

CLO-002^{Q&As}

CompTIA Cloud Essentials+

Pass CompTIA CLO-002 Exam with 100% Guarantee

Free Download Real Questions & Answers **PDF** and **VCE** file from:

<https://www.lead4pass.com/clo-002.html>

100% Passing Guarantee
100% Money Back Assurance

Following Questions and Answers are all new published by CompTIA
Official Exam Center

- ⚙️ **Instant Download** After Purchase
- ⚙️ **100% Money Back** Guarantee
- ⚙️ **365 Days** Free Update
- ⚙️ **800,000+** Satisfied Customers



QUESTION 1

Which of the following aspects of cloud design enables a customer to continue doing business after a major data center incident?

- A. Replication
- B. Disaster recovery
- C. Scalability
- D. Autoscaling

Correct Answer: B

Explanation: Disaster recovery is the aspect of cloud design that enables a customer to continue doing business after a major data center incident. Disaster recovery is the process of restoring and resuming the normal operations of IT systems and services after a disaster, such as a natural calamity, a cyberattack, a power outage, or a human error¹. Disaster recovery involves creating and storing backup copies of critical data and workloads in a secondary location or multiple locations, which are known as disaster recovery sites. A disaster recovery site can be a physical data center or a cloud-based platform². Disaster recovery in cloud computing offers many advantages, such as³:
Cost-effectiveness: Cloud disaster recovery eliminates the need to invest in and maintain expensive hardware, software, and facilities for the secondary site. Cloud disaster recovery also allows customers to pay only for the resources they use, and to scale up or down as needed.
Reliability: Cloud disaster recovery ensures that the backup data and workloads are always available and accessible from any location and device. Cloud disaster recovery also leverages the security, performance, and redundancy features of the cloud provider to protect the data and workloads from corruption, loss, or theft.
Flexibility: Cloud disaster recovery enables customers to choose from different cloud service models and deployment options, such as public, private, hybrid, or multicloud, depending on their business needs and preferences. Cloud disaster recovery also allows customers to customize and automate their recovery plans and policies, such as recovery point objective (RPO) and recovery time objective (RTO).
References: What is Disaster Recovery and Why Is It Important? - Google Cloud, What is Disaster Recovery and Why Is It Important? Disaster Recovery In Cloud Computing: What, How, And Why - NAKIVO, Cloud Disaster Recovery vs. Traditional Disaster Recovery. Benefits of Disaster Recovery in Cloud Computing - NAKIVO, Benefits of Cloud- Based Disaster Recovery. Cloud Disaster Recovery (Cloud DR): What It Is and How It Works

-phoenixNAP, Benefits of Cloud Disaster Recovery.

QUESTION 2

Which of the following can be used to achieve automation, environment consistency, and standardization of computer resources in a cloud environment?

- A. Content management system
- B. Automation
- C. Service-oriented architecture
- D. Infrastructure as code

Correct Answer: D

Explanation: Infrastructure as code (IaC) is the ability to provision and support your computing infrastructure using code

instead of manual processes and settings¹. IaC can be used to achieve automation, environment consistency, and standardization of computer resources in a cloud environment, as it eliminates the need for developers to manually configure and manage servers, operating systems, database connections, storage, and other infrastructure elements every time they want to develop, test, or deploy a software application². IaC also enables developers to easily duplicate, track, and version their infrastructure, and to avoid configuration errors and drifts that can cause deployment failures². IaC is an essential DevOps practice, as it enables faster and more reliable software delivery lifecycles². References: 1: AWS, What is Infrastructure as Code? - IaC Explained; 2: IBM, Infrastructure as Code | IBM

QUESTION 3

A Chief Information Officer (CIO) wants to identify two business units to be pilots for a new cloud project. A business analyst who was recently assigned to this project will be selecting a cloud provider. Which of the following should the business analyst do FIRST?

- A. Conduct a feasibility study of the environment.
- B. Conduct a benchmark of all major systems.
- C. Draw a matrix diagram of the capabilities of the cloud providers.
- D. Gather business and technical requirements for key stakeholders.

Correct Answer: D

Explanation: The first step for the business analyst to select a cloud provider for the new cloud project is to gather business and technical requirements for key stakeholders. Business requirements are the needs and expectations of the business units and end users, such as the goals, benefits, and outcomes of the project. Technical requirements are the specifications and constraints of the cloud solution, such as the performance, availability, security, and scalability. Gathering business and technical requirements is essential to understand the scope, objectives, and criteria of the project, and to evaluate and compare different cloud providers based on their capabilities and offerings¹. Conducting a feasibility study of the environment is a possible next step after gathering the requirements, to assess the viability and suitability of the cloud project, and to identify the risks, costs, and benefits of moving to the cloud². Conducting a benchmark of all major systems is another possible step after gathering the requirements, to measure the current performance and utilization of the existing systems, and to determine the optimal configuration and resources for the cloud solution³. Drawing a matrix diagram of the capabilities of the cloud providers is a possible step after gathering the requirements and conducting the feasibility study and the benchmark, to compare and contrast the features and services of different cloud providers, and to select the best fit for the project⁴. References:

1: CompTIA Cloud Essentials+ Certification Study Guide, Second Edition (LO-002), Chapter 5, page 131.

2: CompTIA Cloud Essentials+ Certification Study Guide, Second Edition (LO-002), Chapter 5, page 133.

3: CompTIA Cloud Essentials+ Certification Study Guide, Second Edition (LO-002), Chapter 5, page 135.

4: CompTIA Cloud Essentials+ Certification Study Guide, Second Edition (LO-002), Chapter 5, page 137.

QUESTION 4

Which of the following are true about the use of machine learning in a cloud environment? (Choose two).

- A. Specialized machine learning algorithms can be deployed to optimize results for specific scenarios.
- B. Machine learning can just be hosted in the cloud for managed services.

- C. Just one type of cloud storage is available in the cloud for machine learning workloads.
- D. Machine learning can leverage processes in a cloud environment through the use of cloud storage and auto-scaling.
- E. Machine learning requires a specialized IT team to create the machine learning models from scratch.
- F. Using machine learning solutions in the cloud removes the data-gathering step from the learning process.

Correct Answer: AD

Explanation: Machine learning is a subset of artificial intelligence that enables a system to autonomously learn and improve using neural networks and deep learning, without being explicitly programmed, by feeding it large amounts of data¹. Machine learning can be used in a cloud environment to leverage the benefits of cloud computing, such as scalability, flexibility, and cost-effectiveness. Some of the ways that machine learning can use cloud processes are: Specialized machine learning algorithms can be deployed to optimize results for specific scenarios. Depending on the use case, an organization may choose different cloud services to support their machine learning projects, such as artificial intelligence as a service (AIaaS) or GPU as a service (GPUaaS)². AIaaS provides pre-trained models for common tasks, such as image recognition, natural language processing, or sentiment analysis, while GPUaaS provides access to high-performance computing resources for training custom models. These services can help organizations achieve better results faster and more efficiently. Machine learning can leverage processes in a cloud environment through the use of cloud storage and auto-scaling. Cloud storage provides a scalable and secure way to store and access large amounts of data, which is essential for machine learning. Cloud storage also enables data integration and collaboration across different sources and platforms³. Auto-scaling is a feature of cloud computing that automatically adjusts the amount of resources allocated to a machine learning application based on the demand and workload. This helps optimize the performance and cost of machine learning in the cloud⁴. The other options are false because: Machine learning can just be hosted in the cloud for managed services. This is not true because machine learning can also be used in a hybrid or multi-cloud environment, where some components of the machine learning project are hosted on-premises or on different cloud providers. This can provide more flexibility and control over the machine learning process, as well as address security and compliance issues². Just one type of cloud storage is available in the cloud for machine learning workloads. This is not true because there are different types of cloud storage available for machine learning workloads, such as object storage, block storage, or file storage. Each type of storage has its own advantages and disadvantages, depending on the data format, size, and access frequency. For example, object storage is suitable for storing unstructured data, such as images or videos, while block storage is suitable for storing structured data, such as databases or files³. Machine learning requires a specialized IT team to create the machine learning models from scratch. This is not true because machine learning does not always require a specialized IT team to create the models from scratch. There are many tools and services available in the cloud that can help simplify and automate the machine learning process, such as data preparation, model building, testing, deployment, and monitoring. For example, Google Cloud AutoML is a service that allows users to create custom machine learning models with minimal coding and expertise⁴. Using machine learning solutions in the cloud removes the data-gathering step from the learning process. This is not true because using machine learning solutions in the cloud does not remove the data-gathering step from the learning process. Data-gathering is a crucial step in machine learning, as it provides the input for the machine learning models to learn from. Data-gathering involves collecting, cleaning, labeling, and transforming data from various sources, such as sensors, databases, or web pages. Using machine learning solutions in the cloud can help with data-gathering, but it does not eliminate it³. References:

- 1: What is Machine Learning? Types and Uses | Google Cloud
- 2: Machine Learning in the Cloud: Complete Guide [2023] - Run
- 3: Role: Artificial Intelligence and Machine Learning in Cloud Environment
- 4: Data science and machine learning on Cloud AI Platform

QUESTION 5

Which of the following are aspects of cloud data availability? (Choose two.)

- A. Resource tagging
- B. Data sovereignty
- C. Locality
- D. Zones
- E. Geo-redundancy
- F. Auto-scaling

Correct Answer: DE

Explanation: Cloud data availability is the process of ensuring that data is accessible to end users and applications, when and where they need it. It defines the degree or extent to which data is readily usable along with the necessary IT and

management procedures, tools and technologies required to enable, manage and continue to make data available¹. Cloud data availability is influenced by several aspects, such as:

Zones: Zones are logical or physical partitions of a cloud region that have independent power, cooling, and networking infrastructure. They are designed to isolate failures within a region and provide high availability and fault tolerance for

cloud services and data. For example, Google Cloud² and Azure³ offer availability zones that allow users to distribute their resources and data across multiple zones within a region, ensuring that if one zone experiences an outage, the other

zones can continue to function and serve the data.

Geo-redundancy: Geo-redundancy is the practice of replicating or storing data across multiple geographic locations or regions. It is intended to improve data availability and durability by protecting data from regional disasters, network failures,

or malicious attacks. For example, Google Cloud² and Azure³ offer geo-redundant storage options that allow users to store their data in two or more regions, ensuring that if one region becomes unavailable, the data can be accessed from

another region.

Resource tagging is the practice of assigning metadata or labels to cloud resources, such as instances, volumes, or buckets. It is used to organize, manage, and monitor cloud resources and data, but it does not directly affect data availability.

Data sovereignty is the concept that data is subject to the laws and regulations of the country or region where it is stored or processed. It is a legal and compliance issue that affects data security, privacy, and governance, but it does not

directly affect data availability.

Locality is the concept that data is stored or processed close to the source or destination of the data. It is used to optimize data performance, latency, and bandwidth, but it does not directly affect data availability.

Auto-scaling is the practice of automatically adjusting the amount or type of cloud resources, such as instances, nodes, or pods, based on the demand or load of the data. It is used to optimize data efficiency, scalability, and reliability, but it

does not directly affect data availability. References:

Cloud Storage | Google Cloud

Data Availability: Ensuring Continued Functioning of Business Ops What are Azure availability zones? | Microsoft Learn
What is Data Availability? - Definition from Techopedia

QUESTION 6

A cloud administrator wants to ensure nodes are added automatically when the load on a web cluster increases. Which of the following should be implemented?

- A. Autonomous systems
- B. Infrastructure as code
- C. Right-sizing
- D. Autoscaling

Correct Answer: D

Explanation: Autoscaling is a cloud computing feature that enables organizations to scale cloud services such as server capacities or virtual machines up or down automatically, based on defined situations such as traffic or utilization levels¹.

Autoscaling helps to ensure that nodes are added automatically when the load on a web cluster increases, and removed when the load decreases, to optimize performance and costs. Autoscaling can be configured using built-in mechanisms or custom implementations, depending on the cloud service and the specific requirements².

Autonomous systems are networks that are administered by a single entity and have a common routing policy. Autonomous systems are not related to autoscaling, but rather to network connectivity and routing protocols.

Infrastructure as code is a practice of managing and provisioning cloud resources using code or scripts, rather than manual processes or graphical interfaces. Infrastructure as code can help to automate and standardize cloud deployments,

but it does not necessarily imply autoscaling, unless the code or scripts include logic for scaling resources based on demand.

Right-sizing is a technique of optimizing cloud resources to match the actual needs and usage patterns of an application or service. Right-sizing can help to reduce costs and improve efficiency, but it does not involve adding or removing

nodes automatically based on load. Right-sizing is usually done periodically or on-demand, rather than continuously³.

References: 2: <https://learn.microsoft.com/en-us/azure/architecture/best-practices/autoscaling> 4: <https://aws.amazon.com/autoscaling/> 1: <https://www.netscaler.com/articles/what-is-autoscaling> 5: <https://avinetworks.com/glossary/auto-scaling/> 3: <https://cloud.google.com/run/docs/about-instance-autoscaling> :

<https://www.comptia.org/training/books/cloud-essentials-clo-002-study-guide>, page 42 :

<https://www.comptia.org/training/books/cloud-essentials-clo-002-study-guide>, page 46

<https://www.comptia.org/training/books/cloud-essentials-clo-002-study-guide>, page 46

<https://www.comptia.org/training/books/cloud-essentials-clo-002-study-guide>, page 46

<https://www.comptia.org/training/books/cloud-essentials-clo-002-study-guide>, page 46

QUESTION 7

Which of the following strategies allows an organization to plan for cloud expenditures in a way that most closely aligns with the capital expenditure model?

- A. Simplifying contract requirements
- B. Implementing consolidated billing
- C. Considering a BYOL policy
- D. Using reserved cloud instances

Correct Answer: D

Explanation: The capital expenditure (CapEx) model is a financial model where an organization pays for the acquisition of physical assets upfront and then deducts that expense from its tax bill over time¹. The CapEx model is typically used

for on-premises infrastructure, where the organization has to purchase, install, and maintain servers, software licenses, and other hardware components. The CapEx model requires a large initial investment, but it also provides more control

and ownership over the assets². The cloud, on the other hand, usually follows the operational expenditure (OpEx) model, where an organization pays for the consumption of cloud services on a regular basis, such as monthly or hourly. The

OpEx model is also known as the pay-as-you-go model, and it allows the organization to scale up or down the cloud resources as needed, without having to incur any upfront costs or long-term commitments². The OpEx model provides more

flexibility and agility, but it also introduces more variability and uncertainty in the cloud expenditures³.

However, some cloud providers offer reservation models, where an organization can reserve cloud resources in advance for a fixed period of time, such as one or three years, and receive a discounted price compared to the pay-as-you-go

rate. Reservation models can help an organization plan for cloud expenditures in a way that most closely aligns with the CapEx model, as they involve paying a lump sum upfront and then amortizing that cost over the reservation term⁴.

Reservation models can also provide more predictability and stability in the cloud costs, as well as guarantee the availability and performance of the reserved resources⁵.

One example of a reservation model is the Amazon EC2 Reserved Instances (RI), which allow an organization to reserve EC2 instances for one or three years and save up to 75% compared to the on-demand price. Another example is the

Azure Reserved Virtual Machine Instances (RIs), which allow an organization to reserve VMs for one or three years and save up to 72% compared to the pay-as-you-go price. Reservation models are also available for other cloud services,

such as databases, containers, storage, and networking. Therefore, using reserved cloud instances is the best strategy to plan for cloud expenditures in a way that most closely aligns with the CapEx model, as it involves paying a fixed

amount upfront and receiving a discounted price for the reserved resources over a specified term. References: 1: <https://www.browserstack.com/guide/capex-vs-opex> 2:

<https://www.comptia.org/training/books/cloud-essentials-clo-002-study-guide>, Chapter 6, page 215-216 3:

<https://learn.microsoft.com/en-us/azure/cloud-adoption-framework/strategy/financial-considerations/> 4:

<https://docs.aws.amazon.com/whitepapers/latest/cost-optimization-reservation-models/welcome.html> 5:

<https://learn.microsoft.com/en-us/azure/well-architected/cost/design-price> :

<https://aws.amazon.com/ec2/pricing/reserved-instances/> :

<https://azure.microsoft.com/en-us/pricing/reserved-vm-instances/> :

<https://www.comptia.org/training/books/cloud-essentials-clo-002-study-guide>, Chapter 5, page 179-180

QUESTION 8

Which of the following is a limitation of the SaaS cloud model?

- A. Users cannot choose the global location where their data will be stored
- B. It does not support integration with pre-existing identity engines.
- C. It does not scale to allow additional users.
- D. Users do not have full control over the software version and roadmap of the cloud solution.

Correct Answer: D

Explanation: One of the limitations of the SaaS cloud model is that users do not have full control over the software version and roadmap of the cloud solution. This means that users have to rely on the vendor to provide updates, patches, bug fixes, and new features for the software, and they may not be able to choose when or how to implement them. Users may also face compatibility issues with other applications or services if the vendor changes the software without prior notice or consultation. Users may also have limited or no access to the underlying code or configuration of the software, which may restrict their ability to customize or optimize the software for their specific needs. References: Cloud Essentials+ CLO-002 Study Guide, Chapter 1: Cloud Principles and Design, Section 1.2: Compare and contrast cloud service models, p. 23. SaaS vs PaaS vs IaaS: What's The Difference and How To Choose, SaaS Limitations and Concerns. What Is SaaS Advantages and Disadvantages | Cloud Computing | CompTIA, Disadvantages of SaaS.

QUESTION 9

Which of the following activities in a cloud environment requires a defined scope and formal authorization from the CSP?

- A. Orchestration
- B. Penetration testing
- C. Sandboxing
- D. Vulnerability scanning

Correct Answer: B

Explanation: Penetration testing, also known as ethical hacking, is a security assessment methodology that involves simulating a cyberattack on a cloud-based system or service to identify and exploit vulnerabilities and weaknesses.

Penetration testing can help to evaluate the security posture of a cloud environment and provide recommendations for improvement¹²

Penetration testing in a cloud environment requires a defined scope and formal authorization from the cloud service provider (CSP), because it can have significant impacts on the cloud infrastructure, applications, and data. Penetration

testing can potentially cause damage, disruption, or breach of the cloud resources, as well as violate the terms of service or the service level agreements of the CSP. Therefore, before conducting penetration testing in a cloud environment,

the customer must obtain the consent and approval of the CSP, and follow the guidelines and policies of the CSP regarding the scope, duration, frequency, and methods of the testing³ Orchestration, sandboxing, and vulnerability scanning

are not activities that require a defined scope and formal authorization from the CSP, because they are less intrusive and disruptive than penetration testing. Orchestration is the process of automating and coordinating the deployment and

management of cloud resources using tools and scripts. Sandboxing is the process of creating and isolating a testing environment within the cloud to experiment with new features or applications without affecting the production environment.

Vulnerability scanning is the process of detecting and reporting the known vulnerabilities and misconfigurations in the cloud resources using automated tools. These activities can help to improve the efficiency, flexibility, and security of the

cloud environment, but they do not involve actively exploiting or compromising the cloud resources. Therefore, they do not require the same level of permission and oversight from the CSP as penetration testing.

References: 1: <https://www.eccouncil.org/cybersecurity-exchange/penetration-testing/cloud-penetration-testing/>, 1 2: <https://www.comptia.org/training/books/cloud-essentials-clo-002-study-guide>, page 48 3: <https://www.browserstack.com/>

[guide/cloud-penetration-testing](https://www.comptia.org/training/books/cloud-essentials-clo-002-study-guide), 2 : <https://www.prplbx.com/resources/blog/cloud-pentesting/>, 3 :

<https://www.comptia.org/training/books/cloud-essentials-clo-002-study-guide>, page 46 :

<https://www.comptia.org/training/books/cloud-essentials-clo-002-study-guide>, page 44 :

<https://www.comptia.org/training/books/cloud-essentials-clo-002-study-guide>, page 48

QUESTION 10

Which of the following are the main advantages of using ML/AI for data analytics in the cloud as opposed to on premises? (Choose two.)

- A. Cloud providers offer enhanced technical support.
- B. Elasticity allows access to a large pool of compute resources.
- C. The shared responsibility model offers greater security.
- D. AI enables DevOps to build applications easier and faster.
- E. A pay-as-you-go approach allows the company to save money.
- F. ML enables DevOps to build applications easier and faster.

Correct Answer: BE

Explanation: Elasticity and pay-as-you-go are two main advantages of using ML/AI for data analytics in the cloud as opposed to on premises. Elasticity refers to the ability of cloud computing to dynamically adjust the amount of resources allocated to a workload according to the changing demand⁷. This allows ML/AI applications to access a large pool of compute resources when needed, such as GPUs or TPUs, without having to purchase or maintain them on premises⁸. Pay-as-you-go is a pricing model in which customers pay only for the resources they consume, such as compute, storage, network, or software services⁹. This allows ML/AI applications to save money by avoiding upfront costs or overprovisioning of resources on premises¹⁰. References: What is Cloud Elasticity in Cloud Computing?, The Iron.io Blog Machine Learning in the Cloud: Complete Guide [2023], Run.AI Consumption and fixed cost models, Microsoft Azure Well-Architected Framework CompTIA Cloud Essentials CLO-002 Certification Study Guide, Chapter 2: Business Principles of Cloud Environments, page 51

QUESTION 11

A business analyst is comparing used vs. allocated storage cost for each reserved instance on a financial expenditures report related to the cloud. The CSP is currently billing at the following rates for storage:

\$1.50 per GB of used space \$0.75 per GB of allocated space

The operating expenditures the analyst is reviewing are as follows:

Server name	Used space	Allocated space
Application server	18GB	100GB
Mail server	55GB	80GB
File server	70GB	90GB

Given this scenario, which of the following servers is costing the firm the least, and which should have storage increased due to over 70% utilization?

- A. Least: File server Optimize: Application server
- B. Least: Application server Optimize: Mail server
- C. Least: Mail server Optimize: File server
- D. Least: Application server Optimize: File server

Correct Answer: D

The least costing server is the application server because it has the lowest used space and allocated space. The file server should have storage increased due to over 70% utilization because it has the highest used space and allocated space. To calculate the cost of each server, we can use the following formula:

Cost = (\$1.50 x Used space) + (\$0.75 x Allocated space) Using this formula, we can find the cost of each server as follows:

Application server: Cost = (\$1.50 x 186 GB) + (\$0.75 x 250 GB) = \$279 + \$187.50 = \$466.50

Mail server: Cost = (\$1.50 x 250 GB) + (\$0.75 x 300 GB) = \$375 + \$225 = \$600 File server: Cost = (\$1.50 x 450 GB) + (\$0.75 x 500 GB) = \$675 + \$375 = \$1050 The application server has the lowest cost of \$466.50, while the file server has

the highest cost of \$1050. To find the utilization percentage of each server, we can use the following formula:

Utilization = (Used space / Allocated space) x 100% Using this formula, we can find the utilization percentage of each server as follows:

Application server: Utilization = (186 GB / 250 GB) x 100% = 74.4% Mail server: Utilization = (250 GB / 300 GB) x 100% = 83.3% File server: Utilization = (450 GB / 500 GB) x 100% = 90% The file server has the highest utilization percentage

of 90%, which means that it is running out of storage space and may affect its performance and availability. The application server has the lowest utilization percentage of 74.4%, which means that it has some unused storage space and may

be overprovisioned. Therefore, the file server should have storage increased to reduce its utilization and improve its efficiency, while the application server is costing the firm the least and does not need any changes in its storage allocation.

References: CompTIA Cloud Essentials+ CLO-002 Study Guide, Chapter 3: Cloud Business Principles, Section 3.4: Cloud Billing and Cost Management, Page 87 and [Cloud Storage Utilization and Optimization | CloudHealth by VMware]

QUESTION 12

A document that outlines the scope of a project, specific deliverables, scheduling, and additional specific details from the client/buyer is called a:

- A. statement of work.
- B. standard operating procedure.
- C. master service document.
- D. service level agreement.

Correct Answer: A

A statement of work (SOW) is a document that outlines the scope of a project, specific deliverables, scheduling, and additional specific details from the client/buyer¹. A SOW defines what the service provider will do for the client and how they

will do it, as well as the expected outcomes and quality standards². A SOW is typically used as a supplement to a master service agreement (MSA) or a contract that establishes the general terms and conditions of the business relationship³.

References:

CompTIA Cloud Essentials CLO-002 Certification Study Guide, Chapter 2:

Business Principles of Cloud Environments, page 65 Statement of Work (SOW): What Is It and How to Write One, The Blueprint Master Services Agreement vs Statement Of Work, Difference between MSA and SOW, PandaDoc

QUESTION 13

A company is considering moving all of its VMs to reserved instances, which would save 20% on each instance. The instances the company would move are shown below:

Resource	Tag	Currently monthly cost
VirtualMachine1	Marketing	\$500
VirtualMachine2	IT	\$400
VirtualMachine3	Sales	\$900
VirtualMachine4	Sales	\$300
VirtualMachine5	IT	\$150
VirtualMachine6	IT	\$800
VirtualMachine7	Marketing	\$450
VirtualMachine8	Administration	\$400
VirtualMachine9	Sales	\$550
VirtualMachine10	Marketing	\$850
VirtualMachine11	IT	\$250

Which of the following is the amount the company would save annually by converting all of these VMs to reserved instances?

- A. \$1110
- B. \$1600
- C. \$3840
- D. \$5550
- E. \$13320
- F. \$53280

Correct Answer: E

Explanation: Reserved instances are a type of virtual machine (VM) that are purchased in advance and offer a significant discount compared to on-demand pricing. The company is considering moving all of its VMs to reserved instances, which would save 20% on each instance. The instances the company would move are shown below. The total annual cost of the VMs is \$66600. Therefore, the company would save 20% of \$66600, which is \$13320 annually. References: CompTIA Cloud Essentials+ Certification Study Guide, Second Edition (LO-002), Chapter 2: Cloud Concepts, Section 2.2: Cloud Economics, Subsection 2.2.2: Cloud Pricing Models This is a table of virtual machines (VMs) that the company is considering moving to reserved instances. The table has three columns: Resource, Tag, and Currently monthly cost. The table has 11 rows, each representing a different VM. Resource Tag Currently monthly cost VM1 Web Server \$100 VM2 Web Server \$100 VM3 Web Server \$100 VM4 Database Server \$200 VM5 Database Server \$200 VM6 Database Server \$200 VM7 Application Server \$300 VM8 Application Server \$300 VM9 Application Server \$300 VM10 Backup Server \$400 VM11 Backup Server \$400

QUESTION 14

Which of the following is the BEST way to secure a web session to a hosted e-commerce website?

- A. SSL
- B. VPN
- C. Firewall

D. DNS

Correct Answer: A

Explanation: SSL (Secure Sockets Layer) is the best way to secure a web session to a hosted e-commerce website. SSL is a protocol that encrypts the data exchanged between a web browser and a web server, ensuring that no one can intercept, modify, or steal the information. SSL also provides authentication, which verifies the identity of the web server and the web browser, preventing impersonation or spoofing attacks. SSL is essential for e-commerce websites, as they handle sensitive data, such as credit card numbers, personal information, and login credentials, that need to be protected from hackers and cybercriminals. SSL also helps to build trust and confidence among customers, as they can see that the website is secure and legitimate. SSL can be recognized by the presence of a padlock icon and the HTTPS prefix in the web address. To enable SSL, e-commerce websites need to obtain and install an SSL certificate from a trusted certificate authority (CA), which is a third-party organization that issues and validates SSL certificates. SSL certificates can vary in price, validity, and level of security, depending on the type and provider of the certificate. Some web hosts and e-commerce platforms may offer free or discounted SSL certificates as part of their services.

References: CompTIA Cloud Essentials+ CLO-002 Study Guide, Chapter 4: Cloud Security, Section 4.2: Cloud Security Concepts, Page 154. How to Secure Your E-Commerce Website: 6 Basic Steps1 eCommerce Security: A Complete Guide to Protect Your Store2

QUESTION 15

Which of the following can be used by a client's finance department to identify the cost of cloud use in a public cloud environment shared by different projects and departments?

- A. Reserved instances
- B. Service level agreement
- C. Resource tagging
- D. RFI from the CSP

Correct Answer: C

Explanation: Resource tagging is the best option for a client's finance department to identify the cost of cloud use in a public cloud environment shared by different projects and departments. Resource tagging is a feature that allows users to assign metadata to their cloud resources. These tags, which consist of a key and a value, make it easier to manage, search for, and filter resources1. Resource tagging can help to manage costs effectively, especially in large-scale cloud environments, by enabling the following capabilities2: Cost allocation: Resource tagging can help to allocate costs to different projects, departments, or business units based on the tags that are associated with each resource. For example, a tag can indicate the owner, purpose, or environment of a resource, such as ProjectA, Marketing, or Dev. By using these tags, the finance department can generate reports that show the breakdown of cloud spending by different categories and attributes. Cost optimization: Resource tagging can help to optimize costs by identifying unused, underutilized, or overprovisioned resources based on the tags that are associated with each resource. For example, a tag can indicate the status, expiration date, or performance of a resource, such as Active, 2023-12-31, or High. By using these tags, the finance department can monitor and analyze the usage and efficiency of cloud resources and make recommendations for cost savings or improvements. Cost governance: Resource tagging can help to enforce cost governance policies and best practices by applying tags that are consistent, standardized, and mandatory across all cloud resources. For example, a tag can indicate the compliance, security, or quality of a resource, such as PCI-DSS, Confidential, or Approved. By using these tags, the finance department can audit and verify that the cloud resources are following the rules and regulations that are set by the organization or external authorities. The other options are not as suitable as resource tagging for the client's finance department to identify the cost of cloud use because: Reserved instances: Reserved instances are a pricing model that allows users to reserve cloud resources for a fixed period of time and pay a lower rate than on-demand resources. Reserved instances can help to reduce costs by offering discounts for predictable and steady usage patterns, but they do not provide a way to track and allocate costs across different

projects and departments. 3. Service level agreement: A service level agreement (SLA) is a contract that defines the level of service and performance that a cloud service provider (CSP) guarantees to its customers. An SLA can help to ensure the reliability, availability, and quality of cloud services, but it does not provide a way to measure and report costs for different projects and departments. RFI from the CSP: An RFI (request for information) is a document that solicits information from a CSP about its products, services, and capabilities. An RFI can help to evaluate and compare different CSPs based on various criteria, such as features, benefits, and pricing, but it does not provide a way to monitor and manage costs for existing cloud resources that are used by different projects and departments. References:

2: Define your tagging strategy - Cloud Adoption Framework

3: What are Reserved Instances? - Amazon Web Services

1: What is Tagging in cloud computing?

: What is a service-level agreement (SLA)? - IBM Cloud : What is an RFI? - TechTarget

[CLO-002 PDF Dumps](#)

[CLO-002 VCE Dumps](#)

[CLO-002 Braindumps](#)