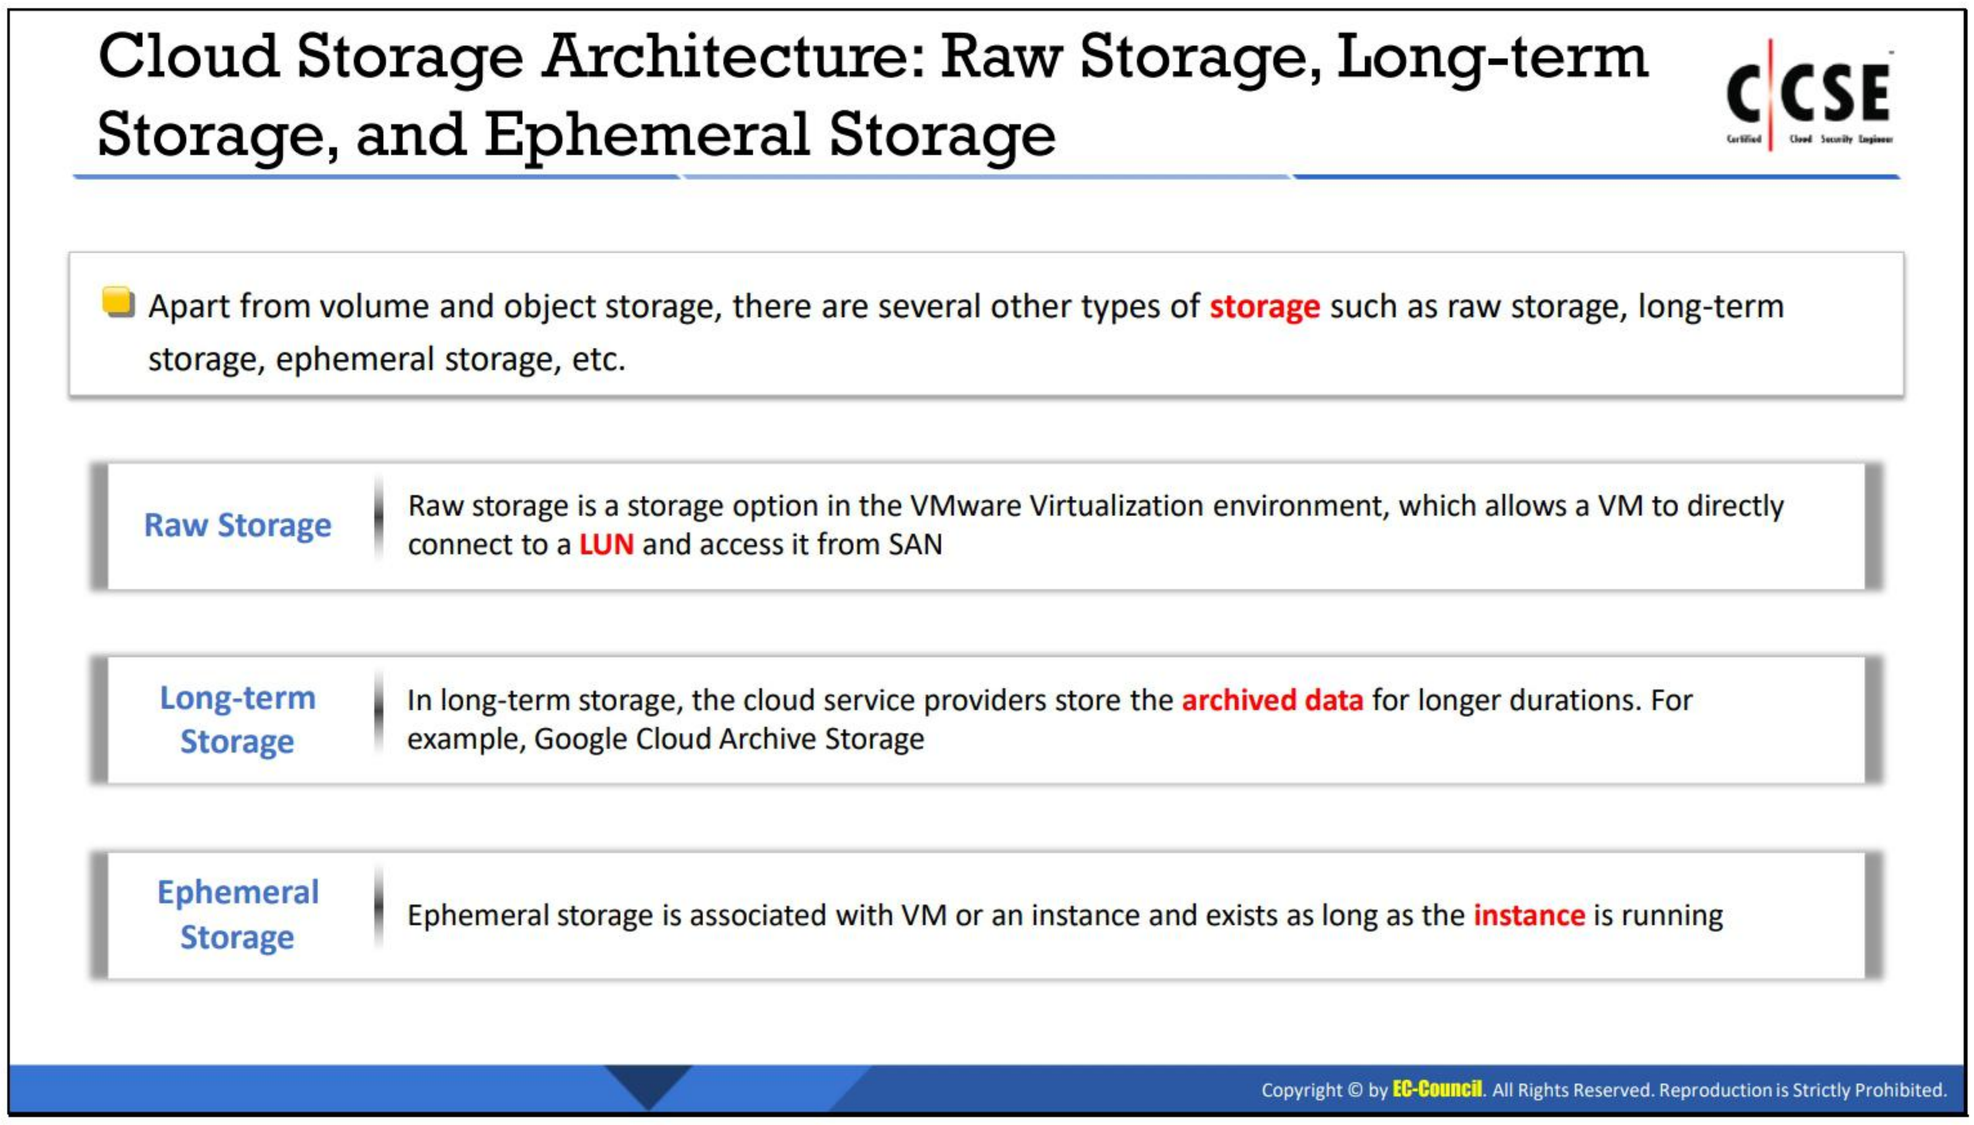
A cloud database runs on the cloud and is accessed by clients using online applications or APIs through the internet. It is the back-end storage in the data center. Databases can be installed on either object storage or file storage.

There are two methods through which users can run a database on the cloud.

- **Database on a VM:** Instead of relying on database service, cloud users can buy virtual machines from any cloud service provider and can run the database on those machines. Users can use either available machine image or their own machine image and then install a database on it. This database works exactly as the on-premises, and the infrastructure is virtualized. Using this, users can either increase or decrease the storage based on the workloads. Cloud provides security, data backup, and recovery options, and the users do not need to worry about availability and security.
- **Database as a service (DBaas):** DBaaS reduces the burden of application owners to install and manage the database. The database service providers are responsible for the installation and management of database infrastructure. This enables the users to focus on their tasks. The application owners will be charged based on the usage of the service. Using DBaaS, business owners can quickly set up their database. DBaaS can be accessed from any location in the world, and it offers high reliability, and performance.

  Cloud database is highly suitable for enterprises that need large volumes of data with their software products distributed in various locations.

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



Cloud Storage Architecture: Raw Storage, Long-term Storage, and Ephemeral Storage

Apart from volume and object storage, there are several other types of storage such as raw storage, long-term storage, and ephemeral storage.

hid991.ir

Raw Storage

Raw storage, also known as raw-device mapping (RDM), is a storage option in the VMware Virtualization environment. It allows a VM to directly connect to a LUN and access it from SAN. RDM is a mapping file that allows the virtual machine to bypass VMFS and directly access the physical storage device.

RDM contains metadata that is responsible for managing and redirecting the disk access to the VM. RDM is often described as mapping a device into disk volumes or datastore.

An RDM file gets created within the VMFS volume when a Vsphere administrator configures the virtual disk (vDisk) of the VM to point to a LUN. Using RDM, any VM in that cluster can access the storage array.

The VM reads the mapping file that contains the physical storage device's raw ID. It refers to the ID and reads the data and writes the data to the storage device. In this way, a VM considers RDM as a SCSI that is used for normal file operations.

RDM offers two compatibility modes:

- Physical compatibility mode that allows minimal SCSI virtualization of the mapped raw LUN. It offers high flexibility to SAN management software. It is useful for running SAN agents.

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

- Virtual compatibility mode offers full virtualization of the mapped raw LUN. It can be used for SAN snapshots and advanced file locking.

Long-Term Storage

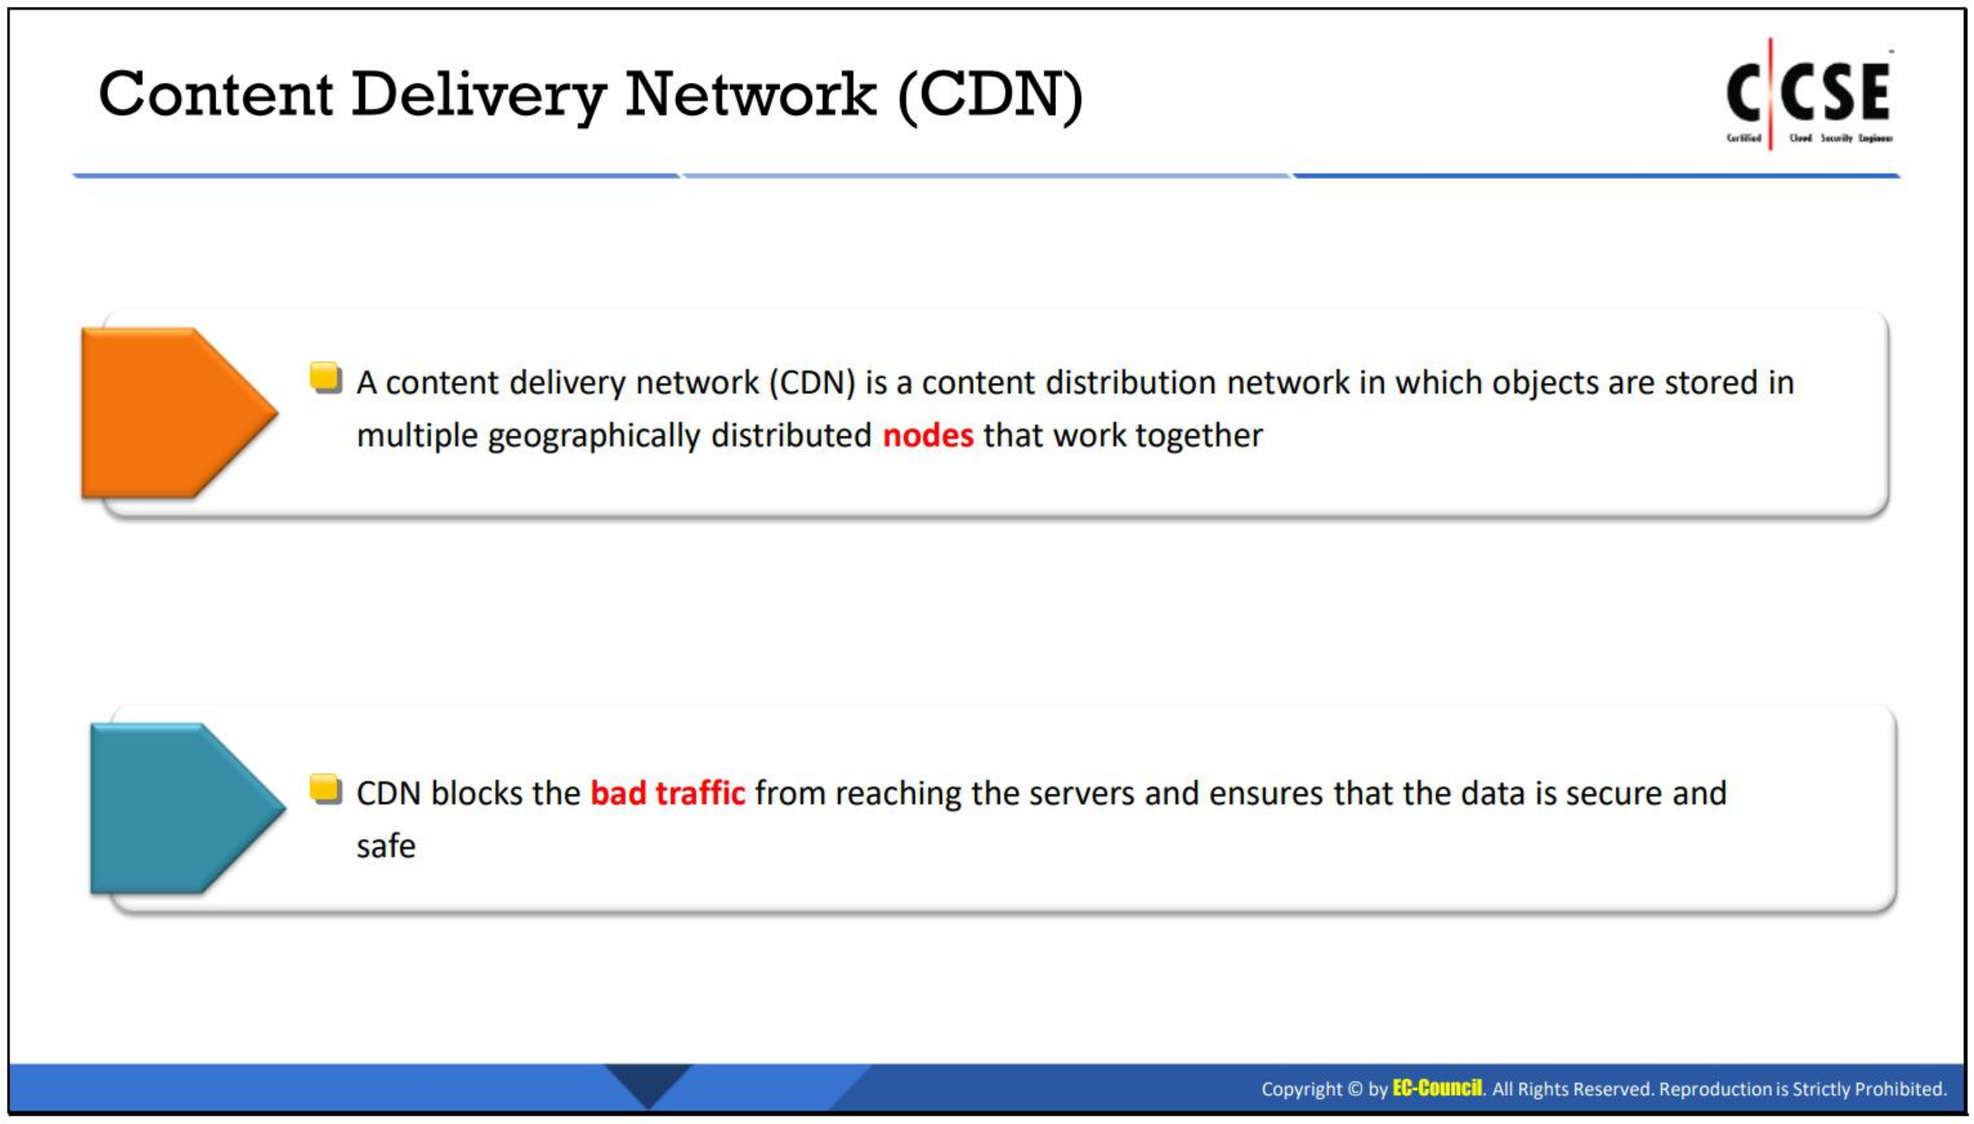
Many CSPs started cloud storage to store the archived data for longer durations. Long-term storage provides various features such as search, data lifecycle management, and data immutability. Examples of long-term storage include Amazon S3 and Google Cloud Archive Storage. Using lifecycle management rules, cloud customers can change the storage class of data that is less frequently accessed and reduce their storage cost. Long-term storage also offers high availability as the data is stored in multiple geographical locations and cloud customers can access it anytime.

Ephemeral Storage

This storage is associated with a VM or an instance. It exists as long as the instance is running. The size of ephemeral storage is based on the instance type and the image used. It cannot be used to store production data or other important data as it is terminated when the instance stops. Ephemeral storage can be used to store swap files and for temporary storage needs such as cache, buffers, etc. There are no separate charges for this storage and its cost is included in the cost of an instance.

hide01.ir

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



Content Delivery Network (CDN)

A CDN is a content distribution network in which objects are stored in multiple geographically distributed nodes that work together. The main goal of CDN is to improve performance and availability. It is different from the traditional method, where data is stored in a single central server. In CDN, a user accesses a copy of the data; this reduces traffic congestion near the server. Whereas in the traditional method, every user request goes to sever, and it affects the loading speed in case of multiple requests.

The distributed servers in the CDN called Points of Presence (PoPs), which store and deliver the cache of the website content. When a user makes a request to access the website, instead of sending the request to the original source, CDN sends the request to the nearest PoP, and delivers the content to the user. With this, it reduces the bandwidth and improves the page load time.

The size of the CDN network varies based on the architecture. The enterprise can have any number of nodes and servers based on the requirement and architecture of the project. To optimize the performance, the cloud consumer can choose the locations that are geographically nearest to the users. To optimize cost, cloud consumers can choose less-expensive locations.

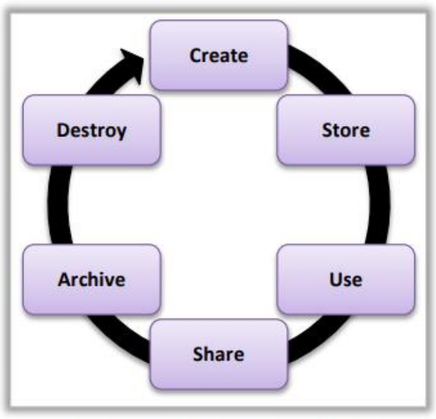
CDN has a feature of blocking access from non-HTTPS ports and thereby block any bad traffic from reaching the servers. It serves only good traffic to the servers. This ensures that the CDN data is safe. A CDN uses the cache eviction method to remove infrequently accessed resources and make space for new ones. Owing to these advantages, CDNs have gained in popularity and most web traffic today is served through CDNs.

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

Cloud Data Life Cycle Phases

- Based on the value and importance of that data, it should be secure throughout its lifecycle and **security controls** must be applied in all the phases
- There are six key phases of the cloud **data lifecycle** as follows:

1. **Create:** New digital content is **created**, existing content is modified or altered
2. **Store:** Data is committed to a **storage repository**
3. **Use:** Data is viewed, **processed**, or used in some other activity
4. **Share:** The data is **accessible** by others such as customers, partners, clients, or users
5. **Archive:** The data is not actively used and is stored for **long-term usage**
6. **Destroy:** When the data is no longer required, it must be permanently **destroyed** using physical or digital means



Cloud Data Life Cycle Phases

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

Cloud Data Life Cycle Phases

Data is an asset for any organization and most organizations store data in the cloud. Based on the value and importance of that data, it should be secure throughout its lifecycle and security control must be applied in all phases.

hide01.ir

Good data management is crucial for safeguarding the data of an organization. There are 6 phases of the cloud data lifecycle.

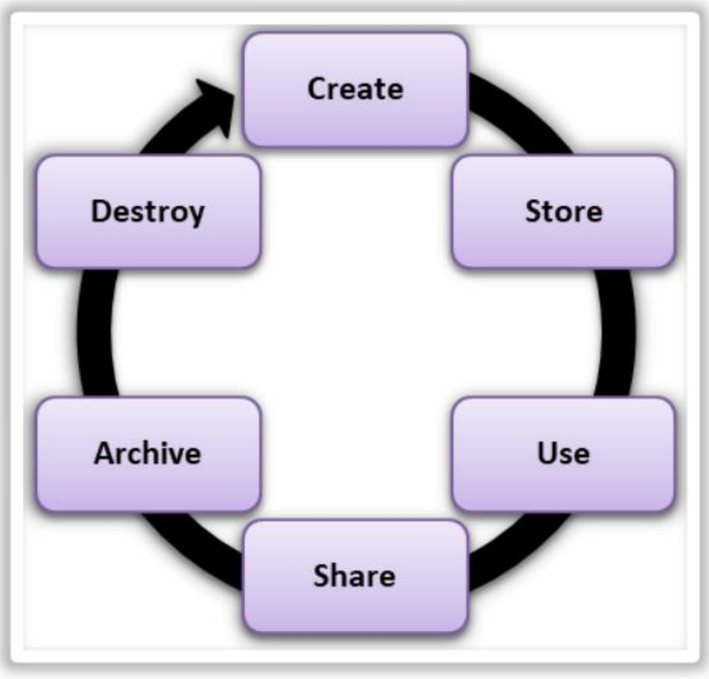


Figure 4.35: Phases of Cloud Data Lifecycle

- **Create:** This is the phase where new digital content is created, or existing content is modified or altered. Data classification is made in this phase based on the sensitivity of the information and its value to the enterprise. The classification should be made very carefully, as weak security controls may be applied if it is not classified properly.

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

- **Store:** In this phase, data is committed to a storage repository. Proper security must be provided to the data based on the sensitivity of the data (including access controls, encryption, monitoring, and backup). If the access controls are not implemented properly, or if the files are not scanned for threats before uploading, the data will be vulnerable to attacks.
- **Use:** Data is viewed, processed, and used in some other activity. It is then moved around the organization and it might be enhanced by the end-users. Data in this phase is more vulnerable to attacks as it might get transported to unsecured locations.
- **Share:** This is the phase where data moves from the organization. The data is accessible to others such as customers, partners, clients, or users. Maintaining data security can be difficult as the data is no longer under the control of the organization. Various technologies can be used to detect unauthorized data sharing and to gain control over data.
- **Archive:** At some point of time, the data is no longer needed for immediate use, and it can be archived if it is required in the future. Data at the archive stage should be protected according to the classification.
- **Destroy:** When the data is no longer required, it must be permanently destroyed using physical or digital means. While deciding the destruction method, the cloud customers should consider regulations, type of cloud service, and data classification. The cloud security professional must ensure that the data is destructed properly.

Though the data lifecycle includes all the phases through which data passes through, it does not address the data location, how it is accessed, functions that can be performed on data, and the authorization process. A secure cloud solution must handle all these aspects.

The data lifecycle should be managed as small lifecycles running on different environments. It is important to keep track of the physical and logical locations of the data for audit and other control requirements. It is also important to control who can access the data, and how they can access it. The cloud solution architect is responsible for implementing all the required controls on the data, irrespective of where the data is located and who is accessing it.

Cloud Data Lifecycle and Retention Policy

Data is an important asset of an organization. An organization must identify where the data is vulnerable. There are five main points where the data is vulnerable, that is, collection, storage, sharing, analysis, and deletion. The attacker targets these vulnerable points and gains access to the data. These vulnerable points are part of a cycle known as the data lifecycle. The data lifecycle identifies where exactly the data resides (on-premises, cloud, third-party vendors, etc.) and steps to be taken to secure the data such as determining user data access depending on their use as well as maintaining least-privileged-model and role-based access controls where the data owners regulate the access of the users to the data.

Whether the data resides in cloud, on-premises, or with third-party vendors, the organization should consider how long the data is going to reside in their system. The data should be retained as long as the purpose of collecting the data is achieved. Data-retention limits, which are

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

measured in years, refer to the duration of time an electronic document is kept. It reduces data security risk. If the user does not need the data, then the data should be deleted. If the size of data is small, then there would be less damage due to a data breach. If the data is sensitive, then ensure that it is accessible to only those who need it. As stale and old files are risky, implementation of data-retention policy helps in quarantine, archiving, and deleting of data.

Retention Policies and Bucket Lock in Cloud Storage

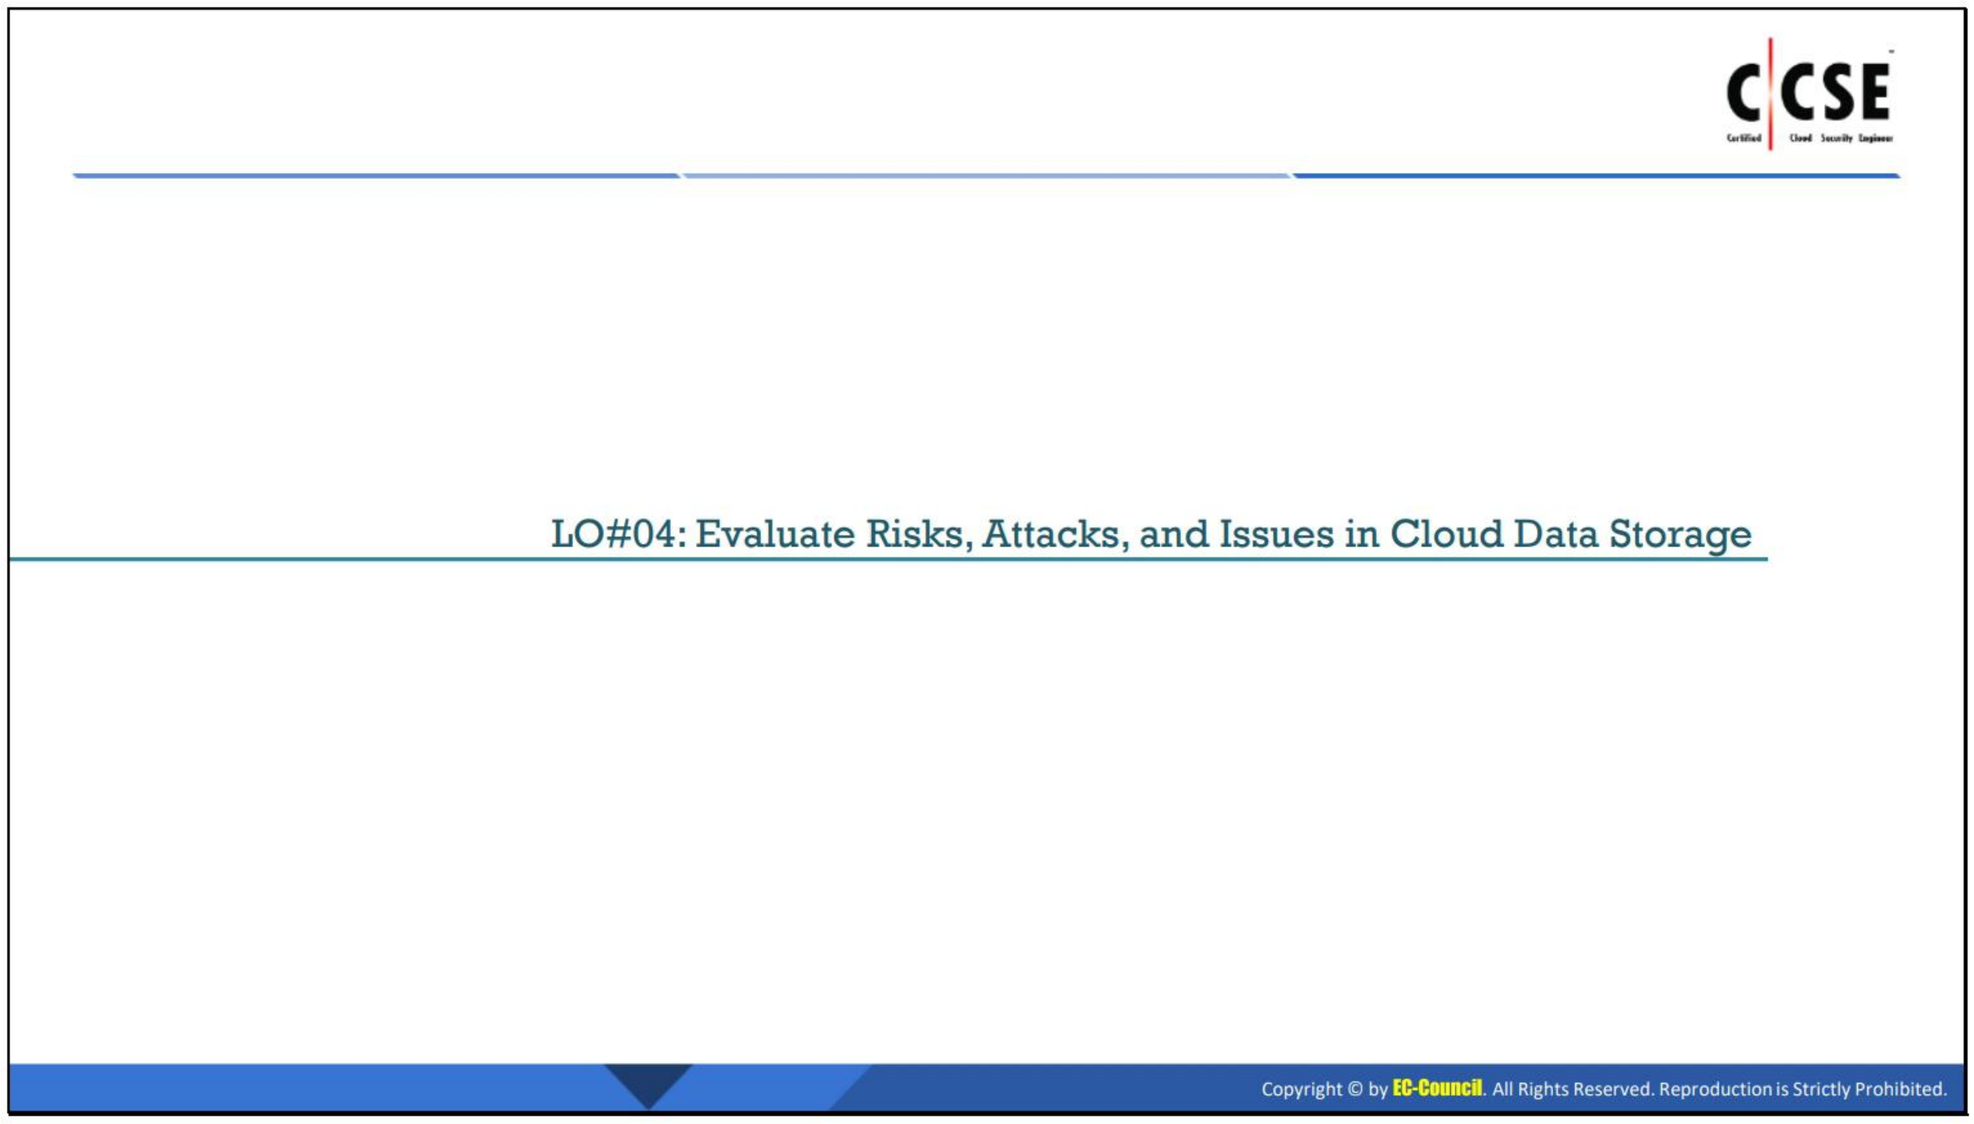
The cloud consumer can utilize the bucket lock feature for configuring the data-retention policy to the cloud storage bucket. The bucket lock feature allows the cloud consumer to decide how long the data should be retained in the bucket. This feature provides immutable storage and allows the cloud consumer to lock the data-retention policy. Bucket lock supports regulatory and compliance requirements.

Retention Policies

The cloud consumer can add a retention policy at the time of creating a new bucket or to an existing bucket. With the implementation of a retention policy on the bucket, all the objects in the bucket cannot be overwritten or deleted unless they reach the defined age in the retention policy. By utilizing the retention expiration time metadata of the retention policy, the cloud consumer can track the deletion time of an object.

hide01.ir

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

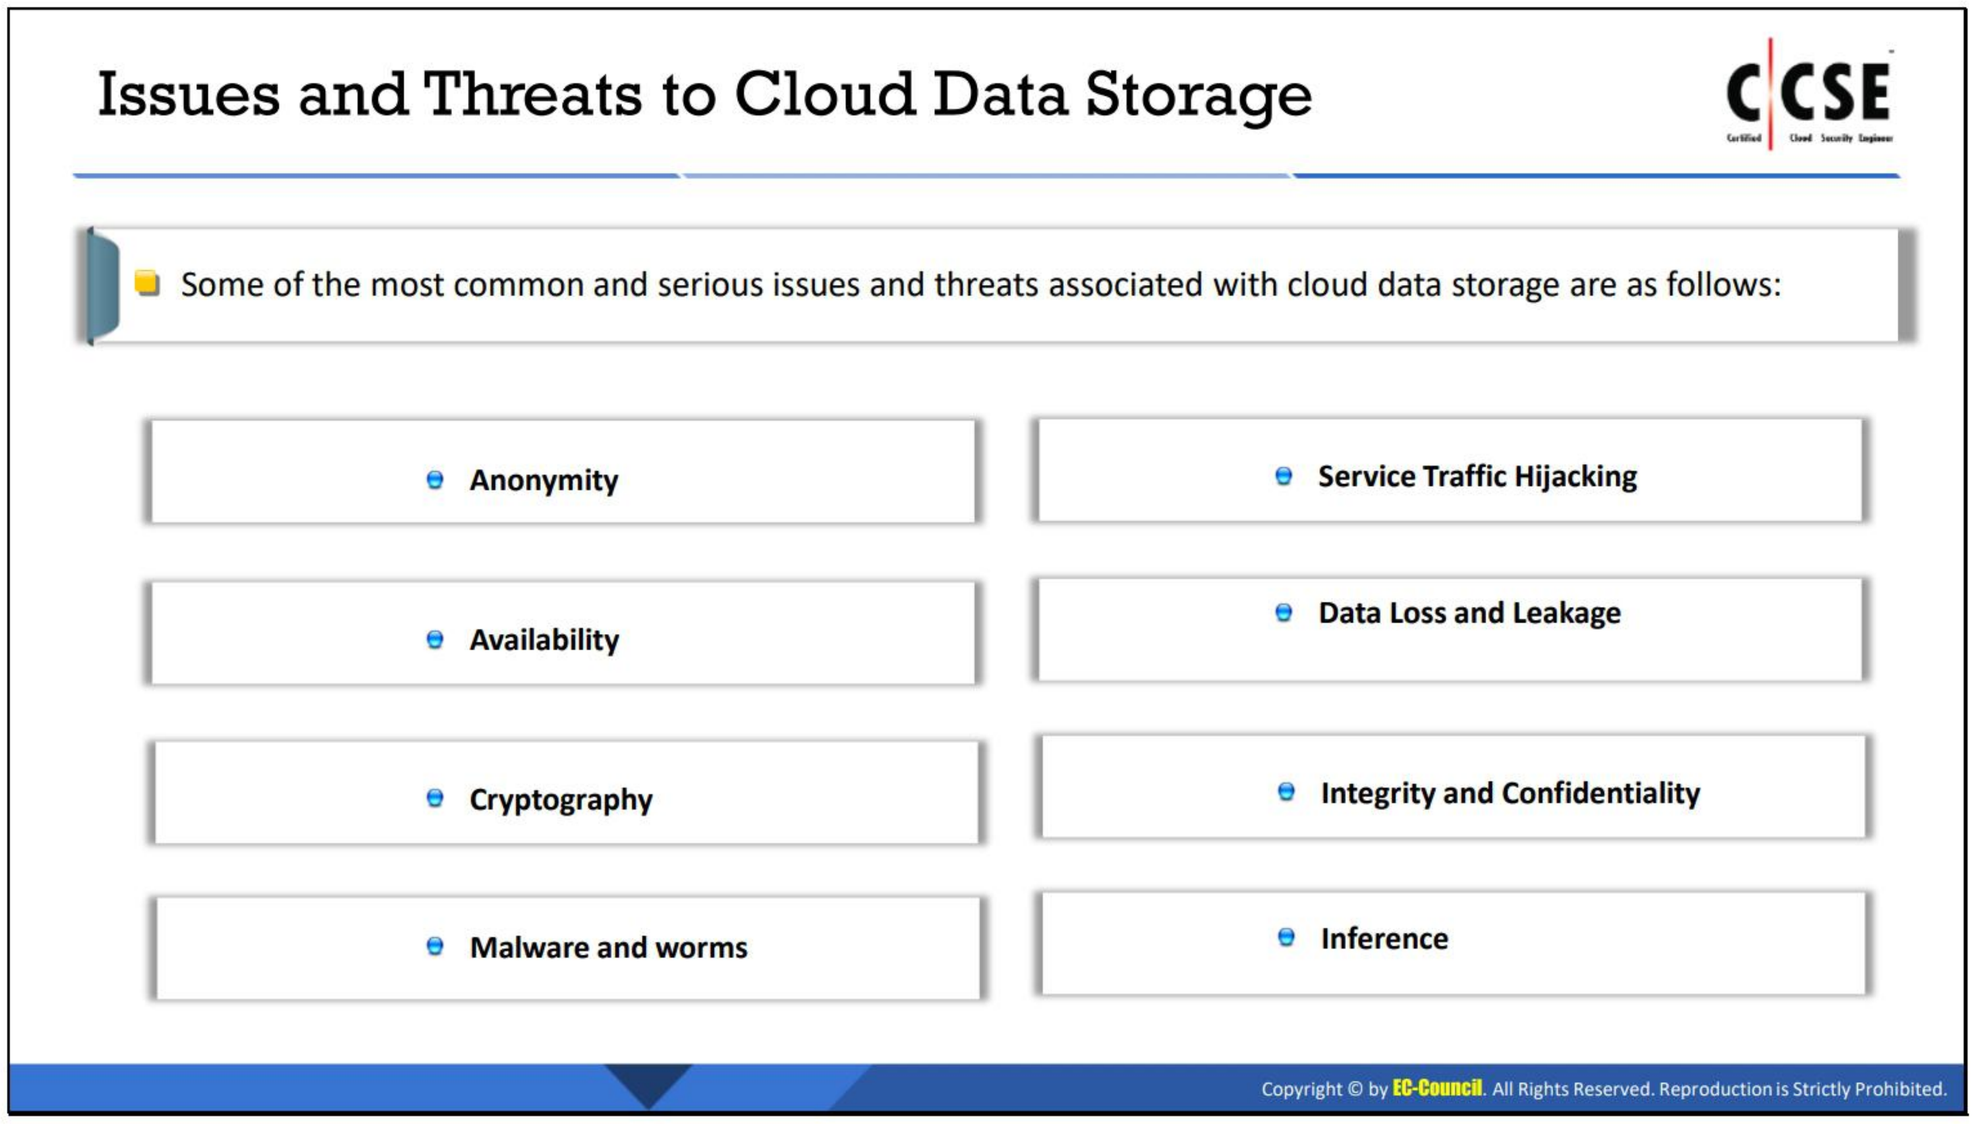


LO#04: Evaluate Risks, Attacks and Issues in Cloud Data Storage

Understanding the risks and threats associated with migration of sensitive data to cloud will help the organization better decide whether to migrate data to cloud. This section addresses the common issues and risks with cloud data storage and the best practices for cloud data storage.

hid001.ir

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



Issues and Threats to Cloud Data Storage

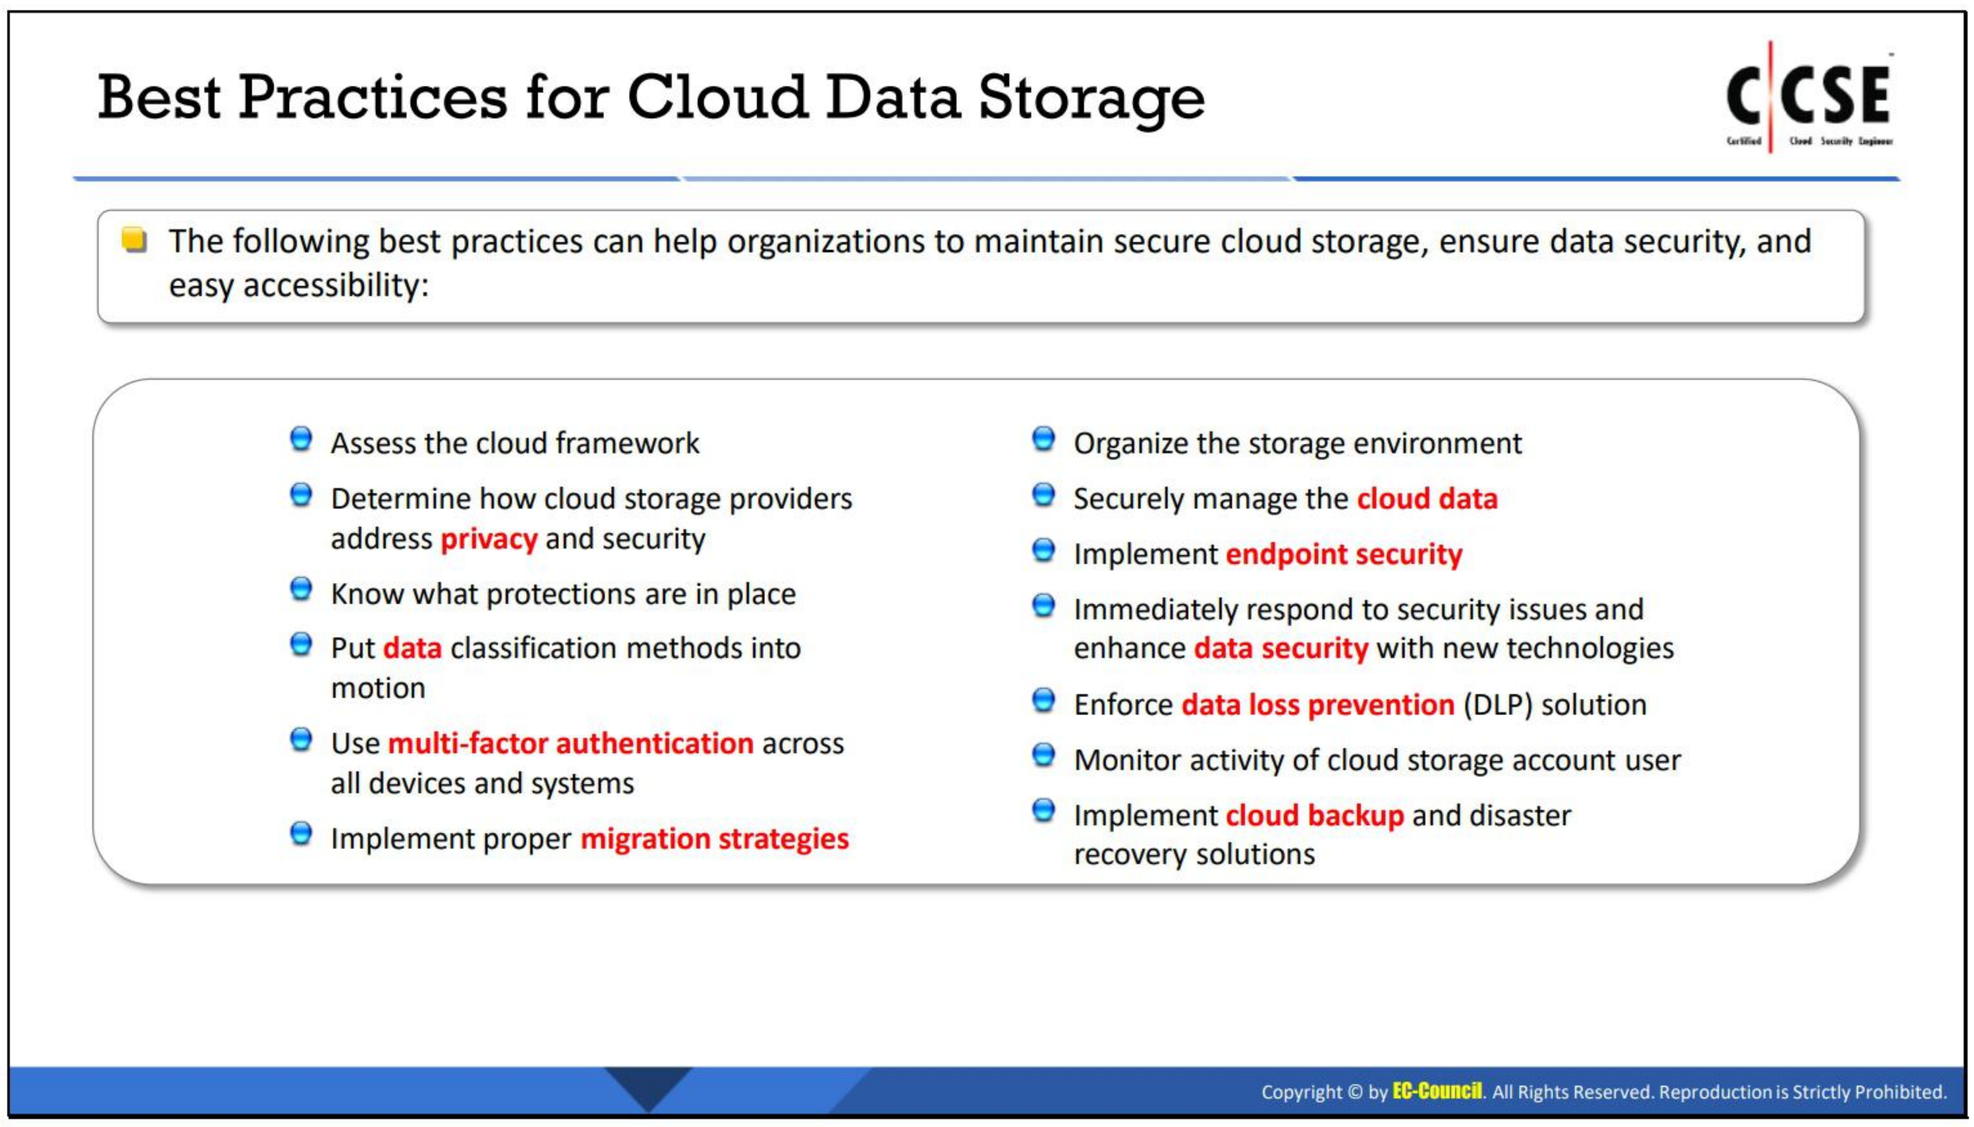
Data storage is one of the most important aspects of cloud storage. The data in the cloud is stored using third parties and is accessed through the internet. Various types of data storage solutions are provided by the cloud environment, and each data storage solution has its own advantage and limitation. The cloud security professional should be well-acquainted with the security threats to safeguard the cloud storage. Some of the most common and serious threats associated with cloud data storage are described below.

- **Anonymity:** Anonymity is a technique where an attacker obscures the key information and the attacker will be untrackable and unreachable. The vulnerabilities of data anonymity include flaws in the process of re-identification and the hidden identity of threats.
- **Service traffic hijacking:** It is one of the biggest threats and is performed by stealing credentials. Using these credentials, an attacker can get access to critical and confidential data such as deployed applications, virtual machines, and storage accounts, etc. This causes loss of confidentiality, integrity, and availability. Using vulnerability exploitations such as phishing, password theft, and buffer overflow, the attacker can gain access to user accounts and cause damage to existing accounts in various ways.
- **Availability:** Cloud services are intended to provide high availability to users, which means that the data can be accessed at any time and from any location. Availability is in terms of both hardware and software. Thus, cloud storage faces security challenges in terms of availability. Network attacks such as denial-of-service (DoS) and distributed denial-of-service (DDoS) are common in multi-tier architecture supported by a load balancer that runs on many servers.

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

- **Data loss and leakage:** When the data is stored in cloud storage, the cloud security professionals expect the data to be safe and secure. A data breach is caused by forceful action and it causes loss of data in the cloud without any backups. A data breach is the main area of leakage where the data can be used in any way by the attacker. This results in a loss of trust and affects the privacy of cloud consumers.
- **Cryptography:** Cryptography refers to the encryption of data stored in the cloud. It ensures that the data can be accessed only by the users who have cryptographic keys to access the server. Cryptography techniques are also applied to deal with loopholes in security areas. Poor key management where secret keys are revealed allows anyone to modify or extract the decrypted data in cloud storage. Additionally, prime factorization in RSA and faulty implementation can cause brute-force attacks.
- **Integrity and Confidentiality:** Integrity refers to protection against unauthorized modification, alteration, and deletion of data stored in the cloud. Misconfiguration of devices such as using incorrectly configured VMs and hypervisors can cause integrity attacks. Man-in-the-middle attack (MIM), session hijacking, packet sniffing, and social engineering are some of the known integrity attacks that affect the confidentiality and integrity of data in cloud storage.
- **Malware and worms:** Malware and worms are malicious applications or documents that can be injected into cloud storage. This will infect the network or application to cause damage or to shut down the systems; thereby, breaking the trust of cloud users.
- **Inference:** Inference is a technique used to attack databases where a malicious user breaches high-level information from complex databases using data-mining techniques.

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



Best Practices for Cloud Data Storage

Cloud services have been used for various purposes by enterprises such as storing data, accessing productivity tools, deploying IT infrastructure, etc. Cloud services are easy to use and can store a large amount of data efficiently. By applying security practices in advance, organizations can reduce the risks related to cybersecurity compliances. It is important to establish cloud storage security standards and cloud storage framework. The following best practices can help organizations maintain secure cloud storage, ensure data security, and easy accessibility.

- **Assess the cloud framework:** For securing the cloud storage, an organization should find all the devices and applications that connect to the cloud. It is also important to comprehend what type of cloud system exists in an organization, who utilizes it, and how they are using it. The organization can secure cloud storage by mapping the flow of data across applications, systems, devices, APIs, and clouds. Most cloud storage applications show the various applications and services that they are connected to, which not only simplifies the mapping task but also allows the organization to disconnect from the service or application.
- **Determine how cloud storage providers address privacy and security:** The general protections provided by the cloud service provider can be known through the terms of service agreements; however, this is not sufficient to provide secure file storage. As cloud service providers frequently update their terms of user and service agreement, there is a chance that a minor change may be overlooked that can have a tremendous impact on security and privacy. Moreover, many agreements would not mention all details such as how security is implemented by the cloud storage provider, what protection it utilizes, what happens at the time of breach, or breakdown. Therefore, the policy and procedures should be closely defined, which would need further discussion and negotiations.

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

- **Know what protections are in place:** Cloud security encryption is a basic requirement, and it is essential to know how a cloud storage provider utilizes encryption, control the encryption keys, and application of encryption keys to a particular data set, etc. Similarly, the companies utilizing the CSP should have complete information regarding access to the systems and the protections in place that can safeguard against DDoS attacks to application security flaws.
- **Put data classification methods into motion:** Due to the accumulation of a large quantity of unstructured data, data security has been becoming complex. Hence, understanding the value of the data plays a significant role, which includes storing the data in the cloud or archiving the data on media such as disk/tape. Data compliance for government regulations such as General Data Protection Regulation (GDPR) in European Union is another factor. The built-in tools provided by a few cloud storage service providers facilitate these processes.
- **Use multi-factor authentication across all devices and systems:** The utilization of MFA prevents unauthorized access and release of malware. MFA can safeguard sensitive information from attackers, disgruntled employees, and other insiders, who could place the data at risk.
- **Implement proper migration strategies:** Most enterprises migrate their data completely; this affects the business if the enterprise does not adapt to the new storage service. Apart from adapting to the new system, the team should also have a basic idea of the new system before they completely operate on it. Hence, before transforming the data to the cloud, the enterprise must plan a migration strategy to test the new storage. In the initial stages, the cloud consumer can move the data that is of less value to the company so that the business is not affected in case of data loss. By testing the storage environment, the team can understand how the data is stored on the cloud. Once the team gains enough knowledge of the storage environment, the enterprise can start migrating the important data to the cloud.
- **Organize the storage environment:** Any enterprise that stores data on the cloud should organize it in a logical and easily manageable layout. The authorized teams/users should be able to easily access the data. However, it may not be possible to replicate an on-premises storage structure; the enterprise should store the data in such a way that it matches the existing storage structure. This helps the organization's employees to understand the new cloud storage by comparing it with the cloud storage they had used.
- **Securely manage the cloud data:** Data security is the main concern for any enterprise that stores the data on the cloud. To achieve data security, cloud consumers must properly classify the data based on the sensitivity of the data. The most sensitive data requires the highest security. A data classification software can help identify the data that requires high security. After classifying the data, the security team can implement a security solution that should be able to identify the sensitive information in databases, endpoints, and cloud storage environments.

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

The enterprise should also focus on data storage and access procedures. The access control permissions associated with the data should also be reviewed and modified on regular basis. The company should also adopt strict data-sharing policies to reduce the risk of malicious insiders or hackers stealing the cloud data.

- **Implement endpoint security:** Though cloud storage provides security, it is important for an organization to enforce endpoint security. Endpoint protection refers to the protection of endpoint devices such as desktops, mobiles, tablets, and laptops. Endpoint devices act as access points to cloud resources and an attacker can exploit them at any time. For this reason, the enterprise should protect the devices that are used to connect to cloud accounts. By implementing endpoint security, the company can prevent activities that can provide access to cloud resources, which helps maintain stronger control.
- **Immediately respond to security issues and enhance data security with new technologies:** As the usage of cloud services increase, the threats and risks also increase with the usage. It is important to follow some best practices to respond to security issues. When a user accesses the sensitive data on the cloud from a new device, ensure that two-factor authentication is required to reduce the high risk associated with accessing sensitive data. It is impossible to predict the cloud service that will be used and handle security for that. For this, the cloud security professional can update the web access policies with information mentioning a cloud service to block the access or to show a warning message.

  When there is malware in a shared folder in the cloud, the associated cloud service automatically gets affected by the malware without any user actions. To avoid this, always scan the files with anti-malware in the cloud storage to prevent data breach attacks. Several technologies such as cloud access security broker (CASB), cloud workload protection, and virtual network security can also be used to enhance data security.
- **Enforce data-loss prevention (DLP) solution:** Enforcing DLP is an effective way to secure cloud data. DLP is responsible for identifying, monitoring, and protecting the data-at-rest and data-in-transit. Enforcing it enables the cloud security team to monitor the usage of data and the location from where it is being accessed. This strengthens the data security against insider attacks. DLP also handles external malware attacks on data and it does not allow the transfer of sensitive information to removable devices.
- **Monitor activity of cloud storage account user:** Though the organization follows a security strategy for using cloud accounts, employees may not adhere to them. Therefore, it is important to monitor their usage activities. With monitoring, the security team can have a clear picture of the services or resources used by the employees. If a user performs a suspicious cloud usage activity, the security team can deny access and protect the cloud data from a breach. By assessing the network firewall, logs in the event management system, and security information, the cloud security professional can determine the risk level. By determining the risk level, they can decide if that user needs full or partial access to the data.

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

It is also important to maintain an effective collaboration between the CSPs and cloud consumers to access and handle the security risk events.

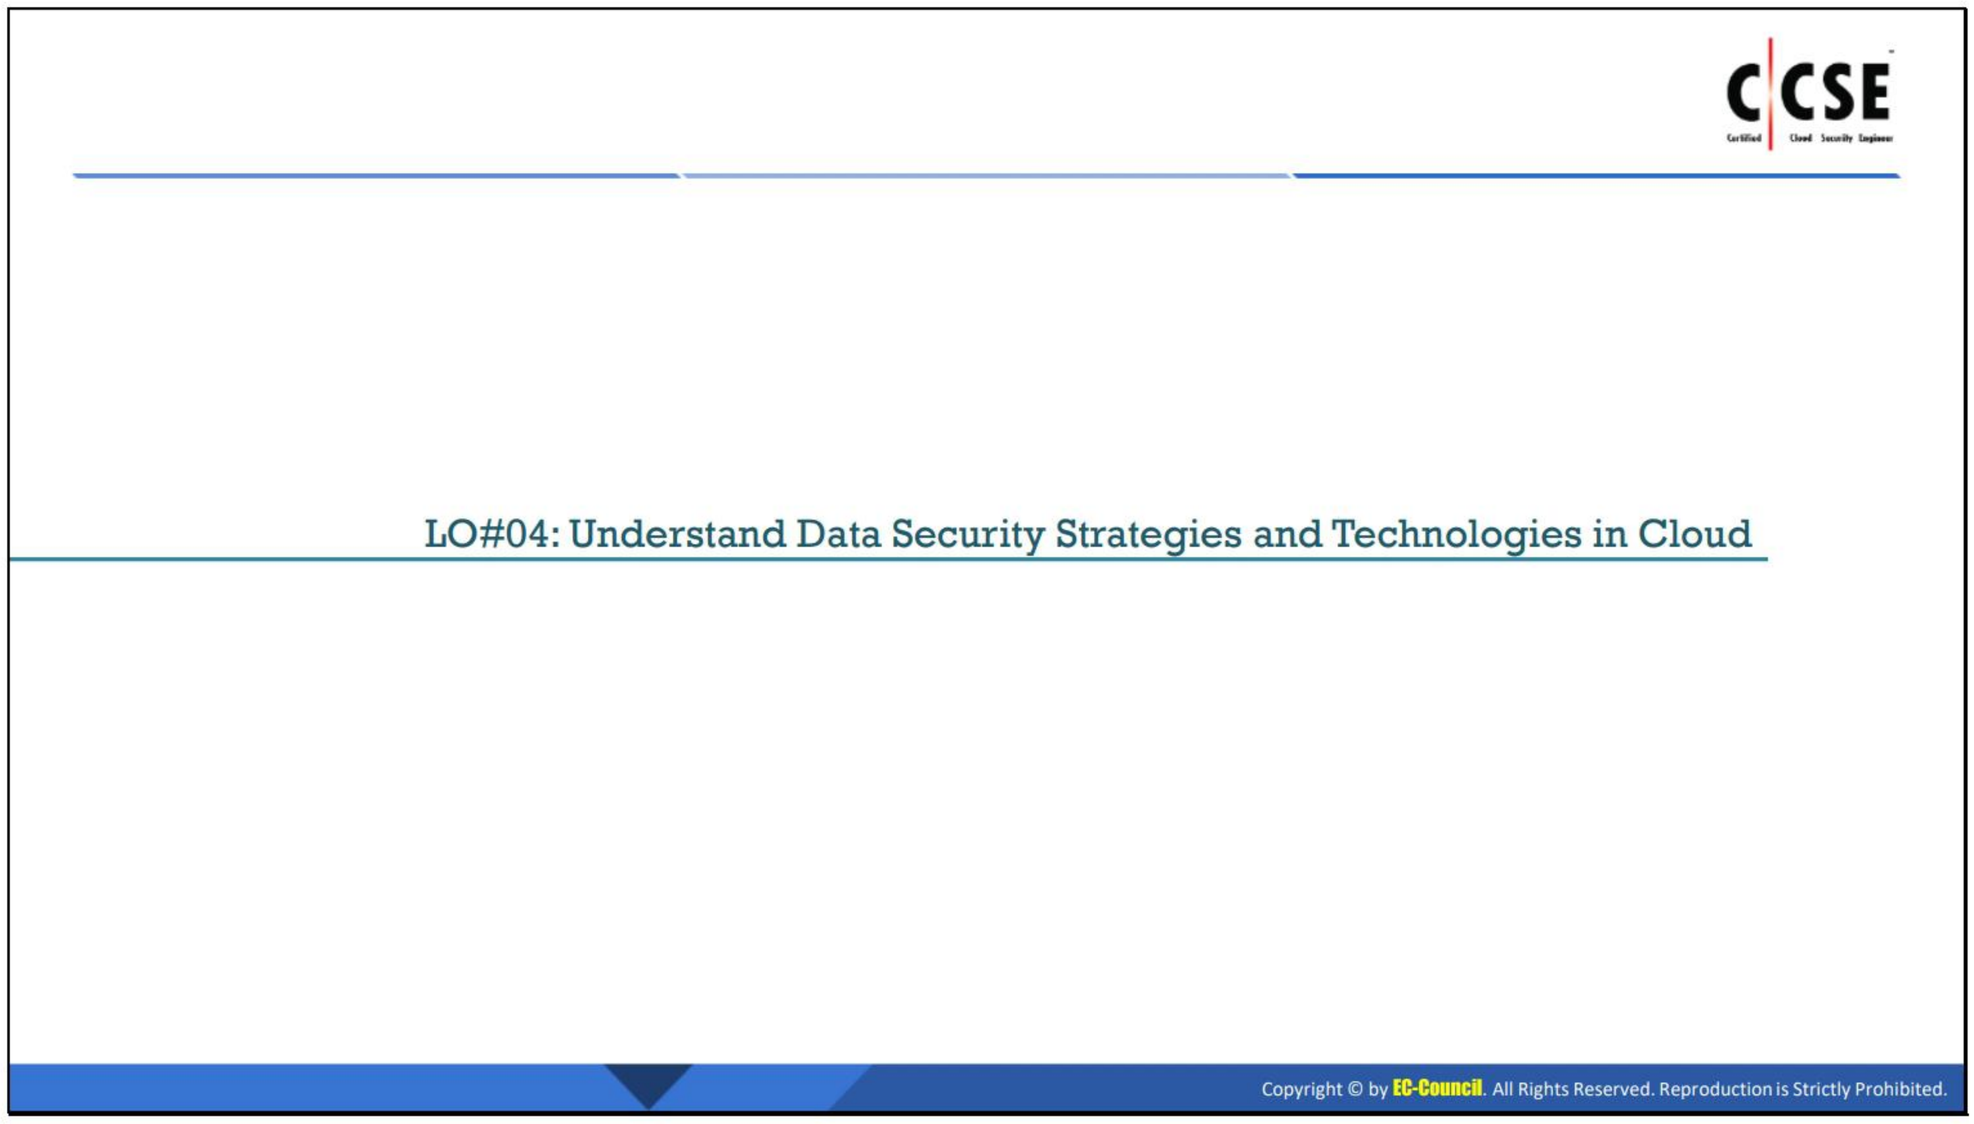
- **Implement cloud backup and disaster recovery solutions:** Storing the data in the cloud is risky as the data security will be handled by someone else other than the organization. Any damage to the server where the data has been stored can affect the entire operations of the organization. It is important to back up the data on a regular basis, which is like storing a second copy of the business-critical data in some other infrastructure. If the main server goes down due to any reason, use the backup server so that all the operations can be performed without any interruption.

  It is also important to have a disaster recovery solution to restore the application access and data recovery. The disaster recovery solution is responsible for redirecting the failover application requests to a virtual data center. With this, users can still use the application in case of a disaster. Once the server is restored, the disaster recovery solution redirects the user requests to the main server.

  Enterprises are using cloud storage services to provide better experience for their users and access the data from any location and at any time. Thus, robust data security policies must be implemented to protect sensitive data, prevent security attacks, and avoid data breaches.

hide01.ir

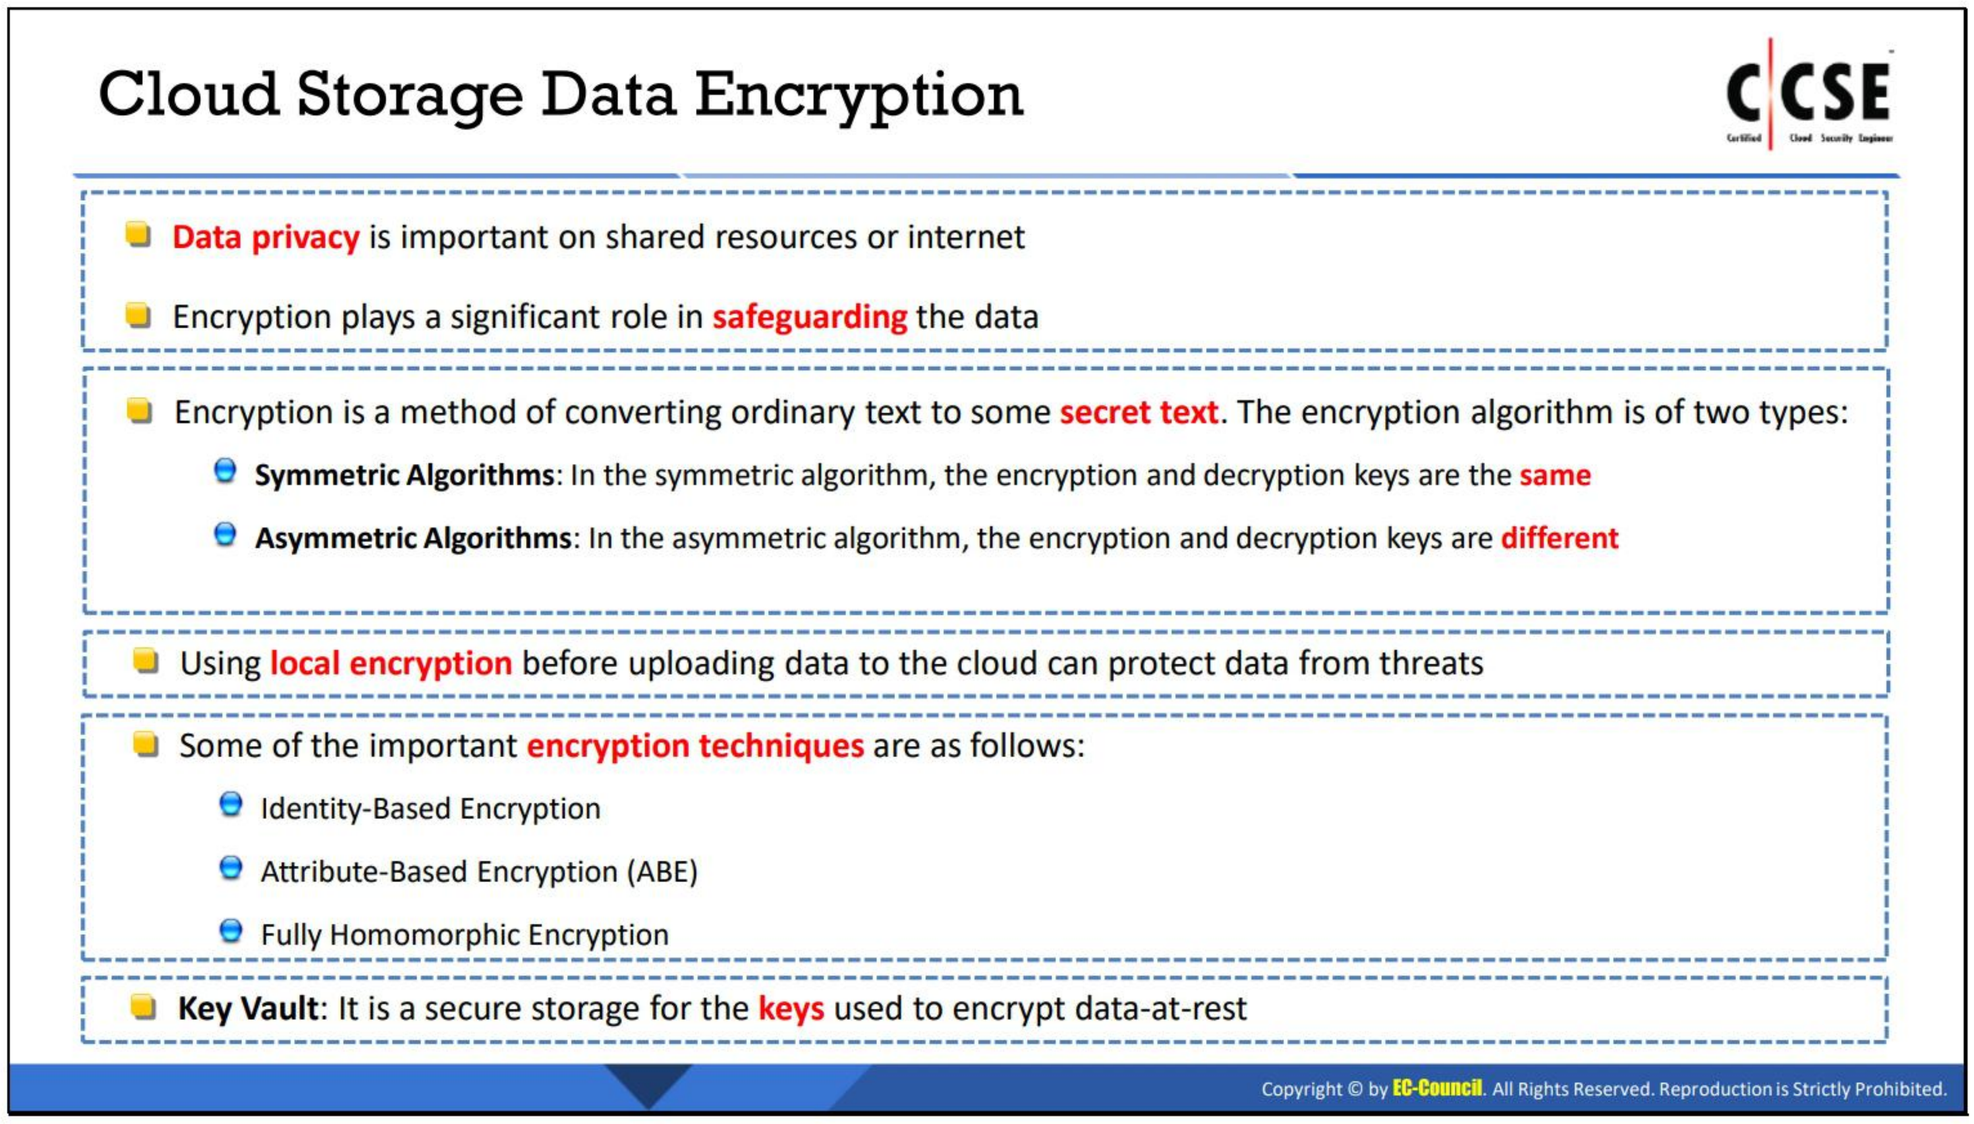
Certified Cloud Security Engineer Copyright © by **EC-Council**
All Rights Reserved. Reproduction is Strictly Prohibited.



LO#05: Understand Data Security Strategies and Technologies

To implement security controls in any cloud service model, it is important to understand the options that are available to secure the sensitive data in the cloud. The objective of this section is to explain different strategies and technologies for securing cloud data. This section includes cloud data encryption, encryption tools, key management, and how to protect the encryption keys. This section also explains cloud data-access controls, different methods for protecting the cloud data, and how to manage data migration in cloud.

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



Cloud Storage Data Encryption

Data privacy is important on shared resources over the internet; here, encryption plays a significant role in safeguarding the data. There are various security issues with cloud computing; thus, multiple cryptographic algorithms have been designed to protect data.

hid001.ir

Encryption is a method of converting ordinary text to some secret text. The encryption algorithm is of two types.

- **Symmetric algorithms**: In these algorithms, the encryption and decryption keys are the same. These algorithms do not consume much computation power. Examples include DES, 3DES, BLOWFISH, etc.
- **Asymmetric algorithms**: In these algorithms, the encryption and decryption keys are different. Examples include DSA, RSA, ELGAMAL, etc.

Using local encryption before uploading data to the cloud can protect data from threats. It is better to opt for a service provider who can provide prerequisite data encryption. Have a preliminary round of encryption for data that can provide consumers extra security if they already have an encrypted cloud service.

Various encryption techniques are implemented to ensure privacy and security in the cloud environment. Some of the important encryption techniques are as follows:

- **Identity-based encryption:** This encryption technique helps in the management of certificates and public key for public cloud infrastructure (PKI). It utilizes public key cryptography where the server of CSP uses simple identifiers such as email addresses for creating a public key that is helpful for encryption and decryption of the electronic message.

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

- **Attribute-based encryption (ABE):** There are various forms for attribute-based encryption. Ciphertext attribute-based and key attribute-based encryption are the common forms of ABE. In ciphertext ABE, the ciphertext attaches to the access policy and is decrypted with the attribute satisfaction. In key policy ABE, the ciphertext attaches to the set attribute and provides the data owner the key pair and policy. The attribute in the ciphertext should comply with the policy of key access for decrypting the entire message.
- **Fully homomorphic encryption:** This encryption technique secures the data in communication and storage. The major disadvantage of the traditional encryption technique is that it requires coding the data before manipulation. In fully homomorphic encryption, the computation is performed on encrypted data and then the data is sent. To perform this, a specific scheme has been prepared in which calculations are performed securely in the cloud and the server does not have any knowledge about the content of the data that is sent or the performed function.

Key Vault

Key Vault serves as a secure storage for the keys used to encrypt data-at-rest. Azure Disk encryption can be used with Azure Key Vault, which will control and manage the encrypted keys. Encryption of virtual disks requires the creation of an Azure Key Vault for storing cryptographic keys, which will encrypt or decrypt the virtual disks.

Azure Key Vault addresses the following concerns:

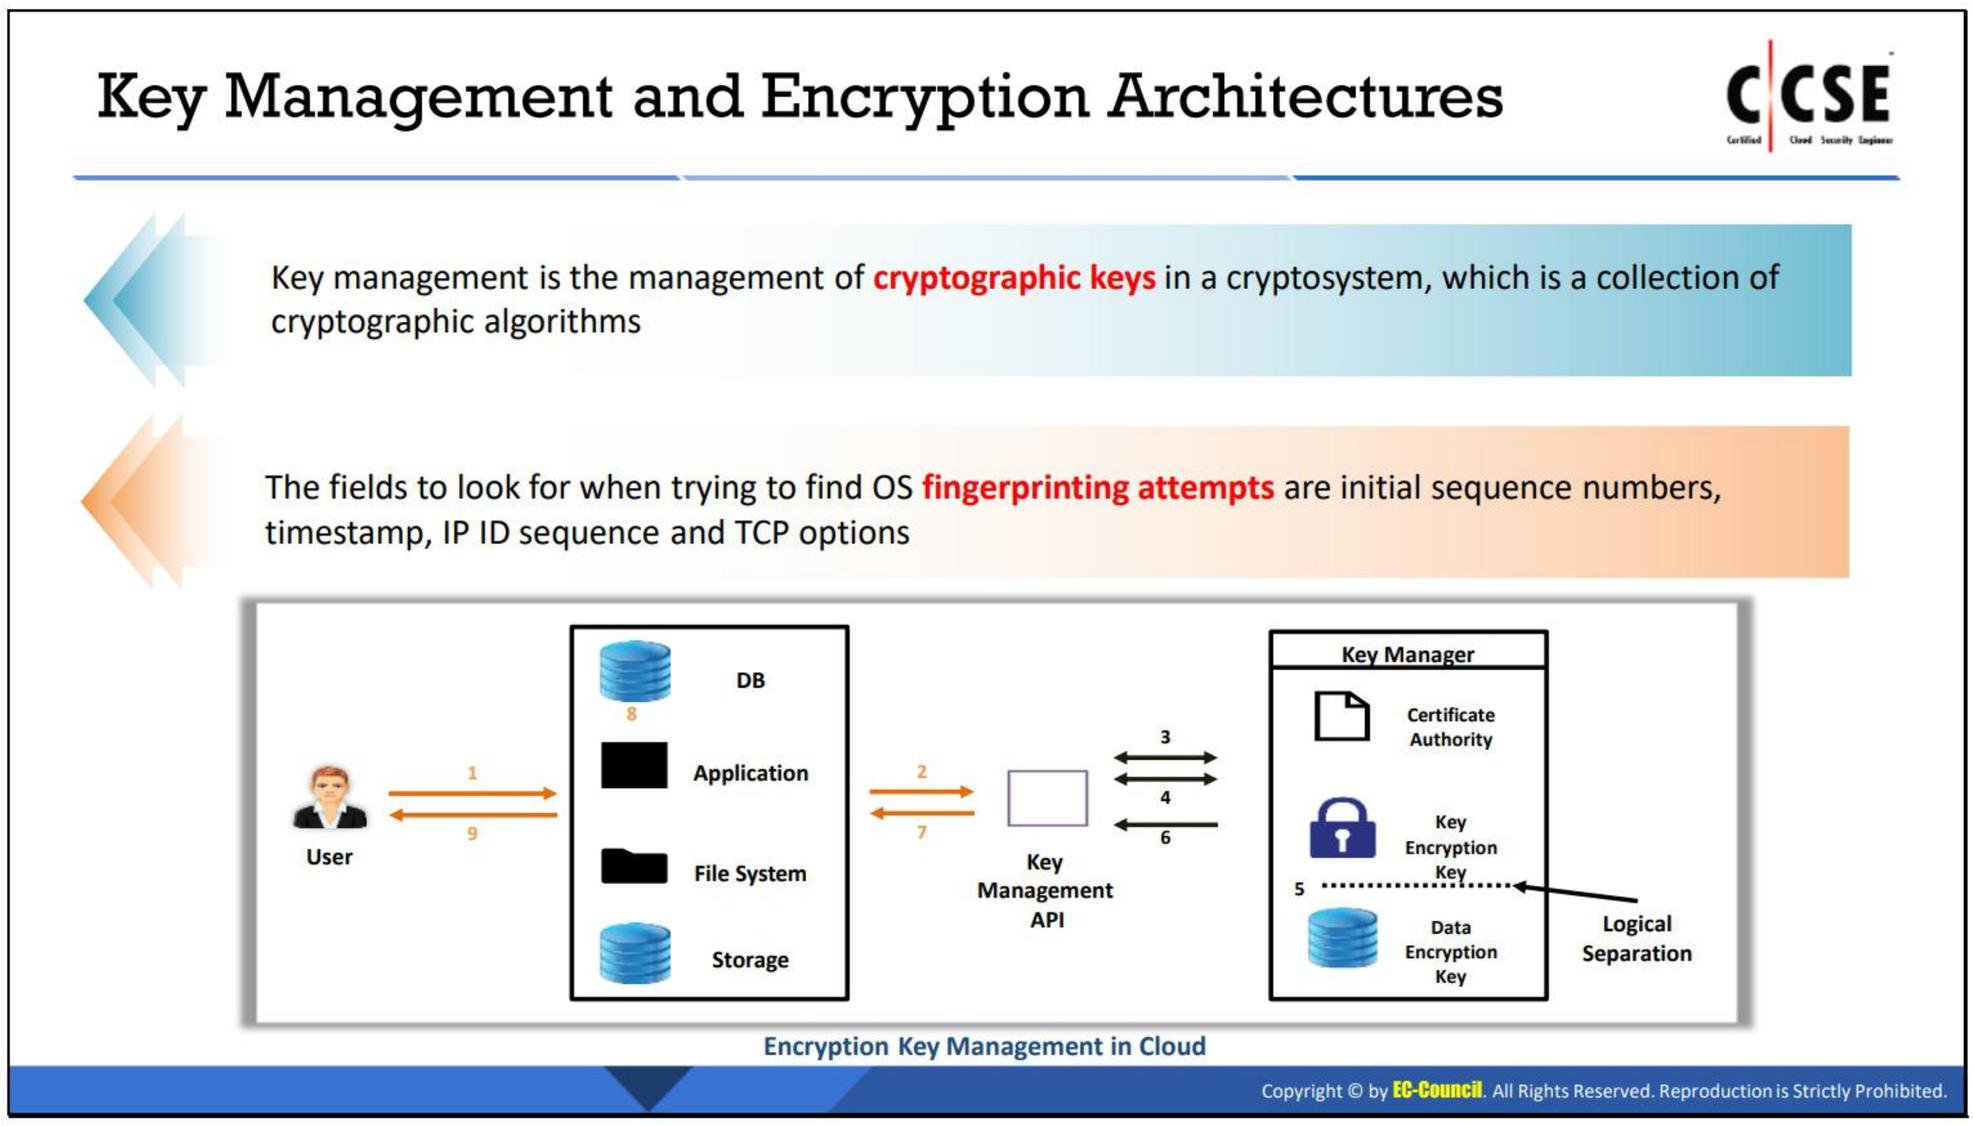
hide01.ir

- **Secrets management**: Azure Key Vault is used to securely store and tightly control access to tokens, passwords, API keys, etc.
- **Key management**: It is used to create and control the encryption keys.
- **Certificate management**: It is used to manage secure socket layers (SSL)/transport layer security certificates to use with Azure and internal connected resources.
- **Store secrets backed by HSM**: The secrets and keys will be protected either by software or FIPS 140-2 Level 2 validated hardware security modules.

Advantage of Azure Key Vault are as follows:

- Centralize application secrets
- Securely store secrets and keys
- Monitor access and use
- Simplified administration of application secrets
- Integration with other Azure services

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



Key Management and Encryption Architectures

Key management generally refers to the management of cryptographic keys in a cryptosystem, which is a collection of cryptographic algorithms. This management includes the generation, storage, exchange, use, replacement, and destruction of keys. Destruction of keys in cryptography is referred to as crypto-shredding. It also includes the design of the cryptographic protocol, user procedures, and key servers. It is critical for a cryptosystem to implement successful key management as it involves many tasks such as training for users, the interaction of multiple departmental and organizations, and mathematical practices for automation.

Types of Keys

Usually, two types of keys are used in cryptography for encryption: symmetric keys and asymmetric keys. Mostly, symmetric keys are used for data-at-rest where the same key is used for encryption and decryption of data. The data that is converted to unreadable format using encryption is known as ciphertext.

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

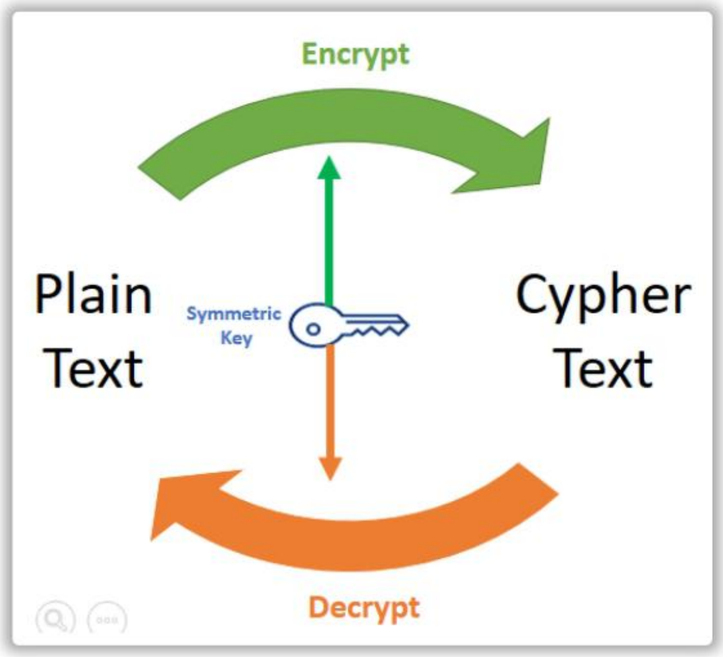


Figure 4.36: Data Encryption and Decryption

Asymmetric keys are used to encrypt data-in-motion; they require two keys: public key and private key. The public key is used for encryption and this key can be shared with anyone who needs to send data to you in an encrypted format. The private key is used to decrypt data that was encrypted by the public key. The private key must be safeguarded securely, as this is the only key that can decrypt data. If it is not stored properly, or an access policy is not enforced, users or people who have access can decrypt data.

Symmetric Key System

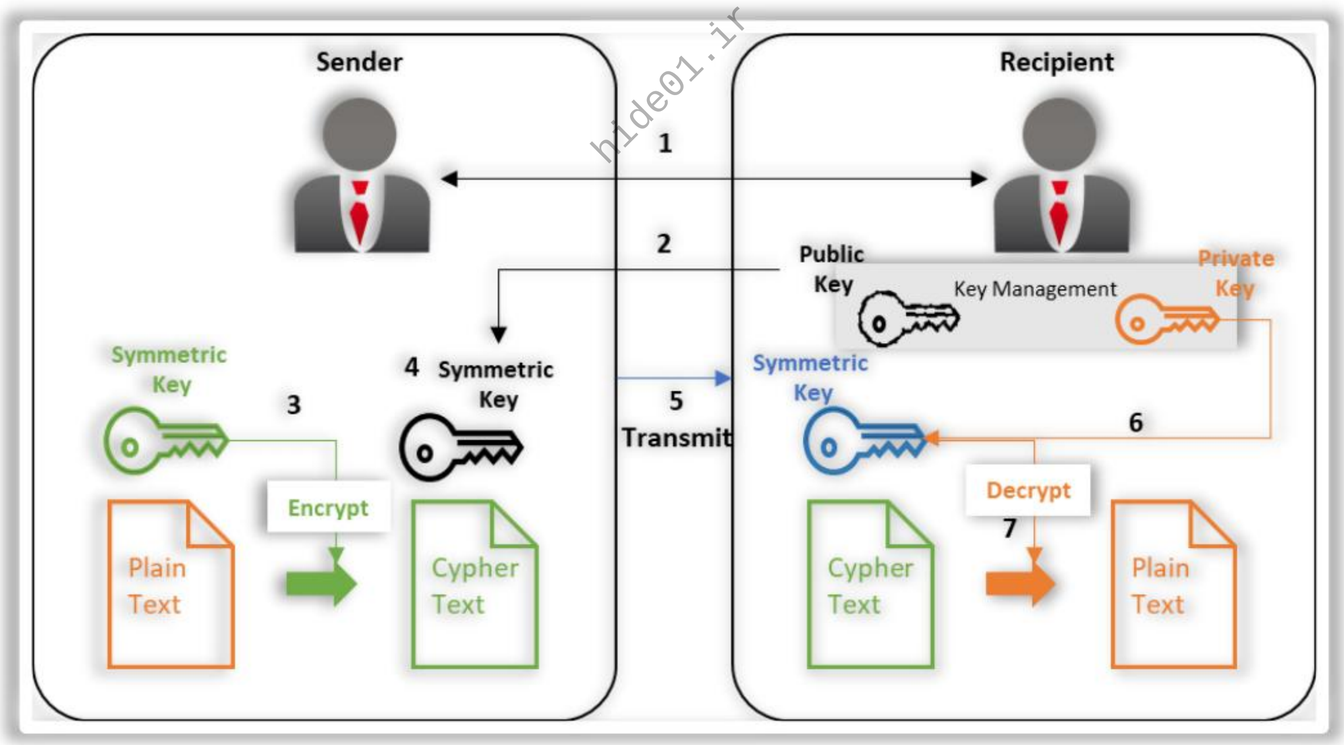


Figure 4.37: Symmetric Key System

There are certain important terminologies to understand regarding the key management system.

- **The data encryption key (DEK):** It is an encryption key that is used to encrypt and decrypt data.
- **The key-encryption key (KEK):** It is an encryption key that is used to encrypt a key, that is, DEK.
- **Key management application program interface (KM API):** It is an application interface used to transfer keys from the key management server to the requesting client securely.

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

- **Certificate authority (CA):** CA creates both public and private keys; apart from that it creates and verifies certificates.
- **Transport layer security (TLS):** The data-in-motion in a network is processed securely through mutual authentication using TLS protocol, which is a cryptographic protocol.
- **Key management system (KMS):** KMS is the environment where the entire key management application is maintained.

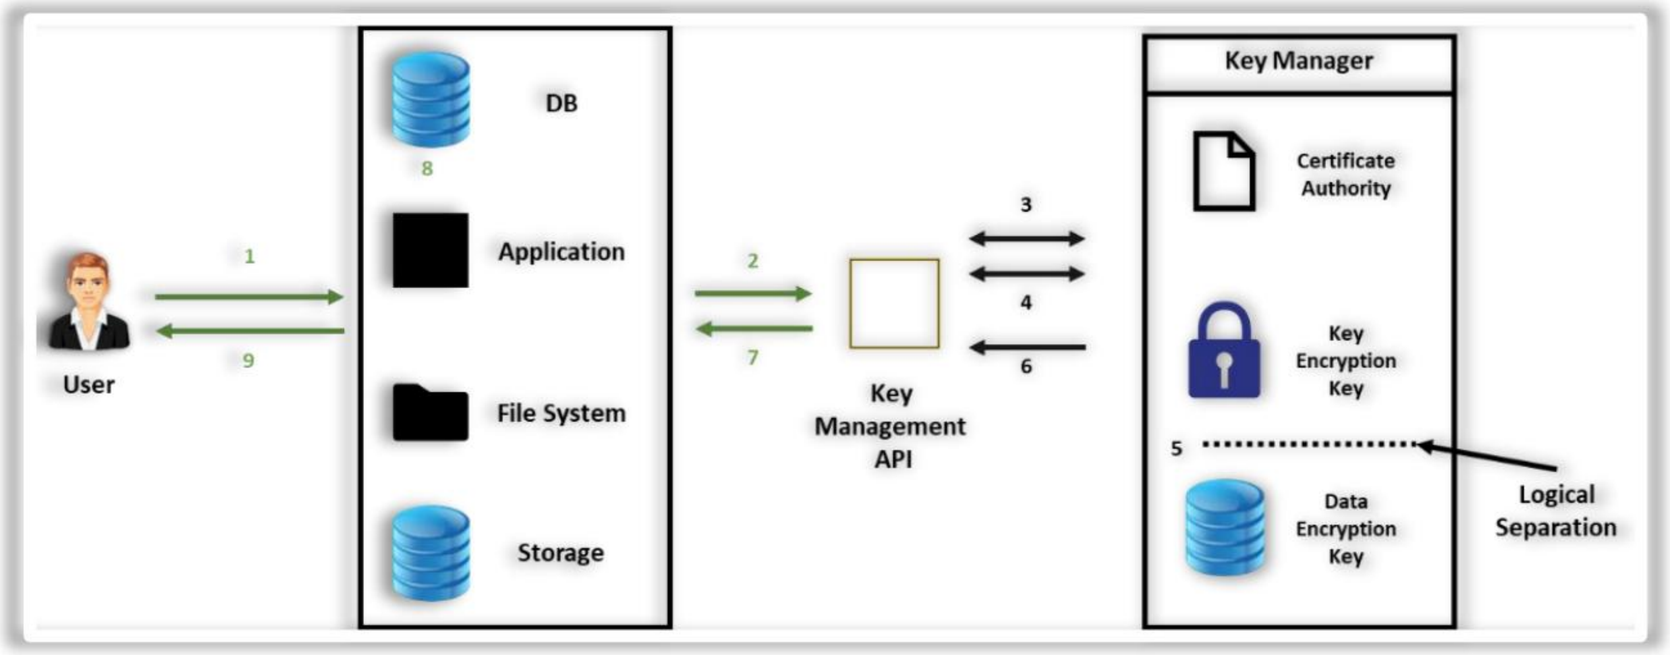


Figure 4.38: Key Management System

- Initially, a user requests access to encrypted data.
- The resource sends a DEK retrieval request to the client (KM API).
- Client KM API and key manager verifies certificates.
- TLS connection for KM API and KM is established once the verification is done.
- Now the key manager starts decrypting the requested DEK with the KEK.
- The key manager sends the DEK to the client (KM API) through an encrypted session using TLS.
- The KM API then sends the DEK to the resource that is required.
- The DEK may be cached in the database's temporary memory.
- The information is not sent to the user from the requested resource in plain text format.

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

Asymmetric Key System

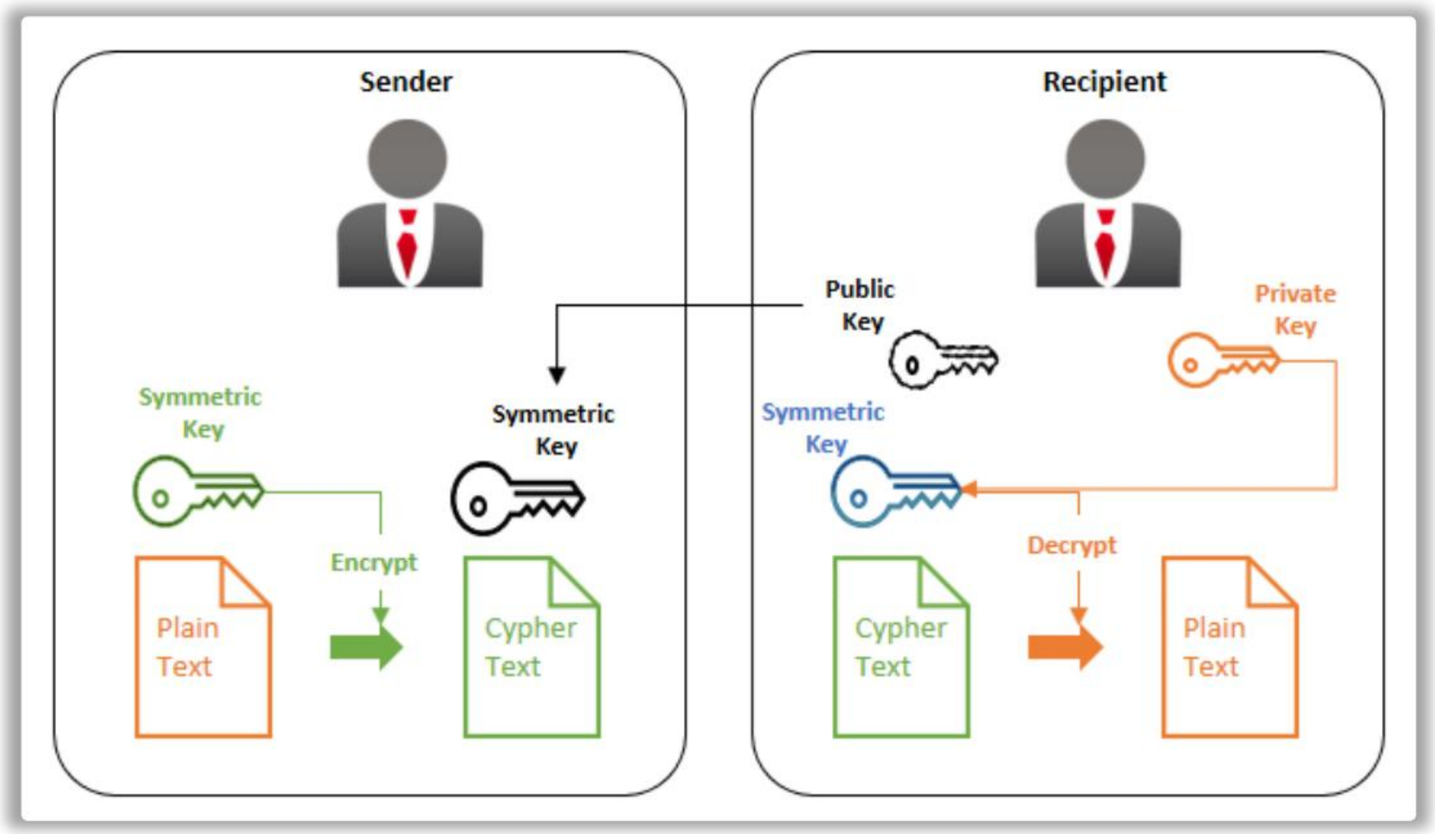


Figure 4.39: Asymmetric Key Flowchart

- Initially, the certificates of both sender and receiver should be verified, where the sender requests for verification by sending a certificate to the recipient; then, the recipient verifies the certificate against their certificate authority for authentication. As soon as the sender certificate is verified, the recipient sends their certificate to the sender for further acceptance and authentication.
- After the completion of mutual authentication, the sender will request the recipient's public key and the recipient will share the public key with the sender.
- The sender will create a transient symmetric key and starts encryption of the data or file that must be shared.
- The symmetric key is encrypted using the public key by the sender.
- Encrypted data and the encrypted symmetric key are sent by the sender.
- The recipient gets the packet and decrypts the symmetric key with the private key.
- Finally, the recipient can decrypt the data with the symmetric key.

hide01.ir

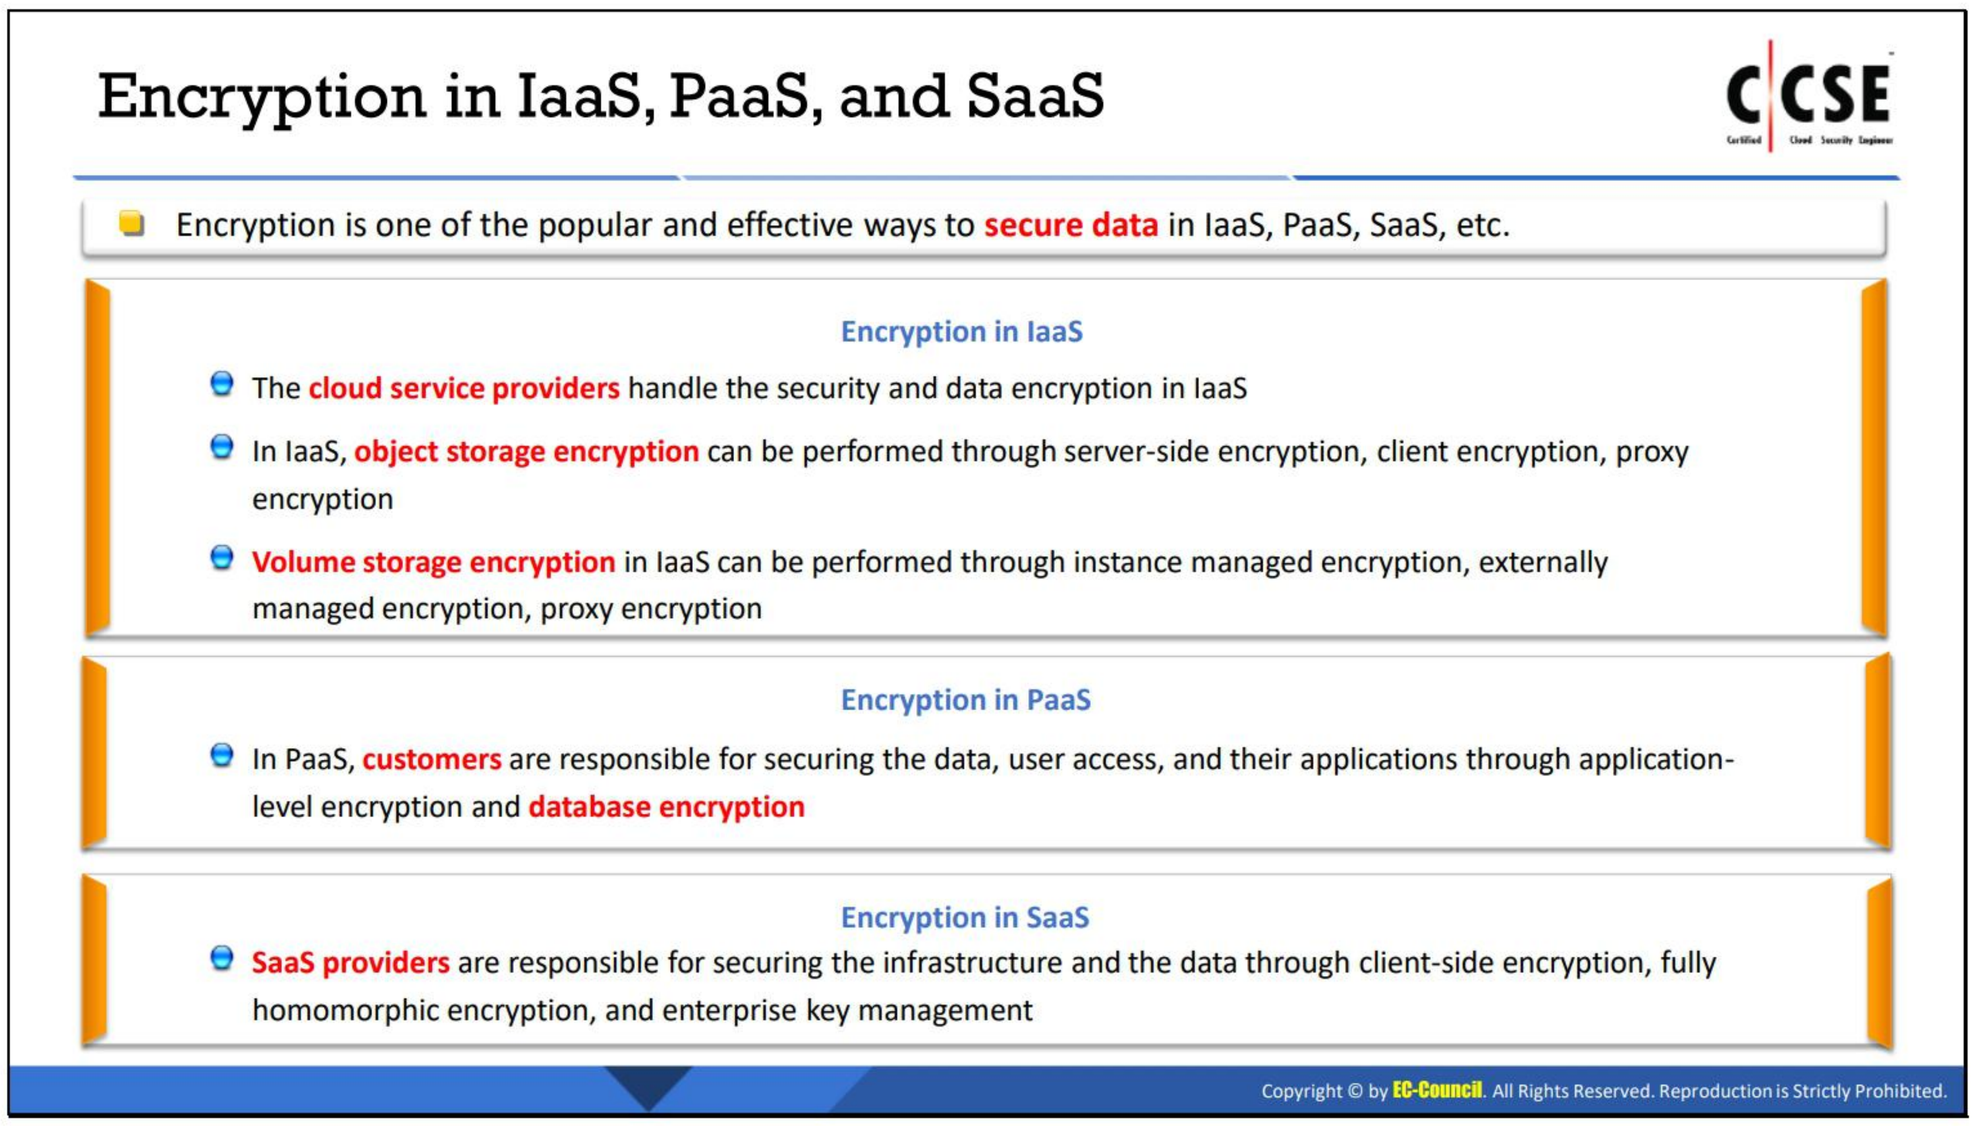
Key Revocation

A key can be revoked using the key manager so that it cannot be used for encryption or decryption. This can be performed by an administrator. The revoked key can be reactivated again in order to decrypt old encrypted data.

Key Deletion

If a key is not needed anymore or if it is being compromised, the key should be completely deleted from the key manager. All its instances or particular instances will be deleted from the key manager. There should be no possible way to recover the key, except in certain cases; for example, any encrypted sensitive data cannot be recovered if the key is not available.

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



Encryption in IaaS, PaaS, and SaaS

As mentioned earlier, encryption is the process of converting data into some other form or code so that people who have a secret key or password can only access the encrypted data. Encryption is one of the popular and effective ways to secure data in IaaS, PaaS, and SaaS.

hidori.ir

The main purpose of encryption is to protect the confidentiality of the data as it is stored in the cloud and is transmitted using the internet, passing through various networks.

Encryption in IaaS

IaaS provides VMs, virtual storage, and virtual networks. Most organizations use IaaS to create a cloud-based IT environment. Security is also a major concern in IaaS. In IaaS, CSPs handle the security and data encryption while customers are also responsible for securing the data, virtual network traffic, and operating system.

Object Storage Encryption

When it comes to object storage, encryption can be performed in 3 ways, namely:

- Server-side encryption
- Client encryption
- Proxy encryption

With server-side encryption, the data is encrypted directly by the cloud platform. All the encryption components reside in the cloud platform and are managed by the cloud administrators. This encryption only protects the data from data loss and the cloud administrator has access to the keys. Therefore, it does not protect the administrators.

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

To deal with this, the customers can encrypt the data before sending it to the cloud. This can be achieved using "virtual private storage." This allows protecting the private data in the public cloud by encrypting it and retaining keys. It can be performed using an encryption agent on the host that connects to the cloud service. The customers only need to install a local encryption agent to encrypt the data before sending it to the cloud and storing the keys. This ensures security to the data as the customers only have the secret keys.

Another method of securing data is using inline or cloud-hosted proxy. There are two options with proxy. In the first option, the proxy is installed on the customer's network and the data runs on that proxy for encryption and decryption. The other option is for the proxy to run on a virtual appliance in the public or private cloud. In the proxy encryption method, all the connections are routed to the proxy, which encrypts the data and moves it to the cloud object storage.

Volume Storage Encryption

Volume storage refers to the volume used while running a VM or an instance in the cloud. Volume storage encryption is performed by encrypting the entire volume attached to the instance. It can be performed in 3 ways, namely:

- **Instance managed encryption:** In this method, the encryption engine runs in the instance and the keys are stored in the volume protected by a keypair.
- **Externally managed encryption:** In this method, the encryption engine runs in the instance and the keys are stored externally. Keys are assigned to the instance if requested.
- **Proxy encryption:** In this method, the volume is connected to the software or a special instance to perform encryption. The proxy performs all encryption and decryption operations and stores keys either externally or onboard.

hidev.ir

Encryption in PaaS

PaaS is a scalable development and deployment environment. In PaaS, customers are responsible for securing the data, user access, and their applications. Whereas the PaaS service provider is responsible for securing the physical infrastructure and operating system.

Encryption in PaaS is different from other services. Encryption in PaaS can be performed at the application layer and database level.

In application layer encryption, data is encrypted in the PaaS application whenever the data passes through the application. The application layer encryption encrypts data at multiple layers such as files, disk, and database. It improves security by reducing the risk of security attacks. As the data is encrypted in multiple layers, the organization can protect the sensitive data before storing it in a database or any cloud platform.

The other option is database encryption. The database is encrypted using built-in encryption that is supported by the database platform.

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

Encryption in SaaS

SaaS applications provide various benefits to organizations such as increased productivity, reduced overheads, ease of use, etc. If the data is not protected in SaaS, it may lead to data loss, which can have significant consequences for the business. SaaS providers are responsible for securing the infrastructure and the data.

Security in SaaS can be provided in 3 ways, namely:

- Client-side encryption
- Fully homomorphic encryption
- Enterprise key management

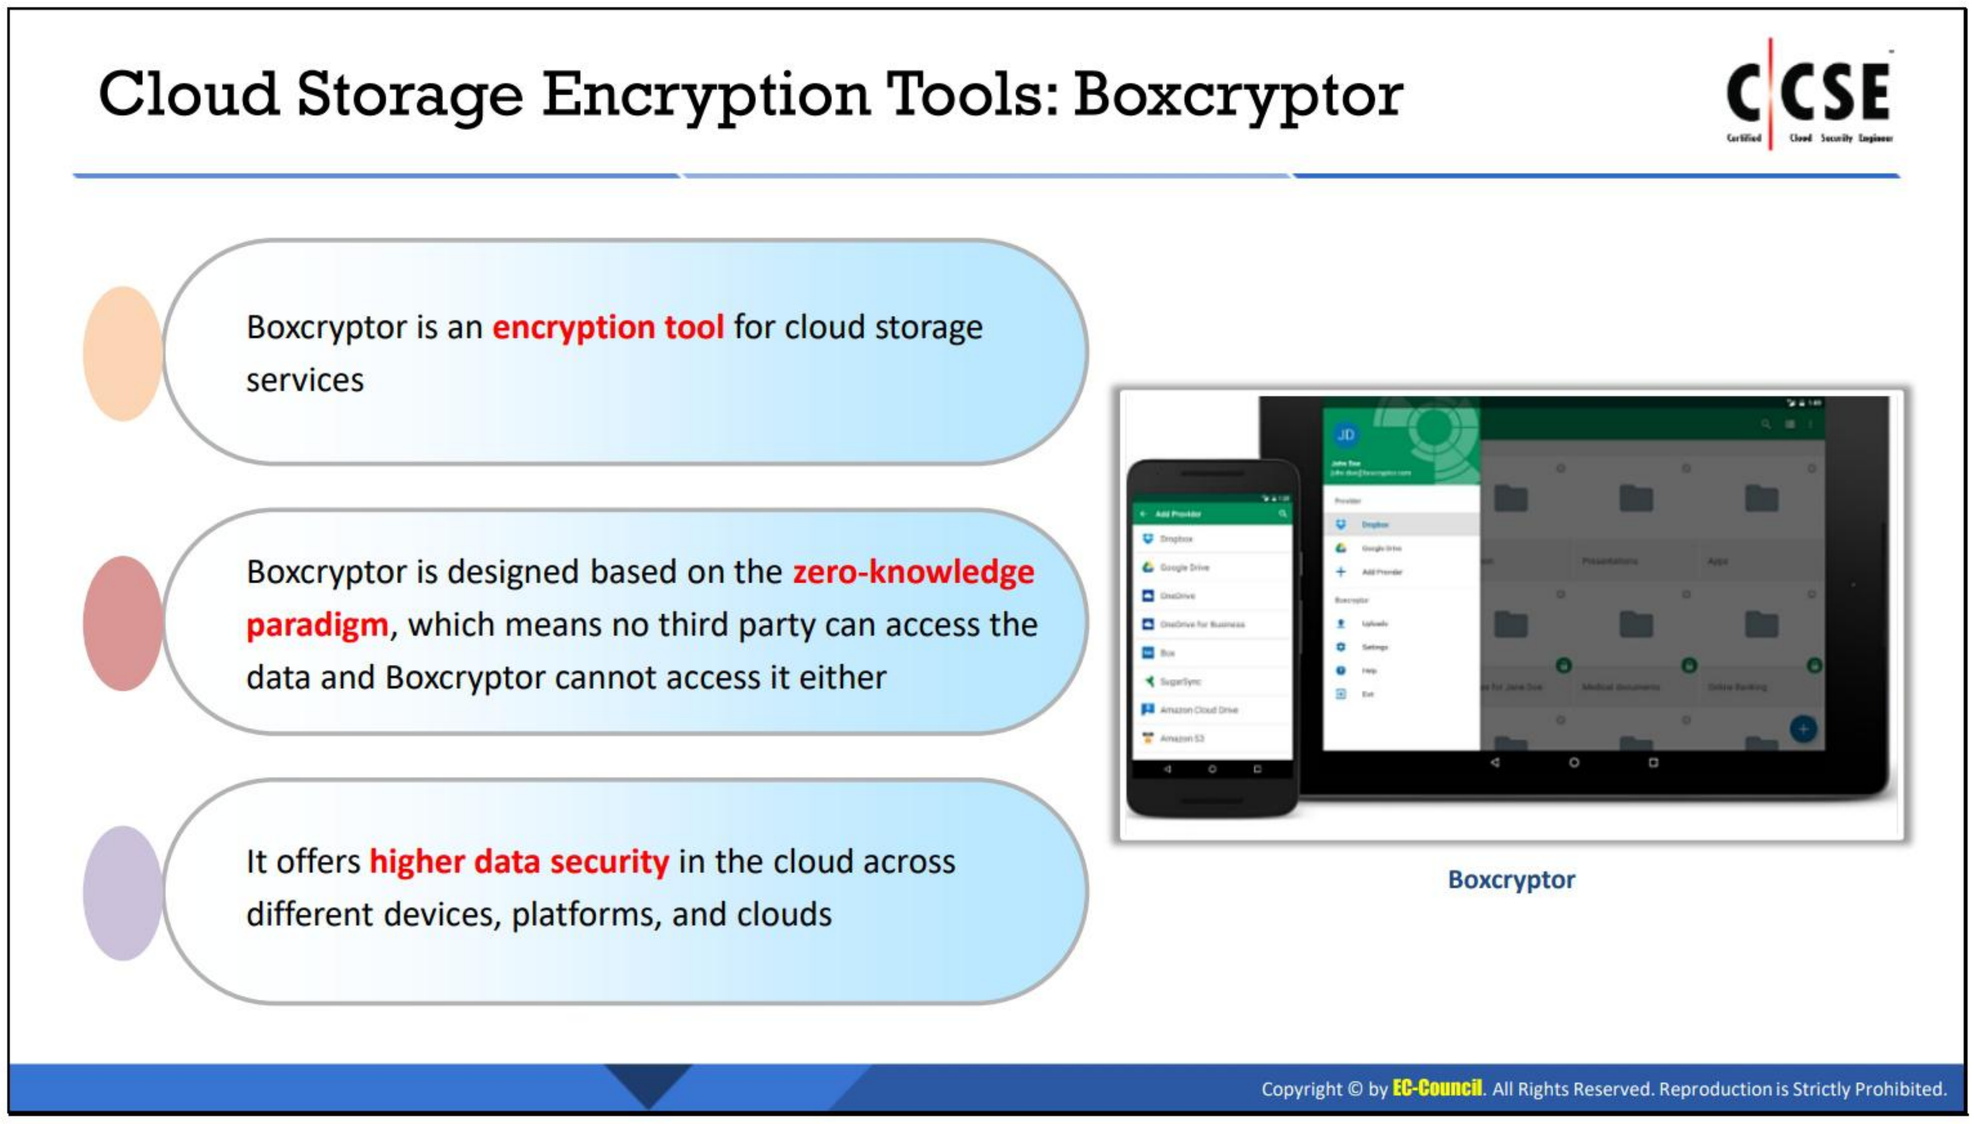
Client-side encryption is a method in which the data is encrypted from the client side before sending it to the SaaS application. It is considered to be a safe and secure method.

In fully homomorphic encryption, however, the operations are performed on the encrypted data rather than decrypting it. When the final data is decrypted, it provides the same output as if the operations were performed on the unencrypted data.

Enterprise key management is a KMS that provides encryption keys across operating systems and databases. It manages the keys throughout their lifecycle, including creation, usage, decryption, and destruction. It protects the keys from SaaS service providers.

When a customer needs selected field encryption, an encryption gateway is enforced within the customer's network. It acts as a proxy server and carefully monitors all the transfer operations. This gateway is implemented following certain rules for encrypting different types of data items. Data is encrypted while sending it to the SaaS application for storage. When it is returned to the client application, it is decrypted to plain text. Through this, only selected people can have access to the plain text as all the data is encrypted in cloud.

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



Cloud Storage Encryption Tools: Boxcryptor

Boxcryptor is an encryption tool for cloud storage services. It supports almost every cloud storage provider. Cloud customers do not need to consider server locations or whether any third parties could access the data. With Boxcryptor, these factors simply do not matter because only the cloud customers can access files in the cloud. Boxcryptor works great with Dropbox, Google Drive, and OneDrive among many others. It provides safety, no matter what storage service is used.

Boxcryptor is designed following the zero-knowledge paradigm. This means that no third party can access the data, and Boxcryptor cannot access it either. It is all under the customer's control and Boxcryptor encrypts all the data on the customer's device before transferring anything to the cloud storage provider.

Boxcryptor automatically detects the installed cloud storage providers and adds them to the virtual Boxcryptor drive. If a user wants to edit an encrypted file, they can only open it, make required changes, and save it. The users do not need to perform any extra work and Boxcryptor handles everything for its users.

Boxcryptor Company and Boxcryptor Enterprise cater to the business requirements of small teams or big companies by adding data security that is compliant with official privacy regulations. It offers premium features such as user management, Active Directory support, two-factor authentication, and many other useful features. Boxcryptor complies with internal and external compliance regulations while keeping the benefits of the cloud.

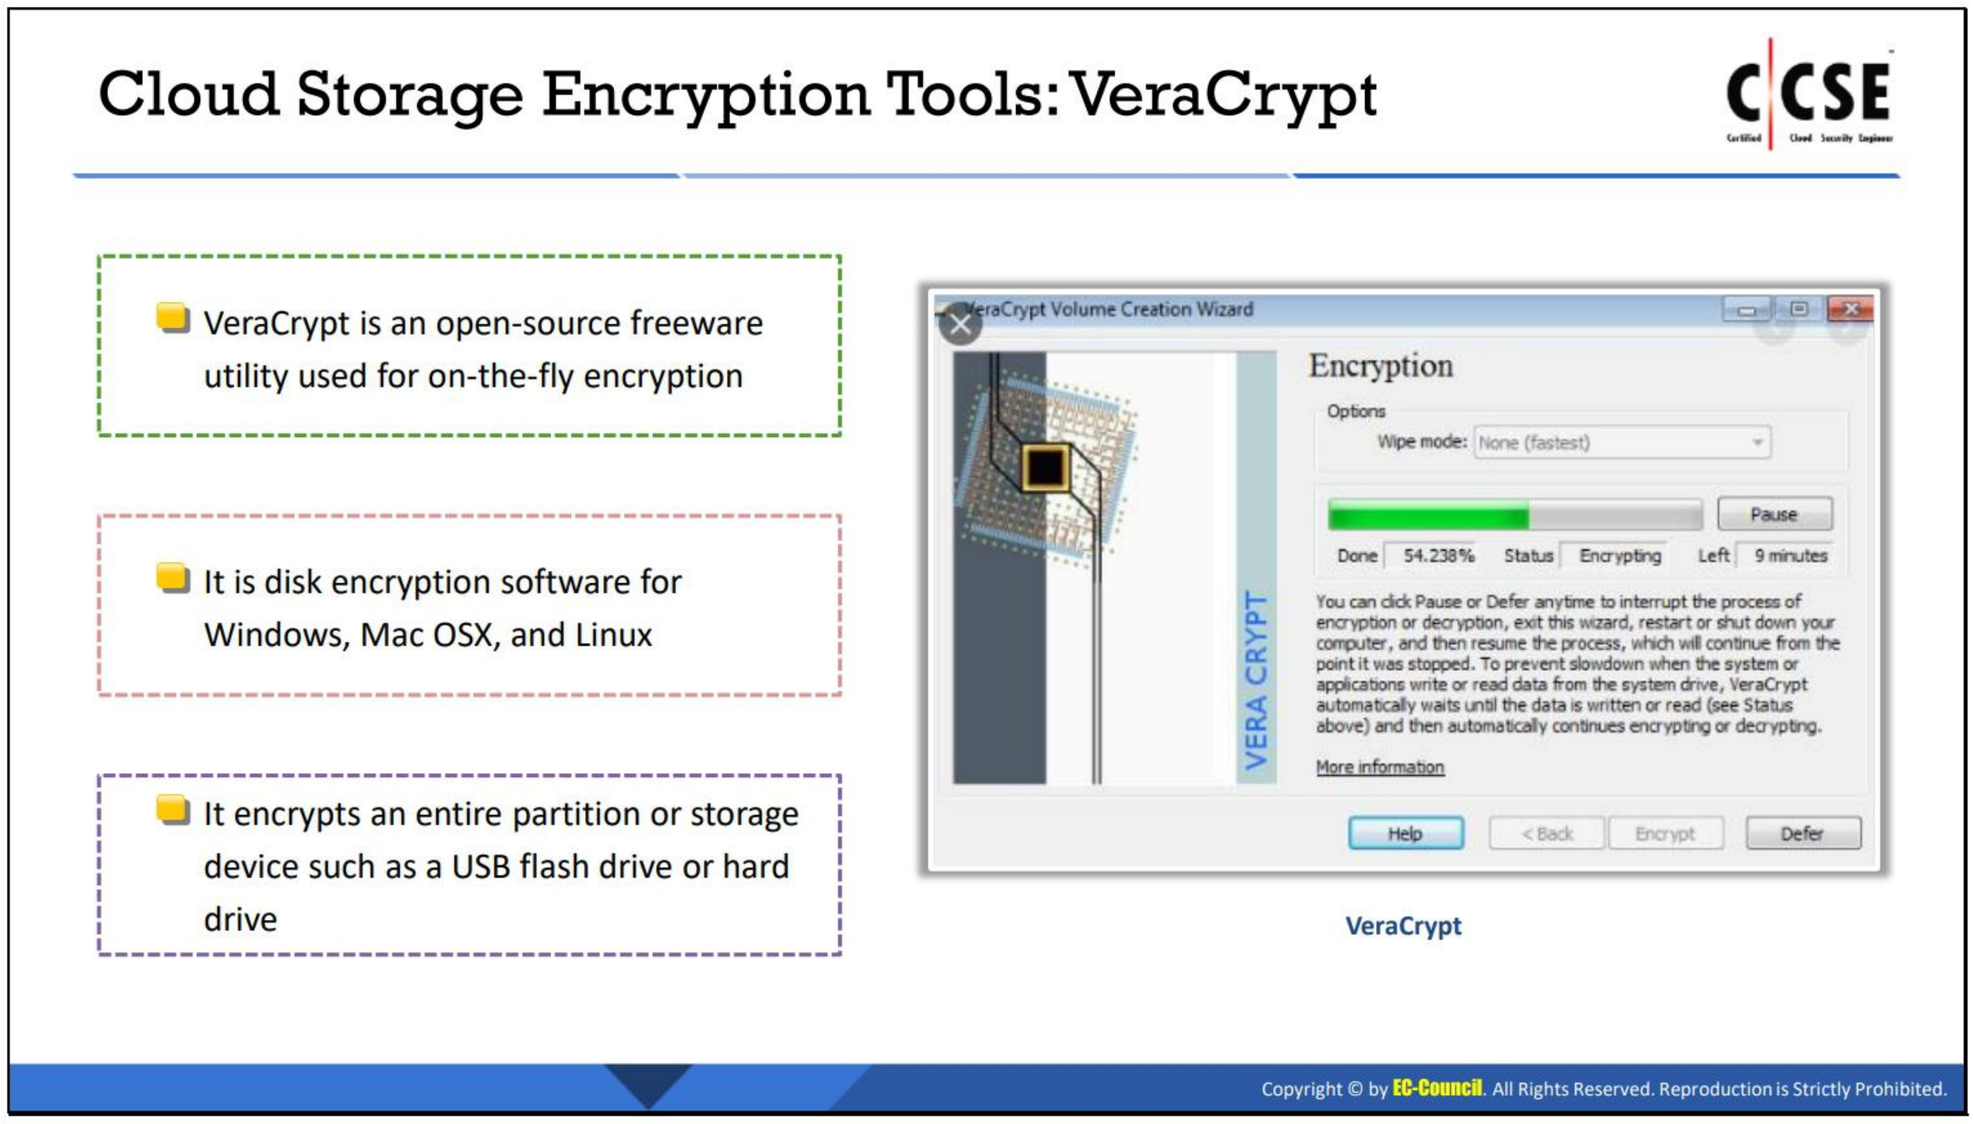
With an additional upgrade, users can enhance data privacy even more with filename encryption. These upgrades are a perfect fit if the users want to add an extra layer of security to the files and folders across multiple devices and cloud storage.

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

Boxcryptor offers an unlimited number of CSPs to store and edit the data securely, based on the selected package. The encrypted files can also be shared with the people who do not use Boxcryptor or the cloud. It offers higher data security in the cloud across different devices, platforms, and clouds. For businesses, Boxcryptor offers various features such as two-factor authentication, single sign-on, custom policies, and a master key provided from an administration panel.

hide01.ir

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



Cloud Storage Encryption Tools: VeraCrypt

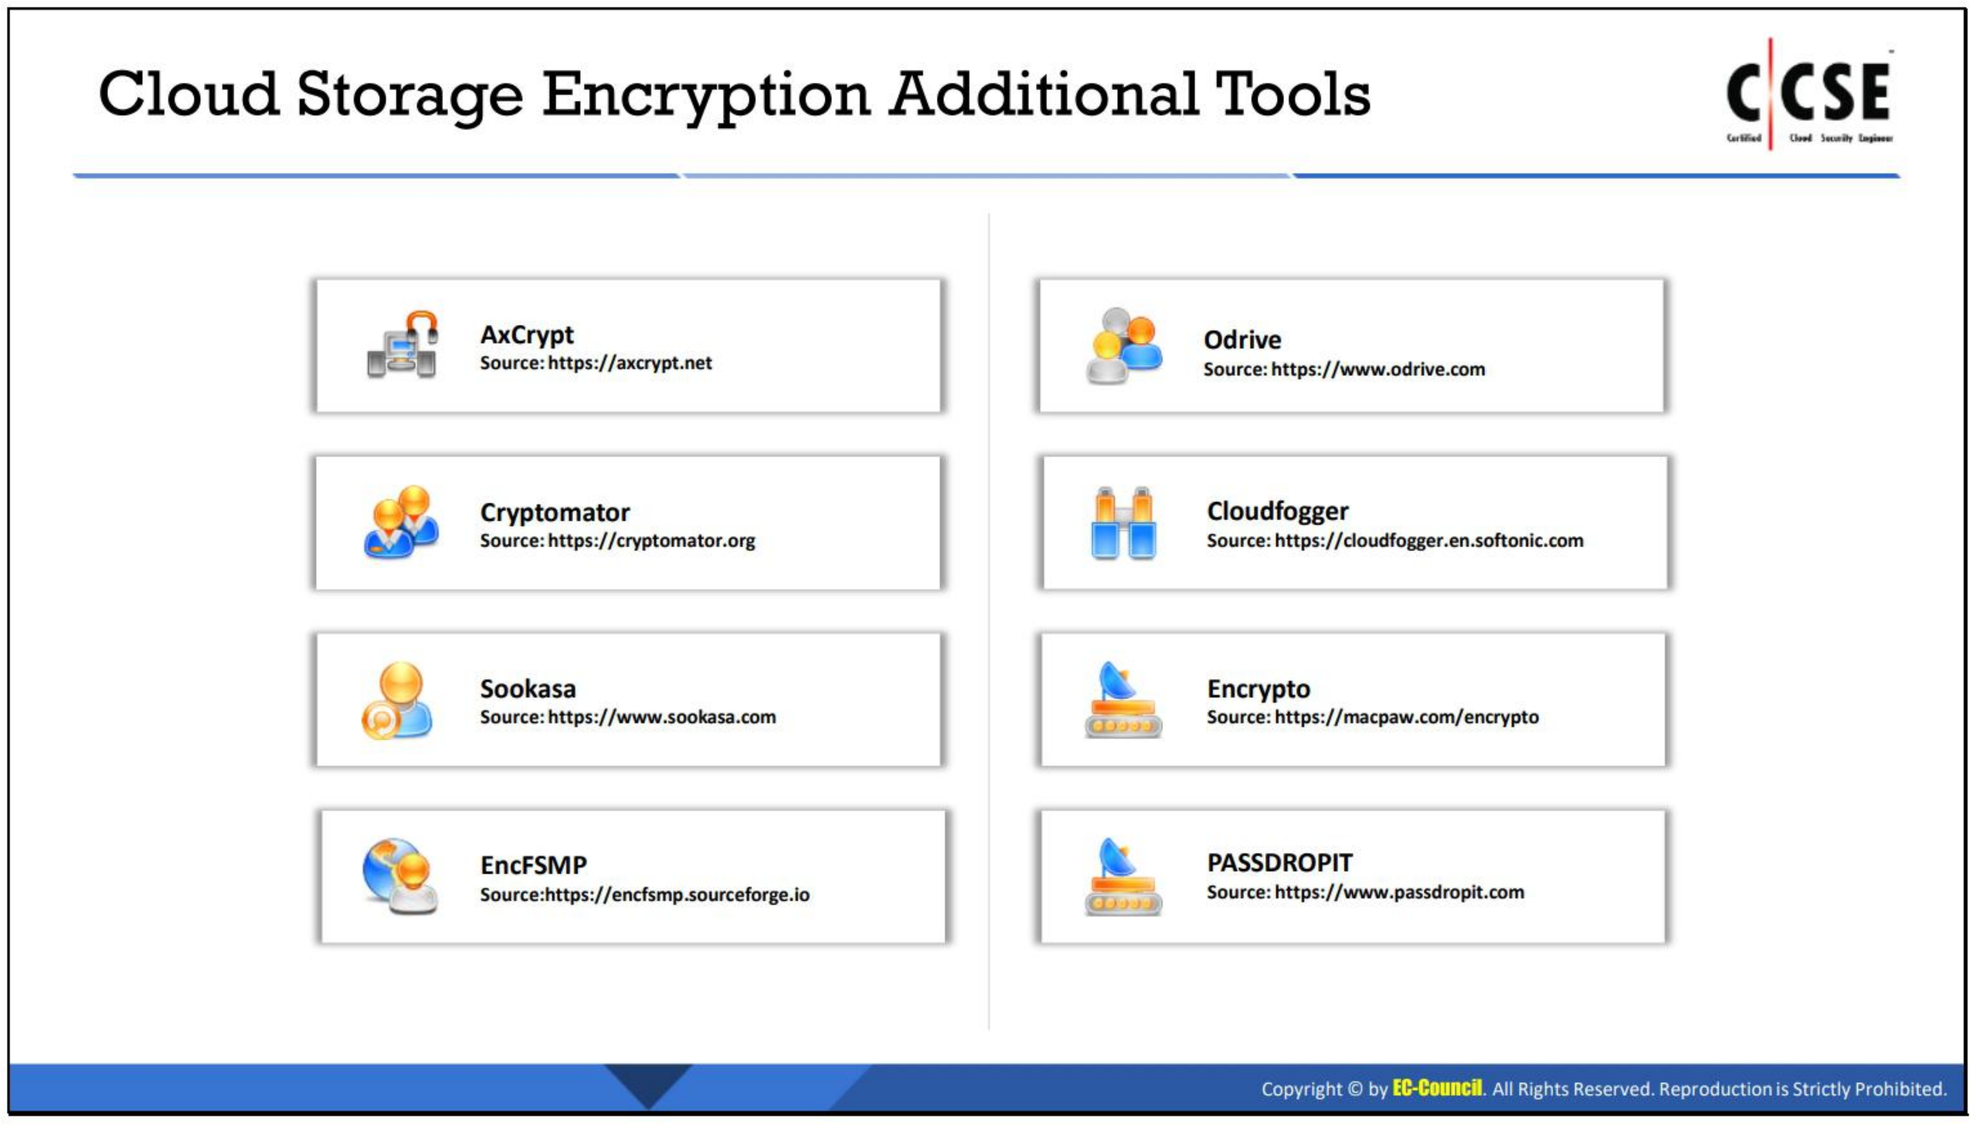
***Source*:** https://www.veracrypt.fr

VeraCrypt is an open-source freeware utility used for on-the-fly encryption. It is disk encryption software for Windows, Mac OSX, and Linux. It encrypts an entire partition or storage device such as a USB flash drive or hard drive.

Main features of VeraCrypt are as follows:

- Creates a virtual encrypted disk within a file and mounts it as a real disk.
- Encrypts an entire partition or storage device such as USB flash drive or hard drive.
- Encrypts a partition or drive where Windows is installed (pre-boot authentication).
- Encryption is automatic, real-time(on-the-fly) and transparent.
- Parallelization and pipelining allow data to be read and written as fast as if the drive was not encrypted.
- Encryption can be hardware-accelerated on modern processors.
- Provides plausible deniability, in case an adversary forces you to reveal the password: hidden volume (steganography) and hidden operating system.

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



Cloud Storage Encryption Additional Tools

There are several other tools available for cloud storage encryption:

- **AxCrypt**

  Source: https://axcrypt.net

  AxCrypt is a software that enables cloud users to easily encrypt files and folders in Dropbox, Google Drive, etc. with 128-bit or 256-bit encryption. It allows secured files to be opened by other AxCrypt users with their own password.

- **Odrive**

  Source: https://www.odrive.com

  Odrive is a tool that allows users to aggregate all the cloud accounts together. The Odrive sync engine makes the cloud files work like local files. Users can open and edit any file, and any changes will be automatically synced with the cloud.

- **Cryptomator**

  Source: https://cryptomator.org

  Cryptomator is a simple tool for digital self-defense that allows users to protect the cloud data independently. It allows the users to access the files from various devices. It is easy to use and integrates seamlessly between the user data and the cloud.

- **Cloudfogger**

  Source: https://cloudfogger.en.softonic.com

hide01.ir

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

Cloudfogger is a file encryption tool that can protect the data when it is sent to the cloud. Cloudfogger uses transparent encryption, which makes daily use fast and simple.

- **Sookasa**

  Source: https://www.sookasa.com

  Sookasa extends client-side encryption, DLP and visibility for the user data in Dropbox, whether it is stored in the cloud or downloaded to end-user devices. It relies on a seamless CASB solution to detect breach risks, encrypt data, and enforce security policies.

- **Encrypto**

  Source: https://macpaw.com/encrypto

  Encrypto encrypts files with AES-256 encryption before sending them to friends or coworkers. Users not only encrypt and share files with Ecrypto but also save them to their own disk.

- **EncFSMP**

  Source: https://encfsmp.sourceforge.io

  With EncFS MP, users can store the data in an encrypted folder. It is important if the sensitive information is stored in cloud services such as Google Drive and Dropbox.

- **PASSDROPIT**

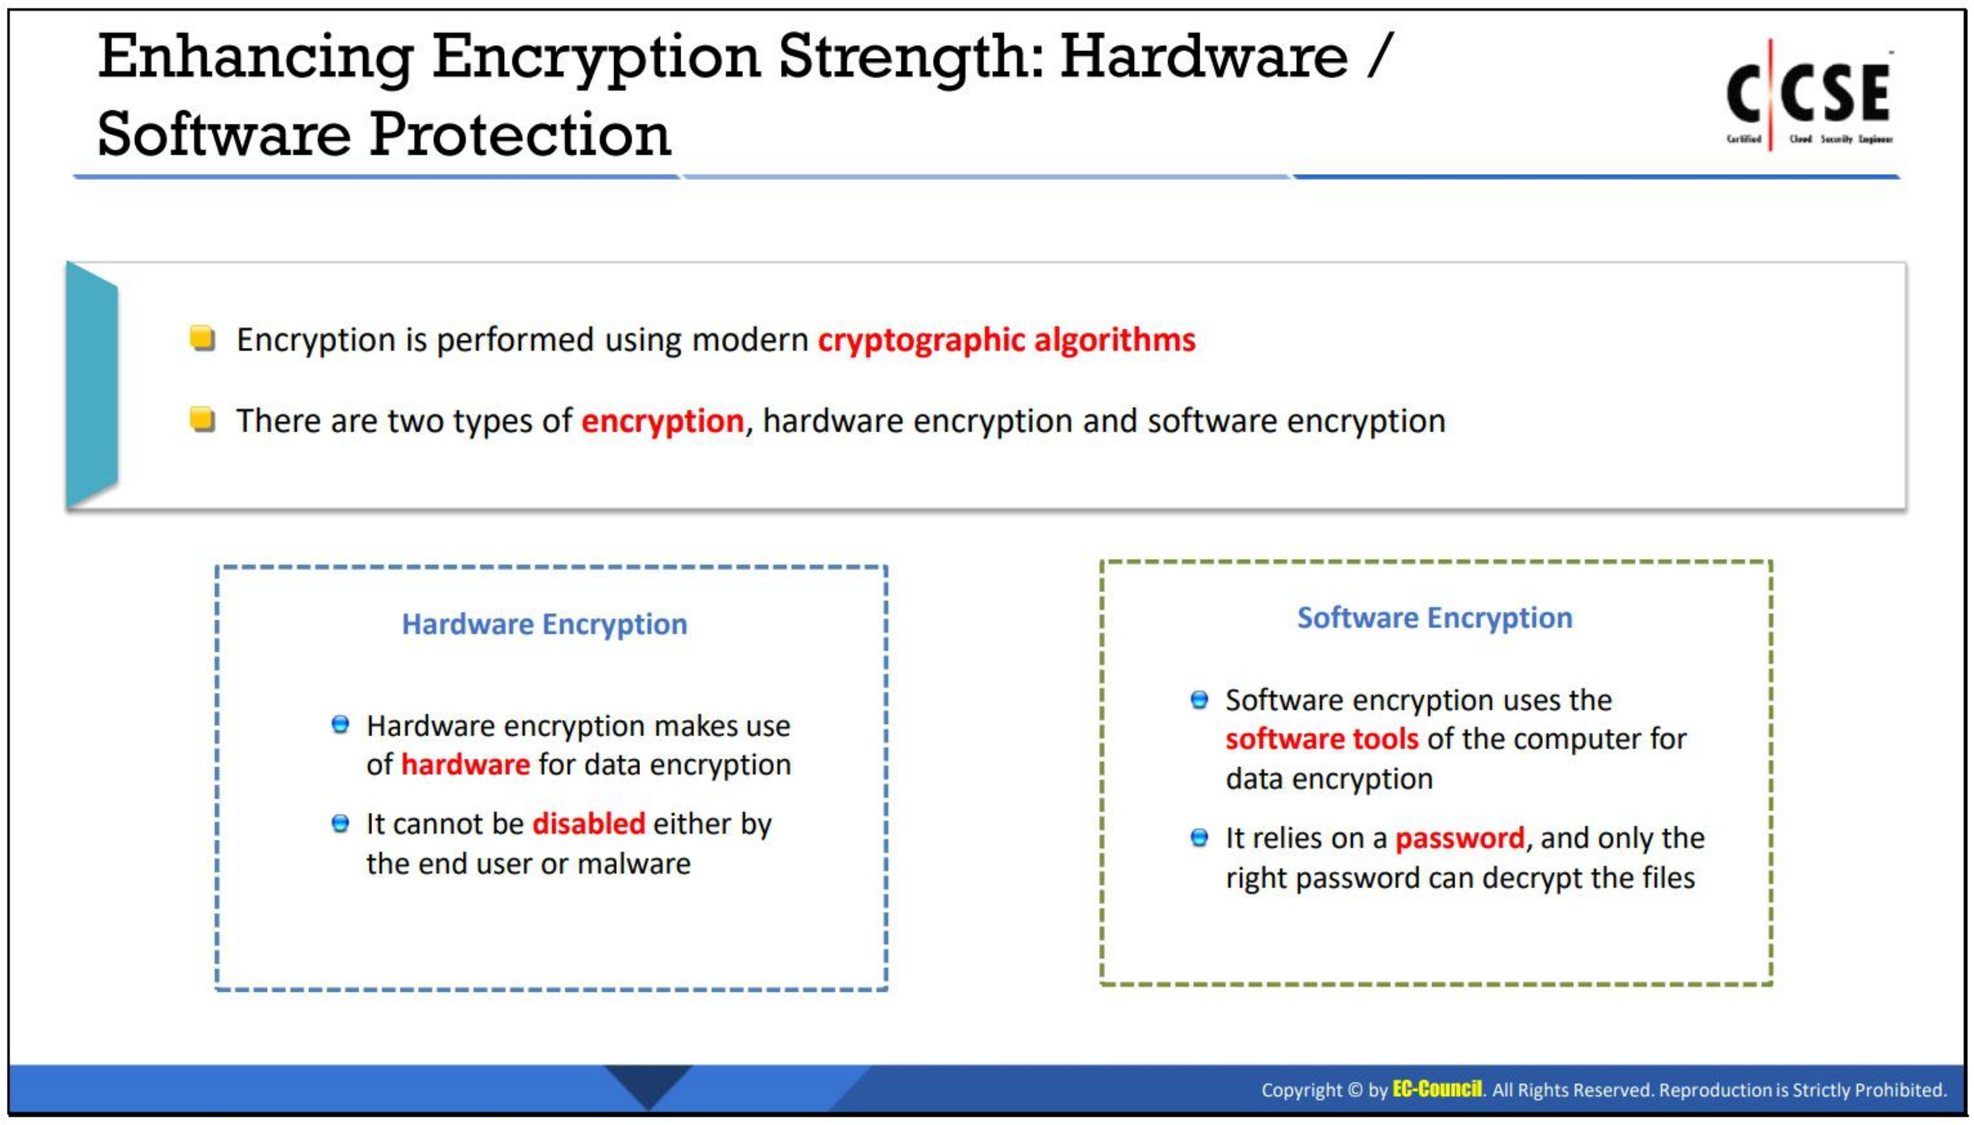
  Source: https://www.passdropit.com

  PASSDROPIT is an encryption tool used for Dropbox. By adding PASSDROPIT to a Dropbox account, it opens various functionalities and saves money. When Dropbox is used along with PASSDROPIT, it offers various special features.

  The features of PASSDROPIT include:

  - Users can share as many secure links as they want
  - Users can customize the URLs
  - The core features are available without an account. The users just need to choose a file and get a link and password

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



Enhancing Encryption Strength: Hardware/Software Protection

Data encryption is the process of converting plain text into an unreadable format or code using cryptographic algorithms. Encryption ensures that the data is secure and cannot be accessed by unauthorized users. Even if an attacker breaches the encrypted files, they will not be able to read the data without a secret key.

Encryption is performed using modern cryptographic algorithms. There are two types of encryption: hardware encryption and software encryption.

- **Hardware encryption:** Hardware encryption makes use of hardware for data encryption. In this method, a separate processor will be assigned for performing encryption tasks. It does not depend on the operating system and is supported by dedicated software to support the processor. It uses a special key to encrypt and decrypt the data, which is a random number generated by the encryption processor.

  Hardware encryption is safe because the encryption method is different from the machine. It can also store security keys and other data on memory locations. Special care is given to hide the secret keys in the hardware system so that they are difficult for the attacker to access. Hardware encryption is also protected against reverse engineering.

  The encryption can never be disabled either by the end-user or malware. It also reduces the burden of relying on the host system as there is a separate processor for encryption and decryption. The encryption and decryption are also faster as they happen in the hardware instead of relying on system processing tools.

- **Software encryption:** Software encryption uses the software tools of the computer for data encryption. Encryption software that runs on a computer generates pseudo-random keys. It uses a symmetric cryptography method in which the same key will be used for

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

encryption and decryption. Although it is safe, any hacker who knows the system and the encryption algorithm can easily get access to the data.

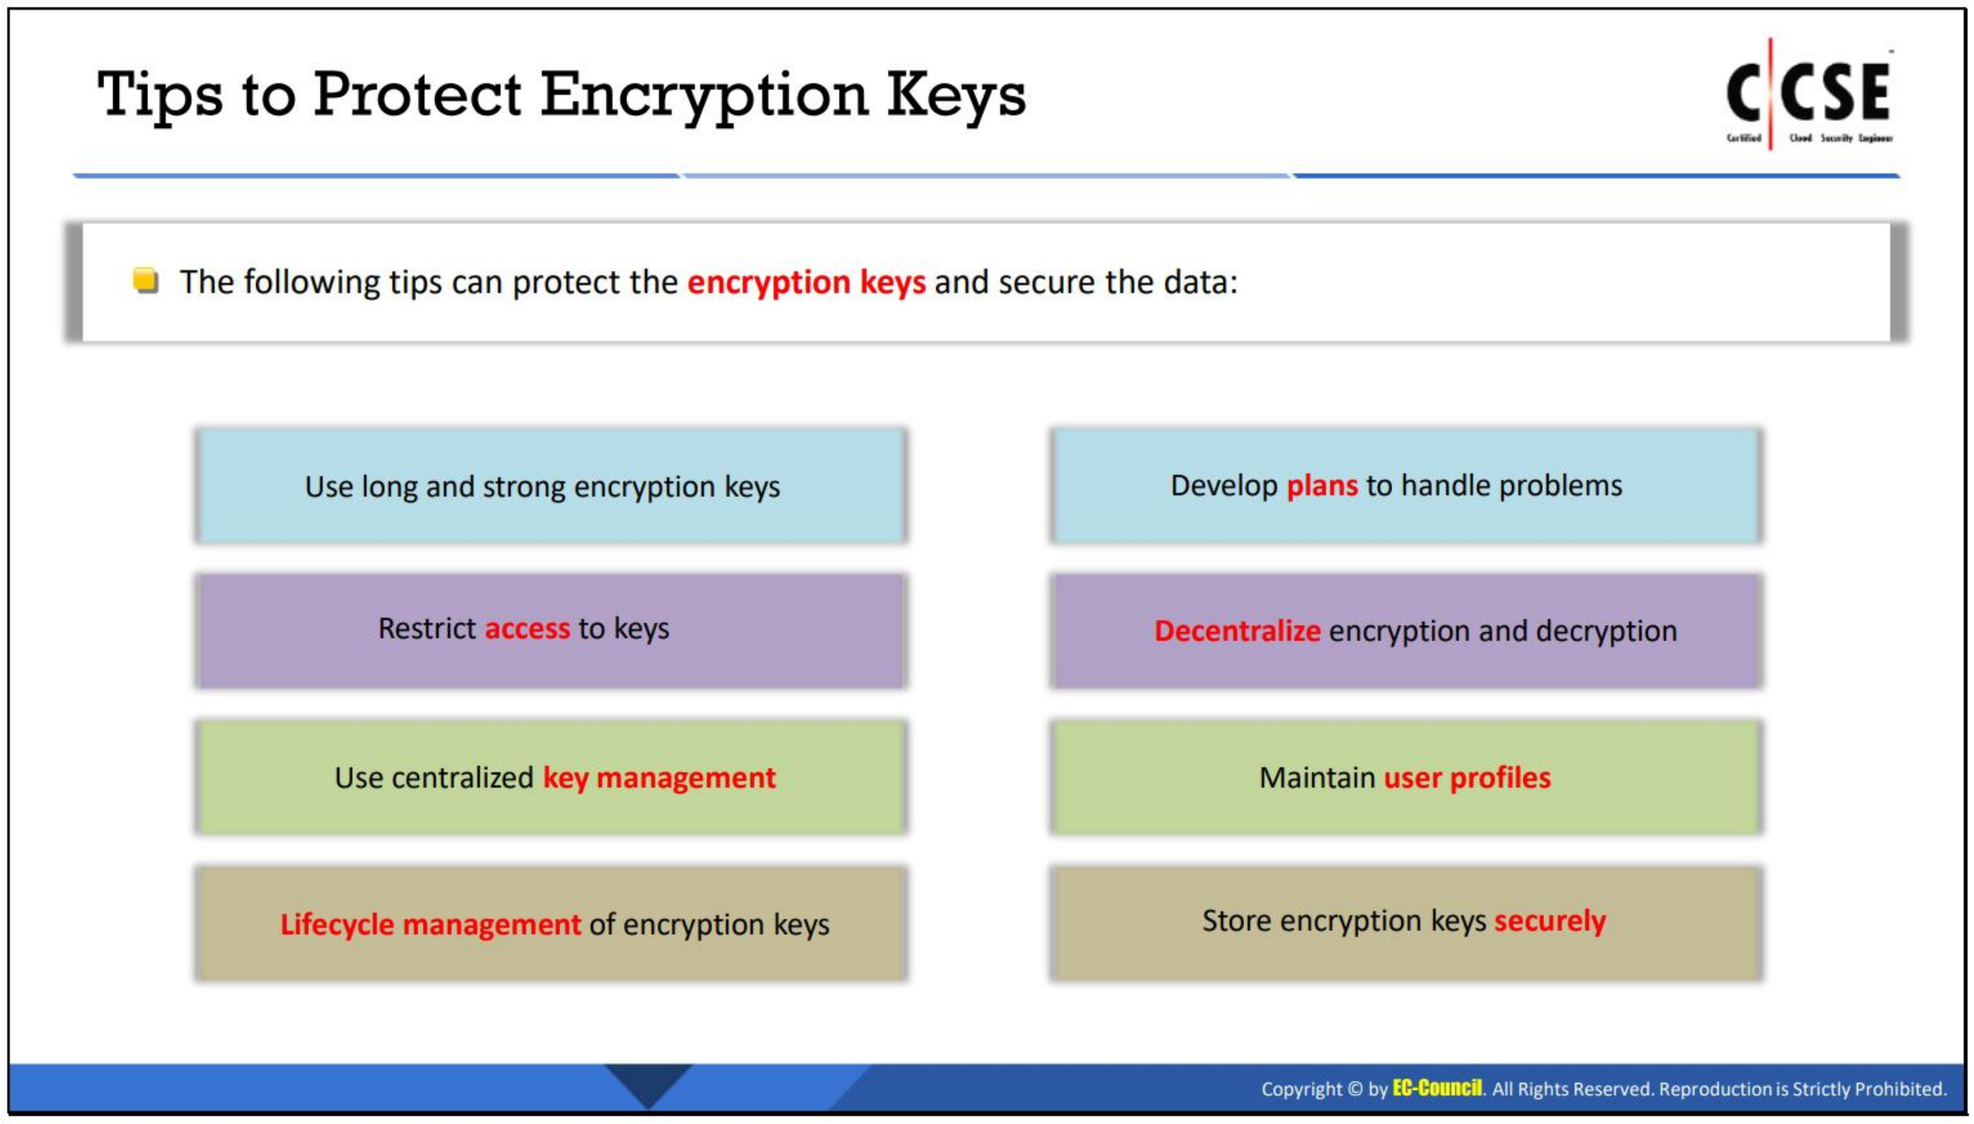
Software encryption relies on a password, and only the right password can decrypt the files. It is dependent on the security and features of the operating system and its supported software. There will not be separate memory locations and data/security keys can be manipulated or stolen.

When software encryption is implemented in a general-purpose system, there are chances of accessing or stealing the information. There is also the possibility of software encryption being disabled by end-users or malware. Further, it is vulnerable to brute-force attacks using any password cracking applications available on the internet.

As software encryption uses the processing resources of the system, the entire system slows down when the data is encrypted or decrypted. There could be a delay while opening or closing the encrypted or decrypted files as it is dependent on the resources.

hide01.ir

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



Tips to Protect Encryption Keys

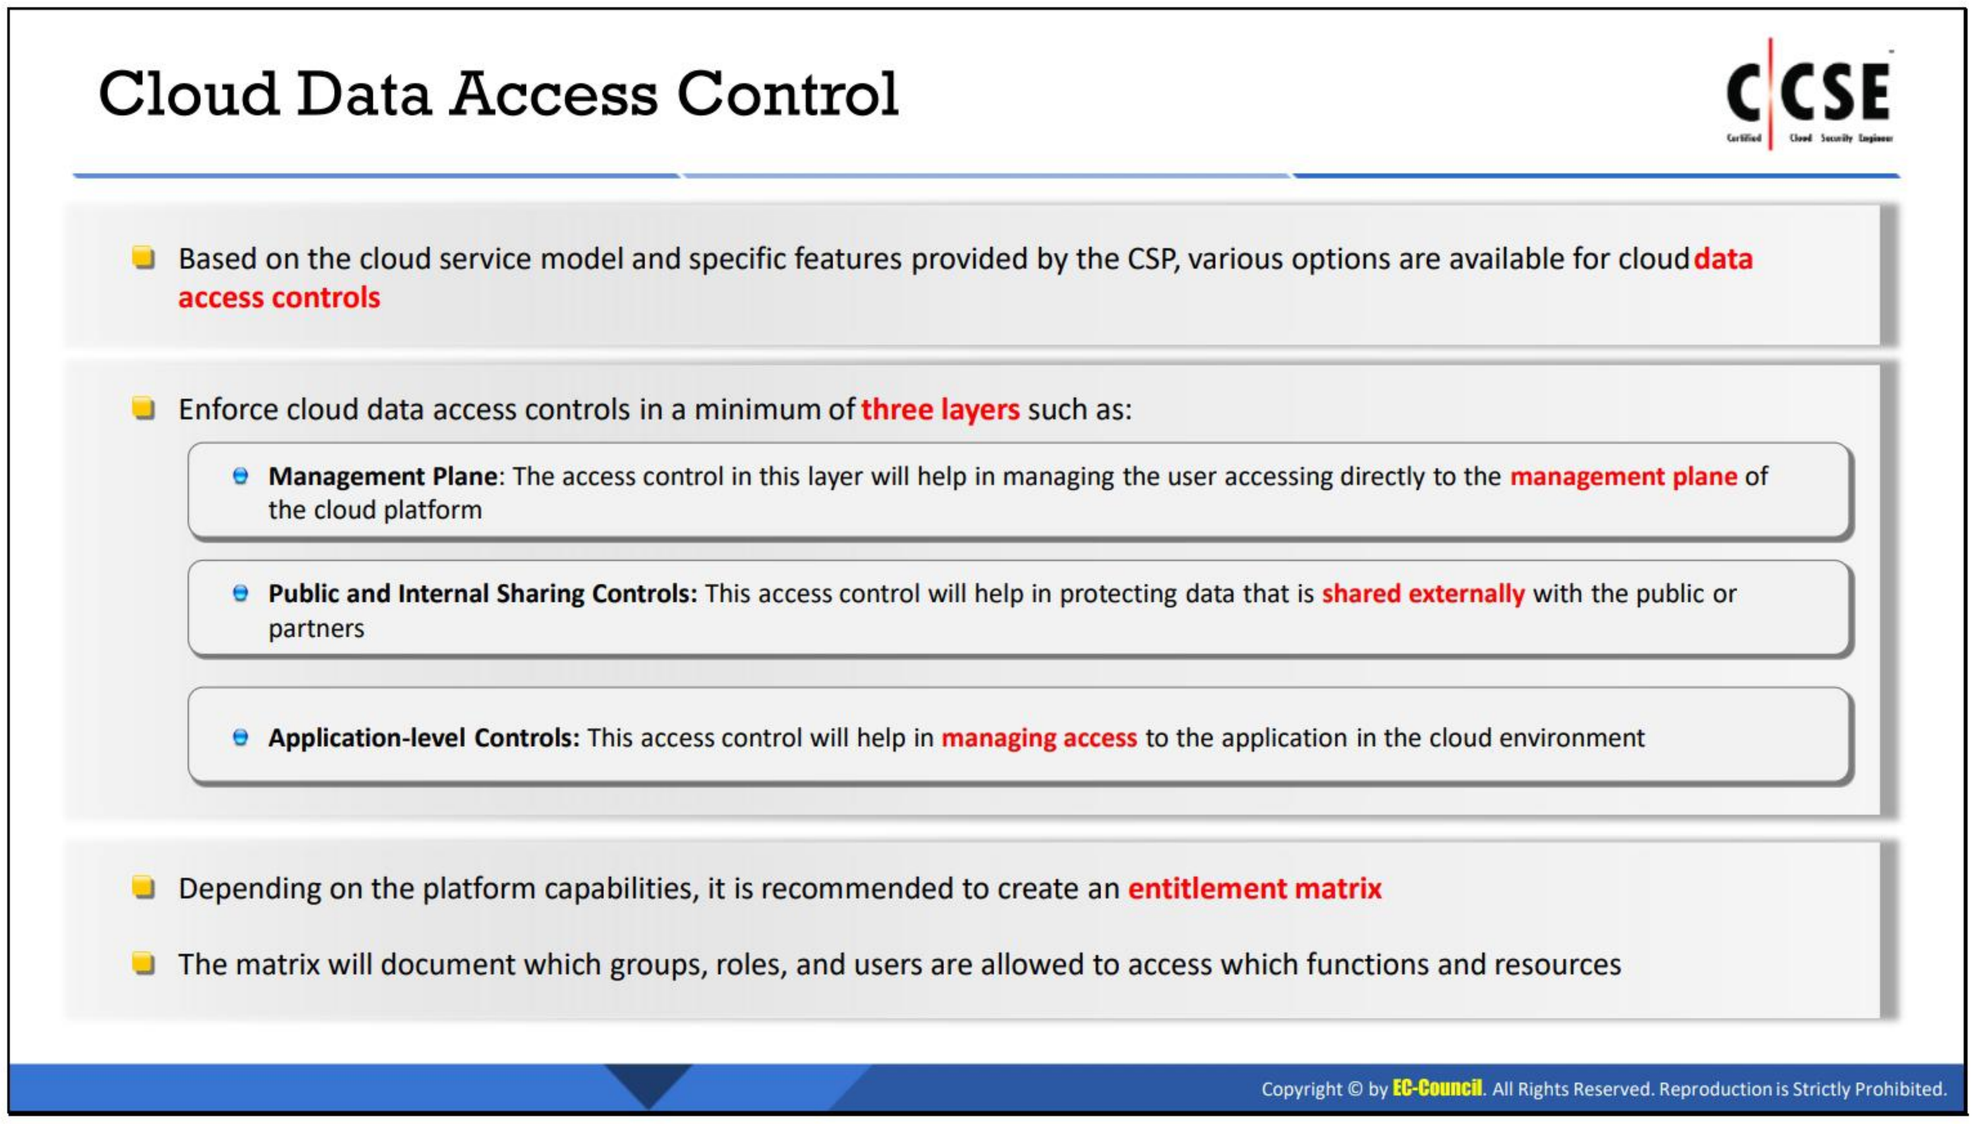
Encryption is an important element to protect data as it is stored and transmitted through various networks. Encryption of data is not difficult but maintaining the encryption keys throughout the lifecycle is indeed difficult. The loss of security keys could cause serious issues and could lead to data becoming inaccessible to the authorized users. Therefore, it is important to follow effective ways to manage the encryption keys throughout the data lifecycle. The following tips could help protect encryption keys:

- **Use long and strong encryption keys:** The encryption key should be of the maximum key length, which would make it difficult for the attacker to find the encryption key. The encryption should be strong so that the hacker needs to go through many layers of encryption to crack the key.
- **Restrict access to keys:** The cloud security professional should ensure that only the required users, administrators, resources, and services should have access to the key. If multiple cloud administrators have access to the keys, an attacker could breach the keys from the storage devices and gain access to the sensitive data. Therefore, the cloud security professional should properly plan regarding who should have access to each key.
- **Use centralized key management:** Protecting the encryption keys could get complicated as there can exist a conflict while controlling access to the keys that need to be distributed throughout the organization. To handle this, the cloud administrator can implement centralized encryption key management services. This will automate and secure all the key management activities such as storage and distribution. Centralized key management reduces the exposure of keys to attackers as the keys are stored in a centralized location.

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

- **Lifecycle management of encryption keys:** All the encryption keys must be protected with a KMS from the moment they are created. The key management solution is responsible for creating, maintaining, and controlling the use of encryption keys throughout their lifecycle. By enforcing key management policies, every stage of the key lifecycle management will be affected. Every encryption key should have a key usage policy mentioning which devices or applications can request the key and what operations can be performed using the keys.
- **Develop plans to handle problems:** A user may forget the credentials to access the key, change job responsibilities, or discover flaws in the cryptographic algorithm. The cloud administrator should discover and implement necessary plans to handle such problems. The administrator should also ensure that the procedures are revised on a regular basis and are accurate; this could save time in case of any problems.
- **Decentralize encryption and decryption:** If the encryption and decryption keys are distributed, it is important to have a secure distribution and management of keys. Encryption at each level such as a file, application, and database level can ensure the highest level of security. This also allows the right users to access the data immediately and ensures high performance and efficient use of bandwidth. It also provides high protection to data by moving it frequently.
- **Maintain user profiles:** Users are the people who access sensitive data. The access to the data should be based on the user profiles. The users should be provided with the credentials to access the encrypted resources based on the user profile. All the user profiles are managed by the administrator role.
- **Store encryption keys securely:** Most attackers target the weak point of the security system. Thus, the cloud security professional must ensure that the KMS is secured. When the encryption keys are stored in cloud storage, the cloud customer should ensure that the keys are kept safe by the cloud service provider. Cloud customer should try to develop an encryption system where they can control the keys to the cloud data. This adds additional security to the encryption keys. Several best practices have been developed to protect the encryption keys and data. Using these methods, cloud consumers can easily secure their keys along with the data in the cloud.

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



Cloud Data Access Control

Based on the cloud service model and specific features provided by the CSP, various options are available for cloud **data access controls**

hade01.ir

Enforce cloud data access controls in a minimum of **three layers**.

- **Management plane**: The access control in this layer will help in managing the user directly accessing the **management plane** of the cloud platform. For example, by default, deny access to IaaS applications web consoles to prevent unauthorized access to the application data
- **Public and internal sharing controls:** This access control helps in protecting data that is **shared externally** with the public or partners.
- **Application-level controls:** This access control helps in **managing access** to the application in the cloud environment

Depending on the platform capabilities, it is recommended to create an **entitlement matrix,** which will document the functions and resources the various groups, roles, and users are allowed to access.

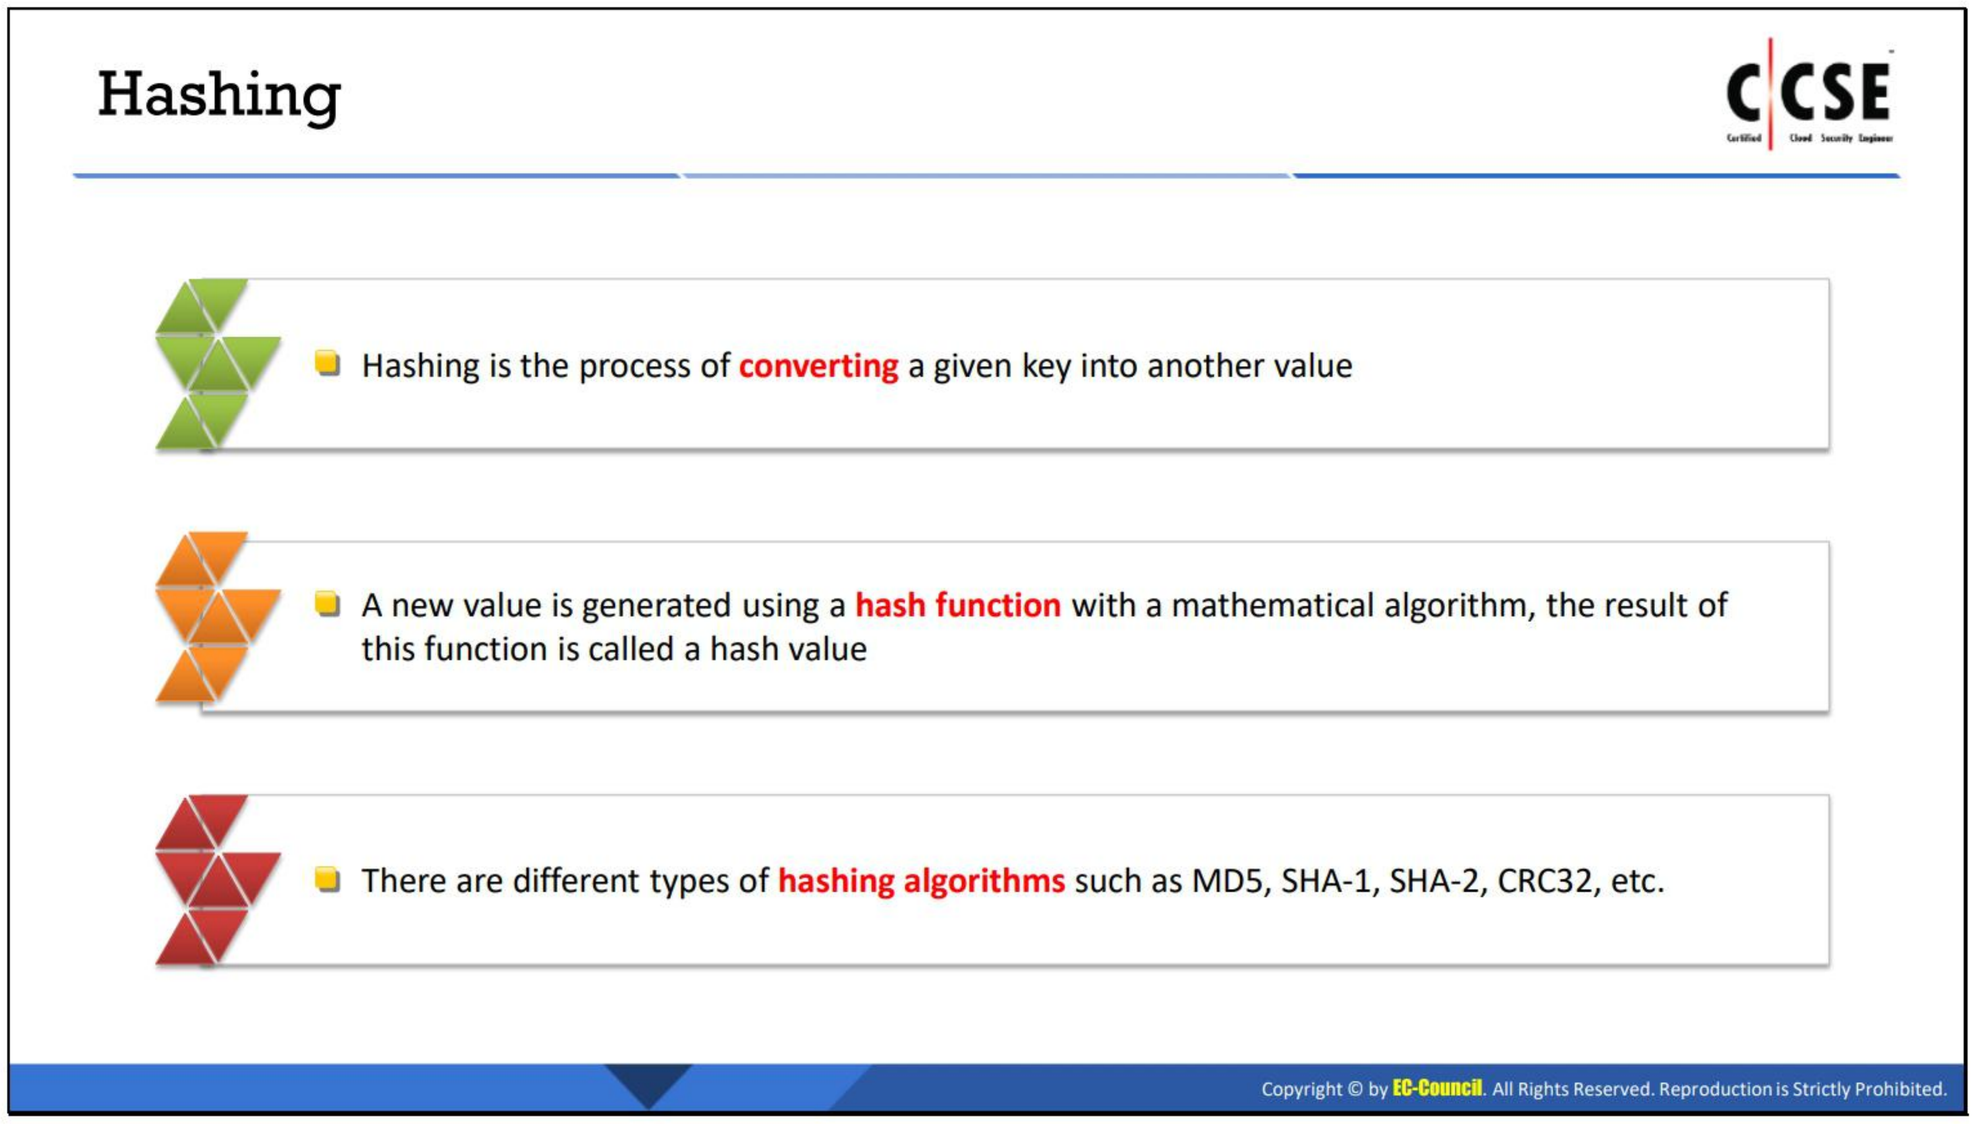
Certified Cloud Security Engineer Copyright © by **EC-Council**
All Rights Reserved. Reproduction is Strictly Prohibited.

Entitlement	Super admin	Admin1 (Service)	Admin2 (Storage)	Developer	Security auditor	Security engineer
Describe volume	Y	Y		Y	Y	Y
Describe objects	Y		Y	Y	Y	Y
Edit volume	Y	Y		Y		Y
Read logs					Y	Y

Table 4.1: Entitlement Matrix

hide01.ir

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



Hashing

Confidentiality of data in cloud environments is a critical factor for every organization. Due to various reasons such as performance, technical abilities, and cost of the services, encryption cannot be the superlative method for the protection of data in every situation. To ensure the confidentiality of data, alternative mechanisms such as masking, obfuscation, and anonymization can be used.

Hashing is the process of converting a given key into another value. A new value is generated using a hash function using a mathematical algorithm, the result of this function is called a hash value. There are different types of hashing algorithms such as MD5, SHA-1, SHA-2, CRC32, etc. Formula generated hashes are always of the same length as MD5 always generates hashes that are the length of 32 characters irrespective of text used in hashing.

Some of the important hashing algorithms are as follows:

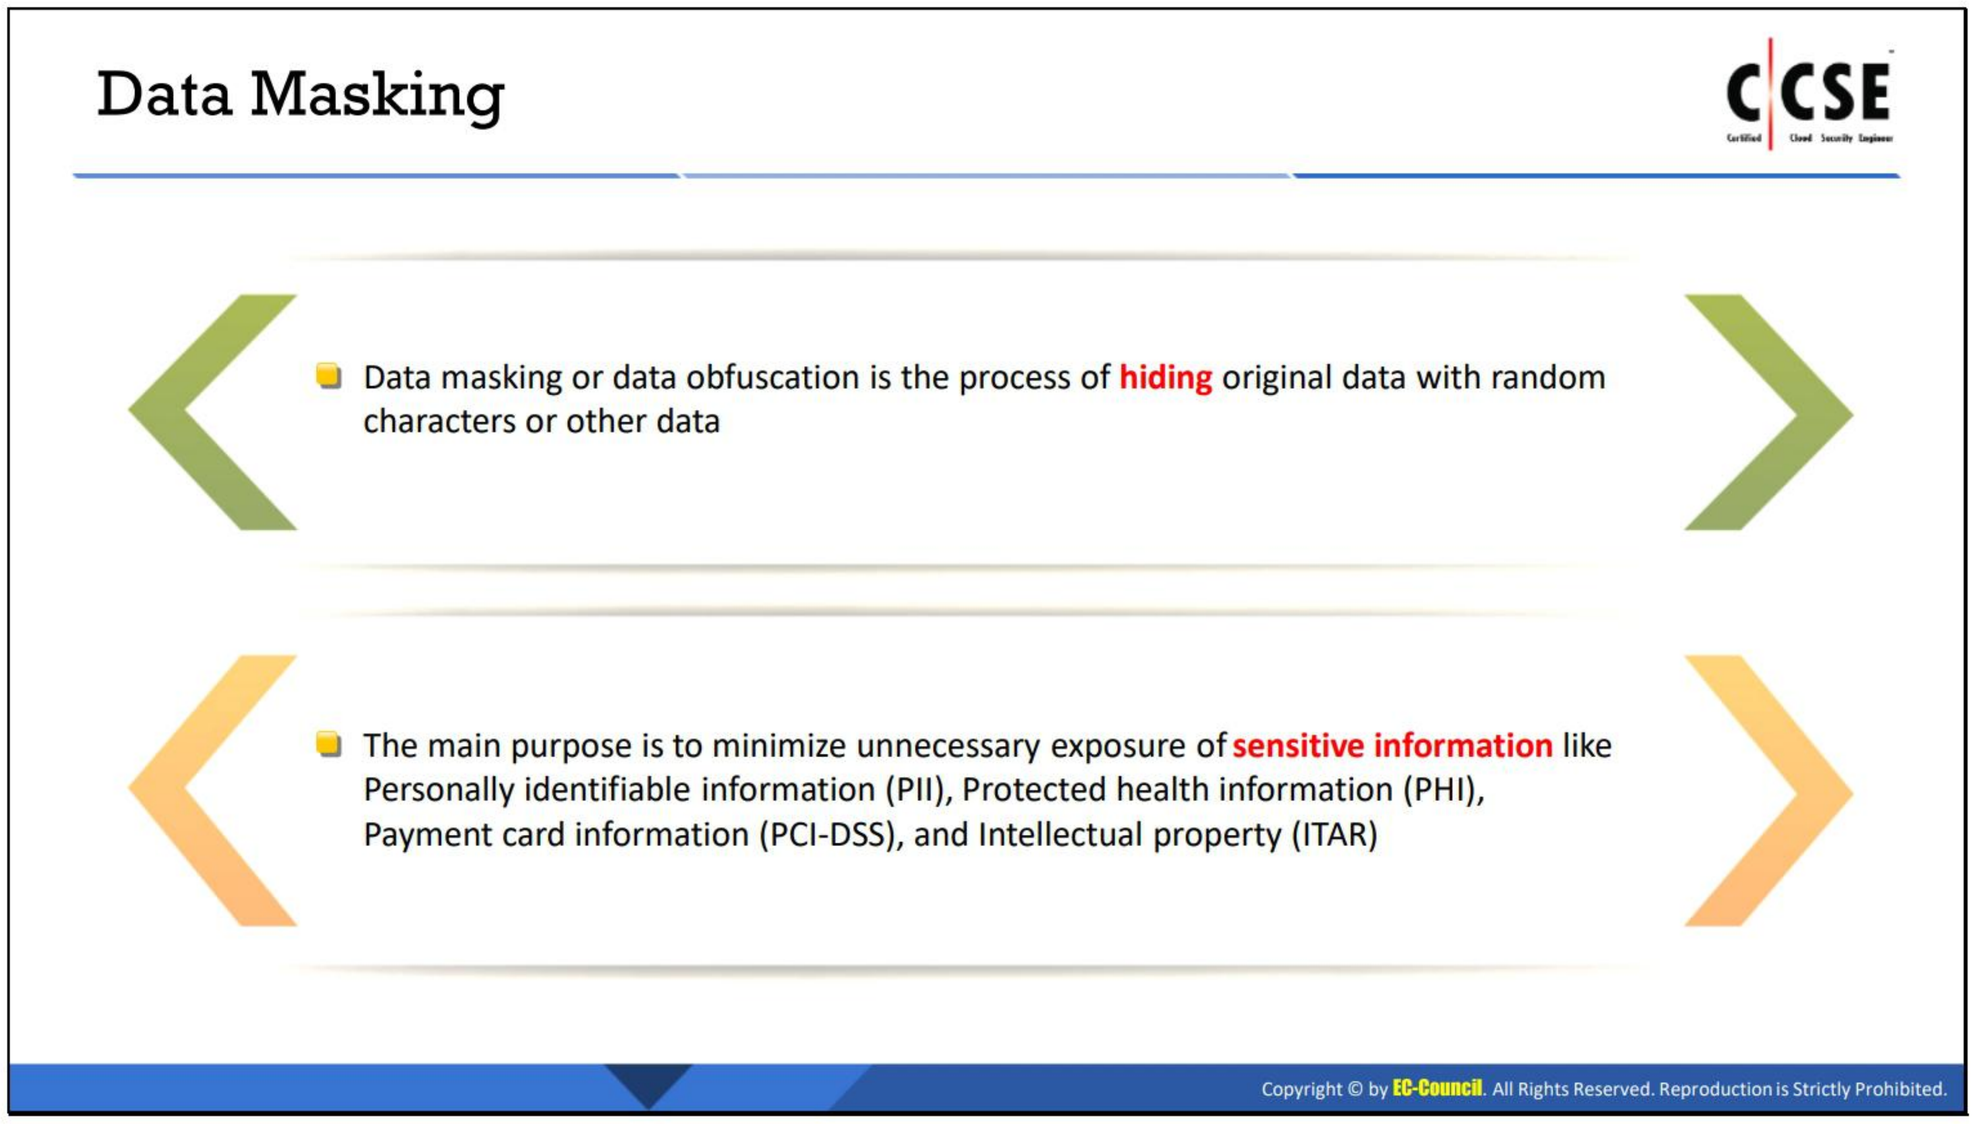
- **MD5**: It is commonly used to verify data integrity; it encodes information into a 128-bit fingerprint. Due to its age, it is vulnerable to collision vulnerabilities (i.e., two keys resulting in the same hash). Collisions can be handled using other methods.
- **SHA-1**: To address the vulnerabilities of MD5, SHA-1 was developed. It became the most renowned hashing technique and the majority of the users use it. It creates an output of 160 bits.
- **SHA-2**: It is a cryptographic hash function that was developed by National Investigation Agency (NSA). SHA-2 consists of six hash functions and they are SHA-224, SHA-256, SHA-384, SHA-512, SHA-512/224, SHA-512/256. It is used for the protection of unclassified information.

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

- **CRC32**: It refers to "cyclic redundancy check" and functions as an error detection code. It is used to verify accidental changes to data. Presently, CRC32 is rarely used for integrity checks.
- **LANMAN**: It was a service provided by Microsoft and functioned as the LAN manager hashing algorithm. It was used to store passwords by legacy windows systems. DES algorithm is used in LANMAN to create a hash. LANMAN implementation was not very secure and was vulnerable to brute-force attacks. Presently, LANMAN is not used for the storage mechanism.
- **NTLM**: It is an NT LAN Manager algorithm developed by Microsoft. It is generally used to hash passwords during authentication. It is the replacement of the LANMAN algorithm. NTML protocol is implemented in a component windows API, which is a security-related authentication; it combines protocols such as NTLMv1, NTLMv2, and NTLM2 in a single package.
- **HMACSHA1**: It provides message authentication using a secret key rather than digital signatures with the use of asymmetric cryptography. As with any MAC, HMACSHA1 may be used for the verification of both the authenticity of a message and for data integrity.

hide01.ir

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



Data Masking

Data masking or data obfuscation is the process of hiding original data with random characters or other data. Its main purpose is to minimize unnecessary exposure of sensitive information such as personally identifiable information (PII), protected health information (PHI), payment card information (PCI-DSS), and intellectual property (ITAR). It replaces vulnerable/sensitive data with fictitious functional data that seems real, allowing it to be used safely in operations where the original data is not required. During data masking, the basic format of the data remains the same, but its key values get changed. For example, a long card number is masked in the following manner:

2424 6789 4545 3421 will be masked as **2424 XXXX XXXX 3421**

Reasons to Include Data Masking in Data Security/Data Masking and Security

- **Protecting non-production data:** In many organizations, multiple copies of production data are required for non-production purposes such as application development and testing, employee training, business analytics modeling, etc. The multiple copies of sensitive data increase the chance of unauthorized exploitation. While allowing the safe sharing of sensitive data, the data masking technique also plays a vital role in securing these data sets and meeting the compliance requirements without affecting business operations.
- **Protection against insider threats:** Data masking allows employees to access the masked data instead of accessing the real production data to get their jobs done, which can mitigate the risk of a data breach from malicious insiders.
- **Protection against third parties:** Use of the data masking technique is suitable to send original personal identifiable data, PCI or PHI to third parties such as market researchers.

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

- **Comply with regulations:** Data masking helps organizations comply with key regulations such as the GDPR passed by the European Union to make personal data secure.

Types of Data Masking

There are different types/strategies of data masking. To select a data masking type, take into consideration the size of the organization, the location (cloud vs. on-premises) and the complexity of the data.

- **Static data masking (SDM)**: SDM type masks sensitive **data-at-rest**, or the data stored in the original database environment. SDM is implemented to provide realistic data for the development and testing of alications. Realism is necessary as it helps development and testing teams to detect the defects in the early stage of the development cycle. SDM also secures data for use in analytics and training, and facilitates compliance with standards and regulations.

  SDM technique applies a series of data transformations on the actual production data to generate high-quality masked data. Various approaches are used to perform SDM.

- **Statistical data masking:** There are alternatives to SDM that rely on the stochastic perturbations of the data that preserve some of the statistical properties of the original data. For example, differential privacy method and DataSifter method.
- **Dynamic data masking (DDM):** Dynamic data masking temporarily masks sensitive data-in-transit without affecting the original data-at-rest. DDM is mainly used to provide role-based security for databases/applications and it uses database and proxy approaches to mask the database. The database proxy alters the SQL query issued by the analyst before transferring it to the database so that the masked data is returned to the analyst.
- **On-the-fly data masking:** Similar to DDM, on-the-fly data masking occurs on demand but for one record at a time. An extract transform load (ETL) process occurs wherever the data is masked within the memory of a given database application, which is helpful for agile companies that target continuous delivery or continuous deployment practices.
- **Data masking and the cloud:** There are different modes of creating test data and moving the test data from on-premises to the cloud, or among environments inside the cloud. If data is masked before being ingested into a cloud-native data repository and if a breach happens, there is nothing the attackers will get from the stolen data. A few cloud-native data services offer built-in masking tools that make the implementation of masking easy in the cloud.

Data Masking Techniques

The various data masking techniques used to obscure confidential data are described below.

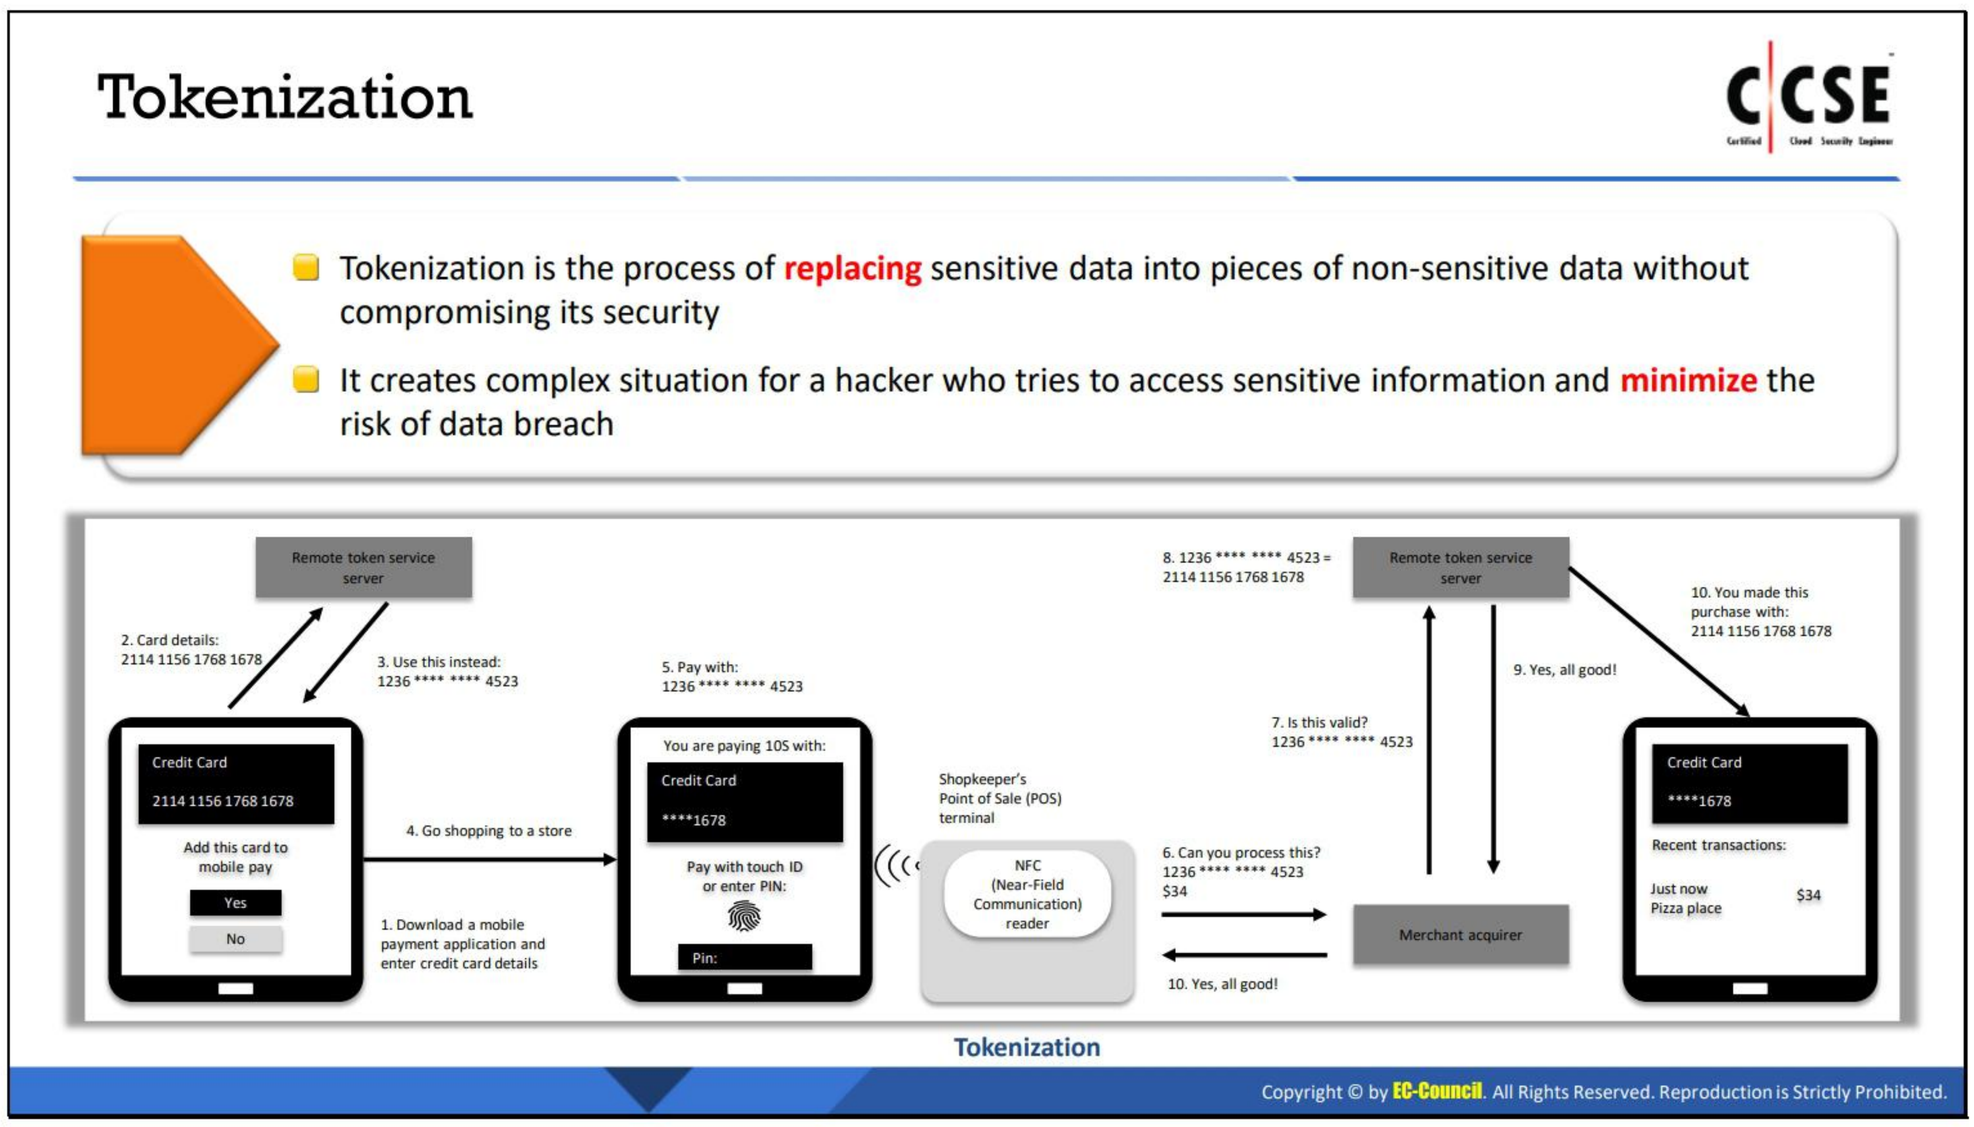
- **Encryption:** Different encryption algorithms are used to mask the data. Only the users who have authorized key can access the data.
- **Character scrambling:** Characters are jumbled into a random order to mask the original data. For example, an emp_id #123456 in a production set of data will be masked as #425613.

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

- **Lookup substitution:** A lookup table is added to provide an alias for the original value. It hides the authentic data and protects the data from security breaches.
- **Nulling out or deletion:** Data becomes null to unauthorized users. This technique is only helpful to avoid the visibility of the data element.
- **Shuffling:** Data in an individual column is randomly shuffled or swapped. The output set seems to be real but does not disclose any sensitive information.
- **Number and date variance:** Number and date variance is helpful for financial and date fields. Each number/date in the data set is masked by a random percentage of its real value.
- **Masking out:** Only a part of the data is masked with mask characters (for example, an X). This technique is mainly used to mask credit card details. For example, when an operator makes a bill for an item to a customer's credit card, the operator will see only last 4 digits while all other digits are masked as XXXX XXXX XXXX 7685.

hide01.ir

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



Tokenization

The process of replacing sensitive data into pieces of non-sensitive data without compromising its security is known as tokenization. It is used in processing data such as bank account numbers, credit card numbers, loan applications, etc.

hide01.ir

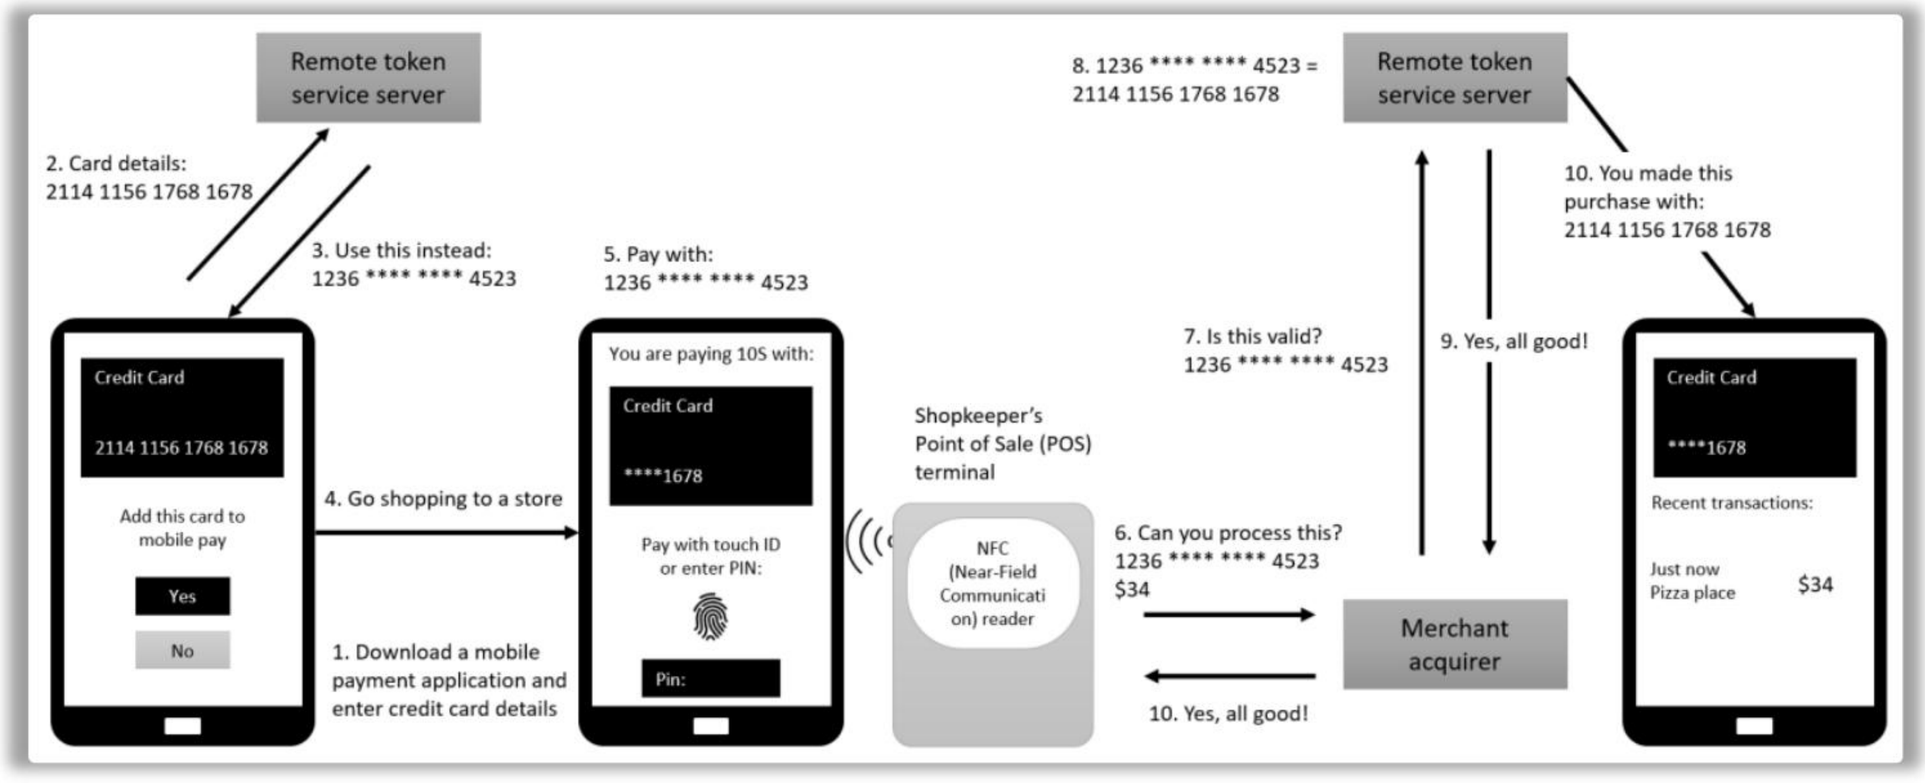


Figure 4.40: Usage of Tokenization

When a user initiates a payment from the mobile application using a credit card, the credit card original number is converted into a token in the remote token service server and the token is used in the transaction. After verification of the token in the token service server, the final confirmation of payment is reflected in the credit card to which the token is mapped.

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

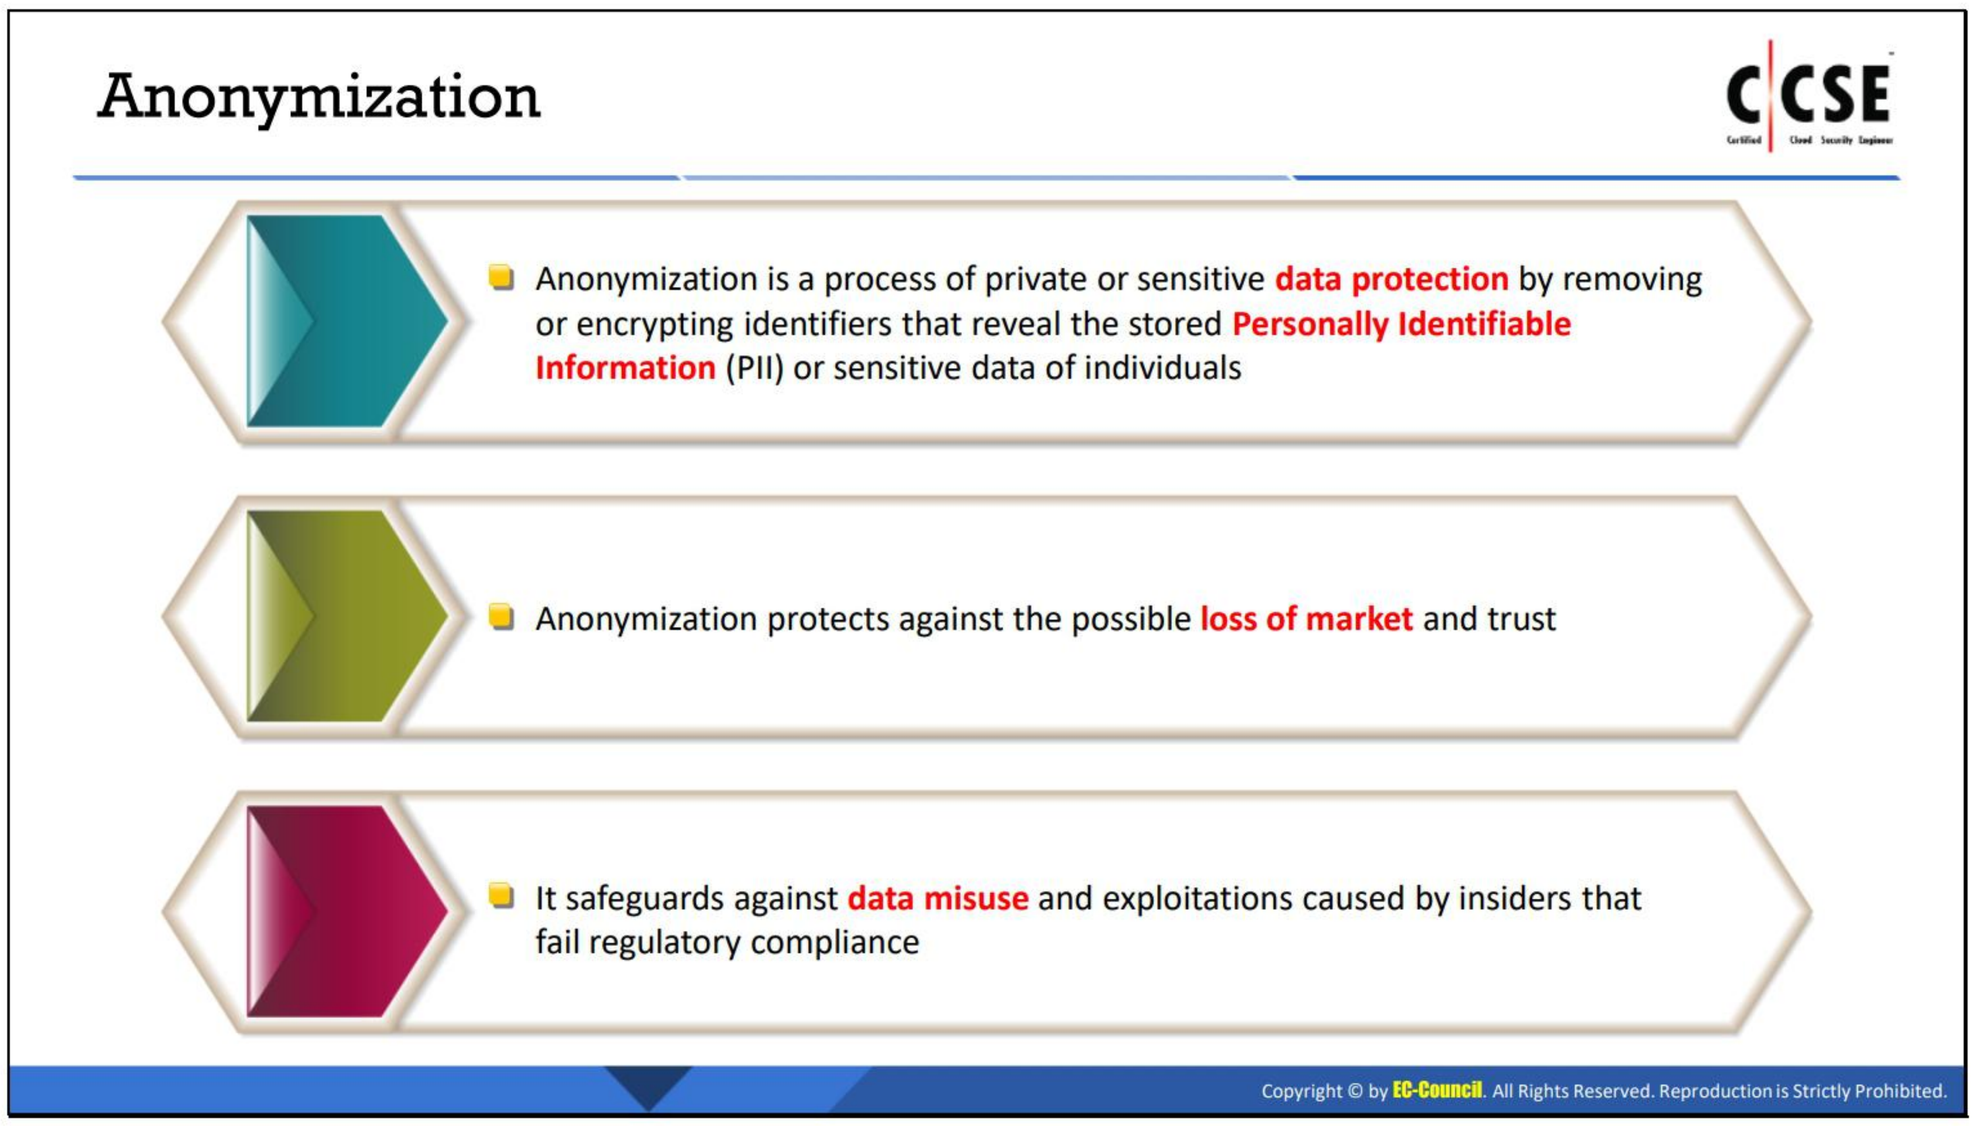
Tokens can be created in multiple ways:

- **Mathematically reversible cryptographic function with a key**: These reversible tokens can be mapped to single or multiple pieces of data and can be protected using strong cryptography.
- **A non-reversible function like a hash function**: A non-reversible or irreversible token is impossible to break down to get the original value from the token. These tokens can be either authenticatable or non-authenticatable. The irreversible authenticatable tokens are generated using a one-way function to verify a known used data whereas irreversible non-authenticatable tokens cannot be used for a specific data.
- **Index function or randomly generated number**: Random tokens can be generated using computer programming languages.

Tokenization protects data by removing original data and replaces it with a mechanism that provides access to required data. Tokenization complies with the laws and regulations that require sensitive data protection. It creates a more complex situation for a hacker who tries to access cardholder data (or respective data) and minimizes the risk of data breach.

hide01.ir

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



Anonymization

Anonymization is a process of private or sensitive data protection by removing or encrypting identifiers that reveal stored PII or sensitive data of individuals. PII such as name or addresses that map to an individual are direct identifiers, whereas socioeconomic information, events, and dates are indirect identifiers. Analysis with a combination of multiple indirect identifiers can be mapped to identify an individual.

Some of the data anonymizations techniques are described below.

- **Data masking:** Replacing the sensitive data with values or characters such as "*" or "X" for protection is called data masking.
- **Pseudonymization:** The process of replacing private identifiers with pseudonyms or fake identifiers is Pseudonymization. Consider a scenario where an identifier is replaced with "John Williams" from "Mark Smith". It provides data integrity and can protect data privacy. This technique can be used for testing, training, and analytics.
- **Generalization:** Generalization is the process of deletion or reduction of the data to make it less identifiable. For example, a house number in the address is removed but not an area or road number so that it cannot be mapped to an individual, and the data retained will be accurate for the intended purpose.
- **Data swapping:** Data swapping is also known as data permutation and shuffling; it is a technique used to rearrange the dataset attribute values. The original data can be protected using this process; for example, swapping of recognized values that may include the date of birth, age, gender, etc., which may have an impact on anonymization.

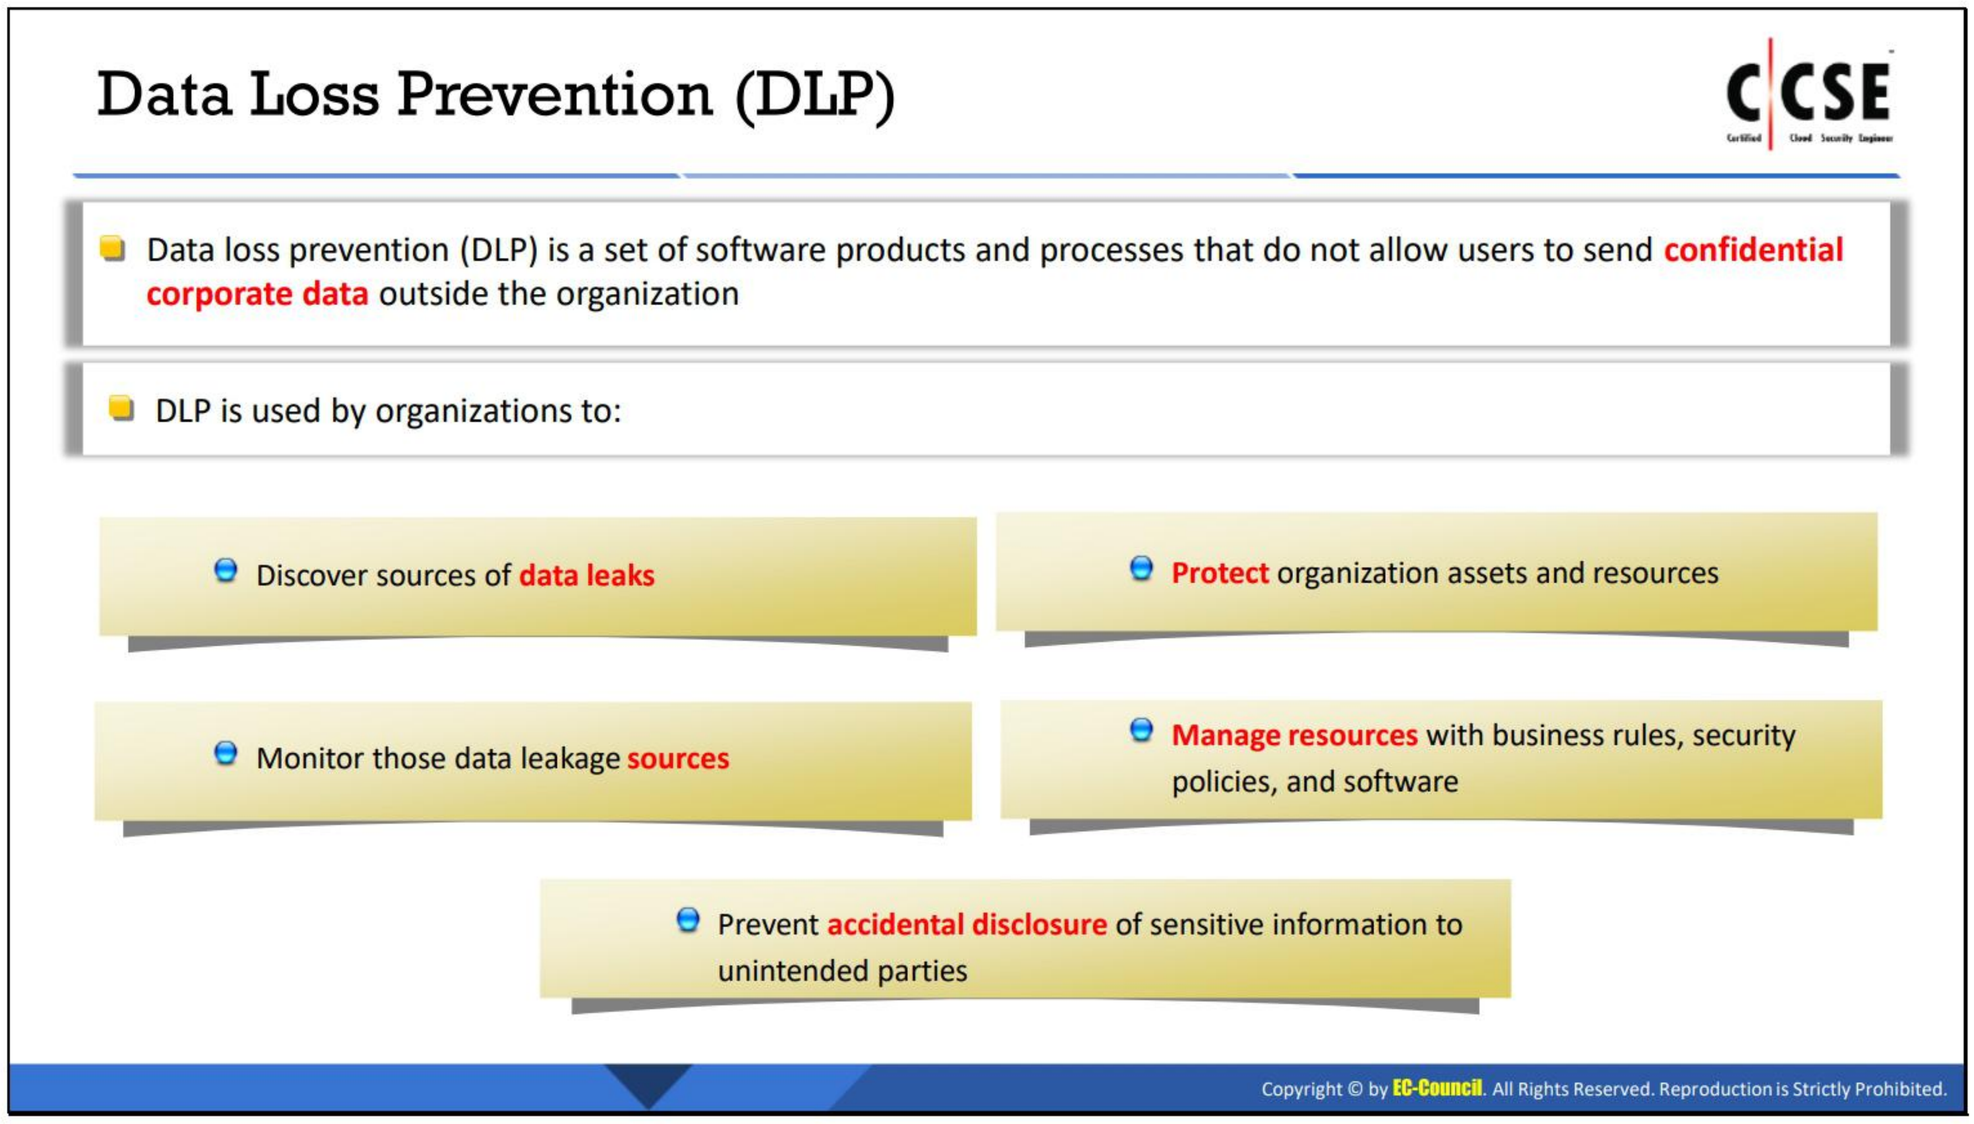
Certified Cloud Security Engineer Copyright © by **EC-Council**
All Rights Reserved. Reproduction is Strictly Prohibited.

- **Data perturbation:** The slight modification of the original dataset by using round number methods along with the addition of some random noise is called data perturbation. The set of values should be in proportion to the perturbation or disturbance. A small base should be used (e.g., 5) for rounding values of age and house number where the credence is not lost and is proportional to the original value; however, a base like 15 cannot be used for ages that do not look genuine but can be used for house numbers that look fine.
- **Synthetic data:** It is algorithmically generated and the generated data has no connection with the actual genuine information. Synthetic data can be used to create artificial datasets, instead of modifying the original dataset and risking privacy and security. Statistical models can be created based on the original dataset. Standard deviations, linear regressions, medians, or other statistical techniques are used to generate synthetic data.

Anonymization protects against the possible loss of market and trust. It can safeguard against data misuse and exploitations caused by insiders that fail regulatory compliance. It increases governance and also helps with result consistency.

hide01.ir

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



Data Loss Prevention (DLP)

DLP is a set of software products and processes that do not allow users to send confidential corporate data outside the organization. These software products help network administrators in controlling what data end-users can transfer. DLP rules block the transfer of any confidential information across external networks. They control any unauthorized access to company information and prevent anyone from sending malicious programs to the organization.

Implement DLP according to the organizational rules set by management. This prevents accidental/malicious data leaks and loss. If an employee tries to forward or even upload company data on cloud storage or even on a blog, the action will be denied by the system.

A DLP policy is adopted by the management when internal threats to a company are detected. DLP is a policy to ensure that none of its employees send sensitive information outside the organization. New emerging DLP tools not only prevent the loss of data but also monitor and control irregular activities from occurring on the system.

There are DLP products available that help administrators determine what data users transfer. DLP products are also known as data-leak prevention, information-loss prevention, or extrusion-prevention products.

DLP is used by organizations to

- Discover sources of data leaks
- Monitor those data leakage sources
- Protect organization assets and resources
- Prevent accidental disclosure of sensitive information to unintended parties

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

- Manage resources with business rules, security policies, and software

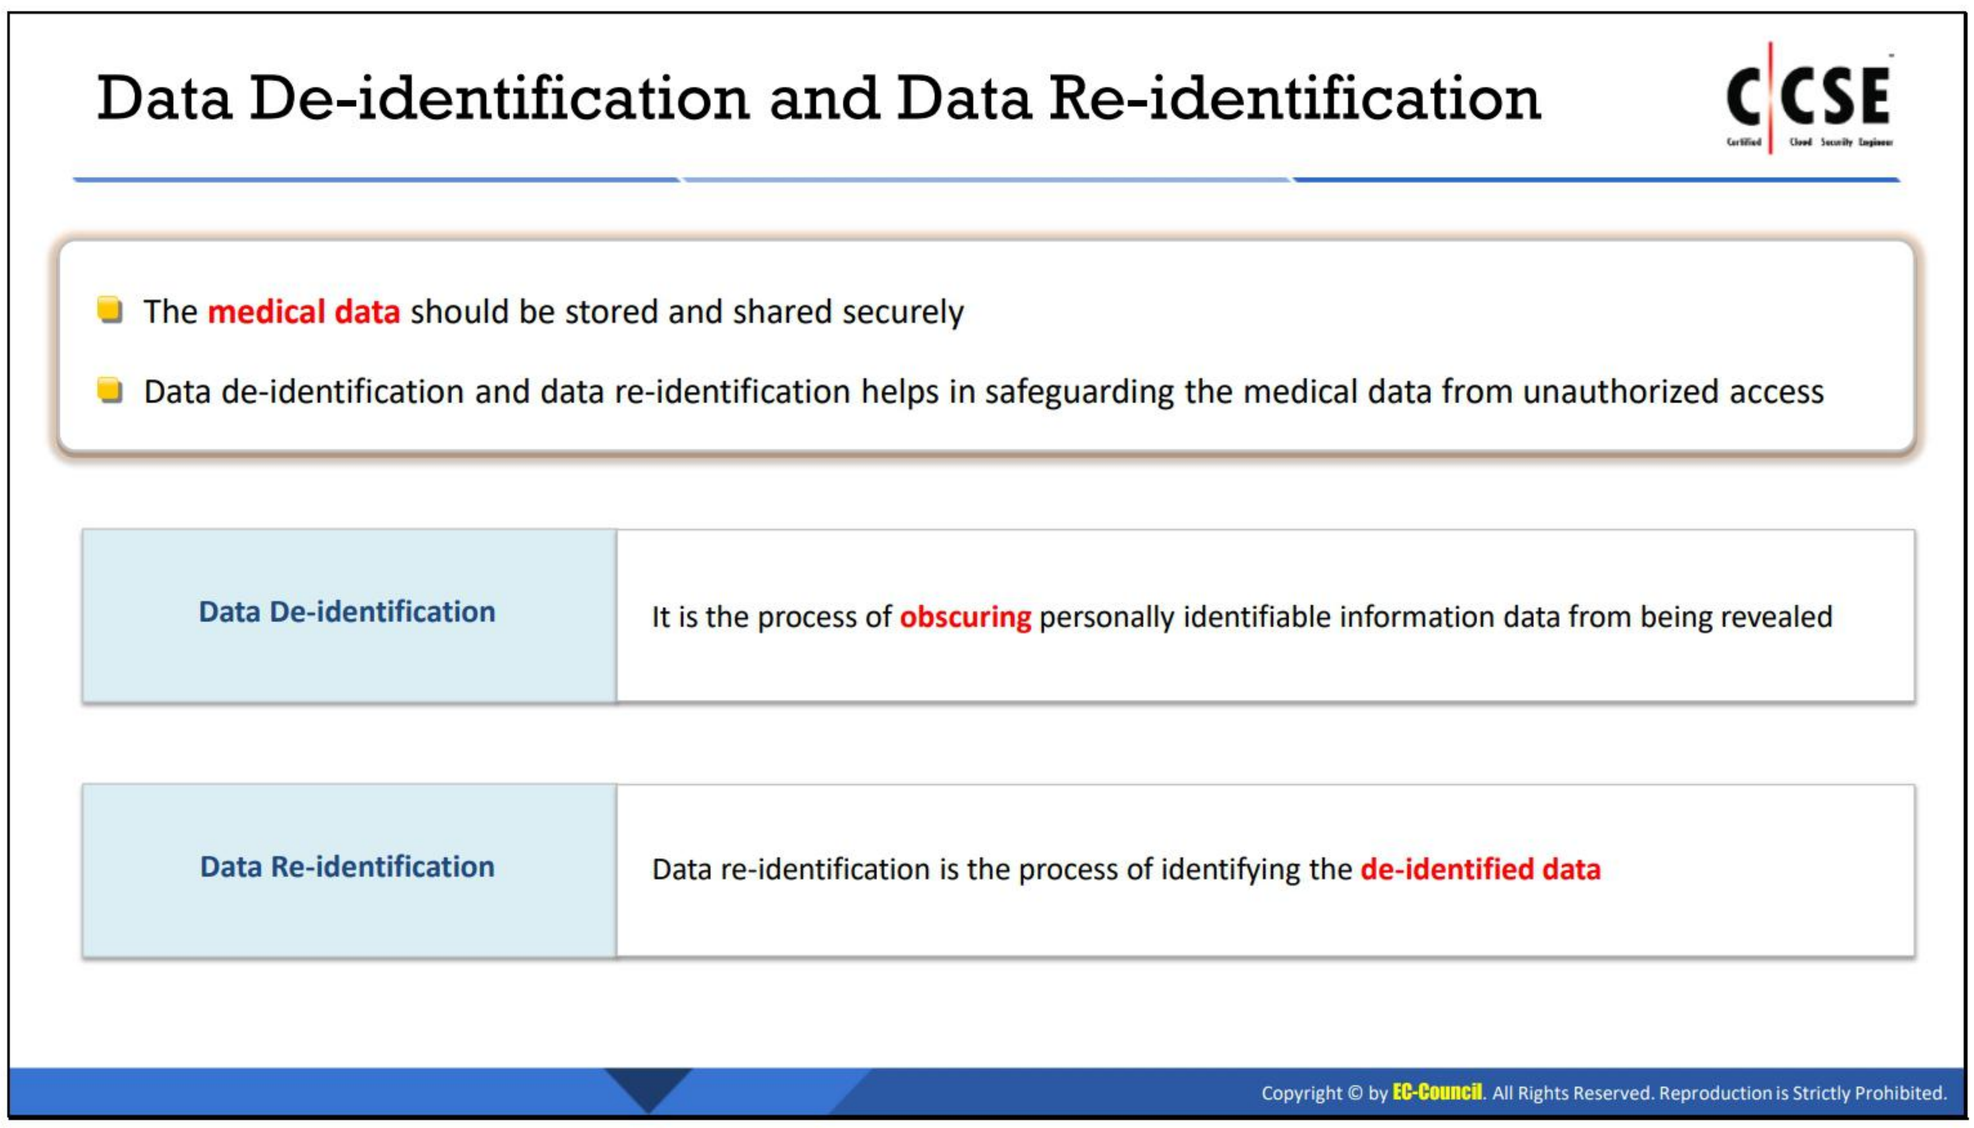
Types of Data Loss Prevention (DLP) Solutions

There are various types of DLP solutions that function differently but have the same objective, that is, to prevent data leakage.

- **Endpoint DLP:** Endpoint data-loss prevention is a solution that monitors and protects PC-based systems such as tablets, laptops, etc. It is used for preventing data leakage through clipboards, removable devices, and sharing applications. The solution includes an agent that monitors specific user operations such as sending an email, copying a file to removable media devices, printing a file, etc. Endpoint DLP protects data in use.
- **Network DLP:** Network DLP is a solution that monitors, protects, and reports all data-in-transit. It is installed at the "perimeter" of organization networks. It helps network administrators in scanning all data moving through ports and protocols of the organization. It may analyze email traffic, social media interactions, SSL traffic, instant messaging, etc. The solution maintains reports containing information such as what data is used, who is using the data, and where the data is sent. Thus, it helps in controlling the flow of data over the organization's networks and meets regulatory compliance. Data collected by the network DLP is stored in a database for later retrieval.
- **Storage DLP:** Storage DLP is a solution that monitors and protects data-at-rest, that is, data stored in the organization's data center infrastructure (file servers, SharePoint, and databases). It identifies the location where sensitive information is stored and helps users in determining whether it is stored securely. It allows authorized users to view and share sensitive files in the organization network.

hiteoi.ir

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



Data De-identification and Data Re-identification

Medical data should be stored and shared securely. Data de-identification and data re-identification helps in safeguarding the medical data from unauthorized access.

hjite01.ir

- **Data De-identification:** Data de-identification is the process of obscuring PII data from being revealed. The separation of PII data and PHI data it stores in the system, makes an organization becomes compliant with Health Insurance Portability and Accountability Act of 1996 (HIPAA) regulations. After separation, the PHI data can be stored anywhere but the PII data must be stored securely and encrypted. Additional care should be taken on how you de-identify data.
- Administrators after getting new data from users can decide the fields which are identifiable and which are not such as first name, last name, age, etc. An administrator can separate this data and send it to third-party service, which stores them securely and provides an ID for this data. Now, the administrator saves this ID along with other de-identified data into the organization's servers.

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

- **Data Re-identification:** It is the process of identifying the de-identified data. When the administrator wants to get back the data, the administrator loads data from third-party in which the data is stored in the form of ID, then merges with data in the organization's server, and provides details.

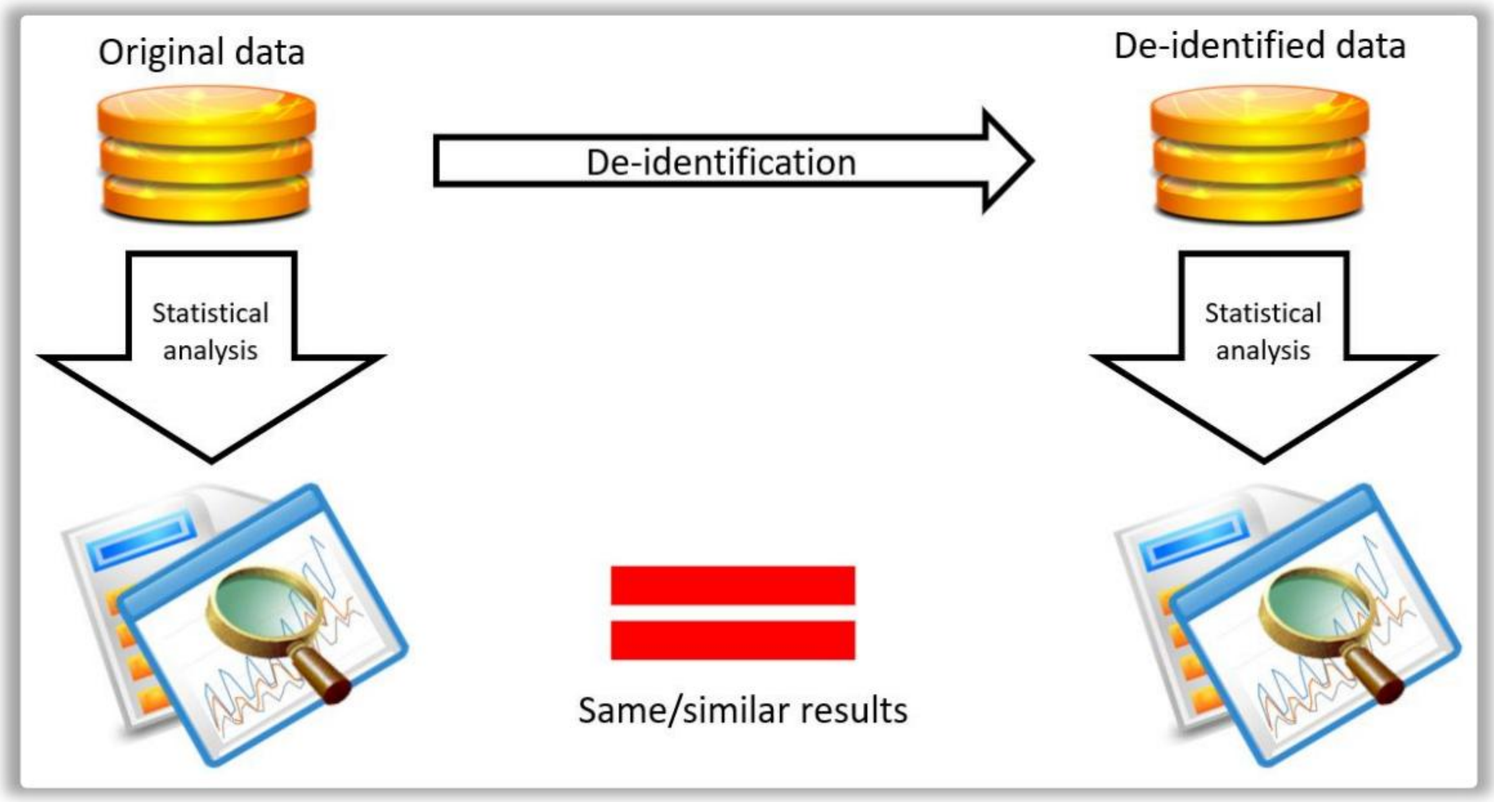
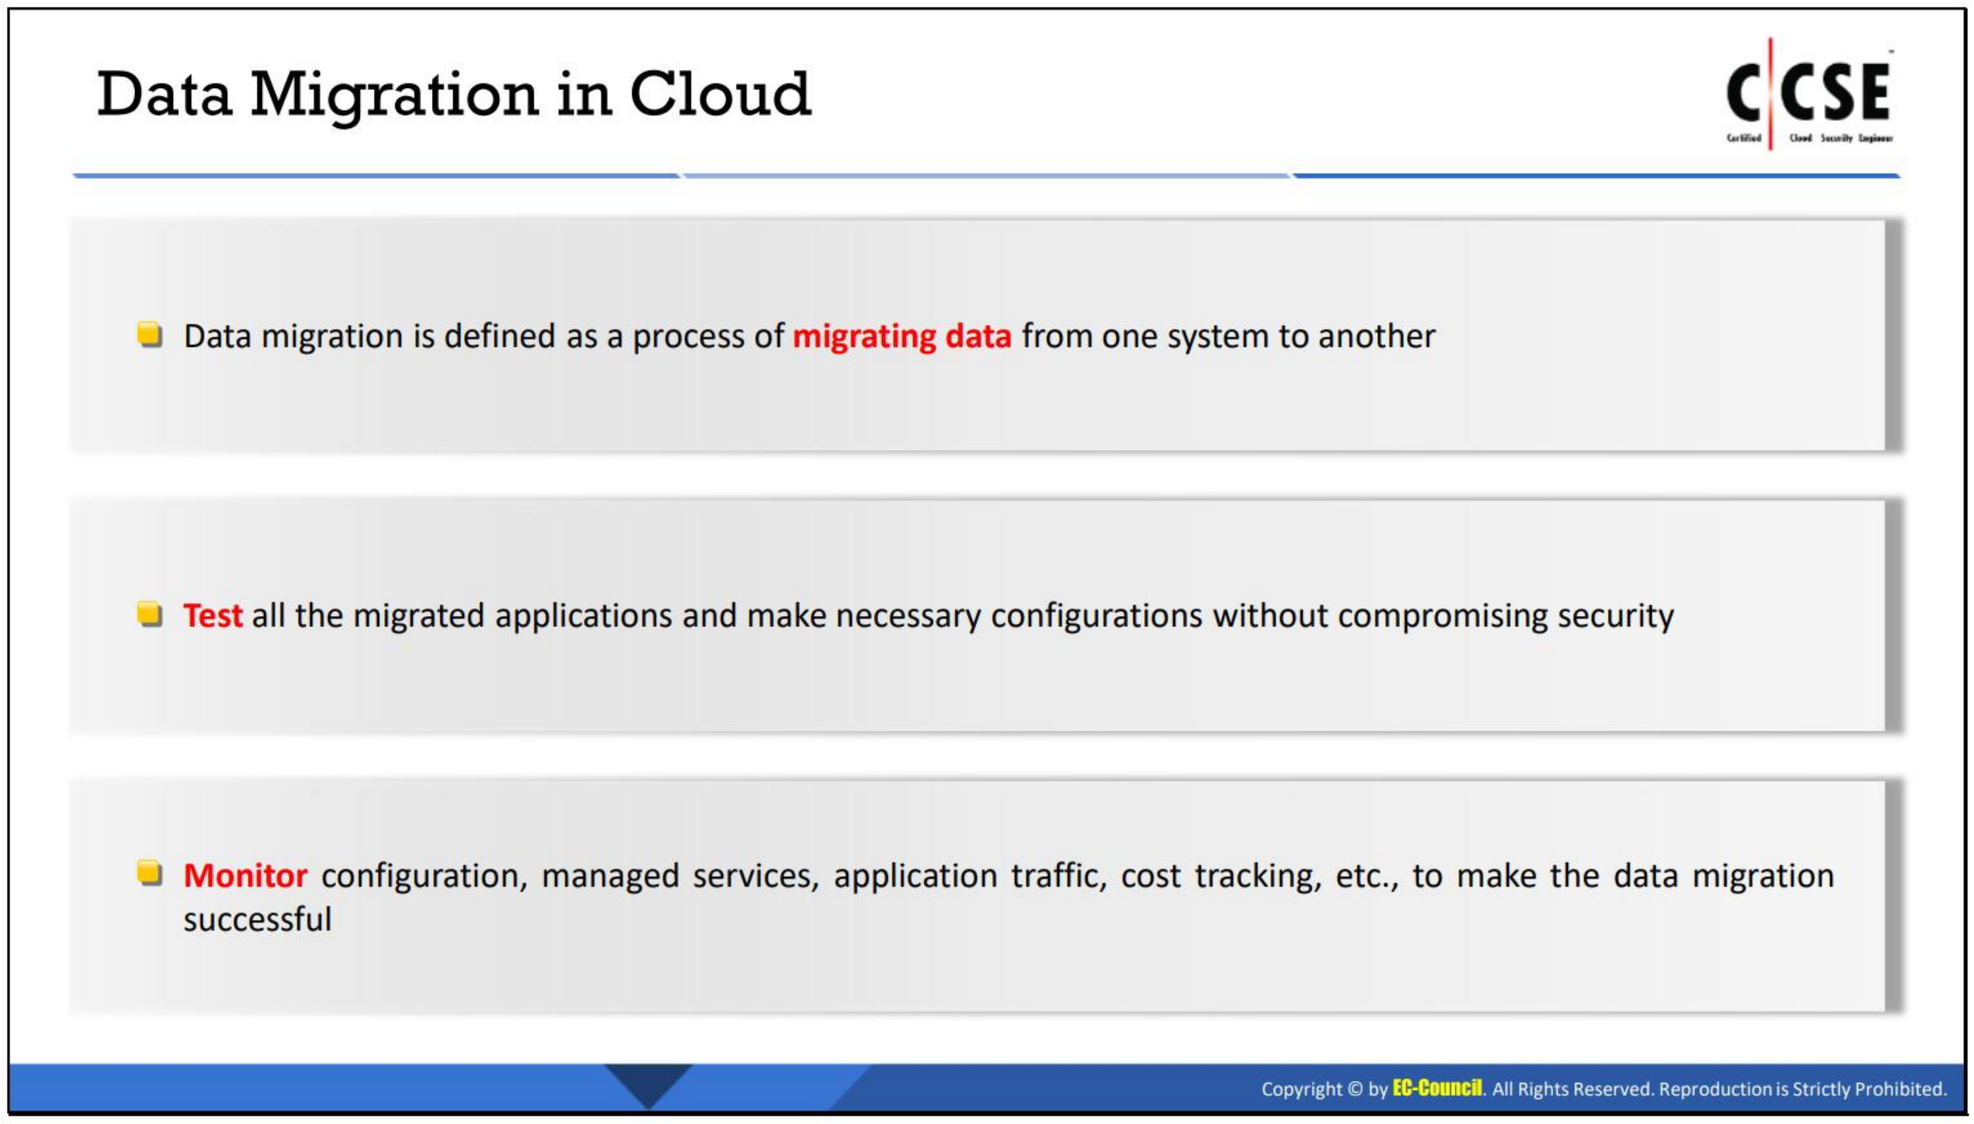


Figure 4.41: Data Re-identification

The original data and PII data are separated using some statistical analysis. Then, the de-identified data is sent to secure HIPAA compliant servers. De-identification and re-identification process are not limited to medical records but also educational records, biospecimen records, etc.

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



Data Migration in Cloud

Data migration is defined as a process of migrating data from one system to another. The strategy for data migration in the cloud environment depends on the complexity of the IT environment. During data migration in cloud, the following recommendations should be considered.

- **Migration planning phase:** Any organization before migrating data to the cloud should assess the current environment. Based on the current resources, it can be easy to assess the cloud server requirements, which avoids excess resources to reduce financial overload. An application performance management (APM) solution, which provides real-time visibility into the environment while migration, should be used.
- **Cloud environment selection:** The migration planning phase provides a clear idea of the requirements. The organization can decide on cloud models such as public, private, or hybrid cloud, based on current or future requirements. APM solution should be able to extend the selected cloud provider.
- **Apps and data migration:** While migrating applications and data, cloud security concerns should be considered such as complying with security policies as well as backup and restoration planning. APM solution should check premigration performance.
- **Validation after migration:** Cloud migration is not successful until the migrated applications work as expected and face no issues in data access. It is recommended to test all the migrated applications and make necessary configurations without compromising security.

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

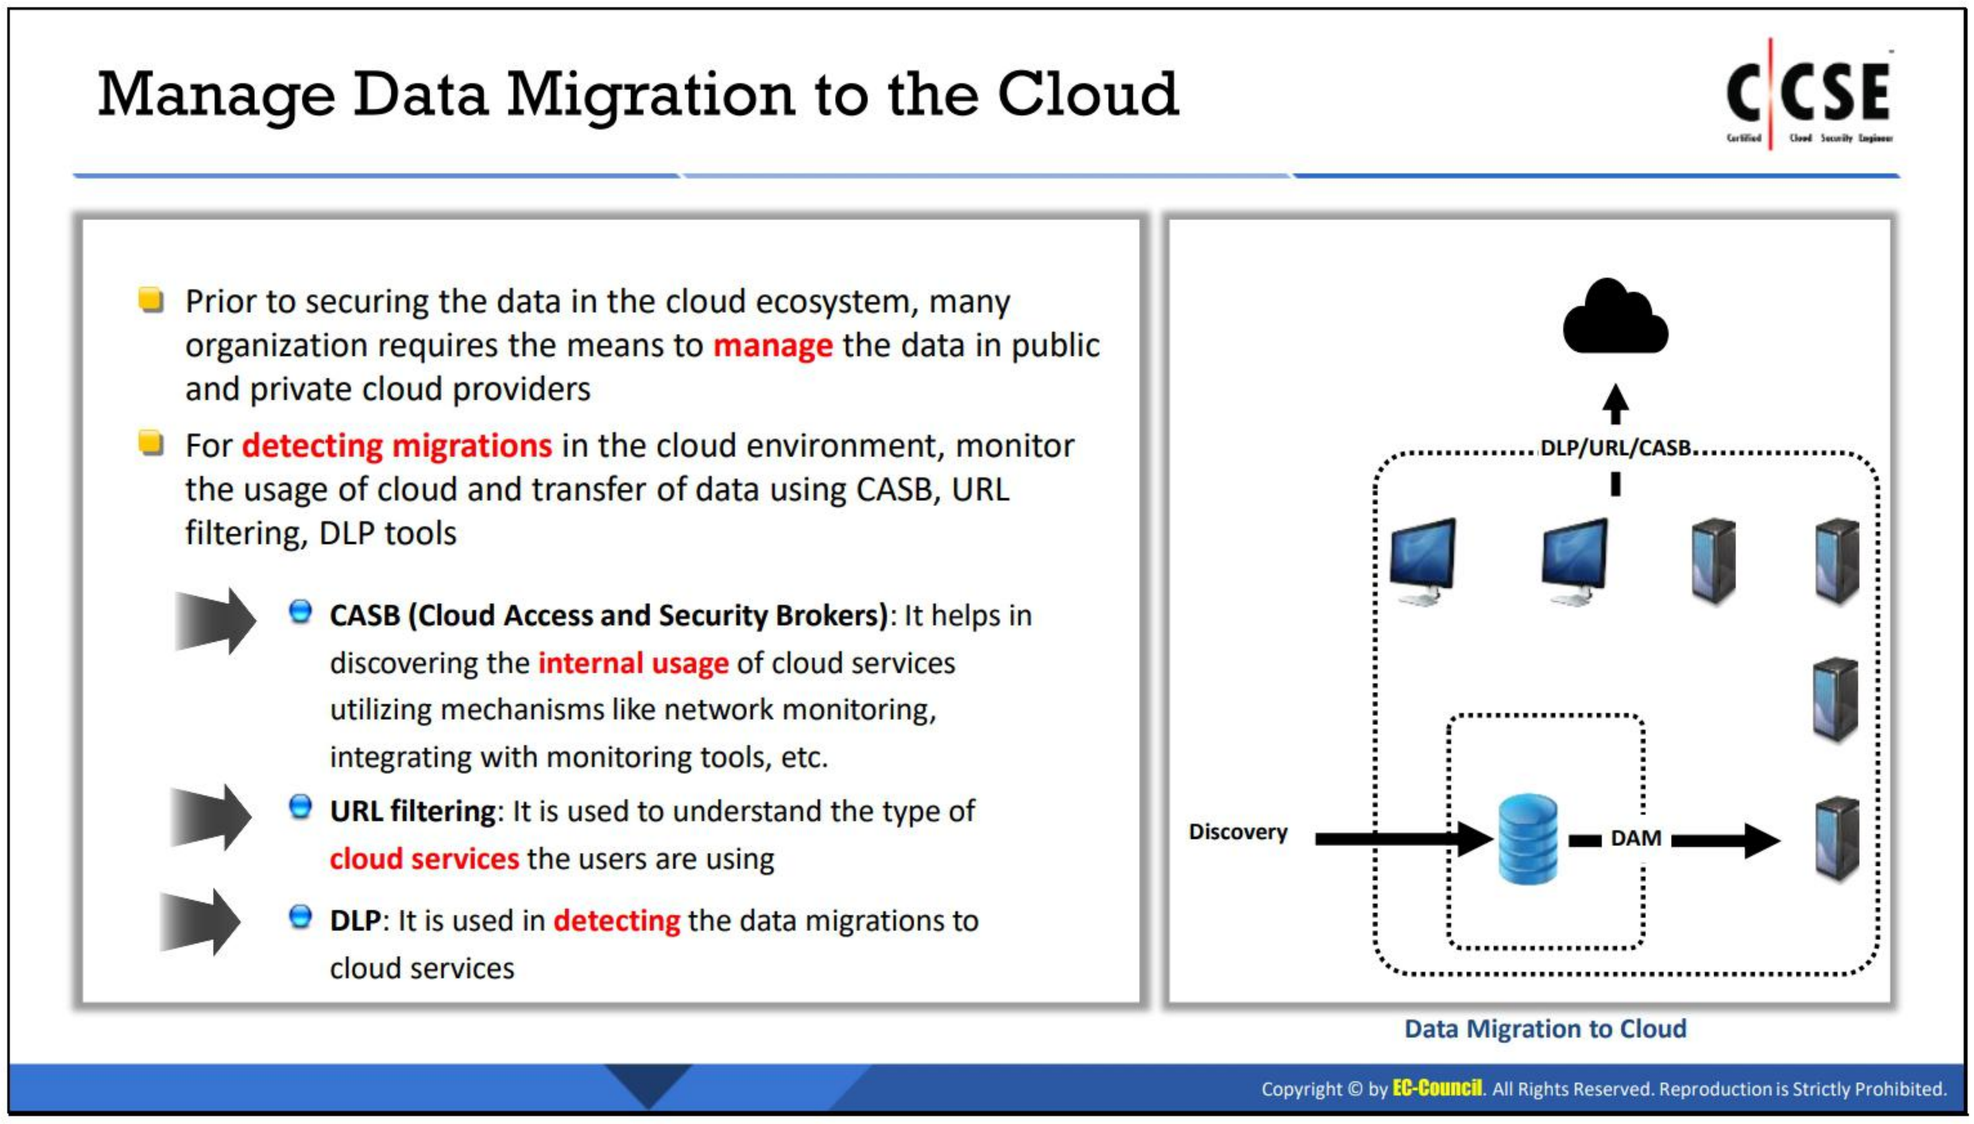
Monitoring

Monitoring reduces the risk of migration failures. All services and infrastructure devices need to be monitored to maintain logs that allow finding root cause analysis. During migration redirections to network traffic may be setup, and certain on-premises resources may be shut down or turned on. Various areas should be monitored to make the migration successful.

- **Configuration:** Initiating a service in a cloud environment requires a network, security, and audit configurations to be properly configured and implemented to make the service work as expected. The access of the application to end-users or clients should be continuously validated. The response time of the configured services should be monitored.
- **Billing:** As the cloud services provide pay as you use services. Usage of resources and activation of the premium services without proper knowledge may cost a lot. Therefore, it is recommended to monitor the migration process and close or remove resources that are not required, as they may generate more billing.
- **Monitoring of managed services:** After the completion of migration, the components that are not under the R&D team's control should be kept under monitoring to analyze the functionality and performance.
- **Application Traffic:** Application traffic tracking is the best to confirm the proper configuration and availability of the applications. It should be verified that during and after migration all the connection requests should be at the new cloud servers. The previous deployment form and the current response times should be similar to the premigration response times without crossing thresholds; only then response rate is maintained. Web server, load balancer, and other logs are used to verify this.
- **Cost tracking:** It is recommended to track the cost during and after the migration process because cost-saving may be the primary reason for organizations to move to the cloud. Any cloud services with improper usage may lead to high billings. Cloud service provider billing systems should be integrated and should be allowed to monitor system visualization. With the use of the monitoring feature, any incorrect billing can be brought to notice.

hiteo.ir

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



Manage Data Migration to the Cloud

Data migration refers to moving of data from one location to another. Cloud data migration refers to transferring of business processes, data, and applications from on-premises to a cloud or from one CSP to another. When data is migrated to the cloud, it becomes essential to ensure how secure it will be after migration.

Prior to securing the data in the cloud ecosystem, many organizations require the means to manage data in public and private cloud platforms. Managing and monitoring data properly in cloud platforms is an essential requirement not only for data security but also for compliance.

Initially, define policies for the allowed data types and where they are allowed and then link these policies to the baseline security requirements. For example, PII is allowed on some services assuming that they meet specific encryption and access control requirements.

Then, figure out the key data repositories and monitor these repositories for large activity/migrations via tools such as File Activity Monitoring and Database Activity Monitoring. This serves as a key data security control for identifying all kinds of major breaches.

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

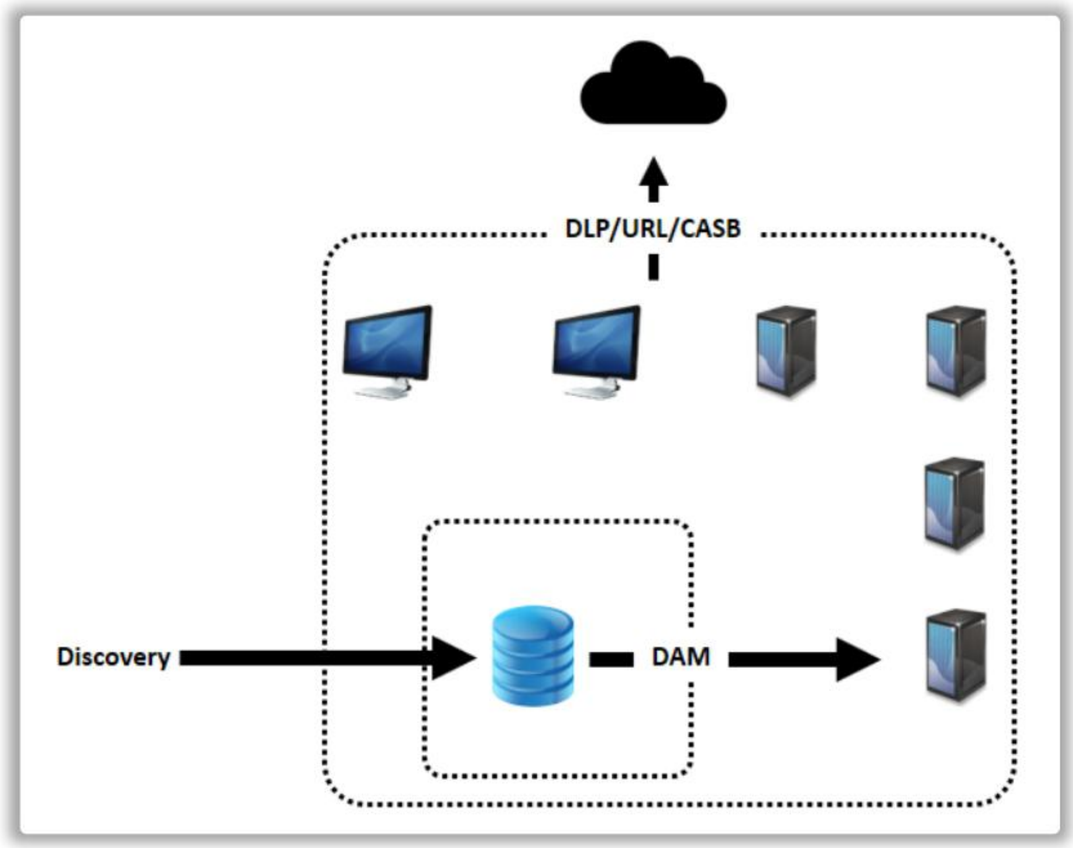
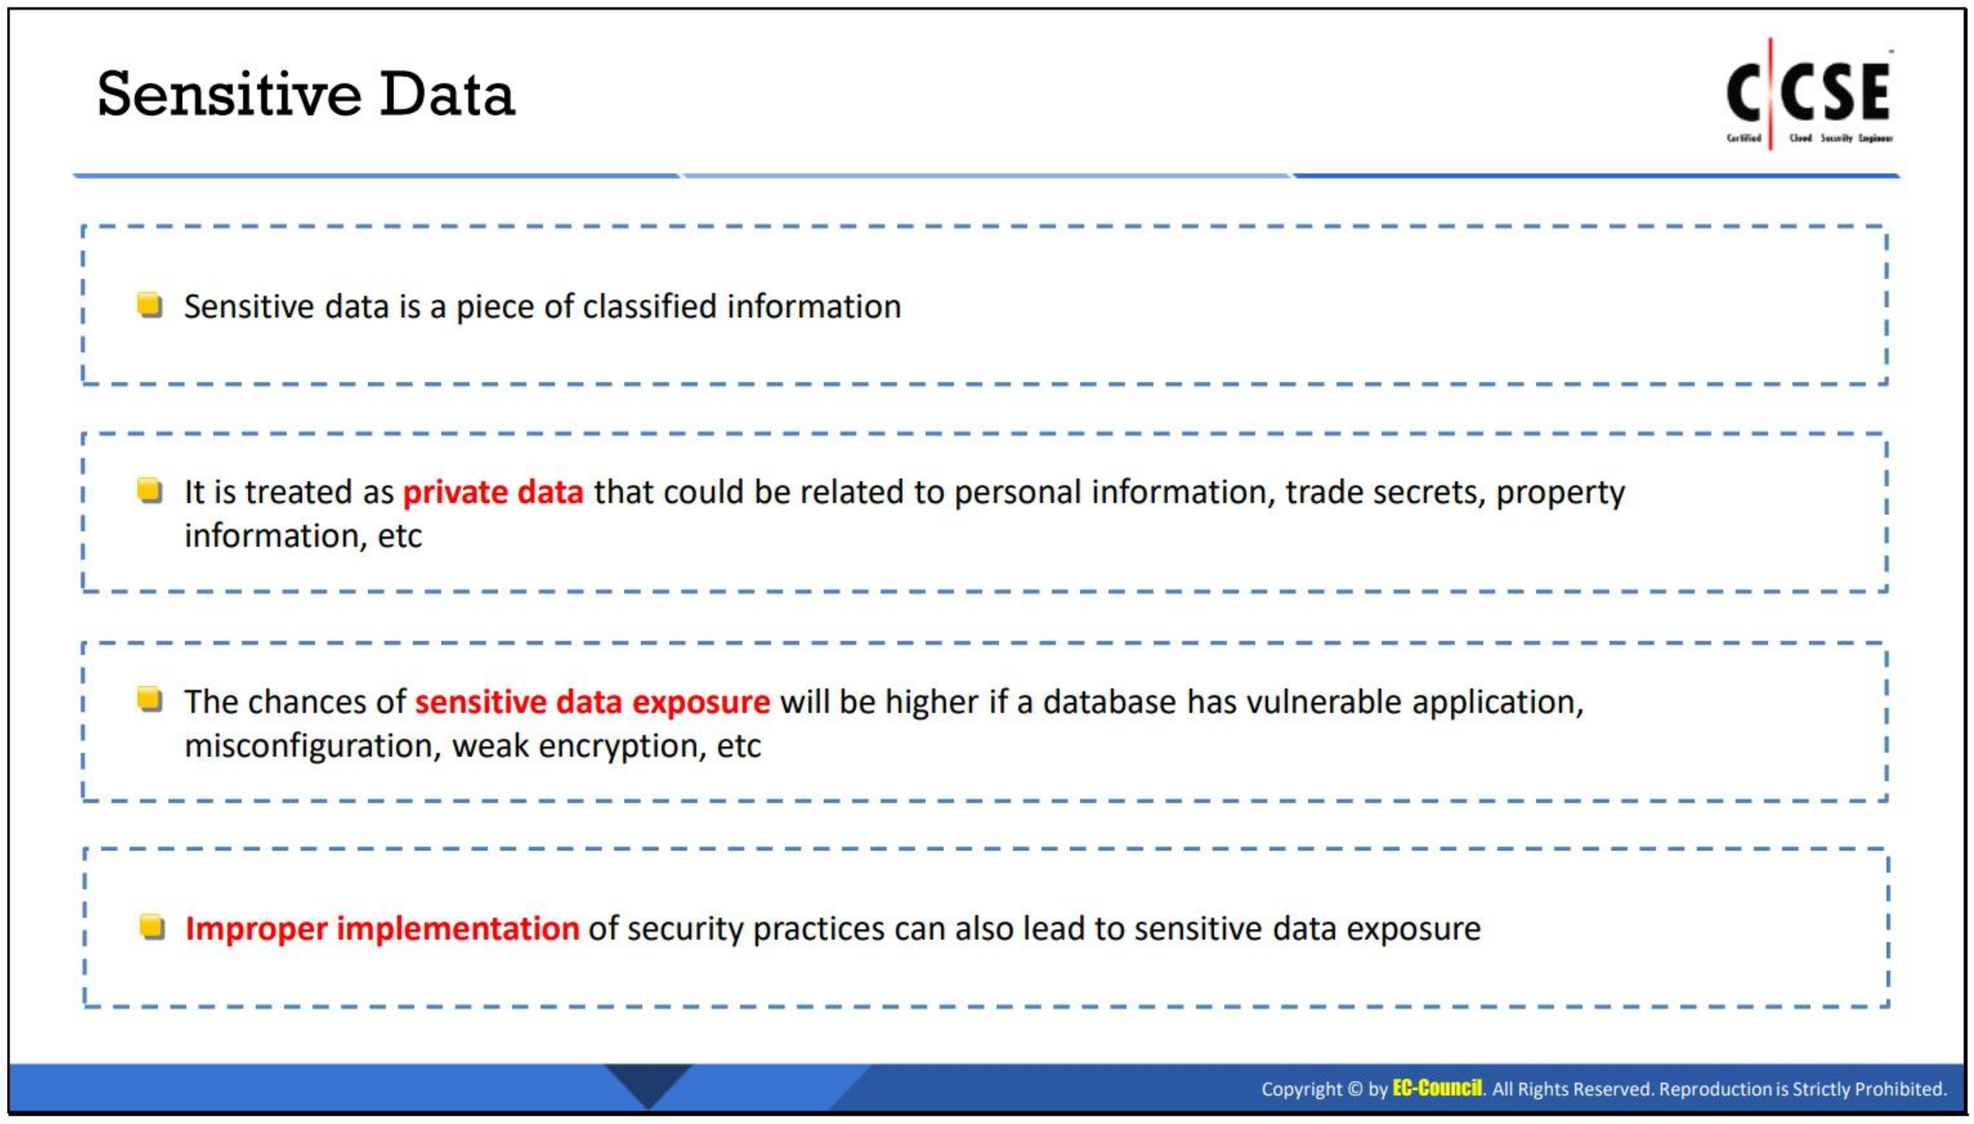


Figure 4.42: Managing Data Migration to the Cloud

For detecting migrations in the cloud environment, monitor the usage of cloud and transfer of data using CASB, URL filtering, and DLP tools.

- **Cloud Access and Security Broker (CASB)**: CASB is a security software (both on-premises or cloud-based) that is used between CSPs and cloud service users to enforce compliance, governance, and security policies for cloud applications. It helps in discovering the internal usage of cloud services by utilizing mechanisms such as network monitoring, integrating with existing monitoring tools, etc. Once the services to which users connect the most are identified, CASBs provide monitoring for those services via API connections or inline interception. Most CASBs provide DLP, security alerting, and controls to better manage sensitive data usage in cloud services.
- **URL filtering**: URL filtering is a technology that controls the ability of accessing certain web pages or URLs. Businesses often use URL filtering to prevent their users or employees from accessing restricted or inappropriate web pages. It is used to understand the type of cloud services the users are using.
- **DLP**: Cloud DLP helps customers to better understand and handle PII and sensitive data to meet their specific compliance requirements. It is used for detecting data migrations to cloud services when monitoring web traffic. Cloud DLP can scan cloud storage buckets, objects, and folders for sensitive information in order to be compliant with regulations and ensure the safety of data.

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



Sensitive Data

Sensitive data is classified information. Only authorized individuals are allowed to access the sensitive data. Thus, sensitive data should be safeguarded from unauthorized access. Sensitive data is treated as private data that could be related to personal information, trade secrets, or property information. The organization should protect this information for fulfilling ethical or legal requirements. The loss of this data brings down the reputation of the organization and the customers also lose their trust on the organization.

Sensitive data includes the following:

- **Personally identifiable information (PII)**: The information containing data that reveals a person's identity with direct or indirect identifiers is PII. Example: Contact number, full name, date of birth, address, etc.
- **Protected health information (PHI)**: PHI is regulated by HIPAA. It includes individually identified health information and may include the records of mental, physical, or any medical condition of an individual.
- **Educational records**: It contains direct information related to students that are maintained by the university or the one acting for the university. Family Educational Rights and Privacy Act of 1974 (FERPA) governs the access of educational data by employers, foreign government, etc.
- **Cardholder data**: It contains the details of credit, debit, or any other payment cards. It is governed by PCI-DSS. Cardholder data may contain the full length of the credit card number, expiry, CVV of the card, and full name of the cardholder.

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

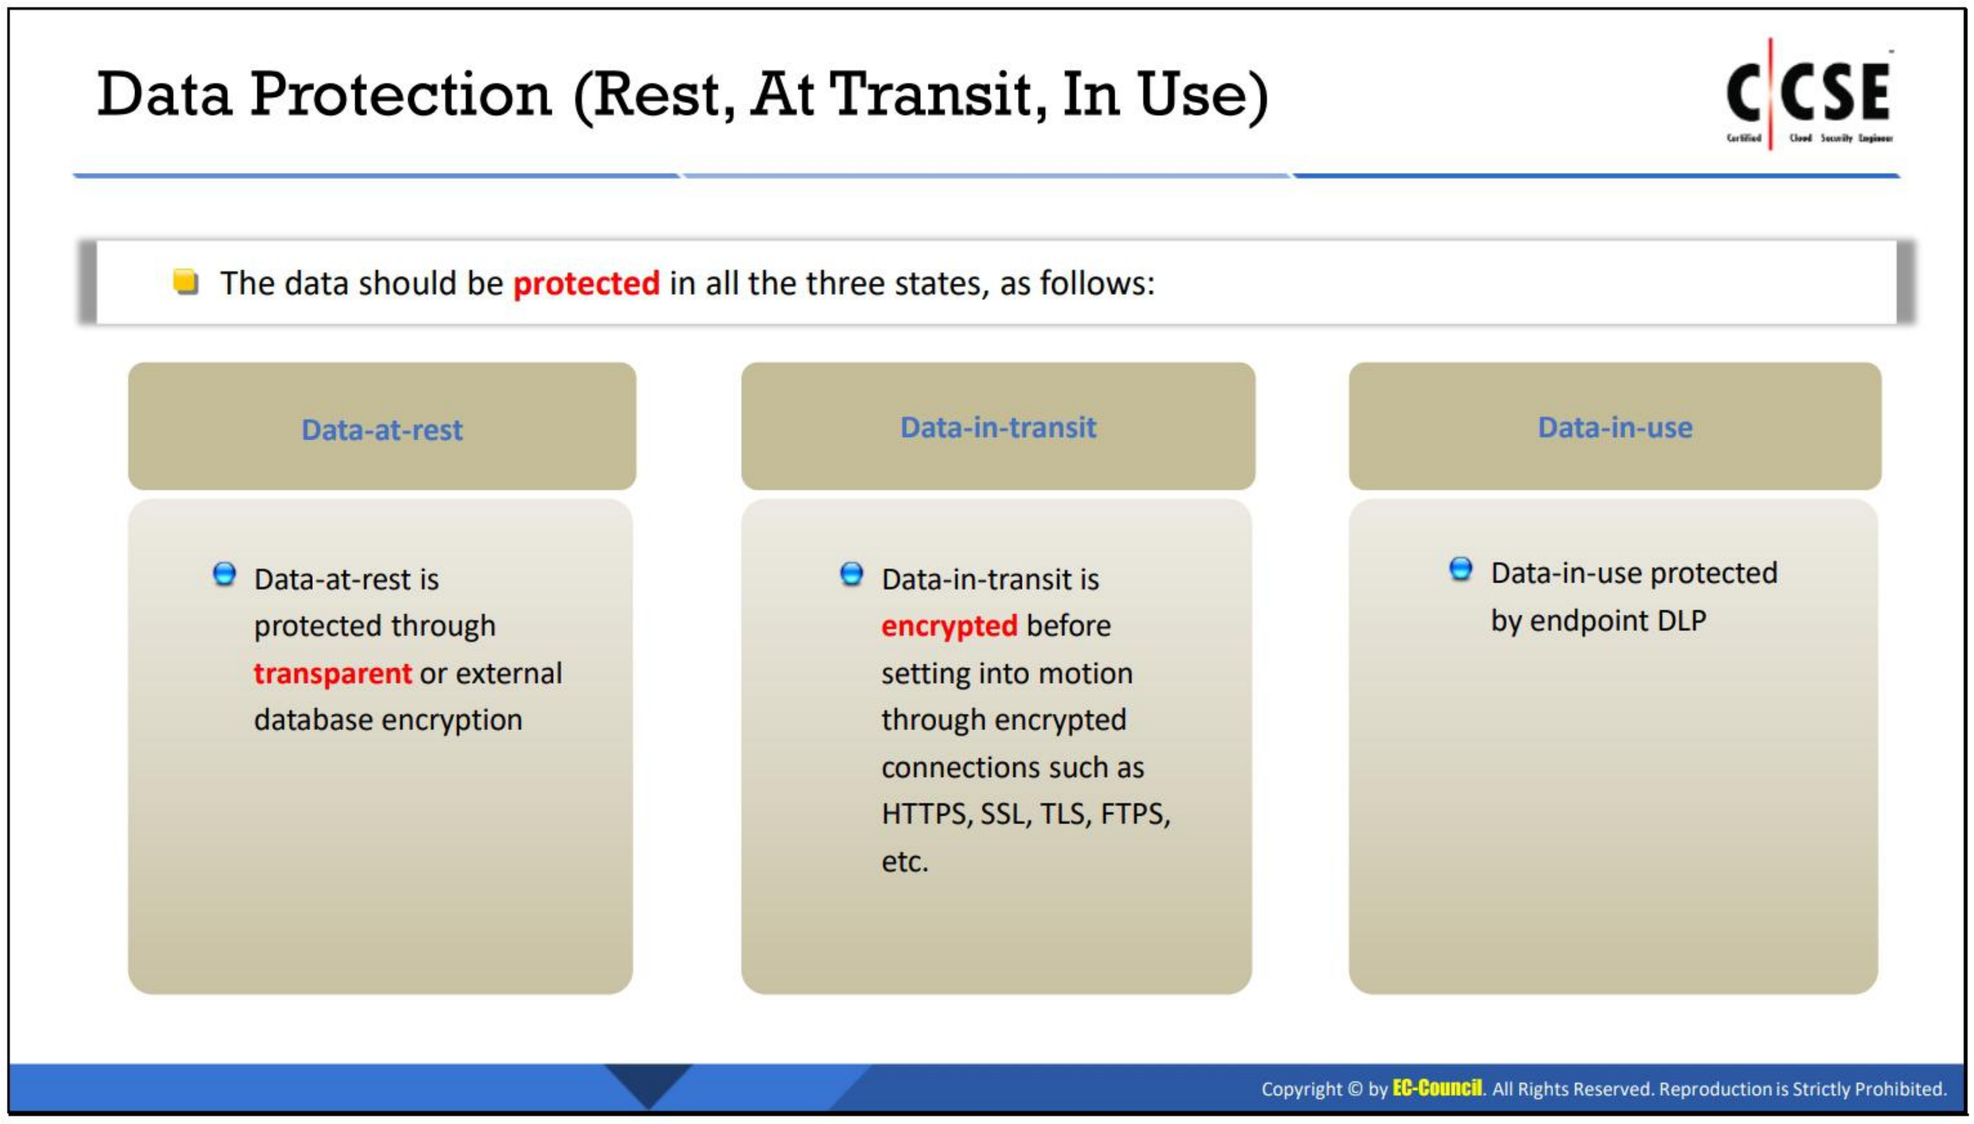
- **Sensitive data exposure**: The chances of sensitive data exposure will be higher if a database has vulnerable application, misconfiguration, weak encryption, etc. Sensitive data exposure can also happen due to improper implementation of security practices.

Secure practices to avoid sensitive data exposure are listed below:

- Use complex passwords and the same password must not be used for another account. It might be difficult for individuals/administrators to remember all of them; therefore, it is recommended to use secure and reputed third-party key vaults.
- Passwords of financial or any other personal accounts should be changed regularly and passwords should not match the previously used 5 to 10 passwords depending on the type of account.
- Make sure to enable text and email notifications for financial accounts if any event takes place.
- If any of your data is previously compromised and a malicious attacker tries to open a credit card or loan against an individual, these may be found in a credit report. Thus, regular verification of credit report helps in identifying and securing financial frauds.
- Use only HTTPS when performing actions related to your sensitive data.
- Use antivirus to identify malicious behavior.

hide01.ir

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

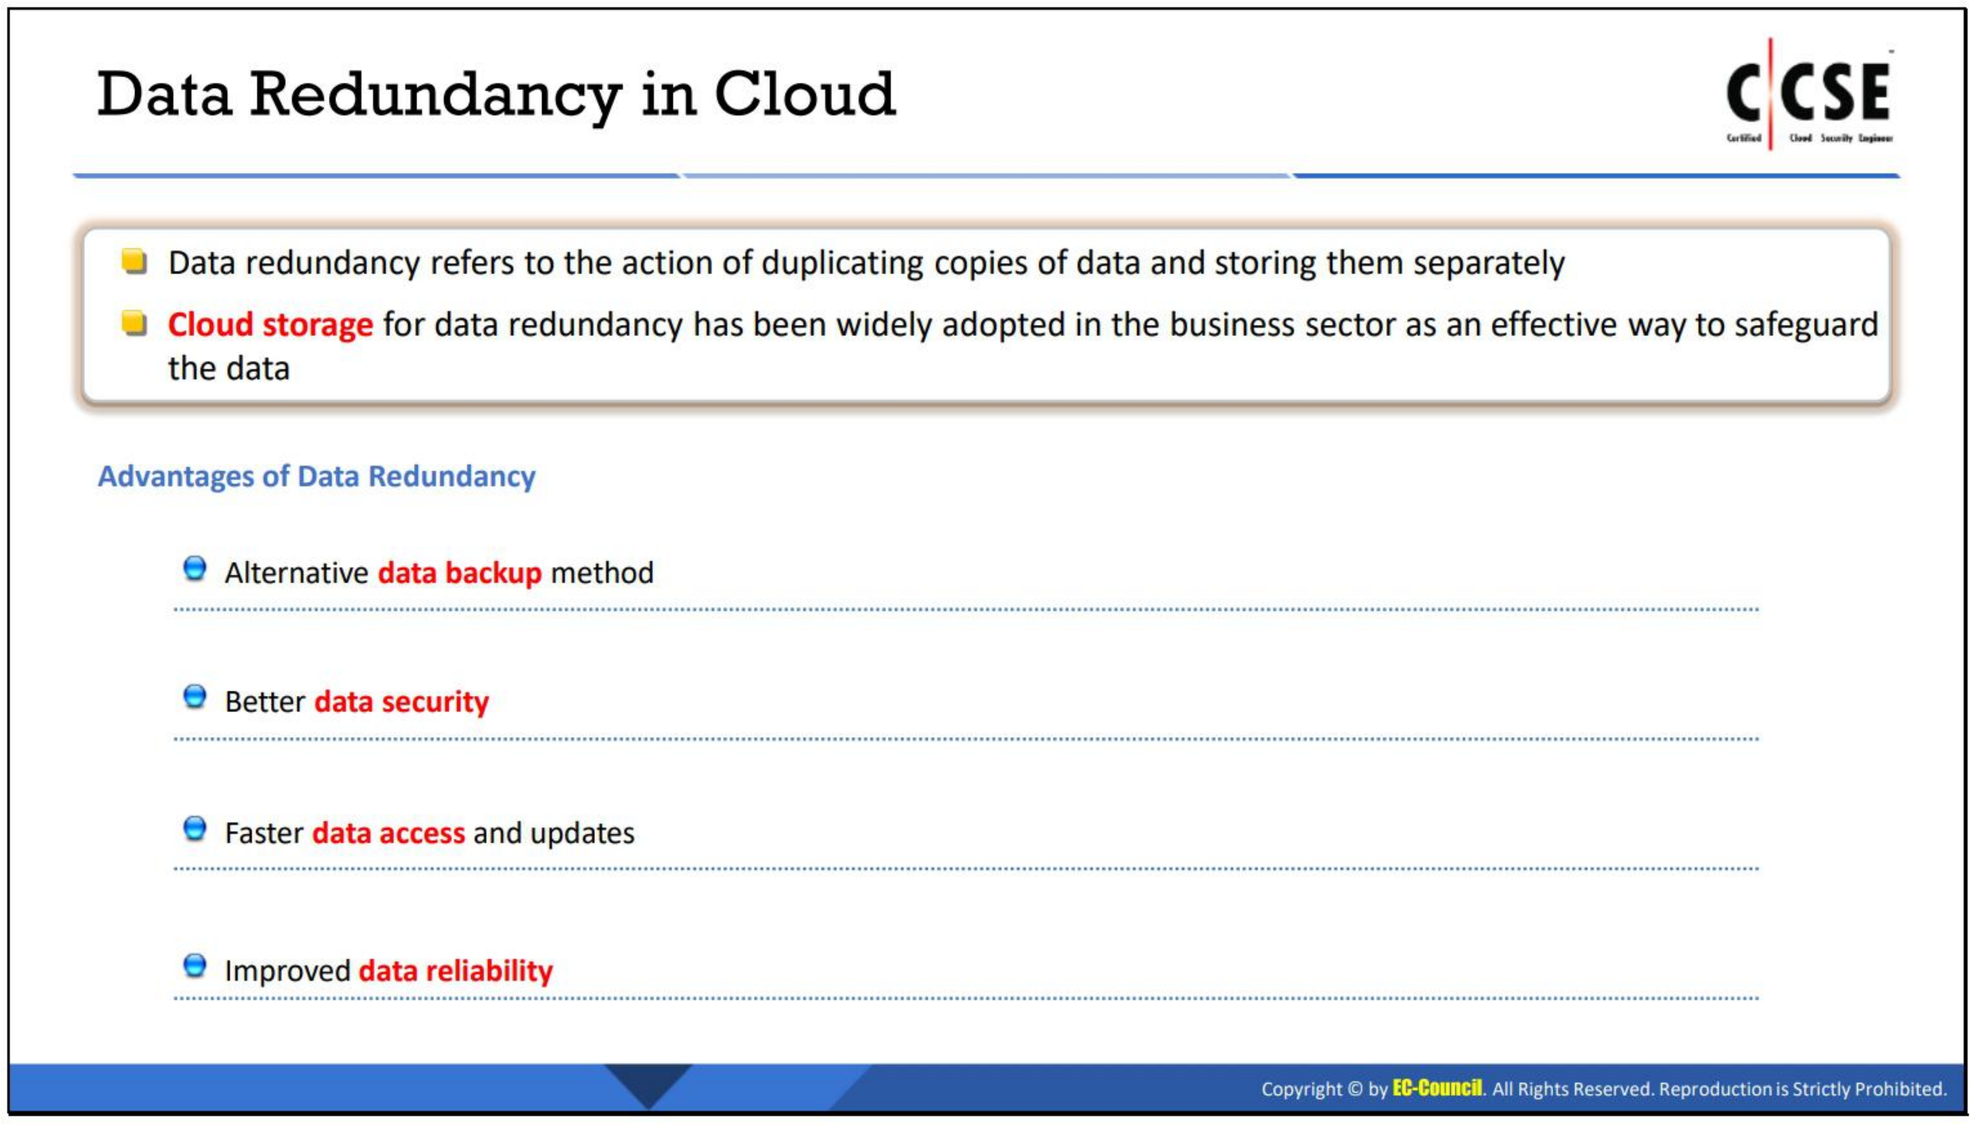


Data Protection (Rest, At Transit, In Use)

The data should be protected in all the three states, as follows:

- **Data-at-rest:** It is the inactive data that is stored on a device or a backup medium such as hard drives, laptops, backup tapes, mobile devices, or in the offsite cloud backup. Data-at-rest remains in a stable state. The data in rest will not move actively in a system or network and cannot be accessed by an application or program. Data-at-rest is protected through transparent or external database encryption.
- **Data-in-transit:** The data-in-transit is the data that is actively moving from one location to another across the network. The data is encrypted before setting into motion and/or transmission through encrypted connections such as HTTPS, SSL, TLS, FTPS, etc.
- **Data in use:** The data in use is the data that is stored or processed by RAM, CPUs, or databases. It is not passively stored on the system but actively moves across IT infrastructure. It is updated, erased, processed, accessed, or read by the system. Endpoint DLP protects data in use.

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



Data Redundancy in Cloud

Data redundancy refers to the action of duplicating copies of data and storing them separately so that if the cloud service becomes unavailable or corrupted, the cloud consumer would have instantaneous secure access to the backed-up version of their data. The cloud storage provider stores the data of the cloud consumer on multiple machines, and in compliance with any regulations that applied on the cloud consumer company (e.g., some organization needs three copies of their data to be stored at different locations)

Cloud storage for data redundancy has been widely adopted in the business sector as an effective way to safeguard the data. With cloud storage, an organization can collect and archive important documents, databases, and other necessary data offsite. An organization can utilize the tools from their cloud service provider to immediately recover the lost data, which makes cloud storage a better alternative against tape and disk-based backup systems. A secure and robust data solution for data redundancy is provided by cloud-based backups.

Advantage of Data Redundancy

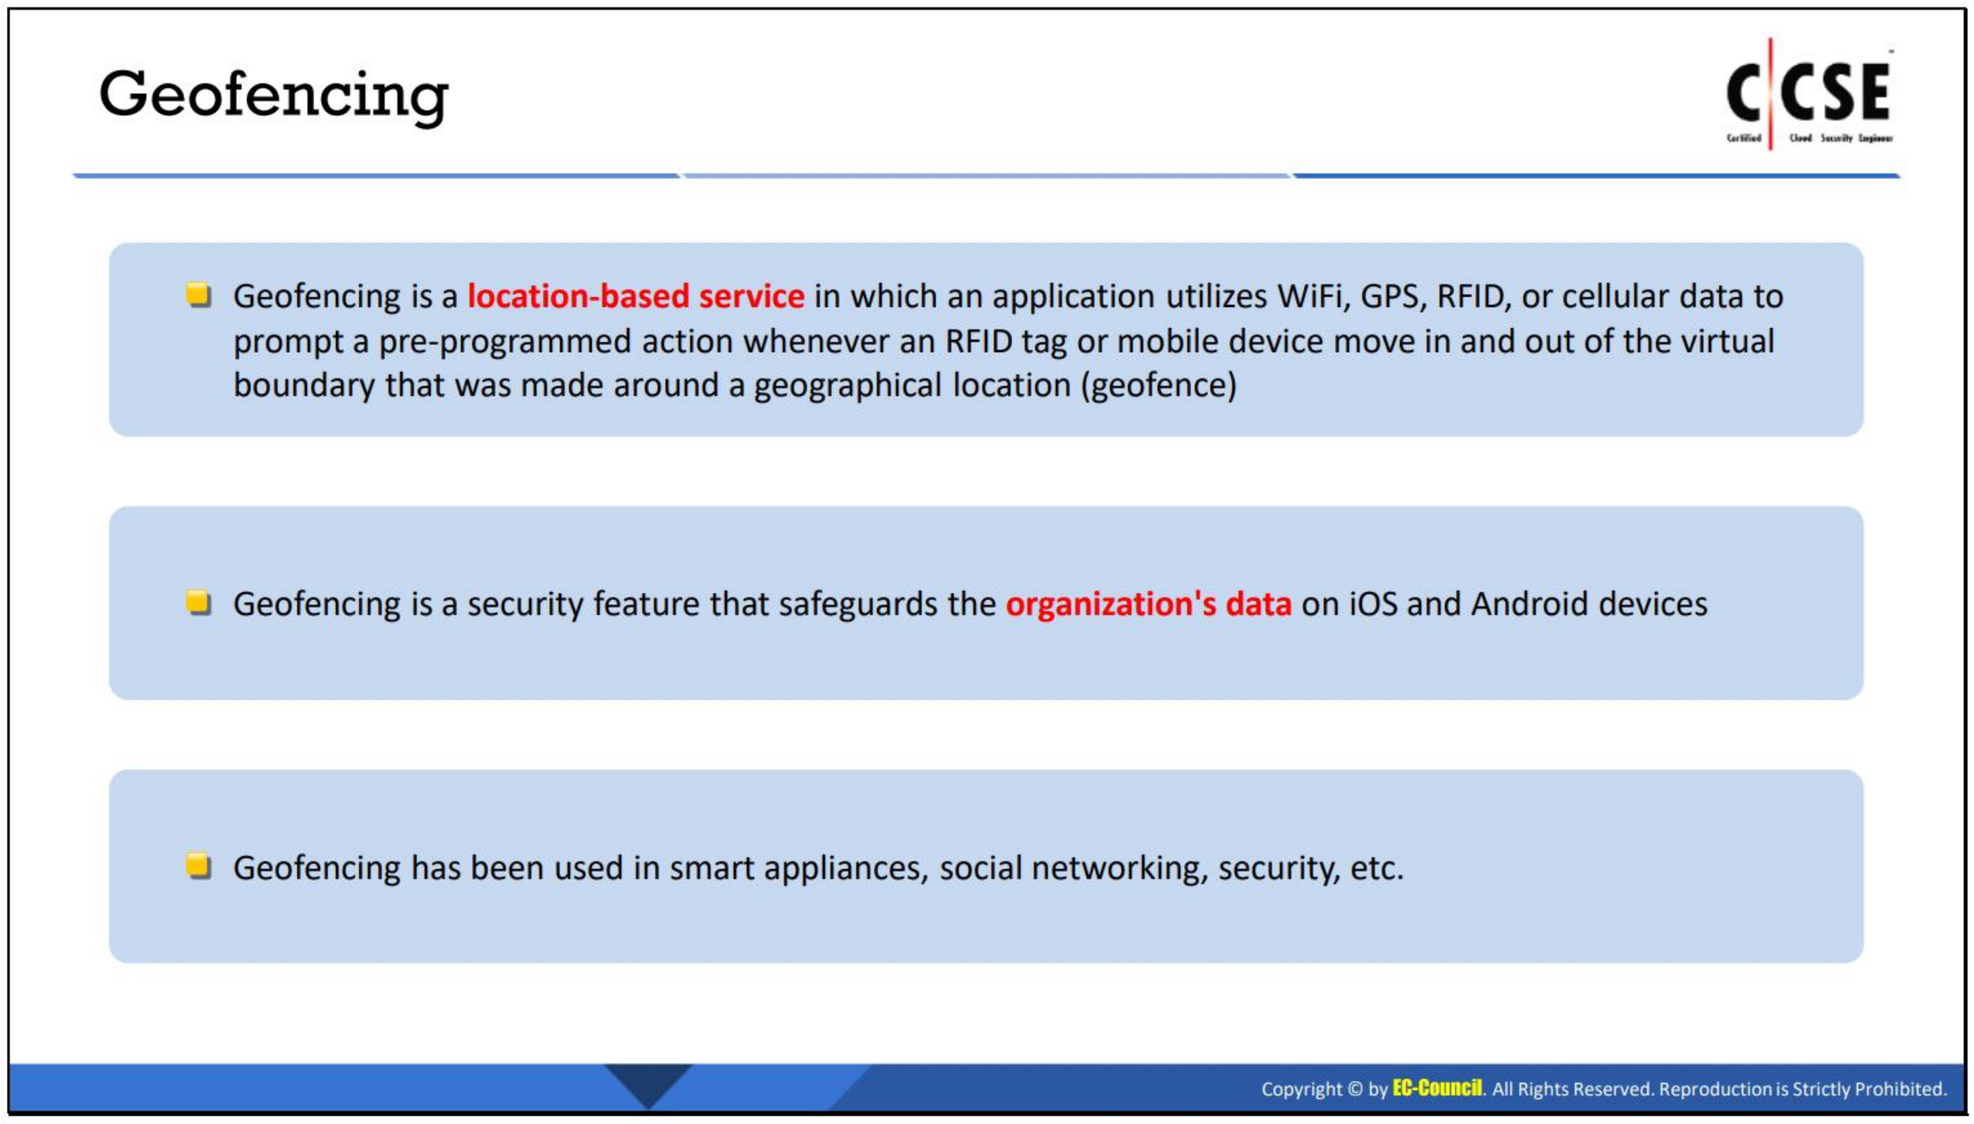
- **Alternative data backup method**: Backing up data includes compressing and encrypting the data and storing it in the cloud storage. Data redundancy provides an additional layer of protection and strengthens the backup by replicating the data. Incorporation of data redundancy in an organization's disaster recovery plan would be beneficial.
- **Better data security**: Date security refers to safeguarding data in the file storage system or databases from unusual activities such as data breaches or cyberattacks. Storing the same data at two or more locations can safeguard an organization during a data breach or cyberattacks.

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

- **Faster data access and updates**: Due to data redundancy, employees can access the data quickly and receive fast updates because the essential information is available on many systems. It plays a significant role in customer-service based organizations wherein the customer anticipates promptness and efficiency.
- **Improved data reliability**: Reliable data is complete and accurate. Companies can utilize data redundancy to double-check the data and verify the correctness. Thus, data redundancy helps in improving data reliability.

hide01.ir

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



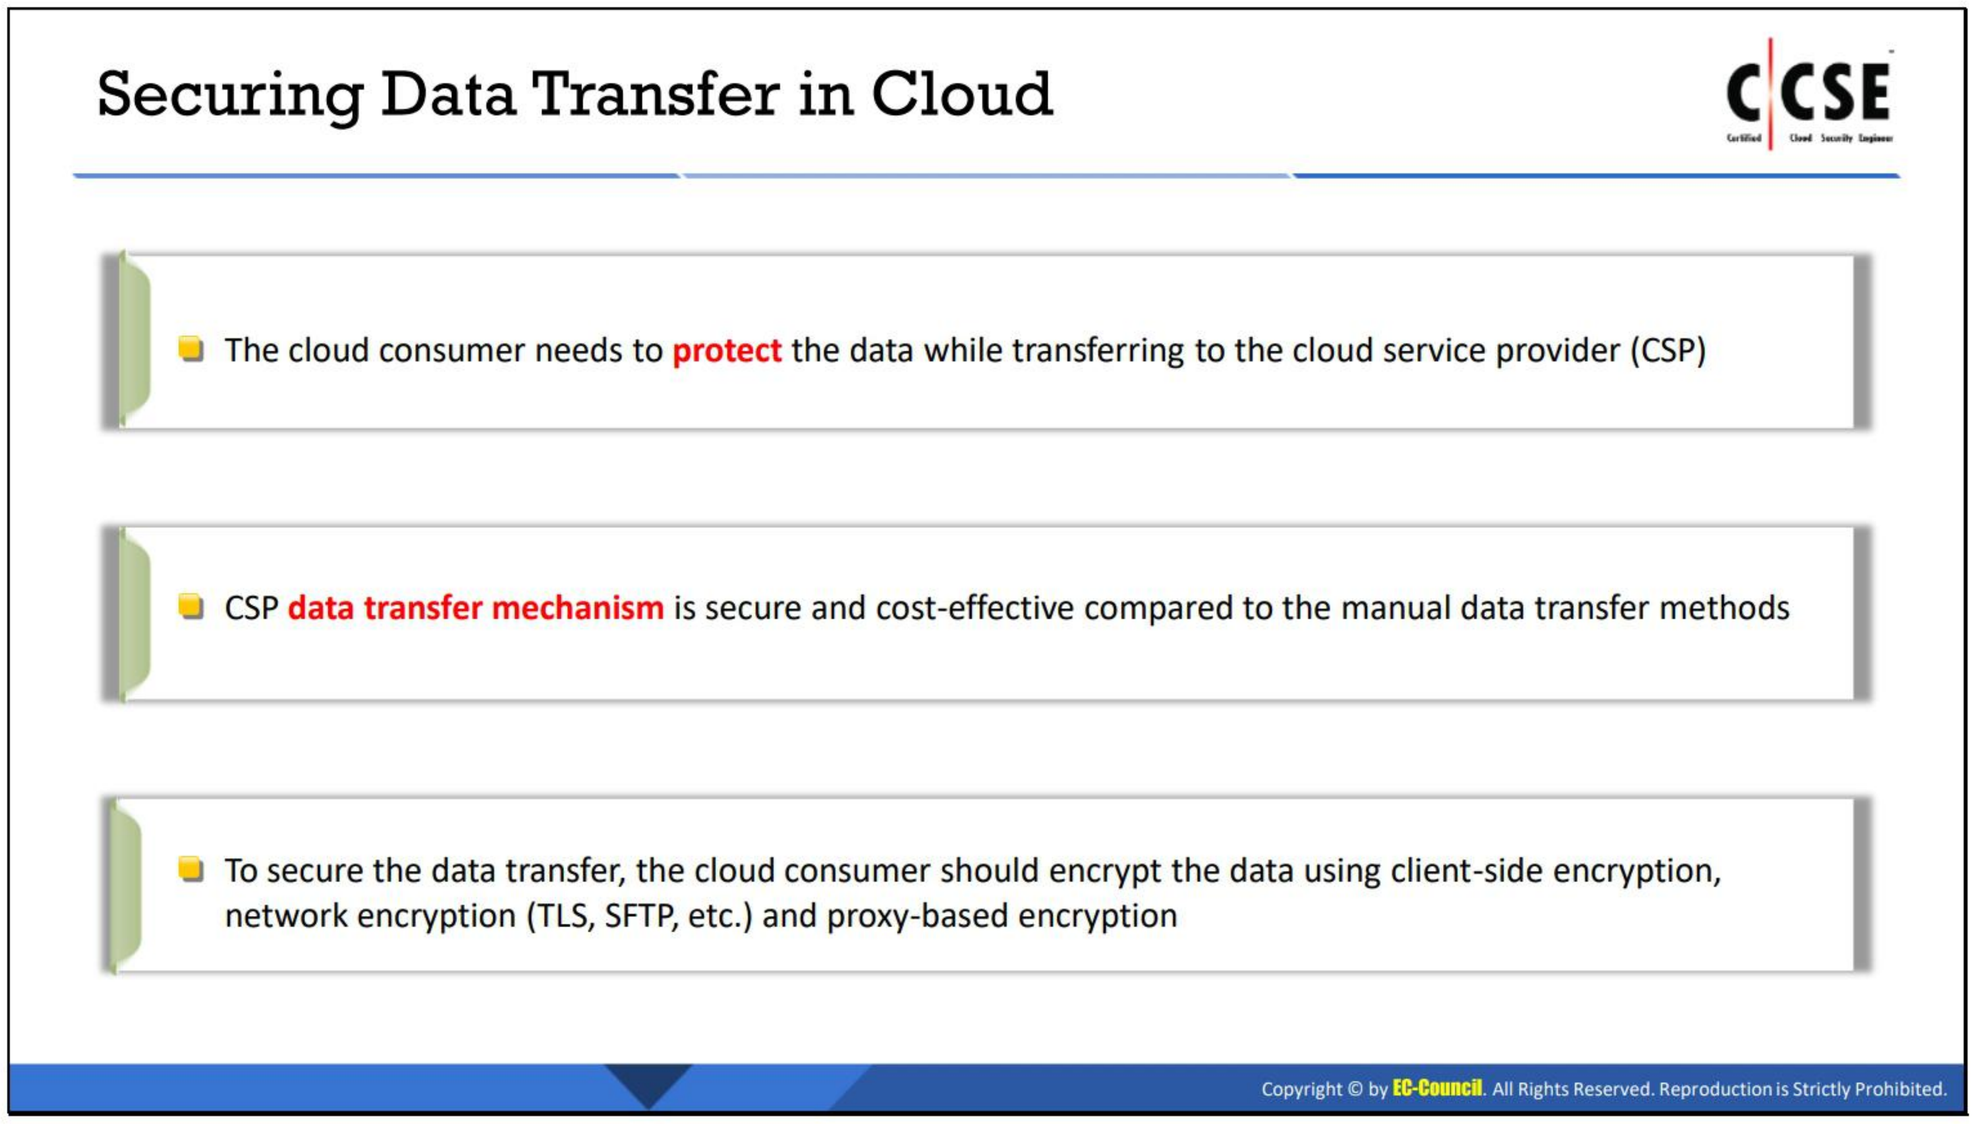
Geofencing

Geofencing is a location-based service in which an application or software utilizes RFID, WiFi, GPS, or cellular data to prompt a pre-programmed action whenever an RFID tag or mobile device moves in and out of the virtual boundary that was made around a geographical location (geofence). Based on the configuration of the geofence, it can trigger text messages or alerts, vehicle fleets can be tracked, technology disabled, location-based marketing data can be sent, etc. Certain geofences help in monitoring activity by sending alerts to the management if someone enters or leaves an area. An organization can utilize geofencing for monitoring employees in the field, automating timecards, etc. Geofencing helps to enhance mobile security; the user can unlock their device at home by using geofence or receive alerts whenever someone enters their home.

For using geofencing, the administrator or developer initially establishes a virtual boundary around a particular location in GPS/RFID enabled software. Based on the specification of the administrator or developer, the virtual geofence prompts a response whenever an authorized device moves in or out of the area. Geofencing has been used in smart appliances, social networking, security, etc.

Geofencing is a security feature that safeguards the organization's data on iOS and Android devices. If the organization security policy disallows the users to access the data outside the corporate premises, then the administrator can configure a geofence policy to detect non-compliant devices along with the action to be taken on them. With a geofence, the administrator can create a virtual fence, thereby limiting the utilization of corporate devices within the restricted area.

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



Securing Data Transfer in Cloud

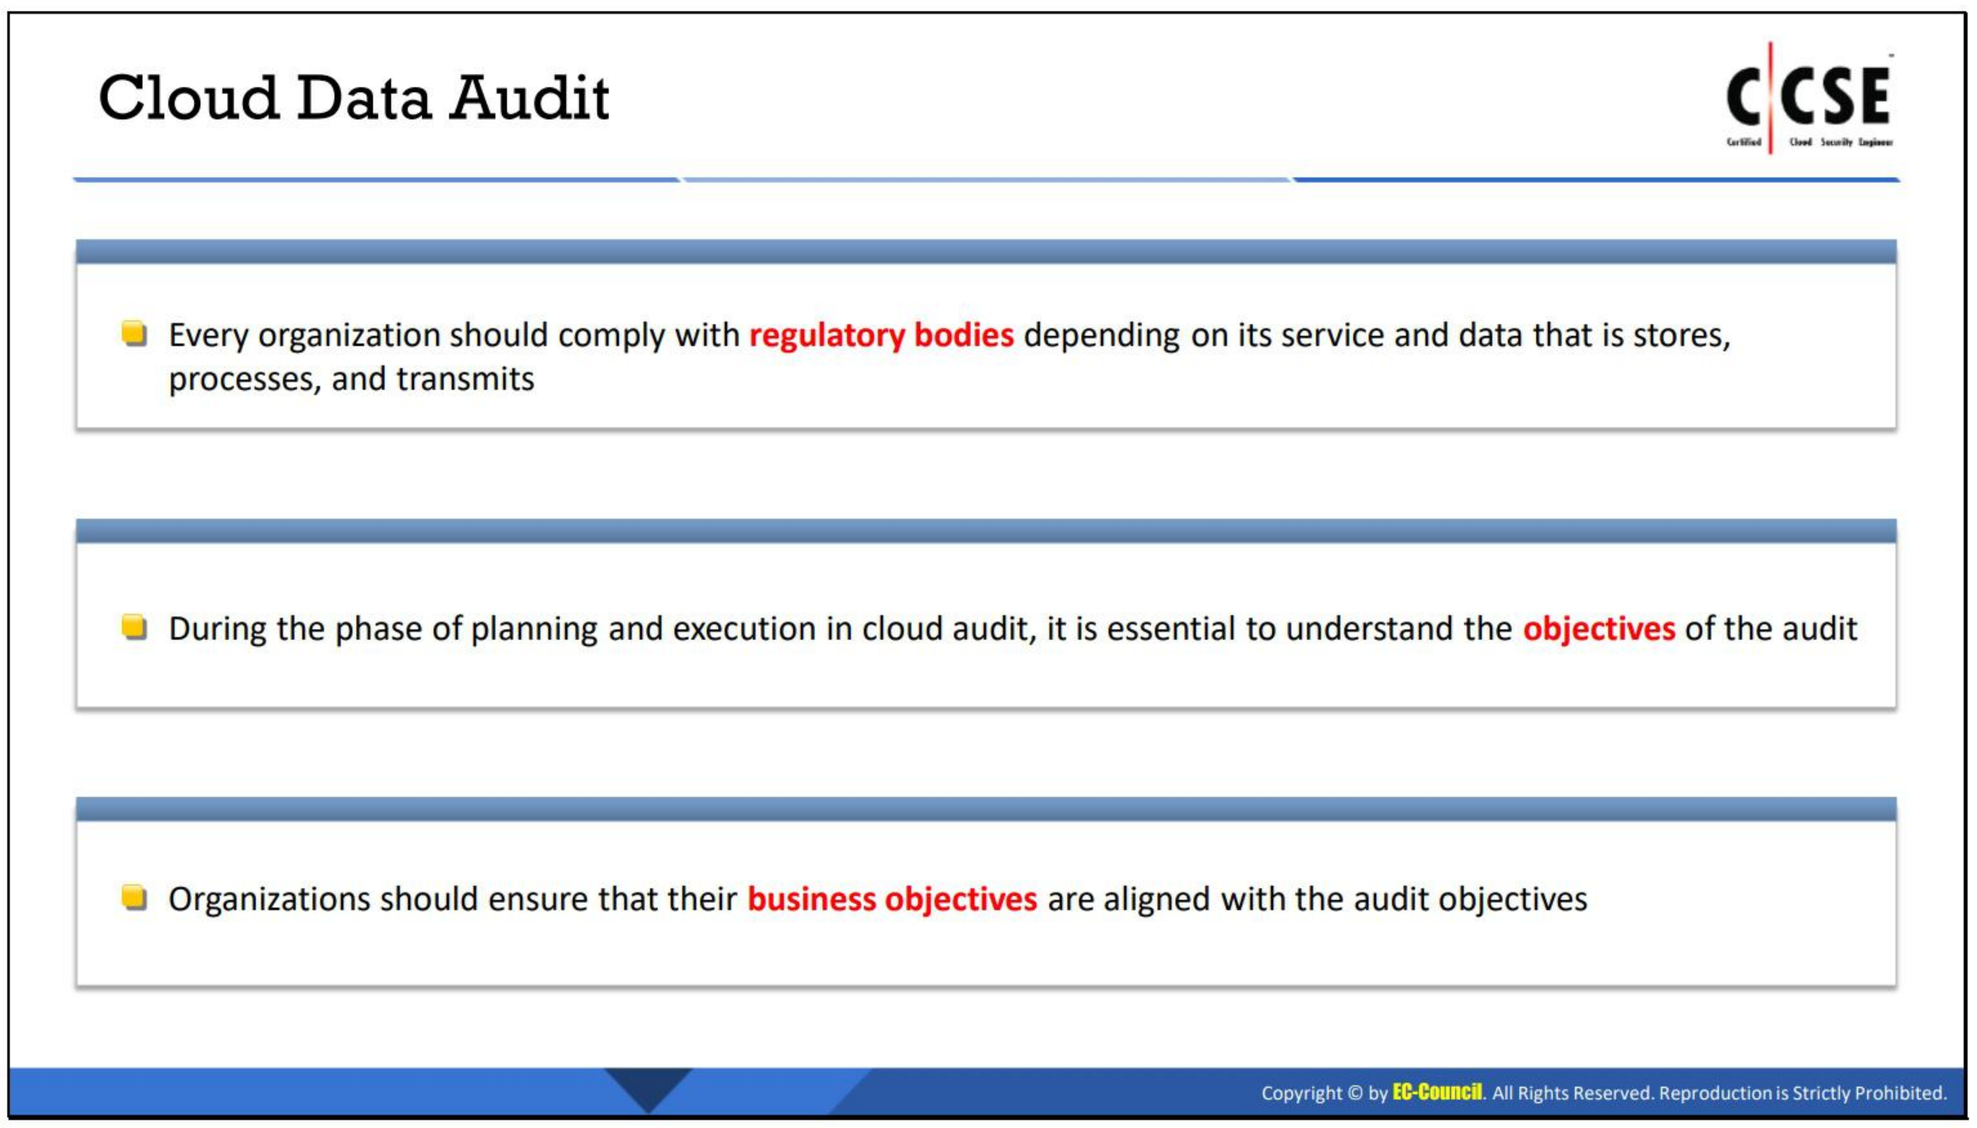
While migrating data to the cloud, it is important to ensure that the data is protected. The data transfer mechanisms of CSPs are more secure and cost-effective than manual data transfer methods. For example, transferring data to the CSP's object storage using an API is more secure and reliable than implementing a secure file transfer protocol (SFTP) on a VM in the same CSP.

Based on the encryption options supported by the cloud platform, there are few options for in-transit encryption.

- The data can be encrypted before sending it to the cloud.
- Network encryption can also be used, where the data over the network is encrypted using encryption algorithms, processes, and standards.
- By default, most of the CSP's API uses TLS. If the cloud provider does not support this, the cloud consumers should shift to other CSP as it is an important security capability.
- Proxy-based encryption is also an option. This type of encryption lets a third-party or proxy to re-encrypt the data before transferring it to the cloud.

There are certain cases where the cloud consumer needs to accept public or untrusted data. In these cases, there should be some security mechanisms to sanitize the data before processing it with the existing data. Before integrating the public data with the existing data, the public data should be isolated and scanned.

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



Cloud Data Audit

An audit, in general, is an independent third-party service that gathers evidence through various approaches such as physical inspection, analysis of procedures, inquiry, or re-performance, etc. Any or all of these combinational approaches are performed to provide an opinion on the design and operational effectiveness of controls, that are discovered in various areas. For example, some of these areas are communications, network security, security incidents, data management, risk management, vulnerability management with remediation, ethical behavior, and top management commitment toward transparency, etc.

Every organization should comply with regulatory bodies depending on its service and data that is stores, processes, and transmits. Different kinds of compliance may be required based on its service. For example, Payment Card Industry Data Security Standard (PCI-DSS) certification is required if cardholder data, bank account information, or any payment-related sensitive data is stored or transmitted. PCI-DSS requires, service providers to encrypt all the sensitive financial data. HIPAA certification is required if any sensitive data that includes health records and patient names are stored or processed. HIPAA requires, institutions or individuals to safeguard patient health records. Service Organization Control (SOC 2) does not have any specific requirement, but has criteria such as implementation of various security controls. SOC 2 provides reports that contain controls that are to be implemented by the service providers such as encryption, password policies, regular patch updates, etc. SOC 2 safeguard end-users' trust. During the phase of planning and execution in cloud audit, it is essential to understand the objectives of the audit. Organizations should try to make sure that their business objectives are aligned with the audit objectives. This pre-planning ensures that time and resources spent will help in achieving control over the internal environment and the risk of a qualified opinion may become low.

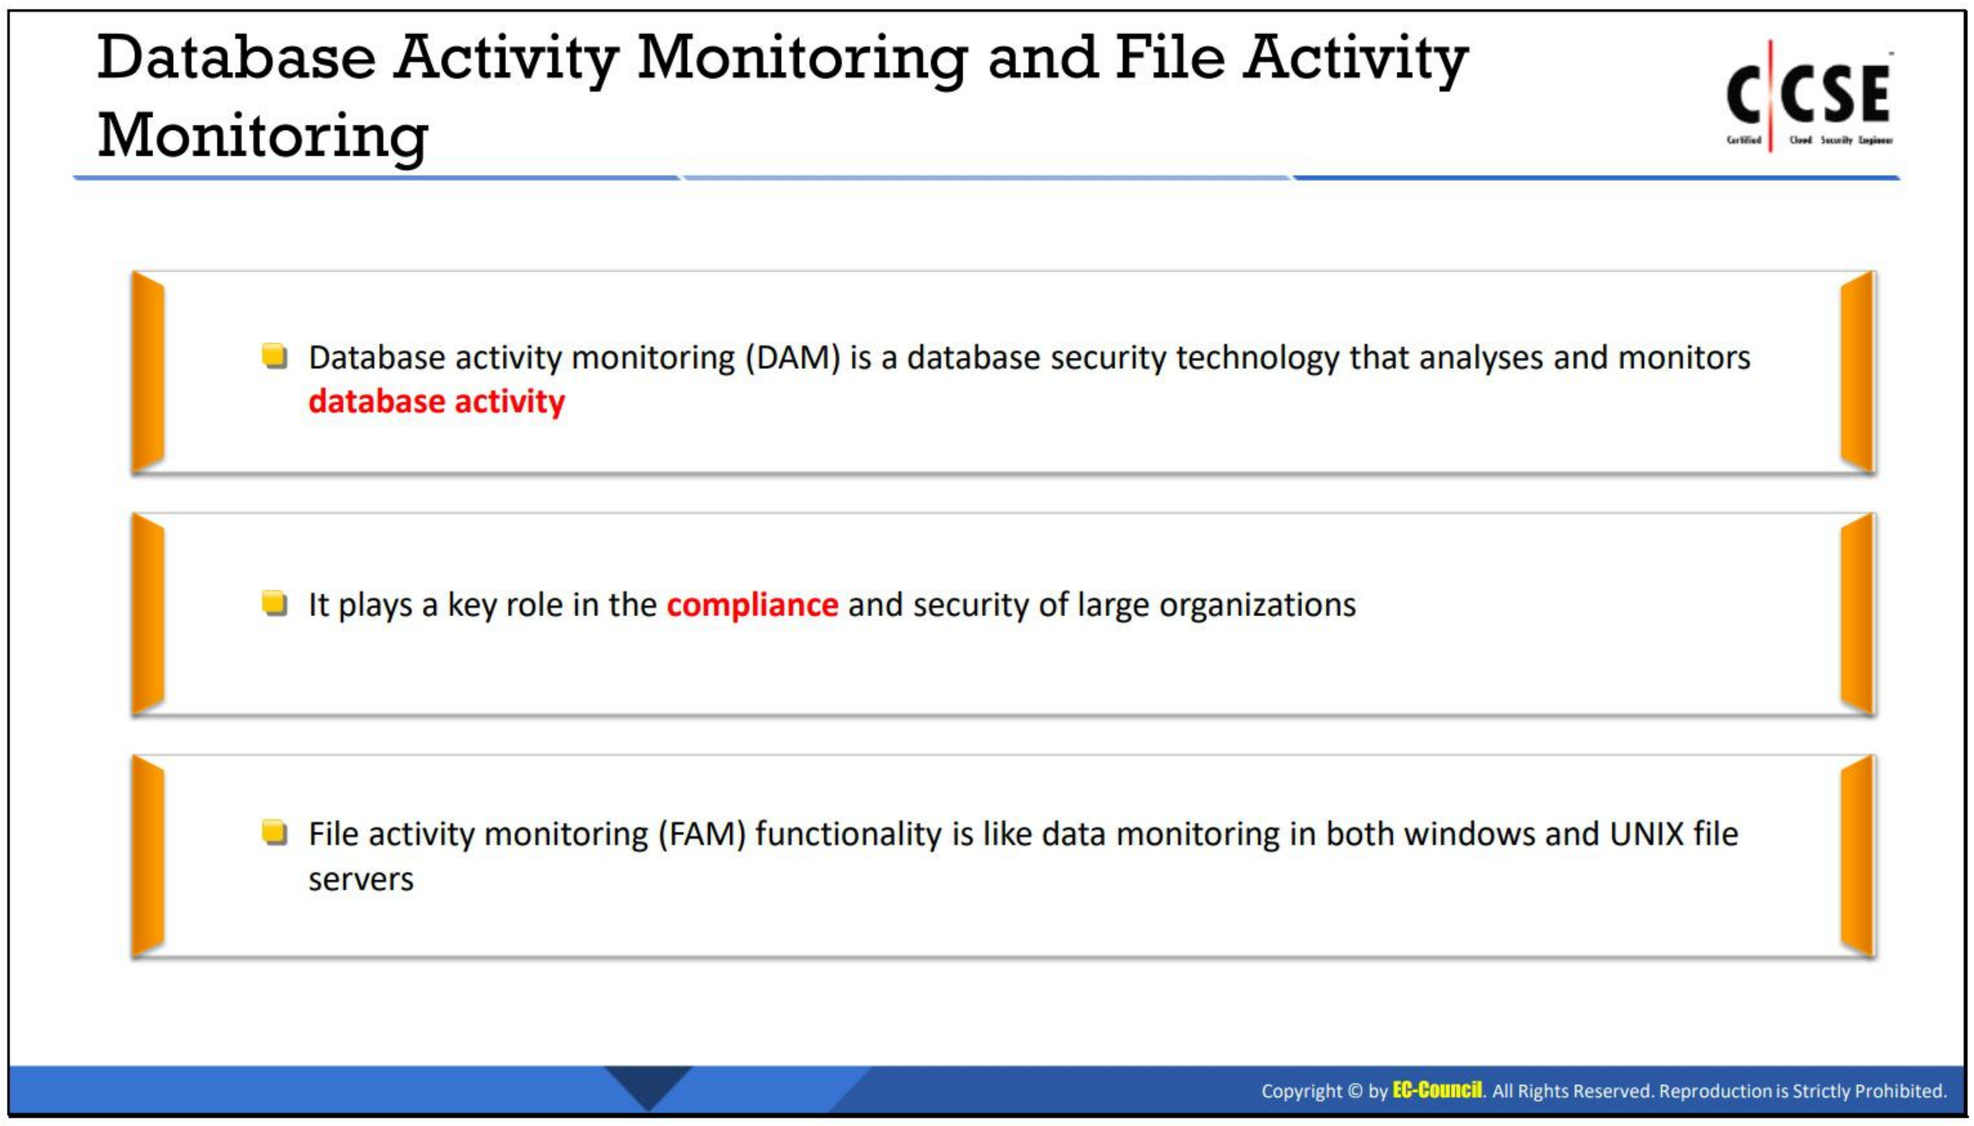
Certified Cloud Security Engineer Copyright © by **EC-Council**
All Rights Reserved. Reproduction is Strictly Prohibited.

Some of the best practices to comply with cloud data audit are listed below.

- **Encryption:** Some of the information might not be personally identifiable but still the data may be sensitive; thus, this kind of data should be encrypted. Any cloud storage, data-at-rest and data in transmission should be encrypted.
- **Reliability:** Even though the applications are built with a reliable process, every organization should have a disaster recovery plan in case of any crucial application failure. If the service is multi-regional, the data can be accessed without any issues.
- **Privacy:** Auditors inspect how organizations are maintaining customer's privacy, whether organizations follow European GDPR or any laws to protect sensitive data, and where is it protected.
- **Data retention:** Any data should have a retention period. The main concern of retention is on sensitive data that includes name, contact numbers, address with house number, etc. architecture
- **Integrity:** Integrity is one of the most crucial areas of security; the end-user's or client's trust can be maintained if the data integrity is not compromised. Possible secure channels and encryption should be enforced to avoid tampering of data.

hide01.ir

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



Database Activity Monitoring and File Activity Monitoring

Database Activity Monitoring (DAM)

DAM plays a key role in the compliance and security of large organizations. It is a database security technology that analyzes and monitors database activity. It monitors real-time database activity continuously. DAM may combine data from network sniffing, system tables, native audit information, and memory scraping, that provides a broad picture of database activity. This extracted data can be used to analyze and report database activity and is used for investigations. DAM strategy also protects sensitive databases from malicious attackers.

Some of the common uses cases for DAM are listed below:

- **Privileged user monitoring:** It is crucial to monitor users such as sysadmins, database administrators, developers, etc. who have the highest level of privileges, for protection against internal and external attacks. This privileged monitoring audits all activities such as the creation of new accounts using root accounts or accessing sensitive data and also monitors authorized changes such as deletion of any tables. As the attackers frequently target privileged user accounts, it is a best-practice to identify any compromised accounts or systems.
- **Application activity monitoring:** The foremost objective of application activity monitoring is providing end-user liability and also detection of fraud, which is usually through enterprise applications instead of database access. Multi-tier applications that are built on middle-tier servers, hide the end-users' identity in the database through a connection pooling mechanism. Using this mechanism all the user traffic is collected in few database connections; these can be identified using the account name of generic service only. Application activity monitoring allows identifying unauthorized activities.

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

- **Cyberattack protection:** Consider a scenario where an application is vulnerable to SQL injection; this vulnerability occurs due to insecure coding that uses relational databases. Here, an attacker tries to inject a malicious SQL statement into the input field, and this statement contains part of the application statement and part of the malicious statement. As the application is insecure and is vulnerable to SQL injection, the malicious SQL injections get executed and allow data access, modification, or deletion. DAM helps prevent SQL injection attacks by identifying normal behavior and attack structure statements. It uses policies for the identification of malicious behavior. DAM helps in regulatory compliance with PCI-DSS, HIPAA, etc.

DAM architectures:

- **Interception-based:** DAM usually intercepts between database client and database server. Interception can be done at numerous locations such as database memory, network, or DB libraries. Packet sniffing can be used if the network traffic is unencrypted.
- **Memory-based:** A lightweight sensor is used to protect databases and collect SQL statements that are being tried to execute. Based on this similar lightweight sensor architecture, the lightweight sensor is used in a host that is attached to OS level for private data structure inspections. This sensor tracks all transactions of the database such as traffic of network, host, etc. This is used to protect databases that are deployed in virtual machines or the cloud.
- **Log-based:** Some of the DAM systems process and isolate information from logs of database transactions. Most of these data will be stored in redo logs but all the information is not required in redo logs (such as SELECT statements executed in the database).
- **File activity monitoring (FAM):** FAM functionality is similar to data monitoring in both windows and UNIX file servers.

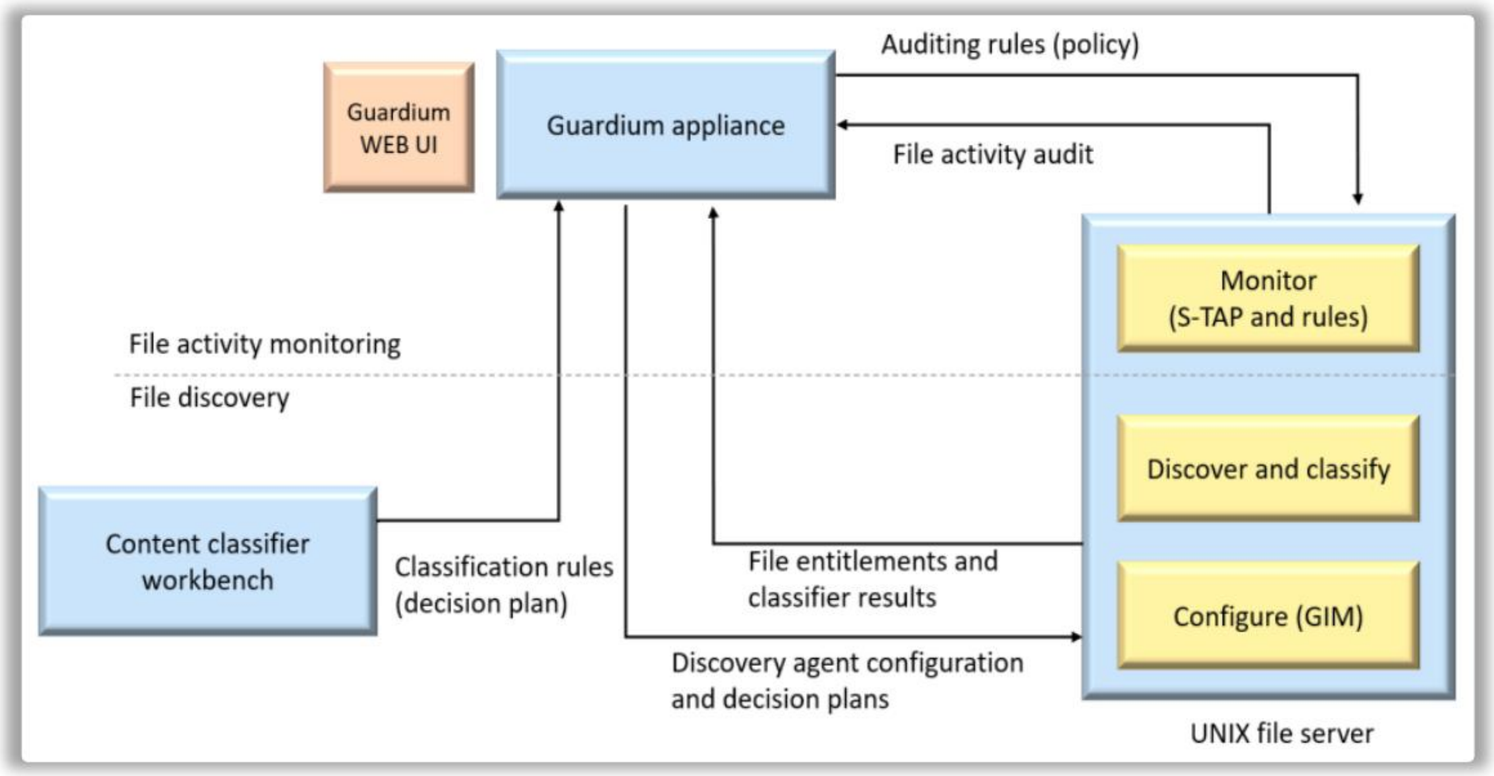


Figure 4.43: FAM Functionality in UNIX File Servers

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

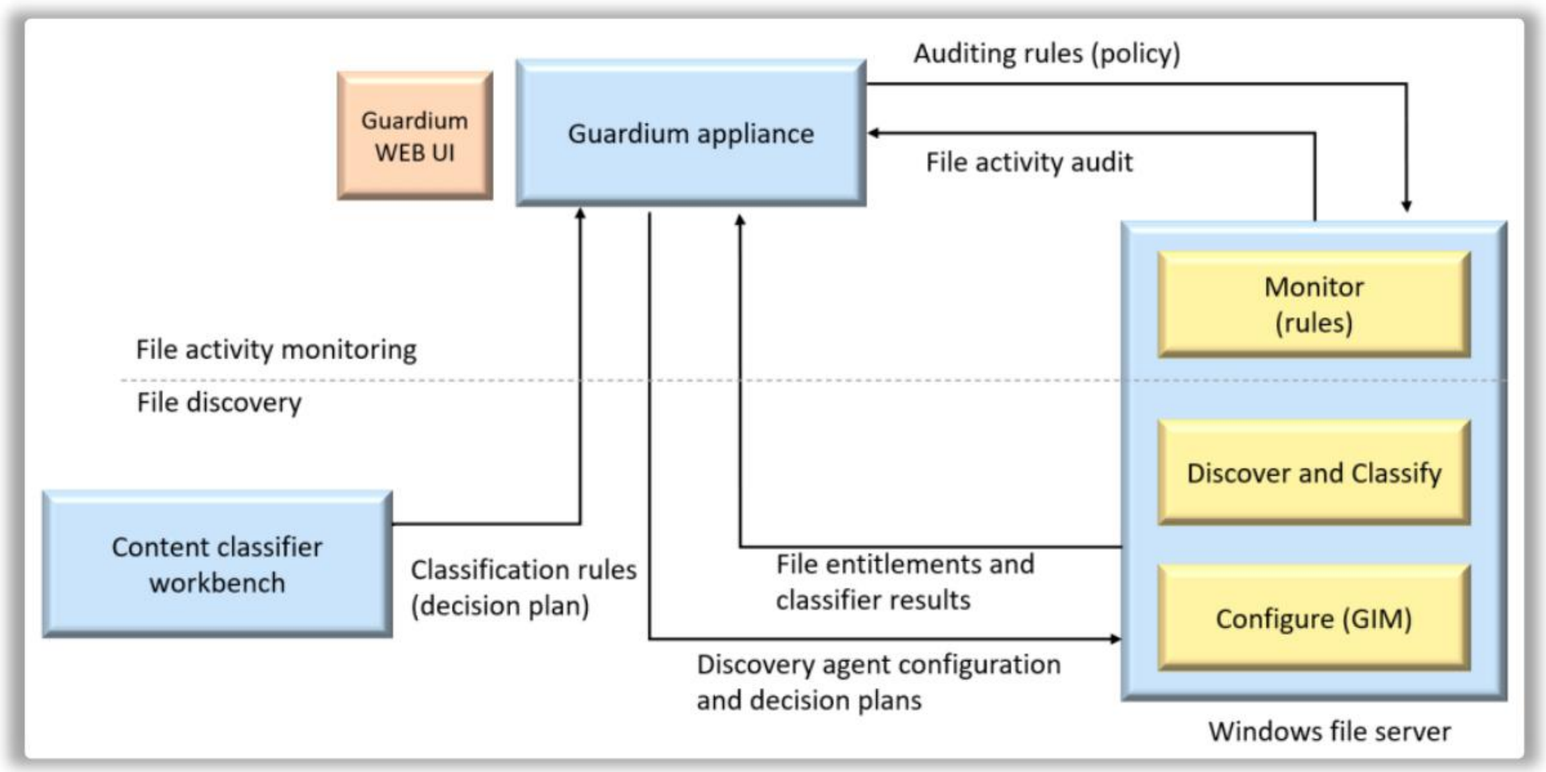


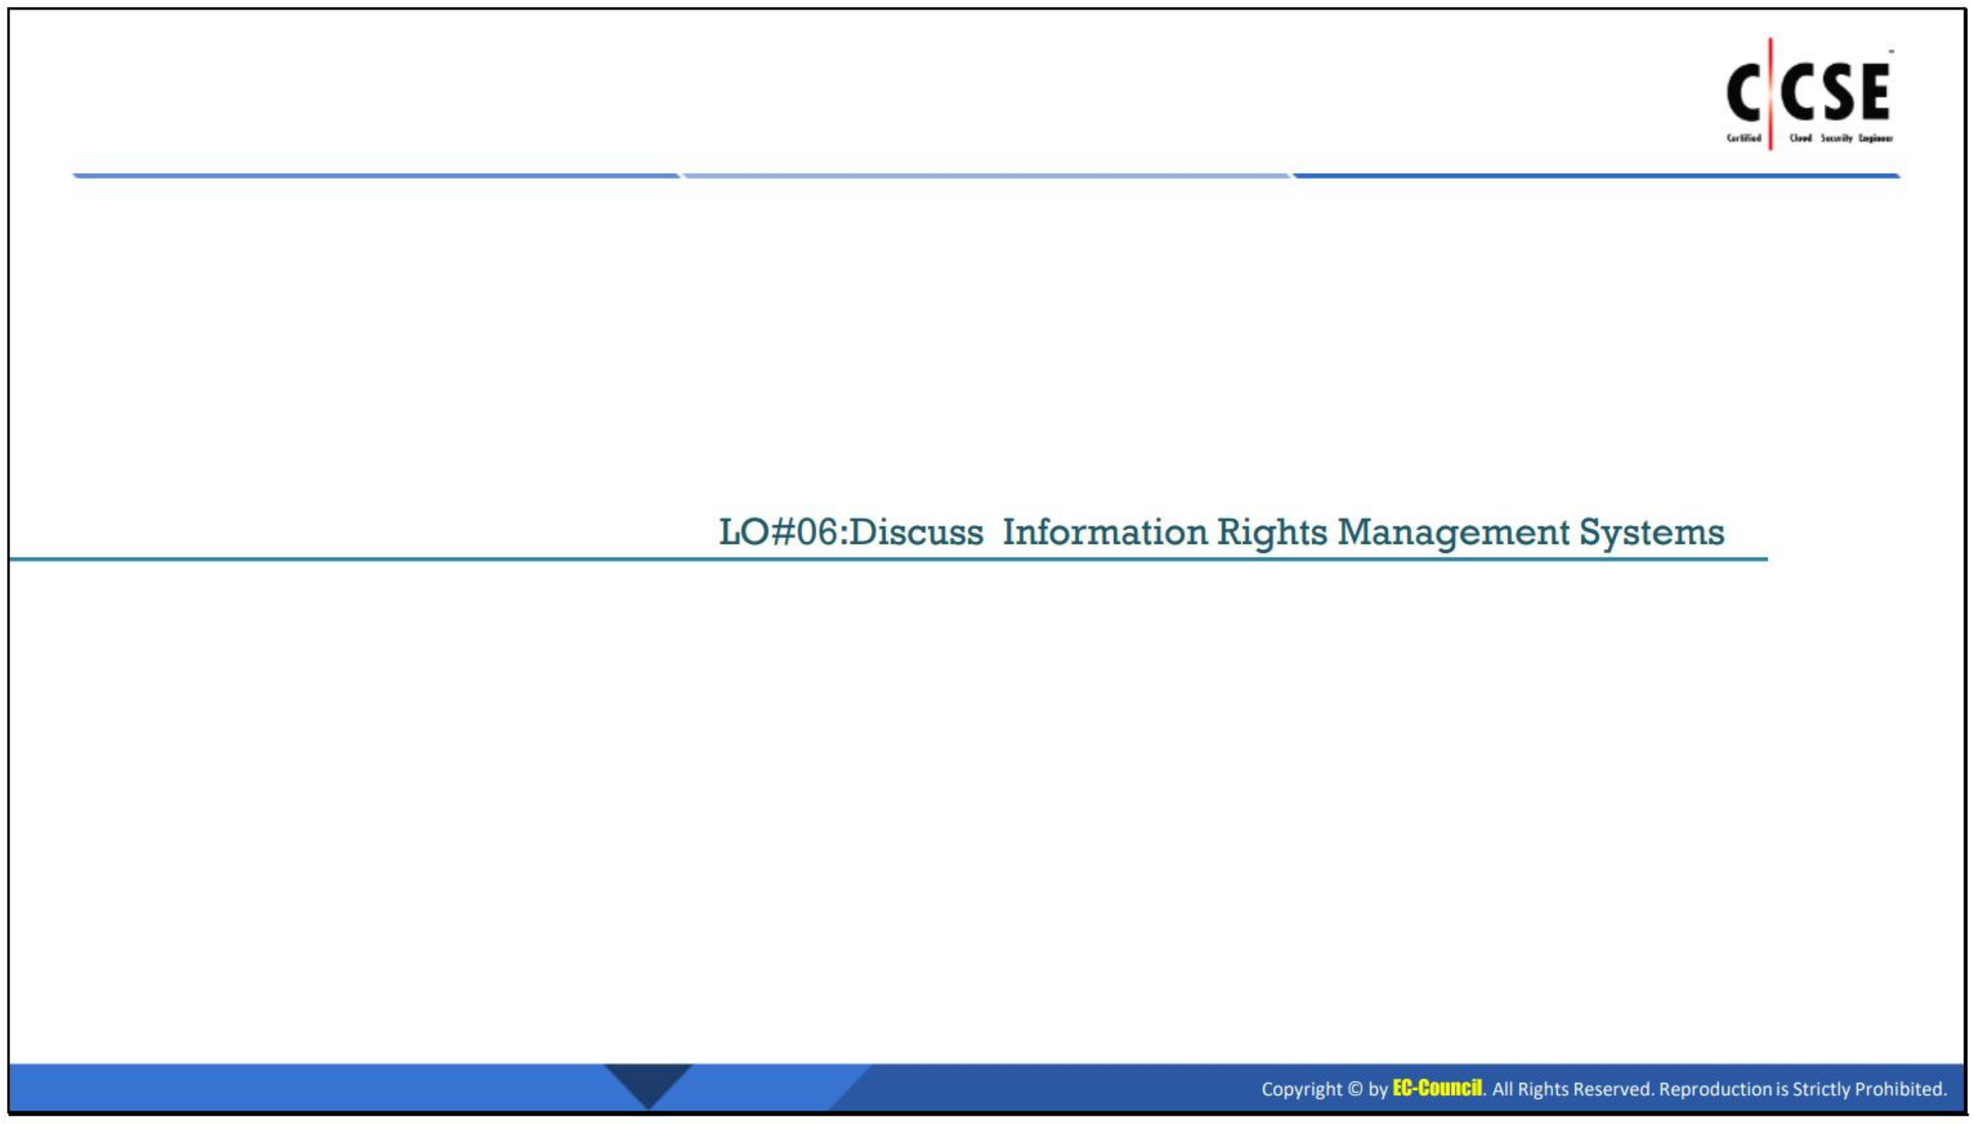
Figure 4.44: FAM Functionality in Windows File Servers

FAM discovery and classification provides the classification of content in the given fileserver folder. Depending on the results and focusing on crucial assets, policies can be created, groups can be defined, and files can be audited. Details such as the list of files and folders with owner name, size of files or folders, timestamp, and access permissions are obtained by performing a basic discovery scan. Discovery and classification can be handled using a discovery agent, which is called a crawler. Metadata and data from discovery and classification are sent to the Guardium system through file crawlers.

Based on the Guardium policy rules, FAM alerts and blocks file access. Rules decide what file servers and files are to be monitored and what actions are to be taken. If there is a policy violation, it specifies an action such as alert, logs the violation, or block access. Read, delete, write, execute, owner, properties, etc., operations are monitored. Activities matching security policy will be sent to Guardium collector, which is stored in Guardium repository.

hiu01.ir

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

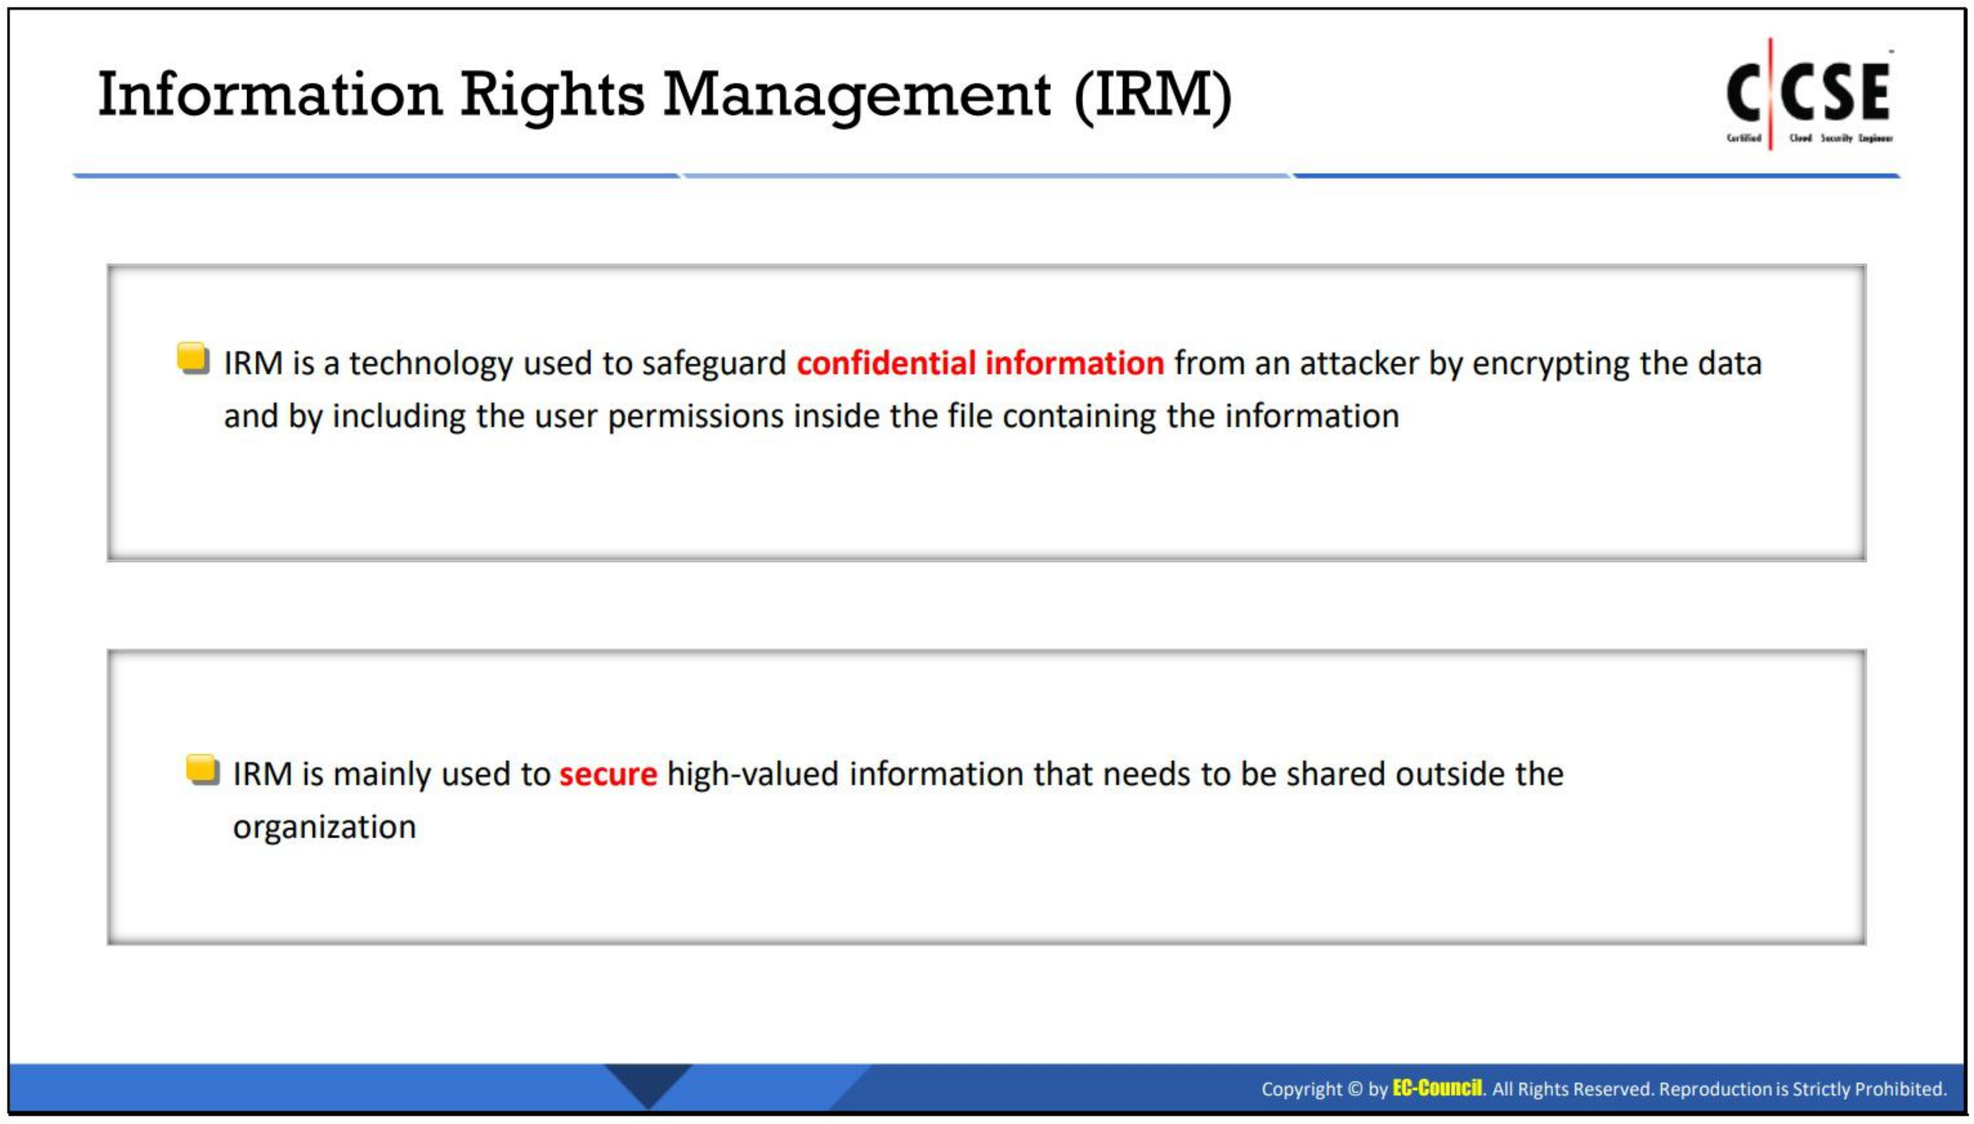


LO#06: Discuss Information Rights Management Systems

Implementing information rights management (IRM) in cloud to control and secure the data as well as to prevent unauthorized users from accessing data. This section explains the challenges with IRM and solutions for successful implementation of IRM.

hidan01.ir

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



Information Rights Management (IRM)

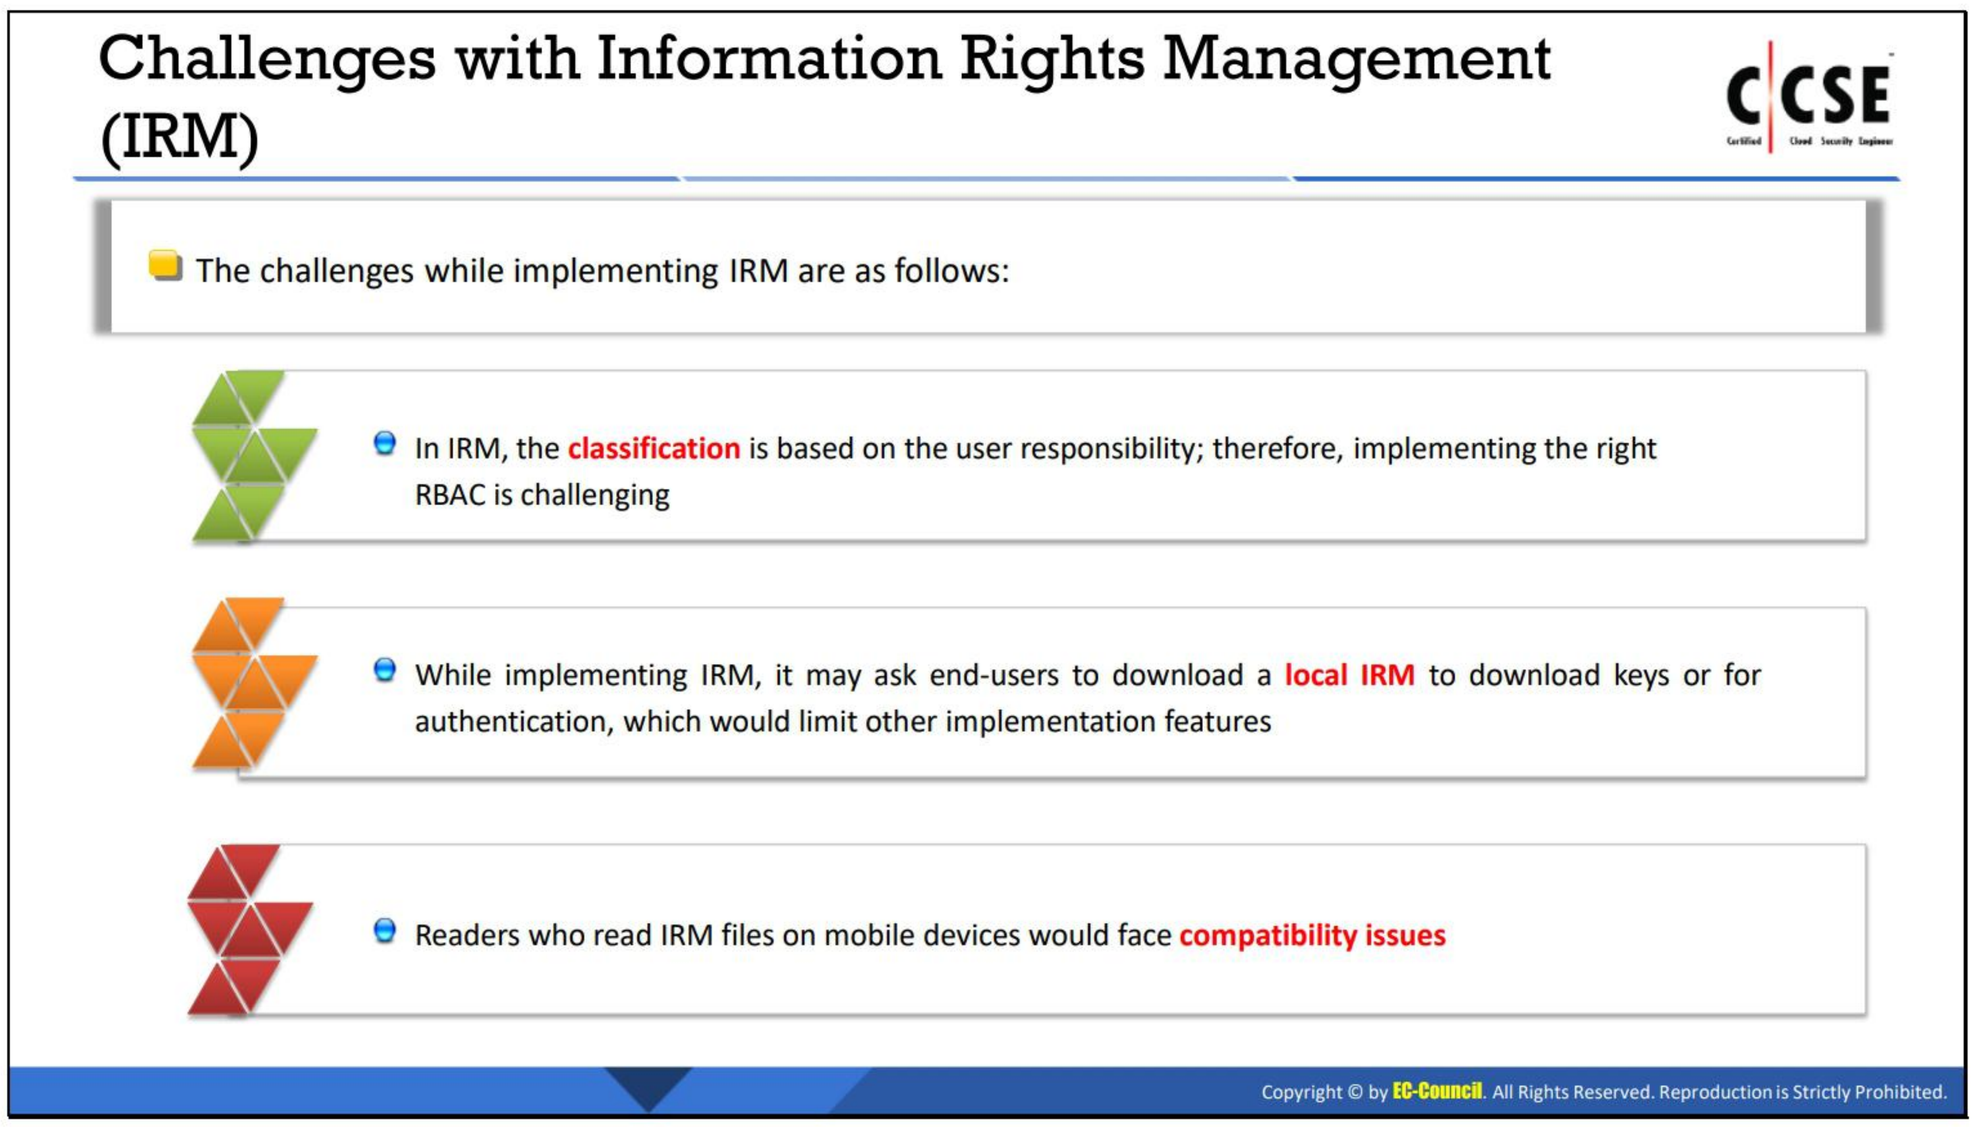
IRM is a technology used to protect sensitive information from an attacker by encrypting the data and by including user permissions inside the file containing the information. Using user permissions, the owner can restrict specific actions on the file. It is different from other security technologies used for securing sensitive files.

Many intruders and hackers try to breach sensitive corporate information. They would sell this sensitive information to traders who use this sensitive information and make trades in the stock market. For this reason, many corporations have started using IRM.

IRM is mainly used to secure high-valued information that needs to be shared outside the organization. When IRM is embedded into those sensitive files, only the authorized person can view the data and the owner can revoke the usage at any time. IRM can also be used to protect emails, web pages, database rows and columns, and other data.

IRM uses an access control with the file or document. The access control list (ACL) defines which users can open the file and what activities they can do with the file such as printing, editing, copying, and saving. IRM also supports file expiration, where the content in a file will be unviewable after the set time. Unlike other security technologies, documents with IRM are protected throughout the lifetime whether it is in motion, rest, or while in use.

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



Challenges with Information Rights Management (IRM)

The challenges while implementing IRM are as follows:

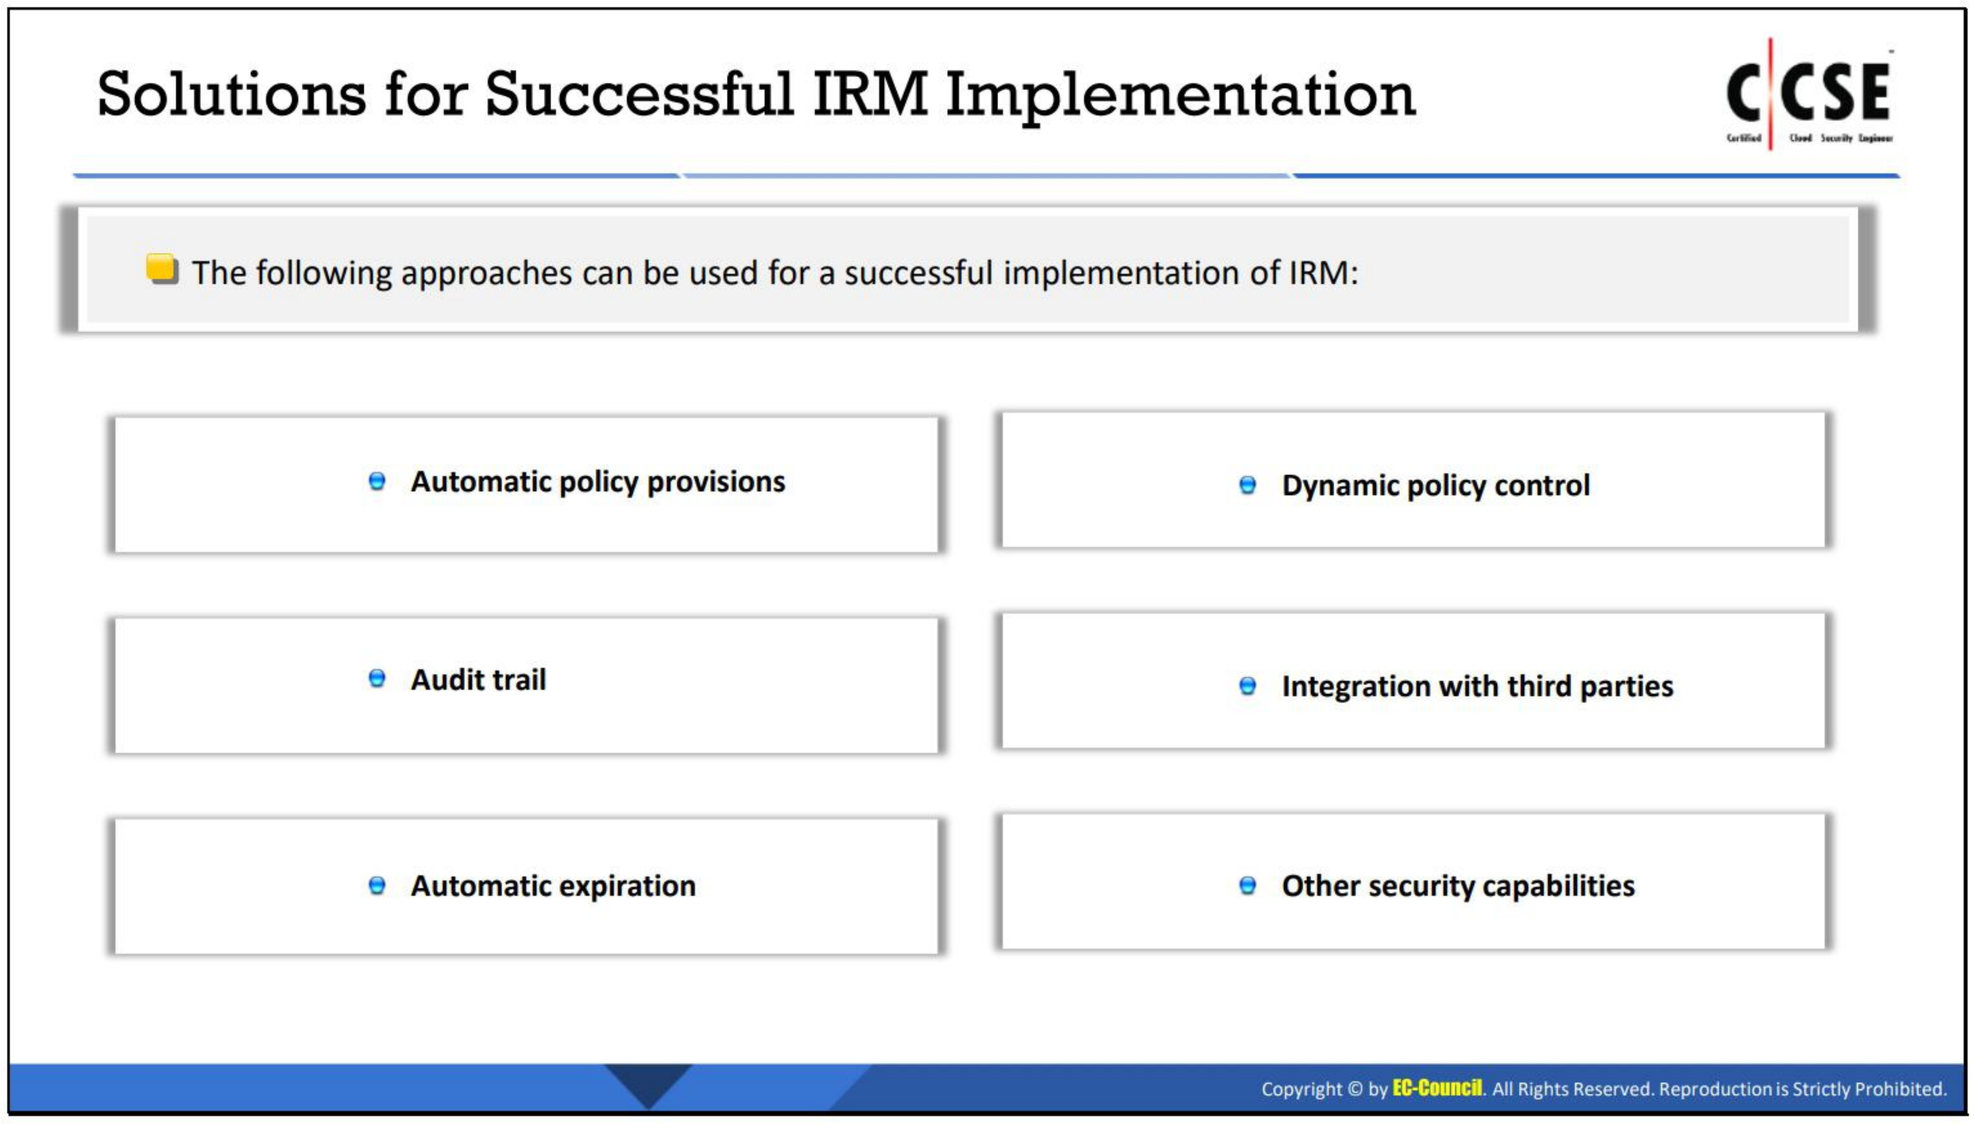
- All the users and data to be embedded with IRM should have the same encryption keys. Implementing IRM requires a strong identity infrastructure and it must expand to customers, users, or others with which the files will be shared.
- Each resource will have an access policy. Each user should have an account and keys to access the resources. All these provisions need to be perfectly made for the successful implementation of IRM. This can be implemented by using automatic policy provision.
- User access to each resource is assigned to the user role by using the RBAC model. All the assignments of roles and users must be integrated with IRM policies. In IRM, the classification is based on the user responsibility; therefore, implementing the right RBAC is challenging.
- The identity infrastructure is established by creating a single location where all the users and access permissions are authorized. Therefore, the cloud administrator must choose the right method based on the security requirements.
- While implementing IRM, it may ask end-users to download a local IRM to download keys or for authentication. This would limit other implementation features. Hence, this should be considered while developing the architecture.
- The readers of IRM-protected files should have IRM-ware reader software. Microsoft products with the latest version support IRM, but other users would face trouble. Therefore, it should be properly tested before deploying.

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

- Readers who read IRM files on mobile devices would face compatibility issues. Thus, the usage of mobile platforms and IRM should be properly tested.

hide01.ir

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



Solutions for Successful IRM Implementation

The following approaches can be used for a successful implementation of IRM:

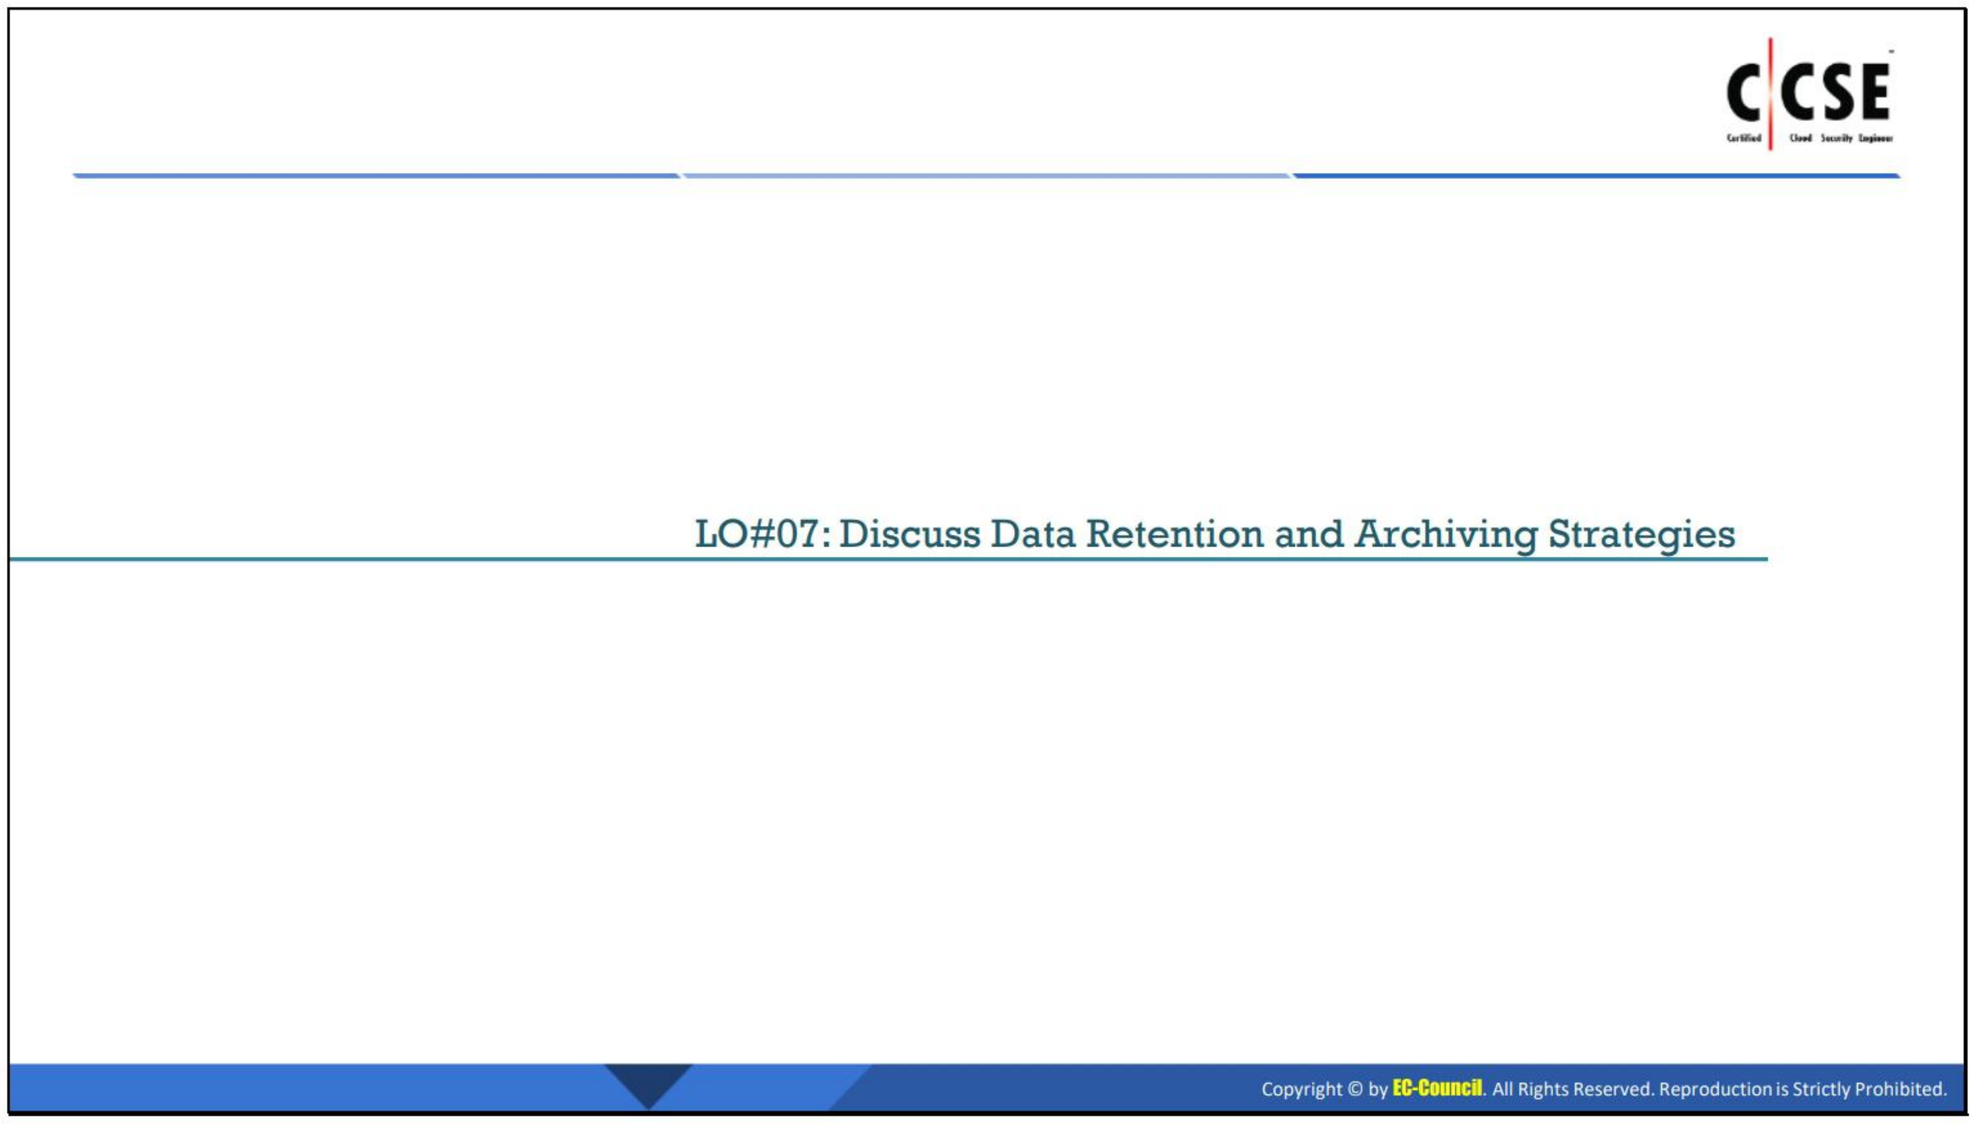
- **Automatic policy provisions:** IRM implementation becomes easier with automatic policy retention. It is based on the file location, document origin, and keywords. Organizations can assign this policy to certain folders, content, or any information that reached at production stage. This saves time and it is also an efficient way to manage a large piece of sensitive information.
- **Dynamic policy control:** IRM policies that govern the content must evolve as per the changes in the business. Using dynamic policy control, the content owners can change the user permissions whenever required, regardless of where the data has been stored. The policies reside under the policy server within the content; therefore, the owner can change or revoke permissions at any time.
- **Audit trail:** Audit trail is an unalterable log of access to the system. It also contains records of modifications, additions, and deletion of any information managed by a system. It also lists the time of access and the actions taken. Using this, a cloud security professional can ensure that the data is delivered and viewed by the users, and it provides proof that all the organization's information security policies are followed.
- **Integration with third parties:** By integrating with third-party email filtering engines, the organization's emails are automatically secured in compliance with the company's information security policies.
- **Automatic expiration:** This feature enables the owner to revoke access to files, attachments, or emails at any point of time, thereby enabling the information security policy when the data is stored or distributed.

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

- **Other security capabilities:** Owners should provide other security capabilities such as restrict printing the entire document or a portion of it, disabling screenshot and copy/paste capabilities, and watermarking pages if print access is granted.

hide01.ir

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

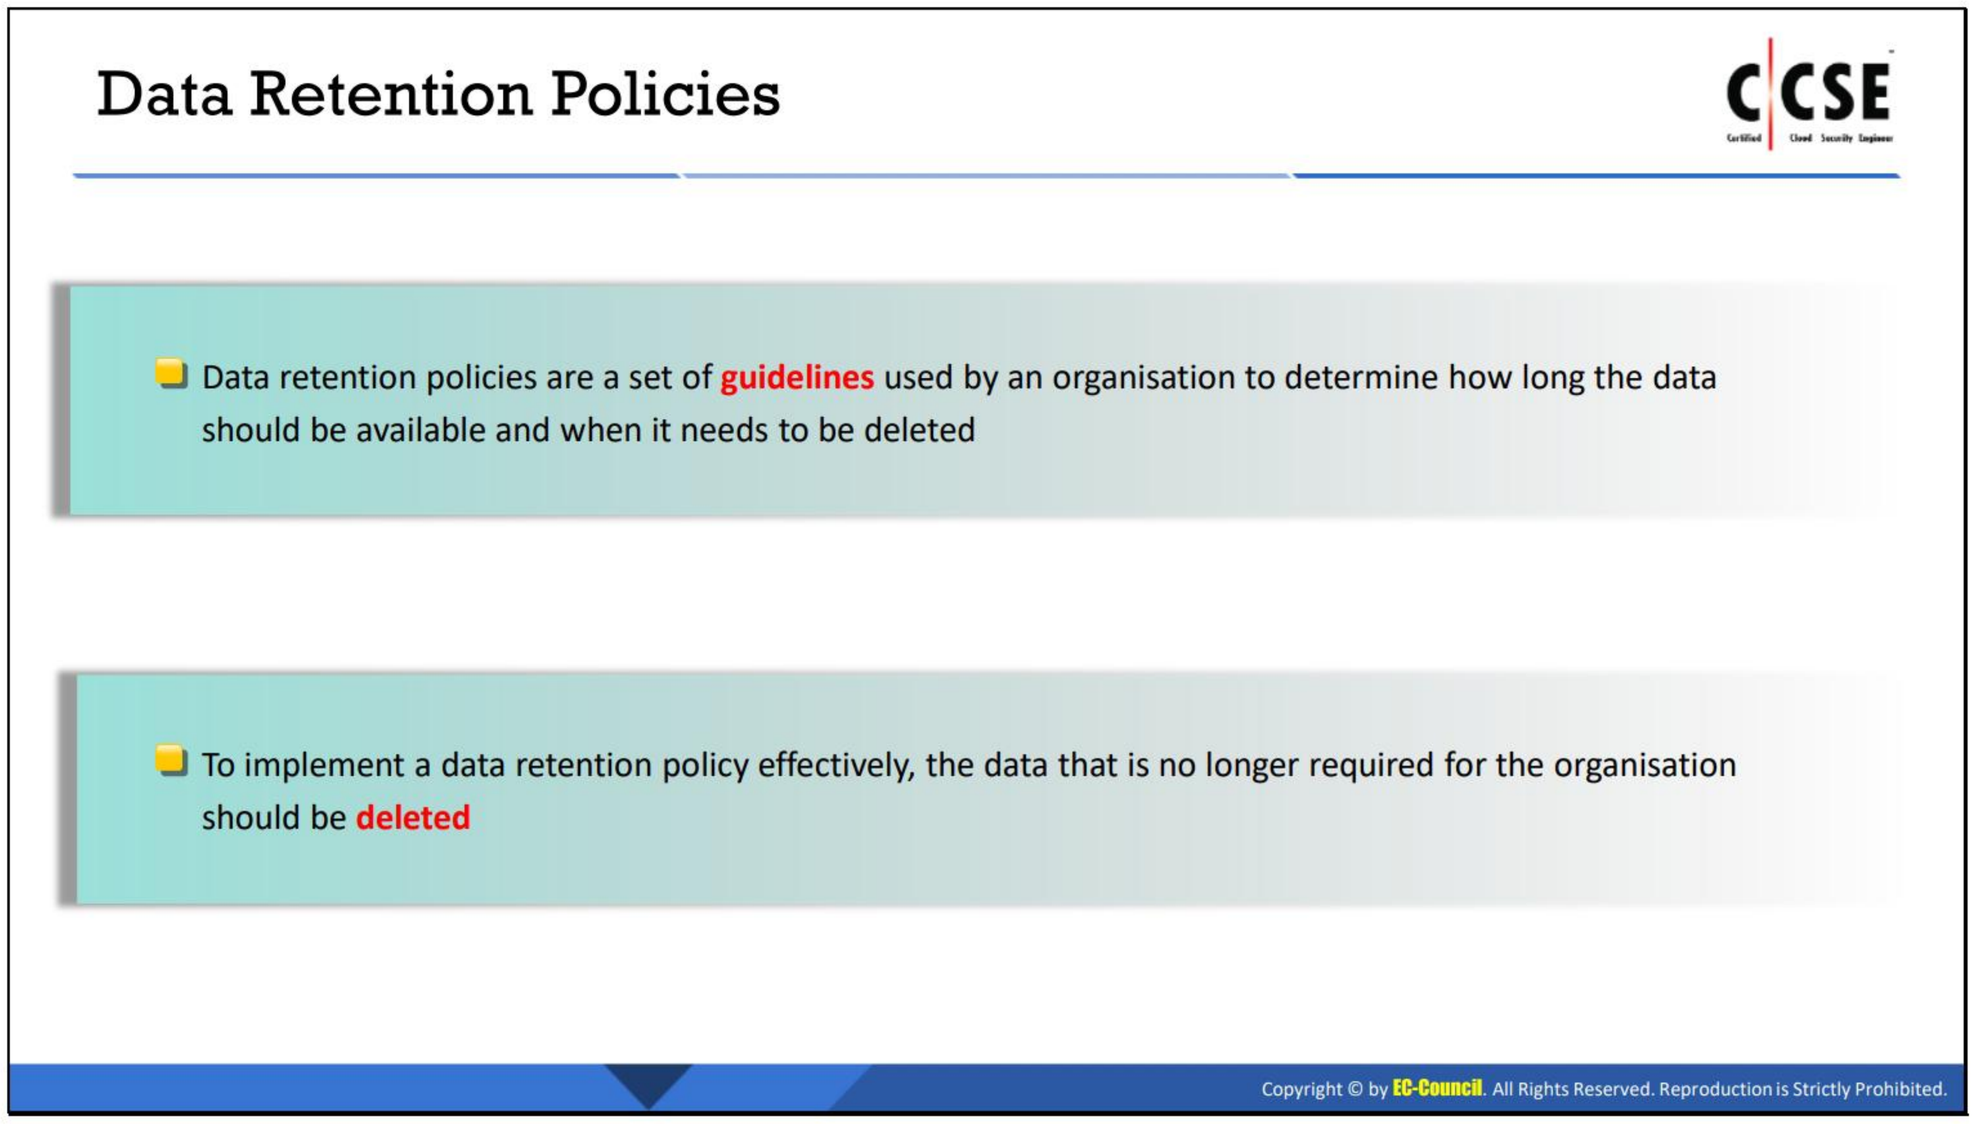


LO#07: Discuss Data Retention and Archiving Strategies

The objective of this section is to explain how to implement data-retention policies, procedures, and mechanisms for deleting the data as well as data archiving procedures to protect SaaS applications and on-premises data.

hide01.ir

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



Data-Retention Policies

Data-retention policies are a set of guidelines used by an organization to determine how long the data should be available and when it needs to be deleted. The main objective of the data-retention policy is to store the data for future reference and organize it in a manner so that it can be easily searched and found whenever required. The policy also includes the disposal of data when it is no longer required.

A data-retention policy must specify the retention period, data types such as names, addresses, financial information, data security, and the data-retrieval procedures. The implementation of the data-retention policy begins by understanding the type of data that the organization stores and then classifying it. All the data-retention policies should ensure that all the local and federal regulations are met. This means that the data should be retained for the specified period and is deleted after completing the specified time.

Every data-retention policy should contain the below components.

- **Legislation, regulation, and standards requirements:** A data-retention policy should consider the data type and should mention the retention period based on the required compliance regimes associated with each type of data.
- **Mapping:** The data-retention policy should include the data mapping method to understand the data types, data formats, and data locations.
- **Data classification:** The data classifications should be made based on their value such as ownership, business usage, and compliance requirements. Based on the classification, the organization would choose the right retention procedures.

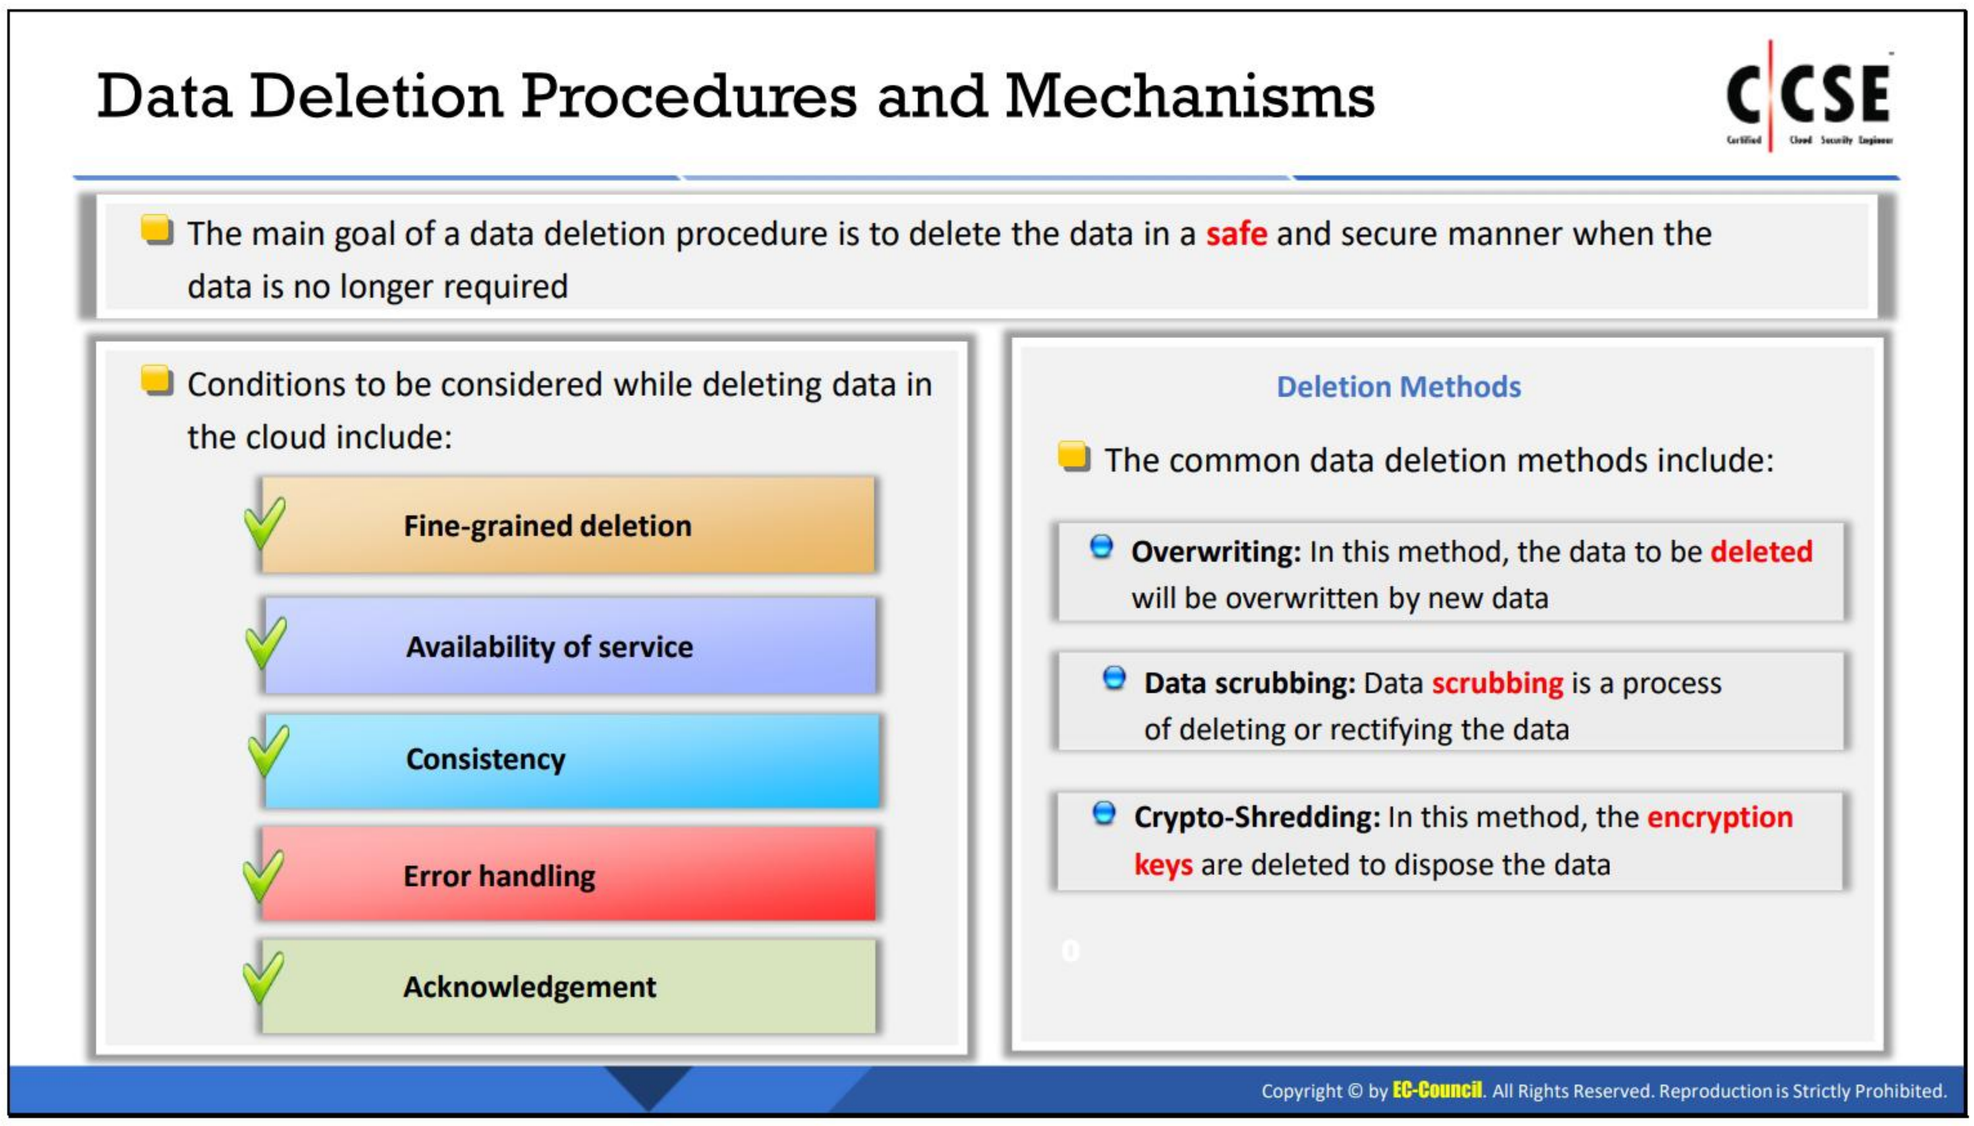
Certified Cloud Security Engineer Copyright © by **EC-Council**
All Rights Reserved. Reproduction is Strictly Prohibited.

- **Data-retention procedure:** A retention procedure should be followed for each data type based on the data-retention policy. The policy should mention how long the data should be stored, its physical location, and how it is stored. The policy should be implemented using a data-retention procedure. The data-retention procedure should also mention backup options and restore procedures.
- **Monitoring and maintenance:** A data-retention policy should also follow procedures to ensure that the process is working well without any changes.

  Many organizations assume that storing the data for a long time than required would be secure so that it can be used if needed. However, storing this misconception can increase the chances of a data breach. To implement a data-retention policy effectively, the data that is no longer required for the organization should be deleted. If the organization stores, collects, and shares the data on regular basis, the data-retention policy should be re-evaluated to ensure security.

hide01.ir

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



Data Deletion Procedures and Mechanisms

The main goal of a data deletion procedure is to delete the data in a safe and secure manner when the data is no longer required. Otherwise, it can lead to a data breach or other compliance failures. The data deletion procedures are designed in such a way that there will not be any files or pointers after deletion. This ensures that the data cannot be restored after deletion.

Many organizations use the data deletion policy to ensure that certain records are properly disposed by following all rules and regulations. Many business policies also insist for the safe deletion of data. Data deletion procedures are also used for the deletion of clear text data after the encryption.

Once the data is deleted in the cloud, it would be challenging for an attacker to retrieve the data because the data in the cloud is distributed and stored in different locations with different pointers.

Deleting data in the cloud can be challenging. The below conditions are to be considered while deleting data in the cloud.

- **Fine-grained deletion:** The deletion procedure should delete only the targeted data and should not affect the remaining portion of the data. The remaining data should be easily accessible by the users.
- **Availability of service:** While deleting the data, the working pattern should not be affected. The service should also work for all the cloud users including the user who has requested for the data deletion.
- **Consistency:** When the data is deleted, it should be deleted in all the locations where the original data or copies are stored. The deleted data should not be accessible to anyone.

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

- **Error handling:** The deletion should take place without any errors in a specified timeframe.
- **Acknowledgment:** Once the data is deleted, the cloud user should receive an acknowledgment that the data deletion is successfully completed.

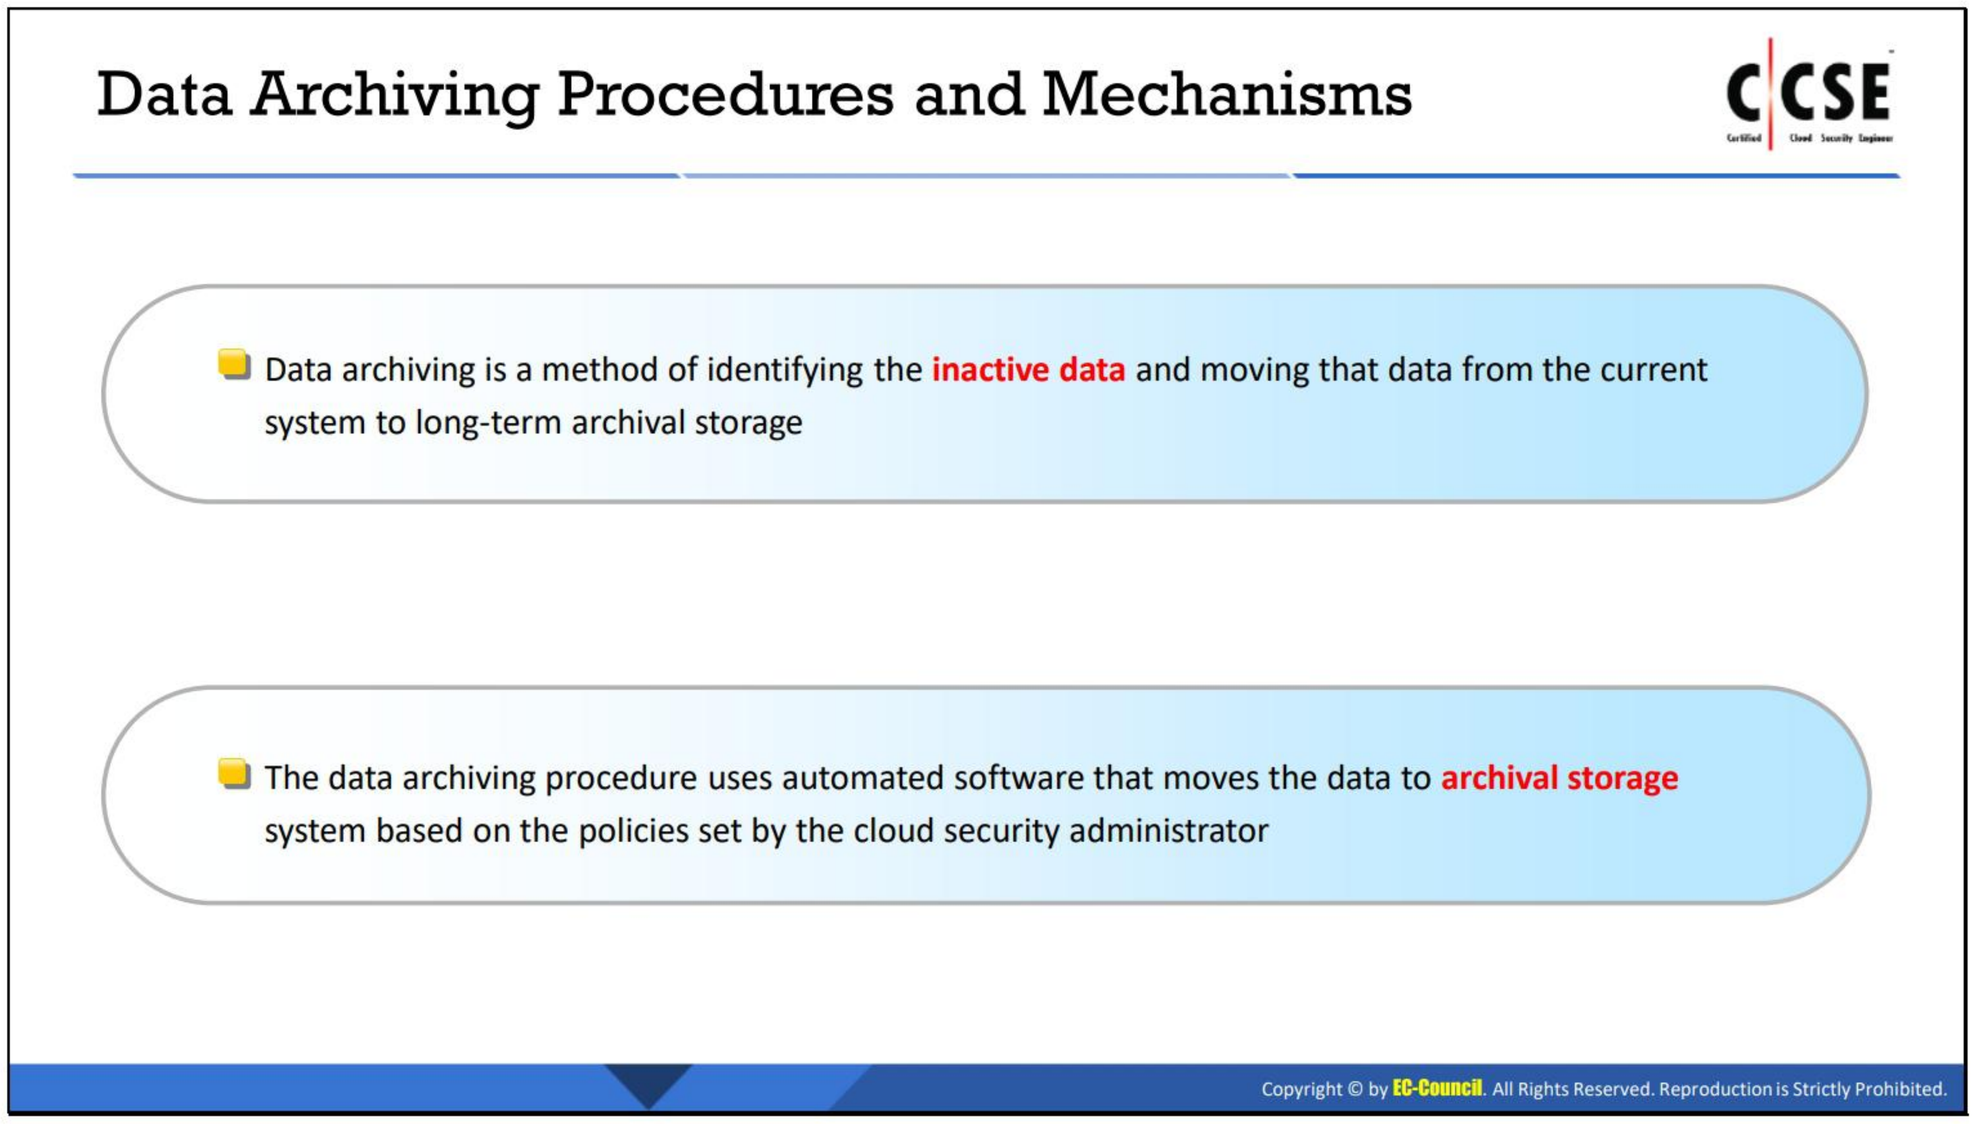
Deletion Methods

It is challenging for the cloud users to delete data in the cloud as the cloud users will not have access to all the locations and they cannot modify the infrastructure. The common data deletion methods are included below.

- **Overwriting:** It is a method in which the data to be deleted will be overwritten by new data. It is a less adopted method as it requires internal modifications, and it is time-consuming. Several repetitions of data overwriting ensure more proper and thorough destruction of the data.
- **Data scrubbing:** Data scrubbing is a process of deleting or rectifying the data. This method ensures the deletion of all the copies of data.
- **Crypto-shredding:** It is the widely used method in which the encryption keys are deleted to dispose the data. In this method, the encryption keys that are used to encrypt the original data are deleted. This makes the data unreadable and unrecoverable. For the successful implementation of crypto-shredding, all the data need to be encrypted without leaving any plain text.

hide01.ir

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



Data Archiving Procedures and Mechanisms

Data archiving is a method of identifying the inactive data and moving that data from the current system to long-term archival storage. Removing the inactive data using data archiving procedures will enhance the performance of the existing resources and enables the users to quickly access the archive storage devices. The data is archived in such a way that it can be restored to service at any time.

The archived data in long-term storage devices provide long-term retention and storage. The data can be stored in any location and it is accessible and provides integrity. It is mainly used by organizations that update the information regularly and still need to retain the old content.

The below elements should be included in a data archiving policy.

- **Data encryption procedure:** To store long-term archive data that is encrypted, it would be challenging for organization to secure the keys. A right encryption policy must be used based on the media used, restoration options, and the threats that need to be handled with encryption.
- **Data tracking procedures:** All the data access and movement should be tracked to ensure that security is applied throughout the lifecycle of the data.
- **Granular retrieval:** The data stored in the archival storage should be in such a way that it can be retrieved whenever required based on factors such as authors, names, subject etc. The archival storage should be able to support eDiscovery to search and retrieve the required data.

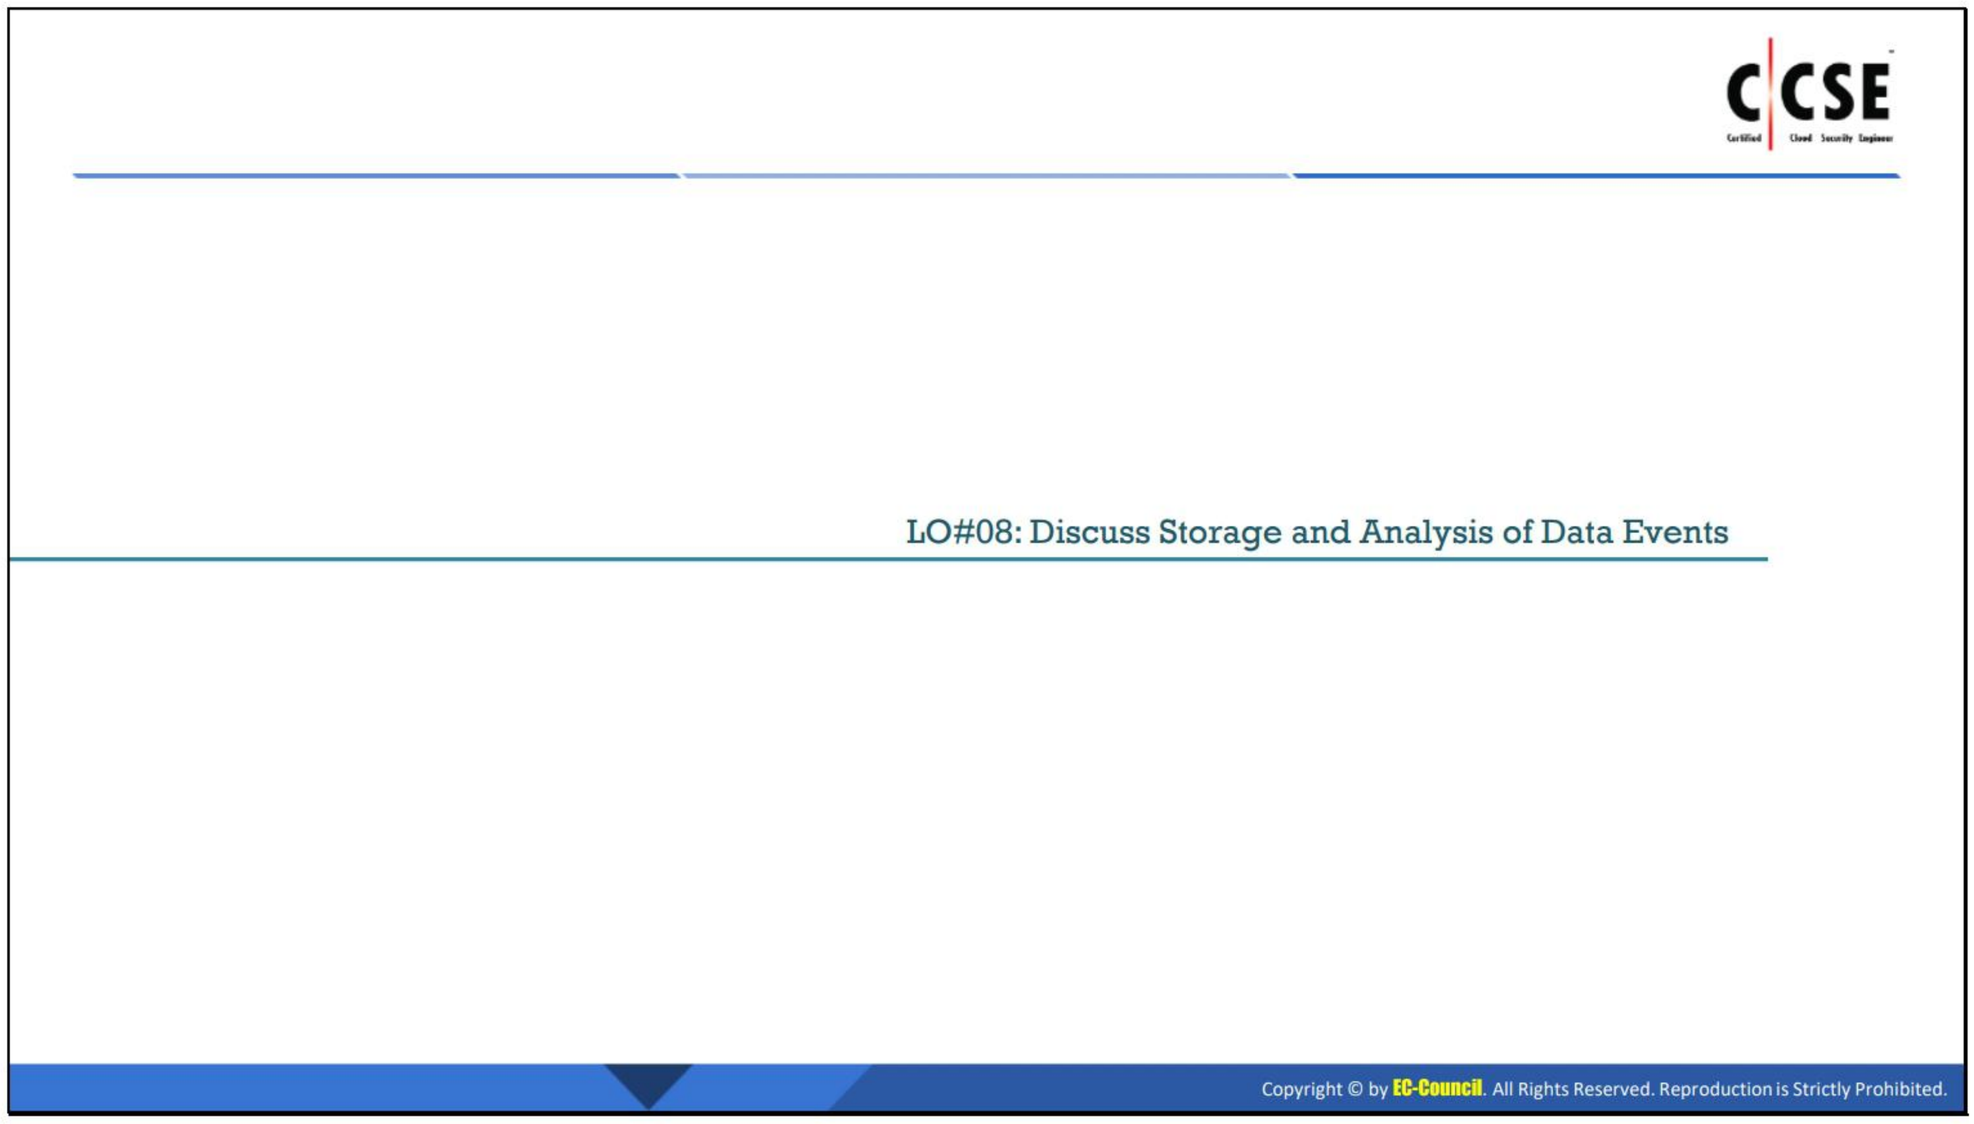
Certified Cloud Security Engineer Copyright © by **EC-Council**
All Rights Reserved. Reproduction is Strictly Prohibited.

- **Backup and disaster recovery:** All the recovery and backup requirements should be specified in the document. The organization should also ensure that the disaster recovery policy is updated.
- **Data formats:** Choosing the right data format is important as the data need to be stored for a long period. The same things should be considered for media storage types.
- **Restoration procedures:** Data restorations procedures should be used to test the data restoration. The test data restoration should be performed in a separate environment to avoid overwriting existing content.

  The data archiving procedure uses automated software that moves the data to archival storage system based on the policies set by the cloud administrator. The archived data will still be accessible by the users at any time and from anywhere.

hide01.ir

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

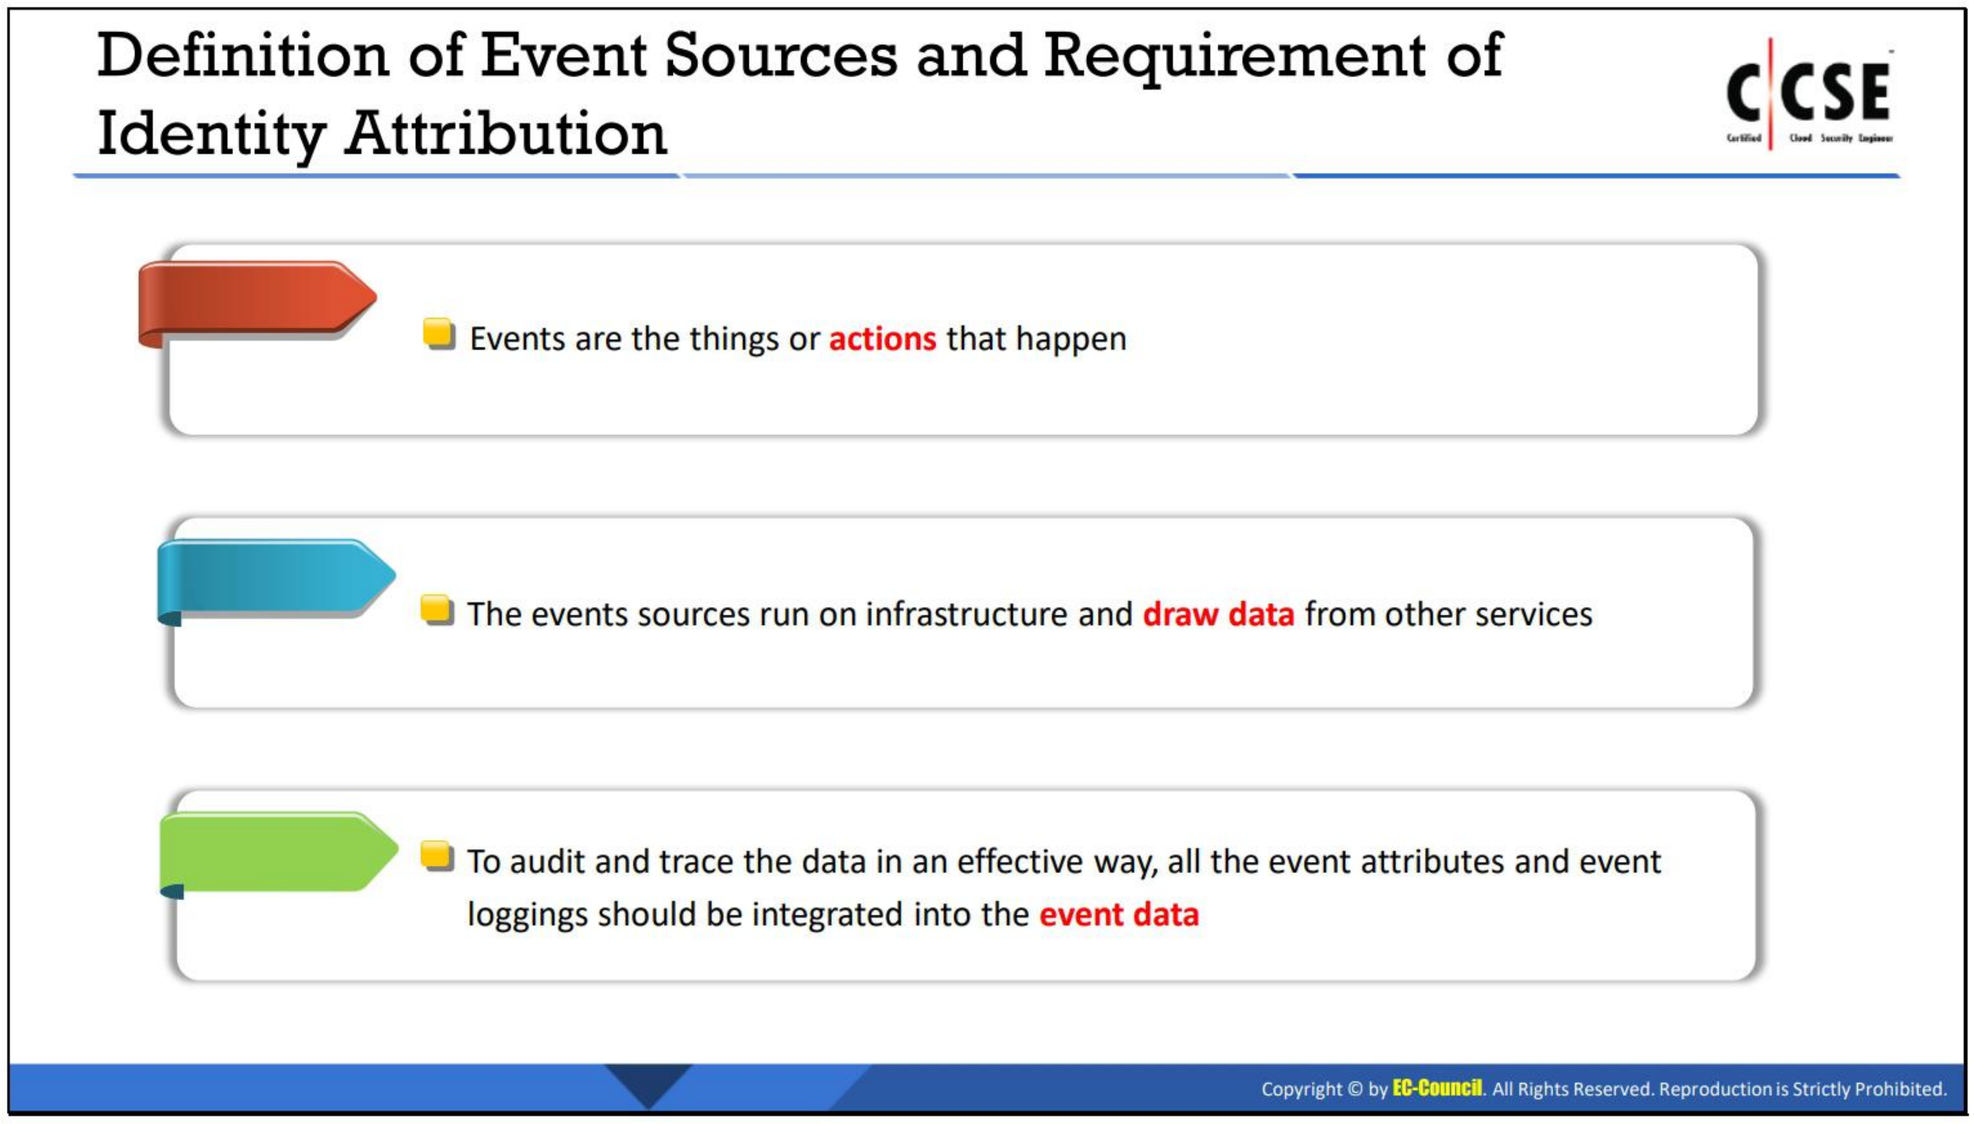


LO#08: Discuss Storage and Analysis of Data Events

This section explains event sources and how to analyze data events. It also explains the important terms related to preserving the event data.

hide01.ir

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



Definition of Event Sources and Requirement of Identity Attribution

Events are the things or actions that happen. A cloud security professional does not need to focus on all the events that happen. Using certain tools, the cloud security professional can filter the events that take place on cloud infrastructure and focus on the relevant and important events.

Events sources run on infrastructure and they draw data from other services. Event attributes define the kind of data that is associated with each type of event. Based on the number of events tracked, a large amount of data will be produced. The collected data will be stored and analyzed to determine any threats or vulnerabilities in the system that need to be focused on.

Event Sources

The event sources from which the data is retrieved depend on the cloud service models used by the organization. The commonly used event sources are SaaS, PaaS, and IaaS. It is important to understand briefly about each type of data source.

SaaS Event Source: SaaS is a service where all the applications will be delivered as services through the internet. SaaS reduces the burden of maintaining complex hardware and software and the users can access the services without installing and maintaining the software.

SaaS provides minimal control and access to the event and diagnostic data. A cloud security professional cannot view the infrastructure level logs. Therefore, to maintain proper investigation capabilities, auditability, and traceability of data, the cloud user needs to mention all the data access requirements in the cloud SLA.

The data sources used for event investigations include the following:

- Application server logs
- Web server logs

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

- Guest operating system logs
- Database and hot access logs
- SaaS portal logs
- Billing records
- Network captures

PaaS Event Sources: PaaS is a cloud computing service that provides a platform for the users. The users can develop, run, and manage their business applications without the need of maintaining the infrastructure. The underlying infrastructure will not be visible to the users and it will be maintained by the cloud service provider.

In PaaS environment, users can have access and control over the event and diagnostic data. The cloud security professionals can view few infrastructure level logs and complete application logs as the applications are built and monitored by the organization. The developers can decide the level of application data that needs to be extracted and monitored.

The cloud security professionals should work with the development team to understand the features of applications that are under development. This helps them to design and implement monitoring regimes to maintain auditability and traceability on data.

IaaS Event Sources: In IaaS, the cloud service provider offers networking, servers, data storage, and virtualization capabilities. The customer will gain access to the data storage and compute power as per their requirements, but the customer should provide their own software platform to run on it such as OS, middleware, data, runtime, and applications.

All the applications of the organization can be used inside the infrastructure provided by the service provider. Based on the organization's processing and storage needs, the infrastructure can be scaled.

All the infrastructure and application-level logs will be visible to the cloud security professional. It is recommended for the cloud users to specify the data access requirements in the SLA, so that the user can maintain auditability and traceability over data.

To perform an effective audit, the cloud security professional should examine the following logs:

- DNS server logs
- API access logs
- Virtual Machine Monitor (VMM) logs
- Management portal logs
- Billing records
- Cloud network logs

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

Requirement of Identity Attribution

To audit and trace the data in an effective way, the event logs must contain the information relevant to the processes that need to be examined. For this, all the event attributes and event loggings should be integrated into the event data. This includes the following:

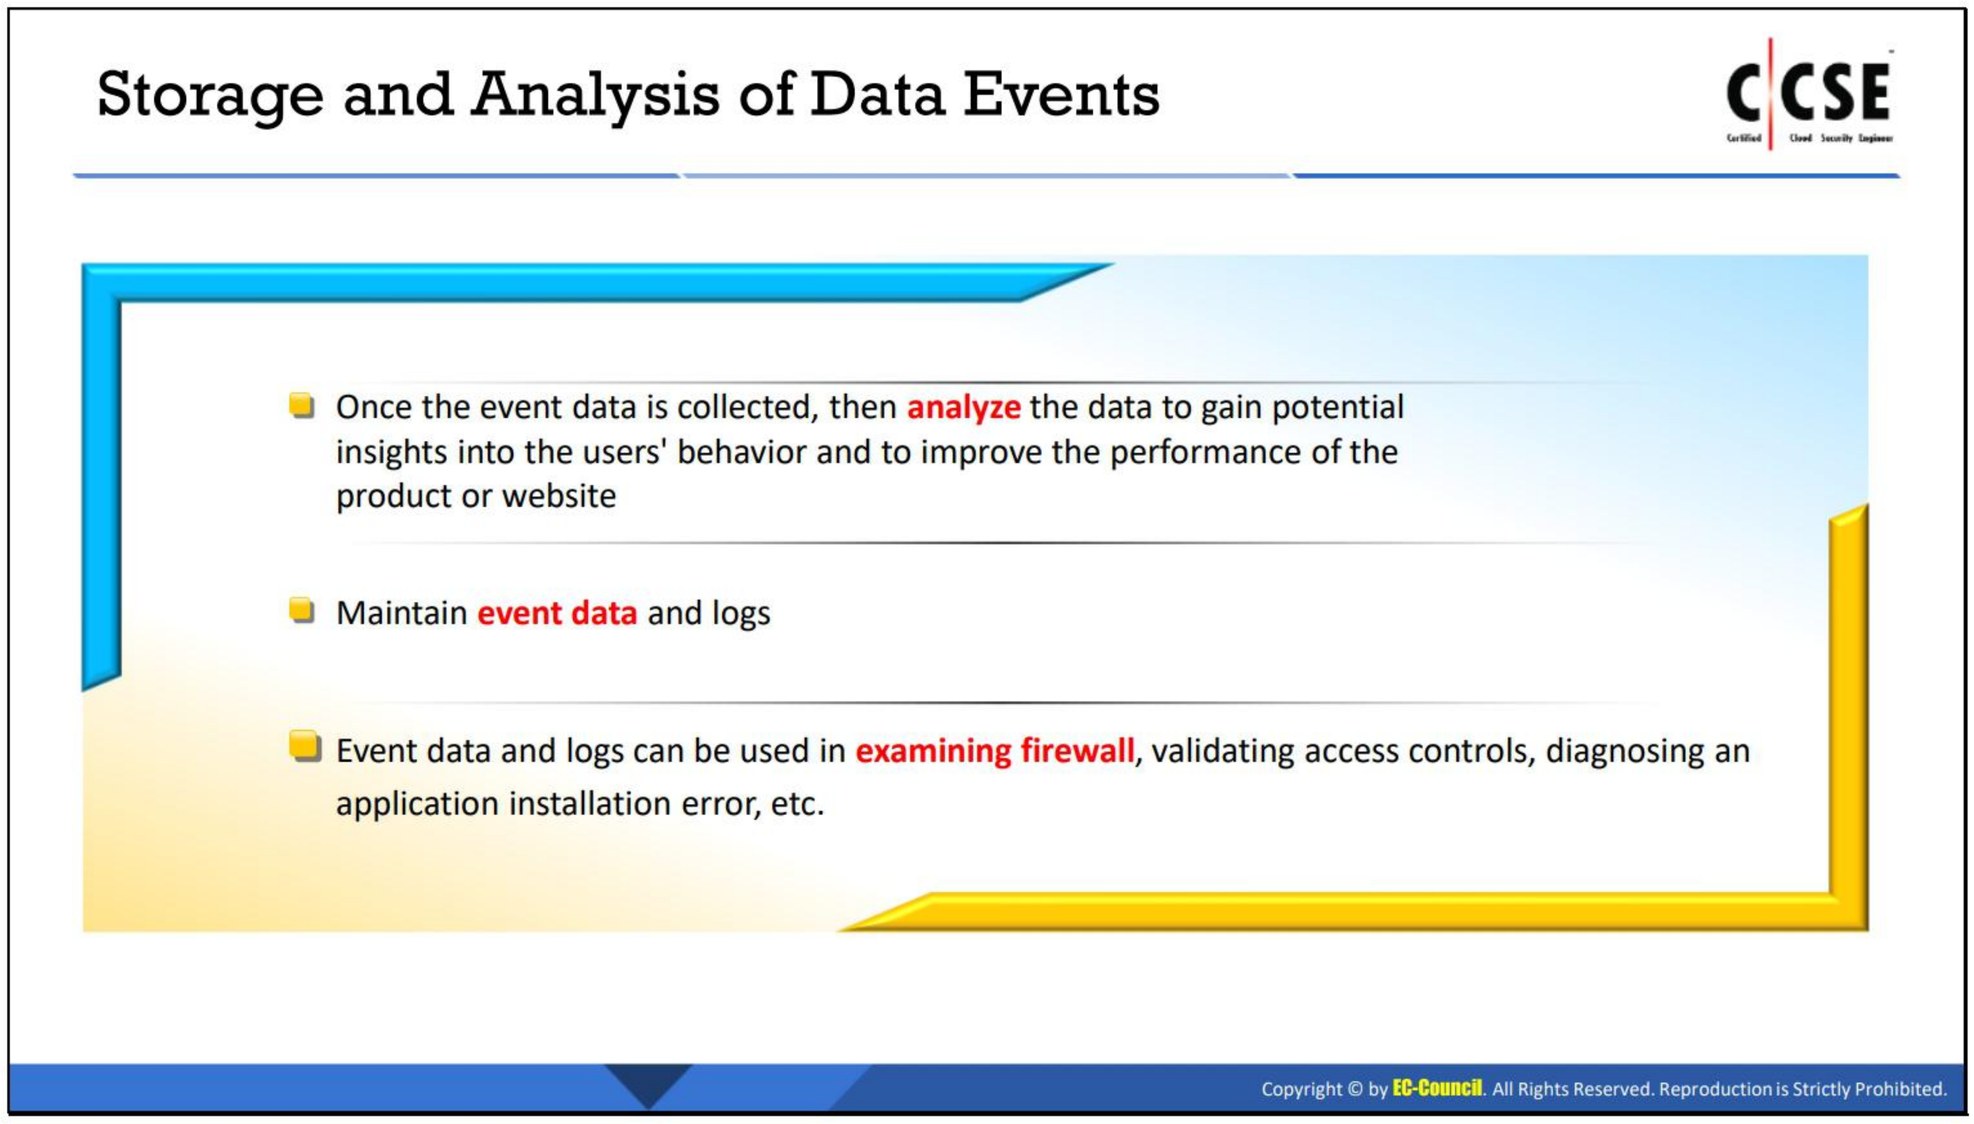
- The event date and time which is different from log time
- Log date and time
- Interactive identifier
- Service name and protocol
- Geolocation
- Application address and an identifier such as server address, port number, name, version, etc.
- Code location
- Source address along with user's IP address and user's device identifier
- User identity
- Type and severity of the event
- Security-related event flags

The following additional event attributes can also be integrated into the event data:

- GPS event date and time
- Action and object
- Result status to check whether the action performed on the object is successful
- HTTP status code
- Responses seen by the users such as session termination and custom text messages

hidedl.ir

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



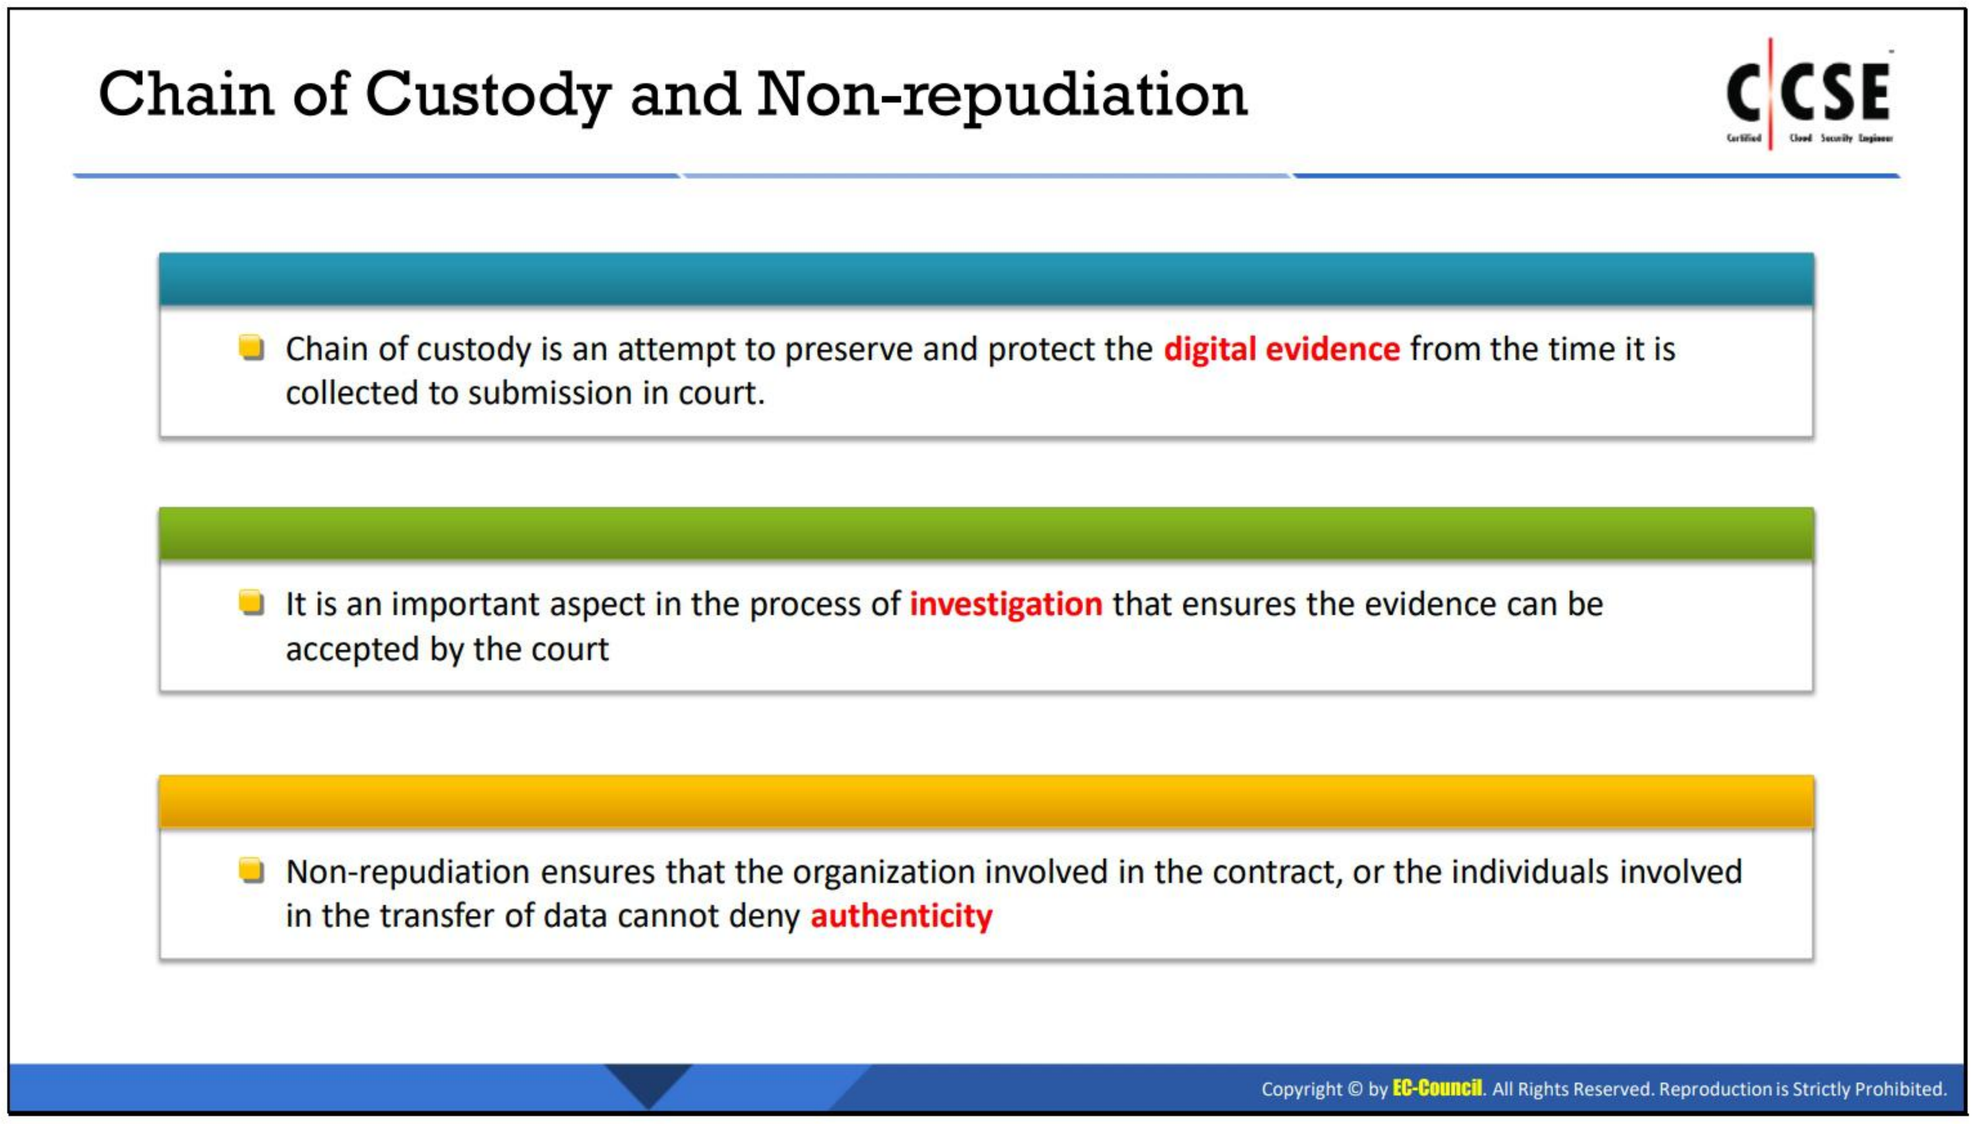
Storage and Analysis of Data Events

Based on the amount of data collected, archiving, and maintenance, event and log data storage and analysis would be costly. The cloud administrator should consider the cost and company regulatory requirements while planning for the storage of event data. Data preservation refers to the storage and maintenance of safety and integrity of the data being stored. Evidence preservations assures that it would be acceptable by the court though it can be easily destroyed or modified. The backlog for forensic laboratories will range from 6 months to one year. The storage of digital evidence would require significant time before it is analyzed. Strict access controls will be required to protect the data from accidental modifications. Additionally, event logging mechanisms should be tamper-proof to reduce the chances of fake event logs.

Once the data is collected, analyzing the data is an important aspect. Analyzing the event data helps in gaining potential insights into the users' behavior. It also helps to improve the performance of the product or website. Though analyzing event data seems to be simple, it is tricky and right tools are required to analyze the event data. Using the right tools, the actions of the customers can be tracked so that the developers can monitor them and test new ideas. This helps to improve the performance of the product. Some of the tools include Looker, Tableau, Periscope, and Mode Analytics.

Maintaining event data and logs is recommended as the data can be used in various scenarios such as examining firewall, validating access controls, and diagnosing an application installation error, etc.

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



Chain of Custody and Non-repudiation

Chain of custody is an attempt to preserve and protect the digital evidence from the time it is collected to submission in court. It is an important aspect in the process of investigation that ensures the evidence can be accepted by the court. To consider the digital evidence as admissible by the court, documentation should be provided for collecting, possessing, transferring, accessing an item, and also analyses performed on the collected data till the date of final disposition. The chain of custody also documents the details of people who have collected, handled, or transferred including the date and time of these activities. It provides trust to the court that the evidence provided has not tampered.

It would be challenging for cloud security professionals to create a verifiable chain of custody in cloud platforms as there are multiple data centers across the world. The best way to provide the chain of custody is to include it in the service agreement itself. This ensures that the cloud service provider will agree to the requests related to the chain of custody. The chain of custody includes all the individuals who are involved in acquiring, processing, and analyzing the evidence.

Stages of Chain of Custody

Chain of custody involves four stages, as mentioned below.

- **Data collection:** It is the stage where the chain of custody process is initiated. The examiner should identify and label each item acquired and should also document all the evidence-related information such as how, who, and when it is created, where it is stored, and who have access to that data.
- **Examination:** During the examination stage, all the information related to the chain of custody is documented including all the forensic steps taken in the process. It is vital to take screenshots to show that the tasks are completed.

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

- **Analysis:** It is the result of the examination stage. Using legal methods and techniques, the useful information will be extracted.
- **Reporting:** It is the stage where chain of custody is documented which includes the source of data, extraction methods used, analysis process, tools used, issues faced, and how they are handled. This is the document that should show that chain of custody is maintained throughout the process and evidence submitted is legally accepted.

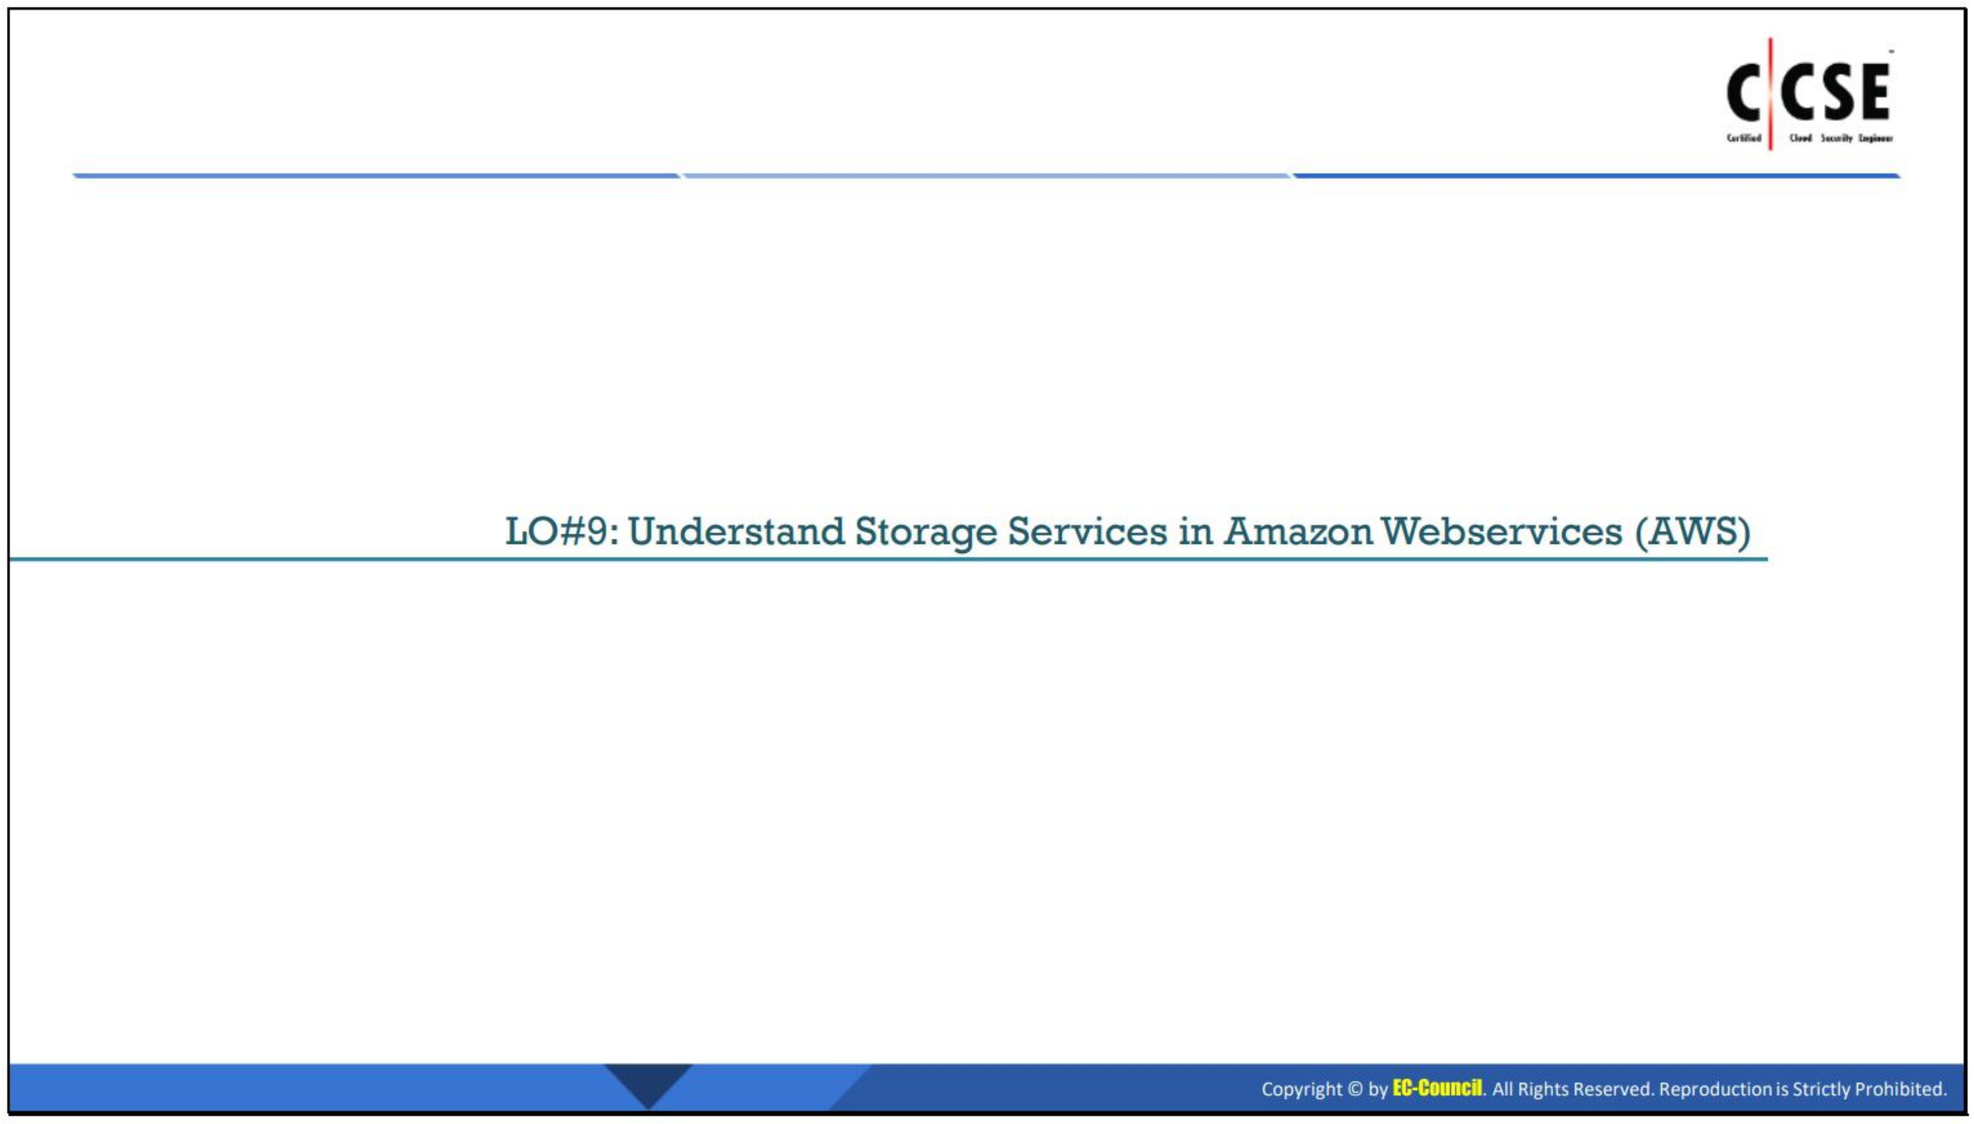
Non-repudiation

Non-repudiation refers to the assurance that the owner of the information provides valid proof and cannot deny the validity. It also provides protection against a person who denies after performing a particular action such as creating data, sending, and receiving message, and approving the information. It also refers to the service that provides proof for the origin and integrity of data.

Non-repudiation makes it difficult to deny in terms of who or from where the data has been collected and the authenticity of the data. In terms of information security, non-repudiation is used in case of a formal contract, any communication, or transfer of data. Non-repudiation ensures that the organization involved in the contract or the individuals involved in the transfer of data cannot deny authenticity.

hide01.ir

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

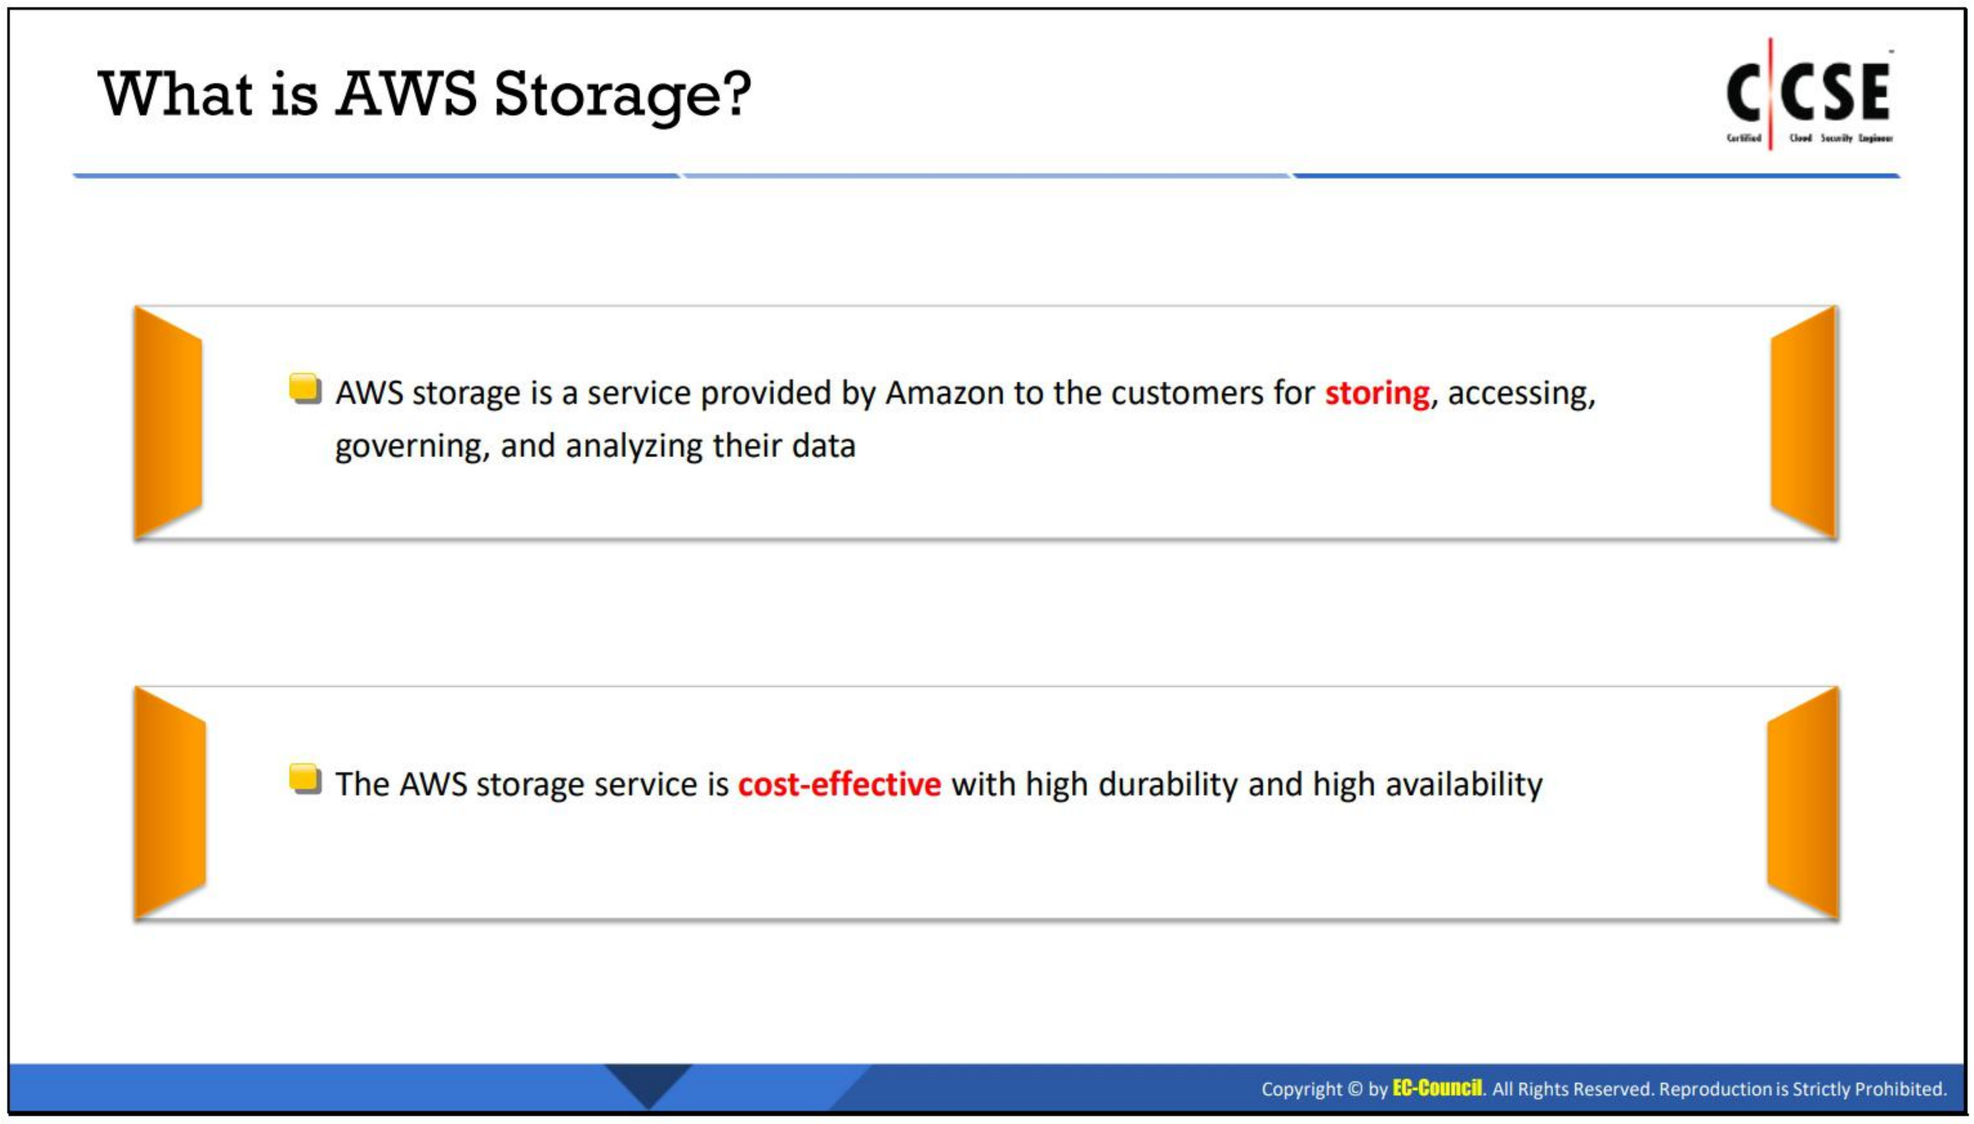


LO#09: Understand Storage Services in Amazon Webservices (AWS)

The objective of this section is to explain the implementation of data security in AWS using different features such as S3 default encryption, customer-provided AWS KMS CMKs, and geo-restriction, etc.

hide01.ir

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



What is AWS Storage?

AWS storage is a service provided by Amazon to customers for storing, accessing, governing, and analyzing their data. The AWS storage service is cost-effective with high durability and high availability. There are various choices provided by AWS for archiving the information, disaster recovery, and backing up.

AWS storage services are listed below.

- **Object storage**: Amazon Simple Storage Service (S3)
- **File storage**: Amazon Elastic File System, Amazon FSx for Windows File Server, Amazon FSx for Lustre
- **Block storage**: Amazon Elastic Block Store
- **Backup**: AWS Backup
- **Data transfer**: AWS Storage Gateway, AWS DataSync, AWS Transfer Family, AWS Snow Family
- **Edge computing and storage**: AWS Snow Family

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

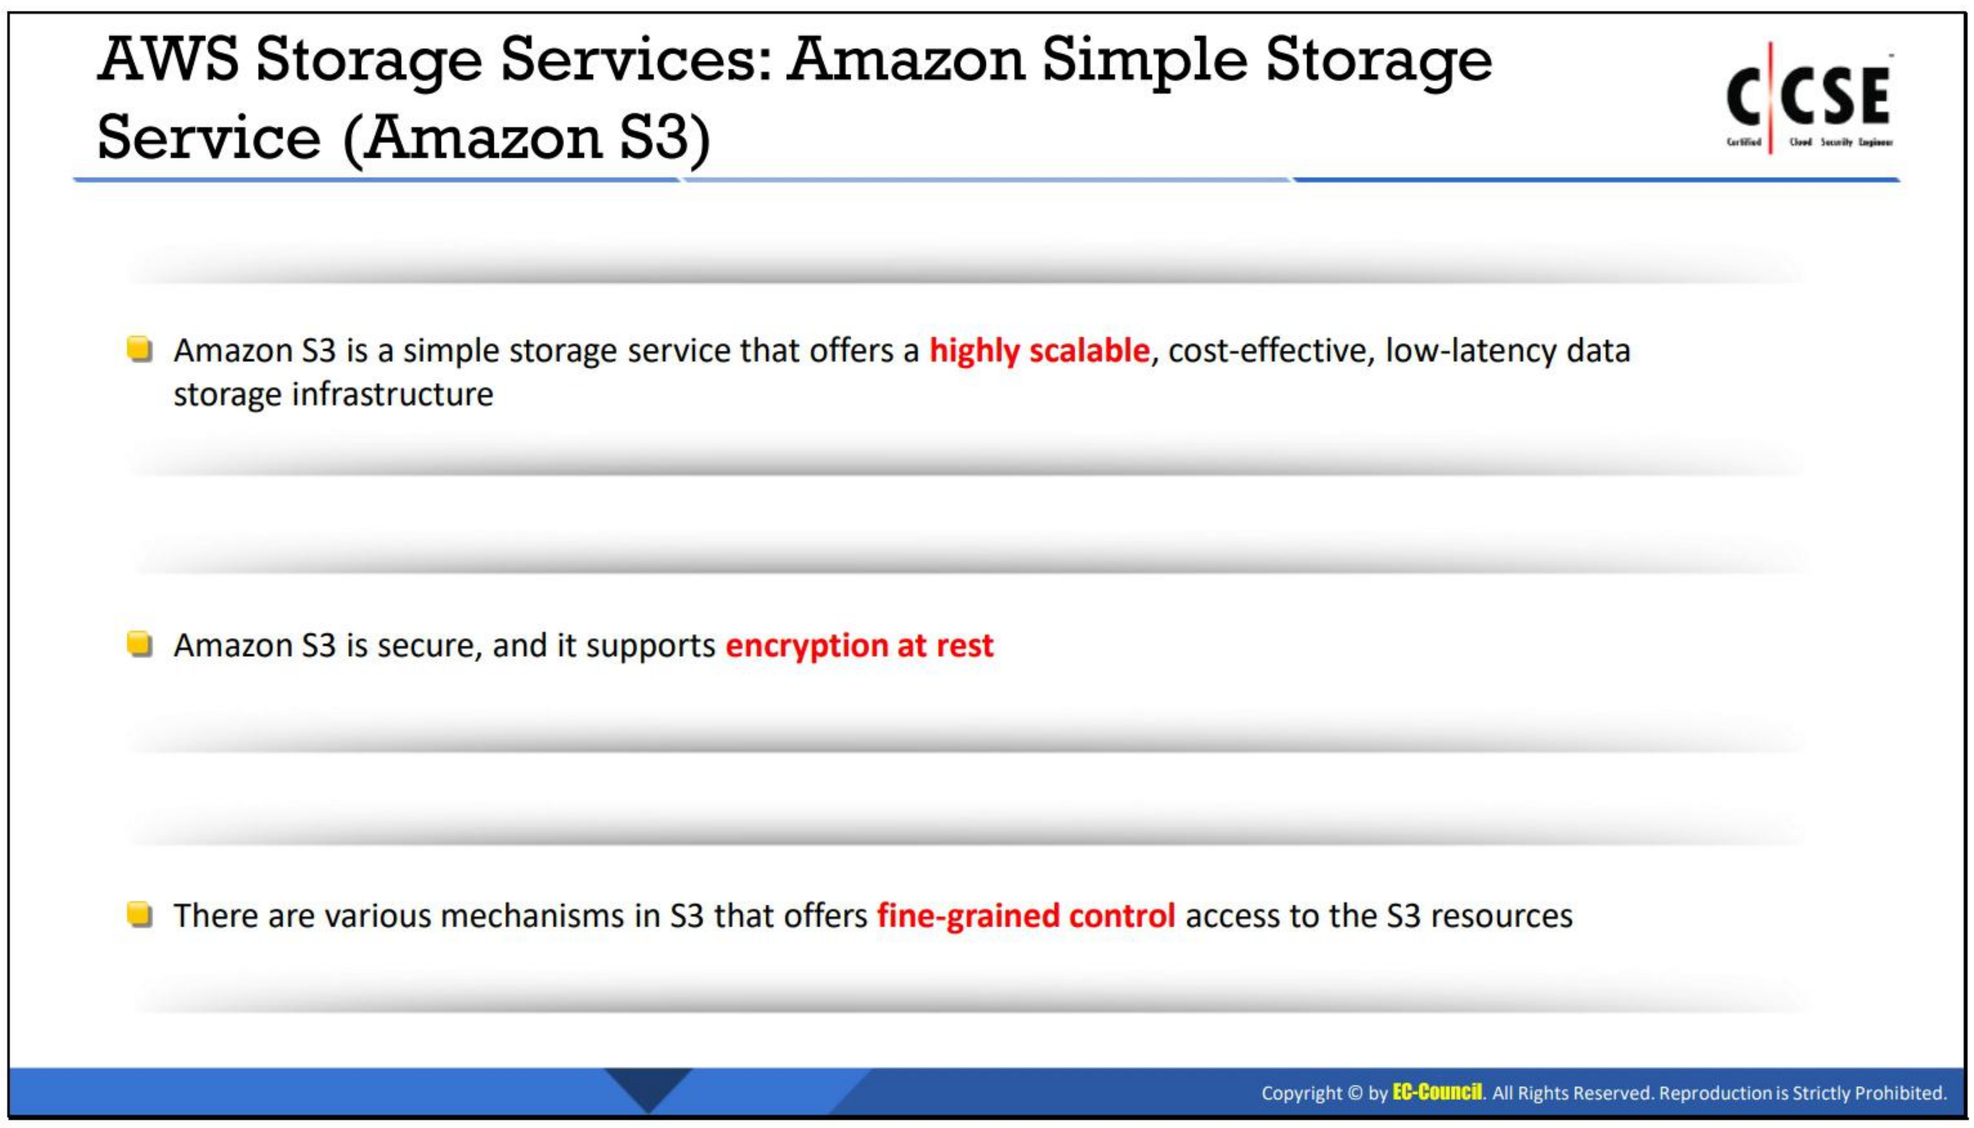
Features of AWS Storage

The features of AWS storage are listed below.

- **Durability and availability**: AWS storage services provide high durability and high availability.
- **Security**: All AWS storage services are highly secure.
- **Accessibility**: The AWS user can easily access the AWS storage service from any location.
- **Cost-effective**: The AWS users need to pay only what they use.
- **Performance**: The performance of AWS storage service depends on the selected storage service. For example, in Simple Storage Service (S3), if you want to access S3 within Amazon EC2 of the same region, the performance will be fast. The performance of Amazon Glacier service is slow, but it is an extremely low-cost amazon storage service.
- **Scalability and elasticity**: Most AWS storage services are designed to provide a high level of scalability and elasticity automatically. All the AWS storage services can meet AWS user business requirement.
- **Disaster recovery**: Most AWS storage services provide disaster recovery for business continuity.

hide01.ir

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



AWS Storage Services: Amazon Simple Storage Service (Amazon S3)

Amazon S3 is a simple storage service that provides a highly scalable, cost-effective, low-latency data storage infrastructure. Using Amazon S3, the cloud consumer can save and recover their data from any part of the world. The AWS user can store an unlimited number of objects in Amazon S3 buckets and can read, write, or delete objects having data that ranges from 1 byte to 5 terabytes.

Amazon S3 is secure, it supports encryption at rest, and there are various mechanisms in S3 that offer fine-grained control access to the S3 resources. Amazon S3 is highly scalable, as it allows multiple clients to have read and write access concurrently to S3 data. With the data lifecycle management capabilities of Amazon S3, the AWS user can define rules to archive S3 data to Glacier automatically or delete the data.

Generally, the Amazon S3 service is used for storing and distributing static web content and media. As every object in Amazon S3 has a distinct HTTP URL address; therefore, this content can be delivered directly from Amazon S3. Due to Amazon S3 elasticity, it is specifically useful for hosting web content having fluctuating bandwidth demands. The cloud consumer can utilize Amazon S3 as a data store to perform large-scale data analytics. Amazon S3's horizontal scalability allows the AWS user to access the data concurrently from various computer nodes without any latency. Accessing Amazon S3 within Amazon EC2 of the same region is fast.

Amazon’s S3 versioning capability safeguards the data against accidental deletion or application failure. It serves as a highly durable and highly secure solution for backing up critical data and to deliver disaster recovery solution for business continuity.

In the AWS platform, Amazon S3 storage offers the highest level of data availability and durability by synchronously and automatically storing the data in multiple devise and facilities of the

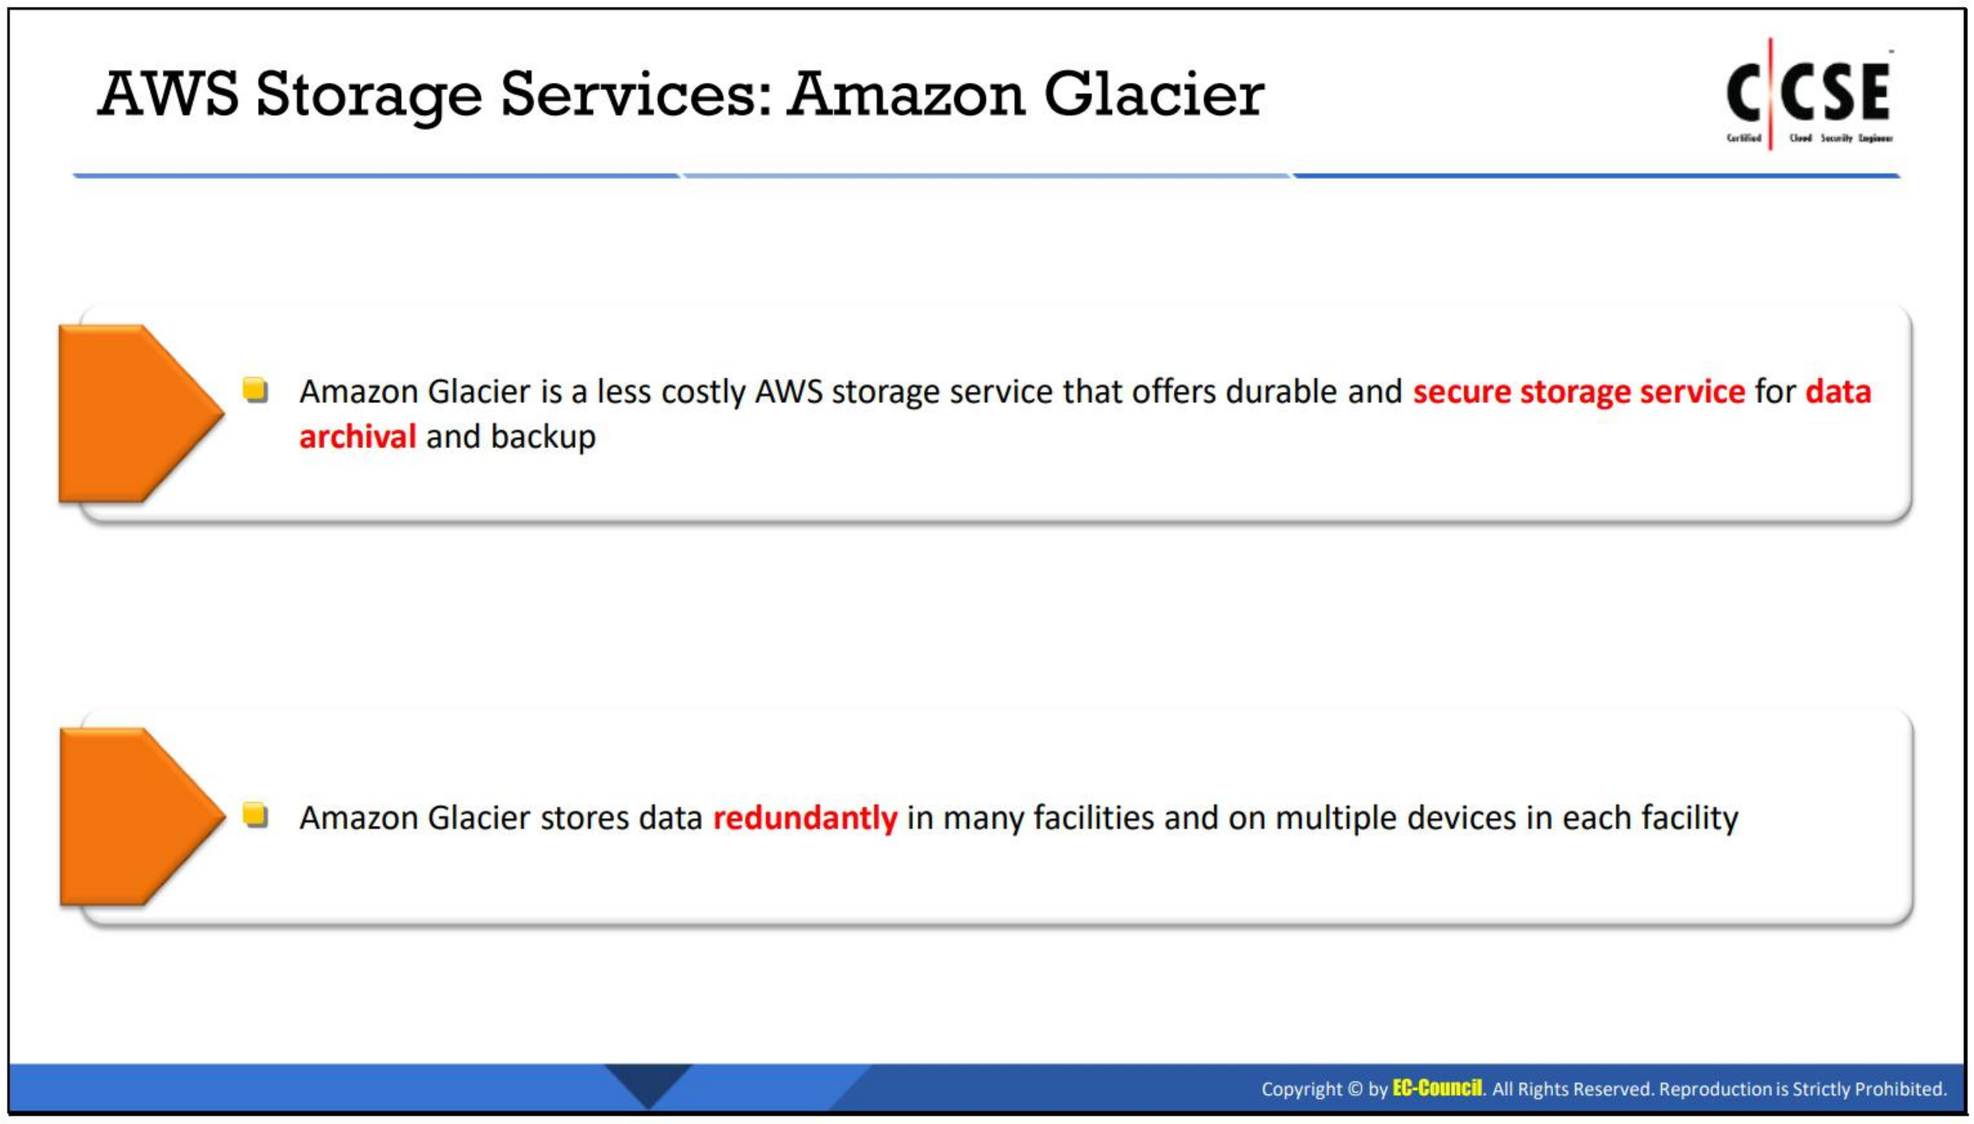
Certified Cloud Security Engineer Copyright © by **EC-Council**
All Rights Reserved. Reproduction is Strictly Prohibited.

selected geographical location. There is no minimum fee to use Amazon S3, the AWS users need to pay for only what they use. Amazon S3 provides robust scalability and elasticity automatically. The AWS consumer can store unlimited number of files virtually in any Amazon S3 bucket.

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. Through versioning and replicating data within an AWS region across several availability zones, it offers 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. By using 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 to recover from outages that happen due to natural disasters, human errors, and system failures.

hide01.ir

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



AWS Storage Services: Amazon Glacier

Amazon Glacier is a less costly AWS storage service that offers durable and secure storage service for data archival and backup. Amazon Glacier is designed for storing data that is not frequently accessed and long-lived. It meets the customers' growing and unpredictable requirements of data storage.

hidn01.ir

The customer data is stored in Amazon Glacier as archives. The archive may contain one or multiple files that are uploaded as single archive. The AWS user can organize their archives in vaults. Utilizing the AWS IAM service, the AWS user can control access to the vaults.

Amazon Glacier can be used with various Amazon web services. The AWS user seamlessly moves the data between Amazon S3 and Amazon Glacier. Amazon Glacier can be used for archiving offsite enterprise data, research and scientific data, media assets, magnetic tape replacement, and digital preservation.

Amazon Glacier stores data redundantly in many facilities and on multiple devices in each facility. Amazon Glacier performs a systemic data integrity check regularly and it has the capability of automatically self-healing.

If customers want to use Amazon Glacier, then they have to pay only what they use; there is no minimum fee to use this service. Amazon Glacier is designed expecting that the retrievals are not frequent and the data will be stored for a long time. 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. Based on the AWS user storage required, Amazon Glacier automatically scales the storage up or down. There is no limit for the data to be stored in Amazon Glacier.

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



AWS Storage Services: Amazon Elastic Block Store (Amazon EBS) Volume

Amazon EBS service is a block storage service that is designed to use with Amazon EC2 instances (VMs). It is a NAS that exists independently from the EC2 instance. The AWS user can use Amazon EBS volume to boot an Amazon EC2 instance. Many EBS volumes can be attached to a single EC2 instance.

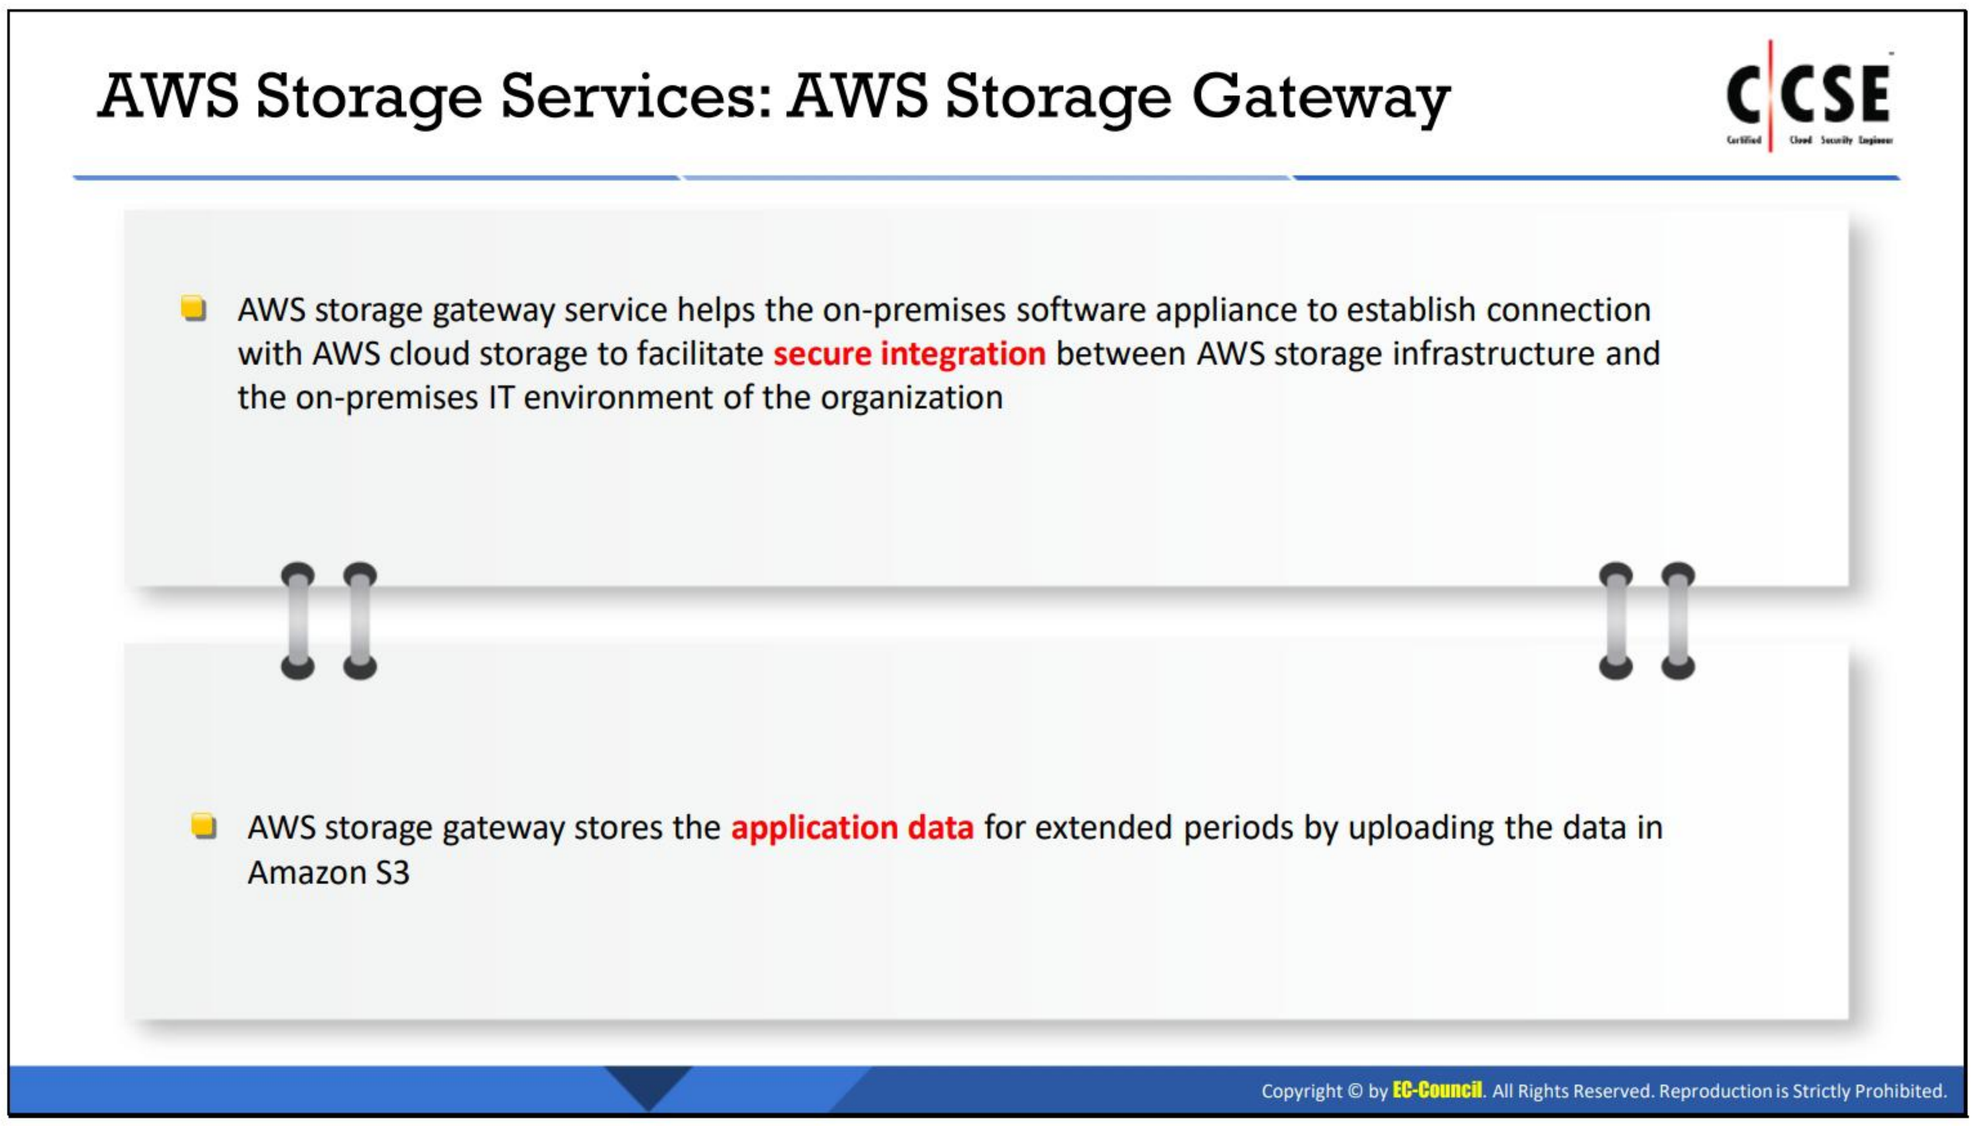
hidoo1.ir

Amazon EBS is designed for data that changes frequently and needs long-term persistence. AWS users can utilize Amazon EBS as primary storage for database and file system. Amazon EBS offers two volume types, i.e., standard volume and provisional IOPS volume. They differ in pricing and storage performance. The standard volume provides low-cost storage for applications with moderate I/O requirements whereas provisional IOPS provide high performance to I/O intensive workloads. The total load on a shared network and other network I/O performed by an instance affect the performance of individual EBS volume.

Amazon EBS 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 EBS volume is highly available and reliable. Its durability relies on volume size and the data that has changed since the last snapshot. To enhance the durability and availability of Amazon EBS data, the AWS user should frequently create snapshots of EBS volume. The AWS user can copy the EBS snapshots from one region to another and can share the EBS snapshots with other user accounts. Thus, Amazon EBS snapshots serve as a disk clone/image for sharing, backup, and disaster recovery.

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



AWS Storage Services: AWS Storage Gateway

This service helps an on-premises software appliance to connect with AWS cloud storage in order to facilitate safe and secure integration between AWS storage infrastructure and the on-premises IT environment of the organization. All industry-standard storage protocols are supported by AWS Storage Gateway. AWS Storage Gateway securely stores all the data encrypted in Amazon S3 and frequently accessed on-premises data; thus, it offers low-latency performance.

The cloud administrator can download the AWS Storage Gateway's software appliance as a VM image and install it on a host in the datacenter. After installing the gateway, associate it with the AWS account through AWS activation process. Create gateway-stored volumes, gateway virtual tape library (VTL), or gateway-cached volumes by utilizing the AWS Management Console. Each of these volumes can be mounted by on-premises applications as an iSCSI device.

AWS Storage Gateway is used in sharing corporate files, making the on-premises backup applications store the primary backup on Amazon S3, data mirroring, and disaster recovery.

AWS Storage Gateway provides disaster recovery with three configurations.

- **Gateway-cached volumes**: It allows storage of primary data into Amazon S3 and storage of frequently accessed data locally (on-premises). 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, but it is asynchronously backed up to Amazon S3. This configuration is used when users need low-latency access to the full data set.

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

- **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).

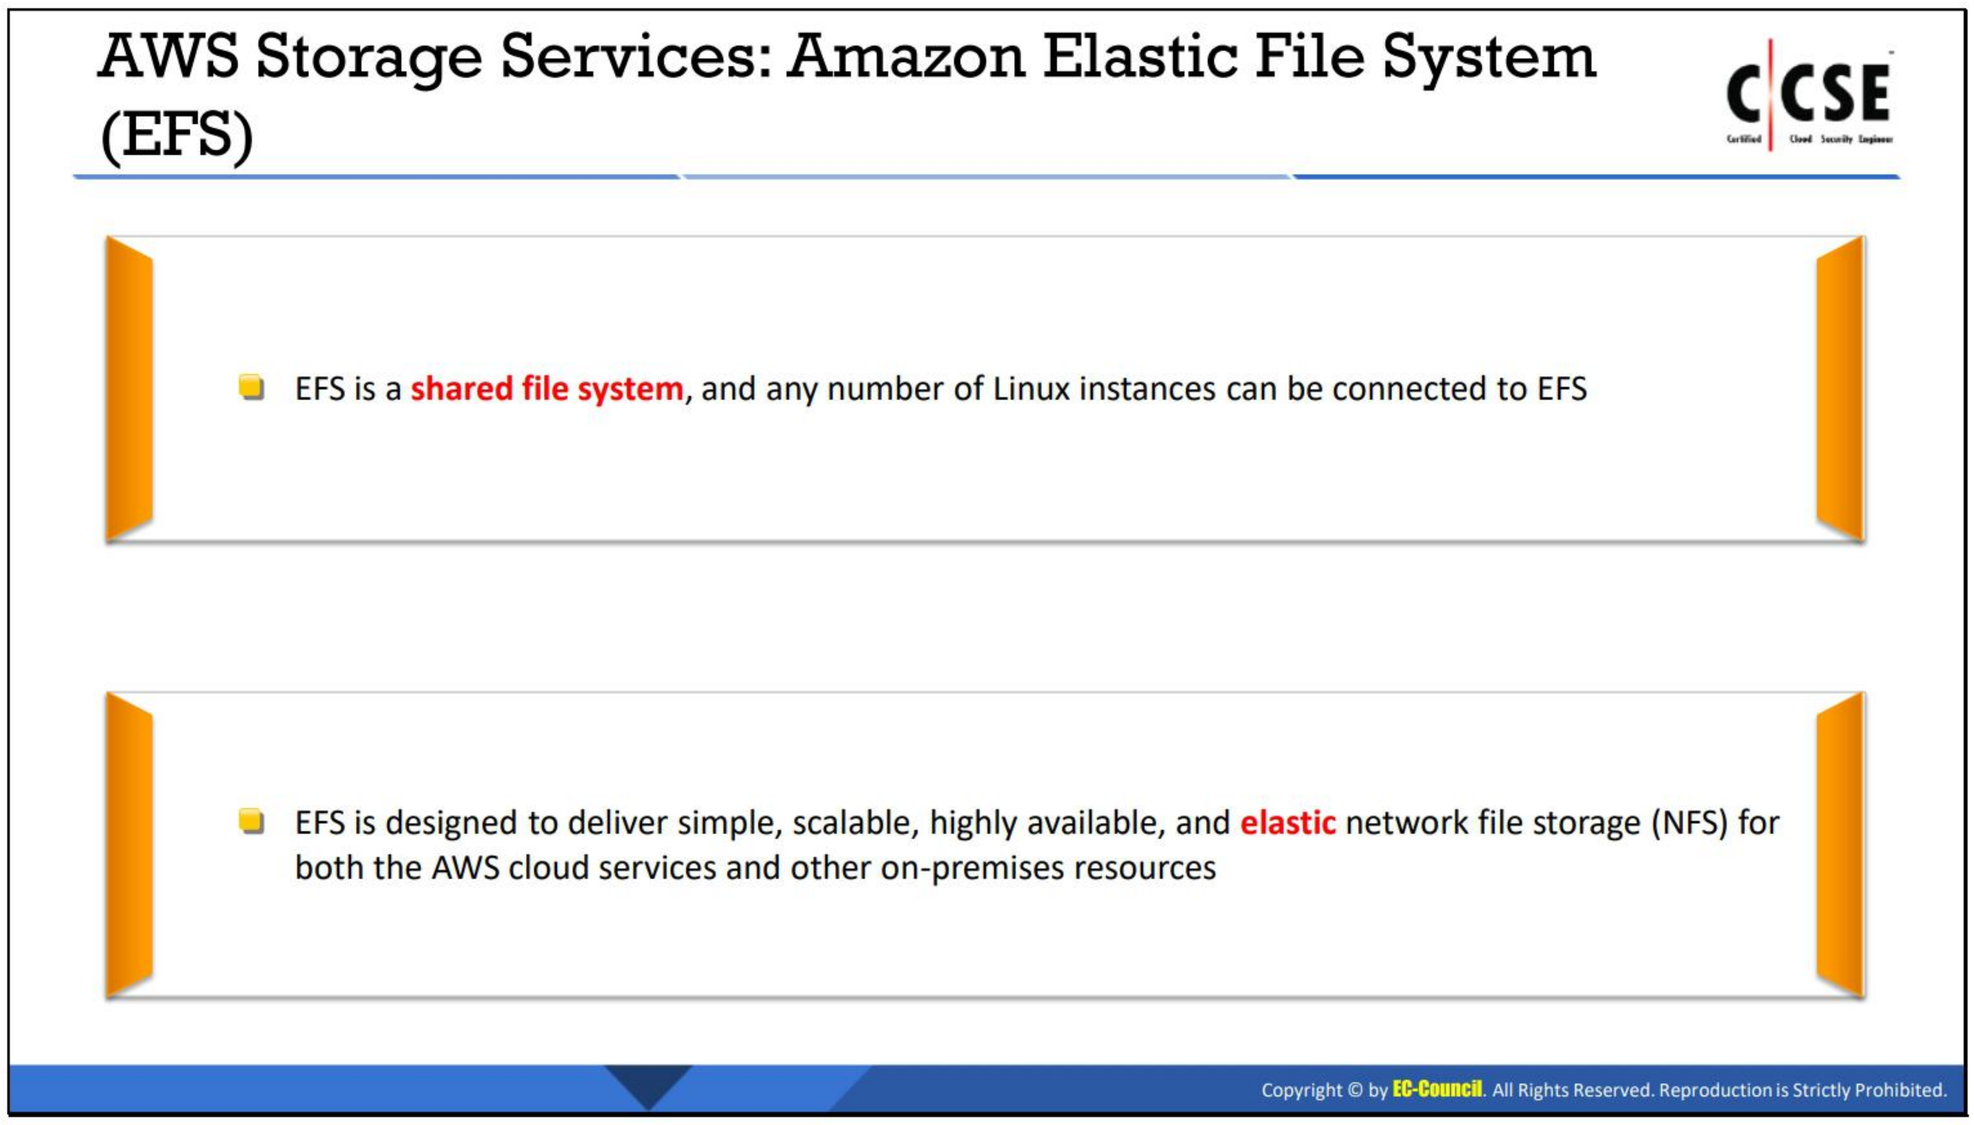
AWS Storage Gateway utilizes the internet bandwidth efficiently to increase the uploading speed of the application data to AWS. The gateway only uploads the modified data and reduces overall data transferred through the internet.

AWS Storage Gateway stores the application data for extended periods by uploading the data in Amazon S3. Amazon S3 stores data in many facilities and on multiple devices in each facility. S3 also performs systemic data integrity check regularly and has the capability of automatically self-healing.

With AWS Storage Gateway, AWS users need to pay for what they use. AWS Storage Gateway stores the data in Amazon S3, which provides robust scalability and elasticity automatically.

hide01.ir

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



AWS Storage Services: Amazon Elastic File System (EFS)

EFS is a cloud storage service utilized for the applications and workloads that run on AWS cloud. It is a shared file system and any number of Linux instances can be connected to it. All the connected instances can perform read and write operations on the files stored in EFS. As the name indicates, EFS is elastic. It means that the storage space can be expanded or reduced when the files are removed and added to the EFS. Hence, it can be used as a shared data source for applications that run on several instances or servers. EFS can be easily created and configured through a web interface.

EFS is designed to deliver simple, scalable, highly available, and elastic NFS for both AWS cloud services and other on-premises resources. AWS is responsible for managing all the EFS infrastructure and the cloud administrator does not need to worry about deploying, combining, and managing the complex file storage configuration.

EFS is mainly designed to provide high availability and durability to all the instances that are connected to services. The files are stored in multiple availability zones in a region so that different users can access the files from different availability zones. EFS also supports periodic backup so that it can be used for disaster recovery.

EFS is cost-effective, as the customer needs to pay only for the storage used. The customer does not need to allocate storage in advance and there will not be any setup cost and minimum charges. The EFS size increase as files are added, and the customer only needs to pay for the used storage space. When the files are deleted from EFS, the storage space shrinks, and the user does not have to pay for the deleted files.

There are two storage classes offered by EFS; namely, Standard Storage Class and Infrequent Access Storage class (IA). The Standard storage class is used to store data or files that are

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

frequently accessed whereas the Infrequent Access storage class (IA) is a low-price storage class used to store data or files that are accessed less frequently. If the EFS lifecycle management is enabled, all the files that are not accessed as per the lifecycle policy will be automatically moved to EFS IA and, hence, the monthly bill will be reduced.

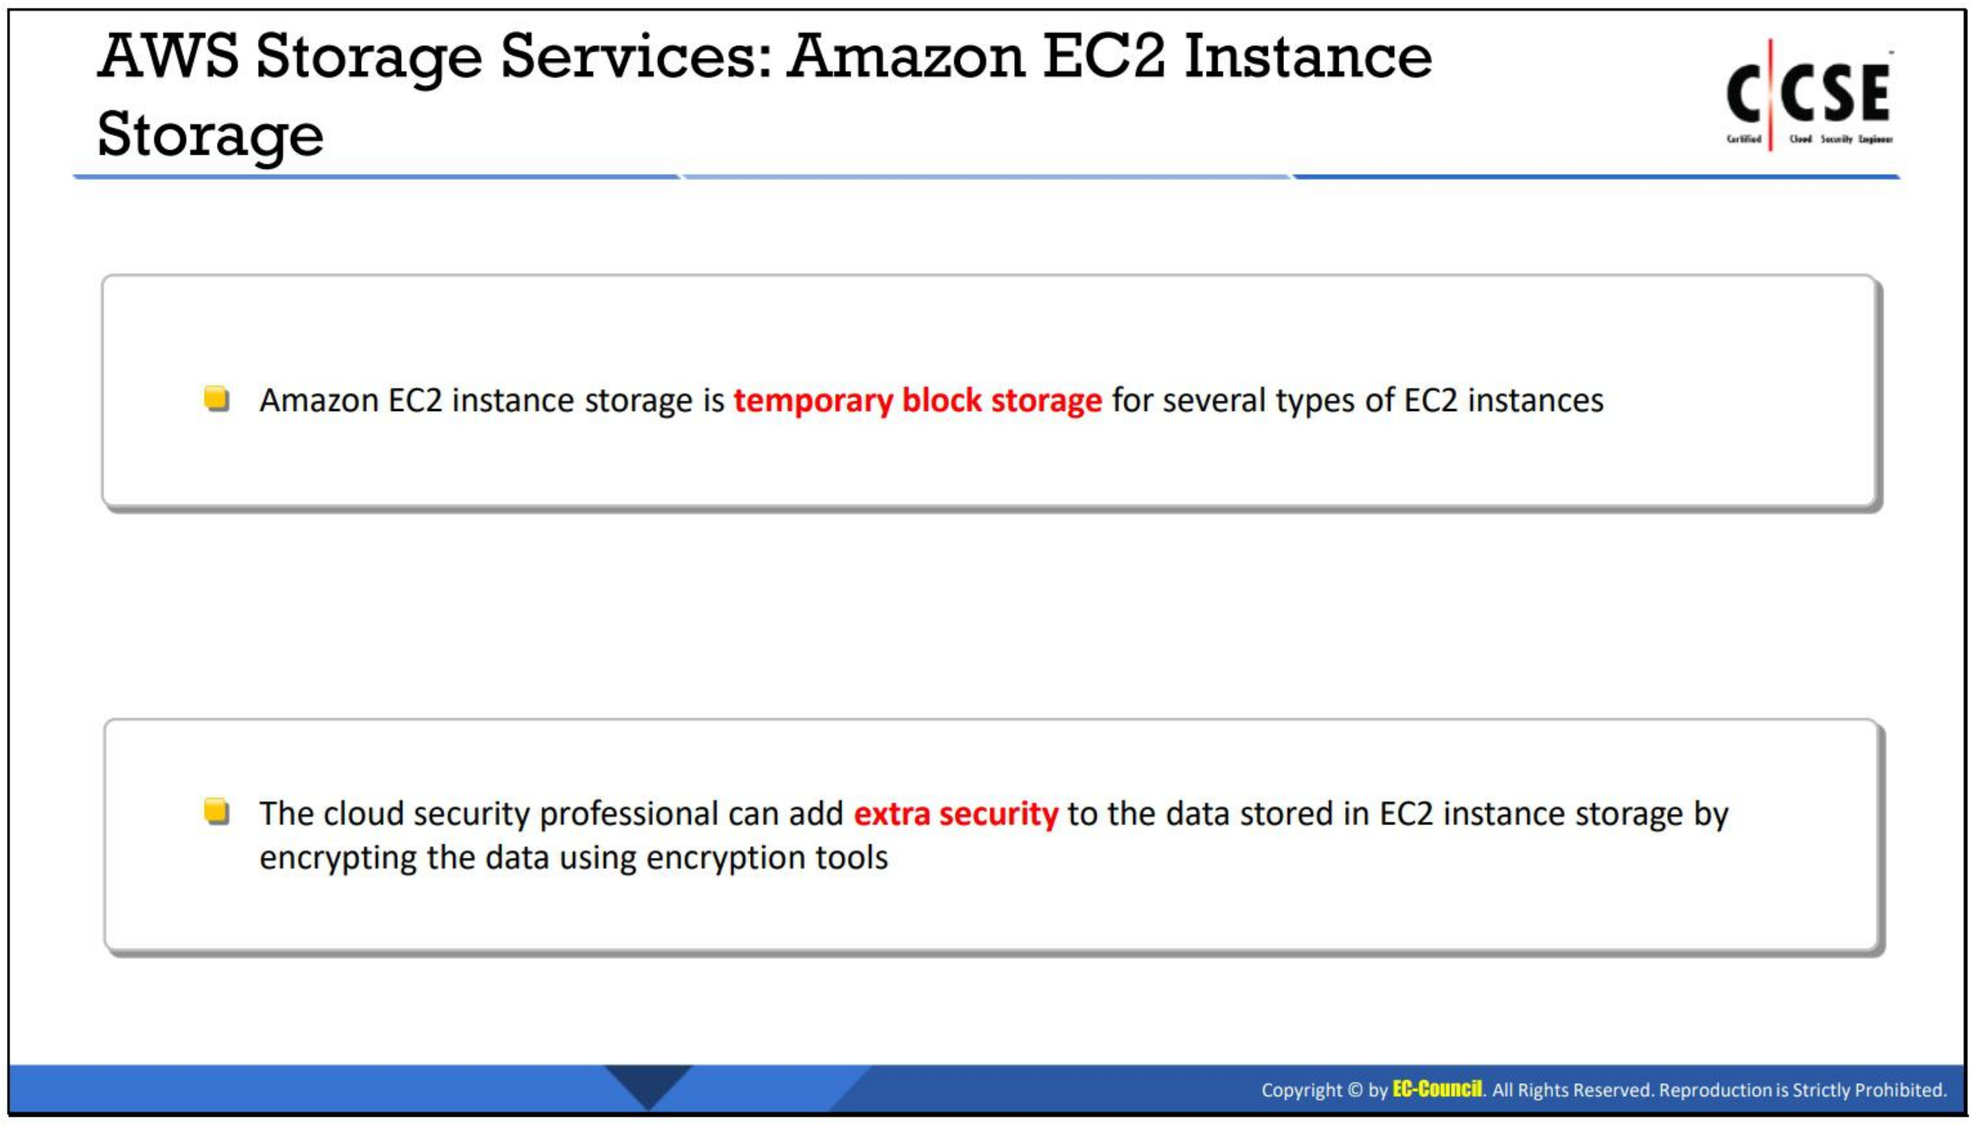
EFS offers different types of performance modes to provide low latency and IOPS for different workloads. The default general-purpose performance mode is suitable for most file systems; the limit is 7000 file operations per second. It is suitable for content management systems and general file serving.

If the file operations exceed 7000 per second, the customers can choose the max I/O performance mode. If the file system is accessed by thousands of EC2 instances, then the file system scaled up to higher levels of throughput every second. The latency per file system is higher than the general-purpose performance mode.

The cloud security professional can control access to the EFS using Portable Operating System Interface (POSIX) permissions, identity and access management (IAM), security groups for EC2 instances, and Amazon VPC. The data stored in EFS can also be secured using different encryption methods.

hide01.ir

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



AWS Storage Services: Amazon EC2 Instance Storage

Amazon EC2 instance storage is temporary block storage for several types of EC2 instances. It is also called as an ephemeral drive. It is one of the cost-effective, easy to use, and flexible data storage. This storage is provided by the disk that is attached to the same physical machine that hosts EC2 instances. Amazon EC2 instance storage is highly suitable for the storage of data that changes frequently such as caches, temporary data, buffers, logs, and scratch information. This storage is also used for applications where the focus is on processing the data instead of storing the data.

The Amazon EC2 instance storage can contain more than one instance store volume. The amount of disk storage and the number of devices is based on the type of the EC2 instance. Larger EC2 instance types tend to provide a large and multiple instance storage volumes.

For micro EC2 instances such as t1 and t2, EBS storage can be used only when the instance store is not available. Additionally, the instances that use EBS for root devices will not expose EC2 instance storage by default. The EC2 instance storage for these instances can be exposed while launching the EC2 instance by mentioning the block device mapping.

The EC2 instance storage cannot be detached and attached from one instance to another. The data in the EC2 instance storage remains only till the EC2 instance exists. The data in the instance storage gets deleted when the underlying disk drive fails and when the instance stops or terminates. When the instance is terminated, stopped, or hibernated, the block storage in the EC2 instance store will be reset. Thus, the data cannot be accessed from the EC2 instance store of another EC2 instance but the data in the instance storage remains in case of instance, reboot whether it is intentional or unintentional. Thus, EC2 instance storage is not suitable for storing long-term and important data.

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

When an Amazon Machine Image (AMI) is created from an instance that has data in its EC2 instance storage volume, the same data will not be replicated in the instance store volume of an instance that is launched from the AMI.

The EC2 instance storage offers extremely high I/O speed as it resides in the same host of the EC2 server. The I/O speed is relatively high compared to the other storage options provided by AWS.

There are different types of EC2 instance storages such as ephemeral, non-volatile memory (NVMe) SSD, and TRIM, and they differ in their performance capabilities. EC2 instance store volumes support ephemeral (0-23) for instances. The default EC2 instance storage volume type is ephemeral while launching an instance. NVMe SSDs are used for high I/O speed operations. This type of instance storage volume is only supported by I3 instances. The instance storage volume type that supports TRIM is used for achieving high performance and to reduce write amplification.

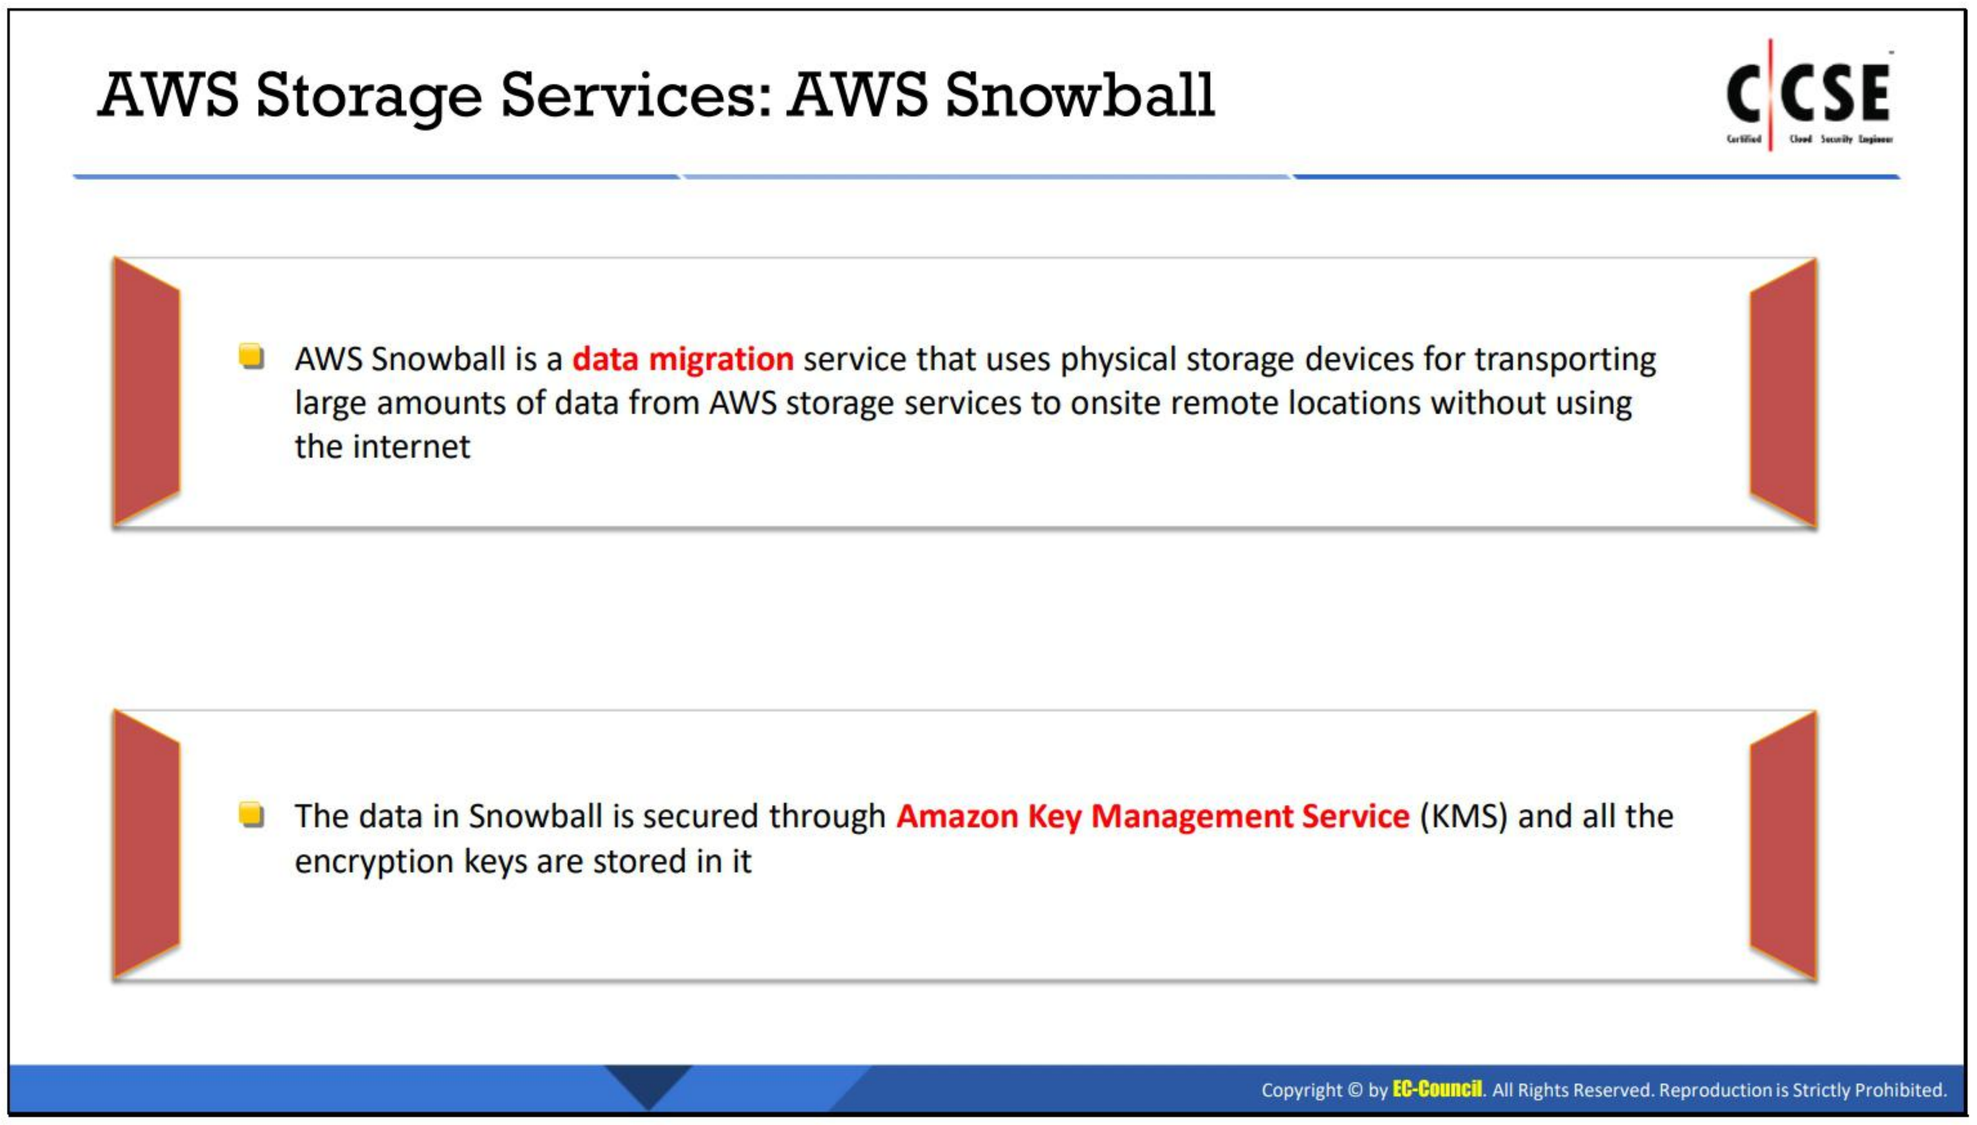
The number of instance store volumes is fixed, and it cannot be increased or decreased for a single instance; however, the storage capacity is still elastic and scalable, which is achieved by increasing or decreasing the number of running instances.

Using IAM, the cloud administrator can control the users to performs operations such as launching and terminating the EC2 instances in the AWS account. Additionally, the data stored in an instance store will be deleted when the instance is stopped or terminated. Thus, the data in that EC2 instance storage cannot be accessed by any other instance. The cloud security professional can add extra security to the data stored in EC2 instance storage by encrypting the data using encryption tools.

hide01.ir

The cost for an EC2 instance also includes EC2 instance storage. Thus, the cloud user does not need to pay extra for the data stored in instance storage; however, network charges would be applicable while transferring the data to instance storage from instances that are in other availability zones.

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



AWS Storage Services: AWS Snowball

AWS Snowball is a data migration service. It uses physical storage devices for transporting large amounts of data from AWS storage services to onsite remote locations by avoiding internet usage. Extremely large amounts of data can be transferred using Snowball secure storage devices. By using AWS Snowball, a cloud user can reduce the cost and risks associated with the transfer of a large amount of data such as network costs, security concerns, and long duration for data transfer, etc.

For example, consider a cloud user with 80 TB of data with a slow internet connection. Using AWS Snowball, the cloud user can send the data to AWS through physical devices and Amazon would transfer the data into AWS storage devices using its own high-speed internet connection. Amazon itself provides the appliance, and the users only need to load the data into the appliance. The data in the Snowball can be imported and exported from and to AWS S3 buckets. The data can be transferred from S3 buckets to other AWS storage services. Cloud users can rely on Amazon for transferring their data securely and in a cost-effective way.

Snowball appliance is waterproof and dustproof. It is available for all the AWS regions and it is lightweight. The backside of the Snowball device contains two SFP+ 10 GigE network connections, one RJ45 1 GigE, and a power cord. The frontside of Snowball device contains control panel and an E Ink display.

Snowball is highly suitable in cases where users do not want to spend money on upgrading their internet to a high-speed internet connection. It is also suitable in scenarios where data transfer over the internet takes more than a week. It is also used to transfer data such as database, medical-related records, backups, content distribution, and disaster recovery.

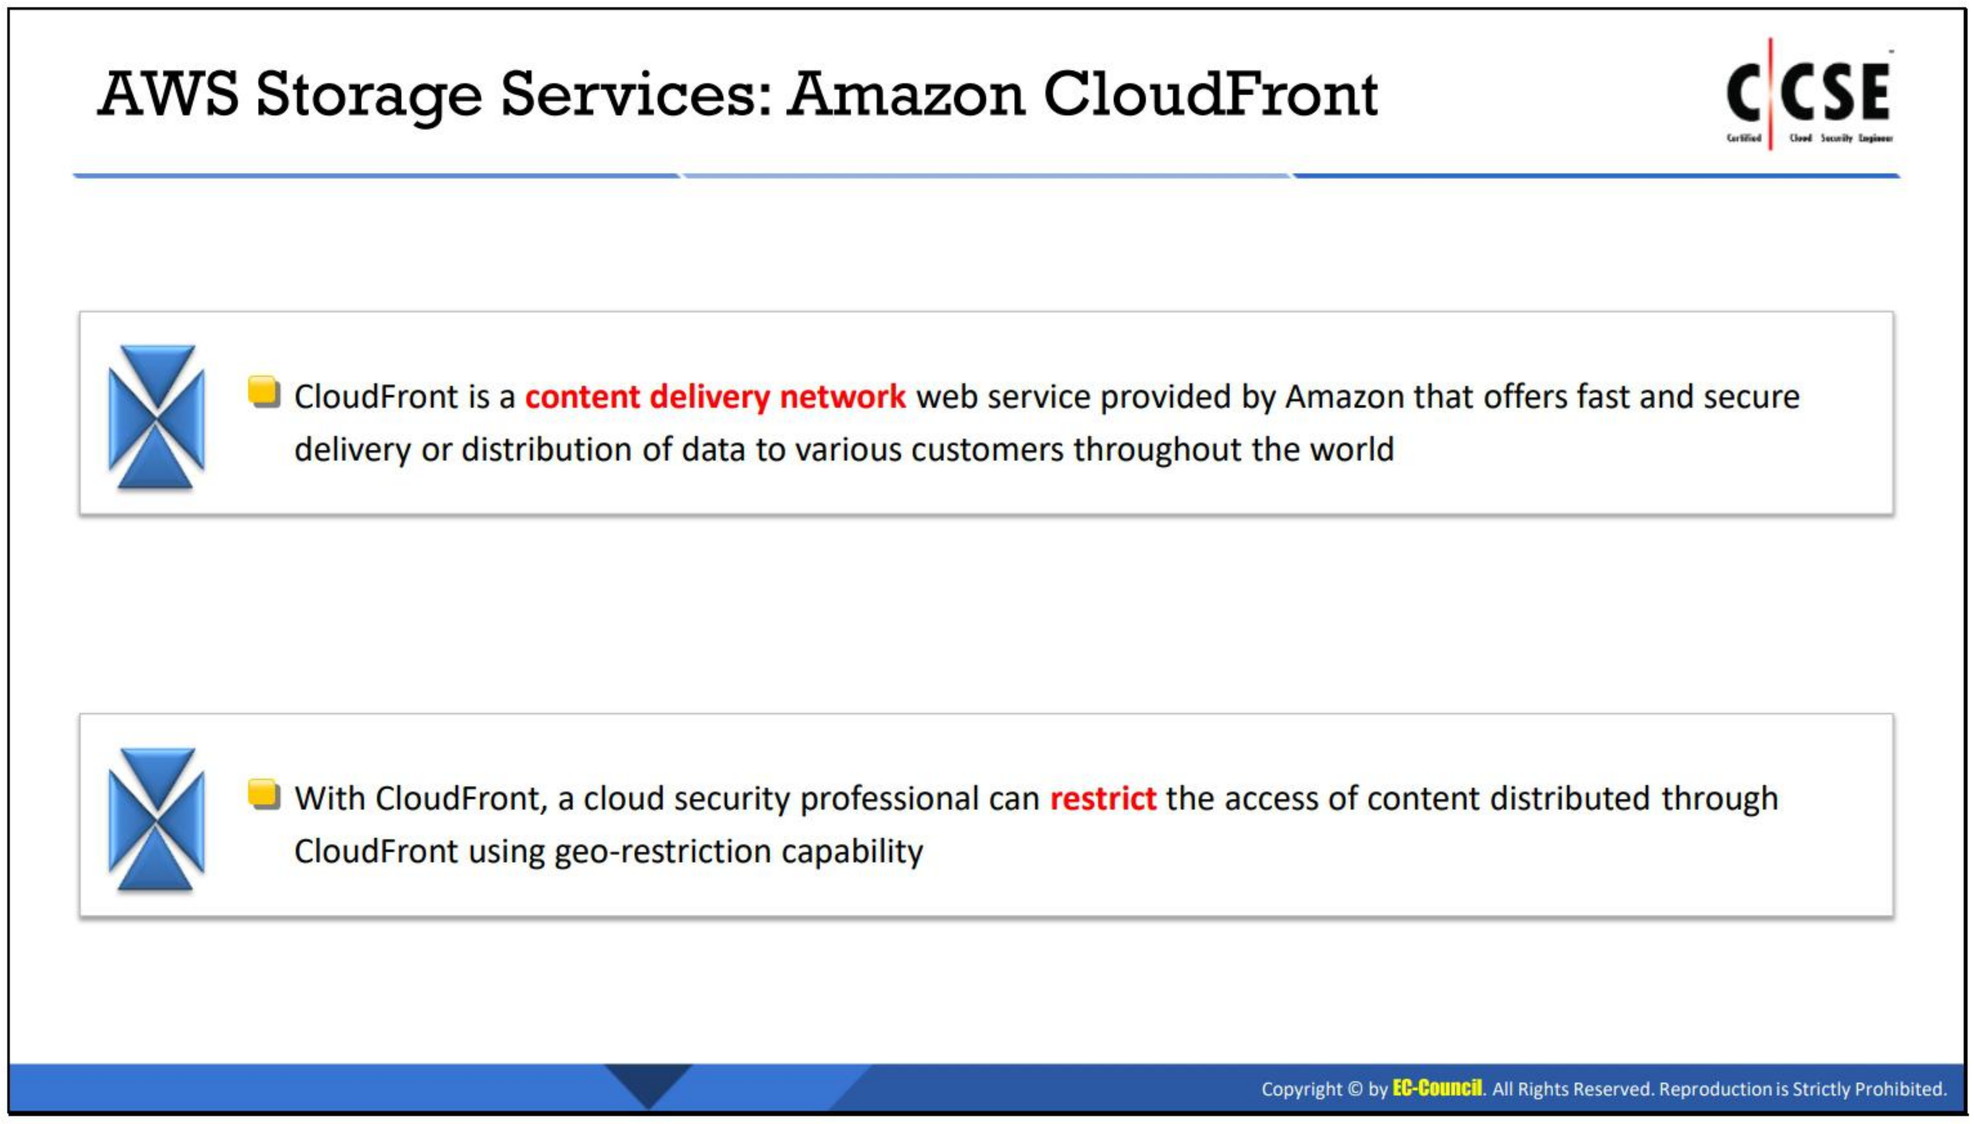
Certified Cloud Security Engineer Copyright © by **EC-Council**
All Rights Reserved. Reproduction is Strictly Prohibited.

Amazon offers three different types of Snowball devices; namely Snowball, Snowball Edge, and Snow Mobile. Snowball can transfer up to 80 TB of data and it is simple, safe, fast, and secure.

Using IAM, a cloud administrator can create users and provide access to perform Snowball actions. IAM users who need to perform the Snowball job must have access to AWS S3 buckets to perform import actions. All the Snowball devices come with tamper-evident enclosures, 256-bit encryption and industry-standard Trusted Platform Modules (TPM) for a secure data transfer. Additionally, the data in Snowball is secured through Amazon KMS and all the encryption keys are stored in it.

hide01.ir

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



AWS Storage Services: Amazon CloudFront

CloudFront is a CDN web service provided by AWS that offers fast and secure delivery for distribution of data to various customers throughout the world. CloudFront offers low-latency and high-speed data delivery through a user-friendly environment. Amazon CloudFront can be integrated with other AWS services such as S3, Route 53, and EC2.

CloudFront can deliver all types of files and content that can be served through HTTP. Examples include HTML pages, website images, audio, videos, and other media files. CloudFront can also deliver live media through HTTP. Amazon CloudFront delivers data using a global network of data centers called edge locations. An edge location is a data center with physical devices where the cache copies of the content are stored to reduce the latency for end-users.

For example, when a user sends a request for the data that is served with CloudFront, the request goes to the nearest edge location to reduce the latency and deliver it as fast as possible with low network traffic. If the edge location has the cache of the requested content, CloudFront delivers it immediately with the lowest latency. If the content is not available in the edge location, it retrieves the content from origin and stores the cache in the edge location for a certain period.

CloudFront is mainly used to stream media files. It is also used to deliver the most accessed static content such as audio, videos, or other media files. It can also be used to deliver web applications that contain static content, dynamic content, or both.

CloudFront is mainly designed to deliver content with low latency and high bandwidth. It speeds up the content distribution by sending the requests of end-users to the nearest edge locations instead of routing to the original source. It also reduces the number of networks through which a user request must pass through and thereby improves performance.

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

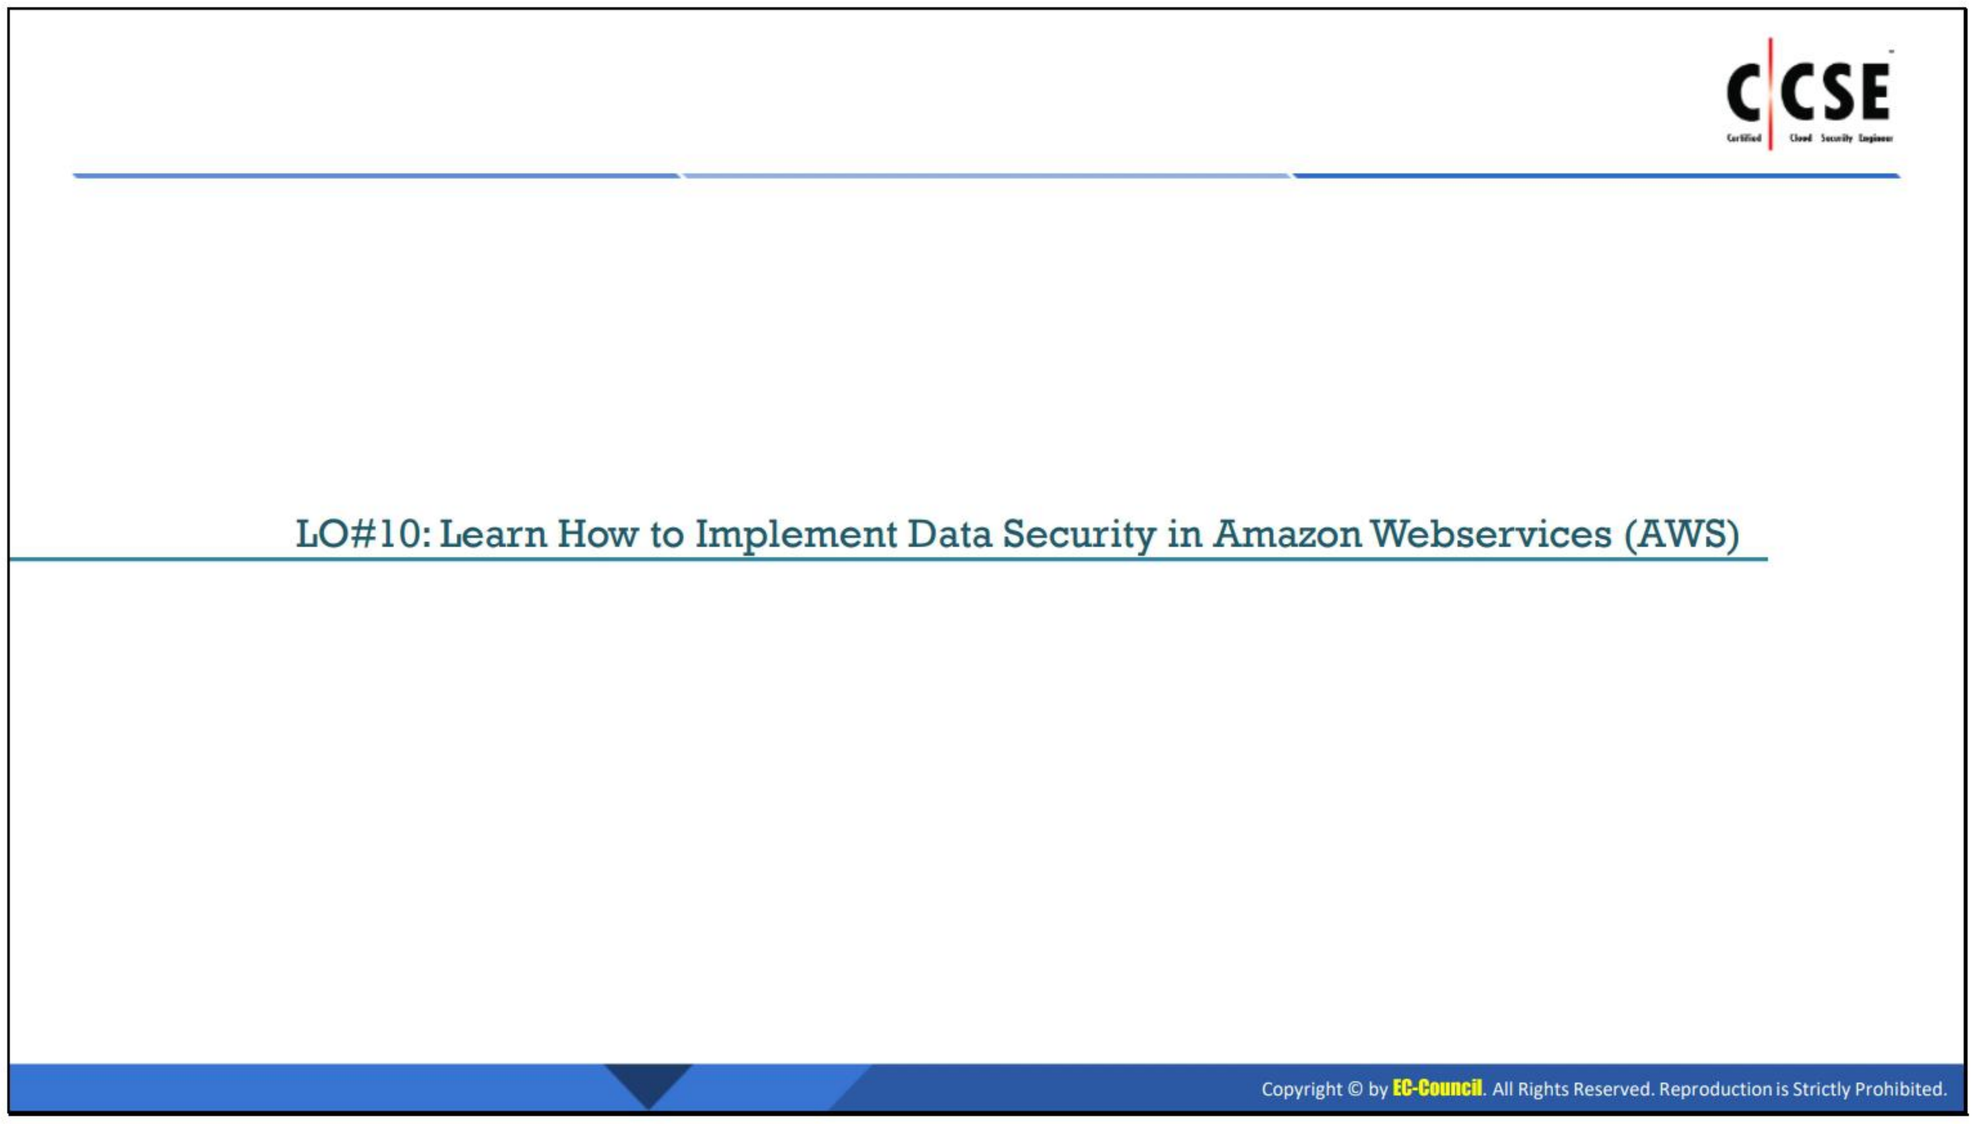
Amazon CloudFront provides both availability and reliability as the content copies are stored in various edge locations throughout the world in the form of cache content. It also provides both elasticity and scalability features. The cloud administrator does not need to worry about maintaining a web server with a high capacity to serve high traffic for the content. CloudFront automatically responds and handles all the traffic. In addition, CloudFront offers multiple layers of caching at every edge location and eliminates the need for scaling the infrastructure when a website becomes popular.

Amazon CloudFront is cost-effective as the cloud users need to pay only for the content that is delivered through the network and there will not be any hidden, minimum, or up-front charges.

With CloudFront, a cloud security professional can restrict the access of content distributed through CloudFront using geo-restriction capability. This feature restricts a user from a particular location from accessing the content. CloudFront provides both network and application-level security. It is integrated with AWS Shield Standard to protect the content and applications from different types of attacks.

hide01.ir

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

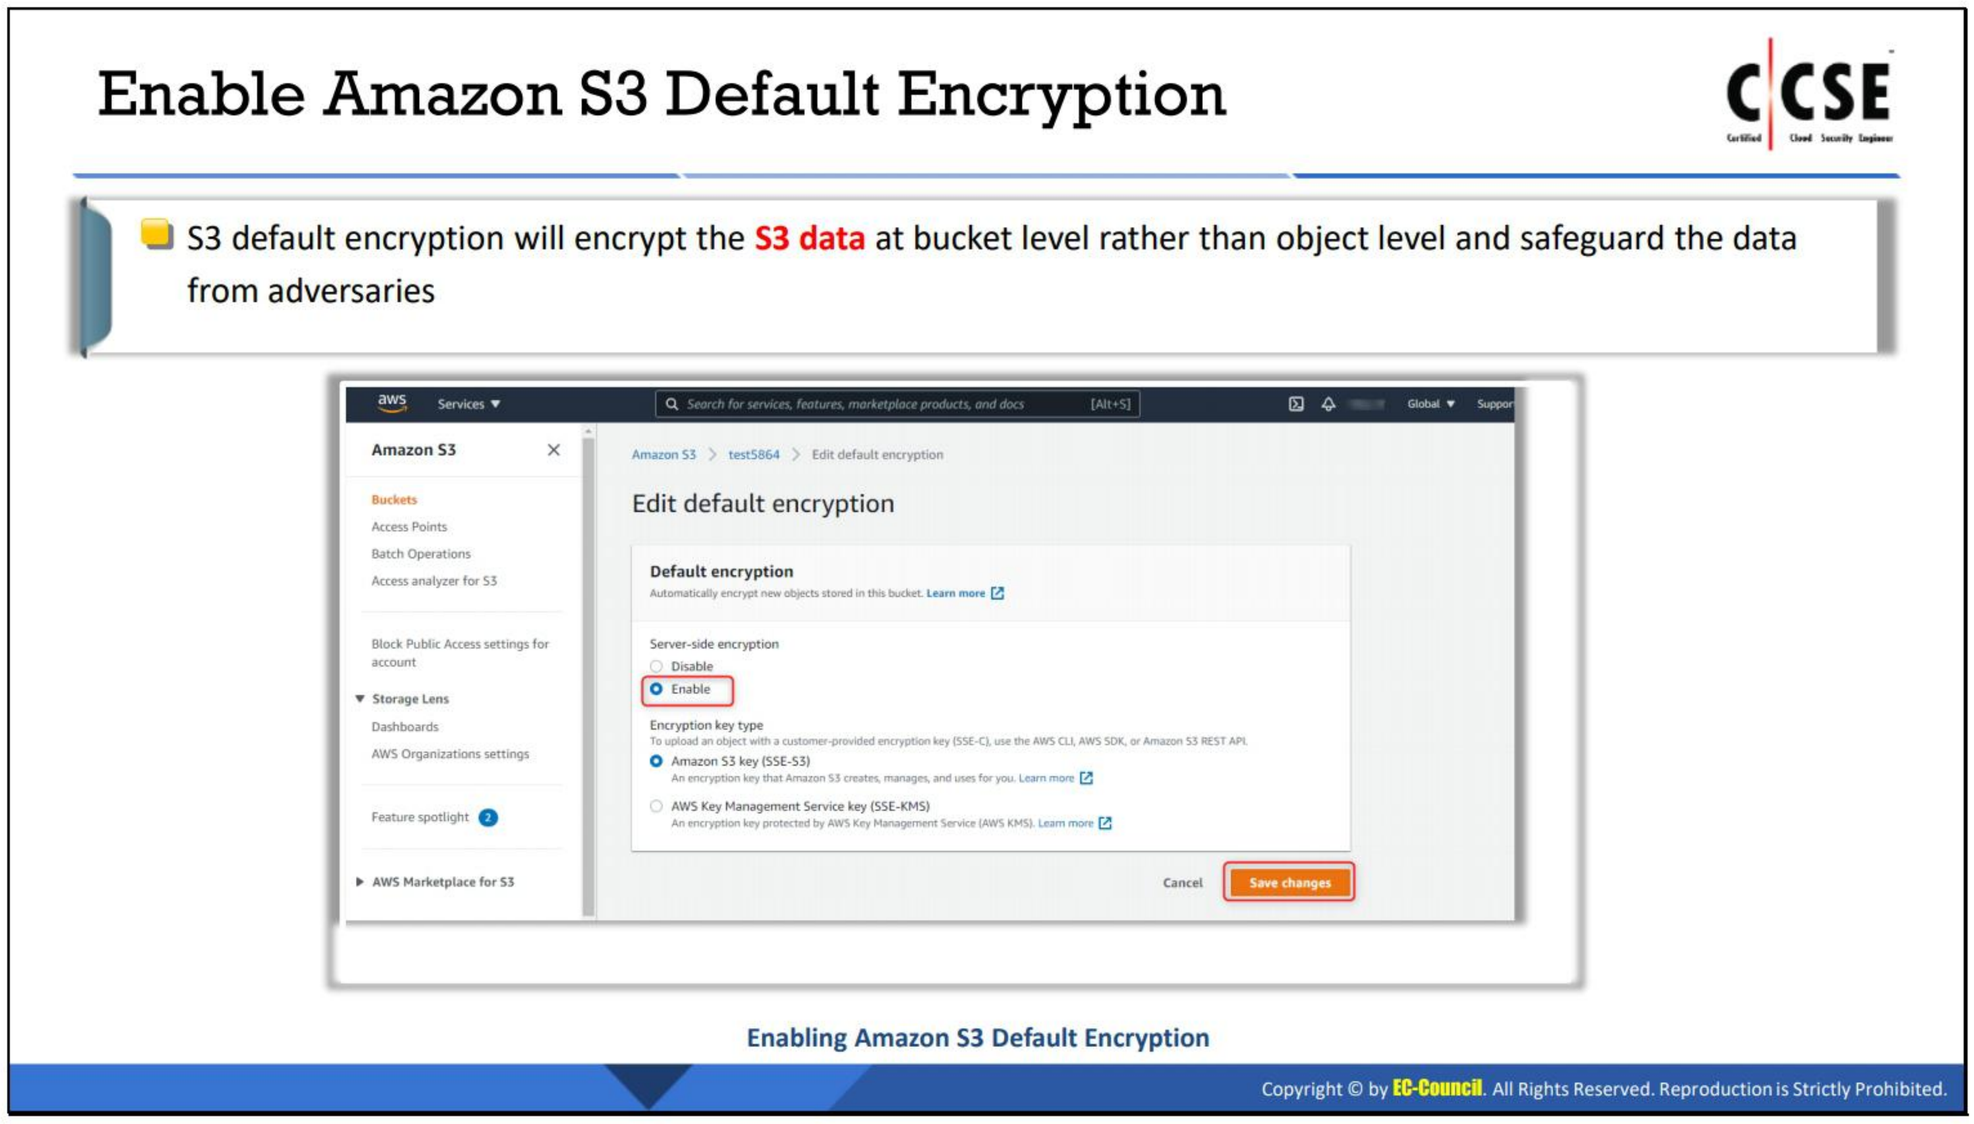


LO#10: Learn How to Implement Data Security in Amazon Webservices (AWS)

The objective of this section is to explain the implementation of data security in AWS using different features such as S3 default encryption, customer-provided AWS KMS CMKs, geo-restriction, etc.

hide01.ir

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



Enable Amazon S3 Default Encryption

Amazon S3's default encryption will encrypt the S3 data at the bucket level rather than object level and safeguard the data from adversaries. Therefore, the cloud security administrator should enable the default encryption of Amazon S3 bucket.

Source: https://aws.amazon.com

To enable default encryption on an Amazon S3 bucket utilizing AWS Console, perform the following steps:

- Sign-in to the **AWS Management Console**.
- Click on **Services** dropdown; then click on **S3** under **Storage.**

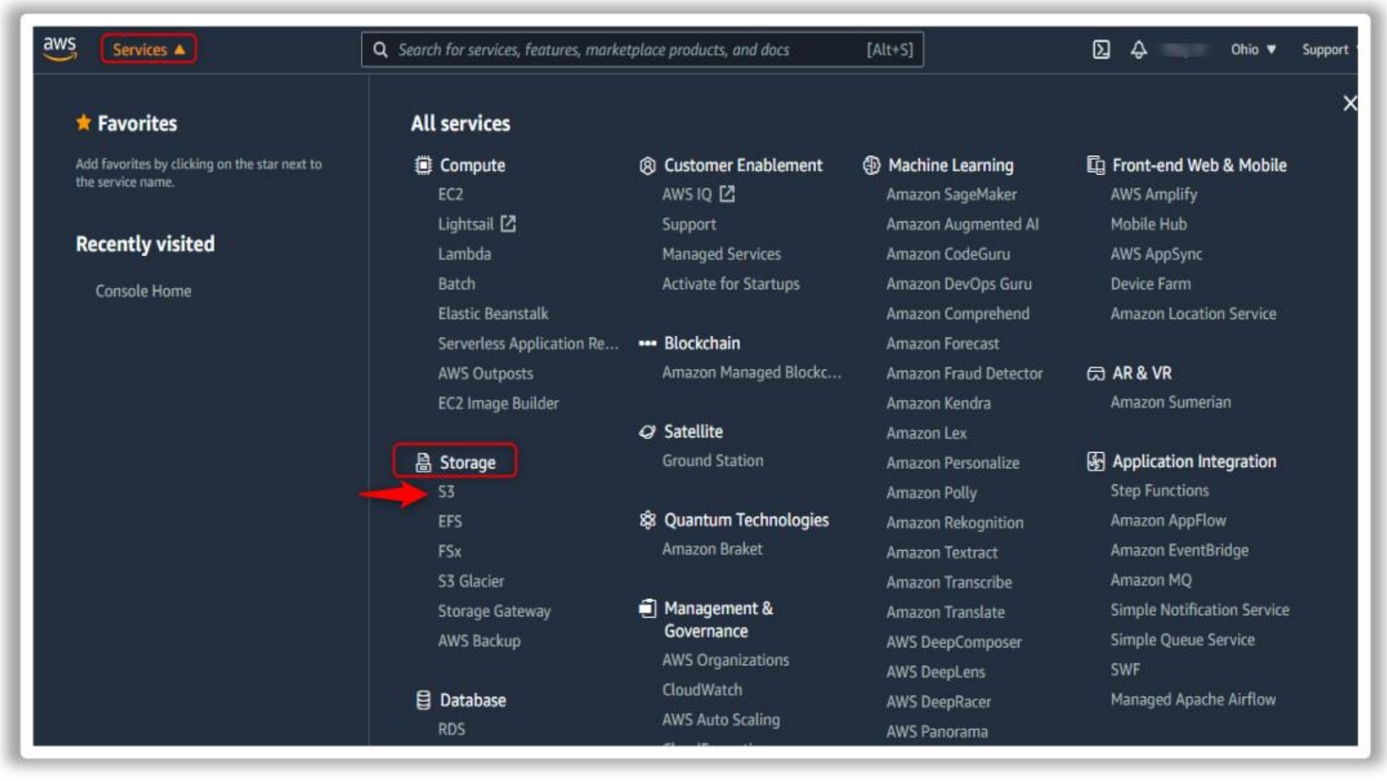


Figure 4.45: Click on Storage

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

- The list of available buckets will be displayed.

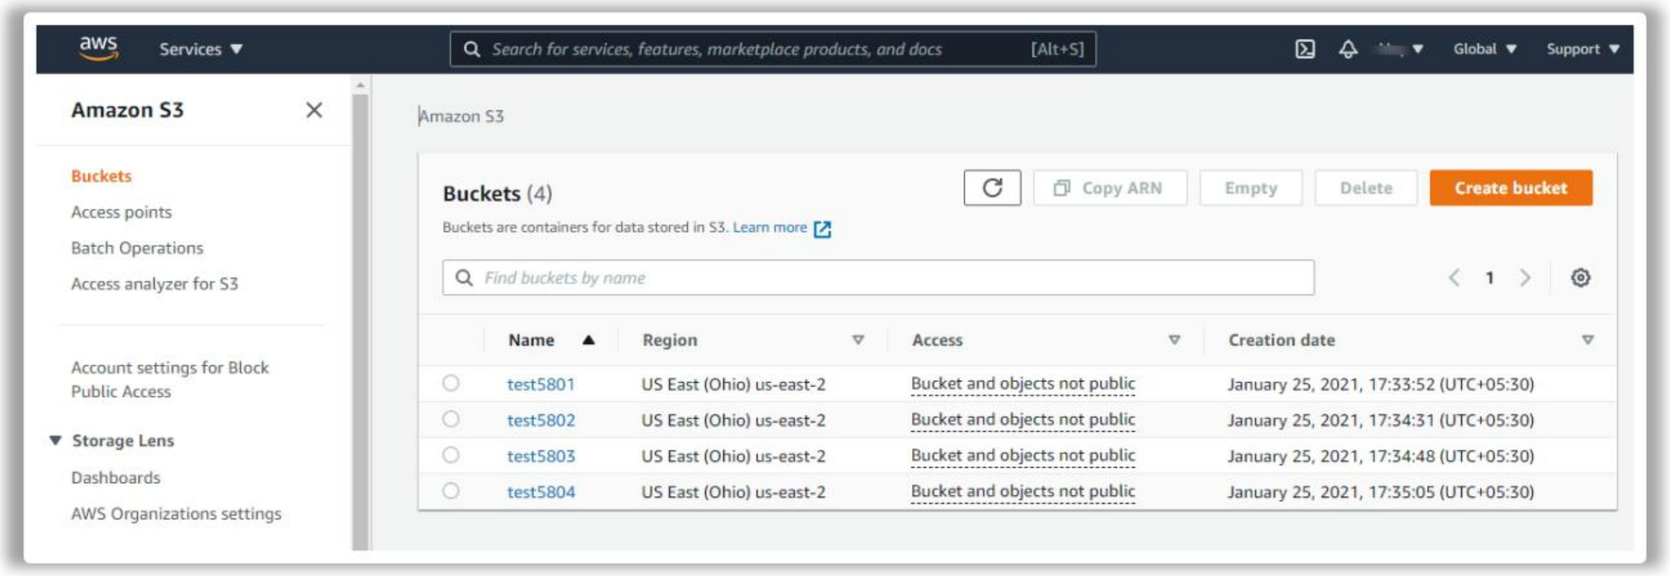


Figure 4.46: List of Buckets

- Select the Bucket that you want to encrypt and click on it. Then, click on **Properties** tab.

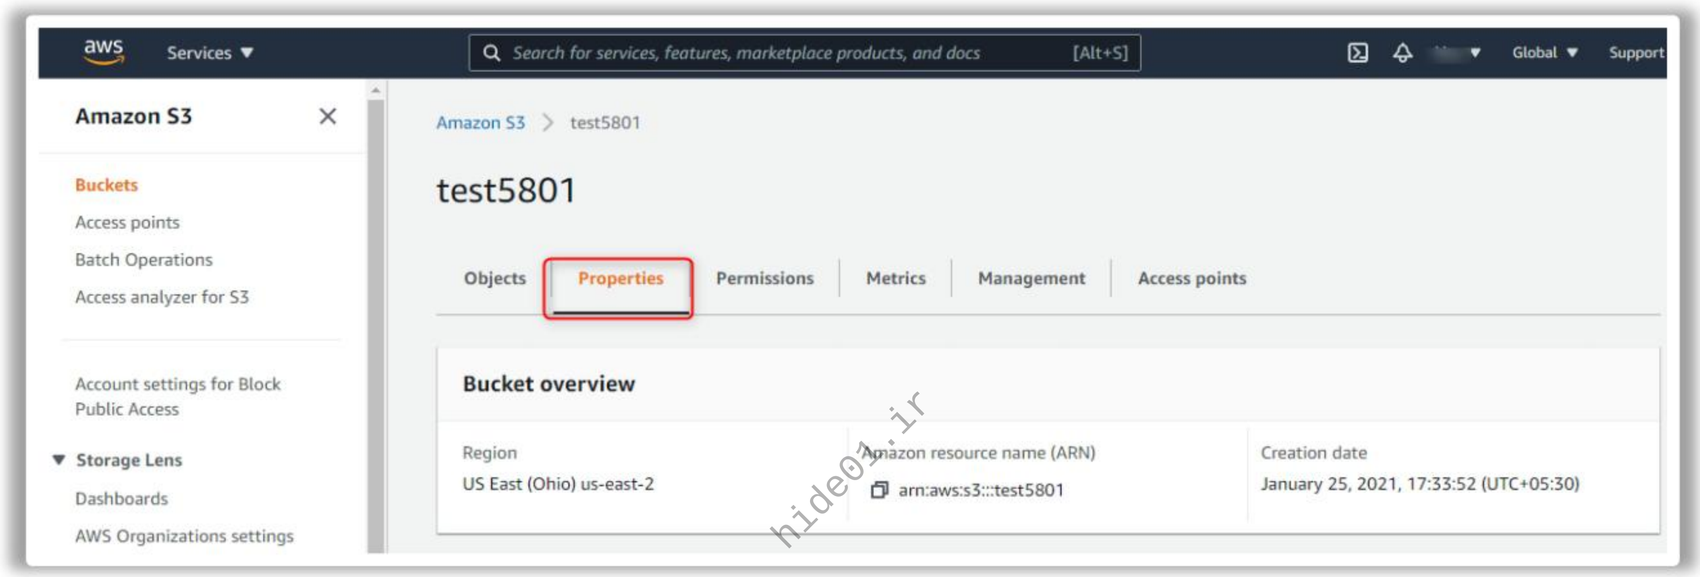


Figure 4.47: Click on Properties

- Scroll down and click the **Edit** button in **Default encryption** section.

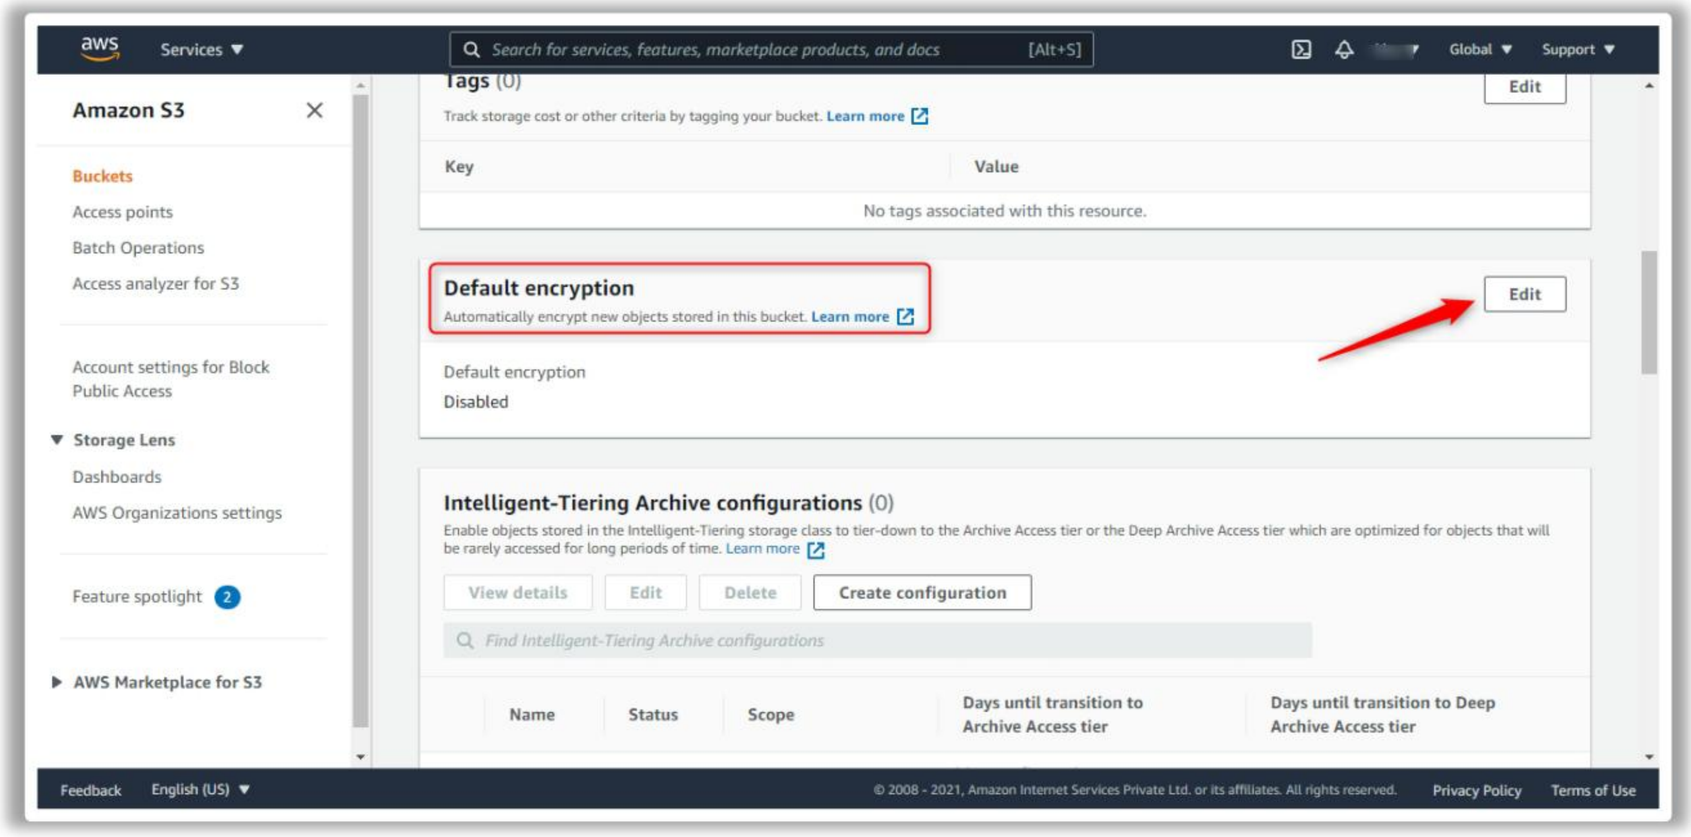


Figure 4.48: Select Edit Button

- Navigate to **Server-side encryption**, check **Enable** radio button.

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

- Based on your requirement, choose the **Encryption key type** and click on **Save changes** button.

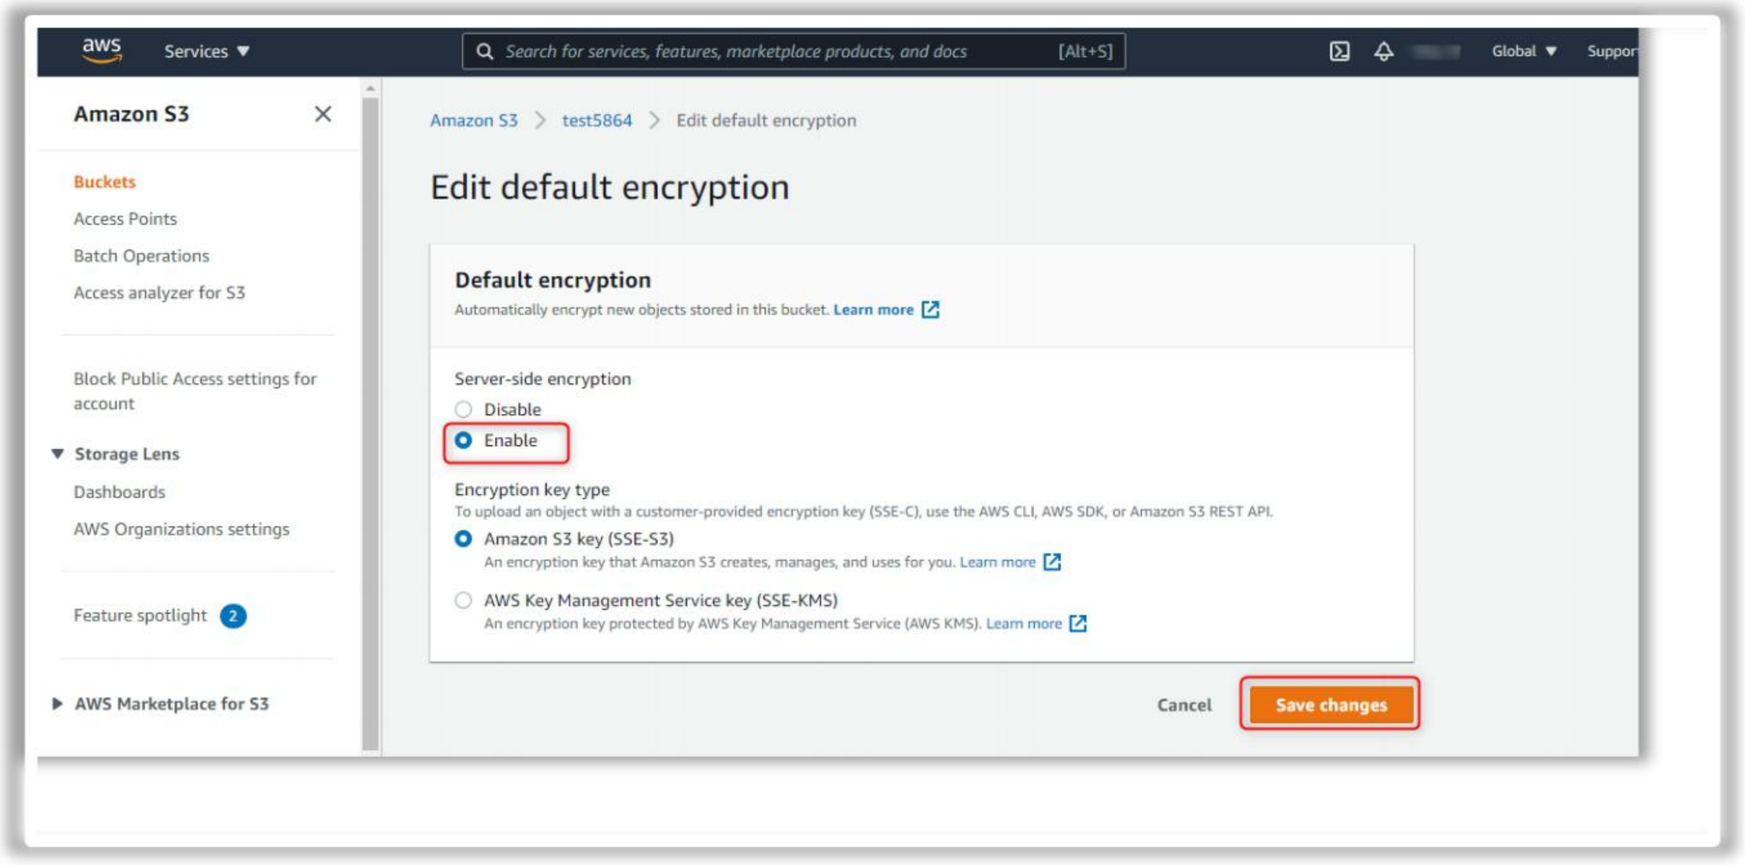


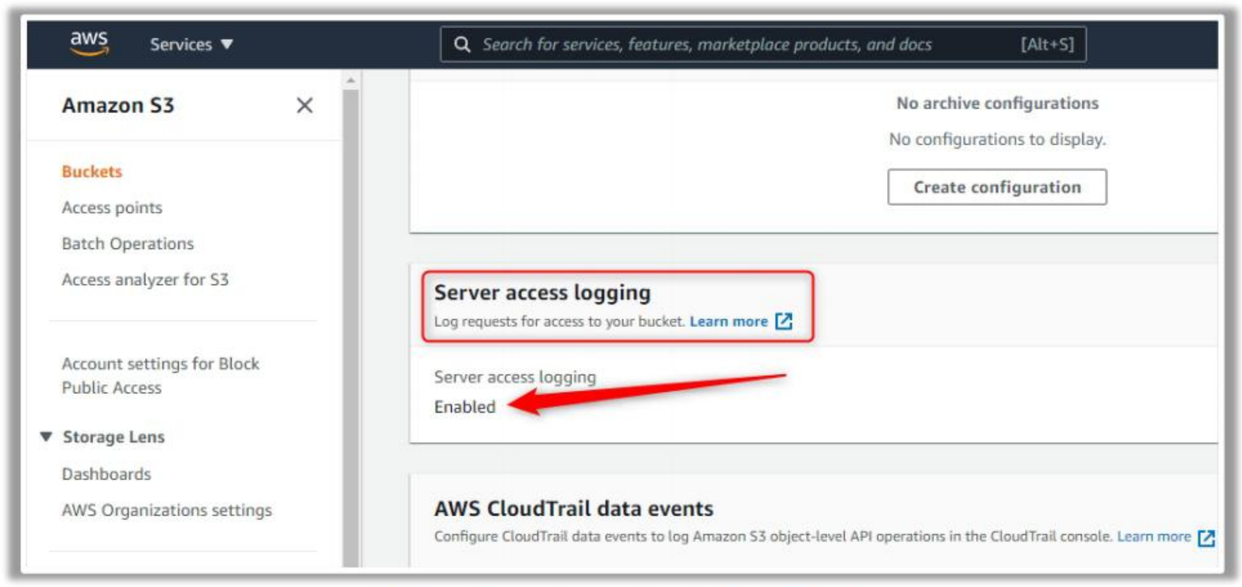
Figure 4.49: Edit Default Encryption Console

hide01.ir

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

Enable Server Access Logging for Amazon S3 Bucket

- Enable server access logging feature for S3 buckets to **track** the requests made to access the S3 buckets and to take measures to protect buckets from unauthorized access



Enabling Server Access Logging for S3 bucket

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

Enable Server Access Logging for Amazon S3 Bucket

Amazon S3 does not allow logging by default; the service user needs to enable logging for S3 buckets. S3 collects logs for the source bucket and target bucket, which reside in the same region; however, ensure that the service user does not have a configured retention period.

Enable server access logging feature for S3 buckets to track the request made to access the S3 buckets and to take measures to protect the buckets from unauthorized access.

Source: https://aws.amazon.com

To enable server access logging for an S3 bucket in AWS Console, perform the following steps:

- Sign-in to the **AWS Management Console**.

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

- Click on **Services** dropdown; then click on **S3** under **Storage.**

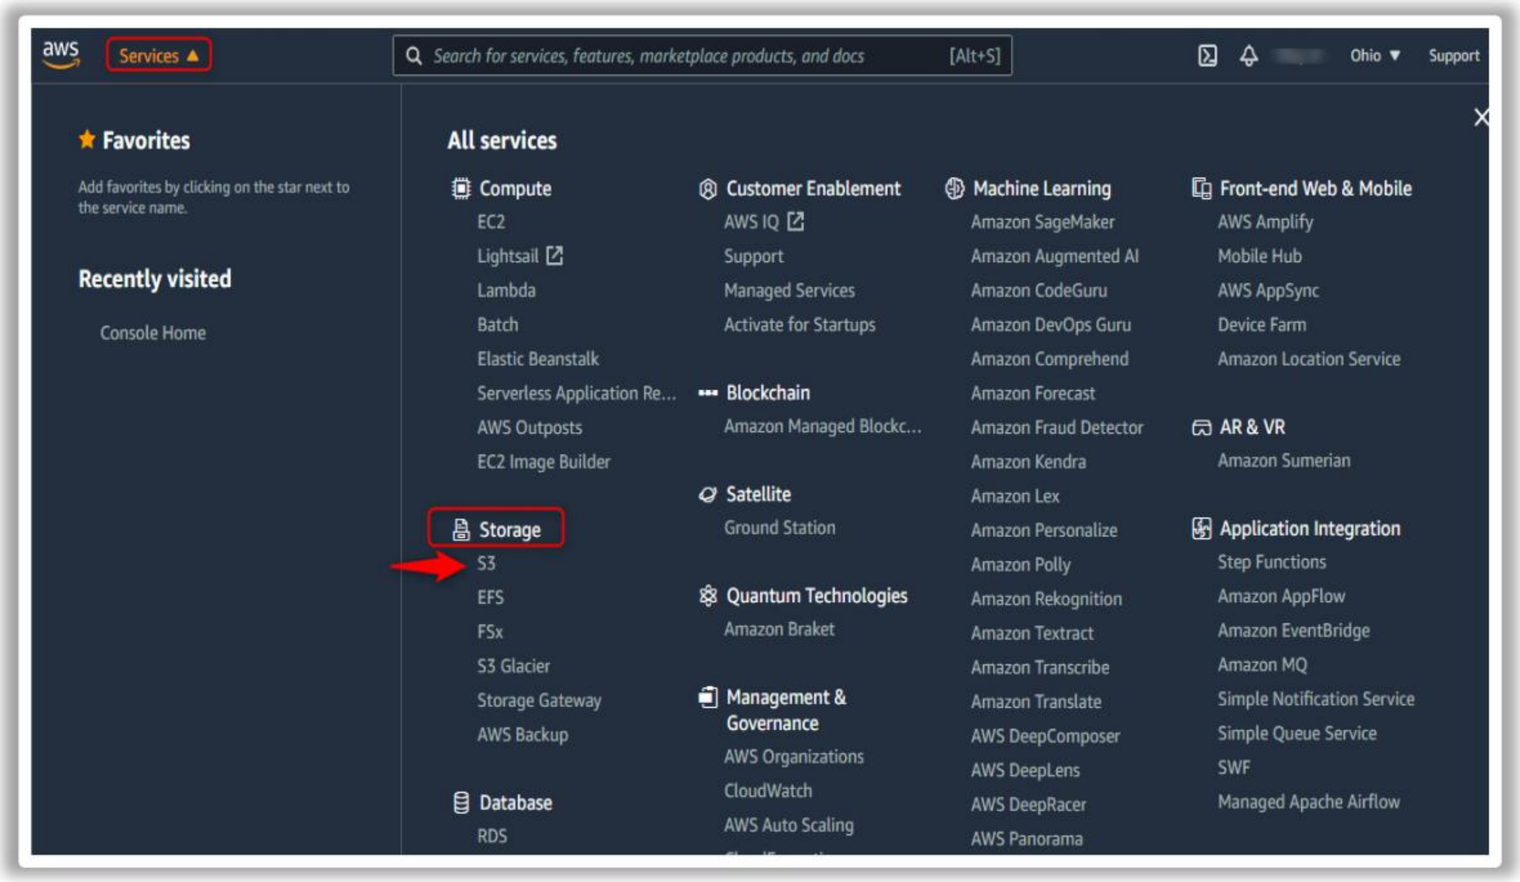
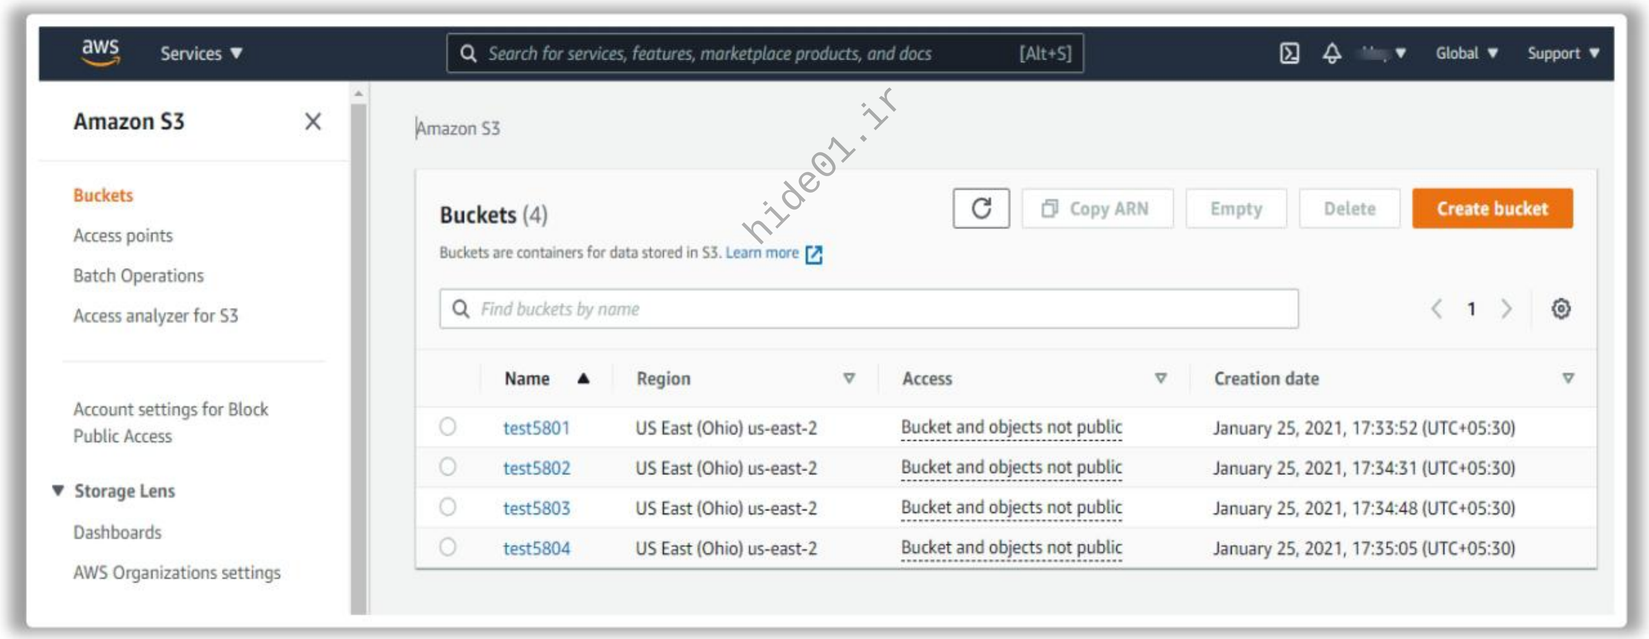


Figure 4.50: Click on Service

- The list of available buckets will be displayed.



hide01.ir

Figure 4.51: List of Bucket

- Select the Bucket that you want to encrypt and click on it. Then, click on **Properties** tab.

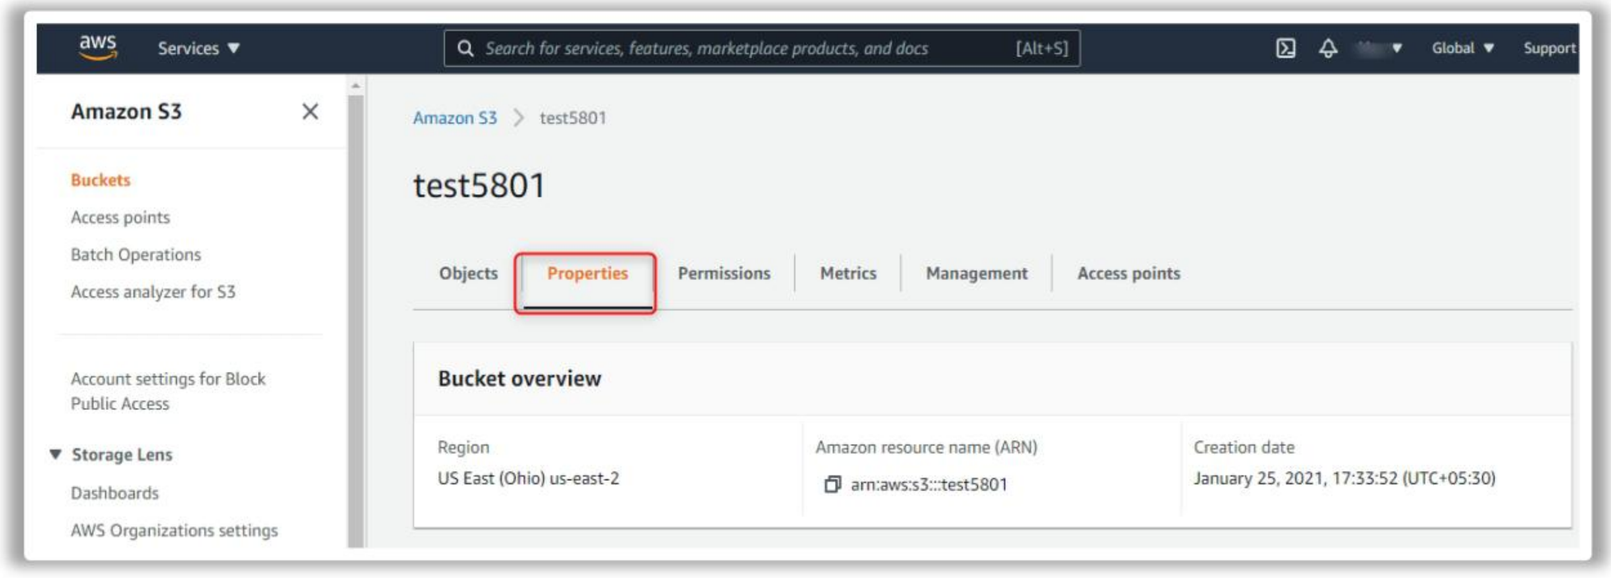


Figure 4.52: Click on Properties

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

- Navigate and click the **Edit** button beside **Server access logging**.

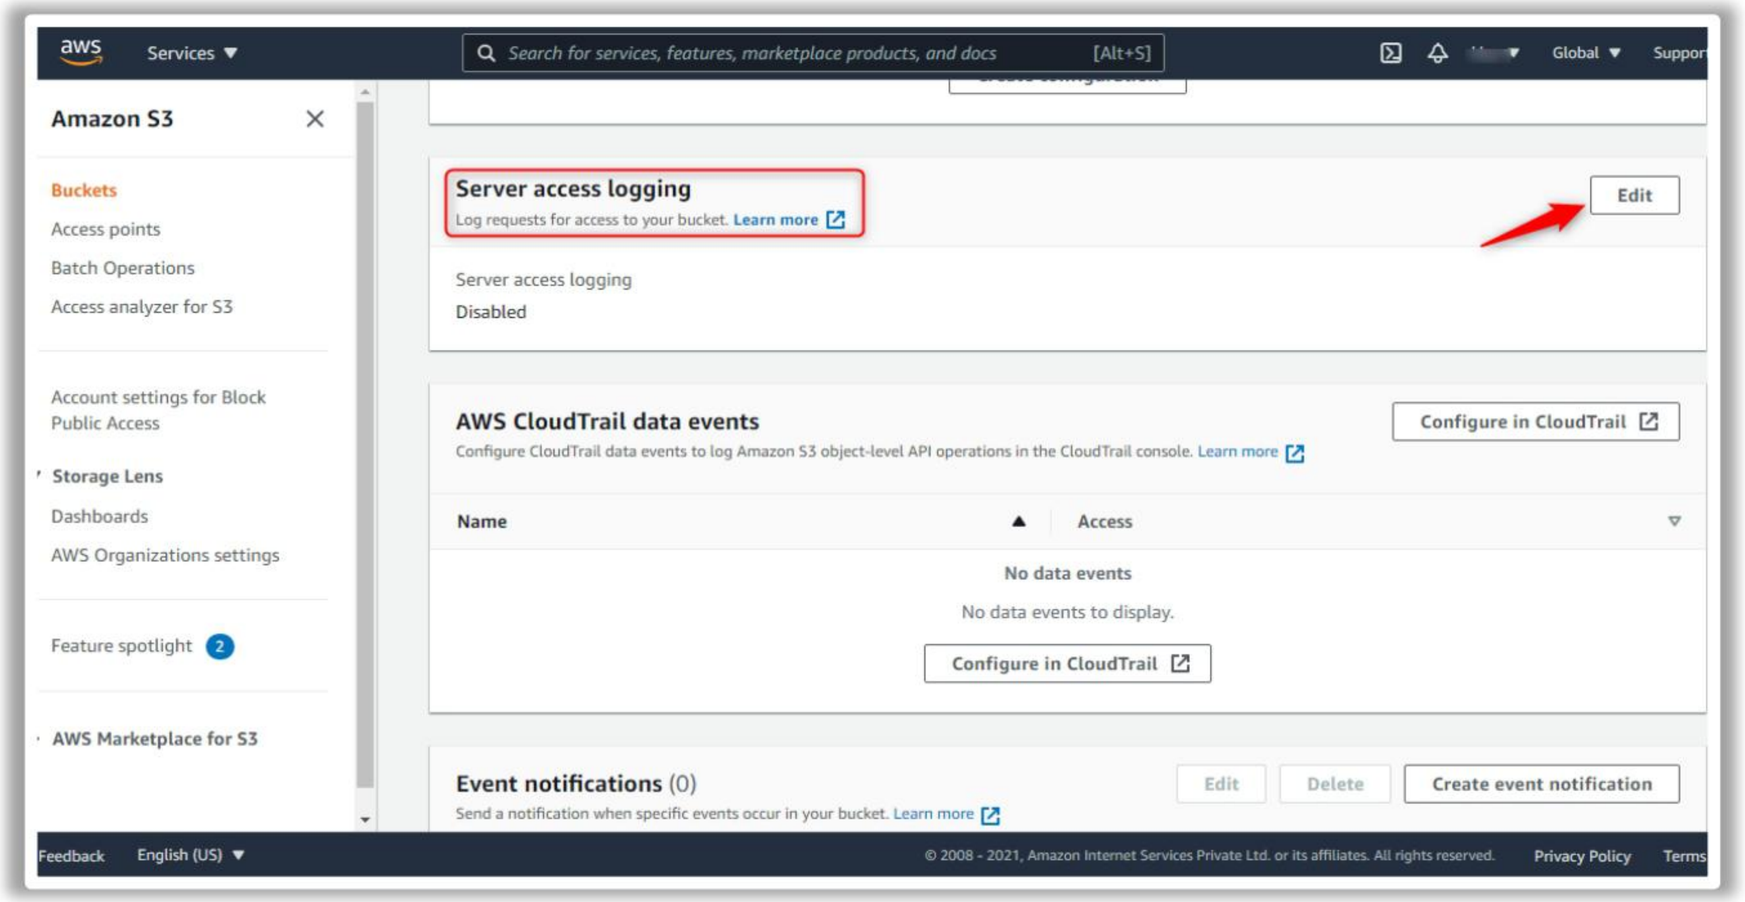


Figure 4.53: Select on Edit Button

- Under **Server access logging**, check the **Enable** radio button. Then, click on **Browse** button.

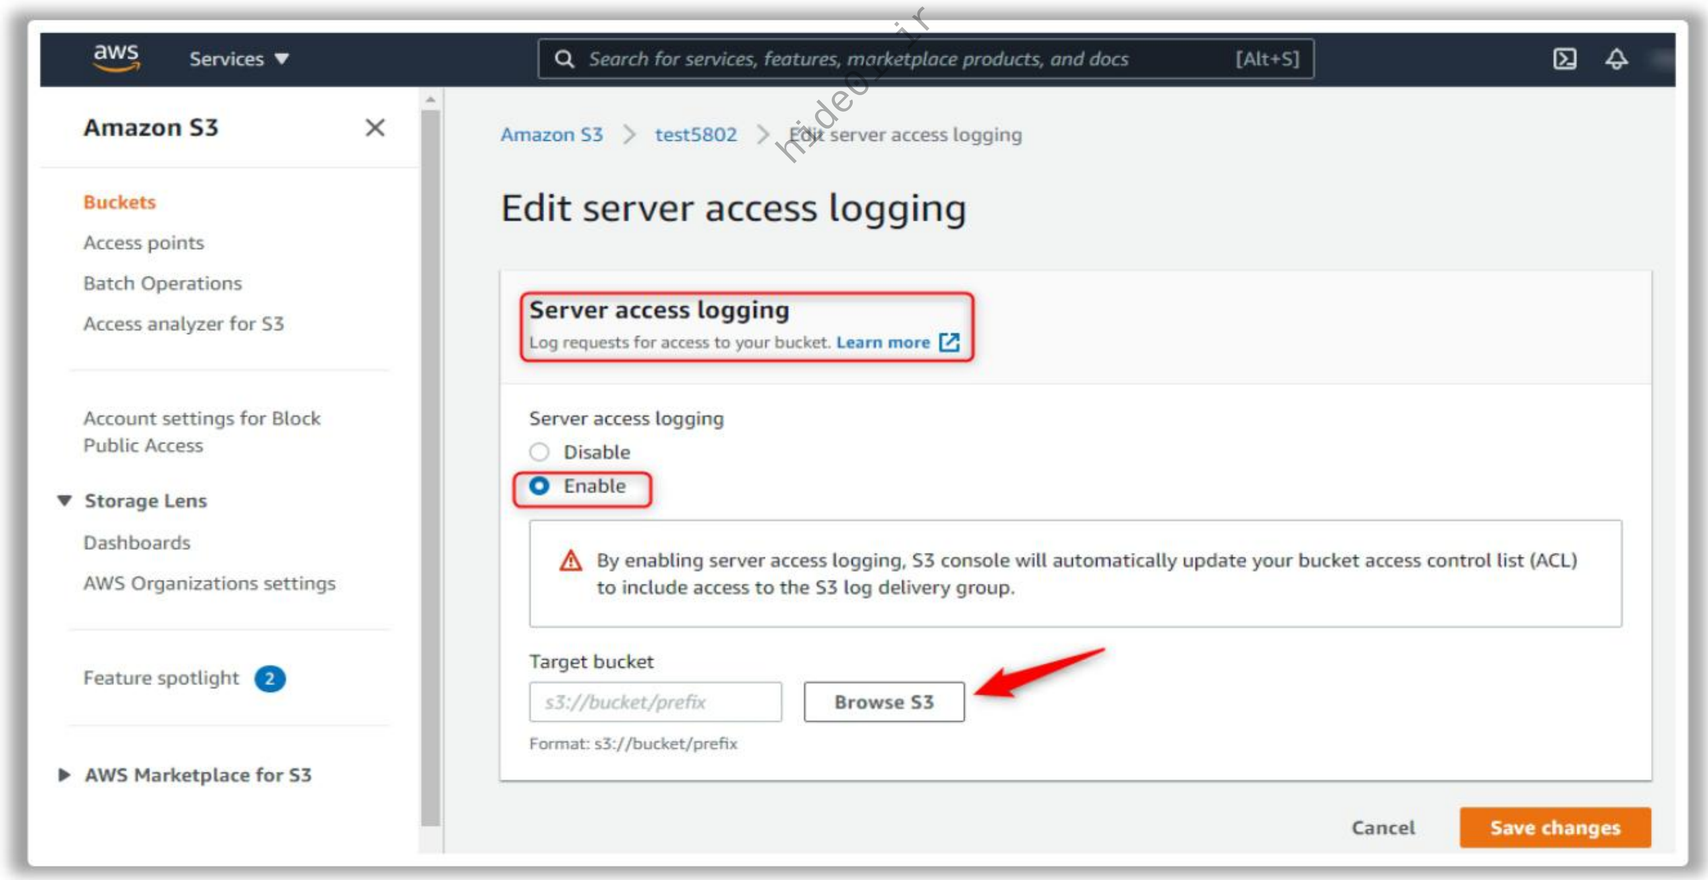


Figure 4.54: Select Enable

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

- Select the **target bucket** from which you want to receive the logs, then click on **Choose path** button

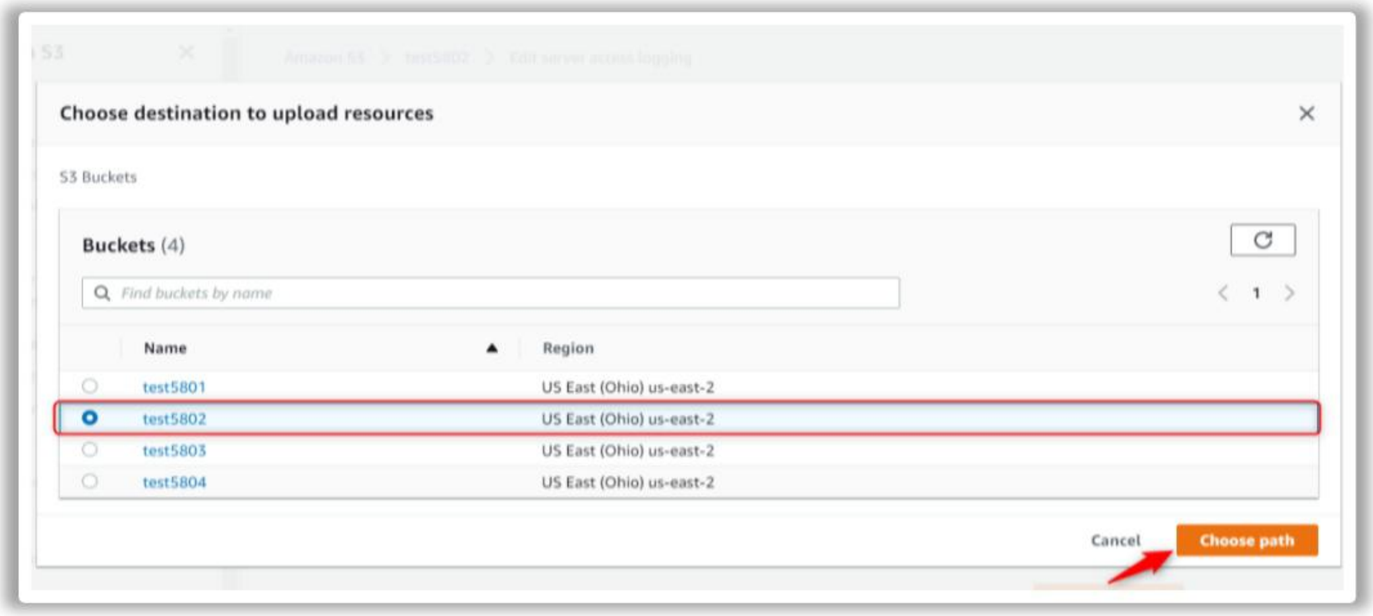


Figure 4.55: Choose Path Button

- Click the **Save changes** button.

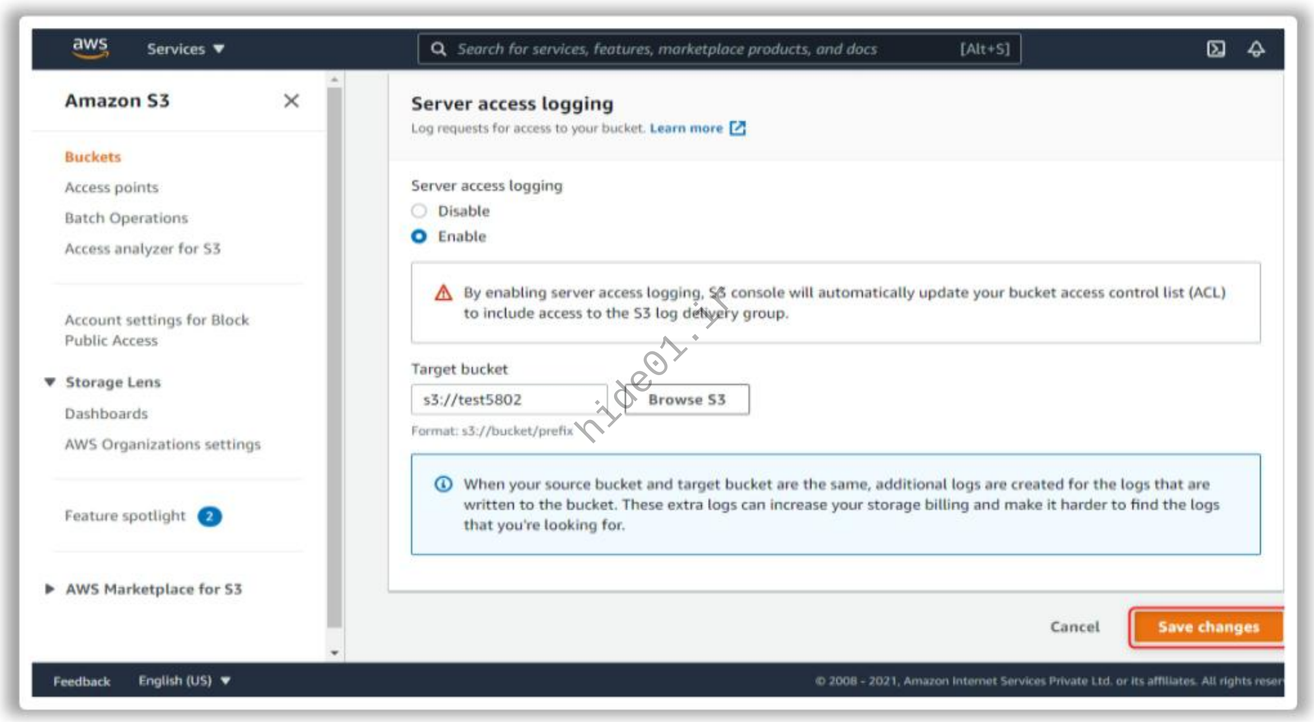


Figure 4.56: Save Changes

- Under **Properties** tab, scroll down to **Server access logging**, you will observe that **Server access logging** is **Enabled**.

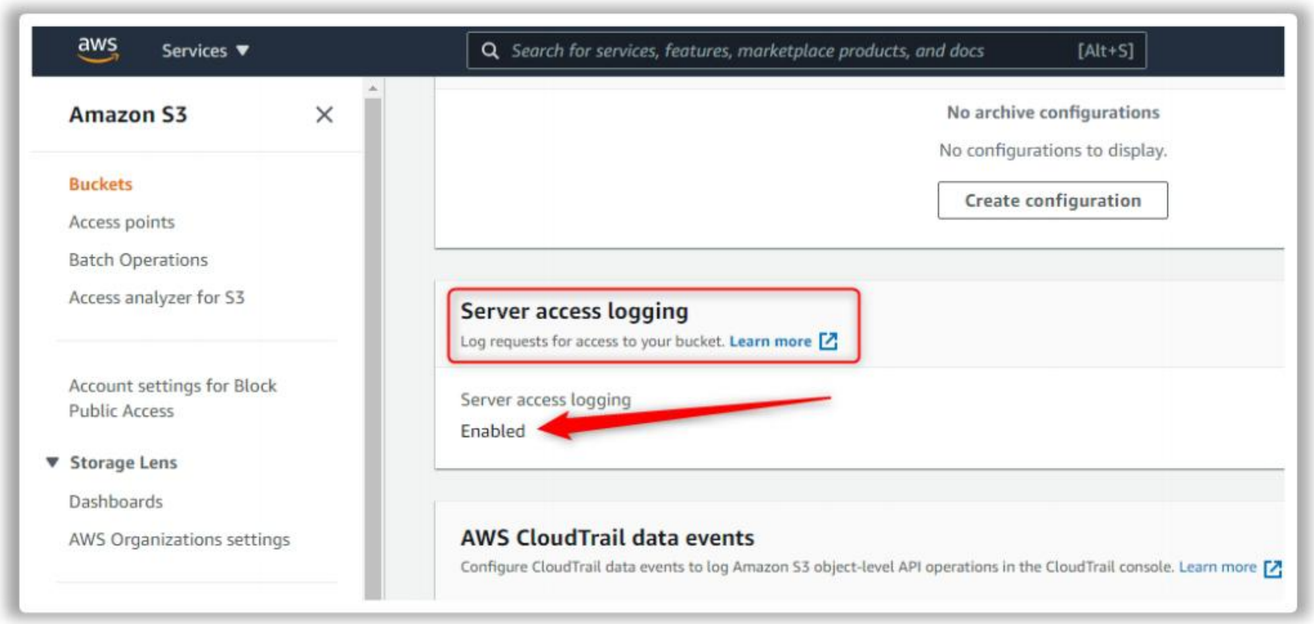
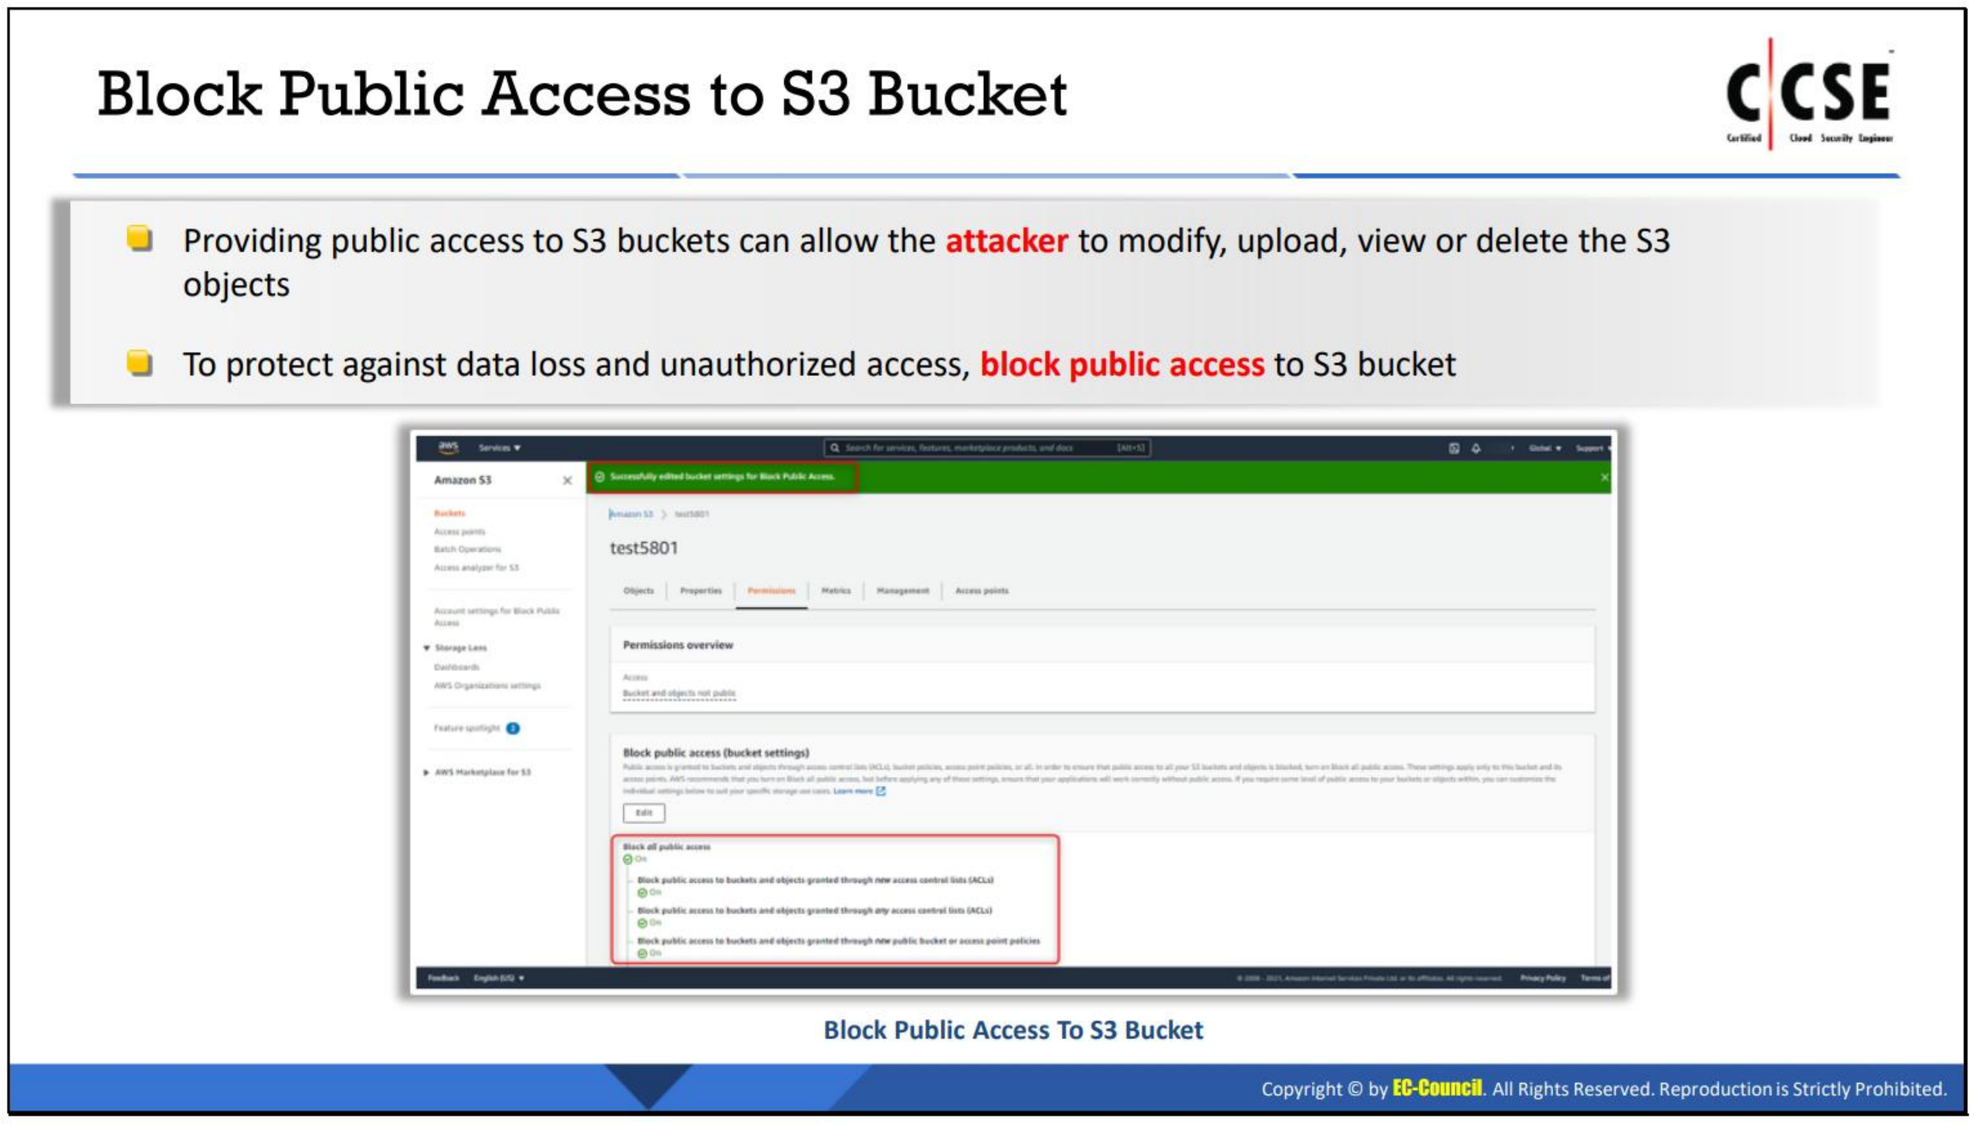


Figure 4.57: Server Access Logging is Enabled

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



Block Public Access to S3 Buckets

Providing public access to S3 buckets can allow the attacker to modify, upload, view, or delete the S3 objects. To protect against data loss and unauthorized access, make sure that there should not be any publicly accessible S3 buckets. Check all the buckets and ensure that public access is blocked in every S3 bucket. If public access to the bucket is not blocked by default or accidentally removed, then block public access to the S3 bucket.

Source: https://aws.amazon.com

To block public access to the S3 bucket, perform the following steps:

- Sign-in to the **AWS Management Console**.
- Click on **Services** dropdown; then click on **S3** under **Storage.**

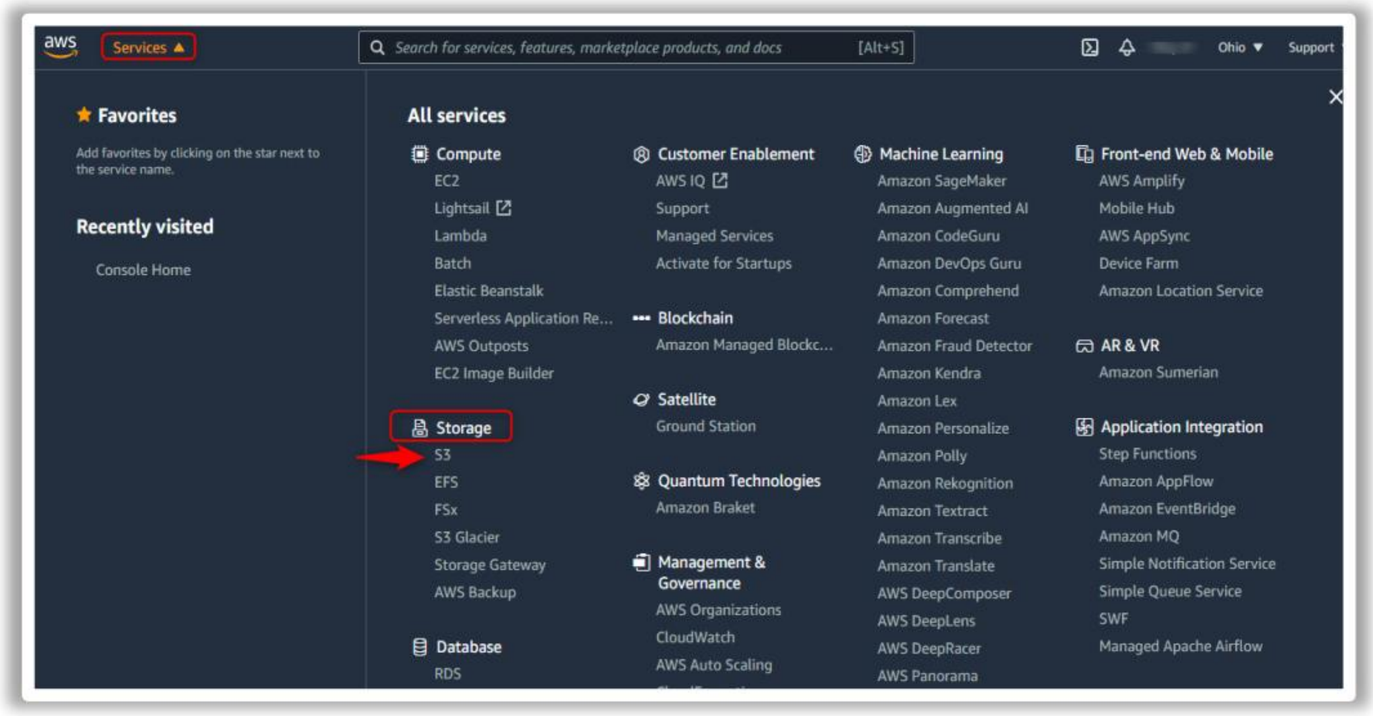


Figure 4.58: Click on Service

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

- The list of available bucketswill be displayed.

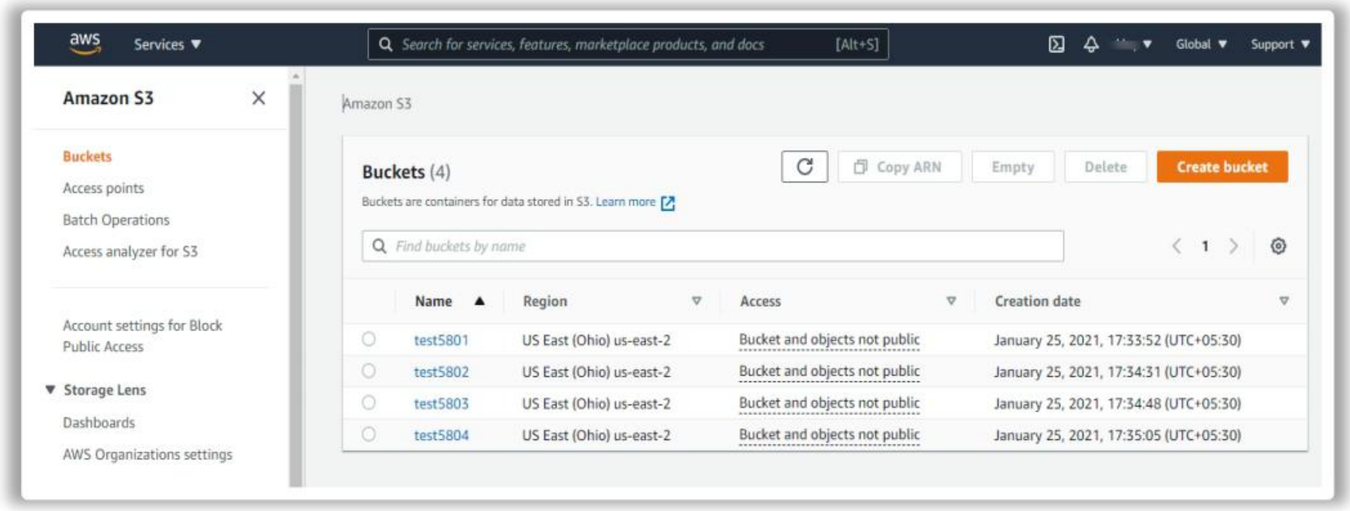


Figure 4.59: List of Buckets

- Select the Bucket that you want to encrypt and click on it. Then, click on **Permissions** tab and click on Edit button.

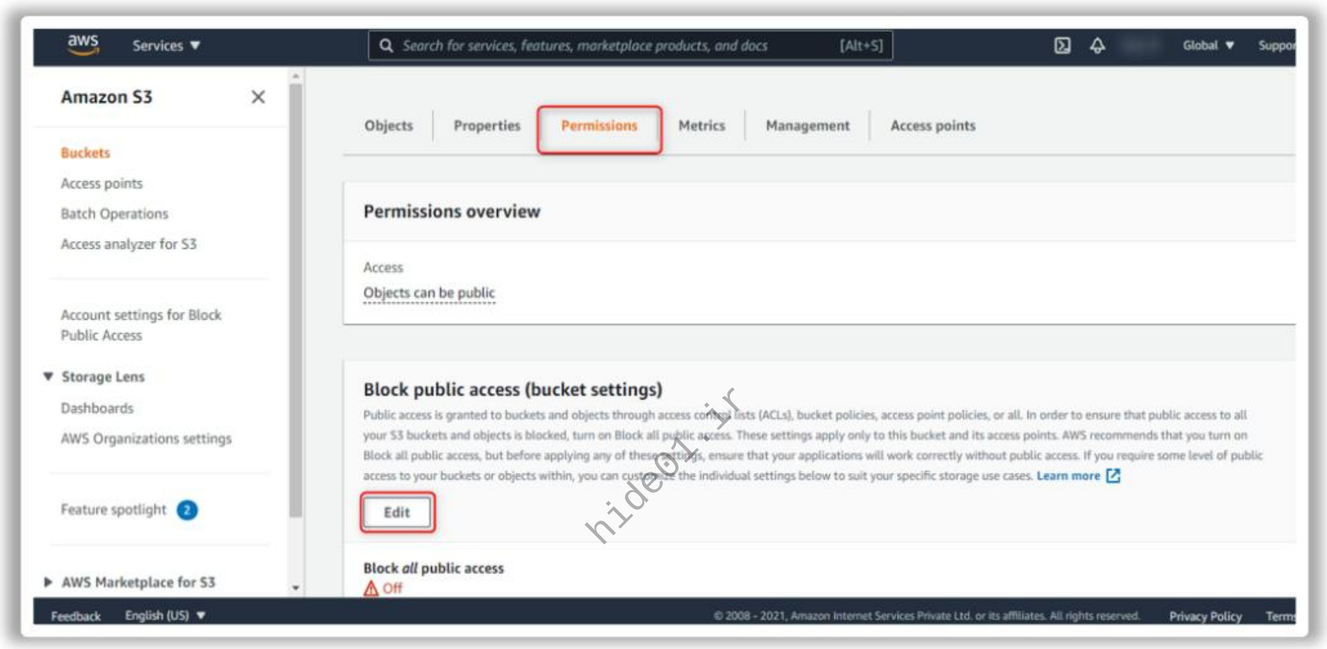


Figure 4.60: Click on Edit

- Under **Edit Block public access (bucket settings)**, check **Block all public access** and click **Save changes** button.

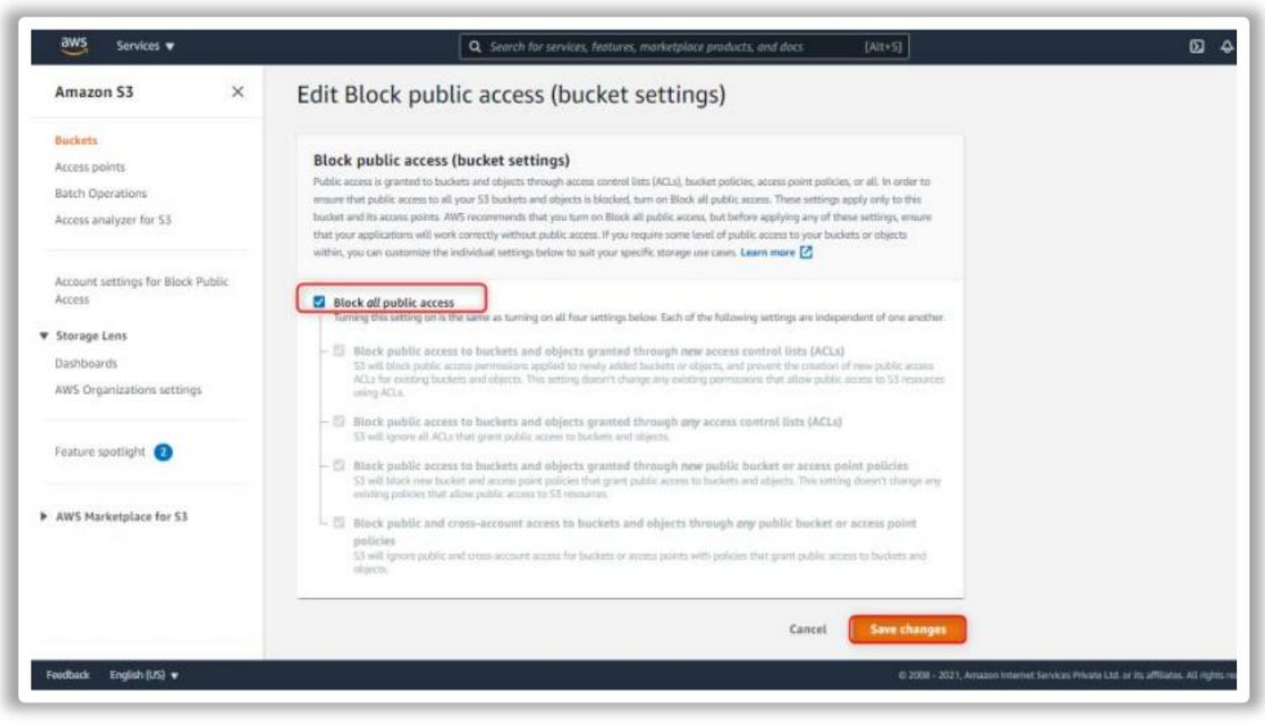


Figure 4.61: Save Changes

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

- **Edit Block public access (bucket settings)** window will open, type **confirm** in the provided field and click the **Confirm** button.

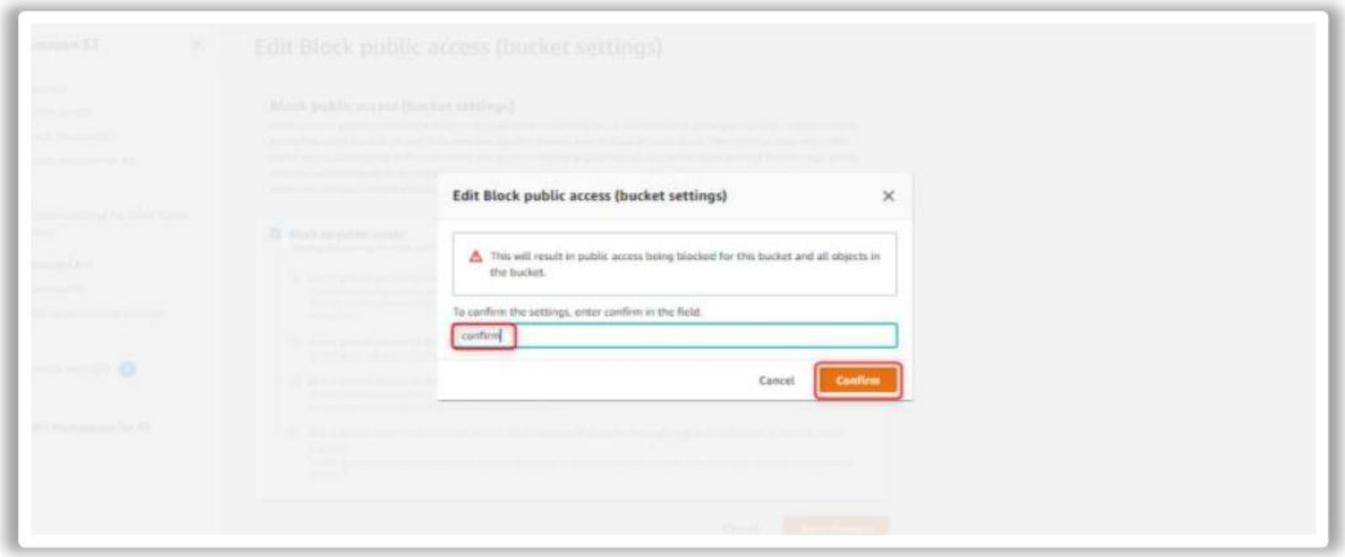


Figure 4.62: Click Confirm

- Public access to the respective bucket is successfully blocked.

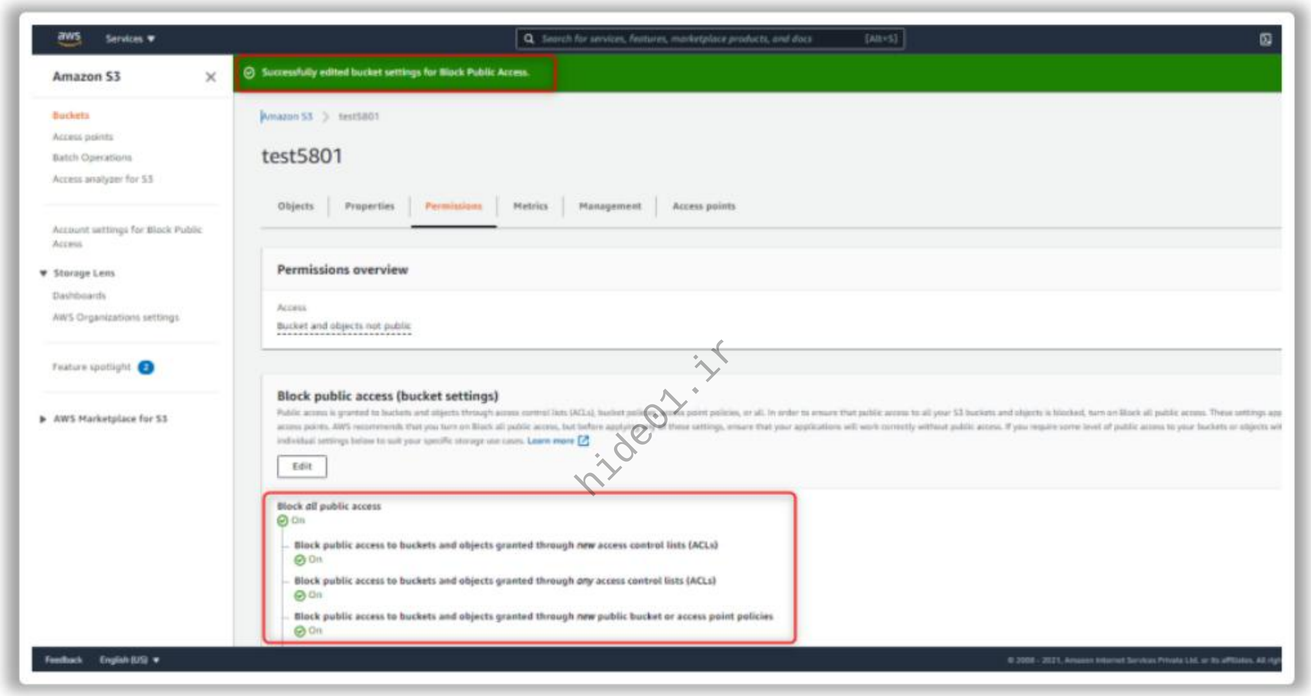
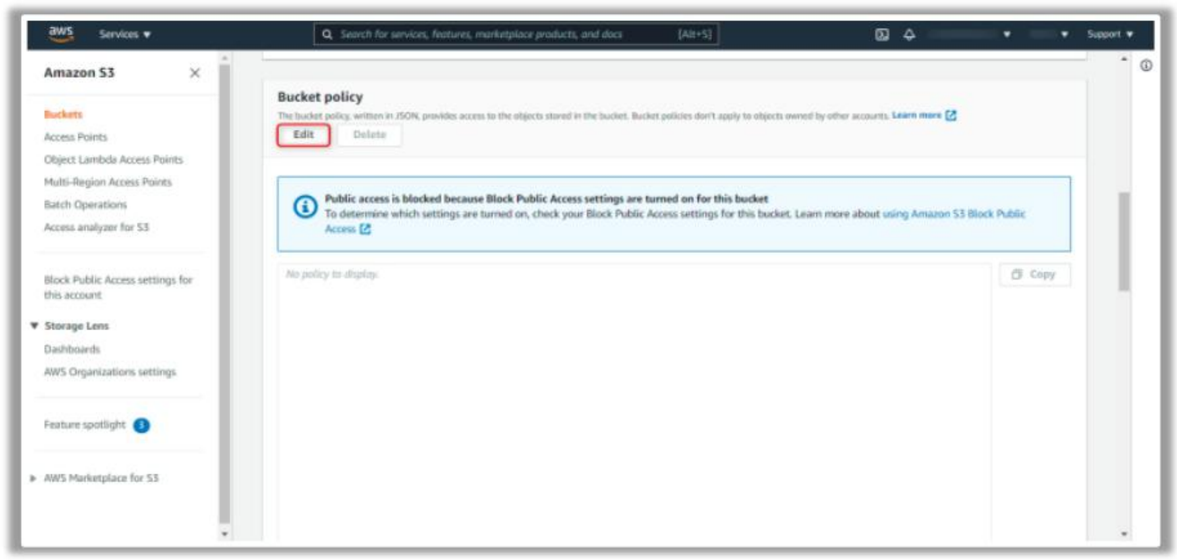


Figure 4.63: Public Access is Blocked

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

Restrict Public Access to S3 buckets via Bucket Policy

- Granting **public access** to S3 buckets can lead to exploitation of S3 objects
- To protect against **unauthorized access**, restrict public access to S3 buckets through bucket policy



Disabling Public Access To S3 Buckets Via Bucket Policy

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

Restrict Public Access to S3 Buckets via Bucket Policies

Providing public access to S3 buckets through bucket policies can allow the attacker to modify, upload, view, or delete the S3 objects. If there is no restriction on accessing bucket policy, then any one can download, upload, delete, modify, or list the objects within the buckets. Hence, implement the bucket policy to restrict access to specific AWS accounts rather than allowing access to everyone on the internet. To protect against unauthorized access, make sure that the S3 buckets cannot have public access through bucket policies.

Source: https://aws.amazon.com

To restrict public access to S3 buckets via bucket policies, perform the following steps:

- Sign-in to the **AWS Management Console**.
- Click on **Services** dropdown; then click on **S3** under **Storage.**

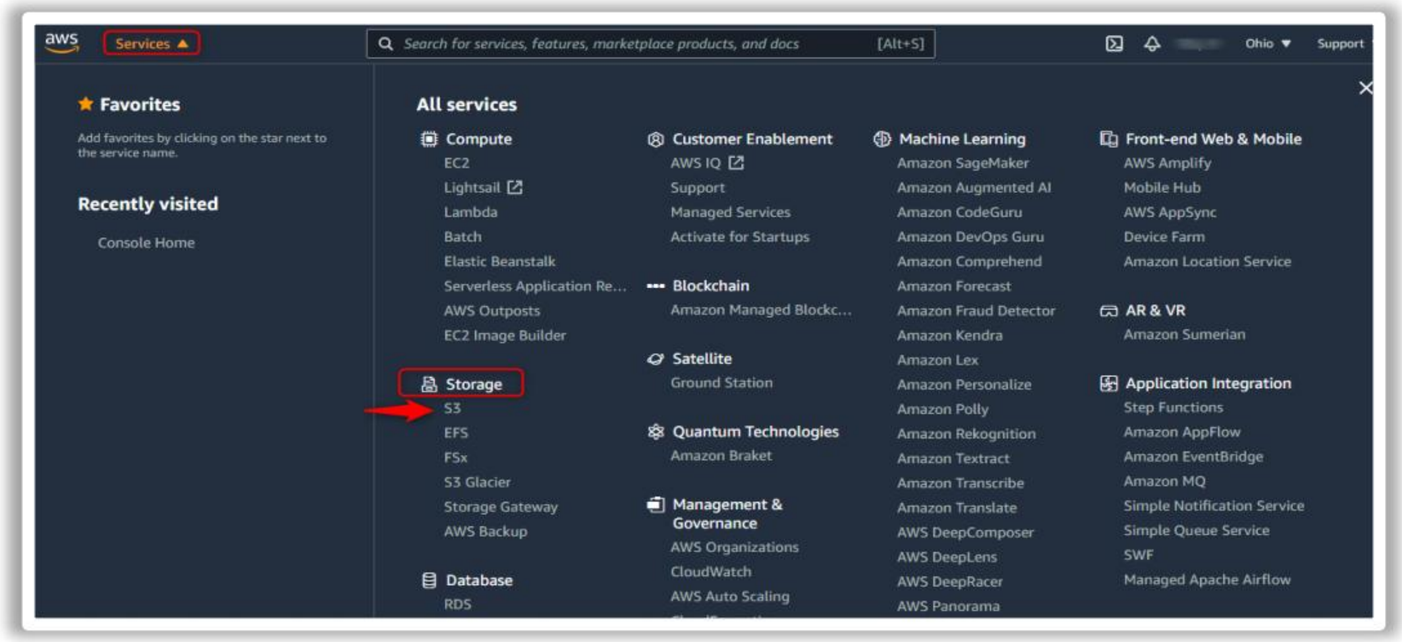


Figure 4.64: Click on Service

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

- The list of available buckets will be displayed.

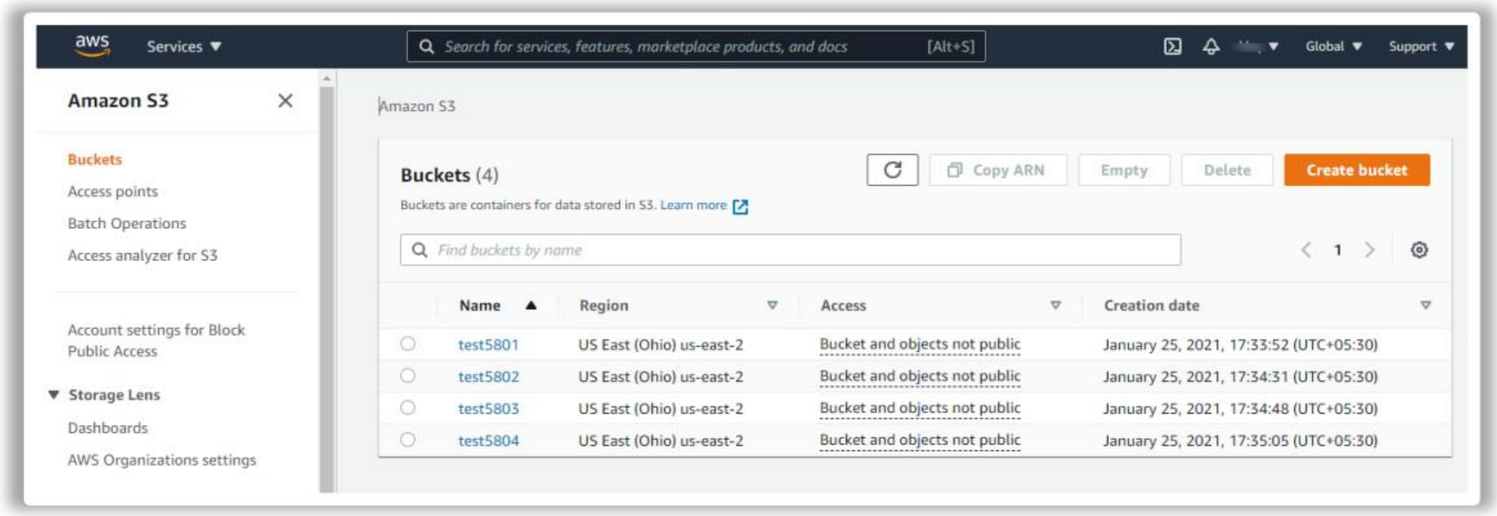


Figure 4.65: List of Buckets

- Select the Bucket that you want to restrict public access via bucket policies. Click on **Permissions** tab and then navigate to **Bucket policy**.

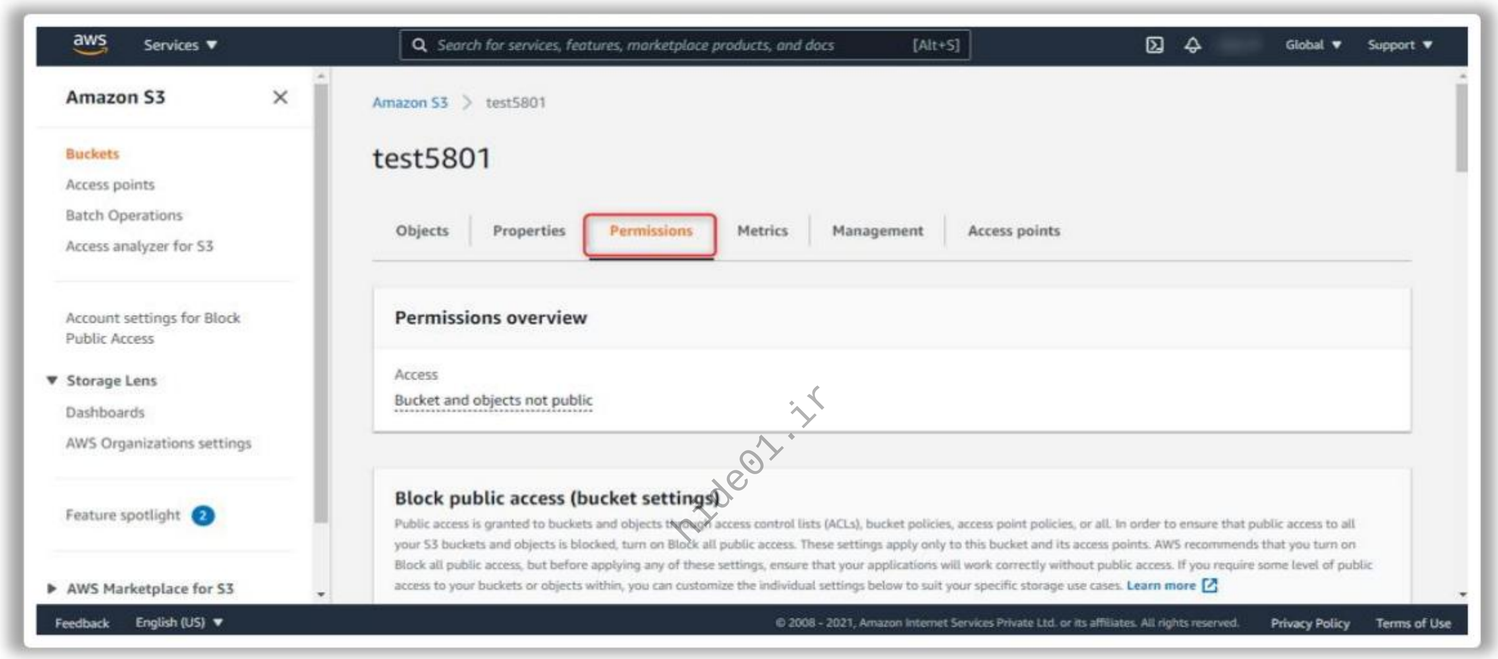


Figure 4.66: Click on Permissions

- Click the **Delete** button to restrict the **Bucket policy**.

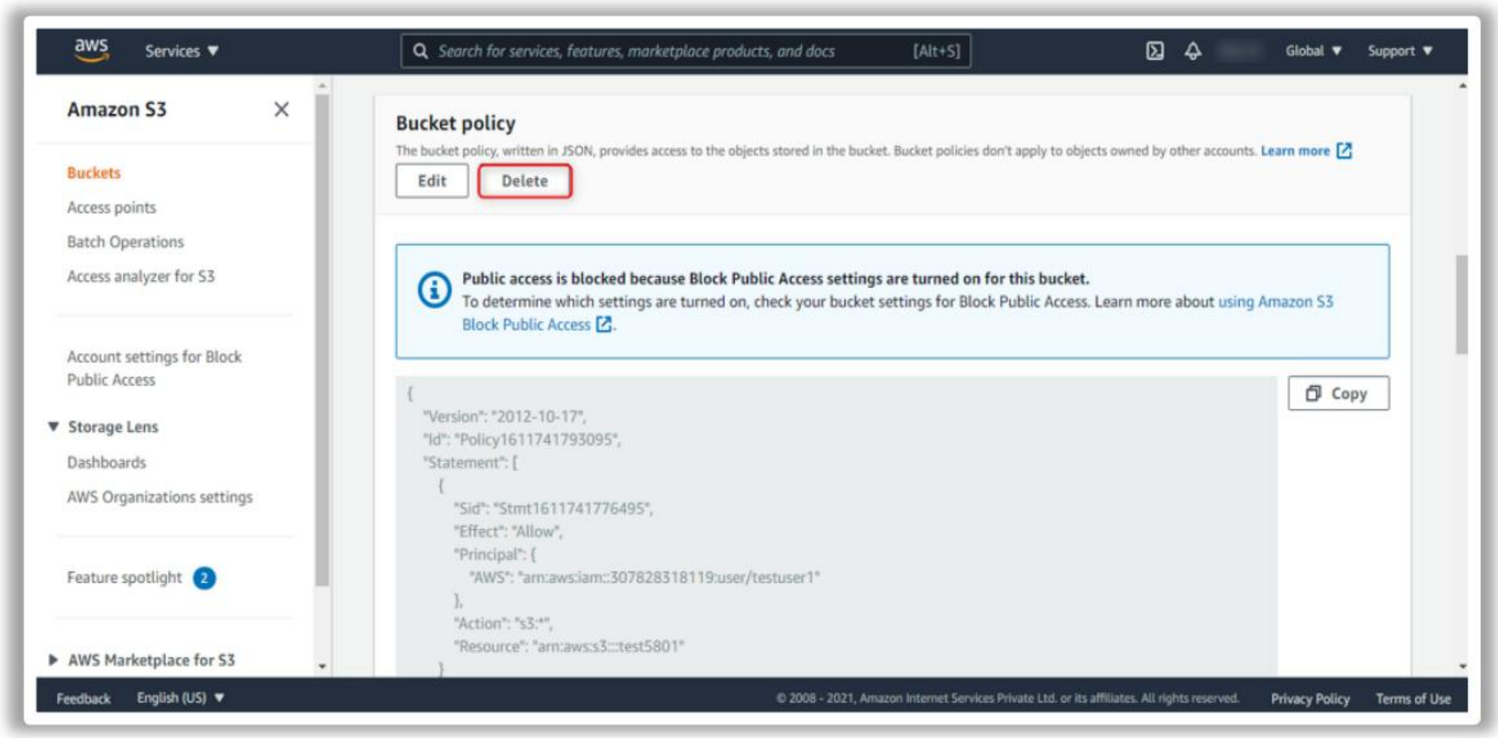


Figure 4.67: Click on Delete

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

- **Delete bucket policy?** window will appear; type **delete** in the field and then click the **Delete** button. (If you want to restrict the public access to a particular AWS IAM user/AWS account, replace the current value of the principal element with Amazon Resource Number (ARN) of the AWS account; then click the **Save** button.)

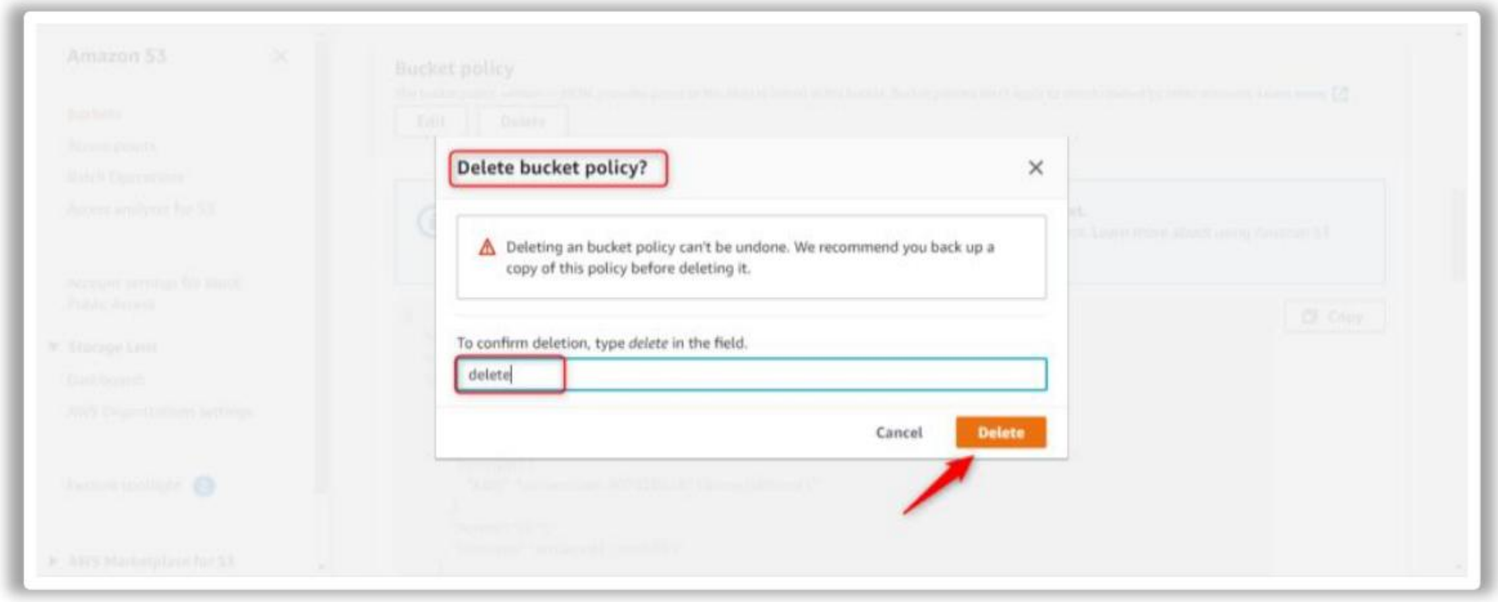


Figure 4.68: Click Save

- Public access to S3 buckets via bucket policies will be restricted.

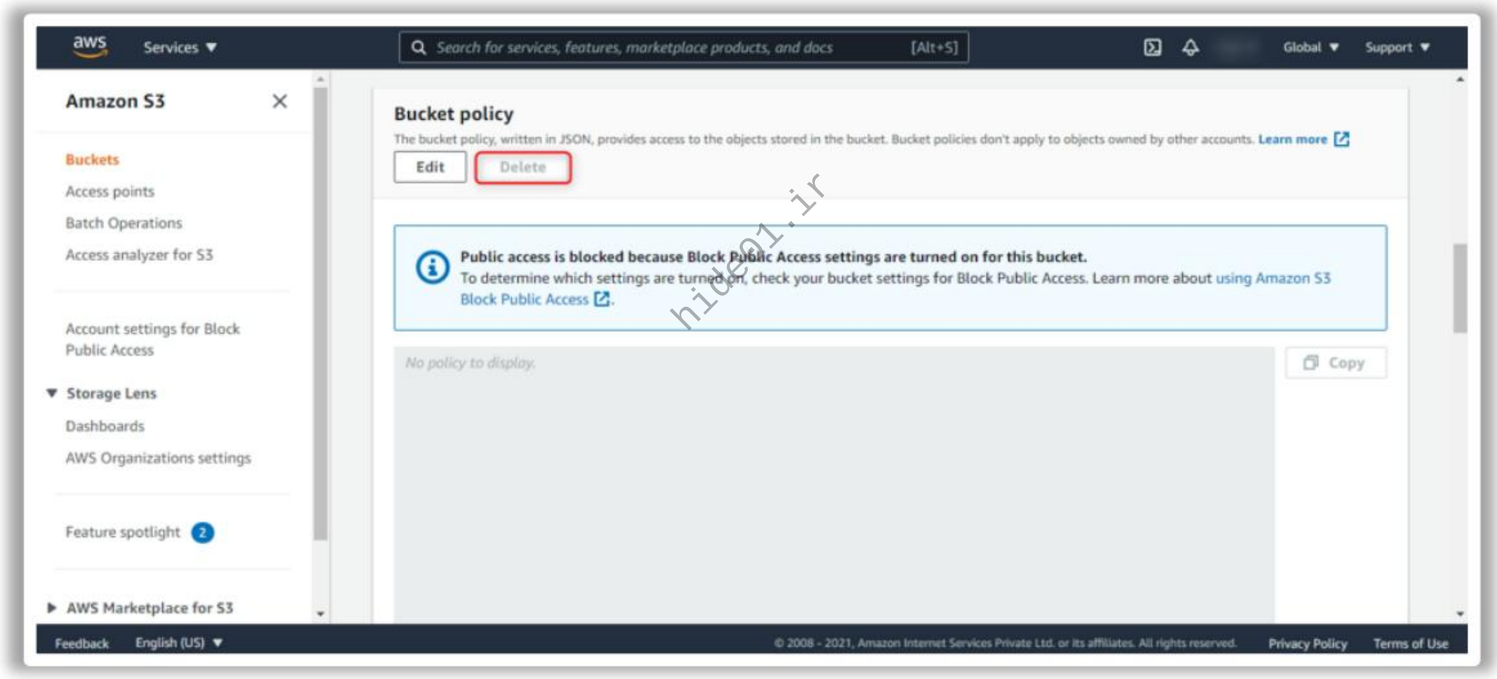
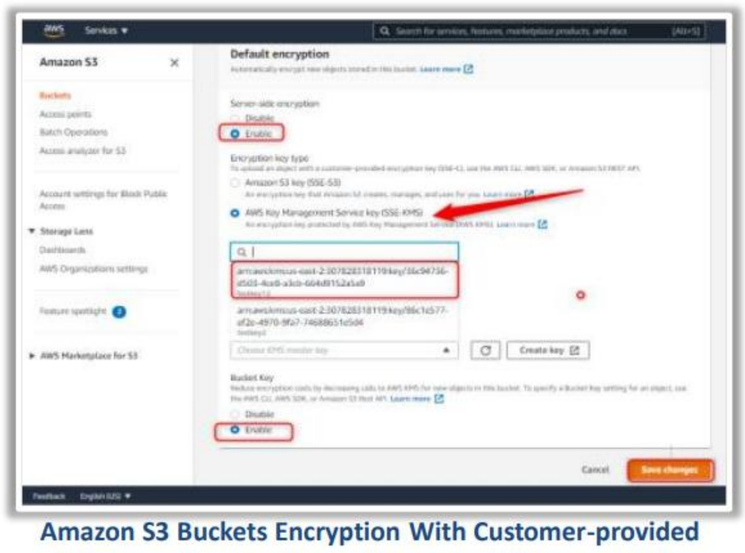


Figure 4.69: Restrict Public Access to S3 Buckets

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

Ensure that Amazon S3 buckets are Encrypted with Customer-Provided AWS KMS CMKs

- The cloud consumer can set their **own encryption keys** utilizing server-side encryption with customer managed CMKs
- For fine-grained control of Amazon S3 data-at-rest encryption and decryption process, encrypt Amazon S3 buckets using customer-provided **AWS KMS CMKs**



Amazon S3 Buckets Encryption With Customer-provided AWS KMS CMKs

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

Ensure that Amazon S3 Buckets are Encrypted with Customer-provided AWS KMS CMKs

The cloud consumer can set their own encryption keys utilizing server-side encryption with customer-managed CMKs. For fine-grained control of Amazon S3, data-at-rest encryption and decryption process, make sure that Amazon S3 buckets are encrypted with customer-provided AWS KMS CMKs.

Source: https://aws.amazon.com

To encrypt Amazon S3 buckets with customer-provided AWS KMS CMKs, perform the following steps:

- Sign-in to the **AWS Management Console**.
- Click on **Services** dropdown; then click on **S3** under **Storage.**

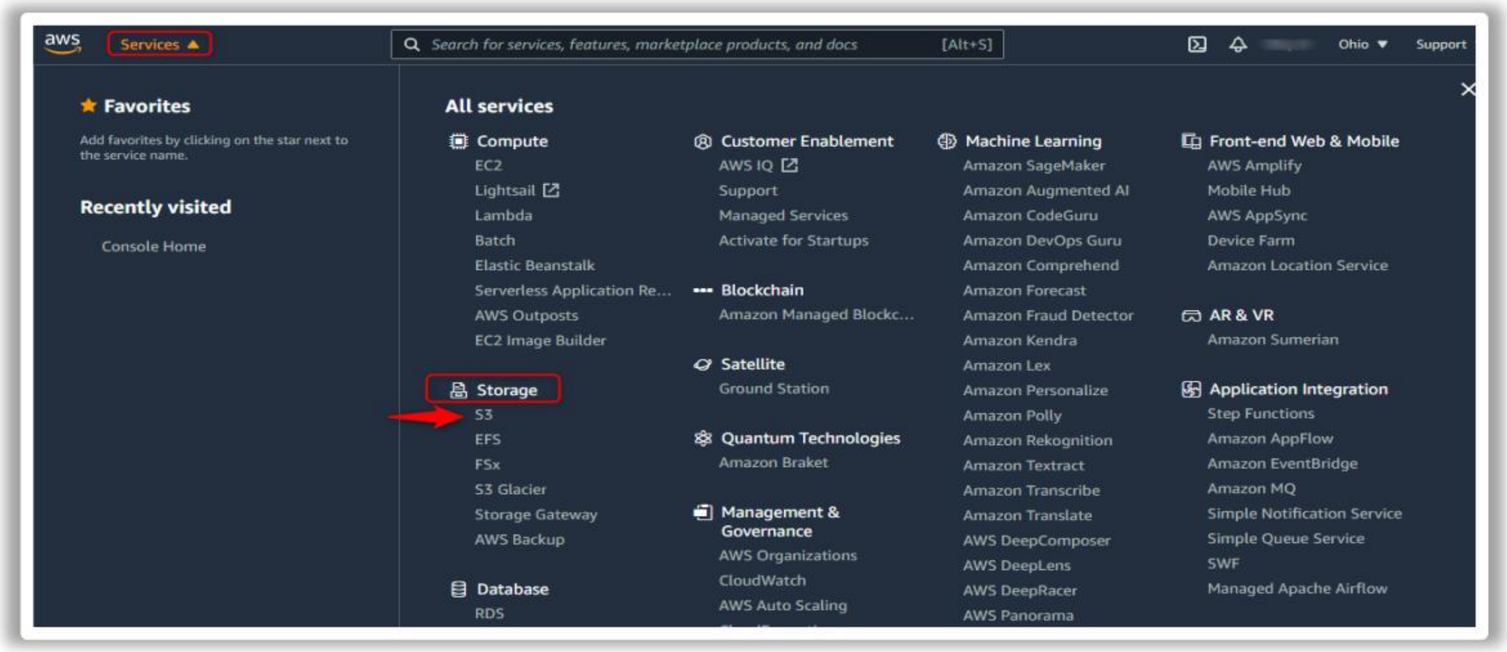


Figure 4.70: Click on Storage

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

- The list of available buckets will be displayed.

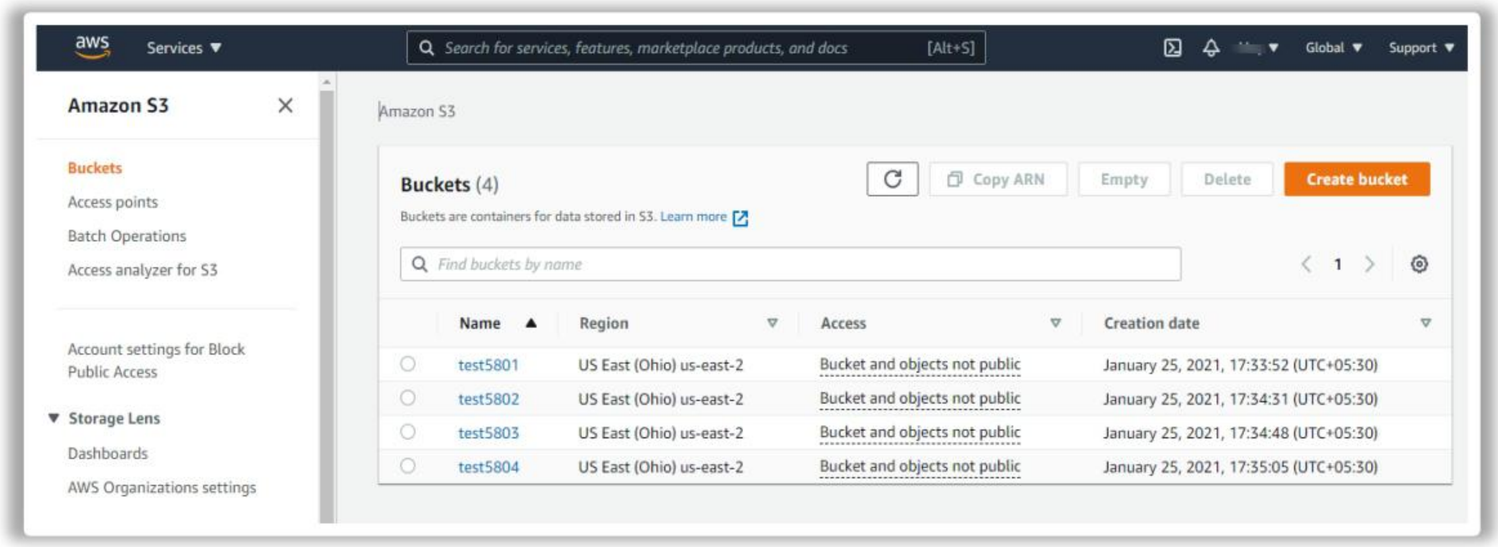
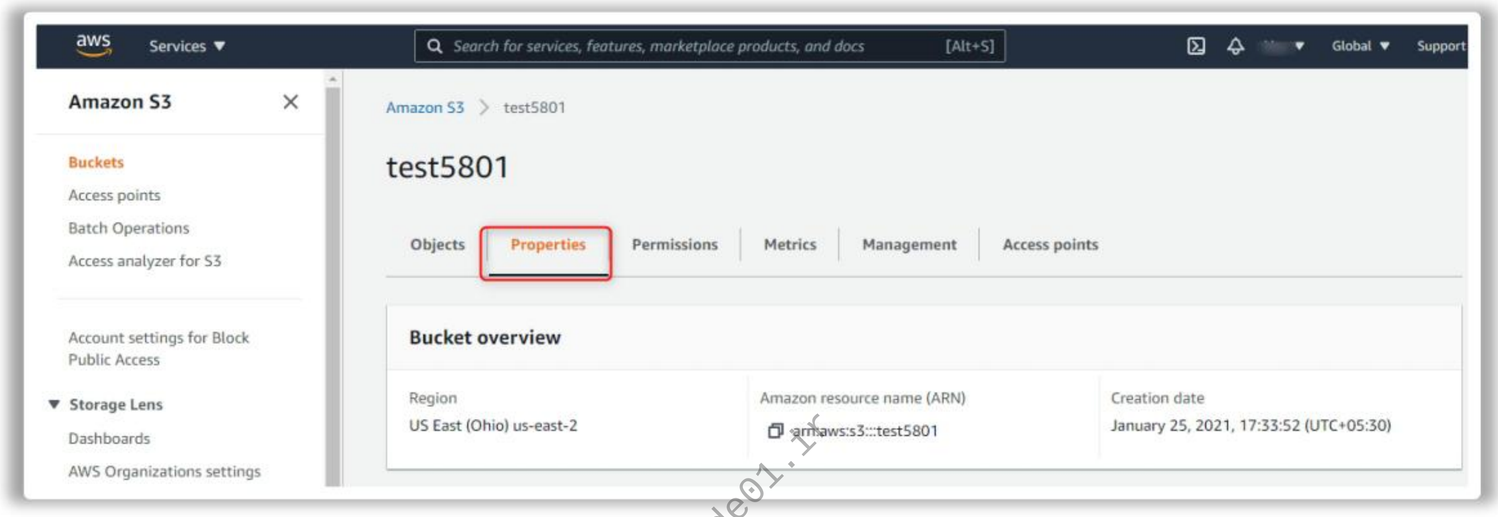


Figure 4.71: List of Buckets

- Select the Bucket that you want to encrypt and click on it. Then, click on **Properties** tab.



hide01.ir

Figure 4.72: Click on Properties

- Scroll down and click the **Edit** button in **Default encryption** section.

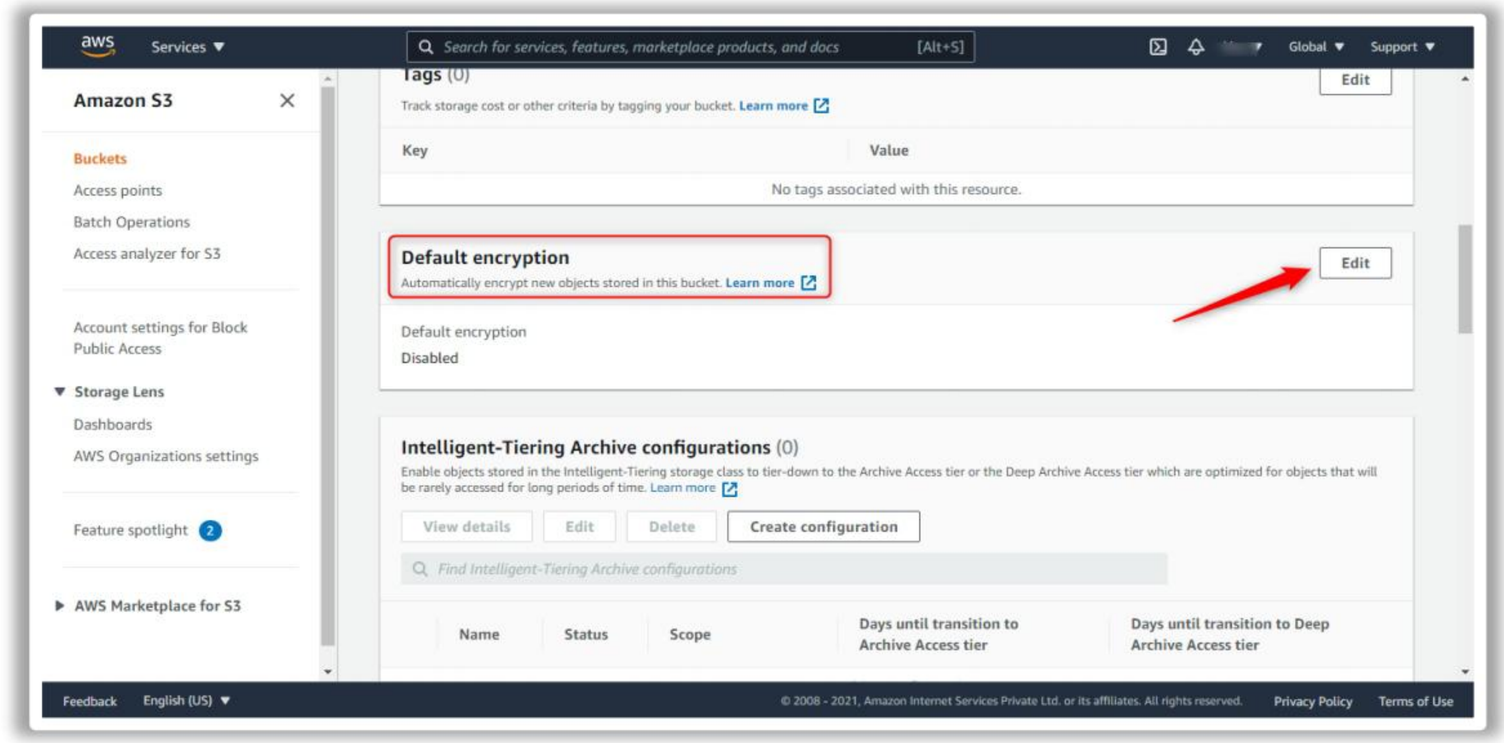


Figure 4.73: Select Edit

- Navigate to **Server-side encryption**; check **Enable** radio button.
- Under **Encryption key type**, check **AWS Key Management Service key (SSE-KMS)**, and then navigate to **KMS master key,** click on dropdown button and choose the KMS master key**.**

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

- Under **Bucket key,** check the **Enable** button and click on **Save changes** button**.**

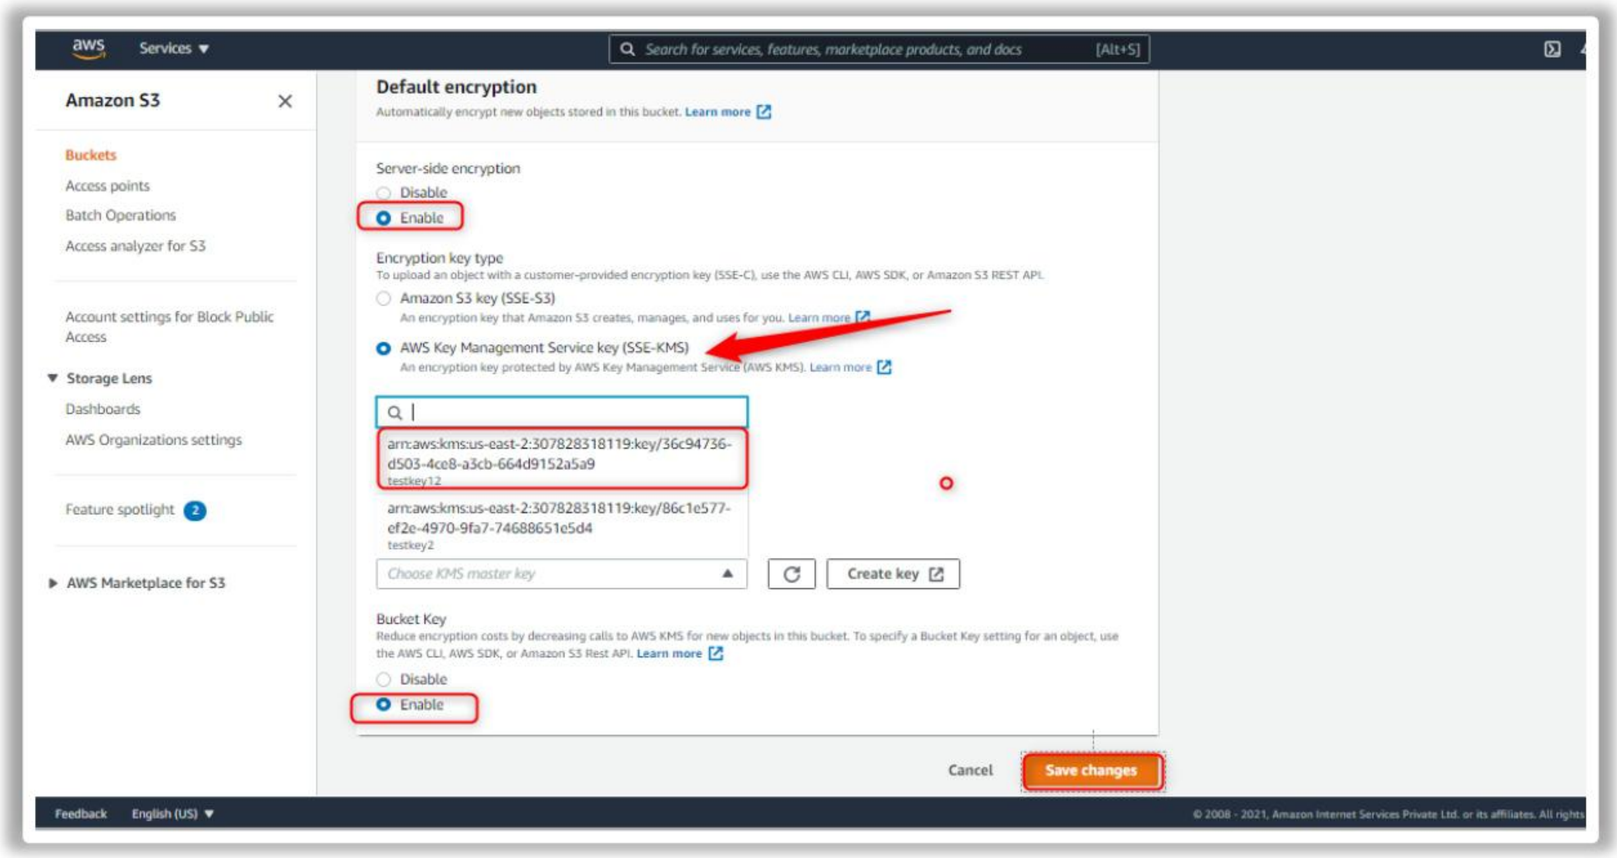
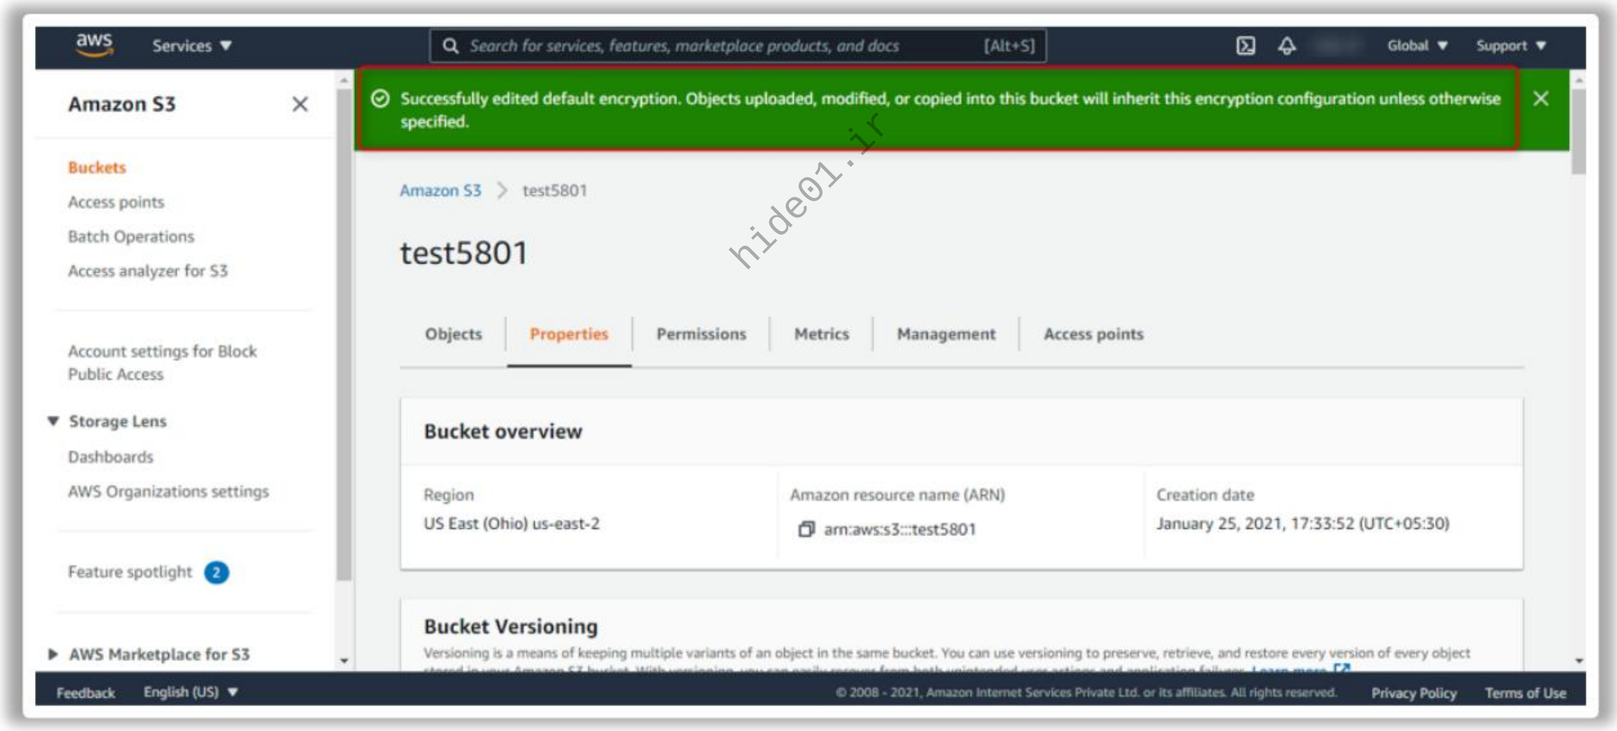


Figure 4.74: Save Changes

- Amazon S3 buckets with customer-provided AWS KMS CMKs is successfully encrypted.



hide01.ir

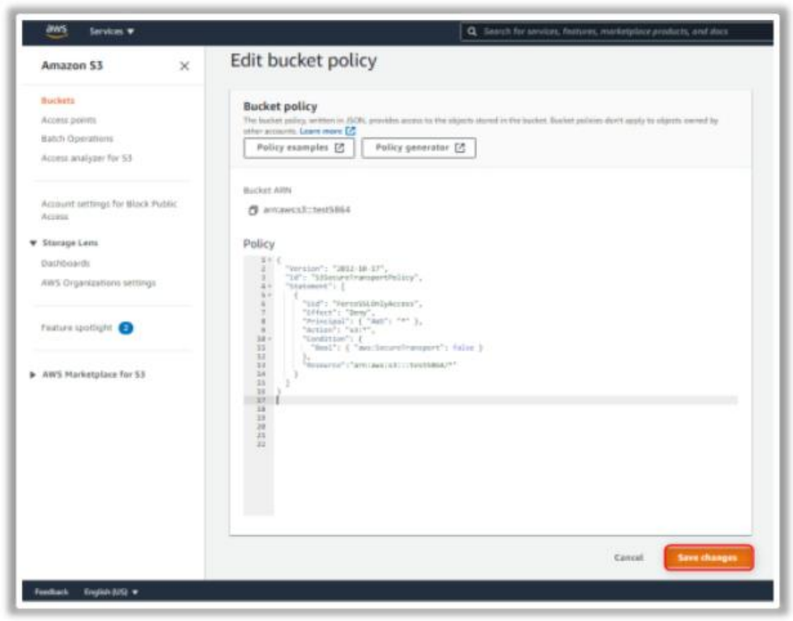
Figure 4.75: AWS KMS CMK is encrypted

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

Restrict Non-SSL S3 Access to all the Objects in Amazon S3 bucket

To safeguard the communication between the S3 buckets and the clients against eavesdropping and man-in-the-middle attack, restrict **non-SSL S3 access** to all the objects in Amazon S3 bucket

```
{
  "Version": "2012-10-17",
  "Id": "S3SecureTransportPolicy",
  "Statement": [
    {
      "Sid": "ForceSSLOnlyAccess",
      "Effect": "Deny",
      "Principal": { "AWS": "*" },
      "Action": "s3:*",
      "Condition": {
        "Bool": { "aws:SecureTransport": false }
      },
      "Resource":"arn:aws:s3:::test5864/*"
    }
  ]
}
```



Restricting Non-SSL S3 access to All Objects in Amazon S3 bucket

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

Restrict Non-SSL S3 Access to All the Objects in Amazon S3 Bucket

To safeguard the communication between the S3 buckets and clients against eavesdropping and MIM attack, ensure that the S3 buckets implement SSL-only access.

hide01.ir

Source: https://aws.amazon.com

To restrict non-SSL S3 access to all the objects in Amazon S3 bucket, perform the following steps:

- Sign-in to the **AWS Management Console**.
- Click on **Services** dropdown; then click on **S3** under **Storage.**

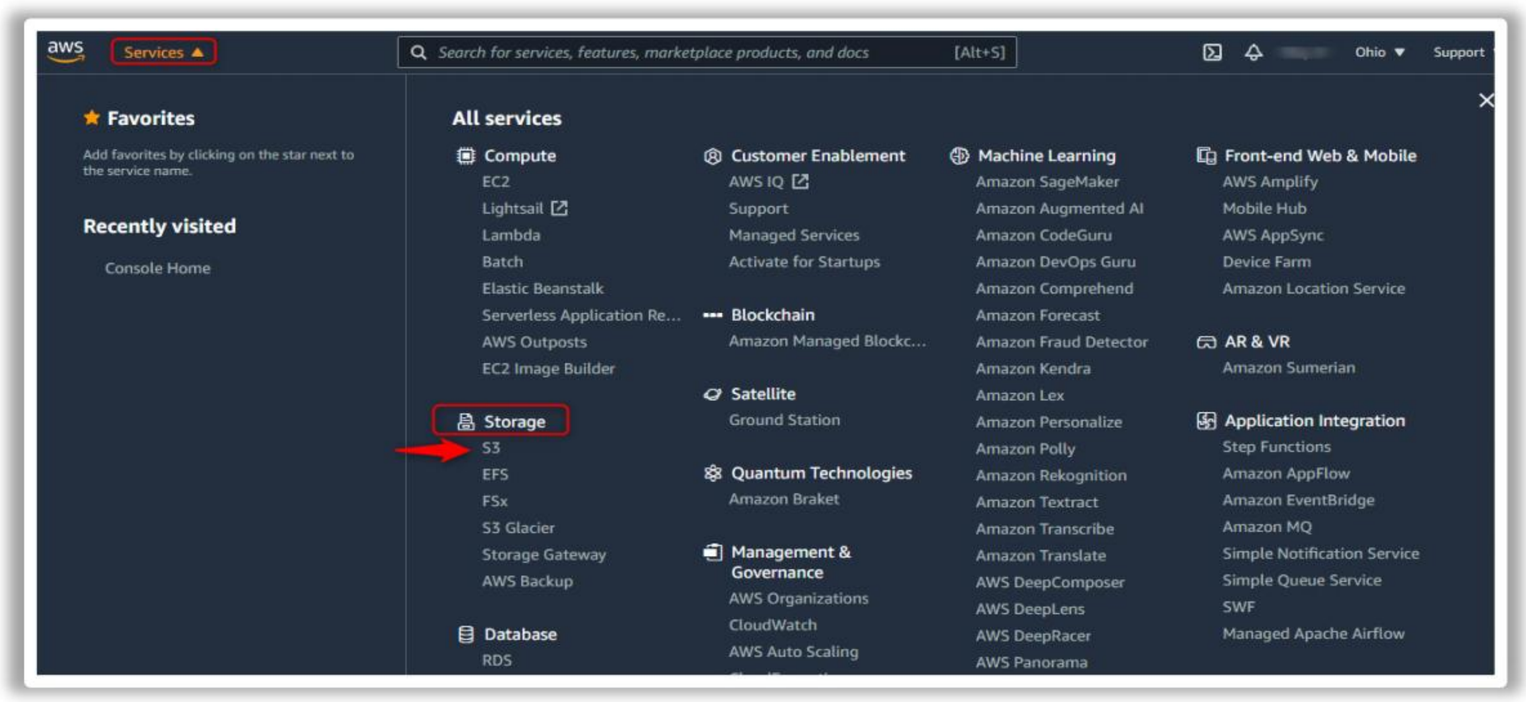


Figure 4.76: Click on Storage

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

- The list of available buckets will be displayed.

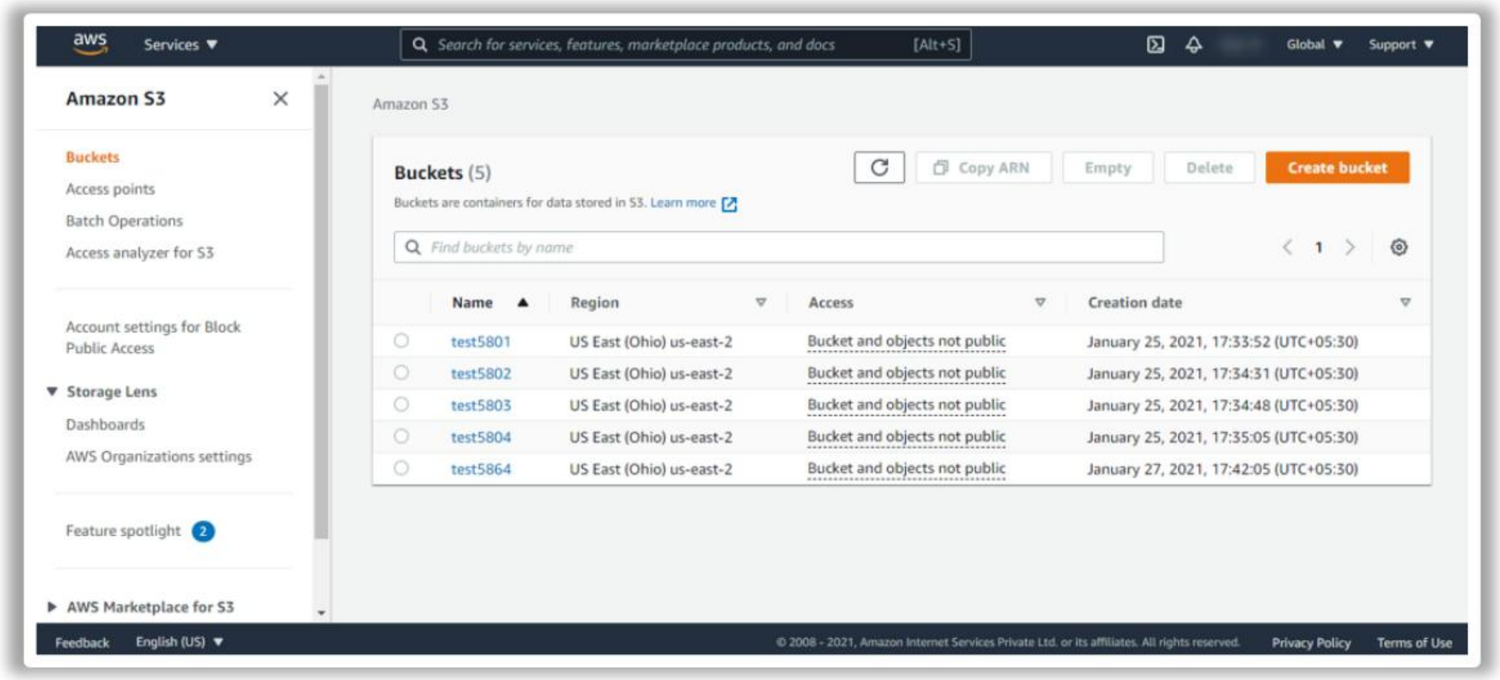
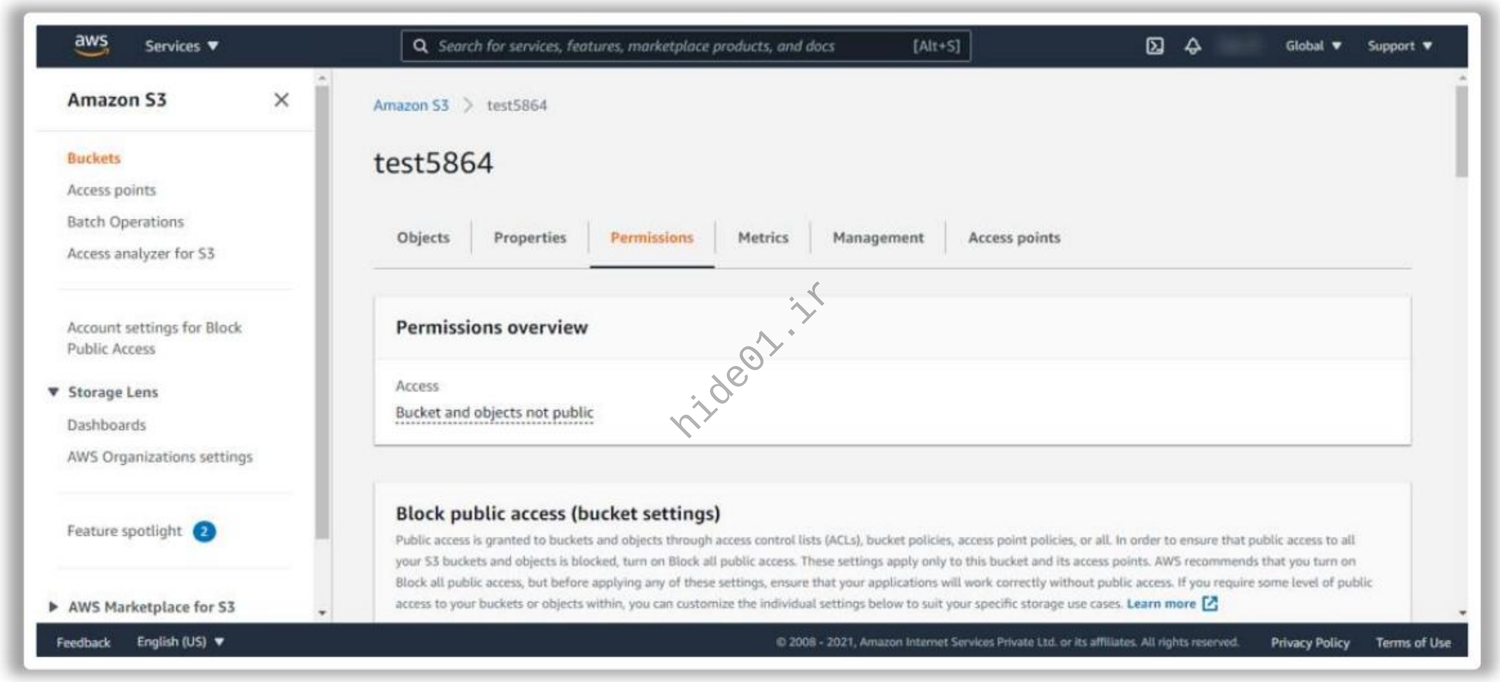


Figure 4.77: List of Buckets

- Select the Bucket that you want to encrypt and click on it. Then, click on **Properties** tab.



hide01.ir

Figure 4.78: Click on Properties

- Under **Bucket policy**, click **Edit** button.

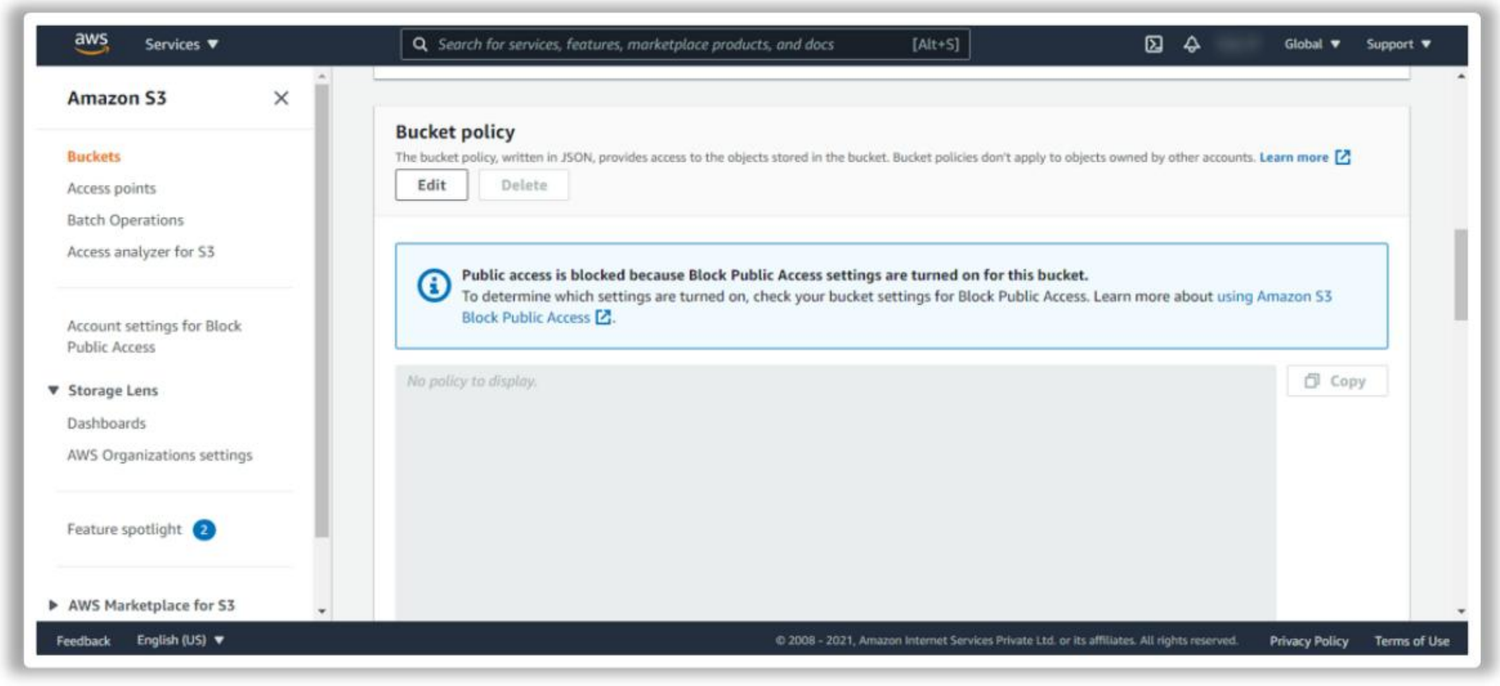


Figure 4.79: Select Edit Button

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

- Currently, if you do not have any access policy in use, then in the bucket policy editor paste the following policy document. Replace the **test5864** bucket name with your own bucket name and click **Save**.

```
{
  "Version": "2012-10-17",
  "Id": "S3SecureTransportPolicy",
  "Statement": [
    {
      "Sid": "ForceSSLOnlyAccess",
      "Effect": "Deny",
      "Principal": { "AWS": "*" },
      "Action": "s3:*",
      "Condition": {
        "Bool": { "aws:SecureTransport": false }
      },
      "Resource":"arn:aws:s3:::test5864/*"
    }
  ]
}
```

hide01.ir

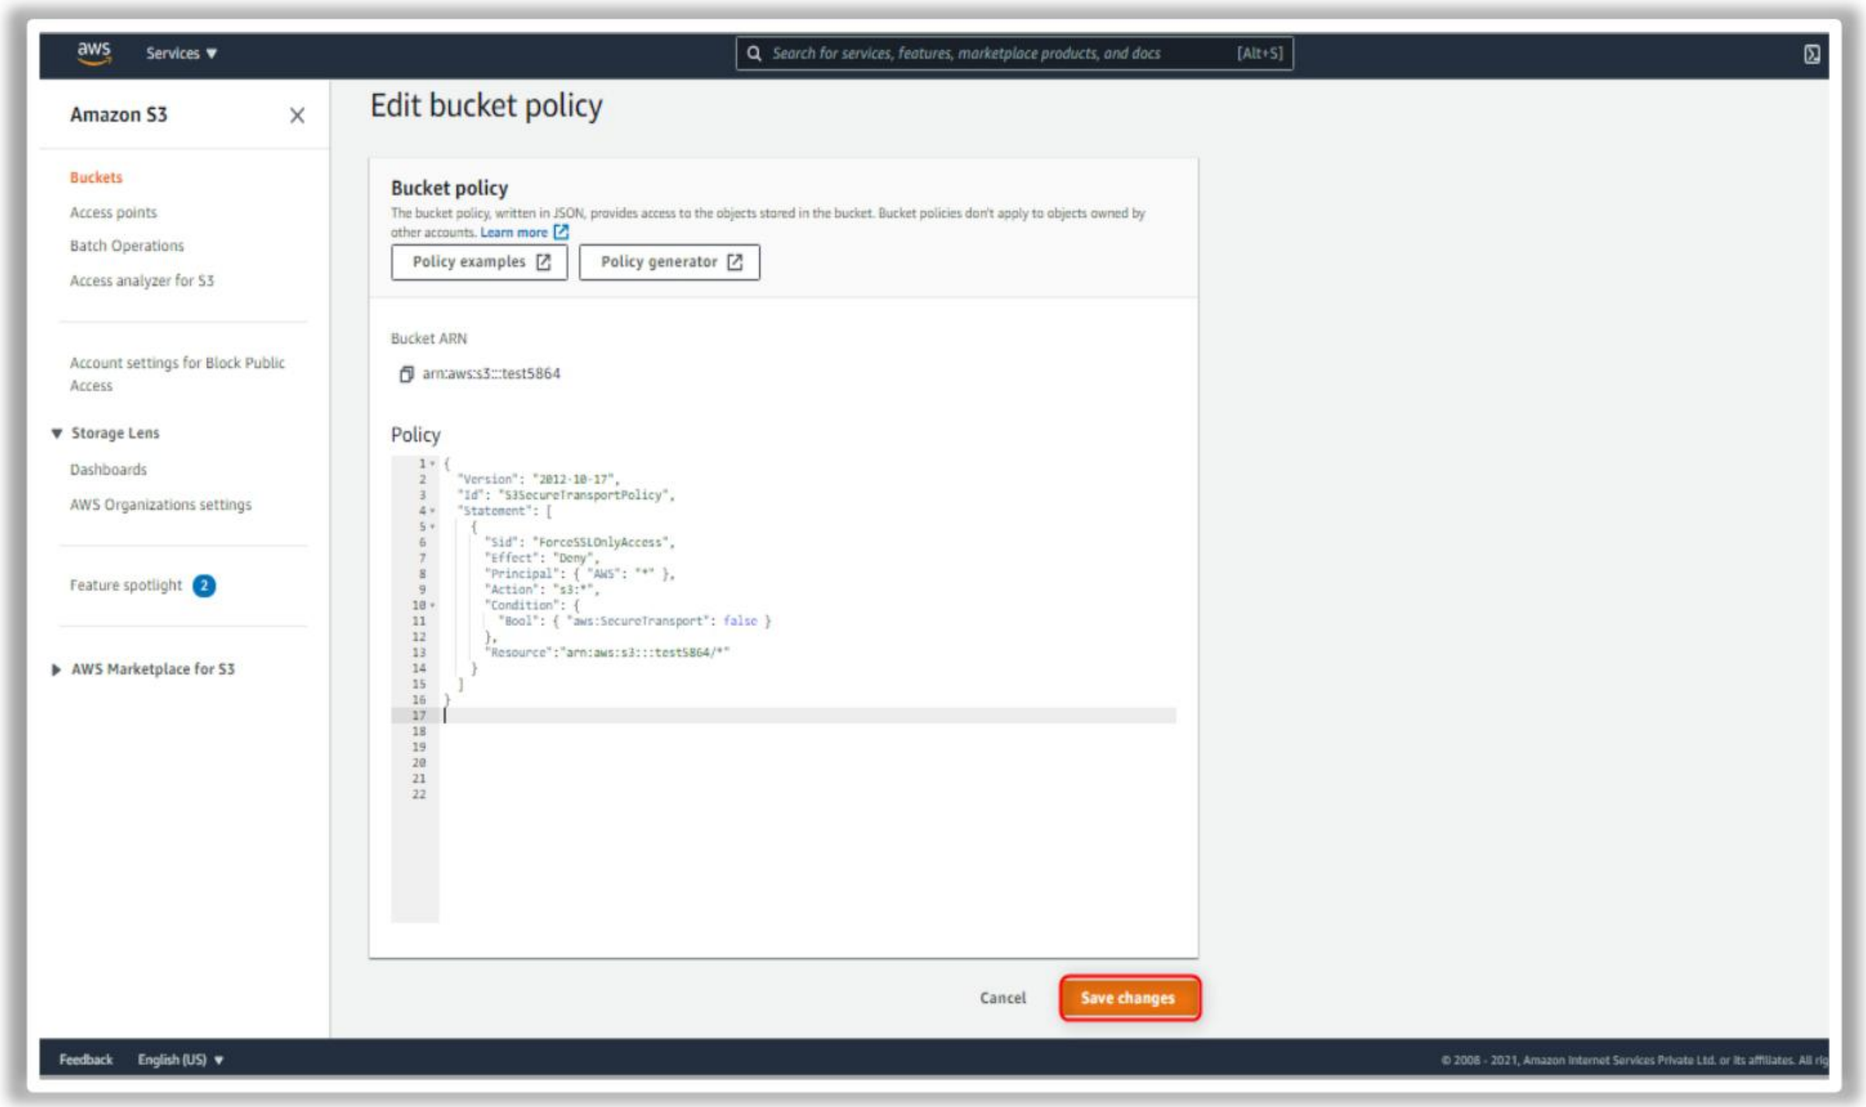


Figure 4.80: Click Save

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

- Non-SSL S3 access to all the objects in the selected bucket will be restricted.

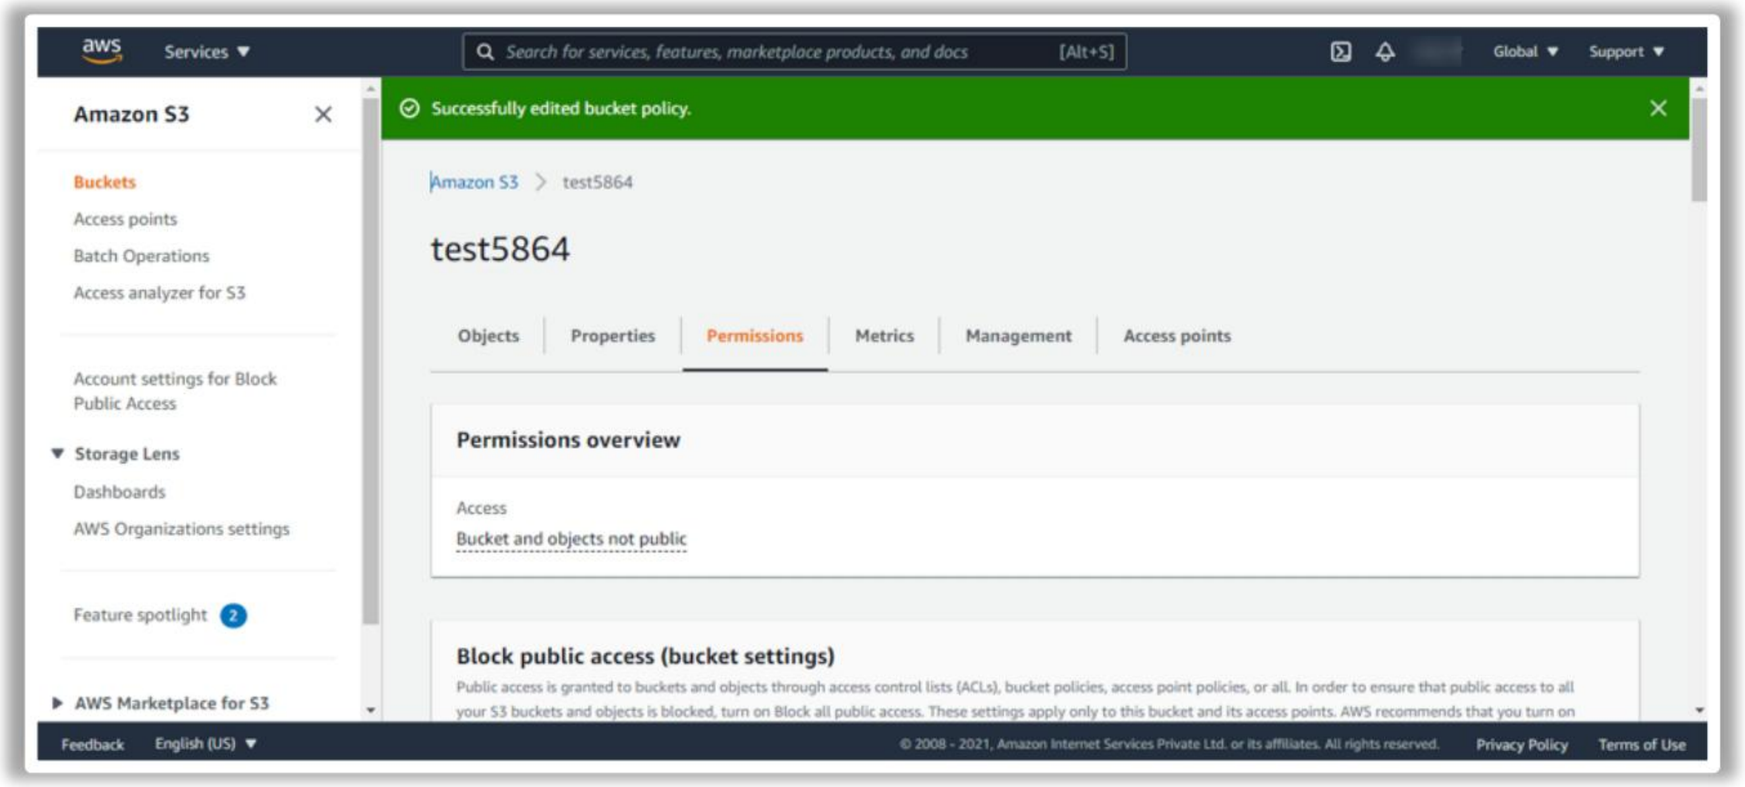


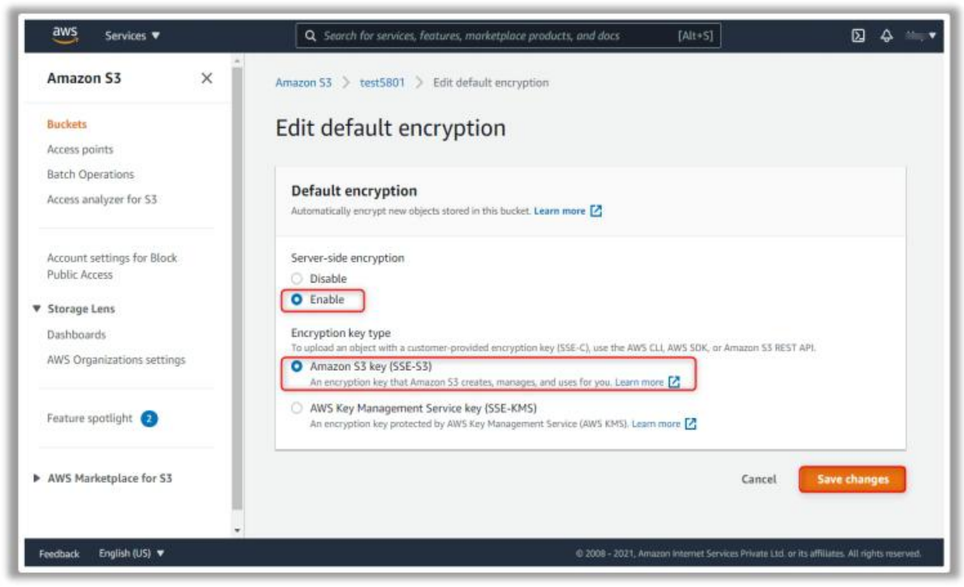
Figure 4.81: Restrict Bucket

hide01.ir

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

Ensure AWS S3 Buckets Enforce Server-Side Encryption (SSE)

- Safeguard the sensitive information at rest in the S3 buckets by implementing **server-side encryption**
- The cloud security professional can enable SSE utilizing **policy conditions** or default encryption



Implementing Service-Side Encryption on AWS S3

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

Ensure AWS S3 Buckets Enforce Server-Side Encryption (SSE)

Safeguard the sensitive information at rest in the S3 buckets by implementing SSE. The cloud security professional can enable SSE utilizing access policy or default encryption.

hide01.ir

Source: https://aws.amazon.com

Steps to enable SSE for S3 buckets are as follows:

- Sign-in to the **AWS Management Console**.
- Click on **Services** dropdown; then click on **S3** under **Storage.**

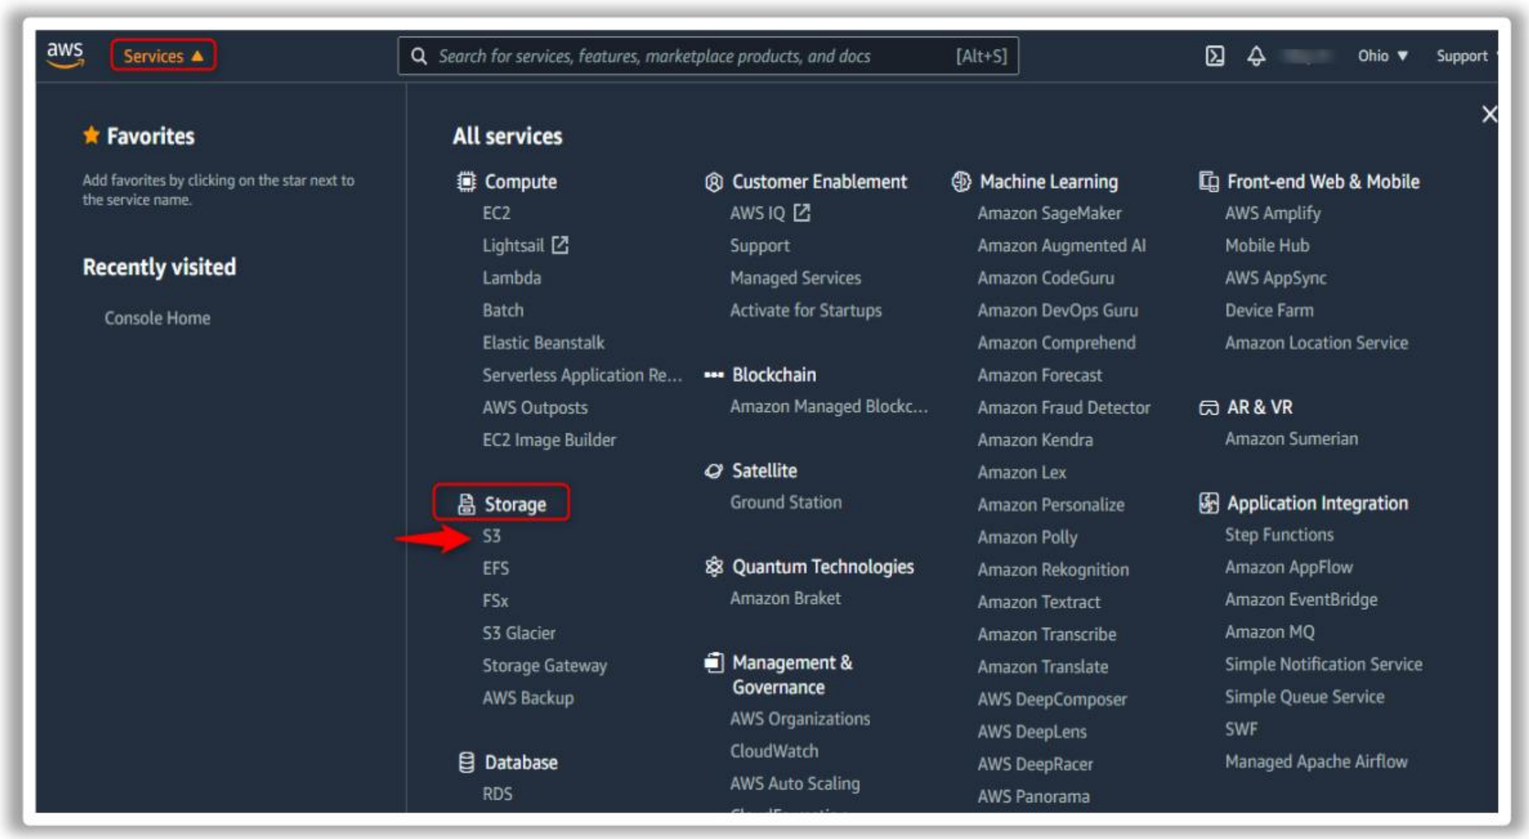


Figure 4.82: Click on Storage

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

- The list of available buckets will be displayed.

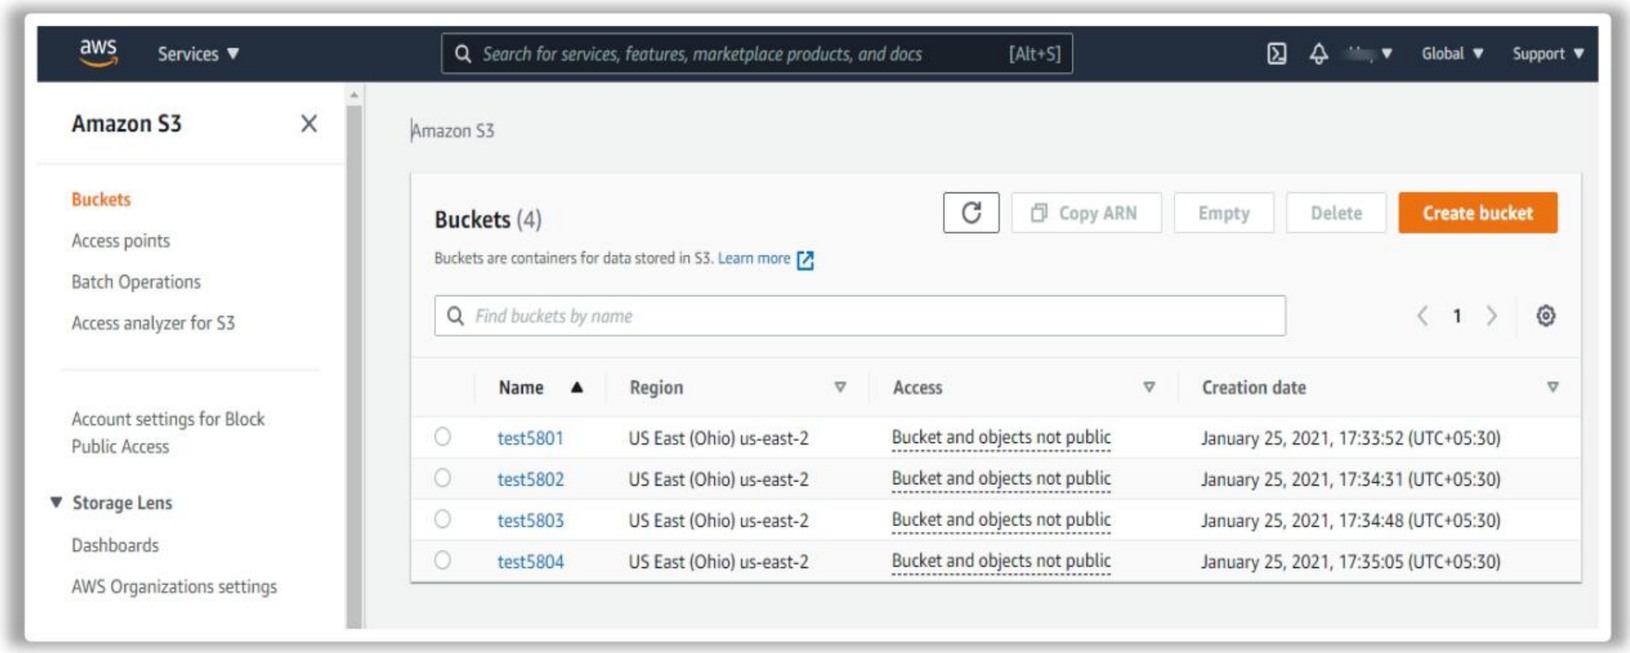


Figure 4.83: List of Buckets

- Select the bucket that you want to encrypt and click on it. Then, click on **Properties** tab.

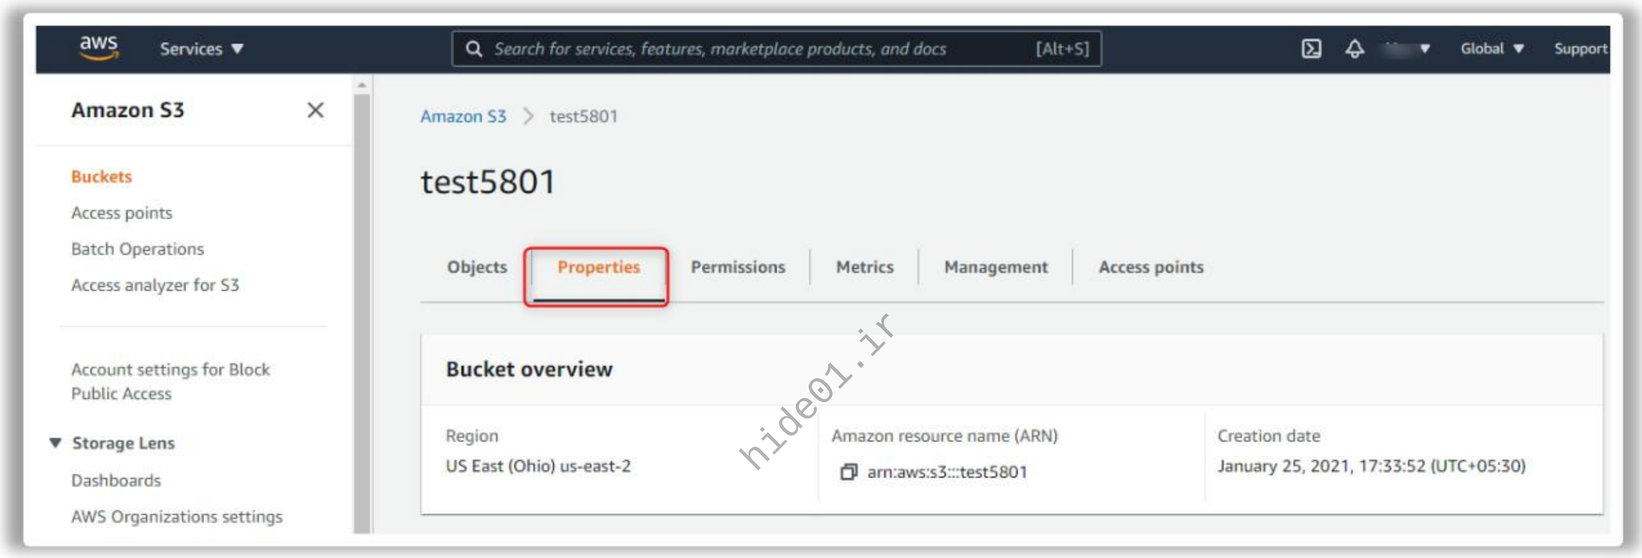


Figure 4.84: Click on Properties

- Scroll down and click the **Edit** button in **Default encryption** section.

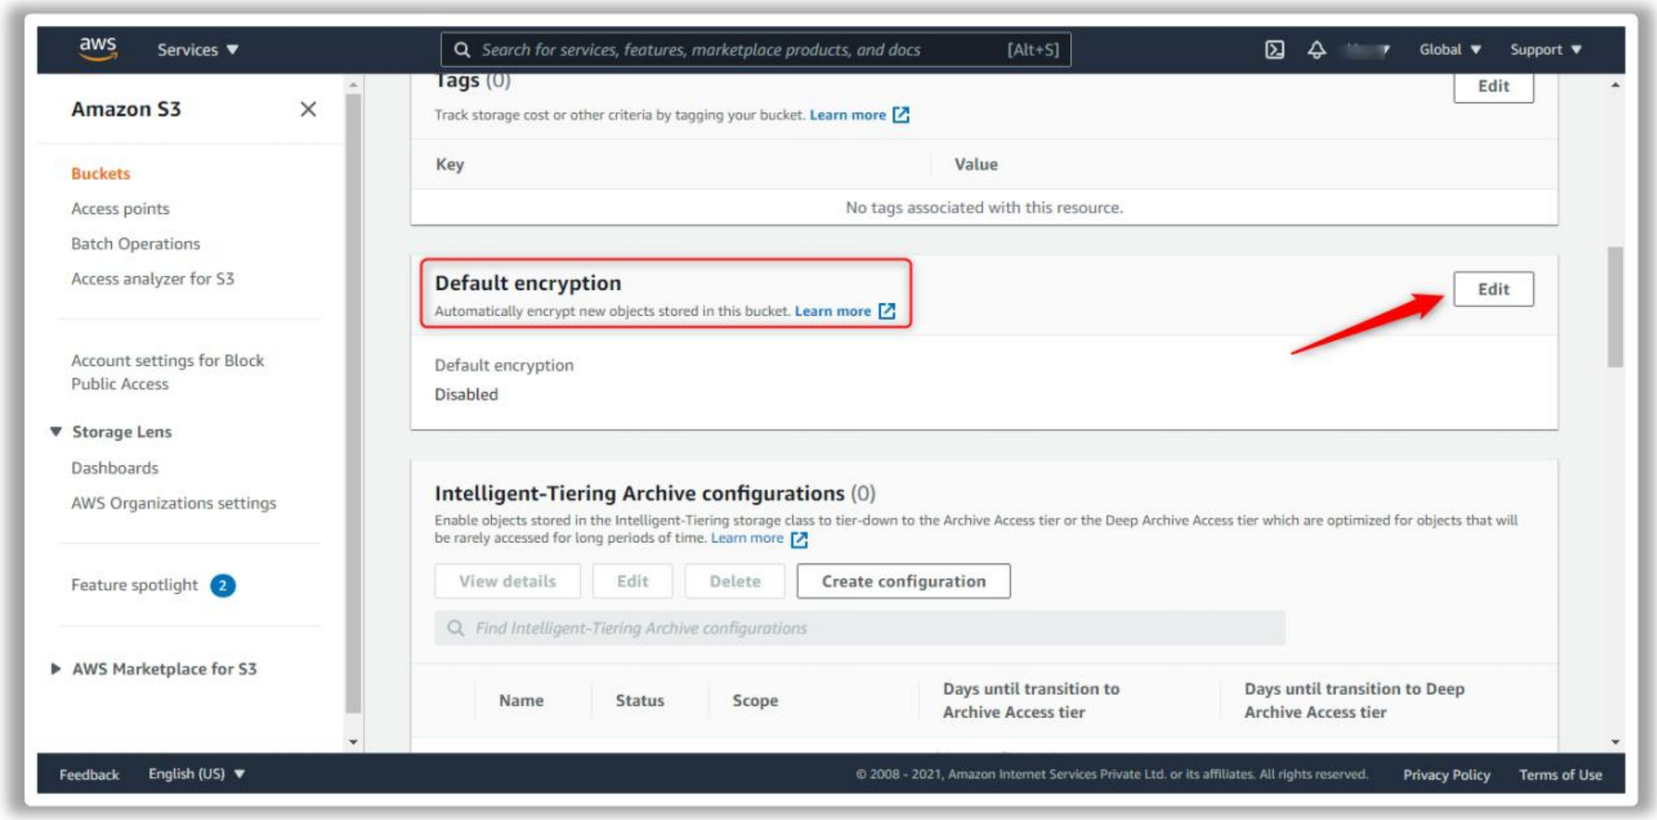


Figure 4.85: Select Edit Button

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

- Navigate to **Server-side encryption**; check **Enable** radio button.
- Then, under **Encryption key type**, check **Amazon S3 key (SSE-S3)**, and then click on **Save changes** button.

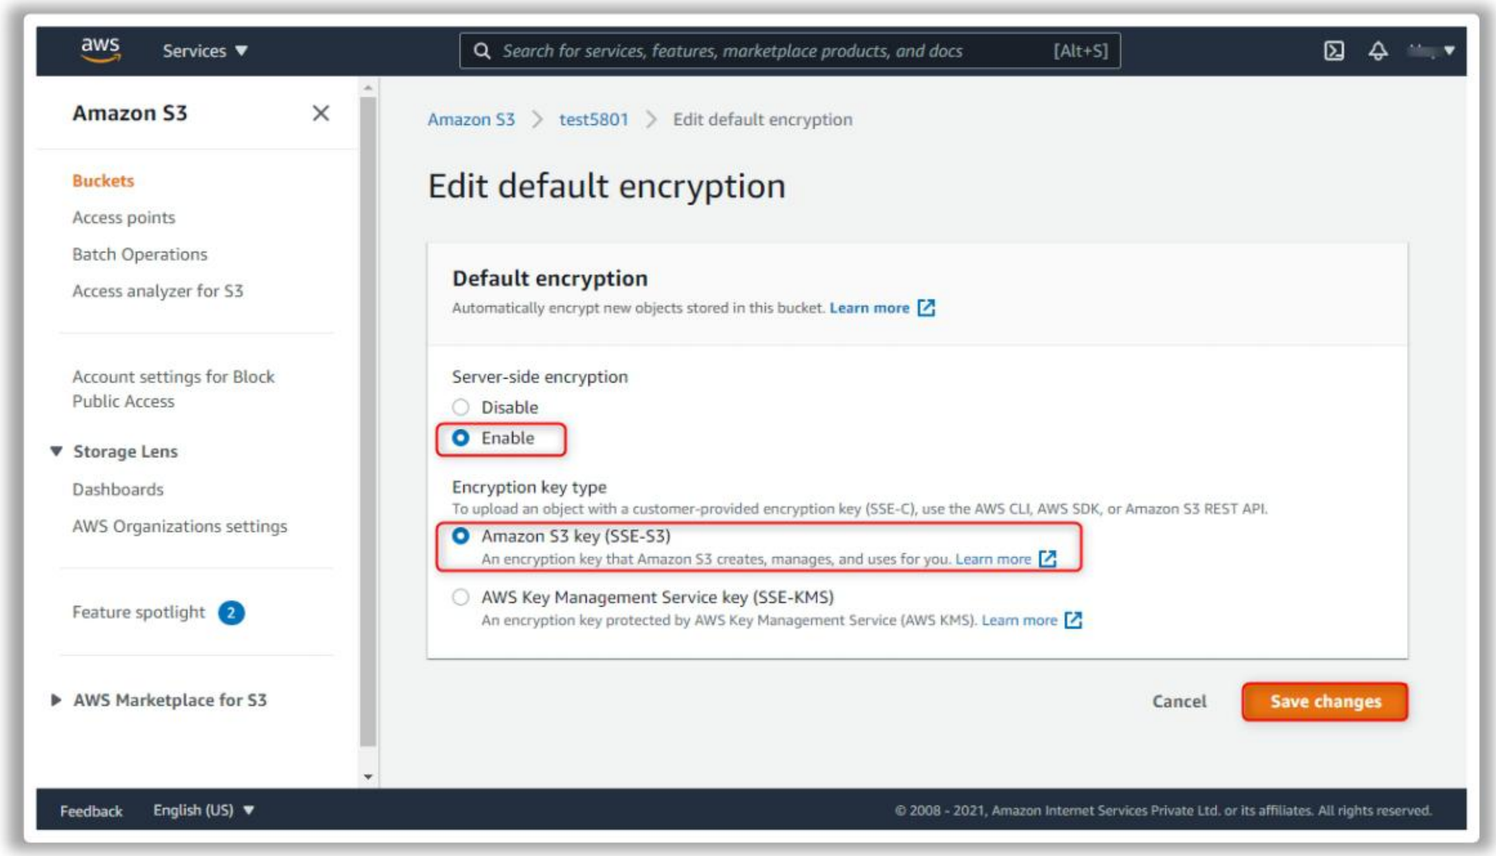


Figure 4.86: Save Changes

Check Authenticated Users Group Access to S3 Bucket

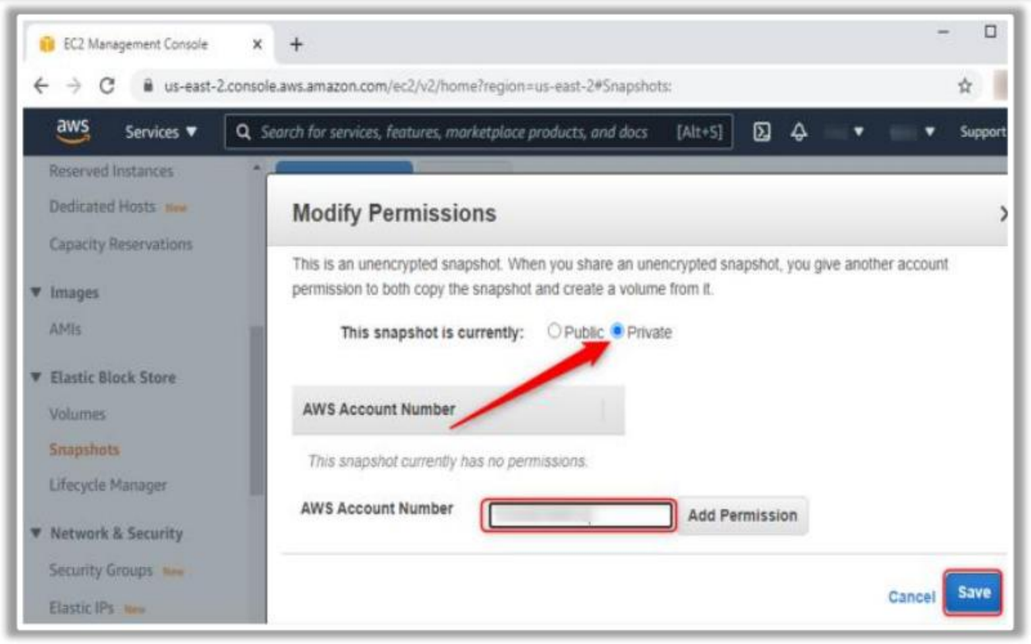
To prevent unauthorized access to S3 buckets, crosscheck authenticated users group (IAM user or AWS account user) to S3 bucket. If the S3 bucket allows FULL_CONTROL access to the authenticated users, then the user can view, modify, upload, or delete objects with ease. It can lead to data loss, data leak, or high billing charge on the S3 service. By default, AWS did not allow S3 bucket access to authenticated users; however, it is recommended to crosscheck authenticated users group and, if List and Read options are checked, then uncheck them to prevent unauthorized access to S3.

h1de01_ir

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

Grant Private Access to EBS Volume Snapshot

To safeguard sensitive information stored on AWS Elastic Block Store, the cloud security professional should **restrict public access** and grant private access to EBS Volume Snapshot to specific AWS user account



Granting Private Access to EBS Volume Snapshot

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

Grant Private Access to EBS Volume Snapshot

To safeguard sensitive information stored on AWS Elastic Block Store, ensure that the EBS volume snapshots are not publicly shared with other AWS account users. Generally, the EBS snapshots contain mirrors of the application as well as data; thus, sharing the EBS volume snapshots publicly allows other AWS account users to copy the snapshot and create a volume. If needed, then you can share your EBS snapshots with a particular AWS account user. It is recommended to restrict public access to the EBS snapshots; instead, provide private access to the EBS volume snapshot by entering the AWS account number of the user with whom you want to share the snapshot.

Source: https://aws.amazon.com

To grant private access to EBS volume snapshot to specific AWS account user, perform the following steps:

- Sign-in to AWS Management Console.

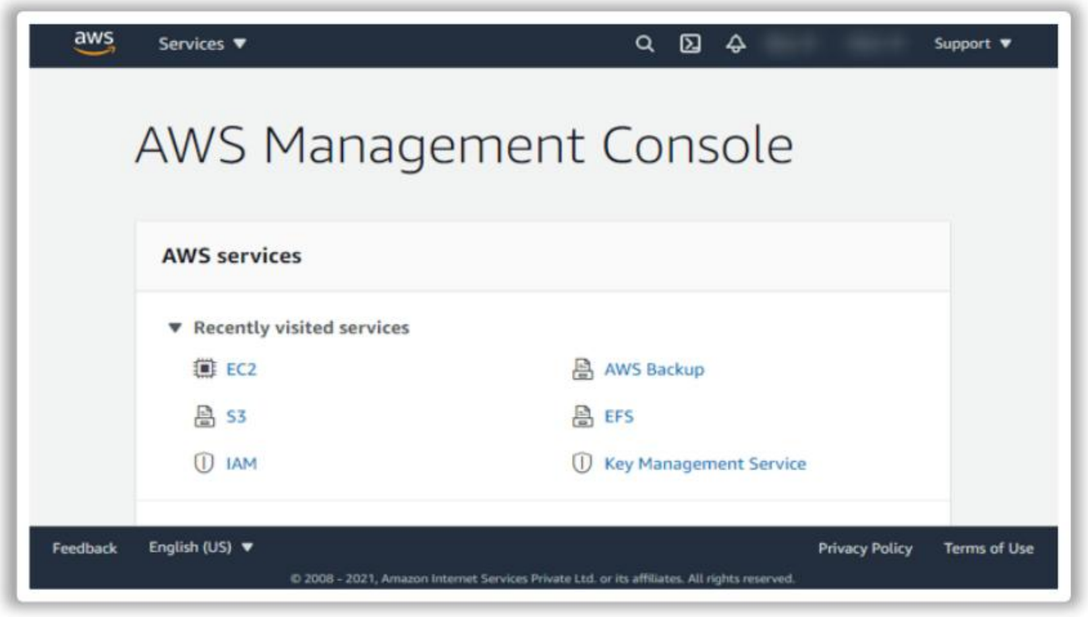


Figure 4.87: AWS Management Console

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/.

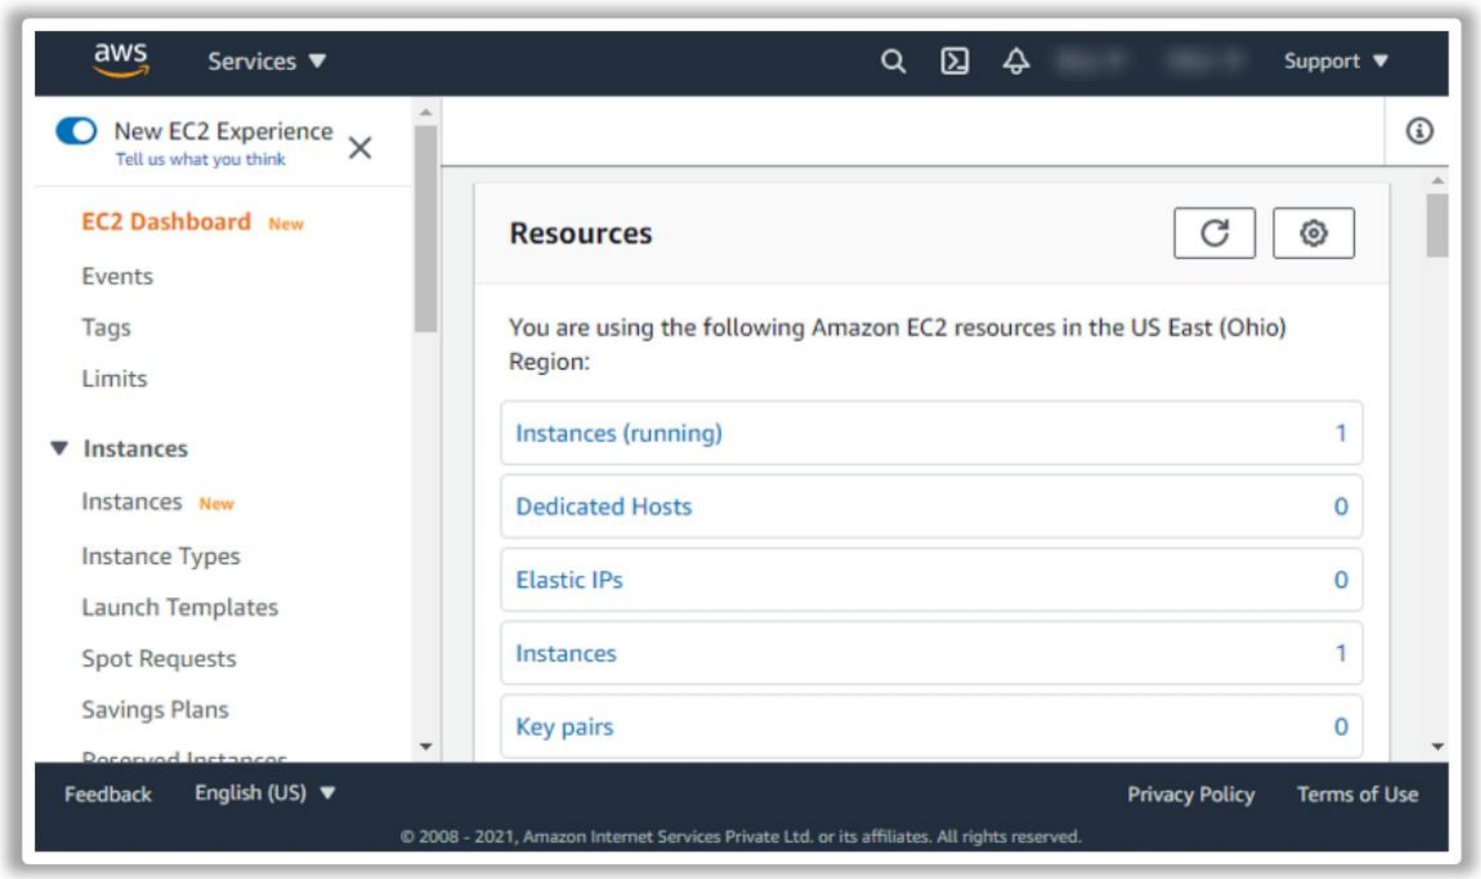


Figure 4.88: ECC2 Dashboard

- Navigate to **Elastic Block Store** and click on **Snapshots.**

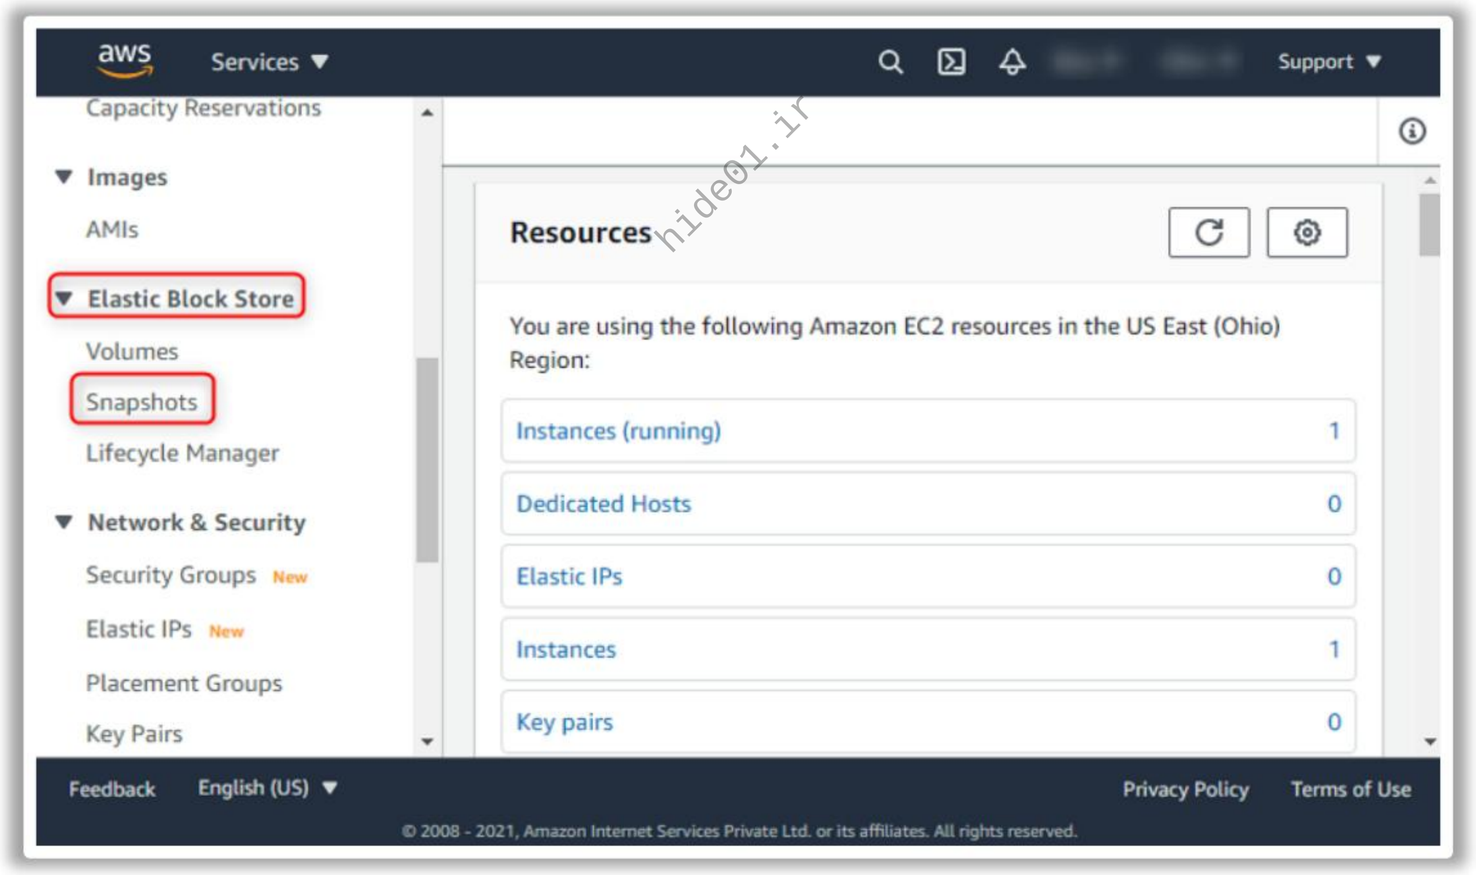


Figure 4.89: Click Snapshots

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

- Choose the screenshot that you want to make private; click on **Actions** dropdown, navigate and click on **Modify Permissions**.

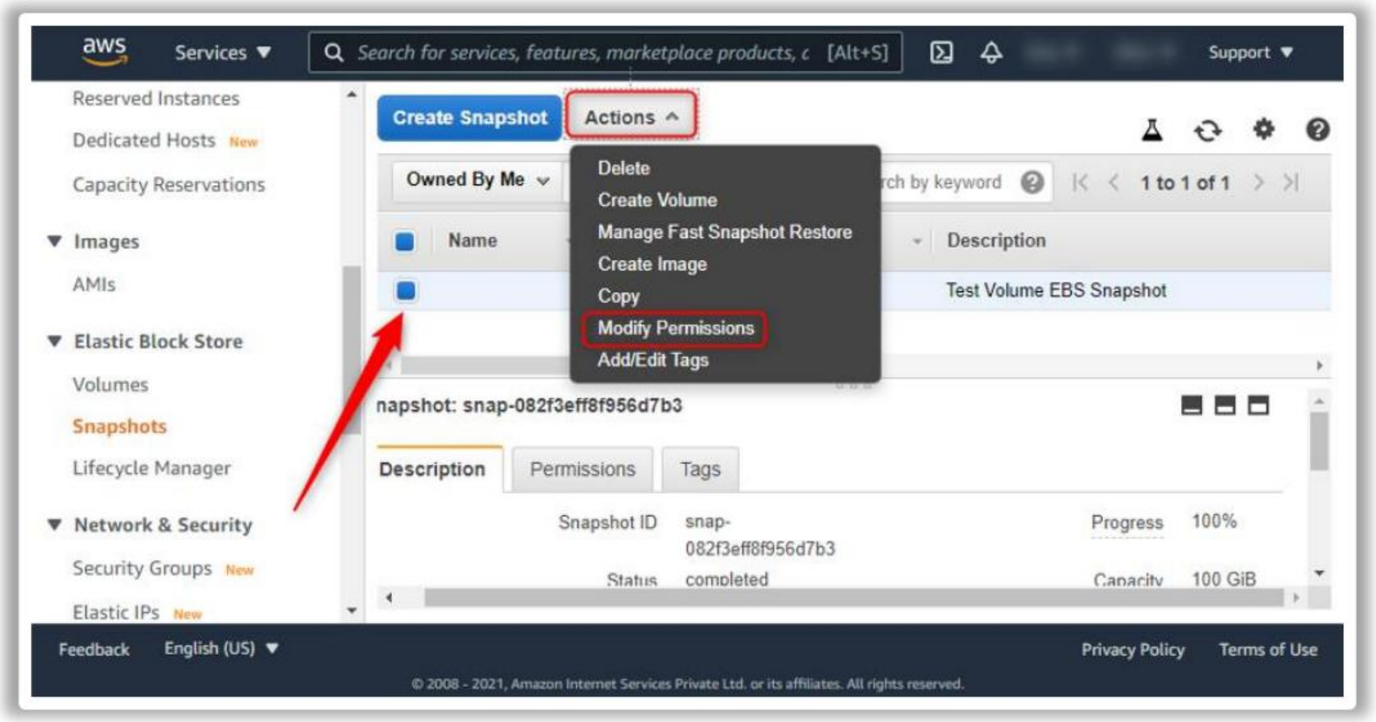


Figure 4.90: Select Modify Permissions

- If the snapshot is currently Public, then select the **Private** radio button and enter the **AWS Account Number** of the user with whom you want to share the snapshot; then click on **Save** button.

  Note: The snapshot shared with the specific AWS account user will be an unencrypted snapshot, which allows the user to copy the snapshot and to create a volume from it.

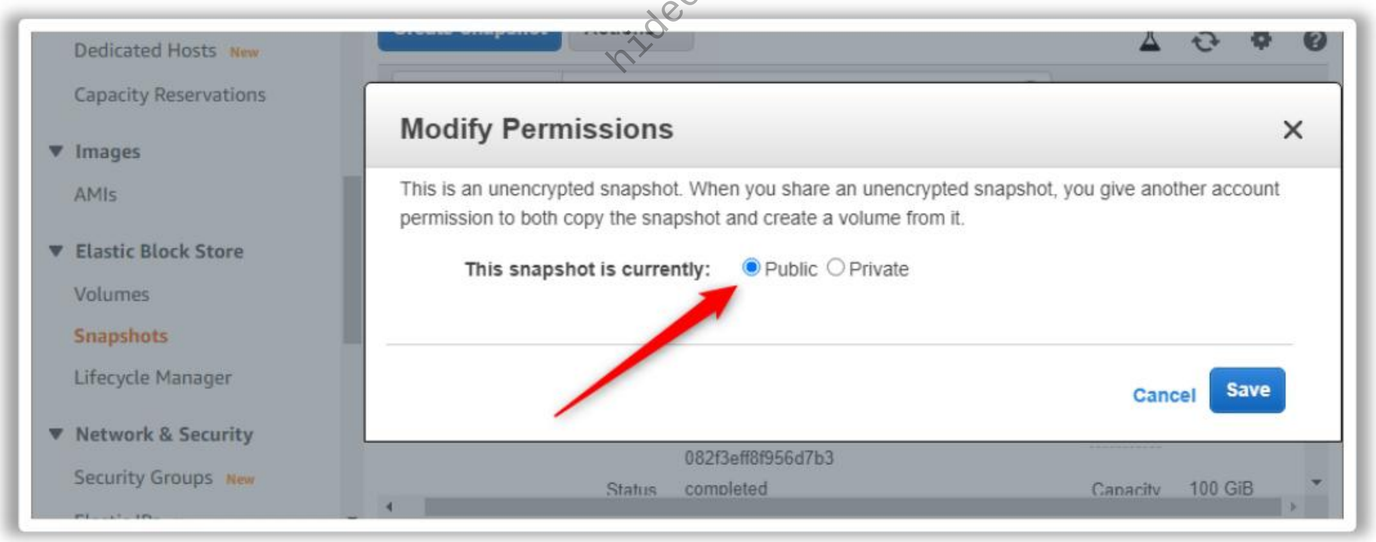


Figure 4.91: Click Private

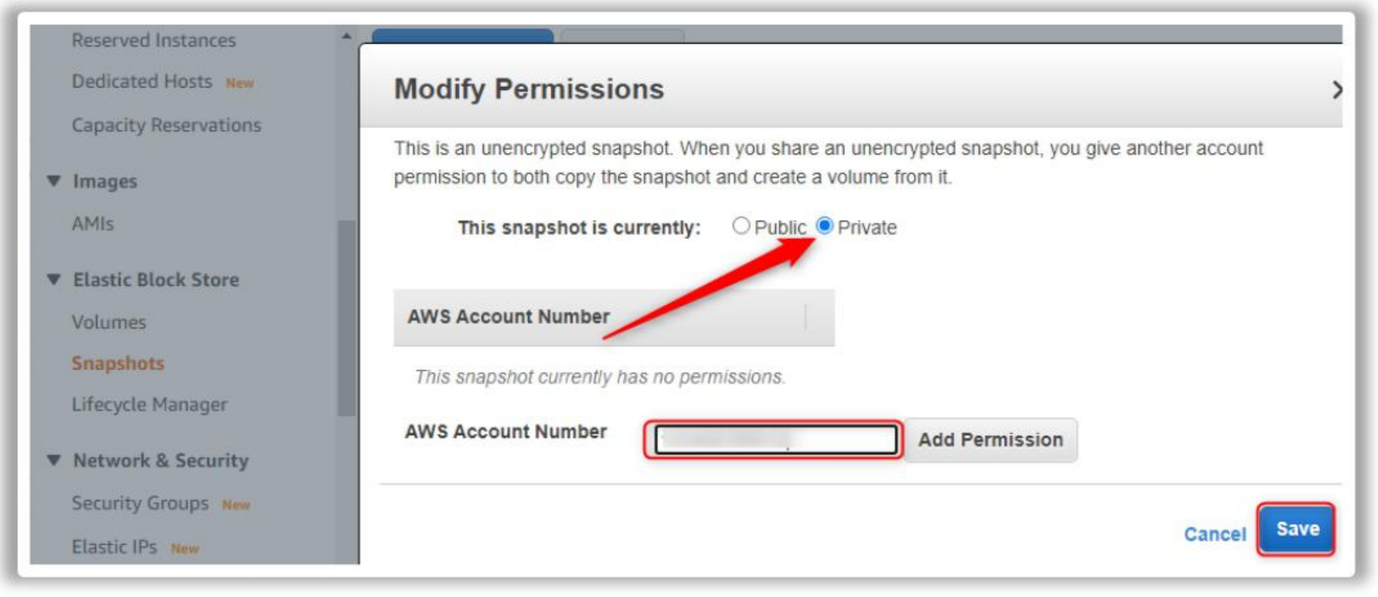
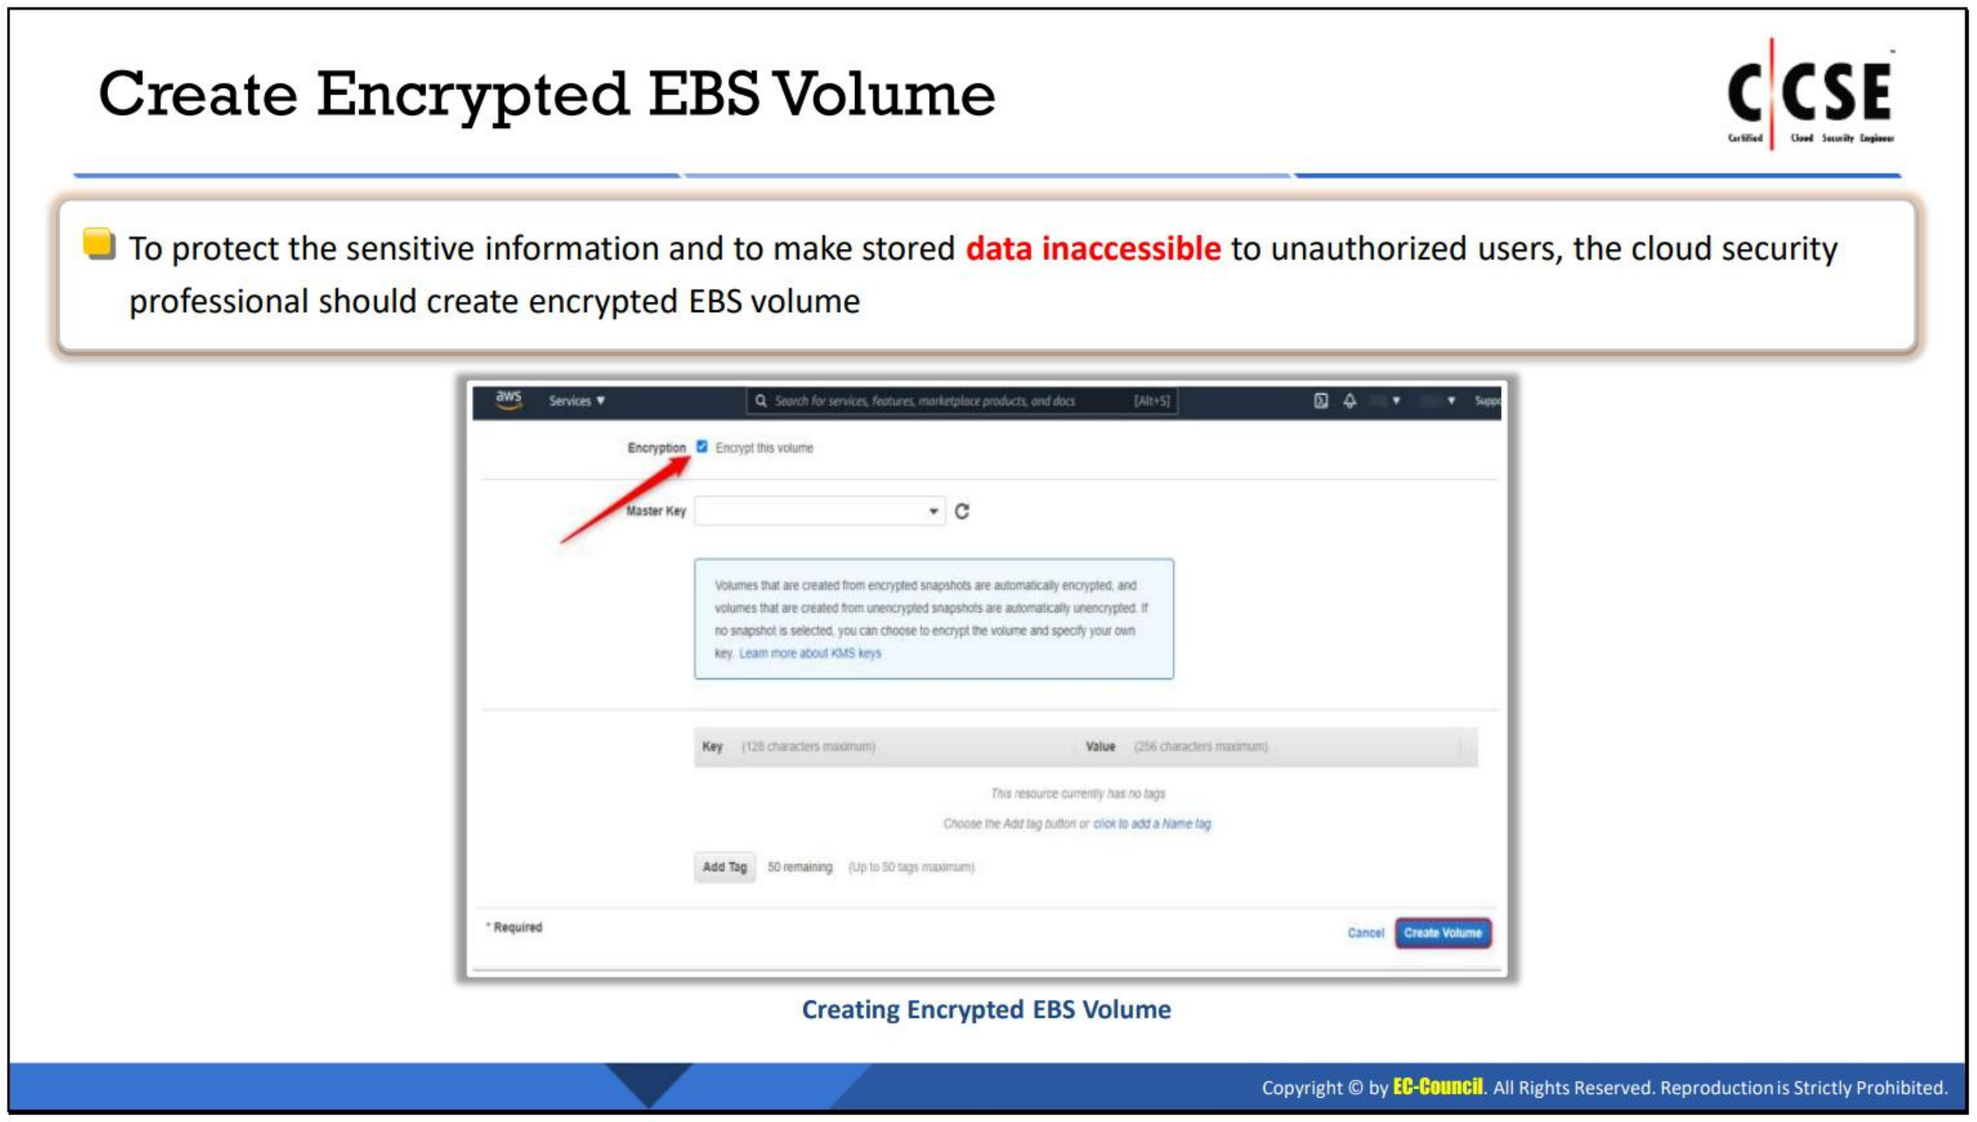


Figure 4.92: Save Changes

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



Creating Encrypted EBS Volume

Create Encrypted EBS Volume

To protect sensitive information and to make stored data inaccessible to unauthorized users, the cloud security professional should create encrypted EBS volume. Encrypted EBS volume will protect corporate secrets, classified information, confidential information, etc. It also provides protection against identity theft.

Source: https://aws.amazon.com

To create encrypted EBS volume, perform the following steps:

- Sign-in to AWS Management Console.

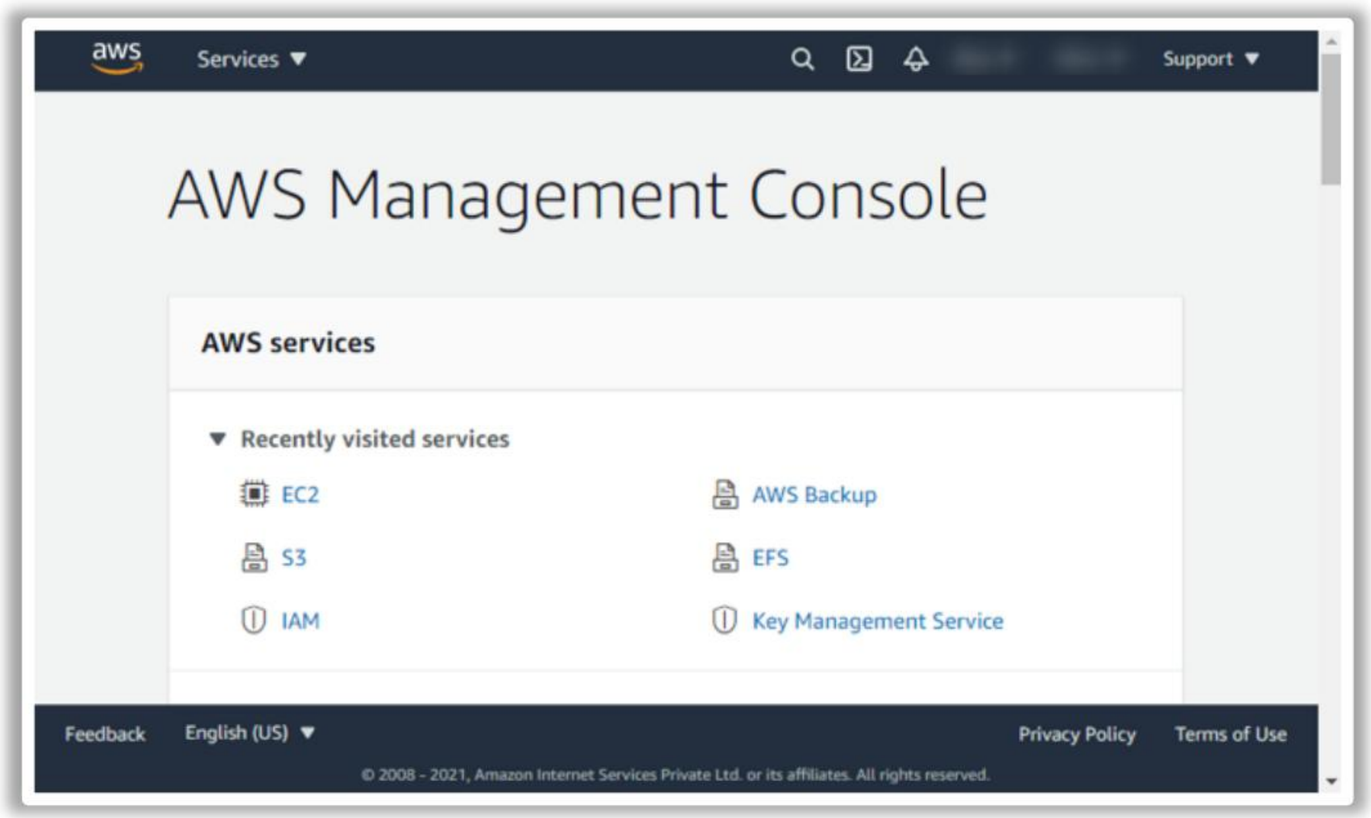


Figure 4.93: AWS Management Console

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/.

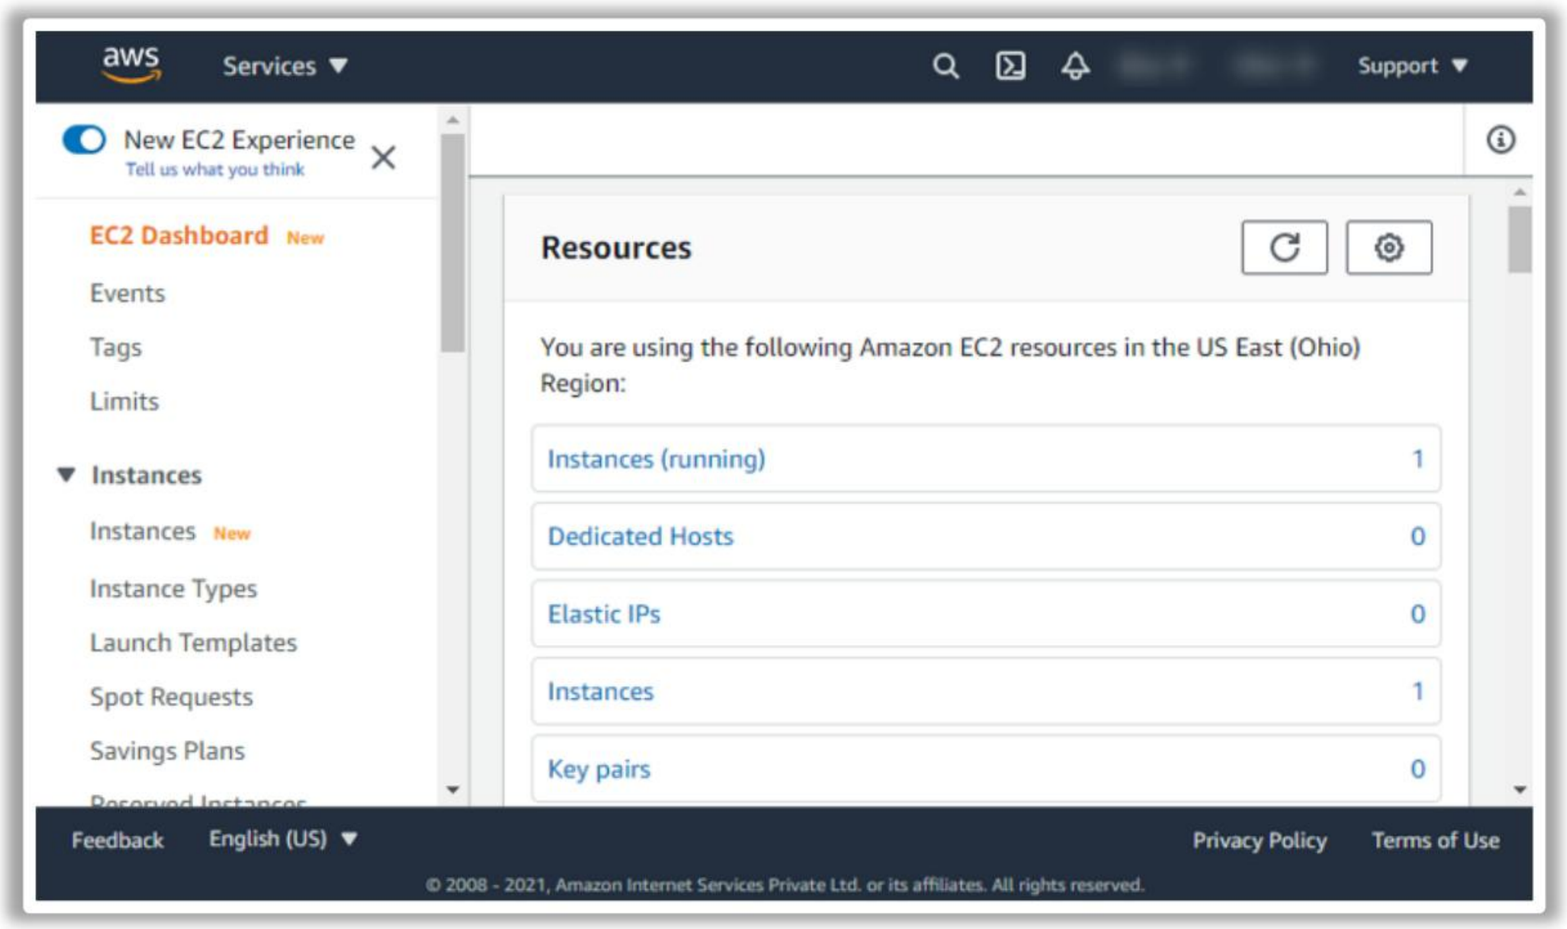


Figure 4.94: Click EC2 Dashboard

- Navigate to **Volume** under **Elastic Block Store** and click on it.

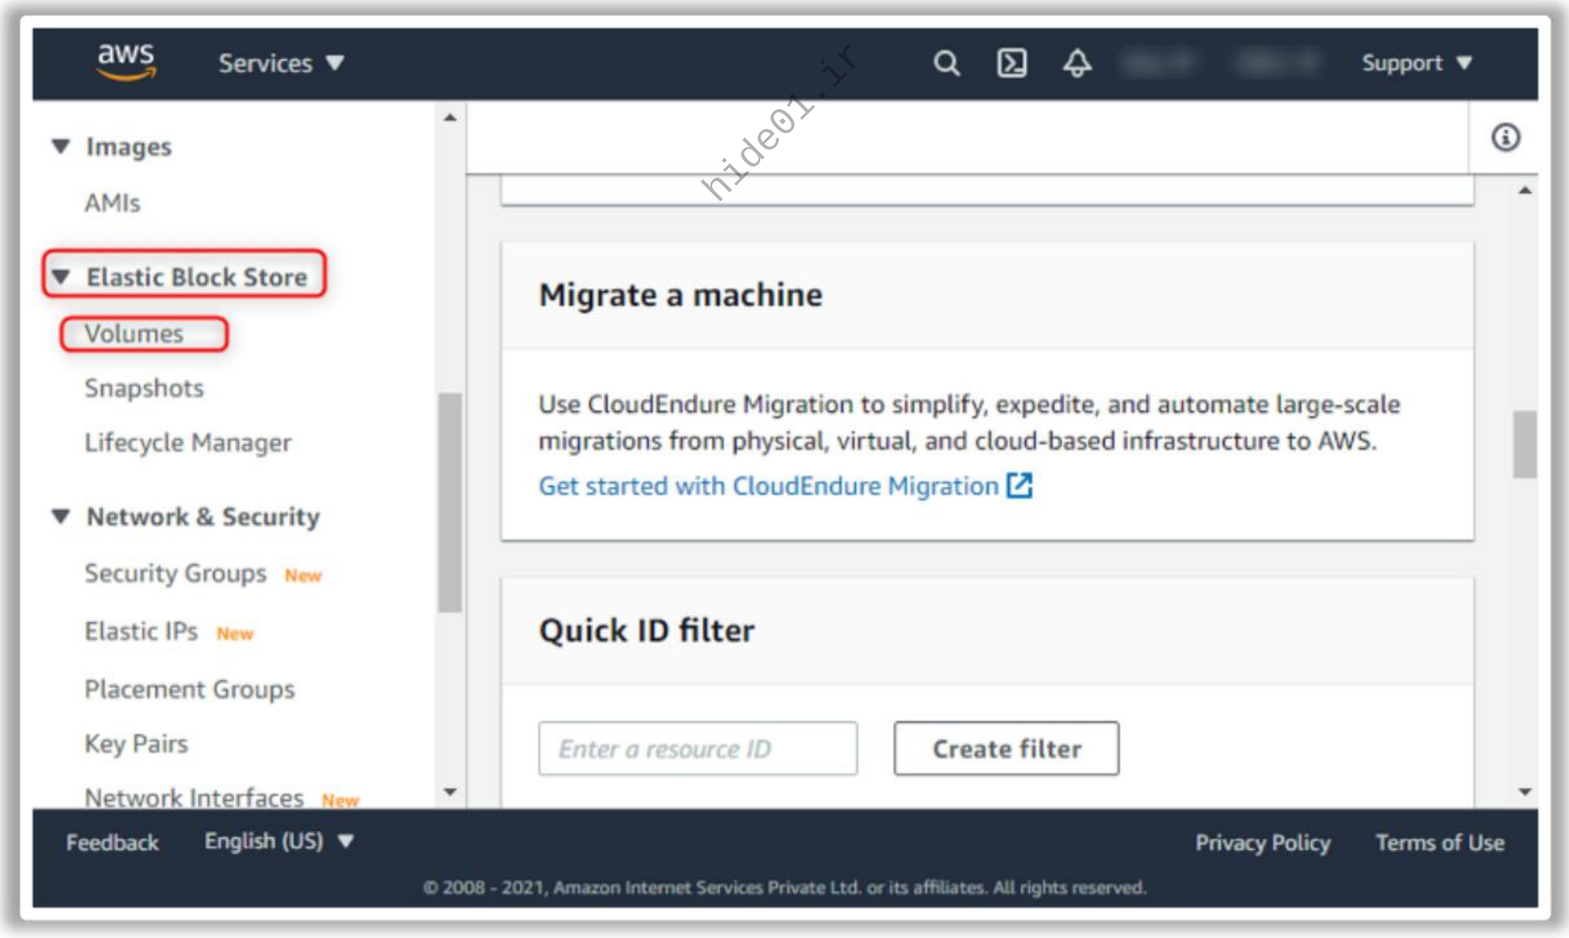


Figure 4.95: Select Volumes

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

- Click on **Create Volume** button.

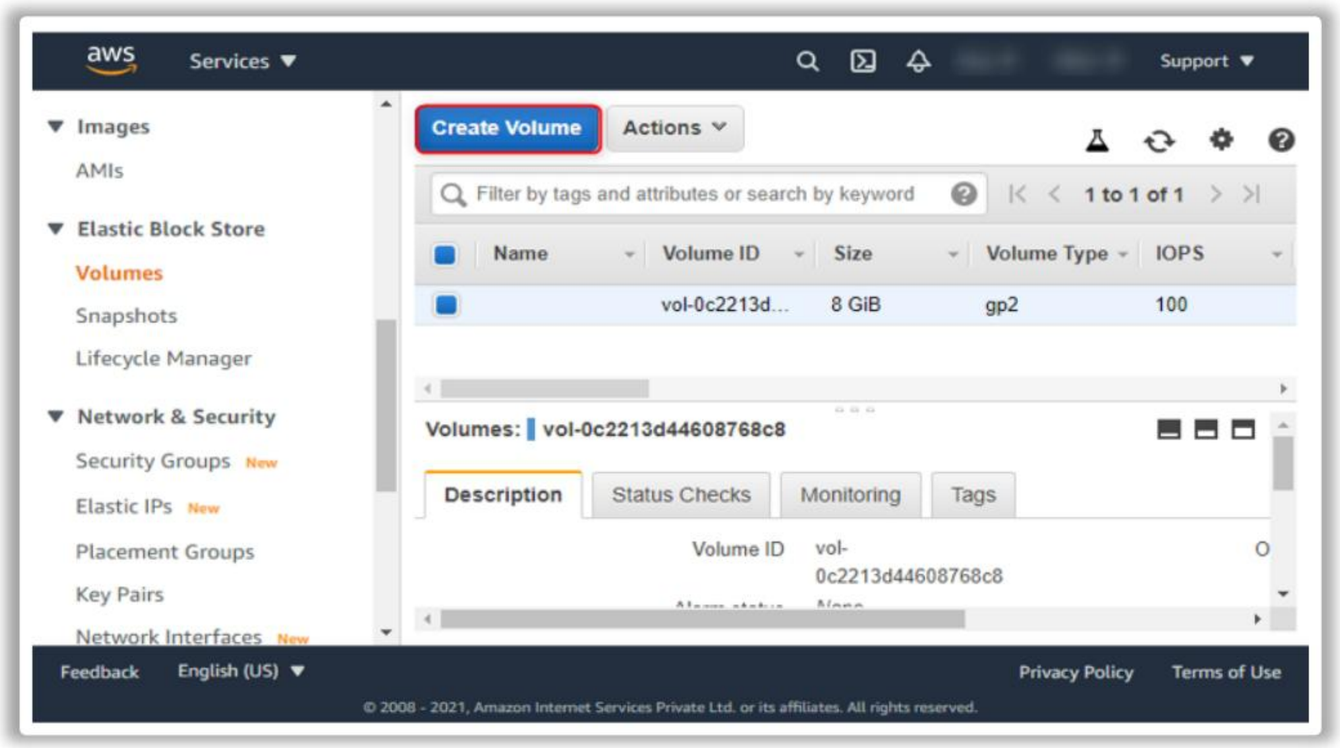


Figure 4.96: Select Create Volume

- Navigate to **Encryption**, check the box of **Encrypt this volume**, and then click on **Create Volume** button.

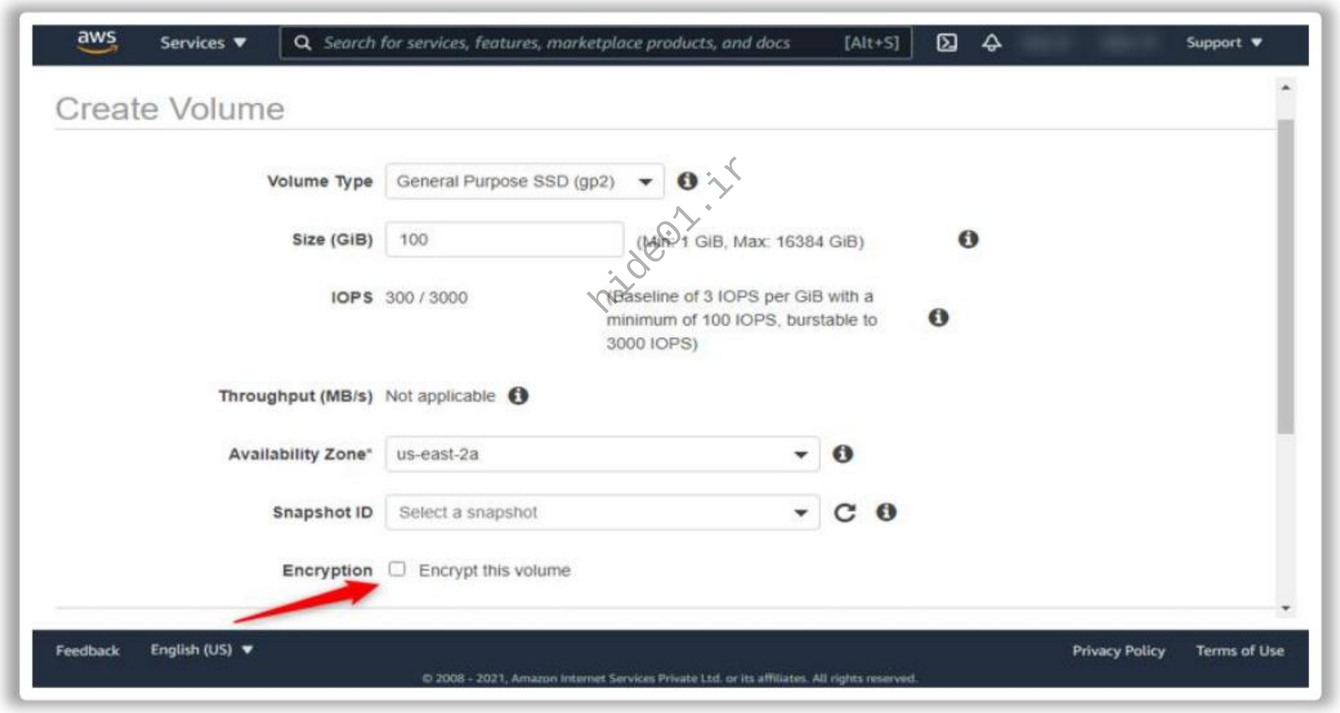


Figure 4.97: Select Encrypt this volume

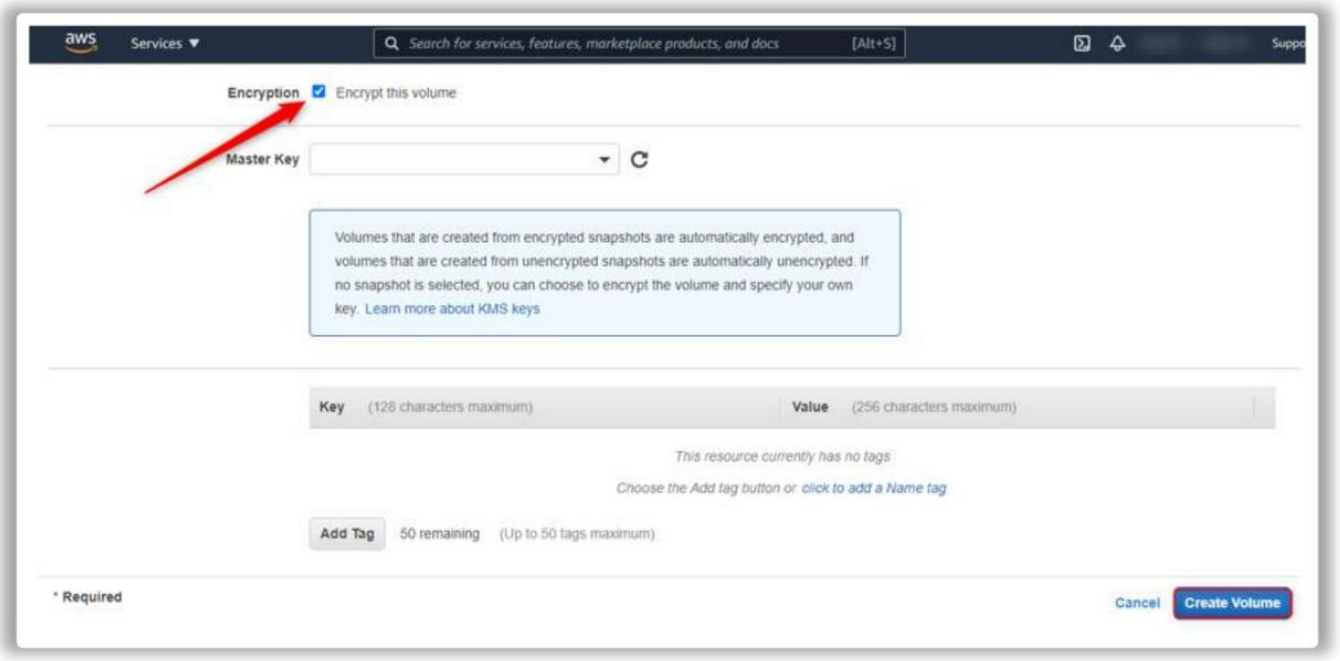


Figure 4.98: Create Volume

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

- **Volume created successfully** window will open; click on **Close** button.

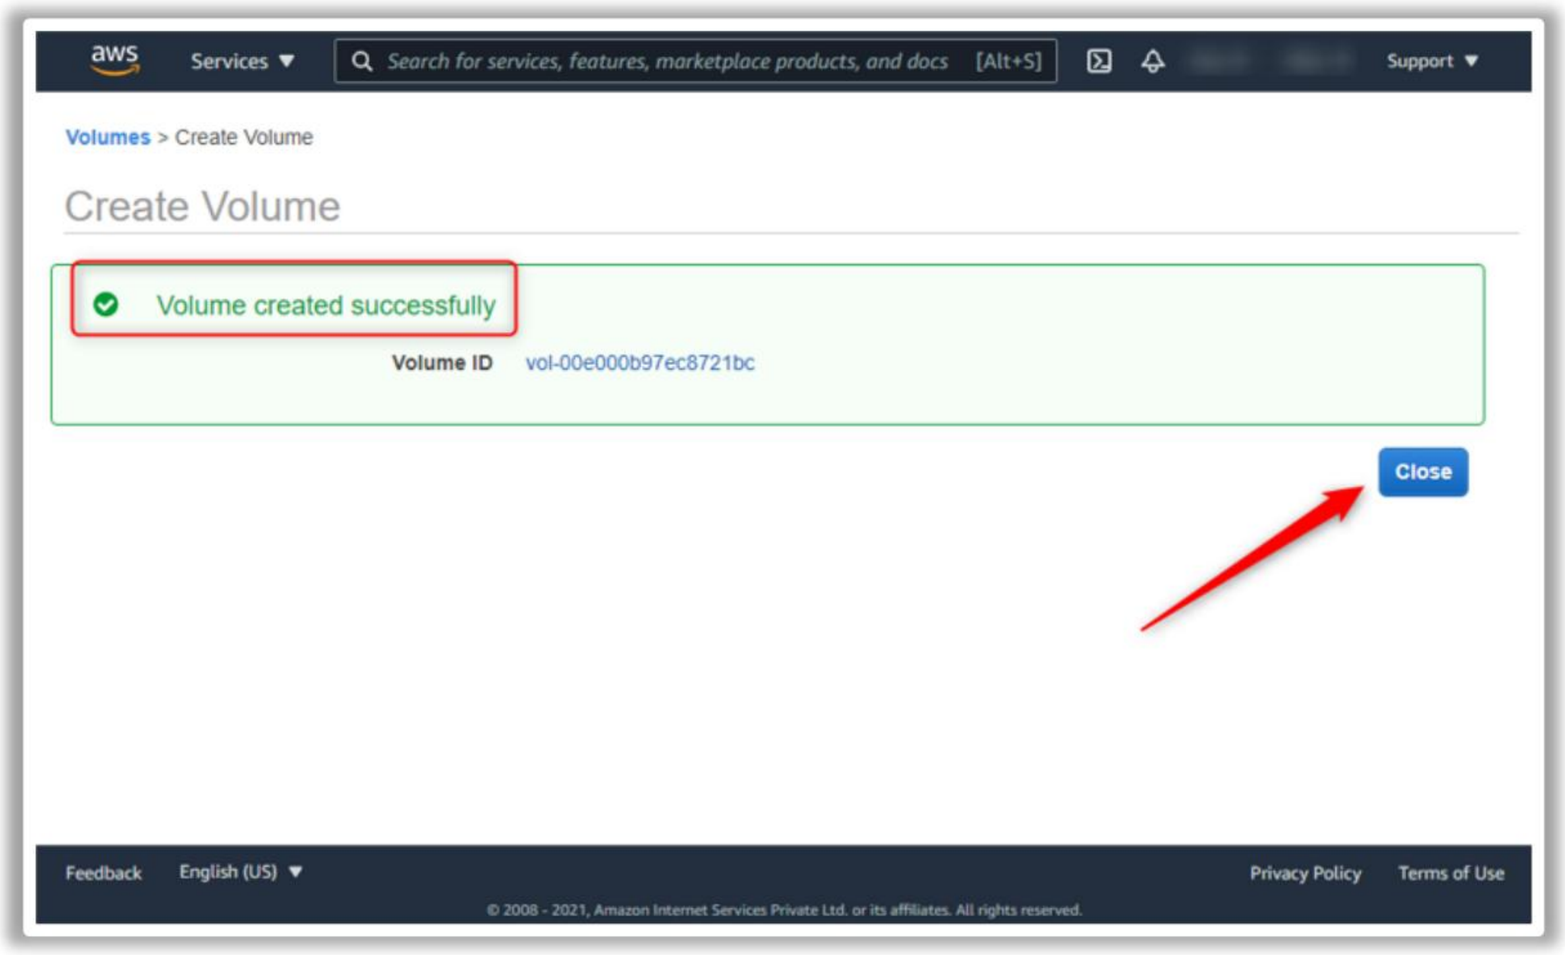


Figure 4.99: Select Close

- Crosscheck the **Volume ID** and **Encryption** status on the main page.

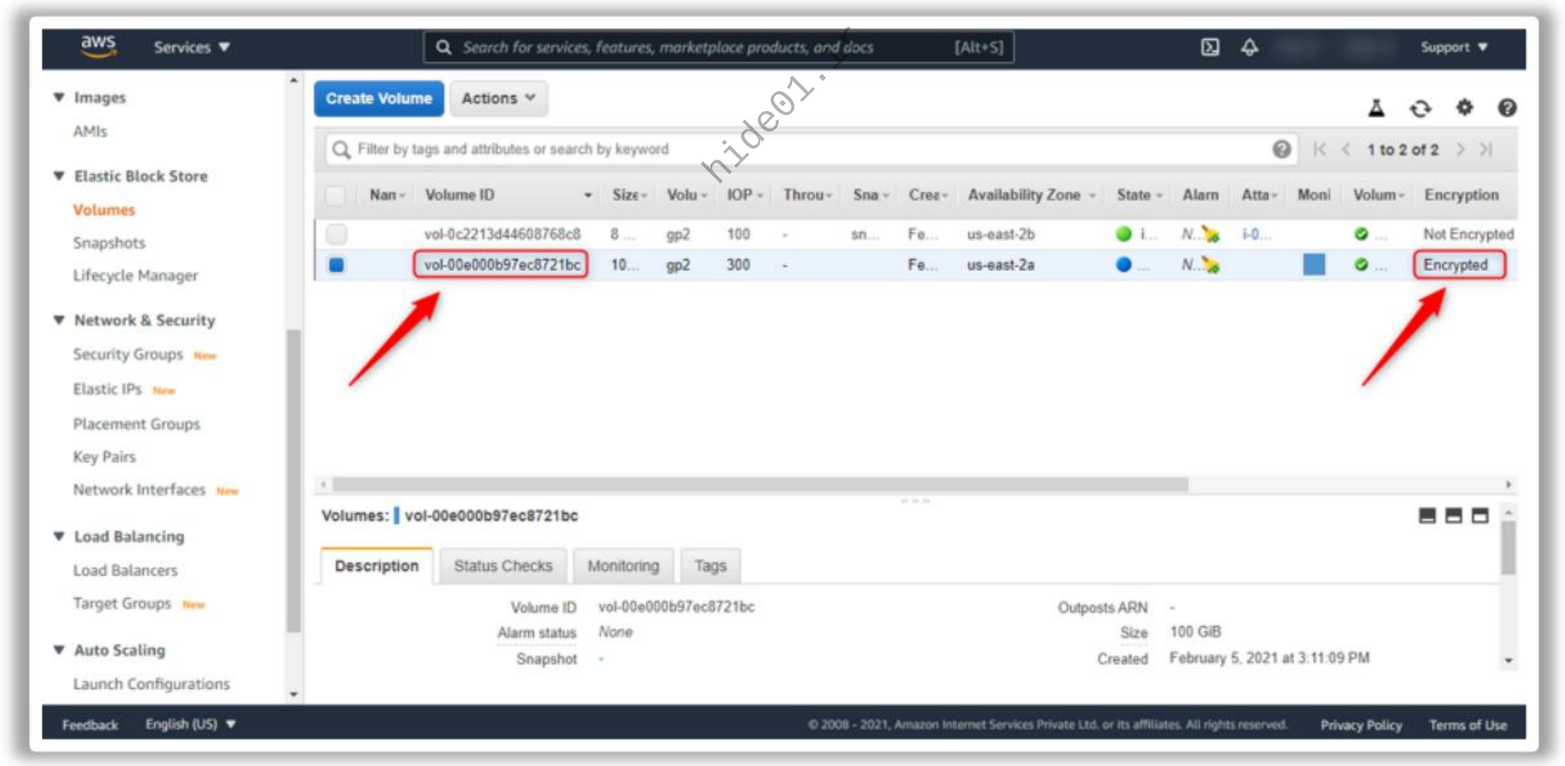
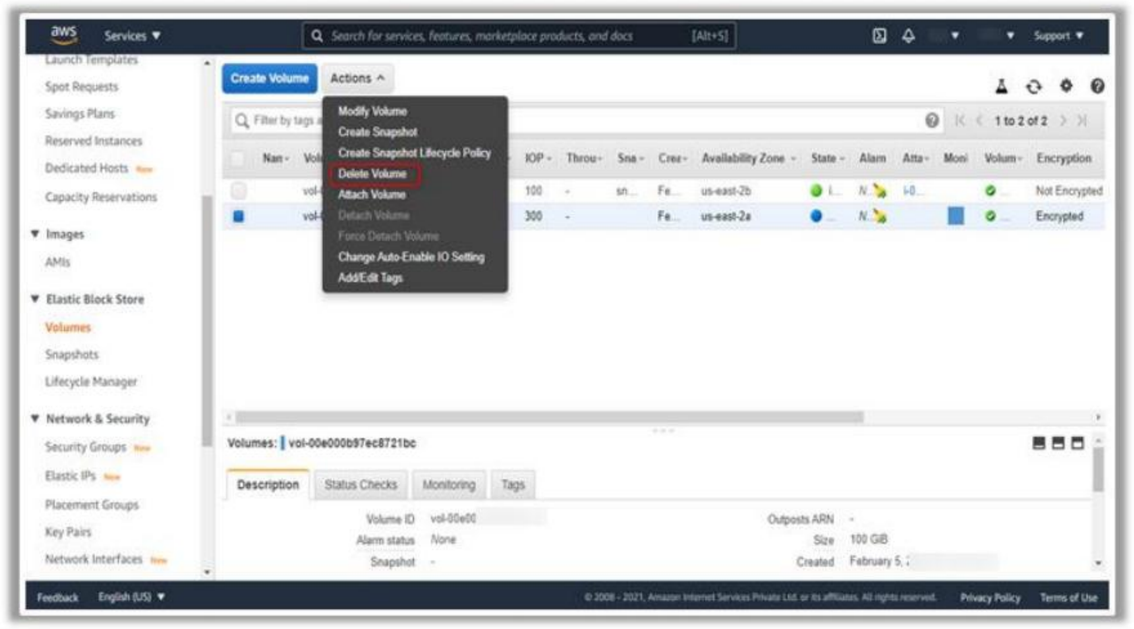


Figure 4.100: Volume Console

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

Delete Unused EBS Volume

- To minimize the risk of confidential or **sensitive data breach** and to reduce the monthly AWS account bill, identify the **unused EBS Volume** present in your AWS account and delete them



Deleting Unused EBS Volume

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

Delete Unused EBS Volume

To minimize the risk of confidential or sensitive data breach and to reduce the monthly AWS account bill, identity the unused EBS volume present in your AWS account and delete them.

The created EBS volumes will be charged monthly, even though the AWS customer is not using them. It is recommended to delete unused/orphaned EBS volumes, which will not only help in reducing the monthly bill but will also minimize the risk of data leaving the premises.

Source: https://aws.amazon.com

To delete unused EBS volume, perform the following steps:

- Sign-in to AWS Management Console.

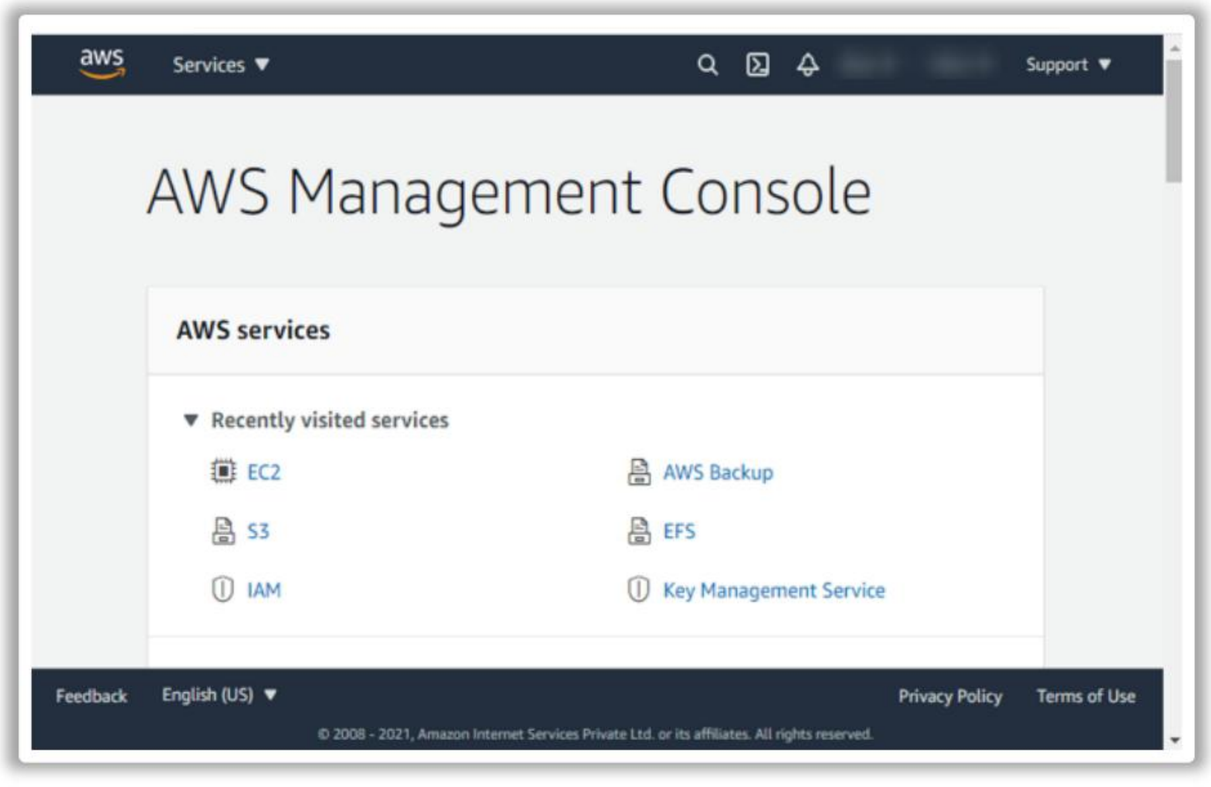


Figure 4.101: AWS Management Console

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/.

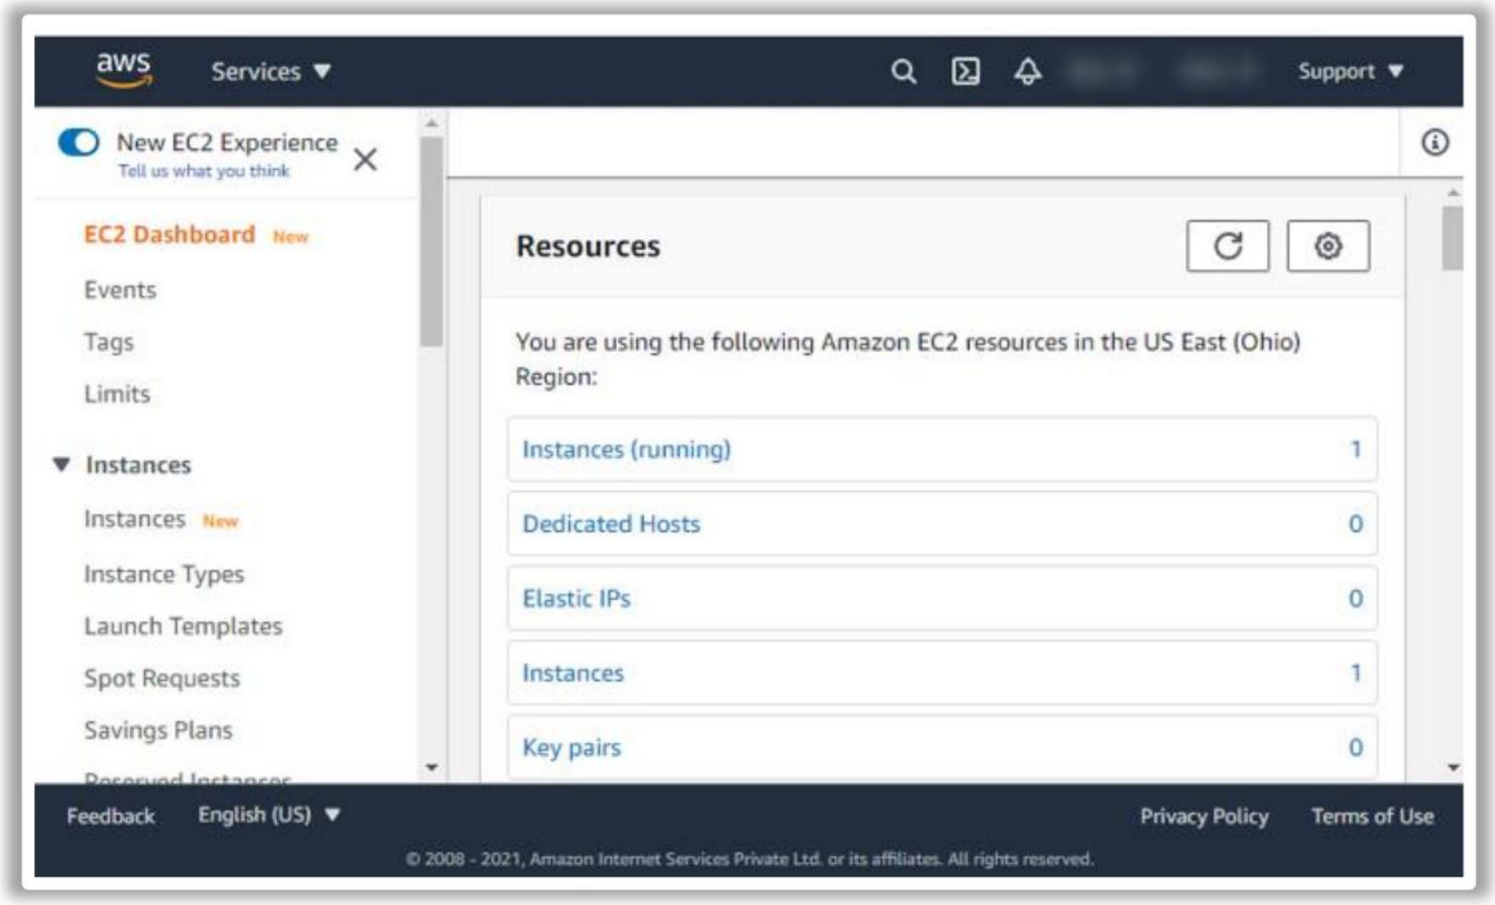


Figure 4.102: Click EC2 Dashboard

- Navigate to **Volume** under **Elastic Block Store** and click on it.

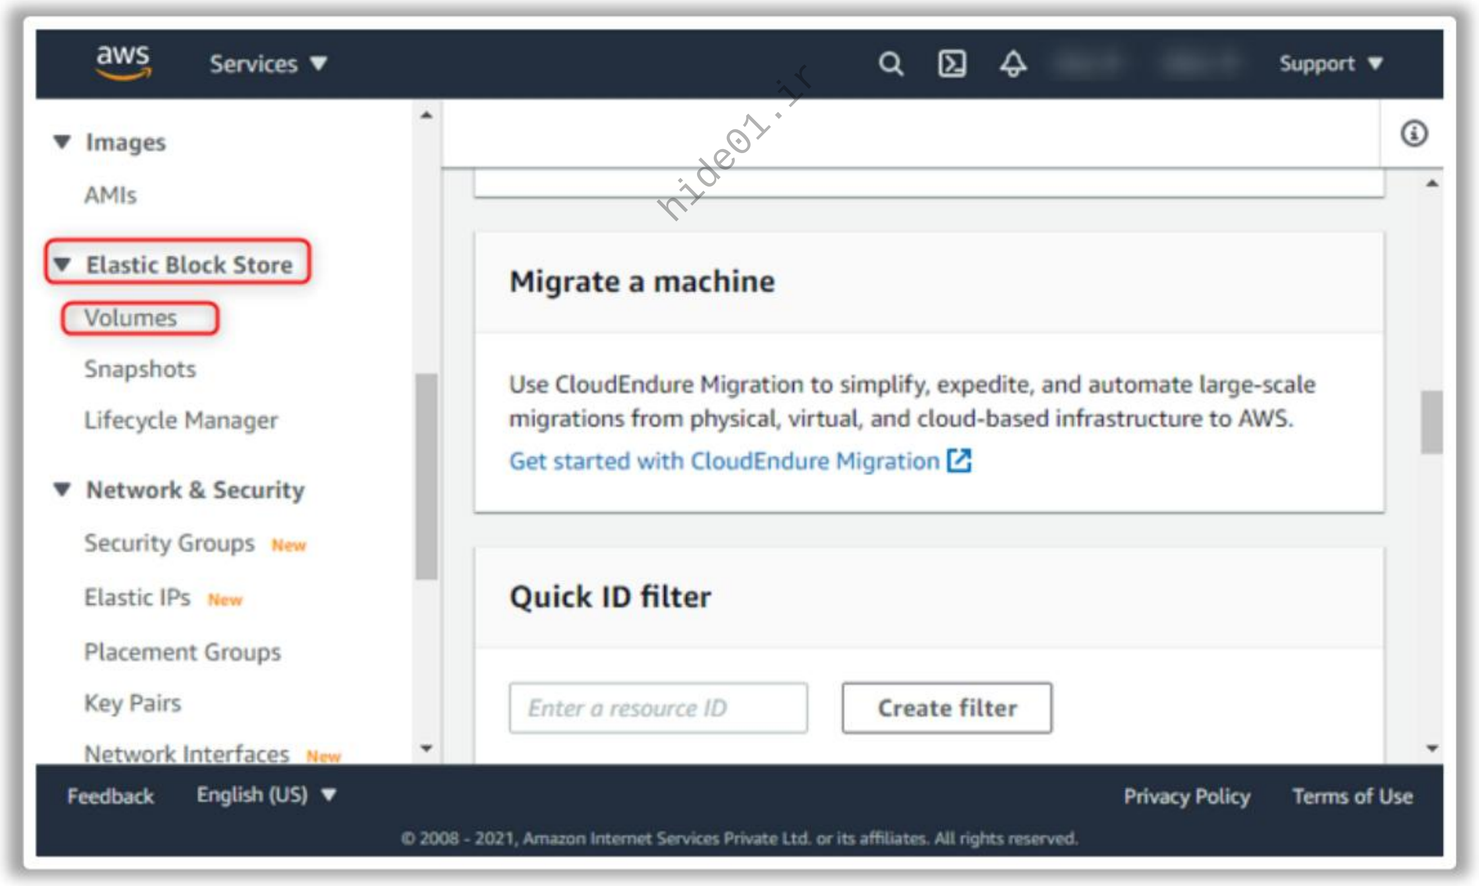


Figure 4.103: Select Volume

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

- Select the unused/orphaned EBS volume.

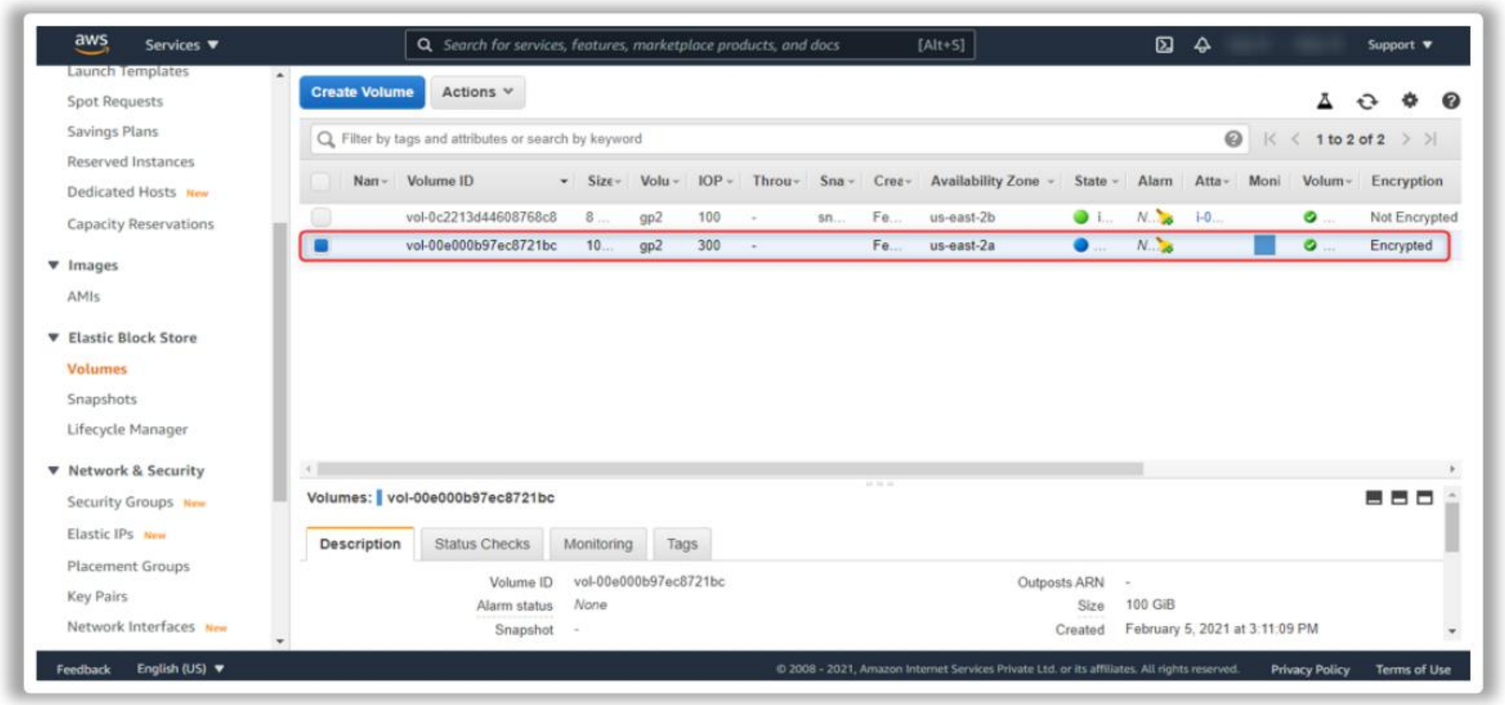


Figure 4.104: Uncheck the Radio Button

- Click on **Actions** dropdown, navigate to **Delete Volume**, and click on it.

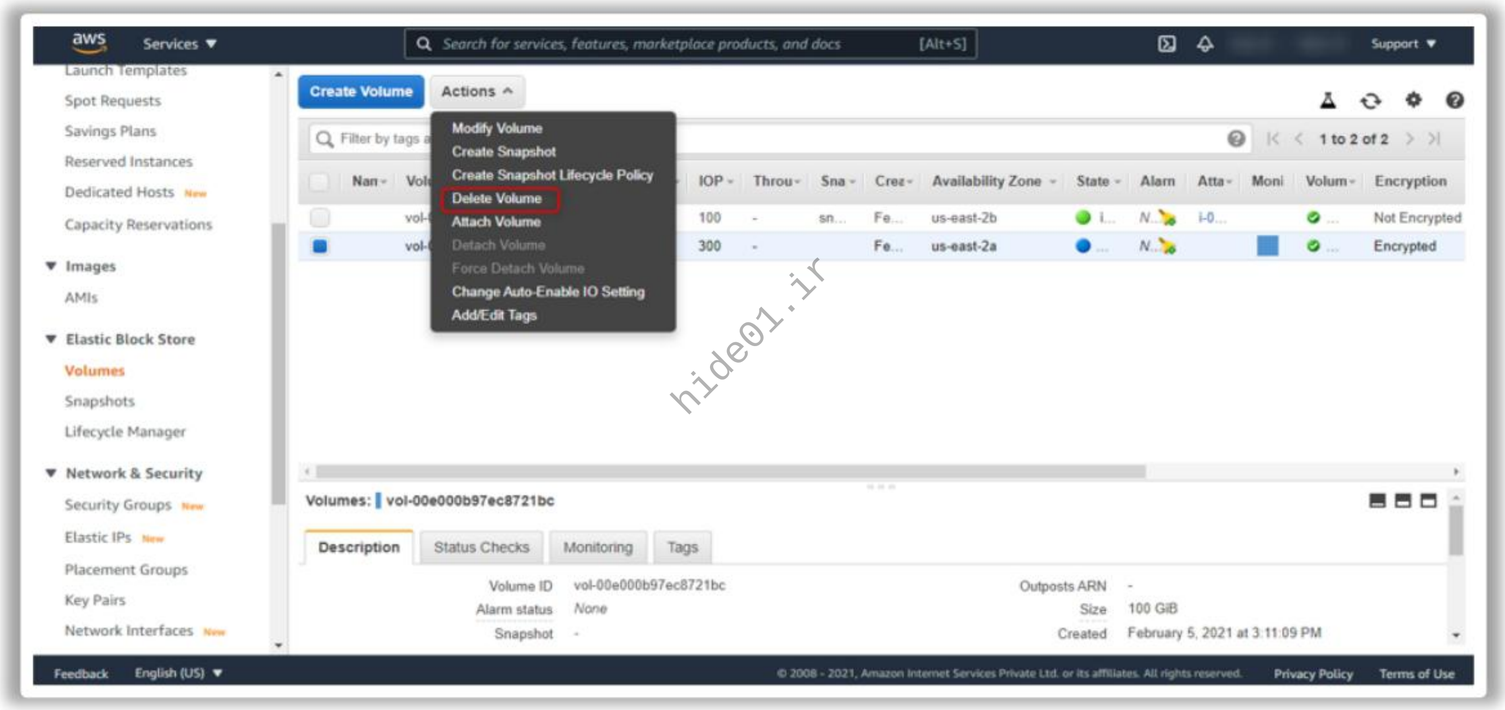


Figure 4.105: Click Delete Volume

- **Delete Volume** window will appear; click on **Yes, Delete** button.

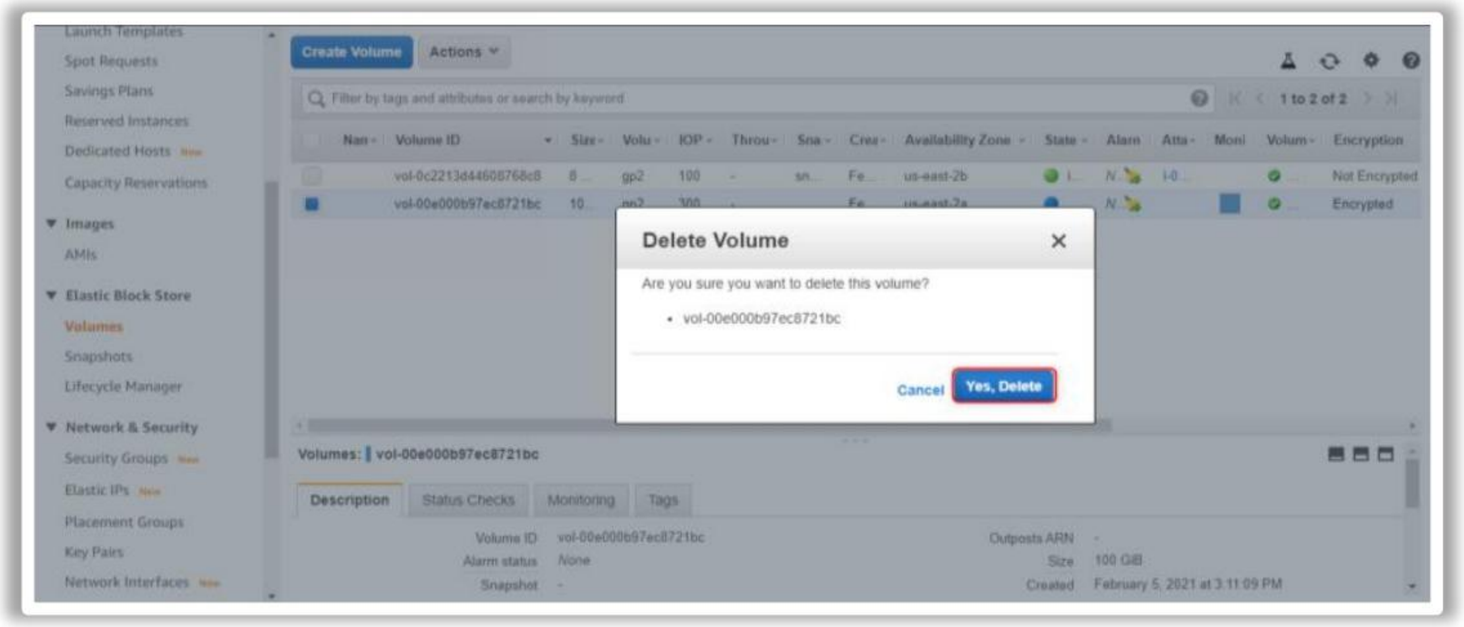
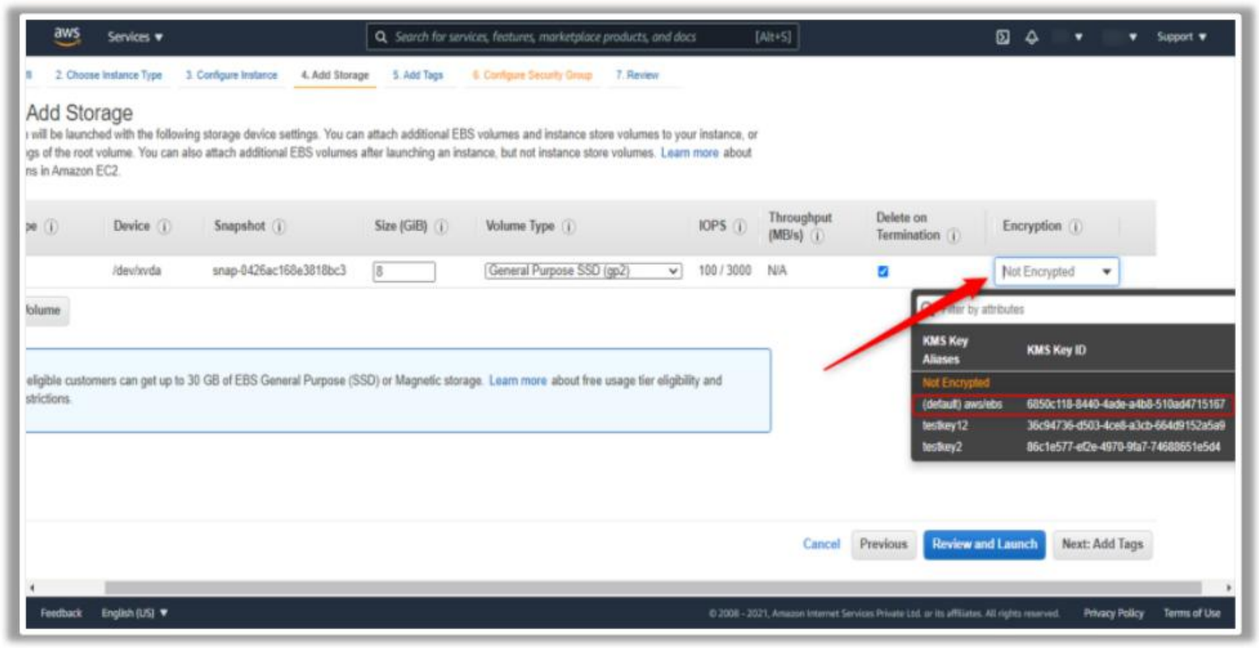


Figure 4.106: Click Yes, Delete

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

Ensure to Launch an EC2 Instance with Encrypted Volume

- To secure the data-at-rest and data-in-transit between the EC2 instance and attached EBS storage, ensure to launch an EC2 instance with **encrypted storage**



Launching EC2 Instance with Encrypted Volume

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

Ensure to Launch an EC2 Instance with Encrypted Volume

To secure the data-at-rest and data-in-transit between the EC2 instance and attached EBS storage, ensure to launch an EC2 instance with encrypted storage. The security administrator can use customer-managed keys or default keys to encrypt the storage while launching the EC2 instance.

Source: https://aws.amazon.com

Ensure to launch an EC2 instance with encrypted volume, perform the following steps:

- Sign-in to AWS Management Console.

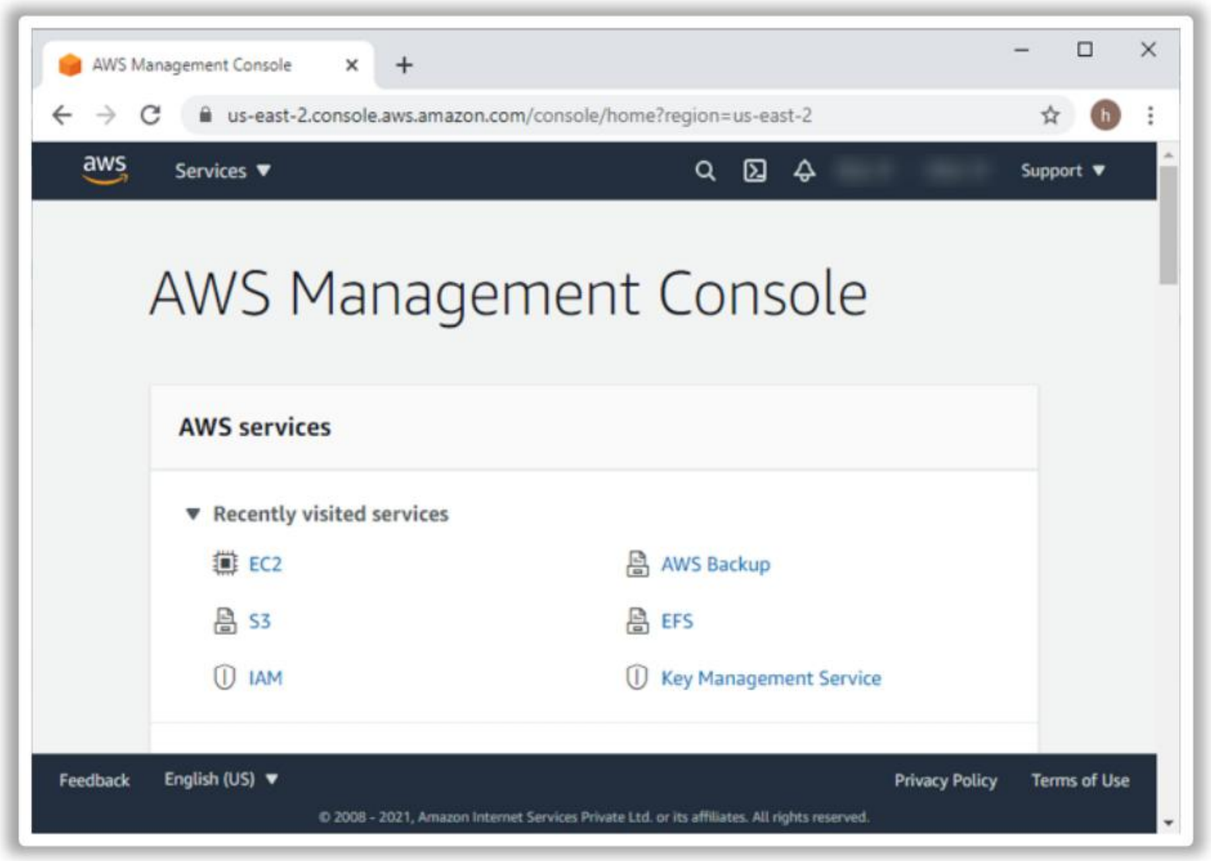


Figure 4.107: AWS Management Console

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/.

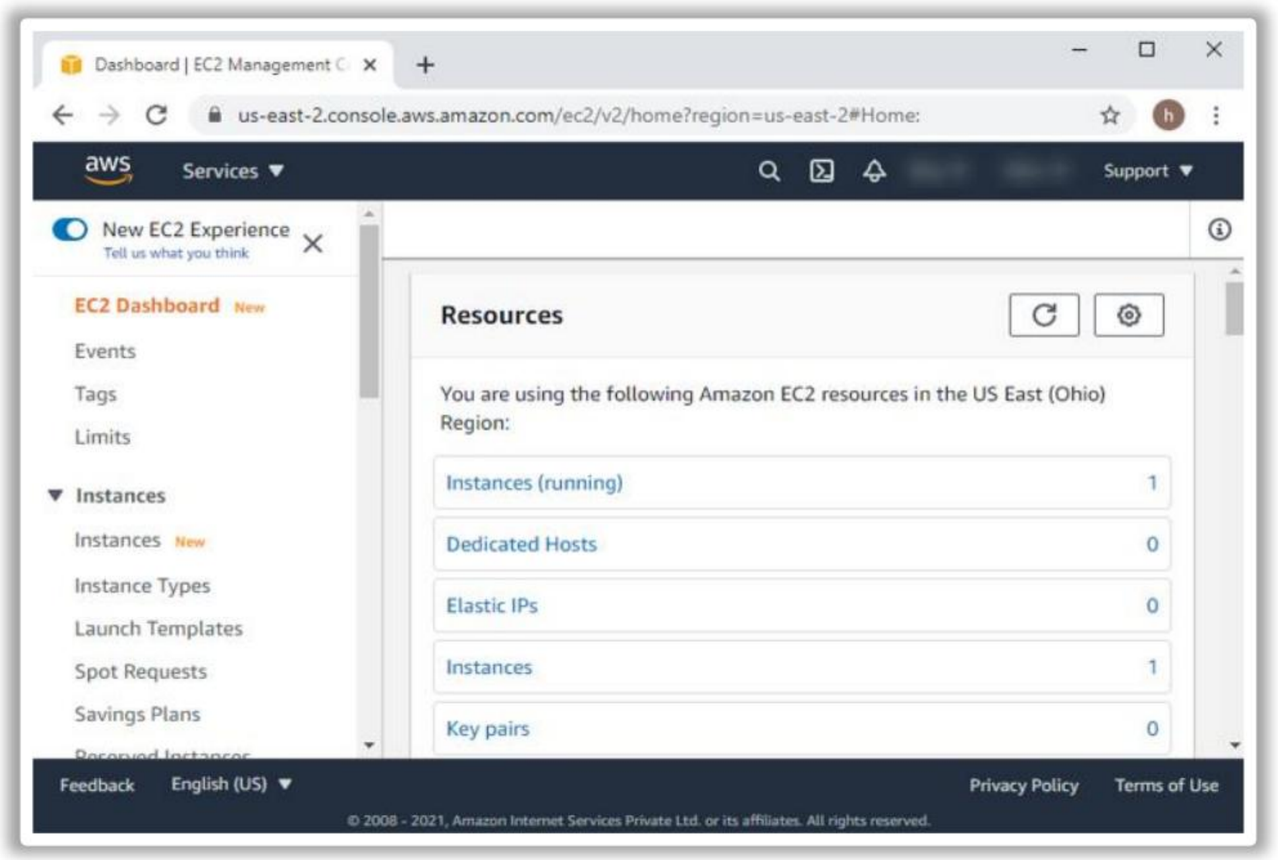


Figure 4.108: EC2 Dashboard

- Navigate and click on **Instances** dropdown; then select and click on **Instances.**

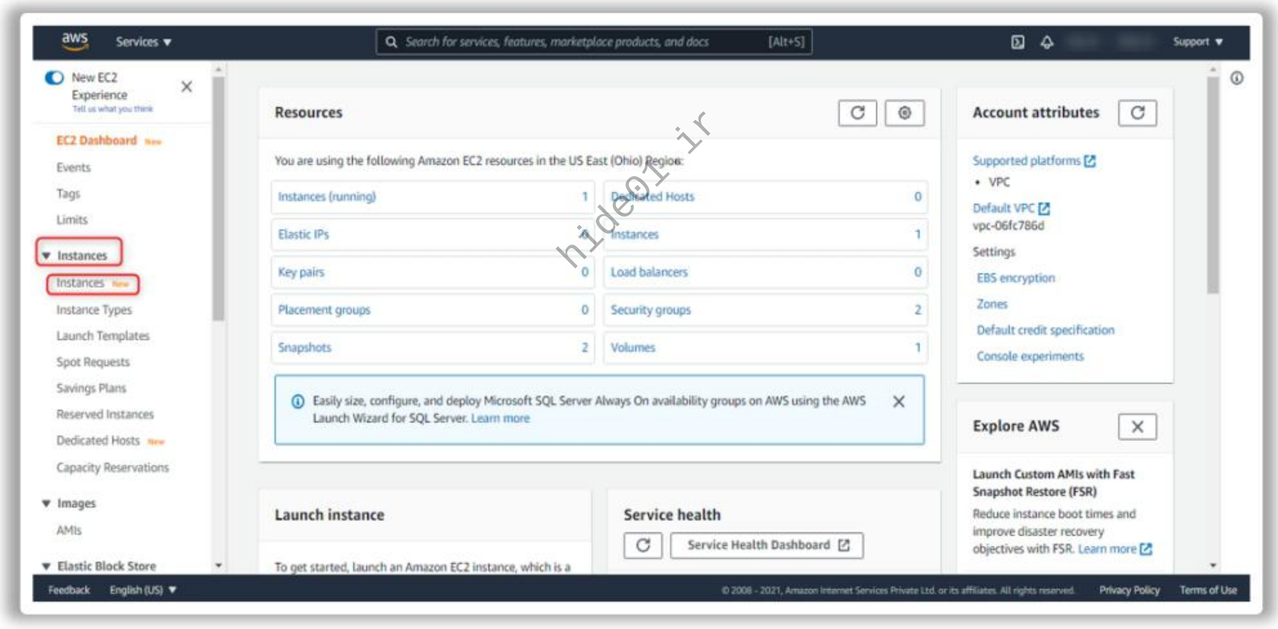


Figure 4.109: Select Instances

- Click on Launch instances button.

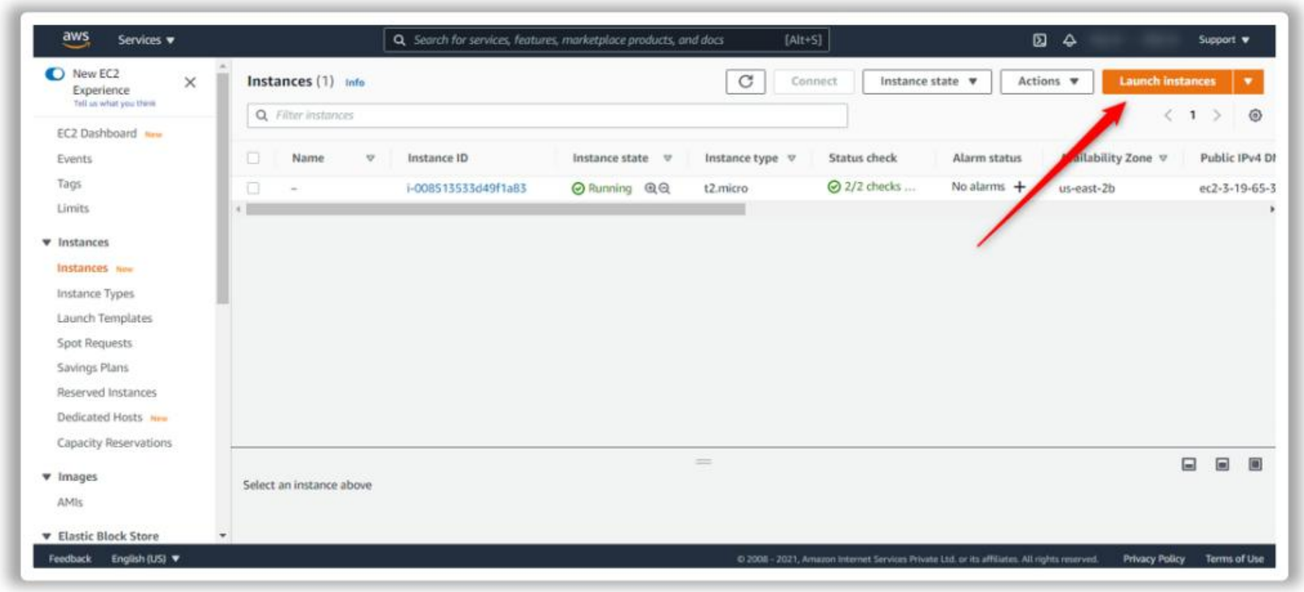


Figure 4.110: Instances Console

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

- Select the AMI as per your requirement.

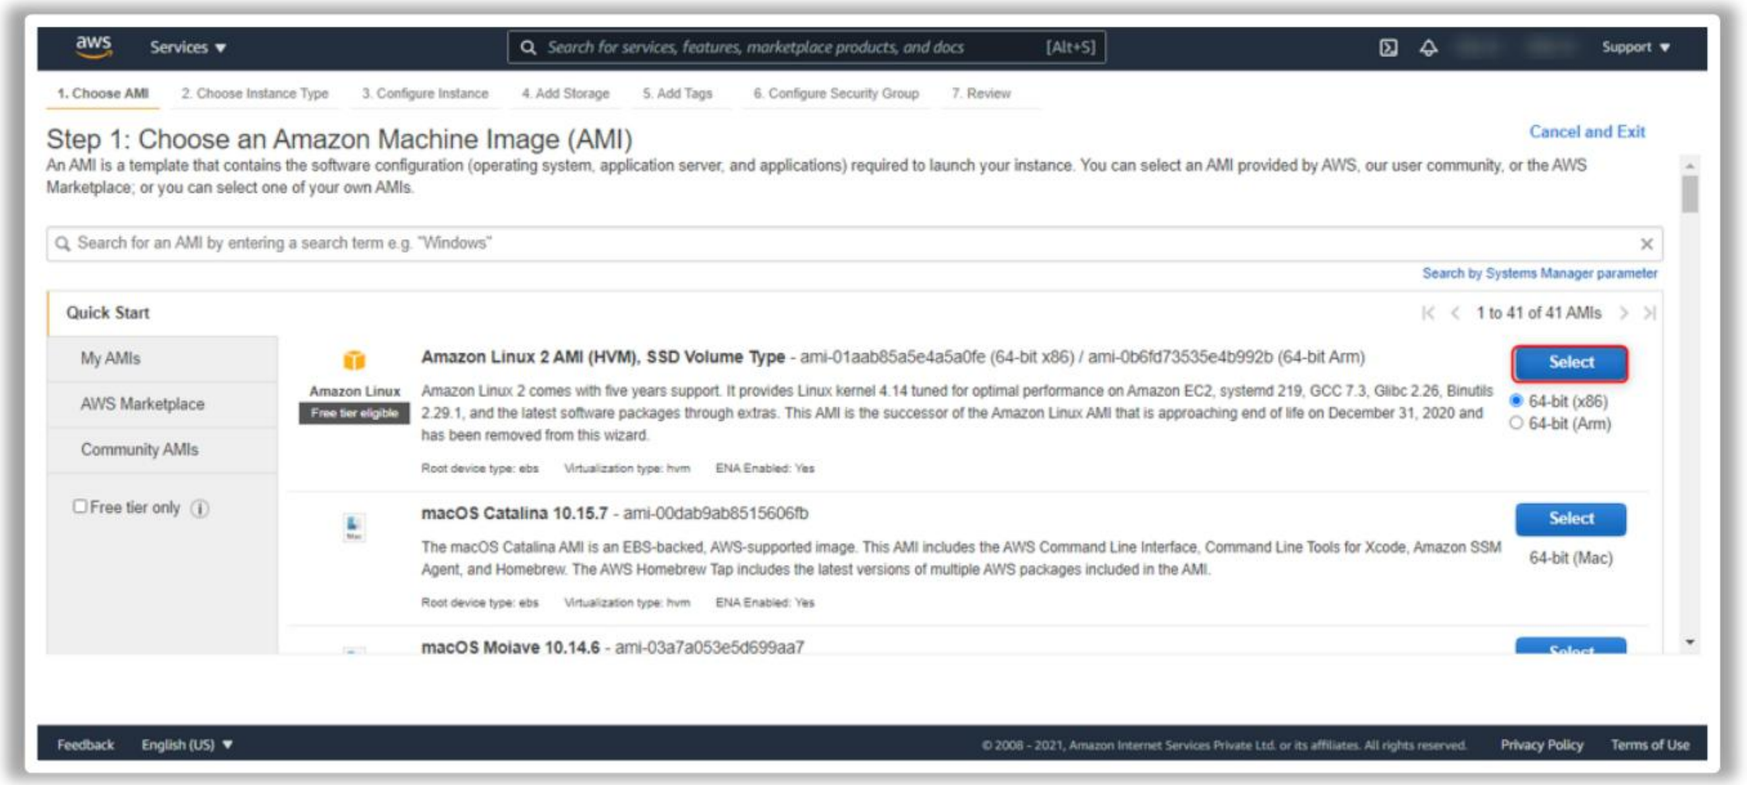


Figure 4.111: Select AMI

- Then, choose an instance type and click on **Review and Launch** button.

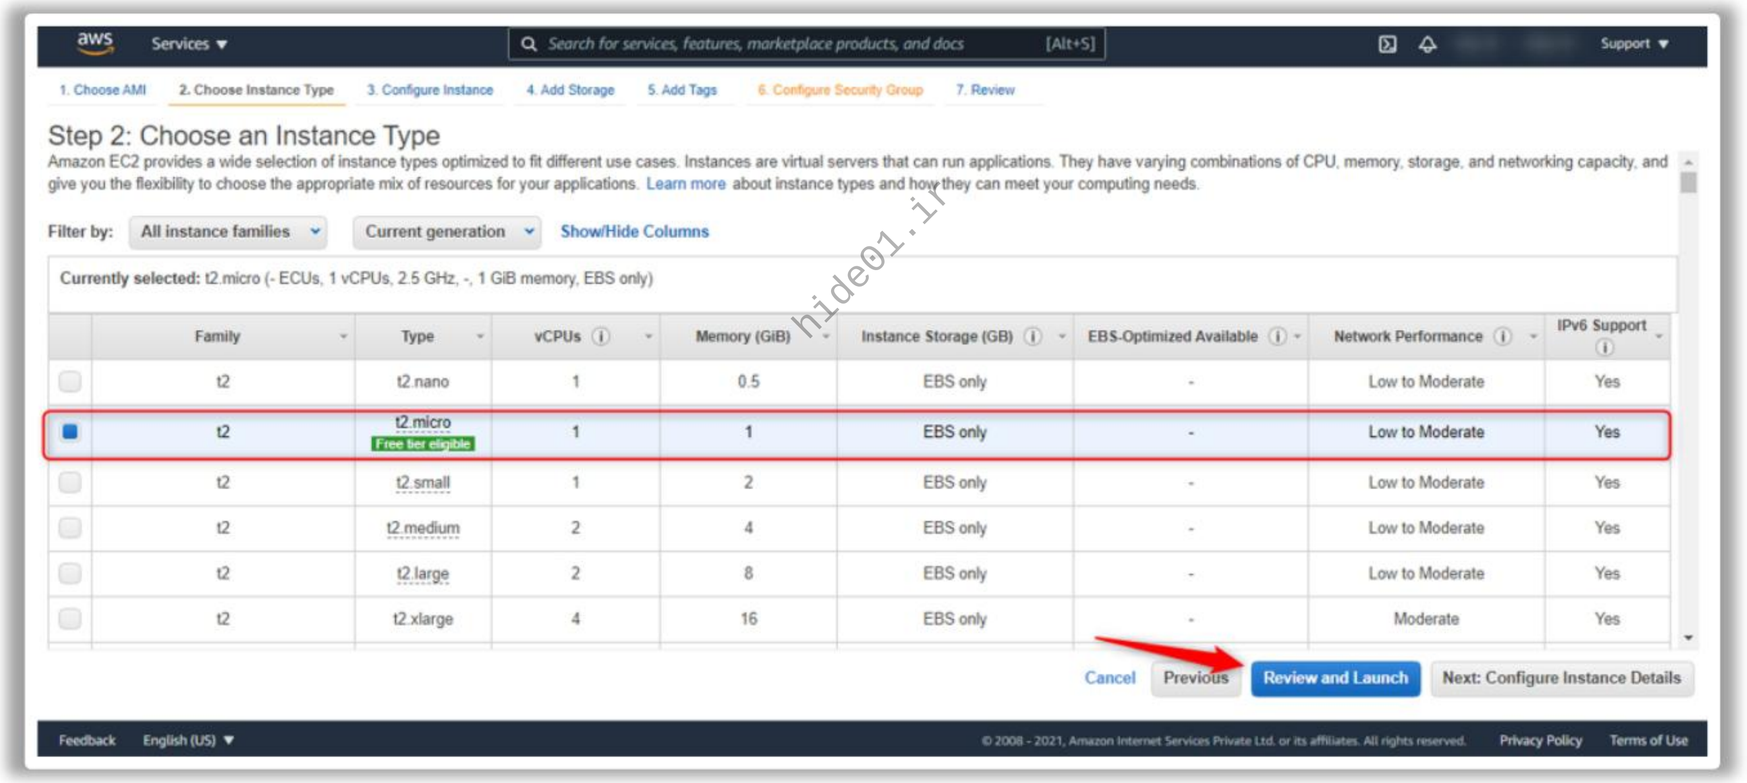


Figure 4.112: Click Review and Launch

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

- In **Step 7: Review Instance Launch**, navigate to **Storage** and click on **Edit Storage.**

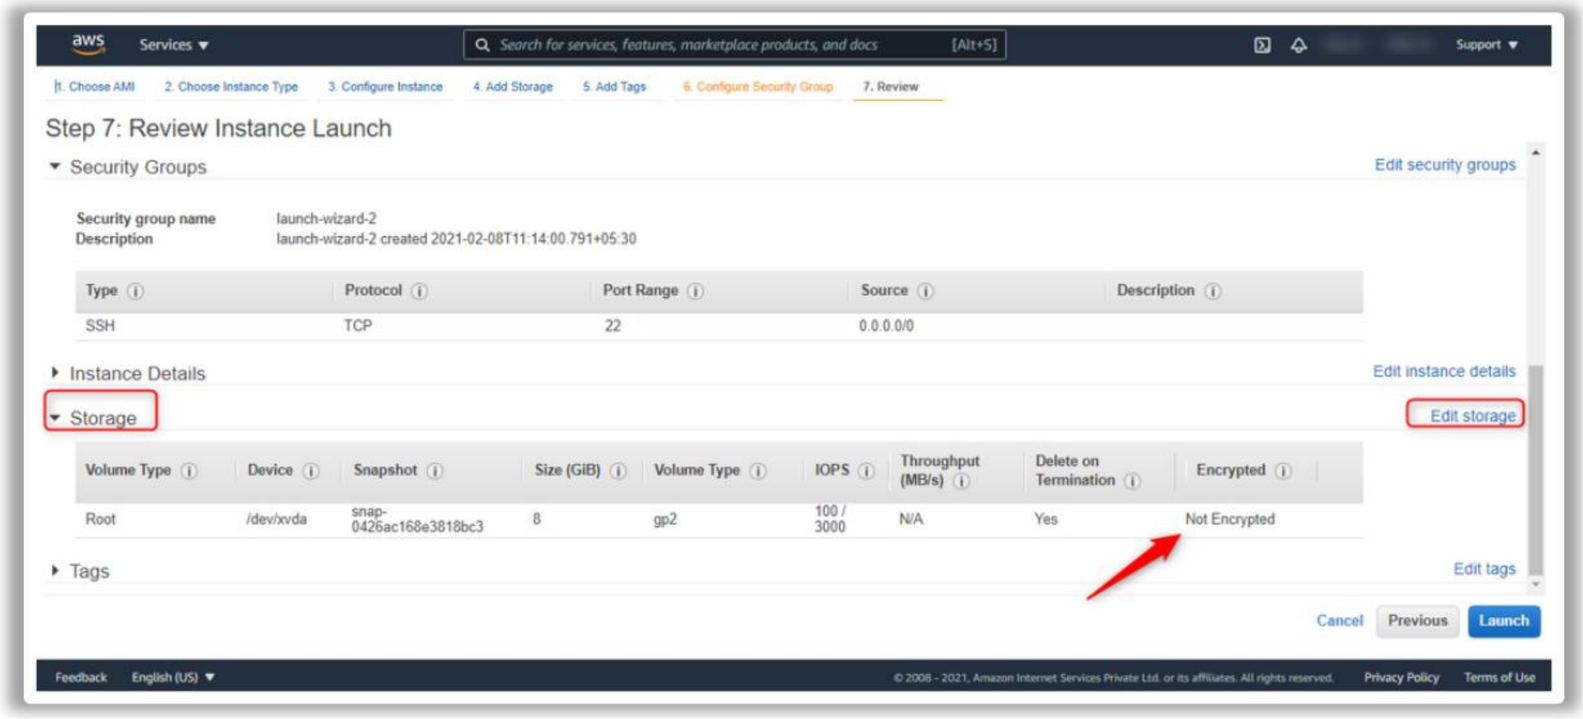


Figure 4.113: Select Edit Storage

- In the **Add Storage** window, under Encryption, click on the dropdown and select the encryption key as per your requirement and click on **Review and Launch** button.

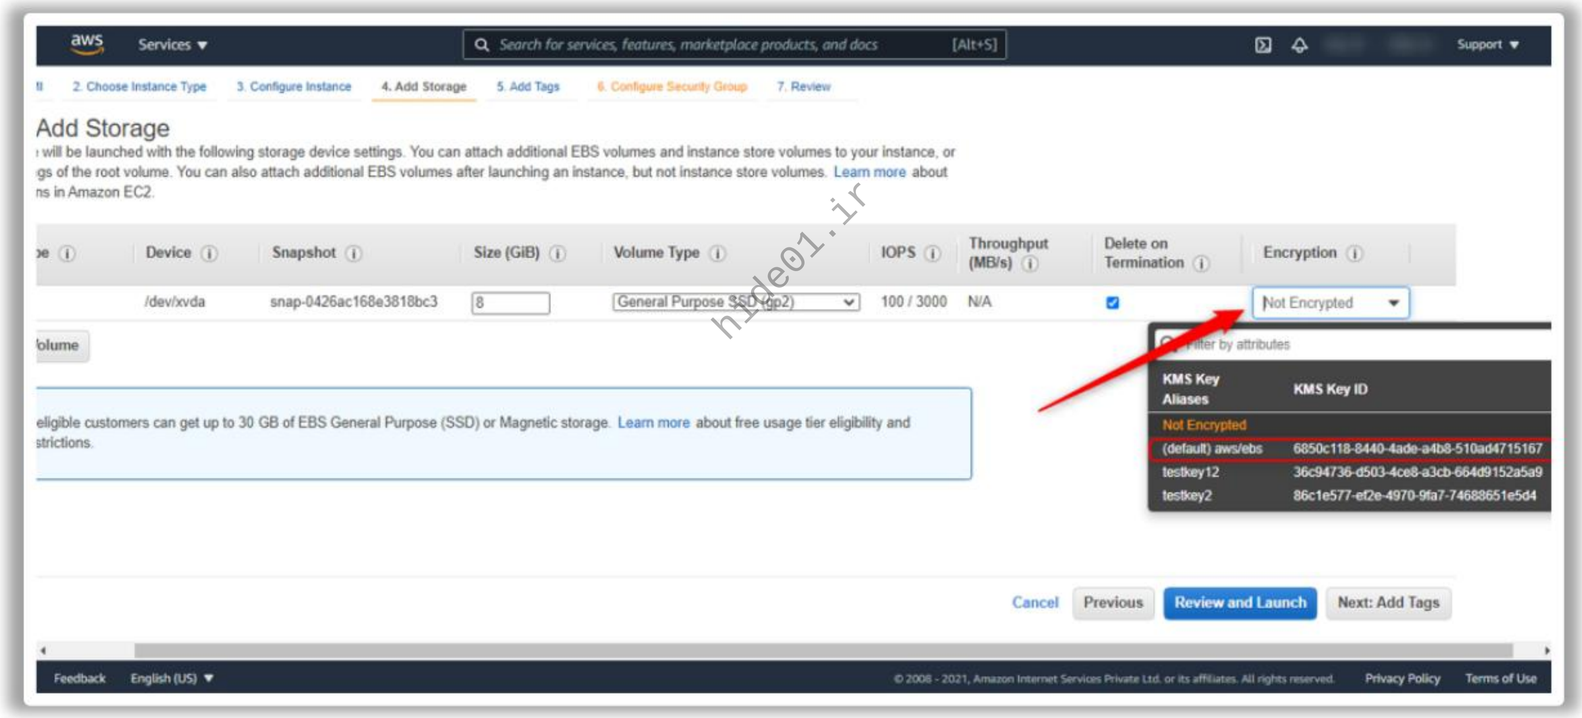


Figure 4.114: Select Add Storage

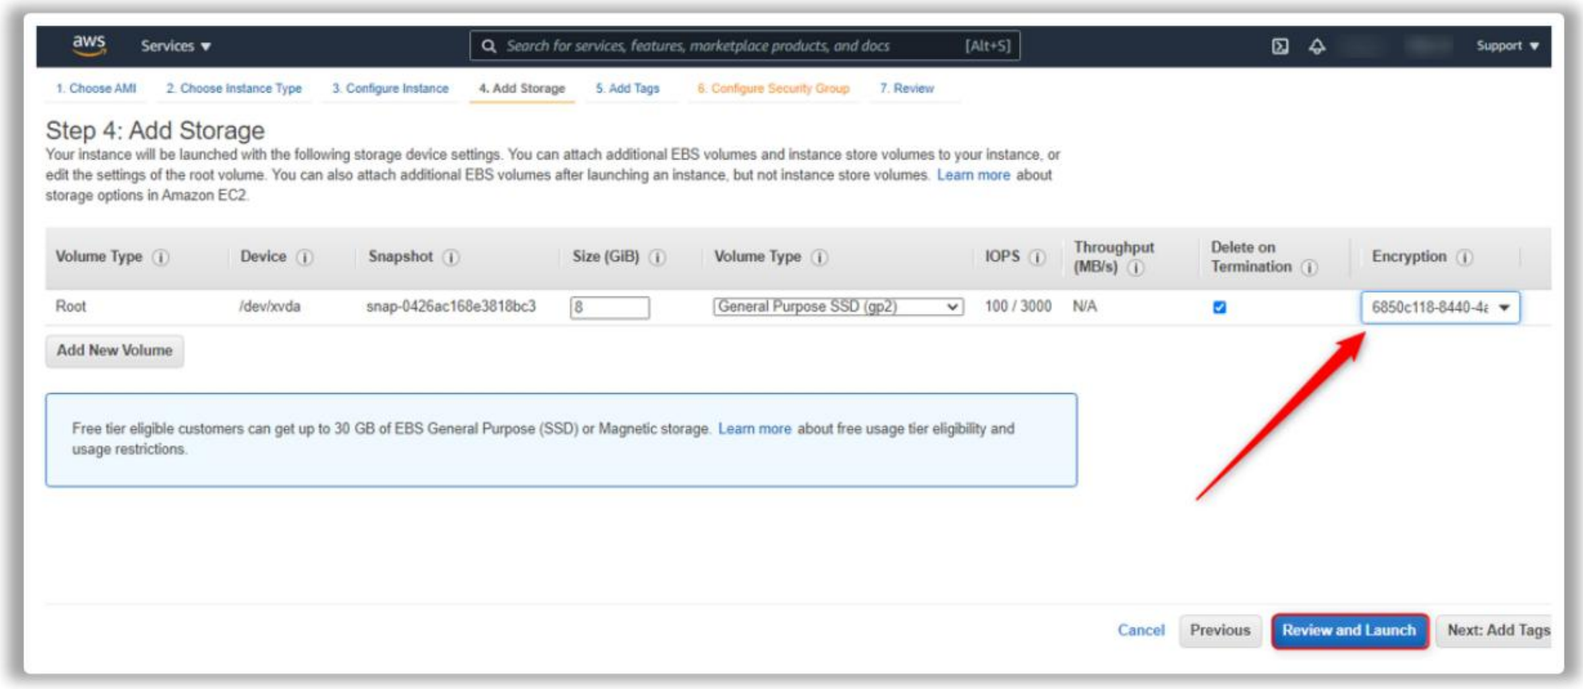


Figure 4.115: Click Review and Launch

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

- You will observe **Encrypted** under **Storage** section; then click on **Launch** button.

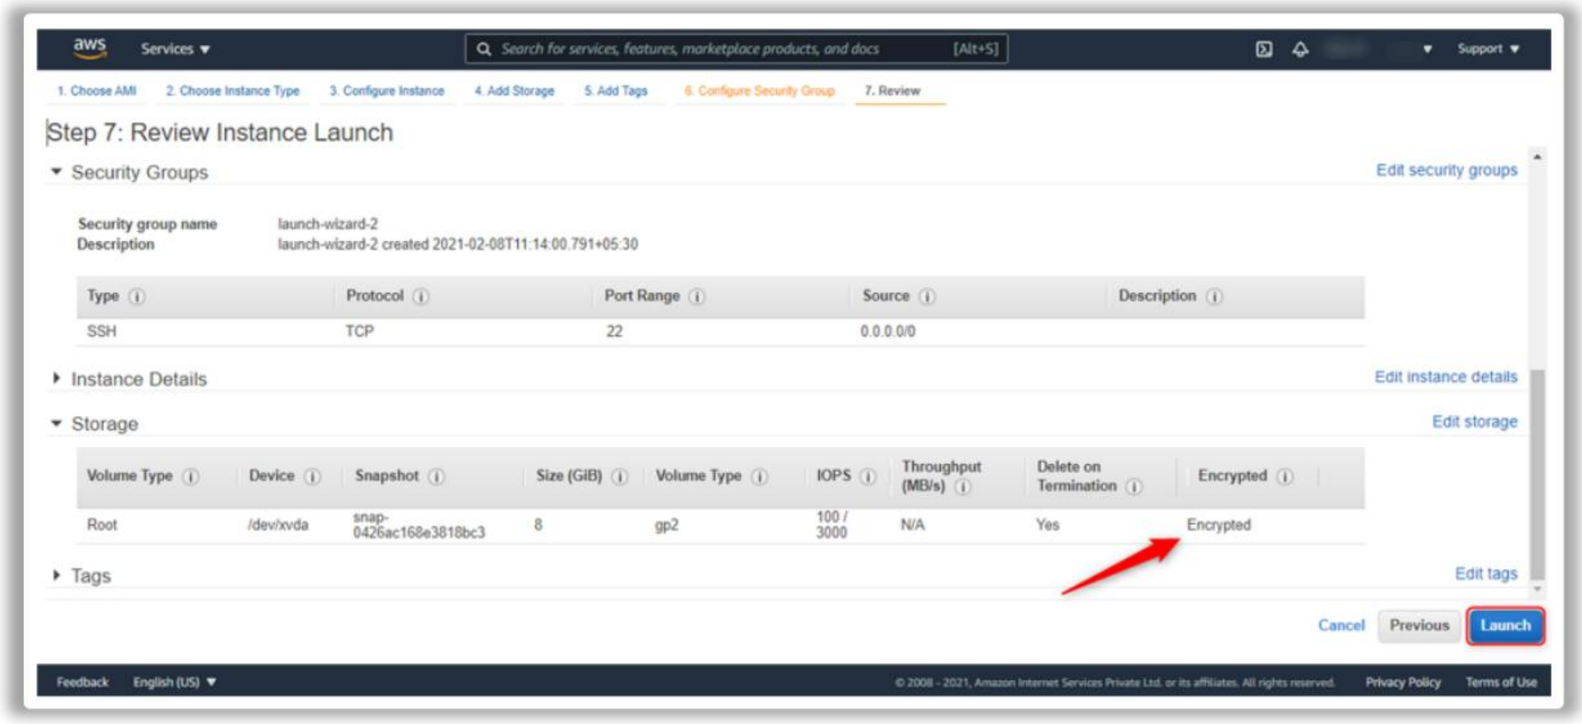


Figure 4.116: Select Launch

- Select an existing key pair or create a new key pair windows will appear, based on your requirement; you can either create a new key pair, choose an existing key pair, or proceed without a key pair. Then, click on Launch instance.

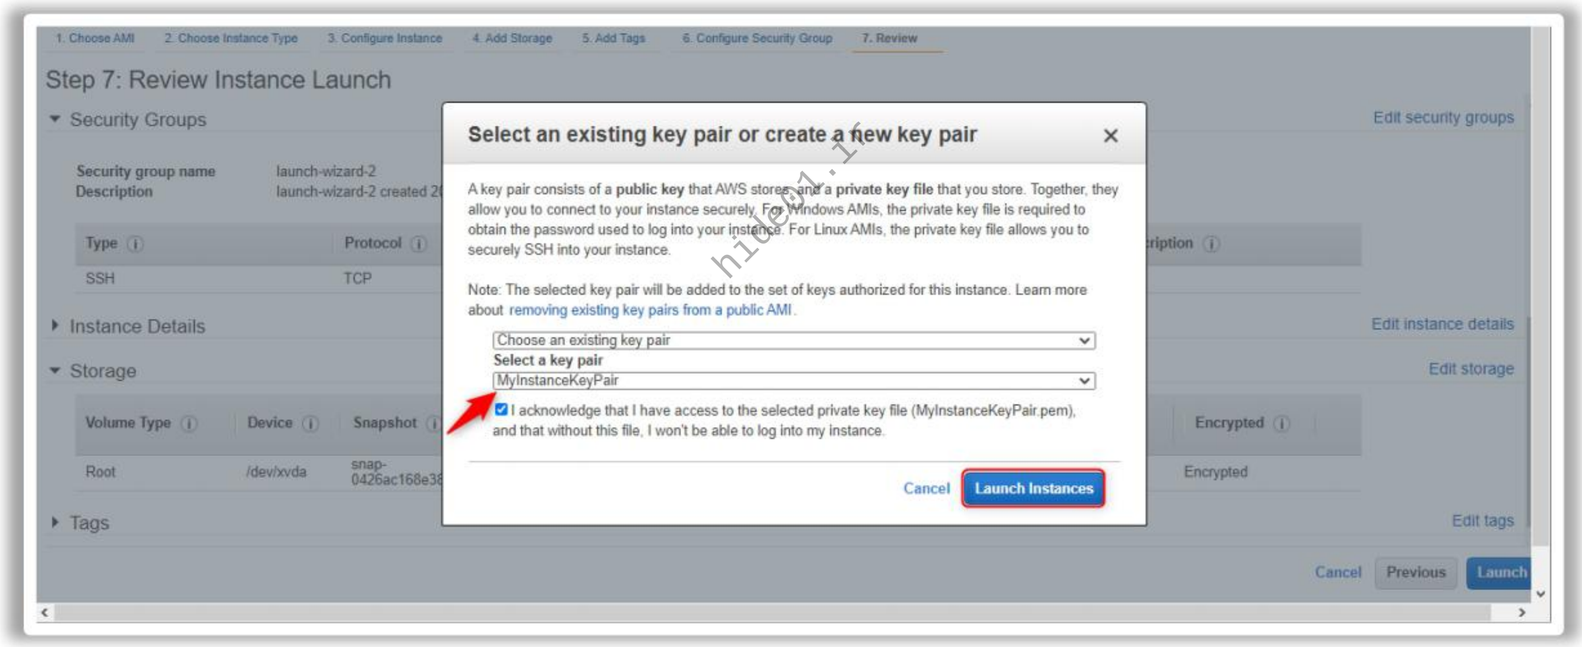


Figure 4.117: Select Existing Key Pair

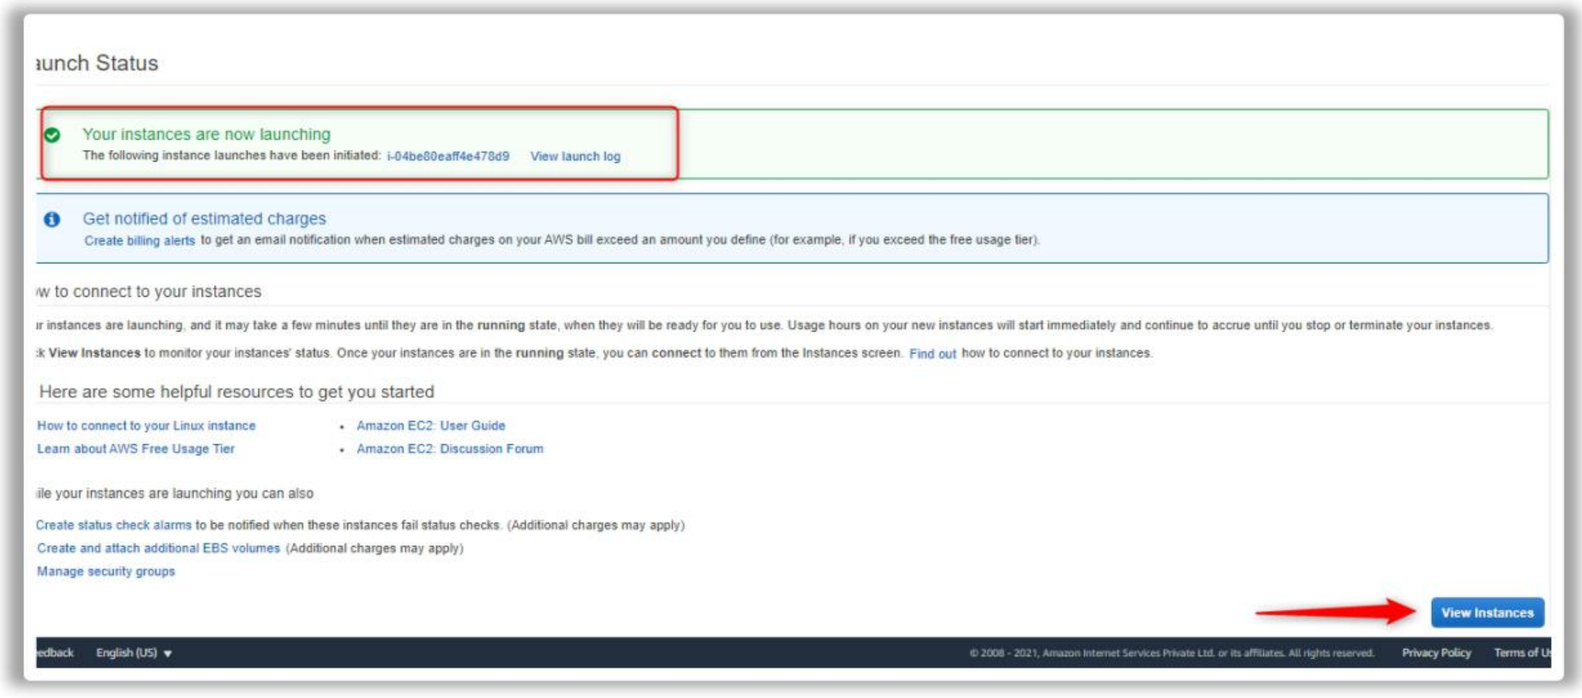


Figure 4.118: Click Launch Instances

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

- After few seconds, you will observe that the instance is launched successfully with encrypted volume.

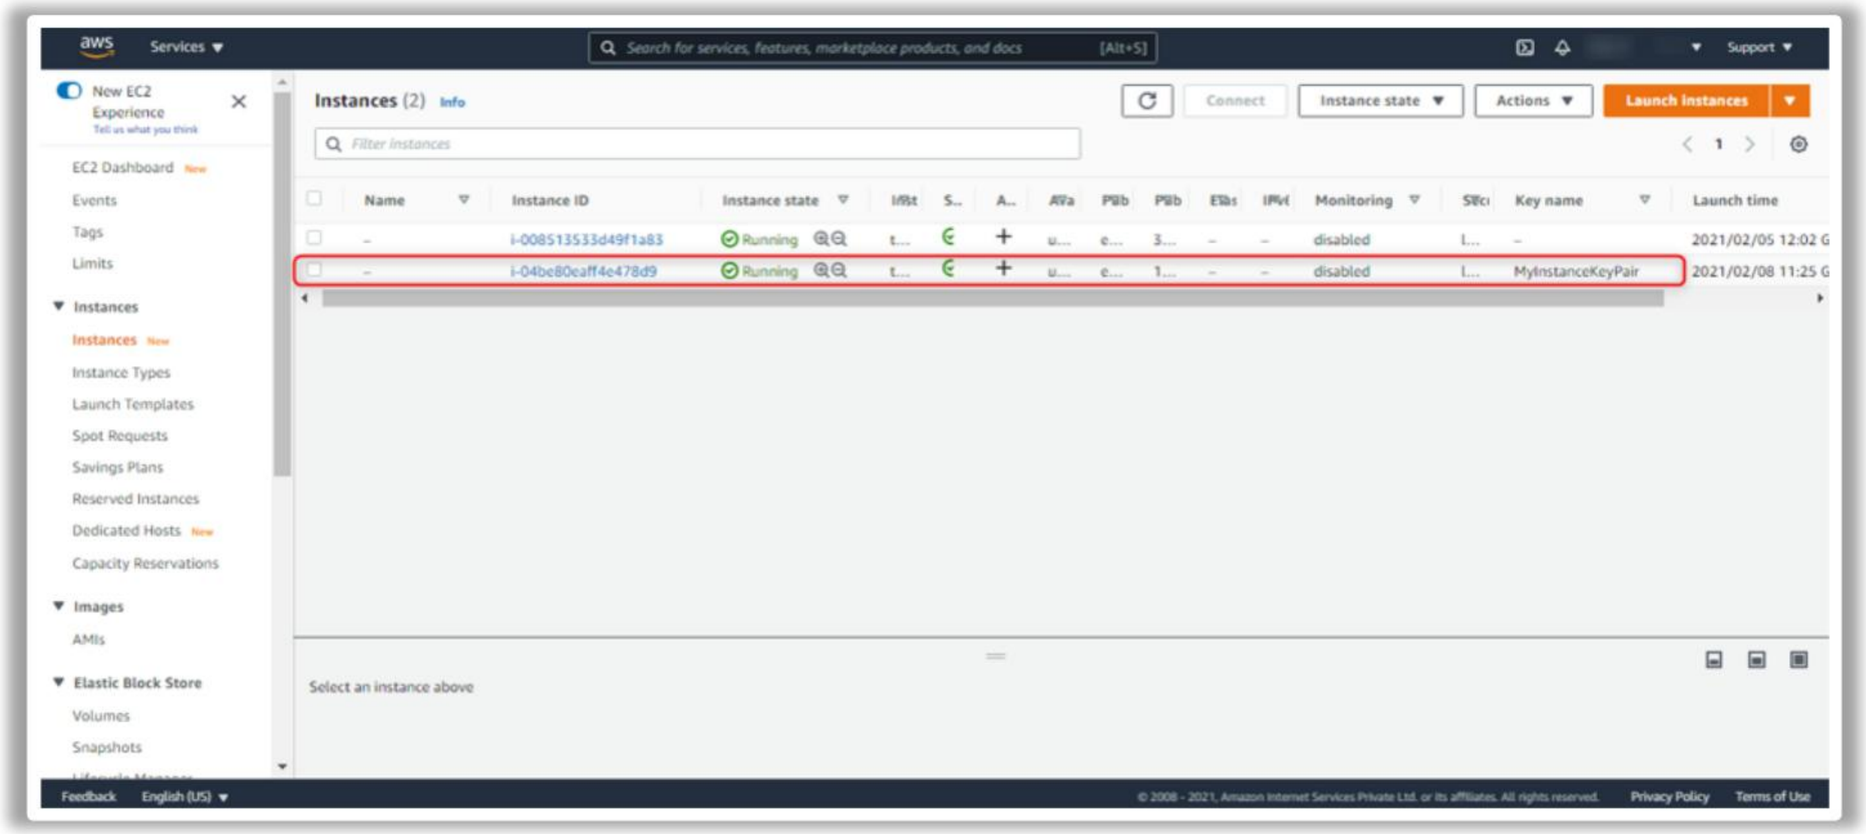
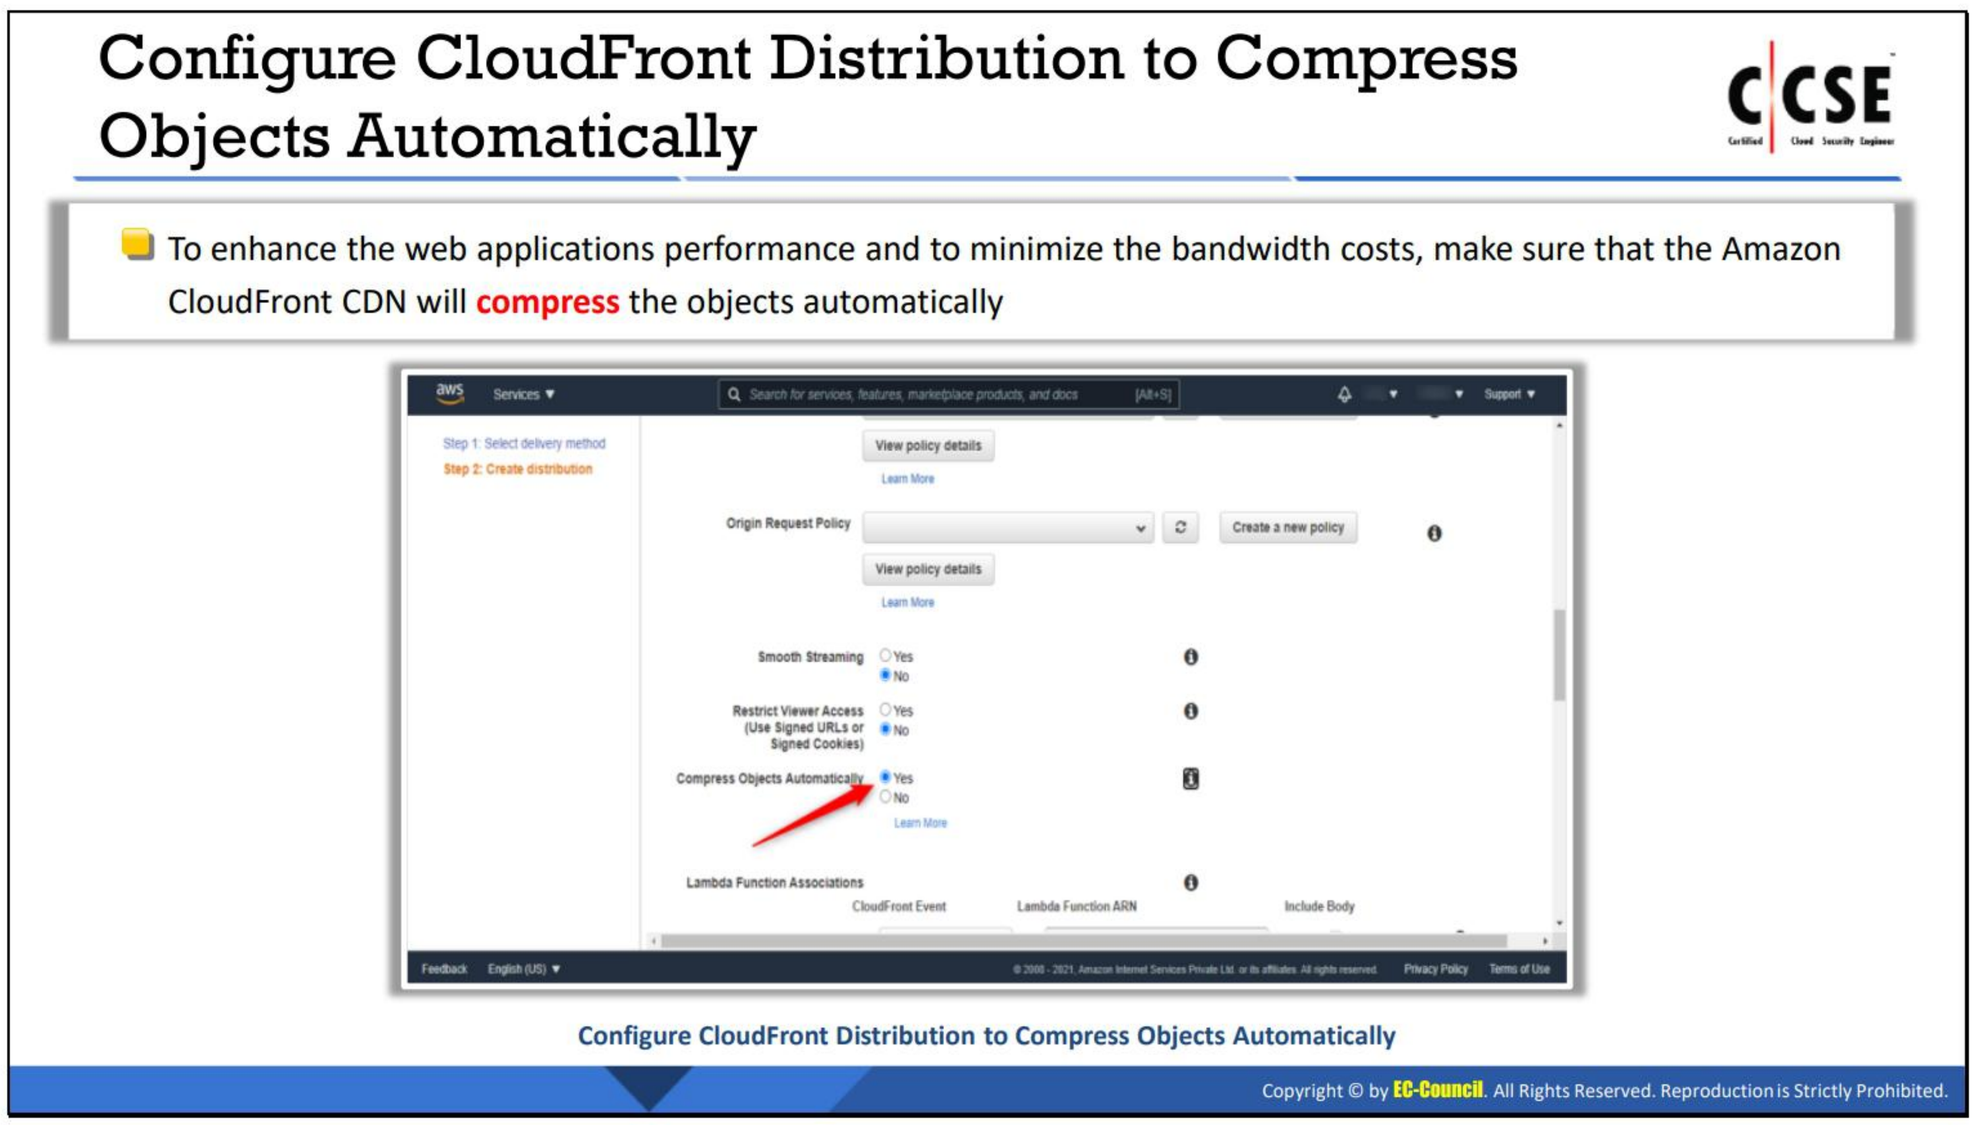


Figure 4.119: Instances Console

hide01.ir

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



Configure CloudFront Distribution to Compress Objects Automatically

To enhance the web application's performance and to minimize bandwidth costs, make sure that the Amazon CloudFront CDN will compress the objects automatically. If the size of the files is smaller, then they can be downloaded with higher speed. This allows fast rendering of the web pages to the user. Apart from this, the charges of CloudFront data transfer depend on the amount of the data that is served. Hence, serving compressed files is less costly than serving uncompressed files.

Source: https://aws.amazon.com

To configure CloudFront distribution to compress objects automatically, perform the following steps:

- Sign-in to AWS Management Console.

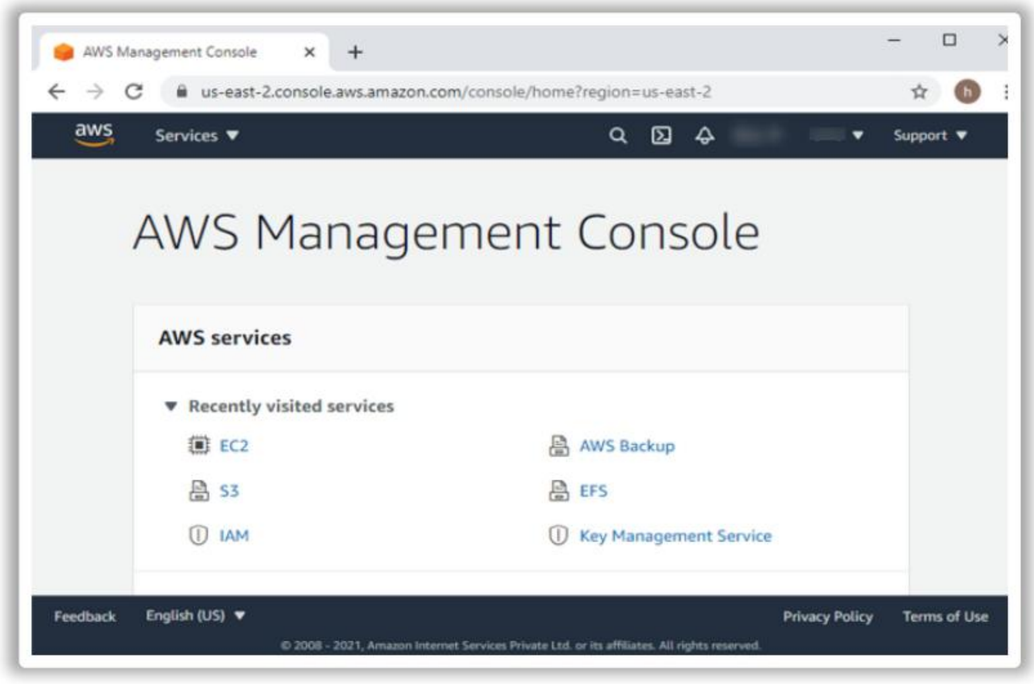


Figure 4.120: AWS Management Console

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

- Click on the following link to open CloudFront dashboard https://console.aws.amazon.com/cloudfront/.
- Click on **Create Distribution** button.

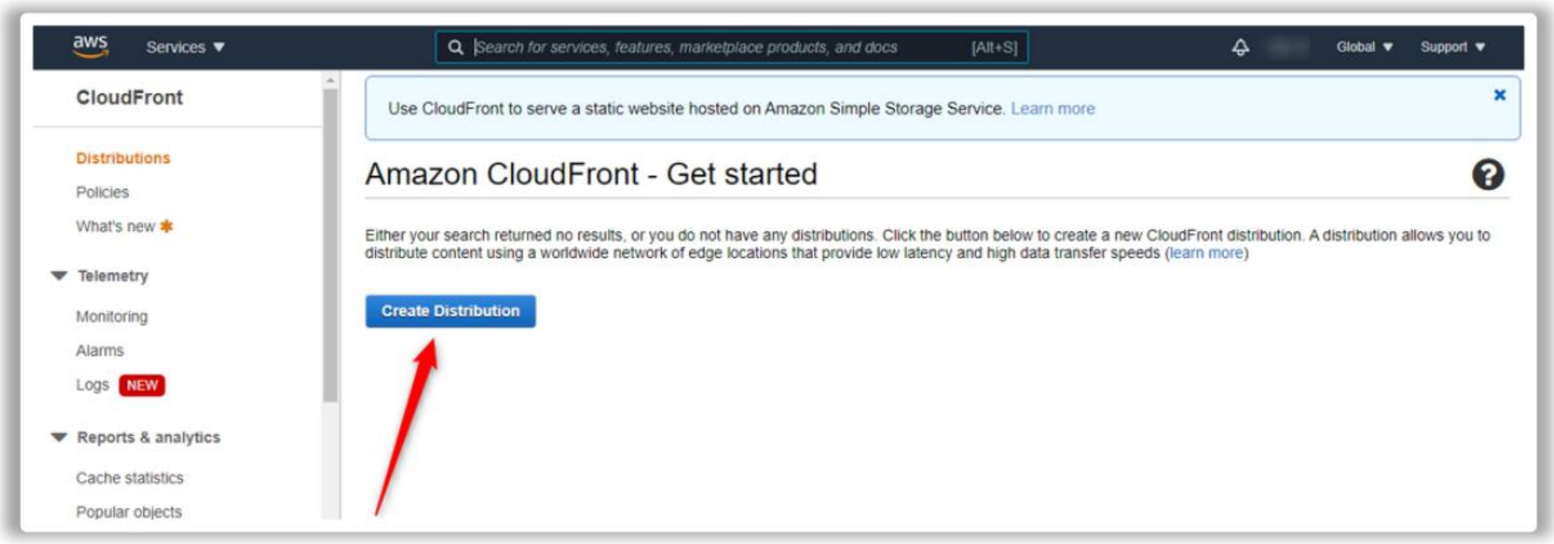


Figure 4.121: Click Create Dashboards

- **Select a delivery method for your content** page will open, click on **Get Started** button.

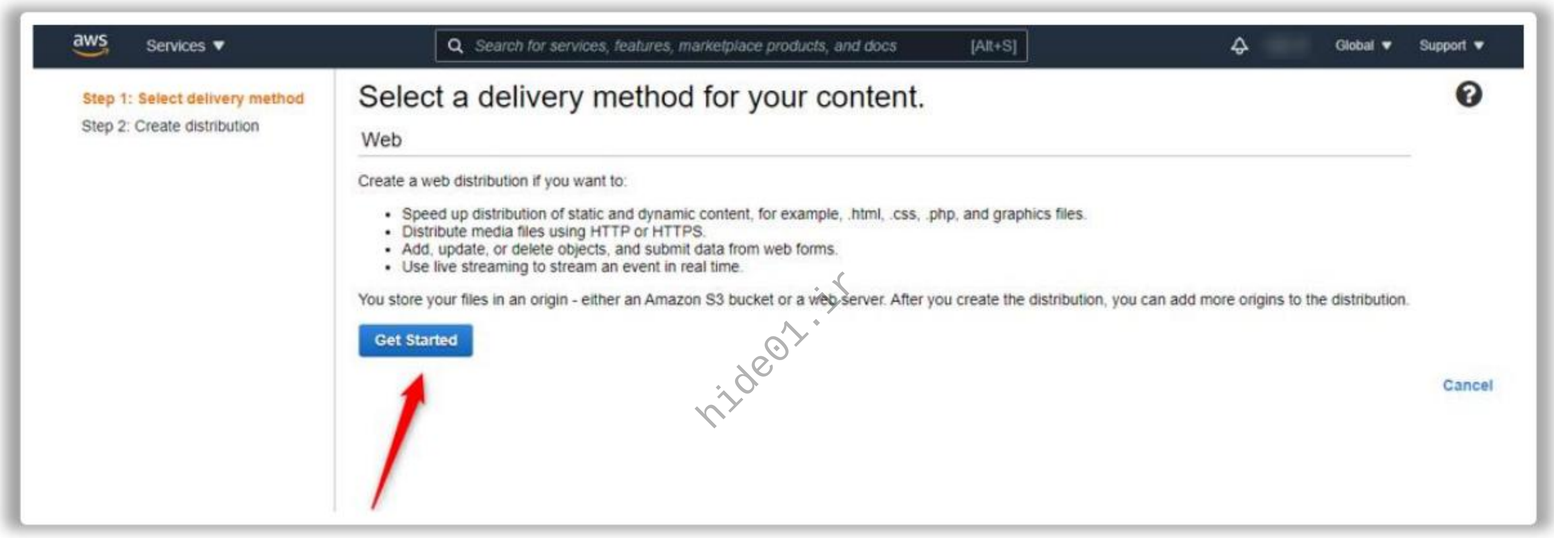


Figure 4.122: Select Get Started

- **Create Distribution** page will open; fill the field and set the settings as per your requirements and feasibility (here we selected **Amazon S3 bucket**).

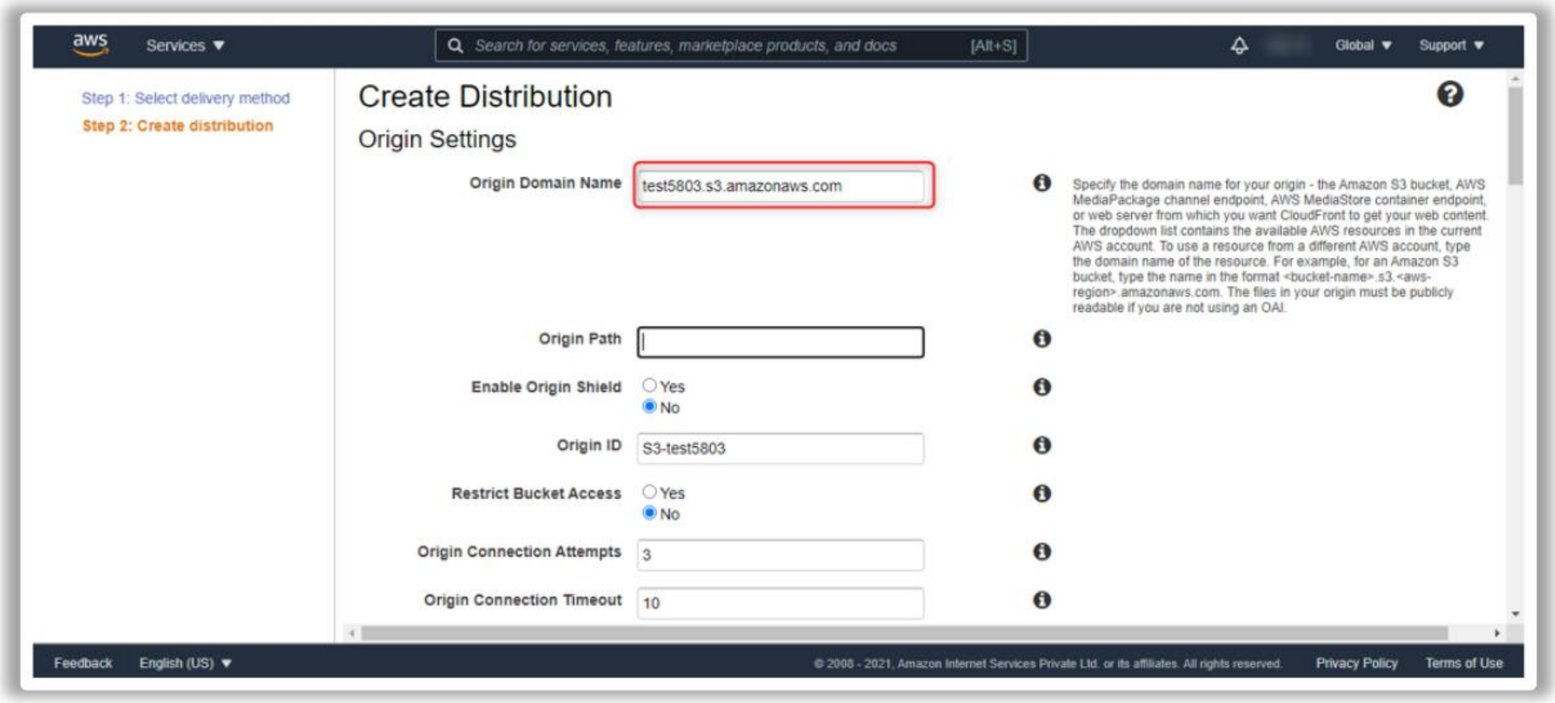


Figure 4.123: Select Amazon S3 Bucket

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

- Under **Default Cache Behavior Settings**, navigate to **Compress Objects Automatically** and select **Yes** radio button.

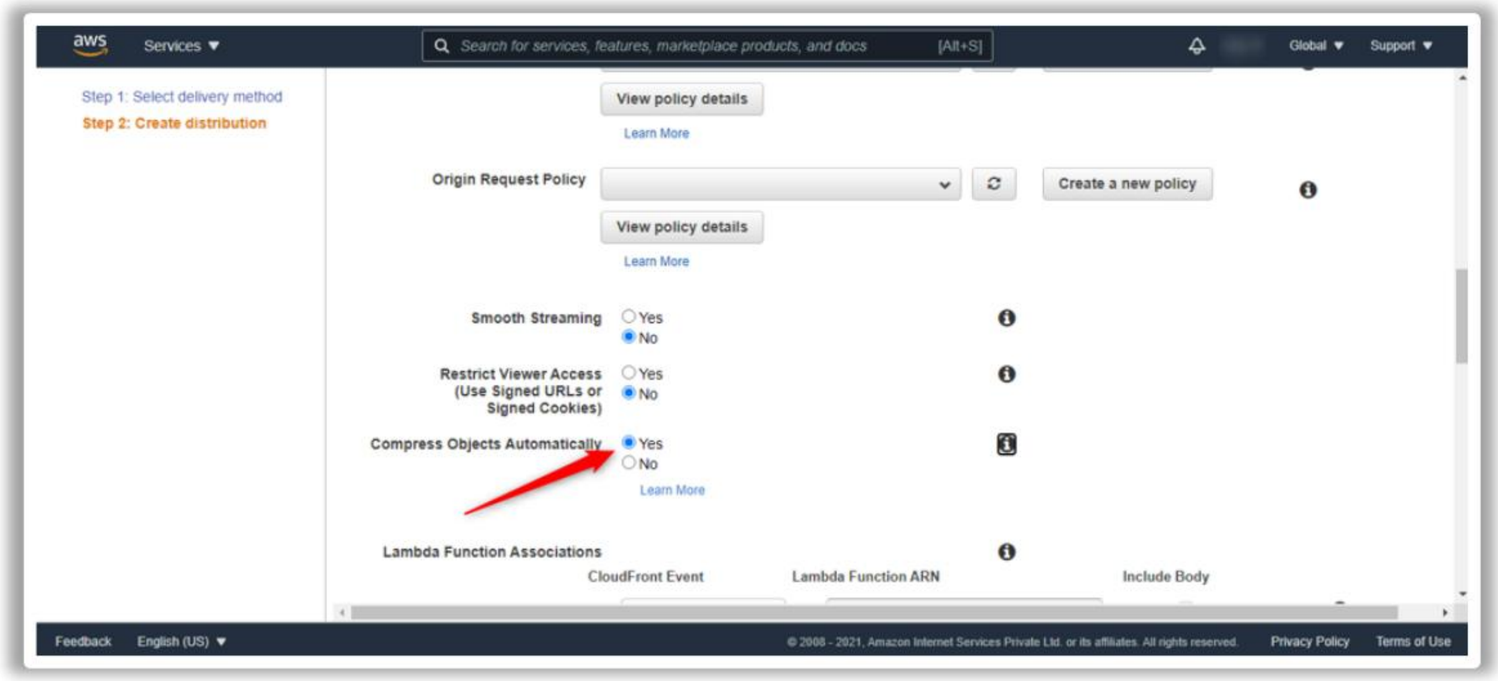


Figure 4.124: Select Yes Radio Button

- Either you can keep default settings or make changes as per your requirements; then click on the **Create Distribution** button.

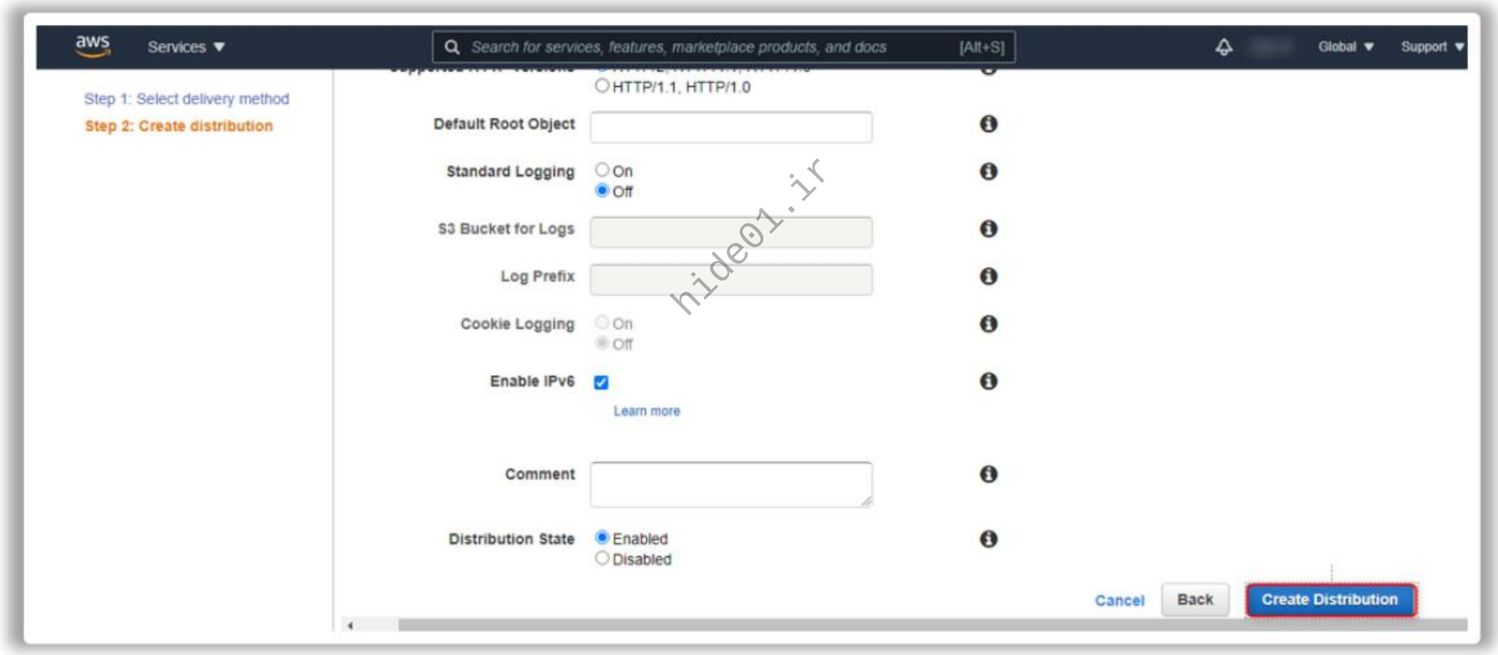


Figure 4.125: Select Create Distribution

- After a few seconds, the CloudFront Distribution will be deployed successfully with compress object automatically feature.

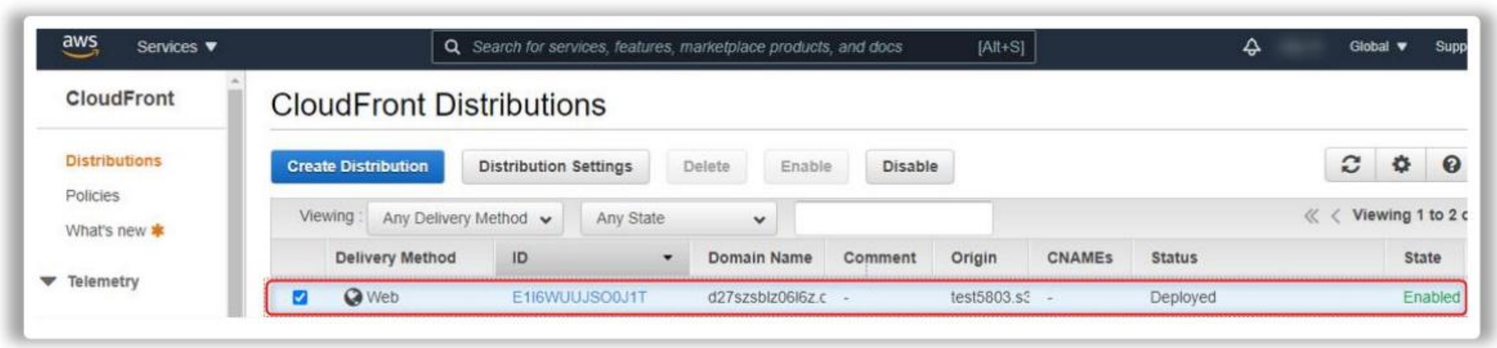
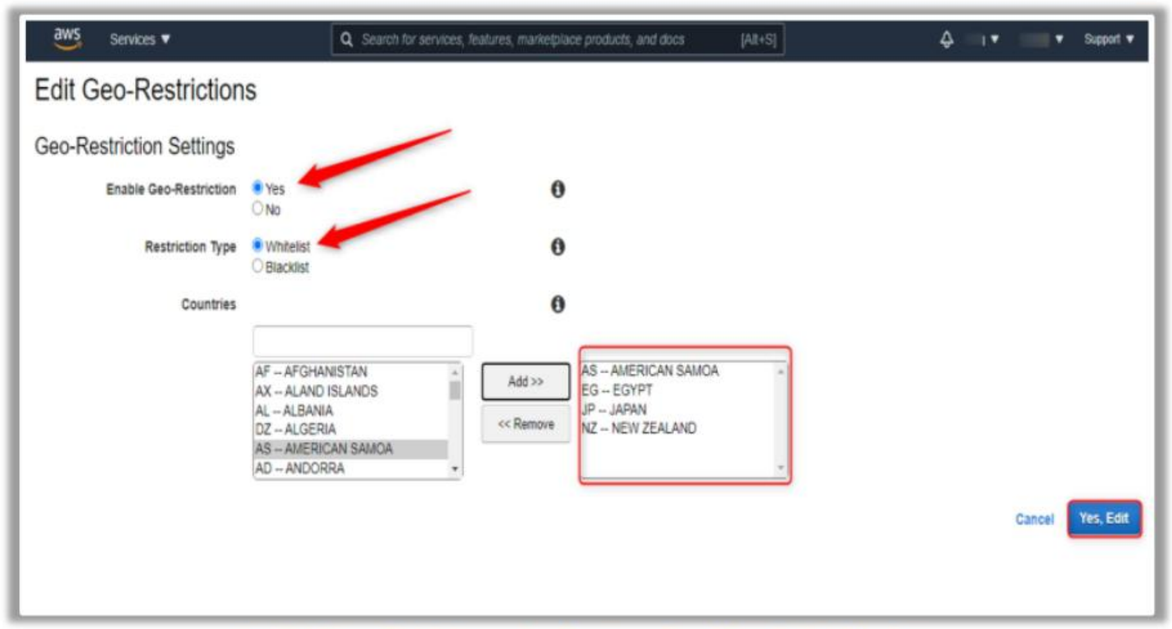


Figure 4.126: CloudFront Distribution Console

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

Enable Geo-Restriction for Amazon CloudFront Distribution

To mitigate DDoS attacks, the cloud security administrator can **whitelist or blacklist** the users from specific locations to access the web application content by enabling geo-restriction



Enable Geo-Restriction for Amazon CloudFront Distribution

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

Enable Geo-Restriction for Amazon CloudFront Distribution

To mitigate DDoS attacks, the cloud security administrator can whitelist or blacklist the users from specific locations to access the web application content by enabling geo-restriction.

hide01.ir

Source: https://aws.amazon.com

To enable geo-restriction for Amazon CloudFront distribution, perform the following steps:

- Sign-in to AWS Management Console.

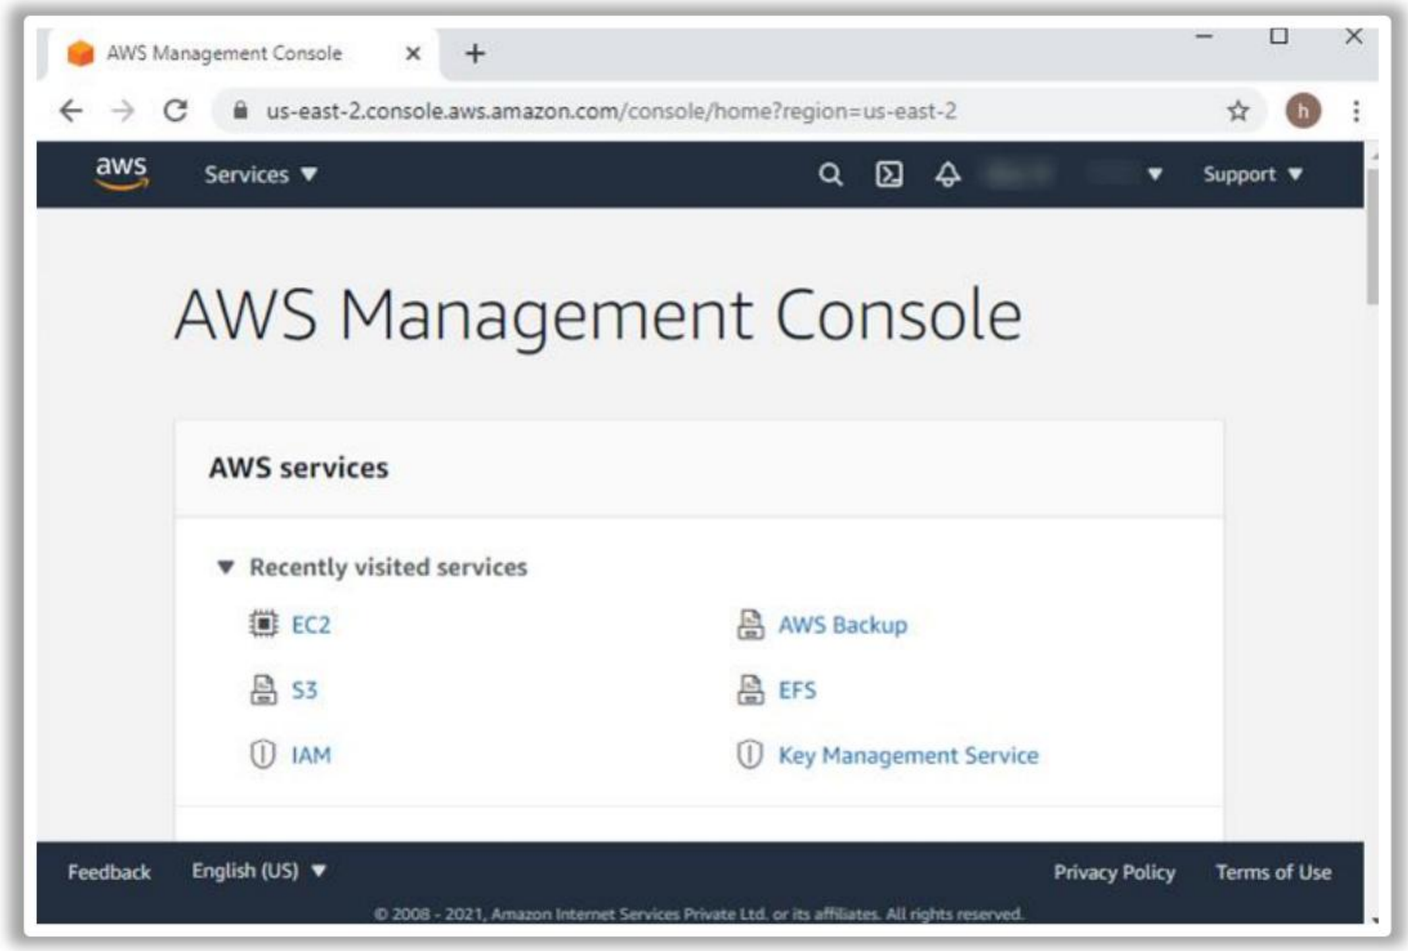


Figure 4.127: AWS Management Console

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

- Click on the following link to open CloudFront dashboard https://console.aws.amazon.com/cloudfront/.
- Select the CloudFront Distribution that needs reconfiguration.

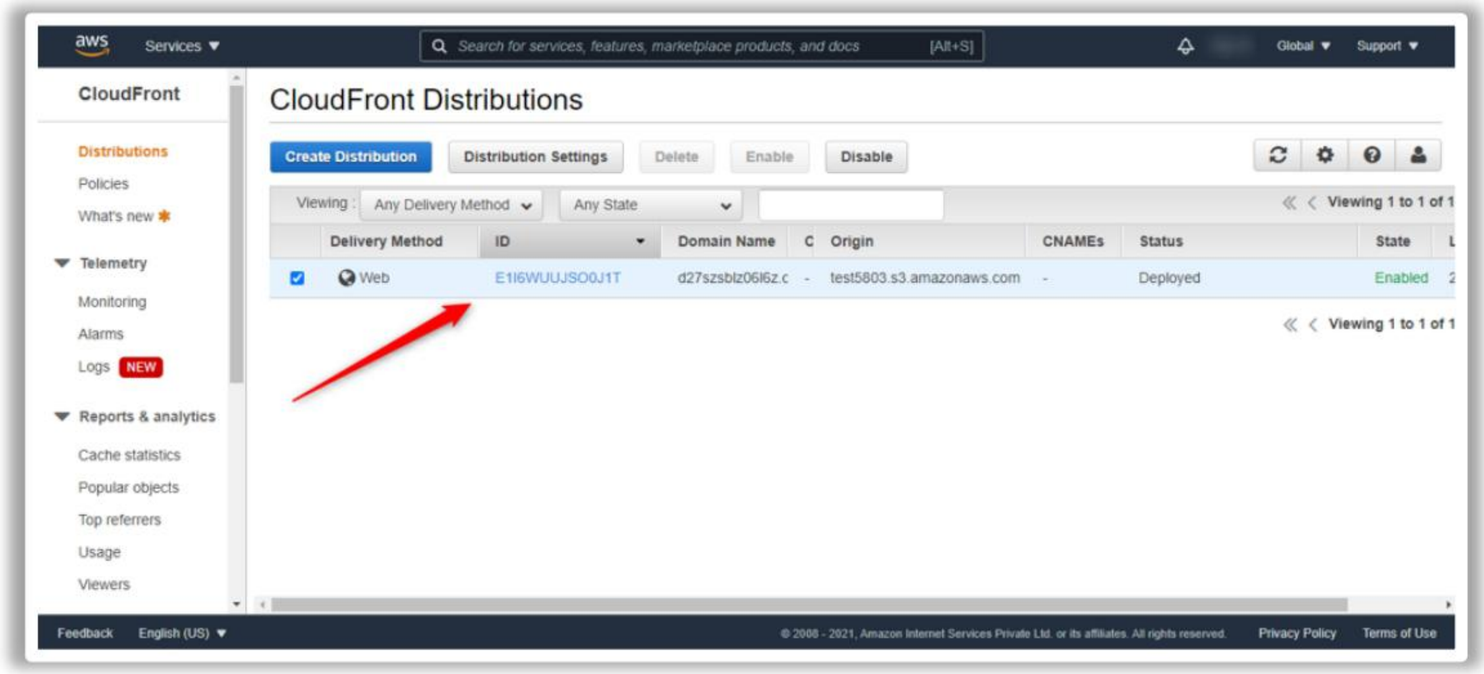


Figure 4.128: CloudFront Distribution Console

- Then, select the Restriction tab and click the **Edit** button.

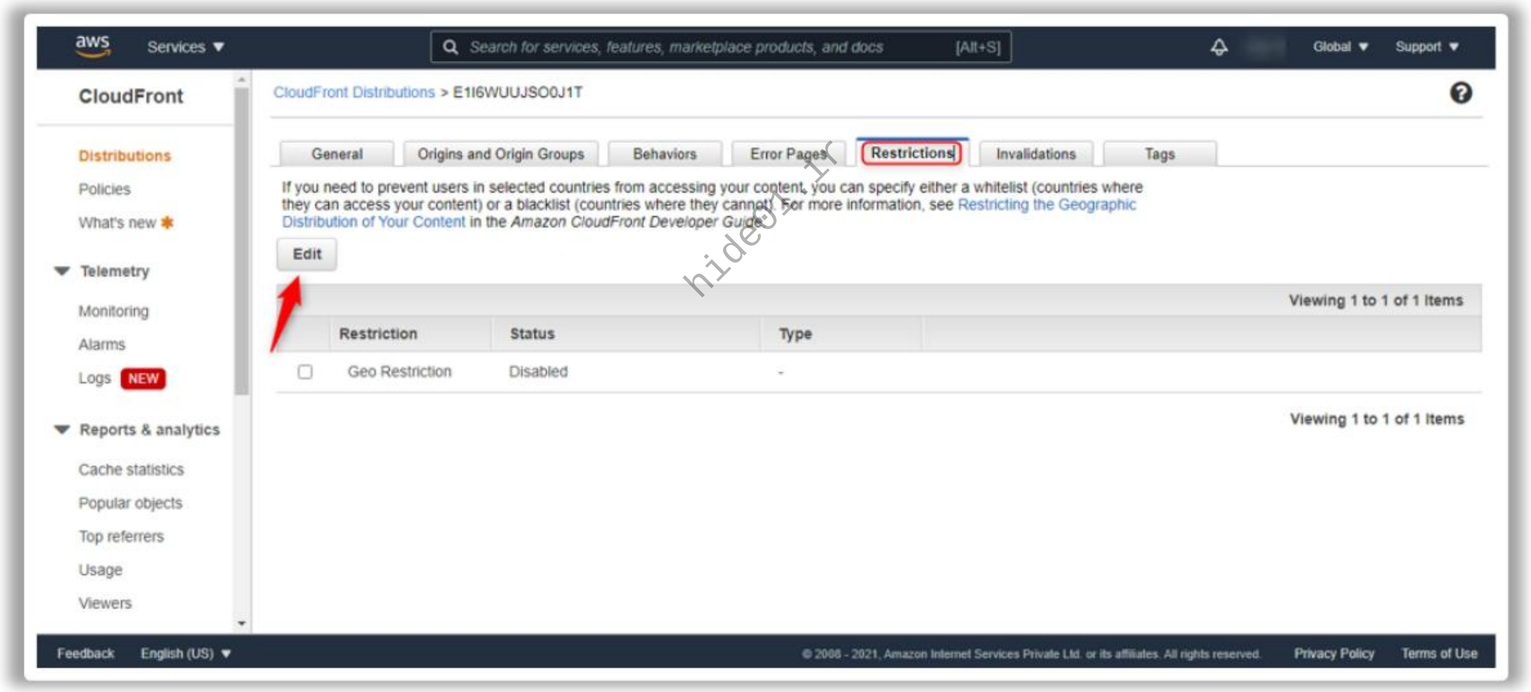


Figure 4.129: Click Edit

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

- Under **Enable Geo-Restrict**, tick the **Yes** radio button. Select the **Restriction Type** based on your requirement; either whitelist or blacklist. Then, choose the name of the countries and click the **Add** button. After adding the countries' name, click **Yes, Edit** button.

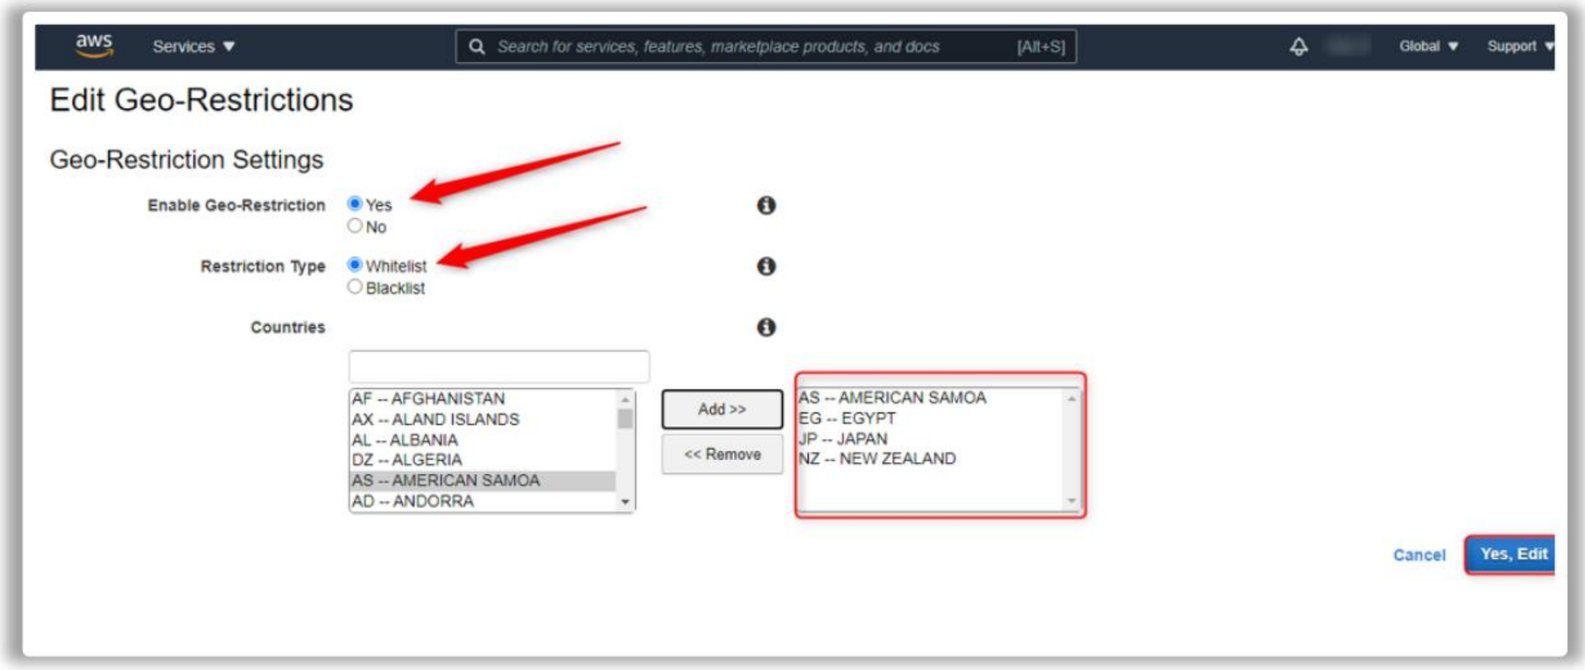


Figure 4.130: Select Yes, Edit

- You will observe that **Geo Restriction** is **enabled** for the selected CloudFront distribution.

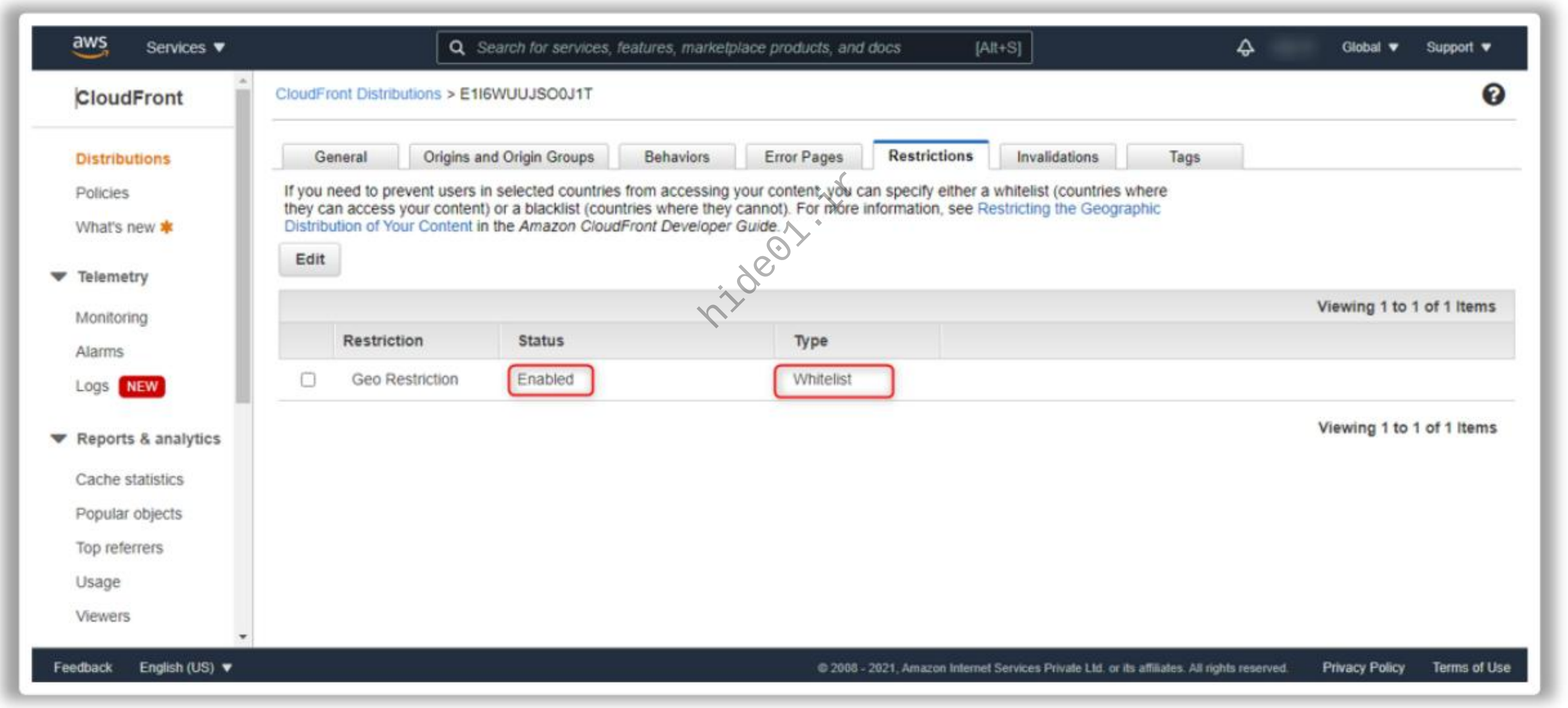
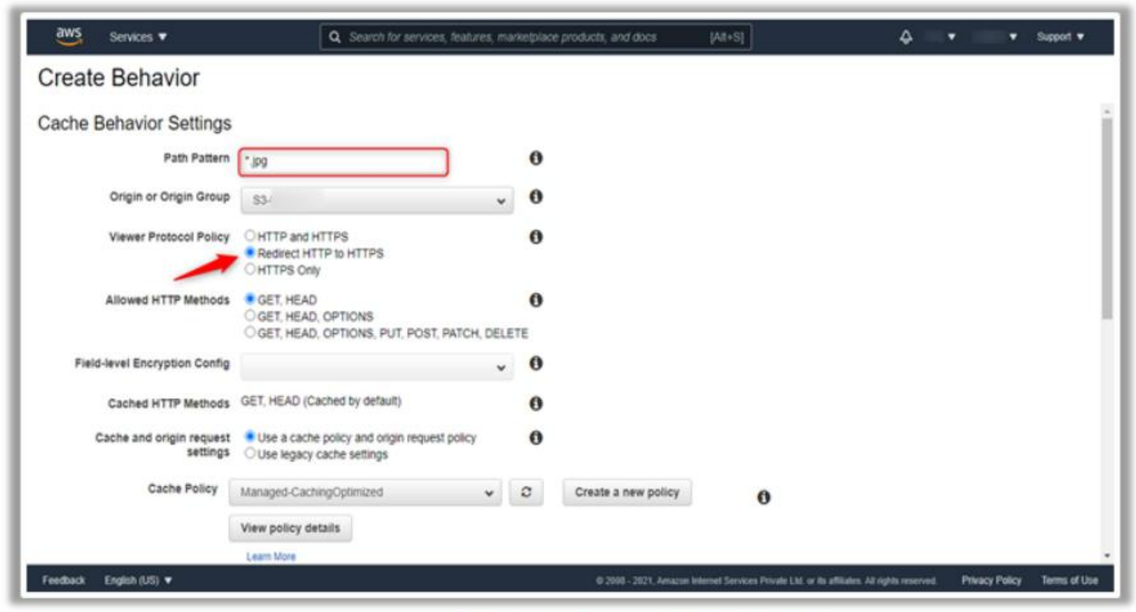


Figure 4.131: Geo Restriction is enabled

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

Configure Viewer Protocol Policy of CloudFront Distribution

- To enhance the security of the web application content and to safeguard the data in motion against interception attacks, the cloud security professional should **encrypt** the communication between Amazon CloudFront CDN distribution and end users utilizing HTTPS



Configure Viewer Protocol Policy

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

Configure Viewer Protocol Policy of CloudFront Distribution

To enhance the security of the web application content, the cloud security professional should encrypt the communication between Amazon CloudFront CDN distribution and end-users utilizing HTTPS. For encrypting the data-in-motion, the administrator should configure the viewer protocol policy either to redirect HTTP to HTTPS or HTTPS only based on the requirement. Utilizing HTTPS for CloudFront CDN distribution assures that the encrypting traffic between the CloudFront distribution and the end-users cannot be decrypted by the adversary, even though the adversary intercepts the sent packets in the CDN distribution network.

Source: https://aws.amazon.com

To configure viewer protocol policy of CloudFront distribution, perform the following steps:

- Sign-in to AWS Management Console.

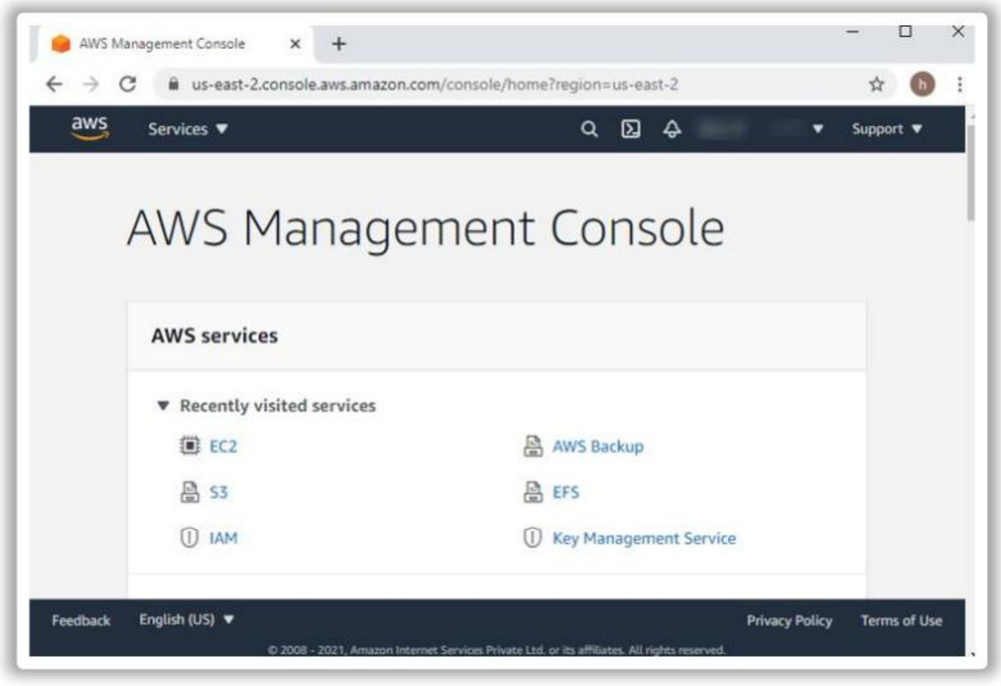


Figure 4.132: AWS Management Console

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

- Click on the following link to open CloudFront dashboard https://console.aws.amazon.com/cloudfront/.
- Select **Distribution** on the left pane to view the existing distributions; then select the distribution on which you want to configure the viewer protocol policy and click on the **Distribution Settings** button.

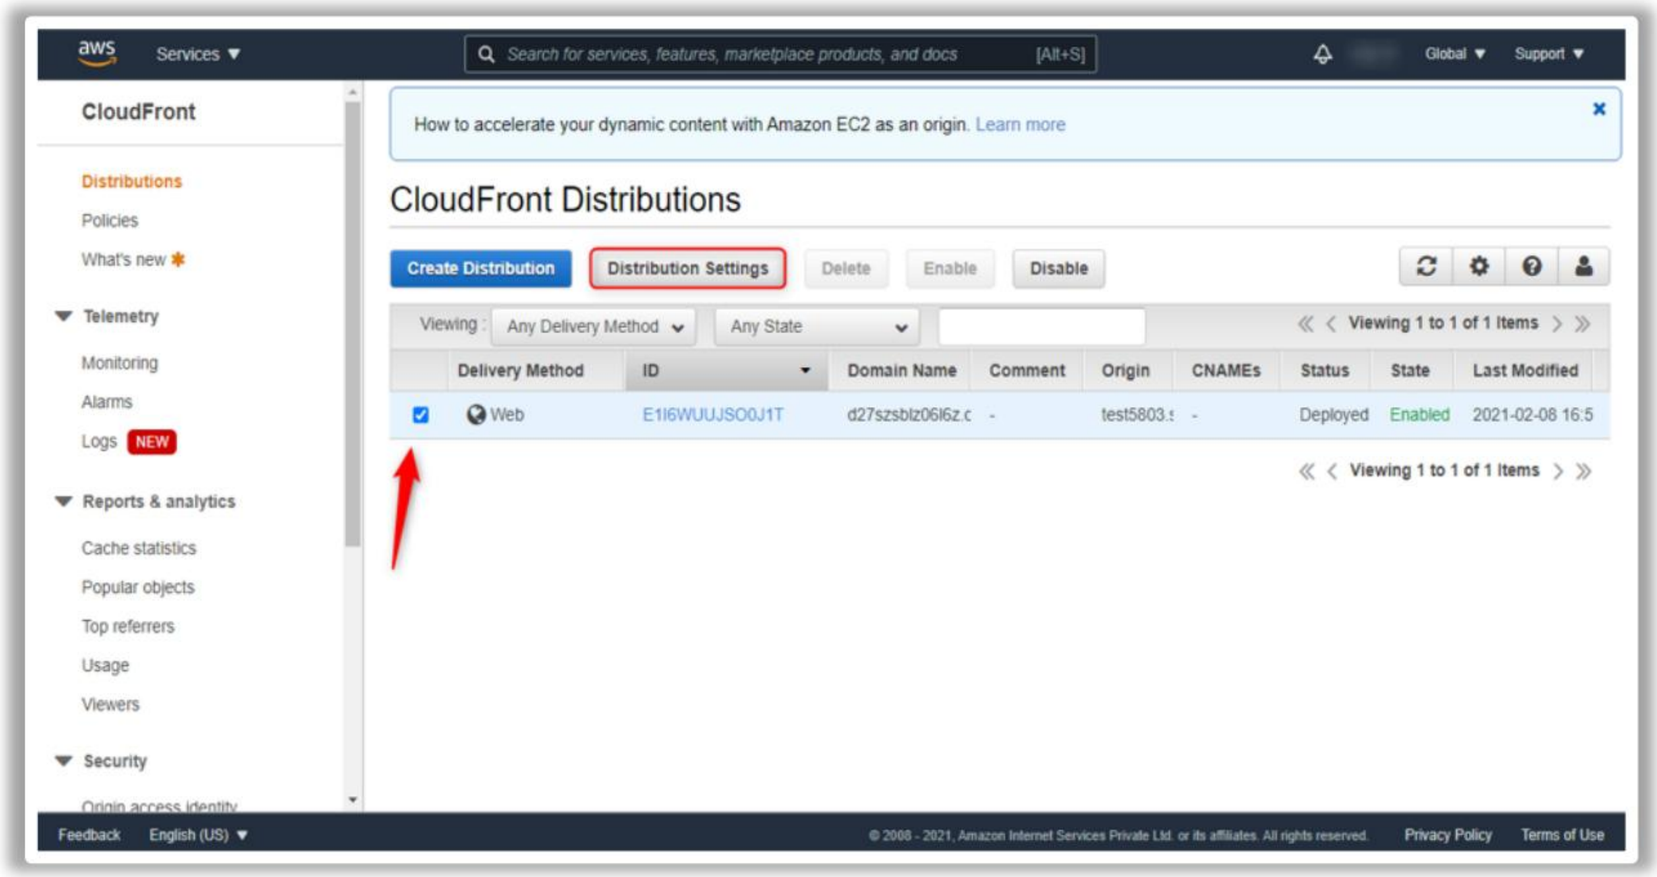


Figure 4.133: CloudFront Distribution Console

- Click on **Behaviors** tab, then click on **Create Behavior** button.

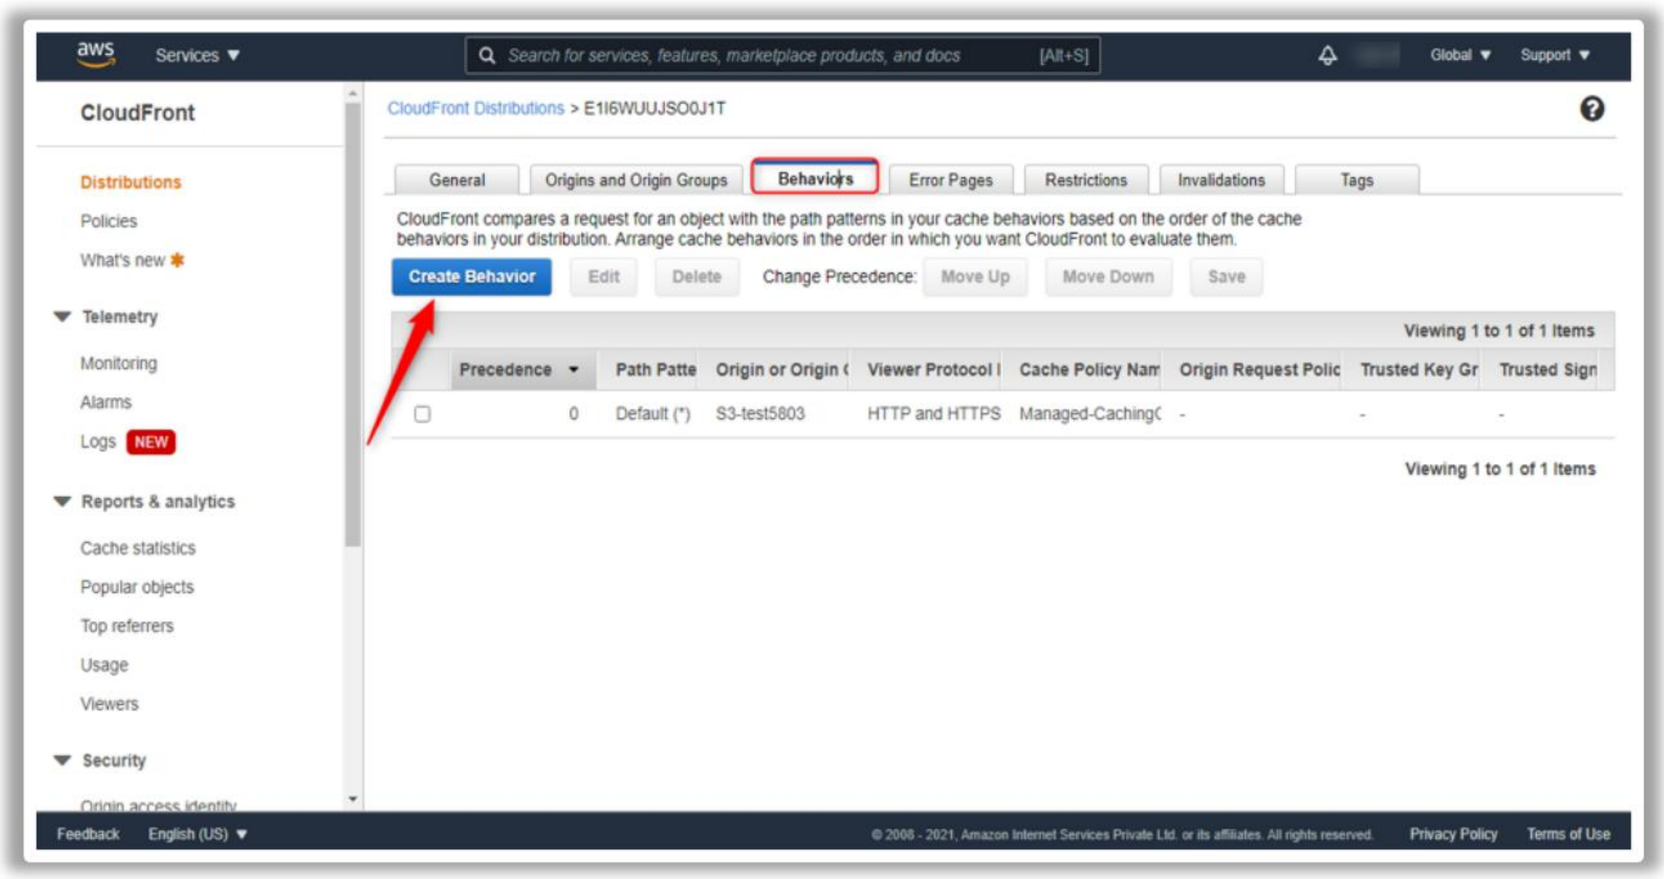


Figure 4.134: Click Create Behavior

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

- In the **Create Behavior** page, enter the Path Pattern as per your requirements and then under **View Protocol Policy** either select **Restrict HTTP to HTTPS** or **HTTPS only** radio button to encrypt the data-in-motion, and then click the **Create** button.

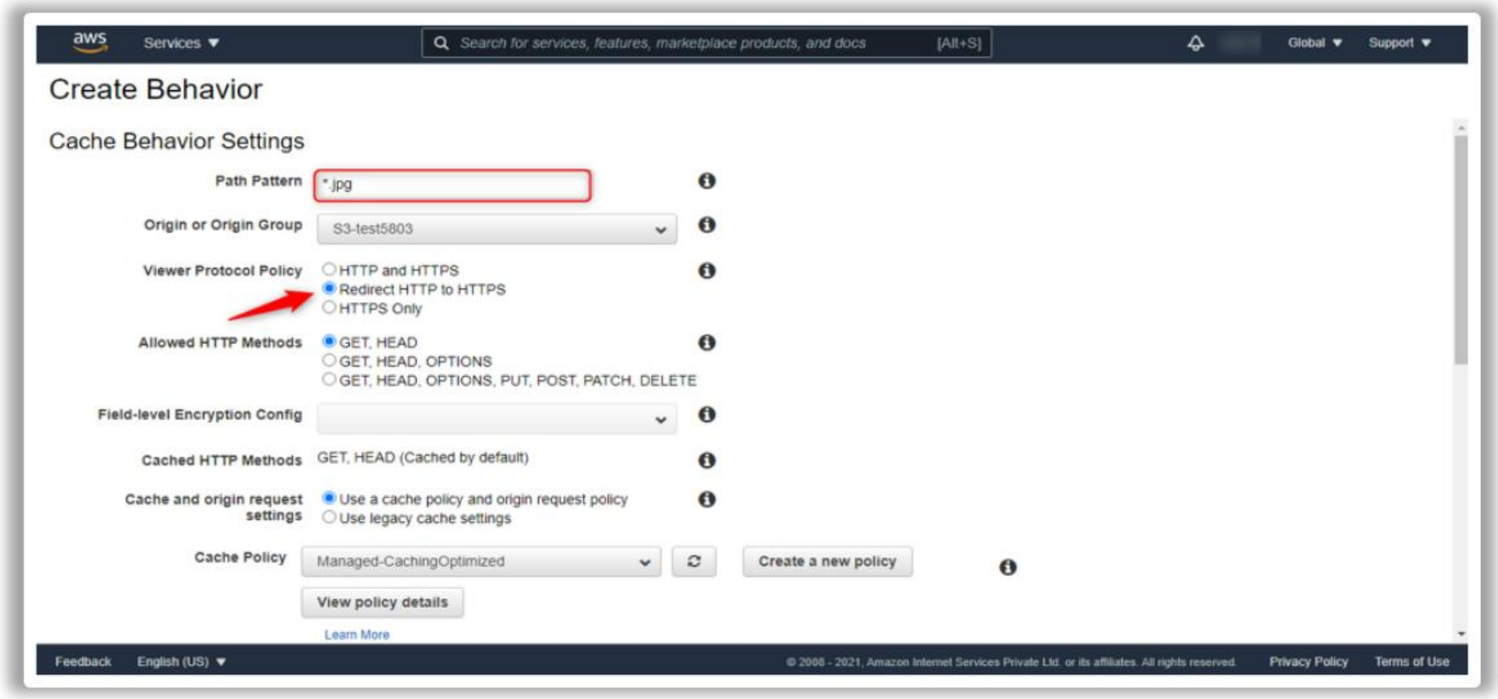


Figure 4.135: Create Behavior Console

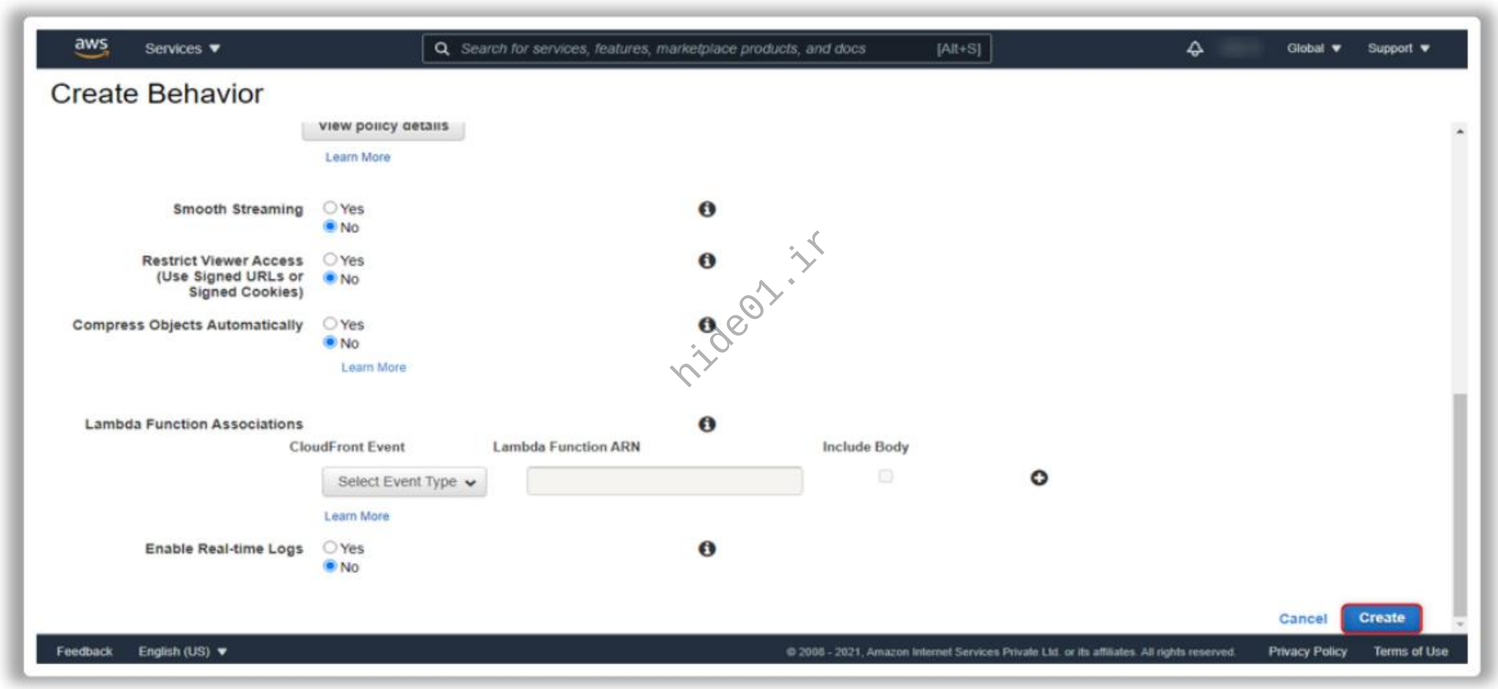


Figure 4.136: Click Create

- Viewer protocol policy is successfully configured to the selected CloudFront distribution.

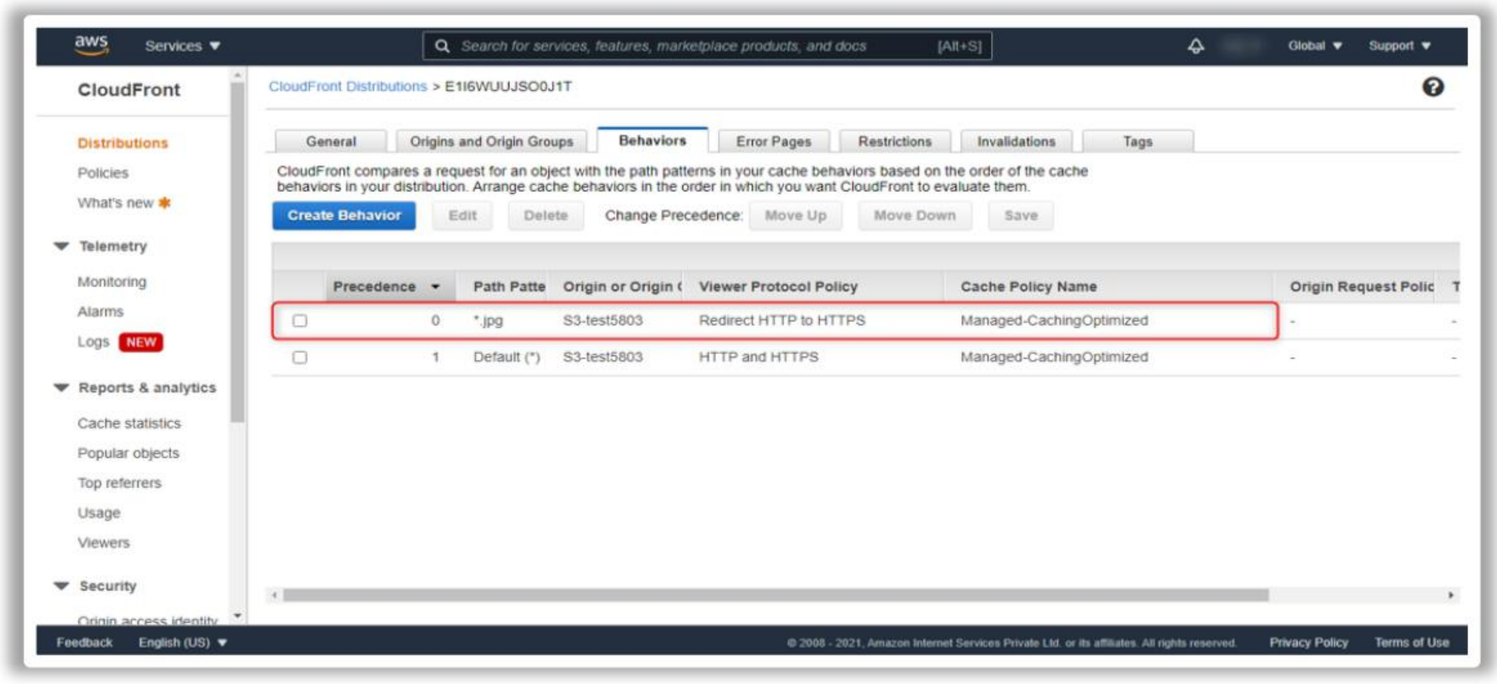
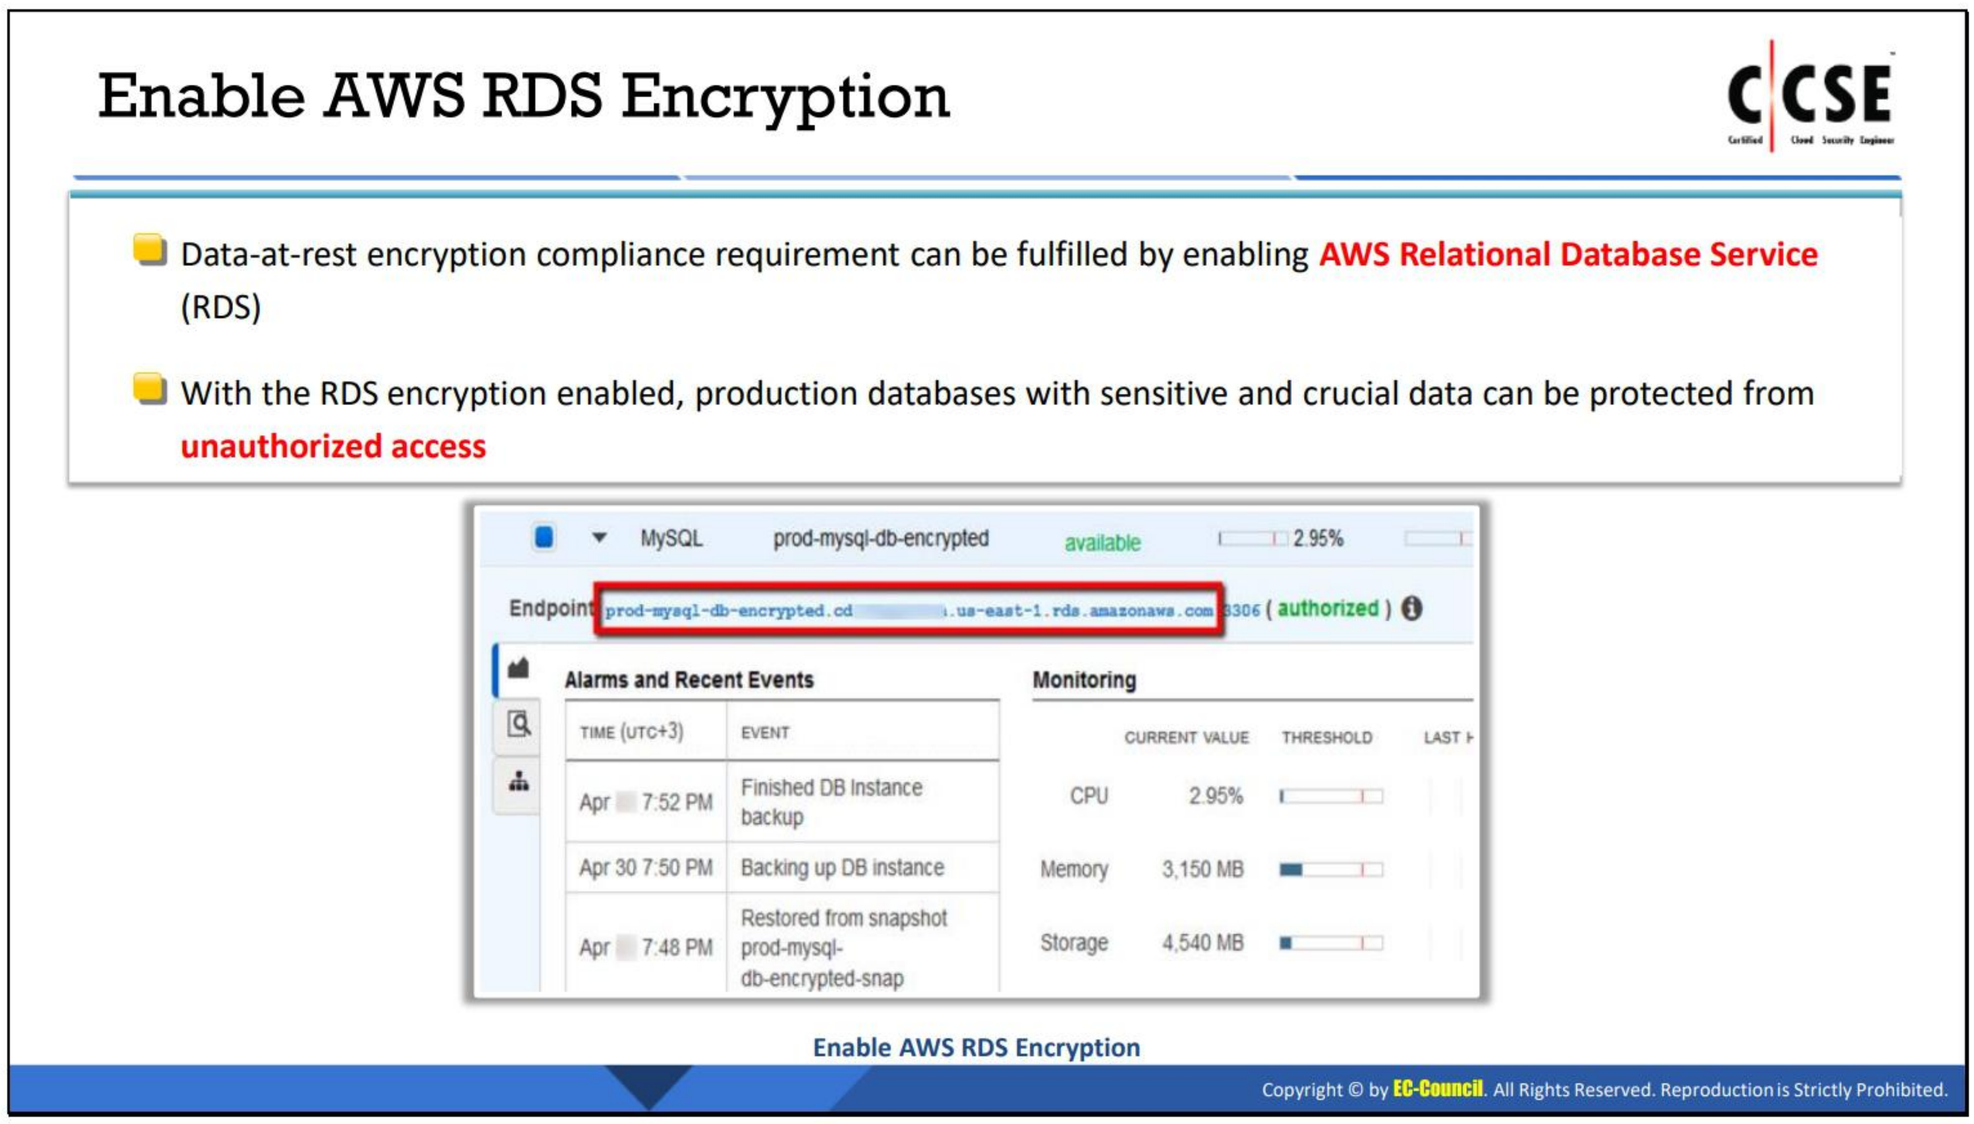


Figure 4.137: List of CloudFront Distribution

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



Enable AWS Relational Database Service (RDS) Encryption

Data-at-rest encryption compliance requirements can be fulfilled by enabling AWS RDS. The encryption and decryption in RDS do not need any additional action from the user or application. It helps in complying with different standards such as PCI-DSS, HIPAA, GDPR, etc. With the RDS encryption enabled, production databases with sensitive and crucial data can be protected from unauthorized access. This RDS encryption encrypts the data stored on the instance underlying storage such as read replicas, backups, and snapshots. It uses the AES-256 algorithm that is managed by AWS KMS. AWS RDS encryption is not available for all types of database instances.

Source: https://aws.amazon.com

To enable RDS encryption, perform the following steps:

- Sign-in to the AWS Management Console.
- Navigate to the **RDS dashboard**; under this section click on **Instances**.
- Now select the RDS database instance in which you want to enable encryption.

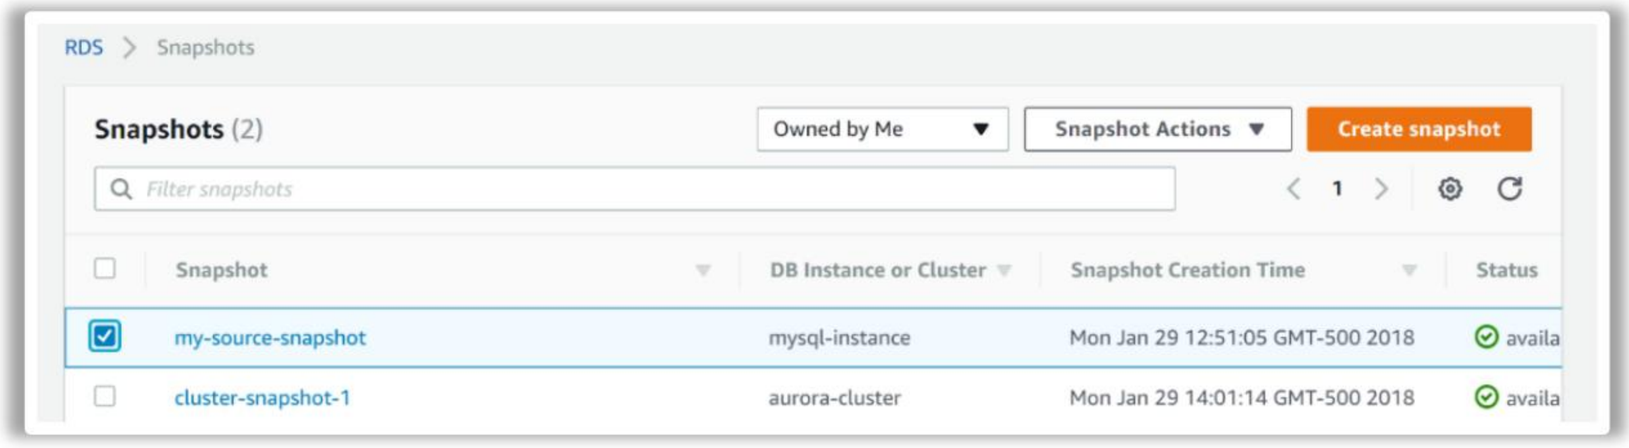


Figure 4.138: Snapshots Console

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

- From the dashboard top menu, click on **Instance Actions** and select **Take Snapshot**.
- In the **Take DB Snapshot** panel, enter a name in the input of **Snapshot Name** and click on **Take Snapshot**.
- Now select the newly created snapshot, click on **Copy Snapshot** from the top menu of the dashboard.
- Next on the **Make copy of DB Snapshot page**, enter the name in **New DB Snapshot Identifier**; make sure that **Copy Tags** have the same source snapshot tags.

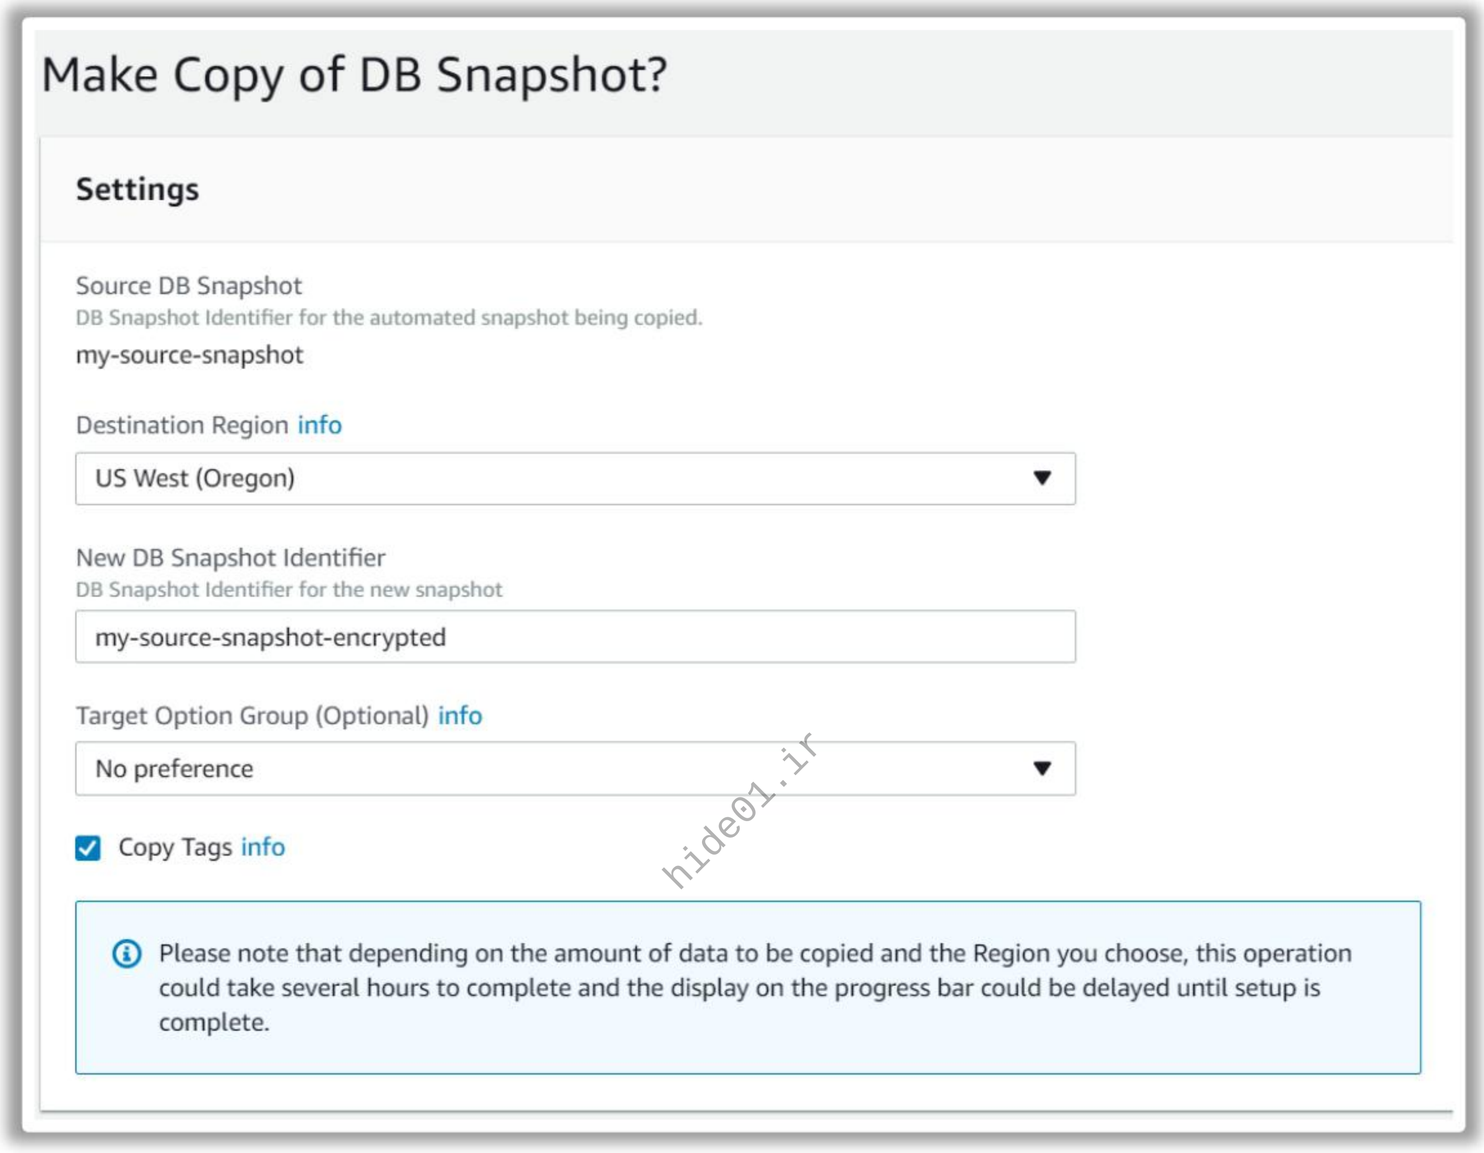


Figure 4.139: Click Copy Tags

- From the dropdown list of **Enable Encryption**, select **Yes**. Select either the default AWS encryption key or custom key (key ARN is required) that you have from the **Master Key** dropdown list.

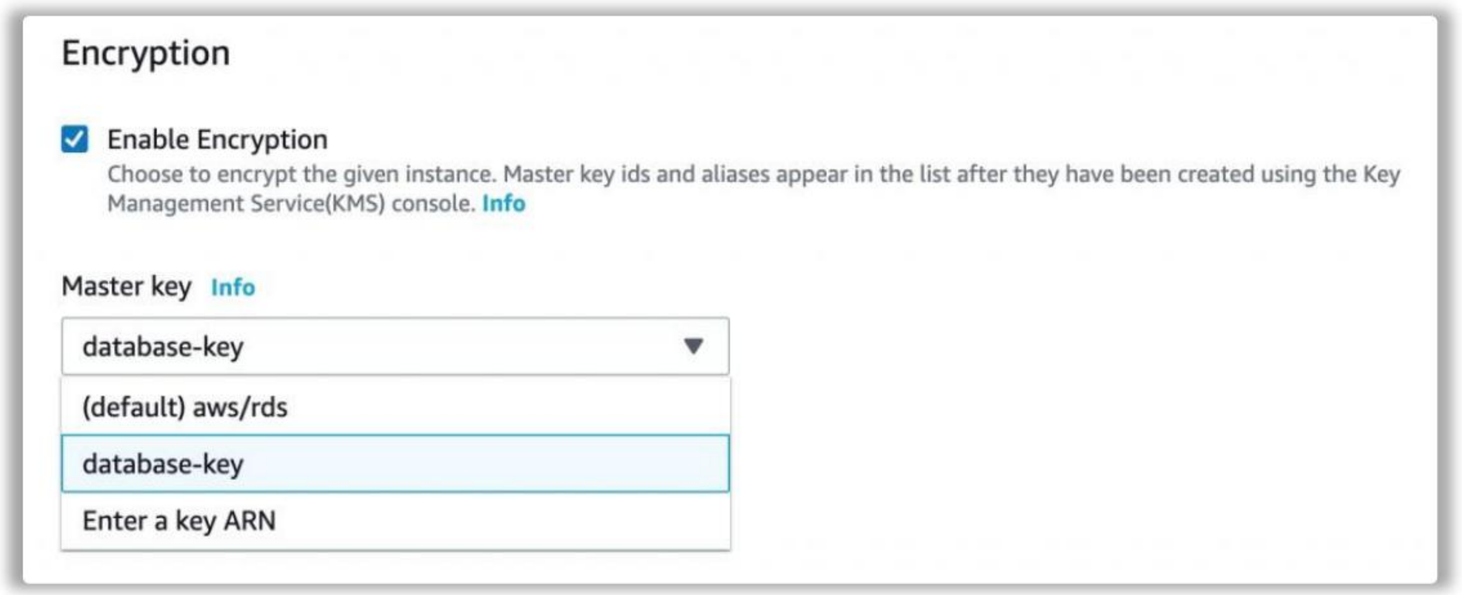


Figure 4.140: Click Enable Encryption

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

- Now click on **Copy Snapshot**; the selected instance snapshot copy will be encrypted.
- Select the new encrypted snapshot copy; click on **Restore Snapshot** from the top menu of the dashboard. It restores the encrypted snapshot to the new database instance.
- In the **Restore DB Instance** page, give a unique name in the **DB Instance Identifier.**
- After reviewing the configuration details of the instance, click on **Restore DB Instance**.
- The new instance provision process is completed. Give it a moment to make it available.
- To refer to an endpoint of the (new) encrypted database instance, update the application configuration. As soon as the database endpoint is changed to the application level. The unencrypted instance can be removed.

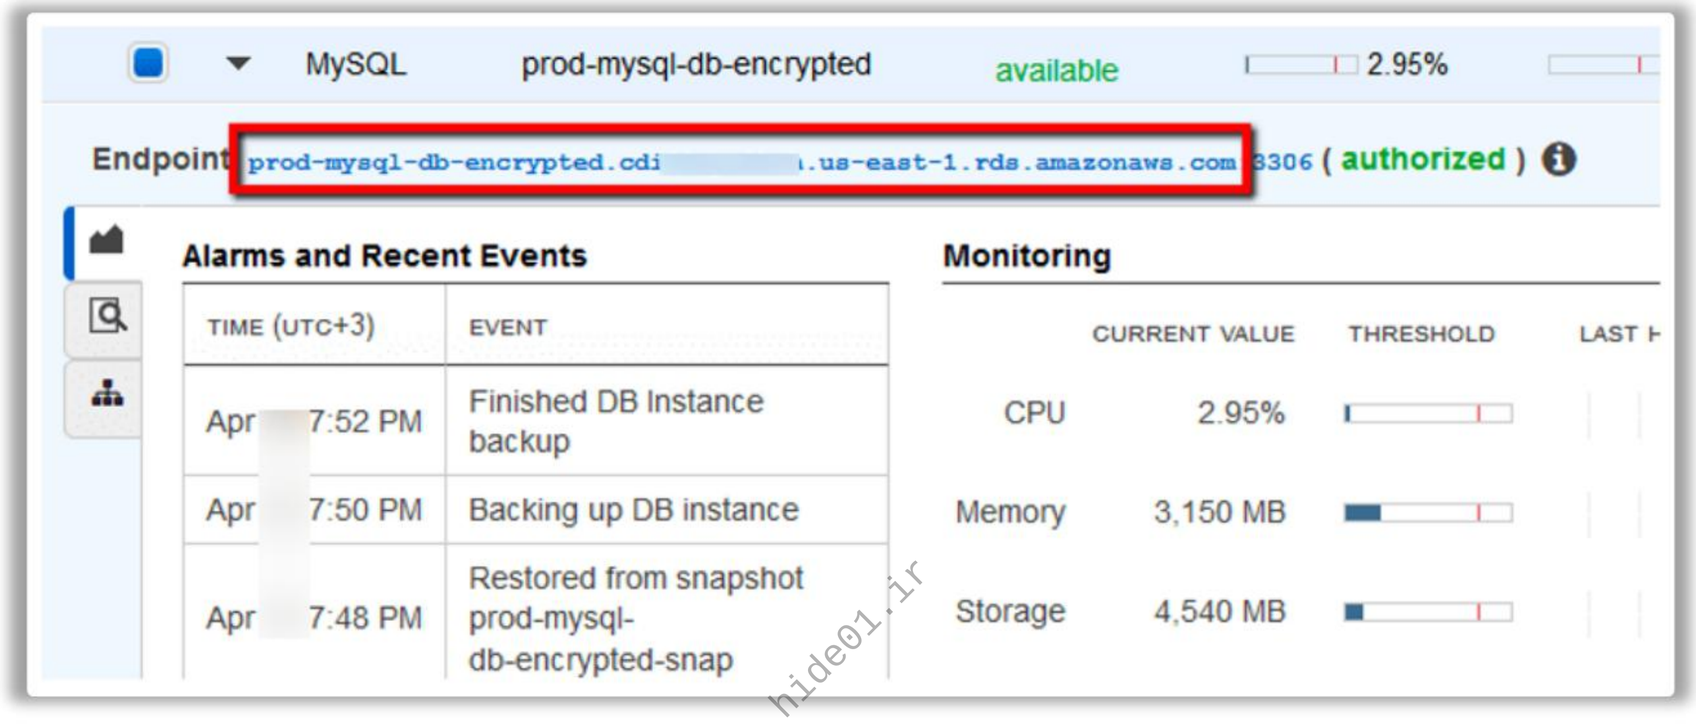
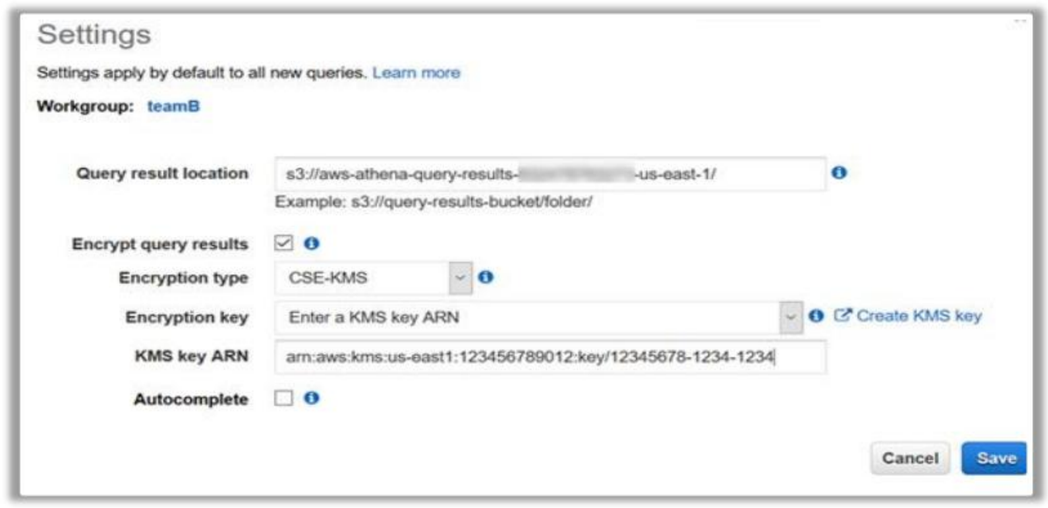


Figure 4.141: MySQL Console

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

Enable Encryption for AWS Athena Query Results

- Athena is an **interactive query service** provided by AWS. It allows the usage of standard SQL for **data analysis** in Amazon S3. By default, it is protected by SSL/TLS, but query results at rest are not encrypted
- To meet the compliance requirement, it is recommended to **enable encryption** for AWS Anthea query results



Enabling encryption for AWS Anthea query results

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

Enable Encryption for AWS Athena Query Results

Athena is an interactive query service provided by AWS. It allows the usage of standard SQL for data analysis in Amazon S3. By default, it is protected by SSL/TLS; however, query results at rest are not encrypted. It is highly recommended to enable encryption for AWS Anthea query results. Enabling encryption allows organizations to meet the compliance requirements. AWS Athena supports three types of S3 encryptions, that is, SSE with an Amazon S3-managed Key (SSE-S3), SSE with an AWS Key Management Service customer-managed key (SSE-KMS), and client-side encryption (CSE) with an AWS KMS customer-managed key (CSE-KMS).

Source: https://aws.amazon.com

To enable encryption for AWS Athena Query results, perform the following steps:

- Sign-in to AWS Management Console.
- Go to Athena dashboard and choose **Settings**.



Figure 4.142: Select Settings

- In Query result location, leave it to default or enter any value. The query results are stored here; this is an Amazon S3 staging directory.
- Select the checkbox of **Encrypt query results**.
- In **Encryption type**, select **CSE-KMS**, **SSE-KMS**, or **SSE-S3** from the dropdown list. If the account has access to any existing AWS KMS customer-managed key (CMK), select its alias or select **Enter a KMS key ARN** from the dropdown list.

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

- If there is no access to the existing AWS KMS CMK; select **Create KMS key**.
- After key creation, get back to the Athena console; specify the key by alias or ARN as in the previous step.
- Click on **Save**.

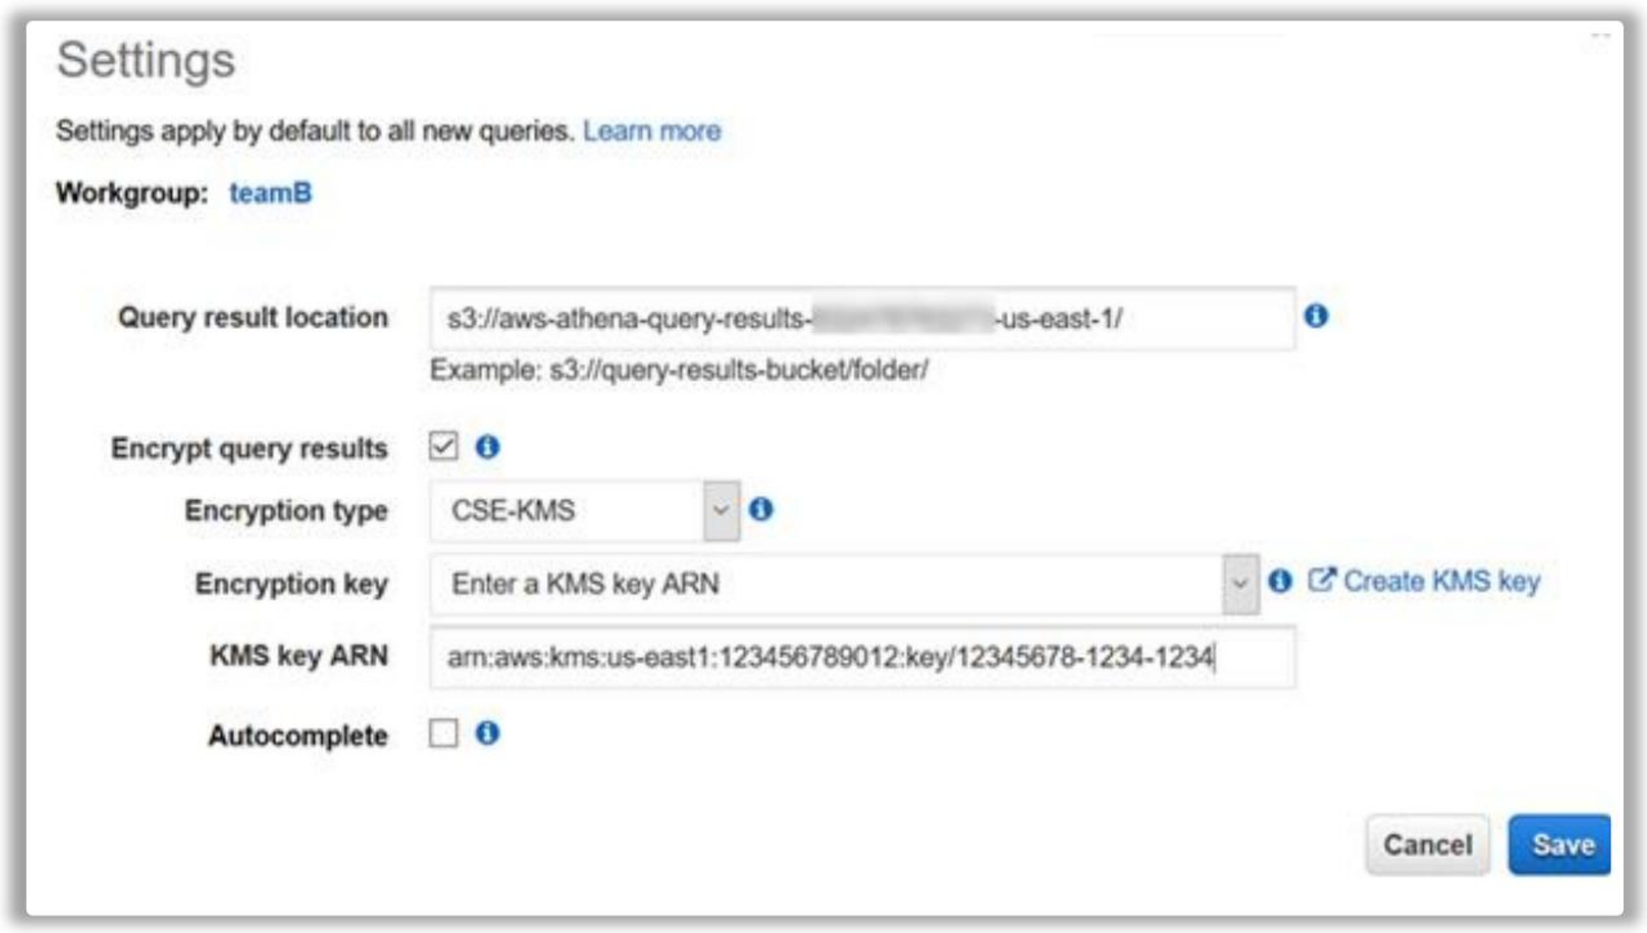
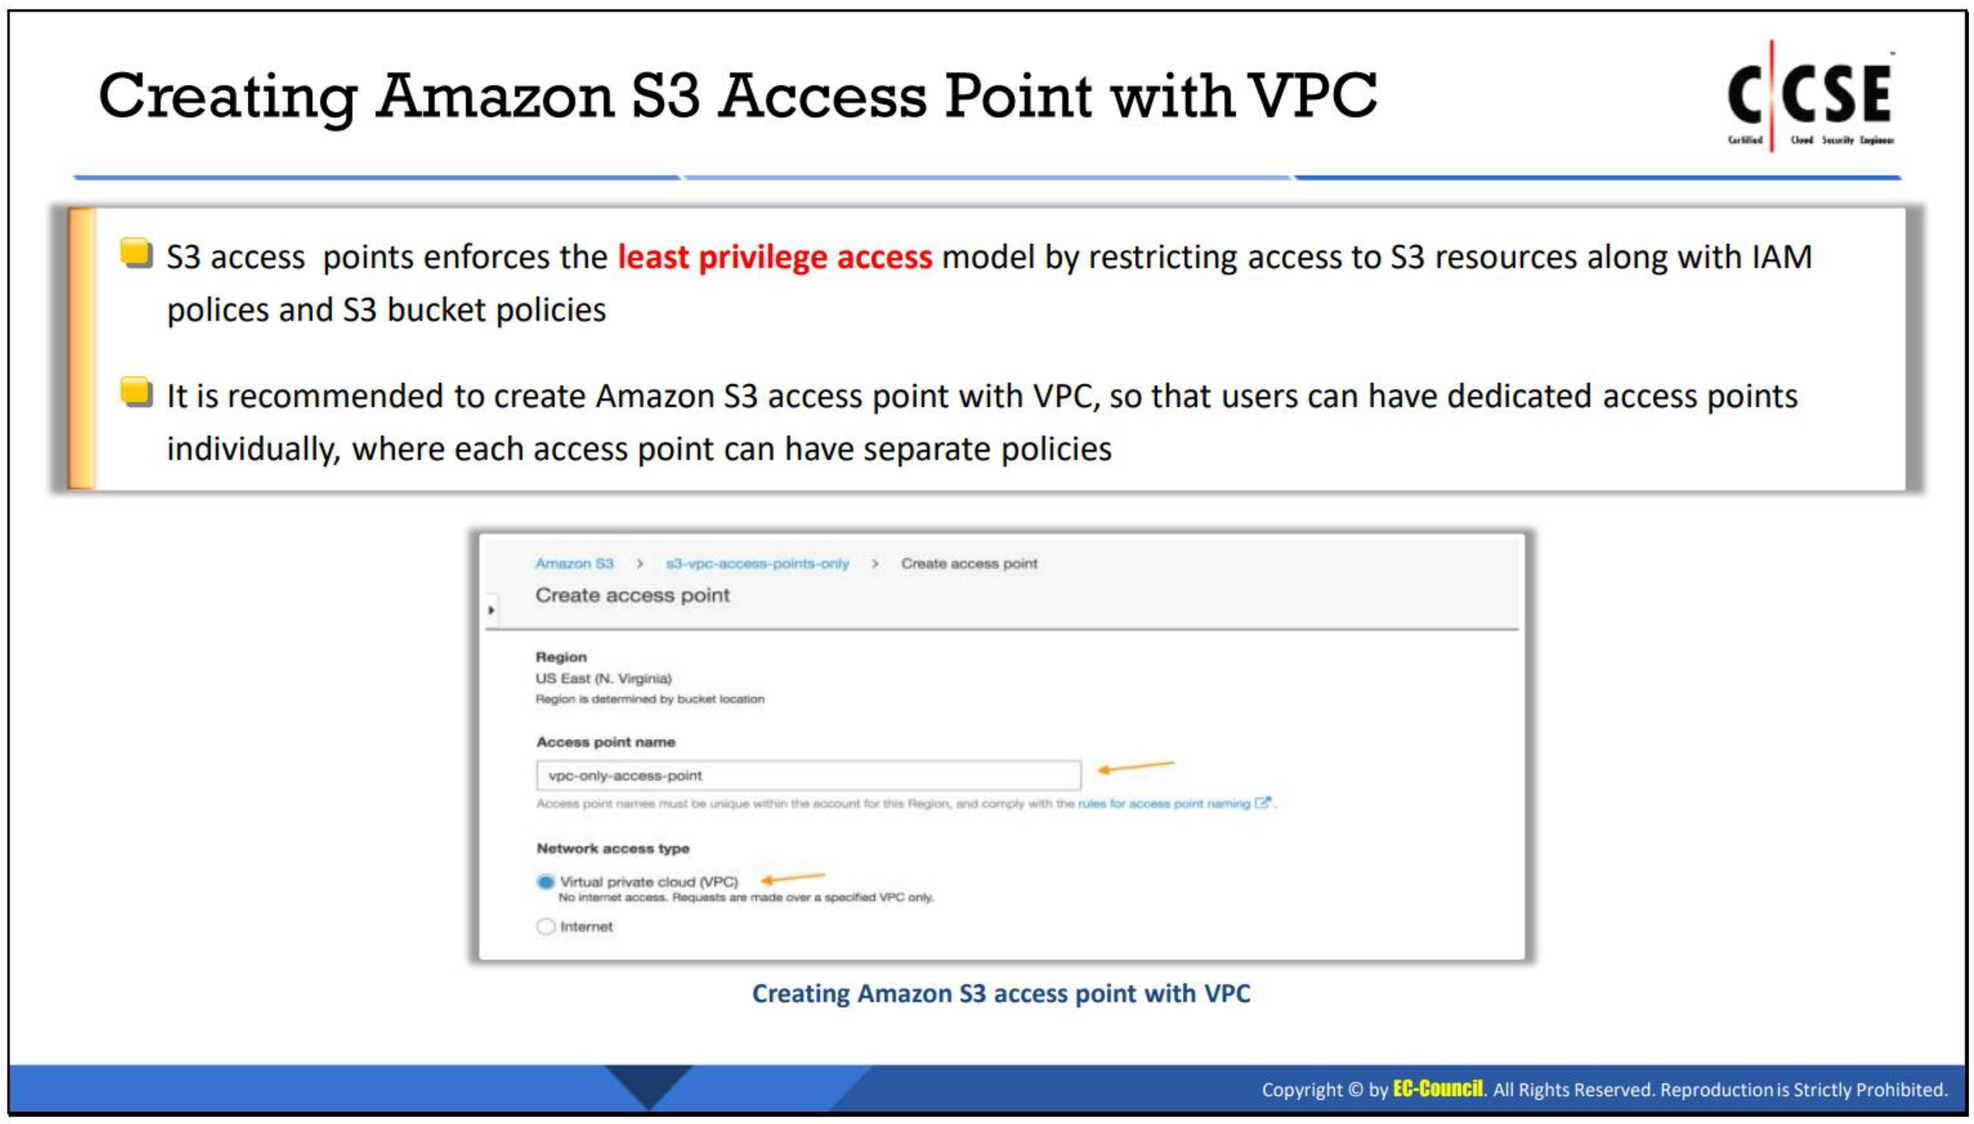


Figure 4.143: Settings Console

hade01.ir

- AWS Athena query results at rest are successfully encrypted.

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



Creating Amazon S3 access point with VPC

Creating Amazon S3 Access Point with Virtual Private Cloud (VPC)

The management of S3 buckets is easy; however, it may become complex with the addition of users and data. The access policies to these users could become even more complicated. To resolve this issue AWS has provided access points on S3. It is very useful when there are large amounts of data in the S3 bucket that needs to be accessed by different users. If the requirement for data security is high, an access point can be restricted to a particular VPC.

S3 access points enforce the least privilege access model by restricting access to S3 resources along with IAM polices and S3 bucket policies. The cloud security professional can restrict a particular VPCs or range of known IP addresses with bucket polices and access point policies. It is recommended to create Amazon S3 access point with VPC so that users can have dedicated access points individually, where each access point can have separate policies.

Source: https://aws.amazon.com

To create Amazon S3 access point with VPC, perform the following steps:

- Sign-in to Amazon AWS Management Console.

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

- Navigate to **Access Point**.

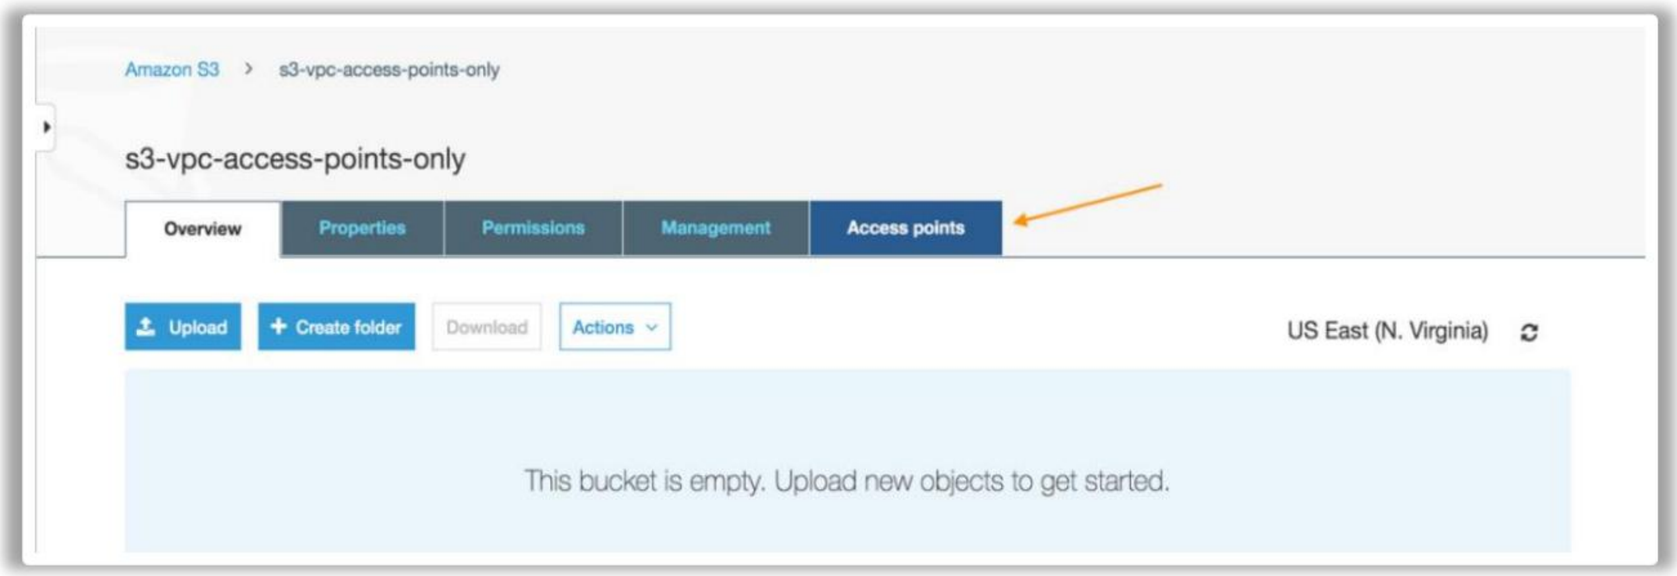


Figure 4.144: Select Access Point

- Click on **Create access point**.

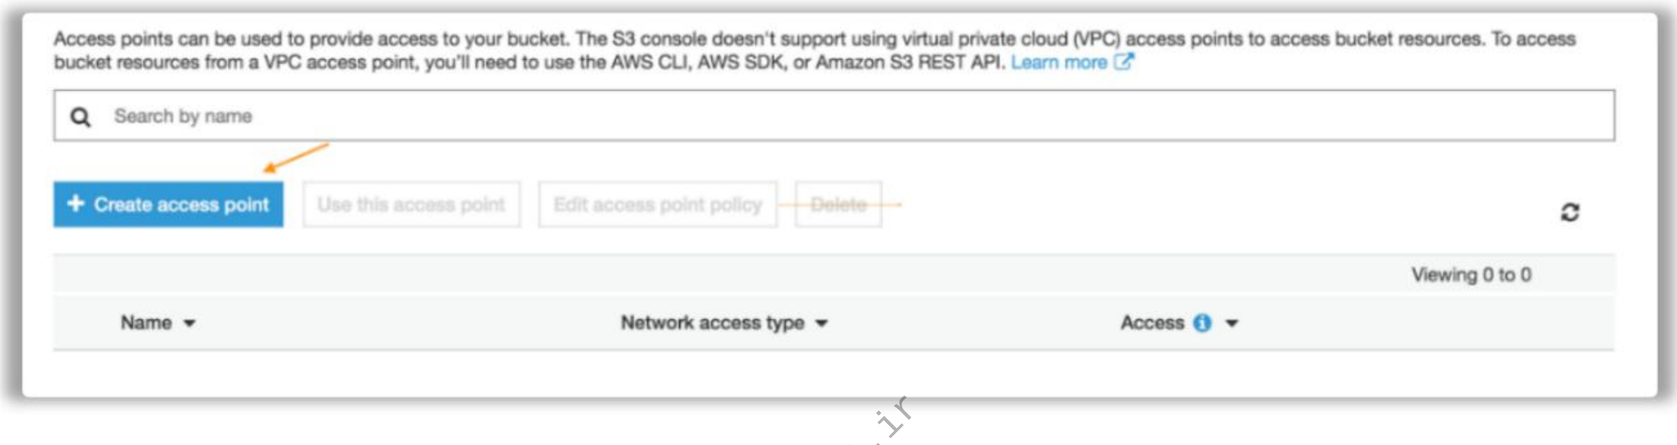


Figure 4.145: Click Create Access Point

hide01.ir

- Enter **Access point name**.
- In the **Network access type**, select **Virtual private cloud (VPC)**.

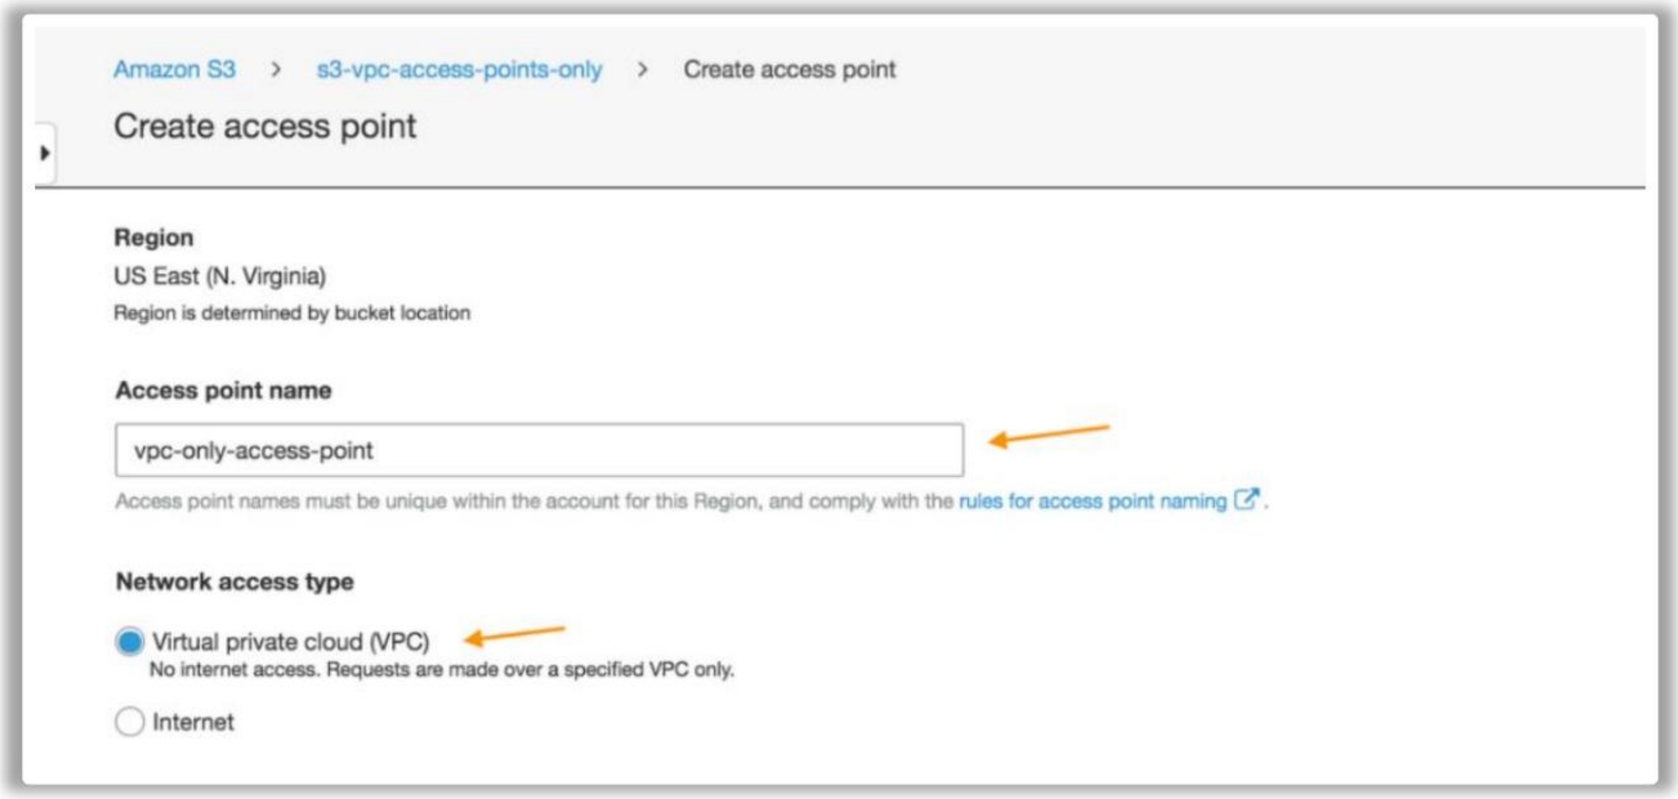


Figure 4.146: Create access point Console

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

- Select the checkbox of **Block all public access**.

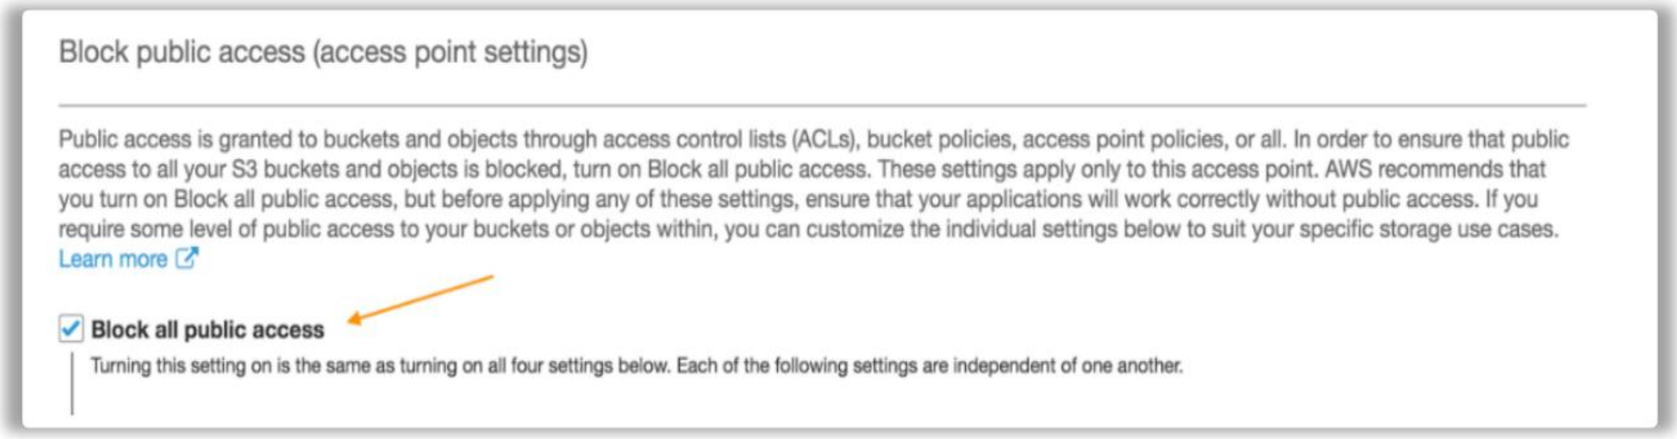


Figure 4.147: Select Block all public access

- Click on **Create access point**.

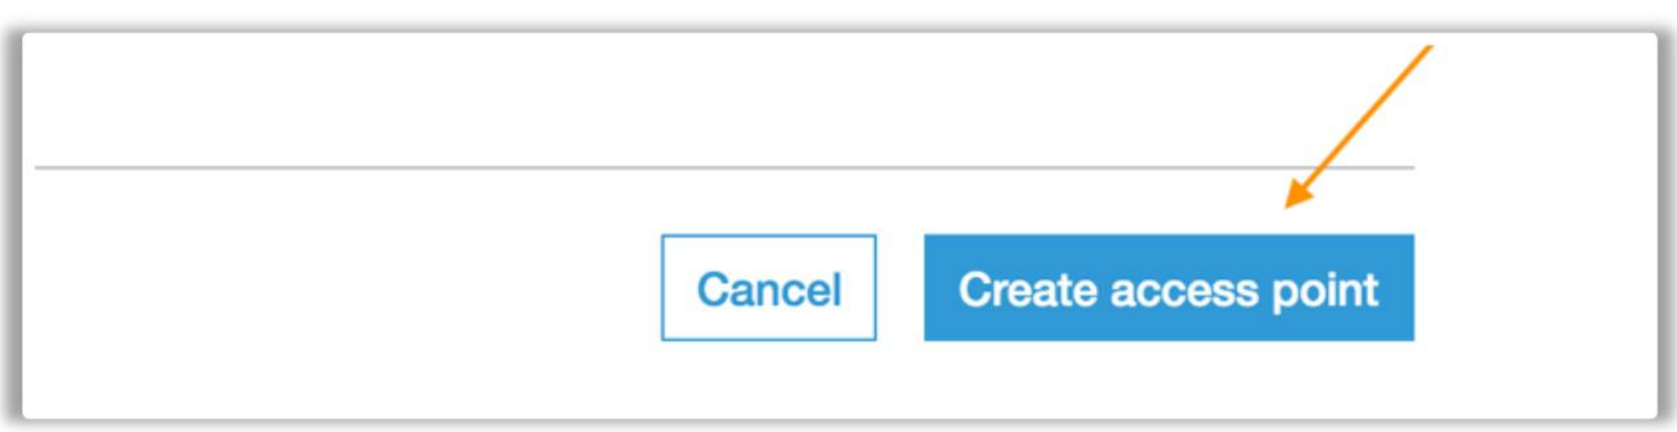


Figure 4.148: Click Create access point

hidoo1.ir

- Return to access point, note the access point ARN.

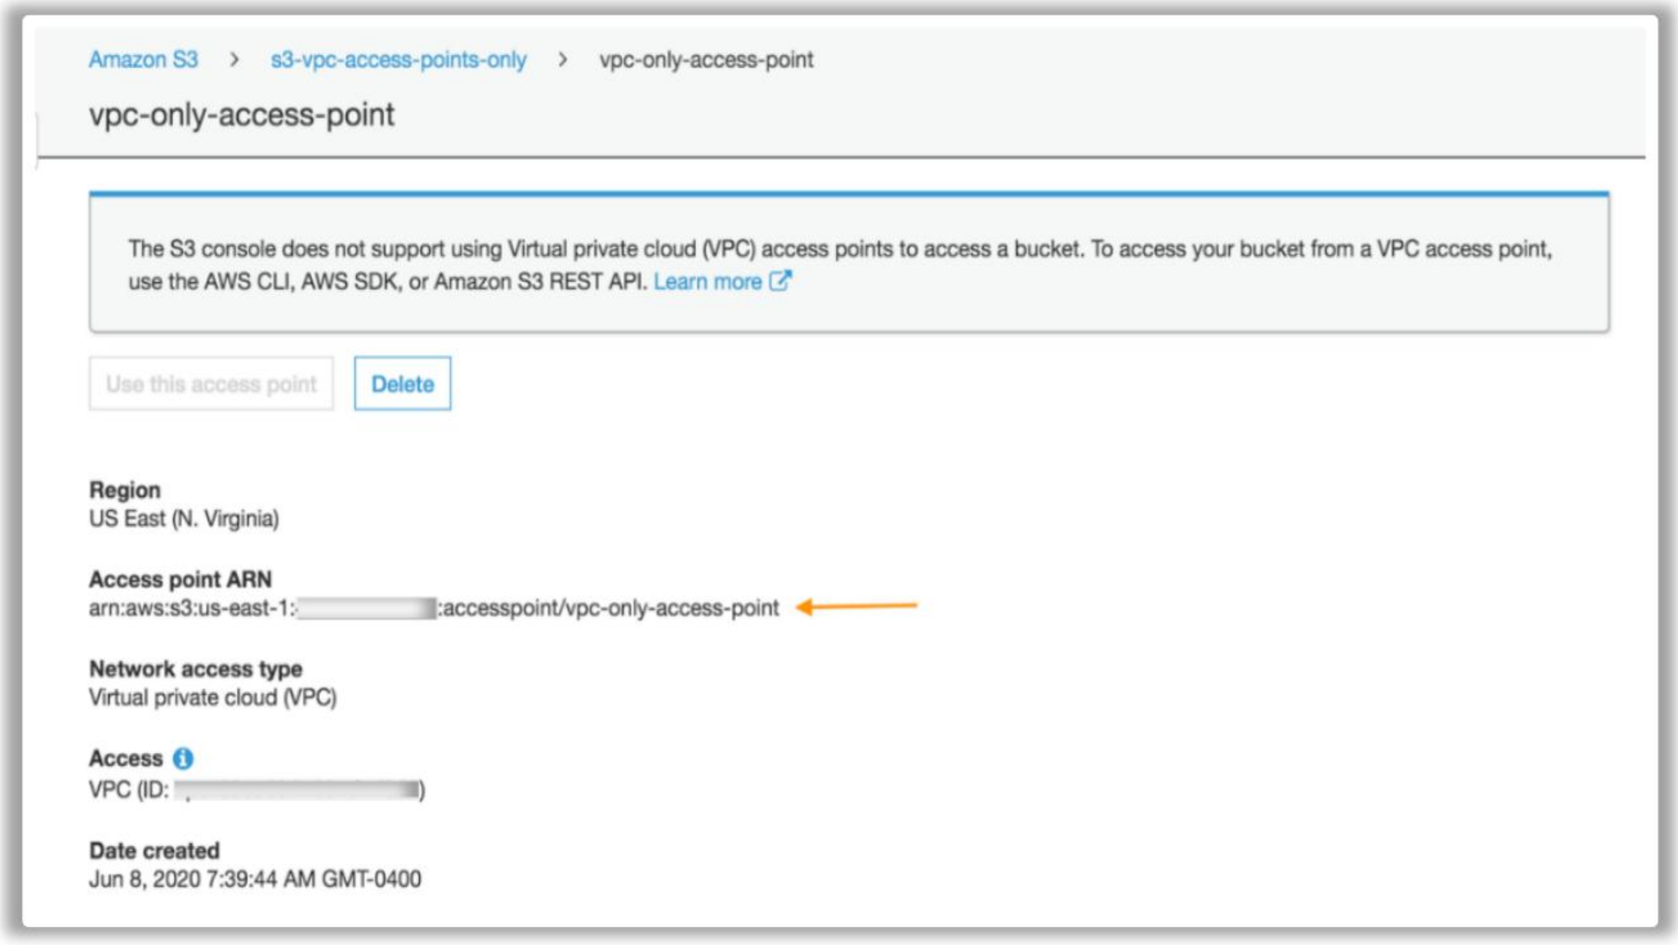


Figure 4.149: Note ARN

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

Amazon S3 Gateway Endpoint Creation in VPC

To ensure that traffic is routed via the S3 VPC endpoint, the Amazon S3 VPC gateway endpoint is created.

- Navigate to the Amazon VPC console.
- Locate **Endpoints** and click on it.
- Click **Create Endpoint** on the right-side panel.

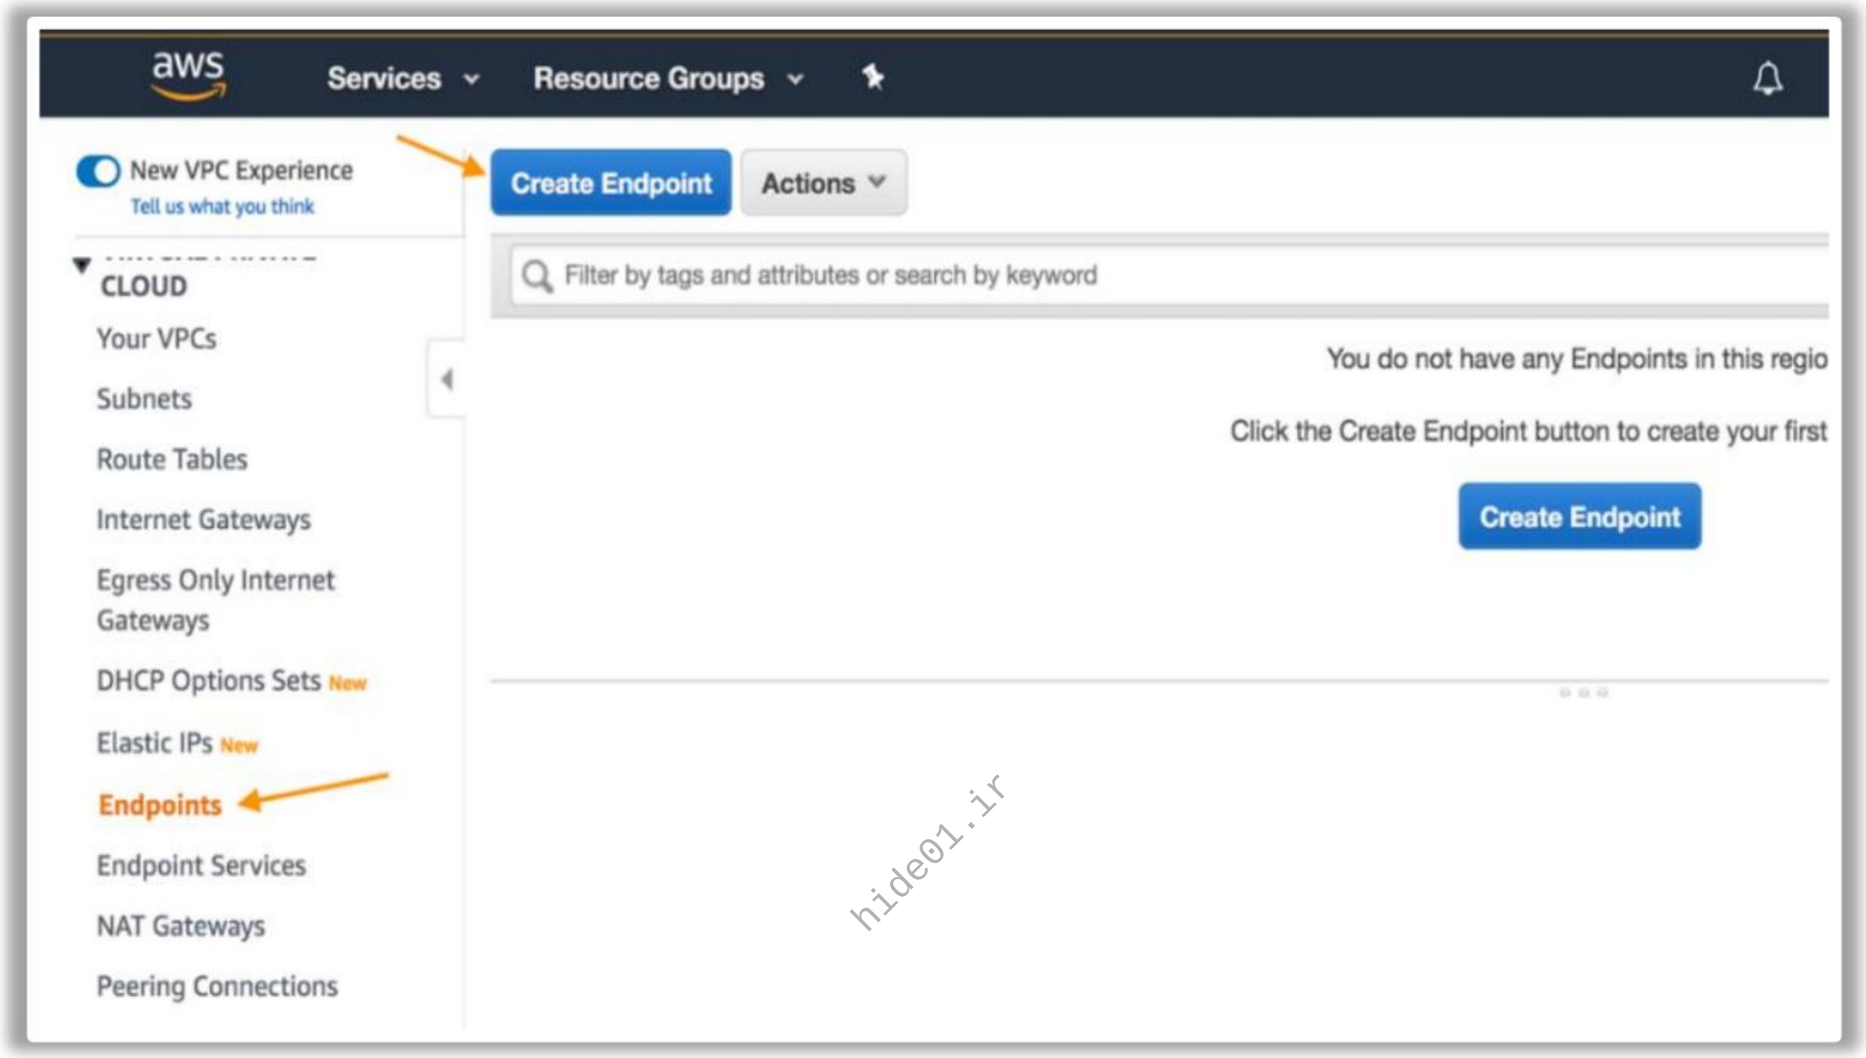


Figure 4.150: Select Create Endpoint

- Select the endpoints for S3 from the search bar.

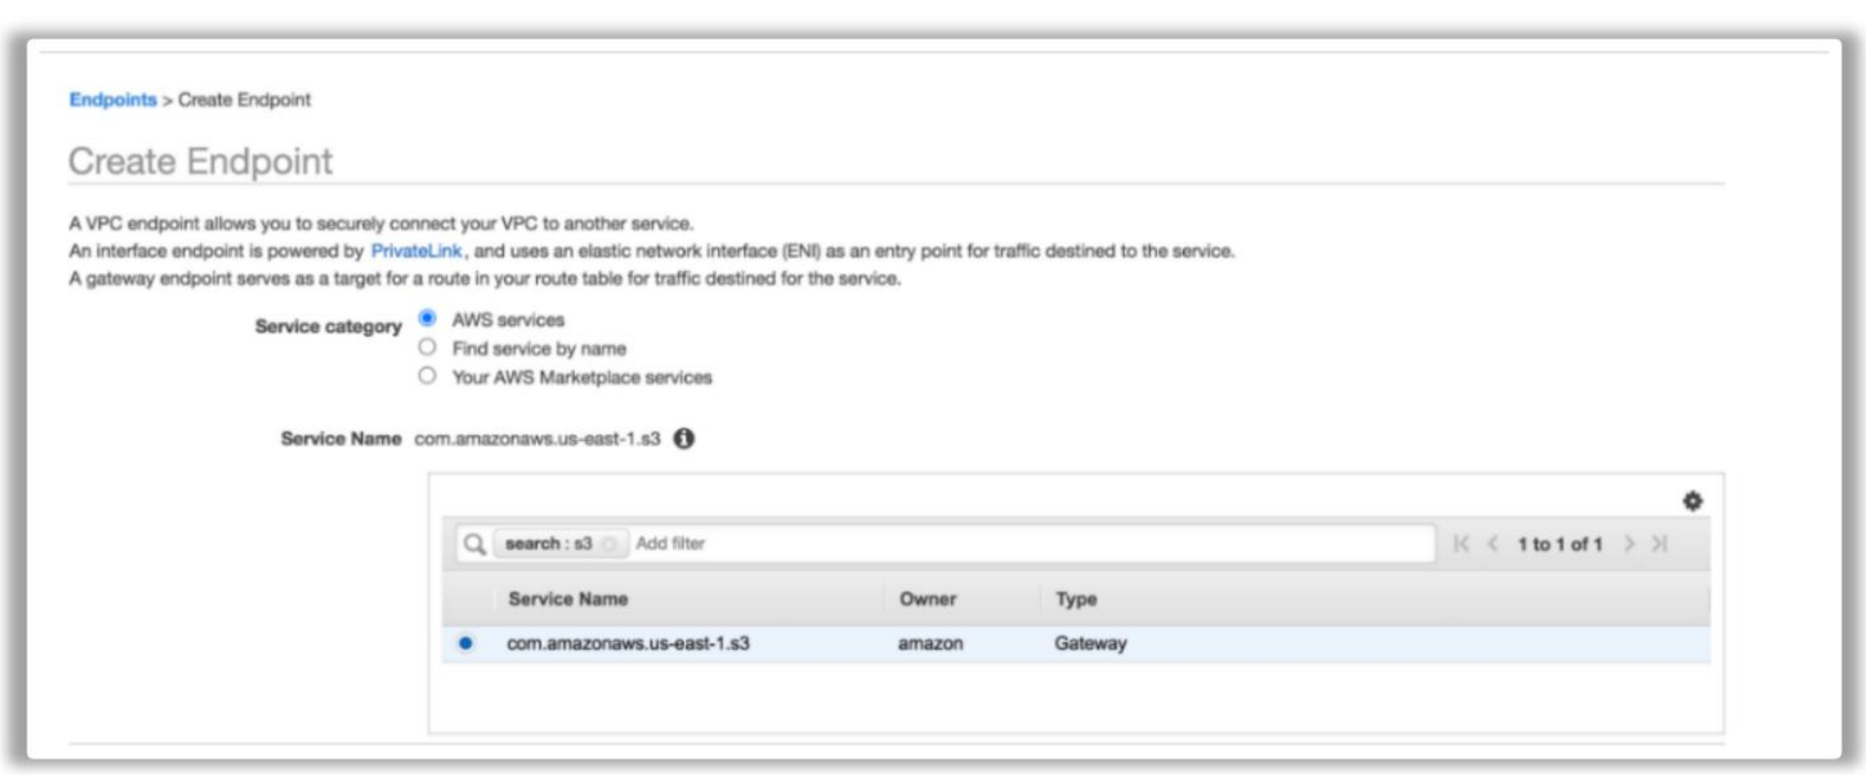


Figure 4.151: Create Endpoint Console

- Select the **VPC** and subnet where you want the endpoint to be created.

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

- Resources in the selected subnets can only access the Amazon S3 endpoint.

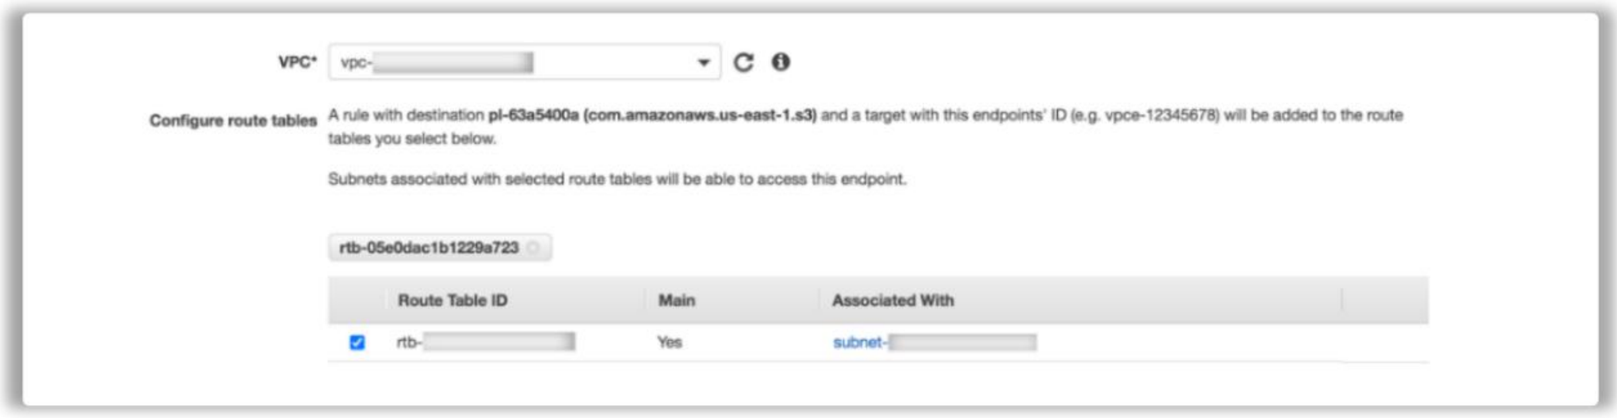


Figure 4.152: Choose Access

- Click on **Create Endpoint**.

Restricting Amazon S3 Bucket to VPC Access Only

To ensure that the Amazon S3 buckets can be accessed only through VPC, it can be done using a bucket policy in the S3 bucket that can restrict to custom VPCs.

- Return to the S3 bucket main page.
- Click on the **Permissions** tab.
- Click on **Bucket Policy**.

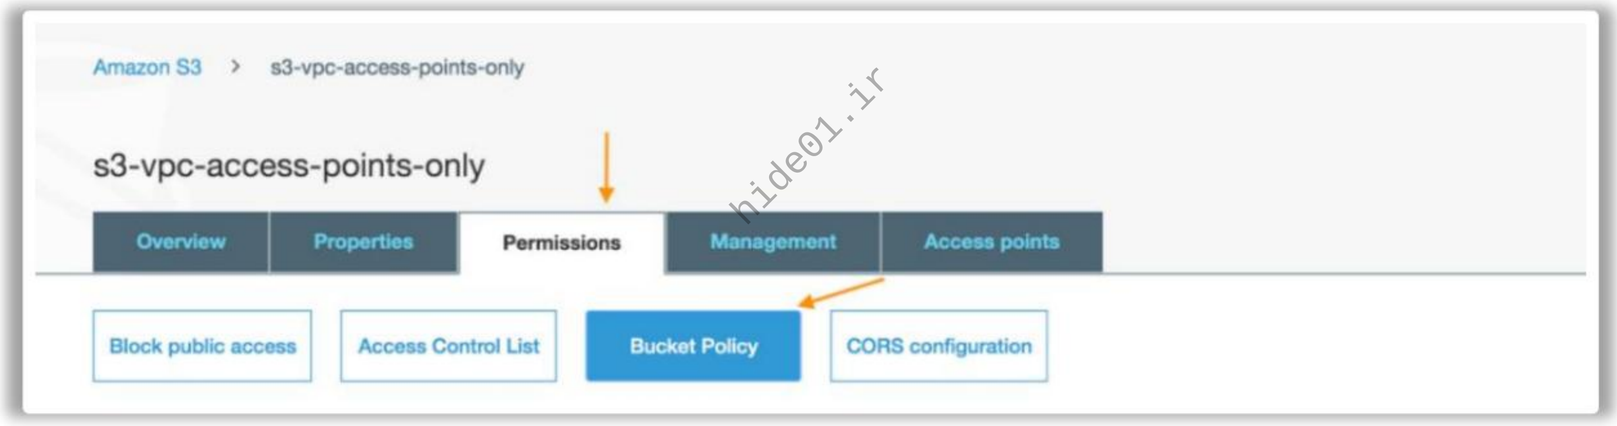


Figure 4.153: Navigate to Bucket Policy

- Create bucket policy.

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



AWS Snowball Best Practices

Implement the following best practices to enhance security, obtain maximum benefits, and satisfaction with the usage of AWS Snowball:

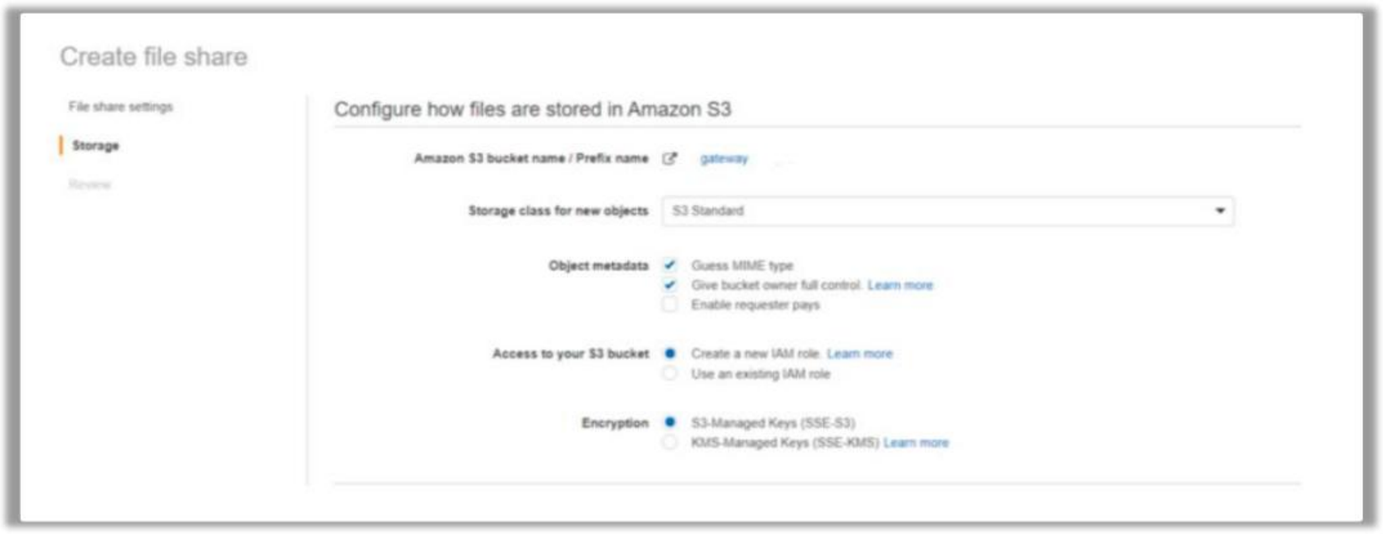
hide01.ir

- Whenever you find anything suspicious in Snowball, it is recommended to contact AWS support before connecting it to the local network. You may get a new Snowball.
- It is highly recommended to avoid saving unlock code in the same place of manifest in the workstation. It prevents unauthorized users from accessing the Snowball. Consider a scenario where the copy of the manifest to the workstation is saved first. Then unlock code is mailed to AWS IAM user, so that the user can perform the transfer of data from the workstation. This will ultimately limit Snowball access to individuals who have access to workstation files and even the email address of IAM users.
- Logs are generated whenever a data transfer is initiated from an on-premises server to Snowball. These logs consist of the file name and its path in plain text, so it is a best-practice to delete these logs after the job is done to protect the sensitive data.
- To transfer data across the network, before initiating a copy to Snowball, batch the local cache, so that the copy operation could be completed as fast as possible.
- Workstation plays a significant role in data transfer. Hence, it is strongly recommended to ensure that the workstation should have robust memory, processing, and networking.
- While transferring data, do not disconnect the Snowball or disturb the network settings, as this may corrupt data.
- At the time of copying, ensure that the file should be in static state. If the files are modified at the time of transfer, then those files would not be imported in Amazon S3.

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

Use KMS CMKs for AWS Storage Gateway File Shares

- To obtain granular control over data-at-rest in AWS storage gateway file shares, use AWS **KMS Customer Master Keys (CMK)** through SSE-KMS



Configure KMS CMKs for AWS Storage Gateway File Shares

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

Use KMS Customer Master Keys for AWS Storage Gateway File Shares

AWS Storage Gateway File Share is a service that provides on-premises access to unlimited cloud storage applications. It makes use of SSL or TLS to encrypt the data while transferring from the customer gateway to AWS storage.

hide01.ir

Cloud customers can also configure their gateway to use SSE with AWS KMS CMK for encrypting the data. For file shares, the cloud customers can make use of AWS KMS CMK through SSE-KMS. This enables them to have granular control over data-at-rest.

Amazon KMS allows the cloud consumers to create, delete, or rotate the master CMKs for encrypting data of file share storage gateway. AWS Storage Gateway is protected by the global network security procedures of AWS. Additionally, all the user requests must be signed using access key ID and secret key. Customers can also sign requests from temporary security credentials created with AWS Simple Token Service (STS).

The storage gateway is integrated with Cloud Trial. All the API calls for storage gateways will be captured by the CloudTrail. By collecting the information from CloudTrail, cloud customers can come to know about the requests made to the Storage gateway, the IP address of the location from which the request was made, time, etc.

Source: https://aws.amazon.com

To use KMS Customer Master Keys for AWS Storage Gateway File Shares, perform the following steps:

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

- Sign into the AWS Management Console.

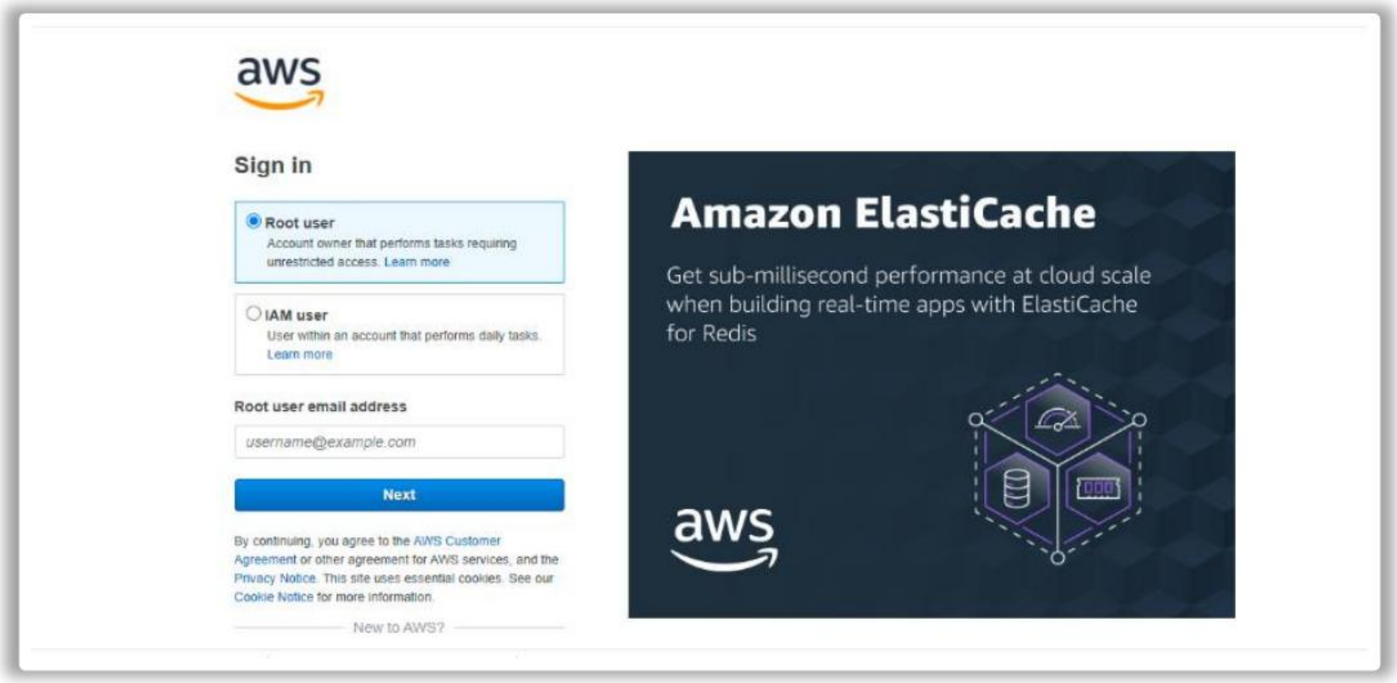


Figure 4.155: Sign in to Console

- Open the "Storage Gateway" service and click on the "Get Started" button.



Figure 4.155: Select Get started

- On the screen, select gateway type as "File gateway".

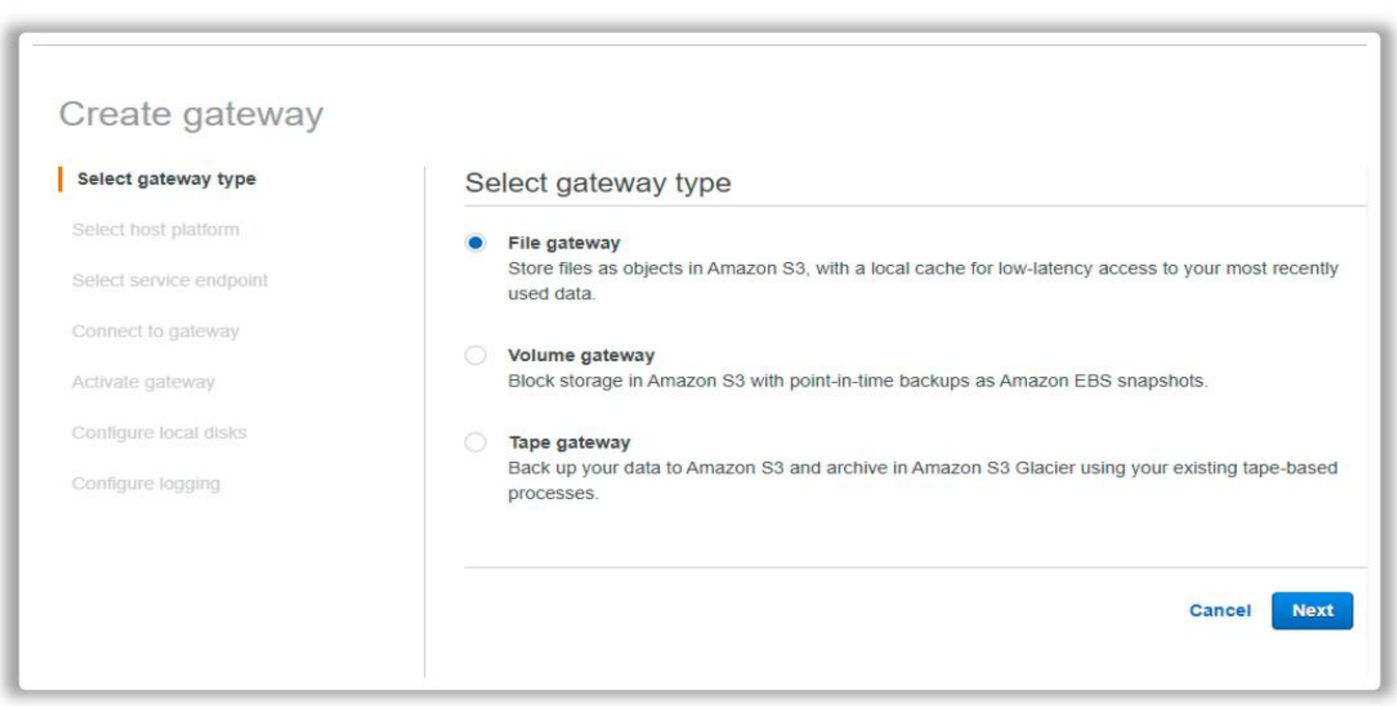


Figure 4.156: Choose File gateway

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

- Now, select the host platform where you want to place your storage gateway.

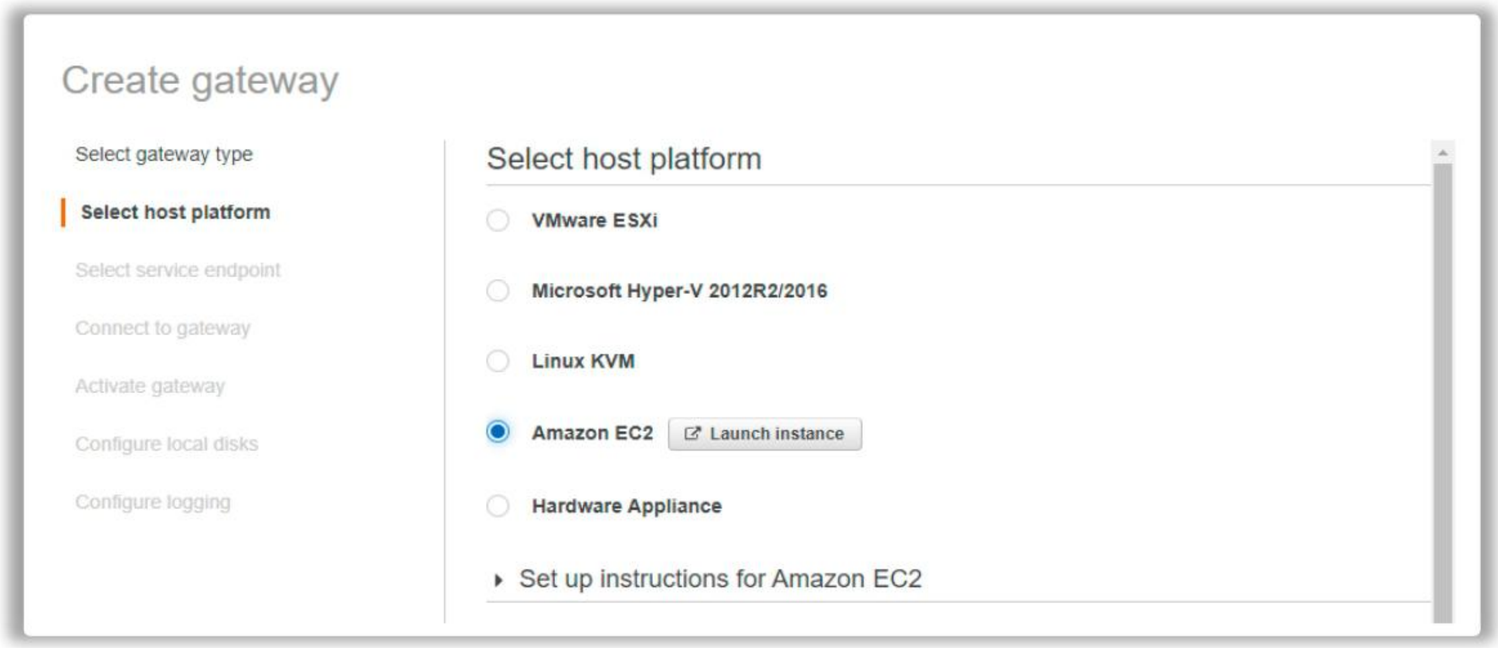


Figure 4.157: Create gateway Console

- Click on “Launch instance” and set up an EC2 instance. Select “m4.xlarge”, which is recommended for storage gateway.

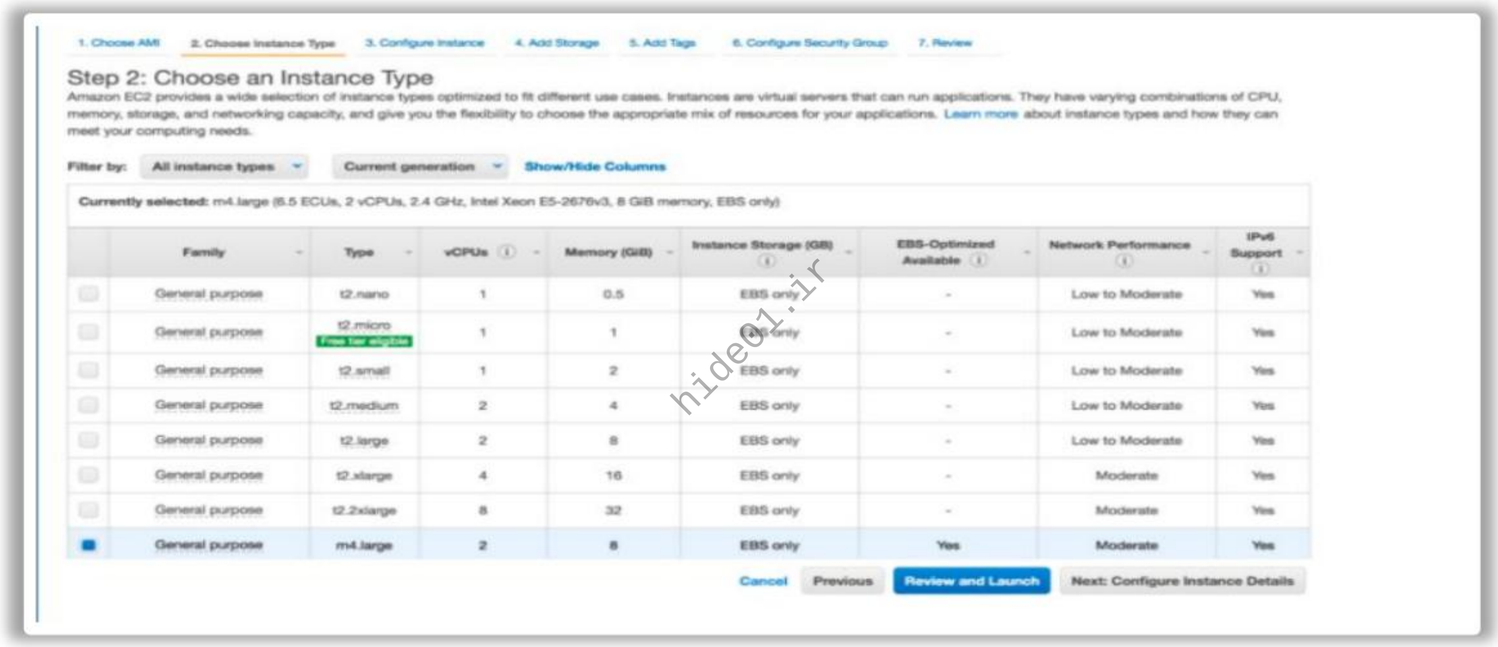


Figure 4.158: Click Launch instance

- In the next step, keep everything as default and click “Add Storage”.

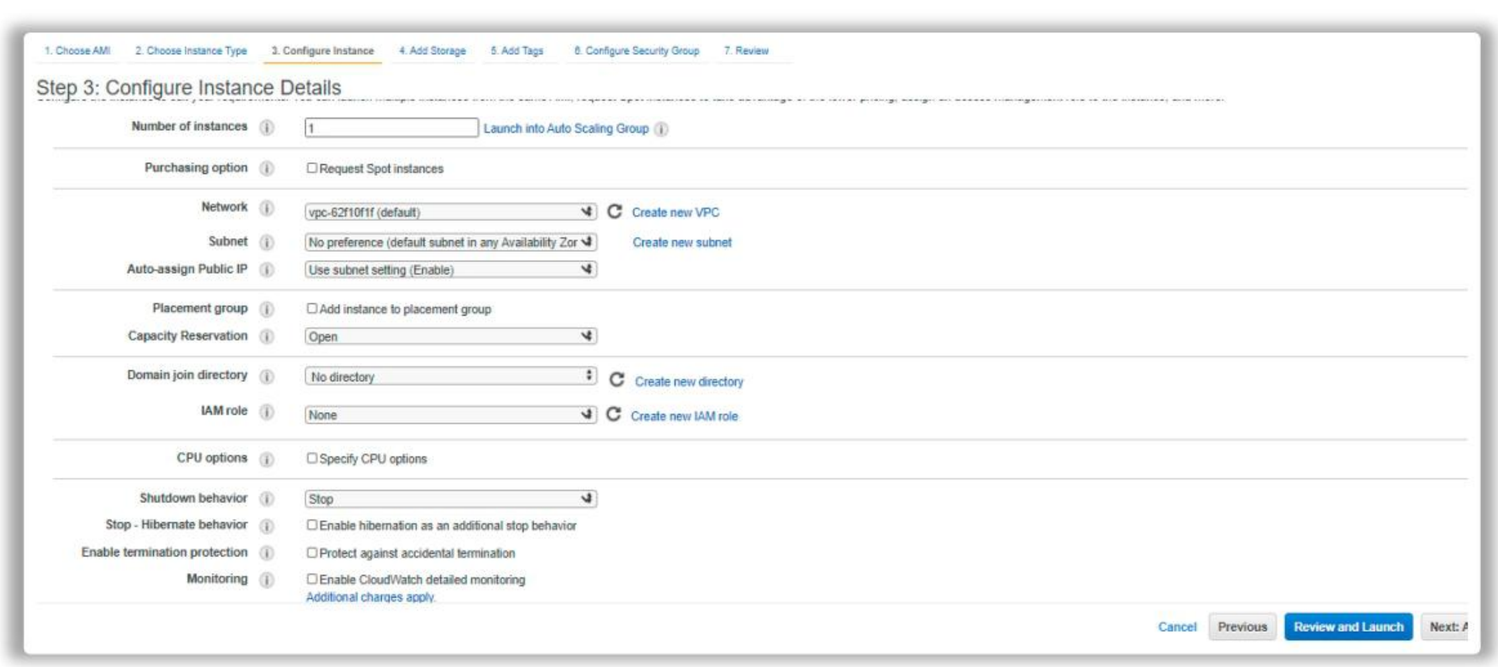


Figure 4.159: Select Add Storage

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

- In storage, add an EBS volume along with the Root for caching. The storage is based on how much cache you want to keep to your S3 bucket.

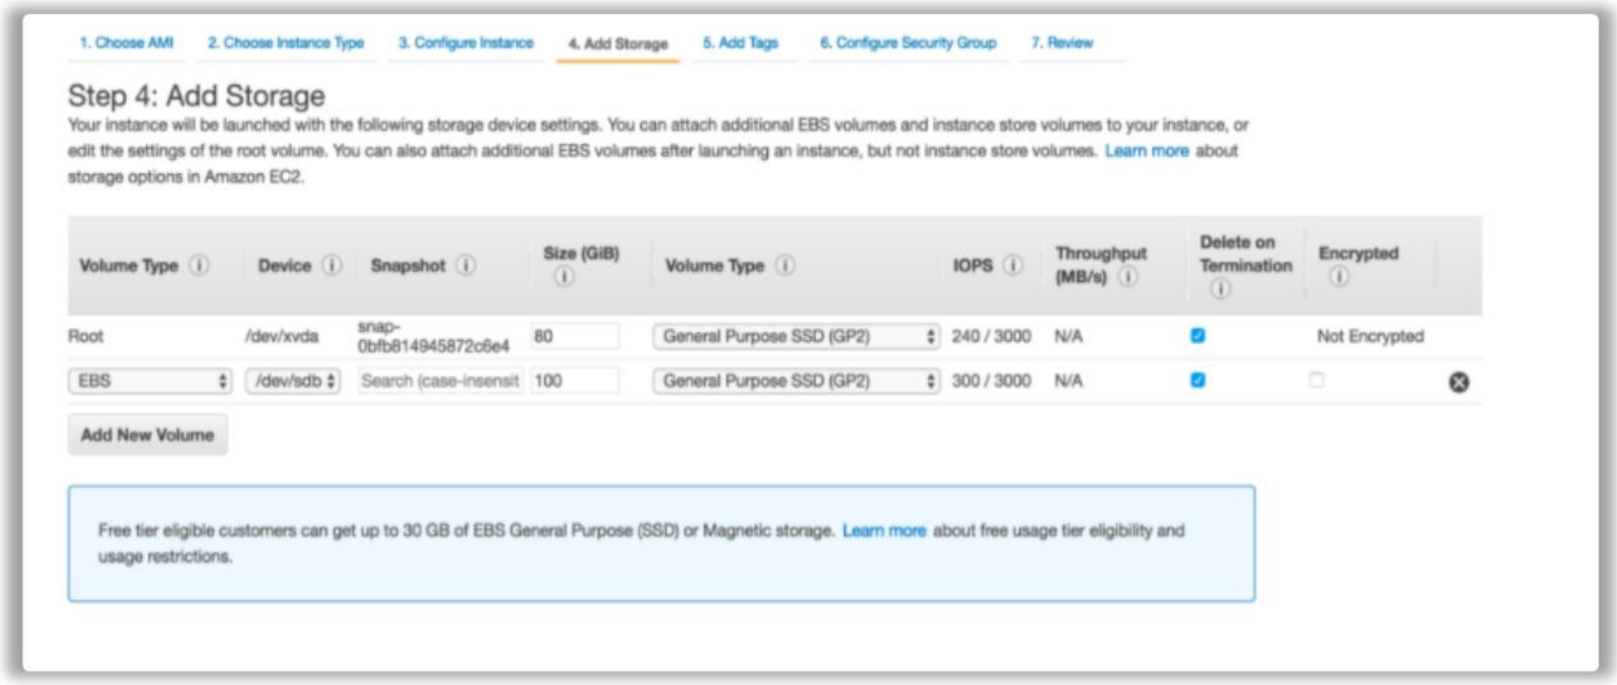


Figure 4.160: Add Storage

- Now, you can add tags to your EC2 instance.

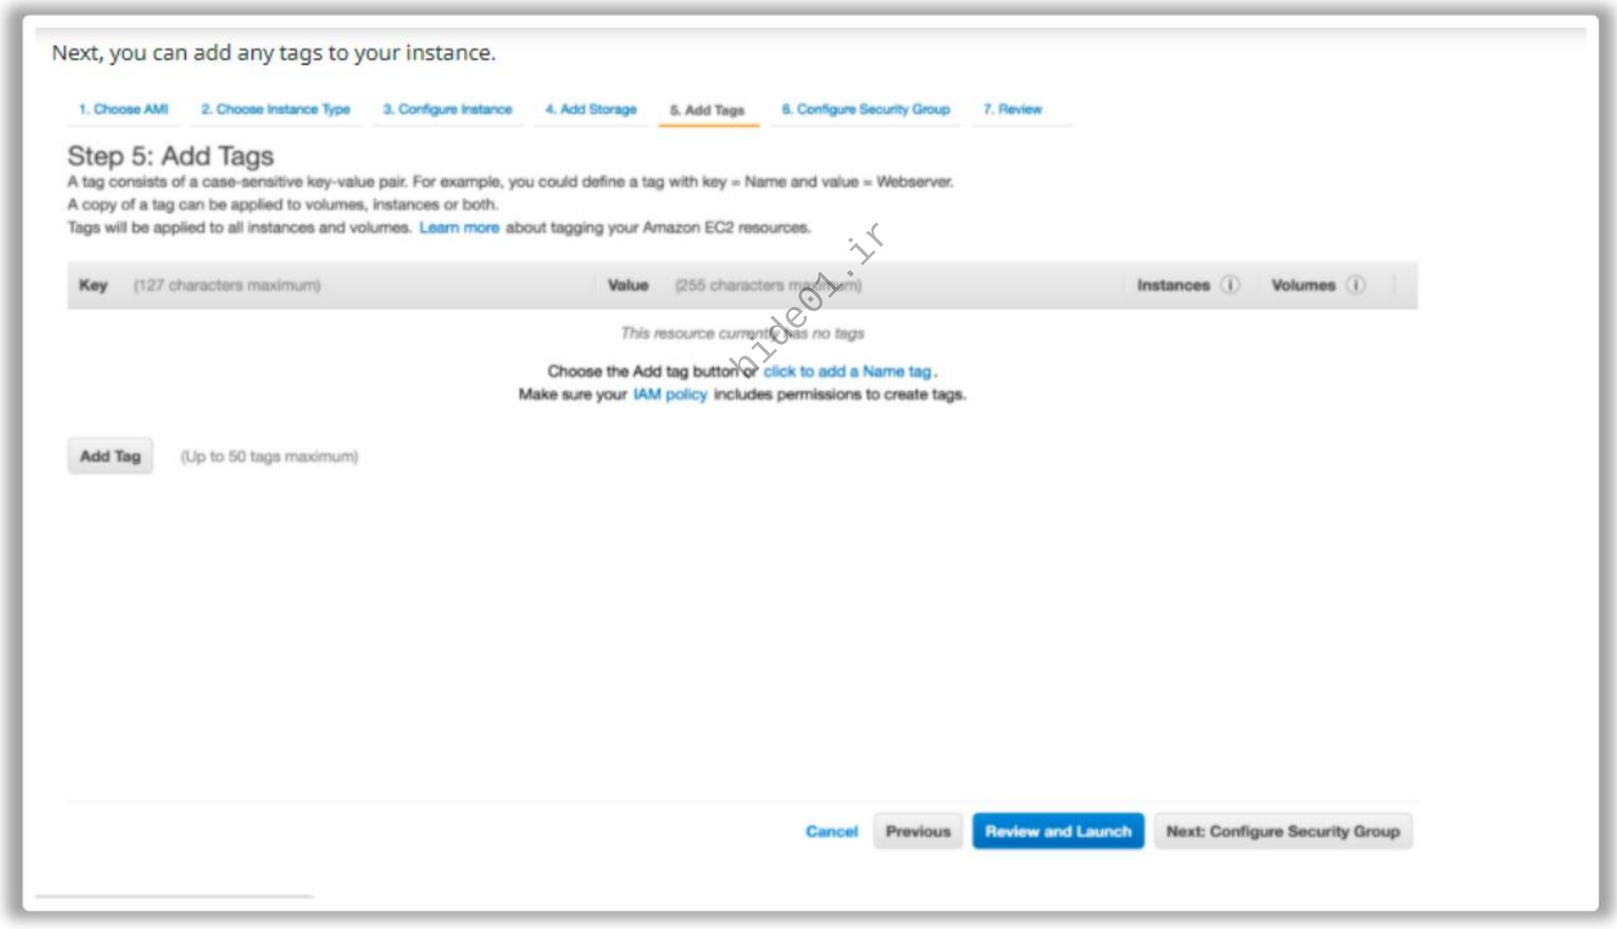


Figure 4.161: Add Tags

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

- On the security groups, you can choose an existing one or create a new security group as required. Now, click on “Review and Launch” to launch the instance.

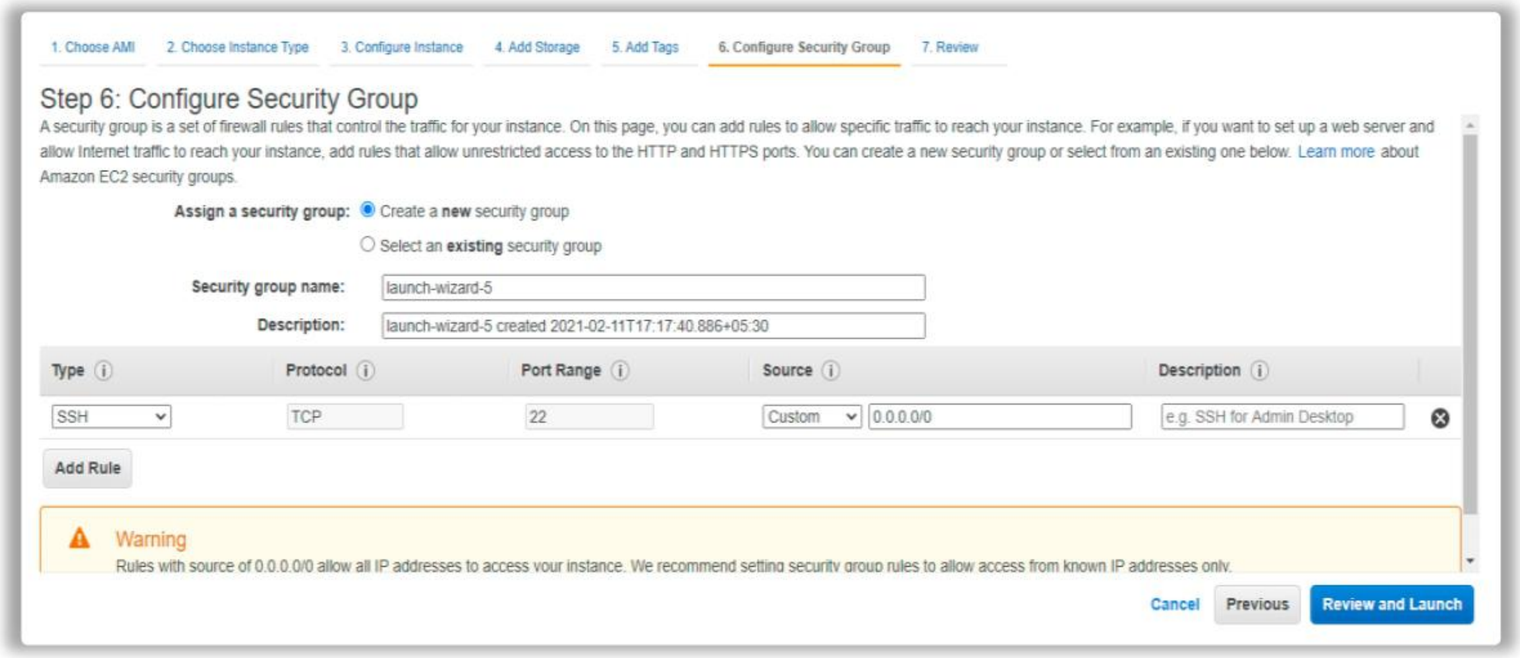


Figure 4.162: Configure Security Groups

- Now, select a key pair. It can be either an existing one or you can create a new pair, or you can proceed without a key pair. Then, click on “launch”. After launching the instance, copy the public IP address.

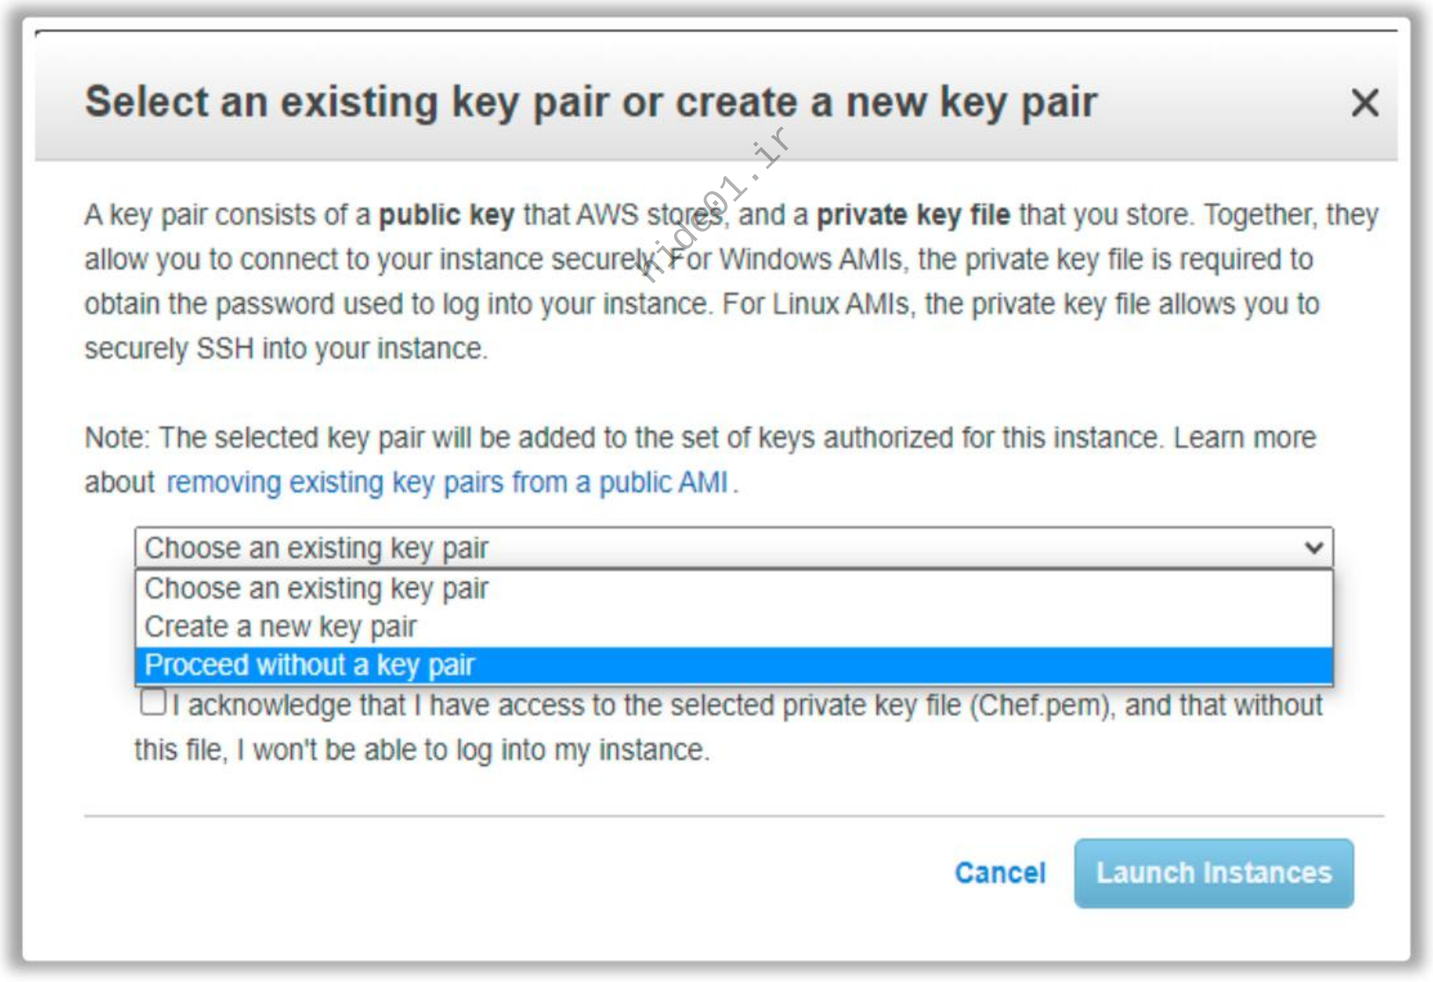


Figure 4.163: Select Launch Instances

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

- Now, jump to the storage gateway wizard and select the “Service End Point” as “Public”. FIPS and VPC can be selected to enhance additional secure transition.

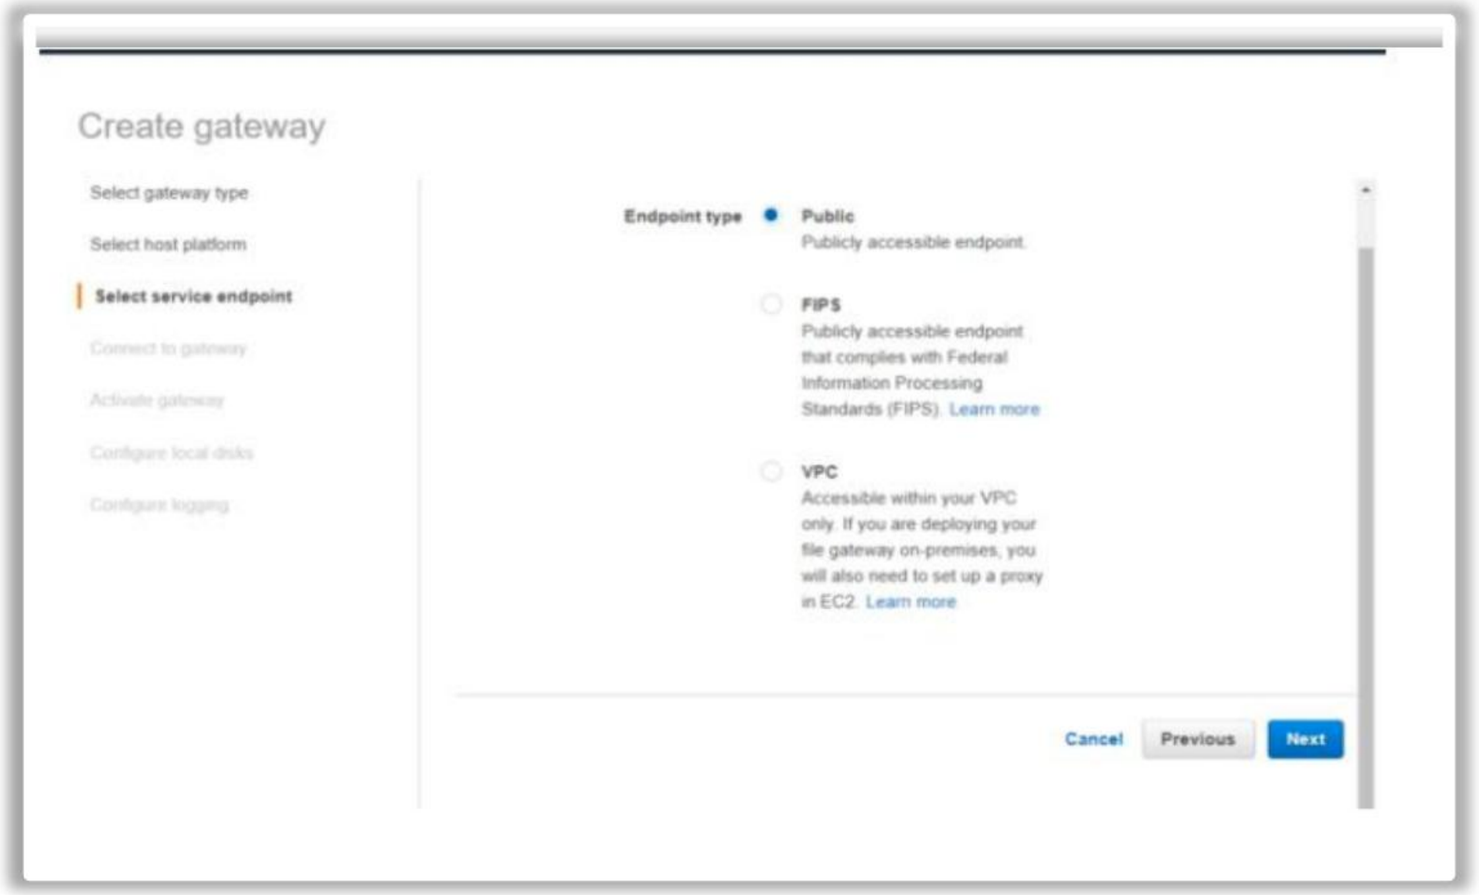


Figure 4.164: Select Public

- Click “Next” and paste the IP address of the above launched instance to connect to the gateway.



hide01.ir

Figure 4.165: Paste IP Address

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

- Once you connect to the gateway, activate the gateway by choosing a time zone and give a name to your gateway.

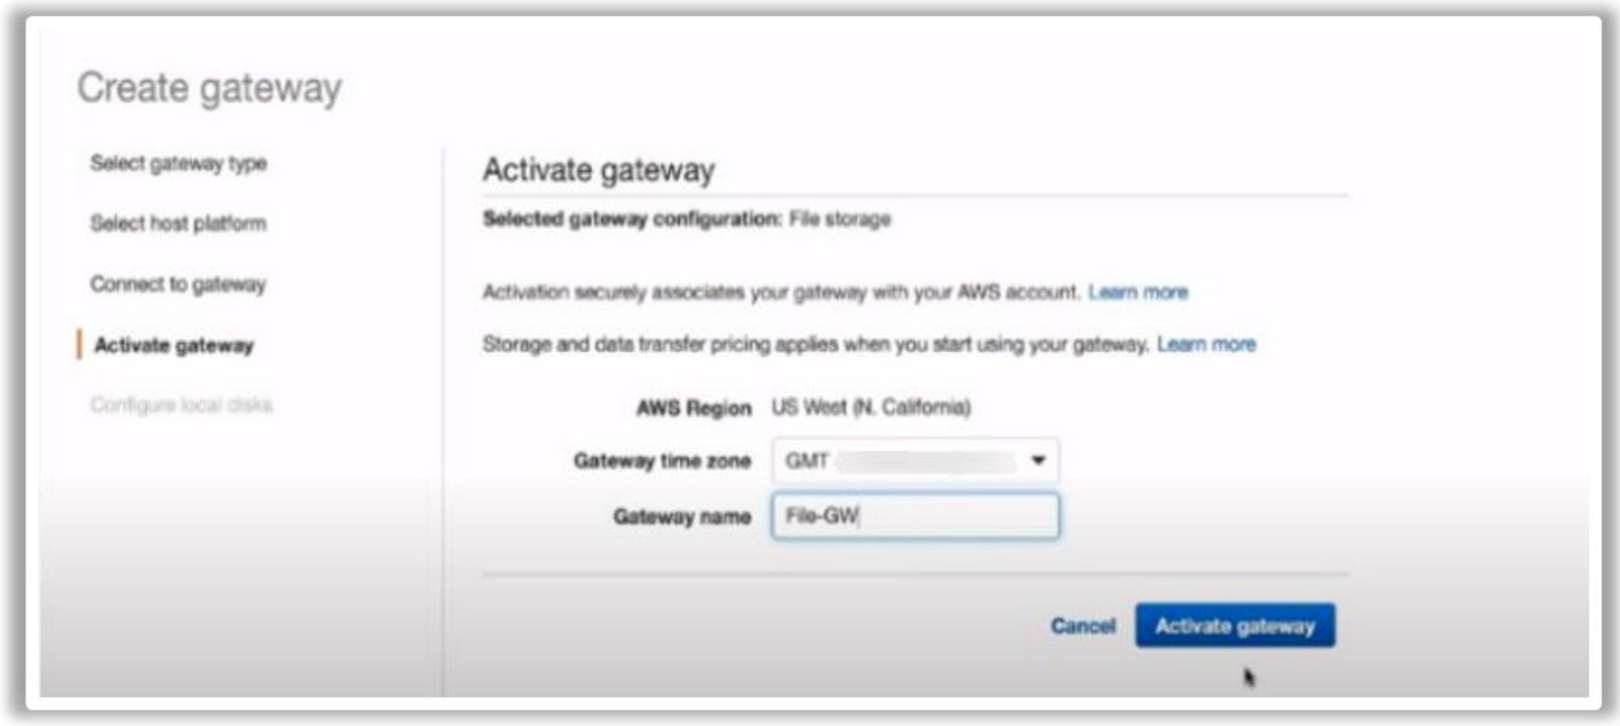


Figure 4.166: Select Active gateways

- Now, click on "Save and Continue" to successfully create a file storage gateway.

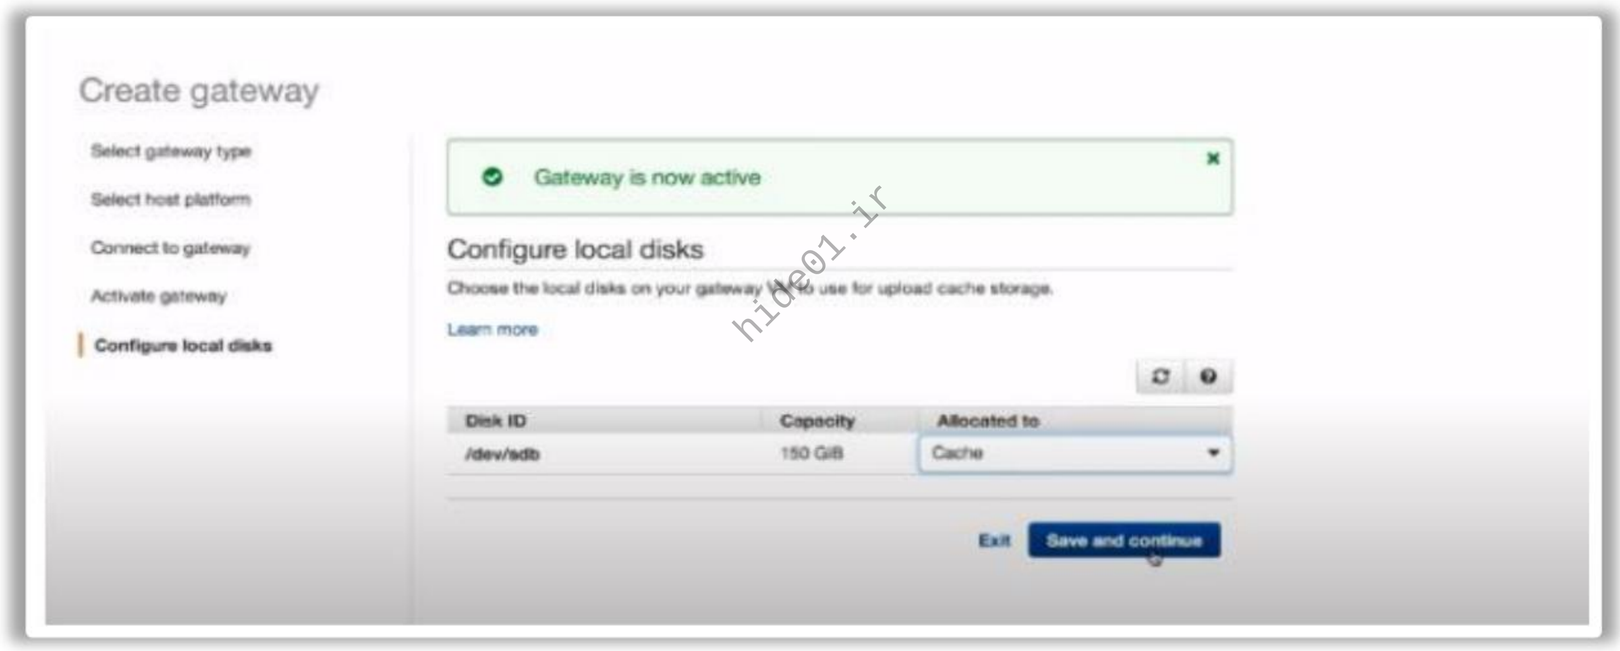


Figure 4.167: Click Save and Continue

- Now, go to the S3 service and create an S3 bucket for the file share. Upload objects to the bucket.

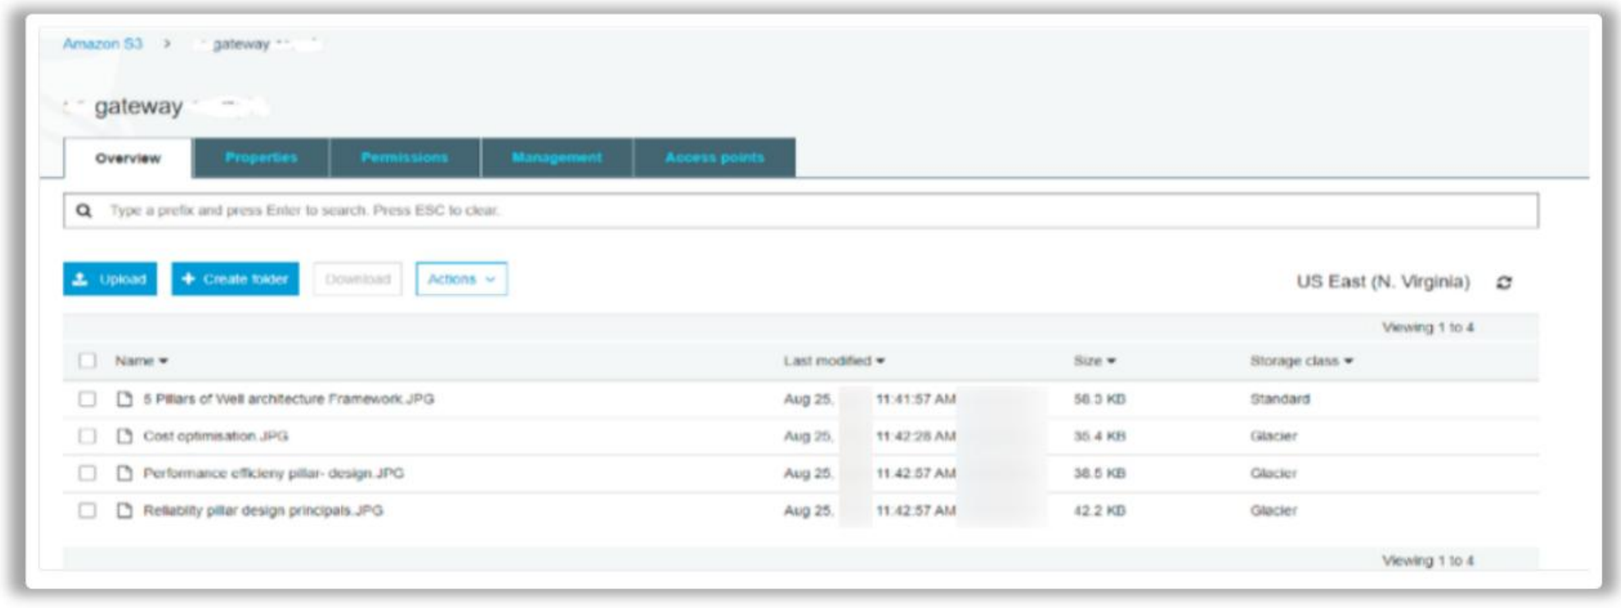


Figure 4.168: Gateways Console

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

- Now, jump to the "Storage gateway" console. On the left panel, go to "File Shares" and select "Create File Share". Type the S3 bucket name you have created and select "Access object using" as NFS. Now, enter the name of the gateway you have created earlier.

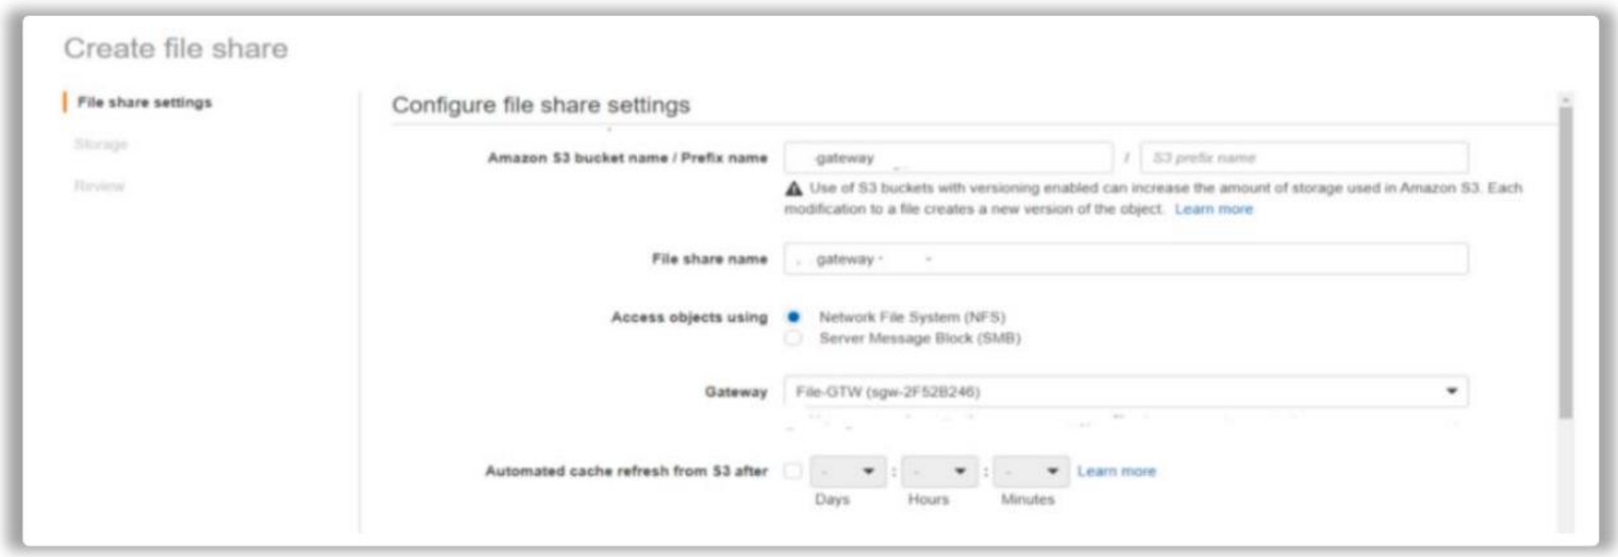


Figure 4.169: Configure File Share Settings

- Select storage class for new objects as "S3 Standard". Select Encryption as "KMS-Managed keys (SSE-KMS)." Review everything and click on "Create File Share".

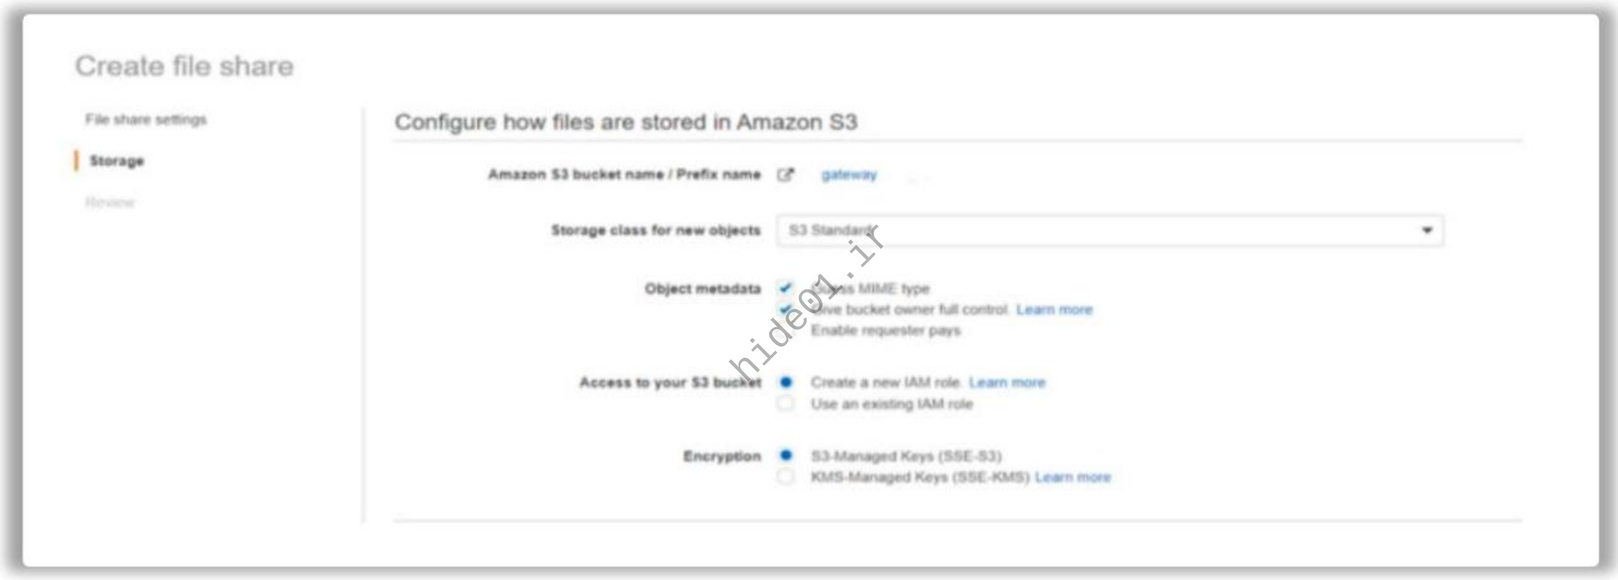
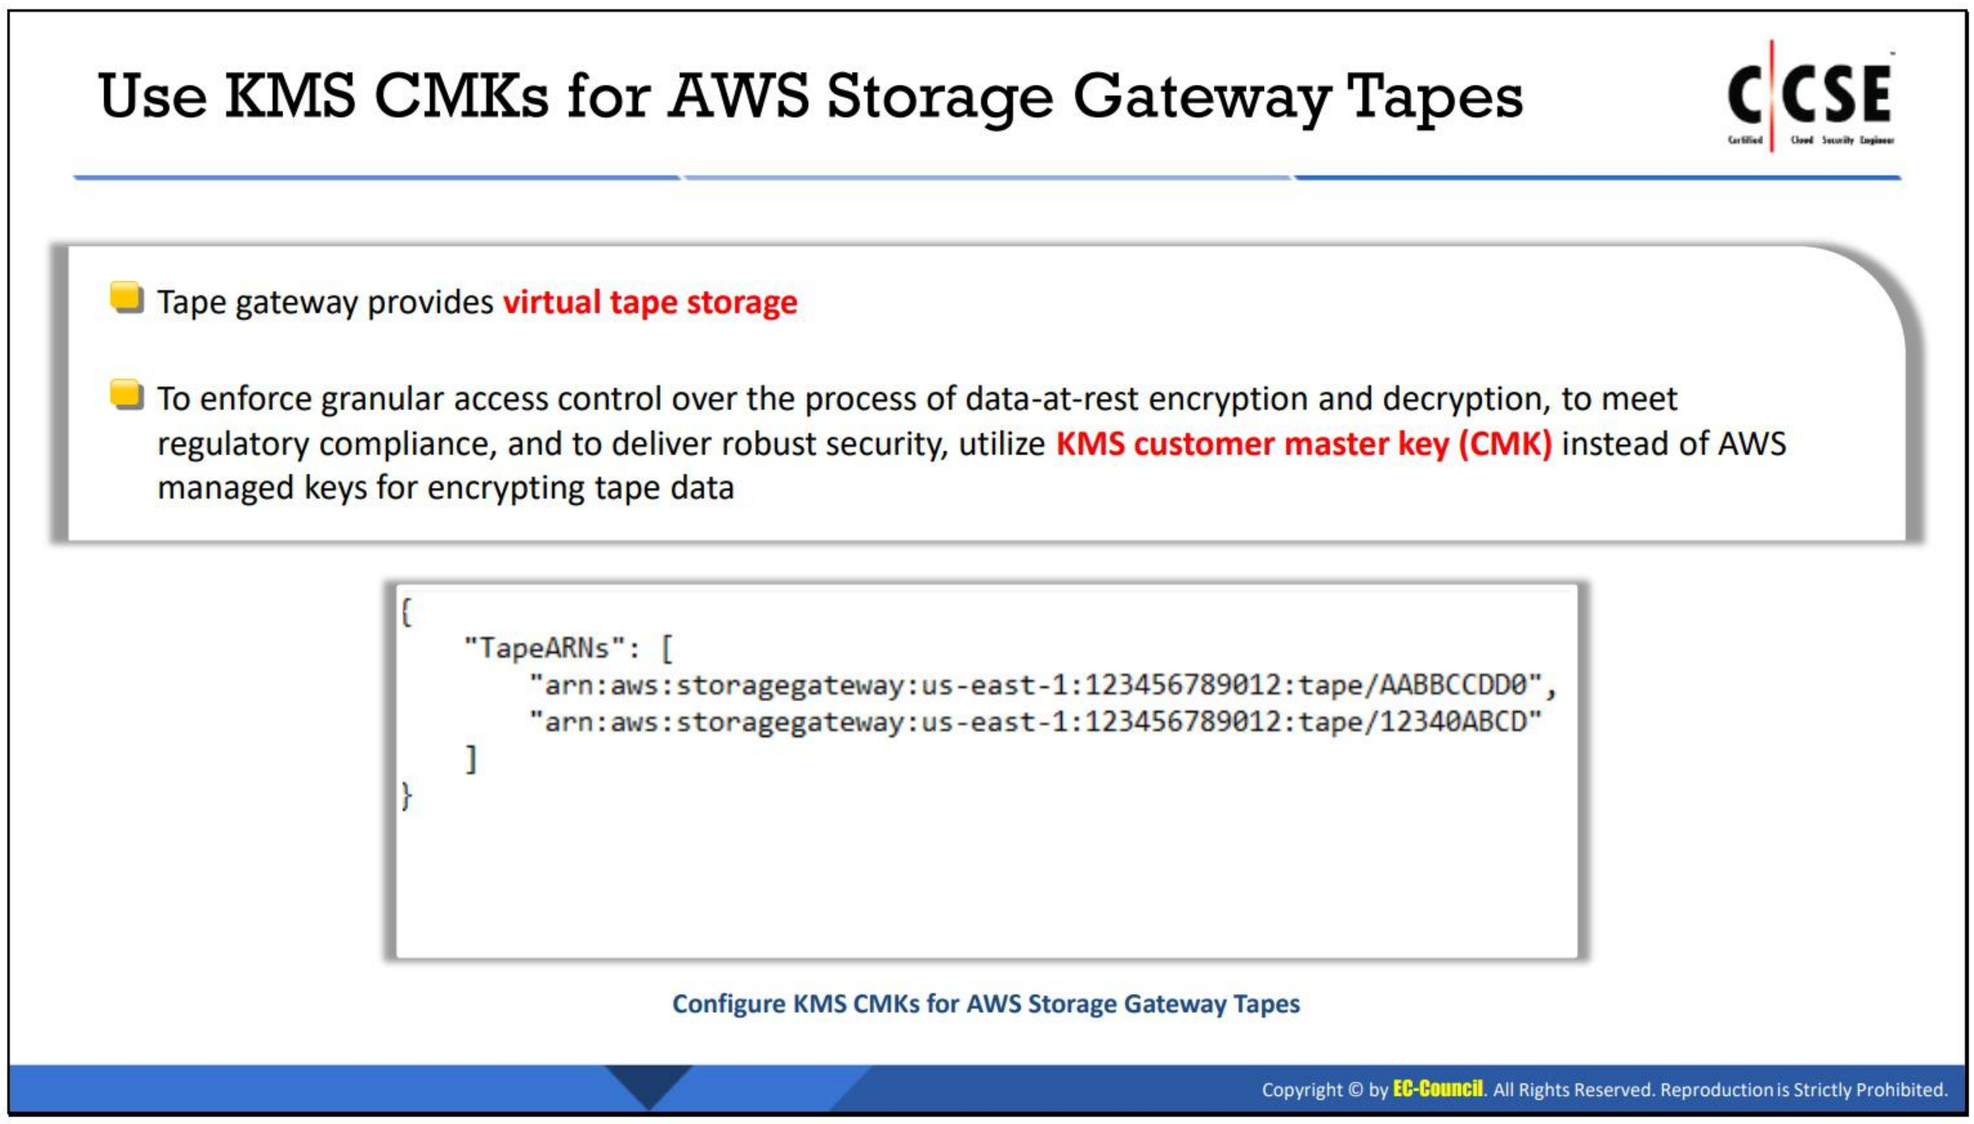


Figure 4.170: Click Create File Share

- When the file share gateway is available, you can mount it to your local on-premises storage.

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



Use KMS Customer Master Keys for AWS Storage Gateway Tapes

Tape gateway provides virtual tape storage. It is cost-effective and cloud customers can durably archive their backup data in AWS storage such as S3 Glacier and Deep Archive. It consists of a VTL interface that can be used along with the backup software infrastructure. Tape gateway reduces the burden of scaling and managing the tapes on-premises. As backup data is archived, it is important to ensure that the data is secure in the cloud. For this, cloud customers should consider using KMS customer master key instead of AWS managed keys for encrypting tape data.

For existing Amazon storage gateway virtual tapes, encryption using CMK cannot be enabled. To encrypt stored tape data with the customer-managed keys, specified tapes must be recreated.

To use KMS customer master keys for AWS Storage Gateway tapes, perform the following steps:

- Define an IAM policy to enable IAM users to manage CMK and perform encryption and decryption on tape data. For this, create a policy and assign a name to it. Paste the following content:

```
{
  "Version": "2012-10-17",
  "Id": "sg-volume-custom-key-policy",
  "Statement": [
    {
      "Sid": "Enable IAM User Permissions",
      "Effect": "Allow",
      "Principal": {
```


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

```
        "AWS": "arn:aws:iam::123456789012:root"
      },
      "Action": "kms:*",
      "Resource": "*"
    },
    {
      "Sid": "Grant access to CMK manager",
      "Effect": "Allow",
      "Principal": {
        "AWS": "arn:aws:iam::123456789012:role/SGTapeManager"
      },
      "Action": [
        "kms:Create*",
        "kms:Describe*",
        "kms:Enable*",
        "kms:List*",
        "kms:Put*",
        "kms:Update*",
        "kms:Revoke*",
        "kms:Disable*",
        "kms:Get*",
        "kms:Delete*",
        "kms:ScheduleKeyDeletion",
        "kms:CancelKeyDeletion"
      ],
      "Resource": "*"
    },
    {
      "Sid": "Allow the use of the CMK",
      "Effect": "Allow",
      "Principal": {
        "AWS": "arn:aws:iam::123456789012:user/SGTapeAdmin"
      },
      "Action": [
        "kms:Encrypt",
```

hide01.ir

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

```
            "kms:Decrypt",
            "kms:ReEncrypt*",
            "kms:GenerateDataKey*",
            "kms:DescribeKey"
        ],
        "Resource": "*"
    },
    {
        "Sid": "Allow attachment of persistent resources",
        "Effect": "Allow",
        "Principal": {
            "AWS": "arn:aws:iam::123456789012:user/SGTapeAdmin"
        },
        "Action": [
            "kms:CreateGrant",
            "kms:ListGrants",
            "kms:RevokeGrant"
        ],
        "Resource": "*",
        "Condition": {
            "Bool": {
                "kms:GrantIsForAWSResource": "true"
            }
        }
    }
]
}
```

hide01.ir

- Now, run **`create-key`** command using the name of policy document to create a new KMS customer master key.

```
aws kms create-key
    --region us-east-1
    --description 'KMS CMK for encrypting virtual tape data.'
    --policy file://virtual-tape-cmk-policy.json
```

- The output of the above command gives metadata of the new KMS CMK. Copy the key ARN as it will be used later.

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

```
            {
              "KeyMetadata": {
                "Origin": "AWS_KMS",
                "KeyId": "abcd1234-abcd-1234-abcd-1234ancd1234",
                  "Description": "KMS CMK for encrypting virtual tape
data.",
                  "Enabled": true,
                   "KeyUsage": "ENCRYPT_DECRYPT",
                    "KeyState": "Enabled",
                     "CreationDate": 1517237395.290,
                    "Arn":                          "arn:aws:kms:us-east-
1:123456789012:key/abcd1234-abcd-1234-abcd-1234ancd1234",
                        "AWSAccountId": "123456789012"
                }
                     }
```

- Run **create-alias** command including the key ARN copied in the previous step. This adds an alias to CMK.

```
aws kms create-alias
    --region us-east-1
    --alias-name alias/SGTapeCustomCMK
    --target-key-id                    arn:aws:kms:us-east-
1:123456789012:key/abcd1234-abcd-1234-abcd-  1234ancd1234
```

- Now, create a virtual tape for the existing tape gateway. Mention ARN of the newly created CMK to enable data encryption using CMK.

```
aws storagegateway create-tapes
    --region us-east-1
    --gateway-arn                  arn:aws:storagegateway:us-east-
1:123456789012:gateway/sgw-12AB34CD
    --tape-size-in-bytes 161061273600
    --client-token AAAABBBB
    --num-tapes-to-create 2
    --tape-barcode-prefix ABC
    --kms-encrypted
    --kms-key    arn:aws:kms:us-east-1:123456789012:key/abcd1234-
abcd-1234-abcd-1234ancd1234
```

- The output is as follows:

```
{
```

hide01.ir

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

```
    "TapeARNs": [
        "arn:aws:storagegateway:us-east-
1:123456789012:tape/AABBCCDD0",
        "arn:aws:storagegateway:us-east-
1:123456789012:tape/12340ABCD"
    ]
}
```

hide01.ir

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

Use KMS Customer Master Keys for AWS Storage Gateway Volumes

- To fulfill the compliance requirement and to have granular control over the data-at-rest encryption and decryption process, make sure that the **AWS storage gateway volume** is using KMS customer master keys instead of default keys.
- By using **KMS customer master keys**, the AWS customer can decide who can have access to the encryption keys.


```
{
    "TargetARN": "arn:aws:storagegateway:us-east-
1:123456789012:gateway/sgw-12341234/target/iqn.1997-10.com.amazon:cc-
stored-volume",
    "VolumeSizeInBytes": 161061273600,
    "VolumeARN": "arn:aws:storagegateway:us-east-
1:123456789012:gateway/sgw-12341234/volume/vol-01234123412341234"
}
{
    "TargetARN": "arn:aws:storagegateway:us-east-
1:123456789012:gateway/sgw-12341234/target/iqn.1997-10.com.amazon:cc-
stored-volume",
    "VolumeSizeInBytes": 161061273600,
    "VolumeARN": "arn:aws:storagegateway:us-east-
1:123456789012:gateway/sgw-12341234/volume/vol-01234123412341234"
}
```


KMS Customer Master Keys for AWS Storage Gateway Volumes

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

Use KMS Customer Master Keys for AWS Storage Gateway Volumes

Storage gateway volumes provide cloud backed storage volumes that can be mounted as iSCSI to the on-premises applications. Volume gateway supports cache and stored volumes. It stores and manages on-premises data in AWS S3 storage service.

To fulfill compliance requirements and to have granular control over the data-at-rest encryption and encryption process, make sure that the AWS Storage Gateway volume is using KMS customer-managed keys instead of default keys. By using KMS CMK, the AWS customer can decide who can have access to the encryption keys.

To Use KMS Customer Master Keys for AWS Storage Gateway Volumes

The following are the steps to create a new AWS KMS customer master key and relaunch the specified volumes using AWS command-line interface (CLI):

- Define an IAM policy to enable the IAM users to manage CMK and perform encryption and decryption on tape data. For this, create a policy and assign a name to it. Paste the following content:

```
{
  "Version": "2012-10-17",
  "Id": "custom-key-policy",
  "Statement": [
    {
      "Sid": "Enable IAM User Permissions",
      "Effect": "Allow",
```


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

```
      "Principal": {
        "AWS": "arn:aws:iam::123456789012:root"
      },
      "Action": "kms:*",
      "Resource": "*"
    },
    {
      "Sid": "Grant access to CMK manager",
      "Effect": "Allow",
      "Principal": {
        "AWS": "arn:aws:iam::123456789012:role/VolumeManager"
      },
      "Action": [
        "kms:Create*",
        "kms:Describe*",
        "kms:Enable*",
        "kms:List*",
        "kms:Put*",
        "kms:Update*",
        "kms:Revoke*",
        "kms:Disable*",
        "kms:Get*",
        "kms:Delete*",
        "kms:ScheduleKeyDeletion",
        "kms:CancelKeyDeletion"
      ],
      "Resource": "*"
    },
    {
      "Sid": "Allow the use of the CMK",
      "Effect": "Allow",
      "Principal": {
        "AWS": "arn:aws:iam::123456789012:user/VolumeAdmin"
      },
      "Action": [
```

hide01.ir

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

```
            "kms:Encrypt",
            "kms:Decrypt",
            "kms:ReEncrypt*",
            "kms:GenerateDataKey*",
            "kms:DescribeKey"
        ],
        "Resource": "*"
    },
    {
        "Sid": "Allow attachment of persistent resources",
        "Effect": "Allow",
        "Principal": {
            "AWS": "arn:aws:iam::123456789012:user/VolumeAdmin"
        },
        "Action": [
            "kms:CreateGrant",
            "kms:ListGrants",
            "kms:RevokeGrant"
        ],
        "Resource": "*",
        "Condition": {
            "Bool": {
                "kms:GrantIsForAWSResource": "true"
            }
        }
    }
}
```

hide01.ir

- Now, run the create-key command using the name of the policy document to create a new KMS customer master key.

```
aws kms create-key
    --region us-east-1
    --description 'KMS CMK for encrypting Storage Gateway volume data.'
    --policy file://cmk-policy.json
```


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

- The output of the command should display the metadata of the new KMS CMK. Copy the ARN.

```
aws kms create-key
    --region us-east-1
    --description 'KMS CMK for encrypting Storage Gateway volume data.'
    --policy file://cmk-policy.json
```

- Run **create-alias** command including the key ARN copied in the previous step. This adds an alias to CMK.

```
aws kms create-alias
    --region us-east-1
    --alias-name alias/VolumeCustomCMK
    --target-key-id arn:aws:kms:us-east-1:123456789012:key/1234abcd-abcd-1234-abcd-1234abcd1234
```

- (A) Run **create-cached-iscsi-volume** command to create a new cache volume for the volume gateway. Use the ARN of newly created CMK to enable encryption using customer master keys.

```
aws storagegateway create-cached-iscsi-volume
    --region us-east-1
    --gateway-arn arn:aws:storagegateway:us-east-1:123456789012:gateway/sgw-12341234
    --volume-size-in-bytes 161061273600
    --target-name cc-cached-volume
    --network-interface-id 172.30.50.15
    --client-token AABBCCDD
    --kms-encrypted
    --kms-key arn:aws:kms:us-east-1:123456789012:key/1234abcd-abcd-1234-abcd-1234abcd1234
```

hide01.ir

- The output should be metadata.

```
{
    "TargetARN": "arn:aws:storagegateway:us-east-1:123456789012:gateway/sgw-12341234/target/iqn.1995-15.com.amazon:cc-cached-volume",
    "VolumeARN": "arn:aws:storagegateway:us-east-1:123456789012:gateway/sgw-12341234/volume/vol-01234AABB1234CCDD"
}
```


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

- (B) To create stored volume for volume gateway, run **create-stored-iscsi-volume** command. Use the ARN of newly created CMK to enable encryption using customer master keys.

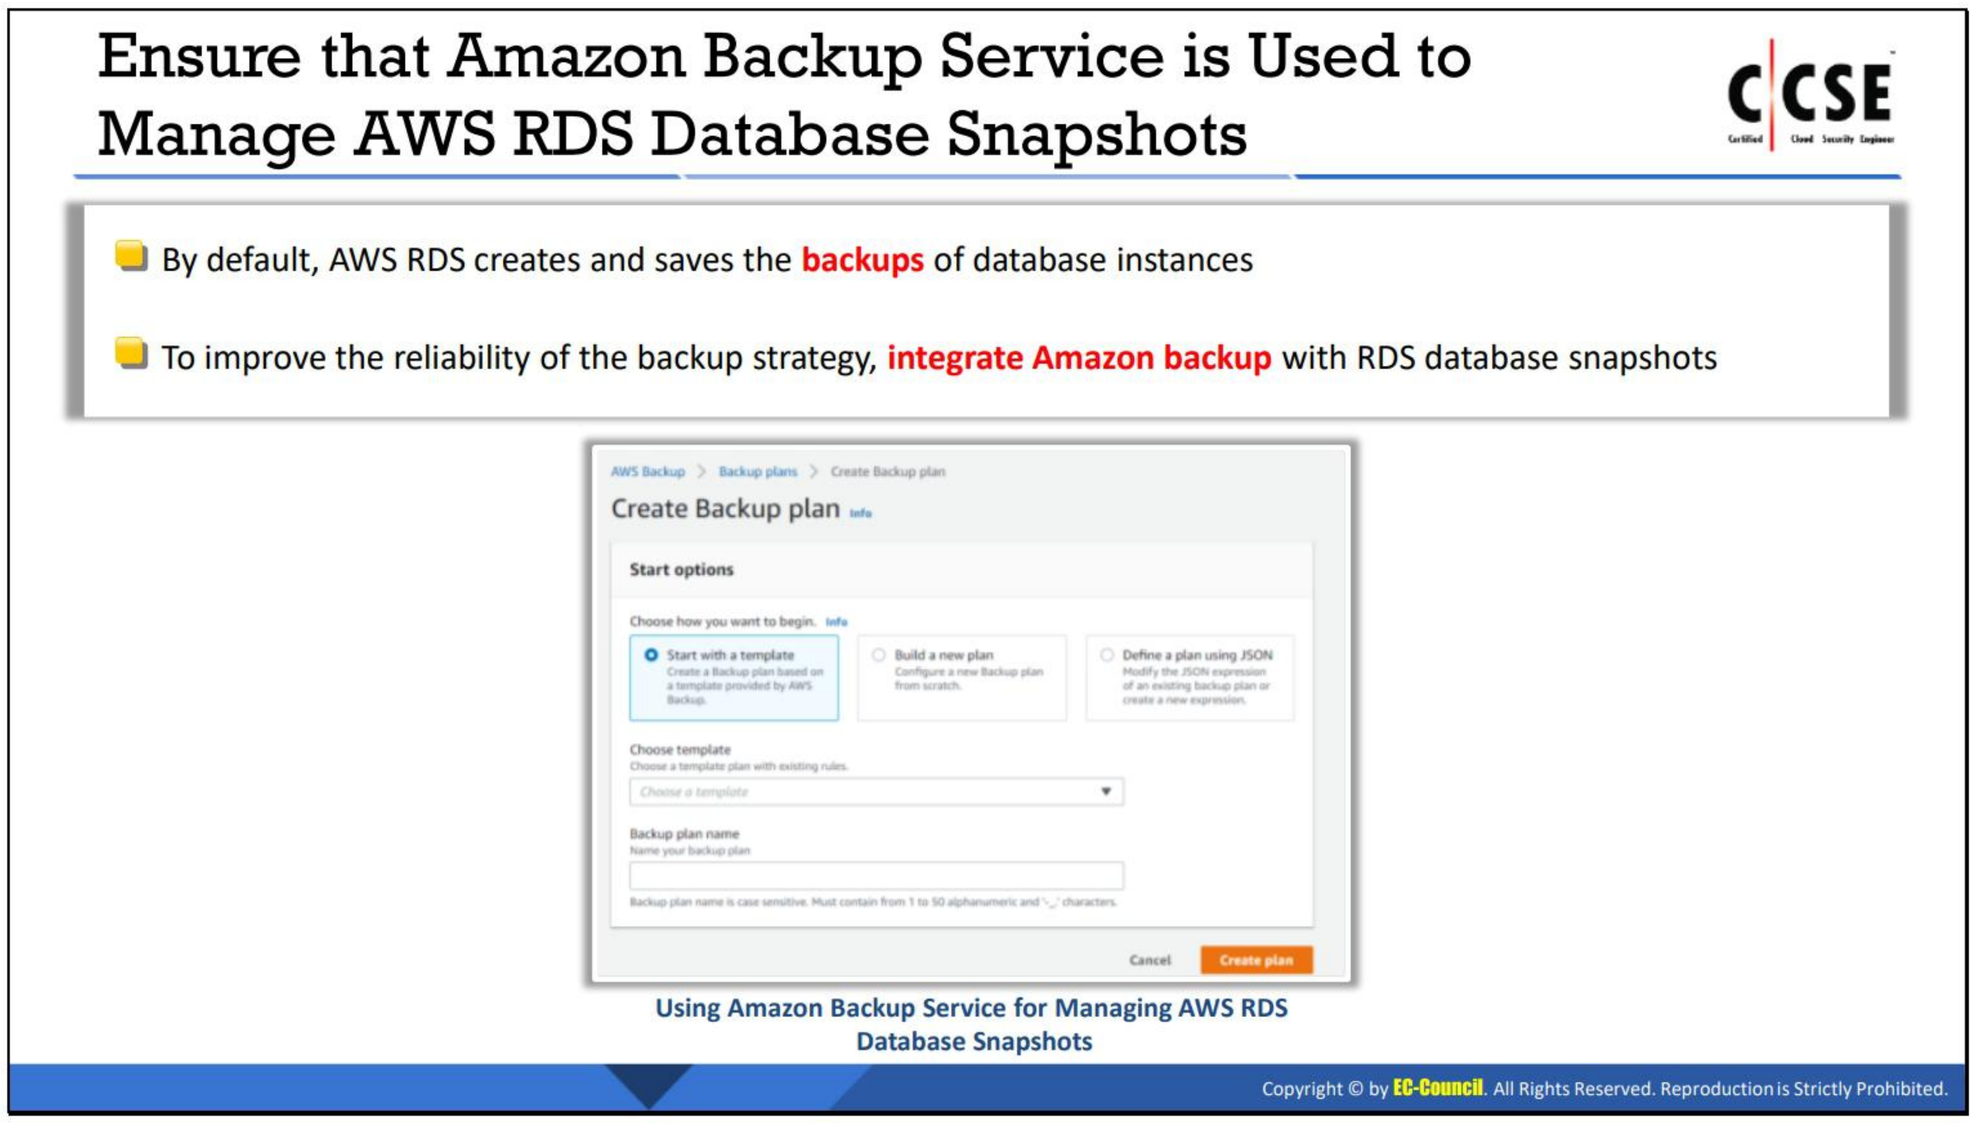
```
aws storagegateway create-stored-iscsi-volume
    --region us-east-1
    --gateway-arn                    arn:aws:storagegateway:us-east-
1:123456789012:gateway/sgw-12341234
    --target-name cc-stored-volume
    --network-interface-id 172.200.10.75
    --preserve-existing-data
    --disk-id aaaabbbb-cccc-aaaa-bbbb-aaaabbbbcccc
    --kms-encrypted
    --kms-key     arn:aws:kms:us-east-1:123456789012:key/1234abcd-
abcd-1234-abcd-1234abcd1234
```

- The output should be the metadata of the **create-stored-iscsi-volume** request.

```
{
    "TargetARN":                    "arn:aws:storagegateway:us-east-
1:123456789012:gateway/sgw-12341234/target/iqn.1997-
10.com.amazon:cc-stored-volume",
    "VolumeSizeInBytes": 161061273600,
    "VolumeARN":                    "arn:aws:storagegateway:us-east-
1:123456789012:gateway/sgw-12341234/volume/vol-0123412341234"
}
```

hi@01.ir

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



Using Amazon Backup Service for Managing AWS RDS Database Snapshots

Ensure that Amazon Backup Service is Used to Manage AWS RDS Database Snapshots

By default, AWS RDS creates and saves the backups of database instances. A storage volume snapshot will be created by RDS for the DB instance. RDS retains the automated backups as per the specified retention period. Cloud consumers can also recover a backup at any point during the retention period.

Amazon Backup is a fully managed service and it automatically creates, deletes, and restores the backup without the user intervention. If Amazon Backup is integrated with RDS database snapshots, it improves the reliability of the backup strategy.

To Use Amazon Backup Service for Managing AWS RDS Database Snapshots

The following are the steps to use the Amazon Backup service to manage RDS database snapshots:

- Sign into the AWS Management Console.
- From the console, navigate to the AWS backup service. Click on “Manage Backup”.
- Now, click the “Create Backup Plan” option and give a name to your backup plan.
- You can choose a predefined plan from the existing templates for backup and retention.
- Assign a name to your backup plan under “Backup plan name”.
- Choose the configurations for your rule and assign a name to your rule.
- For backup vault, select the default one.
- To keep a copy of the backup, you can select a location where your backup will be copied.

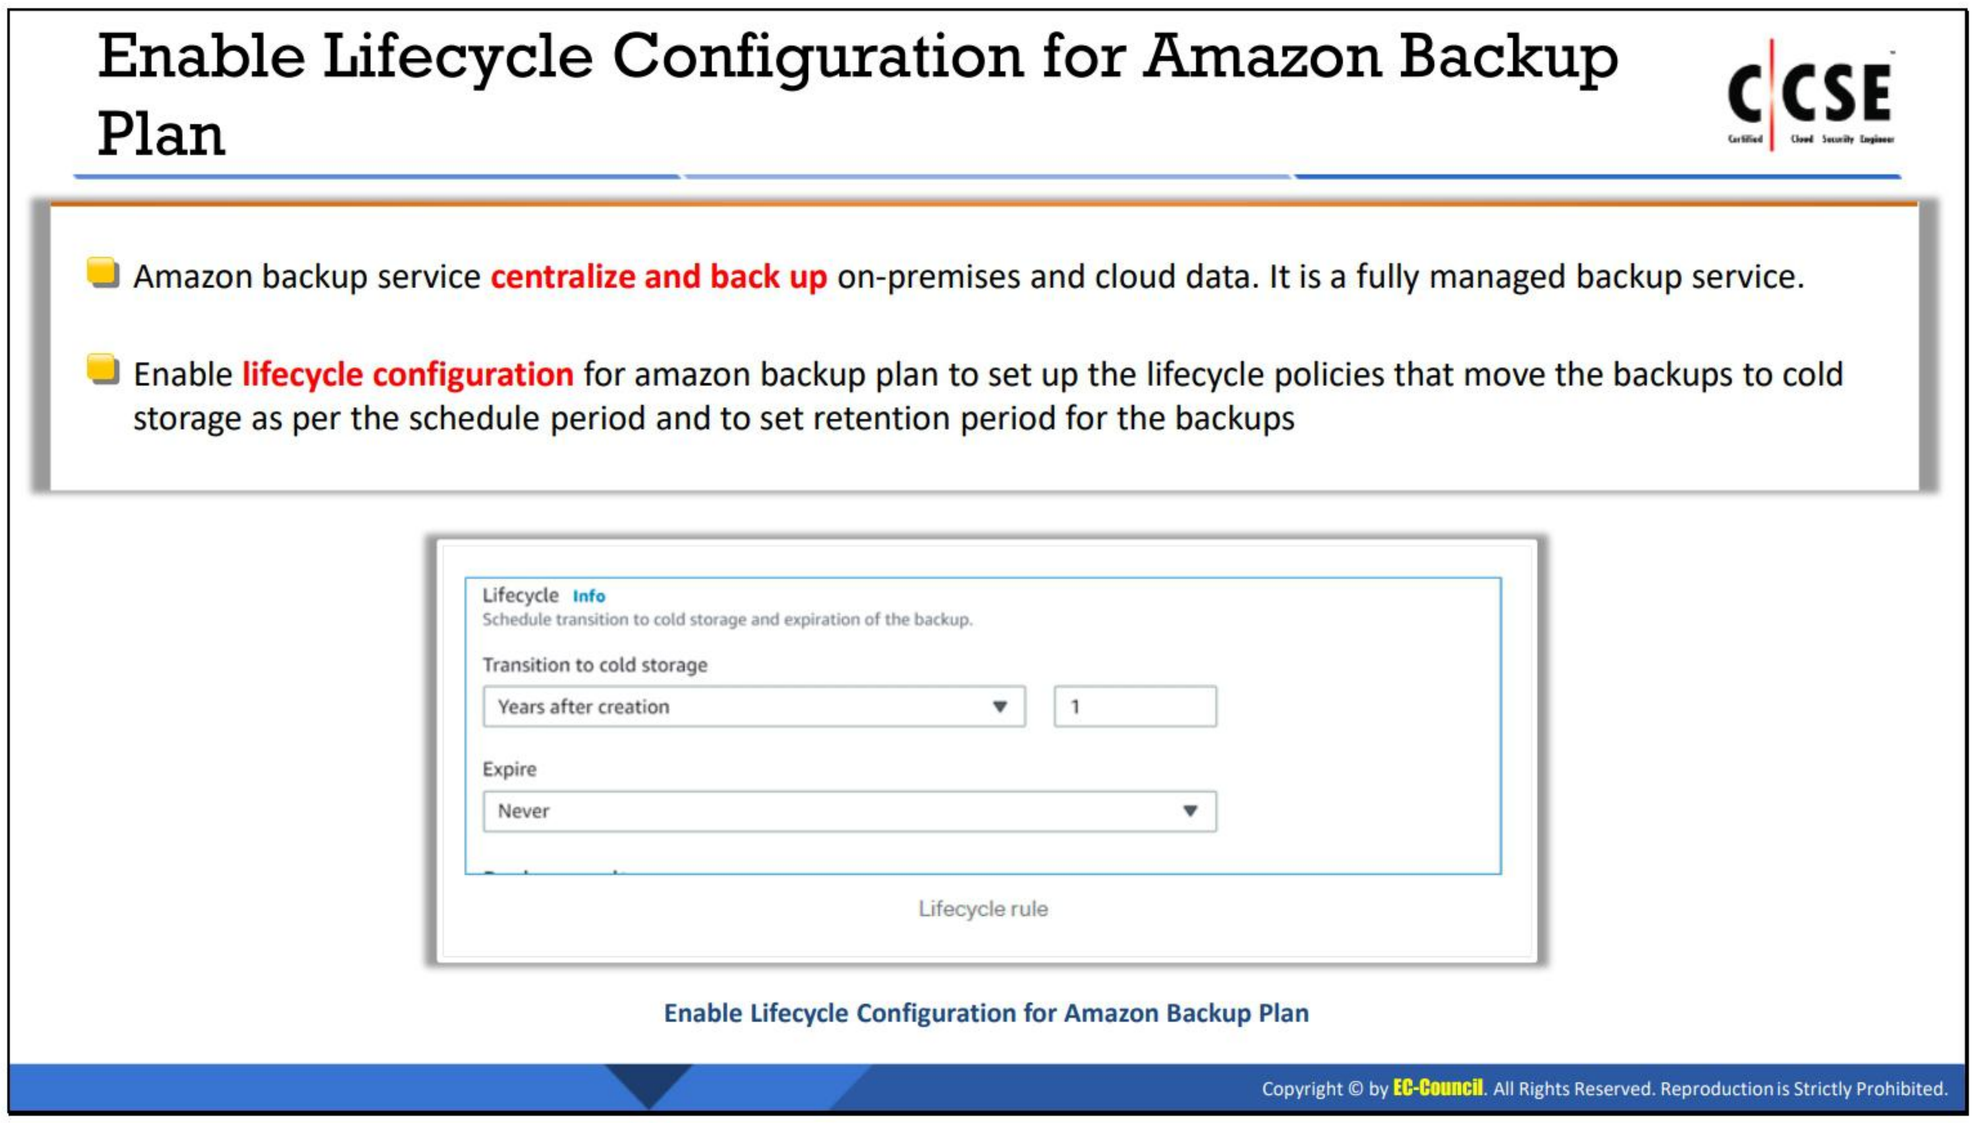
Certified Cloud Security Engineer Copyright © by **EC-Council**
All Rights Reserved. Reproduction is Strictly Prohibited.

- Define any tag sets under the tags section for your backup plan and recovery points. This helps to organize and track your backup plan.
- Now, click on the "Create Plan" to create a backup plan.
- Select the newly created backup plan. Navigate to the "Resource Assignment" section and click on the "Assign Resources" button to apply the new backup plan to AWS RDS snapshots.
- Assign a name for the resource.
- Under the IAM role section, choose the default role.
- Select the resource ID from the dropdown list of "Assign by". Select RDS from the resource list.
- Now, click inside the "Database name" and select the RDS instance which you want to assign to the backup plan. You can add additional resources by clicking "Add Assignment."
- Once you are done with the assignments, click on the "Assign Resources" button.

Thus, RDS database instances use the newly created backup plan to take daily snapshots. The snapshots are also managed using the rules of the AWS backup plan.

hide01.ir

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



Enable Lifecycle Configuration for Amazon Backup Plan

Amazon Backup service centralizes and backs up on-premises and cloud data By utilizing Amazon Backup, cloud consumers can set up backup policies and can monitor the backup activity for all AWS resources. Backup policies can be created easily in few clicks and cloud consumers can automate their backup schedule and retention period.

While using AWS backup, cloud consumers must meet the compliance requirements, which can be achieved using lifecycle management policies. It minimizes storage costs by storing the backup in low-cost storage called cold storage. Cloud consumers can set up the lifecycle policies that move the backups to cold storage as per the schedule period. The lifecycle policies also define the time for which the backups need to expire (retention period). Hence, it is beneficial if the lifecycle configuration is enabled for Amazon Backup plans. This ensures that the backup plans meet the regulatory compliance of the enterprise.

Source: https://aws.amazon.com

The following are the steps to implement Lifecycle Configuration for Amazon Backup plan:

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

- Sign into the AWS Management Console.

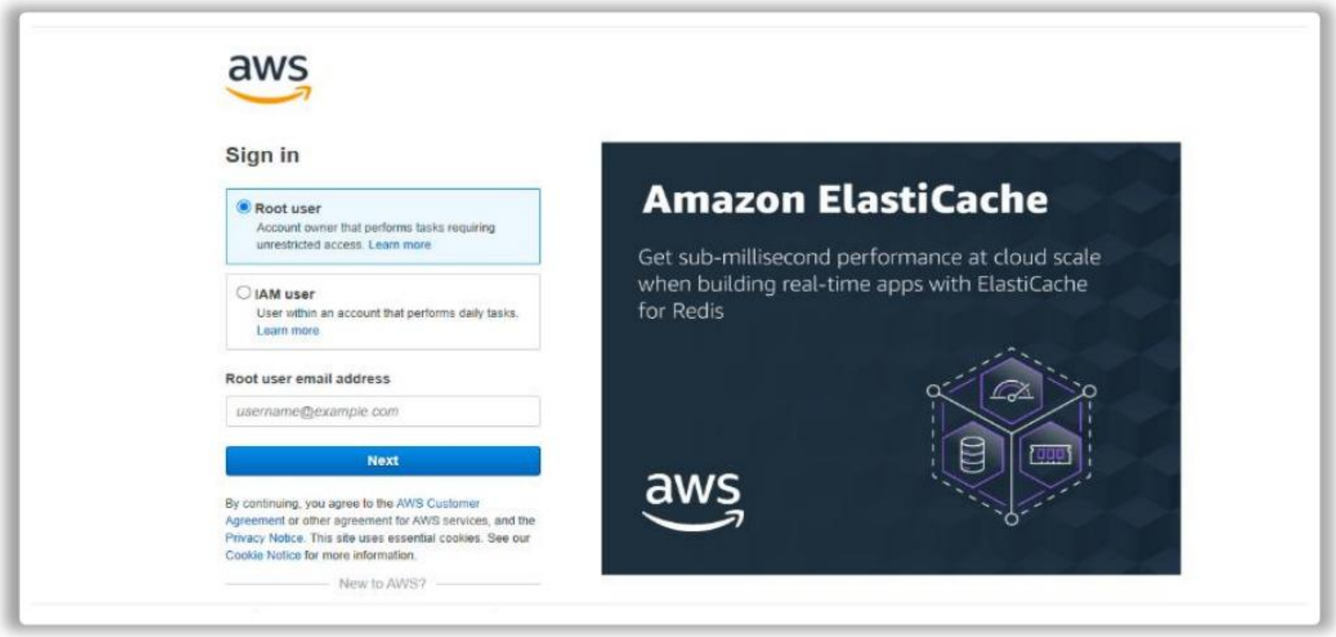


Figure 4.171: Sign in To Console

- From the console, navigate to the AWS backup service. Click on “Manage Backup.”

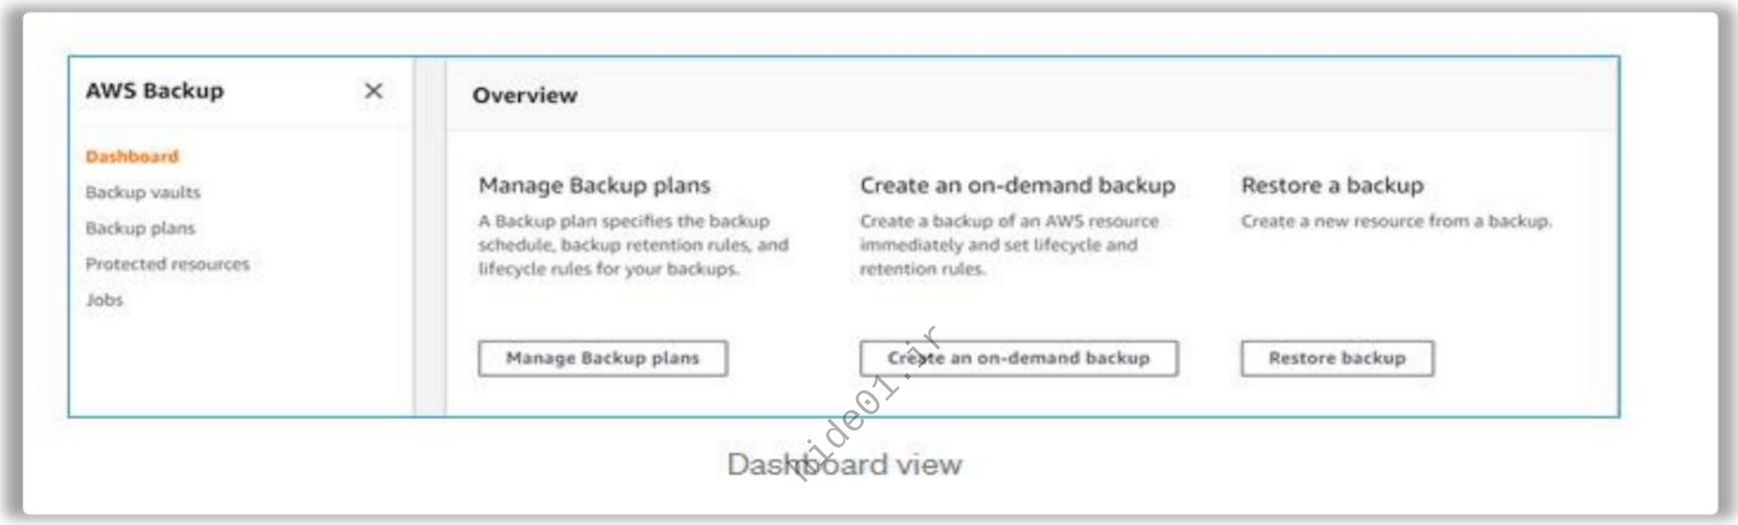


Figure 4.172: Select Manage Backup

- Now, select the “Build a new plan” option. Give a name to your backup plan.

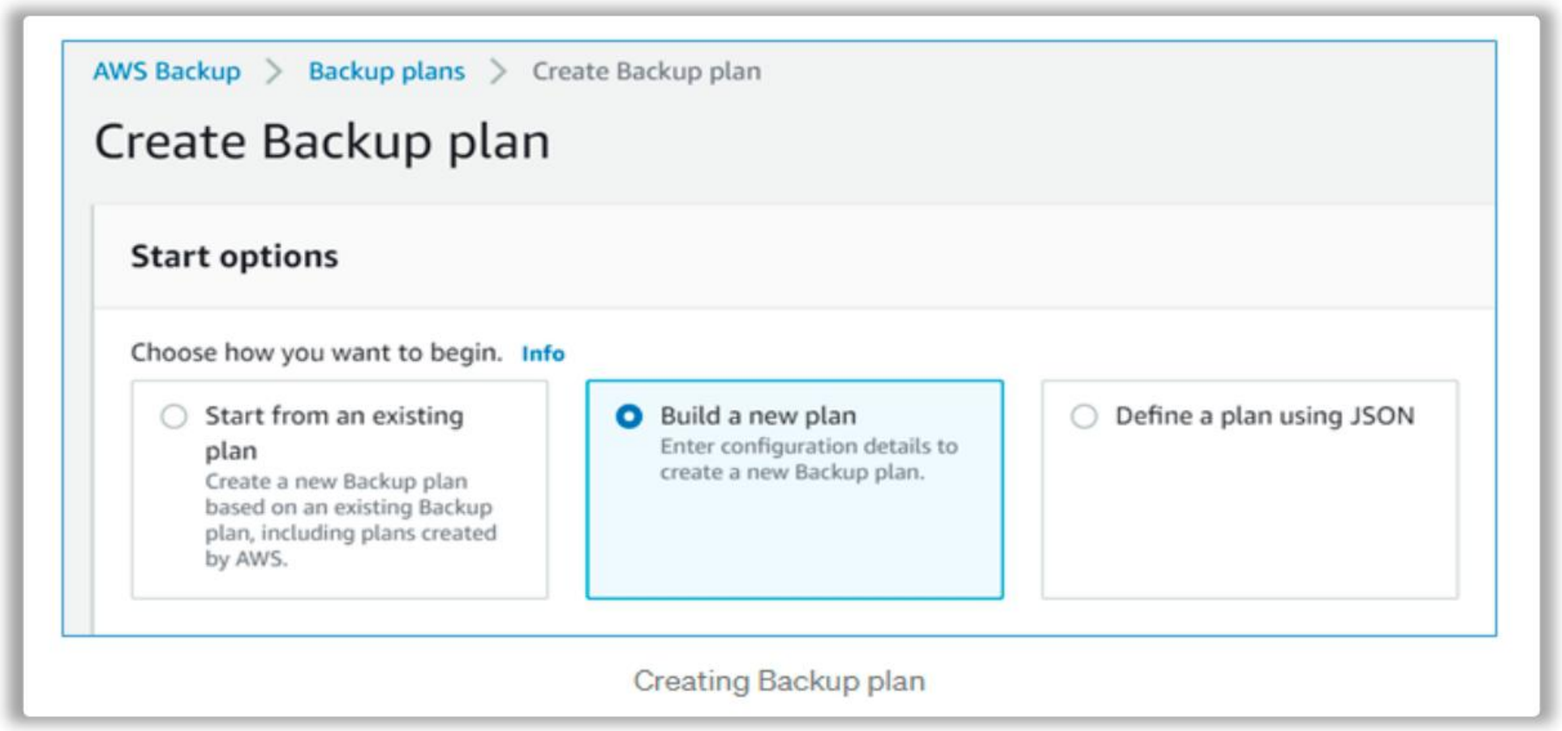


Figure 4.173: Select Build a new plan

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

- Choose a plan from the existing templates for scheduling backup and retention policy.

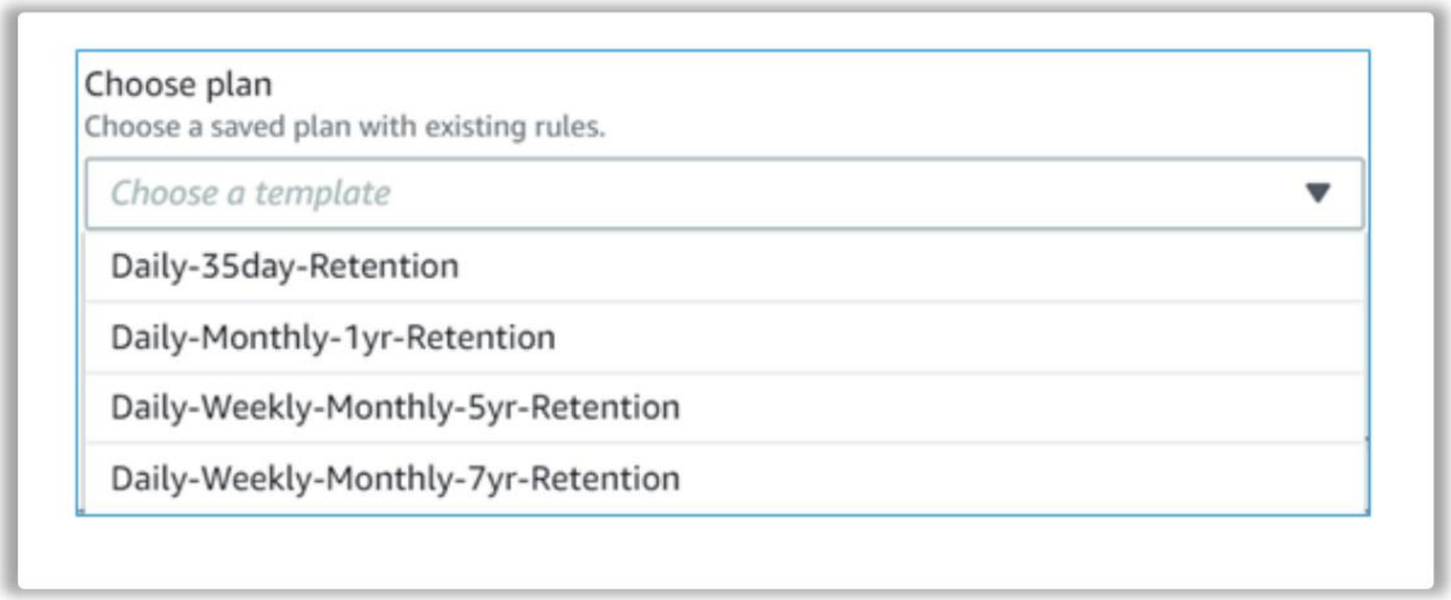


Figure 4.174: Choose Plan

- Under the "Backup rule configuration" section, give a name to your rule. You can add a backup rule by using frequency, and backup window. You can choose either the default backup window or you can set your own window by selecting "Customize backup window".

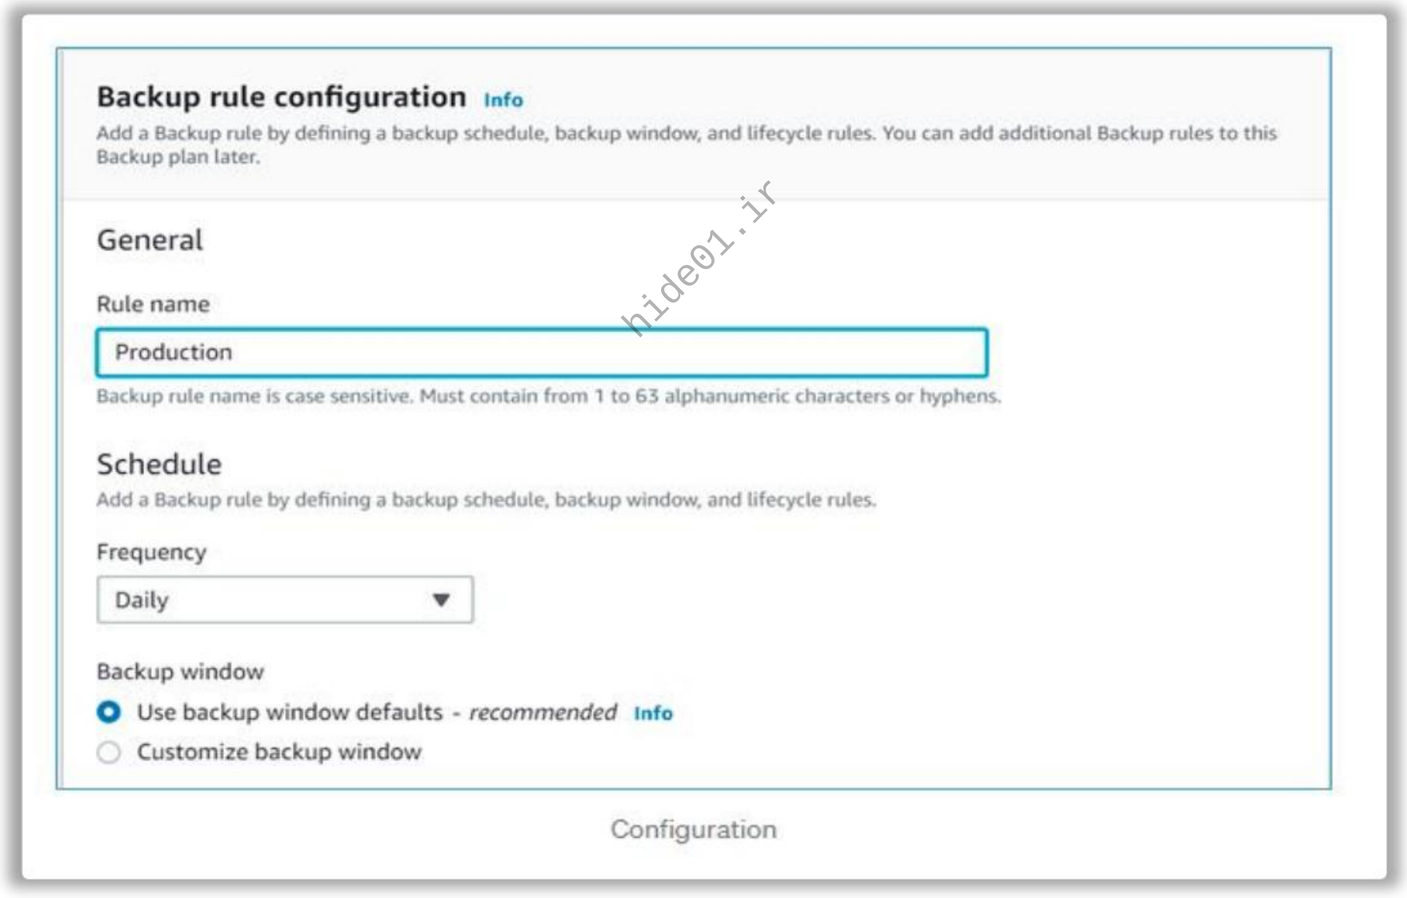


Figure 4.175: Choose Backup Rule Configuration

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

- Now, you can set a lifecycle configuration that sets the time for when the data should be moved to cold storage and its expiry time.

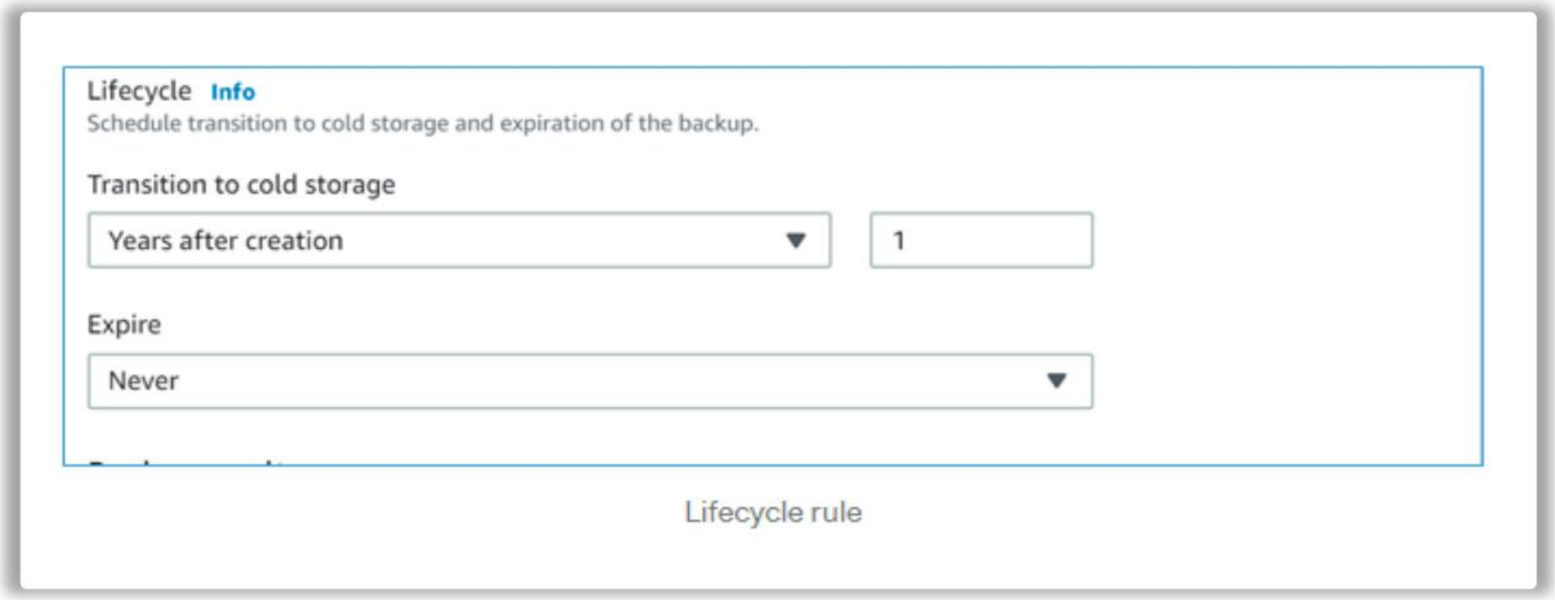


Figure 4.176: Select Lifecycle

- Now, select the "backup vault" to store backups in the default AWS storage vault. You can also create a new backup vault, or you can use the existing backup vault.

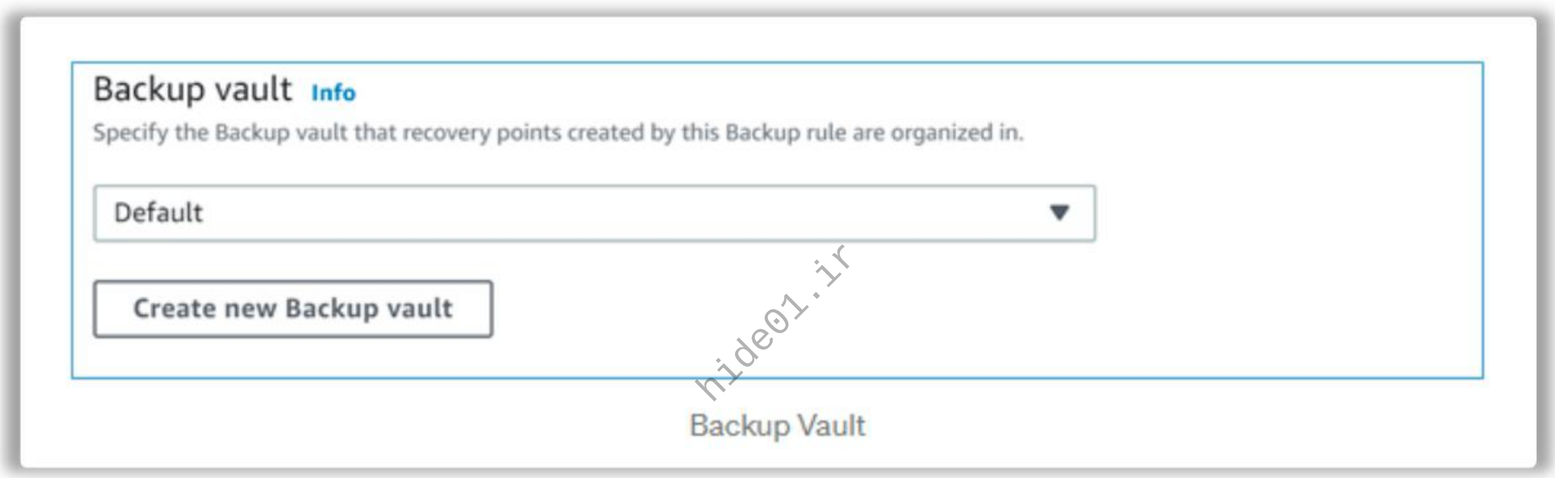


Figure 4.177: Select Backup Vault

- To keep a copy, you can select a location where your backup will be copied.

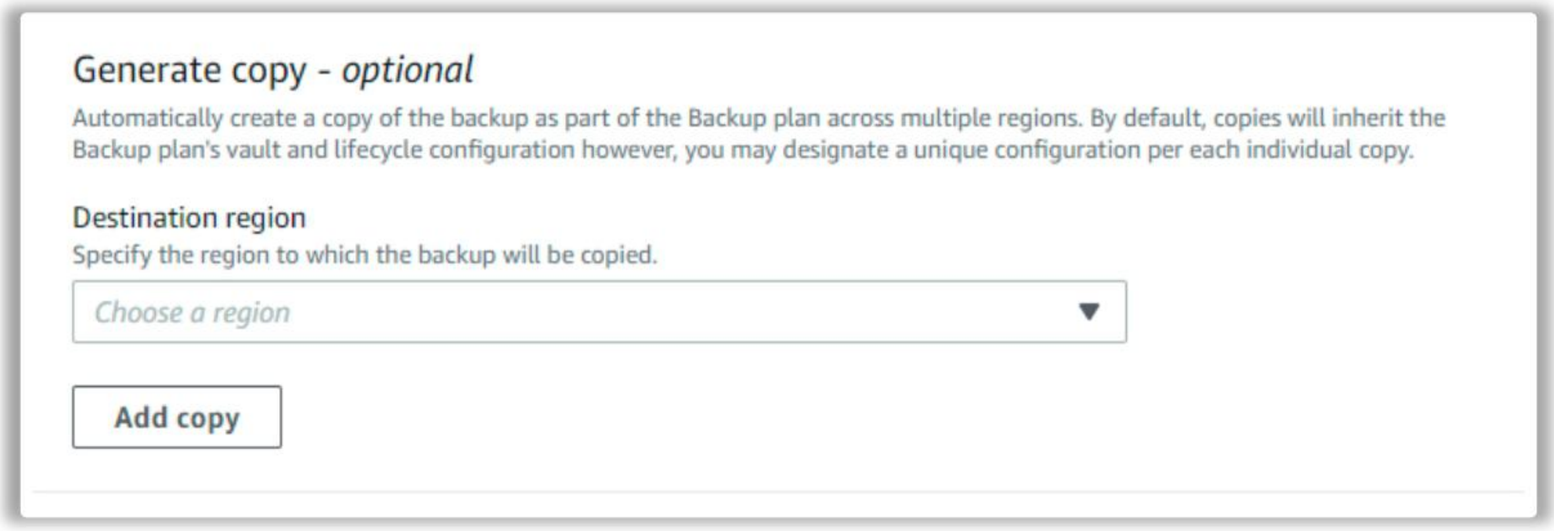


Figure 4.178: Select Location

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

- You can also add tags to the recovery points and backup plan, which is optional.

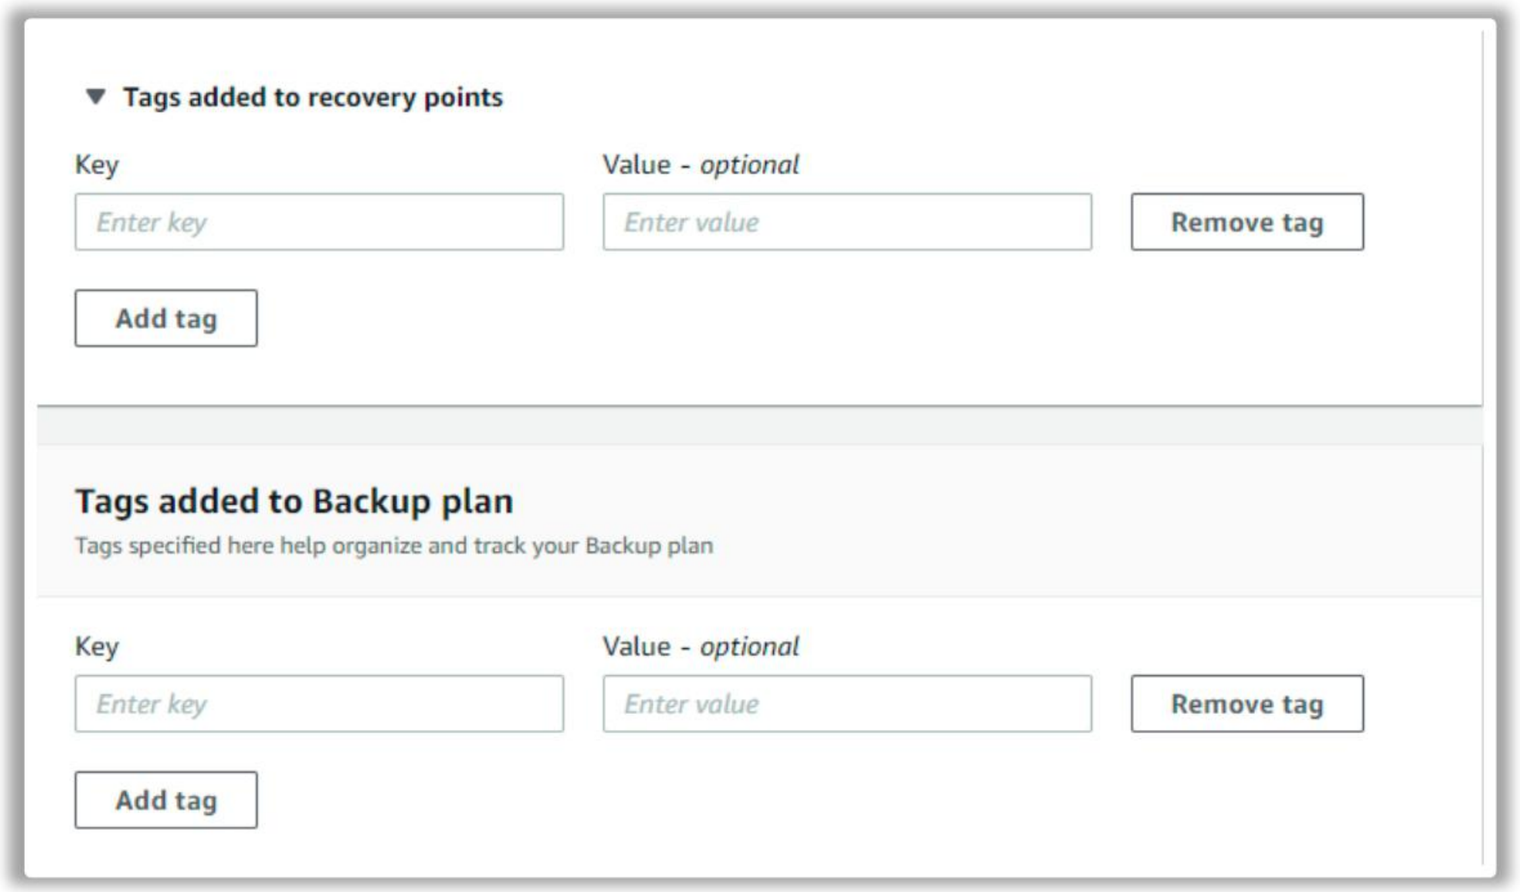


Figure 4.179: Add Tags

- Now, click on "Create Plan" and your backup plan is ready.

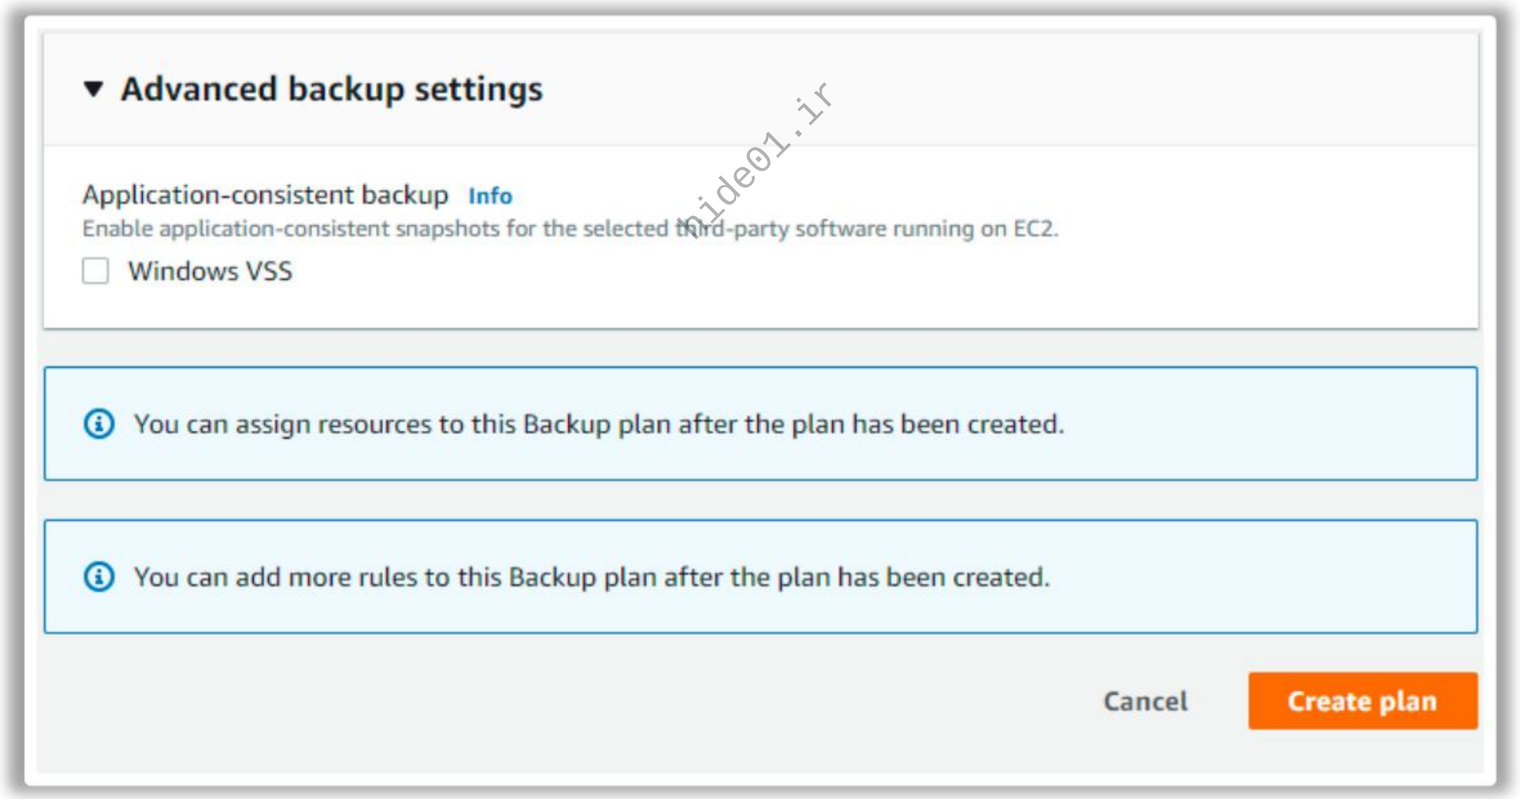


Figure 4.180: Click Create Plan

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

Prevent Deletion of Backups Using an Amazon Backup Vault Resource-based Access Policy

- By using **resource-based access policies** with AWS backup vault, the administrator can specify who can access the backup in the vault and what actions they can perform on that
- Using these policies, the administrator can also restrict the users from **deleting** the backup in the vault

```
{
            "Version": "2012-10-17",
              "Statement": [
             {
            "Sid": "statement ID",
              "Effect": "Deny",
               "Principal": "*",
              "Action": "backup:DeleteRecoveryPoint",
               "Resource": "*",
          "Condition": {
             "StringNotLike": {
                 "aws:userId": [
                     "arn:aws:iam::[
                       "AROAJCG2BTS7KMVVRB2ZC:",
                       "AIDAIECMYGLTVGIBTT5SY",
                       "343734213077"
                     ]
                ]
            }
        }
    ]
}
```

Amazon Backup Vault Resource-Based Access Policy

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

Prevent Deletion of Backups Using an Amazon Backup Vault Resource-Based Access Policy

With AWS backup, the cloud security administrator can assign policies to the backup vault and the related resources. These policies provide access to the users for creating backup plans and can also restrict their ability to delete the recovery points.

By using resource-based access policies with AWS backup vault, the administrator can specify who can access the backup in the vault and what actions they can perform on that. Using these policies, the administrator can also restrict the users from deleting the backup in the vault.

The use of resource-based access policies for Amazon Backup ensures that the backups are protected and can be accessed only by the right users.

To Prevent Deletion of Backups Using an Amazon Backup Vault Resource-based Access Policy

Source: https://aws.amazon.com

The following are the steps to prevent deletion of backup vault using Amazon Backup vault resource-based access policy:

- Sign into the AWS Management Console
- Now, navigate to the AWS Backup console
- On the left navigation panel, select the “Backup Vaults”.
- Select the required backup vault from the list.
- Under the access policy section, mention the below JSON script. This policy restricts users who are not principal from deleting the recovery points in the backup vault. You can replace the statement id and aws: user id with the values of your environment.

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

```
{
          "Version": "2012-10-17",
            "Statement": [
            {
           "Sid": "statement ID",
              "Effect": "Deny",
               "Principal": "*",
              "Action": "backup:DeleteRecoveryPoint",
               "Resource": "*",
         "Condition": {
            "StringNotLike": {
                "aws:userId": [
                    "arn:aws:iam::[
                      "AROAJCG2BTS7KMVVRB2ZC:",
                      "AIDAIECMYGLTVGIBTT5SY",
                      "343734213077"
                    ]
                ]
            }
        }
    }
  ]
}
```

hide01.ir

- If you want to restrict the access to a specific resource type, instead of using “Resource” : “*”, you can mention the recovery point to be denied.
- Now, click on “Attach policy.”

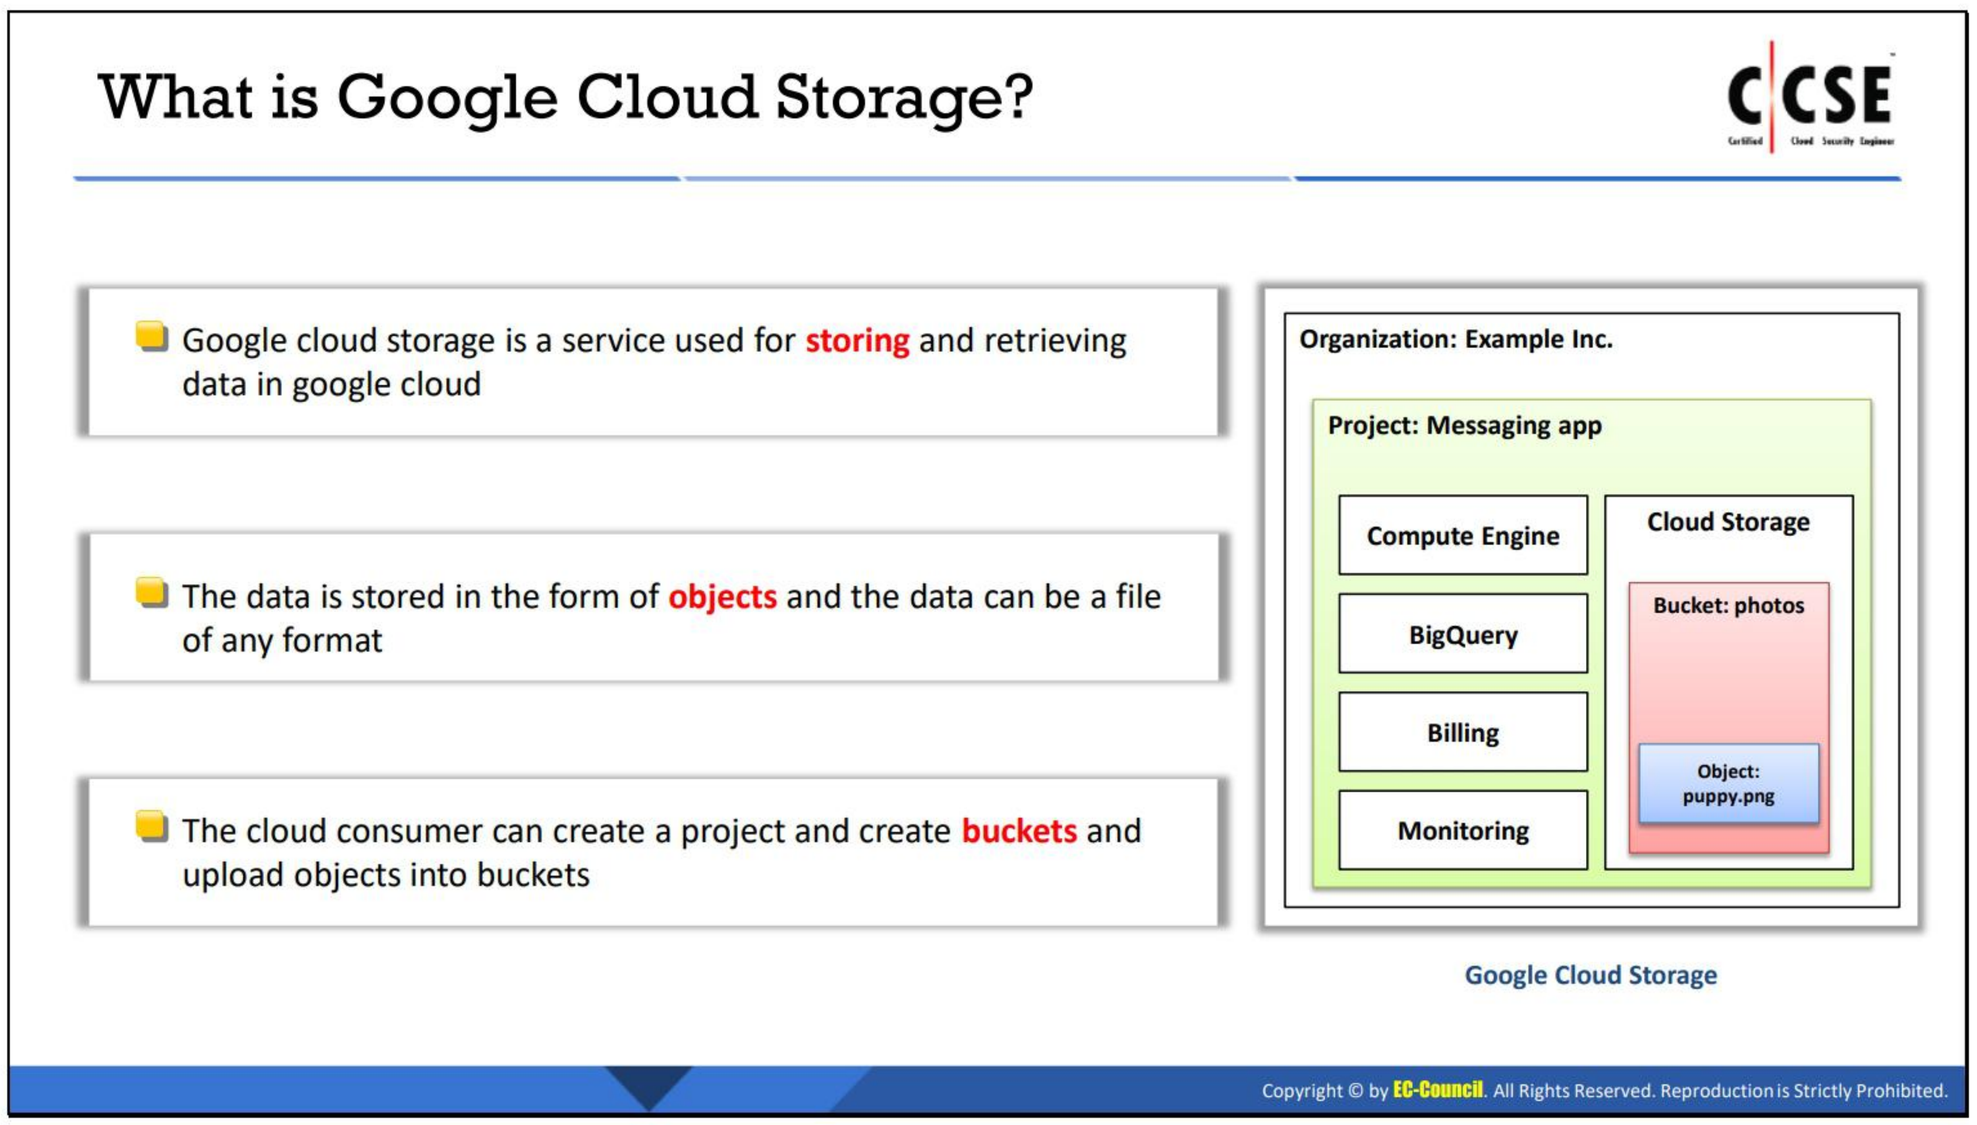
Certified Cloud Security Engineer Copyright © by **EC-Council**
All Rights Reserved. Reproduction is Strictly Prohibited.



LO#11: Understand Storage Services in Google Cloud Platform (GCP)

The objective of this section is to provide a proper understanding of Google Cloud Storage. including features, key terms, interacting tools, and Google Cloud Storage types. This section also includes other Google Storage services namely Data Transfer service, Transfer Appliance, Firebase, and Google Workspace.

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



What is Google Cloud Storage?

Google Cloud Storage is a cloud storage service used for storing and retrieving data on the Google Cloud Platform. Google Storage is an IaaS and it comes with advanced security and sharing abilities. Google Storage is similar to Amazon S3 service and is highly suitable for companies because of its different service specifications. It can be used to store both primary and least accessed data. It can be used for various scenarios such as storing website content or distributing data to multiple users in different locations.

The data is stored in the form of objects and the data can be a file of any format. All the objects are stored in containers called buckets. Based on the project requirement, the customer can create any number of buckets.

Hierarchy of Google Cloud Storage

An example of Google Cloud Storage hierarchy is included below.

- **Organization:** The company named "Test Inc." creates a Google Cloud organization called Test.org
- **Projects:** The above company builds various applications that belong to various projects. Each one of the projects will have its own APIs and other required resources.
- **Buckets:** Buckets are the containers used for storing data in the form of objects. A project can contain any number of buckets for storing different types of data. For example, an "audio" bucket to store all the audio files and a "video" bucket to store all the video files of the project.
- **Objects:** Objects are the files of any format such as images, videos, audio files, etc.

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

An example of Google Cloud Storage is appended below:

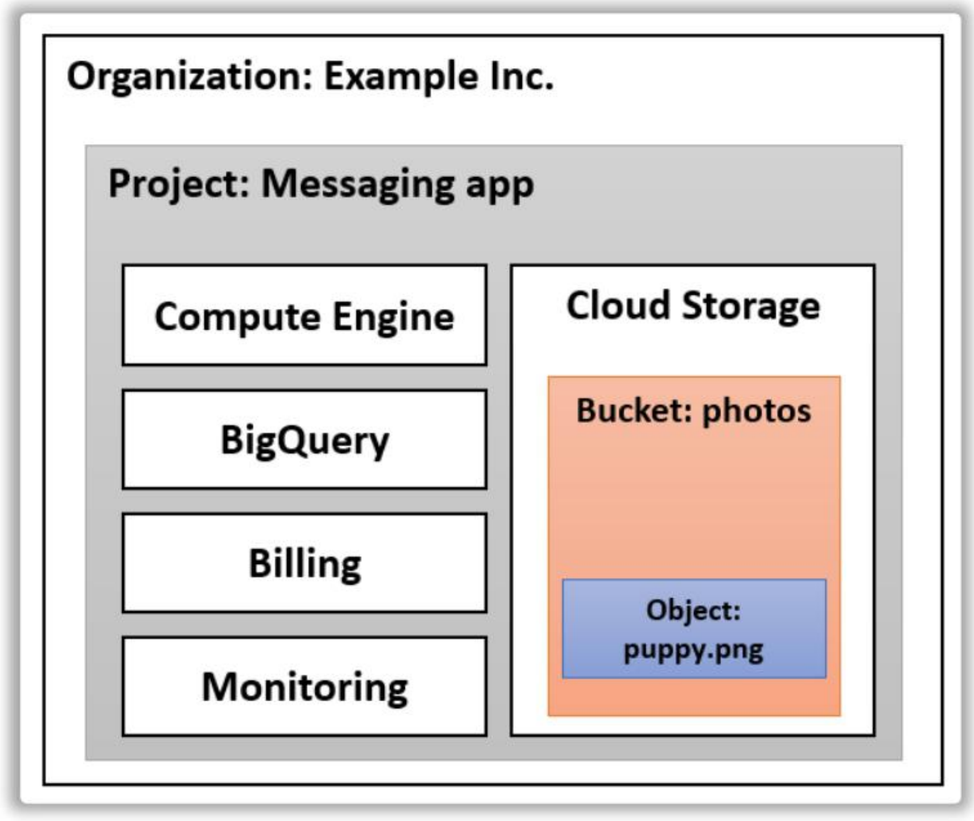
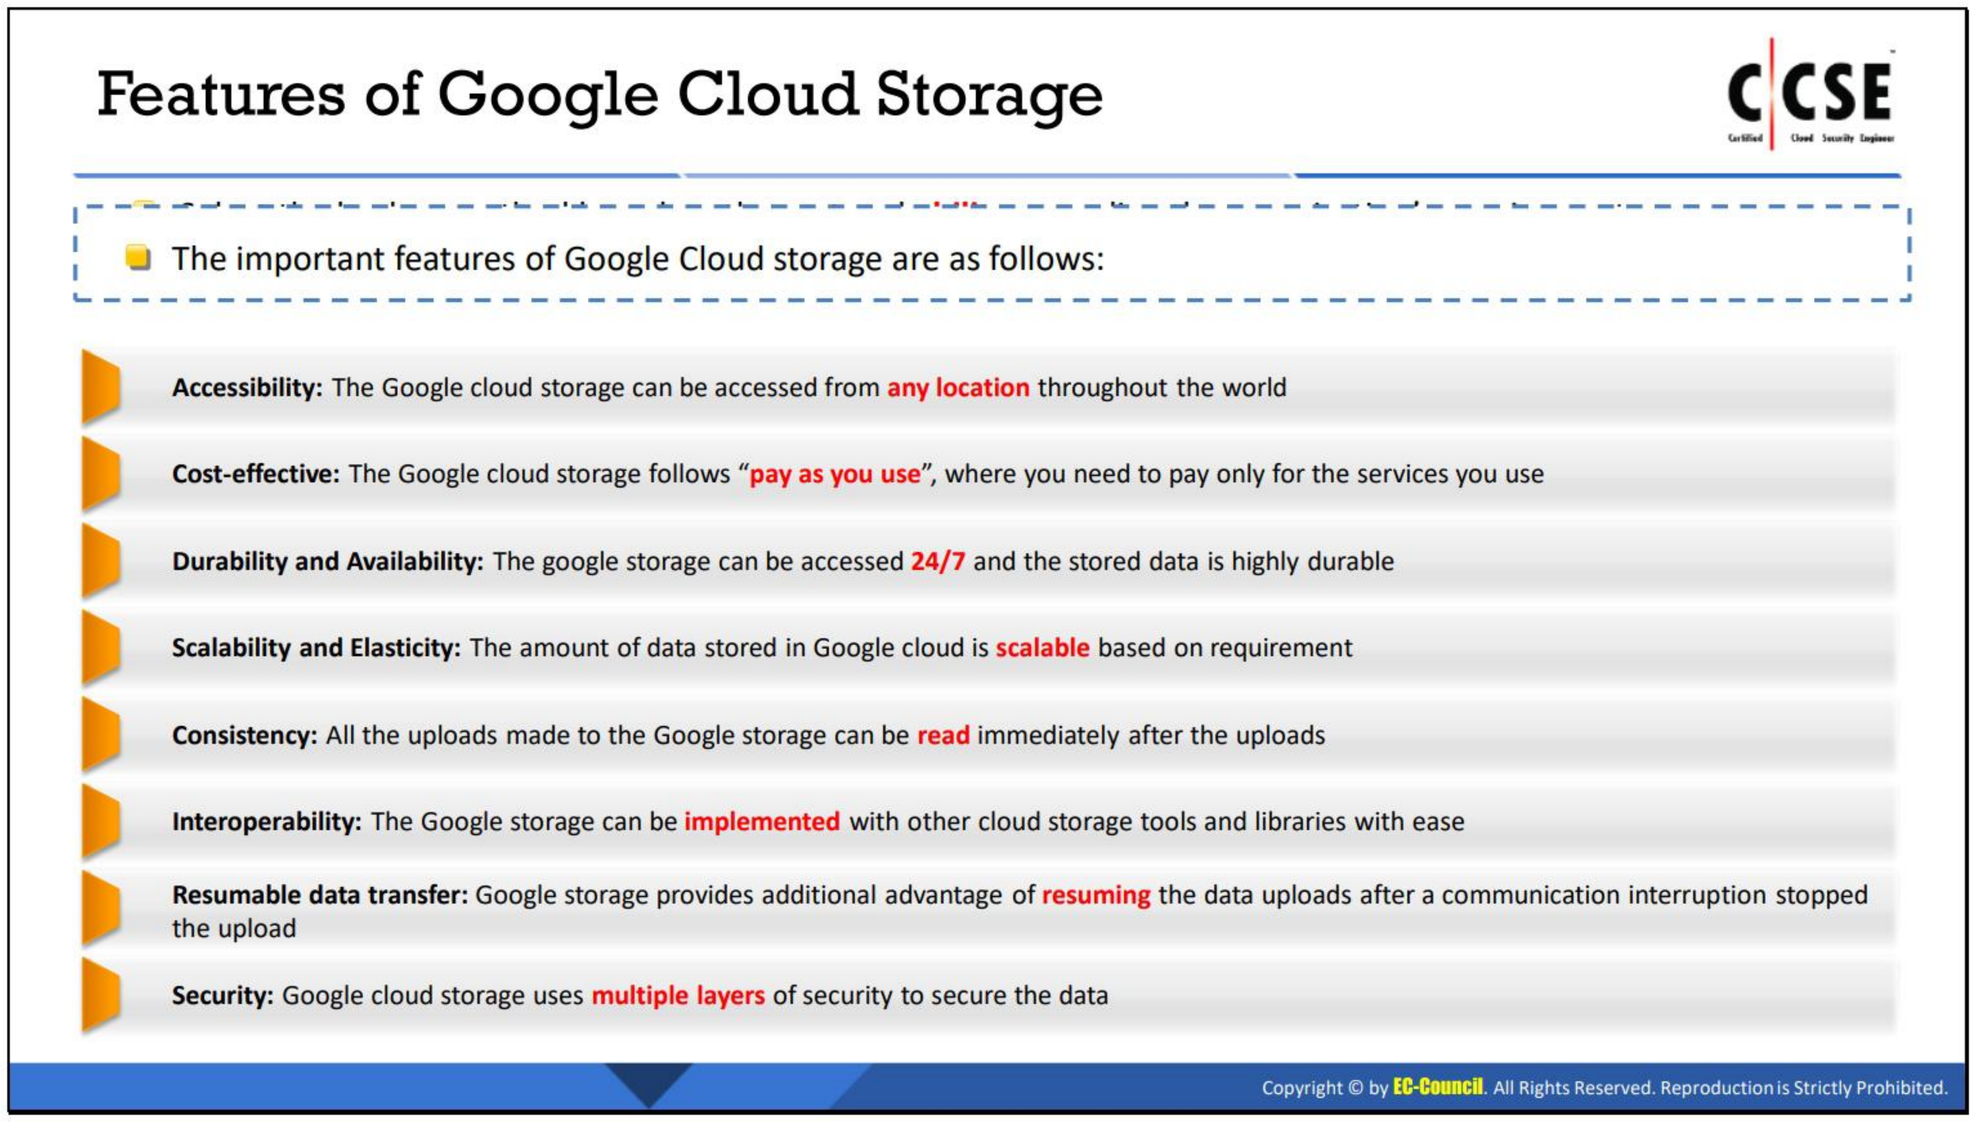


Figure 4.181: Hierarchy of Google Cloud Storage

hide01.ir

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



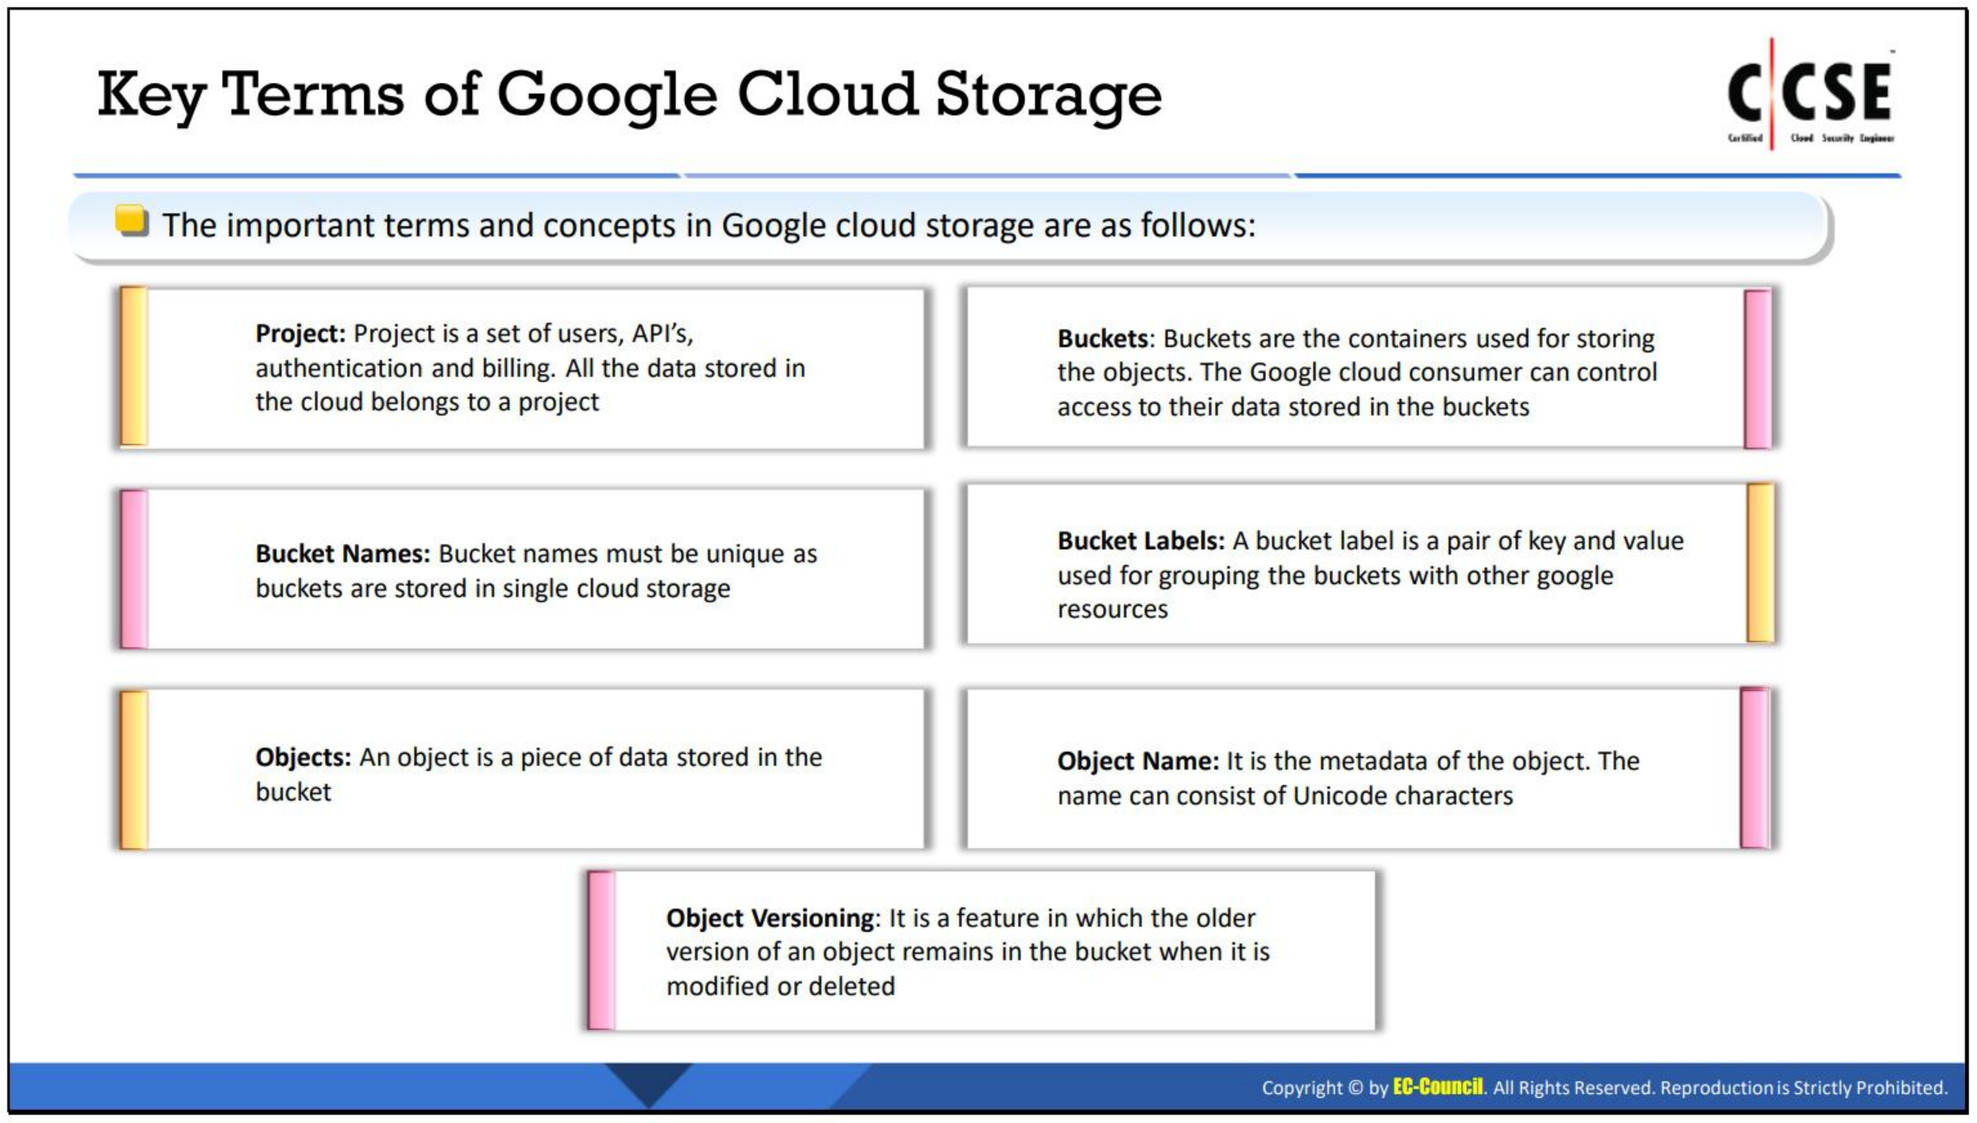
Features of Google Cloud Storage

The important features of Google Cloud Storage are described below.

hide01.ir

- **Accessibility:** Google Cloud Storage can be accessed from any location throughout the world
- **Cost-effective:** It follows the "pay as you use" model, where you need to pay only for the services you use
- **Durability and availability:** Google Cloud Storage can be accessed 24/7 and the data is highly durable.
- **Scalability and elasticity:** The amount of data you store in Google Cloud is scalable based on your requirement.
- **Consistency:** All the uploads made to Google Cloud Storage can be read immediately after the uploads.
- **Interoperability:** Google Storage can be implemented with other cloud storage tools and libraries.
- **Resumable data transfer:** Google Cloud Storage provides the additional advantage of resuming data uploads in case a communication interruption stops them.
- **Security:** Google Cloud Storage uses multiple layers of security—from physical security to advanced cybersecurity technologies.

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



Key Terms of Google Cloud Storage

The important terms and concepts in Google Cloud Storage are described below.

ide01.ir

Project

The data stored in Google Cloud Storage belongs to a project. Every project has a set of APIs, users, billing, authentication, and billing.

Buckets

Buckets are the containers used for storing data. Any data that is stored in Google Cloud must be stored in buckets. Unlike the folders and directories, a Google Cloud user can assign their own set of permissions to users and control access to the data stored in the buckets.

While creating a bucket, one needs to assign a unique name, location where the data need to be stored, and a default storage class. The default storage class can still be modified after bucket creation; however, the bucket name and storage location cannot be modified. The name of a deleted bucket can be used by anyone while creating a new bucket.

Bucket Names

All the bucket names must be globally unique and are publicly visible. The bucket names must be unique as all the storage buckets belong to a single Google Cloud Storage.

Certain guidelines need to be followed for naming a bucket in Google Cloud Storage such as:

- The bucket name can only contain lowercase letters, numbers, dashes, dots, and underscores.
- No spaces are allowed.

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

- All your bucket names must begin and end with either letters or numbers.
- The bucket name must be between 3 to 63 characters.
- Bucket names with dots can extend up to 222 characters.
- Bucket names with dots cannot be represented in dotted decimal notation.
- Bucket names should not begin with “goog”.

Examples of bucket names: abcd_123, test-234, 097ab-09872gre-0000edf

Bucket Labels

A bucket label is a key-value pair used for grouping the buckets with other Google resources. For instance, you can create a “group” as a key with value such as developers, testers, admins. You can assign labels such as group: java, group: developers, and group: administrators to different buckets to indicate the users that are associated with each bucket.

There are few things to remember about bucket labels, namely:

- A bucket can have multiple bucket labels but not more than 64 labels per bucket.
- Key and value should not exceed more than 63 characters.
- Label keys cannot be empty, and a key must begin with a lowercase letter.
- Bucket labels are not associated with specific objects.

ide01.ir

Objects

An object is a piece of data stored in the bucket. Any number of objects can be created in a bucket. Objects consist of 2 components, namely:

- **Object data**: Object data is nothing but the file that is stored in Google Cloud Storage.
- **Object metadata**: The object metadata is a name and value pair for describing the object type.

Object Names

An object name is a piece of object metadata. The object name must be unique within a bucket and must contain a combination of Unicode characters. Objects in different buckets can have the same name.

The following requirements need to be ensured while naming an object:

- The length of the object name should not exceed 1024 bytes.
- Object names should not contain any carriage return or line feed characters.
- Object names cannot start with ‘.’ and must not be named as “.” or “..”.

Object Versioning

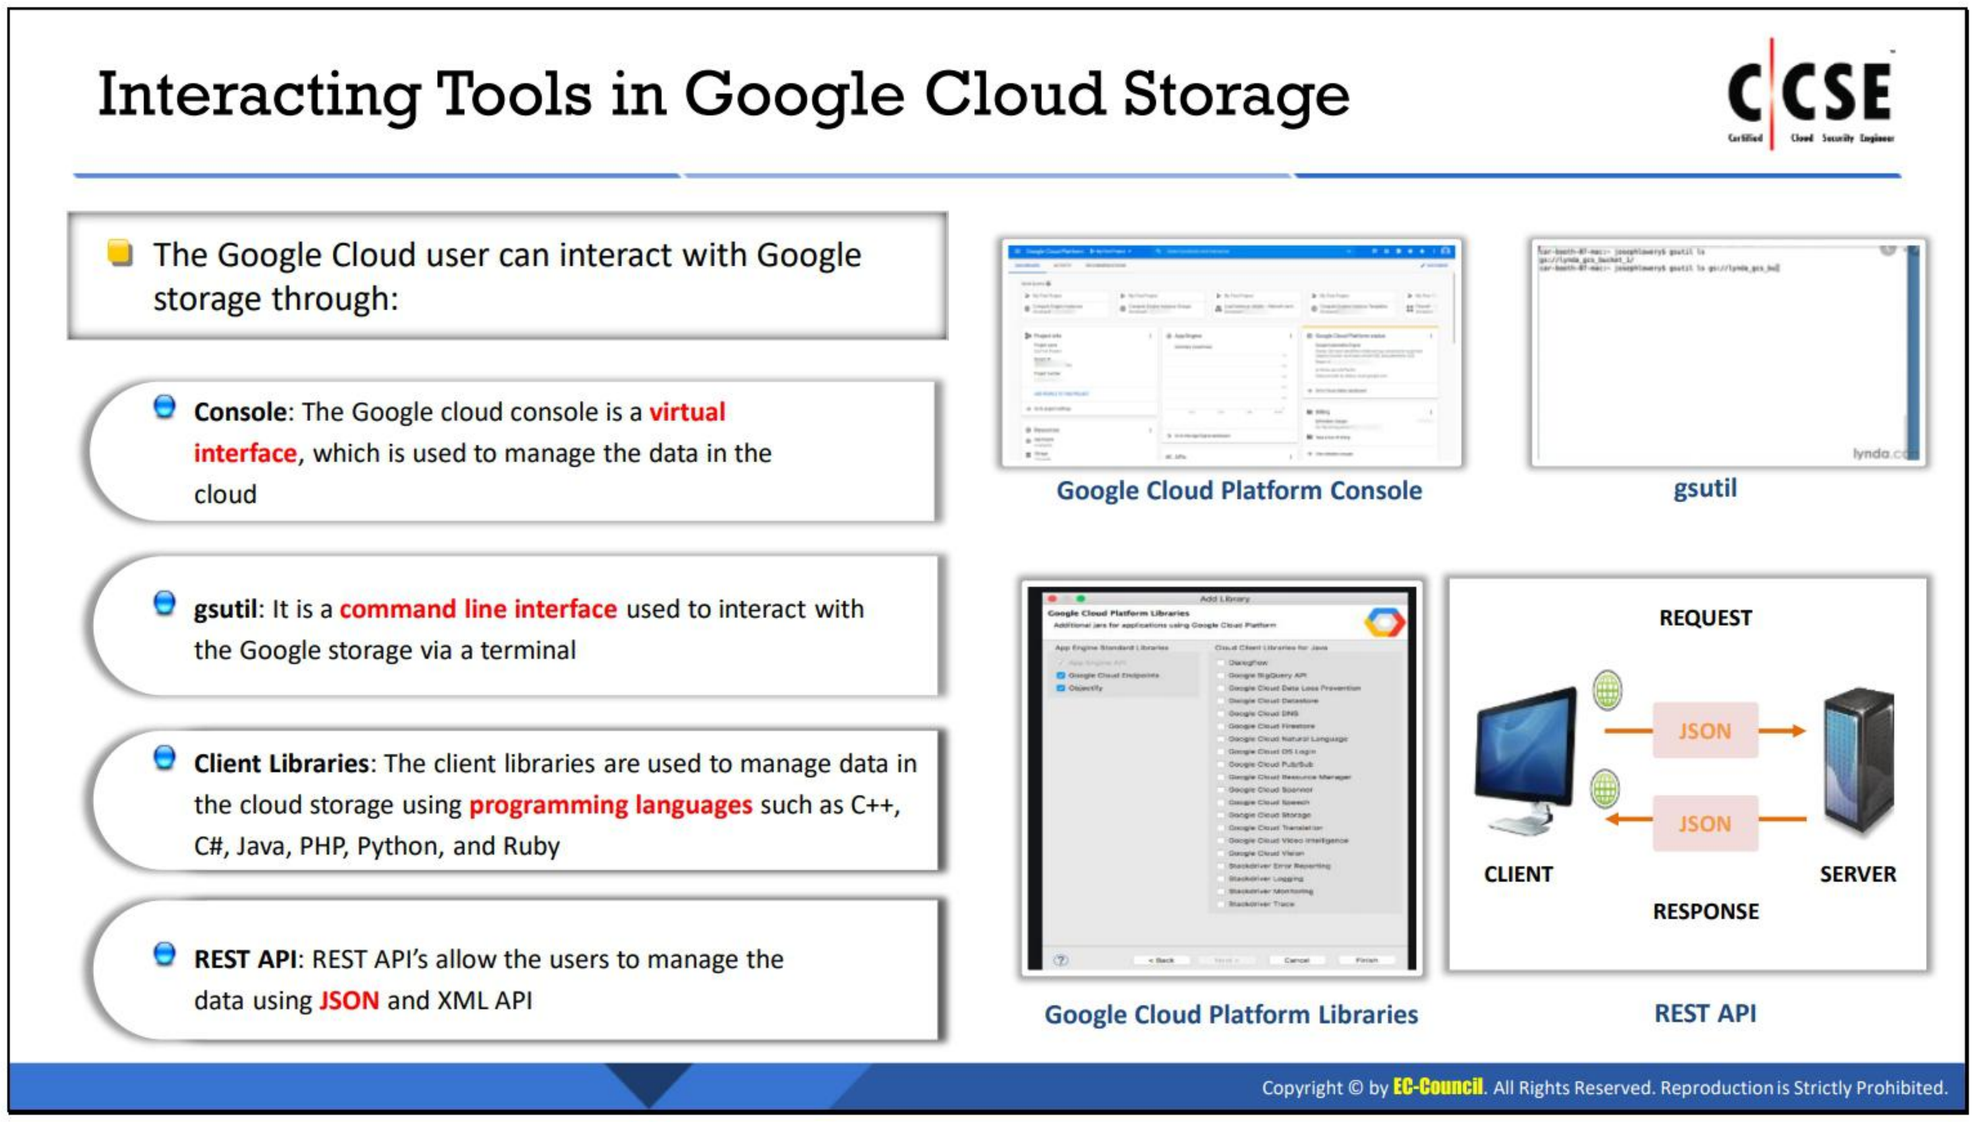
By default, when an object is replaced or updated, Google Cloud Storage deletes the older version and updates it with the new object; however, when object versioning is enabled on a bucket, the older versions remain in the bucket in case of any replacements or deletions.

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

Each object version has its generation number. Using the generation number, you refer to the older versions of an object. This enables the Google Cloud user to access older versions of an object or delete them if they are no longer required.

hide01.ir

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



Interacting Tools for Google Cloud Storage

A Google Cloud user can interact with Google Storage through the following:

Source: https://cloud.google.com

hide01.ir

Google Cloud Console

The Google Cloud Console is a virtual interface that is used to manage the data in the Google Cloud. Using the console, a Google Cloud user can create a bucket, upload objects into the buckets, and can control access to the buckets.

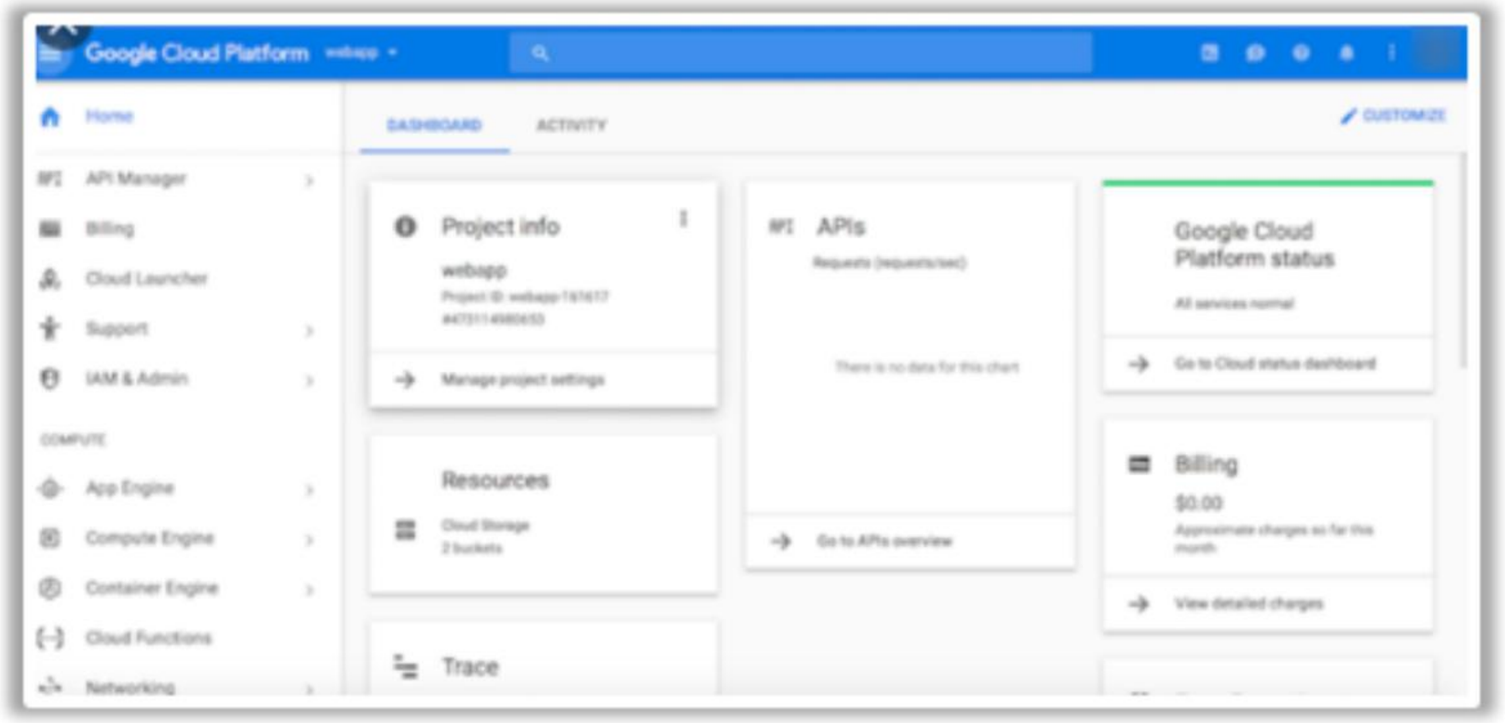


Figure 4.182: Google Cloud Console

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

gsutil

It is a Python language command-line interface through which a cloud practitioner can interact with Google Cloud Storage via a terminal. All the operations in gsutil are performed using the transport layer and HTTPS security.

Using gsutil, various object and bucket management tasks can be performed, including:

- Deletion and creation of Google Storage buckets
- Adding, deleting, and downloading objects in the buckets
- Renaming, copying, and deleting the objects
- Managing bucket ACLs

Client Libraries

The client libraries are used to manage data in Google Cloud Storage using programming languages such as C++, C#, Java, PHP, Python, and Ruby. Using the client library, a cloud user can interact with Google Cloud Storage using their preferred programming language.

The client libraries provide the code and make it easy for the cloud users to interact. Cloud users can interact with the server to handle low-level details of communication. The client library can be easily installed with tools such as pip and npm.

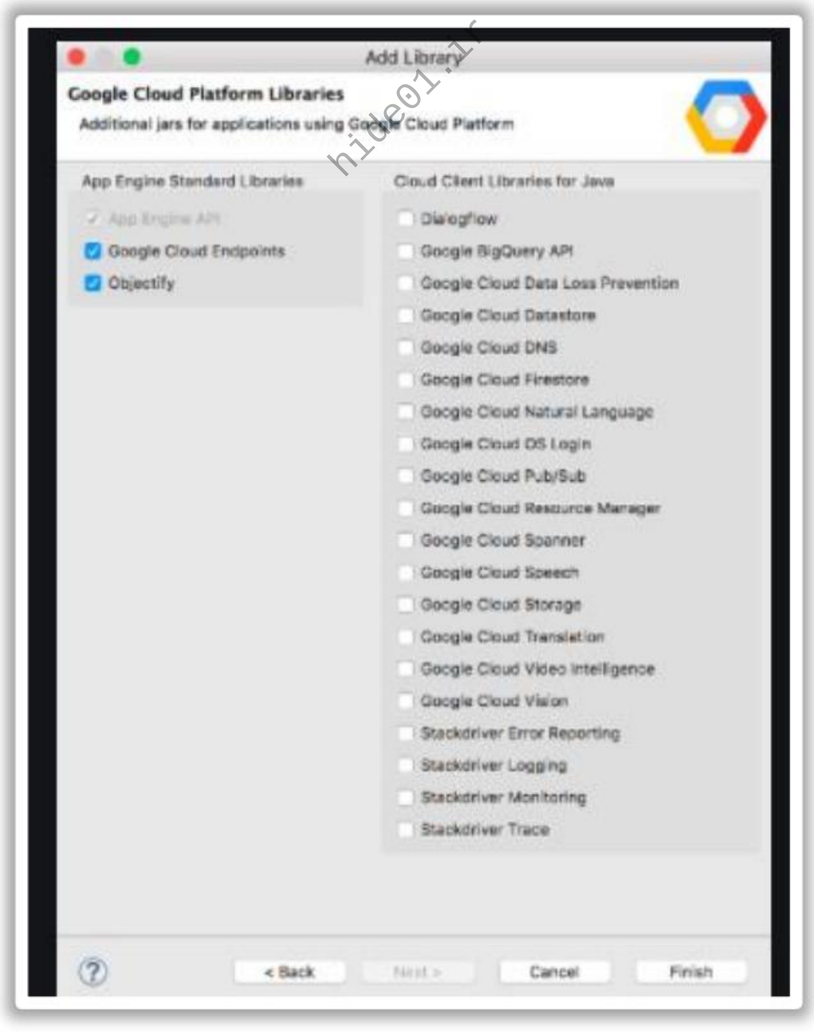


Figure 4.183: Google Cloud Platform Libraries

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

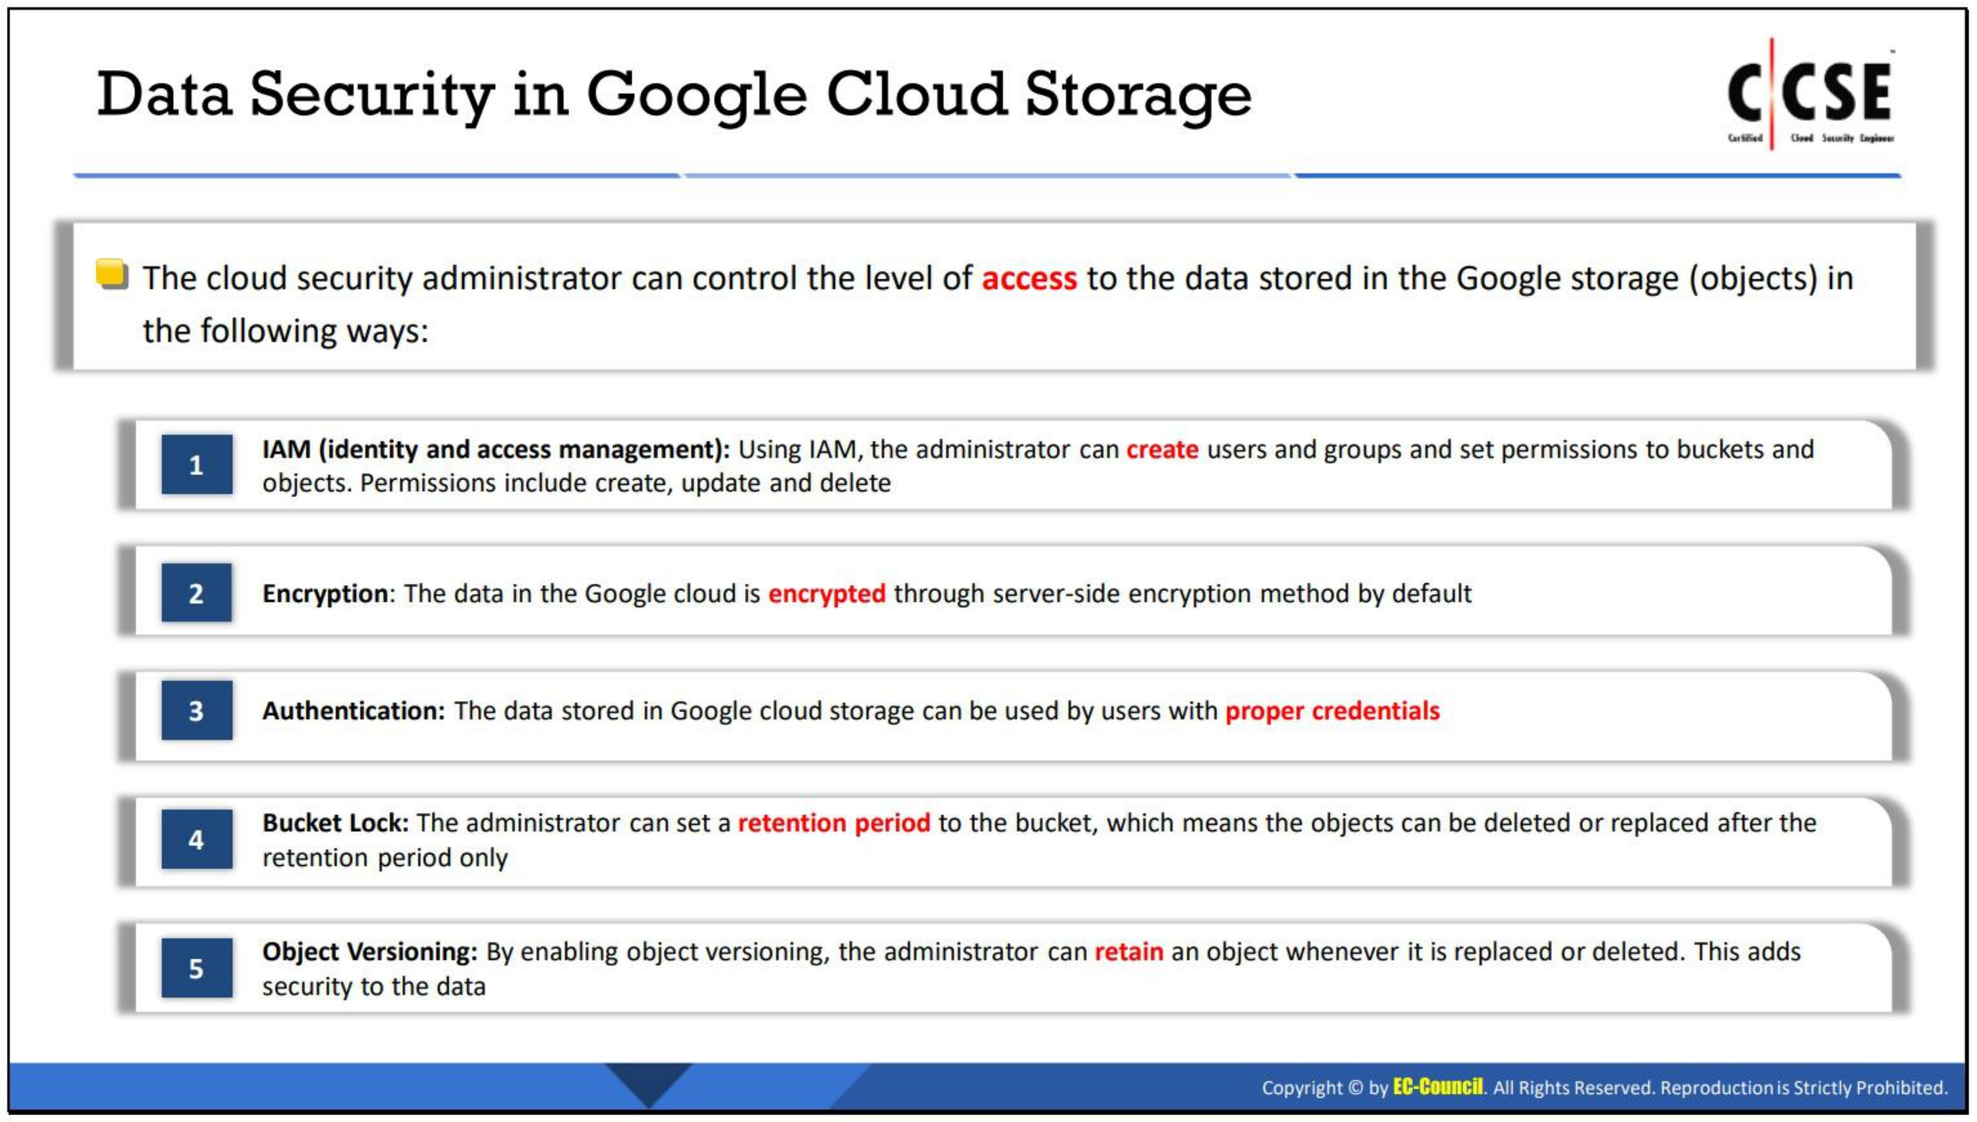
REST API

REST APIs allow the cloud practitioner to manage data using JSON and XML API.

- **JSON API:** The JSON API is simple and used in handling Google Cloud Storage projects through programming. It is mainly used by software developers and they need to be familiar with web programming.
- **XML API:** The XML API is a RESTful interface used by software developers. To use this Google Cloud Storage API, the cloud user must be familiar with RESTful services and web programming. It uses HTTP requests to manage the Google Cloud data. The XML API is useful in case of migration from some other storage to Google Cloud Storage. The migration can be done by making a few changes in the existing libraries to interact with Google Cloud Storage.

hide01.ir

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



Data Security in Google Cloud Storage

The cloud security administrator can control the level of access to the data stored in Google Storage (objects) in the below-mentioned ways.

hide01.ir

- **IAM:** Using IAM, the cloud security professional can control access to Google Cloud resources such as objects and buckets. Permissions can be set to users and groups using IAM policies.

  An IAM policy defines the set of actions that users can perform on objects and buckets. The actions include bucket creation, deletion, read, write, or update permissions on objects in the project.

  For example, the cloud administrator can create an IAM policy to give bucket administration access to one of the users and another IAM policy for viewing the objects in the buckets.

  Using IAM, the cloud administrator can easily set permissions to buckets and objects without changing the ACL.

  - **Permissions**: Permissions are the set of actions that can be performed on Google Cloud Storage.
  - **Roles**: A role is a set of permissions that can be assigned to members for objects and buckets in Google Cloud.
  - **Members**: Members are either groups or users to whom the roles are assigned.

- **Encryption**: By default, the data in Google Cloud is encrypted using the SSE method. The cloud security professional can incorporate additional security by choosing other

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

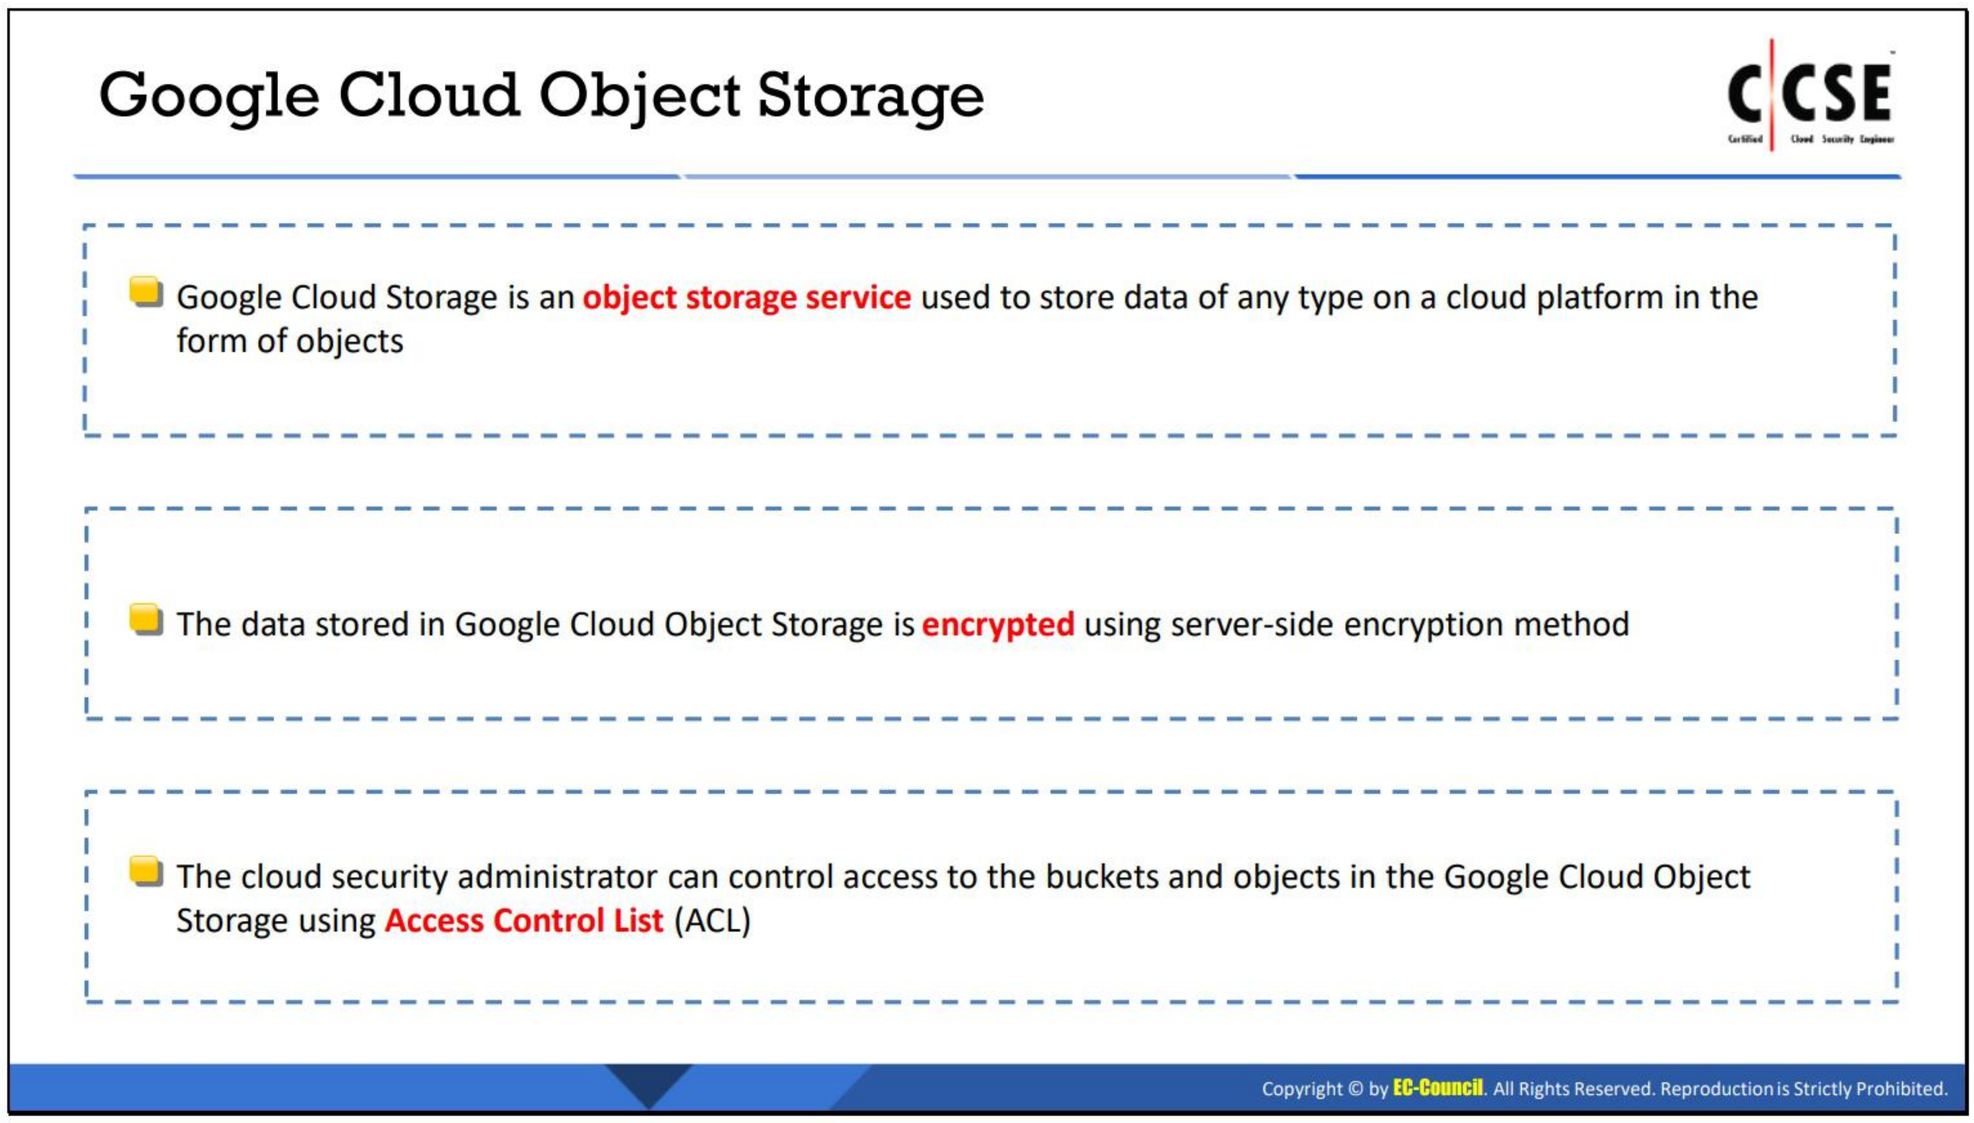
encryption methods such as customer-managed encryption keys and customer-supplied encryption keys.

- **Authentication:** All the actions performed on Google Cloud are authenticated. Authentication is the process of determining the identity of a user to perform actions on Google Cloud resources.
- **Bucket lock**: The cloud user can set the duration of how long an object should stay in the bucket using the retention policy. The bucket lock feature is useful for users who need to write an object once and read its data many times. It prevents the users from modifying or deleting the data for a particular time.
- **Object versioning**: Google Cloud Storage provides additional security by enabling users to retrieve the deleted or replaced object in the storage bucket. This feature can be availed by enabling "object versioning." By enabling this feature, all the older versions of an object are stored in the bucket in case of any replacement or deletion of an object.

All the older versions of the objects have the same object name but will have a different generation number. As mentioned previously, cloud users can delete the older versions of the objects by using the generation number if they are no longer required.

hide01.ir

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



Google Cloud Object Storage

As briefly mentioned earlier, Google Cloud Storage is an object storage service used to store data of any type on a cloud platform in the form of objects. An object can be of any format such as audio, video, text files, etc. All the objects are stored in containers called buckets.

hidoo1.ir

The way how Google Cloud Object Storage works is simple for the end-users. Cloud users can store a small object and can expand its size up to the terabyte scale. The cloud security professional can store an unlimited amount of data in the form of objects without any minimum object size. The objects are accessible globally and the service offers high durability.

The objects are stored in a version format. This means that while replacing the newly modified file with the old file, the previous version of that file will be stored in a backup instead of deleting it. This ensures that cloud users are free from the worry of losing their data.

Google Storage provides four different types of storage classes that a cloud user should choose for the buckets. The storage classes vary in pricing and storage. If the cloud security professional wants to store the data that needs to be accessed on regular basis, they can choose the Standard storage class and can choose to store the buckets in a single geographic location or in different geographic locations across the globe. This offers high availability and high performance as the data is stored in multiple regions. Cloud users can select Nearline storage and Coldline storage if the data does not need to be accessed frequently and it needs to be stored for a longer duration. For Nearline storage, the minimum storage time is 30 days whereas it is 90 days for Coldline storage class.

Over time, the data stored in the cloud may require different storage forms. Google Storage provides this option using lifecycle management rules. For example, if the buckets are used to store application logs, the data need to be available only for the first few months. Later, the data

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

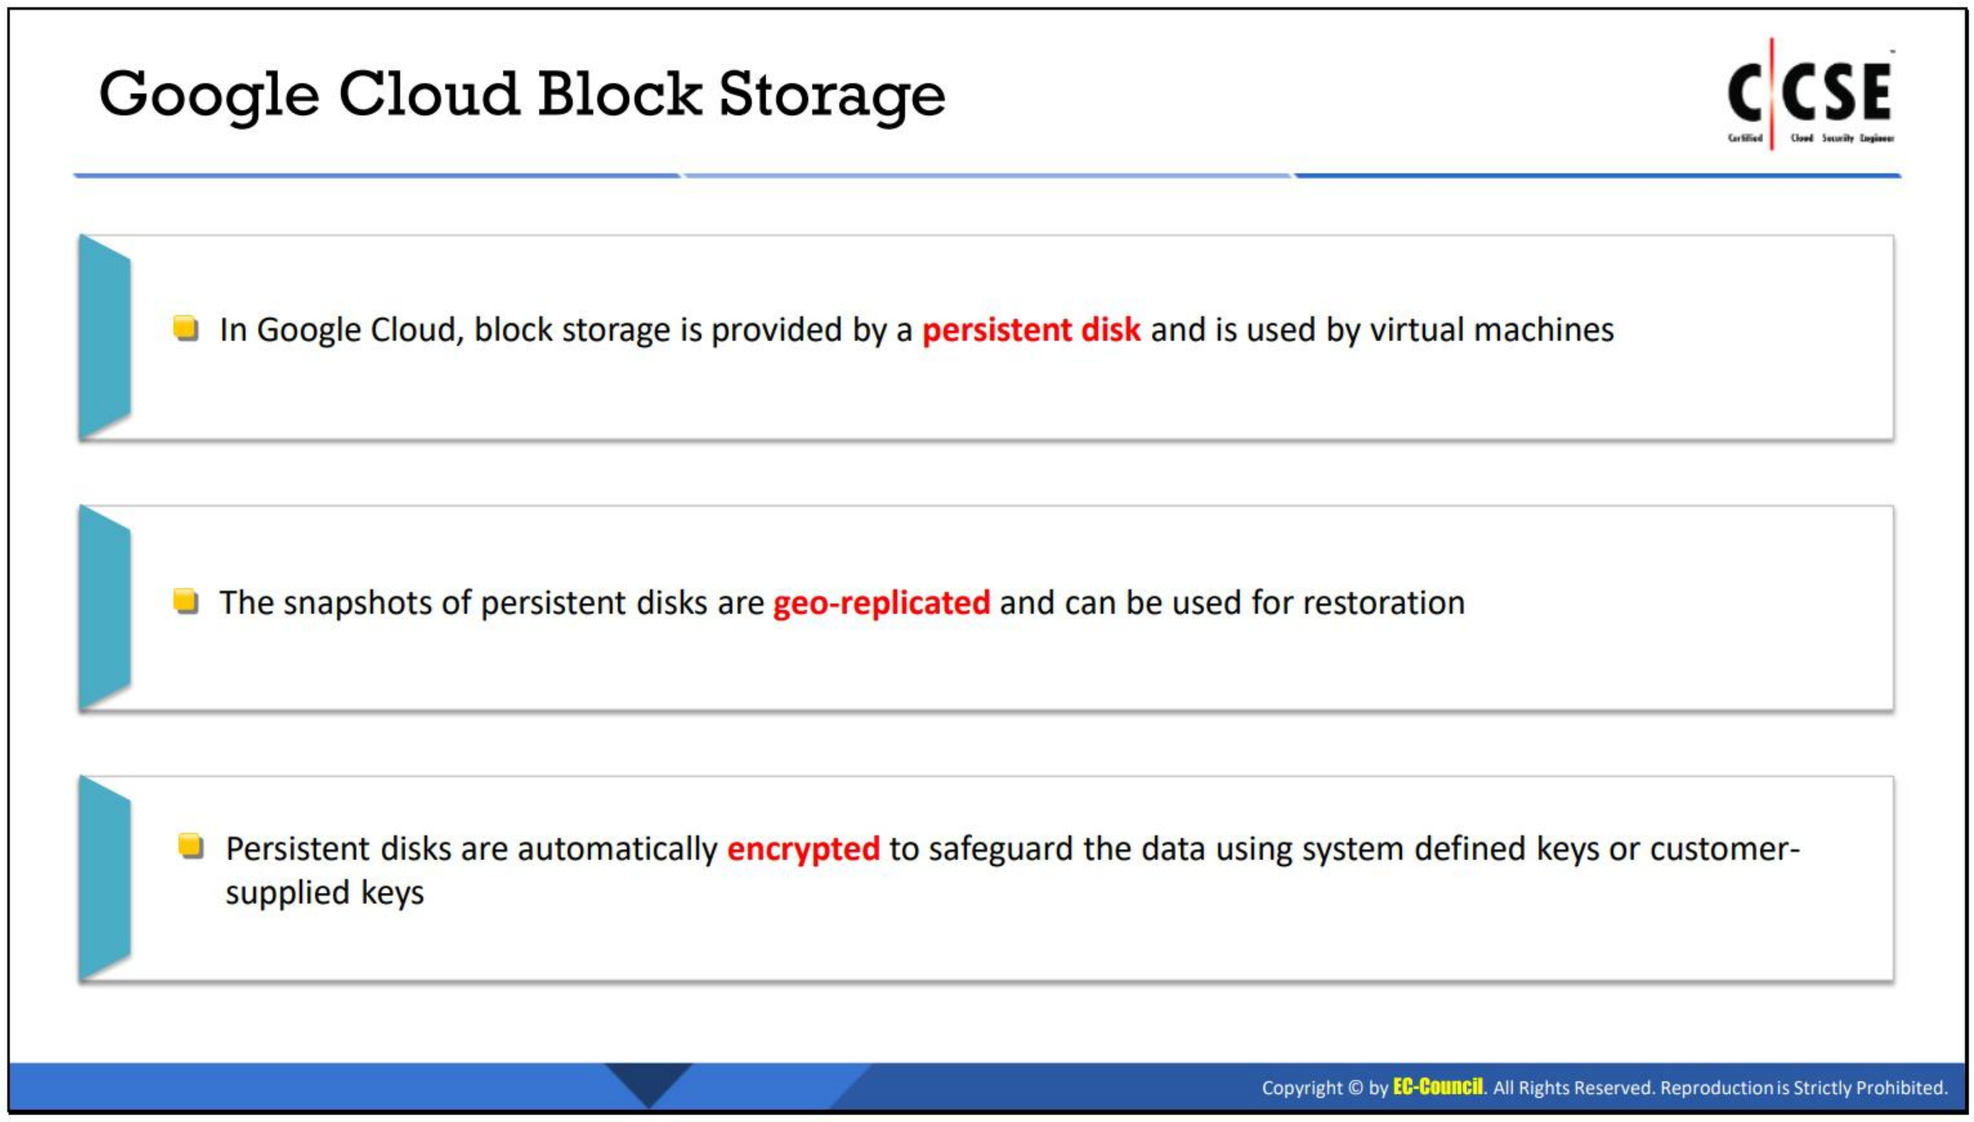
needs less availability. In the end, only a copy of the logs may need to be retained. By using lifecycle management rules, the storage classes can be modified as per requirement and cost efficiency.

With Google Cloud Object Storage, cloud users need to pay only for what they use. The users need to pay only for the data stored in the cloud. The pricing is based on read and write operations performed on the objects and network charges for data movement in buckets.

The cloud security administrator can control access to the buckets and objects in the Google Cloud Object Storage using ACL. ACLs records read, write, or read/write permissions granted to the users. With this feature, a cloud administrator can set his own permission to different users and groups. Using IAM, a cloud administrator can provide additional security to the objects by restricting access to the storage buckets. Using IAM, the cloud administrator can create users, groups, and set permissions to them.

hide01.ir

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



Google Cloud Block Storage

When it comes to Google Cloud Storage, it is important to understand the basic building blocks while developing cloud-based applications. The critical part is to decide how and where the data needs to be stored and used safely without compromising the performance.

Google Cloud provides a popular choice for high-performance storage called persistent disks. A Persistent Disk provides block storage and is used by VMs in Google Cloud. The storage and computing are decoupled in Google Cloud VMs. This means that the block storage is managed separately, and it can be attached and detached from virtual machines as per the requirement. The data in the Persistent Disk can be kept even after deleting the instance. Thus, it provides reusability and availability. The persistent disks can be re-sized while it is still in use and there will not be any downtime.

Persistent Disks can be either HDDs or SSDs. HDD offers cost-effective storage and is used when bulk throughput is required. SSD offers high performance for random workloads and bulk throughout. Both the storages can be extended up to 64TB.

The persistent disks can either be regional, zonal, or local. Regional persistent disks provide high availability by replicating the data in different zones in the same region. Zonal persistent disks are less expensive. It also provides high availability, but it is stored within a single zone. Local persistent disks are the hardware where a virtual machine runs, and it offers the best I/O performance. Local persistent disks are less often recommended as they offer low availability and redundancy.

Persistent disks can be attached to one or more virtual machines without affecting the performance. Multiple virtual machines can read the data from a single Persistent Disk. Disk

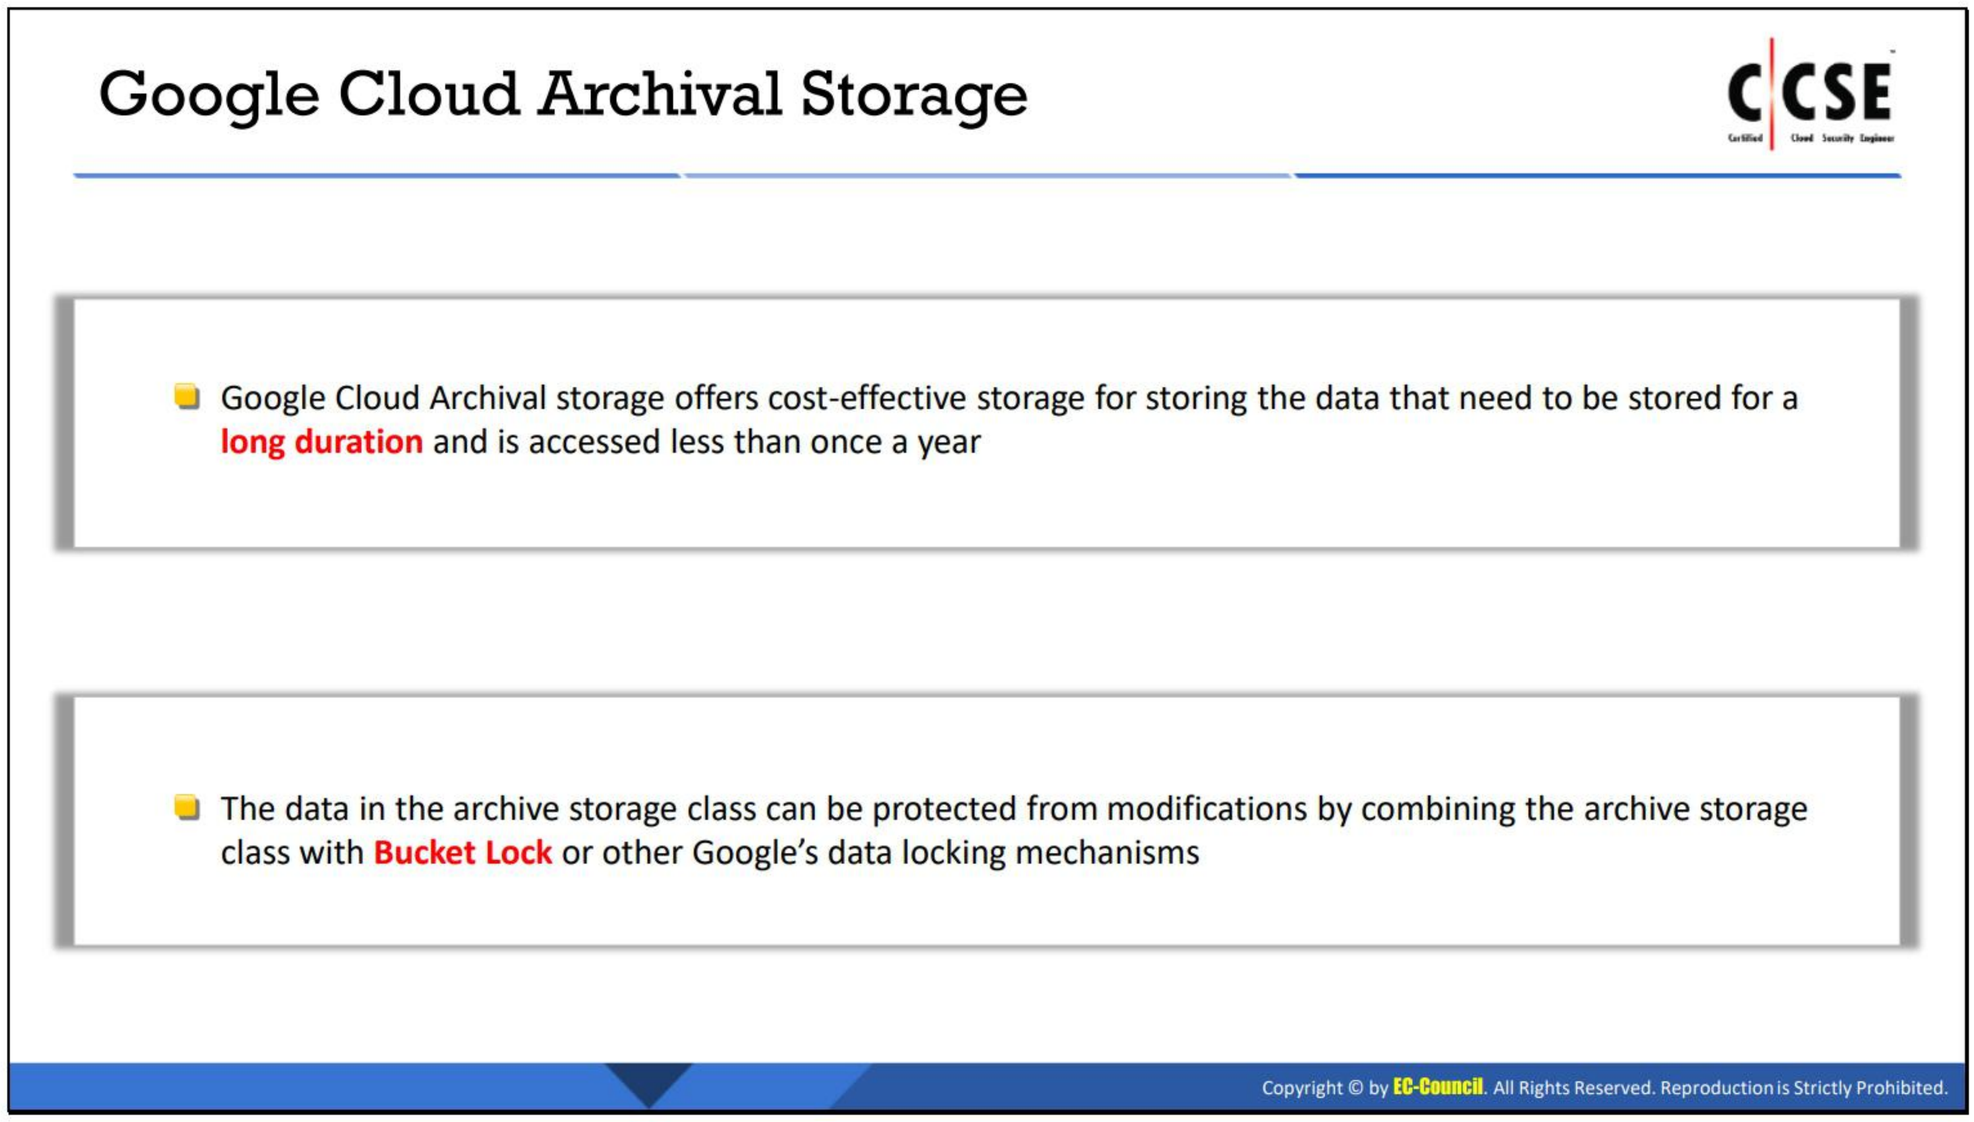
Certified Cloud Security Engineer Copyright © by **EC-Council**
All Rights Reserved. Reproduction is Strictly Prohibited.

management tasks will be handled by the Compute Engine and the cloud consumer does not need to worry about partitioning or sub-volume management.

Snapshots can be created to back up data periodically and it protects the cloud consumers from data loss. The snapshots of persistent disks are geo-replicated and can be used for restoration. The persistent disks are automatically encrypted to safeguard the data using system defined keys or customer-supplied keys.

hide01.ir

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



Google Cloud Archival Storage

Google Cloud Archival storage offers cost-effective storage for storing the data that need to be stored for a long duration and is accessed less than once a year. The minimum duration for archival storage is one year. Access and management are performed with the same APIs that are used for other storage classes. The archive storage class is integrated into object lifecycle management so that the cloud users can move the objects to archive storage and reduce the cost. It offers durability and availability, as the data stored in archive storage is stored in multi-regional locations. It is used to store backups and original copies of videos and images. The archive storage is also helpful in case of disaster recovery with less recovery time.

The data in the archive storage class can be protected from modifications by combining the archive storage class with bucket lock or other Google's data locking mechanisms. Bucket lock combined with retention period can add additional security to the data. It offers high durability. The Google archive storage eliminates the need for a retrieval process and the data can be accessed in milliseconds. The data in archive storage is encrypted by default, both during transit and at rest.

For example, consider an organization that has a bulk amount of data that needs to be retained following all the regulations including data protection. For this, the archive storage class can be used along with the Object Lock to ensure that the objects are stored without any modifications till the specified time. This is especially used by industries where legal retention requirements need to be met. If availability and durability are the major requirements, the cloud customers can choose archive storage classes with multi-region or dual-region locations.

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

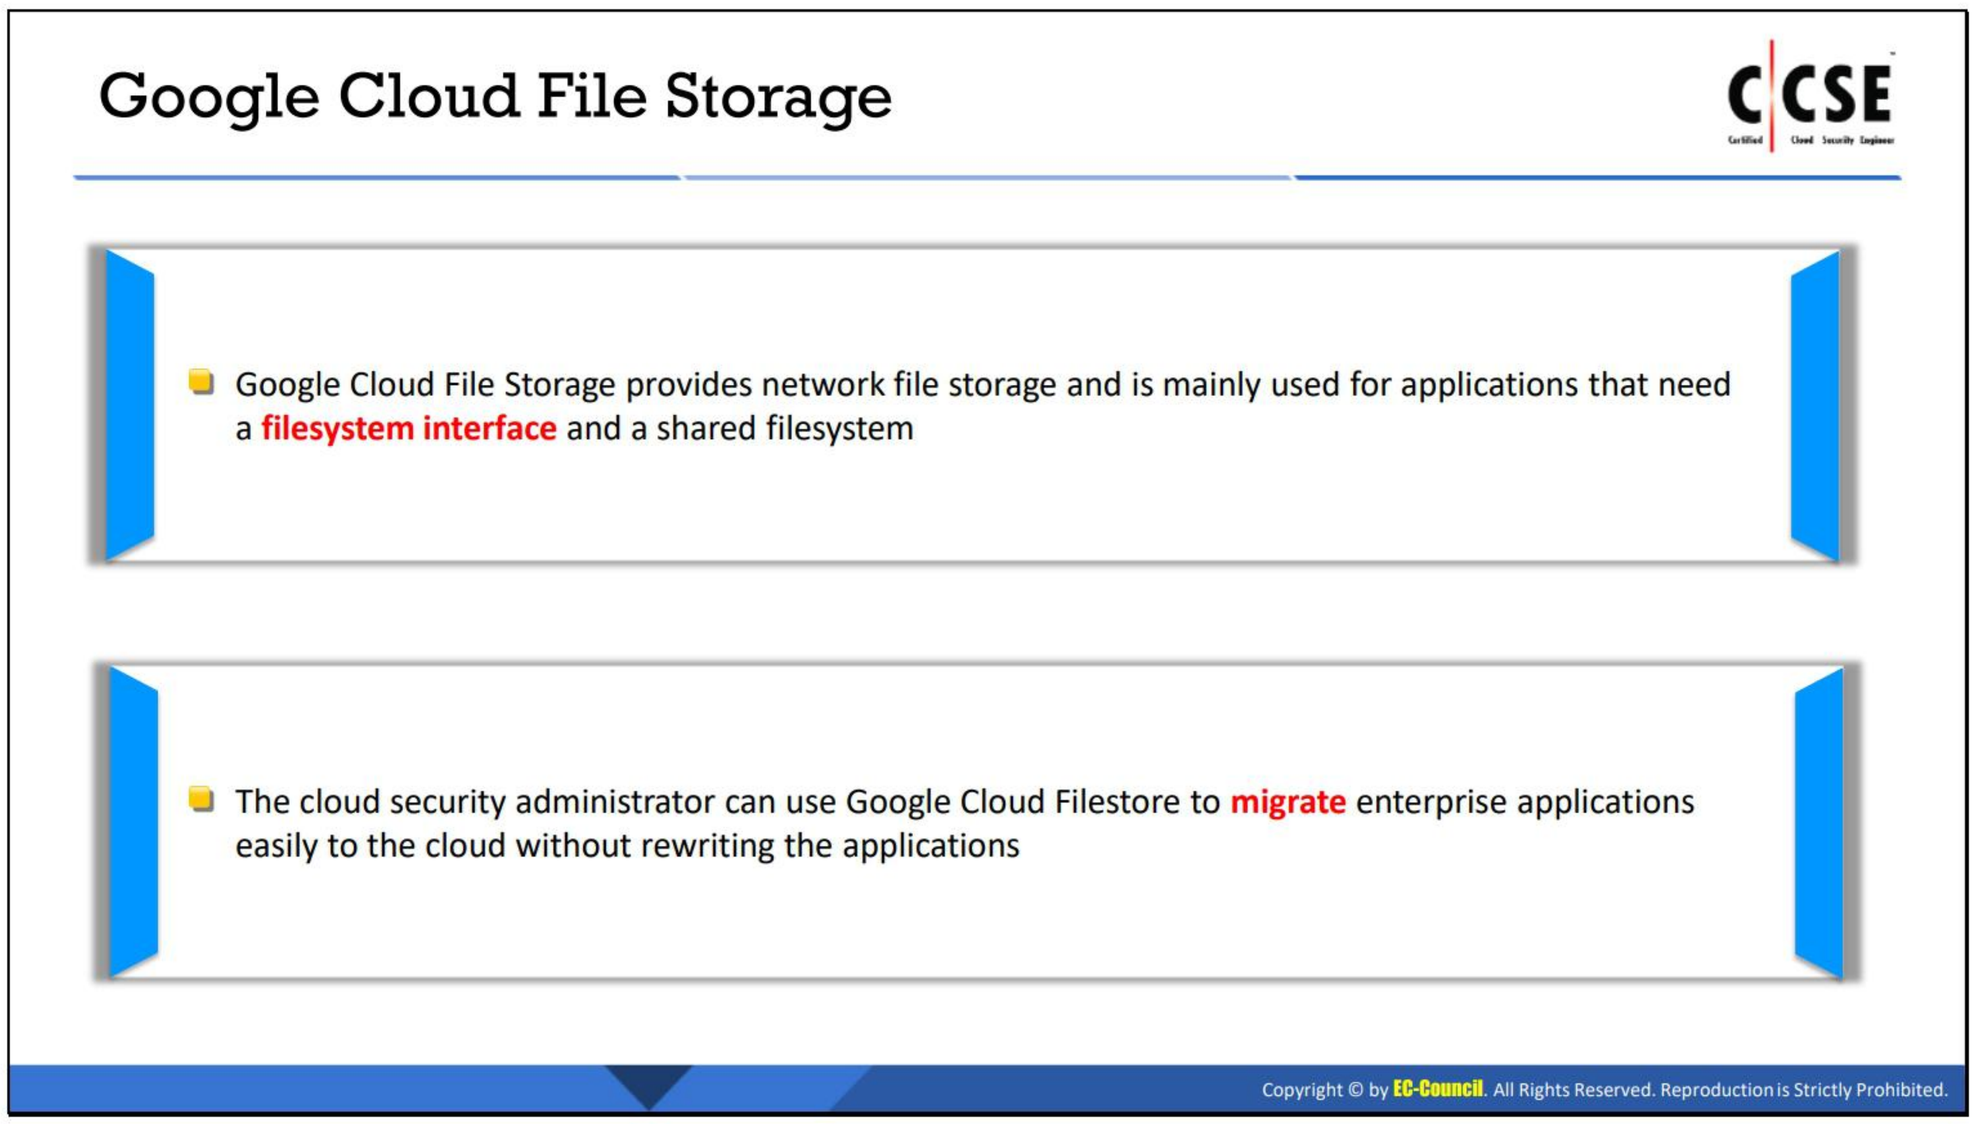
Industries that use archive storage class include:

- Educational sector for storing research data
- Government for storing data for a long time
- Financial industry to store transactions, and audit logs
- Healthcare industry to store medical records
- Insurance for storing claims, and billing
- Media and entertainment industry for storing raw production footages
- Security industry for storing video surveillance data
- Telecommunications for retaining call records, billing, and customer-service records

Though Google archive storage is expensive, it is beneficial as the data can be retrieved in a fraction of seconds.

hide01.ir

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



Google Cloud File Storage

Google Cloud File Storage is a management service that provides NFS. It is mainly used for applications that need a filesystem interface and a shared filesystem. The file storage offers fully managed NAS with Google Compute Engine. As the Filestore is fully managed, the cloud users can easily mount the volumes on Compute Engine Virtual Machines. It provides fast performance for file-based workloads.

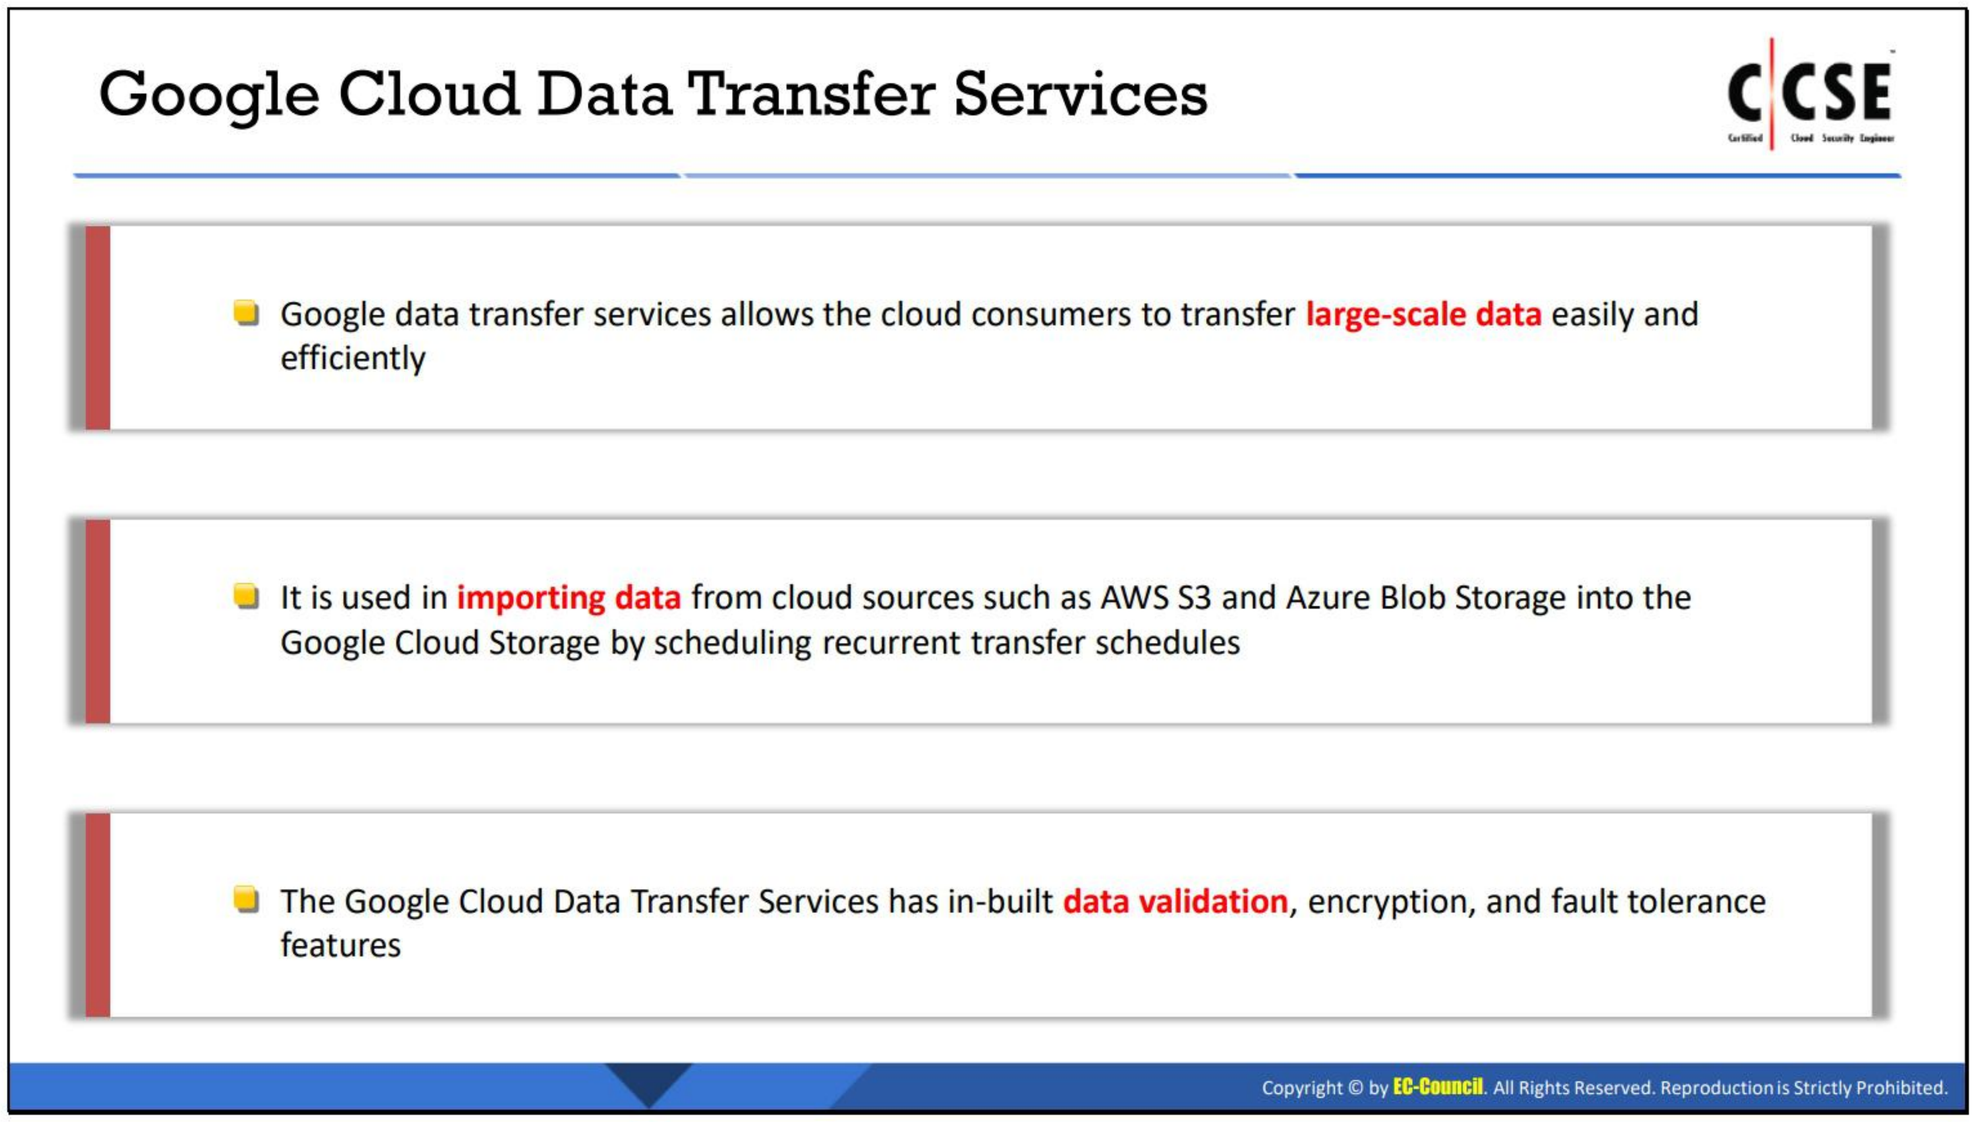
The Google Cloud File Storage provides disk storage through the network and can be increased to 100's of TBs for handling high-performance workloads. The Google Cloud provides the cloud consumers with 2 performance options to handle the workload. The Standard with 5000 IOPS and Premium with a maximum of 30,000 IOPS.

Google File Storage provides high performance with low latency with speed up to 480K IOPS and 16 GB per second. It is highly beneficial for latency-sensitive workloads such as electronic design automation (EDA), data analytics, and media rendering.

Using Google Cloud Filestore, enterprise applications can be easily migrated to the cloud without rewriting the applications. It supports NFSV3, which helps in enabling file lock without installing any plugins or client-side software. With file lock, all the data stored in Filestore is secured. File Storage comes with high scalability. The cloud users can easily increase or decrease the instances as per the business requirements using Google Console, API-based control, or gcloud CLI.

The Google Cloud users can choose the IOPS and storage capacity. This enables them to tune the filesystem for each type of workload. The speed would be consistent for that workload.

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



Google Cloud Data Transfer Services

For cloud consumers, data transfer is the first step for using Google Cloud Storage. To transfer large datasets, the management must develop a team to make the transfer smoothly and securely. This can be easily performed using the Google Cloud Data Transfer service.

hidoo1.ir

Whether the source is on-premises or online, using Google data transfer services the cloud consumers can easily handle large-scale data transfer easily and efficiently.

Using Google Storage Transfer service, cloud customers can import data from cloud sources such as AWS S3 and Azure Blob Storage into the Google Cloud Storage by scheduling recurrent transfer schedules. The data moves on Google's high bandwidth network with reserved resources. The cloud users can set up repeated and incremental synchronization of their data to reduce the transfer time. With this, the data from other cloud providers can be moved to Google Cloud Storage simply, quickly, and safely.

The Storage Transfer service can be used to transfer data from one bucket to another and to move data between storage classes to handle the workload. For data transfer, private credentials are used at source and destination to ensure secure data transfer. The customers are charged when they transfer data from other cloud providers. Charges are also applicable for transferring data from one cloud storage bucket to another. If a list operation needs to be performed for data transfer, the customers will be charged from Google Cloud Storage as well as the source cloud provider.

To transfer on-premises data, Google Transfer Service can be used. The Transfer Service is built with simplicity and security, so that it scales with available bandwidth and makes the transfer in just a few minutes. Using this, petabytes of data can be transferred from on-premises sources to Google Cloud Storage through online networks. It reduces the cost for transferring their data.

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

The customers do not need to engineer their custom software or an off-shelf solution to transfer the data to Google Cloud. The data can be transferred simply without writing even a single line of code.

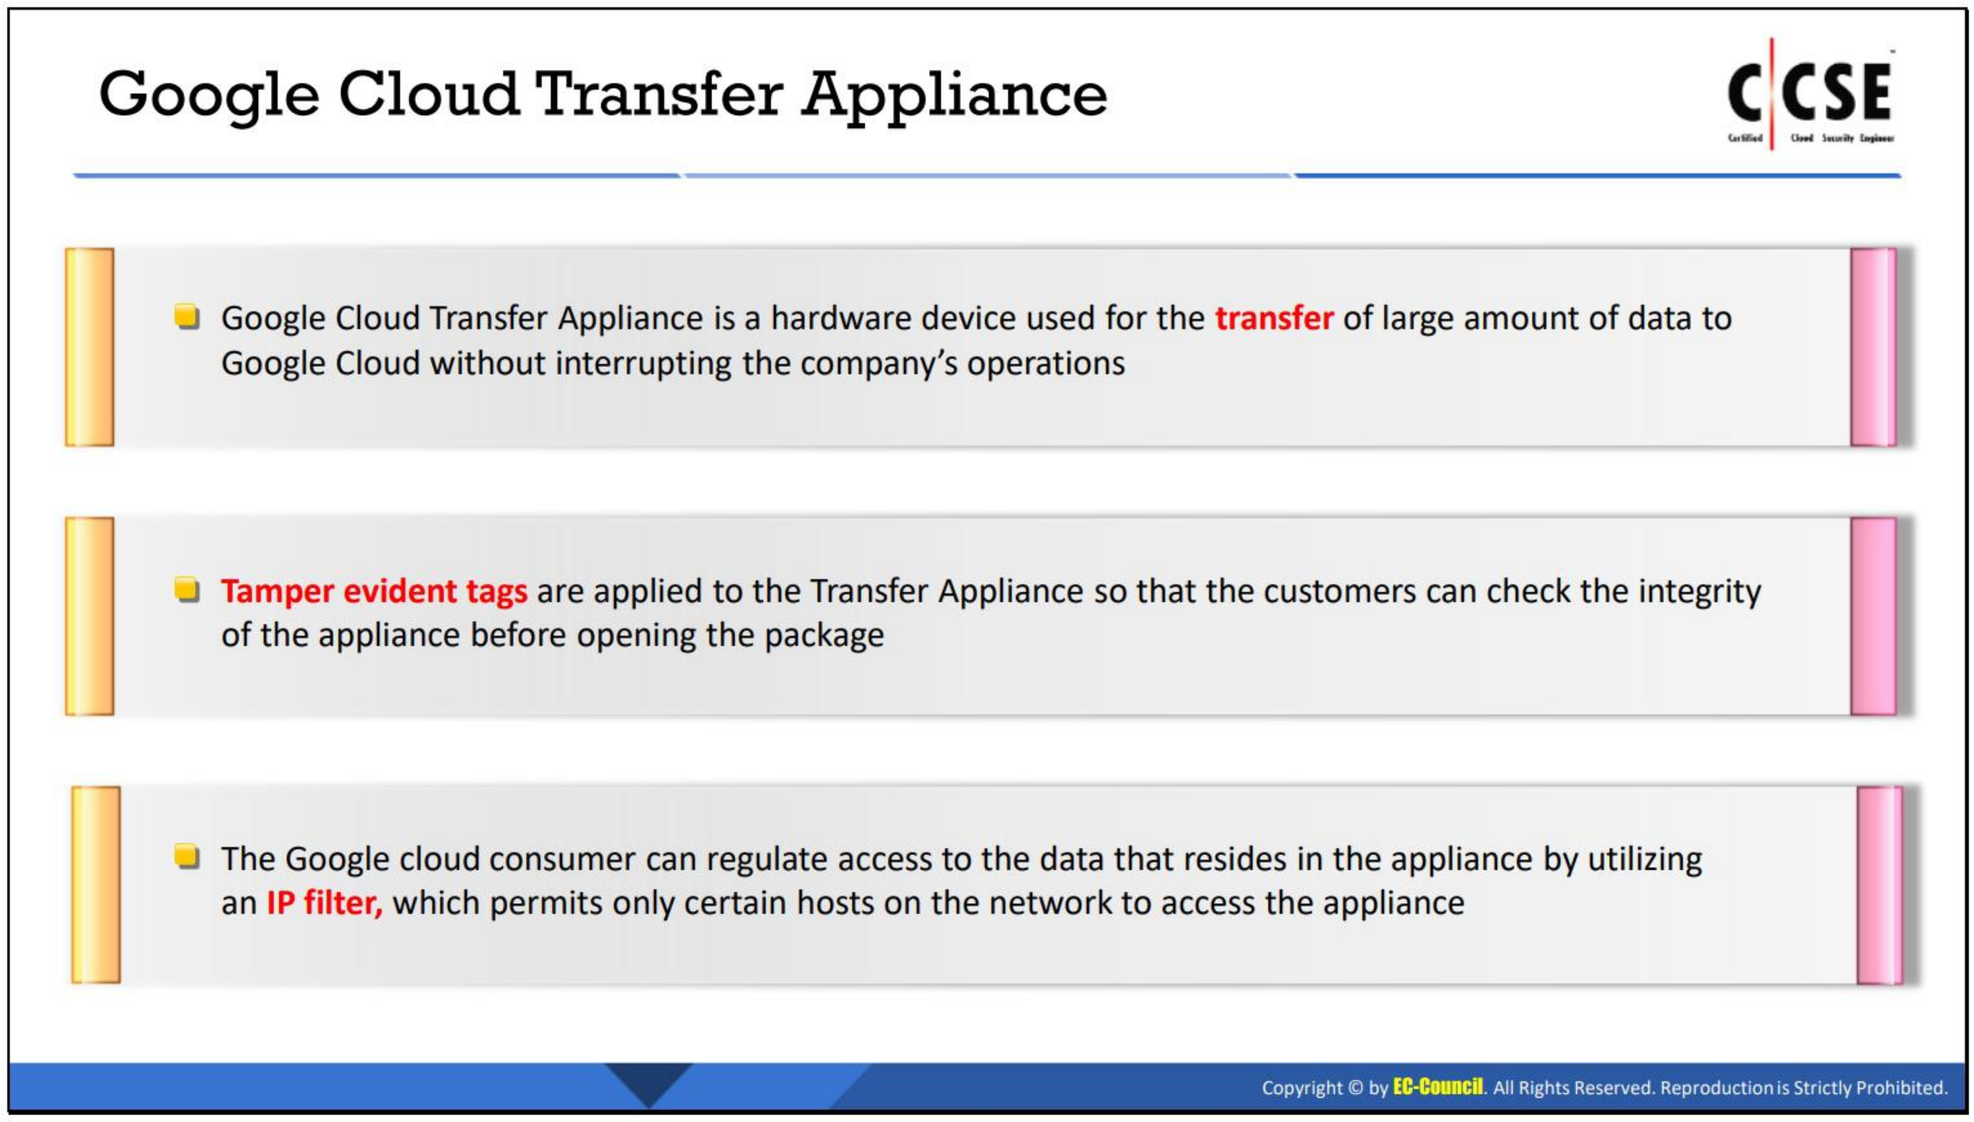
The Google Transfer Services comes with in-built data validation, encryption, and fault-tolerance features. It is reliable and secure so that the transfers will not be impacted in case of an agent failure. The data transfer is also secure as the data is encrypted in transit and byte to byte checksums are performed. The on-premises agent software can be installed in few minutes and is integrated into Stackdriver and hence reduces the operational attention. If the customers have limited bandwidth, they can still use Transfer Service by setting up a bandwidth limit for the Google Cloud project. This bandwidth will be shared with all the transfer jobs.

When a failed object transfer is retried, the customers will not be charged for the initial bytes moved on a failed transfer request. If the customers opt for incremental transfer, the Google Transfer Service will not retransfer the file that has not been changed from the previous transfer.

The customers have several options to protect their resources and data during transfer. They can control the access to their on-premises data by using a restricted user to run the agent container. The customers can also limit the files that are mounted to the agent container. The data in the flight is also encrypted with TLS for both public and private connections. The customers can also use predefined IAM roles such as Storage Transfer Administrator and Storage Transfer User to control, restrict, and manage the transfer operations.

hide01.ir

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



Google Cloud Transfer Appliance

Google Cloud Transfer Appliance is a hardware device used for the transfer of large amount of data to Google Cloud without interrupting the company's operations. It is a high-capacity device that can be used to transfer and ship the business data to the Google Cloud Storage upload facility. From the upload facility, it will be transferred to Google Cloud Storage.

With Transfer Appliance, the transfer process is simple. The customers need to request the appliance. They just need to mount the NFS share exposed by the Transfer Appliance and upload their data. Once the upload is completed, the customers need to seal and ship the appliance. Once it is received, the Google Cloud administrators will upload the data to Google Cloud and will wipe the data from the appliance. The customers can also request a wipe certificate to ensure that the data is wiped from the appliance. The customers can transfer a petabyte of data within 25 days and the data can be accessed on Google Storage within the next 25 days.

No additional software is required to transfer data to the appliance. The customers only need to mount the NFS share on the appliance to the workstation. The customers can order multiple Transfer Appliances to increase the transfer speed.

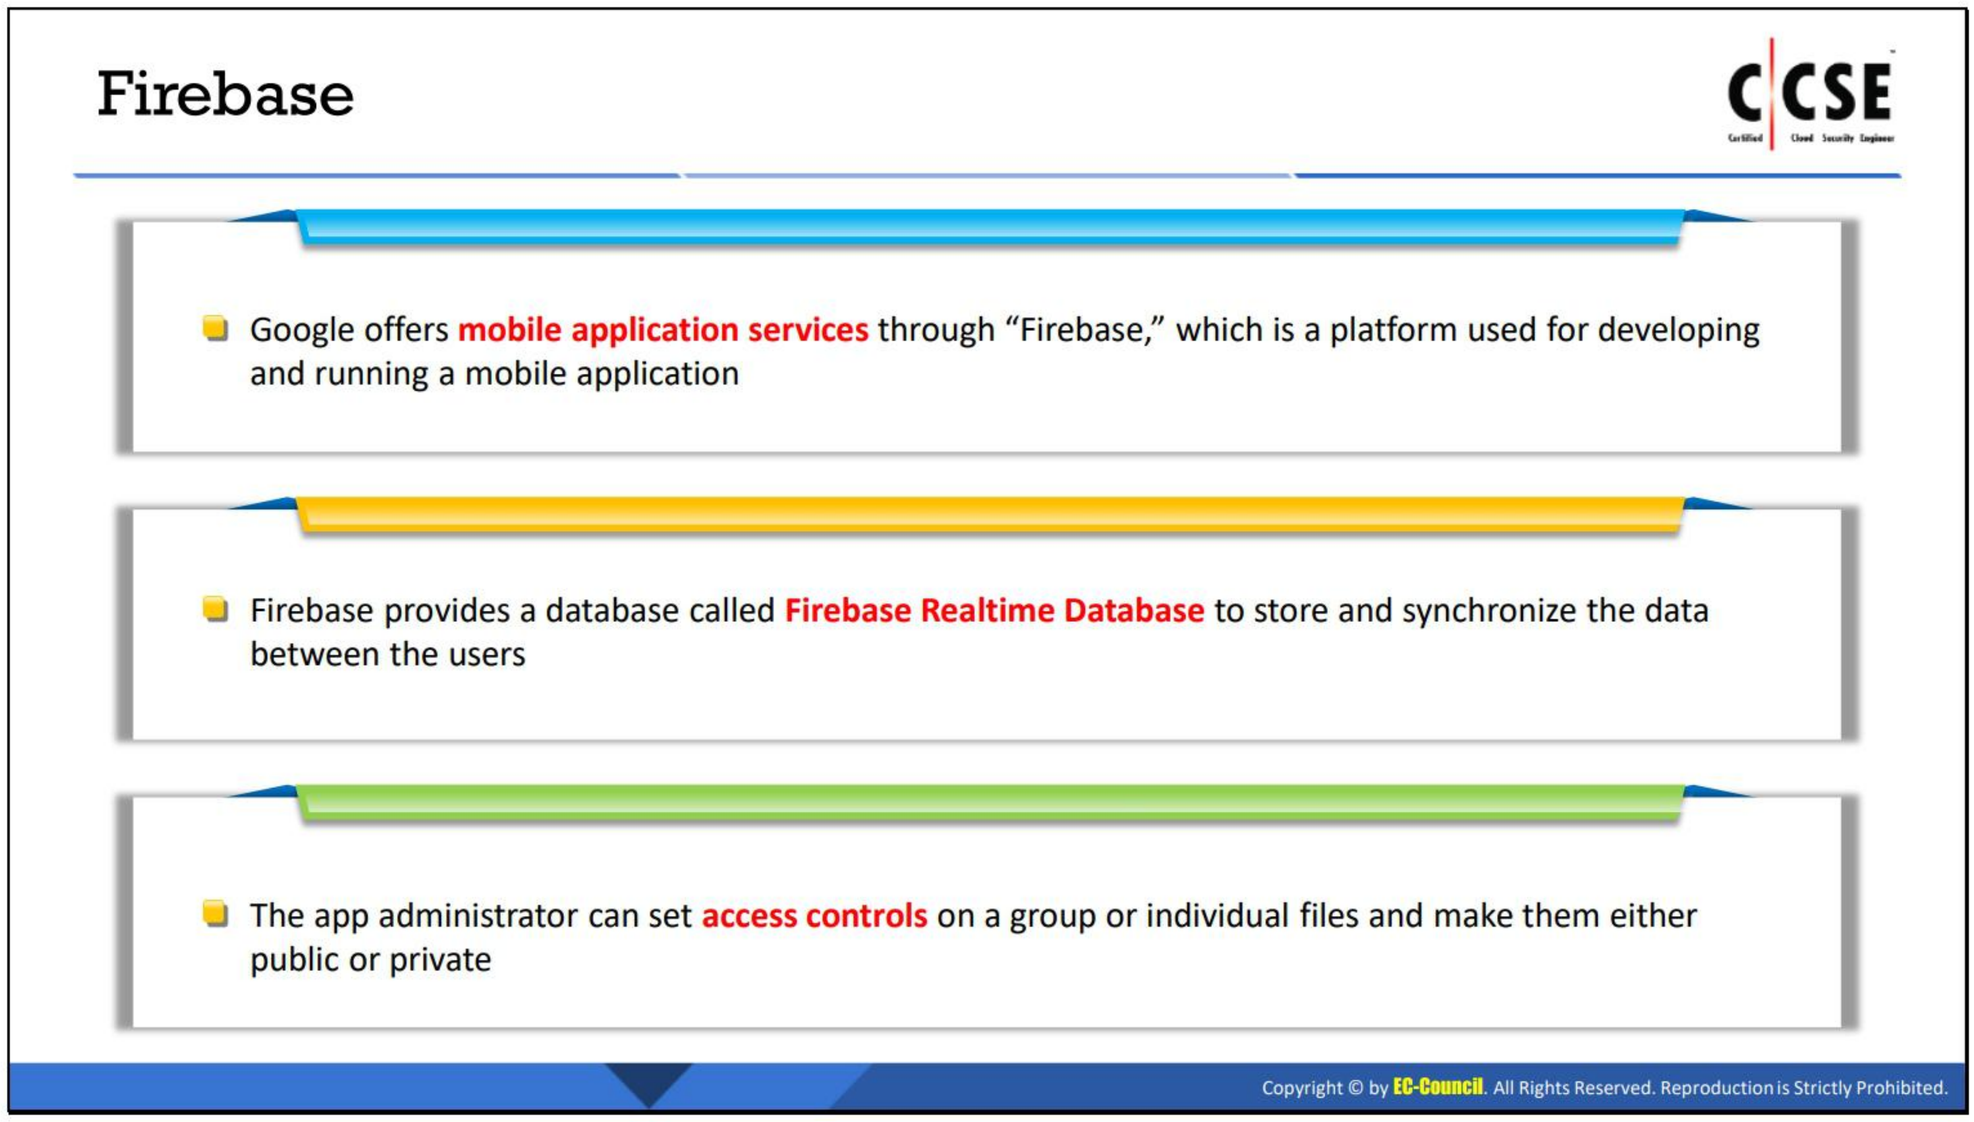
Data and network security are important for the Transfer Appliance. Tamper-evident tags are applied to the Transfer Appliance so that the customers can check the integrity of the appliance before opening the package. The shipping container is also ruggedized to ensure that the data is safe. Google provides a remote attestation process to validate the identity of the appliance before the customers can connect and copy data to it. Once the transfer is completed and is received by Google Cloud administrators, the device is again validated to ensure that it is not being tampered during transit. If the appliance validation shows that the device was tampered, the cloud administrators will invalidate the transfer session and will send a new appliance to the customer.

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

With Transfer Appliance, the customers can trust that their data is safe. The data is encrypted during upload and after uploading to the Google Cloud Storage. The Cloud customers can also restrict access to the data in the appliance by applying an IP filter that allows only specific hosts on their network to access the appliance.

hide01.ir

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



Firebase

Google offers mobile application services through "Firebase." Firebase is a platform used for developing and running a mobile application. Using Firebase, high-quality Android, web, and iOS applications can be developed. It provides features such as auto-synchronization of data, messaging, authentication, file storage services, and many more. It is backed up by the GCP.

With all the services in one place, the developers can focus on insights and the objective of the app rather than wasting time on checking for other third-party services for infrastructure.

The customers do not need to create an account to operate and manage services. They can use the Google Cloud service account to perform the operations. When a Firebase project is created using Google Cloud, a service account will be created in the project. The Firebase SDKs use a default storage bucket, and this allows users to use Google Cloud without using a credit card.

Cloud storage is used to store user content generated from the Firebase application in the form of objects. Using cloud storage, it is easy for the developers to store and provide content to users. Regardless of the network quality, security is provided for Firebase apps file uploads and downloads. The uploads and downloads restart from where they paused in case of any interruption and it reduces time and bandwidth. The data can also be easily shared between Firebase and Google Cloud project. The cloud storage for Firebase scales automatically and the customers do not need to migrate to other storage providers.

As the app progress, the developers can easily integrate it with other Google Cloud services. All the files uploaded through Firebase SDKs can be accessed using Google Cloud Storage APIs. All complex operations such as copying and moving files can be performed. All the requests use the Google Cloud access control option, but not with Firebase Authentication.

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

There are also other ways to access data in the Google Storage buckets based on the actions that need to be performed. For example, to modify the script or to perform administration tasks, a command-line tool can be used. To access data on the server, server-side libraries such as RESTful API can be used.

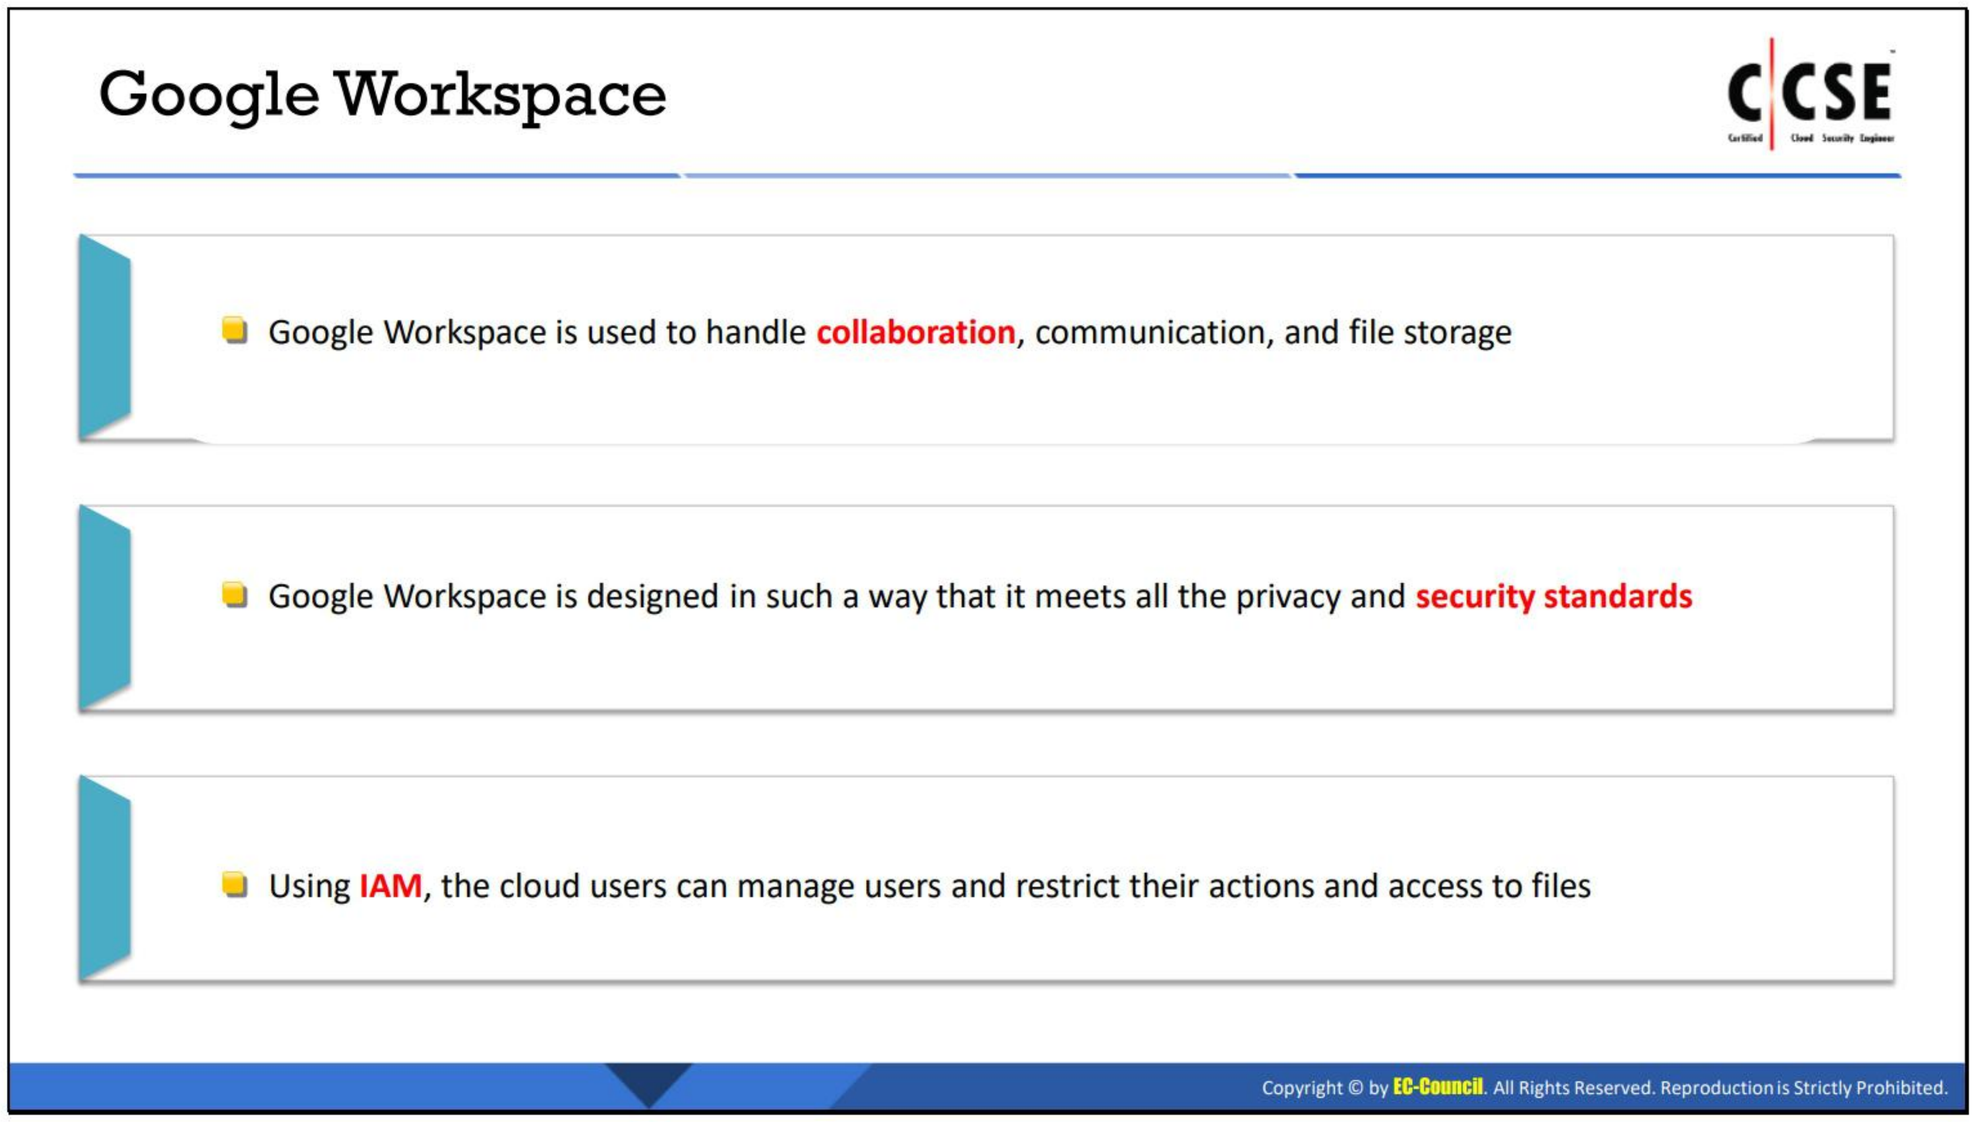
Firebase also provides a database called Firebase Realtime Database to store and synchronize the data between the users. The appearance of the apps can also be changed using "Firebase Remote config." Using this, all the activities such as releasing new content and splitting test for validation improvement can be performed without approval from the app store.

It also provides "Cloud Functions," a serverless framework that runs the JavaScript code to handle the events triggered by Firebase features such as changes to data, creation of new users, and conversion events. It also protects the application logic from tampering by insulating it fully from the client. This ensures that the functions are private and secure.

To handle security and authentication, Firebase provides "Firebase Authentication." It supports phone authentication, Google and social media login, and email and password accounts. The Firebase Security Rules can also be used to handle authorization and validation requests. The authorization is based on the user identities and file properties such as content type and size. It also reduces the work of developers to write the code from scratch for authentication. The app administrator can set access controls on a group or individual files and make them either public or private.

hide01.ir

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



Google Workspace

Google Workspace is used to handle collaboration, communication, and file storage. Google workspace integrates with a set of tools to work together and makes the digital workplace creation easy. This helps people work online and collaborate using new ways and focus efficiently on productive work.

hidoo1.ir

Google Workspace provides an email id for business and it collaborates with tools such as Meets, Chat, Drive, Docs, Sheets, Slides, and many more. Based on the size of the organization, users can select the advanced and additional features as per their requirement. There are different plans, and the users can choose the best suitable plan.

Using Google docs, sheets, and slides, it is easy to interact with the entire team who work in multiple locations and time zones. If the team needs to interact for a business discussion, Google Meet can be used, which provides a simple and reliable video conference. 150 members can participate in the meeting and the meeting can be recorded and uploaded to Google Drive. The recordings can be used by the team members who have not attended the meeting.

Team members can store the data in Google Drive, and it can be shared and accessed from any device. The required files can be uploaded to the shared drive and can be accessed by all the team members.

The Google Drive files can be synced with PC, MAC, or laptop to reduce the disk space. Whenever a team member makes any changes to the synched files, they will be automatically updated on your PC.

Google Workspace can be used from any devices such as mobiles, tablets, or computers, and from anywhere. It also provides offline access, where users can still view, create, and edit files

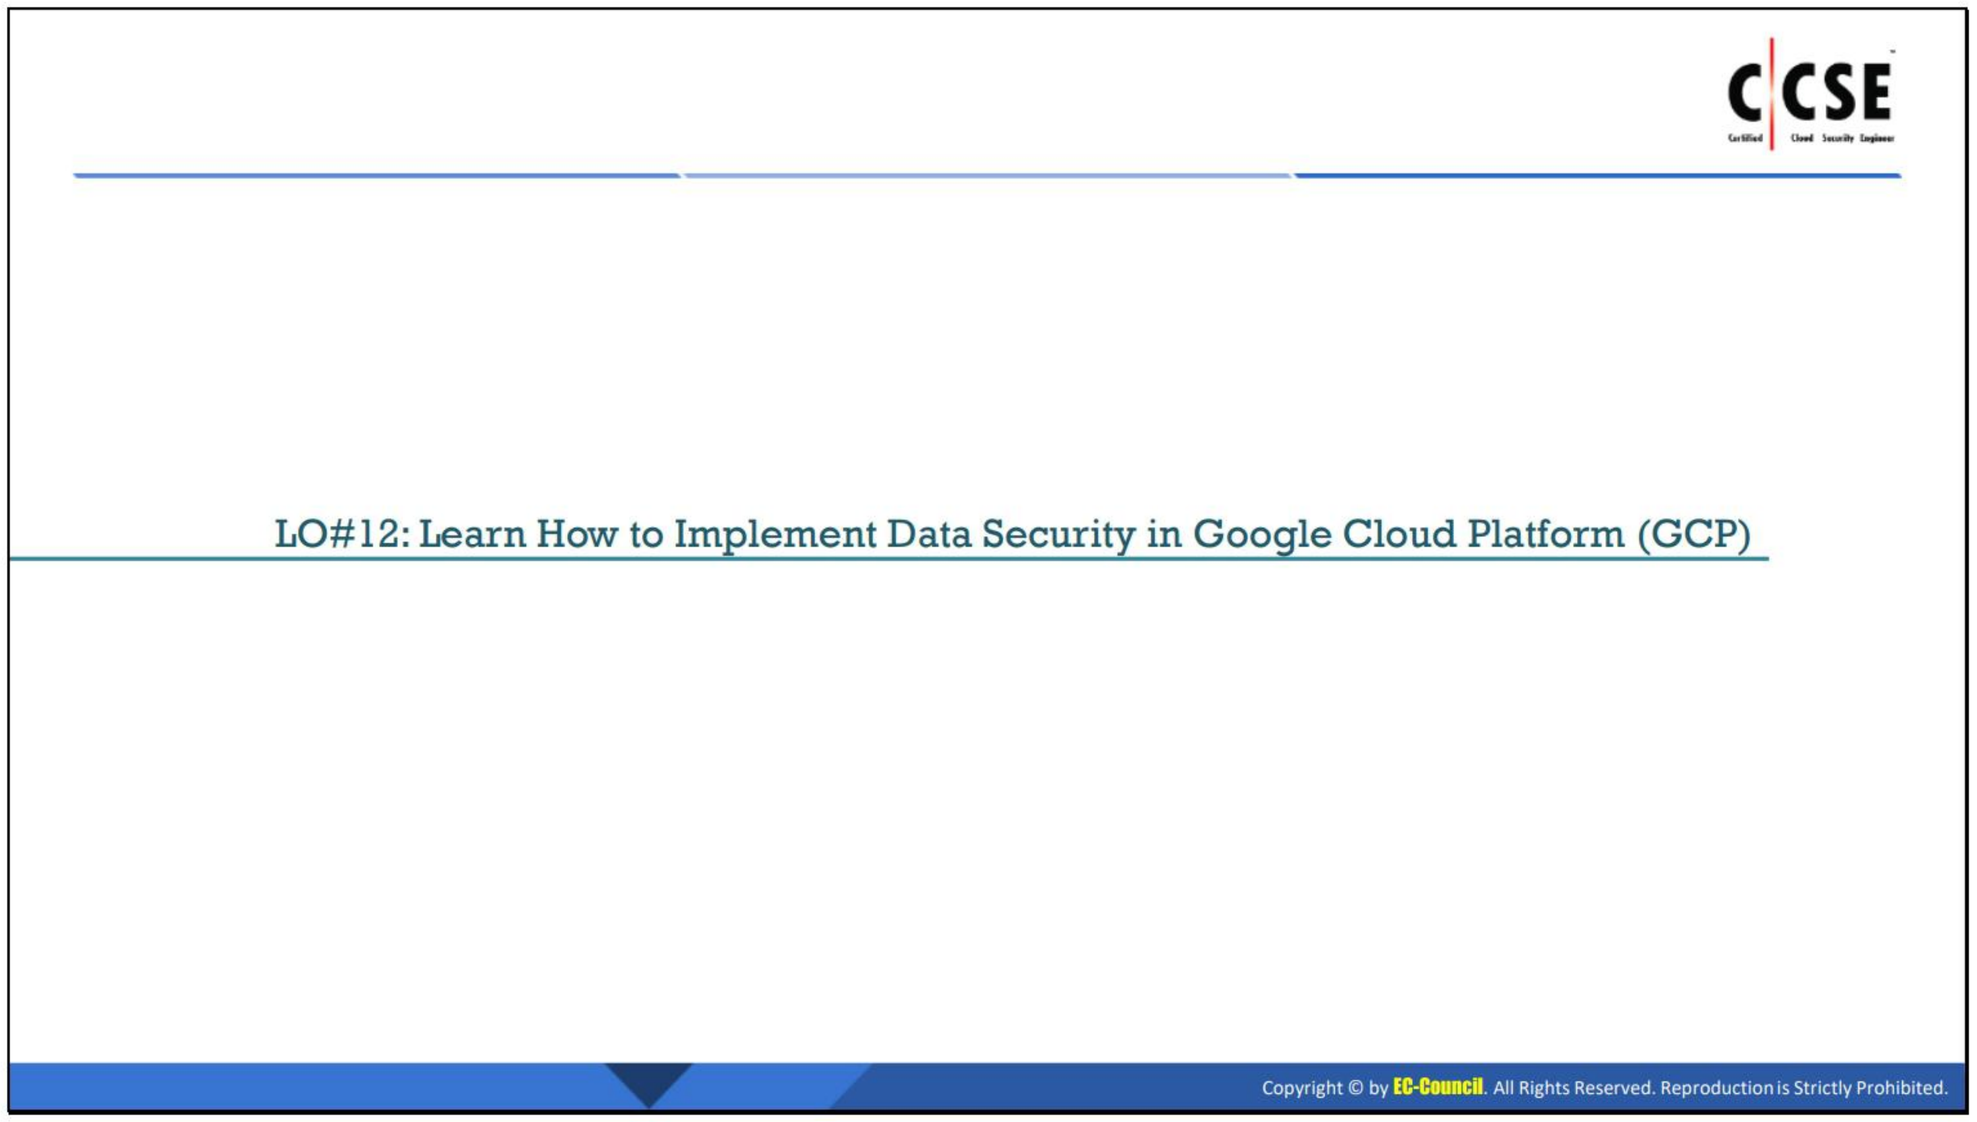
Certified Cloud Security Engineer Copyright © by **EC-Council**
All Rights Reserved. Reproduction is Strictly Prohibited.

when they are not connected to internet. Once the users get connected to the internet, all the changes will be automatically synced.

Google Workspace is designed in such a way that it meets all the privacy and security standards. It also enables the administrators to control the access and handle system configuration and application settings. Authentication, resource protection, operational control can be managed from the dashboard. Using IAM, the cloud users can manage users and restrict their actions and access to files. MFA and security keys can be used to add additional security to resources and files. Google also focuses on transparency. The customers own the data and Google Workspace will not use or sell the data to any third parties, the customer data will always be safe and secured.

hide01.ir

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



LO#12: Learn How to Implement Data Security in Google Cloud Platform (GCP)

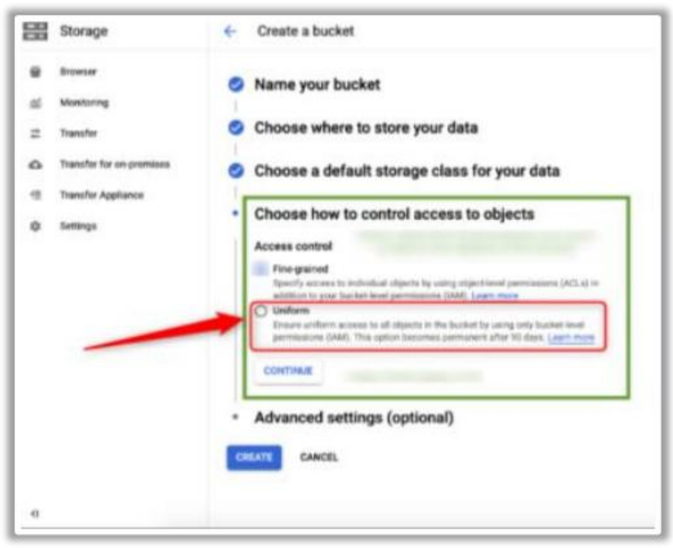
This section focuses on the implementation of data security in Google Cloud using GCP features such as Cloud Audit Logs, VPC Service Controls, Uniform Bucket-level Access, etc.

hide01.ir

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

Turn on Uniform Bucket-Level Access

- To provide additional security to the stored data in the Google Storage, turn on **uniform bucket-level access**
- This feature can be enabled during the bucket creation or on the existing bucket, which would add additional **IAM roles**



Configuring Uniform Bucket-Level Access

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

Turn on Uniform Bucket-level Access

Google Cloud Storage is used by many businesses to store large volumes of data, so that it can be easily accessed by all the users. As the demand increase, Google Cloud started investing more in protecting the customer data in all the Google products. The data stored in the Google Cloud Storage is secured by using robust security controls.

In Google Cloud Storage, users can set permissions to objects and buckets using IAM or ACL. Both systems work in the same way and either of the methods can be used to grant permissions to users.

Uniform bucket-level access can also be used to add additional security to the data stored in Google Storage. By enabling uniform bucket-level access, ACL get disabled and all the permissions to the buckets and objects can be granted only through IAM. Once this option is enabled, the cloud security professional can reverse it in 90 days.

This feature can be enabled either during the bucket creation or explicitly on the existing bucket. When this feature is enabled during bucket creation, additional IAM roles get added to the bucket. Before enabling this option, the cloud administrator should ensure that all the permissions granted through ACLs are migrated to IAM.

The Google Cloud Storage stores the existing ACLs for 90 days, so that the cloud security professionals can disable uniform bucket-level access and use the saved ACLs. The uniform bucket-level access cannot be disabled if it stays on the bucket for 90 days.

Before disabling uniform bucket-level access, all the IAM conditions on that bucket must be deleted. Once this option is disabled, all the objects in the bucket regain the saved ACLs. The objects that are added to the Storage bucket after enabling uniform bucket-level access acquires

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

the default object ACLs used by that bucket. Uniform bucket-level access is highly suitable for enterprises that provide financial services.

To Enable Uniform Bucket-level Access

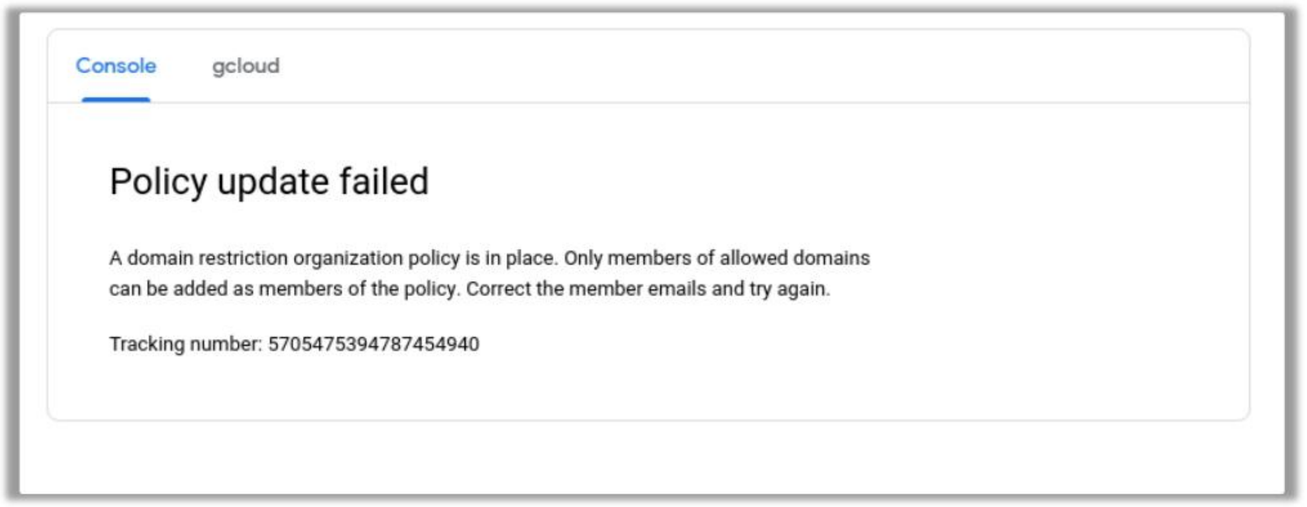
- From Google Cloud Console, open the cloud storage browser.
- Select the desired bucket on which you want to enable uniform bucket-level access.
- From the top of the page, choose the **Permissions** tab.
- Click the **Edit** button near the text box that shows **This bucket uses fine-grained access control**.
- Choose **Uniform** from the pop-up menu, then click **Save**.

hide01.ir

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

Enforce Domain-Restricted Sharing Policy

- Using domain-restricted sharing policy, an organization can control the **access** of domains to the Google Cloud resources
- This feature protects the data of the organization by preventing accidental **public data sharing** or sharing with anyone outside the organization



Policy Update Failed Due to Domain-Restricted Sharing Policy

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

Enforce Domain-Restricted Sharing Policy

When uniform bucket-level access is enabled, an additional feature called "domain-restricted sharing" can also be used. This feature protects the data of the organization by preventing accidental public data sharing or sharing with anyone outside the organization. This enables the organization and development team to focus on the work and the security teams to implement security at a granular level and ensure the right access to the resources.

When domain-restricted sharing is used along with the uniform bucket-level access, the cloud security professional can configure strong access control policies. This constraint can be enforced into the organization's policy and can restrict the resource sharing to one or more workspace domains.

When this constraint is enabled on a Google Workspace, it affects all the Google Workspace users. It means the constraint affects all the users under that domain. To set this constraint, the cloud security professional must have access to edit the organization policies.

Domain-restricted sharing is a list constraint, where Google Workspace user IDs can be included and removed, as per the list of the domain restriction policy. All the primary and secondary domains linked to that user account will get affected by the organization policy.

This policy also shows an error message in case of performing any actions beyond the policy. For example, when the cloud administrator tries to add a user, who is not in the list of domain-restricted constraint, it will display an error message and the action becomes unsuccessful. When a user tries to share a document or OneDrive files with a person who is not in the list of domain restriction sharing, sharing will not be allowed and it shows an error message.

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

Using this domain-restricted sharing policy, companies can control the access of certain domains to the Google Cloud resources. For instance, using this constraint, the organization can grant access to partners or clients and restrict their access to certain projects or folders.

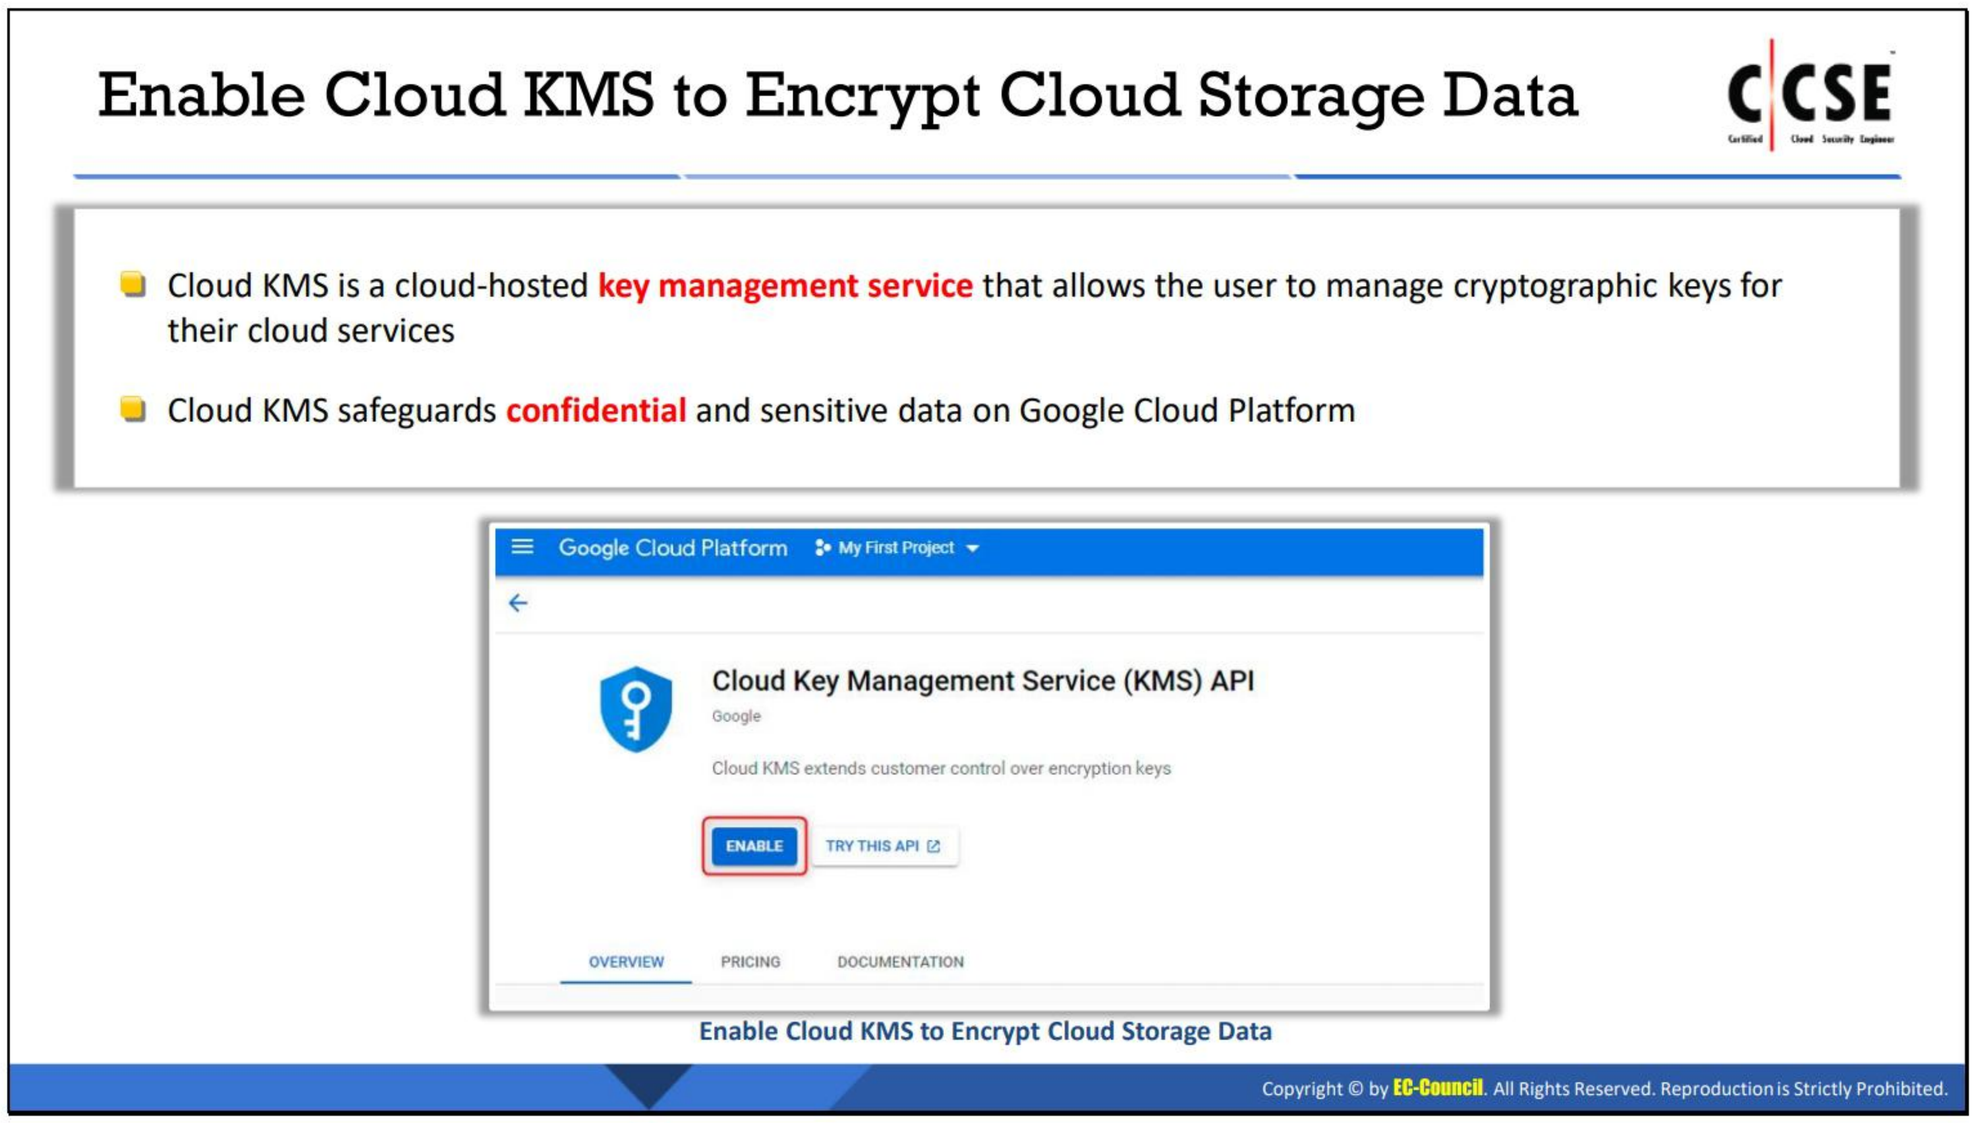
To Enforce Domain-Restricted Sharing Policy

Source: https://console.cloud.google.com

- From **Google Cloud Console**, browse to the **Organization policies** page.
- Click **Select**.
- Choose the name of the organization on which you want to set organization policy and domain restriction constraint.
- Click on **Domain Restricted Sharing.**
- Click on the **Edit** button.
- **Below Applies to**, choose **Customize**.
- Below **Policy values**, choose **Custom**.
- In the **Policy value** text box, type Google Workspace customer ID, then press the **Enter** button. You can enter many IDs in this way.
- Click on **Save**. Then, you will observe a notification that confirm that the policy is updated.

hide01.ir

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



Enable Cloud KMS to Encrypt Cloud Storage Data

Cloud KMS is a cloud-hosted key management service that allows the user to manage cryptographic keys for their cloud services. The user can generate, destroy, or rotate their own cryptographic keys. Cloud KMS safeguard confidential and sensitive data on GCP.

Cloud KMS enables cloud users to take custody of the data and manage the encryption keys in the cloud. By using Cloud KMS, customers can focus on the development of their products and Google takes the responsibility of managing the keys.

Cloud KMS provides additional benefits to users such as access control and monitor, regionality where users can store the keys in their selected region, and durability. It is completely integrated with IAM and it strongly protects the access of keys from an unauthorized user.

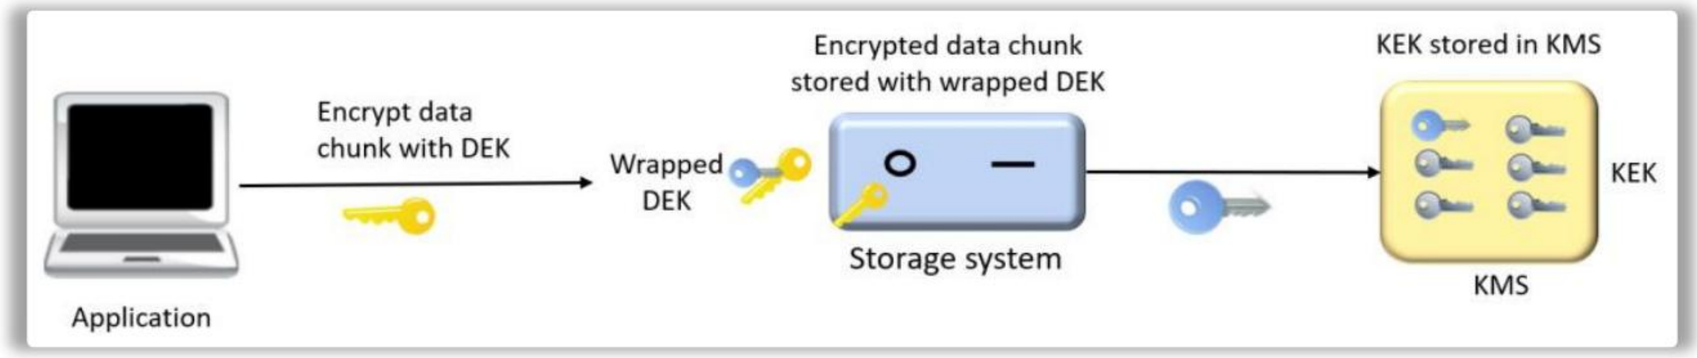


Figure 4.184: Cloud KMS

ACLs allow the user having an authorized role to decrypt the data in Google services. Google stored encrypted keys with data. The DEK are wrapped by KEK for storage. Since one key is enveloping another, it is also termed as Envelope Encryption. These keys are stored and utilized by Google's Central KMS, which tracks and controls data access from the central point. In Google Cloud Storage, the stored customer content is encrypted using Advanced Encryption Standard (AES) such as AES256 or AES 128.

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

Envelope Encryption For generating DEK locally, open the source library such as OpenSSL, mention the cipher type and password for key generation.

- Employ DEK locally to encrypt the data.
- In Cloud KMS, generate a new key or use an existing key that acts as KEK. Use the KEK to wrap DEK.
- Store encrypted data with wrapped DEK.

Envelope Decryption

- Retrieve the stored encrypted data with wrapped DEK.
- To unwrap DEK, utilize the key stored in cloud KMS.
- To decrypt the encrypted data, use the plaintext DEK.

Enabling Cloud KMS

Cloud KMS can be enabled by using GCP User Interface. Before utilizing Cloud KMS, the Google Cloud user should enable this service for their project

Source: https://cloud.google.com

To Enable Cloud KMS with GCP User Interface

- From Google Console dropdown menu, navigate to **Security;** select **Cryptographic Keys** on Menu Tab.

hide01.ir

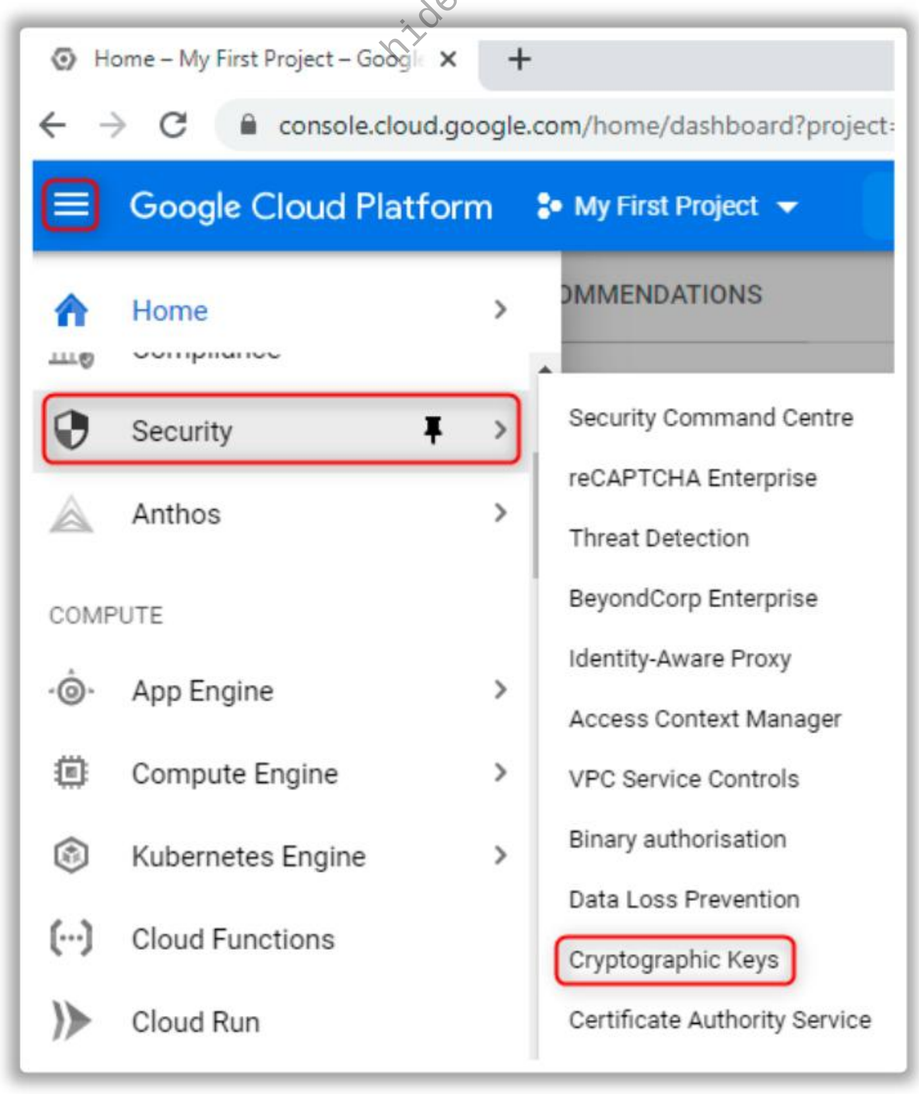


Figure 4.185: Select Cryptographic Keys

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

- Click on **ENABLE** to enable the **Cloud Key Management Service (KMS) API.**

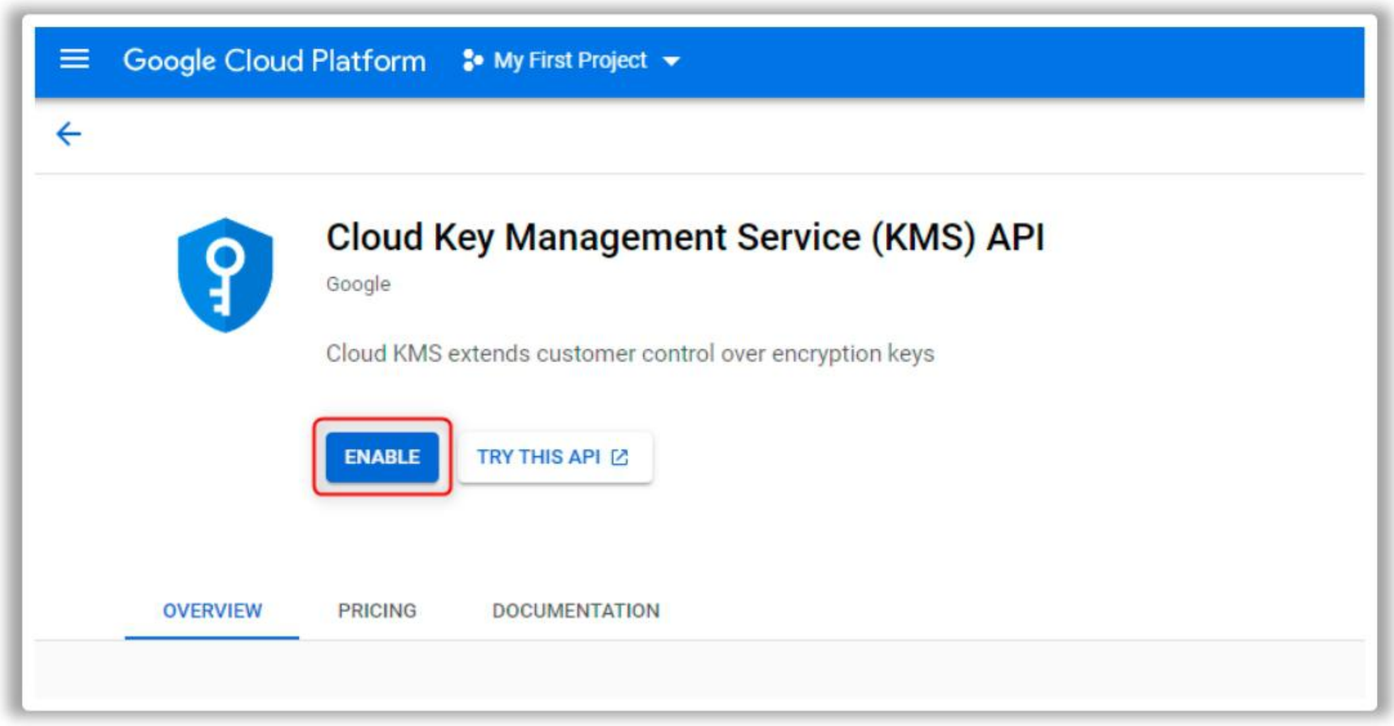
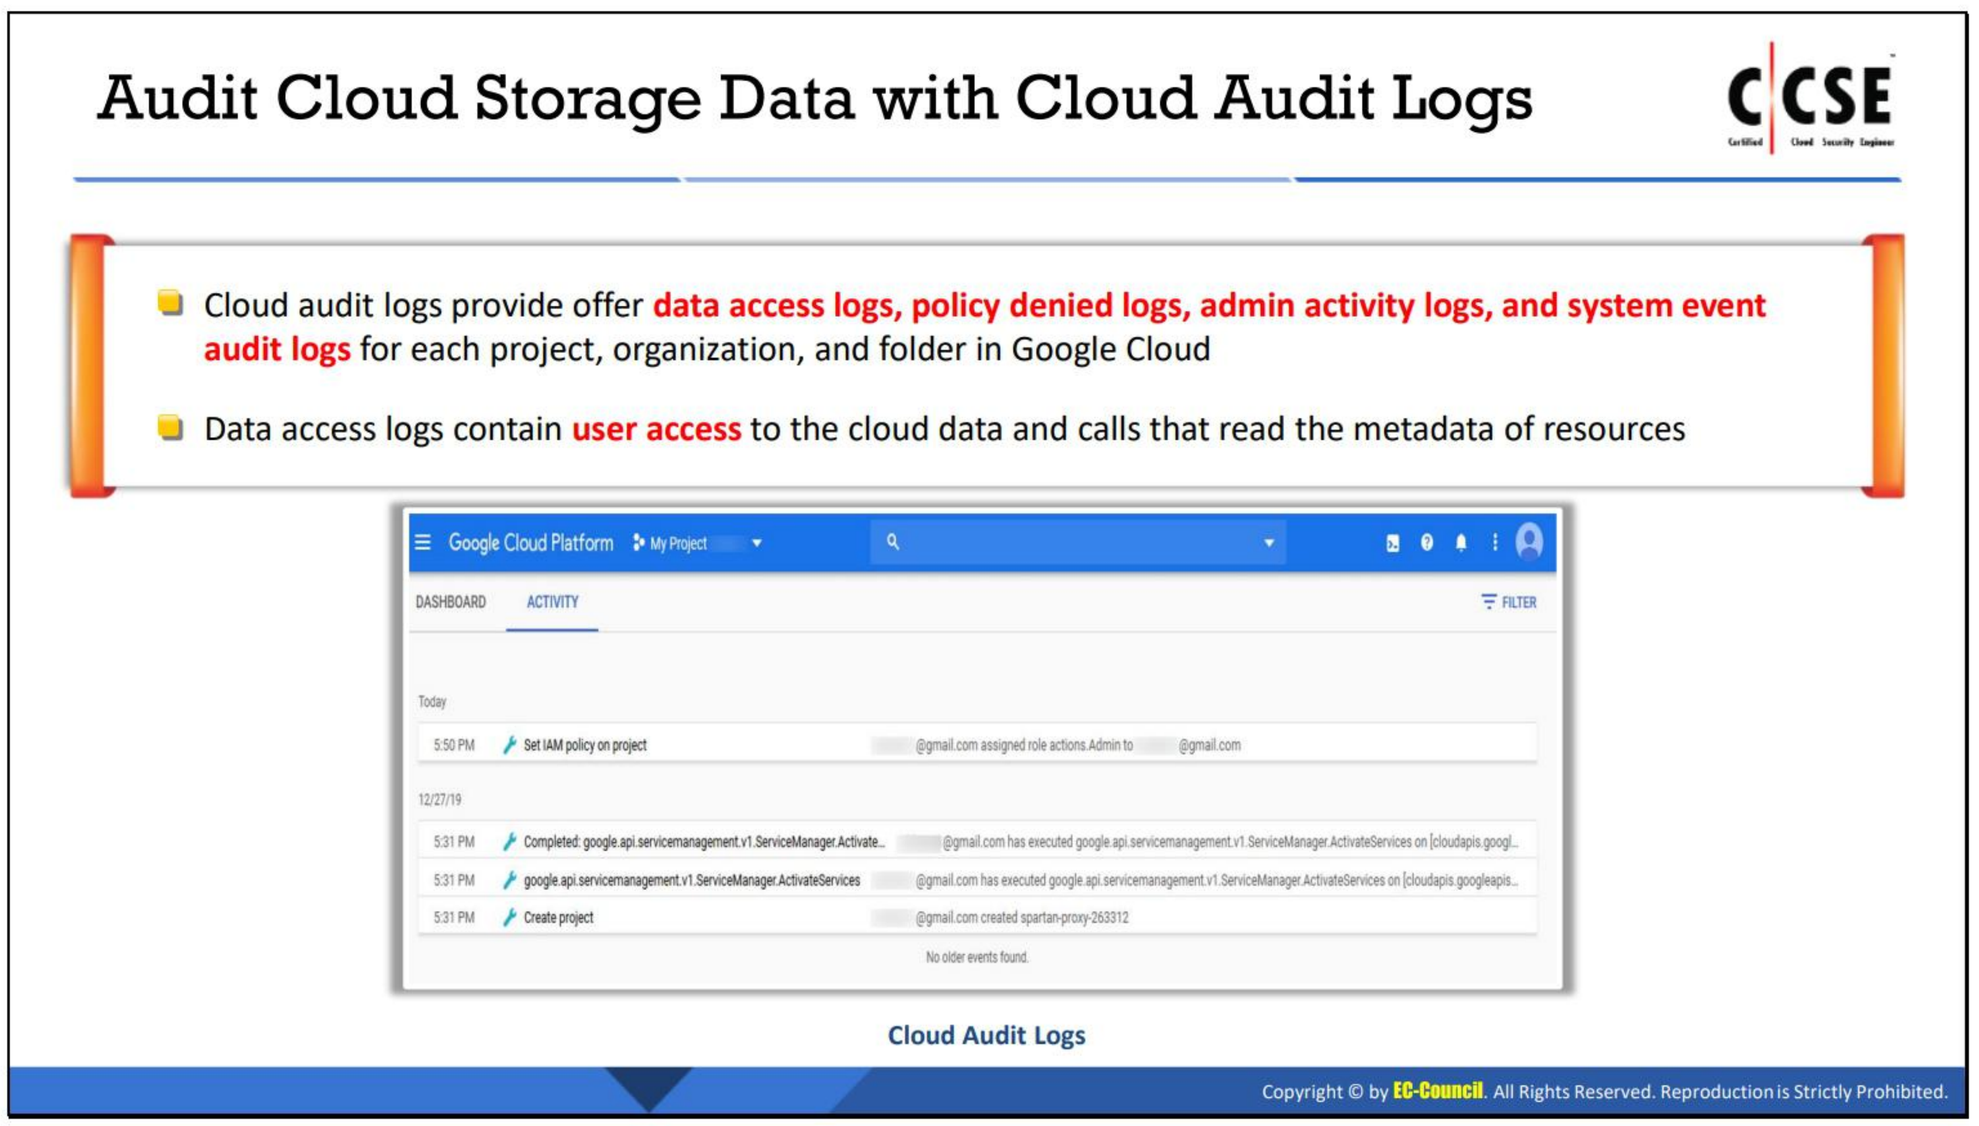


Figure 4.186: Click Enable

hide01.ir

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



Audit the Cloud Storage Data with Cloud Audit Logs

Cloud Audit Logs of Google Cloud Storage help administrator teams maintain the audit details in the GCP. The Cloud Audit logs are stored in highly secured storage space.

h1le01.ir

Cloud Audit Logs provide administrator activity logs, data access logs, system event audit logs, policy denied logs for each project, organization, and folder in Google Cloud. The data access logs can be personalized to meet the organization's needs. Administrator activity logs include administrative activities, data access logs contain user access to the cloud data and calls that read the metadata of resources whereas system event logs include Google Cloud actions that modify the configuration of resources, and policy denied logs include the logs where the Google Cloud denies the user access to resources due to policy violations.

All the Cloud Audit Logs are encrypted at rest. All the audit logs are generated instantly after the event occurrence. It contains the complete information of the users who initiated API calls and from which IP address. This enables the security operations team to assess and respond to any suspicious behavior in an efficient way. Cloud Audit Logs also have incident management tools to monitor and act on incidents.

The security team can export the audit log details into Google Cloud Services or Cloud Logging. Cloud Audit Logs also provide retention, where the audit log can be retained for a specified period and are deleted after that specified time. If the logs need to be stored for a longer duration, the security team can export and store them if required.

Source: https://cloud.google.com

Cloud Audit Logs supports the audit and compliance needs by allowing the user to track the action of administrators in GCP.

- Use Cloud Audit Logs to regularly audit access to service account keys.

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

- Use three audit logs for each GCP project, folder, and organization:
  - Use role **Logging/Logs Viewer** or **Project/Viewer** to know **Administrator Activity Audit** logs.
  - Use roles **Logging/Private Logs Viewer** or **Project/Owner** to know **Data Access Audit** logs.
  - Use roles **Logging/Logs Viewer** or **Project/Viewer** to view **System Event Audit** logs.

Viewing Cloud Audit Logs

- From Google Cloud Console, click **Activity** to view audit logs.

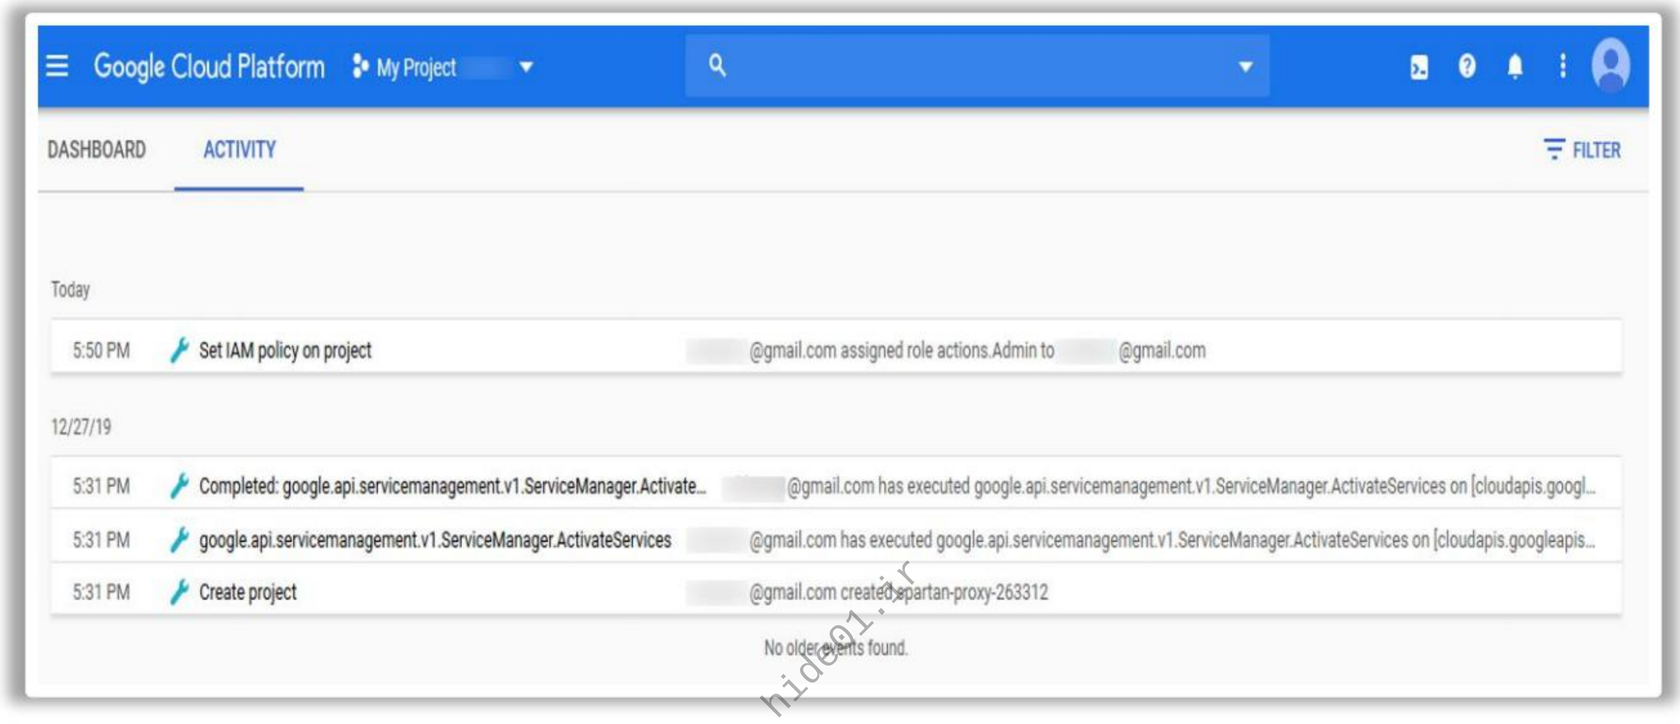
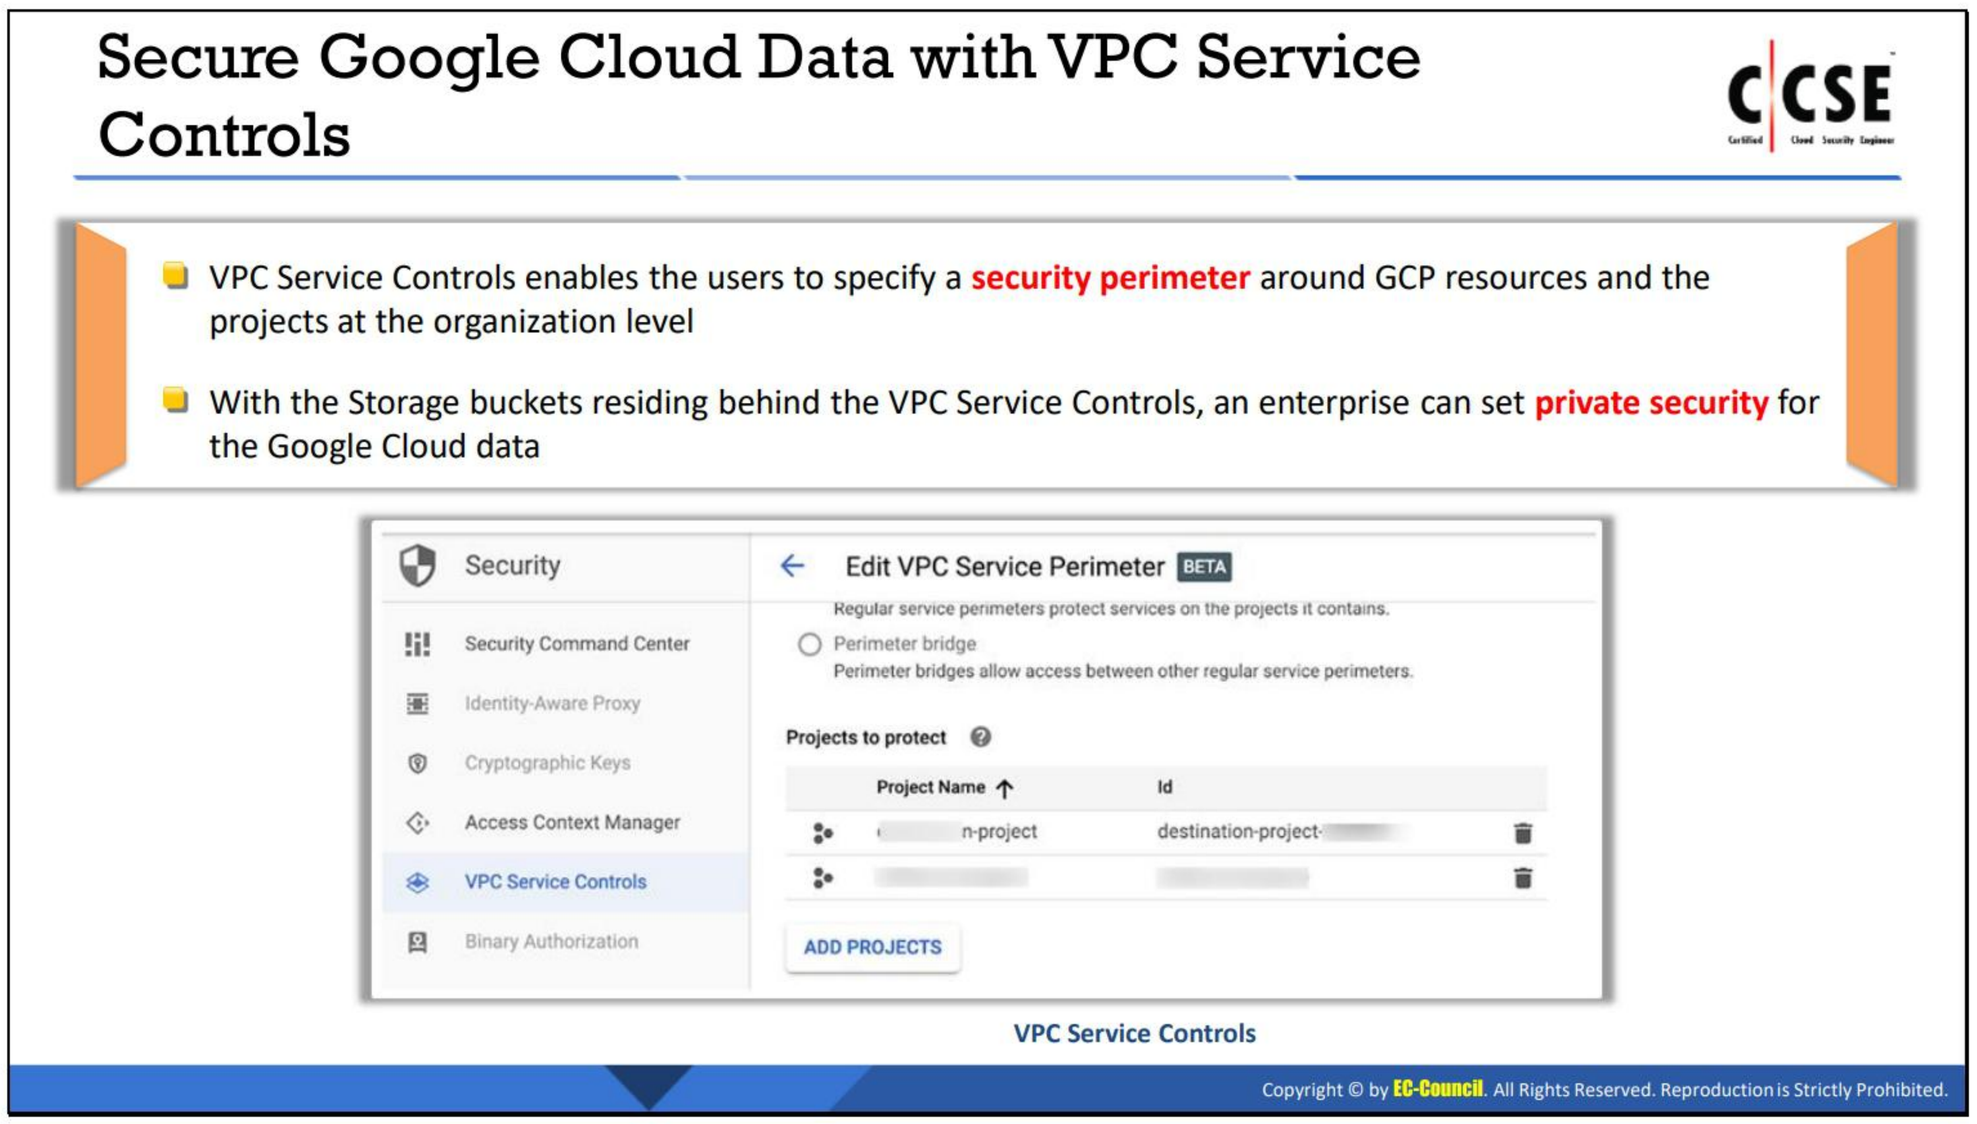


Figure 4.187: Check Activities

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



Secure the Google Cloud Data with VPC Service Controls

VPC Service Controls enables users to specify a security perimeter around GCP resources and the projects at the organization level. It is often described as a firewall for GCP resources. It allows access to the resources only from the authorized IP addresses and verified client devices. Service controls are applied at the network layer of project-owned VPCs.

VPC Service Control is extremely beneficial as it protects the organization from risks such as:

- Access to the resources from unauthorized network
- Data exploitation
- Exposure of private data as public due to wrong IAM policies

Considering all these benefits, an organization can secure its confidential data and can maintain it privately using VPC Service Controls. Using VPC Service Controls, a private connection can be established between Google Cloud resources from VPC and on-premises deployments.

VPC Service Controls allows the security teams to define perimeter controls and achieve security for Google Cloud resources and projects. Users can easily scale their security controls by creating, deleting, and updating resources within the perimeters.

VPC Service Controls also enables the enterprises to set granular access control policies in Google Cloud based on user ID and IP address. This ensures that proper security controls are granted when the resources are accessed through the internet.

It also maintains a log of access denials to identify malicious activity on Google Cloud resources. Flow logs contain the information of all the IP traffic from network interfaces to Compute Engine. It also ensures secure communication as it provides full control over which users can connect to resources. The security teams can also configure service perimeters to manage the

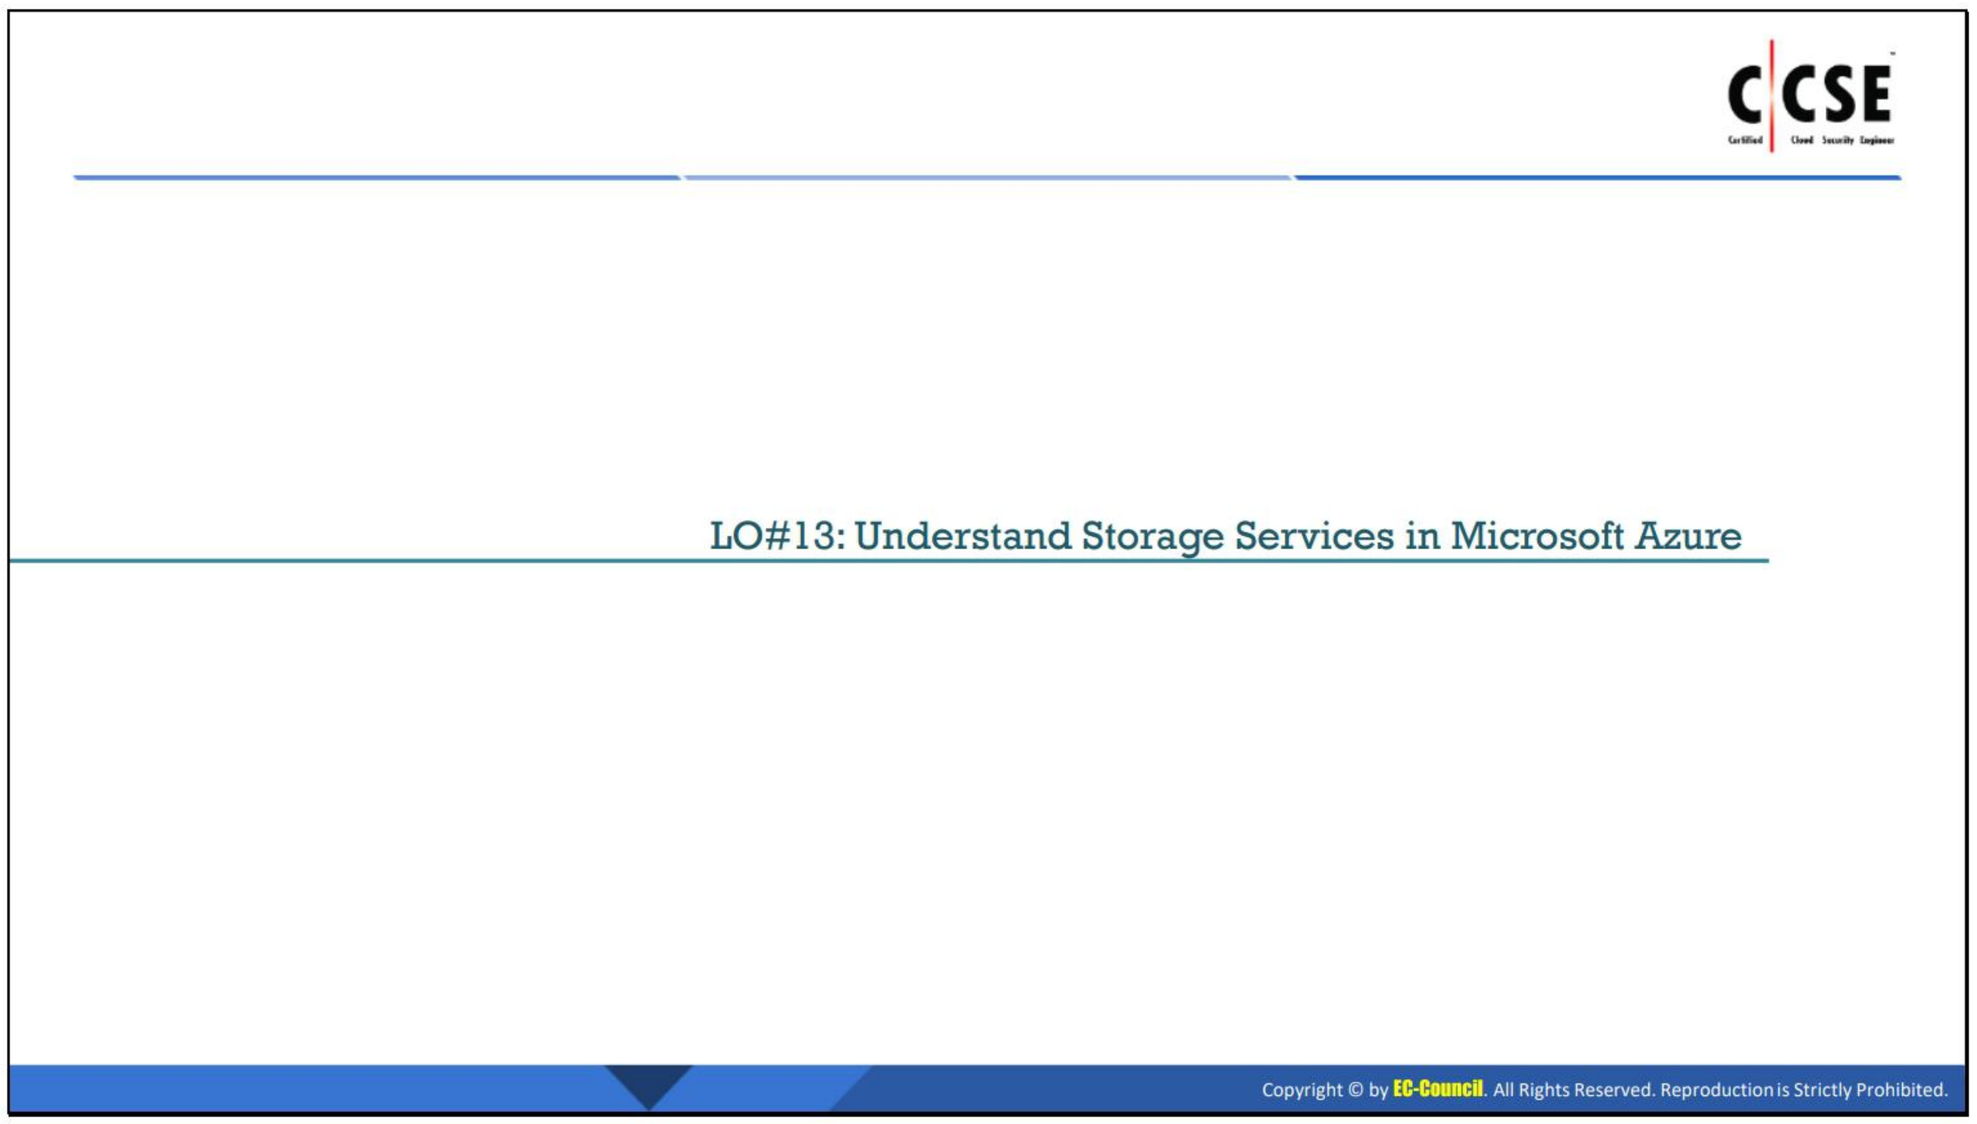
Certified Cloud Security Engineer Copyright © by **EC-Council**
All Rights Reserved. Reproduction is Strictly Prohibited.

communication between virtual machines and Google Cloud resources. This allows free communication within the zone and blocks all the communications outside the perimeters.

With the storage buckets residing behind the VPC Service Controls, an enterprise can set private security for the Google Cloud data.

hide01.ir

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

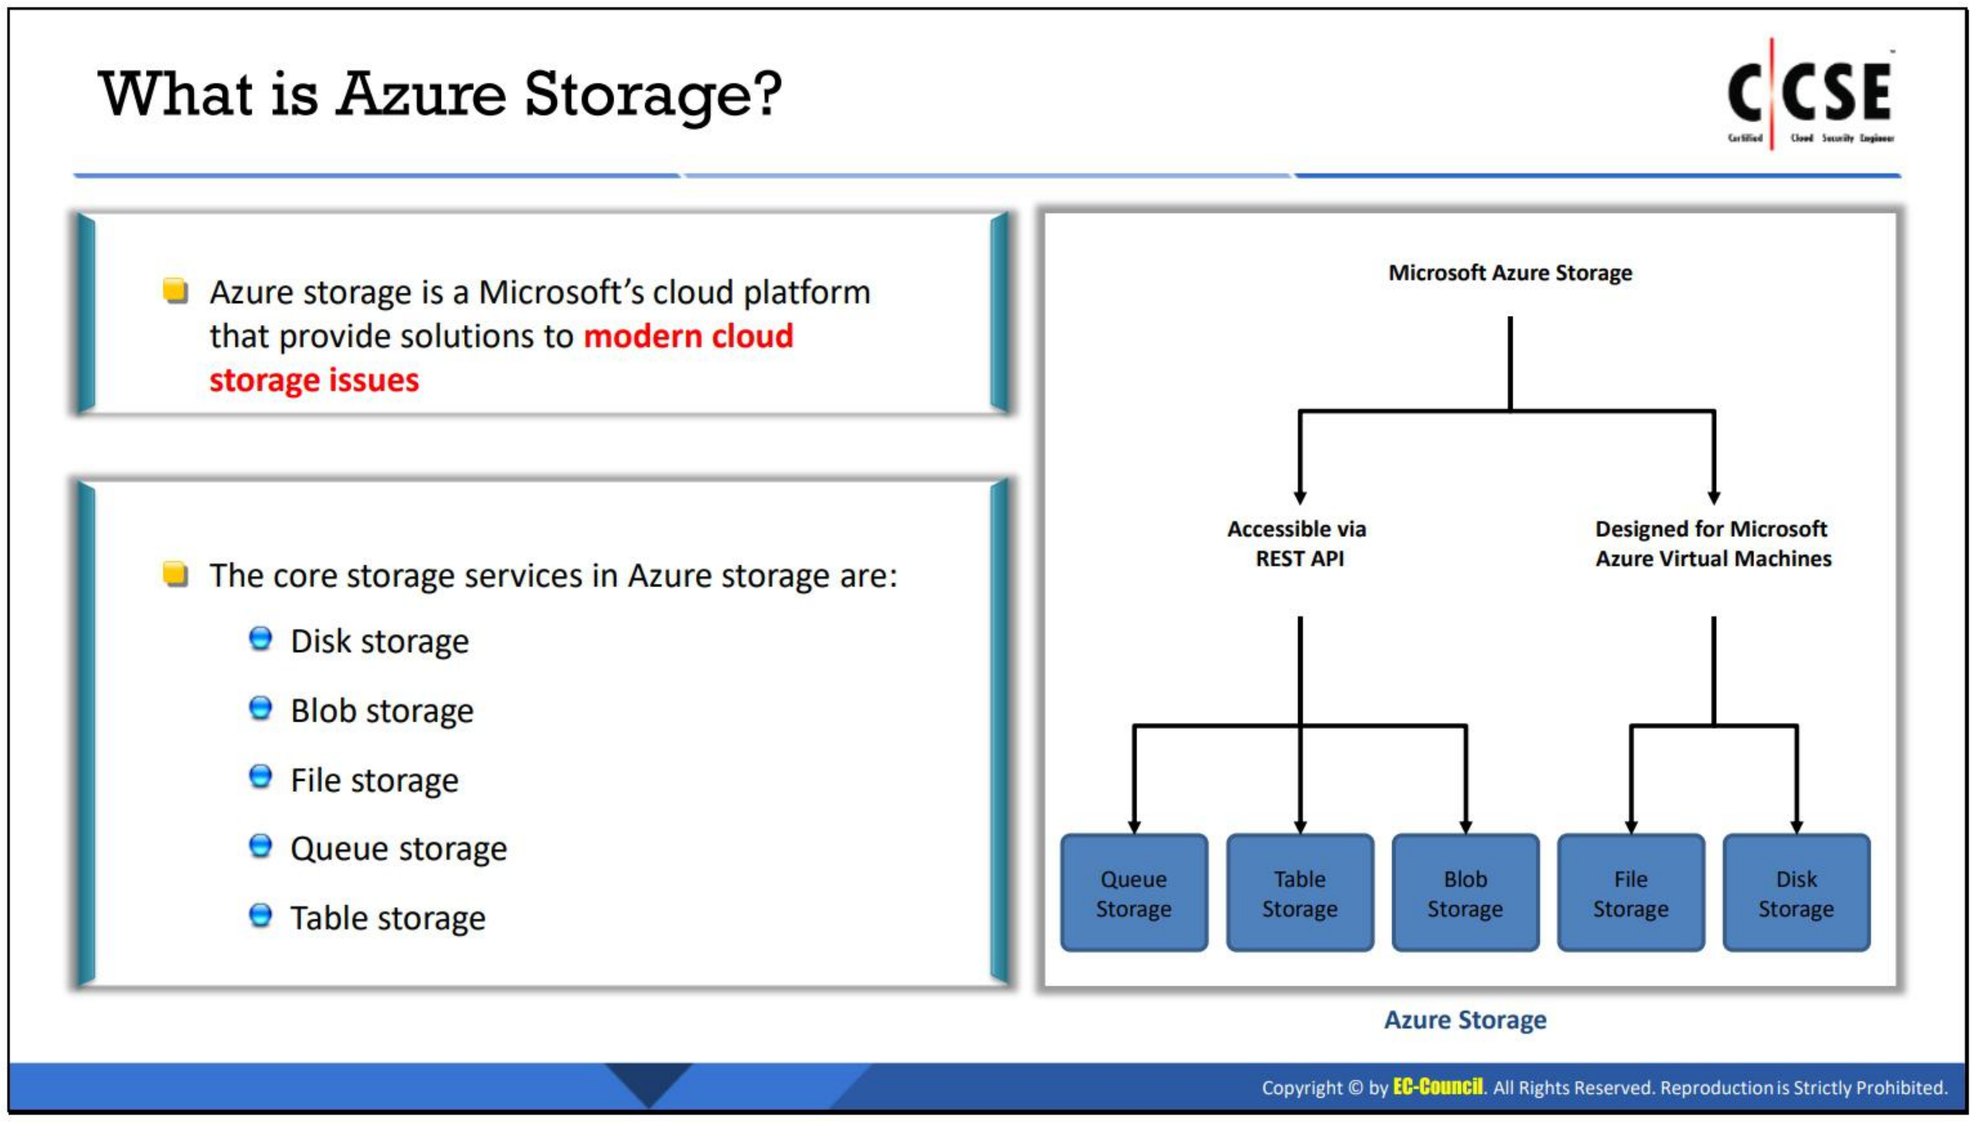


LO#13: Understand Storage Services in Microsoft Azure

The objective of this section is to describe various storage services offered by Microsoft Azure. This section explains Disk Storage, Blob Storage, File Storage, Queue Storage, Table Storage, NetApp Files, and HPC Cache.

hide01.ir

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



What is Azure Storage?

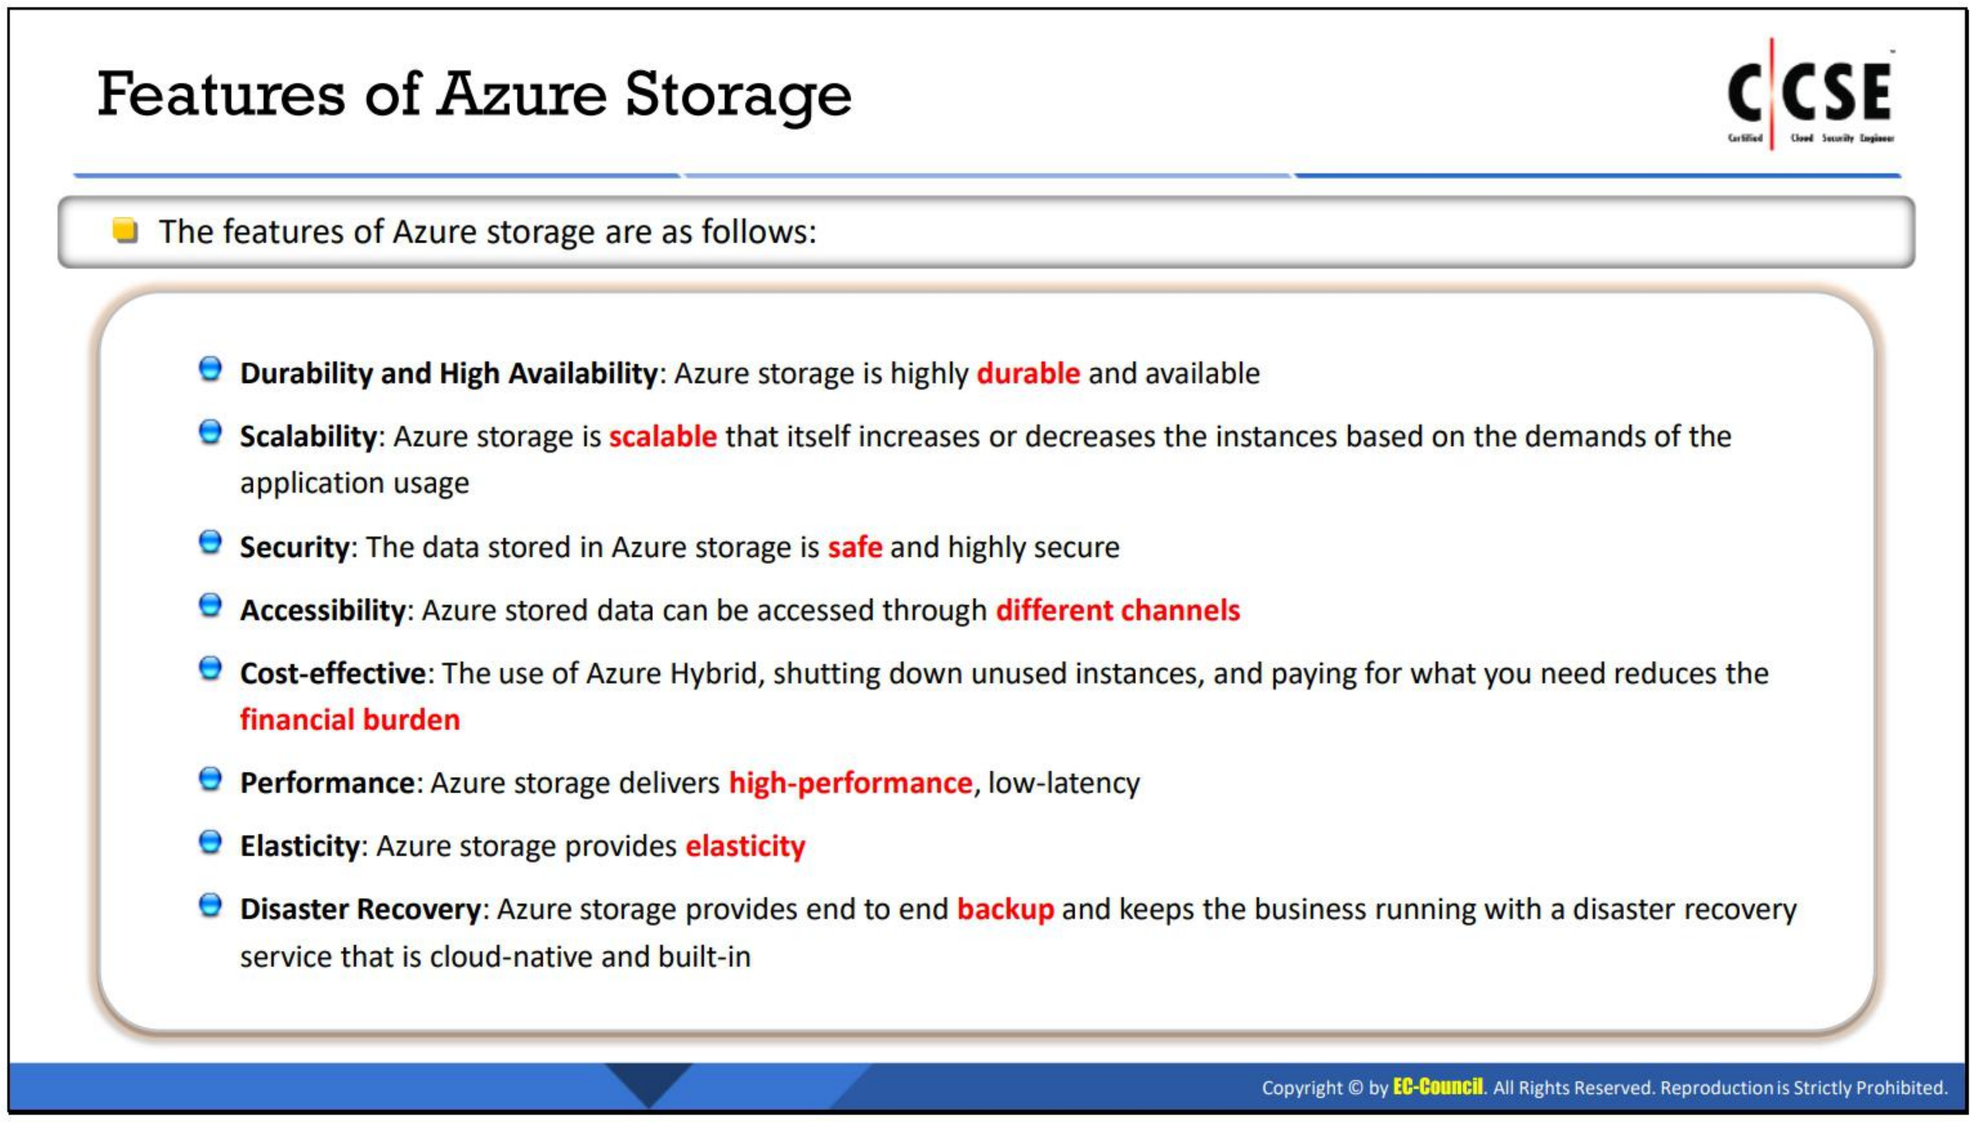
Azure Storage is Microsoft's cloud storage platform that provides solutions to modern cloud storage issues. It offers storage services that are highly available, durable, scalable, and secure. It allows users to store and access data of any format from storage services.

Azure storage services can be used as a pay-as-you-go model. To access storage and services, the Azure storage service user should have an Azure account and access to Azure services depends on the user's subscriptions.

The core storage services in Azure storage are

- Azure Blob
- Azure Files
- Azure Queue
- Azure Tables
- Azure Disks

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



Features of Azure Storage

High availability and durability: Highly available to maintain availability to end-users or customers. Azure storage contains the replica of data as backup in different geo-locations so that even natural disaster occurs the data can be accessed without interruption from a secondary location.

Scalability: Based on the customer's requirement, the data stored in Azure storage is scalable.

Security: Shared access signature (SAS) can restrict data access. It allows the client to have access to the limited required data instead of granting access to the entire storage. Azure uses a shared key model in authenticating a user. Data in storage is encrypted at rest and in transit with respective algorithms. This shows that the data in Azure is safe and secure.

Accessibility: Data can be accessed through different channels, that is, HTTP, HTTPS, or Azure specified applications like Azure PowerShell, Azure portal, etc.

Cost-effective: The use of Azure Hybrid Benefit, shutting down unused instances, and paying for what you need significantly reduces the financial burden.

Performance: Azure storage delivers high performance and low latency, that is, no lag or minimum lag feature.

Elasticity: Azure storage services provide elasticity.

Disaster recovery: Azure provides end to end backup and keeps the business running with a disaster recovery service that is cloud-native and built-in. This allows the Azure Storage service user to recover from the secondary/slave backup servers when a disaster occurs in the primary/master server geolocation.

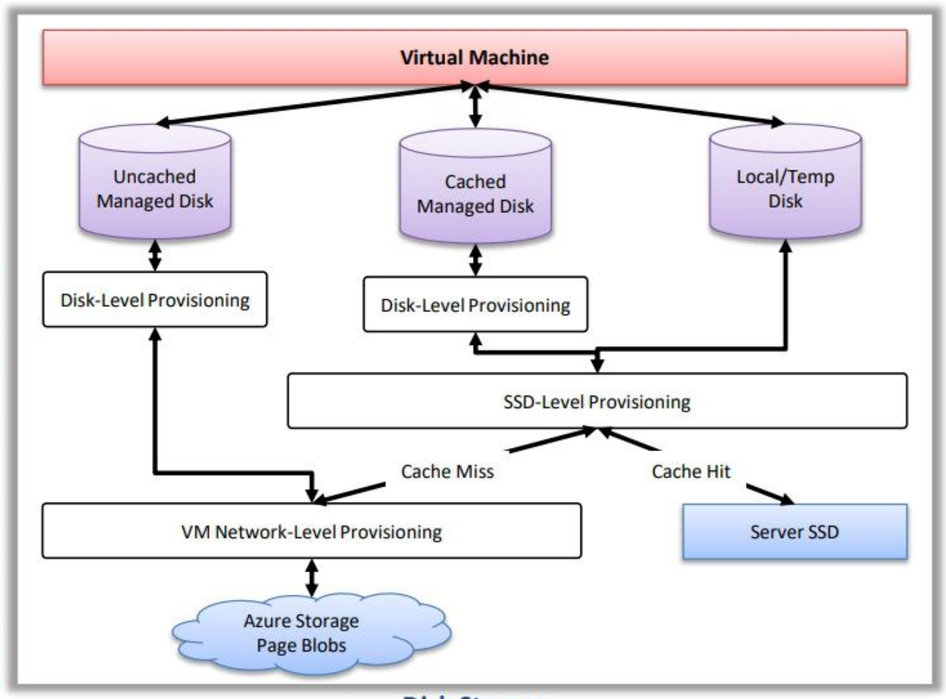
Certified Cloud Security Engineer Copyright © by **EC-Council**
All Rights Reserved. Reproduction is Strictly Prohibited.

Types of Azure Storage

There are various types of Azure Storages as follows:

Disk Storage

- Azure storage provides managed disks that are **block level storage volumes** which are used in virtual machines
- Disk storage is available in Azure Standard Hard-Disk Drives (HDD), Azure Standard Solid-State Drives (SSD), Azure Premium SSD, and Azure Ultra Disk variants



Disk Storage

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

Types of Azure Storage (Cont'd)

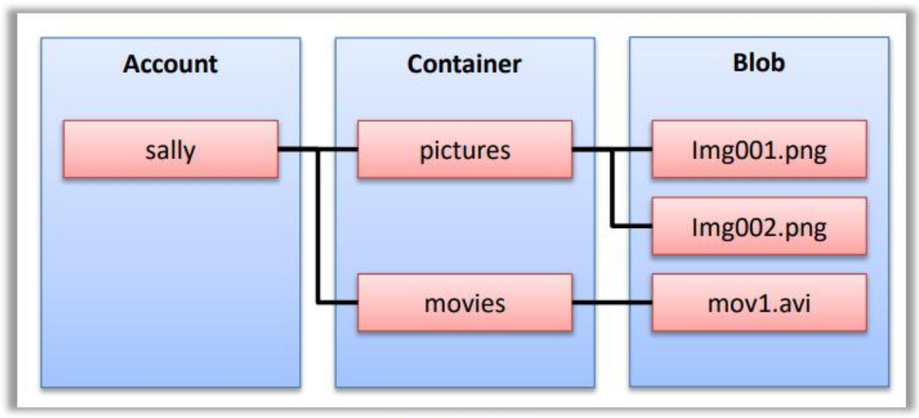
hide01.ir

Blob Storage

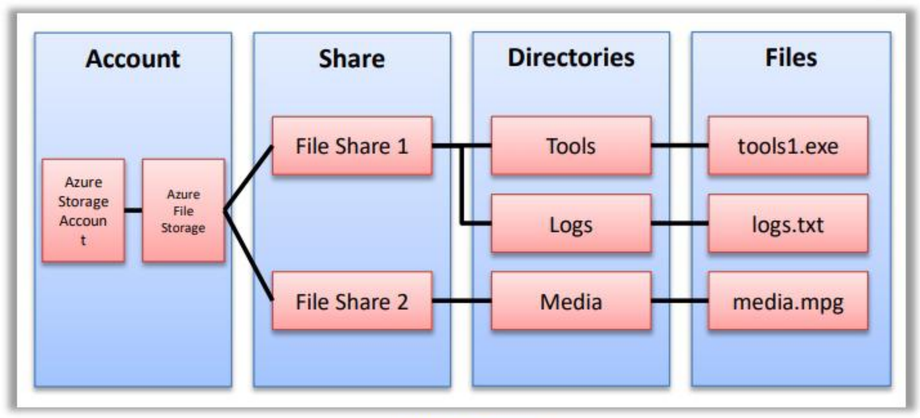
- It is **object storage solution** for the Azure storage cloud
- It stores and **processes** images, documents, log files, backup and restore, disaster recovery, and stream video and audio files

File Storage

- It offers fully **managed file shares** in the Azure storage cloud
- This can be accessed via **different protocols** like server message block (SMB) and network file system (NFS)



Blob Storage



File Storage

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.

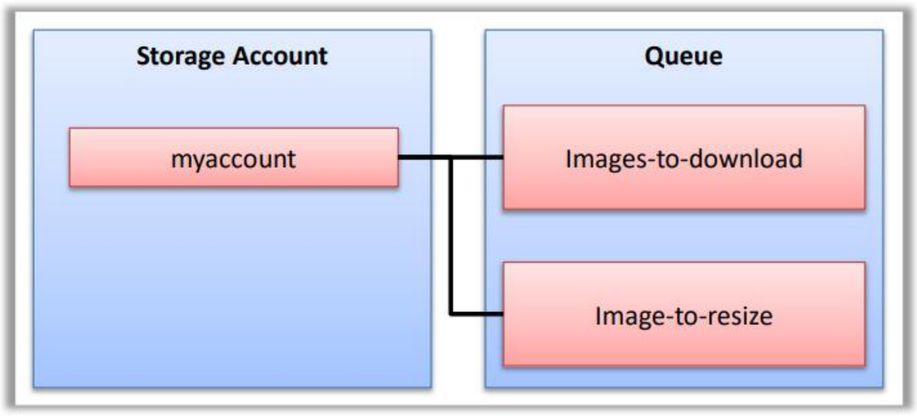
Types of Azure Storage (Cont'd)

Queue Storage

- It stores large number of **messages** and can be accessed from anywhere through HTTP/HTTPS
- It can contain millions of messages using the limit of **storage** assigned to the account

Table Storage

- It stores large amount of **structured data**, non-relational data
- It can process the **data query** quickly using clustered index



Queue Storage

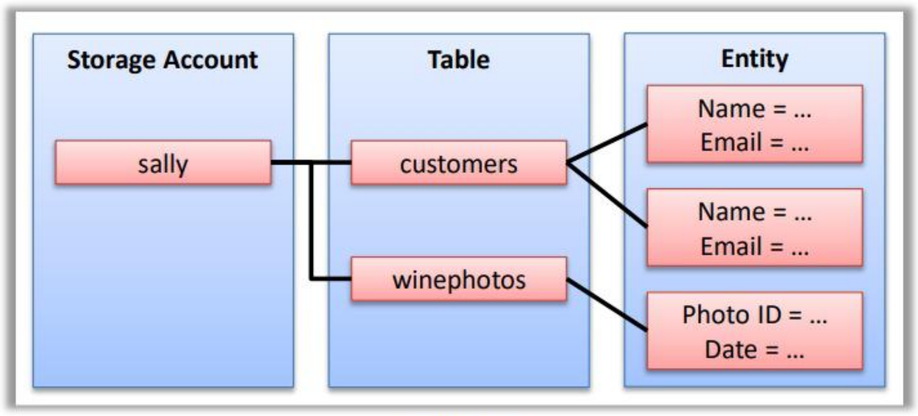


Table Storage

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

Types of Azure Storage (Cont'd)

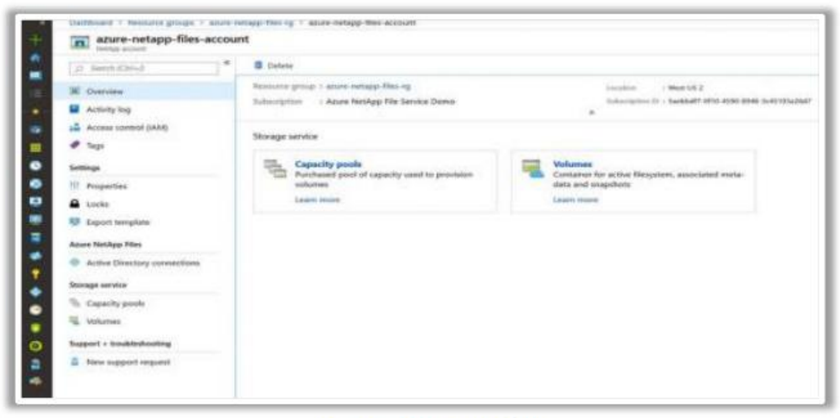
hide01.ir

Azure NetApp Files

- It provides easy way for **storage professionals** to migrate, store, and run complex applications that are required by enterprises
- It can be managed by **Azure specific services** like PowerShell, REST API, etc

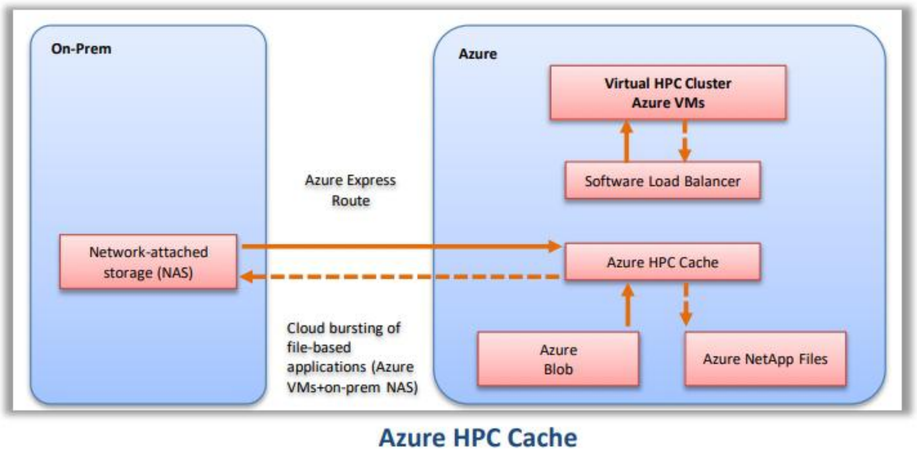
Azure HPC Cache

- It lets storage professionals to **compute** seamlessly with higher capacities and is scalable to meet the peak demands of workload
- This increases the productivity with low **latency**
- It speeds up calculations in risk analysis work loads also



Source: https://www.microsoft.com

Azure NetApp Files



Azure HPC Cache

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

Types of Azure Storage

Azure storage services are described below:

Disk Storage: Disk Storage is a virtual storage service that is available in standard HDD and SSD options, where SSD drives perform operations much faster. Both the storage variants will be having built-in security.

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

It is available in Azure Standard Hard-Disk Drives (HDD), Azure Standard Solid-State Drives (SSD), Azure Premium SSD, and Azure Ultra Disk variants. Any number of disks can be attached to a VM based on scalability and performance.

Managed disks will have security by default that they are encrypted with a platform Managed Key at Rest, which is SSE. The disks in the VMs can be encrypted using Azure disk encryption that may contain OS and data disks.

Azure Key Vault is secure storage for the keys used to encrypt data-at-rest in Azure services. Azure Disk encryption can be done with the Azure Key Vault, which will control and manage the encrypted keys. Encryption of virtual disks requires the creation of Azure Key Vault for storing cryptographic keys, which will encrypt or decrypt the virtual disks. Cloud security professionals can use either platform managed key or CMK.

Blob Storage*:* Azure Blob Storage is designed to store massive unstructured data like images, documents, log files, backup and restore, audio and video streaming files. All the stored unstructured data can be accessed from anywhere through HTTP/HTTPS and with some Azure applications like Azure PowerShell, Azure storage client library where the library supports different computer languages.

The three types of Blob Storage resources are storage account, the container in a storage account, and Blob, which is present in the container.

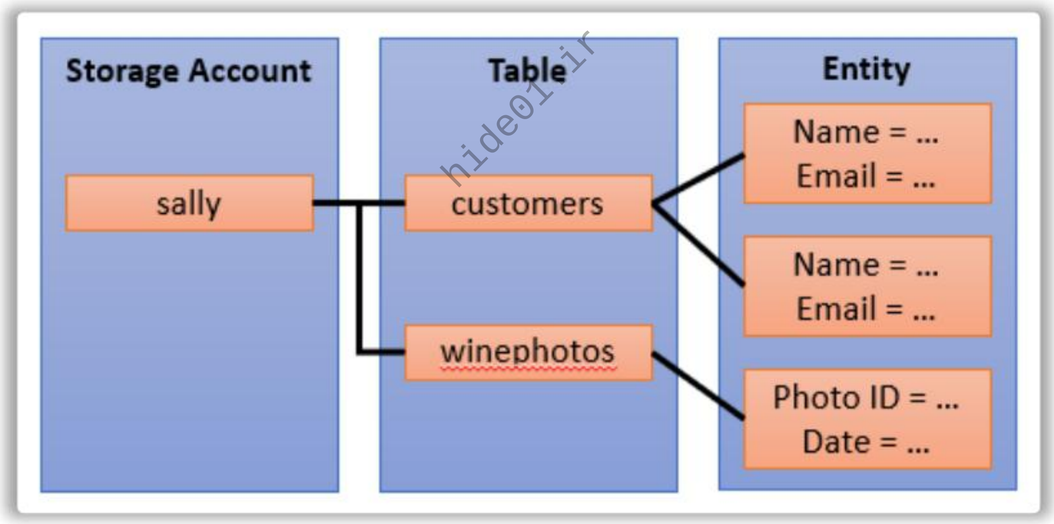


Figure 4.188: Blob Storage Resources

There are three types of Azure storage blobs.

- **Block blobs**: This storage can be used to store images, log files, and more.
- **Append blobs**: It helps to log data from VMs. It is composed of several blocks of different sizes.
- **Page blobs**: Page blobs stores the images of Virtual Hard Drive (VHD) and allows them to act as a disk in the virtual machine.

There are different ways to move existing data to Blob Storage that includes AzCopy, Azure storage Data Movement Library, Azure data factory, etc.

Blob Storage is of three different tiers, they are hot access tier, cool access tier, archive access tier.

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

- **Hot access tier**: It features to store data that require regular access. It has the lowest access cost (read-write) and the highest storage cost.
- **Cool access tier**: It stores data for at least 30 days and is not accessed frequently. Compared to the hot access tier it has higher access cost and lower storage costs.
- **Archive storage***:* Azure Archive storage is used to store files that are rarely used and must be required for a longer period of a minimum of 6 months i.e., 180 days. Archive data is separated from online to cold storage that reduces the cost for storage, but it costs higher for access. The data in the archive cannot be read or modified as it is stored in offline mode.

The data stored in the archive access tier includes

- Backups that are not used frequently.
- Compliance/audit data that is accessed rarely.

Azure Files: It offers shared network file storage in Azure cloud and allows administrators to retrieve the stored data via industry-standard protocols like SMB and Network File System protocol (NFS). Azure files storage can be used as a replacement or additional storage for on-premises file servers or NAS devices. It can ease the process of file sharing such as creating a centralized shared resource where the stored tools, applications, and config files can be retrieved by developers and sysadmins. SMB file shares can be accessed from Windows, Linux, and macOS clients. NFS file shares can be accessed from Linux and macOS clients.

It is easy to manage and is secure as it encrypts data-at-rest and in-transit using SMB and HTTPS.

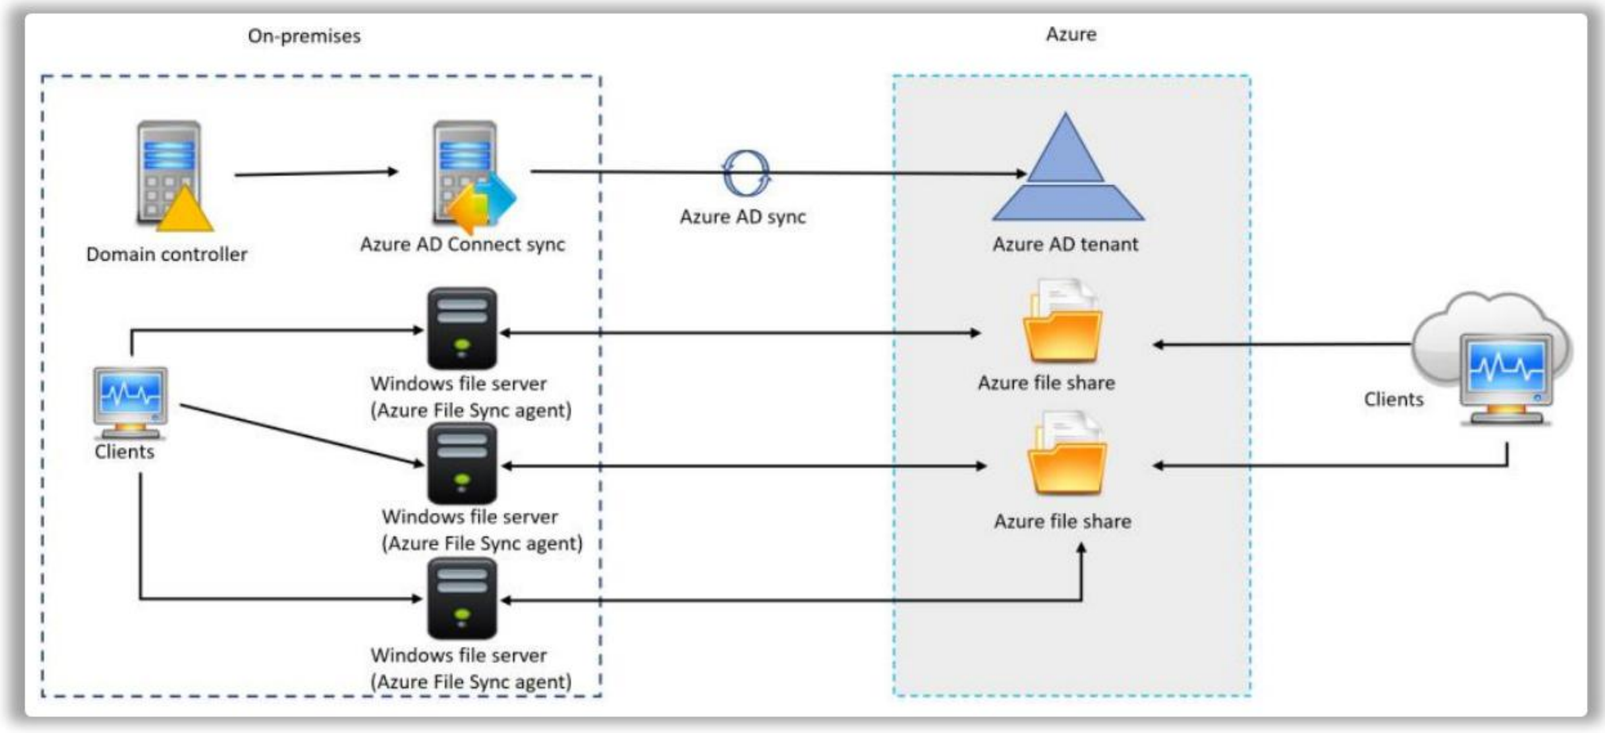


Figure 4.189: Azure Files

Queue Storage: Queue Storage is used to store messages that will be processed asynchronously. It can contain millions of messages using the limit of storage assigned to the account and the messages can be accessed from anywhere through HTTP/HTTPS.

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

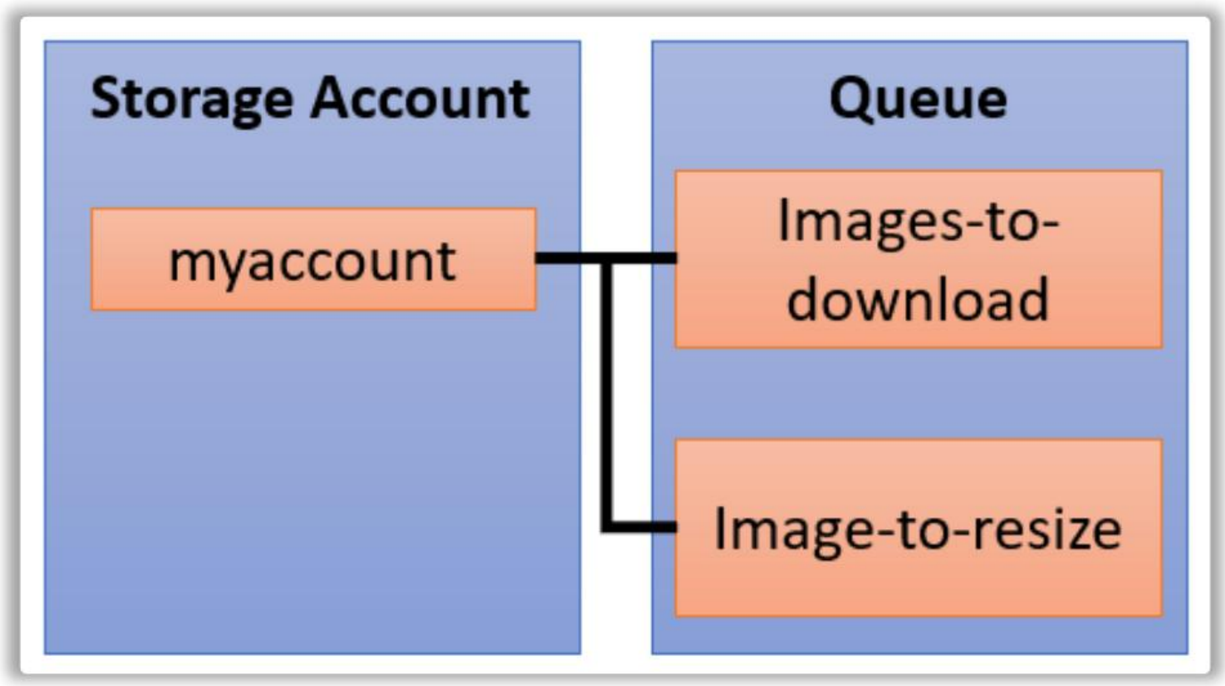


Figure 4.190: Queue Storage

Queue Storage is composed of few elements:

- **Storage account**: This will contain all storage services of Azure that you have subscribed to.
- **Queue**: Queue is a set of messages.
- **Message**: Message can be any information, it may be a text message that triggers an application event or information of the event that happened in the application. The maximum capacity of the message is 64 KB and it can remain in a queue for a maximum of 7 days.

Queue Storage reduces the data loss possibility due to timeouts on a data store or long-running processes.

Data Box*:* Azure data box is a storage device that is designed to transfer terabytes of data. The maximum usable storage capacity is 80TB. It can be used to either import or export data from the servers to the data box or vice versa. It features AES 256-bit encryption for safer transit of data.

The other types of box storage are Data Box Disk, Data Box Heavy, and Data Box Gateway.

- **Data Box Disk**: This storage device can import or export data which is an SSD variant that has a capacity of 8 TB and users get it in a pack of 5 for a total of 40 TB. This storage has 128-bit encryption.
- **Data Box Heavy**: It is a device that is designed to lift 1 PB (Petabyte) of data to the cloud.
- **Data Box Gateway**: It is also used to transfer data to and from Azure but this a virtual appliance that needs signup to the service.

Source: https://azure.microsoft.com

Azure Netapp Files*:* It is a Microsoft Azure native, high-performance file storage service for business applications. It supports multiple protocols of both Linux and Windows applications. Netapp can be set up and managed by Azure services like CLI, PowerShell, or REST API.

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

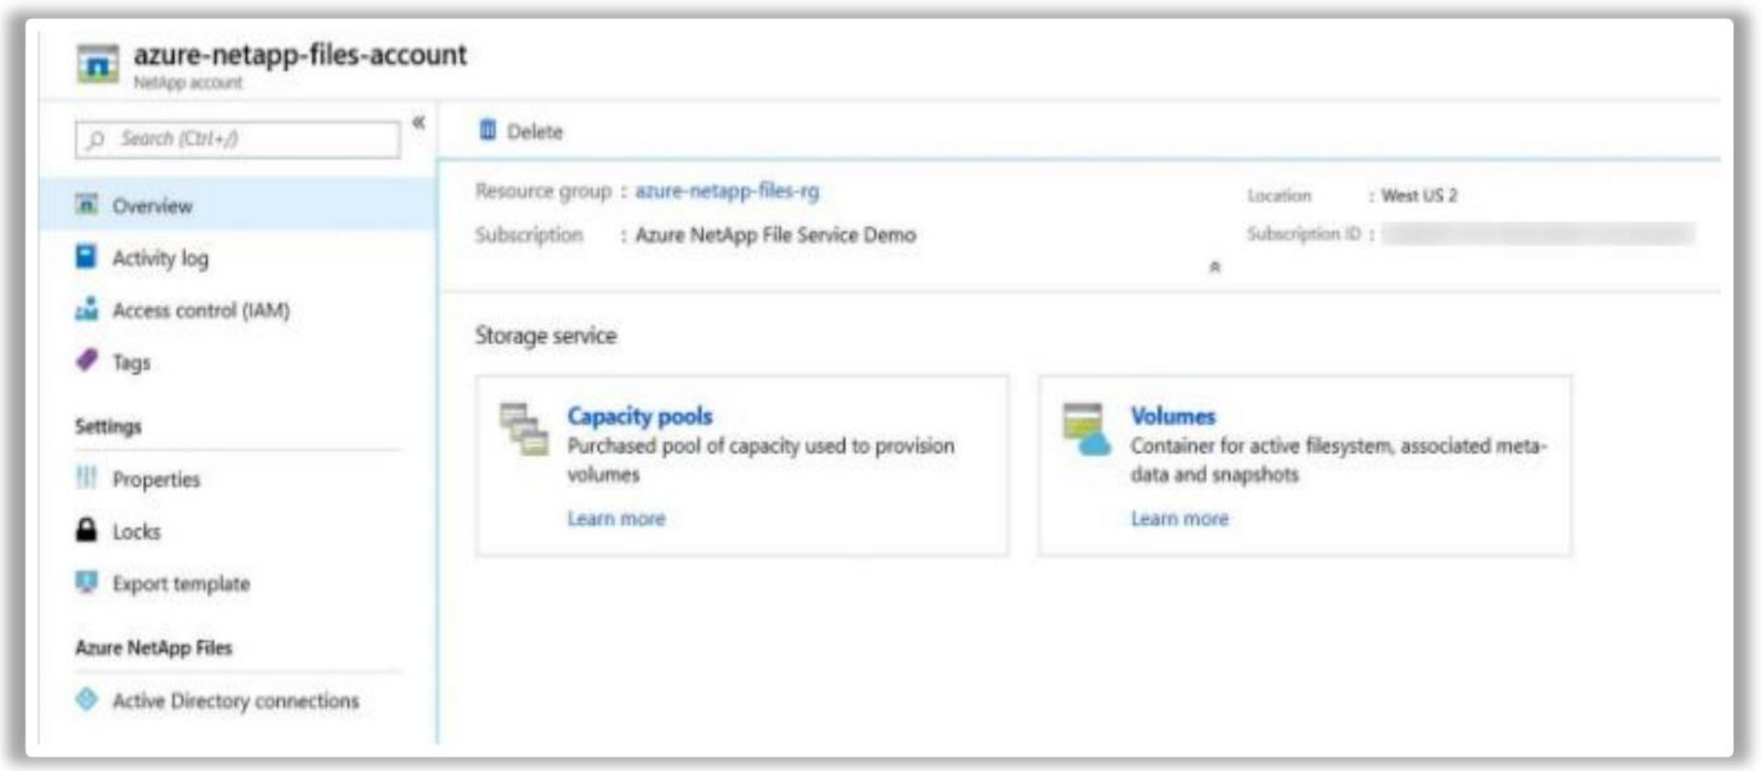


Figure 4.191: Azure Netapp Files Console

Azure HPC Cache: Azure HPC Cache accelerates data for high-performance computing. It can present hot data required by applications in a singular directory structure so that client complexity can be reduced.

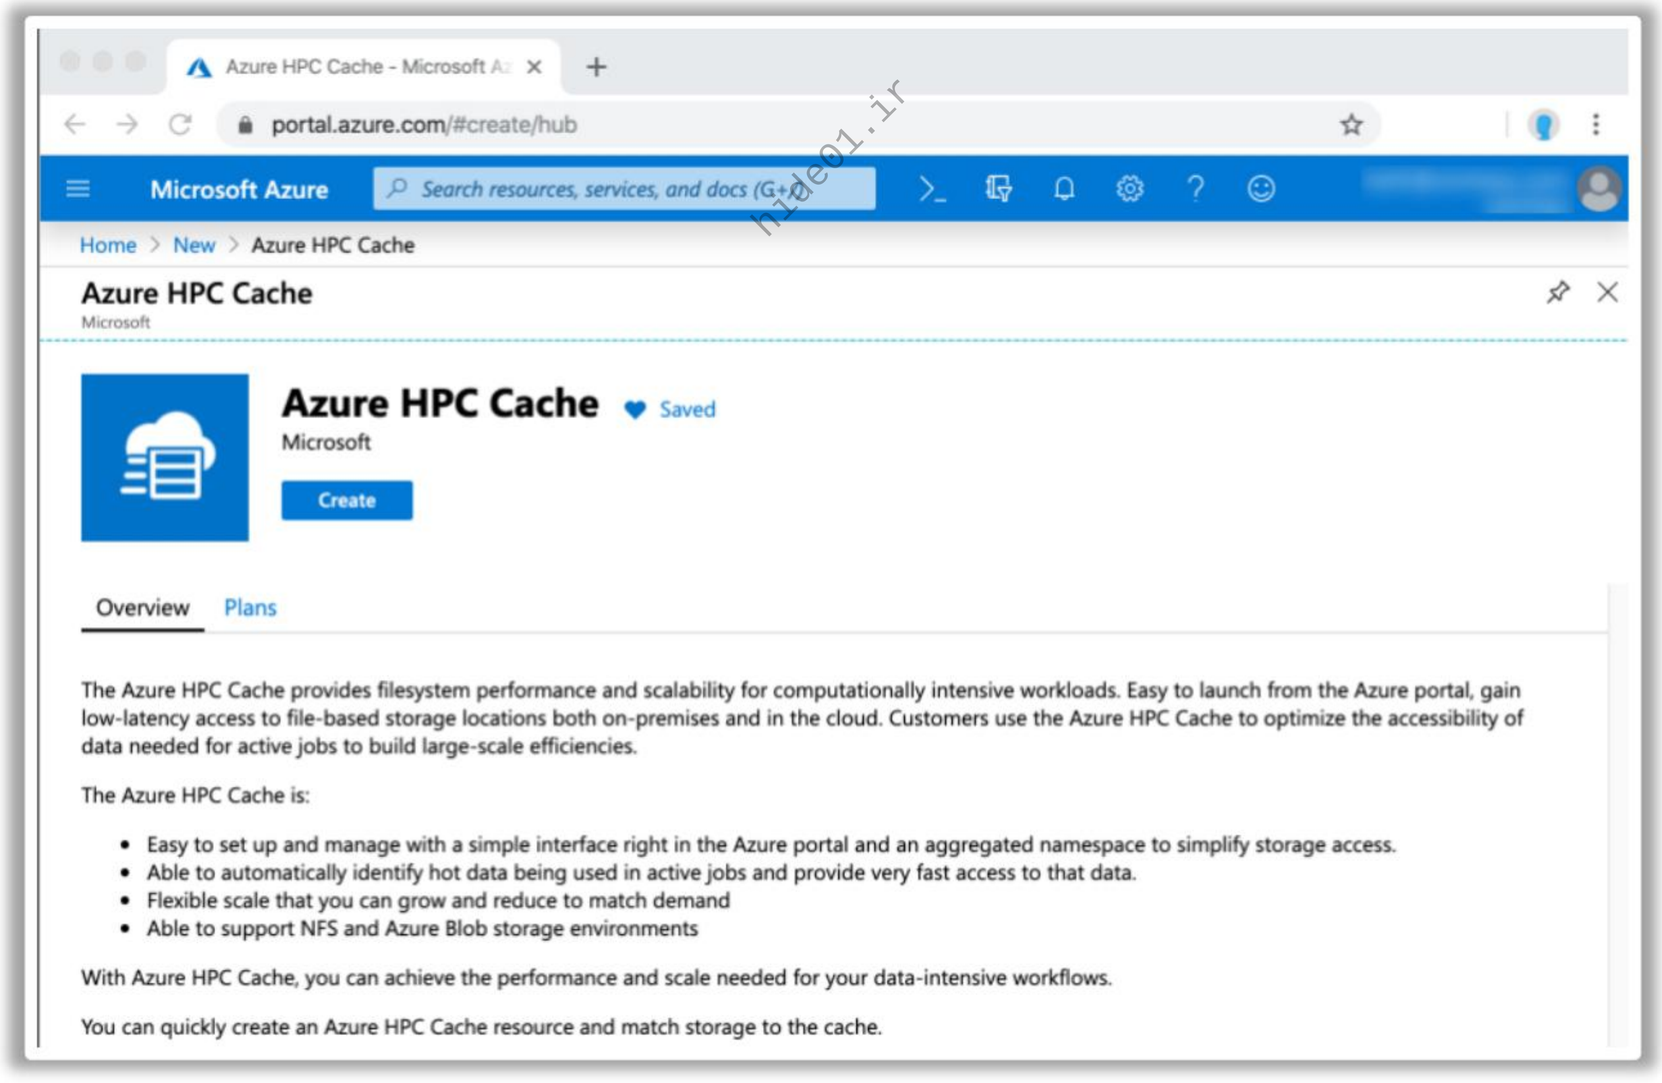
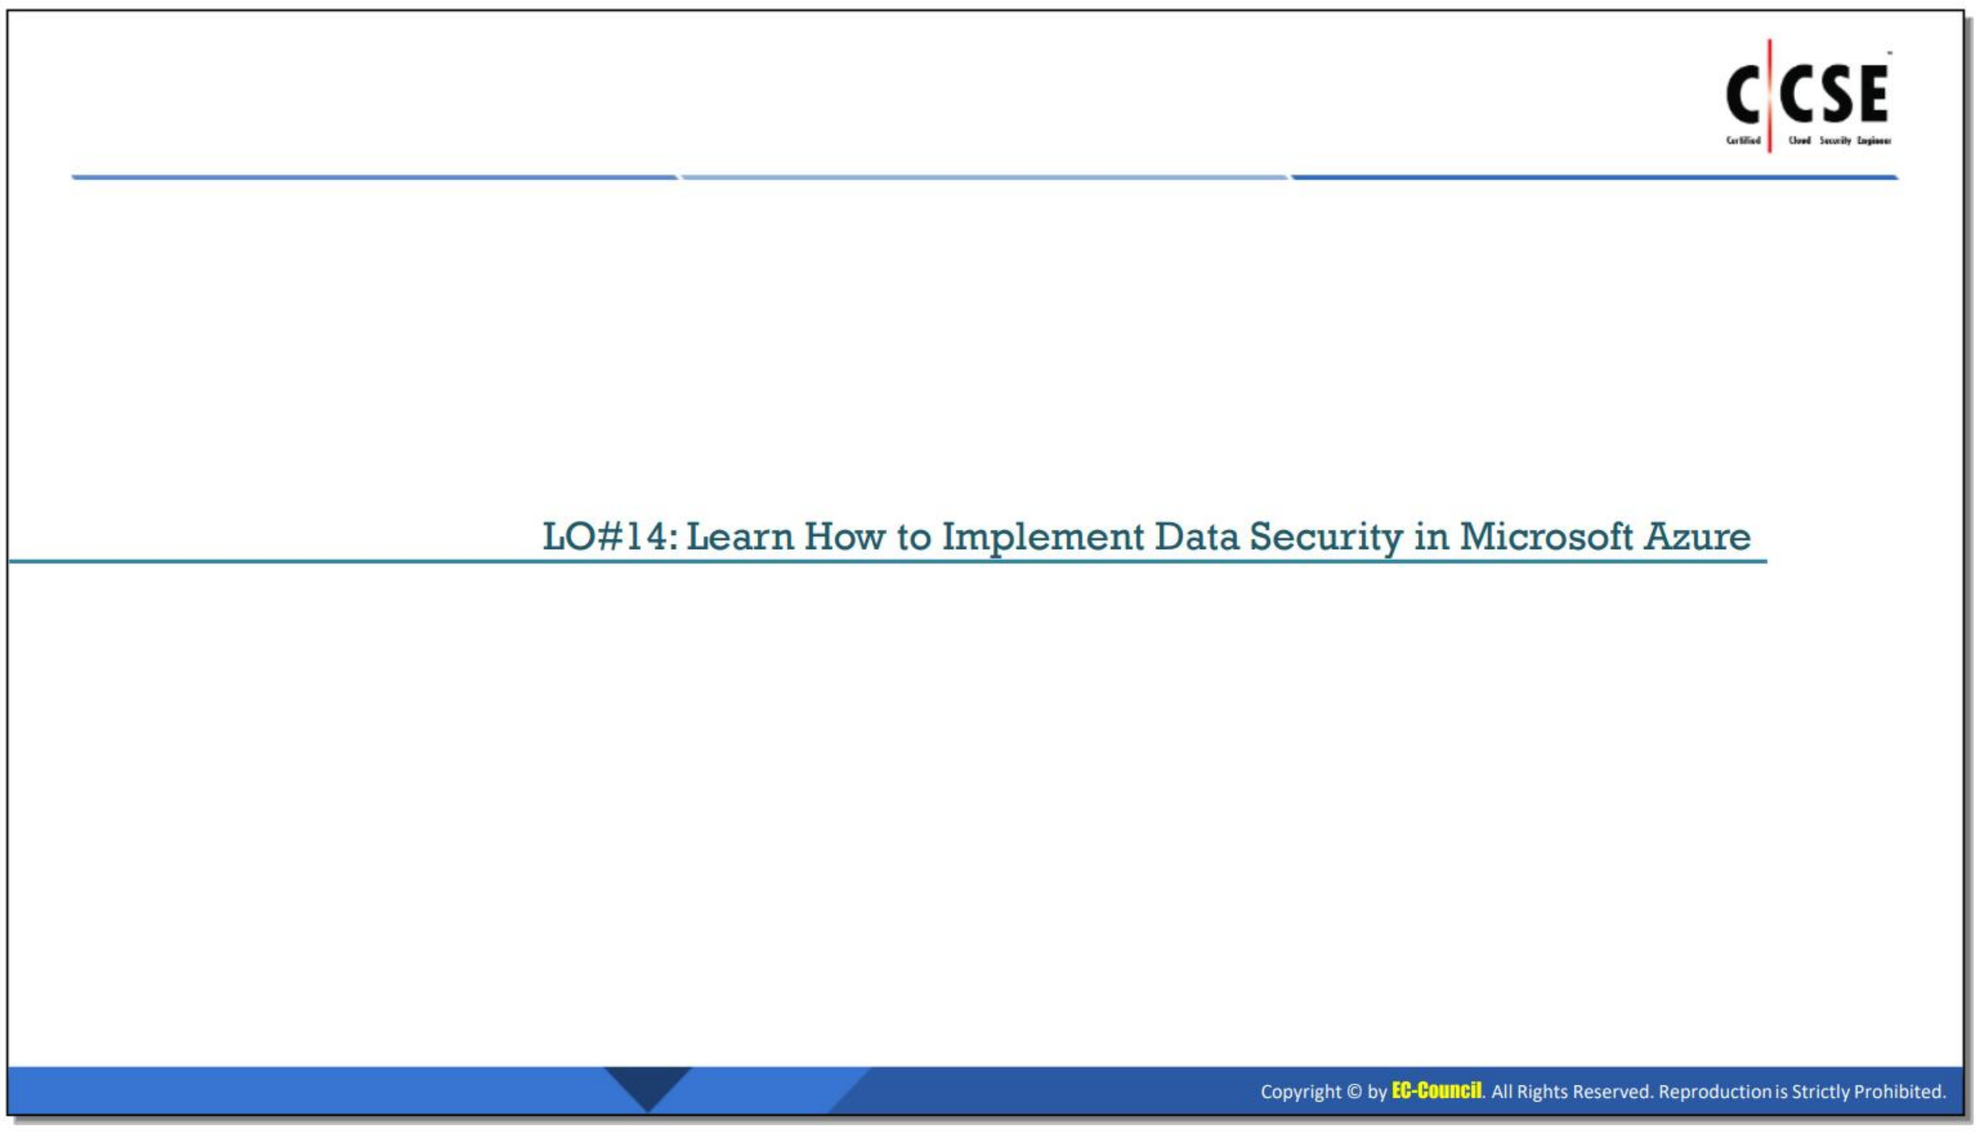


Figure 4.192: Azure HPC Cache

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

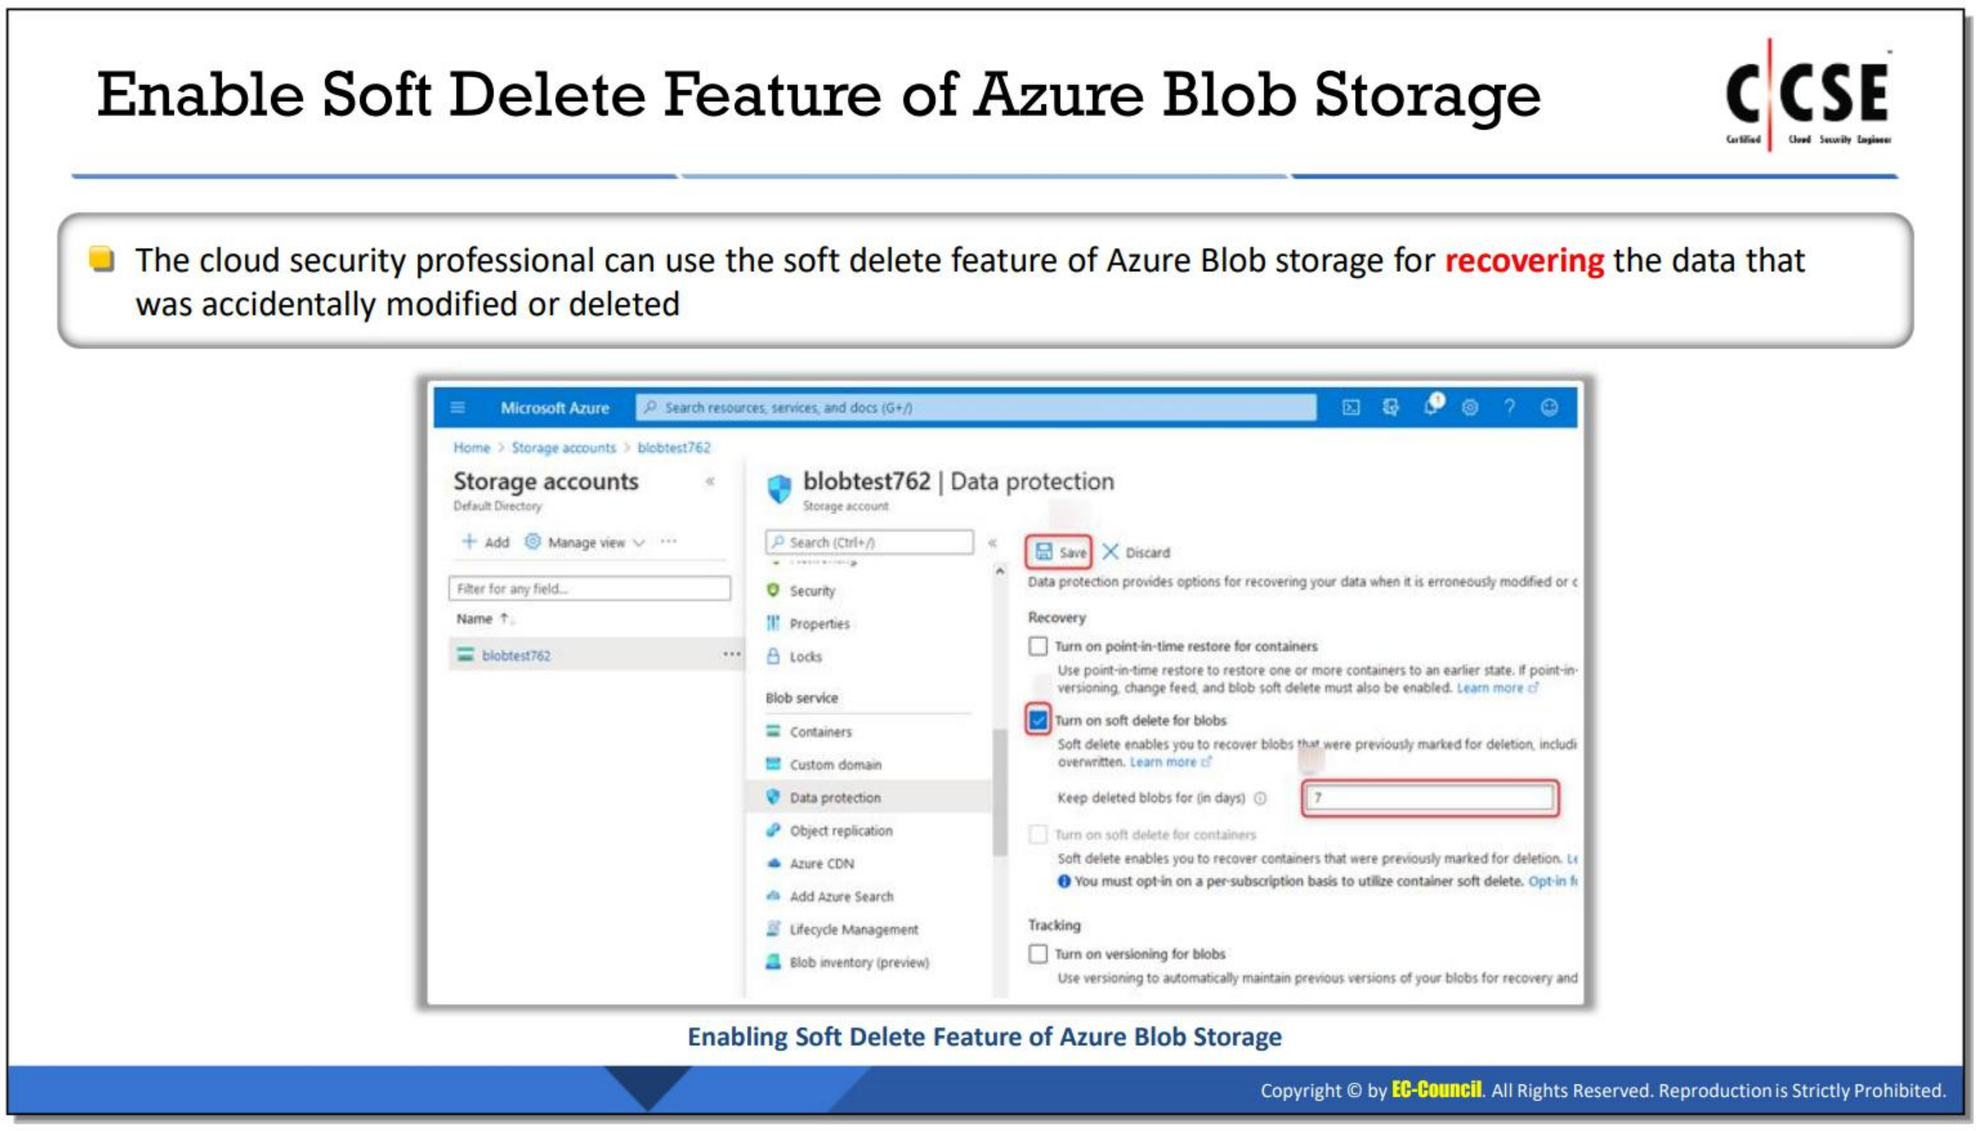


LO#14: Learn How to Implement Data Security in Microsoft Azure

This section explains the implementation of data security in Azure with various Azure features and services.

hide01.ir

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



Enabling Soft Delete Feature of Azure Blob Storage

Enable Soft Delete Feature of Azure Blob Storage

Usually, there are a lot of cases where data may be deleted or modified accidentally by storage account users or anyone who has access to it. Soft delete is one of the solutions for data protection that can help accidental changes or loss of data. It can retain data only for a specific number of days and the user can retrieve the deleted data within the stipulated time.

Soft delete ensures that Blob data can be recovered which was accidentally modified or deleted. Replication tiers like ARS, GRS, and RA-GRS will be locked by Azure Resource Manager. The cloud security professional can use the soft delete feature of Azure Blob Storage for recovering the data that was accidentally modified or deleted

Source: https://azure.microsoft.com

To Enable Soft Delete Feature of Azure Blob Storage

- Sign-in to the Azure portal.
- Under **Azure services**, navigate and select **Storage accounts**

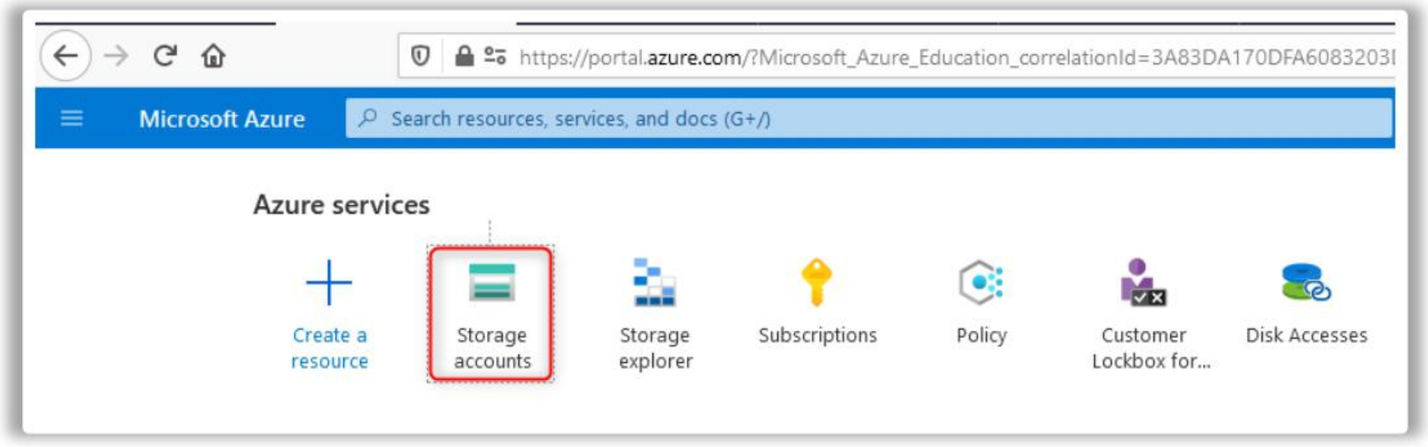


Figure 4.193: Microsoft Azure Console

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

- A list of storage accounts will be displayed.
- Select the storage for which you want to enable soft delete and click on it.

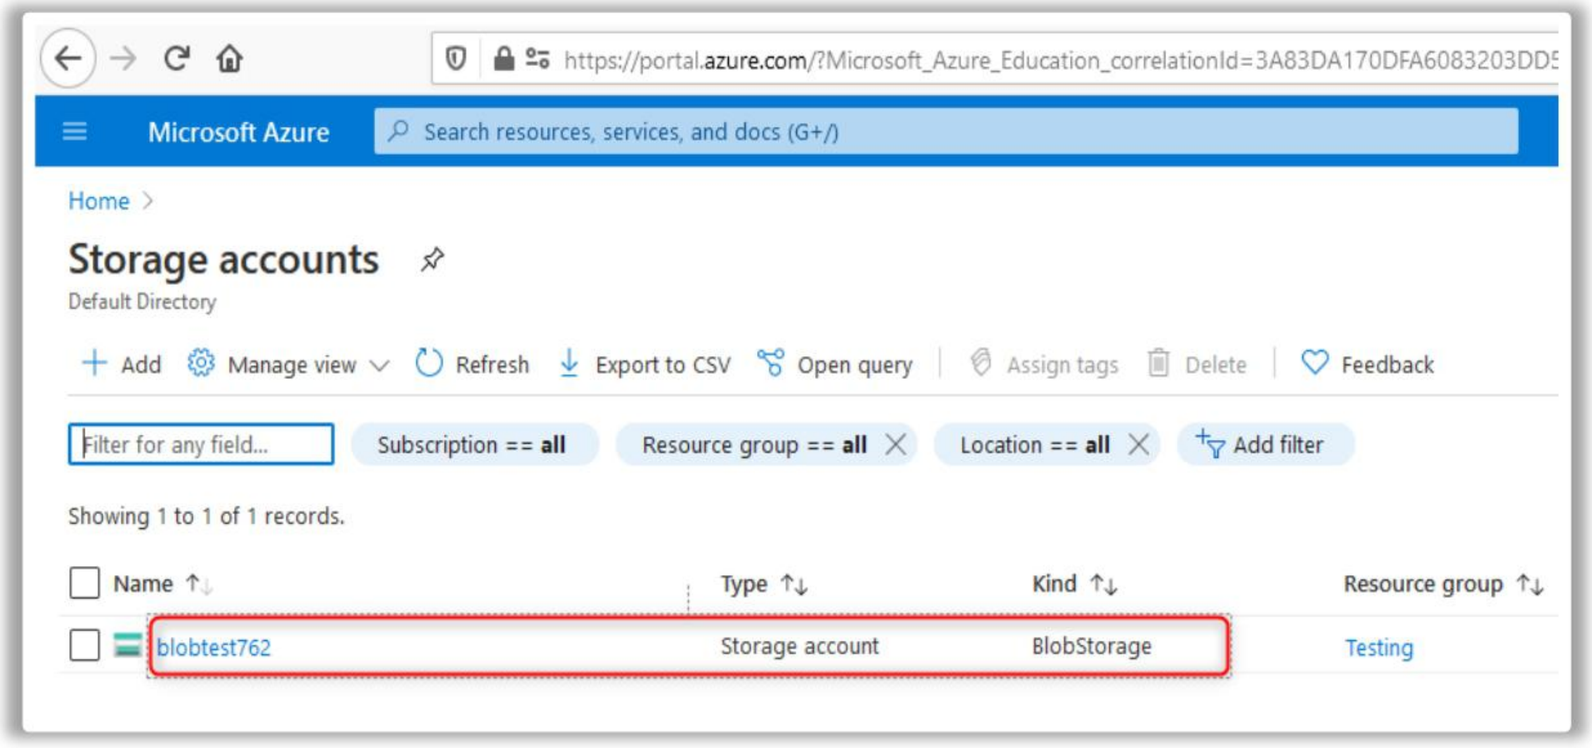


Figure 4.194: Choose Storage Account

- Navigate to **Data protection**, then click on it.

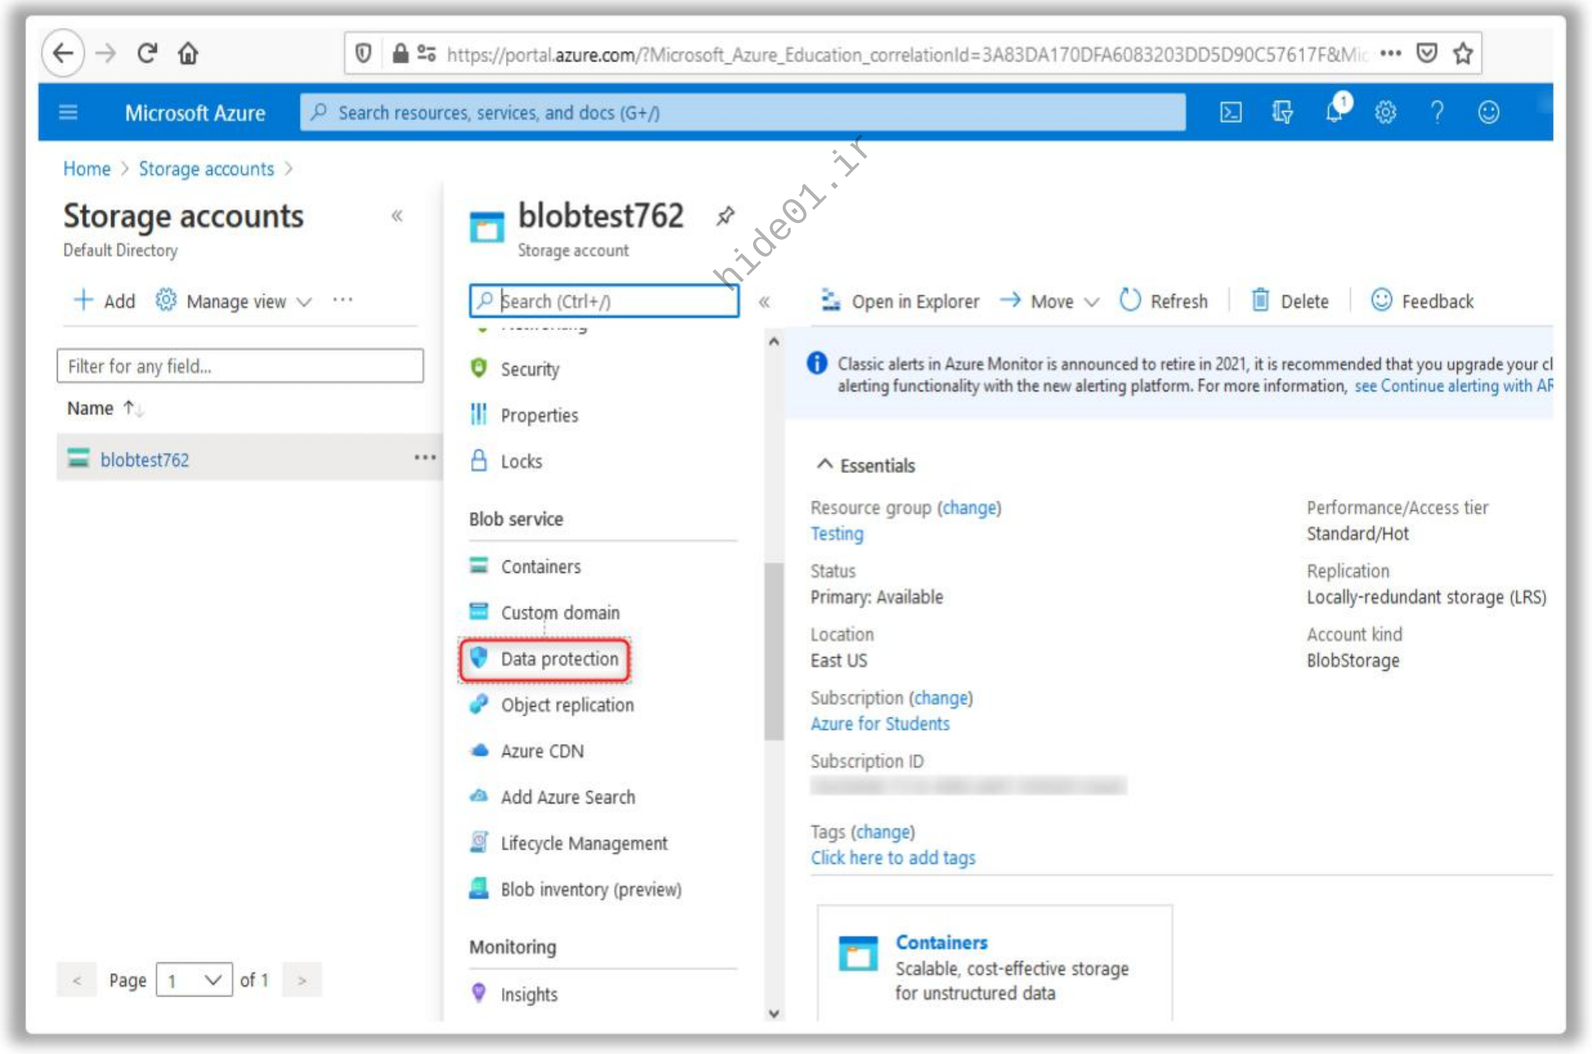


Figure 4.195: Select Data Protection

- Under Recovery, select the checkbox **Turn on soft delete for blobs**.
- Choose the required number of days in "**keep deleted blobs for (in days)**".

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

- Click on **Save** to apply changes.

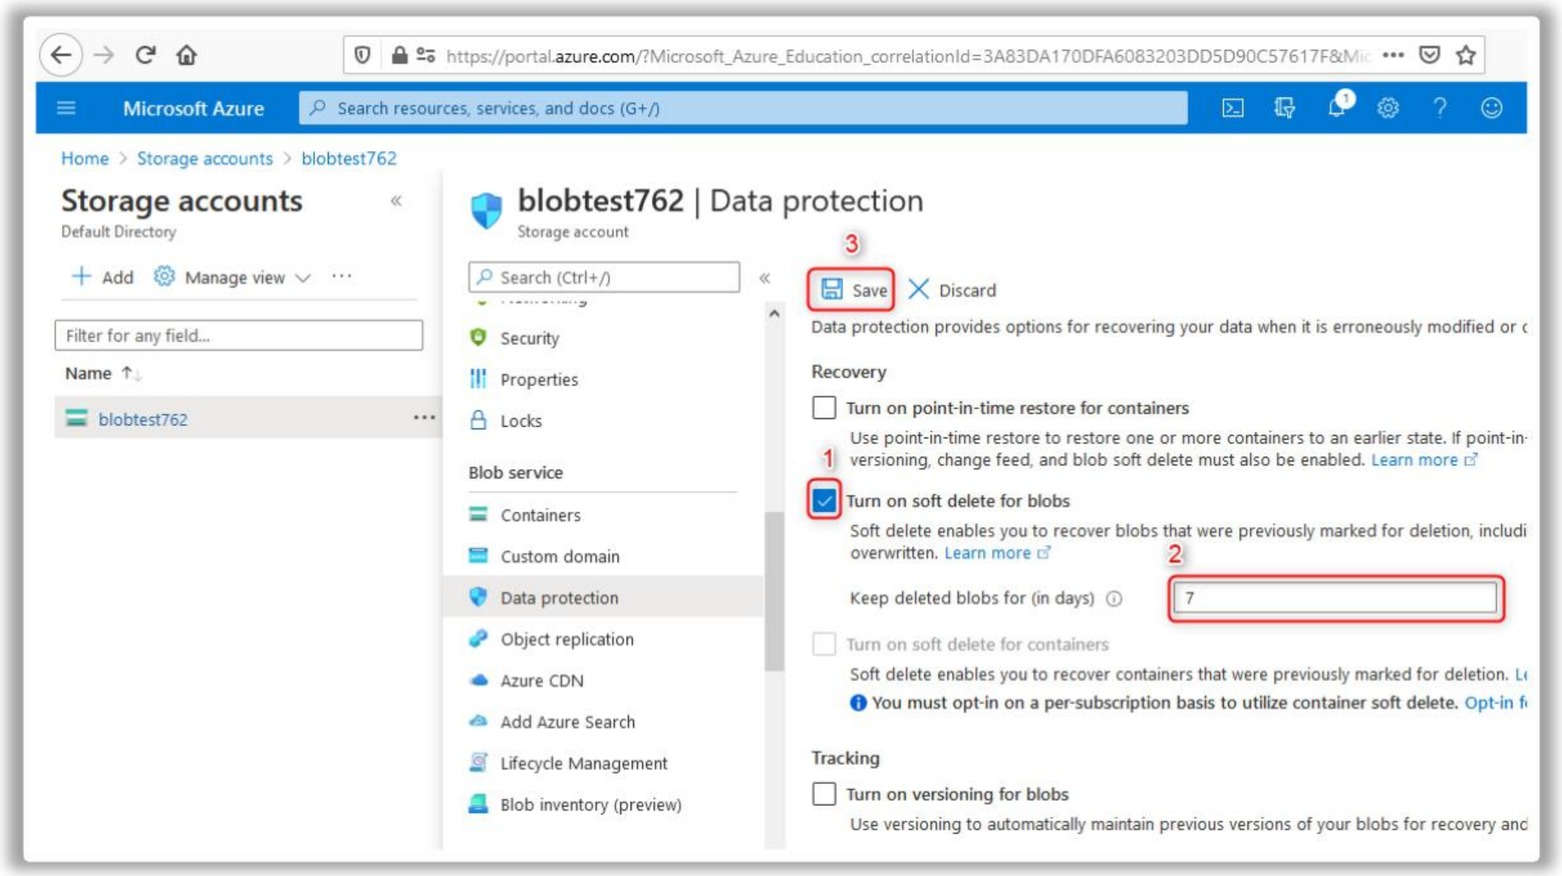


Figure 4.196: Click Save

- Soft delete for Blob protection has been successfully updated.

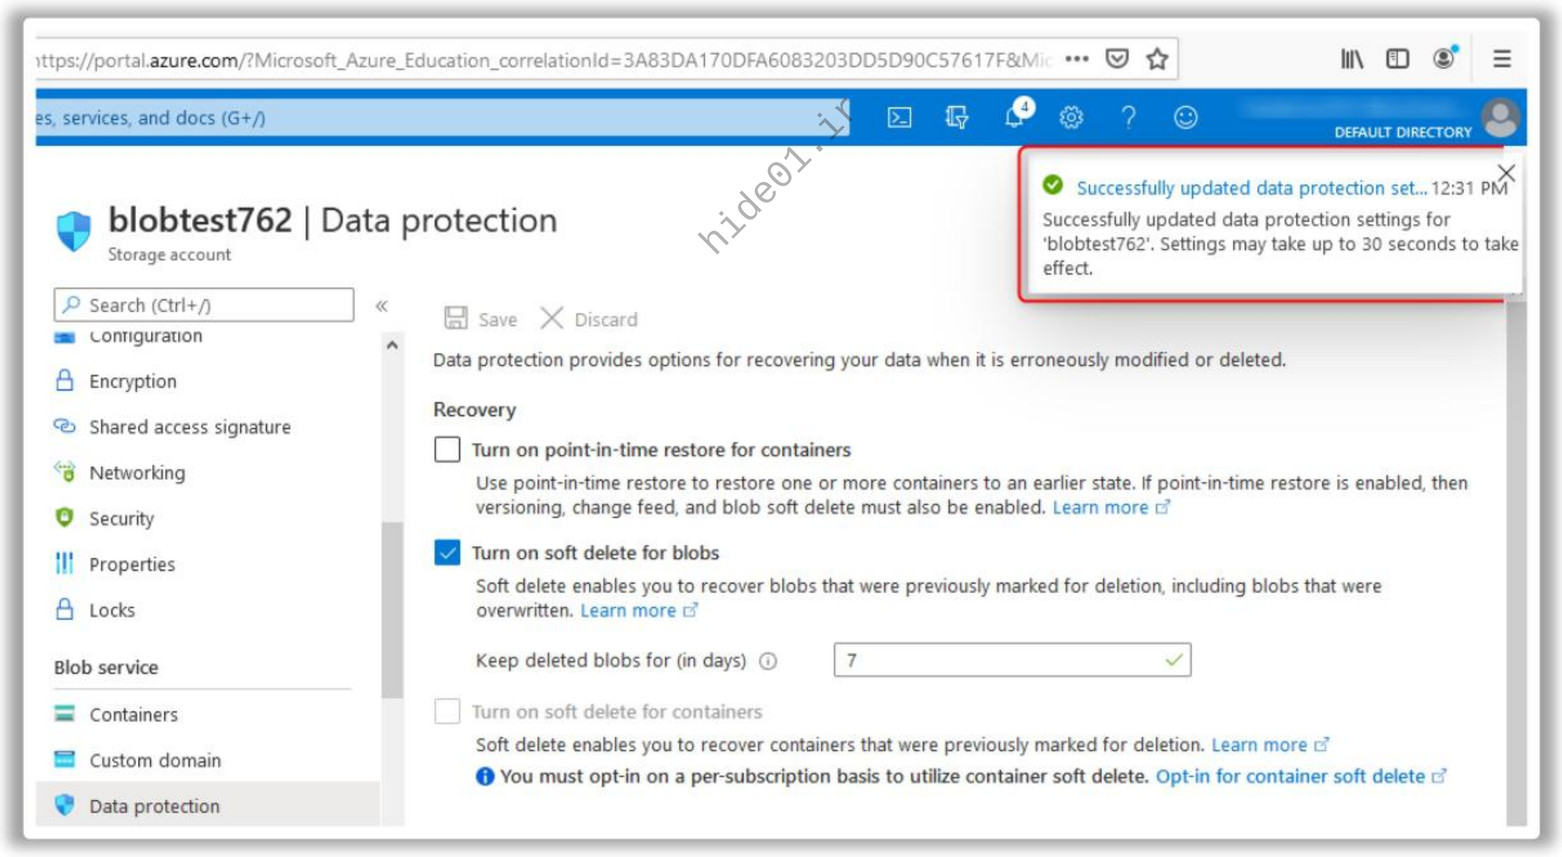
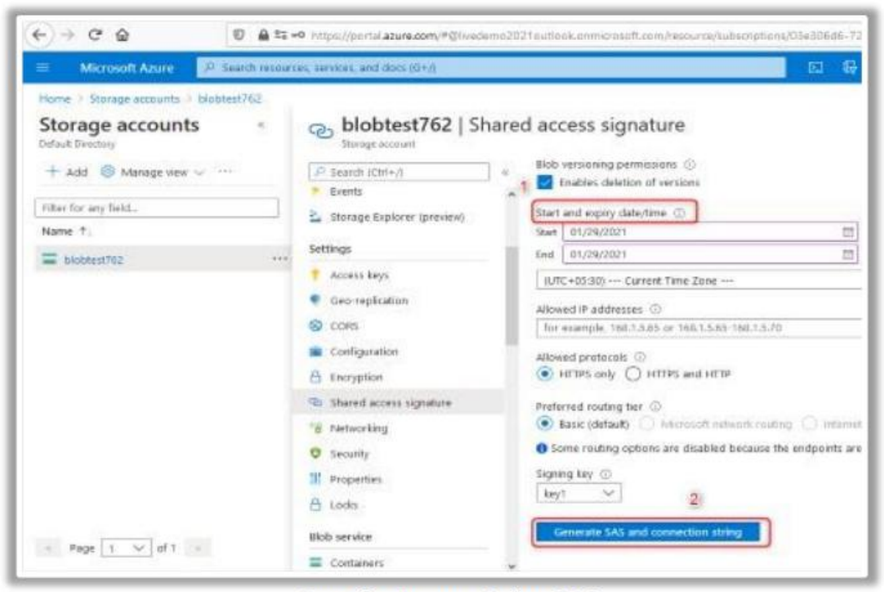


Figure 4.197: Blob Created Successfully

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

Restrict Access Using SAS

- Shared Access Signature (SAS) helps in providing finegrained **access control** to Azure storage account
- The cloud security administrator can restrict access to the Azure storage account by using **SAS tokens** that have specific query parameters



Restrict Access Using SAS

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

Restrict Access Using SAS

SAS helps in providing fine-grained access control to Azure storage account. It restricts access to the Azure storage account with tokens having specific query parameters. The cloud security administrator can configure the SAS token to expire after particular instant of time or limit the access to particular resources.

A SAS token gives you control of data that can be shared with the client; for example:

- What resources that a client can access.
- What permissions they can have for the allowed resources like read-access, write, and delete access, etc.
- Till what duration SAS is valid.

Azure storage supports three different types of shared access signatures:

- **User delegation SAS:** This is secured with the credentials of Azure Active Directory along with specified permissions of SAS. It can be applied to only Blob Storage.
- **Service SAS:** It is secured with a storage account key. It has limitations and can be used only in Queue, Table, Blob, or Azure storages.
- **Account SAS:** It is secured with a storage account key. It can provide access to resources in different storage services.

  To configure Azure storage access, it is recommended to use SAS rather than storage access keys that gives complete access permanently to the whole storage account.

Source: https://azure.microsoft.com

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

To Restrict Access to Azure Storage Account Using Shared Access Signature (SAS)

- Sign-in to the Azure portal.
- Select **Storage accounts** under **Azure services.**

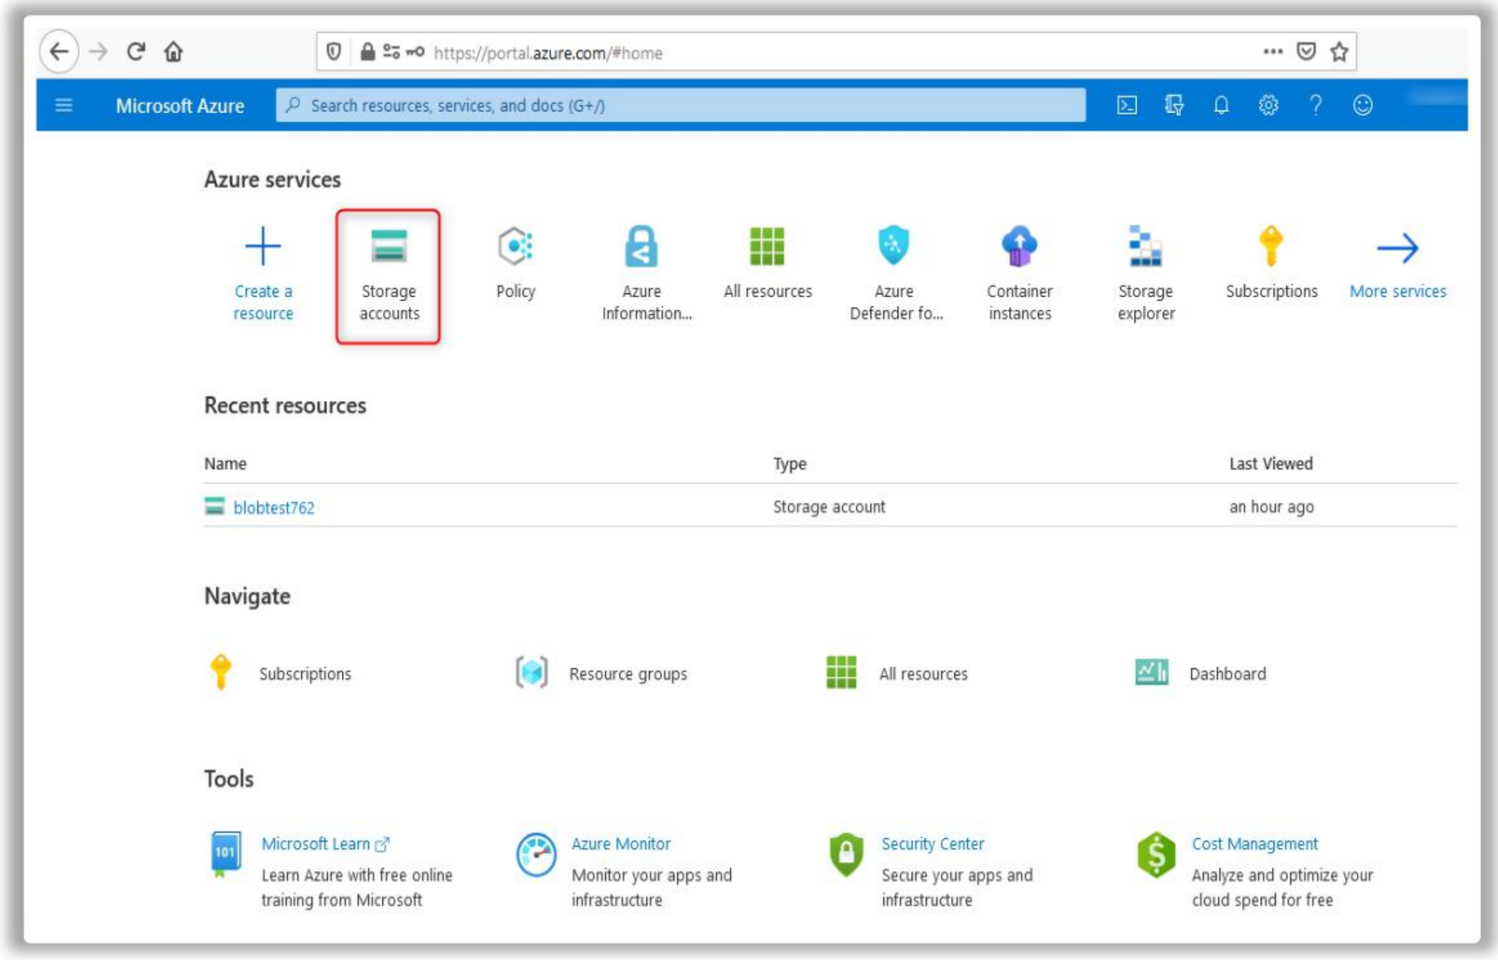


Figure 4.198: Select Storage Account

- A list of storage accounts will be displayed.
- Select the storage to generate a SAS.

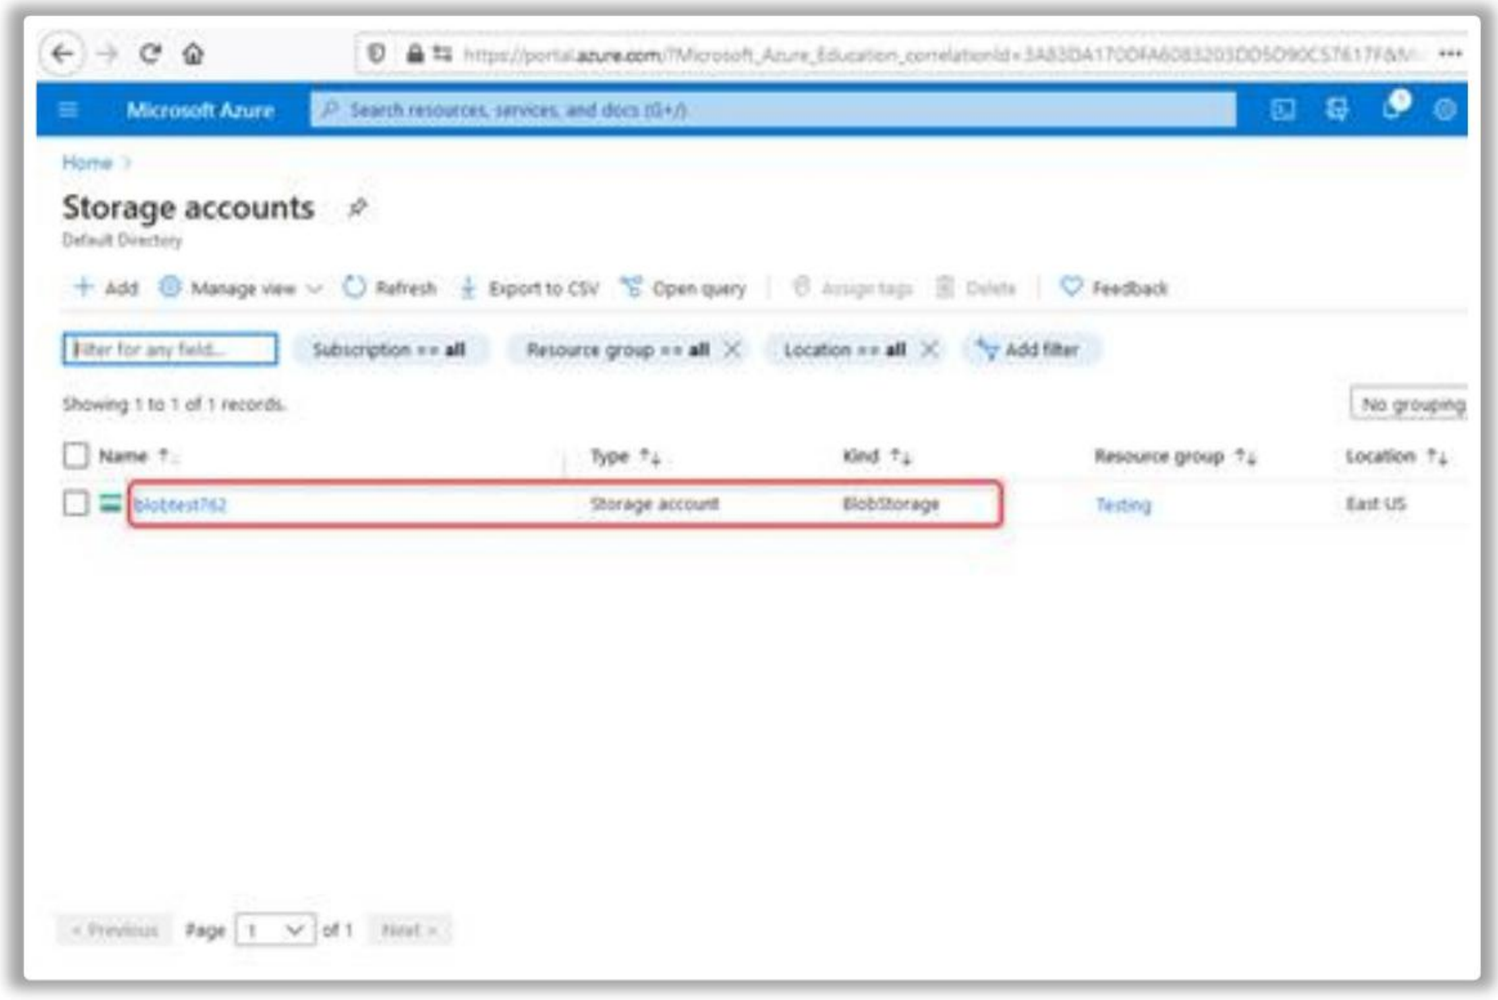


Figure 4.199: Generate SAS

- Under **Settings**, navigate to **Shared access signature** and click on it

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

- Select the **Allowed services** that you want to allow access to.
- Select the **Shared resources types** as per your requirement.
- Select the necessary **Allowed permission** that clients can perform.

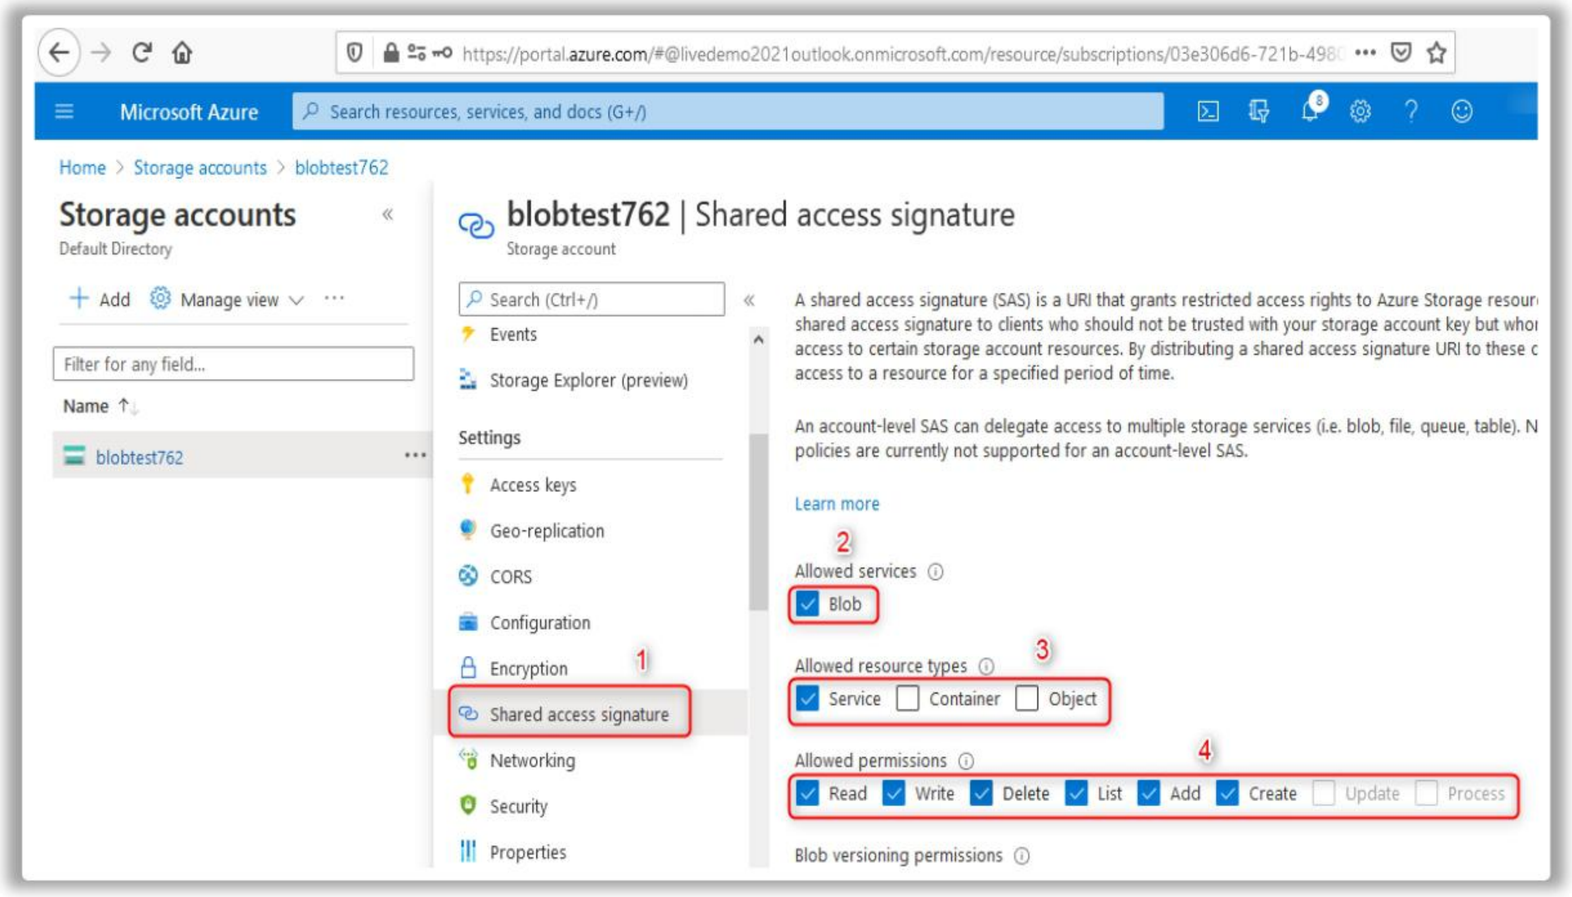


Figure 4.200: Allow Permissions

- Choose the expiry using a calendar or input fields in **Start and expiry date/time**
- Click **Generate SAS and connection string** to generate tokens.

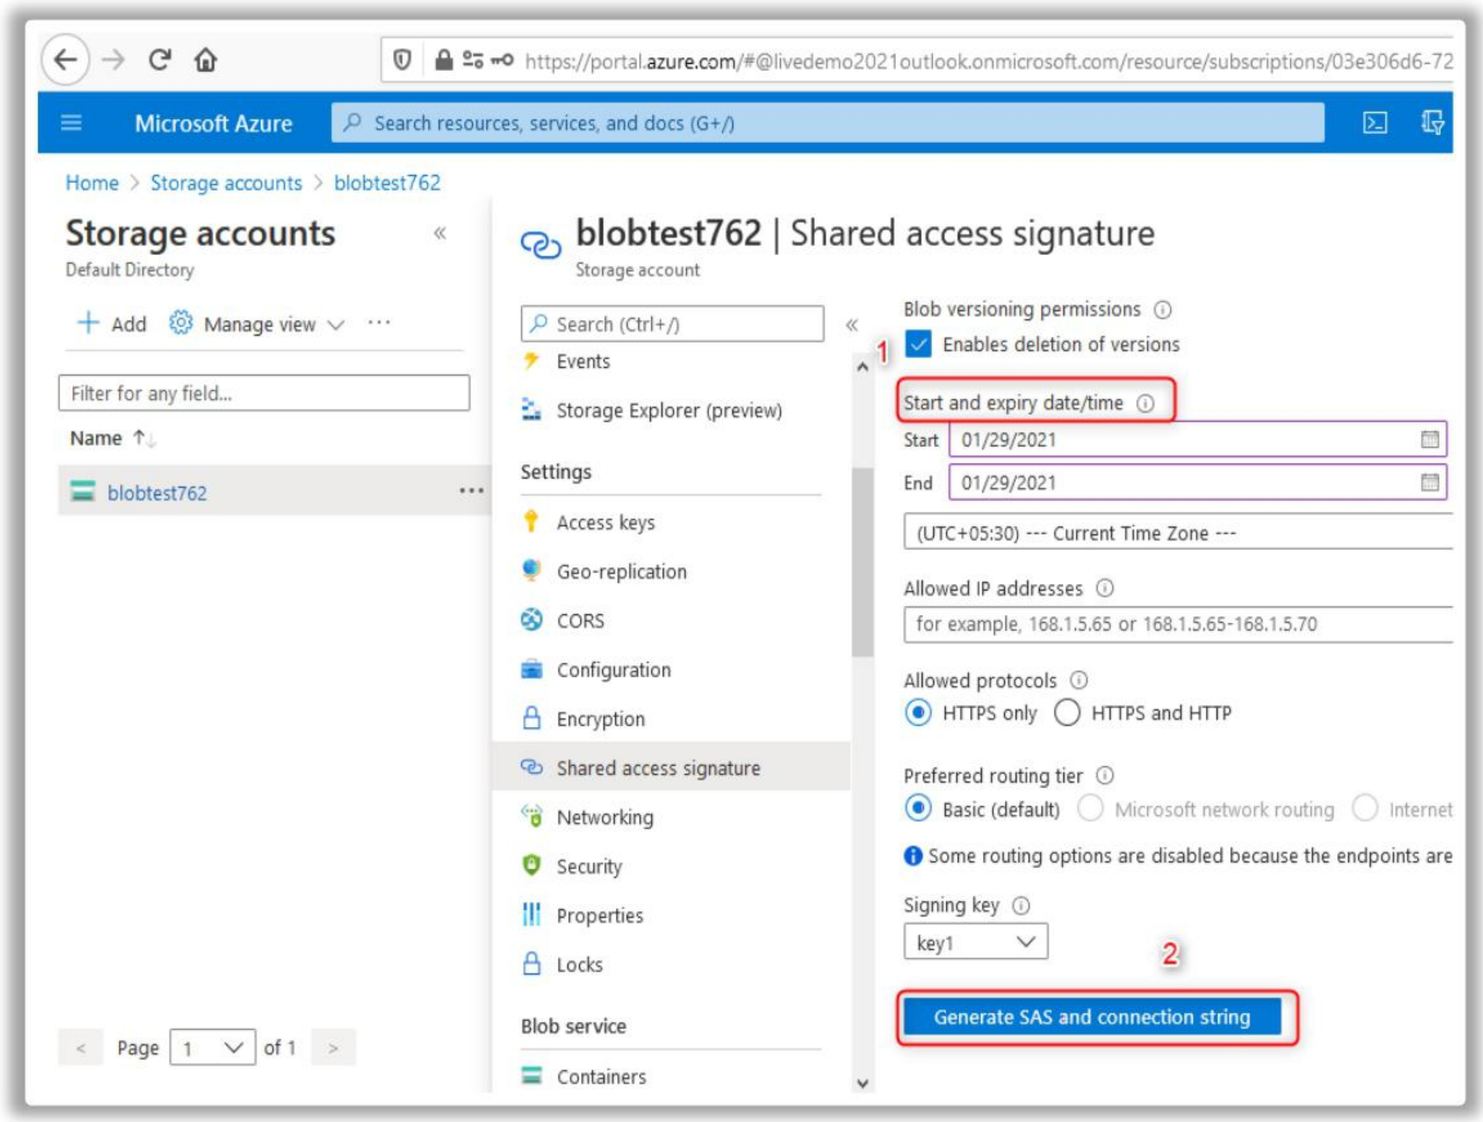


Figure 4.201: Click Generate SAS and connection strings

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

- Connection string, SAS token, Blob service SAS URL are successfully generated. The cloud security administrator can use the SAS token to restrict access to Azure storage account.

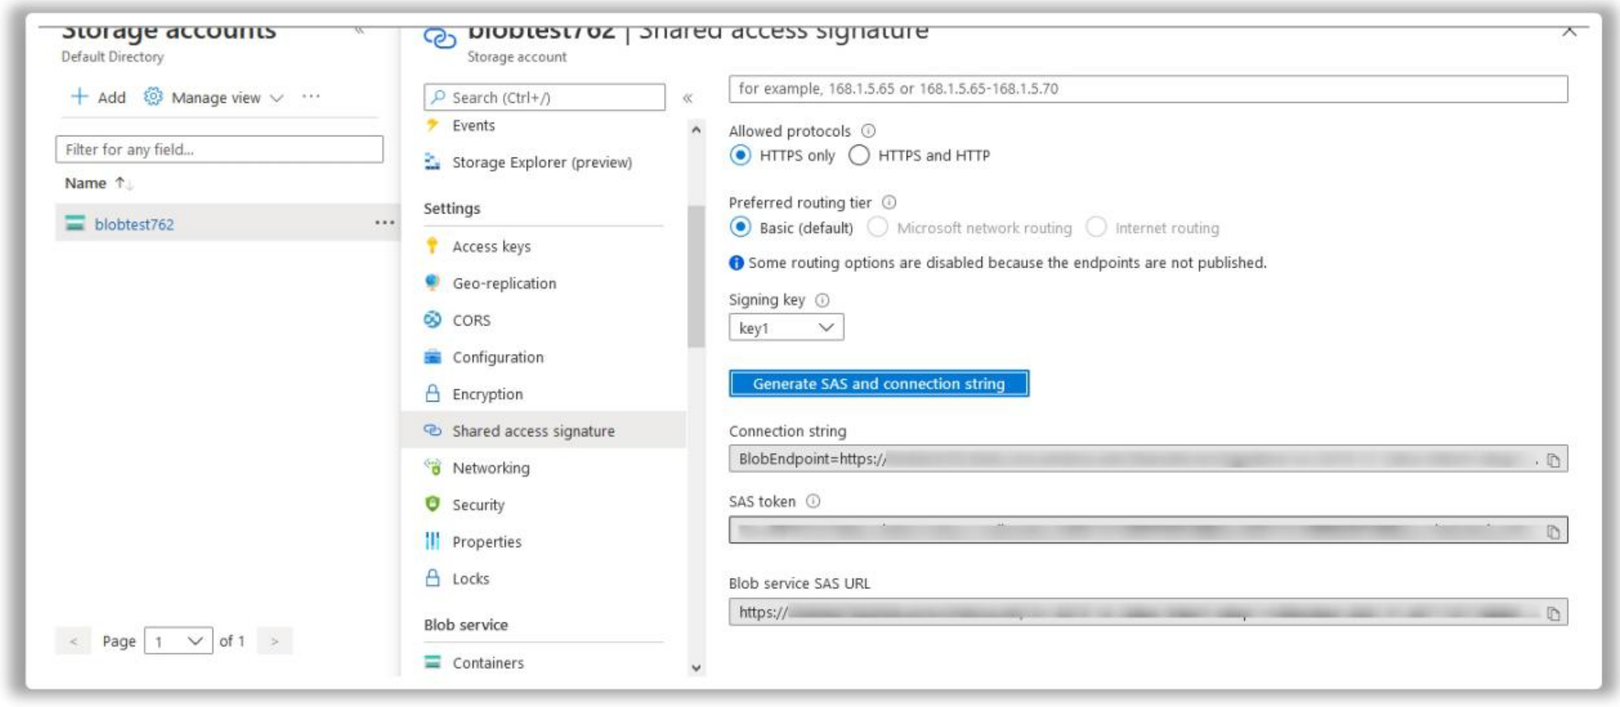
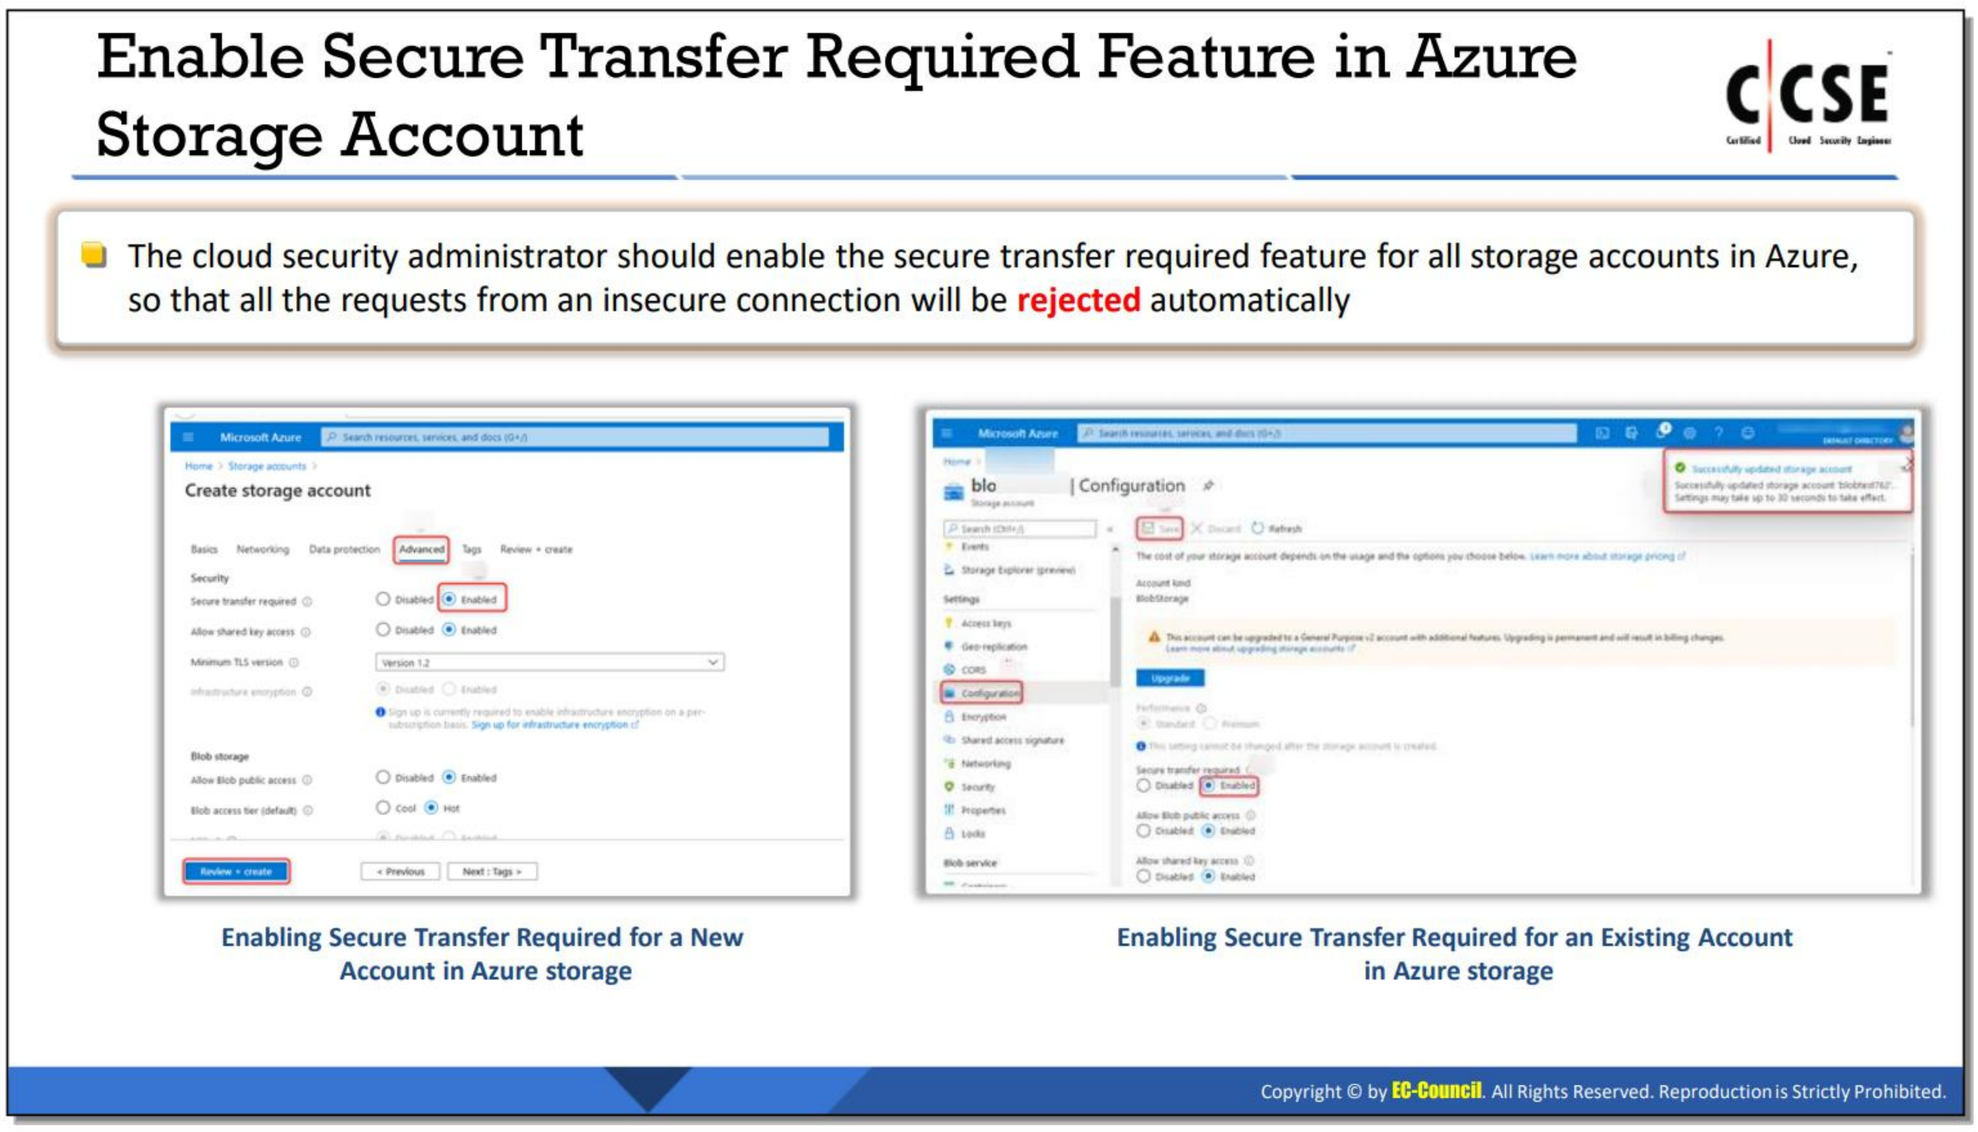


Figure 4.202: Security Administration can use SAS

hide01.ir

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



Enable Secure Transfer Required Feature in Azure Storage Account

To make the data transfer secure, Azure provides a "Secure Transfer Required" facility. An Azure account administrator can configure this while creating a new storage account or for an existing storage account. If it is enabled, all the requests from an insecure connection will be rejected automatically. It is highly recommended to use this feature for all storage accounts.

Any call to Azure Storage REST API actions should be over HTTPS only when a secure transfer is required. All other HTTP requests will be rejected. Any connection established without encryption over SMB like SMB 2.1, SMB 3.0 will be rejected.

Source: https://azure.microsoft.com

To enable Secure Transfer Required for a new account

- Sign-in to the Azure portal.
- Select **Storage accounts** under **Azure services.**

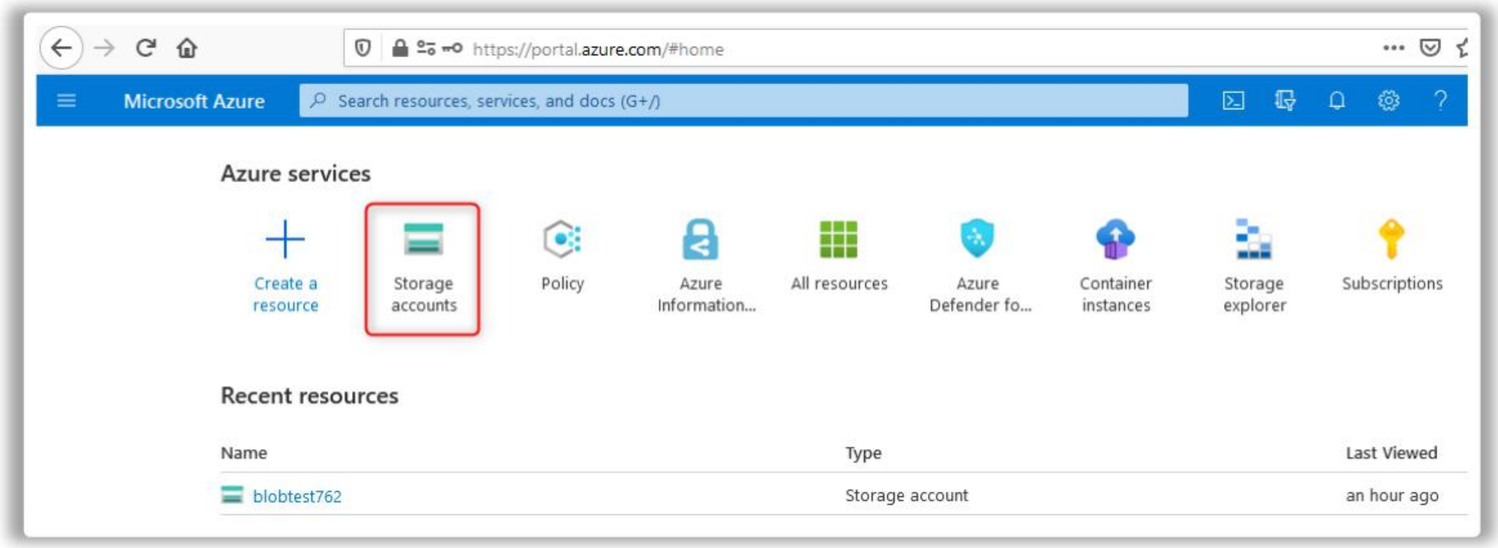


Figure 4.203: Select Storage Accounts

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

- Select **Add** option to create a new storage account

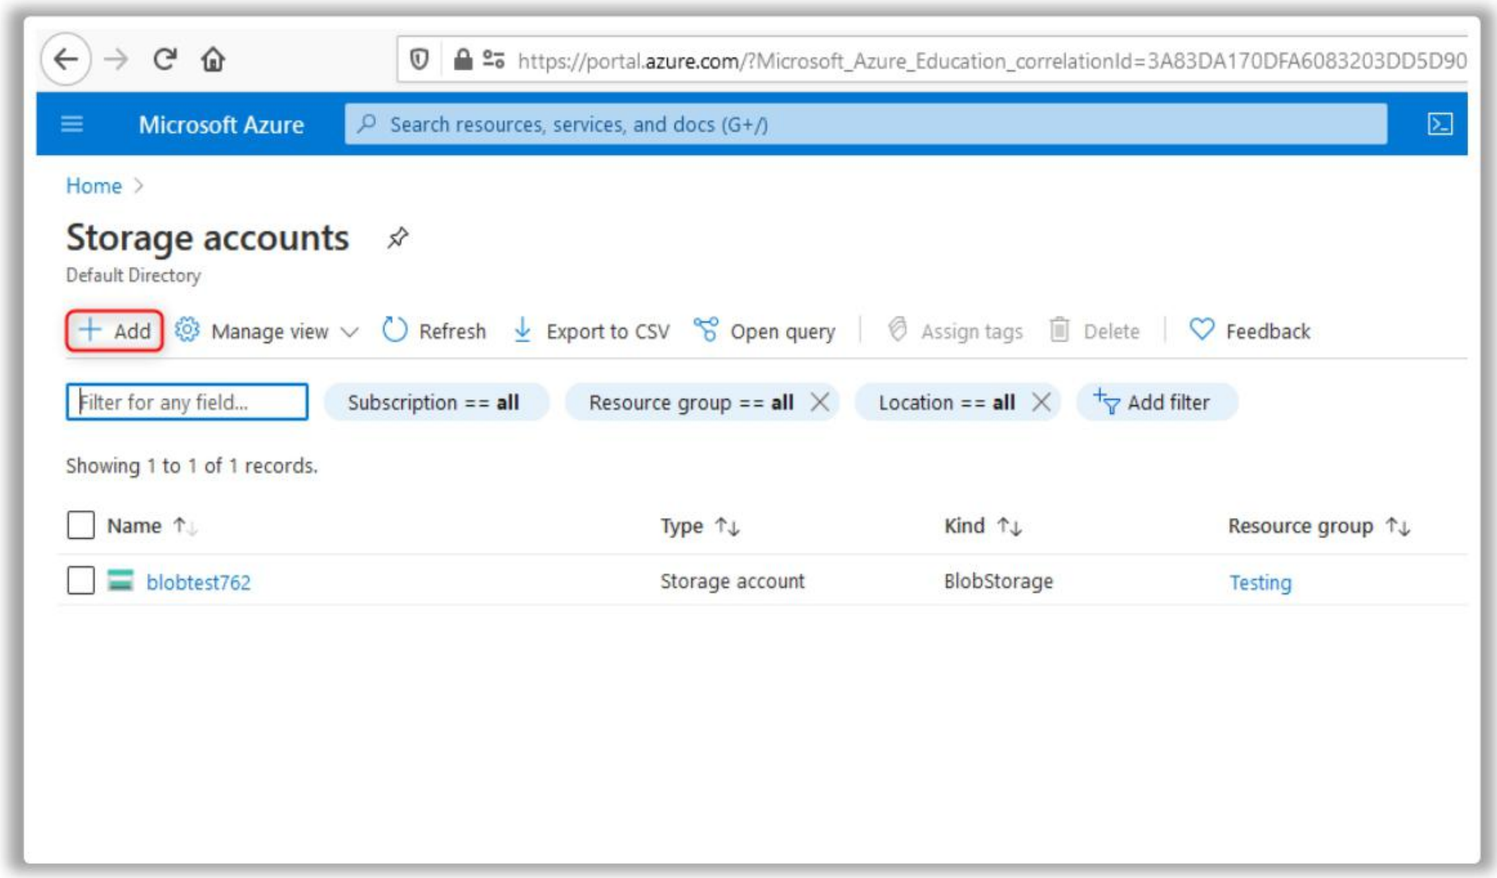


Figure 4.204: Click +Add

- Navigate to the **Advanced** tab.
- Select the **Enabled** checkbox of **Secure transfer required.**
- Click on the **Review + create** button to create a storage account with the secure transfer required feature enabled.

hide01.ir

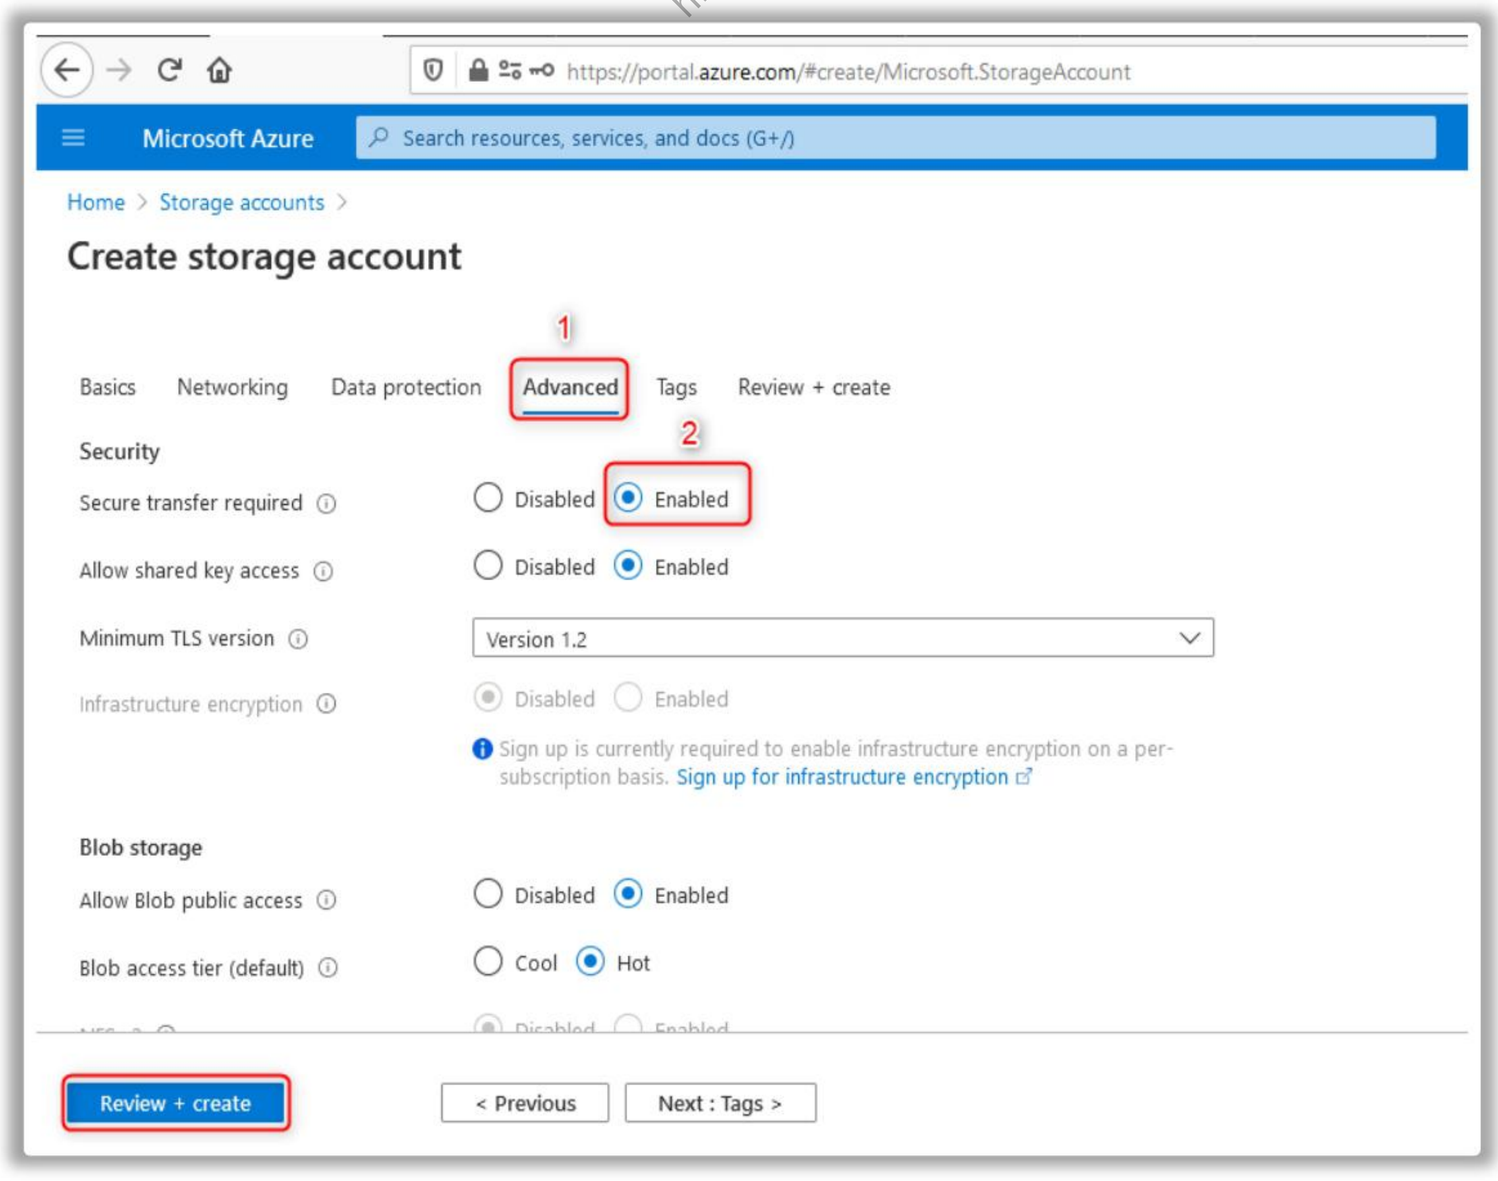


Figure 4.205: Select Review +create

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

Steps to enable secure transfer required for an existing account:

- Sign-in to the Azure portal.
- Select **Storage accounts** under **Azure services.**

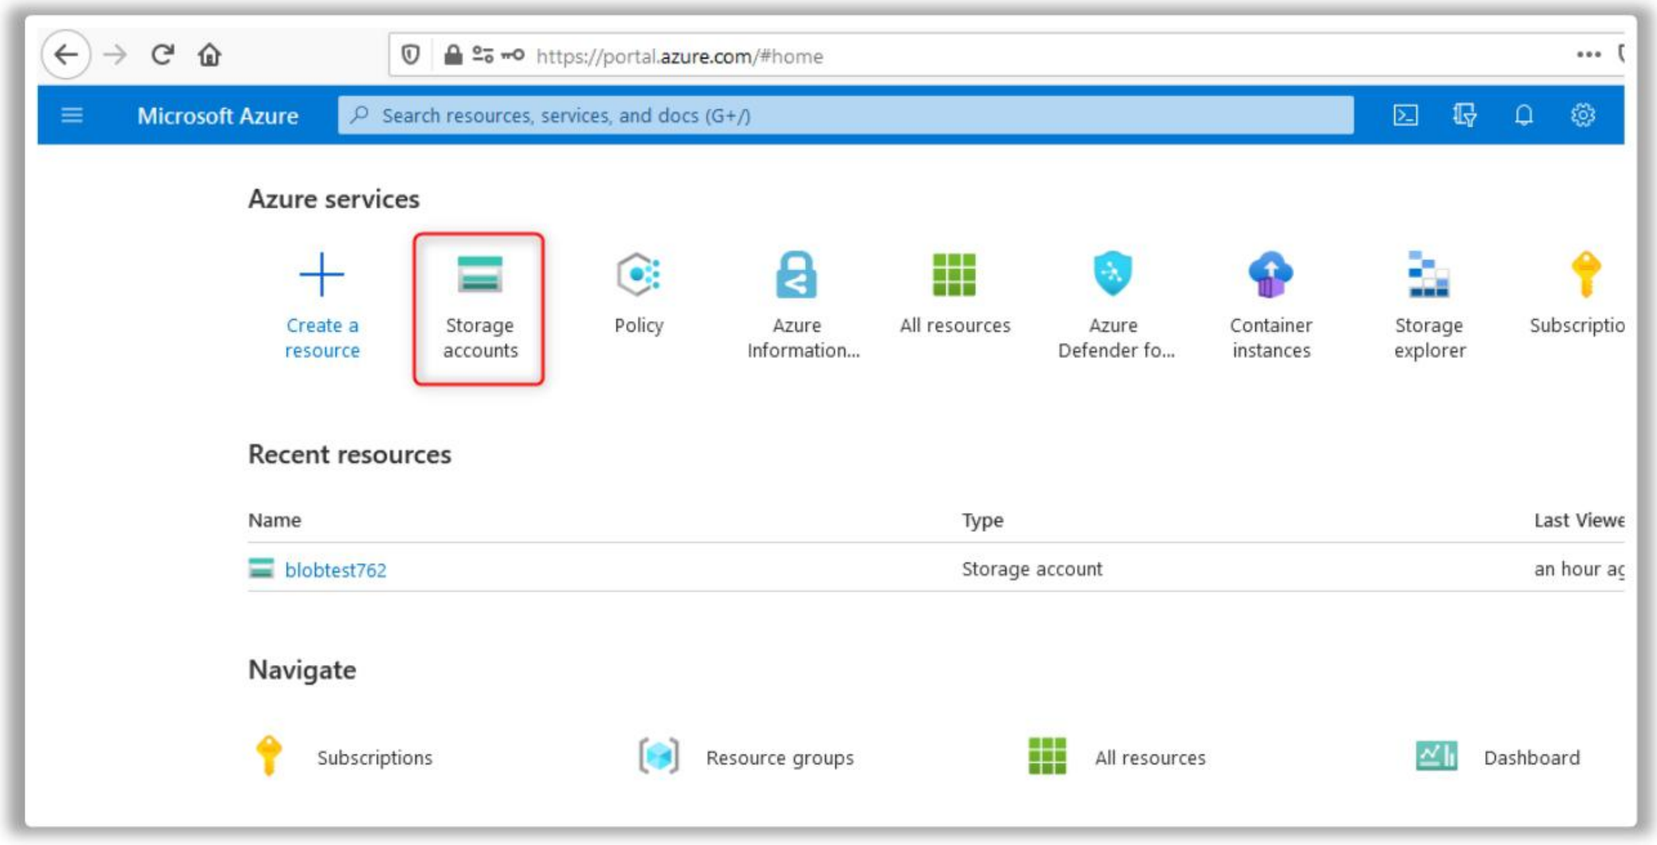


Figure 4.206: Click Storage Accounts

- A list of storage accounts will be displayed.
- Select the storage for which you want to enable Secure Transfer Required.

hidedl.ir

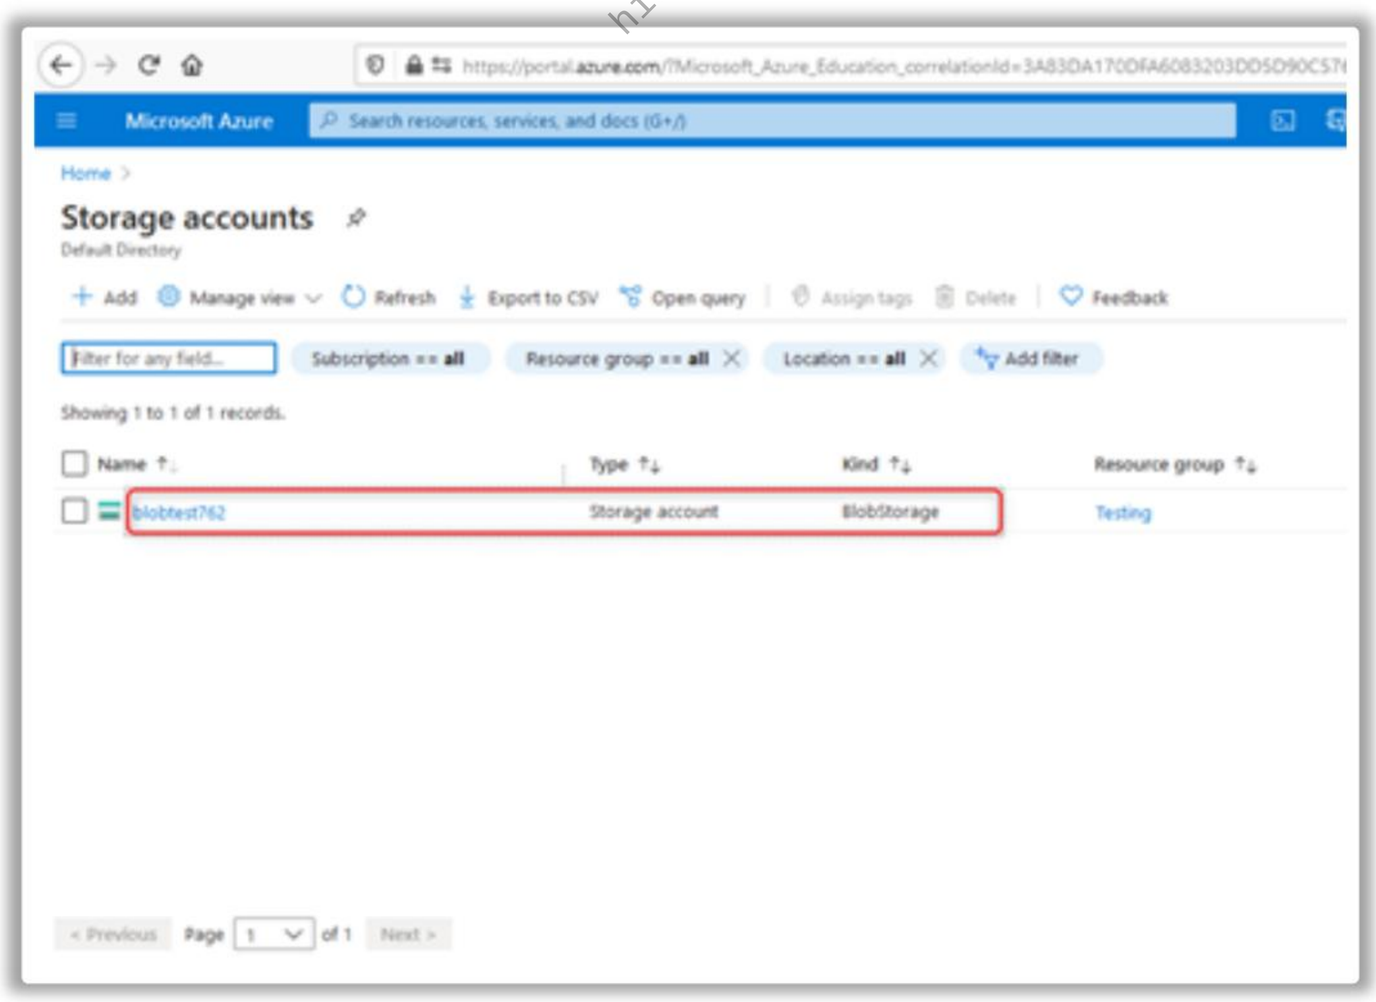


Figure 4.207: Storage Accounts Console

- Locate **Configuration** under **settings**.
- Select the **Enabled** checkbox of **Secure transfer required.**

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

- Click on the **Save** option.
- The secure transfer required has been successfully enabled.

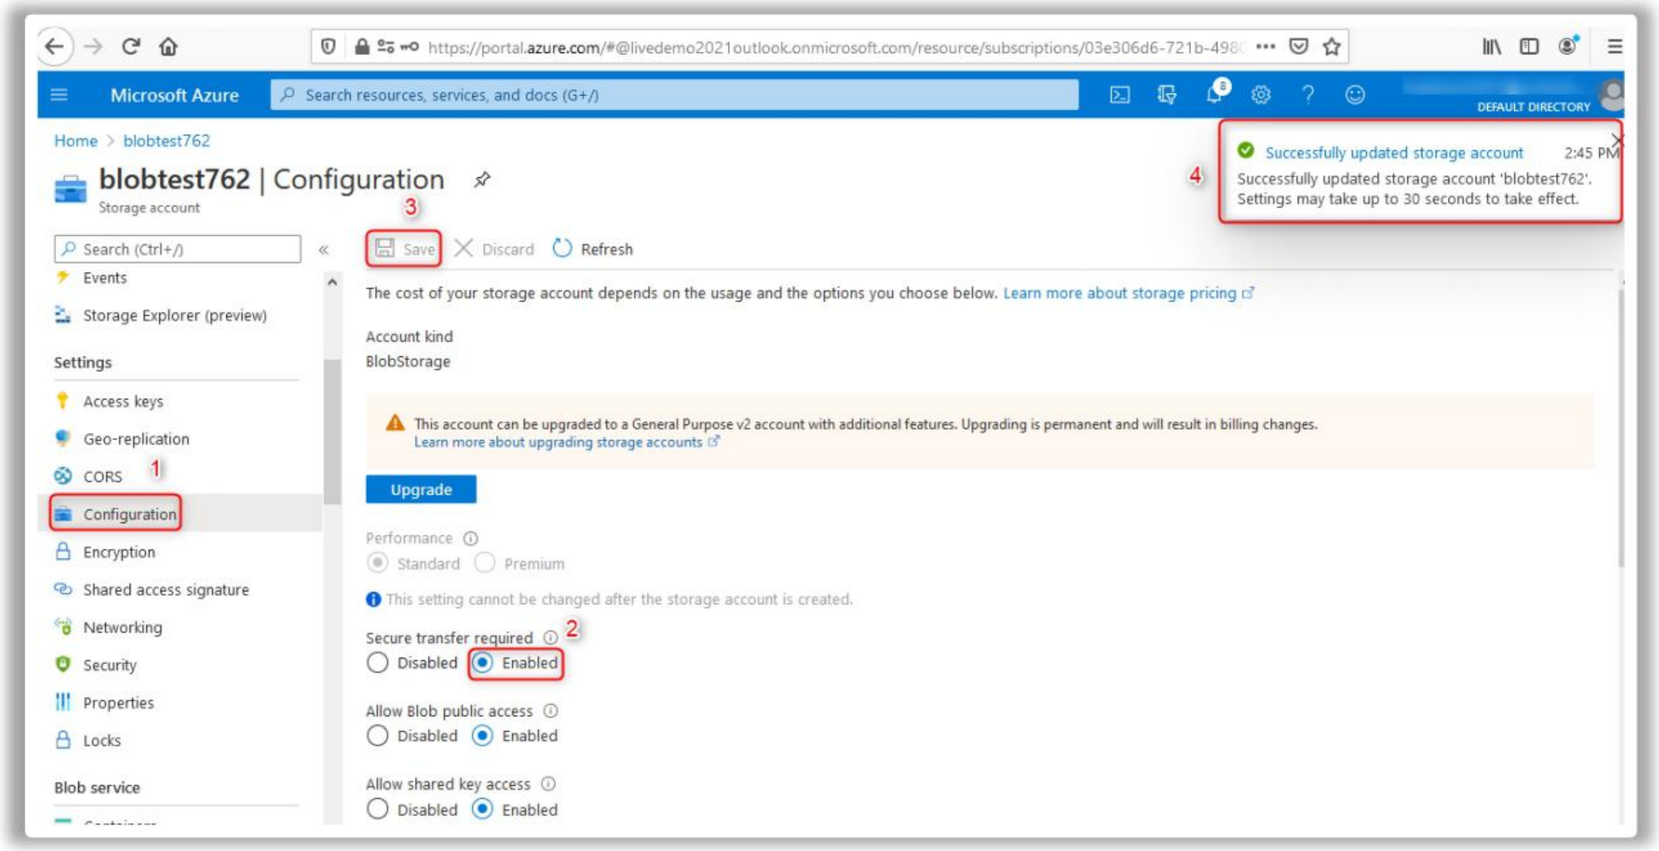


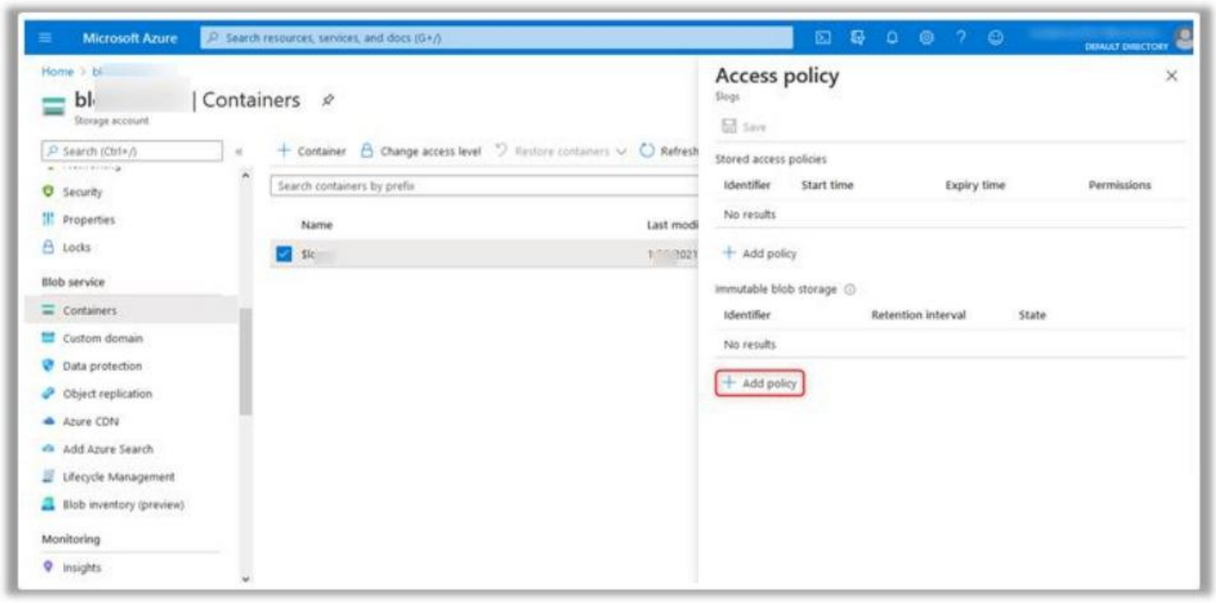
Figure 4.208: Secure Transfer is Enabled

hide01.ir

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

Configure Immutable Policies

- The cloud security administrator should configure the immutable policies to **safeguard** business-critical data that needs to be stored in WORM state (write once and read many)



Configuring Immutable Policies

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

Configure Immutable Policies

Business-critical data that needs to be stored in WORM state (write once and read many) can be configured in immutable policies and legal holds. This is available for both general-purpose v2 as well as Blob storage accounts but not other storage types. These blobs once created cannot be modified and can be read as many times as required, it has some specific time called retention period.

Source: https://azure.microsoft.com

To Configure Immutable Polices

- Sign-in to the Azure portal.
- Select Storage accounts under Azure services.

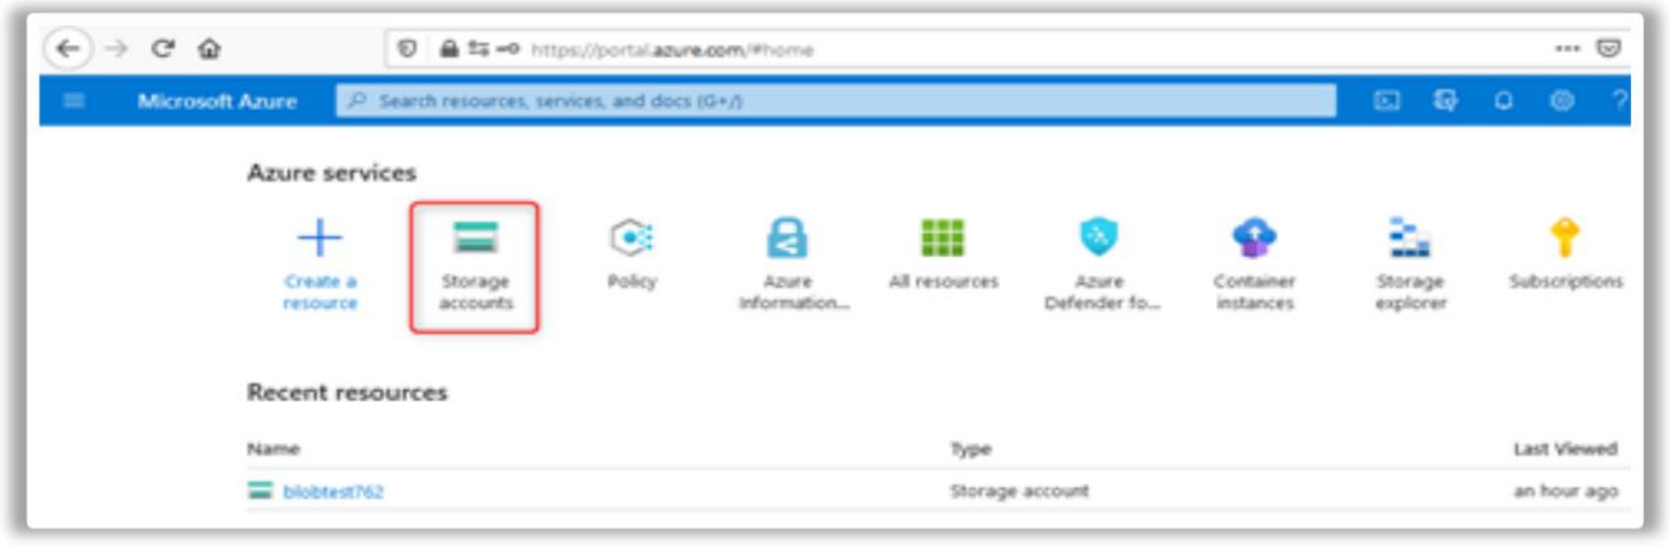


Figure 4.209: Select Storage Accounts

- A list of storage accounts will be displayed.

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

- Select the storage for which you want to configure policies.

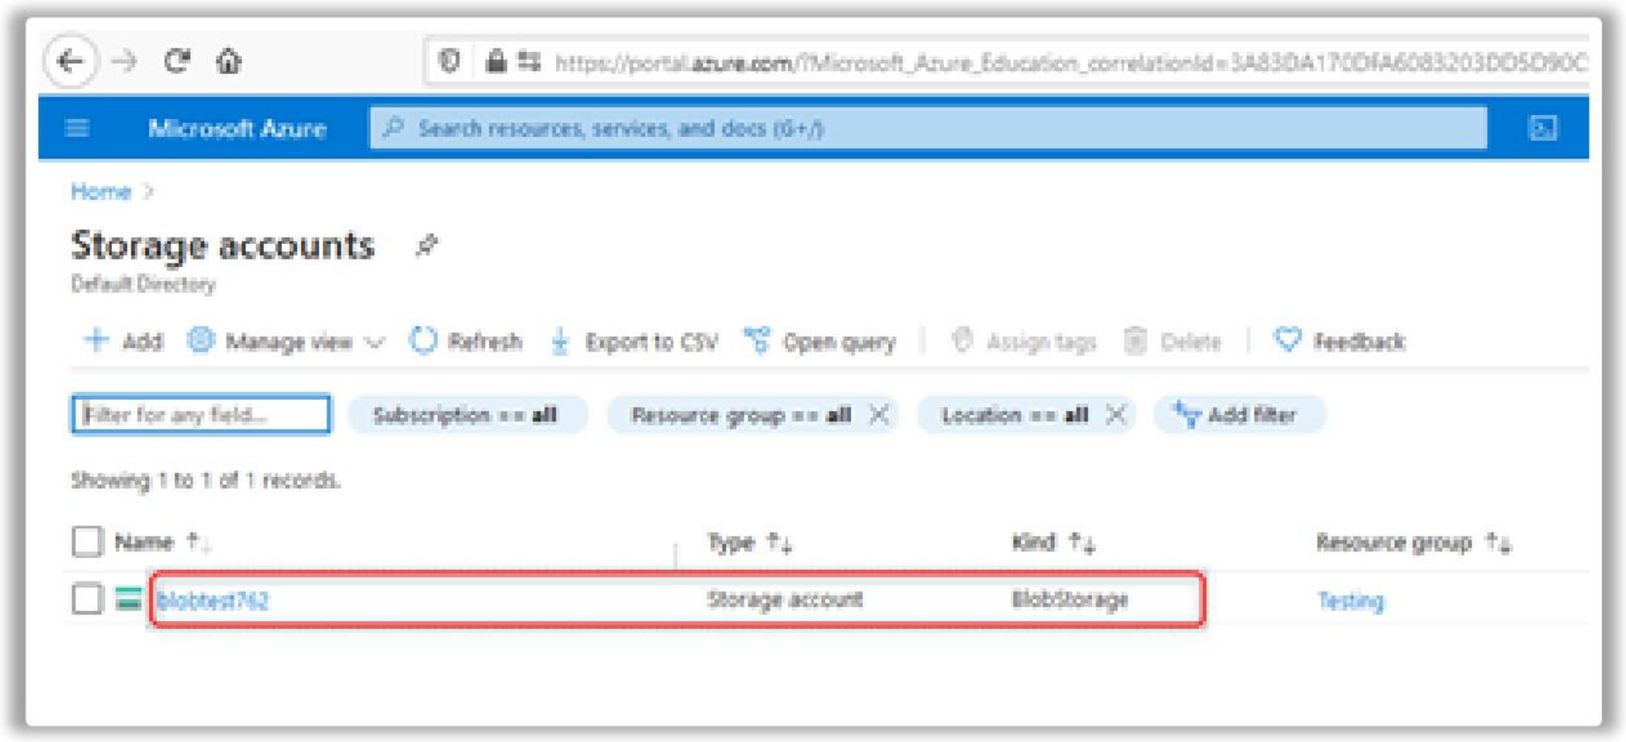


Figure 4.210: Select Storage Accounts

- Locate to Blob service and select Containers.
- Click on the three dots icon and go to the Access policy.

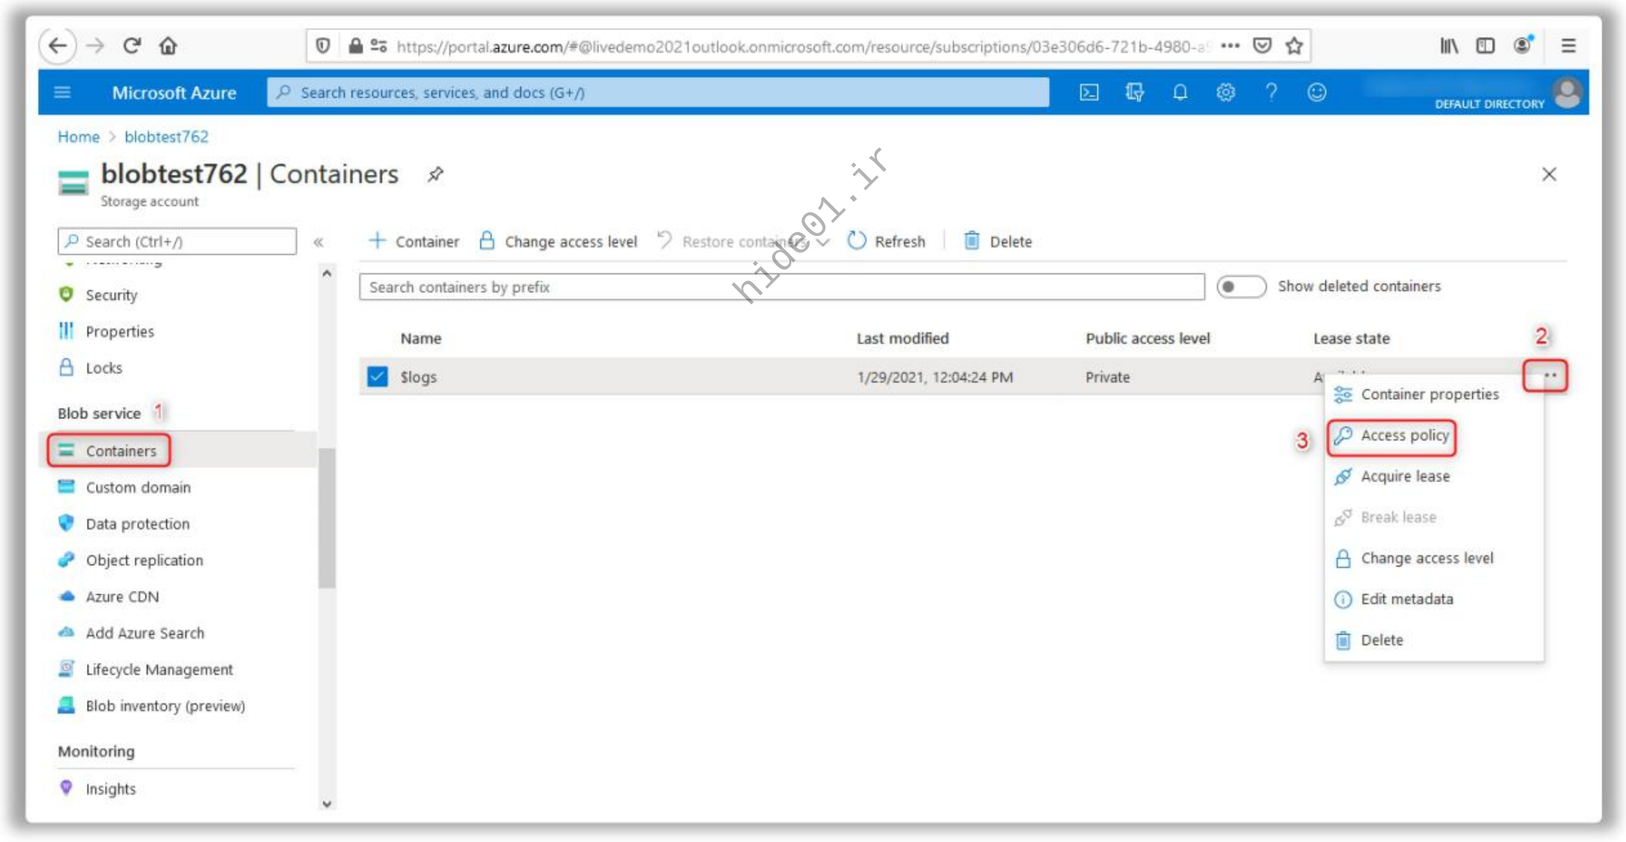


Figure 4.211: Select Access Policy

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

- Select Add policy under Immutable blob storage.

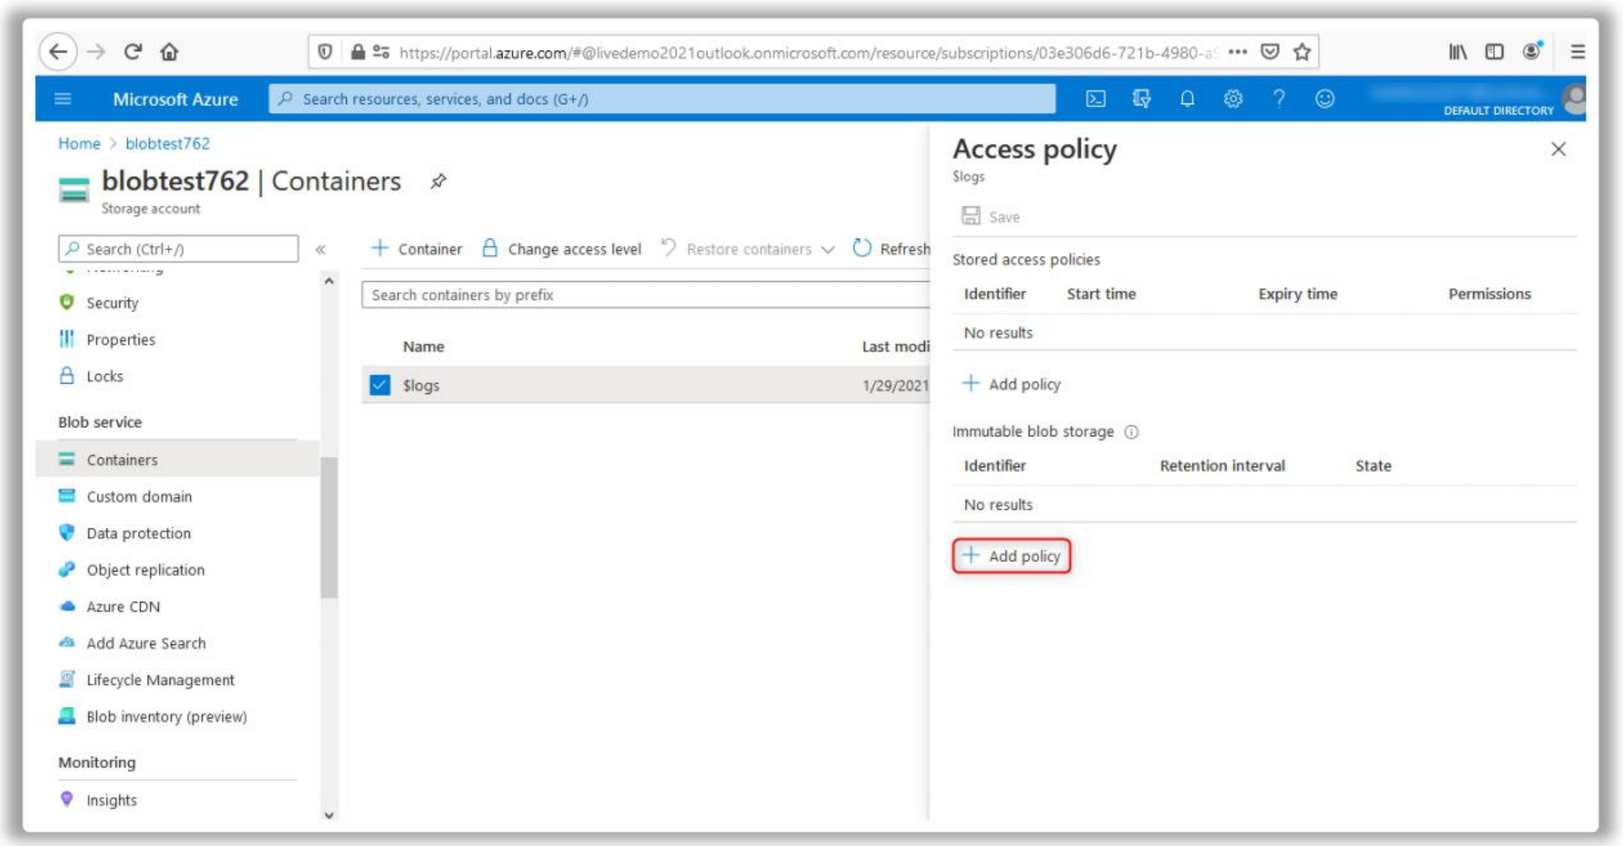


Figure 4.213: Click Add Policy

- Select Time-based retention from the dropdown.
- Set the retention period in days of your requirement.
- Click on **OK**.

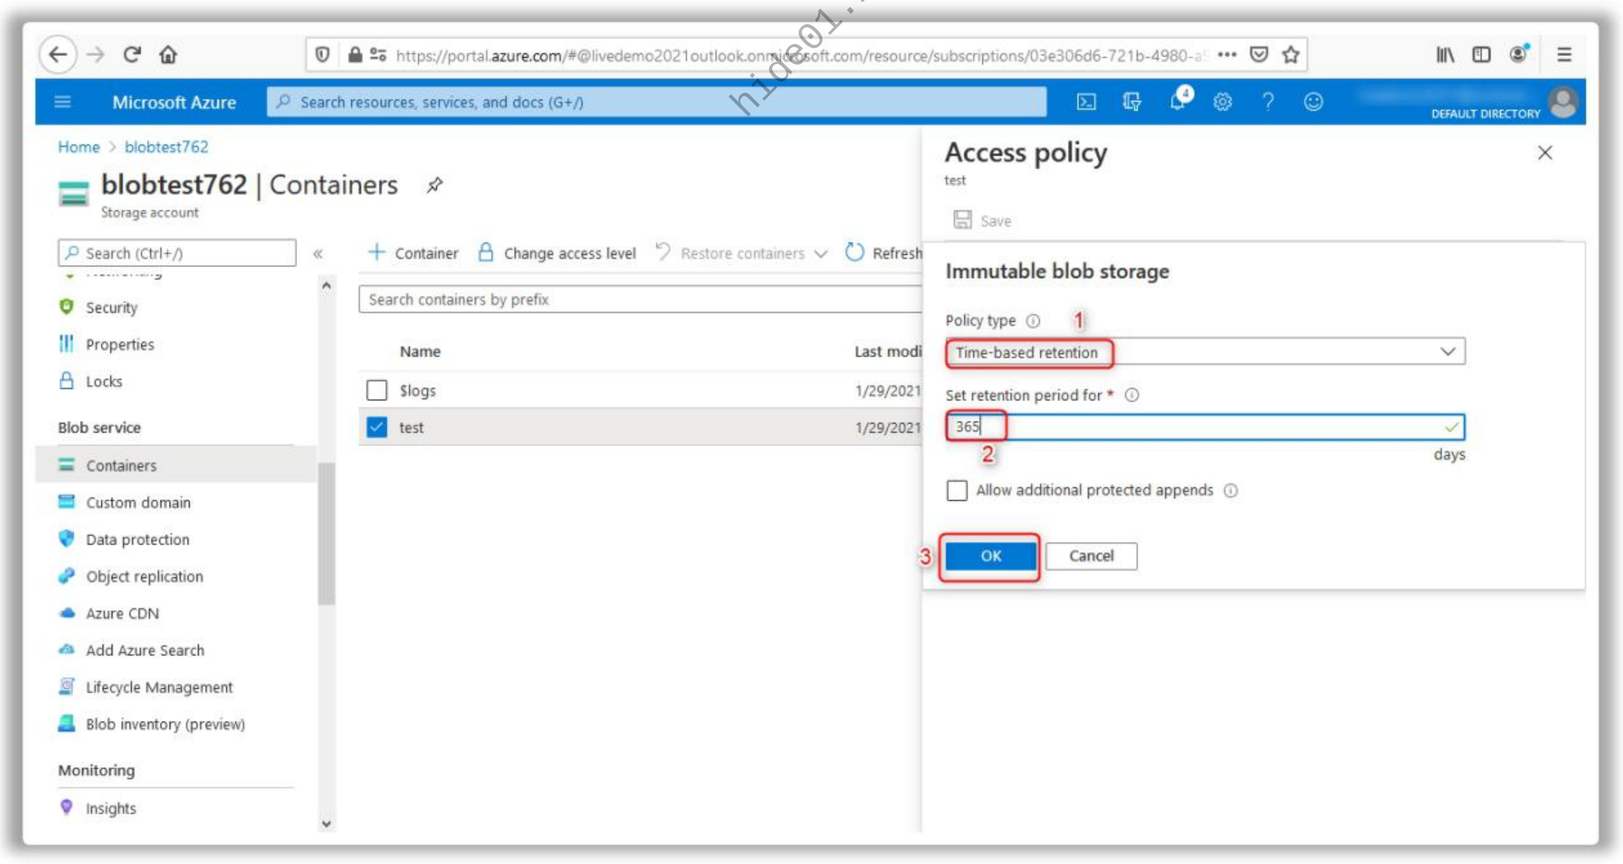


Figure 4.213: Click Ok

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

- To lock the policy, click on the three dots icon of the newly generated policy.

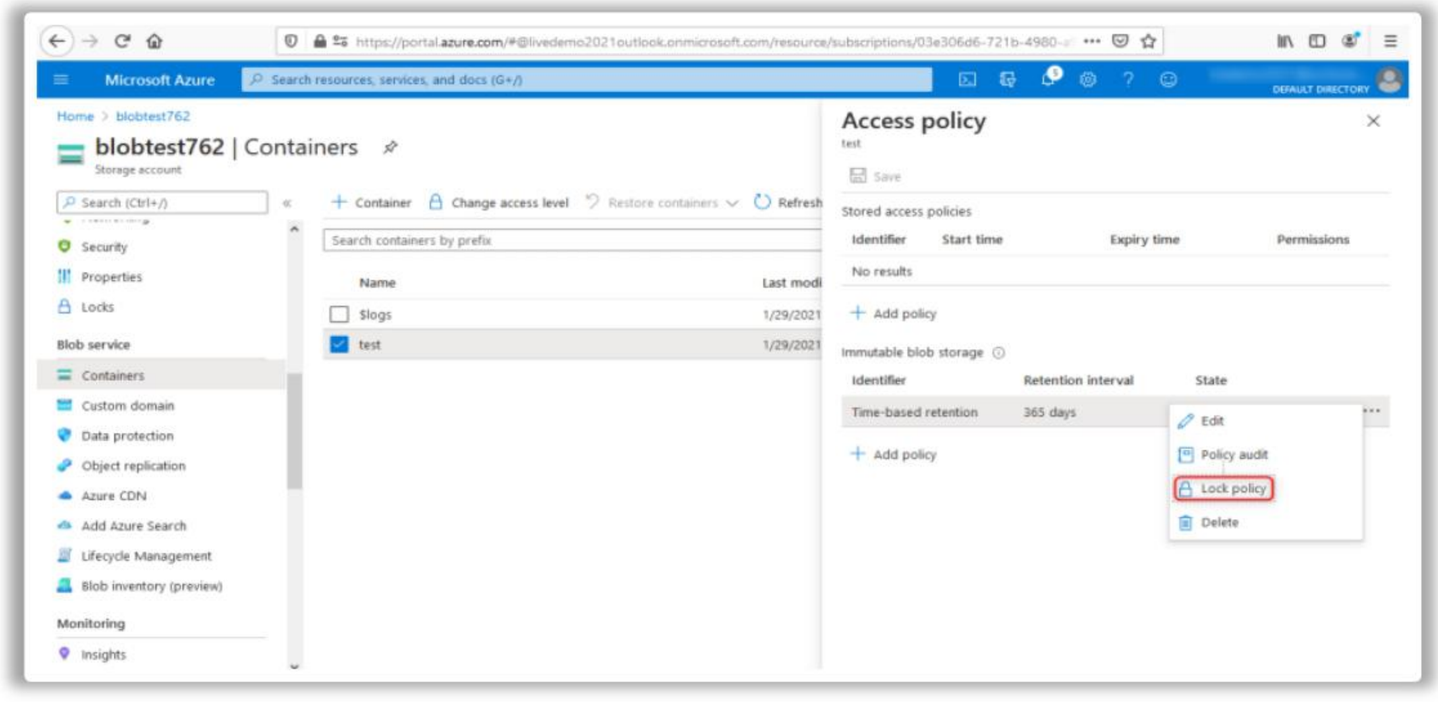


Figure 4.214: Select End Policy

- Type yes in the input field and click OK. The policy will be locked successfully

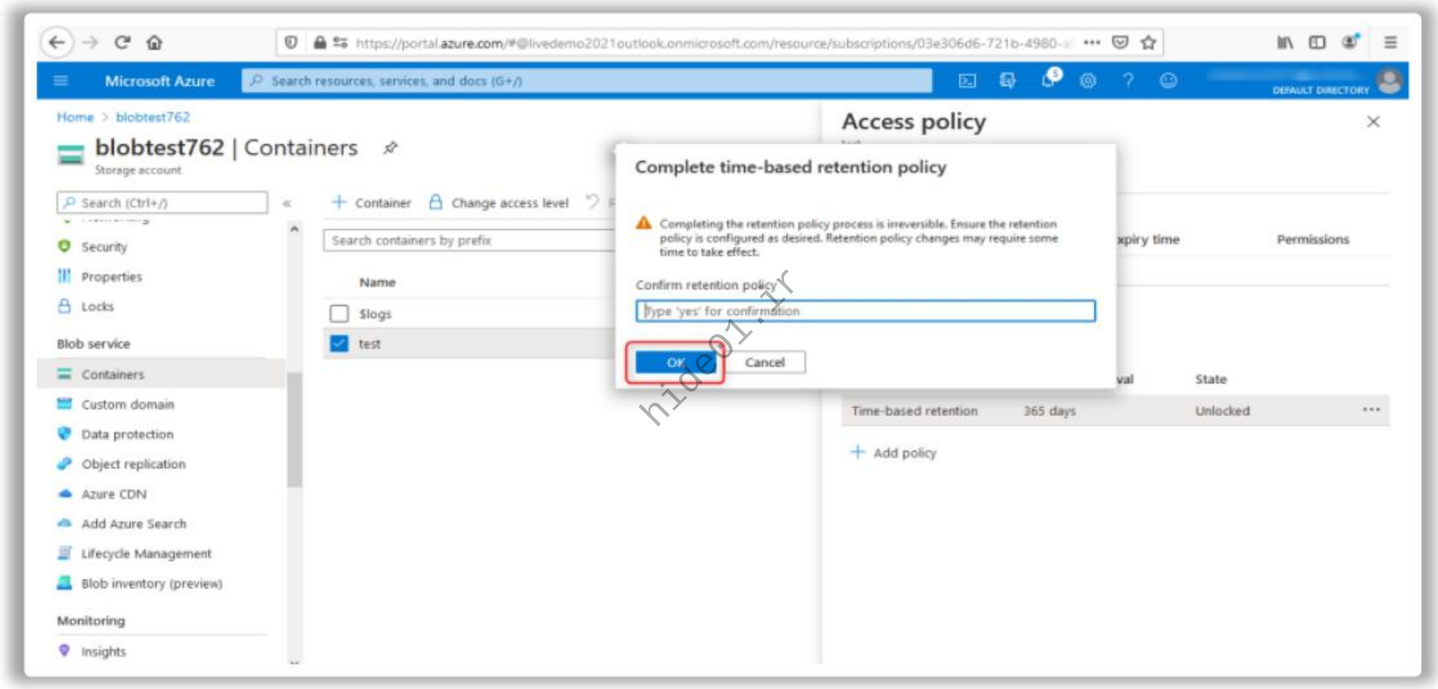
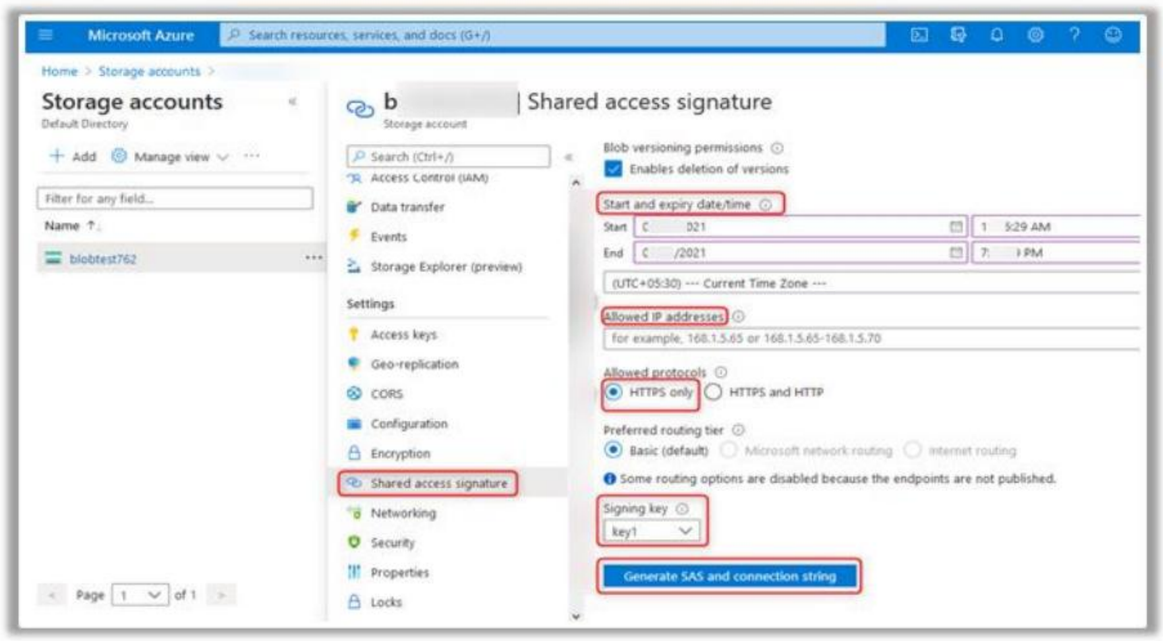


Figure 4.215: Click Ok

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

Allow Shared Access Signature Tokens Over HTTPS Only

- To establish a secure connection and to transfer data securely from Azure storage, it is recommended to allow SAS tokens over **HTTPS** only



Shared Access Signature Tokens Over HTTPS Only

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

Allow Shared Access Signature Tokens Over HTTPS Only

SAS grants limited access to Azure storage features. This can be shared with the people who require access to the storage without compromising security and privacy. This is the best way to provide temporary access to storage account features.

SAS token contains the information of storage services, resources, permissions, and validity of the token. This token authorizes the user to access the permitted storage services. To establish a secure connection and to transfer data securely from Azure storage, it is recommended to allow SAS tokens over HTTPS only.

Any security professionals who have previously generated SAS tokens that accept both HTTP and HTTPS look similar to the below image. These tokens can be replaced with the newly generated tokens that accept HTTPS only protocol, which is created using the below steps.



Figure 4.216: SAS Token

Source: https://azure.microsoft.com

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

To Allow Shared Access Signature Tokens over HTTPS Only

- Sign-in to the Azure portal.
- Select **Storage accounts** under **Azure services.**

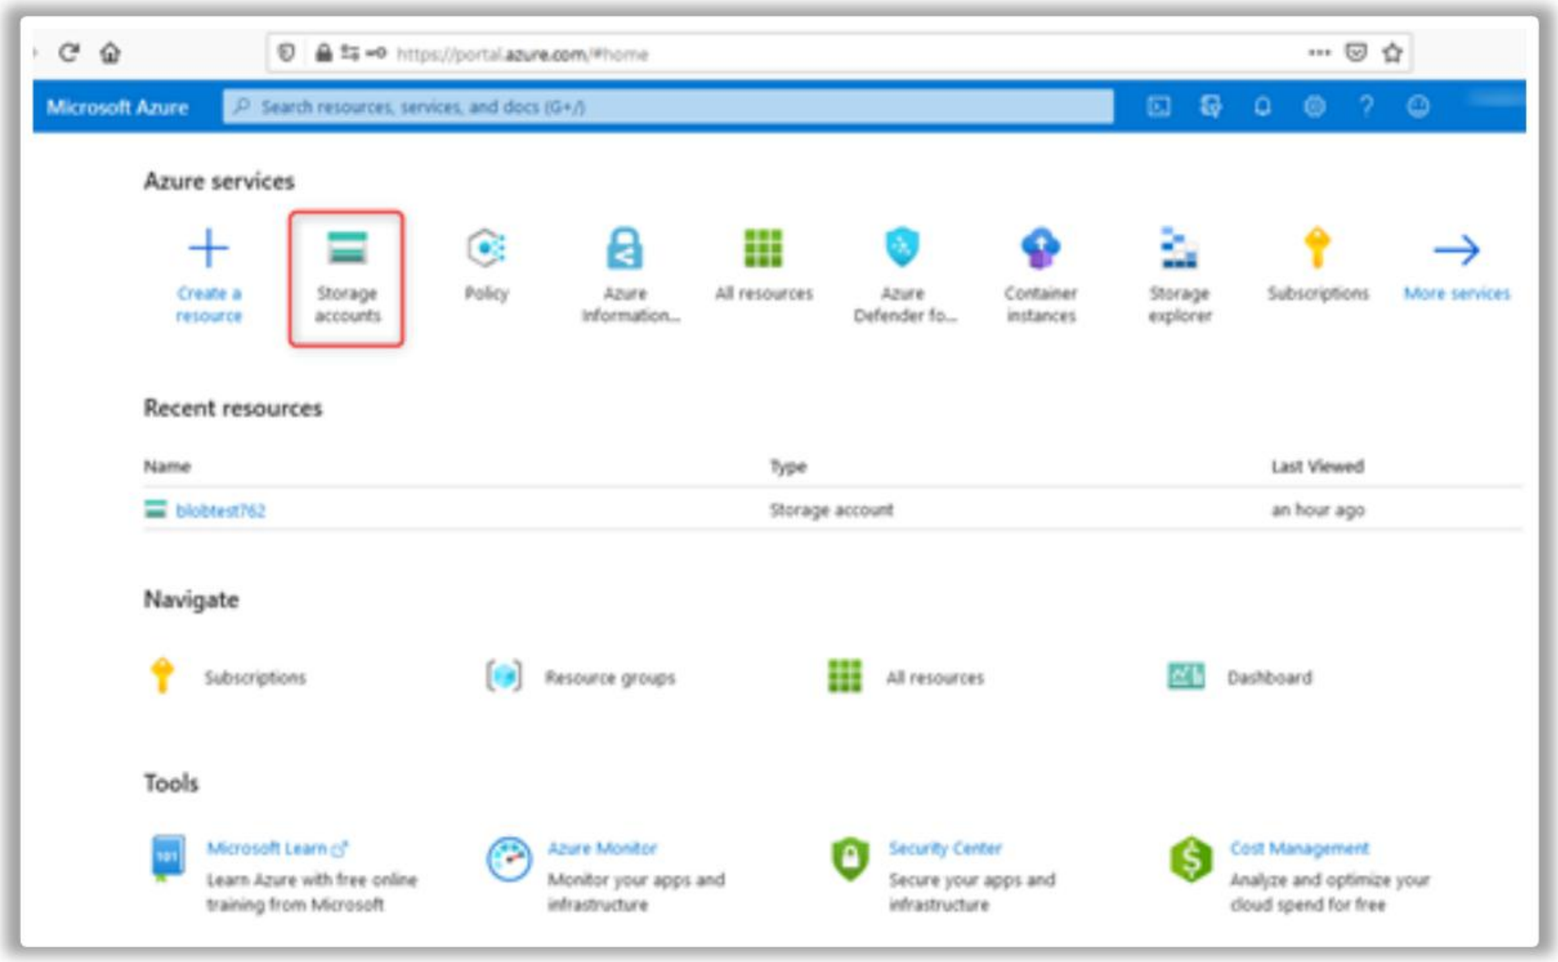


Figure 4.217: Select Storage Account

- A list of storage accounts will be displayed.
- Select the storage to generate a SAS.

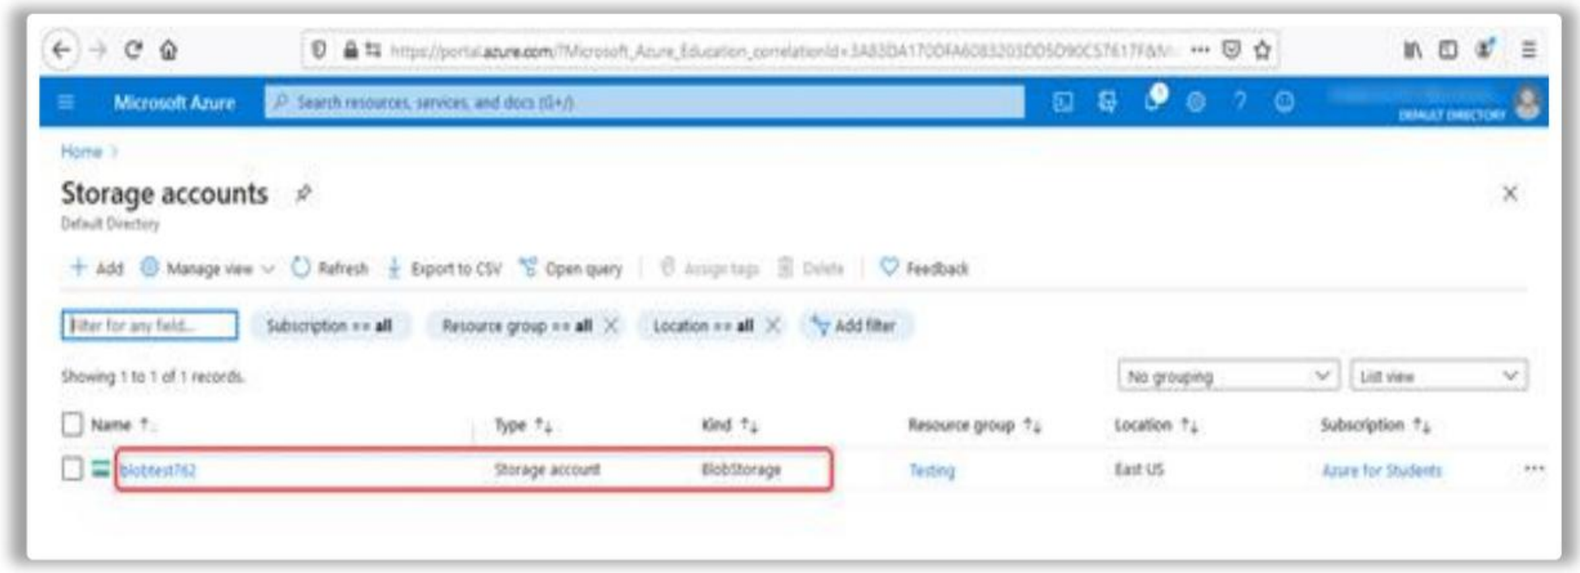


Figure 4.218: Click Storage

- Navigate to **Shared access signature** under **Settings**.
- Select the **Allowed services** that you want clients to access.
- Select the **Allowed resources types** as per your requirement.

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

- Select the necessary **Allowed permission** that clients can perform.

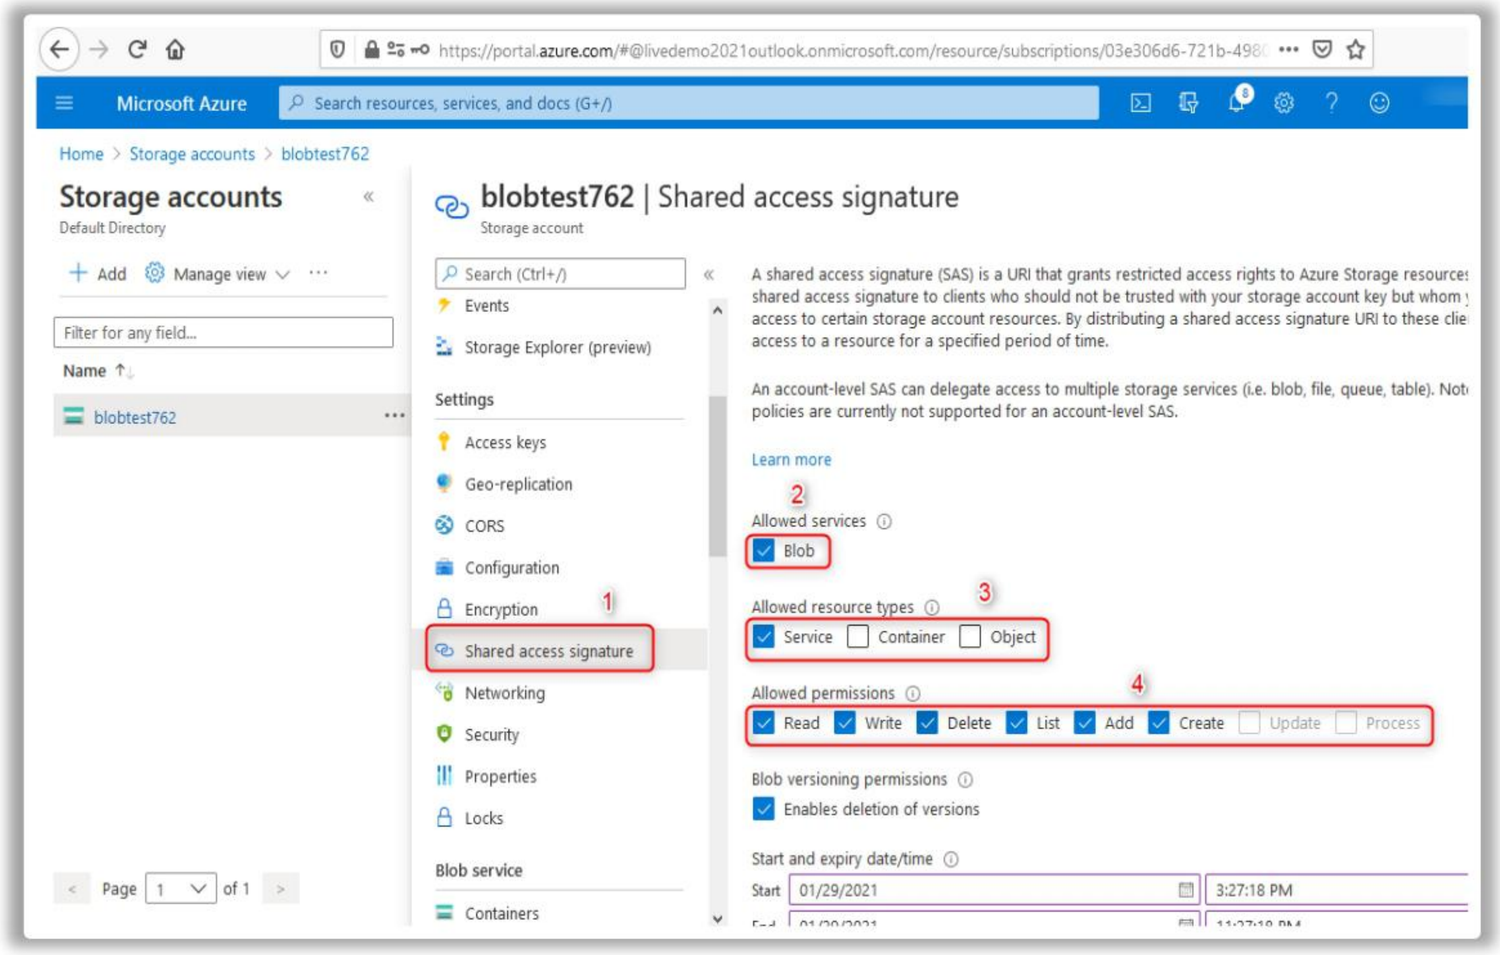


Figure 4.219: Allow Permission

- Choose the expiry using a calendar or input fields in **Start and expiry date/time.**
- Make sure to provide only the required IP addresses in the **Allowed IP addresses** field.
- In the **Allowed protocols** select the radio button of **HTTPS only.**
- Select the preferred key in the **Signing key** dropdown option.
- Click **Generate SAS and connection string** to generate tokens.

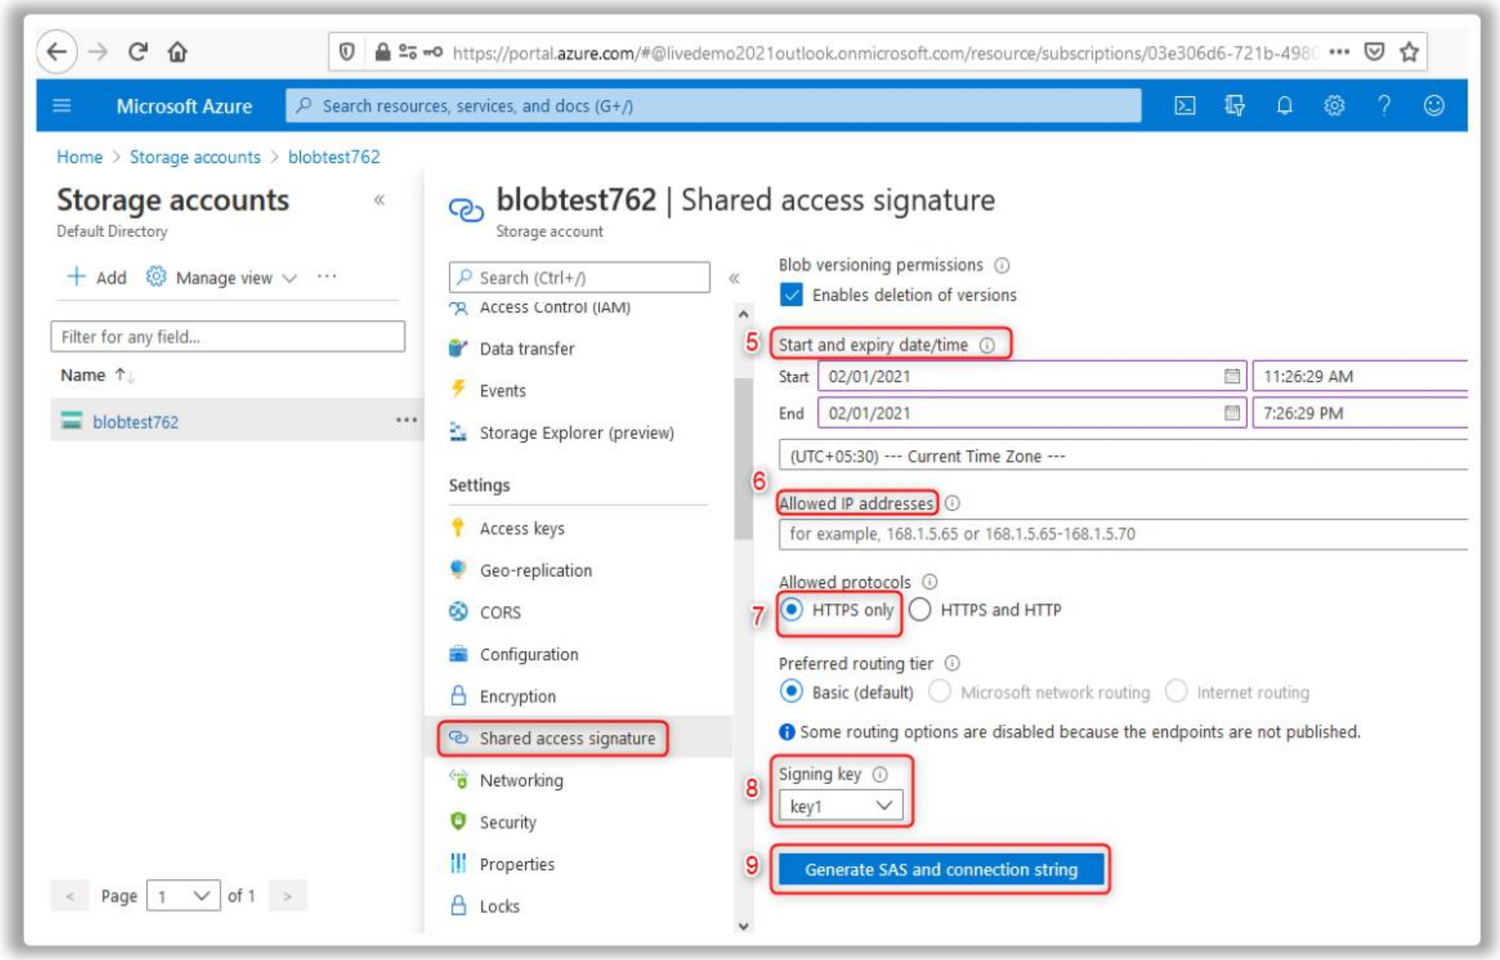


Figure 4.220: Generate SAS and Connection Strings

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

- The SAS is generated successfully and is configured to accept the HTTPS only protocol.

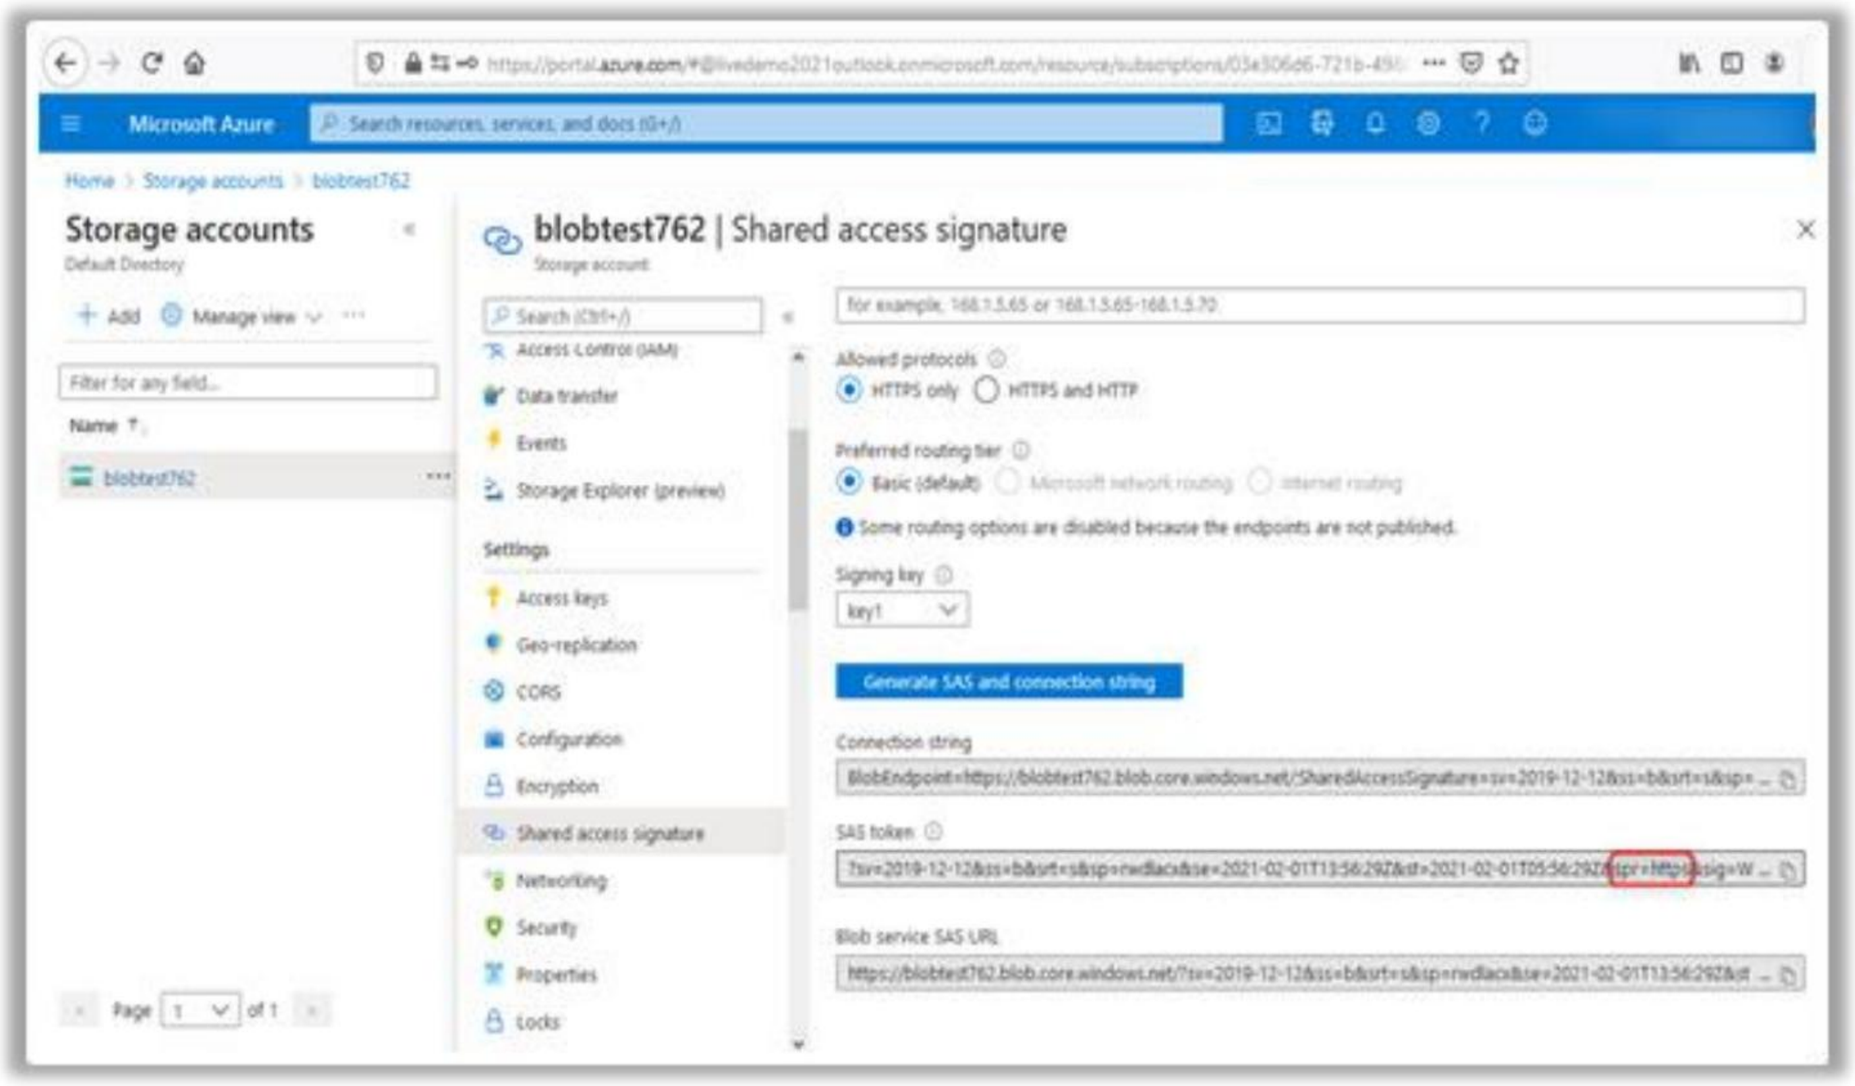


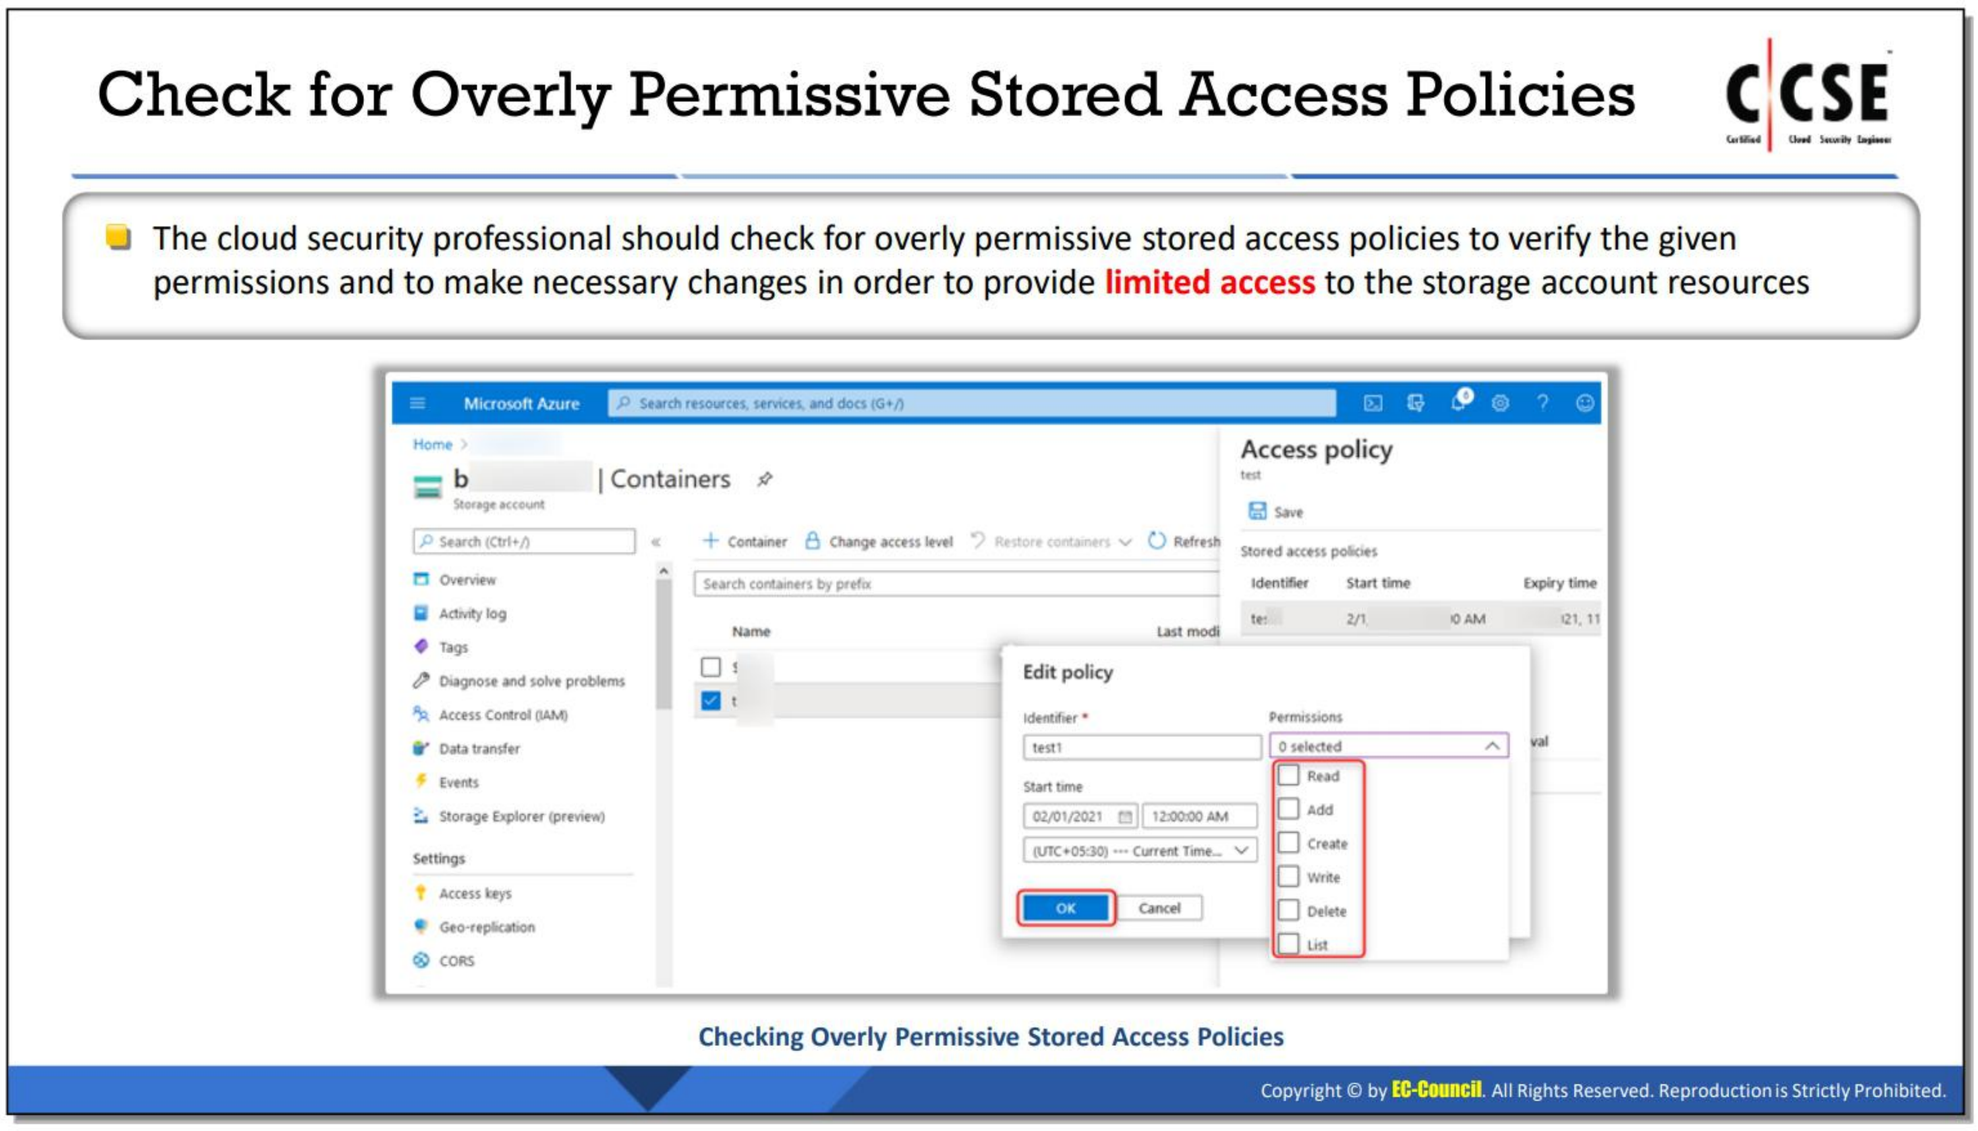
Figure 4.221: Check HTTPS Only Protocol

- The below image clearly shows the SAS token, which is now accepting HTTPS only.



Figure 4.222: SAS Token

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



Check for Overly Permissive Stored Access Policies

A stored access policy is an additional layer of control for service-level SASs. These stored access policies control group shared access signatures by changing or overwriting the parameters like validity (start time and expiry time) and permissions (read, write, list, delete, add, create, and list) as per requirement. Shared access signatures should not be given full access to storage account resources like the Blob, tables, files, and queues, etc. through stored access policies. It is recommended to verify the permissions given and make necessary changes.

Source: https://azure.microsoft.com

To Check for Overly Permissive Stored Access Policies

- Sign-in to the Azure portal.
- Select **Storage accounts** under **Azure services.**

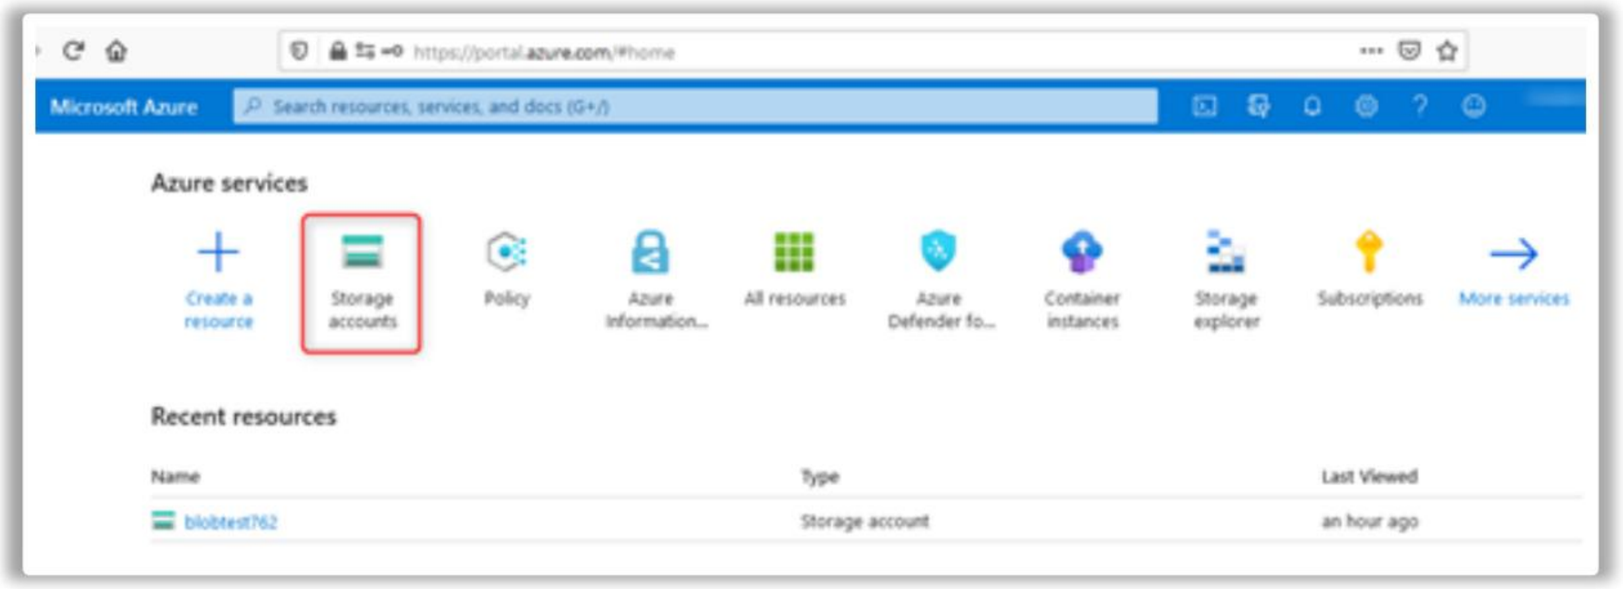


Figure 4.223: Select Storage Accounts

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

- A list of storage accounts will be displayed.
- Select the storage account that you want to check and reconfigure.

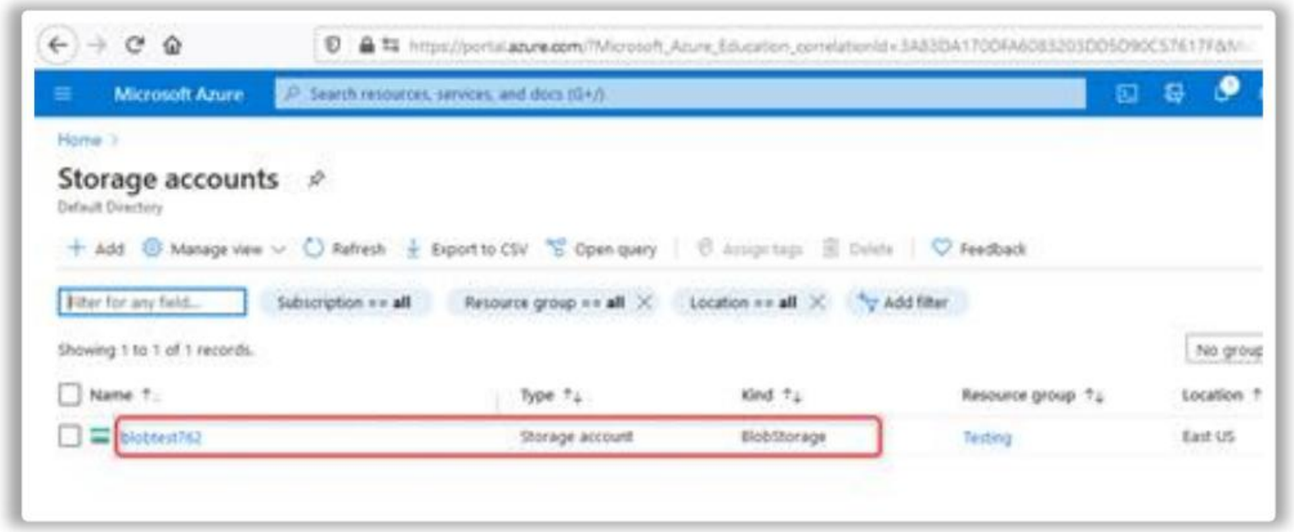


Figure 4.224: Choose Storage

- Navigate to **Containers** under **Blob service**.
- Click on the three dots icon as shown in the image and go to the **Access policy**.

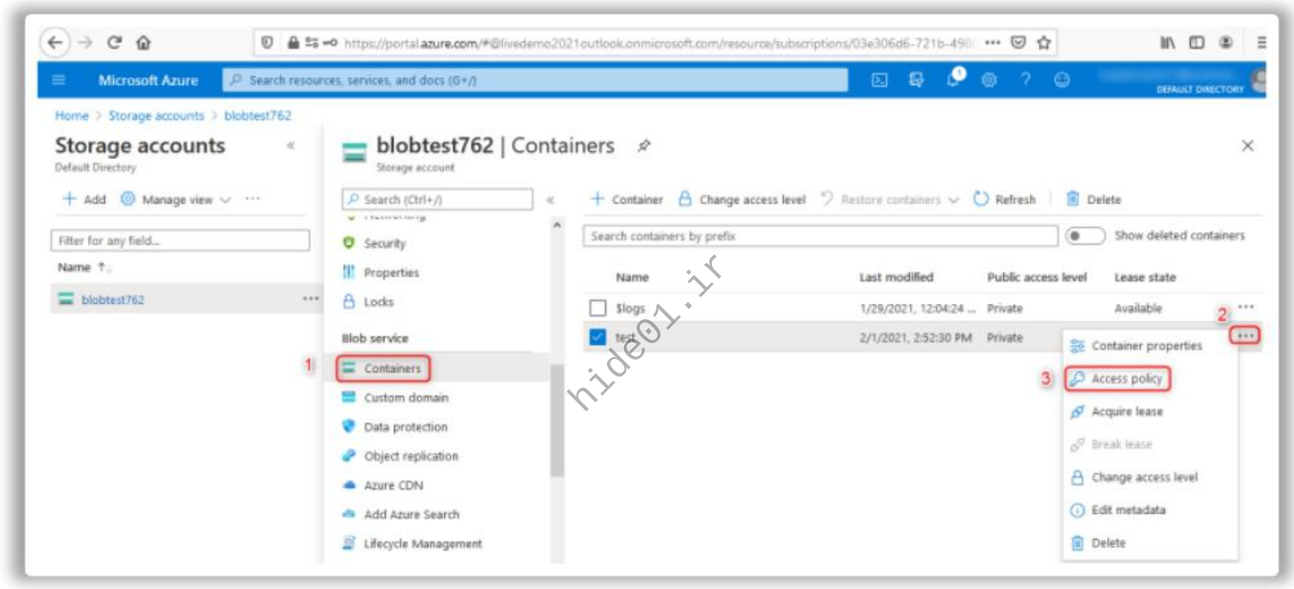


Figure 4.225: Create Access Policy

- Review the **Access policy** i.e. **Start time**, **Expiry time**, and **Permissions**.
- In the below image, all permissions like "**racwdl**" are given, that is, r-Read, a-Add, c-Create, w-Write, d-Delete, and l-List.

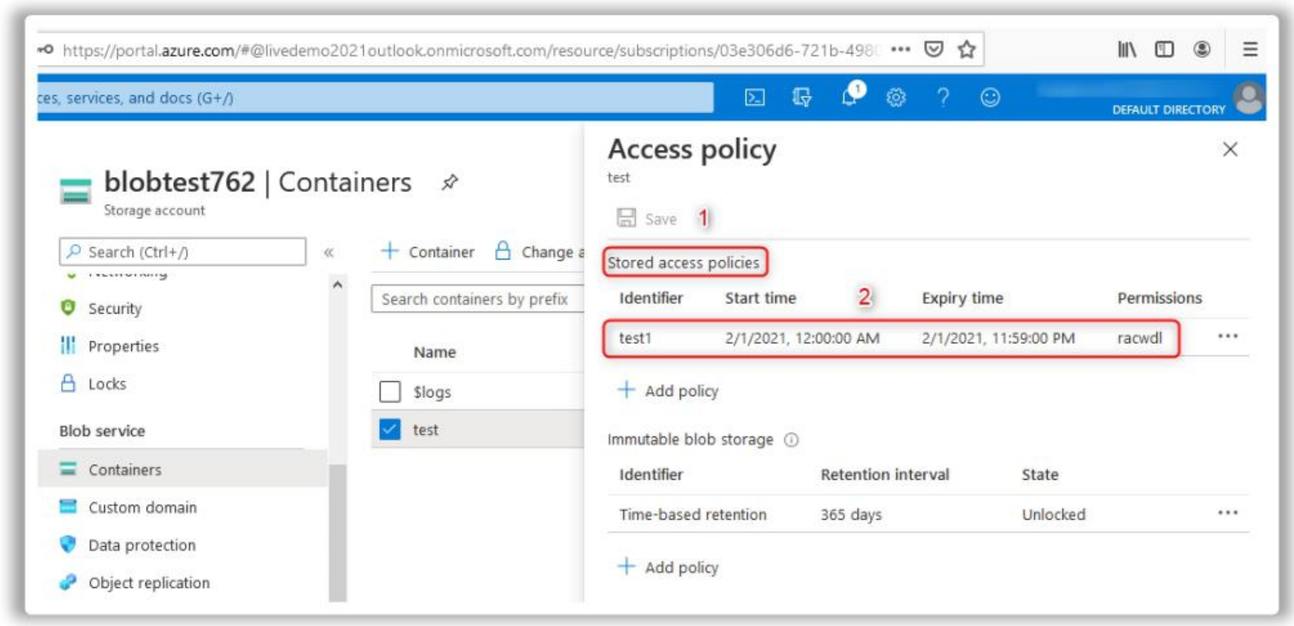


Figure 4.226: Allow Permissions

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

- Click on the three dots icon as shown in the image and go to **Edit**.

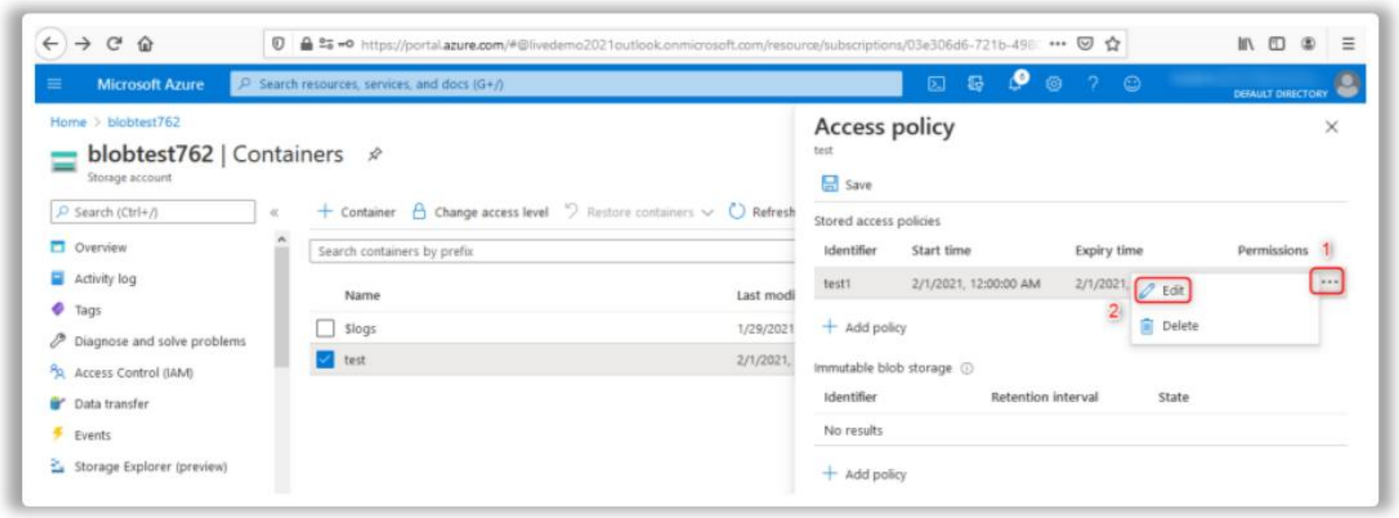


Figure 4.227: Click Edit

- Select the permissions that are required and click **OK**.

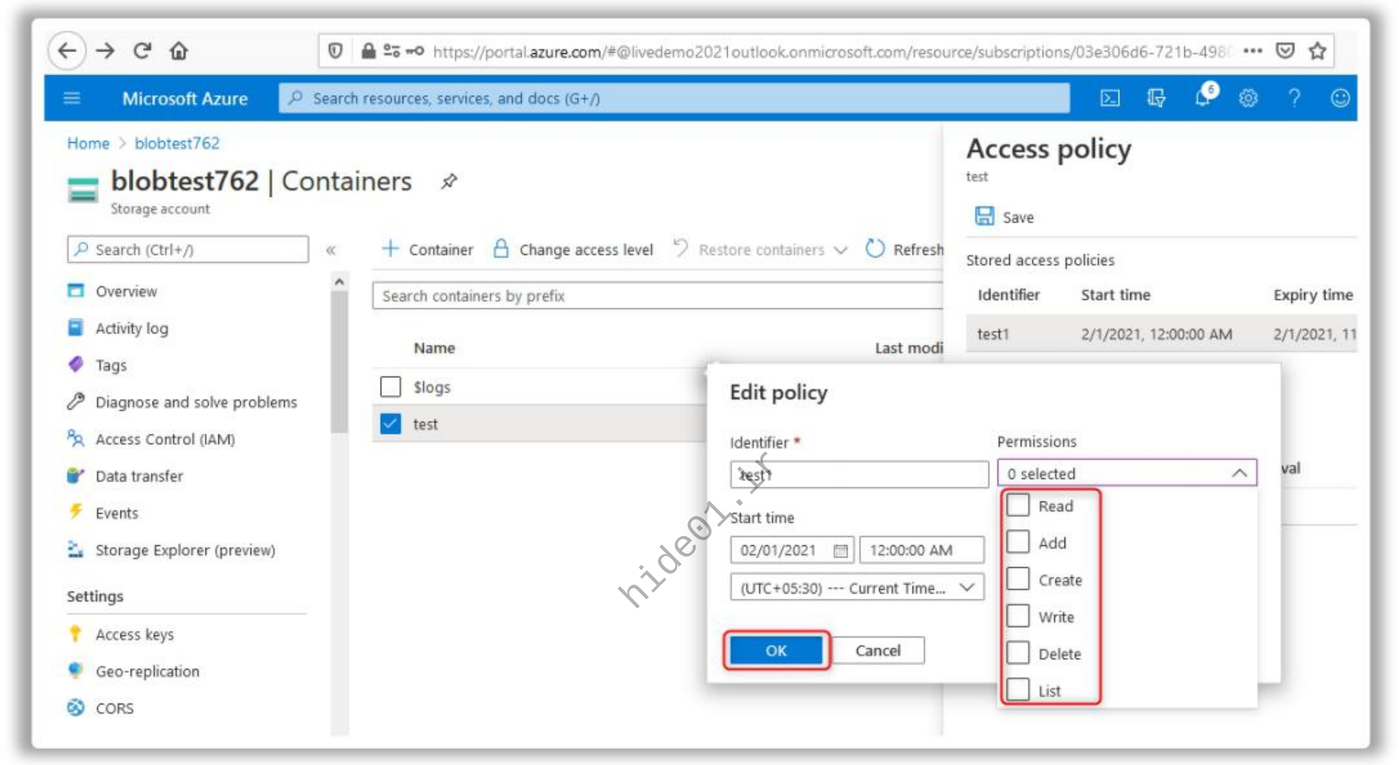


Figure 4.228: Click Ok

- Here in the below image, “racl” permissions are given or modified from full permissions.

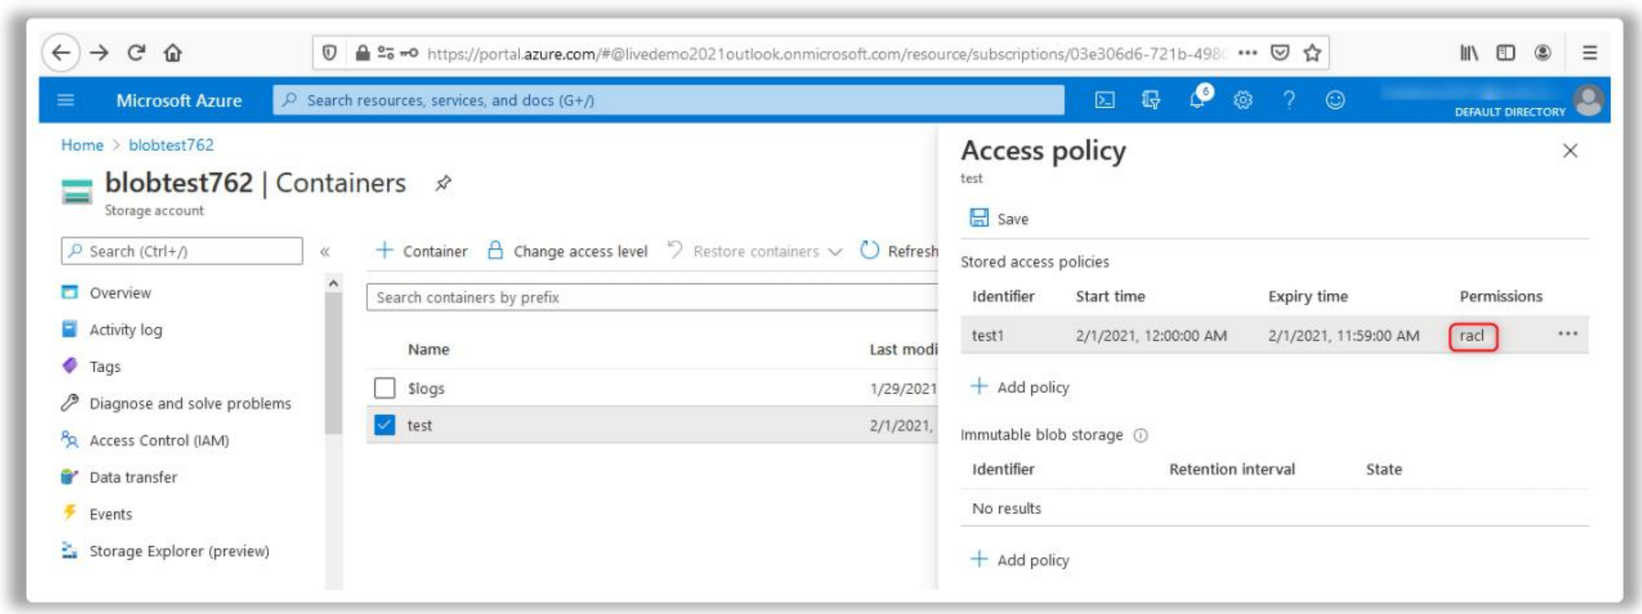
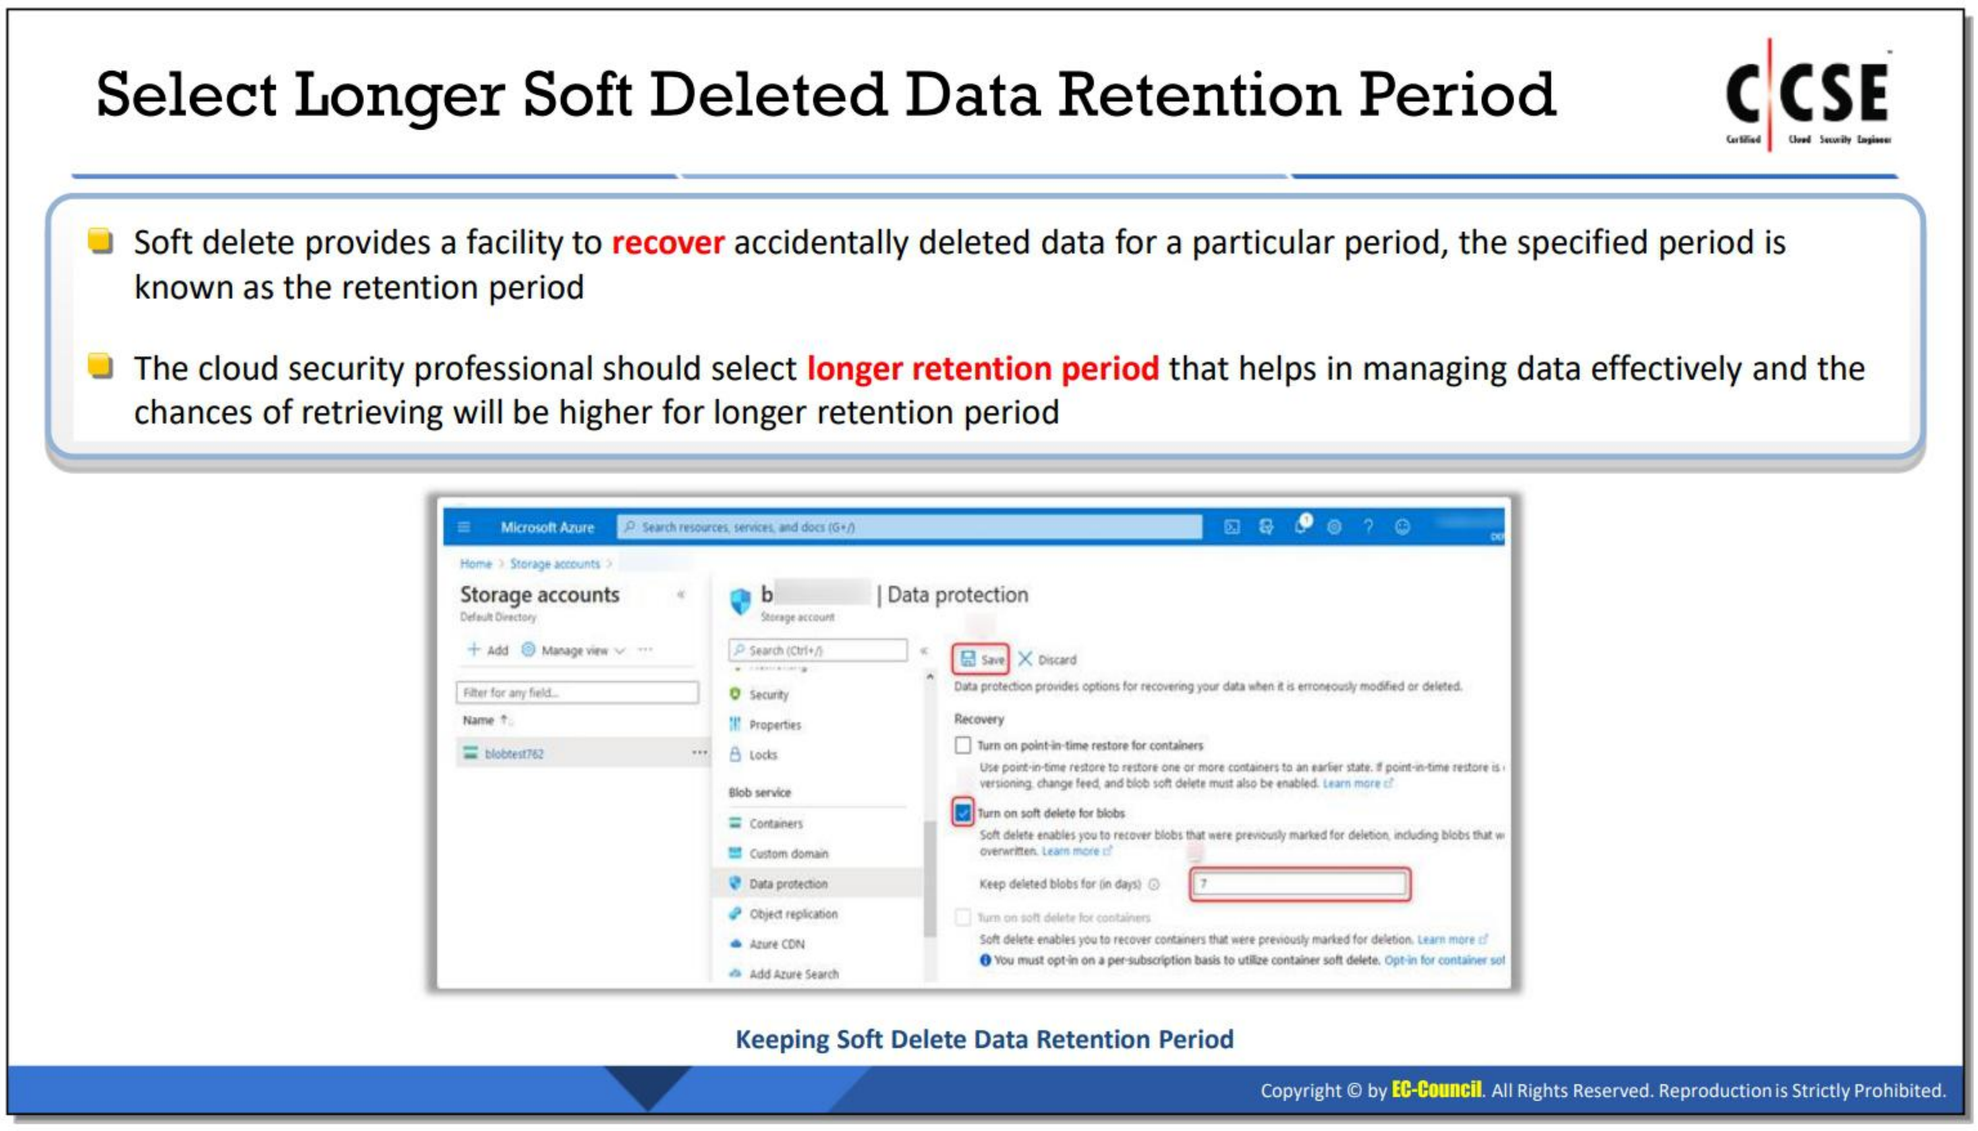


Figure 4.229: Allow Permission

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



Select Longer Soft-Deleted Data-Retention Period

Soft delete provides a facility to recover accidentally deleted data for a particular period, the specified period is known as the retention period. Azure provides a retention period from 1 day to 365 days; customers can opt as per their requirement.

One of the best ways to comply with regulations regarding data retention is having a longer retention period. It helps to manage the data more effectively in case of any failures or data losses.

Source: https://azure.microsoft.com

To Check Soft Delete Data-Retention Period

- Sign-in to the Azure portal.
- Select **Storage accounts** under **Azure services.**

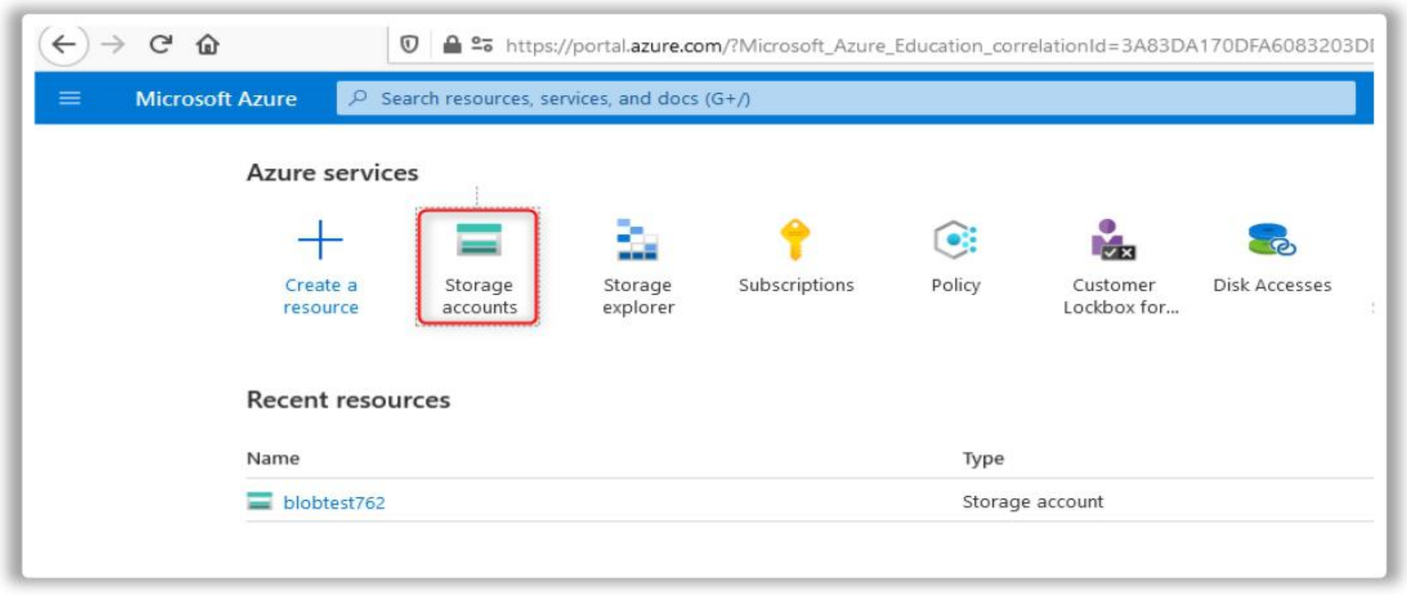


Figure 4.230: Select Storage Accounts

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

- A list of storage accounts will be displayed.
- Select the storage for which you want to check the soft delete retention period.

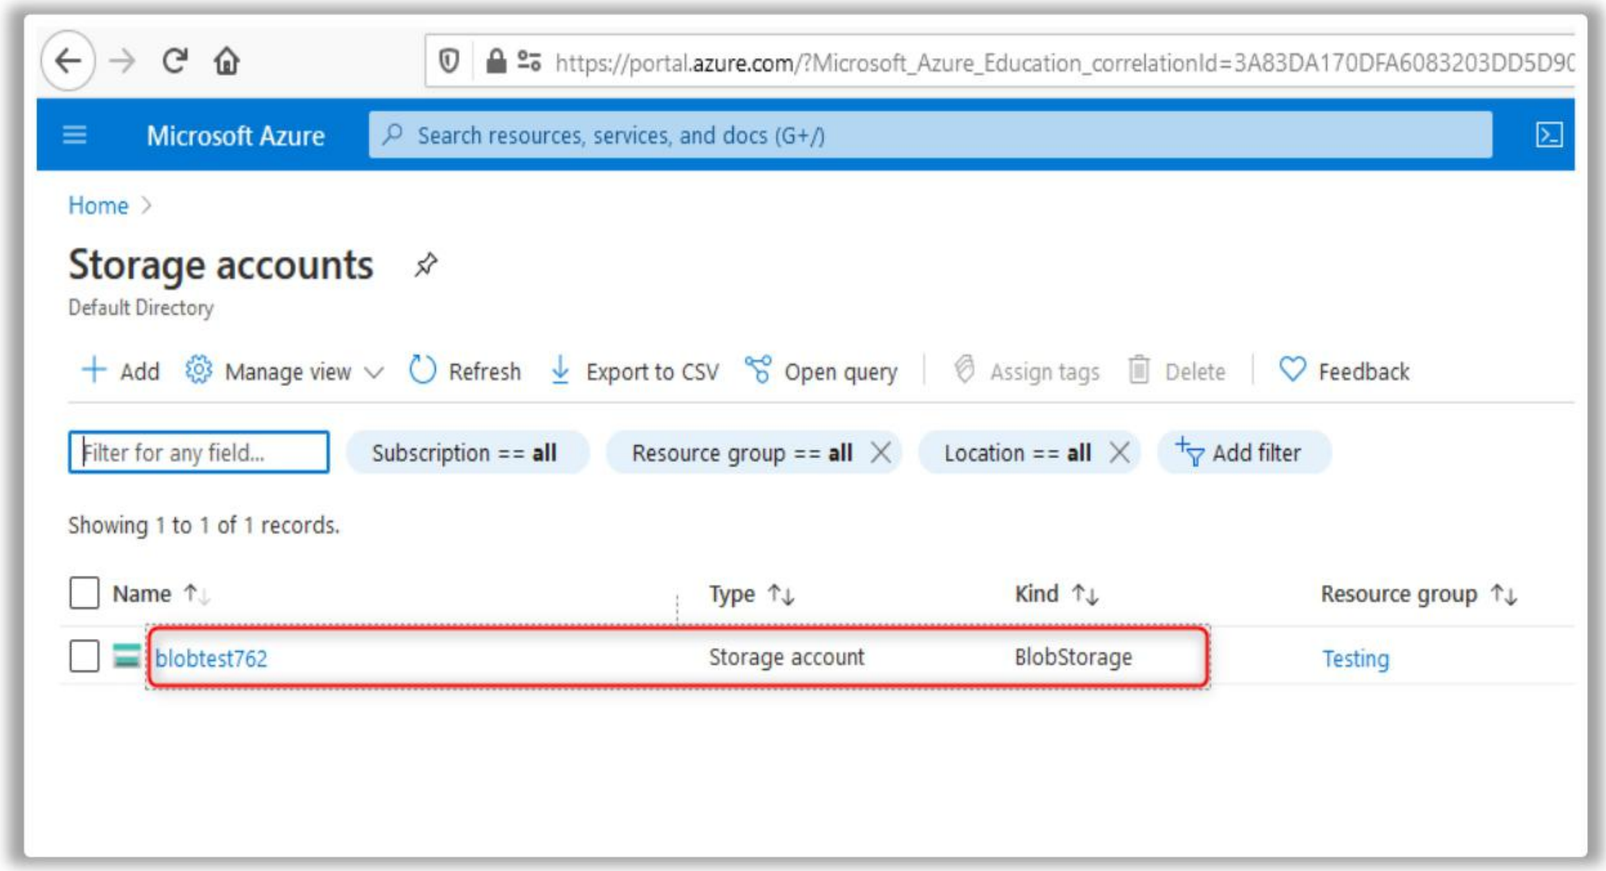


Figure 4.231: Choose Storage

- Locate **Blob service** and navigate to **Data protection**.

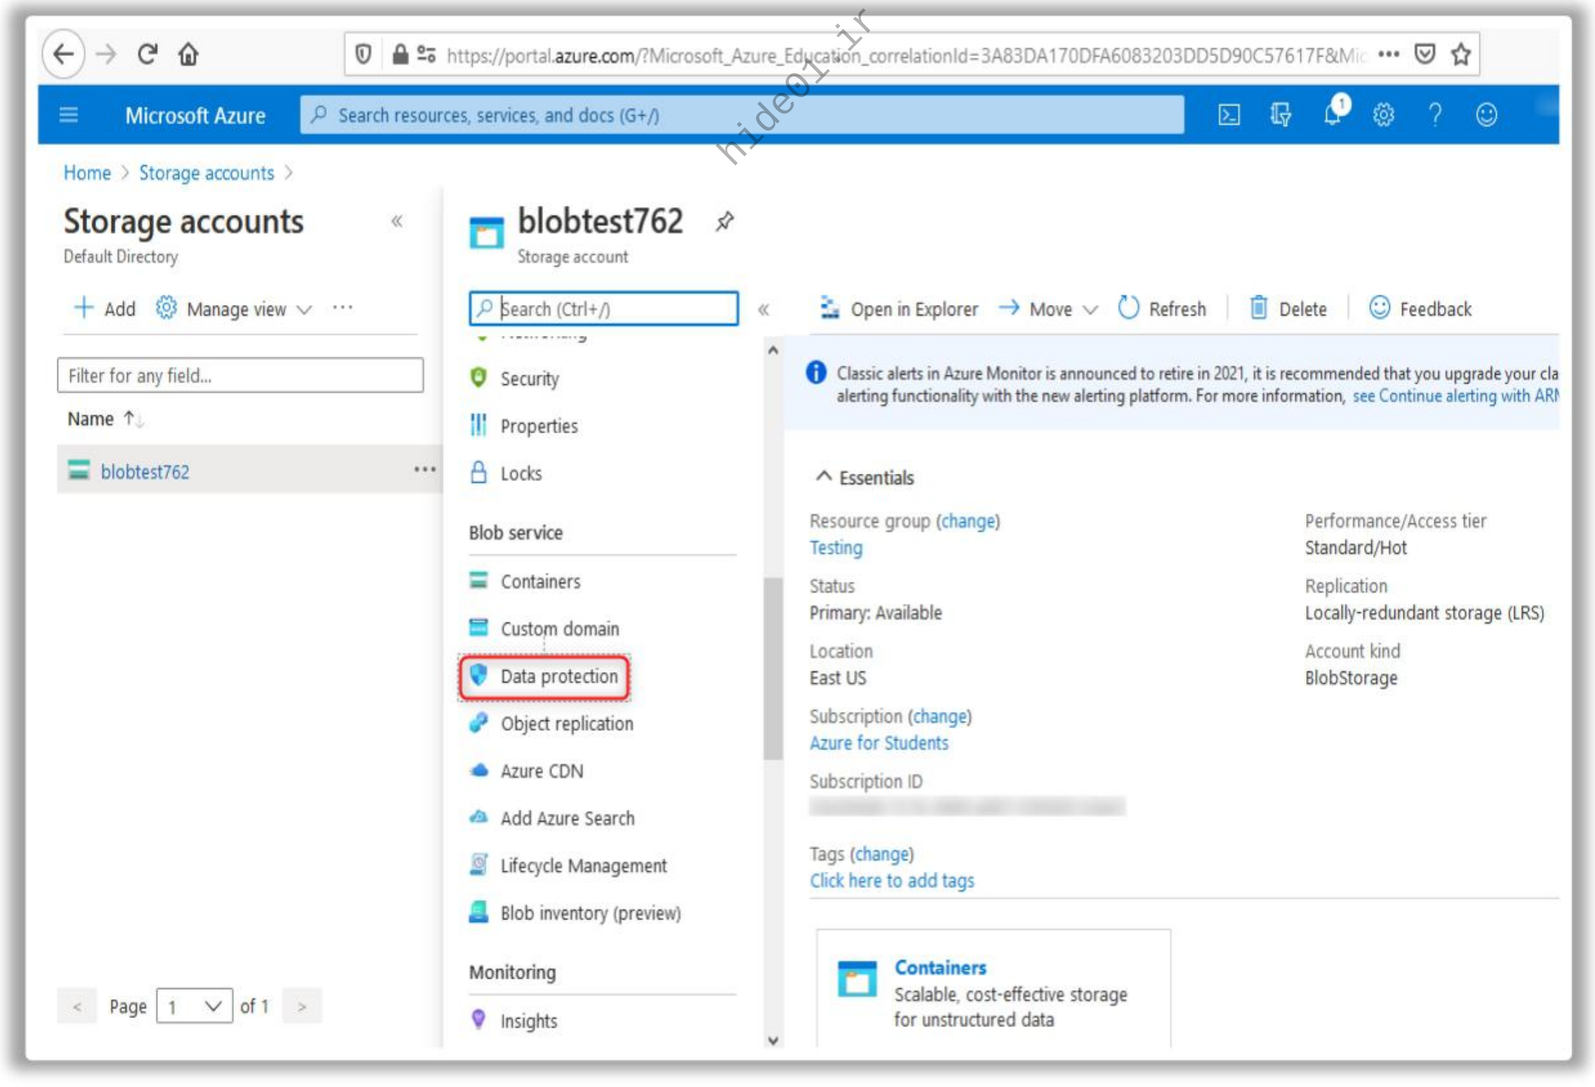


Figure 4.232: Select Data Protection

- Verify the number of days in "**keep deleted blobs for (in days)**".
- Soft delete for Blob protection has been verified.

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

- Change the value of the retention period as per the requirement if required and click **Save**.

Note: Updated retention period applies only to newly deleted data.

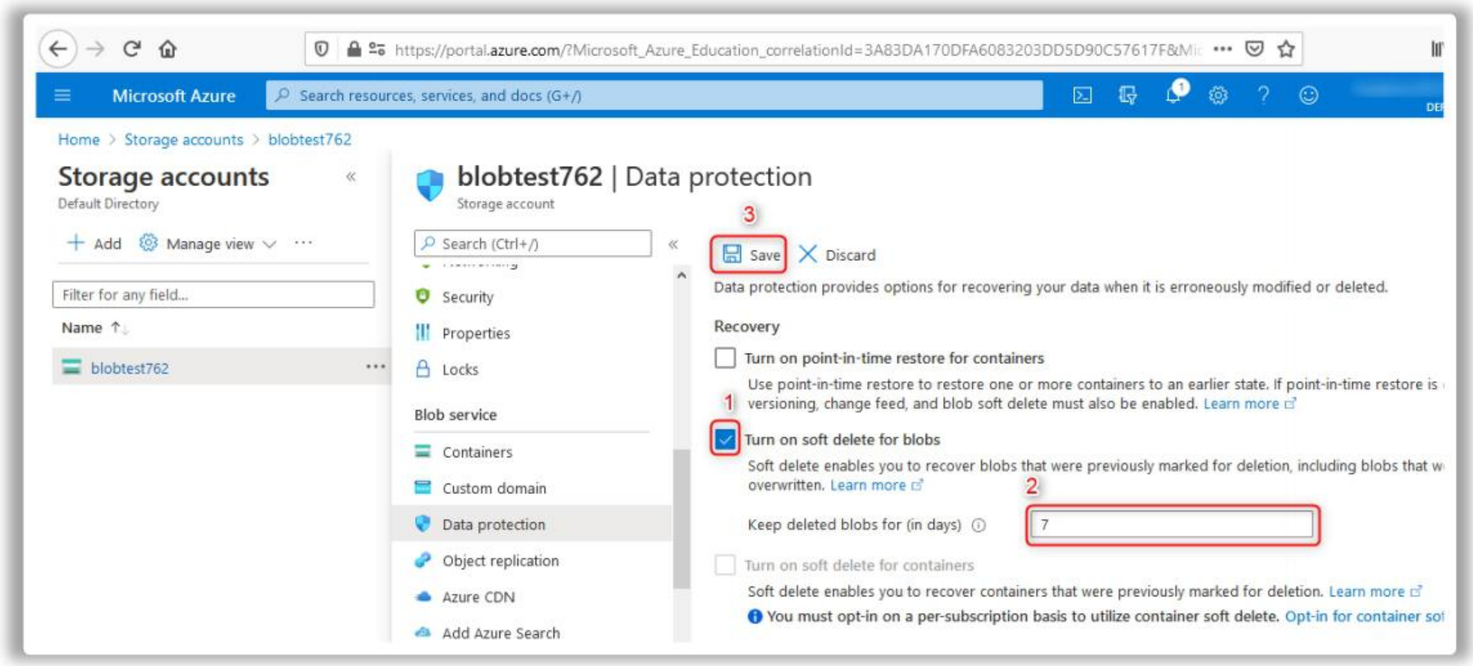
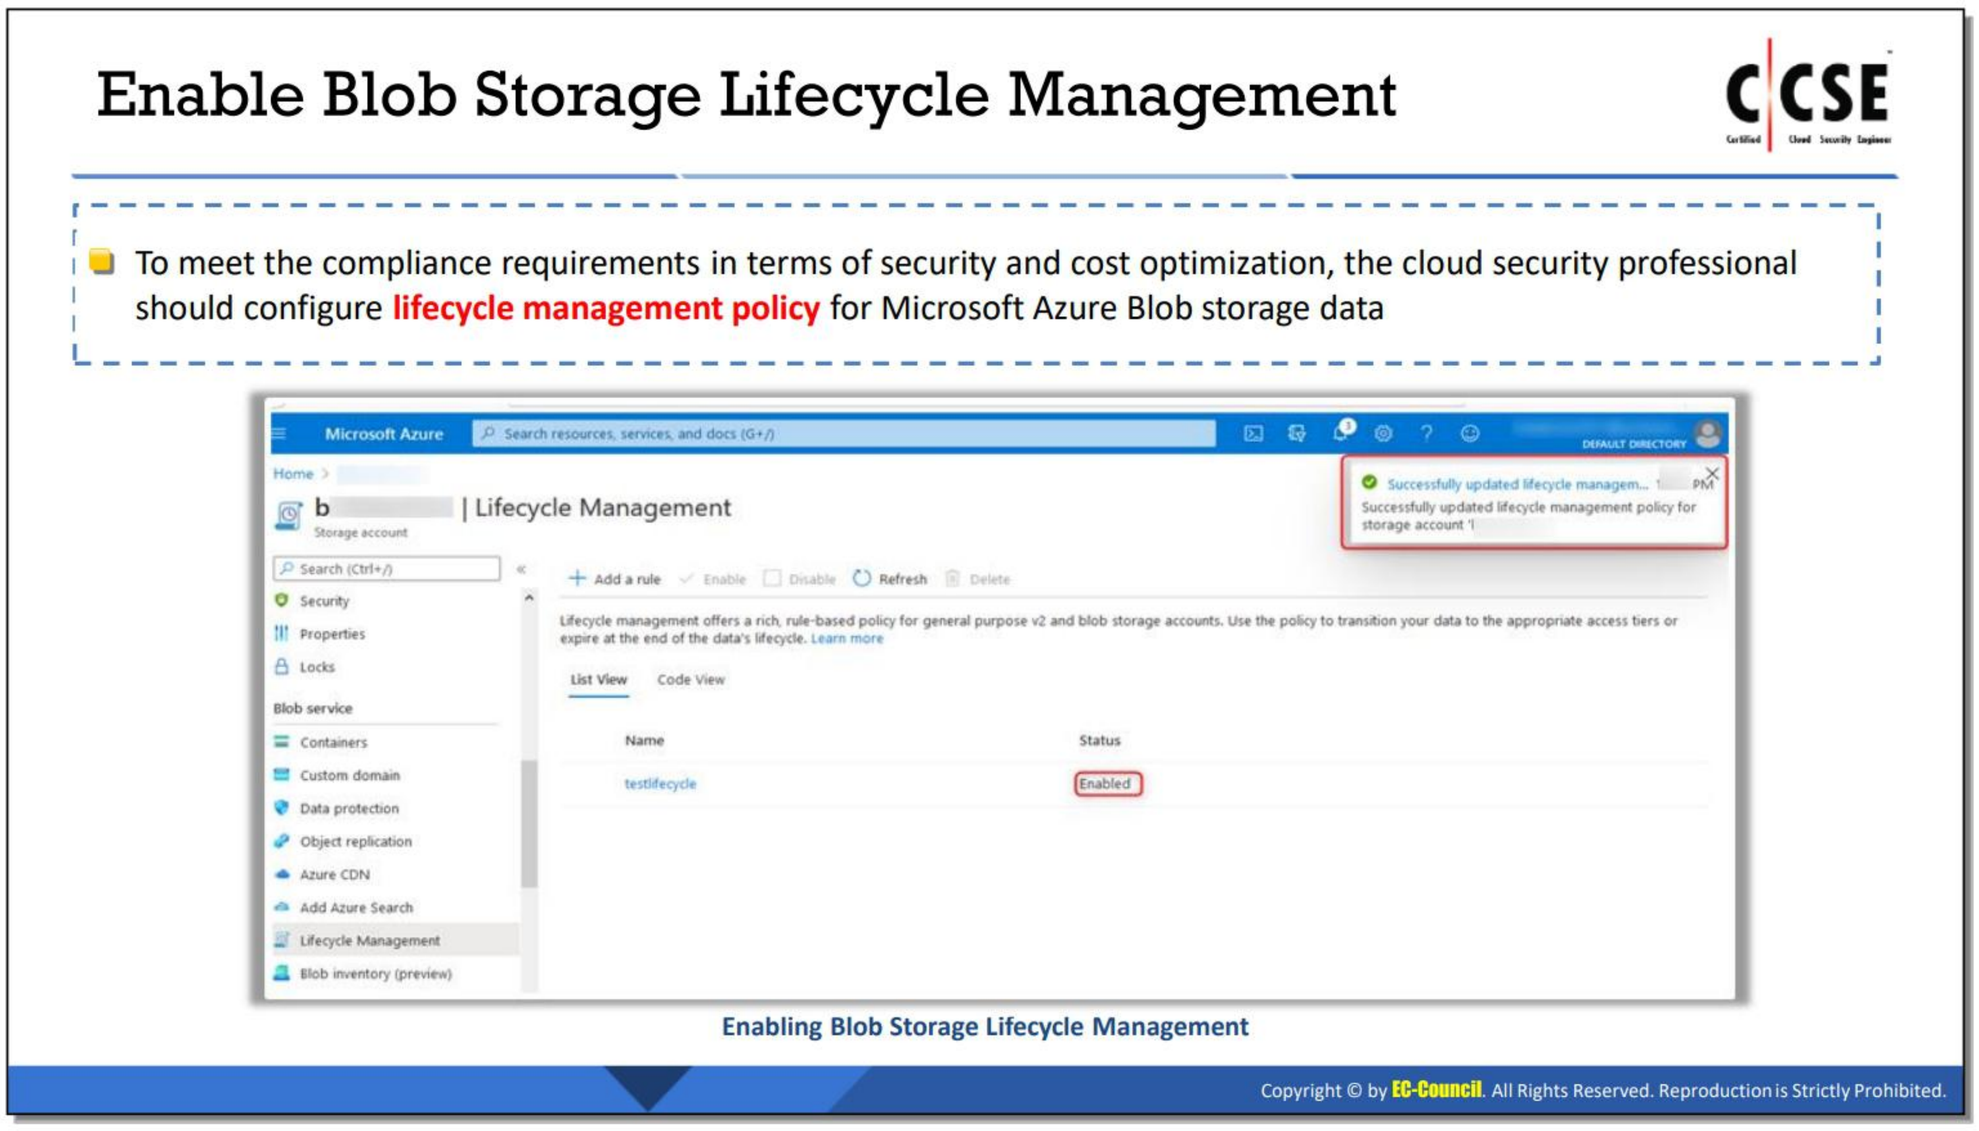


Figure 4.233: Click Save

hide01.ir

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



Enable Blob Storage Lifecycle Management

For meeting compliance requirements in terms of security and cost optimization, the cloud security professional should configure lifecycle management policy for Microsoft Azure Blob Storage data. Depending on the lifecycle management policy defined at the storage account level, the Blob data can be automatically deleted or transit Blob data to appropriate storage tier.

Azure provides Blob Storage lifecycle management that acts according to the requirement using customized rules. Some of the data is accessed frequently and sometimes it may be accessed after longer durations. This lifecycle management helps to store data in the appropriate tiers.

Consider, a particular data is accessed frequently when it is created and later it may be required once in two months and later quarterly or other. Lifecycle management allows users to set up rules as per their requirement, that is, it can be set up to move the data to cool storage, 30 days after the last modified day, set to 90 days to move into an archive, and can set up deletion as well.

Source: https://azure.microsoft.com

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

To enable Blob Storage Lifecycle Management

- Sign-in to the Azure portal.
- Select **Storage accounts** under **Azure services.**

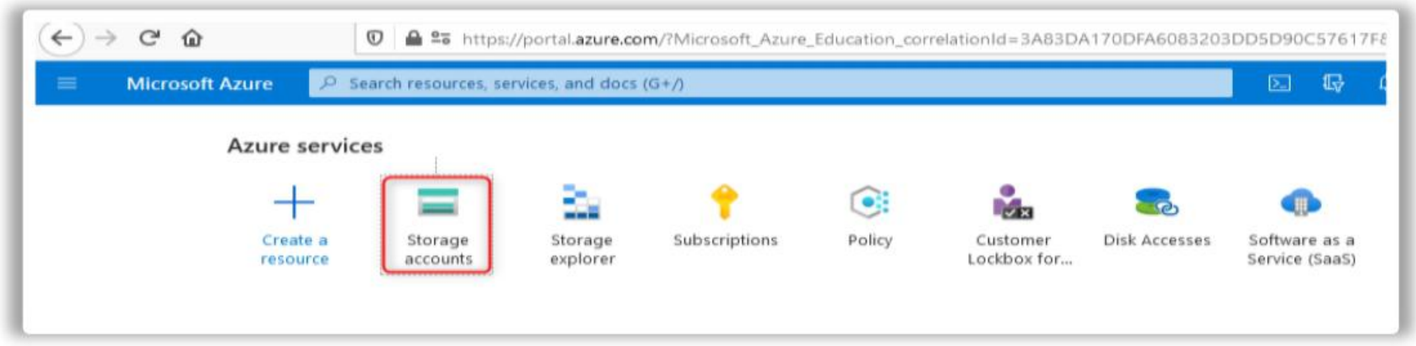


Figure 4.234: Select Storage Accounts

- A list of storage accounts will be displayed.
- Select the storage for which you want to enable Blob storage lifecycle management.

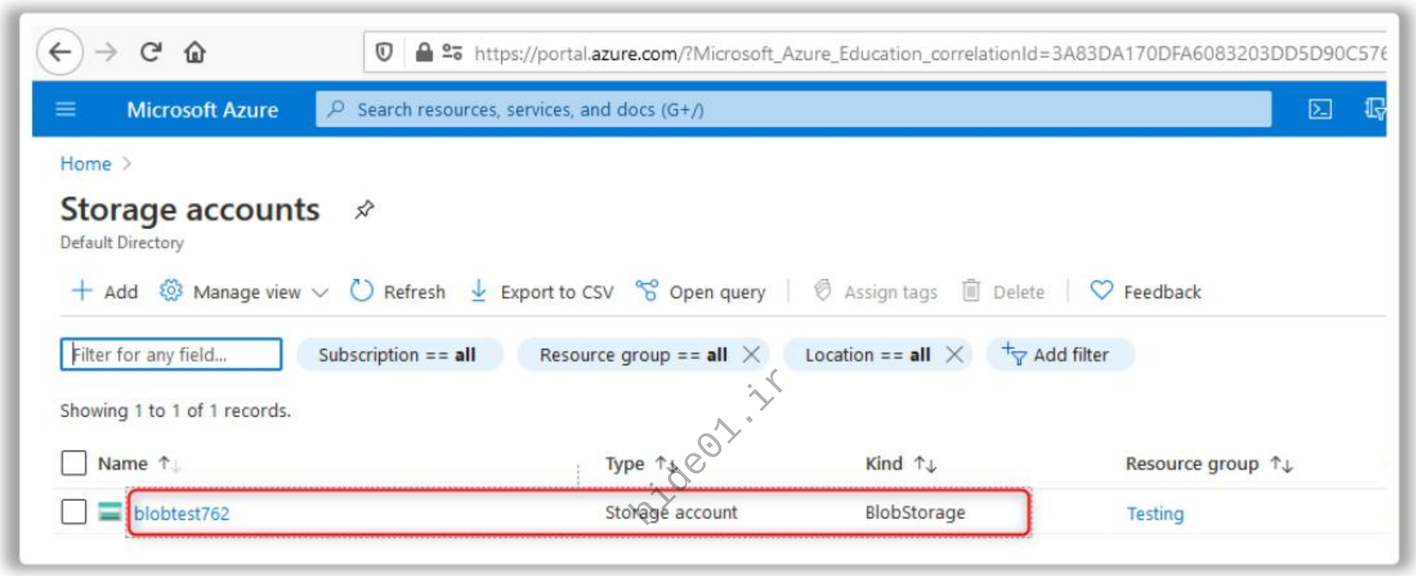


Figure 4.235: Choose Storage

- Locate **Blob service** and go to **Lifecycle Management.**
- Click on **Add a rule**.

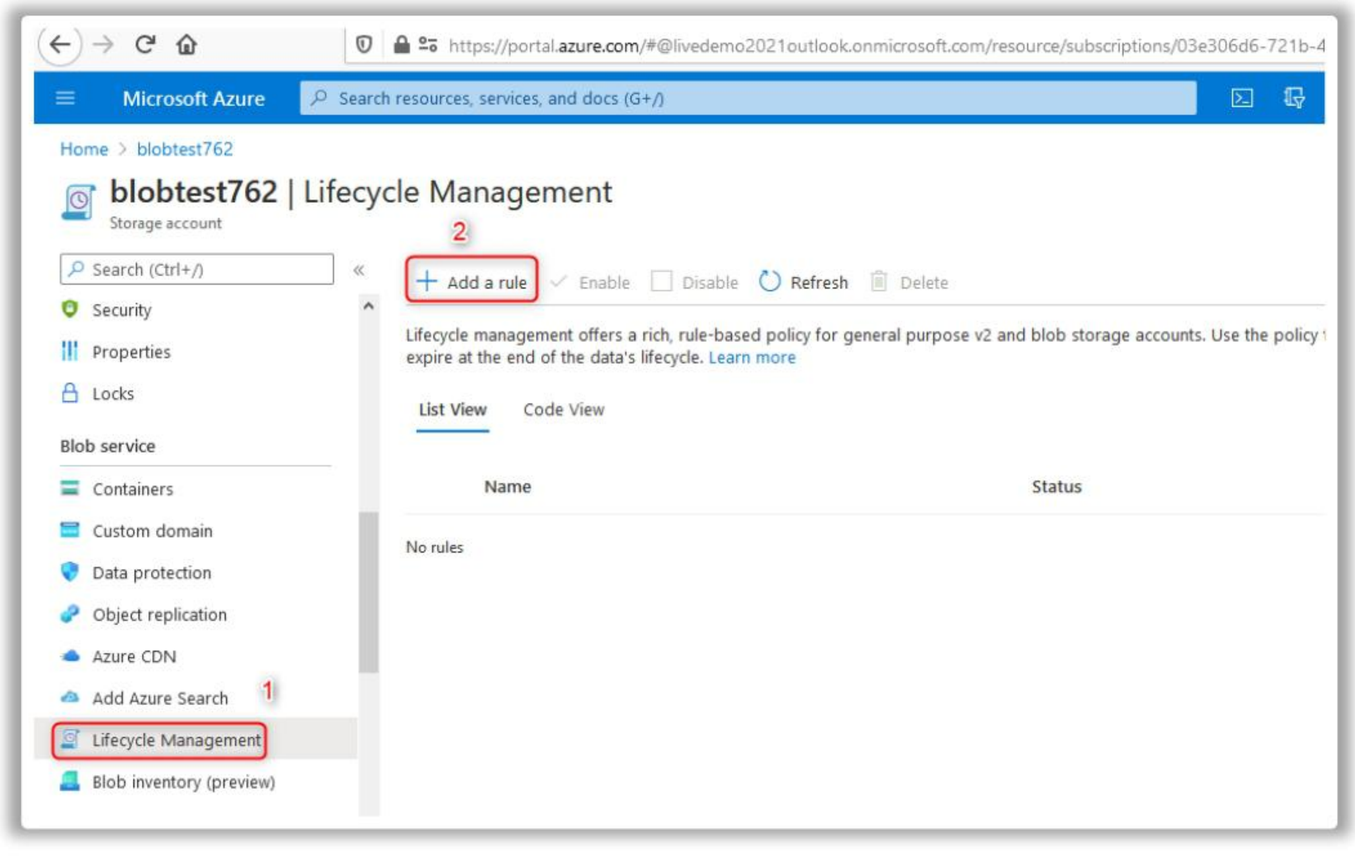


Figure 4.236: Click Add a rule

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

- Enter and select the required details like **Rule name, Rule scope, Blob type, and Blob subtype** click on the **Next** button.

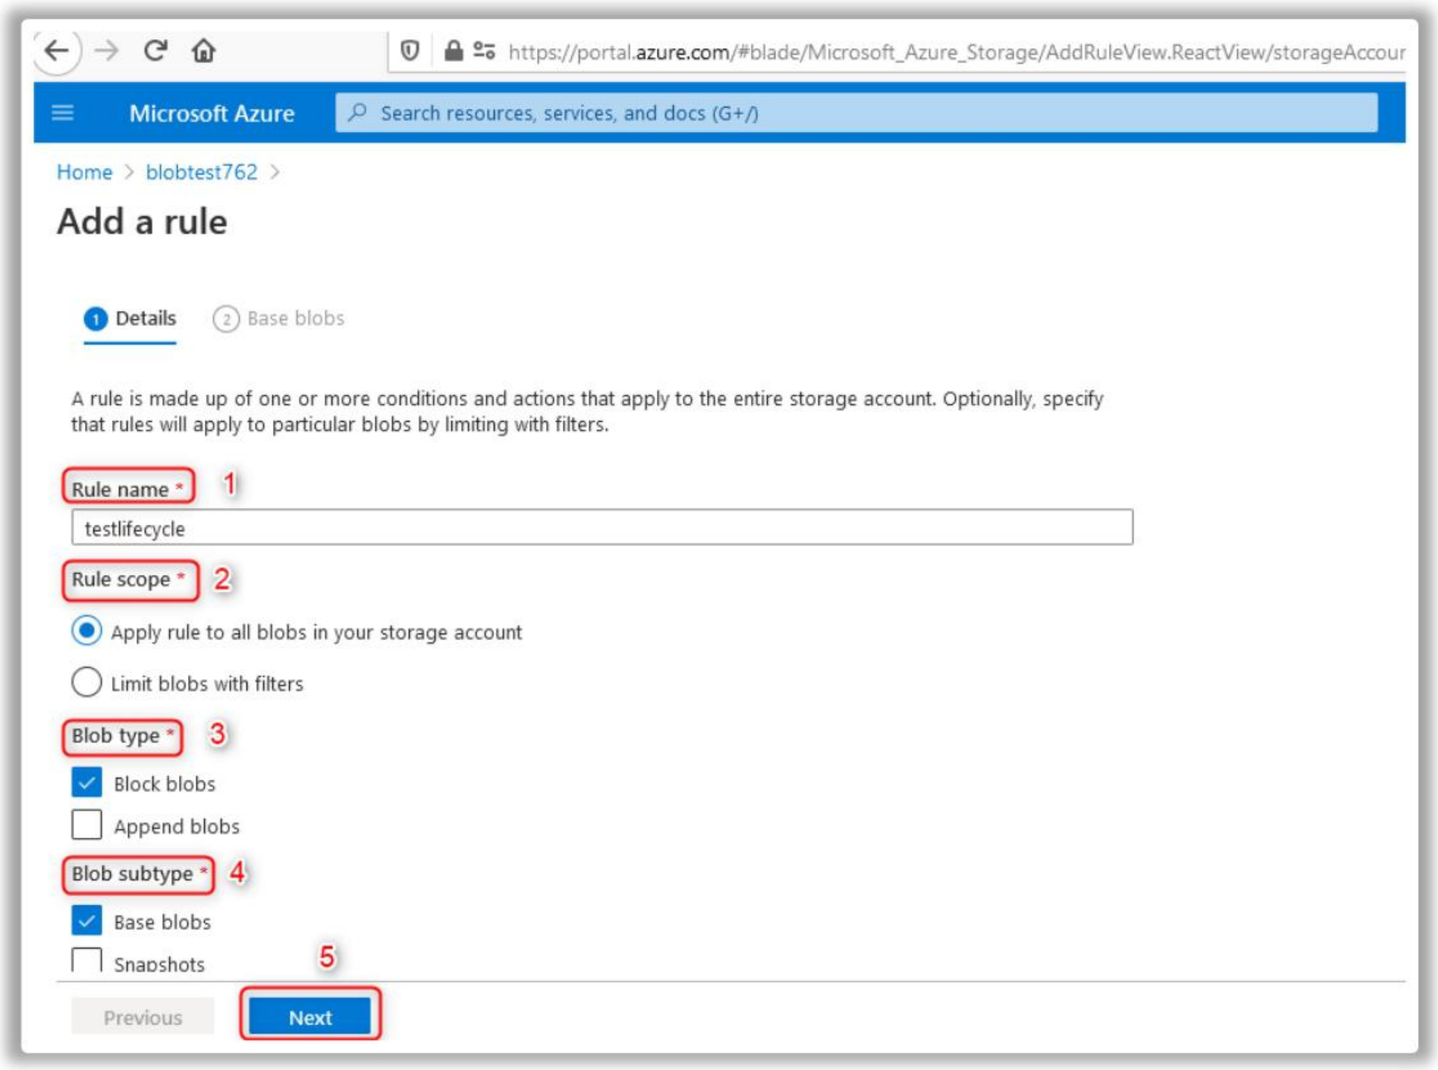


Figure 4.237: Click Next

- Select the dropdown to **Move to cool storage**.

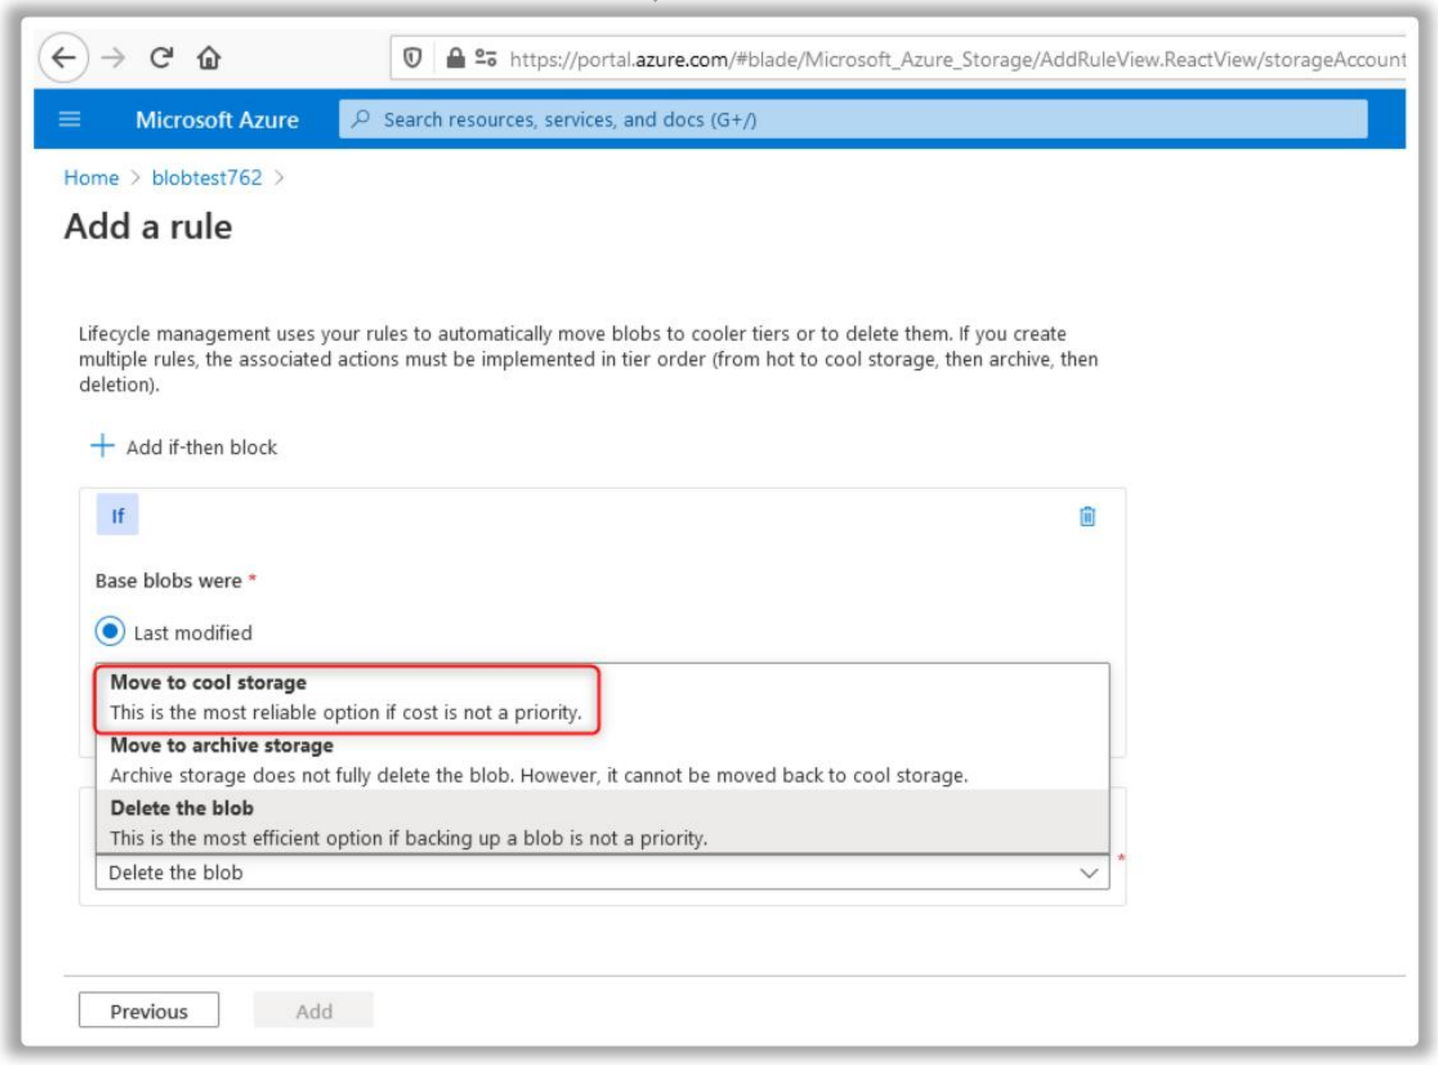


Figure 4.238: Click Move to cool storage

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

- Enter the number of days as per requirement in **Move (than days ago)** field. Thus, it will move to cool storage after a time specified by you.
- Click on **Add**.

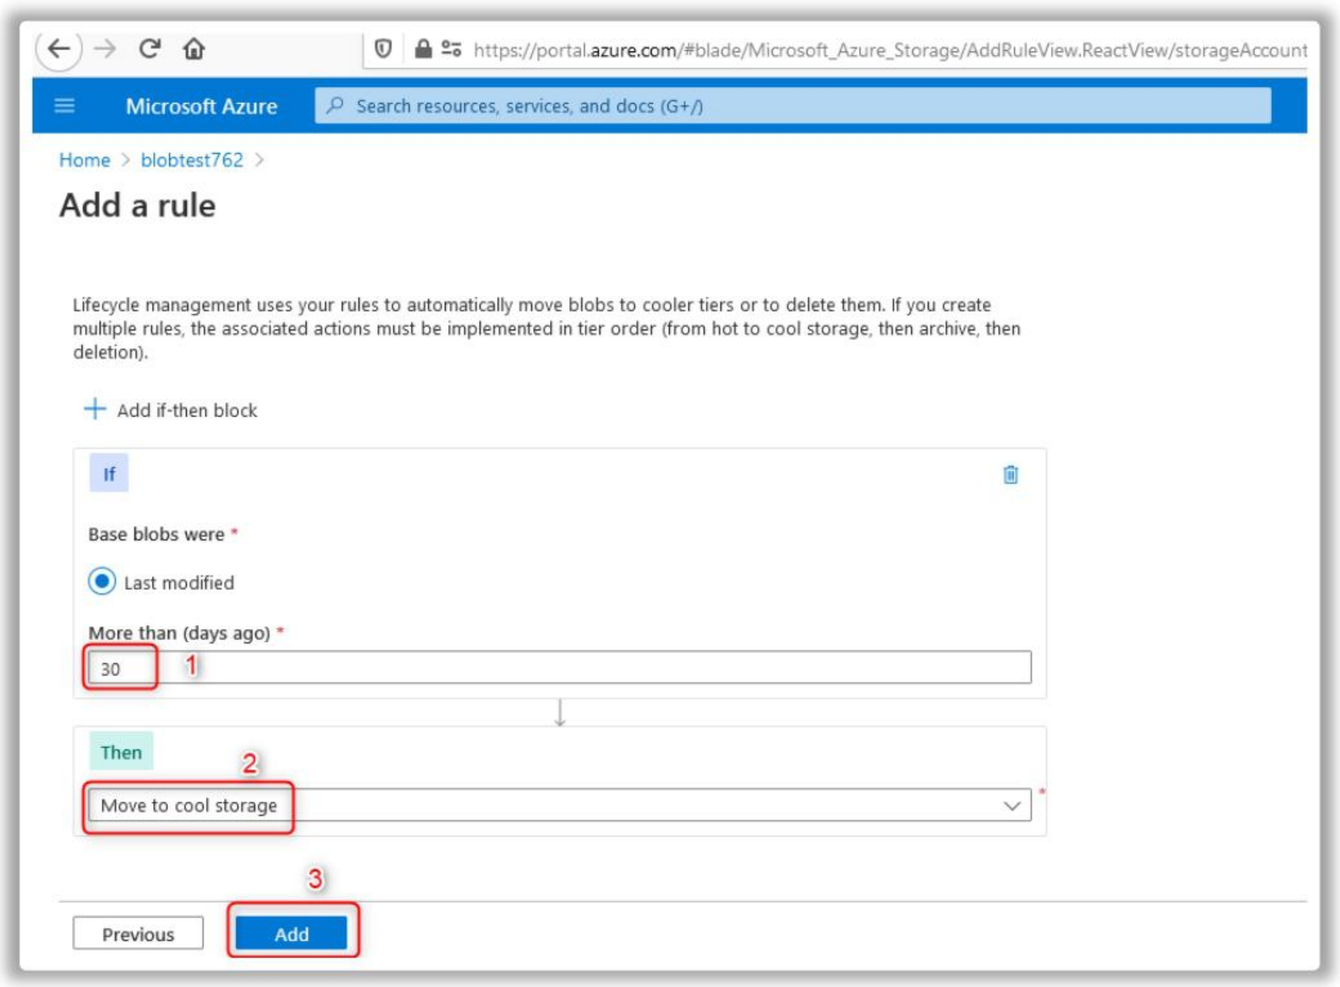


Figure 4.239: Click Add

- Blob Storage lifecycle management was successfully enabled.

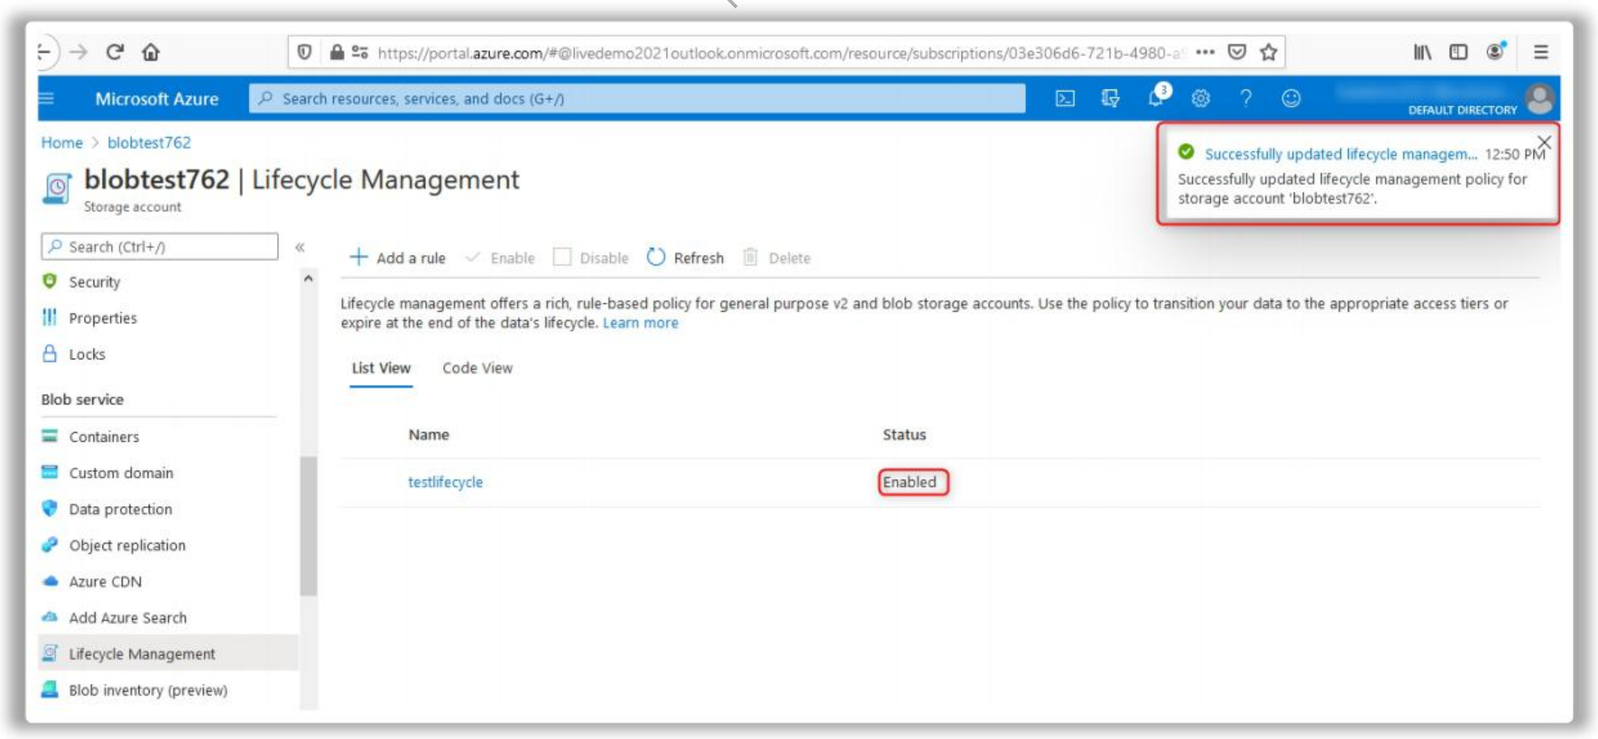
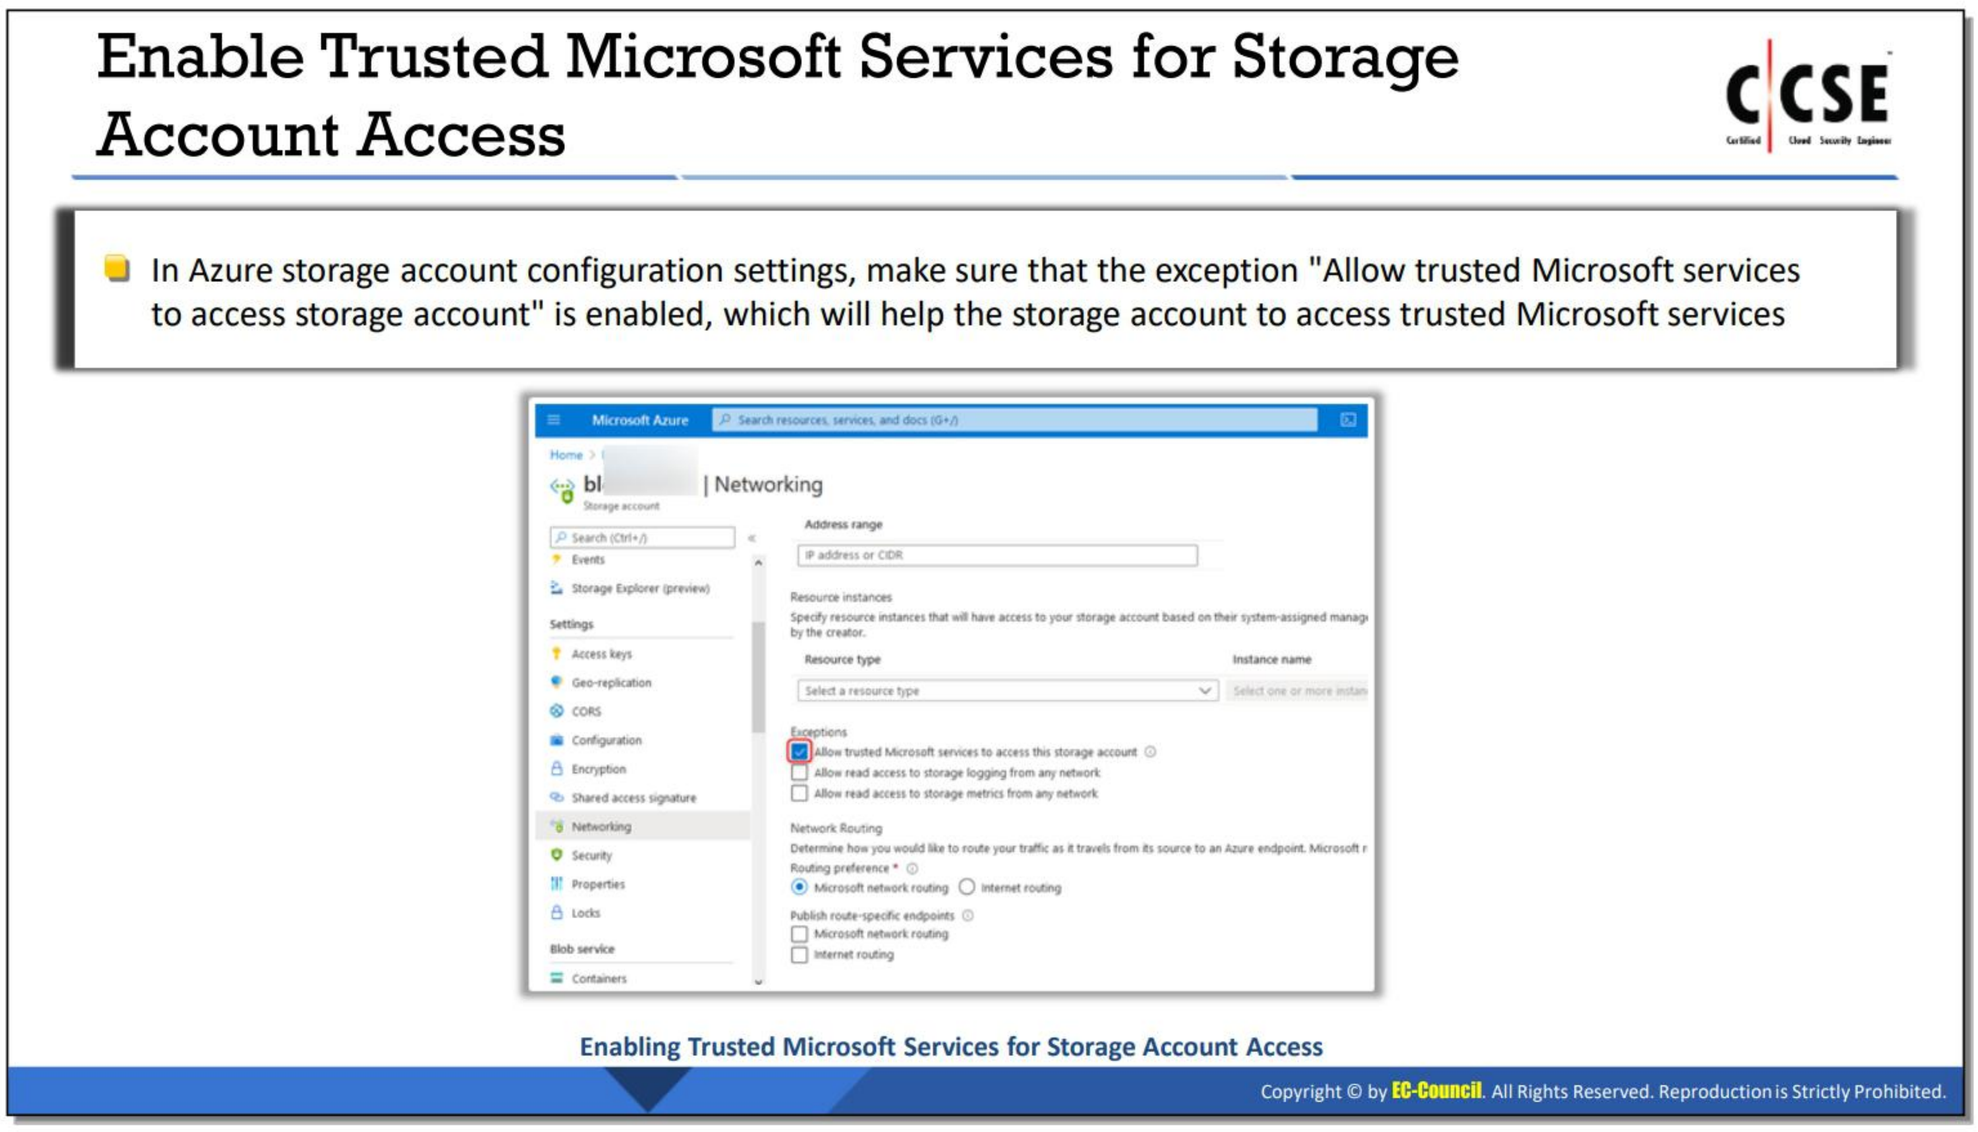


Figure 4.240: Blob Storage Lifecycle is Enabled

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



Enable Trusted Microsoft Services for Storage Account Access

One of the basic security steps in network protection is handling firewall rules properly. Incoming requests to the storage account from any of the Azure services, may be blocked with the usage of a firewall. To make sure that the connections from Azure services are not dropped which fails intended actions, it should allow trusted Microsoft services. Virtual network configuration and firewall rules setup should be applied carefully, as the misconfiguration in networks leads to malfunctions or risks.

In Azure storage account configuration settings, make sure that the exception "Allow trusted Microsoft services to access storage account" is enabled. If the firewall rules are enabled for the storage account, then incoming requests for the data will be blocked including Azure services. If the exception "Allow trusted Microsoft services to access storage account" is enabled, then the access to the storage account resources will be granted to the Azure services.

Source: https://azure.microsoft.com

To Enable Trusted Microsoft Services for Storage Account Access

- Sign-in to the Azure portal.
- Select **Storage accounts** under **Azure services.**

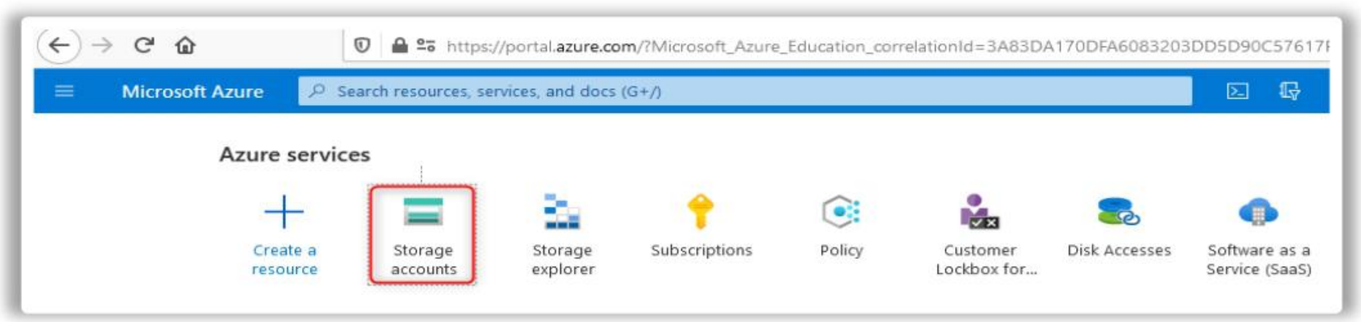


Figure 4.241: Click Storage Accounts

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

- A list of storage accounts will be displayed.
- Select the storage for which you want to enable trusted Microsoft services for storage account access.

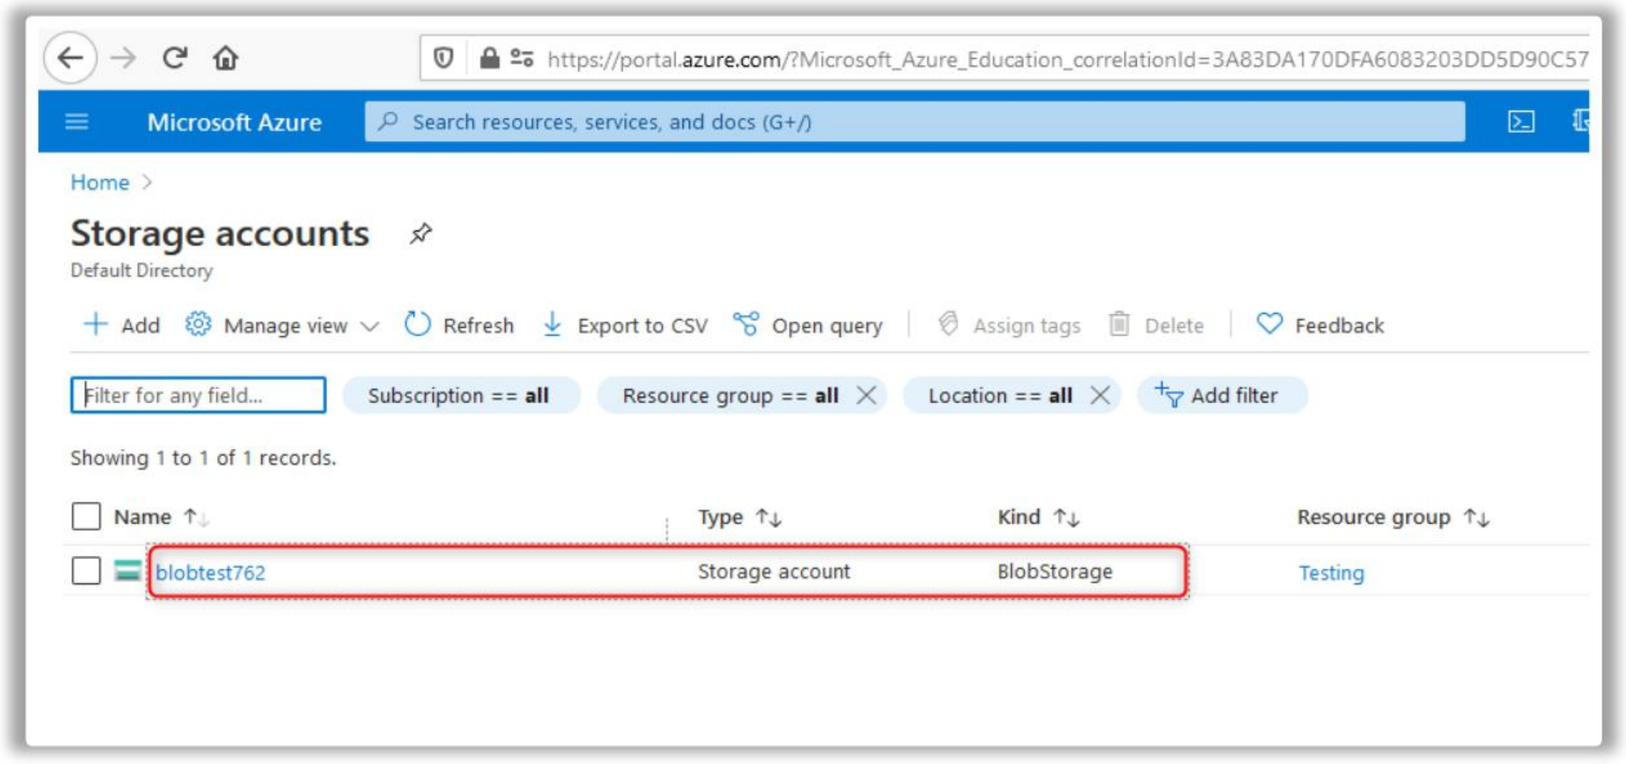


Figure 4.242: Choose Storage

- Locate **Networking** under **Settings**.
- **Allow access from** should be in **Selected networks**.
- Select other details as per requirement.

hide01.ir

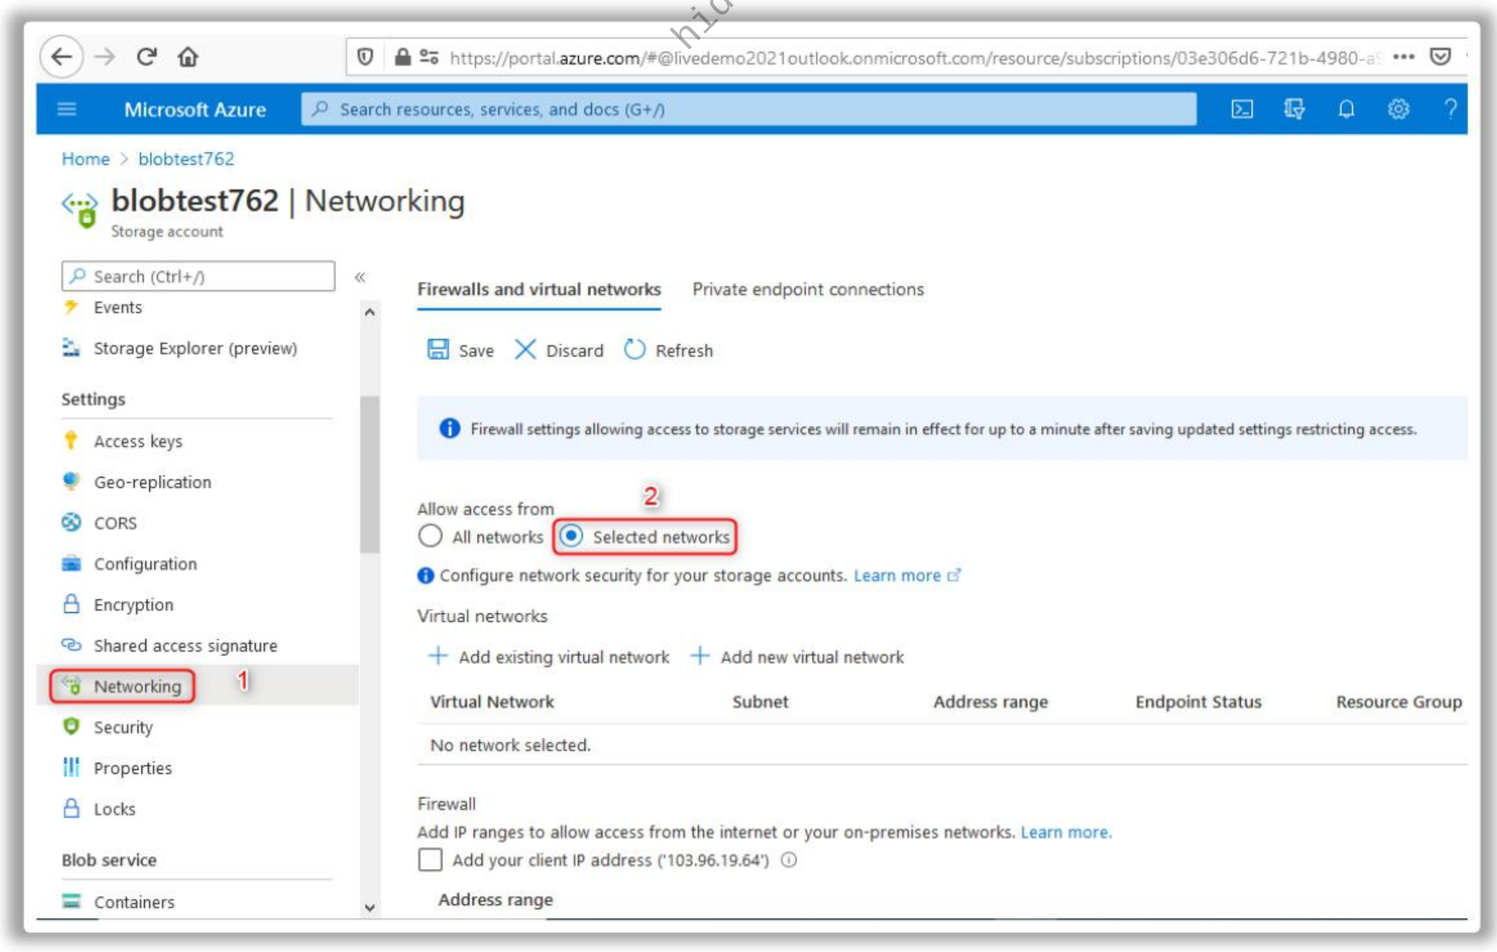


Figure 4.243: Select Details

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

- Scroll down to **Exceptions** and select the check box of **Allow trusted Microsoft services to access this storage account.**

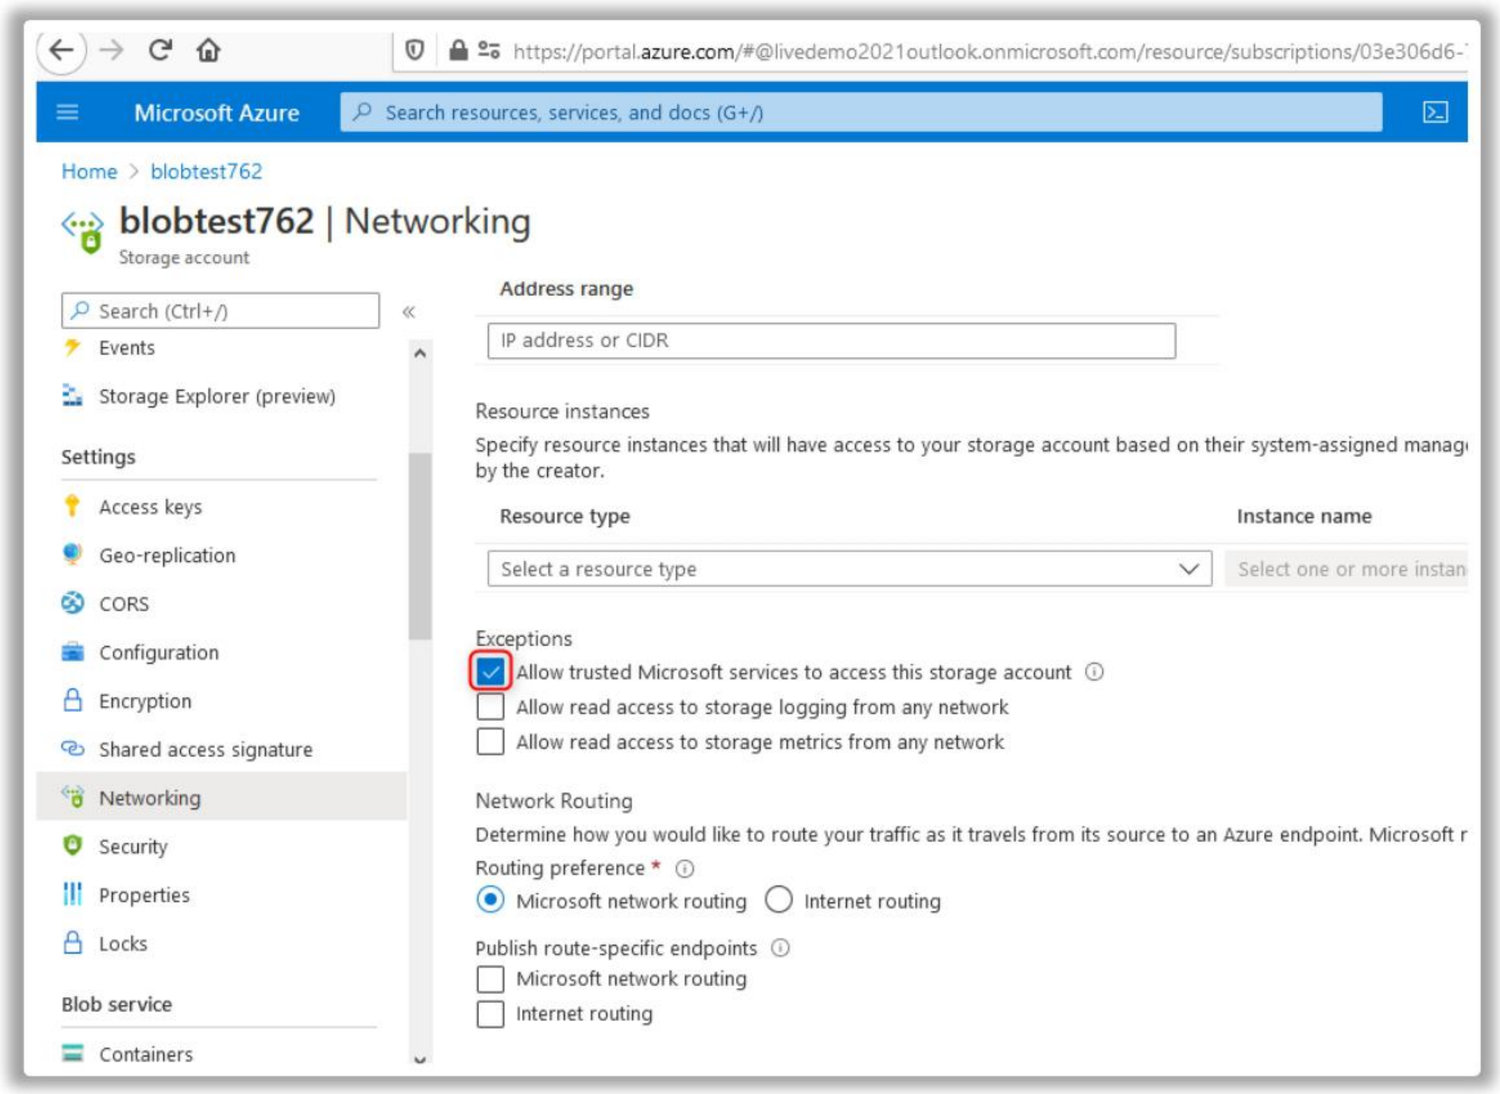


Figure 4.244: choose Exceptions

- Click on **Save** to update the changes.
- Changes will be saved successfully.

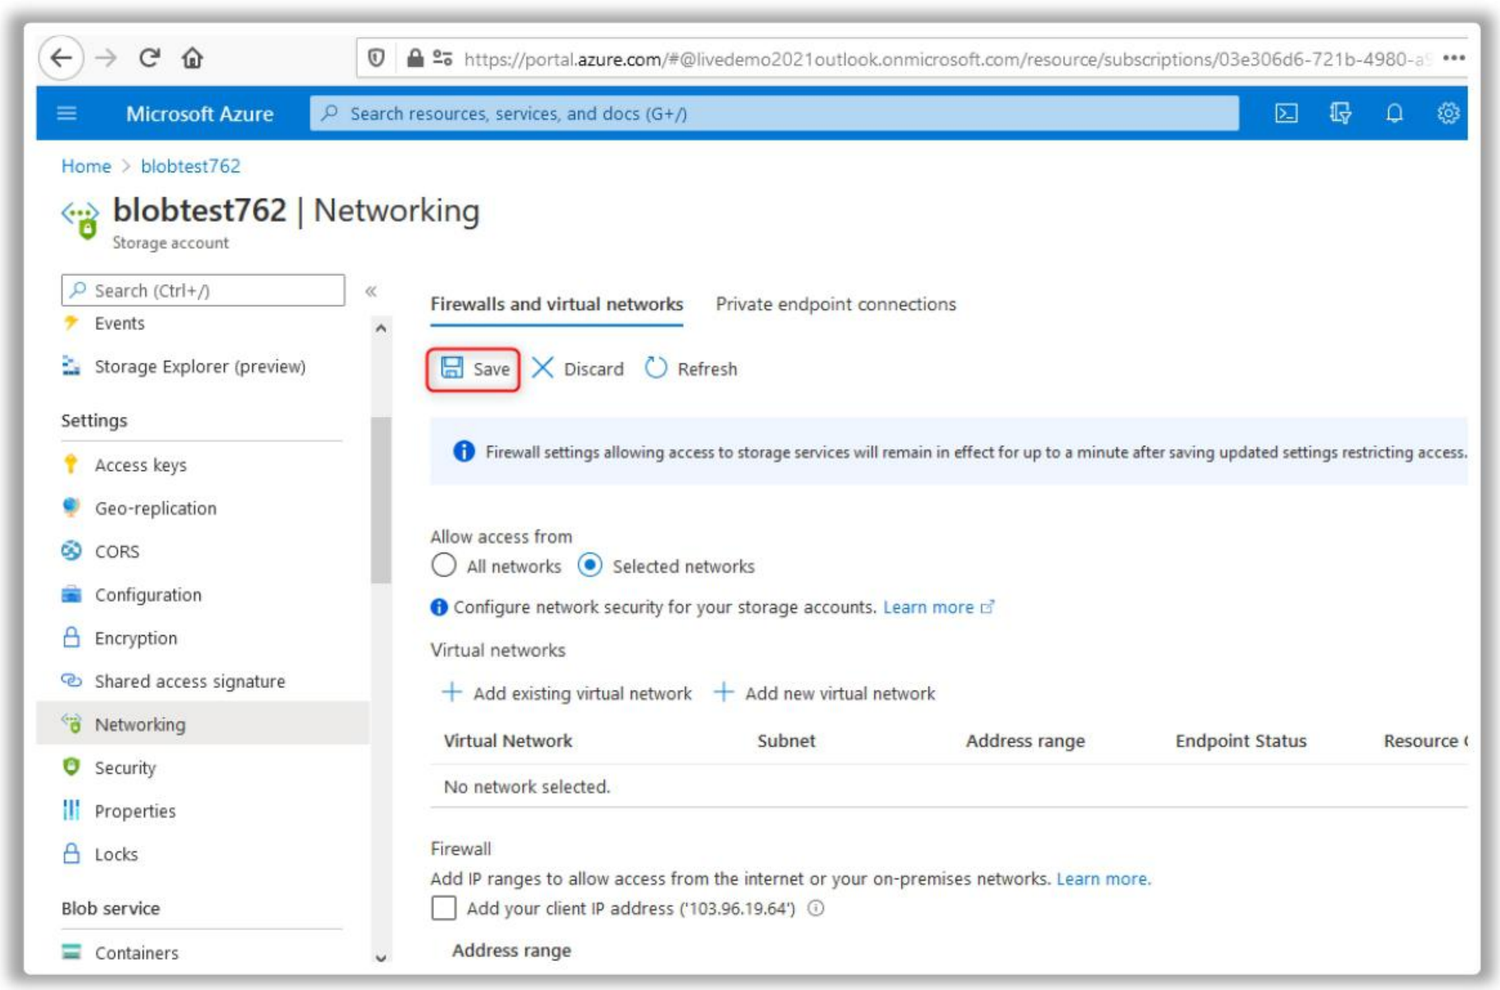
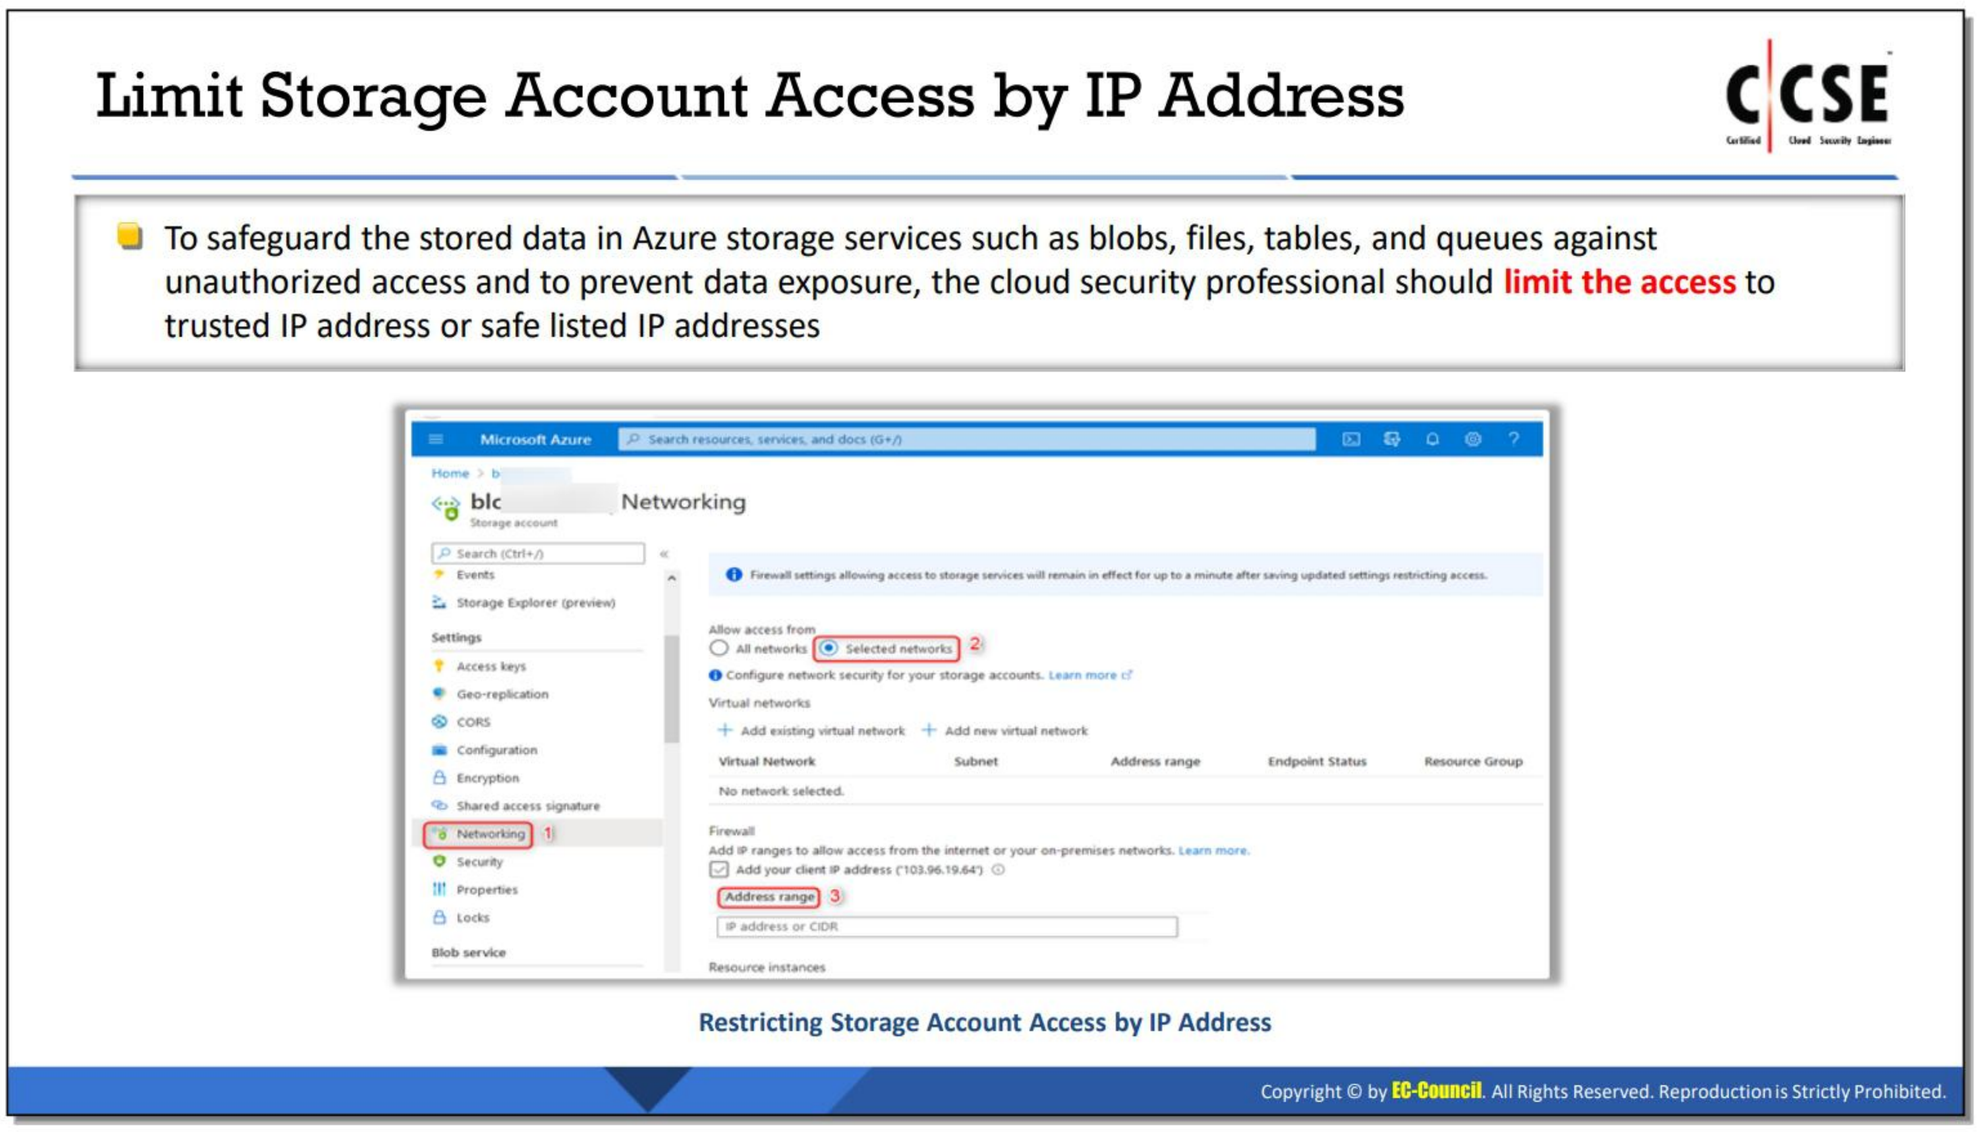


Figure 4.245: Click Save

hide01.ir

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



Limit Storage Account Access by IP Address

To safeguard the stored data in Azure storage services such as blobs, files, tables, and queues against unauthorized access and to prevent data exposure, the cloud security professional should limit the access to trusted IP address or safelisted IP addresses. Ensure the security of the storage accounts by reviewing the IP addresses that are added to the safe list.

Source: https://azure.microsoft.com

To Limit Storage Account Access by IP Address

- Sign-in to the Azure portal.
- Select **Storage accounts** under **Azure services.**

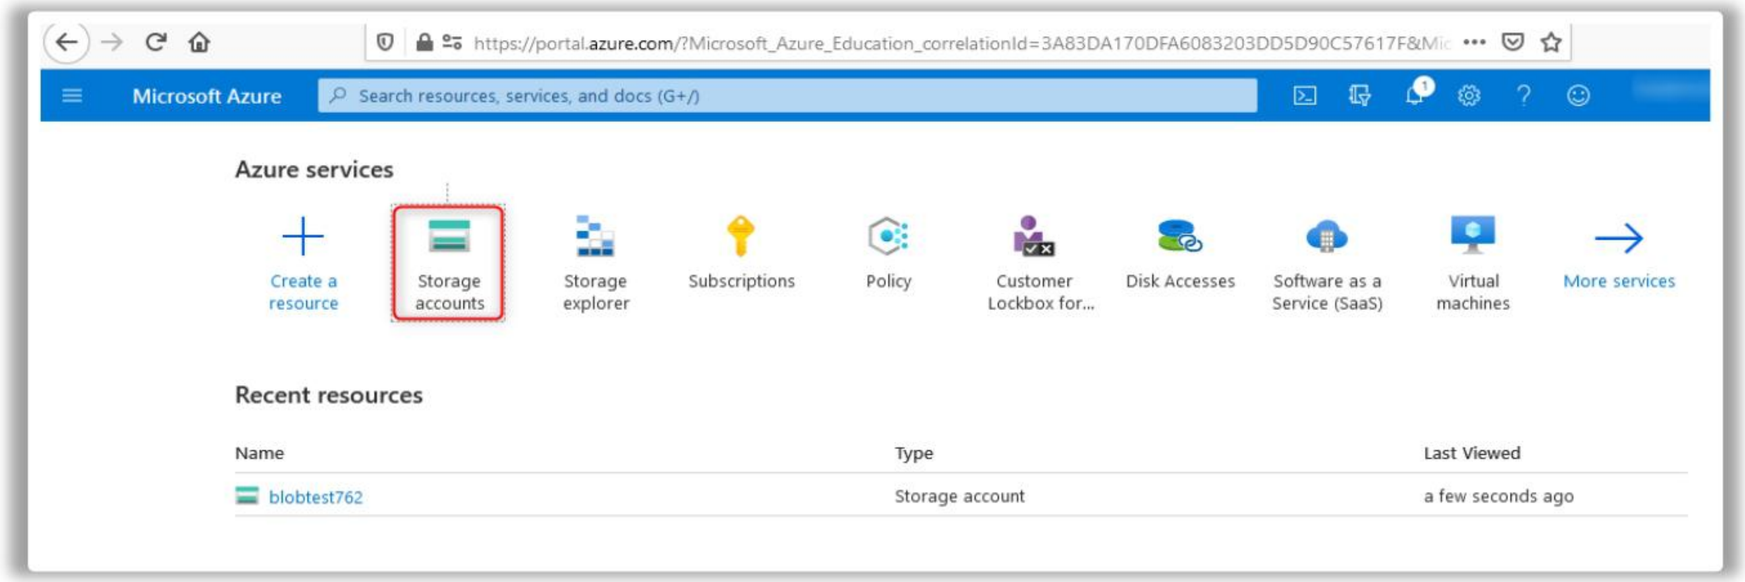


Figure 4.246: Select Storage Accounts

- A list of storage accounts will be displayed.

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

- Select the storage for which you want to restrict access by IP addresses.

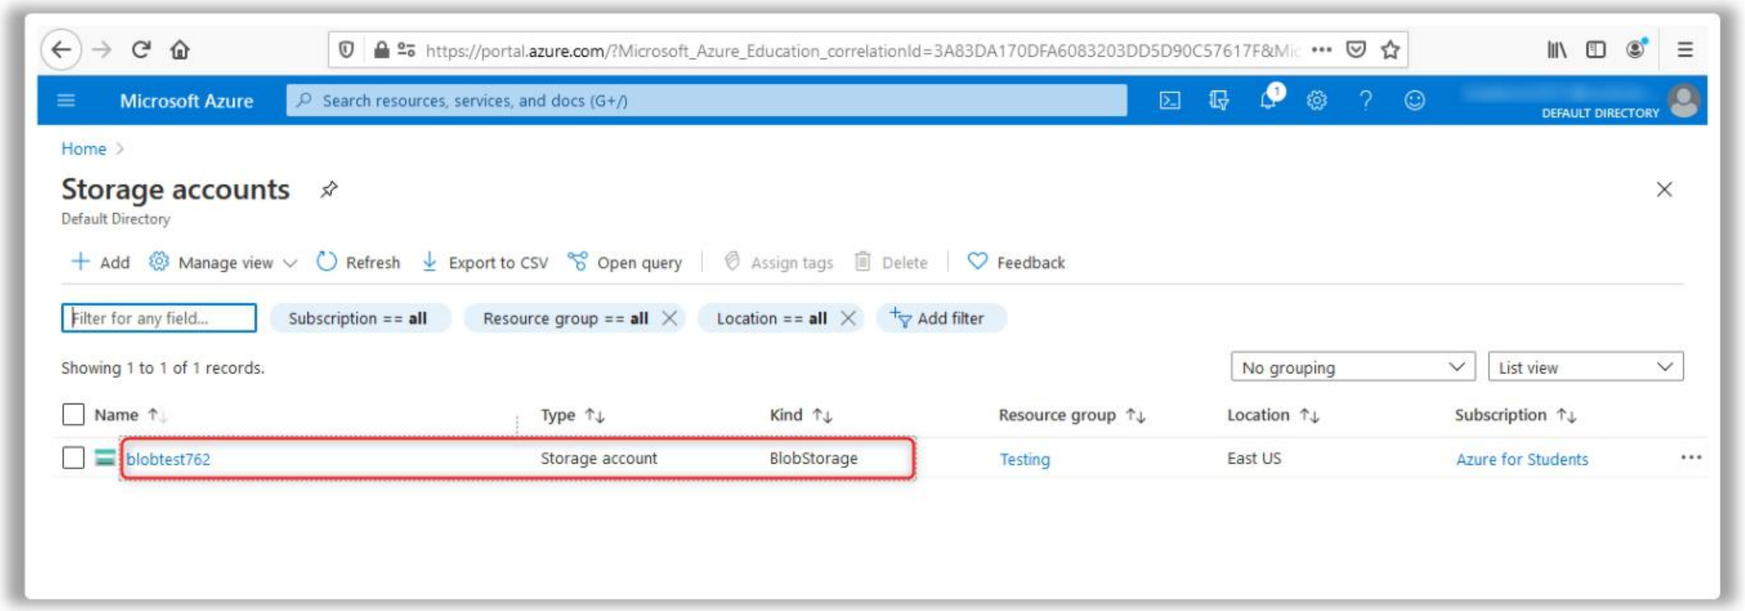
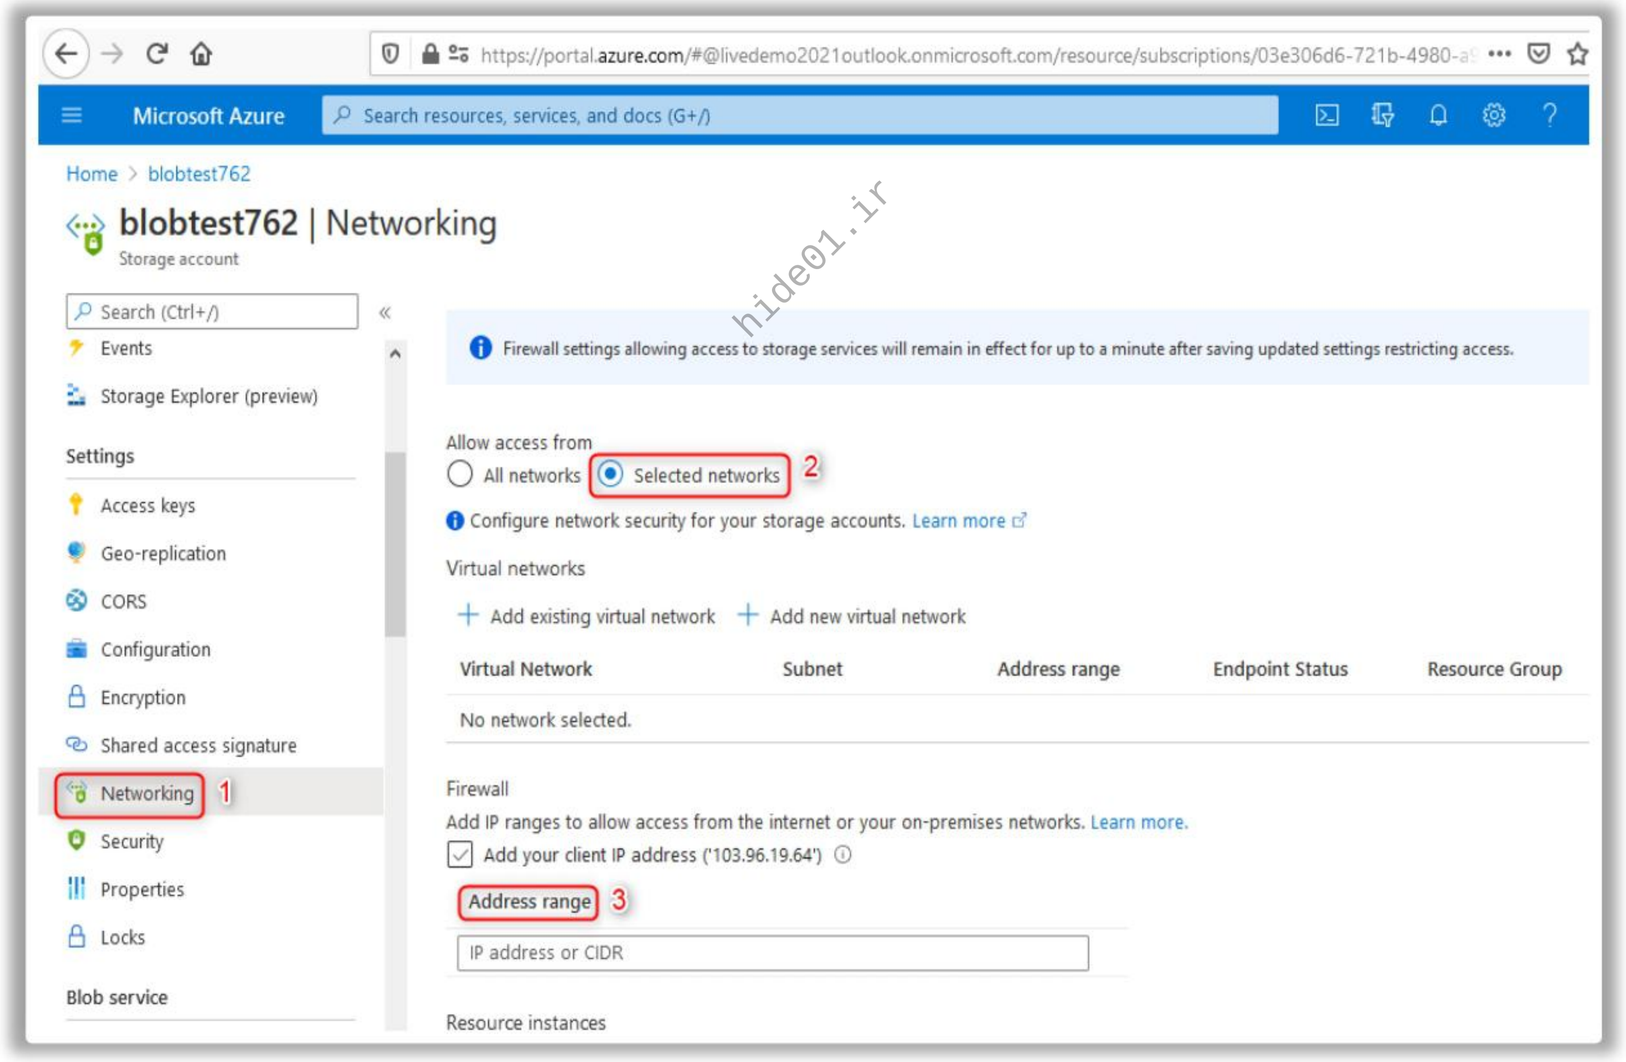


Figure 4.247: Choose Storage

- Locate **Networking** under **Settings**.
- **Allow access from** should be in **Selected networks**.
- Under, **Address range** make sure to enter only IP addresses that are required.



hide01.ir

Figure 4.248: Enter IP Address

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

- Click on **Save** to successfully update the restricted IP addresses.

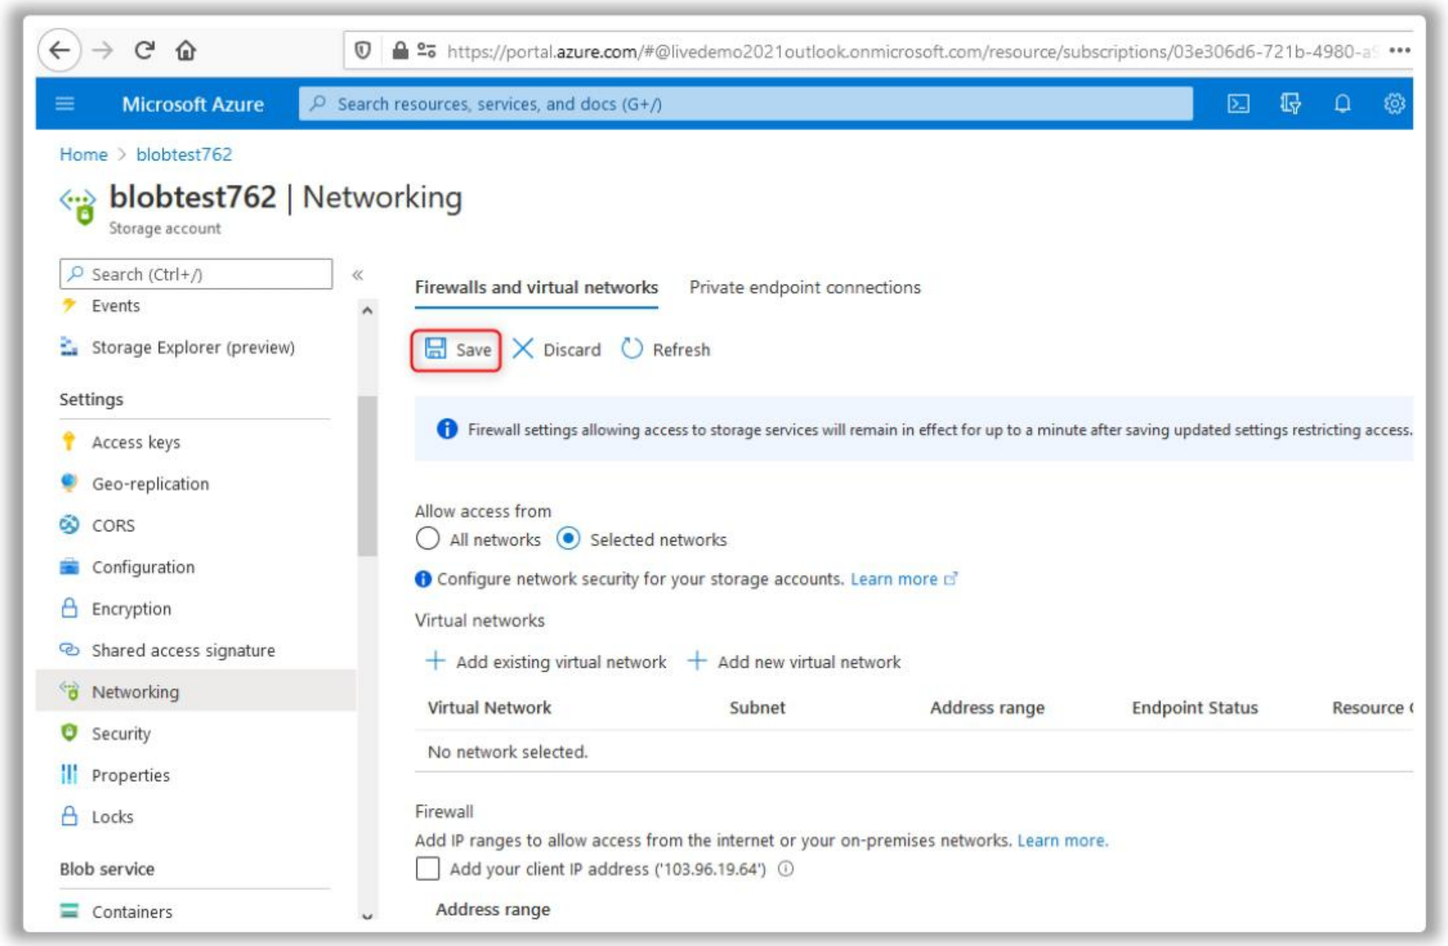
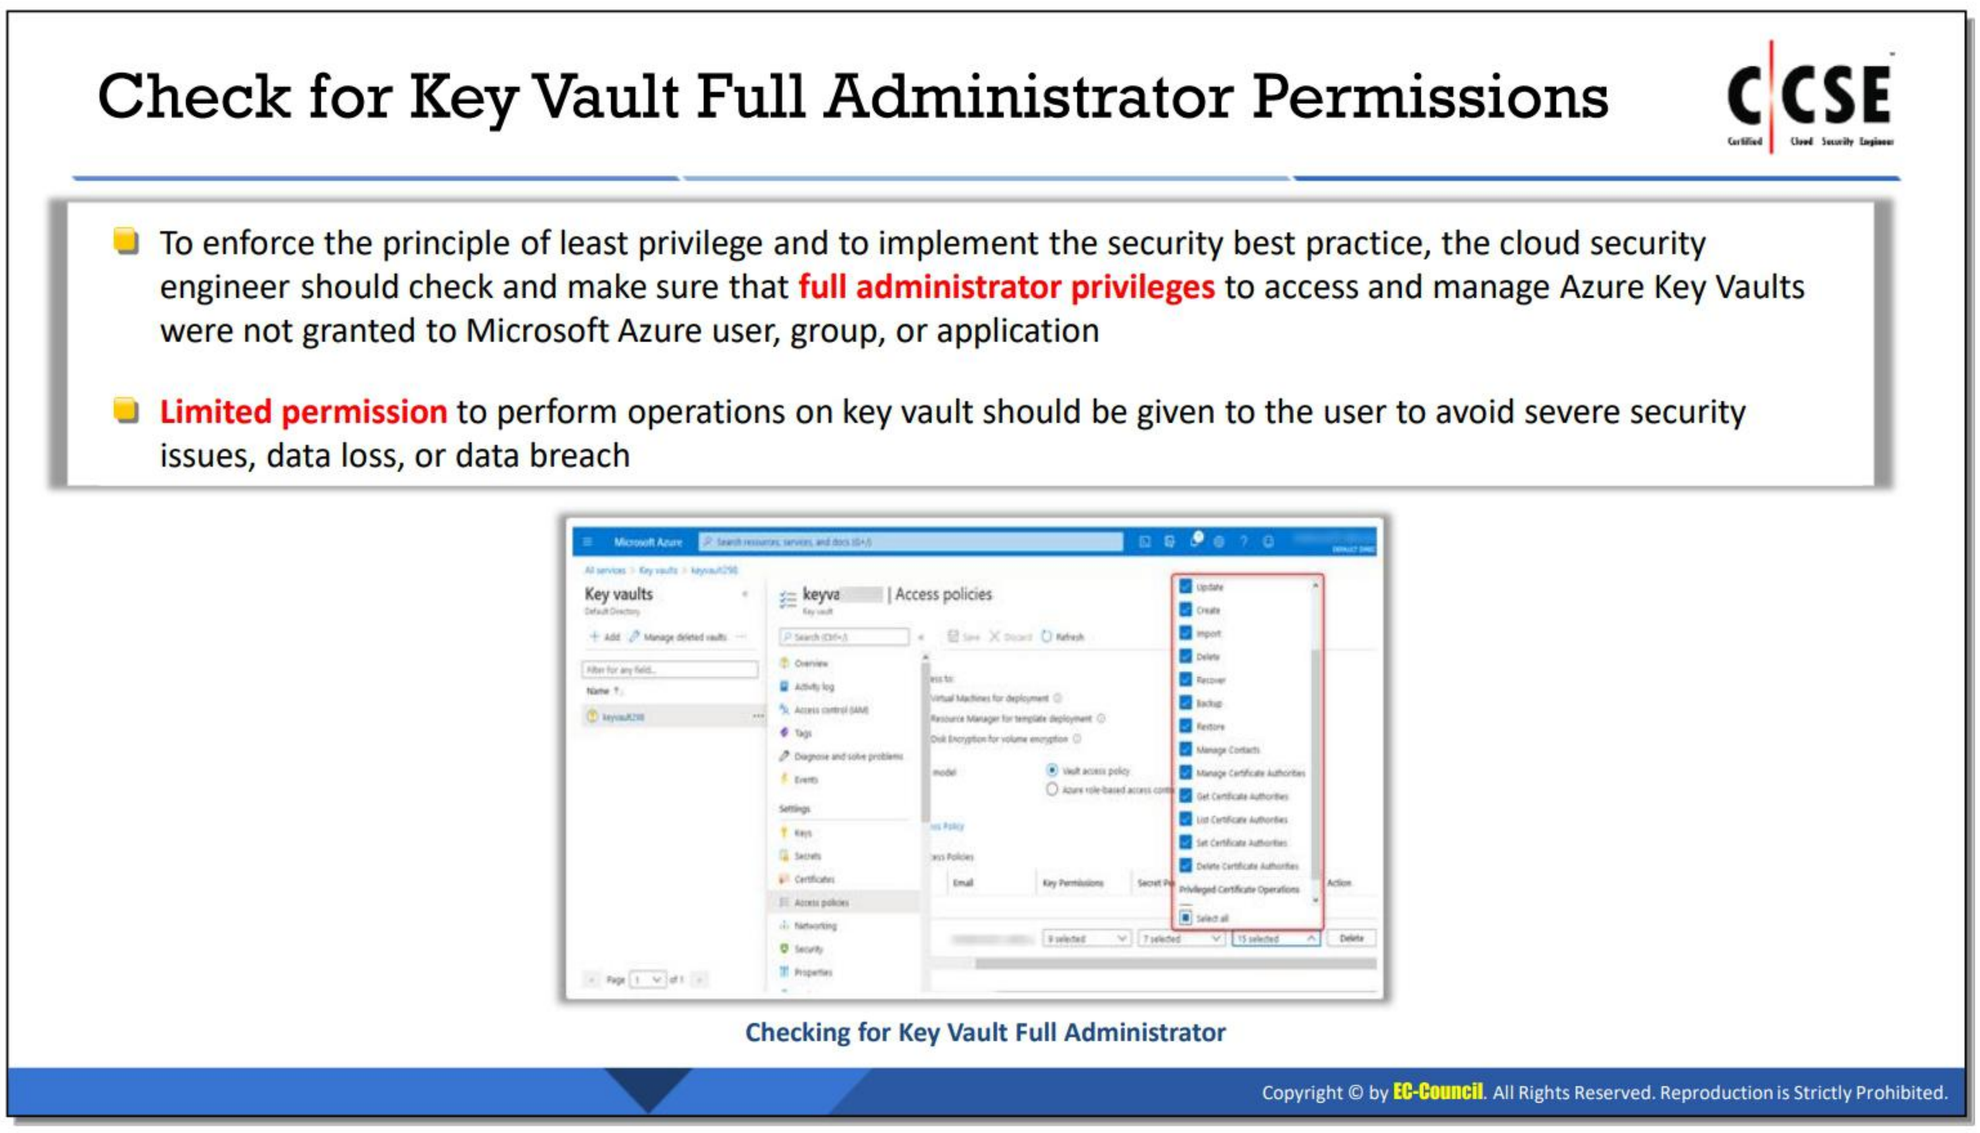


Figure 4.249: Select Save

hide01.ir

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



Check for Key Vault Full Administrator Permissions

To enforce the principle of least privilege and to implement the security best-practice, the cloud security professional should ensure that none of the Microsoft Azure user, group, or application have full administrator privileges for accessing and managing Azure Key Vaults.

Since Azure Key Vaults stores sensitive and business-critical data, it is the responsibility of the cloud security professionals to secure the key vault and the stored data. It is recommended to grant limited access to the key vault based on the requirements. As the Azure Key Vault policy is applicable at the vault level, if the Azure user is given permission to create, delete, or update keys, then the user can perform operations on the entire key vault. Limited permission to perform particular operations on key vault should be given to the user to avoid severe security issues, data loss, or data breach.

Source: https://azure.microsoft.com

To Check for Key Vault Full Administrator Permissions

- Sign-in to Azure Portal.
- Go to **More services.**

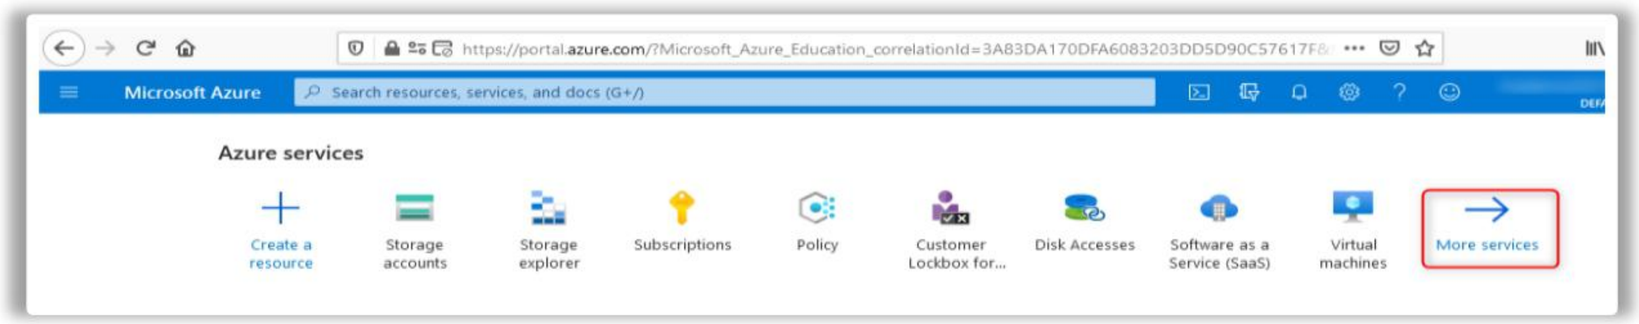


Figure 4.250: Select More service

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

- Search for the **key vault** in the search box beside all services.
- Click on **key vault**.

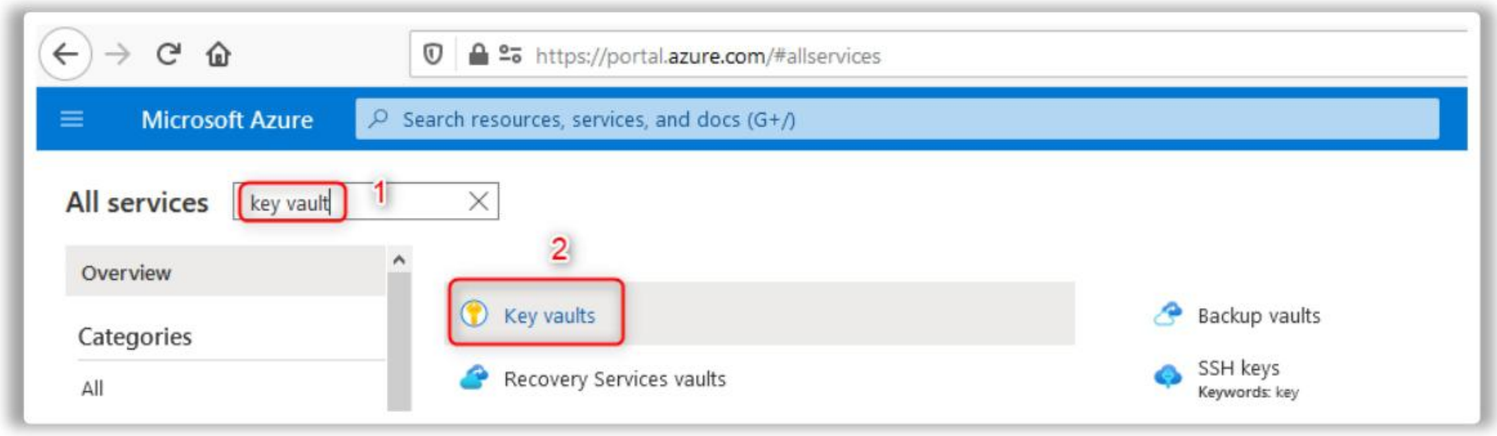


Figure 4.251: Select Key Vault

- A list of all available key vaults will be displayed.
- Select the key vault that you want to check permissions.

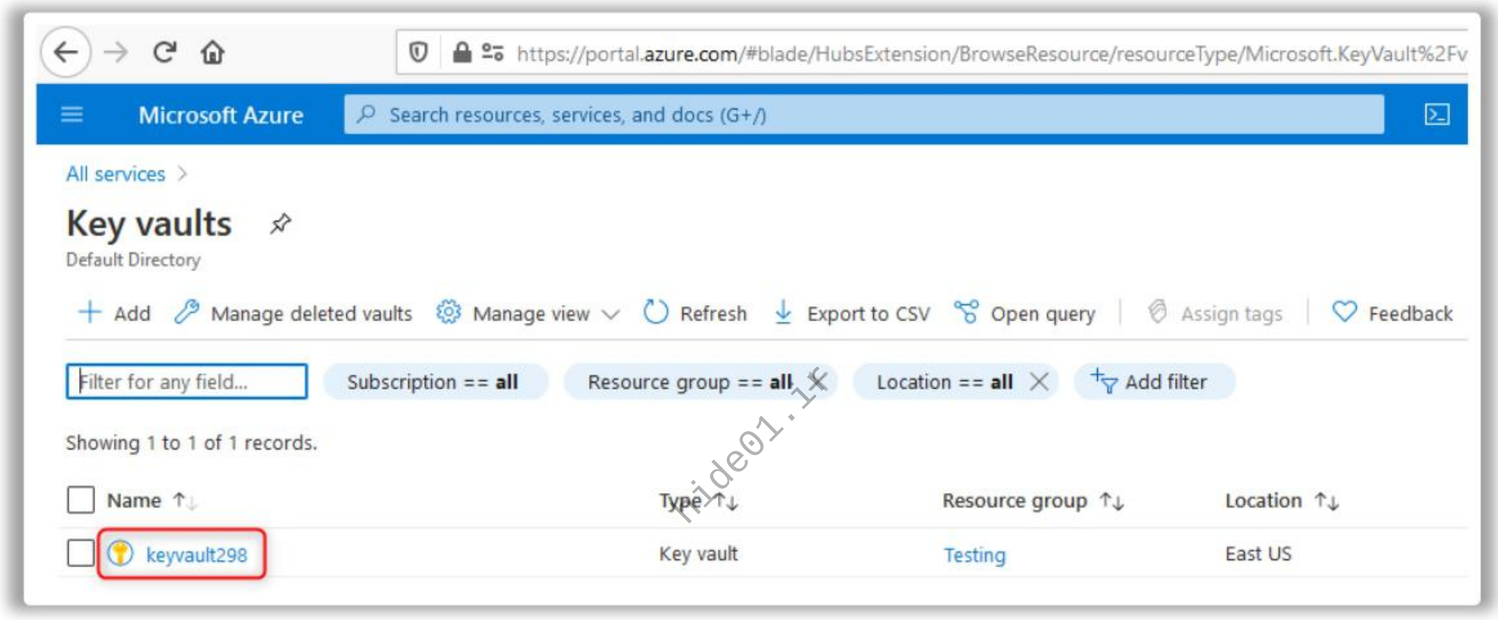


Figure 4.252: Choose Key Vault

- Locate **Settings** and go to **Access policies**.
- Move to **Current Access Policies**.

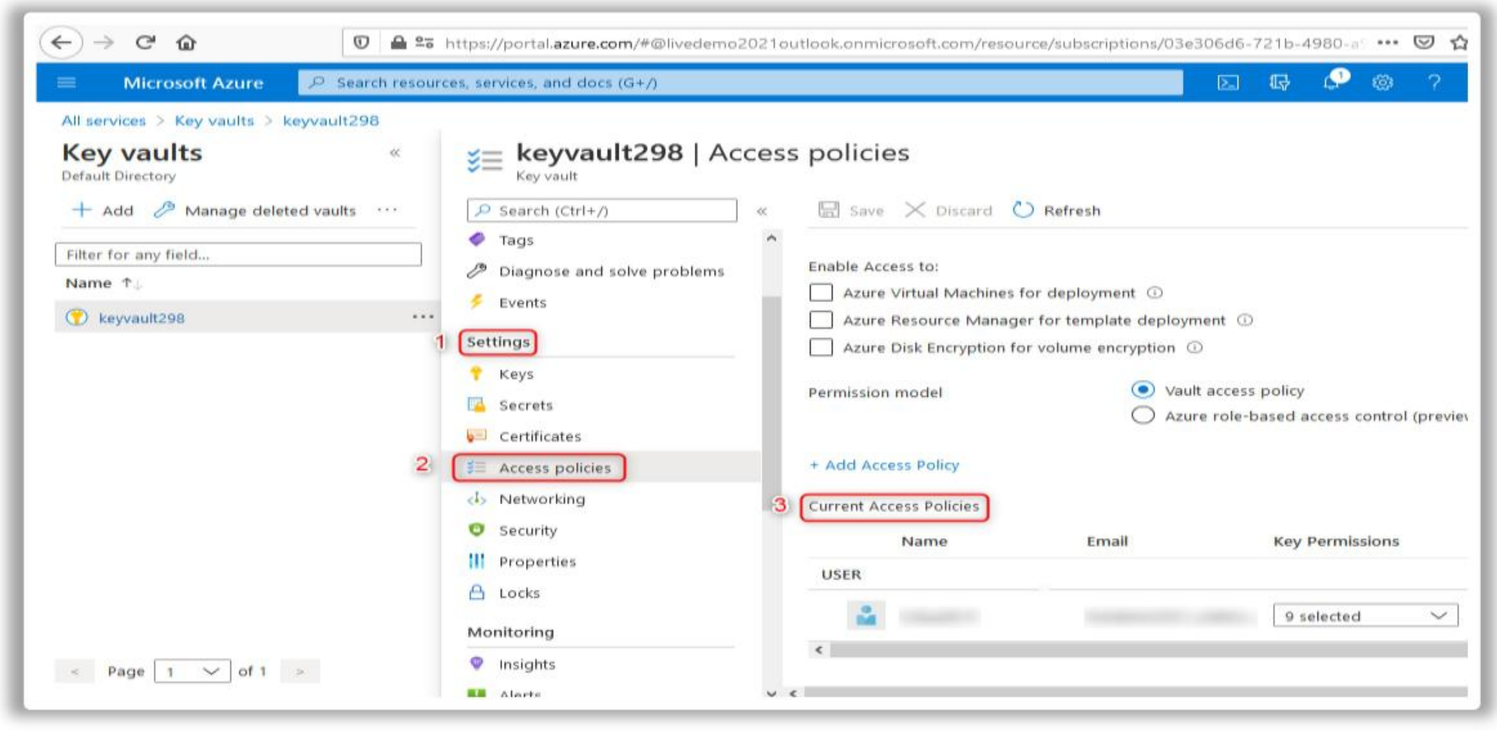


Figure 4.253: Select Current Access Policies

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

- Check all **Key Permissions, Secret Permissions,** and **Certificate Permissions.**
- Open the dropdown list and verify the permissions.

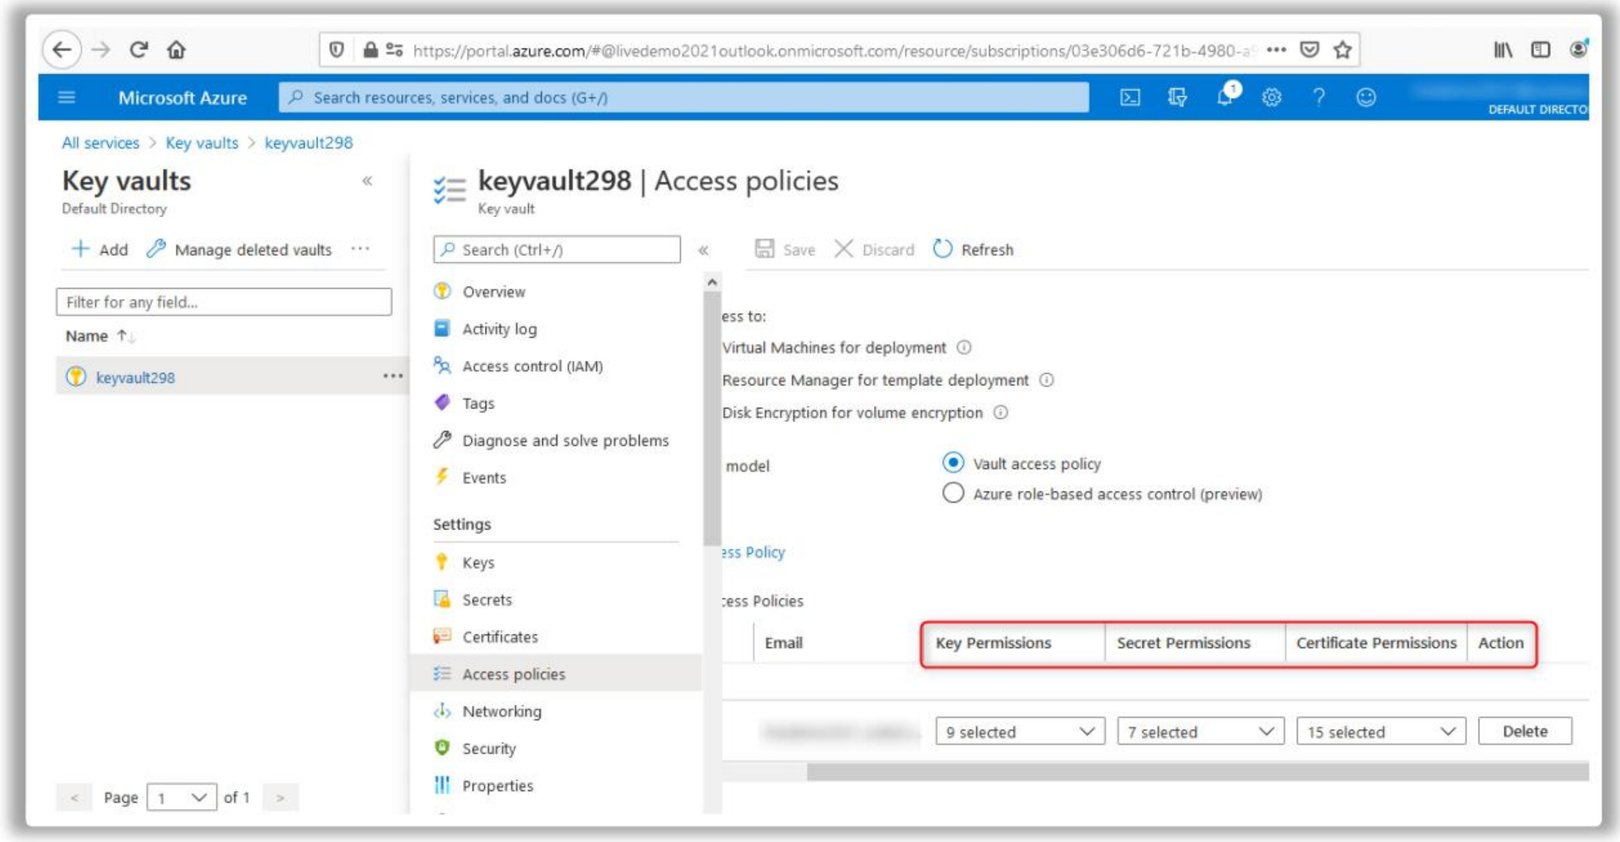


Figure 4.256: Check Permissions

- If all the permissions are granted in key permissions, secret permissions, and certificate permissions then the user, group, or application of validated access policy will be having access to that key vault.
- Enable the permissions as per the requirement and make possible restrictions.

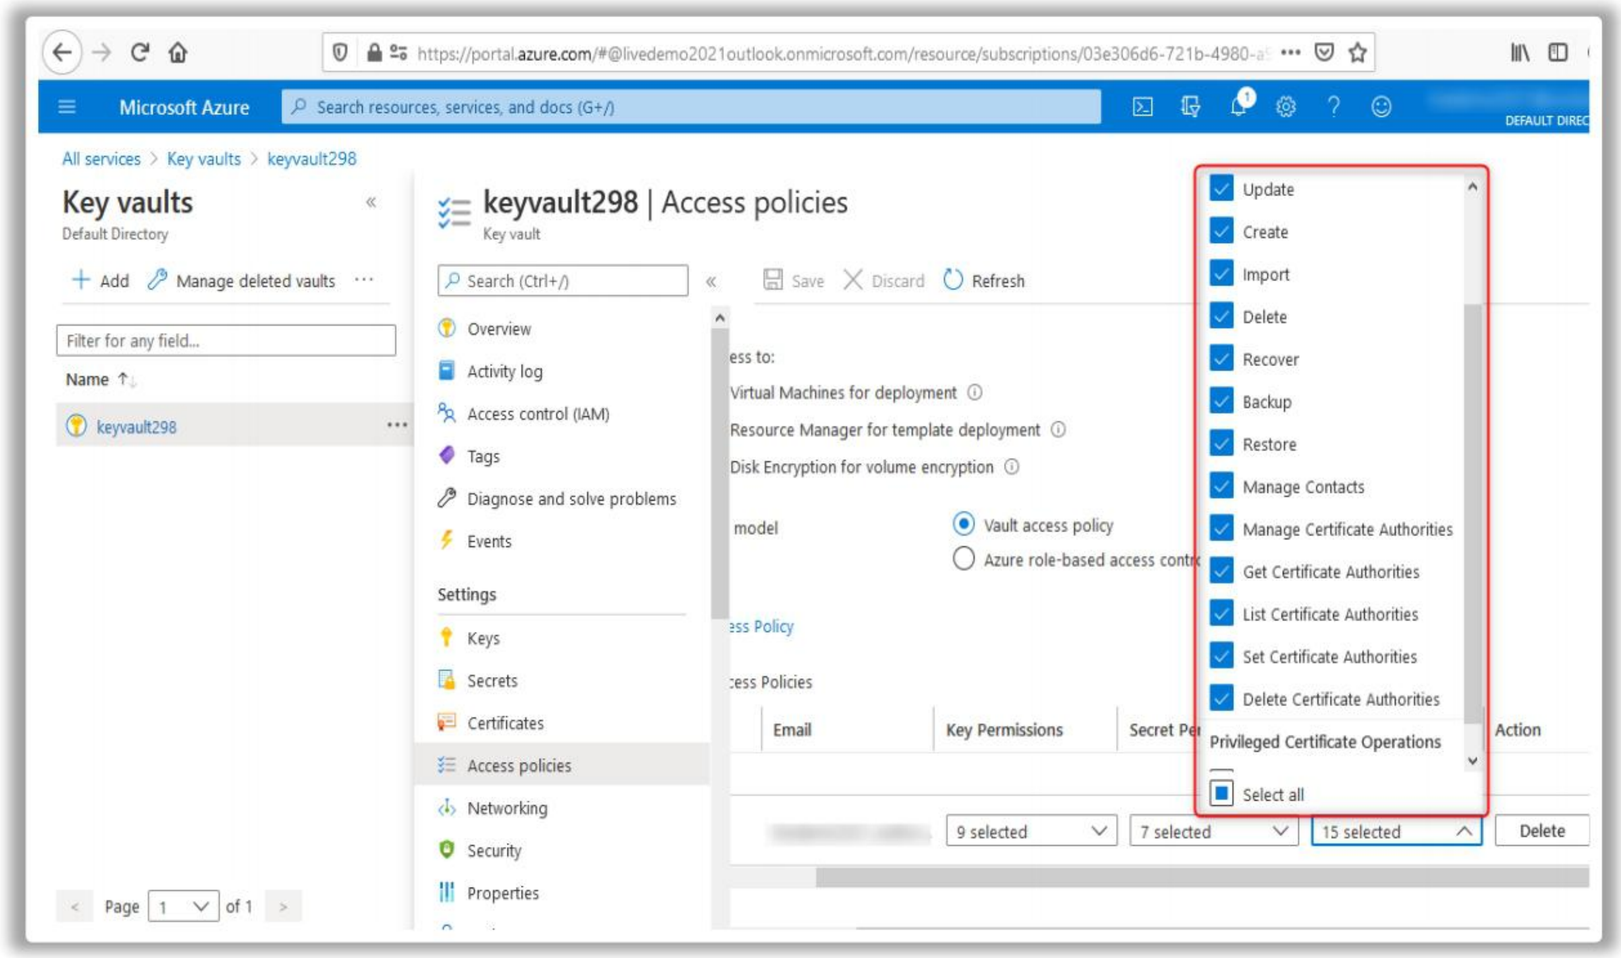


Figure 4.257: Enable Permissions

- Click on **Save** after the modifications.

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

- The permissions will be modified successfully.

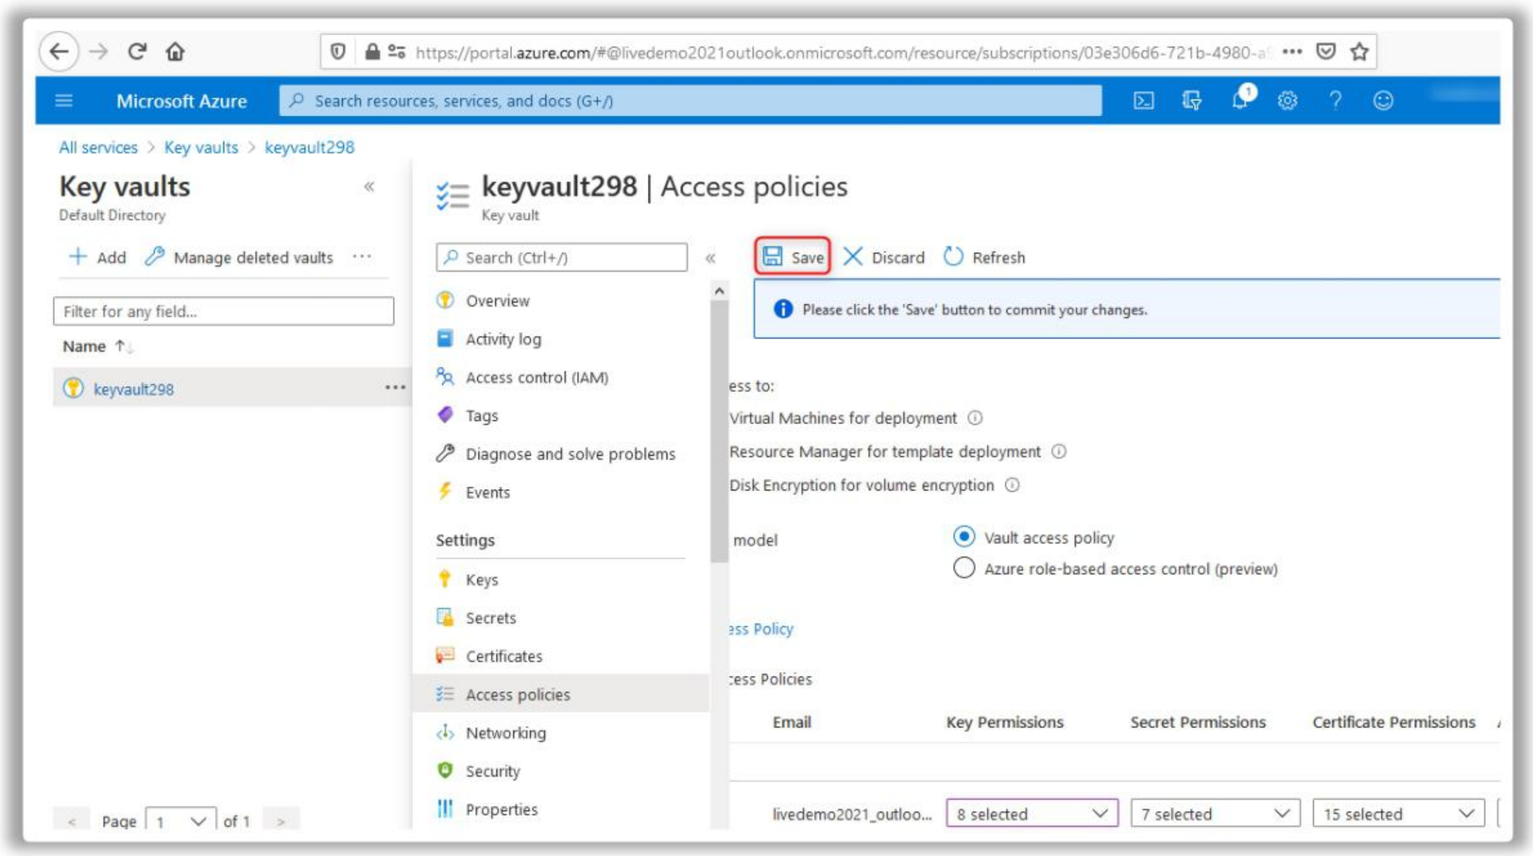
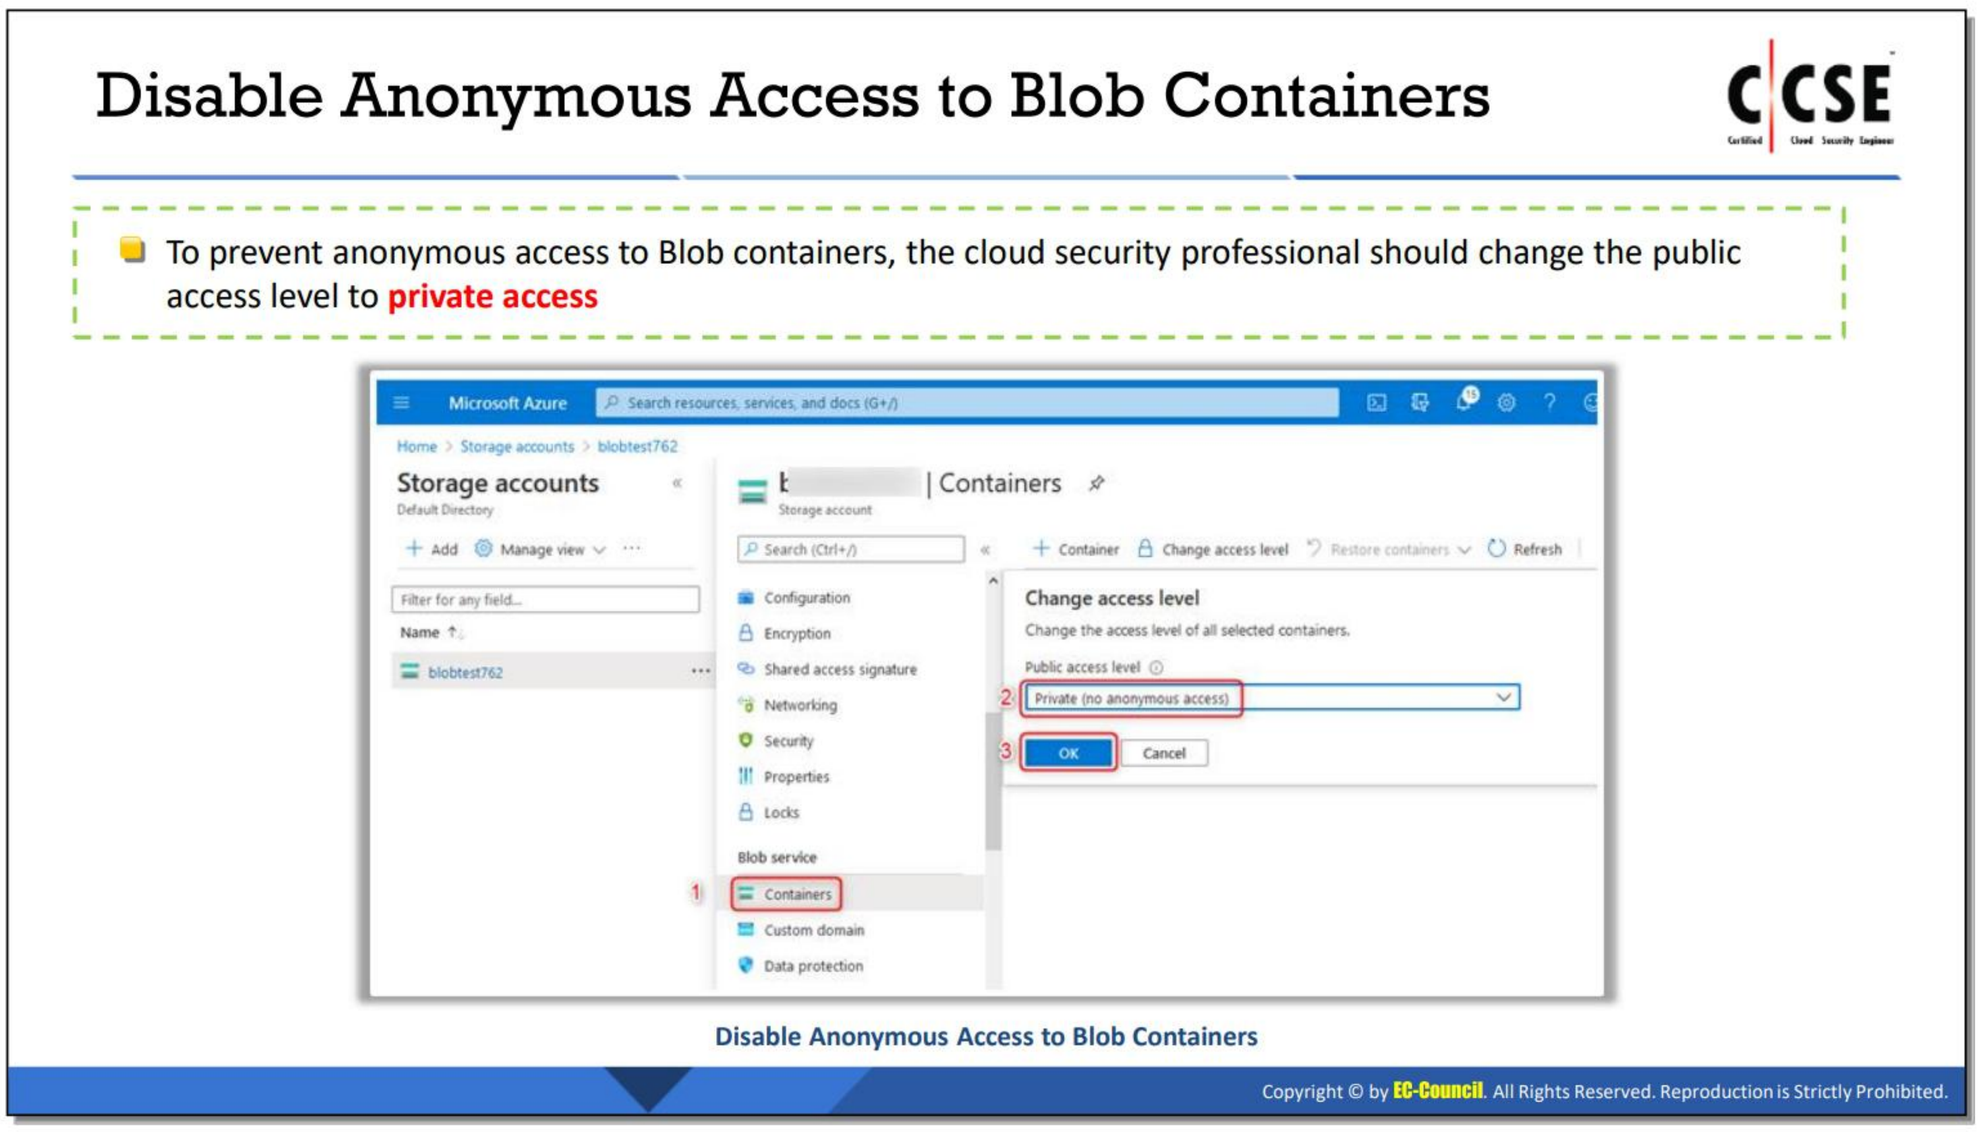


Figure 4.258: Click Save

hide01.ir

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



Disable Anonymous Access to Blob Containers

To prevent anonymous access to Blob containers, the cloud security administrator should change the public access level to private access. With public access level configuration settings, the anonymous user can use constructor to access the Blob containers and access the containers without using any credentials such as SAS. Hence, when the Blob containers access is changed from public to private level access, the anonymous user will not be allowed to access the Blob containers.

A container can be changed to public or private if the storage account level access is public. If the storage account level is set to private, then all the containers become private by default and cannot be changed to public access. Set the containers with private access unless and until there is a requirement to change the settings. It is recommended to allow access using a shared access signature token that helps in controlled access.

Source: https://azure.microsoft.com

To Disable Anonymous Access to Blob Containers

- Sign-in to the Azure portal.

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

- Select **Storage accounts** under **Azure services.**

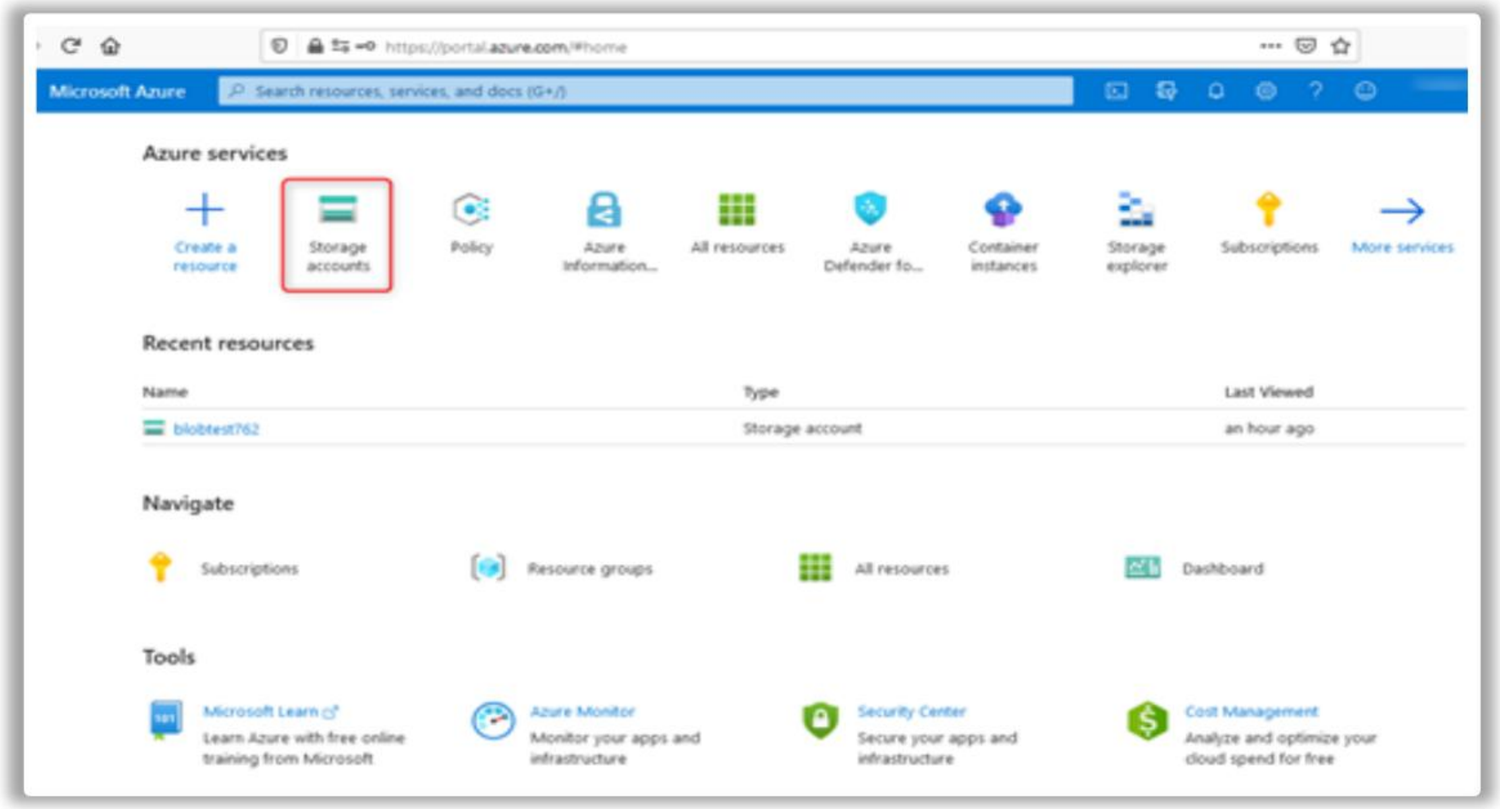


Figure 4.259: Select Storage Accounts

- A list of storage accounts will be displayed.
- Select the storage account that you want to disable public access to Blob containers.

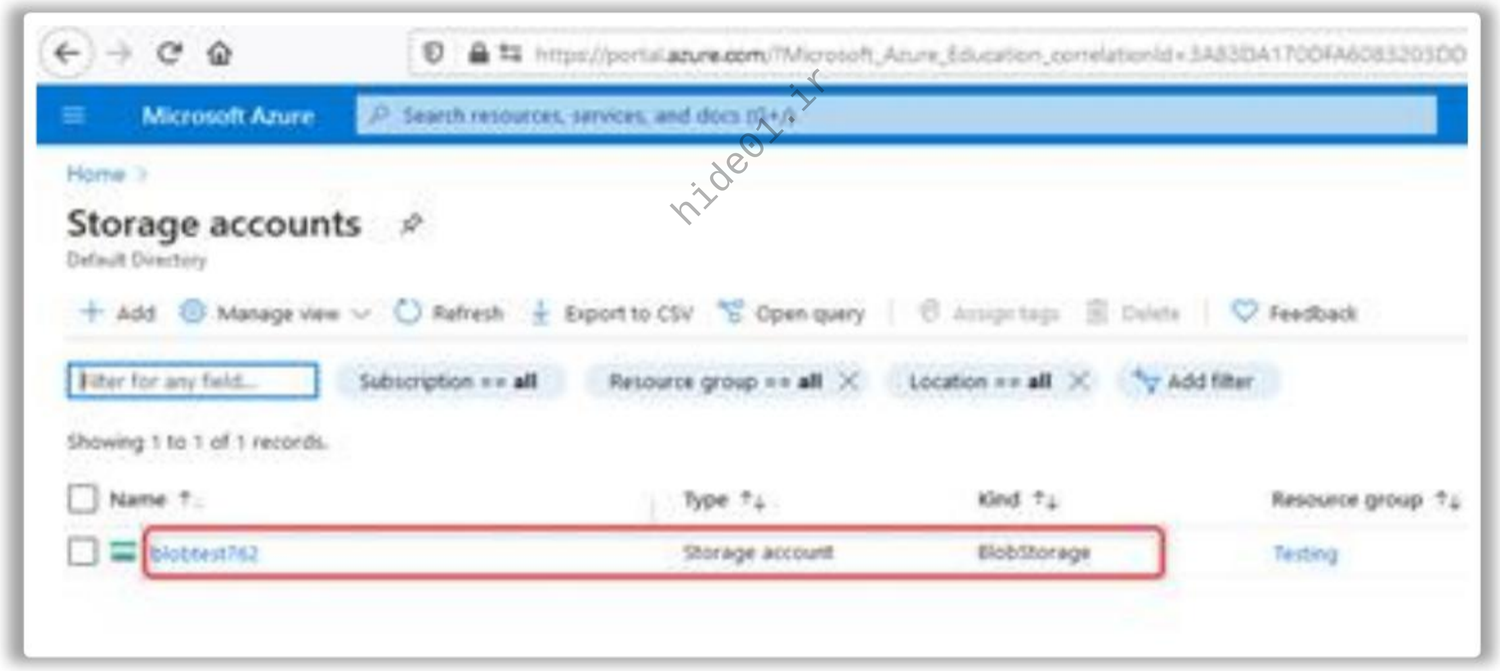


Figure 4.260: Choose Storage

- Navigate to **Containers** under **Blob service**.
- Check the **Container** that has **Public access level**.
- Go to **Change access level.**

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

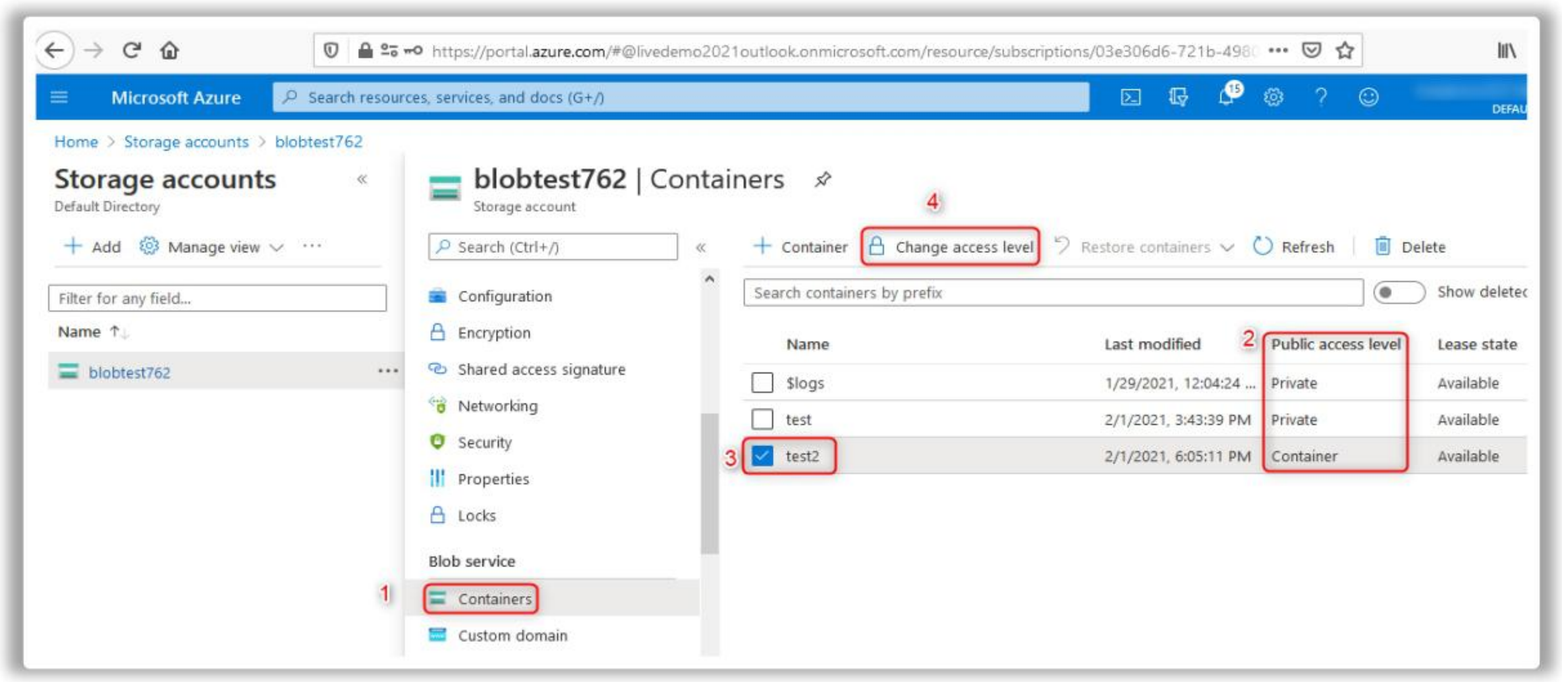


Figure 4.261: Navigate to Change Access Level

- You will not be able to change if changes are made in the storage account level.

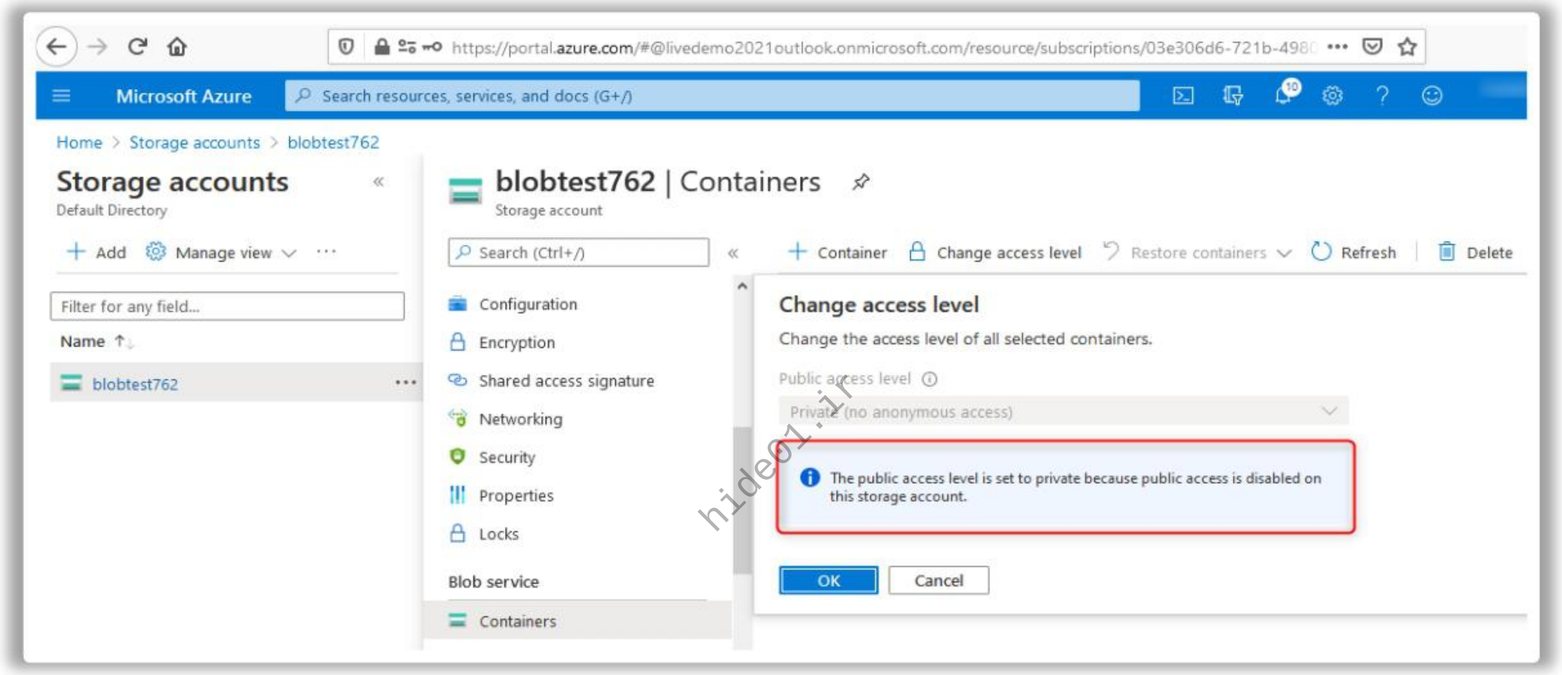


Figure 4.262: Change Access Level Console

- If the access is not changed at the storage level and the container access can be modified.
- Navigate to **Containers** under **Blob service.**
- Search for the **Container** parameter under the **Public access level**, that is, it has public access.

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

- Click on the **Change access level.**

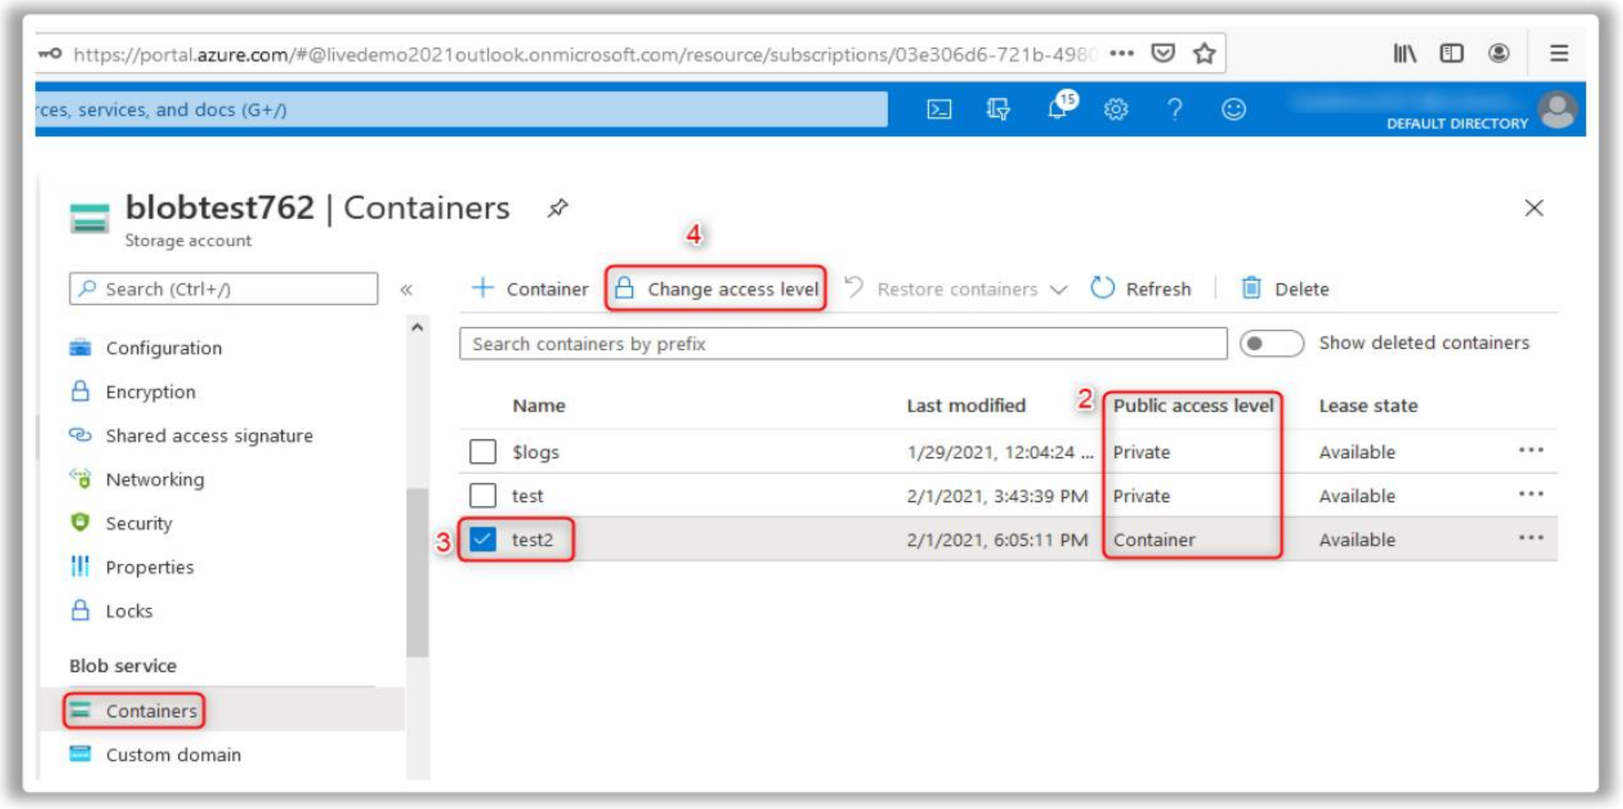


Figure 4.263: Click Change Access Level

- Choose **Private (no anonymous access)** from the dropdown.

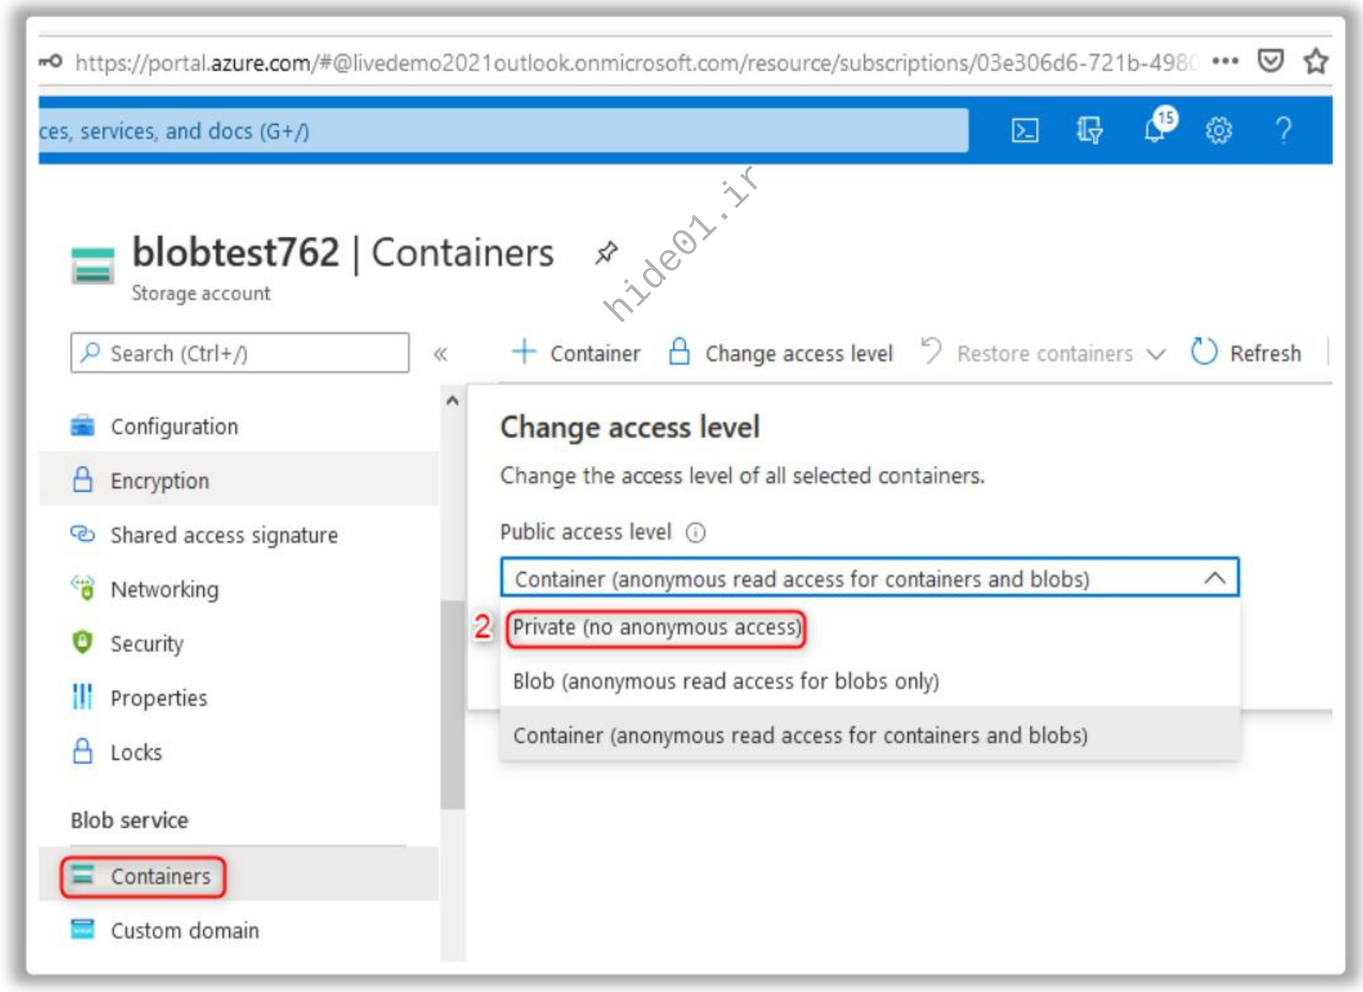


Figure 4.264: Click Private

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

- Click **OK.**

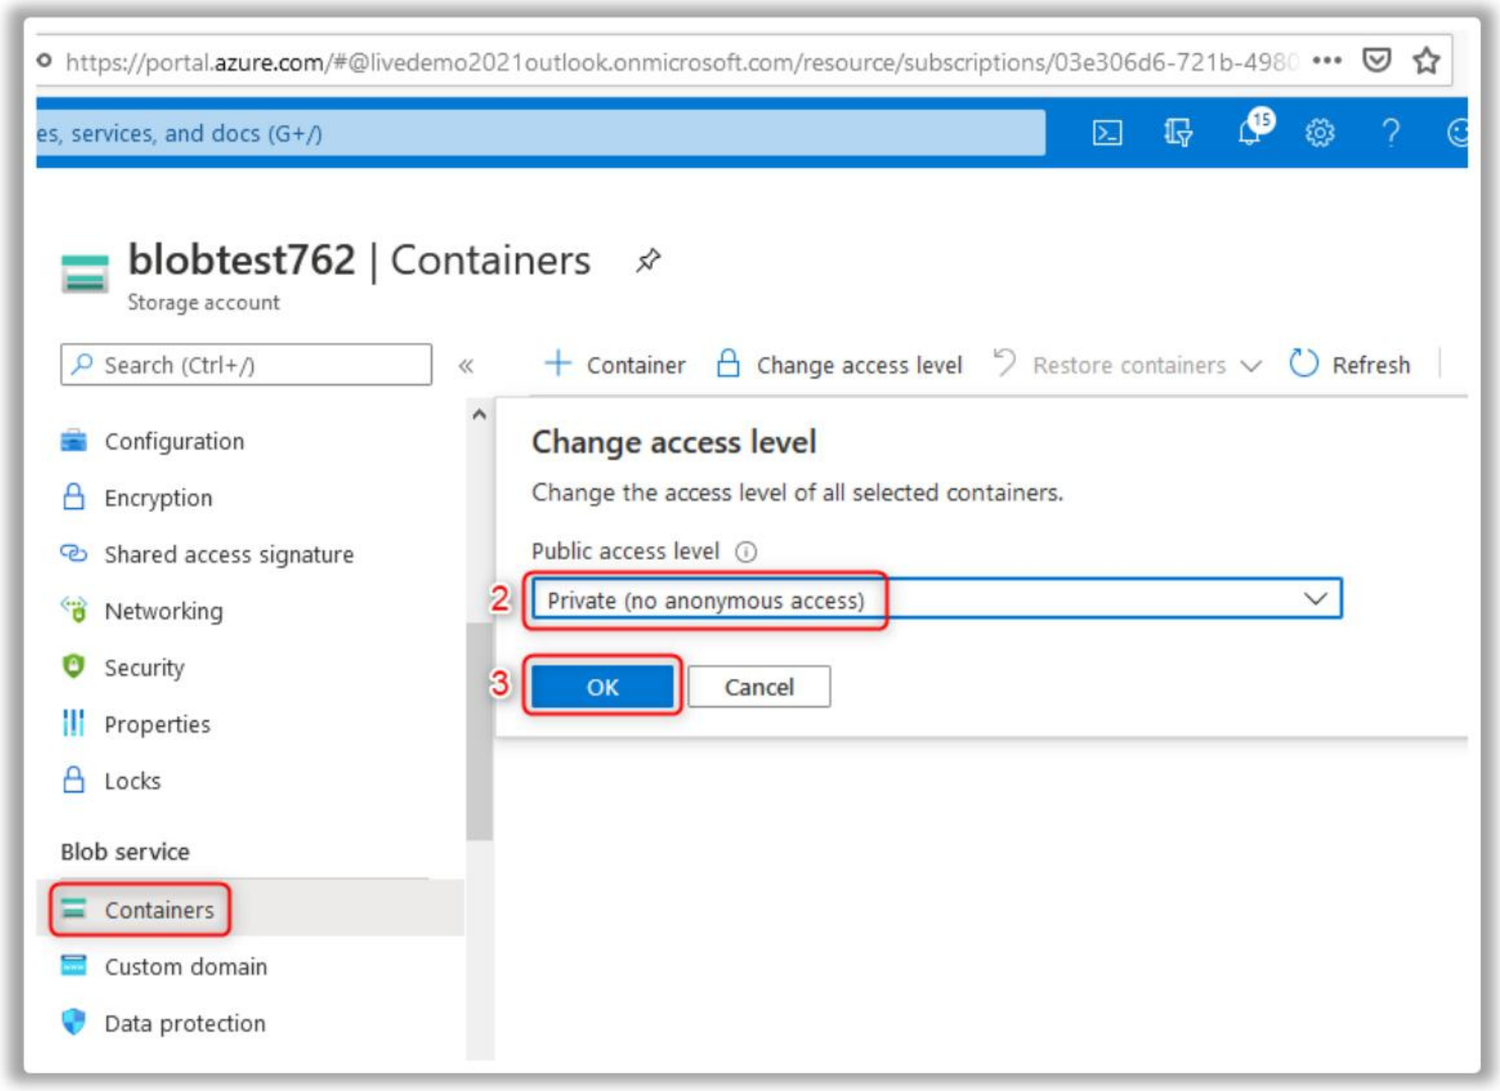


Figure 4.265: Click OK

- The public access level has successfully changed to **Private.**

hide01.ir

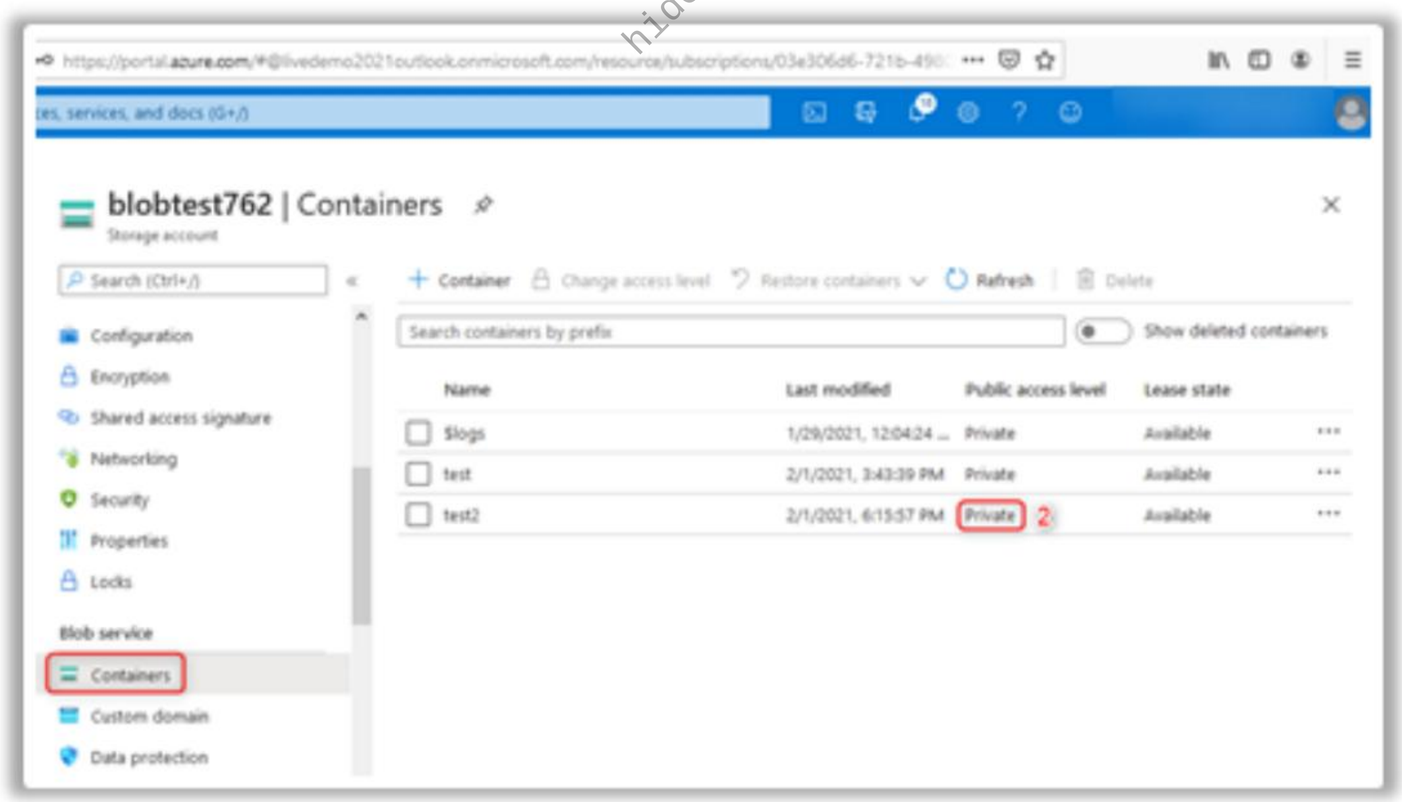
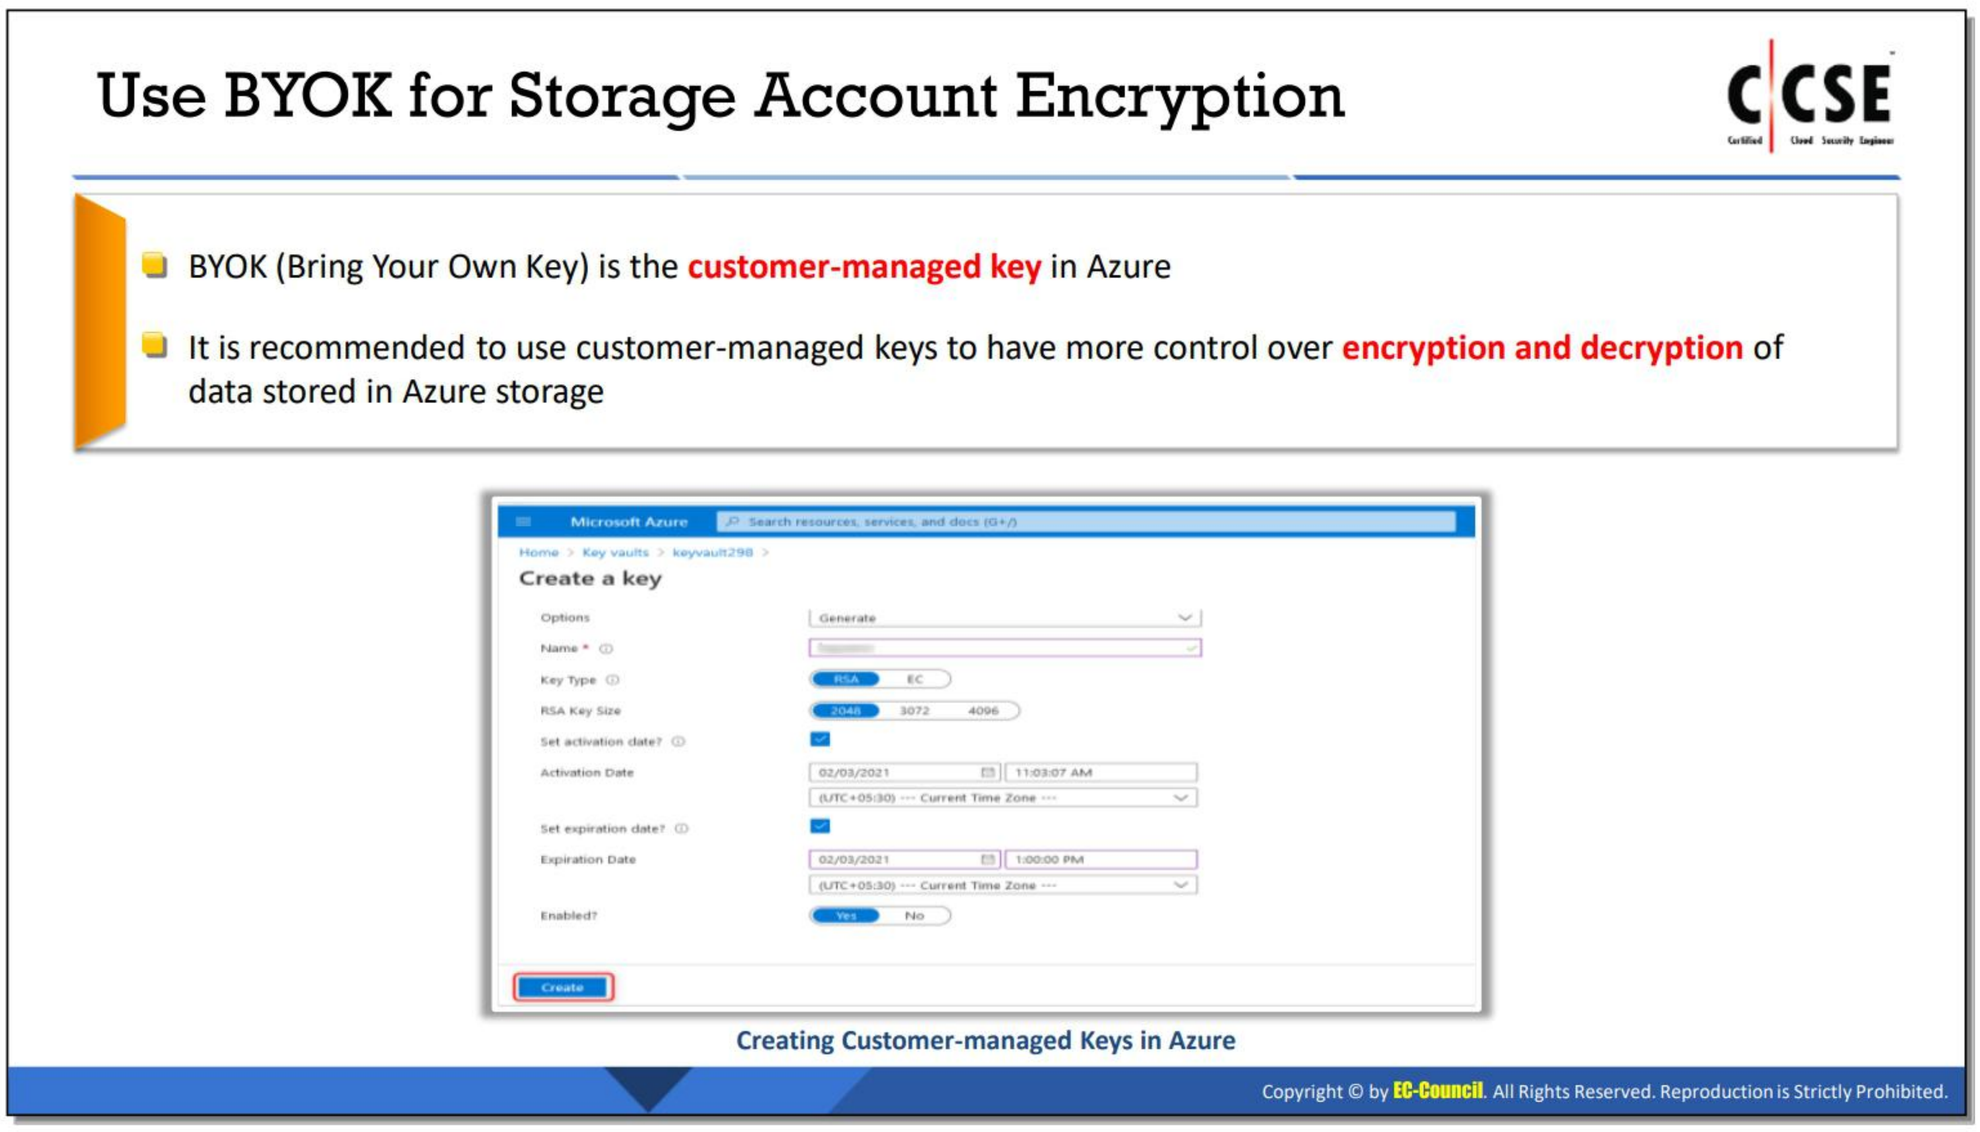


Figure 4.266: Select Private

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



Use BYOK for Storage Account Encryption

BYOK is the CMK in Azure. The stored data in the Azure storage accounts are encrypted by default with Microsoft-managed keys. Microsoft allows customers to encrypt data-at-rest using customer-managed keys. These keys are stored in key vaults. It is recommended to use customer-managed keys to have more control over encryption and decryption of data stored in Azure storage.

Source: https://azure.microsoft.com

To Use BYOK for Storage Account Encryption

- Sign-in to Azure Portal.
- Click on **More services.**

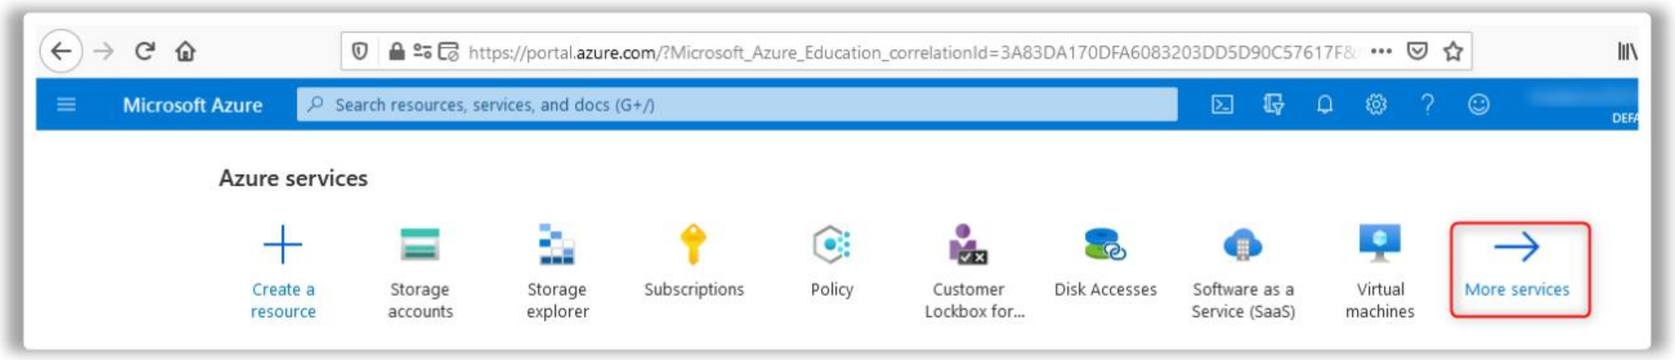


Figure 4.267: Select Storage Accounts

- Search for the **key vault** in the search box beside all services.

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

- Click on **key vault**.

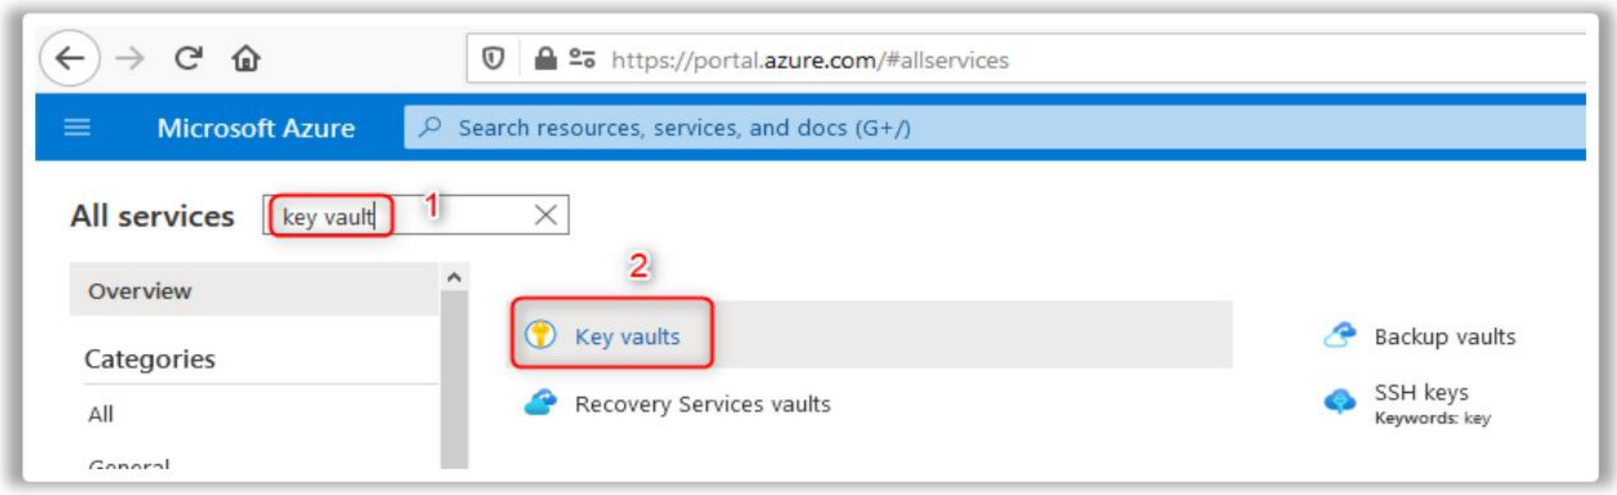


Figure 4.268: Select Key Vault

- A list of all available key vaults will be displayed.
- Select the key vault in which you want to generate keys.

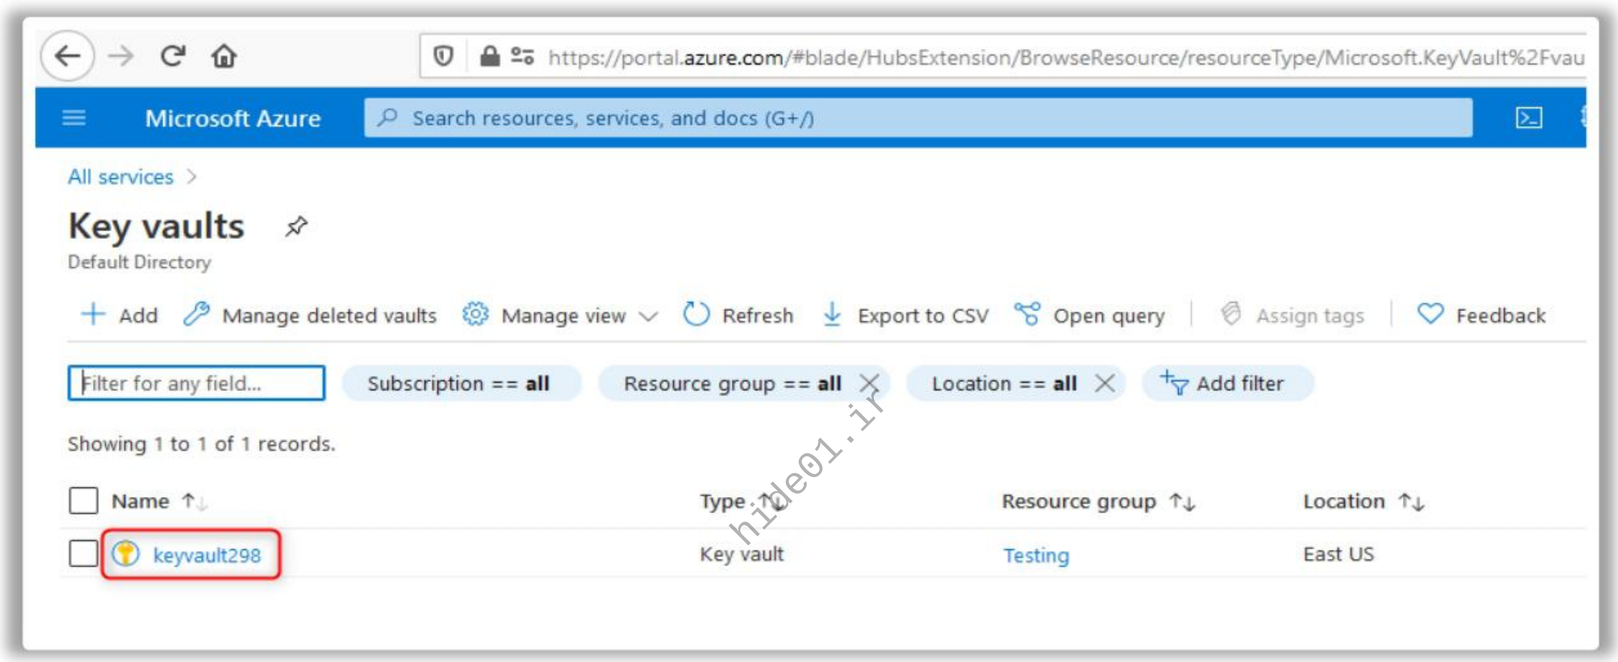


Figure 4.269: Choose Key Vault

- Locate **Settings** and navigate to **Properties**.
- Locate **Purge protection**, select the **Enable purge protection (enforce a mandatory retention period for deleted key vaults and vault objects)**
- Purge protection may be already enabled, while key vault creation.
- Click on **Save**.

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

- This helps in the automatic update of key versions.

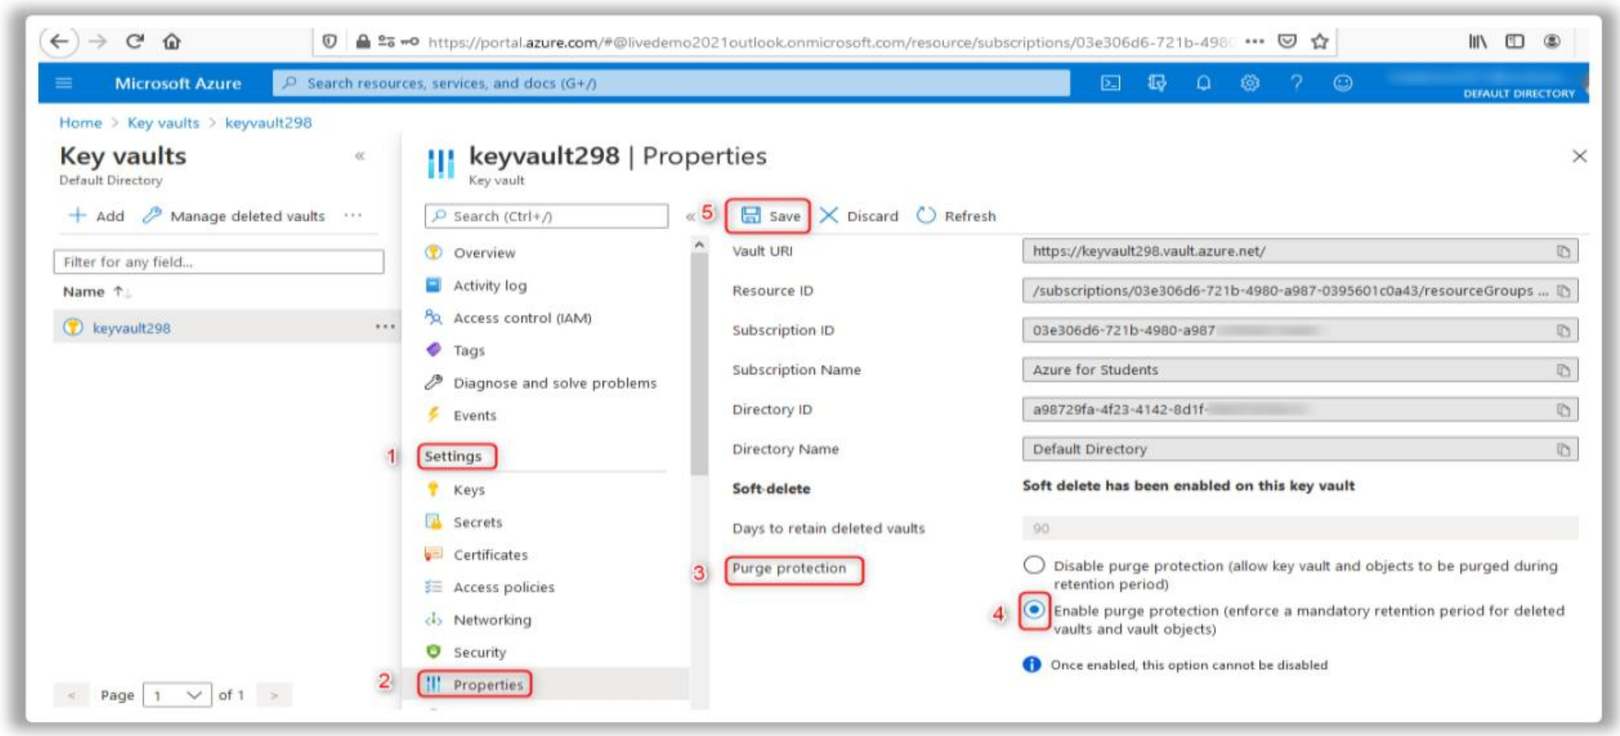


Figure 4.270: Click Save

- Locate **Settings** and navigate to **Keys**.
- Click on **Generate/Import.**

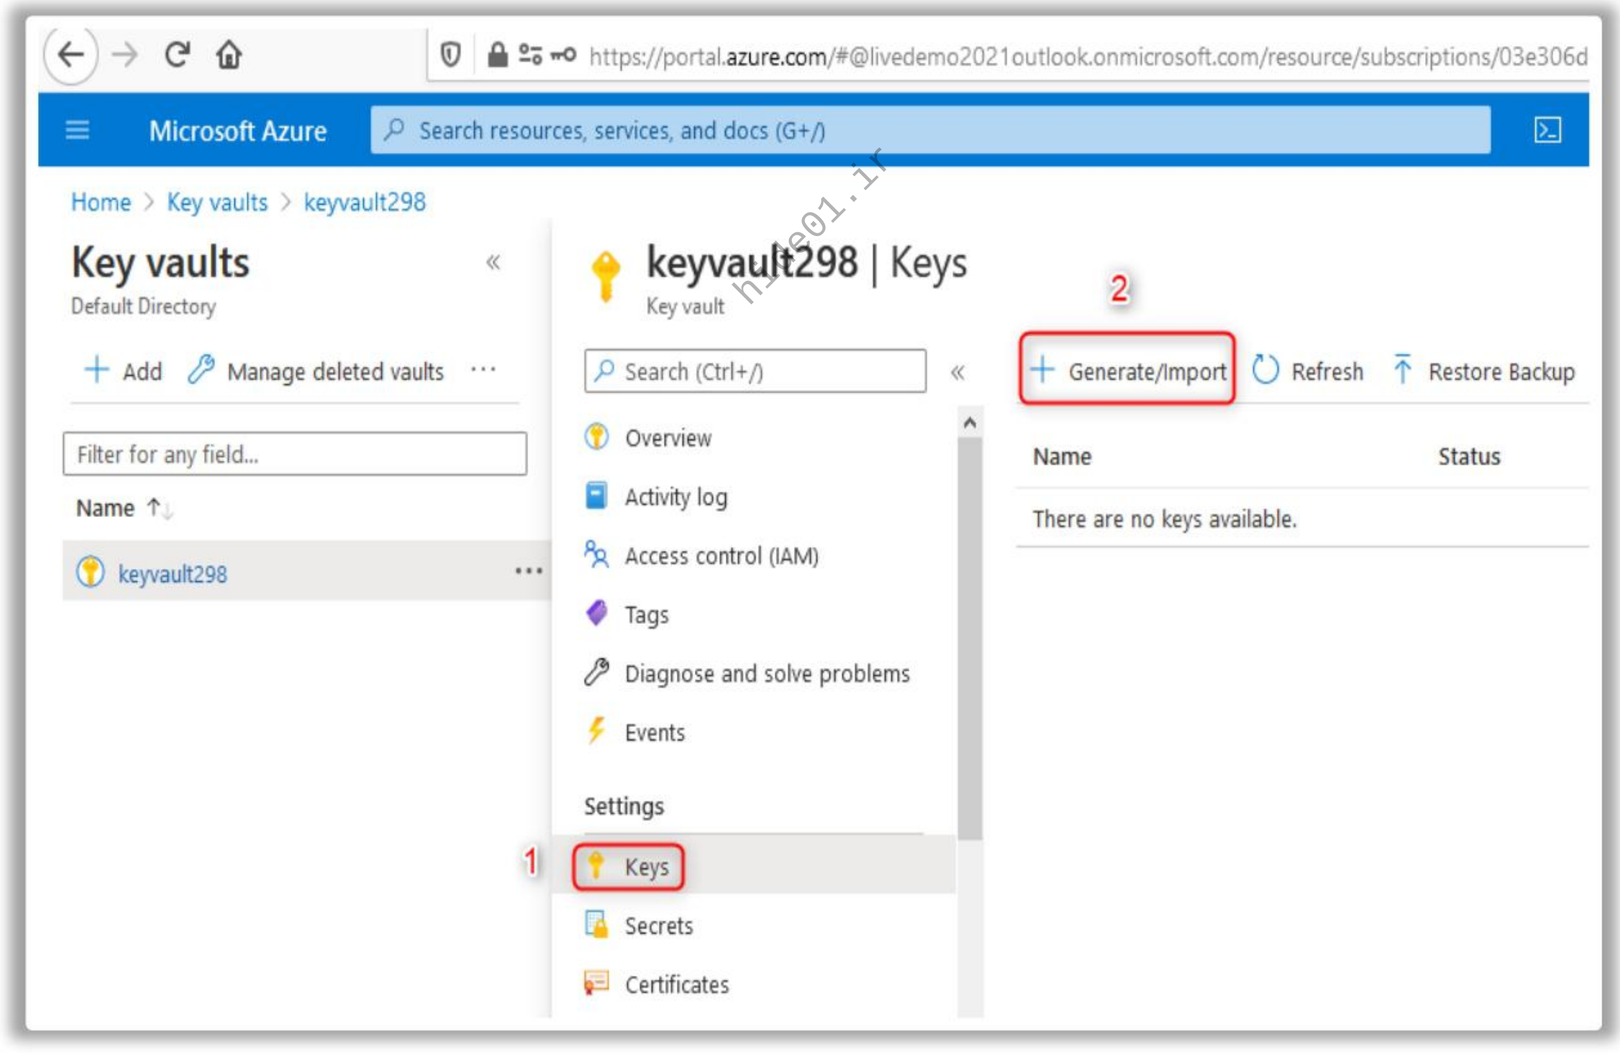


Figure 4.271: Select Generate/Import

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

- Name the key and configure the details like **Key Type, RSA Key Size, Set activation date, Set expiration date** as per requirement.

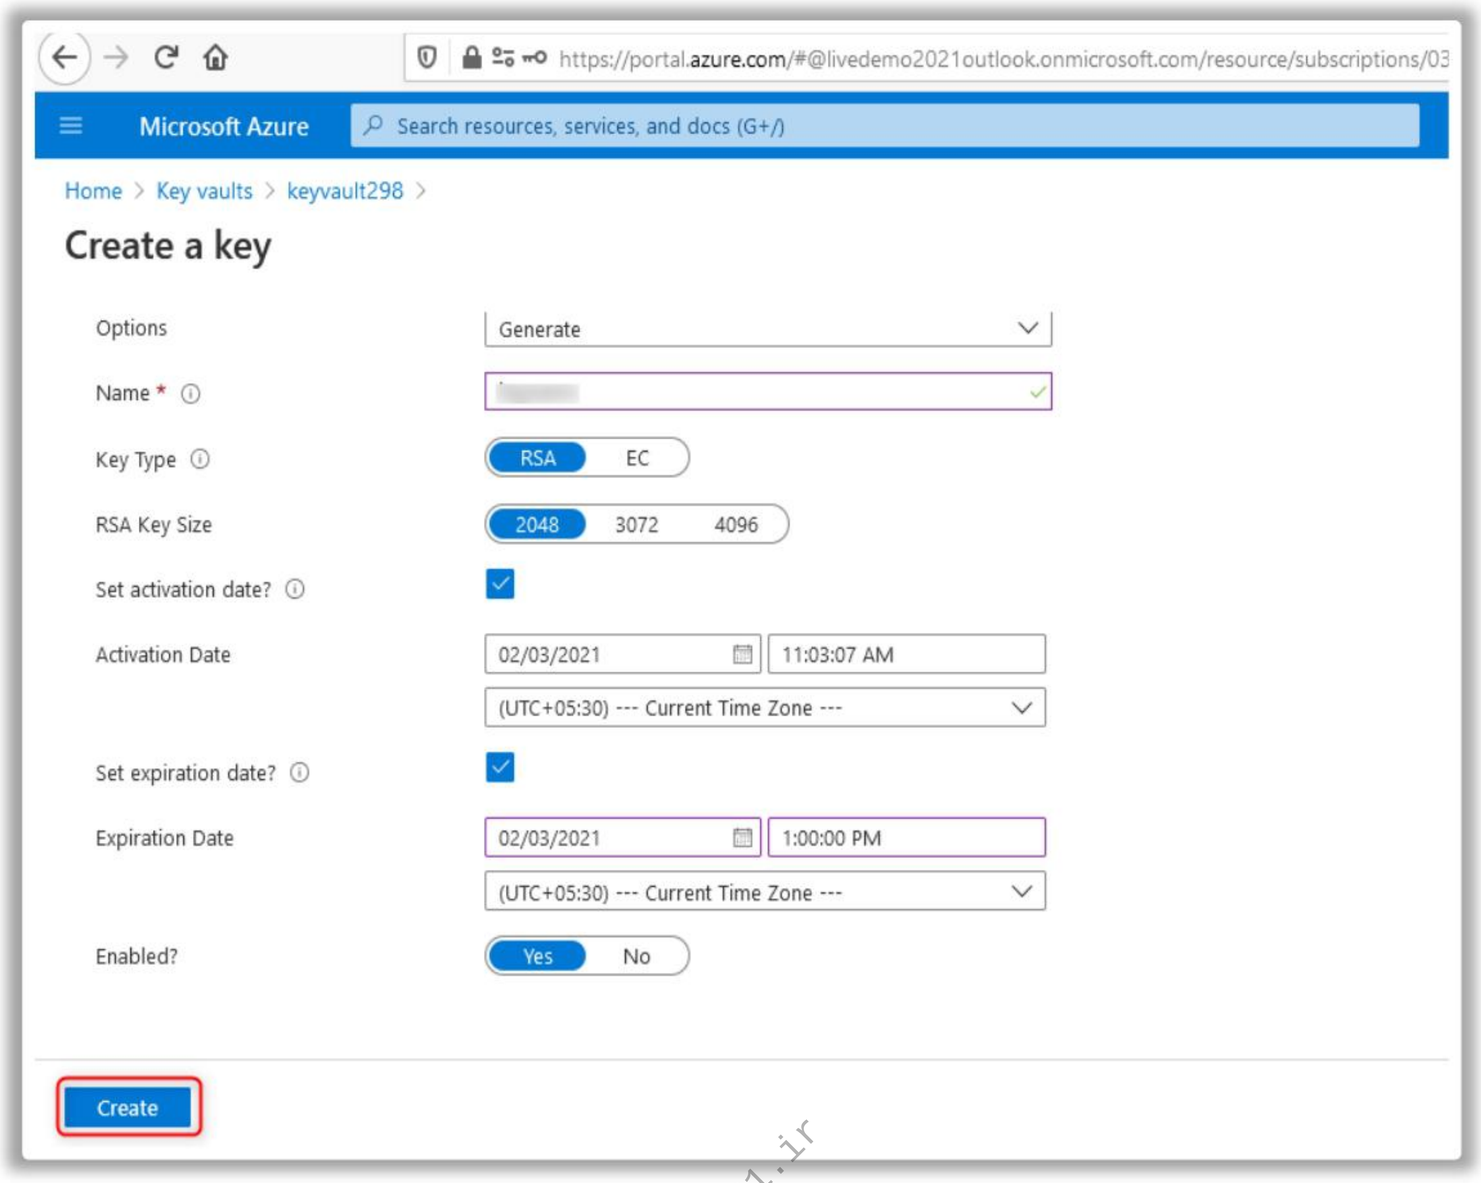


Figure 4.272: Select Details

- The Key is successfully created.

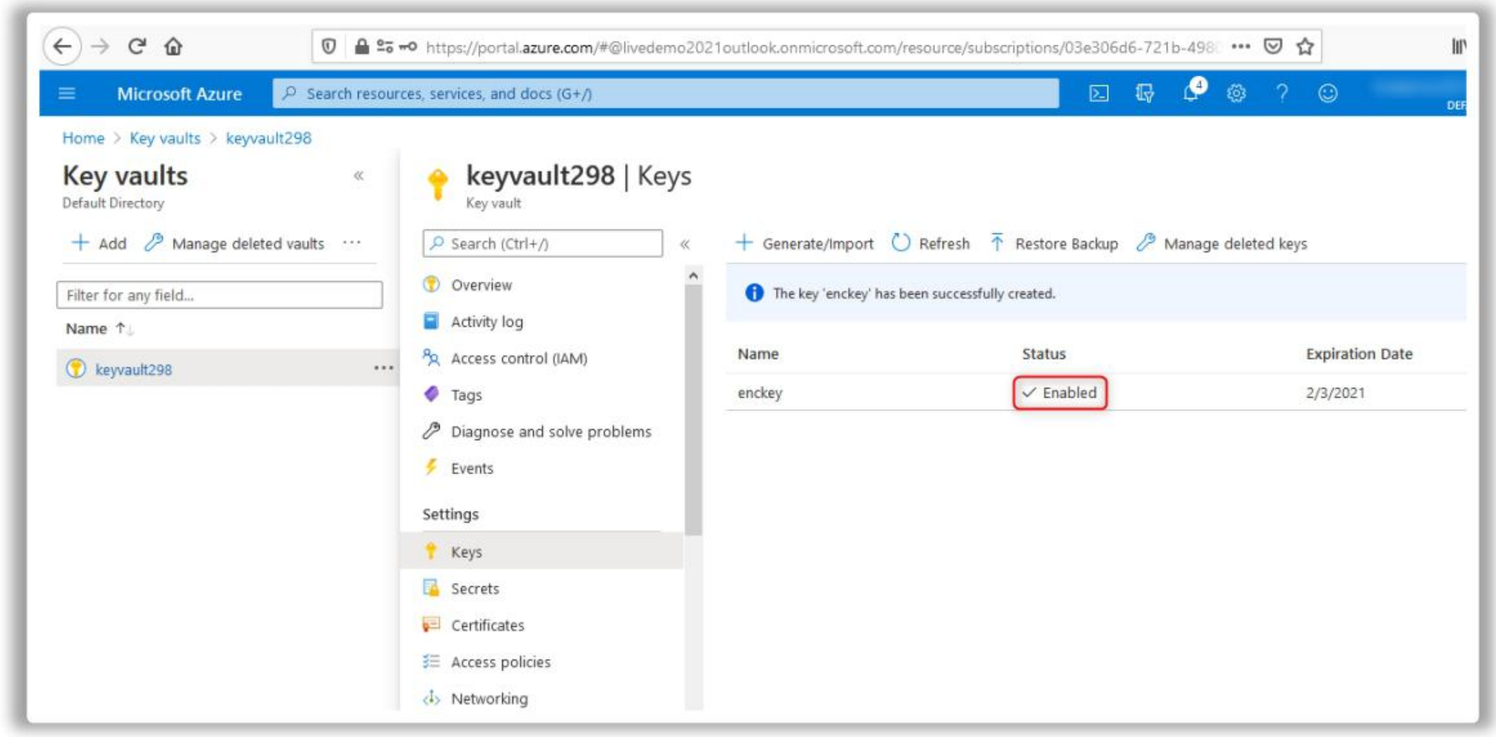


Figure 4.273: Key is Enabled

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

- Now go back to the Azure portal home page and select **Storage accounts** under **Azure services.**

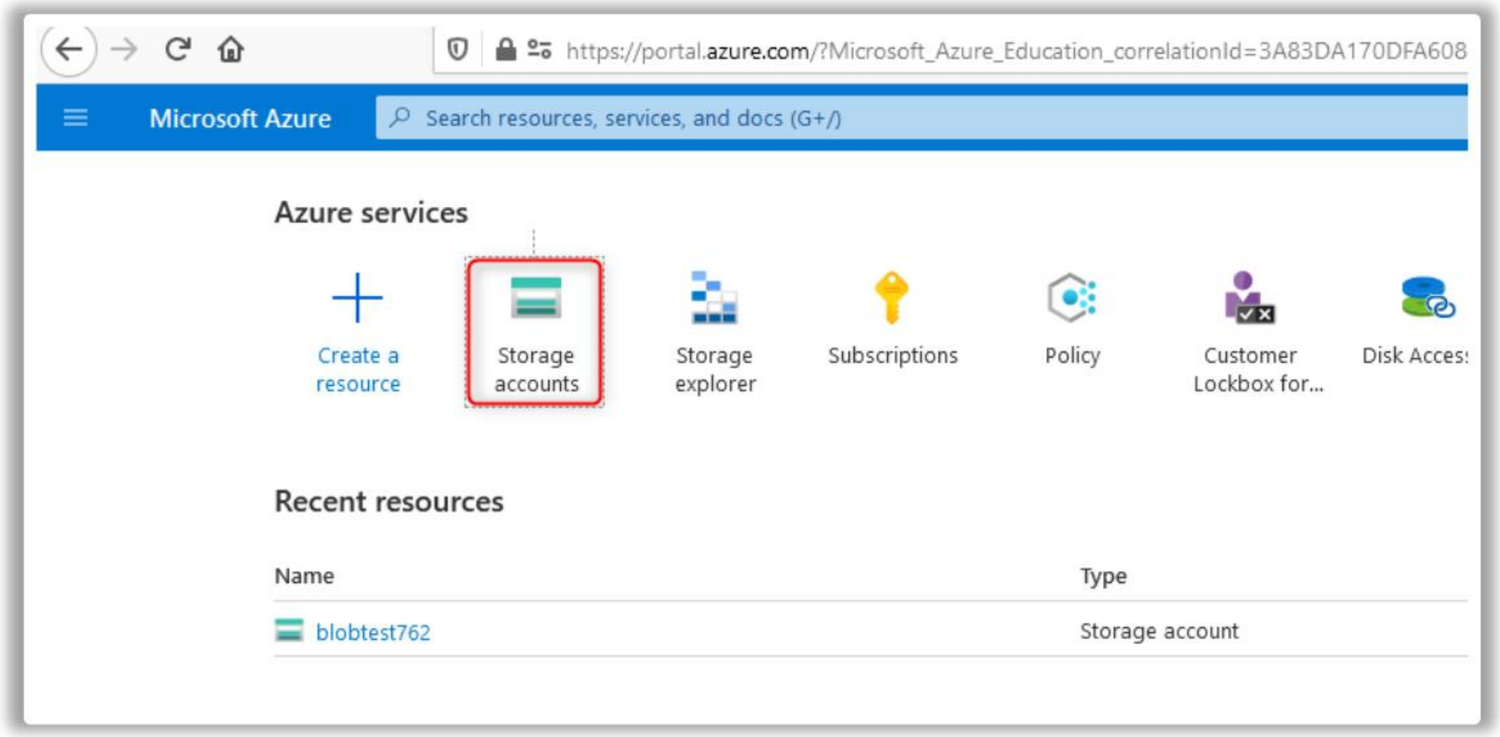
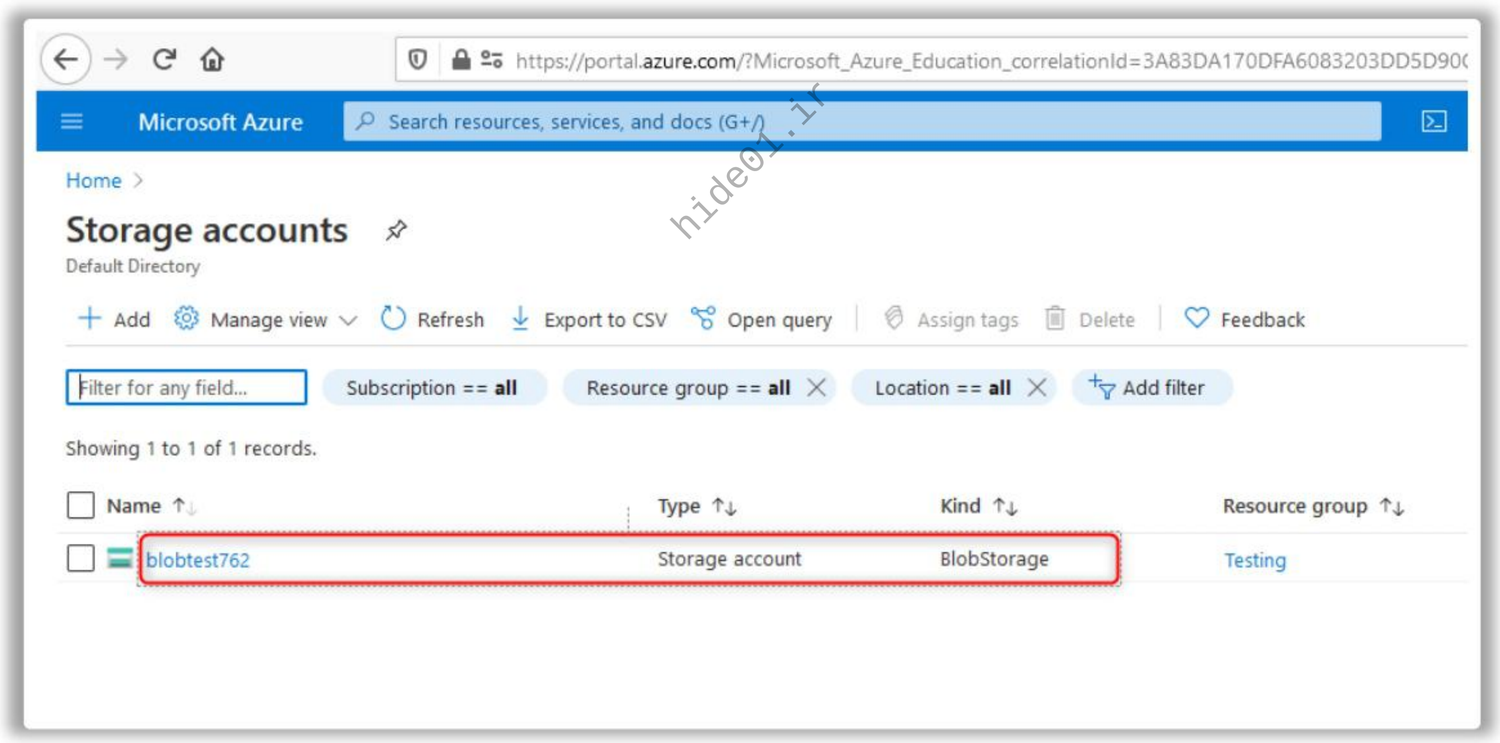


Figure 4.274: Select Storage Accounts

- A list of storage accounts will be displayed.
- Select the storage for which you want to use encryption keys.



hide01.ir

Figure 4.275: Choose Storage

- Locate **Settings** and navigate to **Encryption**.
- Go to **Encryption type** and select **Customer managed keys.**
- Set **Encryption key** to **Select from key vault.**

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

- Click on **Select a key vault and key**.

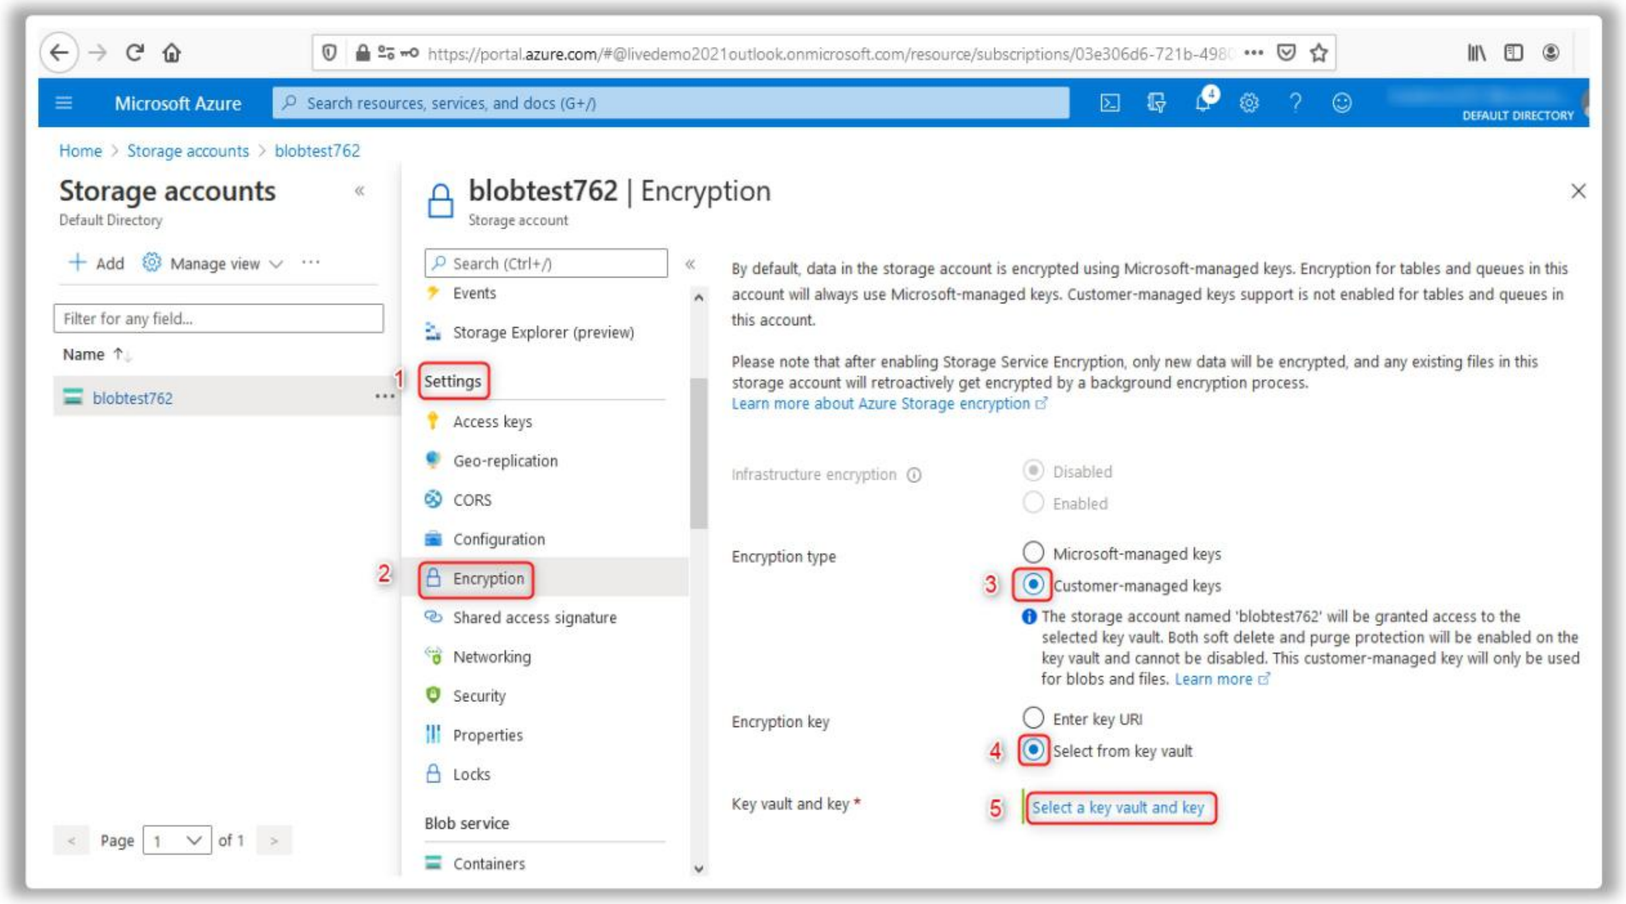


Figure 4.276: Click Select a key vault and key

- Select the **Key vault** from the dropdown list.
- Select the **Key** that was generated.
- Click on **Save**.

hide01.ir

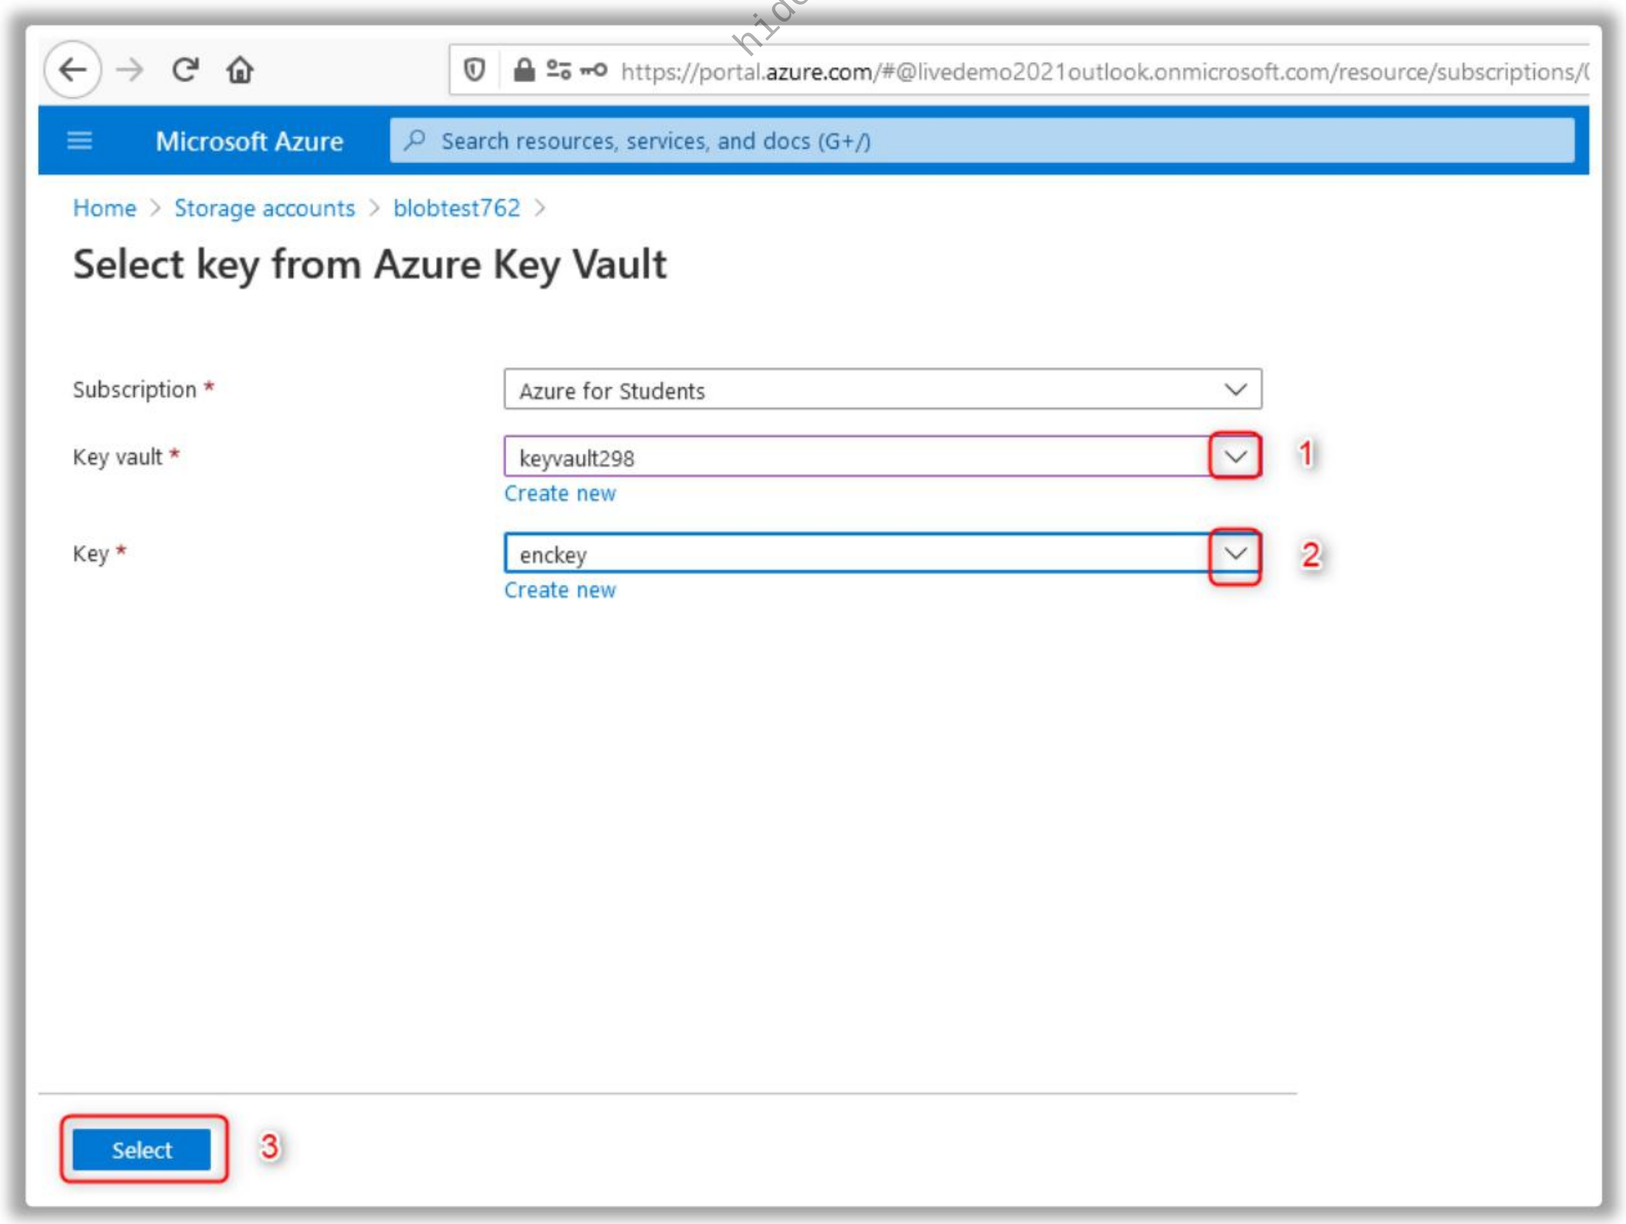


Figure 4.277: Click Select

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

- Selected key vault and key will be displayed.

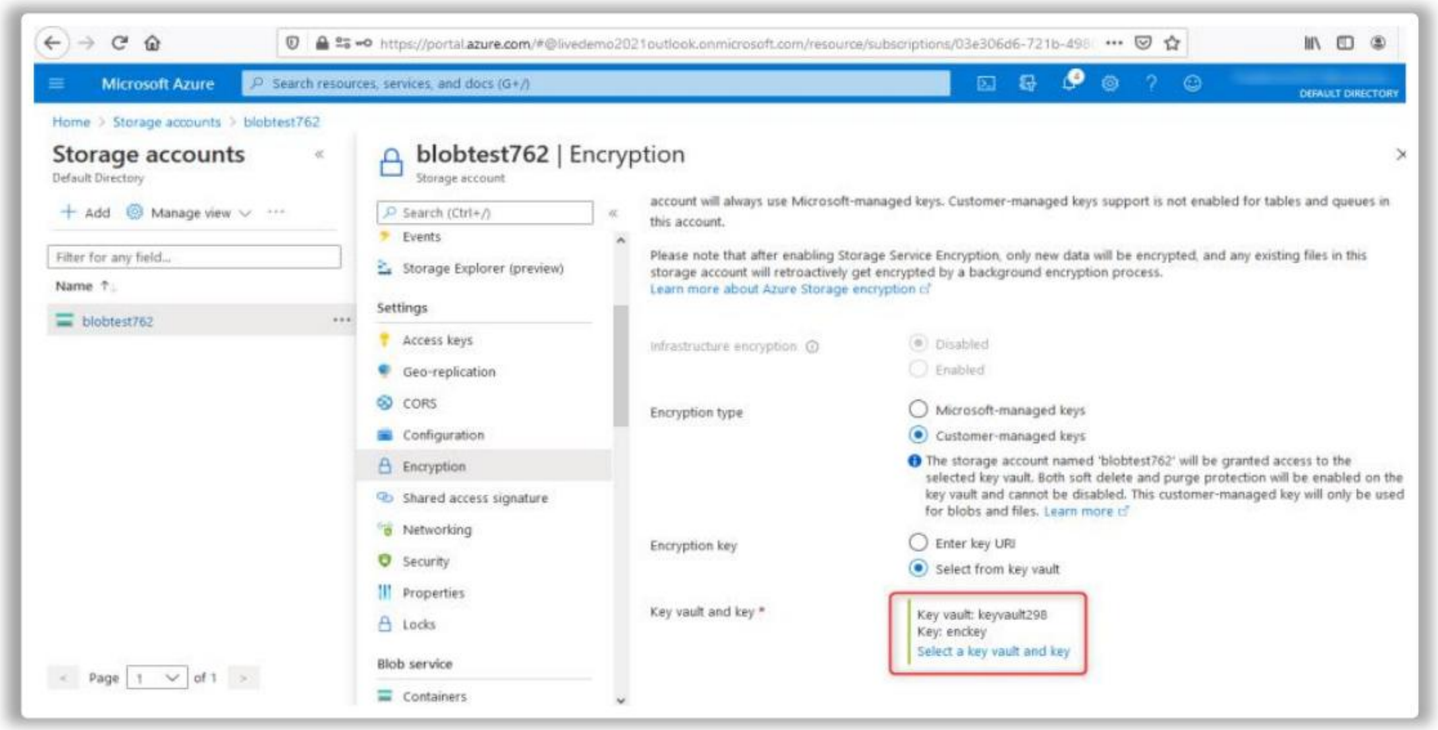


Figure 4.278: Storage Accounts Console

- Scroll to the top and click on **Save.**

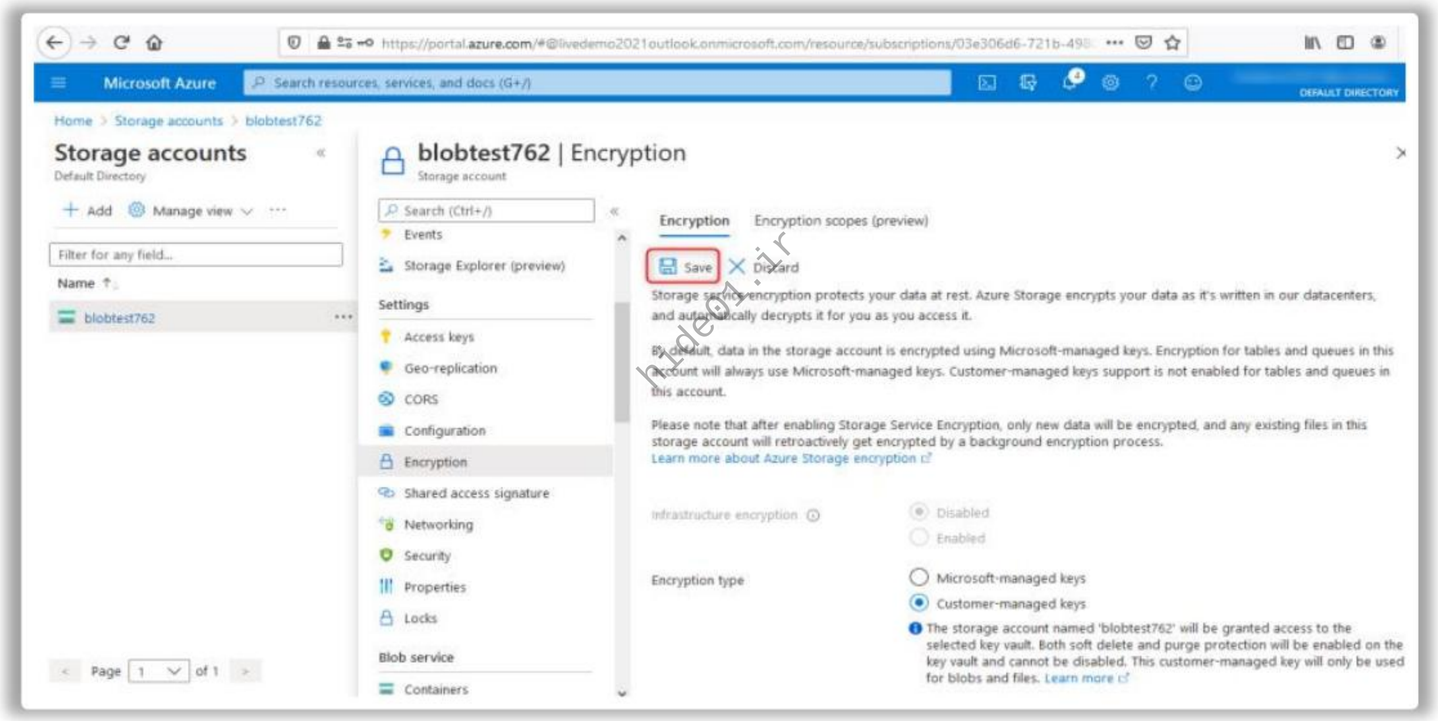


Figure 4.279: Click Save

- SSE is successfully updated.

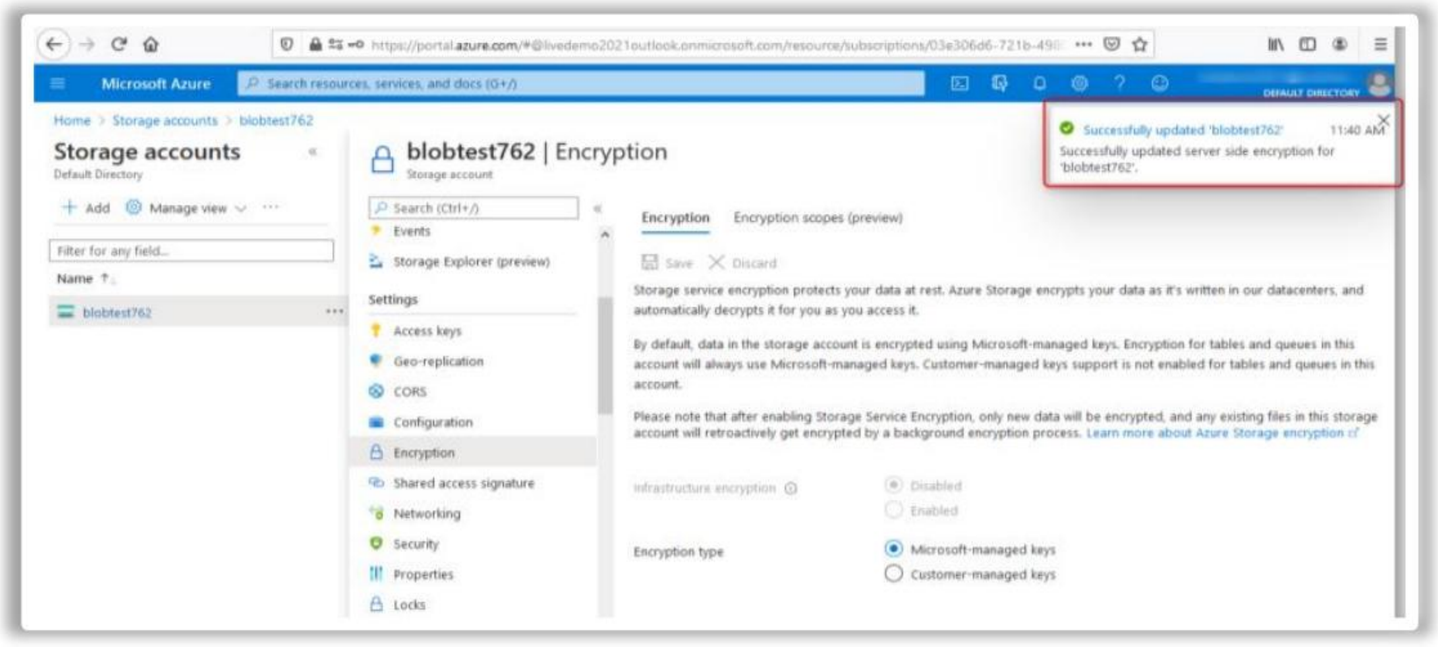
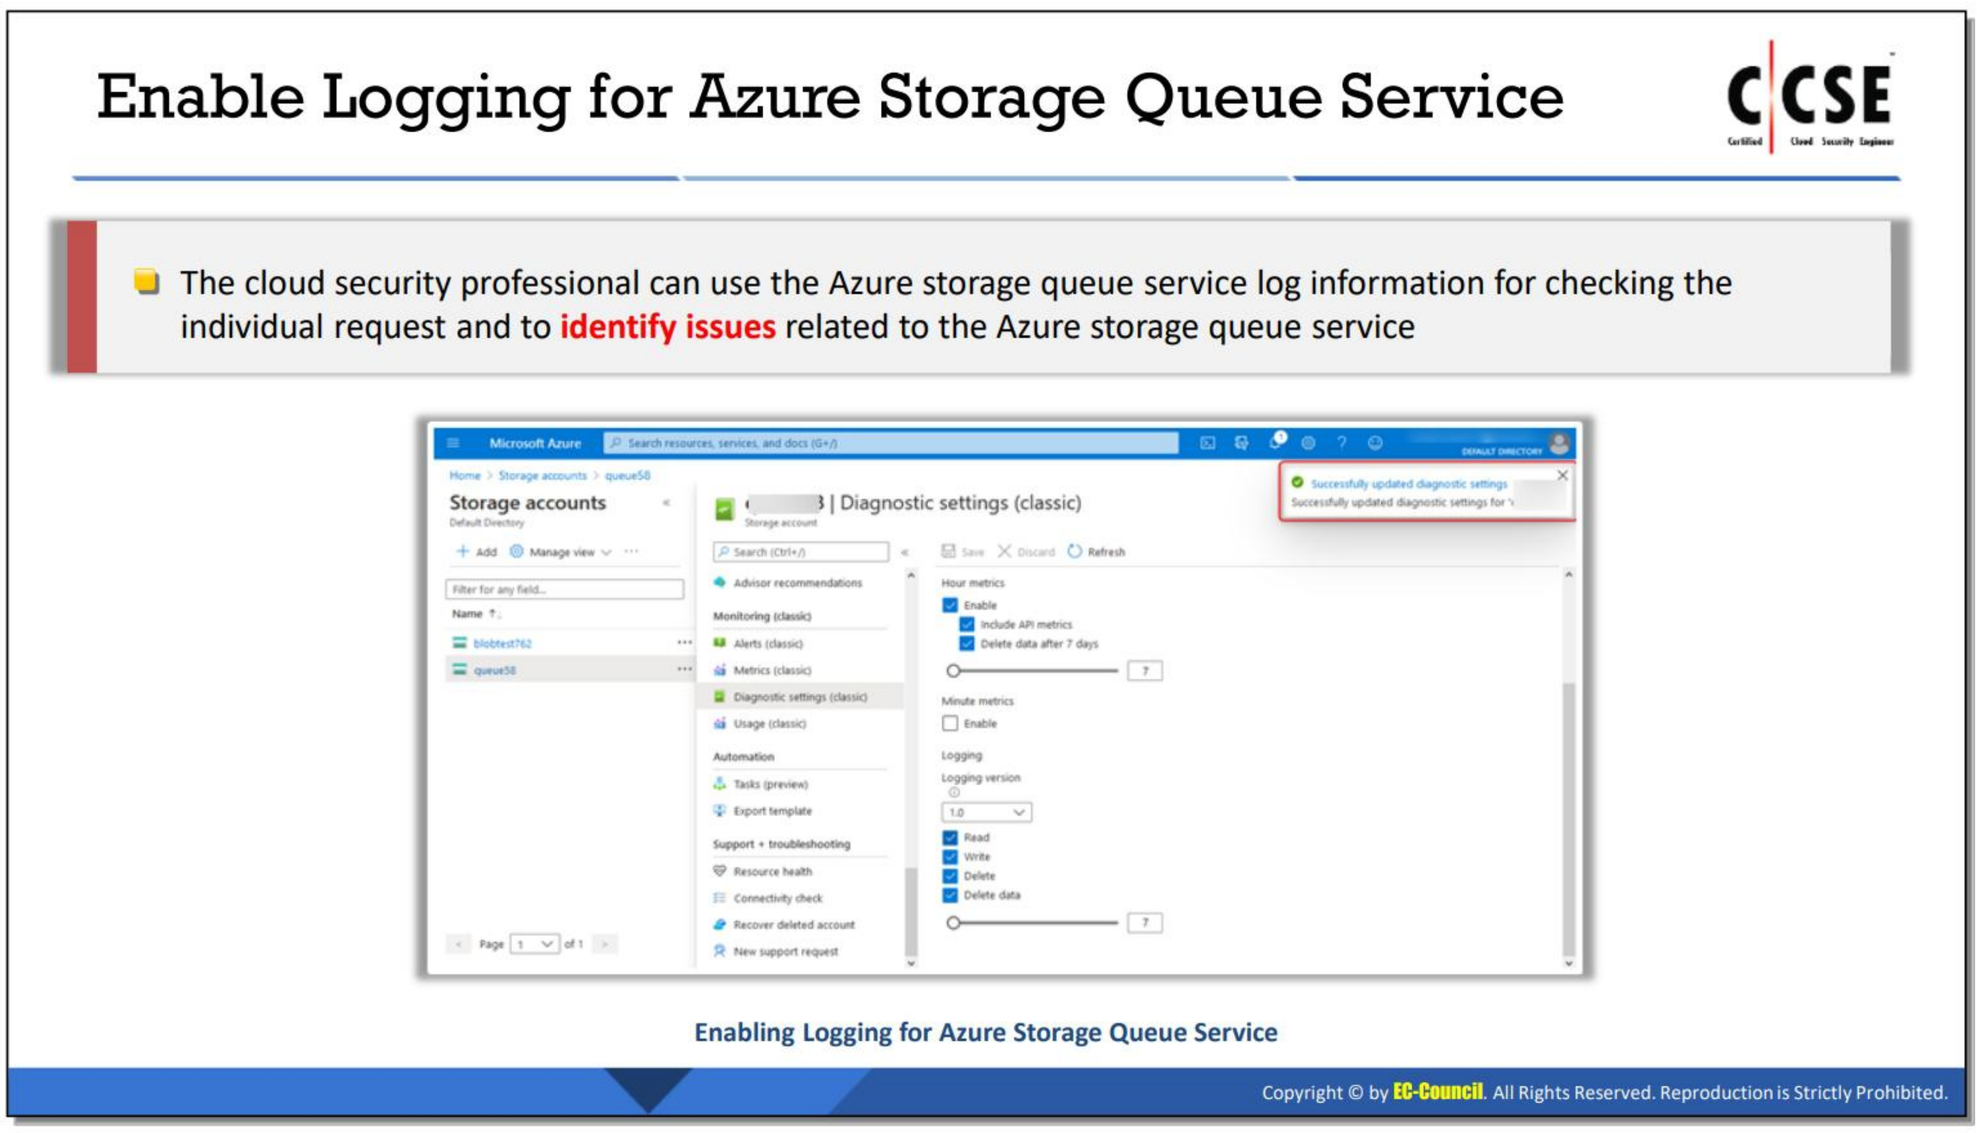


Figure 4.280: SSE is Updated

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



Enabling Logging for Azure Storage Queue Service

Enable Logging for Azure Storage Queue Service

Check logs of any service to verify the proper functioning and availability of the configured services. Logs provide successful and failed requests of selected categories. It is recommended to enable Azure storage Queue service monitoring with "read, write, and delete" categories, so that it logs information of StorageRead, StorageWrite, and StorageDelete. The cloud security professional can use the Azure storage queue service log information for checking the individual request and to identify issues related to the Azure storage queue service.

Source: https://azure.microsoft.com

To Enable Logging for Azure Storage Queue Service

- Sign-in to the Azure portal.
- Select **Storage accounts** under **Azure services.**

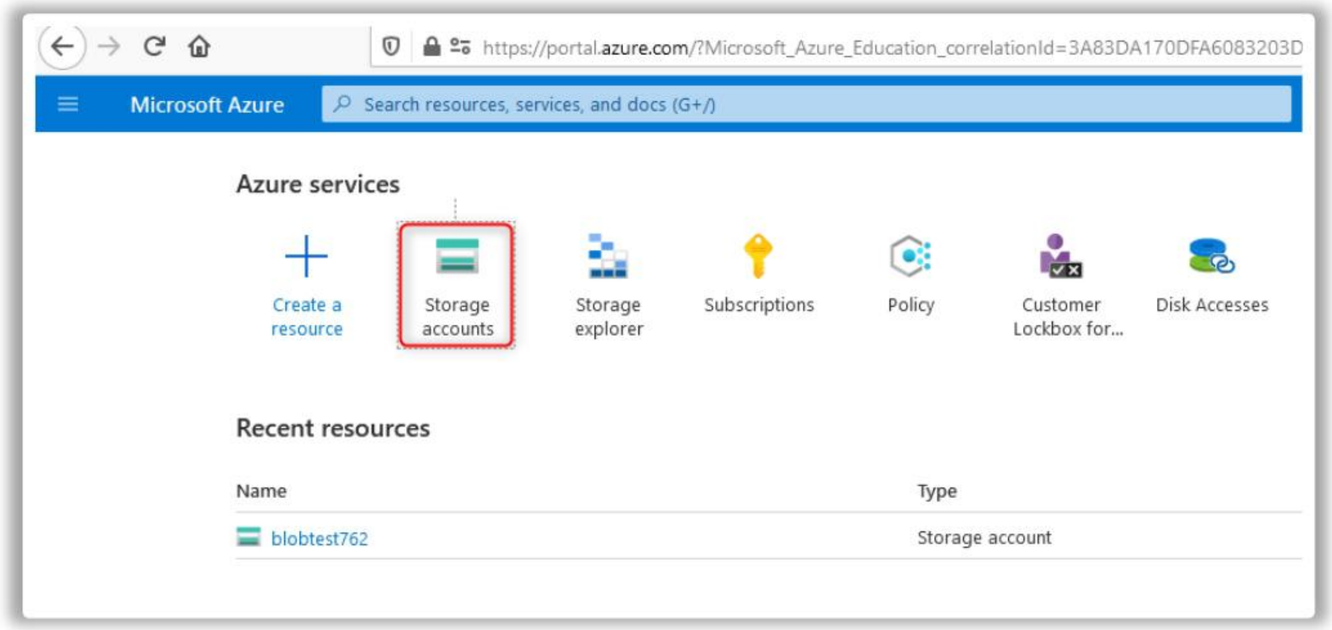


Figure 4.281: Select Storage Accounts

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

- A list of storage accounts will be displayed.
- Select the storage for which you want to enable logging for the Azure storage queue service.

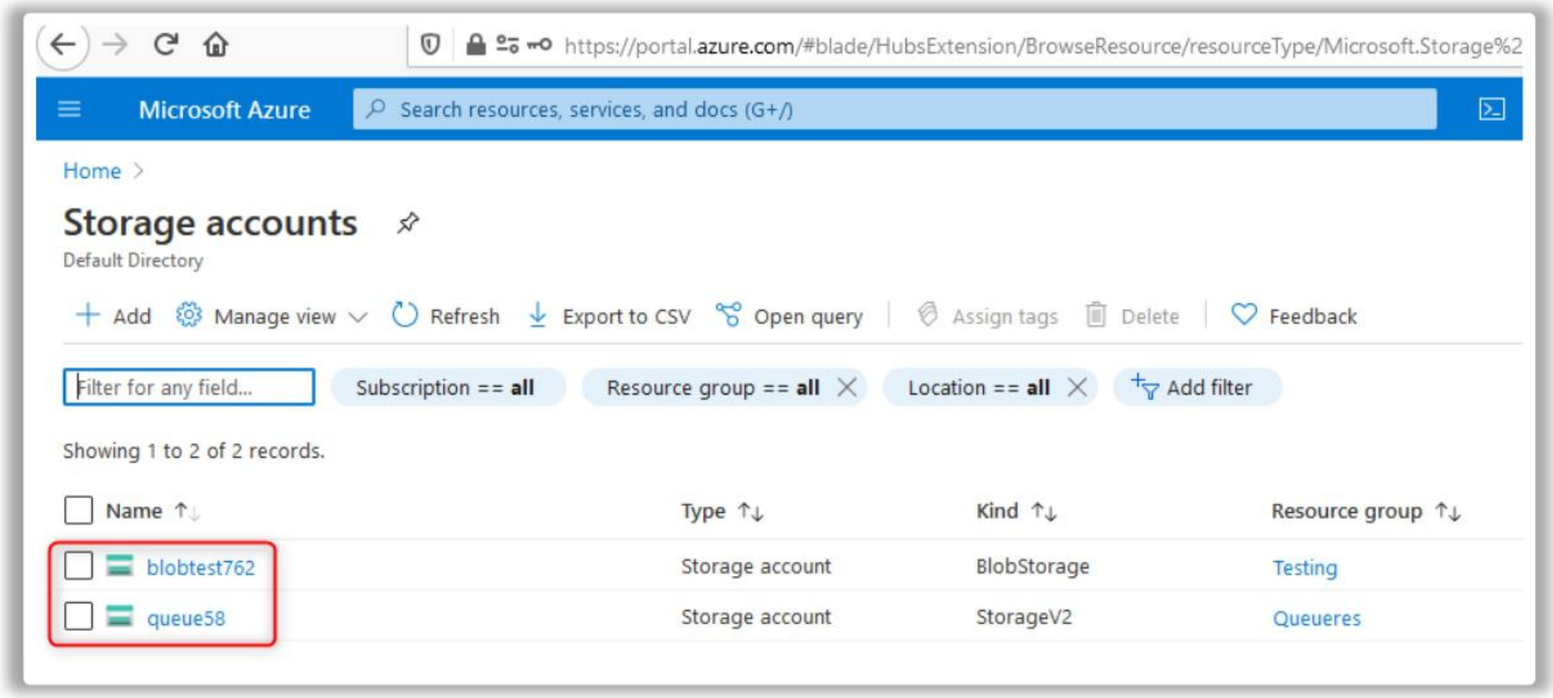


Figure 4.282: choose Storage

- In the navigation panel scroll down and locate **Monitoring (classic)** and go to **Diagnostics settings (classic)**.
- Switch the toggle bar to **On**, if it is off.
- Move to **Queue properties**.
- Select the required number of days in the input field for the retention period.
- Enable **Read, Write,** and **Delete** options for detailed storage logging of Queue services.
- Click on **Save**.

hide01.ir

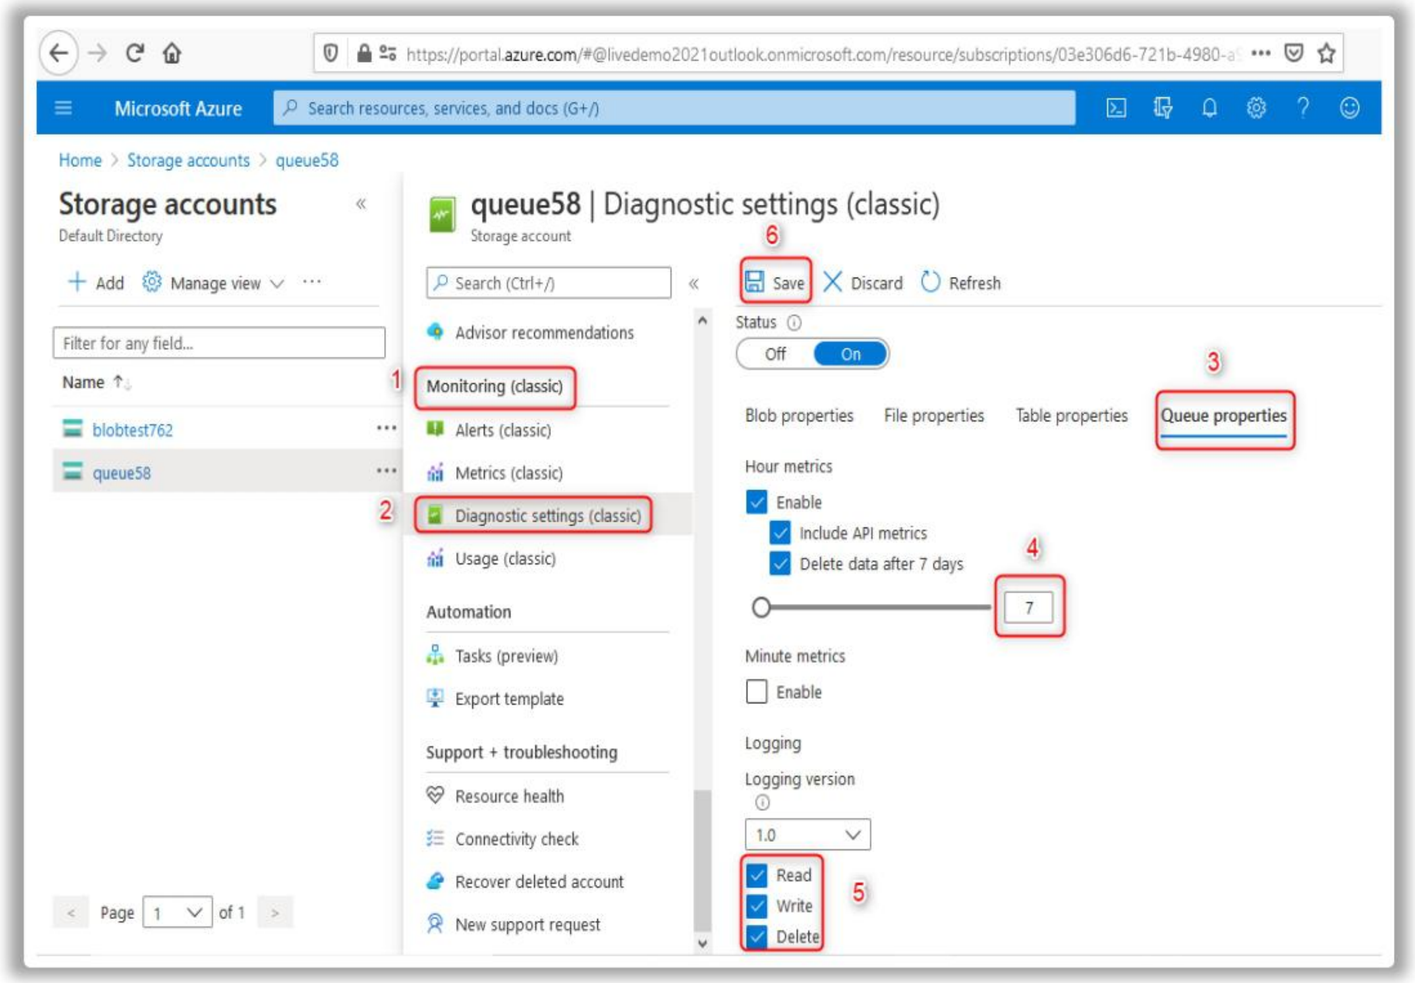


Figure 4.283: Click Save

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

- The logging feature is successfully updated for the Azure storage Queue service.

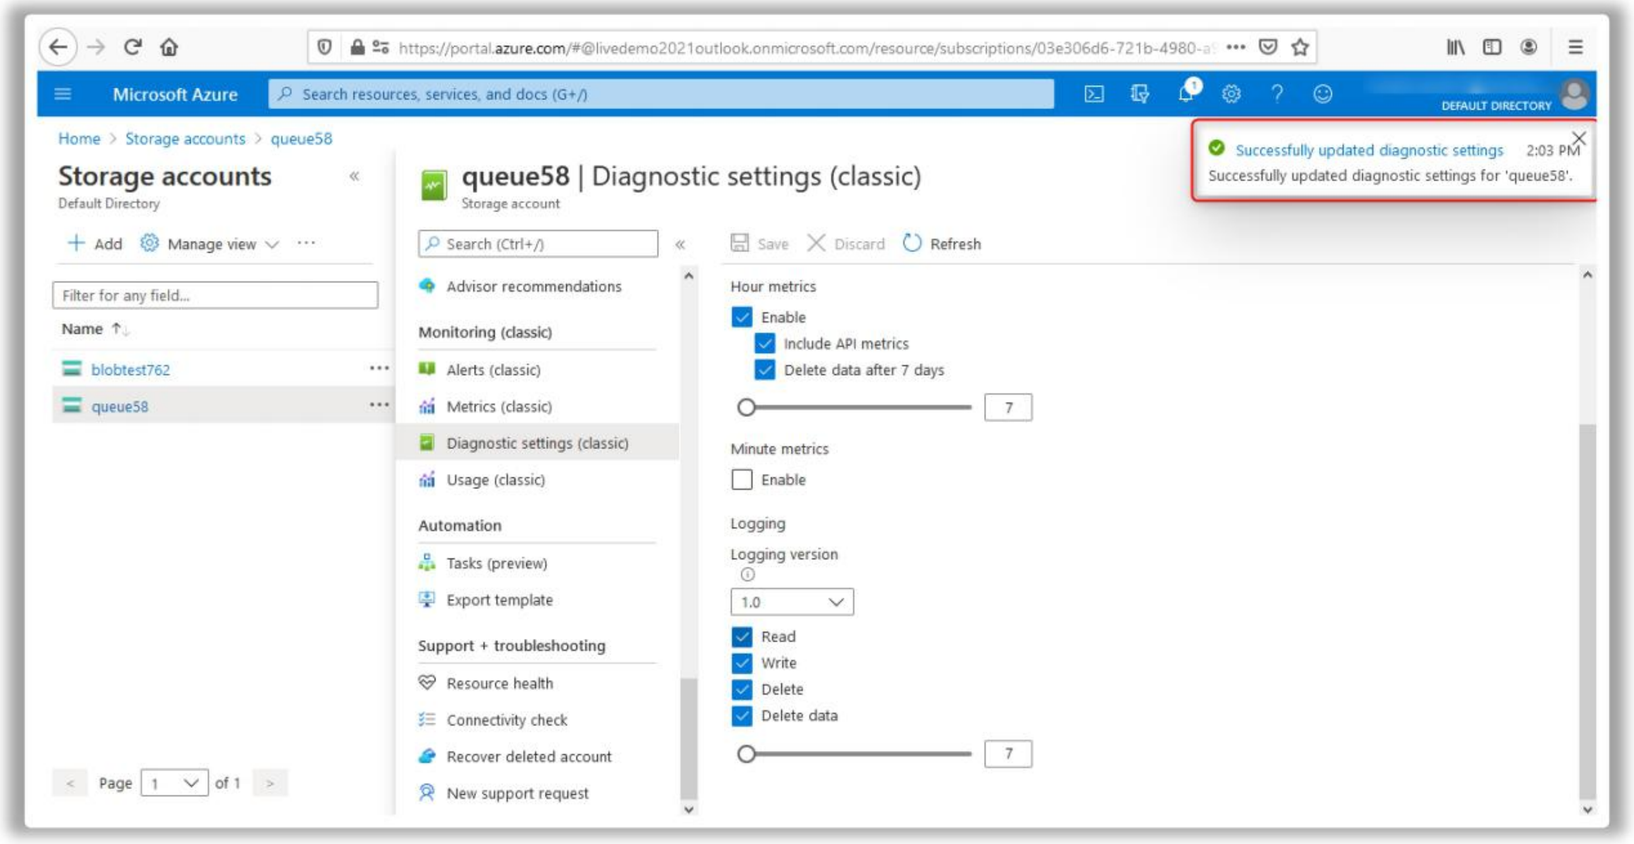


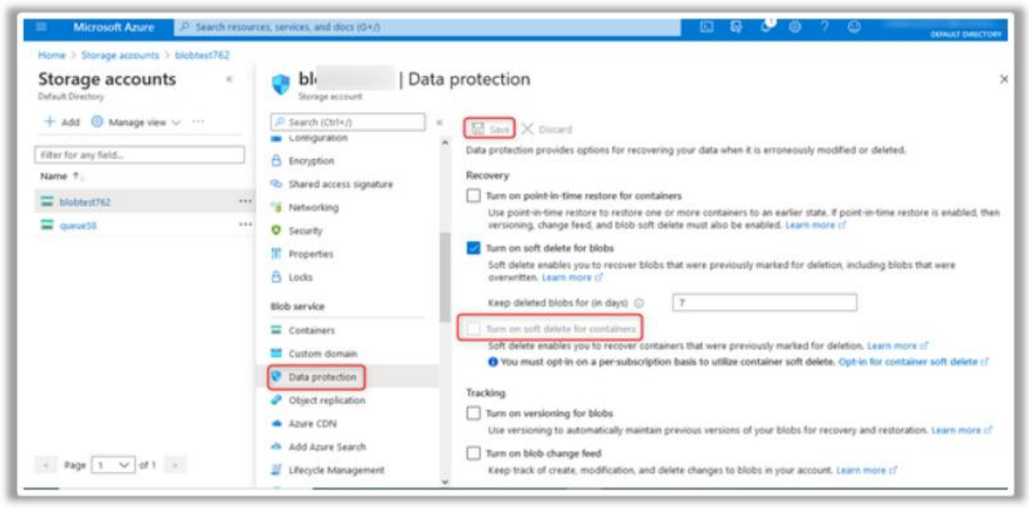
Figure 4.284: Logging Feature is Updated

hide01.ir

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

Use Soft Delete for Containers

- Soft delete enables to **restore** any accidentally deleted containers within the retention period.
- The **retention period** provided by Azure is from 1 day to 365 days
- **Undelete Container option** recovers the deleted container. Once the retention period is completed the container will be deleted permanently



Enabling Soft Delete for Container

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

Use Soft Delete for Containers

Soft delete enables users to restore any accidentally deleted containers within the retention period. The retention period provided by Azure is from 1 day to 365 days. By default, it provides a retention period of 7 days. Undelete Container option recovers the deleted container. Once the retention period is completed the container will be deleted permanently.

If soft delete for the container is disabled, then the previously soft-deleted containers will not be deleted permanently. They are permanently deleted only after the completion of the retention period which was set at the time of enabling.

Source: https://azure.microsoft.com

To Use Soft Delete for Containers

- Sign-in to the Azure portal.

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

- Select **Storage accounts** under **Azure services.**

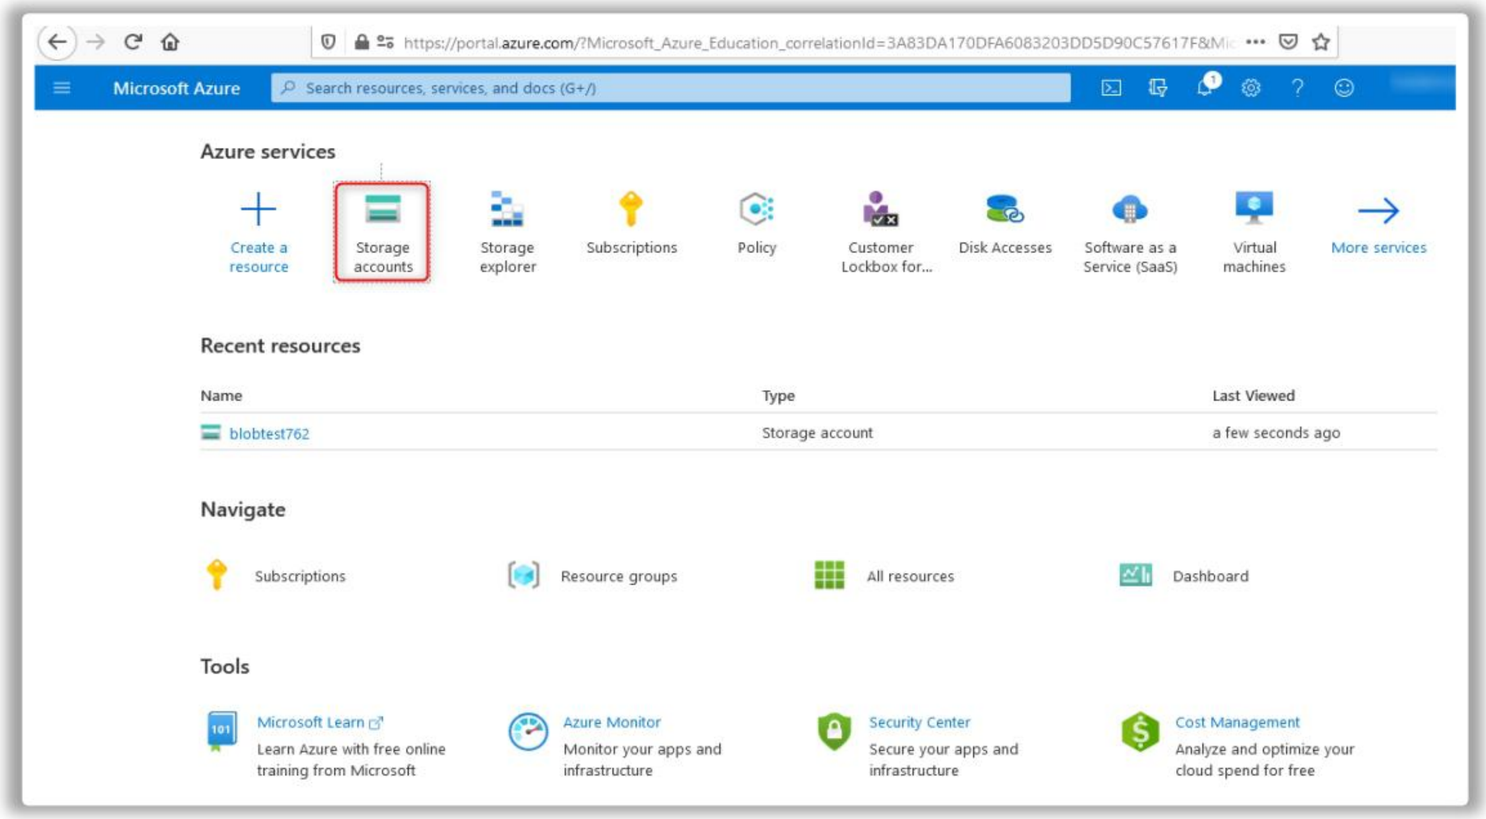


Figure 4.285: Select Storage Accounts

- A list of storage accounts will be displayed.
- Select the storage to enable soft delete for containers.

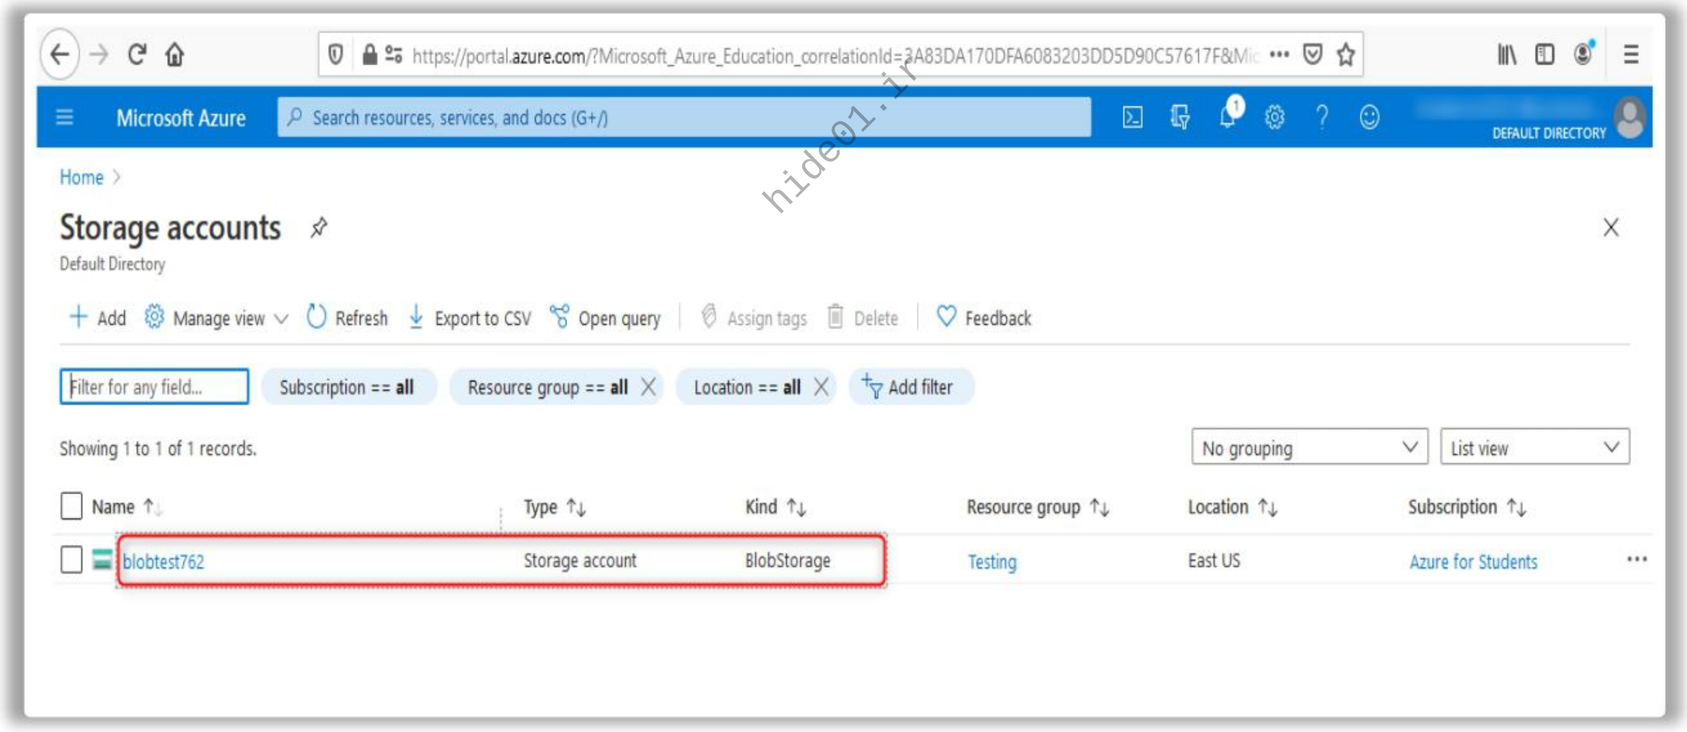


Figure 4.286: Choose Storage

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

- Locate **Blob service** and navigate to **Data protection**.

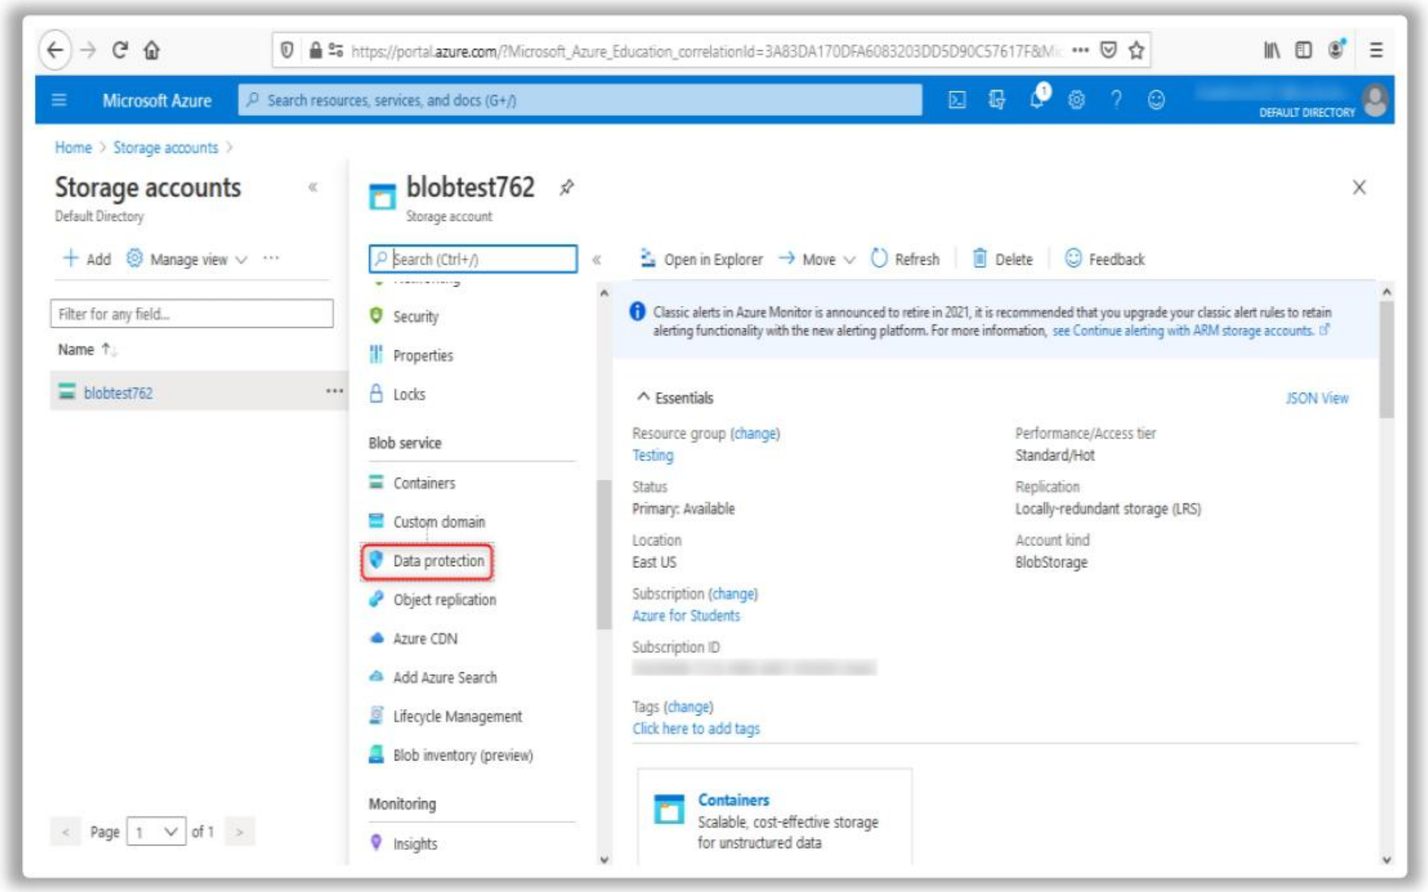


Figure 4.287: Select Data Protection

- Apart from soft delete for Blob, you will find **Turn on soft delete for container** enable it.
- Click on **Save.**

Note: It can be used under Azure subscription-based.

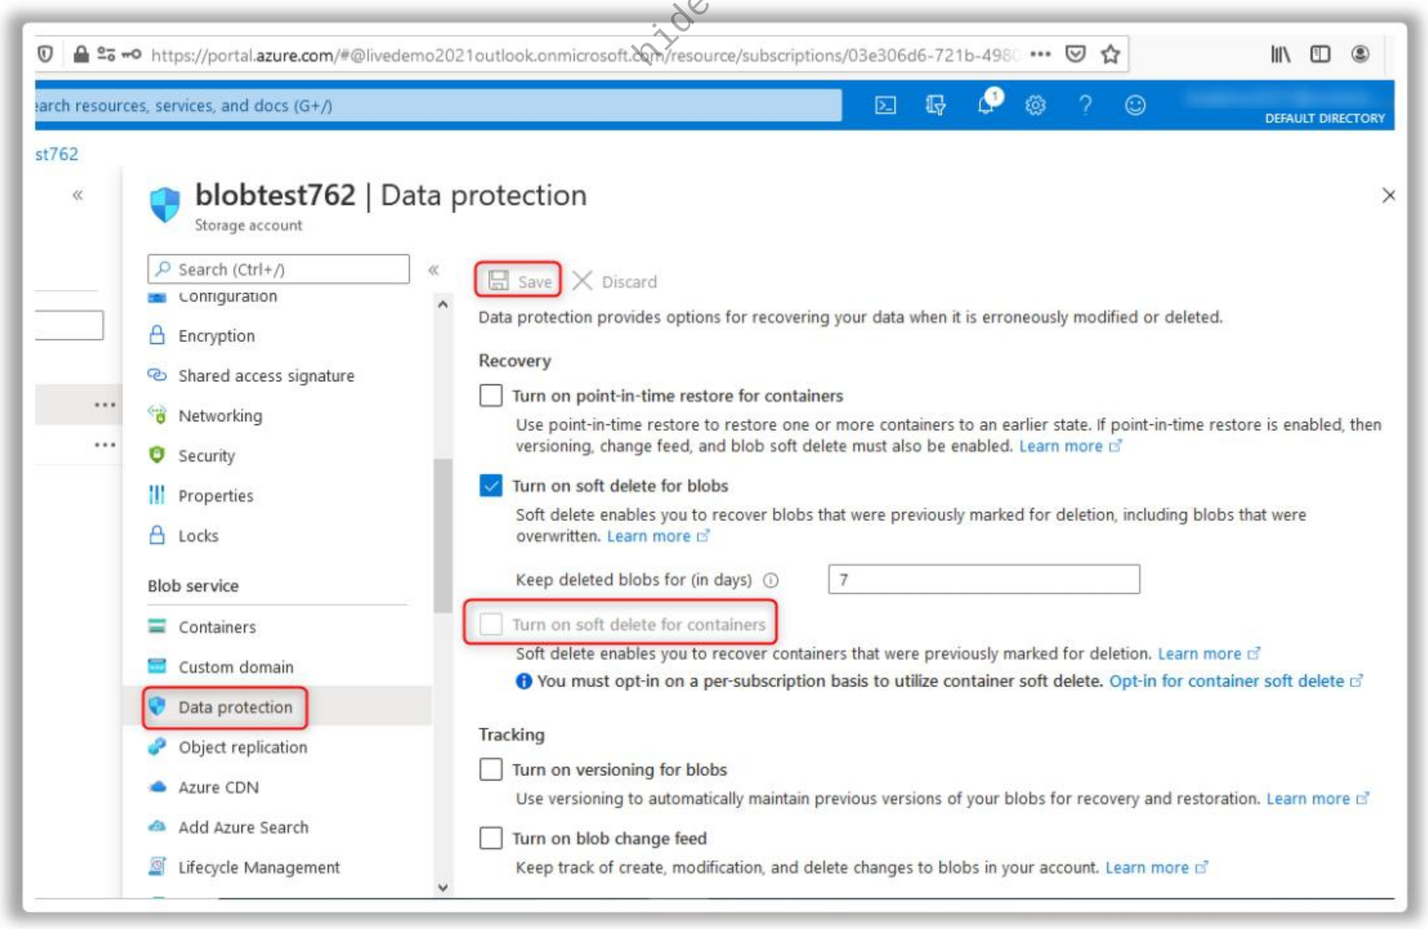
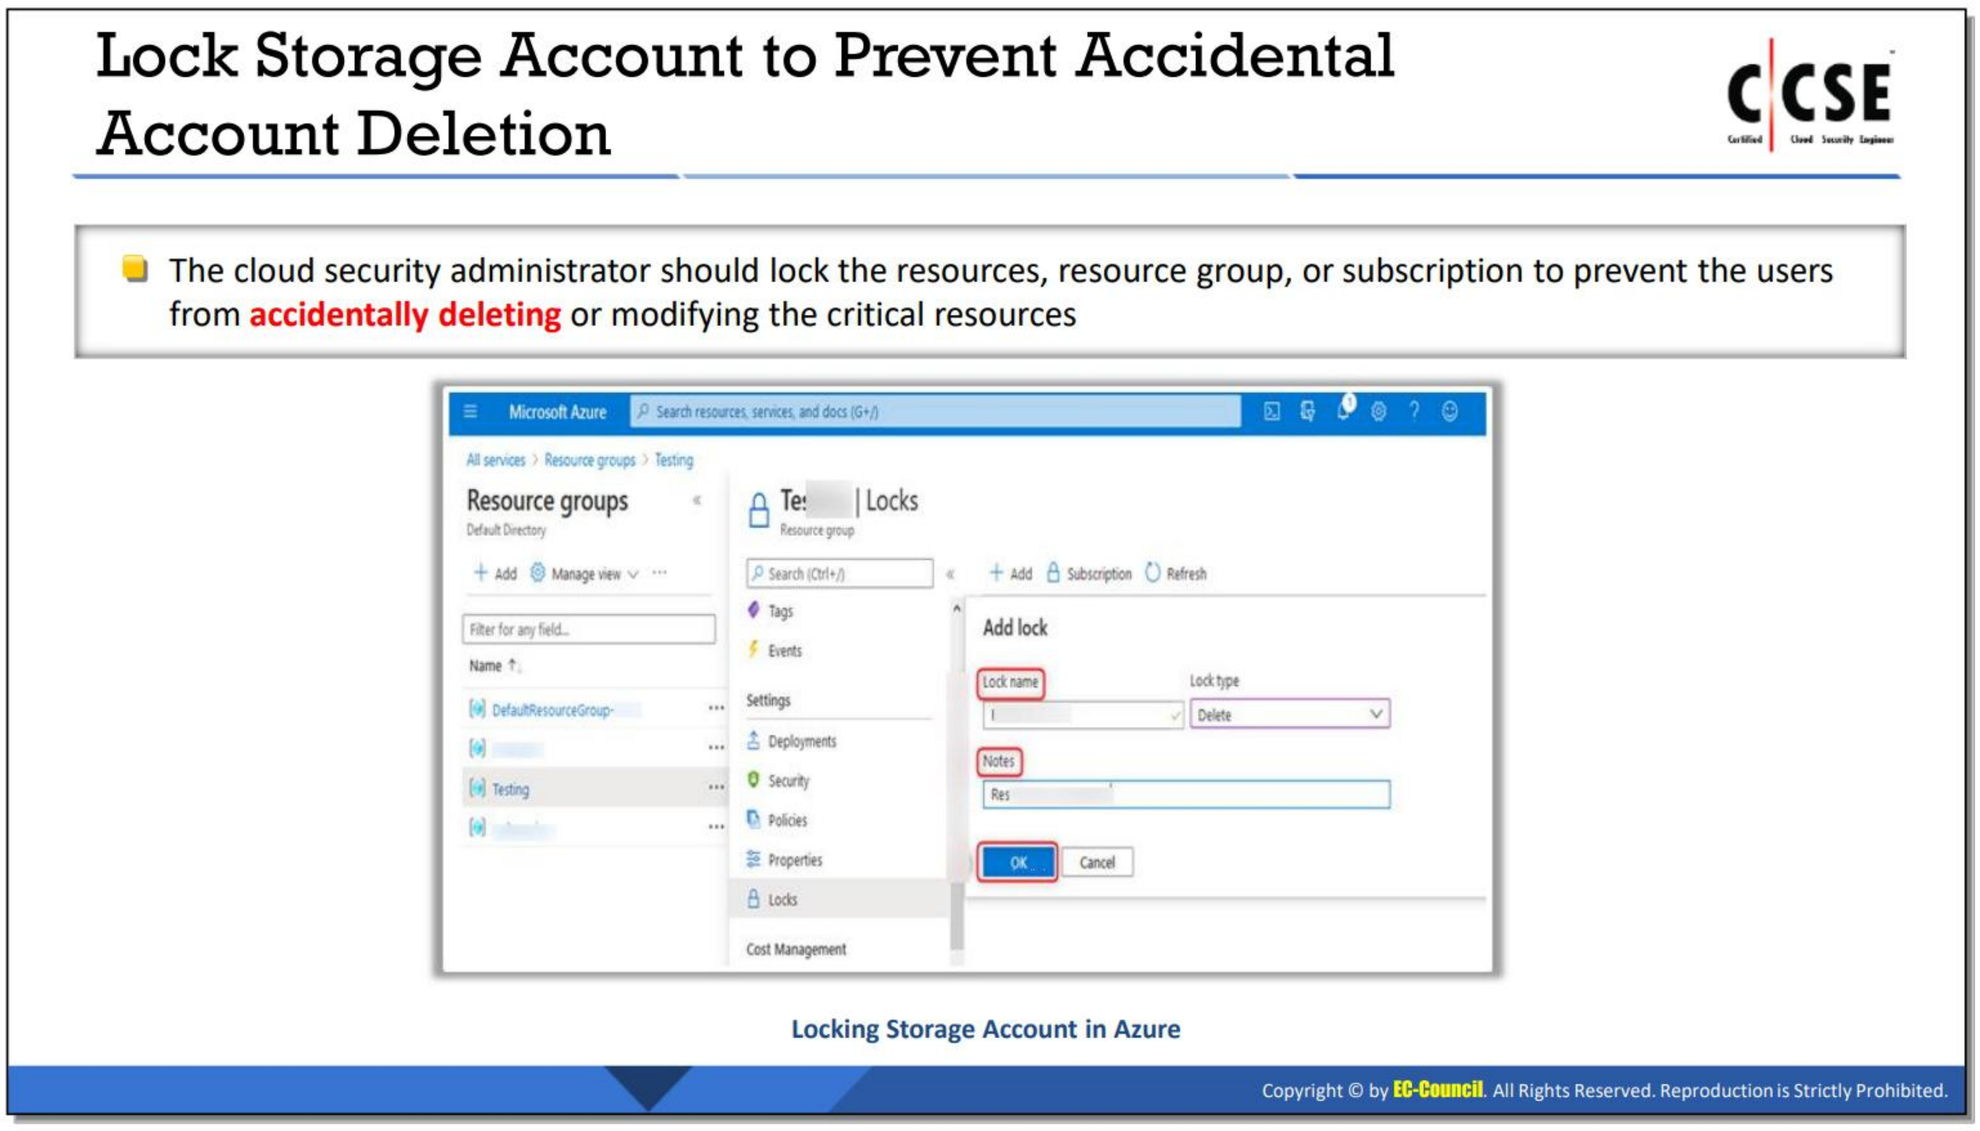


Figure 4.288: Click Save

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



Lock Storage Account to Prevent Accidental Account Deletion

A cloud account is usually handled by multiple users with different kinds of access privileges. An administrator is the one who can control everything and can lock the resource group or resources to prevent accidental deletions. To prevent deletion and to avoid unwanted modifications, an administrator can set read-only access to the users.

If the lock type is set to Delete, users having access can still read and modify data, but cannot perform delete action. If the lock is set to read-only, authorized users can access data to read, but they cannot modify or delete it.

Source: https://azure.microsoft.com

To Lock Storage Account to Prevent Accidental Account Deletion

- Sign-in to Azure portal.
- Click on **More services.**

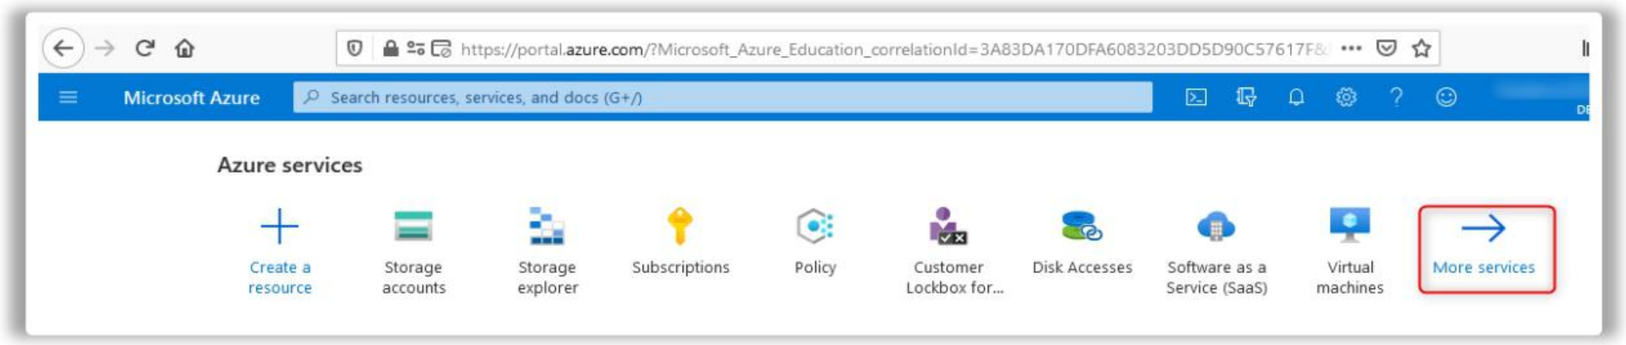


Figure 4.289: Select More Services

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

- Under **Featured**, locate **Resources groups** and click on it.

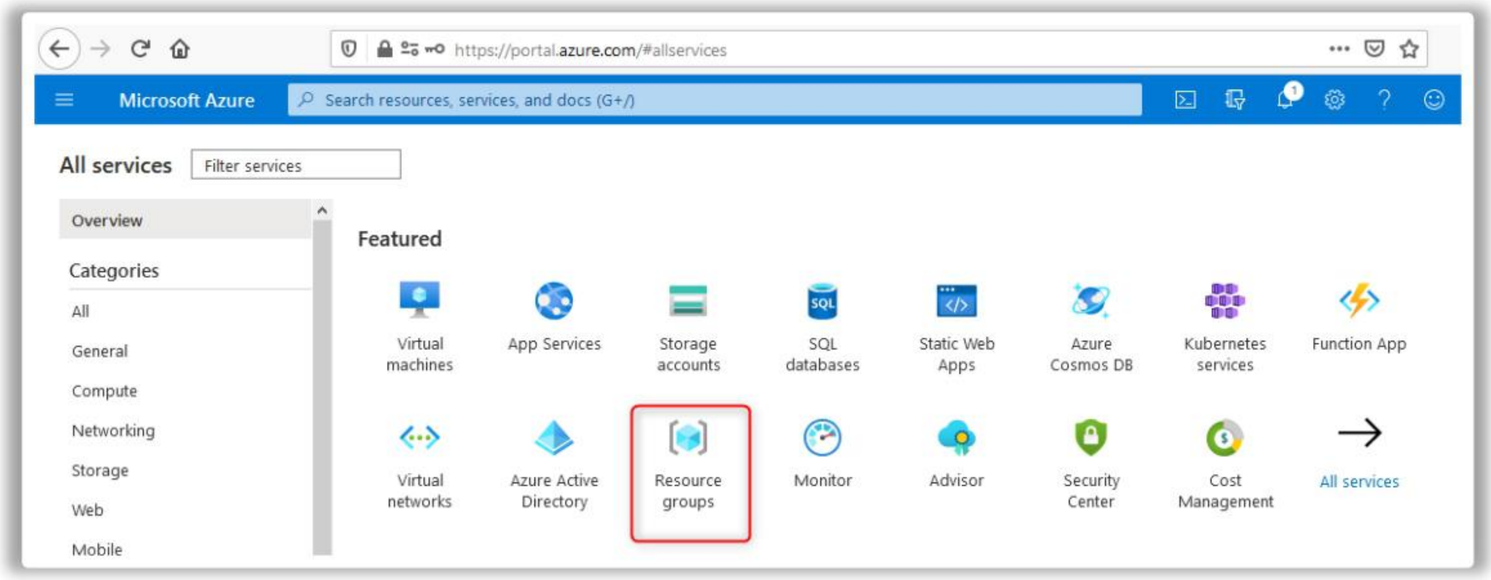


Figure 4.290: Click Resource Groups

- All the services in that resources will be listed.

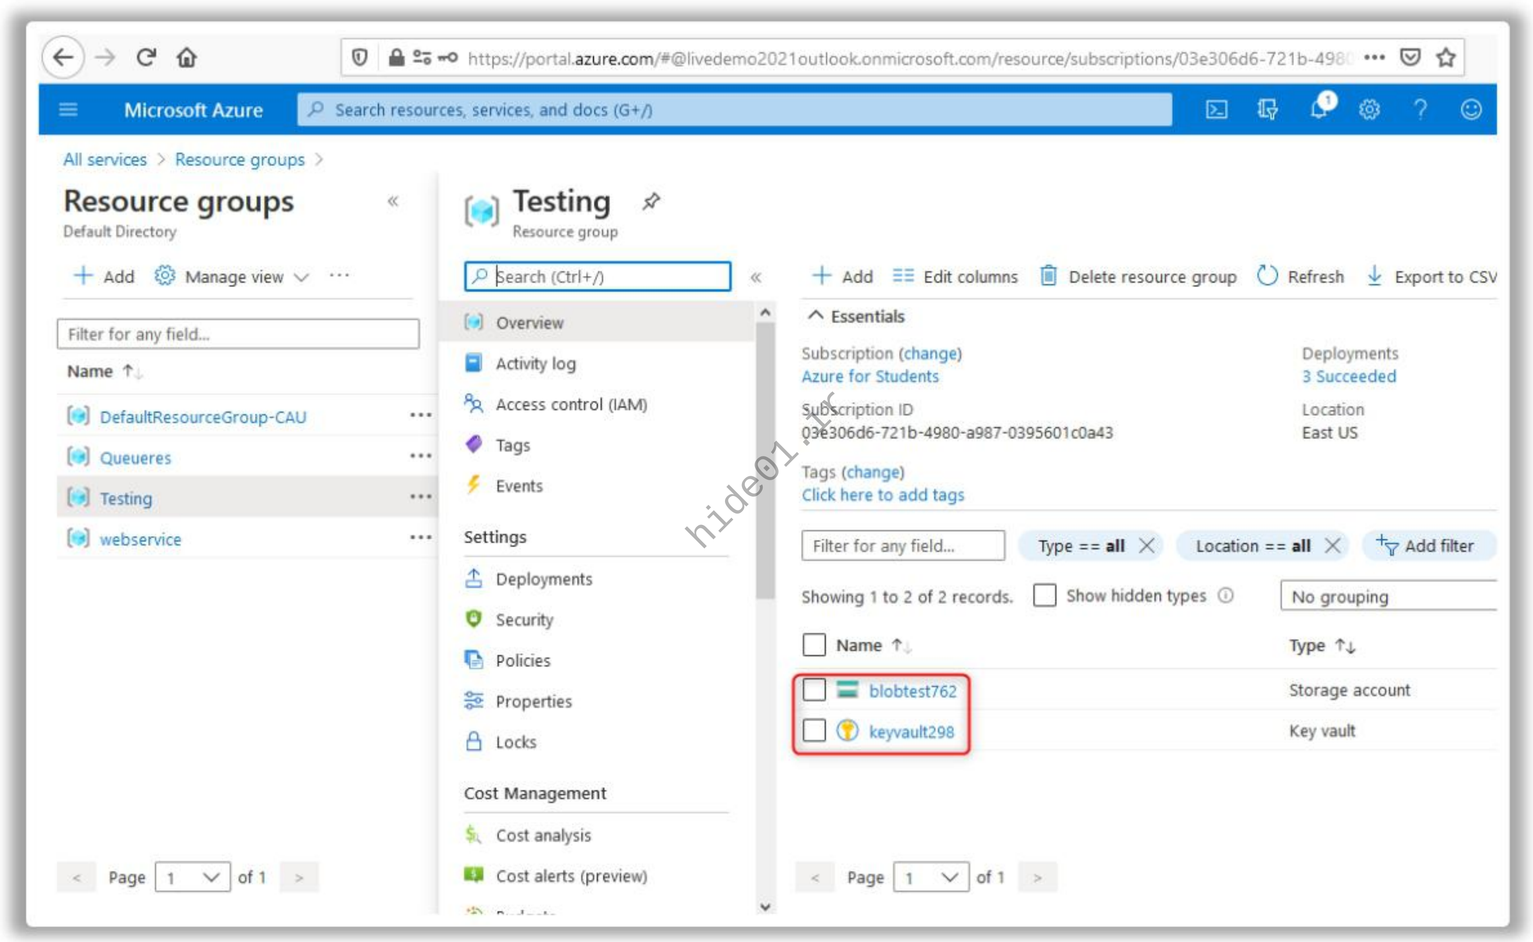


Figure 4.291: Choose Resource Group

- In the navigation panel, under **Settings** go to **Locks.**
- Click on **Add lock.**

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

- Select **Read-only** or **Delete** from the dropdown of the **Lock type.**

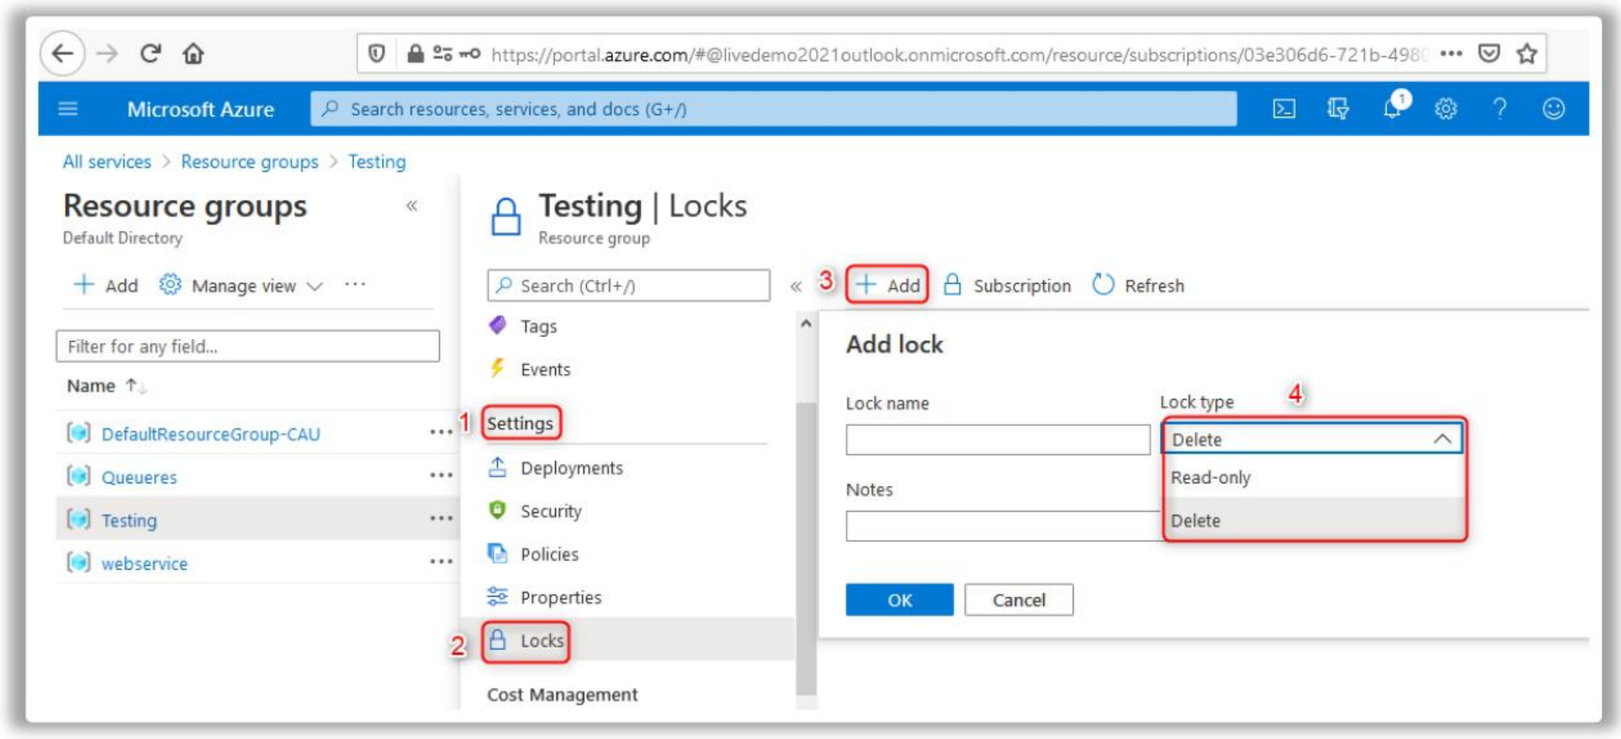


Figure 4.292: Select Settings

- Enter a unique **Lock name**.
- Enter information related to the above lock in **Notes**.
- Click on **OK.**

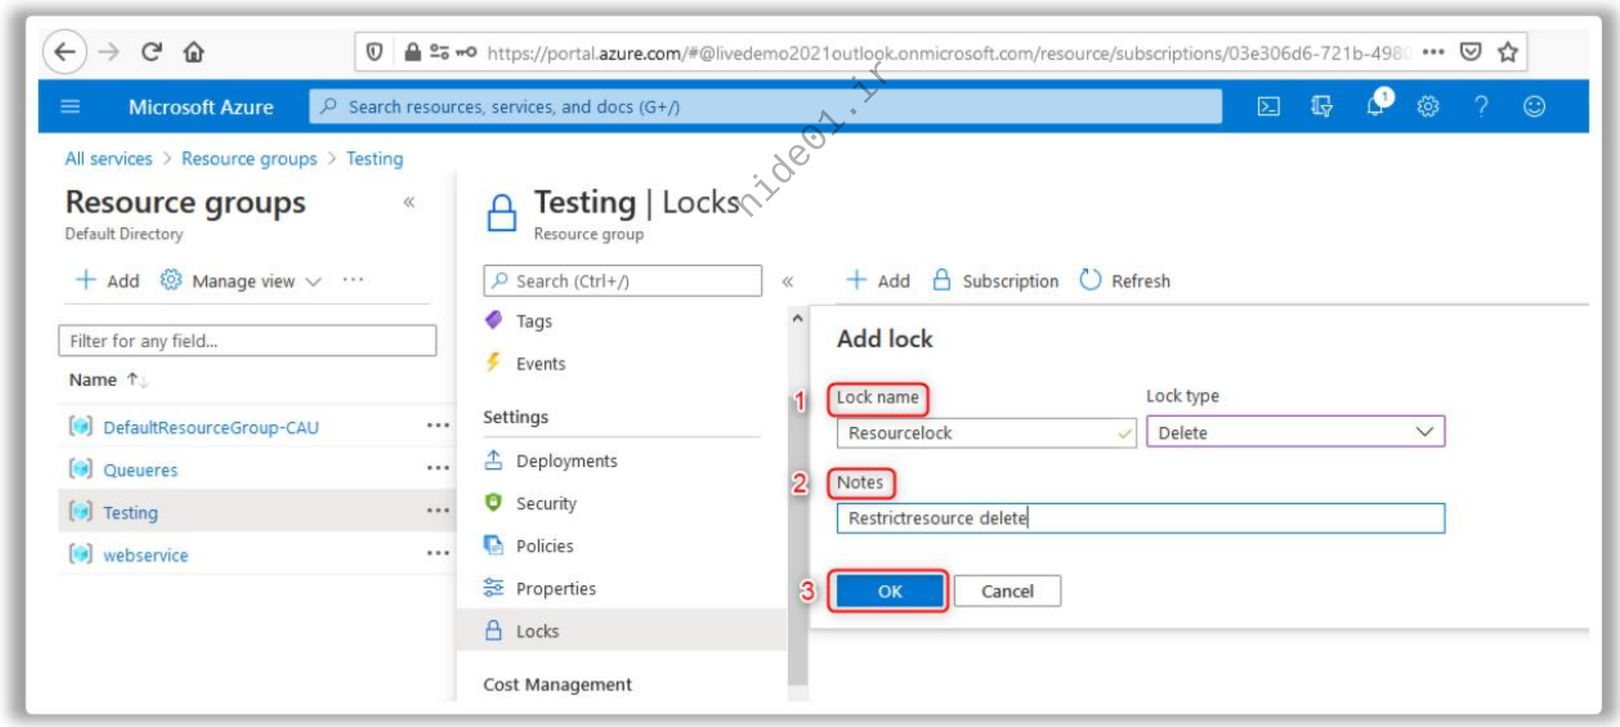


Figure 4.293: Click Ok

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

- The lock has been successfully created.

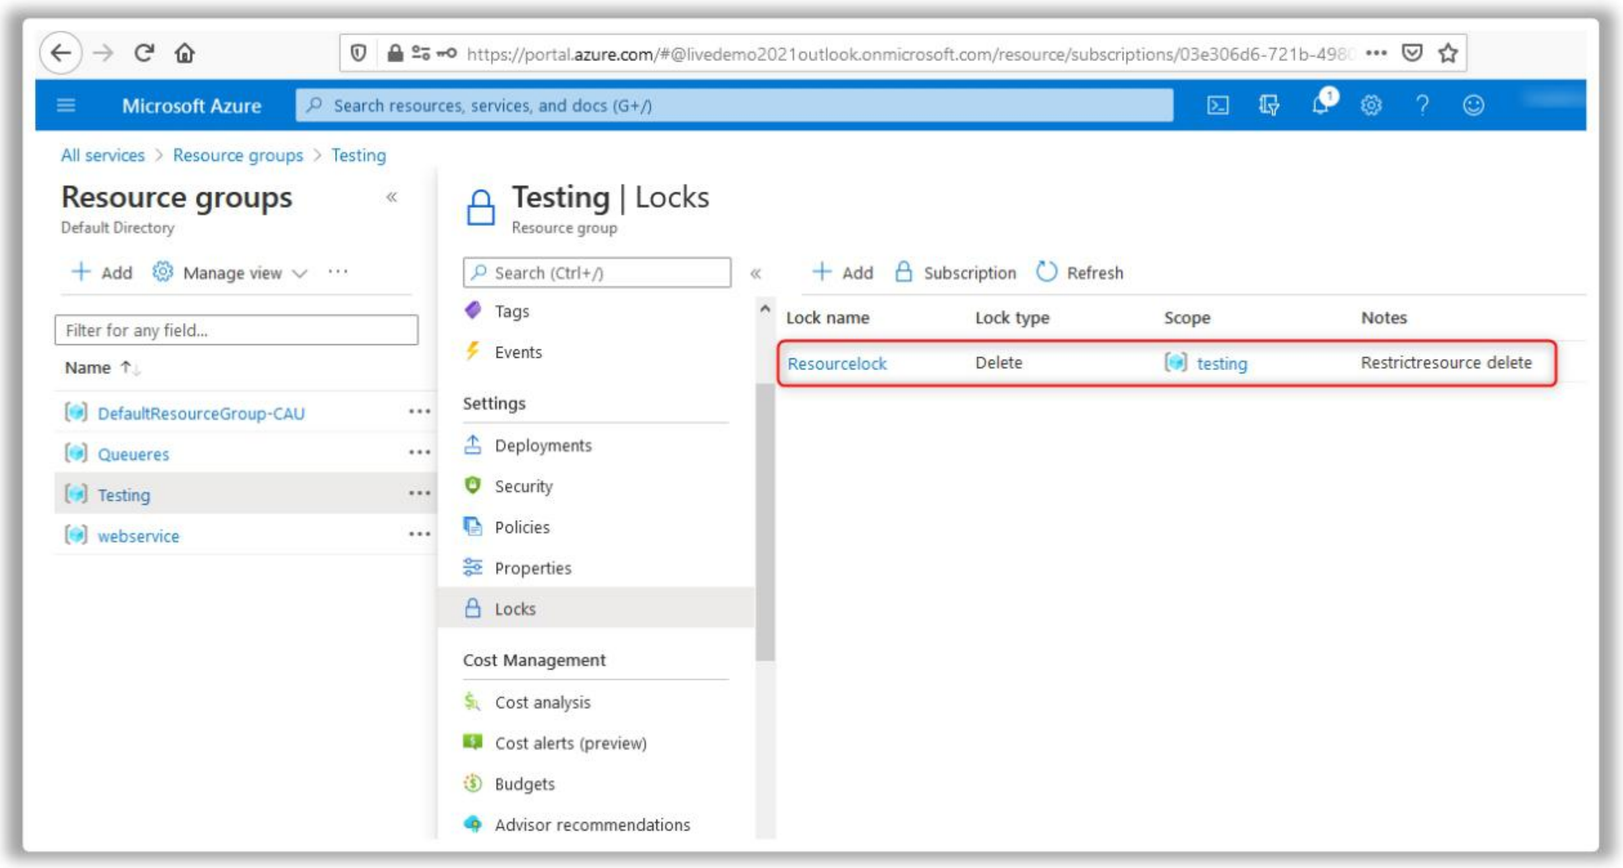
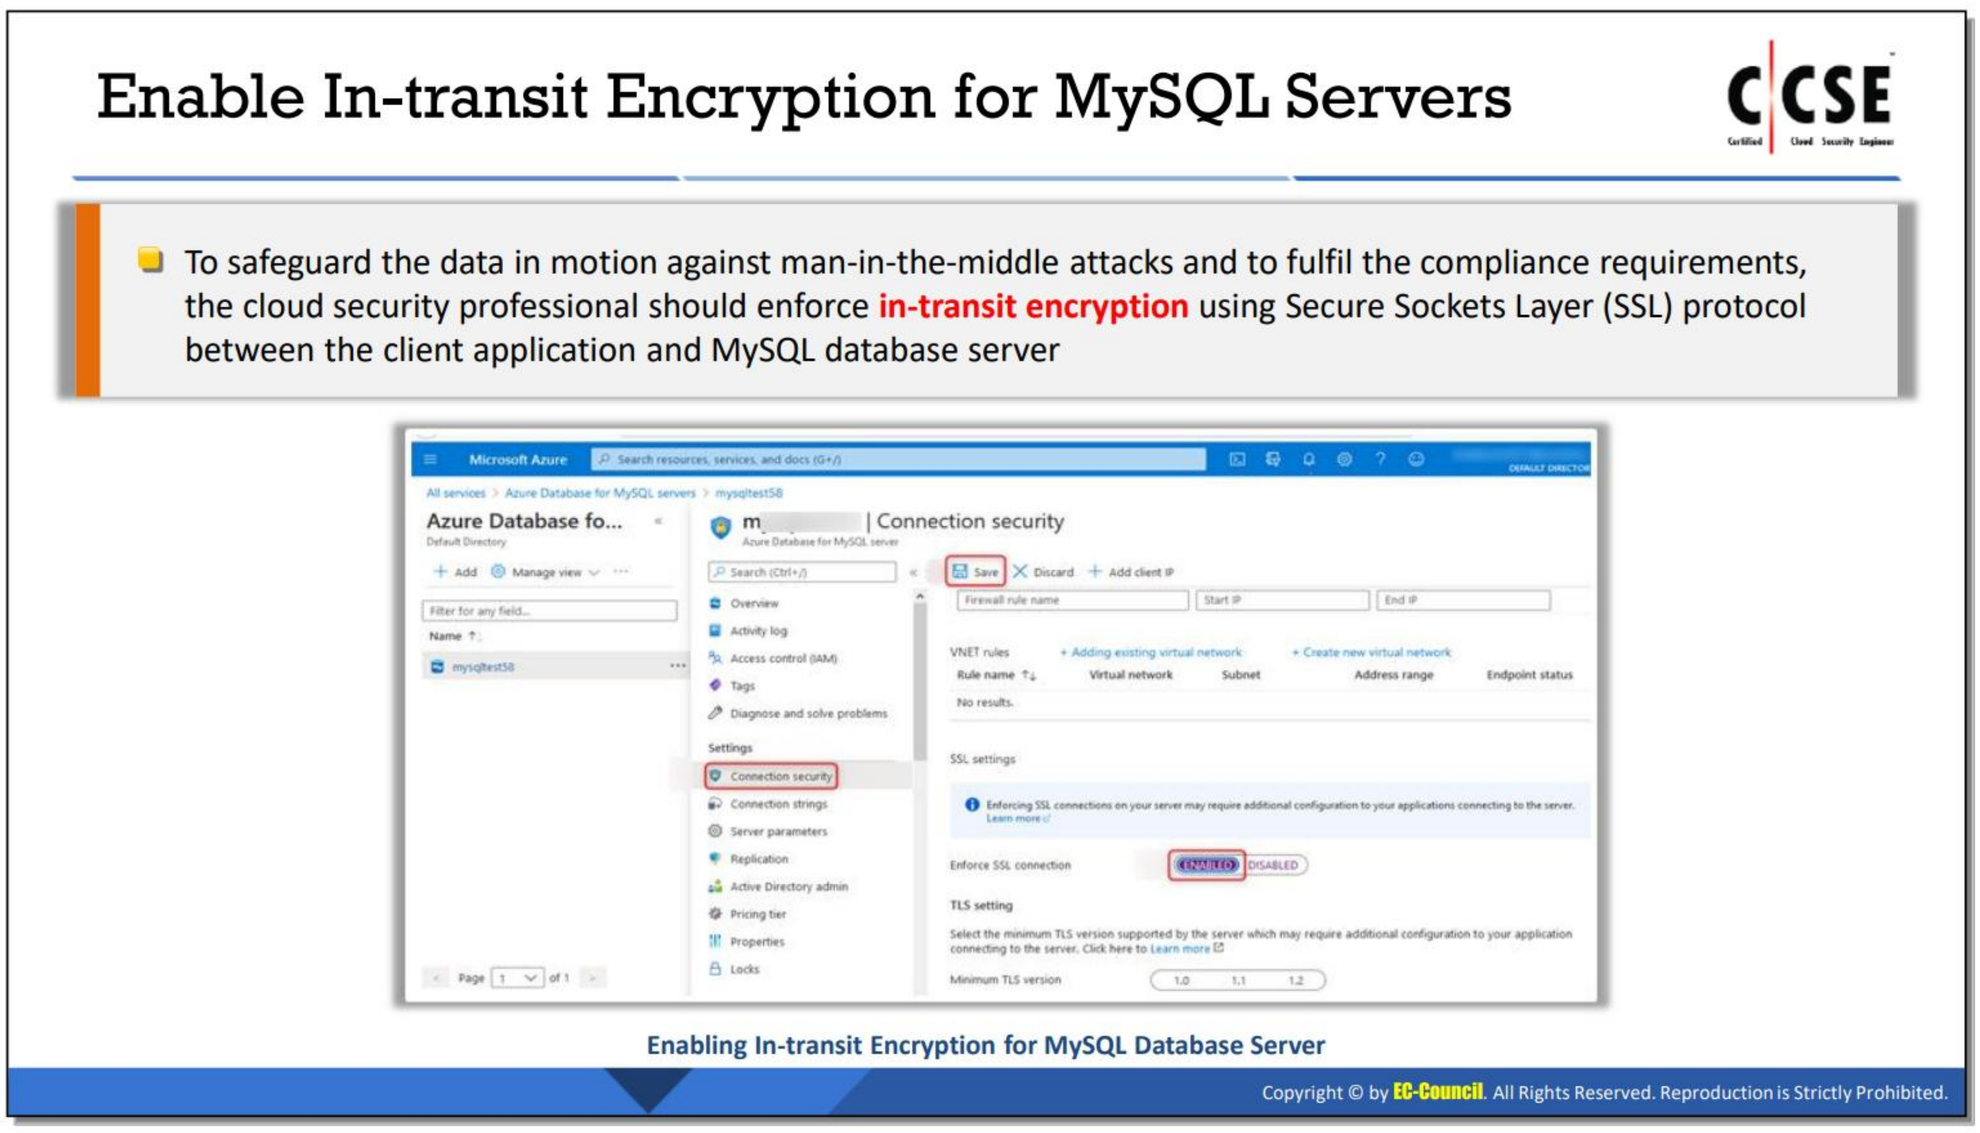


Figure 4.294: Lock is Created

hide01.ir

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



Enable In-Transit Encryption for MySQL Database Servers

To safeguard the data-in-motion against MIM attacks and to fulfill the compliance requirements, the cloud security professional should enforce in-transit encryption using SSL protocol between the client application and MySQL database server. The in-transit encryption on Microsoft Azure MySQL database servers safeguards the sensitive information from unauthorized users.

Source: https://azure.microsoft.com

To Enable In-Transit Encryption for MySQL Servers

- Sign-in to Azure portal.
- Click on **More services**.

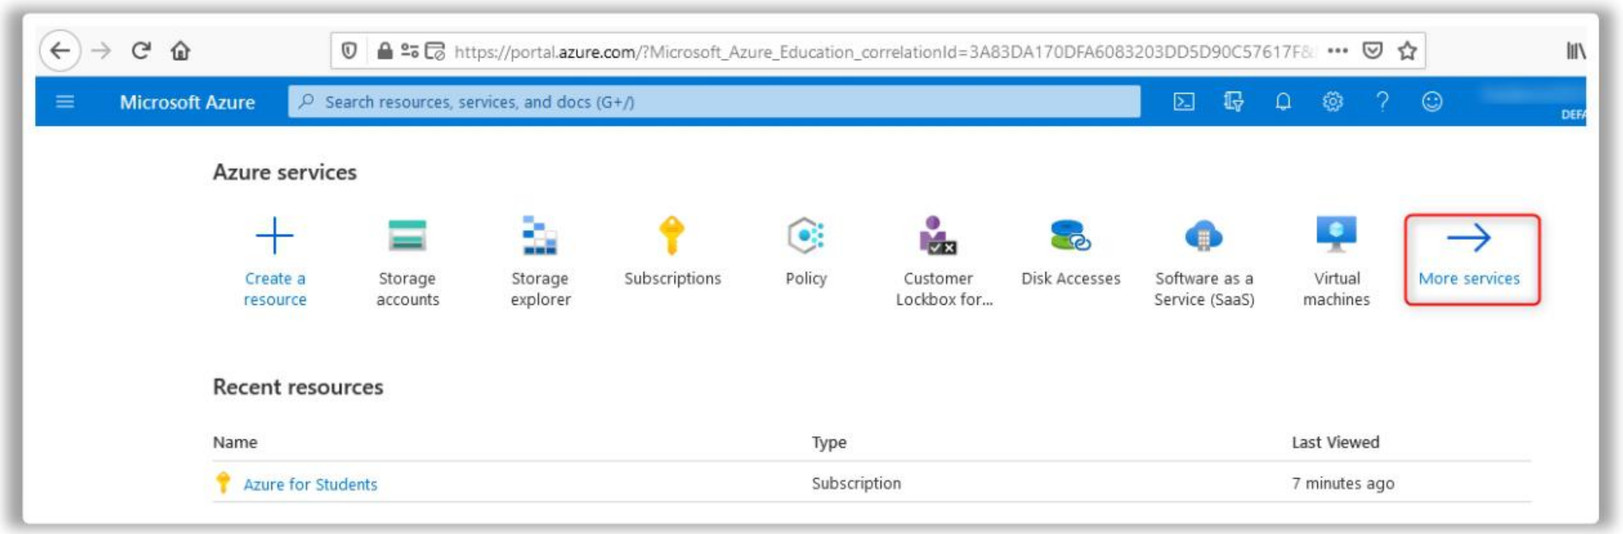


Figure 4.295: Select More Service

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

- Search for **Azure database for MySql servers** using the input field beside **All services**, then click on the same in the right-side panel.

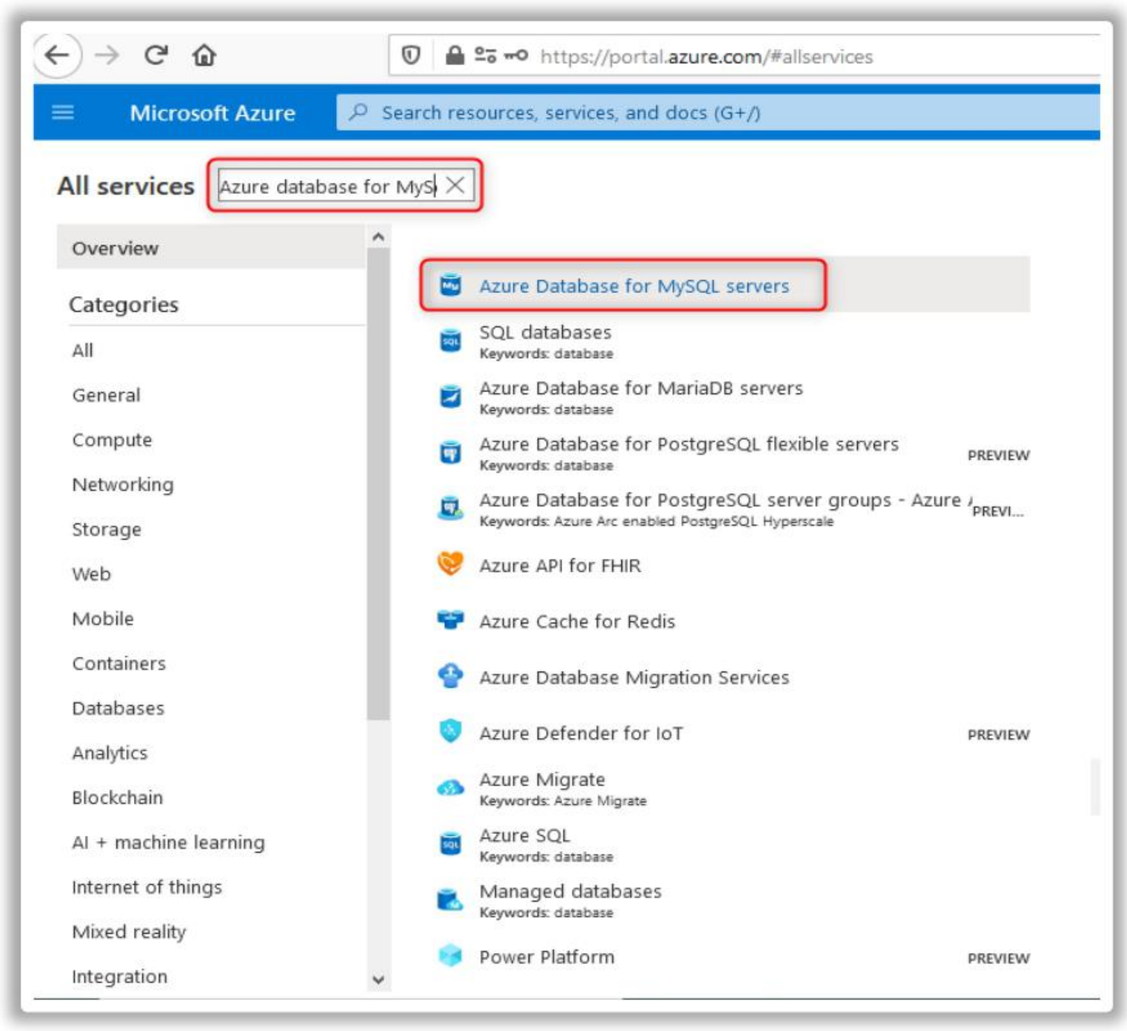


Figure 4.296: Search Azure Database

- A list of MySQL servers will be displayed.
- Select the MySQL server on which you want to enable in-transit encryption.

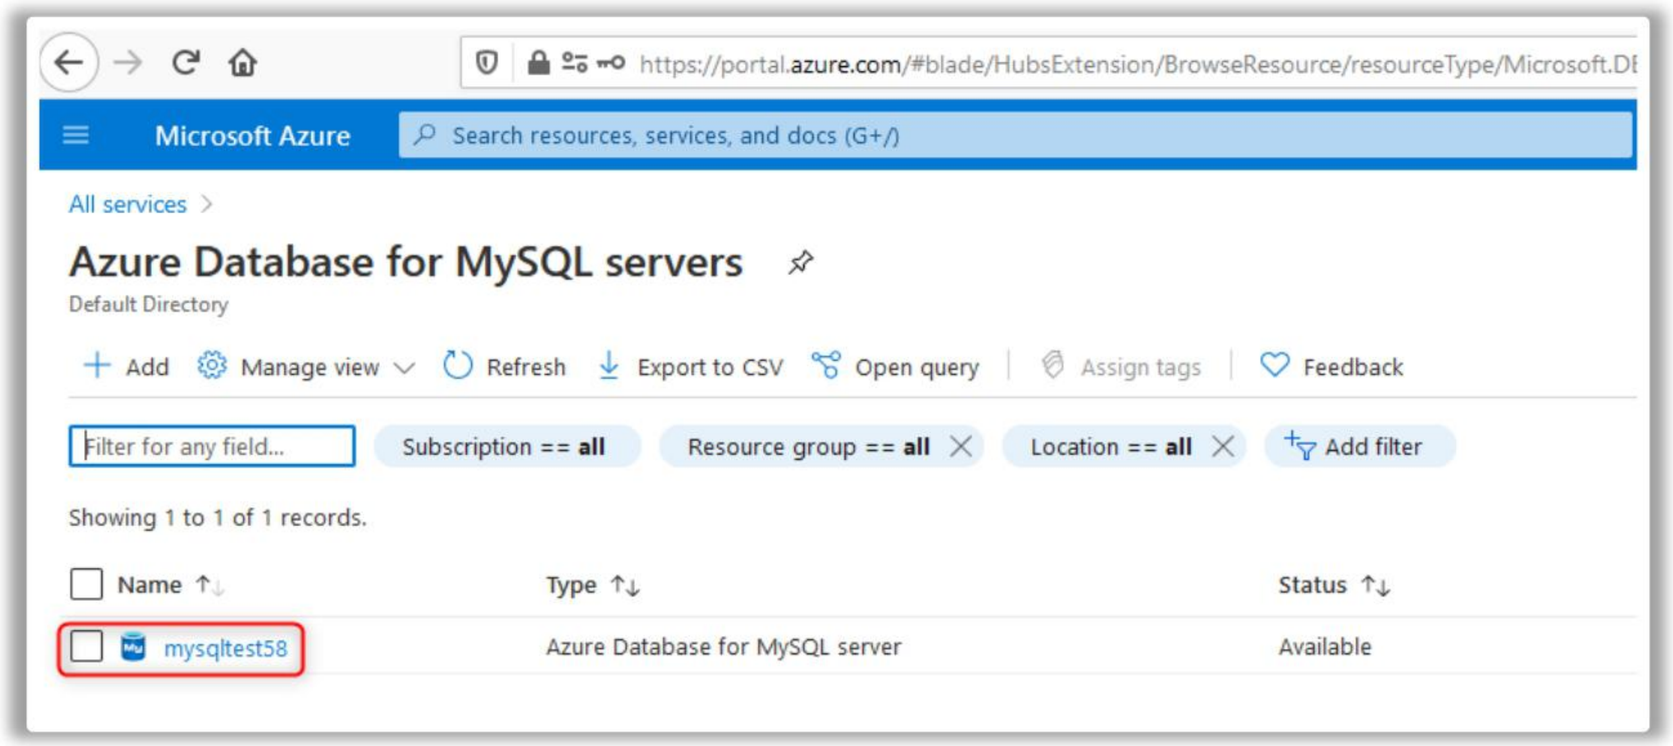


Figure 4.297: Choose MYSQL Server

- Locate **Settings** and navigate to **Connection security**, then click on it.
- Scroll down to **Enable SSL connection** under **SSL settings** and select **Enabled.**

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

- Click on **Save**.

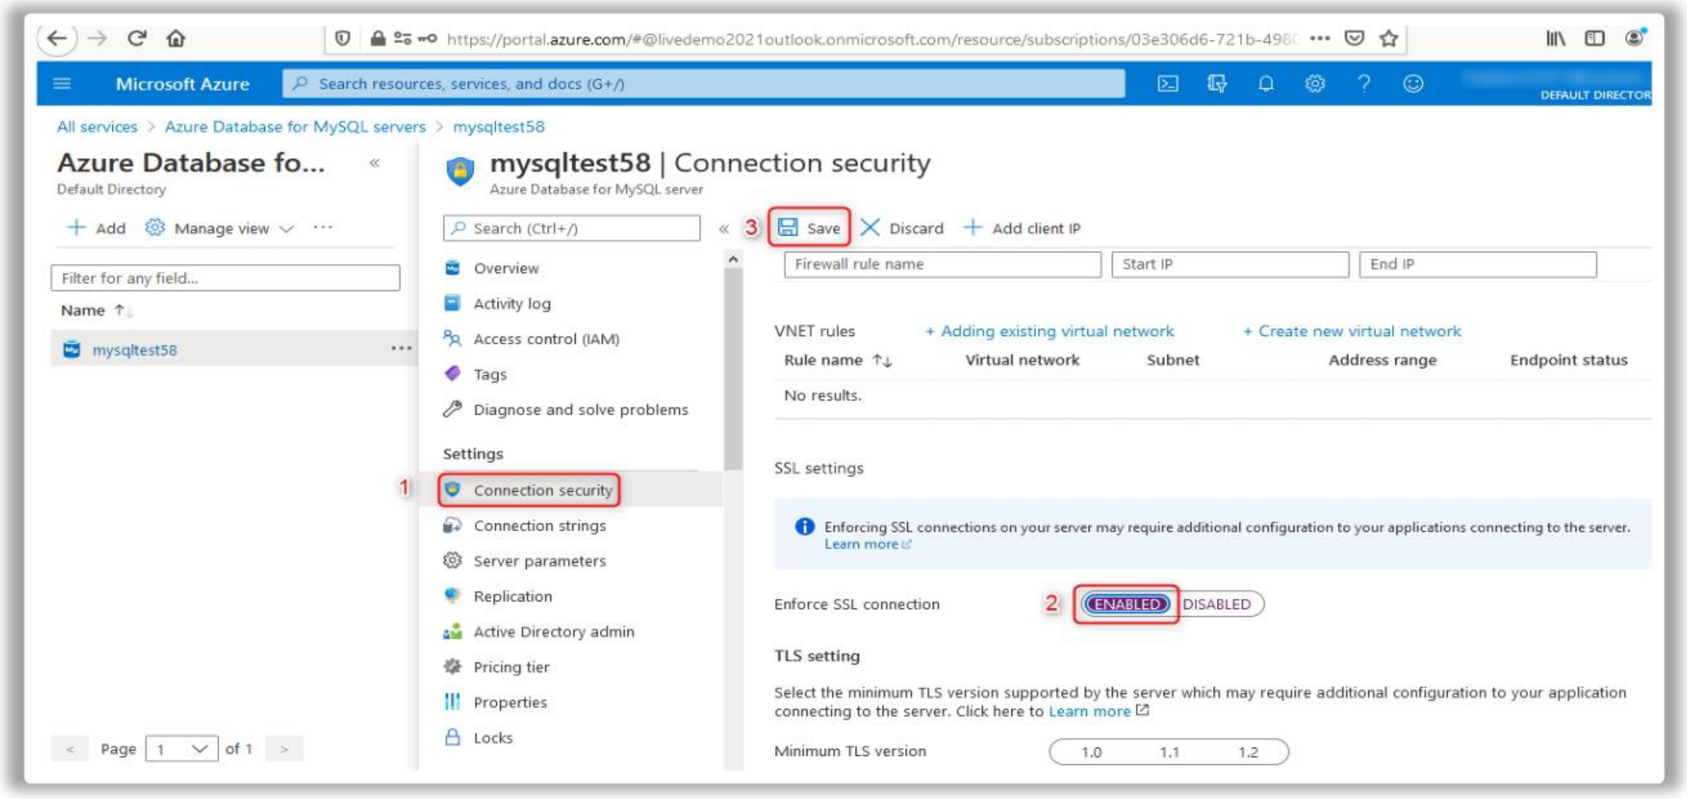


Figure 4.298: Click Save

- The in-transit encryption is enabled successfully.

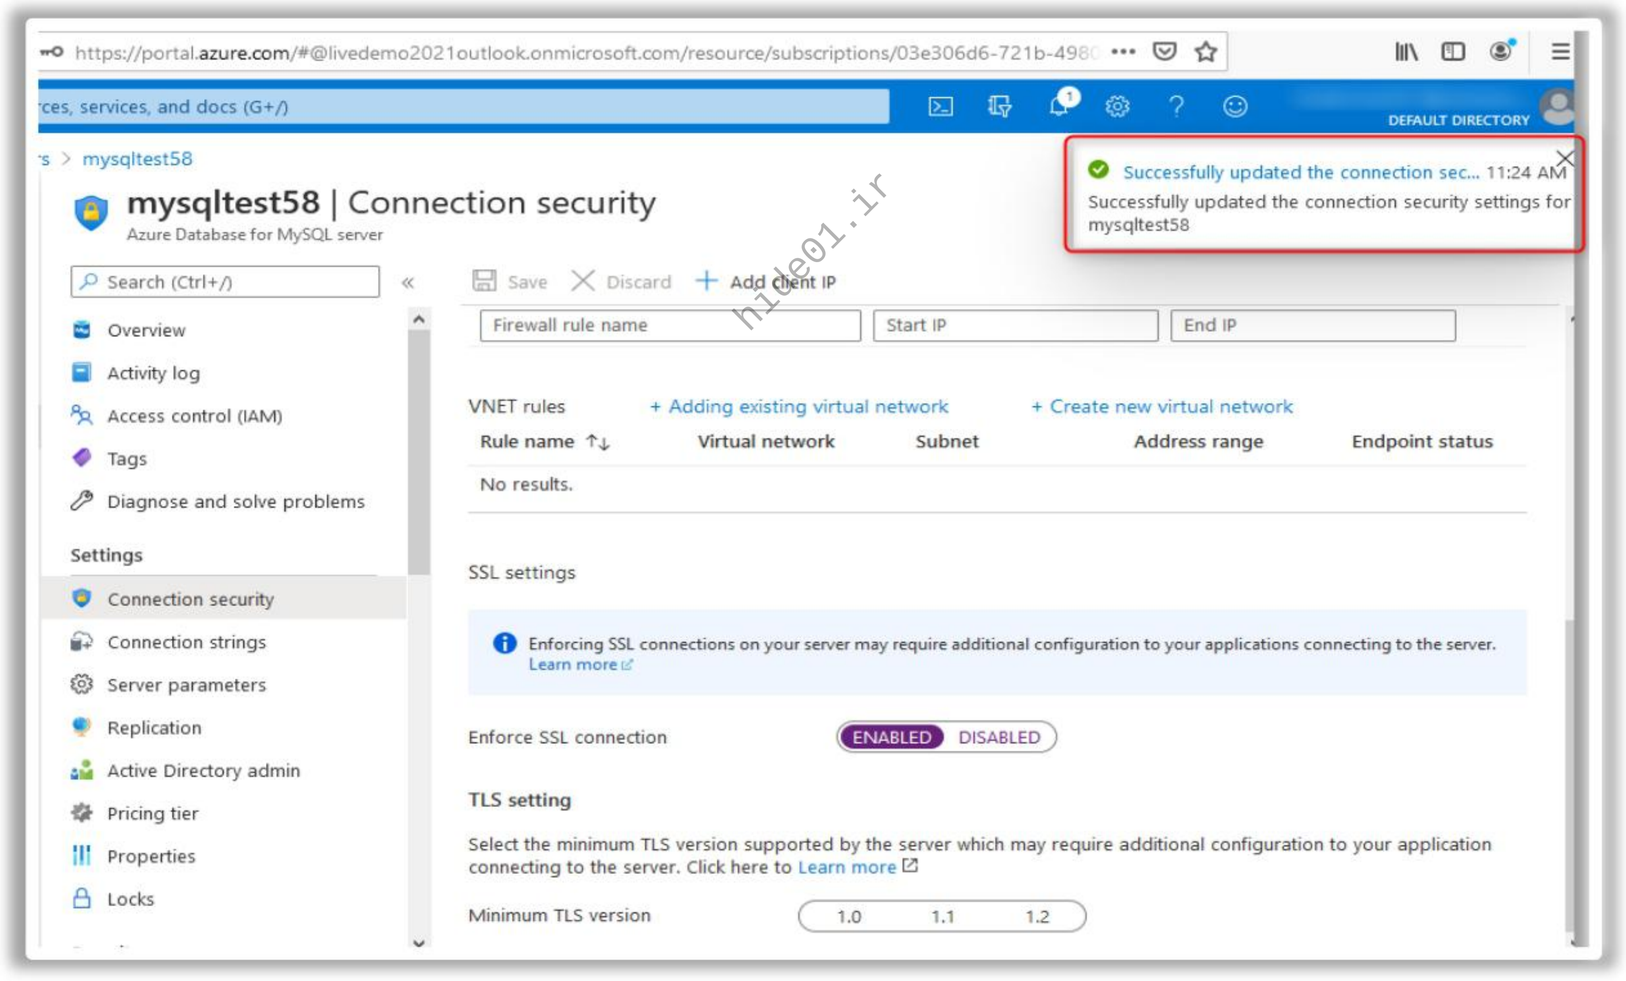
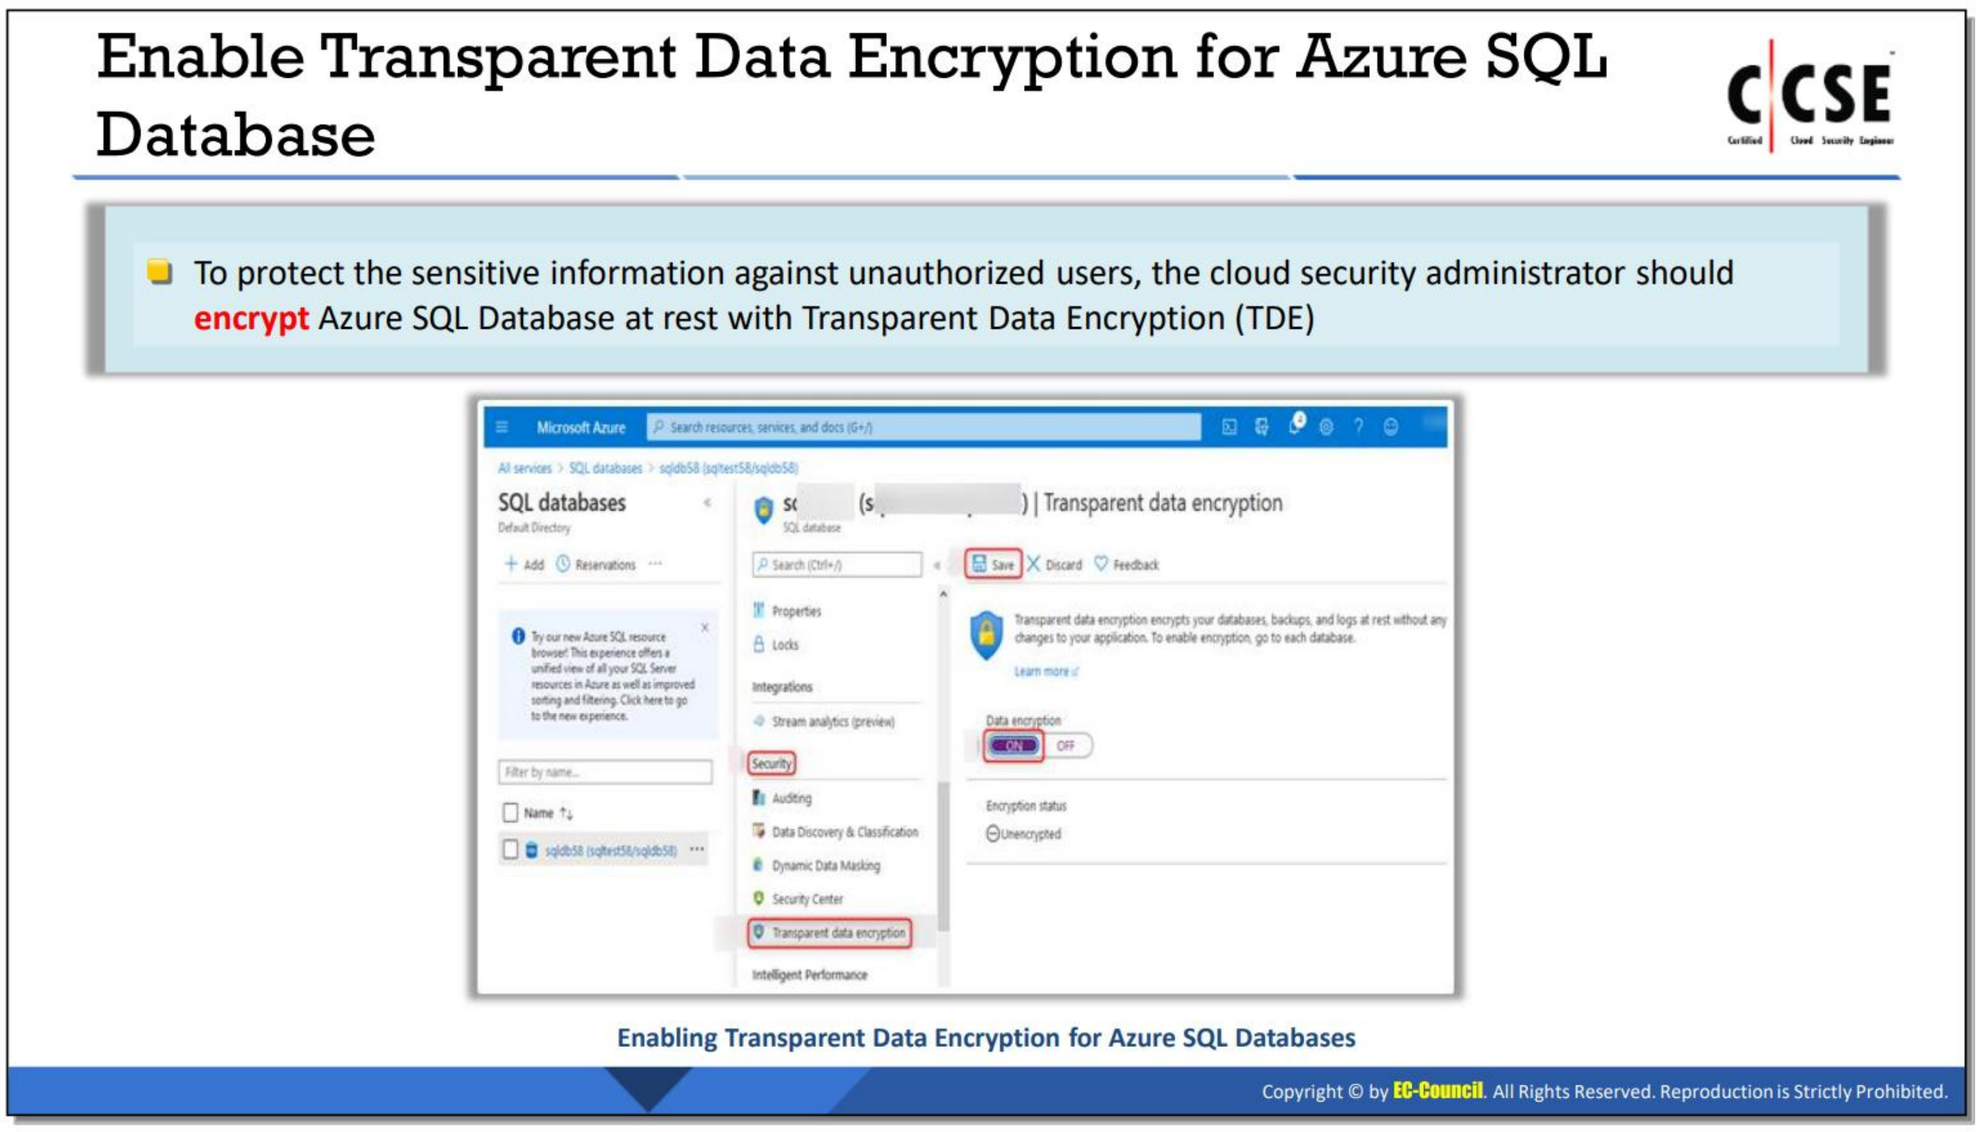


Figure 4.299: Encryption is Enabled

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



Enable Transparent Data Encryption for Azure SQL Database

Transparent Data Encryption (TDE) encrypts services like SQL server and Azure SQL Database. It prevents data loss from malicious attackers by encrypting data-at-rest. It provides real-time encryption and decryption of the databases. This encryption uses a DEK that is a symmetric key. It is stored in the database boot record, so that it can be used during recovery. To protect the sensitive information against unauthorized users, the cloud security administrator should encrypt Azure SQL Database at rest with TDE

Source: https://azure.microsoft.com

To Enable Transparent Data Encryption for SQL Databases

- Sign-in to Azure portal.
- Click on **More services**.

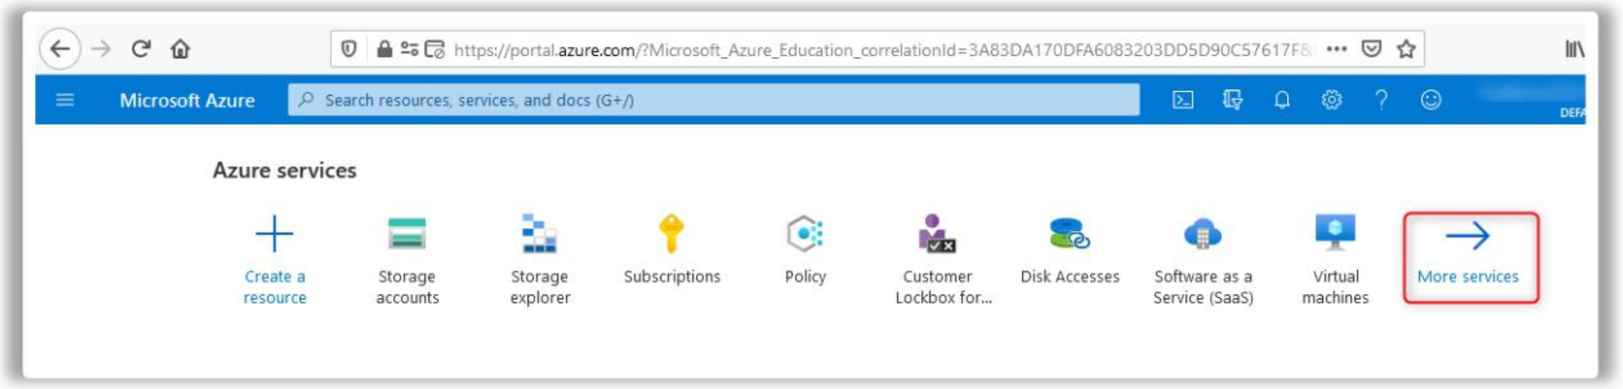


Figure 4.300: Select More Service

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

- Search for **SQL databases** using the input field beside **All services**, then click on the same in the right-side panel.

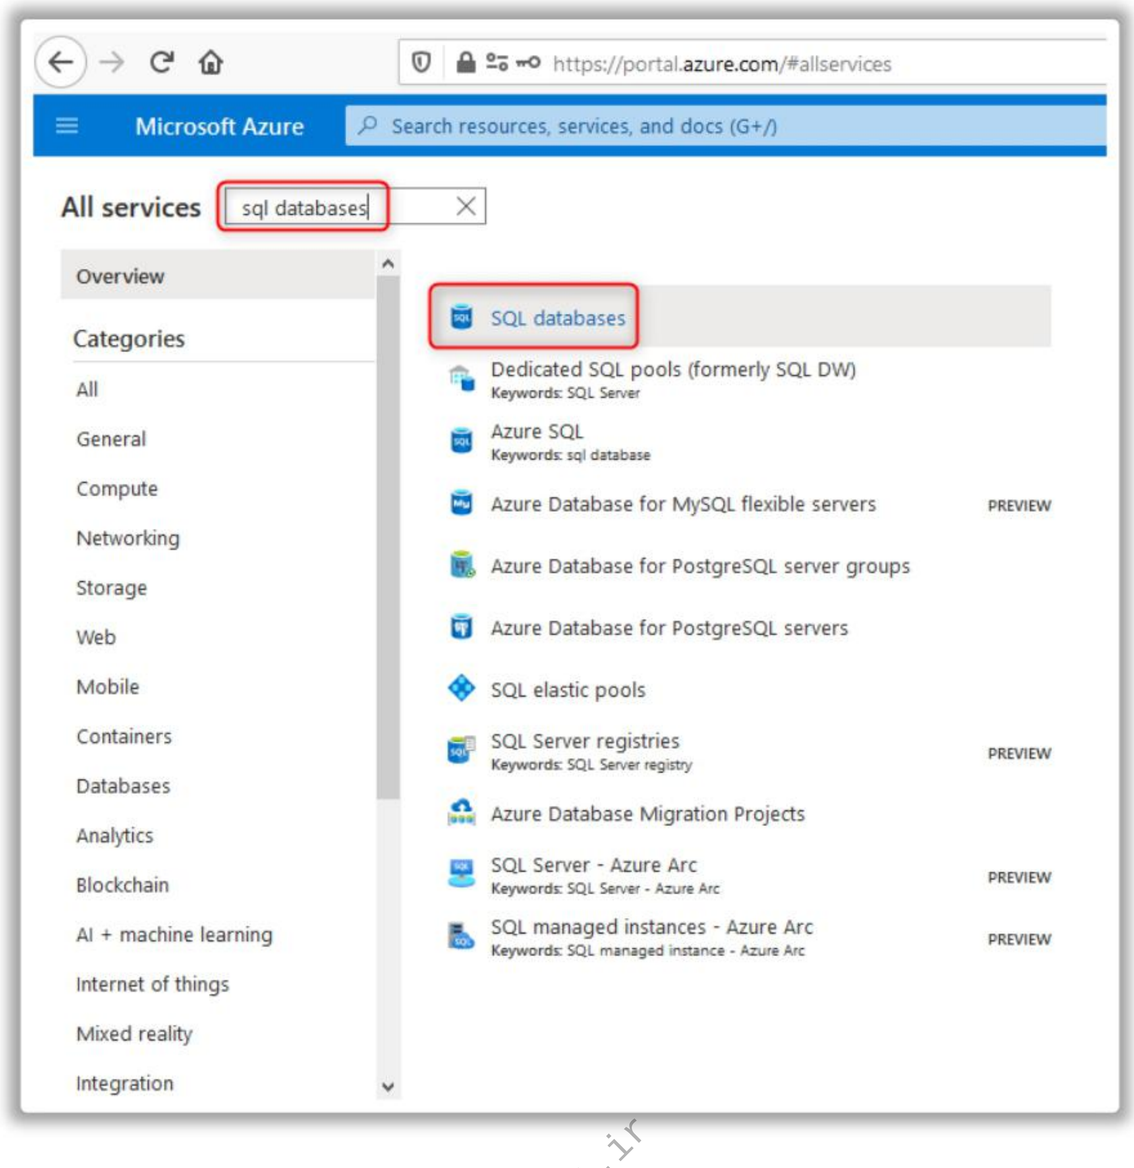


Figure 4.301: Select SQL Databases

- A list of SQL databases will be displayed.
- Select the SQL database in which you want to enable TDE.

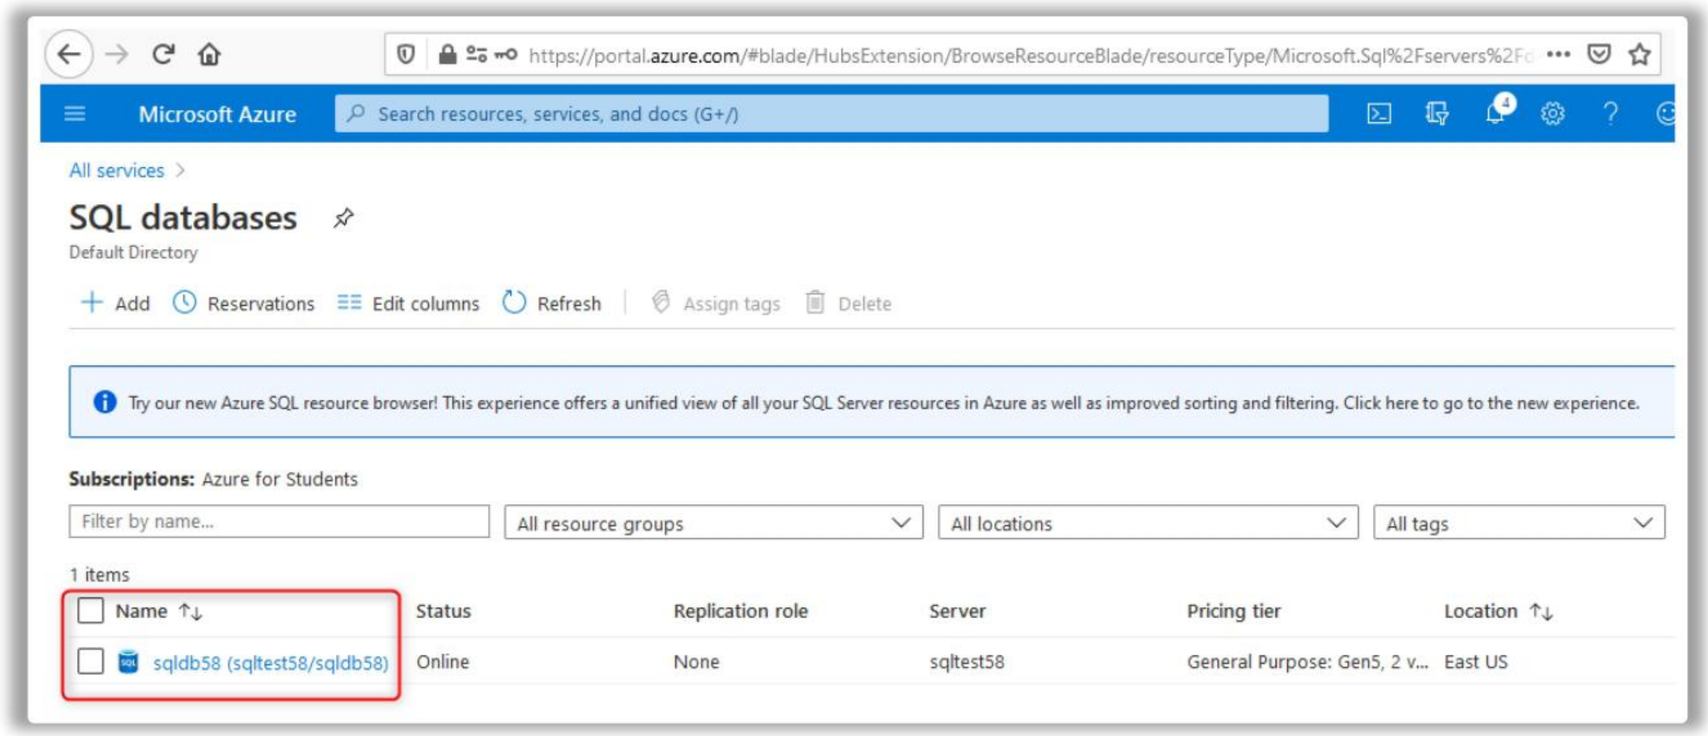


Figure 4.302: Choose SQL Database

- Locate **Security** in the navigation panel and go to **Transparent data encryption**.
- Under **Data encryption,** set it to **ON**.

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

- Click on **Save**.

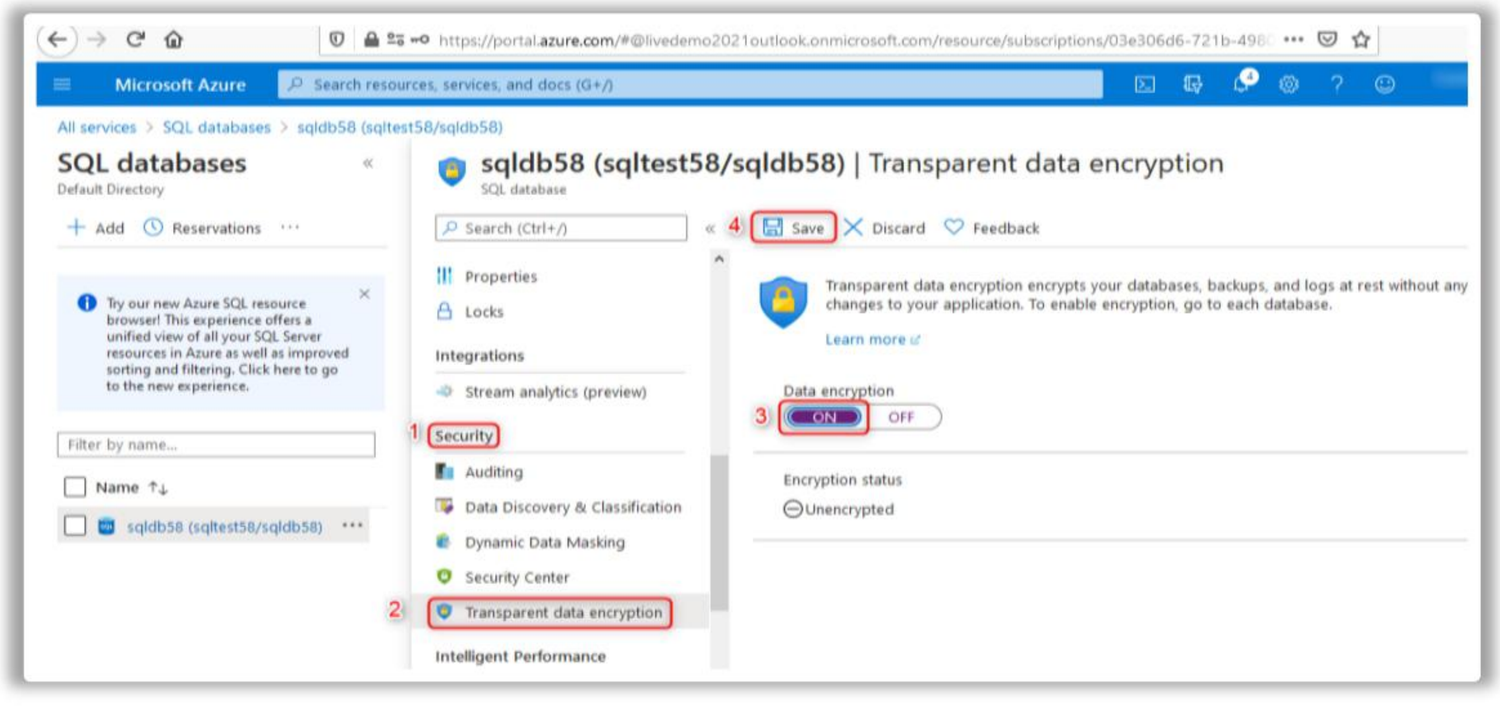


Figure 4.303: Click Save

- TDE is successfully enabled and the status is changed to **Encrypted.**

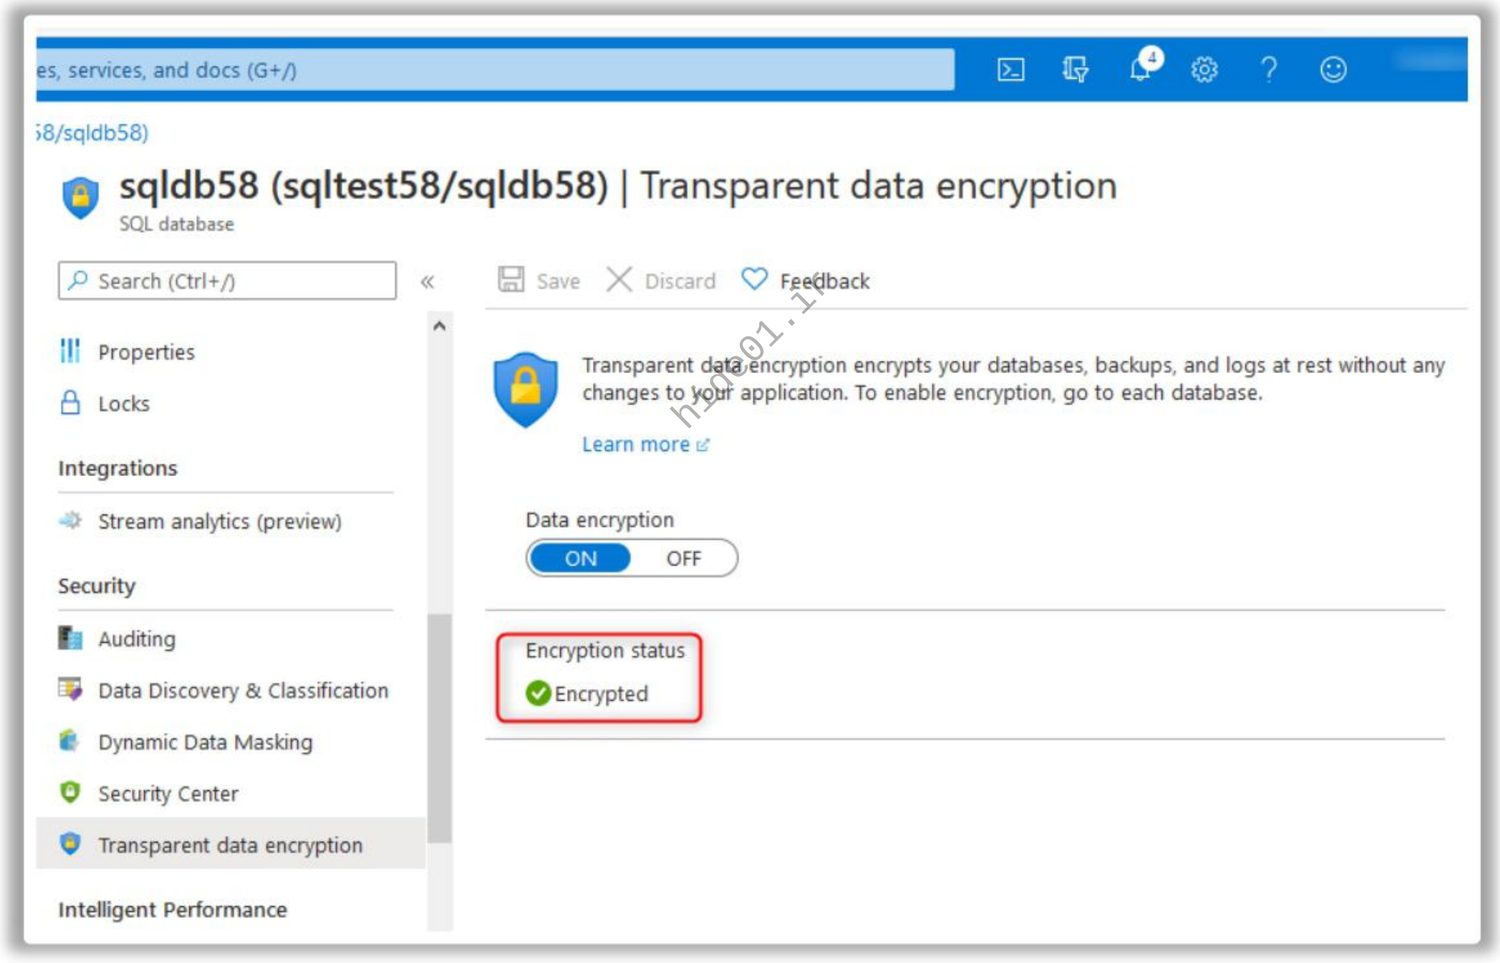
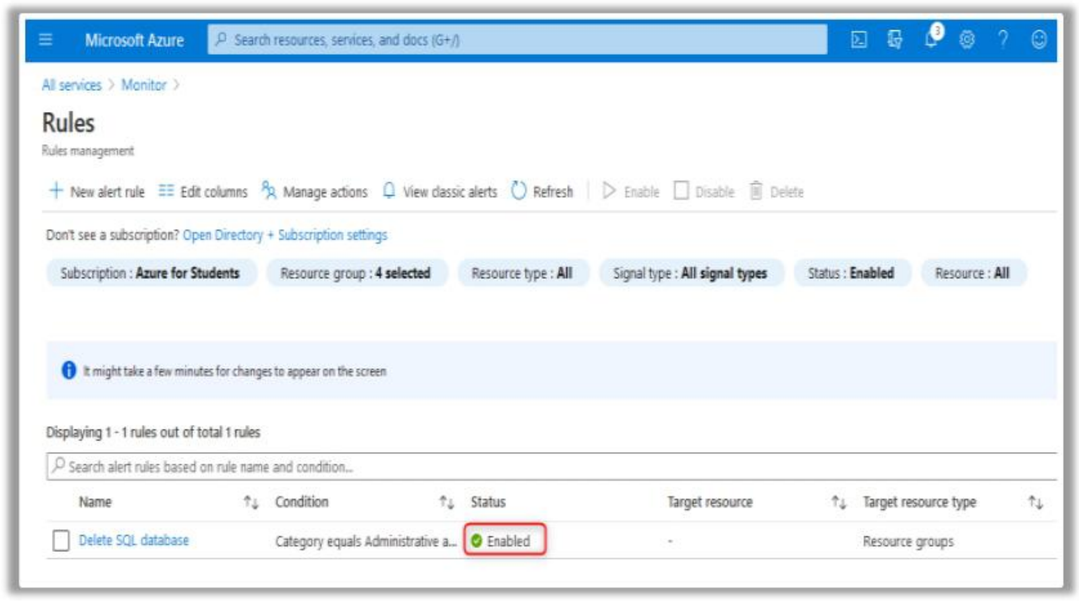


Figure 4.304: Encryption is Enabled

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

Create Alert for "Delete Azure SQL Database" Events

- To enhance Azure SQL database security and availability, the cloud security professional should create **alerts** for "Delete Azure SQL Database" events, which are used in prevention of accidental or intentional deletion of Azure SQL databases



Creating Alert for Delete Azure SQL Database Events

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

Create Alert for "Delete Azure SQL Database" Events

It is highly recommended to create alerts for critical resources. If the delete Azure SQL Database rule is created, when someone tries to delete the Azure SQL Database it triggers the alert when it meets the condition of "Category='Administrative', Signal name='Delete Azure SQL Database (Microsoft.Sql/servers/databases)." Alerts logs are useful in the prevention of accidental or intentional deletion of Azure SQL databases.

Source: https://azure.microsoft.com

To Create Alert for "Delete Azure SQL Database" Events

- Sign-in to Azure portal.
- Click on **More services**.

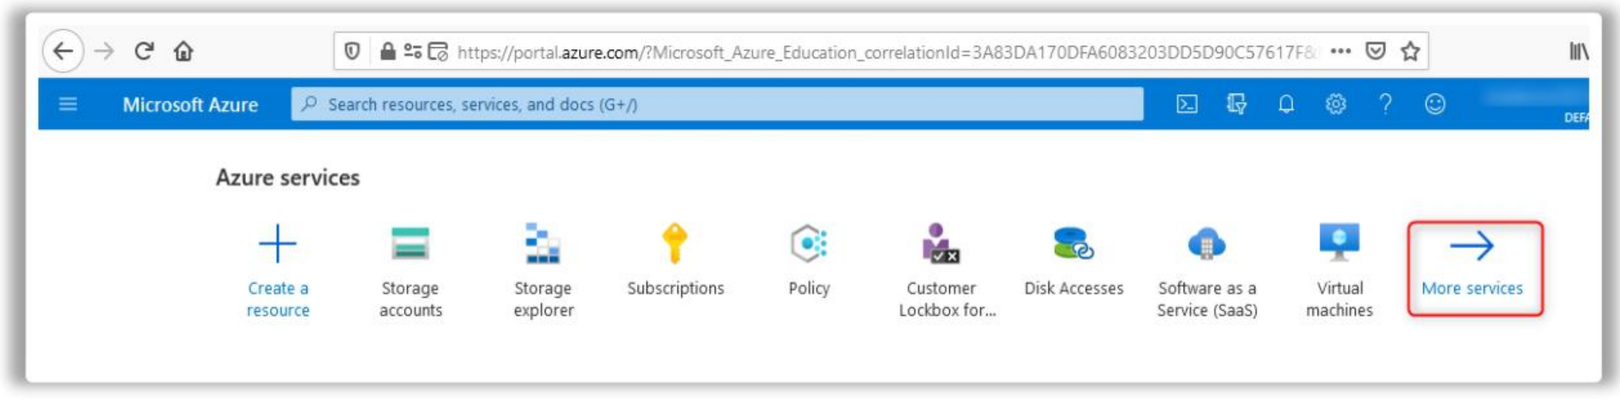


Figure 4.305: Select More Service

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

- Under **Featured** go to **Monitor**.

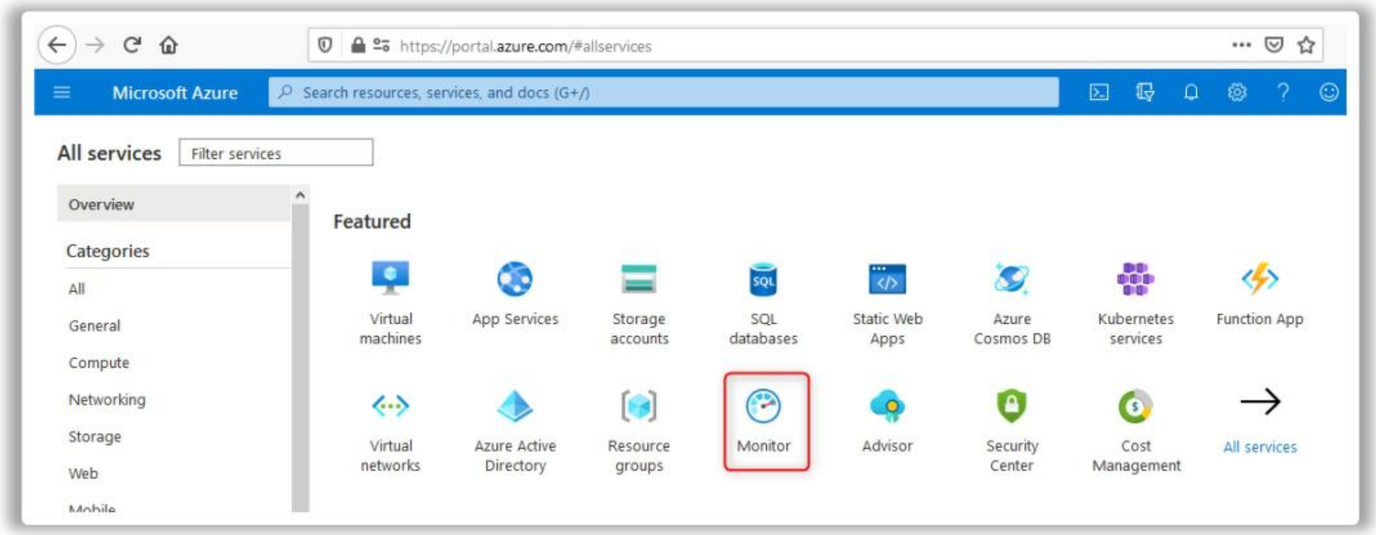


Figure 4.306: Select Monitor

- On the left navigation panel select **Alerts**.
- Go to **Manage alert rules** in the right-side panel.

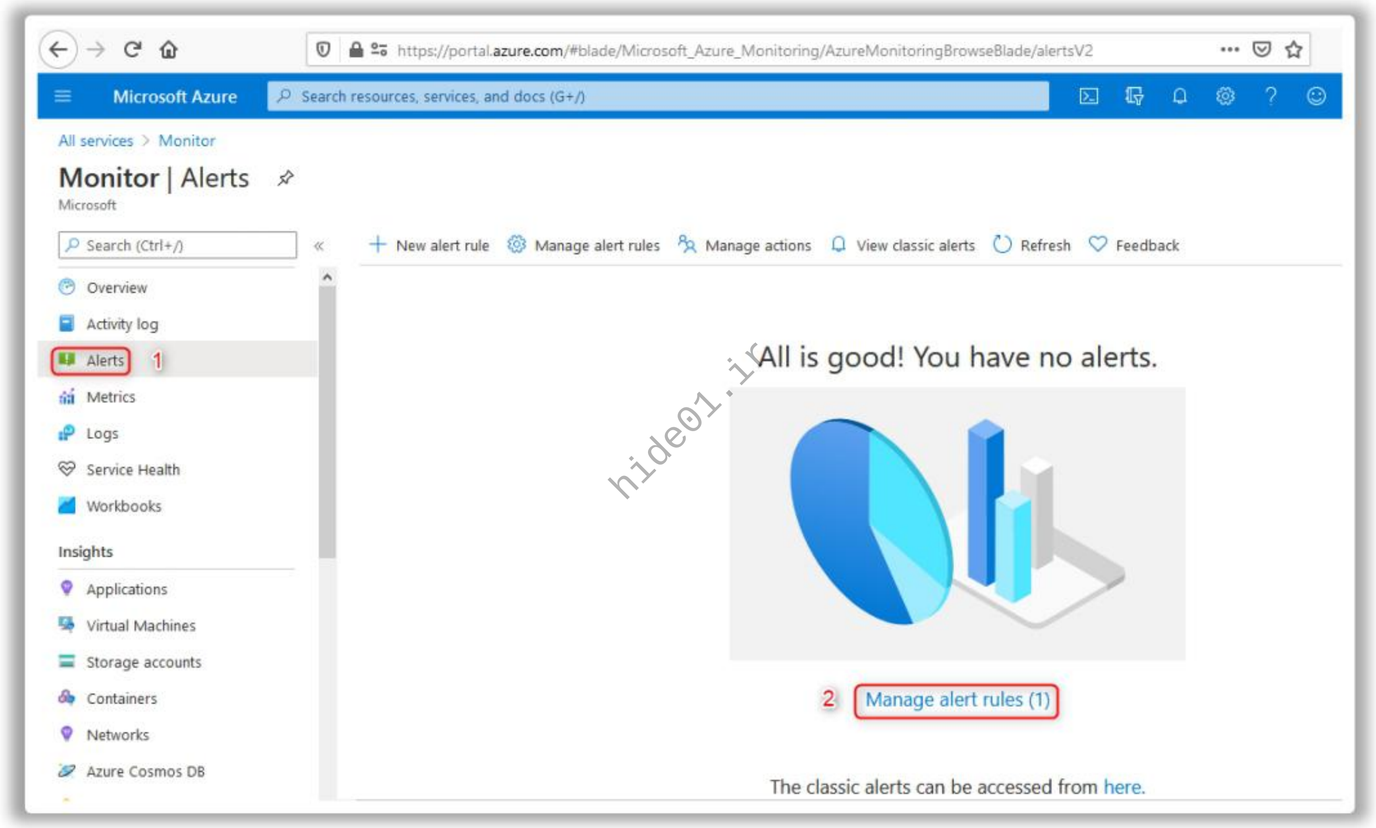


Figure 4.307: Click Manage Alert Rules

- Click on the **New alert rule**.

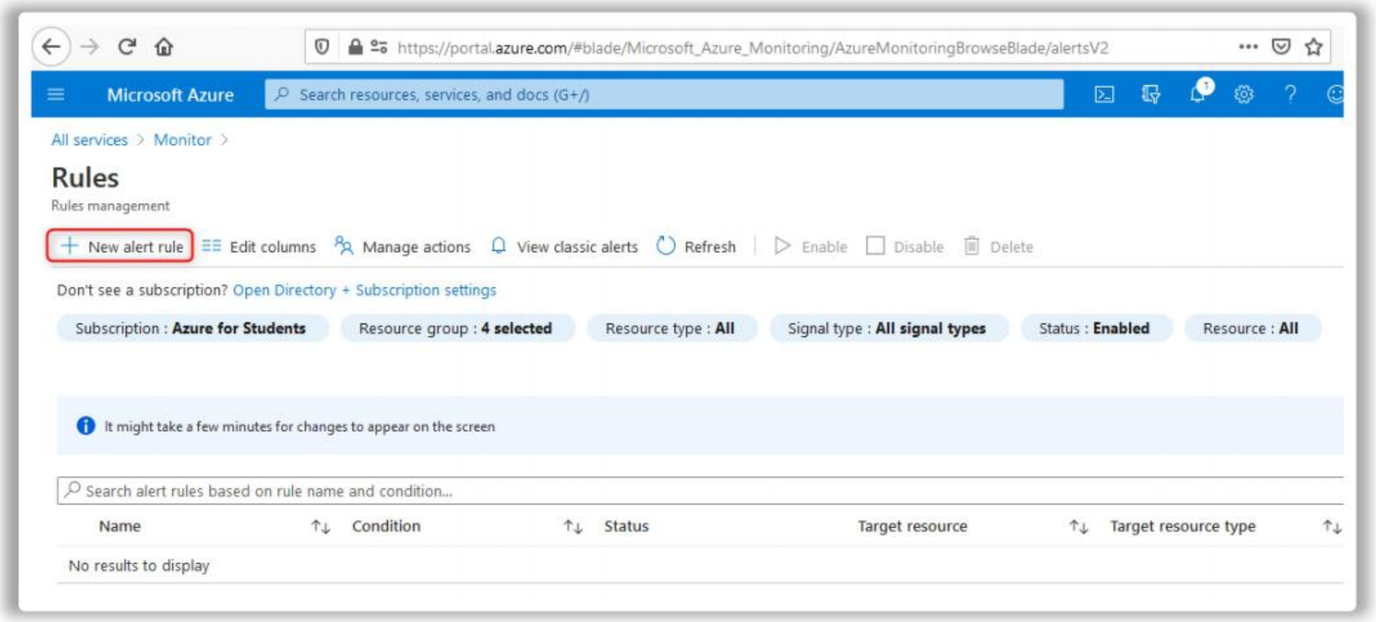


Figure 4.308: Select New Alert Rule

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

- Under scope click on **Select resource**.
- In **Select a resource** panel, select any resource that you want to create alerts.
- You will find the selected resource in the **Selection preview**.
- Click on **Done.**

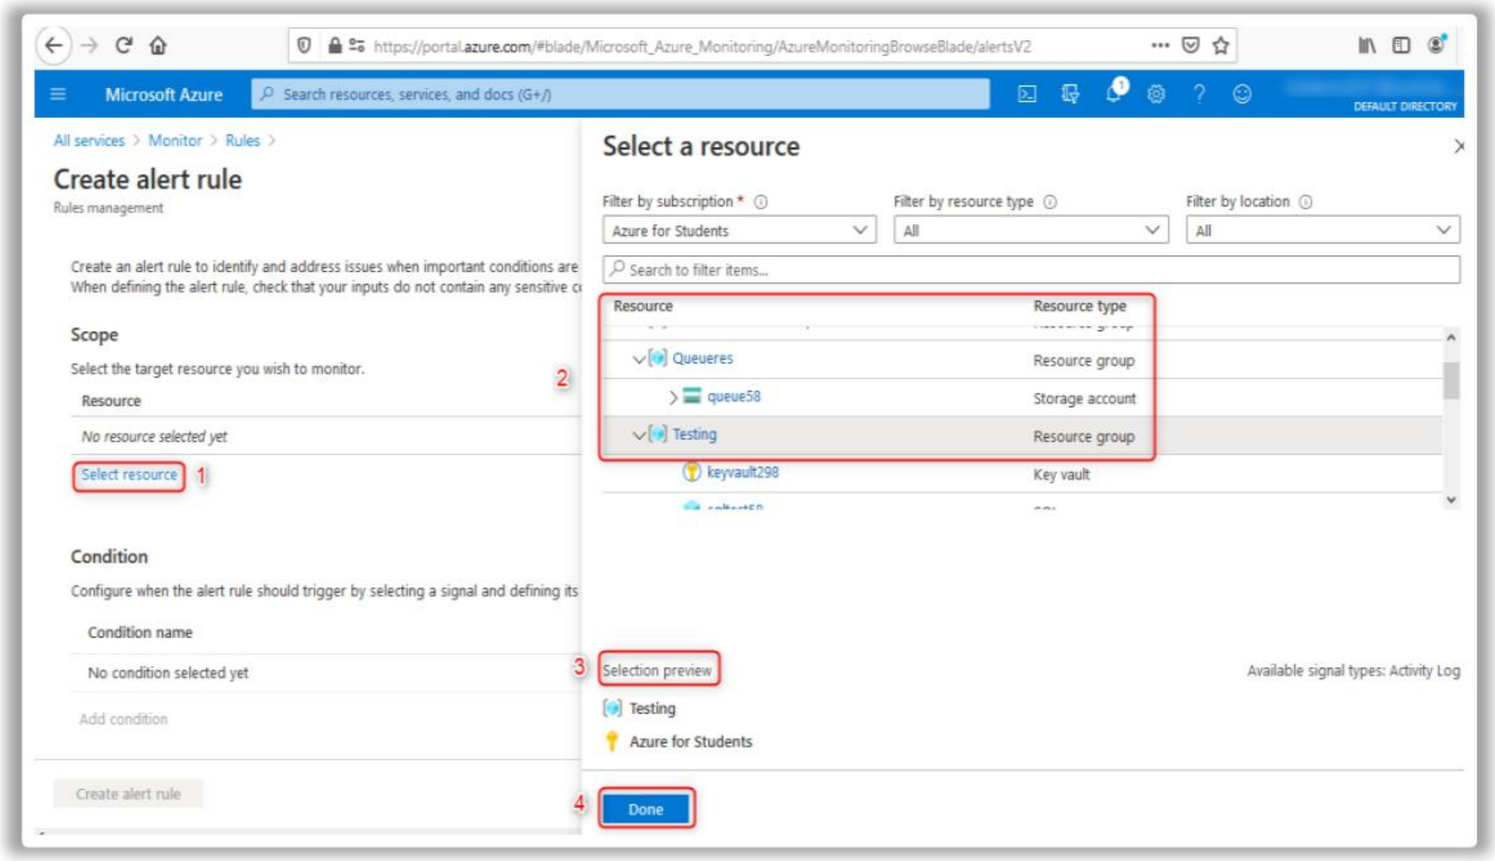


Figure 4.309: Select Save

- Under **Condition** click on **Add condition.**
- In the **Configure signal logic** panel search for **Delete Azure SQL Database (Microsoft.sql/servers/databases)** and select it.

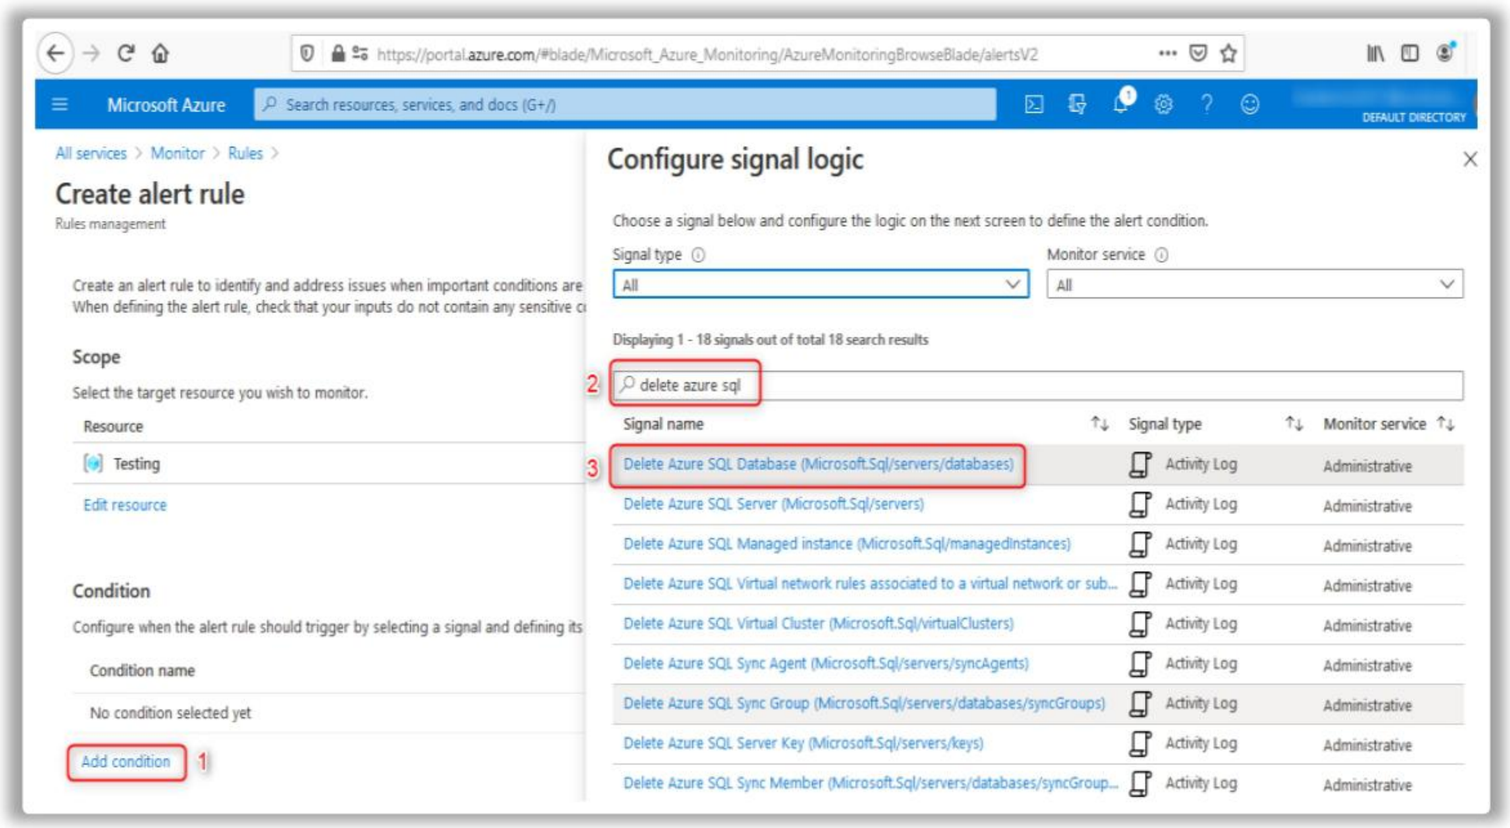


Figure 4.310: Configure Signal Logic Console

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

- Leave all fields default and click on **Done**.

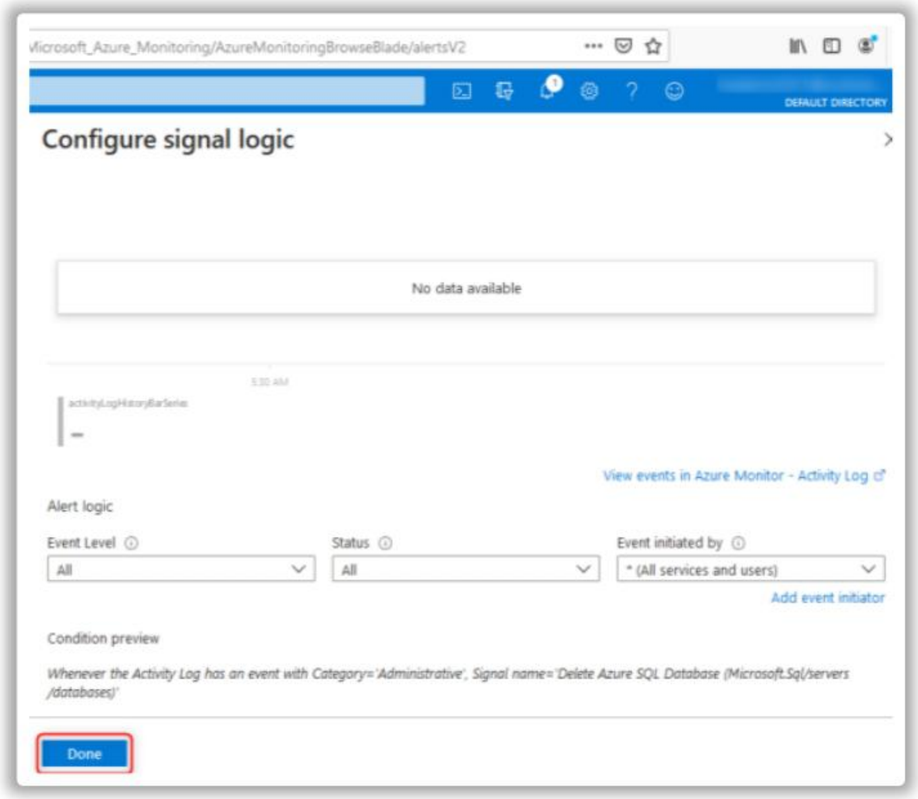


Figure 4.311: Click Done

- Under **Actions**, select **Add action groups** then click on **Create action group**.

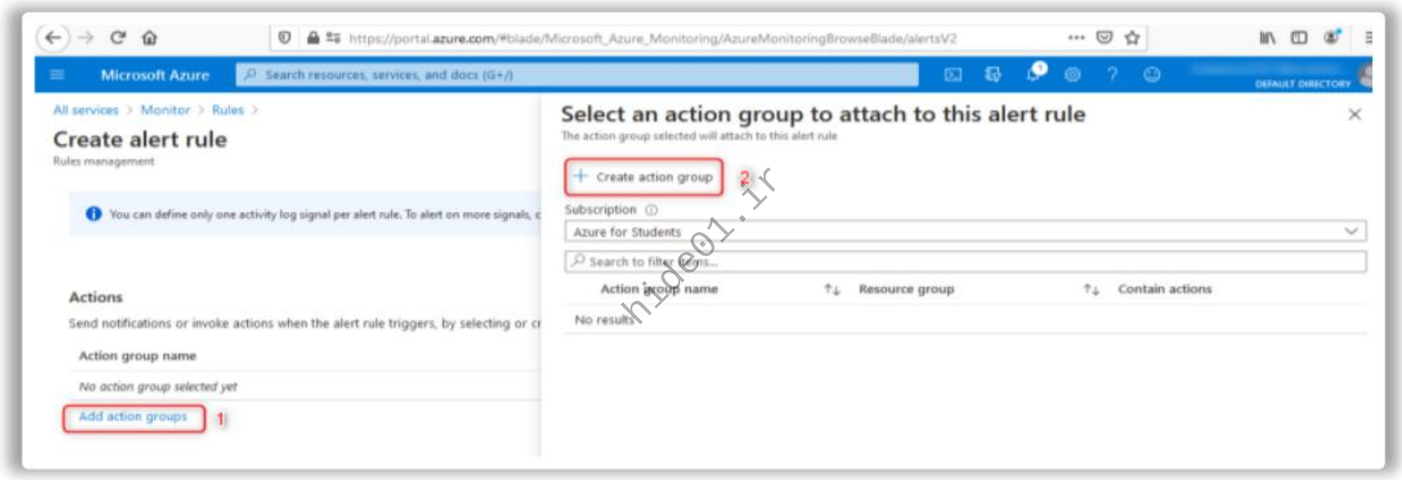


Figure 4.312: Select Create action group

- Enter required details.
- Click on **Review+create.**

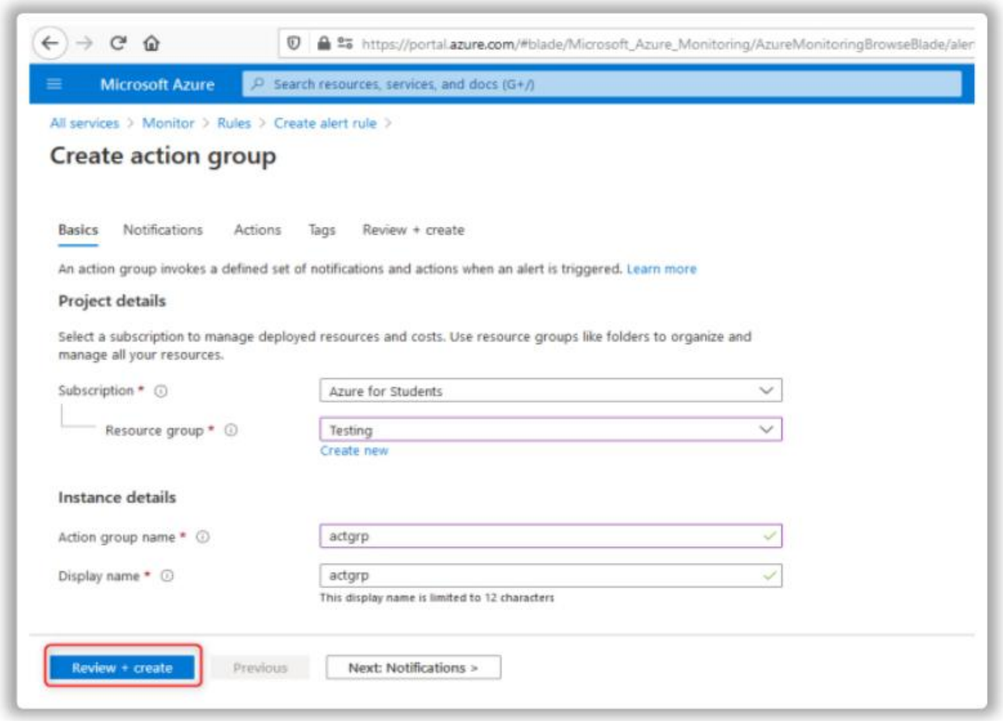


Figure 4.313: Click Review + create

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

- Under **Alert rule details**, fill in the details and Click on **Create alert rule.**

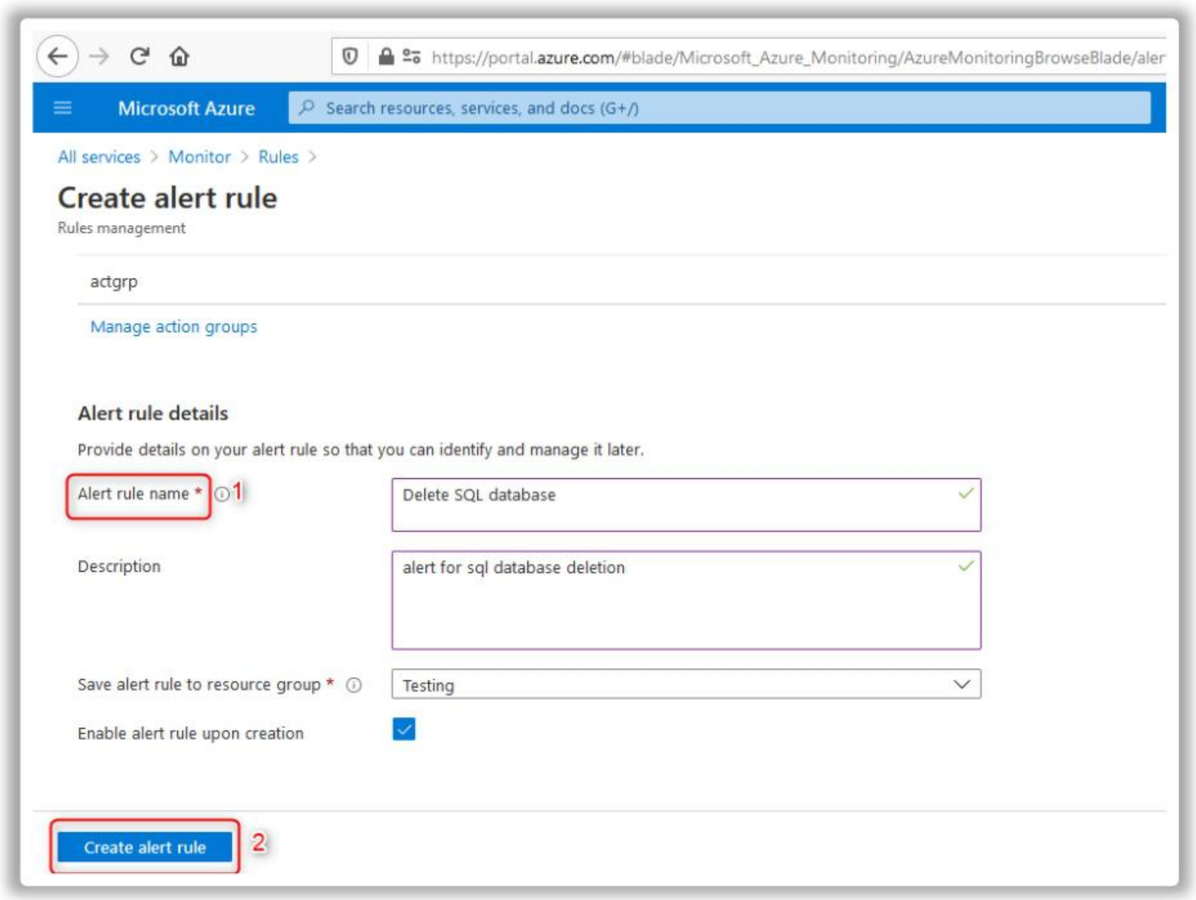


Figure 4.314: Create Alert Rule

- Alert rule is created successfully but it takes time to show in the display.

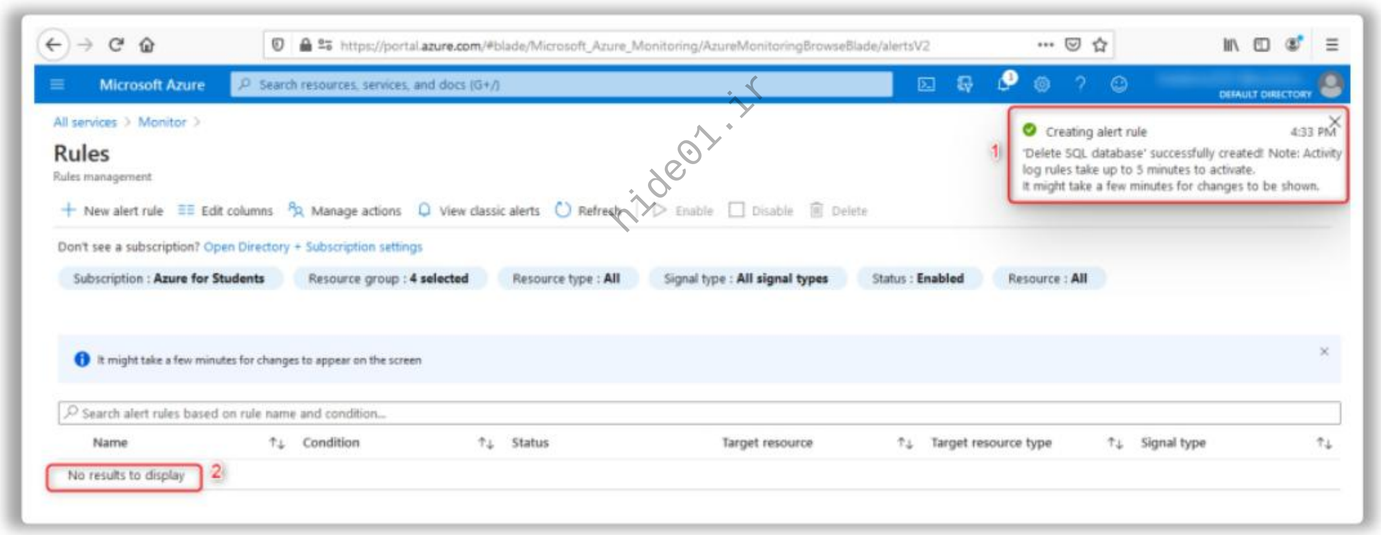


Figure 4.315: Alert Created

- Alert is successfully enabled.

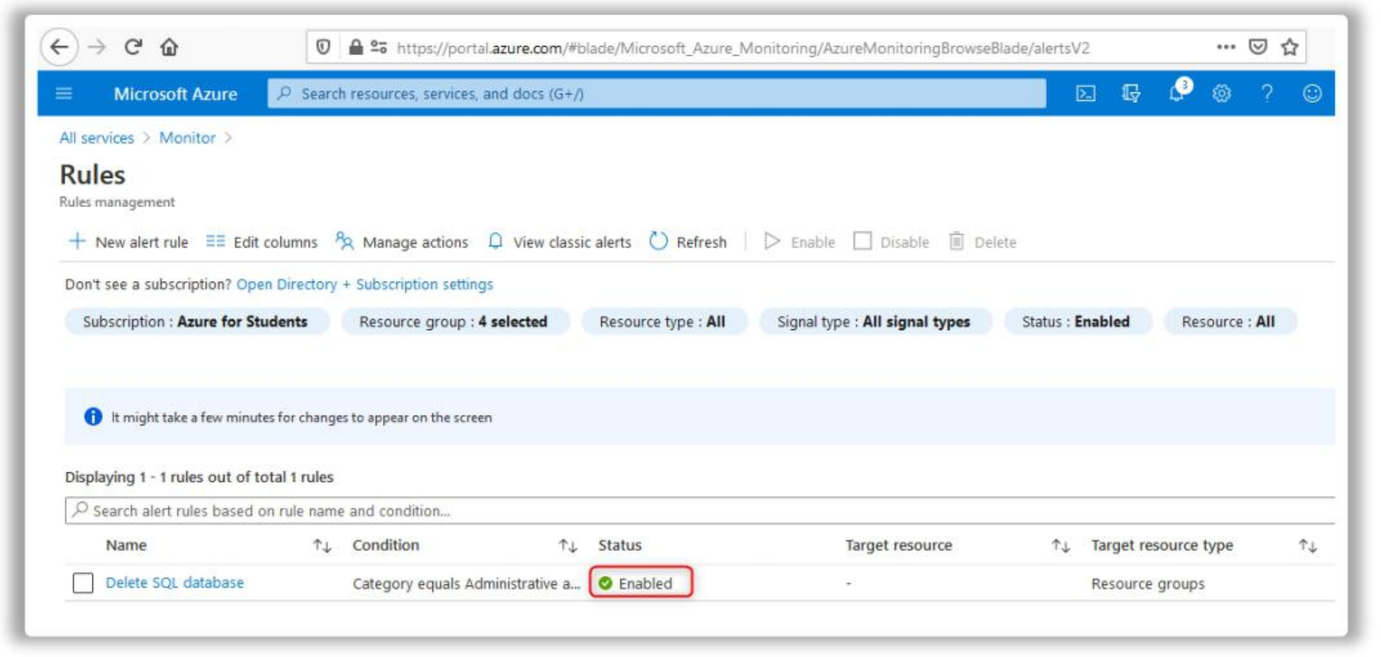
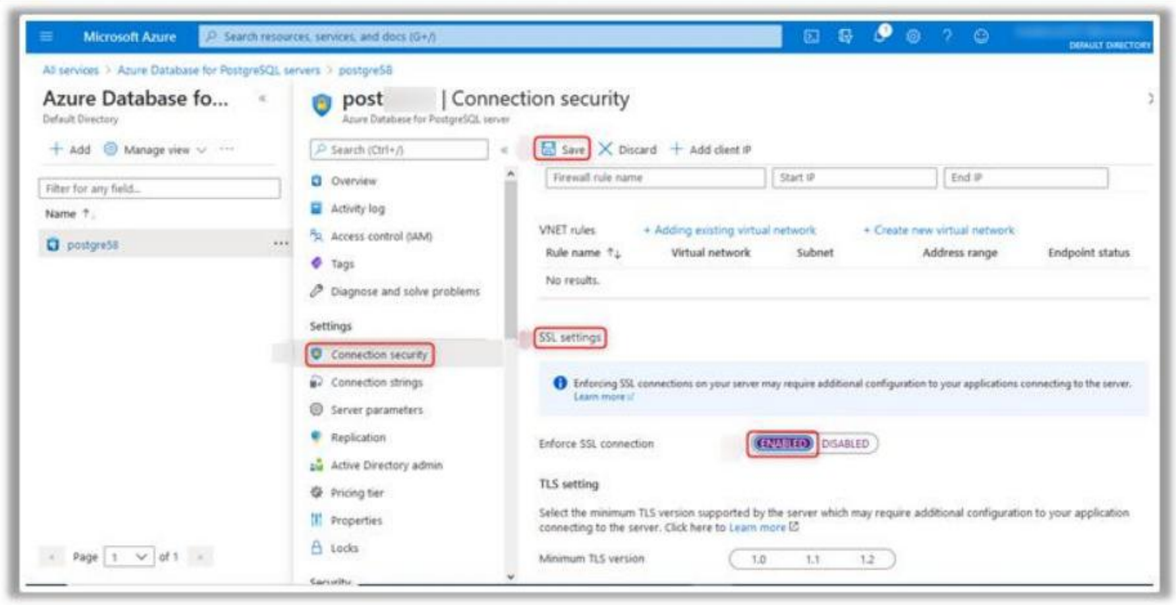


Figure 4.316: Alert Enabled

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

Enable In-transit Encryption for PostgreSQL Database Servers

- The cloud security professional should enable in-transit encryption using SSL protocol between the client application and Microsoft Azure PostgreSQL Database servers to protect the **data in motion** against man-in-the-middle attack and to fulfil compliance requirements



Enabling In-transit Encryption on PostgreSQL Database Server

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

Enable In-Transit Encryption for PostgreSQL Database Servers

The cloud security professional should enable in-transit encryption using SSL protocol between the client application and Microsoft Azure PostgreSQL Database servers to protect the data-in-motion against MIM attack and to fulfill compliance requirements. The in-transit encryption on Microsoft Azure PostgreSQL Database servers protects the confidential information from unauthorized users.

Source: https://azure.microsoft.com

To Enable In-Transit Encryption for PostgreSQL Database Servers

- Sign-in to Azure portal.
- Click on **More services**.

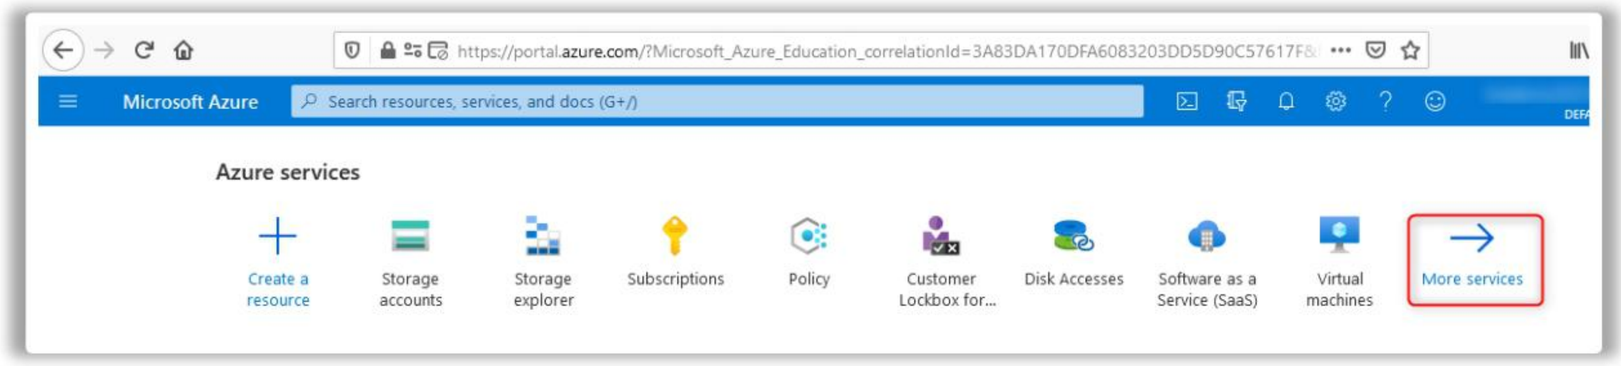


Figure 4.317: More Services

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

- Search for **Azure Database for PostgreSQL servers** using the input field beside **All services**, then click on the same in the right-side panel.

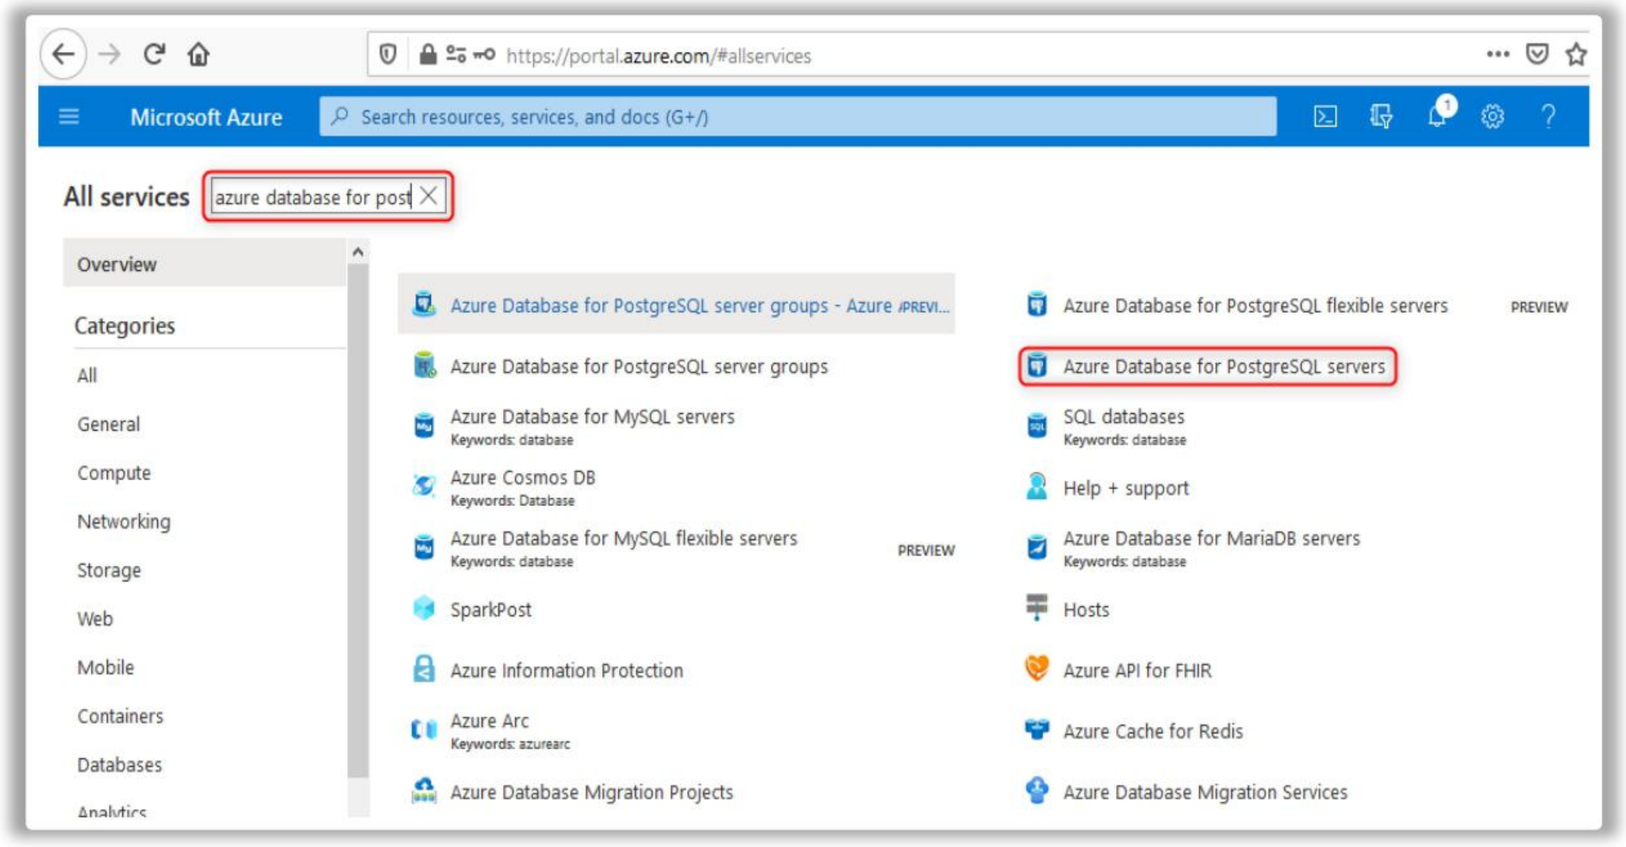


Figure 4.318: Search Azure Database

- A list of PostgreSQL servers will be displayed.
- Select the PostgreSQL server that you want to enable in-transit encryption.

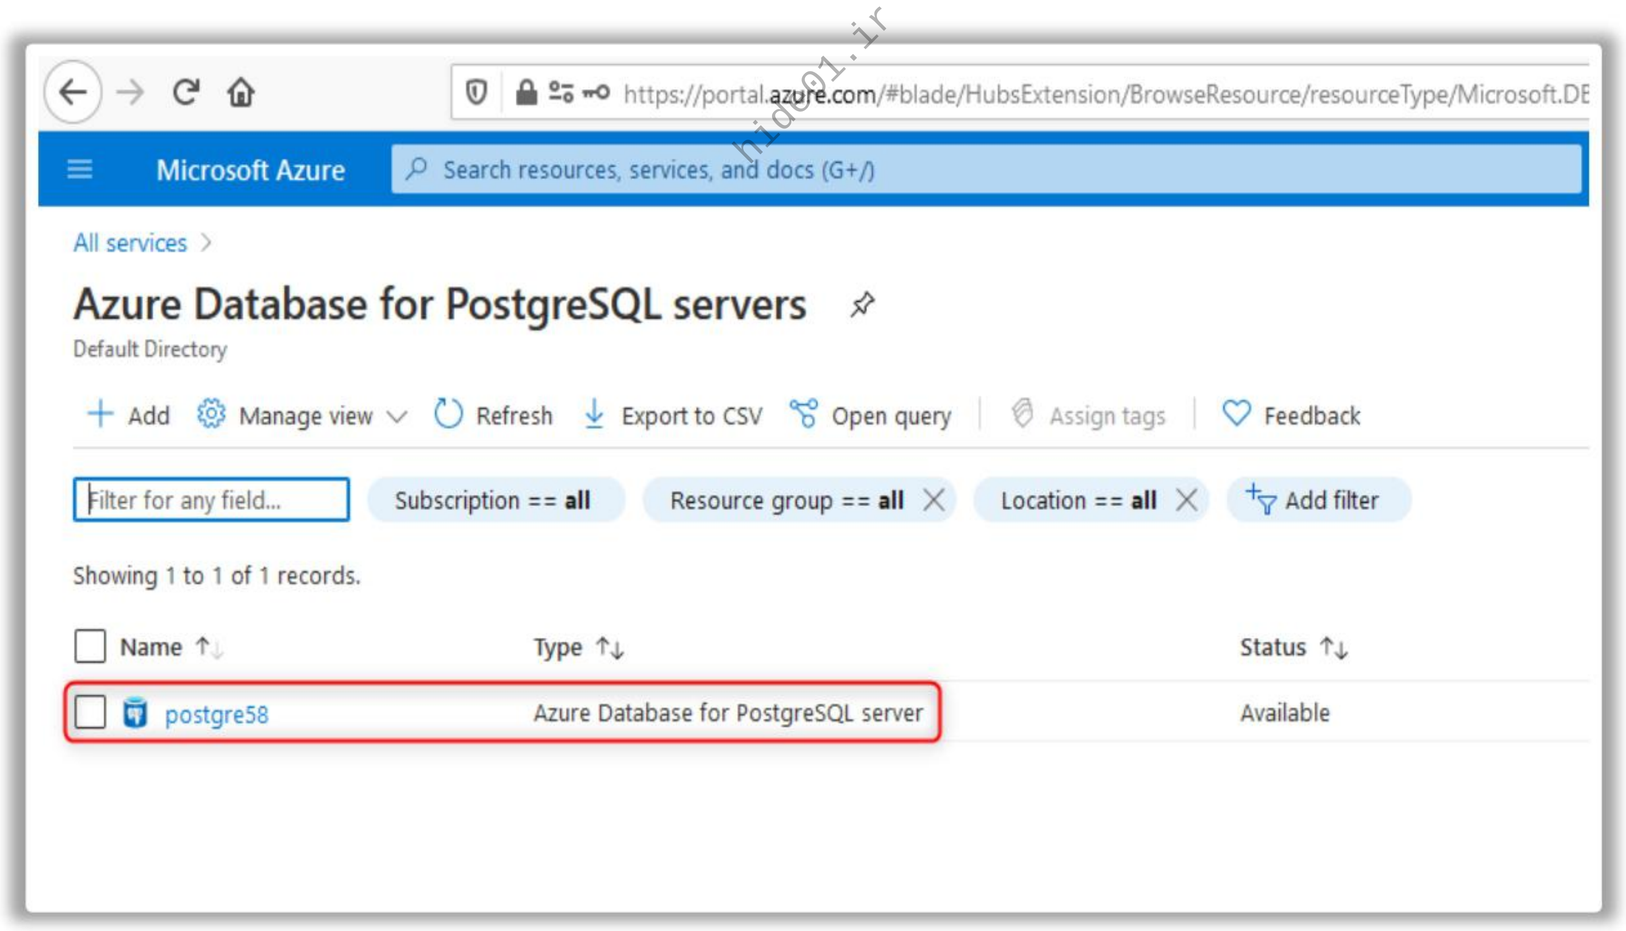


Figure 4.319: Choose PostgreSQL

- Locate **Settings** and go to **Connection security**.
- Scroll down to **SSL settings** on the right-side panel.
- Next to **Enforce SSL connection,** click on **Enabled**.
- Click on **Save**.

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

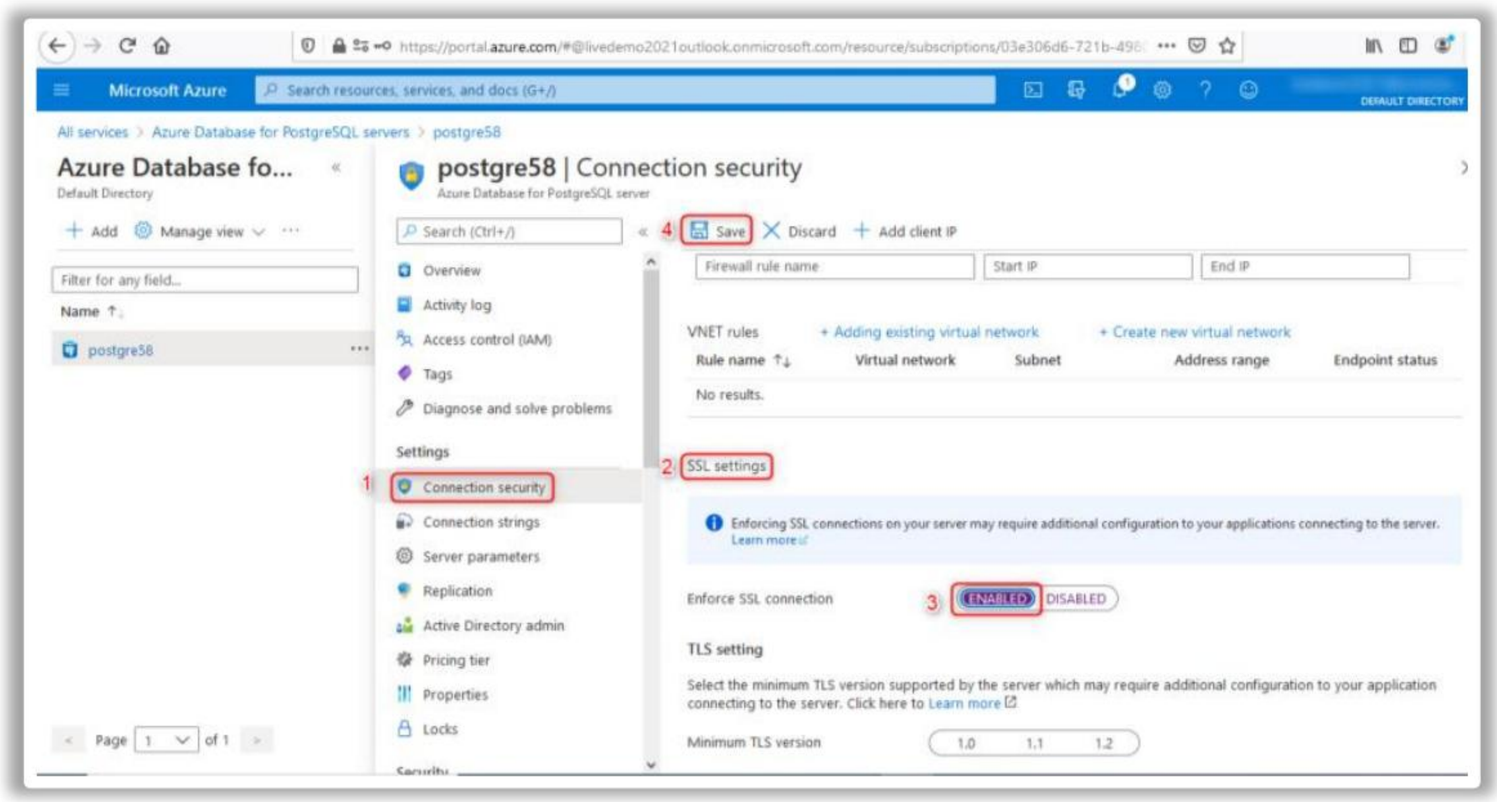


Figure 4.320: Click Save

- In-transit encryption for PostgreSQL servers is successfully enabled.

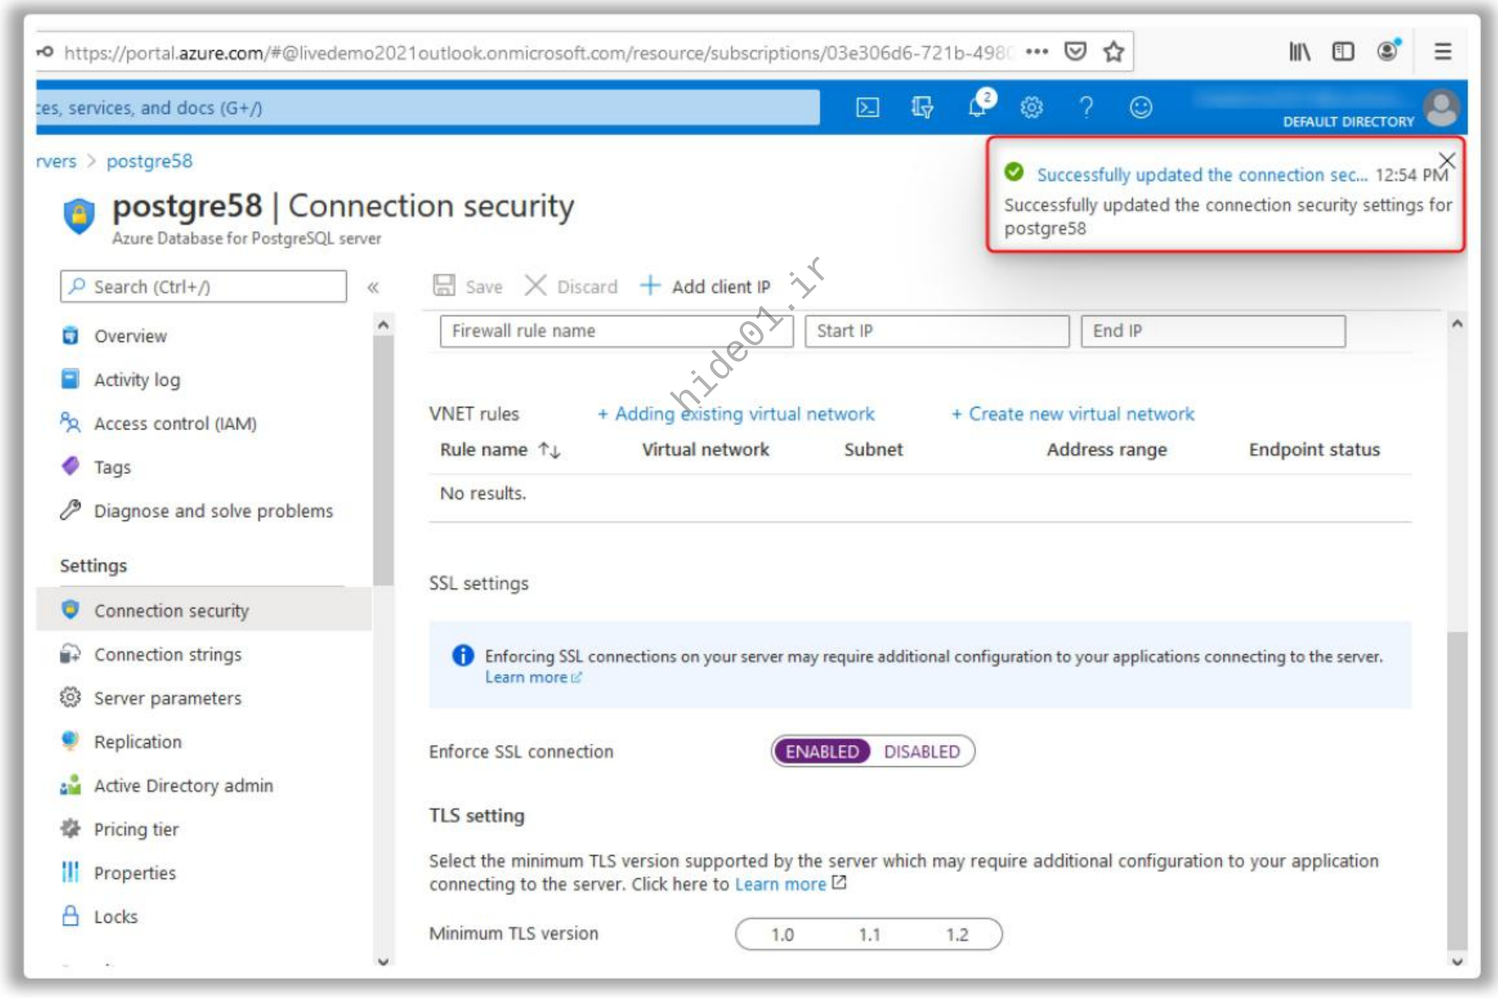
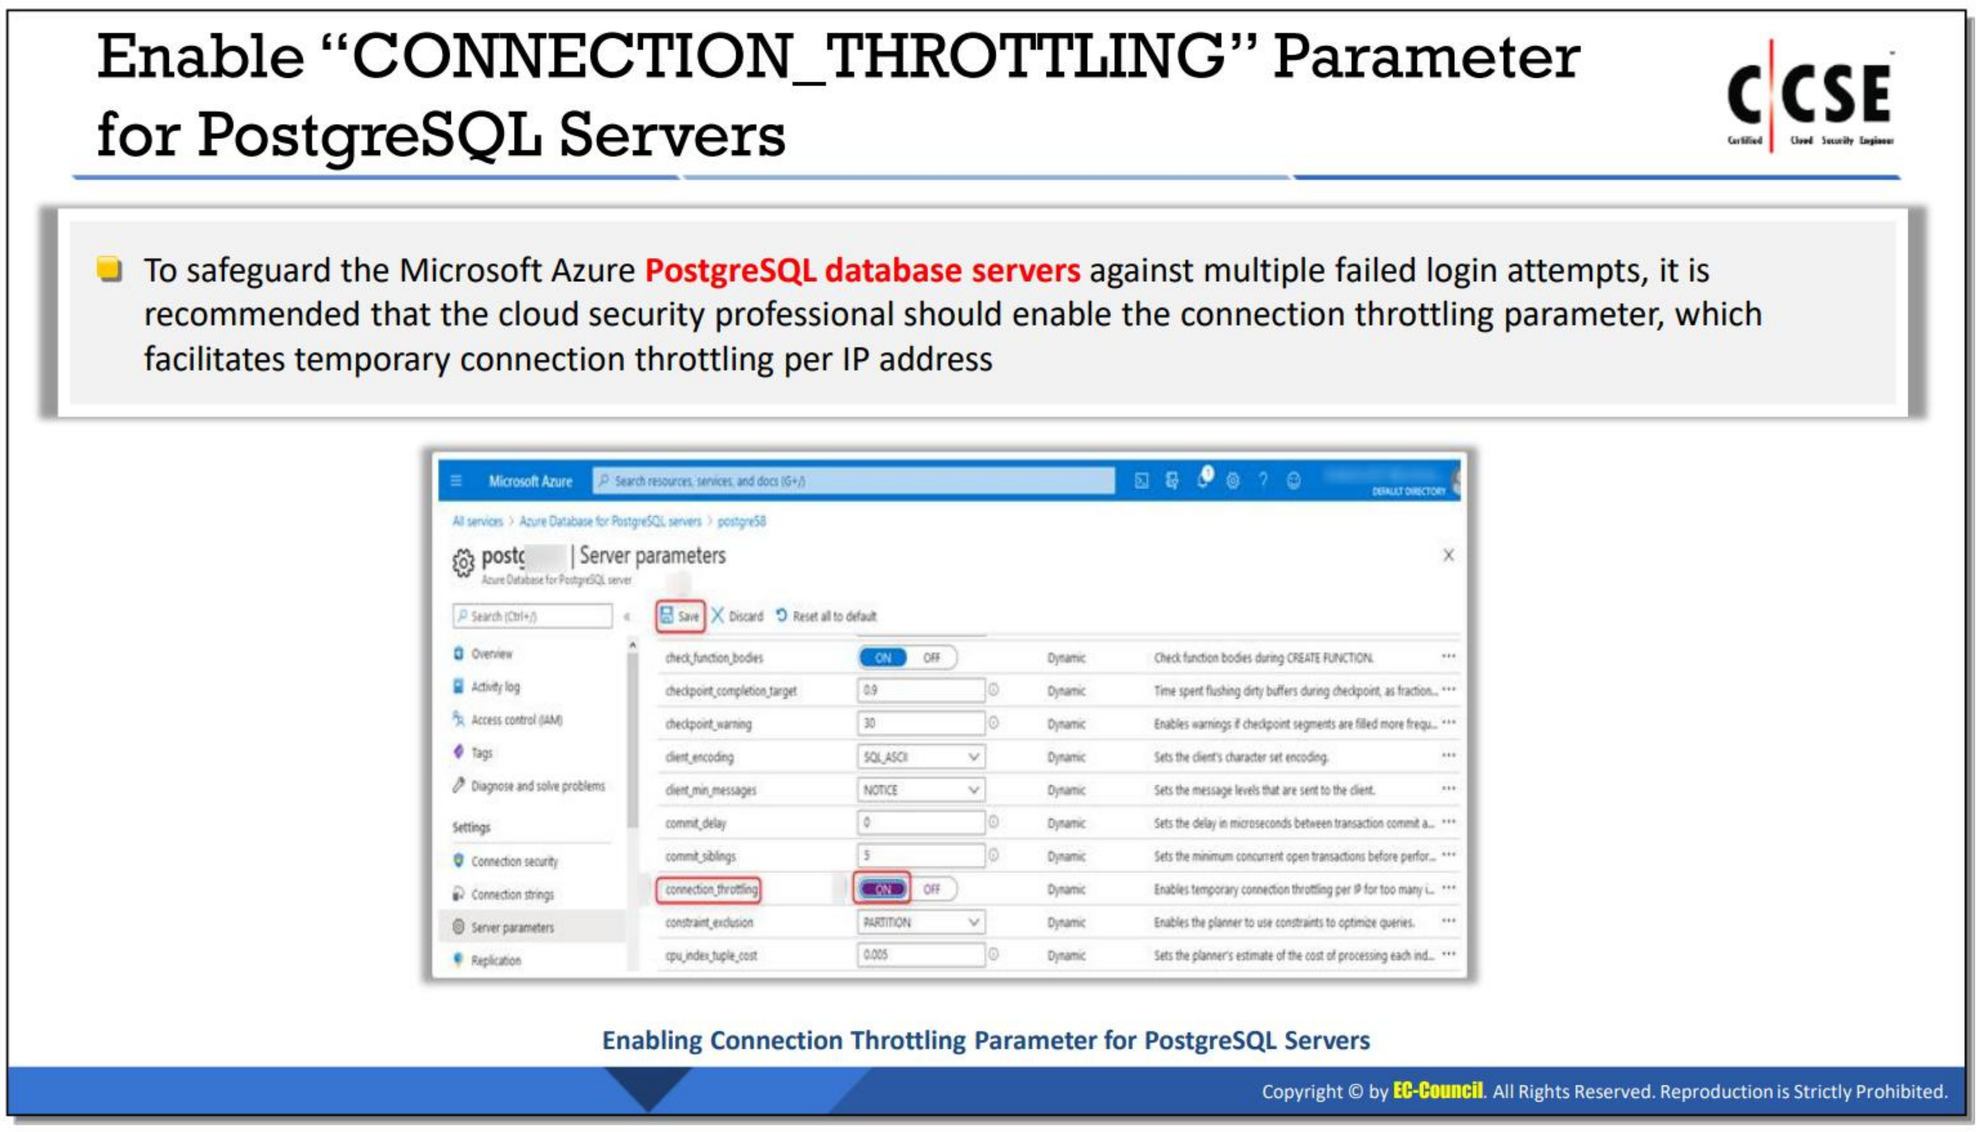


Figure 4.321: Enable Encryption

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



Enable "CONNECTION_THROTTLING" Parameter for PostgreSQL Servers

The connection throttling parameter is useful in the log generation of concurrent connections. This log data helps to identify DDoS attacks in PostgreSQL database servers, which overloads and make resources unavailable for legitimate users.

To safeguard the Microsoft Azure PostgreSQL Database servers against multiple failed login attempts, it is recommended that the cloud security professional should enable the **`connection_throttling`** parameter, which facilitates temporary connection throttling per IP address.

Source: https://azure.microsoft.com

To Enable "CONNECTION_THROTTLING" Parameter for PostgreSQL Servers

- Sign-in to Azure portal.
- Click on **More services**.

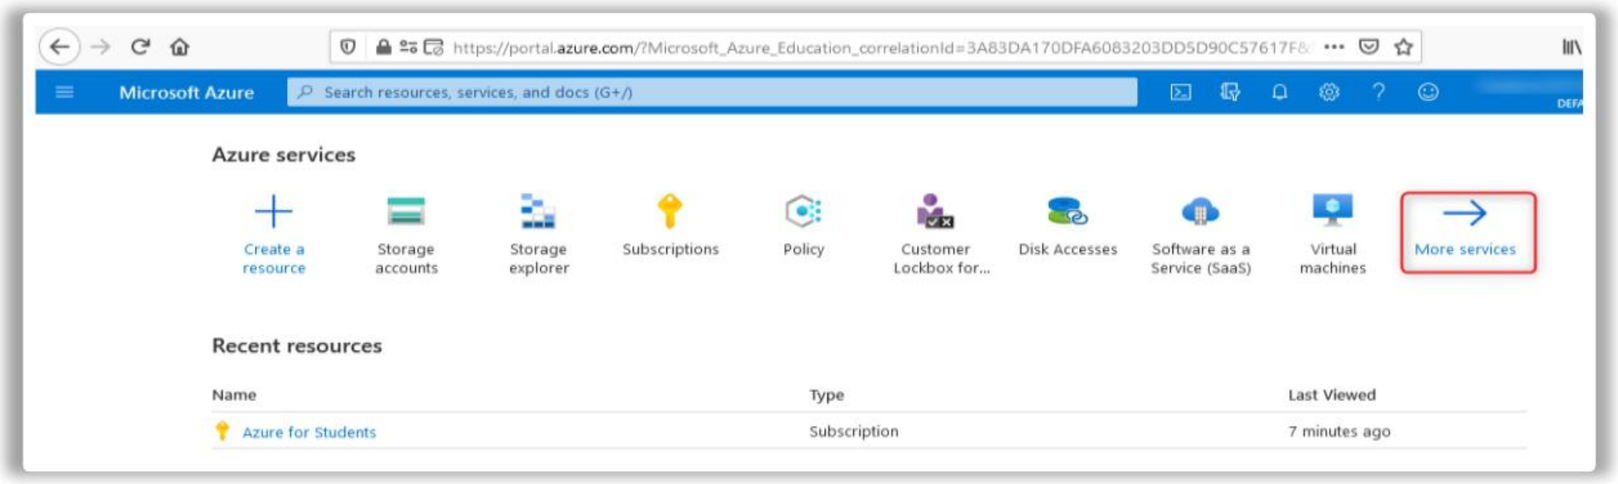


Figure 4.322: More Service

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

- Search for **Azure Database for PostgreSQL servers** using the input field beside **All services**, then click on the same in the right-side panel.

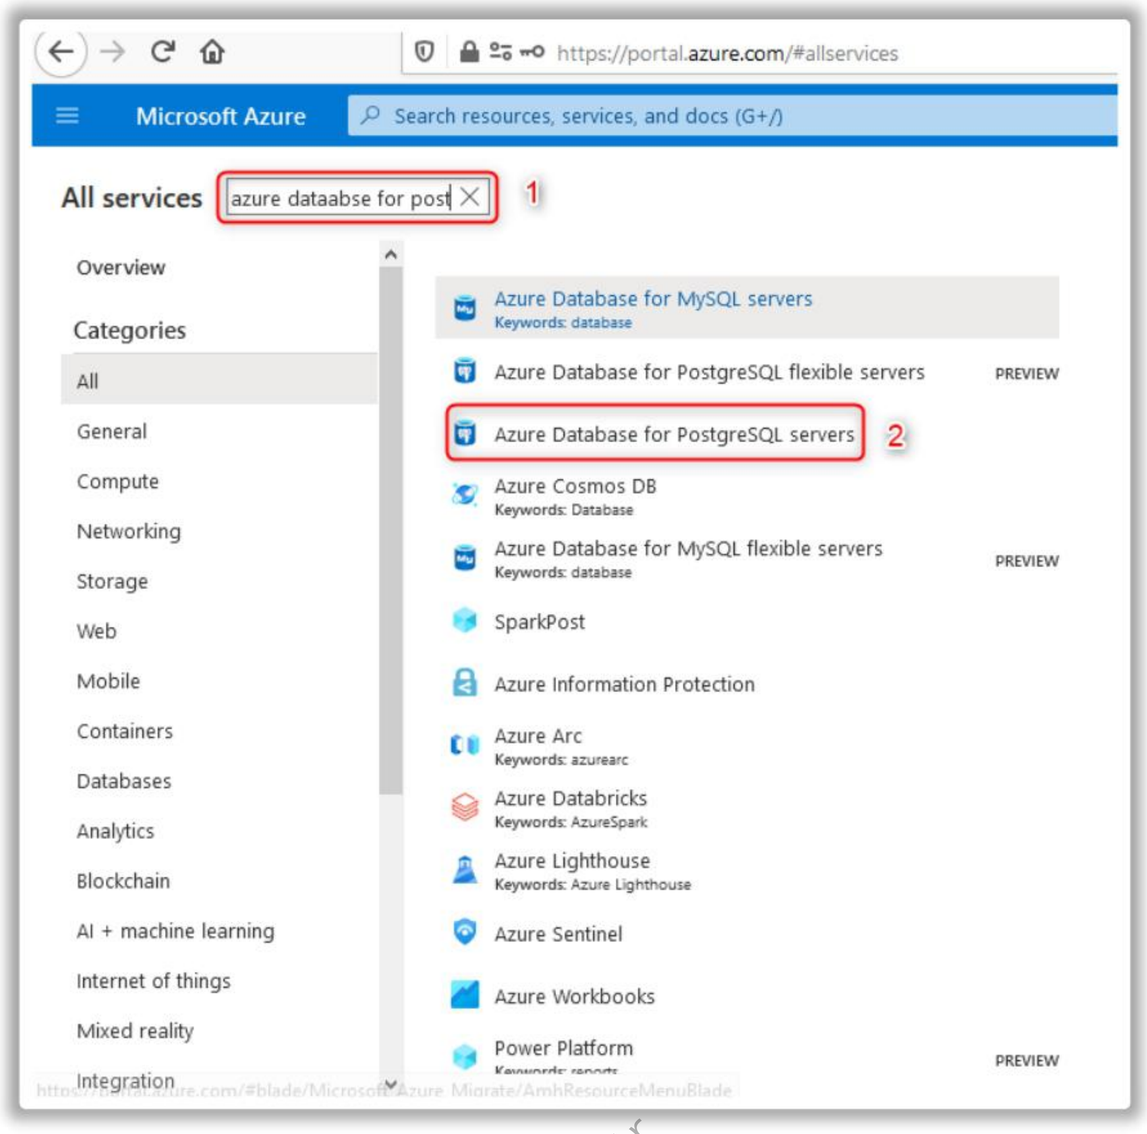


Figure 4.323: Search Database

- A list of PostgreSQL servers will be displayed.
- Select the PostgreSQL server that you want to enable Connection Throttling.

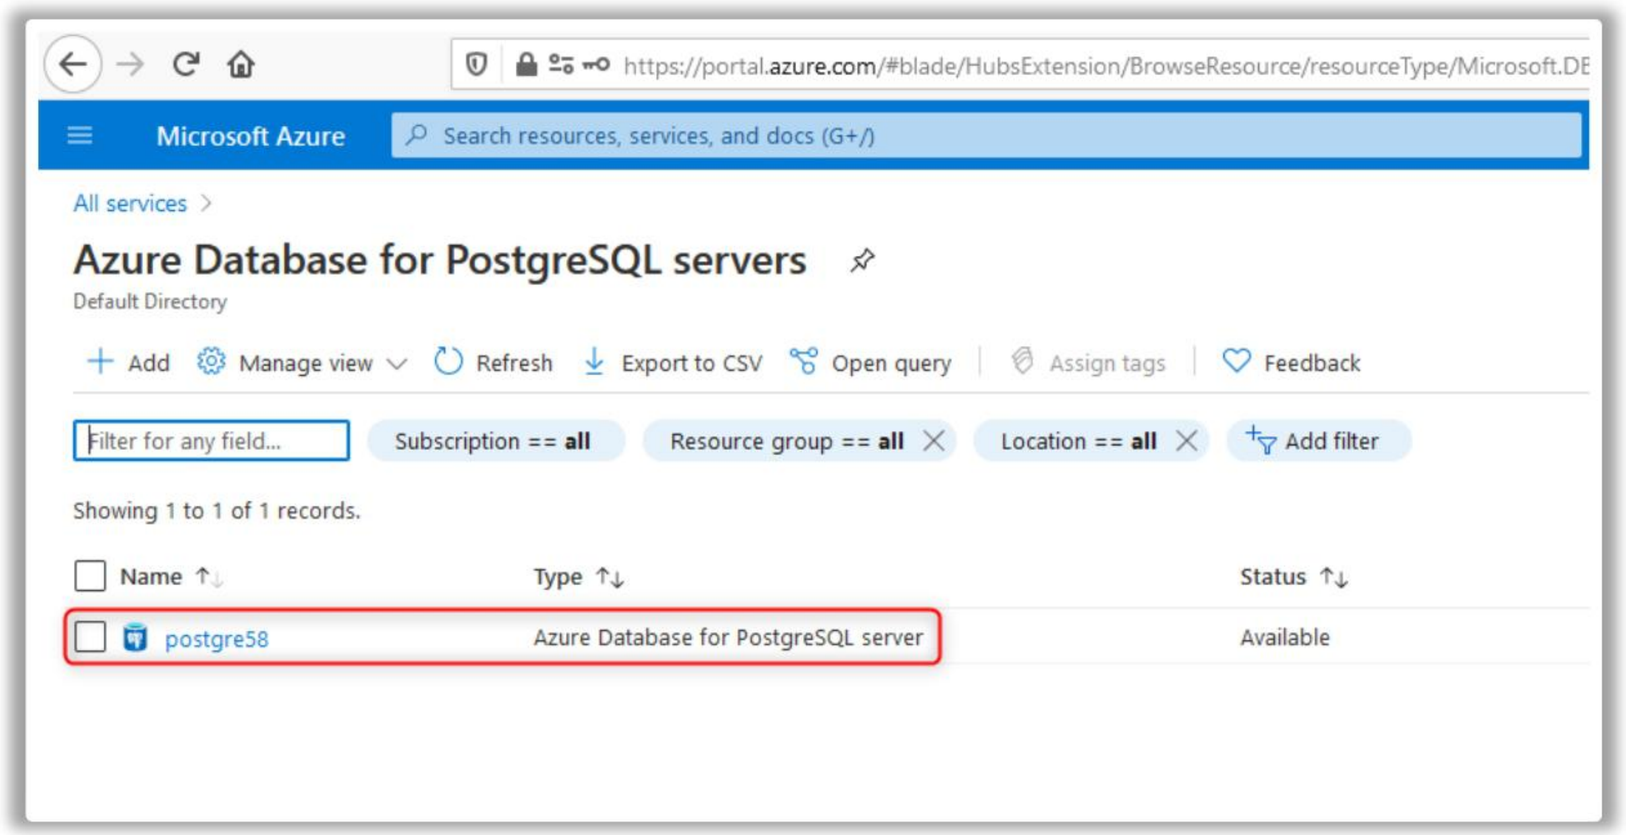


Figure 4.324: Select PostgreSQL

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

- Locate **Settings** in the navigation panel and go to **Server Parameters.**

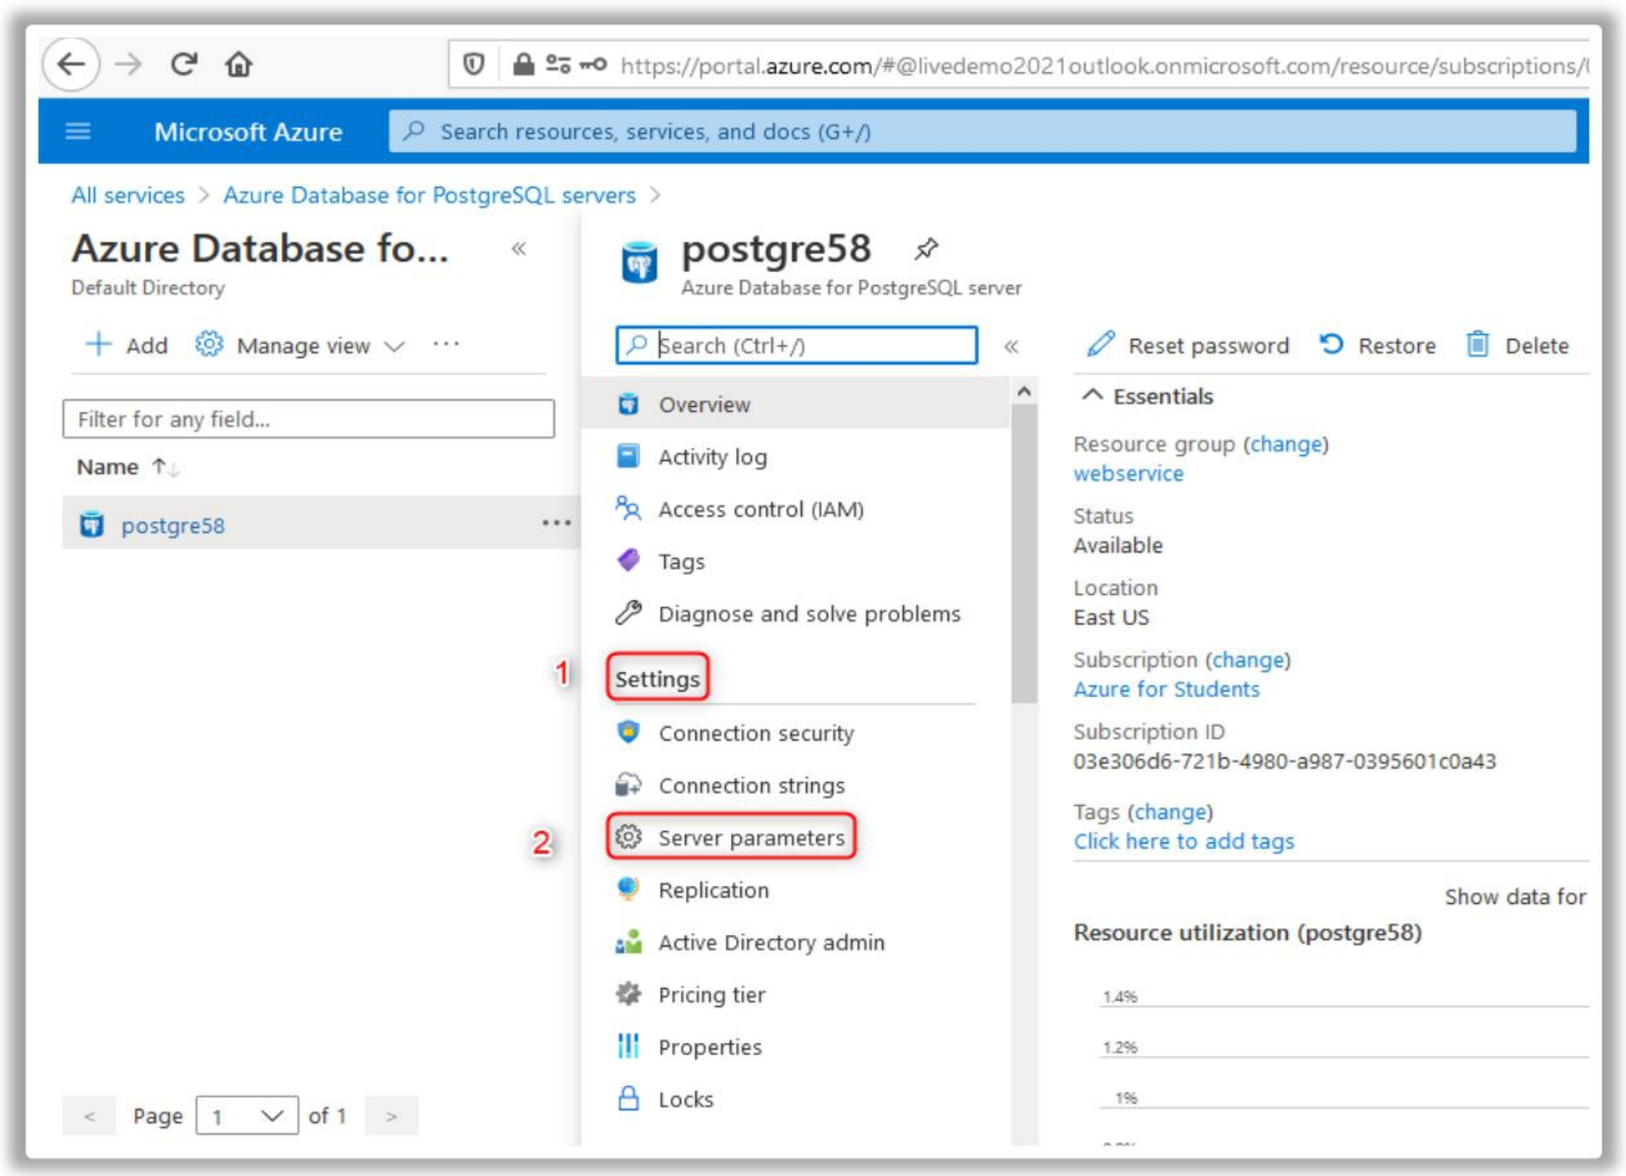


Figure 4.325: Select Server Parameters

- Scroll down to **connection_throttling** and click on **ON**.
- Click on **Save**.

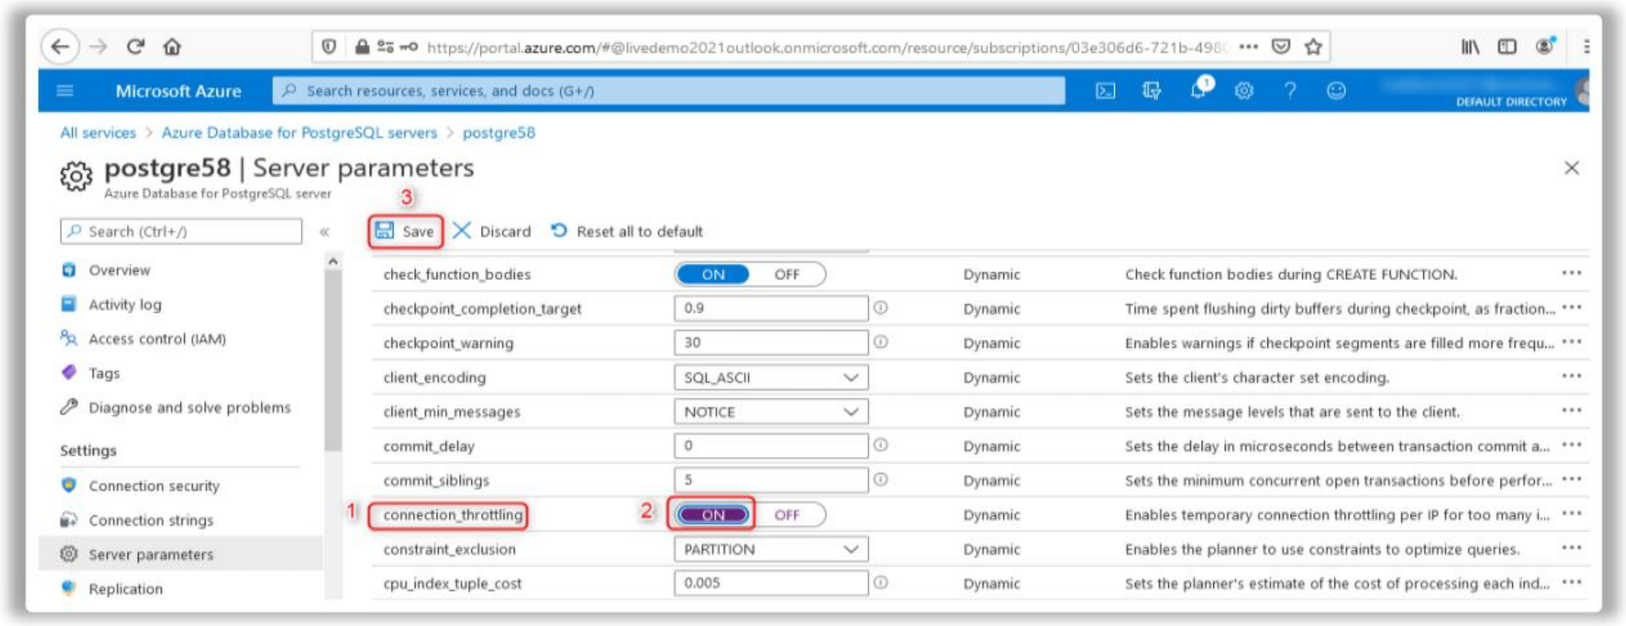


Figure 4.326: Click Save

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

- Connection throttling has been successfully enabled.

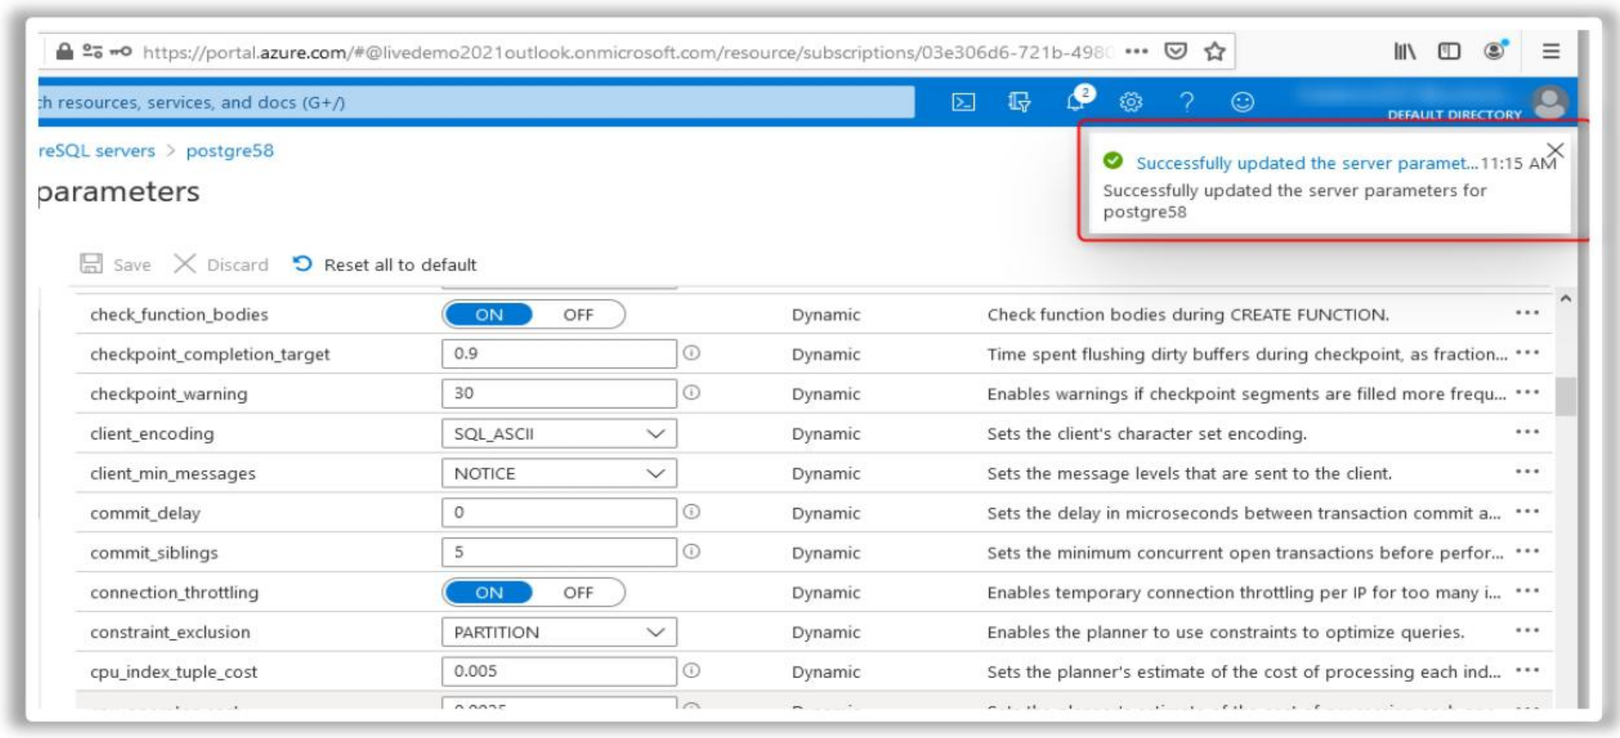
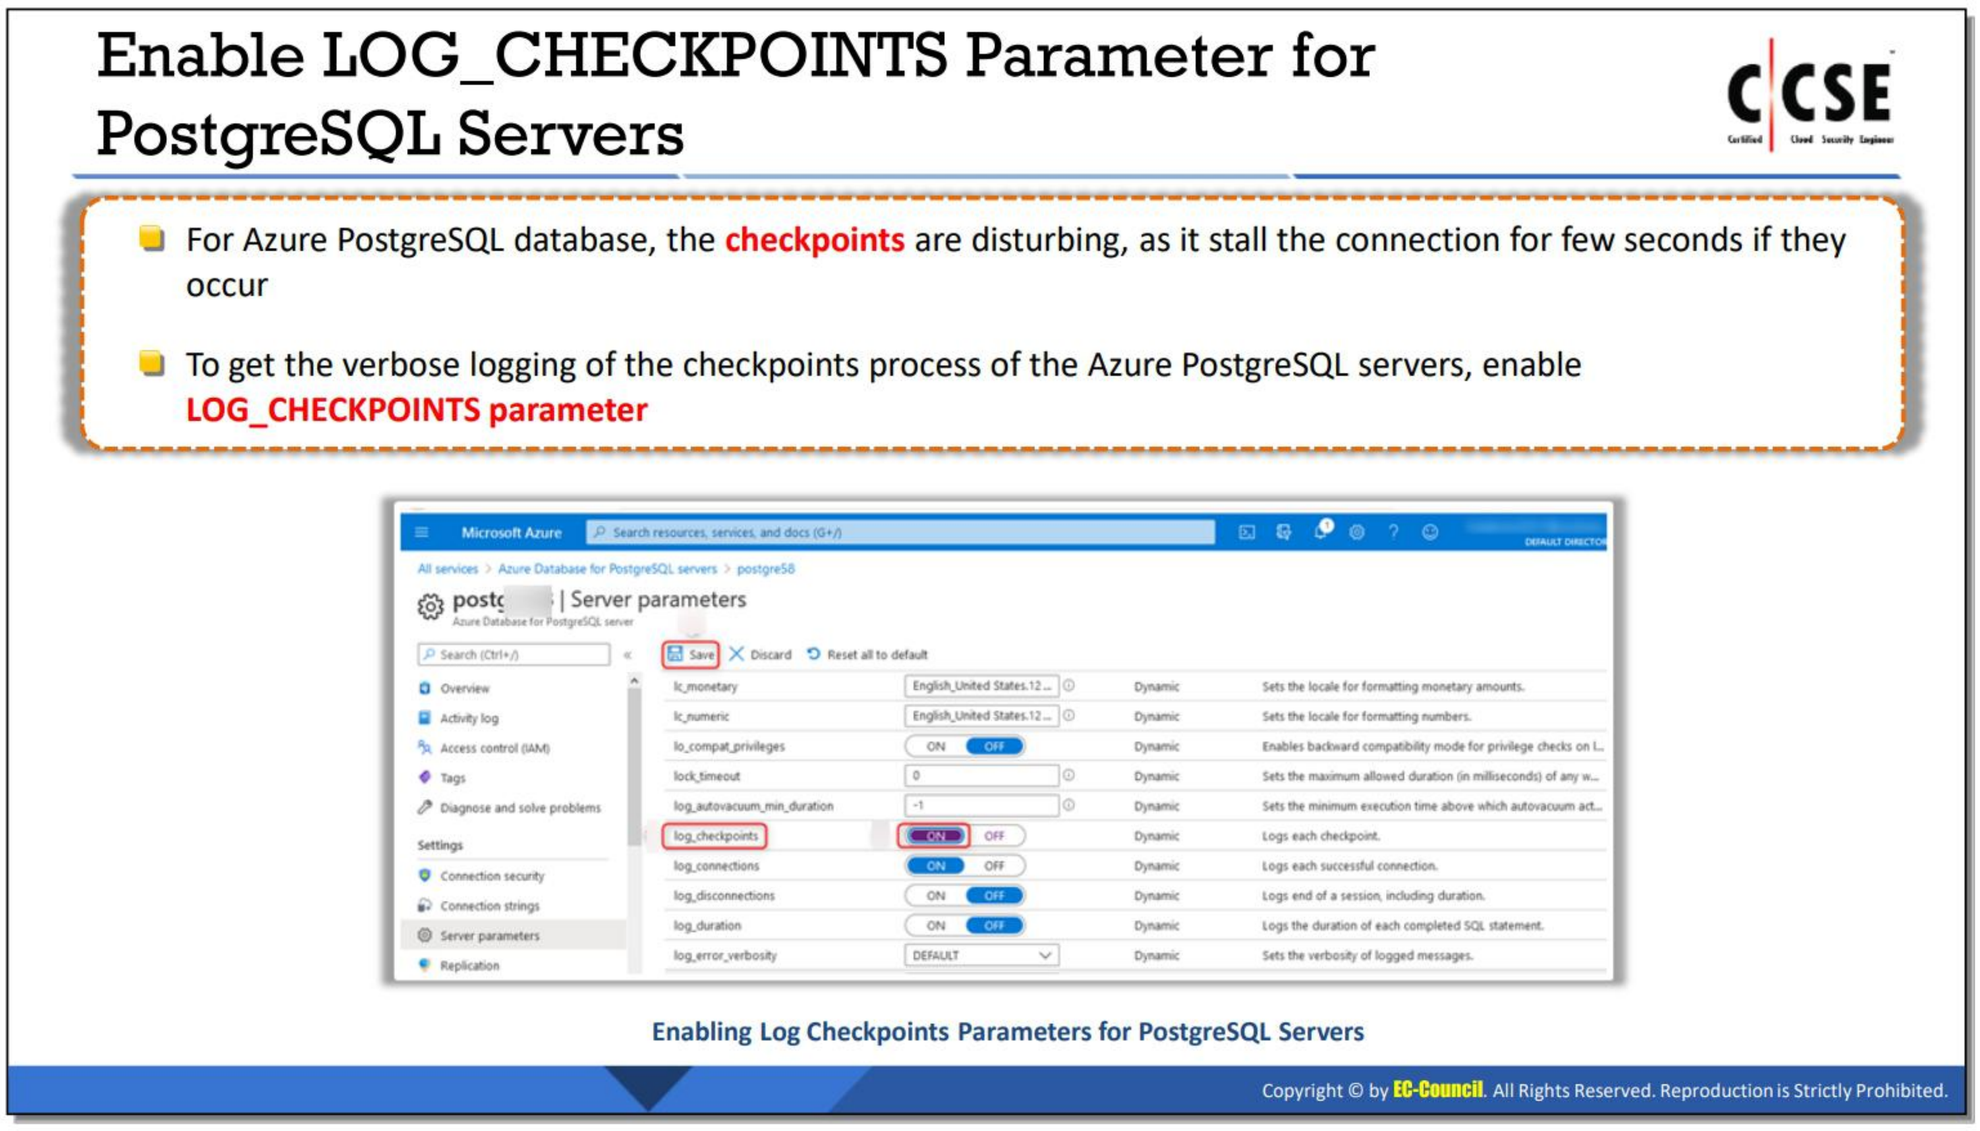


Figure 4.327: Connection Throttling is Enabled

hide01.ir

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



Enable LOG_CHECKPOINTS Parameter for Azure PostgreSQL Servers

For Azure PostgreSQL database, the checkpoints are disturbing, as it stalls the connection for few seconds if they occur. To get the verbose logging of the checkpoints process of the Azure PostgreSQL servers, enable LOG_CHECKPOINTS parameter. The data obtained from this logging will be utilized to troubleshoot and identify the sub-optimal PostgreSQL database performance.

Source: https://azure.microsoft.com

To Enable Log Checkpoints Parameters for PostgreSQL Servers

- Sign-in to Azure portal.
- Click on **More services**.

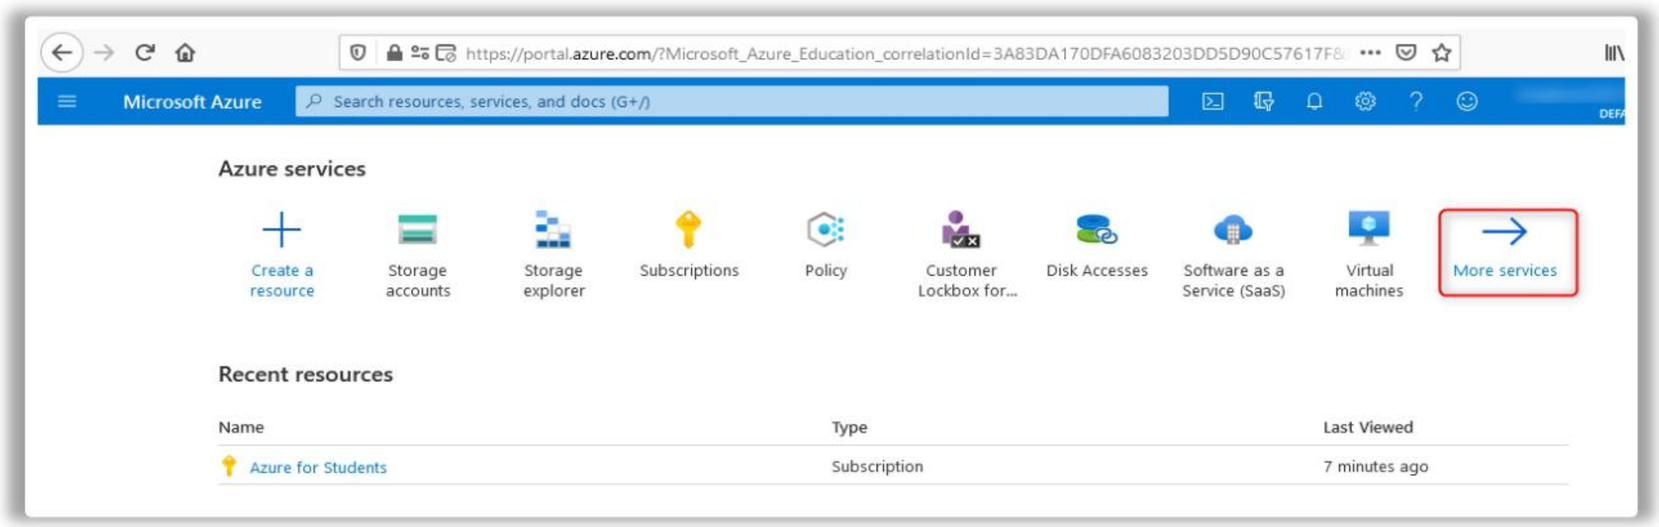


Figure 4.328: Select More Service

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

- Search for **Azure Database for PostgreSQL servers** using the input field beside **All services**, then click on the same in the right-side panel.

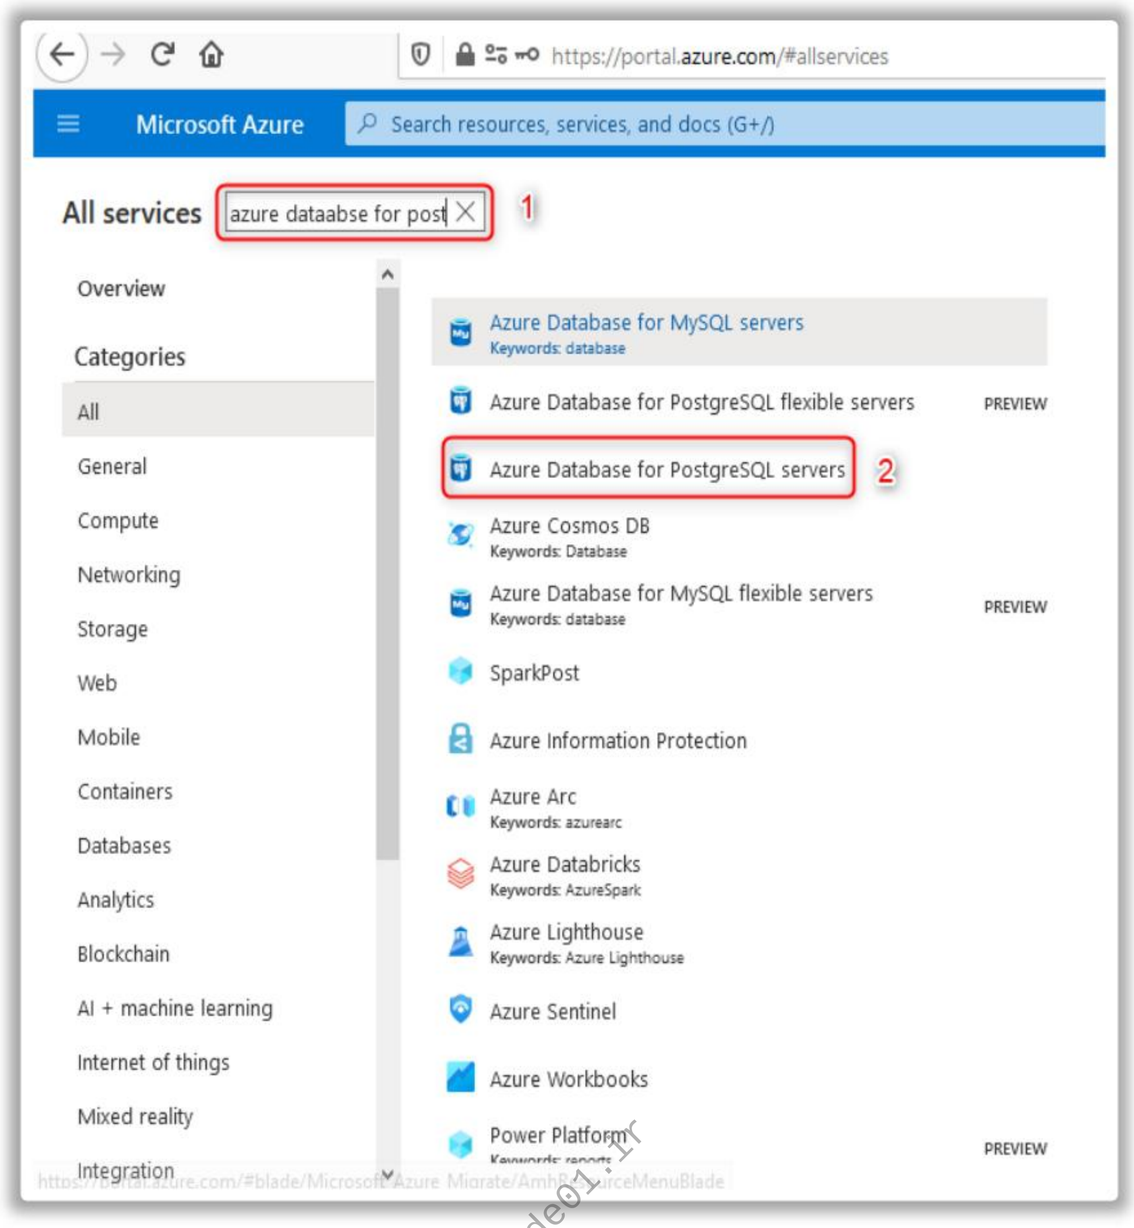


Figure 4.329: Search Database

- A list of PostgreSQL servers will be displayed.
- Select the PostgreSQL server that you want to enable log checkpoints.

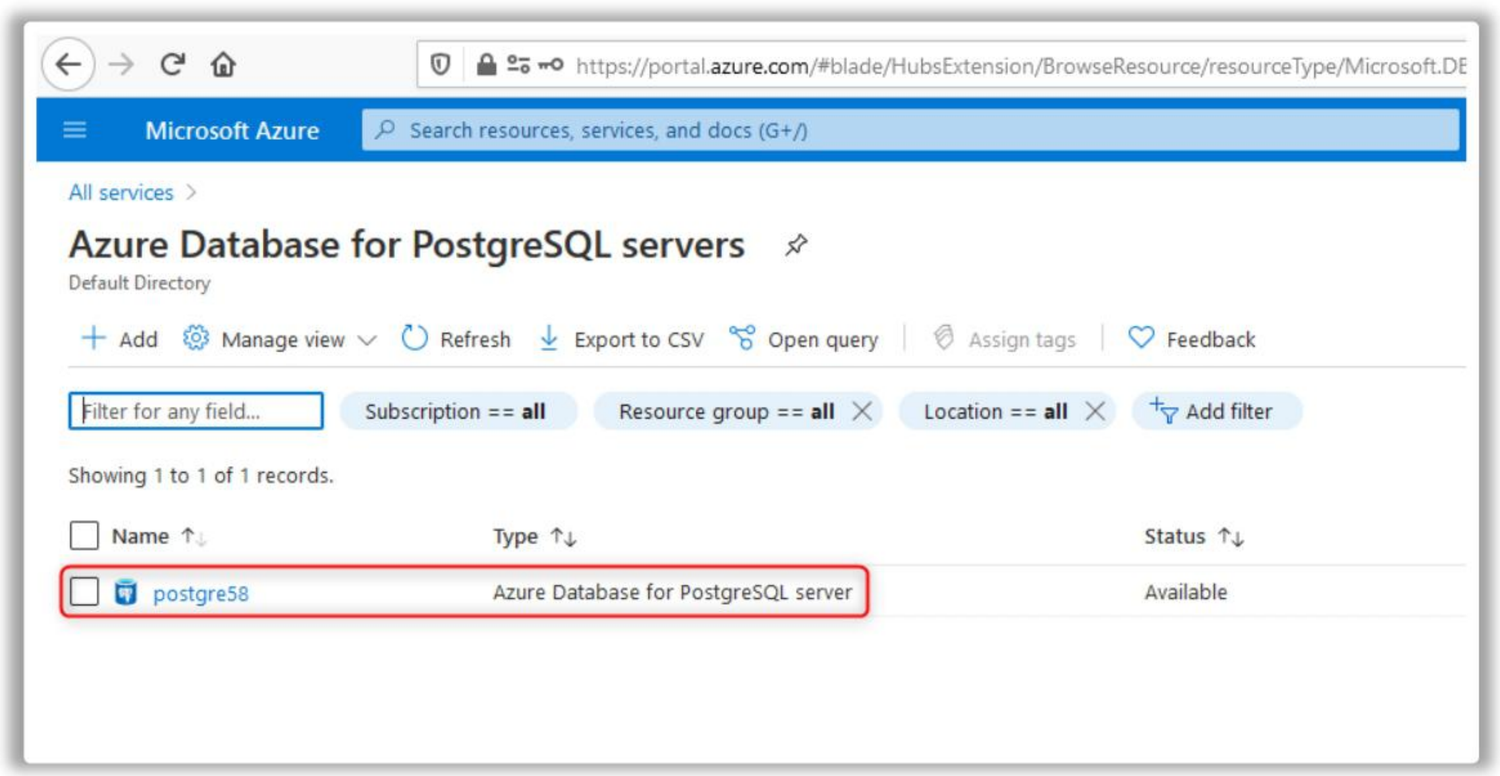


Figure 4.330: Choose Database

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

- Locate **Settings** in the navigation panel and go to **Server Parameters.**

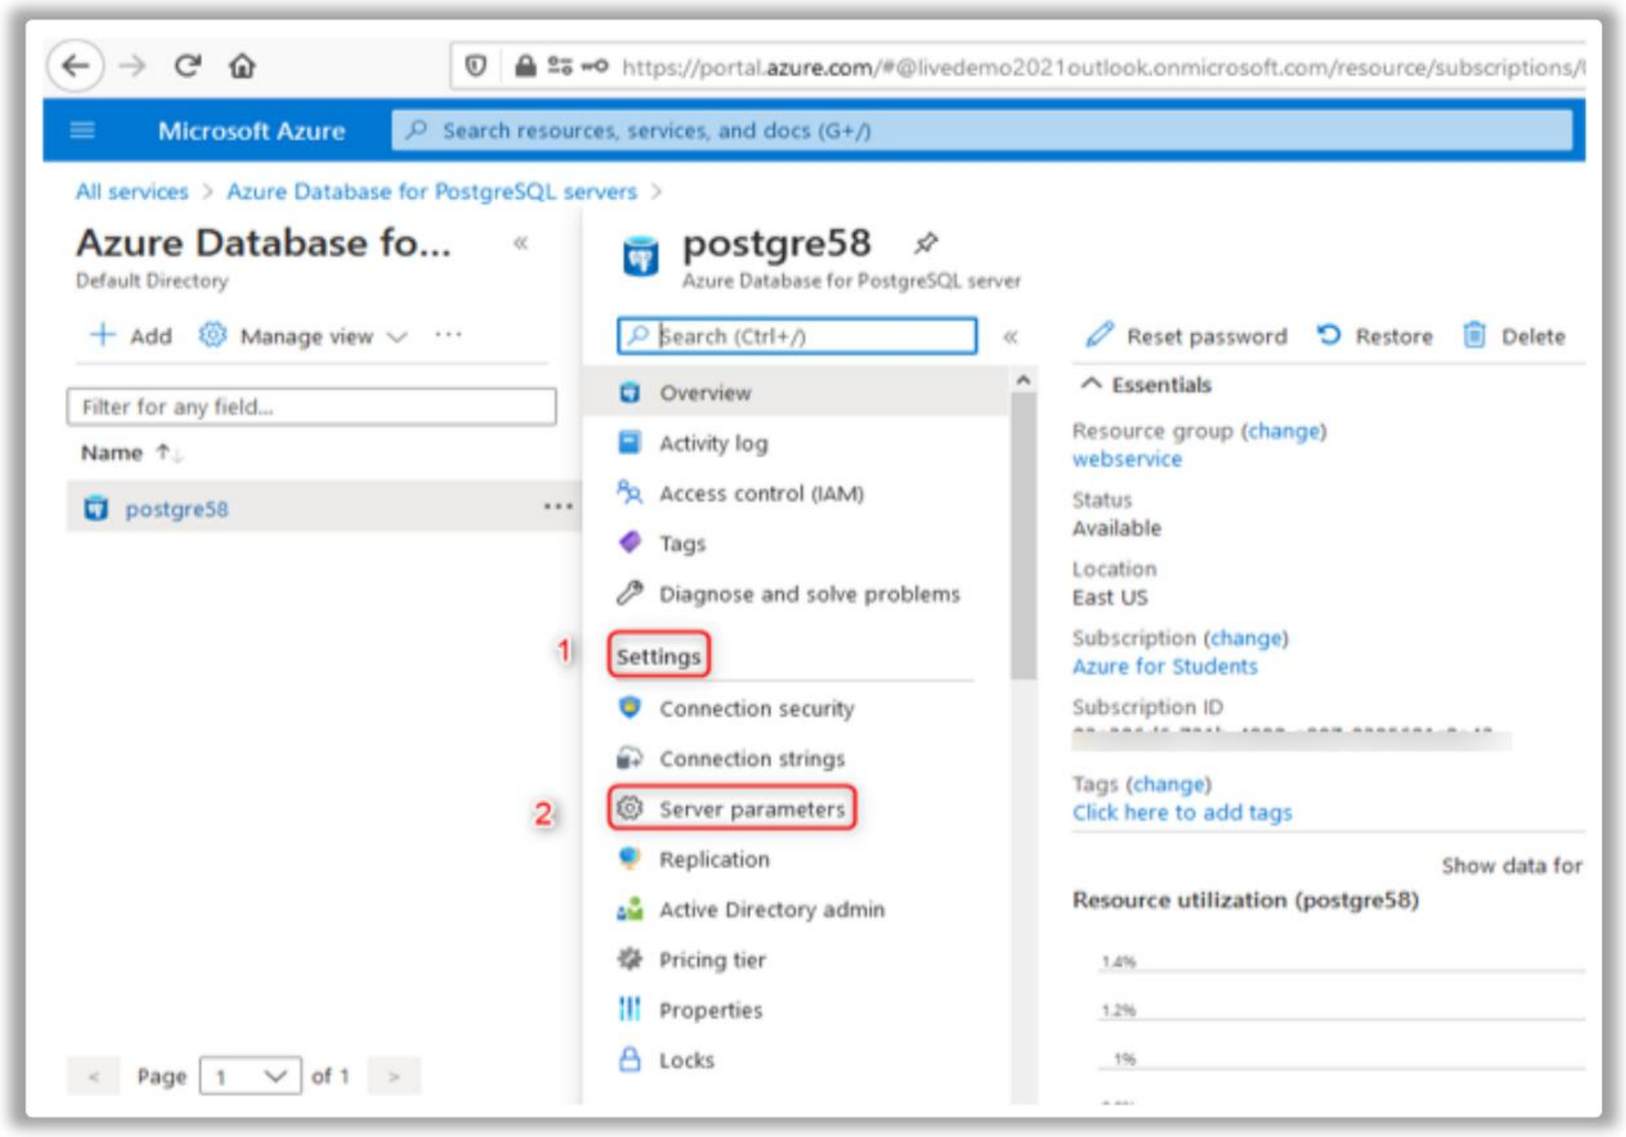


Figure 4.331: Select Server Parameters

- Scroll down to **log_checkpoints** and click on **ON**.
- Click on **Save**.

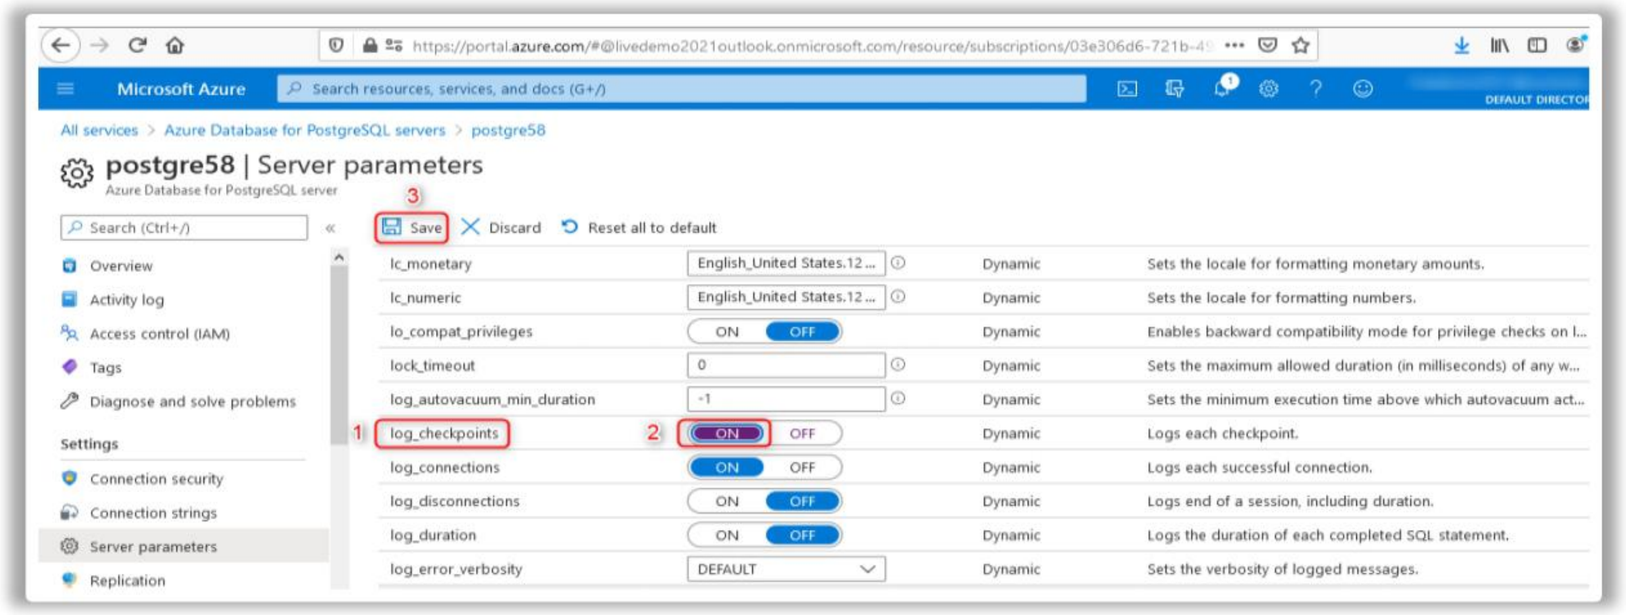


Figure 4.332: Select Save

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

- Log checkpoints has been enabled successfully.

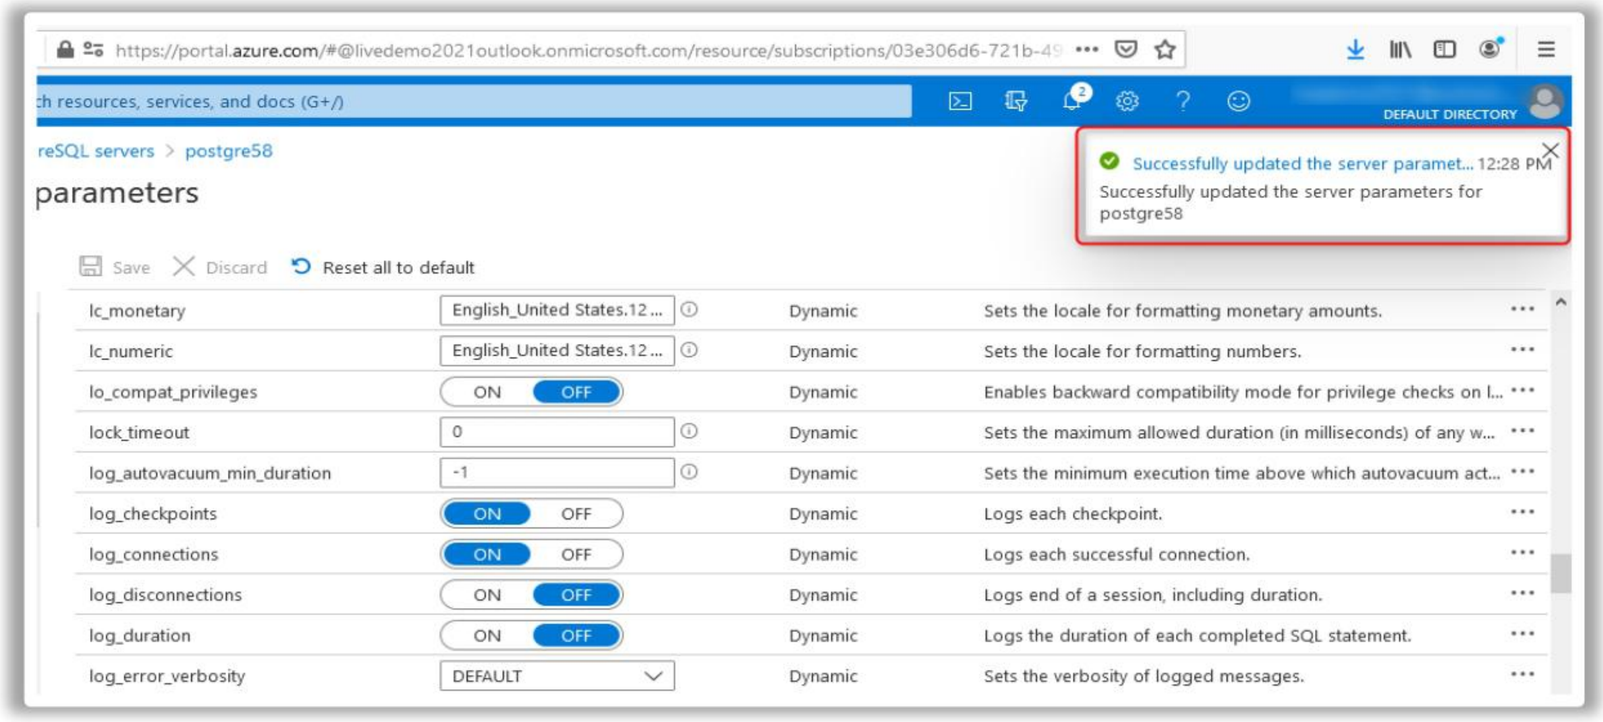


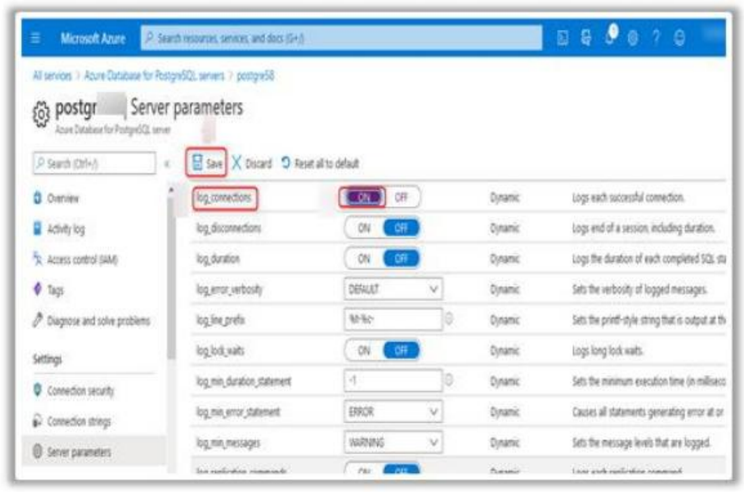
Figure 4.333: Log Checkpoints is Enabled

hide01.ir

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

Enable LOG_CONNECTIONS Parameter for PostgreSQL Servers

- All attempted and successful connections to the **database server logs** are recorded by enabling LOG_CONNECTIONS parameter for PostgreSQL servers
- The **logging data** obtained from LOG_CONNECTIONS parameter helps in identifying, repairing, and troubleshooting all the configuration errors of PostgreSQL database server



Enabling Log Connections Parameters for PostgreSQL Servers

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

Enable LOG_CONNECTIONS Parameter for PostgreSQL Servers

It is recommended to enable log connections server parameter for PostgreSQL database servers in the Azure cloud account. It logs all attempted and successful connections to the database server. It can be changed at starting of the session by Azure users with administrator privileges only. It is not possible to change these logs during the session. The logging data obtained from LOG_CONNECTIONS parameter helps in identifying, repairing, and troubleshooting all the configuration errors of PostgreSQL database server

Source: https://azure.microsoft.com

To Enable LOG_CONNECTIONS Parameters for PostgreSQL Servers

- Sign-in to Azure portal.
- Click on **More services**.

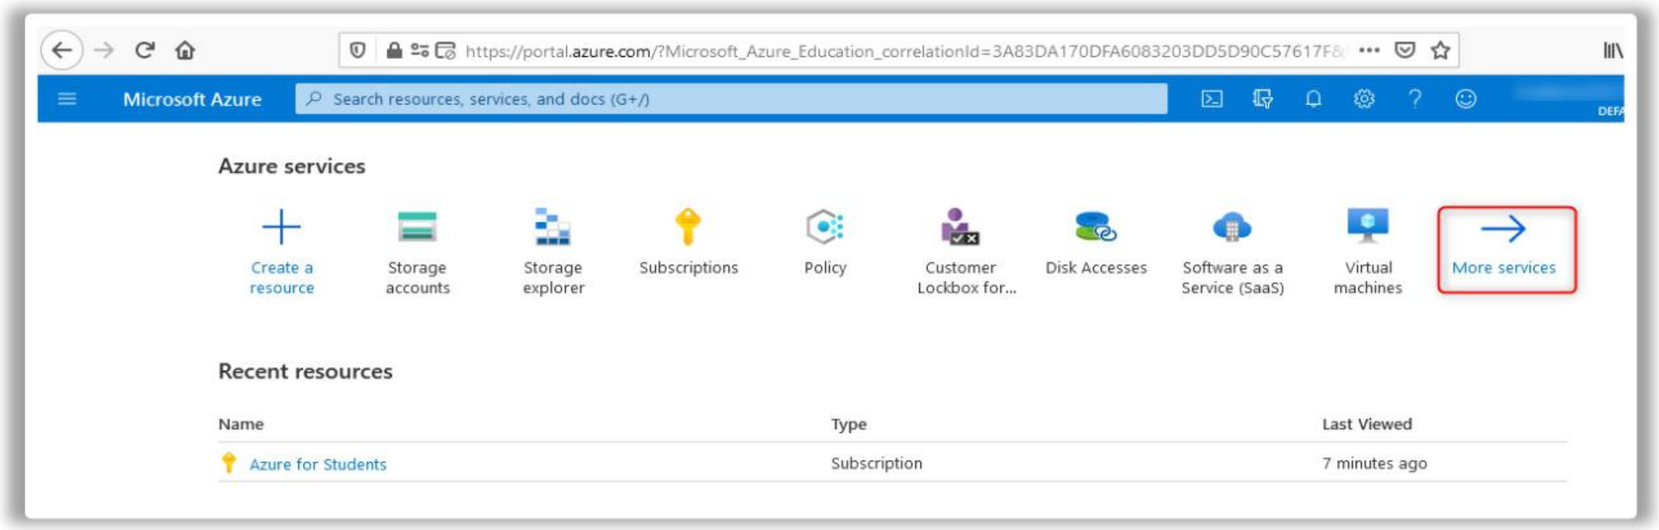


Figure 4.334: Select More Services

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

- Search for **Azure Database for PostgreSQL servers** using the input field beside **All services**, then click on the same in the right-side panel.

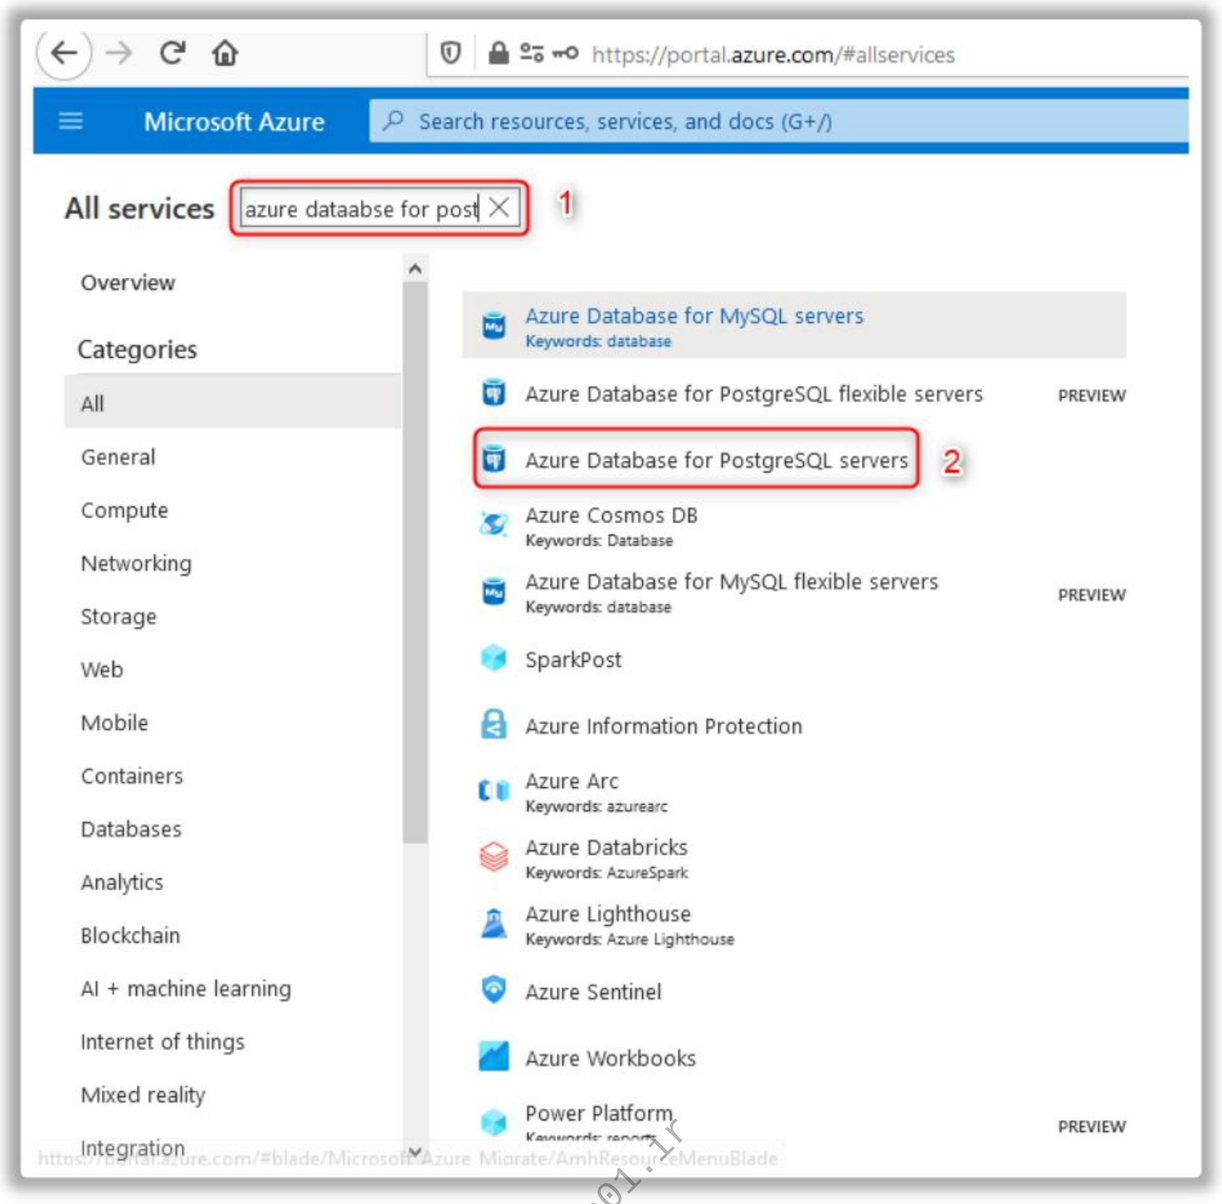


Figure 4.335: Search Database

- A list of PostgreSQL servers will be displayed.
- Select the PostgreSQL server that you want to enable log connections.

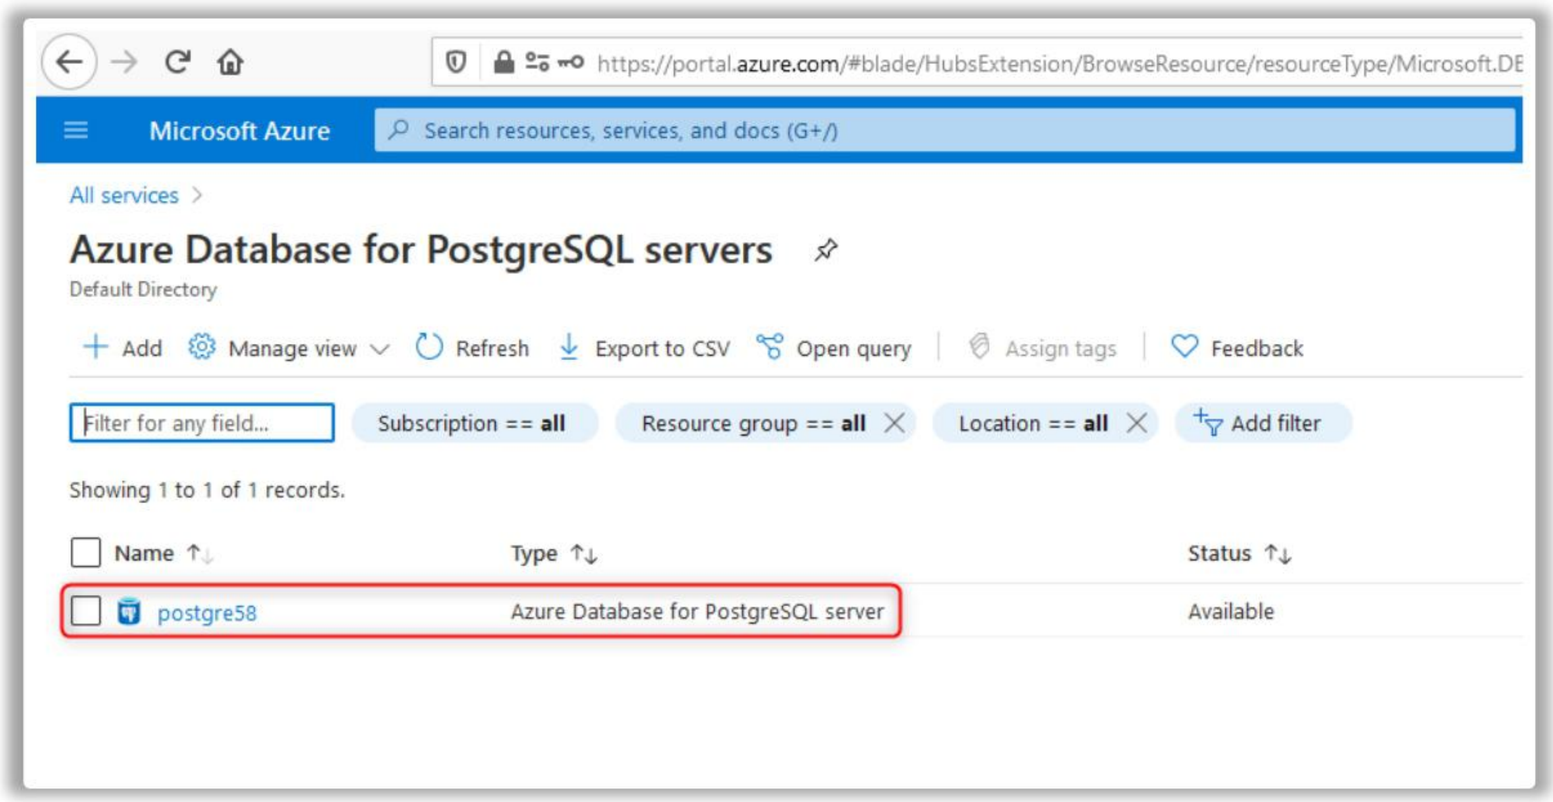


Figure 4.336: Choose Database

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

- Locate **Settings** in the navigation panel and go to **Server Parameters.**

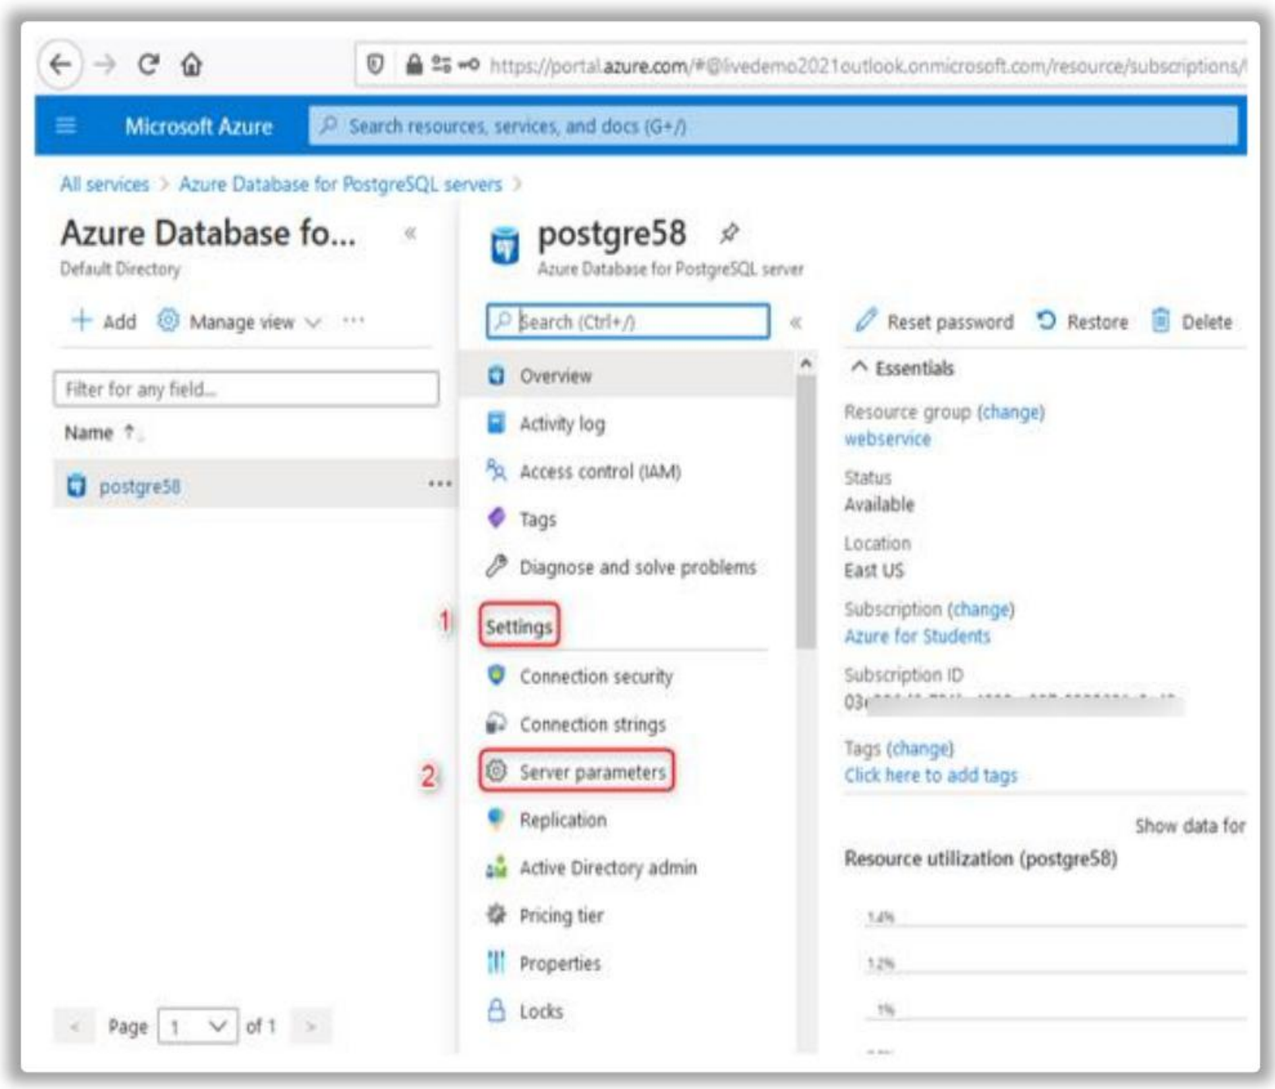


Figure 4.337: Select Server Parameters

- Scroll down to **log_connections** and click on **ON**.
- Click on **Save.**

hide01.ir

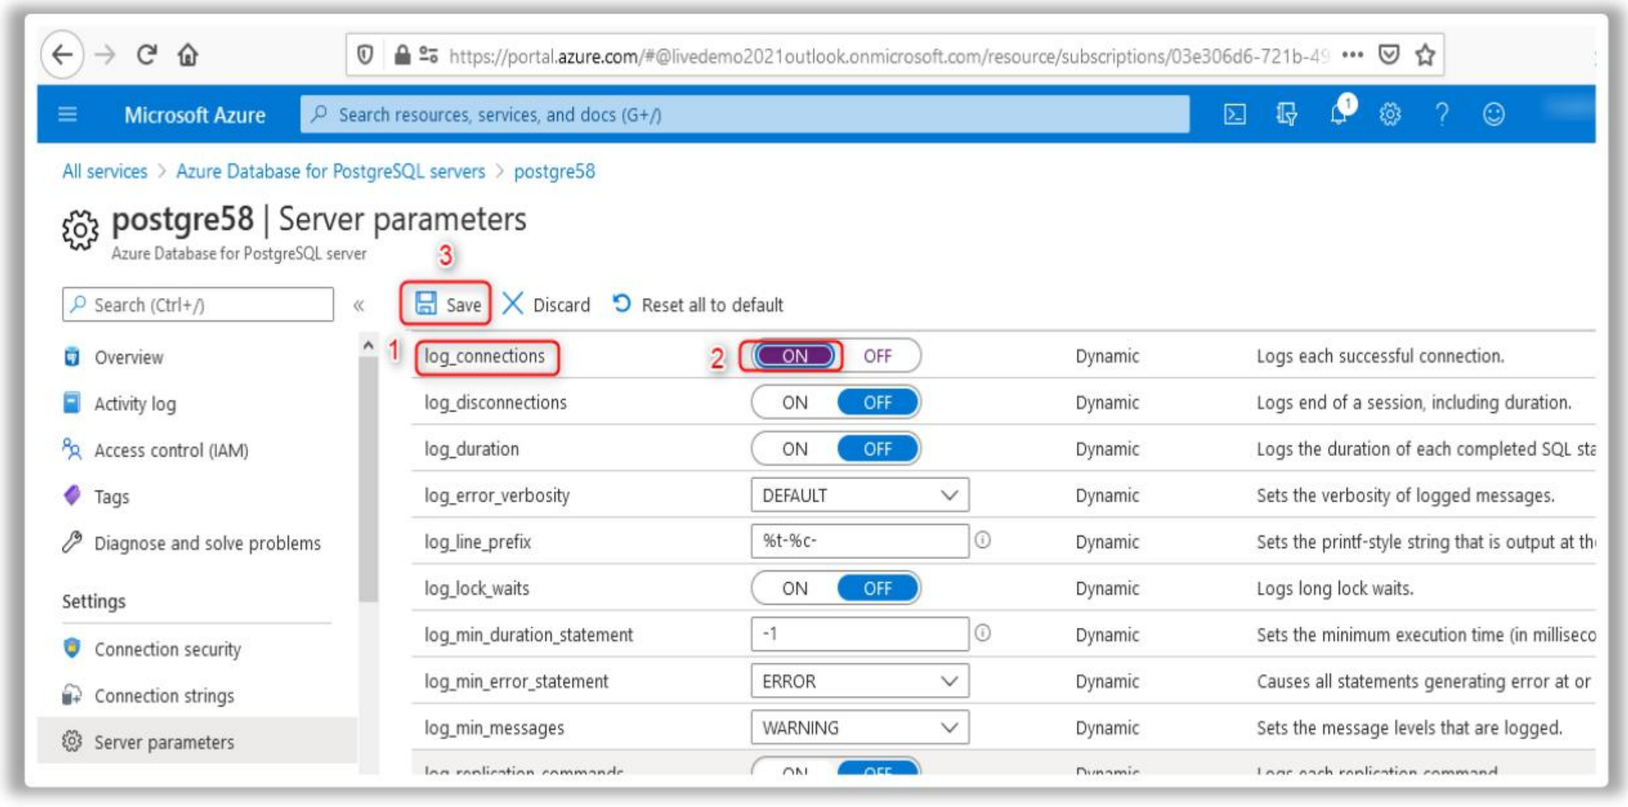


Figure 4.338: Choose Server Parameter

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

- Log connections have been enabled successfully.

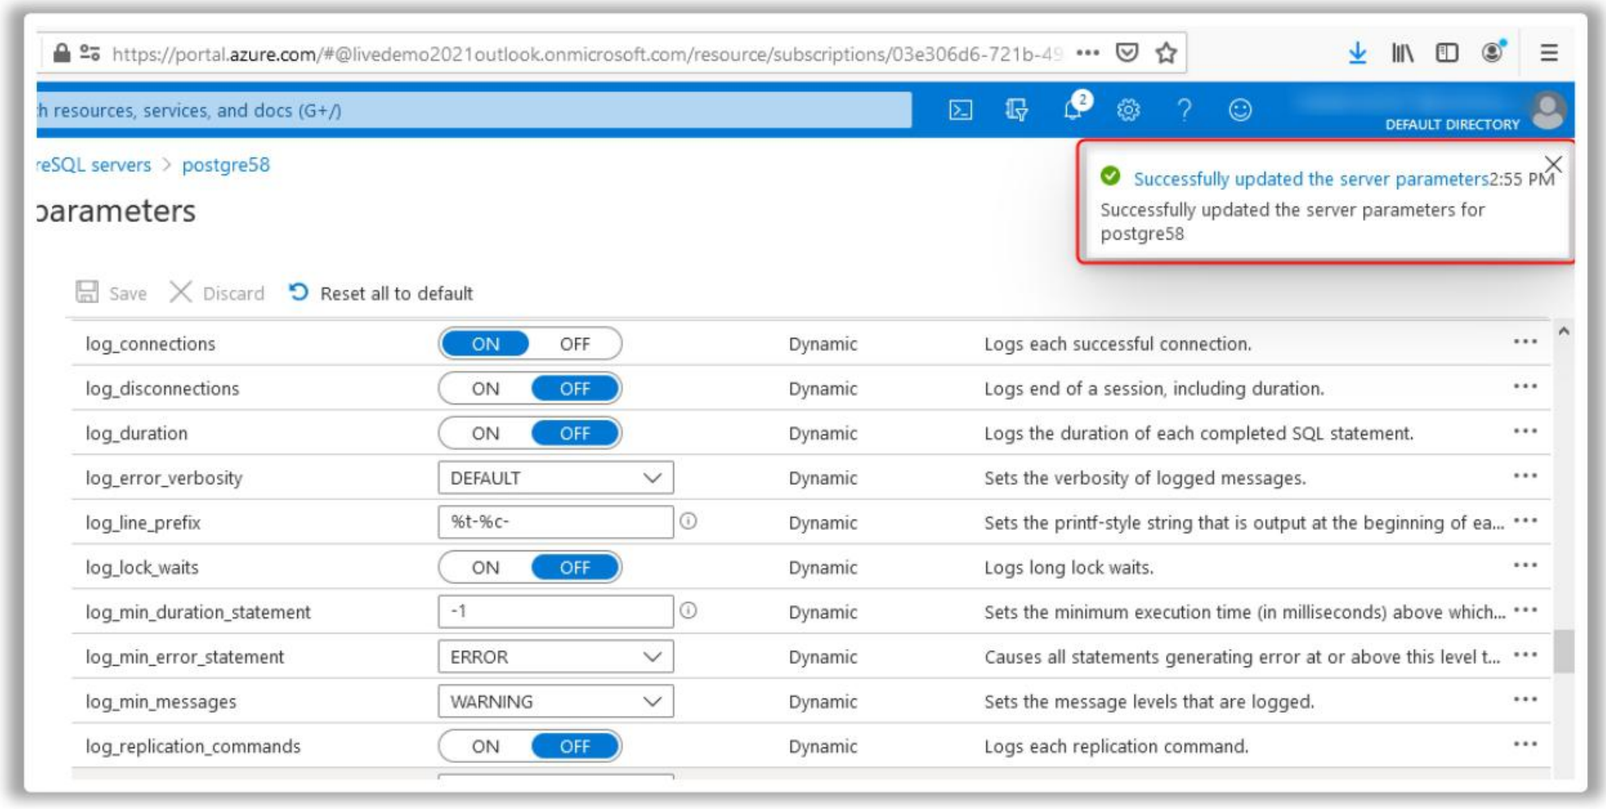


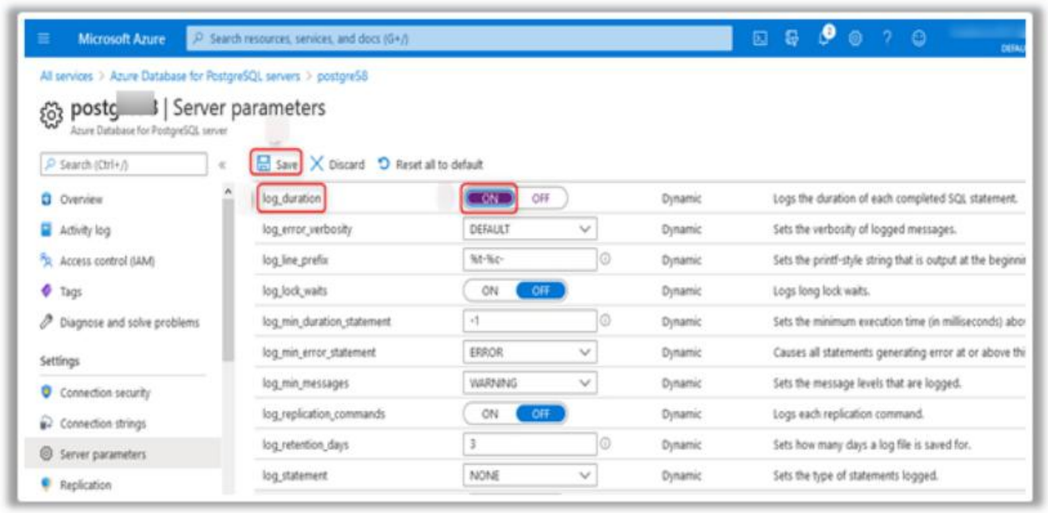
Figure 4.339: Log Connection is Enabled

hide01.ir

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

Enable LOG_DURATION Parameter for PostgreSQL Servers

- Ensure to **enable** LOG_DURATION parameter for PostgreSQL servers to logs the duration of executed PostgreSQL statements
- The LOG_DURATION parameter will generate a **database query** and error logs that are utilized in identifying, repairing, troubleshooting configuration errors and sub-optimal performance of PostgreSQL database server



Enabling Log Duration Parameters for PostgreSQL Servers

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

Enable LOG_DURATION Parameter for PostgreSQL Servers

Ensure to enable LOG_DURATION parameter for PostgreSQL servers to logs the duration of executed PostgreSQL statements. Only administrator privileged users can modify this setting in the Azure PostgreSQL server. Parse, Execute, and Blind steps performed are logged with the duration. The LOG_DURATION parameter will generate a database query and error logs that are utilized in identifying, repairing, troubleshooting configuration errors and sub-optimal performance of PostgreSQL database server

Source: https://azure.microsoft.com

To Enable Log Duration Parameter for PostgreSQL Servers

- Sign-in to Azure portal.
- Click on **More services**.

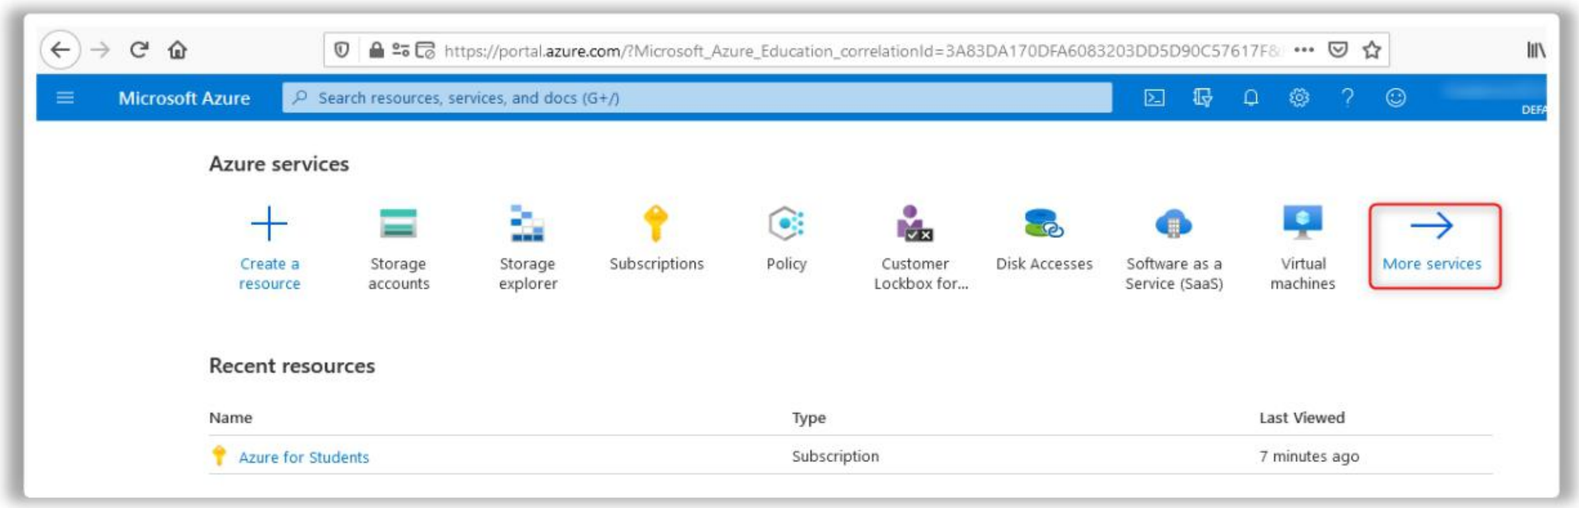


Figure 4.340: Select More Service

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

- Search for **Azure Database for PostgreSQL servers** using the input field beside **All services**, then click on the same in the right-side panel.

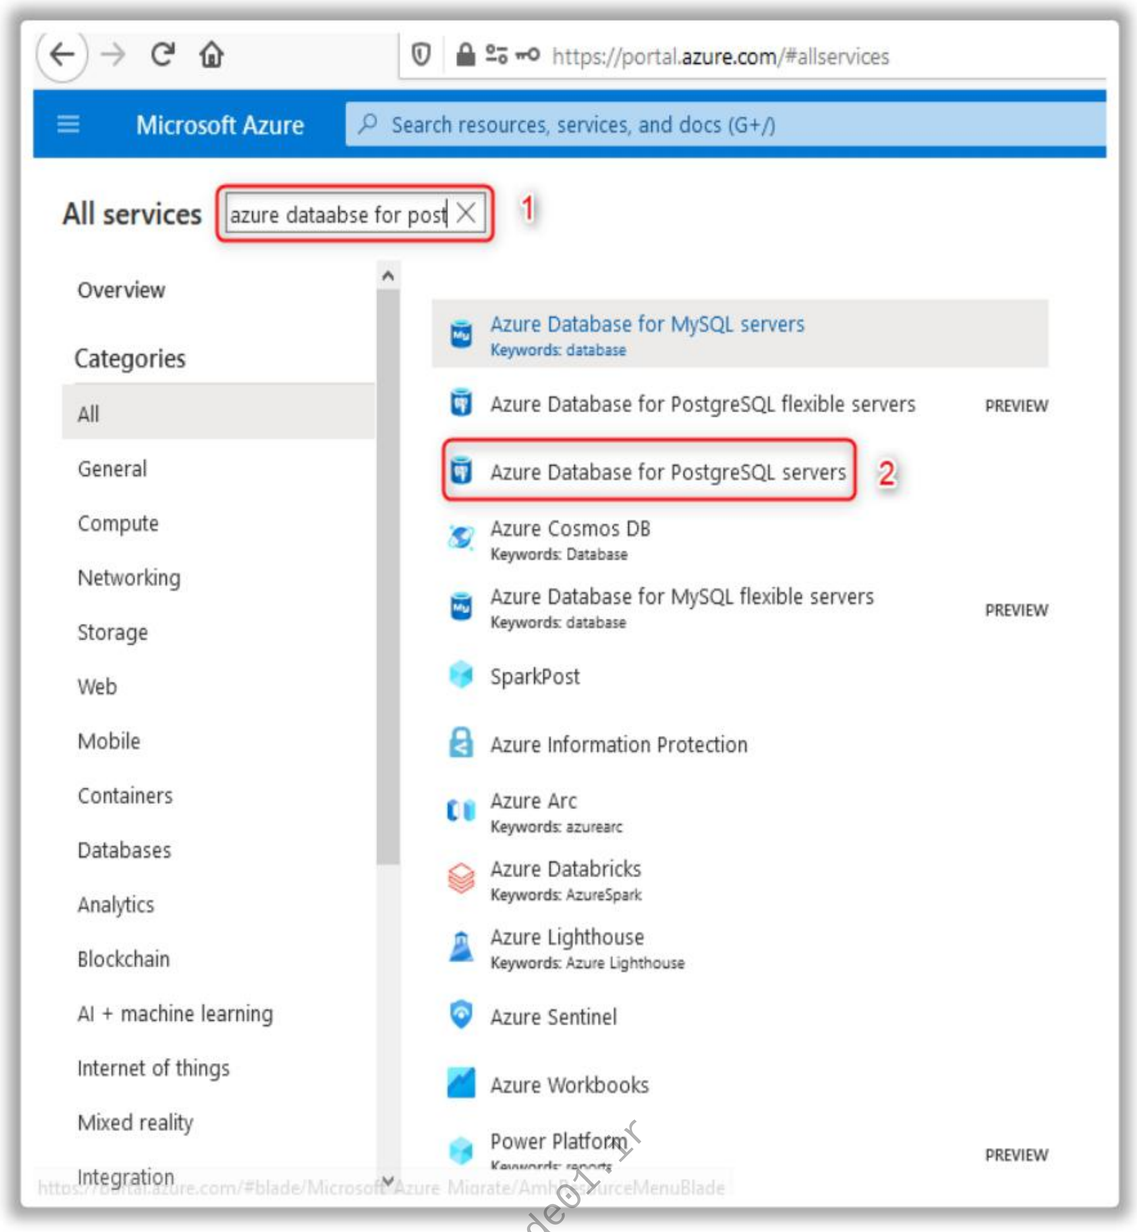


Figure 4.341: Search Database

- A list of PostgreSQL servers will be displayed.
- Select the PostgreSQL server that you want to enable log duration.

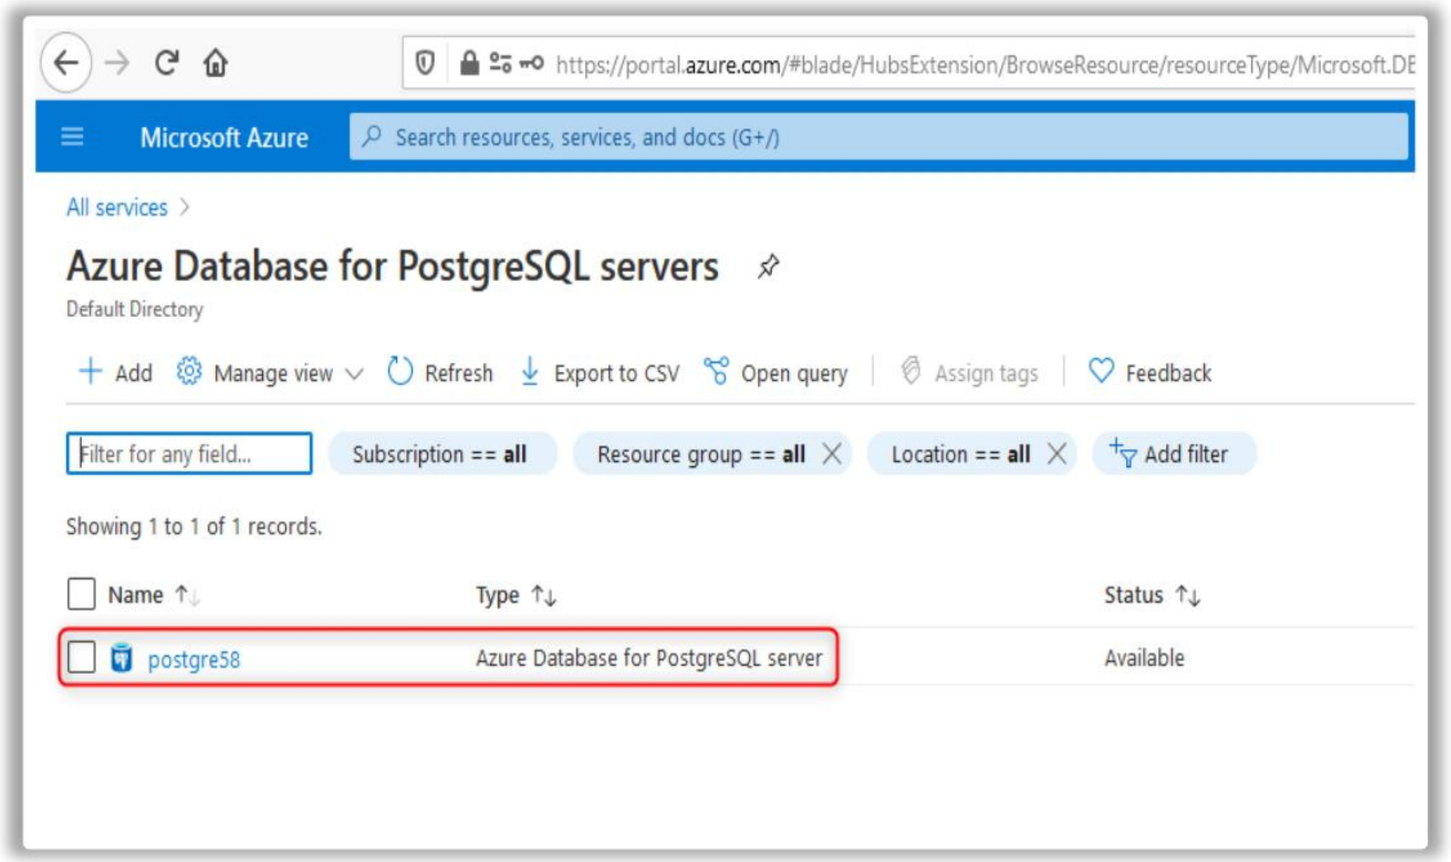


Figure 4.342: Choose Database

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

- Locate **Settings** in the navigation panel and go to **Server Parameters.**

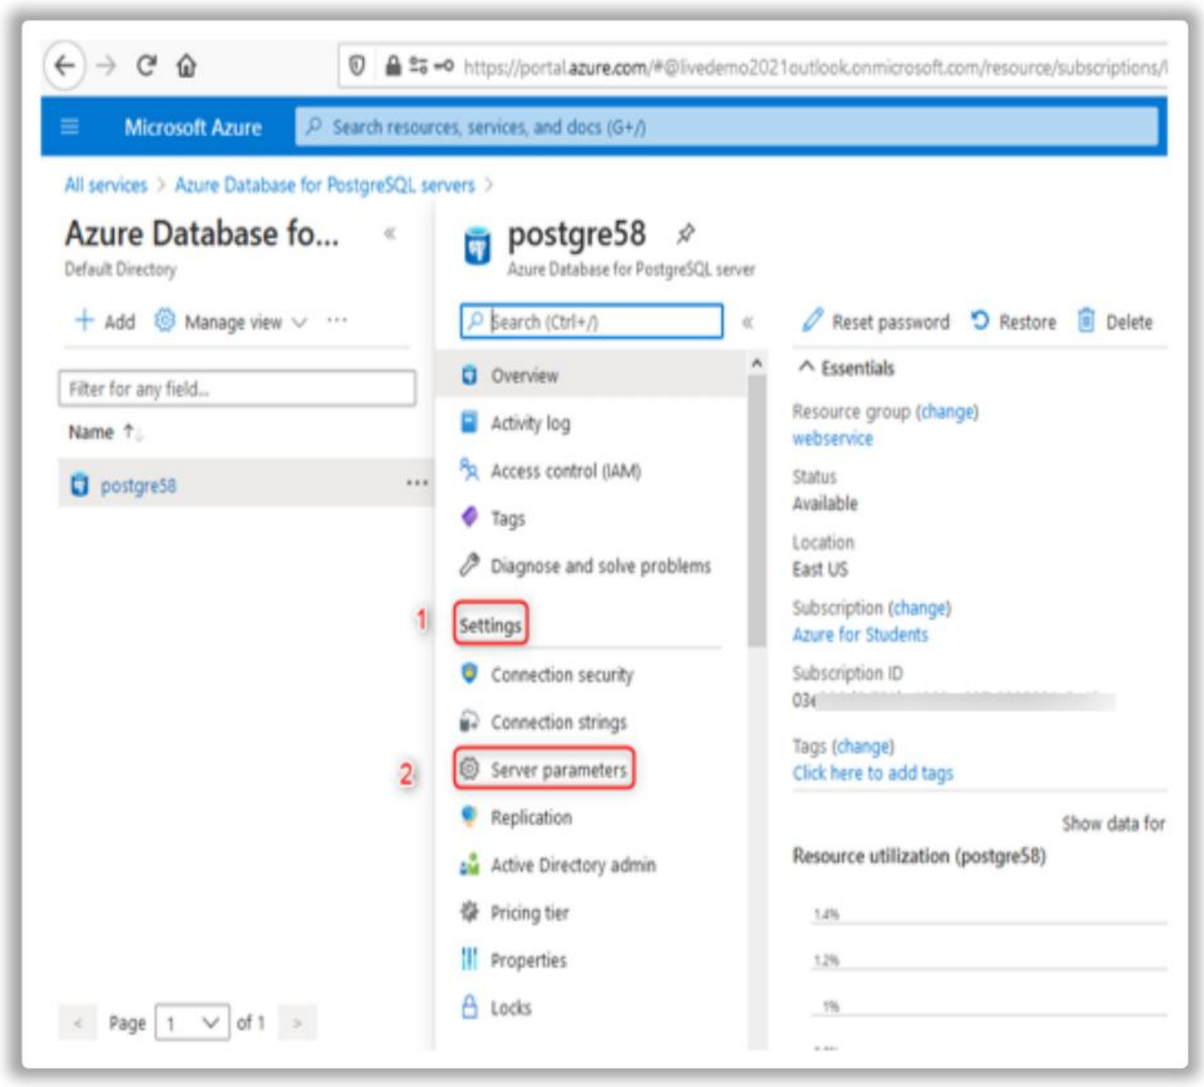


Figure 4.343: Select Parameterrs

- Scroll down to **log_duration** and click on **ON**.
- Click on **Save.**

hide01.ir

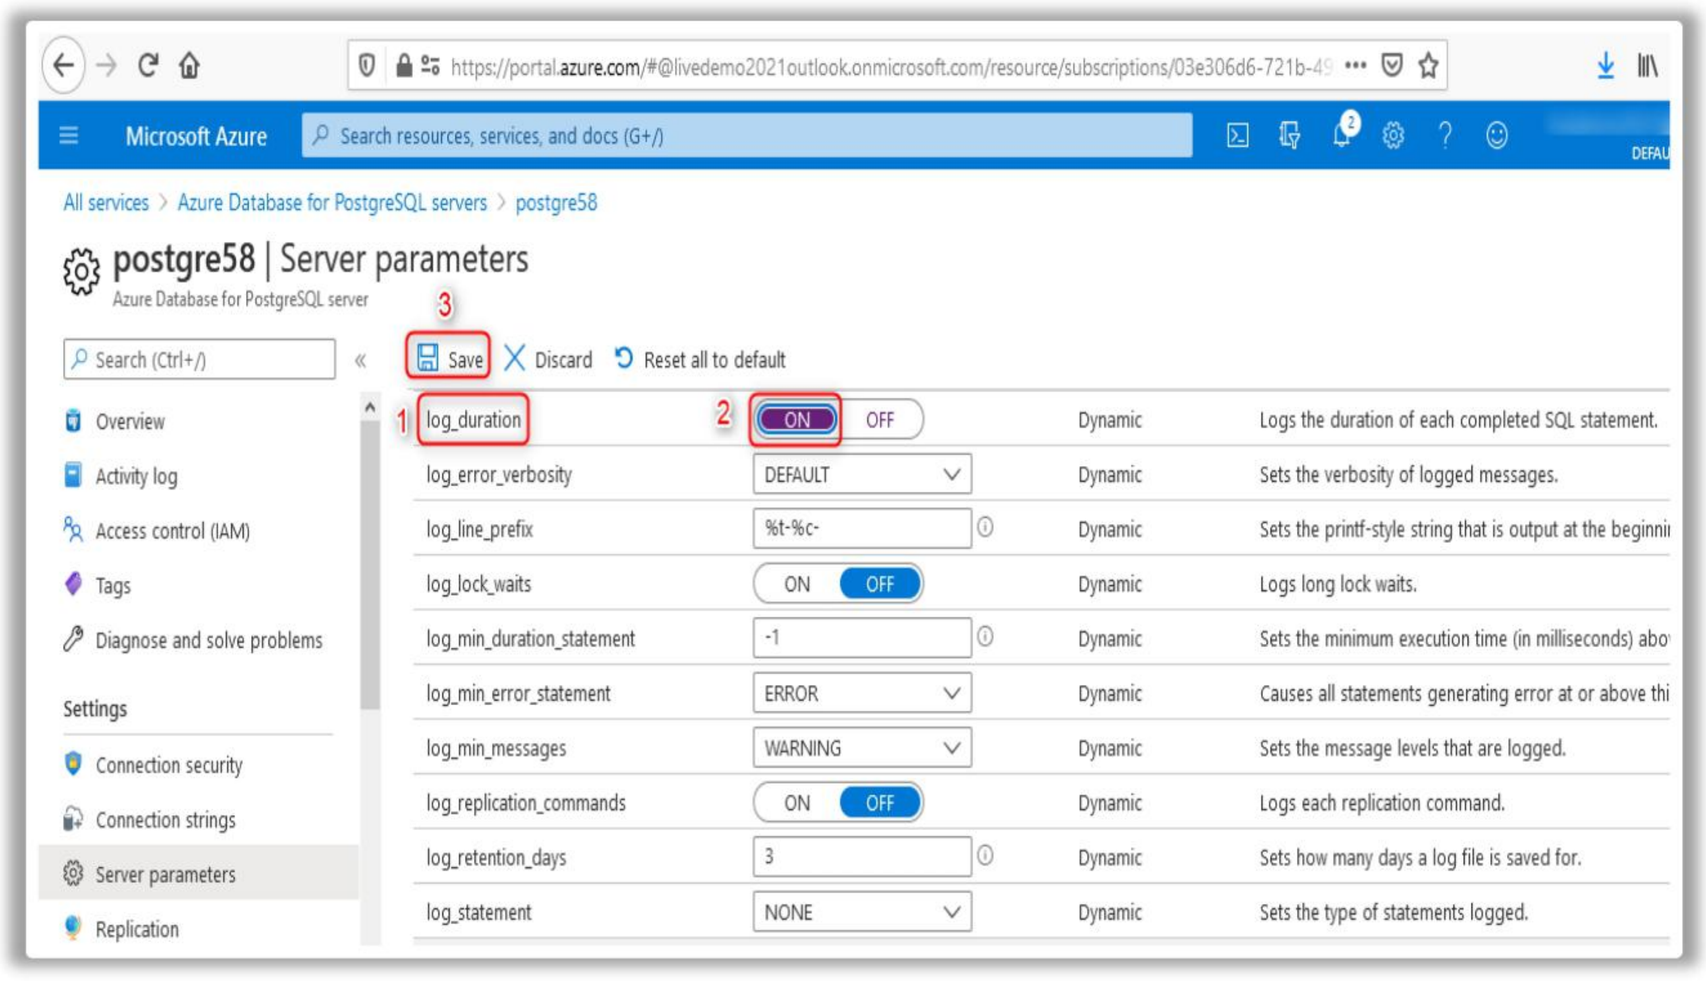


Figure 4.344: Click Save

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

- Log duration has been successfully updated.

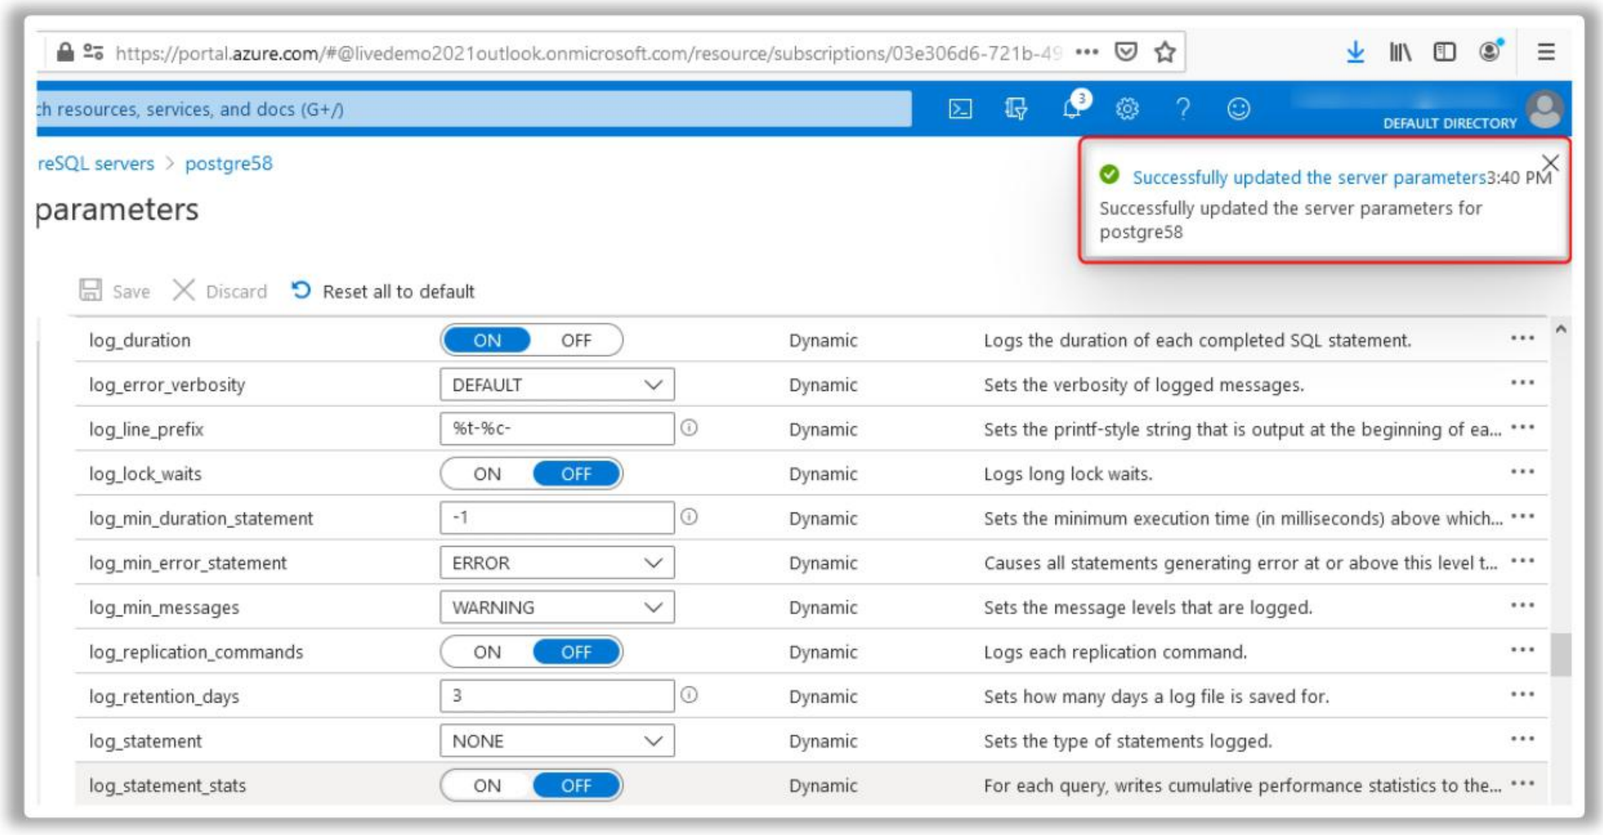
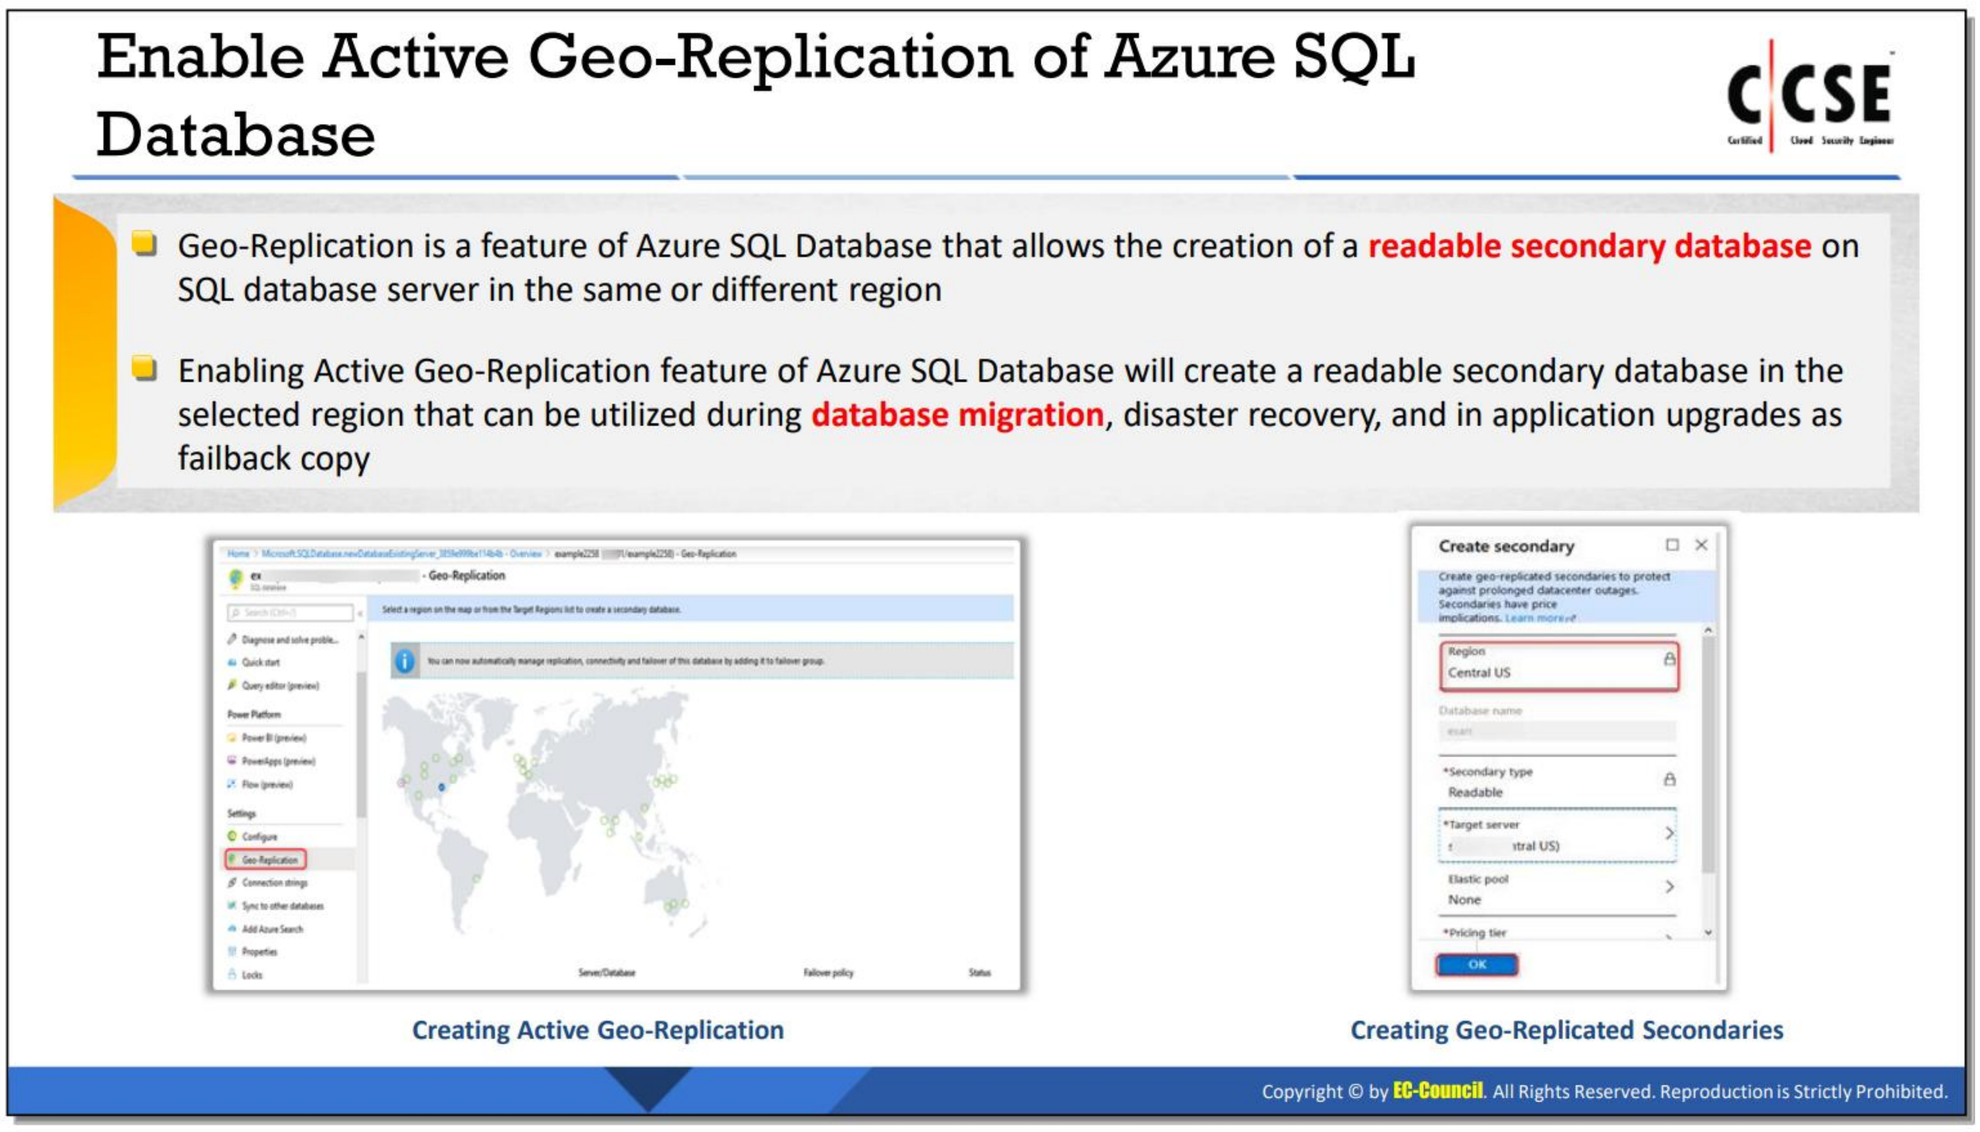


Figure 4.345: Log Duration is Updated

hide01.ir

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



Enable Active Geo-Replication of Azure SQL Database

Active geo-replication is a feature of Azure SQL Database that allows the creation of a readable secondary database on SQL database server in the same or different region. Active geo-replication is a business continuity strategy for quick disaster recovery of individual databases during a regional disaster.

hi2e01.ir

Geo-replication supports four secondary databases, which are used for read-only access queries. Enabling geo-replication allows the application to start failover to a secondary database. After failover, the secondary database will become the primary database with the different connection end points.

Benefits of Geo-replication

- Used in disaster recovery.
- Used in database migration to migrate databases from one server to another with minimum downtime.
- Used in application upgrades, as a failback copy.

Source: https://azure.microsoft.com

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

Enabling Active Geo-Replication in Microsoft Azure

- From Azure Home page, click **SQL database.**

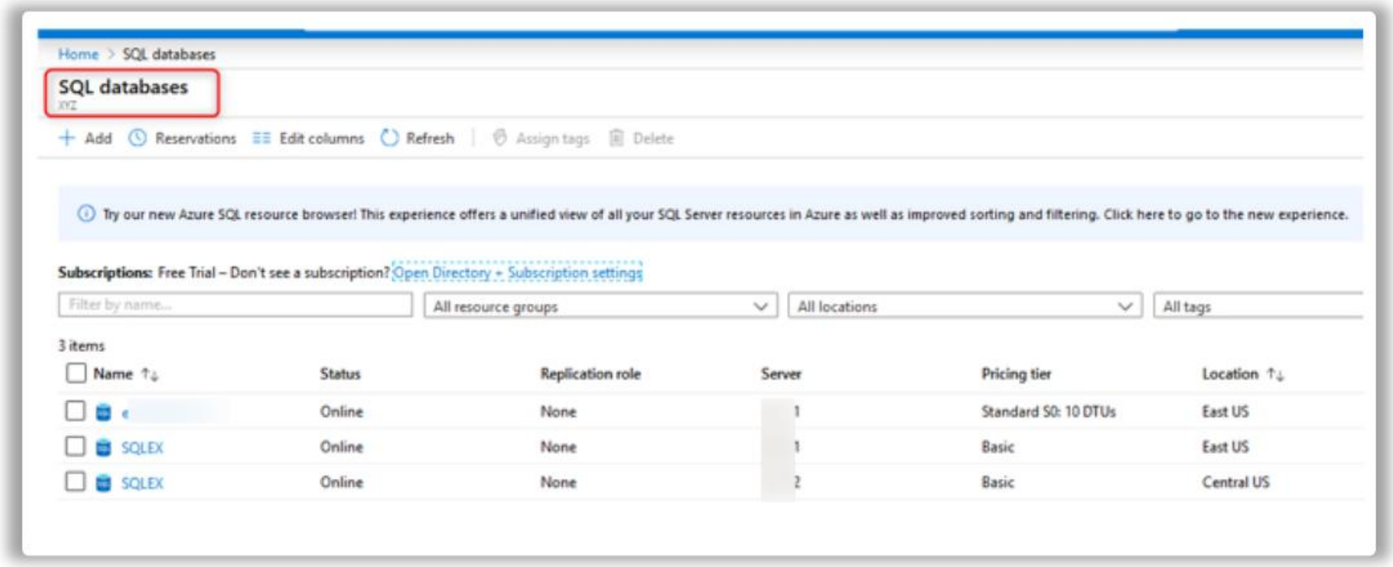


Figure 4.346: Click SQL Database

- **Create SQL Database**, fill the details and click **Review + create.**

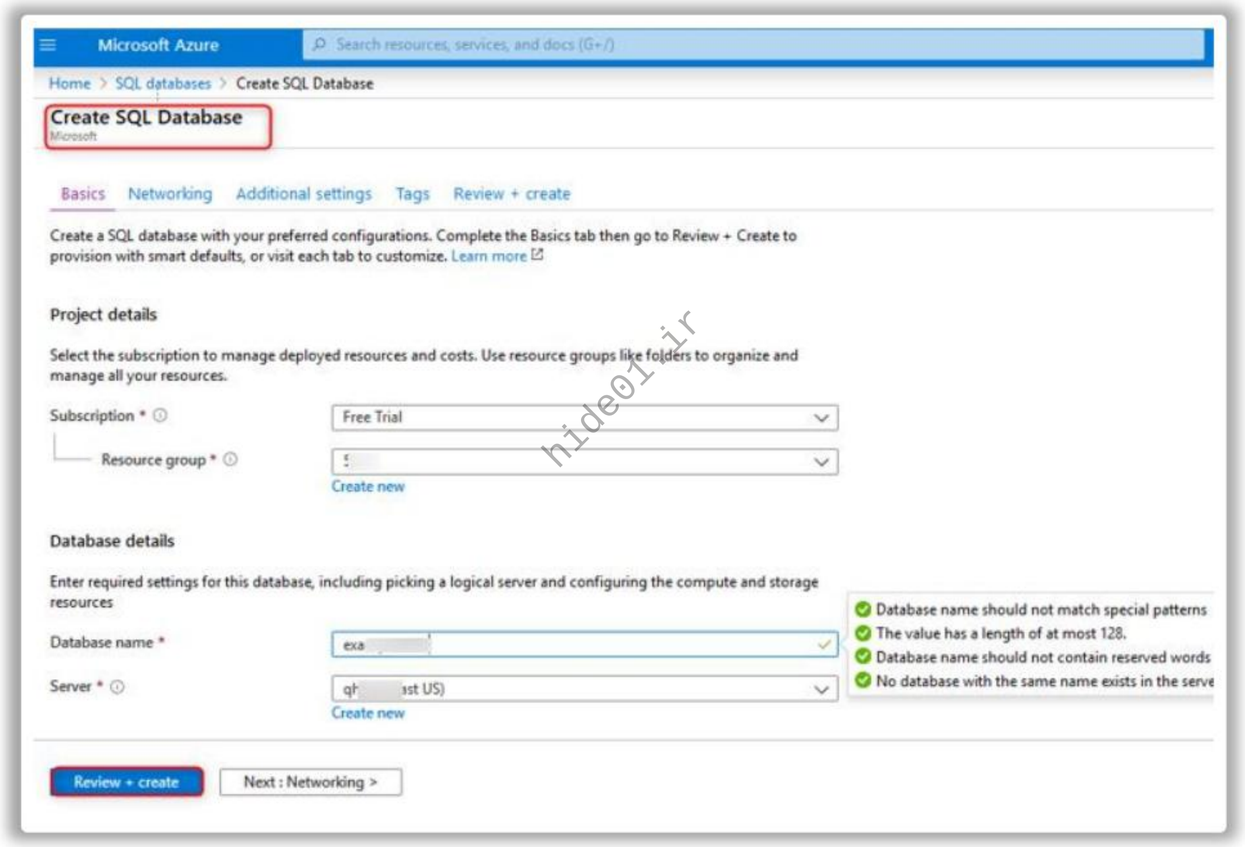


Figure 4.347: Click Review + create

- After deployment, click **Go to Resource.**

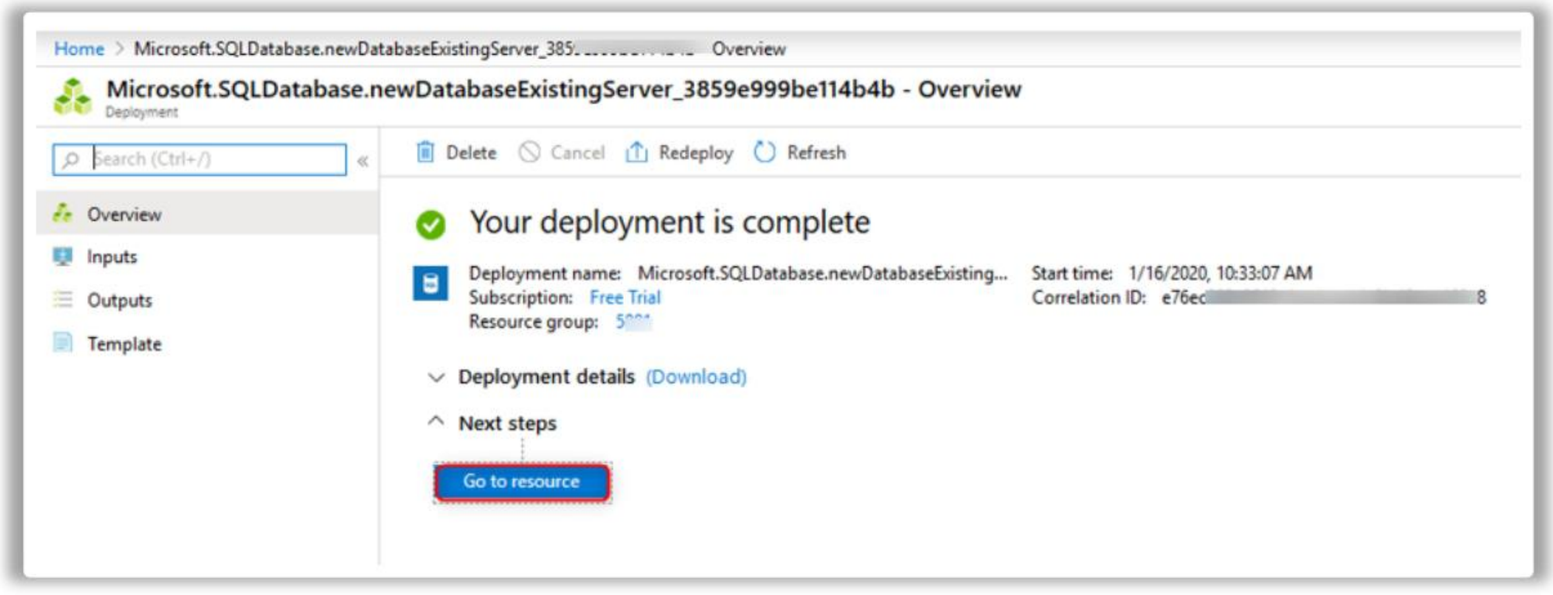


Figure 4.348: Click Go to Resource

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

- In **Settings** section, click **Geo-Replication** to create readable secondary database.

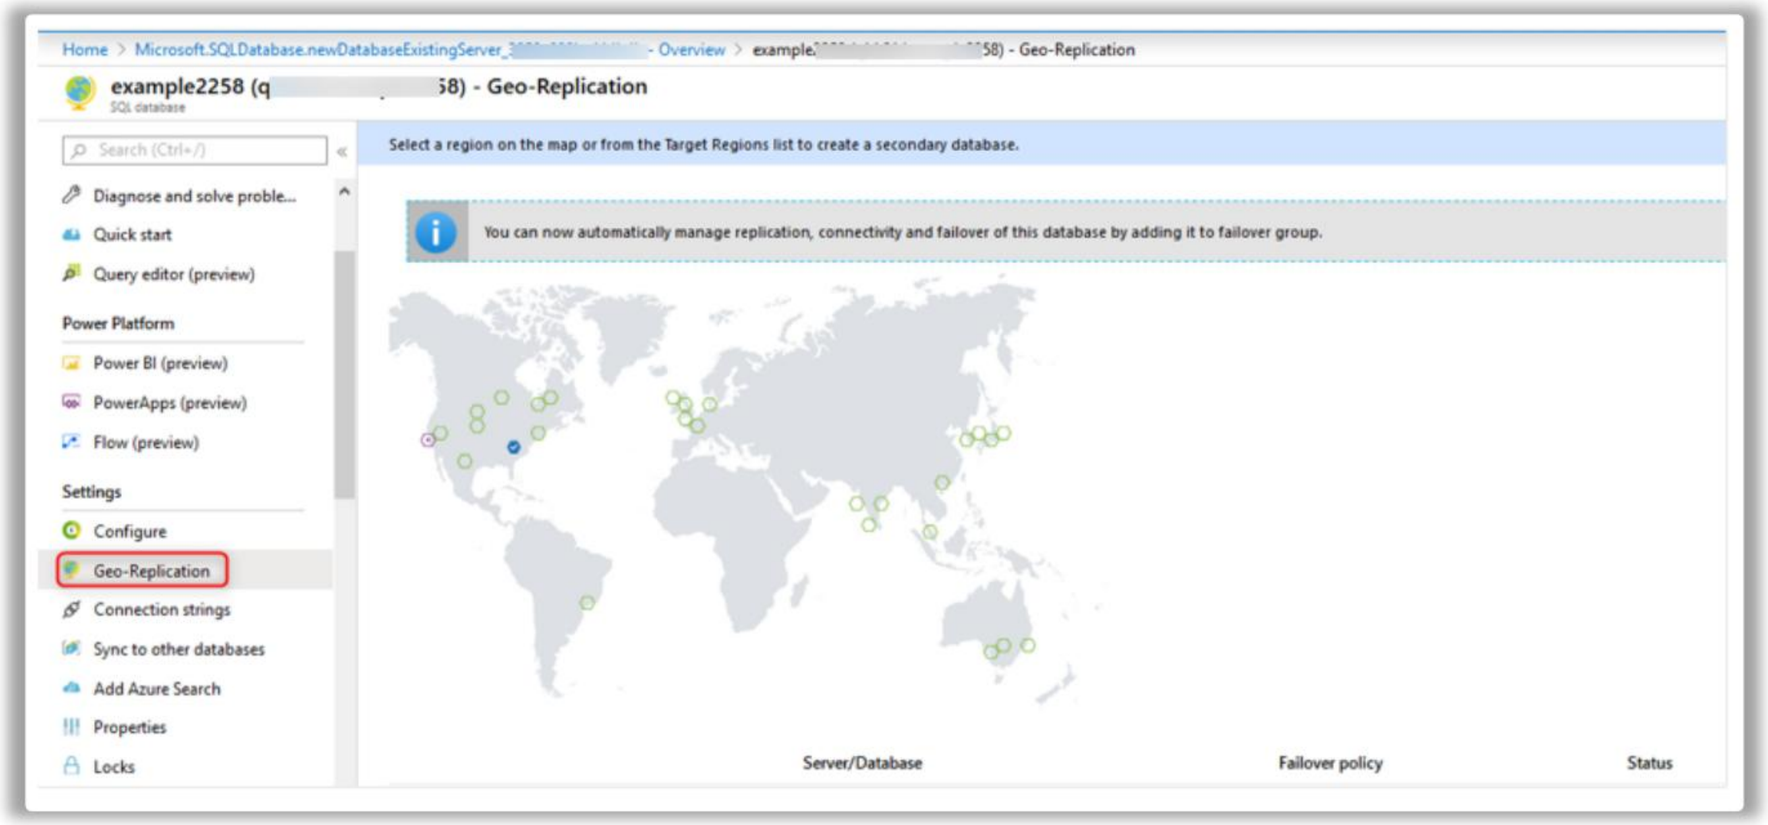


Figure 4.349: Click Geo-Replication

- Select **Target region** and fill details; click **Ok.**

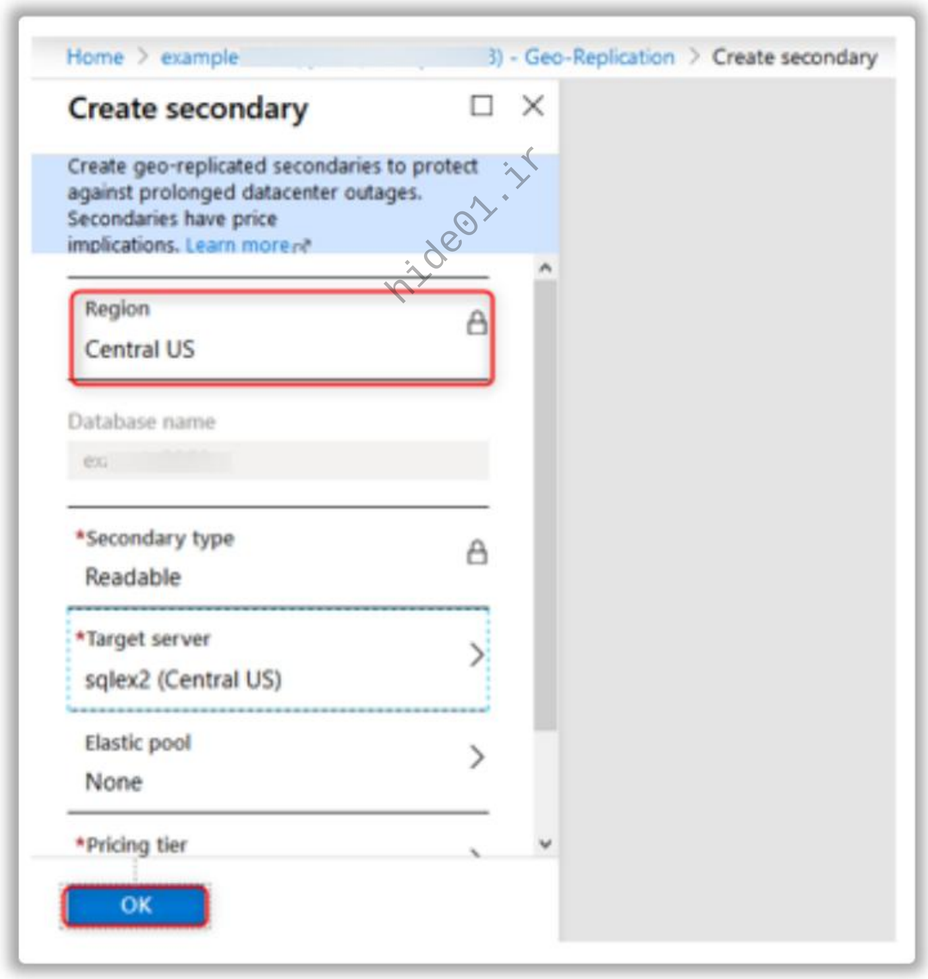


Figure 4.350: Click Ok

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

- Secondary database is created.

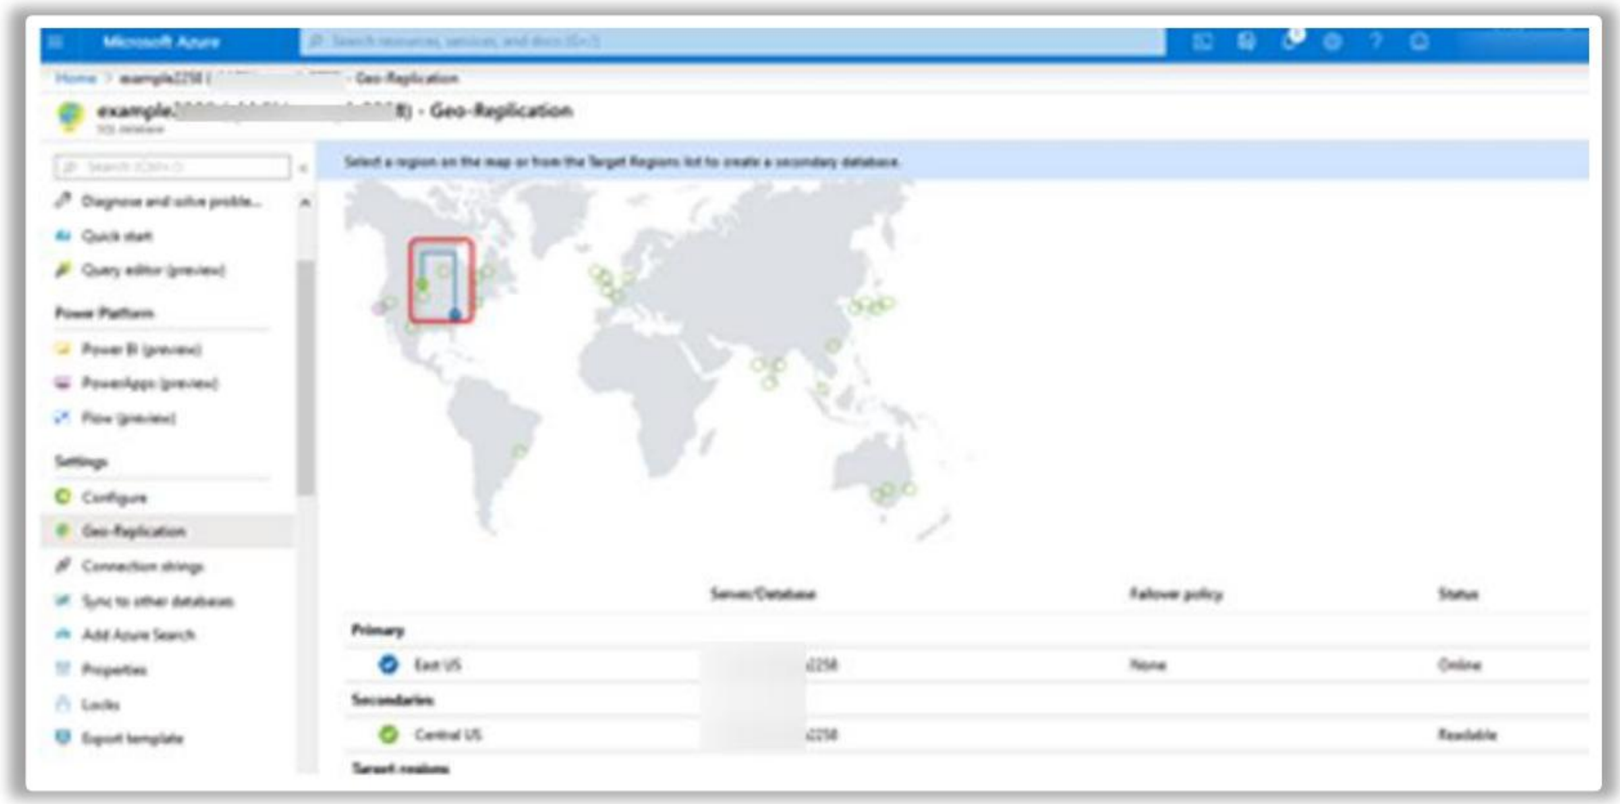


Figure 4.351: Secondary Data Created

hide01.ir

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

Configure Azure Defender for Storage

- Azure Defender for Storage provide an additional layer of **protection** to storage accounts by detecting any unusual/malicious attempts or exploits in storage accounts



Configuring Azure Defender for Storage

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

Configure Azure Defender for Storage

To provide an additional layer of protection for storage accounts, Azure Defender for Storage is used. It detects any unusual, malicious attempts or exploits in storage accounts. This additional layer of security can be used without any expertise in the monitoring of security systems. The security alerts are sent to emails of subscription administrators, which contain details of the unusual activity and remediation recommendations. These alerts can be viewed in storage activity through Storage Analytics Logging.

Source: https://azure.microsoft.com

To Configure Azure Defender for Storage

- Sign-in to the Azure portal.
- Select **Storage accounts** under **Azure services.**

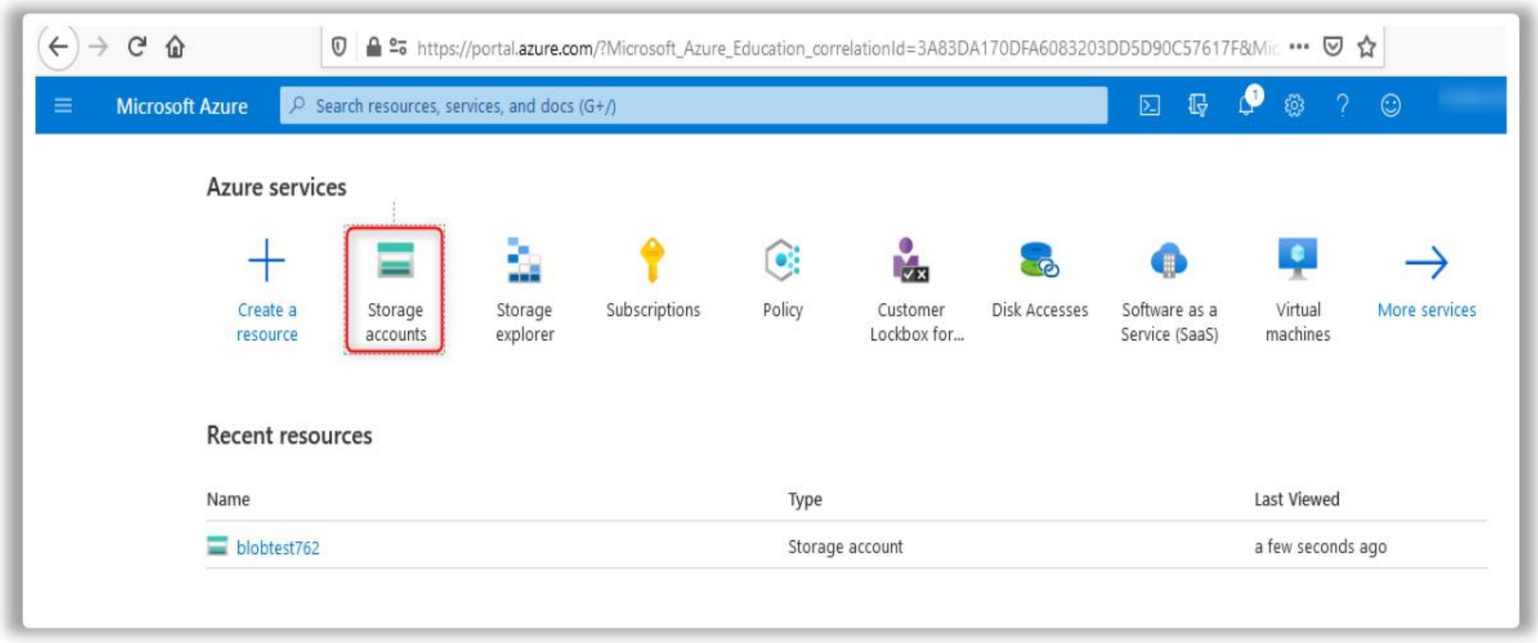


Figure 4.352: Select Storage Accounts

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

- A list of storage accounts will be displayed.
- Select the storage for which you want to enable Azure defender.

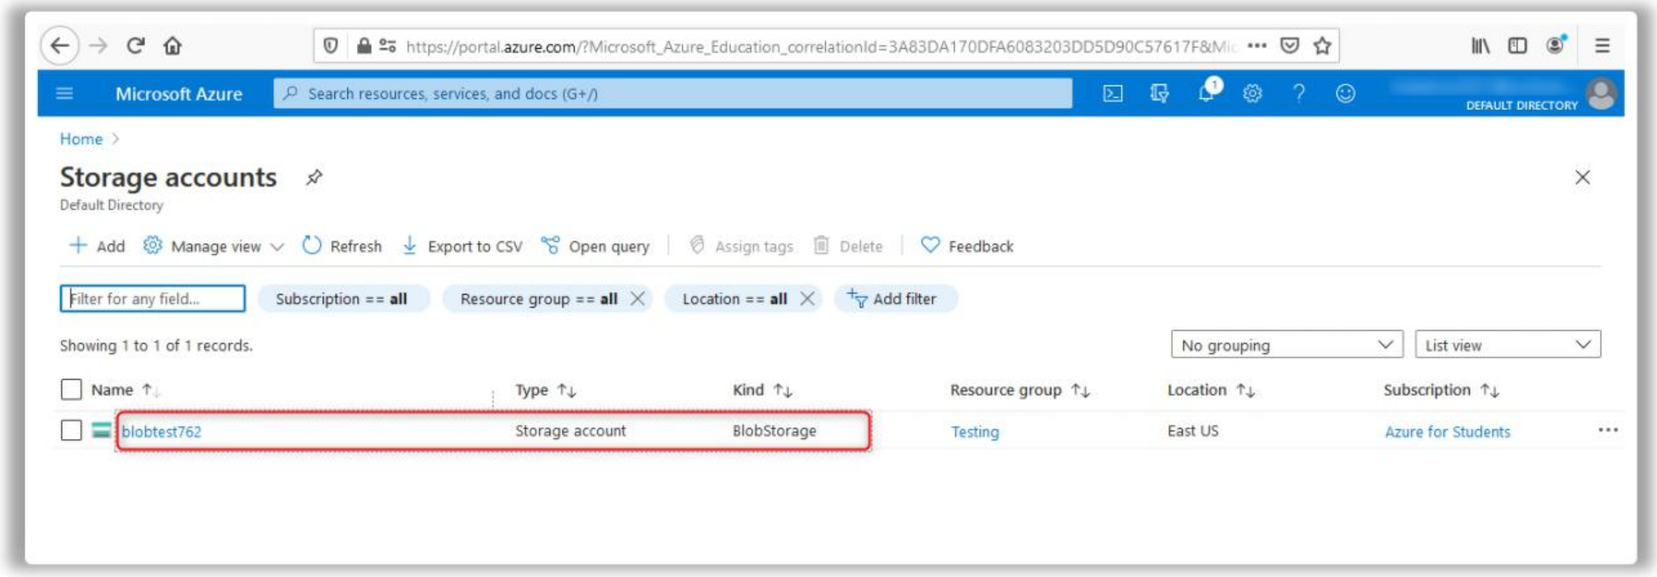


Figure 4.353: Choose Storage

- Under **Settings**, select **Advanced security**.
- Click **Enable Azure Defender for Storage**.
- Azure Defender is now enabled for your storage account.

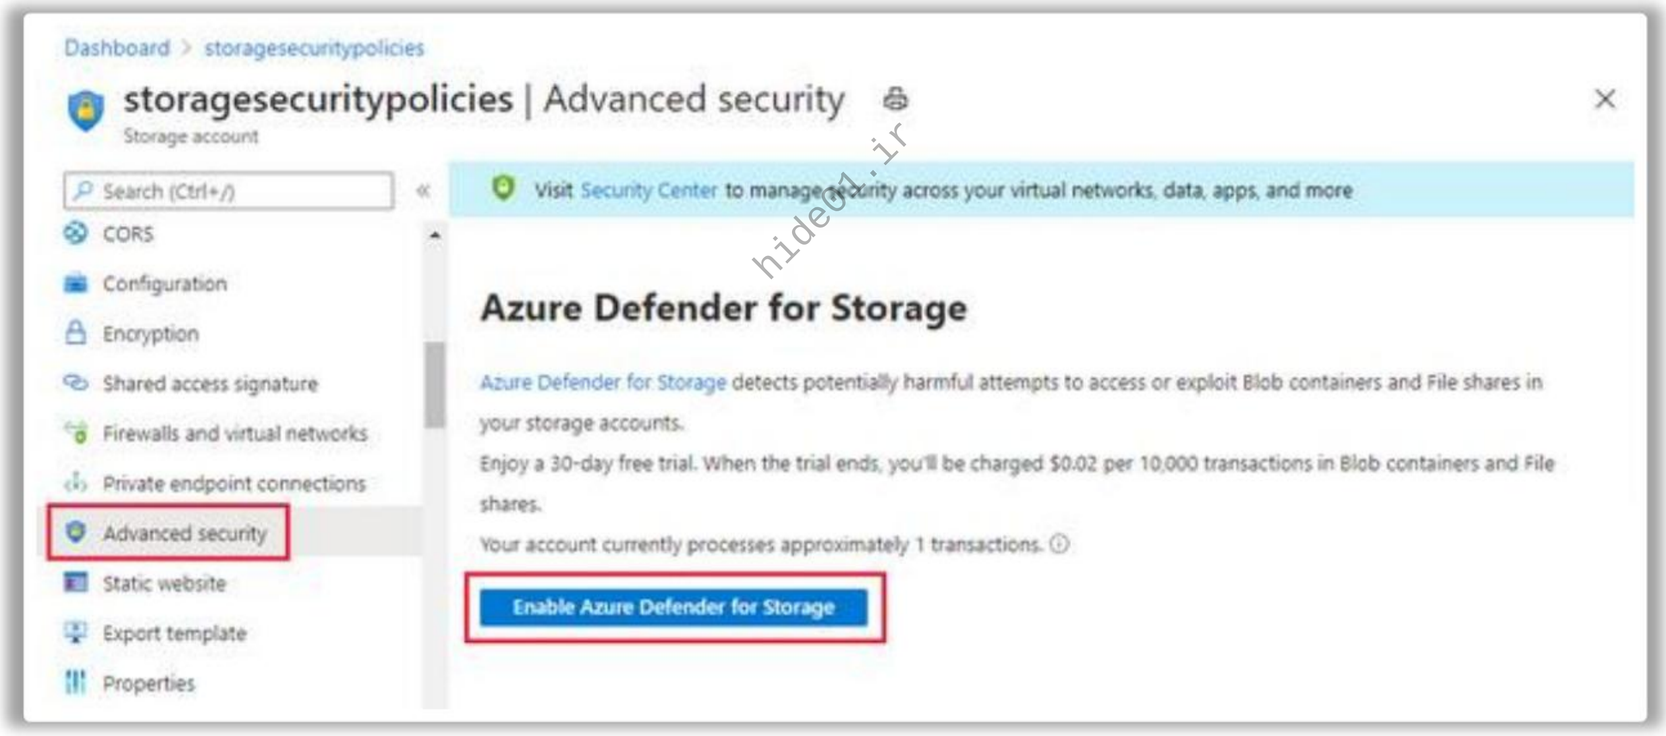
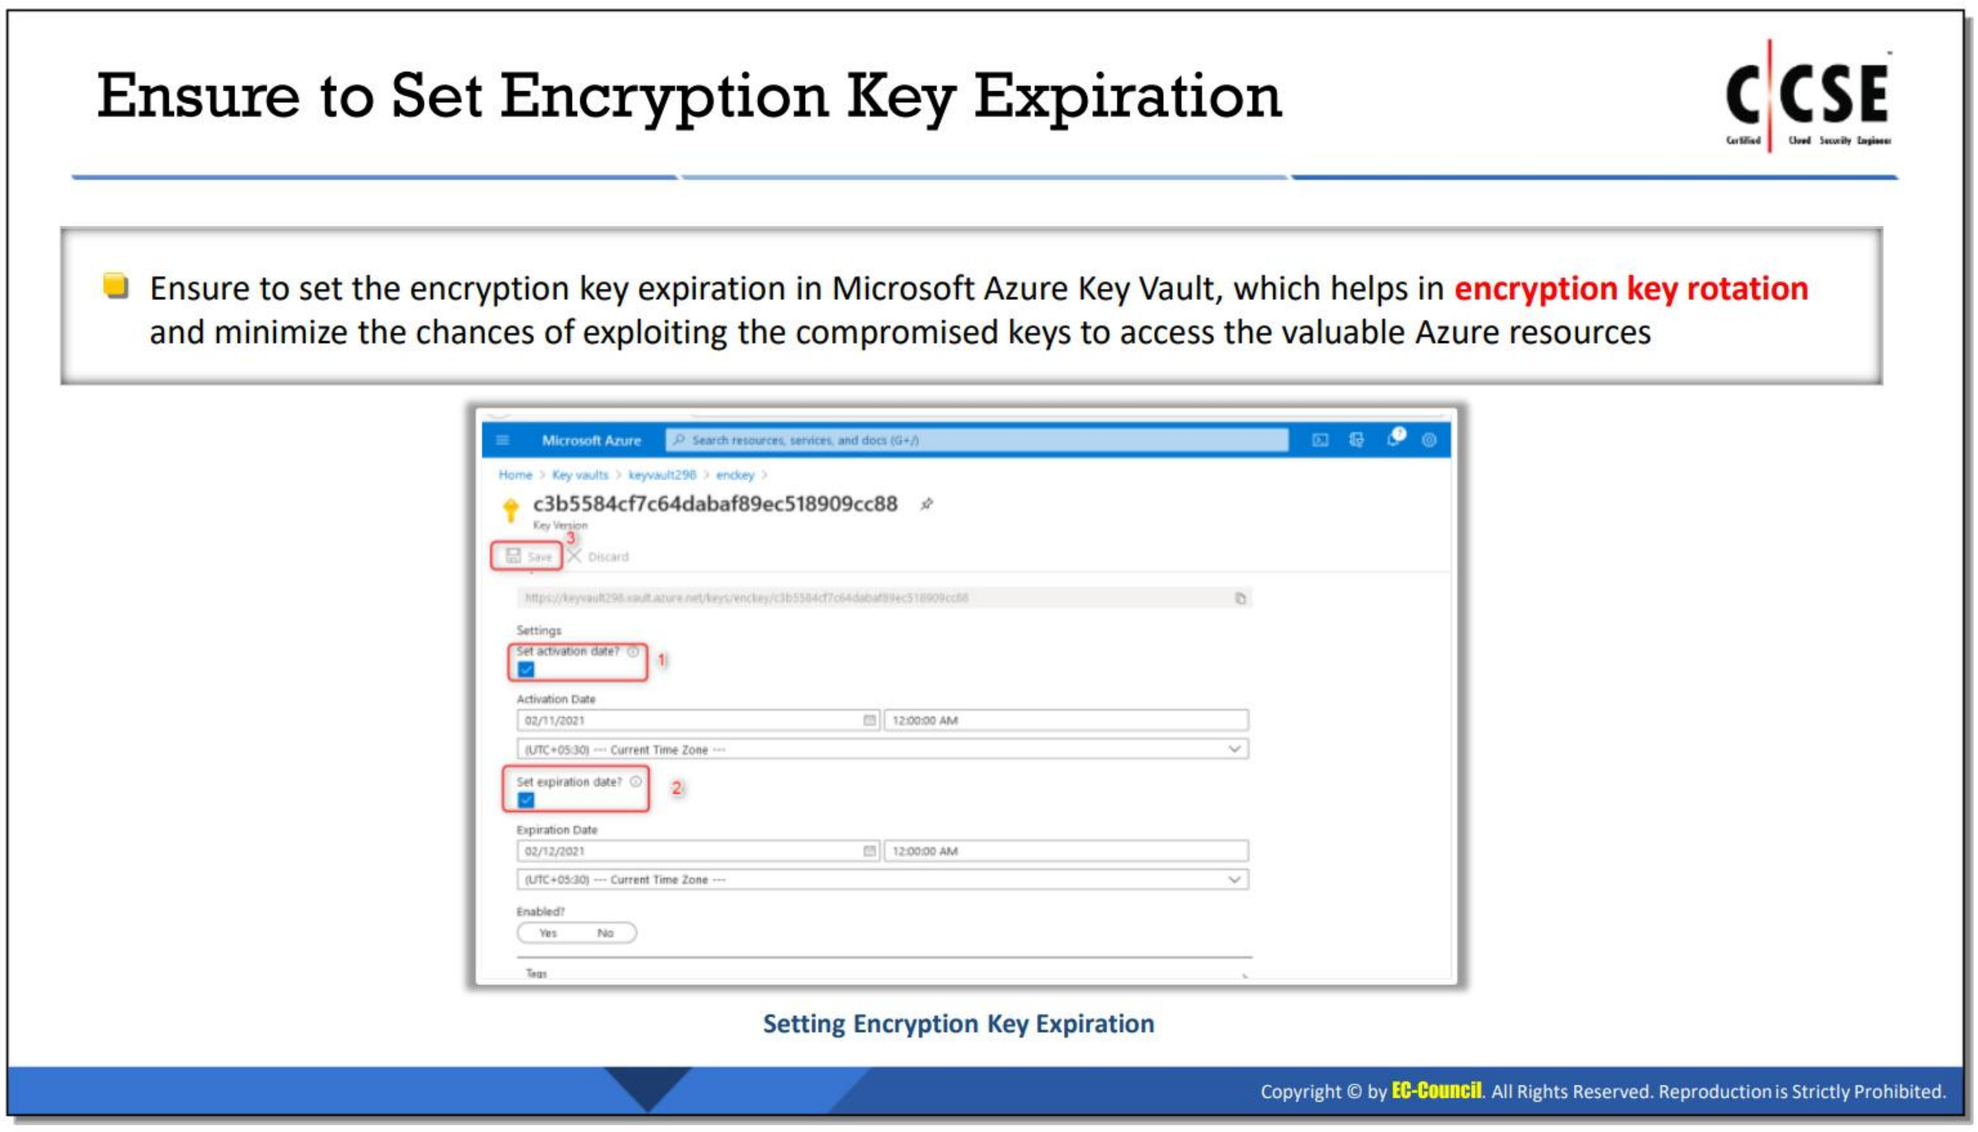


Figure 4.354: Azure Defender is Enabled

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



Ensure to Set Encryption Key Expiration

Microsoft Azure Key Vault is used for storing and using the encryption keys. By default, the stored encryption keys in the vault would not expire. Hence, ensure to set the encryption key expiration in Microsoft Azure Key Vault, which helps in encryption key rotation. After the expiry of the key, it should not be used for cryptographic operations. This will significantly minimize the chances of exploiting the compromised keys to access the valuable Azure resources.

Source: https://azure.microsoft.com

To Set Encryption Key Expiration

- Sign-in to the Azure portal.
- Select Key vault under Azure services.

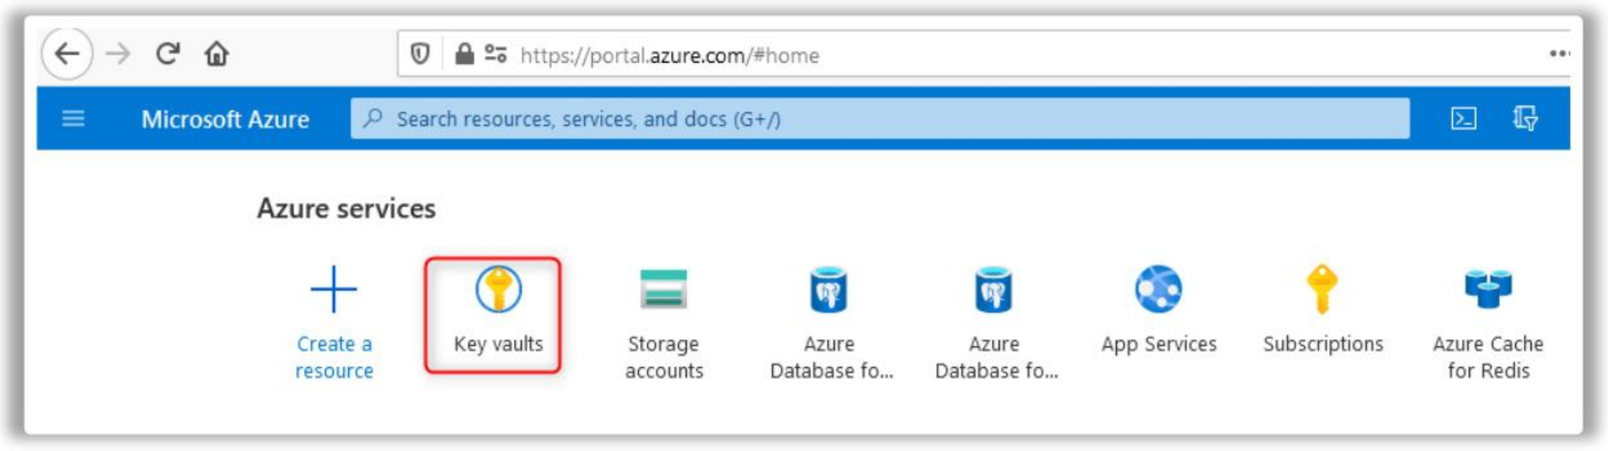


Figure 4.355: Select Key Vault

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

- A list of key vaults will be displayed, select the desired key vault.

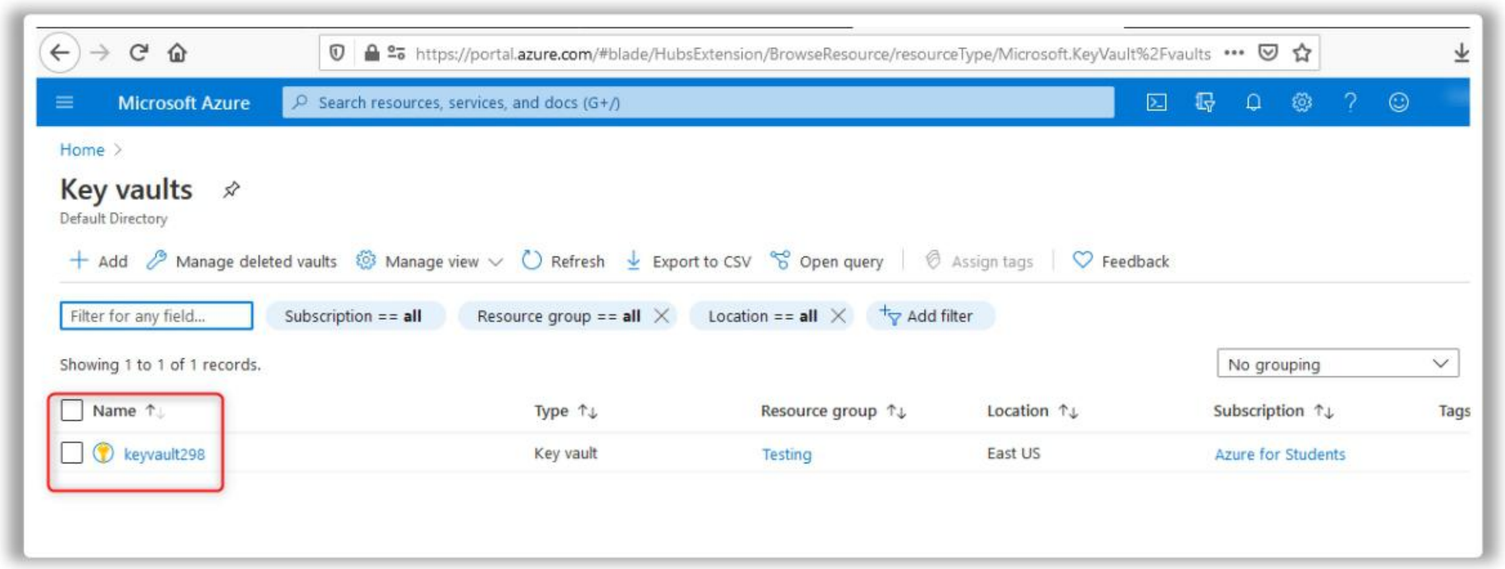


Figure 4.356: Choose Key Vaults

- Select **Keys** under **Settings.**
- Now in the right-side panel, select any encryption key.

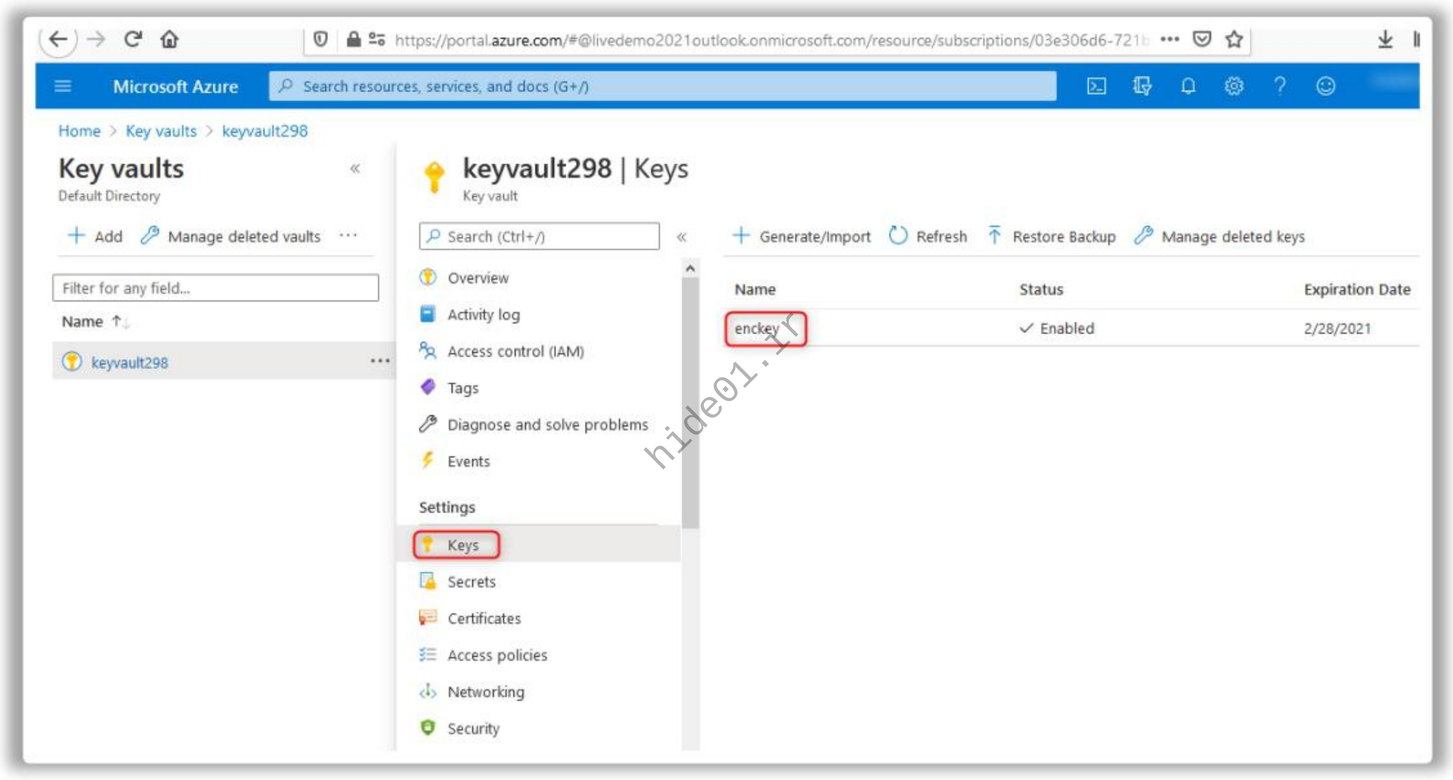


Figure 4.357: Select Encryption

- Select the current version key.

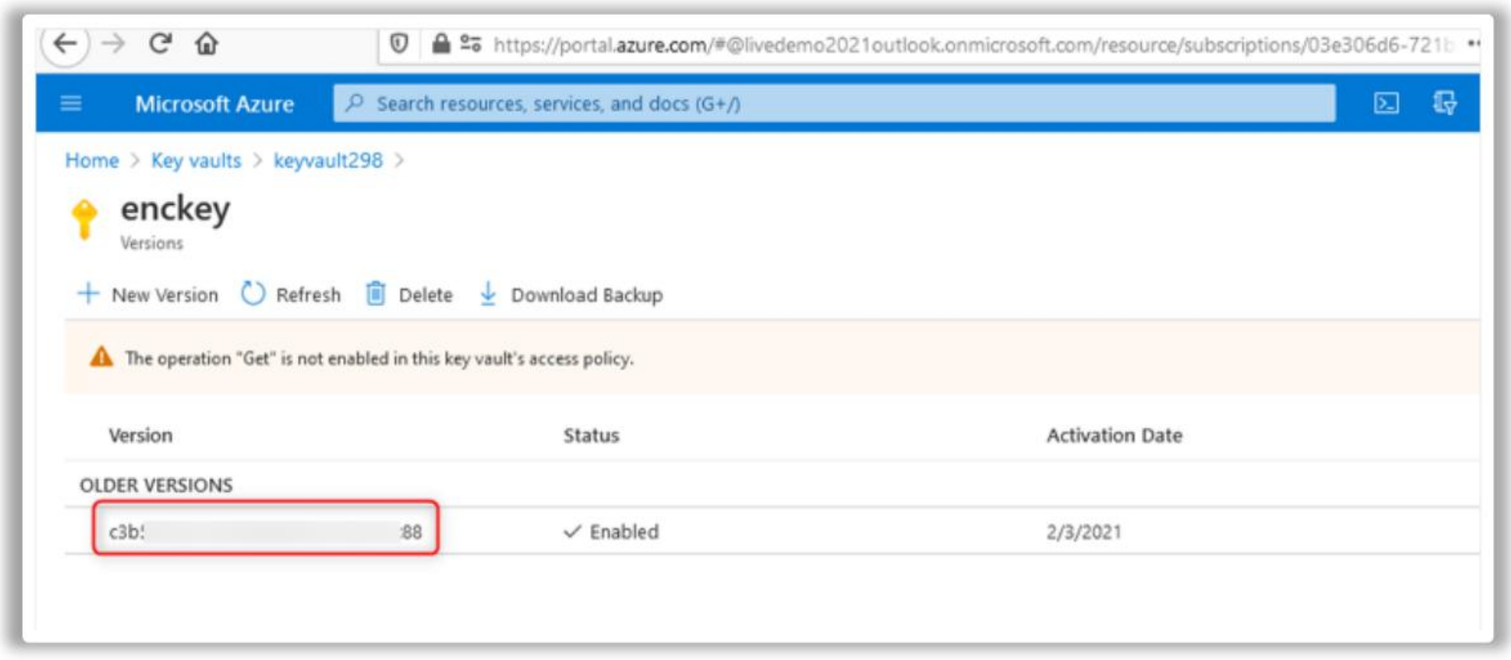


Figure 4.358: Select Key

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

- Modify, **Set expiration date?**
- Click on Save to apply changes successfully.

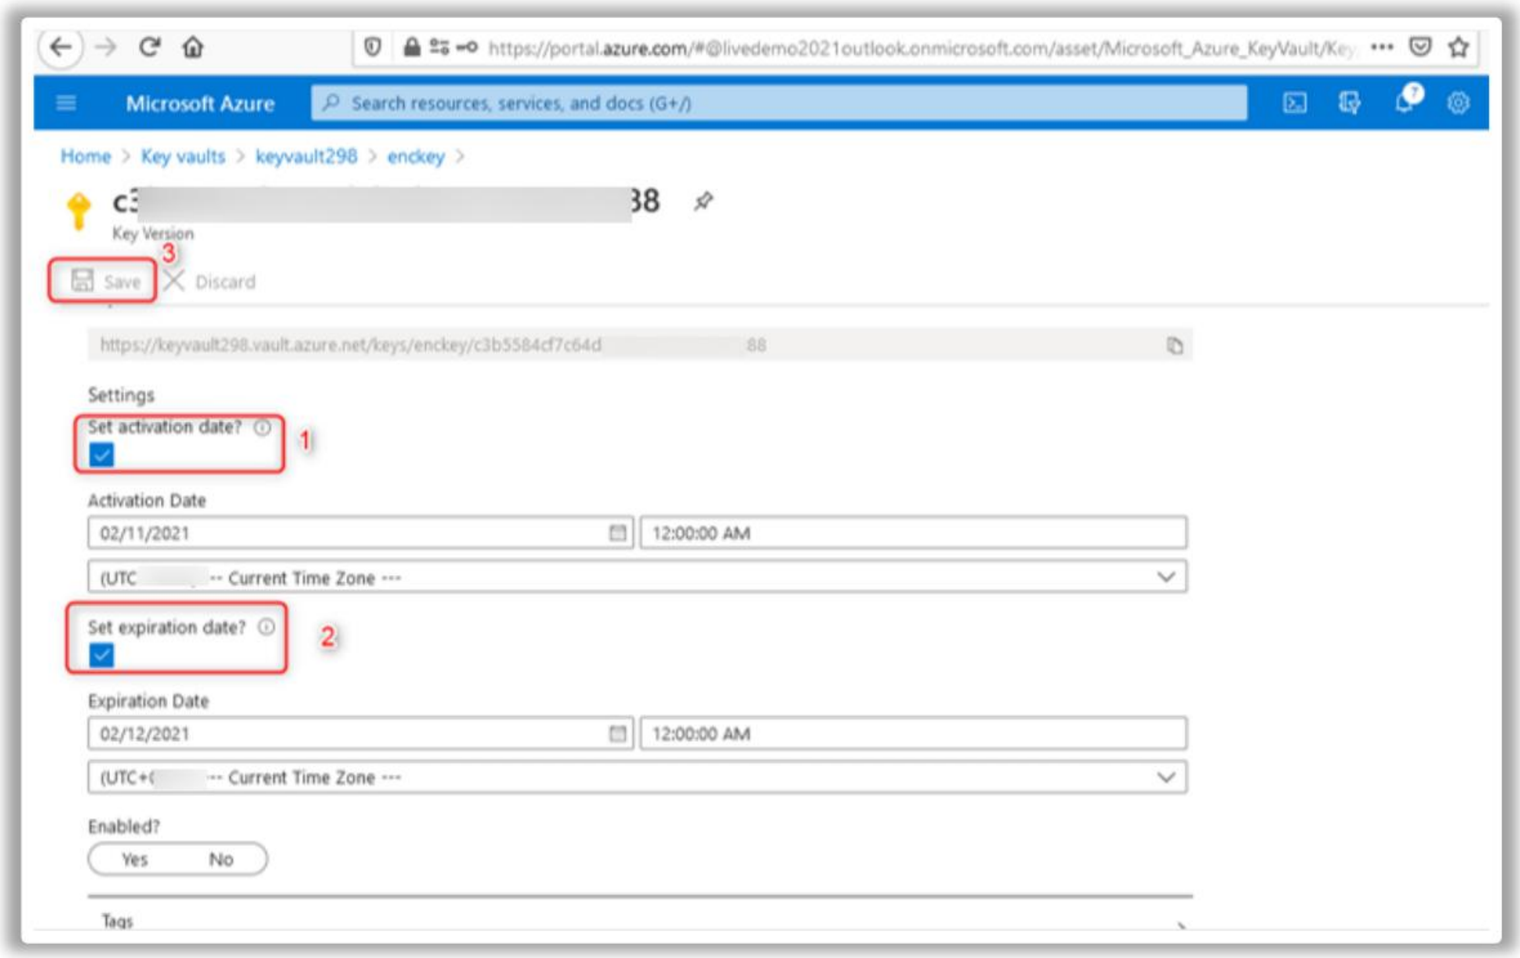
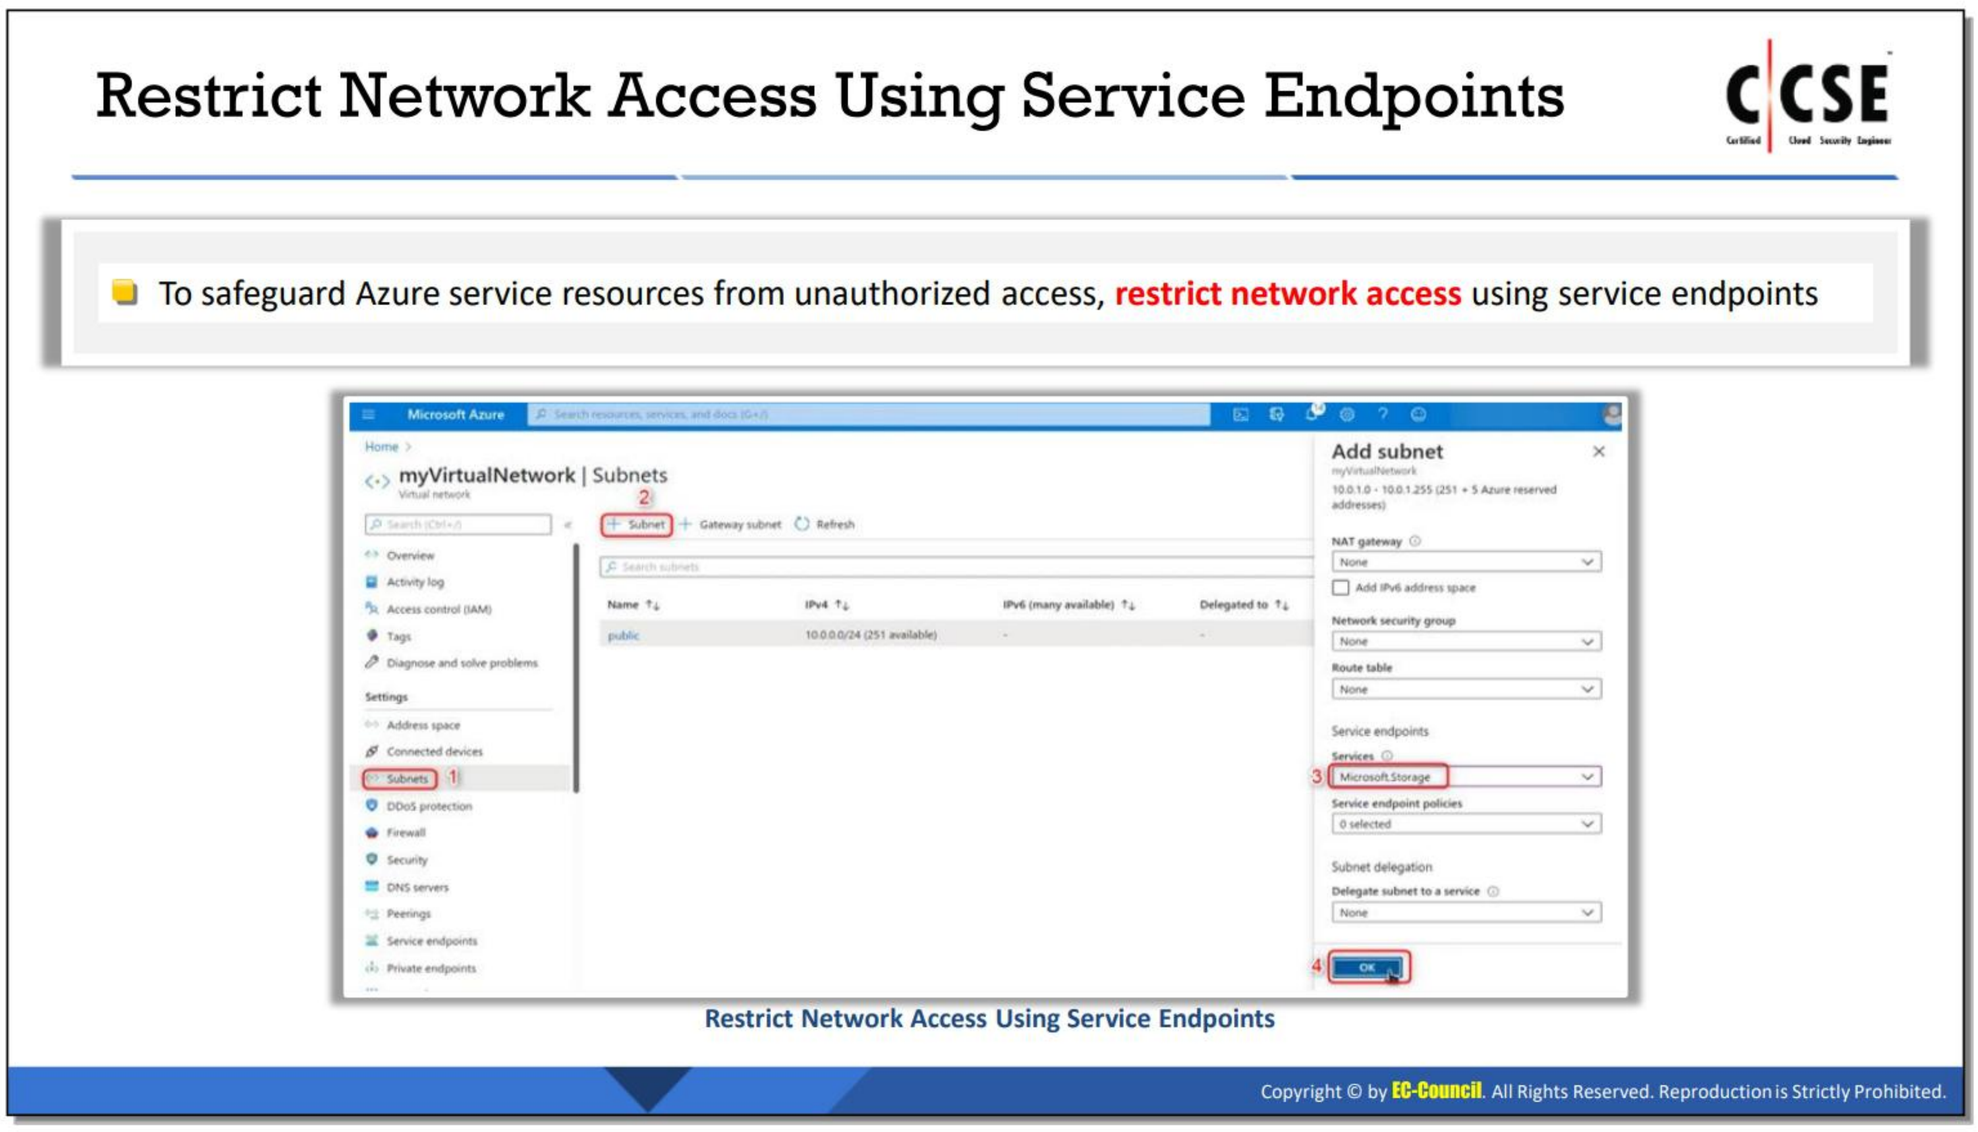


Figure 4.359: Apply Changes

hide01.ir

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



Restrict Network Access Using Service Endpoints

Restrict Network Access Using Service Endpoints

Virtual network service endpoints are used to limit network access to Azure service resources in a virtual network subnet. To safeguard Azure service resources from unauthorized access, restrict network access using service endpoints. Resources can be isolated from the internet. The connection from the virtual network to a list of any supported Azure services are provided by the service endpoint.

Source: https://azure.microsoft.com

Steps to restrict network access using service endpoints are as follows:

- Initially, a virtual network with one subnet should be created.
- A subnet should be added and a service endpoint should be enabled.
- An Azure resource should be created to allow its network access from only a subnet.
- Each VM should be deployed with each subnet.
- Verify access to a resource from a subnet.
- Verify that access to a resource is denied from the internet and a subnet.

Steps to create a virtual network are as follows:

- Sign-in to Azure portal.
- First, create a virtual network.

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

- From the home page, search for **Virtual Networks** as shown below and select it.

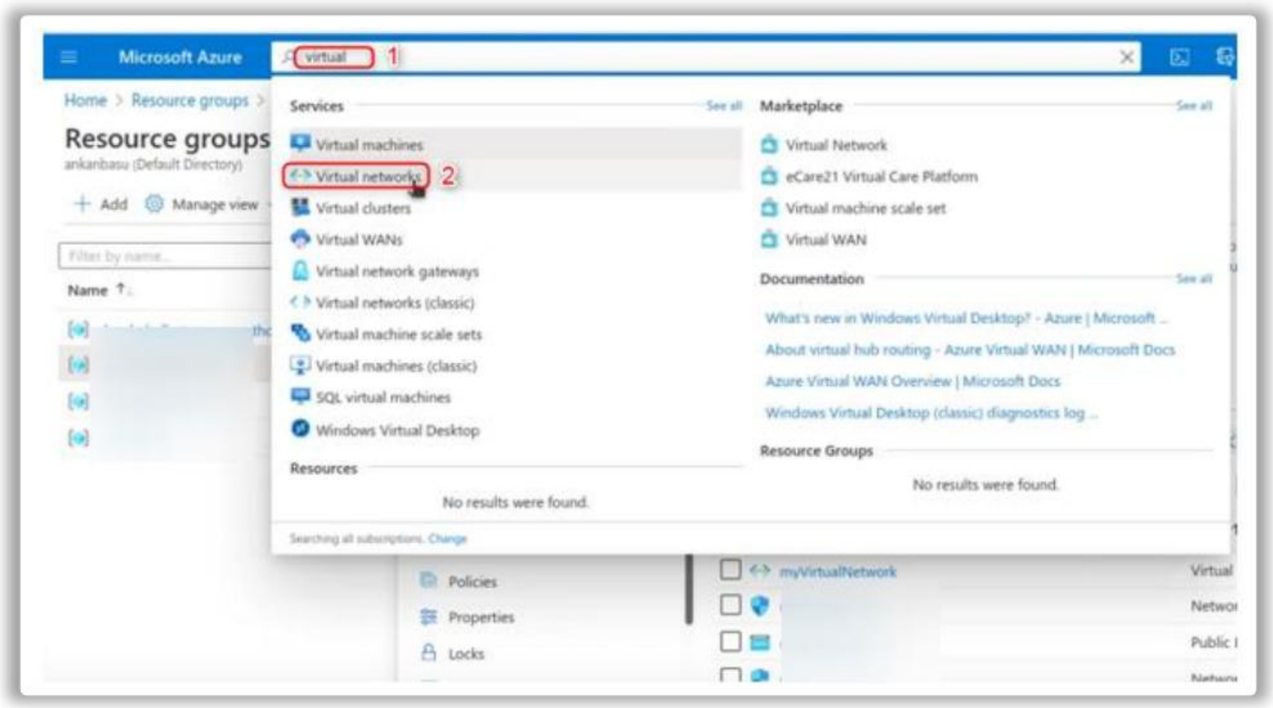


Figure 4.360: Select Virtual Networks

- Click on **Add** to create a virtual network.

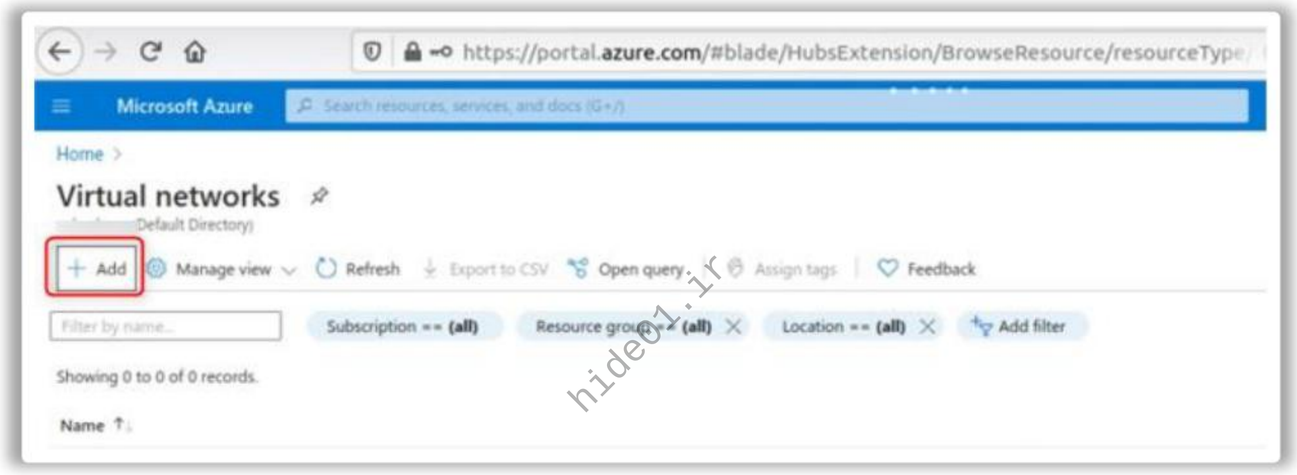


Figure 4.361: Click Add

- Enter required details in the **Basics** tab.

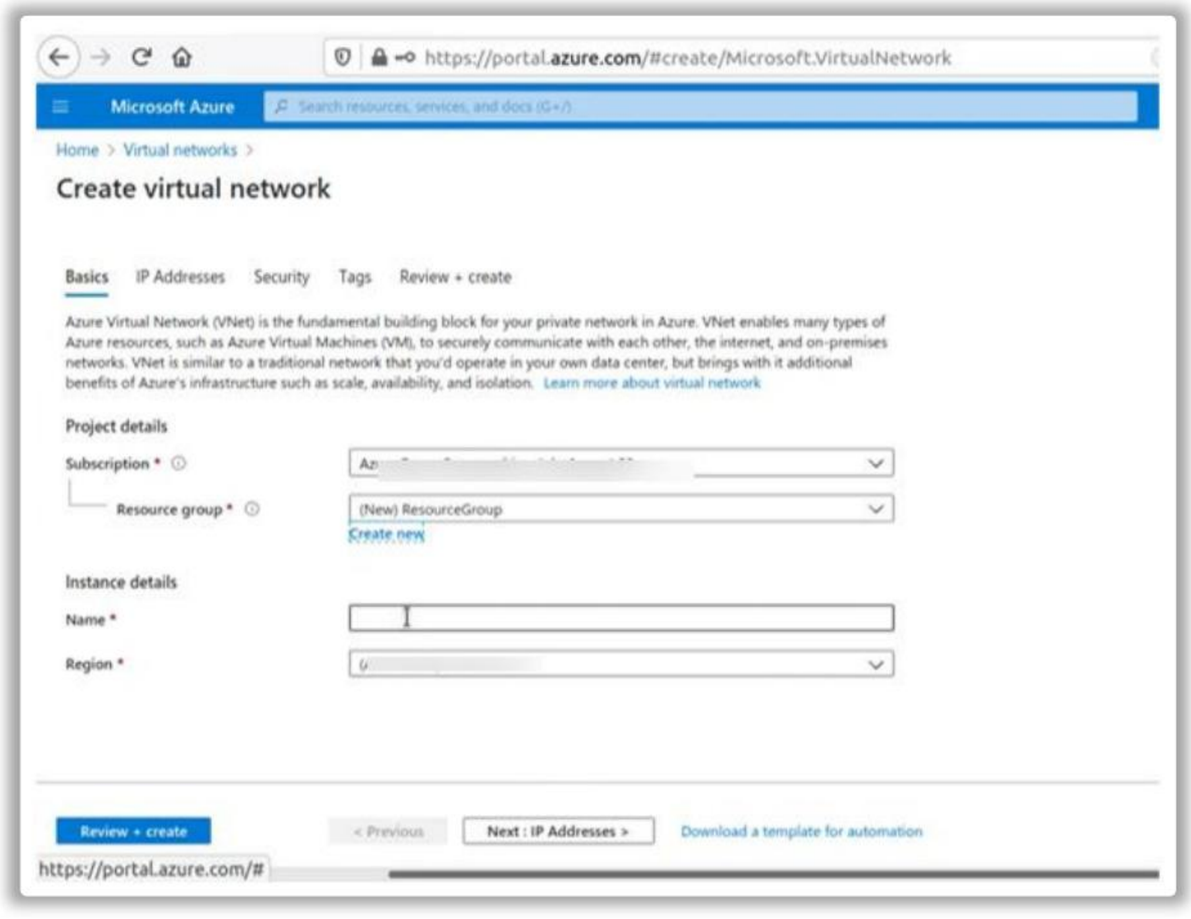


Figure 4.362: Click Basics

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

- Go to the **IP Addresses** tab
- Select **default** under **Subnet name**.
- On the right-side panel, select **Subnet name** and **Subnet address range.**
- Click on **Save.**

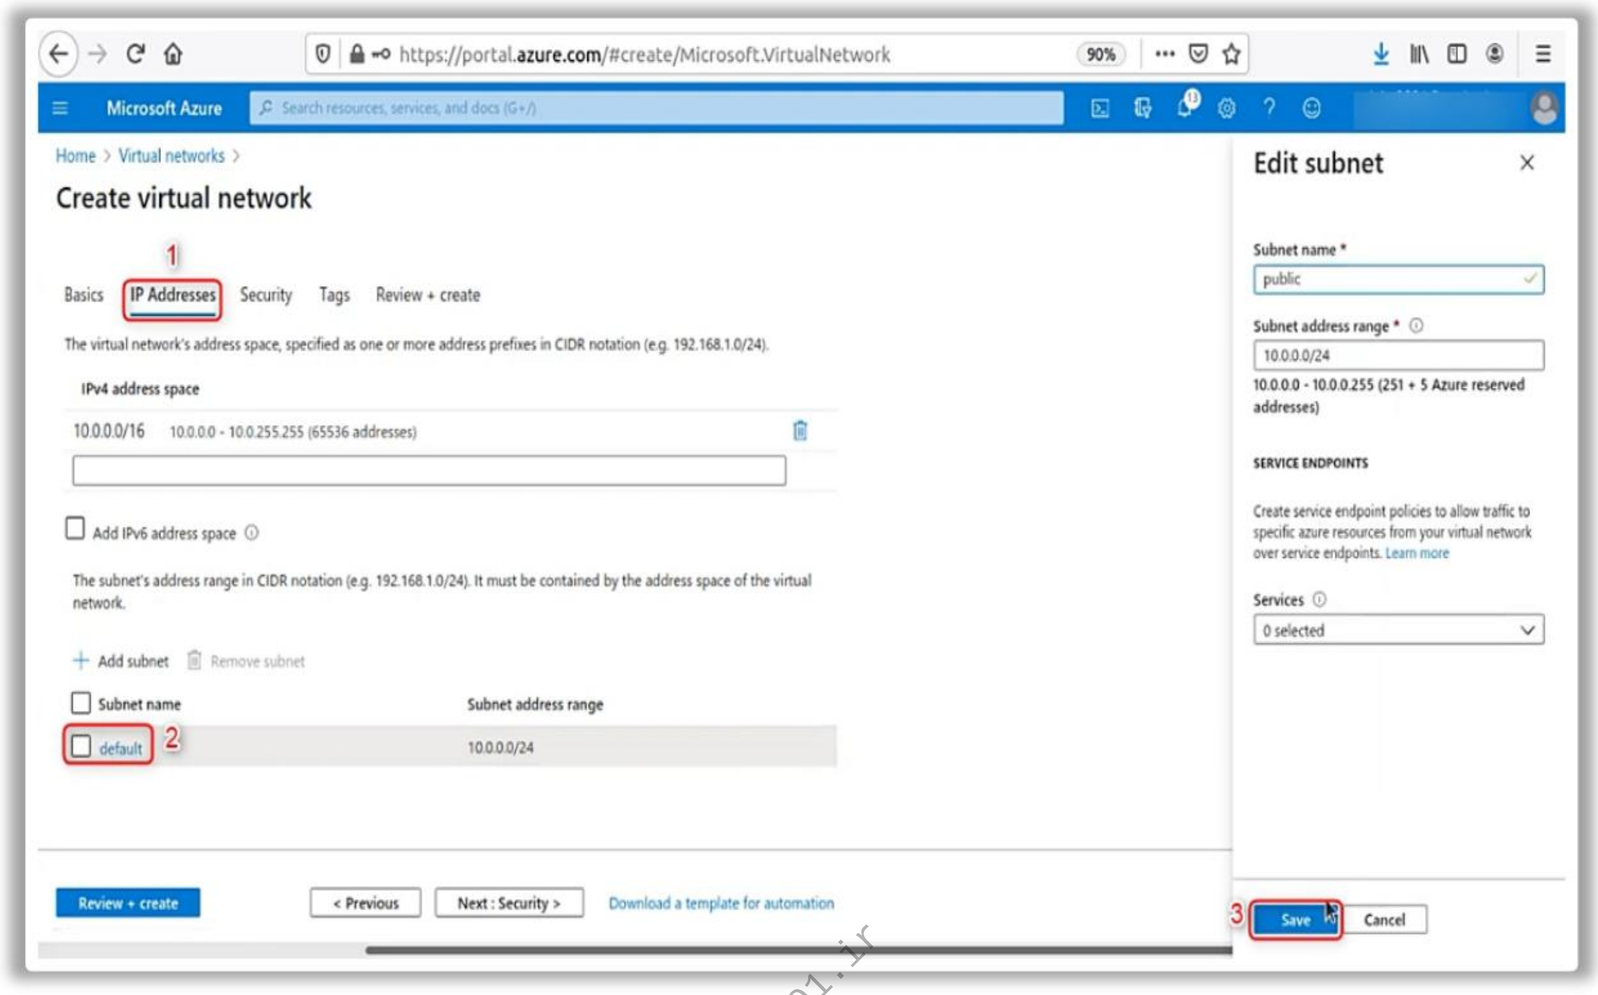


Figure 4.363: Click Save

- Go to the **Security** tab, click on **Review+create**.

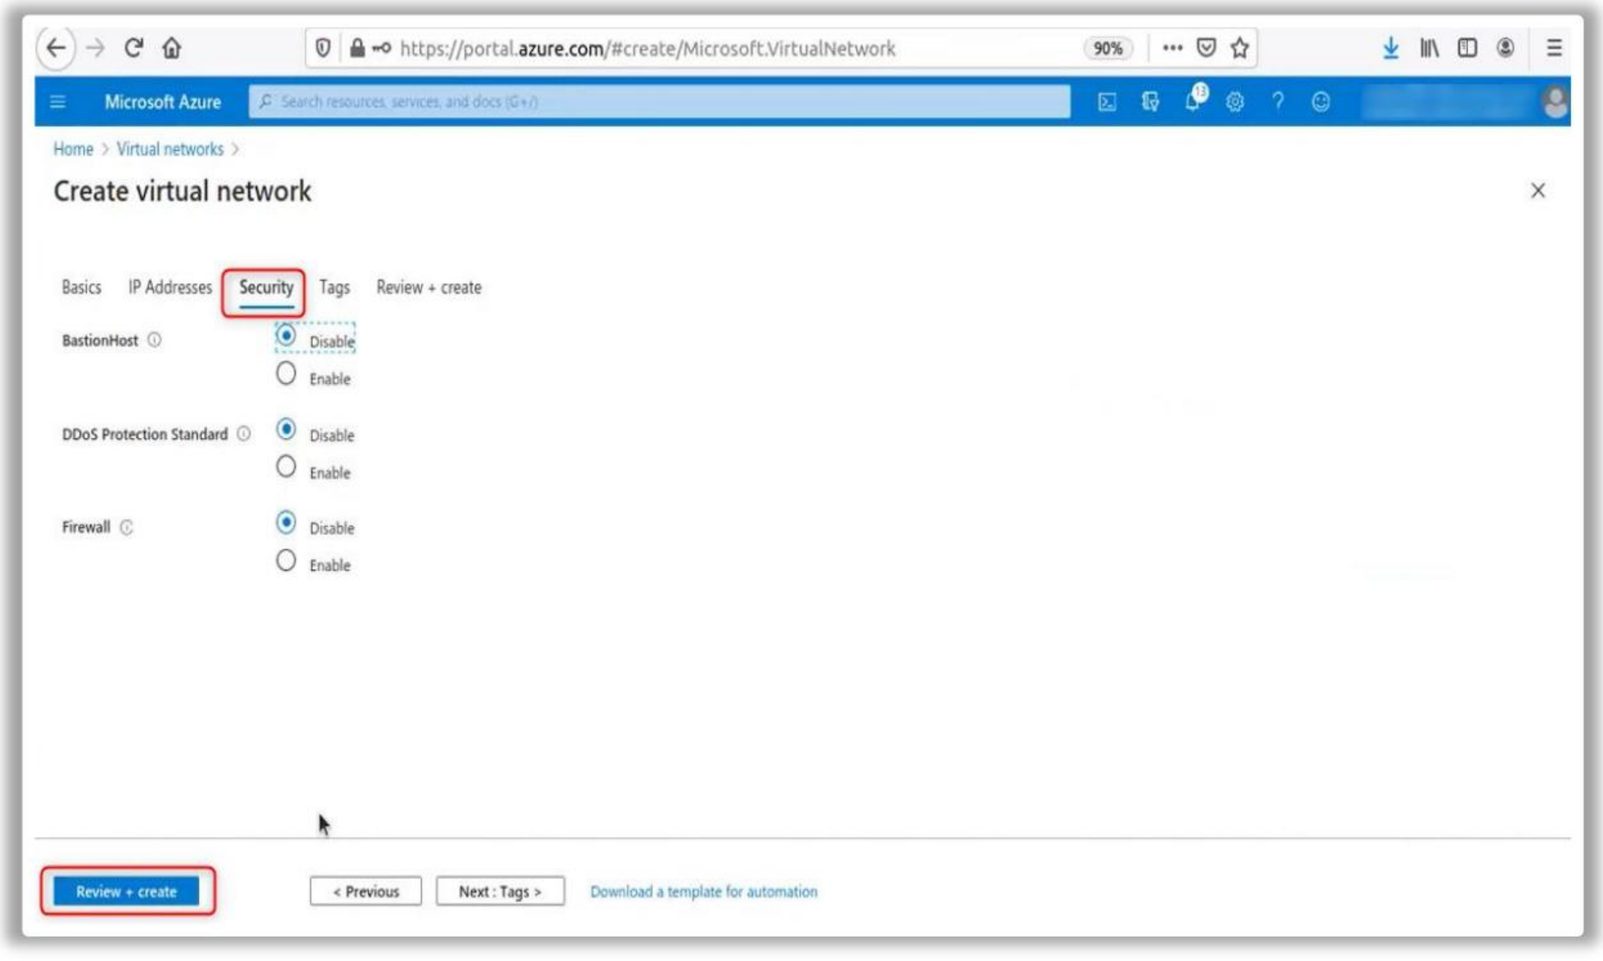


Figure 4.364: Click Review+ create

- The virtual network is successfully created.

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

- Click on "**Go to resource**." Which is displayed under **Deployment succeeded.**

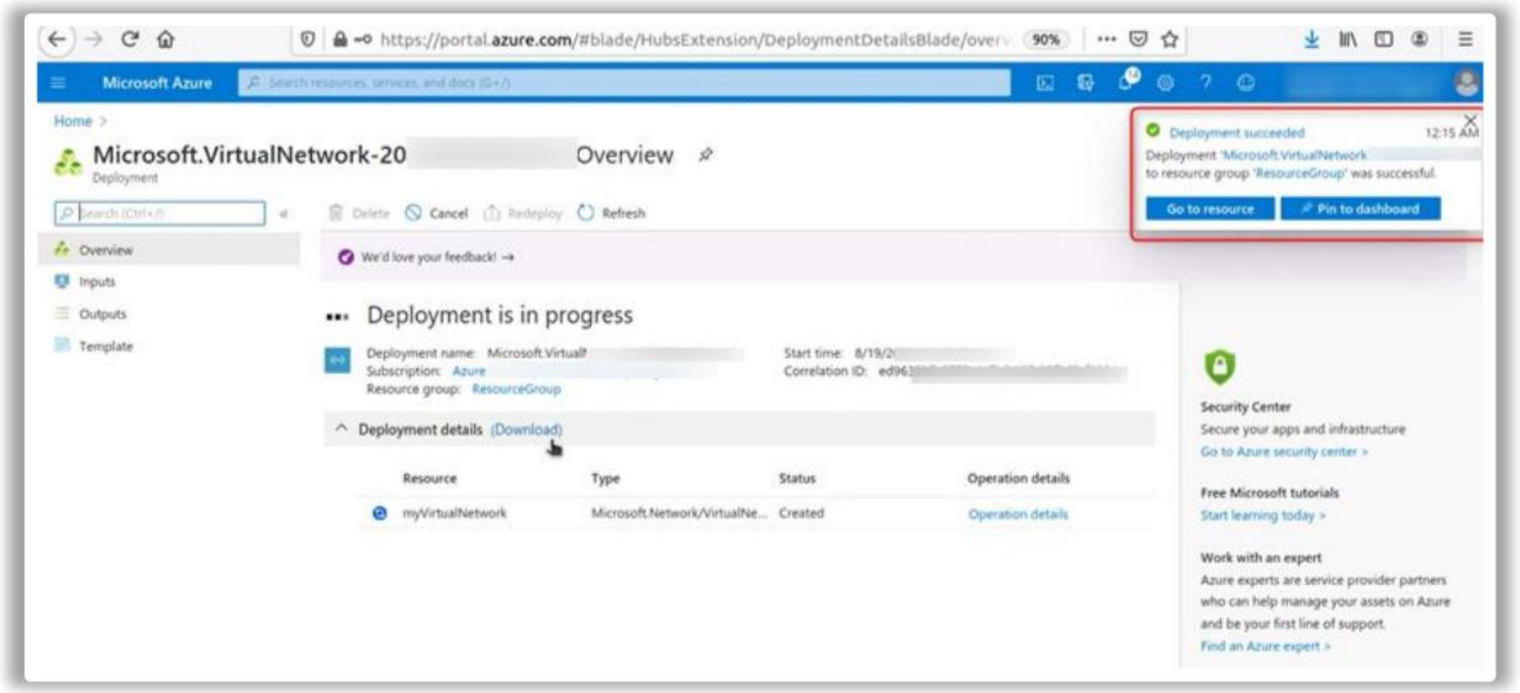


Figure 4.365: Click Go to resource

- Locate and click on **Subnets** in the left navigation panel.
- Click on **+ Subnet**.
- Enter the subnet name as **private.**
- Select **Microsoft storage** from the dropdown under **Services.**
- Click on **OK.**

hide01.ir

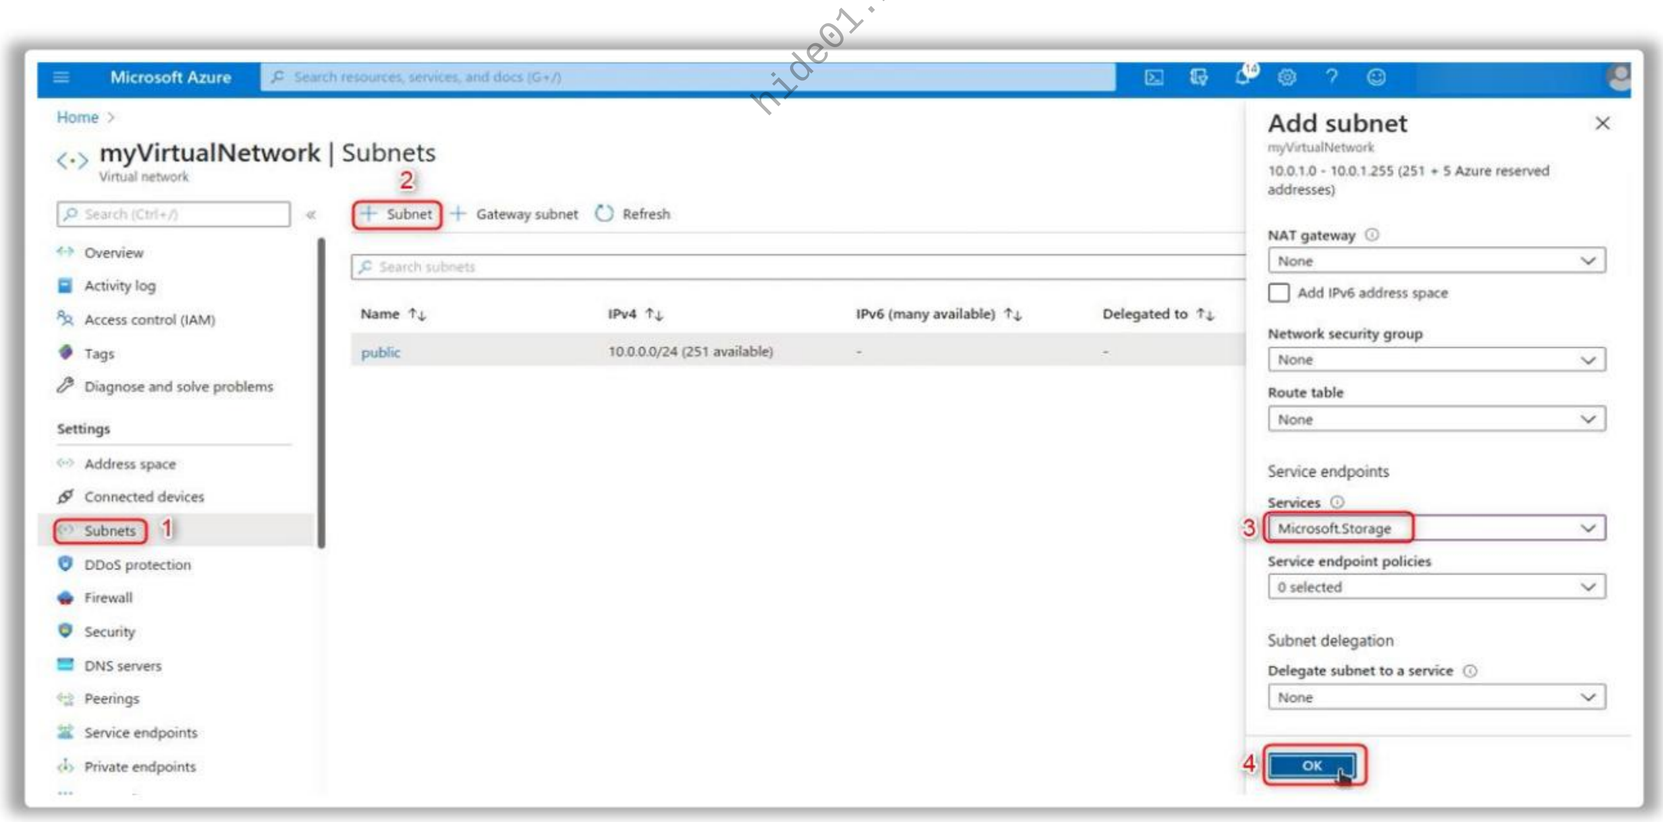


Figure 4.366: Click OK.

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

- The subnet is successfully created.

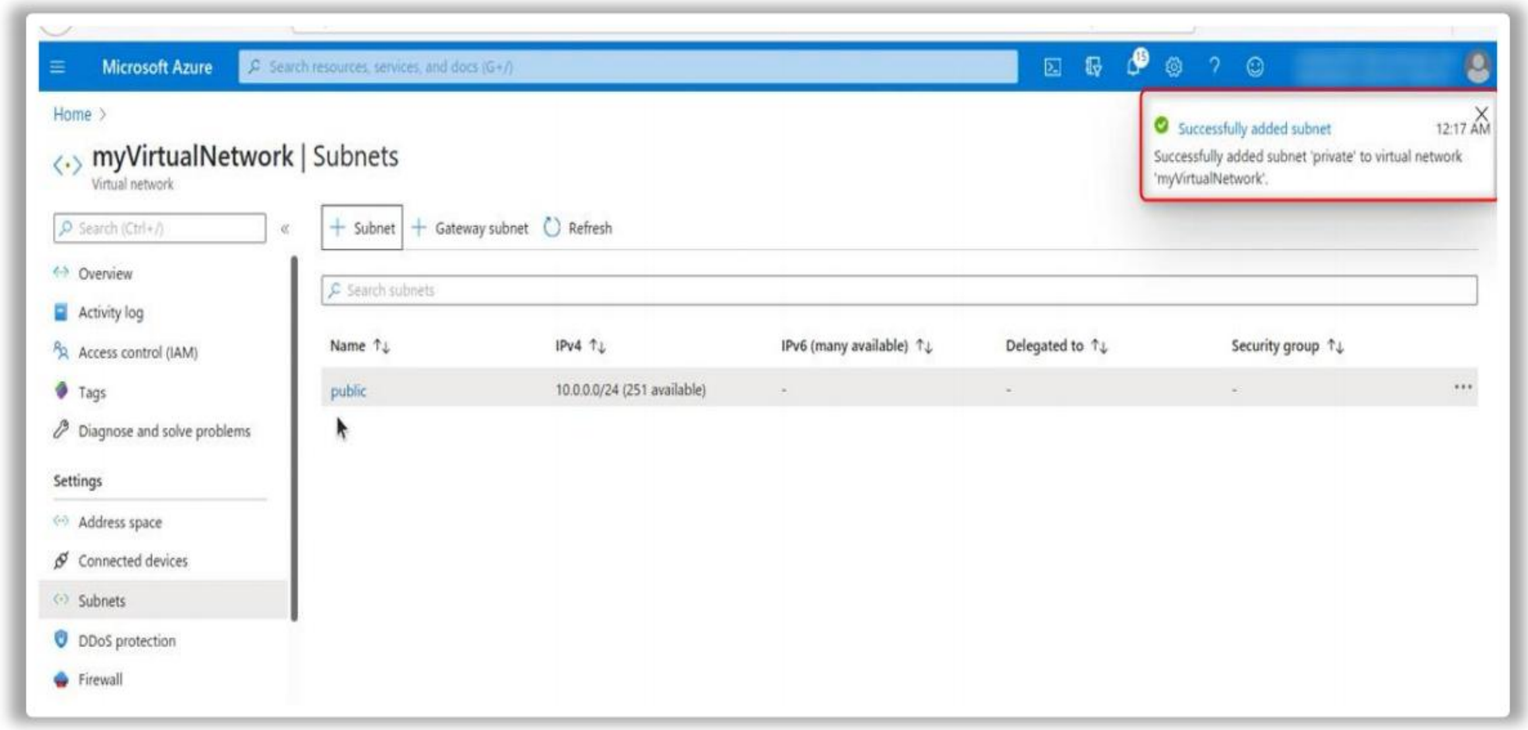


Figure 4.367: Subnet is Created

- Click on **Refresh**, you can find the created subnet below.

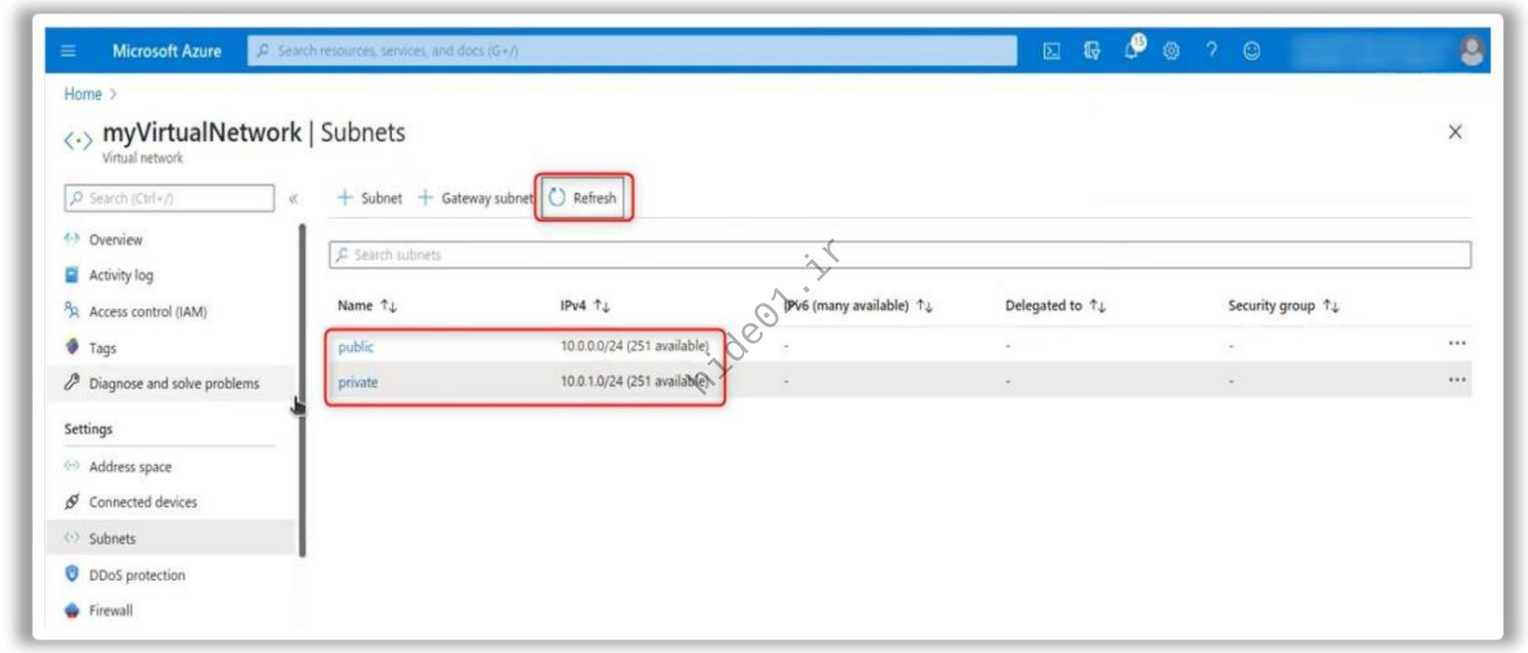


Figure 4.368: Click Refresh

- Search for the **Network security groups** option and click on it, as shown below.

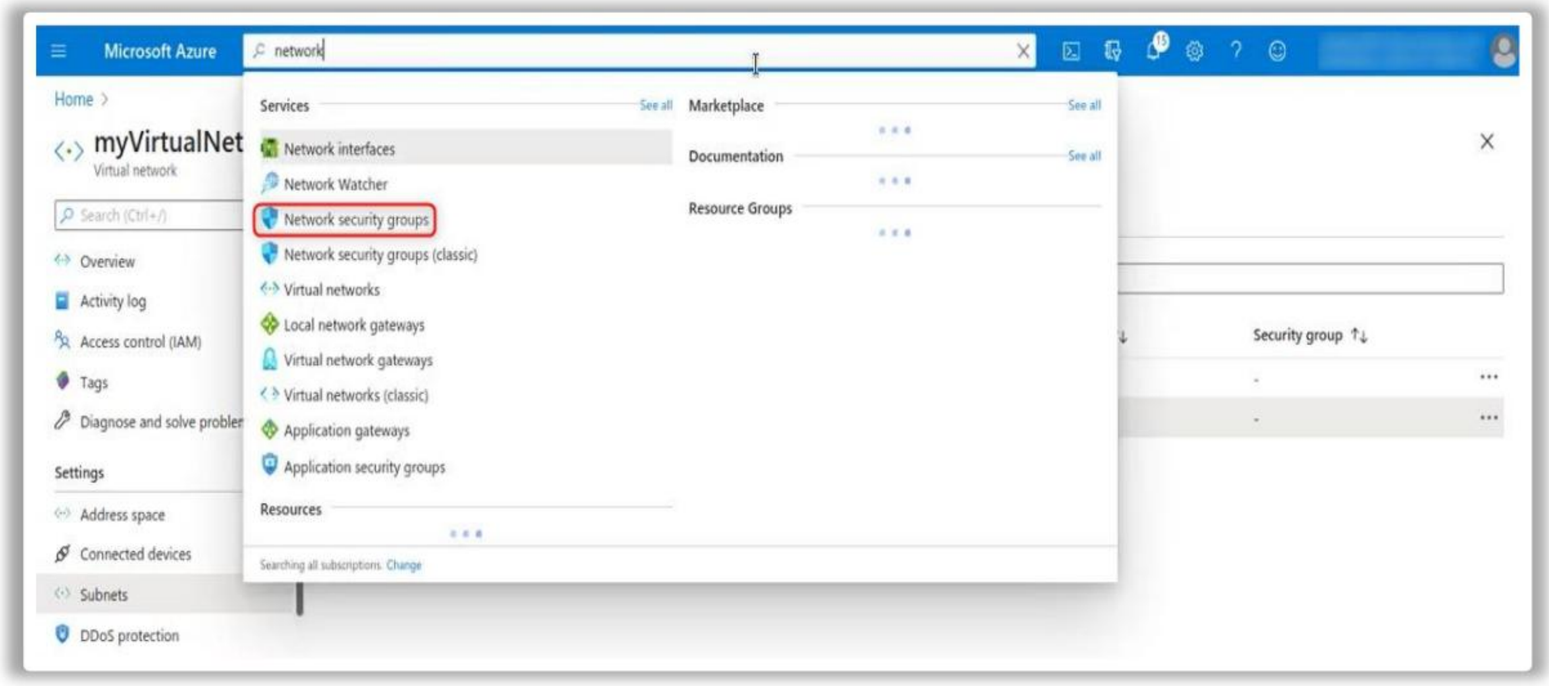


Figure 4.369: Search Network Security Groups

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

- Click on **+ Add**.

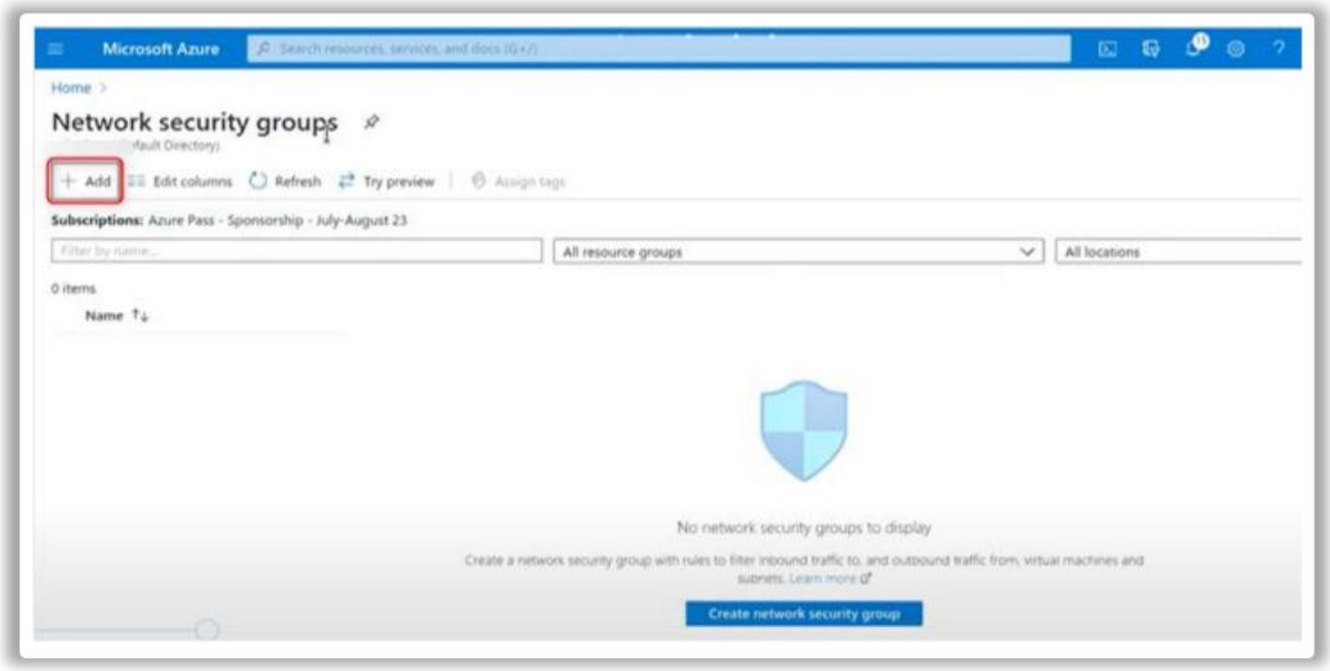


Figure 4.370: Click Add

- Select the required details like **Subscription**, **Resource group**, **Name**, and **Region**.
- Click on **Review + create**.

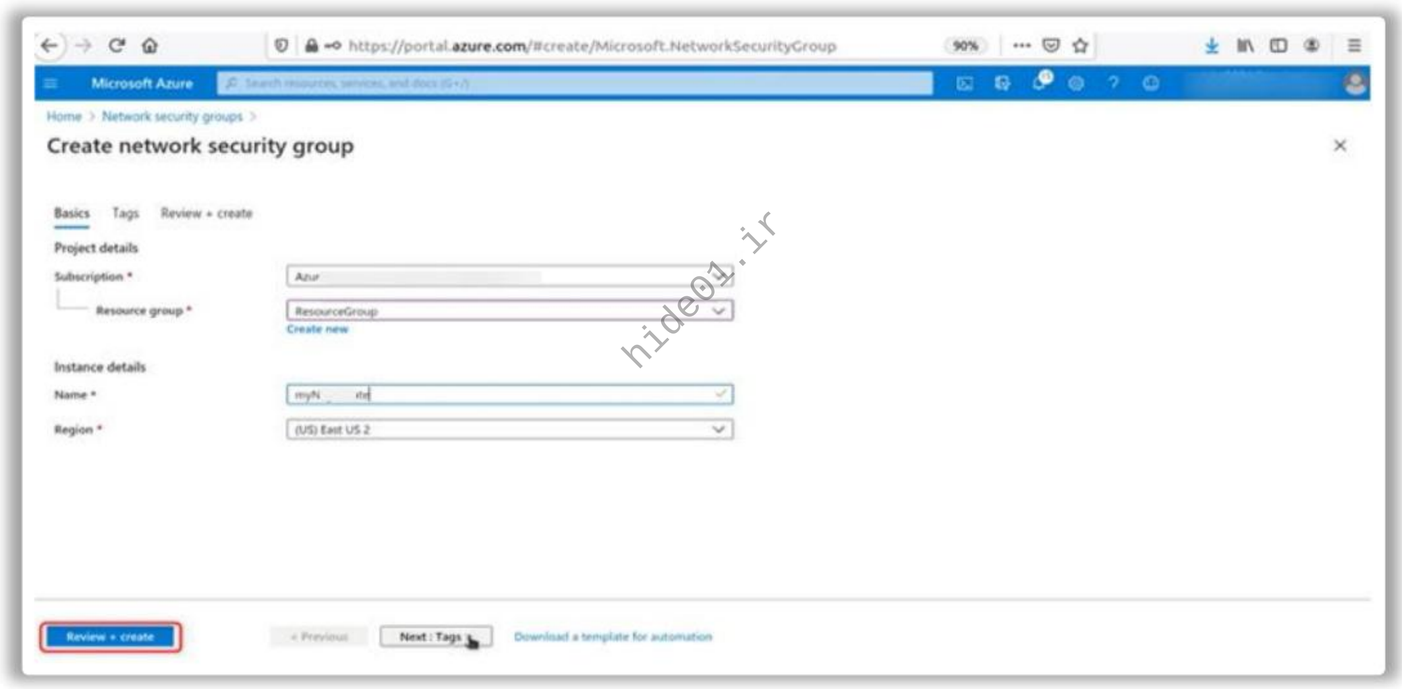


Figure 4.371: Click Review + Create

- The network security group is successfully created.
- Click on **Go to resource**.

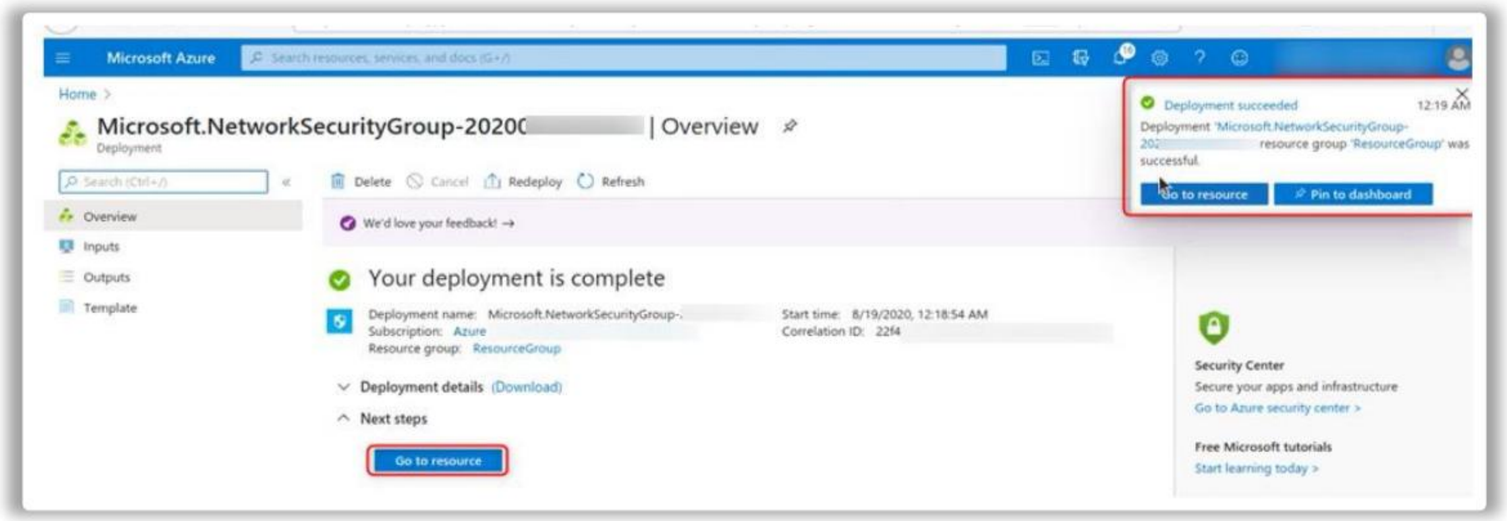


Figure 4.372: Click Go to Resource

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

- Click on **Outbound security rules** from the left navigation panel.
- Click on **Add.**
- Select **Source** as **Virtual Network**, **Destination** as **Service Tag**, **Protocol** to **Any**, **Priority** as **100**, give this rule a name and click on **ADD**.

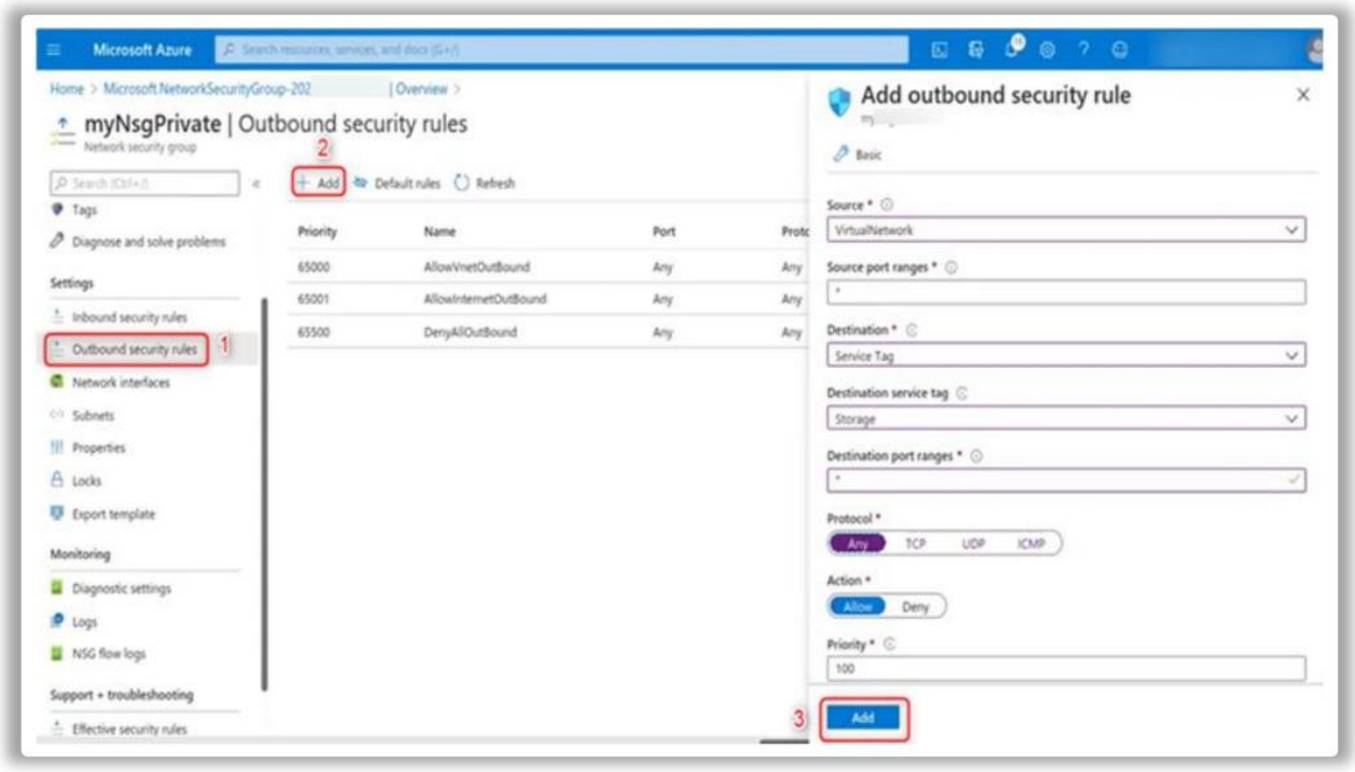
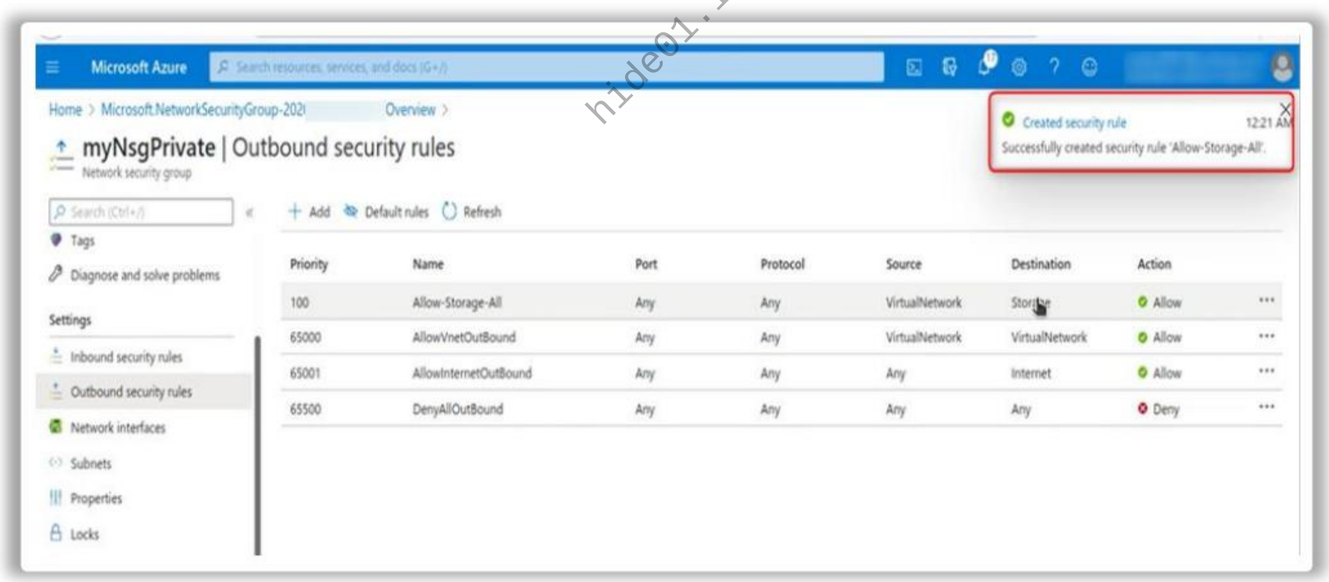


Figure 4.373: Click Add

- Security rule is successfully created.



hide01.ir

Figure 4.374: Security Rule is Created

- Again click on **Add,** select **Source** as **Virtual Network**, **Destination** as **Service Tag**, **Protocol** as **Any**, **Action** to **Deny**, **Priority** as 110, **Name** it as Deny-Internet-All.

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

- Click on **Add**.

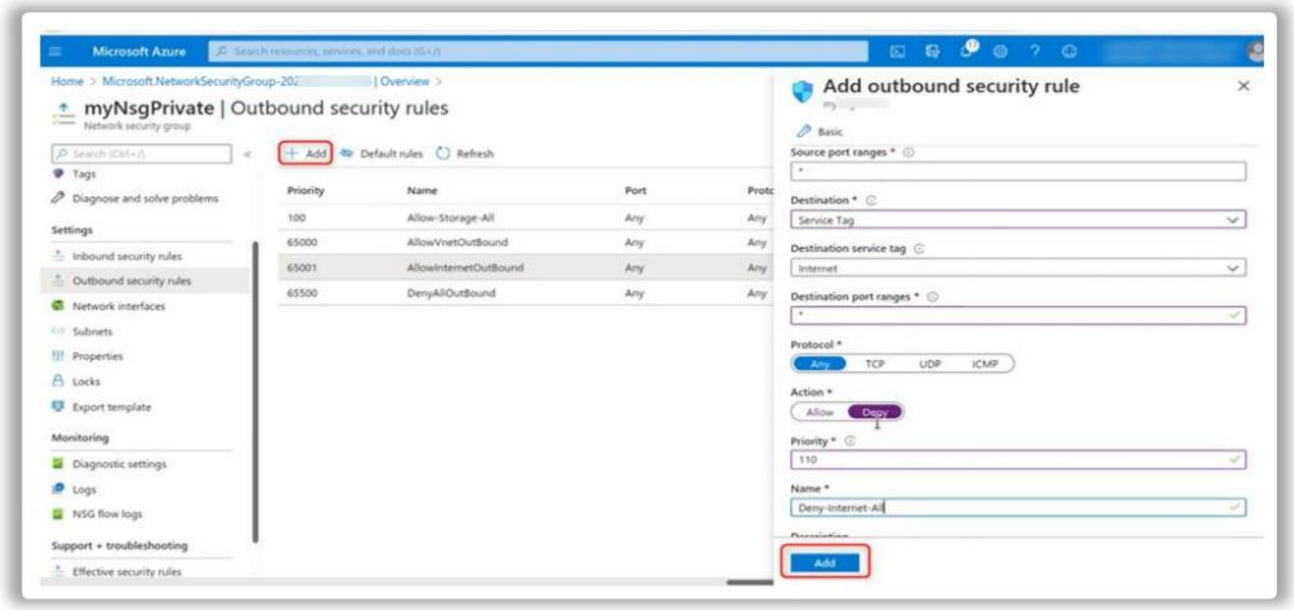


Figure 4.375: Click Add

- Security rule is successfully created.

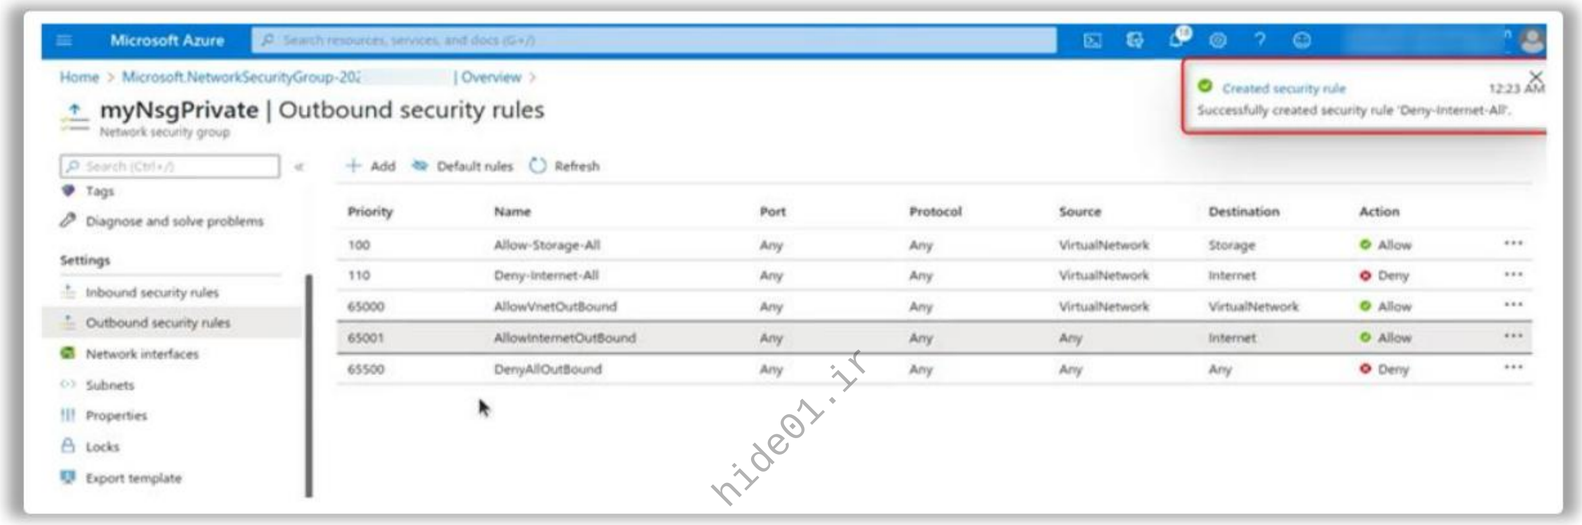


Figure 4.376: Security Rule is Created

- Under **Settings**, click on **Inbound security rules**.
- Click on **Add**.
- Under the **Add inbound security rule**, select **Source** as **Any**, **Destination** as **VirtualNetwork**, **Destination port ranges** as **3389**, **Protocol** as **TCP**, **Action** to **Allow**, **Priority** as **120**, **Name** as **Allow-RDP-All** and click on **Add**.

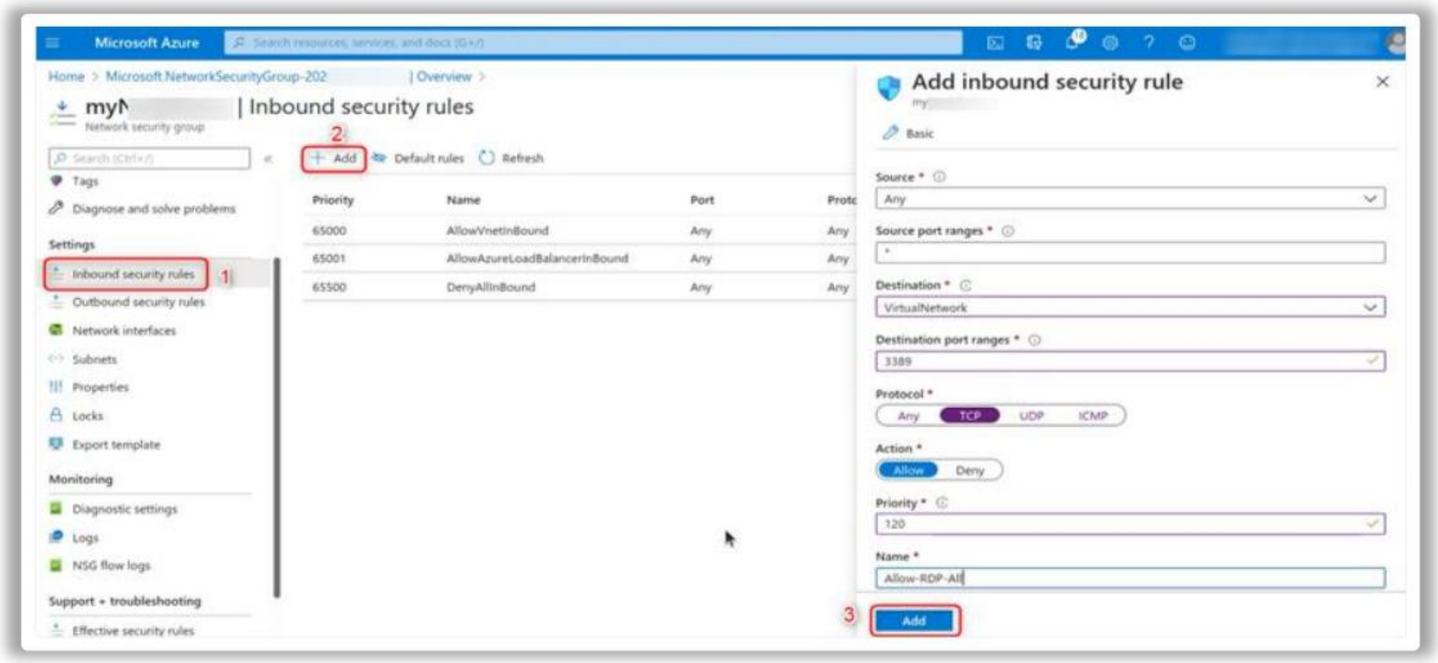


Figure 4.377: Clock Add

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

- Security rule has been created successfully.

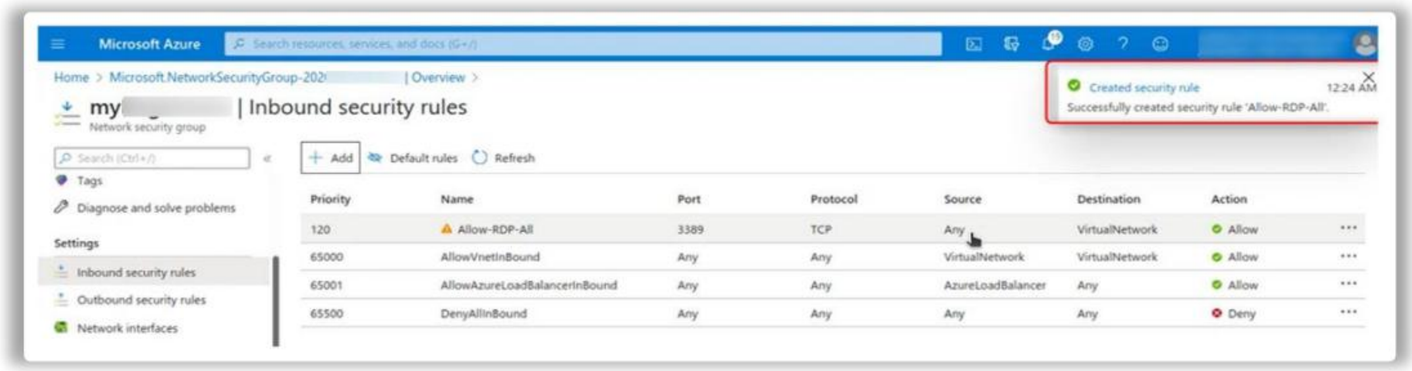


Figure 4.378: Security Rule is Created

- Under **Settings**, click on **Subnets**.
- Click on **Associate**.
- Select **Virtual network** from the dropdown, and **Subnet** as **private**.

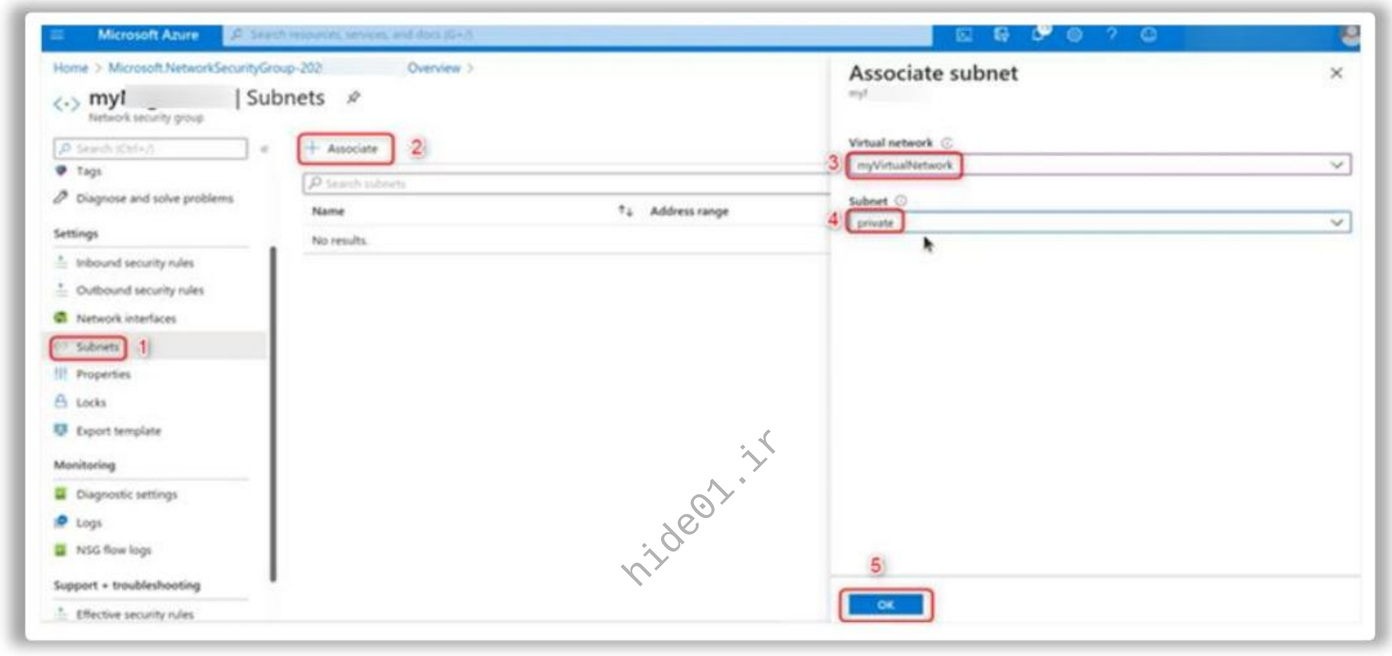


Figure 4.379: Select Private

- Successfully saved network security group for subnet ‘private.’

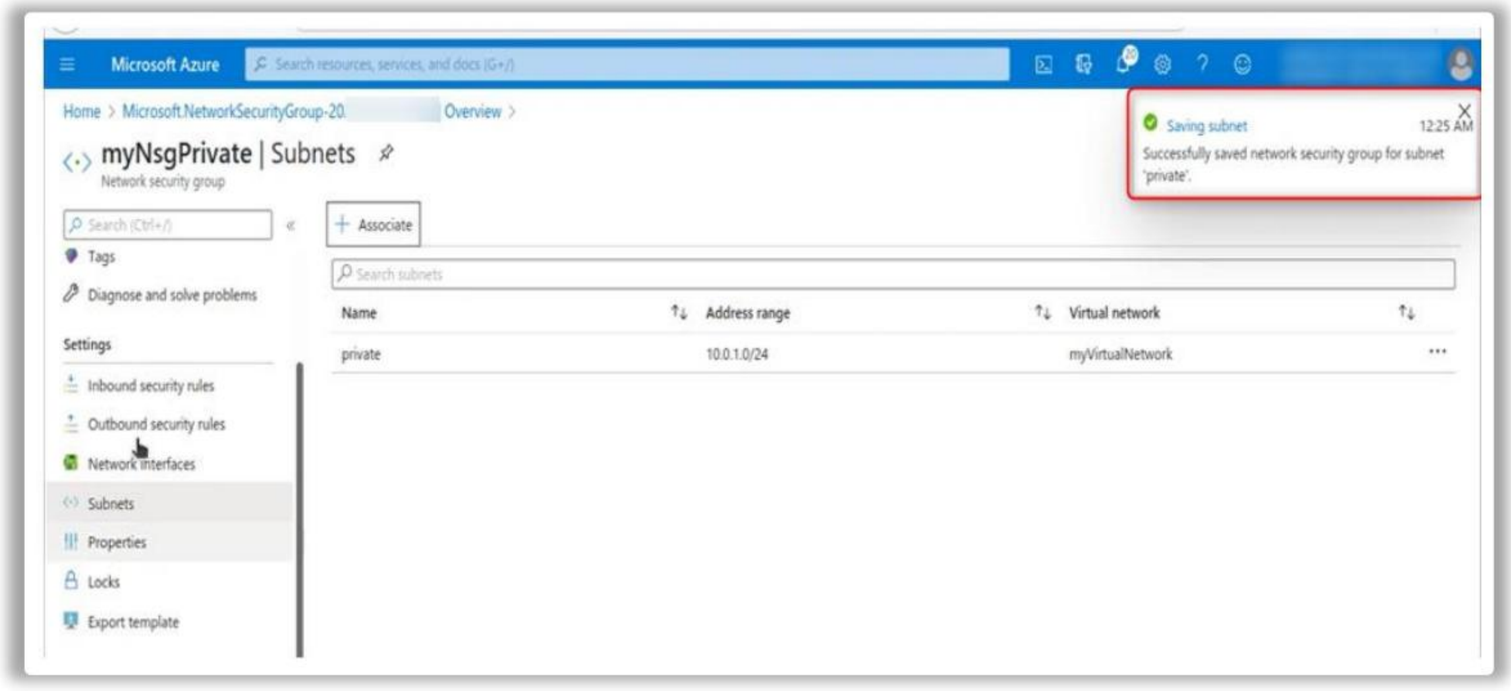


Figure 4.380: Security Group is Saved to Private

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

Create a File Share Storage Account

- Now create a file storage account.
- Go to the home page and click on **Resource group**.



Figure 4.381: Select Resource Group

- Click on **Add.**

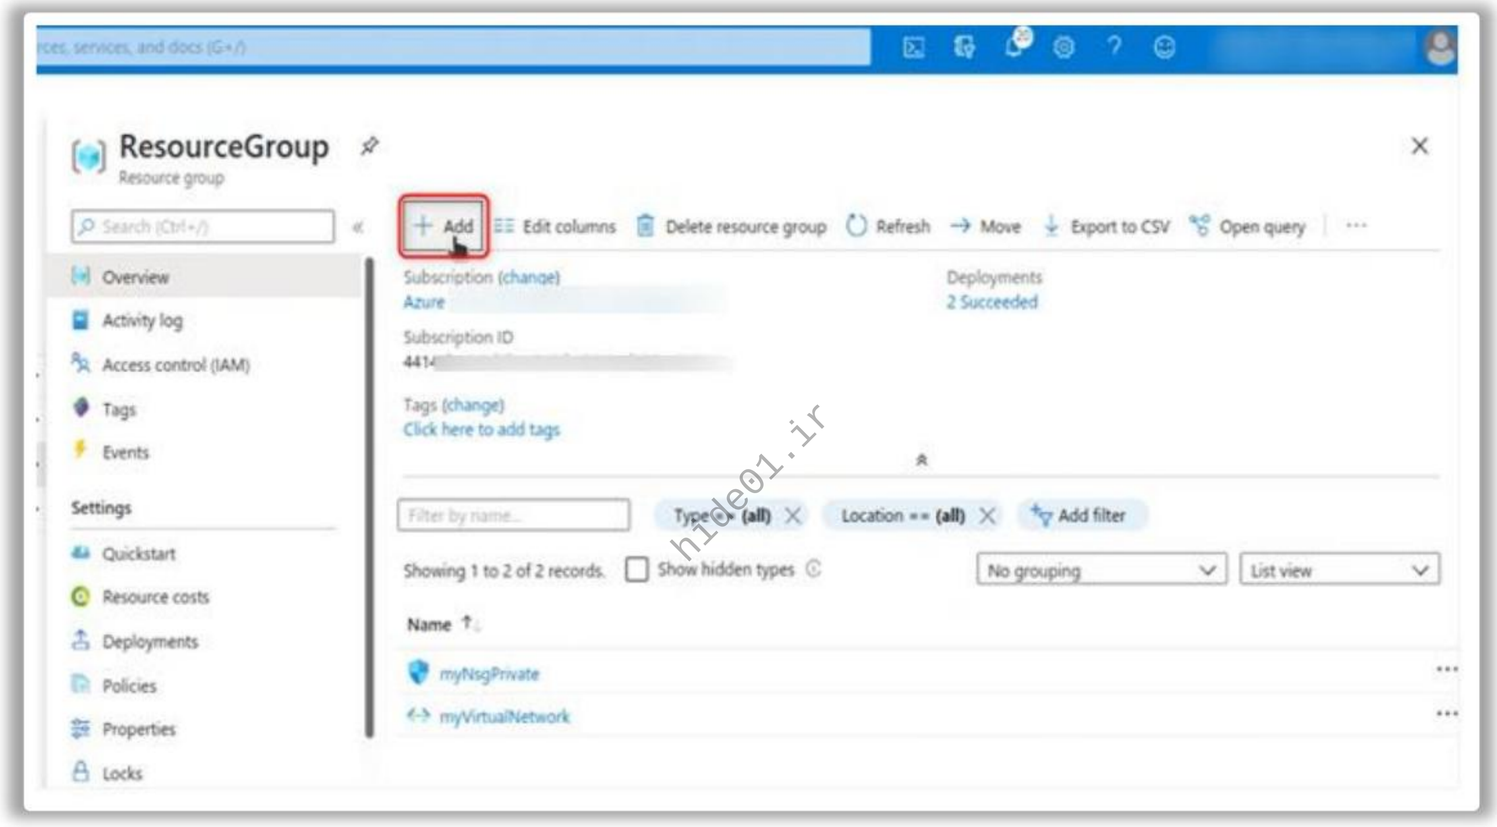


Figure 4.382: Click Add

- In the search field type **storage account** and click on it.

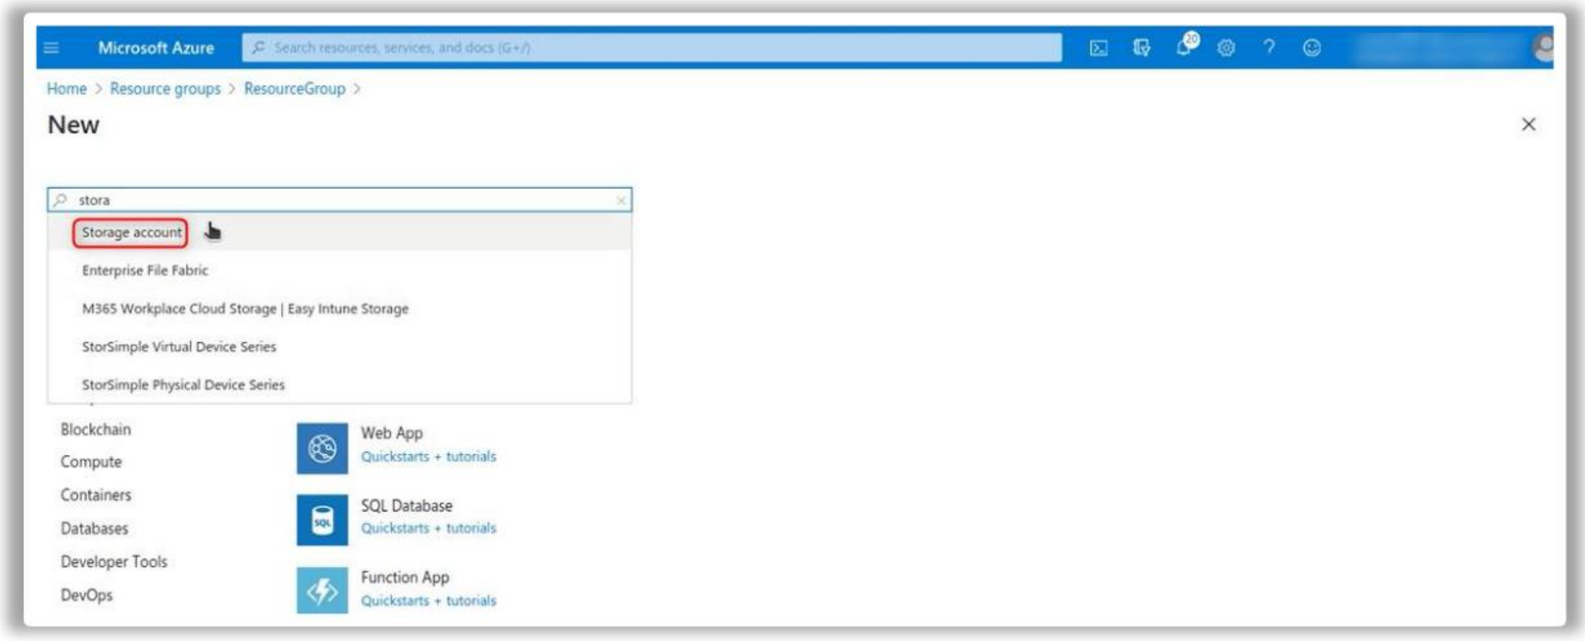


Figure 4.383: Type Storage Account

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

- Click on **create**.

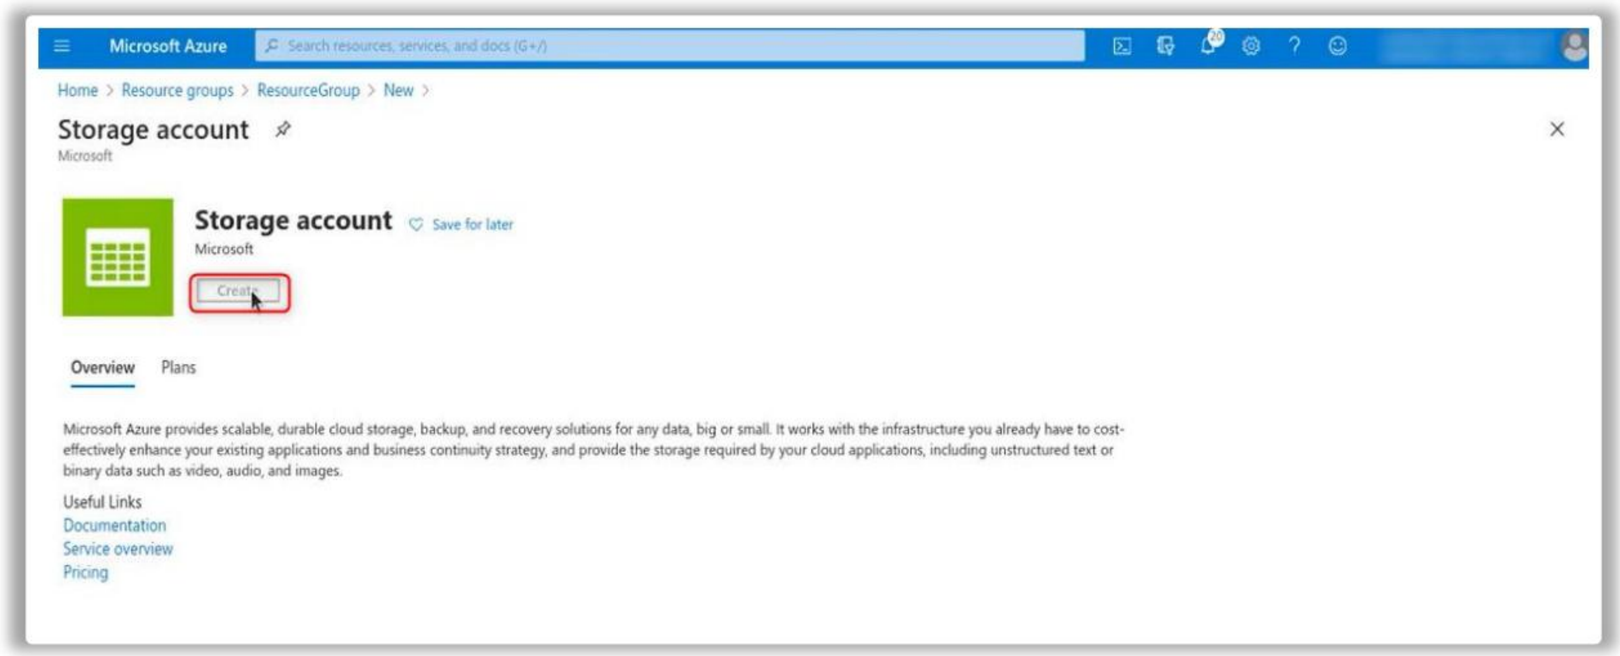


Figure 4.384: Click Create

- Enter the basic details and click on **Next: Networking.**

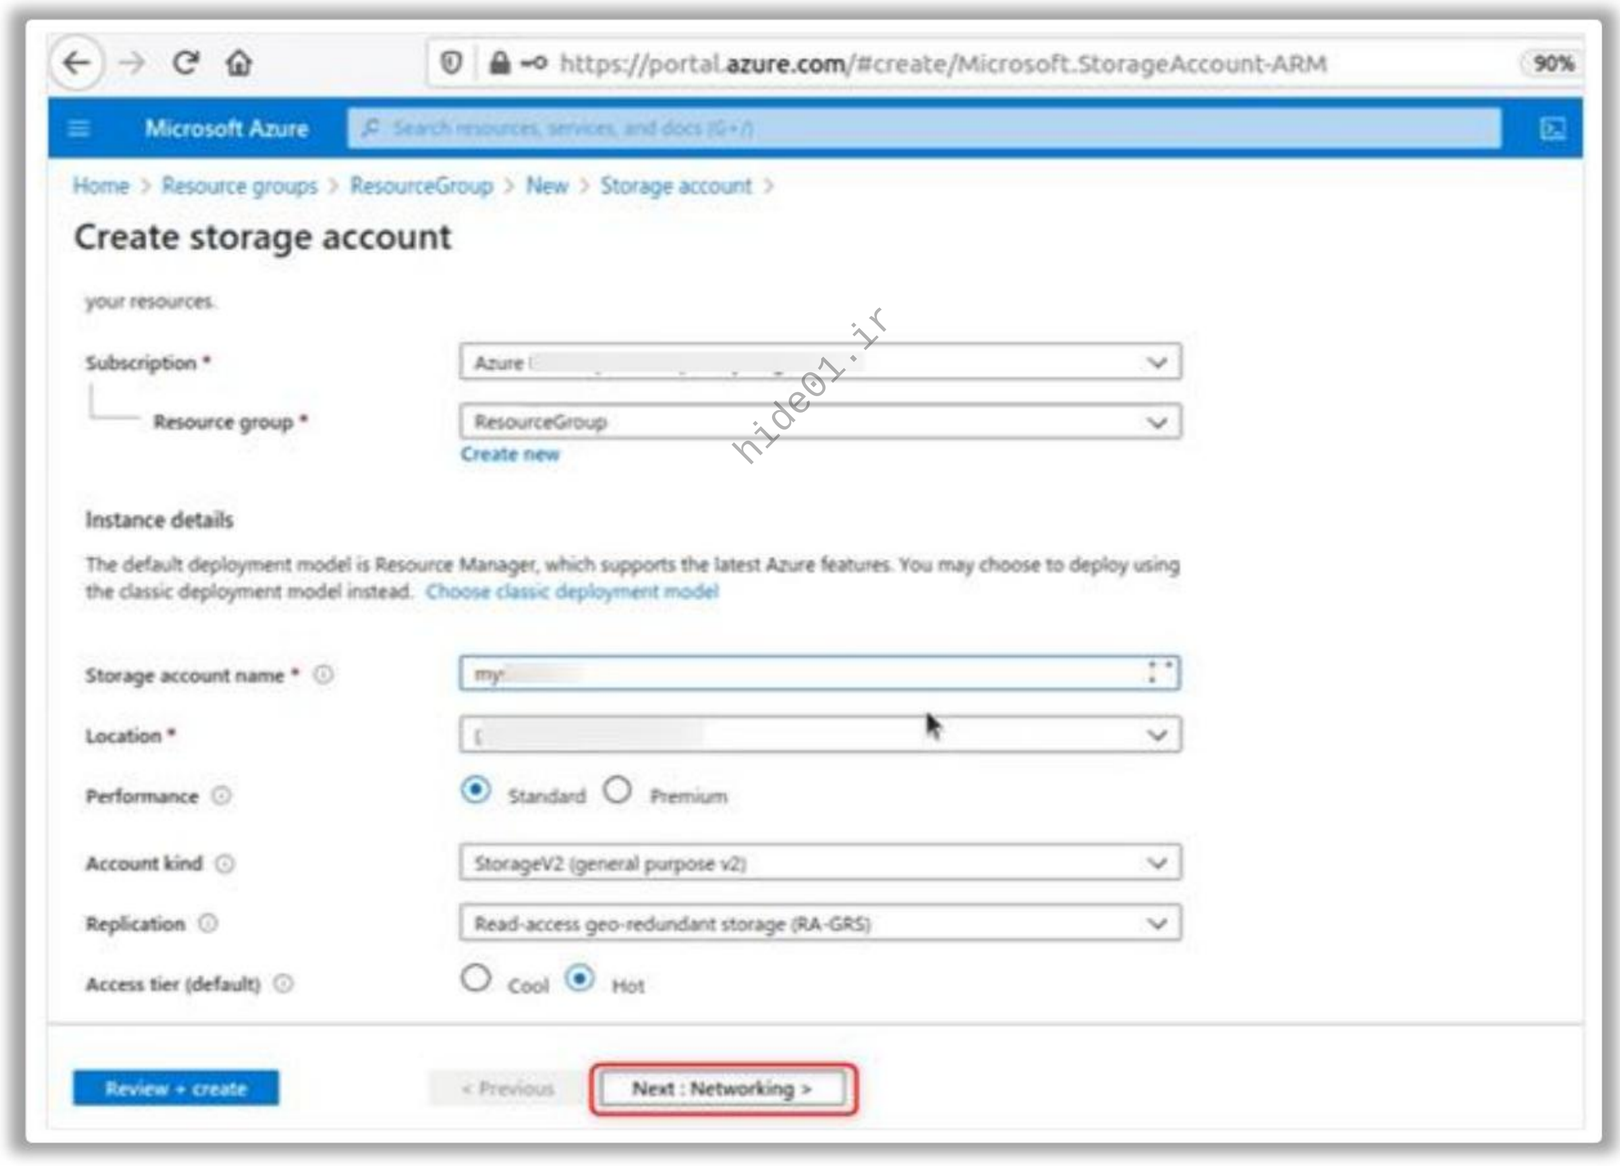


Figure 4.385: Select Next: Networking

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

- Click on **Review + create**.

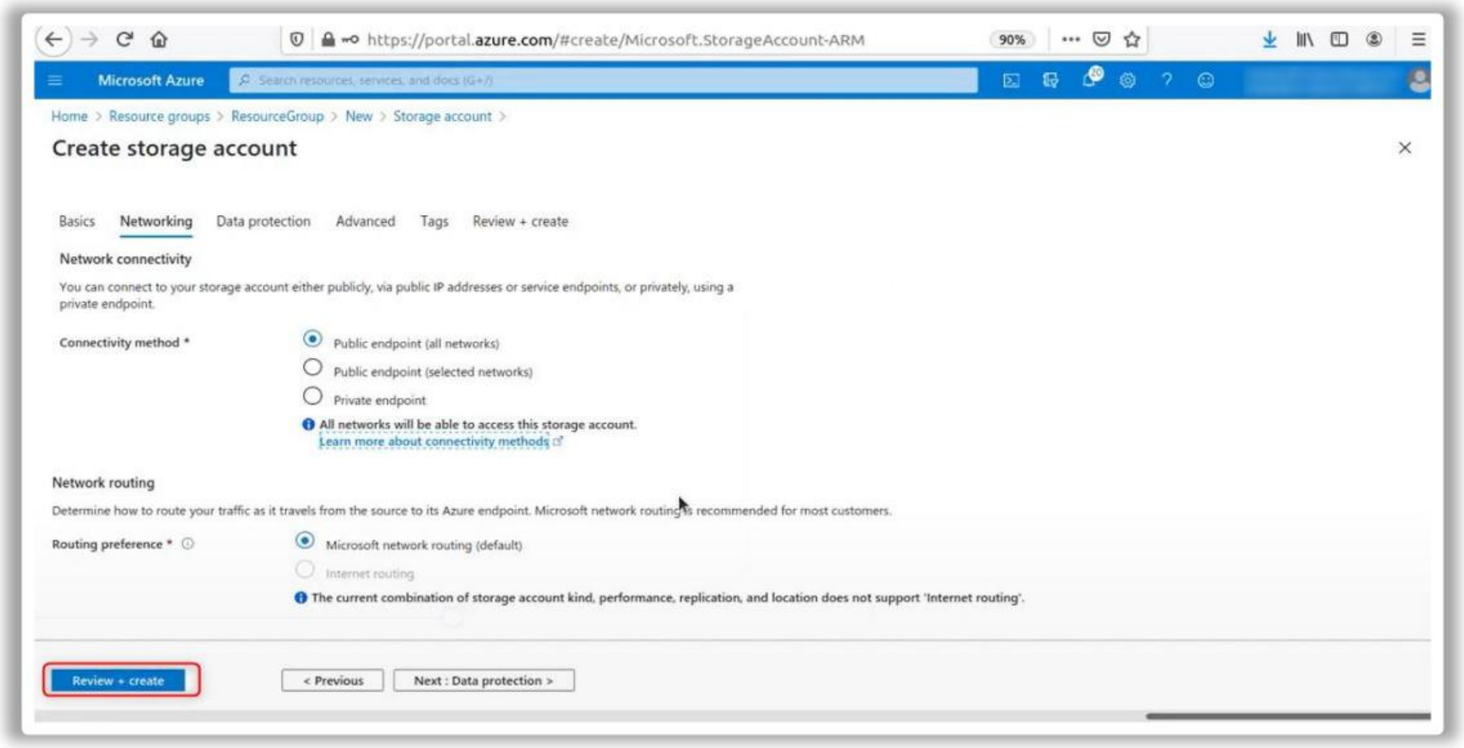


Figure 4.386: Click Review + Create

- Click on **create**.

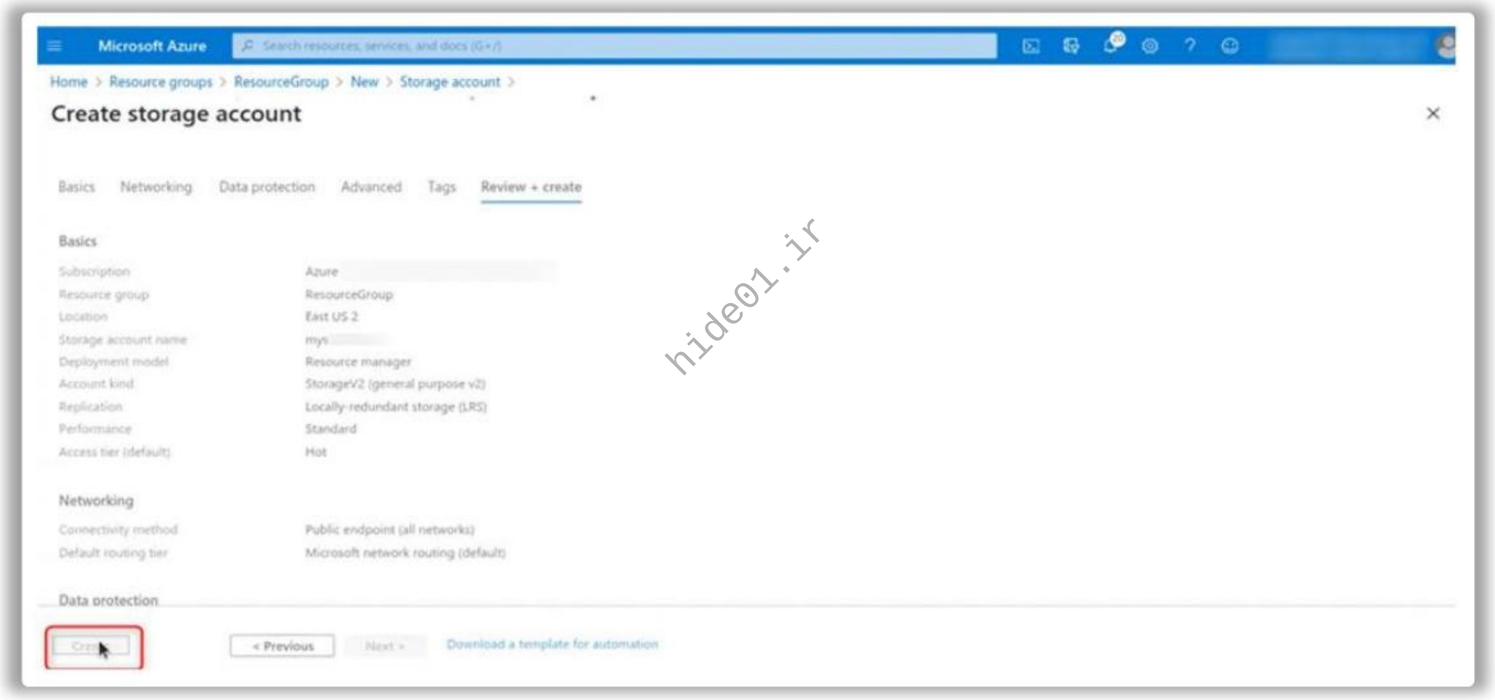


Figure 4.387: Click Create

- A storage account has been successfully created.
- Click on **Go to resource**.

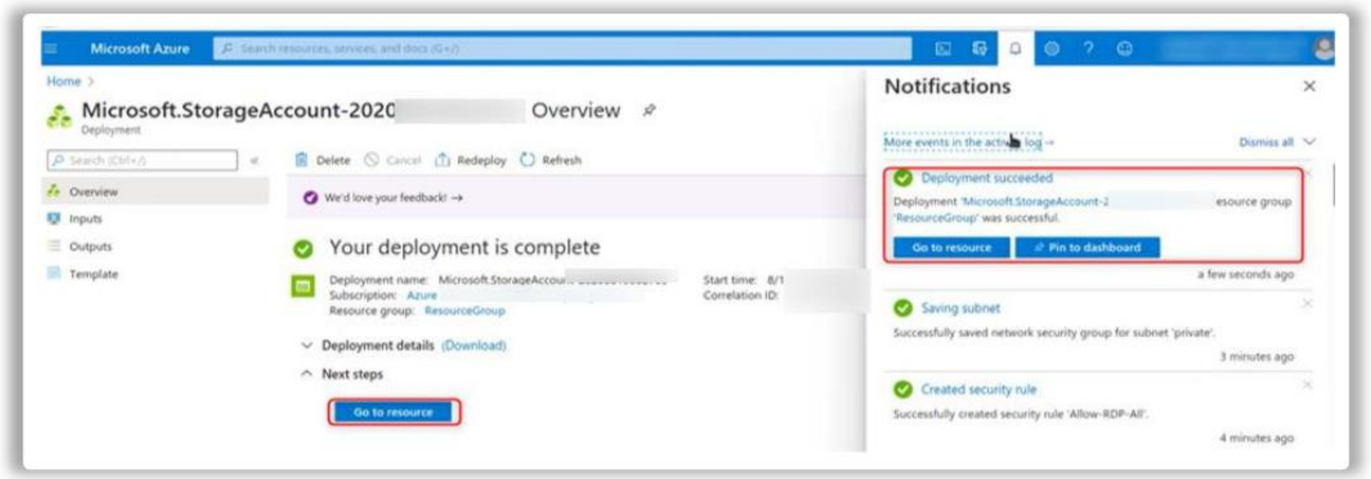


Figure 4.388: Click Go to Resource

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

- Go to **File shares**.

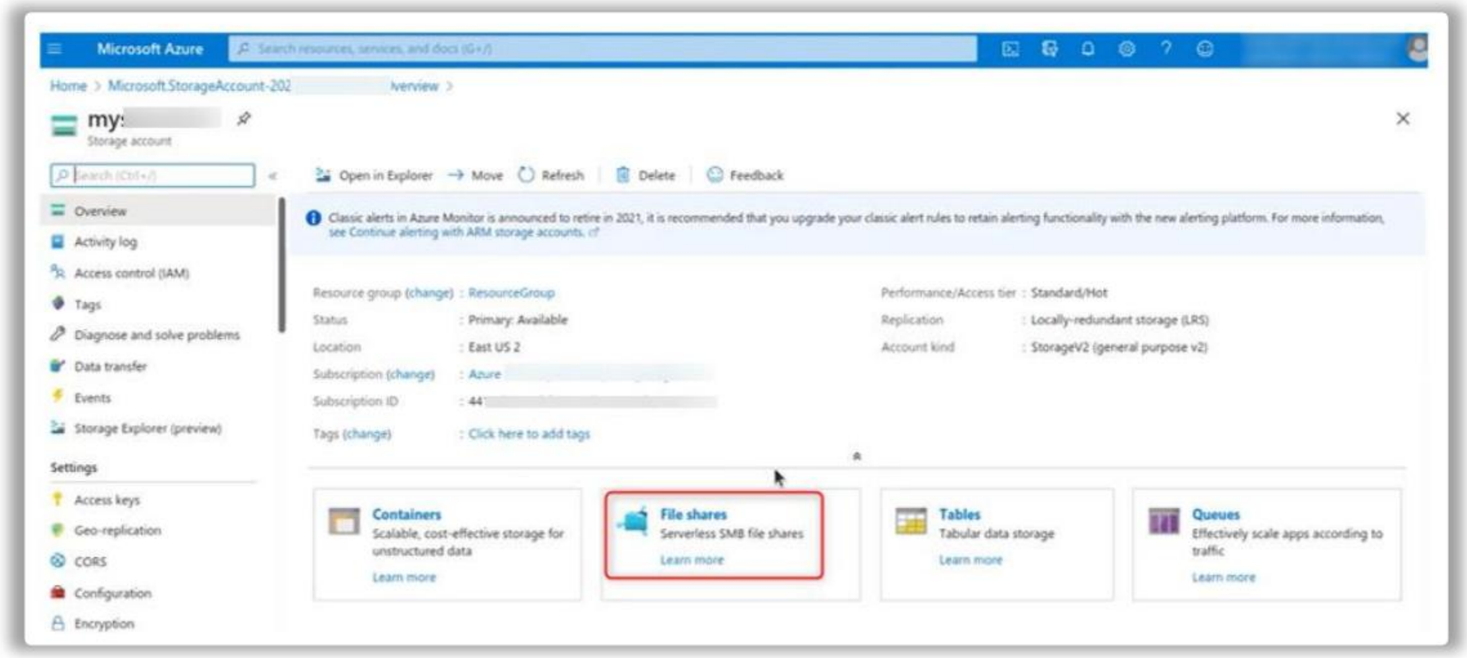


Figure 4.389: Click File Shares

- Select **+ File share**.
- On the New file share panel, select **Name,** and **Quota**.
- Click on **Create**.

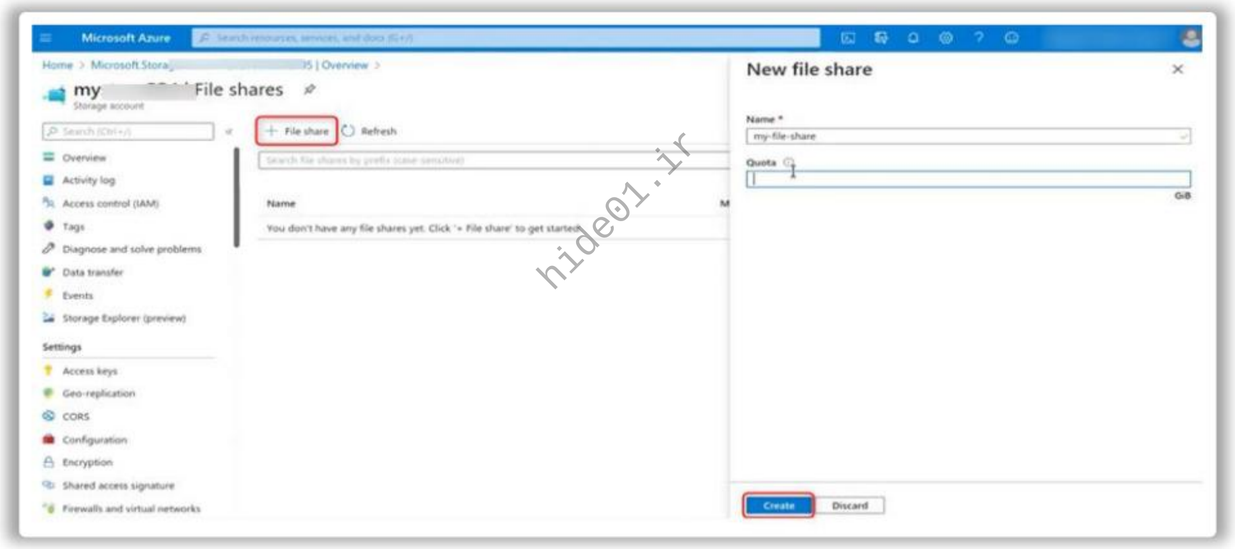


Figure 4.390: Click Create

- File share has been successfully created.

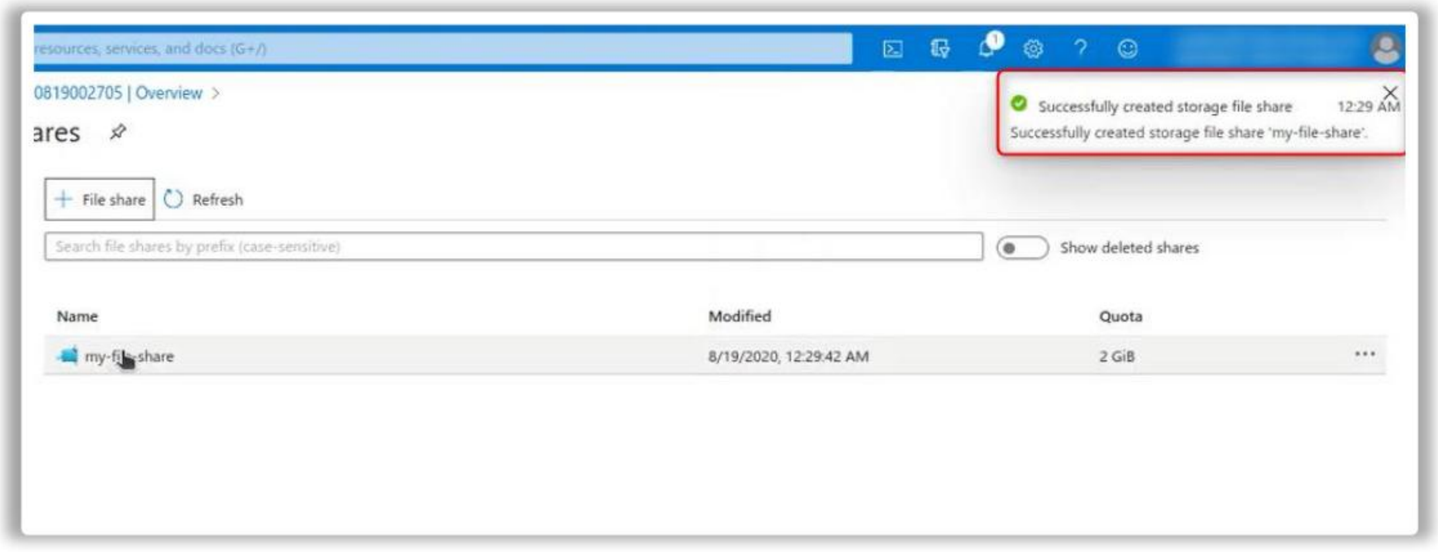


Figure 4.391: File Share is Created

- Under **Settings**, click on **Firewalls and virtual networks**.
- Under **Allow access from,** select the radio button of **Selected networks**.

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

- Next in Virtual networks, select **Add existing virtual network**.
- In the **Add networks panel**, select the **virtual network** from the dropdown, and select **subnet** to **private** which was created.
- Click on **Add**.
- Finally, click on **Save**.

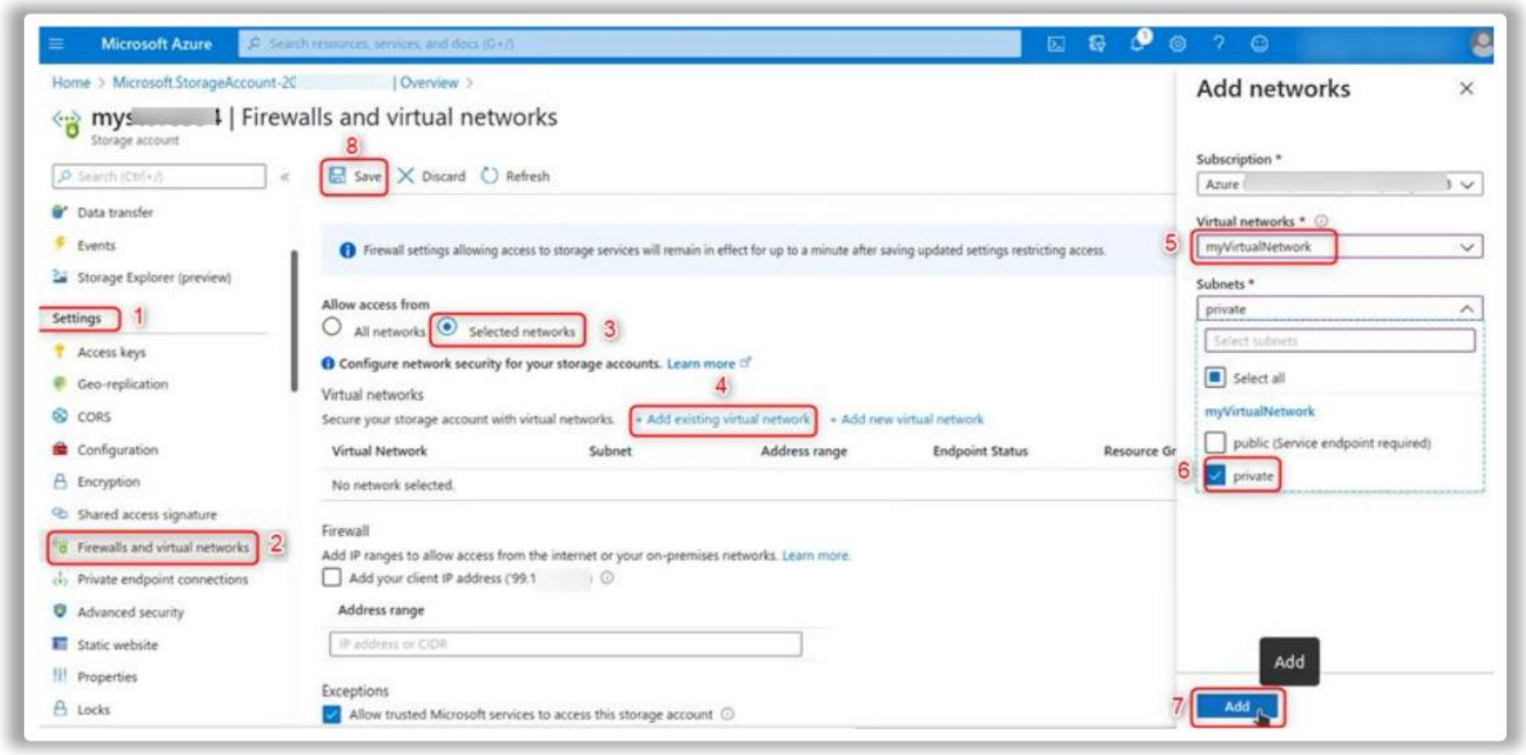


Figure 4.392: Click Save

- Firewall and virtual network settings for selected storage have been successfully updated.

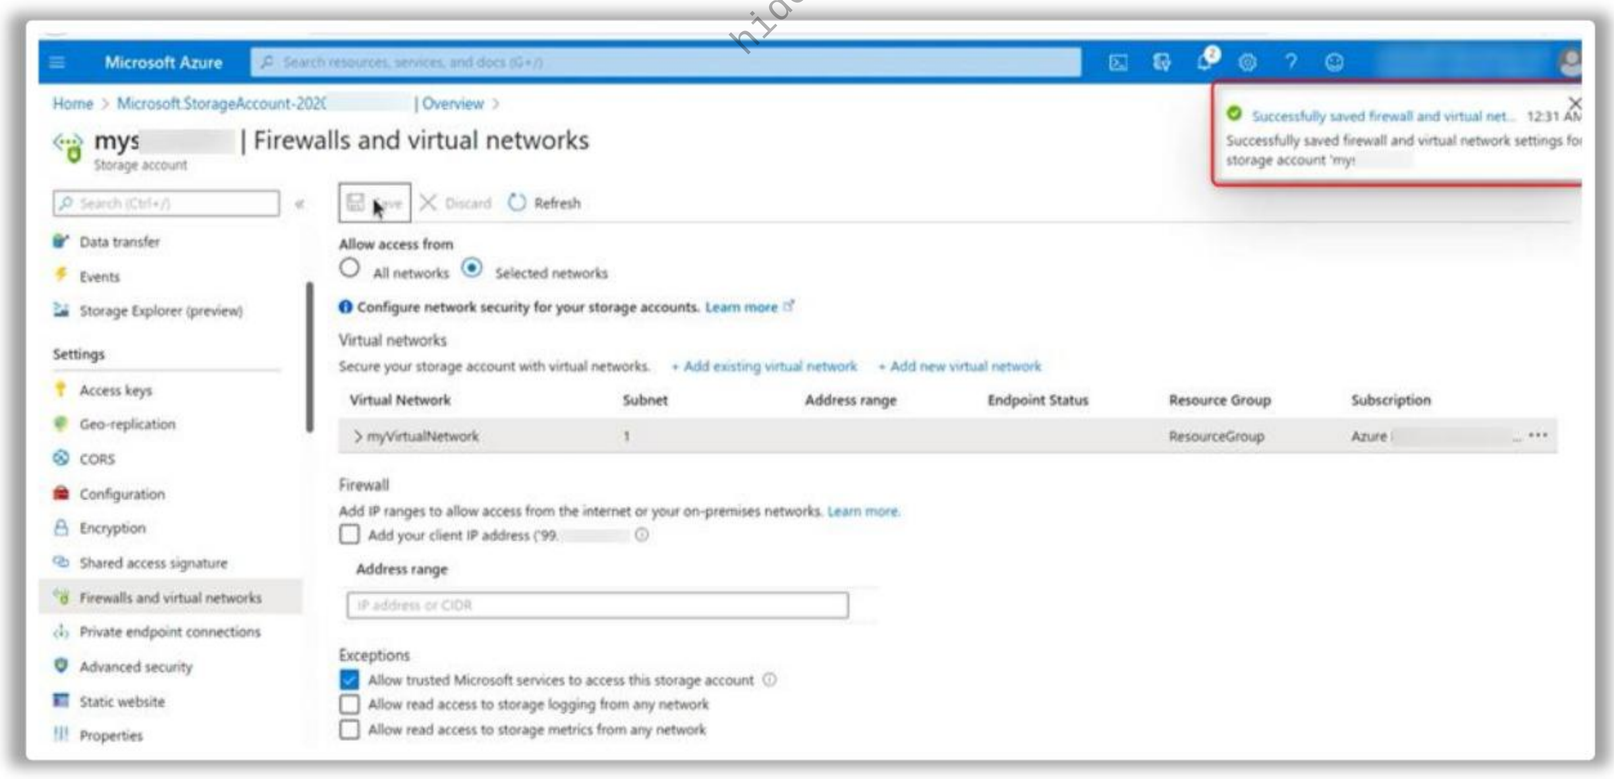


Figure 4.393: Firewall and Virtual Network is Updated

Create Two Virtual Machines

- Now two VMs are required, one with public and the other with private access.
- First, create a public level VM.

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

- Go to the homepage and select **Resource groups**.



Figure 4.394: Click Resource Groups

- Click on **Add**.

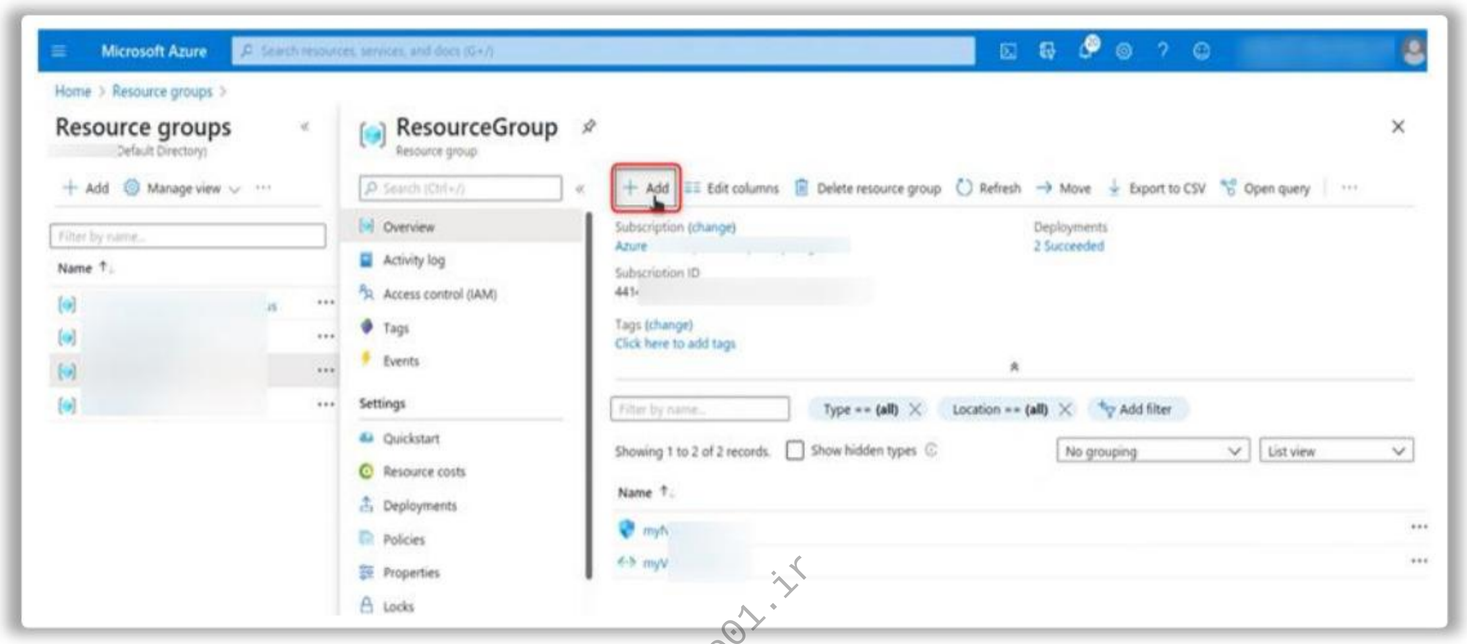


Figure 4.395: Click Add

- Click on **Windows Server 2016 Datacenter**.

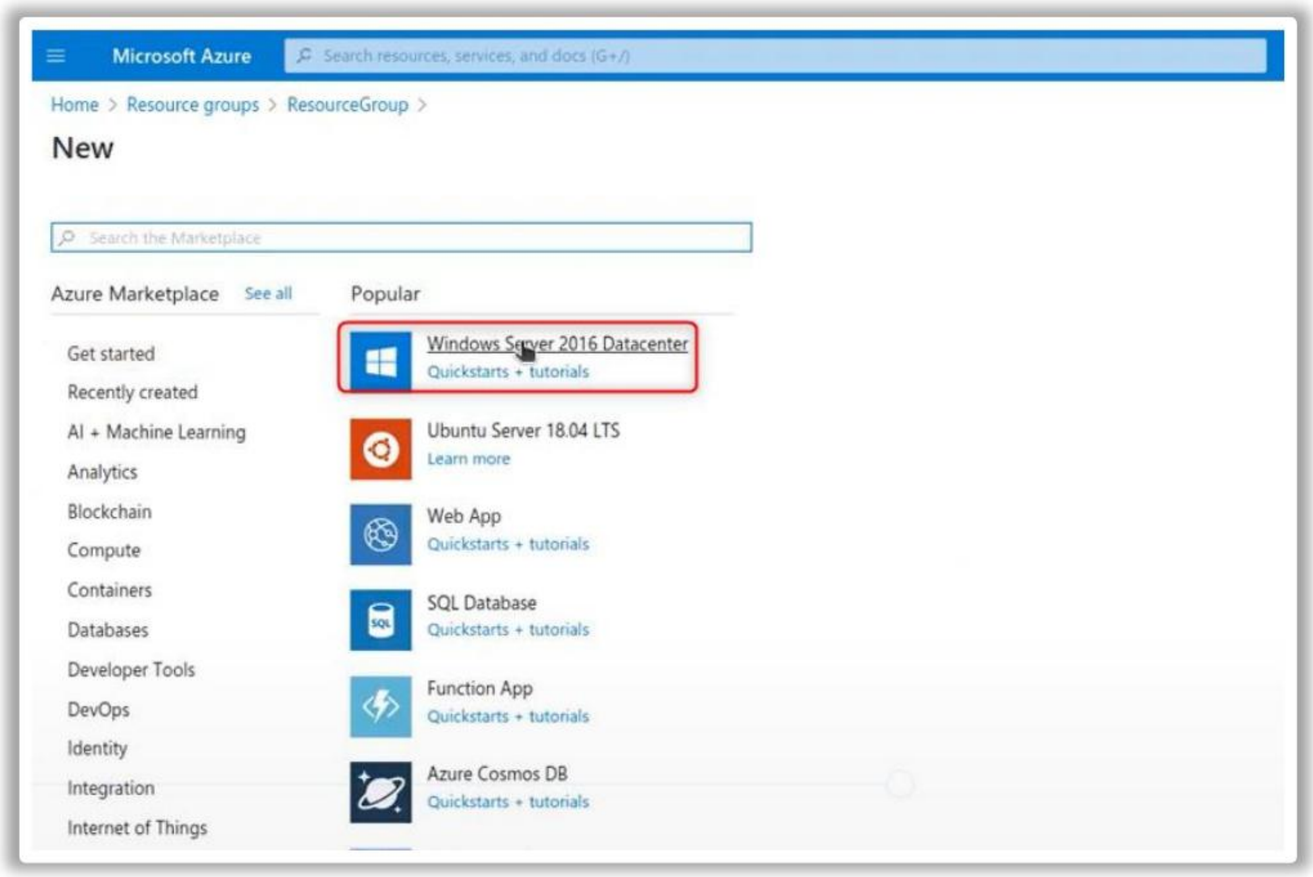


Figure 4.396: Click Windows Server 2016 Datacenter

- Give this **VM name** such as "**myVmPublic**."

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

- Select the remaining fields.

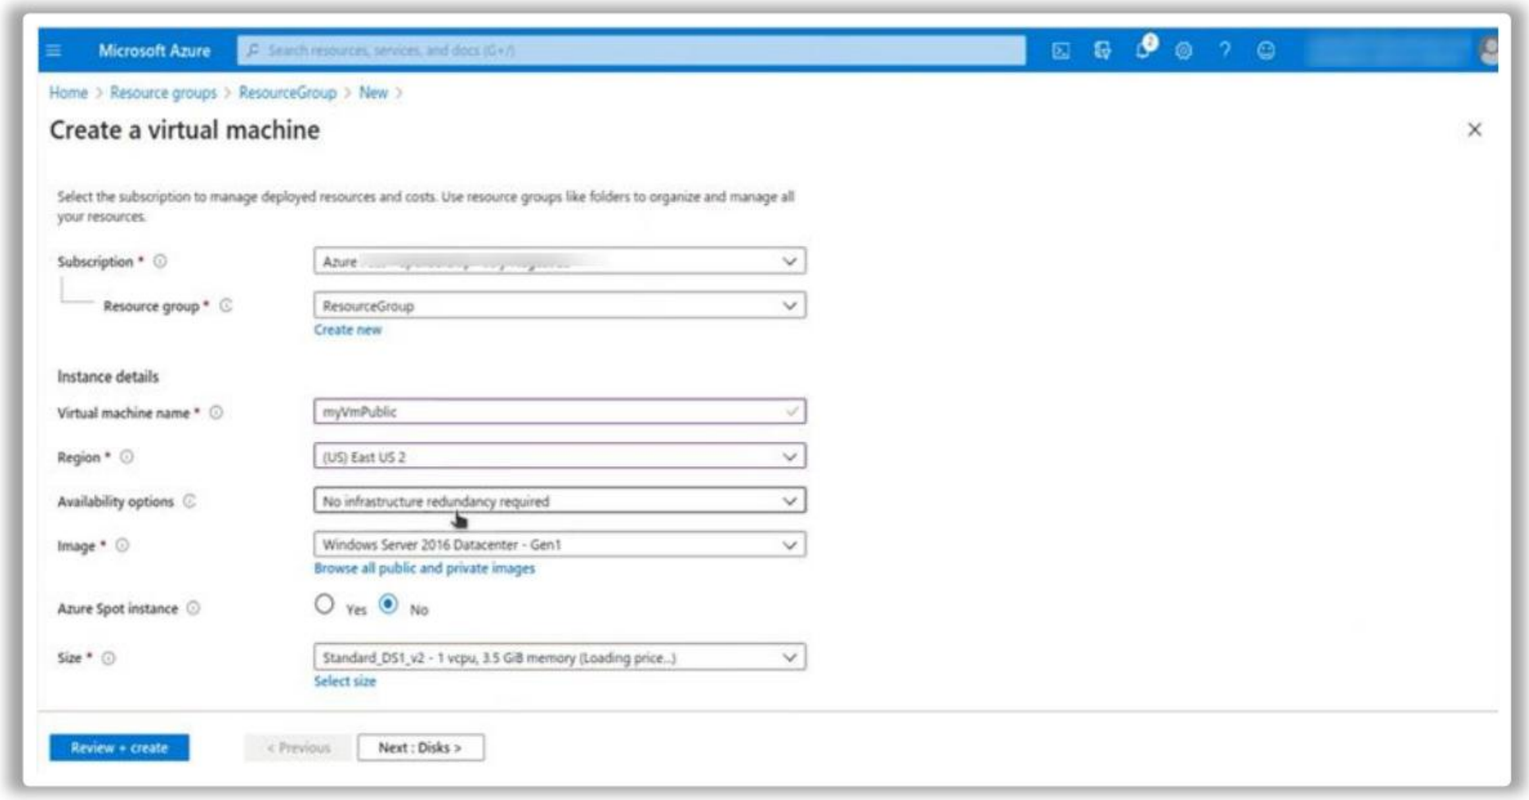


Figure 4.397: Select Required Field

- Scroll down and create a user name and password for this VM.
- Select the **Inbound ports** to **RDP 3389**.
- Click on **Next: Disks**.

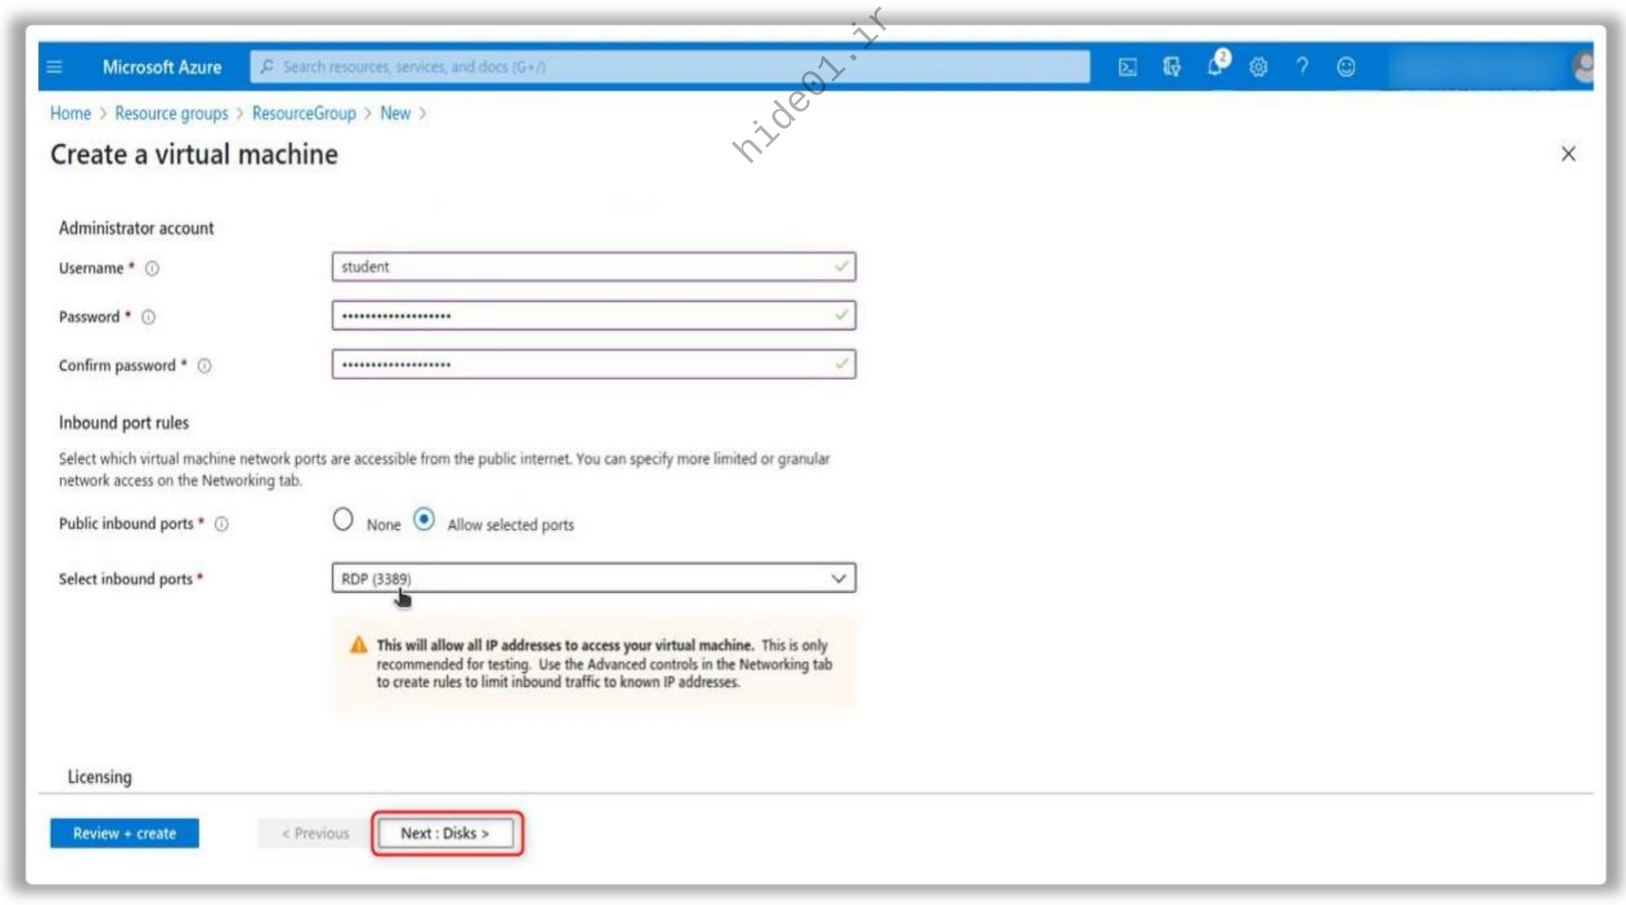


Figure 4.398: Click Next: Disks

- Select the **Subnet** to public.
- Select Basic for NIC network security group.
- Select **RDP 3389** for **Select inbound ports**.

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

- Click on **Next: Management**.

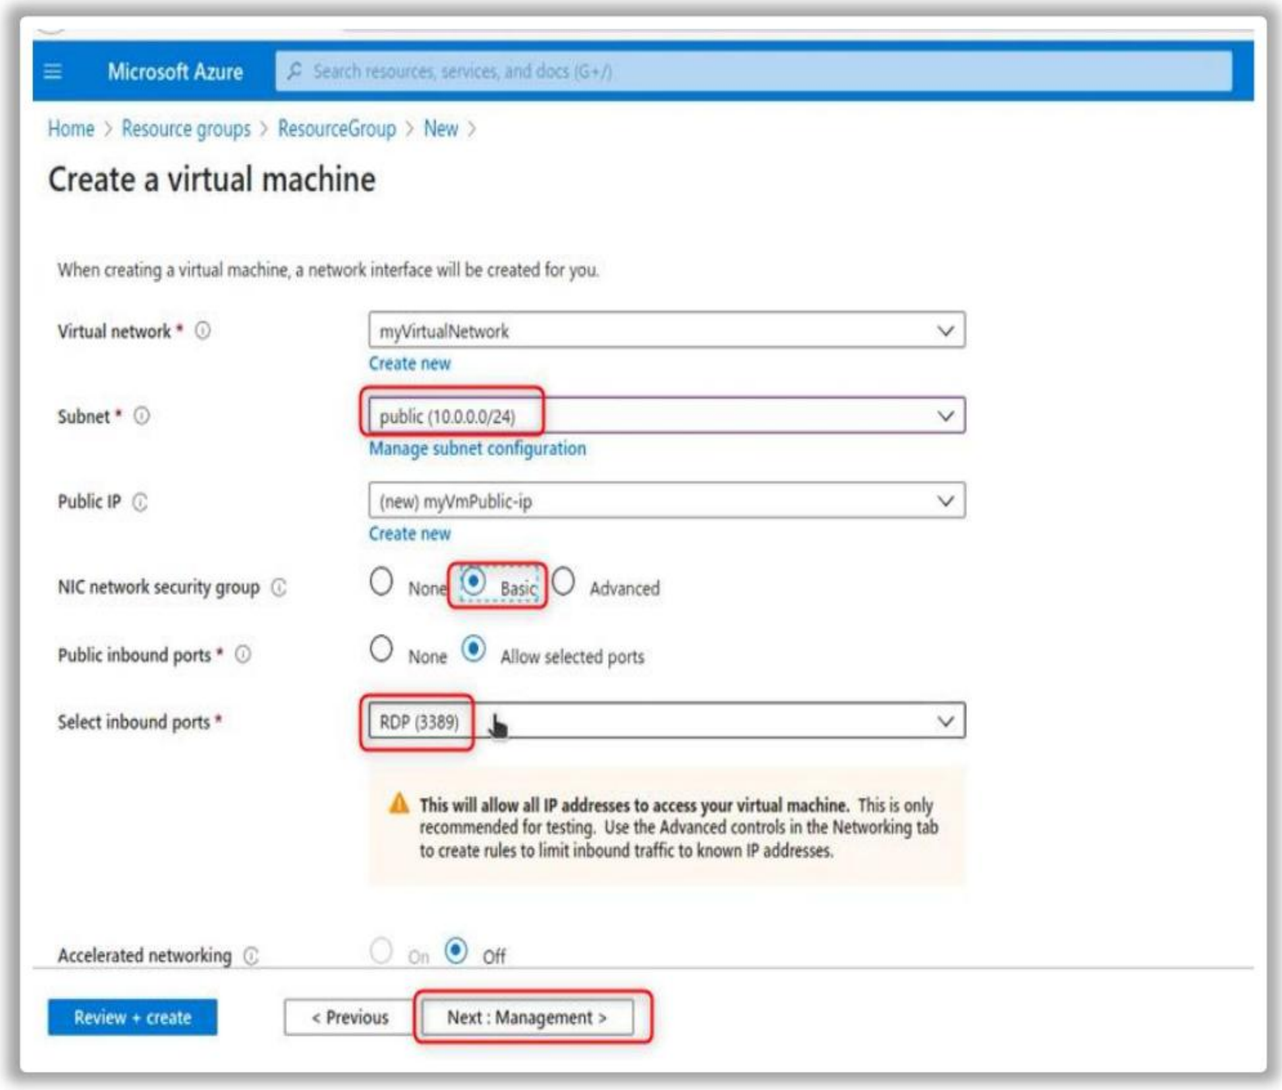


Figure 4.399: Click Next: Management

- Turn off **Boot diagnostics**.
- Click on **Next: Advanced**.

hide01.ir

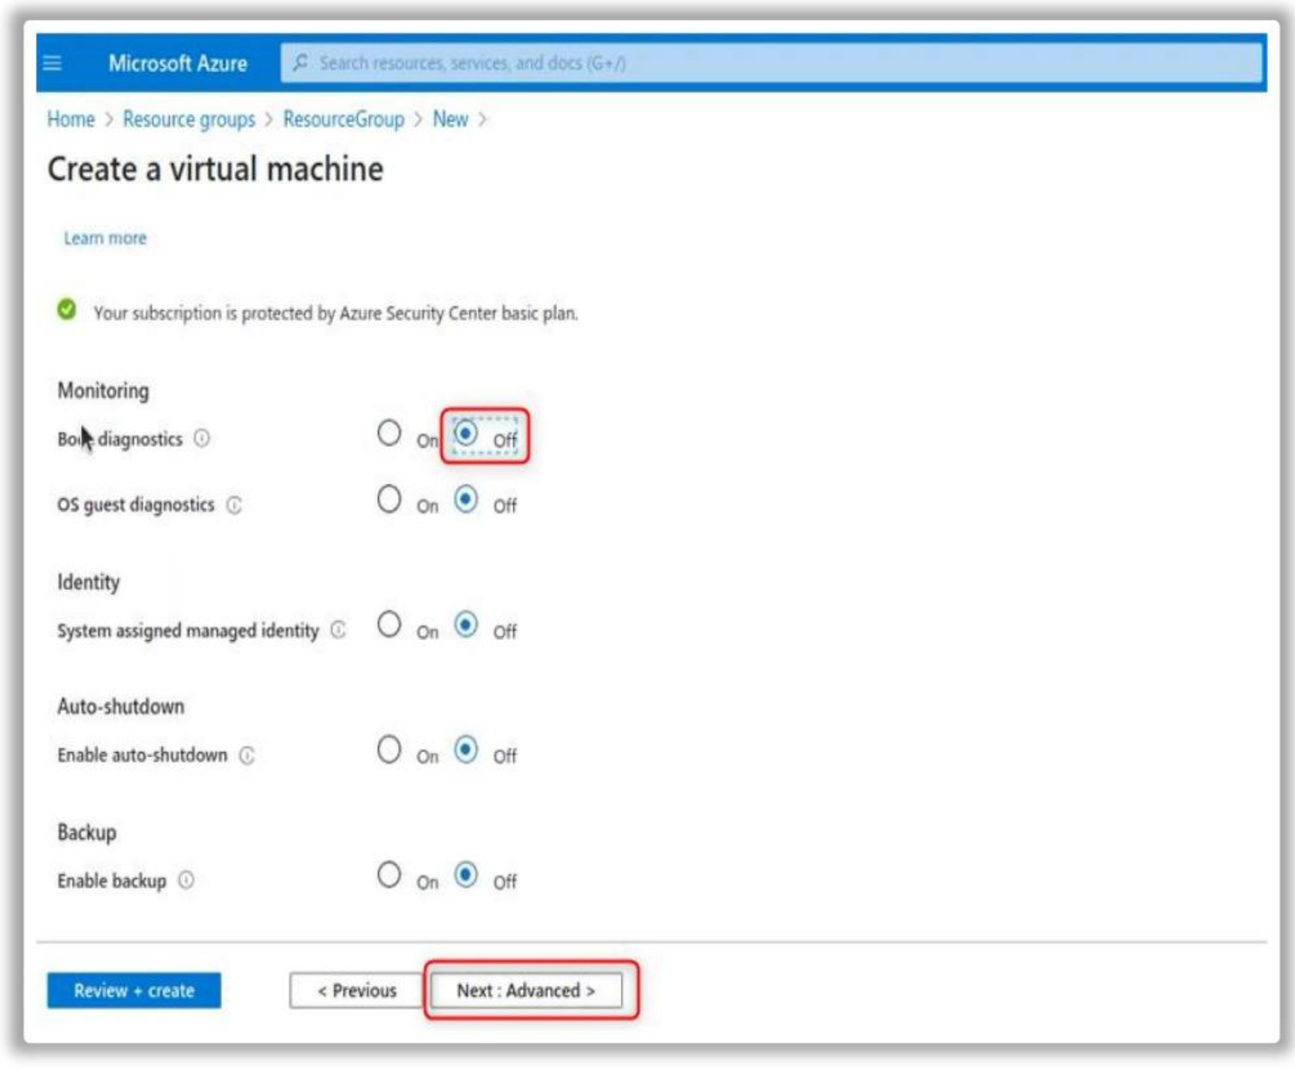


Figure 4.400: Click Next: Advanced

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

- Click on **Review + create**.

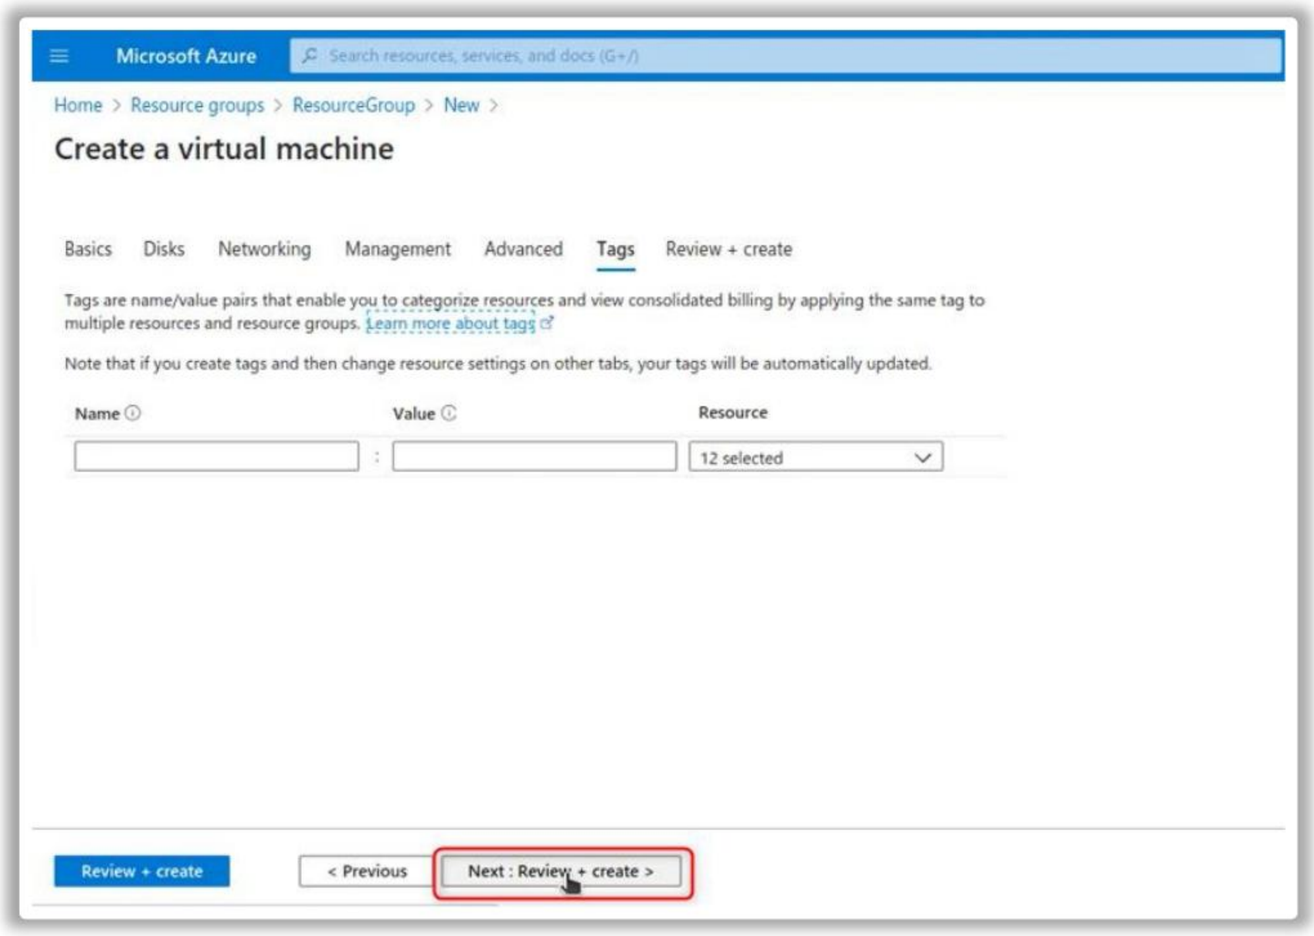


Figure 4.401: Click Review +create

- Click on **Create**.

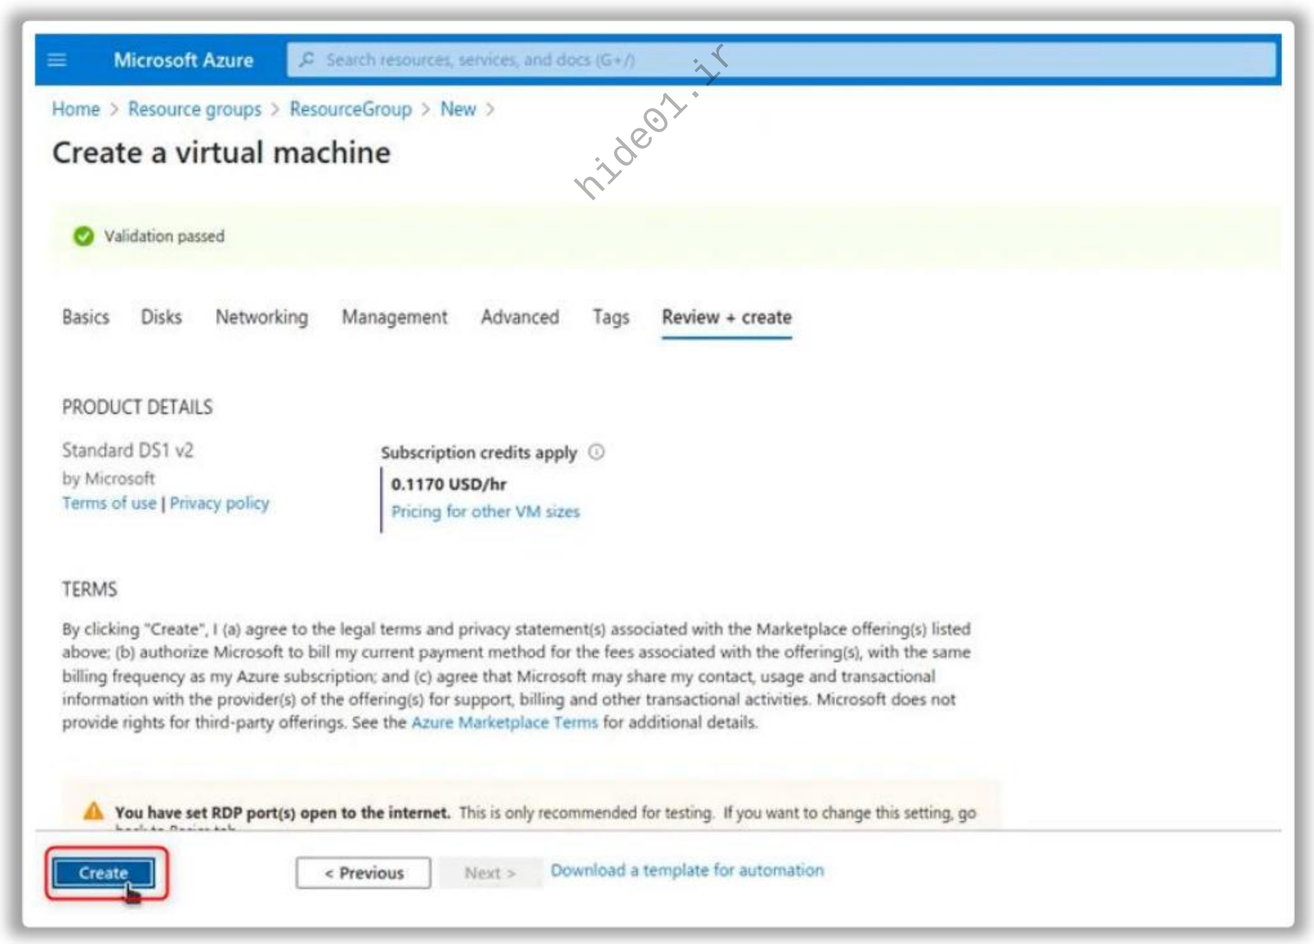


Figure 4.402: Click Create

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

- The VM has been created successfully.

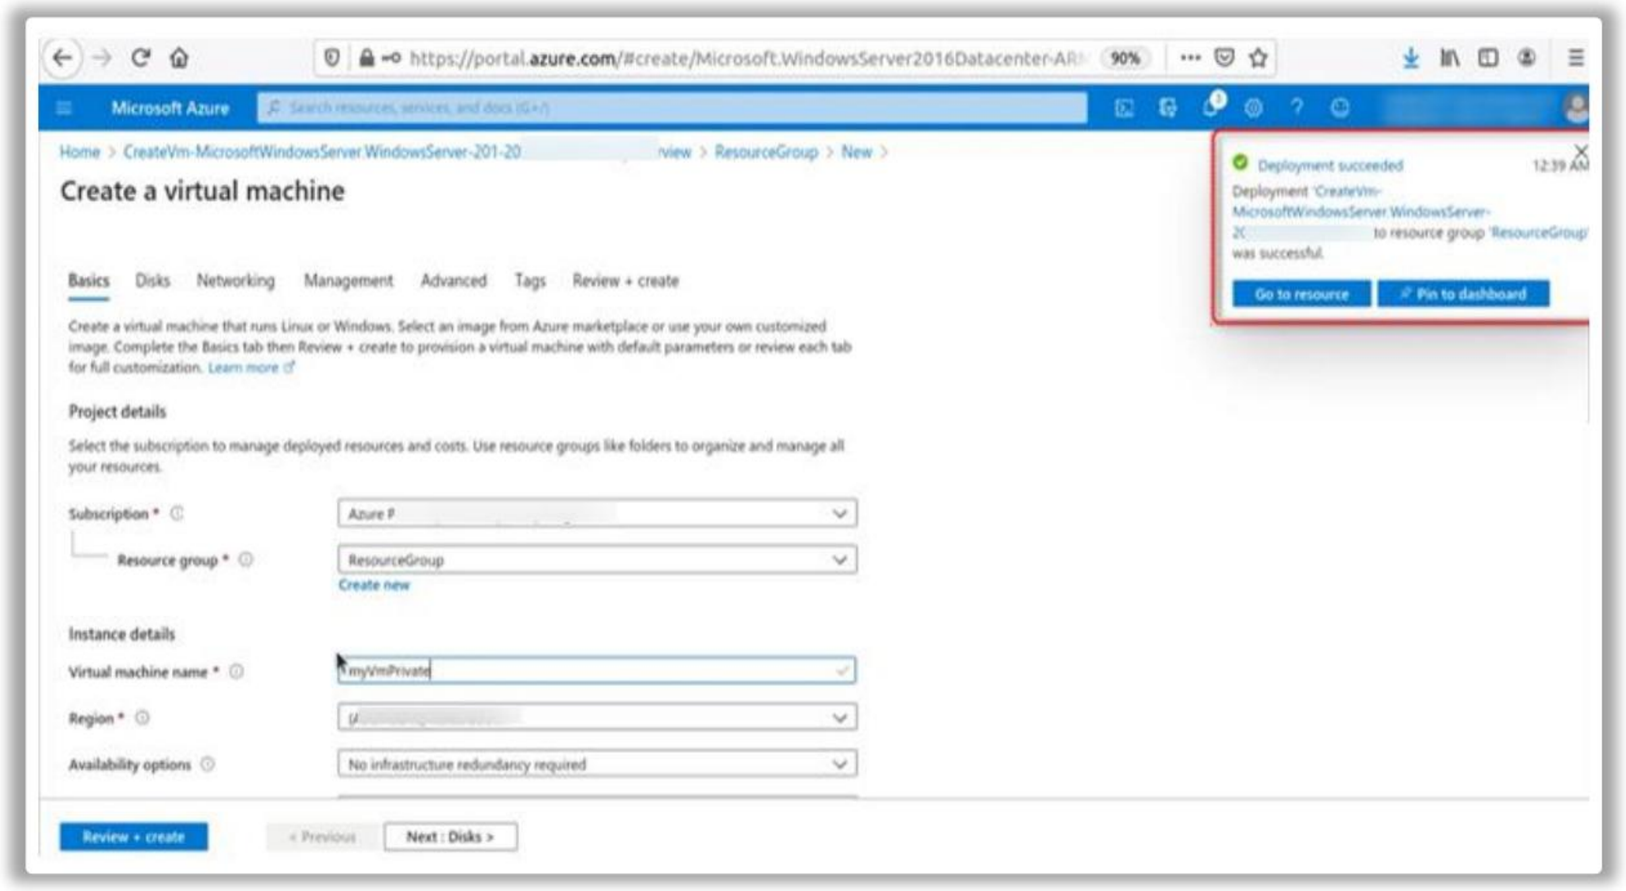


Figure 4.403: VM is Created

- Now create a private level VM.
- Go to the home page, and click on **Resource groups**.



Figure 4.404: Click Resource Groups

- Click on **Add**.

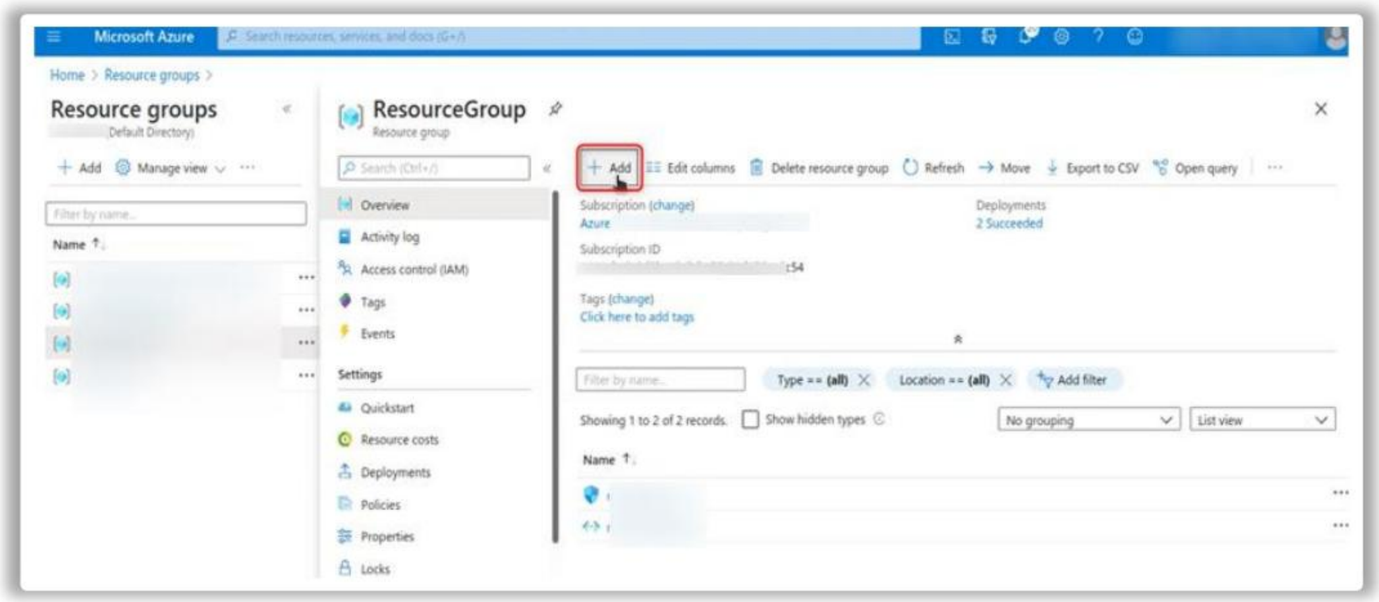


Figure 4.405: Click Add

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

- Select **Windows Server 2016 Datacenter**.

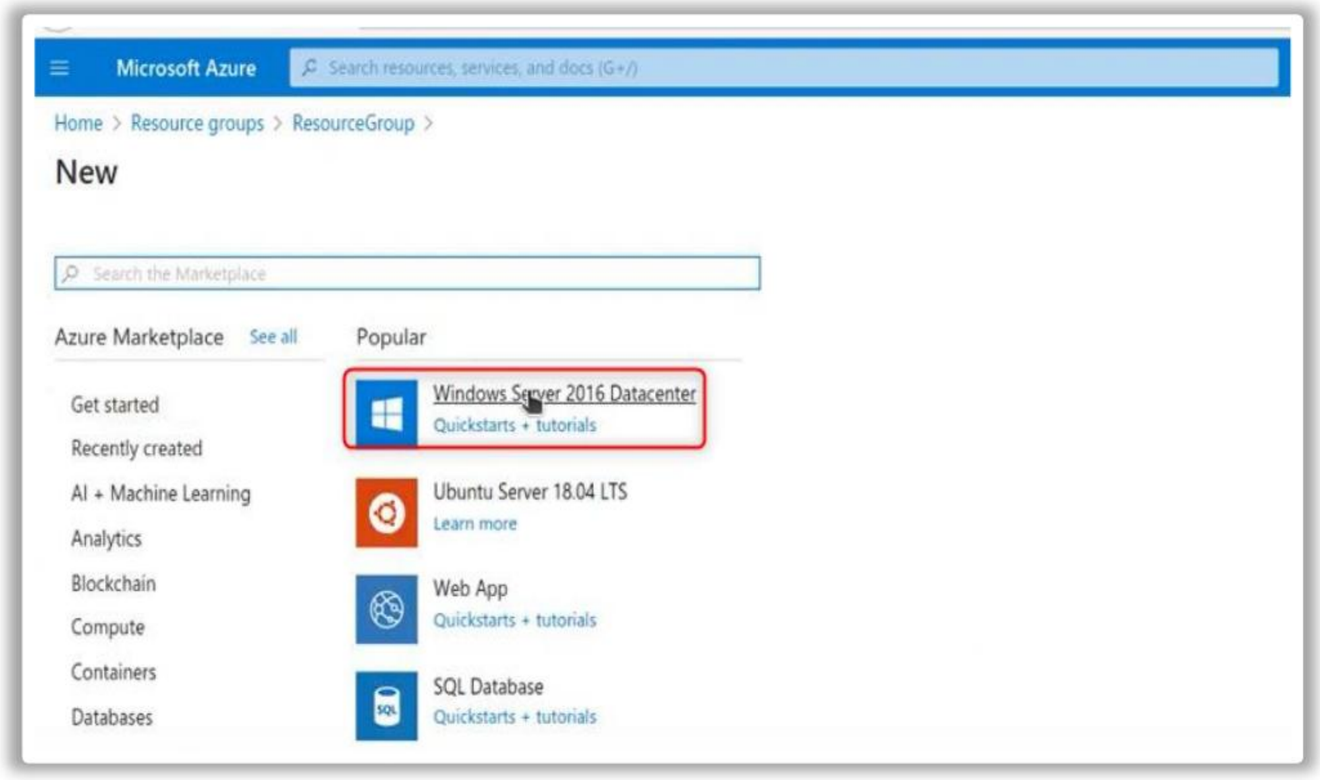


Figure 4.406: Click Windows Server 2016 Datacenter

- Give this **Virtual machine name** like "**myVmPrivate**."
- Scroll down and create a user name and password.

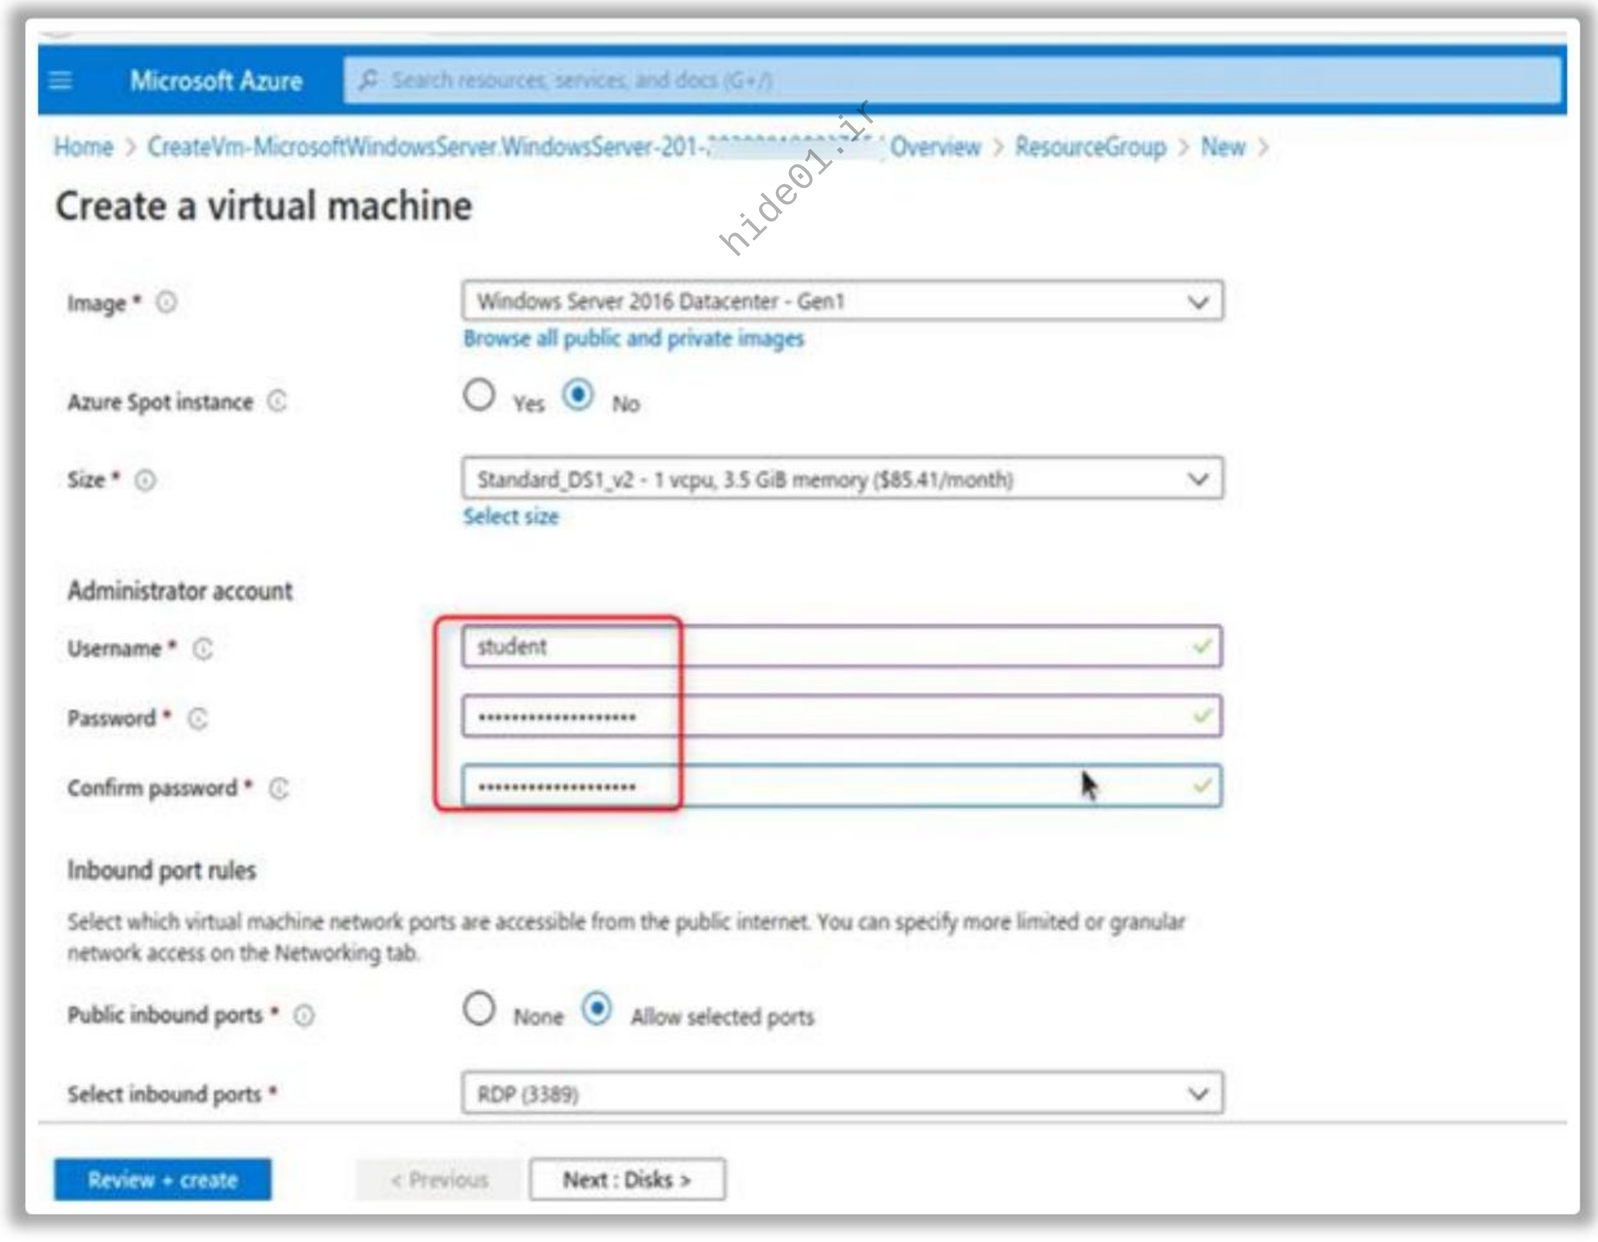


Figure 4.407: Click Username and Password

- Select the **Public inbound ports** to **None**.

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

- Click on **Next: Disks**.

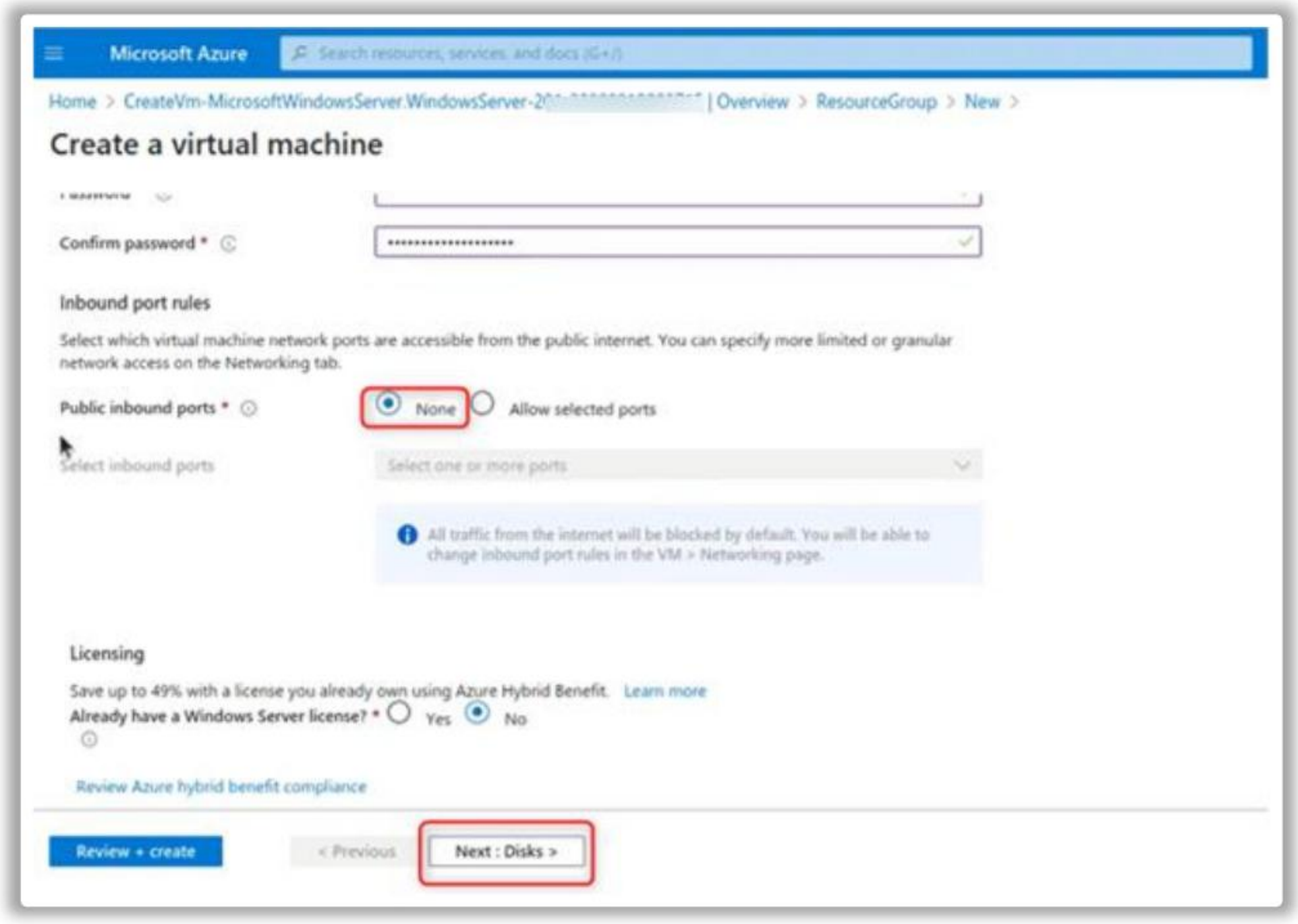
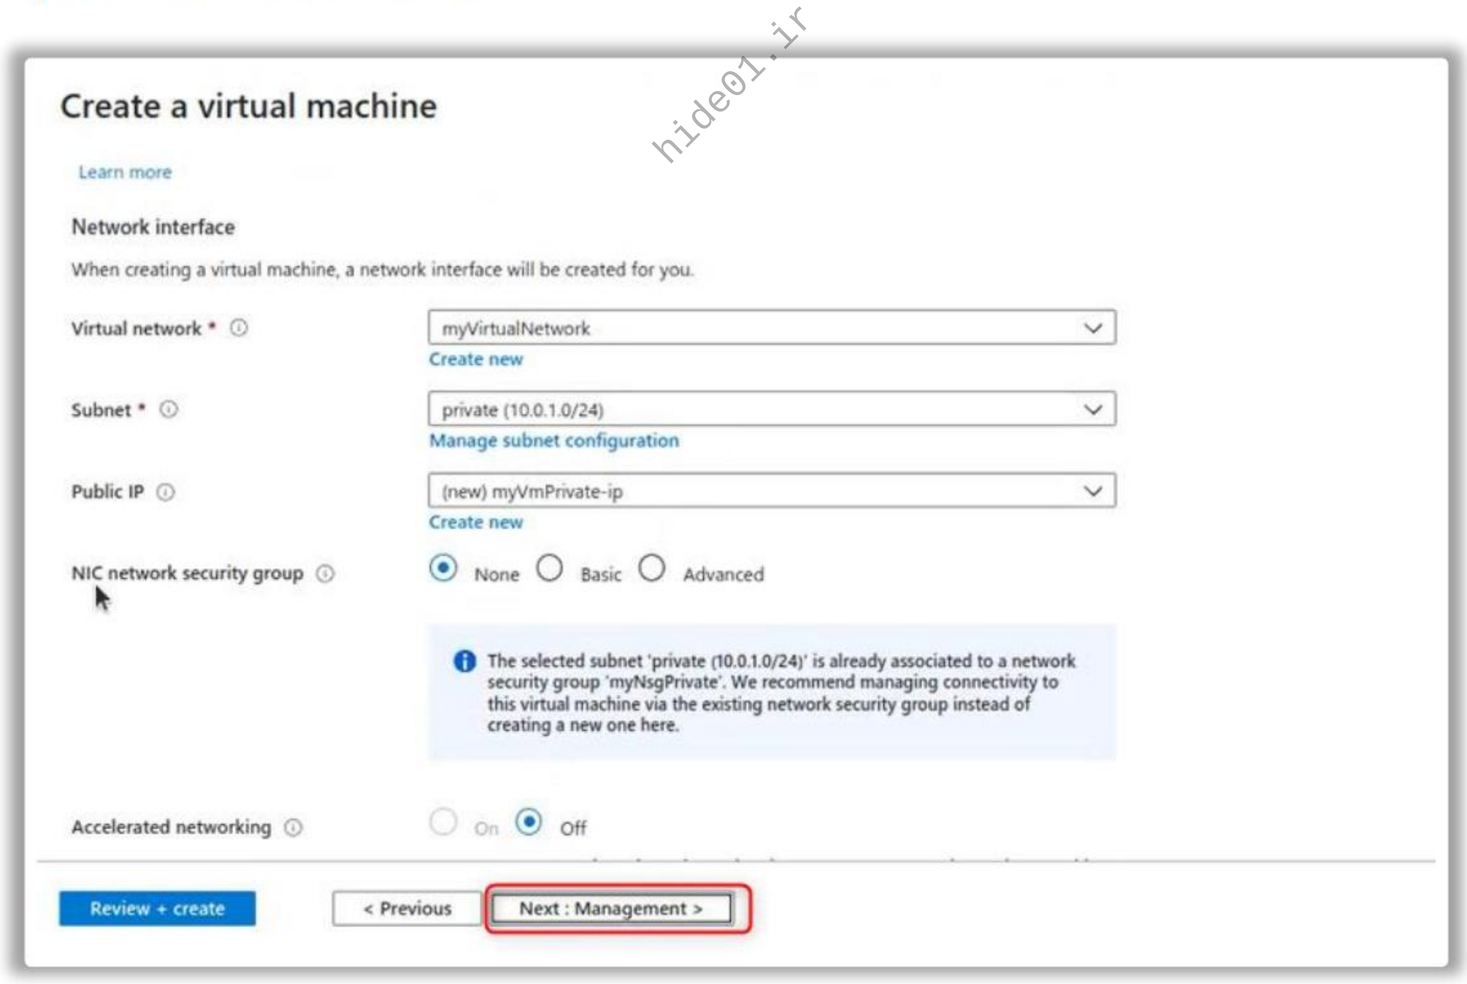


Figure 4.408: Click Next: Disks

- Select the **Virtual network** from the dropdown.
- Select private IP for **Subnet**.



hide01.ir

Figure 4.409: Select Private IP

- Turn off **Boot diagnostics**.

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

- Click on **Next: Advanced**.

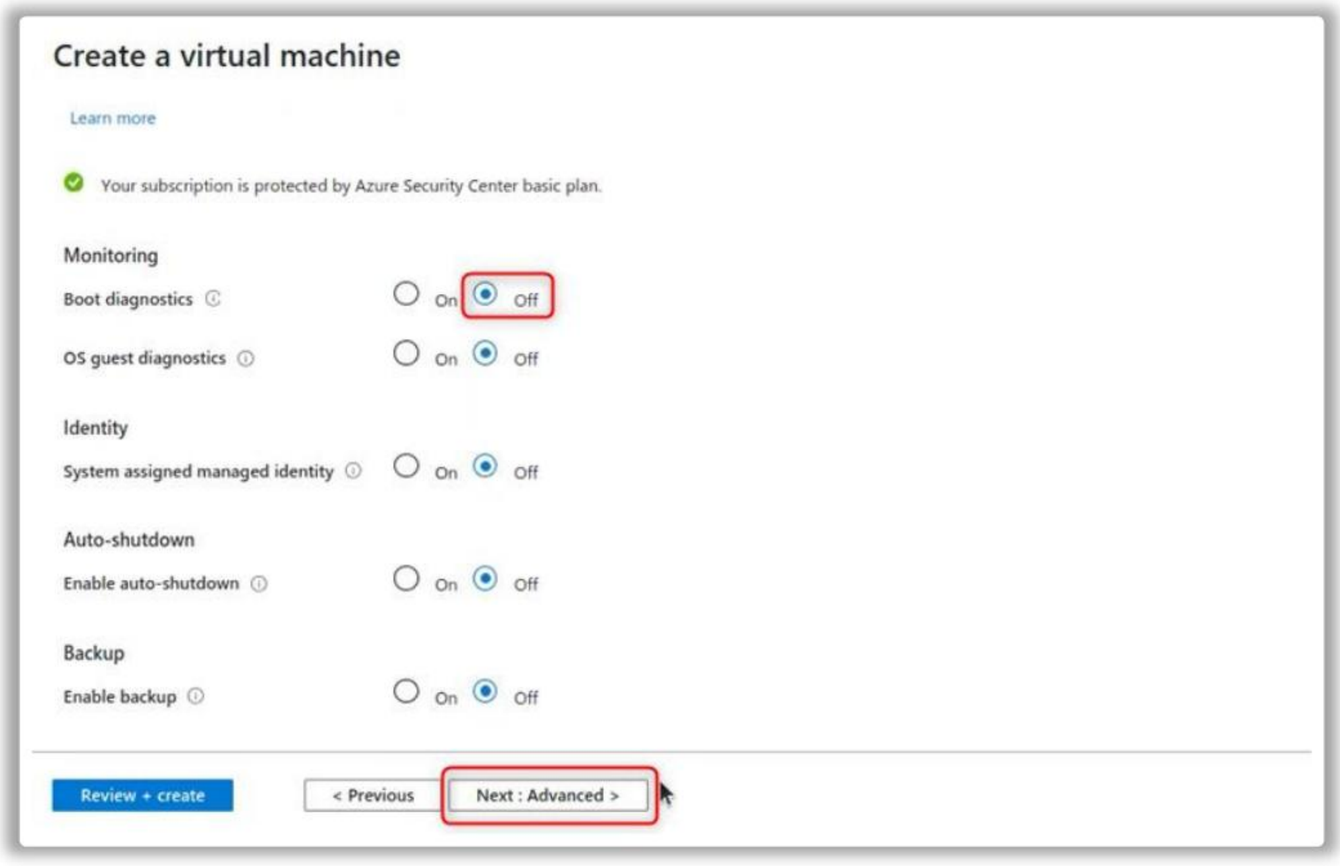


Figure 4.410: Click Next: Advanced

- Click on **Review + create**.

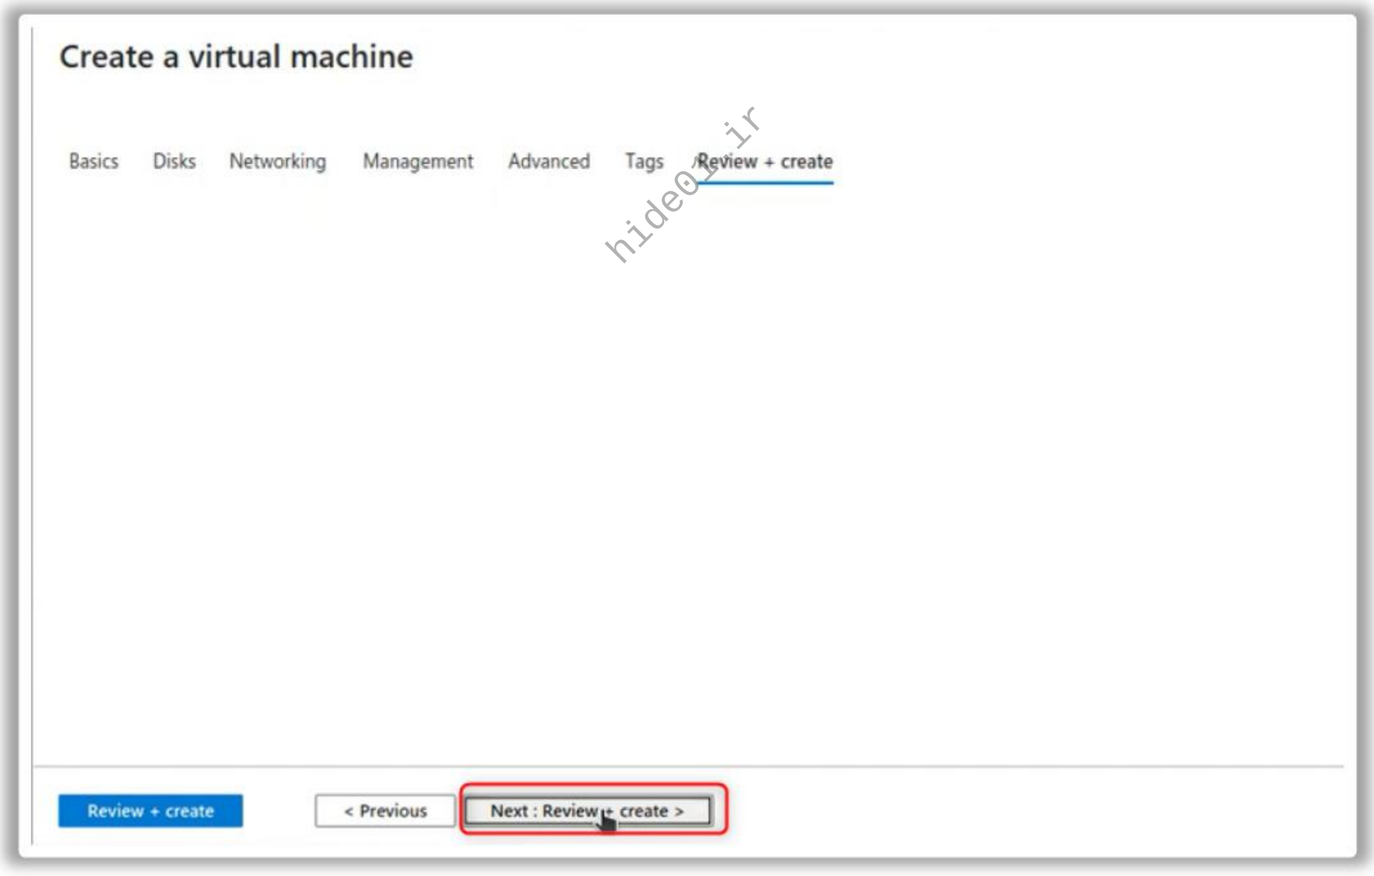


Figure 4.411: Click Review + create

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

- Click on **Create**.

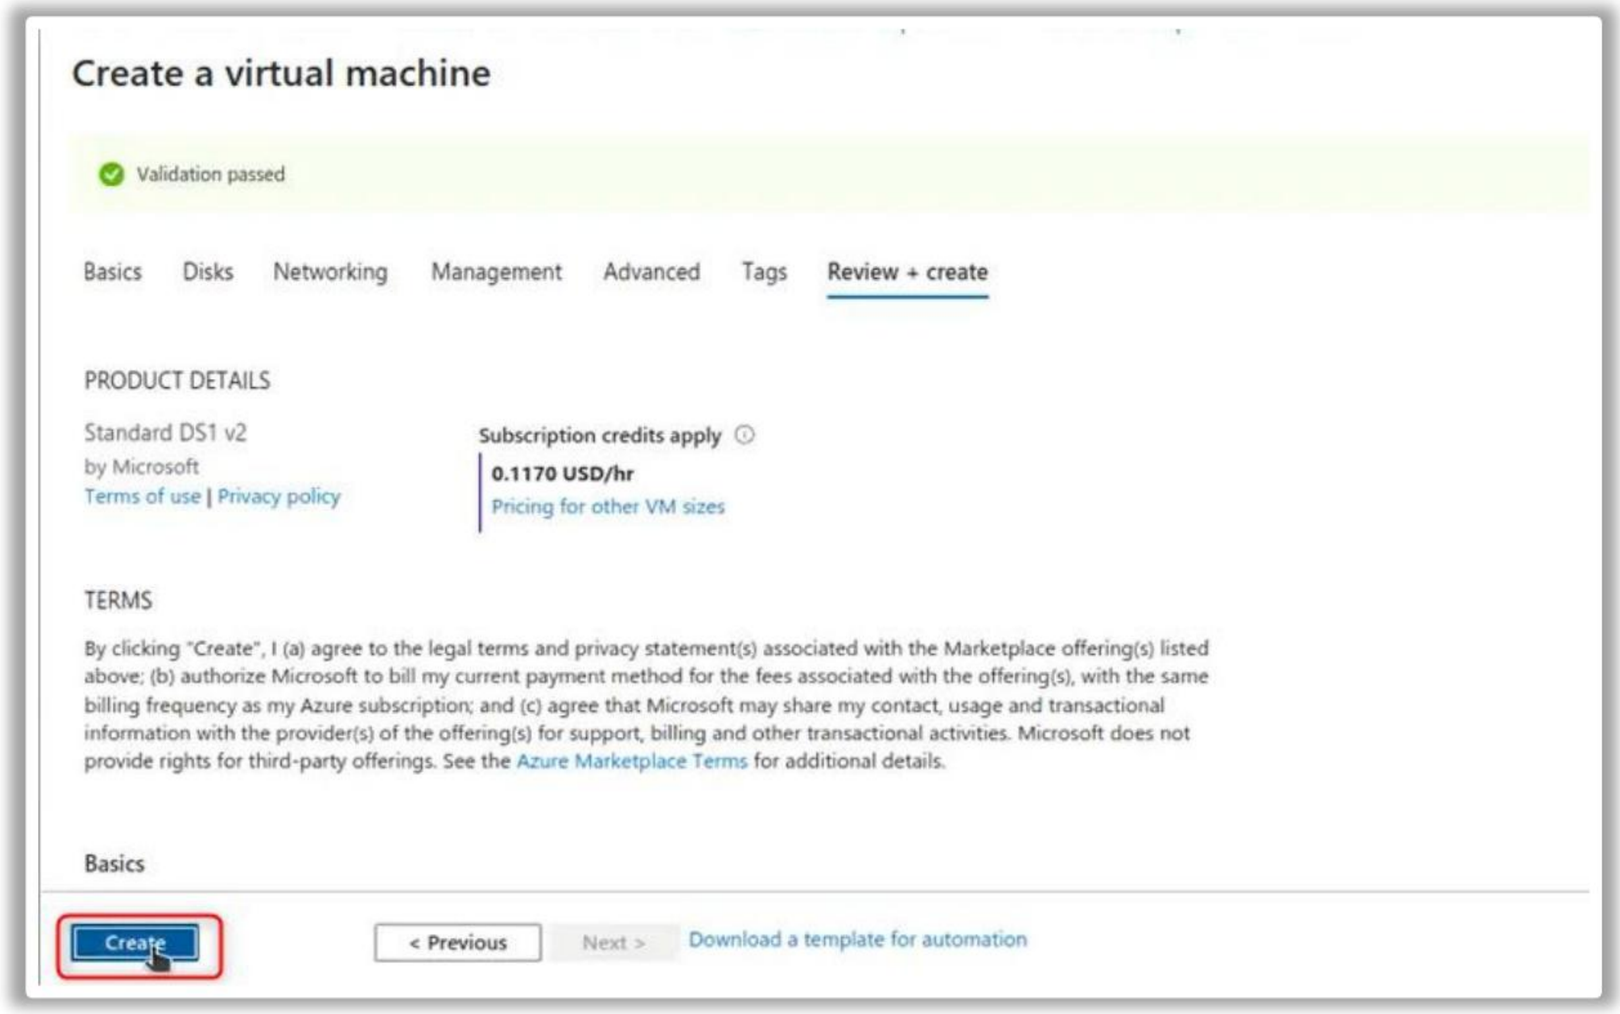


Figure 4.412: Click Create

- The virtual machine has been created successfully.
- Click on **Go to resource**.

hide01.ir

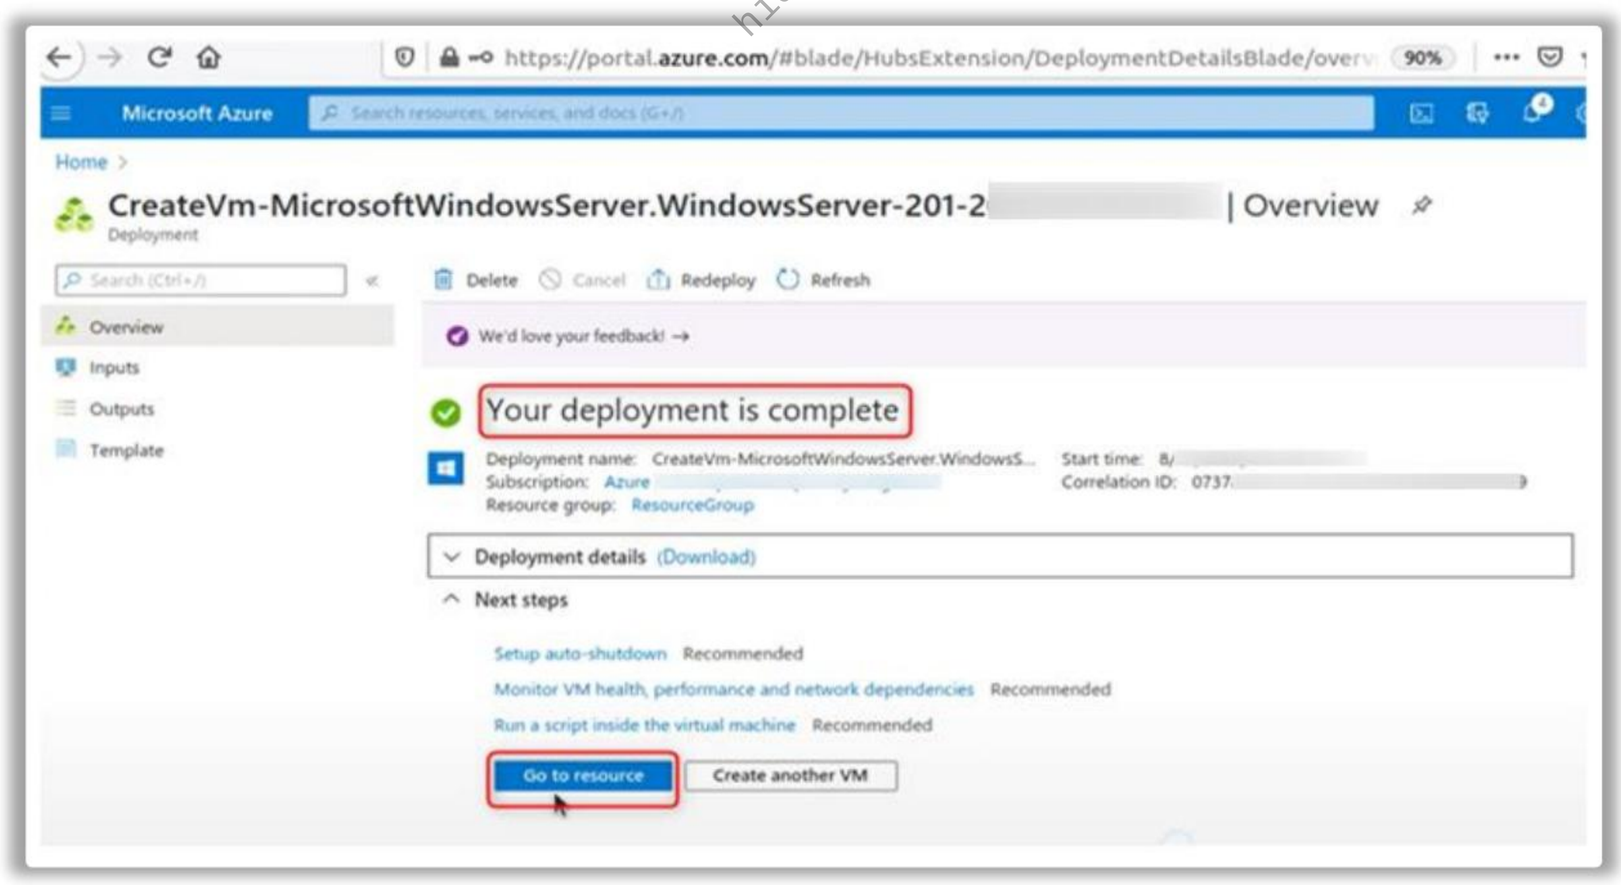


Figure 4.413: Click Go to Resource

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

- Click on **Resource group**.

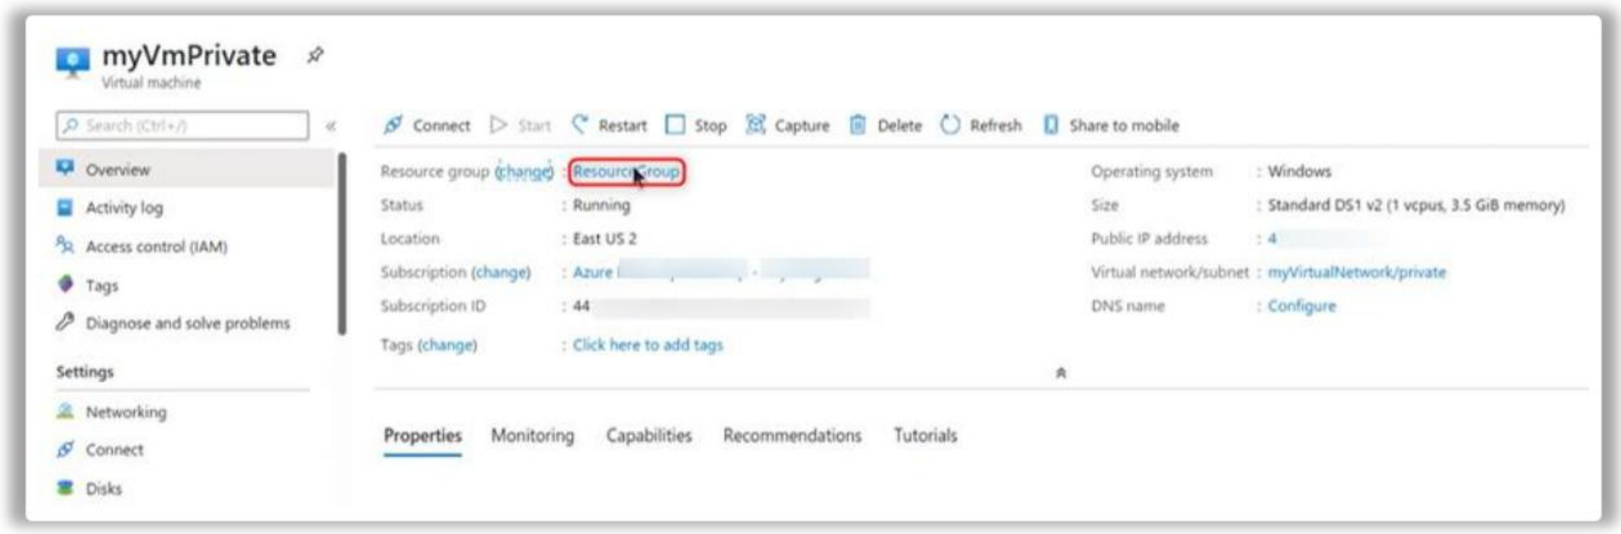


Figure 4.414: Click Resource Groups

- The two virtual machines are now available in the resource group.

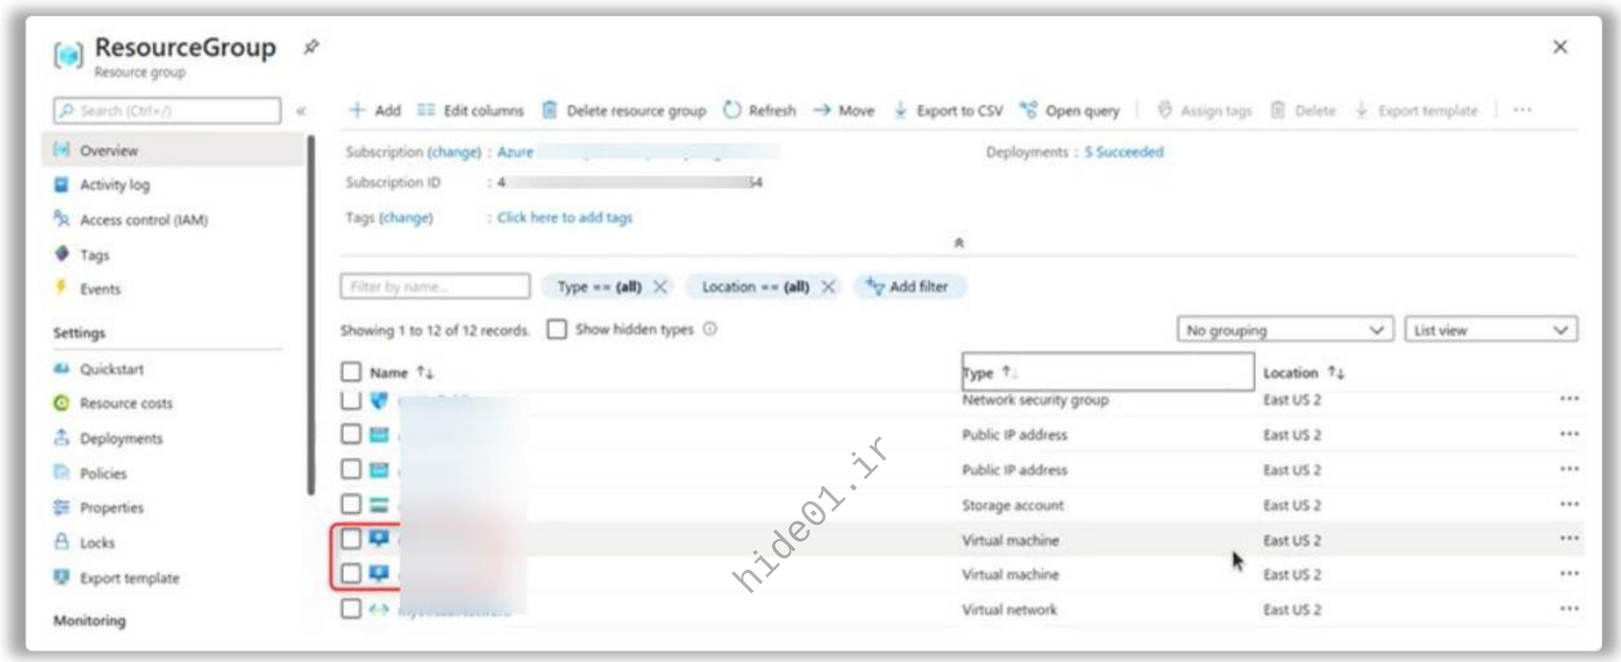


Figure 4.415: list of VM

- In the same list, locate the storage account created previously and click on it.

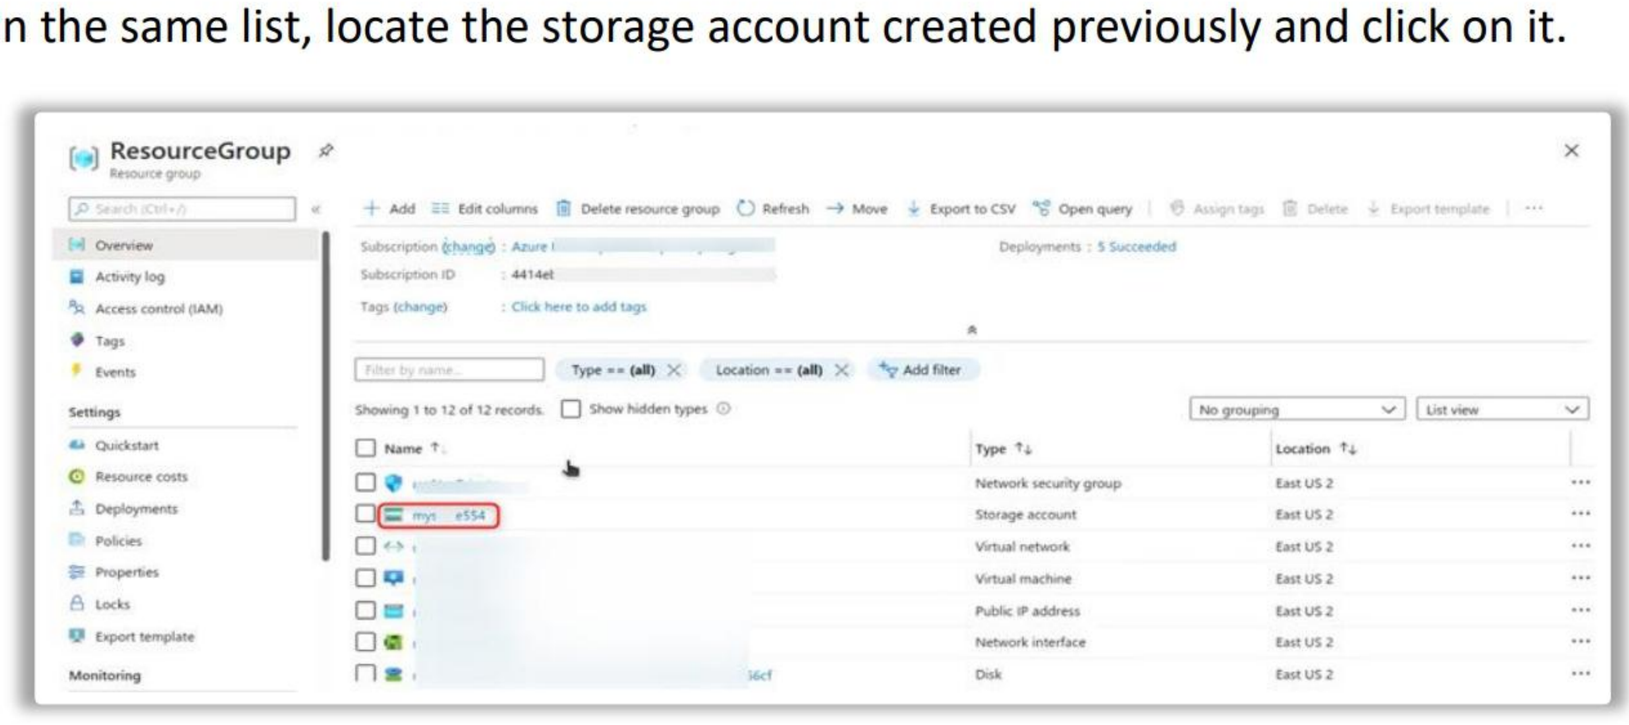


Figure 4.416: Check Location of VM

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

- Click on **File shares**.

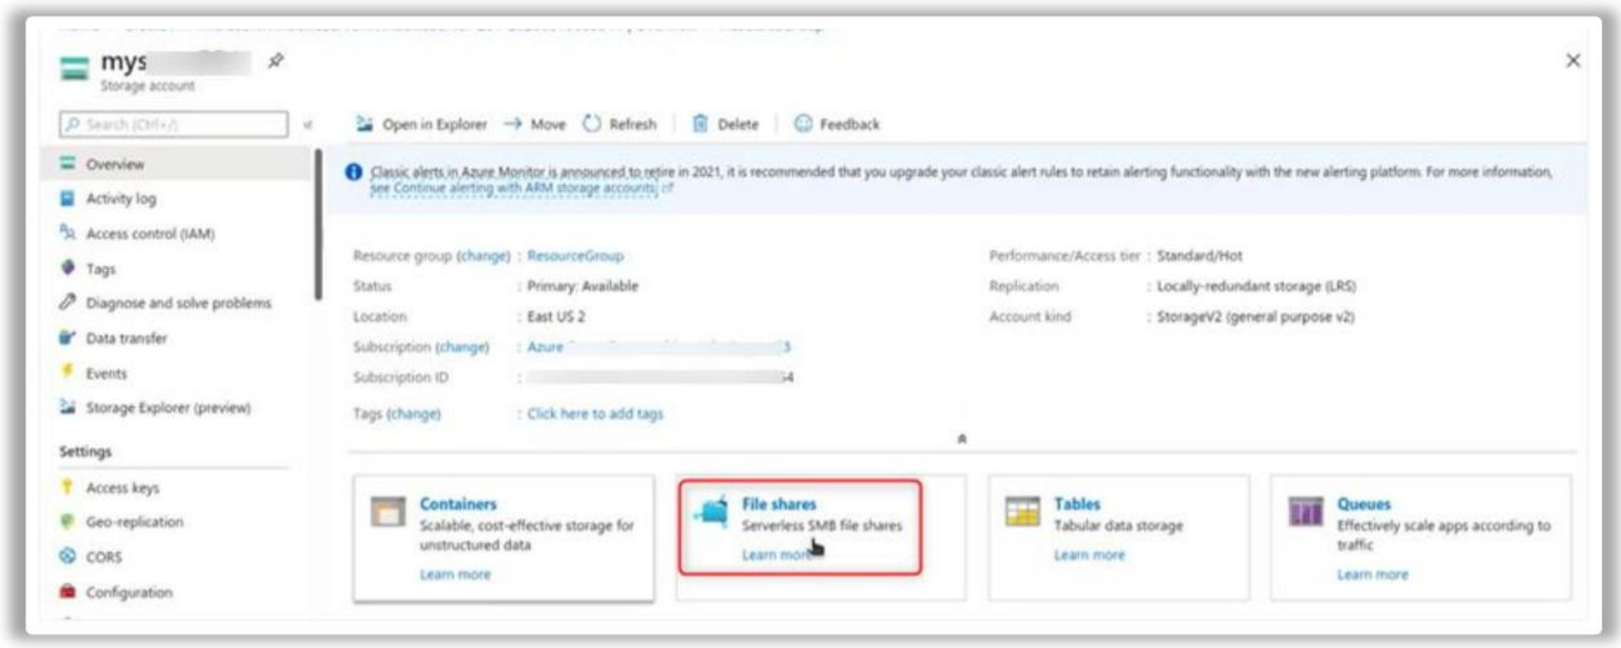
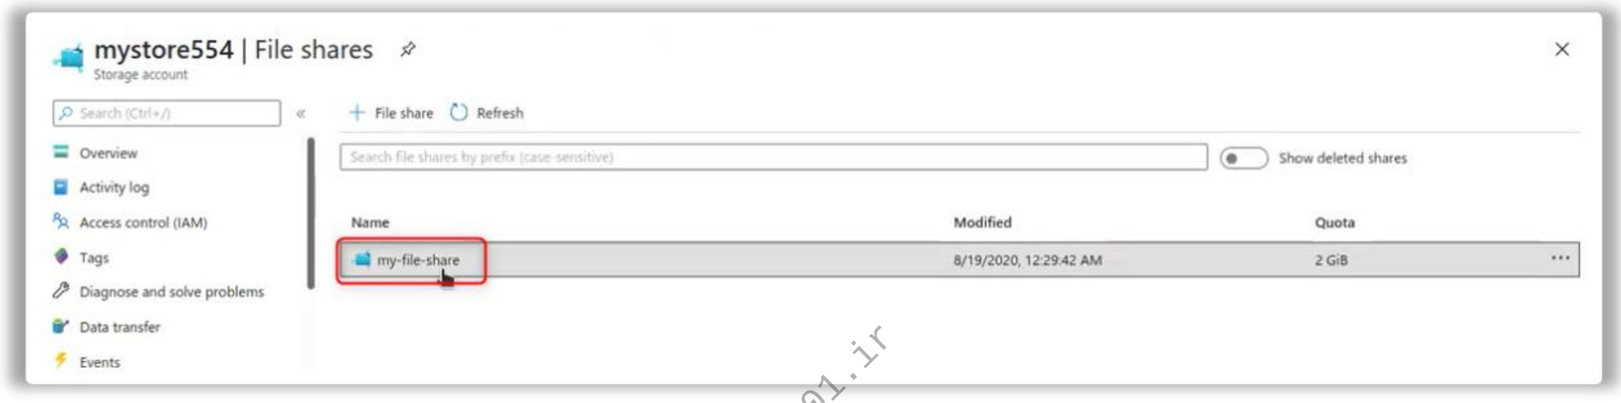


Figure 4.417: Click File Shares

- Click on **my-file-share**.



hide01.ir

Figure 4.418: Click my-file-share

- Access is denied successfully because your computer is not in the *Private* subnet of the *MyVirtualNetwork* virtual network.

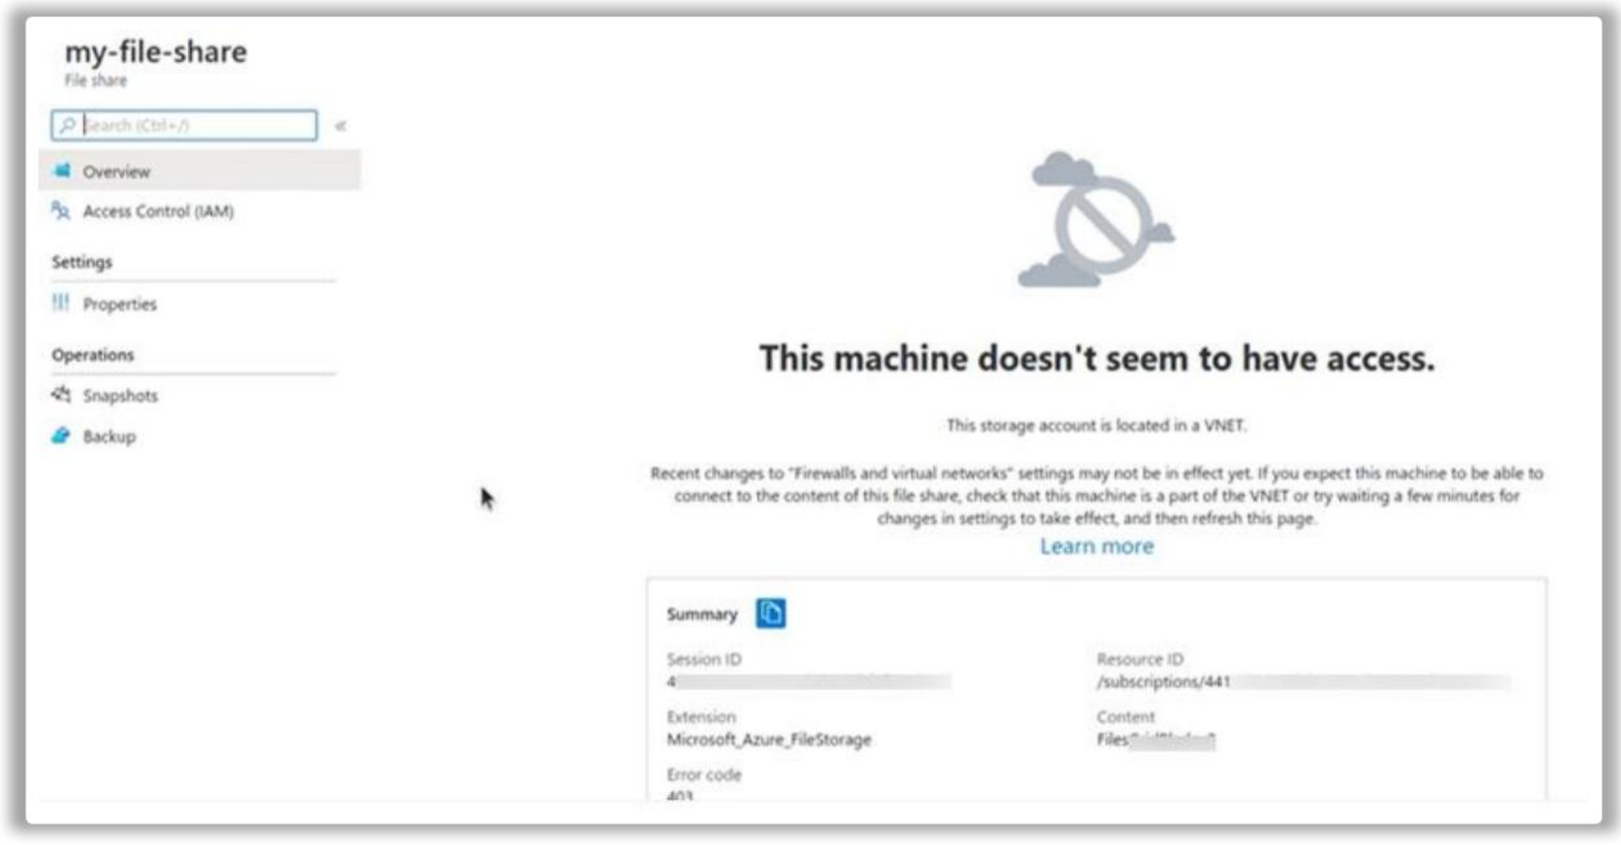
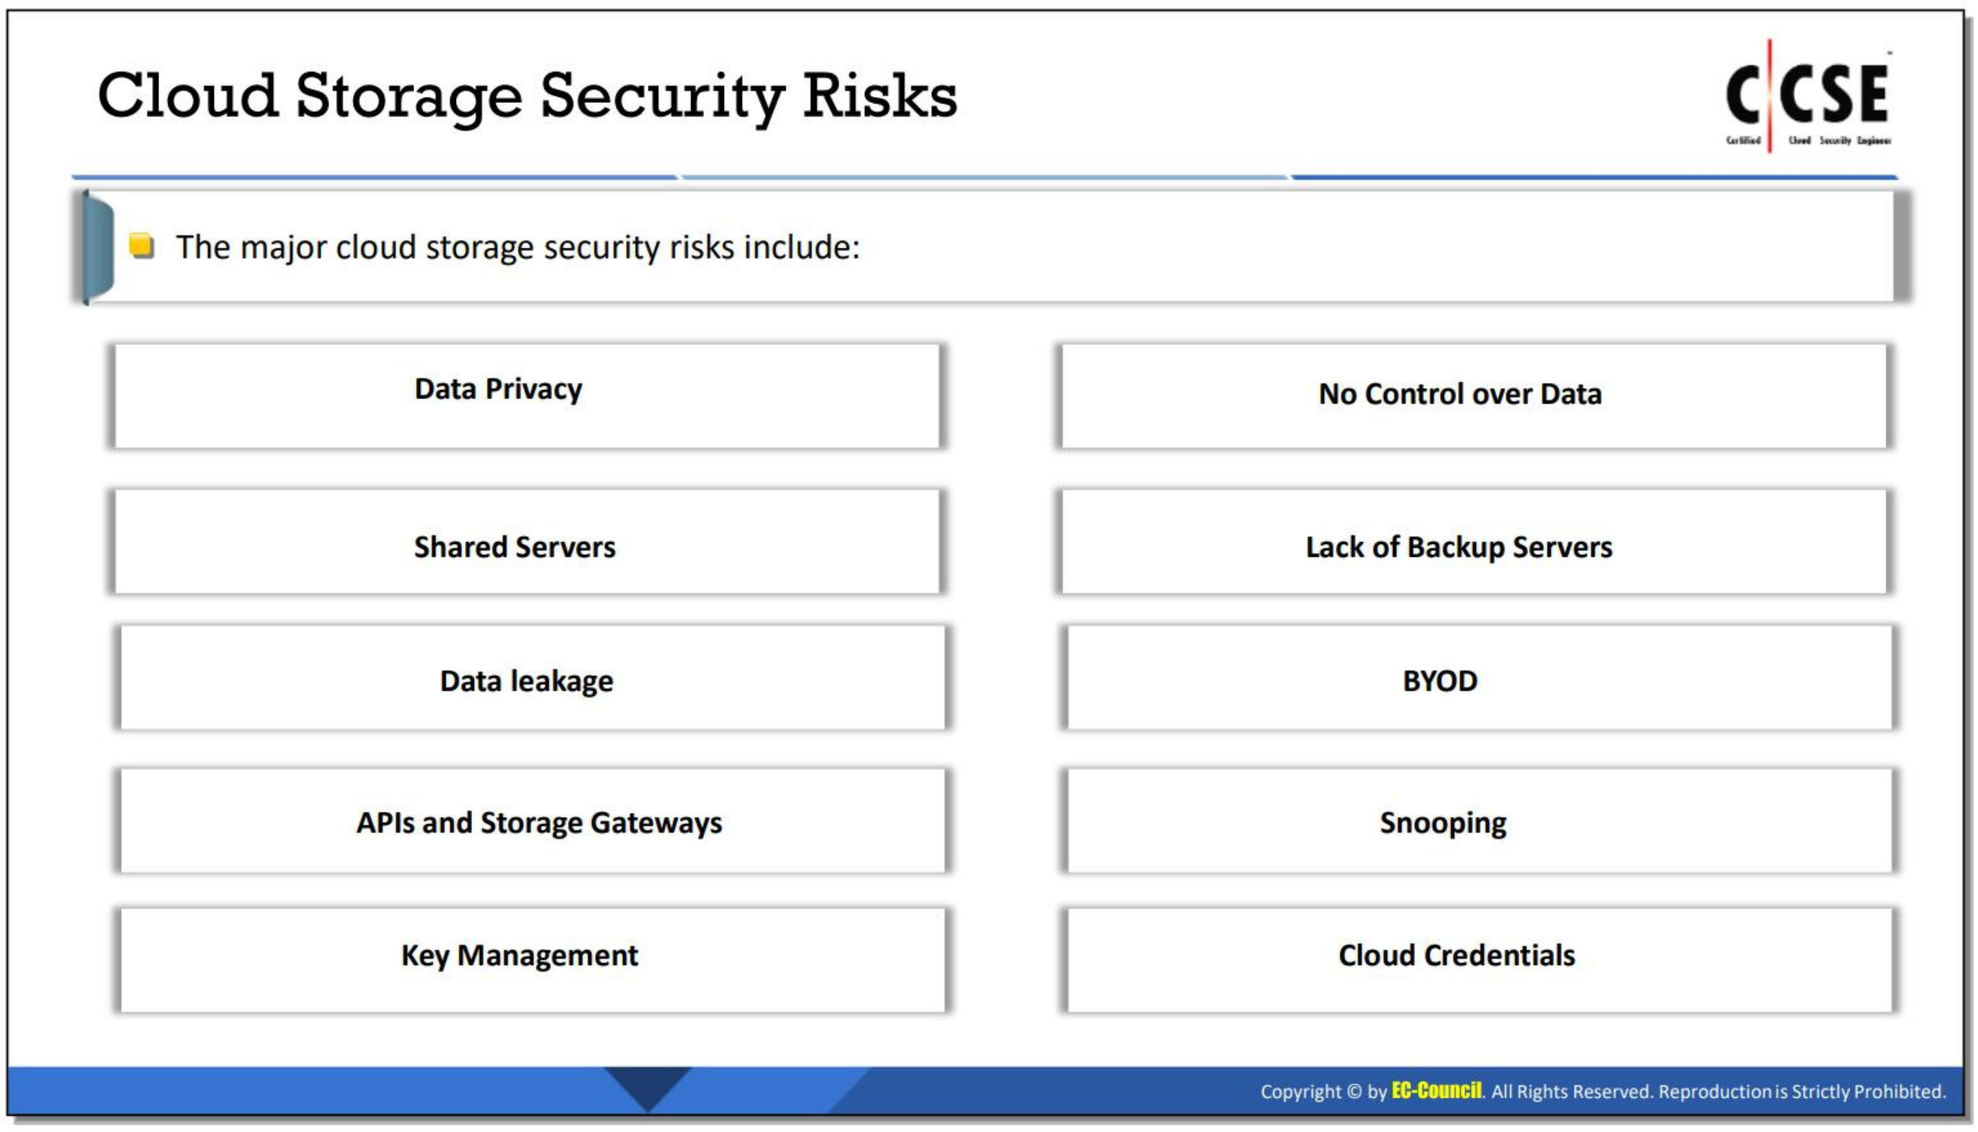


Figure 4.419: Access Notification

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



Cloud Storage Security Risks

Compared to the on-premises storage system, cloud storage is cost-effective and used by most organizations. As most of the employees started using file-sharing/cloud storage applications to share business documents inside the organization and out of the organization, due to lack of enterprise security measures, business information in these third-part applications might be at risk. The major cloud storage security risks are discussed below.

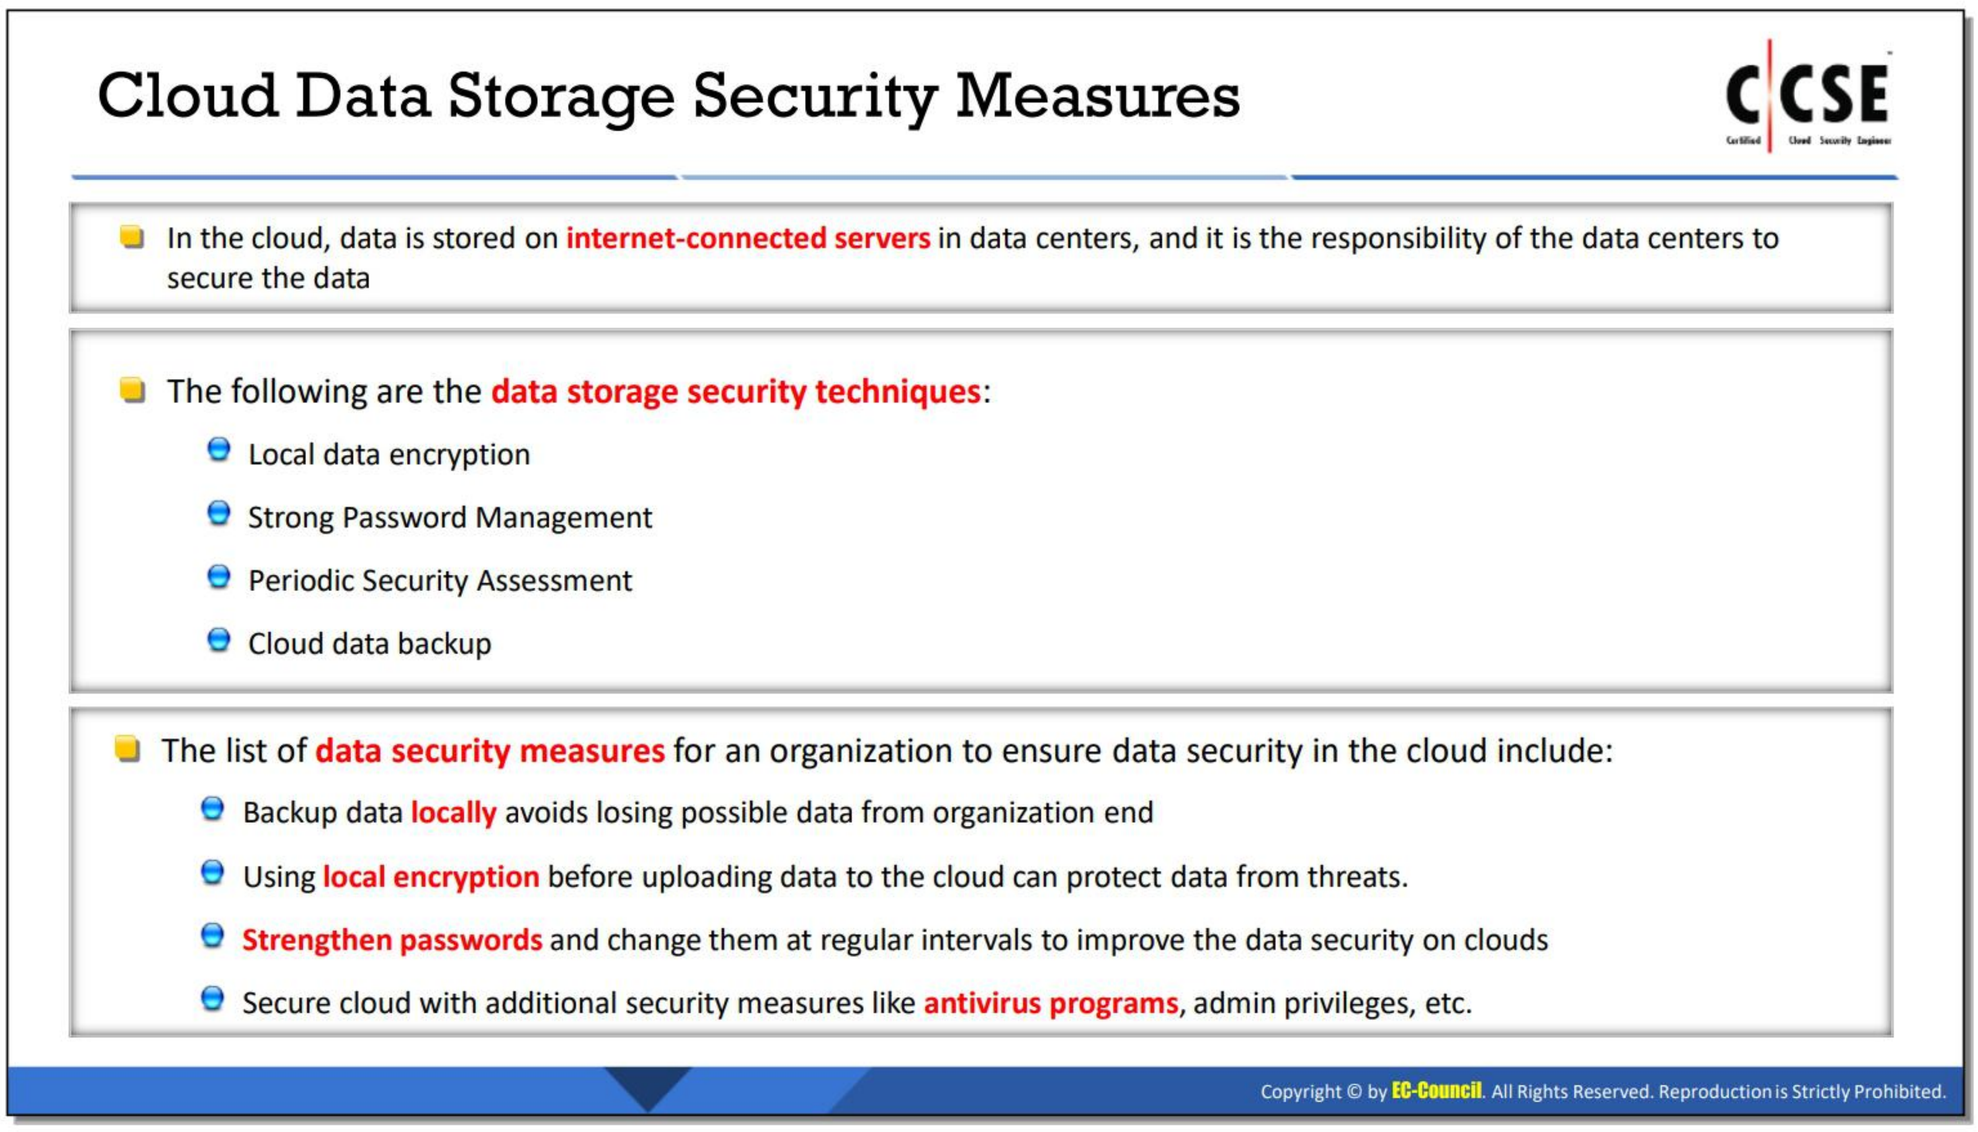
- **Data privacy:** Both physical and virtual factors play a significant role in data privacy and security. Administrators can set the restrictions for accessing the server locations on-premises, but when data is stored in the cloud, administrators are not sure of the privacy provided by the providers. As administrators did not have control over the stored data, organizations need to analyze the risk and implement data privacy controls.
- **No control over data:** Organizations might rely on third-party storage services for storage or backup of on-premises data. Data transmission from organization to third-party cloud storage shows that the data is transmitted out of the organization's IT environment where the data privacy controls set by the company are not efficient. The company must rely on third-party storage providers as long as the data is stored in their servers. There is a high probability that the data can be accessed by an unauthorized individual.
- **Shared servers:** In cloud storage, the cloud consumers share the space with other consumers. If anyone uploads malicious content or files, then there will be a potential risk for the stored data.
- **Lack of backup servers:** Some of the cloud storage providers may provide backup functionality and some others may not. If the cloud provider did not have a backup

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

provision, then there would not be any assurance of data protection in the event of data loss.

- **Data leakage:** As cloud storage is a multi-user environment where different administrators store different kinds of data, there is a high likelihood that the adversary can also store malicious data in the cloud storage and involve in data leakage. Apart from this, cloud storage providers may accidentally expose data to unauthorized persons due to misconfiguration of access controls that would lead to data leakage.
- **BYOD:** Bring Your Own Device (BYOD) refers to a policy that allows an employee to bring their devices such as laptops, smartphones, and tablets at the workplace and use them for accessing the organization's resources as per their access privileges. This may increase the security risk as these personal devices are having access to the organization's paid third-party cloud storage. People using personal devices with administrator access can download and install any software or files, which may be malicious and may execute its malicious nature if security configurations are not enforced properly.
- **APIs and storage gateways:** Organizations at the time of migrating the data from on-premises to the cloud storage may use cloud storage APIs or storage gateways. Due to insecure APIs, there is a high probability of data leakage at the time of transmission.
- **Snooping**: Unauthorized monitoring of traffic between two machines on a network or theft of information during data transmission is known as snooping. An attacker may use snooping attack to extract the information while the data is being sent to cloud storage.
- **Key management:** Key management is one of the crucial factors that decide the security of data stored in cloud storage. If the cryptographic keys are not handled properly, even though the data is encrypted with strong encryption, it would cause enormous damage.
- **Cloud credentials:** The administrators should maintain the cloud account credentials securely. As the compromise of the cloud account credentials not only give access to all the storage accounts, but also allows replication and deletion of data.

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



Cloud Data Storage Security Measures

In the cloud, data is stored on internet-connected servers in data centers and it is the responsibility of the data centers to secure the data. However, customer should protect their data at their end to ensure comprehensive data security.

hits01.ir

Data Storage Security Techniques

- **Local data encryption**: Encryption of local data before uploading it to the cloud.
- **Strong password management**: Using strong passwords and changing them at regular intervals to prevent access to cloud data by unauthorized users.
- **Periodic security assessment**: Implement data security controls that must be continuously monitored and reviewed.
- **Cloud data backup**: Taking backup of cloud data locally avoids losing possible data from the organization's end.

List of data security measures for an organization to ensure data security in the cloud:

- Generally, IT organizations depend on data a lot. Losing data may not limit to financial loss but also legal actions. Backup data locally avoids losing possible data from the organization's end.
- Compromise with storing sensitive information (patents and copyrights) on clouds may create troubles for organizations. Staying away from saving sensitive information on clouds and saving sensitive information are bigger responsibilities for organizations.
- Using local encryption before uploading data to the cloud can protect data from threats. It is better to opt for a service provider who can provide prerequisite data encryption.

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

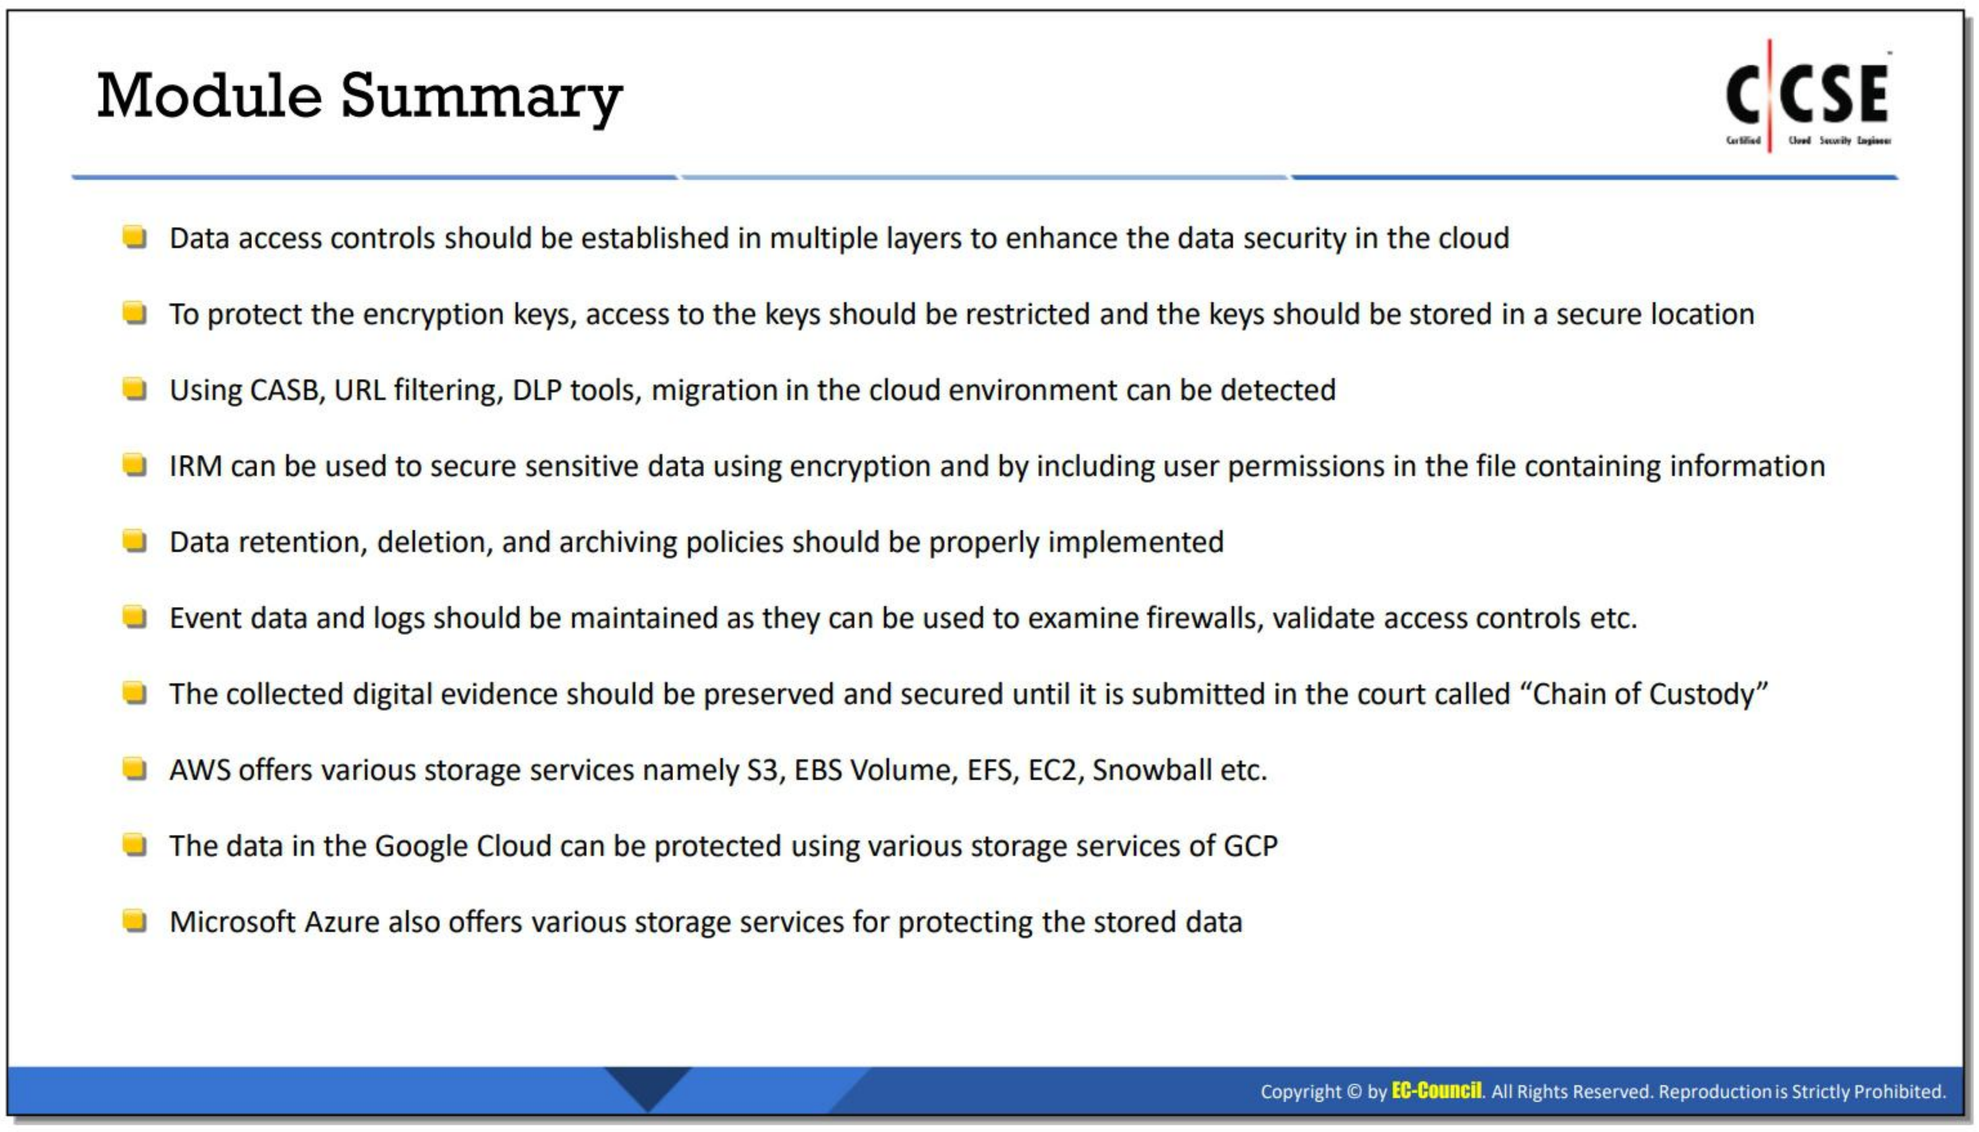
Have a preliminary round of encryption for data that can provide consumers extra security if they already have an encrypted cloud service.

- Strengthen passwords and change them at regular intervals to improve the data security on clouds. Having a two-step verification process and updated patches enhance data security.
- In addition to encrypting data and applying passwords to secure data on clouds, securing cloud with additional security measures like antivirus programs, administrator privileges and local encryption offered by cloud service can add some more protection to data.

Test cloud data security to know how well it is performing. Hire ethical hackers to test the system's security levels whether they have declined over time. This step may also pose attacks from unknown sources through loopholes. Keeping the data secured on the cloud requires constant action.

hide01.ir

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



Module Summary

The module discussed the technical aspects involved in cloud data security. The topics discussed in this module included cloud storage types, storage techniques, data lifecycle phases, cloud storage encryption tools, storage services in AWS, GCP, and Azure.

hi2ol.ir

It highlighted the following points:

- Data access controls should be established in multiple layers to enhance the data security in the cloud.
- To protect the encryption keys, access to them should be restricted and they should be stored in a secure location.
- Using CASB, URL filtering, DLP tools, migration in the cloud environment can be detected.
- IRM can be used to secure sensitive data using encryption and by including user permissions in the file containing information.
- Data-retention, deletion, and archiving policies should be properly implemented.
- Event data and logs should be maintained as they can be used to examine firewalls, validate access controls etc.
- The collected digital evidence should be preserved and secured until it is submitted in the court; this is called ”Chain of Custody.”
- AWS offers various storage services, namely S3, EBS volume, EFS, EC2, Snowball etc.
- The data in Google Cloud can be protected using various storage services of GCP.
- Microsoft Azure also offers various storage services for protecting data stored in the cloud.

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