EC-Council

C|CSE™

Certified Cloud Security Engineer



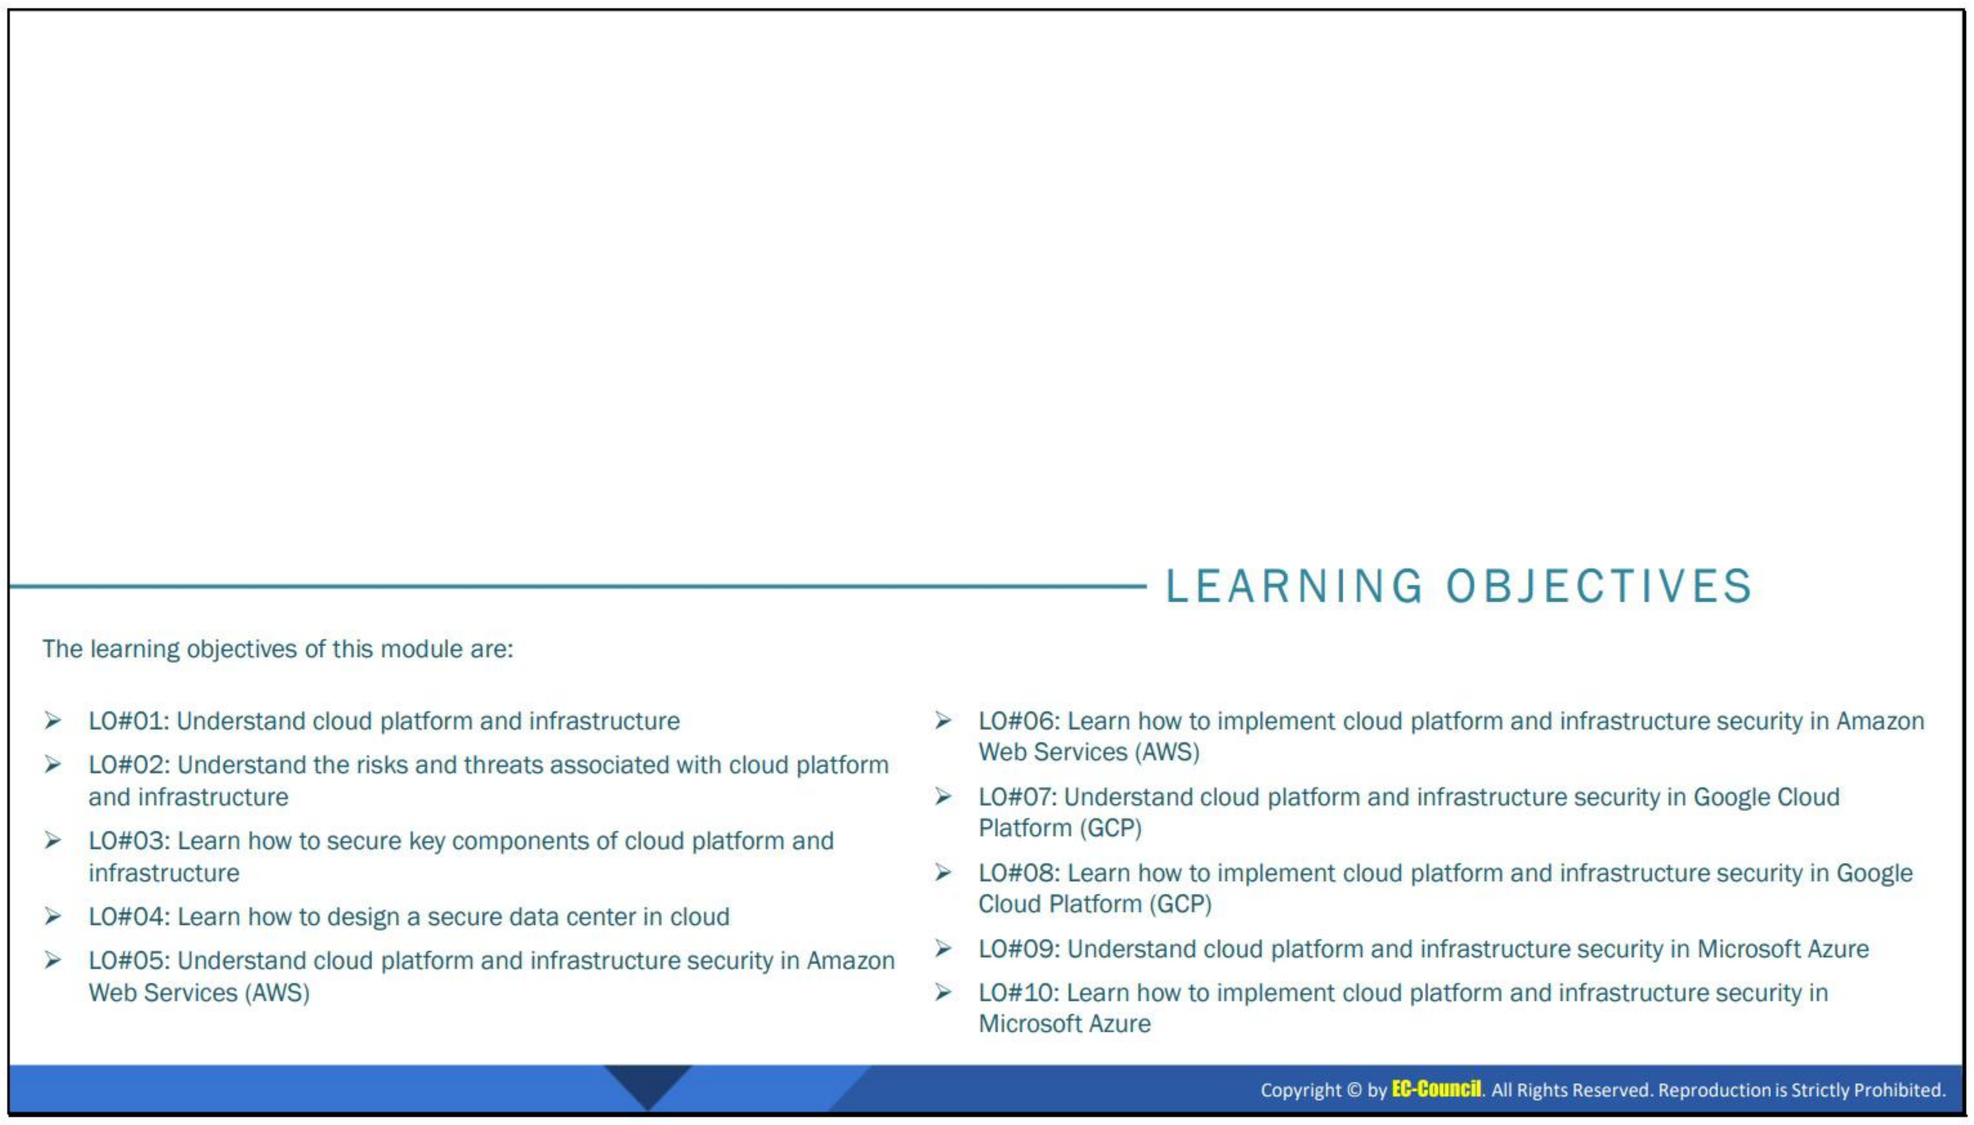
hide01.ir

Platform and Infrastructure Security in Cloud

Module 02

This page is intentionally left blank.

hide01.ir

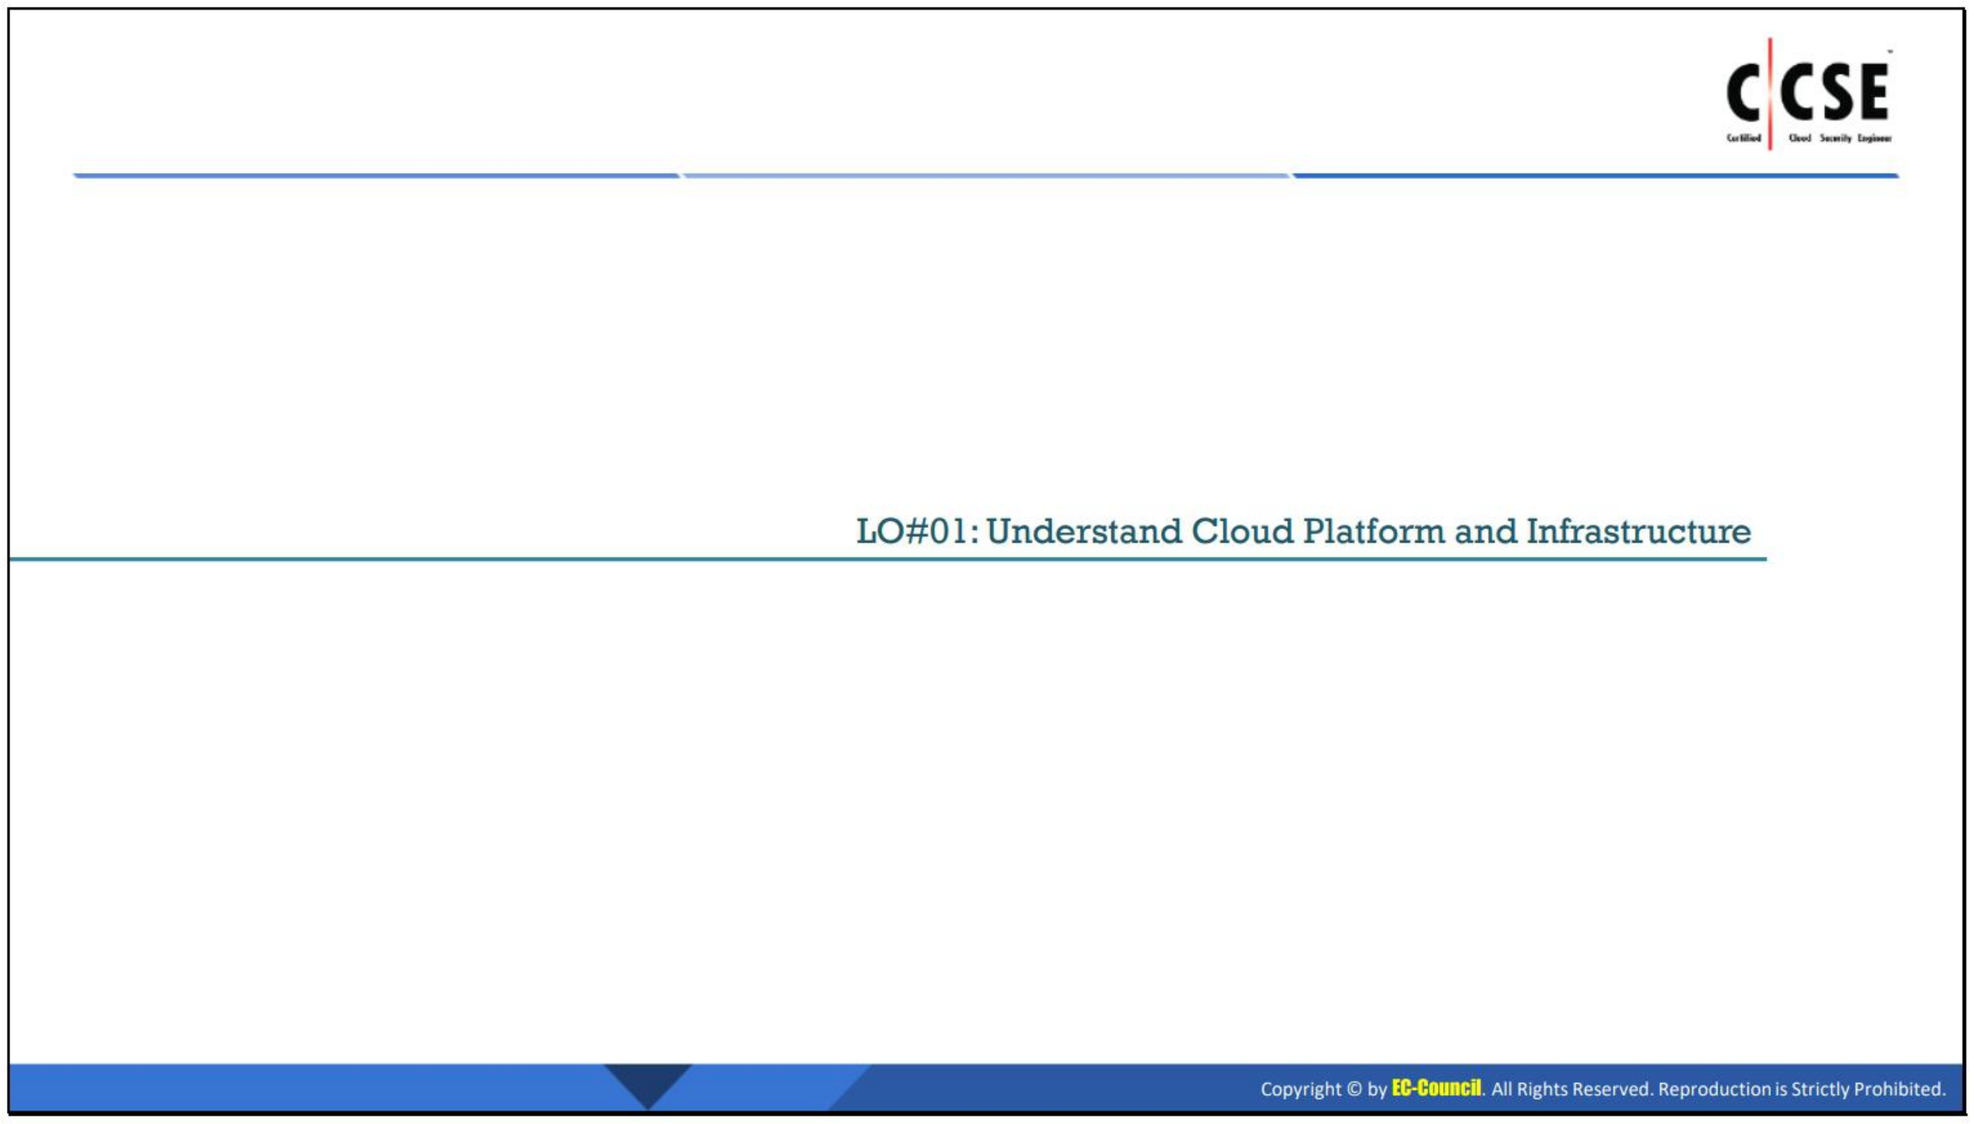


Learning Objectives

Securing the cloud infrastructure is critical for carrying out various operations in the cloud securely. Apart from the computers and networks, it is also important to know how to secure the workloads and storage. This module focuses on the considerations for securing the various components of a data center. It demonstrates the usage of various security features provided by major cloud service providers (Amazon Web Services (AWS), Microsoft Azure, and Google Cloud Platform (GCP)). The learning objectives of this module are as follows:

- Understand cloud platform and infrastructure
- Understand the risks and threats associated with cloud platform and infrastructure
- Learn how to secure key components of cloud platform and infrastructure
- Learn how to design a secure data center in cloud
- Understand cloud platform and infrastructure security in AWS
- Learn how to implement cloud platform and infrastructure security in AWS
- Understand cloud platform and infrastructure security in GCP
- Learn how to implement cloud platform and infrastructure security in GCP
- Understand cloud platform and infrastructure security in Microsoft Azure
- Learn how to implement cloud platform and infrastructure security in Microsoft Azure

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

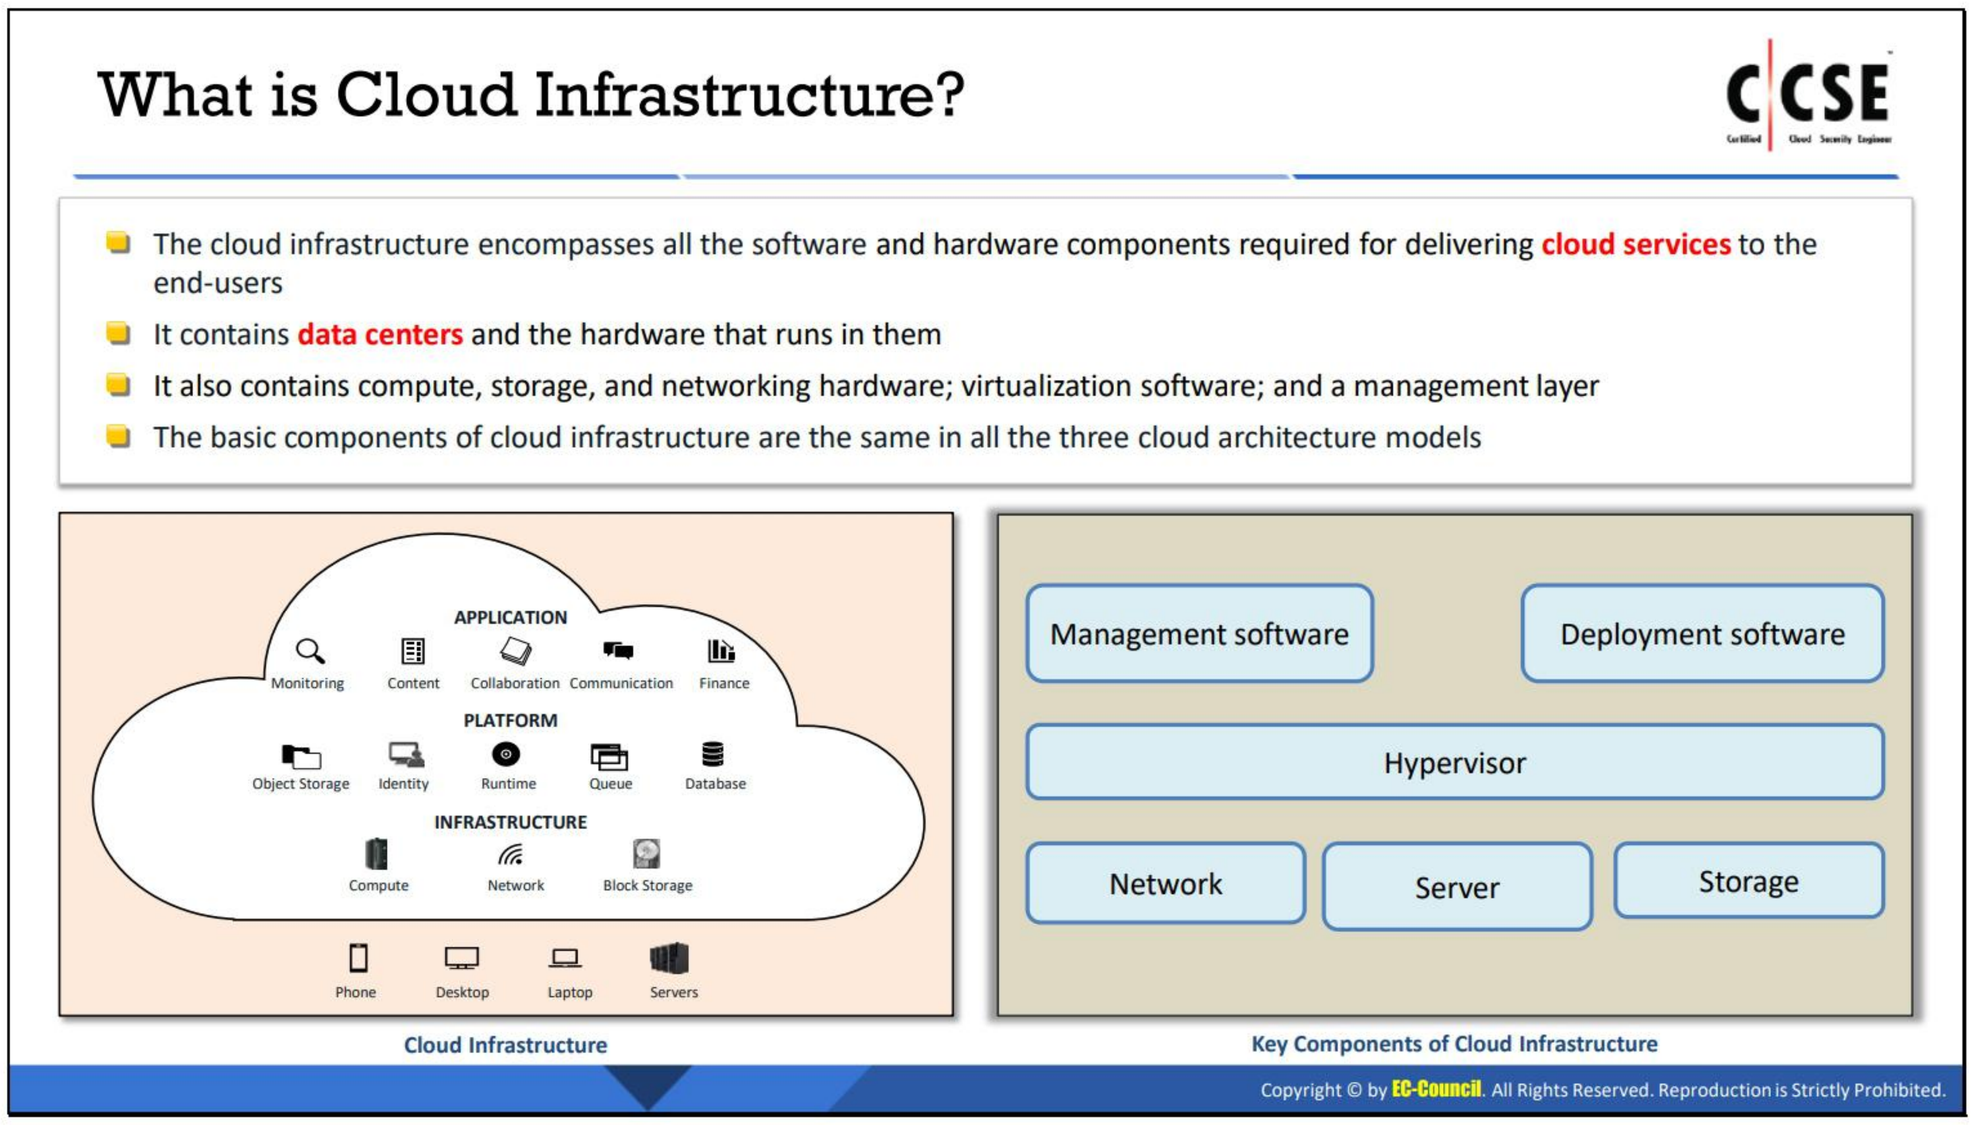


LO#01: Understand Cloud Platform and Infrastructure

The objective of this section is to understand the various components of the cloud platform and infrastructure, which will help you secure cloud storage from cyberattacks and unknown users.

hide01.ir

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



What is Cloud Infrastructure?

Cloud computing provides the facilities for an organization to access the stored data and computing capabilities that they require as per their needs without a significant upfront cost. Instead of establishing new on-premises IT infrastructure, an organization can rent the cloud infrastructure and related capabilities from cloud service providers (CSP) such as Amazon Web Services (AWS), Microsoft Azure, Google Cloud Platform (GCP), etc.

The term "cloud infrastructure" encompasses all the software and hardware components that are required to deliver cloud services to the end-users. Networking equipment, servers, and data storage are the major physical components of cloud infrastructure. It also consists of a hardware abstraction layer that facilitates virtualization of resources.

Cloud Infrastructure in Three Cloud Architectures

The basic components of cloud infrastructure are the same in all the three cloud architecture models for delivering computing services.

- **Private Cloud**: In this model, only a single organization can access cloud infrastructure. The on-site IT staff can develop and maintain the private cloud architecture, or an external service provider can deliver it.
- **Public Cloud**: In this model, a CSP offers cloud resources to paying customers through the internet. The public CSP utilizes a multitenant environment, which reduces the upfront cost of computing resources; however, it can create privacy concerns for companies that store sensitive information.
- **Hybrid Cloud**: In this model, the public and private cloud will interact with each other in a discrete but connected system. By using a hybrid cloud, an organization storing sensitive

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

data can safeguard the data's privacy by storing sensitive data in on-site servers and hosting the rest of the information in the public cloud to lower costs.

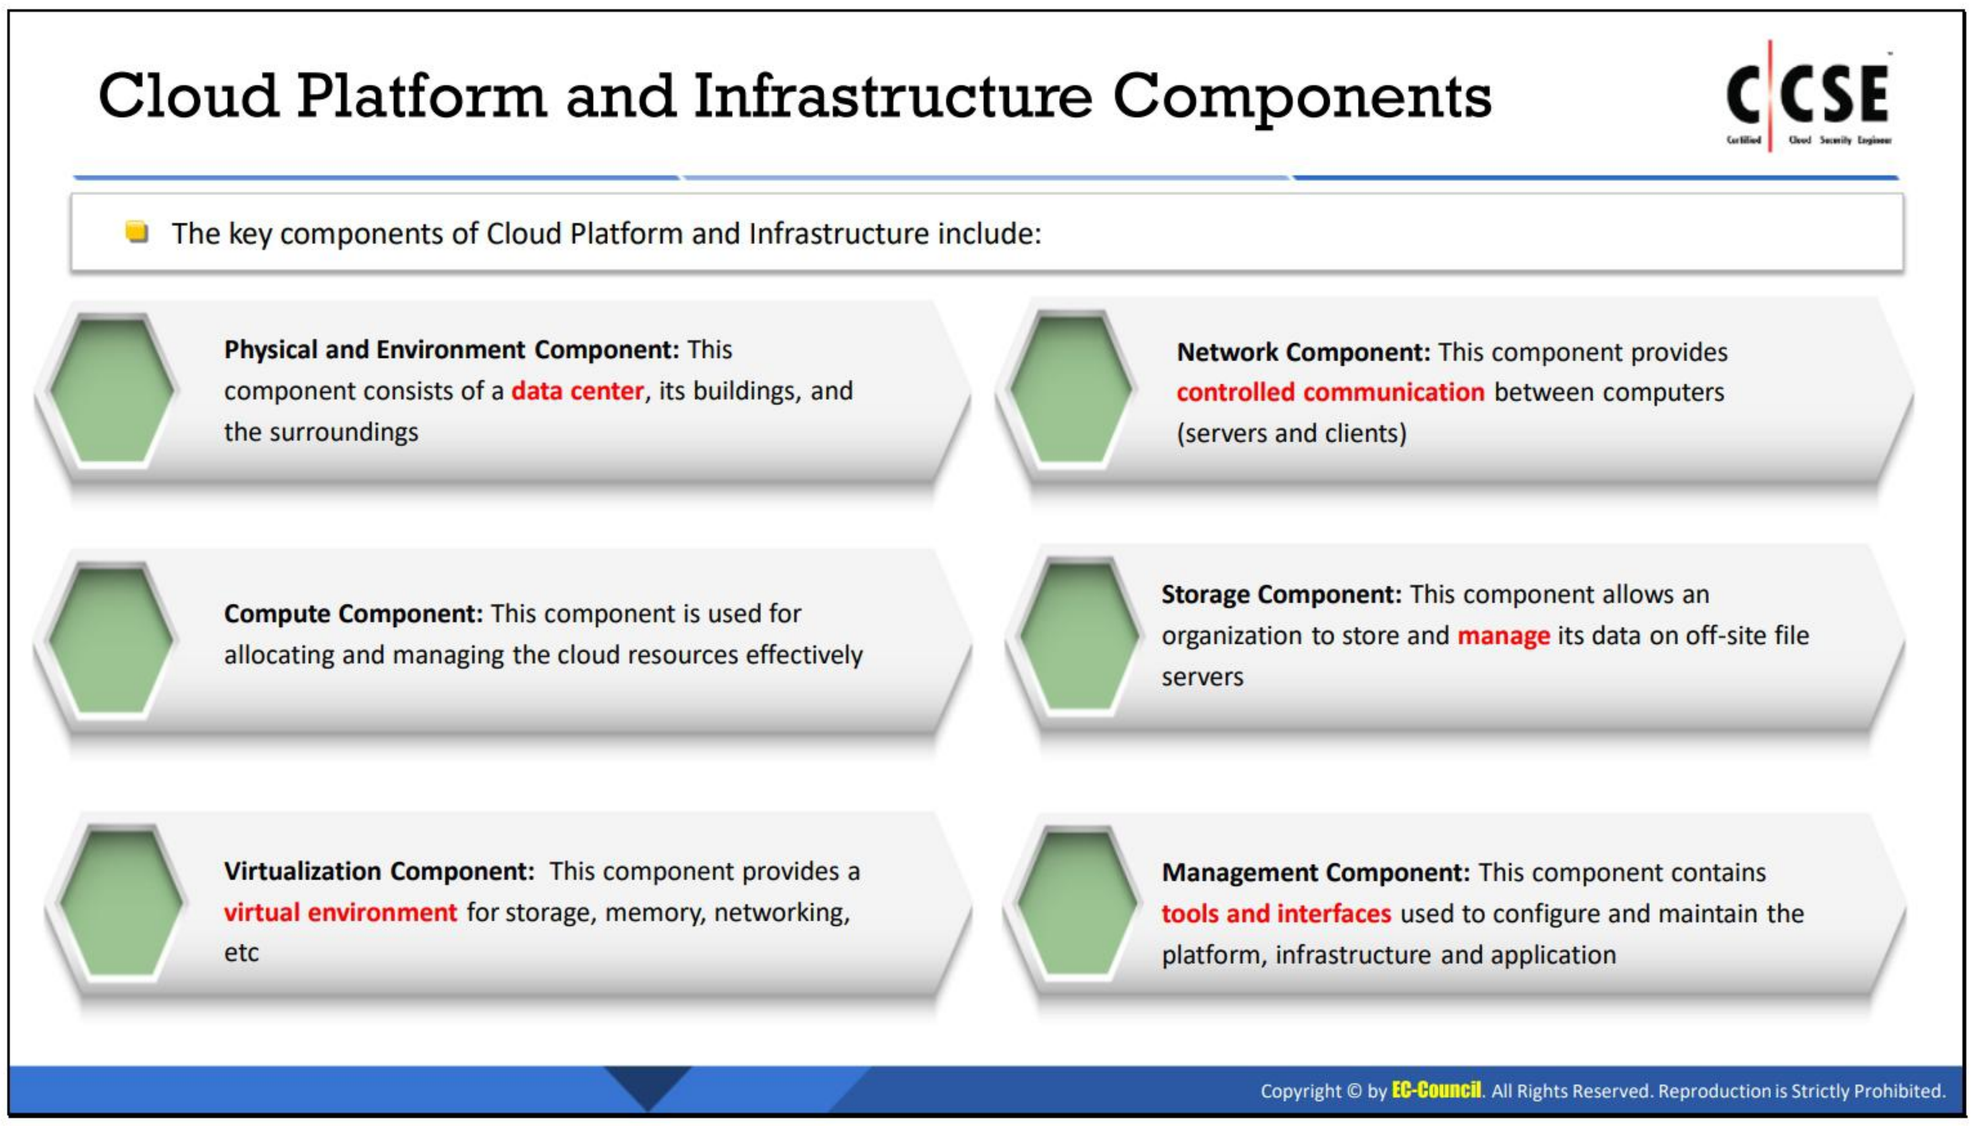
Cloud Infrastructure Delivery Models

Public CSPs deliver cloud infrastructure and related services through the below models.

- **Infrastructure-as-a-Service (IaaS):** In this model, the CSP offers networking, servers, data storage, and virtualization capabilities. The customer buys access to the data storage and compute power as per their requirements but provides their own software platform to run on it such as OS, middleware, data, runtime, and applications.
- **Platform-as-a-Service (PaaS):** In this model, the CSP provides the complete cloud infrastructure (servers, data, networks, virtualization) as well as the software platforms, including the OS, middleware, and runtime. This model offers a platform for customers to develop and test all types of cloud-based applications and lowers costs.
- **Software-as-a-Service (SaaS):** In this model, the CSP offers applications from a web-based portal. There is no need for the customer to store any application data on the local hard disk; all the data is stored on the server of the service provider.

hide01.ir

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



Cloud Platform and Infrastructure Components

When an organization purchases cloud computing services under the IaaS/PaaS model, then the organization is renting the data storage, compute power, virtualization capabilities, etc. from the CSP.

hide01.ir

The key components of cloud platform and infrastructure are as follows:

- Physical and environment component
- Network component
- Compute component
- Storage component
- Virtualization component
- Management component

Every component listed above plays a significant role in deploying and delivering applications and other services.

Many organizations that have used public cloud services have encountered data theft and advanced attacks against their cloud platform and infrastructure. The most common cloud platform and infrastructure threats include data breaches, distributed denial of service (DDoS), malicious insiders, account hijacking, etc. Hence, it is incumbent to implement robust security practices on all the components of a cloud service.

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

Physical and Environment Component

Except an on-premises private cloud, the cloud platform and infrastructure run in a third-party data center. The characteristics of the physical and environmental components include the following:

- A single data center runs hundreds to thousands of servers.
- The data center owners deliver service at multiple levels. The basic level service includes pipe, power, and ping.
- Electrical power and cooling pipes are provided in the data center. The power and pipes restrict the density of the stacked servers in the data center.
- The power density is measured in kW per rack and a data center can accommodate 25 racks per 100 m^2. The standard power density is 100 W per rack, but nowadays power densities of 10 kW or more per rack can also be observed. Advanced cooling engineering should be applied to these densities.
- Datacenter networks should have network connectivity for accessing the storage. To access the wide area network (WAN) resources, external connectivity should be provided.
- The data center should provide floor space, rack space, cages on any level of aggregation.
- The physical environment of the data center should be checked for political and geographical risks such as floods, availability of power, accessibility, seismic activity, etc.

Data Center Design:

A data center design focuses more on redundancy in the design. Anything that can be broken down should be replicated. There should not be any single point of failure that is backup power, multiple independent cooling units, multiple power lines to individual racks and servers, multiple power distribution units (PDUs), multiple entrances to the building, multiple external entry points for power and network, and so on.

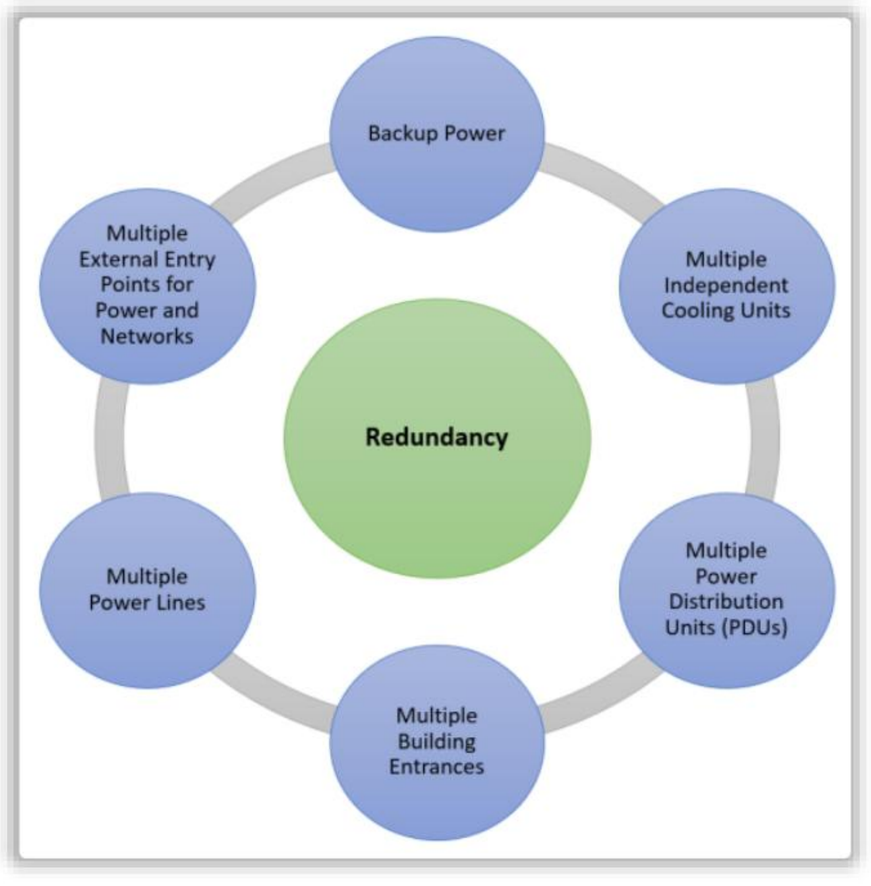


Figure: 2.1: Redundancy of a Data Center Design

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

Network Component

The network is an important component of the cloud infrastructure and its purpose is to provide controlled communication between computers (servers and clients). The users can access the cloud resources through the internet; therefore, the third-party service providers are required to build and maintain the networking infrastructure. The networking infrastructure consists of physical wiring, switches, load balancers, routers, etc.

As per the National Institute of Standards and Technology's (NIST's) "Cloud Computing Synopsis and Recommendations," the key first-level terms are defined as follows:

- **Cloud Service Consumer**: An individual or a company that establishes a business relationship and utilizes the services from the CSPs.
- **Cloud Service Provider (CSP)**: An individual, company, or entity that provides cloud-based services to the cloud service consumer.
- **Cloud Carrier**: It offers the connection and transport of cloud services between the CSP and the cloud user.

In the NIST cloud computing reference model, the cloud carrier is responsible for network and communication functions. This is an IP service and delivered through IPv4 and IPv6.

Network Functionality: This consists of the following:

- **Address allocation:** It is the ability to allocate one or more IP addresses to the cloud resources.
- **Access control:** It is the process of either allowing or disallowing access to a cloud resource.
- **Bandwidth allocation:** It is the ability to allocate a particular bandwidth for system access/use.
- **Rate limiting:** It is the ability to regulate the amount of network traffic received or sent. It is utilized to control the number of API requests made in a particular period.
- **Filtering:** It is the ability to selectively allow or disallow access to the content/resources.
- **Routing:** It is the ability to direct the flow of traffic between endpoints by choosing the best path.

Software-defined Networking: Software-defined networking (SDN) is a networking framework that centralizes the network controller by separating a network's control functions from its data packet forwarding functions. Thus, it centralizes the intelligence and reduces the complexity of the traditional network architecture from applications and services. CSPs are switching to SDN technology to efficiently manage the entire network and devices.

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

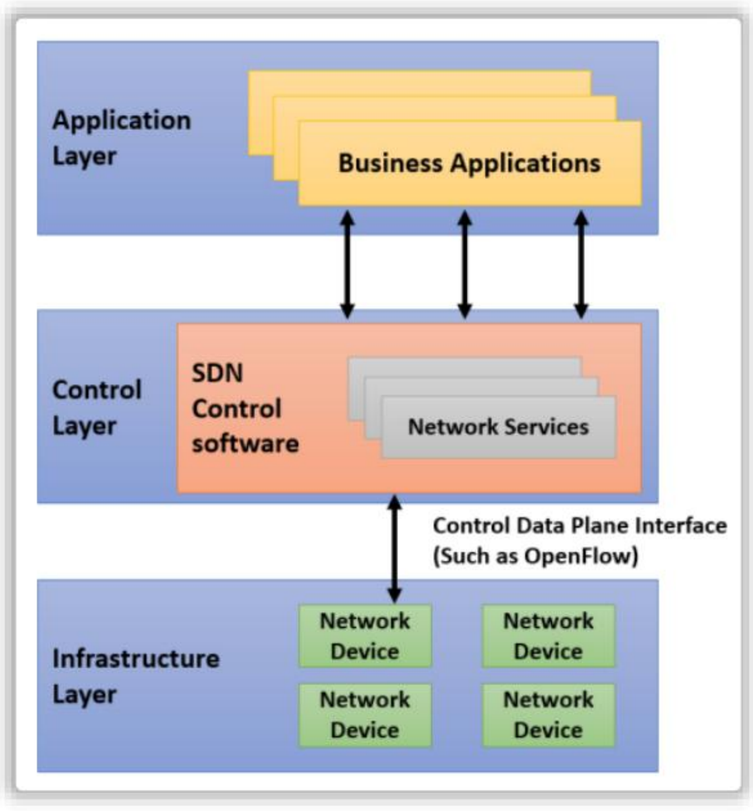


Figure: 2.2: Software-defined Networking

SDN aims to provide a discrete and clearly defined network control plane for managing the network traffic that is separated from the forwarding plane. With this method, network control becomes directly programmable and separated from the forwarding plane, and dynamically adjust the flow of the traffic to address the changing pattern of consumption. SDN allows the user to run the control plane software on general-purpose hardware, permitting the decoupling from particular network hardware configurations and allowing more efficient utilization of commodity servers.

hide01.ir

Compute Component

The cloud server compute parameters include:

- Number of CPUs
- Amount of RAM

The importance of the computing resources of the host is the ability to allocate and manage resources effectively. The utilization of **reservation**, **limit**, and **share** provides the contextual ability to the administrator for allocating the host computing resources.

In reservation, a guarantee minimum resource allocation is generated. The host must meet the minimum resource allocation with its physical compute resources so that the guest will power on and operate. Traditionally, the reservation was available with the CPU or RAM, or both.

In the limit, the maximum ceiling for resource allocation is generated. The ceiling can be fixed or expandable; thus, depending on the customer's requirement, more compute resources can be procured from the borrowing scheme of the host.

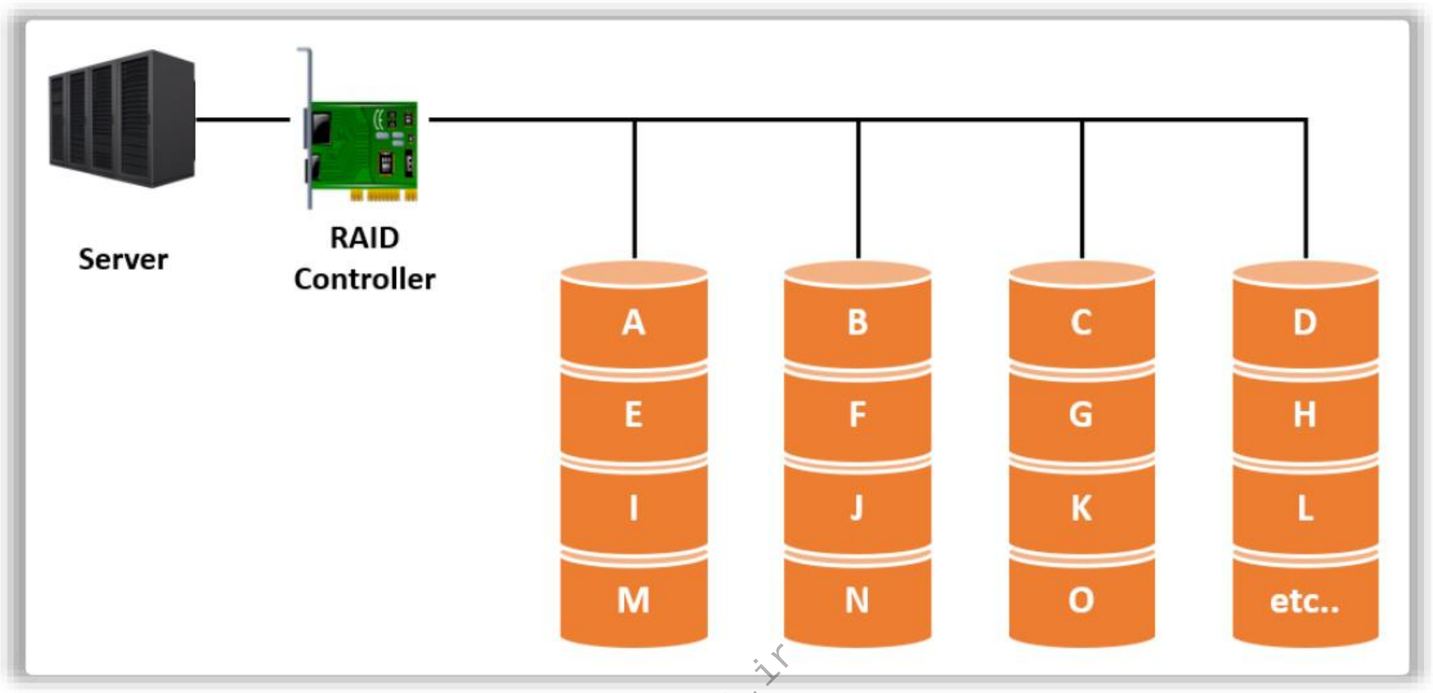
In share, concerns related to compute-resource contention situations are arbitrated. Resource contention denotes the availability of multiple requests for the resource than the actual available resources in the system. If there is resource contention, then share value will prioritize the compute-resource access to all the guests. The share value percentage is taken against all the

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

assigned shares. Hence, at the time of the contention period, if a guest has a higher share value, then a large percentage of the remaining resource access will be provided to that guest.

Storage Component

The storage component of cloud platforms and infrastructure allows an organization to store and manage its data on off-site file servers. Cloud storage providers such as Amazon Simple Storage Service (S3), Google Cloud, Microsoft Azure, etc., not only manage and maintain the data but they also keep remote backups. The cloud-stored data can be accessed through the internet or queried through an application deployed in the cloud. In a cloud infrastructure, cold data (non-actively used data) can also be stored.



hide01.ir

Figure: 2.3: Storage Component of Cloud Platforms and Infrastructure

For data storage in the cloud, solid-state drives (SSDs) or conventional hard disk drives are utilized. To provide redundancy, the disk drives are often grouped. One such approach is RAID (redundant array of inexpensive disks), which is a group of techniques. In RAID groups, redundant disks are configured such that if any disk fails, then the disk controller is still able to retrieve the data. On an average, disk drives have a yearly failure rate of approximately 3-5%. For example, if there are 5000 installed disks, then it can be expected that there will be 1 disk failure each day. The RAID techniques vary in terms of the percentage of redundant disks and the delivery of aggregate performance.

The storage functionality also includes slicing and grouping of disks into logical volumes of arbitrary sizes. There is no file system applied to the storage volume. The operating system on the virtual machine instance applies the file system structure to the storage volume.

Object Storage

Object storage refers to the file system provided by the CSP in which objects (files) are stored along with metadata. The objects can be accessed using APIs or web user interface. Generally, the files are organized in a directory hierarchy; however, in object storage system, the files are stored in containers, which have a flat organization, and each container has a unique ID for retrieval. In Amazon S3, containers are called buckets and the unique IDs are known as keys.

The OS image used by the hypervisor is stored in object storage. To enhance resilience, the object storage implements redundancy by distributing data through fragmenting and duplicating among

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

multiple object storage servers. This not only increases the resilience and performance, but it also reduces the risk of data loss. Using object storage, the client can store, copy, retrieve, and delete files.

Virtualization Component

Virtualization is a process of creating virtual rather than actual versions of networks, storage, servers, operating system, applications, etc. It is a foundational technology in cloud computing. Virtualization is the cost-saving, energy-conserving, and hardware reducing techniques in cloud technology utilized by the cloud provider. Cloud virtualization technology provides a virtual environment for storage, memory, networking, etc.

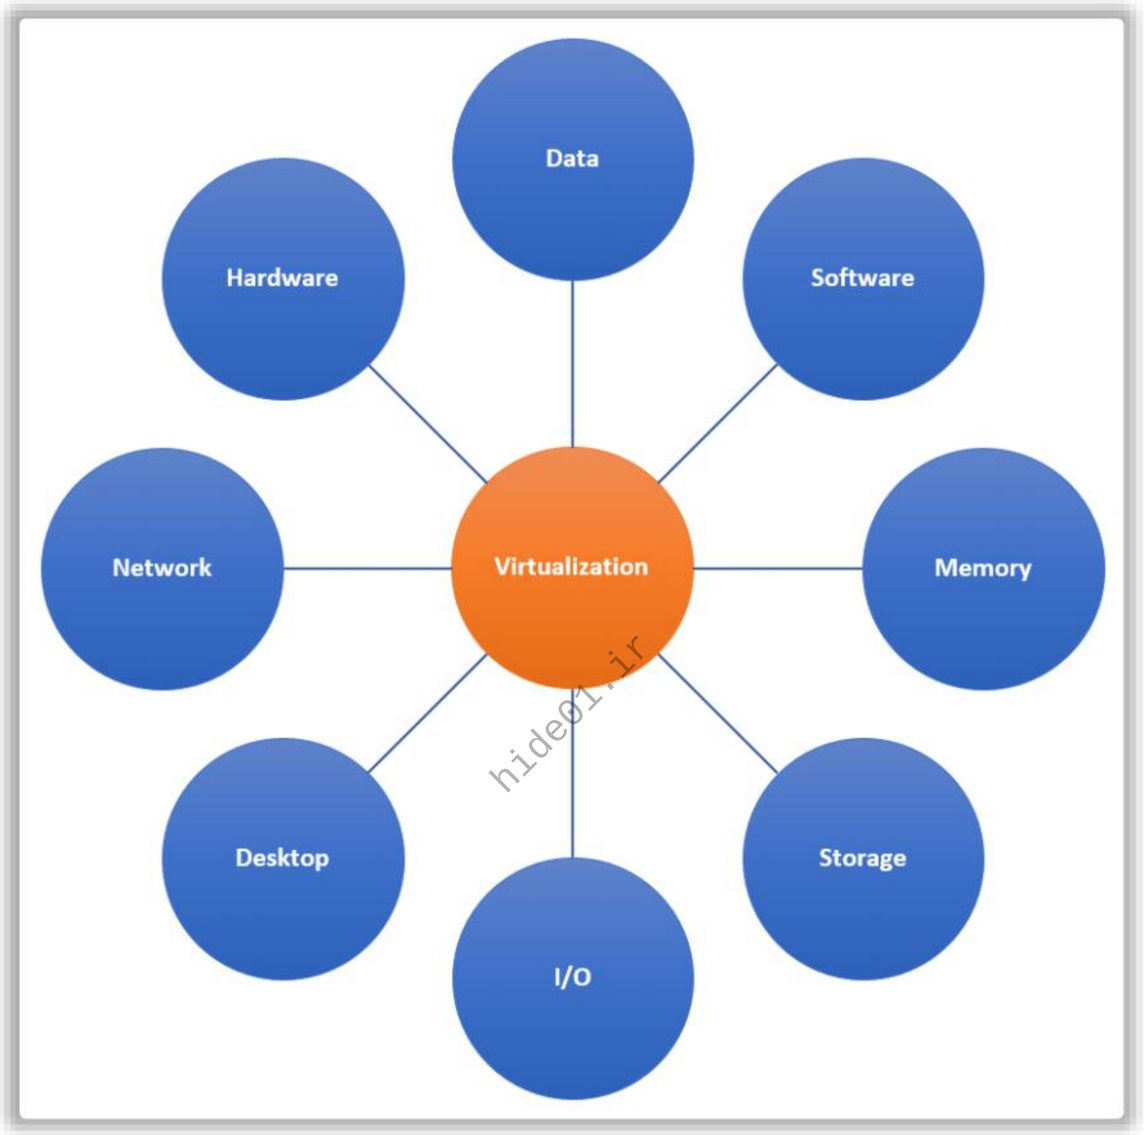


Figure: 2.4: Virtualization Component of Cloud Platform and Infrastructure

Key Components

The goal of virtualization is to create a logical interface by abstracting the underlying infrastructure. The key components of the virtualization environment are listed below.

- **Guest**: The component of the system that interacts with the virtualization layer rather than the host.
- **Host**: The original environment for the management of guests.
- **Virtualization layer**: It will recreate the same or different environment wherein the guest operates. All the resources of virtualization are present in this layer.

Virtualization relies on the utilization of powerful host computers to deliver a pool of resources that can be easily shared and managed to increase the guest OS running on each host. The sharing of underlying resources allows an efficient and agile hardware utilization.

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

Scalability: Using virtualization, customers can run many guests operating systems (VMs) and their associated applications on a single host. On the host machine, the guest (an isolated software instance) runs side by side with other guests, the resource abstraction capabilities of the hypervisor allow the guest to use the resources of the host.

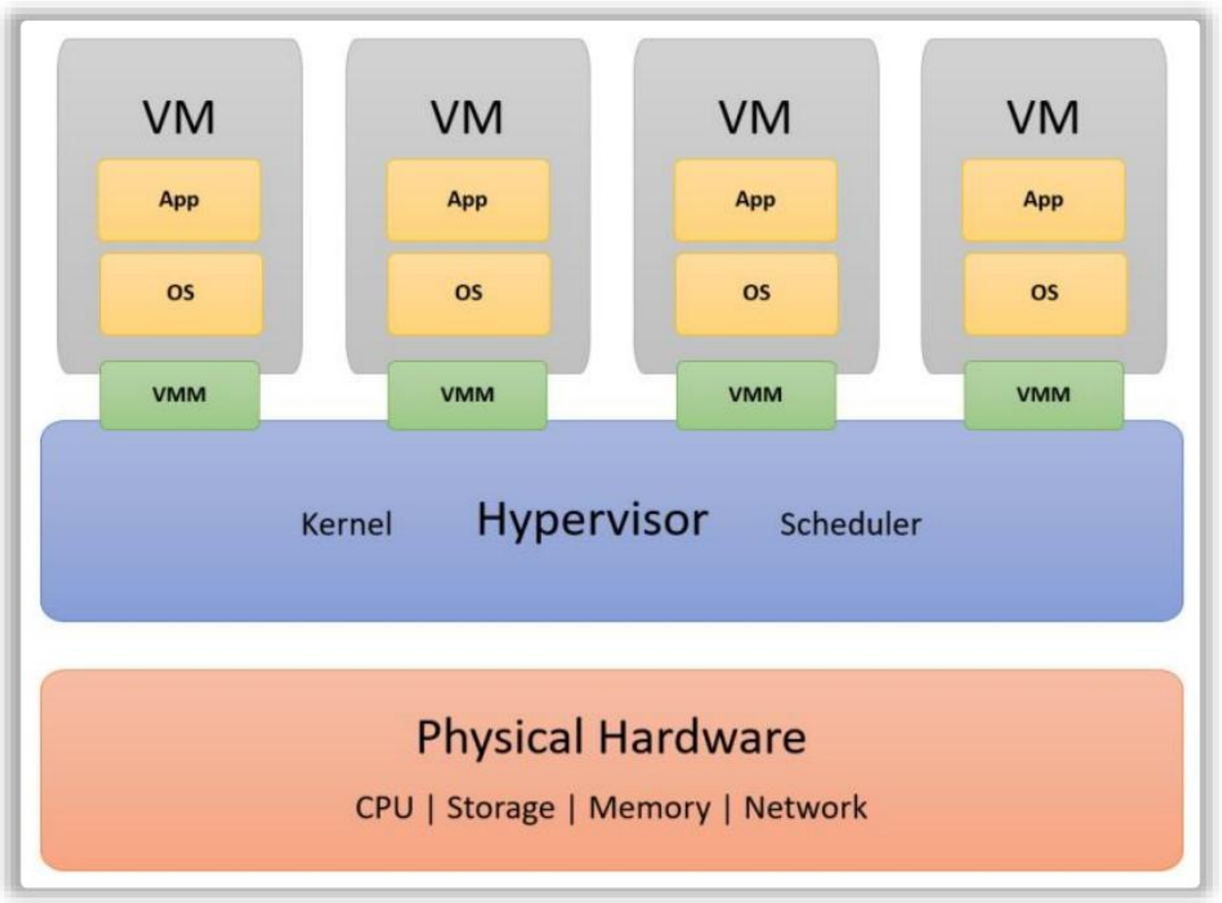


Figure: 2.5: Components of the Virtualization Environment

Hypervisor: It could be a piece of hardware, software, or firmware that provides the impression to the guest OS that it is running directly on the physical hardware of the host. The VM can access the physical hardware resources of the host through the hypervisor. Through the virtual machine manager (VMM) instance, the hypervisor establishes a connection with the VM and securely manages it. There are two types of hypervisors.

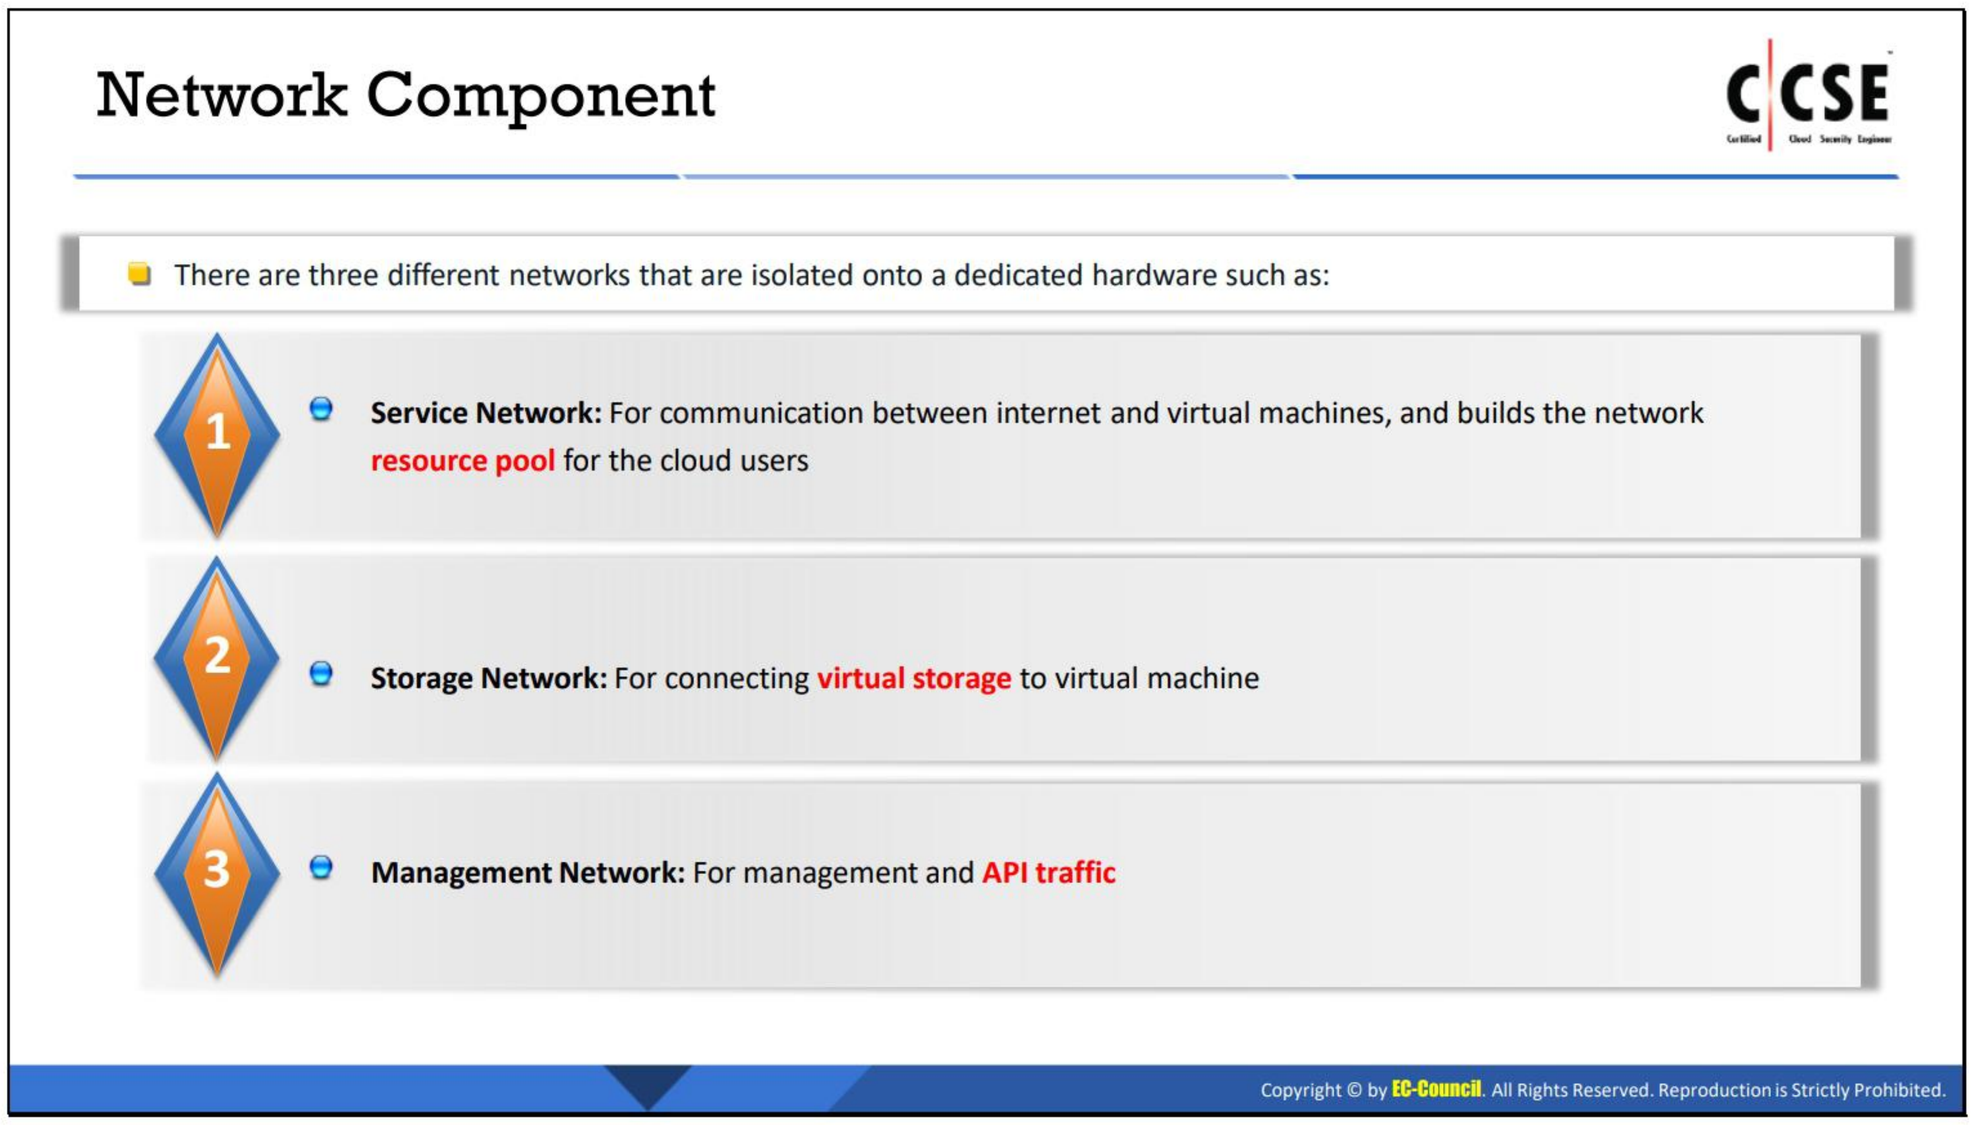
Type 1 Hypervisor: This type of hypervisor has the following characteristics

- It is also called as embedded, native, or bare-metal hypervisor.
- It monitors the OS that runs above it and works directly on the host hardware.
- It is small, as the key role of this hypervisor is to share and manage hardware resources between different guest operating systems.

Type 2 Hypervisor: Type 2 hypervisor has the following characteristics

- The hypervisor is installed on the host's OS and helps the guest operating systems running above it as VMs.
- This hypervisor relies on the host OS for its operations.

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



Network Component

Most clouds make use of virtual networking to create a network resource pool by abstracting the physical network. The cloud user provides the network resources from this pool and they can be configured within the limits of virtualization techniques.

hidoo1.ir

The physical segregation of networks in the cloud is essential due to operational and security reasons. Listed below are the commonly seen networks that are isolated on dedicated hardware.

- **Service Network**: The service network is used for communications between the internet and VMs. The service network builds the network resource pool that can be used by cloud users.
- **Storage Network**: The storage network is used for connecting virtual storage to VMs. Storage networks use high-speed connections to provide fast performance.
- **Management Network**: The management network is used for the management and API traffic. The main goal of a management network is to monitor, create, deploy, and configure the VMs and hosts in the virtual environment.

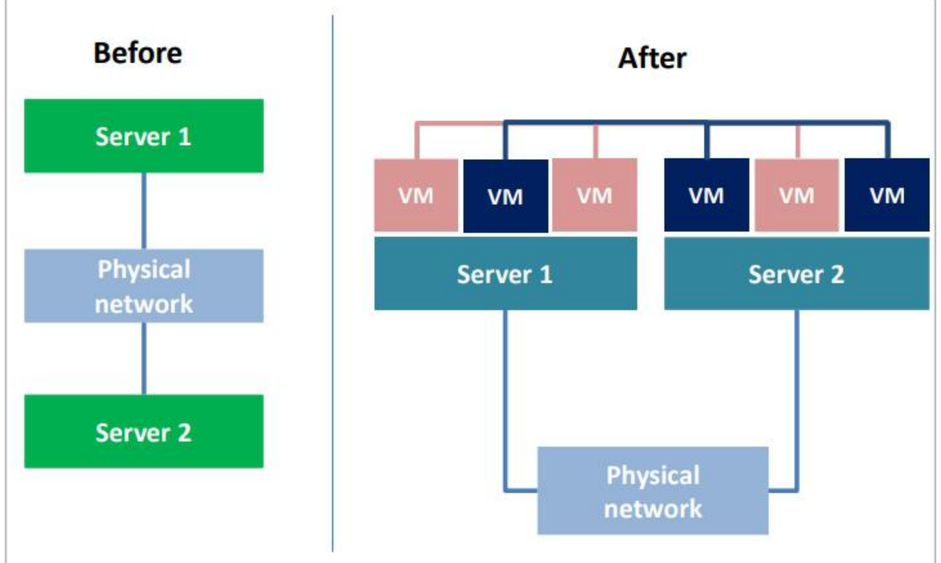
Certified Cloud Security Engineer Copyright © by **EC-Council**
All Rights Reserved. Reproduction is Strictly Prohibited.

Network Virtualization

- Network virtualization is the process of combining the hardware, software, and network functionalities into a single software based network called **virtual network**

Monitoring and Filtering of Network Traffic in Virtual Networks

- Monitoring and filtering of network traffic in virtual networks depends upon how **packets** move around
- Network traffic between two virtual machines in the same physical hardware can be monitored by routing the traffic to a **virtual monitoring tool** on the same hardware
- Customer monitoring and filtering of some communications are not possible in public cloud as the communications occur in **cloud provider's network**
- **Built-in firewalls** are offered by the cloud platforms that operates within the SDN or hypervisor



Network Virtualization

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

Network Virtualization (Cont'd)

hide01.ir

- In the cloud computing ecosystem, there are two important categories of network virtualization namely:

1. **Virtual Local Area Network (VLAN):** It is designed to utilize in **single-tenant network** for separating various business units, functions, etc., and is commonly used in an enterprise network
2. **Software-defined Network (SDN)**: It is a complete **abstraction layer** above the networking hardware that decouples the network control plane from the data plane. It provides high flexibility and isolation

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

Network Virtualization

Network virtualization is the process of combining the hardware, software, and network functionalities into a single software-based network called a virtual network. Cloud computing environments use SDN for virtual networks.

In the cloud computing ecosystem, there are two important categories of network virtualization.

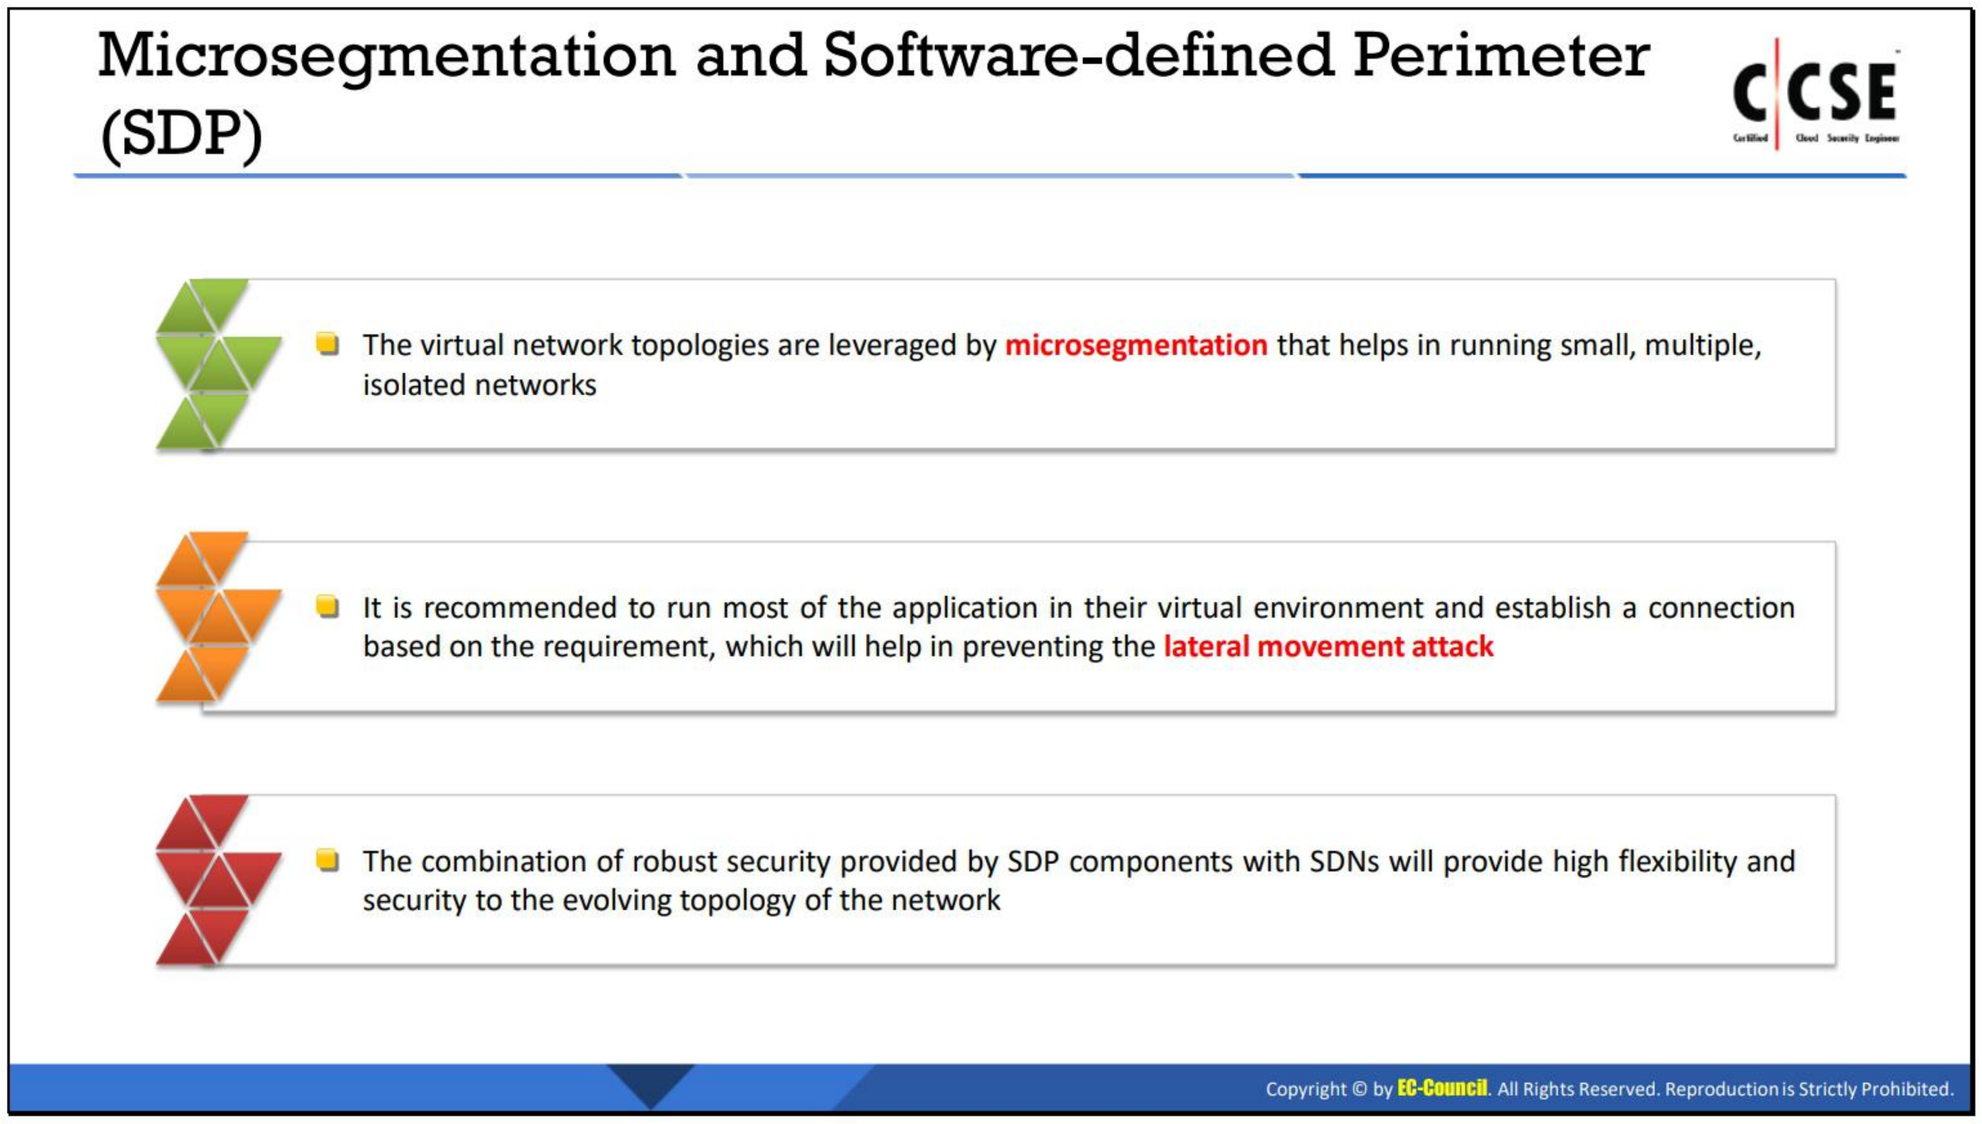
Certified Cloud Security Engineer Copyright © by **EC-Council**
All Rights Reserved. Reproduction is Strictly Prohibited.

- **Virtual Local Area Networks (Vlans)**: It is designed to utilize in *single-tenant networks* for separating various business units, functions, etc., and is commonly used in enterprise networks.
- **Software-Defined Network (SDN)**: It is a complete *abstraction layer* above the networking hardware that decouples the network control plane from the data plane. It provides high flexibility and isolation.

Monitoring and filtering of network traffic in virtual networks depends upon how packets move around. Network traffic between two VMs in the same physical hardware can be monitored by routing the traffic to a virtual monitoring tool on the same hardware. Customer monitoring and filtering of certain communications is not possible in a public cloud as the communications occur in the cloud provider's network. Built-in firewalls are offered by the cloud platforms that operate within the SDN or hypervisor.

hide01.ir

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



Microsegmentation and Software-defined Perimeter

Virtual network topologies are leveraged by micro-segmentation that helps in running small, multiple, isolated networks. Micro-segmentation or hyper-segregation runs isolated networks without any additional hardware cost because it defines these networks in software. As all the virtual networks are defined in software, it is more feasible to run these software-defined environments.

It is recommended to run most applications in their virtual environment and establish a connection based on the requirements, which helps in preventing lateral movement attacks. If an attacker targets and compromises a system, the attacker will be confined to a particular network where the system lies and cannot access other networks. Thus, they will not be able to use this foothold to move across the entire data center.

Cloud micro-segmentation is based on software configurations; therefore, it does not increase capital expenses. However, it may increase operational expenses due to managing several overlapping networks and their connectivity.

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

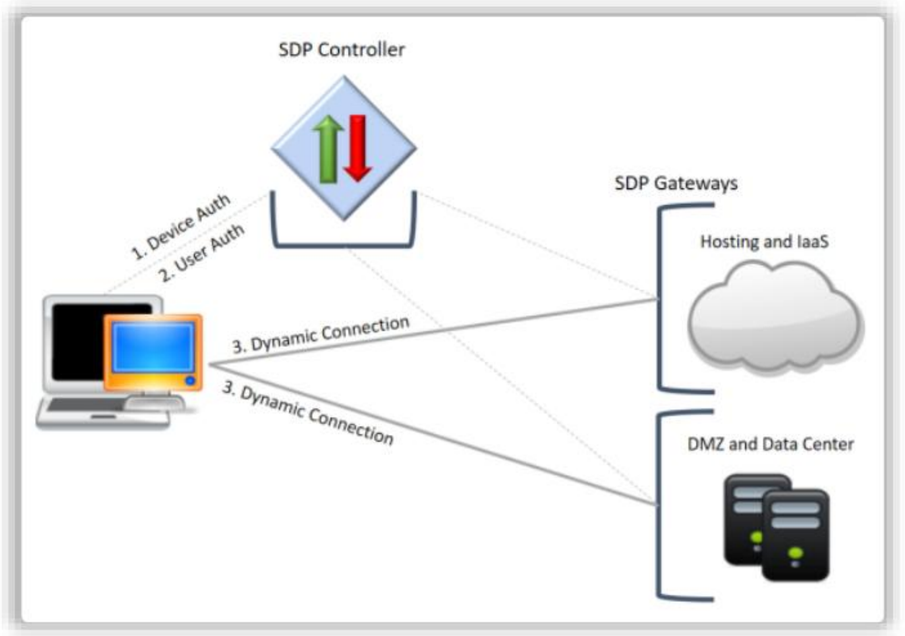


Figure: 2.6: Micro-segmentation and SDP

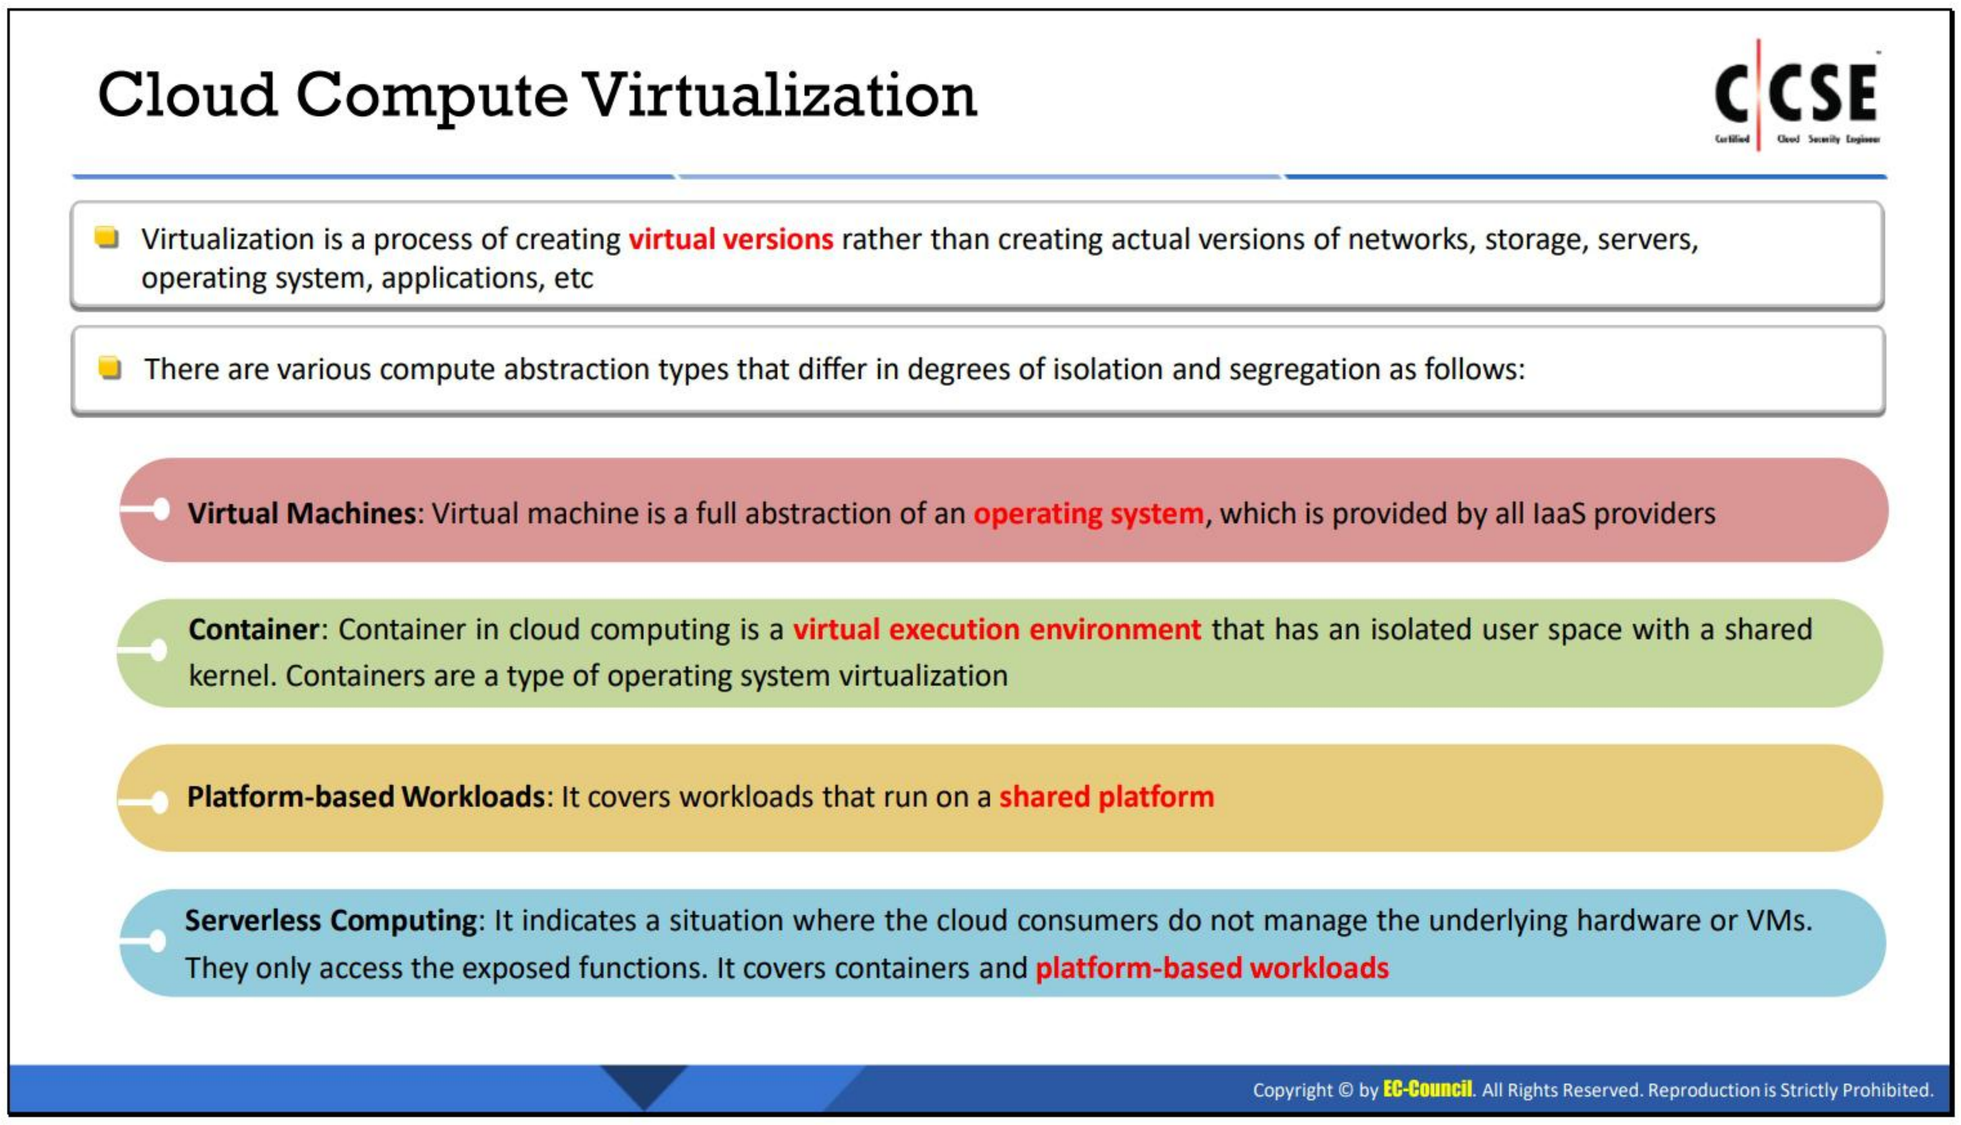
For provisioning network access to resources dynamically and improving security, a model has been developed by the CSA SDP Working Group that combines device and user authentication. There are three components in SDP.

- **SDP Client**: SDP clients are present on connecting assets like laptops.
- **SDP Controller**: It authenticates and authorizes SDP clients and configures connections to SDP gateways.
- **SDP Gateway**: It terminates SDP client network traffic and implements policies along with the SDP controller in communication.

The combination of robust security provided by SDP components with SDNs provides high flexibility and security to the evolving topology of the network. With SDNs, a wide range of criteria can be used rather than just IP packets to make network security decisions.

hi-e01.ir

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



Cloud Compute Virtualization

A workload is a processing unit like a VM, a container, or other abstractions that run on a processor and consume memory. It includes a diverse range of processing tasks ranging from a traditional application on the standard OS of a VM to GPU- or FPGA-based specialized tasks. Compute abstraction types are described below.

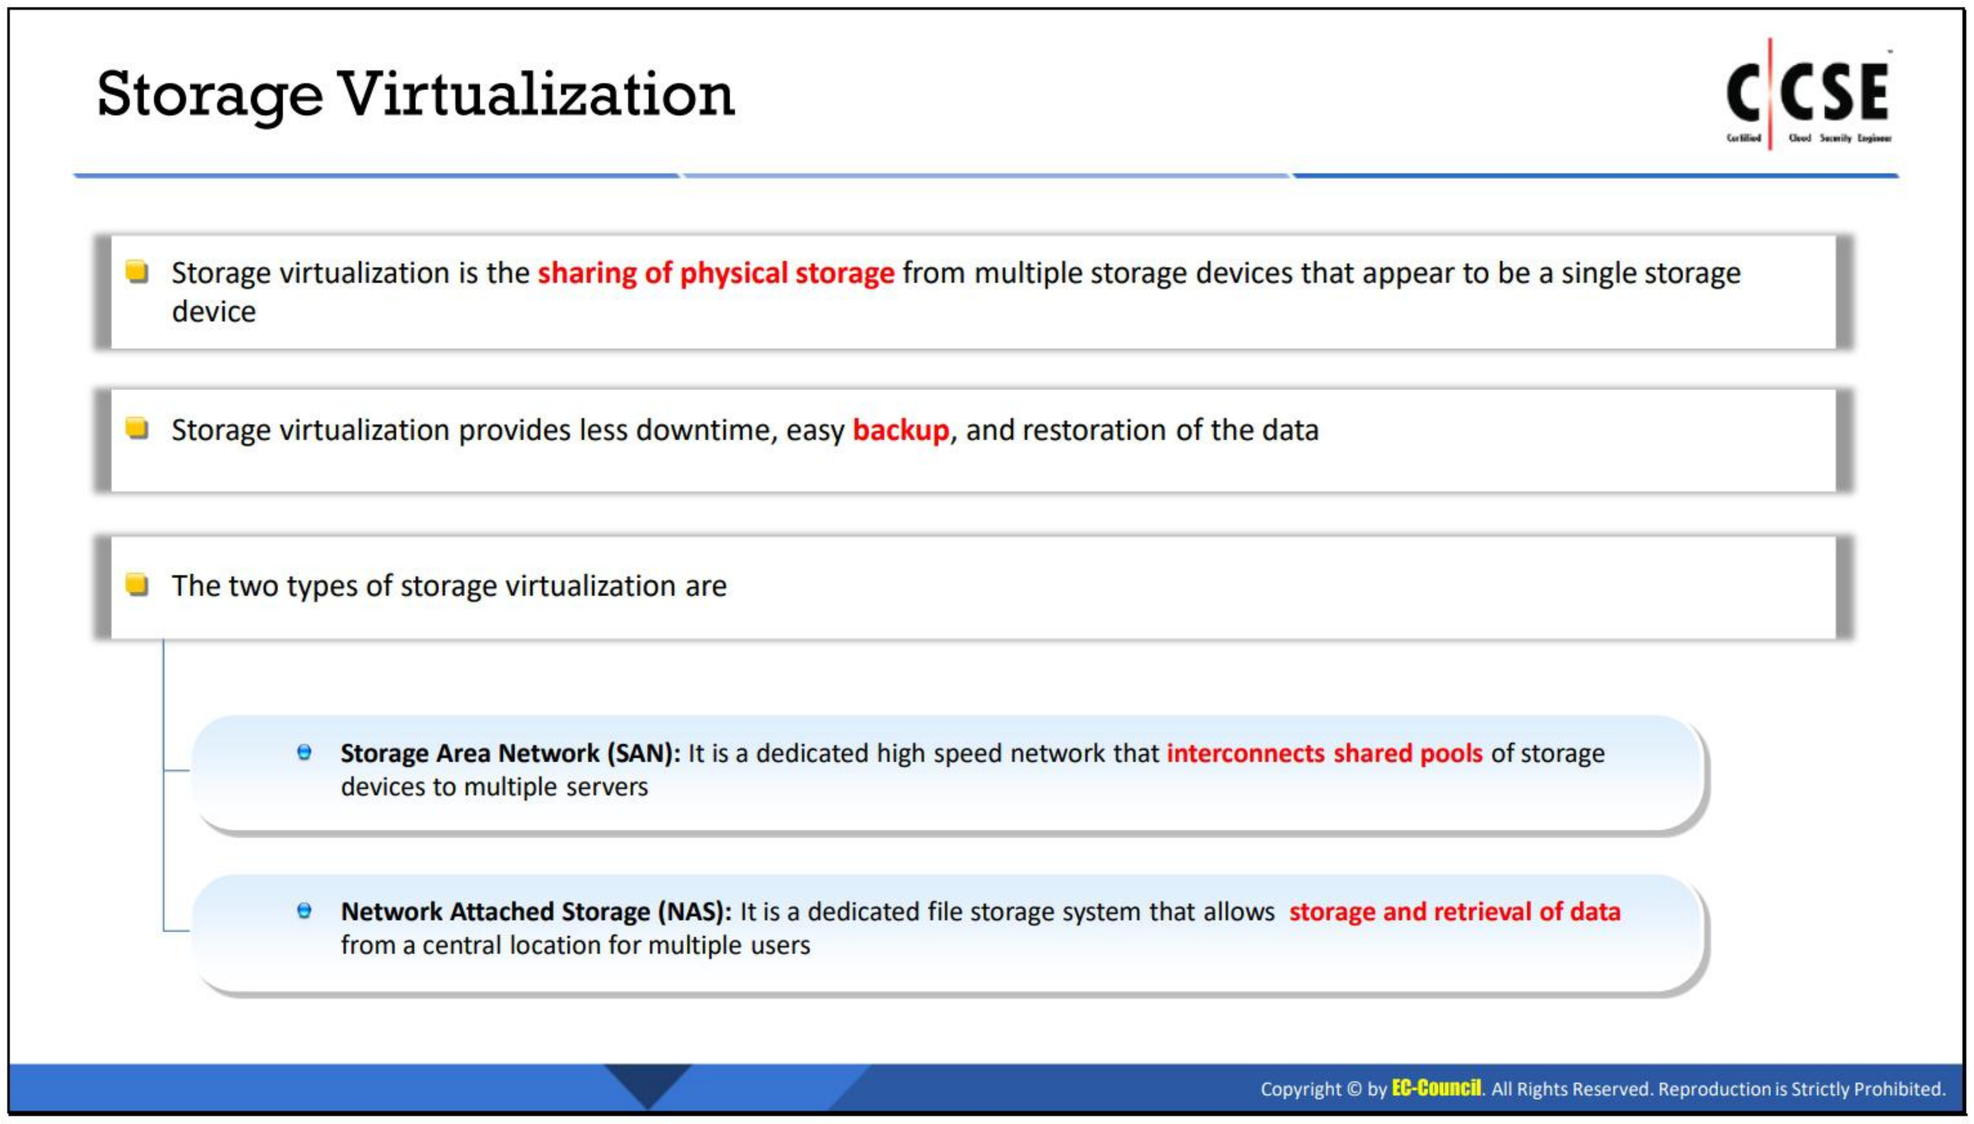
- **Virtual Machines**: VMs offer compute abstraction for all IaaS providers, which are called instances in cloud computing and are created based on pictures. The virtual machine manager (hypervisor) abstracts OS from hardware. VMs are vulnerable to memory attacks, and it is challenging to understand reinforcement of isolation due to ongoing hardware and software enhancements. It effectively controls security and hardware isolation for VMs and secures the execution environment to improve capabilities.
- **Containers**: Containers are the code execution environments that share and leverage resources for the OS. A VM's complete abstraction of the OS is a constrained place to execute segregation processes. Multiple machines run on the same VM in a container or are implemented without VMs to run directly on hardware. They provide code execution inside a restricted environment with access to container configuration for processing and capabilities. This allows containers to launch multiple services rapidly. The containers access running services in the host OS with different isolation capabilities on the platform-dependent environment. This offers evolution on different management systems underlying operating systems and container technologies.
- **Platform-based Workloads:** They enable workloads running on shared platforms that are not VMs or containers, such as when a stored procedure runs inside a multitenant database service. The platform provides isolation and security with specific exposure to security options and controls.

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

- **Serverless Computing:** It offers broad categories of situations in which cloud users do not need to manage hardware or VMs using exposed features, such as serverless platforms for directly executing application code. Capabilities such as containers, VMs, or specialized hardware platforms implemented to utilize in serverless computing. It covers containers and platform-based workloads to manage foundational security functions and controls of underlying layers.

hide01.ir

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



Storage Virtualization

Storage virtualization in cloud computing refers to the allotment of physical storage in multiple storage devices that appears as a single storage device. It tends to be called a group of accessible storage devices and offers less downtime as well as easier backups and recovery of data. Storage virtualization provides better data security as it reduces server traffic and loss of data.

Types

Listed below are the two main types of storage virtualization.

- **Storage Area Network (SAN):** SAN is a highly secured storage device that provides shared pool storage space accessed by multiple servers or computers. It allows computers on the network to access SAN storage by connecting a local disk to the computer and eliminates traditional network bottlenecks. SANs are also referred as SAN storage, SAN network, Network SAN, etc. They provide block-level data storage and are used for accessing storage devices such as tape libraries and disk-based devices from servers. SANs are designed primarily for block storage inside databases, also known as *structured data*.
- **Network-attached Storage (NAS)**: NAS is file storage that allows multiple users and devices to retrieve data from a centralized disk capacity. NAS devices do not have a display or keyboard. They are configured and managed through browsers. NAS consists of independent nodes that define a unique IP address. LAN users access storage via standard Ethernet connections. It provides easy access, high capacity, and cost-efficient storage infrastructure that supports backup, recovery, and multitenancy tasks. NAS can handle unstructured data, such as audio, video, websites, text files, and Microsoft Office documents.

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

Storage Virtualization Methods

Listed below are the methods of storage virtualization.

- **File-based Storage Virtualization**: It uses server message block or filling system protocol to break the reliance during storage exhibition with the help of assistance. It stores the data in actual memory and provides document migration to improve the presentation of data.
- **Block-based Virtual Storage**: It uses logical storage, such as drive partitions, from the physical memory of the storage device to an abstract disk drive or memory devices. It allows virtualization management software to access the available devices and split them into shared resources to assign them to clients or devices.

Shared Responsibility Network Virtualization Security

CSPs and cloud users are both responsible for the security of network virtualization. The shared responsibilities of CSPs and users are discussed below.

Cloud Service Providers Responsibilities

Following are the responsibilities of CSPs in network virtualization security:

- CSPs should implement and configure a secured network infrastructure.
- CSPs should provide network segregation and isolation on network traffic in a multitenant environment to prevent tenants from viewing another tenant's data.
- CSP's responsibility is to disable packet sniffing or metadata leaks that cause data exposure or configuration between multiple tenants and even within tenant's virtual networks to monitor non-virtual networks and reduce compromise of nodes from attackers.
- CSPs should ensure that the taggings and SDN-level metadata are not exposed outside the management planes as attackers could use this compromised data to access the cloud.
- The providers should enable built-in firewalls in the virtual networks to ensure security of user data without any need for a host firewall or other products.
- Providers are responsible for the perimeter security of the cloud for preventing and detecting attacks on physical networks and virtual platforms.

Cloud Service Users Responsibilities

Following are the responsibilities of cloud users in network virtualization security:

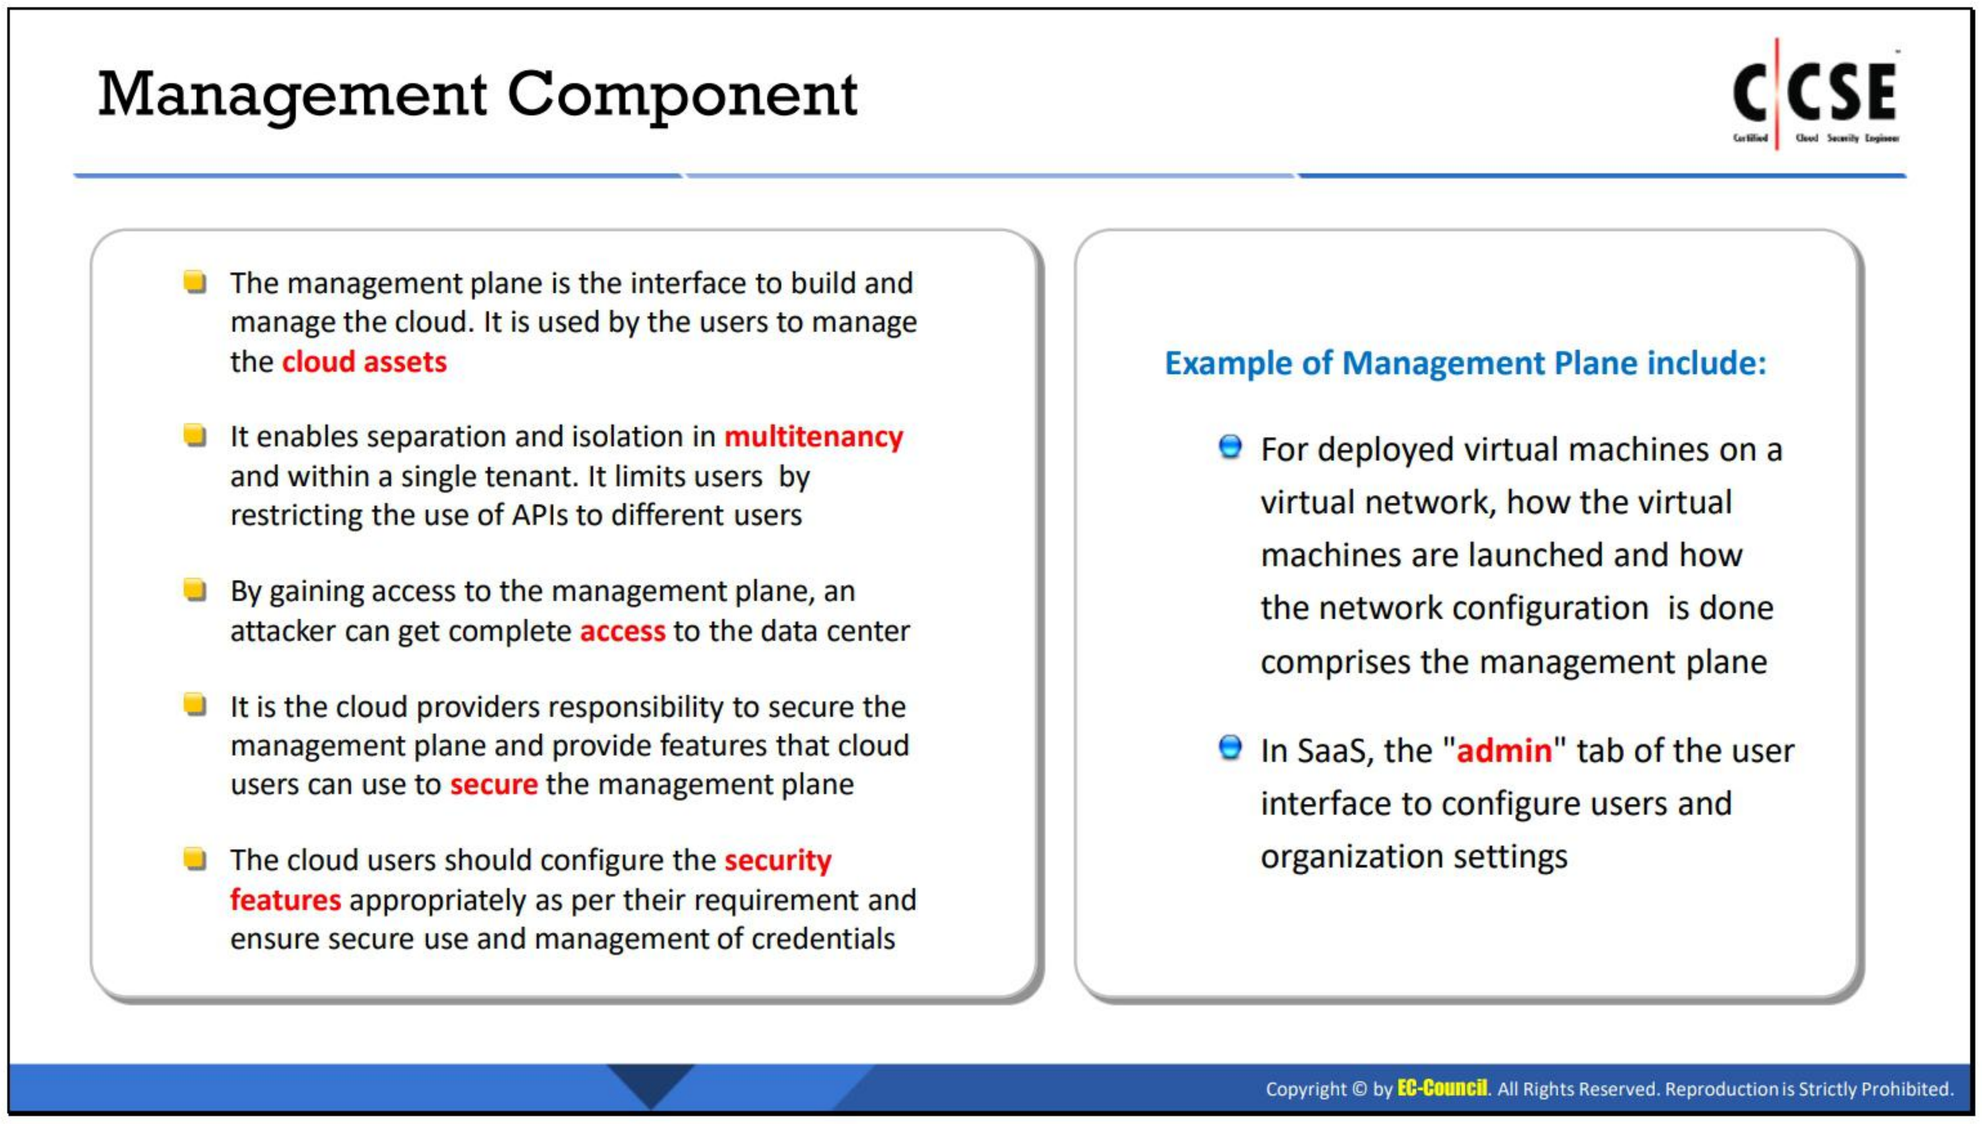
- Cloud users should adequately configure and enable the firewall on their virtual network.
- Network architecture can play a critical role in virtual network security to establish physical connections. Cloud users are responsible for designing a proper network architecture with separate virtual networks to offer advantages that are not possible on a traditional physical network. An equivalent architecture on a physical network is cost-efficient.

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

- Cloud users should define immutable networks using software templates to provide an entire known-good state of the network that can be explained in a template instead of configuring all the settings manually.
- The cloud user is responsible for proper rights management and configuration of exposed controls in the management plane.
- The cloud service users should compensate virtual security appliances or host security agents when virtual firewalls do not meet security needs and configure exposed control on a management plan.

hide01.ir

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

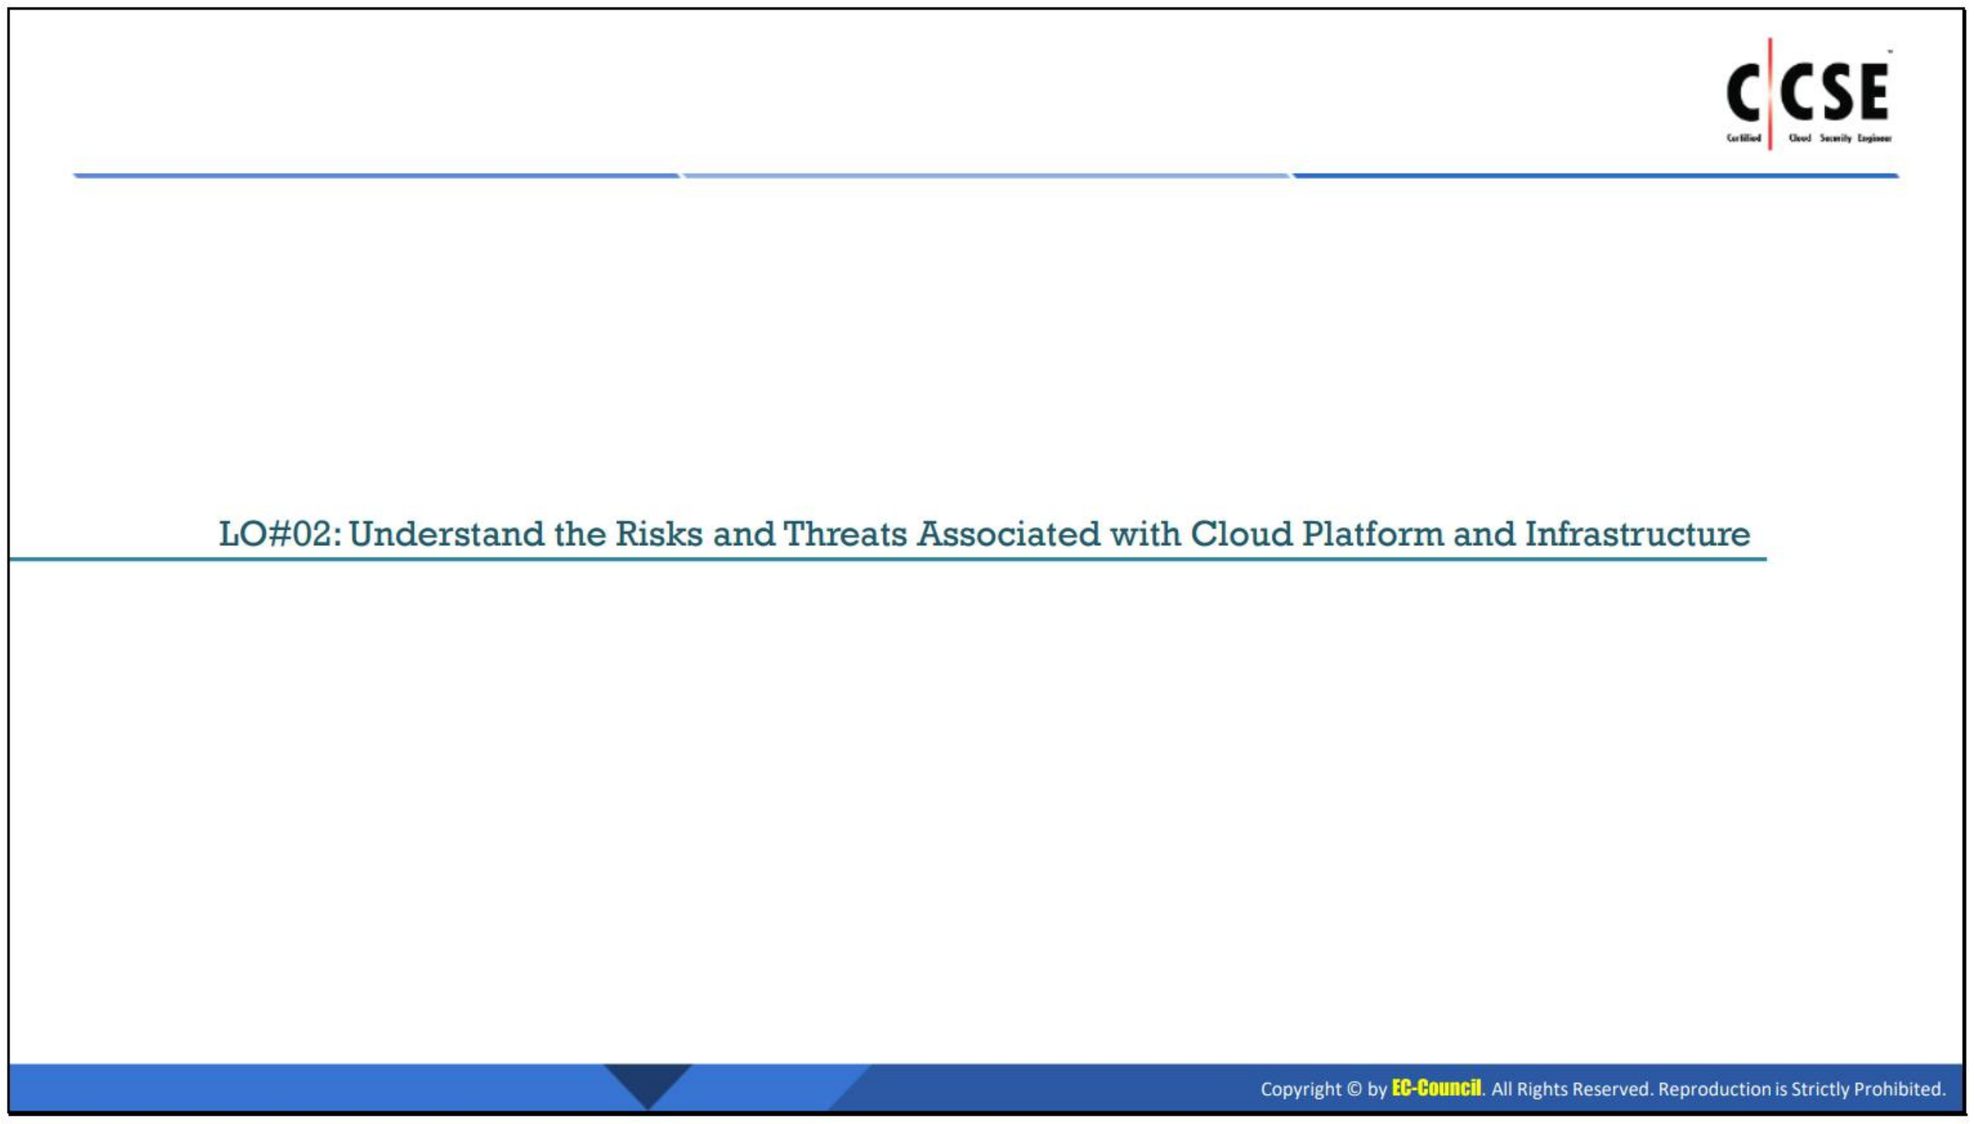


Management Component

The management plane is the interface to build and manage the cloud. It is used by the users to manage the cloud assets. It enables separation and isolation in multitenancy and within a single tenant. It limits users by restricting the use of APIs to different users. By gaining access to the management plane, an attacker can get complete access to the data center. It is the cloud providers responsibility to secure the management plane and provide features that cloud users can use to secure the management plane. The cloud users should configure the security features appropriately as per their requirement and ensure secure use and management of credentials. Following are some examples of Management Plane:

- For deployed virtual machines on a virtual network, how the virtual machines are launched and how the network configuration is done comprises the management plane.
- In SaaS, the "**admin**" tab of the user interface to configure users and organization settings.

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



LO#02: Understand the Risks and Threats Associated with Cloud Platform and Infrastructure

To secure the cloud infrastructure, the cloud security engineer should understand the various risks and threats that exist in a cloud environment. This section provides an understanding of the various security risks and threats associated with the cloud platform and infrastructure, along with challenges posed by virtual appliances.

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

Risks Associated with Cloud Platform and Infrastructure

- Every cloud-provider or cloud-consuming organization is responsible for their **own risk assessment**
- Risks associated with cloud infrastructure and platform include:

Policy and Organization Risks

- Risks are associated with the choice of the customer about the **cloud service provider**
- The key policy and organization risks are provider lock-in, loss of governance, compliance challenges, and provider exit

General Risks

- General risks occur when there is a **failure in meeting requirements** that are expressed in technical terms such as operability, performance, protection, and integration

Virtualization Risks

- Some of the virtualization risks include guest breakout, snapshot and image security, and sprawl

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

Risks Associated with Cloud Platform and Infrastructure (Cont'd)

hide01.ir

Non-Cloud-Specific Risks

- Many IT risks still exist in the cloud environment such as unauthorized facility access, **network attack** on customer and CSP side, social engineering, default passwords, and other non-malicious and malicious actions

Cloud-Specific Risks

Some of the cloud-specific risks are:

- Management plane breach
- Resource exhaustion
- Isolation control failure
- Insecure or incomplete data deletion
- Control conflict risk
- Software-related risks

Legal Risks

From a legal perspective, cloud computing brought various new risks as follows:

- Data protection
- Jurisdiction
- Law enforcement
- Licensing

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

Risks Associated with Cloud Platform and Infrastructure

There are various lists of risks published and maintained by industrial organizations, which provide valuable information. However, every cloud-providing or cloud-consuming organization is responsible for their own risk assessment. Following are the general categories of risks associated with cloud infrastructure and platform:

- Policy and organization risks

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

- General risks
- Virtualization risks
- Cloud-specific risks
- Legal risks
- Non-cloud-specific risks

Policy and Organization Risks

These types of risks are associated with the choice of the customer about the CSP. It might be the consequence of outsourcing IT services. It is also called as third-party risks. Some of the key risks are provider lock-in, loss of governance, compliance challenges, and provider exit.

- **Provider Lock-In**: It refers to a situation wherein the cloud consumer has made substantial vendor-specific investments such as adaptation to data formats, feature sets, and procedures. Thus, these investments can lead to huge costs when switching between providers.
- **Loss of Governance**: It refers to the cloud consumer who is not able to enforce all the required controls. This generally applies to consumers who did not realize the required level of security and compliance risks.
- **Compliance Risks**: The cloud consumer has a compliance obligation while dealing with payment card information, health data, or other personally identifiable information (PII). A particular cloud vendor and solution may not fulfill all such obligations.
- **Provider Exit**: It refers to a situation in which the provider is incapable of delivering the required services. It might occur due to bankruptcy or the restriction of the business.

General Risks

General risks occur when there is a failure in meeting requirements that are expressed in technical terms, such as operability, performance, protection, and integration. The extent of technology associated with the CSP is much higher than the cloud customer and traditional IT departments. This has an impact on risk and the net outcome will depend on the actual situation:

- IT infrastructure consolidation leads to consolidation risk in which a single point of failure has a tremendous impact.
- A larger platform necessitates the CSP to procure more technical skills to maintain and manage the infrastructure.
- Control on technical risks shifts toward the provider.

Virtualization Risks

Some of the risks of virtualization are listed below.

- **Guest Breakout**: It refers to a situation in which the guest OS breaks out to access the hypervisor or other guests. This might occur due to the hypervisor flaw.

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

- **Snapshot and Image Security**: Due to the portability of images and snapshots, one may forget that it contains sensitive information and should be safeguarded.
- **Sprawl**: It occurs when the control on the amount of content is lost on the image store.

Non-Cloud-Specific Risks

Many IT risks still happen in the cloud environment, such as unauthorized facility access, network attack on the customer and CSP side, social engineering, default passwords, and other non-malicious and malicious actions.

Cloud-Specific Risks

Some of the cloud-specific risks are included below.

- **Management Plane Breach**: It is the most important risk and it is also called a management interface breach. Herein, the attacker affects the entire infrastructure that is controlled by the management interface.
- **Resource Exhaustion**: As the cloud resources are shared, resource exhaustion signifies a risk to the cloud consumer. In resource exhaustion, access to the cloud resources that are already provisioned will be denied or will be unable to increase resource consumption. For example, lack of CPU or network bandwidth, which can occur due to overprovisioning to the tenants by the CSP. The following are related to resource exhaustion:
  - Denial-of-service (DoS) attacks in which the network or other resources are saturated and can cause starvation to users
  - Traffic analysis
  - Interception of data-in-transit
- **Isolation Control Failure**: The CSP must realize isolation controls while sharing resources across tenants. Failure or nonexistence of these controls is considered as isolation failure. For example, one tenant VM instance affecting or accessing the instance of another tenant; this can lead to the blacklisting of the IP address because of one tenant's activity.
- **Insecure or Incomplete Data Deletion**: In many operating systems, data erasure is done by deleting directory entries instead of reformatting the used storage. If the storage is reused, then the sensitive data will be at risk due to the recovery and exposure of data.
- **Control Conflict Risk**: In a shared environment, the controls that provide high security to one cloud consumer (blocking traffic) will make it less secure for another (loss of visibility).
- **Software-related Risks**: Each CSP runs multiple software, and each software would have potential vulnerabilities. From the customer's point of view, control is transferred to the CSP to increase security and risk awareness; however, the client is still accountable for compliance.

Legal Risks

From a legal perspective, cloud computing has brought various new risks such as data protection, jurisdiction, law enforcement, and licensing.

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

- **Data Protection**: There are legal requirements for cloud consumers about the way they safeguard the data, specifically PII. CSP controls and actions may be insufficient for the cloud consumer.
- **Jurisdiction**: The data storage locations of the CSPs fall in multiple jurisdictions that can impact other risks and their controls.
- **Law Enforcement**: Due to civil legal activity or law enforcement, the data may need to be handed over to the authorities. As the cloud resources are shared, this would cause exposure risks for the other tenants. For instance, the physical disk examination could expose the data of various customers.
- **Licensing**: If the customer wants to move their software to the cloud environment, then the licensing agreements on the software makes it either a legally impossible or a costly option.

hide01.ir

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

Threats to Cloud Platform and Infrastructure

- Most threats in cloud computing affect the **cloud platform** and infrastructure
- These threats are cloud specific or they are inherited by traditional IT infrastructure
- The common **threats** to cloud platform and infrastructure are data breaches, malware injections, DDoS, regulatory compliance, insecure APIs, advanced persistent threats (APTs), malicious insiders, account hijacking, etc

- **The following table includes the major threats and their mitigations:**

Threats	Mitigations
Natural Disasters Damaging the Cloud Infrastructure: Natural disasters such as fire, hurricanes, earthquakes, etc., are considered serious threats as they can damage the **cloud infrastructure**	Utilize a risk assessment tool to reduce the consequences of natural disasters
Unauthorized Physical Access to Equipment or Facilities: Unauthorized users gaining access to the cloud system facilities or equipment can lead to **denial of service** (DoS)	Use a risk assessment tool to mitigate unauthorized physical access to equipment or facilities
Employee Negligence: The negligence of the employee can lead to **accidental deletion** or loss of back up data. It can also lead to loss of security or operational logs	Use a risk management plan in combination with security policy
Privilege Escalation: Utilizing the shared memory of the virtual machines or the hypervisor of the cloud host, the adversary can use a virtual machine to compromise another virtual machine by escalating the **access right**	Use updated hypervisor and enforce privilege escalation countermeasures
Insecure Data Deletion: Due to some technical glitches, the data may not be completely wiped out and can be **exposed**	Ensure complete deletion of data

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

Threats to Cloud Platform and Infrastructure (Cont'd)

hide01.ir

Threats	Mitigations
Obsolete Cryptography: Inaccurate implementation of the cryptographic or encryption protocols is a **serious threat** to cloud platform and infrastructure	Give high priority to the accurate implementation of contemporary cryptographic techniques
Failure of Cloud Service: By obtaining elevated privileges, the attacker can disable or disturb the services of the cloud platform and infrastructure. This leads to **cloud service termination** or failure	Install multiple types of IDs in various virtual machines
Insufficient Monitoring and Logging: When **contemporary logging** and monitoring mechanism is used, it is quite difficult to retrieve the information relevant to a single user or service	Use robust logging and monitoring tool and, also implement appropriate security controls
Failure of Third-party Suppliers: The cloud service provider outsources various tasks to third-party suppliers, who might implement low-security standards or will not pay attention to the **cloud infrastructure security policy**	The CSP should be cautious with the third-party suppliers or adopt alternate choices for outsourcing the tasks
Lock-In: Based on the infrastructure used by the customer, there are various **IaaS lock-in problems**	Perform thorough research before selecting the appropriate cloud service provider and pay more attention to any change in the cloud policy
Subpoena and e-Discovery: The legal framework safeguarding privacy and **processing of personal data** is different are different countries. In most countries, exposure of private information is not a punishable offense	To overcome privacy-related issues, the cloud consumer must consider the legal framework of the cloud provider

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

Threats to Cloud Platform and Infrastructure

Most of the threats in the cloud computing affects the cloud platform and infrastructure. These threats are cloud specific or they are inherited by traditional IT infrastructure. The major threats to cloud platform and infrastructure are described below.

- **Natural Disasters Damaging the Cloud Infrastructure**: Natural disasters such as fire, hurricanes, earthquakes, etc., are considered serious threats as they can have devastating

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

effects on the cloud infrastructure as well as human life. To mitigate against this threat, utilize a risk assessment tool to reduce the consequences of natural disasters.

- **Unauthorized Physical Access to Equipment Or Facilities**: Unauthorized users try to gain access to the cloud system facilities or equipment. This unauthorized physical access is serious threat to the cloud infrastructure as it threatens the system's devices and equipment and can lead to DoS for an extended period of time. For this threat, use a risk assessment tool to mitigate against unauthorized physical access to equipment or facilities.
- **Employee Negligence**: Negligence by employees can lead to accidental deletion or loss of back up data or it can also lead to loss of security or operational logs. To mitigate against this threat, use a risk management plan in combination with security policy. This will help the employee follow protocol of procedures and would reduce the probability of making serious mistakes.
- **Privilege Escalation**: Utilizing the shared memory of the VMs or the hypervisor of the cloud host, the adversary can compromise VMs by escalating the access right. To mitigate against this threat, use an updated hypervisor and enforce privilege escalation countermeasures.
- **Insecure Data Deletion**: In cloud platform and infrastructure, when there is a need to delete the resource, most of the time the OS fails to delete the data completely or, in some cases, timely data deletion is not possible. Similarly, when the cloud provider applies various changes to the architecture, such as destroying older hardware, reallocating hardware, etc., the data may not get completely wiped out and can get exposed. To mitigate against this threat, ensure complete deletion of data.
- **Obsolete Cryptography**: Inaccurate implementation of the cryptographic or encryption protocols is a serious threat to cloud platform and infrastructure. To mitigate against this threat, give high priority to the accurate implementation of contemporary cryptographic techniques.
- **Failure of Cloud Service**: By obtaining elevated privileges, an attacker can disable or disturb the services of the cloud platform and infrastructure leading to service termination or failure. To mitigate against this threat, install multiple types of IDSs in various virtual machines.
- **Insufficient Monitoring and Logging**: Using contemporary logging and monitoring mechanism, it is quite difficult to retrieve the information relevant to a single user or service. To mitigate against this threat, use robust logging and monitoring tool and also implement appropriate security controls.
- **Failure of Third-party Suppliers**: The CSP outsources various tasks to third-party suppliers, who might implement low-security standards or will not pay attention to the cloud infrastructure security policy. To mitigate against this threat, the CSP should be cautious when employing third-party suppliers or should adopt alternate choices for outsourcing the tasks.
- **Lock-in**: Cloud lock-in occurs when the cloud consumer wants to transfer their products, data, or services to another CSP but is unable to do so due to projected cost, complexity of switching, or duration. This problem is observed in all the three cloud architectures (SaaS, PaaS, and IaaS). The lock-in problem in IaaS depends on the infrastructure utilized

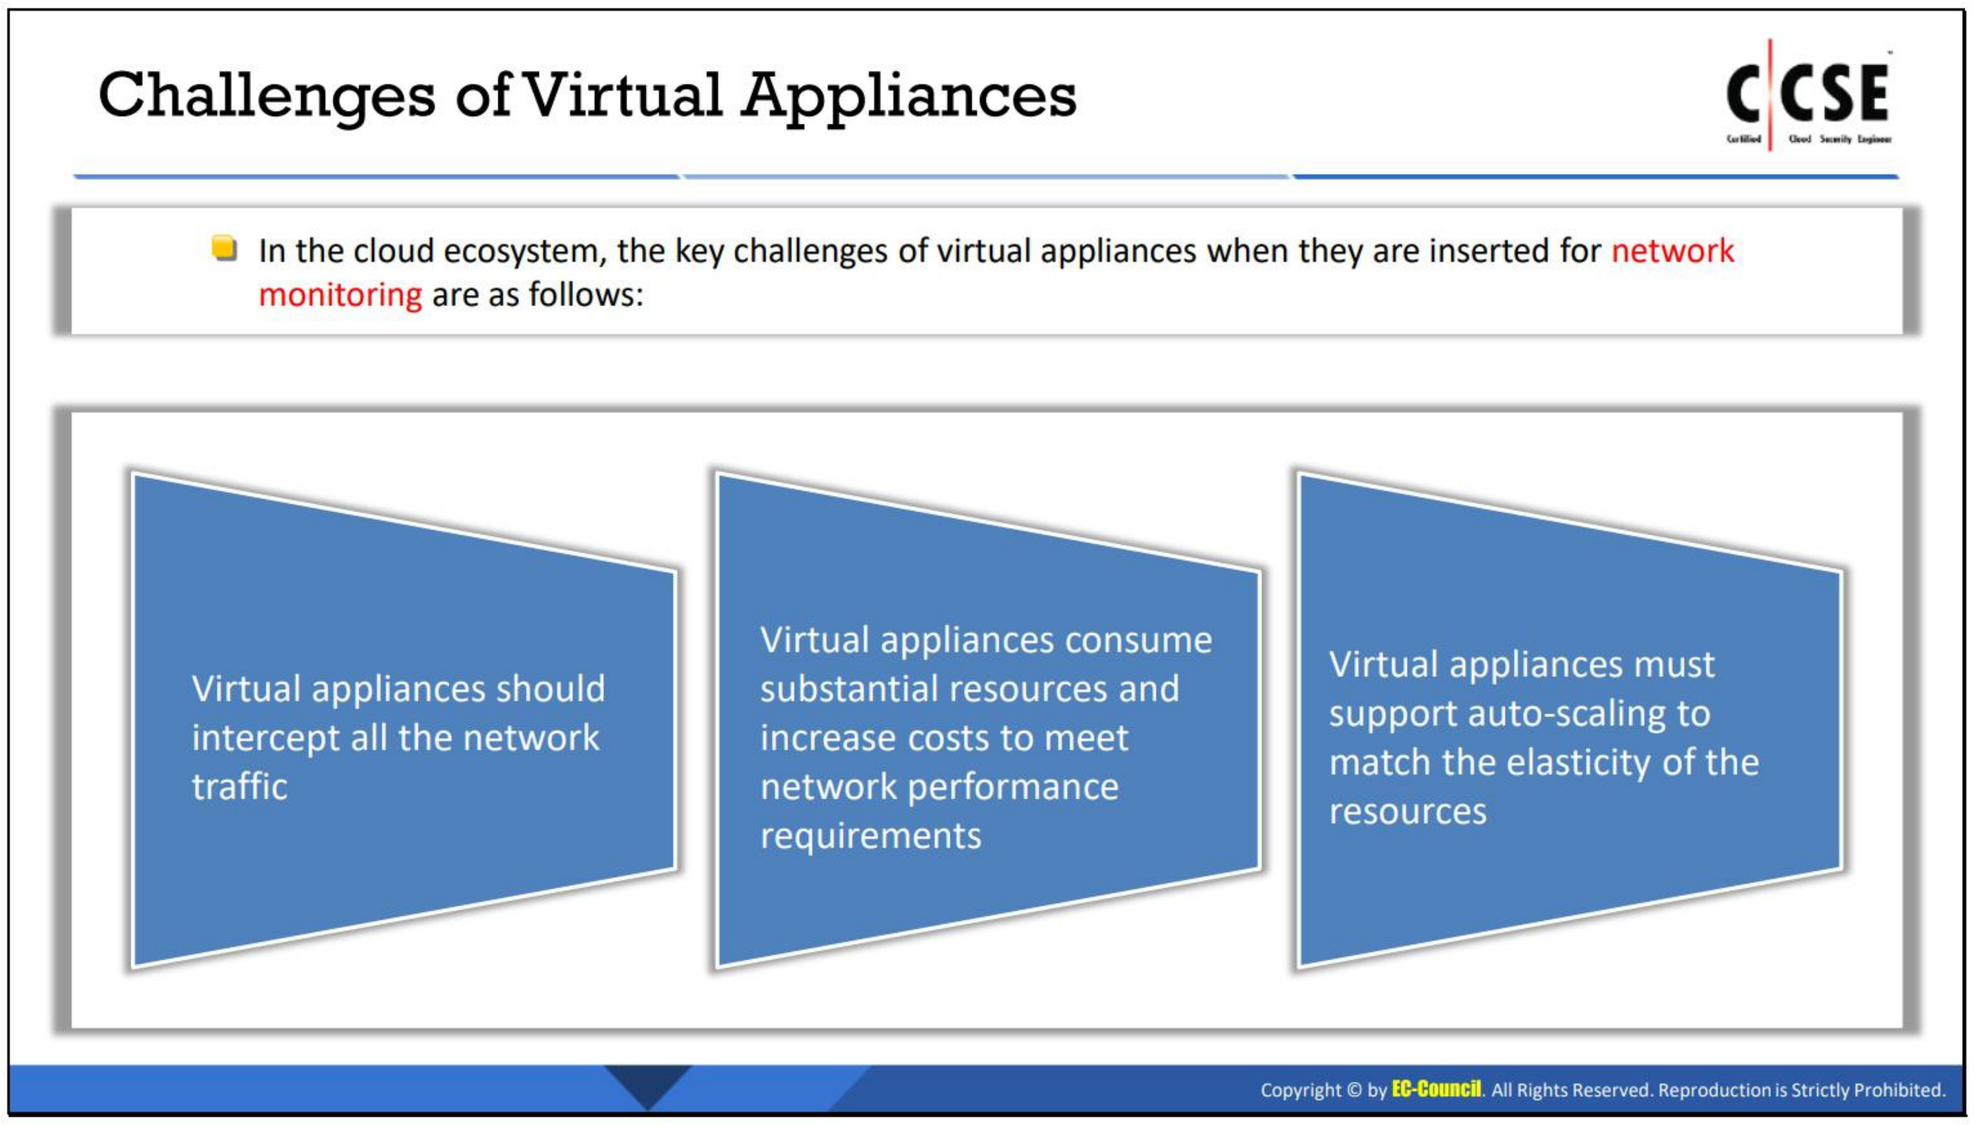
Certified Cloud Security Engineer Copyright © by **EC-Council**
All Rights Reserved. Reproduction is Strictly Prohibited.

by each customer. To mitigate against this threat, perform thorough research before selecting the appropriate CSP and pay more attention to any changes in the cloud policy.

- **Subpoena and e-Discovery**: The legal framework safeguarding privacy and processing of personal data is different in different countries. In most countries, exposure of private information is not a punishable offense. To overcome privacy-related issues, the cloud consumer must consider the legal framework of the cloud provider.

hide01.ir

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

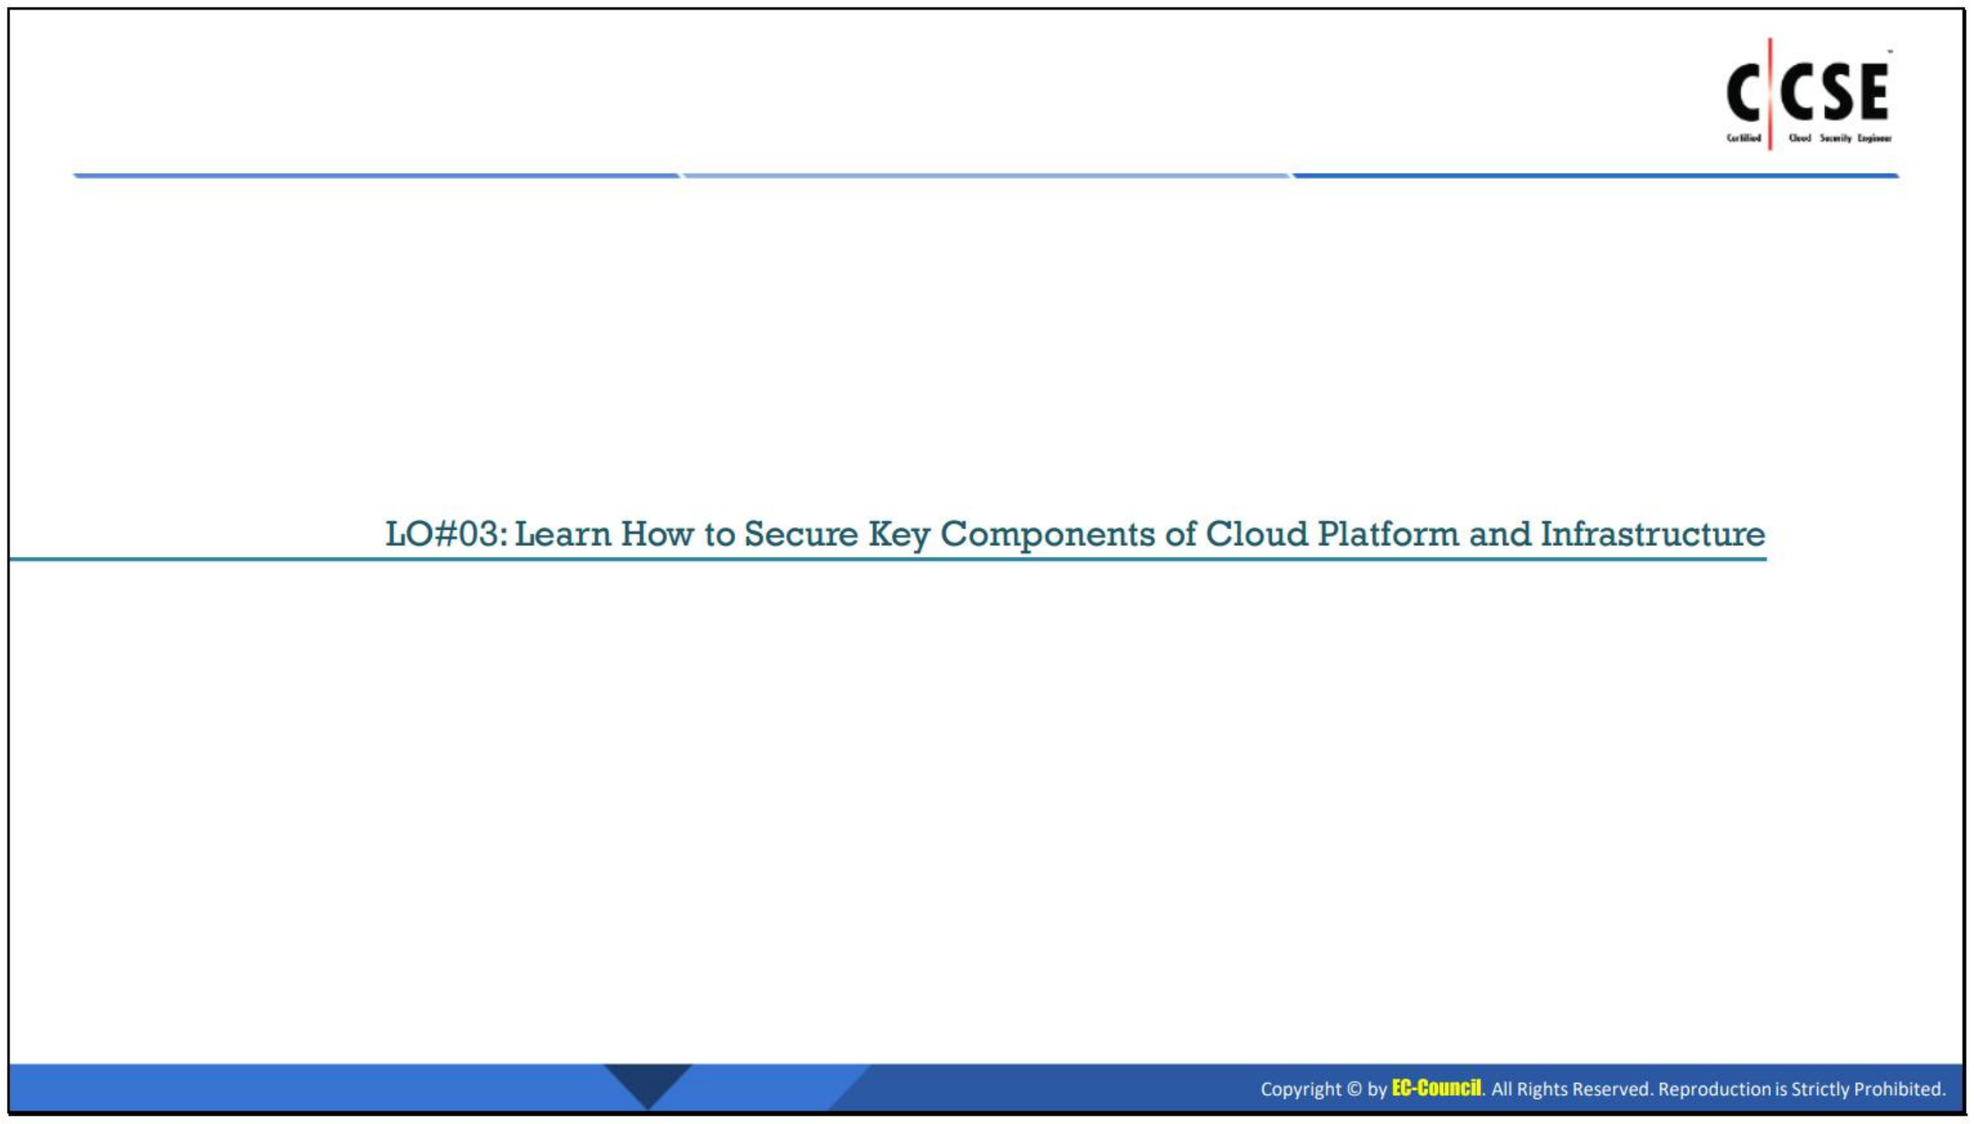


Challenges of Virtual Appliances

Following are the key virtual appliance challenges:

- A virtual appliance offers interception of bottlenecks that cannot fail to open.
- It requires resources to meet network performance requirements at a high cost.
- It supports auto-scaling to achieve elasticity in resource protection, which can cause an issue if a vendor does not support it.
- The virtual appliance should be well-versed with operating in the cloud ecosystem as well as the ability of the instances to move between zones in the cloud. The velocity of change in the cloud network is higher than physical network.
- It tends to improve resiliency due to auto-scaling in virtual servers.
- The running environment for virtual appliances differs subtly.
- Supporting different virtualization environments offers easy deployment of the appliance. The Open Virtualization Format is vendor-independent and allows the appliance to distribute a single file format.
- The appliance runs in a virtual environment that makes it crucial to ensure that the upgraded appliance is a more recent version of the target virtualization engine.
- Establishing completely independent underlying virtualization technology may not be possible.
- Establishing a degree of integration of the virtual appliance with the virtual server management system is complex.

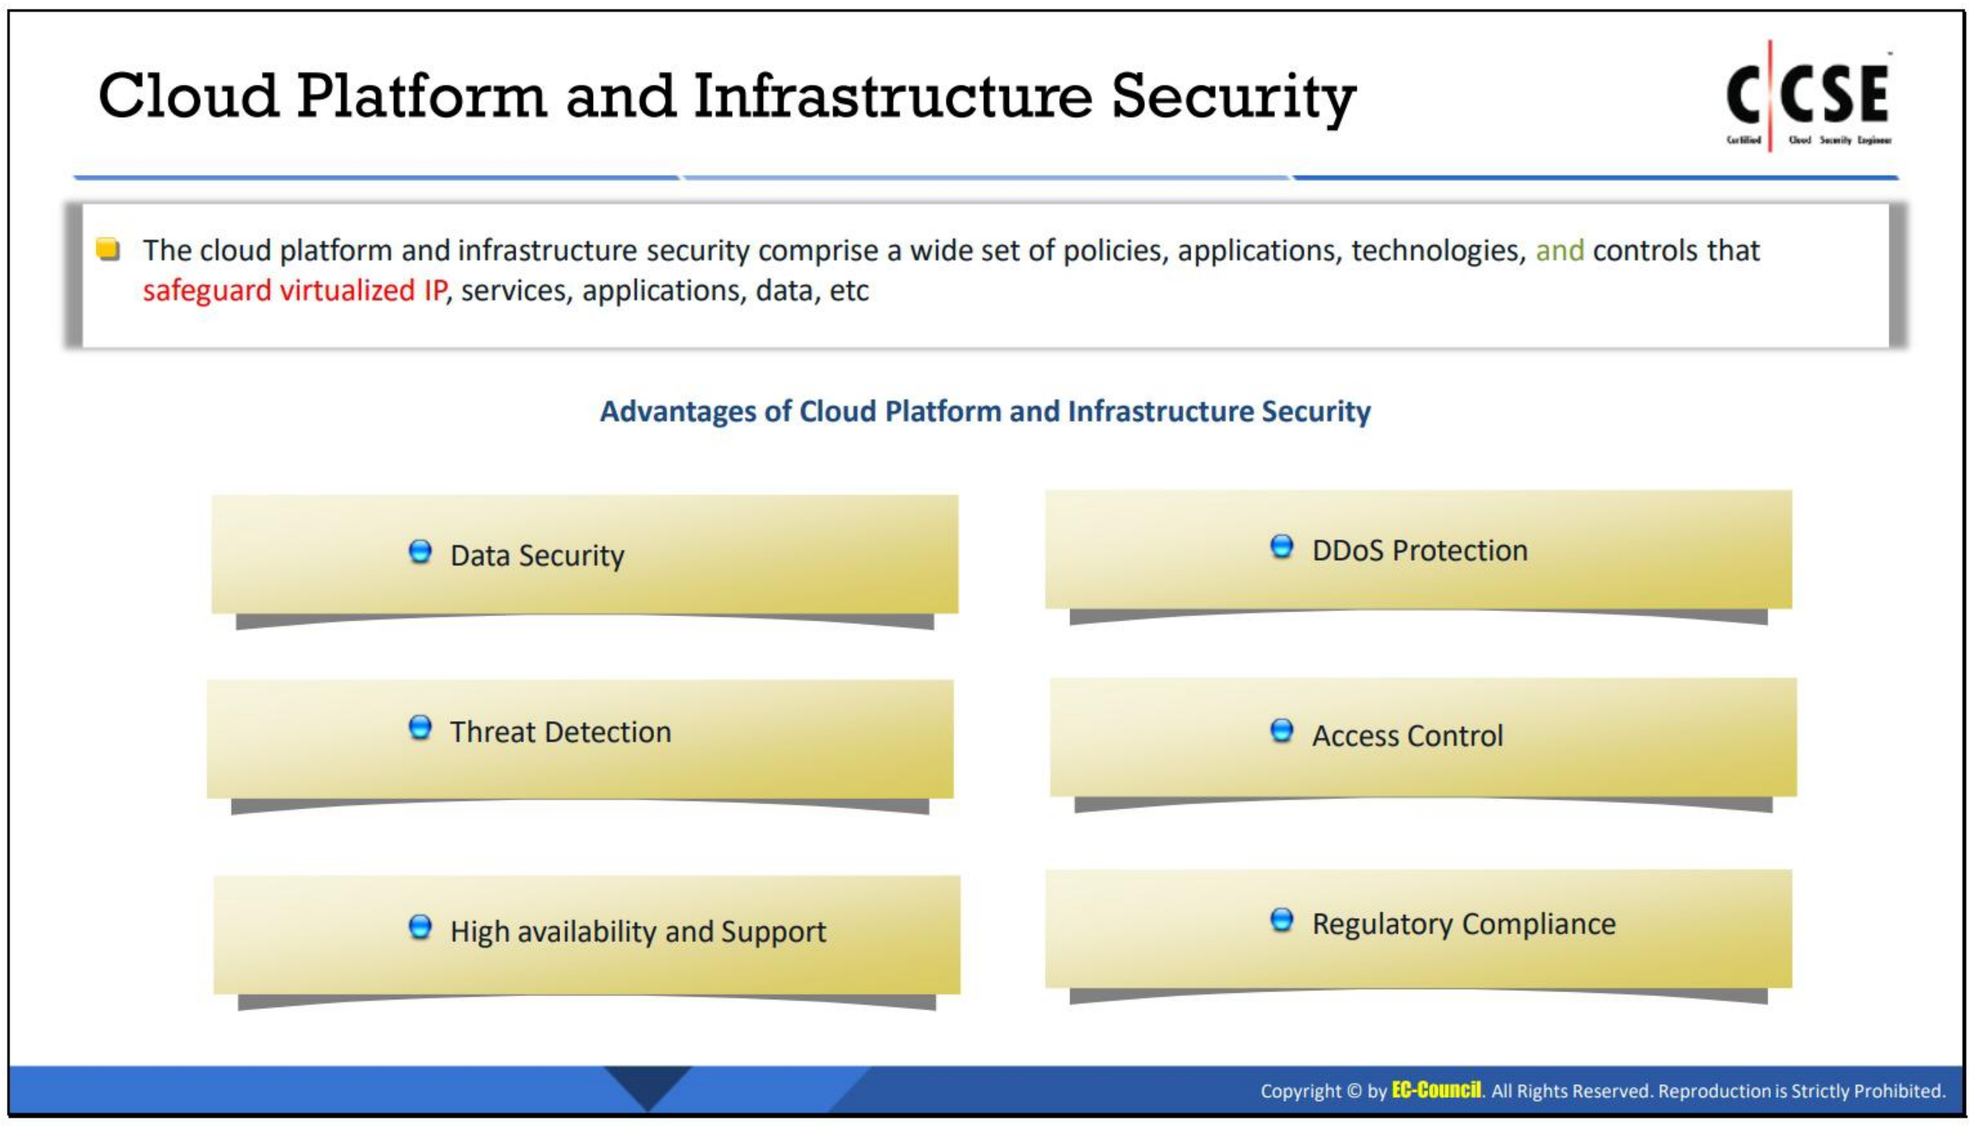
Certified Cloud Security Engineer Copyright © by **EC-Council**
All Rights Reserved. Reproduction is Strictly Prohibited.



LO#03: Learn How to Secure the Key Components of Cloud Platform and Infrastructure

To ensure that the applications and data in the cloud are secure, the cloud security professional needs to understand how to keep the various components of the cloud platform secure. This section deals with securing the cloud platform and infrastructure and shows how to secure the key components of cloud platforms and infrastructure. It will help you understand the importance of authentication, SDN benefits, identity, and access management (IAM), and user management features in detail.

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



Cloud Platform and Infrastructure Security

The most common threats to cloud platforms and infrastructure include data breaches, malware injections, DDoS, regulatory compliance, insecure APIs, advanced persistent threats (APTs), malicious insiders, account hijacking, etc. The cloud platform and infrastructure security comprise a wide set of policies, applications, technologies, and controls that safeguard virtualized IP, services, applications, data, etc. In recent years, most companies have been transferring their data and infrastructure to the cloud and the significance of cloud platforms and infrastructure security has become paramount. Advantages of cloud platform and infrastructure security are as follows:

Data security: Currently, cloud servers are easily susceptible to data breaches. Cloud platforms and infrastructure security protocols safeguard sensitive information and transactions. It will stop the third party from eavesdropping or altering the transmitted data.

DDoS protection: To safeguard against DDoS attacks, a cloud security solution focuses on security measures that prevents bulk traffic that targets the company's cloud server.

Threat detection: Advanced threat detection techniques such as endpoint scanning detect the threat at the device level. Such techniques improve the security of the devices that access the cloud network.

Access control: Organizations that utilize cloud storage may inadvertently leak data to the public, which can compromise their business integrity. Generally, these leaks do not occur due to malicious intent but due to poor security best practices. The best cloud policy is to implement access control on employees and allow access to only specific employees, which will make it difficult for the attacker to infiltrate and will prevent data leaks.

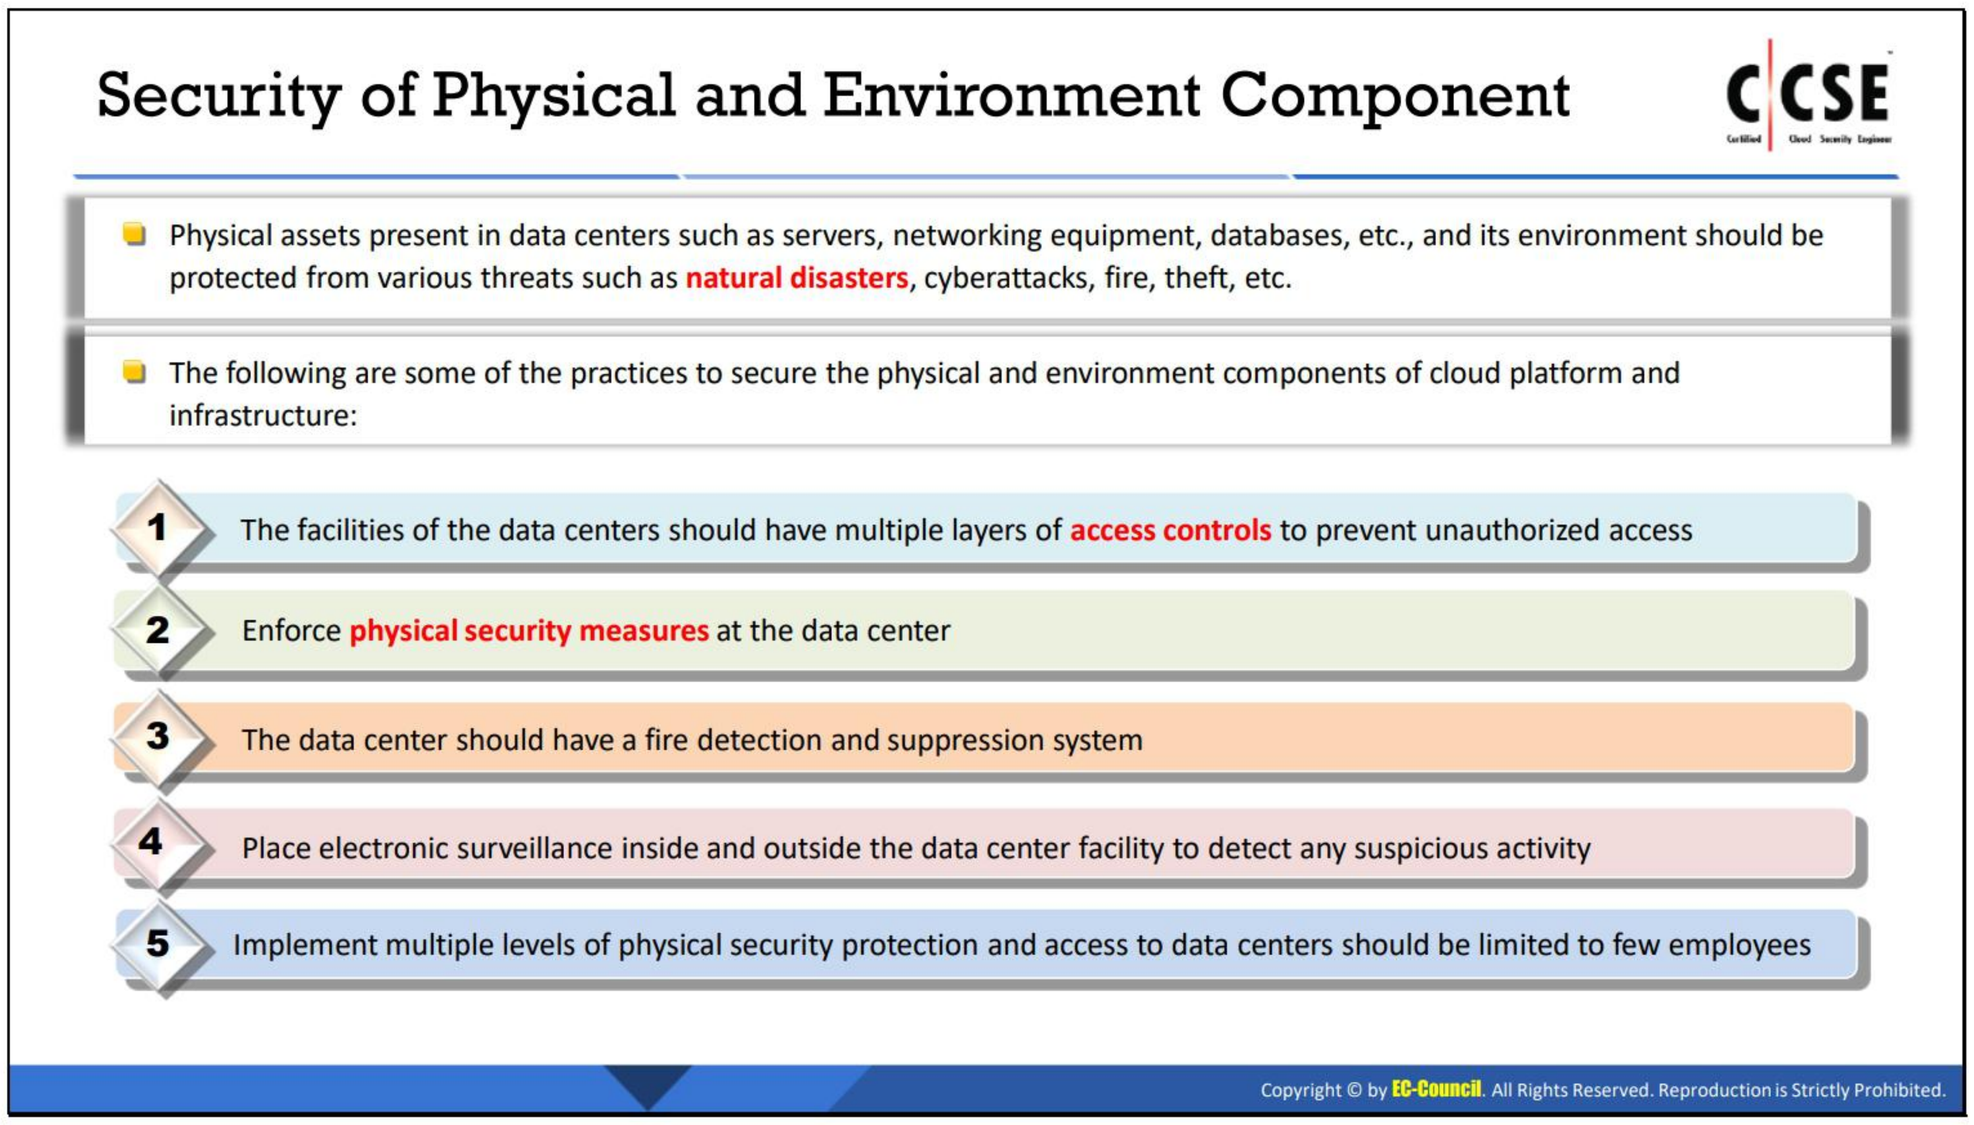
Certified Cloud Security Engineer Copyright © by **EC-Council**
All Rights Reserved. Reproduction is Strictly Prohibited.

High availability and support: A best-practice cloud security solution provides continuous support to the company's assets. It consists of 24/7 live monitoring throughout the year to ensure that the data is secure.

Regulatory Compliance: Cloud security solutions support companies in complying with the various industry regulations, especially those related to personal and financial data.

hide01.ir

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



Security of Physical and Environment Component

- The cloud platform and infrastructure should be build using innovative architecture and engineering approaches.
- Professional security staff should strictly monitor physical access using video surveillance, IDS, etc., at the perimeter and entry points.
- Physical access to the data center should be logged and audited regularly.
- The data center should have a fire detection and suppression system.
- The data center should have 24/7 power supply throughout the year.
- It should utilize climate control to constantly maintain the temperature of servers and hardware of the data center and prevent overheating and service outage.
- After a storage device reaches its end of life, the data center should have a decommissioning process to safeguard customer data from exposure to an unauthorized person.
- A physical security review should be conducted regularly to make sure that the data center addresses security issues appropriately.
- The data center should meet international as well as industry-specific compliance standards. The data center should have multiple levels of physical security protection and access to these data centers should be limited to a few employees. Technologies such as biometric identification, metal detection, etc. should be used.
- The data center should have an automated system to make sure that the servers run updated versions of software stacks with the latest security patches, identify and diagnose hardware and software problems, and remove the machine from the service if necessary.

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

Security of Management Component

- To secure the management component, cloud providers implement their own identity and access management (IAM) for identification, authentication, and authorization for accessing cloud resources
- Providers and customers should implement least privileges for **meta structure access**
- Accounts should implement **multi-factor authentication**

A Good IAM Should Have

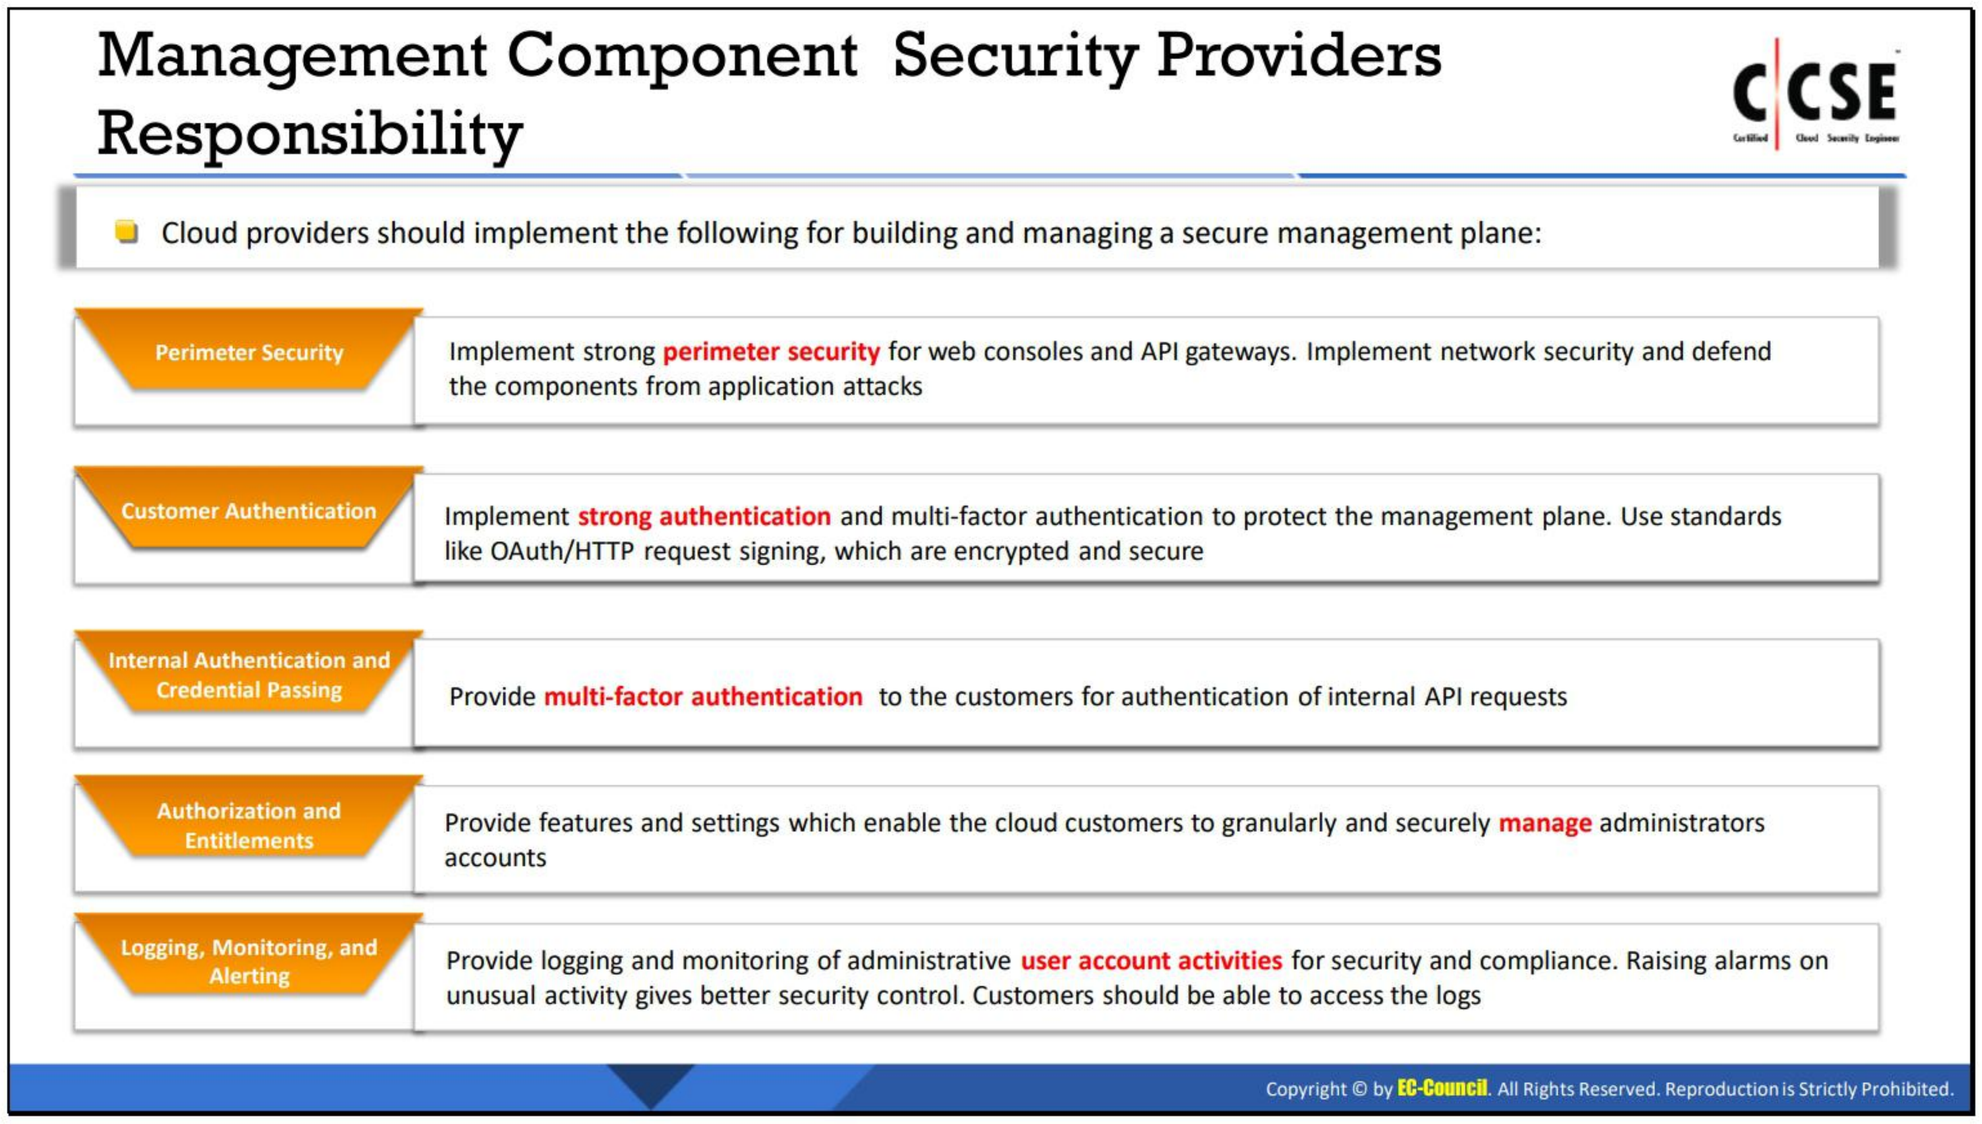
- **Root/master Account:** the enterprise owned account to configure and manage the cloud resources with **super-admin privileges**. Master Account should be tightly locked down and should not be used often
- **Super-admin Account:** for individual **admin users;** should be used sparingly
- **Service Administrator Accounts:** to manage services with **restricted rights** specific to operations

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

Security of the Management Component

- By using identity and access management (IAM), the client can create multiple users and manage the permissions of an individual user with regard to the cloud resources. IAM helps in the implementation of the principle of least privilege, which safeguards the cloud platform and infrastructure from unauthorized access.
- This component provides cloud infrastructure and perimeter security.
- The management component of the cloud platform and infrastructure helps in monitoring cloud resources. It delivers a log of users and system access. With every recorded event, the client can see what services were accessed and actions performed on those services, who has made the request, etc.

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



Management Component Security Providers Responsibility

When an organization uses a cloud service, it is responsible for configuring the management plane's parts that are exposed to it by the provider. However, when the organization acts as a cloud provider and needs to build and manage the cloud (e.g., in a private cloud deployment), it is responsible for everything.

Cloud providers should implement the following for building and managing a secure management plane.

- **Perimeter security**: CSPs must employ perimeter security to prevent attacks that target the management plane's components (web and API servers). They also must implement strong perimeter security for web consoles and API gateways. It includes implementing network security and defending the components against application attacks.
- **Customer authentication**: CSPs should implement strong authentication and multi-factor authentication mechanisms to protect the management plane. They must use standards (such as OAuth/HTTP request signing), which are well-documented, cryptographically valid, and secure.
- **Internal authentication and credential passing**: Internal authentication includes the mechanisms used by employees of an organization to connect to the management plane's non-customer-facing portions, including translations between any internal API requests and the customer's authentication. Enforce multi-factor authentication on the customers for authentication of internal API requests.
- **Authorization and entitlements**: CSPs must consider the entitlements for internal administrators and the entitlements that are for customers. Provide features and settings

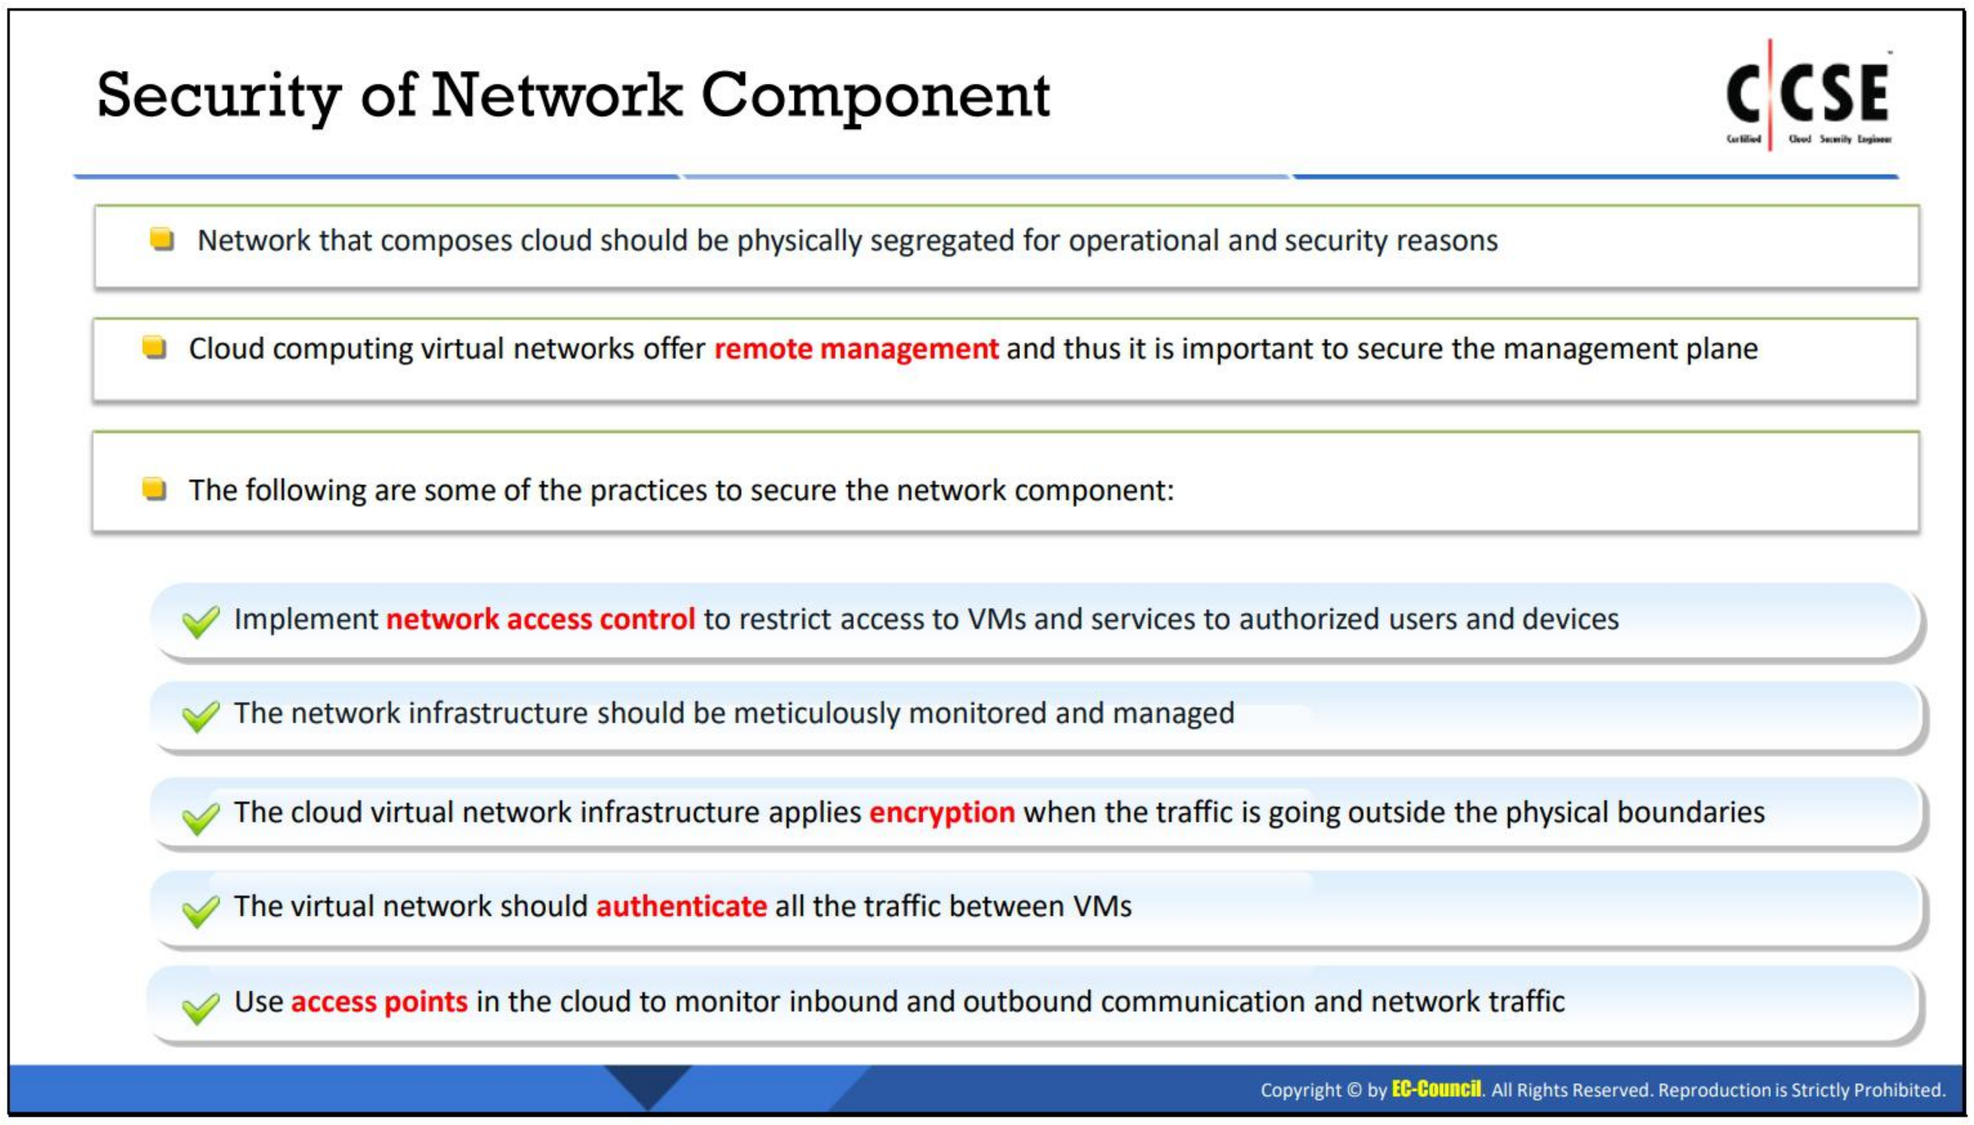
Certified Cloud Security Engineer Copyright © by **EC-Council**
All Rights Reserved. Reproduction is Strictly Prohibited.

that enable the cloud customers to granularly and securely manage administrators' accounts. With granular entitlements, customers can securely manage their administrators and users, and reduce the impact of compromised administrator accounts.

- **Logging, monitoring, and alerting**: CSPs must provide logging and monitoring of administrative user account activities for security and compliance. Robust monitoring and logging apply to what employees do as part of their routine service management as well as what customers do in their accounts. Raising alarms on unusual activity gives better security control. Alerting is a key security control that helps make monitoring actionable and effective. Customers should be able to access their activity logs on the cloud platform using APIs or some other mechanisms.

hide01.ir

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

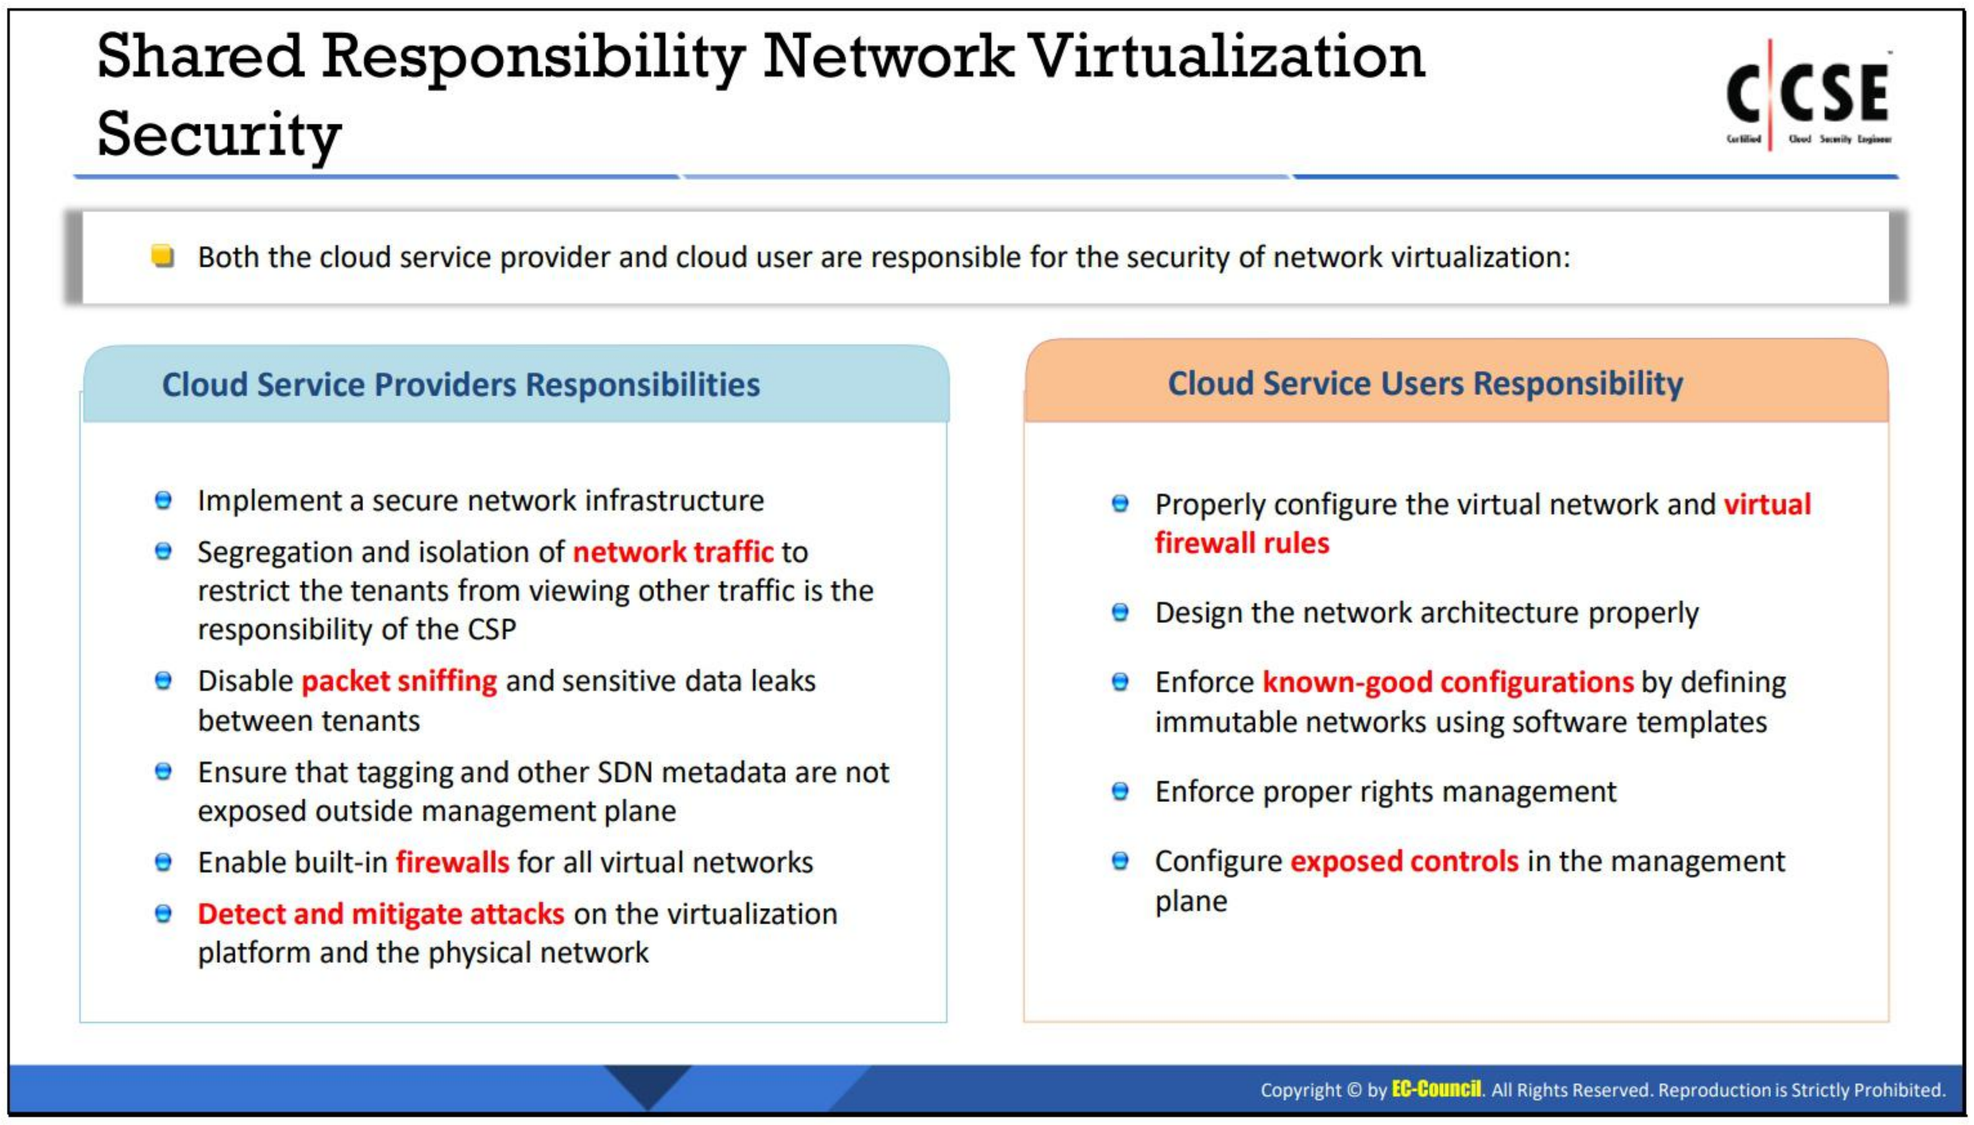


Security of Network Component

- The network infrastructure should be meticulously monitored and managed.
  - Network devices, firewall, and other boundary devices should monitor and manage communication within the network.
  - Implement configuration, set of rules, and access control lists (ACLs) to regulate traffic flow for external and internal devices.
  - Use access points in the cloud to monitor inbound and outbound communication and network traffic.
  - The client can connect to the access point through HTTP or HTTPS utilizing secure socket layer (SSL), which would safeguard against message forgery, tampering, and eavesdropping.
  - Build data centers in various geographical locations so that in case of failure the automated process moves the client data traffic away from the affected area.
  - Utilize automated monitoring systems and tools to monitor network and server utilization, application usage, port scanning activities, and unauthorized access attempts.
  - Implement network access control to restrict access to VMs and services to authorized users and devices.
  - The cloud virtual network infrastructure should apply encryption when the traffic is going outside the physical boundaries. The encryption should be performed at the network layer and should apply to private IP traffic. The virtual network should authenticate all the traffic between VMs.

hidebir

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



Shared Responsibility Network Virtualization Security

The CSP and cloud user are responsible for the security of network virtualization:

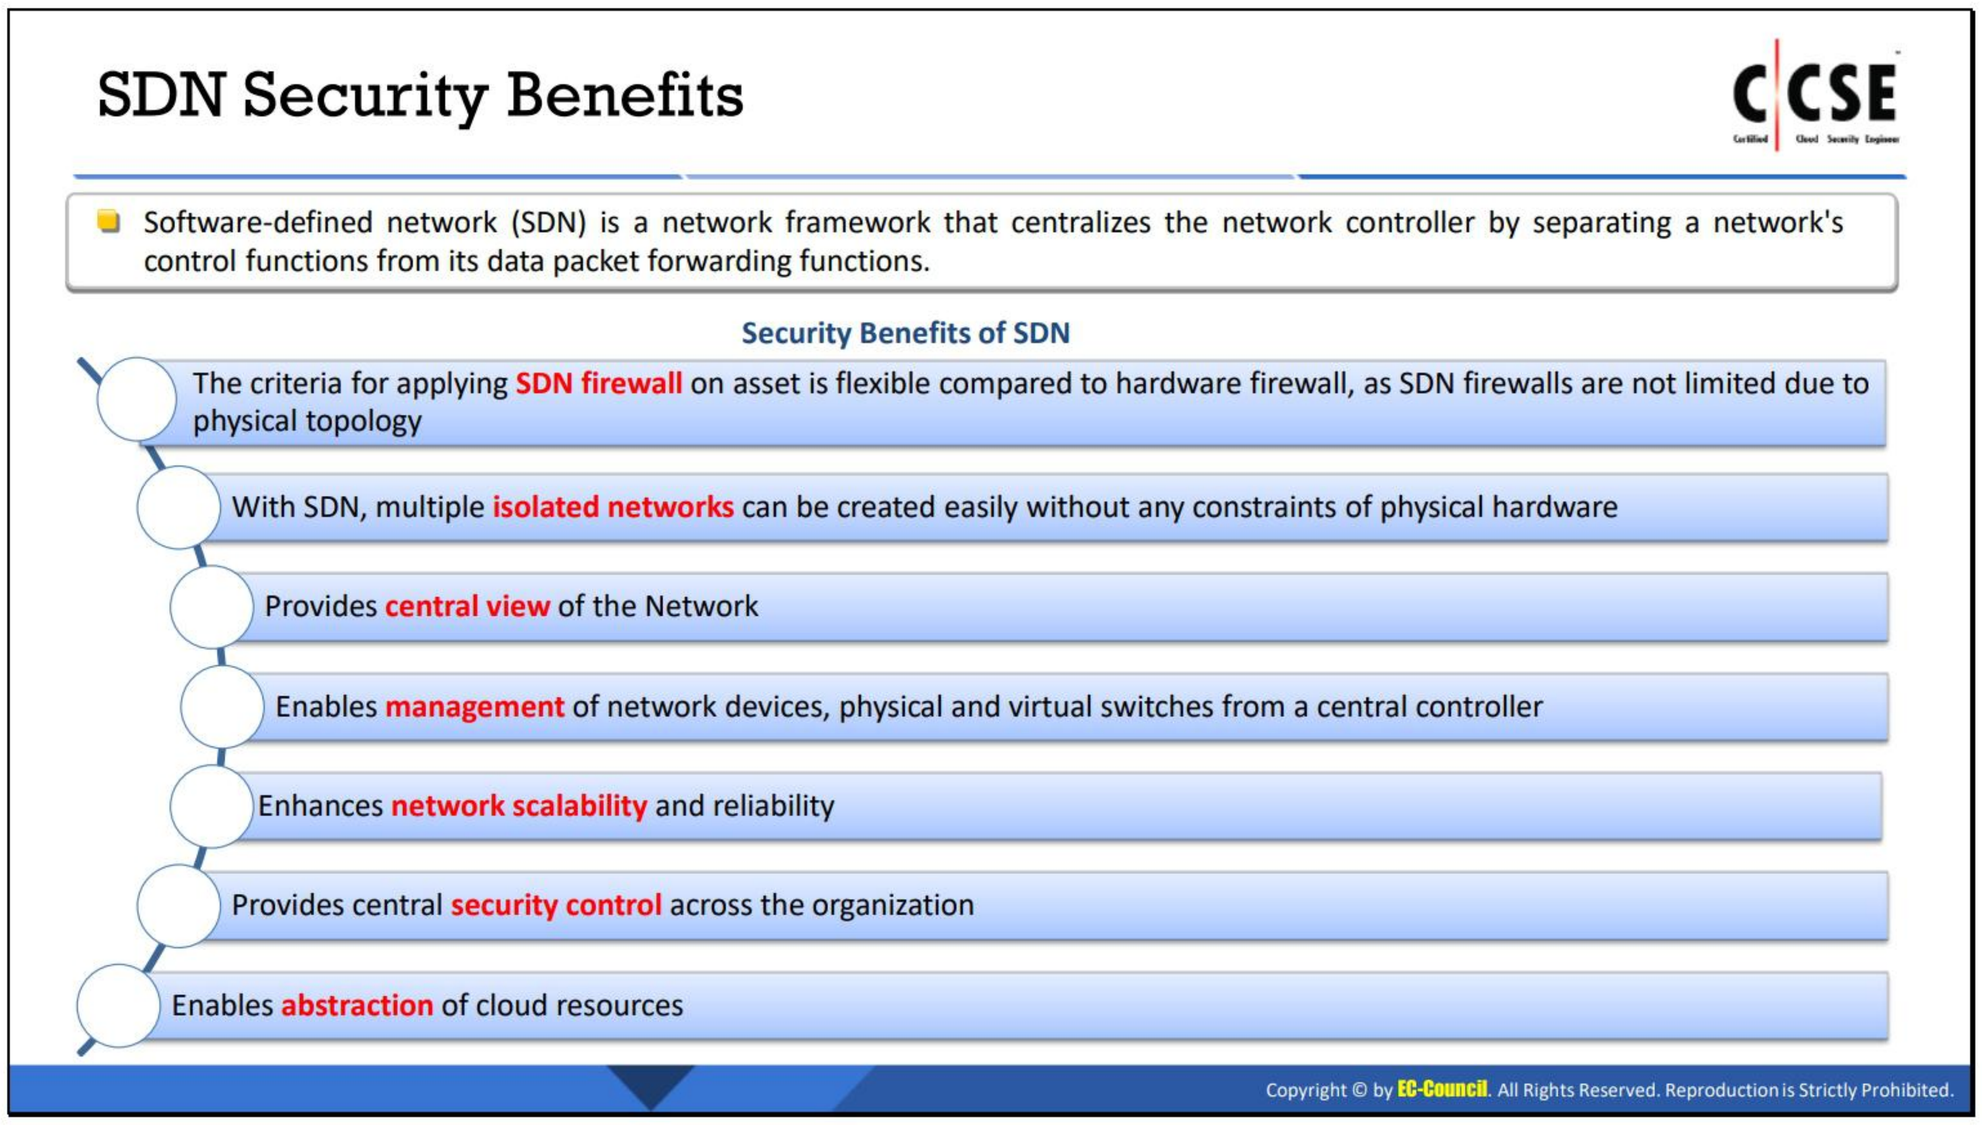
Cloud Service Providers Responsibilities

- Implement a secure network infrastructure
- Segregation and isolation of network traffic to restrict the tenants from viewing other traffic
- Disable packet sniffing and sensitive data leaks between tenants
- Ensure that tagging and other SDN metadata are not exposed outside management plane
- Enable built-in firewalls for all virtual networks
- Detect and mitigate attacks on the virtualization platform and the physical network

Cloud Service Users Responsibilities

- Properly configure the virtual network and virtual firewall rules
- Design the network architecture properly
- Enforce known-good configurations by defining immutable networks using software templates
- Enforce proper rights management
- Configure exposed controls in the management plane

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



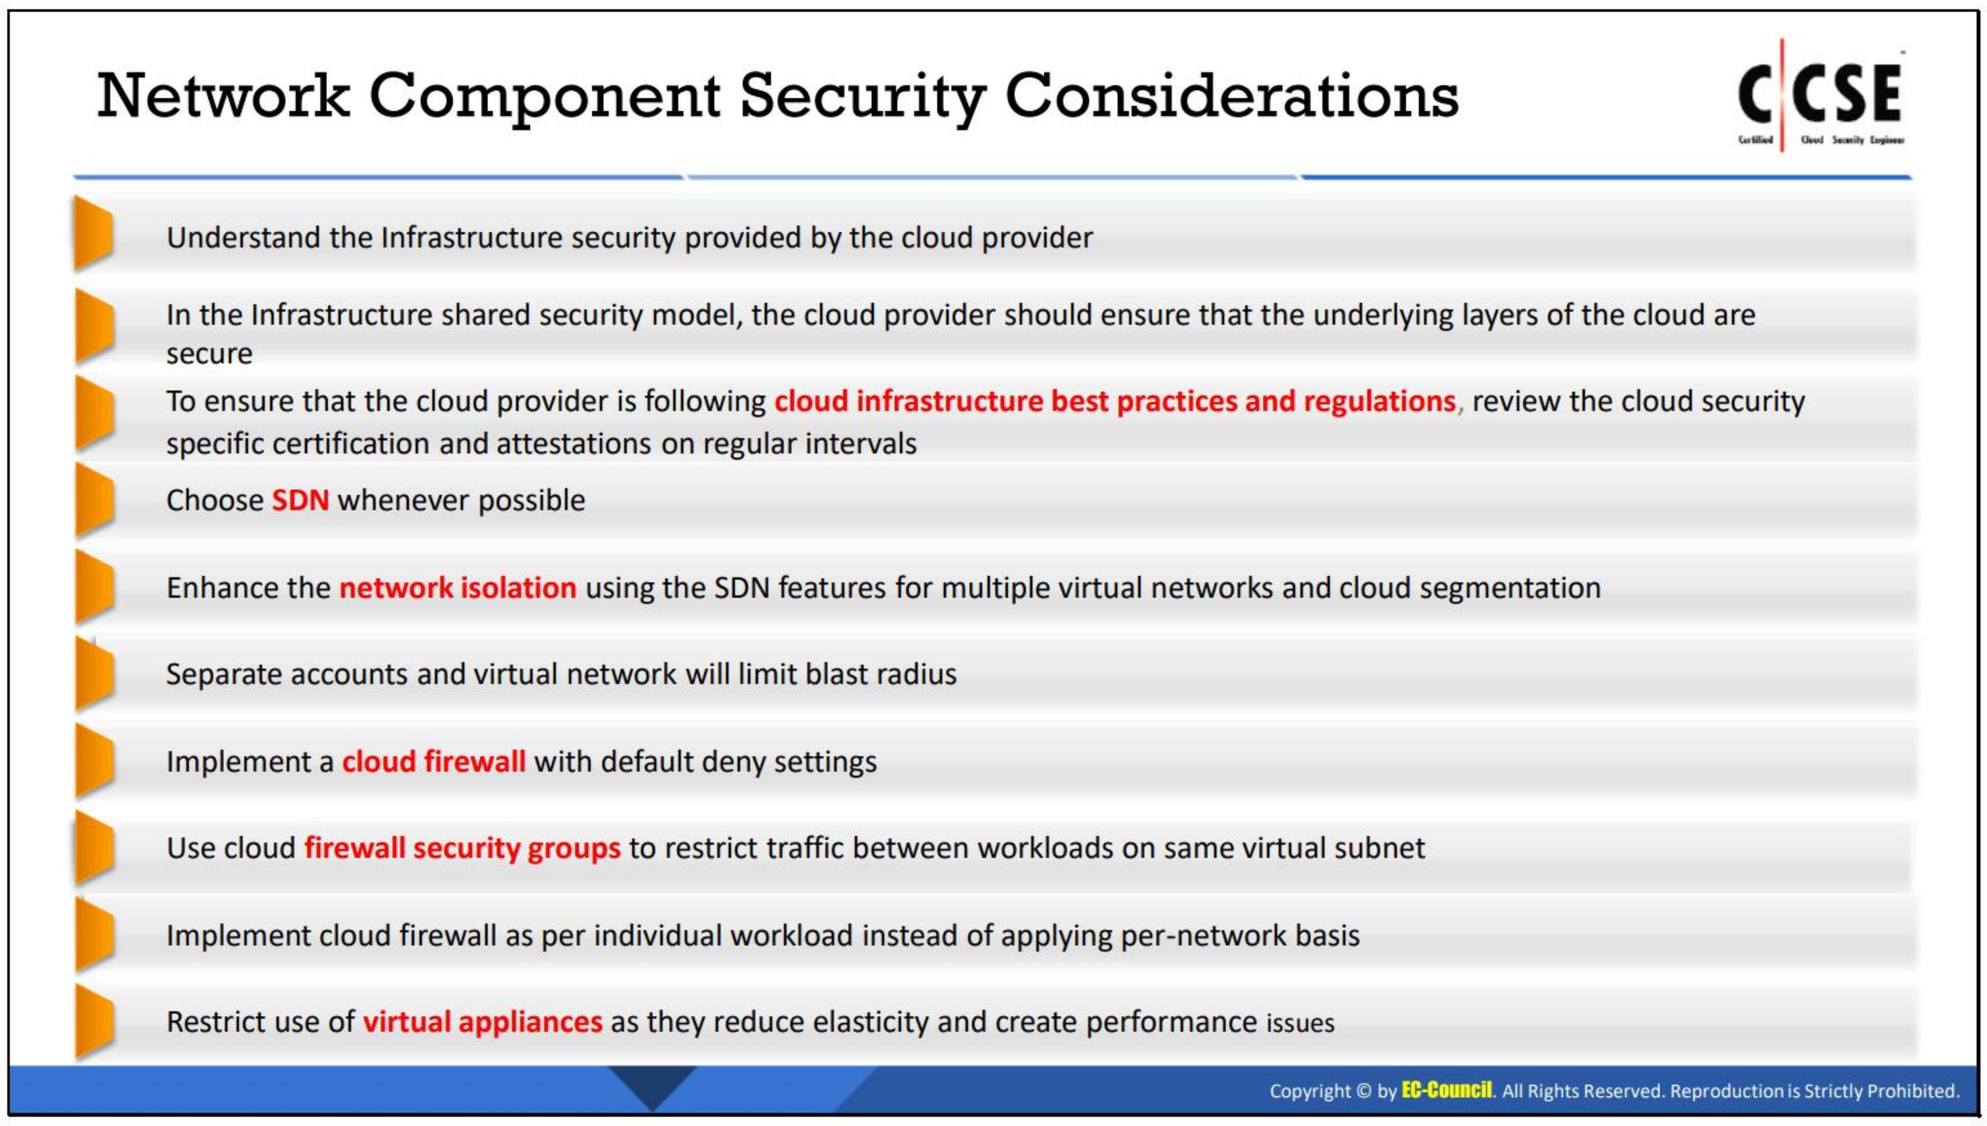
SDN Security Benefits

SDN security benefits include the following:

- An administrator can change network switch rules according to prioritizing, de-prioritizing, or blocking packets with control and security. It enables a multitenant architecture for the administrator to manage the traffic load.
- It provides network management and visibility. Network administrator deals with a centralized controller to distribute policies on connected switches instead of configuring multiple individual devices. The controller can monitor traffic and deploy security policies if it deems traffic suspicious.
- It implements hardware and service virtualization to reduce hardware footprint and lower operational costs.
- It enables software-defined wide area network technology to abstract an organization's connectivity links to WAN by creating a virtual network that connects controller deems fit to traffic.
- Service virtualization offers a reduction in reliance on costly hardware.
- It provides many isolation networks without the constraints of physical hardware.
- It enables dynamic combinations and policies, which require less management than traditional hardware or host-based approach.

h1de01.ir

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



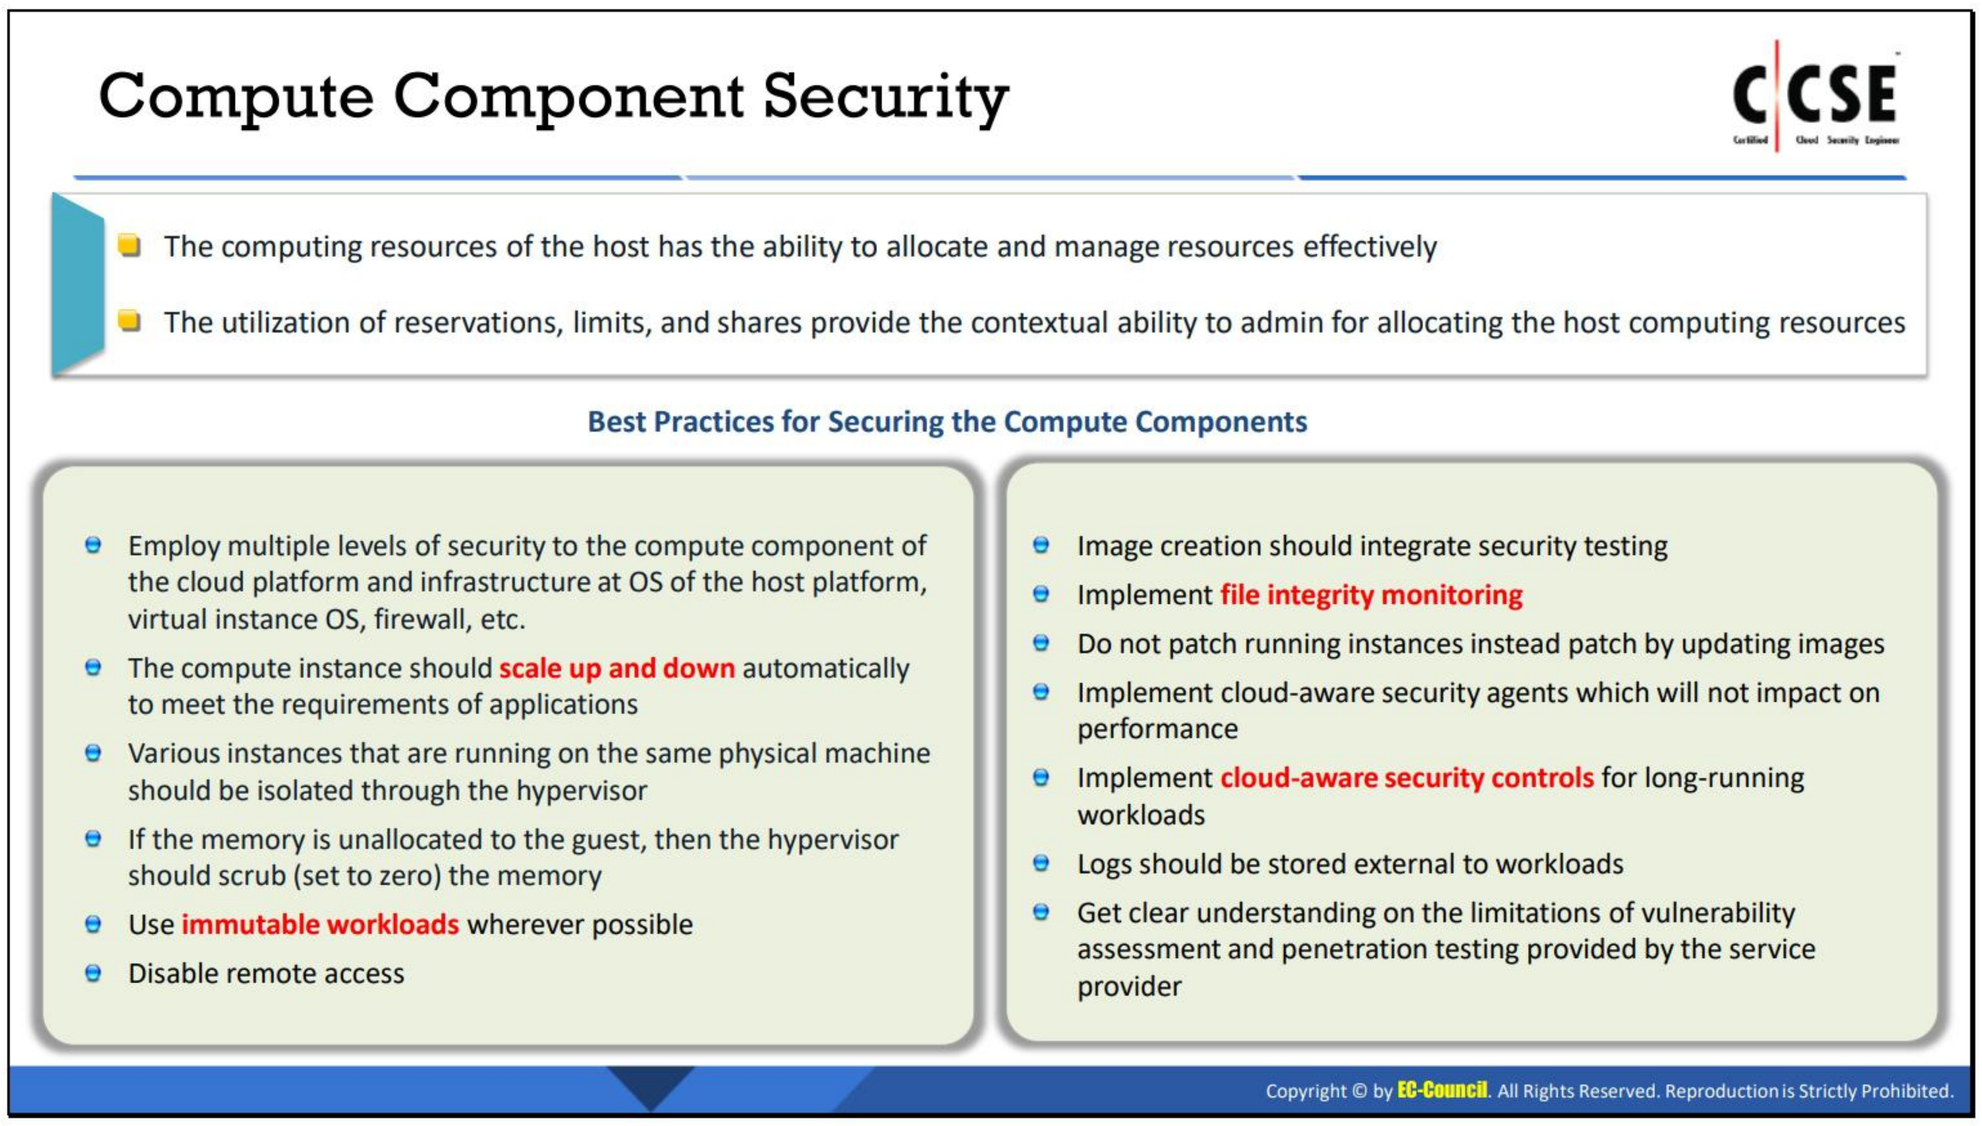
Network Component Security Considerations

Organizations using the cloud must follow these network component security considerations to ensure secure cloud environments:

hidoo1.ir

- Understand the infrastructure security provided by their cloud providers.
- In a shared security model, the cloud provider should ensure that the underlying layers (e.g., physical, orchestration, and abstraction layers) of the cloud are secure.
- To ensure that the cloud provider is following cloud infrastructure best practices and regulations, customers must review the cloud security specific certifications and attestations on regular intervals.
- Choose SDN whenever possible.
- Enhance the network isolation using the SDN features for multiple virtual networks and cloud segmentation.
- Separate accounts and virtual networks will limit the blast radius (i.e., the possible extent of impact of a security event).
- Implement a cloud firewall with default deny settings. Organizations can allow IP addresses that they consider safe by using firewall settings.
- Use cloud firewall security groups to restrict traffic between workloads on the same virtual subnet.
- Implement cloud firewall as per individual workloads instead of applying on a per-network basis.
- Restrict the use of virtual appliances as they reduce elasticity and create performance issues.

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

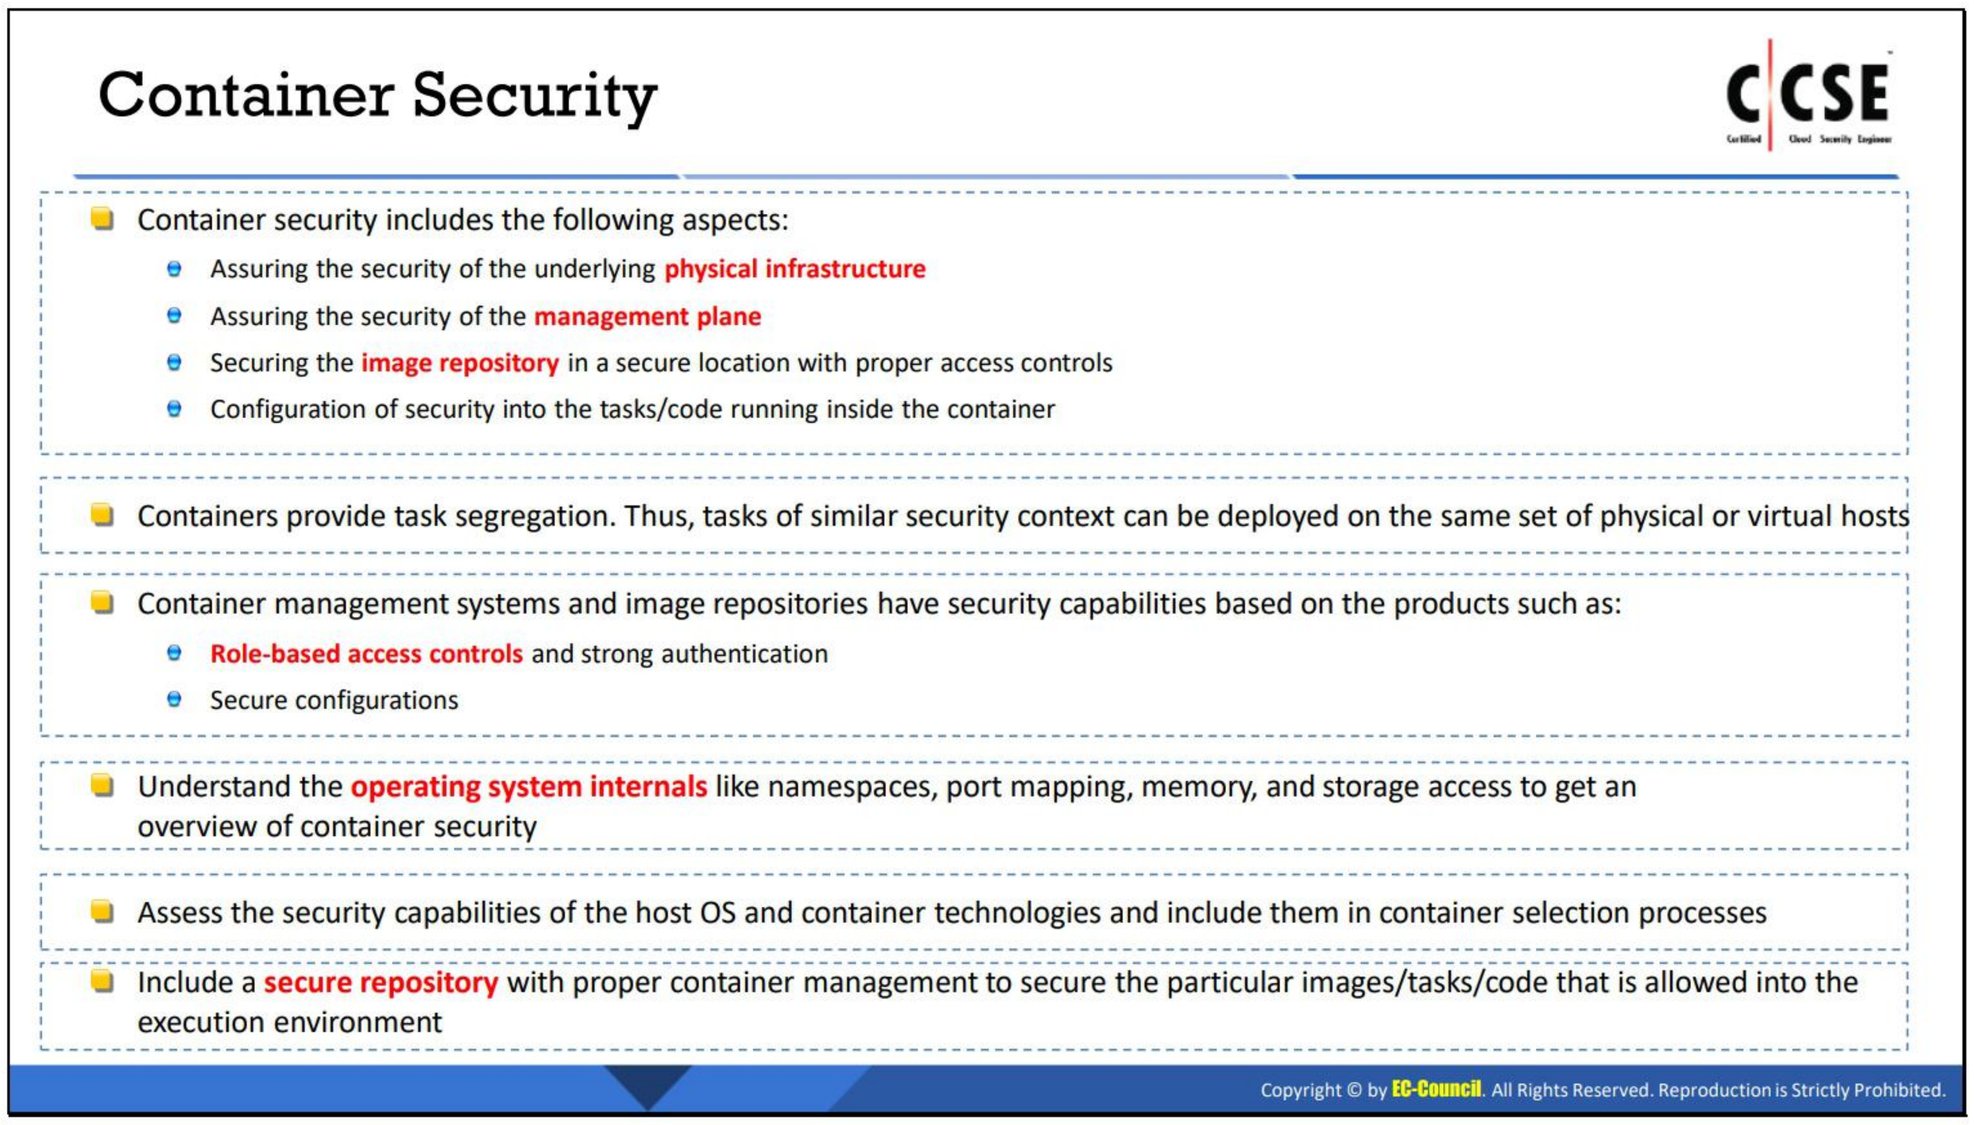


Compute Component Security

- The compute instance should scale up and down automatically to meet the requirements of applications.
- Employ multiple levels of security to the compute component, including at the host platform OS, virtual instance OS, firewall, etc.
- Various instances that are running on the same physical machine should be isolated through the hypervisor.
- Automatically reset every storage block utilized by the client to prevent unintentional exposure of data of one client to another.
- If the memory is unallocated to the guest, then the hypervisor should scrub (set to zero) the memory. Before scrubbing, the memory should not be allowed to move to the pool of free memory available for allocation.
- Enforce multi-factor authentication to access the management plane.
- Along with the security provided by the CSP, the client should run an encrypted file system on the virtualized disk device.

hide01.ir

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



Container Security

Containers provide task segregation. Thus, tasks of similar security context can be deployed on the same set of physical or virtual hosts.

hide01.ir

Container management systems and image repositories have security capabilities based on products such as:

- Role-based access controls and strong authentication
- Secure configurations

Understand OS internals such as namespaces, port mapping, memory, and storage access to get an overview of container security.

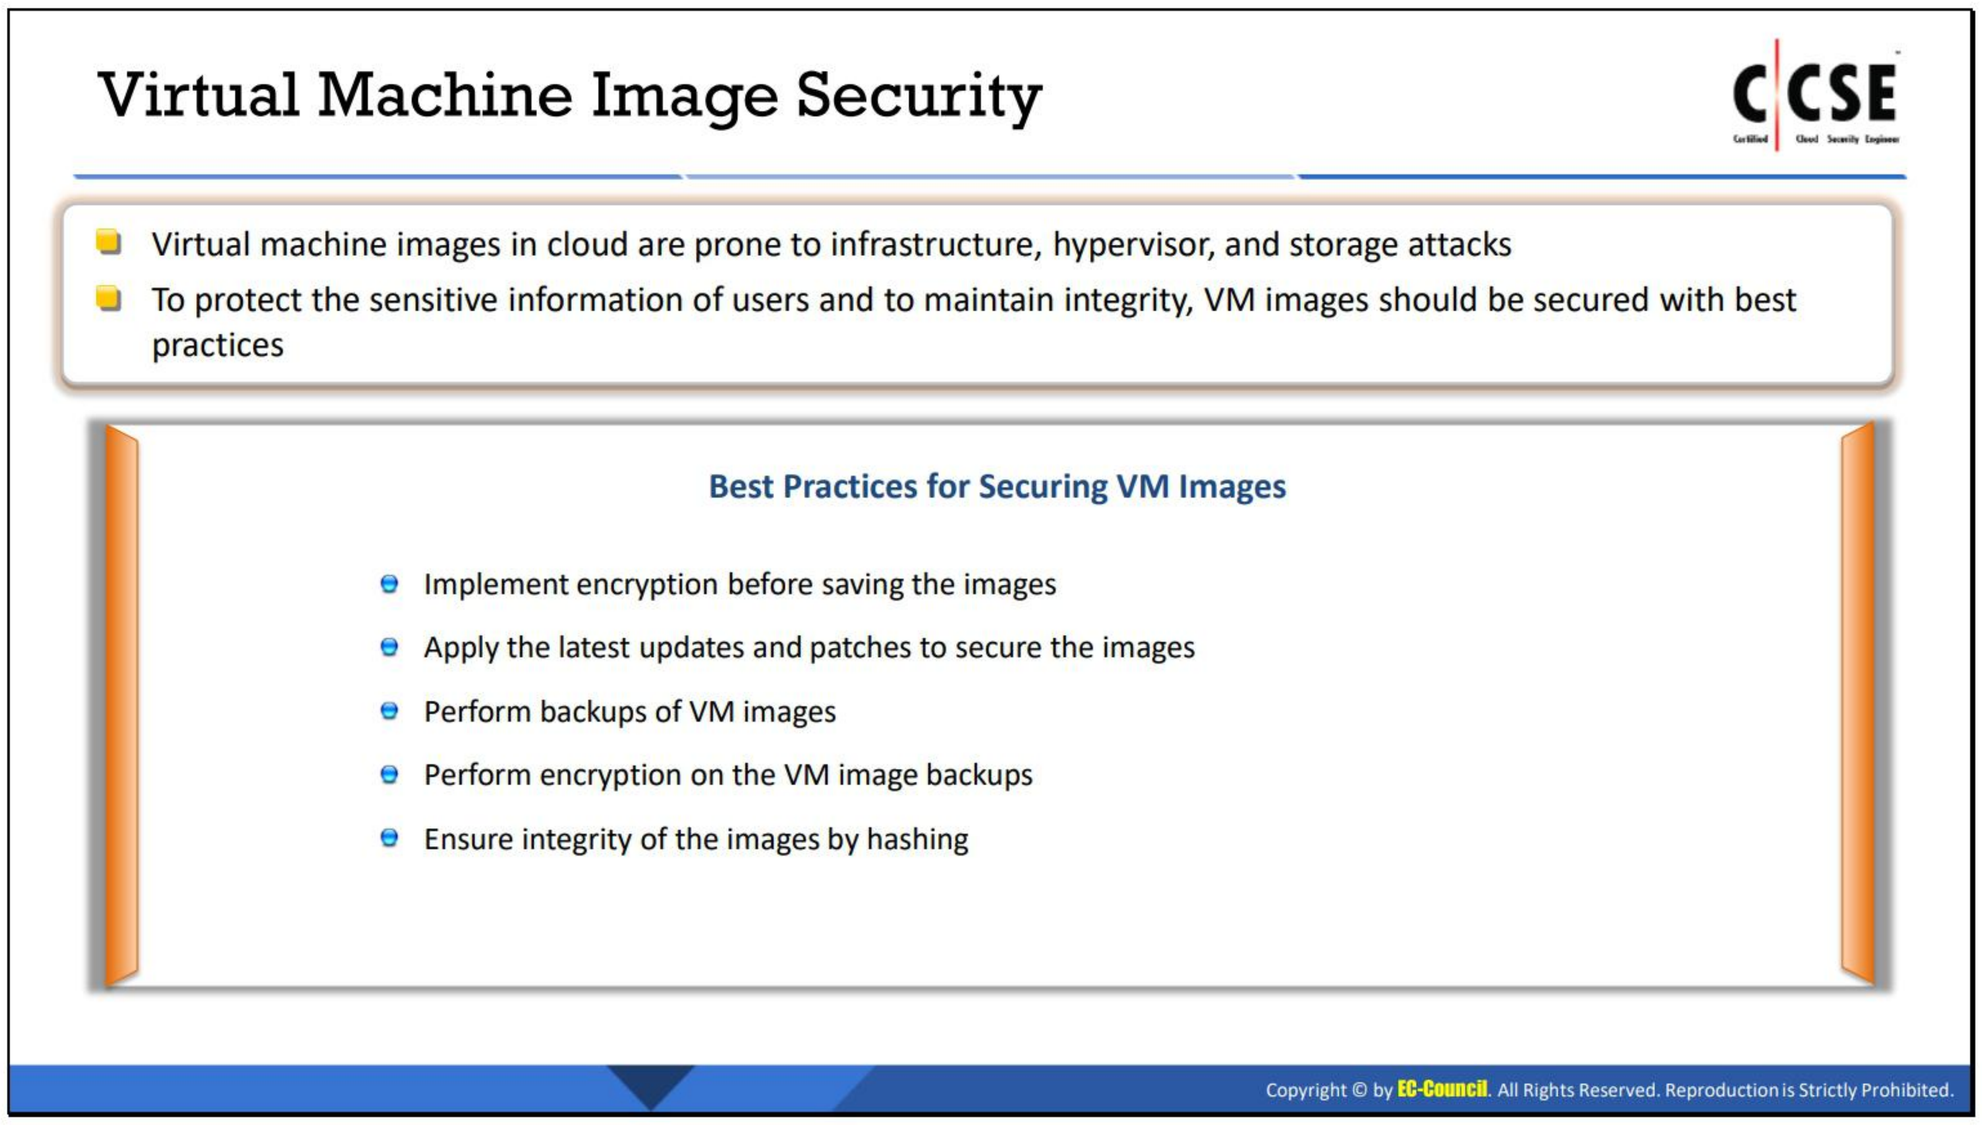
Assess the security capabilities of the host OS and container technologies and include them in container selection processes.

Include a secure repository with proper container management to secure the particular images/tasks/code that is allowed into the execution environment.

Container security includes the following aspects:

- Assuring the security of the underlying physical infrastructure
- Assuring the security of the management plane
- Securing the image repository in a secure location with proper access controls
- Configuration of security into the tasks/code running inside the container

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



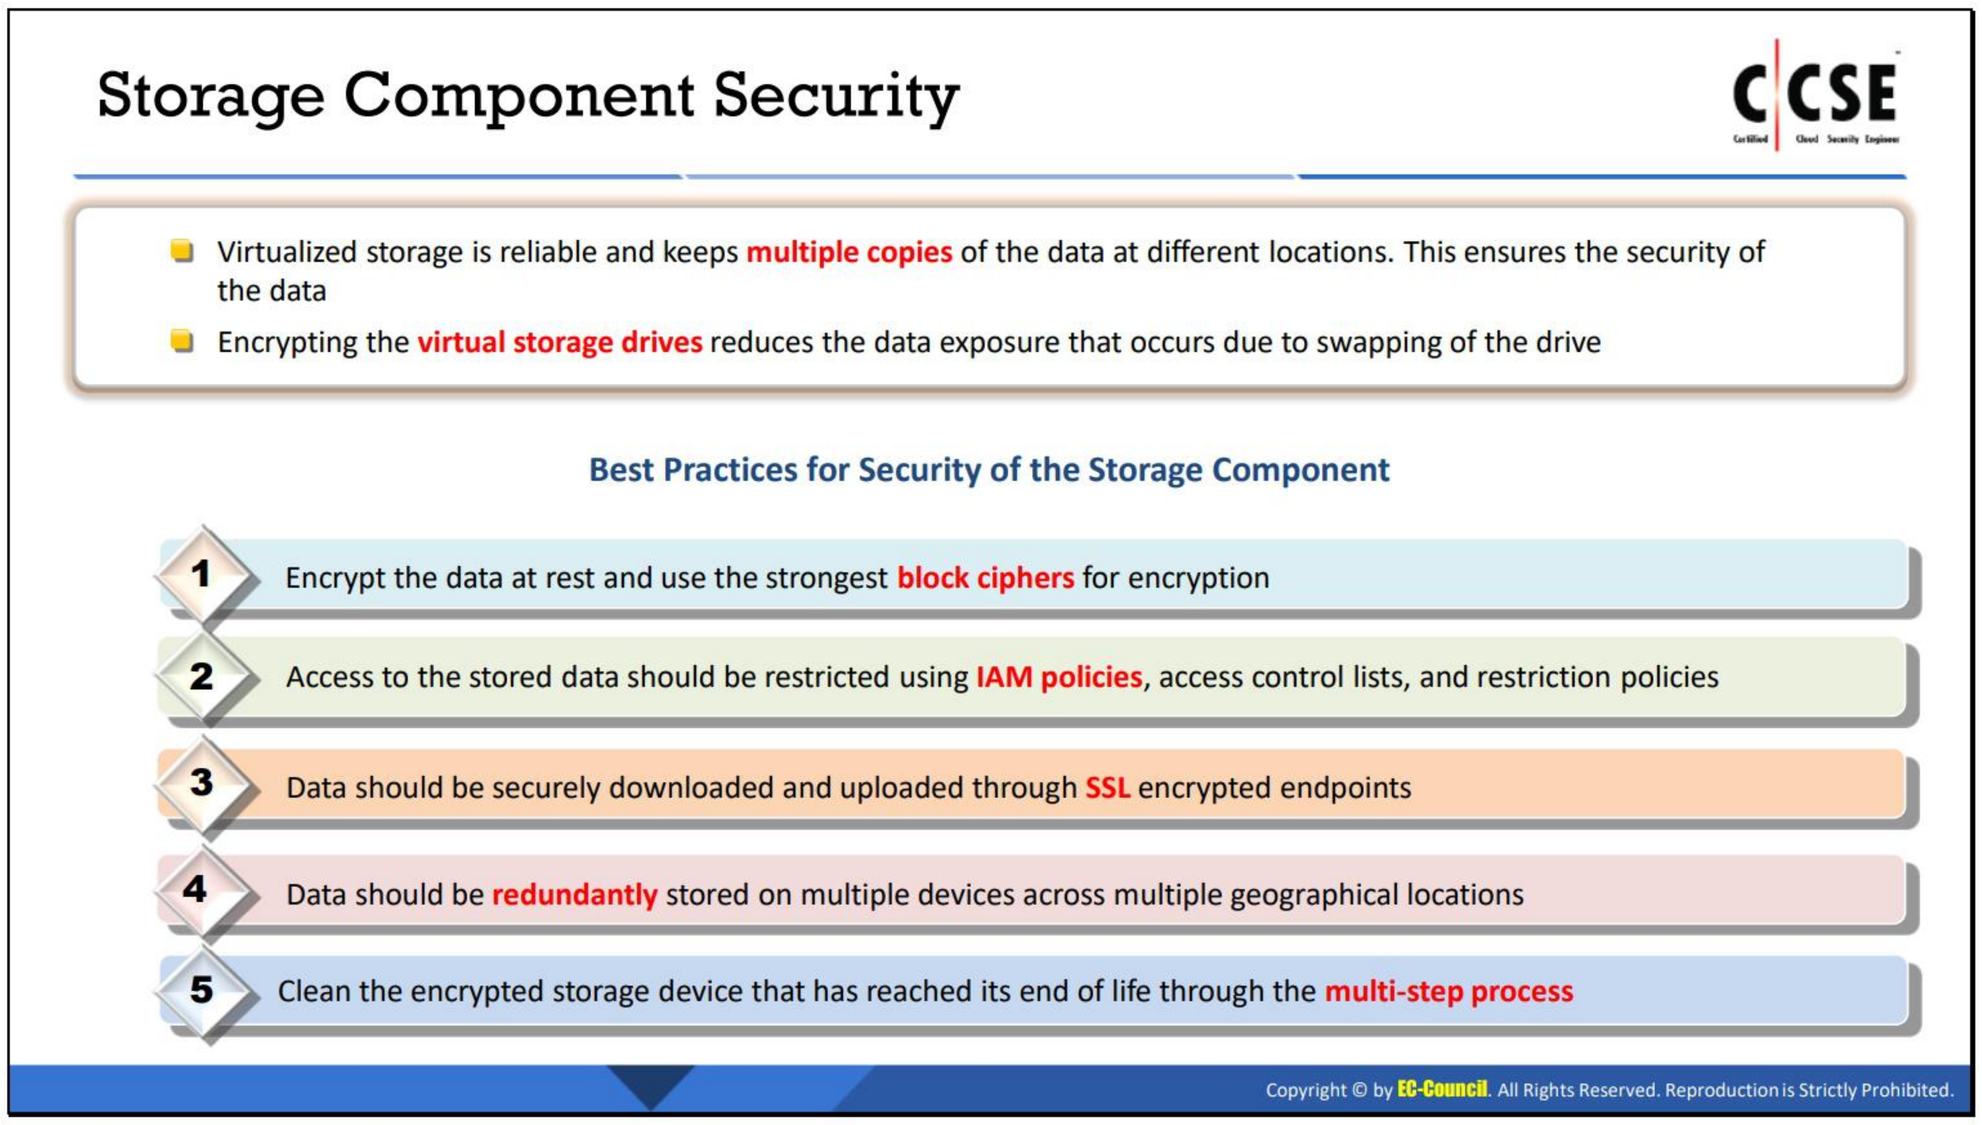
Virtual Machine Image Security

A VM image is a template that can be used to create new instances. A new instance or VM can be created by uploading the image file to the physical machine. VM images provide various benefits to organizations, but they are prone to infrastructure, hypervisor, and storage attacks.

The below best practices should be followed to protect the sensitive data of users and to maintain integrity.

- **Implement encryption**: Before saving the VM image, the image should be encrypted to protect it from attacks.
- **Apply the latest updates and patches**: All the latest updates and patches should be applied to the VM images. This helps to secure the images.
- **Backup VM images**: Backups of the VM images should be performed on a regular basis so that they can be used in case of a disaster.
- **Encrypt VM image backups**: The VM image backups should be encrypted to protect them from being accessed by unauthorized users.
- **Use hashing**: Hashing is a method to ensure that the file is unchanged. To maintain the integrity of the VM image, hashing can be implemented. When the checksum value of the downloaded file is similar to the original file value, it indicates that both files are identical, and no changes have been applied because of file corruptions or any attacks.

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

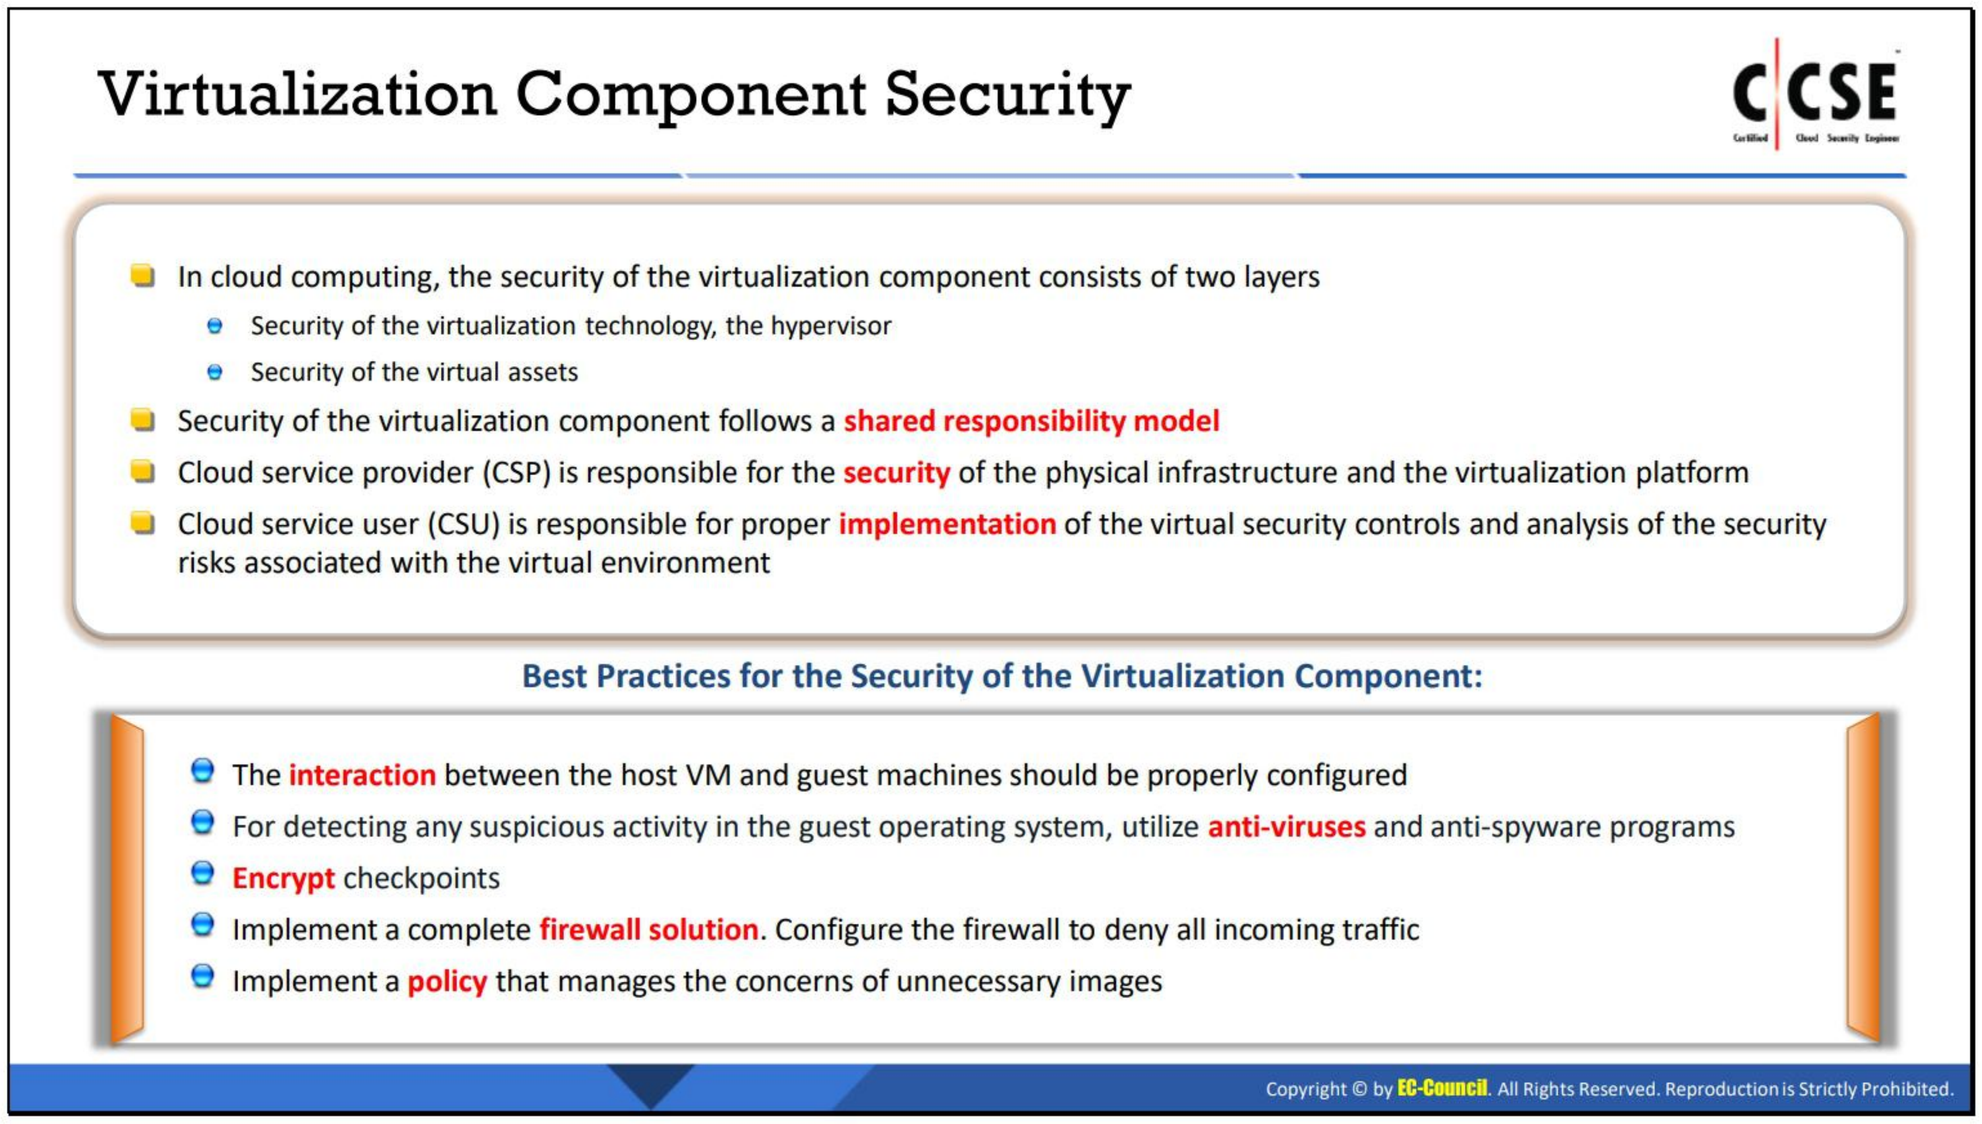


Storage Component Security

- The CSP should provide a storage facility for backup, object storage, archiving, disaster recovery, etc.
- Access to the stored data should be restricted through Identity and Access Management (IAM) policies, Access Control Lists (ACLs), and restriction policies framed by the specific CSP.
- The data should be securely downloaded and uploaded through SSL-encrypted endpoints.
- Encrypt the data-at-rest and use the strongest block ciphers for encryption.
- Data should be redundantly stored on multiple devices across multiple geographical locations.
- Configure the data storage services to use keys from the central key management service for encrypting data before writing it on physical storage.
- Clean the encrypted storage device that has reached its end of life through the multi-step process, which consists of two independent verifications.
- The CSP should provide ample options to protect the data-in-transit such as VPN (using IPsec/IKE encryption), transport layer security, protocols on VMs (such as Windows IPsec or SMB), etc.

hyde01.ir

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

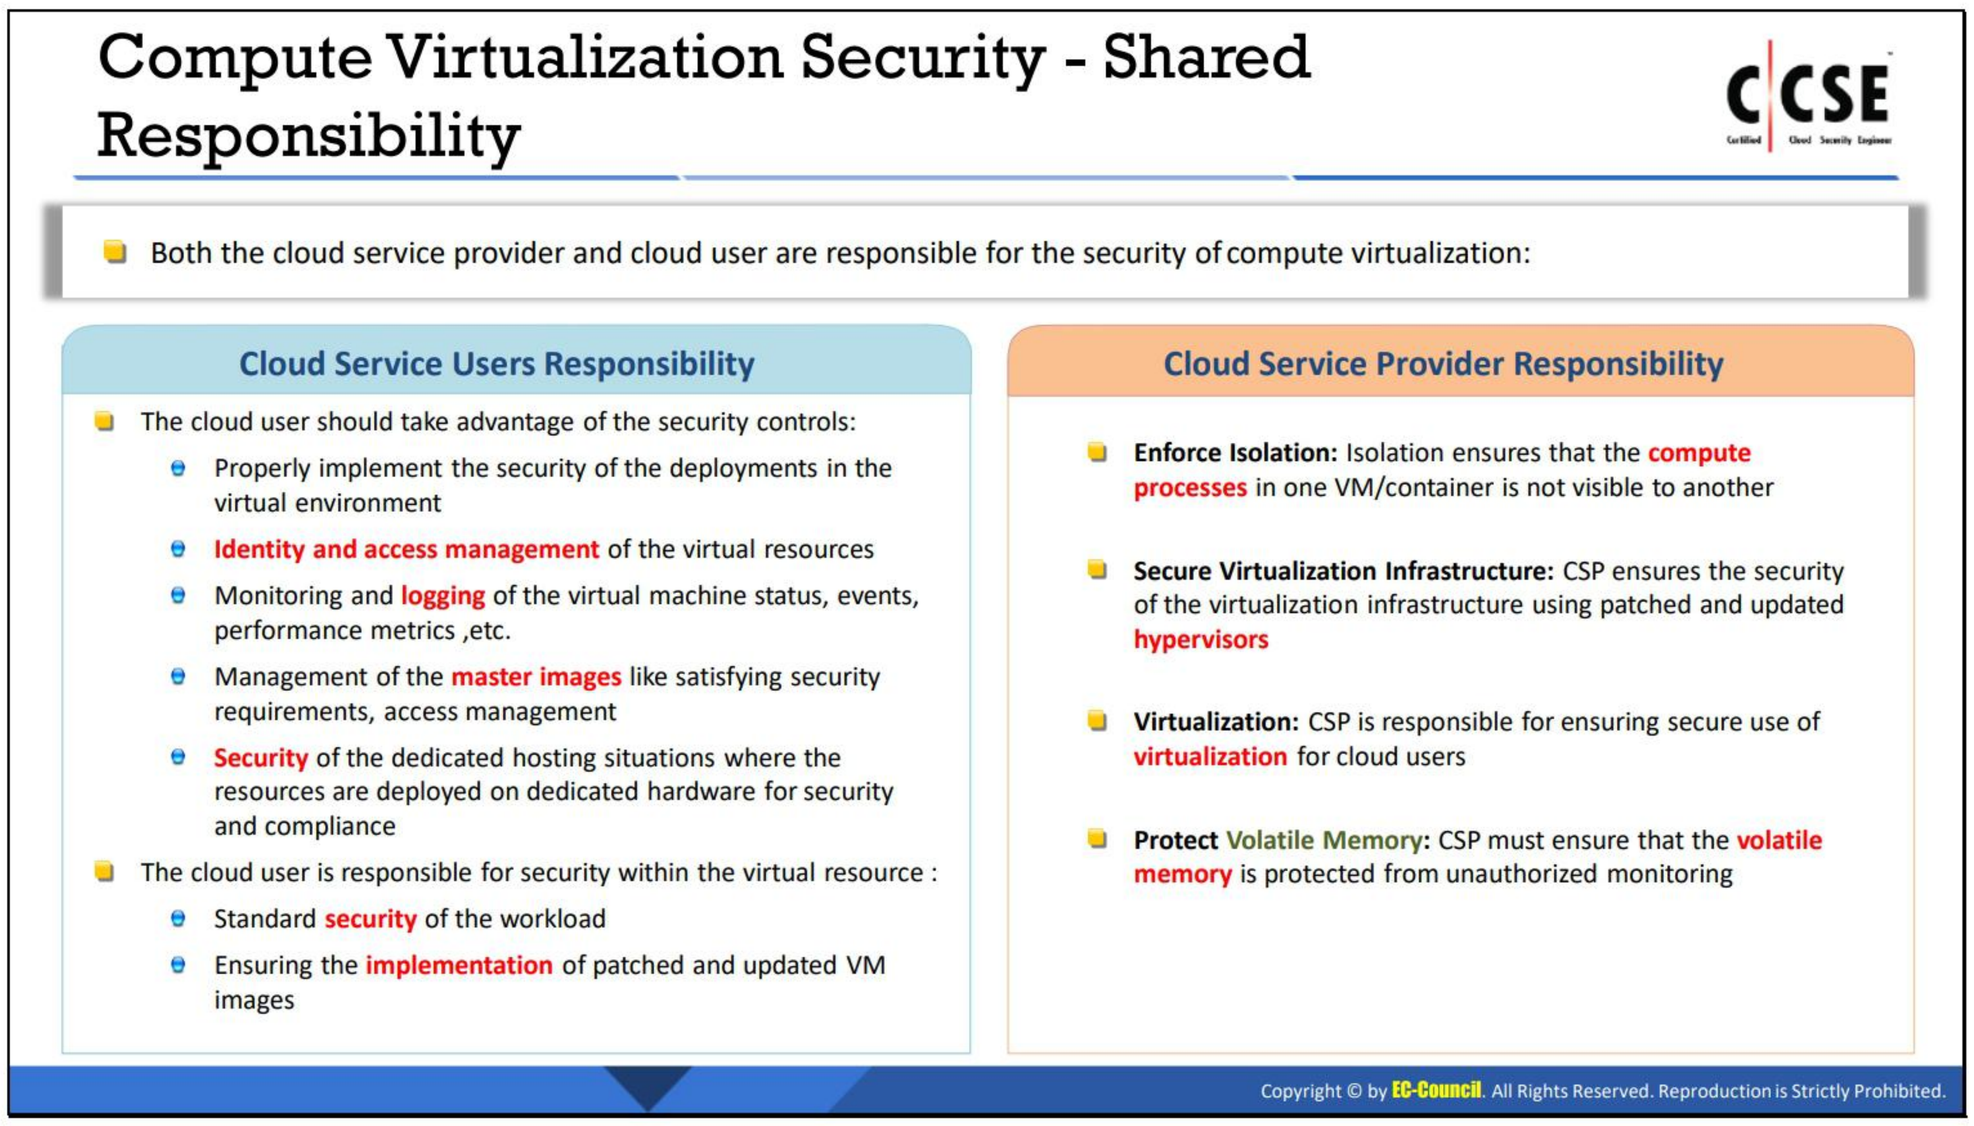


Virtualization Component Security

- The interaction between the host VM and guest machines should be properly configured.
- For detecting any suspicious activity in the guest operating system, utilize anti-viruses and anti-spyware programs.
- Before saving, utilize encryption and hashing of VMs state.
- Checkpoints should be encrypted.
- Organizations utilizing virtualization should implement a policy, which manages the concerns of unnecessary images.
- If there is no need for an employee to access the management plane, then revoke the privilege and access to the host and relevant system.
- The CSP should provide a complete firewall solution. Configure the firewall to deny all incoming traffic.

hide01.ir

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



Compute Virtualization Security- Shared Responsibility

The code execution (e.g., OS) is abstracted from the underlying hardware by compute virtualization. The code runs on top of an abstraction layer rather than on the hardware to allow more flexible usage, for example, when running multiple VMs on the same hardware. Usually, VMs are referred to as compute. However, with ongoing technology evolution and container adoption, compute is also abstracted by containers and certain types of serverless infrastructure.

Both the CSP and cloud user are responsible for the security of compute virtualization:

Cloud Service User Responsibilities

The responsibilities of cloud users are described below.

- The cloud user should take advantage of the security controls:
  - Properly implement the security of the deployments in the virtual environment.
  - IAM of the virtual resources.
  - Monitoring and logging of the VM status, events, performance metrics, etc.
  - Management of the master images, such as satisfying security requirements, access management.
  - Security of the dedicated hosting situations where the resources are deployed on dedicated hardware for security and compliance.
- The cloud user is responsible for security within the virtual resource:
  - Standard security of the workload.
  - Ensuring the implementation of patched and updated VM images.

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

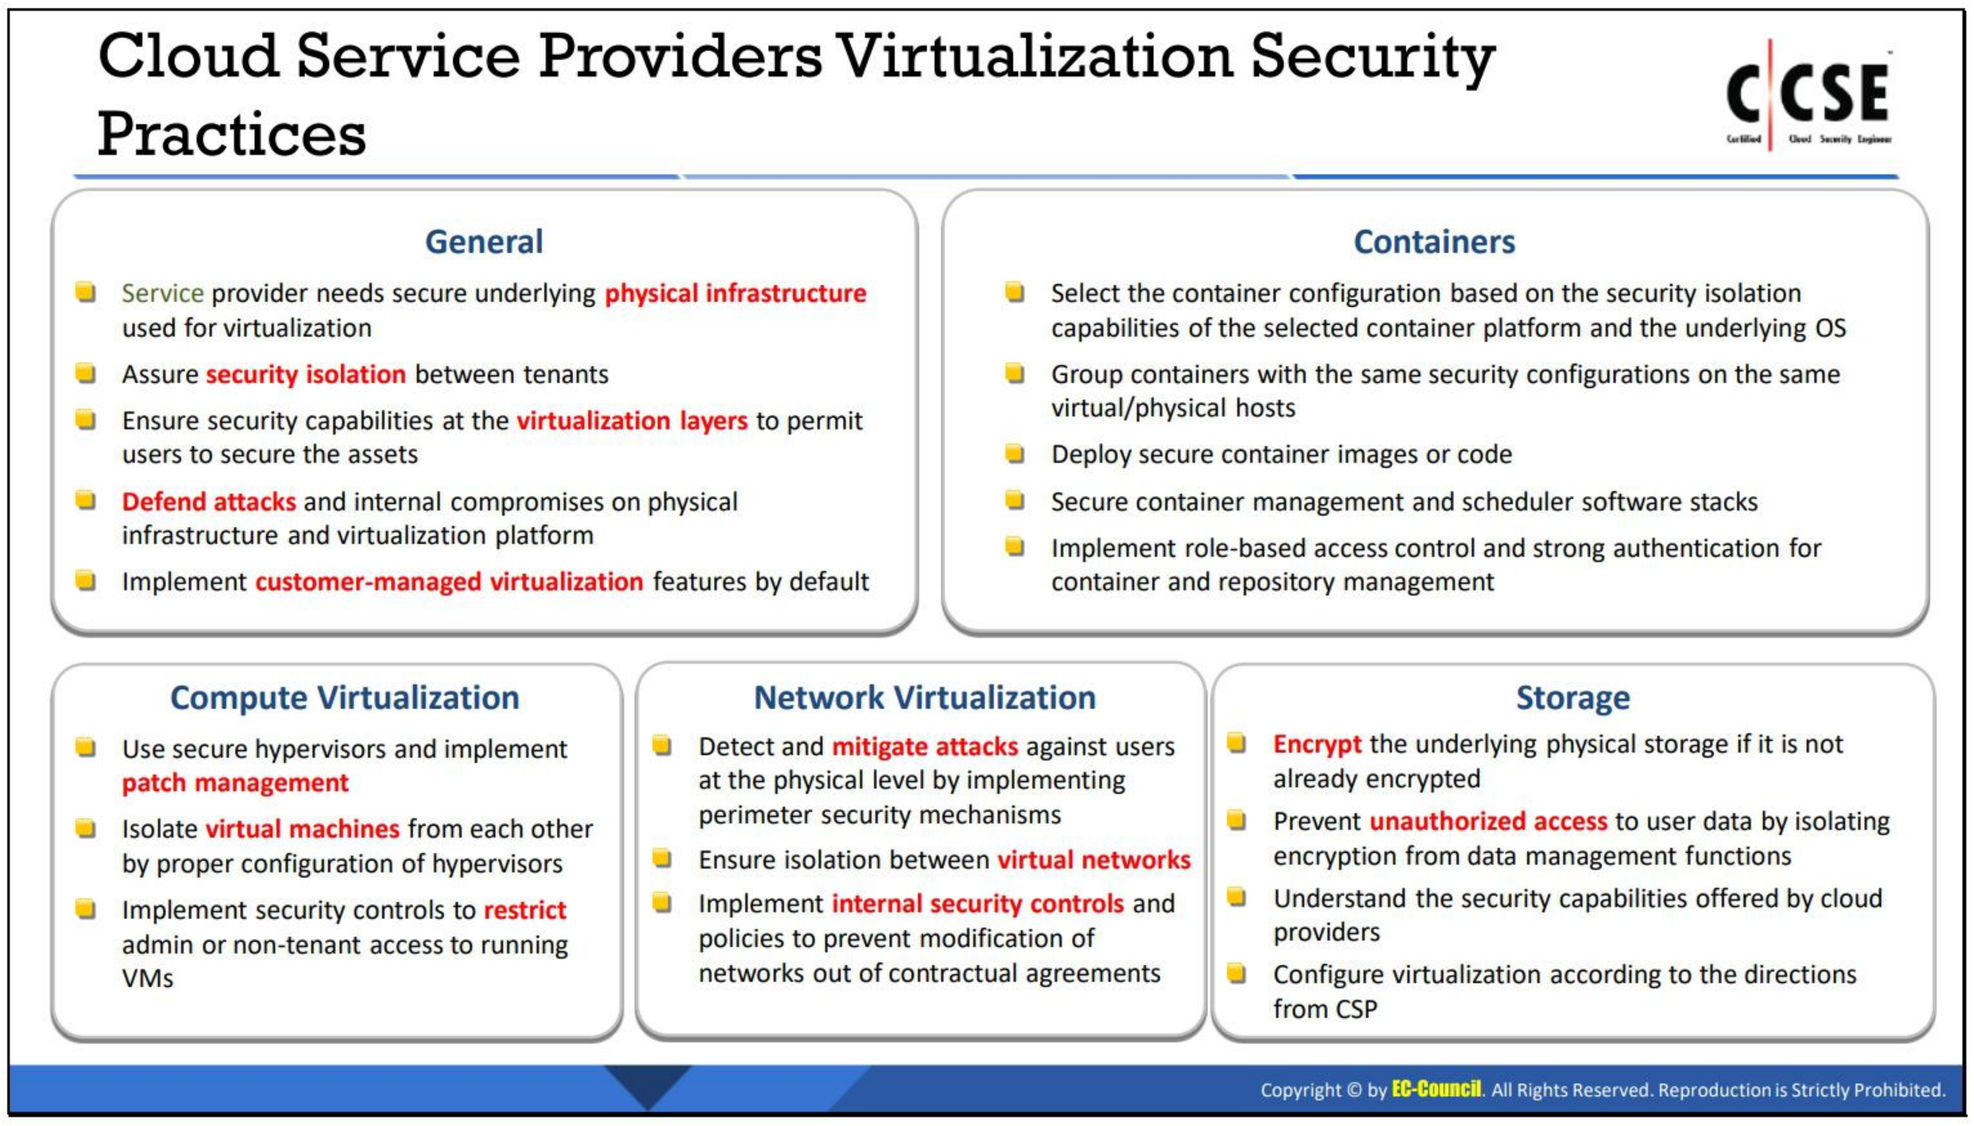
Cloud Service Provider Responsibilities

The responsibilities of CSPs are described below.

- **Enforce isolation:** Isolation ensures that the compute processes in one VM/container is not visible to another.
- **Secure virtualization infrastructure:** CSP ensures the security of the virtualization infrastructure using patched and updated hypervisors.
- **Virtualization:** CSP is responsible for ensuring secure use of virtualization for cloud users.
- **Protect volatile memory:** CSP must ensure that the volatile memory is protected from unauthorized monitoring.

hide01.ir

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



Cloud Service Providers Virtualization Security Practices

The virtualization security practices that CSPs and Cloud Users must follow are described below.

General Best Practices

A CSP must ensure that it implements the following best practices:

- Secure underlying physical infrastructure used for virtualization.
- Assure security isolation between tenants.
- Ensure security capabilities at the virtualization layers to permit users to secure the assets .
- Defend attacks and internal compromises on the physical infrastructure and virtualization platform.
- Implement customer-managed virtualization features by default.

Container Best Practices

A CSP/cloud customer must implement the following best practices related to containers:

- Select the container configuration based on the security isolation capabilities of the selected container platform and the underlying OS.
- Group containers with the same security configurations on the same virtual/physical hosts.
- Deploy secure container images or code.
- Secure container management and scheduler software stacks.

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

- Implement role-based access control and strong authentication for container and repository management.

Compute Virtualization Best Practices

The following best practices must be implemented by a CSP:

- Use secure hypervisors and implement patch management.
- Isolate virtual machines from each other by proper configuration of hypervisors.
- Implement security controls to restrict administrator or non-tenant access to running VMs.

Network Virtualization Best Practices

A CSP must implement the following best practices related to network virtualization:

- Detect and mitigate attacks against users at the physical level by implementing perimeter security mechanisms.
- Ensure isolation between virtual networks.
- Implement internal security controls and policies to prevent modification of networks out of contractual agreements.

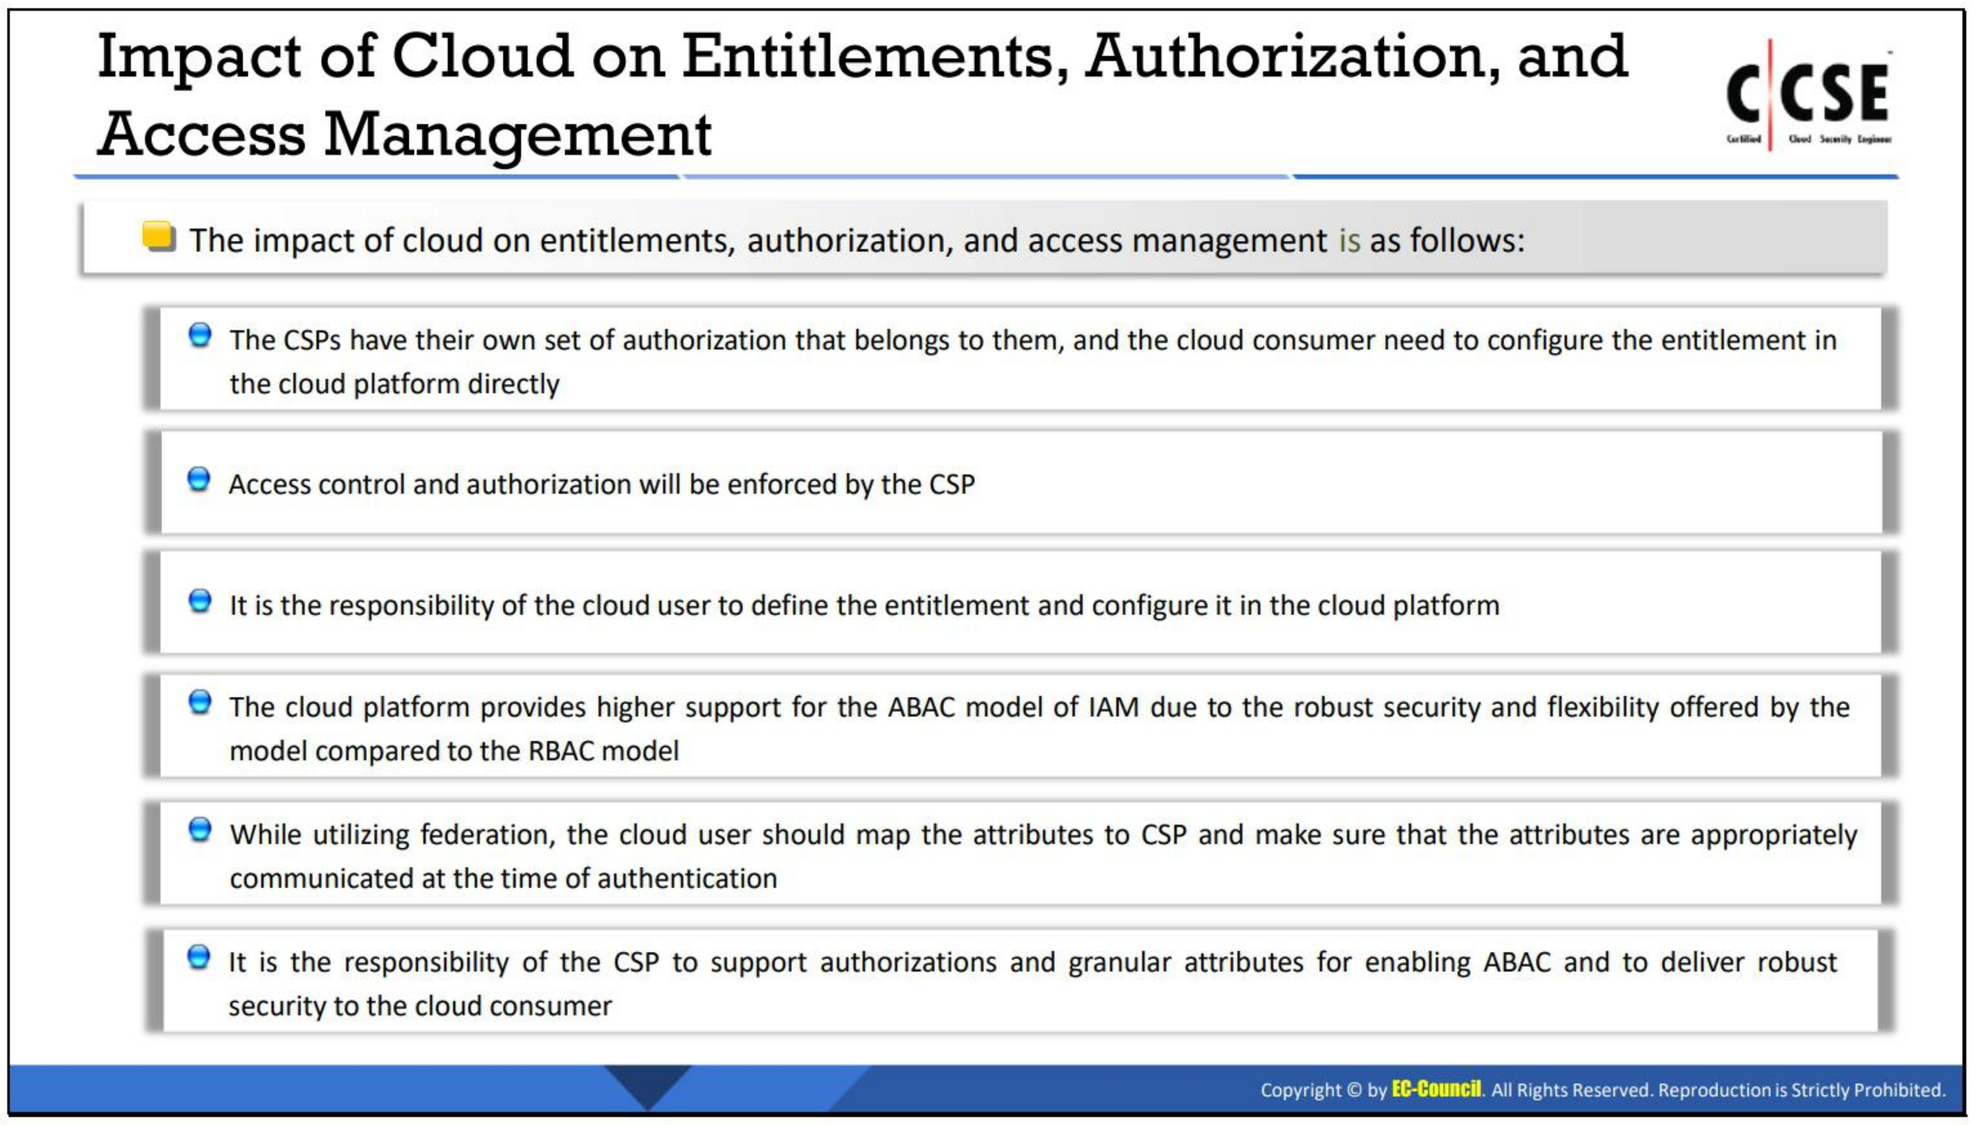
Storage Virtualization Best Practices

A CSP/cloud consumer must implement the following best practices related to storage virtualization:

hide01.ir

- Encrypt the underlying physical storage if it is not already encrypted.
- Prevent unauthorized access to user data by isolating encryption from data management functions.
- Understand the security capabilities offered by cloud providers.
- Configure virtualization according to the directions from CSP.

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



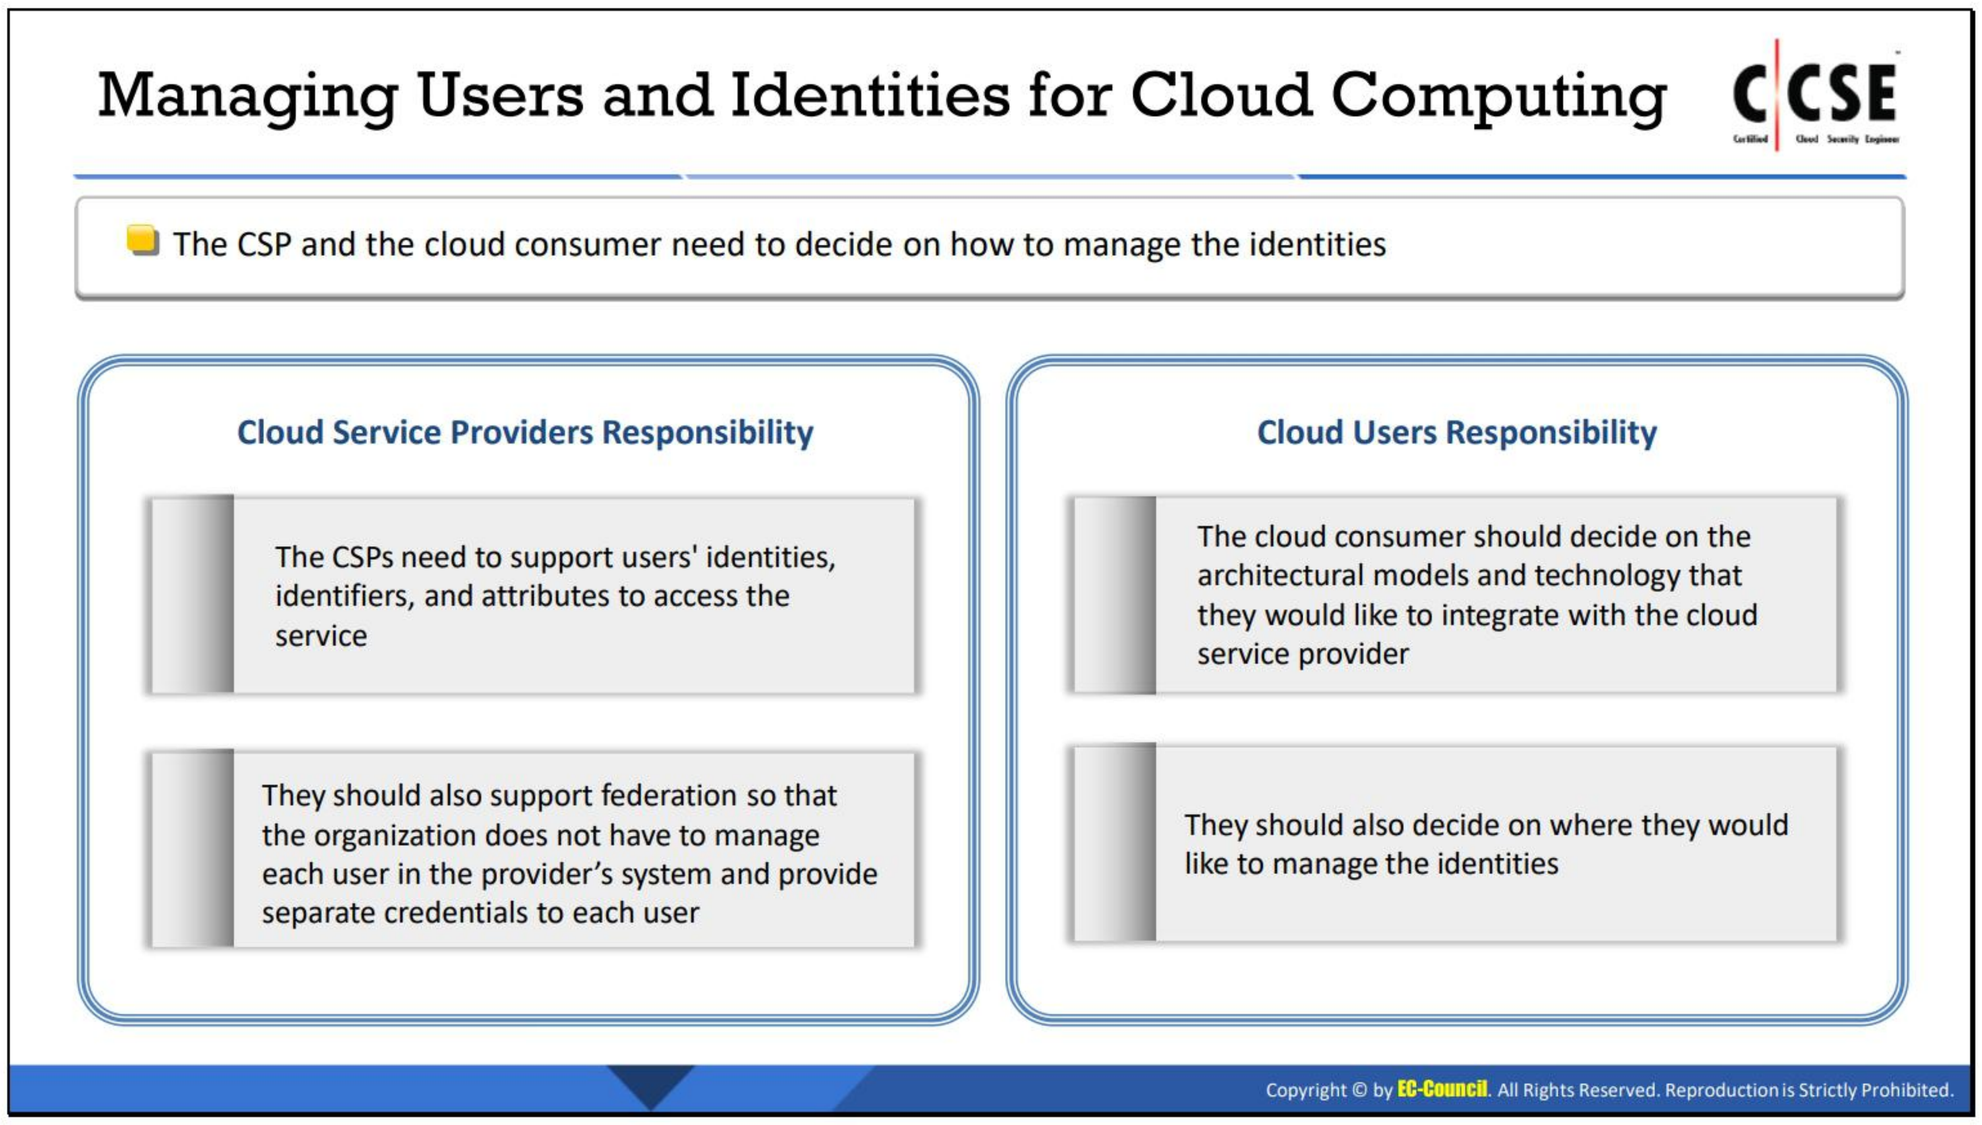
Impact of Cloud on Entitlements, Authorization, and Access Management

Entitlements, authorization, and access management are greatly impacted by the cloud in the following ways:

h3de01.ir

- Like any other technology, the CSPs will also have their own set of authorizations that are specific to them. If the cloud provider does not support XACML, the user needs to directly configure the entitlements in the cloud platform.
- Enforcing authorization and access controls will be handled by the cloud provider.
- It is the responsibility of the cloud user to define and configure entitlements within the cloud platform.
- Cloud platforms provide support for attribute-based access control (ABAC) for IAM. This offers a high level of security and flexibility compared to the role-based access control (RBAC) model. RBAC is a traditional model, which is based on a single attribute, defined as a role. ABAC provides fine granular decisions based on various attributes such as location, authentication method, role, etc. For cloud-based access management, ABAC is the most preferred model.
- When federation is used, the cloud consumer is responsible for mapping the attributes to the cloud provider along with the roles and groups. Cloud consumers should also ensure that roles, groups, and attribute mapping are properly communicated during authentication.
- To enable ABAC and provide high security for cloud consumers, the cloud provider should support granular attributes.

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



Managing Users and Identities for Cloud Computing

Identity management mainly focuses on processes and technologies for managing, registering, and provisioning identities. Management of identities is a major problem for the information technology team. IT administrators used to provision users individually in every internal system. Even with centralized directory servers, single sign-on, users still need to maintain credentials.

Cloud users and cloud providers need to make the following decisions on how to manage identities:

- Cloud providers should provide identifiers, attributes, and internal identities for users who can access the service directly. They should also support federation to prevent the organization from manually provisioning and managing users.
- Cloud users need to decide how and where they want to handle the identities. The users should also decide on the architectural models and technologies that need to be integrated with the cloud provider.

The cloud user can log in to the cloud provider and can create the required identities. Most businesses turn to the federation as this is not scalable. There are certain cases where some of the identities should be isolated from the cloud provider.

When a cloud user makes use of federation, the user needs to identify the authoritative source that holds the identities to be federated. The user should also decide on whether to use the authoritative source directly as the identity provider or use a different identity source.

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

There are two models, namely, the “hub and spoke” model and the “free form” model.

- **Hub and spoke model**: In this model, internal identity providers communicate with a central repository or broker. It then serves as an identity provider for the federation to cloud providers.

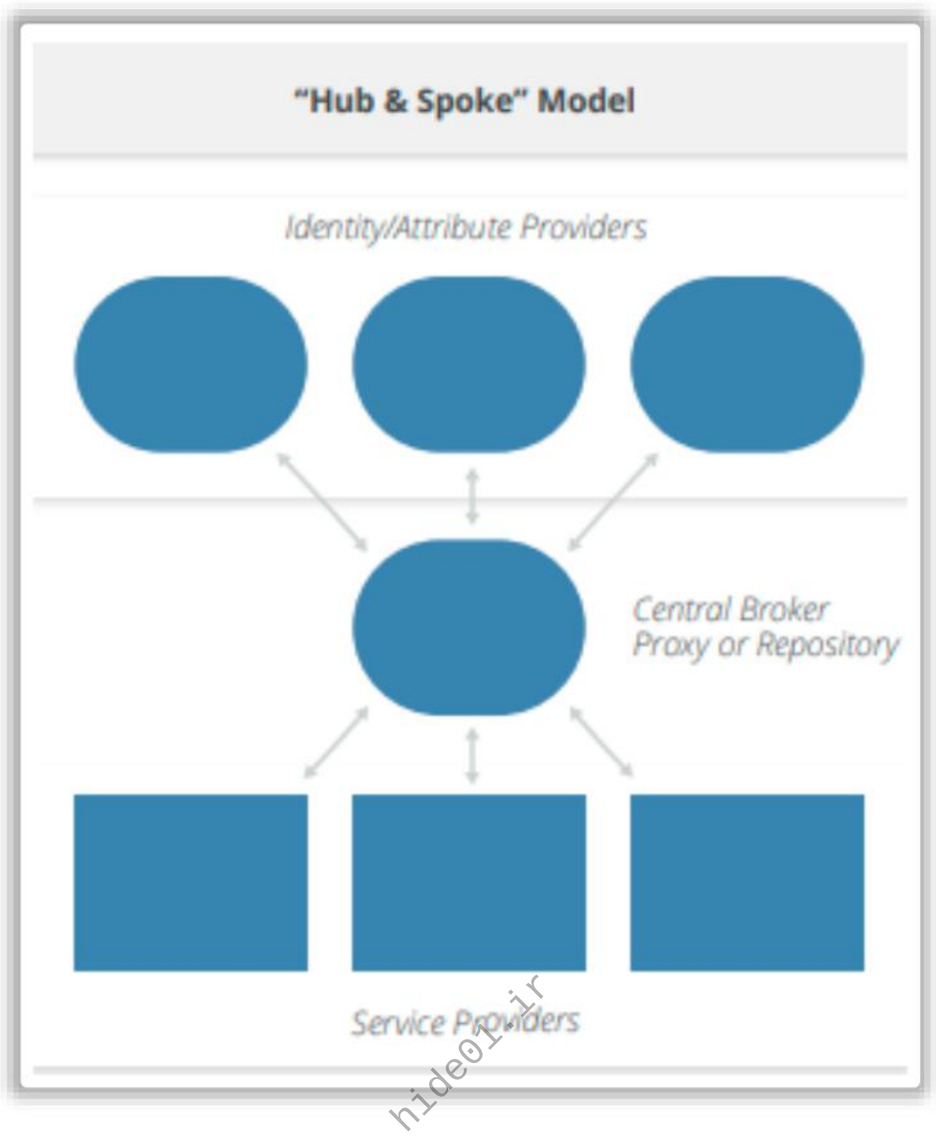


Figure: 2.7: Hub and Spoke Model

- **Free-form model**: In this model, the internal identity provider directly connects to the cloud provider.

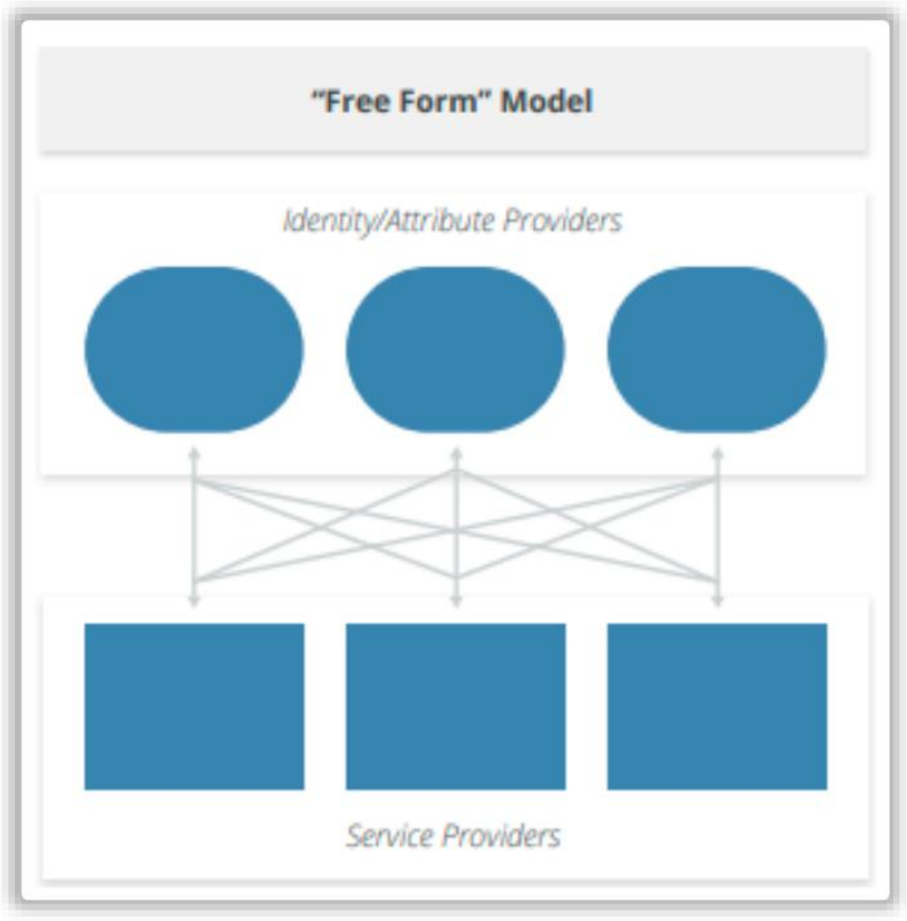


Figure: 2.8: Free Form Model

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

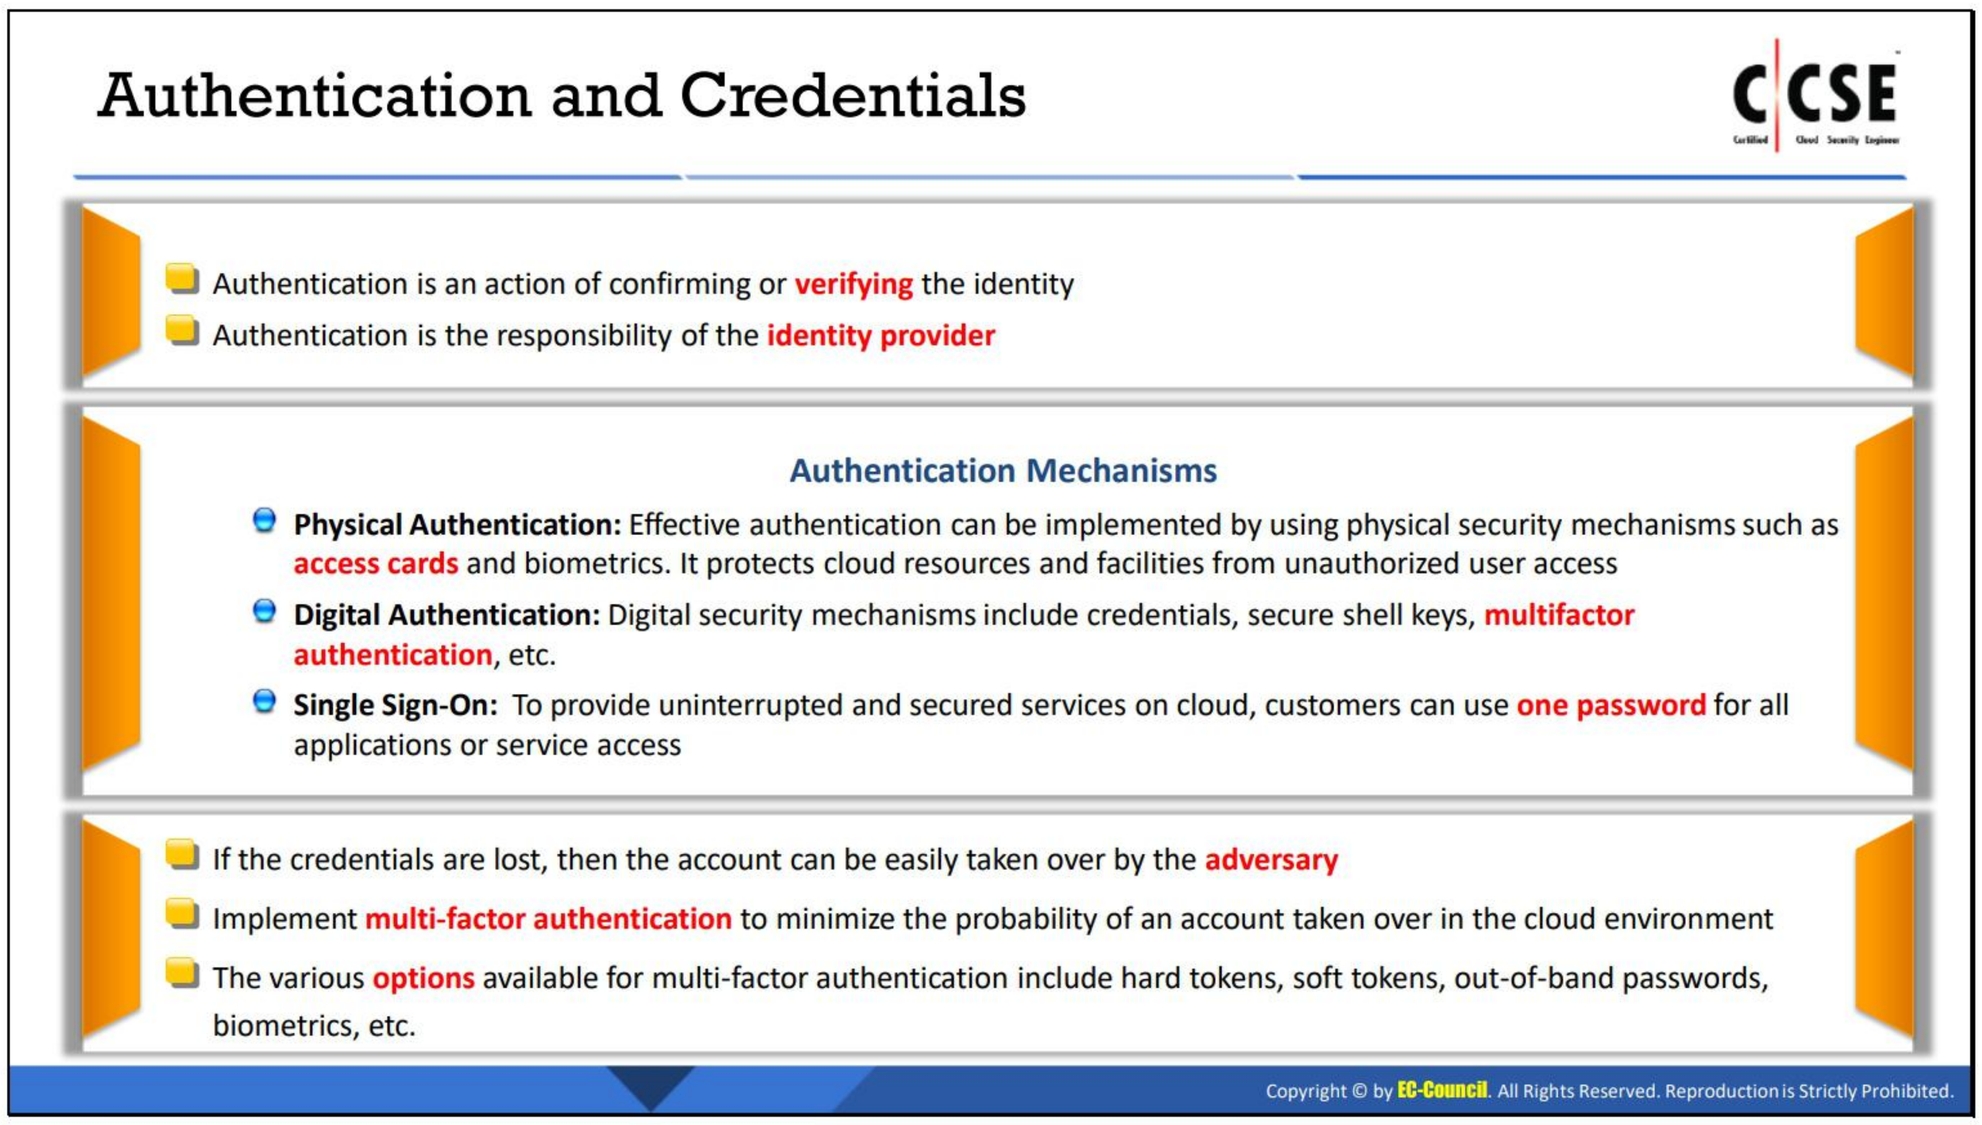
After determining the architectural model, still certain decisions need to be made on the processes required for implementation, such as:

- How identities for applications, codes, devices, and other services should be managed. Here, either the same model and standard or a new approach can be used.
- Defining and integrating the identity provisioning process into cloud deployments. Though there can be multiple identity provisioning processes for different use cases, the goal should be to have a unified process.
- Supporting and provisioning individual cloud providers and deployments. There must be a process to add new providers to the IAM infrastructure.

The CSP should determine what identity management standards they should support. A few providers only support federation whereas certain providers offer multiple IAM standards along with their own internal user management. The cloud providers that serve the enterprise market should support federated identity.

hide01.ir

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



Authentication and Credentials

Authentication is the process of confirming and proving a user identity. The identity provider is responsible for the authentication. In cloud computing, strong authentication using various factors is important because of the following reasons:

- Cloud services are accessed over the network and internet. When a user loses his credentials, an attacker can easily take over the account as the attacker is not restricted to a local network.
- Use of federation for single sign-on uses a single set of credentials. When this single set of credentials is stolen, the attacker can access a greater number of cloud services.

Multi-factor authentication (MFA) is the strongest method for authenticating a user and helps reduce account takeovers. MFA requires a user to provide two or more factors to access a resource. When MFA is used with federation, the identity provider should send the MFA status to the relying party as an attribute.

MFA offers multiple options; they are described below.

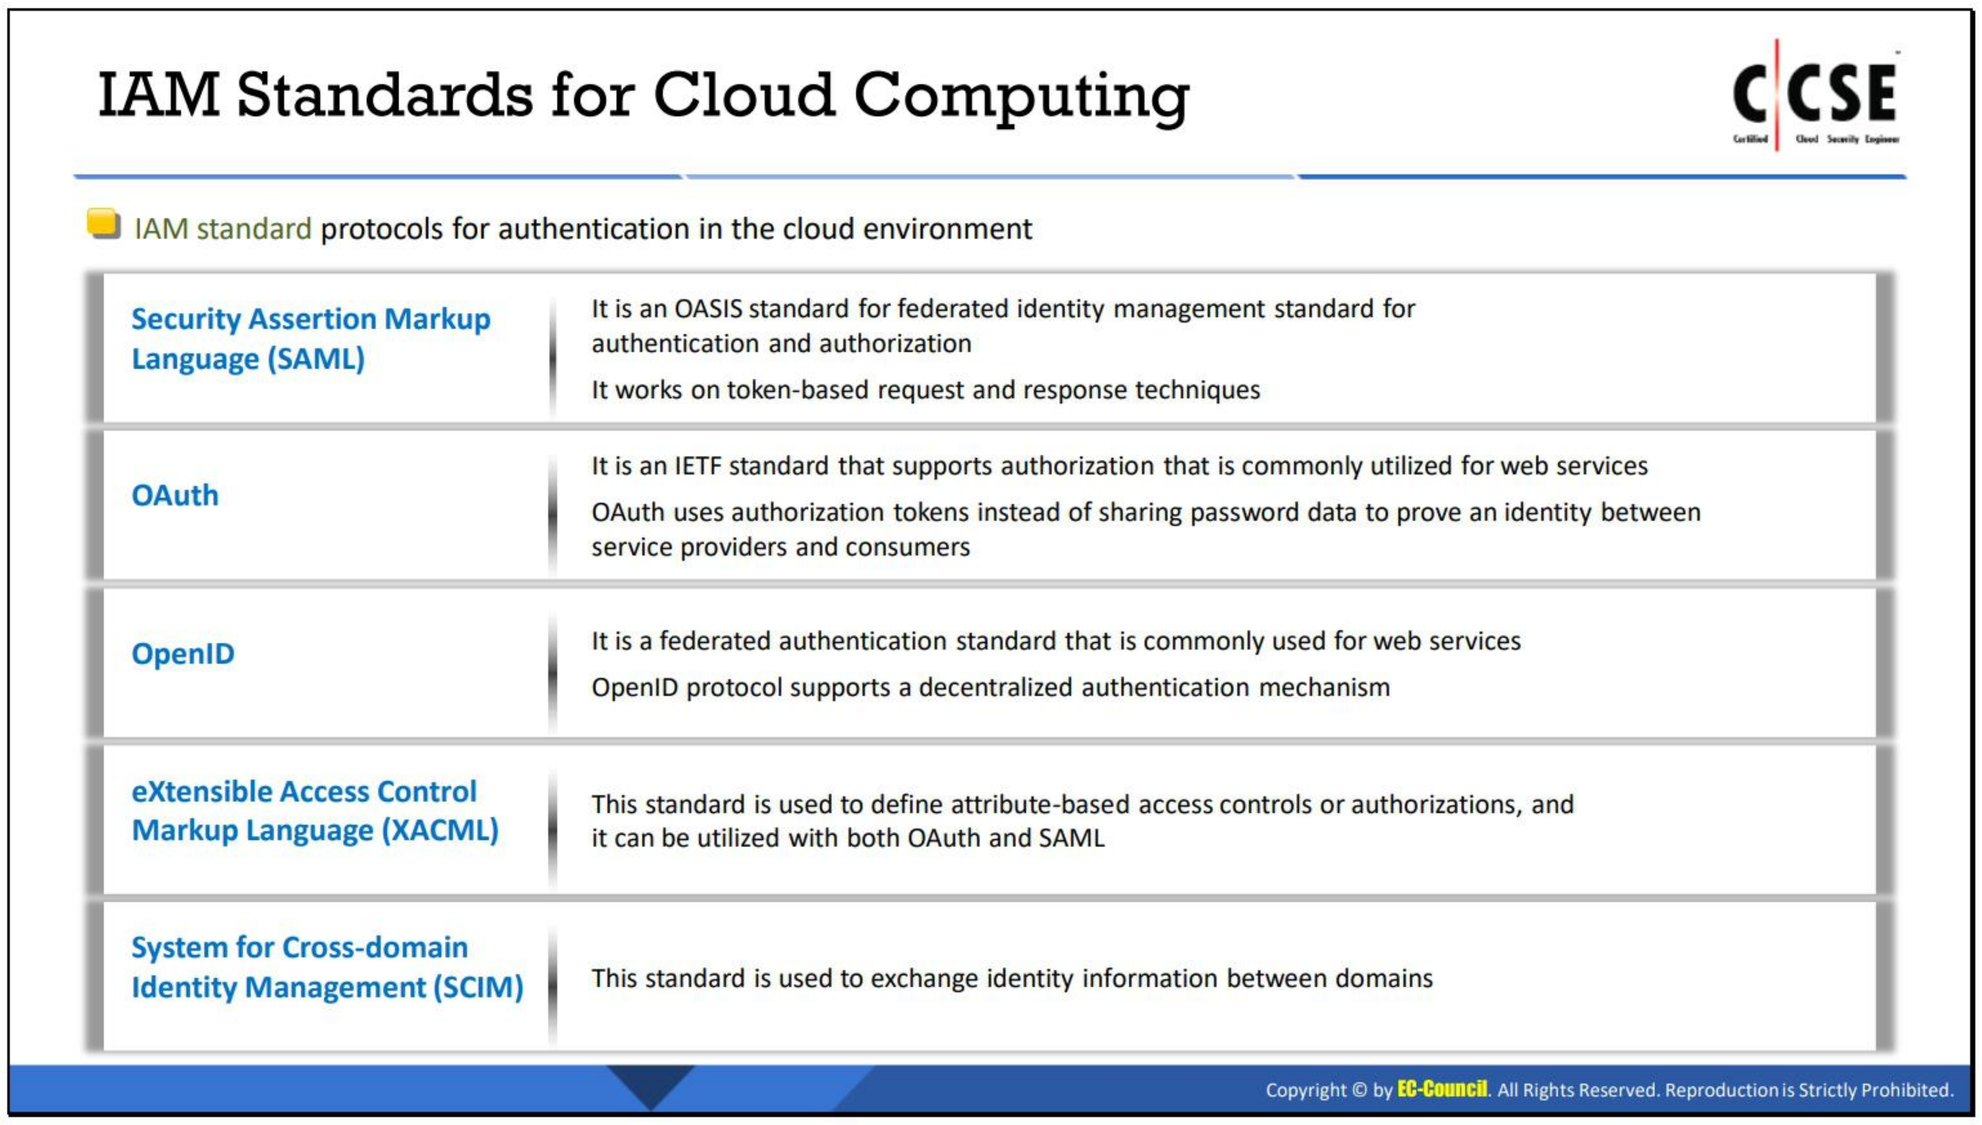
- **Hard tokens:** Hardware devices generate a one-time password (OTP) based on the cryptographic key stored in the device. The same cryptographic key will be held by the server that generates the same OTP. The server verifies that the value entered by the user is correct. Hard tokens are the best option when there is a need for a greater level of security.
- **Soft tokens:** Soft tokens are software applications that run on a device and use cryptographic operations to authenticate the device and user. Soft tokens support advanced cryptography. When the user's device gets compromised, soft tokens can be easily compromised.

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

- **Out-of-bound passwords:** It is a user-friendly method and the users do not need to install any application. The user gets an OTP to the device in the form of an SMS for authentication. Though it is a good option, the threat model should consider message interception.
- **Biometrics:** In cloud computing, biometric is used for local protection and it is not shared with the service provider.

hide01.ir

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



IAM Standards for Cloud Computing

Identity and access management (IAM) helps an organization to meet a wide range of industry standards and regulatory requirements. However, achieving compliance is not always easy even with the IAM.

hide01.ir

Several IAM standards have been designed to securely transfer sensitive information and to protect the data when it passes through a network or servers. It is important to understand what they are and how they are useful for various industrial standards and regulatory requirements. Understanding those standards helps an organization to make use of IAM for compliance.

Discussed below are some of the IAM standards for cloud computing.

Security Assertion Markup Language (SAML)

SAML is an OASIS open standard that is available for all organizations. SAML is commonly used by systems that employ the single sign-on method of access control. One set of credentials can be used for accessing multiple applications. It is used for authentication and authorization of device connections. Assertions between the identity provider and relying party are made using XML. Assertions include attribute statements, authentication statements, and authorization decision statements.

This XML-based markup language helps to improve security and compliance by using digital signature in place of a password for authentication and authorization.

OAuth

OAuth is an Internet Engineering Task Force (IETF) standard used for authorization. OAuth is used by various customer-facing platforms such as Facebook, Twitter, Google to connect third-party applications with the users' permissions. It is designed to work with Hypertext Transfer Protocol

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

(HTTP). OAuth allows approved applications to make use of credentials of a service to access other applications. It authenticates the identity of a user and then connections are authorized between applications.

OpenID

OpenID is a standard used for federated authentication that is used for web applications. It is used when interacting with products from Yahoo and Google. The implementation of the OpenID standard is less complicated than the SAML implementation. This makes it accessible for different applications.

For consumer applications, OpenID helps users maintain a consistent identity through the platform. A user can connect with every service he is authorized to access using a single identifier and password. Many businesses have started using OpenID for cloud applications. OpenID helps to maintain the integrity of user identities in complex systems.

eXtensible Access Control Markup Language (XACML)

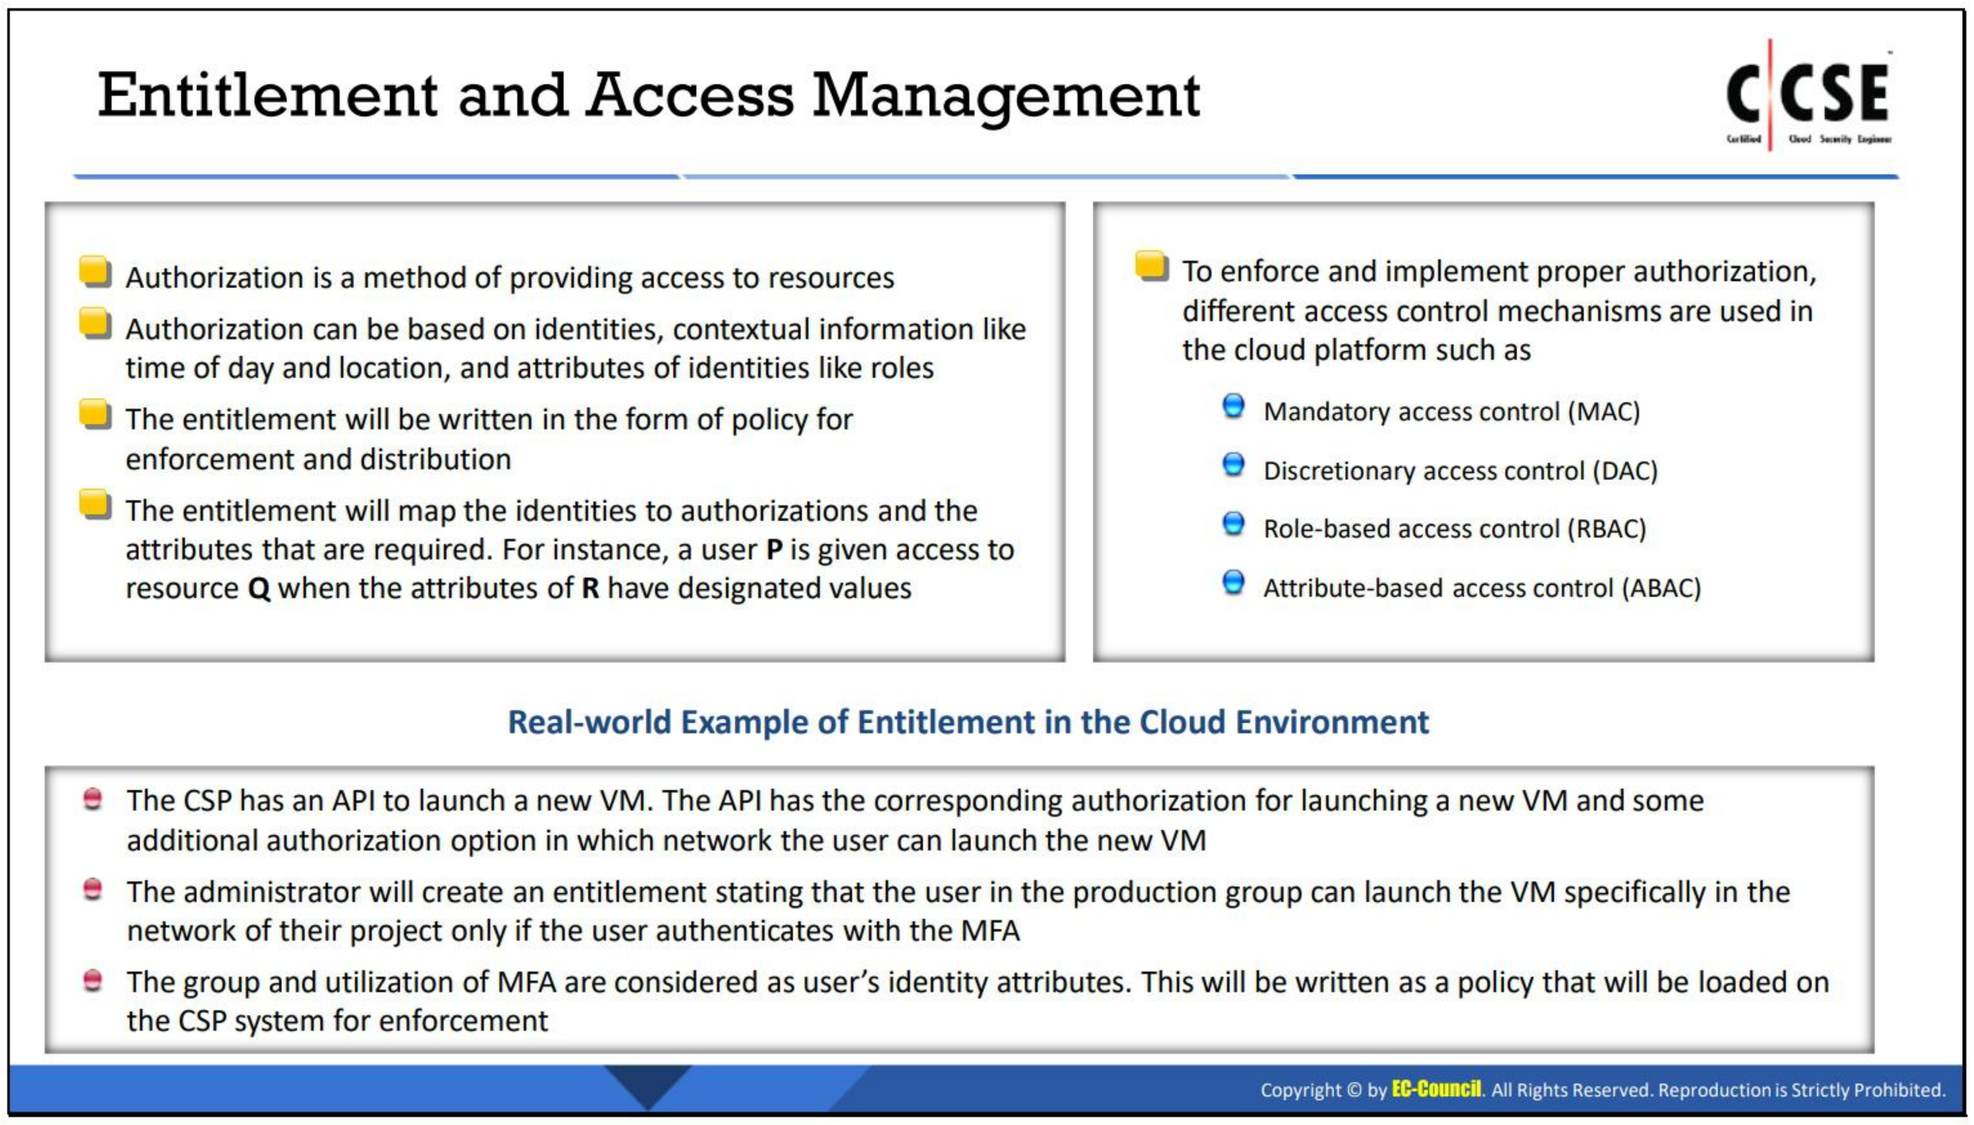
XACML is a standard used for defining fine-grained attribute-based access controls to evaluate access requests. It is an access control policy language used for defining access controls at a policy decision point. The access controls are then passed to a policy enforcement point. It can also be used along with OAuth and SAML.

System for Cross-Domain Identity Management (SCIM)

SCIM is an open standard used for automating user provisions. With SCIM, it is easier to manage user identities in cloud-based services and applications. It can be used for sharing identity information such as user attributes, attribute schema, and group membership between domains.

Lifecycle management becomes easier with SCIM as it can automate provisioning and deprovisioning users when they come in and leave the system. Organizations that adopt SCIM can reduce the risk posed by accounts of former users. This is achieved by ensuring that users who have left the system can no longer log in.

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



Entitlement and Access Management

An entitlement maps the identities to authentications and other required attributes. A map of these entitlements is commonly referred to as the entitlement matrix. An entitlement consists of a set of attributes that mentions the privileges and access rights of an authenticated security principal. For distribution and enforcement, entitlement is encoded in the form of technical policies. The policies will be loaded into the cloud provider's system for enforcement.

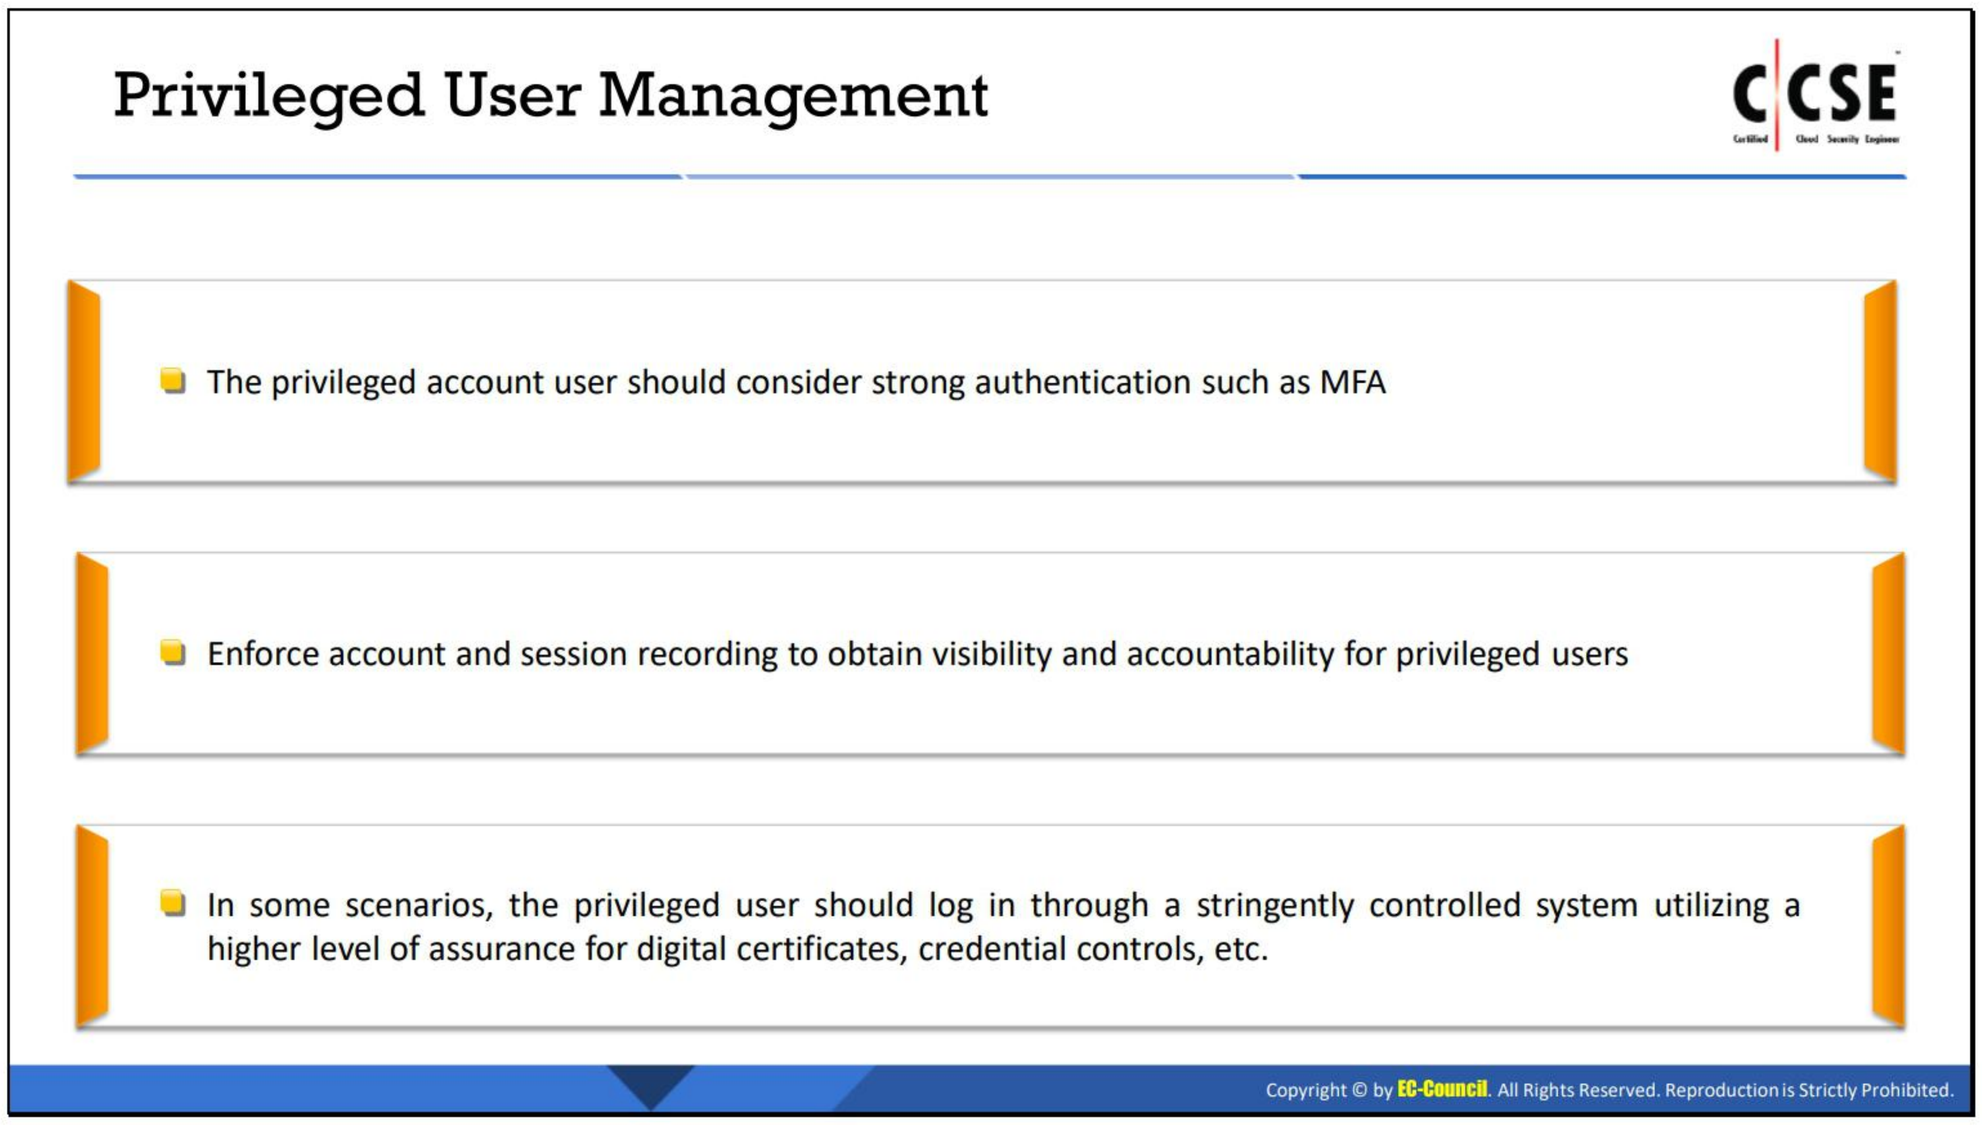
Access management refers to the process of defining and enforcing authorizations.

The concept of entitlement can be easily understood using a sample real-time cloud example. A cloud provider has an API for launching instances or virtual machines (VMs). The API has a related authorization to allow the launching of a VM. It also has some additional authorization options for launching the VM in only a specific virtual network. The administrator creates an entitlement specifying that the user in a developer group can launch VM only in their project network, only when the user is authenticated with MFA. The use of MFA and group are used as attributes for identifying a user.

Entitlement	Super-Admin	Service-1 Admin	Service-2 Admin	Dev	Security-Audit	Security-Admin
Service 1 List	X	X		X	X	X
Service 2 List	X		X	X	X	X
Service 1 Modify Network	X	X		X		X
Service 2 Modify Security Rule	X	X				X
Read Audit Logs	X				X	X

Table 2.1: Entitlement and Access Management

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



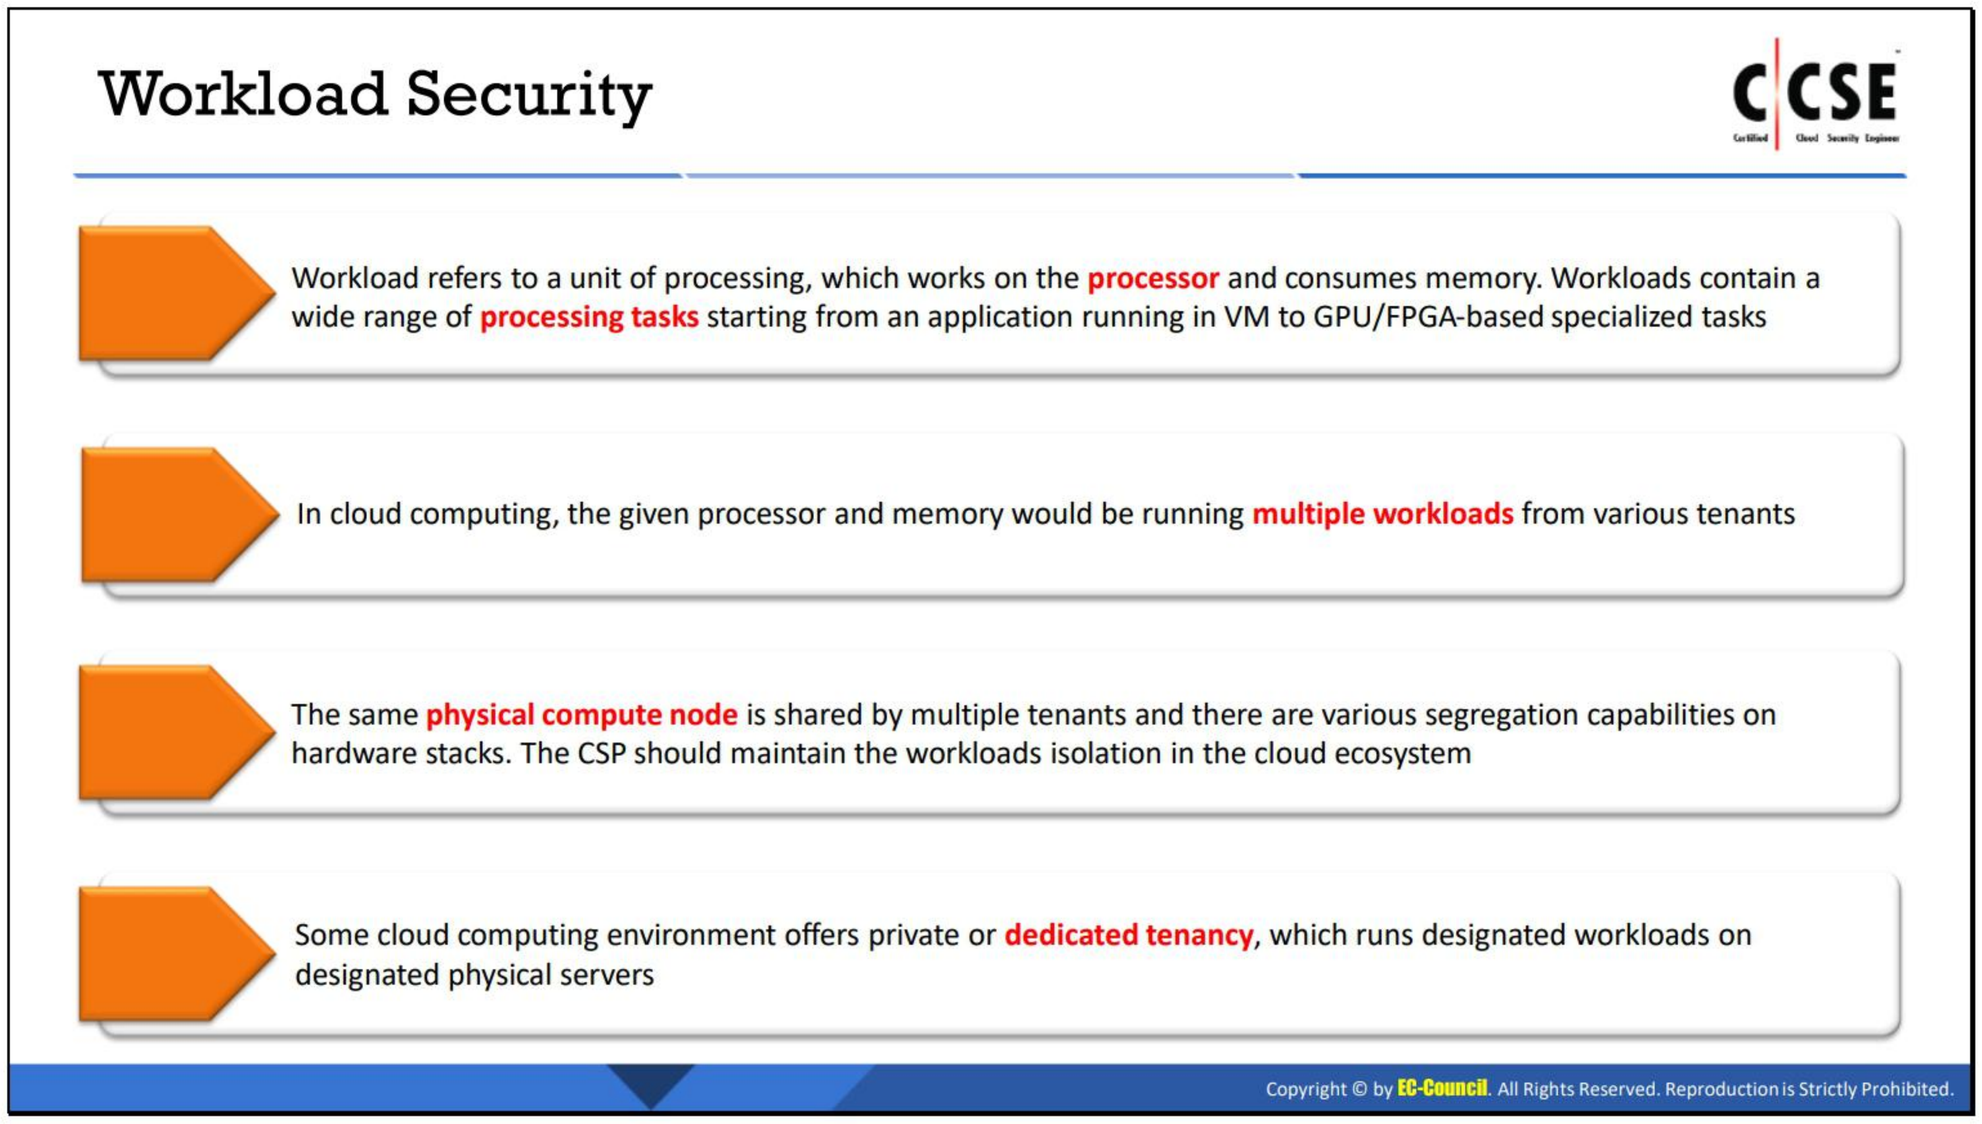
Privileged User Management

Privileged users have elevated permissions to critical resources. These users are authorized to perform security-related operations that a normal user is not allowed to perform. Attackers focus more on stealing the credentials of privileged users. This makes it challenging for the IT team to secure privileged accounts. It is vital for the organization to secure privileged user accounts to protect sensitive information and critical systems from attackers.

Privileged user management (PUM) is the method of handling privileged user accounts that are associated with certain resources or assets. To drive up accountability and visibility for privileged users, account and session re-coding should be implemented.

For certain cases, the privileged users should sign-in from a strongly controlled system that uses a greater level of assurance for credential control, digital certificates, and jump hosts.

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



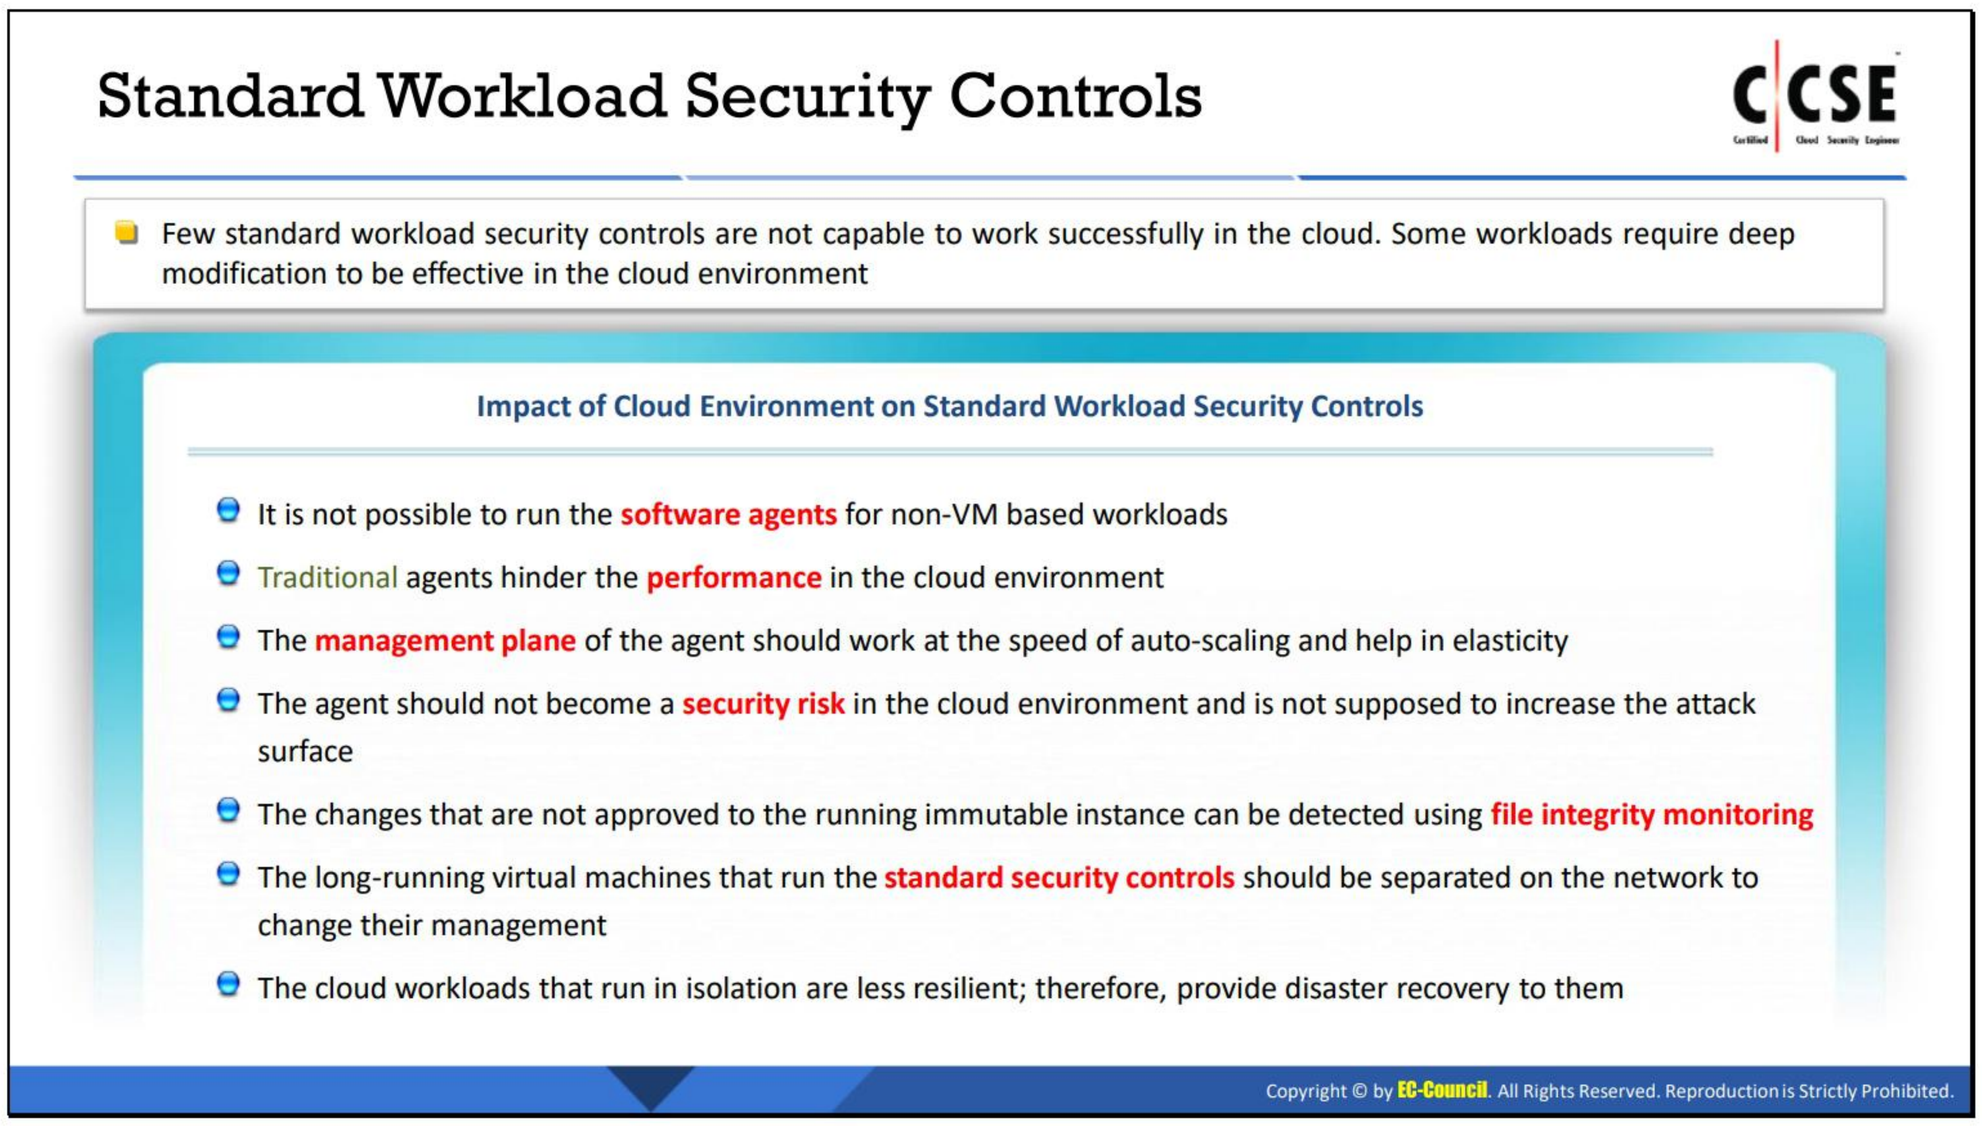
Workload Security

Cloud workload security provides workload visibility by offering control on issues related to the user. It ensures protection against attacks that traditional solutions cannot resolve. It can consolidate events of a single interface to manage security technologies quickly, including a primary security dashboard. Multiple tenants run multiple workloads on processors and memory on physical compute nodes with segregation range capabilities on different hardware stacks. Cloud providers maintain workload isolation with top priorities. This model runs the design workload on a physical server.

Costs in public clouds increase due to suboptimal use of internal resources caused by the resource pool. Cloud users do not control the workload running on a physical platform. However, certain platforms support hardware pools for availability, compliance, and other requirements. Steps to ensure cloud workload security include the following:

- Ensure vulnerability and configuration management
- Implement traffic visibility and network segmentation
- Offer network security management and monitoring
- Prevent exploits and protect the memory
- Encrypt data-at-rest and data-in-transit when using IaaS
- Implement advanced behavioral response and detection
- Install antivirus software

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



Standard Workload Security Controls

Organizations must look for security controls that can ensure workload security in the cloud instead of using the same controls they use on-premises. Because most cloud workloads consist of VMs, a large number of endpoint tools incorporate software agents in VMs for security. Patching and configuration management also helps in securing VM workloads.

A few standard workload security controls cannot work successfully in the cloud (e.g., using an antivirus solution inside a container) and certain workloads are not essentially needed or require deep modification to be effective in the cloud environment.

It is recommended to use cloud-native and services (e.g., AWS Systems Manager) for securing workloads because these services are lightweight, can integrate well with APIs and automation components of the respective cloud platforms, and can be easily used.

The impact of cloud environment on standard workload security controls are as follows:

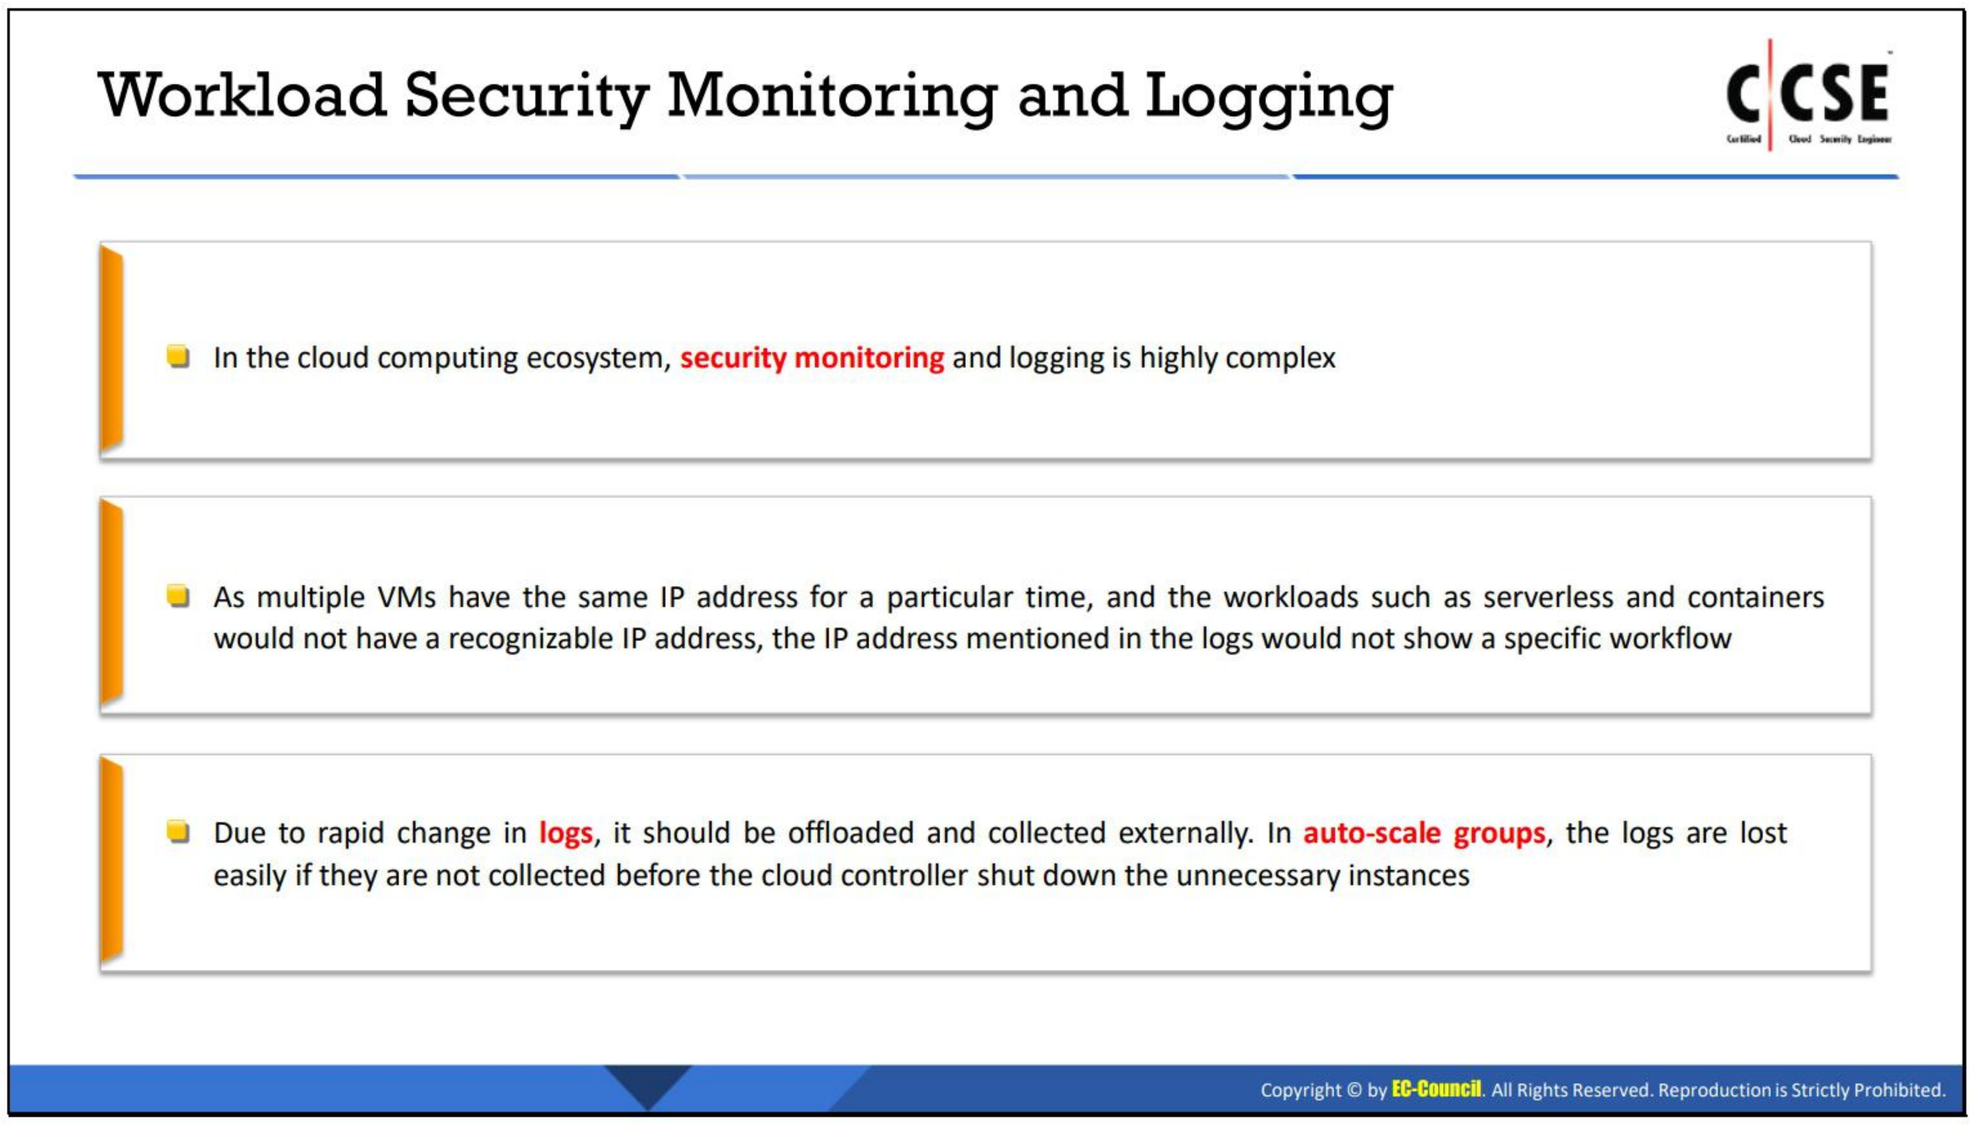
- It is not possible to run software agents for non-VM-based workloads. These agents become useless for serverless applications.
- Traditional agents hinder performance in a cloud environment. To allow efficient use of resources and better workload distribution, it is better to use lightweight agents that have lower compute requirements. Similarly, agents that are not designed for cloud computing will consume more compute capacity than allowed and will increase processing overheads. This will result in increased expenditure.
- Agents in cloud environments must support dynamic cloud workloads and deployment patterns such as auto-scaling. Additionally, agents must be able to discover the control plane and the management plane.

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

- The management plane of the agent should work at the speed of auto-scaling and offer elasticity. Traditional tools cannot provide this degree of velocity.
- The agent should not become a security risk in the cloud environment and is not supposed to increase the attack surface. The agent can increase the attack surface for the following reasons.
  - The agent can be used as an attack vector if it ingests signatures and configuration changes.
  - Additional ports must be opened for certain agents to work, which increases the attack surface.
- The changes that are not approved to the running immutable instance can be detected using file integrity monitoring. File integrity monitoring is a lightweight yet strong security control that can be used for immutable workloads as it provides zero false positives.
- Long-running VMs that run standard security controls should be separated on the network, changing the way they are managed. There might be difficulty in connecting the management tool to a VM deployed in a private network subnet. Running the management tool in the same subnet may increase costs drastically.
- Usually, cloud workloads that run-in isolation are less resilient because of the abstraction. Therefore, it is essential to provide disaster recovery to them.

hide01.ir

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



Workload Security Monitoring and Logging

Workload security plays a crucial role in protecting mission-critical workloads deployed in the cloud because cyberattacks are becoming more sophisticated and the threat landscape is evolving over time. Security teams working in organizations that leverage the cloud for hosting their workloads must have a well-defined strategy along with the right controls to ensure the security of cloud workloads.

A mature cloud security strategy that focuses on the following areas is essential to secure cloud workloads.

- **Reduce the attack surface**: To effectively reduce the likelihood of attacks, understand the attack surface and make efforts to decrease its size. This will help ensure overall cloud security. Continuous logging and monitoring constitute one of the security controls that must be implemented to reduce the attack surface as it helps security teams align their security strategy for ongoing protection.
- **Detect breaches/attacks**: The earlier an attack is detected, the less damage it can cause. The time required to detect an attack can be reduced by implementing a robust cloud security strategy. Logging and monitoring of workloads and using services that can detect breaches/attacks at the network and compute-resource levels are required to detect attacks as early as possible.
- **Respond to attacks**: Respond to attacks after they are detected. Implement an effective threat response strategy to respond to attacks as soon as possible. Security teams need to investigate and identify real threats and eliminate false alarms by analyzing attack vector information recorded by monitoring tools and logs. They can then create the first

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

line of defense by using auto-remediation features such as webhooks and automation scripts that are integrated with alerts.

To collect security logs from multiple sources at the cloud-scale, security teams can use a cloud-native security information event management (SIEM) solution. SIEM solutions can help quickly analyze and detect threats by correlating data from various components of the architecture, no matter where they are deployed. Additionally, these solutions provide better visibility into malicious events, network filtrations, data anomalies, etc., and easy reporting capabilities.

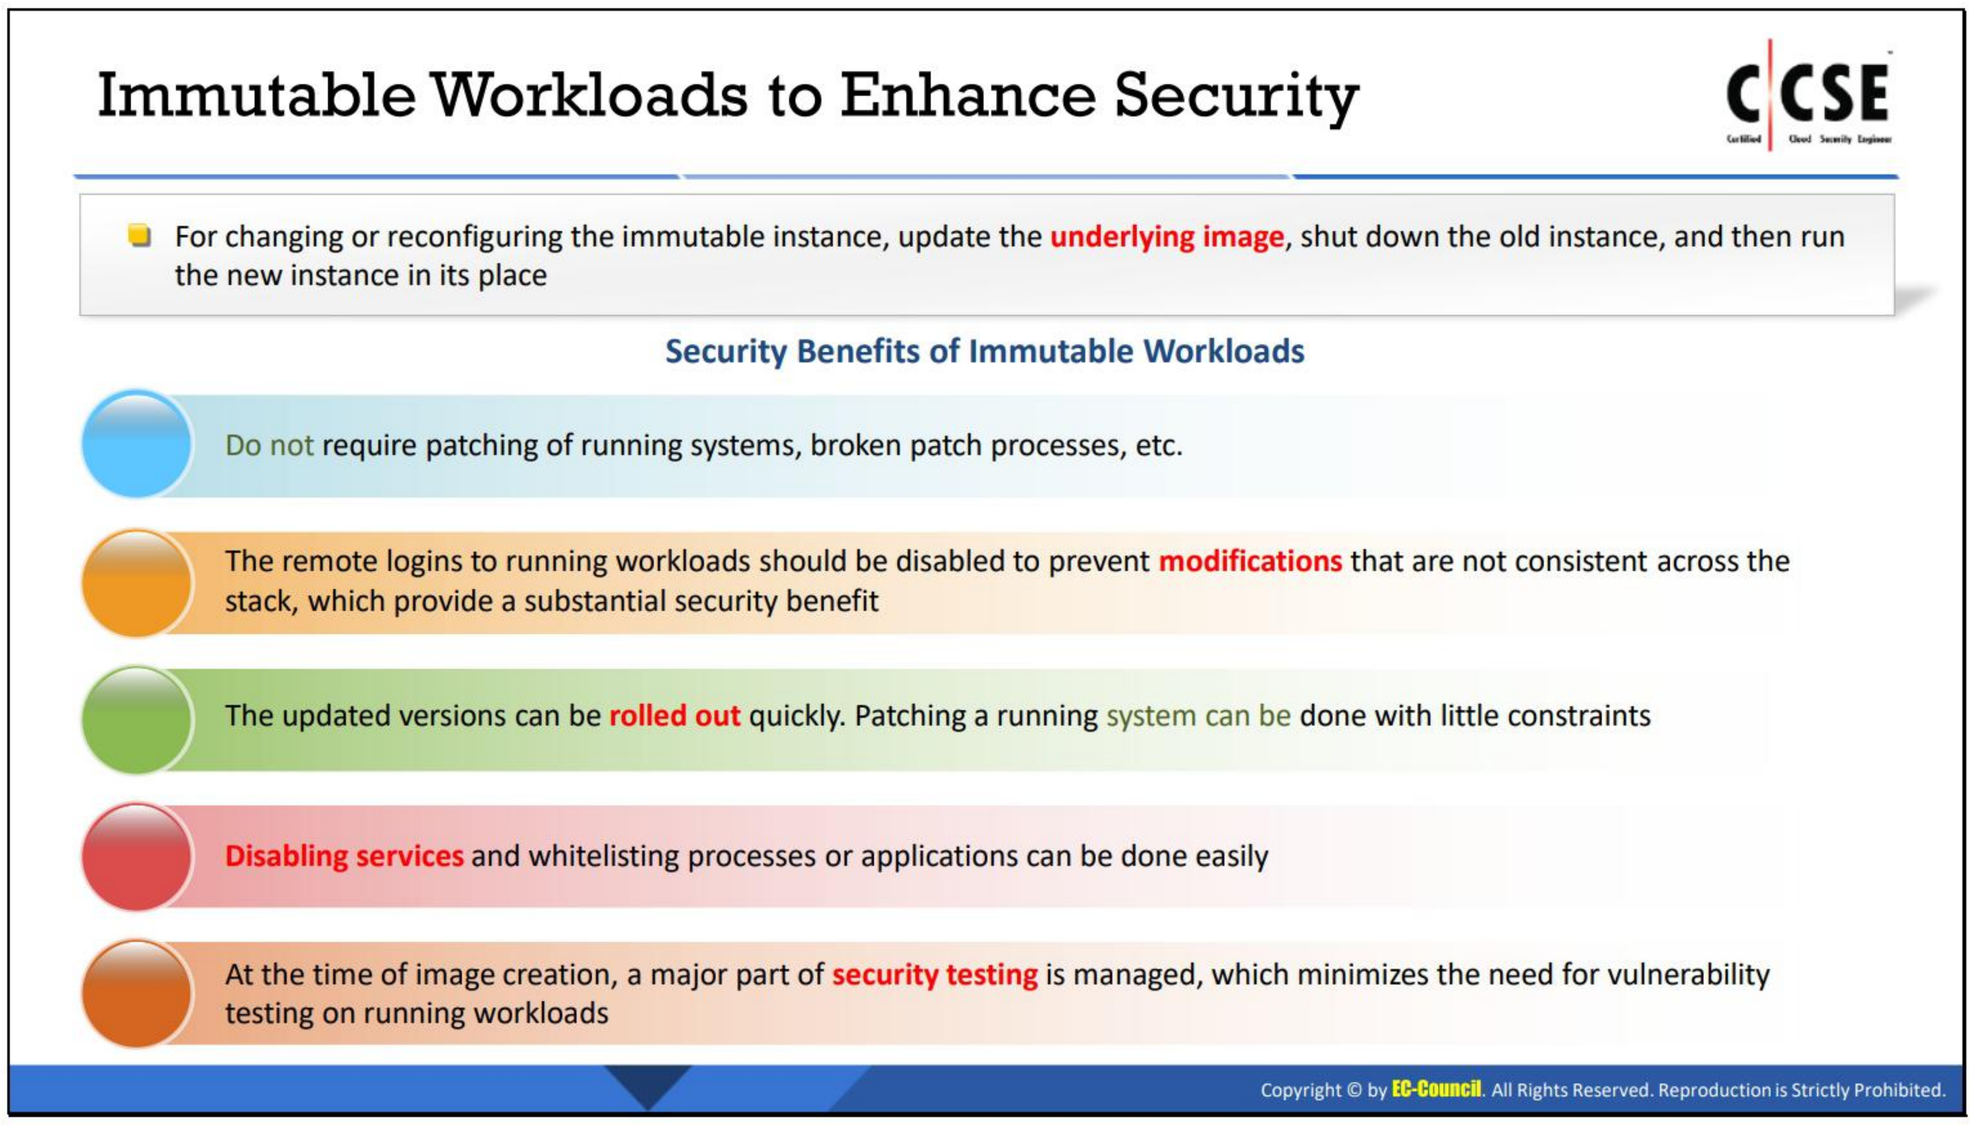
In the cloud computing ecosystem, security monitoring and logging is highly complex:

- As multiple VMs have the same IP address for a particular time and the workloads such as serverless and containers may not possess a recognizable IP address, the IP address mentioned in the logs would not show a specific workflow. In other words, IP addresses become useless in the cloud; thus, unique identifiers must be used in logs to detect assets or to support ephemeral systems.
- Due to rapid changes in logs, the logs should be offloaded and collected externally. In auto-scale groups, the logs are lost easily if they are not collected before the cloud controller shuts down the unnecessary instances.
    - Cloud storage and networking expenses are dependent on logging architectures. For example, it will be expensive to send all logs from public cloud instances to an on-premises SIEM because of the additional fees charged for internal storage and networking.

hide01.ir

Comprehensive logging options are available with all prominent CSPs for relevant service signals. Analyze and monitor the logs of the data plane as well as the logs of the control plane to ensure end-to-end security. These logs are critical while conducting investigations because they include network flows, service activity logs, IAM logs, and data ingress/egress logs.

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



Immutable Workloads to Enhance Security

While running instances launched dynamically that are based on an image, auto-scaling and containers work best as they allow shutting down those instances for capacity when they are not needed without breaking an application stack. Do not apply patches and other changes to a running workload because this will not change the image. Therefore, new images will be out of sync with the changes applied to the running workload manually. These VMs are referred to as immutable. Immutable means not susceptible to change; therefore, VMs that are immutable do not allow any changes to them.

For changing or reconfiguring an immutable instance, update the underlying image, shut down the old instance, and then run the new instance in its place.

The below types of immutable VMs exist.

- **Full mutable VM**: Full mutable VMs allows no changes/patches to running workloads. They can be changed only by completely replacing running instances with a new image.
- **Not completely mutable**: Not completely mutable VMs allows updating the OS by pushing new images. They allow pushing code updates into running VMs by leveraging alternative deployment techniques. These pushes occur automatically via an automated process, and no one logs in to running VMs for making local changes.

Significant security benefits are provided by immutable workloads, as described below:

- No need to patch running systems, broken patch processes, etc. Replace the running systems with a new gold master.
- Remote logins to running workloads should be disabled to prevent modifications that are not consistent across the stack, which provide a substantial security benefit.

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

- Updated versions can be rolled out quickly. Patching a running system is done with little constraints.
- Disabling services and whitelisting processes or applications can be done easily.
- At the time of image creation, a major part of security testing is managed, which minimizes the need for vulnerability testing on running workloads.

Requirements of Immutable Workloads

Immutable workloads need the following requirements:

- A consistent image creation process and automation to help with updating of deployments. Produce new images on a regular basis to apply patch and malware signature updates.
- Integrate security testing into the image creation and deployment process (e.g., source code tests). Integrate vulnerability assessments if using VMs and containers.
- To disable logins and restrict services before images are deployed and used for production VMs, mechanisms are required by image configurations.
- For certain workloads, there may be a need for a process that enables logins to workloads for troubleshooting.

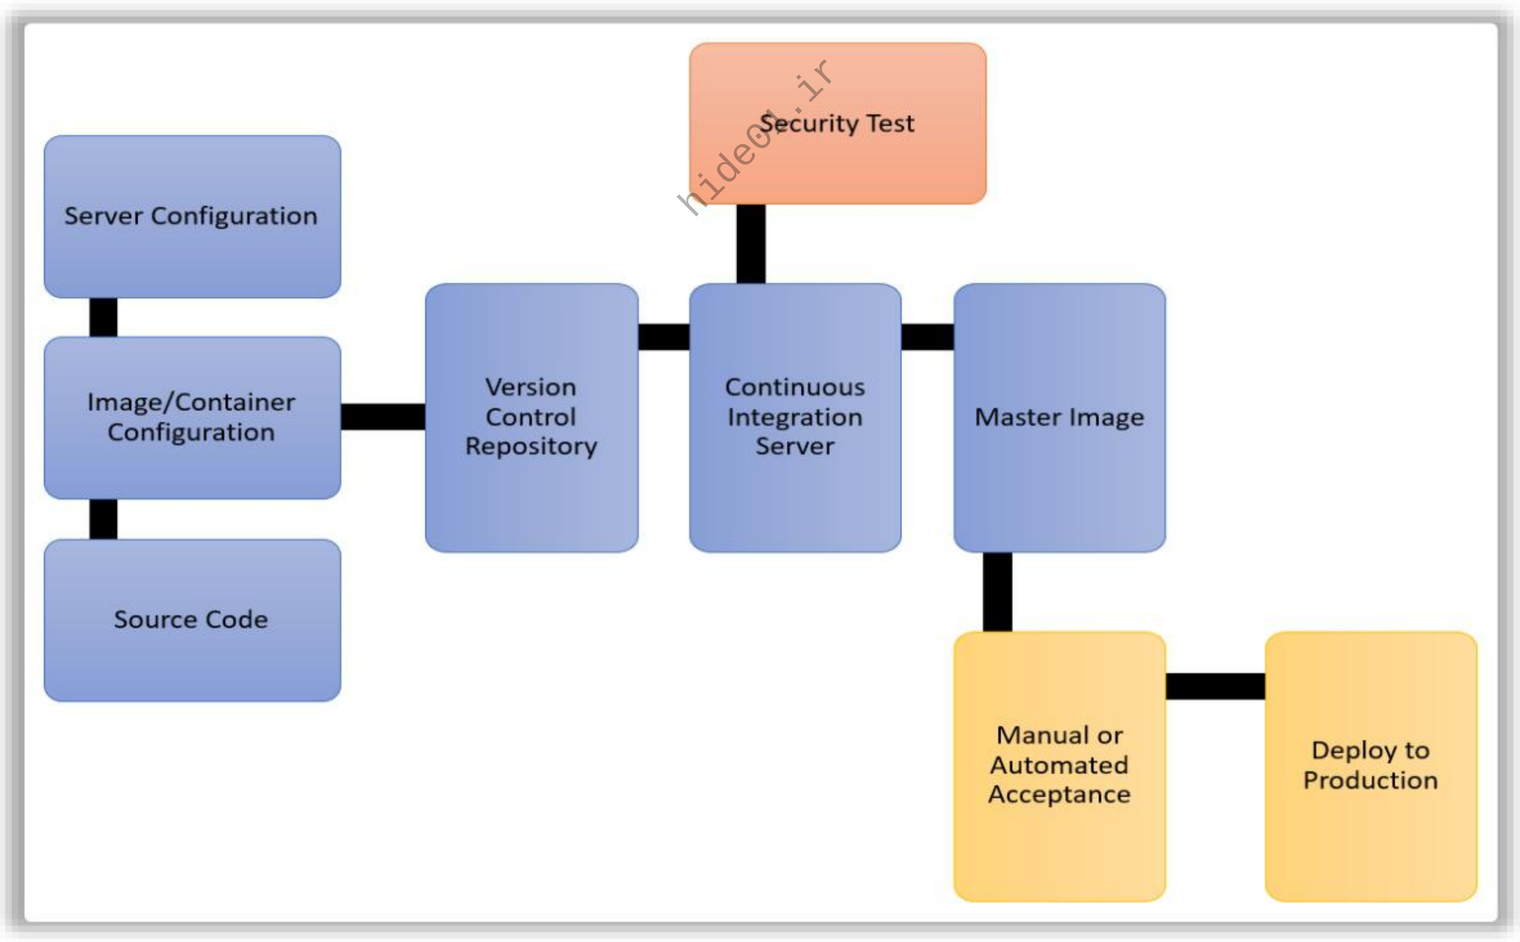
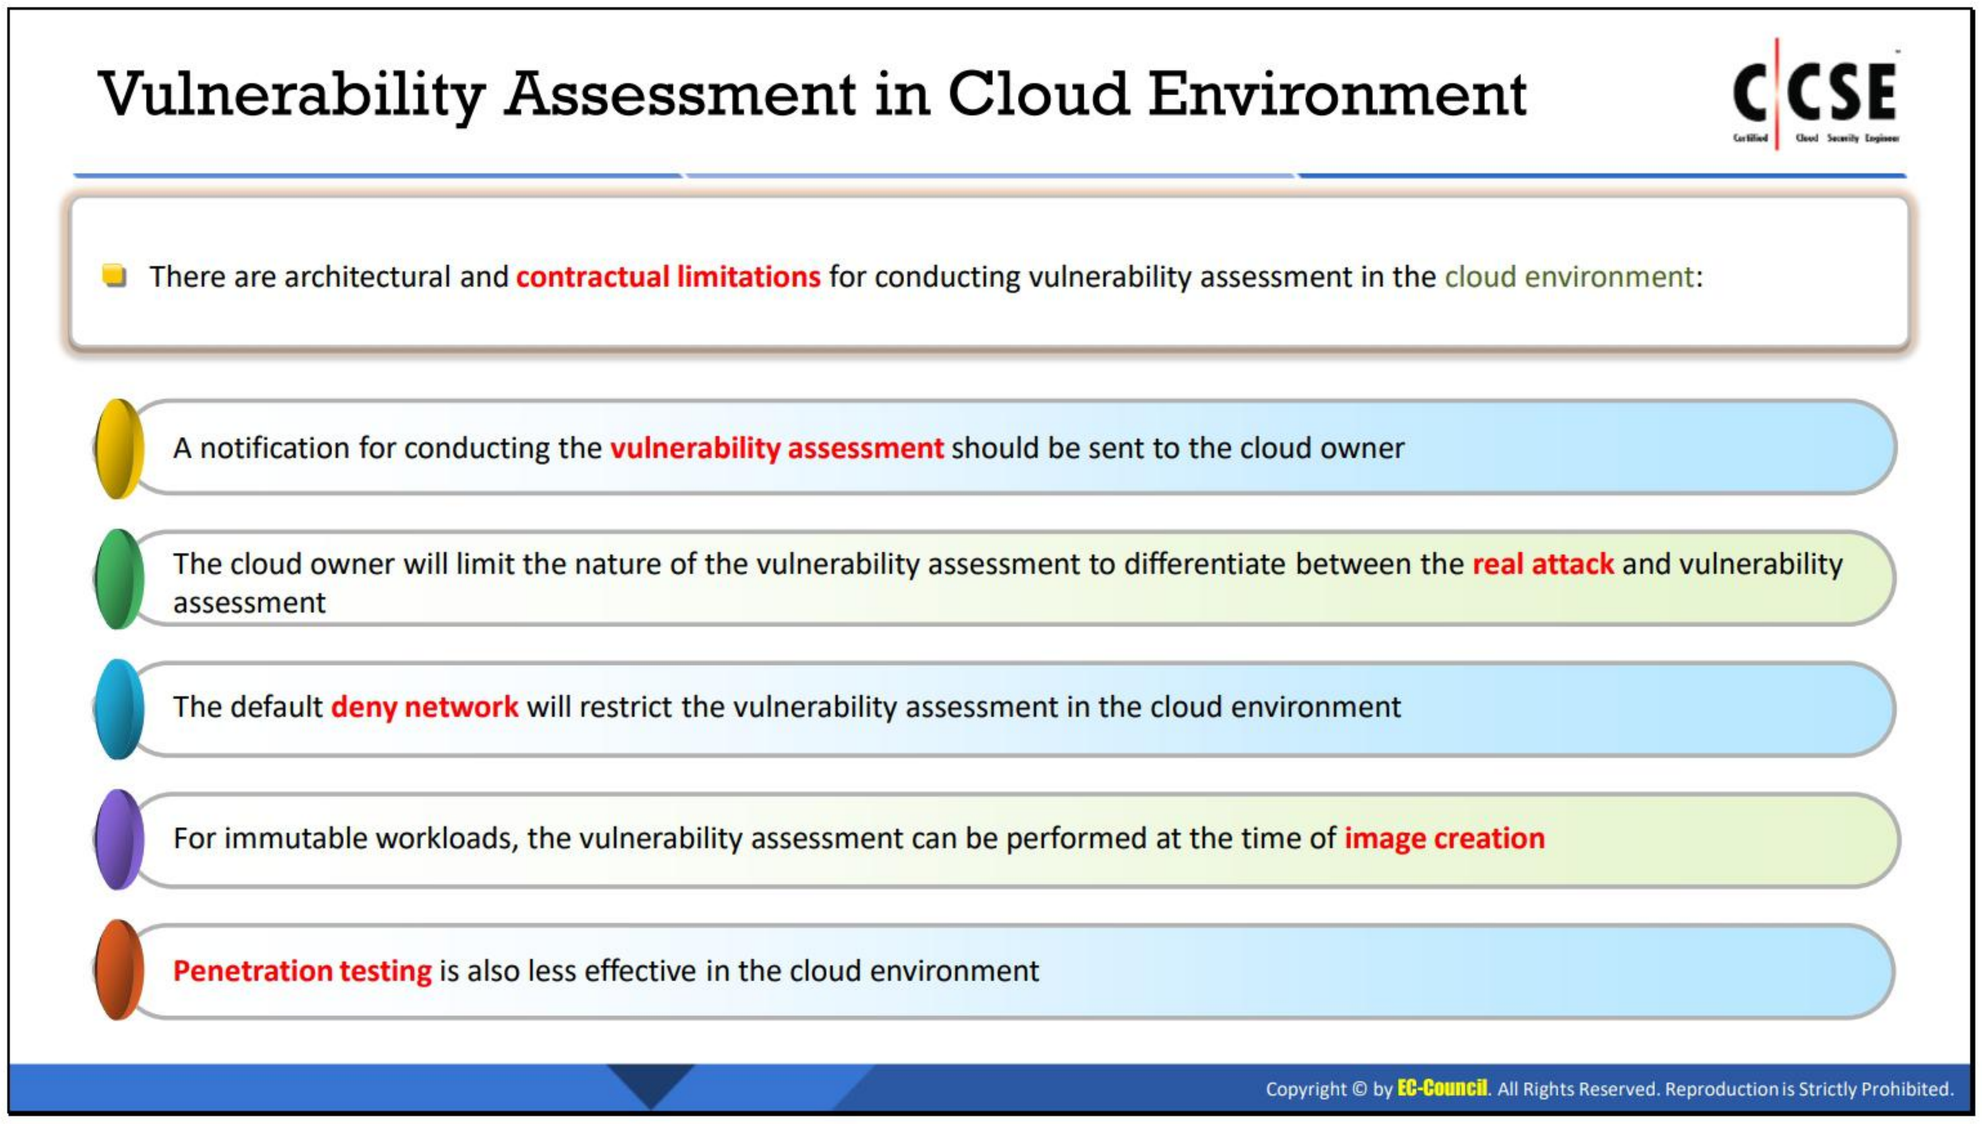


Figure: 2.9: Deployment Pipeline for Creating Images for Immutable Containers and VMs

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



Vulnerability Assessment in Cloud Environment

A vulnerability assessment in the cloud context refers to the process of defining, detecting, categorizing, and prioritizing vulnerabilities in cloud systems, apps, and network infrastructure. Through vulnerability assessments, organizations can understand and react to threats faced by their cloud environments by gaining the necessary knowledge and awareness of the risk the threats pose. Effective vulnerability assessments help organizations make it harder for attackers to compromise their cloud environments.

Any type of weakness or misconfiguration in software that can be exploited or abused by an attacker is known as a vulnerability. An organization evaluates and prioritizes vulnerabilities and designs methods to remedy the situation accordingly via a vulnerability assessment. The assessment helps measure the effectiveness of those remedies. Vulnerability assessments provide a set of findings that help security teams know which vulnerabilities must be remediated first to minimize the risk associated with these vulnerabilities.

Vulnerability assessment is a key element of a vulnerability management framework. These assessments are performed to assess the current security posture of an organization and know what necessary changes to the vulnerability management plan are required. It is considered the first step for enhancing IT security and provides multiple benefits.

Benefits of vulnerability assessment include the following:

- Identifying the key information assets of an organization
- Deciding the vulnerabilities that threaten the security of those assets
- Providing recommendations to strengthen the security posture
- Mitigating risks

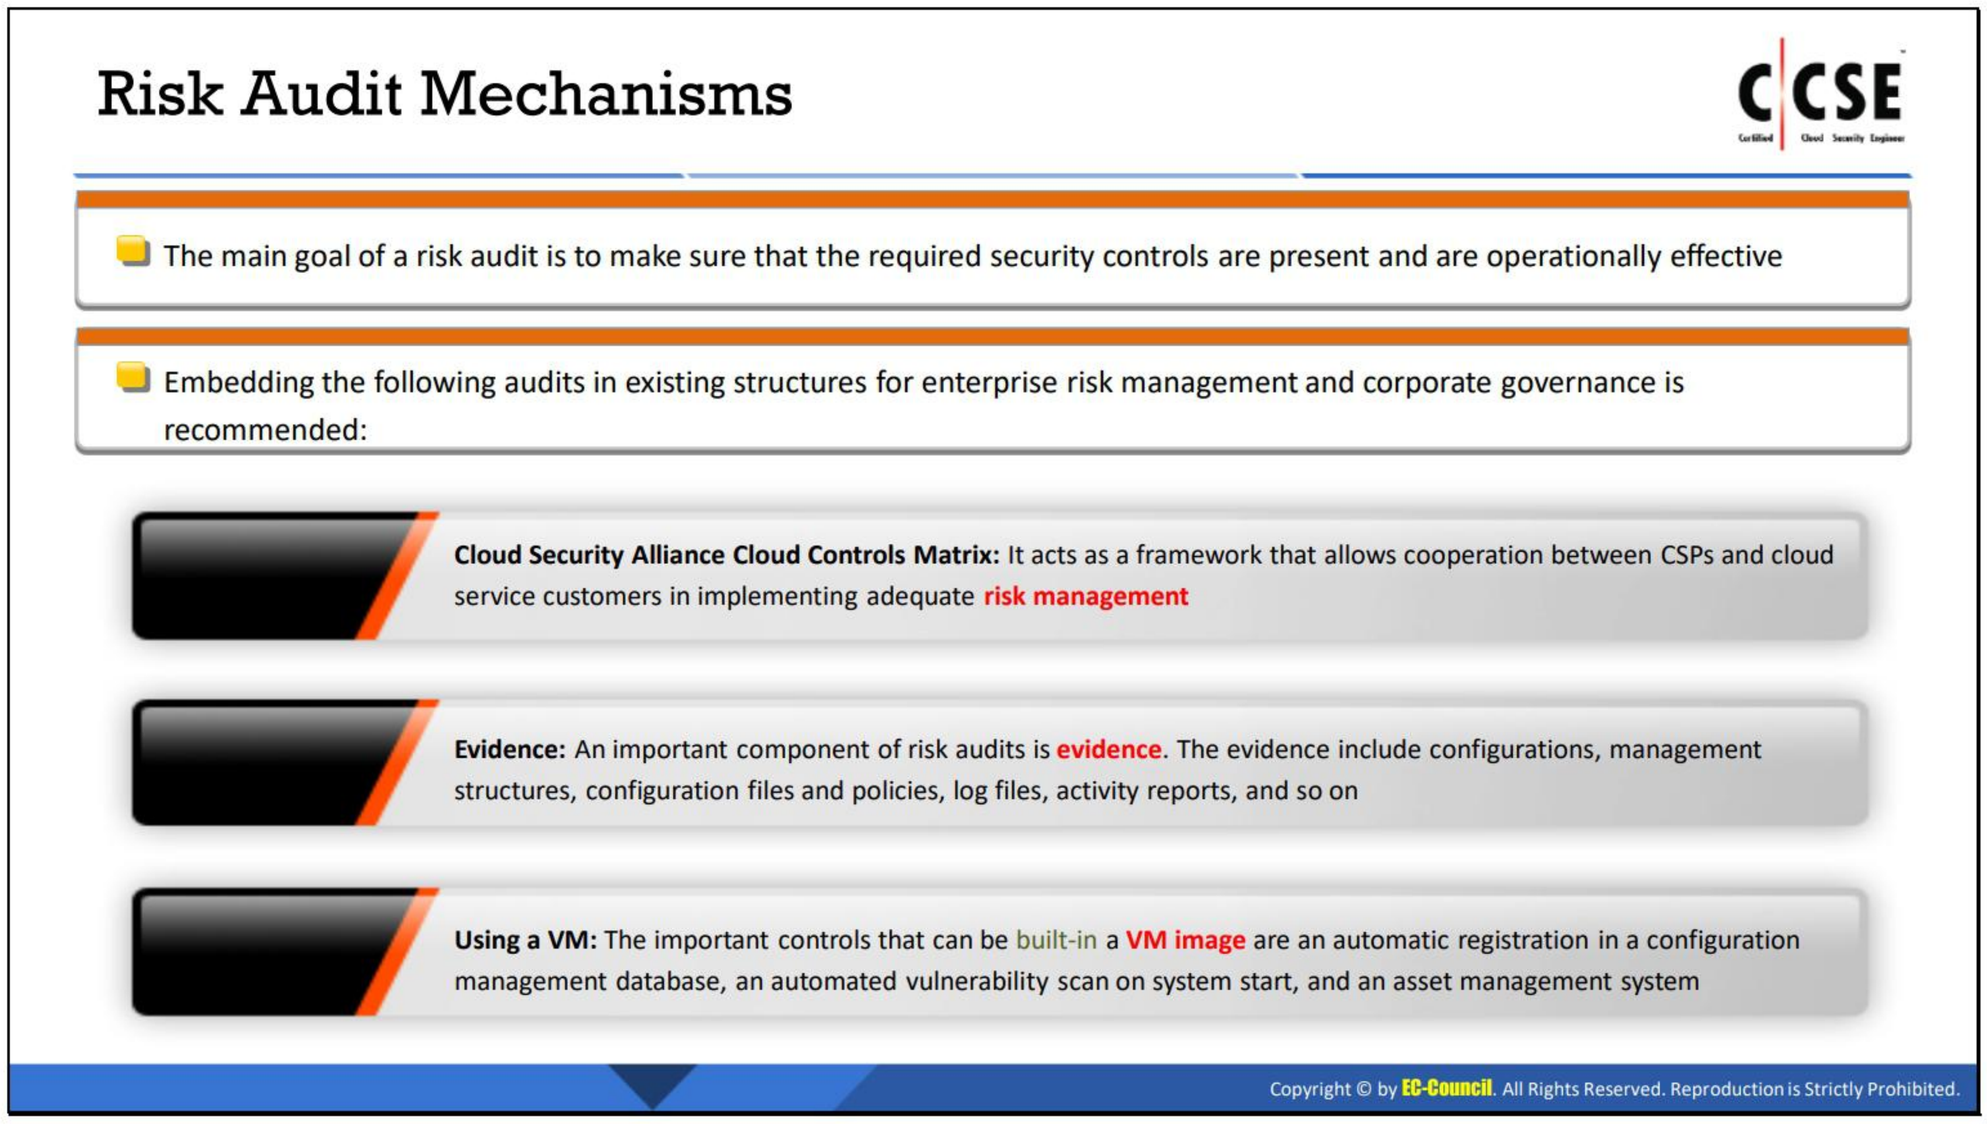
Certified Cloud Security Engineer Copyright © by **EC-Council**
All Rights Reserved. Reproduction is Strictly Prohibited.

Vulnerability assessments can be categorized mainly into two approaches: network-based scans and agent-based scans. Network-based scans help identify potential network security attacks, whereas agent-based scans help detect and locate vulnerabilities in workstations, servers, or other network hosts. The agent-based scan approach best suits cloud environments. However, there are certain architectural and contractual limitations for conducting vulnerability assessments in the cloud environment:

- A notification for conducting the vulnerability assessment should be sent to the cloud owner (public or private). This notification will allow the cloud owner (the CSP) to distinguish between an assessment and a real attack.
- Default deny networks will restrict automated network vulnerability assessments (such as a firewall) in the cloud environment. To allow vulnerability assessments, open up ports or utilize an agent on an instance.
- For immutable workloads, the vulnerability assessment can be performed at the time of image creation. Assessments can be conducted with fewer network restrictions because they are not performed in production and the assessment process is automated. This increases the assessment surface.
- In a cloud environment, penetration testing is less affected because it leverages the same scope as an attacker.

hide01.ir

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



Risk Audit Mechanisms

The main goal of a risk audit is to make sure that the required security controls are present and are operationally effective. There are many reasons for having audits and the most obvious reasons are compliance and regulatory related. Moreover, a quality system employs more internal audits to ensure its quality is up-to-the-mark and cloud customers demand more demonstration of quality.

Embedding audits in existing structures for enterprise risk management and corporate governance is recommended. All the essential requirements for risk controls, such as technical, legal, jurisdictional, contractual, and compliance requirements, can be aggregated. Compliance validates awareness and adherence to organizational obligations such as applicable laws, corporate social responsibility, regulations, policies, and strategies. The compliance process involves assessing the state of that adherence and awareness, assessing the risks and cost of non-compliance against the cost to achieve compliance, and initiating corrective actions that are needed. The key tool that approves or disapproves compliance is a risk audit. It measures how well an organization is adhering to compliance requirements.

The Cloud Security Alliance Cloud Controls Matrix

The Cloud Security Alliance's cloud controls matrix is very important in any cloud environment because it acts as a framework that allows cooperation between CSPs and cloud service customers in implementing adequate risk management.

An important component of risk audits includes evidence to show that the controls are functional. This evidence includes configurations, management structures, configuration files and policies, log files, activity reports, and so on. These are called artifacts and they are needed for audits and compliance. Artifacts are evidence to support compliance activities. Both

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

customers and providers are responsible for producing and managing their artifacts, respectively. The drawback is the gathering of evidence, which is a costly effort. However, cloud computing can give a new angle to the audit process.

Cloud Computing Audit Characteristics

Audit requirements and the audit process are affected by the characteristics of cloud computing in several ways.

Cloud service customers are dependent on CSPs, who in turn rely on other providers that provide hosting services. Cloud infrastructure is located at a hosting facility and it is a dependency. Issues occurring at the hosting facility can affect or impact the cloud infrastructure. From a cloud service customer's perspective, the hosting provider has the required controls; thus, they cannot have audits on their CSPs because of the hosting and CSP restrictions.

Providers consider having on-premise audits done by cloud customers a security risk because they provide multitenant services and cannot allow their customers to conduct individual audits as it may pose security risks.

Customers can make their providers offer independent (third-party) audit reports, so that they can review these report results for the scope and relevance of their requirements. Customers can use their contracts to include third-party audits and attestations to satisfy their compliance requirements. The contract clauses should be clear enough for the validation of controls that are important to cloud customers. Most providers conduct internal and external (third-party) audits to ensure compliance with internal and external requirements.

video01.ir

Providers may disclose these audit reports under a nondisclosure agreement and customers need to enter into a legal agreement with the providers to access the attestations and audit reports for risk assessments or certain other evaluation purposes. Providers make their customers sign an NDA because of legal or contractual requirements with their audit firms. Providers need to provide rigorous third-party attestations to assure their customers that their contractual and regulatory requirements are met. The legal statement of a third party about the audit findings is called an attestation. Attestations should be based on well-defined industry standards with clearly defined scopes and the list of specific controls that are evaluated.

Providers can publicly release these documents on the central repository offered by the Cloud Security Alliance STAR Registry. Standards such as SSAE 16 attest that documented controls are working as they are designed. The scope of controls may not be defined by the standard, but both the scope of controls and the list of specific controls evaluated are required for a full evaluation. Certifications and attestations do not necessarily apply to all services offered by cloud providers; therefore, providers should be clear about the features and services that are covered under their attestations and certifications.

Providers include certain terms of service in their agreement with cloud customers to stop their customers from having technical assessments and audits such as vulnerability testing; thus, customers need to take prior permission when they want to perform such assessments. Providers do this to differentiate between an attack and a legitimate assessment. All attestations and certifications are point-in-time activities and are subjected to expiry. Providers need to keep their

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

published results current (latest) else they may expose their customers to risks of non-compliance. Customers have the responsibility to ensure that they rely on current results and track any changes in their providers' status over a period of time.

If the essential cloud characteristics are optimally used, cloud computing can significantly improve assurance and transparency. Cloud infrastructure involves high degrees of service automation and self-service. If these principles are applied on audit requirements and these requirements are incorporated in developing the cloud infrastructure, reaching new levels of assurance is possible.

Contractual agreements such as hosting agreements and service level agreements distribute the responsibilities and risks between cloud service customers and CSPs. This makes all parties document and understand the liability of failure of any one or more controls and the corresponding risks in a better way.

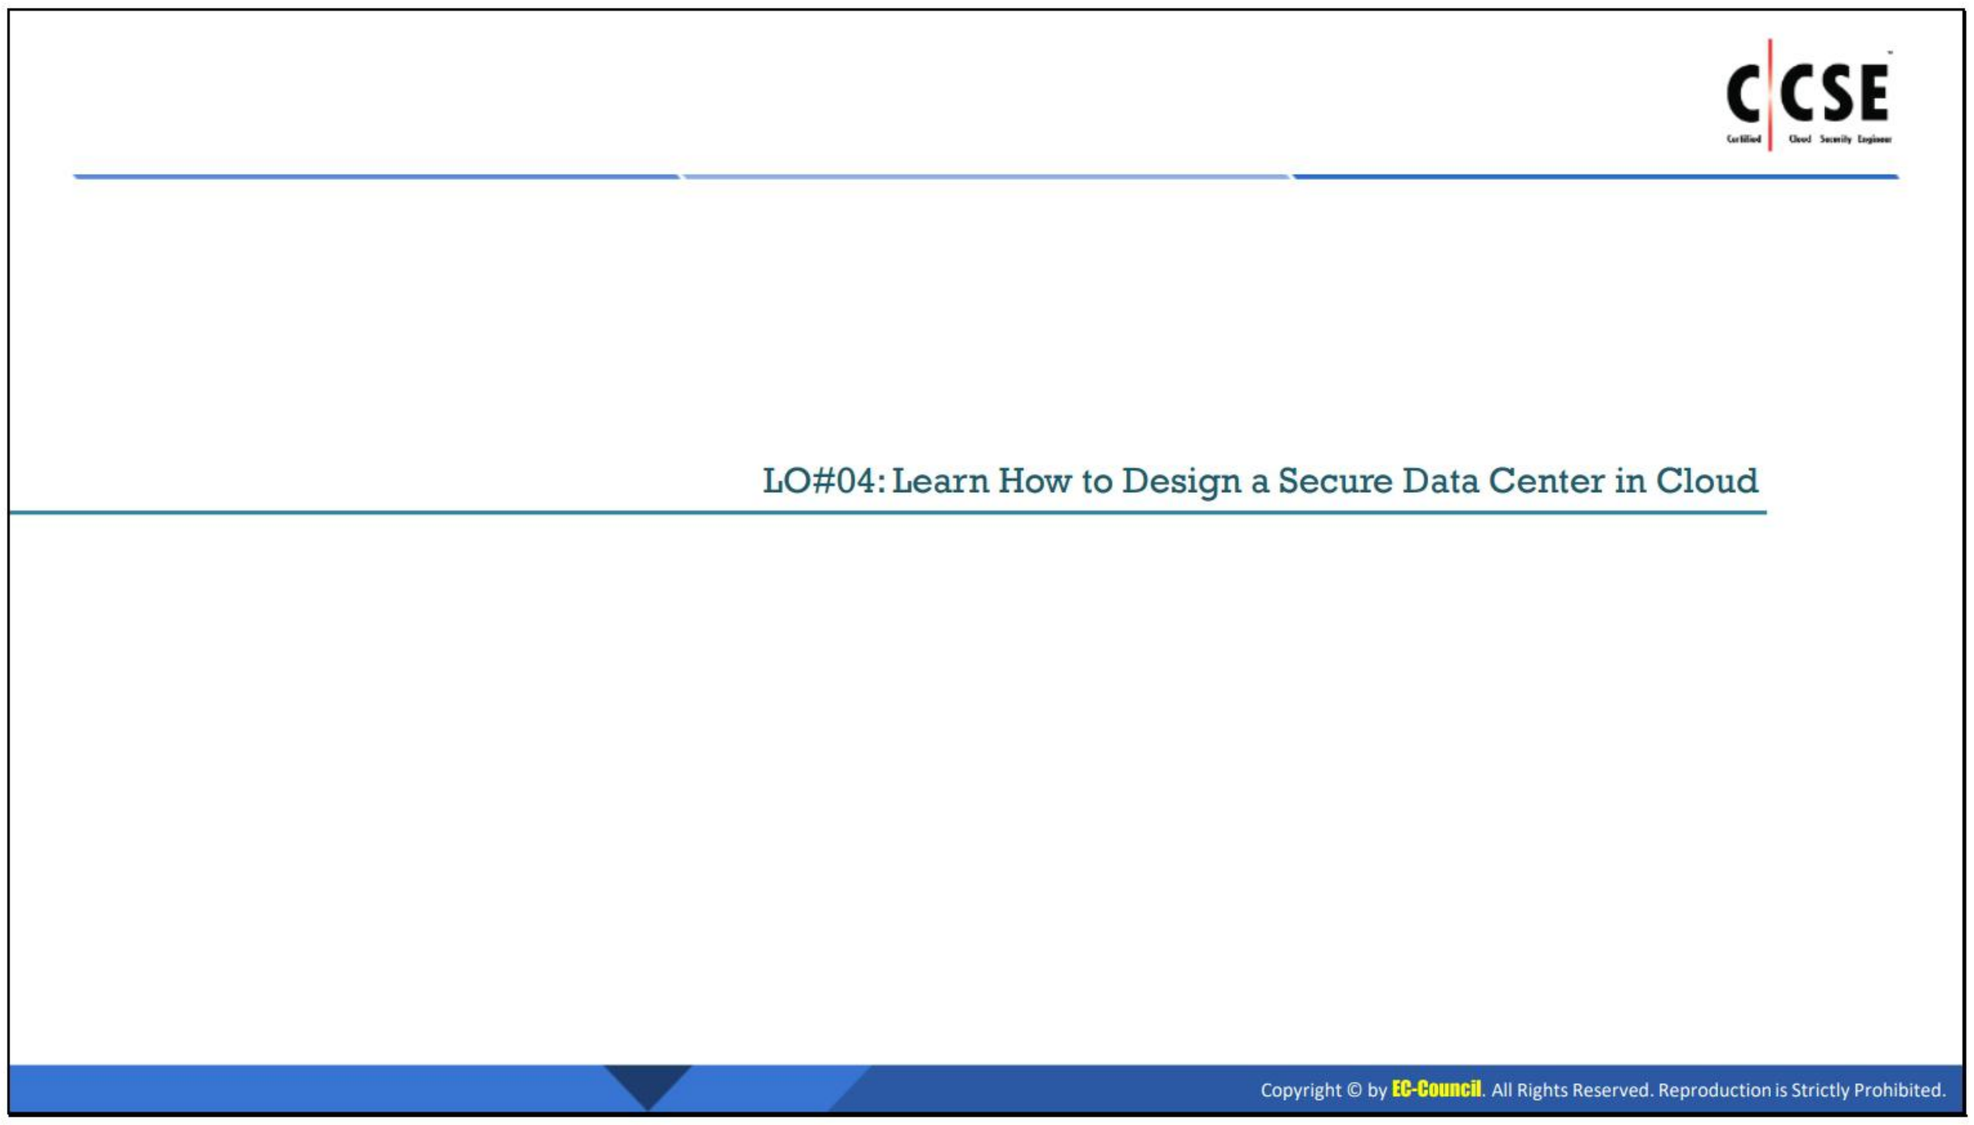
Using a VM

Service automation is very important in generating evidence automatically. For example, a specified configuration can be used to build a VM image. The logs of the build process and configuration baseline then provide evidence that adequate controls will be implemented by all instances of this VM.

The important controls that can be built into a VM image include an automatic registration in a configuration management database, an automated vulnerability scan on system start, and an asset management system. Any number of management and control agents can be implied by using the configuration, such as DLP agents, VM-level firewalls, and automated log file generation. This will result in the automatic generation of evidence.

Cloud computing's self-service provisioning and automation can be improved to perform continuous auditing where the presence and effectiveness of controls can be tested in a near real-time and continuous basis. The evidence can be made accessible to authorized cloud customers in a dashboard style. This allows the customer to collect evidence from the CSP as well as from the upstream CSPs in the delivery chain.

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



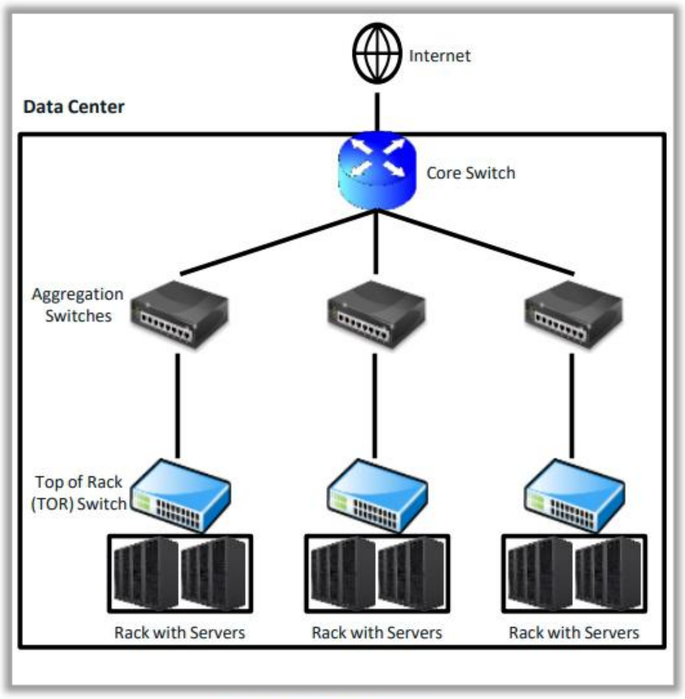
LO#04: Learn How to Design a Secure Data Center in Cloud

Data centers provide centralized data processing, which helps secure cloud platforms and infrastructure from attacks. Integrity and availability of data present in the data center are important security factors. Cloud security engineers need to know how to secure the data center and operations in the cloud. This section explains the data centers in detail and focuses on the physical, logical, and environmental aspects of designing a secure data center.

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

Cloud Data Center

- Cloud data center is a facility containing storage, compute, network, and other resources to provides centralized data processing ability



Data Center Anatomy

Tier Level	Data Center Classification
I	Single non-redundant distribution path serving the IT equipment Non-redundant capacity components Basic site infrastructure with 99.671% availability
II	All Tire I requirements are fulfilled Redundant site infrastructure capacity components with 99.741% availability
III	Fulfills all Tire I and II requirements Multiple independent distribution paths serving IT equipment All IT equipment should be dual-power and fully compatible with the topology of a site's architecture Concurrently maintainable site infrastructure with 99.982% availability
IV	Fulfills all Tier I, Tier II, and Tier III requirements All cooling equipment is independently dual-powered, including chillers and heating, ventilation and air-conditioning (HVAC) systems Fault tolerant site infrastructure with electrical power storage and distribution facilities with 99.995% availability

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

Cloud Data Center

Data centers must be designed based on standards that meet the necessary requirements. The different standards that can be considered are briefly introduced below.

- **Building Industry Consulting Service International Inc. (BICSI):** The BICSI standards are for design of cabling and installation.
- **The International Data Center Authority (IDCA):** The IDCA standard covers data center location, facility structure, and applications.
- **National Fire Protection Association (NFPA):** NFPA standards NFPA 75 and NFPA 76 specify how hot or cold aisle containment is to be performed.

A four-tiered secure, reliable architecture provided by Uptime Institute's Data Center Site Infrastructure Tier Standard Topology, which describes the site level infrastructure topology to handle data center operations, is considered as the baseline for data center design. The four tiers are described below.

- **Tier I (basic data center site infrastructure)**: Tier I infrastructure components include space for IT systems, UPS for handling power outages and spikes, cooling equipment and power production through an engine generator, a fuel cell and fuel storage for power production for 12 hours. The performance and operational impacts of Tier I infrastructure are briefly described below.
  - It has the capacity to meet the site needs.
  - Planned and unplanned works cause systems to be shut down, which affects critical operations.
  - Unplanned outages or failures affect the entire environment.

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

- Maintenance and repair works can be performed only after the complete shutdown of the site infrastructure components. Failure to perform maintenance will cause unplanned outages.

- **Tier II (redundant site infrastructure capacity components)**: Tier II data centers consist of redundant capacity components such as power production engine generators and fuel cells, UPS, energy storage, chillers, heat rejection equipment, pumps, cooling units, and fuel tanks with 12-hour fuel capacity. They also contain a non-redundant distribution path that serves critical infrastructures. The redundant components can be removed from the infrastructure based on a proper plan without shutting down the critical components; however, removal of non-redundant components requires shutdown. The performance and operational impacts of Tier II infrastructure are briefly described below.
  - Disruptions can occur from unplanned and planned activities and from human errors.
  - Unplanned outages of the redundant components can cause failures in critical environments.
  - • Maintenance and repair works can be performed only after the complete shutdown of the site infrastructure components. Failure to perform maintenance will cause unplanned outages.
- **Tier III (concurrently maintainable site infrastructure)**: Tier III data centers have redundant capacity components and multiple independent distribution paths that serve the critical infrastructure. The electrical power back bone covers the electrical power distribution path from the power production systems to the input of the UPS and the critical mechanical equipment. Mechanical distribution path includes the path for heat removal. All the equipment are dual powered and properly installed. To serve critical environments, the electrical power backbone and mechanical distribution path require only one distribution path. The performance and operational impacts of Tier III infrastructure are briefly described below.
  - The critical environment will be impacted from unplanned outages or failures.
  - Operation errors of the components cause service disruptions.
  - The redundant capacity components and distribution paths can be used while planned infrastructure maintenance is performed.
- **Tier IV (fault-tolerant site infrastructure):** Tier IV data center consist of multiple isolated systems that provide redundant capacity components and multiple active distribution paths simultaneously. The equipment are dual powered with fault-tolerant power design and with proper installation to ensure compatibility with the Tier IV architecture. Point of switches can be used for critical environments that does not meet this standard. To prevent a particular event from impacting complementary systems and distribution paths, they are physically isolated in this architecture. Additionally, continuous cooling should be provided within American Society of Heating, Refrigeration, and Air Conditioning Engineers (ASHRAE) temperature guidelines and the duration should be such that it provides cooling until mechanical systems are providing the rated cooling.

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

The performance and operational impacts of the Tier IV infrastructure are

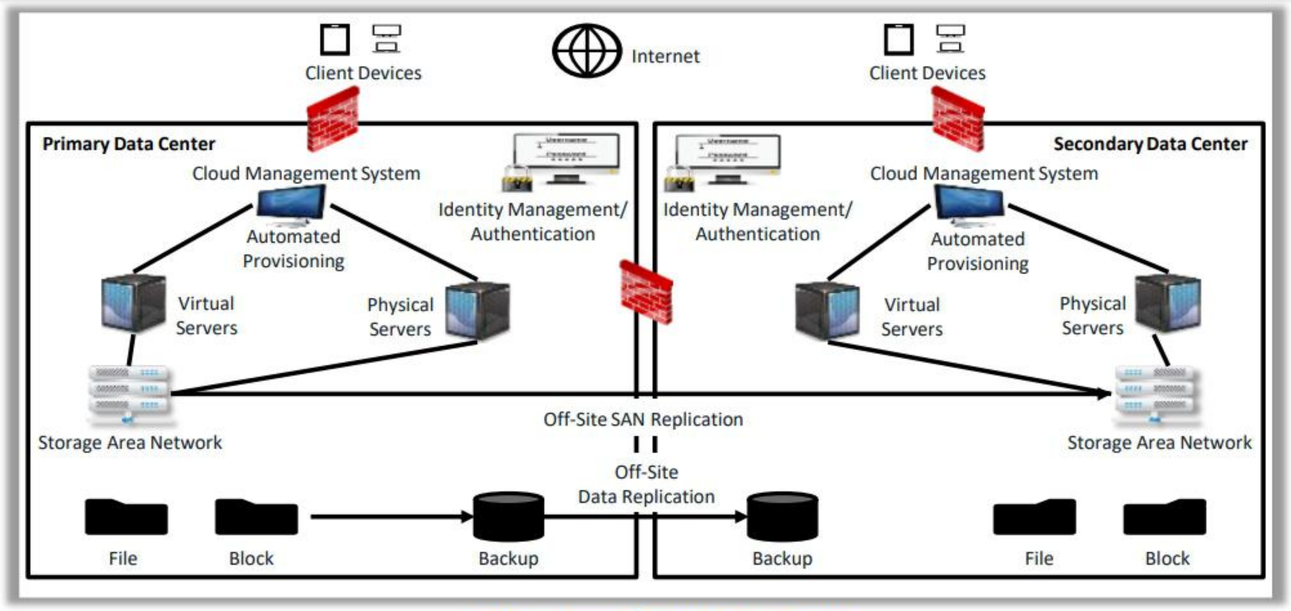
- The Tier IV site will not be impacted by disruptions due to planned or unplanned events.
- Redundant capacity components and distribution paths helps in safe operations during maintenance activities.
- Data center disruption occurs only during the operation of fire alarm, fire suppression, or during an emergency power off.

hide01.ir

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

Architecture of Cloud Data Center

- Cloud data center consists of a primary data center and a secondary data center which consists of servers, SAN, and VMs



Architecture of Cloud Data Center

- In order to secure the data center security, secure the physical, logical, and the environmental design

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

Architecture of Cloud Data Center

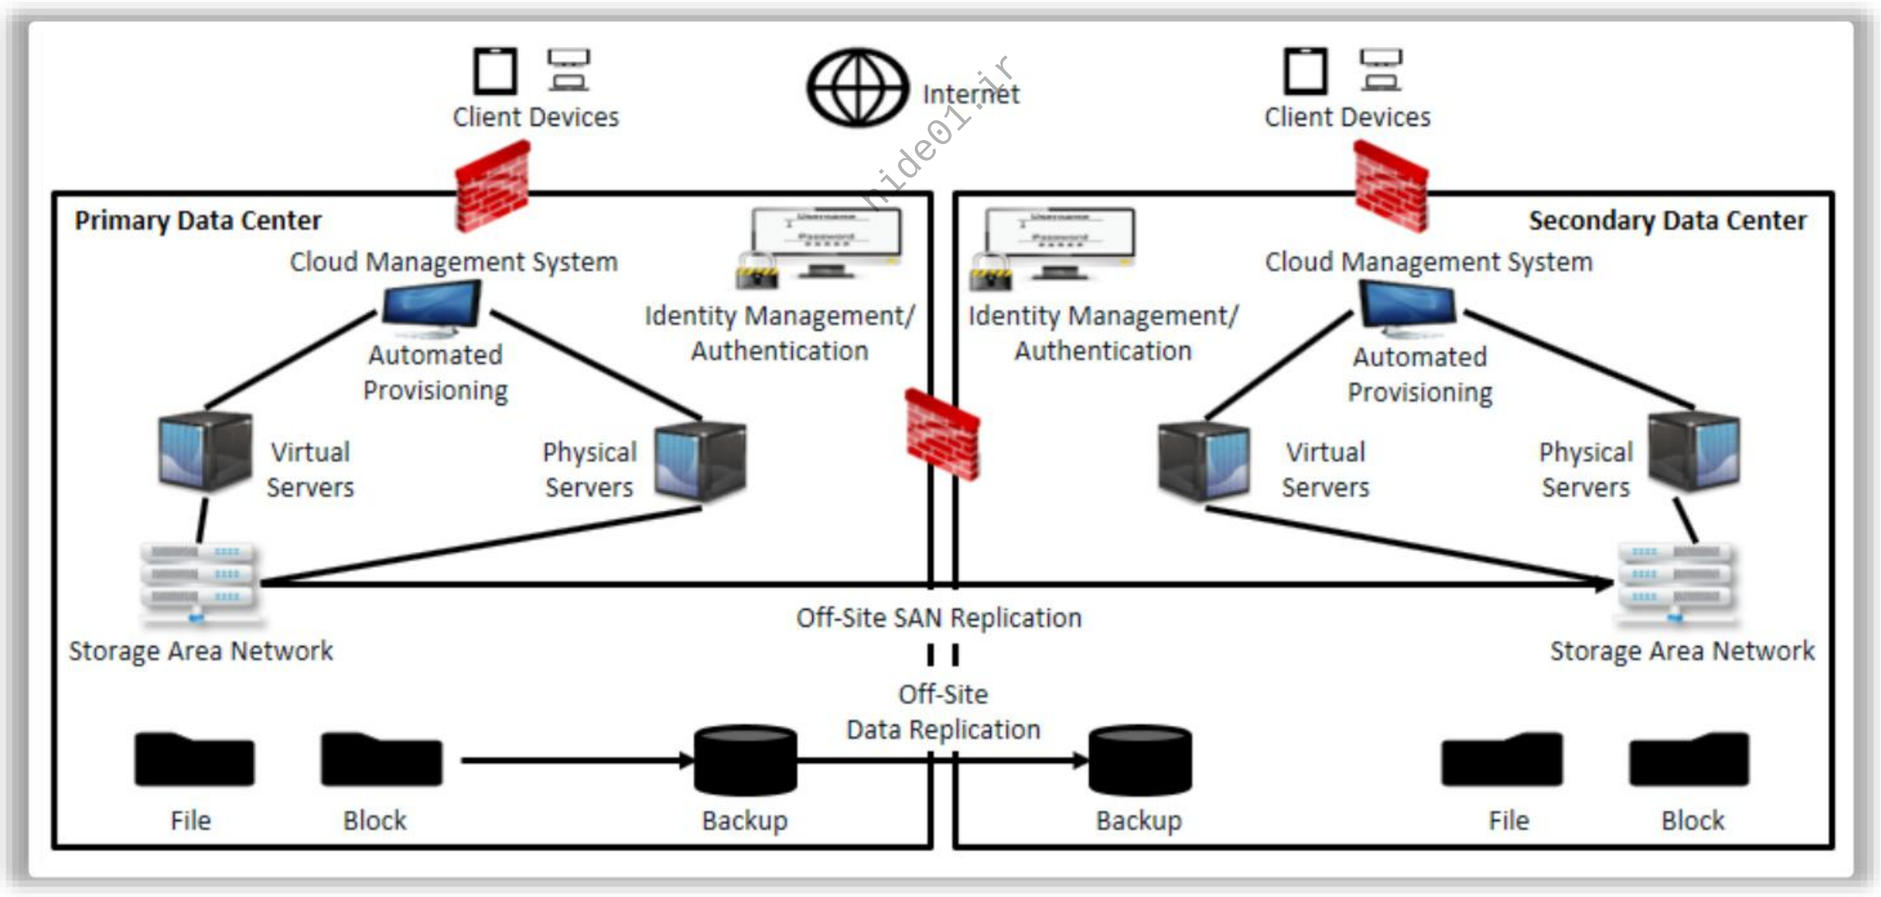
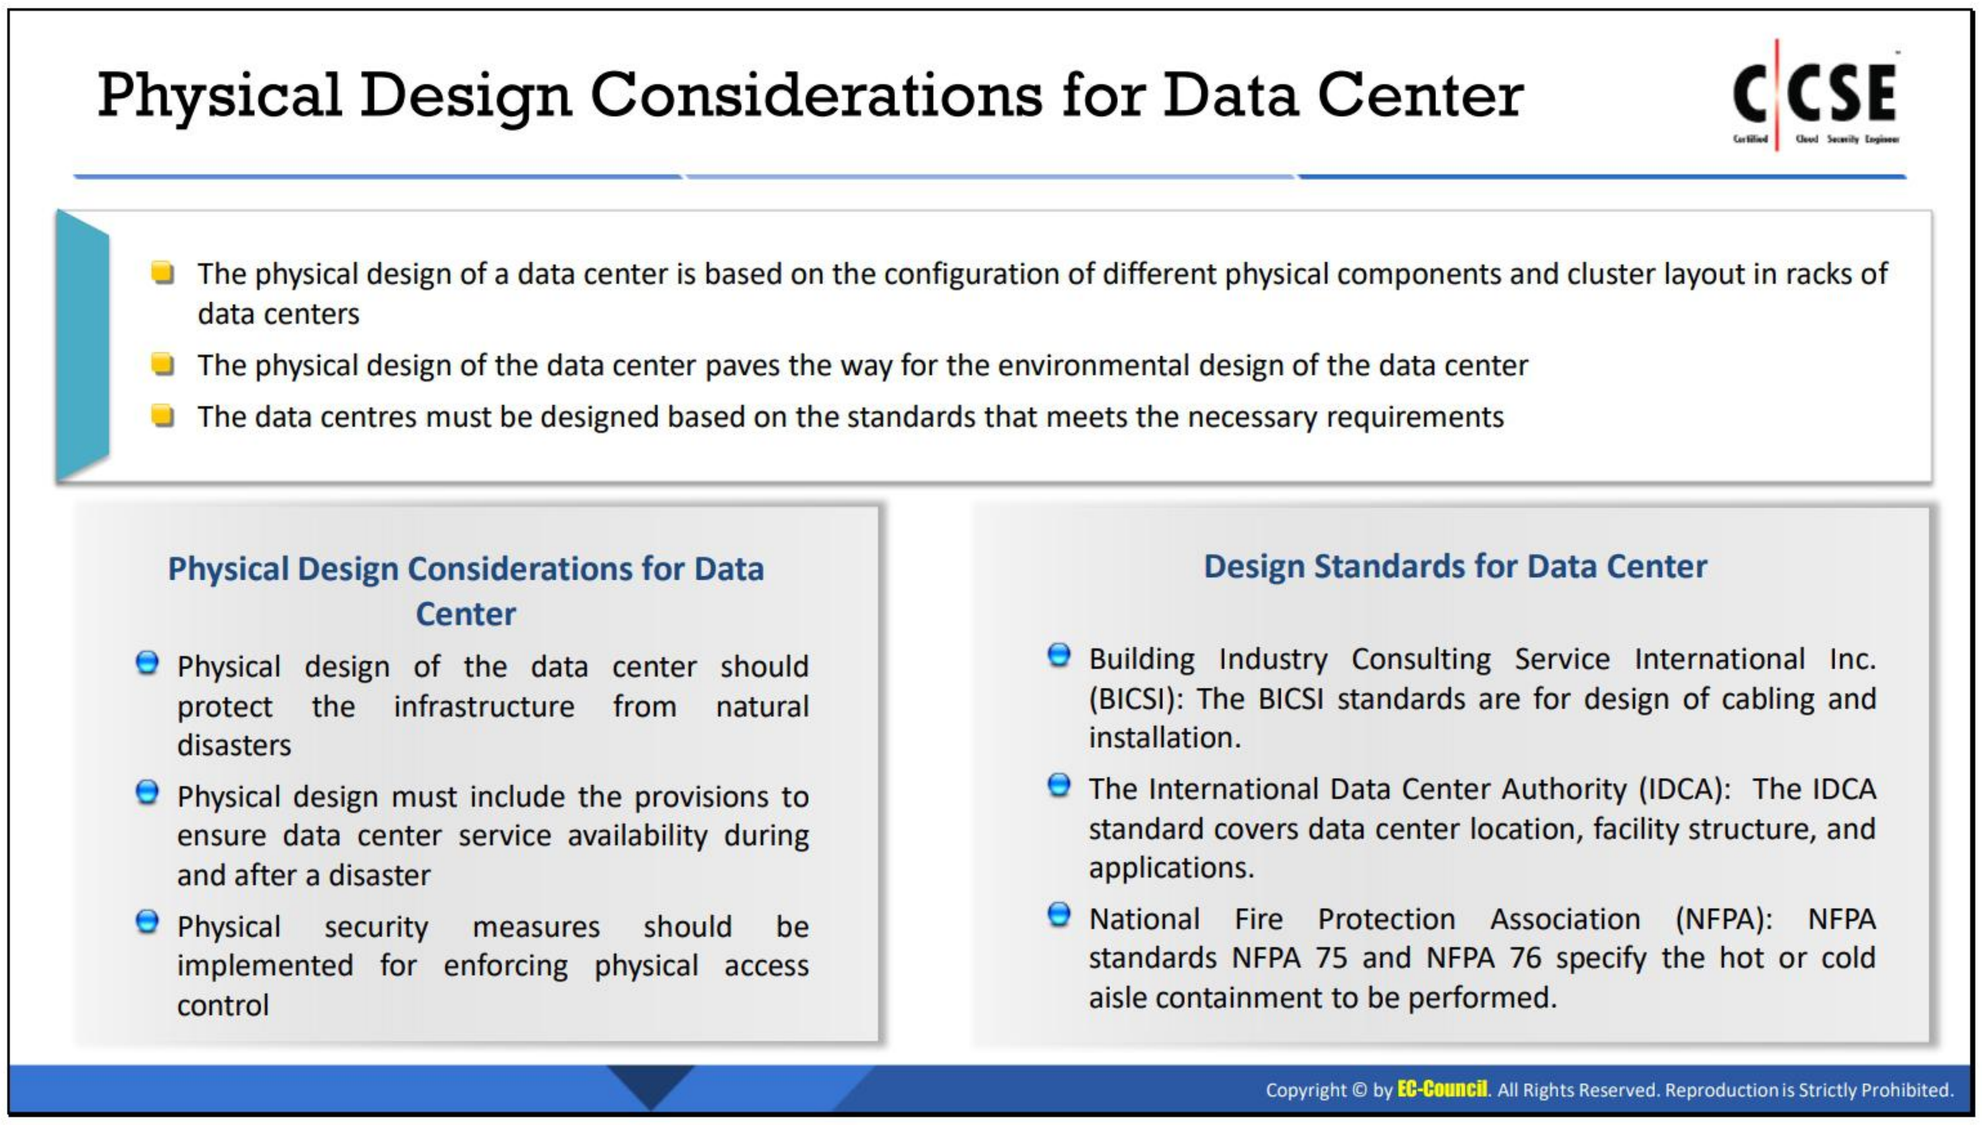


Figure: 2.10: Architecture of Cloud Data Center

Cloud data center infrastructure consists of servers, SAN, and VM. Cloud systems ensure continuous availability by setting up a primary data center along with a secondary data center, which are called as availability zones (AZs). Both the availability zones consist of VMs, management platform, identity authentication system, and back up & recovery systems. The physical and virtual servers are connected to SAN storage. The primary data center is replicated to secondary data center for performing fail over operations to ensure service availability for users. A secure network connection can be established between data centers through the use of a firewall.

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



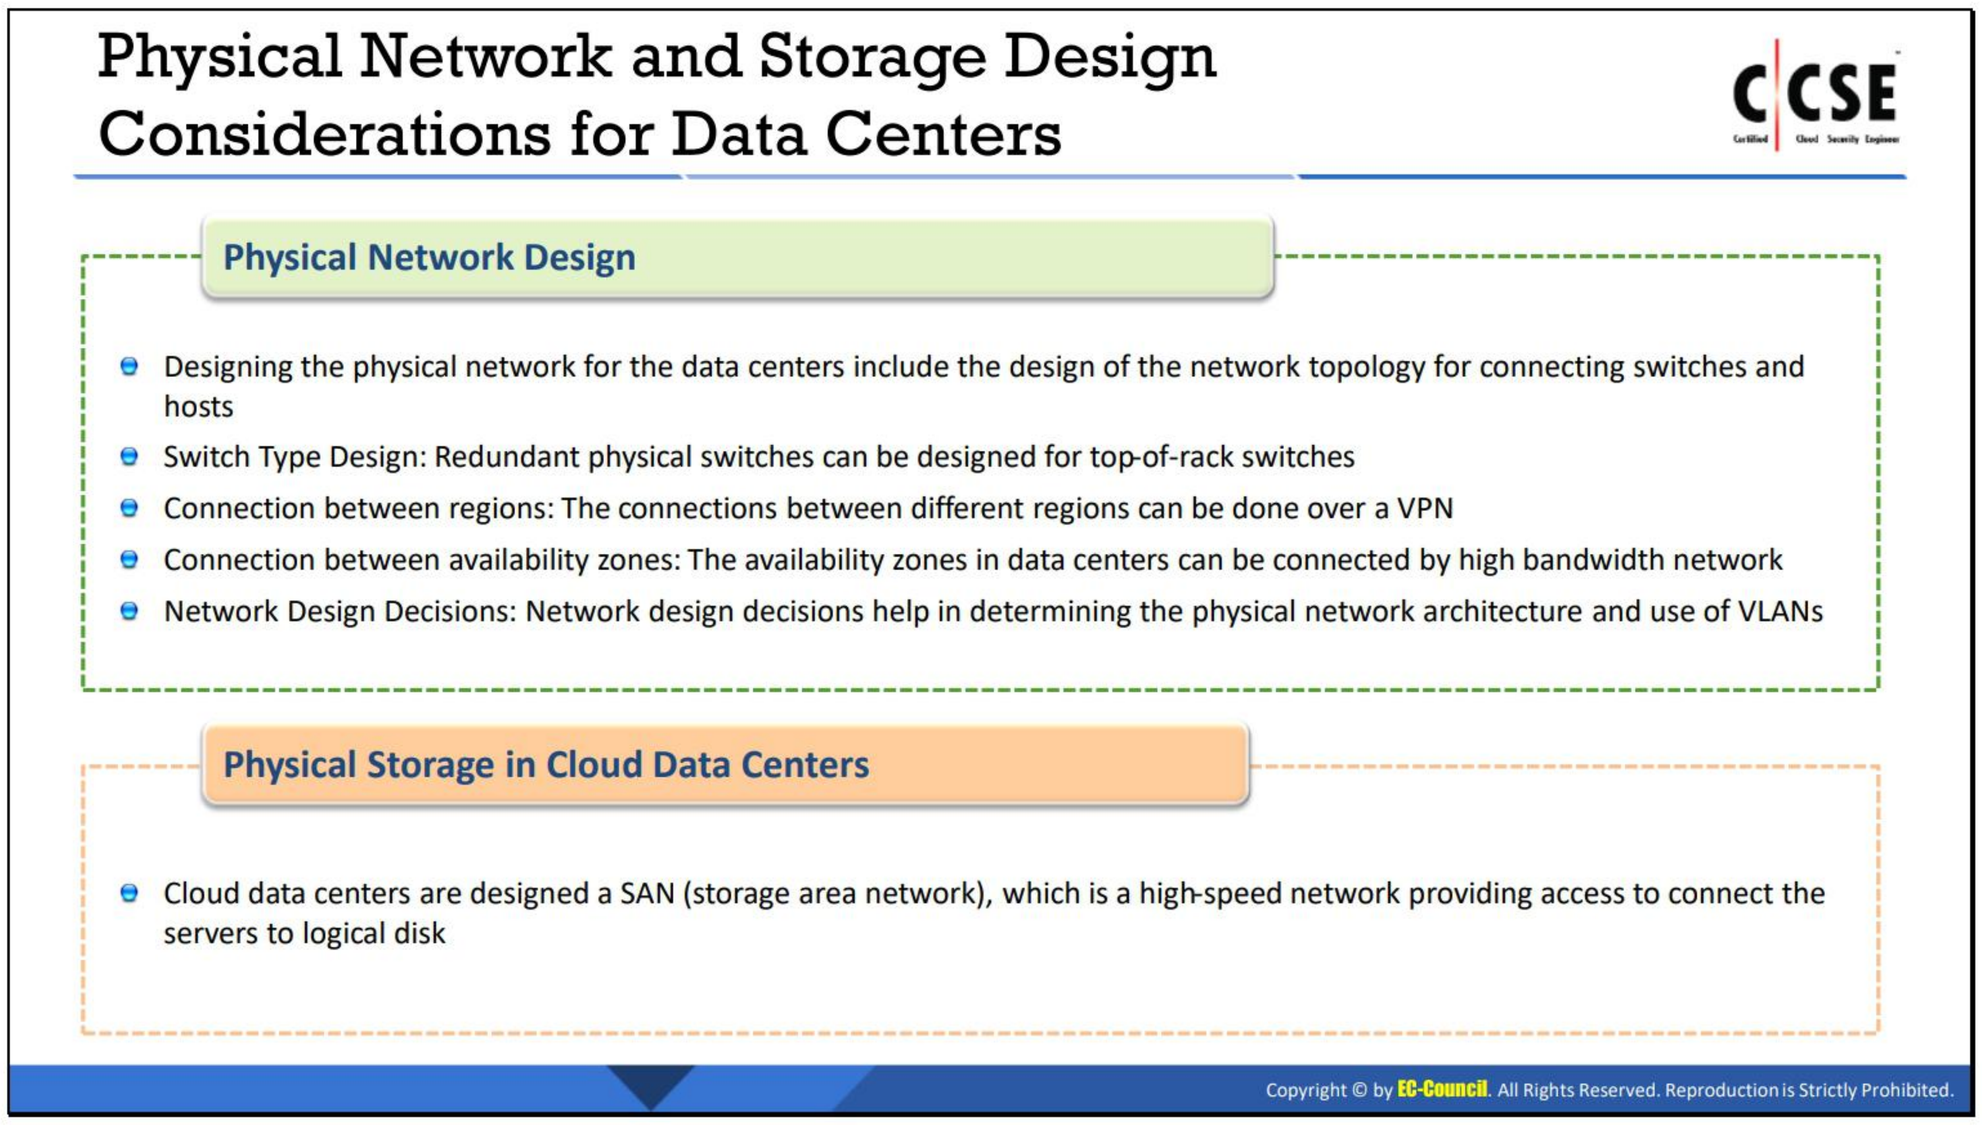
Physical Design Considerations for Data Center

The physical design of a data center is based on the configuration of different physical components and also the cluster layout of the server racks. Broadly, it includes servers, networking, and hardware storage. The physical design of the data center paves the way for the environmental design of the data center. The following factors must be considered for the data center physical design:

- The physical design of the data center should protect the infrastructure from natural disasters.
- It must include provisions to ensure data center service availability during and after a disaster.
- Physical security measures should be implemented for access control such as walls, fences, electronic surveillance, and verification of identity at the entry and exit points.

The physical design of the data center must also follow certain procedures such as tier certification, level of physical security, and the type of usage (whether it is multitenant or dedicated).

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



Physical Network and Storage Design Considerations for Data Centers

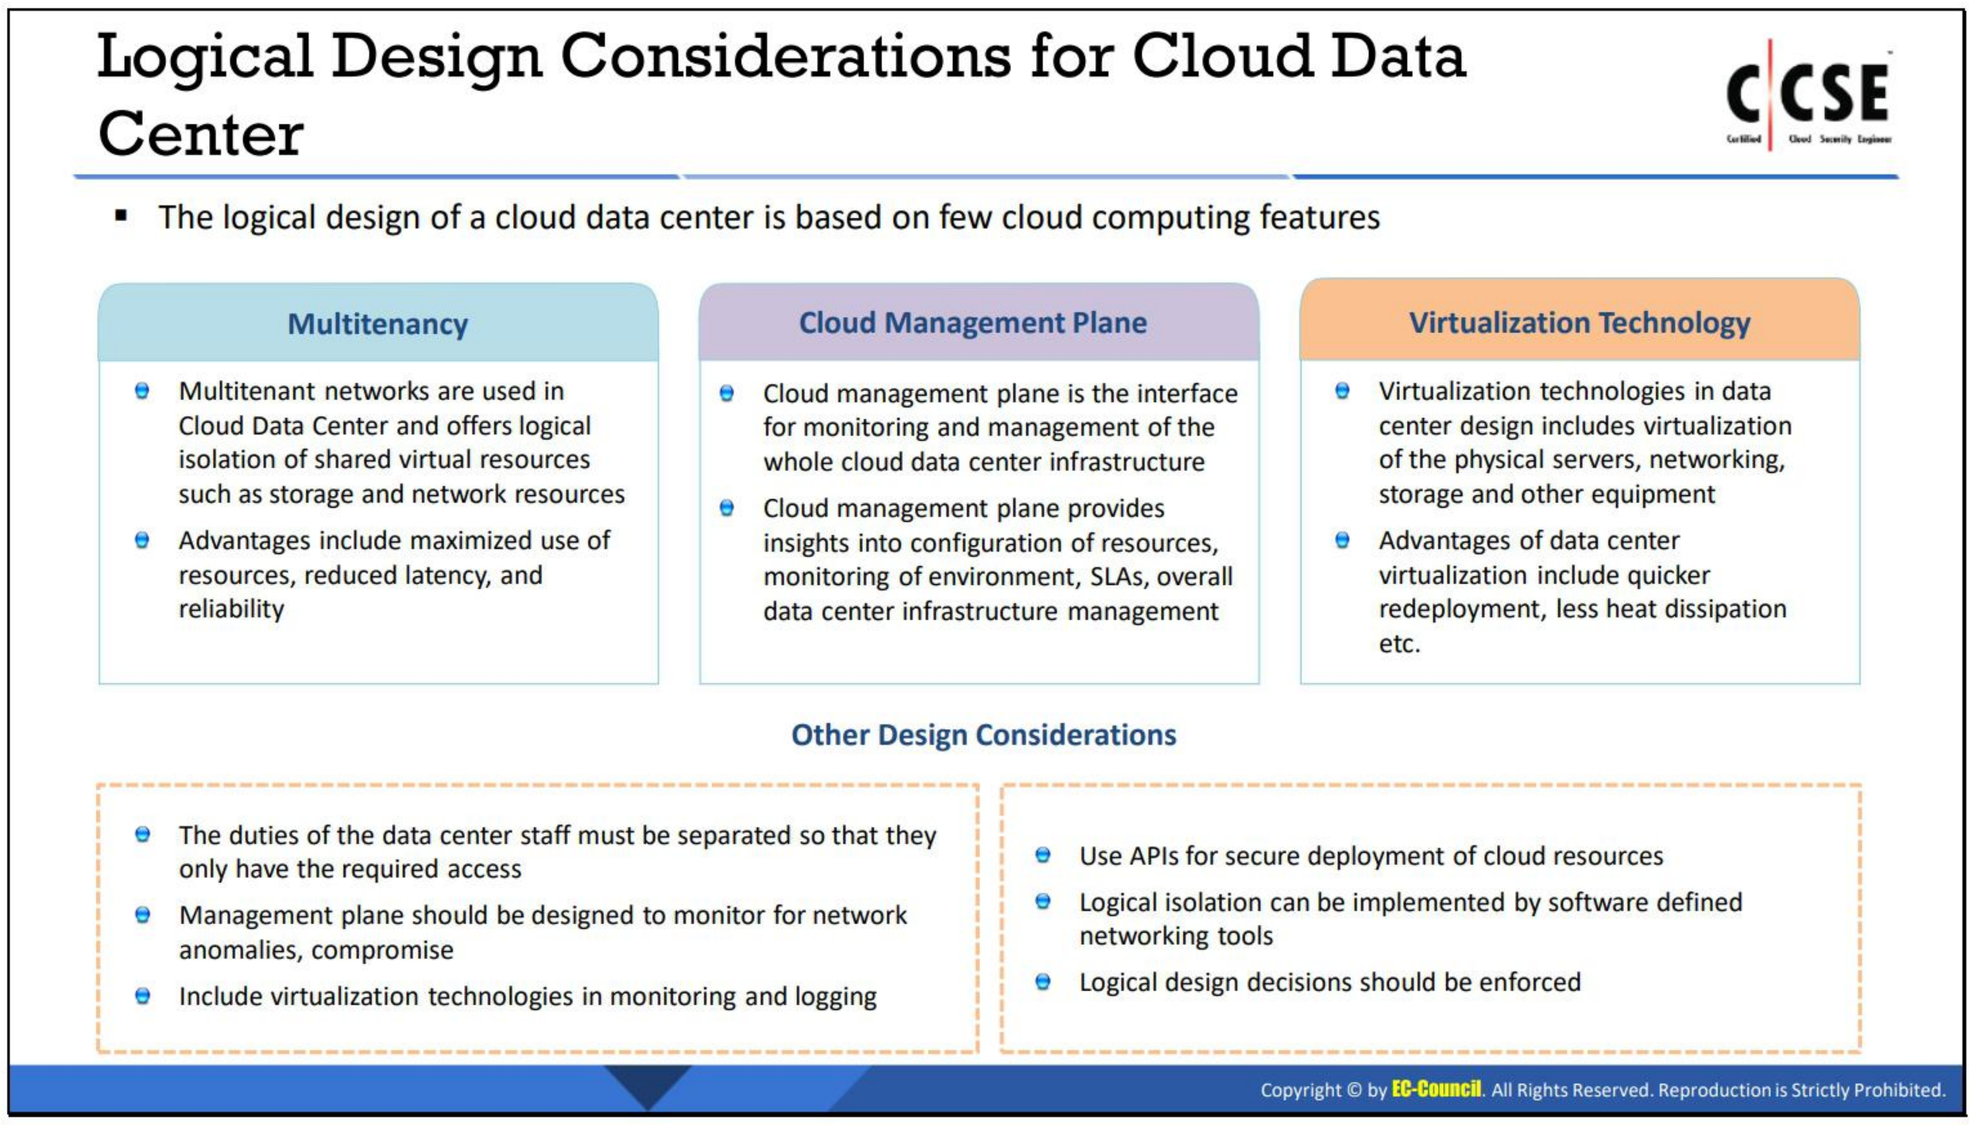
Designing the physical network for the data centers includes the design of the network topology for connecting the switches, hosts, routing designs, VLAN port settings, etc.

- **Switch type design**: While designing top-of-rack switches, redundant physical switches should be designed to ensure availability.
- **Connection between regions:** The connections between different regions should be made over a VPN and the latency should be minimum.
- **Connection between availability zones:** The availability zones in data centers can be connected by high-bandwidth network to ensure low latency.
- **Network design decisions:** The physical network architecture and the use of VLANs are determined by physical network design decisions. The different network designs are
  - Routing protocols: The routing protocols in data centers are selected based on the performance requirements. The routing protocols options are open shortest path first (OSPF), border gateway protocol (BGP), intermediate system to intermediate system (IS-IS).
  - DHCP proxy: The DHCP proxy can be set to point to a DHCP server by IPv4 address.

Physical Storage in Cloud Data Centers

Cloud data centers are designed as a SAN, which is a high-speed network providing access to connect the servers to logical disks. SANs are considered for data centers as they remove single points of failures and thus make storage highly available.

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



Logical Design Considerations for Cloud Data Center

The logical design of a cloud data center is based on the cloud computing features.

- **Multitenancy**: Cloud data centers must provide security for the sensitive data and applications in the virtual server environment. Multitenant networks that are divided logically into smaller networks can be implemented here. Such networks offer logical isolation of shared virtual resources such as storage and network resources. The physical network that the resources operate will be the same; however, they are logically isolated. The multitenant design of the data center requires logical division that separates the sensitive data. VLANs are deployed to separate the machines logically. Multitenant data centers should be scalable to accommodate more resources into the infrastructure. The advantages of implementing a multitenant network for the data centers are
  - Maximized use of data center resources can be achieved by sharing of the resources with a logical isolation.
  - Multitenant design is cost effective, where multiple users can share the resources.
  - Multitenant design of data centers helps in reducing latency.
  - Multitenant design is reliable as cloud data centers have their own security and operation procedures to ensure continuous availability and to operate at maximum efficiency.
- **Cloud management plane:** Cloud management plane is the interface for monitoring and management of the whole cloud data center infrastructure, which can be accessed via APIs and web consoles. This plane is to be logically isolated. It provides insights into the management of the following

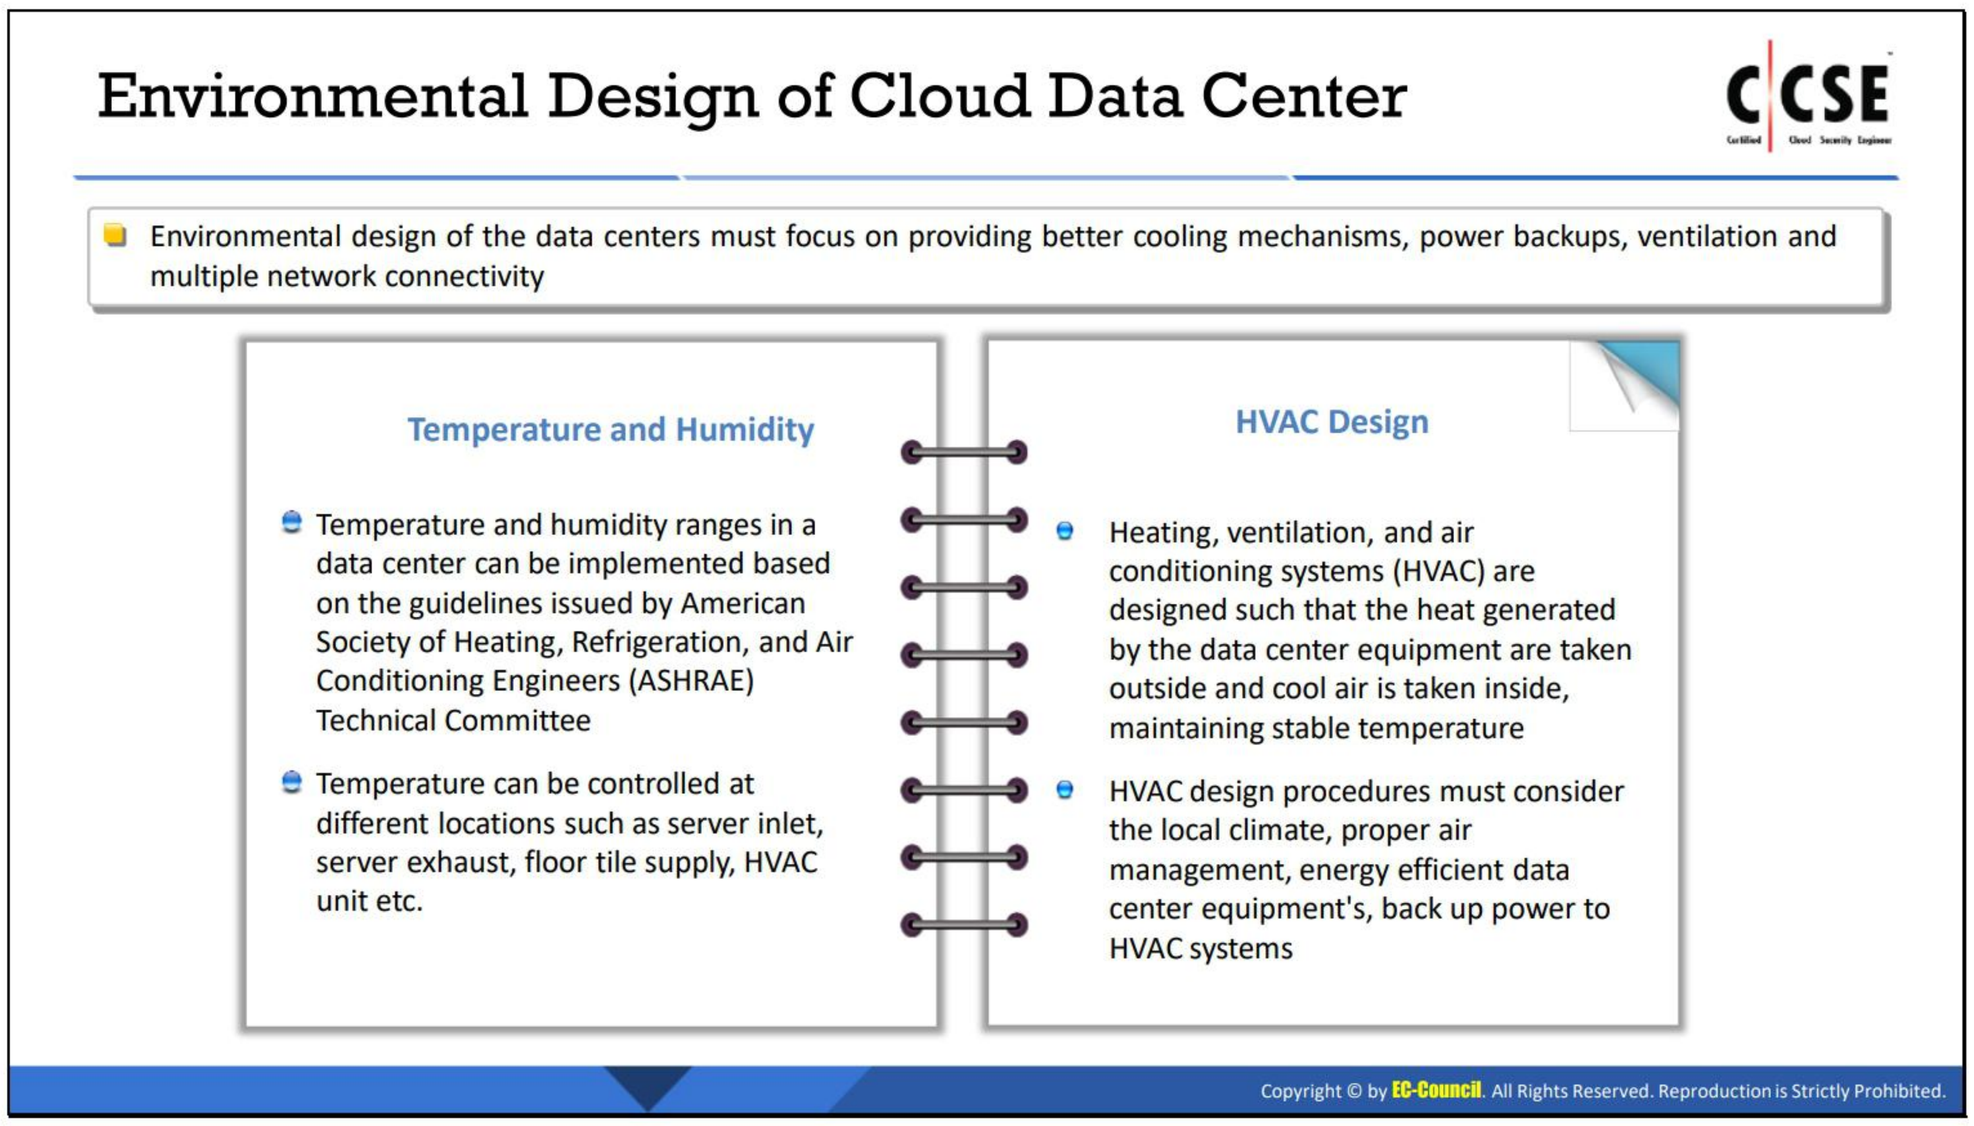
Certified Cloud Security Engineer Copyright © by **EC-Council**
All Rights Reserved. Reproduction is Strictly Prohibited.

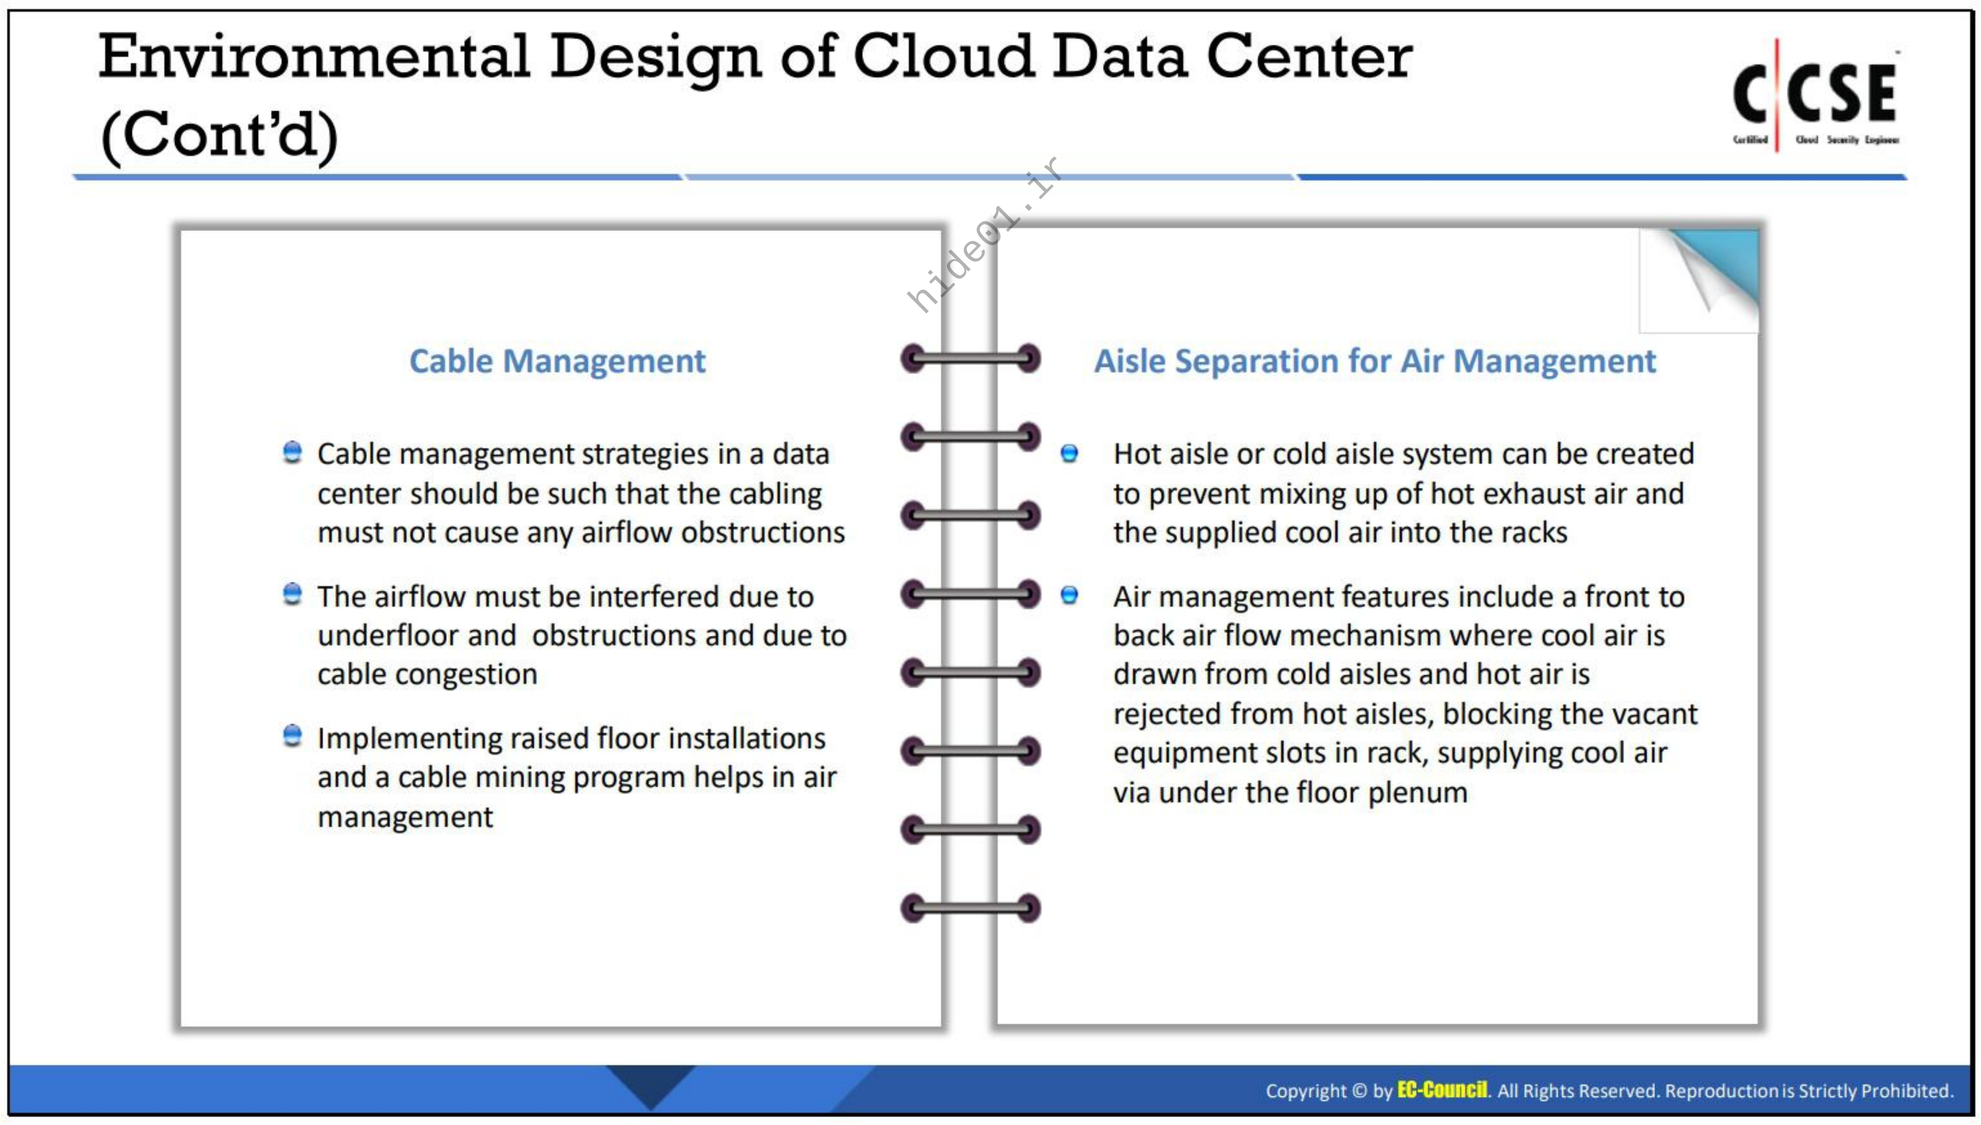
- Configuration management of the resources.
- Monitoring, logging, and auditing of the environment.
- Service level agreement (SLA) and service lifecycle management.
- Management of the overall security of the data center infrastructure.

- **Virtualization technology for logical design:** Virtualization technologies in data center design include virtualization of the physical servers, networking, storage, and other equipment. Virtualization technologies can be used in data center design, which offers many services required for partitioning. The logical design of the data center must include
  - Access for communication based on requirements.
  - Application programming interface (API) access.
  - The communication within the management plane should be secure and encrypted.
  - Storage should be secure by implementing encryption, partition, and secure key management.
  - The logical design must include disaster recovery and business continuity plan via implementation of failover policies, replication, and backup.
- **Logical design best practices:**
  - The logical design must be such that the duties of the data center staff must be separated and assigned with specific roles and responsibilities so that they only have the required access.
  - Management plane should be designed to monitor for network anomalies and system compromise.
  - Hypervisor and virtualization technologies must also be included with monitoring of the data center resources as certain virtualization technologies do not provide functionalities for monitoring of the environment.
  - Logical design must include the use of APIs for a secure cloud deployment, and it must also include loggings for API use.
  - Logical design decisions such as access controls should be such that they can be enforced and logged.
  - Logical isolation can be implemented via software-defined networking tools.

hide01.ir

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





Environmental Design of Cloud Data Center

The environmental design of the data centers must focus on providing better cooling mechanisms, power backups, ventilation, and multiple network connectivity. The different guidelines that can be followed for environmental design of the data centers are discussed below.

- **Temperature and humidity:** The temperature and humidity ranges in a data center can be implemented based on the guidelines issued by ASHRAE Technical Committee. The

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

recommended operating ranges of temperature and humidity for a cloud environment based on the equipment intake temperature are

- Low-end temperature: 64.4° F (18° C)
- High-end temperature: 80.6° F (27° C)
- Low-end moisture: 40% relative humidity and 41.9° F (5.5° C) dew point
- High-end moisture: 60% relative humidity and 59° F (15° C) dew point

The temperature can be controlled in different locations such as server inlet, server exhaust, floor tile supply, HVAC unit, computer room etc.

- **Heating, ventilation, and air conditioning (HVAC) design:** HVAC systems are designed such that the heat generated by the data center equipment is taken outside and cool air to supply the equipment is pulled inside to main maintain a stable temperature. Power requirements for data center cooling is based on the amount of heat removed and temperature difference between inside the data center and outside environment. The best practices for HVAC design in data centers are
  - HVAC design can be based on the local climate of the data center location.
  - HVAC design must provide air management facilities where the cool air supplied to the resources does not get mixed up with the hot air exhaust.
  - Implement energy efficient systems and provide backup power to the systems to ensure continuous cooling.
  - HVAC system should also filter dust.
- **Cable management:** Cable management strategies in a data center should be such that the cabling must not cause any airflow obstructions. Cable management strategies should be implemented to airflow path within the data centers such as rack air intake and discharge areas and also underfloor areas. The air flow may be obstructed and hot spot may rise due to two reasons.
  - Interference in the distribution of cool air and the air flow due to underfloor and overhead obstructions.
  - Interference in air flow due to cable congestion in raised floor plenums.

  For raised floor installations, a minimum height of 24 inches should be implemented. A cable mining program that removes inoperable cables can be implemented for air management to data center cooling systems.
- **Aisle separation for air management:** Hot aisle or cold aisle system can be created to prevent mixing up of hot exhaust air and the supplied cool air into the racks. The data centers are designed for air management such that the equipment are arranged in rows of racks where alternating cold air and hot air aisles between them. A front to back air flow mechanism is implemented in the rack where cool air is drawn in from cold aisles in the front portion and heat gets rejected from the hot aisles in the back portion of the rack. Additionally, the vacant equipment slots in the rack are blocked for air intake to

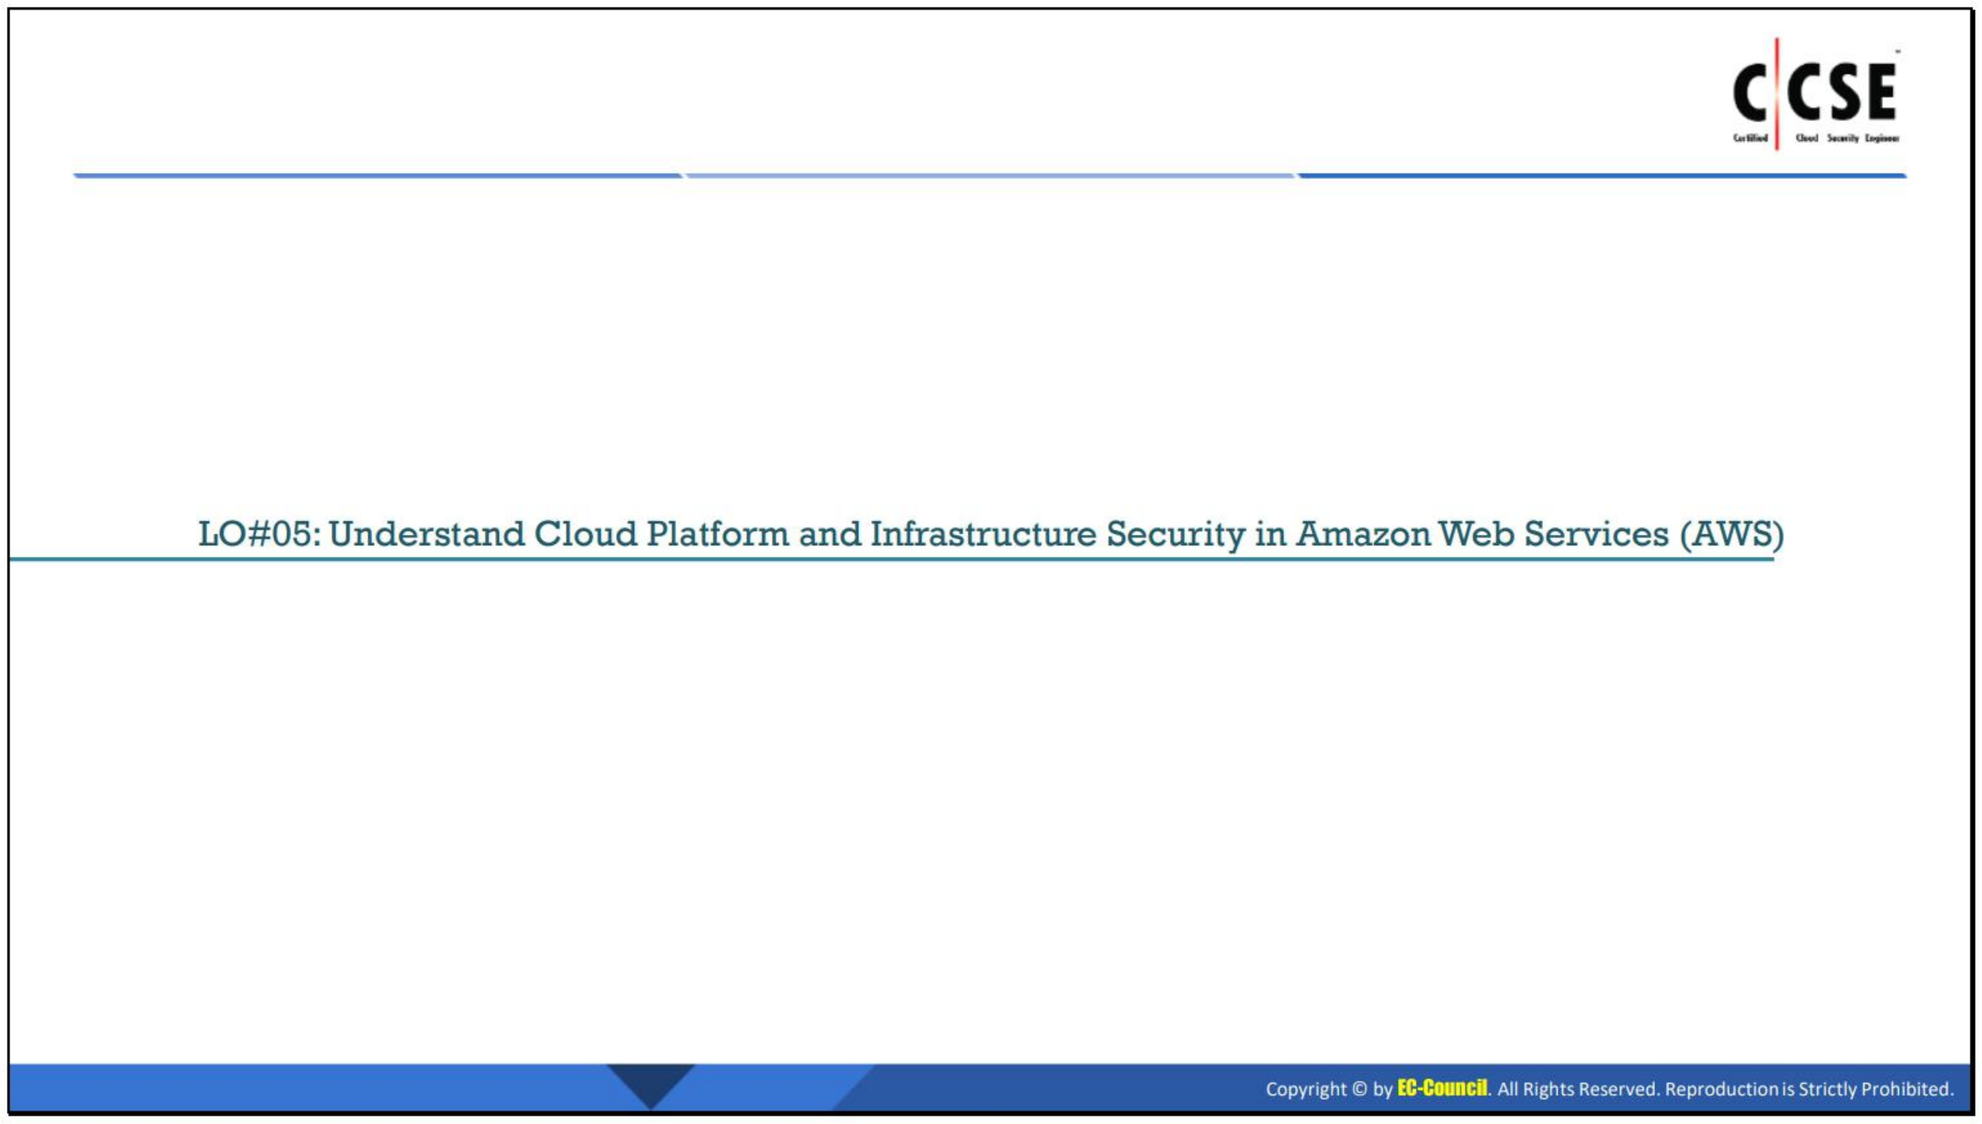
Certified Cloud Security Engineer Copyright © by **EC-Council**
All Rights Reserved. Reproduction is Strictly Prohibited.

prevent recirculation. Thus, if the isolation is proper, the hot aisle becomes a heat exhaust. The cooling system can be designed to supply cold air to cold aisles. This can ensure a better and reliable cooling management of the data center infrastructure.

Another air flow design can be achieved by supplying cool air via an underfloor plenum. The cool air passes through the rack and then passes through a separate area to an overhead plenum. A barrier at the top of the rack and at the hot aisle end prevents mixing of the cold and hot air.

hide01.ir

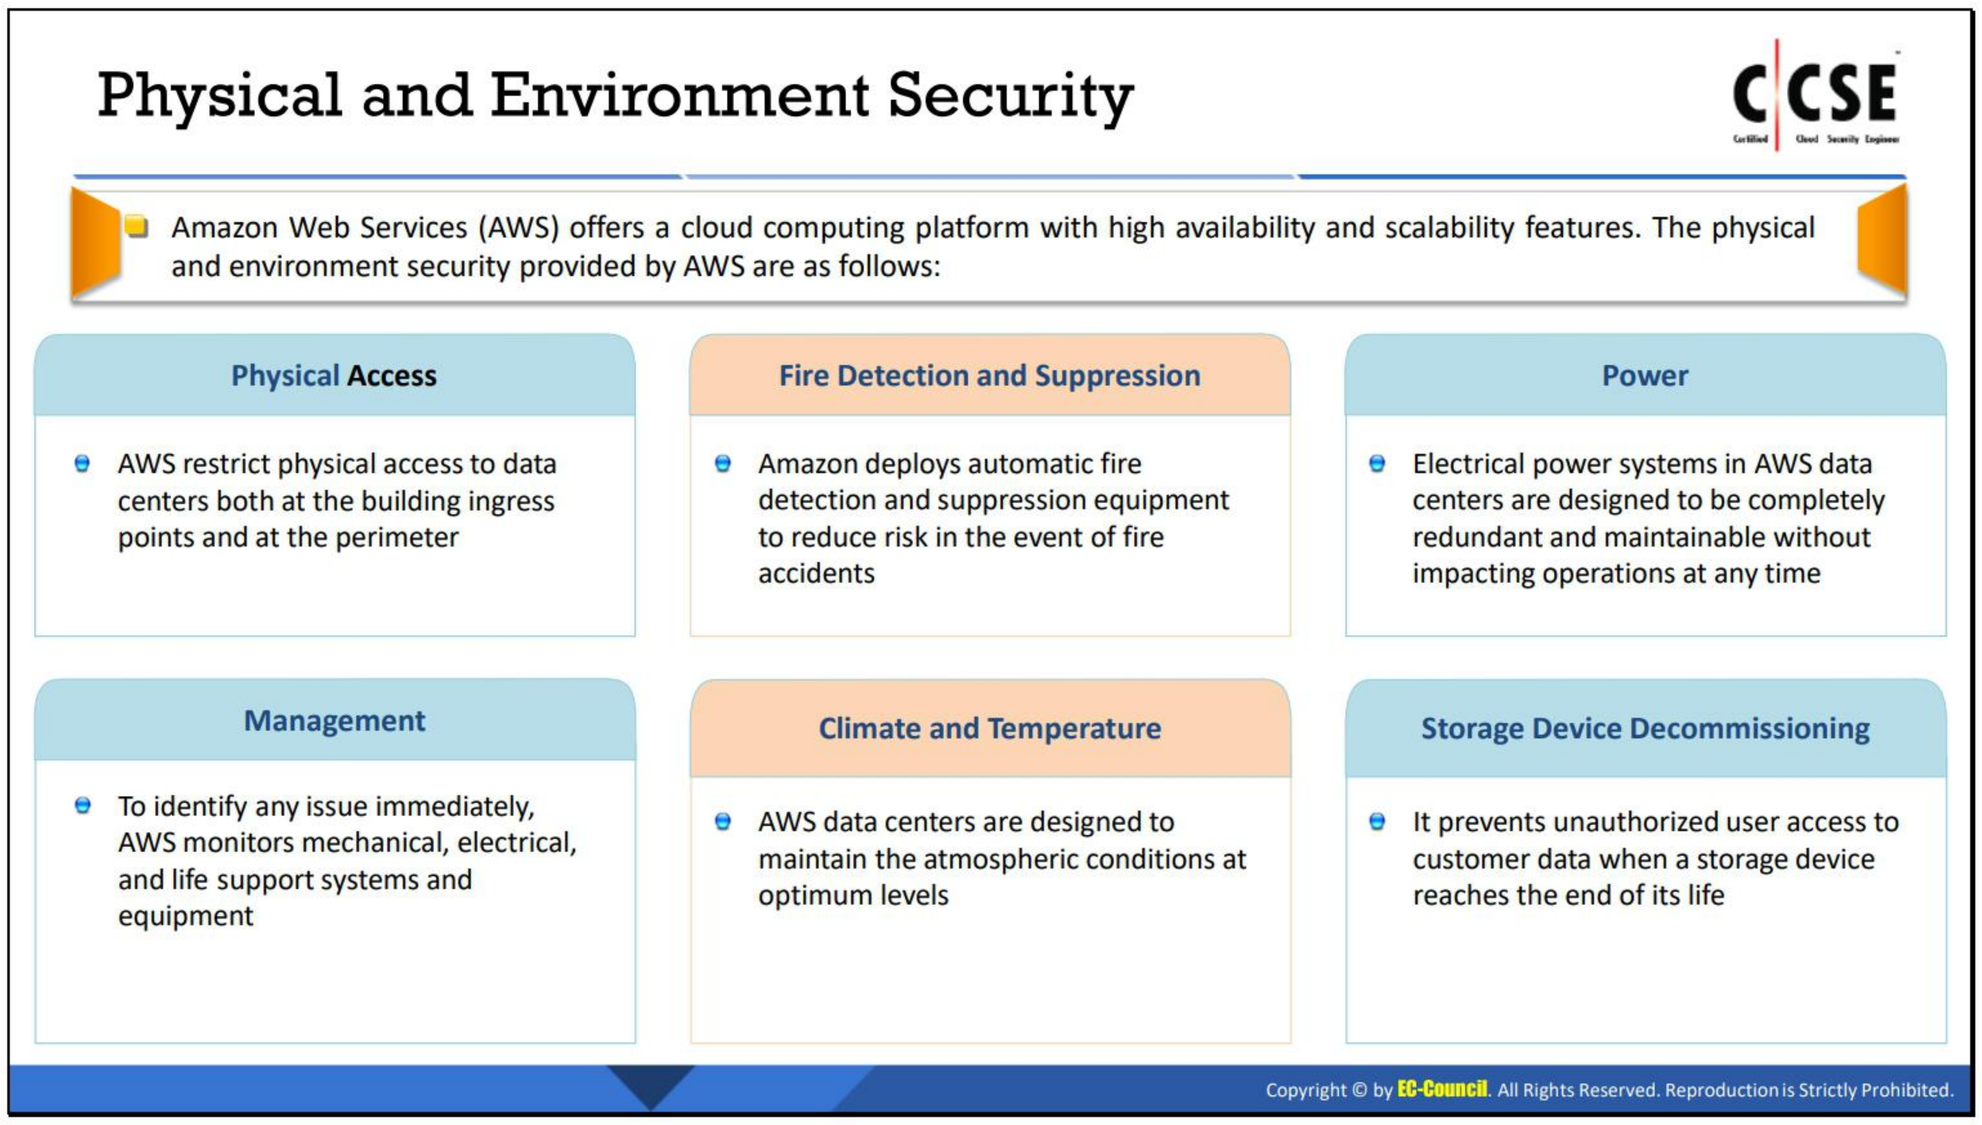
Certified Cloud Security Engineer Copyright © by **EC-Council**
All Rights Reserved. Reproduction is Strictly Prohibited.



LO#05: Understand Cloud Platform and Infrastructure Security in Amazon Web Services (AWS)

Knowledge of the network architecture provided by AWS will help the cloud security engineer build a secure global infrastructure to secure their compute, network, and storage. This section aims to explain the security features and components of AWS cloud resources used to ensure physical and environmental security using network service security.

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



Physical and Environment Security

Cloud Platform and Infrastructure Security in AWS: AWS offers a highly available and scalable cloud computing platform. AWS enables the customer to run applications in a wide range through the provided tools. AWS cloud environment protects the integrity, confidentiality, and availability of its customers' data and systems. AWS provides the below security to cloud platform and infrastructure.

Physical and Environment Security: AWS data centers take advantage of latest engineering and architectural approaches. Amazon used its several years of experience for designing, building, and managing data centers, and this experience is applied to the AWS infrastructure and platform to provide secure cloud services.

Amazon ensures the security of its data centers by implementing strict security measures. Physical access to its data centers is restricted at both the building ingress points and the perimeter by professional security personnel using intrusion detection systems, video surveillance, and other electronic means.

All staff members who have authorized access to AWS data centers must pass two-factor authentication at least two times. Any visitor who wants to go inside a data center must show their identification and then they are allowed to sign-in, and an authorized staff member escorts them all the time when they are inside. AWS allows contractors and employees having legitimate business needs to access its data center and information as per their privileges. It revokes the data center access to the employees who do not have a business need even if they are Amazon employees. Moreover, all physical access to AWS data centers logged with accurate time and audited regularly.

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

Fire Detection and Suppression: Amazon deploys automatic fire detection and suppression equipment to reduce risk in the event of fire accidents. This fire detection system uses smoke detection sensors, chiller rooms, mechanical and electrical infrastructure spaces, and generator equipment rooms. It also uses double-interlocked pre-action, wet-pipe, or gaseous sprinkler systems for protection.

Power: Electrical power systems in AWS data centers are designed to be completely redundant and maintainable without impacting operations anytime. When the electric supply fails for the essential loads, the UPS units provide the backup power supply. In such a case, generators are used to provide backup electric supply to the data center facility in Amazon.

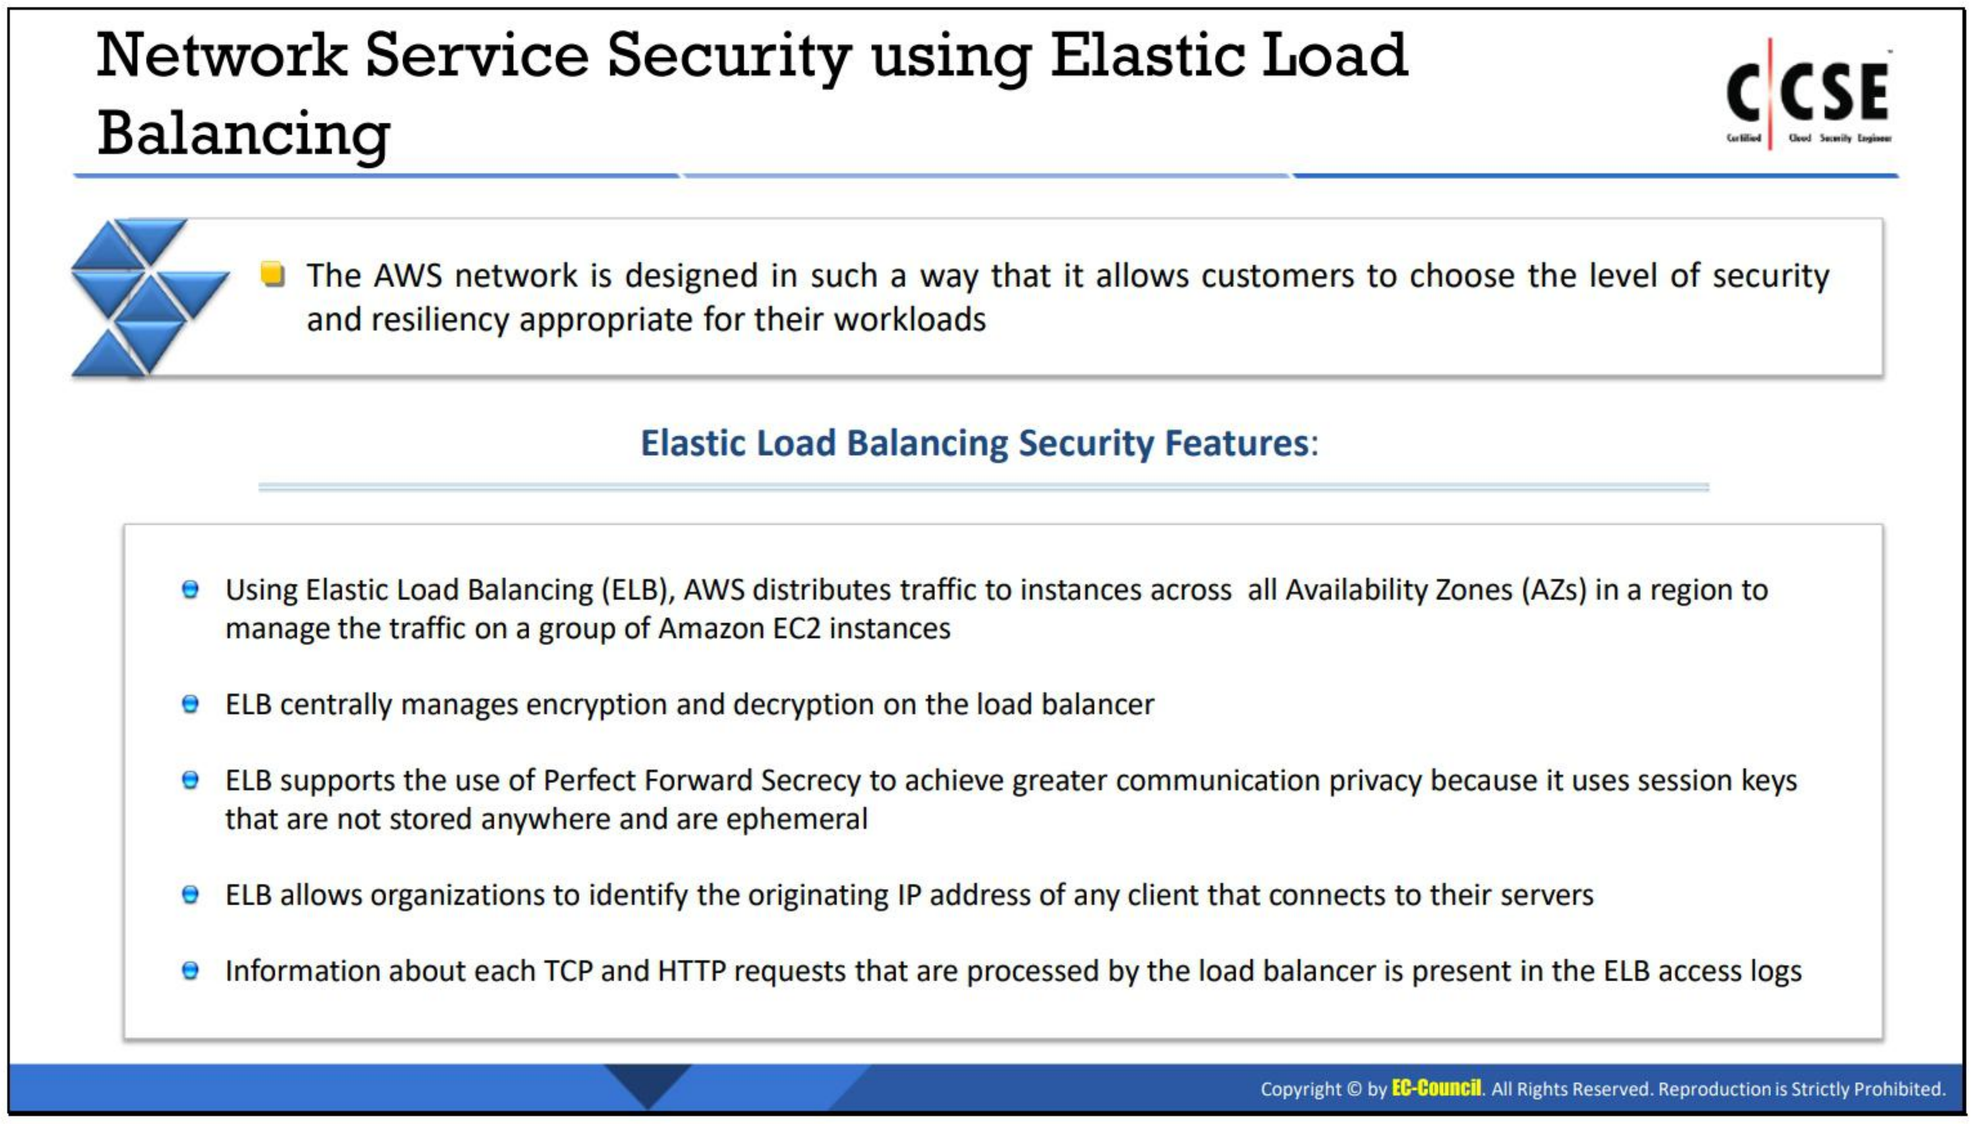
Management: To identify any issues immediately, AWS monitors mechanical, electrical, and life support systems and equipment. AWS performs preventive maintenance to ensure continued operability of equipment.

Climate and Temperature: To maintain a constant operating temperature for hardware components and servers, climate control is essential because it avoids overheating and reduces service outages chances. Data centers are designed in such a way that they maintain atmospheric conditions at optimum levels. The temperature and humidity inside data centers are monitored and controlled by personnel and systems to ensure they are at optimum levels.

Storage Device Decommissioning: AWS procedures contain a decommissioning process designed to prevent the exposure of customer data to unauthorized individuals when a storage device is about to expire. AWS implements the techniques that are detailed in NIST 800-88 with the name “Guidelines for Media Sanitization” as a part of its decommissioning process.

hidden

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



Network Service Security using Elastic Load Balancing

The AWS network is designed in such a way that it allows customers to choose the level of security and resiliency appropriate for their workloads. It has a world-class network infrastructure that is managed and monitored carefully and allows its customers to build fault-tolerant, geographically dispersed web architectures with cloud resources.

Elastic Load Balancing Security

AWS uses Elastic Load Balancing (ELB) for traffic management on Amazon EC2 instances caused by traffic disturbance instances across all regions of AZs. It combines all the advantages of an on-premises load balancer and the benefits in security include:

- ELB takes encryption and decryption work from the EC2 instance and hands it over to the load balancer for the central management.
- It serves as the first line of defense against potential attacks on the network.
- When ELB is used in an Amazon Virtual Private Cloud (VPC), it supports the creation and management of security groups that are associated with the ELB to offer additional networking and security options.
- ELB helps in end-to-end traffic encryption utilizing TLS on those networks that use HTTPS connections.

To create an encryption message between the browser and the server HTTP/TLS deploy a long-term secret key to generate a short-term secret key. When a connection is established between the ELB and a client, ELB configures the load balancer with a predefined cipher set which is used for TLS negotiation. This predefined cipher set uses strong algorithms and provides compatibility with a wide range of clients. To ensure standards are met, some customers may have certain

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

requirements to allow only specific protocols and ciphers such as SOX, PCI, etc. from clients. In such cases, ELB offers options for choosing various configurations for TLS protocols and ciphers. As per specific requirements, customers can enable or disable the ciphers.

Customers can configure the load balancer to ensure that the latest and strong cipher suites are used for establishing a secure connection by selecting a suitable cipher suite during client-server negotiation. The load balancer chooses the appropriate cipher suite based on the server's priority when the Server Order Preference option is selected. This provides control over the security level to organizations and helps their clients to connect to the load balancer.

ELB supports the use of Perfect Forward Secrecy to achieve greater communication privacy because it uses an ephemeral session key that can be stored secretly. Despite the long-term key being compromised, the use of Perfect Forward Secrecy prevents the decoding of captured data.

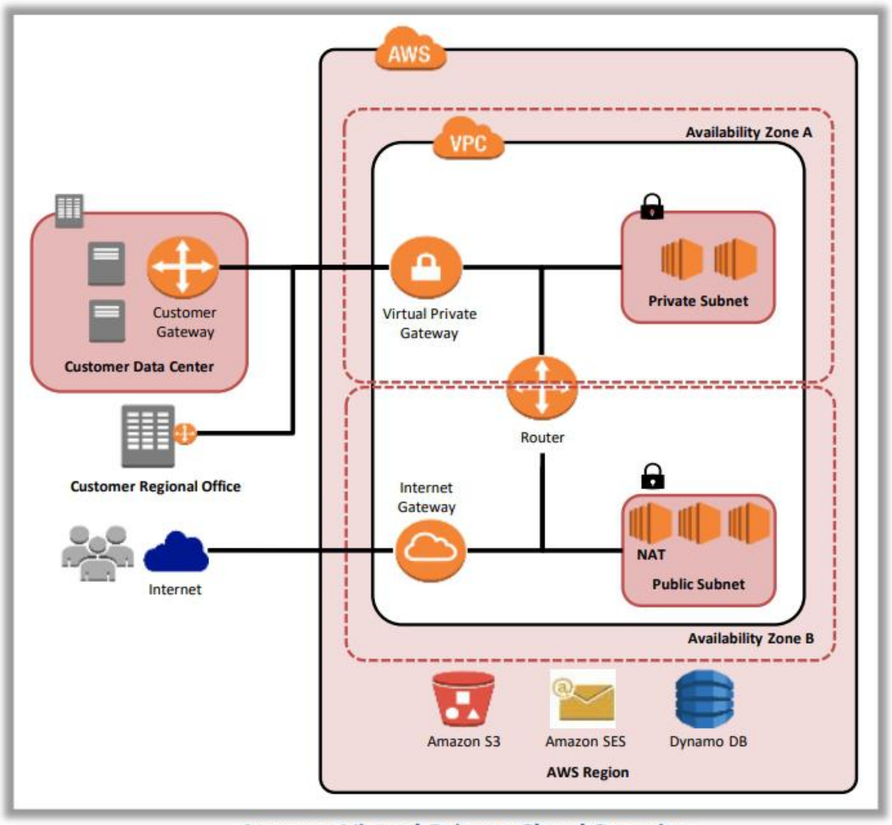
ELB allows organizations to identify the originating IP address of any client that connects to their servers, no matter they use TCP or HTTP load balancing. Generally, when requests are proxied via a load balancer, client connection information (like port and IP address) is lost because the requests are sent by the load balancer to the server on behalf of the client, and displays the load balancer as the requesting client. The client IP address is very useful and helpful to the organizations for the purpose of achieving more information about the client such as gathering connection statics, manage whitelists of IP addresses, or analyze traffic logs.

Information about each TCP and HTTP requests that are processed by the load balancer is present in the ELB access logs. This information includes the size of the request and response, the IP address and port information of the requesting client, the backend IP address of the instance that processed the request, and the actual request line from the client. For example, GET http://www.example.com: 80/HTTP/1.1.All requests sent to the load balancer are logged.

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

Network Service Security: Amazon Virtual Private Cloud (Amazon VPC) Security

- Amazon Virtual Private Cloud (VPC) create an isolated portion of the AWS cloud.
- It enables customers to launch Amazon EC2 instances that have private addresses in the range of their choice
- AWS provides various VPC architecture templates with configurations that offer different levels of public access
  - VPC that has only one public subnet
  - VPC that has both public and private subnets
  - VPC that has public and private subnets including hardware VPN access
  - VPN that has only private subnet along with hardware VPN access
- The security features in an Amazon VPC are security groups, routing tables, external gateways, and network ACLs
- Customers need to create VPC security groups specifically for their Amazon VPCs



Amazon Virtual Private Cloud Security

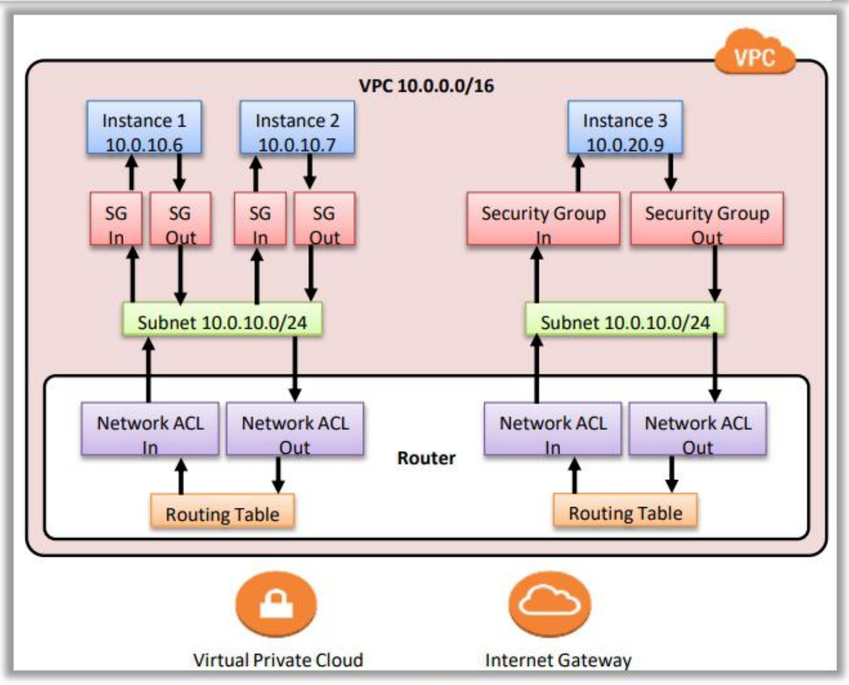
Copyright © by EC-Council. All Rights Reserved. Reproduction is Strictly Prohibited.

Network Service Security: Amazon Virtual Private Cloud (Amazon VPC) Security (Cont'd)

- Network traffic within every Amazon VPC is completely isolated from all other Amazon VPCs
- Enables customers to establish external connectivity by creating and attaching a virtual private gateway, an internet gateway, or both

Amazon VPC Security Controls

- API access
- Subnets and Route tables
- Firewall (Security Groups)
- Network Access Control Lists
- Virtual Private Gateway (VPG)
- Internet Gateway
- Dedicated Instances
- Elastic Network Interfaces



Implementation of AWS Security Controls

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

Network Service Security: Amazon Virtual Private Cloud (Amazon VPC) Security

A public IP address is randomly assigned in the Amazon EC2 instance while it is launched. Amazon VPC enables customers to create an isolated portion of the AWS cloud and launch Amazon EC2 instances that have private addresses in the range of their choice. Customers can define subnets within their VPC by grouping similar types of instances based on the range of IP addresses and set routing and security controls for in and out traffic flow of subnet and instance.

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

AWS provides various VPC architecture templates with configurations that offer different levels of public access.

- **VPC with Only a Single Public Subnet**: Instances run in an isolated, private section of the AWS cloud having direct access to the internet. To have strict control over inbound and outbound network traffic to the instances, customers can use security groups and network ACLs.
- **VPC with Public and Private Subnets:** This configuration contains a public subnet and a private subnet. Instances in the private subnet cannot be addressed from the internet but can use network address translation (NAT) to establish outbound connection to the internet through a public subnet.
- **VPC With Public and Private Subnets and Hardware VPN Access**: An IPsec VPN connection is added between an Amazon VPC and the data center in this configuration. this extends the cloud data centers and establishes direct access to the internet for an Amazon VPC in a public subnets instance. Customers can add a VPN appliance on their corporate data center side for the configuration.
- **VPC with Private Subnet Only and Hardware VPN Access:** Instances run in a private, isolated portion of the AWS cloud with a private subnet. Instances in the private subnet are not addressable from the internet. Customers can connect their corporate data center with private subnet through an IPsec tunnel.

hide01.ir

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

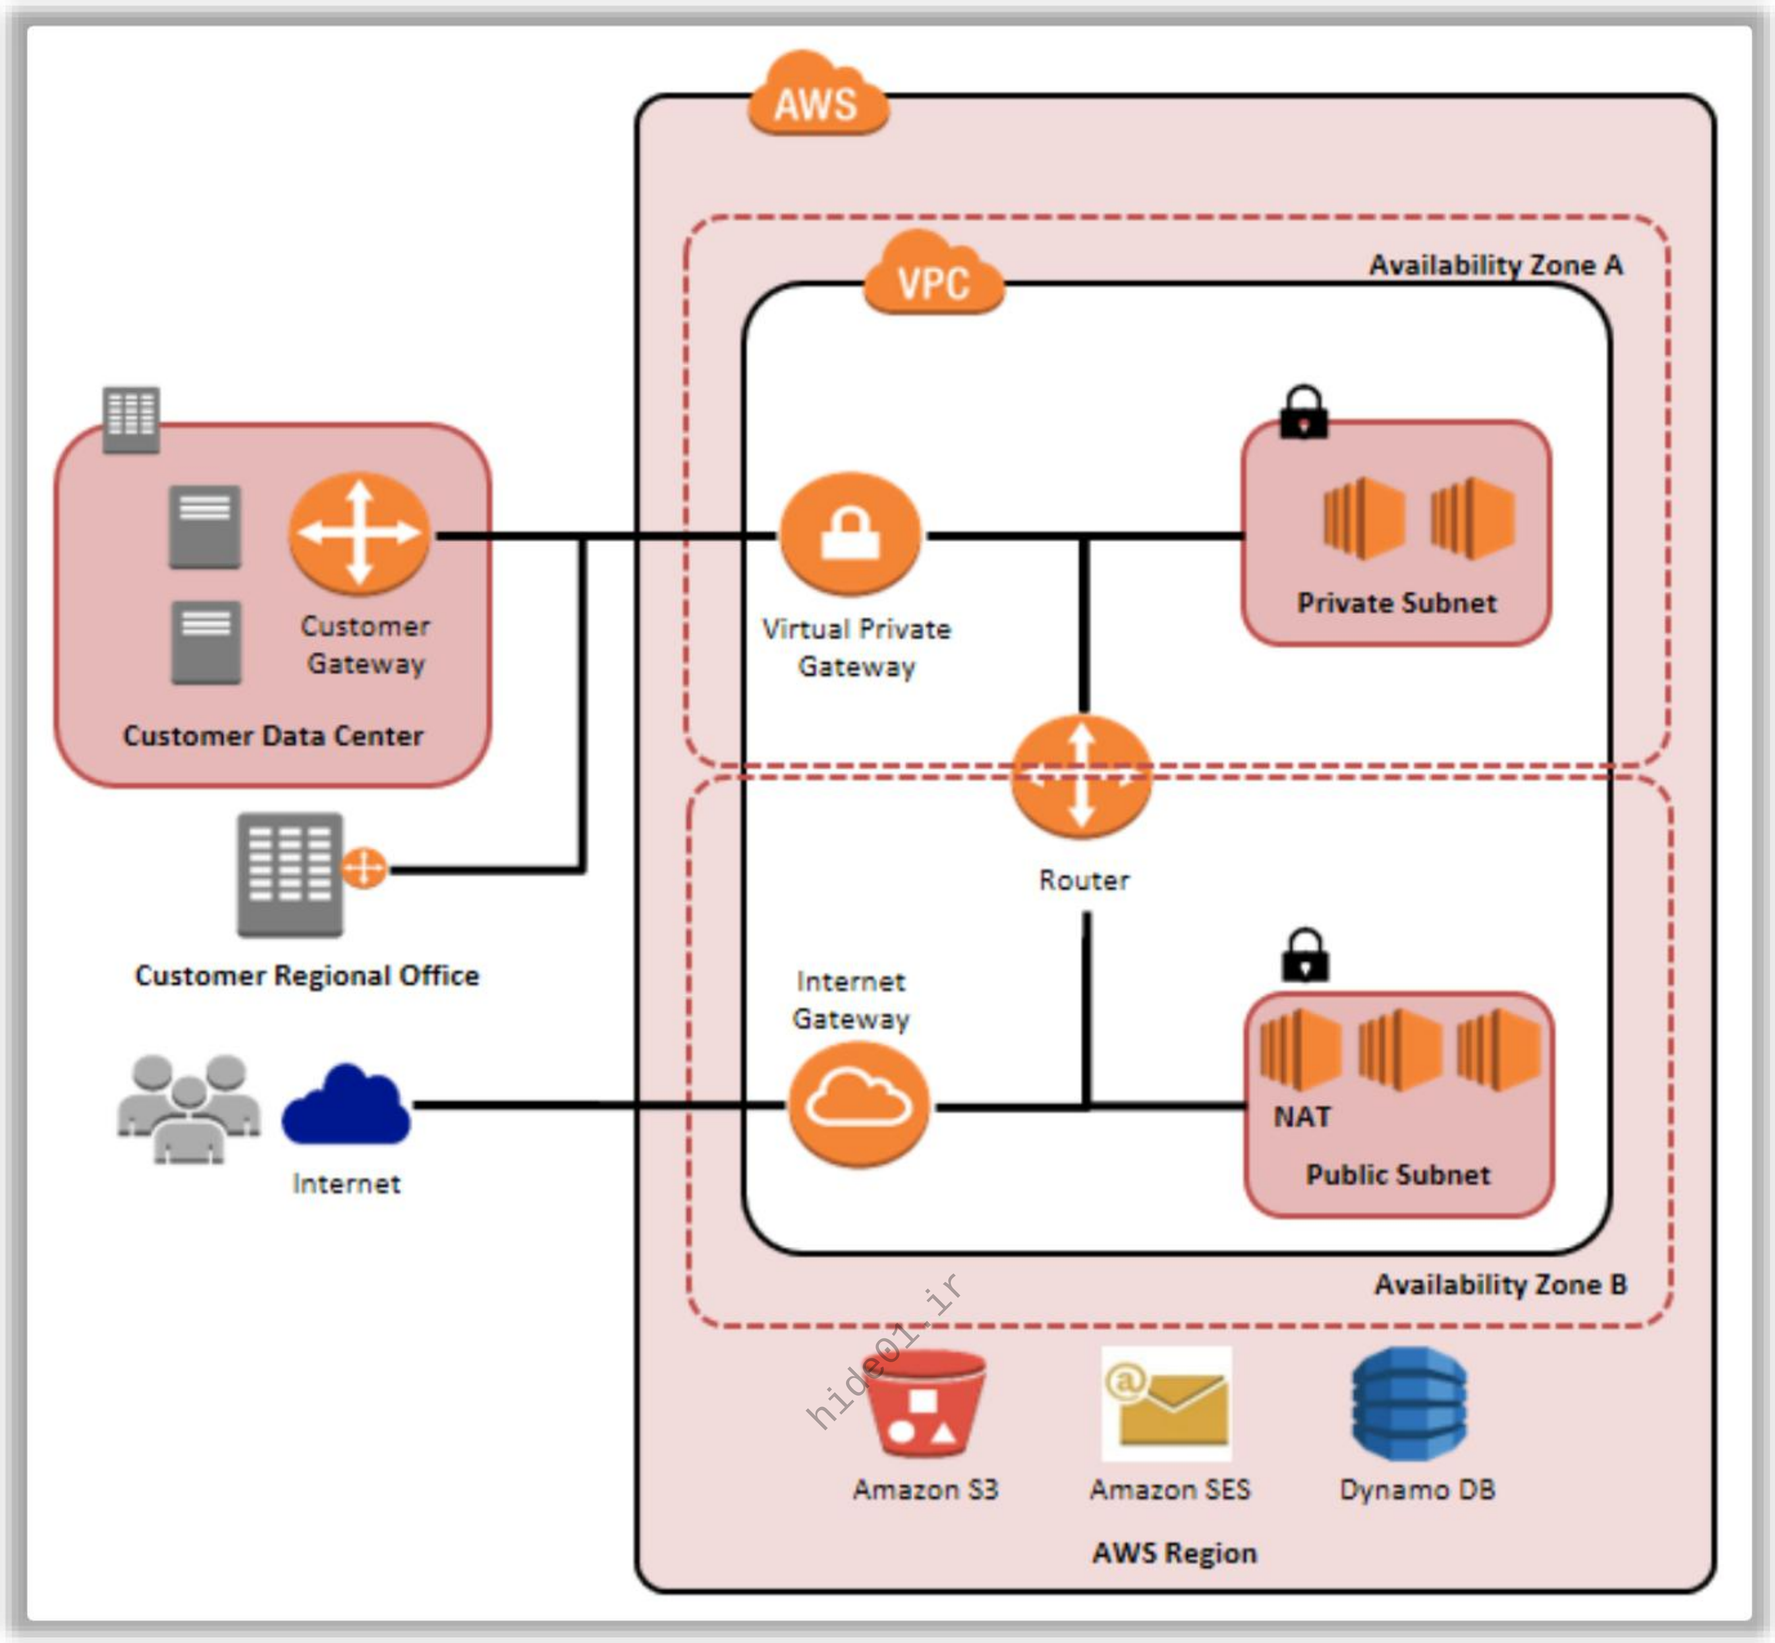


Figure: 2.11: Amazon Virtual Private Cloud Security

It connects two VPCs through a private IP address to establish a communication instance in the same network. Customers can generate a VPC peering connection with each other or other present VPCs from another AWS account of the same region.

The security features in Amazon VPC include security groups, routing tables, external gateways, and network ACLs. Each security feature is complementary to provide a secure isolation network that can be extended by enabling selective private connectivity to another network or through direct access to the internet.

Customers should create VPC security groups for their Amazon VPCs because Amazon EC2 security groups would not work inside Amazon VPC. In addition, Amazon VPC security groups contain additional capabilities that Amazon EC2 security do not possess like the ability to change the security groups after the launch of the instance and the ability to specify any protocol with a standard protocol number. Amazon EC2 instance in an Amazon VPC inherits the benefits of guest's OS and prevents packet sniffing.

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

Every Amazon VPC is a discrete and cloud isolation network. Network traffic for each Amazon VPC is completely isolated from all other Amazon VPCs. Customers can select an IP address range for their Amazon VPCs at the time of creating it. Customers can create and attach a virtual private gateway or internet gateways or both for establishing external connectivity through the below controls.

- **API Access**: Calls to change routing, to create and delete Amazon VPCs, security groups, network ACL parameters, and implementation of other functions are signed by the Amazon Secret Key, which either user or AWS can create. Amazon VPC API are not allowed to make calls on behalf of customers without access to their Secret Access Key. Additionally, SSL can be used to encrypt API calls to maintain confidentiality. The use of SSL-protected API endpoints is recommended by Amazon. Customers can further control what APIs can be called by a newly created user through AWS IAM.
- **Subnets and Route Tables**: Customers create one or more subnets within an Amazon VPC and every instance gets connected to a subnet when it is launched in Amazon VPC. Traditional security attacks of Layer-2 such as ARP spoofing and MAC spoofing are blocked. In an Amazon VPC, every subnet is connected with a routing table that processes network traffic, leaving subnets to detect their destination.
- **Firewall (security groups)**: Amazon VPC offers a complete firewall solution, like Amazon EC2, through egress and ingress traffic filtration from an instance. The default group establishes outbound communication to any destination and inbound communication from members of the same group. Traffic can be restricted or controlled by service port, any protocol, and source/destination IP address. The guest OS cannot control the firewall but can modify it by applying Amazon VPC APIs. Amazon offers customers partial access to the different administrative functions on instance and firewall and allows them to implement extra security by separating duties. The firewall provides security for the function of which ports are used by the customers and its purpose and duration. It still needs security design and well-informed traffic management for a pre-instance basis. Application of per-instance filters with host-based firewalls like the Windows Firewall or IP tables is encouraged by AWS.
- **Network Access Control Lists:** AWS allows customers to add an extra layer of security within Amazon VPC by configuring network ACLs. ACL are considered as stateless traffic filters that are implemented on inbound and outbound traffic subnets in Amazon VPC. They can command orders to allow and restrict traffic based on the IP protocol, by service port, and source/destination IP address. Network ACLs are controlled or managed through Amazon VPC APIs like security groups. These network ACLs provide an extra layer of security and enable additional security via the separation of duties.

  The figure below describes the implementation of security control mentioned above for enabling flexible network topologies that control the traffic flow completely in the network.

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

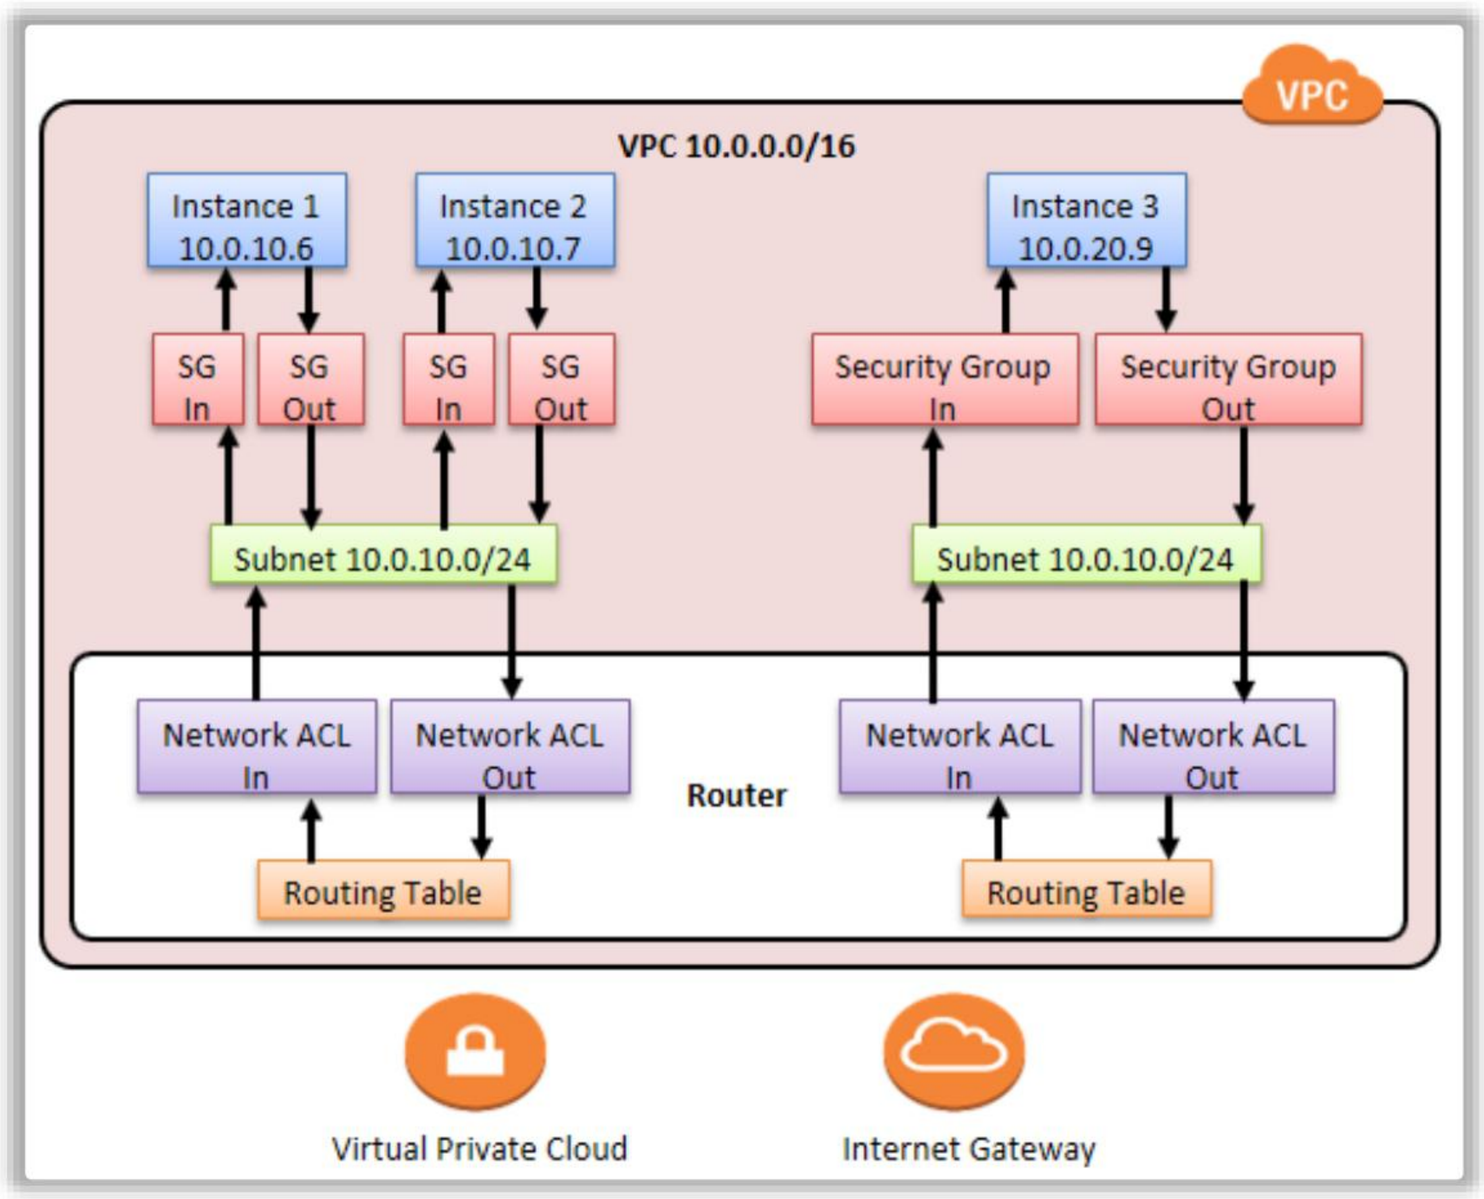


Figure: 2.12: Implementation of AWS Security Controls

- **Virtual Private Gateway (VPG)**: Private connections can be established between an Amazon VPC and another network through a private connectivity is established by the virtual private gateway. The network traffic isolation for every Virtual Private Gateway (VPG) is established from the network traffic within all other VPGs. Customers can create VPN connections from gateway devices at their premises to the VPG. A pre-shared key in conjunction with the customer gateway device's IP address is used to secure each connection.
- **Internet Gateway**: To enable direct connectivity with the internet, Amazon S3, and other AWS services, an internet gateway can be attached to an Amazon VPC. Every instance that wants this access needs to either associate itself with Elastic IP or with route traffic through NAT instance. It provides configuration to direct traffic in the internet gateway. AWS also provides NAT AMIs references that customer can extend to perform deep packet inspection, application-layer filtering, network logging, or other security controls.

  Modification to this access can only be applied by calling Amazon VPC APIs. AWS allows customers partial access to various administrative functions for enabling the implementation of extra security through duties separation in the instance and Internet gateway. To make instances in a private subnet connect to other AWS services and the internet but restrict the internet from initiating a connection with them (instances), customers can use a NAT gateway.
- **Dedicated Instances:** Customers are allowed to launch an Amazon EC2 instance with a high level of physical isolation (which works on single-tenant hardware) within an Amazon

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

VPC. AWS allows the creation of an Amazon VPC with a dedicated tenancy to ensure all instances launched in that VPC use this feature. Similarly, customers can create an Amazon VPC with default tenancy but they can specify dedicated instances for specific instances launched in it.

- **Elastic Network Interfaces**: A default network interface is present in each Amazon EC2 instance assigned with a private IP address on the Amazon VPC network. Customers are allowed to generate and attach an elastic network interface to their Amazon EC2 instances but only two interfaces can be attached per instance. A customer can attach multiple interfaces to an instance for using security and network appliances in their Amazon VPC, to create a management network, or create dual-homed instances with roles/workloads on different subnets. Attributes of a network interface such as the private IP address, MAC address, and elastic IP addresses, follow the network interface because it is attached or detached from an EC2 instance and reattached to another EC2 instance.

hide01.ir

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

Network Service Security: EC2-VPC Network Access Control Features

- All instances will be provisioned automatically in a ready-to-use default VPC when a customer launch instances in a region where there are no instances before the launch of the new AWS EC2-VPC feature
- Security groups for instances in EC2-VPC are different from the security groups for instances in EC2-Classic

Characteristic	EC2-Classic	EC2-VPC (Default VPC)	Regular VPC
		Default IP address until specified by the user otherwise during its launch	Unless the user specifies otherwise during its launch.
Private IP address	The instance gets a private IP address from the EC2-Classic range every time it is launched	The instance gets a static private IP address from your default VPC's address range	The instance gets a static private IP address from the VPC's address range.
Multiple private IP addresses	AWS selects a single IP address for the instance and multiple IP addresses are not supported.	Users can assign multiple IP addresses to their instances.	Users can assign multiple IP addresses to their instances.
Elastic IP address	When the user stops their instance, the elastic IP (EIP) address is disassociated with the instance	When the user stops their instance, the EIP remains associated with the instance	When the user stops their instance, the EIP remains associated with the instance
DNS hostnames	DNS hostnames are by default enabled.	DNS hostnames are by default enabled	DNS hostnames are by default enabled.
Security group	A security group can reference other security groups that are from other AWS accounts.	A security group can reference other security for the user's VPC only.	A security group can reference other security for the user's VPC only.
Security group association	To change the security group, the user must terminate the instance.	The user can change the security group of their running instance.	The user can change the security group of their running instance.
Security group rules	The user can add rules for inbound traffic only.	The user can add rules for both inbound traffic and outbound traffic.	The user can add rules for both inbound traffic and outbound traffic
Tenancy	The user can run the instance on shared hardware and cannot run on single-tenant hardware	The user can run the instance on single-tenant hardware or shared hardware	The user can run the instance on single-tenant hardware or shared hardware

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

Network Services Security: Additional Network Access Control with EC2-VPC

Customers can use Amazon VPC to define a virtual network in their logically isolated area (called VPC) within AWS. Customer can launch Amazon EC2 resources (instances) into the subnets of VPC. Here, the customer's VPC is just like a traditional network that the organization operates in its data center. Customer can configure their VPC, create subnets, select IP address range for VPC, and configure route tables and network gateways. Customers can connect their instances in VPC to the internet or their own data center.

When the customer creates their AWS account, a default VPC will be created for the customer in every region. Here, the default VPC is configured already for customers to use. Customers can create their own nondefault VPC and configure it, or launch their instances immediately into their default VPC.

Differences Between Instances in EC2-VPC and EC2-Classic

Use the security groups created specifically for EC2-Classic when using EC2-Classic. Specify a security group in the same region as the instance while launching an instance in EC2-Classic. Customers cannot specify the security group of VPC created for a VPC when an instance in EC2-Classic is launched.

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

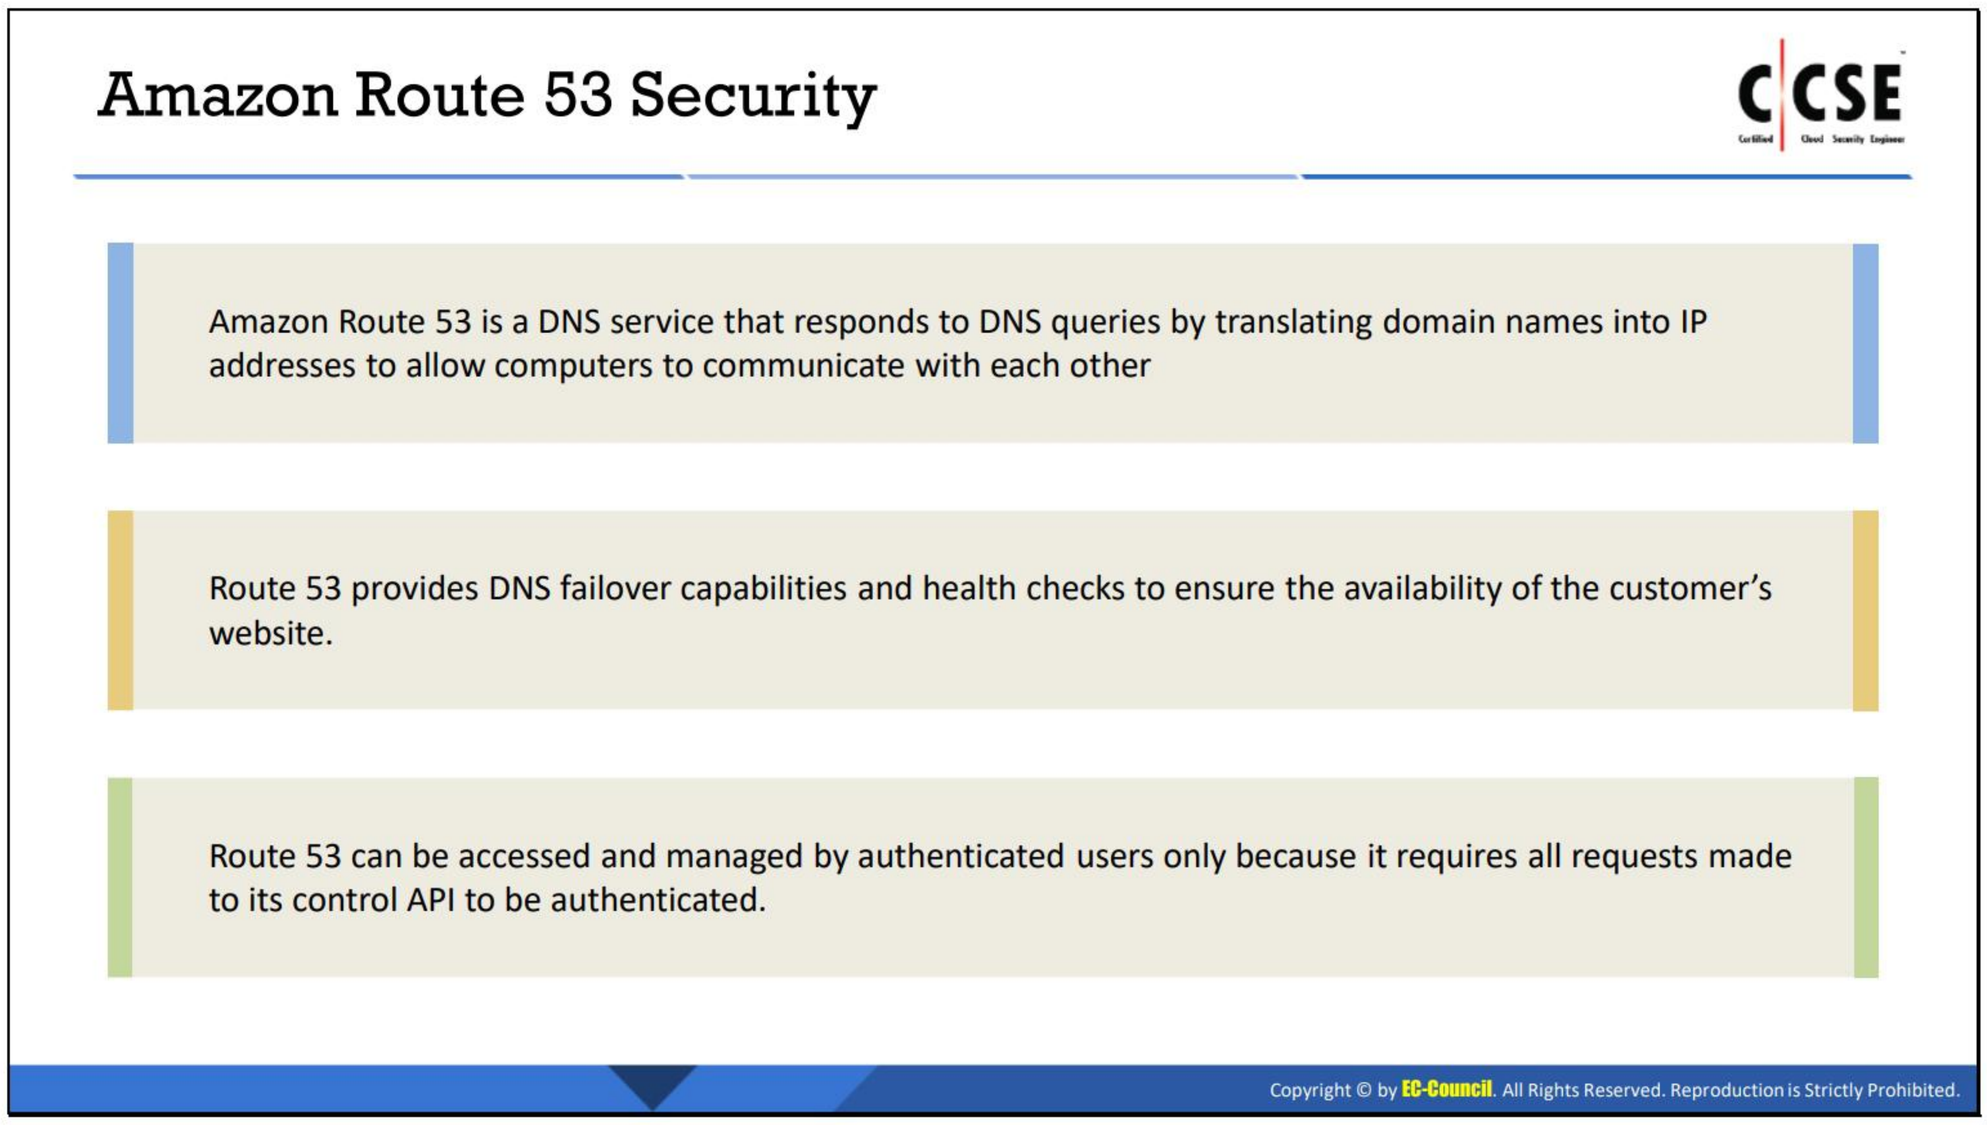
Security groups for instances in EC2-VPC are somewhat different from the security groups for instances in EC2-Classic. The differences between the instances launched in EC2-VPC, EC2-Classic, and nondefault VPC are provided in the below table.

Characteristic	EC2-Classic	EC2-VPC (Default VPC)	Regular VPC
		IP address by default; unless the user specifies otherwise during its launch.	Unless the user specifies otherwise during its launch.
Private IP address	The instance gets a private IP address from the EC2-Classic range every time it is launched	The instance gets a static private IP address from your default VPC's address range	The instance gets a static private IP address from the VPC's address range.
Multiple private IP addresses	AWS selects a single IP address for the instance and multiple IP addresses are not supported.	Users can assign multiple IP addresses to their instances.	Users can assign multiple IP addresses to their instances.
Elastic IP address	When the user stops their instance, the elastic IP (EIP) address is disassociated with the instance	When the user stops their instance, the EIP remains associated with the instance	When the user stops their instance, the EIP remains associated with the instance
DNS hostnames	DNS hostnames are by default enabled.	DNS hostnames are by default enabled.	DNS hostnames are by default enabled.
Security group	A security group can reference other security groups that are from other AWS accounts.	A security group can reference other security for the user's VPC only.	A security group can reference other security for the user's VPC only.
Security group association	To change the security group, the user must terminate the instance.	The user can change the security group of their running instance.	The user can change the security group of their running instance.
Security group rules	The user can add rules for inbound traffic only.	The user can add rules for both inbound traffic and outbound traffic.	The user can add rules for both inbound traffic and outbound traffic
Tenancy	The user can run the instance on shared hardware and cannot run on single-tenant hardware	The user can run the instance on single-tenant hardware or shared hardware	The user can run the instance on single-tenant hardware or shared hardware

Table 2.2: Difference among EC2 Classic, EC2 VPC, Regular VPC

hide01.ir

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



Amazon Route 53 Security

Amazon Route 53 is a DNS service that is highly available and scalable. It responds to DNS queries by translating domain names into IP addresses to allow computers to communicate with each other. It is helps in connecting users request to infrastructure running in AWS such as Amazon S3 bucket and an Amazon EC2 instance or infrastructure that is outside AWS.

Customers can use Amazon Route 53 to manage IP addresses (records) that are listed for their domain names. Amazon Route 53 performs request translation of particular domain names into their corresponding IP addresses. To ensure the lowest latency, queries for a customer's domain are routed to a nearby DNS server automatically by using AnyCast. By using Amazon Route 53, customers can manage traffic globally through a variety of routing types such as Geo DNS, Latency Based Routing (LBR), and Weighted Round-Robin (WRR). All routing types combine with DNS Failover to form various types of low-latency and fault-tolerant architectures. Amazon Route 53 implements the failover algorithms that are designed to route traffic to secure endpoint along with protecting it from worsening disaster recovery scenarios caused by endpoint overloads, partition failures, and misconfigured health checks and applications.

Amazon Route 53 provides domain name registration so that users can buy or maintain domain names, and Amazon Route 53 would prepare DNS settings for them. Users can select the purchase, maintenance, and transfer of domains over a wide range of general or country-specific top-level domains. Users can select privacy protection for their domain during the registration process. The privacy protection option protects the user's personal data from getting compromised through spamming and scraping in public with the help of the Whois database.

AWS has used its highly reliable and available architecture to build Amazon Route 53. AWS DNS servers' distributed nature helps in providing a consistent ability to route end-users to a customer's application. Route 53 provides DNS failover capabilities and health checks to ensure

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

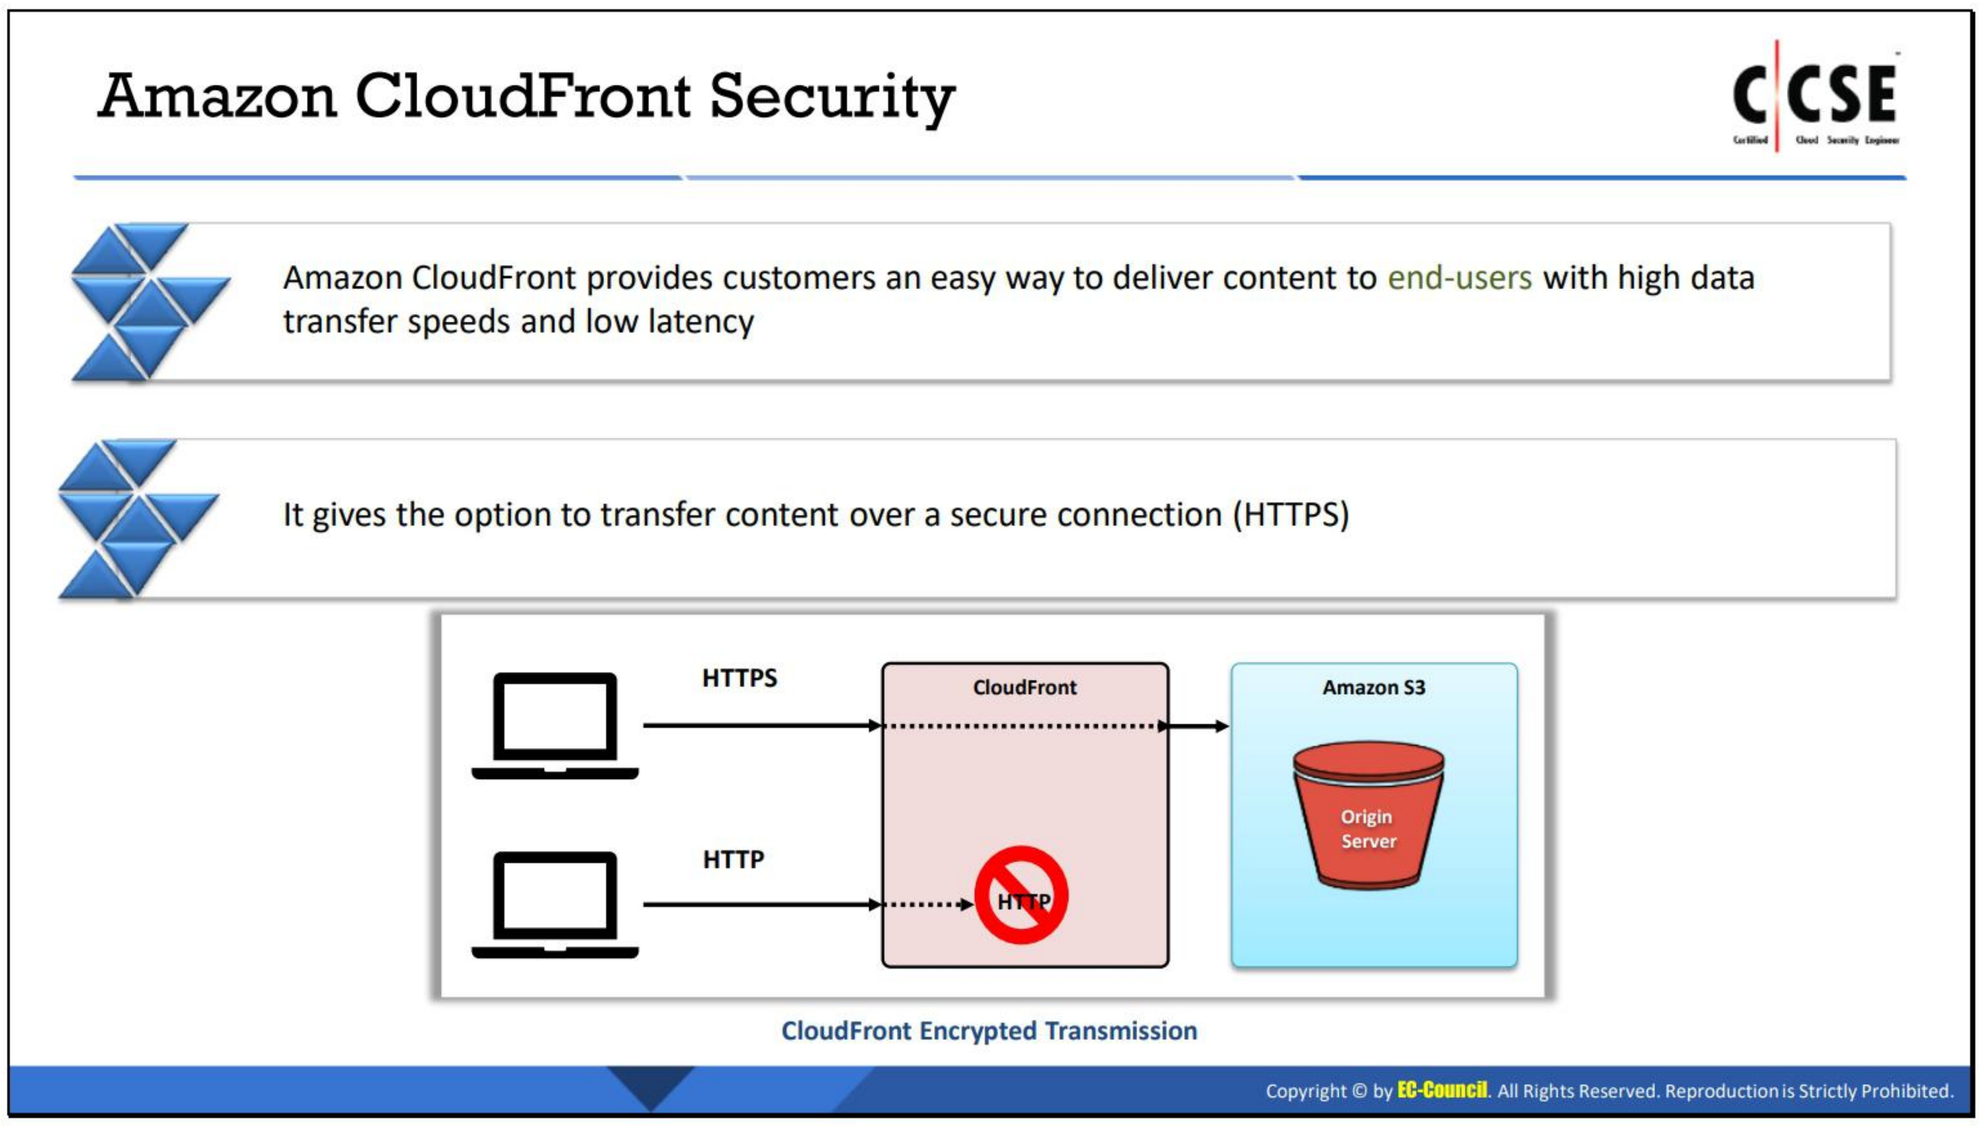
the availability of the customer's website. Customers can configure Route 53 easily to examine their websites' health on a regular basis and to switch (or failover) to a backup site in the event when the primary one becomes unresponsive.

Like all other AWS services, Route 53 can be accessed and managed by authenticated users only because it requires all requests made to its control API to be authenticated. All API requests are signed with an HMAC-SHA256 or HMAC-SHA1 signature, which is calculated from the user's AWS Secret Access Key and the request. Additionally, Amazon Route 53 control API can be accessed by customers only through SSL-encrypted endpoints and assists both IPv4 and IPv6 routing.

An organization can create a user in its AWS Account by employing AWS IAM to access Amazon Route 53 DNS management functions and control user's permission to execute Route 53 operations.

hide01.ir

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



Amazon CloudFront Security

Amazon CloudFront provides customer organizations with a simple way to distribute content with high data speed and low latency to end-users. It works with the global network of edge location to deliver static, dynamic, and streaming content. Content is delivered with the best performance because on customers request their object can be sent to the nearest edge location automatically. It is optimized to ensure that it works well alongside AWS services such as Amazon EC2, Amazon S3, Amazon Route 53, and ELB. It works uninterruptedly with any server that is of non-AWS origin and stores the definitive, original version of customers' files.

Like all other AWS services, it also requires all requests made to its API to be authenticated so that only authenticated are allowed create, delete, or modify Amazon CloudFront distributions. API requests are signed with an HMAC-SHA1 signature, which is calculated from the user's private key and the request. Additionally, Amazon CloudFront control API can be accessed only through SSL-encrypted endpoints.

The durability of the data held in CloudFront edge location cannot be guaranteed. If some objects are not requested frequently, the service may sometimes remove those objects from the edge locations. Amazon S3 works as the origin server for CloudFront and provides durability as it holds the definitive, original copies of objects that are delivered by CloudFront.

Customers can enable the private content feature of the service if they want control over who can download content from Amazon CloudFront. The private contain feature consists of two main components. The first component controls the delivery of the content from Amazon CloudFront edge location to viewers on the Internet. The second component controls how objects are accessed in Amazon S3 by Amazon CloudFront edge locations. It also supports geo restriction, which restricts customers' access to content based on the viewer's geographical location.

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

Customers can create one or more Origin Access Identities by using CloudFront and associate them with their distributions to control access to the original copies of their objects present in Amazon S3. When an Amazon CloudFront distribution gets associated with an Origin Access Identity, it can use that identity for retrieving objects from Amazon S3. The ACL feature of Amazon S3 can be used to prohibit access to Origin Access Identity to protect original copy from getting publicly compromised and read.

The CloudFront service uses a signed-URL verification system to control who can download objects from CloudFront edge locations. A customer can use this system by following the steps given below.

- The customers should generate a public-private key pair and then upload the public key to their respective account using AWS Management Console.
- The customers should generate a public-private key pair and then upload the public key to their respective account using AWS Management Console.
- The customer can set up their CloudFront distribution to identify the AWS account that they would authorize to sign requests. The customer can indicate up to five accounts they trust to sign requests.
- When the customer receives requests, they can create policy documents specifying the terms and conditions they want to follow for the optimize benefit in their context from CloudFront. The policy documents should contain the date and time of the request, the name of the requested object, and the source IP or CIDR range of the requesting client for calculating the SHA1 hash of their policy document and uses their private key to sign it.
- The customer includes both the signature and the encoded policy document as query string parameters when the customer references their objects.
- Amazon CloudFront decodes the signature by using the customer's private key when it receives a request and it only serves the requests that have a matching signature and valid policy document.

Note: AWS offers an optional feature called Private content that must be enabled while setting up a CloudFront distribution. When the content gets delivered without implementing this feature, it can be read publicly.

It provides the choice of transferring content to a secure connection (like HTTPS). It takes calls from both HTTPS and HTTP by default. Although customers can set up HTTPS requests and divert all the HTTP calls to HTTPS. CloudFront distribution can prepare HTTP for some objects, but it needs HTTPS for other objects.

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

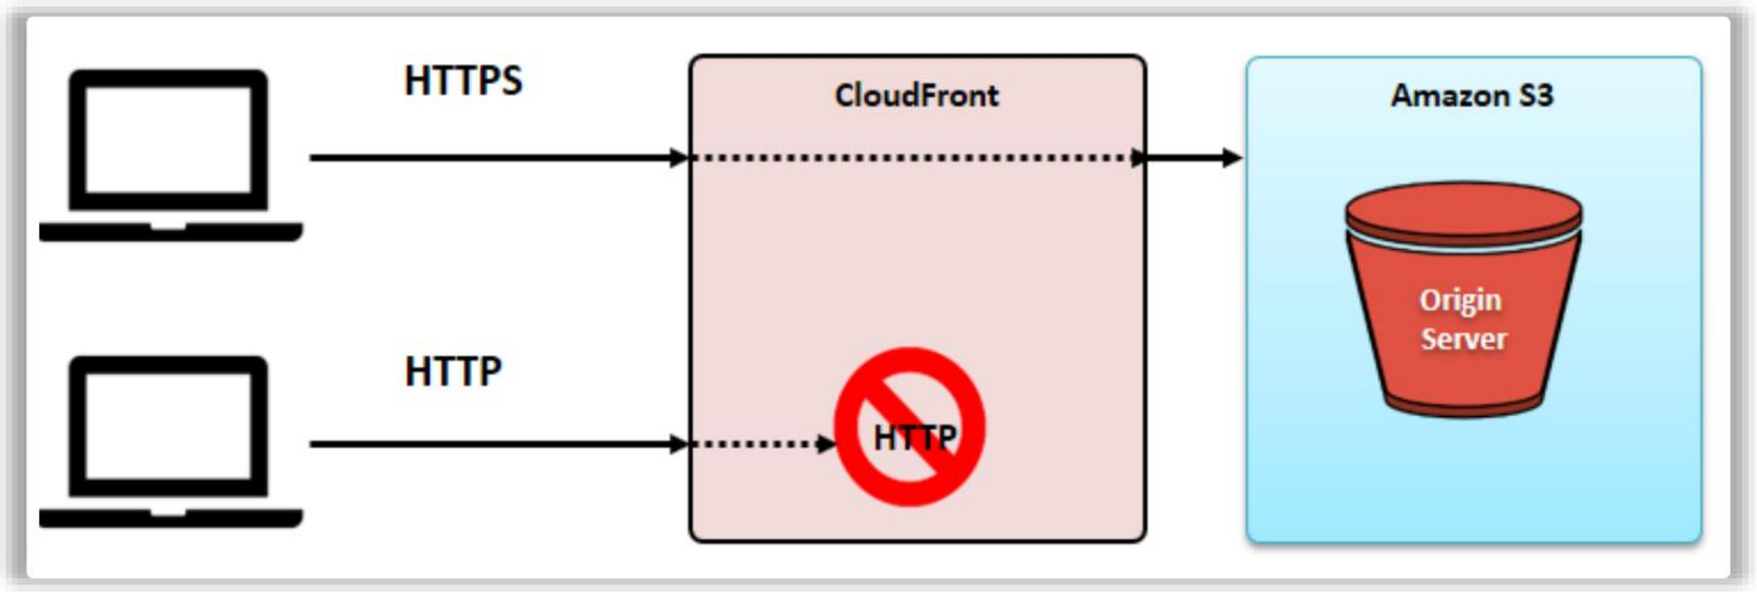


Figure: 2.13: CloudFront Encrypted Transmission

A customer can configure one or more CloudFront origins to make CloudFront fetch objects from the customer's origin by using the protocol that was used by the viewer askes for the object. For example, whenever CloudFront setting is used by the customer and the viewer requests an object from CloudFront using HTTPS, Amazon CloudFront also forwards the request to the customer's origin by using HTTPS.

The TLSv1 or SSLv3 protocols and a range of cipher suites are used by CloudFront (including Elliptic Curve Diffie-Hellman Ephemeral (ECDHE) protocol) on connections to both the origin and viewers. SSL/TLS clients provides Perfect Forward Secrecy through ECDHE protocol, which uses session keys that are not stored by Amazon CloudFront anywhere and are ephemeral. This stops any unauthorized third party from decoding captured data even if the secret long-term key gets stolen or compromised.

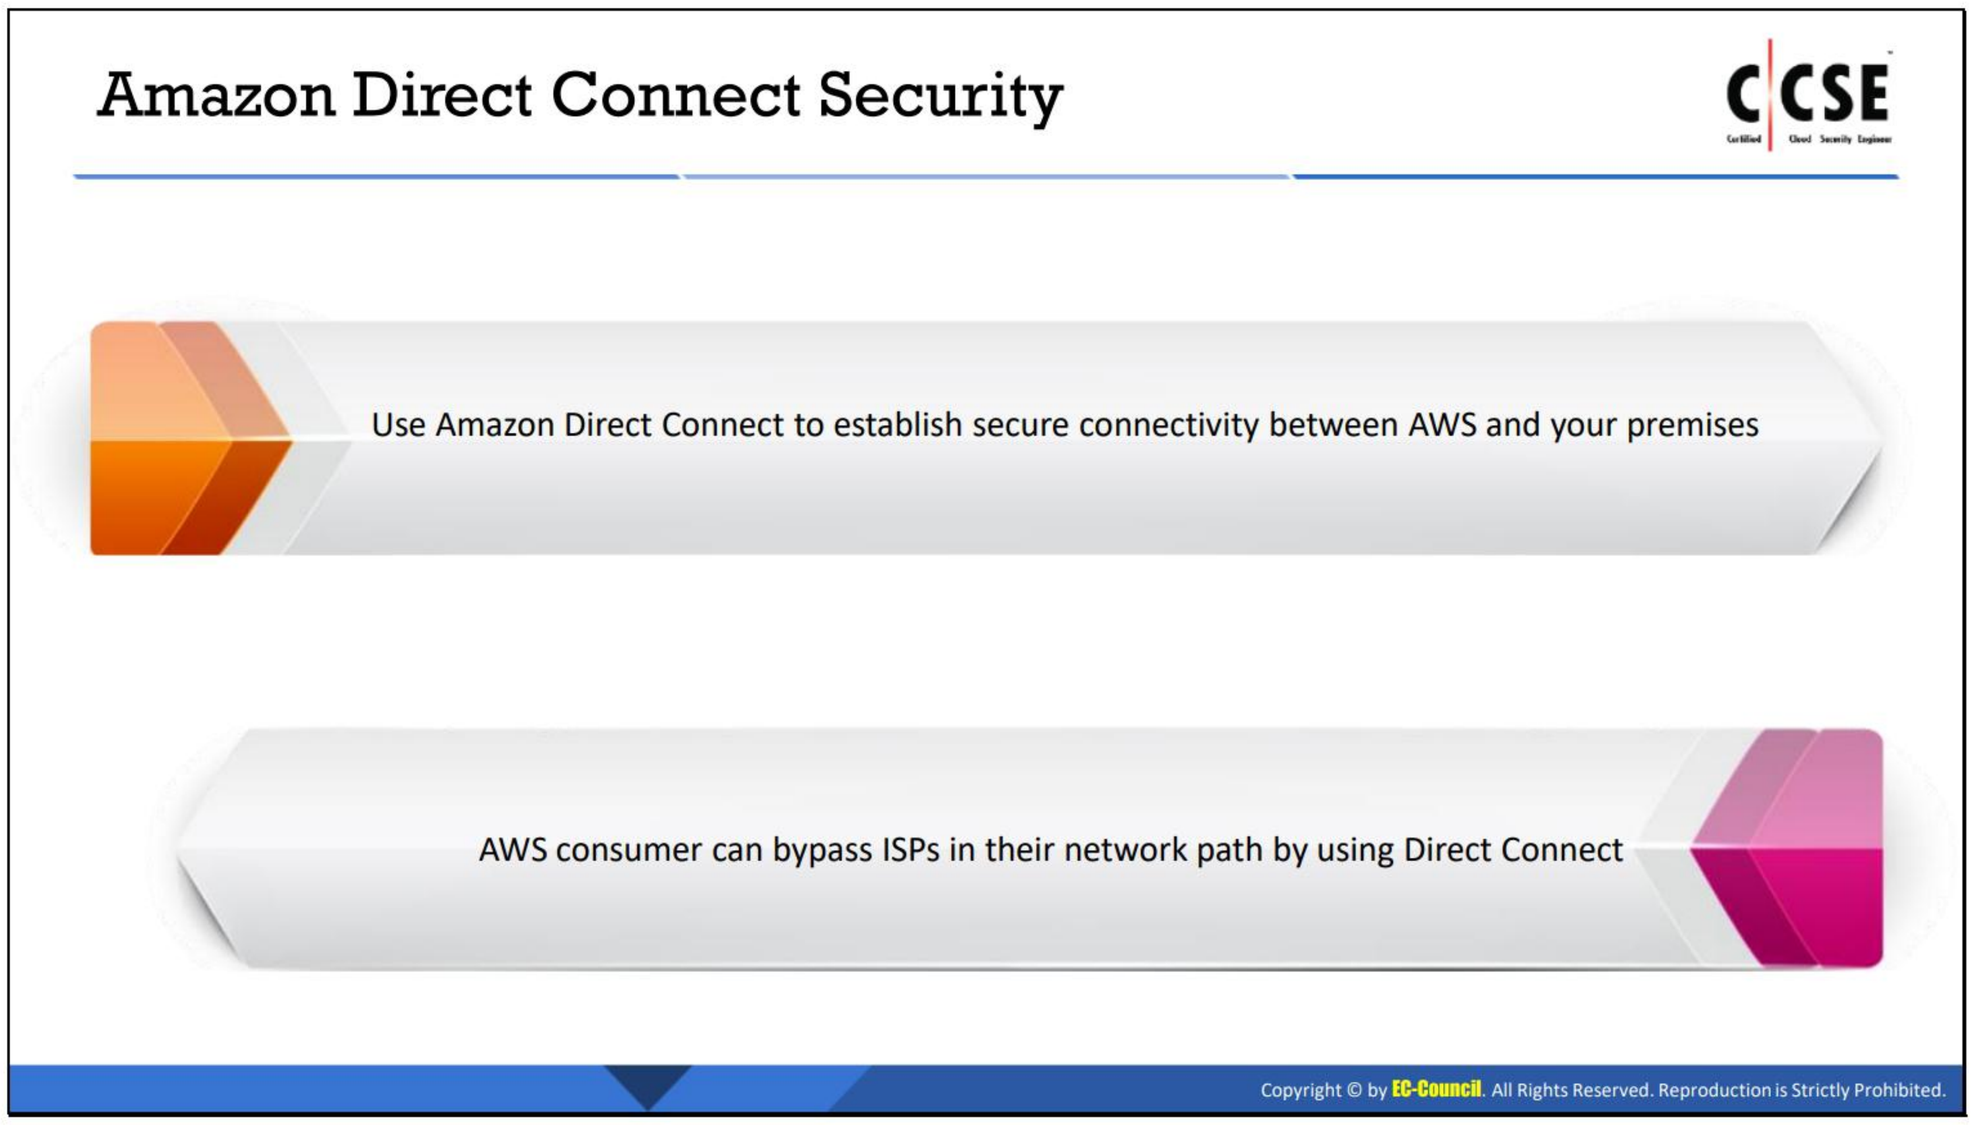
hide01.ir

Note: If a customer is utilizing their server for origin, and willing to use it for HTTPS between CloudFront and viewers and between their origin and CloudFront, they need to install a valid SSL certificate on the HTTP server signed by a third-party certificate authority such as DigiCert or VeriSign.

By default, by using their CloudFront distribution domain name in their URLs, the customer can deliver content to their viewers over HTTPS. For example, https://cxxxx.cloudfront.net/image.jpg. The customer can use dedicated IP custom SSL or SNI custom SSL if they want to deliver their content over HTTPS by using their own domain name and SSL certificate. CloudFront relies on the server name identification (SNI) extension of the TLS protocol (when used with SNI Custom SSL) as it is supported by most modern web browsers. But some users who use old browsers cannot access the customer's content because old browsers do not support SNI. CloudFront dedicates (assigns) IP addresses to the customer's SSL certificate at every CloudFront edge location when it is used with dedicated IP custom SSL so that the incoming requests can be associated with the proper SSL certificate by CloudFront.

CloudFront access logs consist of all the necessary information regarding content requests such as the object requested, the edge location serving the request, the date and time of the request, the referrer, the user agent, and the client IP address. To activate access on the logs while configuring an Amazon CloudFront distribution, the customer needs to determine Amazon S3 bucket name to store the logs in it.

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



Amazon Direct Connect Security

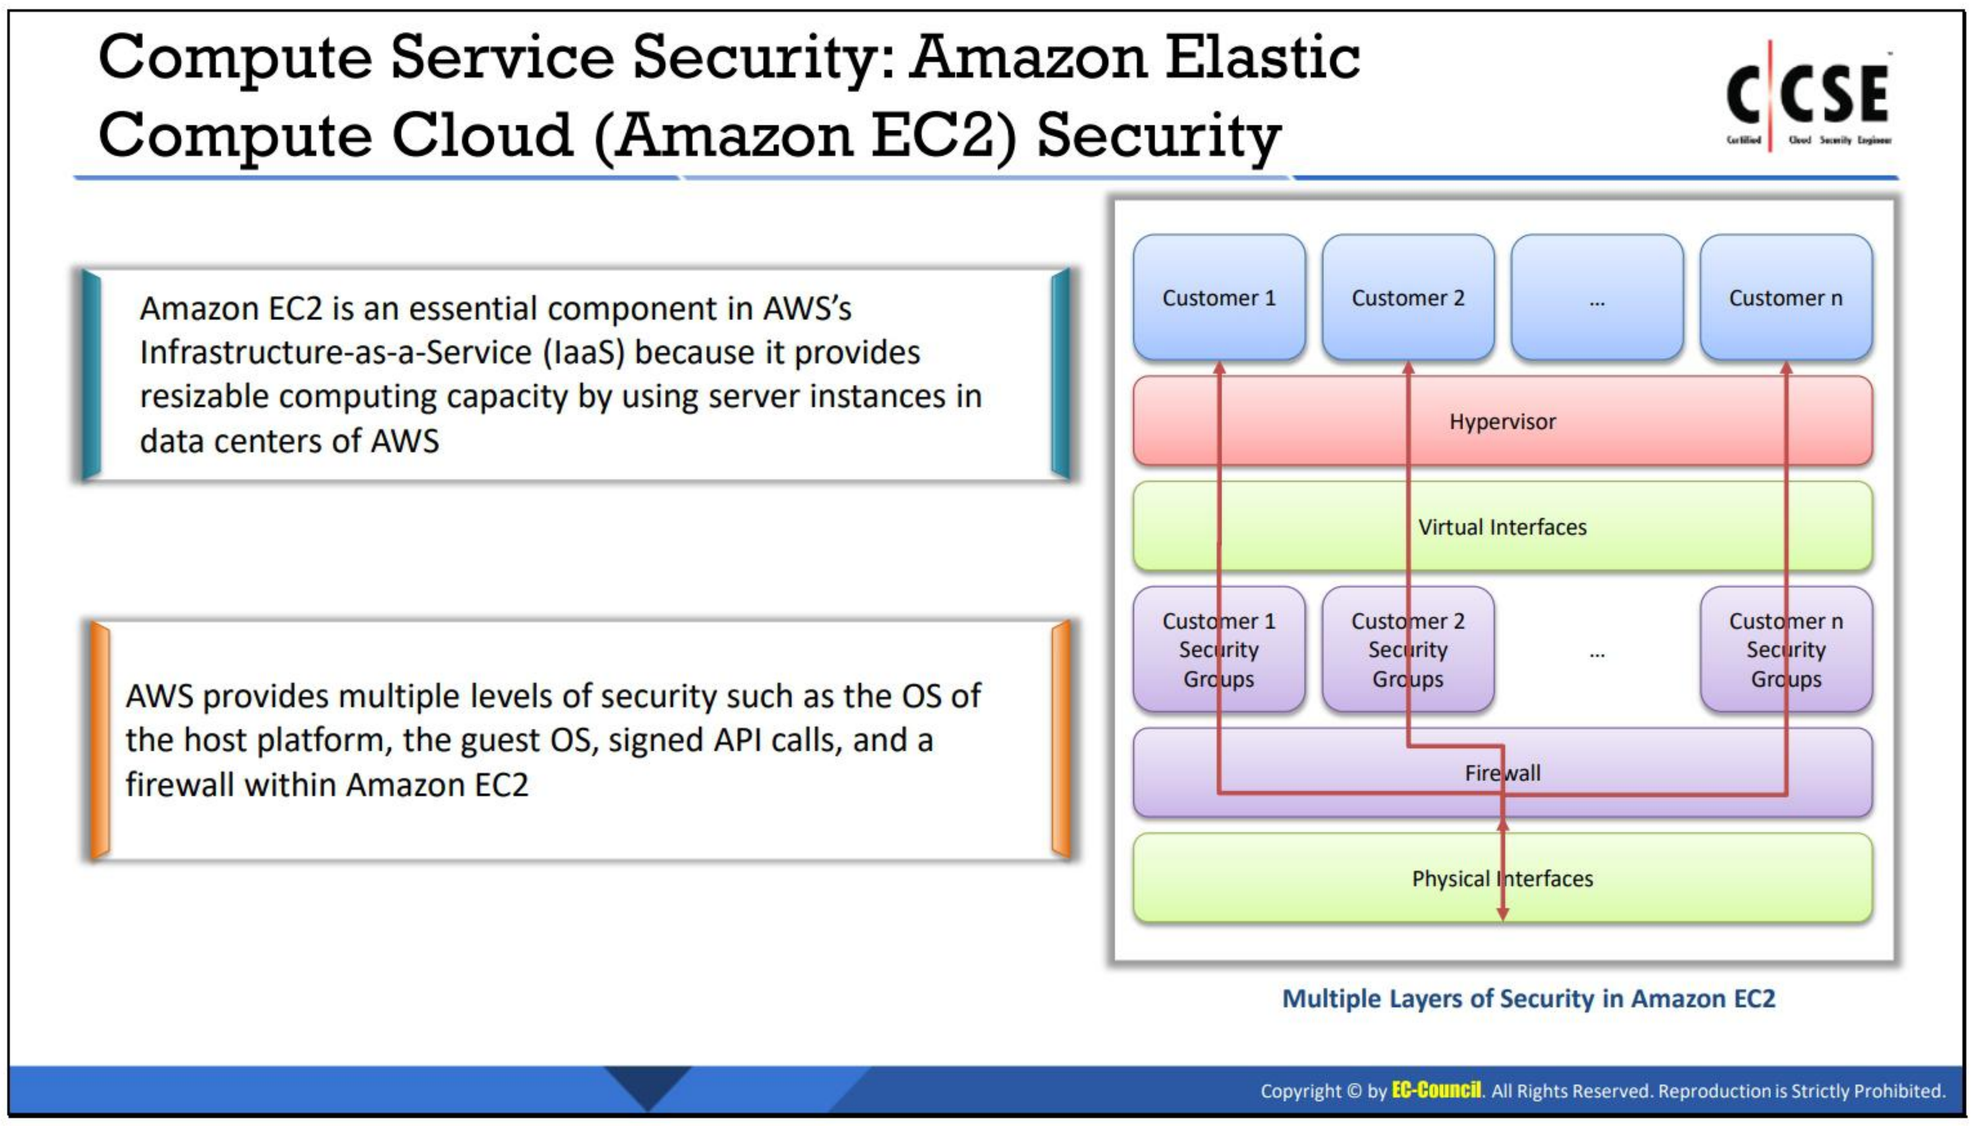
With the help of Amazon Direct Connect, a customer can provide a direct link between an AWS region and their internal network by using a dedicated, high-throughput connection.

hide01.ir

Doing this can help in reducing the customer's network costs, improving throughput, deliver a consistent network experience. Instead of a dedicated connection, the customer can create virtual interfaces to the AWS cloud directly (for example, to Amazon S3 or Amazon EC2) and Amazon VPC. Customers bypass ISPs in their network path by using Direct Connect. They can gain rack space within the facility that contains the AWS Direct Connect location and use it to deploy their equipment. Once their equipment is deployed, they can use a cross-connect to connect that equipment to Amazon Direct Connect. Each Amazon Direct Connect location enables connectivity to the nearest AWS region geographically and access to other US regions. For example, a customer can establish one connection to Amazon Direct Connect location to use it for accessing AWS GovCloud (US) and public AWS services in all US Regions. The dedicated connection utilizes industry-standard 802.1q VLANs for getting divided into different virtual interfaces. This allows customers to utilize the same connection to access private resources like the Amazon EC2 instance utilize private IP space for working in Amazon VPC and public resources like objects stored in Amazon S3 using public IP address space for maintaining segregation between the private and public environment.

Amazon Direct Connect requires the BGP to utilize with an Autonomous System Number. An MD5 cryptographic key can be used by the customer for message authorization to create a virtual interface. By using the customer's secret key, MD5 creates a keyed hash. Customers can provide their own BGP MD5 keys or can have AWS automatically generate them.

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



Compute Service Security: Amazon Elastic Compute Cloud (Amazon EC2) Security

AWS offers several cloud-based computing services that contain an extensive computer instance range that can be scaled automatically up and down according to the requirement of customer or enterprises.

hide01.ir

Amazon Elastic Compute Cloud (Amazon EC2) Security

Amazon EC2 is an important component in AWS's IaaS because it provides resizable computing capacity by using server instances in the data centers of AWS. It makes web-scale computing easier by allowing AWS customers to acquire and configure capacity without any issues. Customers can create and launch the instance that is the platform hardware and software clusters.

Multiple Levels of Security

AWS provides multiple levels of security within Amazon EC2: the OS of the host platform, the guest OS, signed API calls, and a firewall which aids and improves the other's capabilities. The main objective is to secure stored data in Amazon EC2 from unauthorized users and provide high security to Amazon EC2 instances; it also provides on demand flexible configuration for AWS customers.

Hypervisor

Today, Amazon EC2 leverages a highly customized version of Xen hypervisor, benefiting from paravirtualization (in Linux guest case). The guest OS does not have elevated access to the CPU as para-virtualized guests depend on the hypervisor to reinforce an operation that normally requires privileged access. The CPU offers four distinct privilege modes, 0-3, called rings. Of these modes, Ring 0 is the most privileged, while Ring 3 is the least privileged. The host OS executes in

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

the most privileged ring (i.e., Ring 0). Unlike most OS, the guest OS works in a lesser-privileged Ring 1, and applications work in the least privileged Ring 3. The transparent virtualization of the physical resources results in a clear separation between the hypervisor and guest, which results in additional security between them.

Generally, hypervisors protect BIOS, physical hardware, virtualize the CPU, networking, storage, and provide a large number of management capabilities. The Nitro System can separate those functions, offload (unburden) them to dedicated hardware and software, and save money by delivering all resources of a server to customers' instances.

The Nitro Hypervisor delivers consistent performance and increased memory and compute resources for EC2 virtualized instances by eliminating host system software components. The Nitro Hypervisor allows AWS to provide large instance sizes (such as c5.18xlarge) that practically offer all the resources from the server to customers. By moving VPC and Elastic Block Store (EBS) functionality to the hardware that was designed and built by AWS, C3 and C4 instances removed software components previously. This hardware allows the Nitro Hypervisor to be small and uninvolved in tasks related to data processing for storage and networking.

Amazon EC2's use of the Xen Hypervisor will increase as AWS expands its cloud infrastructure globally because Xen will continue to be a core component of EC2 instances for the predictable future.

Instance Isolation

The Xen hypervisor isolates different instances that exist on the same physical machine. An instance's neighbor or other instance cannot access another instance and is treated as if they exist on separate physical hosts. AWS uses similar mechanisms to separate the physical RAM.

Raw disk devices are not accessible to customer instances; however, they are provided with virtualized disks instead. To avoid unintentional exposure of one customer's data to other customers, every block of storage that is used by the customer is reset automatically by the AWS proprietary disk virtualization layer. In addition, when memory is unallocated to a guest, it is cleaned by the hypervisor. Until the memory scrubbing (cleaning) is complete, the memory is not returned to the pool of free memory that is available to customers for new allocations.

To further protect customers' data, AWS recommends them to use appropriate means. Running an encrypted file system on top of the virtualized disk device is one common solution.

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

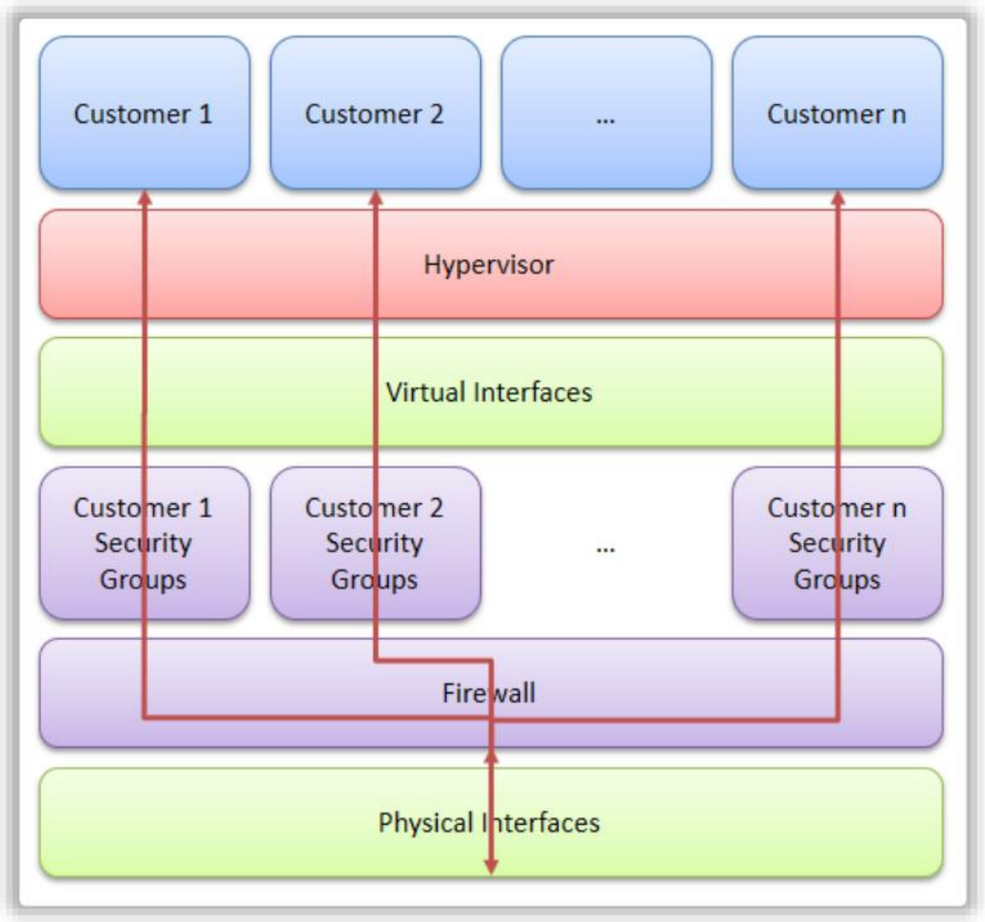


Figure: 2.14: Multiple Layers of Security in Amazon EC2

- **Host Operating System:** Administrators who have a business requirement to access the management plane are required to utilize MFA to get access to purpose-built administration hosts. AWS has specially designed, built, configured, and hardened these administrative hosts to protect the cloud management plane. This access tries to get logged and audited. The privilege and access to these hosts and other relevant systems can be revoked when an employee does not desire to access the management plane for business requirements.
- **Guest Operating System**: Customers have complete control over virtual instances. They have administrative control over services, account, and applications. AWS does not possess access right to customer instance or the guest OS. Some security best practices are recommended by AWS for customer organizations such as disabling password access to their guests and utilizing MFA to access their instances (or minimum certificate-based SSHv2 access). Although, organizations should utilize a privilege escalation mechanism for logging under the user; for example, when Linux is the guest operating system, after security hardening the customer instance, they must use certificate-based SSH version 2 to access the virtual instance, use command-line logging, disable remote login, and use sudo for privilege escalation. Customers must create their key pairs to ensure uniqueness and avoid sharing with AWS customers or AWS.

  To enable customers to log in to their Linux/UNIX EC2 instances, AWS supports the SSH network protocol. To avoid unauthorized access to customer instances, authentication is performed through a public/private key pair for SSH used with AWS. Customers can connect to their Windows instances remotely using Remote Desktop Protocol by using a remote desktop protocol (RDP) certificate generated for their instances.

  Customers are allowed to manage the patching and updating of their guest OS (including security updates). Amazon offered Linux and Windows-based AMIs that are simultaneously updated with the latest patches. When customers do not want to protect customizations or data on their working AMI instances, they can relaunch new instances with the latest updated

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

AMI. Additionally, AWS provides updates for the Linux AMI using the Amazon Linux yum repositories.

- **Firewall**: A complete firewall solution is provided by Amazon EC2, and this firewall is set up in a default deny-all mode for Amazon EC2 customers need to explicitly open the ports they need to allow inbound traffic. The network traffic may be restricted by the service port, protocol, and source IP address (CIDR block or individual IP).

  Customers can configure the firewall in groups, setting different rules for various classes of instances. For example, suppose the case of a traditional three-tier web app. The webserver group contains port 443 (HTTPS) and/or port 80 (HTTP) open to the internet. The webserver group can access the application server group with port 8000. The database server group can access port 3306 (MySQL) only for the application server group. Administrative access would be permitted by all these groups on port 22 (SSH) but from the customer's corporate network only. This expressive mechanism can be used to deploy highly secure applications.

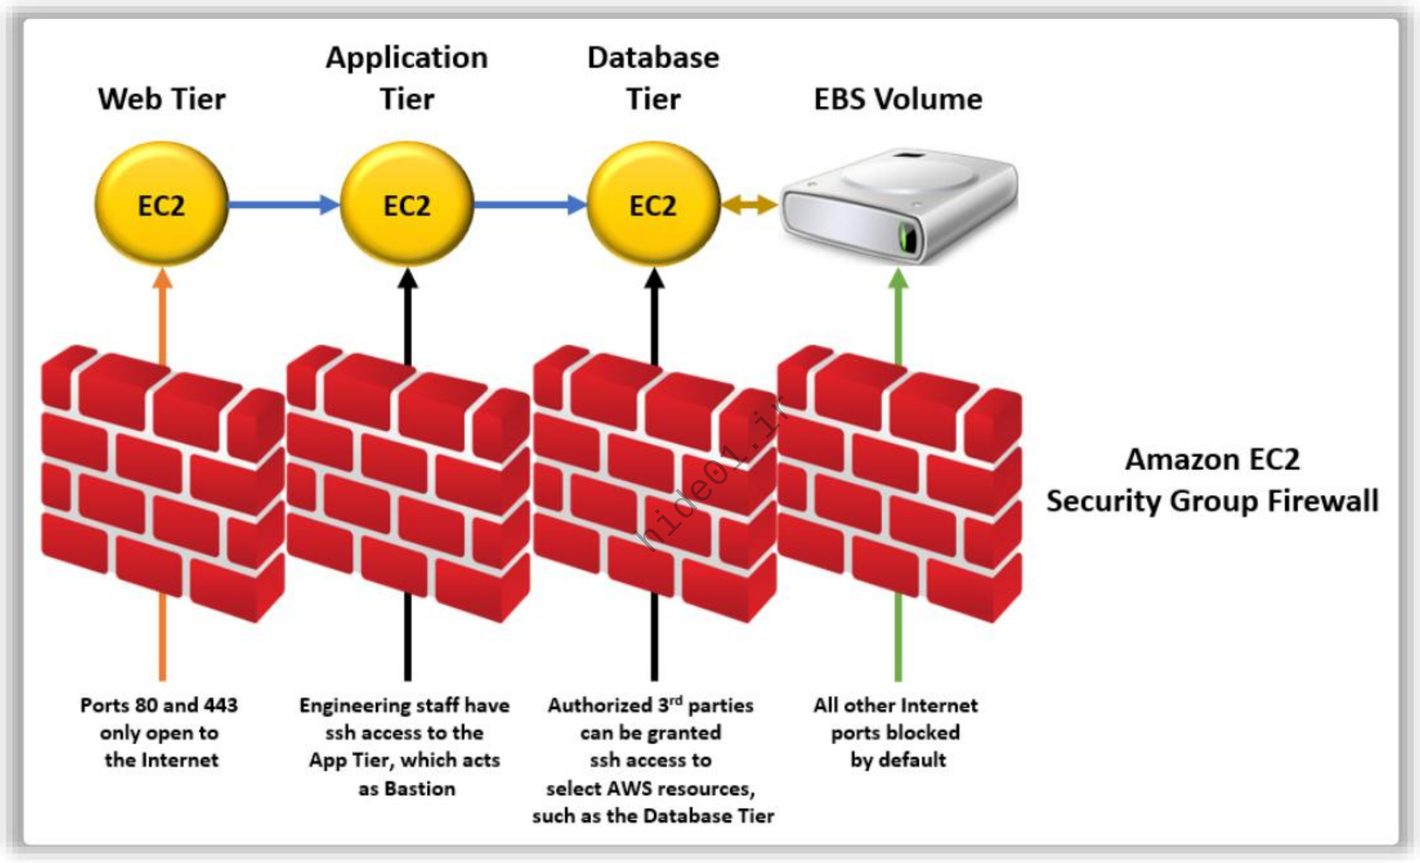


Figure: 2.15: Amazon EC2 Security Group Firewall

Amazon EC2 Security Group Firewall: This firewall cannot be controlled through the guest OS. To authorize changes, it requires the customer X.509 certificate and key, adding an extra layer of security. AWS can grant granular access to administrative functions on the EC2 instances and the firewall, thereby allowing customers to implement additional protection through the separation of duties. The security level for the firewall providers possesses a function in which ports are opened for the customers and ask for the purpose and duration. The firewall blocks all incoming traffic in its default state; thus, customers should plan carefully what ports should be opened by them while building and securing their applications. Security design and well-informed traffic management are still needed on a per-instance basis. The use of additional per-instance filters with host-based firewalls like the Windows Firewall or IP tables and VPNs is encouraged by AWS. By doing this, customers can restrict both incoming and outgoing traffic.

API Access: A customer's Amazon Secret Access Key is used to sign all API calls to change firewall parameters, launch and terminate instances, and execute other functions. Customers cannot

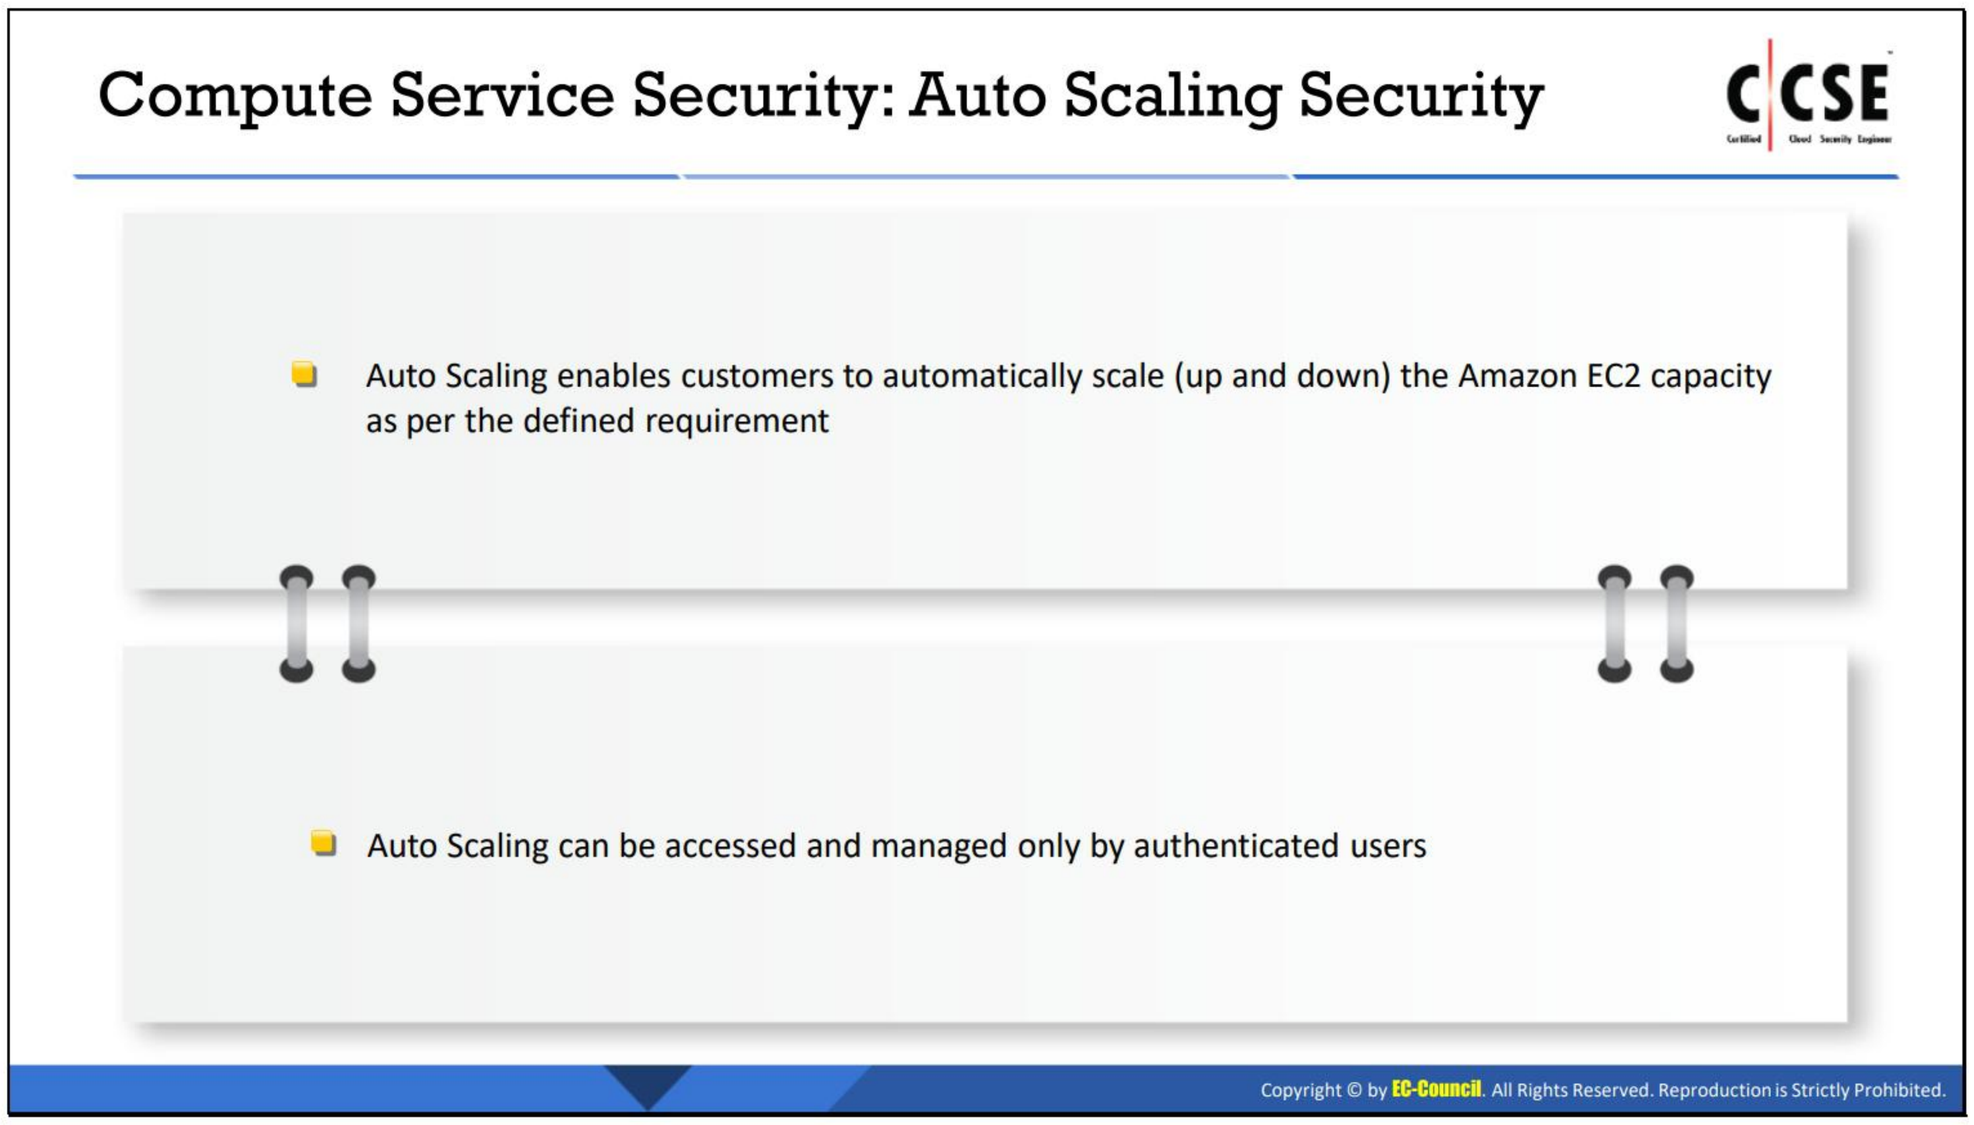
Certified Cloud Security Engineer Copyright © by **EC-Council**
All Rights Reserved. Reproduction is Strictly Prohibited.

make Amazon EC2 API calls on their behalf if they do not have access to their Secret Account Access Keys. To ensure confidentiality, API calls can be encrypted with SSL. Amazon recommends AWS customers to use SSL-protected API endpoints.

Permissions: Customers can use AWS IAM to further control what APIs can be called by users.

hide01.ir

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



Compute Service Security: Auto Scaling Security

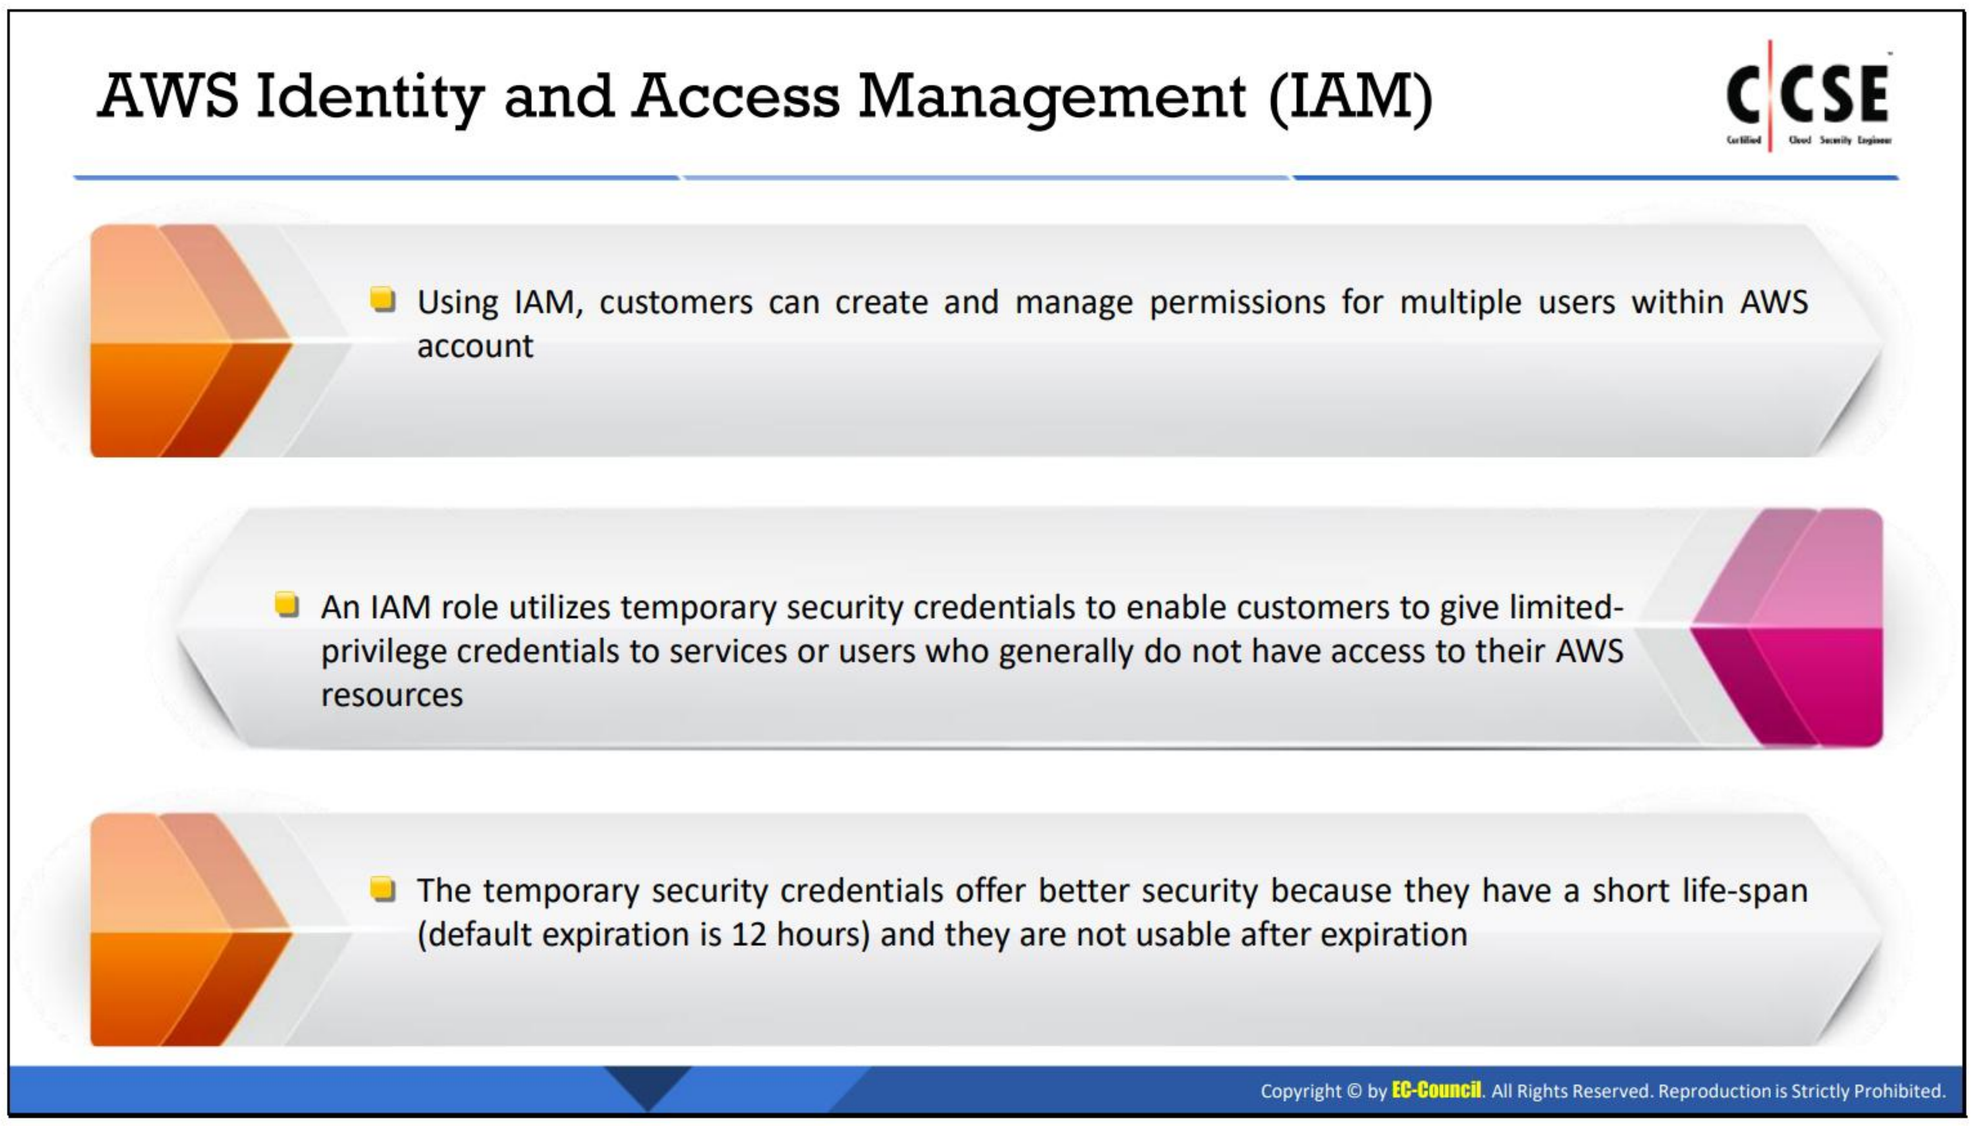
Auto Scaling enables customers to scale their Amazon EC2 capacity up and down automatically as per the specified settings. As a result, the number of EC2 instances customers use scales up seamlessly when demand spikes to maintain performance, and scales down automatically when demand lulls to save costs.

hidoo1.ir

Like all other AWS services, Auto Scaling can be accessed and managed only by authenticated users because it requires all requests made to its control API to be authenticated. All API requests are signed with an HMAC-SHA1 signature, which is calculated from the user's AWS Secret Access Key and the request. However, for large or elastically scaling fleets, getting credentials out to new EC2 instances that are launched with Auto Scaling may be challenging. Customers can use roles within IAM to simplify this process so that credentials are automatically given to any new instances launched with a role. When an Amazon EC2 instance is launched by a customer with an IAM role, temporary AWS security credentials with role-defined permissions are secure delivery to the instance and provide their customer with the EC2 instance Metadata service, which makes new security credentials before the current active credentials expire. Therefore, security credentials get rotated automatically multiple times in a day to provide enhanced security measures.

Using AWS IAM customers can control access to auto scaling they can control what Auto Scaling APIs these users are permitted to call.

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



AWS Identity and Access Management (IAM)

To help customers with the deployment and management of their applications, AWS provides several tools that enable secure access to AWS services such as deploying applications on AWS resources, creating and updating stacks of AWS resources, etc. Customers are provided with other tools that help them manage cryptographic keys by using log AWS API activity and hardware security modules (HSMs) for compliance and security.

AWS Identity and Access Management (IAM)

IAM helps customers to grant access permissions to the various users created within AWS Accounts. An identity is considered an AWS account with a unique security credential using which a customer can access AWS resources.

It enables customers to implement security best practices such as least privilege, by only granting permissions to access the AWS resources and services that are required for users to perform their duties and providing unique credentials to each user within their AWS Accounts. By default, IAM disables access to the created account unless permissions are granted explicitly.

AWS has integrated IAM through AWS Marketplace, which enables it to control and manage organizations' subscriptions to the services and software available in the Marketplace. This is an important access control feature because subscription to specific software in the Marketplace can launch an EC2 instance to start software's working. AWS Account owners can have fine-grained control over software costs and usage by controlling access to Marketplace using IAM.

IAM allows users to minimize the use of their AWS Account credentials. After a customer makes an IAM user account, all interactions with AWS resources and services must be made with IAM user security credentials.

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

Roles

An IAM role utilizes temporary security credentials to establish delegate access customers to services or to users who cannot connect with their AWS resources or granted access to it. A role can be defined as a permission to access particular AWS resources; however, these grants are not associated with a specific IAM group or user.

Temporary security credentials offer better security because they have a short life-span (default expiration is 12 hours) and they are not usable after their expiration. In certain situations, this can be helpful and useful in providing limited, controlled access.

- **Federated (non-AWS) user access**: The users or applications that do not have AWS accounts are called federated users. By using roles, customers can give them (federated users) access to their AWS resources for a limited period. This is useful when a customer has non-AWS users who can be authenticated by using an external service such as LDAP, Microsoft Active Directory, or Kerberos. In a customer's corporate identity and authorization system, identity federation between AWS and their non-AWS users is provided by the temporary AWS credentials used with the roles.

  If the customer organization supports SAML 2.0, they are allowed to generate trust among their organization (as an identity provider) and other organizations (as service providers). Customers can set up AWS as a service provider and utilize SAML to offer federated SSO (single sign-on) to the user in the AWS Management Console or acquire federated control to call AWS APIs in AWS.

  Roles can be used for to control access for mobile or web-based application. The security credentials are required by AWS resources for programmatic invocations. Besides, customers should not embed long-term security credentials in their applications because the application's users can access these credentials and it can be difficult to rotate them.

  Instead of that, customers can let their users sign-in to their (customers') applications using Login with Google, Facebook, or Amazon, and then use their (users') authentication information to assume roles and get temporary security credentials.

- **Cross-account access**: Roles can be setup by organizations who use many accounts for managing AWS resources. Using roles, they can grant permissions to one account to access resources under another account

  Using roles helps ensure that credentials are provided only as needed, temporarily for organizations that have personnel who rarely need access to resources under another account.

- **Applications running on EC2 Instances requiring AWS resource access**: If the application works on an EC2 Instance and asks for permission to make requests for AWS resources such as a DynamoDB table or Amazon S3 bucket, it should have security credentials. Instead of creating individual AWS IAM Accounts for every application on each instance, using roles customers can save a significant time.

  The temporary security credentials include an Access Key ID, a security token, and a Secret Access Key. To permit a user access to any AWS resources, the customer organization

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

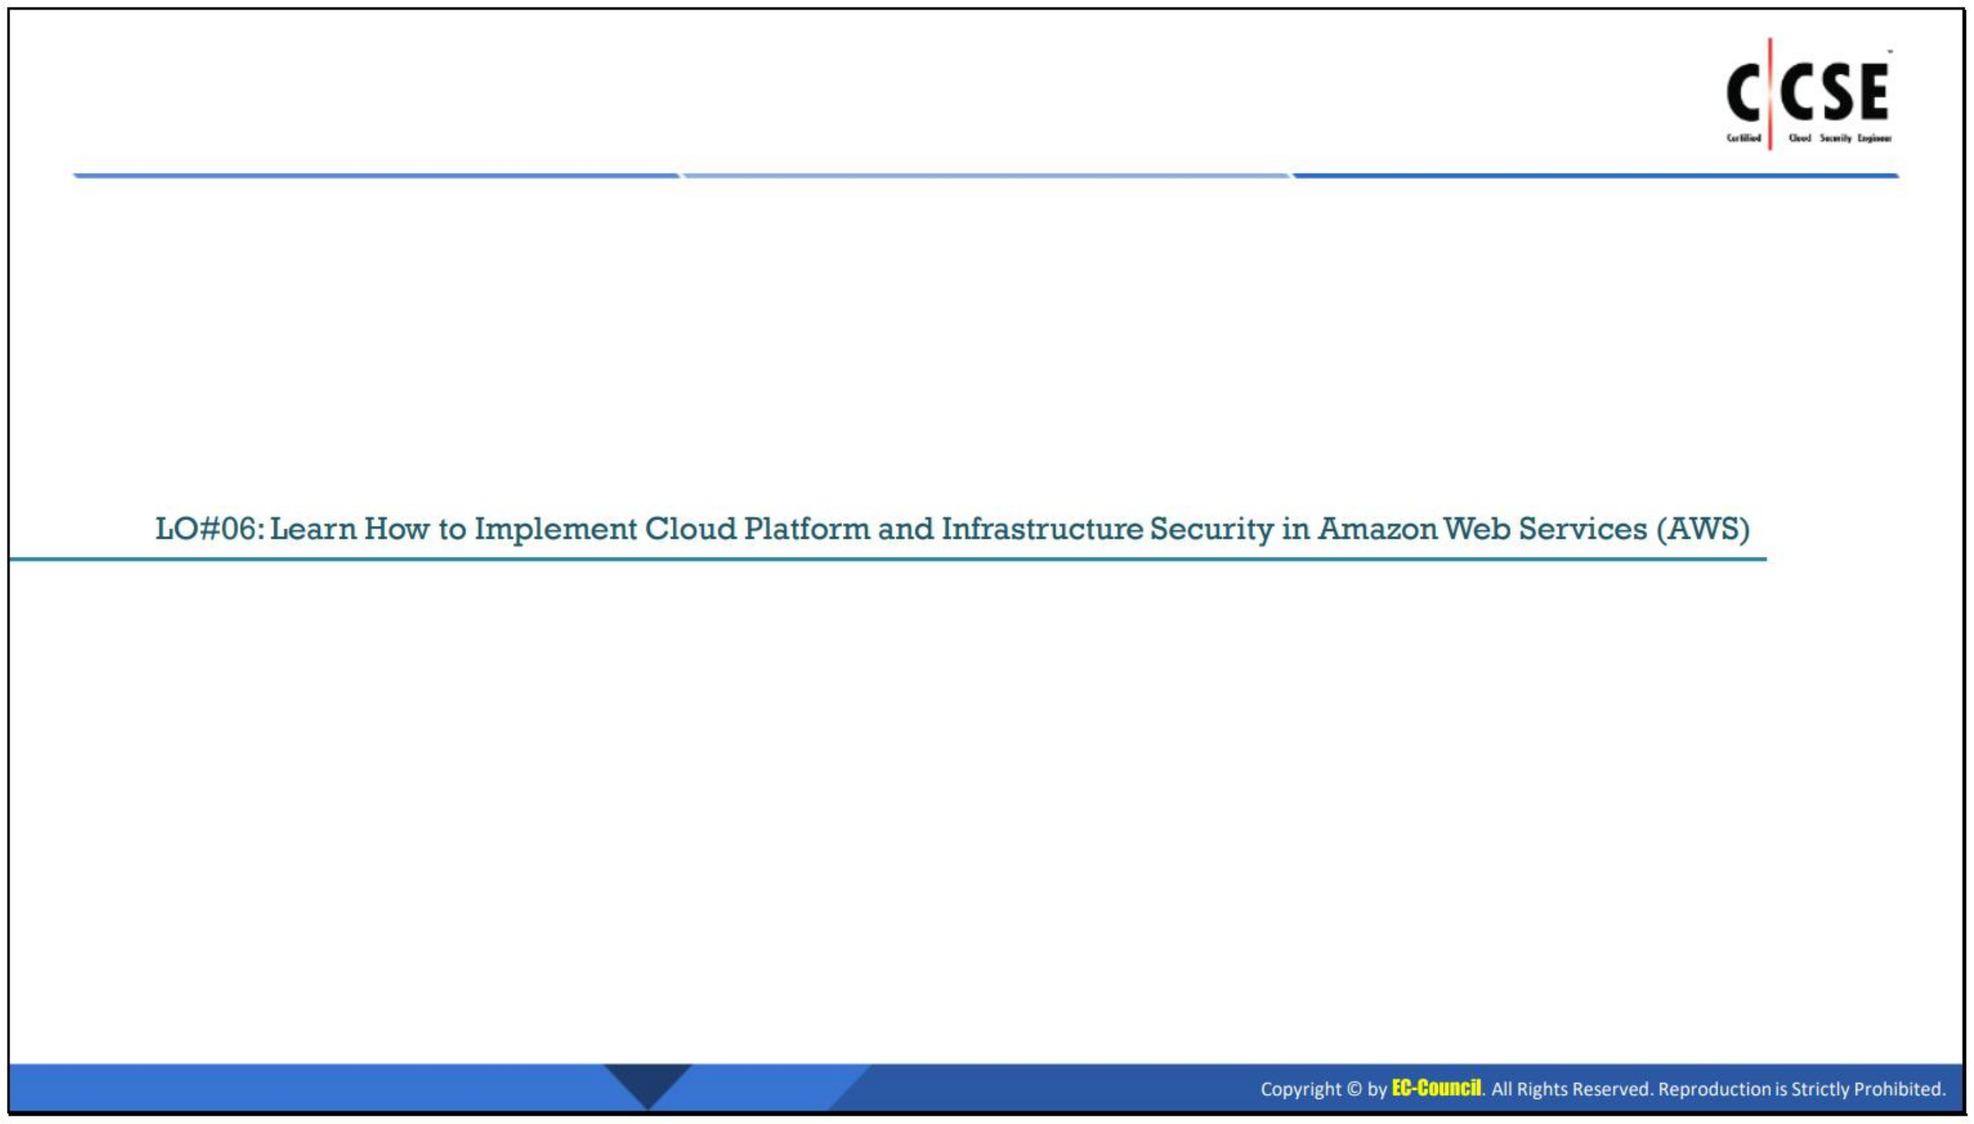
should assign temporary credentials to the user to control temporary access. The user passes in the Access Key ID, the token, and signs the request with the Secret Access Key; if the user makes calls to the organization's resources.

The token does not work with different access keys. The way the token is supplied by the customer depends on the version of the AWS product the API used.

Using temporary credentials offers additional protection for customers because they do not have to manage or provide long-term security credentials to temporary users. Additionally, the temporary credentials are loaded automatically on the target instance. To change the temporary credentials, they are set to automatic rotation, or changed multiple times in a day and securely stored by default without requiring any effort from customers.

hide01.ir

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



LO#06: Learn How to Implement Cloud Platform and Infrastructure Security in AWS

In this section, you will learn how to implement Amazon AWS services security features for isolating and protecting the AWS platform and infrastructure.

hidden1.ir

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



Create Security Group to Secure EC2 Instances

Security groups (SGs) are stateful firewalls that are associated with Amazon EC2 instances and offer protection at the protocol and port access level. SGs are a logical grouping of instances and they work like a firewall that helps in the enforcement of network traffic rules (inbound and outbound) on these instances, irrespective of the subnet where they (instances) reside. For EC2 instances, SGs provide stateful firewalls at the hypervisor level. Regularly review the rules in SGs and create different SGs for different security requirements.

SGs do not have any deny rule like network access control lists (NACLs) so they drop all packets in the absence of any rule that allows a particular packet. Secure the EC2 instances at the protocol and port access level by creating SGs because they restrict access to these instances. A single EC2 instance can have multiple SGs applied.

Enforce the principle of least privilege while designing and implementing the rules in SGs. Allow only the access that is required and never grant more access permissions than required to avoid security vulnerabilities and breaches in the future.

With every EC2 instance, one or more SGs can be attached. The administrator can use the default SGs attached to the instance or create a customized SGs as per the security requirements of that particular instance. Unlike conventional firewall, SGs provide the flexibility to specify rules to allow the movement of traffic. AWS allows creating specific SGs that suit particular applications.

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

Creating a Security Group (via AWS Management Console)

- Sign-in to the Management Console and choose the **EC2 service**. Under the **Network & Security** category, select **Security Groups**.



Figure: 2.16: Selecting Security Groups

- Click on the **Create Security Group** button and enter the name of the SGs and its description. Choose the appropriate VPC and add rules by using the **Add Rule** button (as shown below).

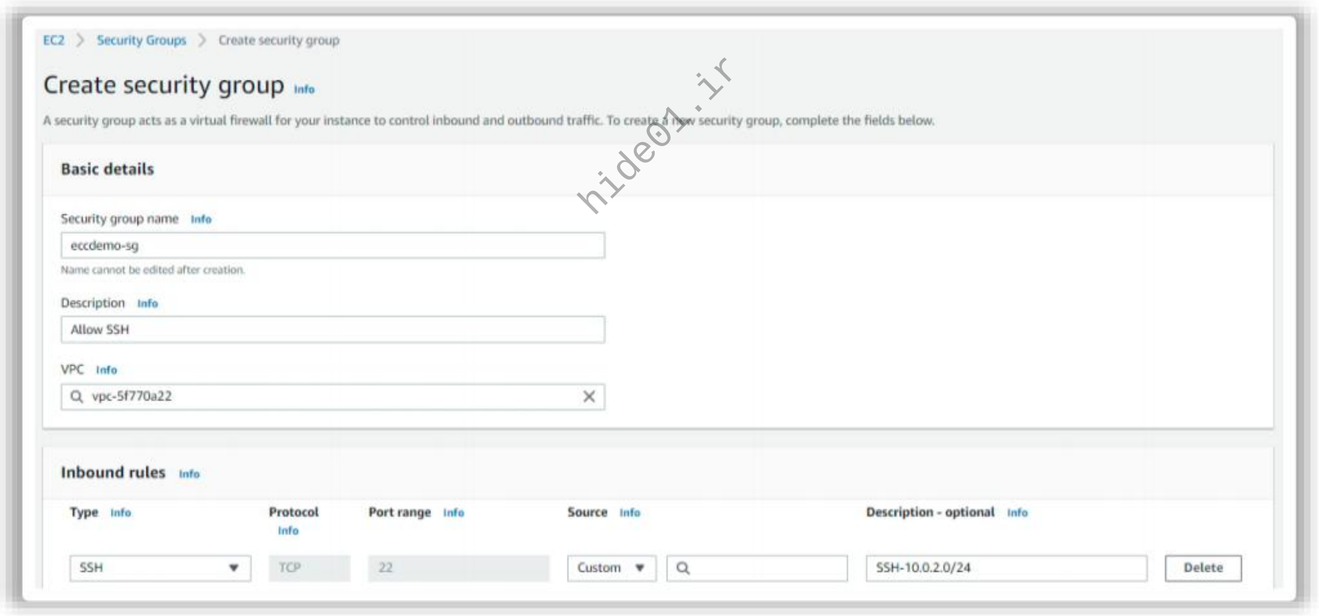


Figure: 2.17: Creating a Security Group

When EC2 instances are created, SGs are usually associated with them by default. However, SGs can be removed or added from running instances. This can be done through the Management Console by following these steps:

- Sign-in to the Management Console and choose the **EC2 service**.
- Under the **Instances** category on the left side, choose **Instances** and select the instance that has to be modified.

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

- Now, click on **Actions**, select **Networking** and then choose **Change Security Groups**.

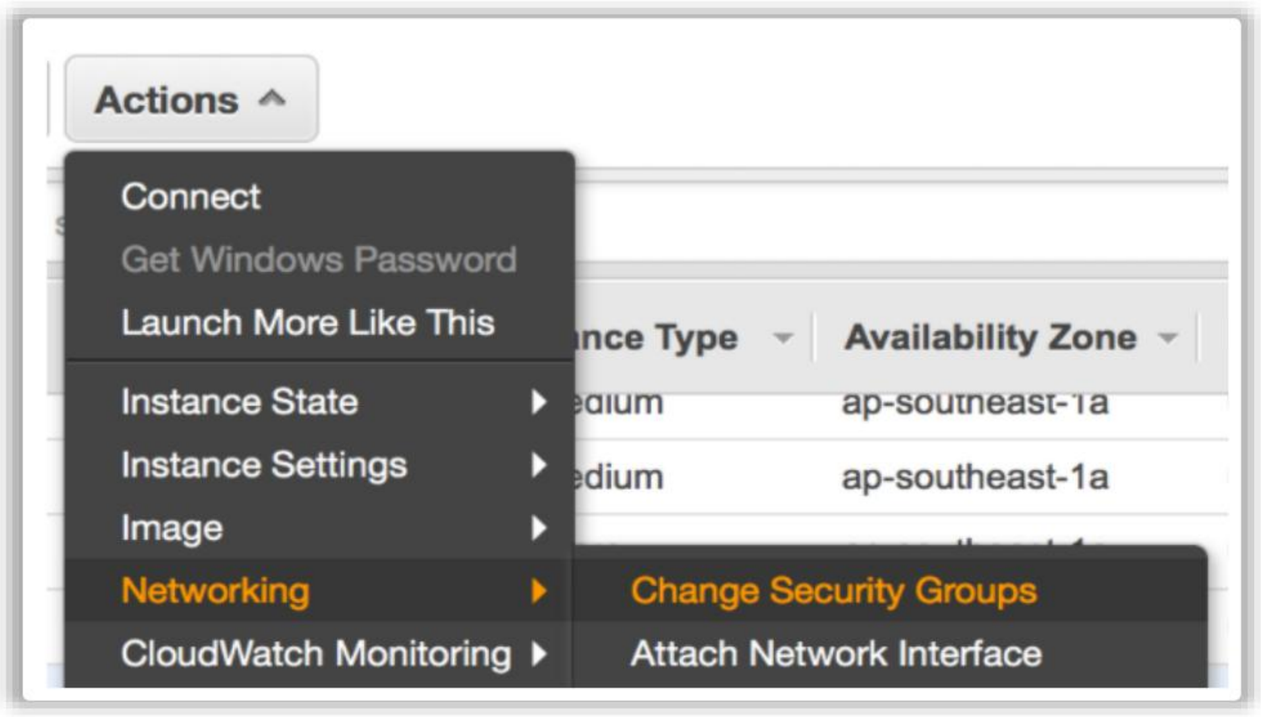
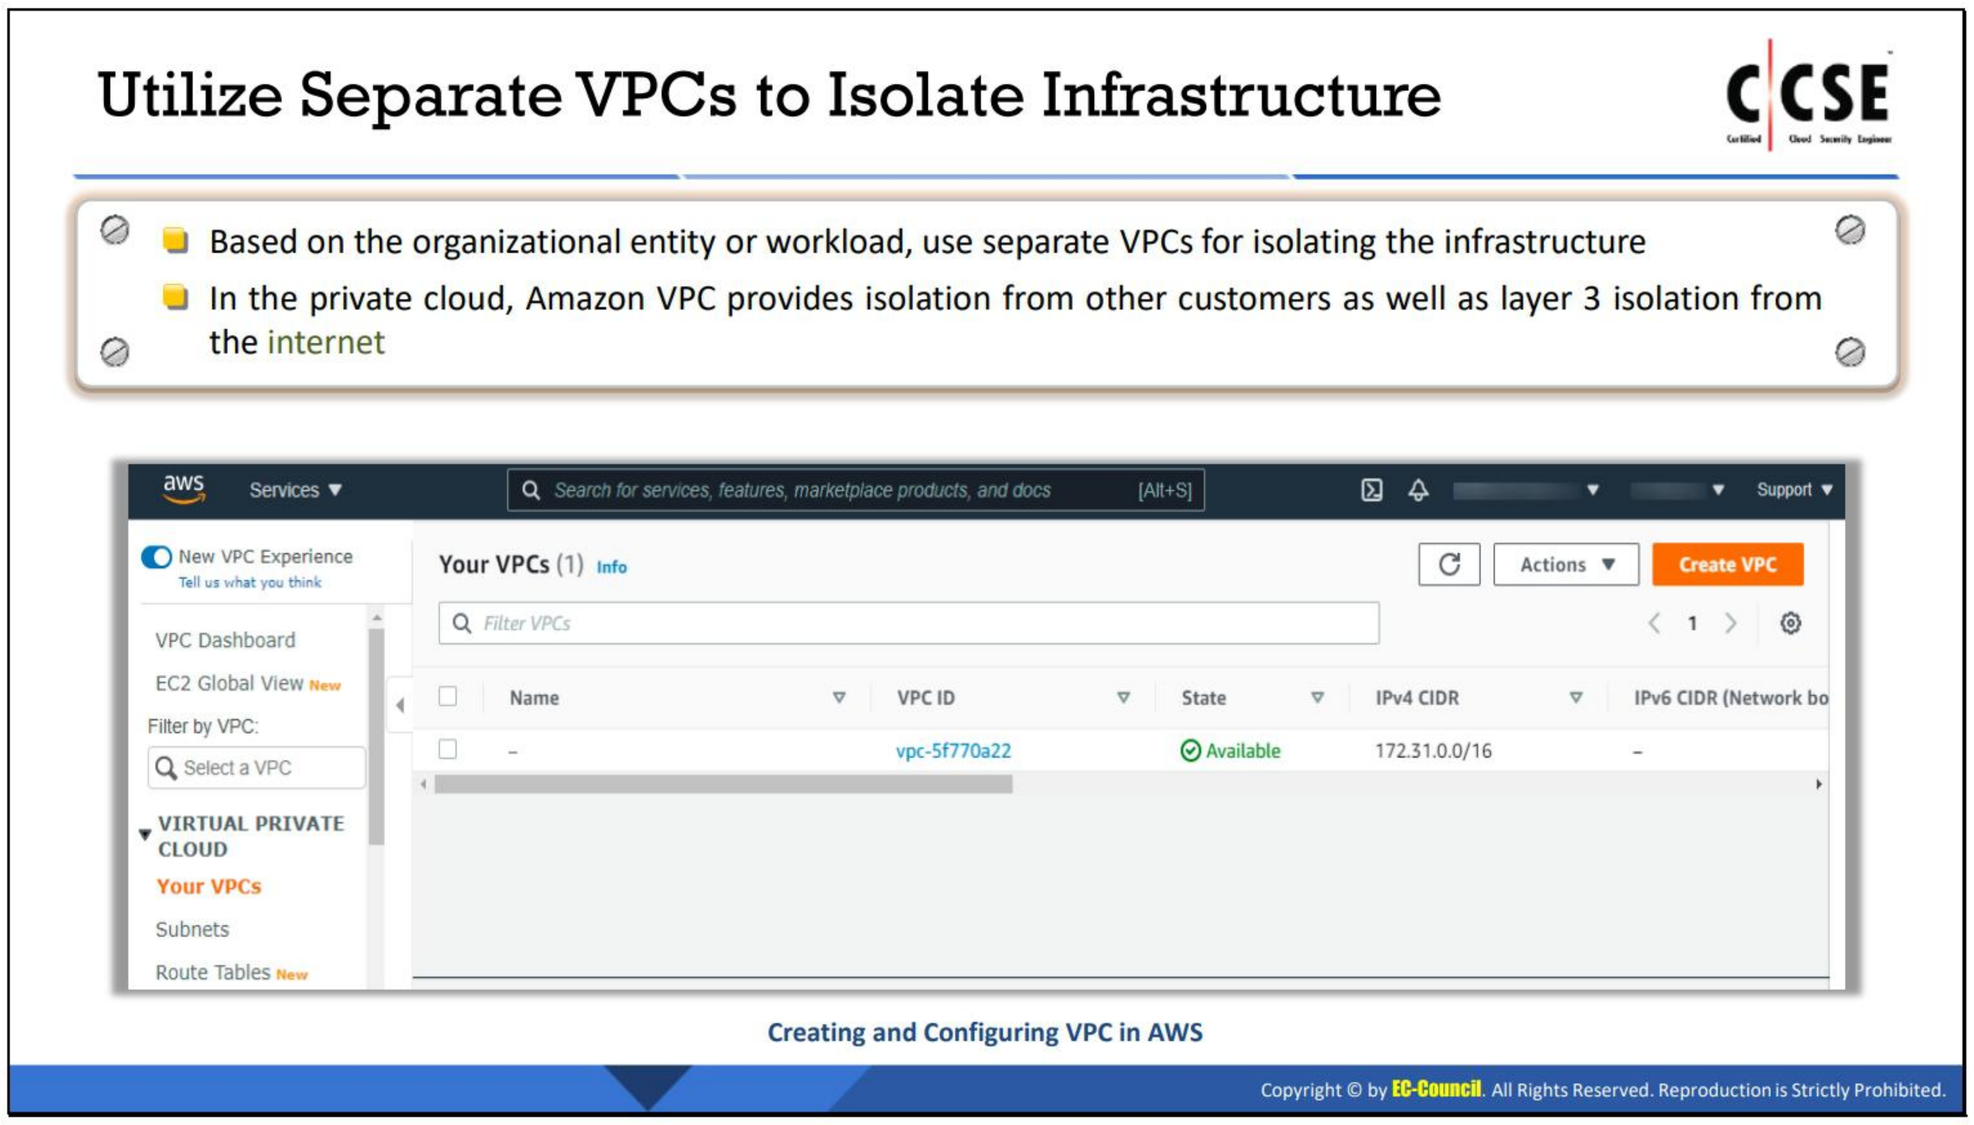


Figure: 2.18: Choosing Change Security Groups Option

hide01.ir

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



Utilize Separate VPCs to Isolate Infrastructure

In the AWS cloud, a VPC refers to a virtual network in a logically isolated area. Based on the organizational entity or workload, use separate VPCs for isolating the infrastructure. VPCs help in creating private clouds inside the AWS public cloud. VPC uses technologies such as subnets, VLANs, and VPNs to isolate computing resources within the public cloud.

In a VPC, a subnet refers to a range of IP addresses. Whenever an instance is launched, it is launched into a subnet of the VPC. To isolate the tiers of the application (e.g., web, database, and application) inside a single VPC, use subnets. If an instance should not be accessed from the Internet directly, use private subnets for the instance. In the private cloud, Amazon VPC provides isolation from other customers as well as layer-3 isolation from the Internet. Amazon VPC helps in improving security of resources and simplifying network deployment in the public cloud by provisioning a logically separate, configurable, and manageable virtual cloud network environment.

VPC allows a cloud administrator to have complete control over the virtual network environment (e.g., network creation and DHCP configuration). To improve the security of the virtual network environment, use SGs.

The logical isolation of VPC is implemented by using virtual network functions and other security features, which help a cloud administrator have granular control over which applications or IP addresses can access specific resources. As Amazon VPC provides a logically isolated network, the data of one customer does not get mixed with the other customers of the cloud provider. The cloud administrator can have complete control over how cloud resources and workloads are accessed and who can access them.

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

Running a normal instance within a public cloud exposes it to a number of potential risks such as DDoS and attacks because the instance is exposed to the Internet directly. The logical isolation provided by the VPC allows the cloud administrator to run instances securely with some specified rules and policies while accessing the Internet. The VPC serves as a wall between the instance and the Internet and allows the administrator to expose only a part of the infrastructure which is necessary to communicate with the Internet. The administrator can also set up IP ranges, subnets, and network configurations as they prefer using Amazon VPC.

VPC security comprises of ACLs and SGs. ACLs allow the administrator to define the applications or set of IP addresses that can access the VPC. SGs allow the administrator to create groups of resources and attach uniform rules to them. These groups of resources may be located in one or more subnets. SGs serve as a firewall that controls the network traffic flow to the virtual server, irrespective of its subnet location.

To create a VPC in AWS using VPC wizard, follow these steps:

- Use https://console.aws.amazon.com/vpc/ to open the Amazon VPC console.
- Click **VPC Dashboard** on the navigation pane. Locate your **VPC area** of the dashboard and click **Get started creating a VPC** if you do not have any VPC resource. If you have one, click **Start VPC wizard**.

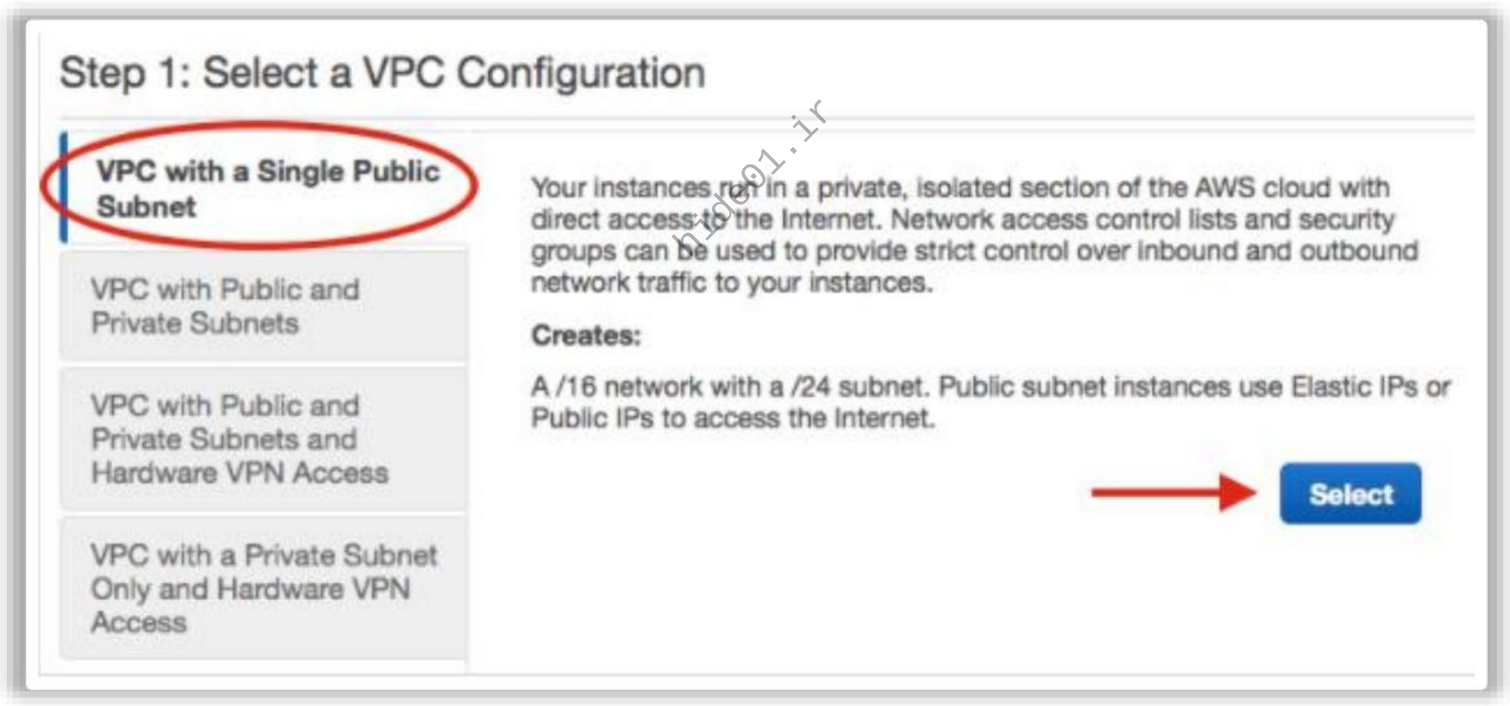


Figure: 2.19: Selecting a VPC Configuration

- Click **VPC with a Single Public Subnet** and then choose **Select**.
- Fill the following data into the wizard and then click Create VPC.
  - **IPCIDR block**: 10.0.0.0/16
  - **VPC name**: ADS VPC
  - **Public subnet**: 10.0.0.0/24
  - **Availability zone**: No preference
  - **Subnet name**: ADS Subnet 1
  - **Enable DNS hostnames**: Leave default selection

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

- **Hardware tenancy**: Default

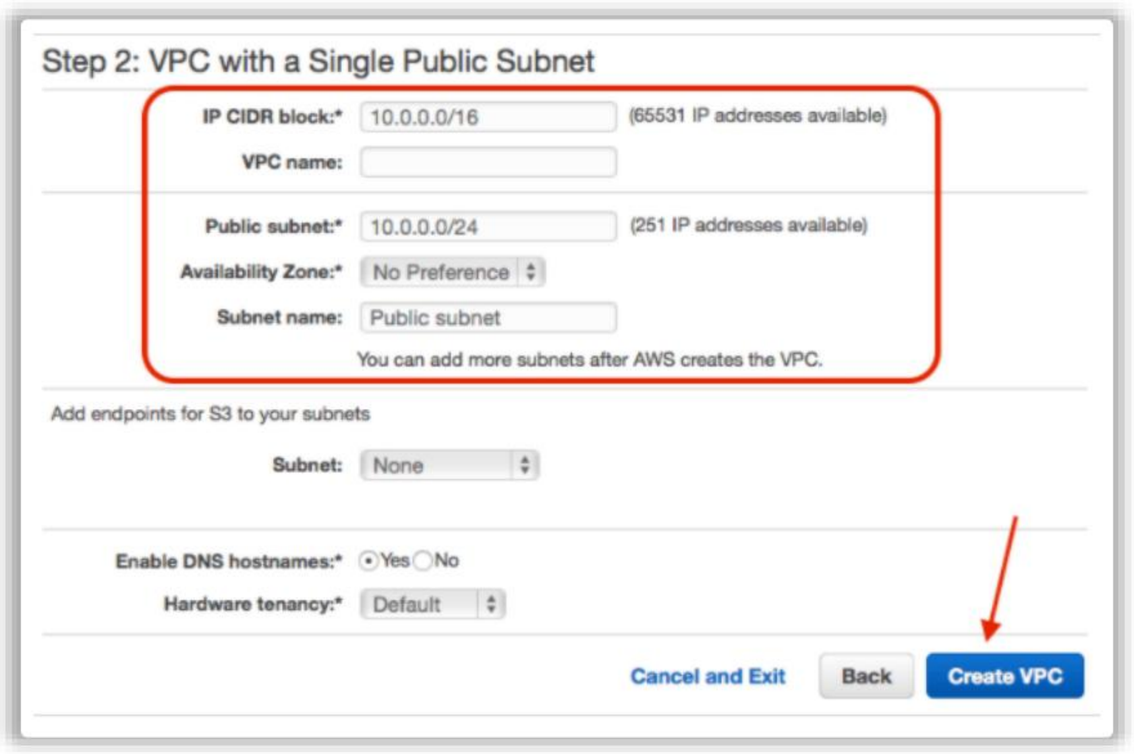


Figure: 2.20: Creating a VPC with Single Public Subnet

- Creating the VPC will take some time. After it is created, add a second subnet.

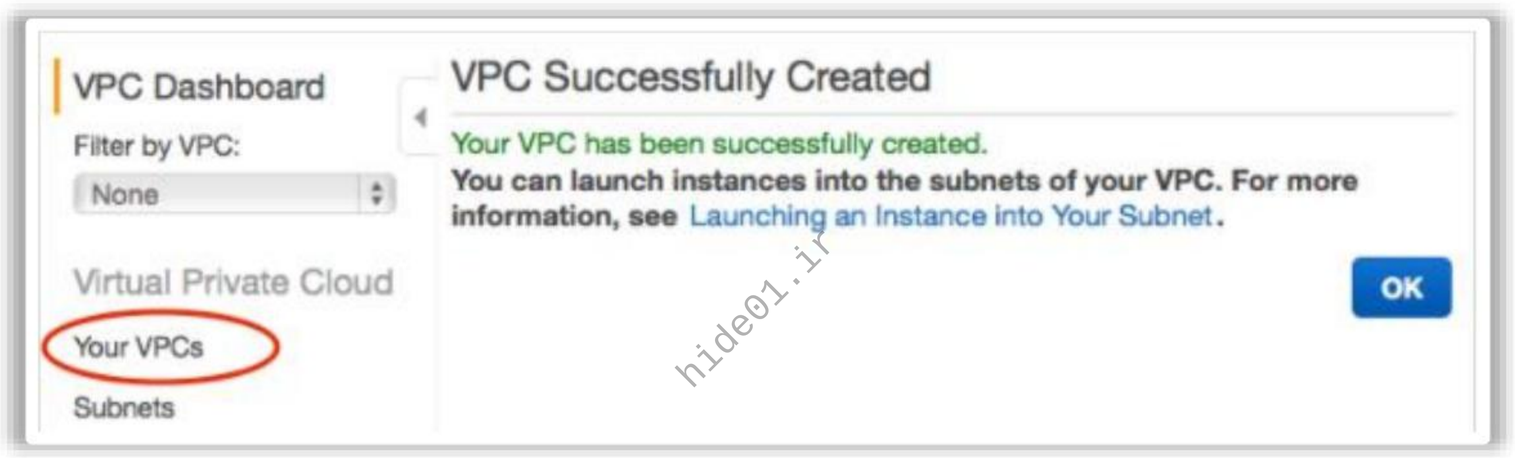


Figure: 2.21: Successful Creation of VPC

AWS Directory Services need two subnets in a VPC and each subnet should be from a different AZ. Only one subnet is created by the VPC wizard; thus, the second subnet must be created manually and should be in a different AZs.

Follow the given steps to add a second subnet to the created VPC.

- Use https://console.aws.amazon.com/vpc/ to open the Amazon VPC console.
- Select Subnets on the navigation pane and then choose the subnet (ADS Subnet 1). At the bottom of the page, select the **Summary** tab and note down the AZ of that subnet.
- Choose **Create Subnet** and fill the following information in the '**Create Subnet**' box and select **Yes, Create**.
  - **Name tag:** ADS Subnet 2
  - **VPC**: Select your VPC (ADS VPC here)
  - **Availability Zone**: Choose any AZ other than the AZ noted in the second step

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

- **CIDR block**: 10.0.1.0/24

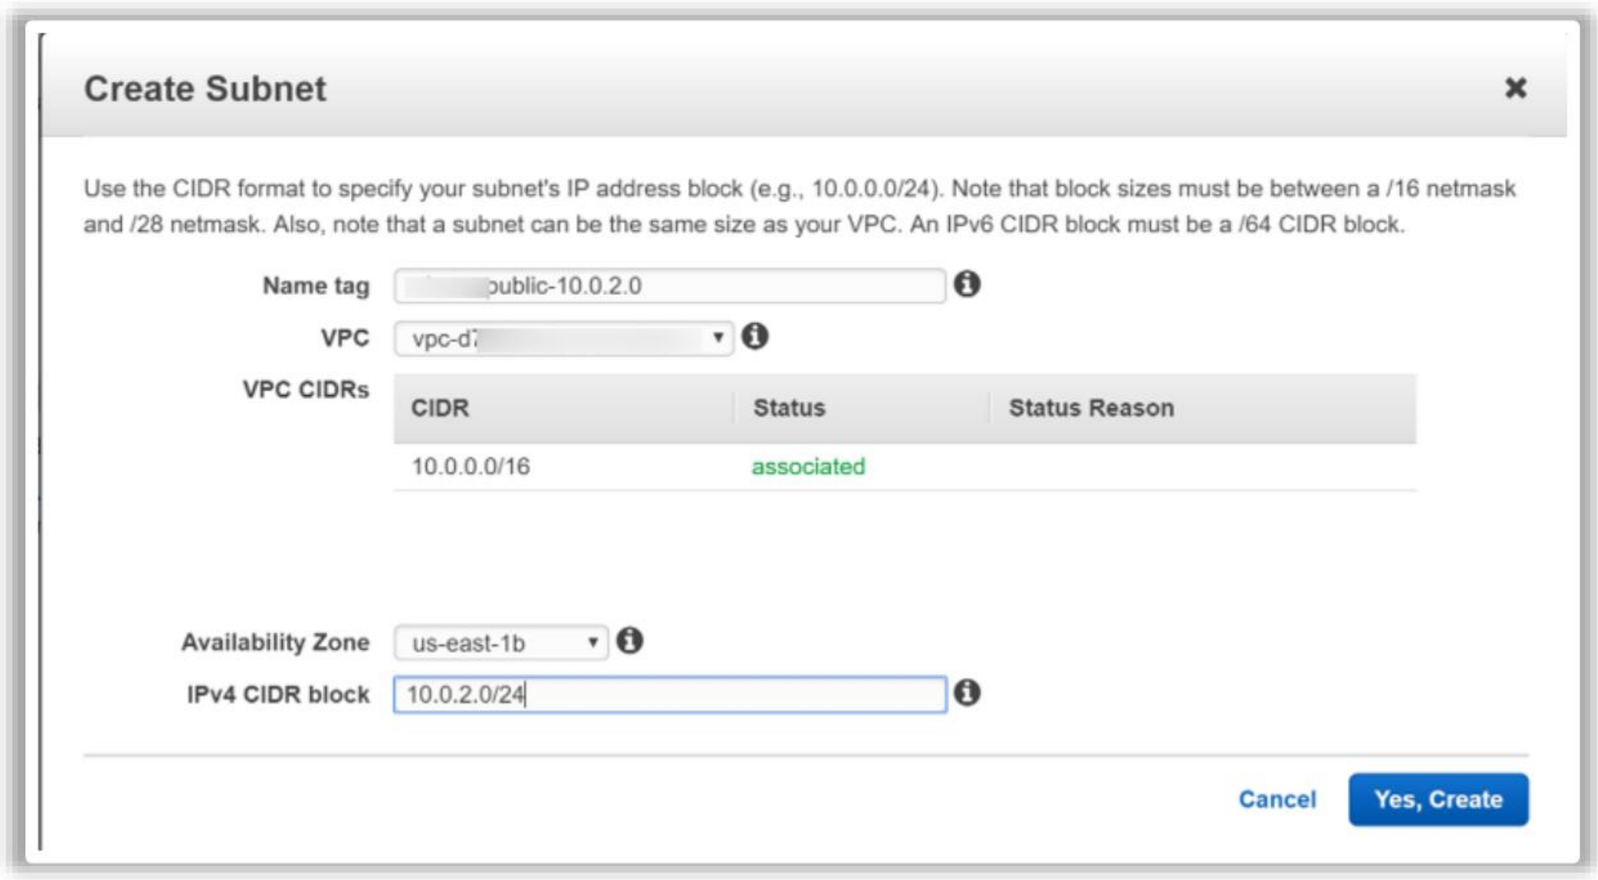
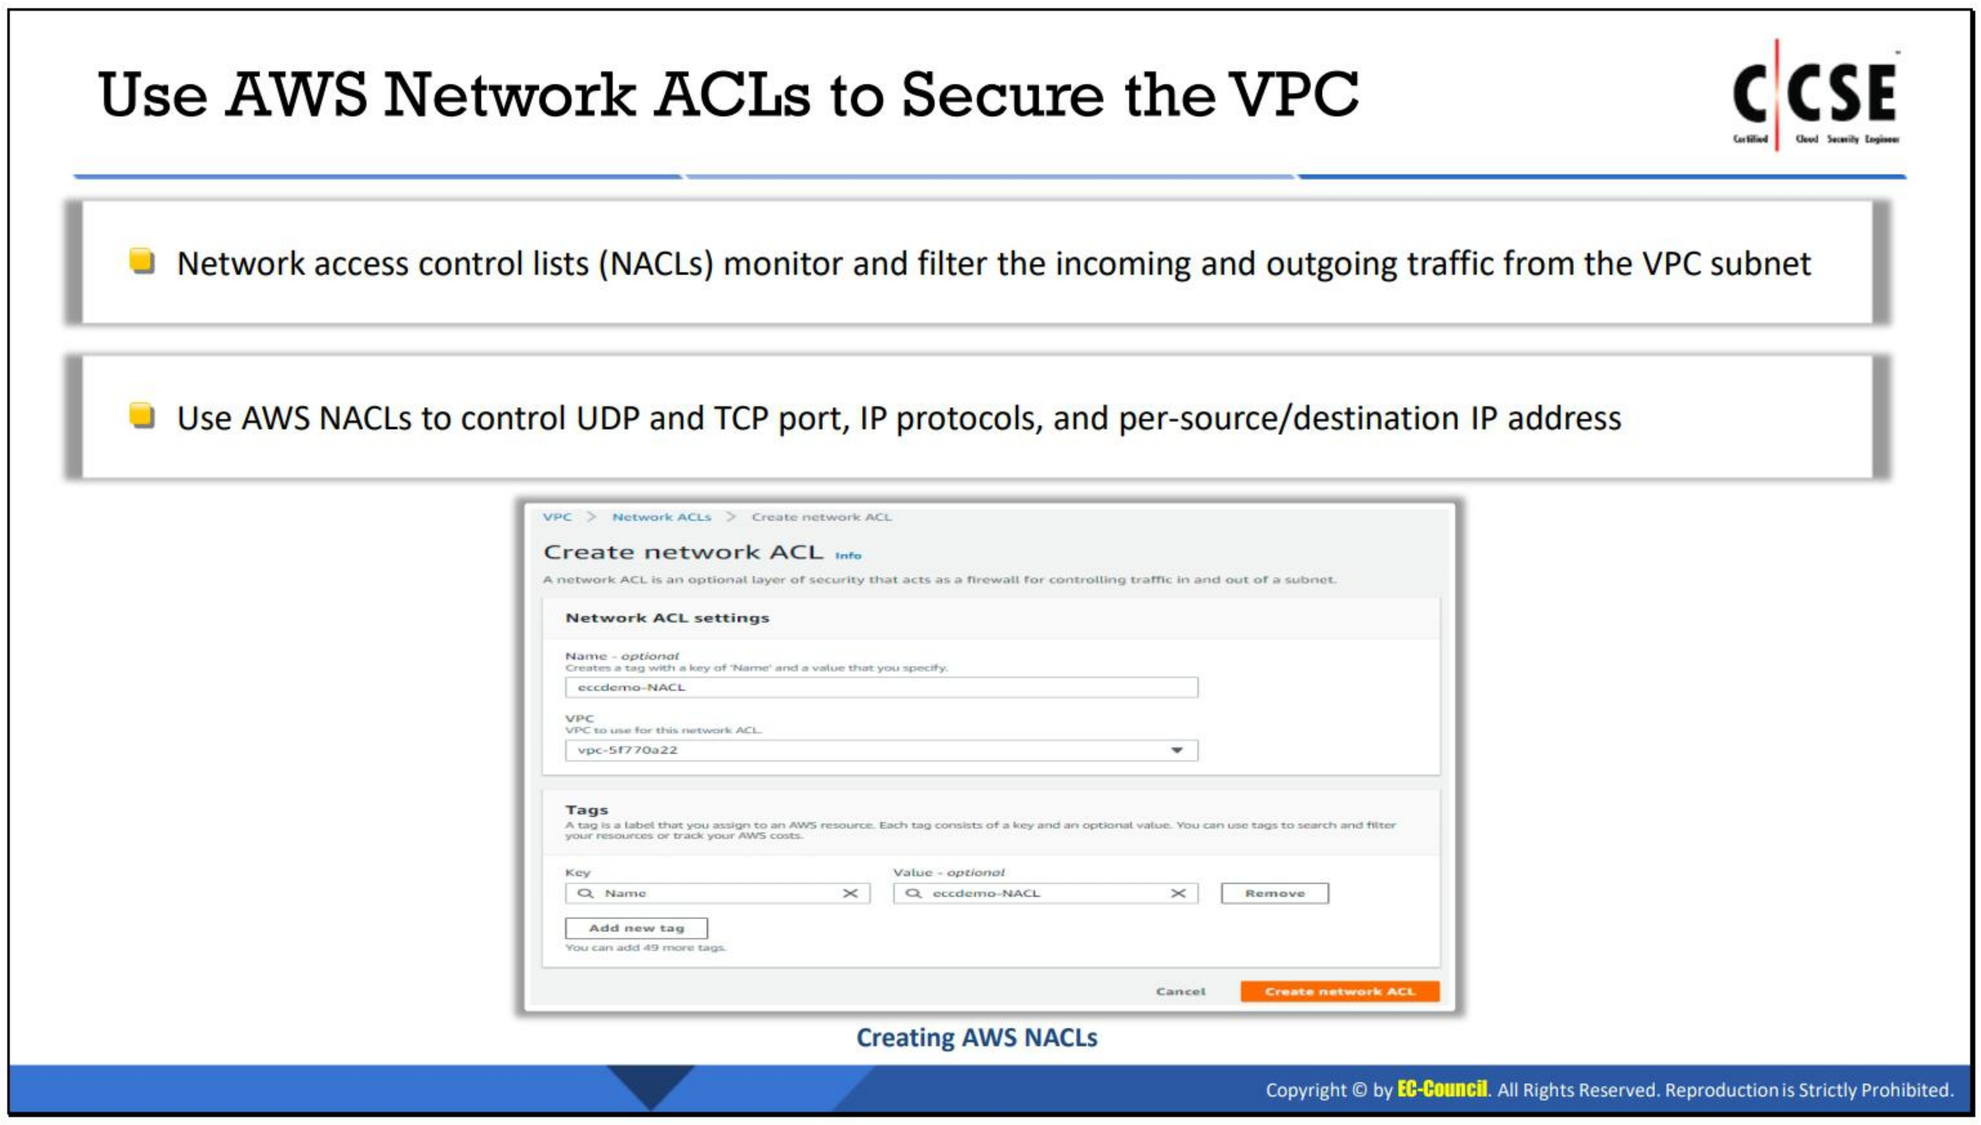


Figure: 2.22: Creating a Subnet

hide01.ir

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



Use AWS Network ACLs to Secure the VPC

Network access control lists (NACLs) monitor and filter the incoming and outgoing traffic from a VPC subnet based on a set of (inbound and outbound) rules. AWS NACLs are part of network-level security that control incoming and outgoing traffic of a VPC subnet. It is an extra layer of security for a VPC that serves as a firewall that controls traffic of one or more VPC subnets.

Unlike SGs, NACLs are stateless; thus, while creating NACL rules, the administrator needs to enforce an outbound reply rule in order to allow responses to incoming requests. Use AWS NACLs to control UDP and TCP port, IP protocols, and per-source/destination IP address. They help the administrator have granular control over IP protocols such as ICMP, IPsec, etc., and are agnostic of UDP and TCP sessions. NACLs are usually employed along with SGs and they control (allow or deny) traffic before it reaches SGs. They support both deny and allow rules whereas SGs support only allow rules.

AWS evaluates NACL rules in order, beginning from the smallest number and one rule is applied at a time. If network traffic matches a NACL rule, that rule is applied and AWS does not evaluate rules further. If the traffic does not match a NACL rule, other NACL rules are evaluated. This allows fine-grained control over traffic before it reaches the SG. NACLs provide protection at the subnet level of the VPC and SGs provide protection at the instance level. DDoS attacks can be effectively prevented by using NACLs.

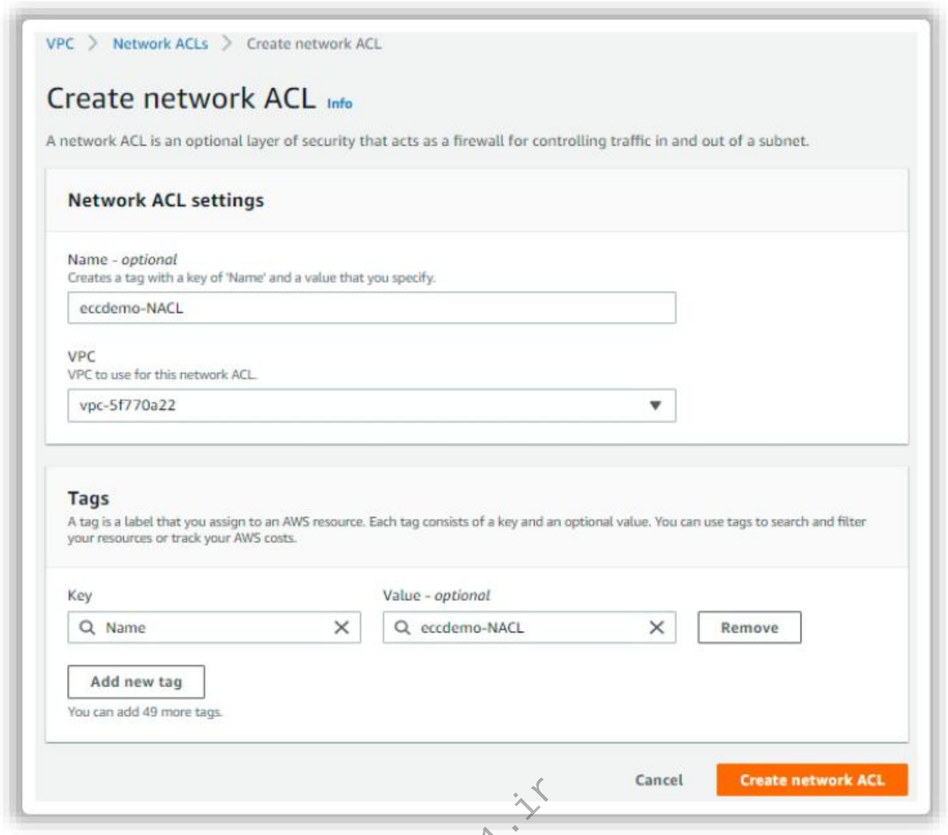
While creating a VPC, a default NACL is also created and attached with it. NACLs are similar or equivalent to SGs. NACLs help the cloud administrator to control traffic at the subnet and protocol level. Within a VPC, the administrator can attach an ACL to one or more subnets. The default NACL allows all (incoming and outgoing) traffic that go to and leave the network. The administrator can change this default NACL with their custom NACL rules.

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

Creating a NACL

To create an AWS NACL, follow these steps:

- Sign-in to the Management Console.
- Choose **VPC**, select **Network ACLs**, and then choose **Create Network ACL**. Fill the name for the ACL and choose the VPC where it must reside.
- Choose **Yes, Create** to create a custom NACL.



hide01.ir

Figure: 2.23: Creating a NACL

Associating a NACL with Subnets

After creating the NACL, it must be associated with single or multiple subnets (it is not yet associated with any subnet as shown above). To associate a NACL with a subnet, follow the steps:

- Select the **Subnet Associations** tab and choose **Edit**.
- Choose the subnet(s) you want it to be associated with and click **Save**.
- The selected subnet(s) will use the NACL for all (incoming and outgoing) traffic.

Creating Custom Rules

Until custom rules are created, the NACL uses the default rule “Deny”, which denies all the traffic (both inbound or outbound). Unlike the default NACL, which allows all traffic, custom NACL blocks all traffic by default.

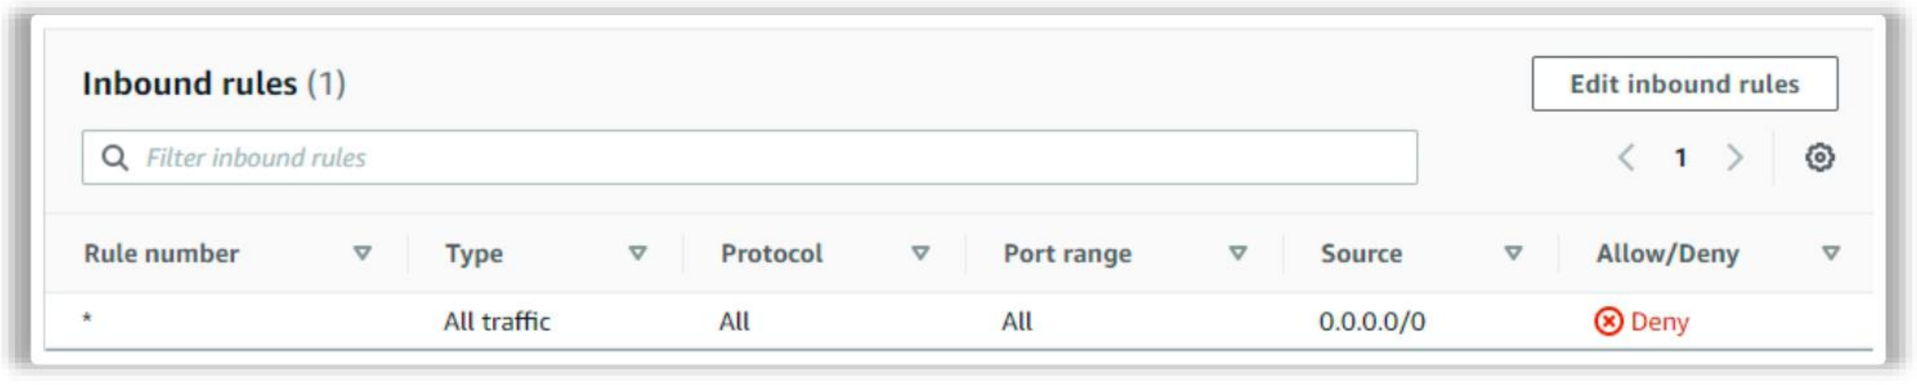


Figure: 2.24: Traffic Blocked by NACL

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

Click on the appropriate tab to configure the ACL's rules and then click **Edit.**



Figure: 2.25: Configuring ACL Rules

The image shows some fields that are to be selected to create custom rules.

Rule#: Rules are read or evaluated in ascending order until a rule matches the packet. Number your rules properly while creating the rule base.

Type: This field has a drop-down list that allows the administrator to select protocol types such as RDP, SSH, POP3, etc. The administrator can specify custom protocols like varieties of ICMP.

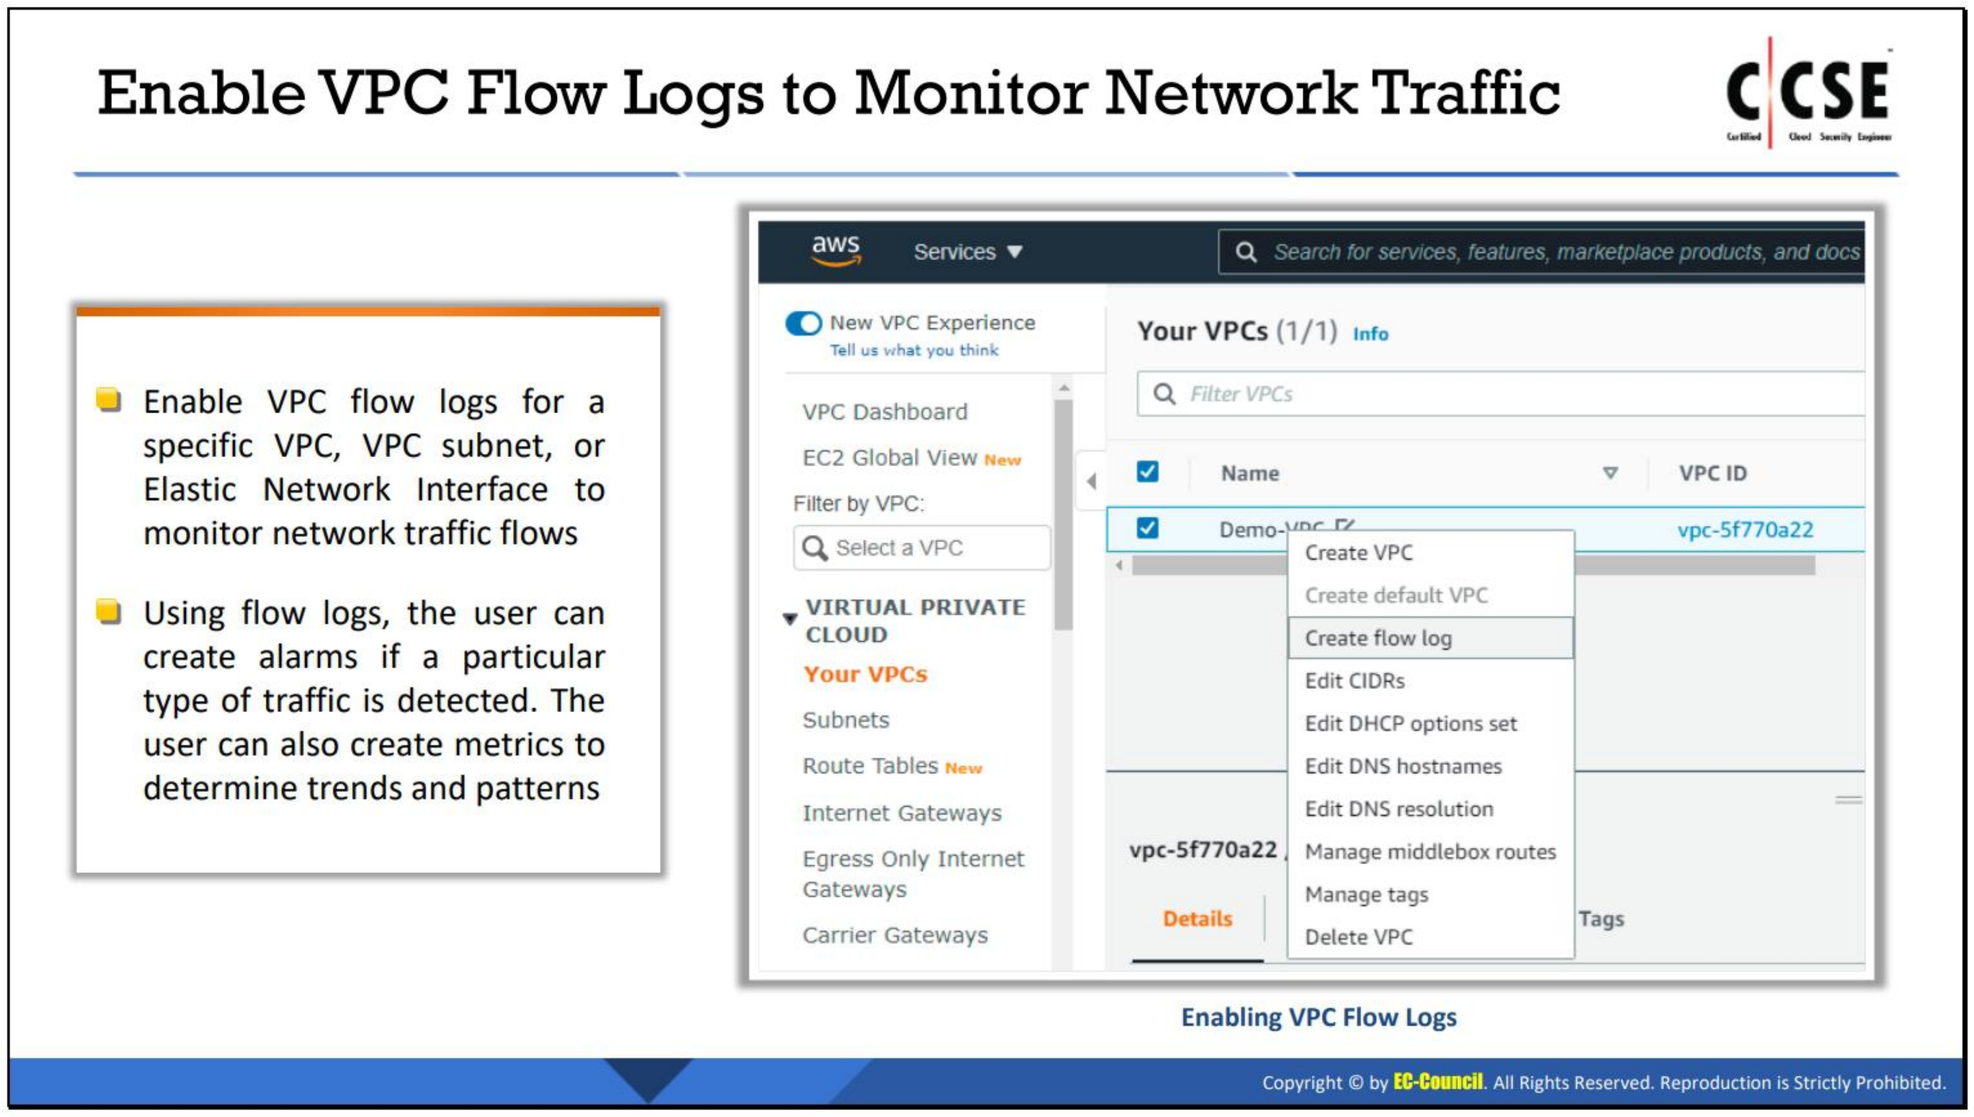
Protocol: The protocol option is highlighted based on the selection for "Type". The administrator must provide a value for custom rules such as TCP/UDP.

Source: The administrator can put a network subnet range or a particular IP address in this field or leave it as such to allow traffic from anywhere.

Allow/Deny: This field helps the administrator to specify if the defined traffic is allowed or blocked.

hide01.ir

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



Enabling VPC Flow Logs

Enable VPC Flow Logs to Monitor Network Traffic

Amazon VPC Flow Logs are logs that provide information about IP traffic that flows between network interfaces of AWS resources inside a VPC. In a VPC, there might be several hundreds of resources that communicate with resources present in different subnets in a VPC and also between different VPCs (via VPC peering connection). The data of VPC Flow Logs helps the cloud administrator in troubleshooting incidents related to traffic flow and network communication as well as in resolving security issues. It helps to identify traffic that is directed toward a prohibited destination.

Enable VPC Flow Logs for a specific VPC, VPC subnet, or Elastic Network Interface to monitor traffic flows. The administrator can use VPC flow logs to rectify connectivity and security issues and ensure that network access rules are working as intended. These logs provide better visibility by allowing the administrator to record complete information of the IP traffic that is going to and coming from network interfaces present in the VPC. The data of these logs is not stored in Amazon S3, but logged and stored in CloudWatch logs, where these logs are analyzed.

Using Flow Logs, the administrator or a user can create alarms if a particular type of traffic is detected. They can also create metrics to determine trends and patterns. The log data contains information about traffic that is allowed and denied as per the NACL and SG rules. Other information present in these logs includes ports, source and destination IP addresses, packet and byte counts, the IANA protocol number, an action (accept or reject), etc.

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

Enabling VPC Flow Logs (via Management Console)

Follow the steps mentioned below to enable VPC Flow Logs.

- Open the **VPC dashboard** and choose **Your VPC.**

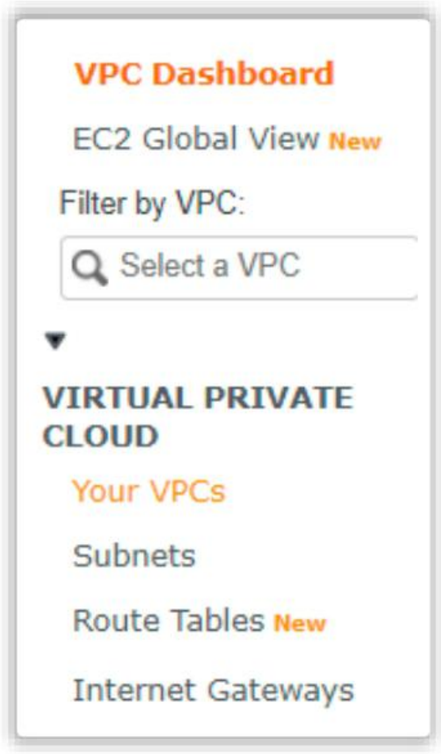


Figure: 2.26: VPC Dashboard

- Choose the **VPC** and then select the **Flow Logs** tab.

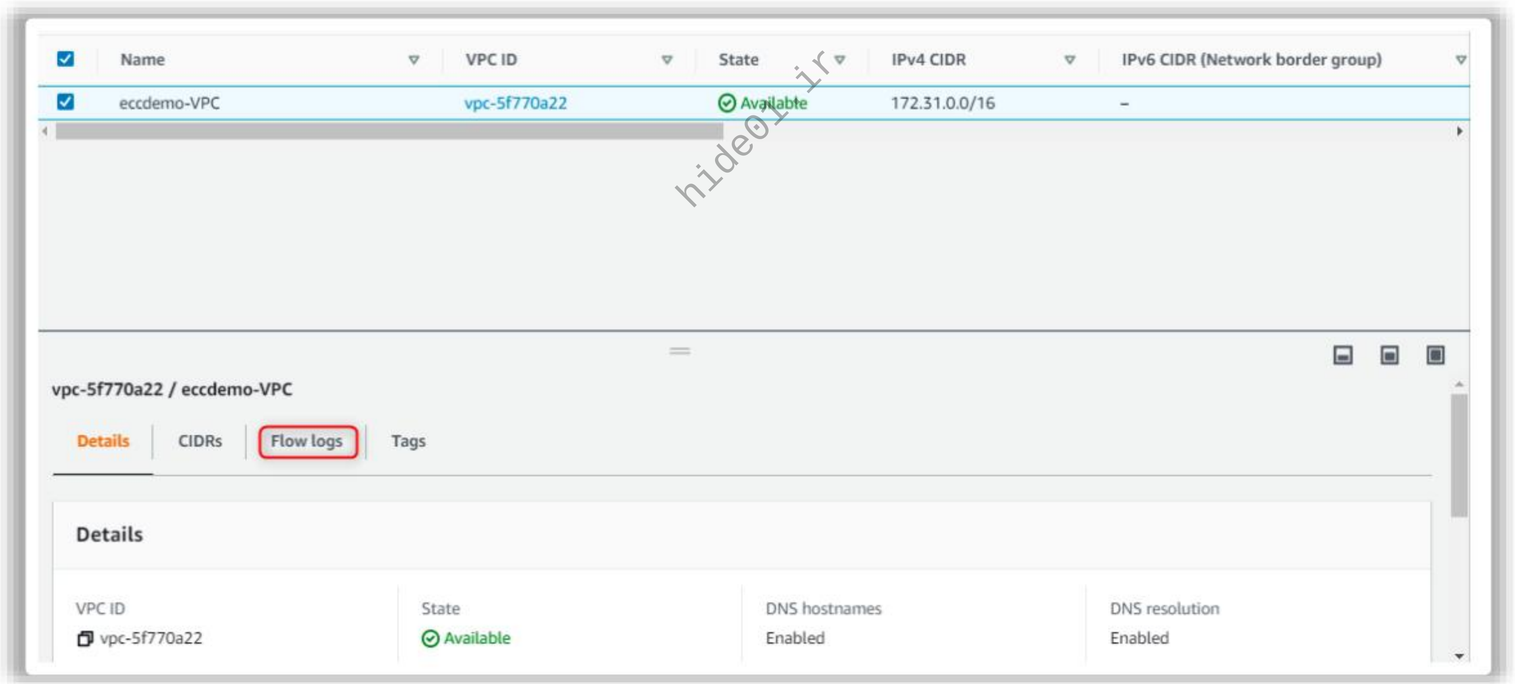


Figure: 2.27: Selecting VPC Flow Logs

- Now, click on Create Flow log.

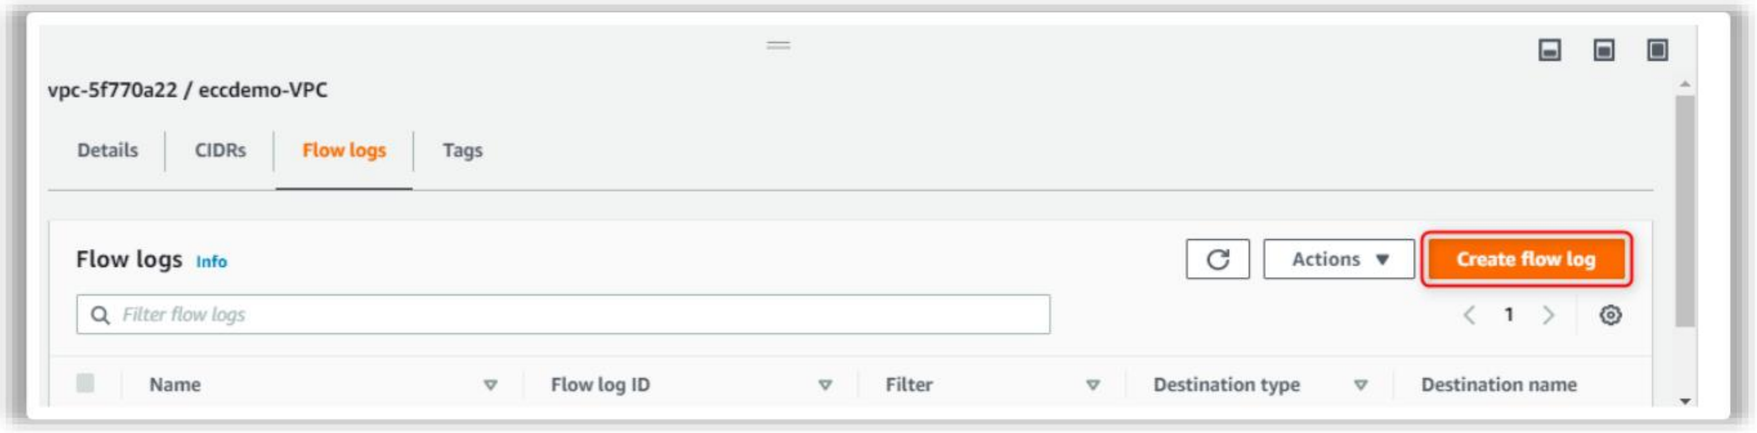


Figure: 2.28: Creating Flow Logs

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

- In the filter field, select All and choose the destination as CloudWatch.

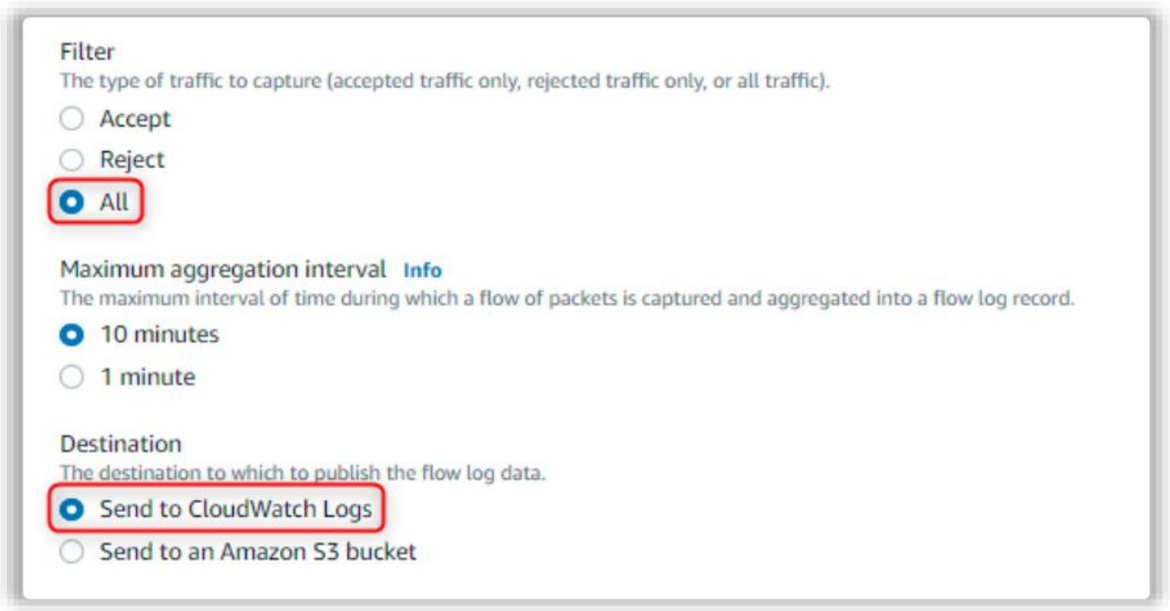
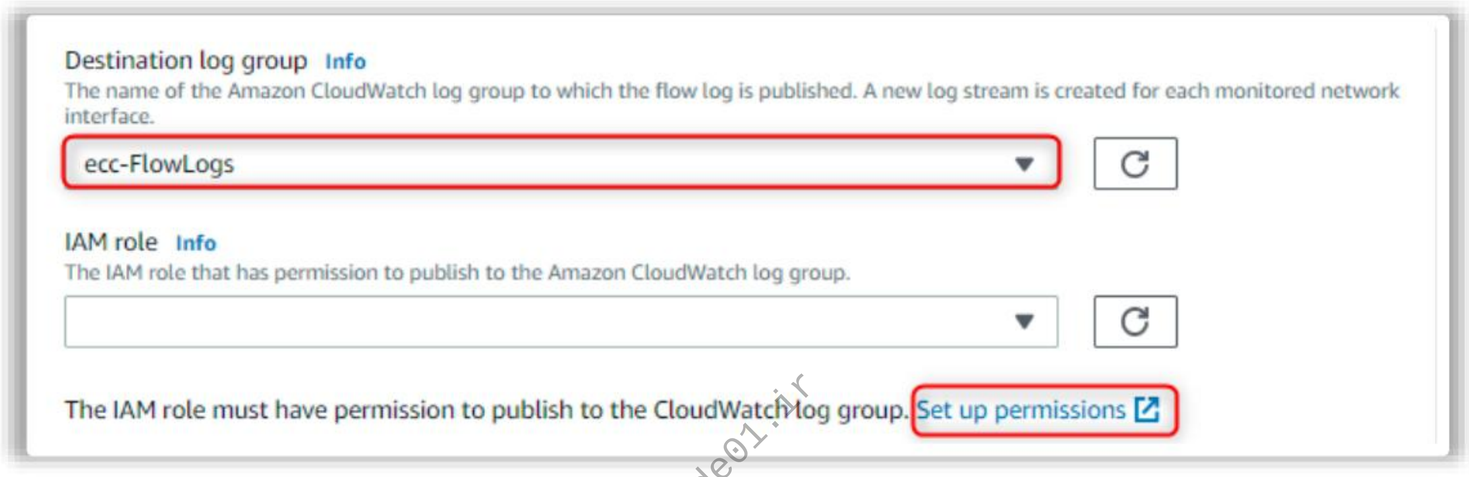


Figure: 2.29: Choosing Filter and Destination

- Select the destination group name and then click on **Set up permission**.



hide01.ir

Figure: 2.30: Entering a Destination Group Name

- Leave the default settings and click on **Allow**

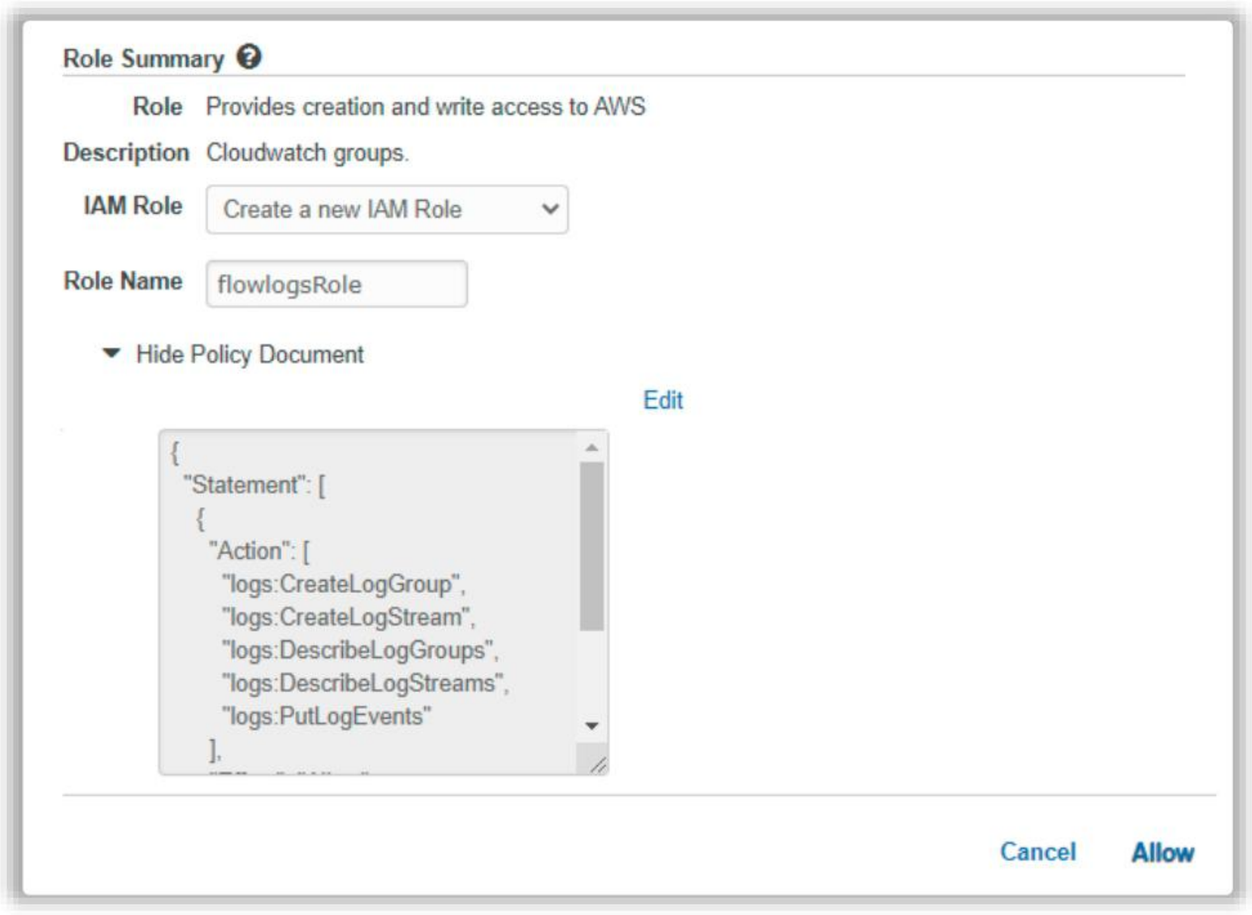


Figure: 2.31: Role Summary

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

- Choose the **flowlogsRole** from the drop-down list of **IAM Role** and click on **Create.**

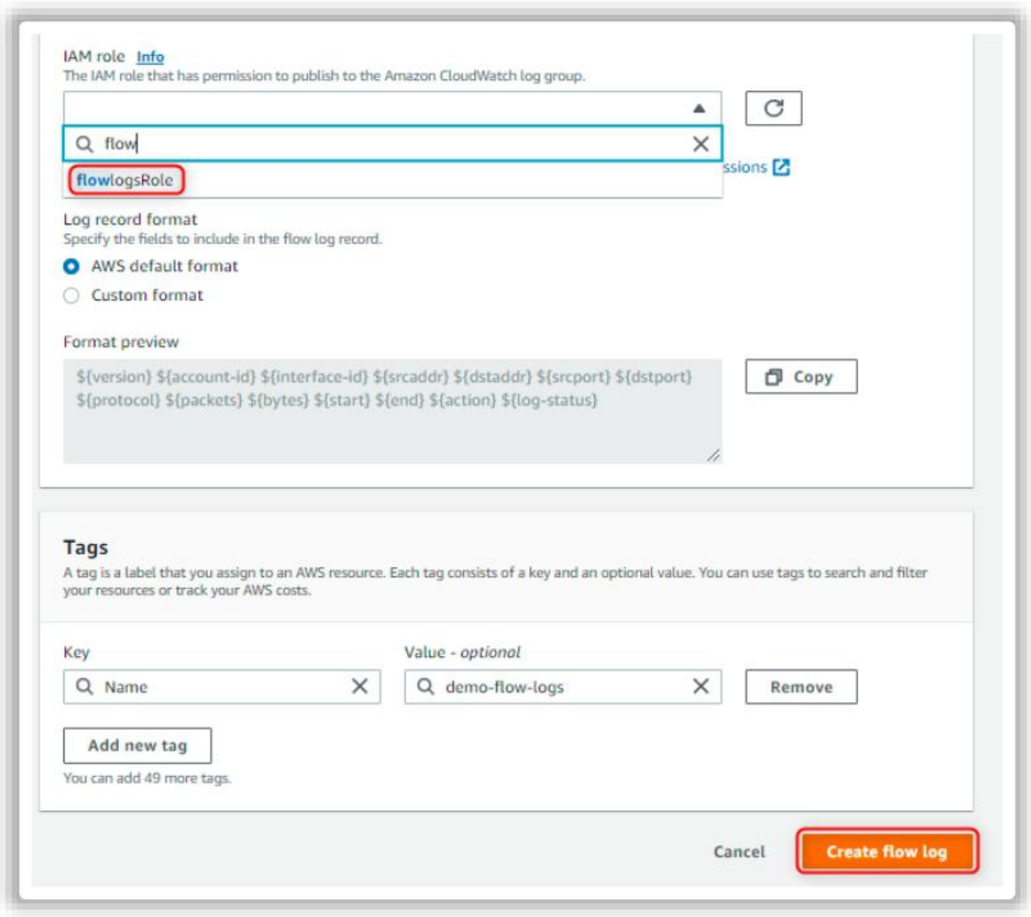


Figure: 2.32: Selecting flowlogsRole from IAM Roles

- Now, the flow logs have been created successfully.

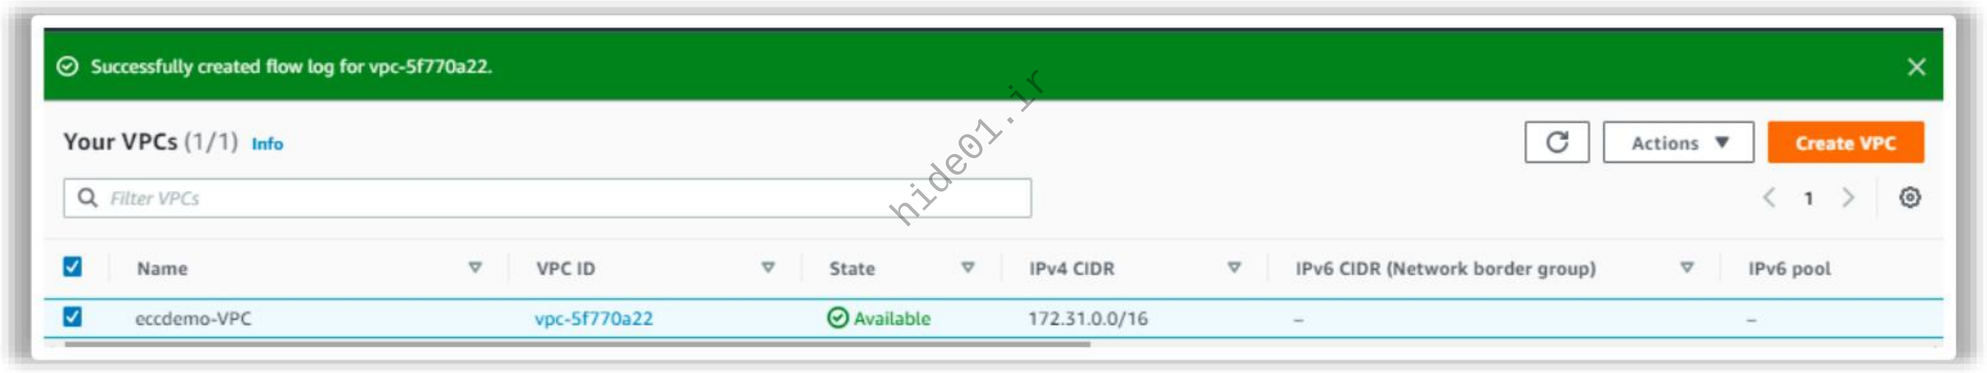


Figure: 2.33: Message Showing Successful Creation of Flow Logs

- The administrator can view the flow logs in **CloudWatch** logs. Go to **CloudWatch** and then select **logs** to see the flow logs.

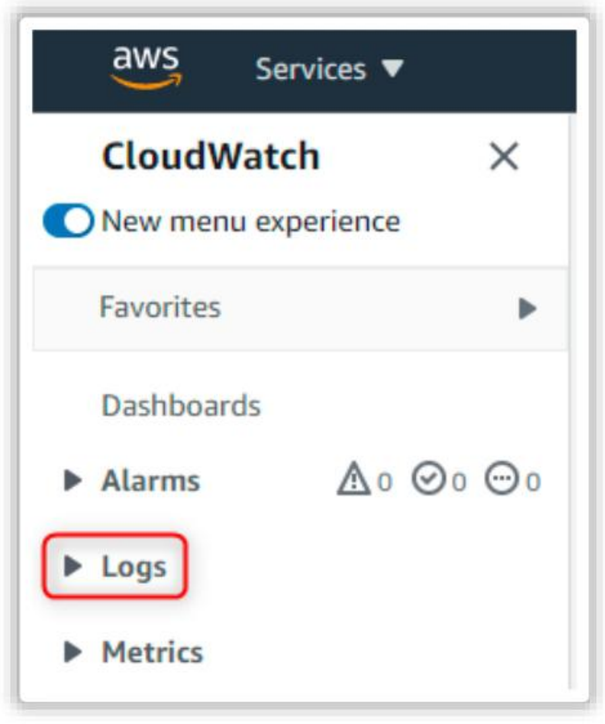


Figure: 2.34: Selecting Logs from CloudWatch

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

- The VPC log group can be seen in CloudWatch.

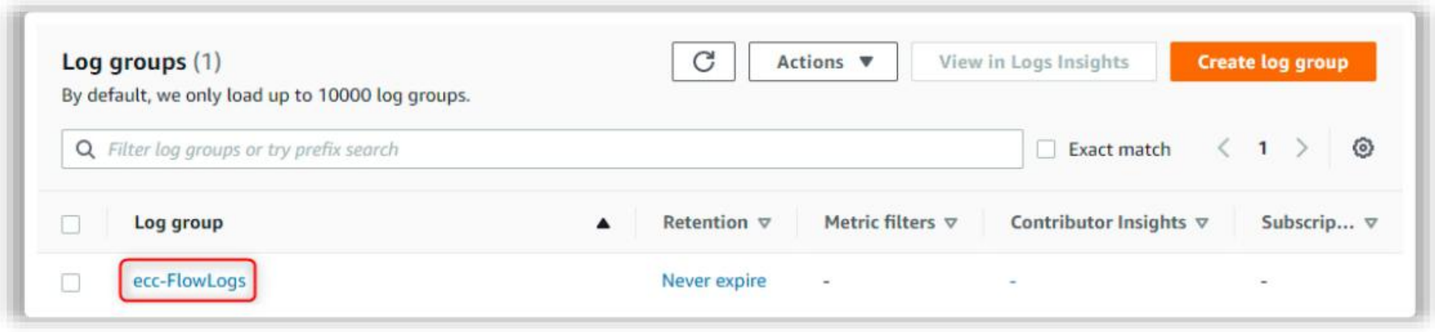


Figure: 2.35: VPC Log Group from CloudWatch

- Now, view the VPC flow logs.

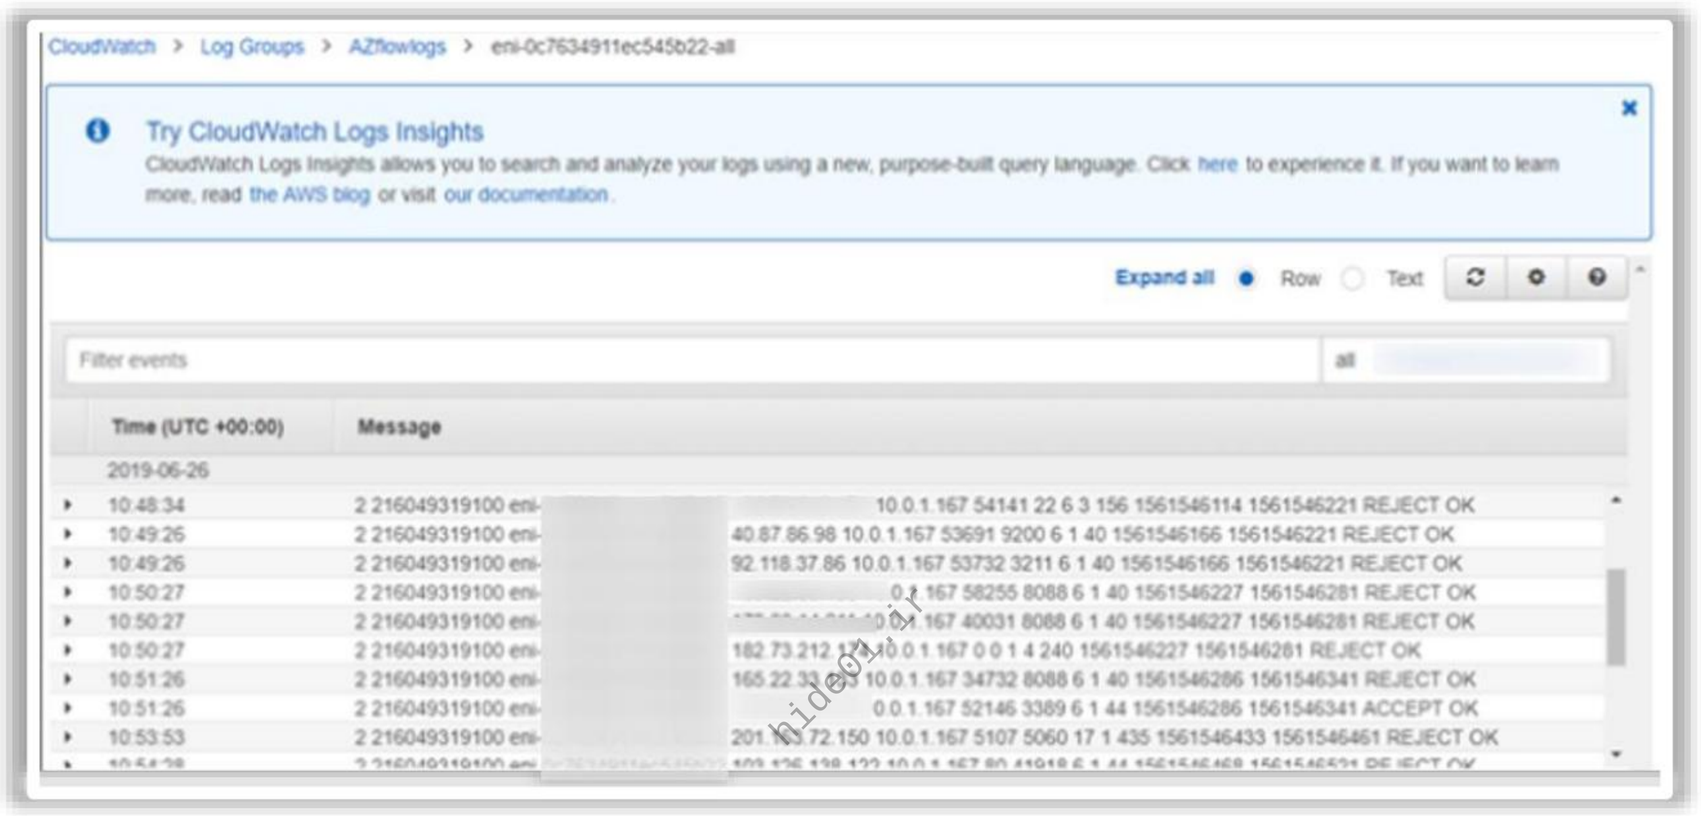
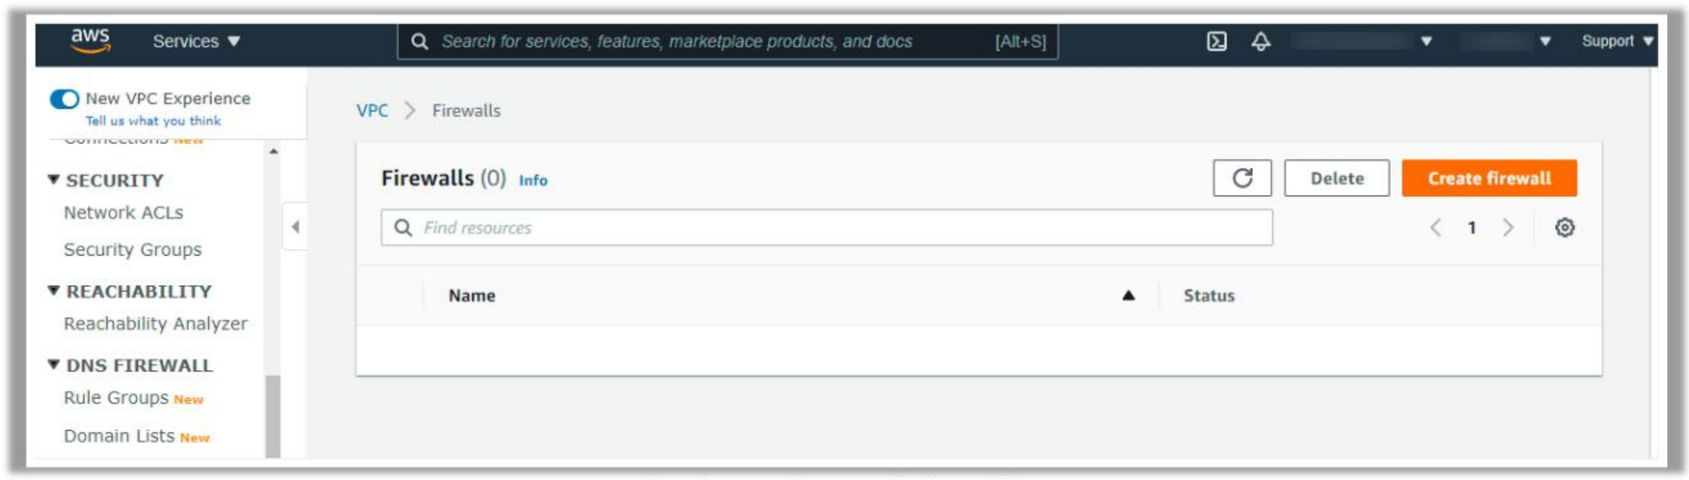


Figure: 2.36: VPC Flow Logs

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

Secure VPC by Deploying AWS Network Firewall

- Using the AWS network firewall, the admin can deploy and manage stateful inspection, perform intrusion detection and prevention, and web filtering to secure the virtual network
- AWS network firewall allows the admin to enforce customized rules that will stop the VPCs from accessing domains that are not authorized, block specific IP addresses, and detect malicious activity



Creating AWS Network Firewall

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

Secure VPC by Deploying AWS Network Firewall

AWS Network Firewall is a managed, stateful, network firewall that provides intrusion detection and prevention service for VPCs. It allows filtering traffic at the perimeter of a VPC, including filtering traffic that passes through an internet gateway, NAT gateway, over AWS Direct Connect or VPN. It utilizes the open-source intrusion prevention system, Suricata, to perform stateful inspection.

Using AWS network firewall, the administrator can deploy and manage stateful inspection, perform intrusion detection and prevention, and perform web filtering to secure the virtual network. To ensure high availability without any additional cost, AWS Network firewall scales automatically with traffic. AWS network firewall allows the administrator to enforce customized rules that will stop VPCs from accessing domains that are not authorized, block specific IP addresses, and detect malicious activity. It uses signature-based detection to identify malicious activity in the VPC. Network firewall uses CloudWatch metrics to make firewall activity visible in real-time.

It provides better visibility of network traffic by sending logs to CloudWatch, S3, and Kinesis Firehose. It runs stateful and stateless traffic inspection rules engines, which use rules and other settings configured within a firewall policy. The administrator can manage the firewall using the below central components.

- **Firewall**: It offers traffic filtering logic for VPC subnets. It protects VPC as defined by the firewall policy. The administrator needs to offer network firewall a public subnet which is dedicated to the firewall endpoint for every AZ that has to be protected. Update the route tables to send traffic via the firewall endpoints for using the firewall.

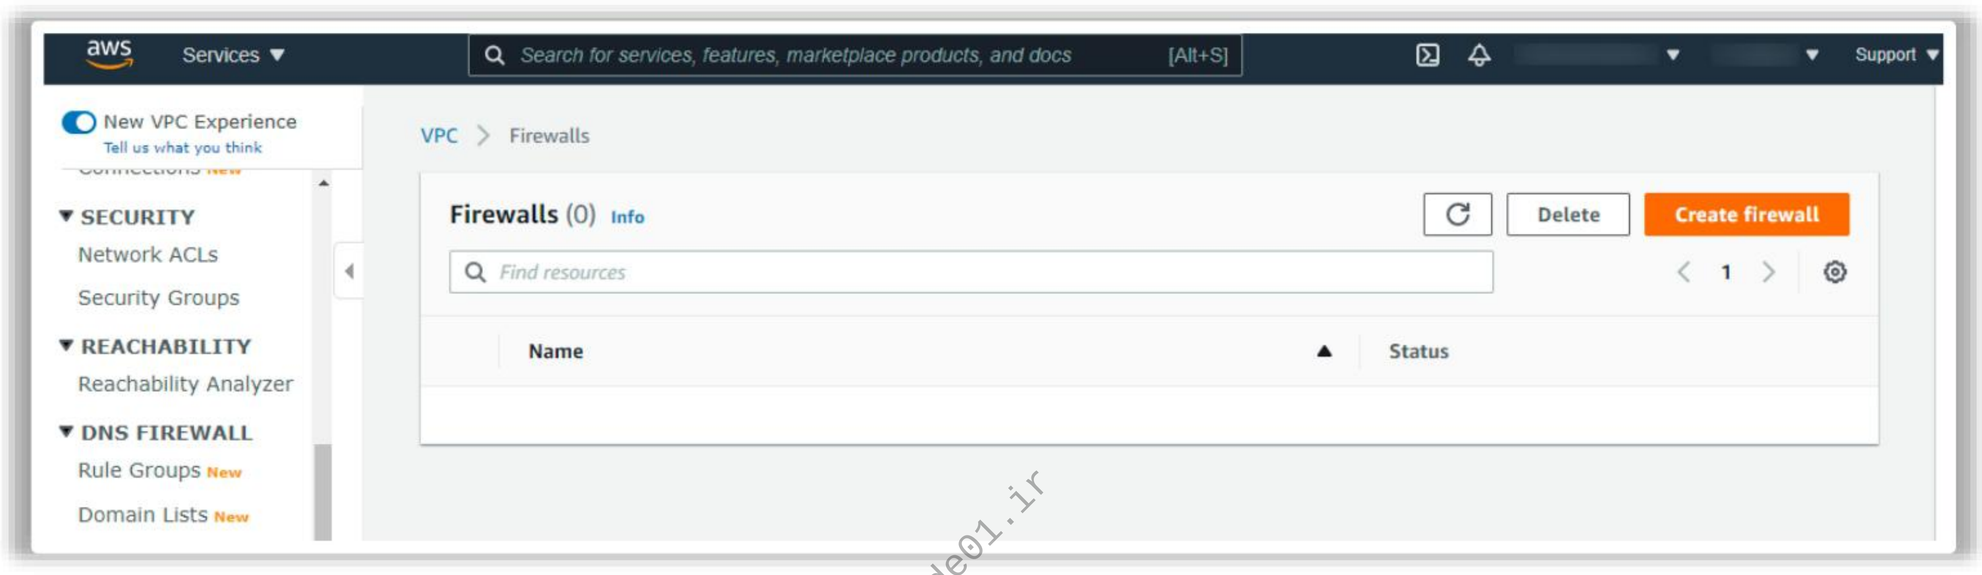
Certified Cloud Security Engineer Copyright © by **EC-Council**
All Rights Reserved. Reproduction is Strictly Prohibited.

- **Firewall policy**: It defines rules and other important settings required for a firewall to filter inbound and outbound traffic in the VPC. The administrator can assign a firewall with only one firewall policy but can associate one firewall policy with multiple firewalls.
- **Rule group**: It is a group of (stateful or stateless) rules that define how to manage and inspect network traffic.

The administrator can use CloudWatch and CloudTrail to monitor the network firewall. AWS CloudWatch gathers raw data and then processes it into simple, readable, near real-time metrics. CloudTrail records all API calls as events for network firewall.

Creating the Network Firewall (via Management Console)

- Sign-in to the Management Console and go to **VPC console**. Expand **AWS Network Firewall** and choose **Create firewall** in the menu.



hide01.ir

Figure: 2.37: Choosing Create Firewall Option

- Enter a name of the firewall and select a VPC. Create a subnet for firewall endpoint and this subnet should have a minimum of 1 IP address and a non-zero size.

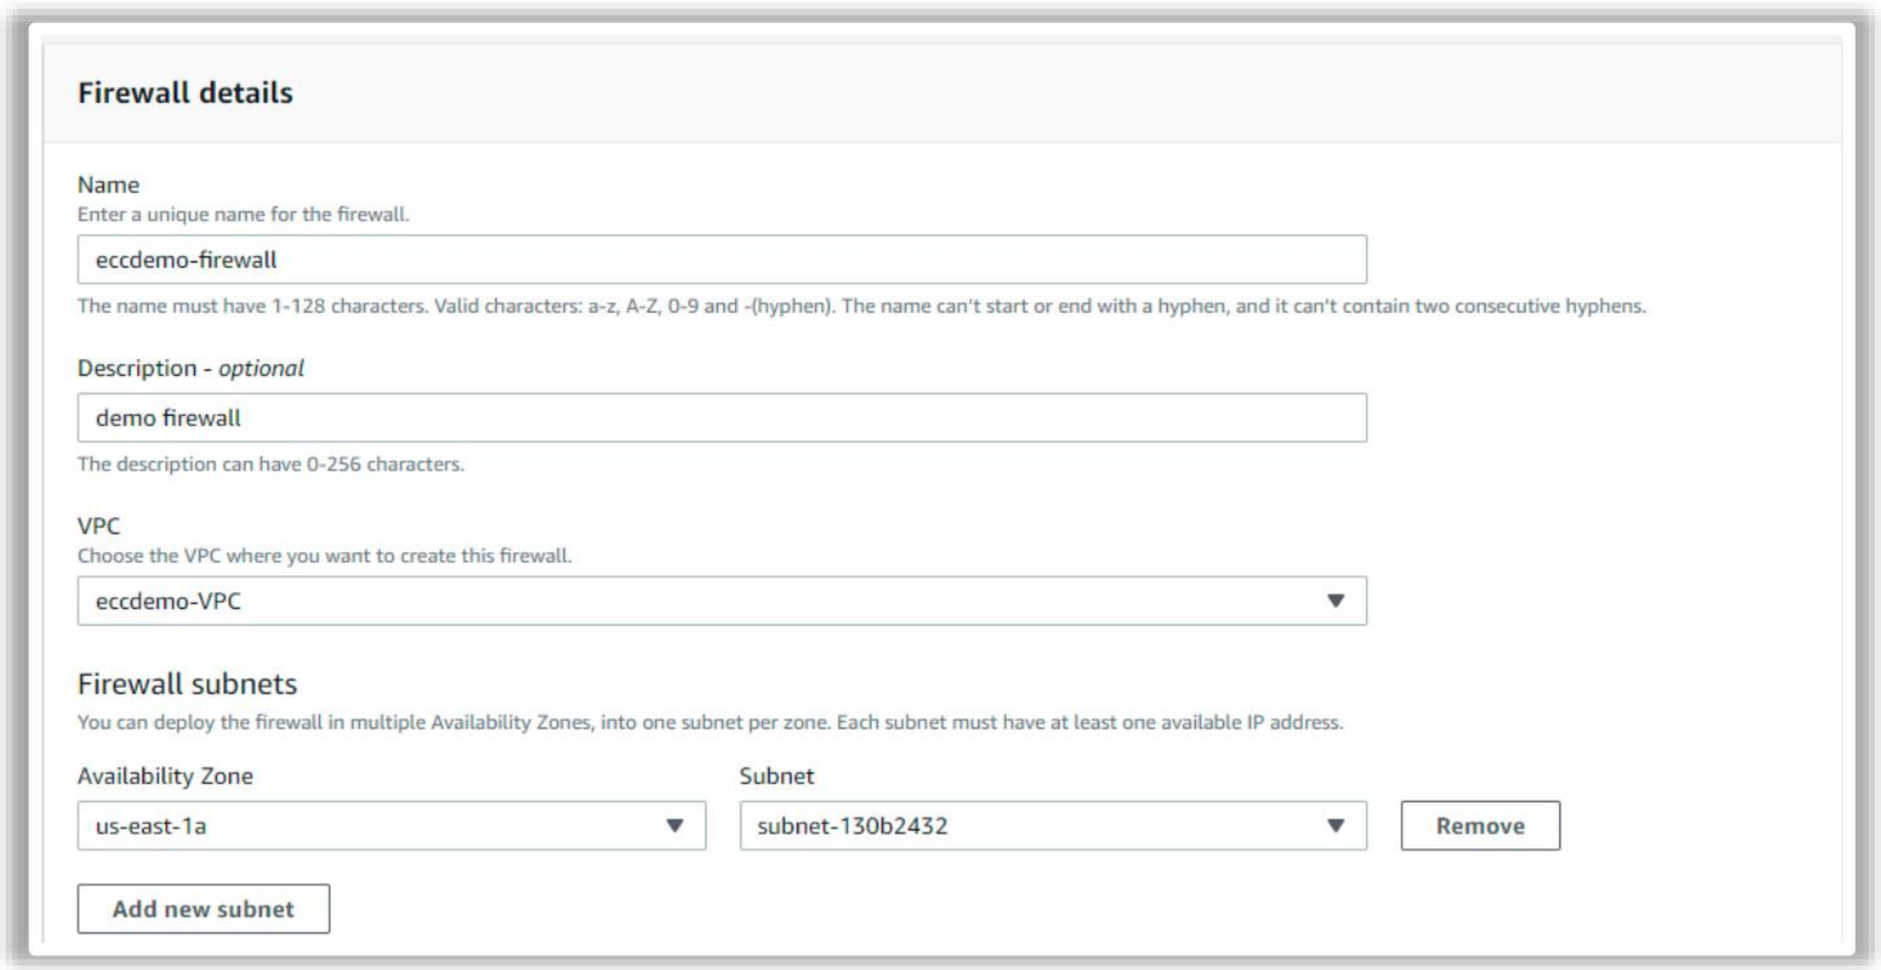


Figure:: 2.38: Choosing Details of the Firewall

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

- Choose **Create and associate an empty firewall policy** for Associated firewall policy and select **Create Firewall**.

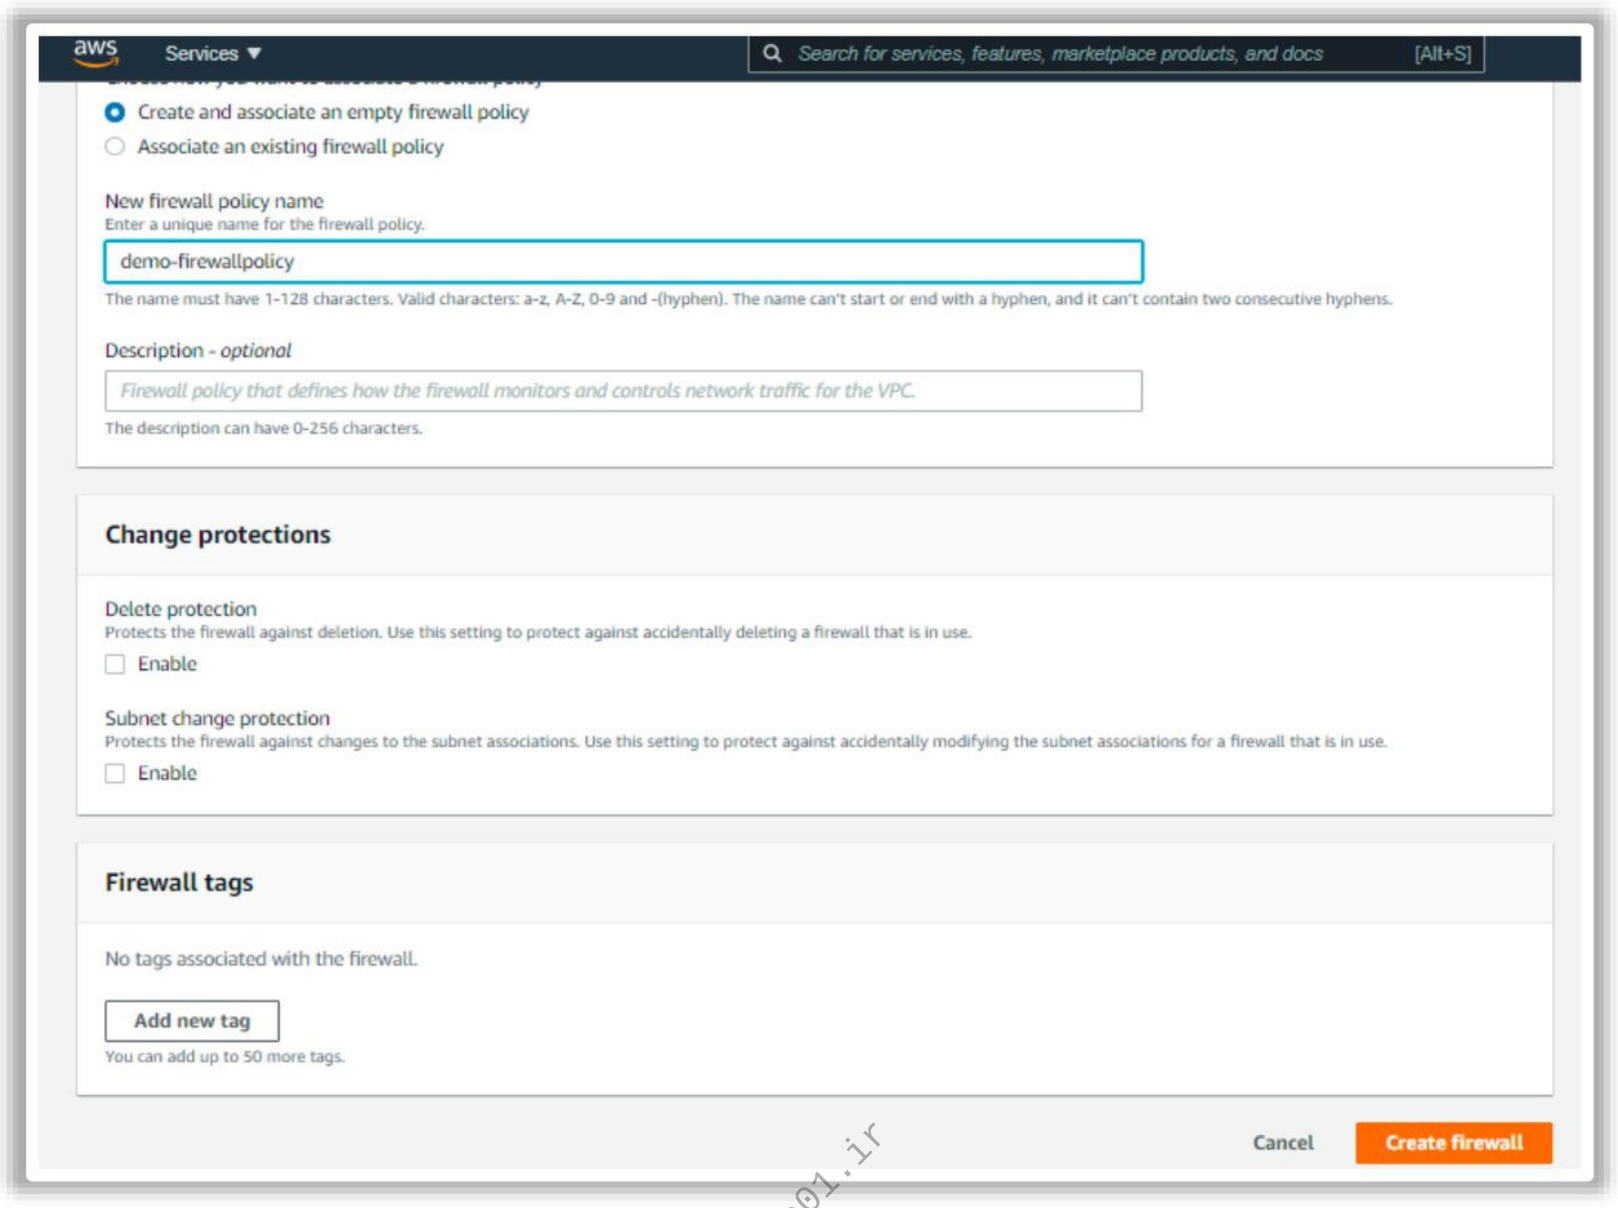


Figure: 2.39: Associating an Empty Firewall Policy for Associated Firewall Policy

- The created firewall has an empty firewall policy. Choose the newly created firewall policy from the **Firewall Policies** menu. The administrator can create or add stateful or stateless rule groups.

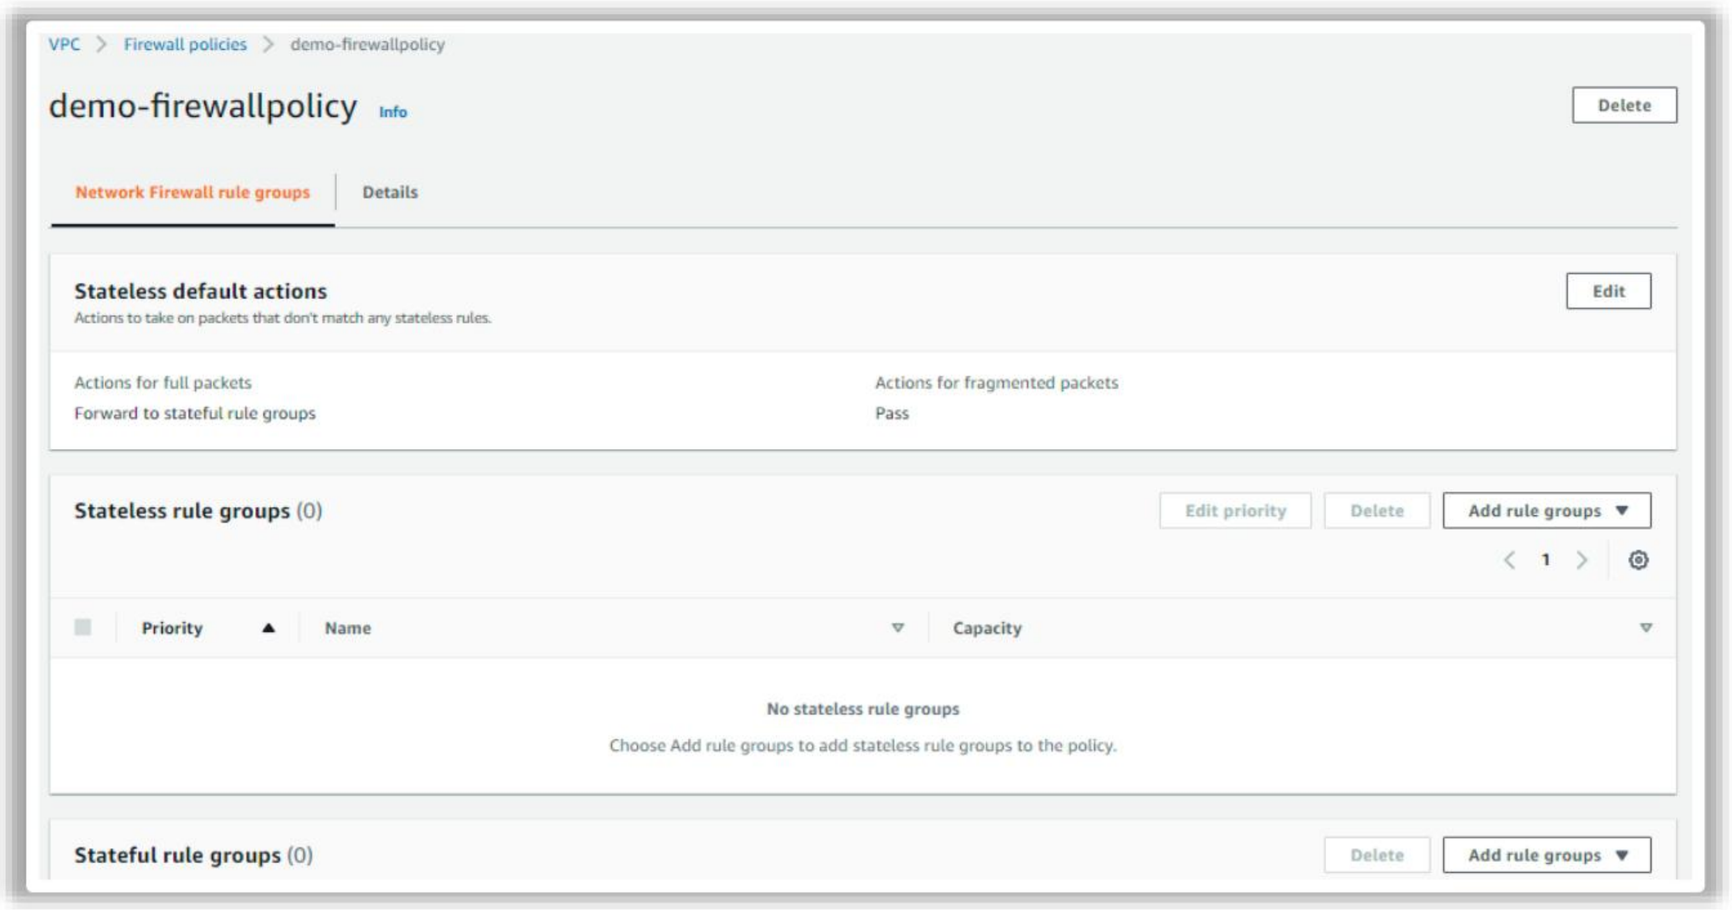


Figure: 2.40: Choosing the Newly Creating Firewall Policy

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

- For stateless default action, the administrator can specify different default settings for packet fragments and full packets. The administrator must specify the below options.
  - **Pass**: Selecting this option will discontinue packet inspection and allow all packets to reach their destinations.
  - **Drop**: Selecting this option will discontinue packet inspection and block all packets from reaching their intended destinations.
  - **Forward to stateful rules**: Choosing this option will stop stateless inspection of packets and direct them to the stateful rule engine.
- After defining the firewall policy, update the VPC route table to include the firewall and to insert the firewall into the VPC traffic flow.

Setting Up Rule Groups

In the Network Firewall rule groups menu, create new stateful or stateless rule groups and choose **Create rule group**. If selecting stateful rule group, choose one of the three options:

- **5-tuple format**: Specify source IP address, destination IP, source port, destination port, and protocol and then choose the action to be taken for matching traffic.
- **Domain list**: Define a list of domain names and actions to be taken for traffic attempts to access a domain.
- **Suricata compatible IPS rules**: Provide advanced firewall rules via Suricata rule syntax.

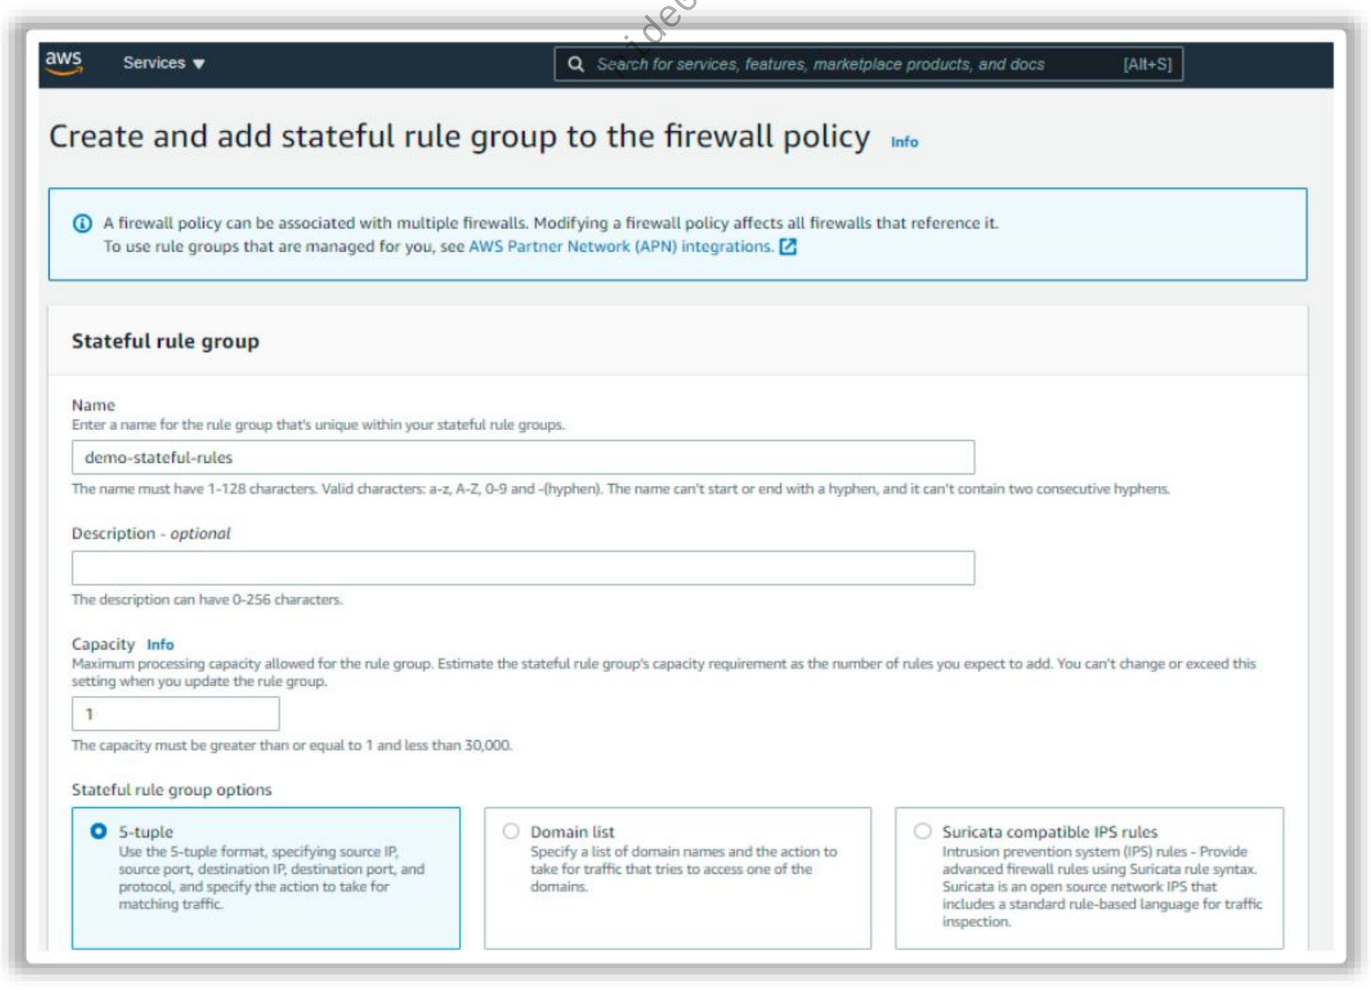


Figure: 2.41: Creating Stateful Rule Group

- The 5-tuple rule specification indicates the processing order of the stateless rule inside the rule group (as shown below).

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

Rules (3) Info — Edit rules

Priority	Protocol	Source	Destination	Source port range	Destination port range	Action	Custom action
10	17	0.0.0.0/1	0.0.0.0/0	80-80 443-443	8001-8001	Pass	-
11	6	0.0.0.0/0	0.0.0.0/0	53-53	5003-5003	Forward	
12	All	0.0.0.0/0	0.0.0.0/0	-	-	Pass	-

Figure:: 2.42: Ordered Stateless Rules in the Rule Group

- The stateful 5 tuple has match settings that specify what stateful rules engine looks for in a data packet. For becoming a match, a data packet should meet all match settings.

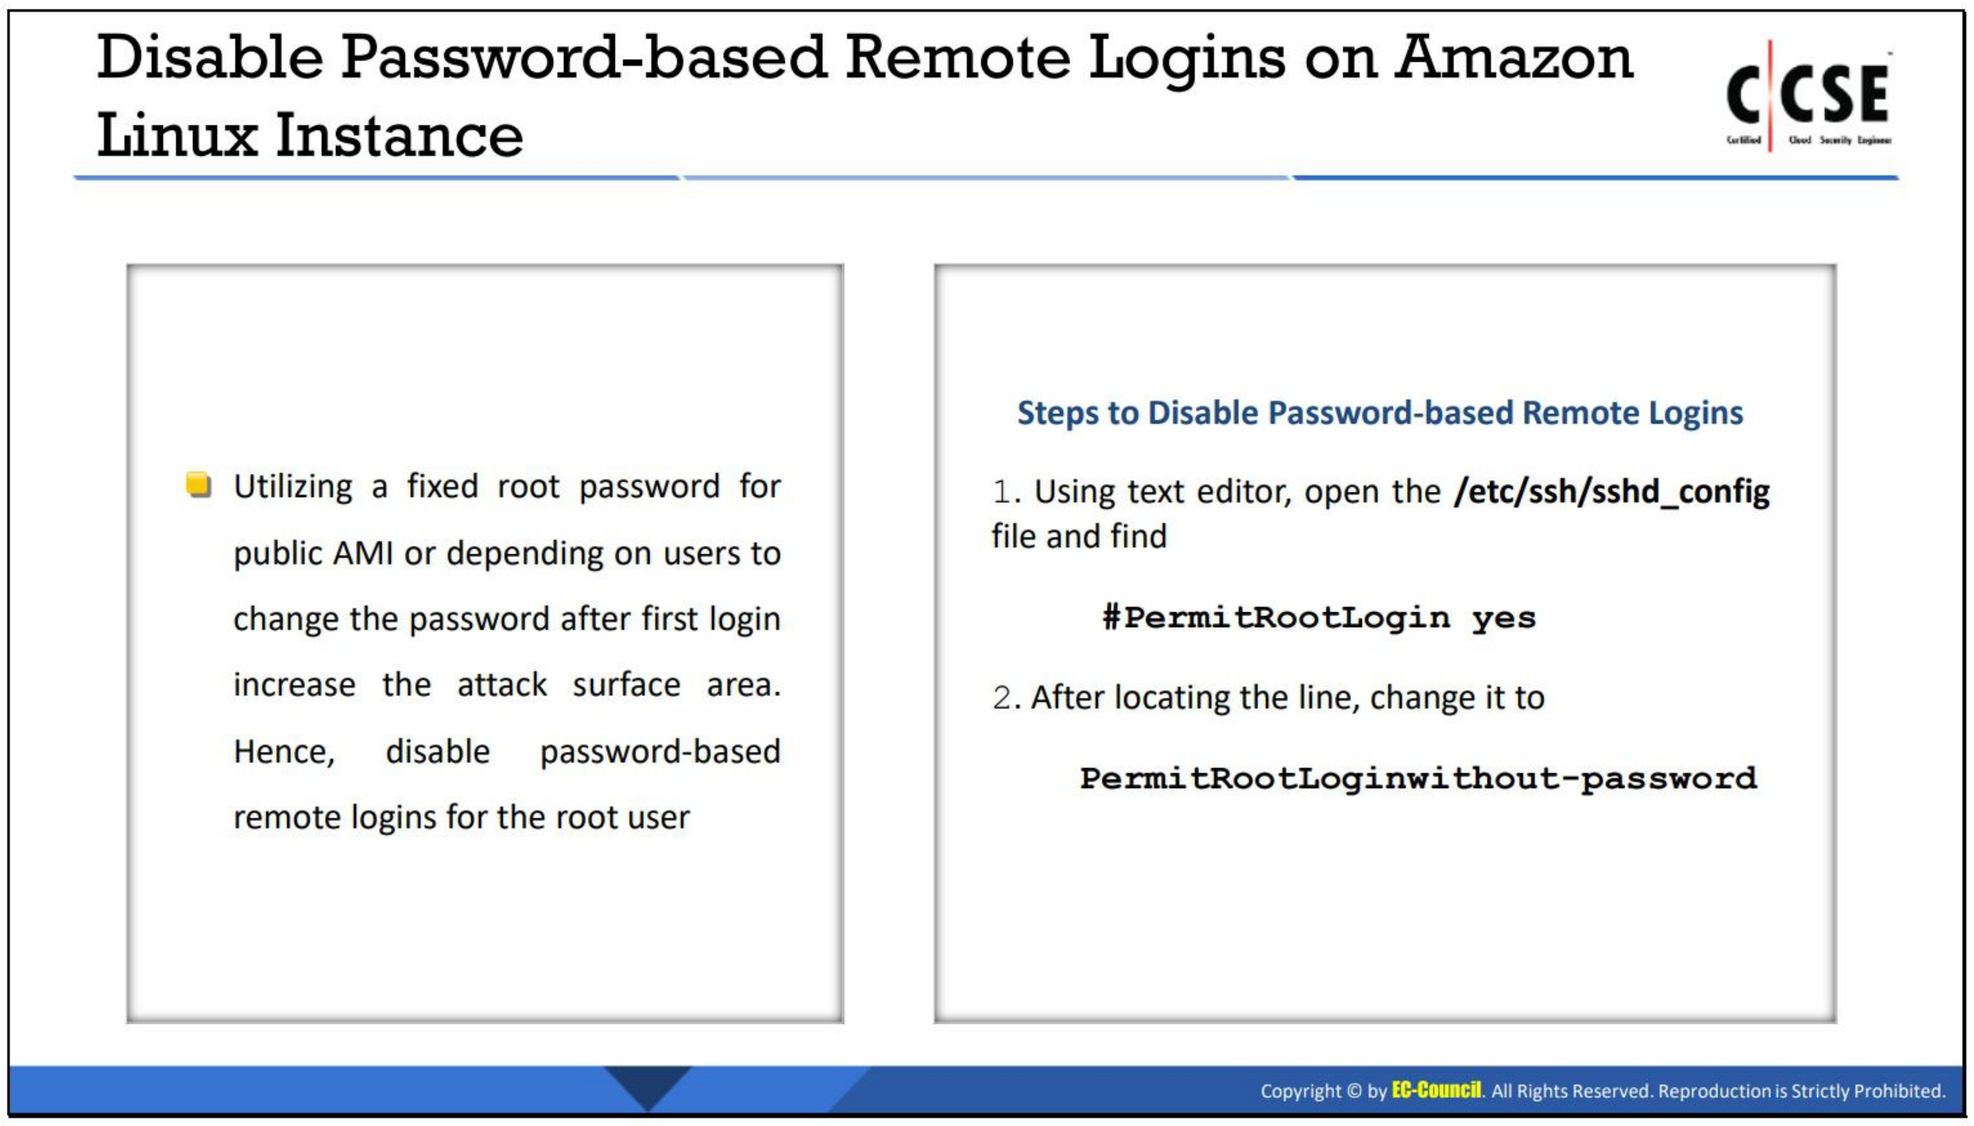
Rules (2) — Edit rule

Protocol	Source	Destination	Source port	Destination port	Action
TCP	Any	Any	443	Any	Pass
SSH	10.0.0.0/16	Any	22	Any	Drop

Figure:: 2.43: Stateless Rules

hide01.ir

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



Disable Password-based Remote Logins on Amazon Linux Instance

For a public AMI, using a fixed password can pose a security risk when it is known to attackers or threat actors. Relying on users to modify the password after they logged in for the first time is also risky and can lead to potential abuse. Users may miss changing their passwords after their first login, which allows attackers to exploit their accounts.

Utilizing a fixed root password for a public AMI or depending on users to change the password after their first login increases the attack surface area. Hence, disable password-based remote logins for remote users. Disabling password-based remote logins will eliminate all possible loopholes through which an attacker can exploit a user account remotely.

The administrator needs to disable remote logins using passwords. To do it, follow the given steps:

- Using the text editor, open the `/etc/ssh/sshd_config` file and find

```
#PermitRootLogin yes
```

- After locating the line, change it to

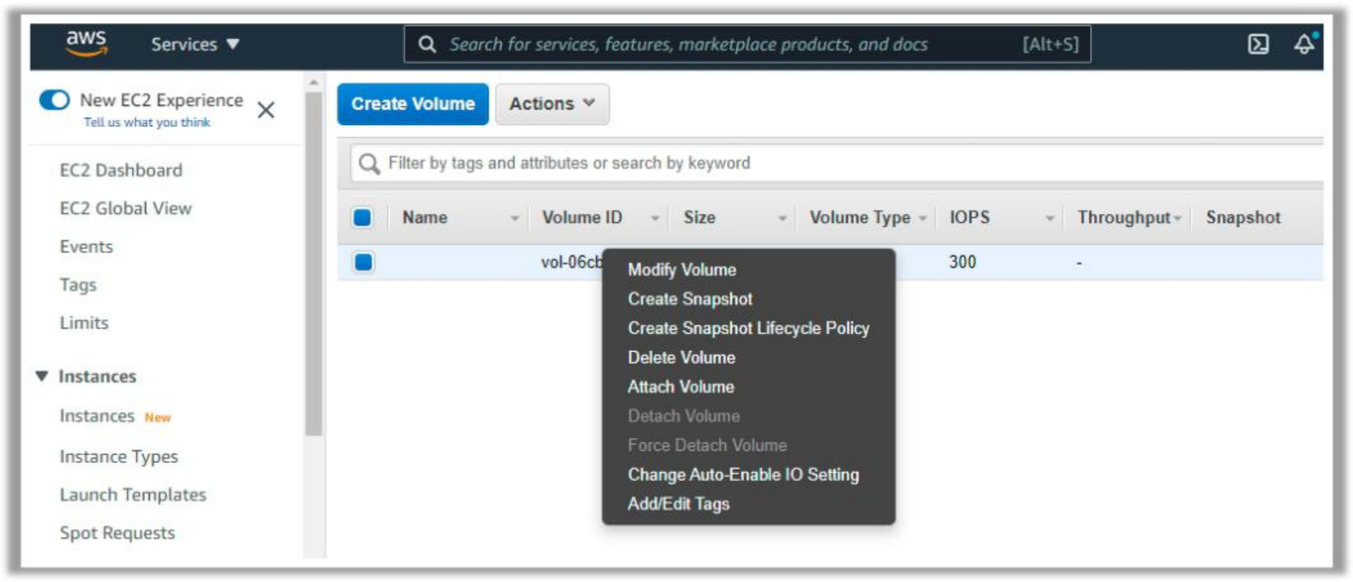
```
PermitRootLogin without-password
```

The file location may be different for different Linux distributions or when not using OpenSSH.

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

Create Amazon EBS Snapshots for Data Backup

- EBS Snapshots are the point-in-time copy of the data that can be utilized in disaster recovery, data migration, and to enhance backup compliance
- The cloud administrator can use Amazon EBS snapshots to safeguard the block storage data



Creating Amazon EBS Snapshots for Data Backup

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

Create Amazon EBS Snapshots for Data Backup

AWS has introduced a simple and much secure solution in the form of EBS snapshots to protect block storage data. EBS snapshots not only protect EBS volumes, boot volumes but also protect on-premises block data. EBS snapshots are point-in-time copy of the data that can be utilized in disaster recovery, data migration, and to enhance backup compliance. EC2 instances use EBS volumes to store data. Snapshots work like an incremental backup which includes the copy of storage blocks that have been modified or changed since the creation of the previous snapshot. Only the initial snapshot of a volume contains a complete copy of the volume.

The cloud administrator can use Amazon EBS snapshots to safeguard the block storage data. EBS snapshots can be integrated with Amazon Data Lifecycle Manager (ADLM). With ADLM, the administrator can define policies to automate snapshot lifecycle management. ADLM also helps the administrator to create, retain and remove snapshots manually.

Snapshots taken from encrypted volumes get encrypted automatically and the volumes created from these snapshots are also encrypted. The data stored in the encrypted volumes or snapshots is protected both in transit and at rest. These encrypted snapshots can be shared only with specific AWS accounts. If the administrator wants to share these encrypted snapshots with others, they need to share the CMK key used to encrypt snapshots with other users. The administrator can use EBS snapshots for ensuring business continuity during disasters, while migrating data, avoiding data loss, improve backup compliance, data backup, etc. The administrator can create multi-volume snapshots for all EBS volumes that are attached to a particular EC2 instance.

Creating an EBS Snapshot (via Management Console)

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

- Sign-in to EC2 Console (via https://console.aws.amazon.com/ec2). Under **Elastic Block Store**, choose **Snapshots** and select **Create Snapshot**.
- Choose Volume for **Select resource type** and choose the volume for **Volume**.
- Right-click on the volume and select **Create Snapshot** from the menu.

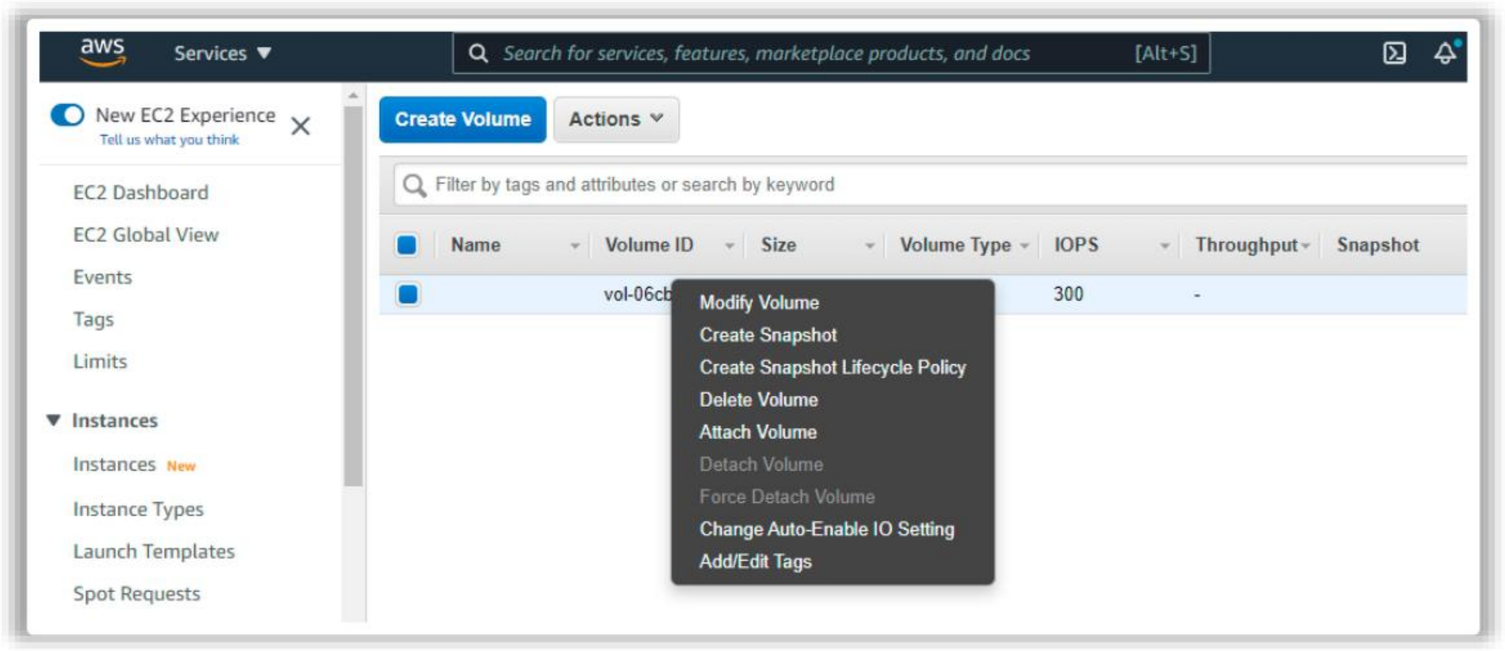


Figure: 2.44: Menu Showing Create Snapshot Option

- Enter the information asked (snapshot name and description) and choose **Create** for creating the snapshot.

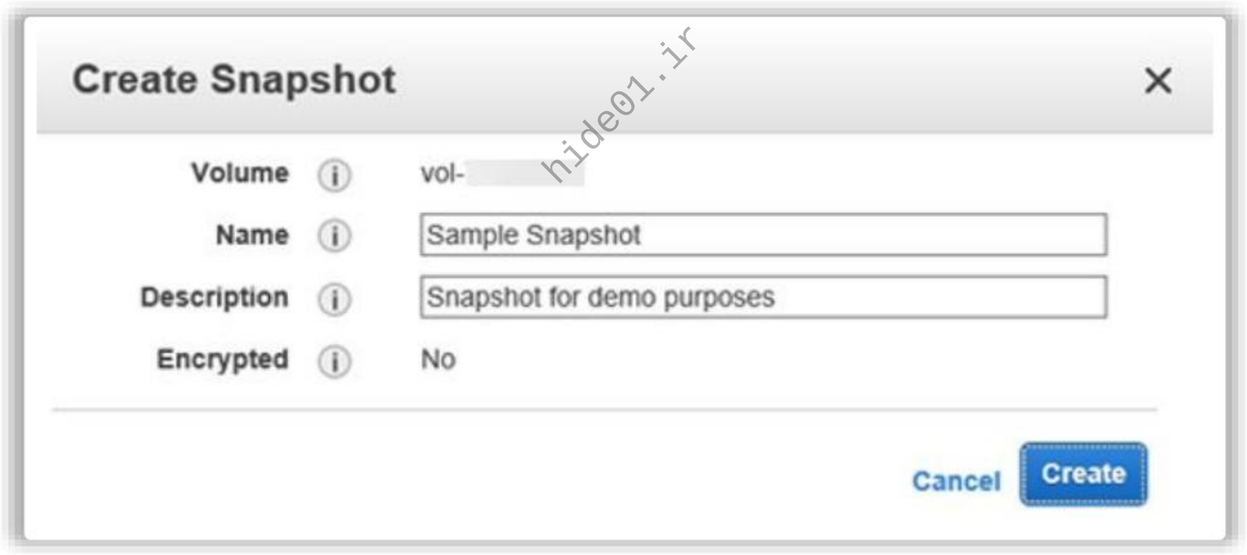


Figure: 2.45: Entering the Details of the Snapshot

- After clicking **Create**, the process of creating the snapshot starts.

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

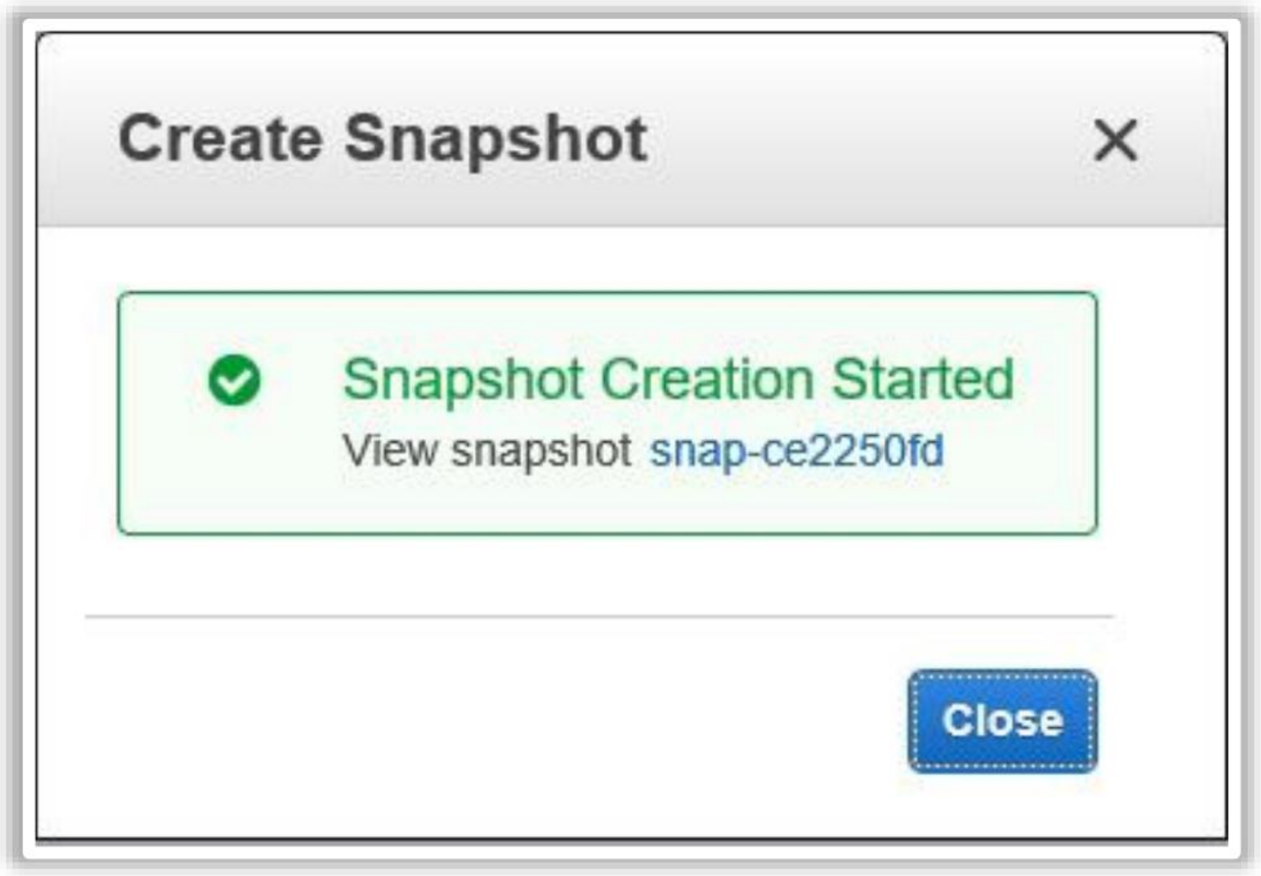


Figure: 2.46: Beginning of Snapshot Creation

- The status of the snapshot can be viewed on the console's **Snapshot** tab.

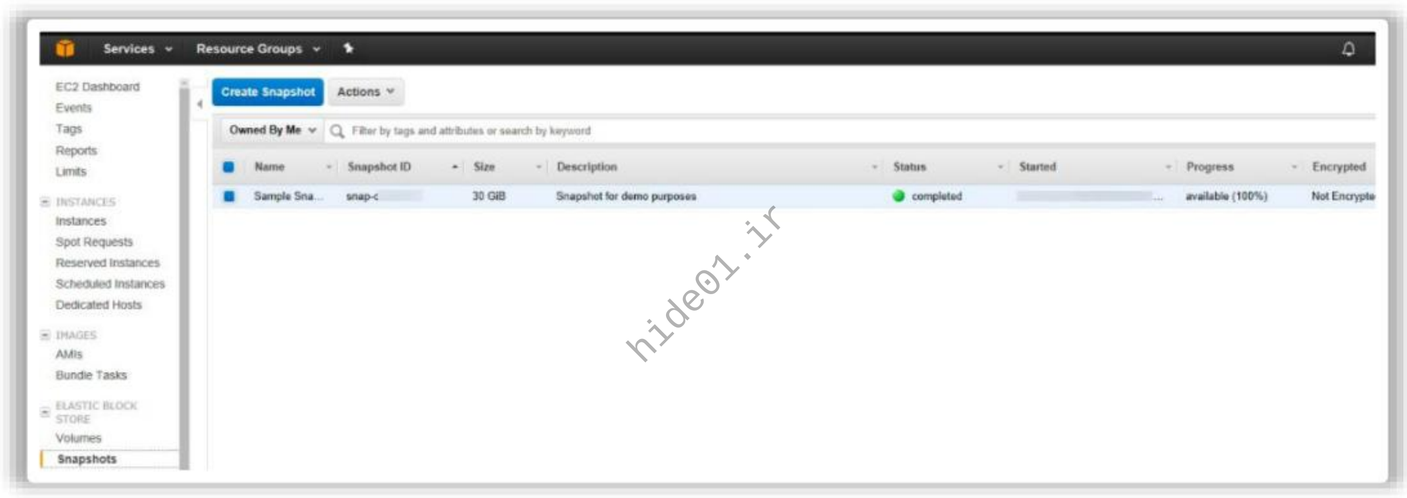
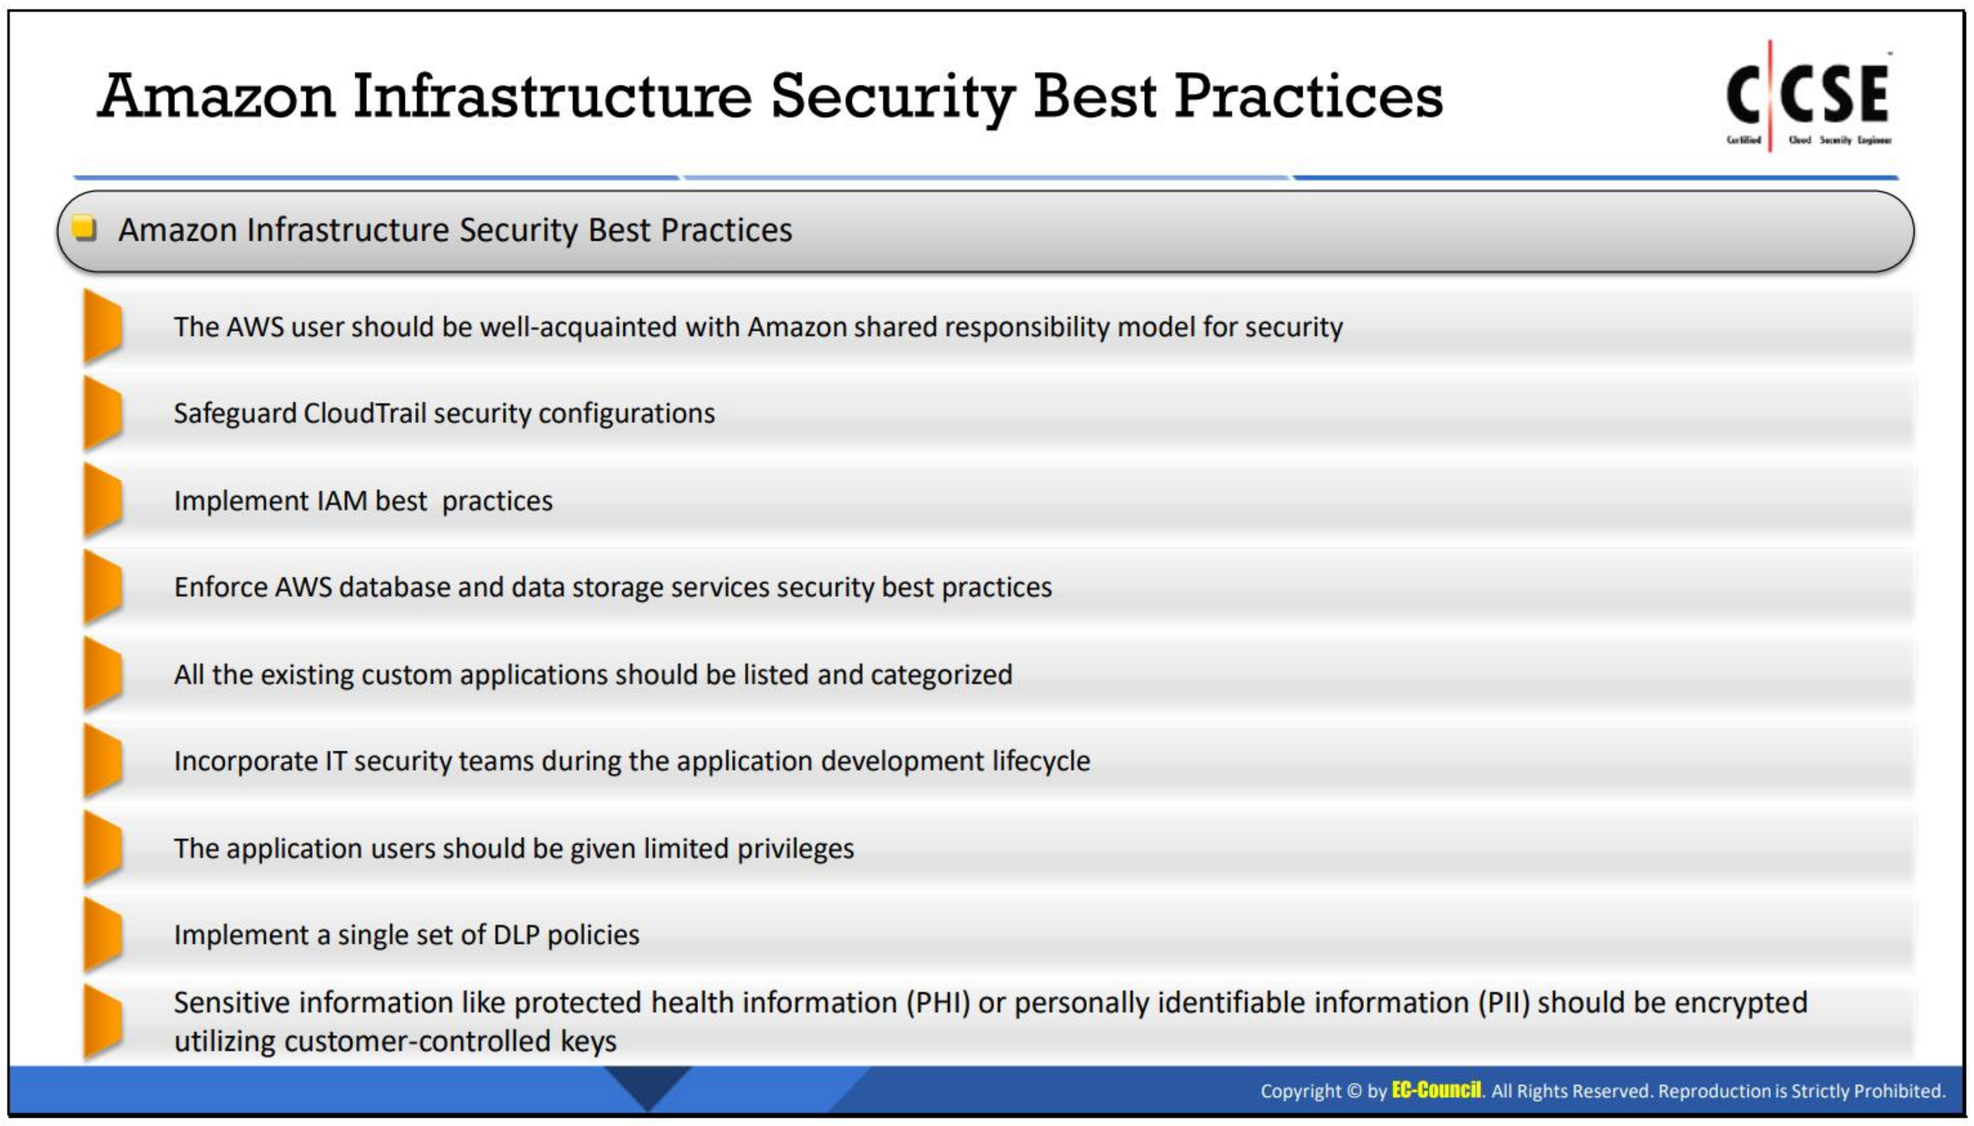


Figure: 2.47: Snapshot Tab Showing the Status of Snapshots

Creating Multi-volume EBS Snapshots

- Sign-in to EC2 Console (via https://console.aws.amazon.com/ec2). Under **Elastic Block Store**, choose **Snapshots** and select **Create Snapshot**.
- Choose **Instance** for **Select resource type** and choose the instance ID for which the multi-volume snapshots of all attached EBS volumes must be created.
- Choose **Create Snapshot.**

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



Amazon Infrastructure Security Best Practices

Security best practices for AWS infrastructure are listed below.

- **The AWS User Must be Well-Acquainted with Amazon Shared Responsibility Model for Security**: Amazon is responsible for the security of its infrastructure and platform whereas the AWS customer is responsible for secure configuration of their cloud environment and data security. A brief description of the responsibilities is as follows.
  - **Amazon's Responsibilities**: It is responsible for the security of its infrastructure, which includes all services such as storage, networking, database, etc. Amazon is also liable for protecting its physical facilities, software, and hardware on which AWS services run. Ensuring the security of its managed services such as RDS, DynamoDB, etc., is the responsibility of Amazon.
  - **Customer's Responsibilities**: Mainly AWS customers are responsible for secure handling of their data on the platform by using security features provided by Amazon such as MFA.
- **Safeguard CloudTrail Security Configurations**: CloudTrail generates logs of all API calls that are made inside AWS. It helps organizations to monitor activities that happen within AWS for post-incident forensic investigations and compliance auditing. Attackers try to disable this AWS service and delete all logs when they get access to an AWS account. To secure CloudTrail, following security measures should be taken:
  - Ensure that CloudTrail is enabled across all geographic areas and AWS services to have continuous activity monitoring.

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

- Enable CloudTrail log file validation to identify and track any changes made to the log file after being delivered to the S3 bucket and ensure log file integrity.
- Turn on access logging for the CloudTrail S3 bucket. This helps in tracking access requests and detecting potentially unauthorized access attempts.
- Enable MFA to delete the CloudTrail S3 bucket and ensure all log files are encrypted at rest and in transit.

- **Implement IAM Best Practices:** AWS provides access control and user provisioning capabilities through the IAM service. This service helps the administrator to manage users and groups and to apply permission rules to them for controlling access to resources and APIs. Follow these IAM best practices to make the most of IAM:
  - Attach IAM policies to groups or roles instead of assigning them to users to reduce the risk of an individual user getting higher or unnecessary privileges.
  - Allow access to resources by using IAM roles rather than through credentials as the credentials can be misplaced or exploited.
  - Grant minimum access privileges to IAM users that are needed to fulfill their duties.
  - Ensure MFA is enabled for all IAM users and limit IAM users who have administrator privileges.
  - Regularly rotate IAM access keys and fix a number of days for which a password is valid.
  - Implement a strong password policy that requires at least 14 characters of which one should be an upper-case letter, a number, and a symbol. Apply a password reset policy that prohibits users from using passwords they used in their past 24 password resets.

- **Enforce AWS Database and Data Storage Services Security Best Practices**: AWS Provides various database services (e.g., Amazon RDS, DynamoDB, Redshift, etc.) and data storage services (S3 and EBS) to its customers. These services must be secured to avoid any customer data leakage and other complications. Follow the given best practices to secure these services:
  - No S3 bucket must be Publicly readable or writable unless it is a part of business needs.
  - Enable Redshift audit logging. It supports post-incident forensic investigations and auditing for a particular database.
  - Encrypt all EBS data.
  - Encrypt RDS data.
  - Enable require_ssl parameter to reduce the risk of man-in-the-middle attacks in all Redshift clusters.

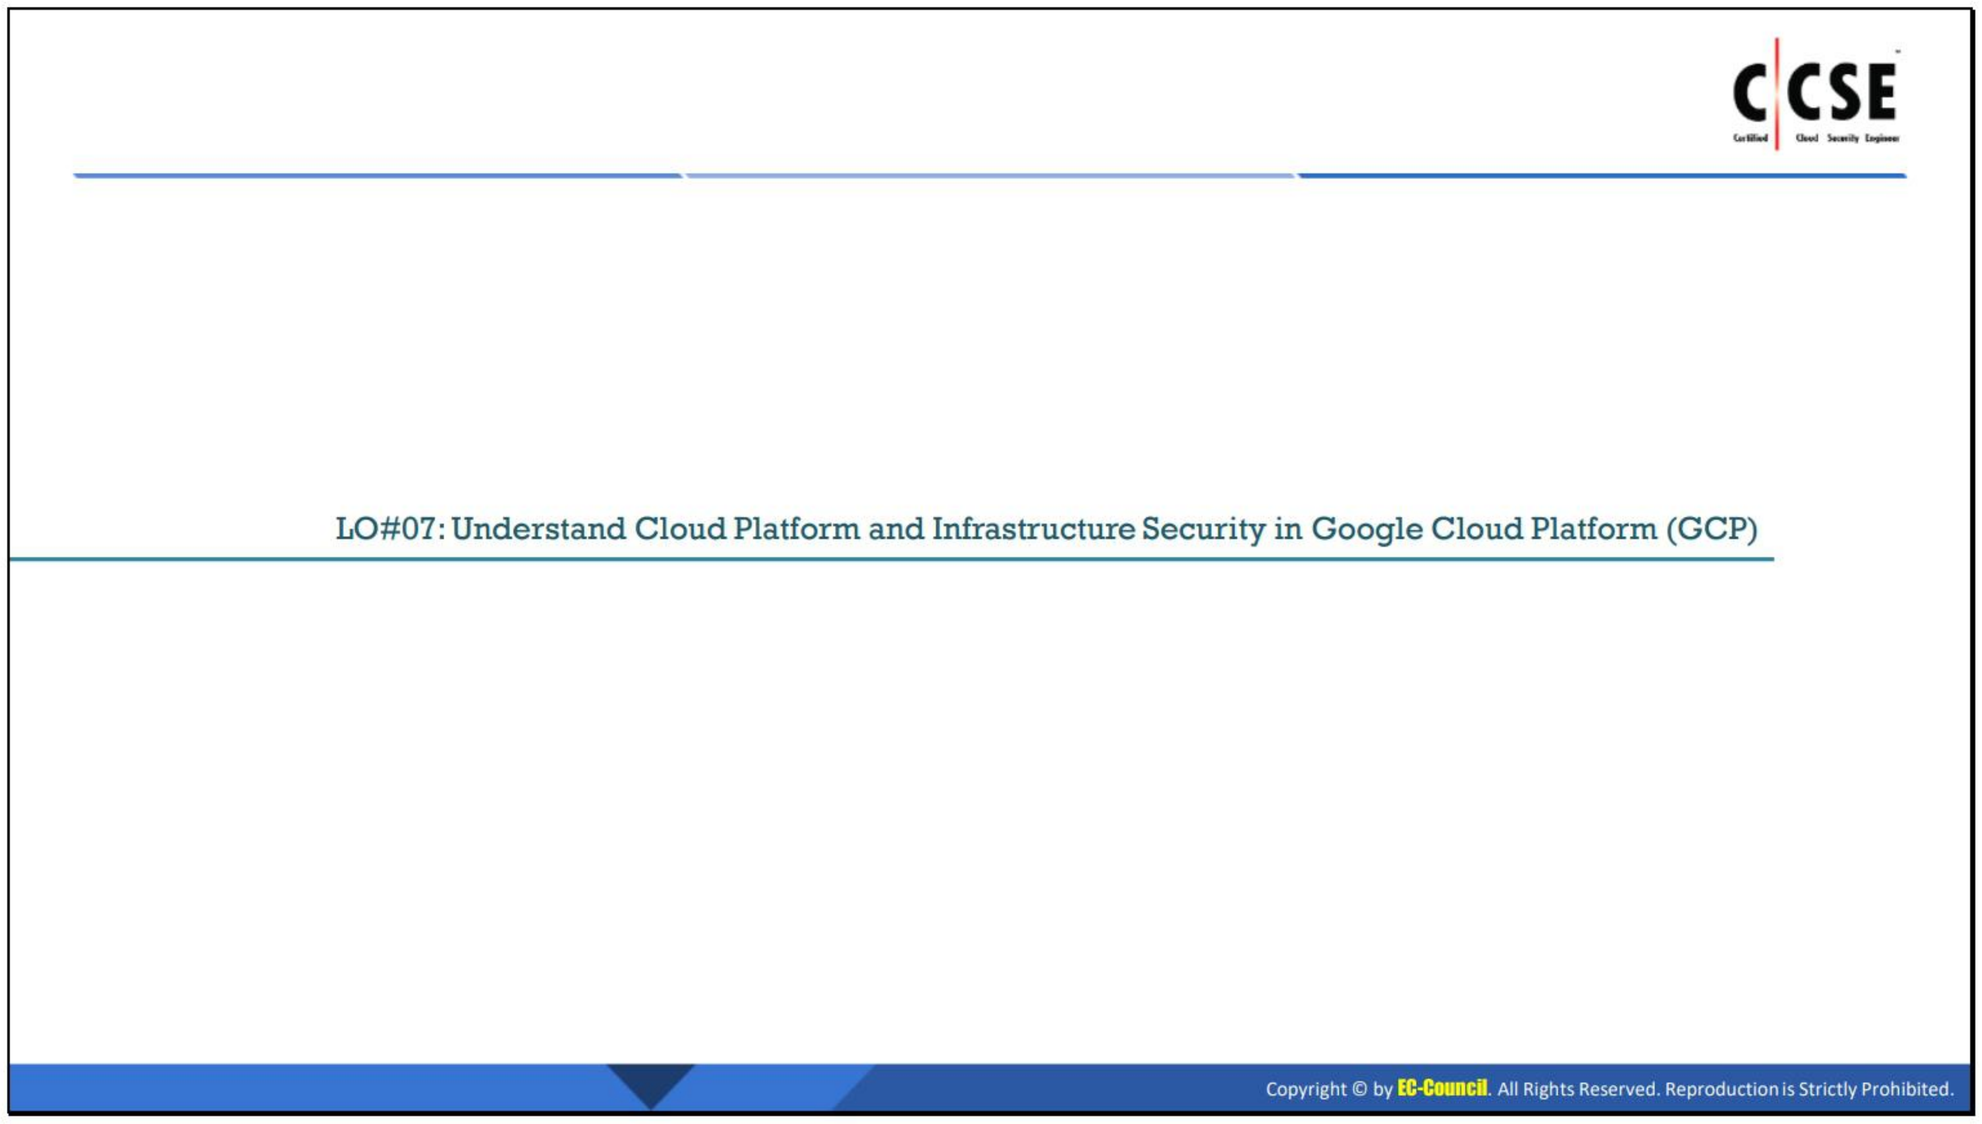
Certified Cloud Security Engineer Copyright © by **EC-Council**
All Rights Reserved. Reproduction is Strictly Prohibited.

- Limit access to RDS instances to reduce the risk of malicious activities such as DoS attacks, brute-force attacks, SQL injection, etc.

- **All the Existing Custom Applications Should be Listed and Categorized**: Enlist and categorize custom applications by the data stored in them, possible threats to them, and compliance requirements. If the security team has clear visibility of sensitive data, it will help them apply internal and external regulations on the apps and their data and apply proper security controls to protect them.
- **Involve IT Security Teams During the Application Development Lifecycle**: IT security teams must be involved in the stages of the application development lifecycle. They should be allowed to use their testing tools to test the security of the application under development. The IT security team must ensure that the application can be used by end-users in a secure manner.
- **The Application Users Should be Given Limited Privileges**: Application administrators must limit permissions or privileges granted to users and should provide the minimum privileges that are needed for them to perform their duties. Providing users with unrestricted permissions or more than required permissions can increase the risk of various threats.
- **Implement a Single Set of DLP Policies**: Use a single set of DLP policies across all custom applications and cloud services. Implementing a single DLP policy engine helps in incident reporting and remediation workflow. It also helps organizations to avoid policy enforcement gaps between various services so they can achieve higher operational efficiency.
- **Encrypt Sensitive Information**: Sensitive information like protected health information (PHI) or personally identifiable information (PII) should be encrypted utilizing customer-controlled keys.

hidden1.js

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

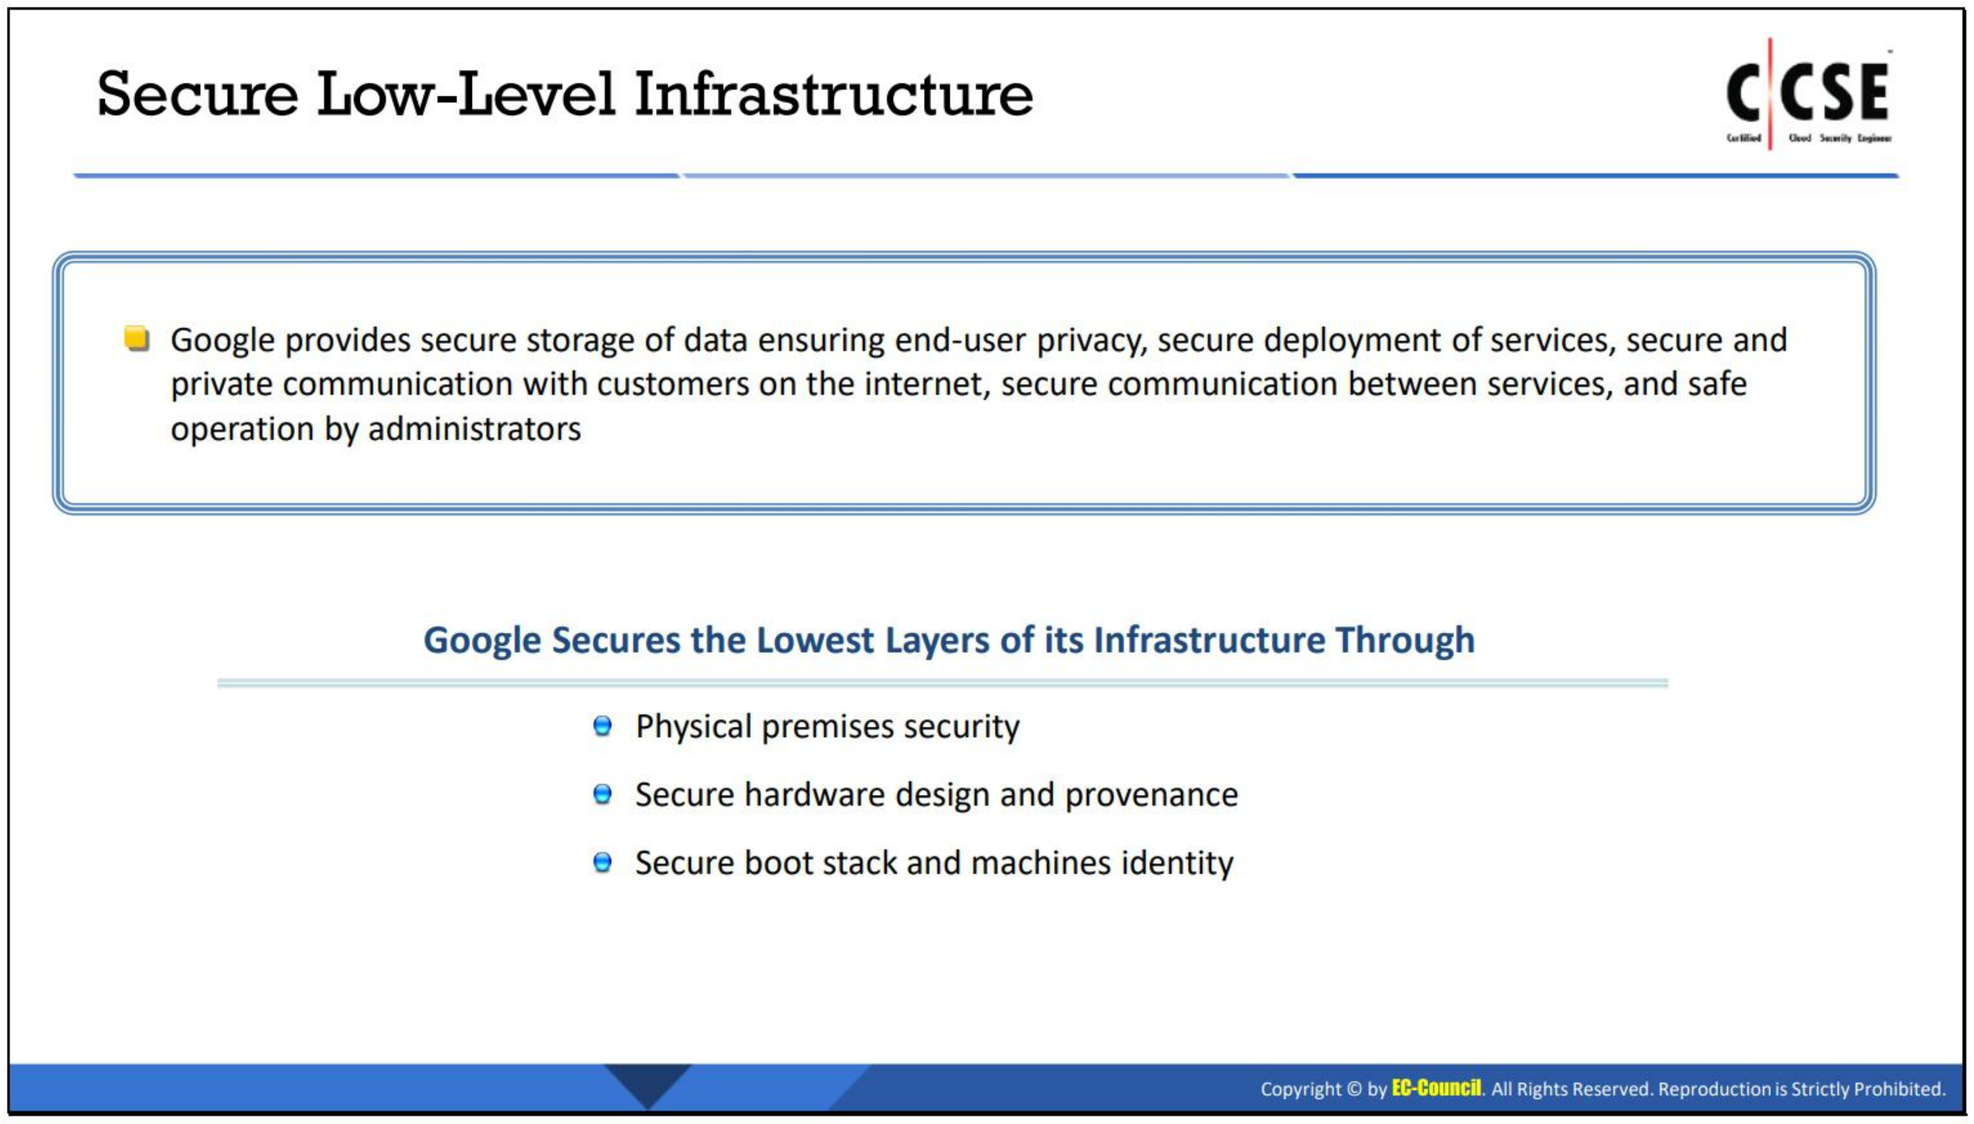


LO#07: Understand Cloud Platform and Infrastructure Security in Google

Understanding the infrastructure security features to secure the deployed services, stored data, and network communication between services helps the cloud security engineer secure the services hosted in GCP. This section will describe the resources that Google provides to end-users to ensure secure storage and communication in their infrastructure.

hid001.ir

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



Secure Low-Level Infrastructure

Google's infrastructure provides secure storage of data, thereby ensuring end-user privacy, secure deployment of services, secure and private communication with customers on the internet, secure communication between services, and safe operation by administrators.

hid01.ir

Google leverages this infrastructure to build its internet services, which include both enterprise services (such as Google Workspace and Google Cloud) and consumer services (such as Gmail, Search, and Photos).

Google secures the lowest layers of its infrastructure through the security of physical premises, hardware design and provenance, and secure boot stack and machine identity.

Security of Physical Premises: Google has designed and built its own data centers that have multiple layers of physical security protections. Access to its data centers is restricted to only a very small fraction of its employees. Google uses several physical security layers to secure its data center floors and leverages technologies such as metal detection, biometric identification, cameras, laser-based intrusion detection systems, and vehicle barriers. In third-party data centers, Google hosts some servers additionally, where it makes sure that Google-controlled physical security measures are in place on top of the security layers offered by the data center operator. For example, Google may operate cameras, independent biometric systems, and metal detectors in such sites.

Hardware Design and Provenance: A Google data center has thousands of server machines that are connected to a local network. Both the network equipment and the server boards are custom-designed by Google. Google scrutinizes or checks the background of component vendors it works with. It takes utmost care while selecting components when working with vendors to audit and validate the components' security properties. It also designs custom chips. It has

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

designed a hardware security chip that can be deployed with both peripherals and servers. At the hardware level, these chips help Google to securely identify and authenticate legitimate Google devices.

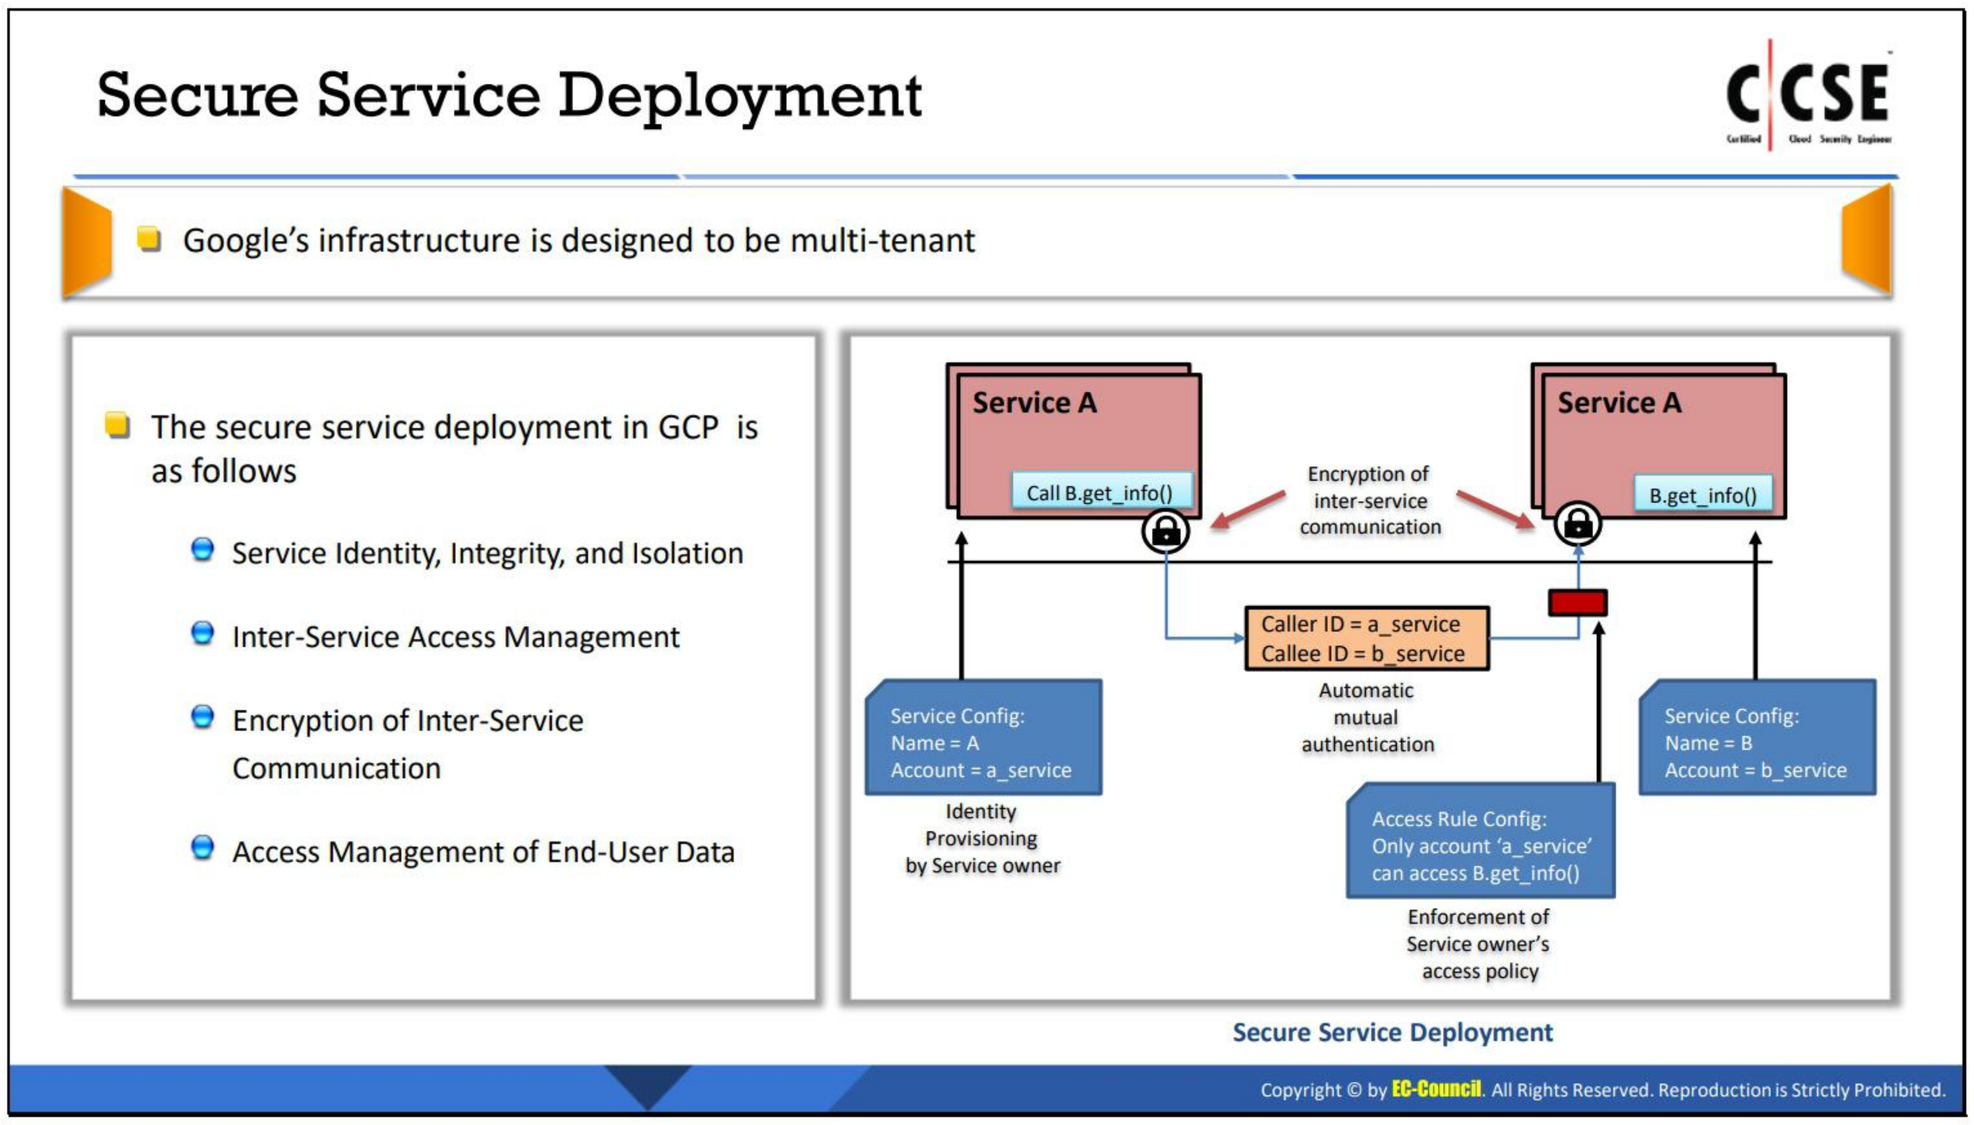
Secure Boot Stack and Machine Identity: To ensure that Google server machines boot the correct software stack, they utilize several technologies. Cryptographic signatures are utilized by Google over low-level components such as the kernel, bootloader, BIOS, and base OS image. During each boot and update, these signatures can be validated. All these components are Google-controlled, built, and hardened. Google continuously tries to improve security with every new hardware generation. For example, Google roots the trust of the boot chain in either a micro-controller running Google-written security code, a lockable firmware chip, or the Google-designed security chip depending on the generation of server design.

In the data center, every server machine has its specific or unique identity that is tied or linked to the hardware root of trust and the software that helps the machine to boot. For authenticating API calls to and from low-level management services on the machine, this identity is useful.

Google's automation systems are authored to ensure servers run up-to-date versions of their software stacks (as well as security patches), remove machines from service if required, and detect and diagnose hardware and software issues.

hide01.ir

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



Secure Service Deployment

"Service" means an application binary written by a developer who wants to run it on Google's infrastructure; for example, a Bigtable storage server, a Gmail SMTP server, a YouTube video transcoder, or an App Engine sandbox that is running a customer application. To handle the required scale of the workload, there may be several machines running copies of the same service. A cluster orchestration service, known as Borg, controls the services running on the infrastructure. Google's infrastructure is basically designed to be multitenant.

Service Identity, Integrity, and Isolation: For inter-service communication at the application layer, Google uses cryptographic authentication and authorization. This offers strong access control at an abstraction level and granularity that can be understood naturally by administrators and services.

Though Google utilizes ingress and egress filtering at multiple points in its network to avoid IP spoofing as an additional security layer, it does not rely on firewalling or internal network segmentation as its primary security mechanism. This approach aids in maximizing its network performance and availability as well.

Every service running on the infrastructure has a service account identity associated with it. Cryptographic credentials are provided to every service to establish or prove its identity while receiving or making remote procedure calls to other services. Clients use these identities to make sure they are talking to the right server, and servers use them to limit access to methods and data for specific clients.

Google stores its source code in a central repository in which both past and current versions of the service are auditable. Additionally, the infrastructure can be configured to require that a service's binaries are built from specific source code that is reviewed, checked, and tested.

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

Inspection and approval are required from at least one engineer other than the developer who wrote it for such code reviews, and the system mandates that any code modification to any system must be approved by the system owners. These requirements provide a forensic trail from a service back to its source and also limit the ability of insiders or adversaries to make malicious modifications to source code.

To protect a service from other services that run on the same machine, Google provides multiple isolation and sandboxing techniques, including language and kernel-based sandboxes, hardware virtualization, and normal Linux user separation. Generally, Google utilizes more layers of isolation for riskier workloads; for example, while running user-supplied code for products such as Google Compute Engine or Google App Engine or while running complex file format converters on user-supplied data. It provides some very sensitive services as an extra security boundary such as some key management services and the cluster orchestration service to run exclusively on some dedicated machines.

Inter-service Access Management: By using access management features that the infrastructure provides, a service owner can exactly specify which other services can communicate with that service. For example, a service may wish to offer some APIs only to a particular whitelist of other services. With the whitelist of the allowed service account identities, that service can be configured, and the infrastructure then enforces this access restriction automatically.

Individual identities are issued to Google engineers who access services; thus, services can be configured similarly to deny or allow their accesses. All types of identities (employee, machine, and service) are in a global name space, which is maintained by the infrastructure. End-user identities are separately handled.

hide01.ir

A good identity management workflow system is provided by the infrastructure for the internal identities. This workflow system includes logging, approval chains, and notification. For example, these identities can be allocated to access control groups through a system that allows two-party control where any change (to the group) proposed by one engineer must be approved by another engineer (who is also an administrator of the group). This system enables secure management processes to scale to several hundreds of services running on the infrastructure.

The infrastructure also allows services to read from group databases and central ACL along with providing the automatic API-level access control mechanism. This enables them to implement their custom and fine-grained access control when required.

Encryption of Inter-service Communication: In addition to the RPC authentication and authorization capabilities, cryptographic privacy and integrity for RPC data are provided by the infrastructure on the network. Google encapsulates other application-layer protocols like HTTP inside its infrastructure RPC mechanisms to provide the above-mentioned security benefits to them. This offers application-layer isolation and eliminates any dependency on the network path security. Even if a network device is compromised or the network is tapped, encrypted inter-service communication can be secure.

For each infrastructure RPC, services can configure the level of cryptographic protection they need, for example, they can configure only integrity-level protection for low-value data inside data centers. To defend against sophisticated attackers who try to tap Google's private WAN

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

links, all infrastructure traffic that goes over the WAN between data centers is automatically encrypted by the infrastructure, without needing any explicit configuration from the service. To extend this default encryption within its data centers for all infrastructure RPC traffic, Google deploys hardware cryptographic accelerators.

Access Management of End-user Data: A normal Google service is written (coded) to perform a task for end-users. For example, Gmail allows end-users to store their emails on it. The interaction of end-users with an application such as Gmail spans other services inside the infrastructure. For example, an API provided by the Contacts service may be called by the Gmail service to access the address book of an end-user.

The Contacts service can be configured in a way that only RPC requests from the Gmail service are allowed (or from any specific service that the Contacts service wishes to allow). However, still, this is a very broad set of permissions. The Gmail service can request any user's contact at any time within the scope of this permission.

The infrastructure requires the Gmail service to present an "end-user permission ticket" as a part of an RPC because the Gmail service, on behalf of a particular end-user, makes the RPC request to the Contacts service. The ticket endorses that the Gmail service is making a request on behalf of that particular end-user. This allows the Contacts service to implement a safeguard by returning only the data for the end-user named in the ticket.

These "end-user permission tickets" are issued by a central user identity service provided by the infrastructure. The central user identity service verifies the end-user login and issues a user credential, such as OAuth token or cookie, to the user's client device. This user credential must be present in every subsequent request to Google from the client device.

A service passes the user credential to the central user identity service for verification after it receives the credential. If the user credential is verified correctly, a short-lived "end-user permission ticket" is returned by the central user identity service, which can be used for RPCs related to the request. In the above example, the service that receives the "end-user permission ticket" is the Gmail service, which passes the ticket to Contacts service. After that, that ticket can be handed down by the calling service to the callee for all cascading calls as part of the RPC call.

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

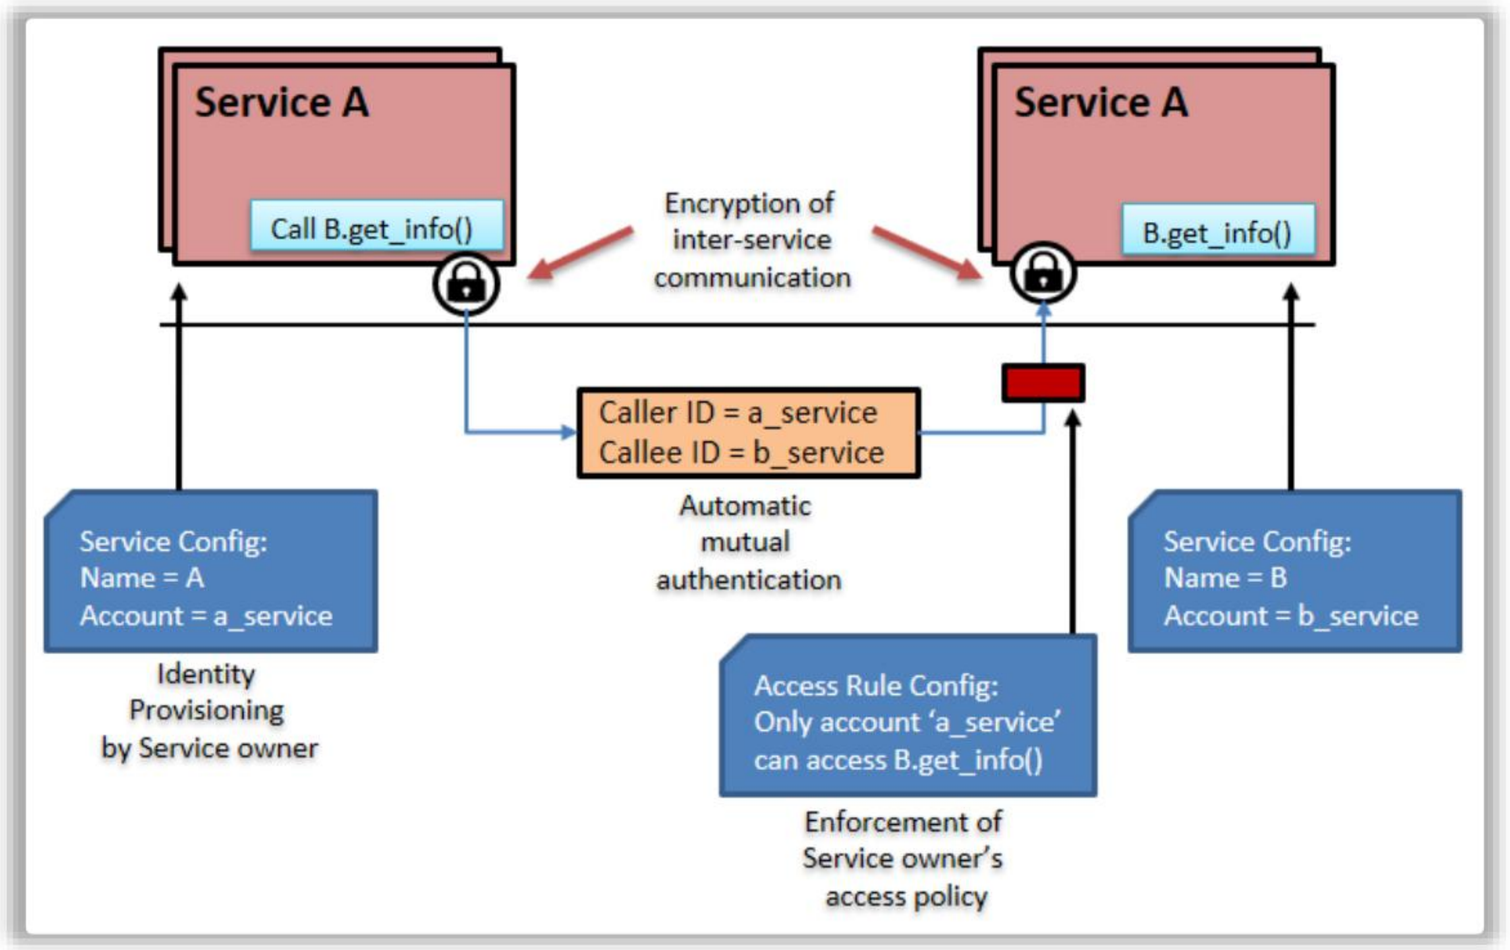
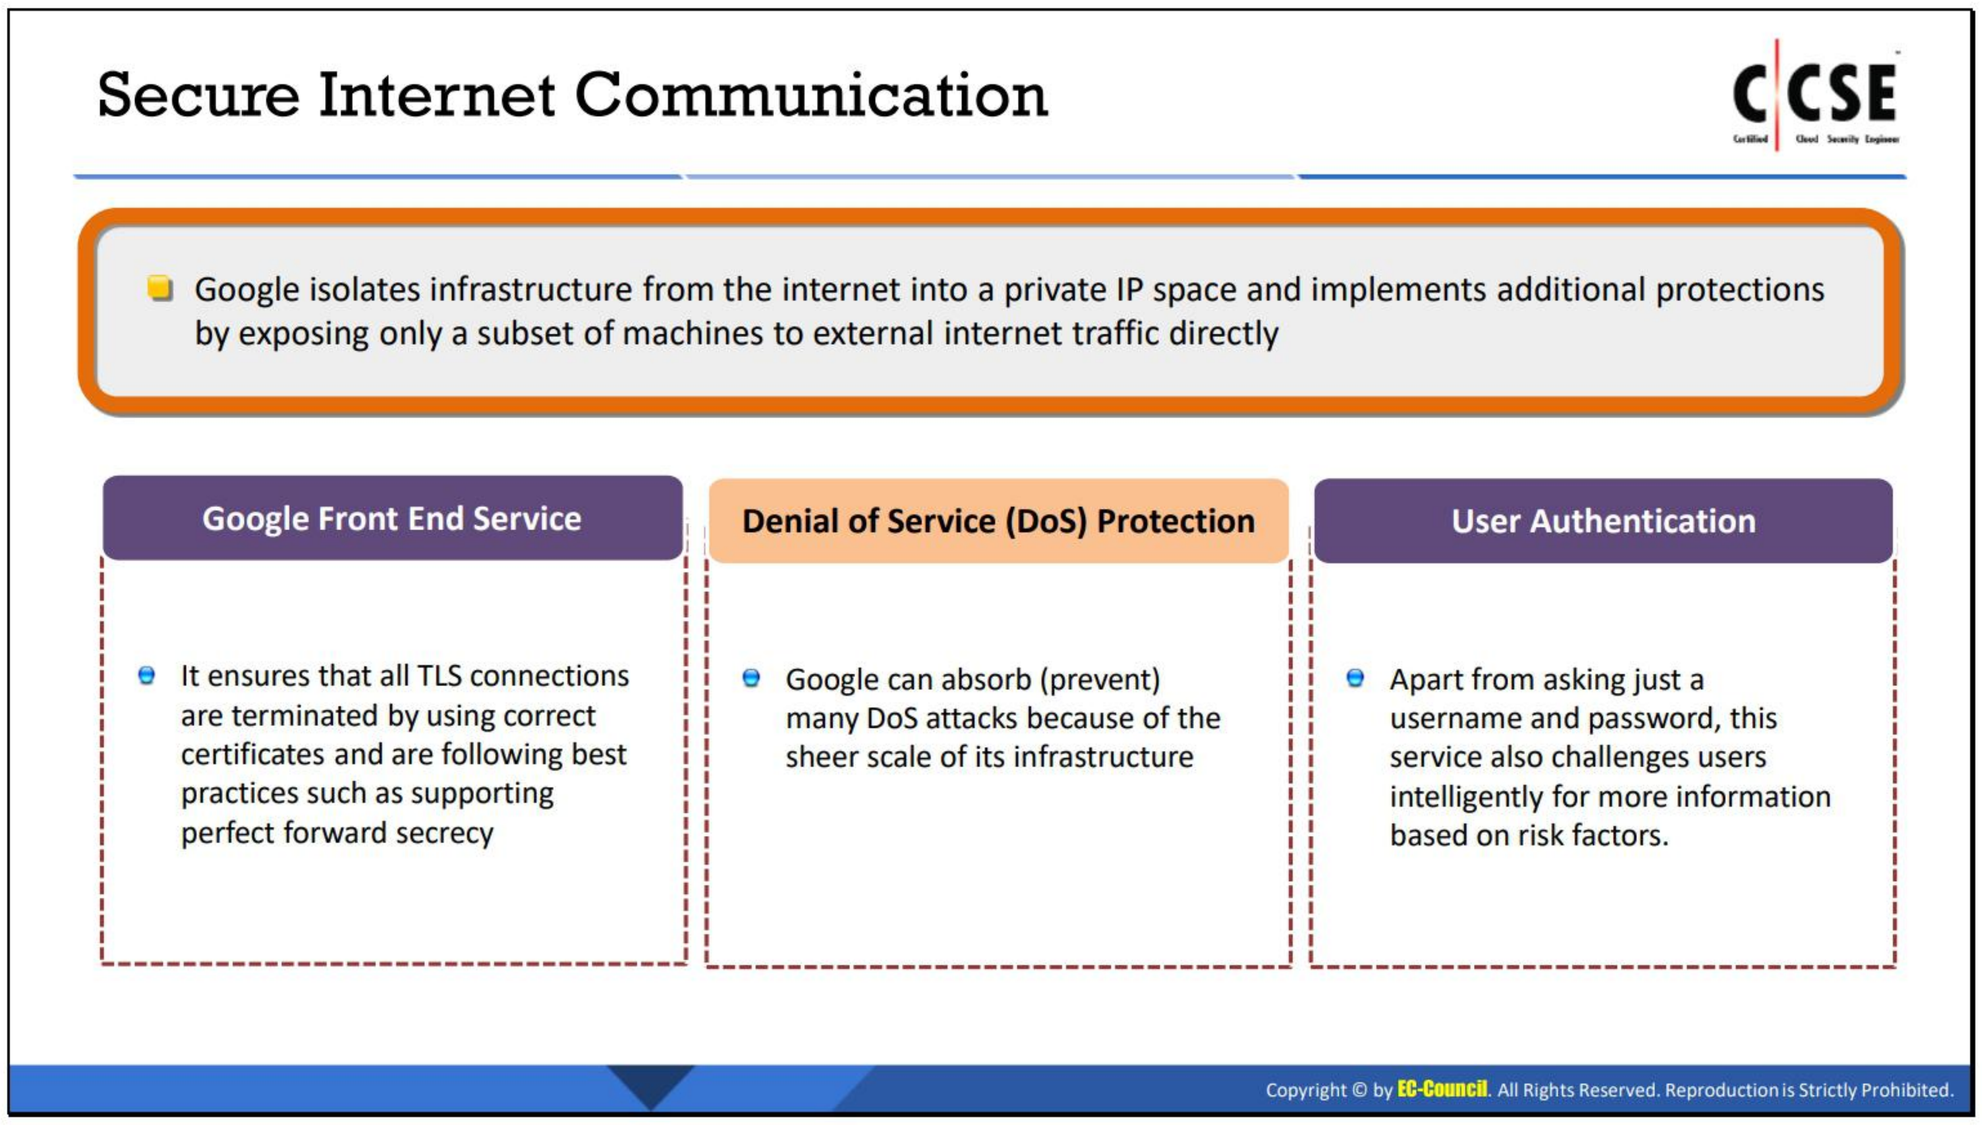


Figure: 2.48: Secure Service Deployment

As shown in the above image, Google's infrastructure provides service identity, encrypted inter-service communication, automatic mutual authentication, and enforcement of access policies defined by the owner of the service.

hide01.ir

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



Secure Internet Communication

Google's infrastructure contains a large number of physical machines that are interconnected over the LAN and WAN and the inter-service communication security is independent of the network security. However, Google isolates its infrastructure from the internet into a private IP space. This helps Google in implementing additional protections such as defenses against DoS attacks more easily by exposing only a subset of machines to external internet traffic directly.

Google Front End (GFE) Service: If a service wants itself to be available for users on the internet, it may register itself with the Google Front End (an infrastructure service). GFE makes sure that all TLS connections are terminated by using correct certificates and are following best practices such as supporting perfect forward secrecy. Additionally, it applies protection against DoS attacks. It uses the RPC security protocol to forward requests for the service.

Any internal service that wants itself to be published externally utilizes the GFE as a smart reverse-proxy front end. This front end provides DoS protection, public IP hosting of its public DNS name, and TLS termination. Like other services, GFEs also run on the infrastructure and have the ability to scale to match (handle) incoming request volumes.

Denial-of-service (DoS) Protection: Google can absorb (prevent) many DoS attacks because of the sheer scale of its infrastructure. Besides, it provides multi-layer, multi-tier DoS protections that reduce the risk of DoS impact on a service that runs behind a GFE even more.

After an external service is delivered to one of Google's data centers, it goes through multiple layers of hardware and software load balancing. Information about incoming traffic is reported by these load balancers to a central DoS service running on the infrastructure. If the central DoS service detects a DoS attack, it configures the load balancers to throttle or drop the traffic associated with the attack.

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

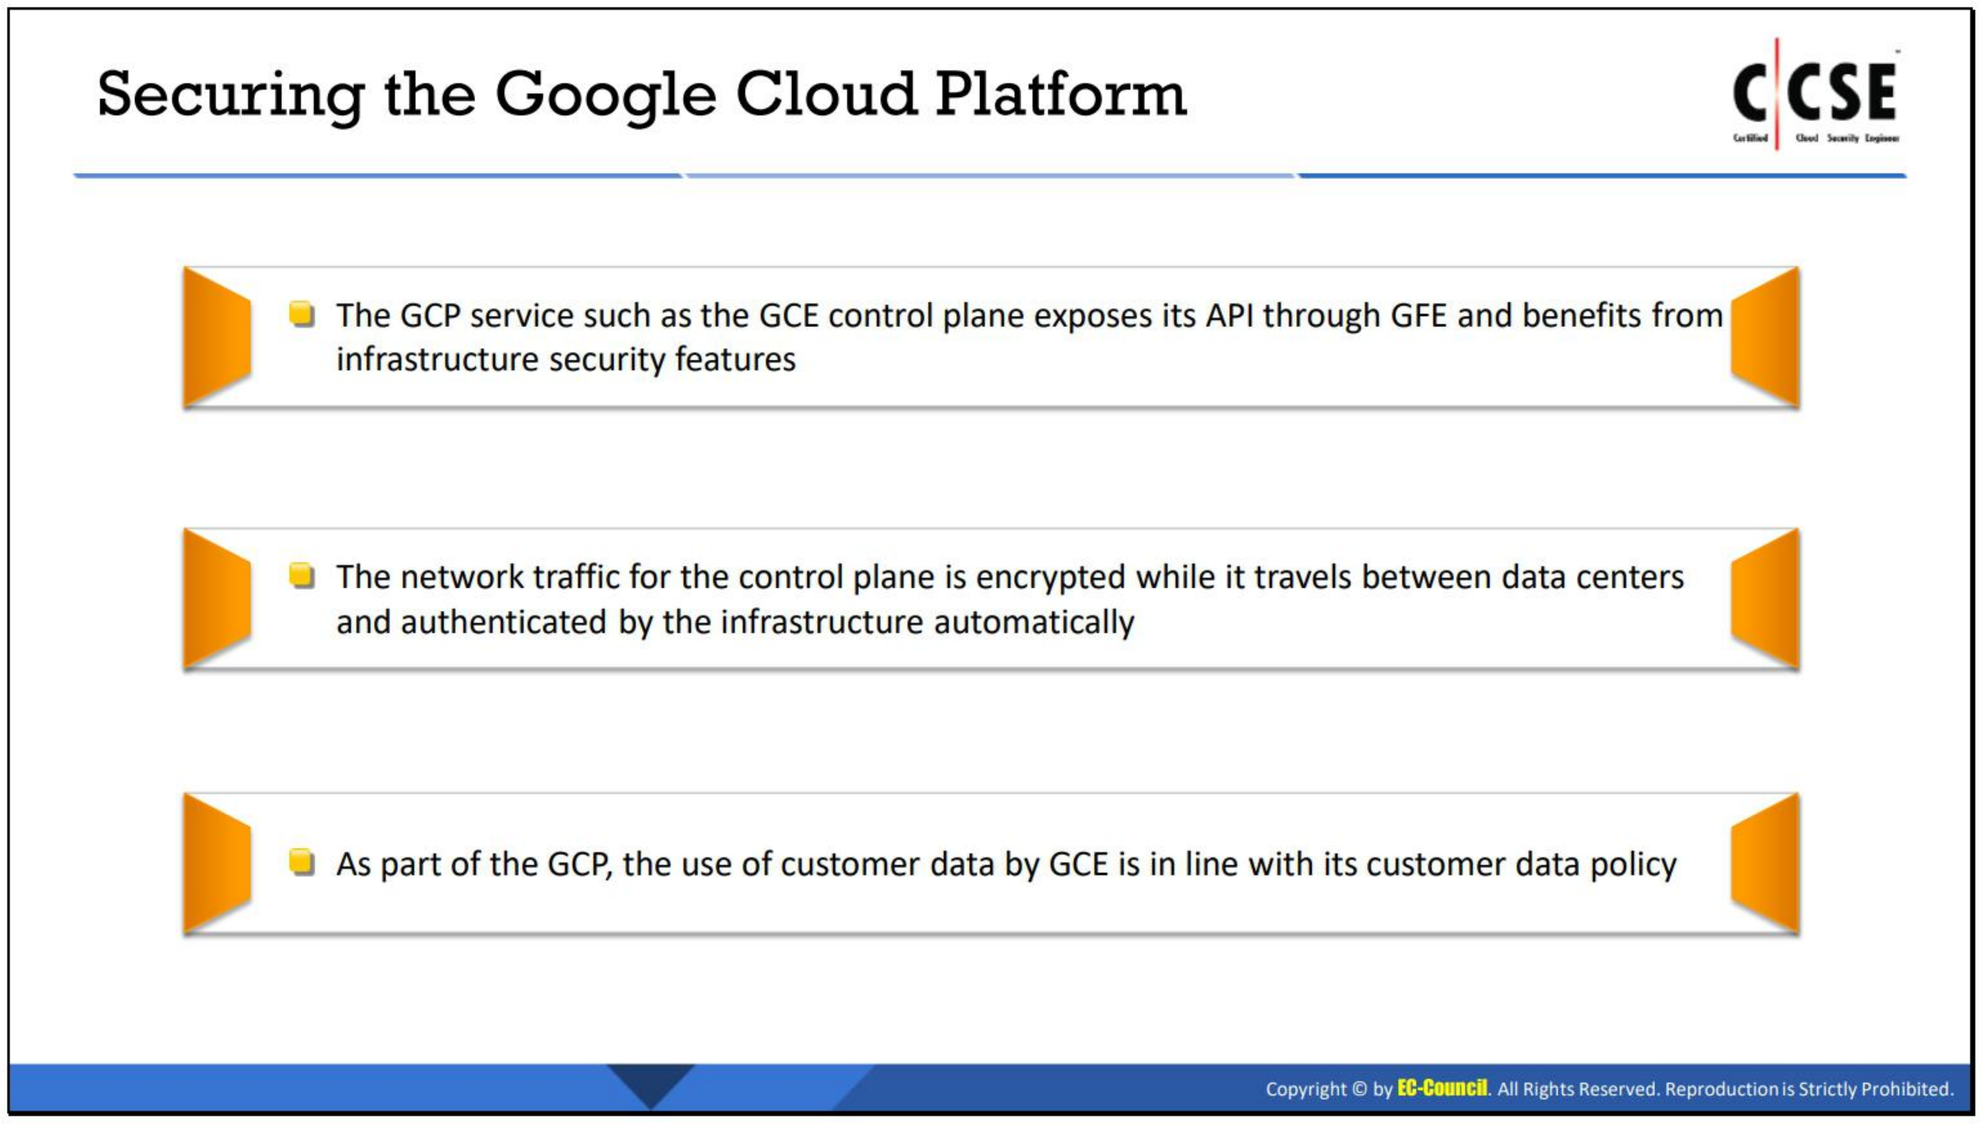
At the next layer, GFE instances report information about requests (they get) to the central DoS service. The central DoS service then can configure these instances also to throttle or drop the attack traffic.

User Authentication: Google's central identity service provides the next layer of defense after DoS protection. It generally appears to end-users as the Google login page. Apart from asking just a username and password, this service also challenges users intelligently for more information based on risk factors such as whether they logged in from a similar location or the same device in the past. Once the user is authenticated, this service issues credentials such as OAuth tokens and cookies that are used for future calls.

Users are given the option of using second factors, for example, OTPs or phishing-resistant security keys, while signing in. Additionally, Google has worked in the FIDO Alliance with many device vendors to develop the U2F (universal 2nd factor) open standard. The devices that support this standard are available in the market and other major web services are also implementing U2F support.

hide01.ir

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



Securing the Google Cloud Platform

GCP benefits from the security of the underlying infrastructure. Consider Google Compute Engine (GCE) as an example service. It allows customers to run their VMs on Google's infrastructure. Many logical components are involved in the GCE implementation. The most notable logical components involved include VMs themselves and the management control plane.

The management control plane orchestrates tasks such as VM creation and migration and exposes the external API surface. On the infrastructure, GCE runs as a variety (collection) of services so it gets foundational integrity features like a secure boot chain automatically. To grant each service only the permissions it requires while making remote procedural calls to the rest of the control plane, individual services run under discrete internal service accounts. The central Google source code repository stores the code for all these services and an audit trail is conducted between this code and the binaries which are finally deployed.

Through the GFE, the GCE control plane exposes its API so it benefits from the infrastructure security features such as centrally managed SSL/TLS support and DoS protection. Customers can obtain similar protections for applications that run on their GCE VMs by opting to utilize the optional Google Cloud Load Balancer service built on top of the GFE and can prevent various types of DoS attacks.

Google's centralized identity service offers security features such as hijacking detection, thus, it is used to authenticate end-users to the GCE control plane API. The central Cloud IAM service is used for authorization.

The network traffic for the control plane (between other control plane services and from the GFEs to the first service behind it) is encrypted while it travels between data centers and

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

authenticated by the infrastructure automatically. Additionally, Google has configured its infrastructure to encrypt some of the control plane traffic inside the data center, too.

Every VM runs with a VM manager service instance associated with it. Two identities are provided to these services by the infrastructure. The VMM service instance uses one identity for its own calls and another identity for calls (it makes) on behalf of the customer's VM. This enables Google to further divide the trust placed in calls that come from the VMM.

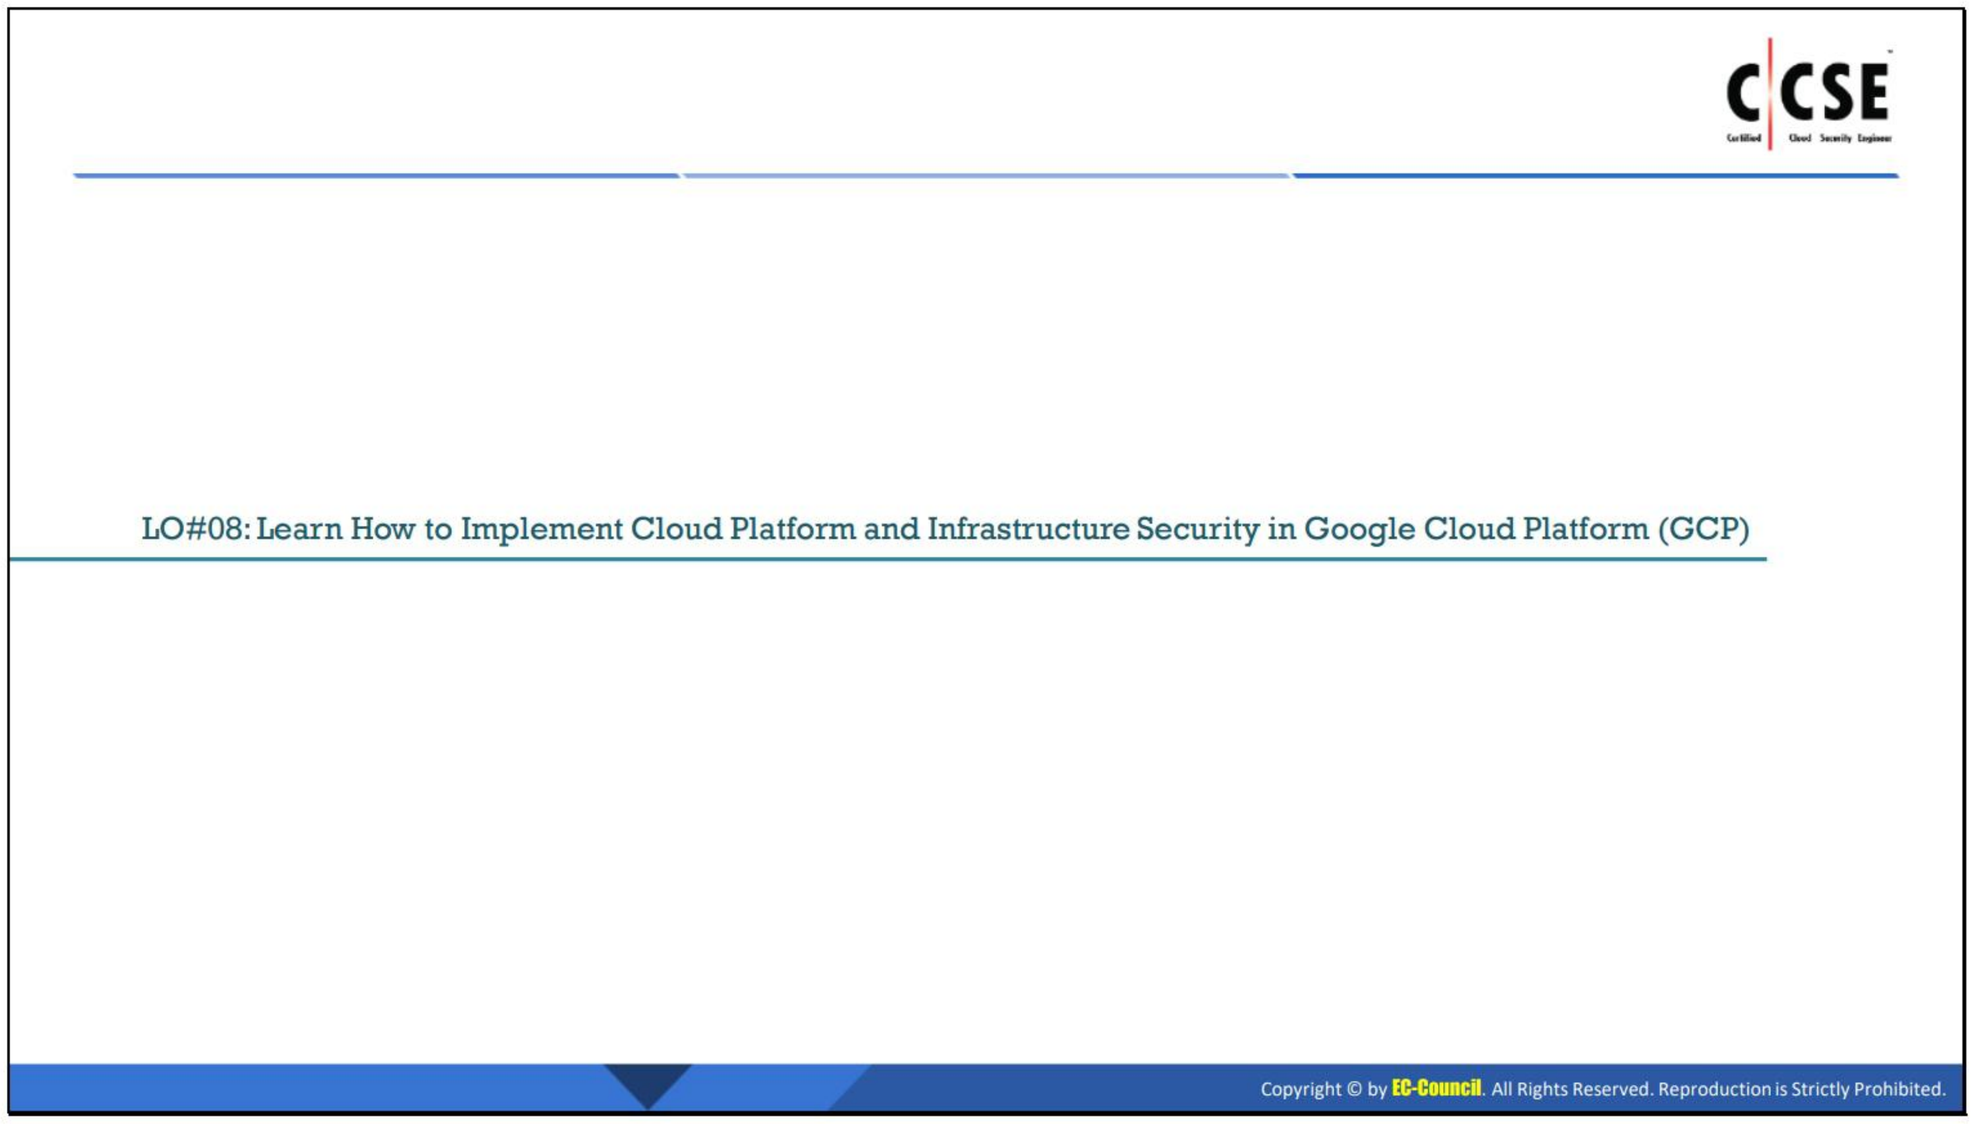
By using keys that are protected by the central infrastructure key management system, GCE persistent disks are encrypted at rest. This facilitates automated rotation and central auditing of access to these keys.

It provides customers with the option to send traffic from their VMs to other VMs or the internet without encryption or use any encryption they want for this traffic. For the WAN traversal hop of customer VM to VM traffic, Google has begun to roll out automatic encryption. Within the infrastructure, all control plane WAN traffic is already encrypted. To encrypt inter-VM LAN traffic inside the infrastructure, Google plans to make use of the hardware-accelerated network encryption in the future.

VMs are provided isolation based on hardware virtualization by using the open-source KVM stack. By shifting hardware emulation stack and some of the control into an unprivileged process outside the kernel, Google has hardened its particular implementation of KVM further. It has also tested the core of KVM extensively by using techniques such as static analysis, fuzzing, and mutual code review. The majority of vulnerabilities recently disclosed publicly that are upstreamed into KVM are discovered by Google.

Finally, Google uses its operational security controls as a key part to ensure that accesses to data are as per its policies. As part of the GCP, the use of customer data by GCE is in line with its customer data policy, that is Google shall not access or use its customer data except when necessary to offer services to customers.

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

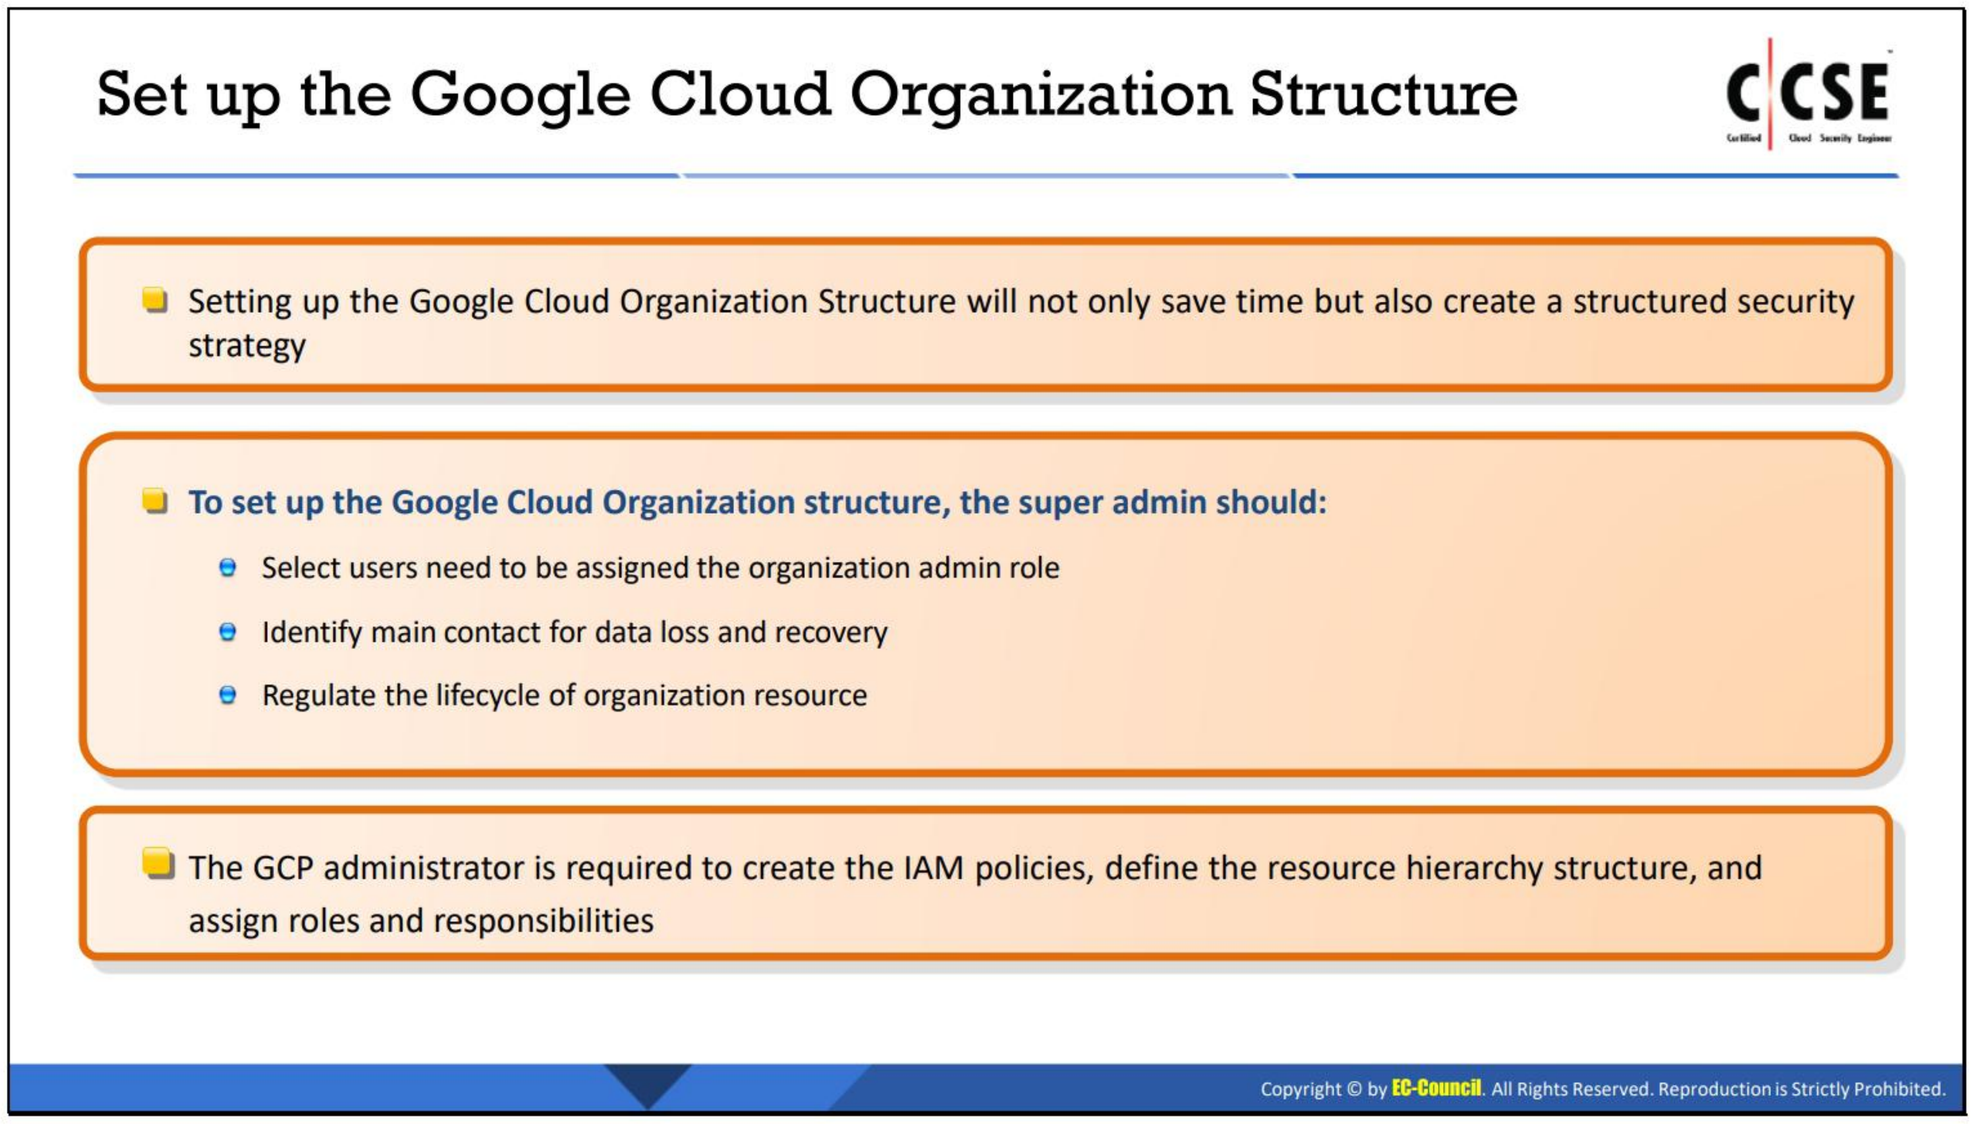


LO#08: Learn How to Implement Cloud Platform and Infrastructure Security in Google

This section shows you how to implement the security features within GCP to ensure secure storage of customer data through various in-built resources.

hidra1.ir

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



Set up the Google Cloud Organization Structure

A cloud administrator gets everything grouped or organized into a single organizational unit when they log in to Google Administrator Console. The settings applied to this organizational unit/group affect all devices and users in the organization. It is essential for a cloud administrator of an organization to plan how to manage their organizational units and hierarchy. For this, they need to set up the Google Cloud Organizational structure. Setting up the Google Cloud Organizational structure will not only save time but also create a structured security strategy.

The core part of Google Cloud (GC) hierarchy is an Organization resource. All Cloud Identity customers as well as Google Workspace super administrators have access to it. This Organization resource allows an enterprise to create a structure for teams and projects inside the enterprise.

To set up the GC Organization structure, the super administrator should

- Assign the Organization administrator roles to suitable users.
- For data loss and recovery, serve as the main contact.
- Regulate the lifecycle of organization structure.

Later, the GCP Organization administrator must create IAM policies, set up the Resource Hierarchy structure, and assign roles and responsibilities.

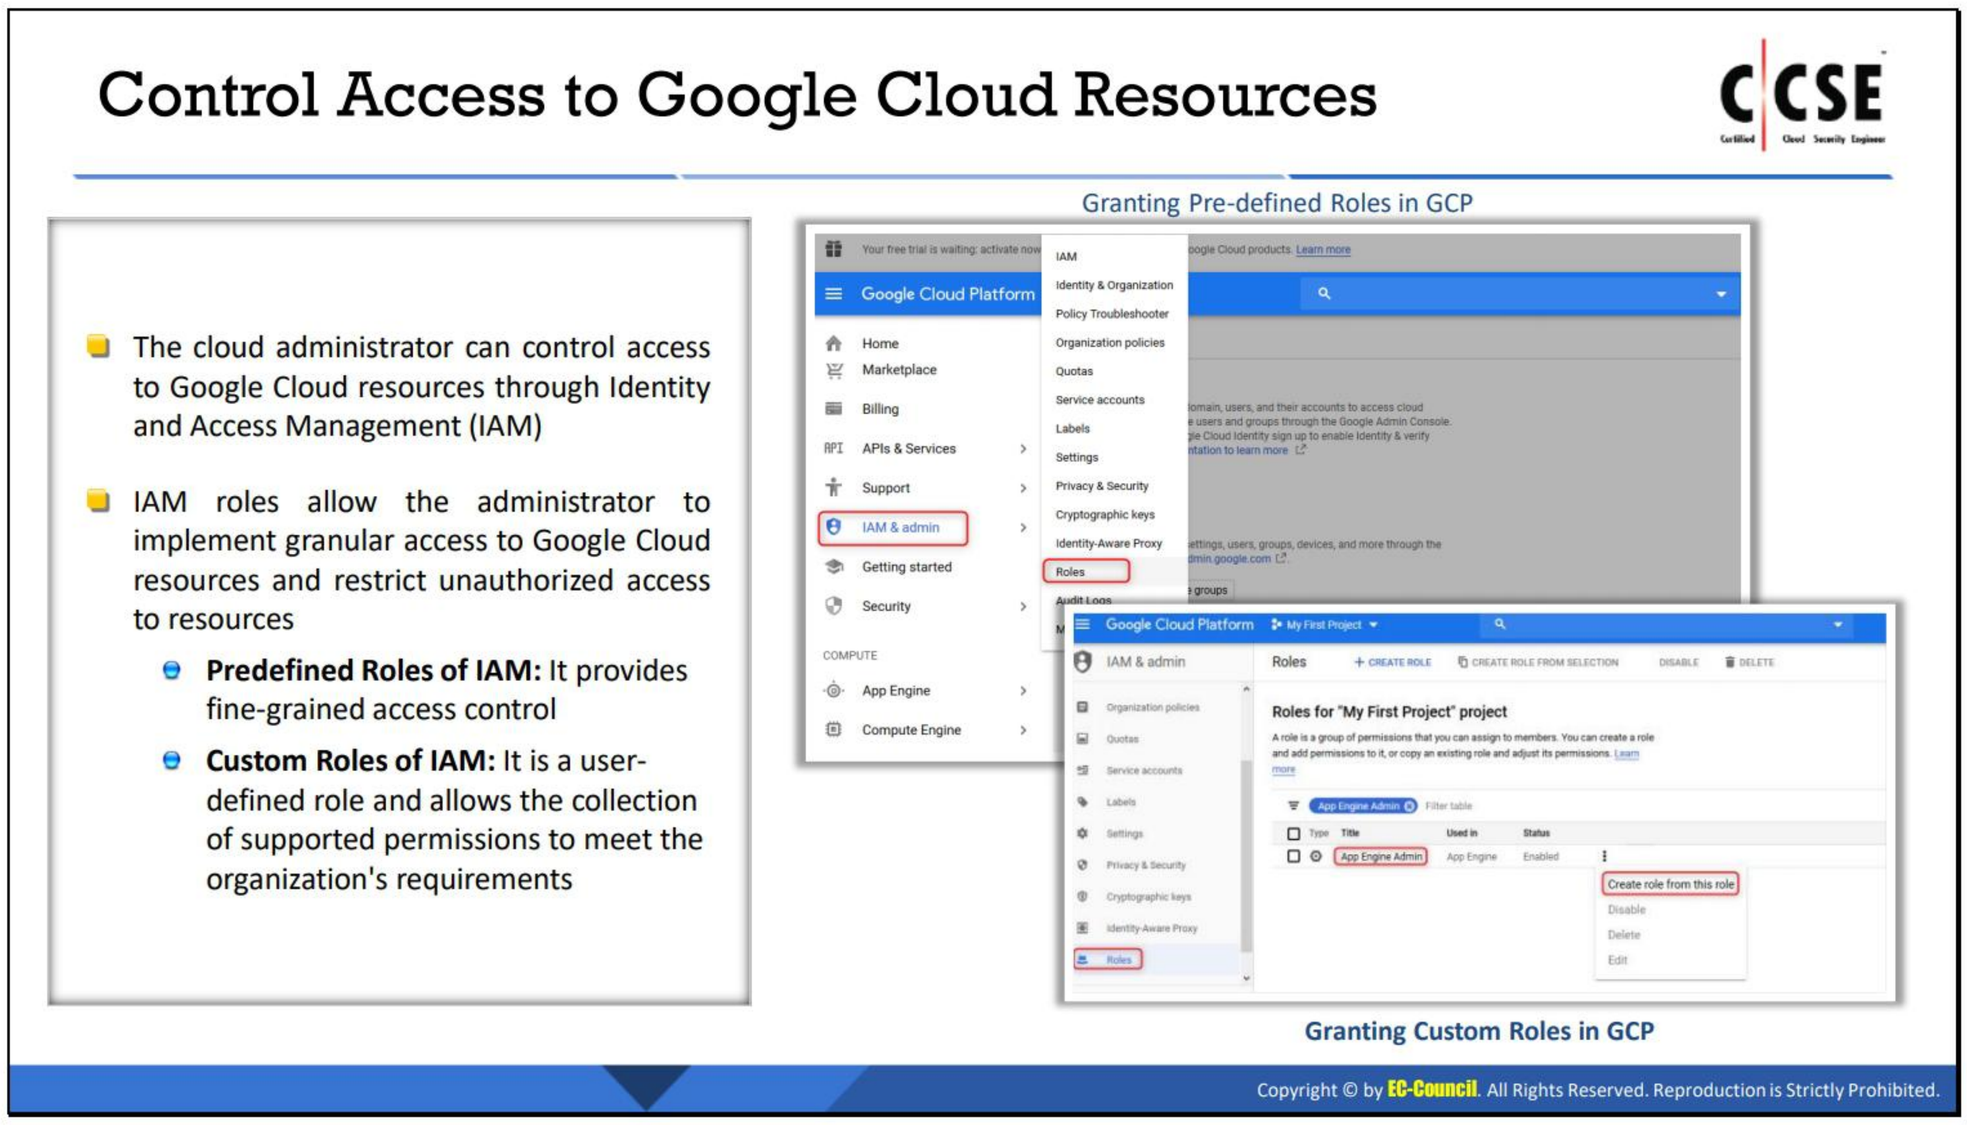
Certain important things to consider for structuring an enterprise on GCP.

- **Ensure Setting Up Identity with Roles and Users**: Always implement the least privilege policy and never allow normal users to create projects at the organization-level. Grant this ability to only Organization, DevOps, and Billing administrators. Never add individuals to the IAM. Instead, grant necessary privileges to groups and roles. This can help in controlling quotas and access to enable easier cost control.

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

- **Utilize Shared VPC to Centralize Networks**: Utilize Shared VPC for centralizing networks and projects in the organization. Encapsulate service projects by a host project and ensure that network configuration is centralized at one place. Allow only network administrators to access these projects.
- **Organize Resources on GCP**: Organize all organizational resources on GCP for easier maintenance of users and groups.
- **Tag Resources**: Ensure that all resources are tagged properly to avoid confusion while reporting, accessing, cost control and billing.
- **Setup Monitoring Alerts**: Use the in-built features provided by GCP and set up monitoring alerts to ensure a secure cloud environment.
- **Secure Data**: Though GCP provides encryption service for data-at-rest, identify the encryption key needs (BYOK, self-managed, or Google-managed) for securing data-in-transit and data-at-rest. Use VPC Service Controls to isolate access to data inside the organization.
- **Make Use of Cloud Security Command Center**: Utilize the service of Cloud Security Command Center provided by GCP for threat and anomaly detection, asset management, vulnerability scanning, and WAF.
- **Test the Documented BC/DR Plans**: Document the BC/DR plans to prevent any business interruption and data loss during disasters. Test the documented plans every year or more often to ensure that the plans are reliable and can work out in such situations without any issues. Use an automated backup and restore approach for better implementation of BC/DR plans.
- **Partner with Industry Experts**: It is recommended to be in touch with industry experts who have good exposure in handling cloud-related issues effectively.

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



Control Access to Google Cloud Resources

GCP IAM provides granular access to a user with principle of least privilege (PoLP) security to essential Google Cloud resources. The cloud administrator can control access to Google Cloud resources through IAM.

hide01.ir

Users must not be assigned access privileges individually. The administrator should assign permissions to Roles rather than Users in GCP. IAM Roles are a set of permissions that allow users under them to perform certain actions on certain cloud resources.

GCP IAM not only offers various predefined roles for handling a large number of common use cases but also provides the administrator with an option to create custom roles for handling roles that require customized access privileges.

The IAM roles allow the administrator to grant granular access to specific Google Cloud resources and prevent unwanted access to other resources.

- **Predefined Roles in IAM**: They provide fine-grained access control and are meant specifically to support common control patterns and use cases.
- **Custom Roles of IAM**: They are user-defined roles and allow the collection of supported permissions to meet an organization's requirements. They provide granular access to cloud resources based on the user-specified list of permissions. Because Google do not maintain custom roles, these roles are not updated automatically by Google Cloud. It is possible to create custom roles at the organization and project level; however, it is not possible to create the roles at the folder level. When predefined roles do not meet the need of the organization, then custom roles that are created and used to meet the security requirements.

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

Granting Predefined Roles in GCP

- From Google console drop-down menu, navigate to **IAM & admin,** select **Roles** in the menu.

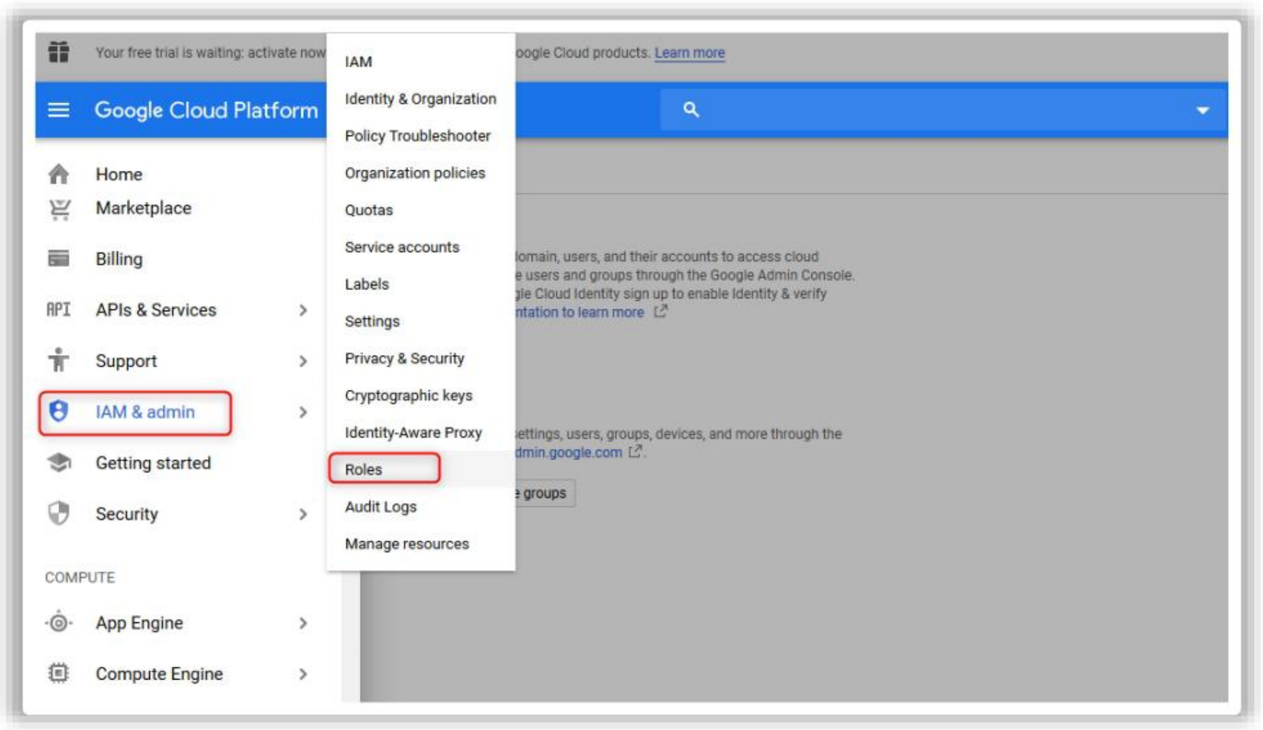


Figure: 2.49: Choosing Roles from the Menu

- Select the filter from filter table (**Name**) and search the role (**actions.Admin**). Click the context menu to see the assigned permissions.

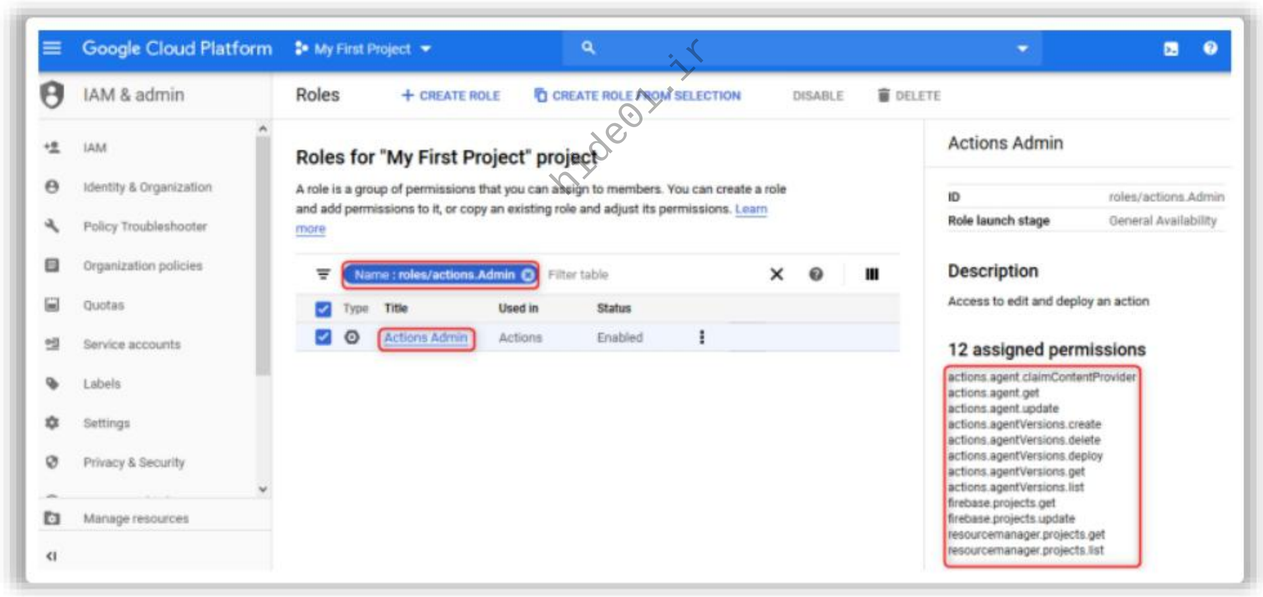


Figure: 2.50: Searching for Action.Admin Role

- Click **IAM** on the Menu tab, select **Add** to assign role, enter email ID of the user, select Actions Administrator from the **Role** drop-down menu and click **Save.**

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

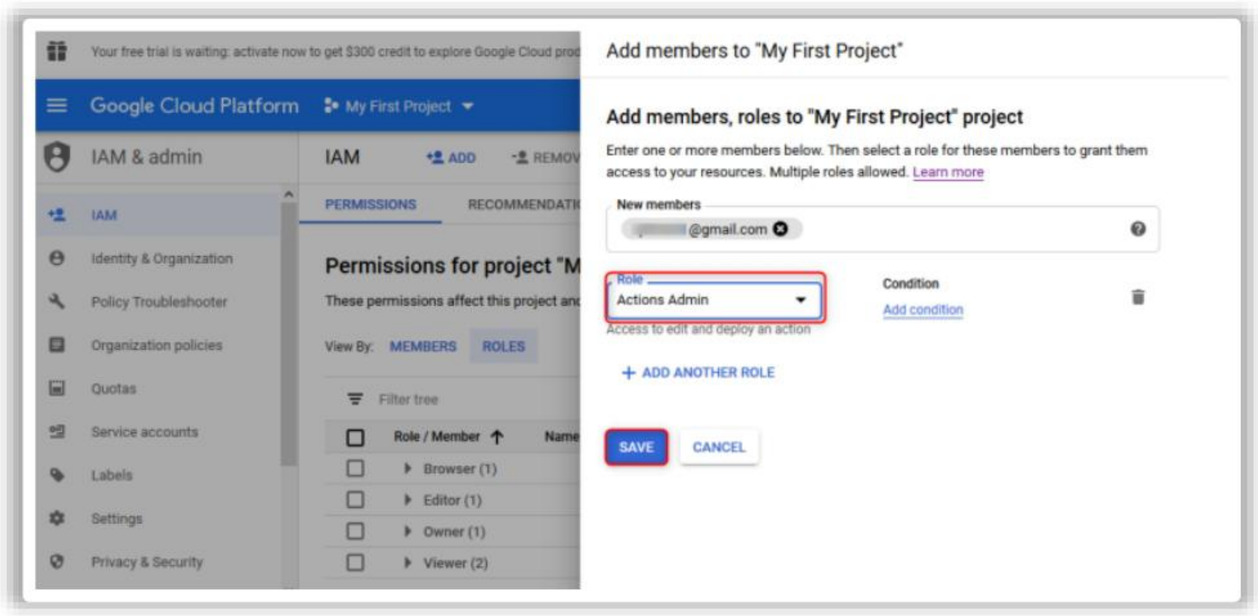
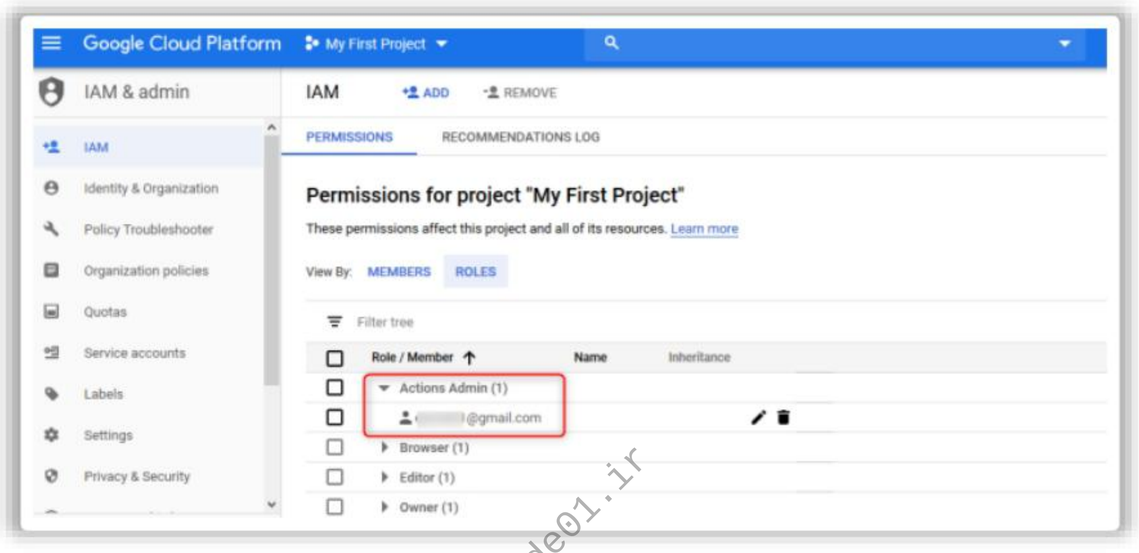


Figure: 2.51: Adding Roles and Users

- Predefined role is granted in GCP.



hide01.ir

Figure: 2.52: Grating Predefined Role to GCP

Granting Custom Roles in GCP

- From Google console drop-down menu, navigate to **IAM & Admin,** select **Roles** in the menu.
- Select the role from filter table, click the context menu and select **Create role from this role.**

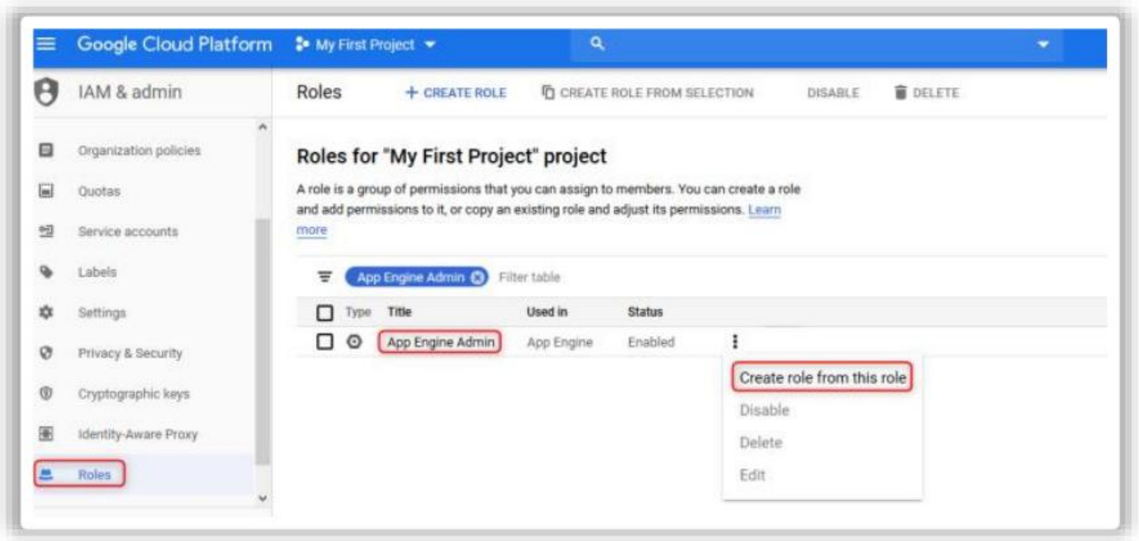


Figure: 2.53: Granting Custom Roles in GCP

- Create custom role, click **Add Permissions.**

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

- Select the required permissions and Click **Add.**

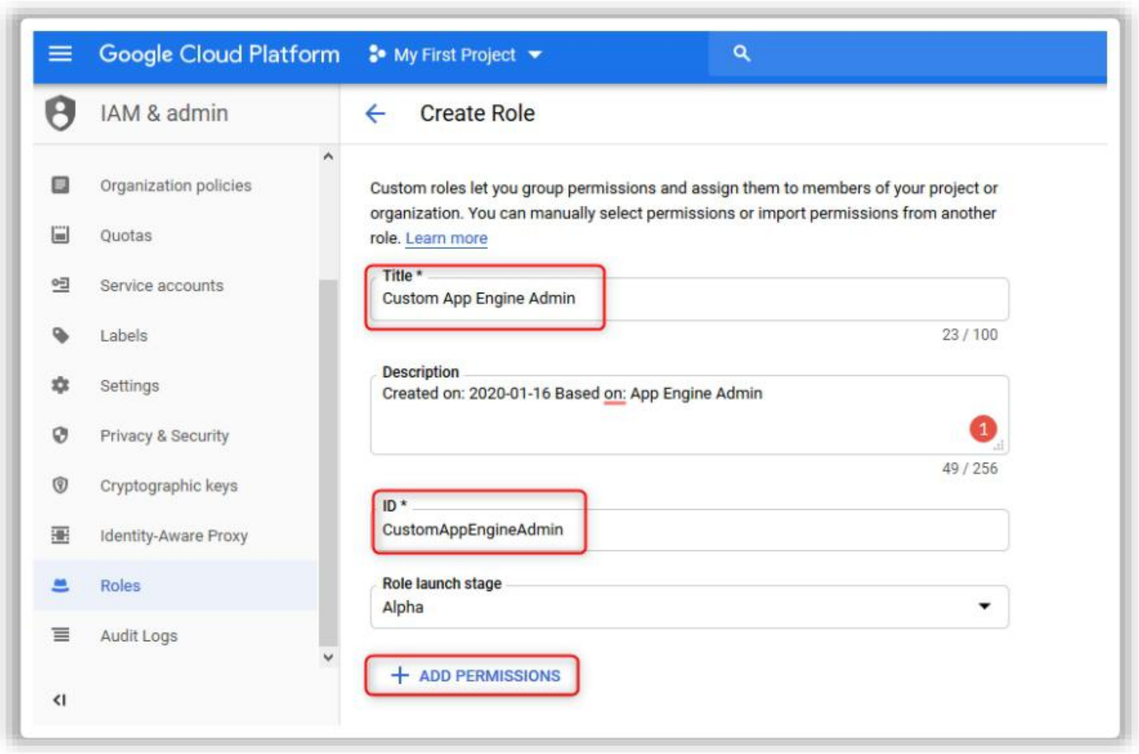


Figure: 2.54: Selecting Required Permissions

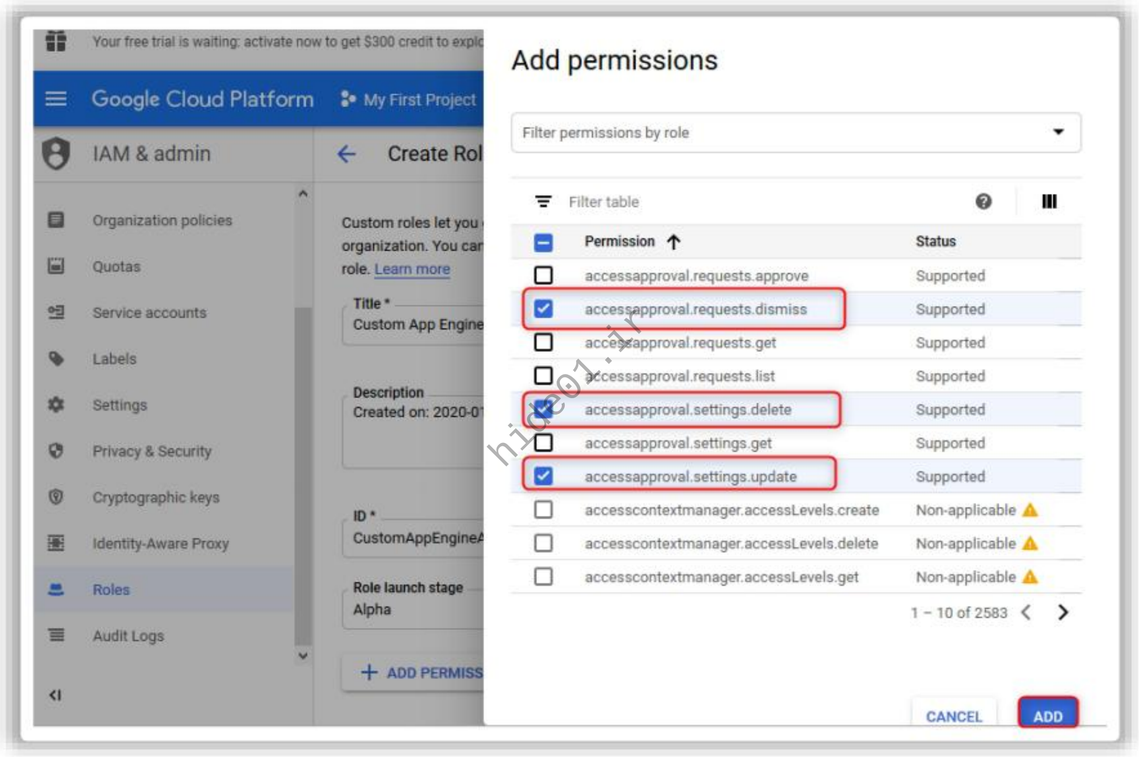


Figure: 2.55: Adding Required Permissions

- The custom App Engine Administrator roles are created.

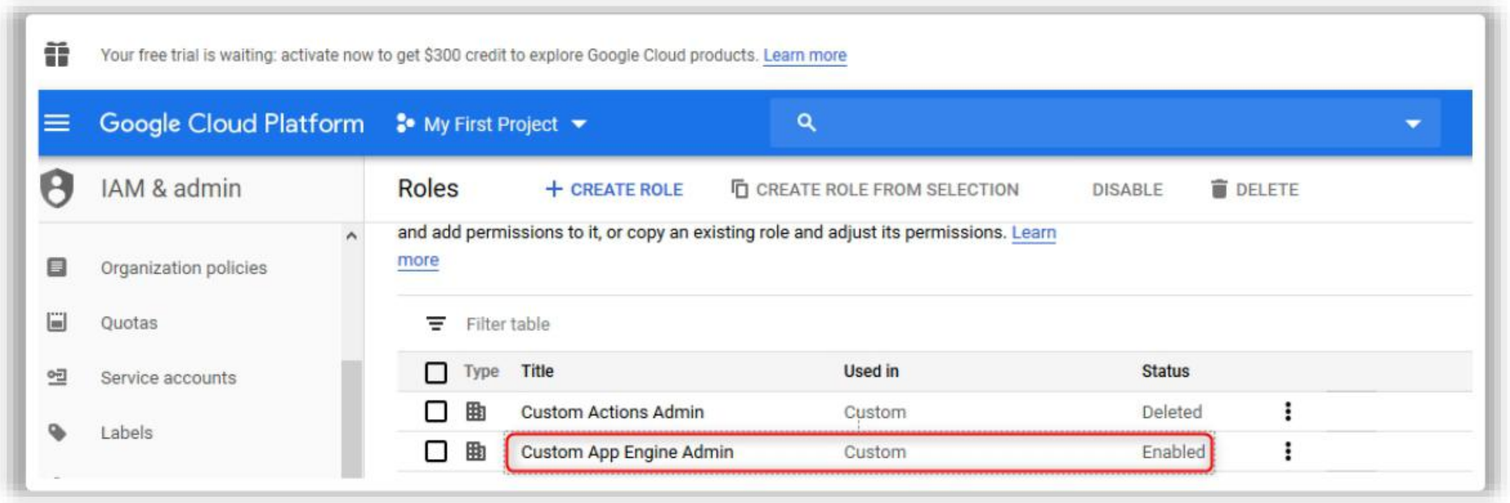


Figure: 2.56: Successful Creation of Custom App Engine Administrator Roles

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

- Click **IAM** on the **Menu** tab, select **Add** to assign role, enter email ID of the user, select Custom App Engine Administrator from the **Role** drop-down menu and click **Save.**

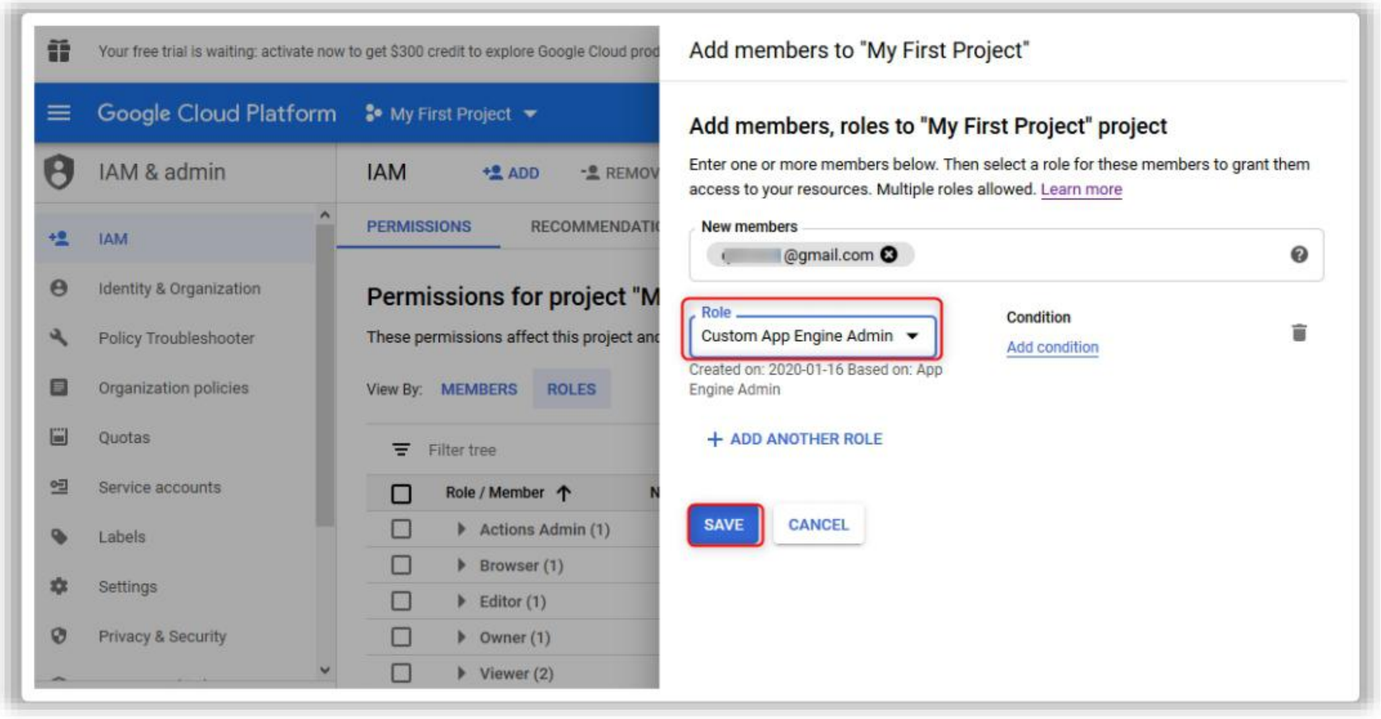


Figure: 2.57: Assigning Roles and Members to the Project

- Custom roles are granted to the user.

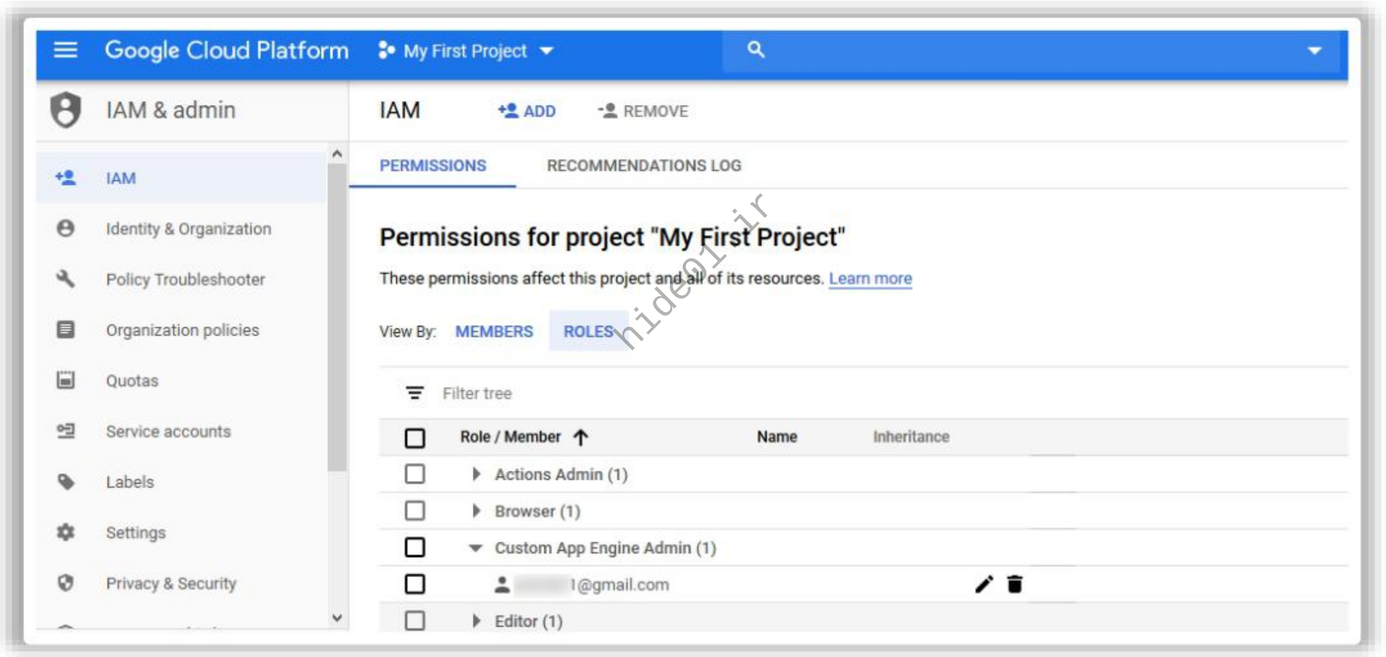
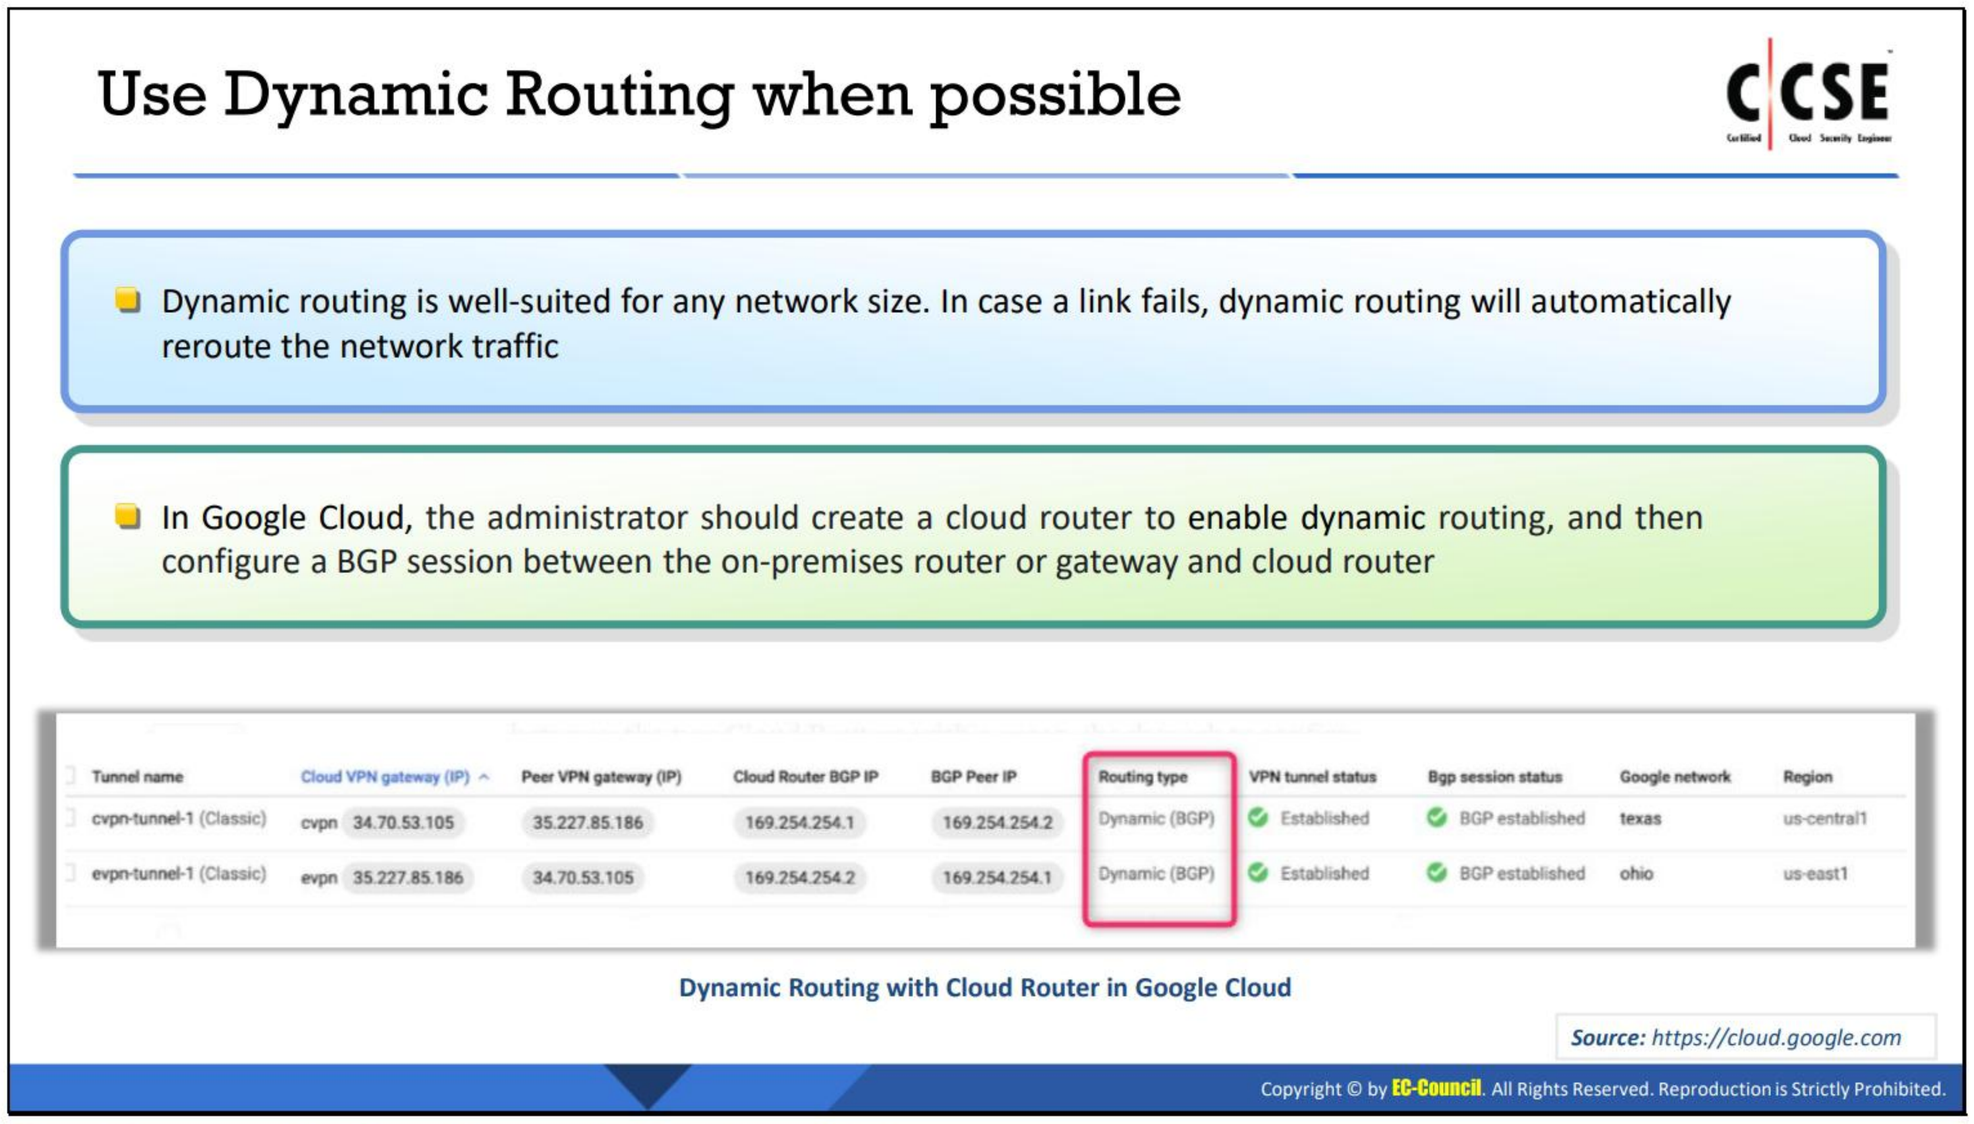


Figure: 2.58: Roles Granted to Users

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



Use Dynamic Routing when Possible

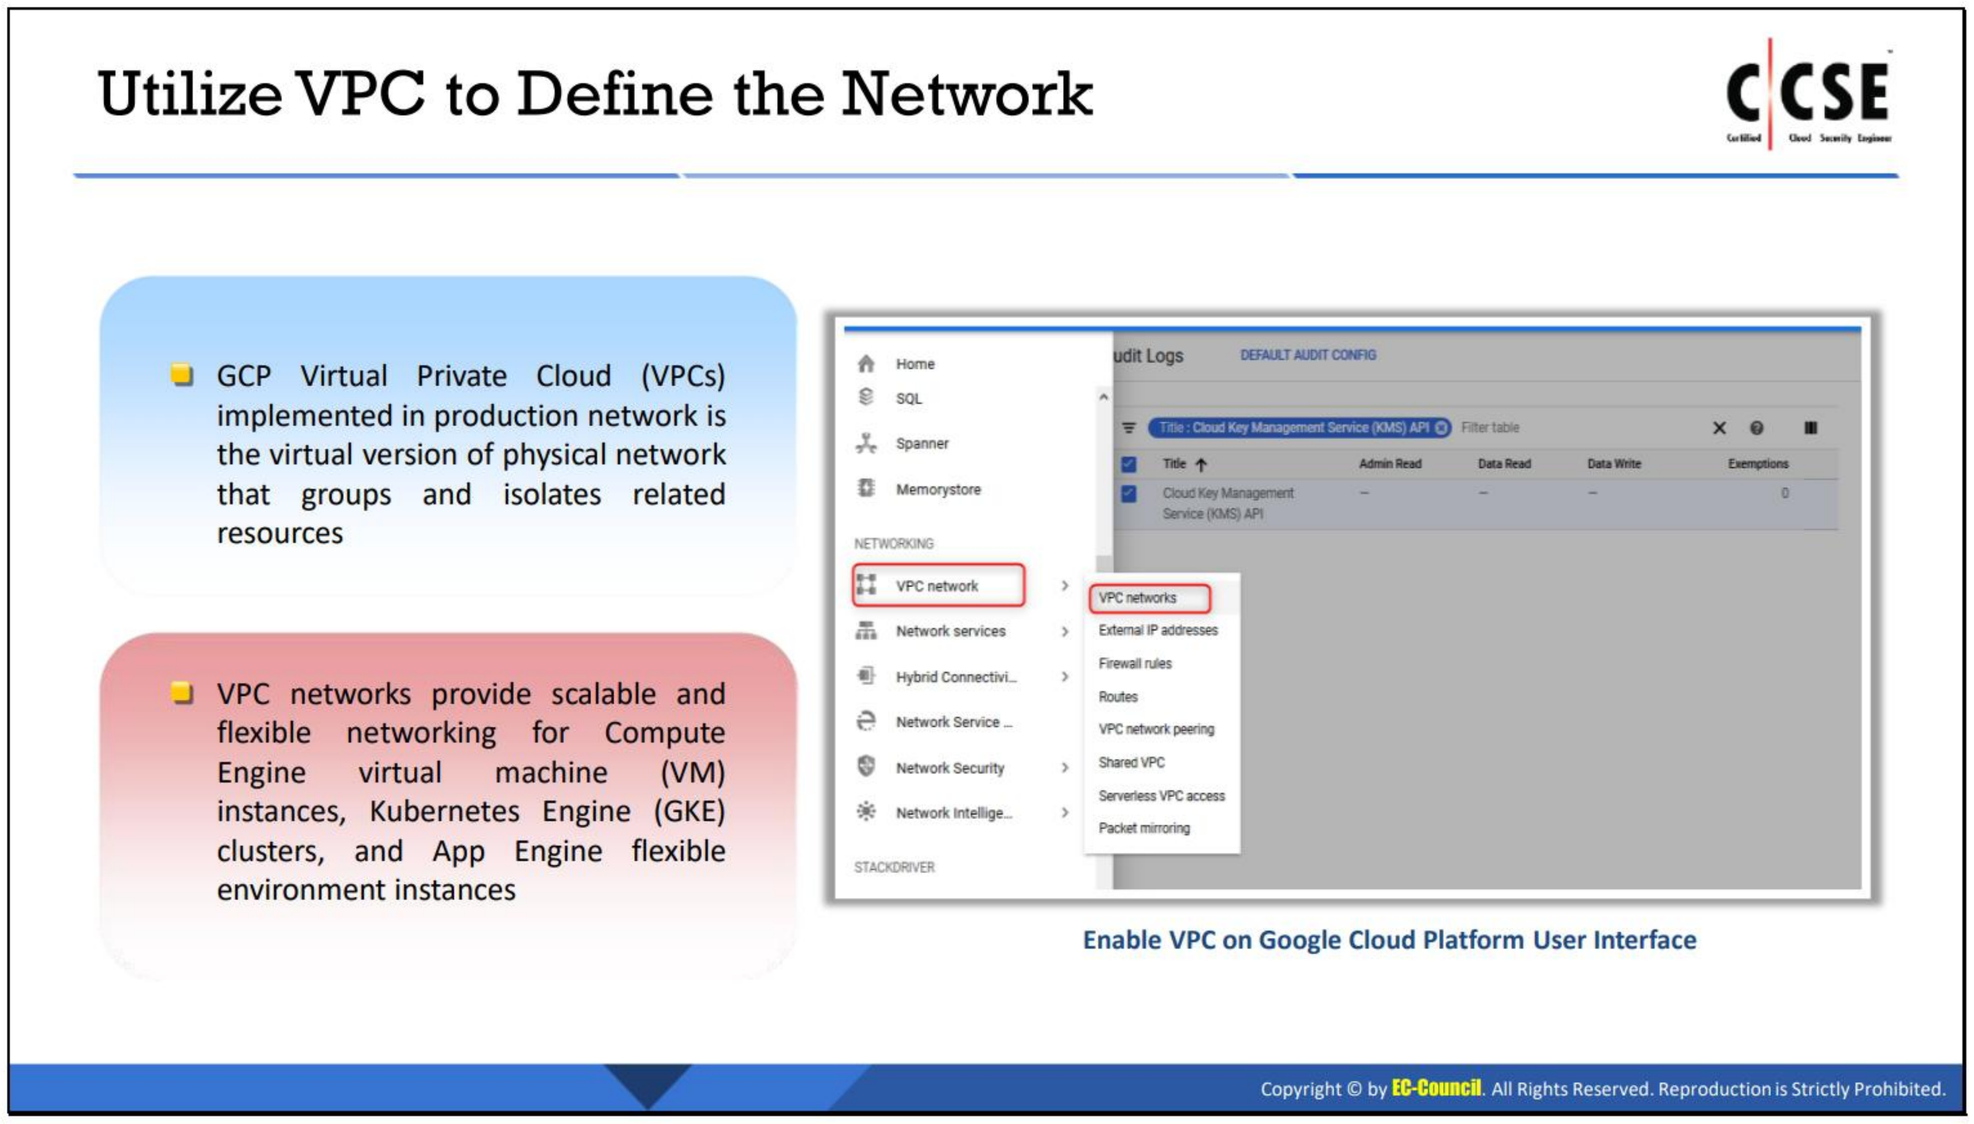
Cloud Router (CR) is a service that is fully managed and distributed by Google. CR creates custom dynamic routes and scales (up or down) with the traffic. It works with VPC networks as well as legacy networks. It uses border gateway protocol (BGP) to offer dynamic routing over Cloud VPN or interconnect connections. It is used to dynamically, securely exchange routes between on-premises networks and GCP networks. CR connects with on-premises VPN gateways or routers that use BGP to exchange topology information.

While using static routing, the administrator must create and manage a routing table. Static routing requires them to manually change and update static routes if there is any change in topology. Moreover, static routing requires strict control over routing tables and is suitable only for small networks that have stable topology. Dynamic routing is well-suited for any network size. In case, a link fails, then dynamic routing will automatically re-route the network traffic. Unlike dynamic routes, static routes are not fault-tolerant and require manual intervention if any problem occurs with them. They work well with systems and networks that are always stable.

Dynamic routing involves advertising route information to all routers on the network, which helps them to decide on the best route to exchange information. Routers make intelligent decisions on how to send the data from a source to the intended destination by using a routing protocol (BGP). It provides redundancy of routing paths. When one route fails, the data gets re-routed to the destination through an alternate route.

The administrator can exchange routing information with the help of CR using BGP between networks. BGP helps networks to automatically and immediately discover any topology changes, which eliminates the need of manually configuring static routes. Exchanging routes using BGP is referred to as dynamic routing. In Google Cloud, the administrator should create a CR to enable the dynamic routing, and then configure a BGP session between the on-premises router or gateway and CR.

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



Enable VPC on Google Cloud Platform User Interface

Utilize VPC to Define the Network

A VPC network is a virtual representation of physical network, which connects VM instances, GKE clusters, App Engine, and other resources of a project. A project could have multiple VPC networks depending on the set organizational policy.

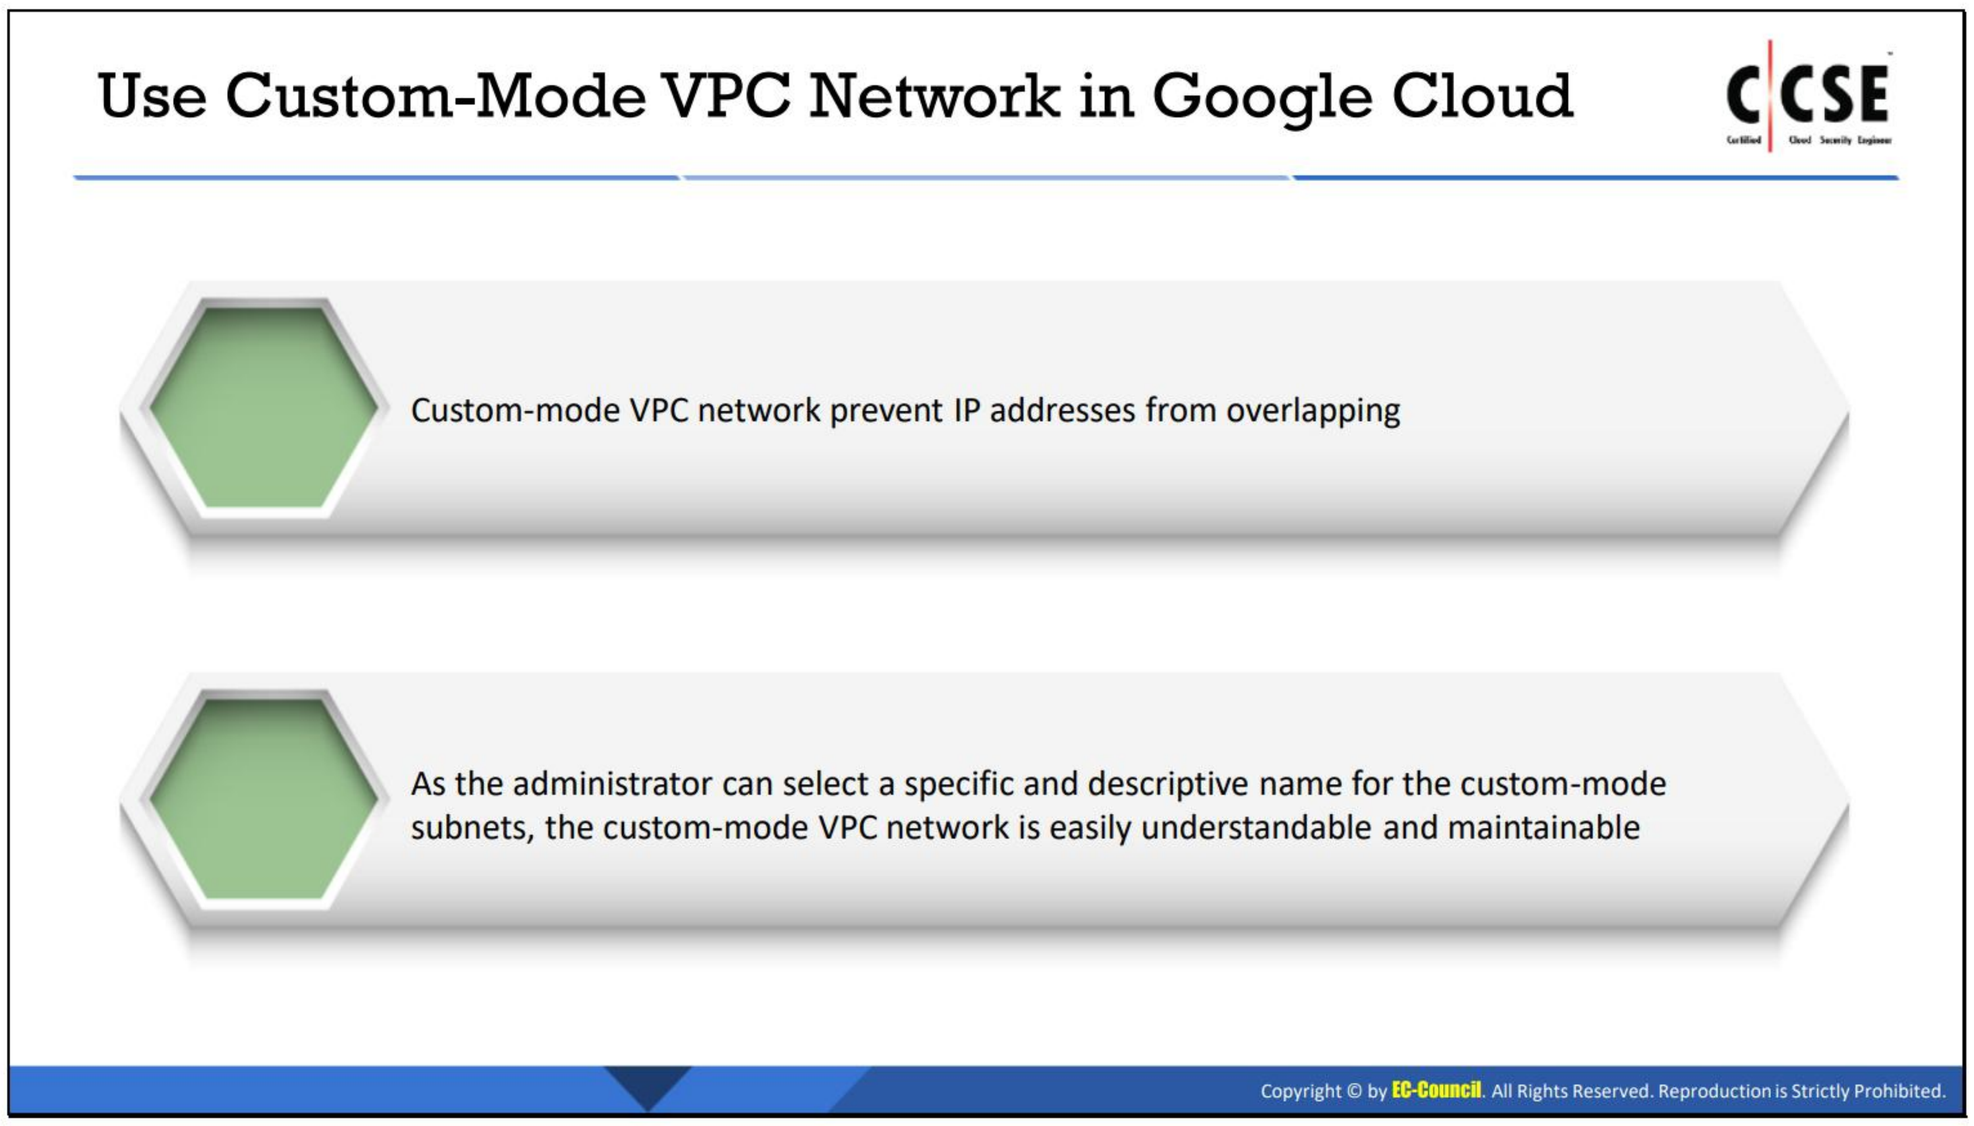
VPC networks are not associated with a particular zone or region and are global resources. Each VPC network contains subnets that define a set of IP addresses. In a VPC, network firewall rules control traffic that moves in and out of instances. Resources present inside a VPC network use IPv4 addresses to communicate with one another, as per the applicable firewall rules.

In a hybrid environment, the administrator can use Cloud Interconnect or Cloud VPN to securely connect VPC networks. To connect VPC networks from different projects, use VPC network peering. Use IAM roles to secure network administration in VPC networks.

Following are the security benefits offered by VPC networks:

- Use GCP VPC and subnets to map the network, and to group and isolate related resources.
- VPC networks provide scalable and flexible networking for the Compute Engine VM instances.
- A VPC network provides internal TCP/UDP Load Balancing.
- It establishes secure connection with on-premises networks via cloud interconnect attachments and cloud VPN tunnels.
- It helps in distributing traffic from Google Cloud external load balancers to backend.

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



Use Custom-Mode VPC Network in Google Cloud

A VPC network comes with two modes: auto mode or default mode and custom mode. VPC auto mode networks create subnets and subnets routes automatically in each Google Cloud region by using a set of RFC 1918 address ranges and they include some pre-populated firewall rules automatically. As every organization has different security requirements, custom-mode VPC networks suit them well. With custom-mode VPC networks, organizations can have better control over their networks with enhanced security. Default security controls in auto mode VPC networks are often predictable and can be abused by attackers; thus, it is better to create a custom-mode VPC network.

Organizations must use custom-mode VPC networks because

- They integrate well into present IP address management schemes. Unlike auto mode VPC networks, custom-mode VPC networks prevent IP address overlapping.
- Two auto mode VPC networks cannot be connected using VPC network peering because their subnets have the same primary IP addresses.
- The administrator can select specific and descriptive names for the custom-mode subnets; thus, they are easily understandable and maintainable.
- With custom-mode VPC networks, the administrator gets better flexibility for planning and avoiding overlapping addresses.

The administrator can delete a default VPC network after creating a custom-mode VPC network. However, the default or auto mode VPC network cannot be removed or deleted unless all resources that depend on it are removed. Minimize communication using public IP addresses in

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

hybrid cloud environments to keep them secure and use VPC peering or shared VPC for all internal communication among VPCs.

To create a custom-mode VPC network, the administrator must first create a VPC network and then create subnets inside a region. The subnets can be created while creating the VPC network or can be added afterward.

Creating a Custom-Mode VPC Network

Follow these steps to create a custom-mode VPC network:

- Sign-in to the GCP Console and Click on **VPC Network**.
- On the **VPC Networks** page, Choose **Create VPC network**.

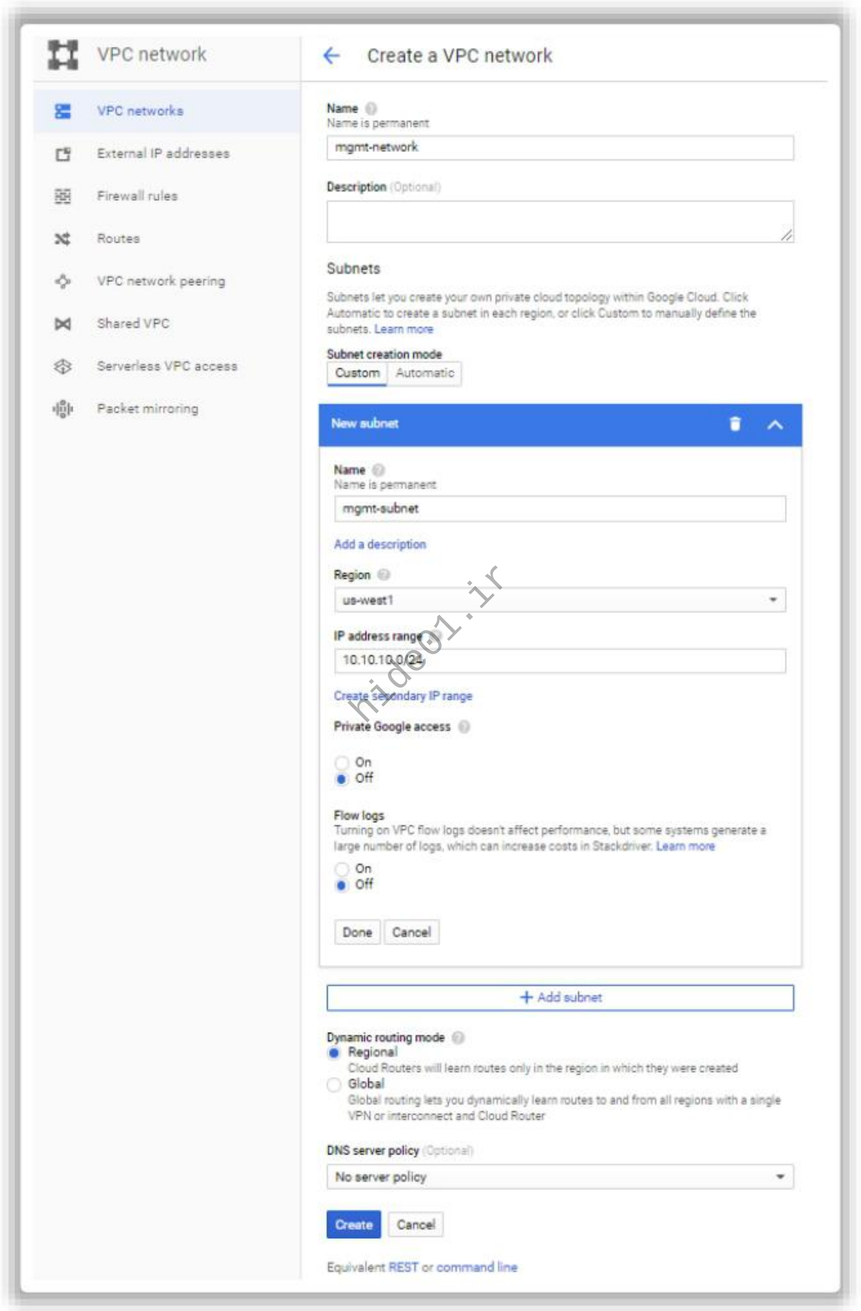
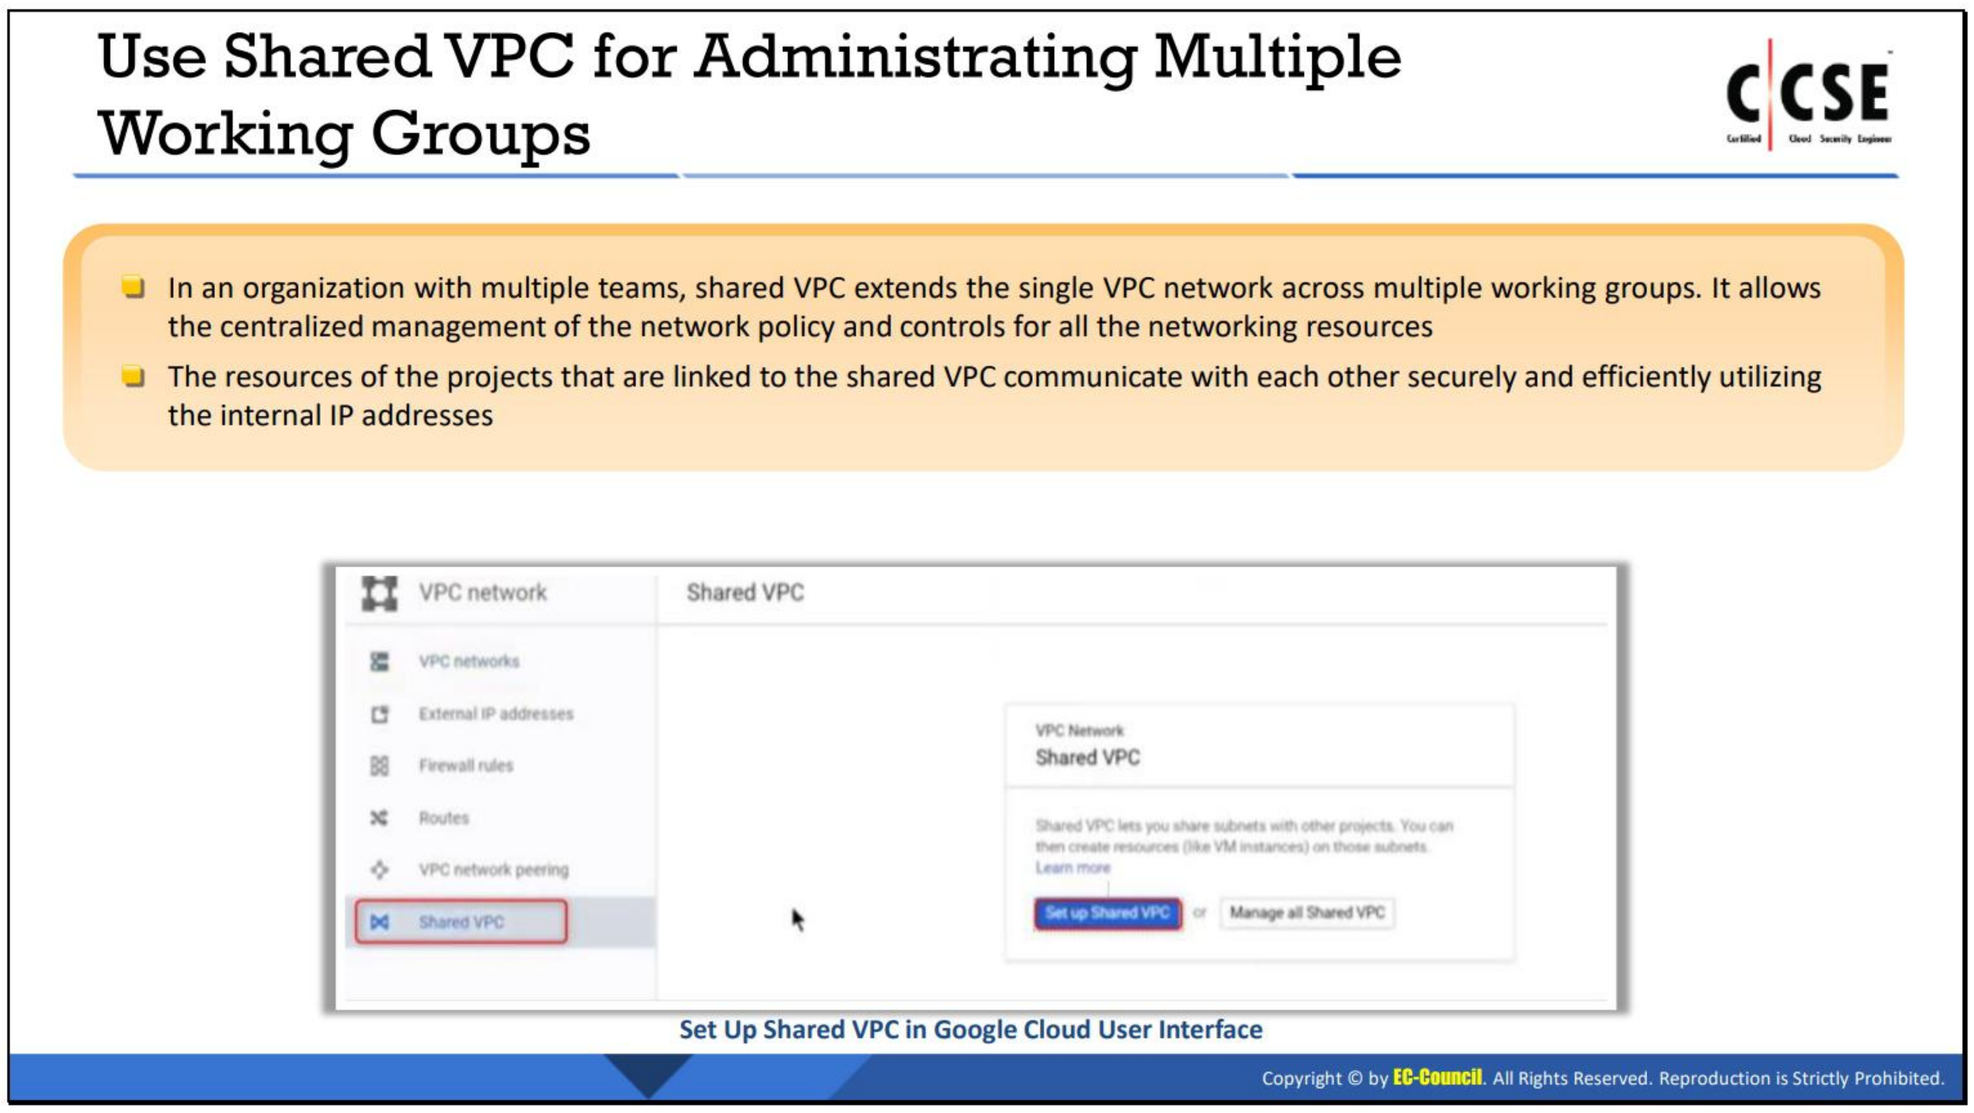


Figure: 2.59: Creating a VPC Network

- Provide a name in the **Name** field for the VPC network.
- Choose **Custom** as the **Subnet creation mode** under the **Subnets**.
- Specify the required configuration information in the **New subnet** field. After filling the information, click **Done.**
- Choose **Add subnet** to add more subnets and repeat Step 5.
- Select the **Dynamic routing mode** and click on **Create**.
- The custom-mode VPC network is created successfully.

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



Use Shared VPC for Administrating Multiple Working Groups

In organizations with multiple teams, Shared VPC extends the single VPC network across multiple working groups. It allows centralized management of network policy and control for all networking resources. Shard VPC helps enterprises to connect resources from several projects to a common VPC network. The resources of the projects that are linked to the Shared VPC communicate with each other securely and efficiently by utilizing the internal IP addresses.

While using Shared VPC, the administrator selects a project as a host project and assigns one or more service projects to it. Here, the VPC networks present in the host project are referred to as Shared VPC networks. Minimize communication using public IP addresses in hybrid cloud environments to keep them secure and use VPC peering or shared VPC for all internal communication among VPCs.

Shared VPC helps the administrator in various ways:

- It can be used to connect to a common VPC network.
- Shared VPC and IAM controls enable separating network administration from project administration.
- The administrator can centrally supervise the creation of firewalls, routes, subnet IP ranges, VPN connections, etc.
- The administrator can create security policies in the host project such as controlling which VMs can possess public IPs and can access the internet through having a single organizational policy.

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

- Organization administrators can assign administrative responsibilities (e.g., creating and managing instances) to service project administrators, while having centralized control over all network resources at the same time.

Setting Up Shared VPC in Google Cloud User Interface

- From the Google console drop-down menu, navigate to **VPC network** and select **VPC networks** on Menu Tab.
- Navigate to **Shared VPC**.

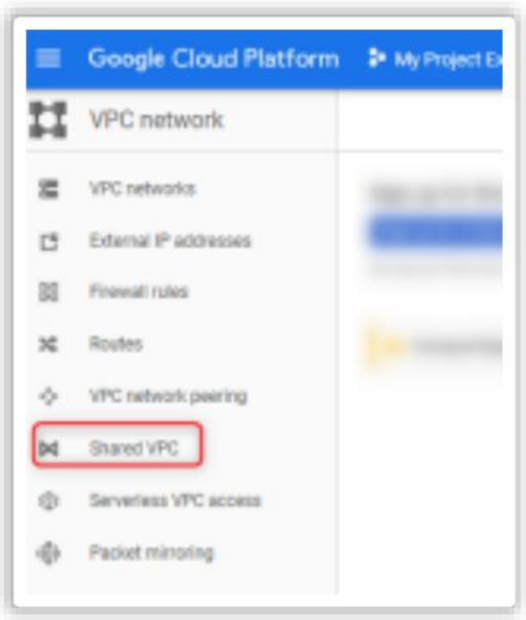


Figure: 2.60: Navigating to Shared VPC

- Select **Set up Shared VPC.**

hide01.ir

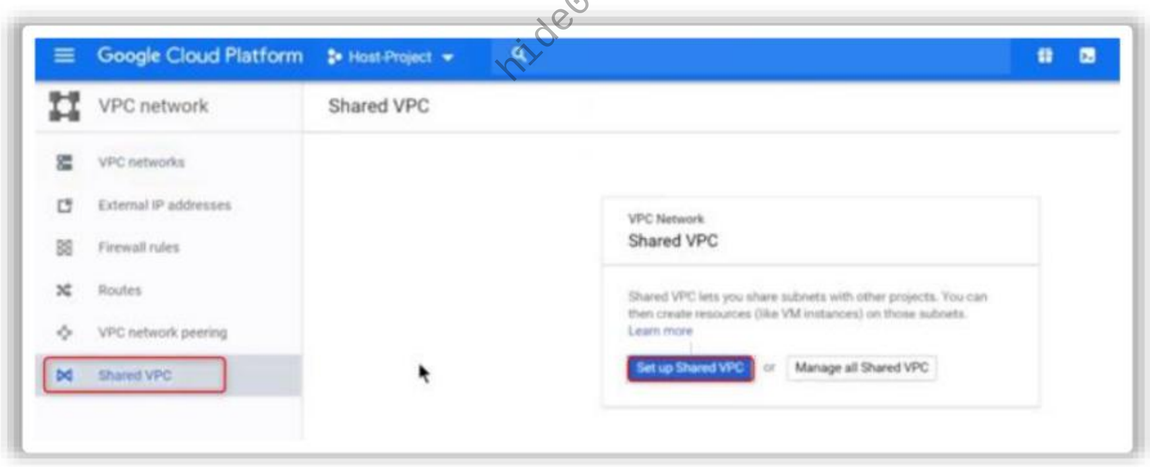


Figure: 2.61: Setting Up of a Shared VPC

- Click **Save & continue.**

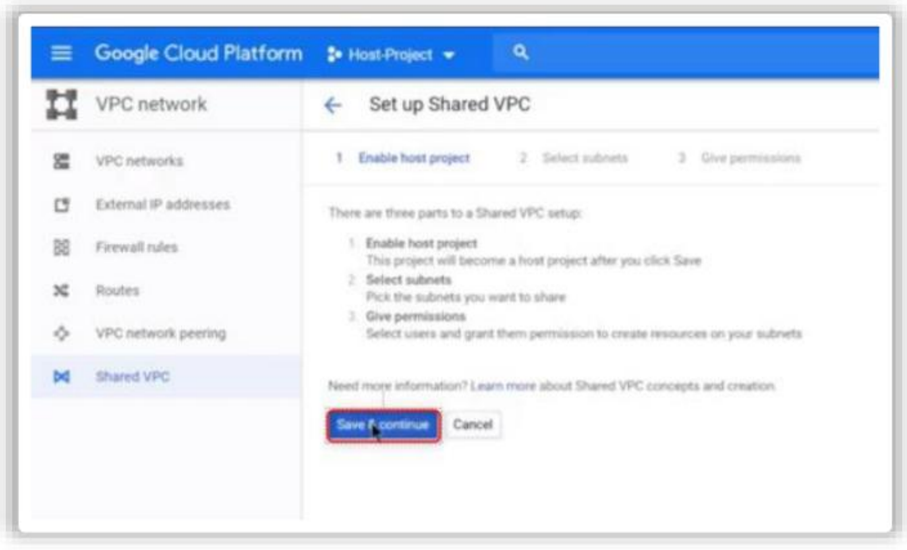


Figure:: 2.62: Saving the Changes of Shared VPC

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

Creating a Shared VPC Using Cloud Shell

- Using Cloud Shell, create a host project and two service projects; for example, prod and dev. Instead of [xxx], use a unique set of numbers.

```
gcloud projects create host-project-[xxx] – name=”demo-
   project” \
--enable-cloud-apis
gcloud projects create dev-project-[xxx] – name=”dev” \
--enable-cloud-apis
gcloud projects create prod-project-[xxx] – name=”prod” \
--enable-cloud-apis
```

- Now, replace the project ID with the hostproject ID.

```
gcloud config set project [hostproject ID]
```

- Make a custom VPC with two subnets

```
gcloud compute networks create vpc1 --subnet-mode=custom
gcloud compute networks subnets create dev --network=vpc1 –
   range=192.168.1.0/24
gcloud compute networks subnets create prod --network=vpc1
   – range=192.168.25.0/24
```

hide01.ir

- Next, create two firewall rules (one rule to allow SSH to all instances and the other rule to allow ICMP traffic between various instances in these subnets)

```
gcloud firewall-rules create allow-ssh --allow tcp:22 --
   source-ranges 0.0.0.0/0
gcloud firewall-rules create allow-icmp --allow tcp:80,icmp
   --target-tags dev,prod
```

Configuring Shared VPC

- On the left panel, choose **Networks** and then select **Shared VPC.**
- Choose **Set up Shared VPC** and click **Save & Continue**.
- Enable sharing mode on Individual Subnets.
- Choose the two subnets (**prod and dev**) as the subnets that are to be shared and click on **Continue**.

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

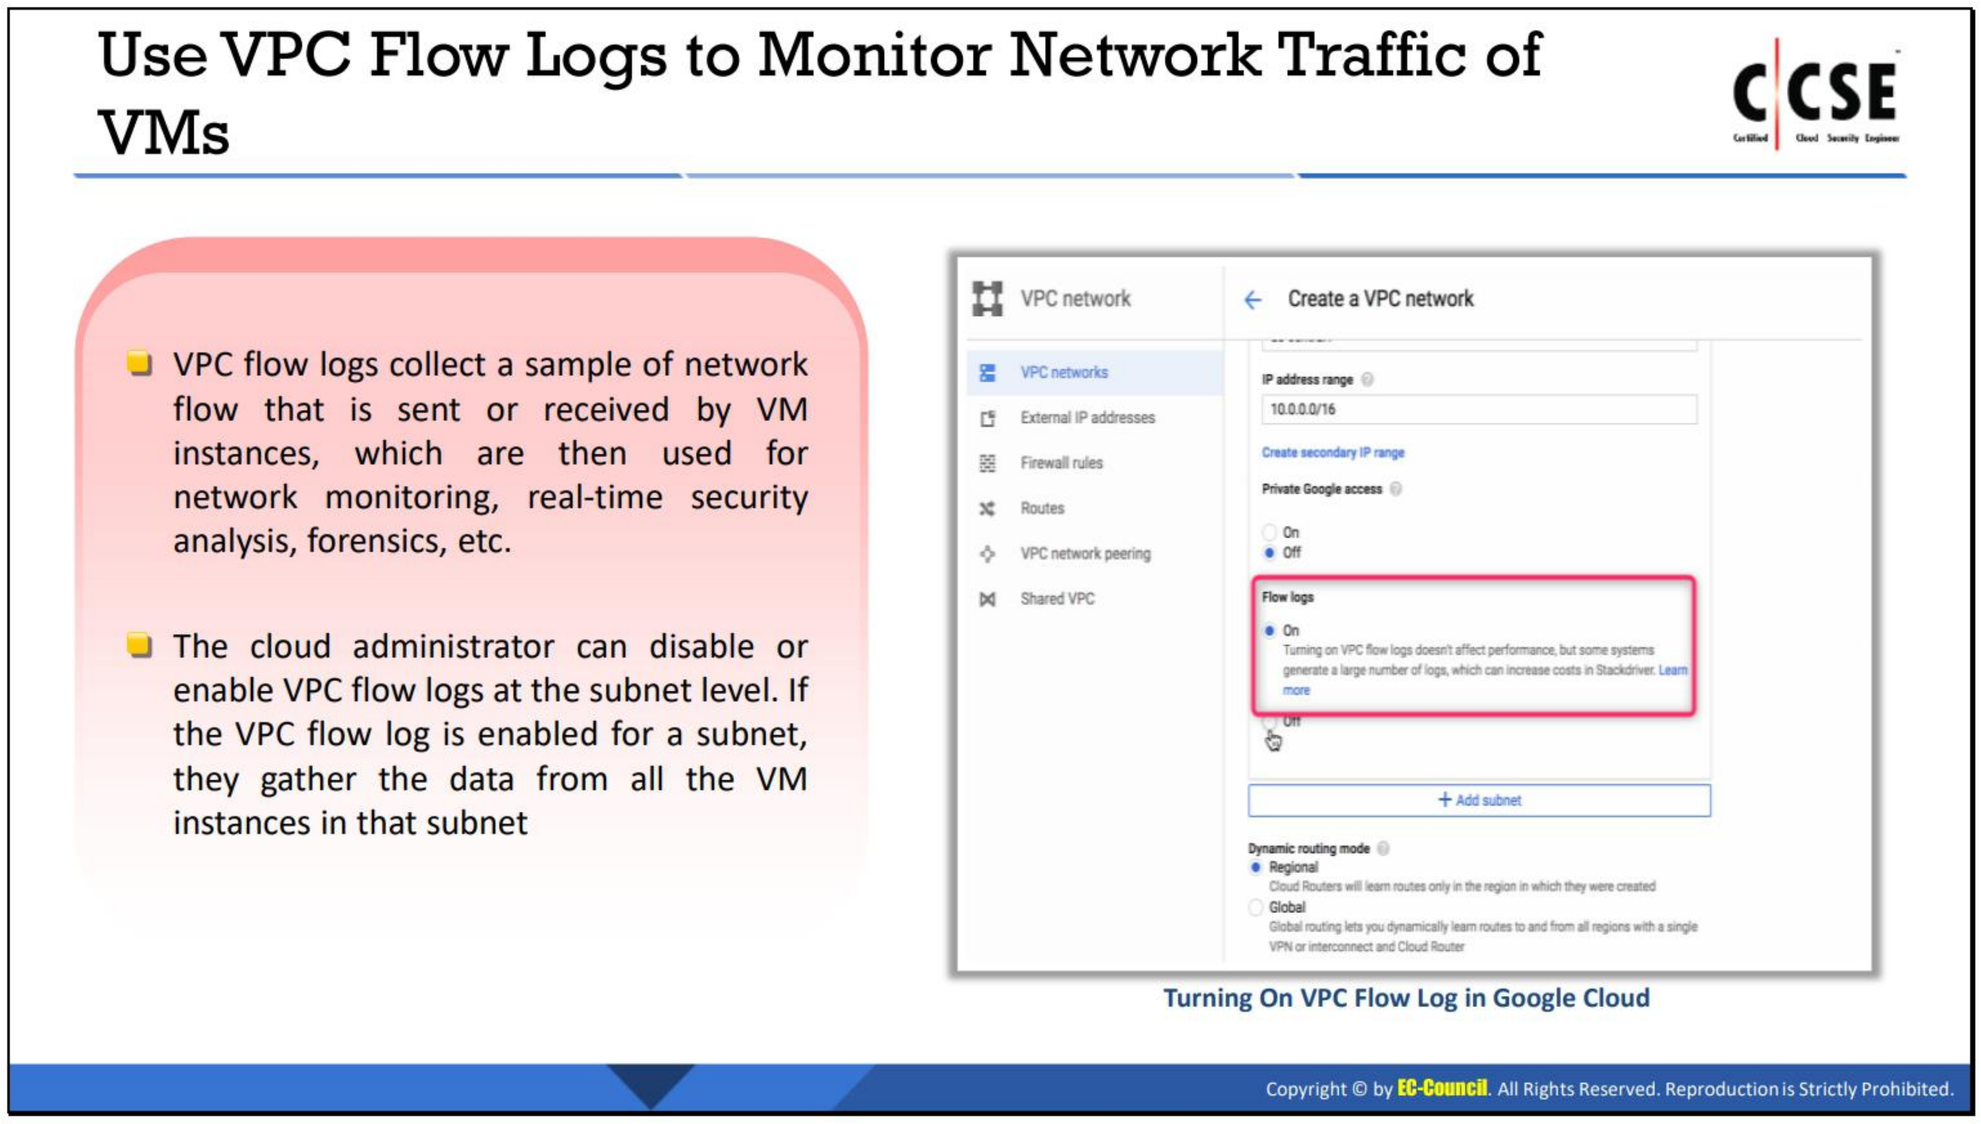
- Choose the two service projects (prod and dev) to attach on the next page. Keep the user roles as 'default' and click on **Save**.



Figure: 2.63: Configuring Sahred VPC

hide01.ir

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



Turning On VPC Flow Log in Google Cloud

Use VPC Flow Logs to Monitor Network Traffic of VMs

VPC flow logs collect a sample of network flow that is sent or received by VM instances, which will be used for network monitoring, real-time security analysis, forensics, etc. These logs help in network logging and monitoring for VMs. The cloud administrator can use them to identify suspicious traffic and access patterns that might pose security or operational risks. For GCP environments, VPC flow logs offer sensible flow-level network telemetry by creating logs within five seconds. With VPC flow logs, network telemetry can be collected at various levels.

The cloud administrator can disable or enable VPC flow logs at the subnet level. If the VPC flow logs are enabled for a subnet, then they can gather the data from all the VM instances in that subnet. The administrator can also gather telemetry for a specific VPC network or a subnet to monitor a particular virtual interface or VM instance. Use VPC Flow Logs to capture telemetry data from several sources. For example, it can track the following:

- Flows between Google services and a server
- Flows between an Internet endpoint and a server
- Internal VPC traffic
- Flows between the on-premises deployments and the VPC over both GC Interconnects or VPNs

The administrator can use VPC flow logs to enhance a variety of operational tasks, including

- **Network Monitoring**: With VPC flow logs, the administrator can monitor applications, their performance, perform debugging and troubleshooting, and implement necessary changes.

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

- **Real-time Security Analysis**: The administrator can use Cloud Pub/Sub API to export logs easily to an existing SIEM ecosystem.
- **Network Forensics and Security Analytics**: VPC flow logs help the administrator to perform network forensics while investigating suspicious activity and ensure compliance.
- **Optimizing Network Usage**: With VPC flow logs, the administrator gets better visibility into inter-region traffic and global traffic usage of an application, which helps them in optimizing network bandwidth utilization, content distribution, and load balancing.

Source: https://cloud.google.com

Enabling VPC Flow Logging while Creating a New Subnet

- On the GCP console, go to the **VPC networks** page.
- Click on the network where the subnet is to be added.
- Click on **Add subnet** and choose **On** under **Flow logs**.



Figure: 2.64: Enabling VPC Flow Logs

- For adjusting log sampling and aggregation, choose **Configure logs** and then adjust
  - **Aggregation interval**
  - Include **metadata** or not
  - **Sample rate**. Keep it 100% to include all entries

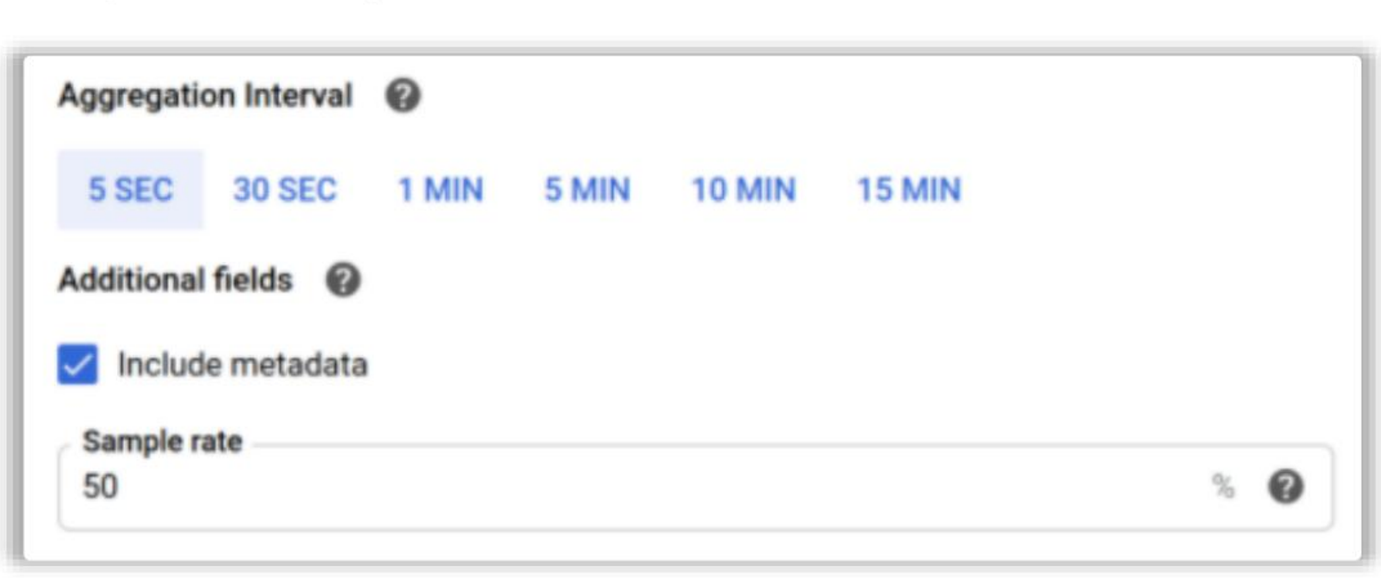


Figure: 2.65: Adjusting Log Sampling and Aggregation

- Fill other relevant fields as appropriate and click **Add**.

Turning on VPC Flow Logging for an Available Subnet

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

- On the GCP console, go to the **VPC networks** page.
- Choose the subnet that must be updated and click Edit.
- Choose **On** under **Flow logs**.

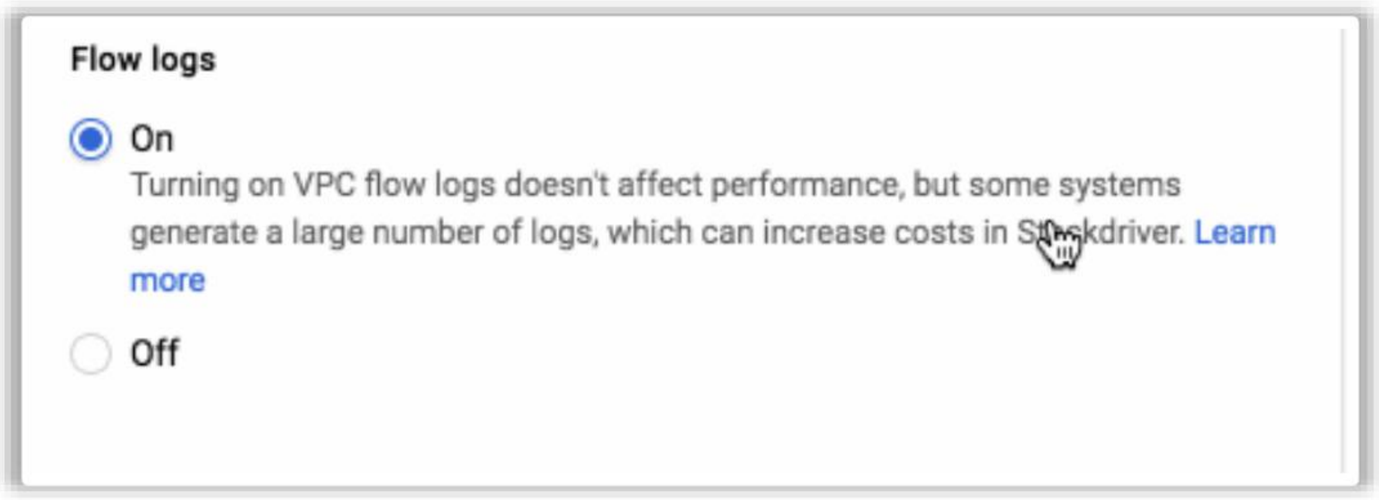


Figure: 2.66: Turning On Flow Logs

- For adjusting log sampling and aggregation, choose **Configure logs** and then adjust
  - **Aggregation interval**
  - Include **metadata** or not
  - **Sample rate**. Keep it 100% to include all entries



Figure: 2.67: Adjusting Log Sampling and Aggregation

- Click **Add**

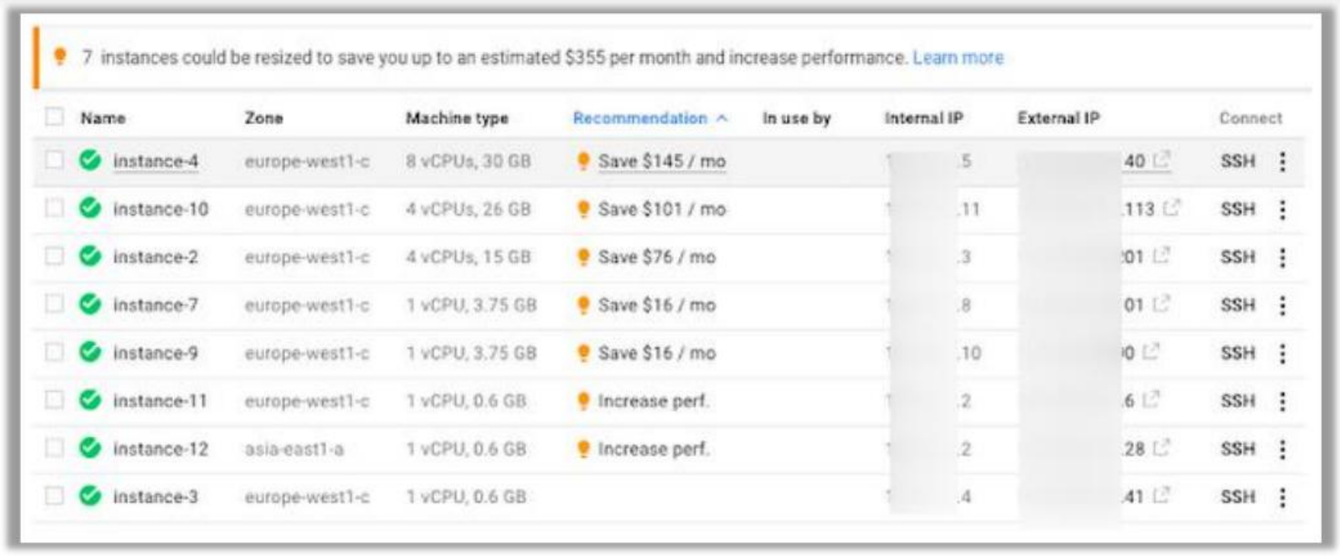
Disabling VPC Flow Logging

- On the GCP console, go to the **VPC networks** page.
- Choose the subnet that must be updated and click Edit.
- Choose **Off** under **Flow logs** and click **Save.**

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

Enforce Compute Engine Rightsizing Recommendations

- The cloud administrator should check the compute engine rightsizing recommendations to optimize the resource utilization of VM instances
- Utilizing the custom machine type feature of Compute Engine, the administrator can create customize machine based on the workloads



Compute Engine Rightsizing Recommendations

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

Enforce Compute Engine Rightsizing Recommendations

Compute Engine in GCP offers machine type recommendations to help the cloud administrator optimize the resource utilization of VM instances. Based on system metrics that were collected by the Cloud Monitoring service or Stackdriver Monitoring service for the past 8 days, these recommendations are automatically generated. The cloud administrator should check the compute engine rightsizing recommendations to optimize the resource utilization of VM instances. For resizing an instance's machine type, use these recommendations to use the instance's resources more efficiently.

Utilizing the custom machine type feature of the Compute Engine, the administrator can create customized VMs based on the workloads. It is easier to view, resize, and perform other actions from the console due to this feature provided by Compute Engine.

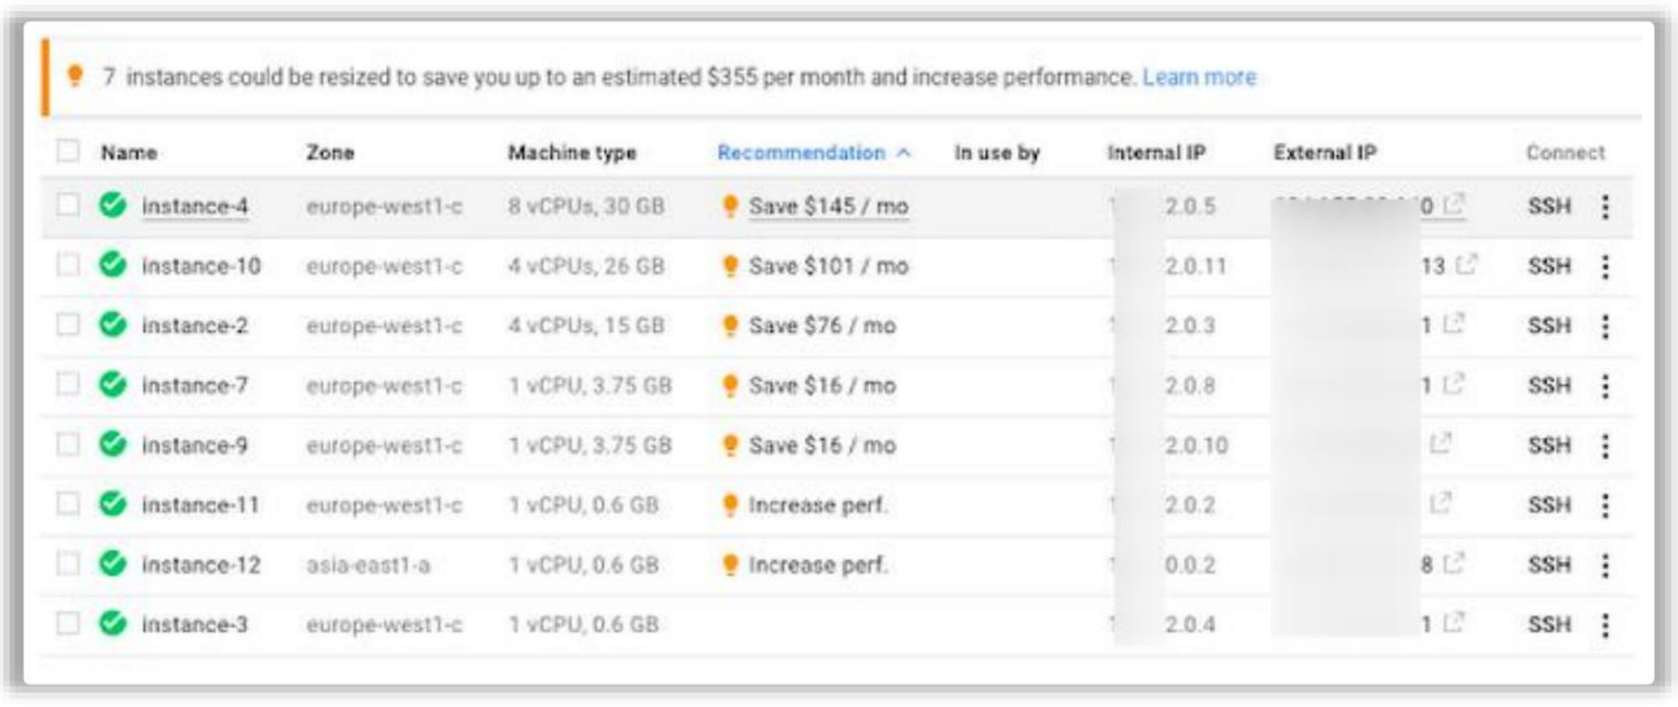


Figure: 2.68: Compute Engine (CE) recommendations

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

Compute Engine (CE) recommendations are largely based on CPU utilization and memory usage. If an instance has had low CPU usage over the previous days, CE recommends a machine that has fewer virtual CPUs. If an instance has had high CPU usage over the previous days, CE recommends a machine that has more virtual CPUs. Similarly, it recommends a machine type that has less memory to the instance that has not used much of its memory over the past days and vice-versa.

Checking the Sizing Recommendations

The administrator can see the recommendations after 24 hours of creating an instance.

For viewing the recommendations, follow the steps:

- Go to the **VM instance** page in the GCP console.
- Choose Recommendation to view the recommendations for individual instances.

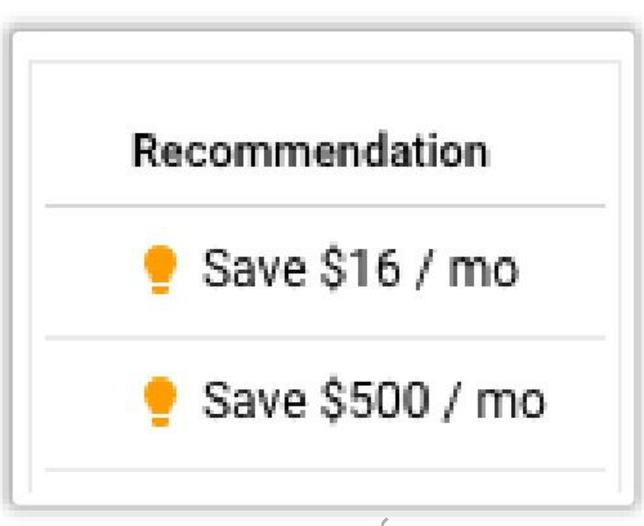


Figure: 2.69: Recommendations for Individual Instances

Resizing Instances as Per Recommendations

If the administrator wants to apply resizing recommendations, they can do so from the Recommendations page directly. Follow these steps to apply resizing recommendations:

- Go to the **VM instance** page in the GCP console.
- Check if any recommendation is present for a particular instance in the **Recommendation** column. Then, choose the recommendation for that instance to resize the instance.

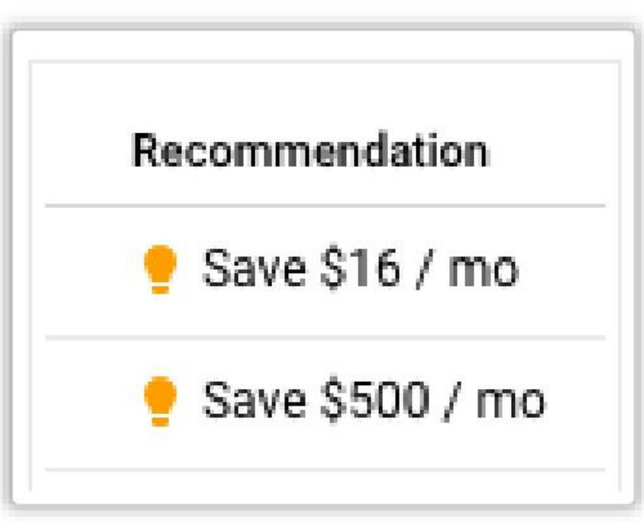


Figure: 2.70: Recommendations for Instances

- Click **Apply** on the next window that contains the recommendations for the instance.

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

- To customize the instance as desired, click the **Customize** link and click on **Apply** to apply the requested changes (optional).

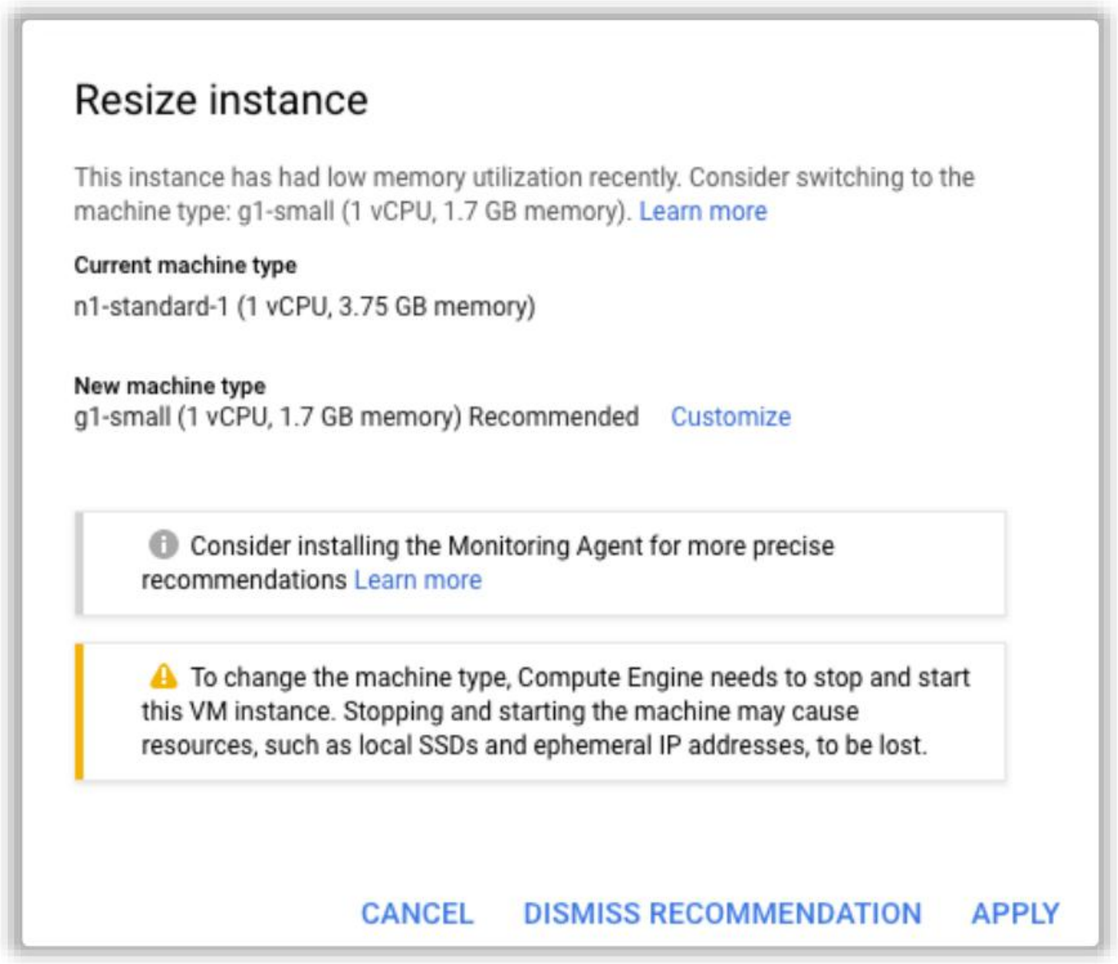
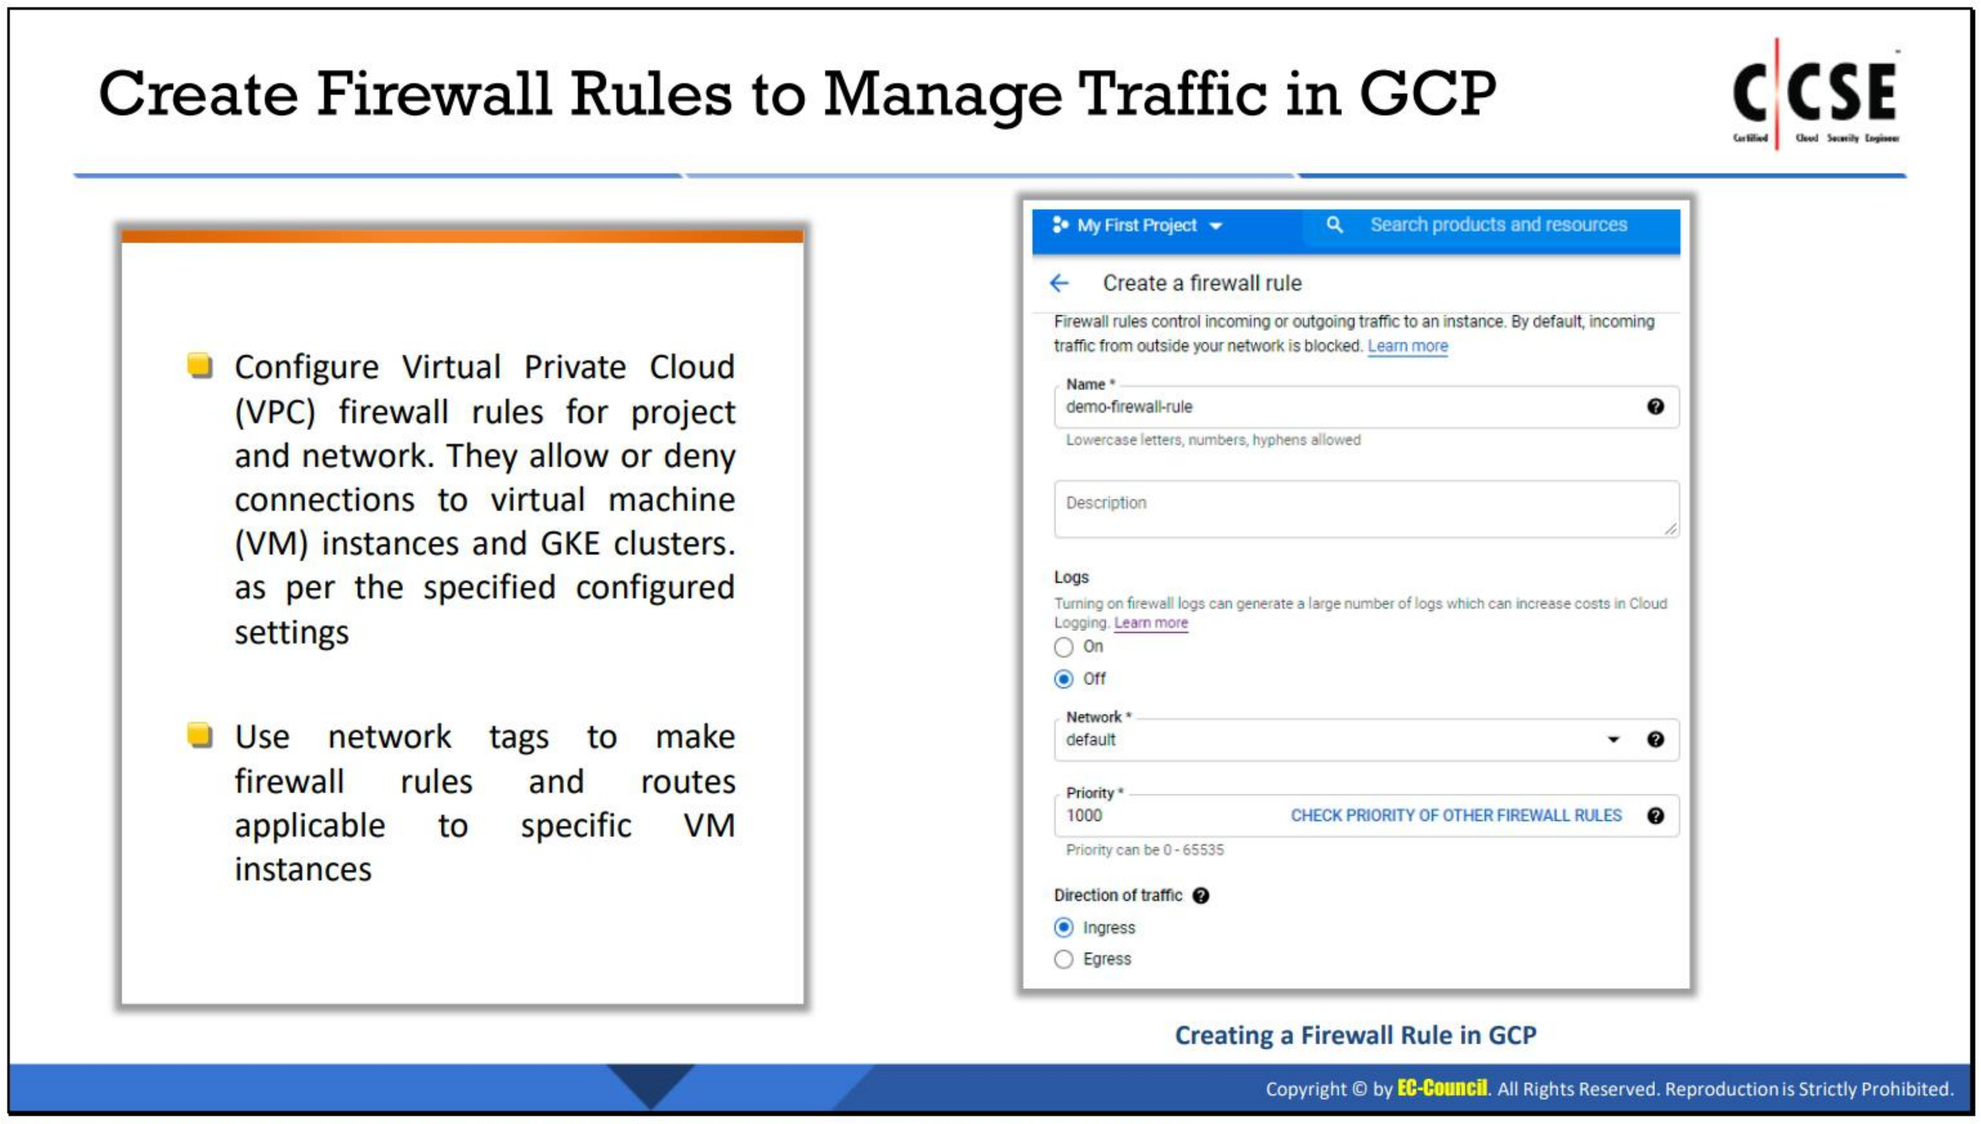


Figure: 2.71: Resizing an Instance

hide01.ir

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



Create Firewall Rules to Manage Traffic in GCP

Firewall rules are applied in a VPC network of GCP. VPC firewall rules are applied to a particular project or network. The firewall rules that are enabled are strictly enforced, thereby securing instances within the VPC network, irrespective of their configuration and OS. A VPC firewall serves as a distributed firewall and network-level rules are created to deny or allow connection for each instance.

Set up firewall rules to allow or deny traffic forward and backward for the resources attached to the VPC, including Compute Engine VM instances and GKE clusters. Firewall rules exist between the instances of a VPC and other networks as well as between the individual instances within the same VPC.

Use network tags to make firewall rules and routes applicable to specific VM instances. The administrator can apply firewall rules to instances that are tagged. In GCP, each firewall rule should have a target that specifies the instances to which it applies. By default, a firewall rule targets all instances present in the network, but it can be customized by the administrator using target network tags. These tags specify the VMs in GCP to which they apply.

Firewall rules provide the following security benefits:

- The administrator can use firewall rules to secure the network and instances in GCP from unauthorized and unwanted access.
- They help to monitor traffic for suspicious activity and block potentially dangerous packets or items based on a set of security rules.
- They help prevent attacks, malware, viruses by serving as the first line of defense to secure the network.

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

Enabling Firewall Rules (GC User Interface)

- From Google console drop-down menu, navigate to **VPC network;** select **VPC networks** on Menu Tab.
- Navigate to **Firewall rules.**

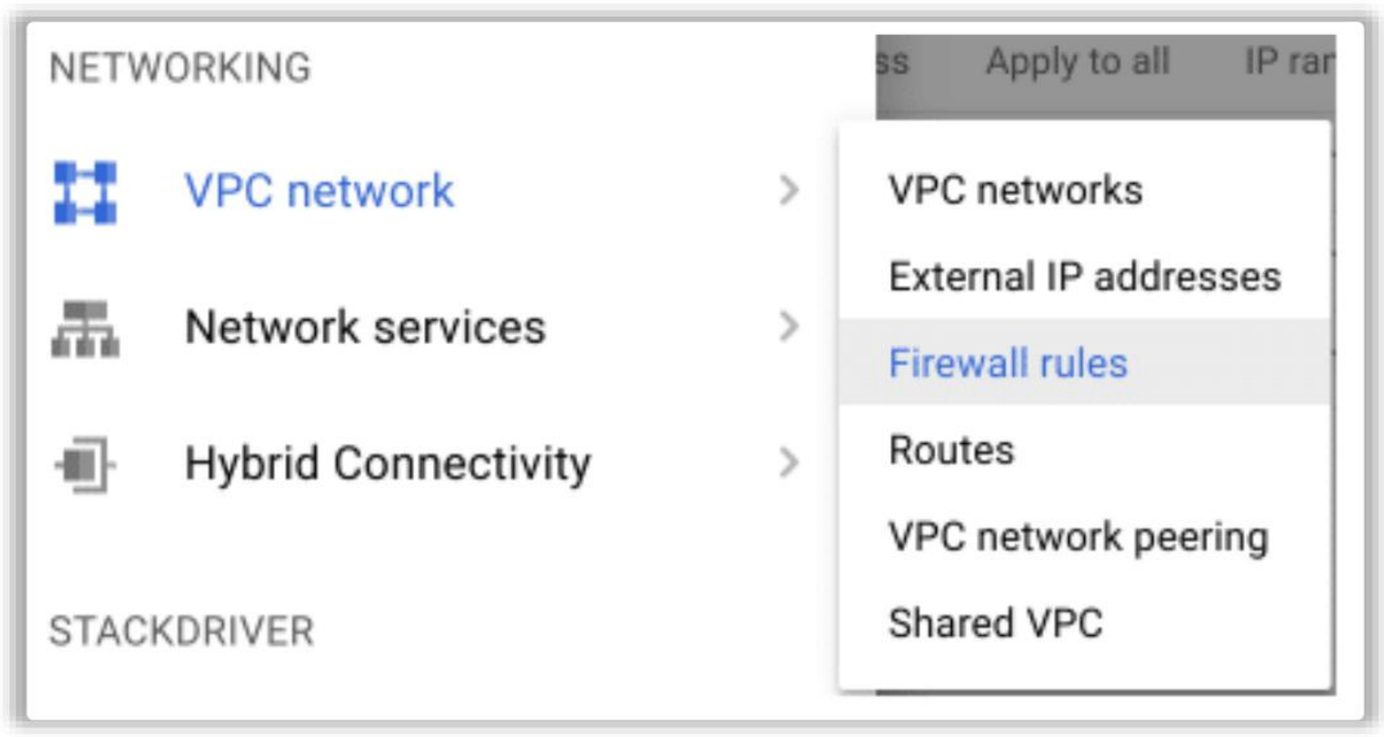


Figure: 2.72: Navigating to Firewall Rules

- Click **Create firewall rule**, set the required rules, then click **Create.**

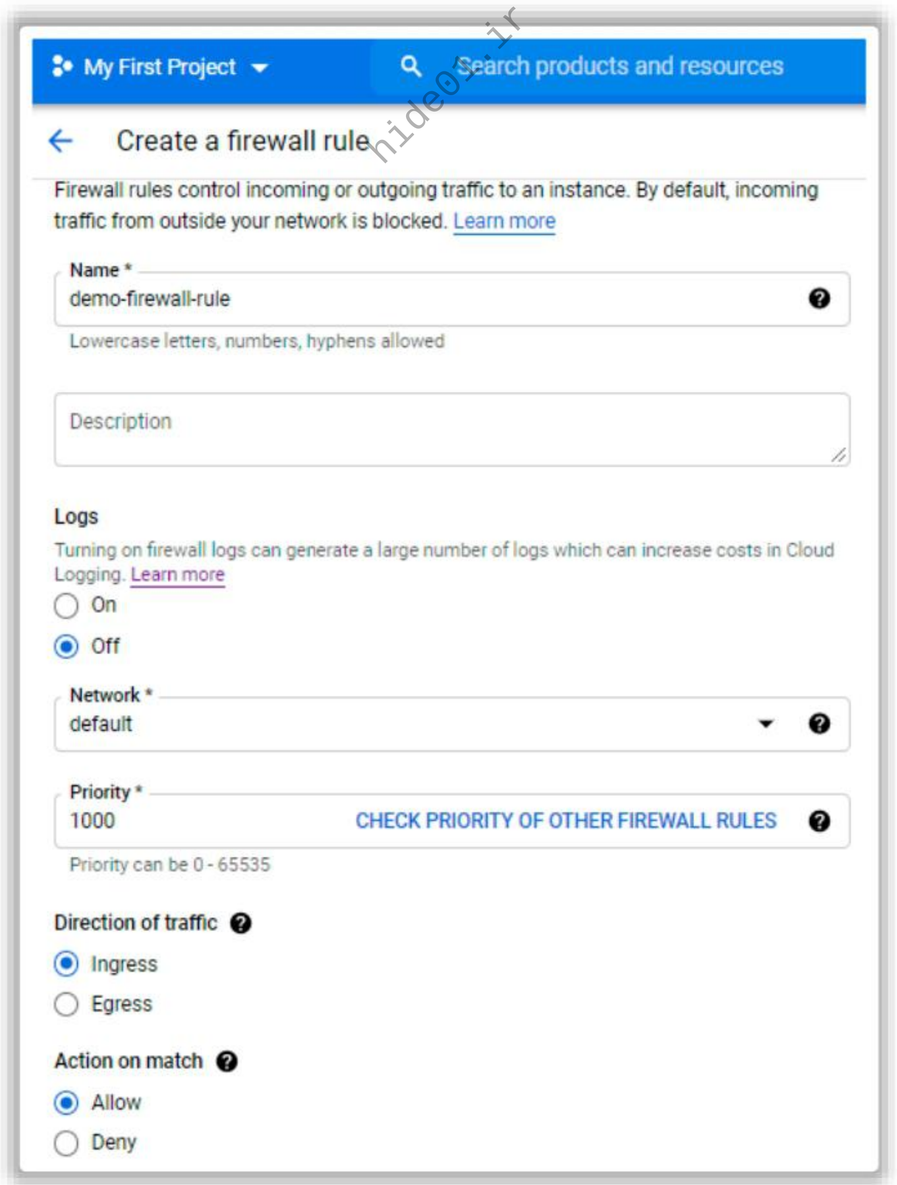


Figure: 2.73: Creation of a Firewall Rule

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

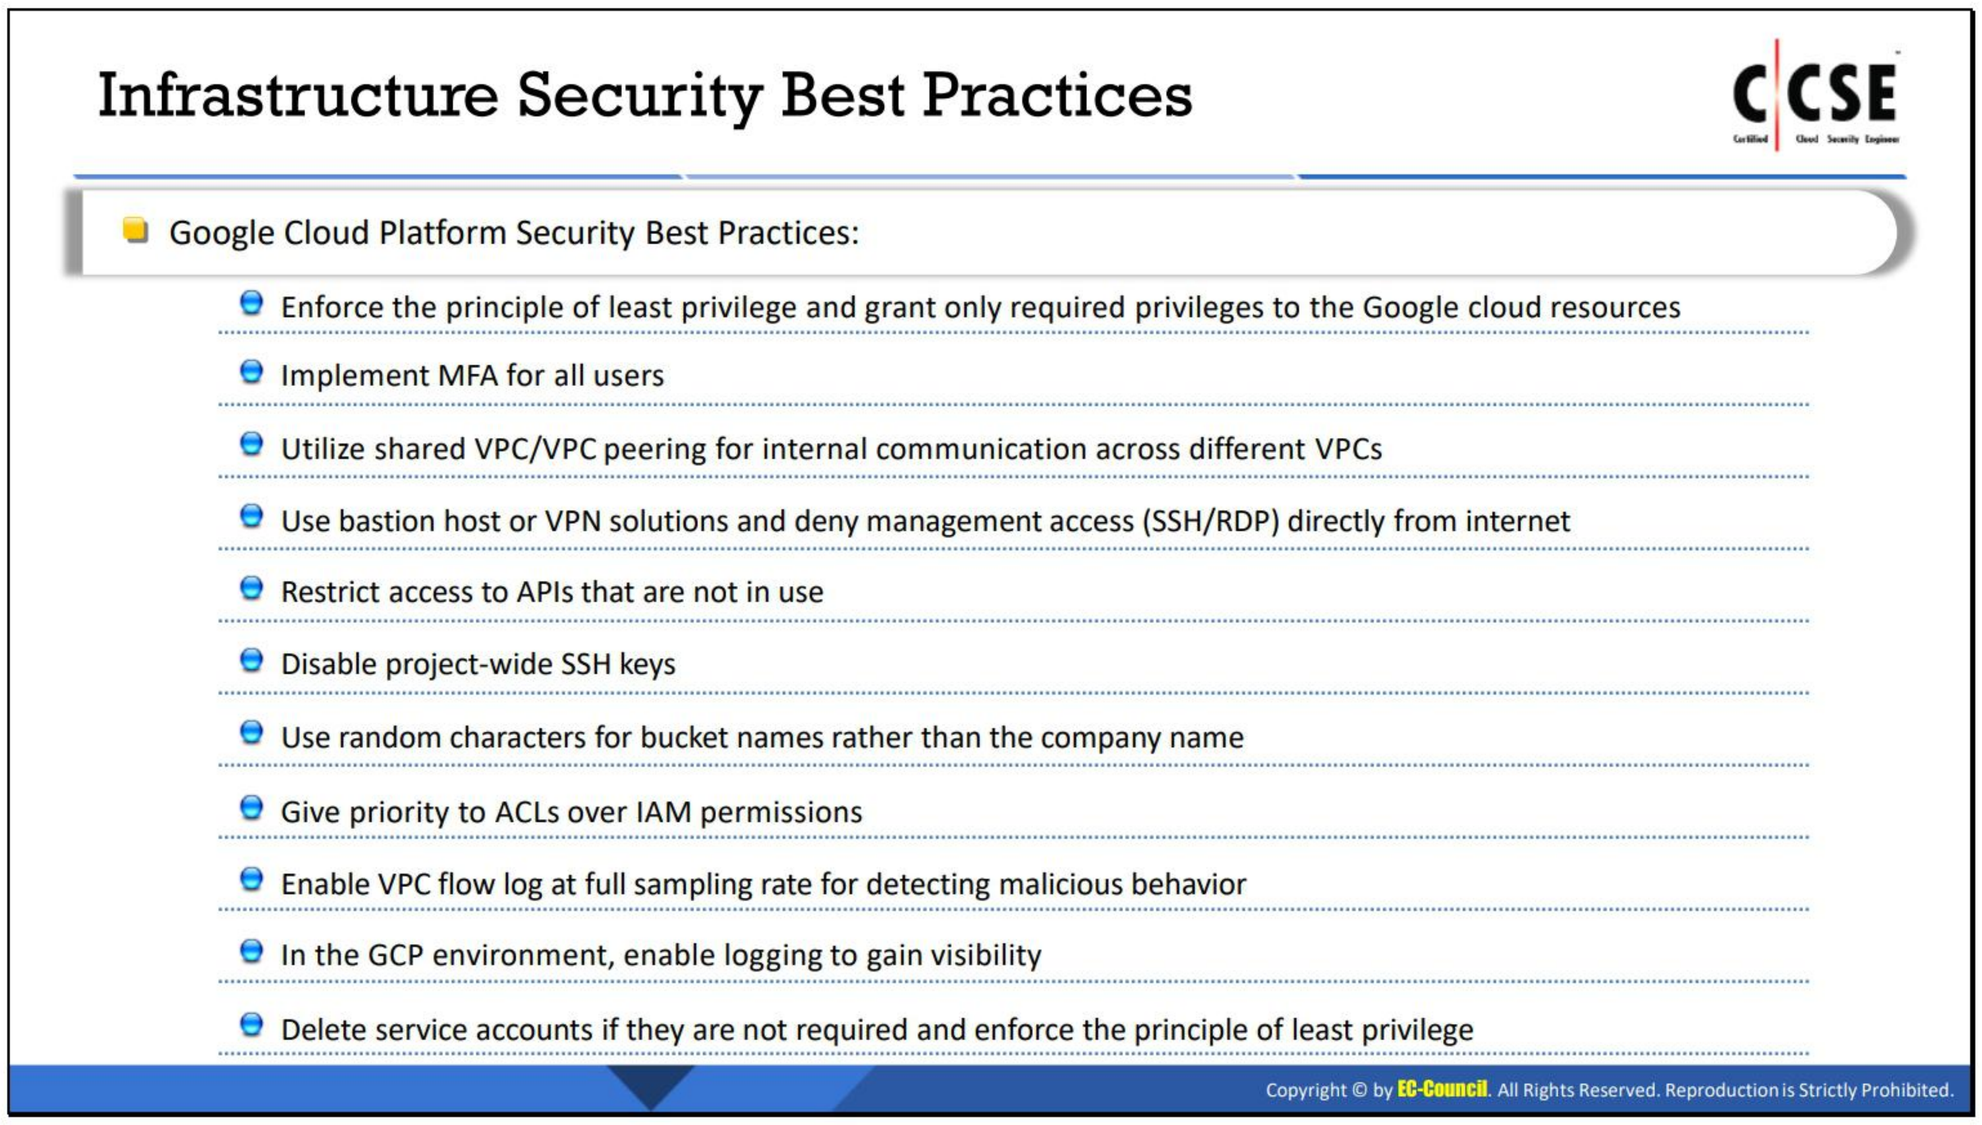
Creating Firewall Rules (Cloud Shell)

Make a firewall rule that allows all IP addresses present in **subnet A** (**10.240.10.0/24**) to access **port 80** on instances that are tagged with **webserver**.

```
gcloud compute firewall-rules create vm1-allow-ingress-tcp-
                    port80-from-subnetA \
                  --network my-network \
                    --action allow \
                  --direction ingress \
                    --rules tcp:80 \
               --source-ranges 10.240.10.0/24 \
                    --priority 50 \
               --target-tags webserver
```

hide01.ir

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



Infrastructure Security Best Practices

Security best practices for GCP are described below.

- **Enforce The Principle of Least Privilege (PoLP) and Grant Only Required Privileges to the Cloud Resources:** Use Cloud Identity or Google Workspace to administer user accounts and assign them to suitable groups. Utilize Cloud IAM to assign privileges inside GCP. Use accounts that are easy to administer.

  Use predefined roles (predefined sets of privileges) to implement fine-grained privileges for all the services. Assign the roles to users or groups, enforce PoLP, and grant only required privileges to the cloud resources.

  When predefined roles did not meet the need of the organization, then use custom roles. Custom roles are user-defined roles and allow the collection of supported permissions to meet the organization's requirements. Avoid assigning Primitive roles.

- **Implement MFA for All Users:** MFA is used to control access to cloud service APIs. With the implementation of MFA, the user account can be secured even if the credentials or access keys are compromised.

  Enforce MFA for all users. It is necessary for the administrators on internet exposed platforms. GCP supports MFA for corporate entities and Cloud Identity.

- **Utilize Shared VPC/VPC Peering for Internal Communication Across Different VPCs:** Create and manage network resources (routes, subnets, firewall rules, etc.) within a VPC. Instead of the default network and associated firewall rules, create your VPC and firewall rules when a new project is created in GCP.

  Allow only the required egress and ingress traffic for the functioning of the environment.

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

In hybrid environments, reduce the attack surface area and keep the communication utilizing public IPs to a minimum extent.

In the hybrid environments, use Google Cloud VPN and Cloud Interconnect between two sites to maintain communication using public IPs to a minimum extent and to reduce the attack surface area. Similarly, utilize a shared VPC/VPC peering for internal communication across different VPCs.

- **Deny Management Access (SSH/RDP) Directly from The Internet and Utilize Bastion Host or VPN Solution:** Disable public IPs in a project because all VMs in GCP, by default, are accessible directly from the internet if firewall rules allow it (like the default ones). If external communication is required, assign public IP to a particular instance, or utilize an internet gateway such as Cloud NAT.

  It is recommended to deny management access (SSH/RDP) directly from the internet and use bastion host or VPN solution to manage VMs in the cloud.

- **Restrict Access to APIs That are Not in Use:** Disable or restrict access to APIs that are not in use through IAM policies and VPC Service Controls to reduce the attack surface.

- **Disable Project-Wide SSH Keys:** In GCP, the Compute Engine manages the VMs. By default, there are some features in VMs that are less secure such as "Project-wide SSH keys." Project-Wide SSH keys allows administrators to provision public SSH keys of all instances in a project. Disable Project-Wide SSH keys to prevent the SSH key holder privileged access to all Linux instances.

- **Use Random Characters for Bucket Names Rather Than The Company Name:** Ensure the buckets' names are unique as they are accessible globally. For example, name buckets as "mybucket-customer-y" and "mybucket-customer-z" for Customer Y and Customer Z, respectively.

  Use random characters for bucket names rather than the company name to make the attacker hard to locate buckets. For example, use "prod-logs-a1b2d456z4567e355c5321f89g48g35".

- **Give Priority to ACLs Over IAM Permissions:** Understand which policy permissions and ACL rules will fit the use cases. Give priority to ACLs over IAM permissions as they provide more granular control on an object level.

- **Enable VPC Flow Log at Full Sampling Rate for Detecting Malicious Behavior:** The cloud has no physical firewalls or switches. Enable VPC flow log at full sampling (100%) rate for detecting port scans or malicious behavior within the VPC network. The VPC flow log consist of fields such as source and destination IP, protocol, source and destination port, etc. Although, this could create a large amount of data and requires extra storage, it could also provide full visibility into the GCP environment. If the storage is an issue, enable full logging on specific subnets.

- **Enable Logging to Gain Visibility:** Use Cloud Audit Logs to supervise firewall or IAM configuration rules. Cloud Audit Logs comprises many sources. Two such sources are described below.

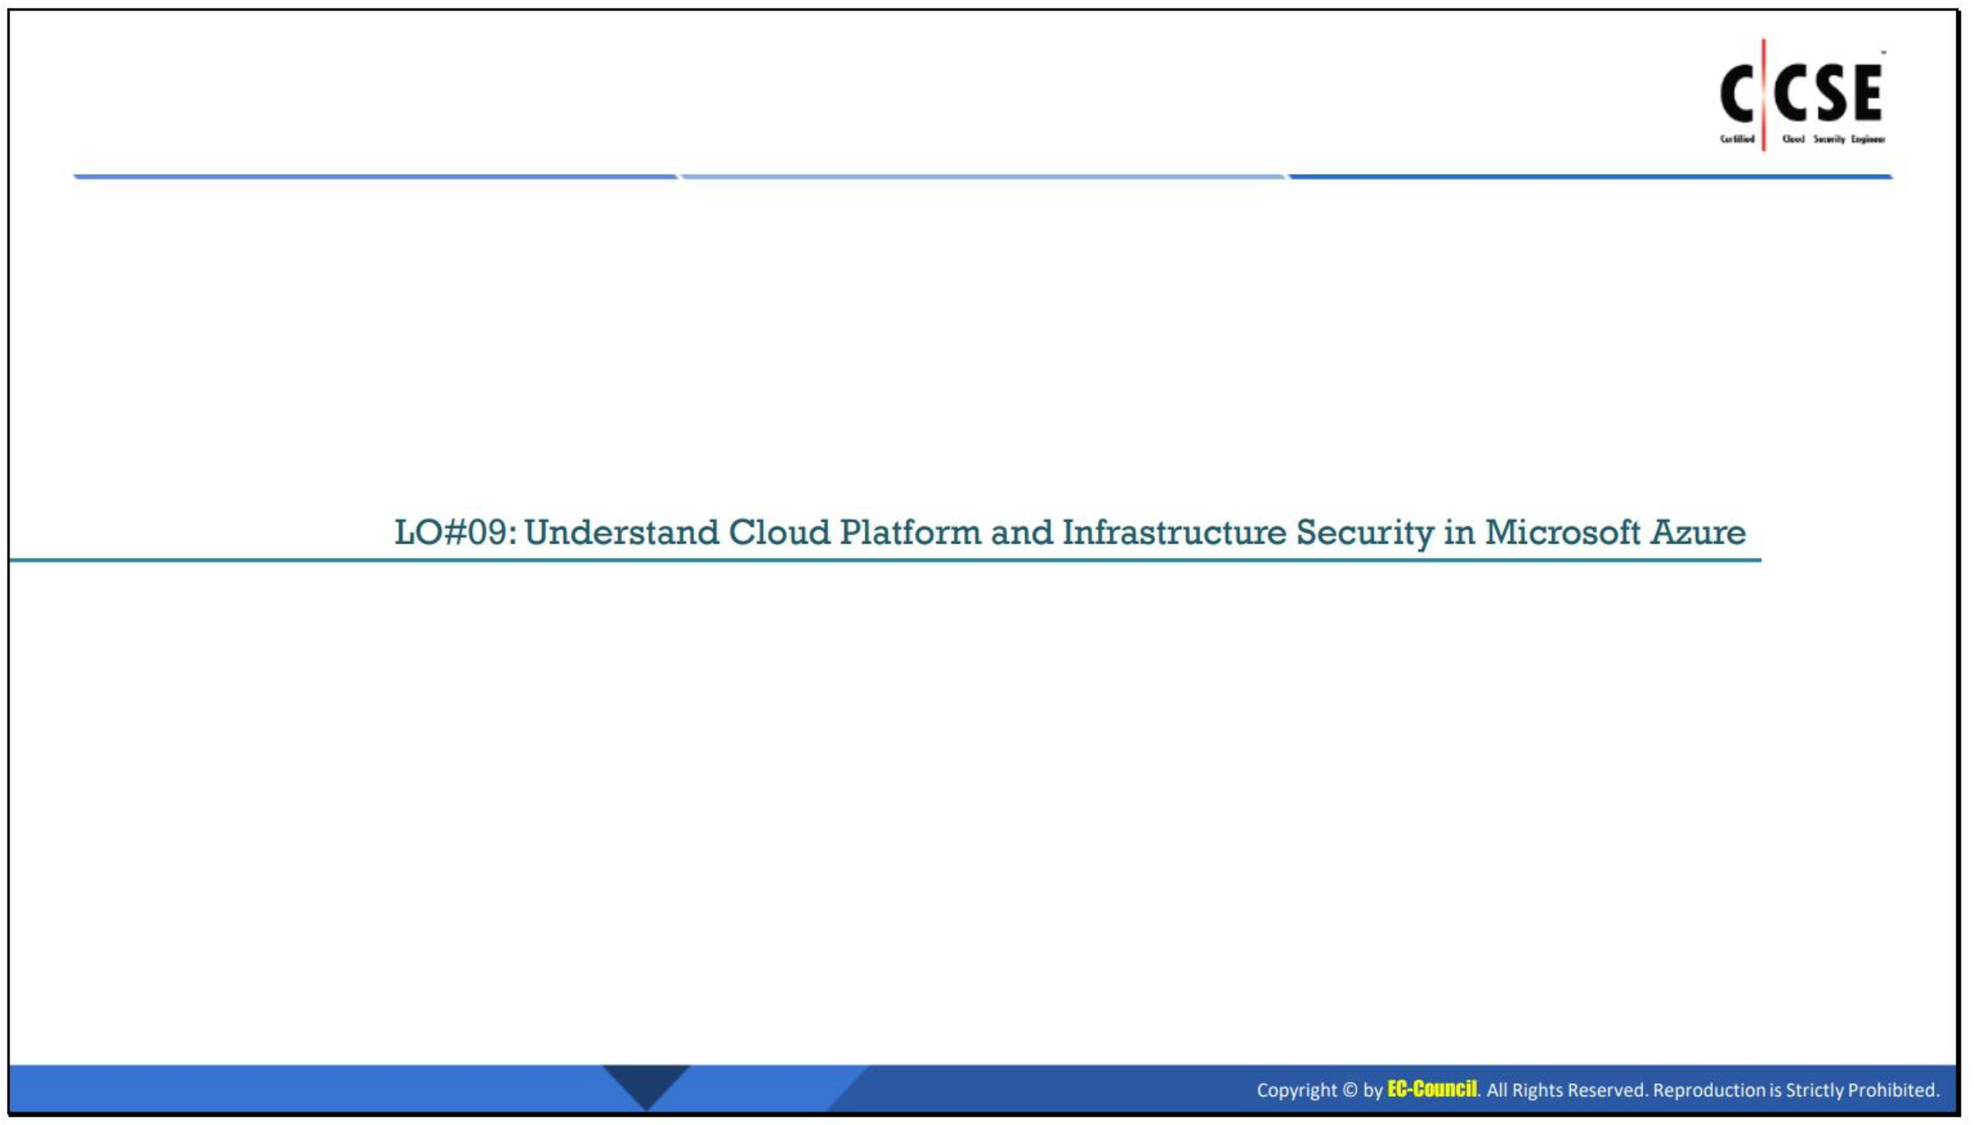
Certified Cloud Security Engineer Copyright © by **EC-Council**
All Rights Reserved. Reproduction is Strictly Prohibited.

- **Administrator Activity Log:** It is enabled by default and record operations that modify metadata or configurations of a resource. For example, an administrator modifying an IAM policy or creating a VM instance.
- **Data Access Log:** It is disabled by default as it generates many entries. It records operations that read information about a resource. Service account reading an IAM policy is an example of data access log.

- **Delete Service Accounts If They Are Not Required and Enforce The Principle Of Least Privilege:** Delete service accounts if they are not required and enforce PoLP. In case, the attacker got hold of the service account, then the attacker may issue requests to Google APIs by impersonating a service account and could cause serious damage.

hide01.ir

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

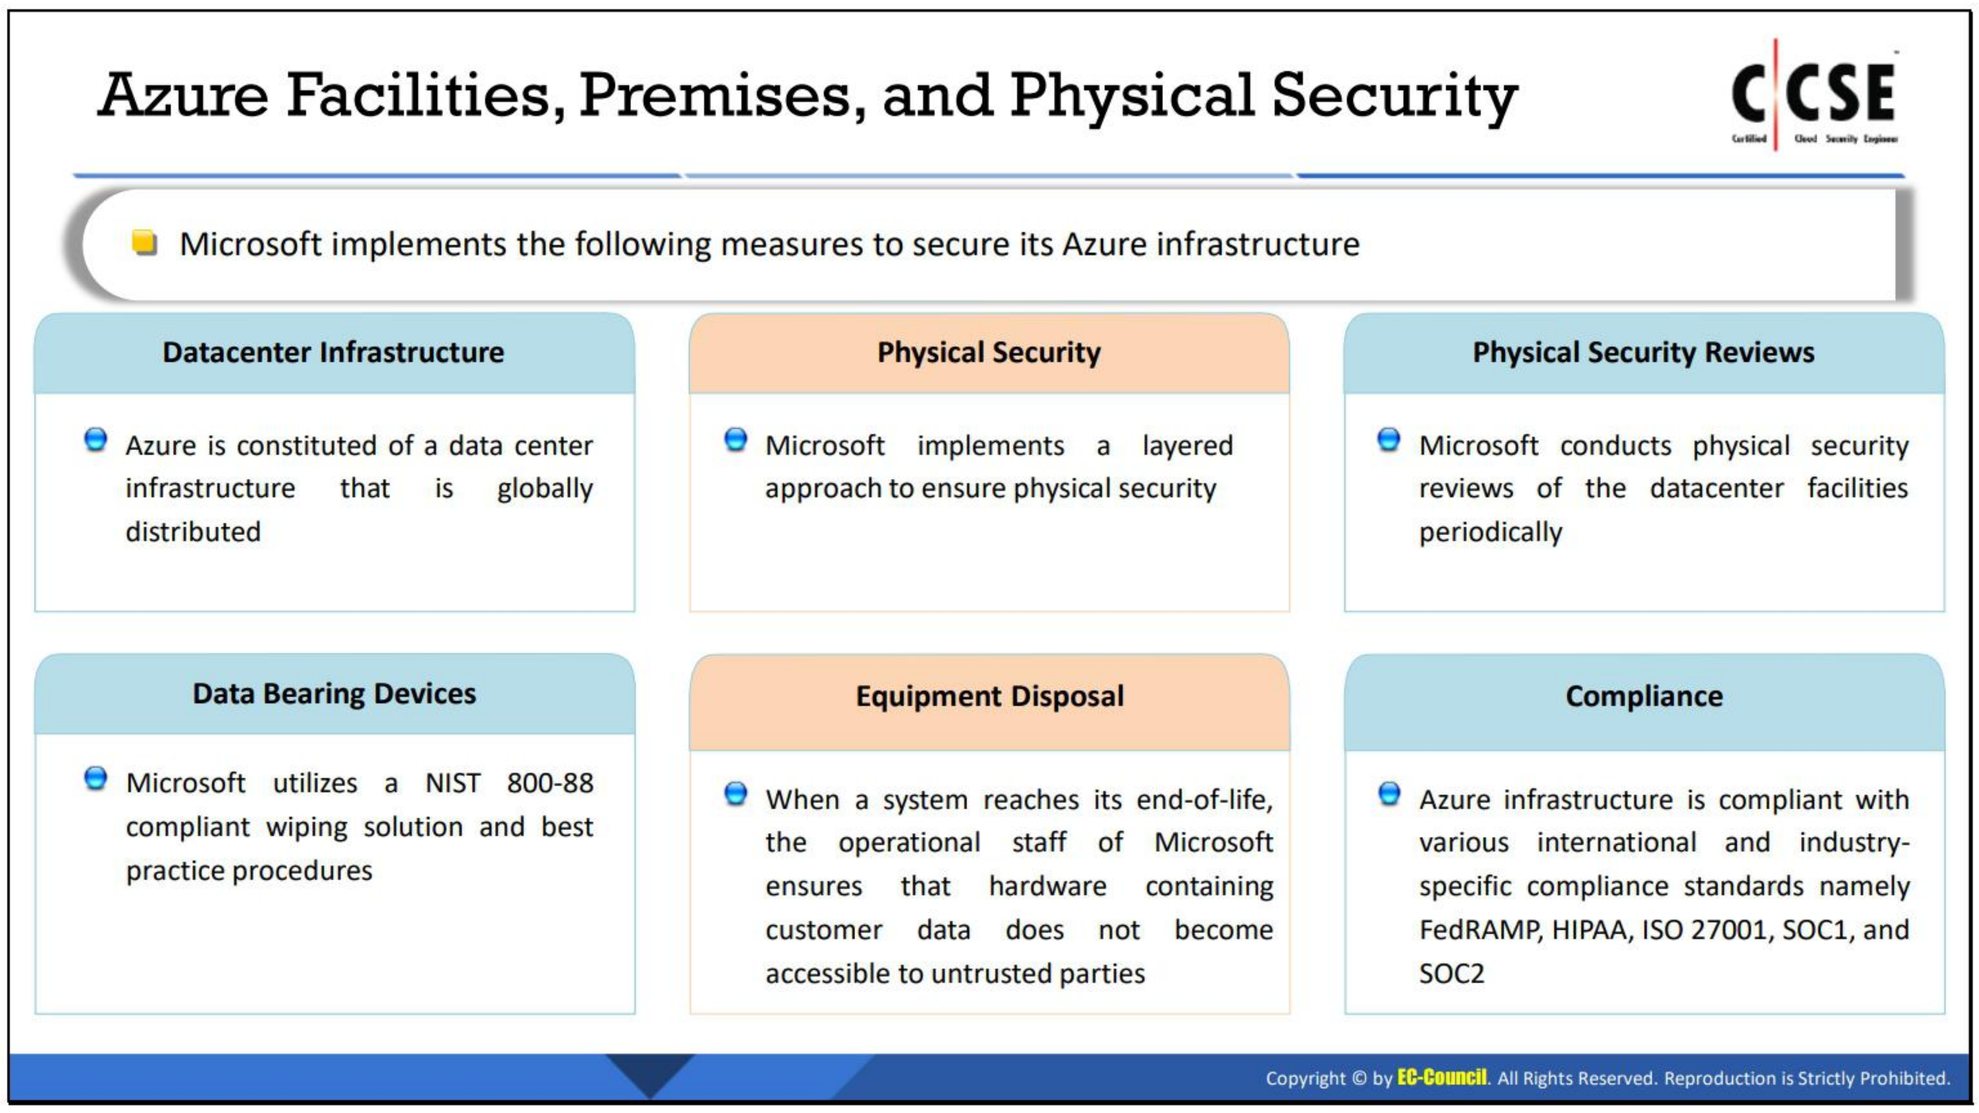


LO#09: Understand Cloud Platform and Infrastructure Security in Microsoft Azure

This section will provide you a detailed explanation of security measures and system components used to ensure cloud platform and infrastructure security in Azure. It will help you understand commonly used Azure resources by CSPs and customers in Microsoft Azure.

vid01.ir

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



Azure Facilities, Premises, and Physical Security

Microsoft operates and manages the data centers that are used by Microsoft Azure. Microsoft data centers are spread throughout the world in multiple locations and comply with all key industry standards such as NIST SP800-53 and ISO/IEC 27001:2013 for reliability and security. Microsoft operations staff is responsible for managing, monitoring, and administering these data centers. Its operational staff has years of experience in providing the world's largest online services around the clock.

Microsoft implements the following measures to secure its Azure infrastructure.

Datacenter Infrastructure

Azure is made up of a data center infrastructure that is distributed globally. Its global datacenter infrastructure spans more than a hundred highly secured facilities around the world and provides several online services.

Microsoft's infrastructure is designed to bridge the gap between applications and users around the world and make applications easily reachable to users. It preserves data residency and offers comprehensive and resiliency options for its customers. Microsoft Azure is available in 140 countries or regions and has 58 regions worldwide.

A region is a set of interconnected data centers through a massive and resilient network. The network consists of load balancing, content distribution, redundancy, data-link layer encryption by default for all Azure traffic traveling between regions or within a region. Azure provides customers with the flexibility to deploy applications where they want because it has more global regions than any other cloud provider.

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

Azure regions are organized into geographies and each Azure geography ensures that requirements such as data residency, compliance, sovereignty, and resiliency are honored (met) within geographical boundaries.

Geographies help Azure customers who have specific data-residency and resiliency requirements to keep their data close. Geographies are designed to be fault-tolerant and can withstand complete region failure by using their connection to the high-capacity, dedicated networking infrastructure.

Within an Azure region, AZs are physically separated locations. Every AZ is composed of one or more data centers that are equipped with independent cooling, power, and networking. With the help of AZs, customers can run mission-critical applications with low-latency replication and high availability.

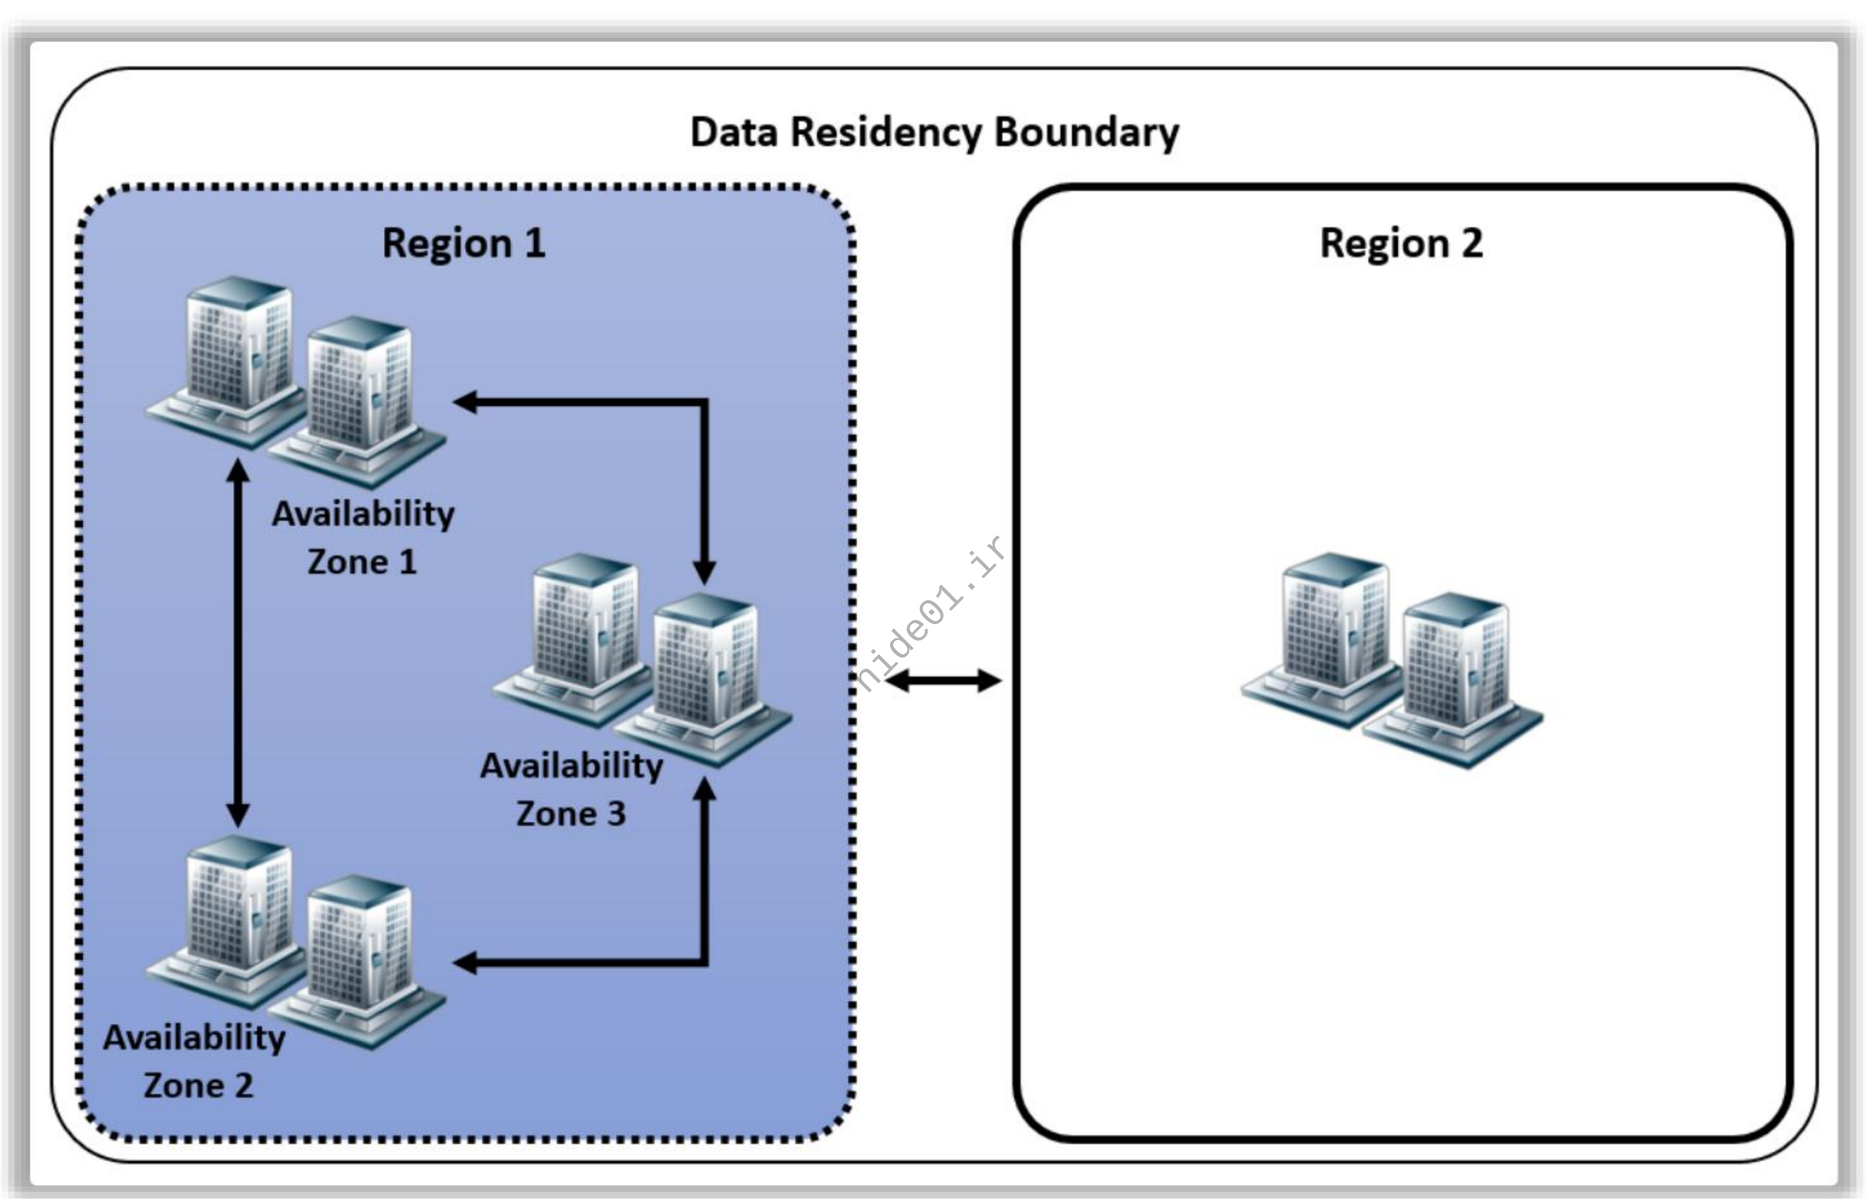


Figure: 2.74: Data Residency Boundary

The above figure shows how Azure global infrastructure connects regions and AZs within the same data-residency boundary to provide backup, disaster recovery, and high availability.

Data centers that are distributed geographically allow Microsoft to be close to its customers, enable geo-redundant backup and failover, and reduce network latency.

Physical Security

Microsoft data centers are designed, built, and maintained in a way that strictly controls physical access to the areas where customer data is stored. Microsoft knows the importance of securing customer data and is always committed to ensuring the safety and security of data centers that contain customer data. It has a separate division that takes care of designing, building, and

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

operating the data center facilities that support Azure. The sole purpose of that team is to maintain state-of-art physical security.

Microsoft implements a layered approach to ensure physical security, reducing the risk of unauthorized users getting access to the data centers or their resources. There are extensive layers of protection in Microsoft data centers, including access approvals at various locations such as the facility's perimeter, the building's perimeter, within the building, and on the data center floors. These layers of physical security are described below.

- **Access Request and Approval:** Before arriving at the data center, a visitor must request access. The visitor must give a valid business justification for their visits like auditing or compliance purposes. Microsoft employees approve all requests on a need-to-access basis because it helps them keep the number of visitors who need to perform a task in the data center to the bare minimum. Once an individual gets permission to get inside, they are provided access to only the required discrete area of the data center, based on their approved business justification. Permissions granted to individuals tend to expire after a certain time.
- **Facility's Perimeter**: When the visitor arrives at the data center, they are made to go through a well-defined access point. Generally, tall steel fences and concrete surround every inch of the perimeter. Microsoft has cameras around its data centers and dedicated security teams who monitor the video recording all the time.
- **Building Entrance**: The entrance of every Microsoft data center has security officers who have gone through rigorous background checks and security training. They also patrol the data center routinely and monitor video recordings of cameras inside the data center all the time.
- **Inside the Building:** A visitor must pass two-factor authentication with biometrics after they enter the building to continue moving within the data center. If the validation is successful, the visitor can only go to the portion of the data center for which they got access permission. They can be there only for the approved duration of time.
- **Datacenter Floor**: Visitors are allowed onto the floor only after getting the approval. They need to pass the full-body metal detection screening before going on the floor. Microsoft allows only approved devices into the data center floor to reduce the risk of unauthorized data entering or leaving the data center facility without its knowledge. In addition, Microsoft uses video cameras to monitor the front and back of each server rack. The visitor must pass the full-body metal detection screening when they leave the data center floor. When the visitor leaves the data center, they must pass through another security scan. As per Microsoft's rules, all visitors must return badges when they leave a Microsoft data center facility.

Physical Security Reviews

To ensure whether data centers are addressing Azure security requirements properly or not, Microsoft conducts physical security reviews of the datacenter facilities periodically. Azure

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

service management is not provided by the datacenter hosting provider staff and they cannot sign-in to Azure systems and have no physical access to the Azure collocation rooms and cages.

Data Bearing Devices

Microsoft utilizes a NIST 800-88 compliant wiping solution and best-practice procedures. It uses a destruction process for hard drives that cannot be wiped, which destroys them and makes the data recovery impossible. This process can be used to shred, disintegrate, incinerate, or pulverize. According to the asset type, Microsoft determines the means of disposal and retains the records of the destruction.

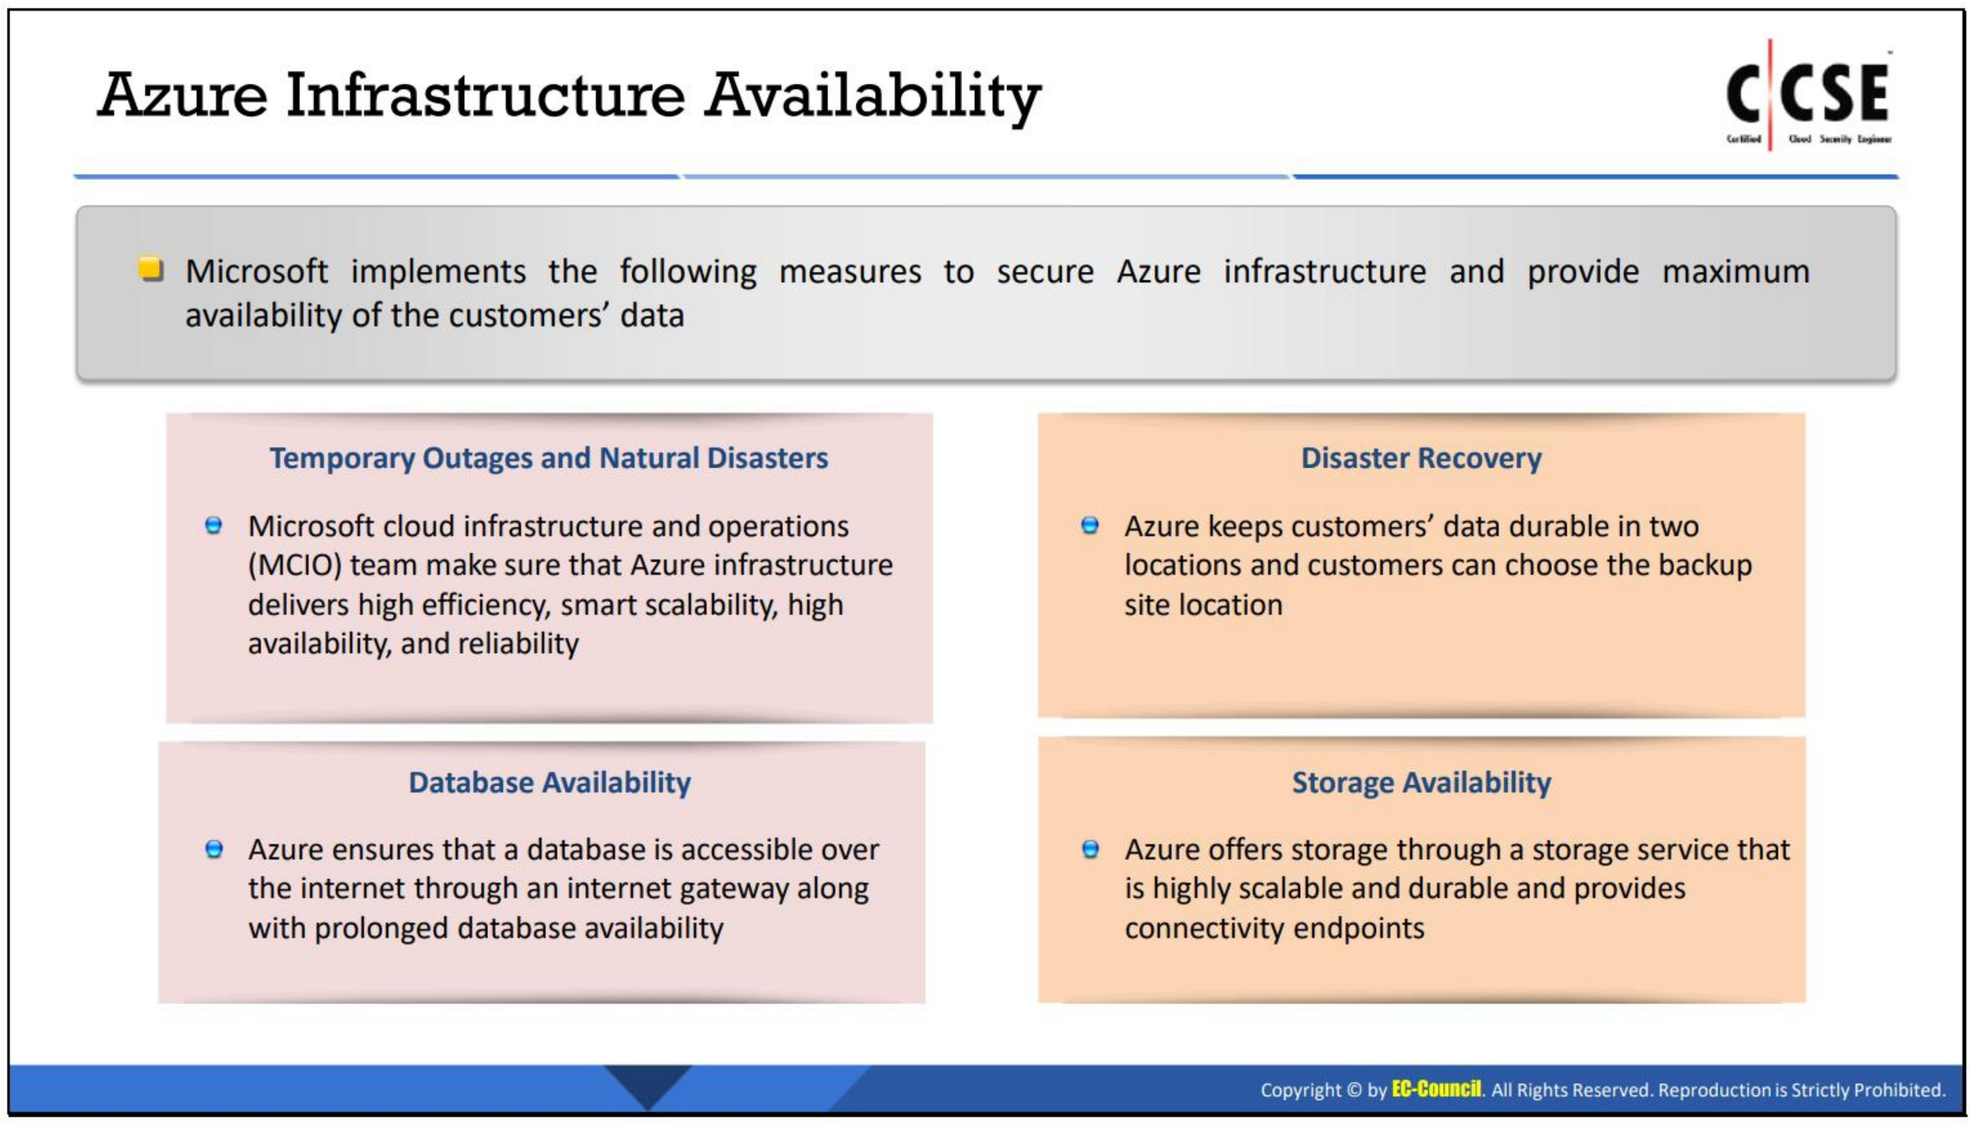
Equipment Disposal

When a system has reached its end of life, the operational staff of Microsoft ensures that hardware containing customer data does not become accessible to untrusted parties by following rigorous data handling and hardware disposal procedures. It uses a secure erase approach for hard drives that support that method. It uses a destruction process for hard drives that cannot be wiped, which destroys them and makes the data recovery impossible. This process can be used to shred, disintegrate, incinerate, or pulverize. According to the asset type, Microsoft determines the means of disposal and retains the records of the destruction. Only approved media storage and disposal management services are used for all Azure services.

Compliance

Azure infrastructure has been designed and managed by Microsoft to meet a wide range of international and industry-specific compliance standards such as FedRAMP, HIPAA, ISO 27001, SOC1, and SOC2. Besides, it meets other country-or-region-specific standards such as UK G-Cloud, Australia IRAP, and Singapore MTCS. Third-party audits, like those done by the British Standards Institute, verify adherence to the mandated security controls by these standards.

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



Azure Infrastructure Availability

Microsoft implements the multiple measures to secure its Azure infrastructure and provide maximum availability of its customers' data. These measures are discussed below. Based on the extensive redundancy achieved with virtualization, Azure offers robust availability.

hidoo1.ir

Temporary Outages and Natural Disasters

The Microsoft cloud infrastructure and operations (MCIO) team designs, builds, manages, and improves the security of Microsoft cloud infrastructure. This team makes sure that the Azure infrastructure delivers smart scalability, high efficiency, high availability, and reliability. The team offers a more secure, trusted, and private cloud.

The continuous flow of electricity is ensured by using uninterrupted power supplies and vast banks of batteries when a short-term power failure occurs. For any planned maintenance and extended outages, emergency generators provide backup power. The data center can utilize the on-site fuel reserves if a natural disaster happens.

To connect data centers with other major hubs and internet users, robust and high-speed fiber-optic networks are used. Workloads closer to users are hosted by compute nodes to provide geo-redundancy, reduce latency, and increase overall service resiliency. A group of engineers continuously work around the clock to make sure services are available persistently.

By using service support, backup failover capability, and advanced monitoring and incident response, Microsoft ensures high availability. Microsoft operations centers that are distributed geographically operate 24/7/365 (continuously throughout the year). Data centers and edge nodes are connected through the fiber-optic and content distribution network to provide high performance and reliability. The Azure network is considered as one of the largest networks around the world.

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

Disaster Recovery

Azure keeps customers' data durable in two locations and customers can choose the backup site location. Azure maintains three healthy replicas of customers' data constantly in both locations.

Database Availability

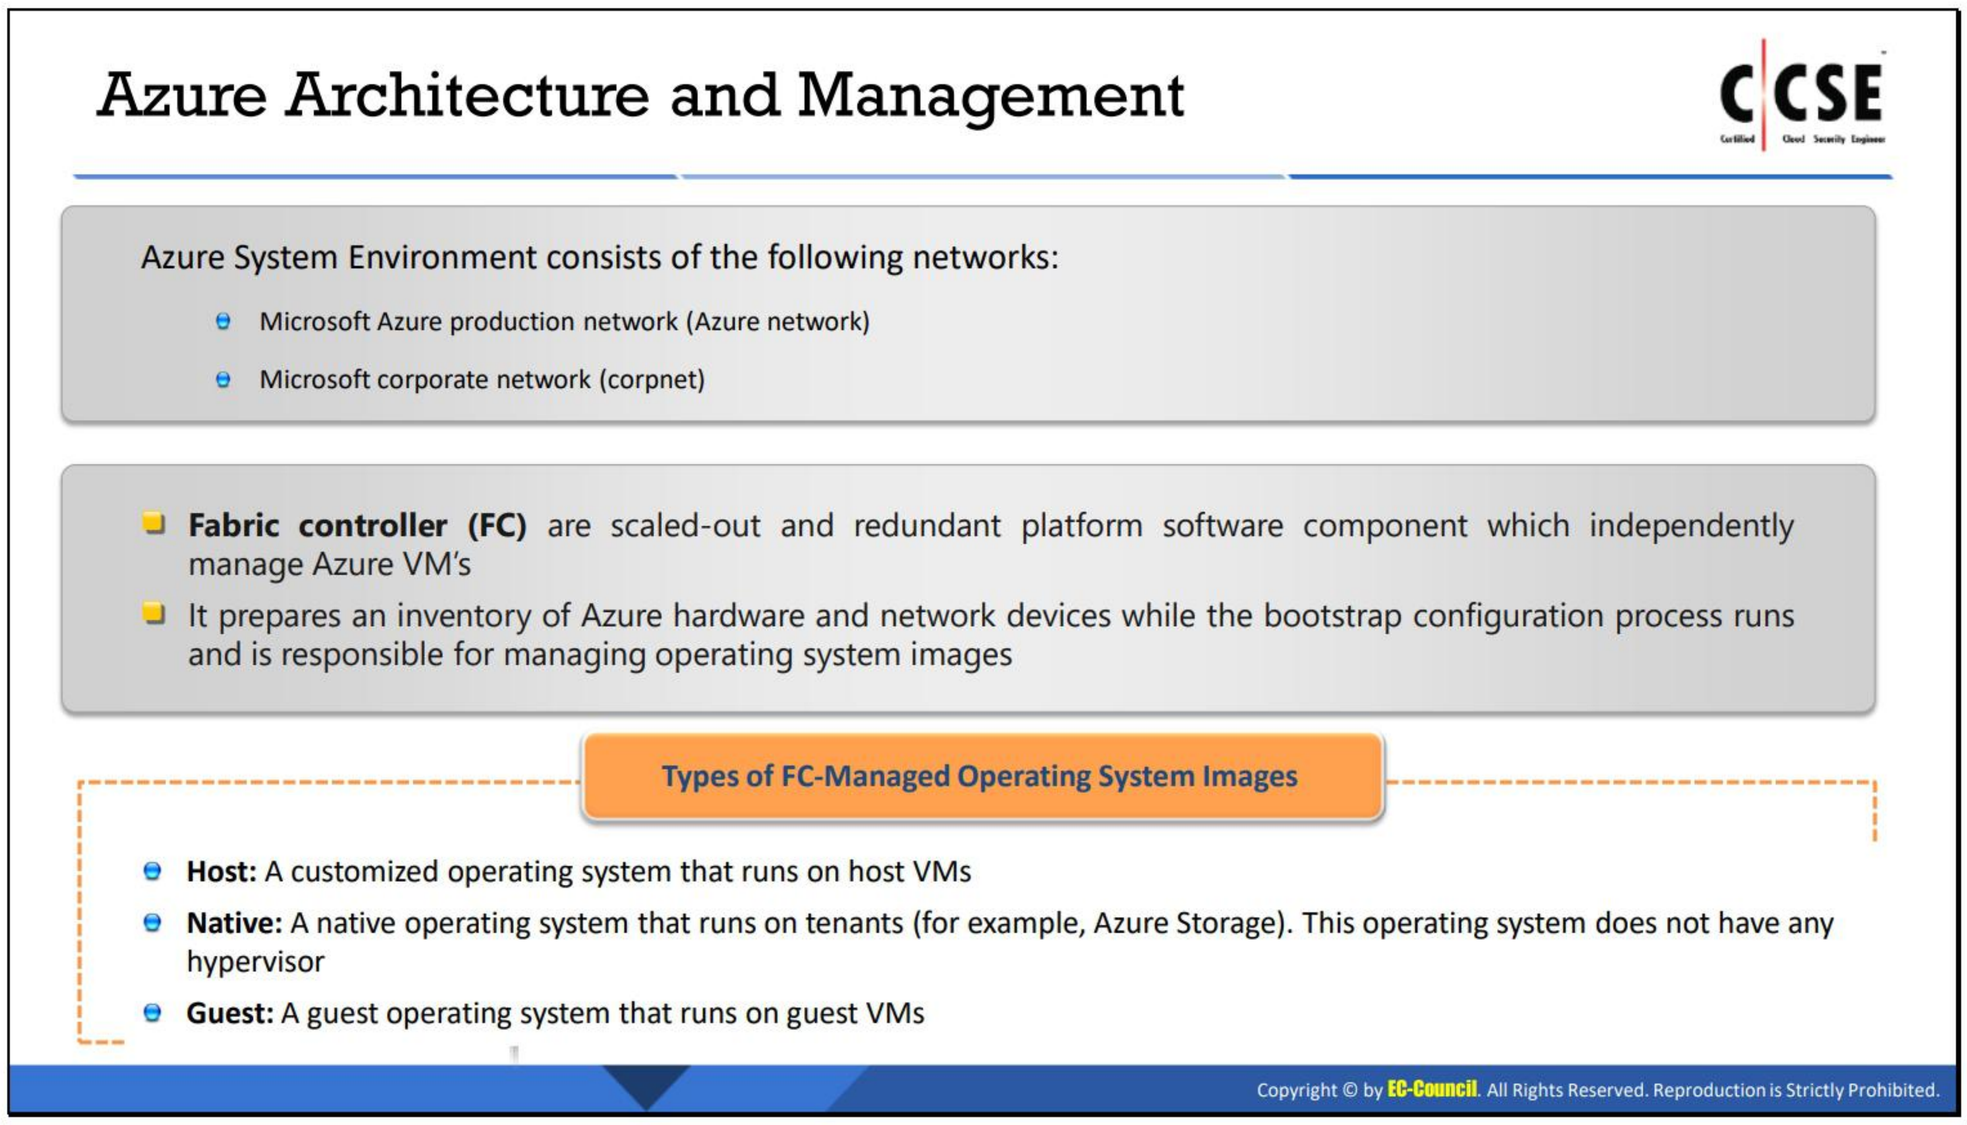
Azure ensures that a database is accessible over the internet through an internet gateway along with prolonged database availability. The health and state of active databases are assessed by monitoring them every five minutes.

Storage Availability

Azure offers storage through a storage service that is highly scalable and durable and provides connectivity endpoints. This means the storage service can be directly accessed by an application. Incoming storage requests are processed with transactional integrity by the storage service.

hide01.ir

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



Azure Architecture and Management

The Azure system environment is composed of two networks:

- Microsoft Azure production network (Azure network)
- Microsoft Corporate network (corpnet)

Separate IT teams hold the responsibility for operating and maintaining these networks.

Azure Architecture

Azure is a cloud computing platform and infrastructure that is utilized to build, deploy, and manage services and applications through data centers. These data centers are managed by Microsoft. Azure creates VMs based on the number of resources specified by customers. These VMs run on an Azure hypervisor which cannot be accessed by the public and is designed for use in the cloud.

A hypervisor runs directly over the hardware on each Azure physical server node, which divides the node into several guest VMs. In addition, every node has a root VM that runs the host OS and each VM has Windows Firewall enabled. By configuring the service definition file, customers define which ports are addressable. The customer-defined ports are only open and addressable internally or externally. The hypervisor and root OS mediate all traffic and access to the network and disk.

Azure VMs run a hardened and customized version of the latest Windows Server at the host layer. The version of Windows Server used by Azure includes only the components that are necessary to host VMs. This reduces the attack surface and improves performance. The hypervisor enforces machine boundaries, and it is not dependent on OS security.

Azure Management by Fabric Controllers

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

VMs that run on physical servers (nodes/blades) in Azure are grouped into clusters of approximately 1000. A redundant and scaled-out platform software component known as the fabric controller (FC) manages the VMs independently.

Every FC manages the lifecycle of applications that run in its cluster and provisions and monitors the hardware health under its control. An FC runs automatic operations, for example, when it determines that a server has failed, it reincarnates VM instances on healthy servers. It also performs application-management operations like updating, deploying, and scaling out applications.

The datacenter is segmented into clusters that isolate flaws or faults at the FC-level and stop certain error classes from affecting servers outside the cluster they occur in. FCs serving a specific Azure cluster are grouped into an FC cluster.

Hardware Inventory

During the bootstrap configuration process, the FC makes an inventory of Azure hardware and network devices. The boot configuration process must be followed by every new hardware and network component that enters the Azure production environment. Managing the complete inventory listed in the datacenter.xml configuration file is the responsibility of the FC.

FC-managed Operating System Images

In the Azure production environment, the operating system team provides images, in the form of virtual hard disks (VHDs) that are deployed on all host and guest VMs. The team uses an automated offline build process to construct base images. A version of the OS where the kernel and other core components are modified and optimized to support the Azure environment is called a base image.

There are mainly three kinds of fabric-managed OS images.

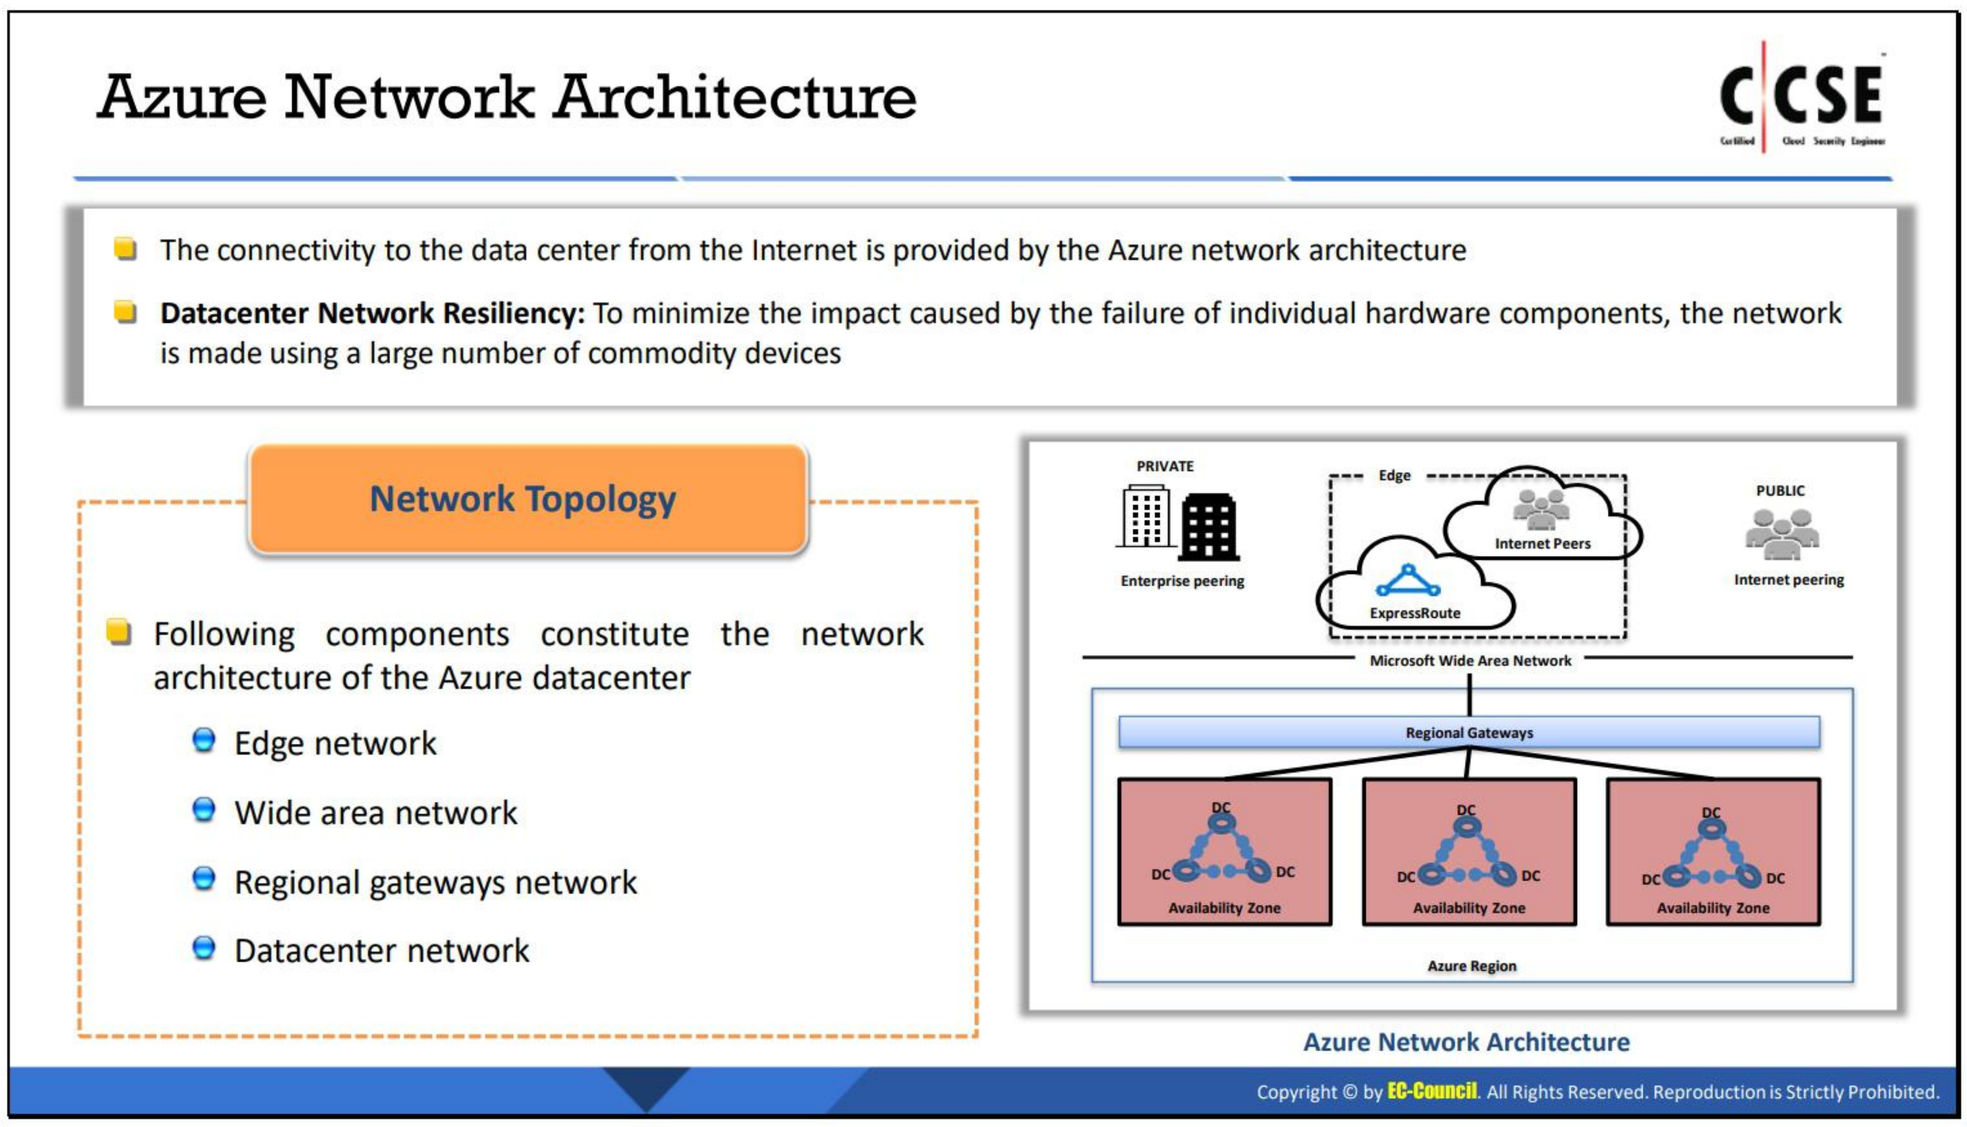
- **Host**: It is a customized OS that runs on host VMs.
- **Native**: It is a native OS that runs on tenants (e.g., Azure storage) and it does not have any hypervisor.
- **Guest**: It is a guest OS that runs on guest VMs.

The host and native FC-managed OSes are not publicly accessible and are designed for use in the cloud.

Host and Native Operating Systems: Host and native are hardened OS images, which host the fabric agents. They run on storage nodes and a compute node (runs on the node as the first VM). Using the optimized base images of native and host has the advantage of reducing the surface area exposed by unused components and APIs because they can cause high-security risks and increase the OS footprint. Operating Systems that have reduced footprint include only the components that are essential to Azure.

Guest Operating System: Azure internal components running on guest OS VMs would not run RDP. The changes required to baseline configuration settings will be through the change and release management process.

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



Azure Network Architecture

The connectivity to the data center from the Internet is provided by the Azure network architecture. Workloads deployed (PaaS, IaaS, and SaaS) on Azure leverage the Azure datacenter network.

hide01.ir

Network Topology

Following are the components that constitute the network architecture of the Azure datacenter.

- Edge network
- Wide area network
- Regional gateways network
- Datacenter network

Network Components

Edge Network: It serves as a demarcation point between Microsoft networking and other networks such as an enterprise network and the internet.

Wide Area Network: It is the intelligent backbone network of Microsoft that covers the whole world and provides connectivity between Azure regions.

Regional Gateway: A regional gateway provides connectivity between datacenters inside an Azure region and serves as a point of aggregation for all data centers in the Azure region.

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

Datacenter Network: The datacenter network offers connectivity between servers with low oversubscribed bandwidth inside the datacenter.

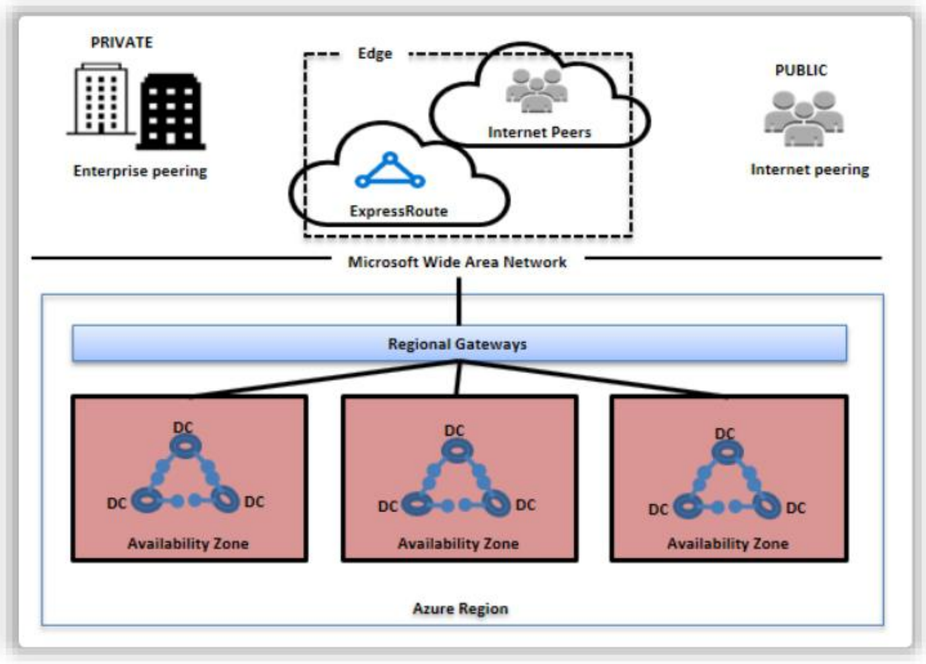


Figure: 2.75: Azure Network Architecture

Microsoft has designed all the above-mentioned components to provide maximum availability to support its always-on, always-available cloud business. Microsoft has designed and built redundancy into the network (ranging from the physical aspect to control protocol).

Datacenter Network Resiliency

The Azure datacenter network is an enhanced version of a Clos network that provides high bi-directional bandwidth for cloud-scale traffic. To minimize the impact caused by the failure of individual hardware components, the network is made using a large number of commodity devices. These commodity devices are located strategically in different physical locations where they have separate power and cooling domains to reduce the impact of any environmental event. All network devices run as OSI Model Layer 3 routing mode on the control plane, which solves or eliminates the issue of traffic loop. All paths between different tiers are functional to offer high redundancy and bandwidth by utilizing equal-cost multi-path (ECMP) routing.

The below diagram shows that different tiers of network devices are used to construct the datacenter network. The bars shown in the diagram represent groups of network devices that offer high-bandwidth connectivity and redundancy.

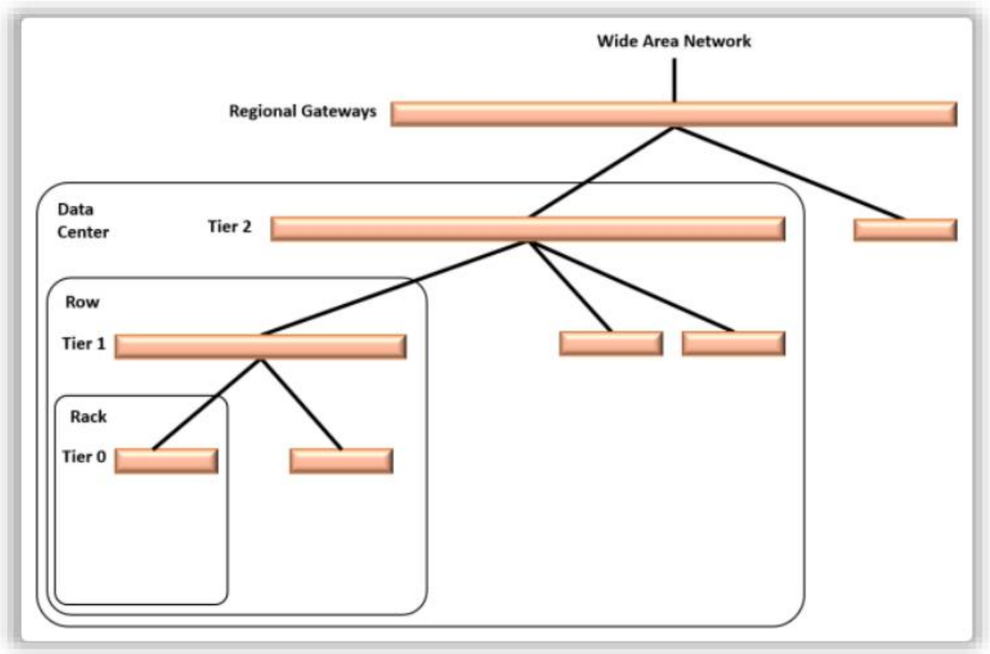
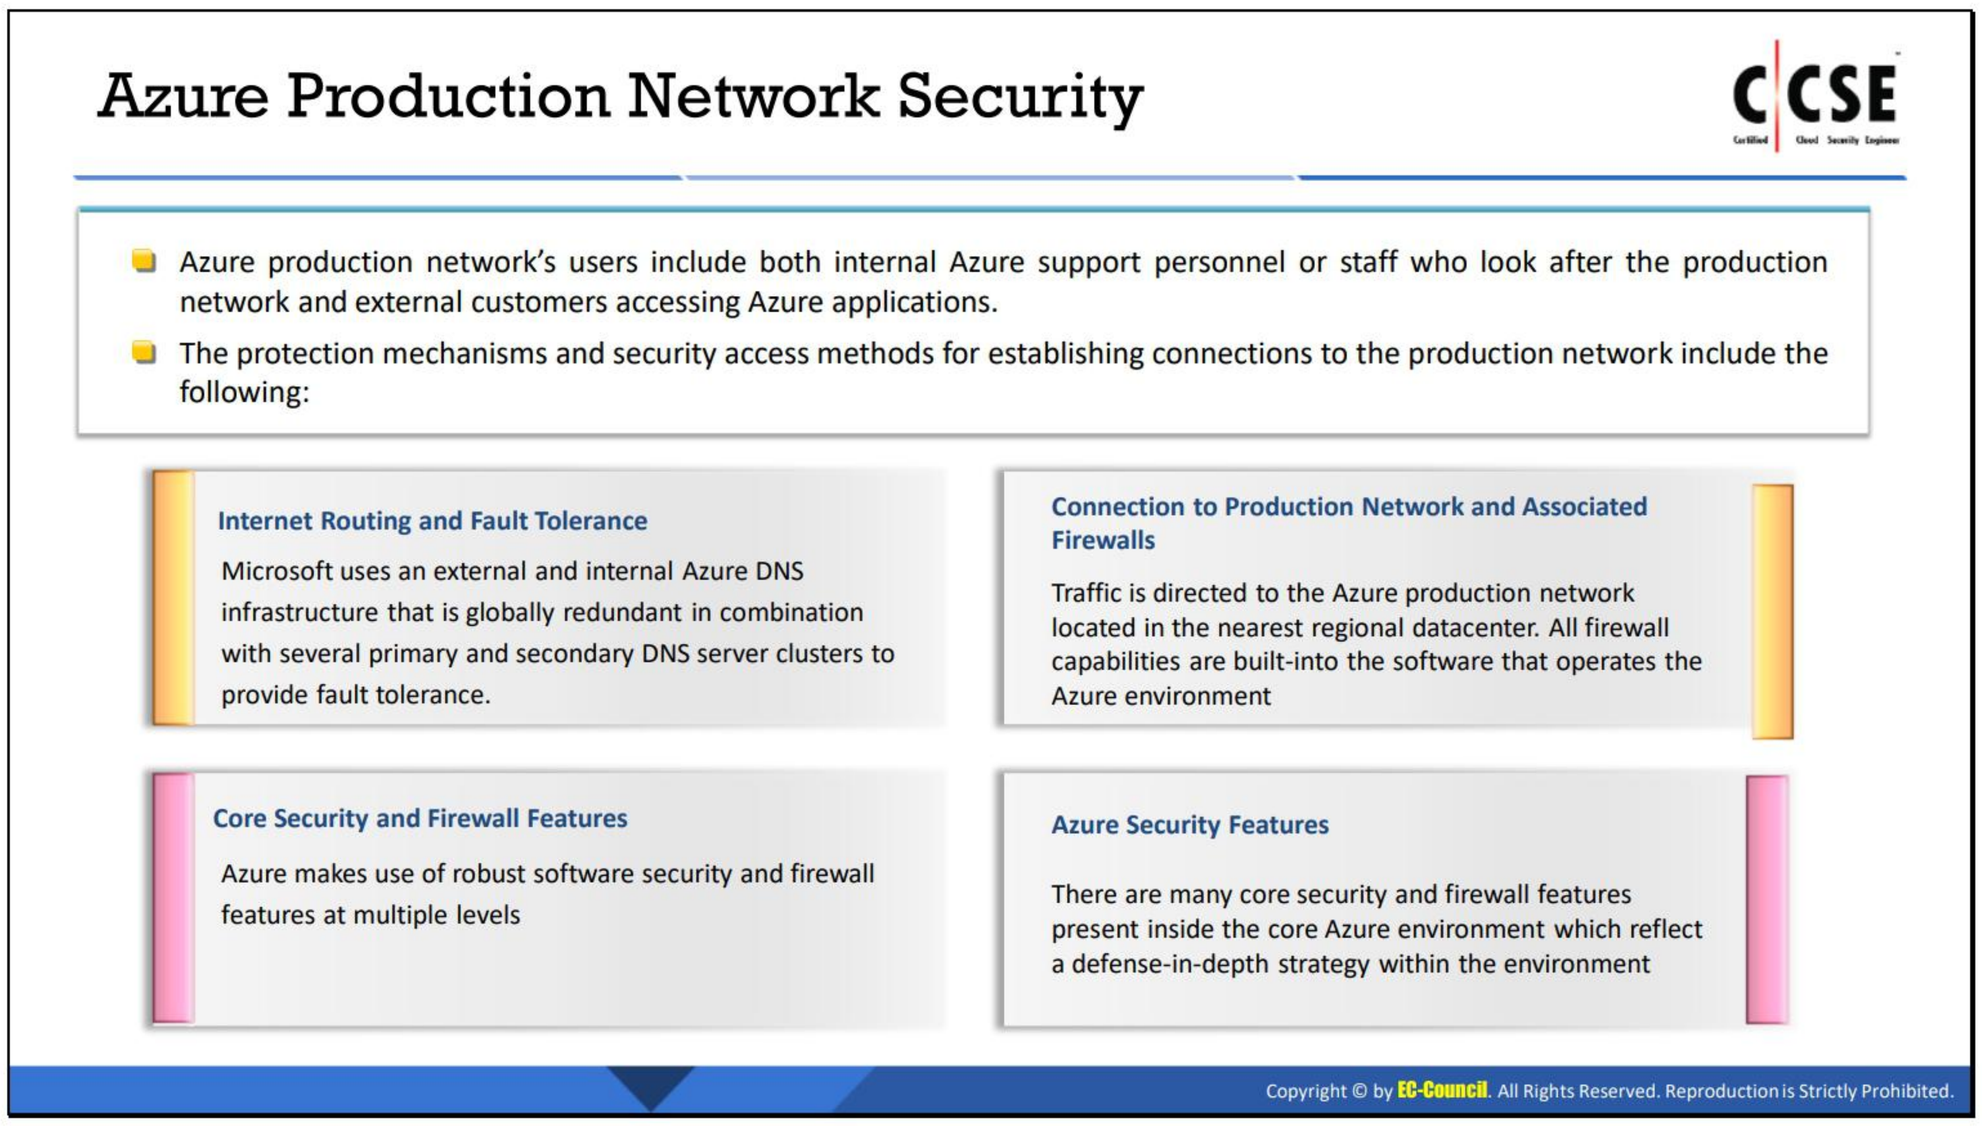


Figure: 2.76: Network Topology in Azure

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



Azure Production Network Security

Azure production network's users include both internal Azure support personnel or staff who look after the production network and external customers accessing Azure applications. The protection mechanisms and security access methods for establishing connections to the production network include.

hide01.ir

Internet Routing and Fault Tolerance

Microsoft uses an external and internal Azure DNS infrastructure that is globally redundant in combination with several primary and secondary DNS server clusters to provide fault tolerance. It uses additional network security controls such as NetScaler at the same time to avoid DDoS attacks and safeguard the Azure DNS services' integrity.

Multiple datacenter facilities host Azure DNS servers. To resolve Azure customer domain names publicly, a hierarchy of primary and secondary DNS servers is incorporated in the Azure DNS implementation. Usually, the domain names resolve to a CloudApp.net address. The virtual IP (VIP) address is wrapped by this address for the customer's service.

Azure is hosted within the US in geographically distributed data centers and built on advanced routing platforms that follow scalable, robust architectural standards. Its notable features include

- It has multiprotocol label switching (MPLS)-based traffic engineering feature that offers efficient link utilization. If there is an outage, it provides graceful degradation.
- All Azure networks are implemented with N+1 redundancy architectures or better.

Externally, high-bandwidth, dedicated network circuits serve Azure datacenters and connect properties redundantly with over 1200 ISPs worldwide at several peering points. This offers more than 2,000 Gbps of edge capacity across the network.

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

As all network circuits among data centers are owned by Microsoft, it helps the Azure offering to achieve 99.9+% network availability without requiring any external or third-party ISPs.

Connection to Production Network and Associated Firewalls

Traffic is directed to the Azure production network located in the nearest regional datacenter inside the US by the Azure network internet traffic flow policy. The traffic flow applies to all datacenters consistently because the production datacenters maintain consistent network hardware and architecture.

After routing the internet traffic for Azure to the nearest datacenter, a connection to the access routers is established. The access routers isolate traffic between customer-instantiated VMs and Azure nodes. Network devices that are deployed at access and edge locations serve as boundary points where egress and ingress filters are applied. Microsoft configures these routers by using a tiered ACL to apply traffic rate limits (if necessary) and prevent unwanted network traffic. Traffic allowed by ACL is routed to the load balancers. Distribution routers allow Microsoft approved IP addresses only, offer anti-spoofing, and establish TCP connections using ACLs.

Located behind the access routers, external load-balancing devices perform NAT from internet-routable IPs to Azure internal IPs. These load-balancing devices also help to route packets to valid production internal IPs and ports and limit exposing the internal production space by acting as a protection mechanism.

For all traffic transmitted to web browsers of customers, Microsoft enforces HTTPS by default. This includes sign-in and all traffic thenceforth. Using TLS v1.2 provides a secure tunnel for traffic to pass through. ACLs applied on core and access routers make sure that the traffic source is consistent with what is expected.

h1e0.ir

An important difference in this architecture is that there are no specialized intrusion that or prevention devices, dedicated hardware firewalls, or other security appliances which are expected normally before the establishment of connections to the Azure production environment. Usually, customers expect hardware firewall devices in Azure but it does not employ any. These security features are built into the software that operates the Azure environment to offer multi-layered, robust security mechanisms that include all firewall capabilities. Additionally, it is easier to manage the scope of the boundary and associated sprawl of essential security devices because the software that runs Azure manages it.

Core Security and Firewall Feature

Azure makes use of robust software security and firewall features at multiple levels to enforce expected security features in a traditional environment for the protection of the core security authorization boundary.

Azure Security Features

Inside the production network, Azure uses host-based software firewalls. There are many core security and firewall features present inside the core Azure environment, which reflect a defense-in-depth strategy within the environment. The firewalls that protect customer data in Azure are described below.

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

- **Hypervisor Firewall (Packet Filter)**: In the hypervisor, the firewall is implemented and FC configures it. The tenant that operates inside the VM is protected by this firewall from any unauthorized access. When a VM is created, all (incoming and outgoing) traffic is blocked by default. Rules and exceptions to allow authorized traffic are added in the filter by the FC agent. Two types of rules are programmed by the FC agent.
  - **Machine Configuration or Infrastructure Rules**: All communication is by default blocked. Some exceptions allow a VM to send and receive DNS information and DHCP communications and send traffic to the public internet outbound to the OS Activation server and other VMs present within the FC cluster. These rules act as a defensive layer for other Microsoft properties and Azure router subnets because they are not included in the VMs' allowed list of outgoing destinations.
  - **Role Configuration File Rules**: These rules define the inbound ACLs based on the service model of tenants. For example, on a particular VM, if a tenant has a web front end on port 80, this port is opened for all IP addresses. If a worker role is running in the VM, the role is opened only to the VM that is within the same tenant.
- **Native Host Firewall**: A native OS has no hypervisor. Azure Storage and Azure Service Fabric run on the native OS so the Windows Firewall is configured with the above-mentioned two sets of rules.
- **Host Firewall**: The host partition that runs the hypervisor is protected by the host firewall. The rules of the host firewall allow only jump boxes and the FC to talk to the host partition on a particular port. Other exceptions allow DNS replies and DHCP responses. For the host partition, Azure utilizes a machine configuration file that contains a template of firewall rules. VMs are allowed by a host firewall exception to communicate with host components, the metadata server, and the wire server via specific ports/protocols.
- **Guest Firewall**: It is the Windows Firewall piece of the guest OS that customers can configure on their VMs and storage.

Certain additional features/capabilities that are included in Azure are as follows:

- IP addresses are assigned to infrastructure components from DIPs. Attackers cannot direct traffic to these addresses because they cannot reach Microsoft. They cannot enter the production network because internet gateway routers filter out packets that are addressed only to internal addresses. Load balancers are the only components that accept traffic directed to VIPs.
- For any given scenario, there are three basic security architecture considerations in firewalls that are implemented on the internal nodes.
  - Firewalls accept packets from anywhere in this security consideration and are located behind the load balancer. These packets are meant to be exposed externally and correspond to the open ports in a normal perimeter firewall.
  - In this consideration, packets are accepted by firewalls from a limited set of addresses as part of the defensive in-depth strategy against denial-of-service attacks, and such connections are authenticated cryptographically.

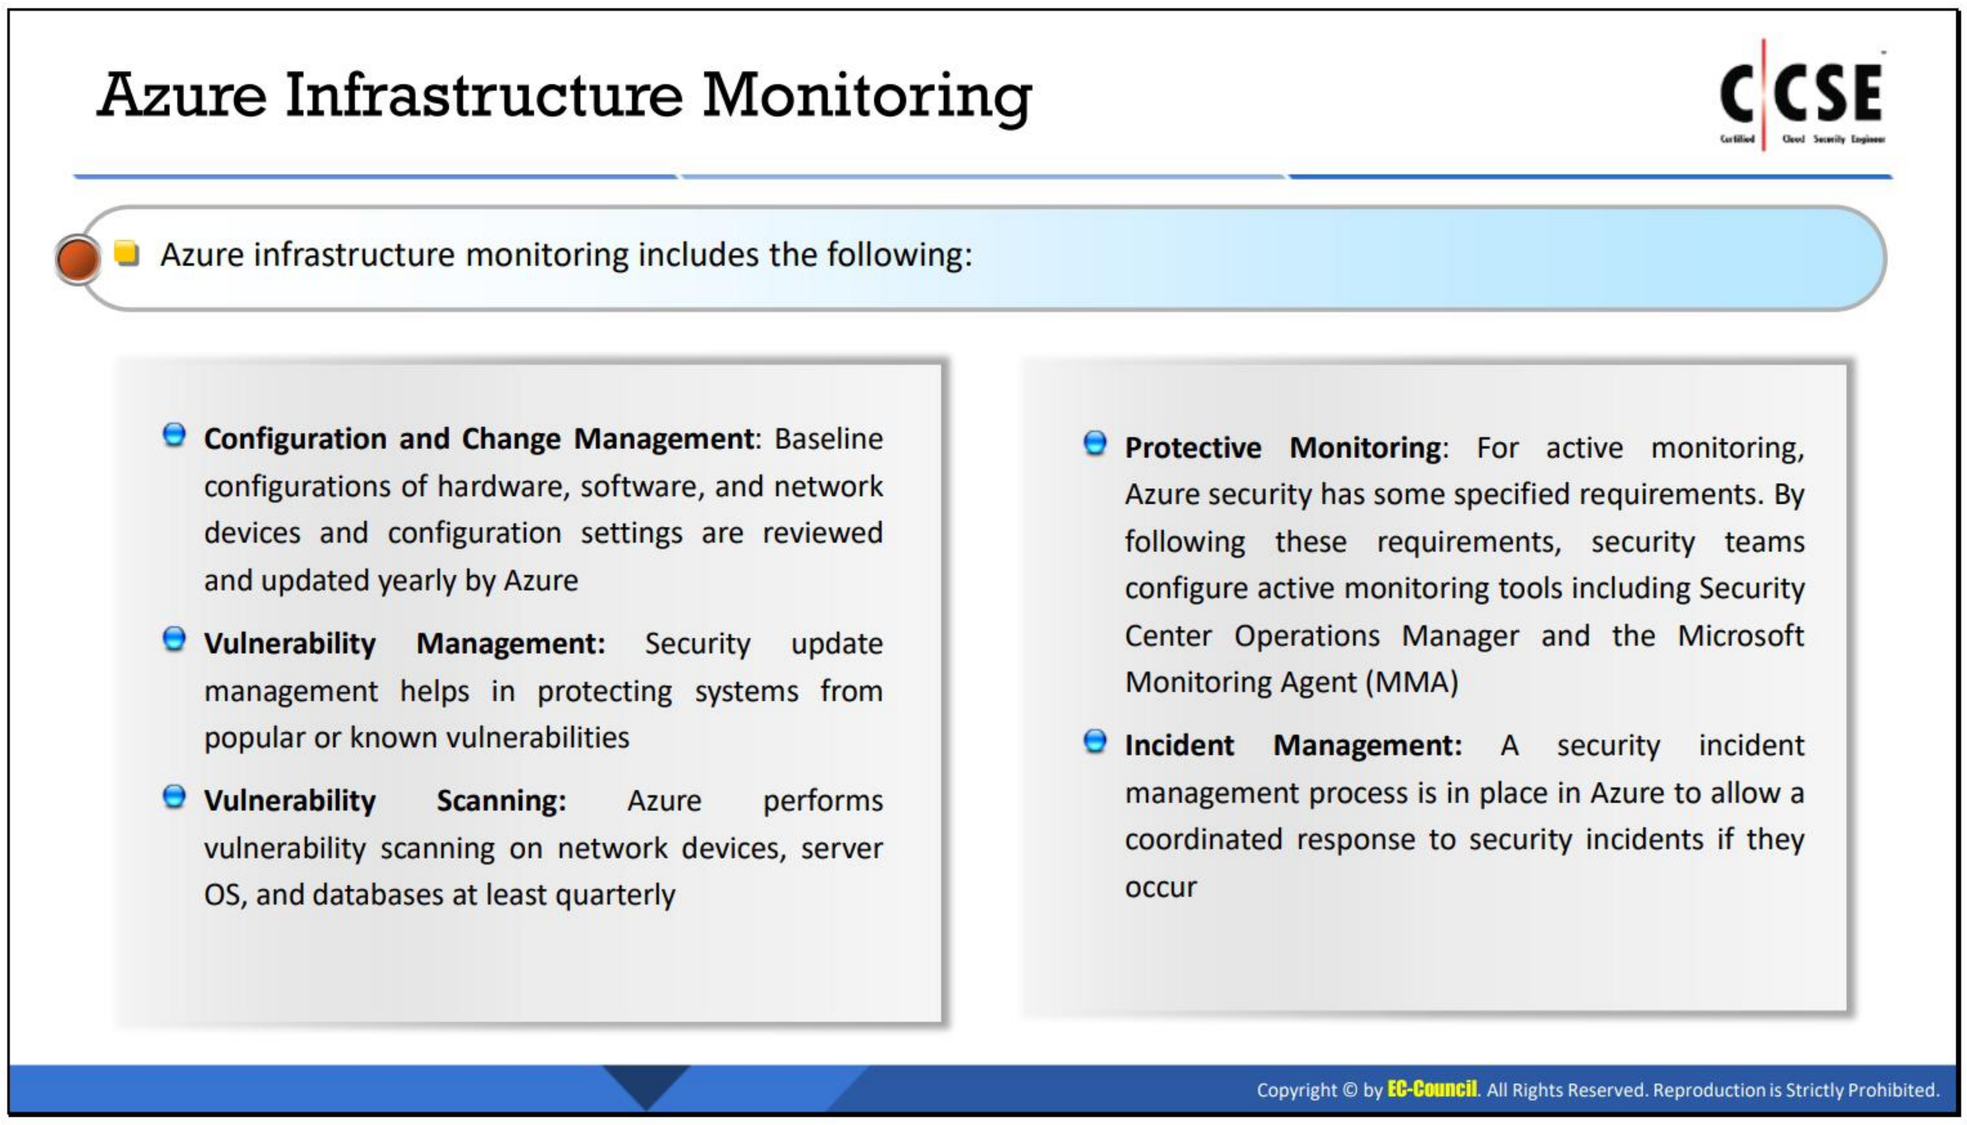
Certified Cloud Security Engineer Copyright © by **EC-Council**
All Rights Reserved. Reproduction is Strictly Prohibited.

- In this consideration, only select internal nodes can access firewalls, which accept packets only from a specified list of source IP addresses. All the source IP addresses are DIPs within the Azure network. For example, an attack on the network can direct requests to these addresses but it will be blocked because its source IP address is not specified in the list. By default, all attacks are blocked within the Azure network unless the packet's source IP address is present in the specified list.

At the perimeter, the access router blocks outbound packets that are directed to an address that is present within the Azure network due to its configured static routes.

hide01.ir

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



Azure Infrastructure Monitoring

Azure infrastructure monitoring includes the below-mentioned points.

- **Configuration and Change Management:** Baseline configurations of hardware, software, and network devices and configuration settings are reviewed and updated yearly by Azure. In the development environment, changes are developed, tested, and approved before they are implemented in the production environment.

  Azure security and compliance team and service teams review the required baseline configurations for Azure-based services. The service teams conduct reviews as part of testing that is done before their production services are deployed.

- **Vulnerability Management:** Security update management helps in protecting systems from popular or known vulnerabilities. Azure manages the distribution and installation of security updates for Microsoft software by using integrated deployment systems. Azure can also use Microsoft Security Response Center (MSRC) resources. The MSRC detects, monitors, responds to, and mitigates cloud vulnerabilities and security incidents around the clock, throughout the year.

- **Vulnerability Scanning:** Azure performs vulnerability scanning on network devices, server OS, and databases at least quarterly. Azure performs penetration testing of its cloud boundary with the help of independent assessors. It also performs red-team exercises from time to time and security improvements are done on the basis of their results.

- **Protective Monitoring**: For active monitoring, Azure security has some specified requirements. By following these requirements, security teams configure active monitoring tools, including Security Center Operations Manager and the Microsoft

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

Monitoring Agent (MMA). Azure configures these tools to provide timely alerts to Azure security staff in cases where immediate action is to be taken.

- **Incident Management:** A security incident management process is in place in Azure to allow a coordinated response to security incidents if they occur.

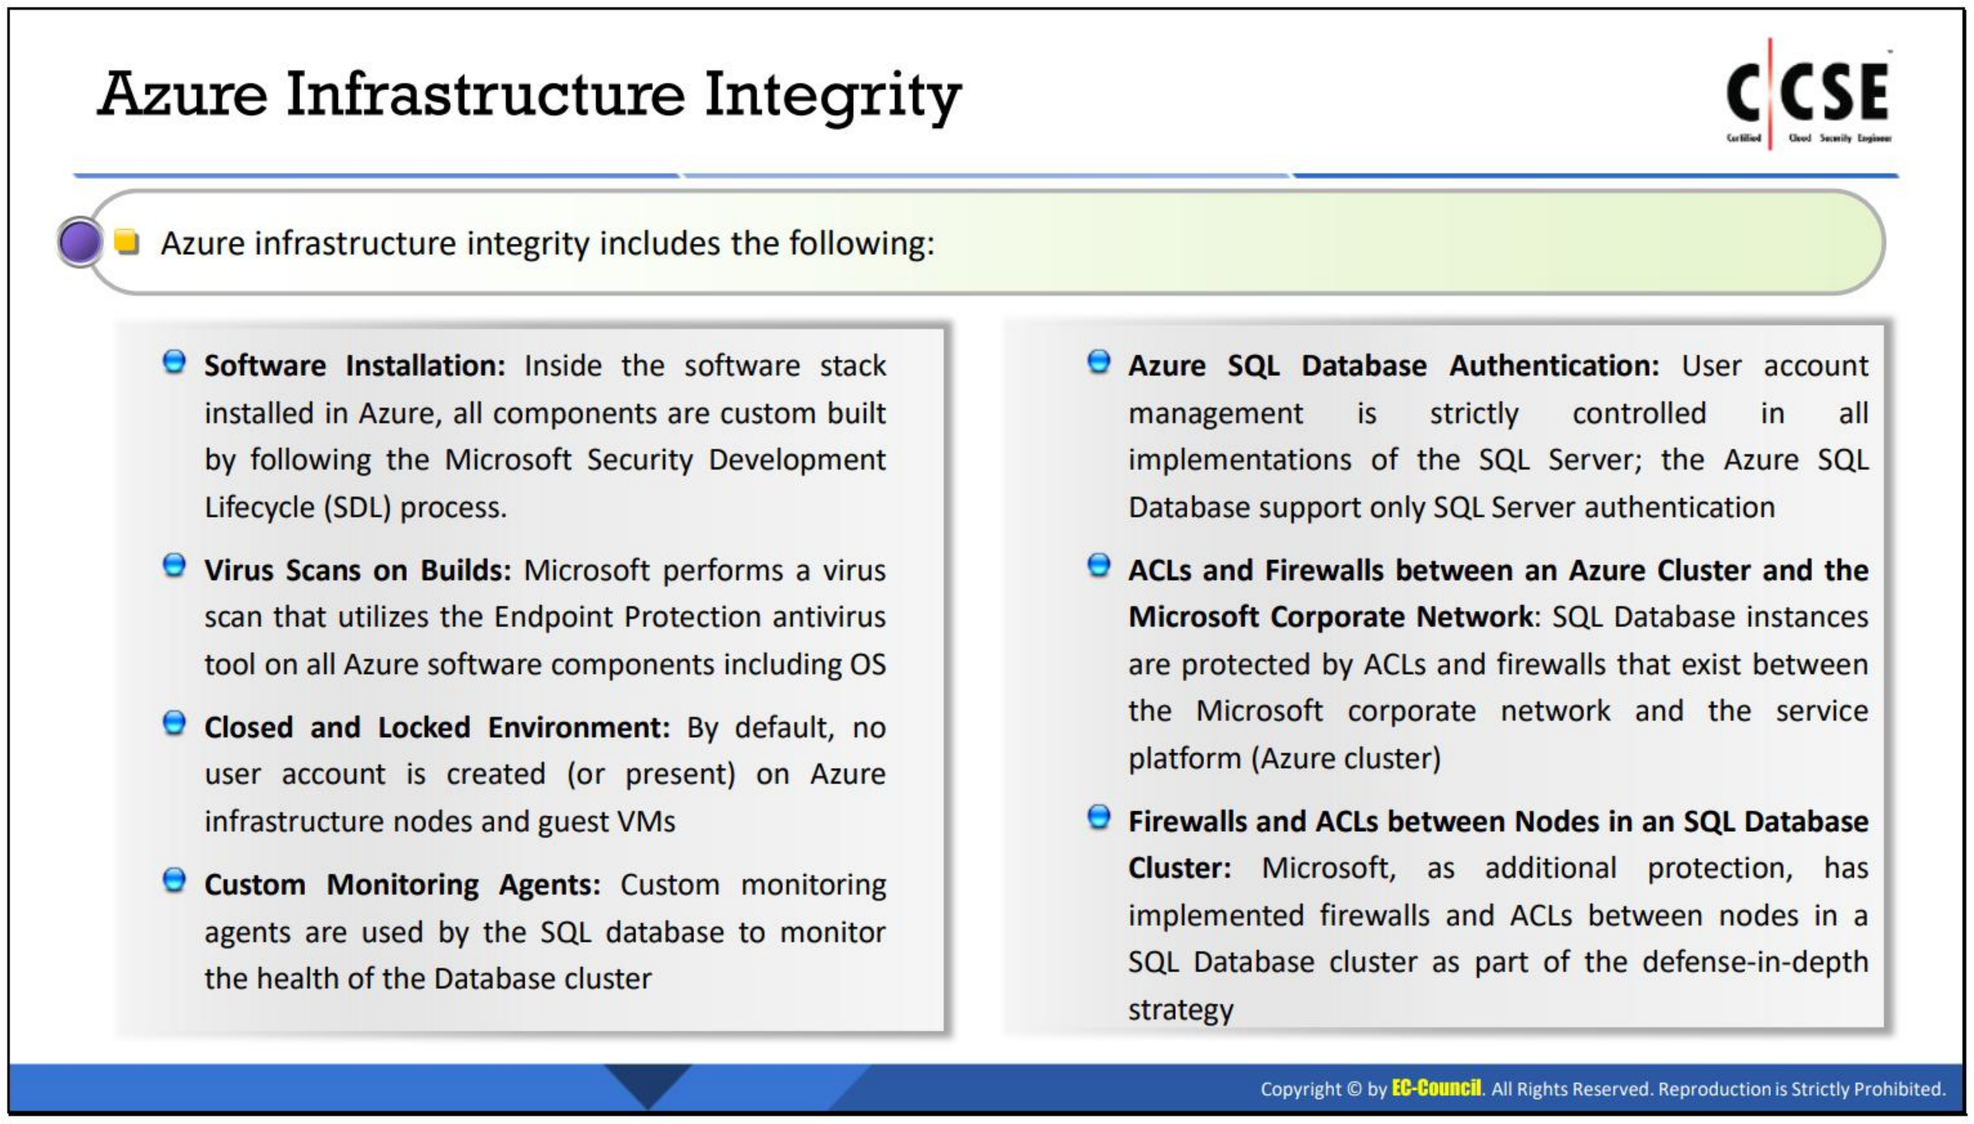
  If Microsoft detects any unauthorized access to customer data stored on its infrastructure or detects any unauthorized access to its infrastructure, it performs the following actions.

  - It alerts the customer about the security incident promptly.
  - It quickly investigates the incident and provides detailed information to customers about the incident.
  - It takes prompt and reasonable measures to mitigate the effects and reduce the damage caused due to the incident.

Microsoft has established an incident management framework that specifies roles and assigns responsibilities. The responsibility of managing security incidents (that include escalation and confirming the involvement of specialist teams when required) lies with the Azure security incident management team. The responsibility of overseeing (supervising) all the investigations and resolutions of security and privacy incidents lies with Azure operations managers.

hide01.ir

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



Azure Infrastructure Integrity

Azure infrastructure integrity involves the below-mentioned points.

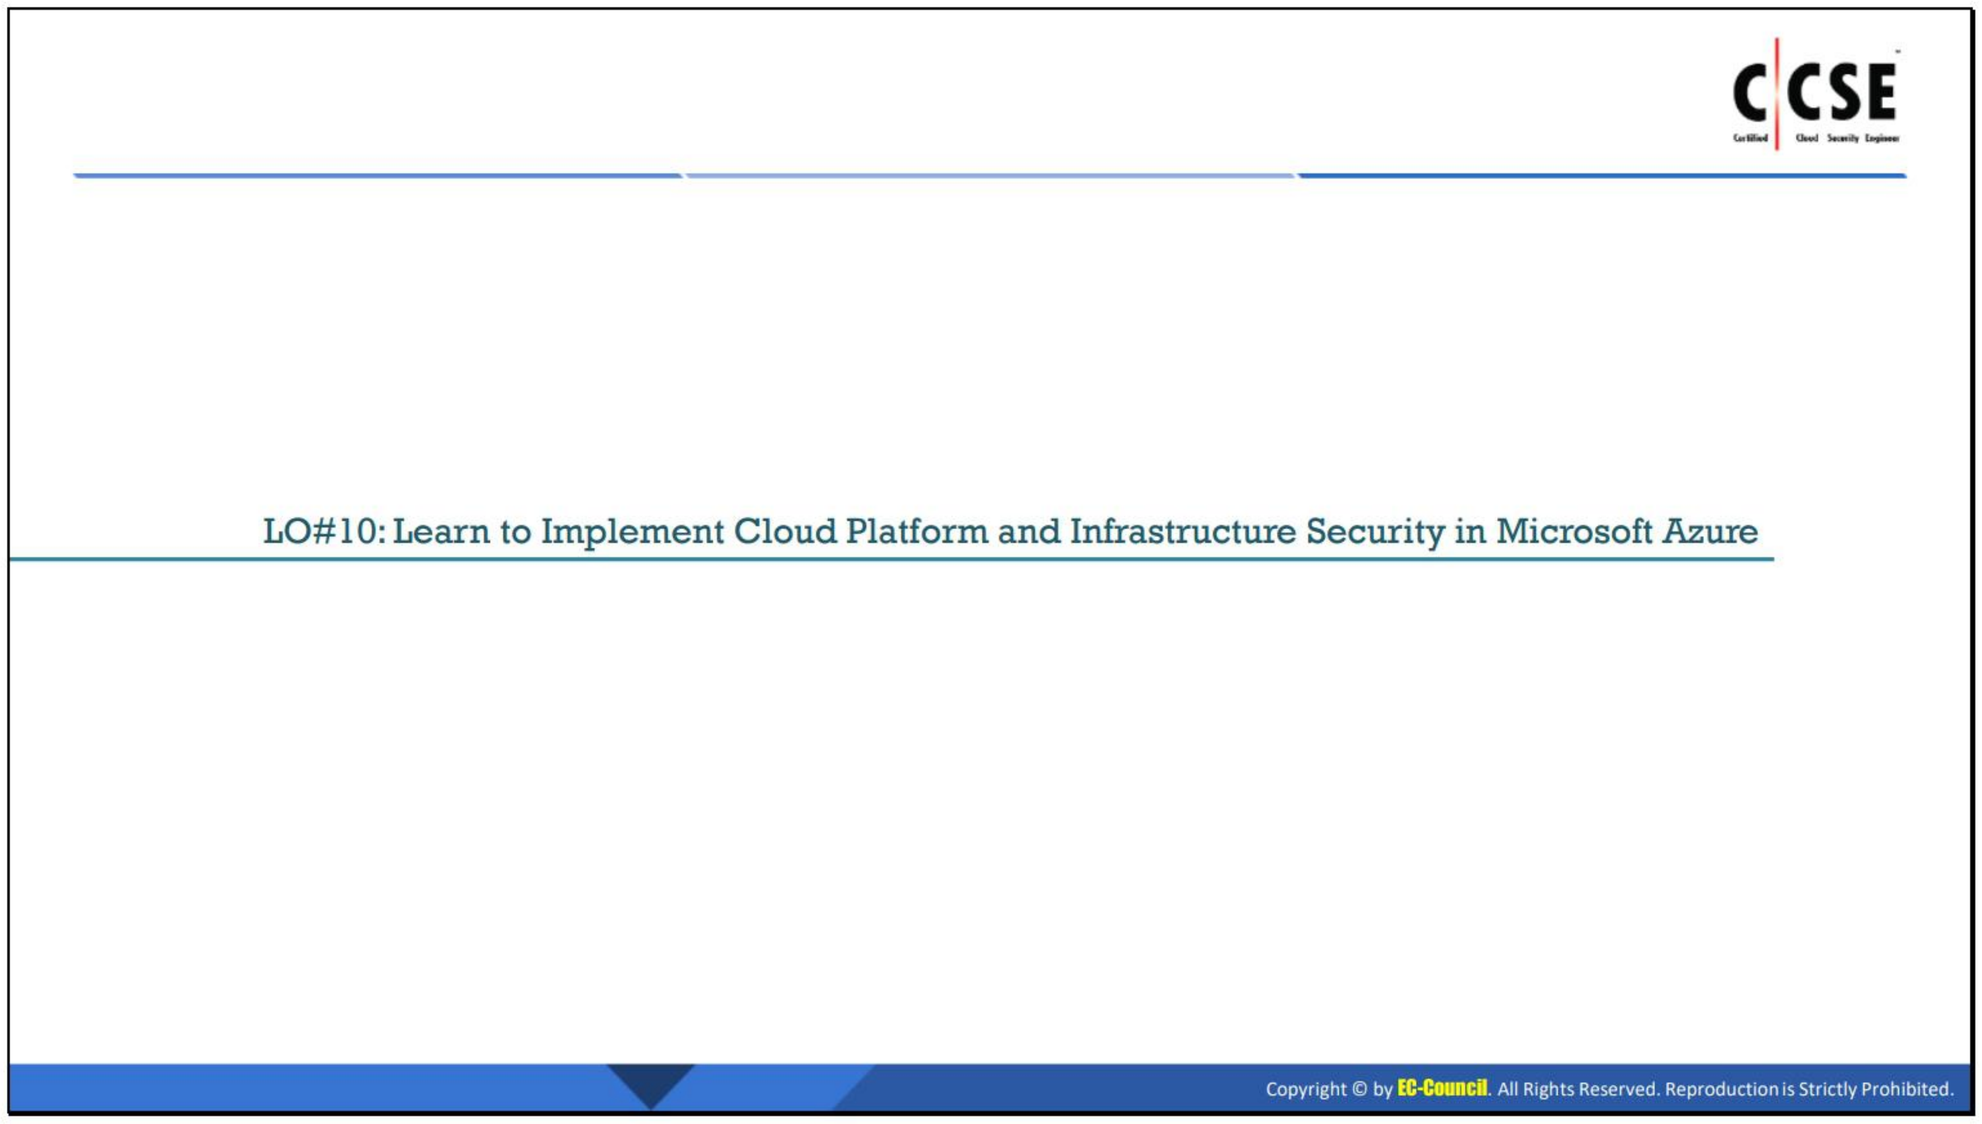
- **Software Installation:** Inside the software stack installed in Azure, all components are custom built by following the Microsoft Security Development Lifecycle (SDL) process. As part of the change management and release management process, all software components like OS images and the SQL Database are deployed. A customized OS version of Windows Server 2008 or Windows Server 2012 runs on all nodes. The FC chooses the exact version of the OS depending on the role it wants the OS to play. Additionally, the host OS prevents the installation of any unauthorized software components.
- **Virus Scans on Builds**: Microsoft performs a virus scan that utilizes the Endpoint Protection antivirus tool on all Azure software components, including OS. A log is created after every virus scan inside the associated build directory that gives detailed information about what was scanned and the scan results. For each component in Azure, the virus scan is a part of its build source code. Code cannot be moved to production until it has a successful and clean scan. If any issues are found in a build, the build is stopped or frozen and sent to the security teams inside Microsoft security to detect where the rogue code got into the build.
- **Closed and Locked Environment:** By default, no user account is created (or present) on Azure infrastructure nodes and guest VMs. Additionally, default Windows administrator accounts are disabled as well. For emergency repairs, administrators from Azure live support can log into the machines after proper authentication and administer the Azure production network.

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

- **Azure SQL Database Authentication:** User account management should be strictly controlled in all implementations of the SQL Server. The only authentication supported by Azure SQL Database is SQL Server authentication. To enhance the customer's data security, customers should have user accounts configured with special rights and strong passwords.
- **ACLs and Firewalls Between an Azure Cluster and The Microsoft Corporate Network:** SQL Database instances are protected by ACLs and firewalls, that exist between the Microsoft corporate network and the service platform (Azure cluster), from unauthorized insider access. In addition, Windows Fabric platform-management endpoint can be accessed only by users who use specific IP address ranges from the Microsoft corporate network.
- **Firewalls and ACLs Between Nodes in A SQL Database Cluster:** Microsoft, as additional protection, has implemented firewalls and ACLs between nodes in a SQL Database cluster as part of the defense-in-depth strategy. All communication that happens inside all running code and the Windows Fabric platform is trusted.
- **Custom Monitoring Agents:** Custom monitoring agents, also known as watchdogs, are used by the SQL Database to monitor the health of the Database cluster.

hide01.ir

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

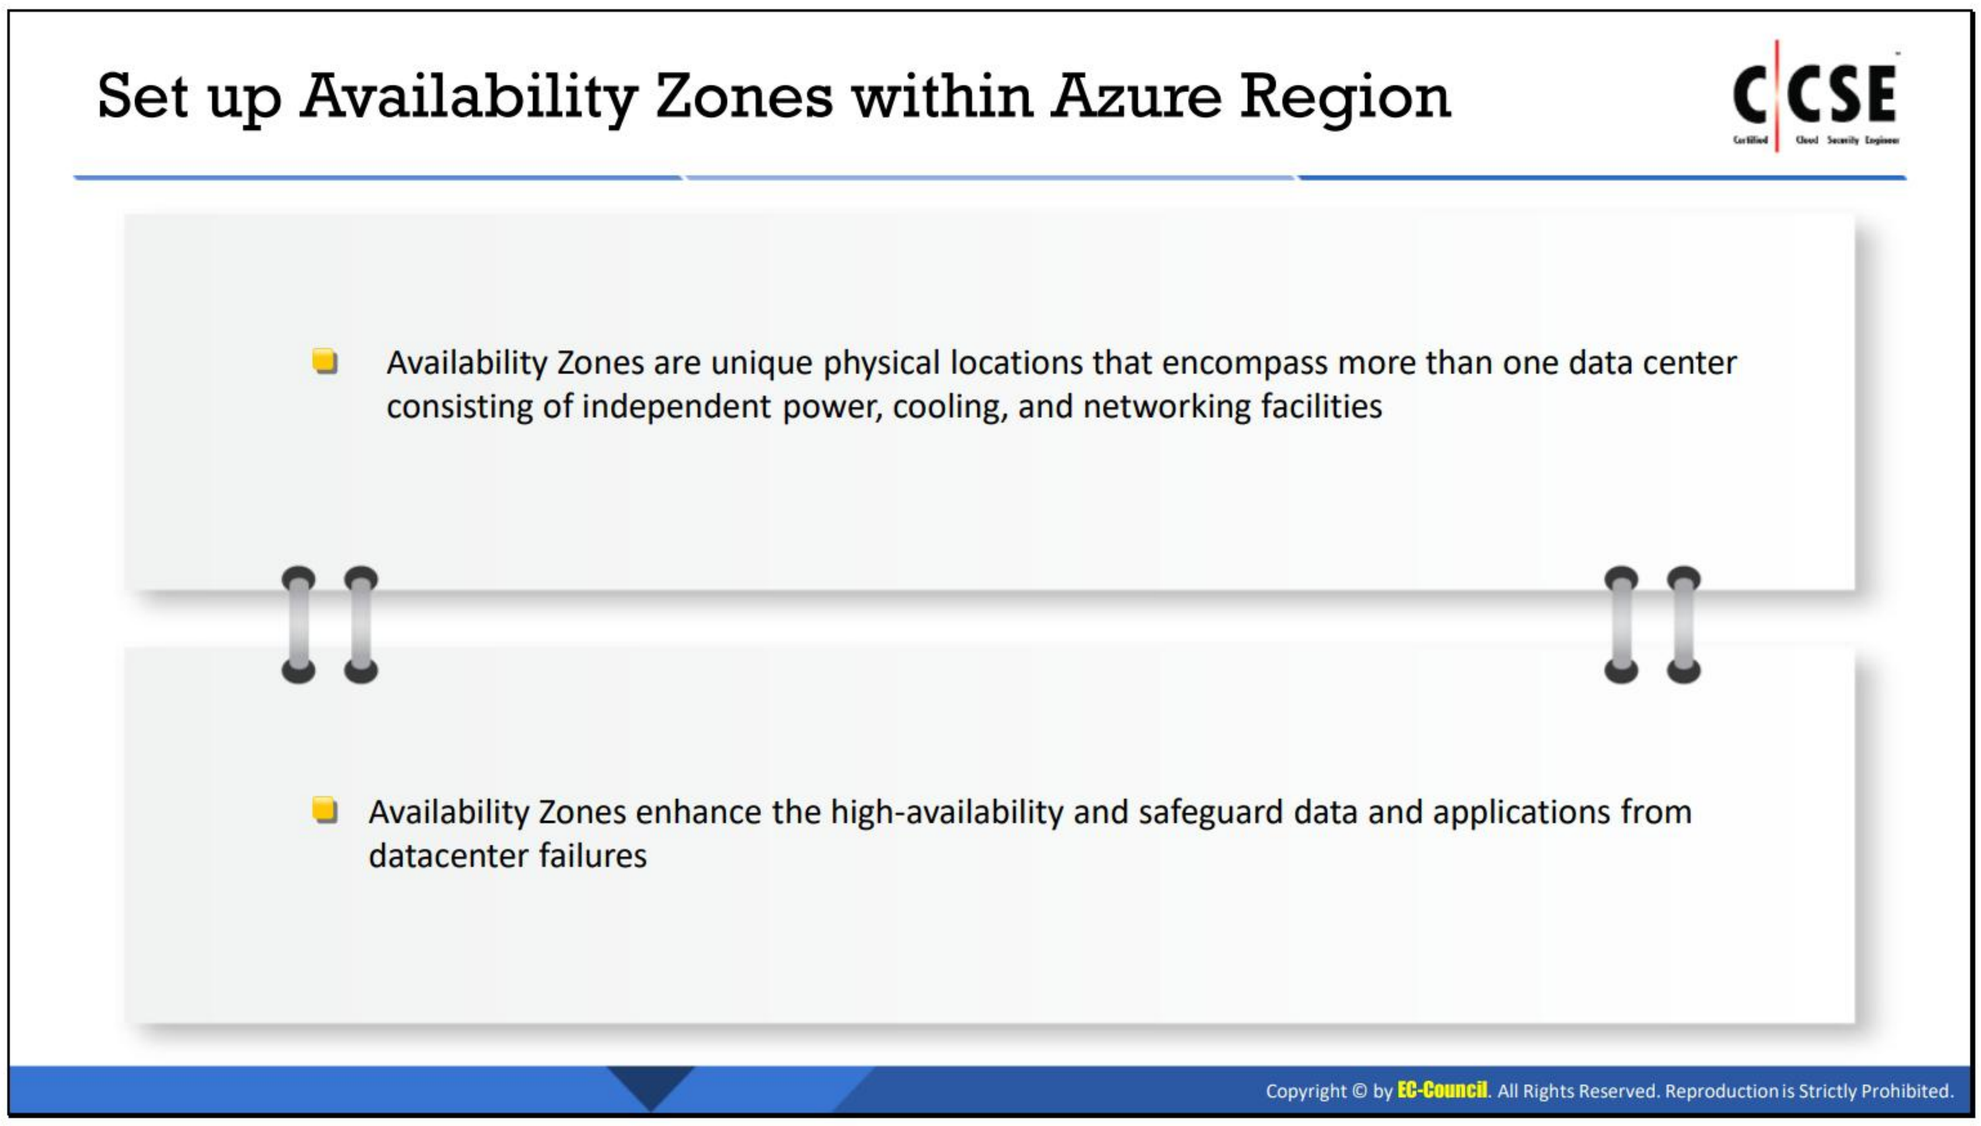


LO#10: Learn to Implement Cloud Platform and Infrastructure Security in Microsoft Azure

In this section, you will learn how to set up security features in cloud platforms and infrastructure with the help of Azure resources.

hide01.ir

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



Set Up Availability Zones within Azure Region

AZs are unique physical locations that encompass more than one data center consisting of independent power, cooling, and networking facilities. AZs enhance the high availability and safeguard data and applications from datacenter failures. Azure maintains a minimum of 3 separate AZs in a region to provide resiliency.

These zones provide protection to applications and data as they are physically isolated so any failure occurring in a zone has no impact on the other zones within the same region. With the help of AZs, Azure successfully provides 99.99% VM uptime SLA.

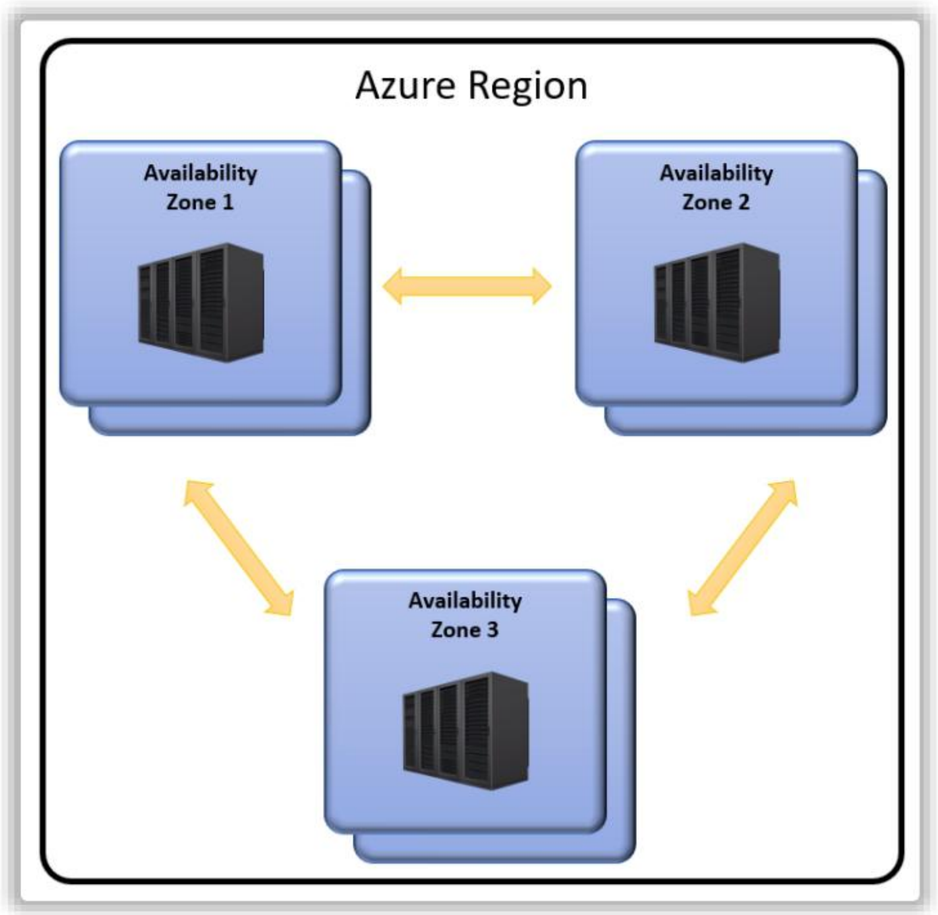


Figure: 2.77: Availability Zone within Azure Region

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

Azure has zero-redundant services that play an important role in offering 99.99% VM uptime SLA by replicating customer data and applications across all AZs within the region, thereby protecting them from single point of failure.

Organizations should set up AZs to increase the availability of data and applications and to protect them from datacenter failures. They should make strategy and incorporate high availability for building migration architecture through sharing storage, compute, networking, and data resources within a zone and replicating it to other zones.

The main categories to explain the services in Azure that supports AZs are listed below.

- **Zonal Services:** The services that are associated with a particular zone like managed disks, VMs, or IP addresses fall in this category.
- **Zero-redundant Services**: The resources that automatically replicate across multiple zones, for example, Azure SQL Database zone-redundant storage fall in this category.

A cloud administrator can deploy an Azure load balancer instance with internet-facing application tiers or workloads to offer zonal fault tolerance.

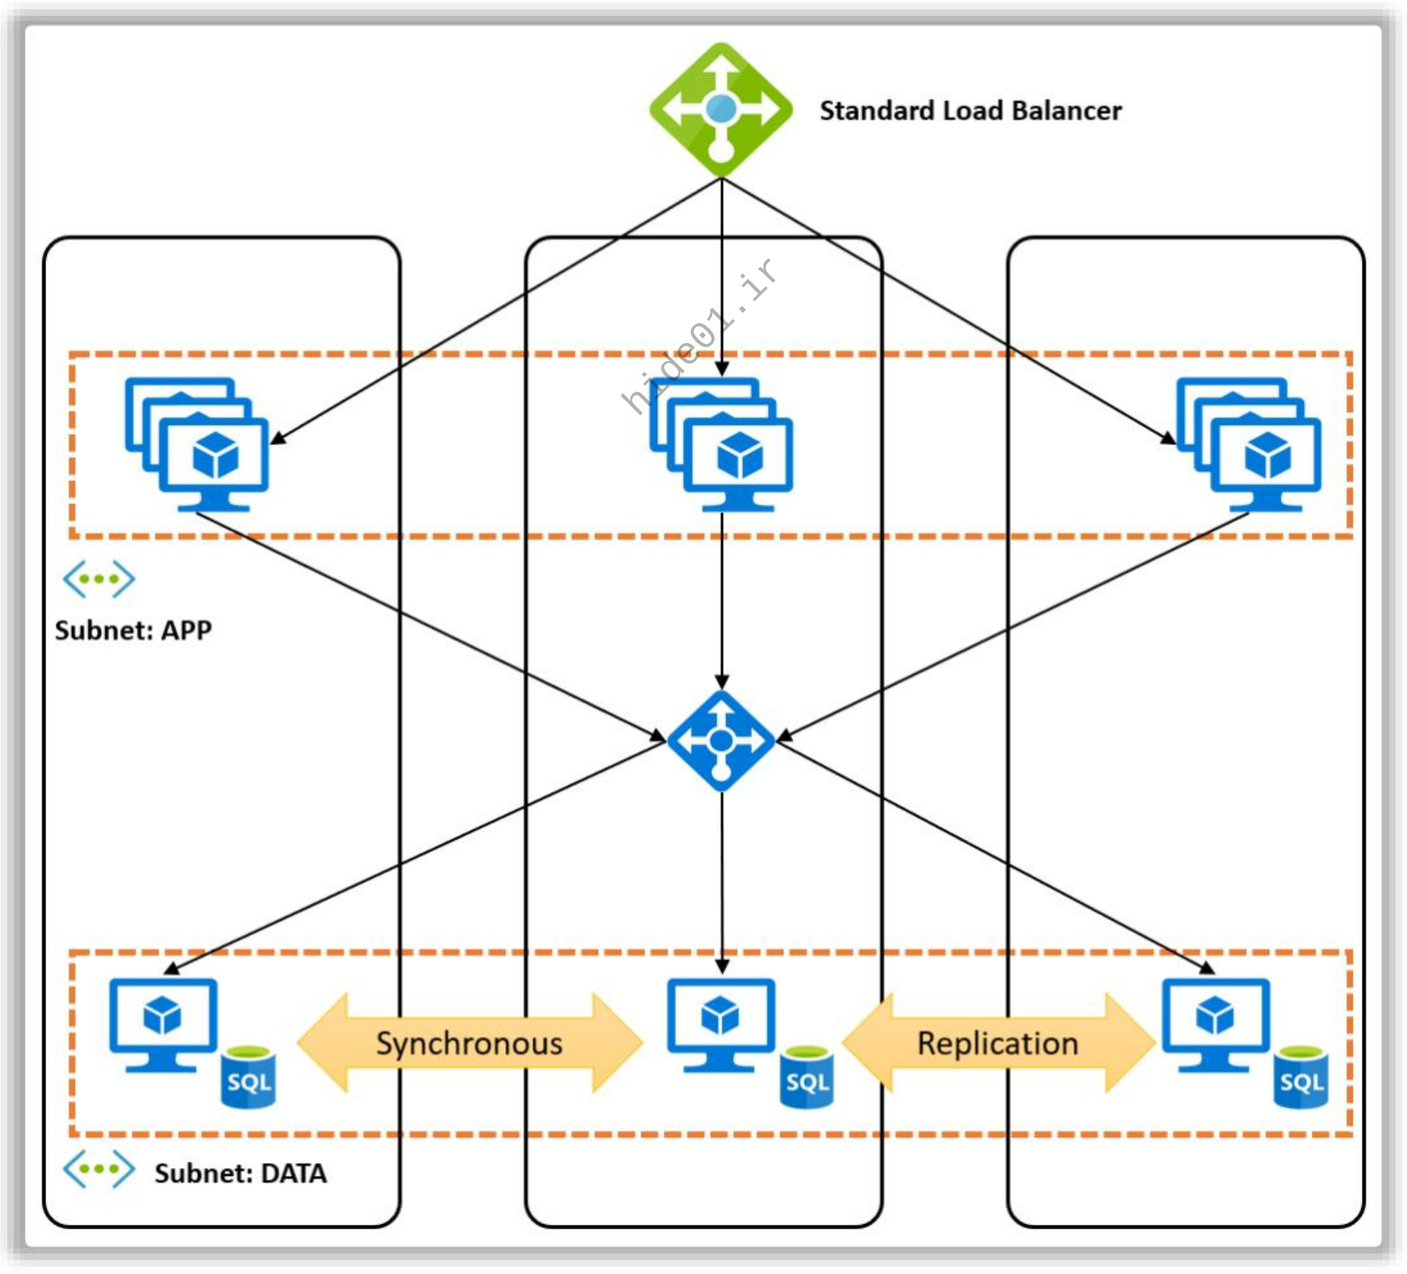
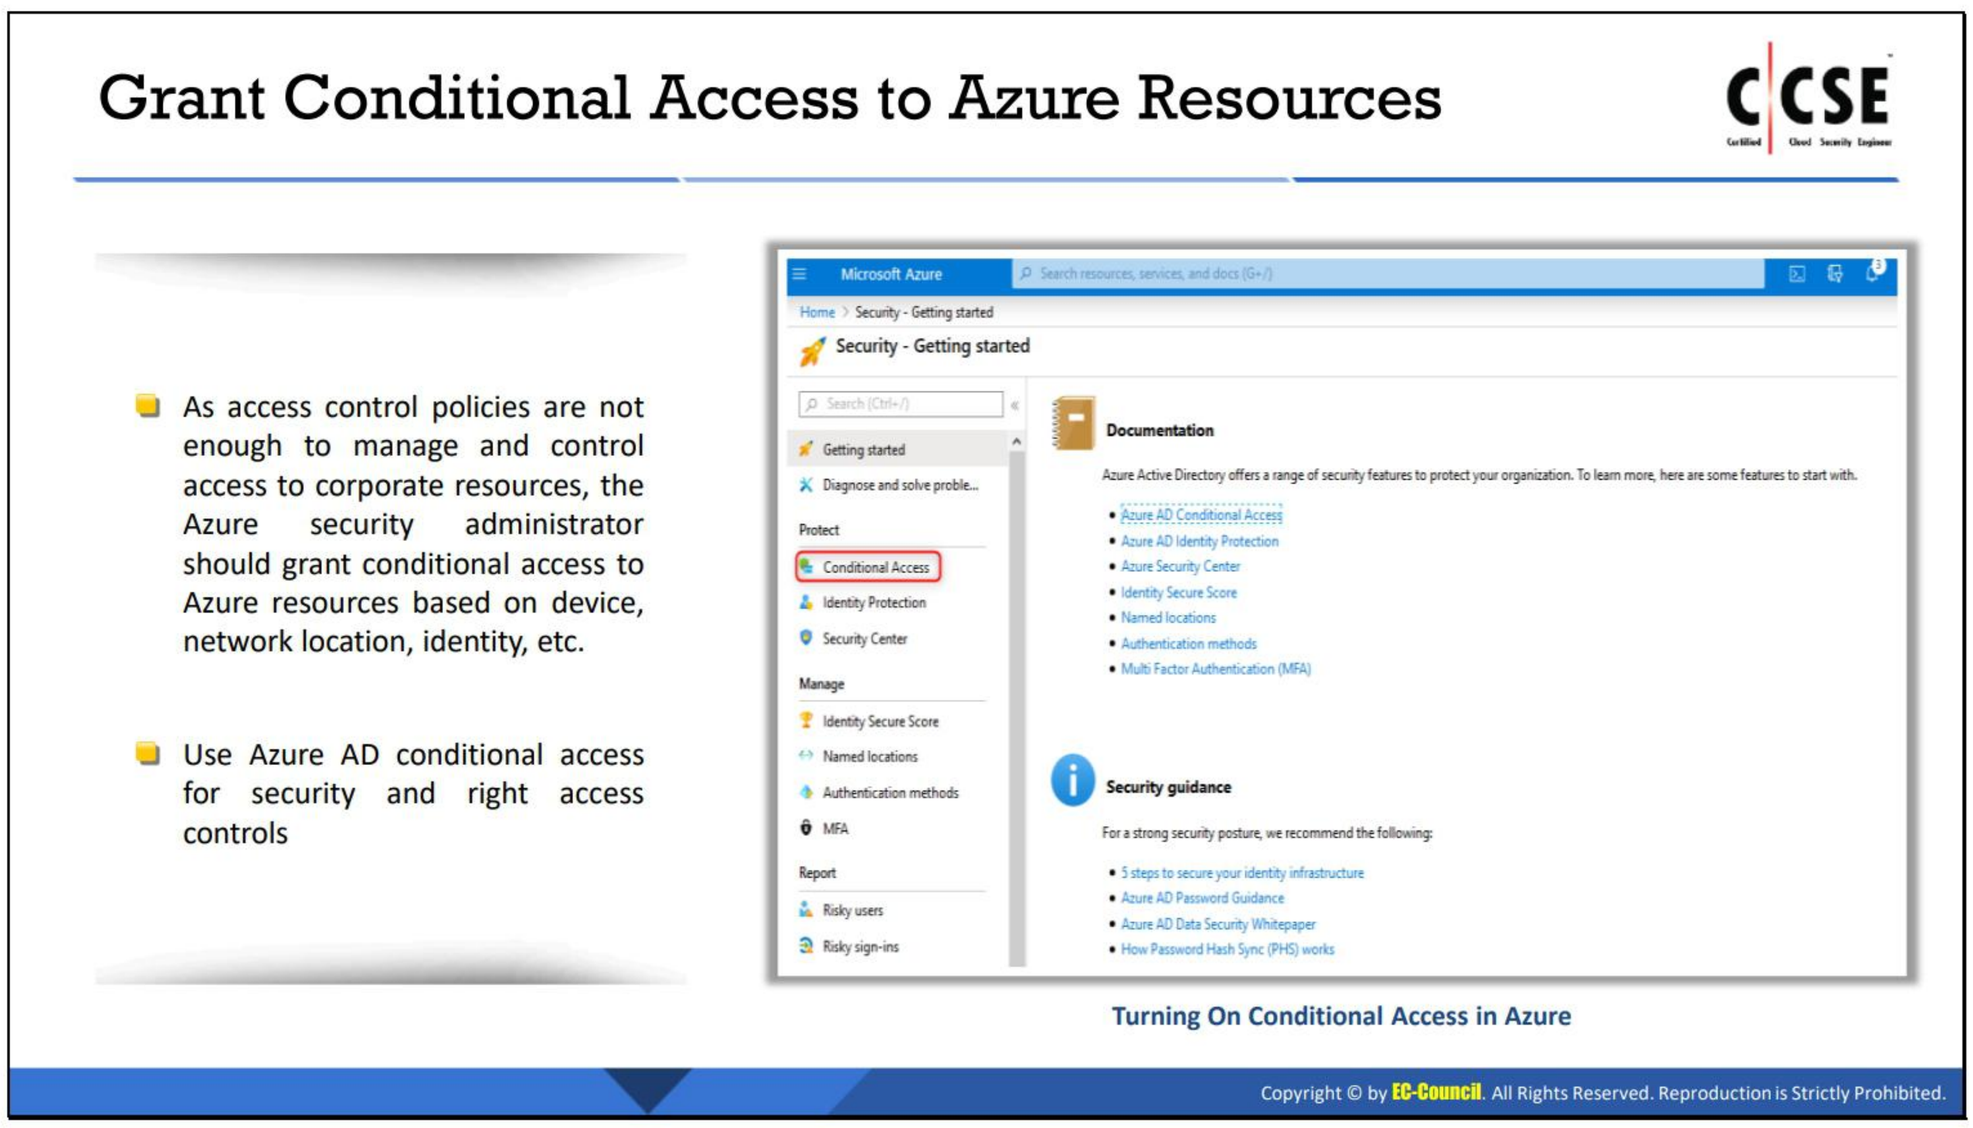


Figure: 2.78: Azure Load Balancer

In Azure, business continuity can be achieved by making the application architecture using a combination of AZs with Azure region pairs. The administrator can replicate applications and data synchronously using AZs inside an Azure region to ensure high availability and can replicate them across regions asynchronously for disaster recovery purposes.

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



Grant Conditional Access to Azure Resources

The Azure Active Directory (AD) tool, Conditional Access, helps to enforce the right access controls using the implementation of automated access control decisions when required. Enforce Conditional Access to Azure resources based on assurance, device, network location, identity, etc. Implementing Conditional Access is required as the access control policies are not enough to manage and control access to the organization's resources on Azure.

Source: https://www.microsoft.com

Steps to Create a Conditional Access Policy

The below steps help to create a Conditional Access policy for those who need to access the Microsoft Azure Management app to enforce MFA.

- **Create Conditional Access Policy**
  - Sign-in to the **Azure portal.**
  - Navigate to **Azure Active Directory** -> **Security** -> **Conditional Access**, select **New policy**.

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

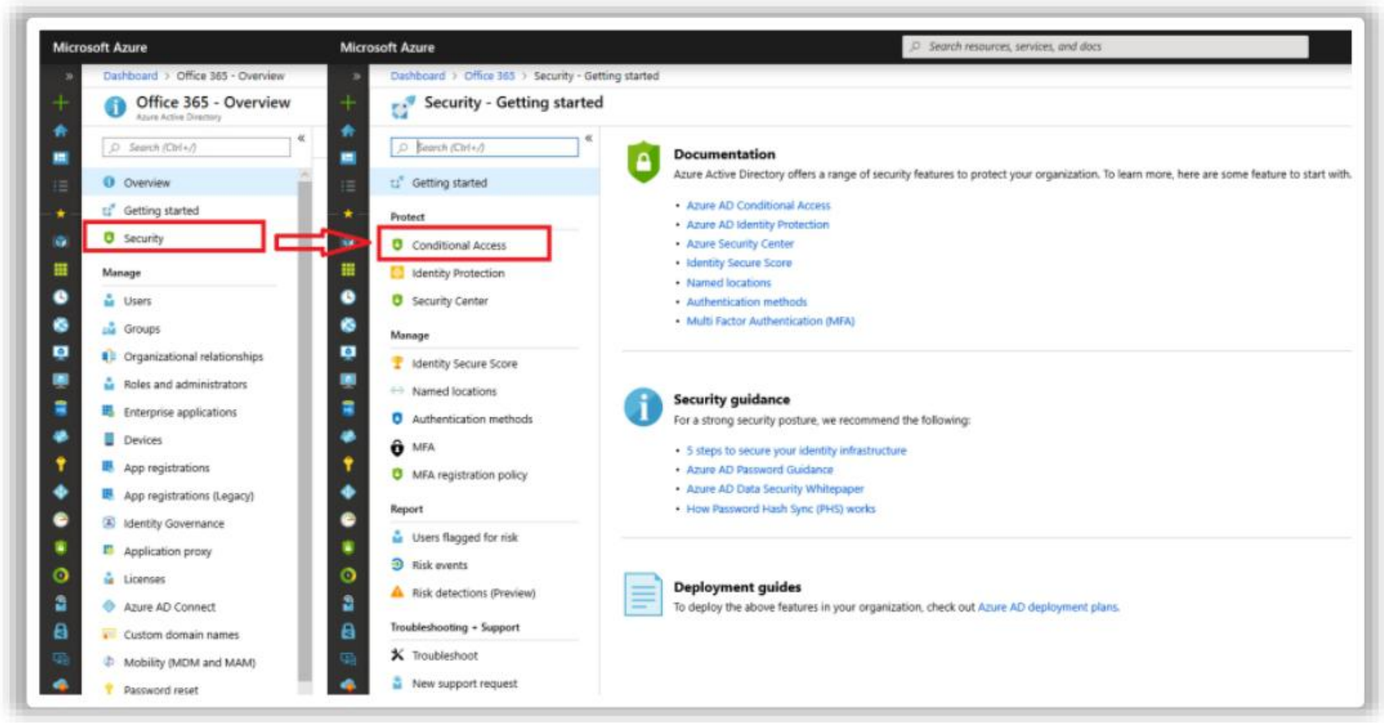


Figure: 2.79: Navigating to Azure Active Directory

- Enter the policy name. Create a meaningful standard for the names policies. For example: MFA Pilot.
- Under **Assignments**, select **Users and groups. Select the following:**
  - Under **Include**, select **All users**.
  - Under **Exclude**, check **Users and groups.** Select the organization's emergency access or break-glass accounts.
  - Select **Done**.

hide01.ir

- **Configure the Conditions for Multi-factor Authentication (MFA)**
  - Go to **Cloud apps or actions** -> **Include.** Select **Select apps and select Microsoft Azure Management**, and select **Select** then **Done**.
  - Under **Conditions** -> **Client apps (Preview)**, under **Select the client apps this policy will apply to,** leave all defaults selected and select **Done**.

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

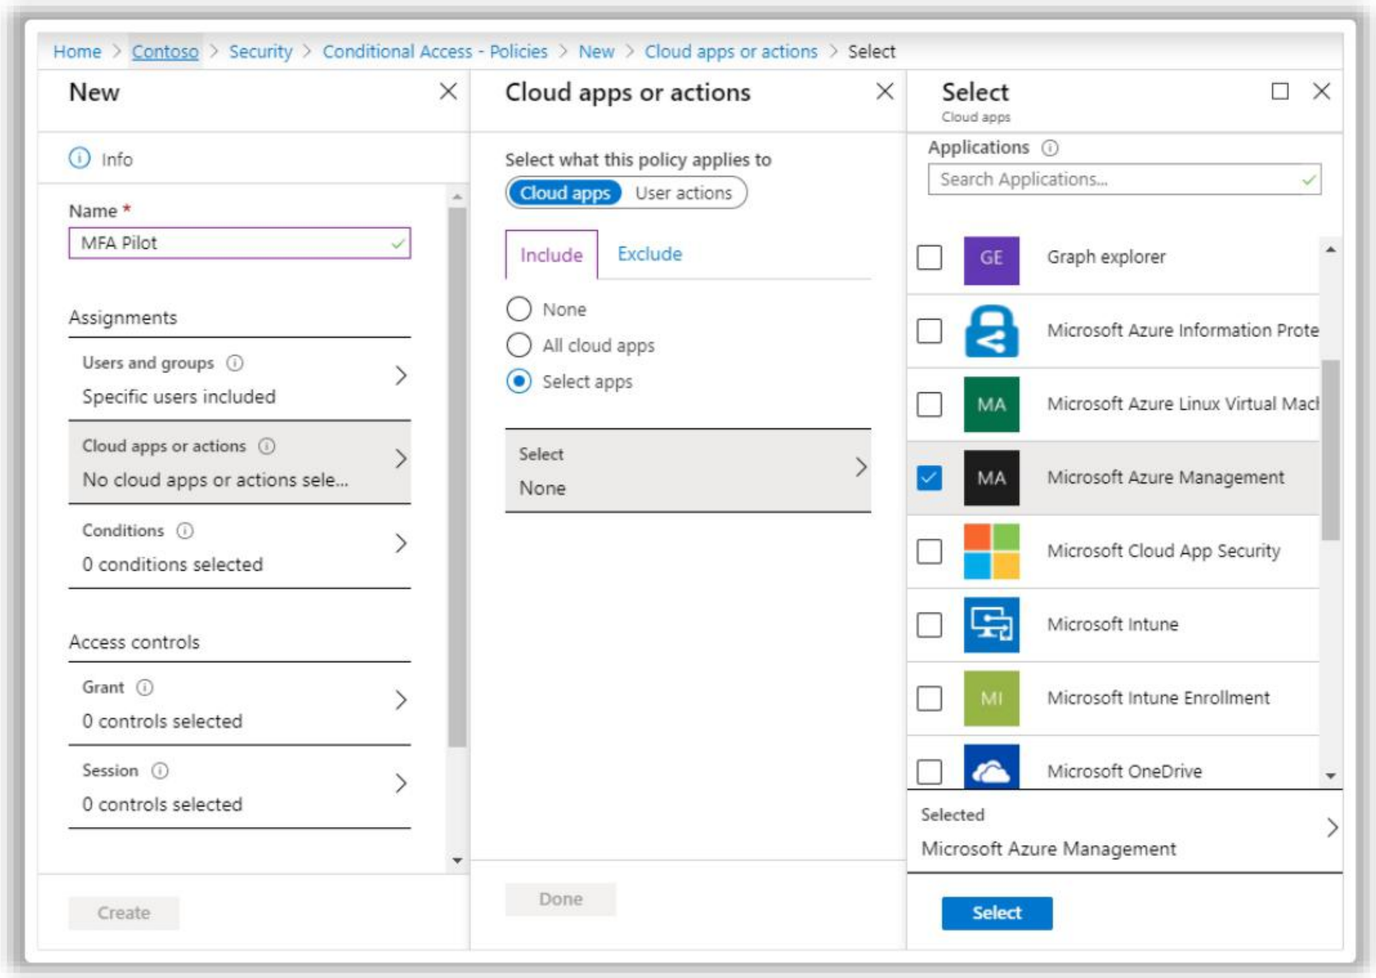
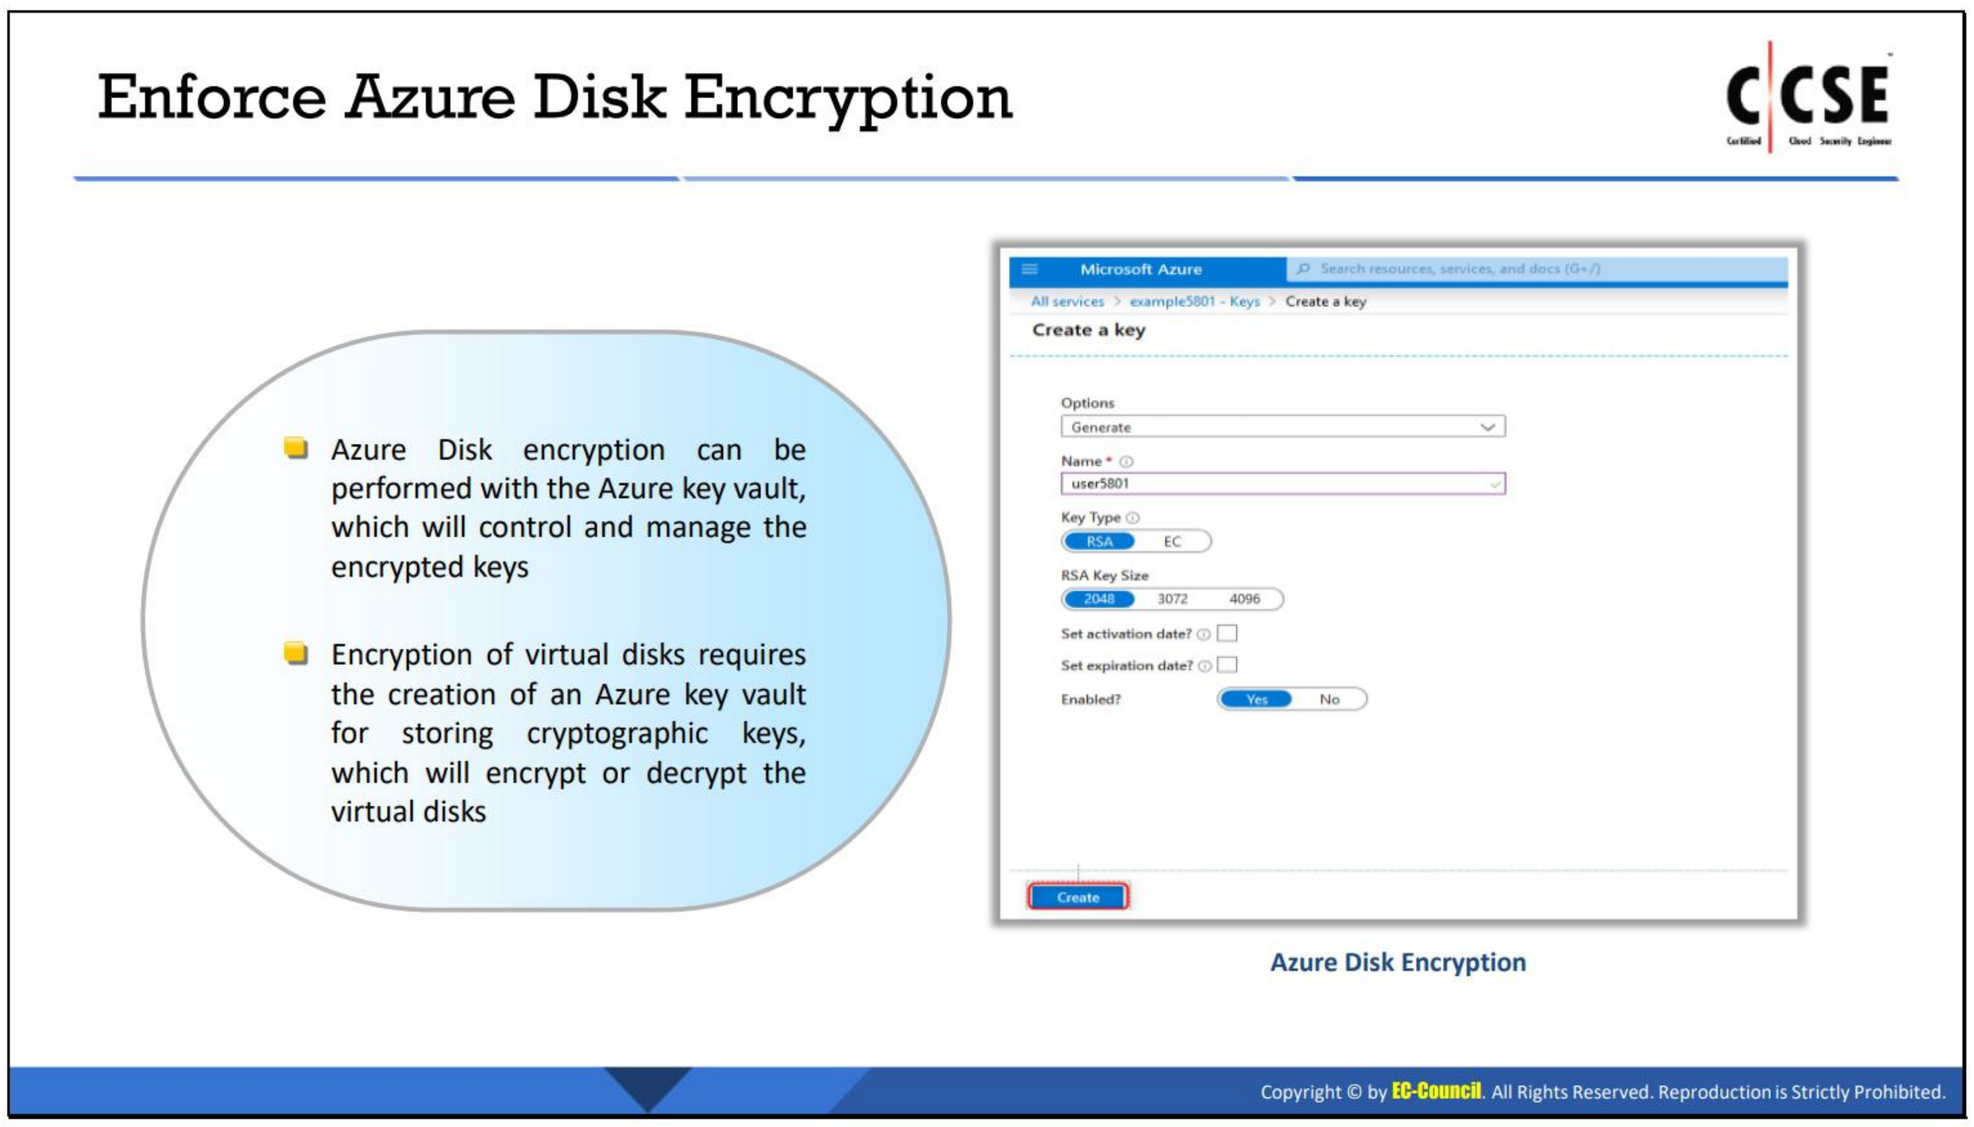


Figure: 2.80: Selecting Client Applications

- Under **Access controls -> Grant**, select **Grant access**, **Require multi-factor authentication**, and select **Select**.
- Confirm all settings and set **Enable policy** to **On**.
- Select **Create** to create to enable the policy.

hide01.ir

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



Enforce Azure Disk Encryption

The cloud administrator should use Azure Disk Encryption to secure data and meet security and compliance requirements of the organization. Use Azure Disk Encryption for all virtual machines (Windows and Linux) as well as VM scale sets. Azure Disk Encryption can be done with the Azure key vault, which will control and manage the encrypted keys. This key vault helps with the below.

- **Secrets Management**: The administrator can use the Azure key vault for strictly controlling access to passwords, tokens, API keys, and other secrets while storing them securely.
- **Key Management**: The administrator can use the key vault as a key management solution because it offers easier creation of encryption keys and control over them.
- **Certificate Management**: The administrator can utilize this key vault for easy provision, management and deployment of private and public TLS/SSL certificates that are used for Azure and other internally connected resources.

Encryption of virtual disks requires the creation of Azure key vault for storing cryptographic keys, which will encrypt or decrypt the virtual disks. Azure key vault is required due to below security reasons.

- **Store Secrets and Keys Securely**: Secrets and keys can be stored securely in Azure key vault. Any user or application that attempts to access the key vault needs to be properly authenticated and authorized before they access it. The key vault uses Azure AD for performing authentication and uses Azure RBAC for performing authorization.
- **Centralize Application Secrets**: With Azure key vault, the chances of secrets getting leaked accidentally are significantly reduced. Key vault helps to centralize the storage of

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

application secrets, which eliminates the need of storing security information in applications. Applications can use URIs to securely access the information from the vault. It also simplifies the administration of application secrets.

- **Monitor Access**: The administrator can use the key vault to monitor when and how the encryption keys and secrets are accessed after creating the vault. The administrator can also secure logs by restricting access to them and can delete logs which are not necessary.
- **Integration with Other Azure Services**: Azure key vault can integrate with event hubs, storage accounts, and log analytics.

Creating Azure Key Vault

- Log in to the Azure portal. Choose **Create a resource**, enter '**key vault**' and then select **Create**.
- On the new wizard, configure all basic settings required to begin utilizing the new key vault with disk encryption. Under **Basics**, specify the region, name, and resource.
- Choose **NEXT:Access Policy** and in that tab, select **Azure Disk Encryption for volume encryption** and click **review + create** to provision the key vault.

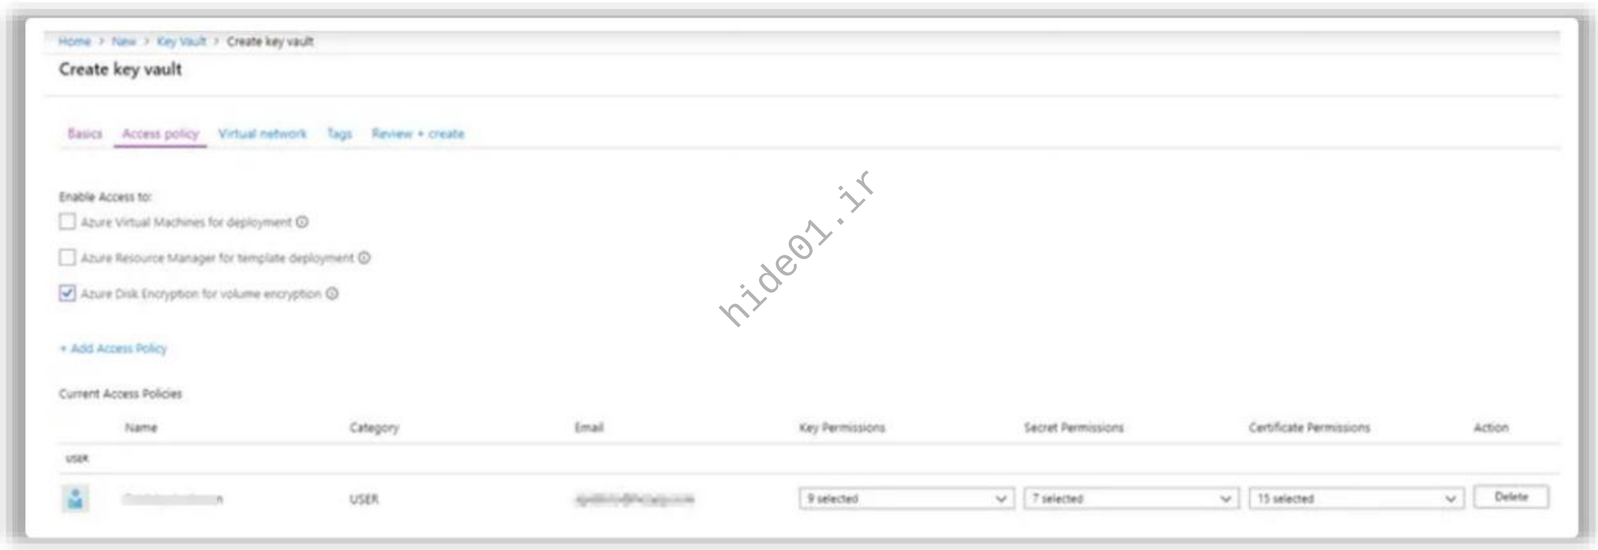


Figure: 2.81: Provisioning the Key Vault

- When a key vault is created without the above configuration, enable **Disk Encryption** by going to the **Key Vault properties** and then choosing **Access Policies**.

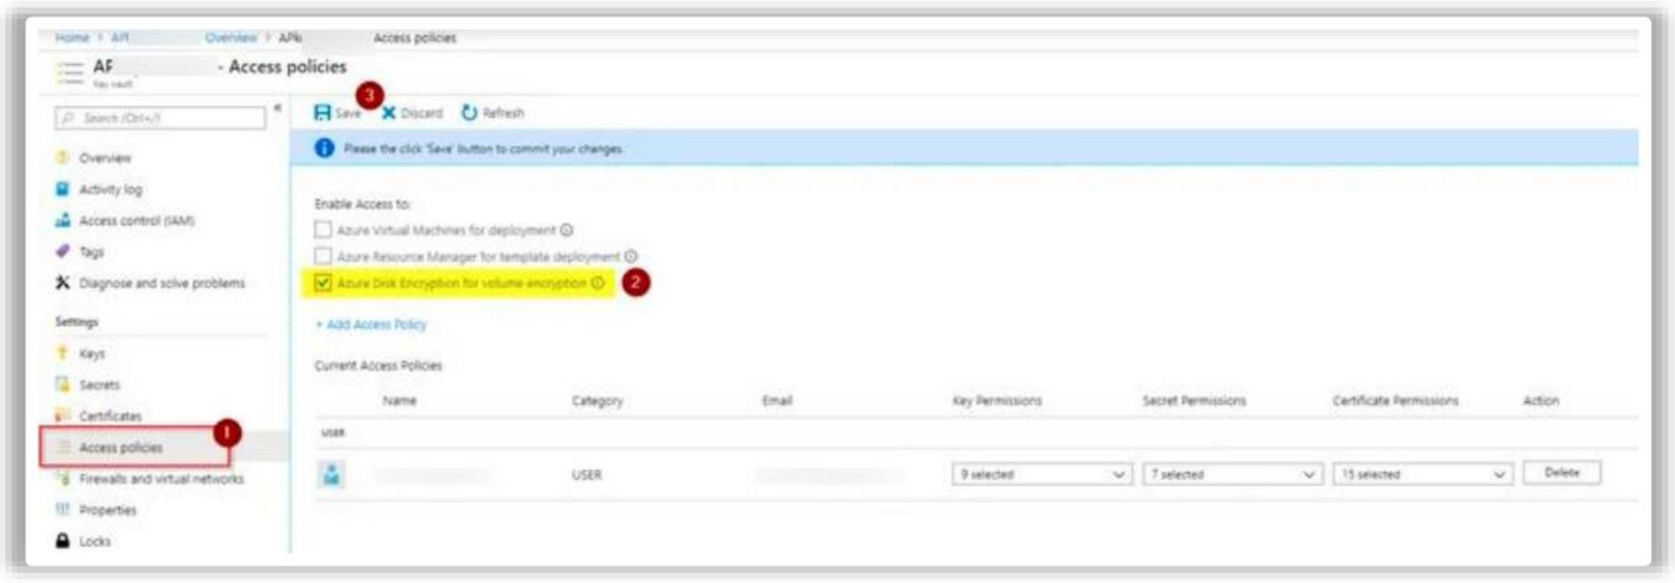


Figure: 2.82: Enabling Disk Encryption

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

Creating Encryption Keys

- After opening the newly created key vault, choose **Keys** and then select **Generate/Import**.
- On the Create a key page, choose **Generate**, enter a name for the key, specify the key type and key size, and set activation and deactivation data for the key.
- Choose **Create**.

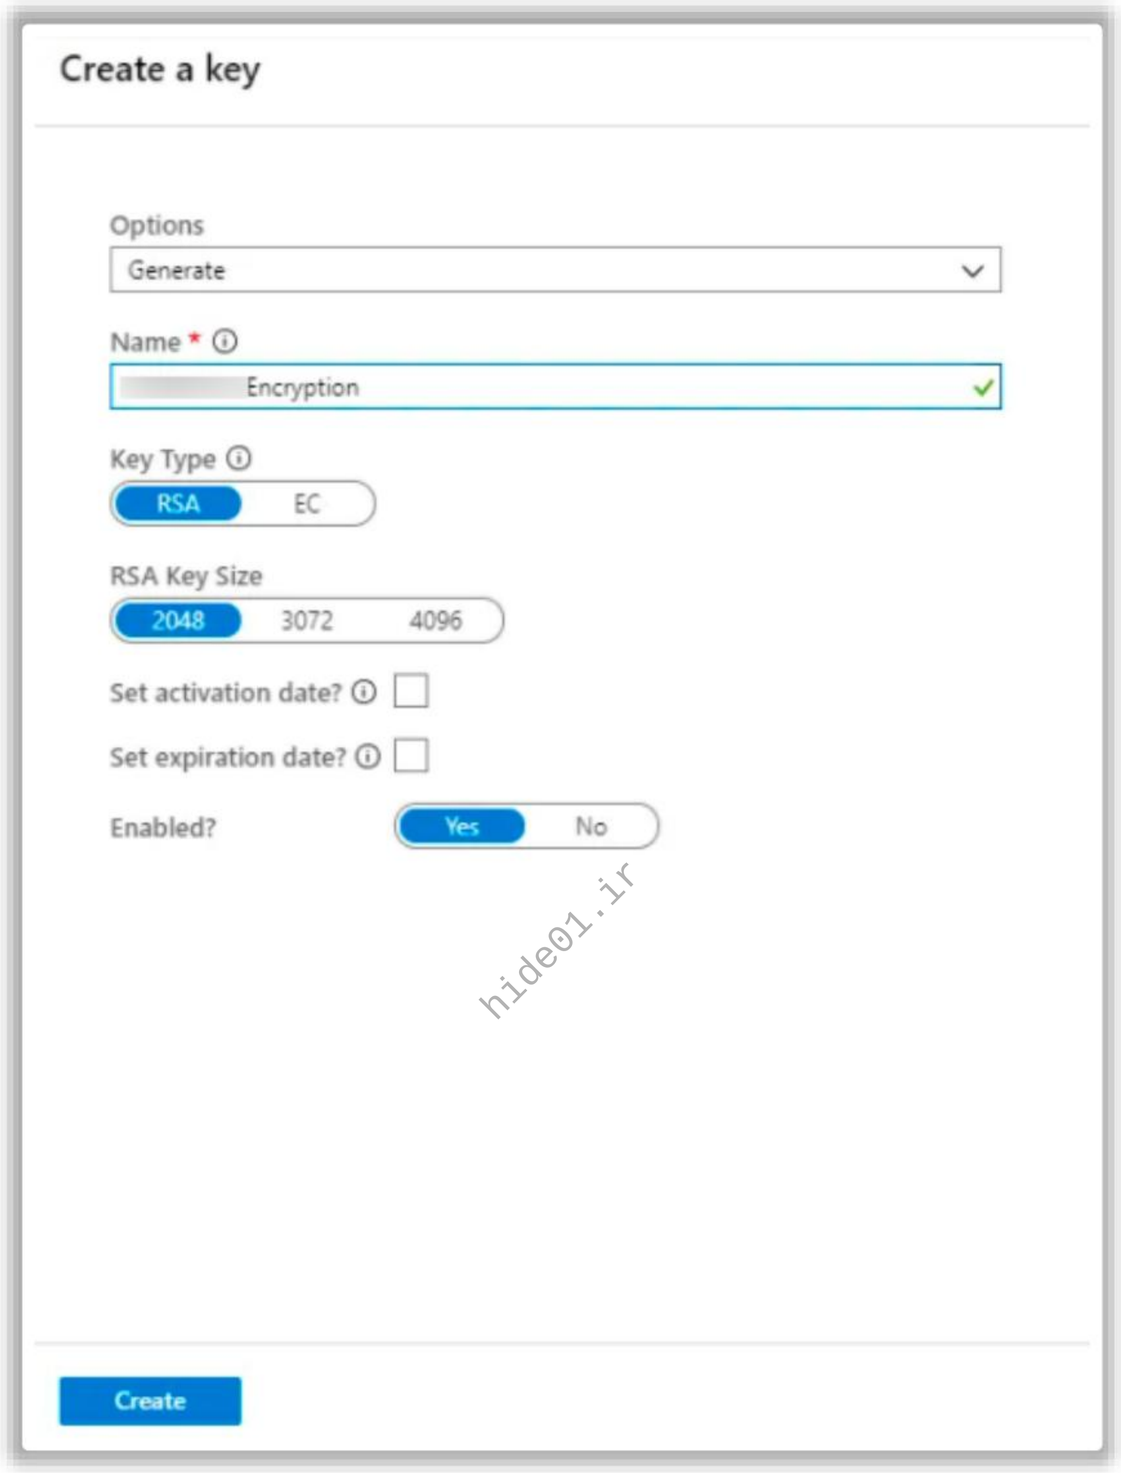


Figure: 2.83: Creating a Key

Encrypting a VM (Via Azure Portal)

- Go to the VM where the disk must be encrypted and select **Disks** to view the list of disks that are associated with the VM.
- Click on **Encryption** to begin the encryption process. Three options are present on the new window: **None** - no encryption; **OS disk** - encrypts just the OS; and **OS and data disks** - encrypts all disks in the VM.

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

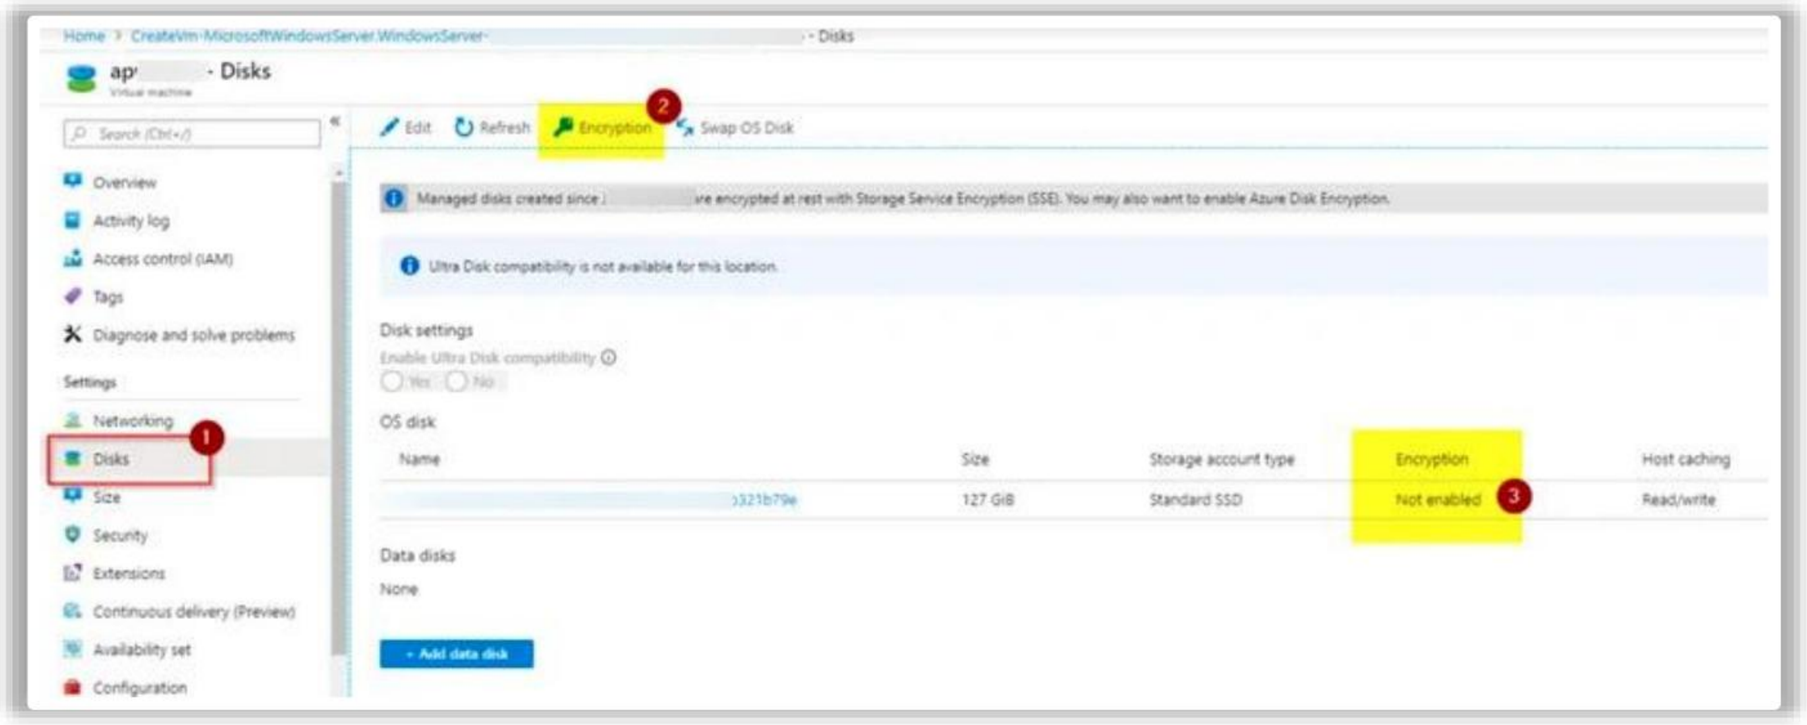


Figure: 2.84: Different Options for Encryption

- Select **OS disk** as an encryption option for example. Enter the requested information (keys, key vault, and version). The **Select a key vault and key for encryption link** can be used as it provides all keys, key vault, and version that are available in the existing environment. Choose the key that was just created; click **Select** and then choose **Save**.

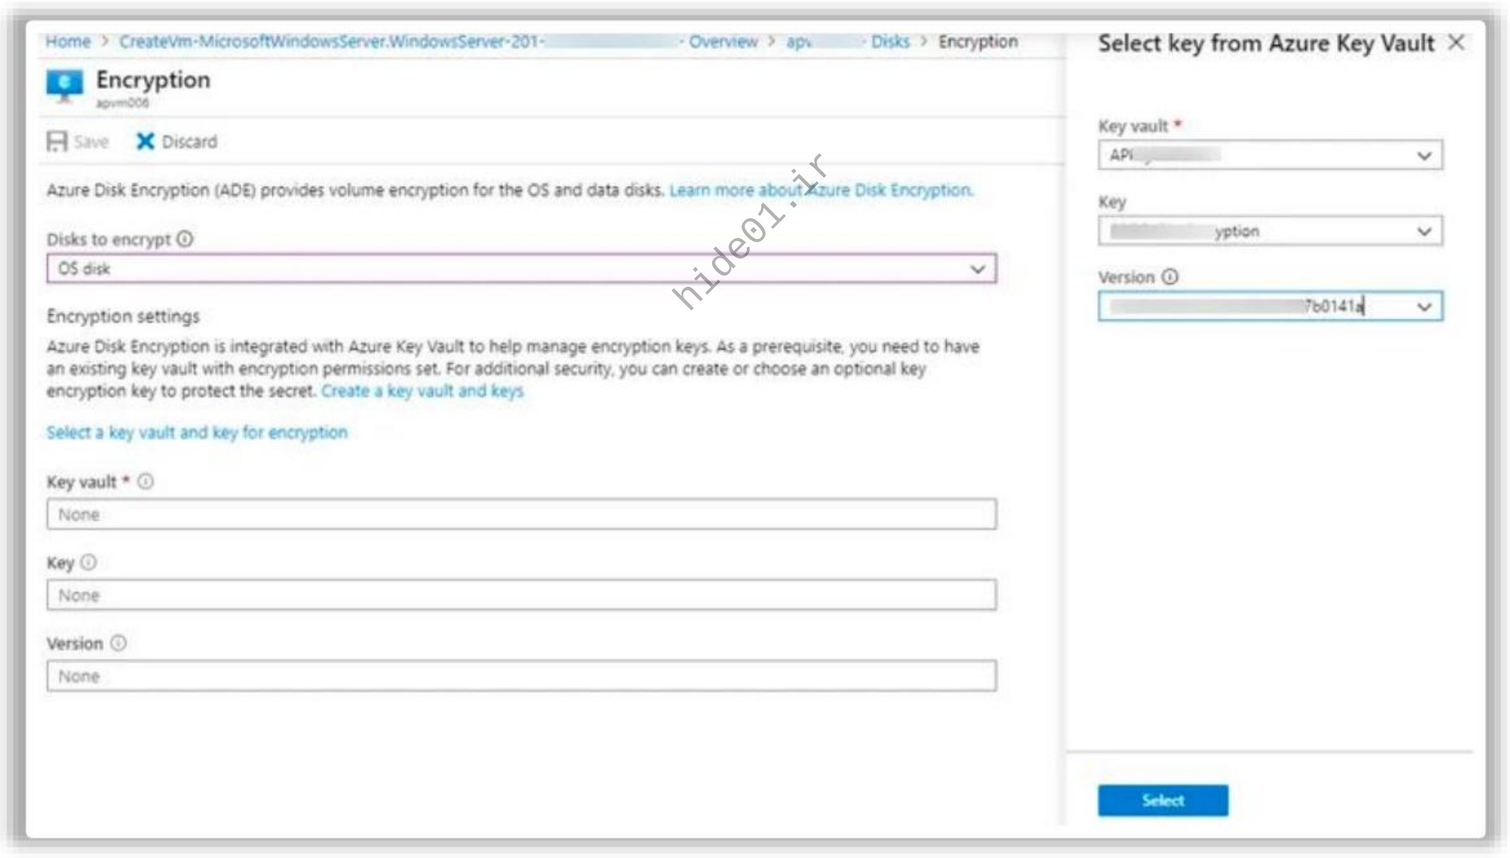


Figure: 2.85: Choosing the Created Key

- Click **Yes** on the confirmation dialog box. The VM will restart after this process. This will add AzureDiskEncryption to the VM.

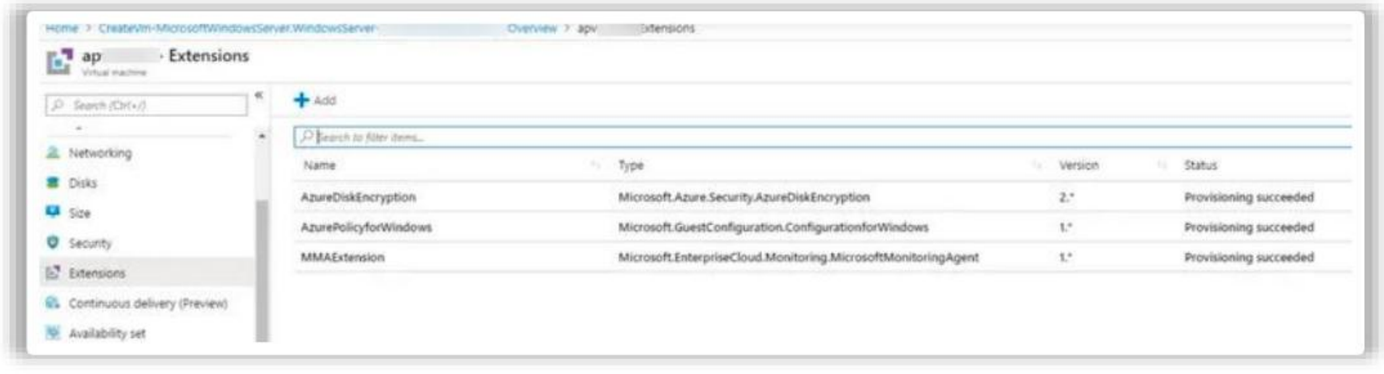
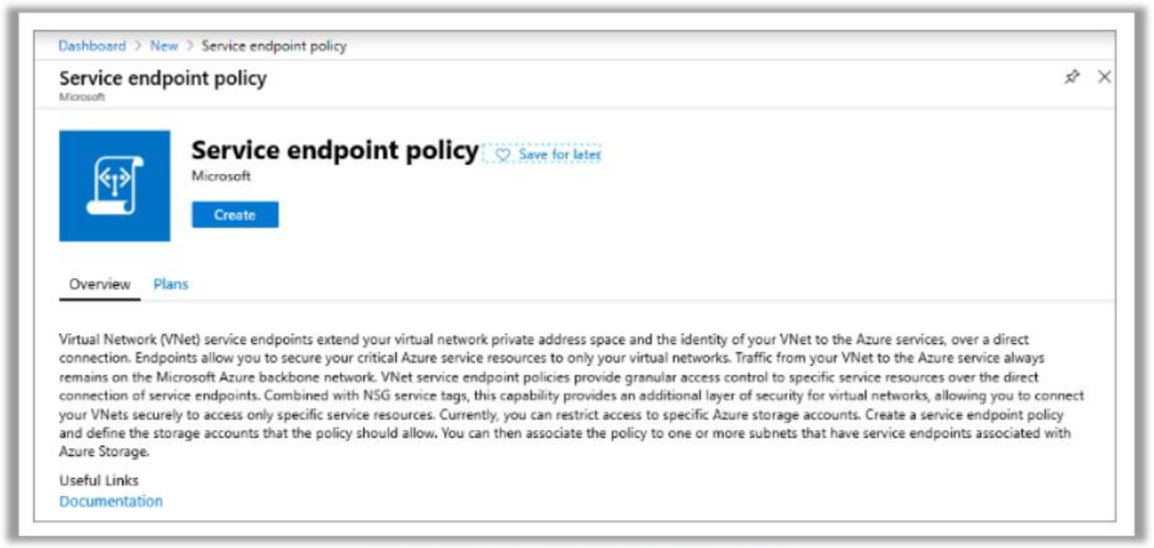


Figure: 2.86: Adding AzureDiskEncryption to the VM

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

Implement Virtual Network Service Endpoints

- The cloud security administrator can use the virtual network service endpoints to secure and directly connect the Azure service resources through an optimized route the in Azure backbone network
- In the virtual network, the service endpoints enable the private IP addresses to reach the azure services endpoint ruling out the need for public IP address on the virtual network



Implement Service Endpoint policy

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

Implement Virtual Network Service Endpoints

An Azure service endpoint helps Azure resources to connect and communicate directly and securely over the Azure backbone network (instead of connecting over the Internet). The cloud security administrator can use the virtual network service endpoints to secure and directly connect the Azure service resources through an optimized route in Azure backbone network.

In the virtual network, the service endpoints enable the private IP addresses to reach the Azure service's endpoint, ruling out the need of a public IP address on the virtual network. The administrator can restrict the service resource to make it accept traffic only from the subnet that is linked to the service endpoint. By using a network security group (NSG), the administrator can lock down or restrict the virtual network to only allow outbound traffic to the desired Azure service and block all other outbound traffic.

Suppose a VM present in a virtual network wants to communicate with a storage account in Azure. The administrator can club the storage account, a service endpoint, and NSG to ensure that traffic from a VM present in a private subnet does not use the internet to reach the storage account. All traffic to the storage account is blocked unless it is from that subnet and outbound traffic to the Internet from the subnet is restricted by NSG. In the absence of the service endpoint, it would be required to assign a public IP address to the VM, which directly exposes it to the Internet and risks.

With the help of service endpoint policies, the administrator can filter or block VNet traffic to Azure services. These policies can be used to have granular access control for VNet traffic to Azure services.

Service endpoints offer the below security benefits.

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

- **Enhanced Security**: With the use of service endpoints, no VM in the virtual network is assigned a public IP address so it is not easy for attackers to scan open ports of a VM for vulnerabilities and exploit them.
- **Optimized Routes**: There is no traffic delay or slow down as the service endpoint offers a direct route to the service resource from the virtual network over the Azure backbone network.
- **Simpler Network Architecture**: Using service endpoints does not require assigning a public IP address to the VM or configuring a gateway device or NAT because the traffic flows over the Azure backbone network from the virtual network resource to the service resource. Moreover, the administrator can enable service endpoints very easily.

Source: https://www.microsoft.com

Creating a Virtual Network (VNet) and Enabling a Service Endpoint (via Azure Portal)

- Sign-in to the Azure portal using https://portal.azure.com

Creating a VNet: Follow the steps to create a VNet.

- To create a virtual network, choose **+ Create a resource** on the upper left side of the portal.
- Choose **Networking** and then click on **Virtual networks**.
- Choose **+ Add** and fill the requested information as
  - **Subscription**: Choose your subscription
  - **Resource group**: Choose **Create new** and type **myResourceGroup**
  - **Name**: Type **myVirtualNetwork**
  - **Region**: Choose **East US**

h1de01.ir

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

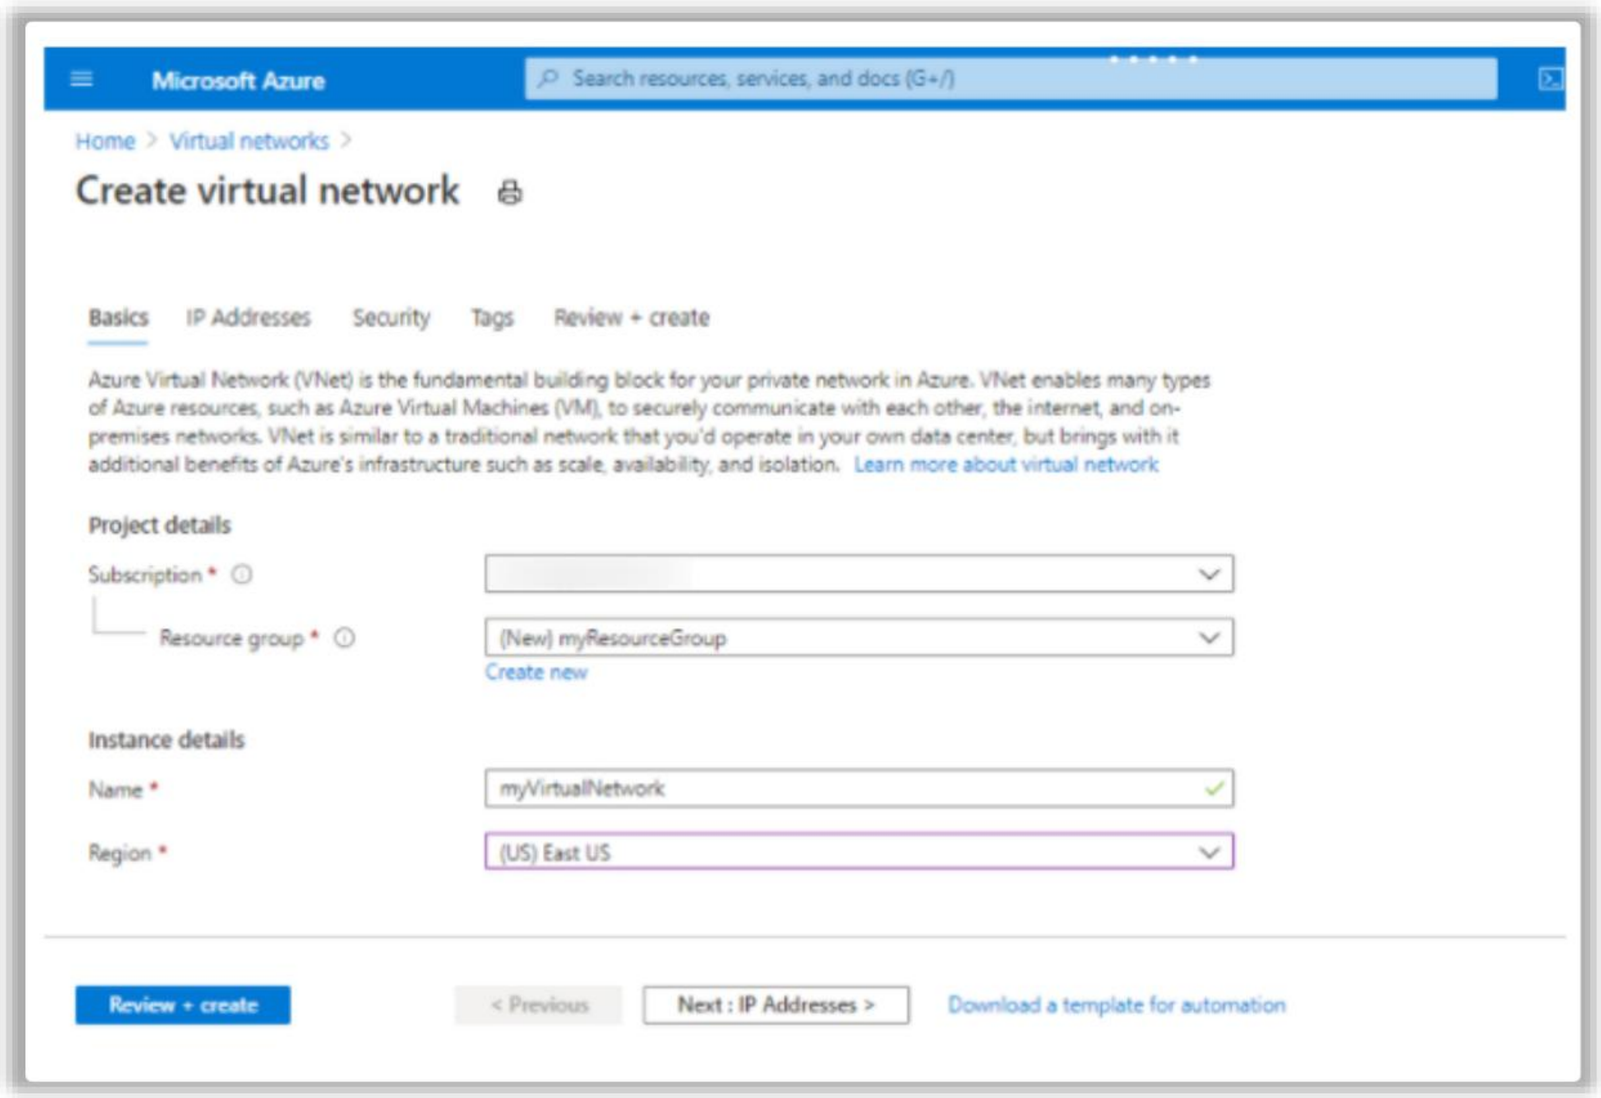


Figure: 2.87: Filling Details for Creating a Virtual Network

- Choose **Next:IP addresses>** and leave all settings as default but keep the name as **Public** under the subnet name.
- Choose **Next:Security>** and then select **Disable** in all settings.
- Then, click **Review and create**. After the validation check, choose **Create**.
- Click on **Go to resource** after the deployment

Enabling The Service Endpoint: Create a subnet and turn on the service endpoint by following these steps.

- Enter **myVirtualNetwork** in the search box on the portal and choose the network from the list.
- Choose **Subnets** from the **Settings** menu on the left side. Then, choose **+ Subnet**.

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

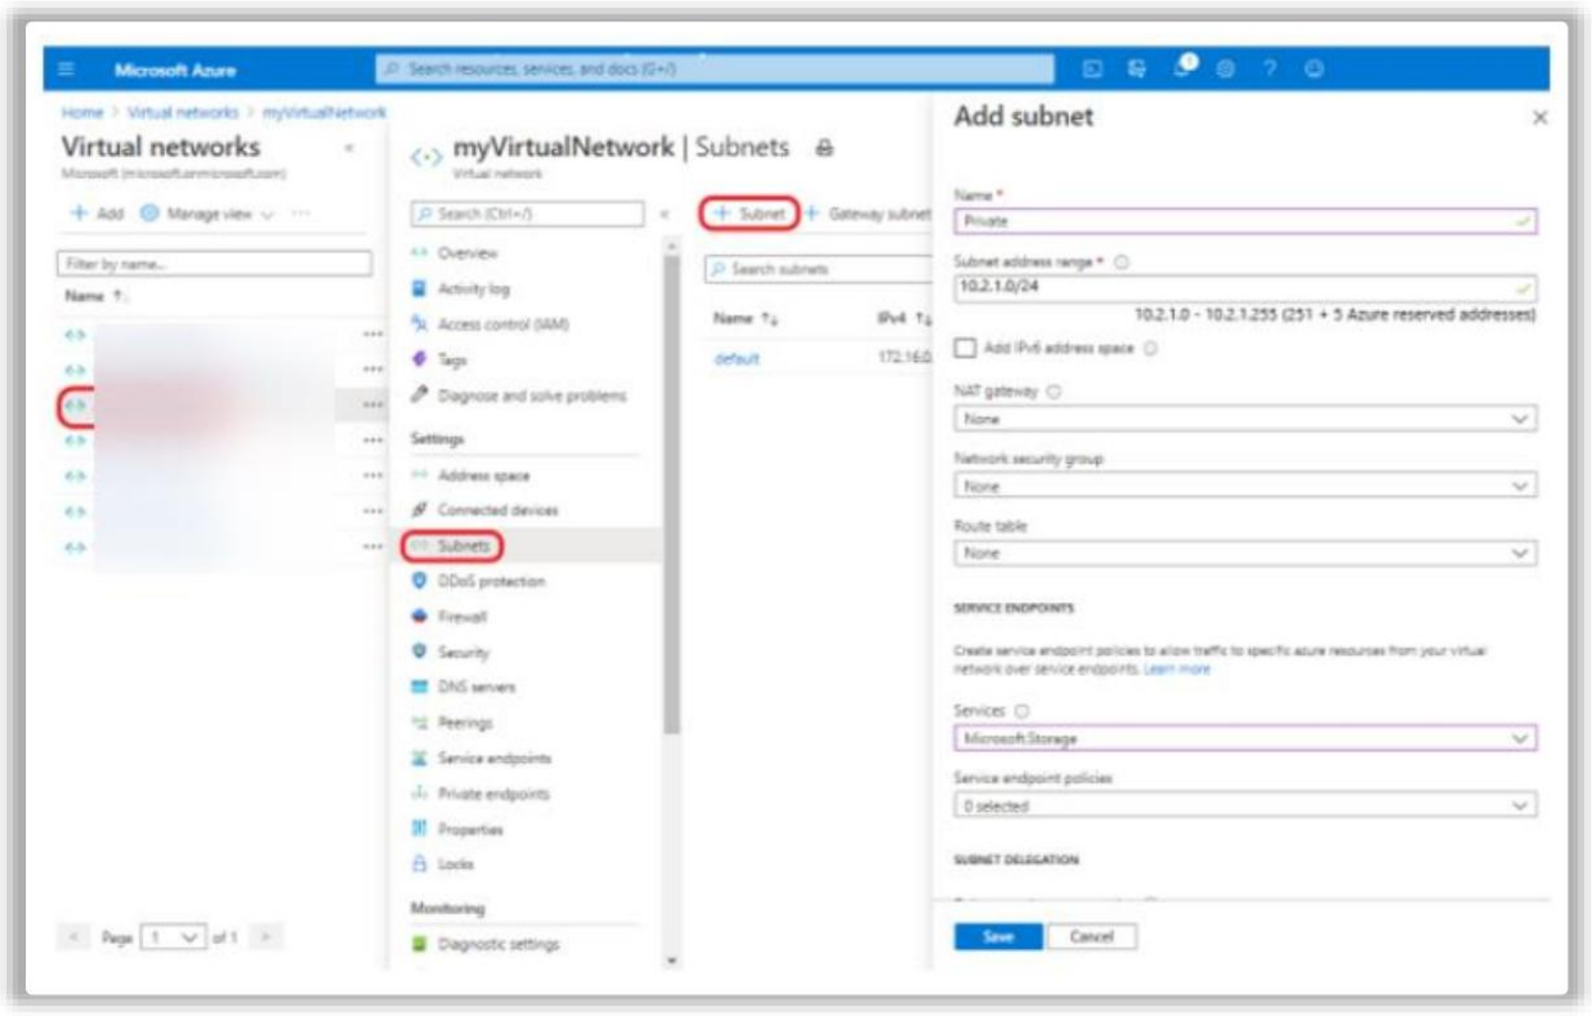


Figure: 2.88: Adding a Subnet

- Enter the given information under **Add subnet** and then click **OK.**
    - **Name**: Private
    - **Address**: Keep it default
    - **Service endpoints**: Microsoft.Storage
    - **Service endpoint policies**: default (0)

Creating an Endpoint Policy: Follow the steps for creating an endpoint policy.

- Sign-in to the Azure portal using https://portal.azure.com
- Choose **+ Create a resource** on the upper left side of the portal.

hade01.ir

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

- Type '**service endpoint policy**' in the search pane and choose **Service endpoint policy**. Then, select **Create**.

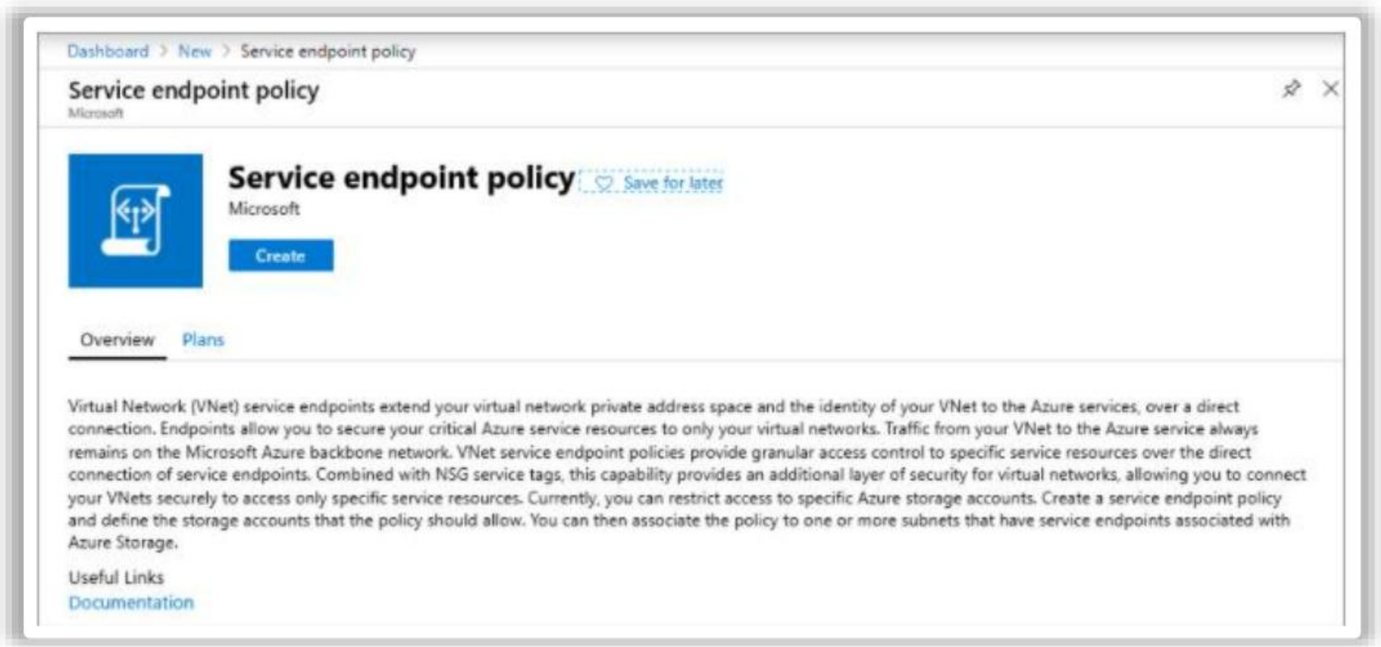


Figure: 2.89: Choosing Service Endpoint Policy Option

- Under **Basics**, enter the following information
  - **Subscription**: Choose a policy subscription
  - **Resource group**: Choose **Create new** and type myResourceGroup
  - **Name**: myEndpointPolicy
  - **Region**: Central US

hide01.ir

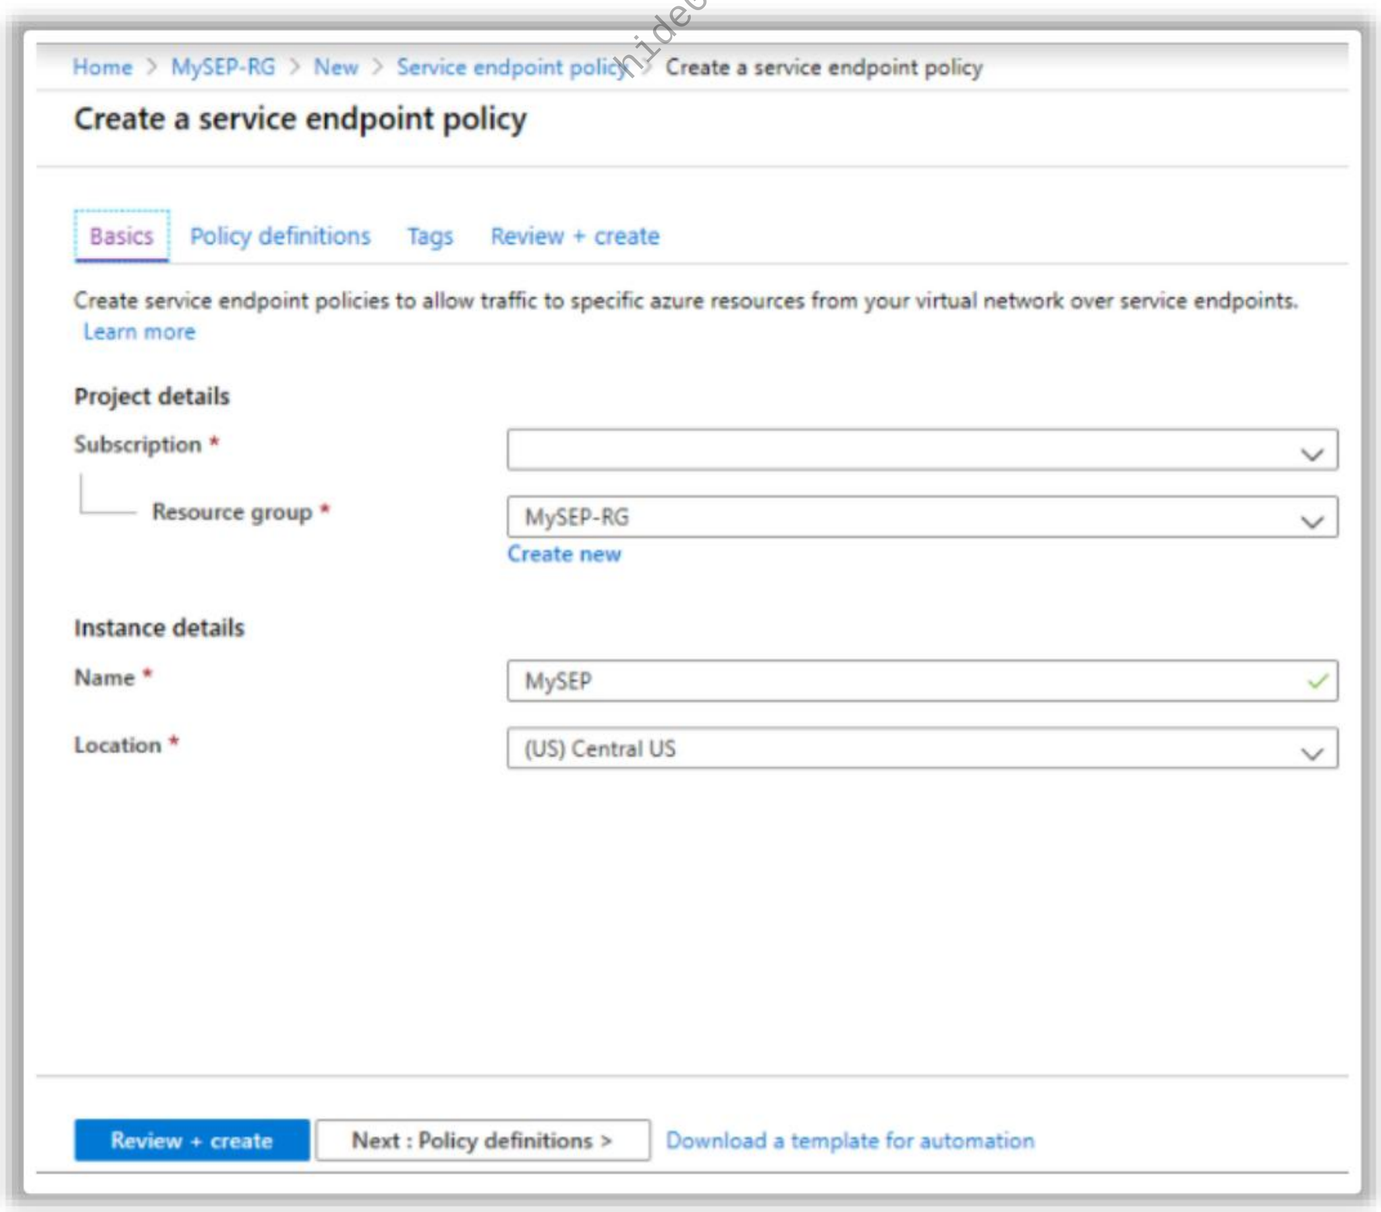


Figure: 2.90: Creating a Service Endpoint Policy

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

- Under **Resources**, choose **+ Add** and in the **Add a resource** pane, fill the following information:
    - **Service**: Microsoft.Storage
    - **Scope**: Choose an appropriate option.
    - **Subscription**: Choose the subscription for the storage account.
    - **Resource group**: Choose the resource group if the scope is set as 'Single account' or 'All accounts in resource group'.
    - **Resource**: Choose the Storage resource under the selected **Resource Group** or **Subscription**.
- To finish adding the resource, choose **Add**.

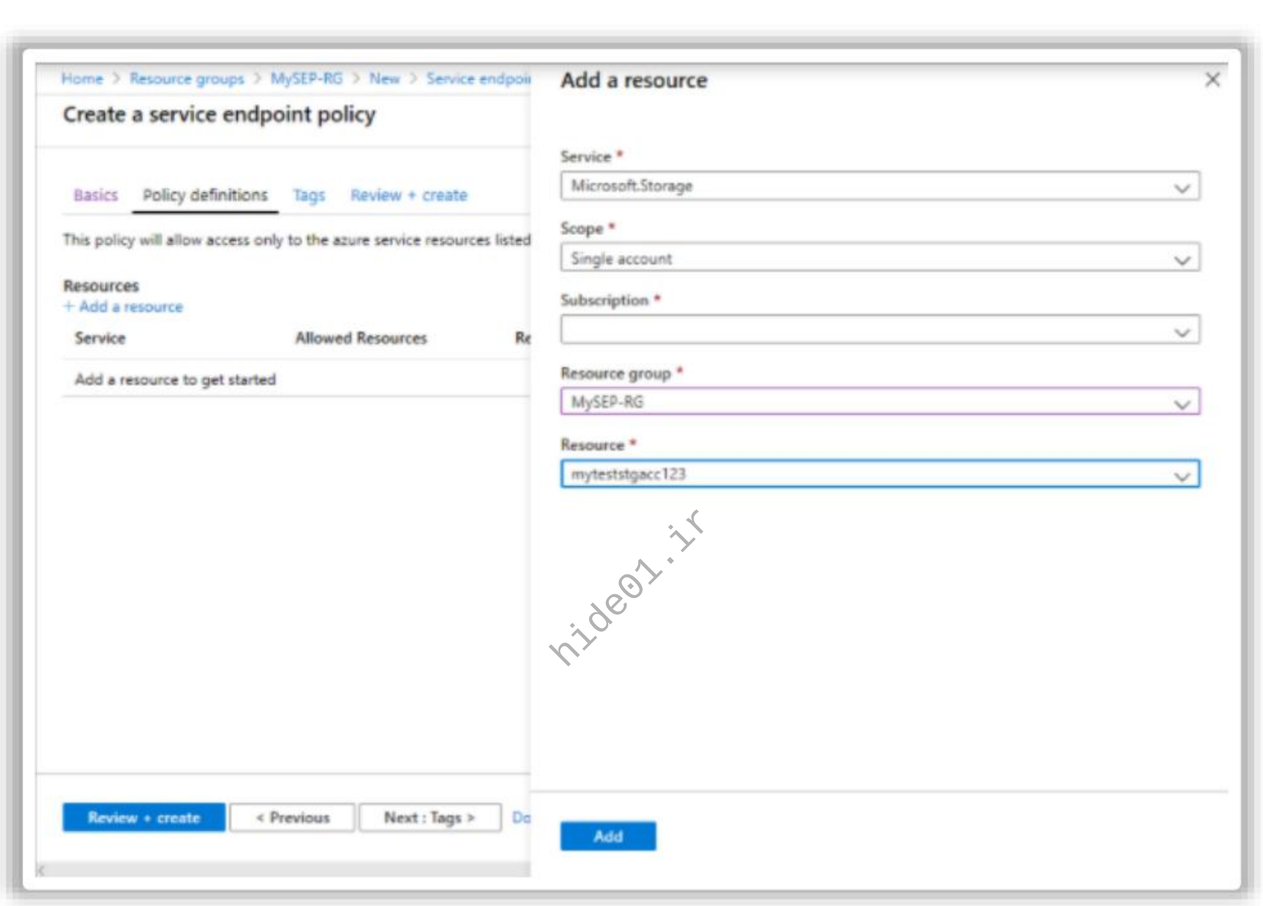


Figure: 2.91: Adding a Resource

- Repeat the above-mentioned steps to add more resources.
- Enter information in **Tags** (key and value fields). This step is optional.
- Choose **Review + Create**. Select **Create** after validating the information. Use the **Previous** button to edit previous information.

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

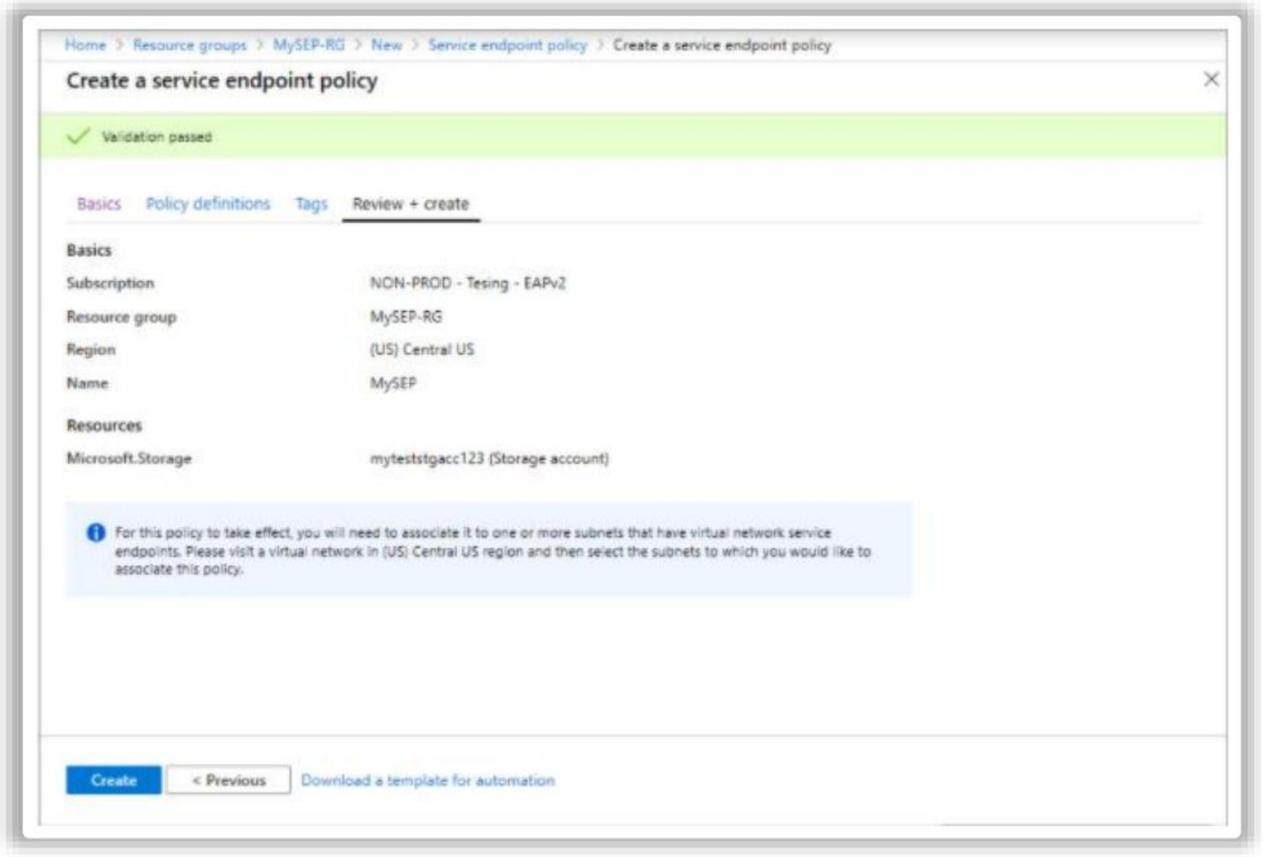
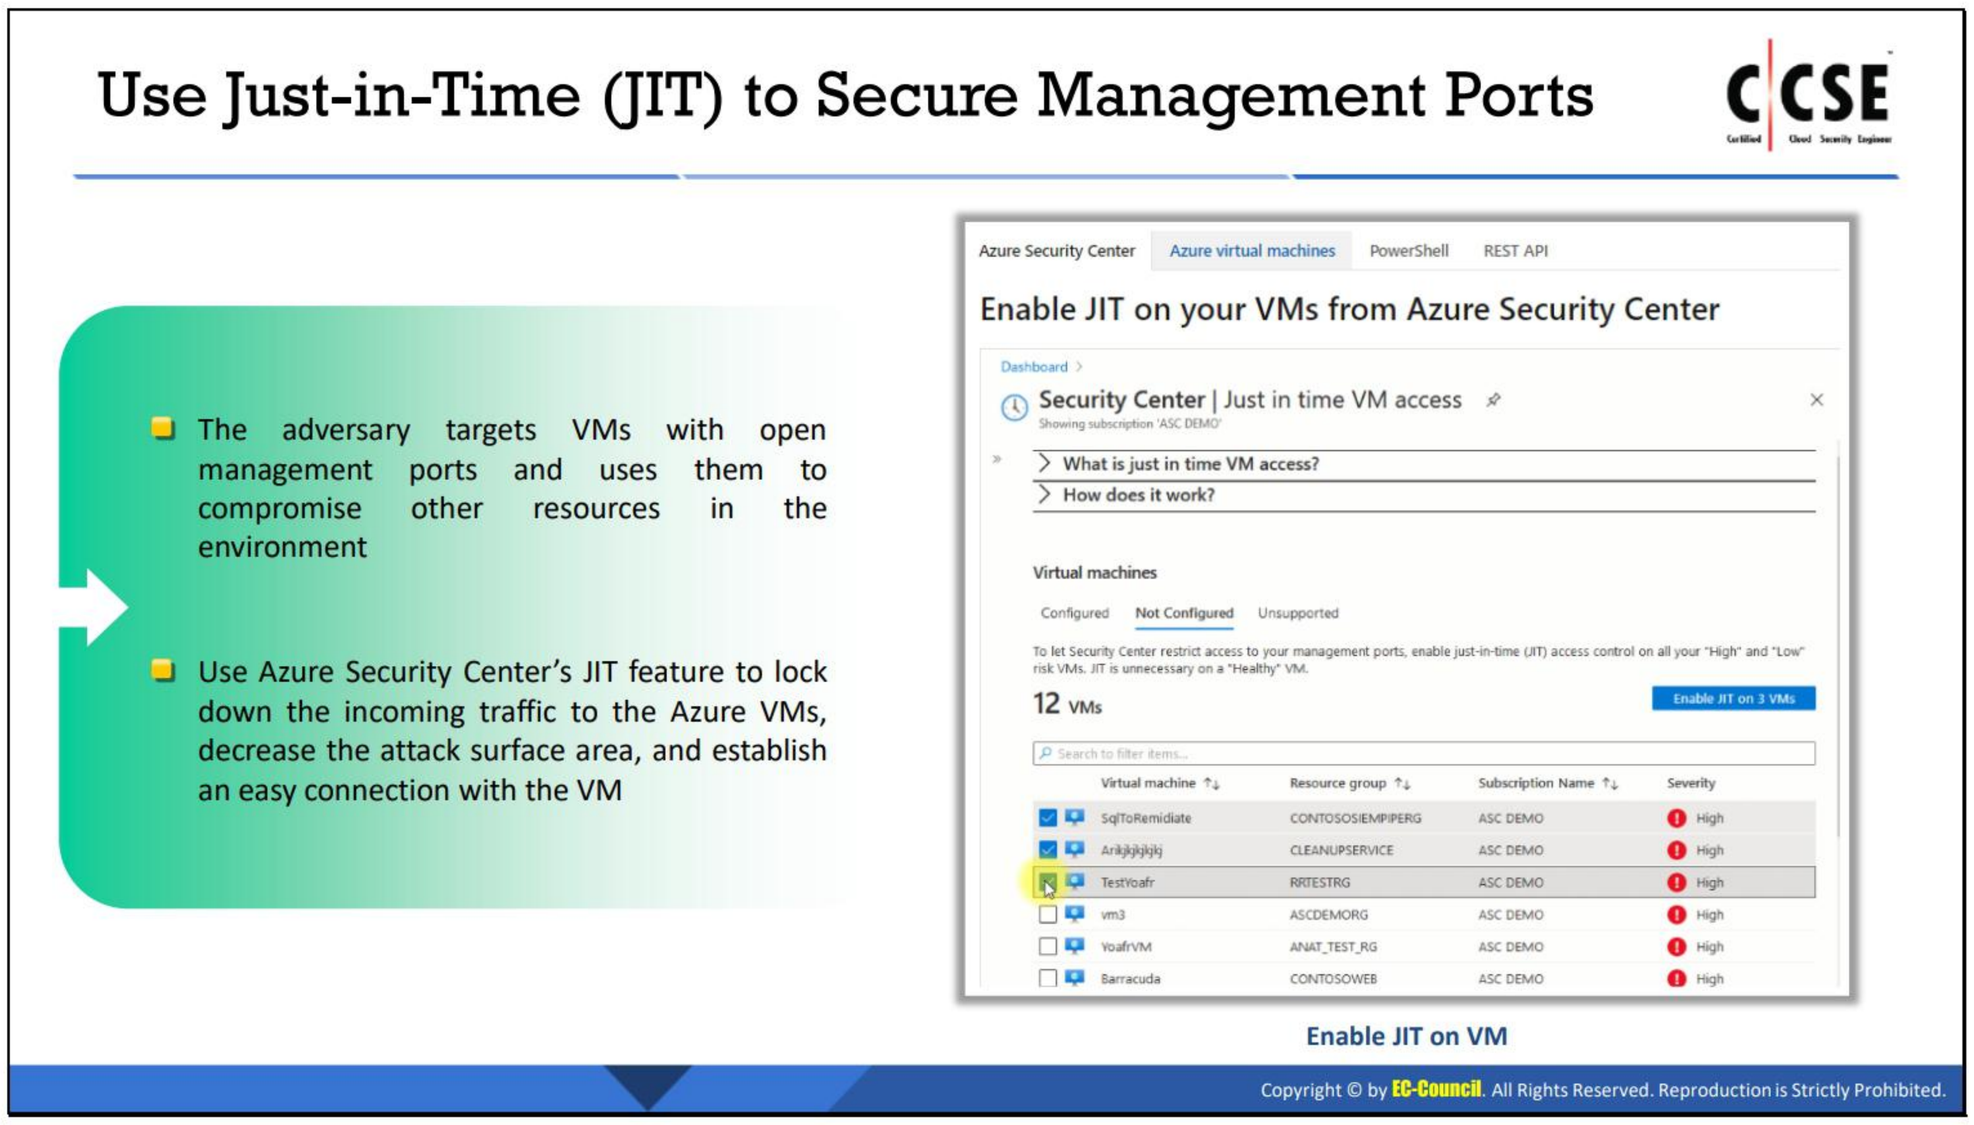


Figure: 2.92: Validation of Information

hide01.ir

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



Use Just-in-Time (JIT) to Secure Management Ports

Threat actors search for Azure VMs with open management ports RDP/SSH and try to attack the VMs. If a VM is compromised, it can be used as an entry point to attack other Azure resources within the organization's environment. Ensure having fewer open pots (especially management ports) and so reducing the attack surface. However, it is not practical to close the management ports as legitimate users also use those ports. In this case, use Azure Security Center's Just-in-Time (JIT) to lock down the inbound traffic to VMs, provide easy access to VMs when required and so reduce exposure to attacks.

Source: https://www.microsoft.com

Steps to Enable JIT on VMs From Azure Security Center

Use custom options or default options for one more VMs to enable and configure JIT VM access using the Security Center.

- Open the **Azure Defender** dashboard.
- Under the **Advanced protection** area, select **Just-in-time VM access**.

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

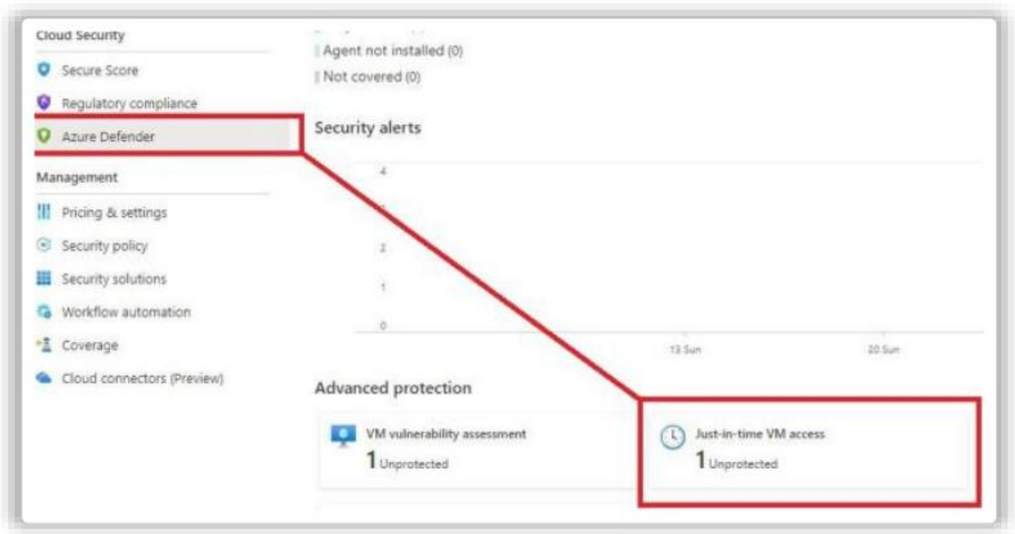


Figure: 2.93: Selecting Just-in-time VM Access

- The **Just-in-time VM access** page appears with the VMs grouped into the below tabs.
  - **Configured:** Configured VMs that support just-in-time VM access.
  - **Not Configured:** VMs without JIT enabled. However, they can support JIT.
  - **Unsupported**: VMs without JIT enabled. They do not support the JIT feature.

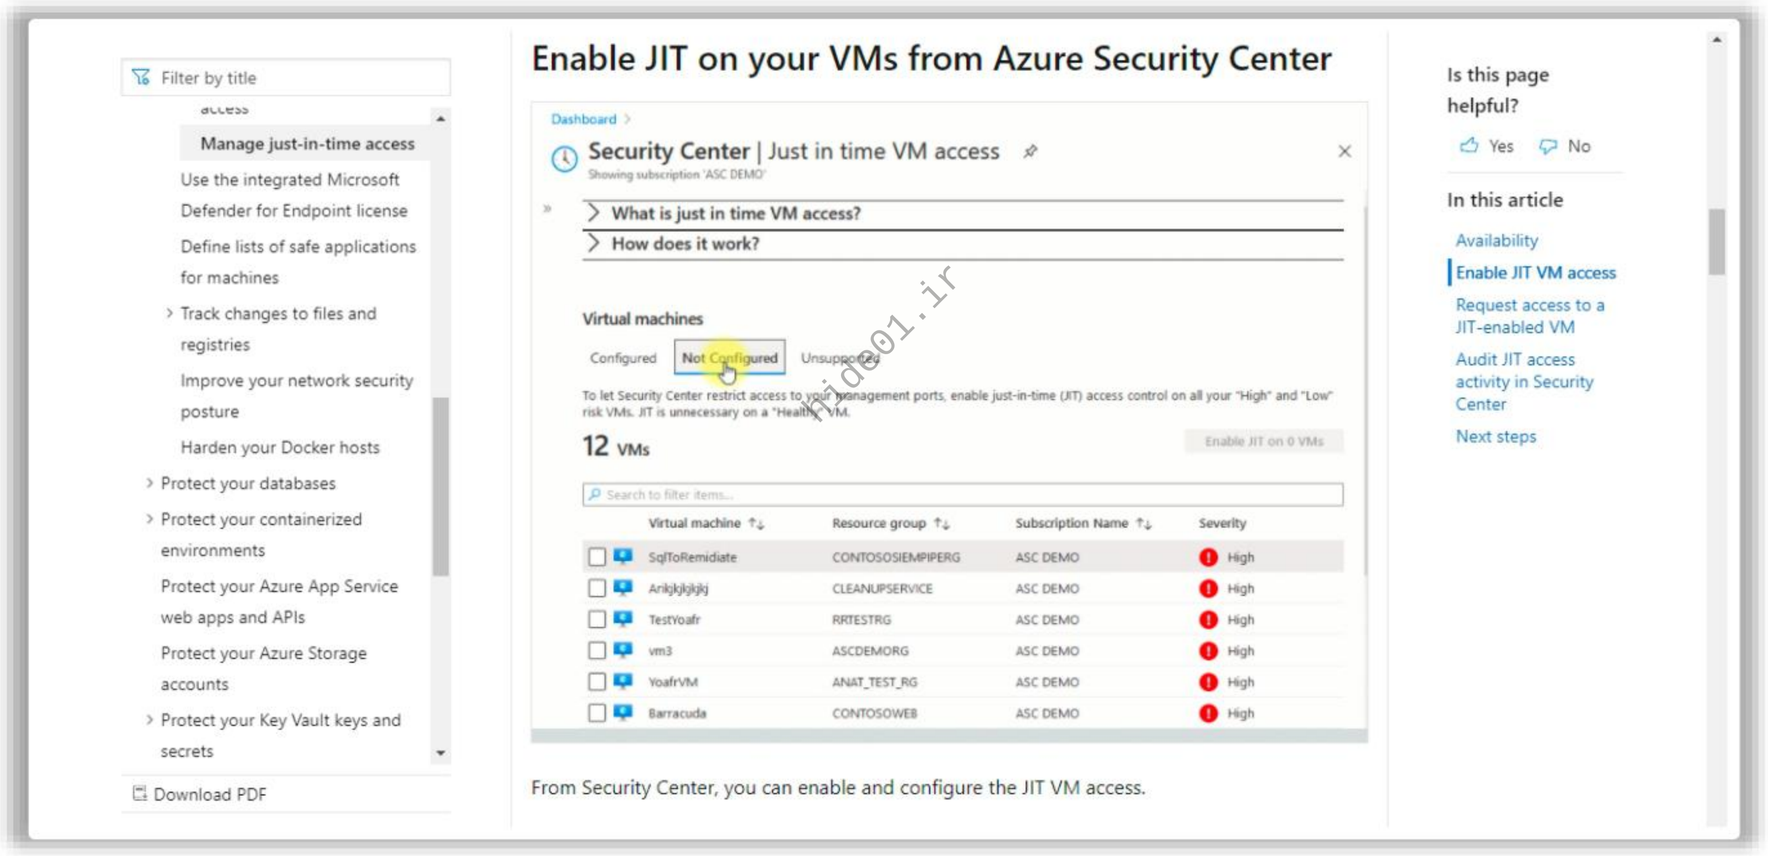


Figure: 2.94: VMs Grouped as Different Options

- Under **Not Configured** tab, check the VMs to protect with JIT and select **Enable JIT on VMs**.

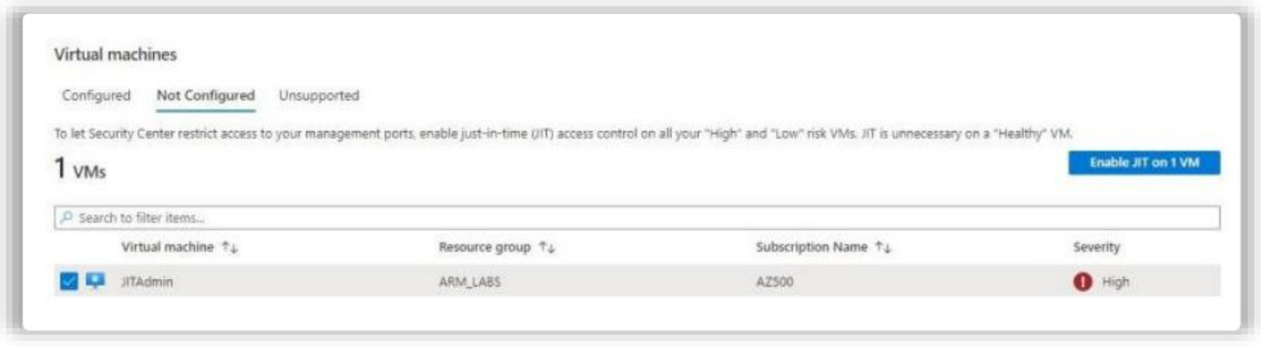


Figure: 2.95: Enabling JIT on VMs

- The JIT VM access configuration page appears. The page lists the ports that Security Center recommends protecting:

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

- 22: SSH
- 3389: RDP
- 5985: WinRM
- 5986: WinRM

Select **Save** to accept the default settings.

- To customize the JIT options:
  - Click **Add** to add custom ports.
  - Modify any default port by selecting it from the list.
    - The **Add port configuration** pane appears. Provide the required for **Protocol, Allowed source IPs, Max request time and IP address.**
    - Select **OK**.

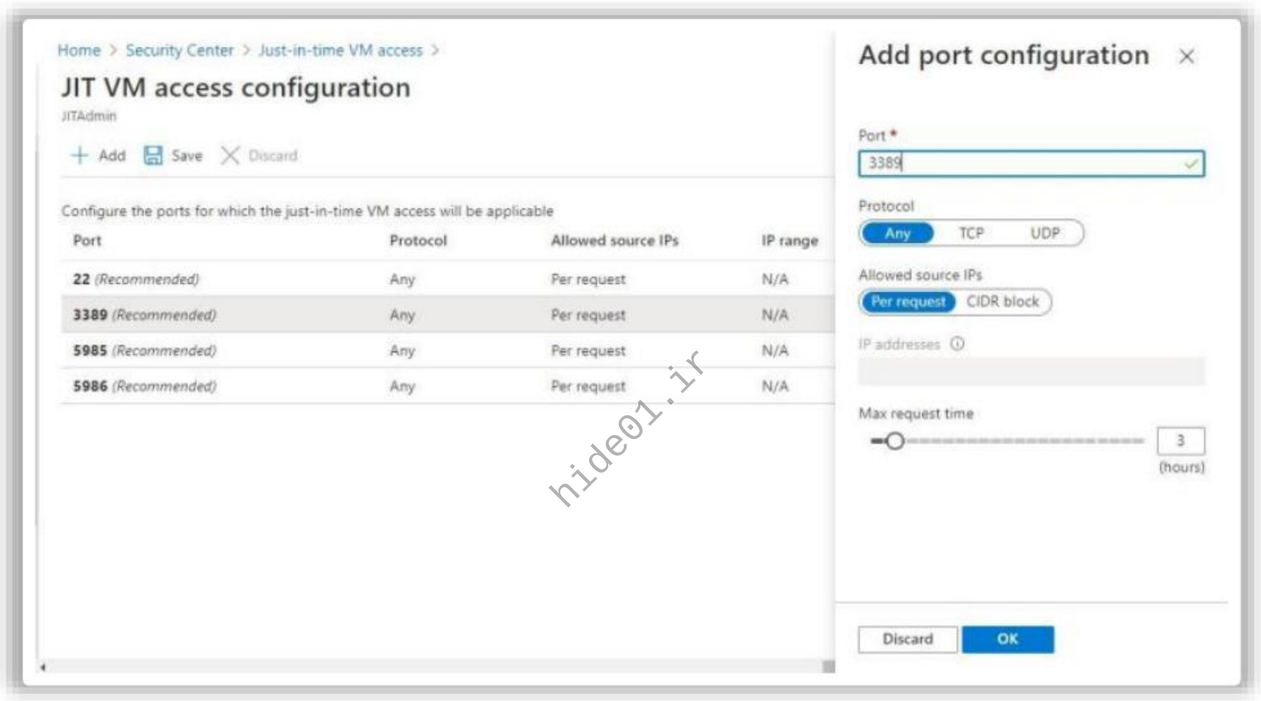


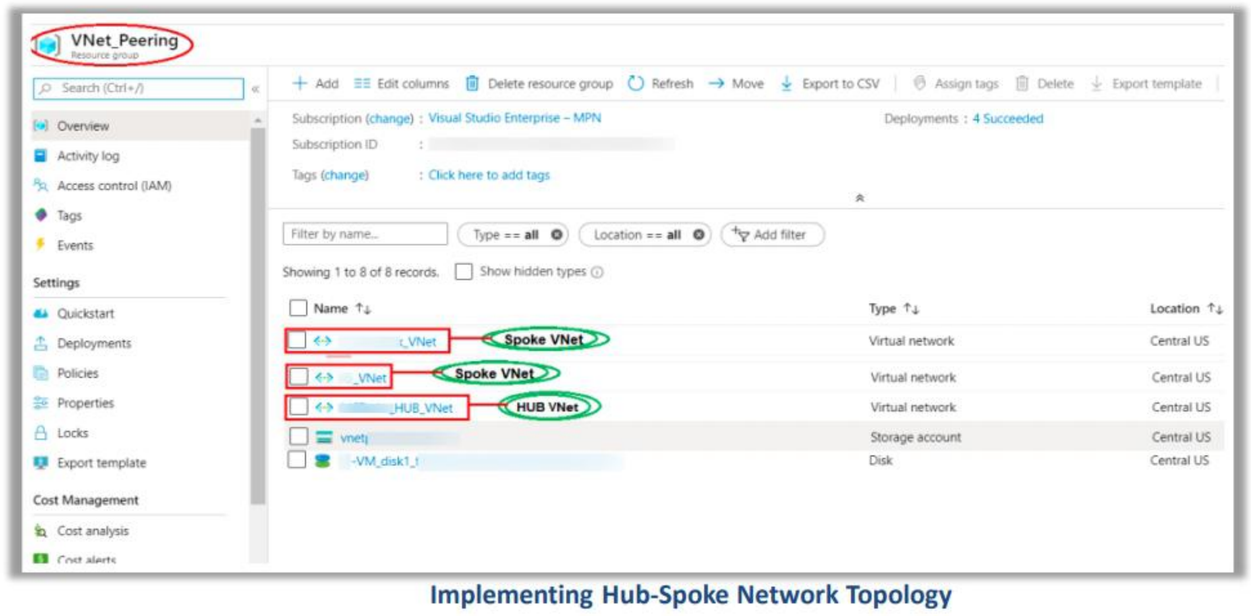
Figure: 2.96: Customizing JIT Options

- Select **Save**.

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

Implement Hub-Spoke Network Topology

- The Azure security administrator can implement hub-spoke network topology to isolate the workloads in various resource groups
- In this topology, Azure virtual network is the hub that serves as a central point of connectivity, and the spokes are the virtual network that connects to the hub virtual network by utilizing peering
- The administrator can deploy the shared services in the hub and individual workloads as spokes



Implementing Hub-Spoke Network Topology

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

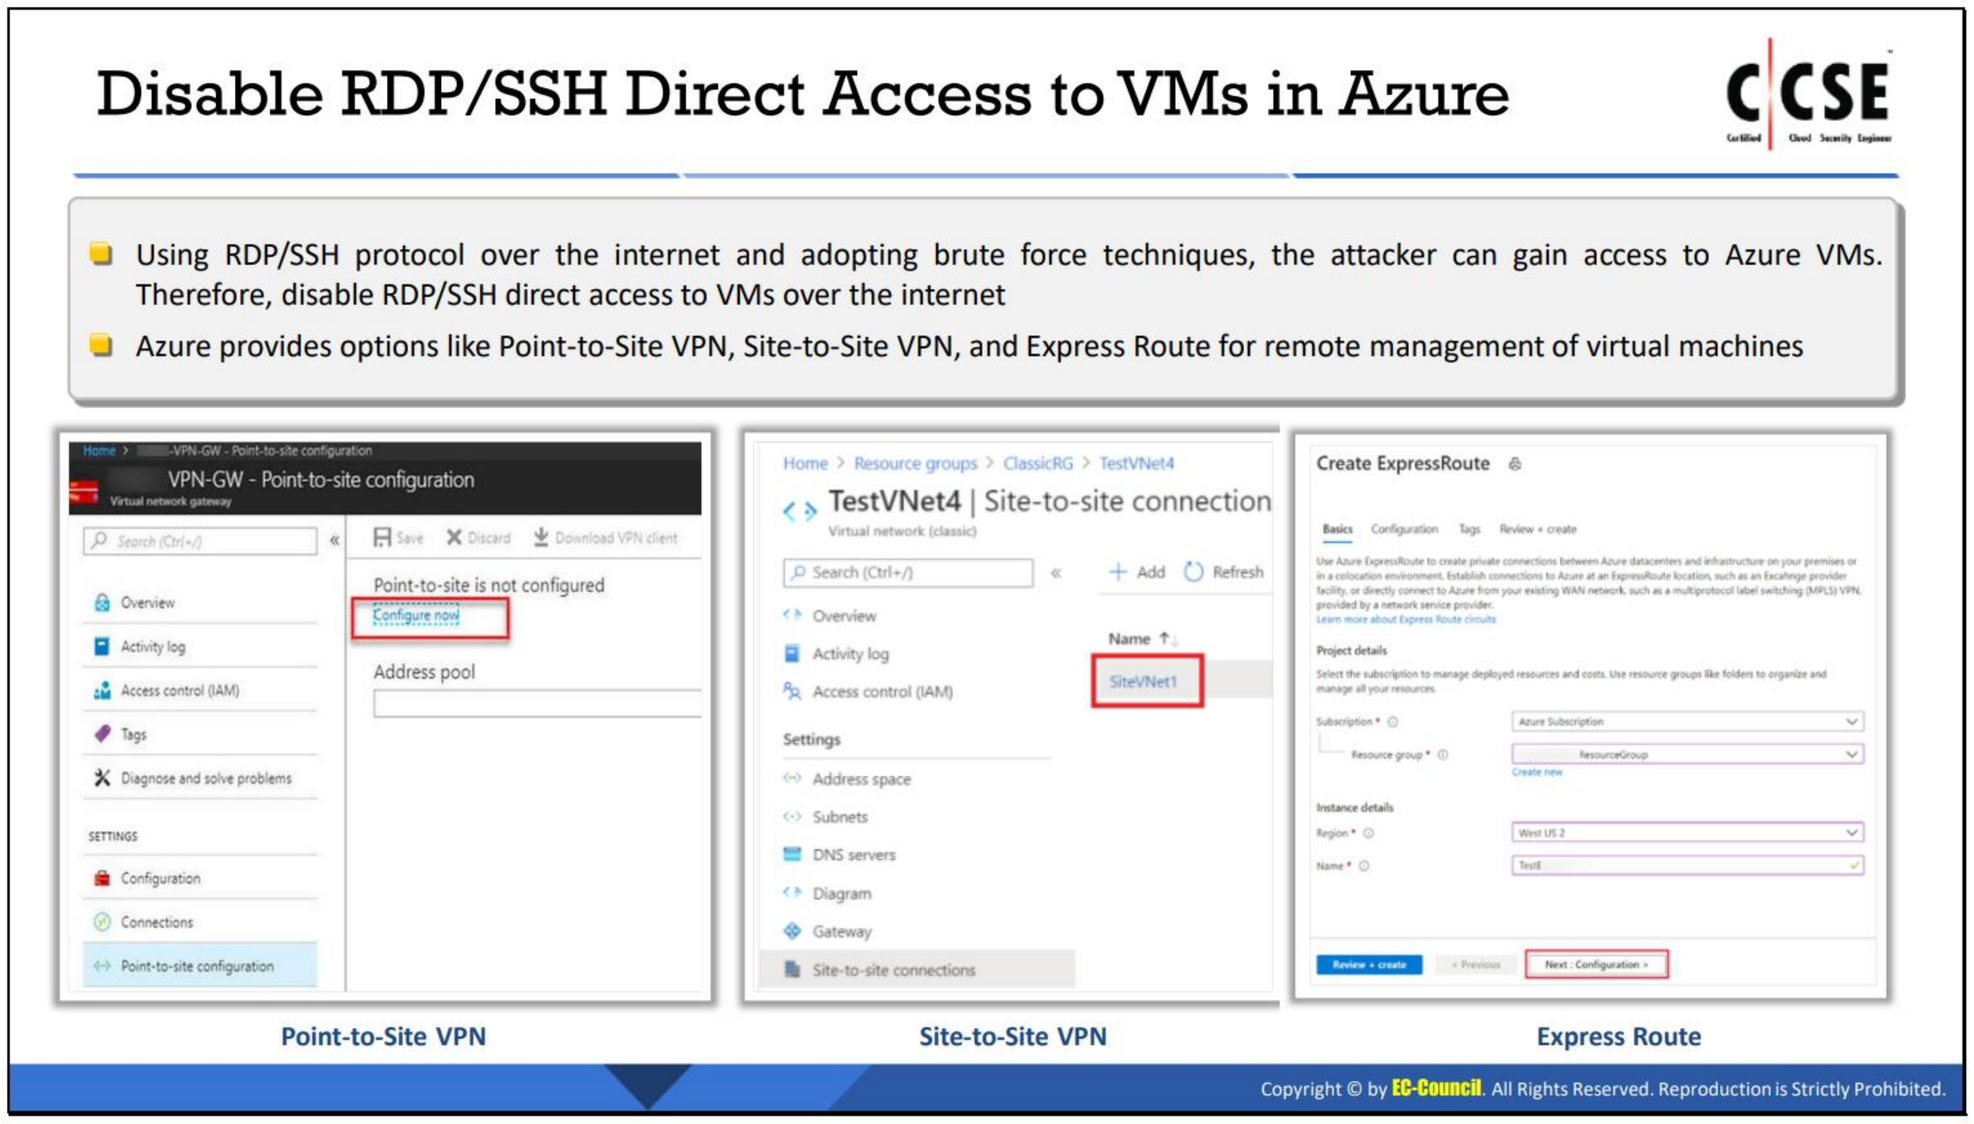
Implement Hub-Spoke Network Topology

Hub-Spoke network topology is intended to isolating workloads when sharing services such as identity and security. Here, the Hub is a virtual network in Azure, which is treated as a central point of connectivity and the Spoke is also a virtual network, which connects to the hub virtual network through peering. The Hub helps to deploy shared services and the Spoke helps to deploy individual workloads. Generally, large organizations use a hub-spoke configuration while small networks prefer simpler design to save costs and complexity.

Centralize Azure services containing connections to on-premise networks, firewalls, and isolations among virtual networks using hub-spoke topology. Here, the centralized hub virtual network offers connectivity to on-premises networks and hosts services used by the workloads that are hosted in spoke virtual networks.

Use spoke virtual networks to isolate workloads. Here, each workload will have multiple ties and subnets connected to Azure load balancers. Establish hub-spoke topology in various resource groups and subscription, which allows for decentralized management of workloads.

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



Disable RDP/SSH Direct Access to VMs in Azure

The remote desktop protocol (RDP) and secure shell (SSH) protocol help to connect with Azure VMs remotely. Using RDP/SSH protocol over the internet and adopting brute-force techniques, the attacker can gain access to Azure VMs. If attackers access a VM, then they could use it as a launch point to compromise other VMs on the virtual network or attack network devices outside the Azure cloud. Therefore, disable RDP/SSH direct access to VMs over the internet.

Cloud consumers can use other options such as Point-to-Site VPN, Site-to-Site VPN, and Express Route to access VMs for remote management.

Point-to-Site VPN

Use Point-to-site VPN (a remote access VPN client/server connection) to allow a single user/organization to connect to an Azure virtual network over the internet. Initially, establish the point-to-site connection. Later, the user who has connected to Azure virtual network via Point-to-Site VPN can access Azure VMs. Point-to-Site VPN uses protocols such as Secure Socket Tunneling Protocol (SSTP), Open VPN Protocol, and IKEv2 VPN.

Point-to-Site VPN is beneficial for people who want to connect to Azure virtual networks (VNets) from remote locations like a conference and home. It is also useful in situations where a few clients need to connect to a VNet.

Site-to-Site VPN

A Site-to-Site VPN method enables an entire on-premises network to access VMs on Azure virtual network over the Internet. It allows the users/organizations to access Azure VMs using the RDP/SSH protocol instead of direct RDP/SSH access over the Internet.

Site-to-Site VPN provides the below security benefits.

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

- **Secure Connectivity:** Site-to-Site VPN provides secure connection for all traffic that flows across it by encrypting it. This helps to protect the traffic against data modification and eavesdropping.
- **Simplified Network Architecture**: Site-to-Site VPN simplifies the network architecture by eliminating the need of converting internal IP addresses to external IP addresses. Devices used inside an organization usually use internal IP addresses which are converted to external IP addresses to make them accessible from the Internet.
- **Access Control**: The administrator can define access control rules in a simpler and secure way in the site-to-site VPN because S2S VPN users are internal users so traffic entering the VPN can be blocked or denied from accessing the resources.

ExpressRoute

ExpressRoute works like Site-to-Site VPN. Use a dedicated WAN link to get its functionality. This dedicated WAN link does not go through the internet and they are stable and perform better than Site-to-Site VPN.

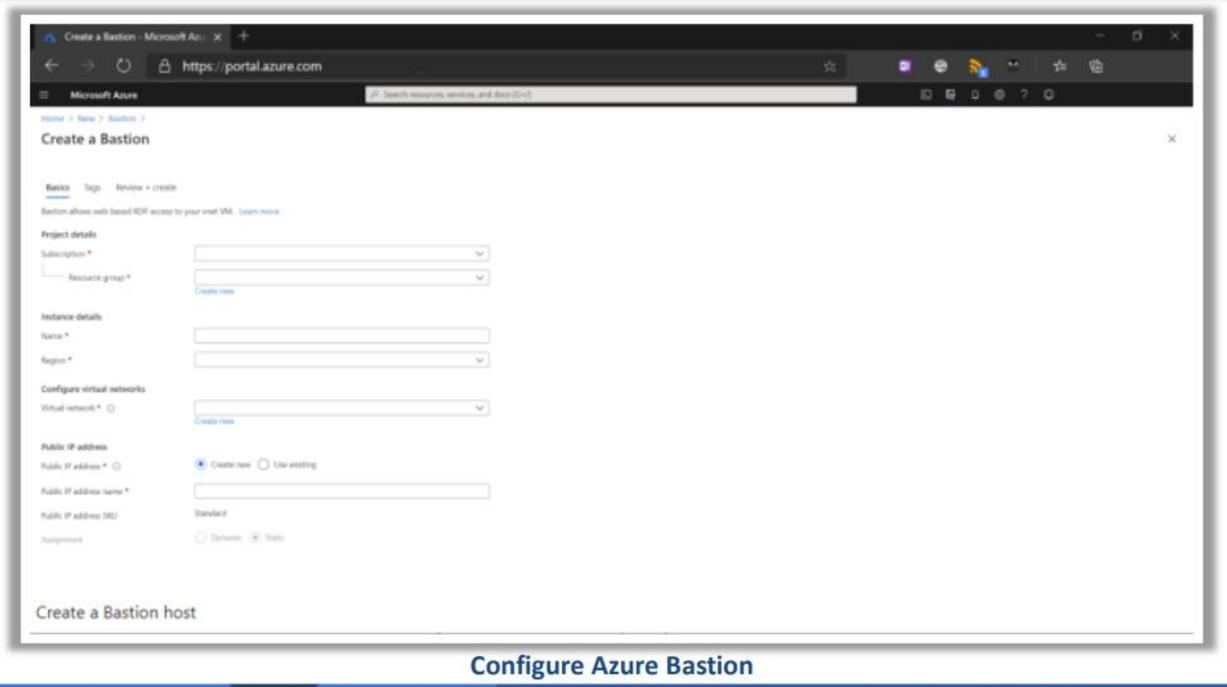
Azure provides Express route service to let customers exchange data between their on-premises or co-locating (shared) infrastructure and Azure data centers by creating private connections. It offers faster speed, lower latency, and higher reliability than other Internet connections as it does not use the Internet to connect to Azure data centers. Organization can get cost benefits by using Express Route connections for transferring data between Azure and on-premises infrastructure.

hide01.ir

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

Configure Azure Bastion to Secure RDP/SSH Access to Azure VMs

- Azure Bastion Service is a fully managed PaaS service which you deploy in the virtual network
- It safeguards the VMs from exposing the RDP/SSH ports to the outside world and establishes secure RDP/SSH Access to the virtual machines



Configure Azure Bastion

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

Configure Azure Bastion to Secure RDP/SSH Access to Azure VMs

Using SSH or RDP for logging in to a VM opens ports on the VM, which makes it vulnerable to attacks. Deploy VMs in a VNet and make use of Azure Bastion to access the VMs. Secure RDP or SSH connectivity is possible by using Azure Bastion without the need of opening ports on the VMs. The Azure Bastion service is a completely platform-managed service that can be provisioned inside the virtual network.

It safeguards the VMs from exposing the RDP/SSH ports to the outside world and establishes secure RDP/SSH access to the virtual machines. Azure Bastion uses TLS to provide secure connectivity to VMs from the Azure portal. Using this service eliminates the need of having a public IP address, client software, or agent.

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

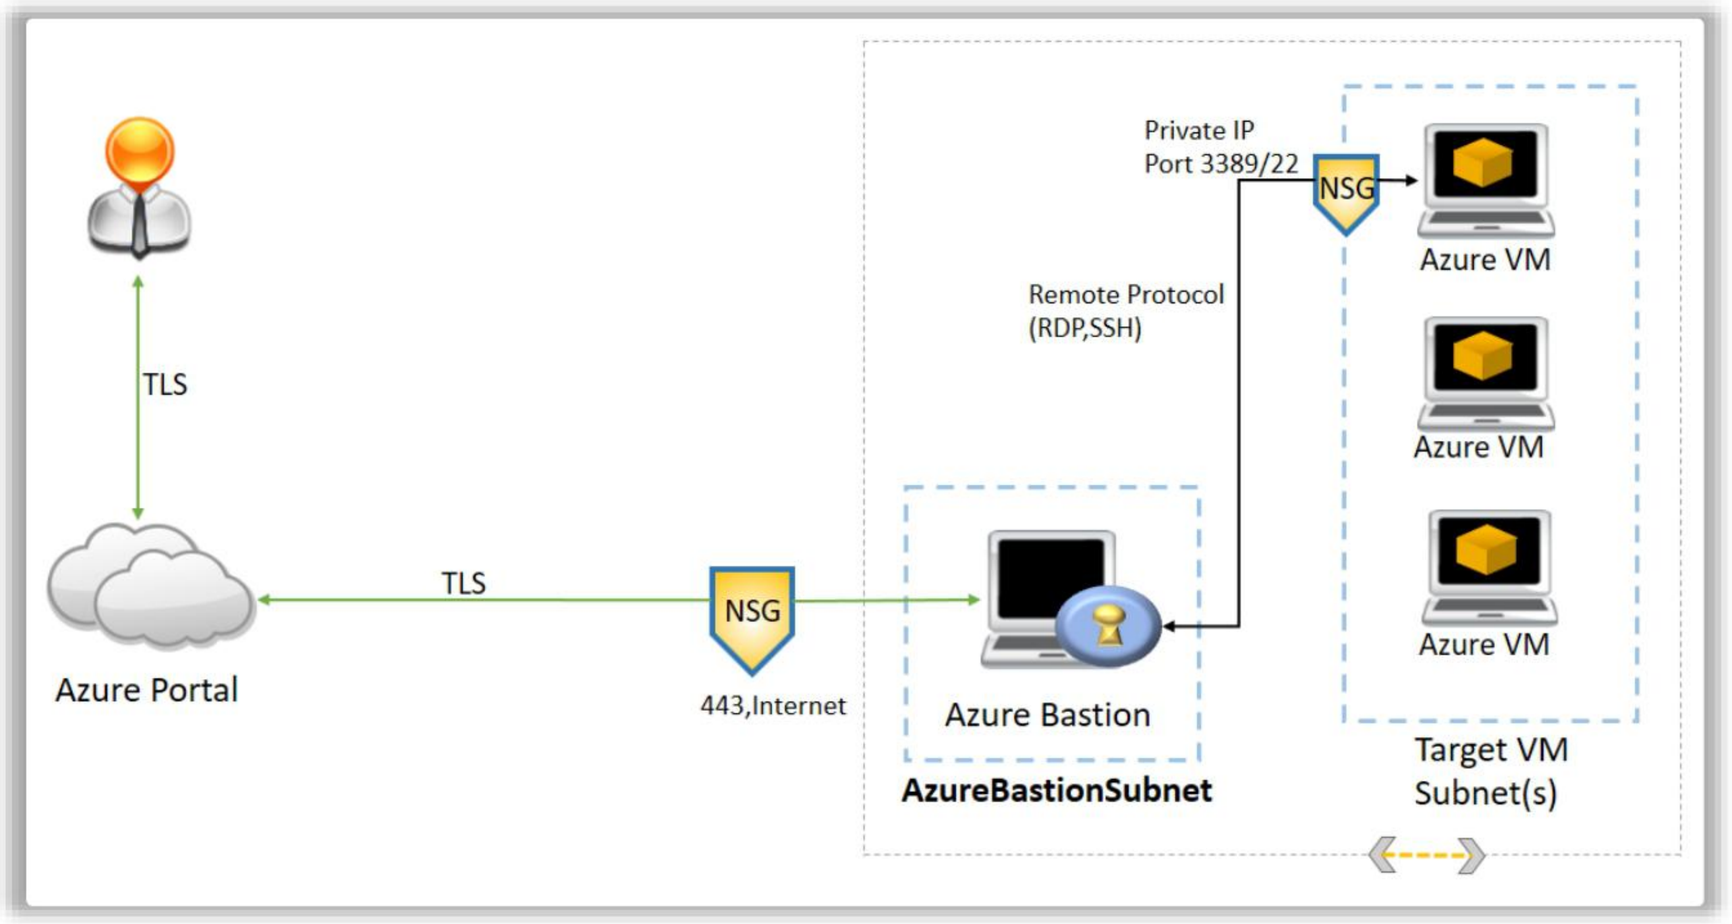


Figure: 2.97: Azure Bastion Service

The administrator can use this service to prevent VMs from exposing SSH/RDP ports, while maintaining secure RDP/SSH access to VMs. Azure Bastion offers the following security benefits:

- Protects from zero-day exploits
- No need of public IP on the VM
- Prevents port scanning
- Hassle-free NSG management

hide01.ir

Source: https://www.microsoft.com

Creating a Bastion Host

- Log in to the Azure portal. Choose **+ Create** a resource on the **Home** page.
- Type **bastion** in the search box on the new page and press **Enter**. Confirm the publisher is Microsoft from the search result for Bastion.
- Choose **Create**. Configure a new Bastion resource on the **Create a Bastion** page. Fill all the requested details in relevant fields.

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

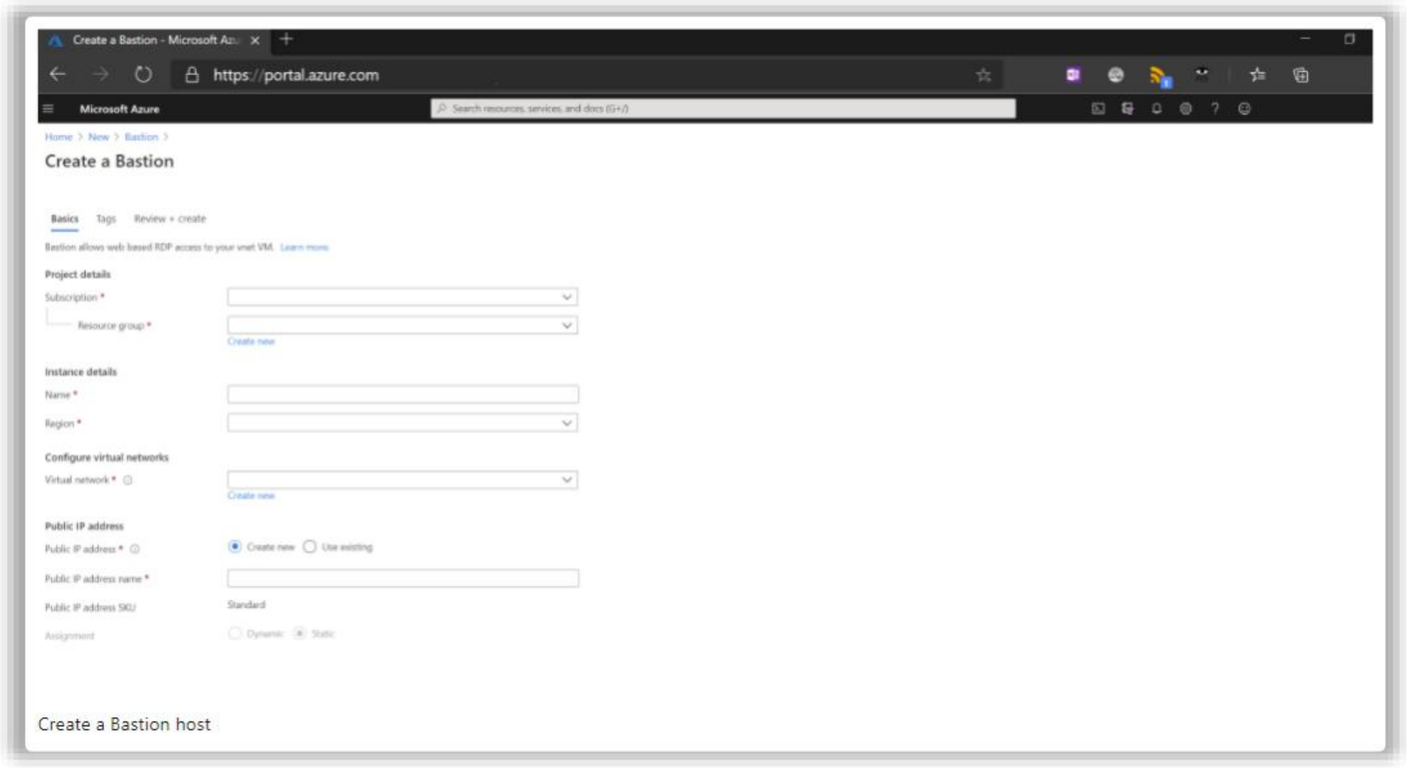


Figure: 2.98: Configuring a New Bastion Resource

- Now, choose **Review + Create** to validate the values. Bastion can be created after the validation process is successful.
- Choose **Create**. Normally, it may take up to 5 minutes to complete the creation and deployment process.

Connecting to a VM

- On the Azure portal, select the VM that must be connected. Select **Connect** and choose **Bastion** from the drop-down page.

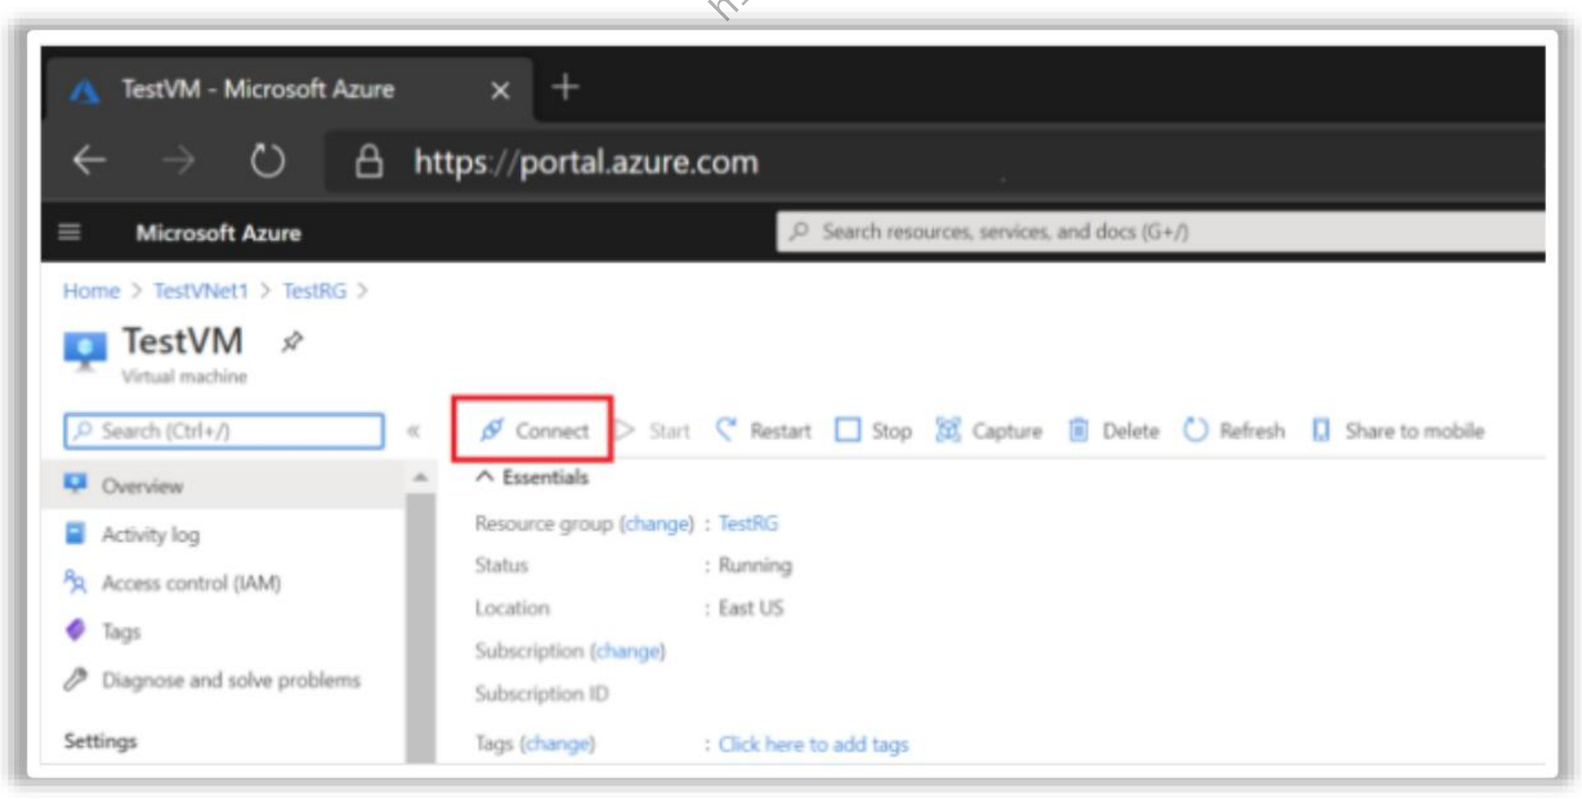


Figure: 2.99: Selecting the VM to be Connected

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

- On the next page, choose **Use Bastion**.

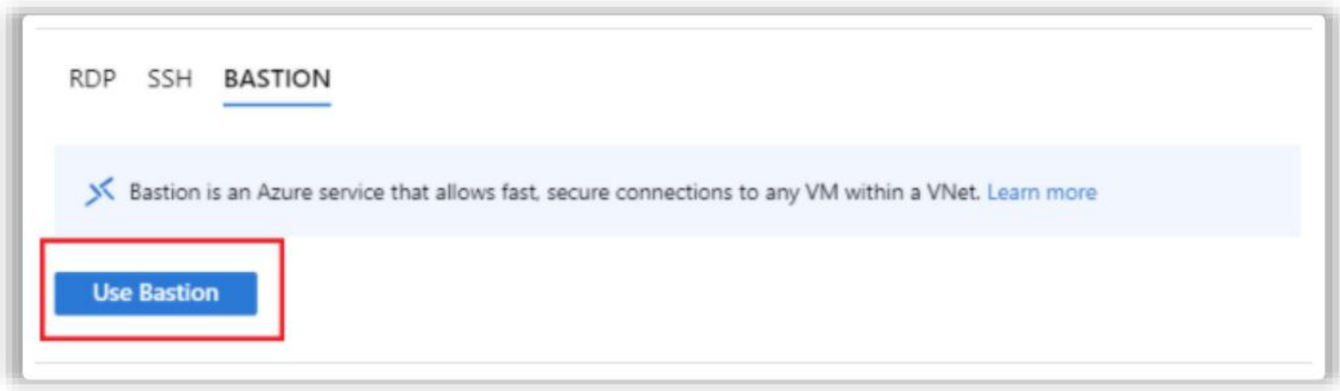


Figure: 2.102: Selecting Bastion

- Enter user credentials (username and password) for the VM on the **Connect using Azure Bastion** window. Select **Connect**.

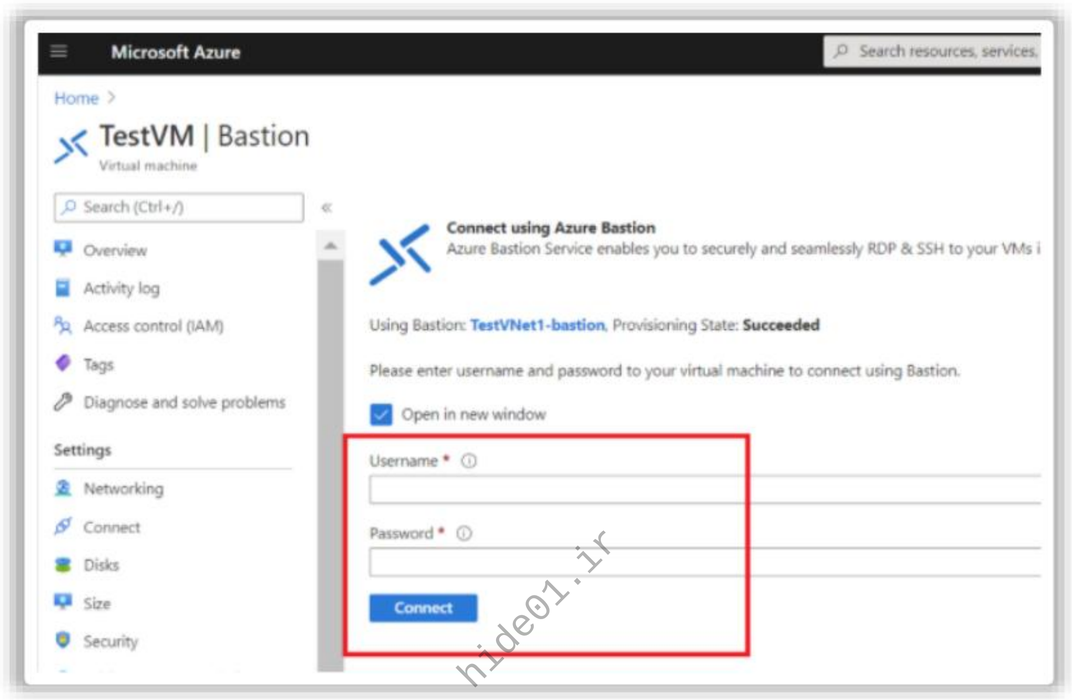


Figure: 2.100: Entering the User Credential for the VM

- This establishes an RDP connection to the VM through Bastion. The RDP connection can be accessed via the Azure portal using the Bastion service and port 443.

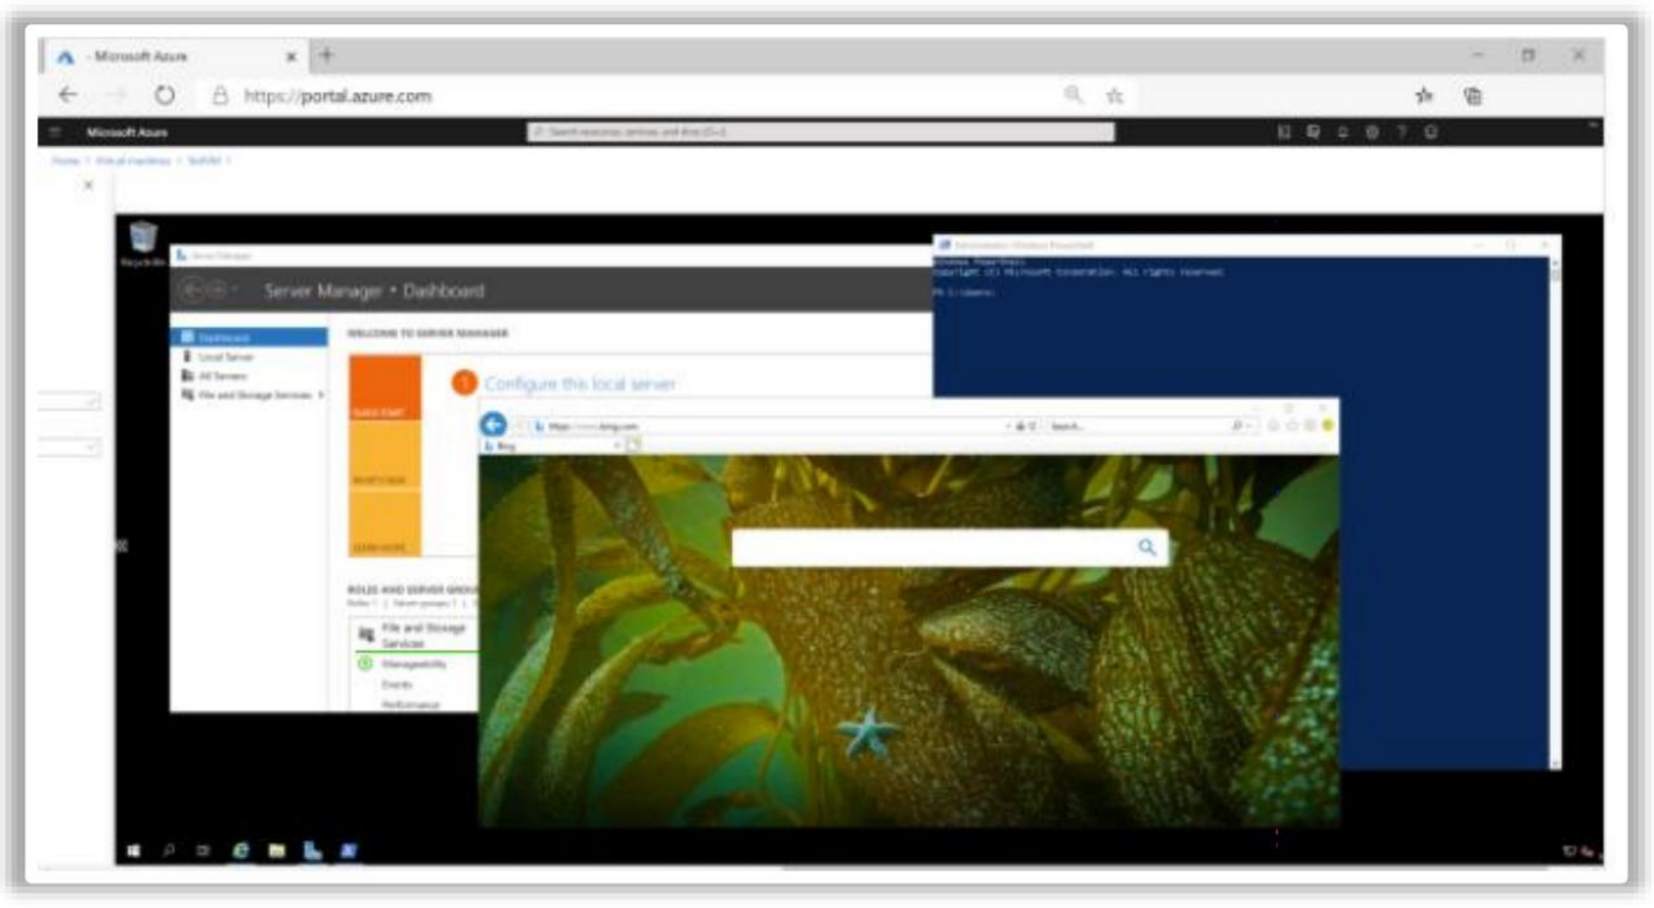
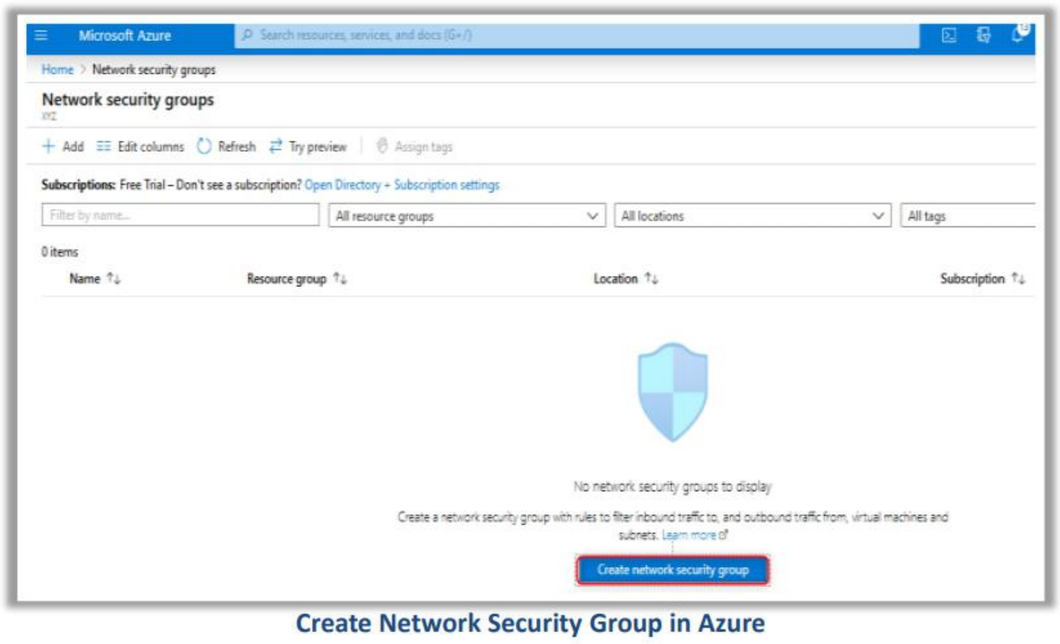


Figure: 2.101: Accessing RDP Connection via the Azure Portal

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

Use Network Security Group (NSG) to Filter Virtual Network Traffic

- The security rules of NSG allows the cloud security administrator to filter incoming and outgoing traffic to the Azure resources
- The administrator can perform the filtering by source and destination IP addresses, port, protocol



Create Network Security Group in Azure

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

Use Network Security Groups (NSG) to Filter Virtual Network Traffic

Network security groups (NSGs) are simple, stateful packet inspection devices. They comprise inbound and outbound security rules. These rules can filter traffic going from and to Azure resources. The filtering/stateful packet inspection is done using the 5-tuple approach: source IP address destination IP address, layer 4 protocol, source port, and destination port. Use NSGs to allow/deny traffic to and from a single IP address to and from multiple IP addresses (or) to and from all subnets.

For existing connections, create a flow record initially. The traffic is allowed/denied depending on the flow of the connection state of record. Here, the flow record provides a disturbed NSG. For example, generate an outbound rule to address through port 80. Now, an inbound rule is not required to respond to outbound traffic. Only create an inbound rule if traffic comes from outside. Instead, allow inbound traffic over a port. Here also, the outbound rule is not required to respond to traffic over the port.

Remove a security rule that allowed communication if you want to interrupt the existing connections. Or, stop the connections if you want to interrupt the traffic flow in any direction for a few minutes.

Source: https://www.microsoft.com

Steps to Use Network Security Groups (NSGs) to Filter Virtual Network Traffic

- **Create a Virtual Network**
  - On the Azure **Home** page, select **Create a resource**.
  - Select **Networking -> Virtual network**.

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

- Enter the below information. Leave the defaults for the remaining settings. Select **Create.**
  - **Name:** myVirtualNetwork
  - **Address space:** 10.0.0.0/16
  - **Subscription:** Select subscription
  - **Resource group:** Select **Create new** and enter **myResourceGroup**
  - **Location:** Select **East US**
  - **Subnet name:** mySubnet
  - **Subnet - Address range:** 10.0.0.0/24

- **Create Application Security Groups**
  - On the Azure portal menu/from the **Home** page, select **Create a resource**.
  - In **Search the Marketplace** box, enter *Application security group*. Select when **Application security group** appears in the search results. Select **Application security group** under **Everything** and select **Create**.
  - Enter the below information and then select **Create.**
    - **Name:** myAsgWebServers
    - **Subscription:** Select subscription
    - **Resource group:** Select **Use existing** and select **myResourceGroup**
    - **Location: East US**
  - Complete the above step again.
    - **Name:** myAsgMgmtServers
    - **Subscription:** Select subscription
    - **Resource group:** Select **Use existing** and select **myResourceGroup**
    - **Location: East US**

- **Create a Network Security Group**
  - On the Azure portal menu/from the Home page, select **Create a resource**.
  - Select **Networking -> Network security group**.
  - Enter the below information and select **Create**.
    - **Name:** myNsg
    - **Subscription:** Select subscription
    - **Resource group:** Select **Use existing** and **myResourceGroup**

hidden1.ir

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

- **Location: East US**

- **Associate Network Security Group to Subnet**
  - In the *Search resources, services, and docs* search box, search for *myNsg*. Select it when **myNsg** appears in the search results.
  - Under **SETTINGS**, select **Subnets** and select **+ Associate.**

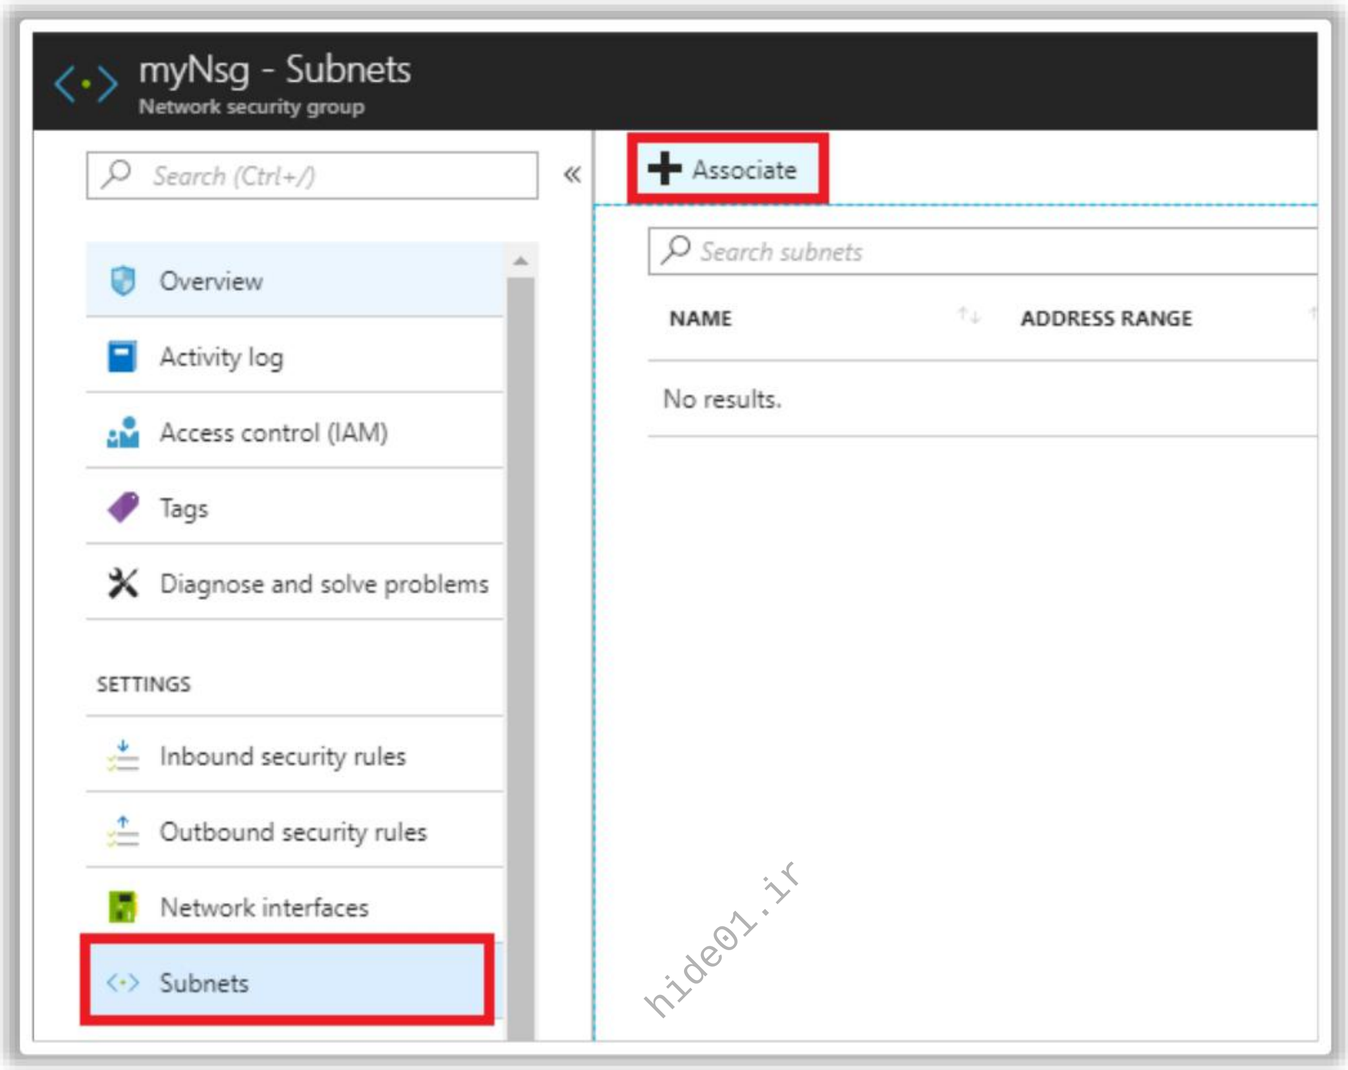


Figure: 2.103: Associating Subnets

  - Under **Associate subnet**, select **Virtual network** and **myVirtualNetwork**.
  - Select **Subnet**, select **mySubnet**, and select **OK**.

- **Create Security Rules**
  - Under **SETTINGS**, select **Inbound security rules** and select **+ Add.**

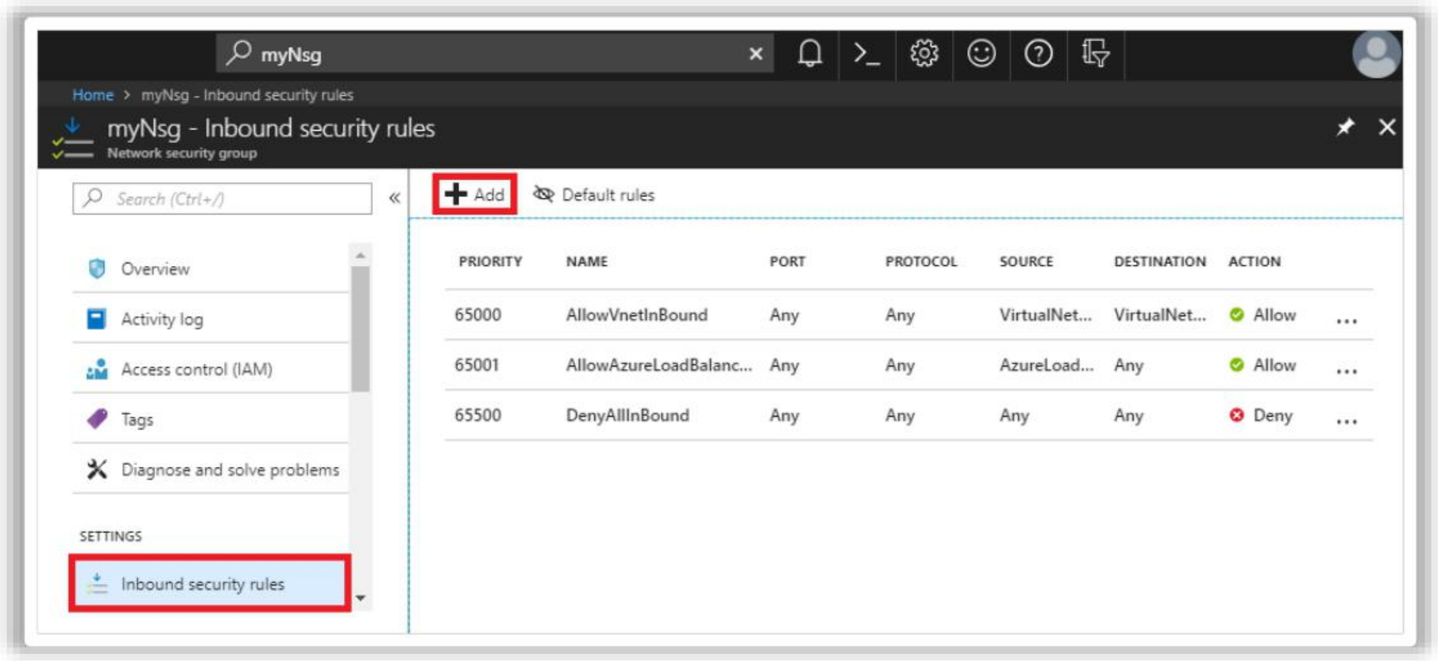


Figure: 2.104: Adding Inbound Security Rules

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

- Create a security rule that can allow port 80 and port 443 to the **myAsgWebServers** ASG. Under **Add inbound security rule**, enter the below values, accept the remaining defaults, and then select **Add.**
  - **Destination**:Select **Application security group**, and select **myAsgWebServers** for **Application security group**
  - **Destination port ranges:** Enter 80,443
  - **Protocol**: Select TCP
  - **Name:** Allow-Web-All
- Complete the above step again using the below values.
  - **Destination**: Select **Application security group** and select **myAsgMgmtServers** for **Application security group**
  - **Destination port ranges**: Enter 3389
  - **Protocol:** Select TCP
  - **Priority**: Enter 110
  - **Name**: Allow-RDP-All

In this exercise, the RDP (port 3389) is exposed to the internet for the VM assigned to the application security group **myAsgMgmtServers**. However, for production environments, connect to the Azure resources to be managed using a VPN/private network connection.

After finishing the above three steps, the rules should look like the below.

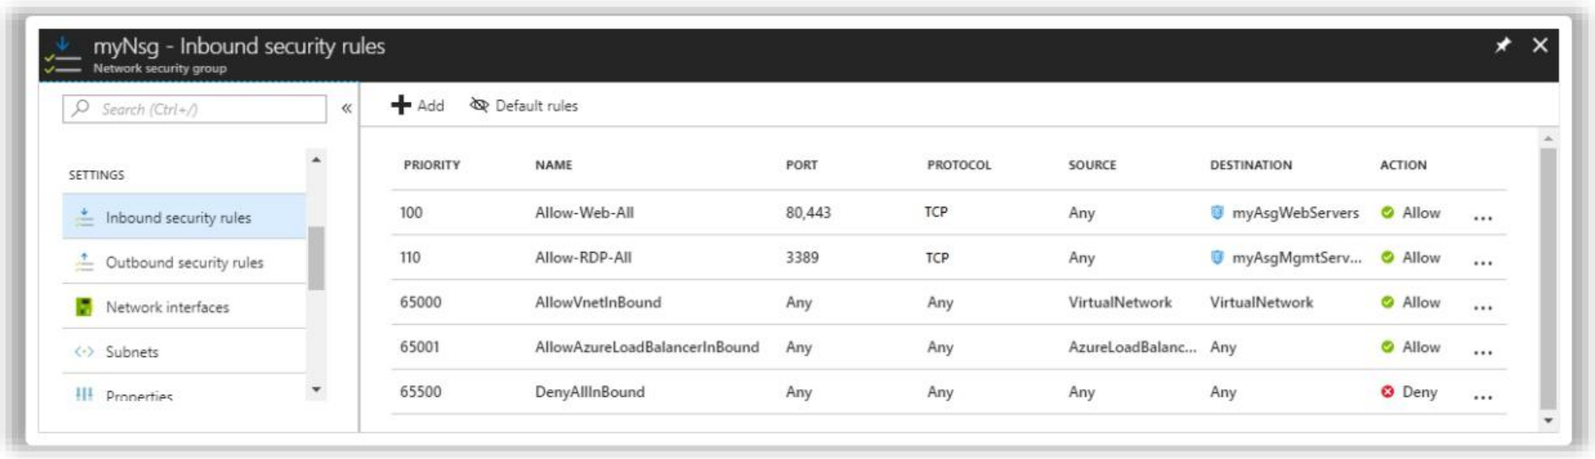


Figure: 2.105: Inbound Security Rules

- **Create Virtual Machines** (create two VMs in the virtual network)
  - **Create First VM**
    - On the Azure portal menu/from the **Home** page, select **Create a resource**.
    - Select **Compute** and select **Windows Server 2016 Datacenter**.
    - Enter the below information and leave the defaults for the remaining settings.
      - **Subscription**: Select subscription
      - **Resource group:** Select **Use existing** and select **myResourceGroup**

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

- **Name: myVmWeb**
- **Location**: Select **East US**
- **User Name:** Enter a user name
- **Password:** Enter a password

- Select a size for the VM and select **Select**.
- Under **Networking**, enter the below values and accept the remaining defaults:
  - **Virtual Network:** Select myVirtualNetwork
  - **NIC Network Security Group:** Select None**.**
- Select **Review + Create**. Select **Create** to start VM deployment.

- **Create Second VM**

  Execute the steps used for creating the first VM. However, use the VM's name as myVmMgmt. The VM takes some time to deploy. Wait until the deployment of VM is finished.

- **Associate Network Interfaces to an ASG**

  A network interface is created and attached to each VM. Add the network interface for each VM to any one of the created application security groups.

  - In the *Search resources, services, and docs* search box**,** search for *myVmWeb*. Select it when the **myVmWeb** VM appears.
  - Under **SETTINGS**, select **Networking**.
  - Select **Configure the application security groups**, select **myAsgWebServers** for **Application security groups**.
  - Select **Save.**

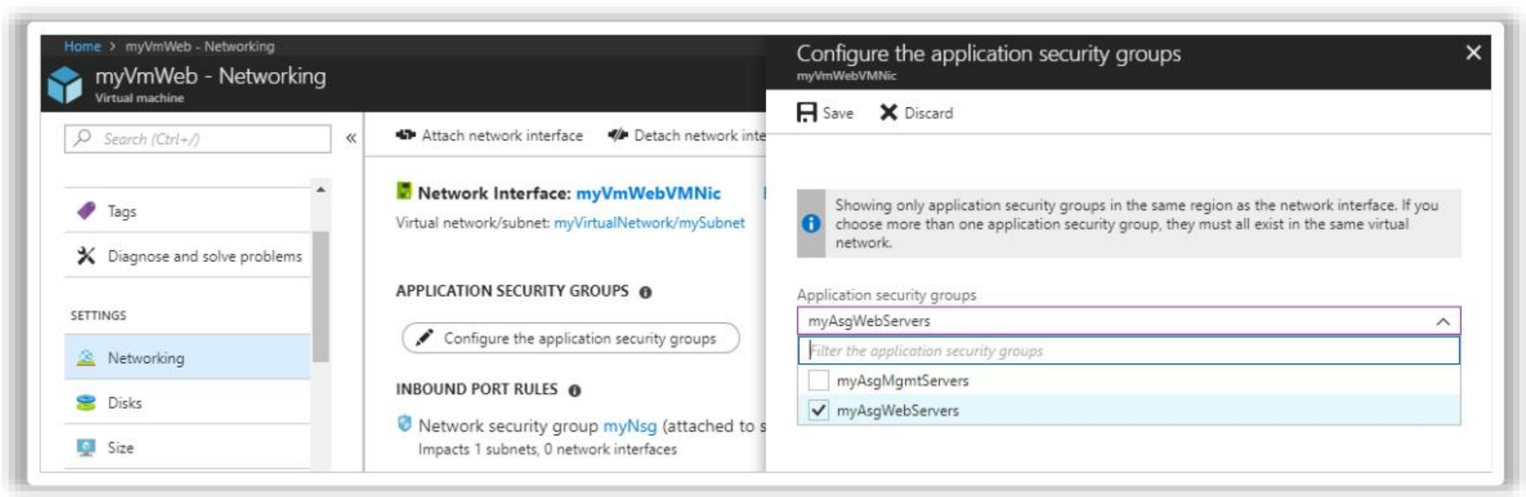


Figure: 2.106: Configuring the Application Security Groups

- In the *Search resources, services, and docs* search box**,** search for *myVmMgmt*. Select it when the **myVmMgmt** VM appears.
- Under **SETTINGS**, select **Networking**.

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

- Select **Configure the application security groups**, select **myAsgMgmtServers** for **Application security groups.**
- Select **Save.**

○ **Test Traffic Filters**

- Connect to the *myVmMgmt* VM. Enter *myVmMgmt* in the search box. Select it when **myVmMgmt** appears in the search results. Select the **Connect** button.
- Select **Download RDP file,** open the downloaded RDP file, and select **Connect**.
- Enter the specified username and password when creating the VM. Select **More choices and select Use a different account** to specify the credentials used when creating the VM.
- Select **OK**.
- A certificate warning displays during the sign-in process. Select **Yes** or **Continue** to continue with the connection.
- The connection is established because port 3389 is allowed inbound from the Internet to the *myAsgMgmtServers* ASG that the network interface attached to the *myVmMgmt* VM is in.
- Enter the below command in a PowerShell session to connect to the **myVmWeb** VM from the **myVmMgmt** VM.

```
mstsc /v:myVmWeb
```

hide01.ir

It establishes a successful connection to the **myVmWeb** VM from the **myVmMgmt** VM. Because, by default, the VMs that exist in the same virtual network can connect with each other over any port. However, a remote desktop connection cannot be created to the **myVmWeb** VM from the Internet. Because the security rule of **myAsgWebServers** disallows port 3389 inbound from the Internet. By default, the inbound traffic from the Internet is denied to all resources.

- Enter the below command from a PowerShell session on the *myVmWeb* VM to install Microsoft IIS on the **myVmWeb** VM.

```
Install-WindowsFeature -name Web-Server -IncludeManagementTools
```

- The IIS installation is completed. Disconnect from the **myVmWeb** VM, which will leave the user in the **myVmMgmt** VM remote desktop connection.
- Additionally, disconnect from the **myVmMgmt** VM.

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

- In the **Search resources, services, and docs** search box, search for **myVmWeb** from the user's PC. Select **myVmWeb** when it appears in the search results. Note the **Public IP address** of VM. Here, it is **137.135.84.74** (it may be different in your environment)**.**

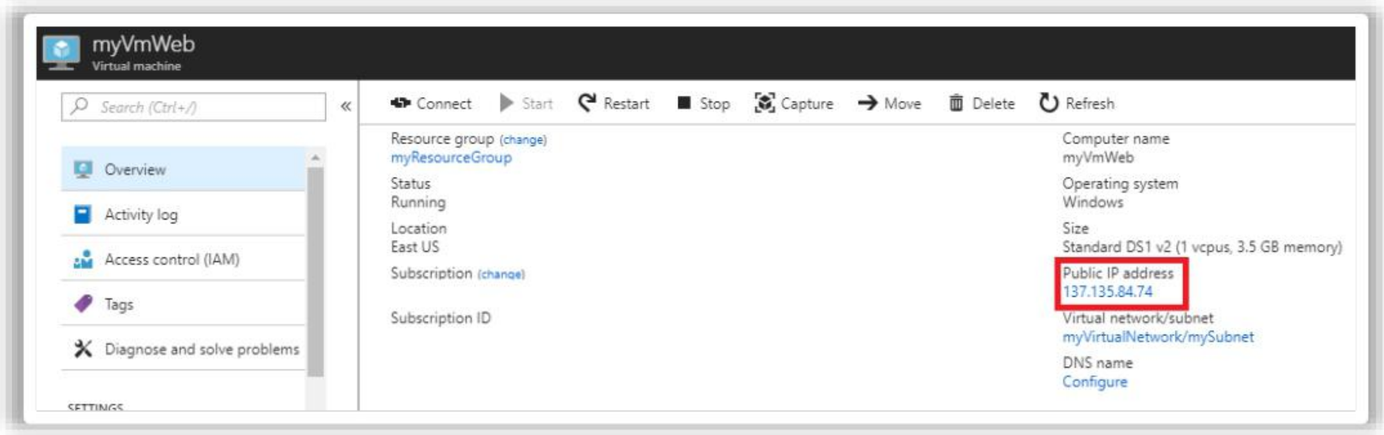
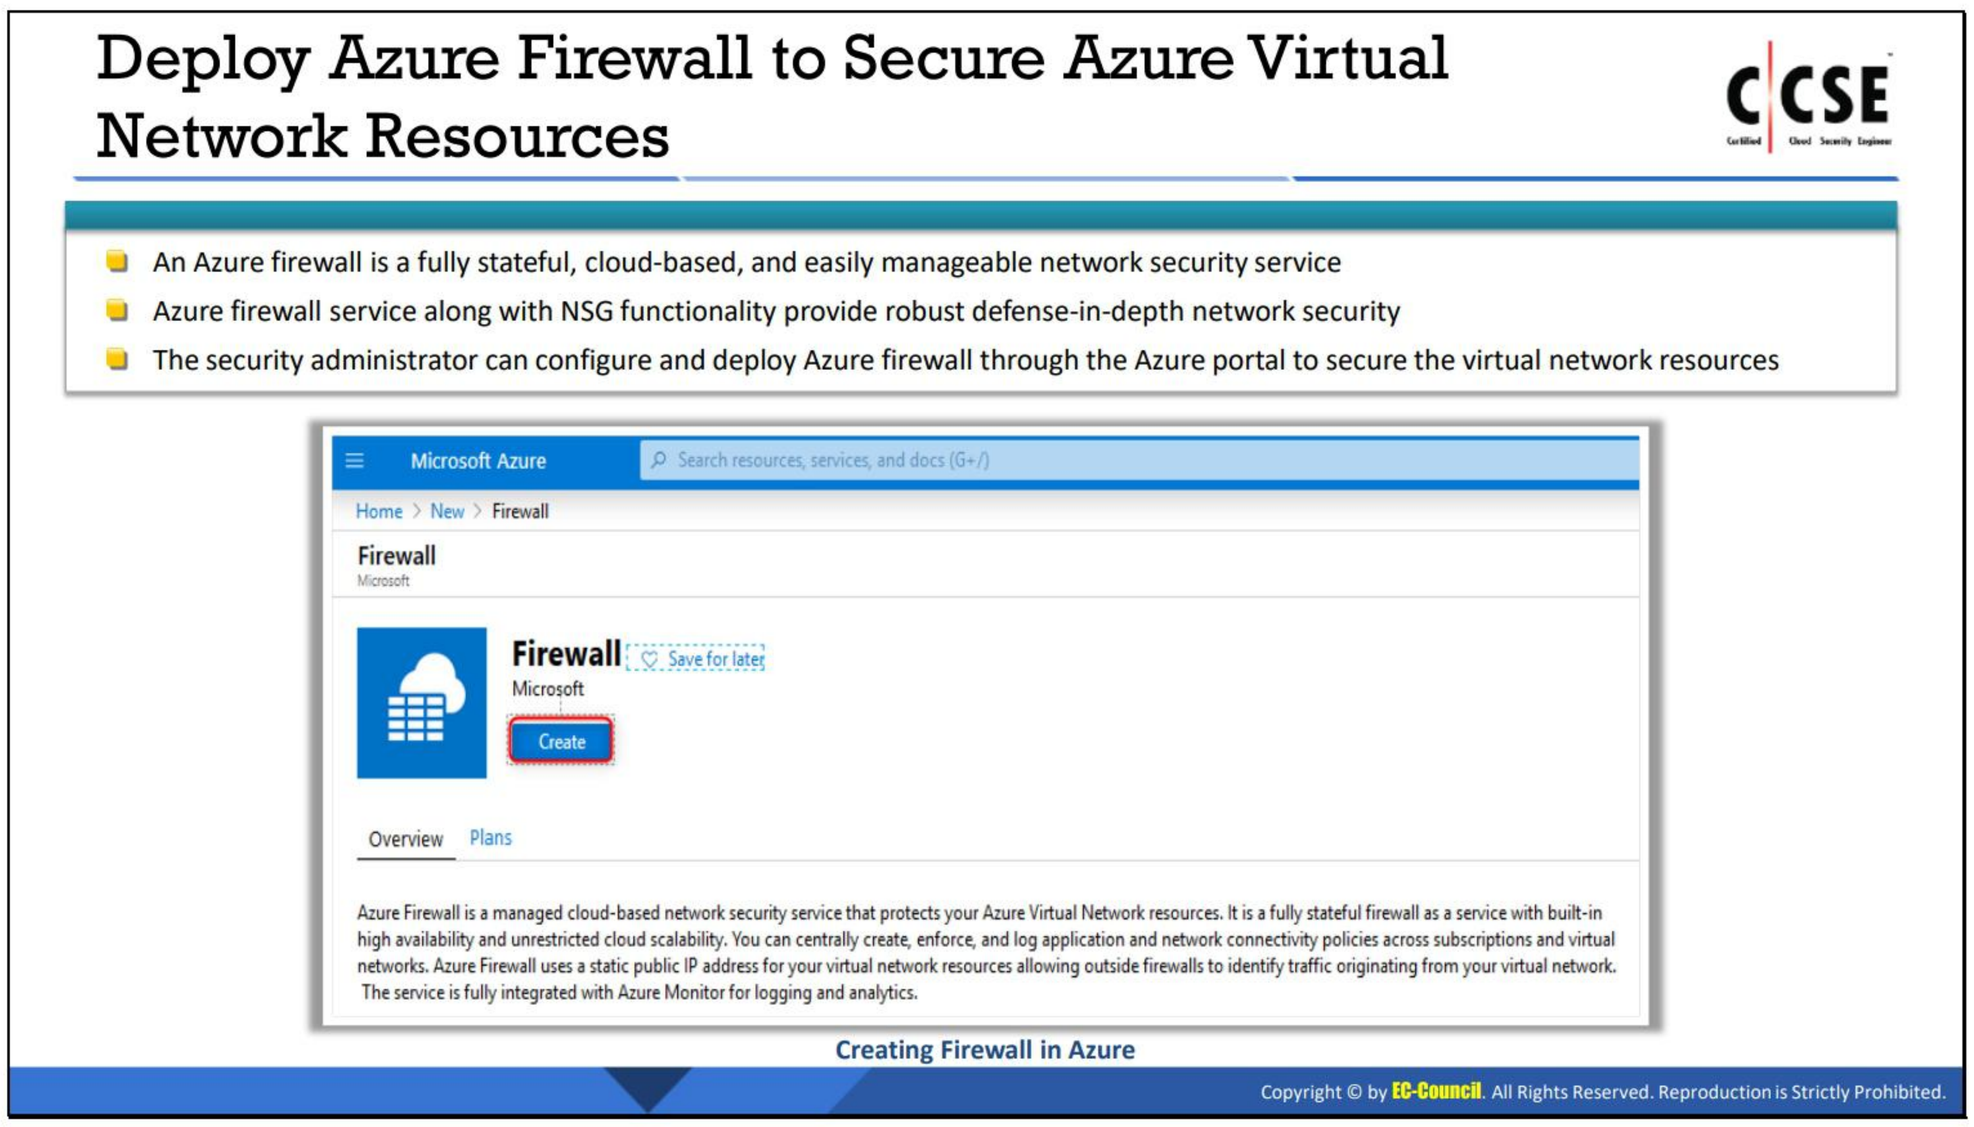


Figure: 2.107: Public IP Address of VM

hide01.ir

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



Deploy Azure Firewall to Secure Azure Virtual Network Resources

Azure Firewall is a fully stateful, cloud-based, and easily manageable network security service. It protects all virtual network resources in Azure and offers high availability and cloud scalability without restrictions. A cloud administrator can use Azure Firewall to centrally create, implement, and log network and application connectivity policies in all subscriptions and VNets.

Azure Firewall service along with NSG functionality provides robust defense-in-depth network security. It allows external firewalls to detect outbound traffic coming from a VNet as it uses a static public IP address for the VNet resources. Additionally, it is completely integrated with Azure Monitor for the purpose of logging and analytics.

Azure Firewall offers the following important benefits:

- Network traffic filtering rules to control traffic
- Threat intelligence-based filtering to block known malicious IP addresses
- Azure Monitor logging to log events
- In-built high availability without load balancers
- Application Fully Qualified Domain Names (FQDN) filtering rules

Security administrators can configure and deploy Azure Firewall through the Azure portal to secure the virtual network resources. For deploying the firewall, a resource group has to be created which will contain resources that need the firewall. After that, a VNet, subnets, and a test server must be created.

Source: https://www.microsoft.com

Creating a Resource Group

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

All resources are present in this resource group.

- Log in to the Azure portal by using https://portal.azure.com
- On the menu, choose **Resource groups** or search for it and select it. Choose **Add**.
- Fill the following information
  - **Resource group name**: Test-FW-RG
  - **Subscription**: choose a subscription
  - **Resource group location**: choose a location; ensure that all resources that are to be created should be in the same location
- Choose **Create.**

Creating a VNet

There will be three subnets in this VNet.

- Choose **Create a resource** from the **Home** page.
- Choose **Networking** and then select **Virtual network**.
- Fill the following information
  - **Subscription**: choose your subscription
  - **Resource group**: Test-FW-RG
  - **Name**: Test-FW-VN
  - **Region**: previously selected location
- Choose Next: IP addresses and enter the following information
  - **IPv4 Address space**: 10.0.0.0/16
  - **Subnet**: default
  - **Subnet name:** AzureFirewallSubnet
  - **Address range**: 10.0.1.0/26
- Choose **Save.**

hide01.ir

Now, creating a subnet

- Choose **Add Subnet** and enter the given information
  - **Subnet name**: Workload-SN
  - **Subnet address range**: 10.0.2.0/24
- Choose **Add** and then click on **Review + Create**
- Select **Create**

Creating a VM

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

A workload VM is to be created that must be placed in the subnet 'Workload-SN'.

- Choose **Create a resource** from the **Home** page.
- Choose **Compute** and then click on **Virtual machine.**
- Choose **Windows Server 2016 Datacenter** from the list. Fill the following information
  - **Resource group**: Test-FW-RG
  - **VM name**: Srv-Work
  - **Region**: Same as previously selected
  - **Image**: Windows Server 2016 Datacenter
  - **Administrator user name**: enter a user name
  - **Password**: enter a password
- Choose **None** under the **Inbound port rules**, **Public inbound ports**
- Leave other things as default and click on **Next: Disks.**
- Leave other things as default and click on **Next: Networking.**
- Ensure the following information is selected
  - **Virtual network**: Test-FW-VN
  - **Subnet**: Workload-SN
  - **Public IP**: None
- Leave other things as default and click **Next: Management**.
- Click **Off** for disabling the boot diagnostics and accept other default settings. Choose **Review + Create.**
- On the **summary** page, review the settings and choose **Create**.

hide01.ir

Deploying the Firewall

Deploying the Azure Firewall in the VNet.

- Choose **Create a resource** from the **Home** page or Azure portal.
- In the search box, type 'firewall' and press **Enter**. Choose the **Firewall** and then select **Create**.
- Configure the firewall on the **Create a Firewall page** as follows
  - **Subscription**: your subscription
  - **Resource group**: Test-FW-RG
  - **Name**: Test-FW01
  - **Location**: previously selected location
  - **Choose a virtual network**: Test-FW-VN

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

- **Public IP address**: add new with the name 'fw-pip'

- Choose **Review + Create** and review the summary
- Choose **Create** for creating the firewall. It may take some time to deploy.
- After deployment, navigate to the resource group (**Test-FW-RG**) and choose the firewall (**Test-FW01**).
- Note the private and public IP addresses of the firewall to use later.

Creating a Default Route

To go through the firewall, configure the outbound default route for the Workload-SN subnet.

- Select **All Services** on the Azure portal menu. Search and select it if not found directly.
- Choose **Route tables** under **Networking**.
- Choose **Add** and enter the following information
  - **Name**: firewall-route
  - **Subscription**: your subscription
  - **Resource group**: Test-FW-RG
  - **Location**: same as selected before
- Choose **Create.**
- Click on **Refresh** and then choose the route table '**Firewall-route**'
- Choose **Subnets** and then click on **Associate**.
- Select **Virtual network** and then click on **Test-Fw-VN**
- Choose **Workload-SN** for **Subnet** and click on **OK**.
- Choose **Routes** and then select **Add**. Fill the following information
  - **Route name**: fw-dg
  - **Address prefix**: 0.0.0.0/0
  - **Next hop type**: Virtual appliance
  - **Next hop address**: firewall's private IP address that was noted down
- Click on **OK.**

hide01.ir

Configuring an Application Rule

This application rule allows outbound access or connectivity to www.google.com

- Go to the **Test-FW-RG** tab and choose the firewall (**Test-FW01**).
- On the **Test-FW01** page, choose **Rules** under **Settings**.
- Choose the **Application rule collection** tab and select **Add application rule collection**. Enter the following details

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

- **Name**: App-Coll01
- **Priority**: 200
- **Action**: Allow

- Under **Rules, Target FQDNs**, type **Allow-Google** for **Name**
  - **Source type**: choose **IP address**
  - **Source**: 10.0.2.0/24
  - **Protocol:port**: http, https
  - **Target FQDNs**: www.google.com
- Choose **Add**

Configuring a Network Rule

This network rule provides outbound access to 2 IP addresses at DNS port 53

- Choose the **Network rule collection** tab.
- Choose **Add network rule collection** and enter the following details:
  - **Name**: Net-Coll01
  - **Priority**: 200
  - **Action**: Allow
- Under **Rules, IP addresses**, type **Allow-DNS** for **Name**
  - **Protocol: UDP**
  - **Source type**: choose **IP address**
  - **Source**: 10.0.2.0/24
  - **Destination type:** choose **IP address**
  - **Destination address**: 209.244.0.3, 209.244.0.4 (these are DNS servers of CenturyLink)
  - **Destination Ports:** 53
- Click on **Add.**

Configuring a DNAT Rule

A DNAT rule helps in connecting a remote desktop to Srv-Work VM via the firewall.

- Click on the **NAT rule collection** tab.
- Choose **Add NAT rule collection** and fill the following information
  - **Name**: rdp
  - **Priority**: 200

ide01.ir

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

- Under **Rules,** type **rdp-nat** for **Name**
  - **Protocol: TCP**
  - **Source type**: choose **IP address**
  - **Source**: *
  - **Destination address**: enter the firewall's public address
  - **Destination Ports**: 3389
  - **Translated address**: enter the private IP address of Srv-Work
  - **Translated port**: 3389
- Click on **Add.**

Testing the Firewall

Test the firewall to ensure that it is working as intended.

- Connect a remote desktop to the IP address of the firewall and log in to the VM (Srv-Work).
- Go to the **Internet Explorer (IE)** and try to access www.google.com
- Choose **OK** and **Close** on the IE security alerts. The Google homepage can be accessed.
- Now, try to access www.microsoft.com. The firewall blocks access to it.
- It has been verified that the firewall is allowing only the configured FQDN and blocking any other IP address.

h4de01.ir

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



Use Azure Security Center

Azure Security Center (ASC) is a security management system that helps in securing the data center and offers advanced threat protection to workloads in the cloud and on-premises. ASC is the most important service that should be used to ensure a secure cloud environment on Azure. It helps organizations in implementing security best practices by not only indicating the places where they need to be implemented but also allowing to implement these practices directly from the security center.

The cloud administrator can use the ASC to check the security of the resource and to be alerted if there is anything that requires attention. The administrator can utilize it to enforce security policies across workloads, get security alerts and security recommendations about the environment, reduce the threat exposure, and identify and respond to cyberattacks.

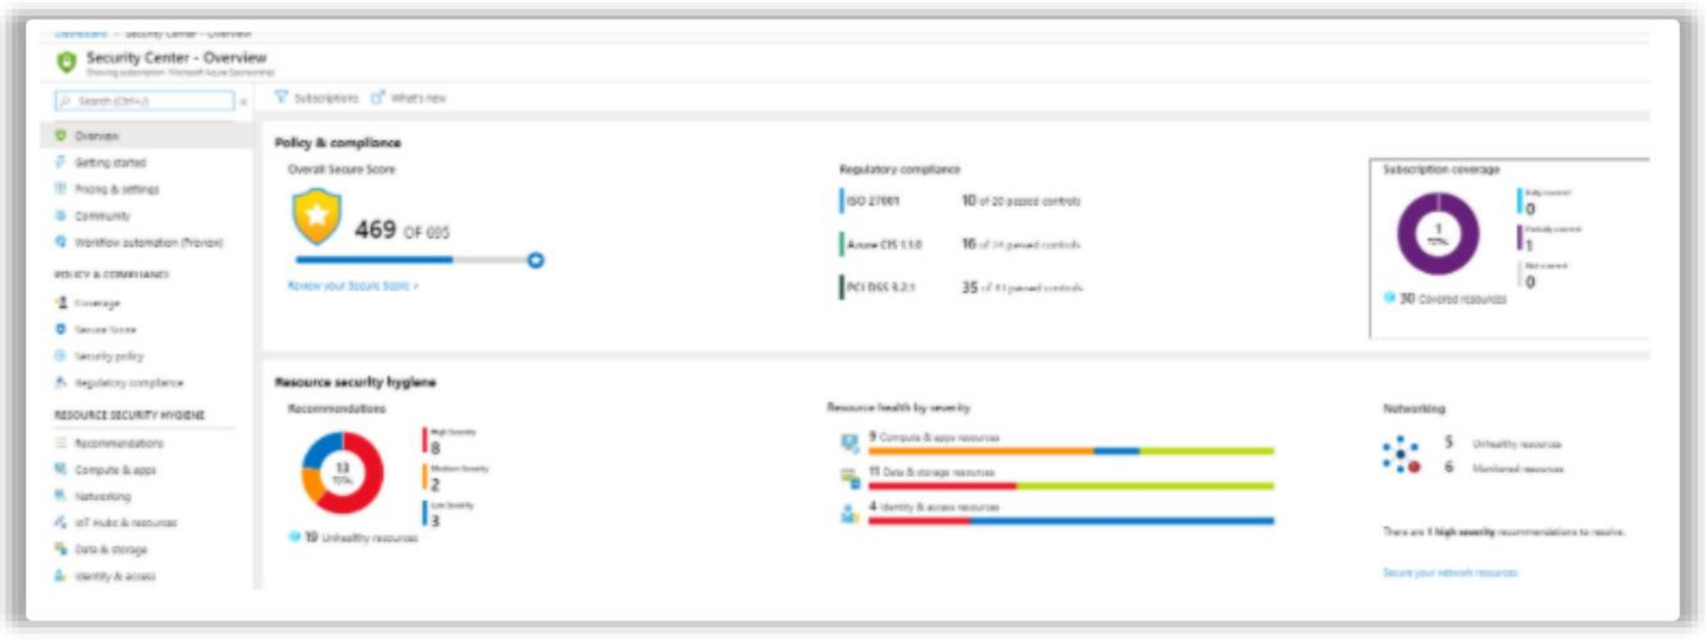


Figure: 2.108: Overview of Azure Security Center

Source: https://www.microsoft.com

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

ASC addresses the most important challenges such as sophisticated attacks, rapidly changing workloads, and shortage of skilled security staff by offering tools for the below.

- **Strengthening Security Posture**: ASC evaluates the environment and helps the administrator to know the status of resources and their security posture.
- **Protecting Against Attacks**: ASC evaluates workloads and provides threat prevention recommendations along with security alerts.
- **Better and Faster Security**: As ASC is integrated with Azure services, its deployment is much easier and faster. The administrator can use it for auto-provisioning and protecting Azure services. It discovers and on-boards resources automatically.

Source: https://www.microsoft.com

Enabling ASC

Log in to the Azure portal and choose **Security Center** from the menu.

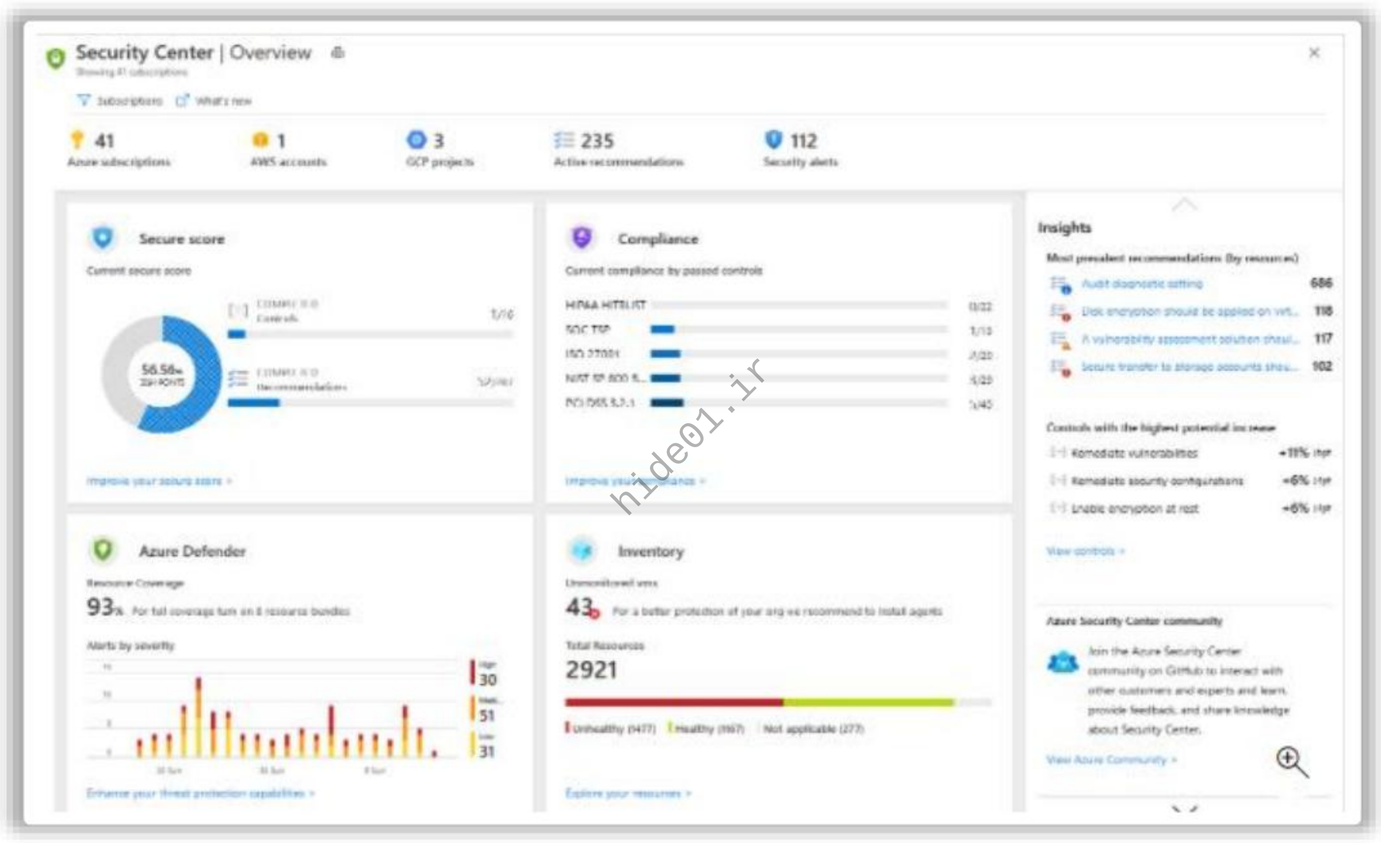


Figure: 2.109: Azure Security Center

Customer security recommendations across all workloads are reviewed by the ASC. It determines how critical each recommendation is by applying advanced algorithms and then calculates the Secure Score based on them. The Secure Score helps customers to assess or evaluate their workload security posture. It can be accessed by customers from their dashboard. The Overview page in the Security Center dashboard shows the overall secure score.

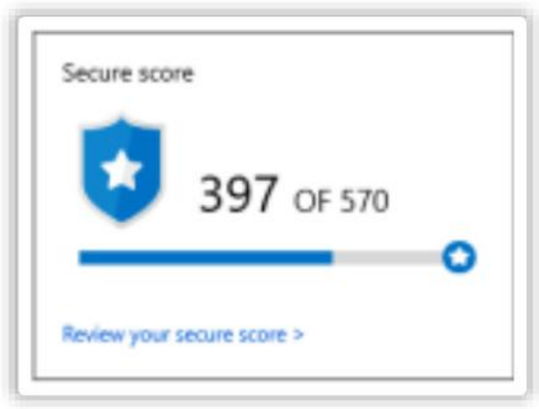


Figure: 2.110: Overall Secure Score

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

Customers can know more about their Secure Score from their Secure Score dashboard by clicking the **Review your secure score** button.

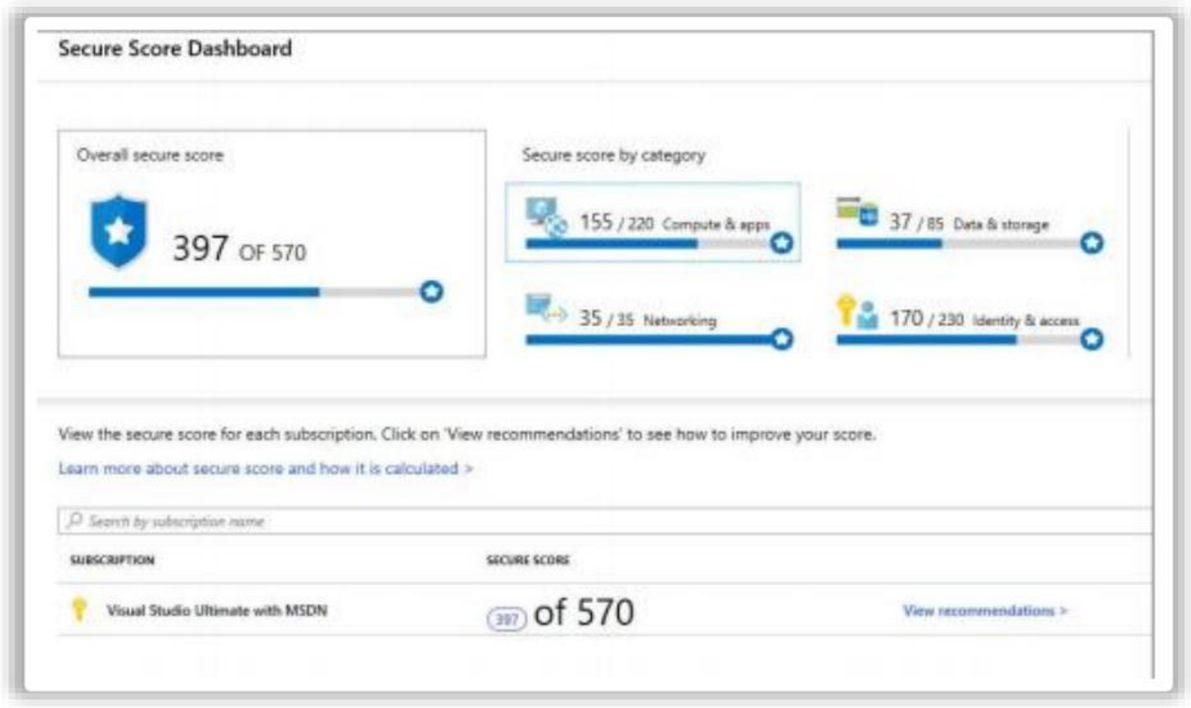


Figure: 2.111: Reviewing Secure Score

This dashboard provides customers with better visualization of how their workloads influence their Secure Score. It shows which workloads are fully compliant and which ones are not. Customers can see the active recommendations by clicking the **View recommendation** option on the dashboard. Besides offering security recommendations, it maps assessments to many regulatory compliance standards as well. It shows which resources are not passing compliance and provides steps for remediation for every assessment under regulatory compliance.

hide01.ir

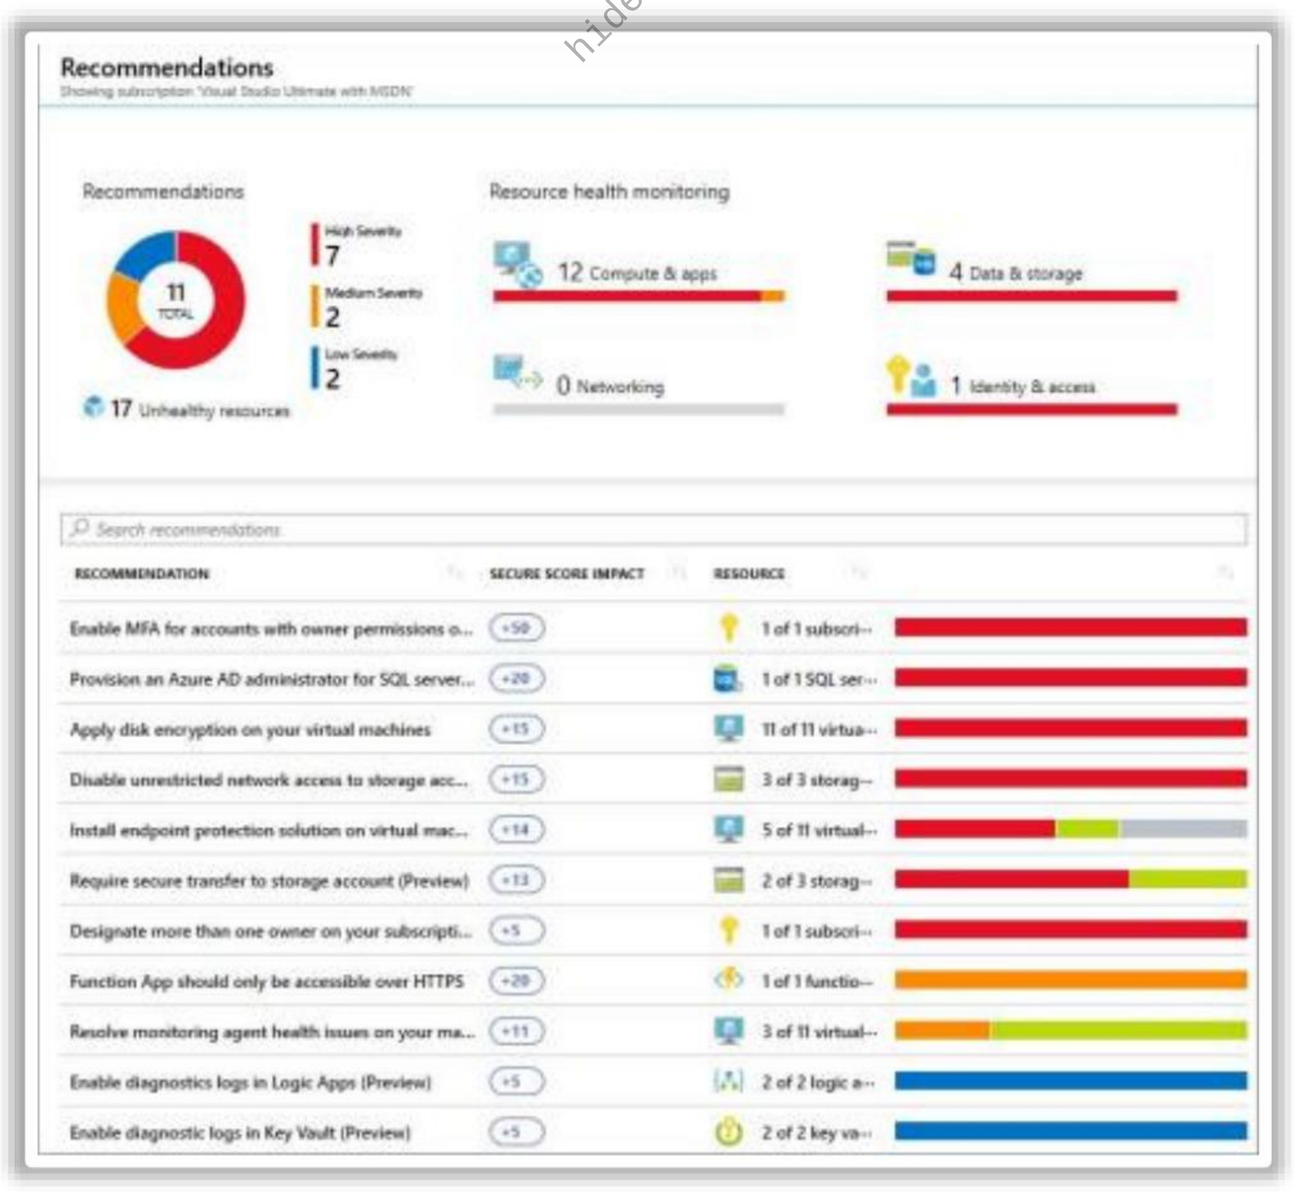


Figure: 2.112: Active Recommendations for Customers

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

Enabling Azure Defender and installing log analytics are essential to take complete advantage of ASC.

hide01.ir

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

Source: https://www.microsoft.com

Enabling Azure Defender

Follow the given steps to enable Azure Defender from the ASC dashboard.

- Choose **Getting started** on the Security Center's sidebar. Choose the **Upgrade** tab that shows all workspaces and subscriptions that can be on-boarded.

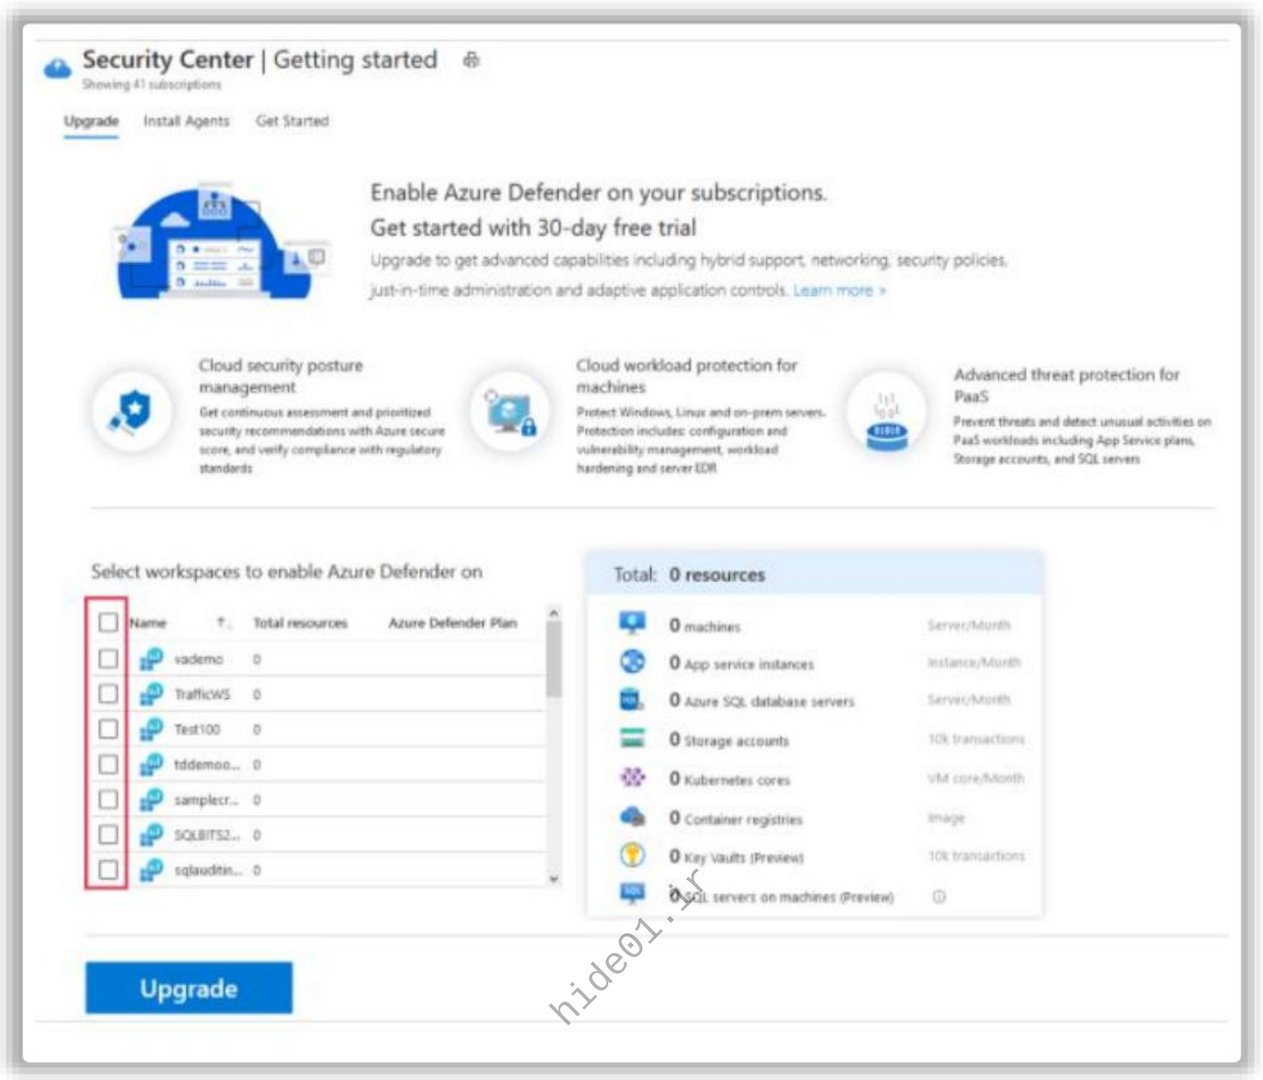


Figure: 2.113: View of Workspaces and Subscriptions that can be On-boarded

- Choose the workspaces that are to be upgraded under **Select workspace to enable Azure Defender on**.
- Click on **Upgrade** for enabling Azure Defender.

Enabling Automatic Data Collection

Log Analytics agents are installed on all supported machines by ASC when auto-provisioning is enabled. Through Log Analytics agents, data is collected from machines by ASC to monitor for threats and vulnerabilities. To enable auto-provisioning of the Log Analytics agent on machines, follow these steps:

- Choose **Pricing & settings** from the ASC's menu.
- Choose the relevant subscription.

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

- Enable auto-provisioning by setting **Auto provisioning** to **On** on the **Data collection** page and click on **Save**.

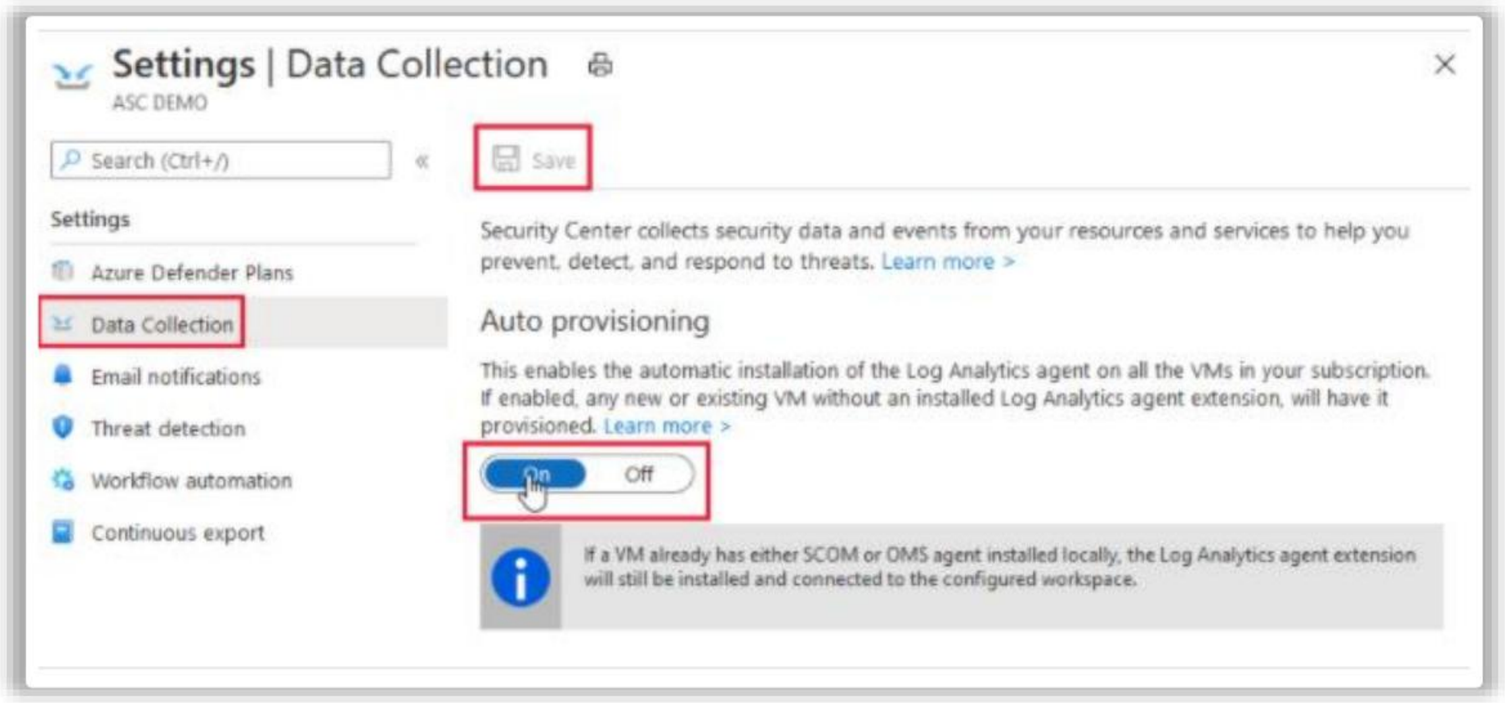


Figure: 2.114: Enabling Auto-provisioning

hide01.ir

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



Infrastructure Security Best Practices

Infrastructure best practices for Azure are listed below.

- **Comprehend the Shared Responsibility Model**: Cloud security professionals need to be aware of the responsibilities that are shared between Azure and the Azure customer (themselves). For every Azure service, the division of responsibilities differs but largely the customer is responsible for their data and controlling access to it. There may be additional responsibilities for them depending on the service they consume. Understanding the shared responsibility model is essential for customers who plan to move to the cloud. Though the CSP provides reasonable security, customers cannot rely completely on it for securing their applications, services, and users.
- **Utilize Dedicated Workstations**: Organizations often become victims of cyberattacks because their users perform activities that may compromise the network like Internet browsing. This exposes them to ransomware and malware that allows attackers to exploit systems and access the business. Microsoft has provided Privileged Access Workstations to deal with it. These workstations have a dedicated OS that offers protection against external attacks. By using separate workstations for work and every-day use, users can prevent phishing, OS vulnerabilities and application vulnerabilities.
- **The Data-At-Rest and In-Flight Should Be Encrypted**: Encrypting data is required to protect it against various cyberattacks. Azure provides Transparent Data Encryption by default in Azure Synapse Analytics and Azure SQL and other equivalent technology is enabled for databases. Encryption keys are managed by Azure but customers can use their own keys. Azure provides encryption for both data-at-rest and data-in-flight and manages the encryption keys. Customers can always use their own keys if they want. It provides TLS encryption for data-in-flight. Customers can use Azure Storage Client Library

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

for encrypting data inside client apps prior to uploading them to Azure storage and decrypting data when downloading.

- **Access To the Database Should Be Restricted**: Utilize database firewalls in Azure to protect databases against various threats. Use Azure Private Link in production where customers want their databases to be fully protected from the outside world. Azure Private Link can be used to allow access to web apps and VMs in virtual networks. Always use Azure AD accounts for database access in Azure because it allows the administrator to manage password policies, rotate passwords, authenticate identities on the database-level, and enable MFA.
- **Safeguard the Virtual Machines**: Using SSH or RDP for logging in to a VM opens ports on the VM, which makes it vulnerable to attacks so disable the SSH or RDP access to VMs. Deploy VMs in a VNet and make use of Azure Bastion to access the VMs. Secure RDP or SSH connectivity is possible by using Azure Bastion without the need of opening ports on the VMs.
- **Enforce Multiple Authentication to Secure Microsoft Azure Account**: The cloud administrator needs to enable multiple authentication to secure Azure account. This authentication is used additionally with the password. This best practice can eliminate unauthorized access followed by password theft. Azure provides MFA that can secure application and data access as well as meet the user demand. Users can choose a method they prefer for MFA from a variety of methods such as phone calls, SMS, app notifications, etc.
- **Administrator Access and User Access Should Be Restricted**: Administrators should not use privilege access while performing standard tasks. Azure provides an administration system where accounts are configured with more privileges for a limited time to allow administrators to do their tasks. The administrator can use Privileged Identity Management for analyzing, managing, and controlling access in the organization.

When Azure is used for application development and labs, the organization is exposed to internal threats. This could be due to assigning overly permissive rights to some users. To handle such problems, use Azure DevTest Labs service, which utilizes RBAC to divide responsibilities among users within a team by permitting access rights that are necessary for them to perform their tasks.

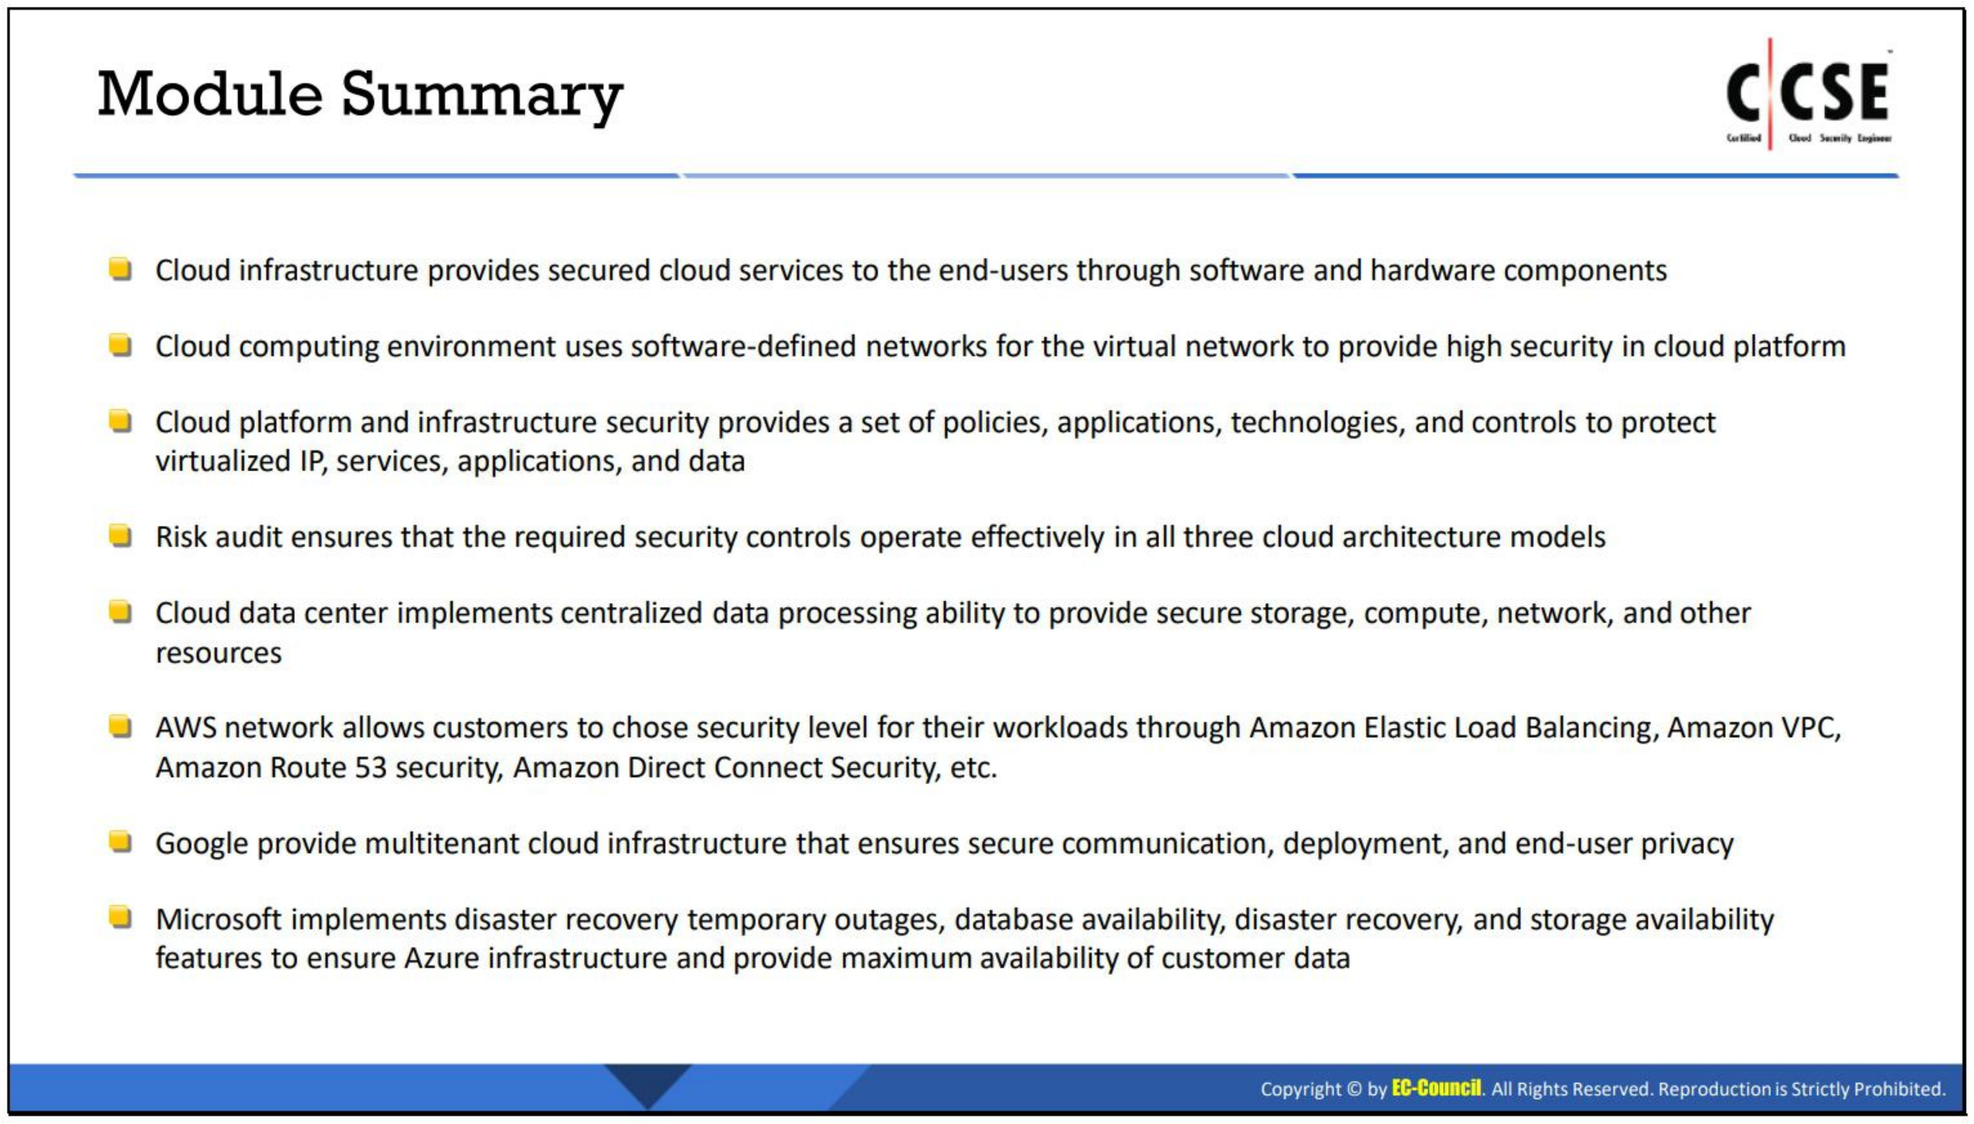
- **Use Azure AD for Securing Identity**: Securing Identity is easier with Azure AD. Azure completely relies on its AD for authenticating users. Centralizing Identity into one Authoritative source is recommended. Use Azure AD Connect to integrate cloud and on-premises directories in a hybrid identity situation to manage identities at one place. Implement two-factor authentication for all users who use Azure AD. This will restrict access to Azure resources and increase security. Two-factor authentication can be used along with SSO to provide the same level of security.
- **Limit Subscription Owners**: This best practice says that there must not be more than three subscription owners. Generally, having two product owners or administrators is good.

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

- **Network Access Should Be Stringently Controlled**: Control network access by using multiple layers of security in-between and around protected resources. Use Azure Firewall or any external virtual network solution, IDS/IPS, Web content filtering, DDoS prevention systems, and vulnerability management solutions (like antivirus, application controls, etc.) as the first layer of security. Use an NSG that filters network traffic in an Azure virtual network. NSG applied to a subnet forms the second layer of security which prevents unwanted traffic from coming or leaving the subnet and this NSG can be used to control access between subnets. Apply NSG to VM network interfaces as the third security layer to control inbound and outbound traffic from VMs. Moreover, use site-to-site VPN or ExpressRoute instead of exposing the network directly to the Internet.
- **Periodically Update the Virtual Machine**: Azure provides a simple and automated update management system to detect and apply important security updates. Use this update management system to periodically update VMs on Azure. Additionally, use Windows Defender Advanced Threat Protection and Windows anti-malware as these both can integrate well with ASC. This will help to manage VM security at one place.
- **Protect Sensitive Data**: Use Azure Key Vault to protect sensitive data such as certificates, encryption keys, and secrets which are used by cloud services and applications. There is a unique ACL in every Azure Key Vault that uses RBAC.

hide01.ir

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



Module Summary

This module described the key components of cloud platforms and infrastructure. It provided an insight into the components of cloud infrastructure security that organizations should follow to ensure security in a cloud platform. It highlighted the importance of evaluating the cloud service providers before consuming cloud services and comparing major CSPs' features. It discussed the security features provided by Amazon AWS, Microsoft Azure, and GCP to secure their respective infrastructure. It also showed how to implement the services offered by these CSPs to secure cloud infrastructure.

hid001.ir

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

This page is intentionally left blank.

hide01.ir