Question #1 *Topic 1*

Which of the following roles is responsible for creating cloud components and the testing and validation of services?

A. Cloud auditor

B. Inter-cloud provider

C. Cloud service broker

D. Cloud service developer

Suggested Answer: *D*

The cloud service developer is responsible for developing and creating cloud components and services, as well as for testing and validating services.

Community vote distribution

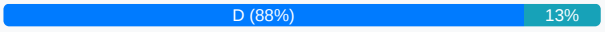

squeeze_talus0y Highly Voted 4 months, 1 week ago

Just wanted to let everyone know that I passed the exam today and no questions from the 512 available were in the actual exam. While the quality of the questions here is questionable, I can tell you for sure the ones on the actual exam are serious.

upvoted 5 times

sweetykaur Most Recent 3 days, 13 hours ago

Selected Answer: A

No questions from here on actual exam.

upvoted 1 times

gmuyir 1 week, 4 days ago

Selected Answer: D

D. Cloud service developer

A cloud service developer is responsible for creating cloud components, including services, applications, and other resources. They also conduct the testing and validation of these services to ensure they function correctly and meet the necessary requirements https://ln.run/1hwqN

upvoted 1 times

mjjjj 2 weeks, 6 days ago

Selected Answer: D

Best answer form the question but there isn't a single question provided that resembles or is actually on the test.

upvoted 1 times

nanja4 2 months, 4 weeks ago

Selected Answer: D

D. Cloud service developer is the role responsible for creating cloud components and performing testing and validation of services.

upvoted 1 times

starc5 2 months, 4 weeks ago

Selected Answer: D

D. Cloud service developer

A cloud service developer is responsible for creating cloud components, including services, applications, and other resources. They also conduct the testing and validation of these services to ensure they function correctly and meet the necessary requirements.

upvoted 1 times

Poxahex 4 months, 3 weeks ago

Selected Answer: D

d

upvoted 1 times

tuantmb 1 year, 1 month ago

Should be D - Cloud Service Developer

A - Cloud Auditor: assess the security, performance and compliance of cloud services
B - Inter-Cloud Provider: offer cloud services across multiple cloud platform
C - Cloud Service Broker: help customers find and use the best cloud services for their needs

D - Cloud Service Developer: build and maintain cloud-based applications and infrastructure --> Correct
upvoted 3 times

dev46 1 year, 1 month ago

Selected Answer: D

My explanation below but a correction - the broker is not in context with CASB
upvoted 1 times

HeatMiser 1 year, 1 month ago

se;OEO;IUH
upvoted 1 times

ndeen 1 year, 2 months ago

Guys, i did attend CCSP exam on 30-DEC 2023. Out 515, only 2 questions were there in actual exam but diff format. Dont reply on this questions. text books are the best
upvoted 2 times

74gjd_37 1 year, 5 months ago

Selected Answer: D

D. Cloud service developer is responsible for creating cloud components and testing and validating services from an ISC2 CCSP perspective.
upvoted 2 times

NJALPHA 2 years ago

D -The cloud service developer is responsible for developing and creating cloud components and services, as well as for testing and validating services.
upvoted 1 times

Pika26 2 years ago

Selected Answer: D

A cloud service developer is responsible for creating cloud components, as well as testing and validating services. This role involves designing, building, and deploying cloud-based applications and services, ensuring they meet the necessary requirements and function properly.
upvoted 1 times

luckflying 2 years ago

Selected Answer: B

CSP is responsible for cloud component development/testing. E.g. Security Center on Azure. Definitely Microsoft is responsible for it.
upvoted 1 times

akg001 2 years, 10 months ago

Selected Answer: D

D. Cloud service developer
upvoted 3 times

Question #2 *Topic 1*

What is the best source for information about securing a physical asset's BIOS?

A. Security policies

B. Manual pages

C. Vendor documentation

D. Regulations

Suggested Answer: *C*
Vendor documentation from the manufacturer of the physical hardware is the best source of best practices for securing the BIOS.

Community vote distribution
C (100%)

nestorreveron Highly Voted 4 years, 9 months ago
C Correct
upvoted 8 times

akg001 Highly Voted 2 years, 9 months ago
Selected Answer: C
C. Vendor documentation
upvoted 5 times

nanja4 Most Recent 2 months, 4 weeks ago
Selected Answer: C
The best source for securing a physical asset's BIOS is C. Vendor documentation, as it provides detailed and hardware-specific security guidance.
upvoted 1 times

tuantmb 6 months ago
C. Vendor documentation.

As NIST said:
Vendor documentation provides the most accurate and up-to-date information on how to configure, update and protect the BIOS firmware for a specific device or model, also contains instructions on how to access the BIOS settings, enable password protection, and troubleshoot common issues.
upvoted 1 times

74gjd_37 1 year, 5 months ago
Selected Answer: C
C. Vendor documentation
upvoted 1 times

NJALPHA 2 years ago
C Vendor documentation
upvoted 1 times

Pika26 2 years ago
Selected Answer: C
C is correct.
upvoted 2 times

Question #3 *Topic 1*

Which of the following is not a component of contractual PII?

A. Scope of processing

B. Value of data

C. Location of data

D. Use of subcontractors

Suggested Answer: *C*
The value of data itself has nothing to do with it being considered a part of contractual

Community vote distribution
B (100%)

estarisbourne Highly Voted 5 years, 4 months ago
Contractual PII components are as follows: (correct answer is B)
Scope of processing
Use of Subcontractors
Removal/deletion of data
Appropriate/Required data security controls
Locations of Data
Return of data/restitution of data
Audits/right to audit subcontractors
upvoted 27 times

ctux 5 years, 3 months ago
Yes, correct anser is B.
Wrong answer but correct explanation, funny... :)
upvoted 10 times

ArizonaClassics Highly Voted 5 years, 2 months ago
Correct Answer is B
upvoted 9 times

74gjd_37 Most Recent 6 months ago
Selected Answer: B
B. Value of data
The components of contractual PII include the scope of processing, location and duration of processing, the purpose and use restrictions, confidentiality and security safeguards, data ownership and return or disposal upon contract termination, sub-processing controls if any sub-processors are used, audit rights and obligations, indemnification or liability clauses in case there is a breach or non-compliance with regulations that would result in fines or litigation costs.
upvoted 1 times

tuantmb 6 months ago
The answer is B: Value of data

Contractual PII: refers to personal information that is processed, transmitted, or stored by an organization or entity as part of its business services, typically includes the following components (as csonline):
+ Scope of processing: purpose, nature and extent of the data processing, activities.
+ Location of data: the geographic regions where the data is stored, transferred, or accessed.
+ Use of subcontractors: The third parties that are involved in the data processing activities, such as cloud providers, data processors, or data controllers.
+ Data protection measures: the technical and organizational safeguards that are implemented to protect the data from unauthorized or unlawful access, use, disclosure, alteration or destruction.

The "Value of Data" is not a component of contractual PII, also may depend on the context, market or demand for data but does not affect the obligations or responsibilities of the data processors or controllers.

upvoted 2 times

coentror 8 months, 1 week ago

B of course :)

upvoted 1 times

cloudenthusiast 1 year, 1 month ago

Selected Answer: B

Value of data changes constantly

upvoted 1 times

joeee7 1 year, 8 months ago

Correct Answer is B

upvoted 1 times

FATWENTYSIX 1 year, 10 months ago

B: Official ISC2 Guide to the CCSP CBK, 3rd Edition, UNDERSTANDING PRIVACY ISSUES: Components of a Contract

upvoted 1 times

Pika26 2 years ago

Selected Answer: B

Value of data does not process personal data.

upvoted 1 times

Frito_Chip 2 years, 5 months ago

Selected Answer: B

Correct answer is B

upvoted 1 times

Eric0223 2 years, 5 months ago

i confirm it s B.

upvoted 1 times

ckho 3 years, 1 month ago

Selected Answer: B

Wrong answer but correct explanation

upvoted 2 times

nelombg 3 years, 2 months ago

Selected Answer: B

correct

upvoted 2 times

Pegasus_orb 3 years, 3 months ago

True , B

upvoted 1 times

kns20 3 years, 7 months ago

please correct the answer, it should be B

upvoted 1 times

Bancoco 3 years, 10 months ago

B is correct

upvoted 1 times

JKCY 3 years, 12 months ago

B is correct.

upvoted 1 times

Question #4 *Topic 1*

Which of the following concepts refers to a cloud customer paying only for the resources and offerings they use within a cloud environment, and only for the duration that they are consuming them?

A. Consumable service

B. Measured service

C. Billable service

D. Metered service

Suggested Answer: *B*

Measured service is where cloud services are delivered and billed in a metered way, where the cloud customer only pays for those that they actually use, and for the duration of time that they use them.

Community vote distribution

B (100%)

rjt Highly Voted 4 years, 9 months ago

NIST says "Measured service".

upvoted 14 times

vega Highly Voted 4 years, 6 months ago

measured service is the correct answer, it is about the "proper" term used in the way to define the benefits of cloud. Always remember it's a "measured service in a metered way"

upvoted 10 times

NotanAdmin Most Recent 1 month, 3 weeks ago

Selected Answer: B

Although Microsoft refers to and bills for Consumption (of cloud resources)

upvoted 1 times

sweetykaur 5 months, 2 weeks ago

This ensures customers pay based on their actual usage, often measured in time or consumption of resources. It's a fair and flexible approach to cloud billing.

upvoted 1 times

FATWENTYSIX 6 months ago

B: NIST Special Publication 800-145 says "Measured service. Cloud systems automatically control and optimize resource use by leveraging a metering capability1 at some level of abstraction appropriate to the type of service (e.g., storage, processing, bandwidth, and active user accounts). Resource usage can be monitored, controlled, and reported, providing transparency for both the provider and consumer of the utilized service."

CCSP CBK also states under the UNDERSTAND CLOUD COMPUTING CONCEPTS, Key Cloud Computing Characteristics: "Measured Service

Metering service usage allows a CSP to charge for the resources used. In a private cloud, this can allow an organization to charge each department based on their usage of the cloud. For a public cloud, it allows each customer to pay for the resources used or consumed. With a measured service, everyone pays their share of the costs."

upvoted 1 times

globy118 9 months, 3 weeks ago

Selected Answer: B

both B or D are correct

upvoted 1 times

canonigo 1 year, 11 months ago

In the Study guide, both measured and metered service are equivalent.

upvoted 1 times

nelombg 1 year, 11 months ago

Measure service is the correct answer.

upvoted 1 times

Pika26 2 years ago

Selected Answer: B

A measured service allows for better resource management and cost control, as it aligns with the pay-as-you-go approach that is a key characteristic of cloud computing services.

upvoted 1 times

DA95 2 years, 3 months ago

The correct answer is D. Metered service. In a metered service model, a cloud customer pays only for the resources and offerings they use within a cloud environment, and only for the duration that they are consuming them. This is also known as a pay-as-you-go or pay-per-use model. Consumable, measured, and billable services are not specific terms related to this type of pricing model.

upvoted 2 times

Ramnik 4 years ago

Correct answer is D. As per NIST document publish by Ajani.
Measured service. Cloud systems automatically control and optimize resource use by leveraging a metering capability1 at some level of abstraction appropriate to the type of service (e.g., storage, processing, bandwidth, and active user accounts). Resource usage can be monitored, controlled, and reported, providing transparency for both the provider and consumer of the utilized service.

upvoted 3 times

Ajani 5 years, 1 month ago

Guessing we stick with the NIST 800-145 definitions.

upvoted 2 times

ArizonaClassics 5 years, 2 months ago

Excuse me I meant both " B and D"

upvoted 3 times

ArizonaClassics 5 years, 2 months ago

Both A & D are correct

upvoted 3 times

Question #5 *Topic 1*

Which of the following roles involves testing, monitoring, and securing cloud services for an organization?

A. Cloud service integrator

B. Cloud service business manager

C. Cloud service user

D. Cloud service administrator

Suggested Answer: *D*

The cloud service administrator is responsible for testing cloud services, monitoring services, administering security for services, providing usage reports on cloud services, and addressing problem reports

Community vote distribution

D (100%)

SadioMane Highly Voted 5 years, 1 month ago

100% D

upvoted 11 times

dev46 Most Recent 6 months ago

Selected Answer: D

Which of the following roles involves testing, monitoring, and securing cloud services for an organization?

A. Cloud service integrator - doesn't sound right

B. Cloud service business manager - nope, they don't test or secure

C. Cloud service user - user consume the service, not monitor or secure

D. Cloud service administrator - sounds right

upvoted 2 times

Pika26 2 years ago

Selected Answer: D

It is D.

upvoted 1 times

bp339 2 years, 9 months ago

Selected Answer: D

Cloud service administrator

upvoted 2 times

Question #6 *Topic 1*

What is the only data format permitted with the SOAP API?

A. HTML

B. SAML

C. XSML

D. XML

Suggested Answer: *D*

The SOAP protocol only supports the XML data format.

Community vote distribution

D (100%)

akg001 Highly Voted 3 years, 2 months ago

D:- correct answer

upvoted 9 times

Pika26 Most Recent 6 months, 1 week ago

Selected Answer: D

D is correct.

upvoted 1 times

Question #7 *Topic 1*

Which data formats are most commonly used with the REST API?

A. JSON and SAML

B. XML and SAML

C. XML and JSON

D. SAML and HTML

Suggested Answer: *C*
JavaScript Object Notation (JSON) and Extensible Markup Language (XML) are the most commonly used data formats for the Representational State Transfer
(REST) API, and are typically implemented with caching for increased scalability and performance.

Community vote distribution
C (100%)

akg001 Highly Voted 3 years, 2 months ago
C: correct answer
upvoted 9 times

Pika26 Most Recent 6 months ago
Selected Answer: C
Answer is C.
upvoted 1 times

bp339 1 year, 3 months ago
Selected Answer: C
The REST API supports the following data formats: application/json. application/json indicates JavaScript Object Notation (JSON) and is used for most of the resources. application/xml indicates eXtensible Markup Language (XML) and is used for selected resources.
upvoted 2 times

Question #8 *Topic 1*

Which of the following threat types involves an application that does not validate authorization for portions of itself after the initial checks?

A. Injection

B. Missing function-level access control

C. Cross-site request forgery

D. Cross-site scripting

Suggested Answer: *B*

It is imperative that an application perform checks when each function or portion of the application is accessed, to ensure that the user is properly authorized to access it. Without continual checks each time a function is accessed, an attacker could forge requests to access portions of the application where authorization has not been granted.

Community vote distribution

B (100%)

akg001 Highly Voted 3 years, 2 months ago

Correct

upvoted 9 times

serget12 Highly Voted 1 year, 2 months ago

CSRF - This allows the attacker to force the victim's browser to generate

requests the vulnerable application thinks are legitimate requests from the victim. Don't like B as the answer, but it is the best option.

upvoted 5 times

Pika26 Most Recent 6 months ago

Selected Answer: B

B is correct.

upvoted 1 times

Question #9 *Topic 1*

Which of the following roles involves overseeing billing, purchasing, and requesting audit reports for an organization within a cloud environment?

A. Cloud service user

B. Cloud service business manager

C. Cloud service administrator

D. Cloud service integrator

Suggested Answer: *B*
The cloud service business manager is responsible for overseeing business and billing administration, purchasing cloud services, and requesting audit reports when necessary

ArizonaClassics Highly Voted 3 years, 8 months ago
B is correct https://blogs.vmware.com/cloudops/2015/12/cloud-business-management-series-part-one.html
upvoted 9 times

NJALPHA Most Recent 6 months ago
B- The cloud service business manager is responsible for overseeing business plans and customer relationships as well as processing financial transactions.
upvoted 1 times

Ramnik 2 years, 6 months ago
B is correct. Question clearly ask about "roles involves" "for an organization" within cloud env.

https://blogs.vmware.com/services-education-insights/2015/12/cloud-business-management-series-part-one.html
Under heading "Cloud Business Manager Role: Run Cloud Like a Business".
upvoted 3 times

CertMaster 3 years ago
Are we assuming that B is not with the Cloud Service Provider? The answer should be something that is not with a CSP. Could be A--only one that is clear?
upvoted 1 times

jwfrancis 3 years, 4 months ago
B is correct
upvoted 3 times

Question #10 *Topic 1*

What is the biggest concern with hosting a key management system outside of the cloud environment?

A. Confidentiality

B. Portability

C. Availability

D. Integrity

Suggested Answer: *C*

When a key management system is outside of the cloud environment hosting the application, availability is a primary concern because any access issues with the encryption keys will render the entire application unusable.

Community vote distribution

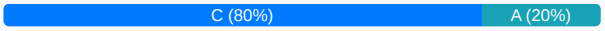

vitoscotorro Highly Voted 4 years, 7 months ago

C - should be the answer

upvoted 7 times

MaciekMT Most Recent 1 month ago

Selected Answer: C

Why Not the Others?

A. Confidentiality → External KMS solutions are typically designed to be highly secure, and proper key management practices can maintain confidentiality.

B. Portability → Moving key management between providers is difficult, but it is not the biggest concern compared to availability.

D. Integrity → The integrity of cryptographic keys is usually ensured by secure key management practices, regardless of location.

upvoted 1 times

MaciekMT 1 month ago

Selected Answer: C

The biggest concern with hosting a Key Management System (KMS) outside of the cloud environment is availability. If the external KMS becomes unreachable, cloud-based applications and services that depend on those keys for encryption and authentication could fail to function, resulting in downtime or loss of access to critical data.

 Why Availability is the Primary Concern:

If the KMS is unavailable, encrypted data cannot be accessed or decrypted.

Network dependencies between cloud services and an external KMS introduce risks of latency or outages.

Cloud services require constant access to encryption keys for ongoing operations.

upvoted 1 times

globy118 3 months, 3 weeks ago

Selected Answer: A

should be A. availability is a concern, but not the primary concern. the answer is confidentiality because external management introduces risk related to unauthorized access and exposure of sensitive keys.

upvoted 2 times

dmo_d 1 year, 4 months ago

Selected Answer: C

A is a key concern for both external/on-prem hosting and in-cloud hosting.

But C is the unique (additional) biggest concern that comes with an externally hosted key-management.

upvoted 3 times

Pika26 1 year, 4 months ago

Selected Answer: A

A: Confidentiality

upvoted 1 times

NJALPHA 1 year, 6 months ago

When a key management system is outside of the cloud environment hosting the application, availability is a primary concern because any access issues with the encryption keys will render the entire application unusable

upvoted 2 times

DA95 1 year, 9 months ago

The biggest concern with hosting a key management system outside of the cloud environment is likely confidentiality. A key management system is a type of security system that is used to securely store and manage keys, which are used to encrypt and decrypt data. If the key management system is hosted outside of the cloud environment, it may be more vulnerable to unauthorized access, which could compromise the confidentiality of the keys and the data they protect. This could also affect the integrity and availability of the data, but confidentiality is likely the most significant concern in this situation.

upvoted 1 times

Voldamort 2 years, 8 months ago

Selected Answer: C

My first thought was that it should be Confidentiality but then I thought that I could be hosting the Key Management on Premise and then that would not be the problem. I think that the correct answer is C Availability.

upvoted 2 times

Zeezee2 2 years, 10 months ago

Selected Answer: C

correct

upvoted 3 times

Ramnik 3 years, 6 months ago

C is correct.

upvoted 4 times

stevrod 3 years, 7 months ago

C - Availability is the correct answer.

upvoted 2 times

JKCY 3 years, 7 months ago

confidentiality should be the answer

upvoted 1 times

Benojojo 4 years, 4 months ago

Confidentiality looks like the closest answer to me

upvoted 2 times

echo_cert 4 years, 1 month ago

Outside the cloud env could be on premise. And that does not imply loss of confidentiality. Whereas Availability of the key is always a question when needed for any operation

upvoted 10 times

CL888 4 years ago

Exactly

upvoted 1 times

xaccan 2 years, 11 months ago

Confidentiality will be an issue when you host the key on a cloud provider, not the opposite.
The availability is the real issue.

upvoted 2 times

nelombg 2 years, 8 months ago

availability

upvoted 1 times

tngx2020 3 years, 11 months ago

risk to key confidentiality could lead to data breach, while risk to key availability leads to data loss. The question here is of the biggest concern and C should be correct.

upvoted 1 times

Question #11 *Topic 1*

Which of the following approaches would NOT be considered sufficient to meet the requirements of secure data destruction within a cloud environment?

A. Cryptographic erasure

B. Zeroing

C. Overwriting

D. Deletion

Suggested Answer: *D*

Deletion merely removes the pointers to data on a system; it does nothing to actually remove and sanitize the data. As such, the data remains in a recoverable state, and more secure methods are needed to ensure it has been destroyed and is not recoverable by another party.

Community vote distribution

ArizonaClassics Highly Voted 5 years, 2 months ago

Agreed with D

upvoted 13 times

Seke Highly Voted 3 years, 3 months ago

Selected Answer: C

Chapter 3: Data Classification (CCSP Official Study Guide)

Method 1: Physical Destruction of Media and Hardware

Any hardware or portable media containing the data in question can be destroyed by burning, melting, impact (beating, drilling, grinding, and so forth), or industrial shredding. This is the preferred method of sanitization, since the data is physically unrecoverable.

Method 2:Degaussing

This involves applying strong magnetic fields to the hardware and media where

the data resides, effectively making them blank. It does not work with solid-state drives.

Method 3:Overwriting

Multiple passes of random characters are written to the storage areas (particular disk sectors) where the data resides, with a final pass of all zeroes or ones. This can

be extremely time-consuming for large storage areas.

Method 4:Cryptoshredding (AKA Cryptographic Erasure)

This involves encrypting the data with

a strong encryption engine, and then taking the keys generated in that process, encrypting

them with a different encryption engine, and destroying the keys.

upvoted 7 times

MaciekMT Most Recent 1 month ago

Selected Answer: D

Deletion alone is NOT considered a sufficient method for secure data destruction in a cloud environment because it typically only removes file pointers rather than erasing the actual data. The data remains on the storage medium and can potentially be recovered using forensic techniques.

 Secure Data Destruction Methods:

A. Cryptographic Erasure → Deletes encryption keys, making encrypted data permanently unreadable.

B. Zeroing → Overwrites storage sectors with zeros, making data recovery difficult.

C. Overwriting → Replaces existing data with random patterns or predefined values to ensure secure erasure.

upvoted 1 times

SANKETBANG123 8 months, 3 weeks ago

Answer should be C as Zeroing is not data destruction method.

upvoted 1 times

BuckLee 1 year, 6 months ago

Selected Answer: D

D for delete

upvoted 1 times

dmo_d 1 year, 10 months ago

Selected Answer: D

D is correct.

All other options will likely have less data remanence.

upvoted 2 times

ikamalbhatt 1 year, 11 months ago

Selected Answer: D

D is correct

upvoted 1 times

Pika26 2 years ago

Selected Answer: D

D is correct.

upvoted 1 times

secisfun 2 years, 3 months ago

Selected Answer: D

Agree with D

upvoted 2 times

nigthwish 2 years, 6 months ago

The answer is D.

upvoted 2 times

ay_caramba24 2 years, 6 months ago

Selected Answer: D

D is the correct answer.

upvoted 1 times

ggx 2 years, 10 months ago

Selected Answer: D

D is the right answer

upvoted 3 times

Ramnik 4 years ago

D is correct. Data Remanence is a concern here.

upvoted 4 times

RBa001 4 years, 8 months ago

This is a negative Question. Option D is the least preferred method and Option A is the most preferred. Thus Correct answer is Option D

upvoted 4 times

guest999 4 years, 9 months ago

The only INCORRECT answer is A.

Either the question is incorrect or there should be multiple answers. Because B,C & D are correct in a Cloud environment.

upvoted 4 times

evilwizardington 4 years, 1 month ago

Nope. In cloud environments, usually you dont have ways to overwrite the information at physical level, so the best way is to use cryptographic erasure. And the worst is just trusting the Deletion features, because it does not actually delete the information; only the pointers to it.

upvoted 5 times

vitoscotorro 5 years, 1 month ago

D seems right

upvoted 5 times

Question #12 *Topic 1*

Which of the following cloud aspects complicates eDiscovery?

A. Resource pooling

B. On-demand self-service

C. Multitenancy

D. Measured service

Suggested Answer: *C*

With multitenancy, eDiscovery becomes more complicated because the data collection involves extra steps to ensure that only those customers or systems that are within scope are turned over to the requesting authority.

Community vote distribution

C (100%)

ArizonaClassics Highly Voted 3 years, 8 months ago

Agreed with C "Multi-tenancy". Explanation= E-discovery refers to a process in which an electronic data is sought, located, secured and searched with the intent of using its evidence in civil or criminal case. It can be carried out online or offline

E-Discovery Challenges>>> You receive a call from a legal advisors or from a third party advising of potential or unlawful activities across the infrastructure and resources that employees access. Given that your systems are no longer on premises but on cloud (Multi-tenancy). At this point you cannot tell where your data is hosted in the cloud or with the CSP. Hence validates the answer MULTI-TENANCY

upvoted 14 times

MaciekMT Most Recent 1 week, 3 days ago

Selected Answer: C

Multitenancy in cloud computing complicates eDiscovery because multiple customers share the same physical infrastructure. This raises legal, technical, and compliance challenges when trying to identify, collect, and preserve data for legal proceedings.

 Why Multitenancy is a Challenge for eDiscovery?

✔ Data Segregation Issues → Customer data is stored together, making it difficult to isolate specific records.

✔ Legal Jurisdiction Complications → Data may be spread across multiple geographic regions, each with different regulations.

✔ Access Control Limitations → Cloud providers manage underlying infrastructure, limiting how investigators can access data.

✔ Chain of Custody Concerns → Ensuring data integrity and preventing unauthorized access is more complex in a shared environment.

upvoted 1 times

dmo_d 4 months ago

Selected Answer: C

Yes, C it is.

upvoted 1 times

Question #13 *Topic 1*

What does the management plane typically utilize to perform administrative functions on the hypervisors that it has access to?

A. Scripts

B. RDP

C. APIs

D. XML

Suggested Answer: *C*

The functions of the management plane are typically exposed as a series of remote calls and function executions and as a set of APIs. These APIs are typically leveraged through either a client or a web portal, with the latter being the most common.

Community vote distribution

C (100%)

ArizonaClassics Highly Voted 3 years, 2 months ago

API is CORRECT!

upvoted 5 times

akg001 Most Recent 9 months, 2 weeks ago

Selected Answer: C

C. APIs

upvoted 3 times

lmpdlnar 1 year, 1 month ago

The functions of the management plane are typically exposed as a series of remote calls and function executions and as a set of APIs. These APIs are typically leveraged through either a client or a web portal, with the latter being the most common.

upvoted 3 times

Question #14 *Topic 1*

What is a serious complication an organization faces from the perspective of compliance with international operations?

A. Different certifications

B. Multiple jurisdictions

C. Different capabilities

D. Different operational procedures

Suggested Answer: *B*

When operating within a global framework, a security professional runs into a multitude of jurisdictions and requirements, and many times they might be in contention with one other or not clearly applicable. These requirements can include the location of the users and the type of data they enter into systems, the laws governing the organization that owns the application and any regulatory requirements they may have, as well as the appropriate laws and regulations for the jurisdiction housing the IT resources and where the data is actually stored, which might be multiple jurisdictions as well.

Community vote distribution

B (100%)

akg001 9 months, 2 weeks ago

Selected Answer: B

B. Multiple jurisdictions

upvoted 3 times

deegadaze1 2 years ago

B correct answer!

upvoted 2 times

Mars_oo1 2 years, 4 months ago

B is correct

upvoted 3 times

Question #15 *Topic 1*

Which networking concept in a cloud environment allows for network segregation and isolation of IP spaces?

A. PLAN

B. WAN

C. LAN

D. VLAN

Suggested Answer: *D*

A virtual area network (VLAN) allows the logical separation and isolation of networks and IP spaces to provide enhanced security and controls.

Community vote distribution

D (100%)

bp339 Highly Voted 9 months, 3 weeks ago

Selected Answer: D

VLAN is correct

upvoted 6 times

deegadaze1 Most Recent 2 years ago

D Correct

upvoted 3 times

Question #16 *Topic 1*

Which of the following standards primarily pertains to cabling designs and setups in a data center?

A. IDCA

B. BICSI

C. NFPA

D. Uptime Institute

Suggested Answer: *B*
The standards put out by Building Industry Consulting Service International (BICSI) primarily cover complex cabling designs and setups for data centers, but also include specifications on power, energy efficiency, and hot/cold aisle setups.

TimothyF Highly Voted 5 months, 2 weeks ago
BISCI is correct: The standard is focused on cabling design and setups, and also includes specifications on power, energy efficiency, and hot/cold aisle setups.
upvoted 9 times

ssurmeds Highly Voted 1 year, 4 months ago
BISCI is correct because that specifically deals with cabling standard. Uptime provides Tiers for DCs.
upvoted 6 times

Preacher Most Recent 1 year, 6 months ago
...Through our globally respected Tier Standards, and other program offerings, we've helped enterprise and vendor organizations around the globe build and maintain business-critical infrastructure to optimize performance, reliability, and efficiency. We have awarded over 1700 Tier Certifications in over 100 countries and trained thousands of professionals with our Accredited Tier Training programs.....
upvoted 2 times

evilwizardington 1 year, 1 month ago
Uptime Institute only has standards for Datacenters, not cabling.
upvoted 3 times

Chibabest 1 year, 6 months ago
What about the uptime Institute?
upvoted 1 times

evilwizardington 1 year, 1 month ago
Uptime Institute only has standards for Datacenters, not cabling.
upvoted 2 times

Question #17 *Topic 1*

Which of the following publishes the most commonly used standard for data center design in regard to tiers and topologies?

A. IDCA

B. Uptime Institute

C. NFPA

D. BICSI

Suggested Answer: *B*

The Uptime Institute publishes the most commonly used and widely known standard on data center tiers and topologies. It is based on a series of four tiers, with each progressive increase in number representing more stringent, reliable, and redundant systems for security, connectivity, fault tolerance, redundancy, and cooling.

Community vote distribution

B (100%)

akg001 9 months, 2 weeks ago

Selected Answer: B

B. Uptime Institute

upvoted 4 times

deegadaze1 2 years ago

B correct

upvoted 4 times

Question #18 *Topic 1*

What type of segregation and separation of resources is needed within a cloud environment for multitenancy purposes versus a traditional data center model?

A. Virtual

B. Security

C. Physical

D. Logical

Suggested Answer: *D*

Cloud environments lack the ability to physically separate resources like a traditional data center can. To compensate, cloud computing logical segregation concepts are employed. These include VLANs, sandboxing, and the use of virtual network devices such as firewalls.

Community vote distribution

akg001 Highly Voted 2 years, 3 months ago

Selected Answer: D

D. Logical

upvoted 5 times

globy118 Most Recent 3 months, 4 weeks ago

Selected Answer: D

while cloud environments lack the physical separation of traditional data centers, they compensate with robust logical and security segregation mechanisms12. Keep in mind that the choice between these models depends on factors like scalability, cost, and security requirements. : by chatgpt

upvoted 1 times

Kneebee 7 months, 1 week ago

D. Logical... The segregation and separation of resources in the cloud are typically achieved via techniques such as virtualization, software-defined networking, and access controls applied at the network, application, and storage layers.

upvoted 1 times

nelombg 1 year, 2 months ago

Logical is the correct answer

upvoted 1 times

ikamalbhatt 1 year, 4 months ago

Selected Answer: D

Page 162 Multitenancy refers to the concept that multiple cloud customers may share the same services, systems, networks, and other underlying resources. It is important that configurations be made in such a way as to ensure logical isolation of the various tenants, otherwise data leakage and corruption could occur.

upvoted 1 times

Pika26 1 year, 6 months ago

Selected Answer: A

A. Virtual. Within a cloud environment, virtual segregation and separation of resources are needed for multitenancy purposes, as opposed to a traditional data center model that may rely more on physical segregation. Virtual separation is achieved through various techniques, such as virtualization and software-defined networking, that allow multiple customers to share the same infrastructure while ensuring isolation of their workloads, data, and network traffic.

upvoted 1 times

deegadaze1 3 years, 4 months ago

Why not (A) -- Virtual?

upvoted 2 times

Amits3 3 years, 1 month ago

We are segregating logically by using virtualization.

upvoted 9 times

NastyNutsu 2 years, 6 months ago

Virtualization is a method while logical is the terminology used?

upvoted 4 times

Question #19 *Topic 1*

Which United States law is focused on data related to health records and privacy?

A. Safe Harbor

B. SOX

C. GLBA

D. HIPAA

Suggested Answer: *D*

The Health Insurance Portability and Accountability Act (HIPAA) requires the U.S. Federal Department of Health and Human Services to publish and enforce regulations pertaining to electronic health records and identifiers between patients, providers, and insurance companies. It is focused on the security controls and confidentiality of medical records, rather than the specific technologies used, so long as they meet the requirements of the regulations.

Community vote distribution

D (100%)

bp339 8 months, 4 weeks ago

Selected Answer: D

HIPAA is correct

upvoted 3 times

luca8818 10 months, 3 weeks ago

Selected Answer: D

D - is correct

upvoted 4 times

Question #20 *Topic 1*

What is used for local, physical access to hardware within a data center?

A. SSH

B. KVM

C. VPN

D. RDP

Suggested Answer: *B*
Local, physical access in a data center is done via KVM (keyboard, video, mouse) switches.

khindinikorse 8 months, 3 weeks ago
B is Correct: KVM
upvoted 1 times

Question #21 *Topic 1*

Within an Infrastructure as a Service model, which of the following would NOT be a measured service?

A. CPU

B. Storage

C. Number of users

D. Memory

Suggested Answer: *C*

Within IaaS, the number of users on a system is not relevant to the particular hosting model in regard to cloud resources. IaaS is focused on infrastructure needs of a system or application. Therefore, a factor such as the number of users that could affect licensing requirements, for example, would apply to the SaaS model, or in some instances to PaaS.

Community vote distribution

C (100%)

Mhiar Highly Voted 2 years ago

Correct

upvoted 5 times

Tamtout Most Recent 1 month, 2 weeks ago

Selected Answer: C

c is the correct answer

upvoted 1 times

akg001 9 months, 2 weeks ago

Selected Answer: C

C. Number of users

upvoted 4 times

Question #22 *Topic 1*

Which of the following is NOT a criterion for data within the scope of eDiscovery?

A. Possession

B. Custody

C. Control

D. Archive

Suggested Answer: *D*

eDiscovery pertains to information and data that is in the possession, control, and custody of an organization.

Community vote distribution

davidpcm Highly Voted 4 years ago

D is correct

The Electronic Discovery Reference Model (EDRM) divides the legal eDiscovery process into six stages: identification, preservation, collection, processing, review and production.

upvoted 12 times

vaolo Most Recent 3 months, 2 weeks ago

Selected Answer: D

The correct answer is: D. Archive

The criteria for data within the scope of eDiscovery include:

- Possession: Refers to the physical or legal ownership of the data.
- Custody: Indicates who has the responsibility or authority over the data.
- Control: Refers to the ability to access, manage, or direct the data.

Archive, while relevant to data management and storage, is not a core criterion for determining whether data falls within the scope of eDiscovery. It pertains to how and where data is stored, rather than its legal applicability or relevance.

upvoted 1 times

escaprix 9 months, 2 weeks ago

Selected Answer: D

D - While possession, custody, and control are criteria commonly associated with data within the scope of eDiscovery, "archive" is not typically considered a criterion for eDiscovery. Archiving refers to the long-term storage and preservation of data for historical, compliance, or reference purposes, but it does not directly relate to the possession, custody, or control of data during the eDiscovery process.

upvoted 4 times

Pika26 1 year ago

Selected Answer: D

Answer is D.

upvoted 1 times

luckflying 1 year ago

Selected Answer: B

Custody is not in scope of eDiscovery.

upvoted 1 times

Question #23 *Topic 1*

Which United States law is focused on accounting and financial practices of organizations?

A. Safe Harbor

B. GLBA

C. SOX

D. HIPAA

Suggested Answer: *C*

The Sarbanes-Oxley (SOX) Act is not an act that pertains to privacy or IT security directly, but rather regulates accounting and financial practices used by organizations. It was passed to protect stakeholders and shareholders from improper practices and errors, and it sets forth rules for compliance, regulated and enforced by the Securities and Exchange Commission (SEC). The main influence on IT systems and operations is the requirements it sets for data retention, specifically in regard to what types of records must be preserved and for how long.

Community vote distribution

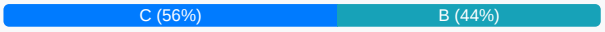

TheFivePips 9 months, 1 week ago

Selected Answer: C

Both GLBA and SOX have to do with financial practices. But C is the better answer.

B. GLBA (Gramm-Leach-Bliley Act): This U.S. federal law, also known as the Financial Services Modernization Act of 1999, focuses on protecting consumers' personal financial information held by financial institutions. It includes provisions to protect personal data, requires financial institutions to explain their information-sharing practices to their customers, and safeguards sensitive data.

C. SOX (Sarbanes-Oxley Act): This is the correct answer. The Sarbanes-Oxley Act of 2002 is a U.S. federal law that established comprehensive auditing and financial regulations for public companies. It was enacted in response to financial scandals such as Enron and WorldCom to protect investors from fraudulent accounting activities and improve the accuracy and reliability of corporate disclosures.

upvoted 3 times

globy118 10 months ago

Selected Answer: C

GLBA focuses on customer data protection within financial institutions

upvoted 1 times

Eltooth 11 months, 3 weeks ago

Selected Answer: C

C is correct answer - SOX

upvoted 1 times

FranklinG 1 year ago

financial practices of organizations? The answer is "C"

SOX protects financial information of public companies, and GLBA protects the data of financial institution customers. We are talking about organizations or public companies and not financial industry like a Bank

upvoted 1 times

rand1220 1 year ago

C is the right answer - SOX is focused on accounting practices, while GLBA is focused on information privacy & security within FSI

upvoted 1 times

ikamalbhatt 1 year, 11 months ago

Selected Answer: B

Answer is B GLBA which pertains to financial institutions, Sarbanes Oxley is for publicly listed orgs.

upvoted 4 times

Question #24 *Topic 1*

What type of masking strategy involves making a separate and distinct copy of data with masking in place?

A. Dynamic

B. Replication

C. Static

D. Duplication

Suggested Answer: *C*

With static masking, a separate and distinct copy of the data set is created with masking in place. This is typically done through a script or other process that takes a standard data set, processes it to mask the appropriate and predefined fields, and then outputs the data set as a new one with the completed masking done.

Community vote distribution

C (100%)

glenpharmd 12 months ago

Static

Static Data Masking (SDM) involves creating a sanitized version of the data where sensitive information is obscured, replaced, or scrambled. This "masked" data resides in a separate and distinct copy, ensuring that the original data remains unchanged. The masked copy can then be used in environments where sensitive information should not be exposed, such as development or testing environments.

upvoted 2 times

TraceSplice 6 months, 1 week ago

As per AWS

Static data masking is the process of applying a fixed set of masking rules to sensitive data before it's stored or shared. It's commonly used for data that does not change frequently or remains static over time. You predefine the rules and consistently apply them to the data, which ensures consistent masking across multiple environments.

upvoted 1 times

Pika26 1 year, 6 months ago

Selected Answer: C

Answer is C. Static.

upvoted 1 times

DA95 1 year, 9 months ago

The type of masking strategy that involves making a separate and distinct copy of data with masking in place is called duplication. In this approach, a copy of the data is created and the sensitive information is replaced with fake, but realistic, values. This copy of the data can then be used for testing, development, or other purposes without exposing the sensitive information.

upvoted 4 times

Mhiar 3 years, 6 months ago

correct

upvoted 4 times

Question #25 *Topic 1*

Which of the following storage types is most closely associated with a database-type storage implementation?

A. Object

B. Unstructured

C. Volume

D. Structured

Suggested Answer: *D*
Structured storage involves organized and categorized data, which most closely resembles and operates like a database system would.

Community vote distribution
D (100%)

akg001 4 months ago

Selected Answer: D
D. Structured
upvoted 4 times

Question #26 *Topic 1*

Which of the following roles is responsible for overseeing customer relationships and the processing of financial transactions?

A. Cloud service manager

B. Cloud service deployment

C. Cloud service business manager

D. Cloud service operations manager

Suggested Answer: *C*
The cloud service business manager is responsible for overseeing business plans and customer relationships as well as processing financial transactions.

Community vote distribution
C (100%)

MaciekMT 1 month ago

Selected Answer: C

The Cloud Service Business Manager is responsible for overseeing customer relationships and managing financial transactions within a cloud service environment. This role ensures that billing, pricing models, revenue management, and customer contracts are handled efficiently.

 Key Responsibilities of a Cloud Service Business Manager:

Managing customer relationships → Ensuring service quality and satisfaction.

Processing financial transactions → Handling billing, pricing, and cost management.

Negotiating SLAs (Service Level Agreements) → Ensuring customers receive agreed-upon service levels.

upvoted 1 times

Thelo26 5 months, 1 week ago

Selected Answer: C

Cloud service business manager manage business plan, manage customer relationships and financial processing

upvoted 1 times

akg001 1 year, 10 months ago

C. Cloud service business manager

upvoted 3 times

Question #27 *Topic 1*

Which protocol does the REST API depend on?

A. HTTP

B. XML

C. SAML

D. SSH

Suggested Answer: *A*
Representational State Transfer (REST) is a software architectural scheme that applies the components, connectors, and data conduits for many web applications used on the Internet. It uses and relies on the HTTP protocol and supports a variety of data formats.

evilwizardington Highly Voted 1 year, 7 months ago
A is correct. REST works over HTTP, always. But the information can be transferred in XML or JSON format.
upvoted 13 times

khos77 Most Recent 8 months, 1 week ago
HTTP and SSH are the only protons listed. XML and JSON are formats. So A is the only logical answer
upvoted 4 times

[Removed] 9 months ago
XML is not a protocol, correct answer Http (which is protocol)
upvoted 3 times

NobleGiantz 1 year, 7 months ago
A is indeed correct
upvoted 3 times

NobleGiantz 1 year, 7 months ago
REST API: XML & JSON
upvoted 1 times

Rangakarthik 1 year, 9 months ago
SOAP only works with XML formats whereas REST work with plain text, XML, HTML and JSON
upvoted 4 times

Question #28 *Topic 1*

Which United States program was designed to enable organizations to bridge the gap between privacy laws and requirements of the United States and the
European Union?

A. GLBA

B. HIPAA

C. Safe Harbor

D. SOX

Suggested Answer: *C*
Due to the lack of an adequate privacy law or protection at the federal level in the United States, European privacy regulations generally prohibit the exporting or sharing of PII from Europe with the United States. Participation in the Safe Harbor program is voluntary on behalf of an organization, but it does require them to conform to specific requirements and policies that mirror those from the EU. Thus, organizations can fulfill requirements for data sharing and export and possibly serve customers in the EU.

donny555 3 months, 2 weeks ago
Safe Harbor was nullified and replaced with Privacy Shield and now we've got the Trans atlantic data privacy framework.
upvoted 2 times

Question #29 *Topic 1*

What is the biggest benefit to leasing space in a data center versus building or maintain your own?

A. Certification

B. Costs

C. Regulation

D. Control

Suggested Answer: *B*

When leasing space in a data center, an organization can avoid the enormous startup and building costs associated with a data center, and can instead leverage economies of scale by grouping with other organizations and sharing costs.

Community vote distribution

B (100%)

click1 Highly Voted 3 years, 2 months ago

The quiz is about the benefit of the cloud, then answer should be B

upvoted 7 times

Pika26 Most Recent 4 months, 1 week ago

Selected Answer: B

B: Costs

upvoted 2 times

awscnna3 2 years, 8 months ago

B of course

upvoted 2 times

Ahbey_911 2 years, 8 months ago

Of the options, B (Costs) is the most correct.

upvoted 2 times

kjjcraigskel 2 years, 11 months ago

It's B. If control was the main benefit of leasing then you may as well host on prem. The main benefit is to save on buying and maintain that gear and licensing. You actually LOSE control not gain more by leasing.

upvoted 2 times

ichnos 3 years ago

B is correct, D is wrong

The main drawback to buying or leasing space in a data center is the lack of control over the design and features that it will have.

upvoted 4 times

guest999 3 years, 3 months ago

Per ISC2 D is the correct answer. It is Control. Yes, cost is a factor, but overall, having own DC will provide for full Control of the facility.

upvoted 1 times

evilwizardington 2 years, 8 months ago

Nope. You could have lack of Control in any. But for sure you are reducing costs.

upvoted 2 times

xaccan 2 years, 6 months ago

IN CISSP or CCSP, try always to prioritize the costs.

This is how they need you to think.

upvoted 4 times

Question #30 *Topic 1*

Which of the following security measures done at the network layer in a traditional data center are also applicable to a cloud environment?

A. Dedicated switches

B. Trust zones

C. Redundant network circuits

D. Direct connections

Suggested Answer: *B*
Trust zones can be implemented to separate systems or tiers along logical lines for great security and access controls. Each zone can then have its own security controls and monitoring based on its particular needs.

deegadaze1 Highly Voted 2 years, 6 months ago
B correct
upvoted 7 times

MaciekMT Most Recent 1 month ago
Selected Answer: B
Trust zones, also known as network segmentation or security zones, are a network security measure used in both traditional data centers and cloud environments. Trust zones logically separate different workloads, users, or services based on security policies and access controls to minimize attack surfaces and reduce lateral movement.

 Why Trust Zones Apply to Cloud Environments:

Cloud providers support Virtual Private Clouds (VPCs), Security Groups, and Subnets to implement trust zones.
Helps enforce least privilege access by isolating workloads based on security requirements.
Reduces risk of unauthorized access between environments (e.g., public vs. private zones, production vs. development).
upvoted 1 times

khindinikorse 8 months, 3 weeks ago
answer is correct: B. Trust zones
upvoted 3 times

Question #31 *Topic 1*

Which aspect of cloud computing will be most negatively impacted by vendor lock-in?

A. Elasticity

B. Reversibility

C. Interoperability

D. Portability

Suggested Answer: *D*

A cloud customer utilizing proprietary APIs or services from one cloud provider that are unlikely to be available from another cloud provider will most negatively impact portability.

Community vote distribution

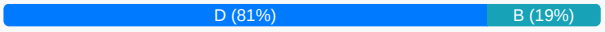

kepalon Highly Voted 3 years ago

Selected Answer: D

d is the right one

upvoted 9 times

kamin123 Highly Voted 3 years, 6 months ago

D is correct

upvoted 6 times

hacker1919 Most Recent 8 months, 2 weeks ago

B reversibility

as portably helps to move to other cloud vendor

upvoted 1 times

namtel 9 months, 2 weeks ago

B

Imagine can not move out to other

upvoted 1 times

BuckLee 1 year, 6 months ago

Selected Answer: B

B foh sho

upvoted 1 times

Lenell 2 years, 3 months ago

Selected Answer: D

To mitigate negative impacts, carefully evaluate and plan for a future need to change, try not use proprietary file formats, maintain organically internal backups that can be used for

migration, and consider a multi-cloud infrastructure.

upvoted 3 times

Eric0223 2 years, 5 months ago

Selected Answer: D

it has to be D. B is just partial

upvoted 1 times

nigthwish 2 years, 6 months ago

The answer is D. Portability encompasses moving your data. So reversibility would be wrapped up in Portability. Reversibility is a distractor.

upvoted 1 times

samsom 2 years, 9 months ago

Portability is the ability to move applications, containers, code, and associated data from one CSP to another or between legacy on-prem environments and the cloud

upvoted 2 times

tblazeen 2 years, 11 months ago

Selected Answer: B

B is the right answer - Reversibility

https://www.ibm.com/garage/method/practices/run/reversibility-in-the-cloud/

upvoted 1 times

Joadeika 2 years, 2 months ago

Reversibility isn't the same thing as portability, which is the ability to move projects or workloads from one environment to another. Reversibility includes technical and operational considerations to ensure that the new alternative state, which might include a new vendor, platform, or operating structure, can maintain your projects, workloads, or environments so that users aren't disrupted. So, this question refers to the portability part of reversibility from that link above. Reversibility is the Set while portability is the subset of it.

upvoted 1 times

[Removed] 3 years, 2 months ago

Selected Answer: B

Portability enables reversibility.

upvoted 1 times

AWSPro24 3 years, 3 months ago

I think this might be B. reversability is the ability to move between clouds.

upvoted 1 times

NastyNutsu 2 years, 11 months ago

Reversibility is the ability to move away from clouds (by removing your data from CSP).

upvoted 2 times

Question #32 *Topic 1*

Which of the following APIs are most commonly used within a cloud environment?

A. REST and SAML

B. SOAP and REST

C. REST and XML

D. XML and SAML

Suggested Answer: *B*

Simple Object Access Protocol (SOAP) and Representational State Transfer (REST) are the most commonly used APIs within a cloud environment. Extensible
Markup Language (XML) and Security Assertion Markup Language (SAML) are both standards for exchanging encoded data between two parties, with XML being for more general use and SAML focused on authentication and authorization data.

Community vote distribution

B (100%)

ikamalbhatt 4 months, 4 weeks ago

Selected Answer: B

B is correct, only APIs in the listed options.

upvoted 1 times

kepalon 1 year, 6 months ago

SOAP & REST

upvoted 4 times

Zeezee2 1 year, 10 months ago

XML and SAML are no APIs leaving B the only logical choice.

upvoted 3 times

Question #33 *Topic 1*

Which of the following attempts to establish an international standard for eDiscovery processes and best practices?

A. ISO/IEC 31000

B. ISO/IEC 27050

C. ISO/IEC 19888

D. ISO/IEC 27001

Suggested Answer: *B*
ISO/IEC 27050 strives to establish an internationally accepted standard for eDiscovery processes and best practices. It encompasses all steps of the eDiscovery process: identification, preservation, collection, processing, review, analysis, and the final production of the requested data.

bayeslife Highly Voted 1 year, 9 months ago
ISO/IEC 27001 - Information Security for Systems
ISO/IEC 31000 - Risk Management Standard
ISO/IEC 19888 - Not an iso standard.
upvoted 8 times

nelombg Most Recent 5 months, 3 weeks ago
ISO/IEC 27050 - is correct
upvoted 1 times

akg001 1 year, 4 months ago
B. ISO/IEC 27050
upvoted 3 times

Question #34 *Topic 1*

Which of the following roles is responsible for obtaining new customers and securing contracts and agreements?

A. Inter-cloud provider

B. Cloud service broker

C. Cloud auditor

D. Cloud service developer

Suggested Answer: *B*
The cloud service broker is responsible for obtaining new customers, analyzing the marketplace, and securing contracts and agreements.

akg001 Highly Voted 1 year, 10 months ago
B. Cloud service broker
upvoted 5 times

Thelo26 Most Recent 5 months, 1 week ago
Cloud Service Broker : Acquire and assess customers , assess marketplace ans set up legal agreement
upvoted 1 times

Brittle 1 year, 1 month ago
I go for B
upvoted 1 times

serget12 1 year, 8 months ago
I don't believe any of the available answers are correct, a CSB is a 3rd party entity that helps evaluate CSP and matches it's customer with the best option.
upvoted 3 times

Question #35 *Topic 1*

Which term relates to the application of scientific methods and practices to evidence?

A. Forensics

B. Methodical

C. Theoretical

D. Measured

Suggested Answer: *A*
Forensics is the application of scientific and methodical processes to identify, collect, preserve, analyze, and summarize/report digital information and evidence.

akg001 Highly Voted 10 months, 2 weeks ago

A. Forensics

upvoted 5 times

Ciuciaro Most Recent 9 months ago

Probably should be "digital evidence" in the question.

upvoted 2 times

Question #36 *Topic 1*

Which of the following roles involves the provisioning and delivery of cloud services?

A. Cloud service deployment manager

B. Cloud service business manager

C. Cloud service manager

D. Cloud service operations manager

Suggested Answer: *C*
The cloud service manager is responsible for the delivery of cloud services, the provisioning of cloud services, and the overall management of cloud services.

Community vote distribution
C (100%)

Corrector Highly Voted 2 years, 8 months ago
This URL has a good breakdown of all roles - https://cloudgal42.com/cloud-computing-activities/
upvoted 8 times

0b68b9e Most Recent 9 months ago
Selected Answer: C
the key word is "provisioning" here. Since yes a deployment manager obviously handles deployment but a CSM handles Deployment AND provision
upvoted 1 times

zuko1989 1 year, 2 months ago
Delete that first comment... The term "Cloud Service Deployment Manager" isn't formally defined in a single, universally recognized standard like those from ISO or NIST. Instead, it's a role that you might find described in industry best practices, job descriptions, or organizational charts of companies that provide or heavily use cloud services. The specific title and responsibilities can vary between organizations but generally encompass the management of deploying cloud services.
upvoted 1 times

zuko1989 1 year, 2 months ago
A.) Cloud service Manager is a general role and could encompass a variety of responsibilities related to cloud services, including aspects of deployment, operations, or business management. Without additional context, it's less specific than the Cloud Service Deployment Manager in terms of focusing solely on provisioning and delivery. The Cloud Service Deployment Manager is typically responsible for the deployment, management, and operational aspects of cloud services. This includes provisioning (allocating necessary resources), configuring (setting up services according to requirements), and delivering cloud services to end-users. They ensure that the cloud services are correctly deployed and fully operational for use.
upvoted 1 times

JKRowlings 2 years, 6 months ago
Cloud Service Manager
Responsible for policy design, business agreement, pricing models, and SLAs. This role interacts with cloud management and customers. In addition, the role will work with the cloud administrator to implement SLAs and policies
upvoted 2 times

ichnos 4 years, 6 months ago
C is correct

Cloud service manager: Delivers, provisions, and manages the cloud services
upvoted 4 times

rafnex 5 years, 3 months ago
should this be the deployment managers job?
upvoted 1 times

Ajani 5 years, 1 month ago

Deployment manager is not defined as a role at least not on the CBK

- Cloud service manager Delivers, provisions, and manages the cloud services (

CCSP All-in-One Exam Guide, Second Edition, 2nd Edition by Daniel Carter)

Cloud service manager: This person is typically responsible for policy design,
business agreement, pricing model, and some elements of the SLA (not necessarily the legal components or amendments that require contractual amendments).
This role works closely with cloud management and customers to reach agreement and alongside the cloud administrator to implement SLAs and policies on
behalf of the customers.

upvoted 12 times

Question #37 *Topic 1*

What is the primary reason that makes resolving jurisdictional conflicts complicated?

A. Different technology standards

B. Costs

C. Language barriers

D. Lack of international authority

Suggested Answer: *D*
With international operations, systems ultimately cross many jurisdictional boundaries, and many times, they conflict with each other. The major hurdle to overcome for an organization is the lack of an ultimate international authority to mediate such conflicts, with a likely result of legal efforts in each jurisdiction.

akg001 4 months, 1 week ago
D. Lack of international authority
upvoted 4 times

Question #38 *Topic 1*

GAAPs are created and maintained by which organization?

A. ISO/IEC

B. AICPA

C. PCI Council

D. ISO

Suggested Answer: *B*
The AICPA is the organization responsible for generating and maintaining what are the Generally Accepted Accounting Practices in the United States.

Community vote distribution
B (100%)

bayeslife Highly Voted 3 years, 3 months ago
AIPCA is short for American Institute of Certified Public Accountants
upvoted 10 times

TheFivePips Most Recent 9 months, 1 week ago
Generally Accepted Accounting Principles (GAAPs) are created and maintained by:

B. AICPA (American Institute of Certified Public Accountants)

Explanation:

A. ISO/IEC: The International Organization for Standardization (ISO) and the International Electrotechnical Commission (IEC) develop and publish international standards, but they are not responsible for accounting standards like GAAP.

B. AICPA (American Institute of Certified Public Accountants): The AICPA plays a significant role in the development and maintenance of accounting standards in the United States. However, it's worth noting that the Financial Accounting Standards Board (FASB) is the primary organization that establishes GAAP. The AICPA supports the FASB in this role.

C. PCI Council: The Payment Card Industry Security Standards Council (PCI SSC) is responsible for the PCI Data Security Standards (PCI DSS), which relate to payment card security, not accounting standards.

D. ISO: The International Organization for Standardization (ISO) develops international standards across various industries, but it does not create or maintain GAAP.
upvoted 2 times

z0rr02020 2 years, 11 months ago
Selected Answer: B
The answer given states "Practices" instead of "Principles". This needs to be corrected.

GAAP=Generally Accepted Accounting Principles
https://en.wikipedia.org/wiki/Generally_Accepted_Accounting_Principles_(United_States)
upvoted 3 times

Question #39 *Topic 1*

Which of the following roles is responsible for preparing systems for the cloud, administering and monitoring services, and managing inventory and assets?

A. Cloud service business manager

B. Cloud service deployment manager

C. Cloud service operations manager

D. Cloud service manager

Suggested Answer: *C*
The cloud service operations manager is responsible for preparing systems for the cloud, administering and monitoring services, providing audit data as requested or required, and managing inventory and assets.

TheFivePips 9 months, 1 week ago
every one of these types of questions reads like "which one the following roles is a kind of apple?"

a. apple
b. apples
c. green apples
d. red apple

I hate it
upvoted 2 times

khindinikorse 2 years, 2 months ago
c. Answer is Correct
upvoted 1 times

EdwardLeeBurtle 3 years, 3 months ago
What is the reference for this question?
upvoted 2 times

Question #40 *Topic 1*

Which protocol allows a system to use block-level storage as if it was a SAN, but over TCP network traffic instead?

A. SATA

B. iSCSI

C. TLS

D. SCSI

Suggested Answer: *B*
iSCSI is a protocol that allows for the transmission and use of SCSI commands and features over a TCP-based network. iSCSI allows systems to use block-level storage that looks and behaves as a SAN would with physical servers, but to leverage the TCP network within a virtualized environment and cloud.

ay_caramba24 6 months, 2 weeks ago
B is the correct answer.

https://www.techtarget.com/searchstorage/definition/iSCSI

How iSCSI works
ISCSI works by transporting block-level data between an iSCSI initiator on a server and an iSCSI target on a storage device. The iSCSI protocol encapsulates SCSI commands and assembles the data in packets for the TCP/IP layer. Packets are sent over the network using a point-to-point connection. Upon arrival, the iSCSI protocol disassembles the packets, separating the SCSI commands so the operating system (OS) will see the storage as if it was a locally connected SCSI device that can be formatted as usual.
upvoted 3 times

akg001 10 months, 2 weeks ago
B. iSCSI
upvoted 2 times

Question #41 *Topic 1*

Which of the cloud deployment models is used by popular services such as iCloud, Dropbox, and OneDrive?

A. Hybrid

B. Public

C. Private

D. Community

Suggested Answer: *B*
Popular services such as iCloud, Dropbox, and OneDrive are all publicly available and are open to any user for free, with possible add-on services offered for a cost.

akg001 4 months, 1 week ago

B. Public

upvoted 3 times

Question #42 *Topic 1*

Why does a Type 2 hypervisor typically offer less security control than a Type 1 hypervisor?

A. A Type 2 hypervisor runs on top of another operating system and is dependent on the security of the OS for its own security.

B. A Type 2 hypervisor allows users to directly perform some functions with their own access.

C. A Type 2 hypervisor is open source, so attackers can more easily find exploitable vulnerabilities with that access.

D. A Type 2 hypervisor is always exposed to the public Internet for federated identity access.

Suggested Answer: *A*

A Type 2 hypervisor differs from a Type 1 hypervisor in that it runs on top of another operating system rather than directly tied into the underlying hardware of the virtual host servers. With this type of implementation, additional security and architecture concerns come into play because the interaction between the operating system and the hypervisor becomes a critical link. The hypervisor no longer has direct interaction and control over the underlying hardware, which means that some performance will be lost due to the operating system in the middle needing its own resources, patching requirements, and operational oversight.

JKRowlings 6 months ago

Type 2 hypervisors also require a means to share folders, clipboards and other user information between the host and guest OSes. Sharing data increases the risk of hacking and spreading malicious code, so VMs demand a certain level of trust from Type 2 hypervisors.

upvoted 3 times

akg001 10 months, 2 weeks ago

A. A Type 2 hypervisor runs on top of another operating system and is dependent on the security of the OS for its own security.

upvoted 4 times

Question #43 *Topic 1*

Which is the appropriate phase of the cloud data lifecycle for determining the data's classification?

A. Create

B. Use

C. Share

D. Store

Suggested Answer: *A*
Any time data is created, modified, or imported, the classification needs to be evaluated and set from the earliest phase to ensure security is always properly maintained for the duration of its lifecycle.

MarshalLaw 4 months ago
CCSP chapter 3 - Cloud Data Lifecycle

Regardless of where data is created or resides, the Create phase should include the activities you reviewed in Chapter 2, “Data Classification,” including categorization and classification, labeling, tagging, marking, and assigning metadata.
upvoted 1 times

akg001 1 year, 4 months ago
A. Create
upvoted 3 times

Question #44 *Topic 1*

Which of the following is the optimal temperature for a data center, per the guidelines established by the America Society of Heating, Refrigeration, and Air
Conditioning Engineers (ASHRAE)?

A. 69.8-86.0degF (21-30degC)

B. 64.4-80.6degF(18-27degC)

C. 51.8-66.2degF(11-19degC)

D. 44.6-60-8degF(7-16degC)

Suggested Answer: *B*
The guidelines from ASHRAE establish 64.4-80.6degF (18-27degC) as the optimal temperature for a data center.

akg001 4 months, 1 week ago
B. 64.4-80.6degF(18-27degC)
upvoted 4 times

Question #45 *Topic 1*

Which of the following is not a risk management framework?

A. COBIT

B. Hex GBL

C. ISO 31000:2009

D. NIST SP 800-37

Suggested Answer: *B*

Hex GBL is a reference to a computer part in Terry Pratchett's fictional Discworld universe. The rest are not.

Community vote distribution

B (100%)

TheFivePips 9 months, 1 week ago

Selected Answer: B

A. COBIT: COBIT (Control Objectives for Information and Related Technologies) is a framework created by ISACA for IT management and governance. It provides a comprehensive framework that assists enterprises in achieving their objectives for the governance and management of enterprise IT. It includes aspects of risk management.

B. Hex is an elaborate, Heath Robinson/Rube Goldberg-esque, magic-powered computer housed at Unseen University (UU) in the city of Ankh-Morpork. The main structure works through the movements of large numbers of ants through the complex pipes and tubing which make up the main quantity of Hex's infrastructure.

Hex is a computer unlike any other the Disc has ever seen (which is not particularly hard since all other 'computers' on the Disc consist of druidic stone circles). Hex runs and evolves under the watchful eyes of wizard Ponder Stibbons, who becomes the de-facto IT manager at UU because he's the only one who understands what he's talking about.

C. ISO 31000:2009: ISO 31000:2009 is an international standard for risk management.

D. NIST SP 800-37: NIST Special Publication 800-37 provides guidelines for a risk management framework used by federal agencies in the United States.

upvoted 3 times

akg001 2 years, 10 months ago

Selected Answer: B

B. Hex GBL

upvoted 4 times

CISSPCCSPMP 3 years, 7 months ago

duplicate question

upvoted 1 times

VSAN888 2 years, 5 months ago

doesnt matter

upvoted 3 times

Question #46 *Topic 1*

Which of the following threat types involves the sending of untrusted data to a user's browser to be executed with their own credentials and access?

A. Missing function level access control

B. Cross-site scripting

C. Cross-site request forgery

D. Injection

Suggested Answer: *B*

Cross-site scripting (XSS) is an attack where a malicious actor is able to send untrusted data to a user's browser without going through any validation or sanitization processes, or where the code is not properly escaped from processing by the browser. The code is then executed on the user's browser with the user's own access and permissions, allowing an attacker to redirect their web traffic, steal data from their session, or potentially access information on the user's own computer that their browser has the ability to access.

Community vote distribution

Fpaul Highly Voted 3 years, 9 months ago

This is 100% cross site request forgery.... the basic way it works is someone sends you a link to execute commands on a website you are already authenticated with (hence it runs with users credentials being the key)

upvoted 17 times

xroxro 2 years, 7 months ago

see my comment below

I thing that "sending of untrusted data to a user's browser" correspond more to a stored XSS than a CSRF

In CSRF you do not directly send forged data to the browser

upvoted 6 times

Zeezee2 Highly Voted 3 years, 4 months ago

Certainly B.

XSS means user loads a site he/she trusts which has an injected script of attacker and that user's browser executes that attacker's script with the authentication of that user because he/she may be logged in to the site at that time.

With CSRF, an attacker tricks a user's brower into issuing requests which are first sent by the attacker to the user, and the website executes the attacker's requests. Here, the website trusts the user not the other way around. There is no script from the attacker that is stored on the website in this case.

Both attacks relate to lack of authenticating/validating user input.

upvoted 6 times

TheFivePips Most Recent 9 months, 1 week ago

Selected Answer: B

B. Cross-site scripting (XSS): XSS vulnerabilities occur when an application includes untrusted data on a web page without proper validation or escaping, allowing attackers to execute malicious scripts in the user's browser. These scripts run in the context of the user's session, potentially using their credentials and access.

C. Cross-site request forgery (CSRF): CSRF attacks involve tricking a user into making unwanted actions on a web application where they are authenticated. It doesn't involve injecting and executing untrusted scripts in the user's browser.

upvoted 2 times

8a21350 1 year ago

CSRF is an attack that tricks the victim into submitting a malicious request. It inherits the identity and privileges of the victim to perform an undesired function on the victim's behalf (though note that this is not true of login CSRF, a special form of the attack described below). For most sites, browser requests automatically include any credentials associated with the site, such as the user's session cookie, IP address, Windows domain credentials, and so forth. Therefore, if the user is currently authenticated to the site, the

site will have no way to distinguish between the forged request sent by the victim and a legitimate request sent by the victim.
Correct Answer is C
upvoted 1 times

cloudenthusiast 1 year ago
Selected Answer: C
key word is trig the user
upvoted 1 times

TheFivePips 9 months, 1 week ago
thats never used in the question....
upvoted 1 times

nikhilborle 1 year, 5 months ago
Selected Answer: C
The answer is C, CSRF.
https://owasp.org/www-community/attacks/csrf
upvoted 1 times

Squidly888 1 year, 6 months ago
good discussion here. My first thought was XSS but you may have convinced me that it should be CSRF. I hope I don't have that question on my test tomorrow.
upvoted 1 times

bu3oof 1 year, 6 months ago
Answer is C (CSRF) . Just have a look on any digram of CSRF, you will notice the hacker will send a users with phishing link where the user will enter his credential.
upvoted 1 times

Krishna2637 1 year, 7 months ago
Selected Answer: C
untrusted data is the key, Forgery is the one I pick.
upvoted 1 times

nelombg 1 year, 8 months ago
Cross site forgery oils the answer
upvoted 1 times

SamDavid 1 year, 8 months ago
Selected Answer: C
CSRF is the coorect answer
upvoted 1 times

earlyDev 1 year, 9 months ago
B.Cross-Site Request Forgery (CSRF) is a type of attack that tricks the victim into submitting a malicious request. It infiltrates a victim's browser and then forces it to send an HTTP request to a target site to which the victim is already authenticated.

The crucial difference is that with CSRF, the malicious request is sent to the site with the victim's credentials, meaning it's the site that's fooled into thinking the request is legitimate. With Cross-Site Scripting (XSS), malicious scripts are executed in the user's browser, not the server.
upvoted 2 times

Joe09 1 year, 10 months ago
Selected Answer: B
B definitaley
upvoted 1 times

ikamalbhatt 1 year, 11 months ago
Selected Answer: C
C definitely
upvoted 1 times

nachoqueen 2 years, 4 months ago
Selected Answer: B

B. XSS

CSRF attacks require the authenticated user to be in an active session, while the XSS attack does not. In an XSS attack, payloads can be stored and delivered whenever the user logs in.

upvoted 4 times

quagga 2 years, 5 months ago

Selected Answer: B

B: XSS

upvoted 2 times

serget12 2 years, 5 months ago

Answer is B, XSS

Cross-Site Request Forgery (CSRF): A CSRF attack forces a logged-on vic□tim's browser to send a forged HTTP request, including the victim's session cookie

and any other automatically included authentication information, to a vulnerable

web application.

upvoted 3 times

Question #47 *Topic 1*

How is an object stored within an object storage system?

A. Key value

B. Database

C. LDAP

D. Tree structure

Suggested Answer: *A*

Object storage uses a flat structure with key values to store and access objects.

Community vote distribution

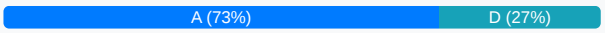

DA95 Highly Voted 2 years, 3 months ago

An object is typically stored in an object storage system as a collection of data associated with a unique identifier, known as the "key". The key is used to retrieve the object from the system. The object itself can be stored in a variety of ways, such as in a flat file system, a database, or a tree structure. The specific method used to store the object will depend on the particular object storage system being used.

upvoted 7 times

MaciekMT Most Recent 1 month ago

Selected Answer: A

In an object storage system, data is stored in a key-value structure, where each object consists of:

A unique key (identifier)

The actual data (value)

Metadata describing the object

This structure allows scalability, durability, and efficient retrieval of unstructured data in cloud storage solutions like AWS S3, Google Cloud Storage, and Azure Blob Storage.

Why Not the Others?

B. Database → Object storage does not use relational databases; it is optimized for scalable unstructured data storage.

C. LDAP (Lightweight Directory Access Protocol) → Used for directory services and authentication, not object storage.

D. Tree structure → Traditional file systems use a hierarchical tree structure, but object storage stores objects in a flat namespace with unique keys.

upvoted 1 times

el3ctronick 7 months ago

Selected Answer: A

key value, any other answer is terribly wrong.

upvoted 1 times

TheFivePips 9 months, 1 week ago

Selected Answer: A

A. Key value: In an object storage system, each object is stored with a unique identifier, known as a key. This key-value pairing allows for the storage and retrieval of objects based on their unique key.

D. Tree structure: A tree structure is typically used in hierarchical file systems and directories, not in object storage systems. Object storage is flat, meaning there isn't a hierarchy of folders and subfolders.

In the context of object storage systems, objects themselves are typically stored in a flat namespace without inherent hierarchical organization. Here's why:

Unlike traditional file systems that use a hierarchical directory structure (tree structure), object storage systems do not impose directories or folders on objects.

Are designed to scale horizontally across multiple storage nodes. A flat namespace simplifies the management and distribution of objects across these nodes, allowing for efficient scaling and access.

Objects in object storage systems often include metadata (attributes or tags) that describe the object's characteristics or usage. Metadata provides flexibility and additional context without the need for hierarchical organization.
upvoted 1 times

rodlee 1 year, 7 months ago
FALSEEEEE its D and bnot A
upvoted 1 times

nelombg 1 year, 8 months ago
Key value is the correct answer
upvoted 1 times

nelombg 1 year, 10 months ago
Key-value
upvoted 1 times

Pika26 1 year, 10 months ago
Selected Answer: A
A. Key value
upvoted 1 times

nachoqueen 2 years, 5 months ago
Selected Answer: A
"In the case of object storage, data is saved as objects by associating metadata with it. This “Key + Value” (metadata and object) format is compatible with modern data formats and retrieval needs."
Reference: https://www.ridge.co/blog/what-is-object-storage/
upvoted 3 times

ggx 2 years, 6 months ago
Selected Answer: D
Databases use key-value store.
upvoted 3 times

akg001 2 years, 9 months ago
Selected Answer: A
A. Key value
upvoted 2 times

tomres 3 years, 3 months ago
Object storage - data is stored as objects, arranged in a hierarchical structure, like a file tree, so D is correct
upvoted 2 times

kns20 3 years, 7 months ago
https://www.redhat.com/en/topics/data-storage/file-block-object-storage
upvoted 1 times

kns20 3 years, 7 months ago
this should be D, object storage typically stores in hierarchical , not key-value which is typical of database
upvoted 2 times

Warriors 3 years, 5 months ago
Object storage is flat not hierarchical. key-value is correct
upvoted 1 times

Bobbiy 3 years, 3 months ago
The link you shared says "Object storage, also known as object-based storage, is a flat structure in which files are broken into pieces and spread out among hardware." so key-value is correct
upvoted 3 times

Question #48 *Topic 1*

Which of the following is NOT a regulatory system from the United States federal government?

A. PCI DSS

B. FISMA

C. SOX

D. HIPAA

Suggested Answer: *A*
The payment card industry data security standard (PCI DSS) pertains to organizations that handle credit card transactions and is an industry regulatory standard, not a governmental one.

kepalon 6 months, 1 week ago
Agree A - this is a card industry regulatory
upvoted 2 times

Question #49 *Topic 1*

Which jurisdiction lacks specific and comprehensive privacy laws at a national or top level of legal authority?

A. European Union

B. Germany

C. Russia

D. United States

Suggested Answer: *D*
The United States lacks a single comprehensive law at the federal level addressing data security and privacy, but there are multiple federal laws that deal with different industries.

akg001 4 months, 1 week ago

D. United States

upvoted 4 times

Question #50 *Topic 1*

Which United States law is focused on PII as it relates to the financial industry?

A. HIPAA

B. SOX

C. Safe Harbor

D. GLBA

Suggested Answer: *D*
The GLBA, as it is commonly called based on the lead sponsors and authors of the act, is officially known as "The Financial Modernization Act of 1999." It is specifically focused on PII as it relates to financial institutions. There are three specific components of it, covering various areas and use, on top of a general requirement that all financial institutions must provide all users and customers with a written copy of their privacy policies and practices, including with whom and for what reasons their information may be shared with other entities.

akg001 4 months, 1 week ago
D. GLBA
upvoted 3 times

Question #51 *Topic 1*

Which of the following threat types can occur when encryption is not properly applied or insecure transport mechanisms are used?

A. Security misconfiguration

B. Insecure direct object references

C. Sensitive data exposure

D. Unvalidated redirects and forwards

Suggested Answer: *C*
Sensitive data exposure occurs when information is not properly secured through encryption and secure transport mechanisms; it can quickly become an easy and broad method for attackers to compromise information. Web applications must enforce strong encryption and security controls on the application side, but secure methods of communications with browsers or other clients used to access the information are also required. Security misconfiguration occurs when applications and systems are not properly configured for security, often a result of misapplied or inadequate baselines. Insecure direct object references occur when code references aspects of the infrastructure, especially internal or private systems, and an attacker can use that knowledge to glean more information about the infrastructure. Unvalidated redirects and forwards occur when an application has functions to forward users to other sites, and these functions are not properly secured to validate the data and redirect requests, thus allowing spoofing for malware or phishing attacks.

akg001 4 months, 1 week ago
C. Sensitive data exposure
upvoted 2 times

Question #52 *Topic 1*

What is the best approach for dealing with services or utilities that are installed on a system but not needed to perform their desired function?

A. Remove

B. Monitor

C. Disable

D. Stop

Suggested Answer: *A*

The best practice is to totally remove any unneeded services and utilities on a system to prevent any chance of compromise or use. If they are just disabled, it is possible for them to be inadvertently started again at any point, or another exploit could be used to start them again. Removing also negates the need to patch and maintain them going forward.

Community vote distribution

C (100%)

evilwizardington Highly Voted 4 years, 1 month ago

Remember that all questions are looking for the BEST option. Disable is possible, but removing the unused services is the best.

upvoted 16 times

saucehozz Highly Voted 4 years, 4 months ago

I agree with the answer of removing services and utilities.

-Think of all services CSPs offer, if you've provisioned services that aren't needed then they are an attack vector; e.g., S3 buckets (*ding ding* (remove, restrict, encrypt, etc)), public API Gateways, any other publicly exposed service, etc.

-Think of Linux and other OS' that allow you to install and remove services and utilities; e.g., SMB, NFS, iSCSI, NTP, CUPS, DNS, LDAP... etc.

-Think of rampant or forgotten instances themselves that run services that are not longer needed; decommission and remove

upvoted 11 times

MaciekMT Most Recent 1 month ago

Selected Answer: A

The best approach for dealing with unnecessary services or utilities on a system is to remove them completely. This reduces the attack surface, eliminates potential security vulnerabilities, and minimizes resource consumption.

 Why "Remove" is the Best Option?

Eliminates potential exploits → Attackers can't target what isn't there.

Frees up system resources → Improves performance and efficiency.

Simplifies maintenance → Reduces the need for patching and monitoring unused services.

upvoted 1 times

MartinRB 1 month, 2 weeks ago

Selected Answer: A

Remove is better than disable, disabled can be enabled again by malicious actor

upvoted 1 times

ra1paul 1 month, 3 weeks ago

Selected Answer: A

A because it reduces Attack surface, eliminates the maintanenance headache and also some frameworks such as ISO 27001,CIS benchmarks recommends removing the unused services or ulities. If removal is not possible then disabling is the next best option.

upvoted 1 times

sweetykaur 5 months, 2 weeks ago

Remove. The best approach is to remove unnecessary services or utilities to reduce potential attack surfaces.

upvoted 1 times

globy118 9 months, 3 weeks ago

Selected Answer: C

Disabling unnecessary services reduces resource usage and minimizes potential security risks.

upvoted 1 times

Pika26 1 year, 10 months ago

Selected Answer: C

C. Disable

upvoted 1 times

ikamalbhatt 1 year, 11 months ago

Selected Answer: C

Unneeded services are disabled not removed.

upvoted 1 times

Chibabest 4 years, 6 months ago

-All guest accounts are removed

-All unnecessary services are disabled(Correct answer)

upvoted 3 times

gouhaha 4 years, 6 months ago

Answer should be disabled service not remove as system service we cannot remove

upvoted 4 times

nelombg 3 years, 8 months ago

Remove- See above for other explanations.

upvoted 1 times

Wumza 2 years, 5 months ago

All non essential services should be stopped and set to disabled to ensure that they do not run. Al non essential software should be removed from the system. So key point here is services. So disable is the best (note you must disable after stopping the service as stopping alone cannot prevent from running again.)

upvoted 1 times

Question #53 *Topic 1*

Which of the following actions will NOT make data part of the "create" phase of the cloud data lifecycle?

A. Modifying metadata

B. Importing data

C. Modifying data

D. Constructing new data

Suggested Answer: *A*

Although the initial phase is called "create," it can also refer to modification. In essence, any time data is considered "new," it is in the create phase. This can come from data that is newly created, data that is imported into a system and is new to that system, or data that is already present and modified into a new form or value. Modifying the metadata does not change the actual data.

Community vote distribution

lolanczos 3 months, 1 week ago

Selected Answer: A

It's A. Metadata refers to data about data, such as file size, creation date, and author information. Modifying metadata involves changing these descriptive attributes without creating new data. According to the (ISC)² CCSP Official Study Guide, the Create phase is focused on the generation or importation of new data, not on altering existing descriptive information.

upvoted 1 times

sweetykaur 5 months, 2 weeks ago

You

Which of the following actions will NOT make data part of the "create" phase of the cloud data lifecycle?

A. Modifying metadata

B. Importing data

C. Modifying data

D. Constructing new data

Copilot

Modifying metadata. Modifying metadata doesn't create new data; it only changes the information about existing data.

upvoted 1 times

zxccvbnm 1 year, 11 months ago

Selected Answer: C

per Chatgpt

upvoted 1 times

PedroAsani 5 days, 14 hours ago

I'm inclined to ignore hallucinating robots

upvoted 1 times

kylesam2017 2 years, 2 months ago

CREATE: Data is either created from scratch, generated, inputted, or modified into a new form and value. Source: CCSP All In One Exam Guide, 2nd Edition, page 76. Based on this definition, I think the answer should be 'A' modification of metadata because I think, importing data may be covered under data generation/data input clauses of the aforementioned definition. Data modification is certainly covered in the definition so that is not the correct answer.

upvoted 2 times

Joadeika 2 years, 2 months ago

Selected Answer: C

Metadata are formed as part of Create state. Modifying seemed to the best option which occurs after creation

upvoted 1 times

xroxro 2 years, 7 months ago

Create phase includes assigning metadata...
The key is the "data part" in the question, right ?
upvoted 2 times

akg001 2 years, 10 months ago
A. Modifying metadata
upvoted 2 times

Question #54 *Topic 1*

What are the two protocols that TLS uses?

A. Handshake and record

B. Transport and initiate

C. Handshake and transport

D. Record and transmit

Suggested Answer: *A*
TLS uses the handshake protocol to establish and negotiate the TLS connection, and it uses the record protocol for the secure transmission of data.

TheFivePips 9 months, 1 week ago
Handshake: The TLS handshake protocol allows the client and server to authenticate each other, negotiate encryption algorithms, and establish session keys before any data is transmitted.

Record: The TLS record protocol is responsible for encapsulating higher-level protocol data (such as HTTP, FTP, etc.) into TLS records, encrypting them, and then transmitting them securely over the network.
upvoted 1 times

akg001 2 years, 10 months ago
A. Handshake and record
upvoted 1 times

Question #55 *Topic 1*

Which type of cloud model typically presents the most challenges to a cloud customer during the "destroy" phase of the cloud data lifecycle?

A. IaaS

B. DaaS

C. SaaS

D. PaaS

Suggested Answer: *C*
With many SaaS implementations, data is not isolated to a particular customer but rather is part of the overall application. When it comes to data destruction, a particular challenge is ensuring that all of a customer's data is completely destroyed while not impacting the data of other customers.

APTLYTiC 1 year ago
The question should say cloud service provider for SAAS to be correct. It currently says cloud customer, why would a cloud customer data pose a problem to a customer that already has access to if even if they opt to destroy it?
upvoted 3 times

TheFivePips 9 months, 1 week ago
However, in the context of the Certified Cloud Security Professional (CCSP) certification, the "destroy" phase of the cloud data lifecycle refers to securely deleting data and ensuring that it is irrecoverable once it is no longer needed.

In a SaaS model, customers use applications that are hosted and managed by the cloud service provider. The challenge during the "destroy" phase for SaaS customers is ensuring that their data within the application is properly deleted when it is no longer needed. This includes ensuring that all copies and backups of the data are securely erased to prevent unauthorized access or data leakage.

While IaaS customers have direct control over virtualized infrastructure resources, including data storage and virtual machines, securely deleting data in IaaS can be more straightforward compared to SaaS because customers have direct access to the underlying infrastructure and can manage data deletion processes more directly.
upvoted 2 times

akg001 2 years, 10 months ago
C. SaaS
upvoted 3 times

Question #56 *Topic 1*

Which of the following may unilaterally deem a cloud hosting model inappropriate for a system or application?

A. Multitenancy

B. Certification

C. Regulation

D. Virtualization

Suggested Answer: *C*
Some regulations may require specific security controls or certifications be used for hosting certain types of data or functions, and in some circumstances they may be requirements that are unable to be met by any cloud provider.

MaciekMT 1 month ago

Selected Answer: C

Regulation can unilaterally deem a cloud hosting model inappropriate for a system or application due to legal, compliance, or industry-specific requirements. Some regulations impose strict data residency, security, or sovereignty rules that restrict certain cloud deployments.

 Examples of Regulatory Restrictions on Cloud Hosting:

GDPR (General Data Protection Regulation - EU) → Requires personal data to be stored within the EU or in countries with equivalent protections.

HIPAA (Health Insurance Portability and Accountability Act - U.S.) → Enforces specific security and privacy requirements for healthcare data.

ITAR (International Traffic in Arms Regulations - U.S.) → Prohibits storing defense-related data in non-U.S. cloud environments.

upvoted 1 times

akg001 4 months, 1 week ago

C. Regulation

upvoted 3 times

Question #57 *Topic 1*

Which of the following is considered an internal redundancy for a data center?

A. Power distribution units

B. Network circuits

C. Power substations

D. Generators

Suggested Answer: *A*

Power distribution units are internal to a data center and supply power to internal components such as racks, appliances, and cooling systems. As such, they are considered an internal redundancy.

Community vote distribution

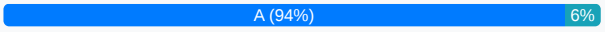

pwxfuchxaubgkfcqay Highly Voted 2 years, 9 months ago

Selected Answer: A

Generators are considered an external redundancy to a data center. Power distribution units (PDUs) are internal to a data center, and as such they are considered internal redundancies.

upvoted 14 times

certifiedgeek Highly Voted 2 years, 10 months ago

As the word "internal" is assumed to be defined as inside the data center, PDUs are located internal in the DC which are recommended to be redundantly sourced (direct powered and UPS powered).

upvoted 5 times

MaciekMT Most Recent 4 weeks, 1 day ago

Selected Answer: A

without specifying which part of redundancy (power generation, power distribution, or other aspects)—then PDUs could also be a valid answer because they provide internal redundancy for power distribution within the data center.

Re-evaluating the Answers in That Context:

✔ Generators → Internal redundancy for power backup/generation.

✔ PDUs → Internal redundancy for power distribution.

 Network Circuits → Typically external (provided by ISPs).

 Power Substations → External, controlled by utility companies.

upvoted 1 times

TheFivePips 9 months, 1 week ago

Selected Answer: A

Location: Internal redundancy duplicates components within the same facility or physical location, whereas external redundancy duplicates components across different locations or facilities.

Scope of Protection: Internal redundancy protects against component-level failures within a facility, ensuring continuous operation and minimizing downtime. External redundancy protects against broader risks such as facility-level disasters or regional outages.

Implementation: Internal redundancy typically involves redundant hardware or systems installed within the same data center or building. External redundancy requires duplicate infrastructure in geographically diverse locations, often involving additional costs and logistical considerations.

upvoted 1 times

nelombg 1 year, 10 months ago

A is the correct answer

upvoted 1 times

akg001 2 years, 10 months ago

D. Generators

upvoted 1 times

MariaGabiGabriela 3 years, 1 month ago

Selected Answer: D

I would say only generators provide a source of redundancy. All others would cause availability issues if they fail.

upvoted 1 times

MariaGabiGabriela 3 years, 1 month ago

I would say only generators provide a source of redundancy. All others would cause availability issues if they fail.

upvoted 2 times

Question #58 *Topic 1*

Which of the following represents a control on the maximum amount of resources that a single customer, virtual machine, or application can consume within a cloud environment?

A. Share

B. Reservation

C. Provision

D. Limit

Suggested Answer: *D*
Limits are put in place to enforce a maximum on the amount of memory or processing a cloud customer can use. This can be done either on a virtual machine or as a comprehensive whole for a customer, and is meant to ensure that enormous cloud resources cannot be allocated or consumed by a single host or customer to the detriment of other hosts and customers.

akg001 4 months, 1 week ago
D. Limit
upvoted 1 times

Question #59 *Topic 1*

Which of the following roles is responsible for peering with other cloud services and providers?

A. Cloud auditor

B. Inter-cloud provider

C. Cloud service broker

D. Cloud service developer

Suggested Answer: *B*
The inter-cloud provider is responsible for peering with other cloud services and providers, as well as overseeing and managing federations and federated services.

Community vote distribution
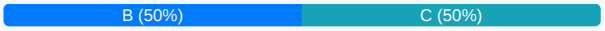

MaciekMT 4 weeks, 1 day ago

Selected Answer: B
In the (ISC)² Certified Cloud Security Professional (CCSP) framework, the Inter-cloud Provider is responsible for peering with other cloud services and providers, as well as overseeing and managing federations and federated services. This role involves establishing and maintaining relationships between different cloud environments to ensure seamless integration and interoperability
According to the CCSP study guide, the Inter-cloud Provider is tasked with "peering with other cloud services and providers, as well as overseeing and managing federations and federated services."
This role is crucial for enabling interoperability and seamless integration across diverse cloud platforms, ensuring that services can work together efficiently in a multi-cloud environment.
upvoted 1 times

krauzo 1 month, 1 week ago

Selected Answer: B
Inter-cloud provider. Peering is CSP responsibility.
upvoted 1 times

lolanczos 3 months, 1 week ago

Selected Answer: C
It's C.

This term is not widely recognized or standardized in cloud computing terminology, including within the (ISC)² CCSP framework. While it might intuitively seem to fit the role described, it doesn't align with established roles and definitions commonly used in the industry or within CCSP study materials.
upvoted 1 times

lolanczos 3 months, 1 week ago

nter-cloud Provider: This term is not widely recognized or standardized in cloud computing terminology, including within the (ISC)² CCSP framework. While it might intuitively seem to fit the role described, it doesn't align with established roles and definitions commonly used in the industry or within CCSP study materials.
upvoted 1 times

hacker1919 8 months, 3 weeks ago

B.
The CSP:inter-cloud provider is a sub-role of cloud service provider that relies on one or more peer cloud service providers to provide part or all of the cloud services offered to cloud service customers by that CSP:inter-cloud provider. The CSP:inter-cloud provider's main activities are the intermediation, aggregation, arbitrage, peering or federation of peer cloud service providers' cloud services and their business and administration capabilities from the cloud service customer viewpoint so that the cloud service customer only uses the service, business and administration interfaces of the inter-cloud service provider.
upvoted 2 times

TheFivePips 9 months ago

Selected Answer: B

C. Cloud Service Broker

The cloud service broker acts as an intermediary between cloud consumers and cloud providers, offering services such as aggregation, integration, and customization.

While cloud service brokers may facilitate the use of multiple cloud services and provide value-added services, they do not specifically handle the technical aspects of peering and direct integration between cloud providers.

B. Inter-cloud provider

An inter-cloud provider is a service provider that facilitates the interconnection and interoperability between multiple cloud environments. This role involves enabling communication and data exchange between different cloud services and providers, ensuring seamless integration and cooperation across various cloud platforms.

upvoted 1 times

globy118 9 months, 3 weeks ago

Selected Answer: C

C. Cloud service broker

upvoted 1 times

akg001 2 years, 10 months ago

B. Inter-cloud provider

upvoted 2 times

Question #60 *Topic 1*

Which of the following does NOT relate to the hiding of sensitive data from data sets?

A. Obfuscation

B. Federation

C. Masking

D. Anonymization

Suggested Answer: *B*
Federation pertains to authenticating systems between different organizations.

akg001 4 months, 1 week ago
B. Federation
upvoted 2 times

vega 2 years ago
the answer to this question does not correspond with question.
upvoted 1 times

kjjcraigskel 1 year, 11 months ago
It does. it's saying federation is the one incorrect answer and proceeds to define it.
upvoted 4 times

Question #61 *Topic 1*

Which of the following are the storage types associated with IaaS?

A. Volume and object

B. Volume and label

C. Volume and container

D. Object and target

Suggested Answer: *A*

Community vote distribution

A (100%)

Seke Highly Voted 2 years, 7 months ago

absolutely

upvoted 6 times

ACNgo Most Recent 3 weeks, 4 days ago

Selected Answer: A

In Infrastructure as a Service (IaaS), the primary storage types are:

Volume Storage: This refers to block storage, which is typically used for structured data and is often attached to virtual machines (e.g., Amazon EBS, Azure Disk Storage).

Object Storage: This is used for unstructured data and is scalable, often accessed via APIs (e.g., Amazon S3, Azure Blob Storage).

The other options are not standard storage types in IaaS:

B. Volume and label: "Label" is not a storage type.

C. Volume and container: "Container" refers to a runtime environment for applications, not a storage type.

D. Object and target: "Target" is not a storage type.

upvoted 1 times

BuckLee 6 months ago

Selected Answer: A

On the A team

upvoted 2 times

Question #62 *Topic 1*

Which technology can be useful during the "share" phase of the cloud data lifecycle to continue to protect data as it leaves the original system and security controls?

A. IPS

B. WAF

C. DLP

D. IDS

Suggested Answer: *C*
Data loss prevention (DLP) can be applied to data that is leaving the security enclave to continue to enforce access restrictions and policies on other clients and systems.

Banzaaai 6 months, 1 week ago
Answer C is correct.

DLP might be based on tags managed by IRM or manually.
upvoted 4 times

David_S 10 months, 3 weeks ago
IRM is not a listed choice, but would be a better answer. https://en.wikipedia.org/wiki/Information_rights_management
upvoted 2 times

Mhiar 12 months ago
https://digitalguardian.com/blog/what-data-loss-prevention-dlp-definition-data-loss-prevention
upvoted 2 times

Question #63 *Topic 1*

Which of the following storage types is most closely associated with a traditional file system and tree structure?

A. Volume

B. Unstructured

C. Object

D. Structured

Suggested Answer: *A*
Volume storage works as a virtual hard drive that is attached to a virtual machine. The operating system sees the volume the same as how a traditional drive on a physical server would be seen.

Community vote distribution
A (87%) 13%

nachoqueen Highly Voted 1 year, 7 months ago

Selected Answer: A

Volume Storage. Reference: page 233 of CCSP Official (ISC)2 Student Guide 4th Edition

upvoted 8 times

krauzo Most Recent 1 month, 1 week ago

Selected Answer: A

Volume storage.
Volume storage – Volume storage is like a virtualized version of a physical hard drive.
Object storage – Data is stored as objects, which are basically just collections of bits with an identifier and metadata.

upvoted 1 times

MaciekMT 2 months, 1 week ago

Selected Answer: A

A. Volume. This type of storage emulates the structure of physical hard drives, allowing data to be organized hierarchically using directories and subdirectories (a tree structure). It is commonly used in operating systems like Windows, Linux, and macOS.

Object storage organizes data as objects, which include metadata and a unique identifier. It does not use a hierarchical file system; instead, objects are stored in a flat namespace.

upvoted 1 times

rodlee 7 months, 3 weeks ago

ok volume stoage but why about tree !!! Question is not good enough !!!

upvoted 2 times

Pika26 10 months, 1 week ago

Selected Answer: A

A: Volume

upvoted 2 times

TheGinjaNinja 11 months, 1 week ago

Selected Answer: C

I'd go with Object storage

upvoted 1 times

pwxfuchxaubgkfcqay 1 year, 9 months ago

Selected Answer: A

Volume Storage

upvoted 3 times

F34 1 year, 10 months ago

Selected Answer: C

Object Storage

upvoted 1 times

tomres 2 years, 3 months ago
Object Storage
upvoted 2 times

kabbra 2 years, 4 months ago
This should be Object Storage.
According to CSA, this is their definition: “ Object storage: Object storage is similar to a file system…”
upvoted 2 times

Seke 2 years, 7 months ago
Should be Object Storage.

Object: An object is file storage that can be accessed directly through an API or web interface, without being attached to an Operating System. Data kept in object storage includes the object data and metadata and can store any kind of information, including photos, videos, documents and more. Many CSPs have interfaces that present object storage in a similar fashion to standard file tree structures (like a Windows directory), but the files are actually just virtual objects in an independent storage structure that rely on key values to reference and retrieve them.
upvoted 1 times

Xindydo 2 years, 7 months ago
You actually prove that it is not Object Storage in your blurb:
"Many CSPs have interfaces that present object storage in a similar fashion to standard file tree structures (like a Windows directory), but the files are actually just virtual objects in an independent storage structure that rely on key values to reference and retrieve them"

While a CSP may present it like tradition storage, it is not actually stored in a tree like structure. So the correct answer is Volume.
upvoted 7 times

Seke 2 years, 6 months ago
I think you are right. I do some research again about it.
File storage is meeting the requirement and it is under the volume storage.

" Volume storage architecture can take different forms; there is a great deal of discussion among cloud professionals about what type of volume might be preferable: file storage or block storage.

File Storage (Also “File-Level Storage” or “File-Based Storage”) The data is stored and displayed just as with a file structure in the legacy environment, as files and folders, with all the same hierarchical and naming functions. File storage architectures have become more popular with cloud technology and big data analytical tools and processes.
upvoted 4 times

duracell 2 years, 10 months ago
the questions asks about traditional file storage... not cloud based storage. So IMO volume based storage is correct
upvoted 4 times

vishwab 3 years ago
Believe it should be Object storage
upvoted 1 times

Question #64 *Topic 1*

Which of the following represents a prioritization of applications or cloud customers for the allocation of additional requested resources when there is a limitation on available resources?

A. Provision

B. Limit

C. Reservation

D. Share

Suggested Answer: *D*

The concept of shares within a cloud environment is used to mitigate and control the request for resource allocations from customers that the environment may not have the current capability to allow. Shares work by prioritizing hosts within a cloud environment through a weighting system that is defined by the cloud provider. When periods of high utilization and allocation are reached, the system automatically uses scoring of each host based on its share value to determine which hosts get access to the limited resources still available. The higher the value a particular host has, the more resources it will be allowed to utilize.

Community vote distribution

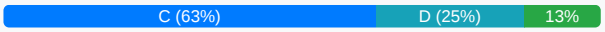

Sven007 Highly Voted 2 years, 4 months ago

C is correct, because all other options do not prioritize one customer over another. With "reservations", the customer has a predefined minimum of resources that will be given to him when he demands them.

upvoted 11 times

MaciekMT Most Recent 4 weeks, 1 day ago

Selected Answer: D

In the context of cloud computing, "shares" are utilized to prioritize resource allocation among applications or cloud customers, especially when resources are limited. This mechanism assigns a relative weight to each application or customer, determining their priority level for accessing additional resources during high-demand periods.

Why Not the Others?

A. Provision: Refers to the process of allocating resources to a customer or application but does not inherently involve prioritization during resource contention.

B. Limit: Sets a maximum threshold on the amount of resources an application or customer can use, restricting usage but not prioritizing access when resources are scarce.

C. Reservation: Guarantees a minimum amount of resources for an application or customer, ensuring baseline availability but not addressing prioritization for additional resources beyond the reserved amount.

upvoted 1 times

sweetykaur 5 months, 2 weeks ago

Share. This prioritization ensures that the available resources are allocated based on the predefined importance or needs of applications or cloud customers.

upvoted 4 times

Dasccsp 6 months, 1 week ago

Answer should be D, Share prioritize the application to allocate resource on top of the limit assigned.

upvoted 3 times

globy118 9 months, 3 weeks ago

Selected Answer: C

D. Share: Sharing resources involves dividing them among multiple users or applications. While sharing is essential, it doesn't inherently prioritize specific workload by chatGPT.

so the answer is C. Reservation

upvoted 3 times

Mo22 11 months, 1 week ago

Selected Answer: C

I agee that C is the right answer here!

upvoted 1 times

TraceSplice 1 year ago

C reservation is correct..

From GPC

To make sure that Compute Engine resources are available when you need them, use reservations. Reservations provide a very high level of assurance in obtaining capacity for Compute Engine zonal resources. You can use reservations to help ensure that your project has resources for future increases in demand

upvoted 1 times

cloudenthusiast 1 year ago

Selected Answer: C

Reservation

upvoted 1 times

Schuiram 1 year, 10 months ago

Selected Answer: A

Provisioning makes the most sense here. Many of these questions are simply asked in a very bad way.

upvoted 1 times

ikamalbhatt 1 year, 11 months ago

Selected Answer: C

Answer is C

Quita Limits are prioritised with Reservation

upvoted 1 times

Pika26 1 year, 11 months ago

Selected Answer: D

Answer is D.

upvoted 3 times

kylesam2017 2 years, 2 months ago

Option "C", Reservation, seems to be the correct answer here.

upvoted 3 times

DA95 2 years, 3 months ago

The correct answer is A. Provision. This refers to the process of setting aside or allocating resources for a specific purpose or use. In the context of cloud computing, provisioning can refer to the allocation of additional resources to applications or customers when there is a limitation on available resources. This process involves prioritizing the allocation of resources based on pre-defined criteria, such as the importance or urgency of the applications or the priorities of the customers.

upvoted 2 times

Question #65 *Topic 1*

Which type of audit report does many cloud providers use to instill confidence in their policies, practices, and procedures to current and potential customers?

A. SAS-70

B. SOC 2

C. SOC 1

D. SOX

Suggested Answer: *B*

One approach that many cloud providers opt to take is to undergo a SOC 2 audit and make the report available to cloud customers and potential cloud customers as a way of providing security confidence without having to open their systems or sensitive information to the masses.

Community vote distribution

Kneebee 6 months, 3 weeks ago

B correct - SAS70 was superseded by the SSAE 16 auditing standard in 2011

upvoted 1 times

hanyahmed 1 year, 9 months ago

Selected Answer: B

SOC2 is the right answer

upvoted 3 times

certifiedgeek 2 years, 4 months ago

SAS70 is already defunct and SOX is not applicable for cloud providers. The question has "instill confidence" doesn't mean the report will shared. SOC1 and SOC2 reports both provide this confidence. SOC2 is the better choice among the two. Not among the choices, I would prefer SOC3 report as "potential" costumers are listed in the question.

upvoted 4 times

Voldamort 2 years, 8 months ago

Selected Answer: A

It would have to be SAS-70 (now defunct) SOC 1 is financial, SOC 2 would be good but a cloud provide is not going to give that to you. SOX is there to throw you. SAS-70 Type 1 had an auditors report 'Opinion' and a SOC 2 - Type 1 style report. SOC 3 would be best but is not there.

upvoted 1 times

serget12 1 year, 11 months ago

Don't think it can be SAS-70, This report(SOC 1) is the replacement of the Statement on Auditing Standards No. 70(SAS 70). SAS-70 is deprecated. I agree that SOC 2 is usually considered a restricted report but not sure that is important to answering the question.

upvoted 1 times

deegadaze1 3 years, 6 months ago

B correct - SOC 2 is an auditing procedure that ensures your service providers securely manage your data to protect the interests of your organization and the privacy of its clients. For security-conscious businesses, SOC 2 compliance is a minimal requirement when considering a SaaS provider.

upvoted 2 times

CISSP_Wannabe 3 years, 6 months ago

Is this correct - I would have thought based on this list SOX is the best answer - can't think why SAS-70 (replaced by SSAE No 16, which is actually, SOC-1, SOC-2, and SOC-3. SOC-1 is financial and SOC-2 is effectiveness of controls and these are not on offer to potential customers. So that leaves SOX as the potential (best) answer?

upvoted 2 times

deegadaze1 3 years, 6 months ago

Sarbanes-Oxley (SOX); is U.S. law meant to protect investors from fraudulent accounting activities by corporations. SOC-2 is the best option...

upvoted 3 times

Question #66 *Topic 1*

Which of the following statements accurately describes VLANs?

A. They are not restricted to the same data center or the same racks.

B. They are not restricted to the name rack but restricted to the same data center.

C. They are restricted to the same racks and data centers.

D. They are not restricted to the same rack but restricted to same switches.

Suggested Answer: *A*
A virtual area network (VLAN) can span any networks within a data center, or it can span across different physical locations and data centers.

Community vote distribution
A (100%)

MaciekMT 4 weeks, 1 day ago

Selected Answer: A
Virtual Local Area Networks (VLANs) are designed to segment network traffic logically, independent of physical location. This means that devices assigned to the same VLAN can communicate as if they are on the same physical network, regardless of their actual physical placement.

Key Points:

VLAN Flexibility: VLANs can span across multiple switches, racks, and even data centers, allowing for flexible network design and efficient resource utilization.
CISCO.COM

Layer 2 Adjacency: In data centers, it's common to extend VLANs across various switches to meet application requirements that need Layer 2 adjacency, such as high availability clusters.
upvoted 1 times

Mdorgham 6 months, 3 weeks ago
I think this is purely wrong or poor verbiage for the questions itself.

VLANs can span DCs in general ,but a single VLAN wont! The technology that is used to do that is actually VxLAN!
upvoted 1 times

Cloud_baby 1 year, 4 months ago
ANS A. A. They are not restricted to the same data center or the same racks. But the verbiage in Statement for Correct answer contradicts answer A.
upvoted 1 times

Pika26 1 year, 10 months ago

Selected Answer: A
A. They are not restricted to the same data center or the same racks.
upvoted 1 times

Chriscrown 1 year, 7 months ago
Looks like the answer is A but in the real world we typically don't extend VLANs across Datacenters. In a cloud environment two Datacenters are "well Connected" so extending VLANs works.
upvoted 1 times

DA95 2 years, 3 months ago
VLANs, or Virtual Local Area Networks, are a way of segmenting a physical network into logical sub-networks. VLANs are typically used to isolate different groups of devices or applications on a network, such as to provide security or performance benefits. VLANs are not restricted to the same data center or the same racks, meaning that they can span multiple data centers or racks within a network. Instead, VLANs are typically defined based on logical criteria, such as the type of device or application being used, or the VLAN ID assigned to the device or application.

upvoted 2 times

Cyril_the_Squirl 2 years, 3 months ago

Correct answer is B.

You cannot stretch a VLAN over WAN, the moment you do, you're using other technologies such as PW, vXLAN, etc. to transport your Layer2 vlan....but VLANs by themselves ca only span within a datacentre or campus.

upvoted 2 times

GregP 2 years, 8 months ago

Selected Answer: A

Plenty of people use layer 2 vlans across datacentres in a stretched storage and/or hypervisor cluster active/active mode. it's a high availability strategy (as opposed to DR)

upvoted 4 times

Zeezee2 3 years, 4 months ago

A is correct, even though it is a shitty practice and I would say B is highly preferable.

upvoted 4 times

Fpaul 3 years, 9 months ago

A stretched VLAN is a VLAN that spans multiple physical data centers.

https://download3.vmware.com/vcat/vmw-vcloud-architecture-toolkit-spv1-webworks/Hybridity/Architecting%20a%20Hybrid%20Mobility%20Strategy/Architecting%20a%20Hybrid%20Mobility%20Strategy.2.14.html#

upvoted 1 times

manowwww 3 years, 10 months ago

Looks like Answer is correct. VLAN could be from different datacenter as well.

upvoted 2 times

Bhuvy 3 years, 11 months ago

Given answer is incorrect. Correct Answer is B.

VLANs are always restricted to the same data center. It can be spread across to multiple racks in the same datacenter.

For example when a packet wants to leave a datacenter it has to use a WAN connection. And VLAN name itself Virual local area network. It cannot travel through WAN.

upvoted 1 times

duracell 3 years, 10 months ago

actually you could span Layer 2 across 2 data centers... you shouldn't but you could ;)

upvoted 1 times

Question #67 *Topic 1*

What must be secured on physical hardware to prevent unauthorized access to systems?

A. BIOS

B. SSH

C. RDP

D. ALOM

Suggested Answer: *A*
BIOS is the firmware that governs the physical initiation and boot up of a piece of hardware. If it is compromised, an attacker could have access to hosted systems and make configurations changes to expose or disable some security elements on the system.

Community vote distribution
A (100%)

TraceSplice 6 months, 1 week ago

Selected Answer: A

A - BIOS
Based on NIST 800-147B (BIOS Protection Guidelines for Servers)
https://nvlpubs.nist.gov/nistpubs/specialpublications/nist.sp.800-147b.pdf
upvoted 1 times

deegadaze1 3 years ago

D - Correct answer: https://docs.oracle.com/cd/E19102-01/n440.srvr/817-5481-11/intro.html
upvoted 1 times

Zeezee2 2 years, 10 months ago

Answers a way too specific physical hardware element. This question talks on hardware in general, and securing BIOS brings far more value.
upvoted 7 times

Question #68 *Topic 1*

What type of PII is regulated based on the type of application or per the conditions of the specific hosting agreement?

A. Specific

B. Contractual

C. regulated

D. Jurisdictional

Suggested Answer: *B*
Contractual PII has specific requirements for the handling of sensitive and personal information, as defined at a contractual level. These specific requirements will typically document the required handling procedures and policies to deal with PII. They may be in specific security controls and configurations, required policies or procedures, or limitations on who may gain authorized access to data and systems.

MaciekMT 2 months, 1 week ago

Selected Answer: B

B. Contractual PII is regulated based on the specific application or conditions of a hosting agreement. This type of PII may not necessarily fall under jurisdictional or industry-specific laws but is governed by terms agreed upon between the parties involved (e.g., between a cloud service provider and a customer).

The hosting agreement or contract outlines how this type of PII should be handled, protected, and managed, often including security requirements and compliance obligations tailored to the specific use case.

upvoted 1 times

akg001 4 months, 1 week ago

B. Contractual

upvoted 1 times

Question #69 *Topic 1*

Which of the following security technologies is commonly used to give administrators access into trust zones within an environment?

A. VPN

B. WAF

C. IPSec

D. HTTPS

Suggested Answer: *A*
Virtual private networks (VPNs) are commonly used to allow access into trust zones. Via a VPN, access can be controlled and logged and only allowed through secure channels by authorized users. It also adds an additional layer of encryption and protection to communications.

BigWatch Highly Voted 8 months, 2 weeks ago
The word "Administrator" is not necessary - VPN will allow access to all authorized users
upvoted 5 times

akg001 4 months, 1 week ago
A. VPN
upvoted 3 times

ra1paul Most Recent 1 month, 3 weeks ago
Selected Answer: A
VPN is specifically designed for secure access into trusted zones.
upvoted 1 times

Question #70 *Topic 1*

Which concept BEST describes the capability for a cloud environment to automatically scale a system or application, based on its current resource demands?

A. On-demand self-service

B. Resource pooling

C. Measured service

D. Rapid elasticity

Suggested Answer: *D*
Rapid elasticity allows a cloud environment to automatically add or remove resources to or from a system or application based on its current demands. Whereas a traditional data center model would require standby hardware and substantial effort to add resources in response to load increases, a cloud environment can easily and rapidly expand to meet resources demands, so long as the application is properly implemented for it.

MaciekMT 4 weeks, 1 day ago

Selected Answer: D

Rapid Elasticity refers to the cloud's ability to automatically scale resources up or down based on real-time demand and workload changes. This feature ensures that applications and systems efficiently use computing resources without manual intervention.

Automatically adds or removes resources (e.g., CPU, memory, storage) based on demand.
Optimizes cost efficiency by ensuring resources are available when needed but not over-provisioned.
Common in cloud computing models such as Auto Scaling (AWS), Azure Virtual Machine Scale Sets, and Kubernetes Horizontal Pod Autoscaling.
Why Not B. Resource Pooling → Enables multi-tenant sharing of resources, but it does not handle dynamic scaling.
upvoted 1 times

akg001 4 months, 1 week ago

D. Rapid elasticity
upvoted 2 times

Question #71 *Topic 1*

If you're using iSCSI in a cloud environment, what must come from an external protocol or application?

A. Kerberos support

B. CHAP support

C. Authentication

D. Encryption

Suggested Answer: *D*
iSCSI does not natively support encryption, so another technology such as IPsec must be used to encrypt communications.

Community vote distribution
B (100%)

ra1paul 1 month, 3 weeks ago

Selected Answer: D

iSCSI provides methods for data transport and block-level access to storage.
upvoted 2 times

MaciekMT 2 months, 1 week ago

Selected Answer: D

iSCSI (Internet Small Computer Systems Interface) does not natively include encryption for securing data in transit. Instead, encryption must be provided by an external protocol or application, such as:

IPsec (Internet Protocol Security): For securing data at the network layer.
TLS (Transport Layer Security): For securing application-layer communication.
Without encryption, iSCSI traffic is vulnerable to interception and eavesdropping, making external encryption essential for secure deployments in cloud environments.
upvoted 2 times

lolanczos 3 months, 1 week ago

Selected Answer: D

It's 100% D.

Encryption: iSCSI does not provide built-in encryption for data in transit. To secure the data being transmitted, encryption must be handled by external protocols or applications, such as IPsec (Internet Protocol Security) or TLS (Transport Layer Security). These protocols can encrypt the data packets, ensuring confidentiality and integrity during transmission.
upvoted 2 times

sweetykaur 5 months, 2 weeks ago

Authentication. When using iSCSI in a cloud environment, authentication must be handled by an external protocol or application to ensure secure connections.
upvoted 1 times

globy118 9 months, 3 weeks ago

Selected Answer: B

B. Use CHAP (Challenge Handshake Authentication Protocol) to ensure each host has its own password.
upvoted 1 times

ST42 1 year ago

Data Encryption: the data exchanged between the iSCSI initiator and iSCSI target is not encrypted, which is why an attacker who is able to sniff the data of the wire will also be able to reconstruct the data, or more precisely, the files and directories transferred between the two parties. Additionally, an attacker might also be able to inject his own data into the traffic, thus creating/modifying/deleting arbitrary files on the iSCSI target. To mitigate the risk, the IPsec ought to be used to properly encrypt the communication between the iSCSI endpoints.

https://resources.infosecinstitute.com/topics/cloud/iscsi-security-considerations-cloud/
upvoted 1 times

akg001 2 years, 10 months ago

D. Encryption

upvoted 4 times

Question #72 *Topic 1*

Which of the following pertains to a macro level approach to data center design rather than the traditional tiered approach to data centers?

A. IDCA

B. NFPA

C. BICSI

D. Uptime Institute

Suggested Answer: *A*
The standards put out by the International Data Center Authority (IDCA) have established the Infinity Paradigm, which is intended to be a comprehensive data center design and operations framework. The Infinity Paradigm shifts away from many models that rely on tiered architecture for data centers, where each successive tier increases redundancy. Instead, it emphasizes data centers being approached at a macro level, without a specific and isolated focus on certain aspects to achieve tier status.

MaciekMT 2 months, 1 week ago

Selected Answer: A

The International Data Center Authority (IDCA) takes a macro-level approach to data center design, focusing on the overall effectiveness, efficiency, and resilience of the data center as a whole. Unlike the traditional tiered approach (such as that defined by the Uptime Institute), IDCA emphasizes a broader perspective, considering factors like sustainability, adaptability, and alignment with business objectives rather than just reliability tiers.

Uptime Institute: The organization behind the tiered approach to data center design, which focuses on reliability and redundancy rather than broader considerations.

upvoted 1 times

sweetykaur 5 months, 2 weeks ago

BICSI. BICSI's methodologies offer a comprehensive view of data center design, focusing on the entire infrastructure at a macro level, as opposed to the traditional tiered approach.

upvoted 1 times

akg001 2 years, 10 months ago

A. IDCA

upvoted 1 times

Question #73 *Topic 1*

What does the REST API support that SOAP does NOT support?

A. Caching

B. Encryption

C. Acceleration

D. Redundancy

Suggested Answer: *A*

The SOAP protocol does not support caching, whereas the REST API does.

Community vote distribution

A (100%)

MaciekMT 1 month, 1 week ago

Selected Answer: A

REST is built on HTTP and naturally leverages its built-in caching mechanisms, allowing responses to be stored and reused—speeding up performance and reducing server load. SOAP, with its rigid XML-based messaging, doesn't have native support for caching. So, when it comes to caching, REST takes the cake while SOAP is left out in the cold.

upvoted 1 times

globy118 3 months, 2 weeks ago

Selected Answer: A

SOAP relies on XML for all messages

upvoted 1 times

akg001 2 years, 4 months ago

A. Caching

upvoted 1 times

Question #74 *Topic 1*

Why does a Type 1 hypervisor typically offer tighter security controls than a Type 2 hypervisor?

A. A Type 1 hypervisor also controls patching of its hosted virtual machines ensure they are always secure.

B. A Type 1 hypervisor is tied directly to the bare metal and only runs with code necessary to perform its specific mission.

C. A Type 1 hypervisor performs hardware-level encryption for tighter security and efficiency.

D. A Type 1 hypervisor only hosts virtual machines with the same operating systems as the hypervisor.

Suggested Answer: *B*
Type 1 hypervisors run directly on top of the bare metal and only contain the code and functions required to perform their purpose. They do not rely on any other systems or contain extra features to secure.

akg001 4 months, 1 week ago

B. A Type 1 hypervisor is tied directly to the bare metal and only runs with code necessary to perform its specific mission.
upvoted 2 times

Question #75 *Topic 1*

Which of the following are the storage types associated with PaaS?

A. Structured and freeform

B. Volume and object

C. Structured and unstructured

D. Database and file system

Suggested Answer: *C*

Community vote distribution

akg001 Highly Voted 2 years, 10 months ago

Selected Answer: C

C. Structured and unstructured

upvoted 8 times

MaciekMT Most Recent 2 months, 1 week ago

Selected Answer: C

In Platform as a Service (PaaS), the storage types most commonly associated are:

Structured storage: Used for data with a defined schema, such as databases (e.g., relational databases like SQL).
Unstructured storage: Used for data without a predefined structure, such as documents, images, videos, or other media files (e.g., object storage like Azure Blob Storage or AWS S3).

upvoted 1 times

lolanczos 3 months, 1 week ago

Selected Answer: C

It's C. Platform as a Service (PaaS) provides storage types that cater to application development and deployment needs. These typically include:

Structured storage: For managing organized data, such as relational databases.
Unstructured storage: For handling files, logs, media, and other types of data that do not follow a strict schema.

upvoted 1 times

lolanczos 3 months, 1 week ago

Selected Answer: C

It's definitely C.

Platform as a Service (PaaS) offerings typically provide data storage options that cater to both structured and unstructured data types:

Structured data: Information organized into defined fields and schemas, often stored in relational databases or other structured data stores. PaaS services commonly offer managed database options (e.g., SQL-based databases) as a core feature.

Unstructured data: Information that does not follow a predefined schema, such as documents, images, videos, or large binary objects. PaaS platforms frequently incorporate object storage or other non-relational storage solutions for these data types.

upvoted 1 times

vaolo 3 months, 2 weeks ago

Selected Answer: C

Answer is C
Database and file system: While relevant, they are more specific types of storage implementations rather than broader categories.

upvoted 1 times

sweetykaur 5 months, 2 weeks ago

Structured and unstructured. These storage types are typically associated with Platform as a Service (PaaS) offerings.

upvoted 1 times

TheFivePips 9 months ago

Selected Answer: D

Database: PaaS often includes managed database services, where developers can deploy and manage databases without worrying about the underlying infrastructure details.
File System: PaaS platforms may also provide file storage services or filesystem APIs that allow applications to store and retrieve files or data objects.

A. Structured and freeform: This option is not typically used to categorize storage types associated with PaaS. "Freeform" is not a standard storage type classification in the context of PaaS.
B. Volume and object: While these are valid storage types, they are not typically specifically associated with PaaS. "Volume" and "object" are more commonly associated with IaaS (Infrastructure as a Service) for managing storage resources directly.
C. Structured and unstructured: These terms describe types of data rather than storage types. They refer to how data is organized and managed, not specific storage services or technologies provided by PaaS platforms.

upvoted 3 times

FranklinG 1 year ago

This is a tricky one. Because C and D can be correct. However, answer is D. Here's why; The PaaS service model utilizes two categories of storage, Structured and Unstructured. But the question ask, which of the following are the storage types associated with PaaS? Key words: "STORAGE TYPES ASSOCIATED WITH PAAS," that would be Database and file system.

upvoted 2 times

cloudenthusiast 1 year, 1 month ago

Selected Answer: D

PaaS is associated with database according it ISC2.

upvoted 2 times

Onimole 1 year, 3 months ago

C
The PaaS service model utilizes two categories of storage. Structured and Unstructured
https://www.dummies.com/article/academics-the-arts/study-skills-test-prep/ccsp/cloud-data-storage-architectures-274125/

upvoted 2 times

badjo 1 year, 3 months ago

Structured and Unstructured, Answer C

upvoted 1 times

Kneebee 1 year, 5 months ago

The correct answer is "C"
PaaS service model uses 2 categories of storage: Structured and Unstructured

IaaS service model users 2 categories of storage: Volume and Object

SaaS service model commonly utilizes 2 METHODs (not types)of data storages - those are:
Information Storage and Management & Content and File Management.

upvoted 1 times

BuckLee 1 year, 5 months ago

Selected Answer: D

D for Delta

upvoted 1 times

Krishna2637 1 year, 7 months ago

Structured and unstructured is not the storage type.

upvoted 2 times

kamin123 1 year, 8 months ago

Selected Answer: D

Database and file system

upvoted 3 times

Purespace 1 year, 11 months ago

Selected Answer: C

The reason is that "files" is contained within the concept of unstructured data storage. That is, unstructured storage is not organized into database format, meaning that "structured" storage must follow strict rules (called the Schema) about how data is stored (e.g., in tables in a database). Thus it follows that data which is "loose" and comprised of traditional files and folders is "unstructured" (even if you have your desktop folders arranged in lovely rows, it's still unstructured data).

upvoted 1 times

luckflying 2 years ago

Selected Answer: D

PaaS supports DB + files.

upvoted 4 times

Question #76 *Topic 1*

Which of the following threat types can occur when baselines are not appropriately applied or unauthorized changes are made?

A. Insecure direct object references

B. Unvalidated redirects and forwards

C. Security misconfiguration

D. Sensitive data exposure

Suggested Answer: *C*

Security misconfigurations occur when applications and systems are not properly configured or maintained in a secure manner. This can be caused from a shortcoming in security baselines or configurations, unauthorized changes to system configurations, or a failure to patch and upgrade systems as the vendor releases security patches.

Community vote distribution

C (100%)

MaciekMT 1 month, 1 week ago

Selected Answer: C

When baselines aren't properly applied or unauthorized changes occur, you end up with security misconfiguration—basically, your system's settings are off-kilter, creating vulnerabilities. This can lead to a host of issues like open ports, default settings, or unpatched software that attackers can exploit. The other options target more specific vulnerabilities, but misconfiguration is the broader issue resulting from not sticking to your secure baselines.

upvoted 1 times

akg001 4 months, 1 week ago

Selected Answer: C

C. Security misconfiguration

upvoted 3 times

akg001 4 months, 1 week ago

C. Security misconfiguration

upvoted 1 times

Question #77 *Topic 1*

What is the data encapsulation used with the SOAP protocol referred to?

A. Packet

B. Envelope

C. Payload

D. Object

Suggested Answer: *B*

Simple Object Access Protocol (SOAP) encapsulates its information in what is known as a SOAP envelope and then leverages common communications protocols for transmission.

Community vote distribution

B (100%)

jonclem Highly Voted 1 year, 8 months ago

Handy to know. Also found this link to clarify the answer:
https://en.wikipedia.org/wiki/SOAP

upvoted 5 times

Mobi333 Most Recent 4 months, 1 week ago

B. Envelope

upvoted 3 times

akg001 4 months, 1 week ago

Selected Answer: B

B. Envelope

upvoted 3 times

Question #78 *Topic 1*

Which of the following threat types can occur when an application does not properly validate input and can be leveraged to send users to malicious sites that appear to be legitimate?

A. Unvalidated redirects and forwards

B. Insecure direct object references

C. Security miscomfiguration

D. Sensitive data exposure

Suggested Answer: *A*

Many web applications offer redirect or forward pages that send users to different, external sites. If these pages are not properly secured and validated, attackers can use the application to forward users off to sites for phishing or malware attempts. These attempts can often be more successful than direct phishing attempts because users will trust the site or application that sent them there, and they will assume it has been properly validated and approved by the trusted application's owners or operators. Security misconfiguration occurs when applications and systems are not properly configured for security--often a result of misapplied or inadequate baselines. Insecure direct object references occur when code references aspects of the infrastructure, especially internal or private systems, and an attacker can use that knowledge to glean more information about the infrastructure. Sensitive data exposure occurs when an application does not use sufficient encryption and other security controls to protect sensitive application data.

Community vote distribution

SCha81 2 weeks, 6 days ago

Selected Answer: A

Unvalidated redirects and forwards occur when an application does not properly validate user input before redirecting or forwarding requests. Attackers exploit this weakness to redirect users to malicious sites that appear legitimate, leading to phishing attacks, malware downloads, or credential theft.

upvoted 1 times

MaciekMT 1 month, 1 week ago

Selected Answer: A

When an application doesn't validate input properly, it may inadvertently allow attackers to craft URLs that redirect users to malicious sites. This is known as unvalidated redirects and forwards, and it poses a significant risk because users might be tricked into believing they're navigating within a legitimate environment.

upvoted 1 times

akg001 4 months, 1 week ago

Selected Answer: A

A. Unvalidated redirects and forwards

upvoted 1 times

Question #79 *Topic 1*

Which publication from the United States National Institute of Standards and Technology pertains to defining cloud concepts and definitions for the various core components of cloud computing?

A. SP 800-153

B. SP 800-145

C. SP 800-53

D. SP 800-40

Suggested Answer: *B*

NIST Special Publications 800-145 is titled "The NIST Definition of Cloud Computing" and contains definitions and explanations of core cloud concepts and components.

Community vote distribution

B (100%)

ST42 6 months, 3 weeks ago

B. SP 800-145: The NIST Definition of Cloud Computing

SP 800-153: Guidelines for Securing Wireless Local Area Networks (WLANs)
SP 800-53: Security and Privacy Controls for Information Systems and Organizations
SP 800-40: Guide to Enterprise Patch Management Planning: Preventive Maintenance for Technology

upvoted 1 times

akg001 2 years, 4 months ago

Selected Answer: B

B. SP 800-145

upvoted 1 times

Question #80 *Topic 1*

What is the biggest negative to leasing space in a data center versus building or maintain your own?

A. Costs

B. Control

C. Certification

D. Regulation

Suggested Answer: *B*

When leasing space in a data center, an organization will give up a large degree of control as to how it is built and maintained, and instead must conform to the policies and procedures of the owners and operators of the data center.

Community vote distribution

B (100%)

Dasccsp 6 months, 1 week ago

Correct answer is B, the questions says what is the biggest negative in leasing that is control as we loose control compare to our own data center, cost is the positive so the answer should be B.

upvoted 1 times

Lenell 2 years, 3 months ago

Selected Answer: B

This question is a bit vague but if you take only what it provides, you loose physical control when you lease. You can't manage the physical assets as well as if you owned the DC. If more information on the situation is added, then there might be a great cause to select "cost." So this is about wat does Comptia say...cost is a positive to leasing and control is a positive to building.

upvoted 1 times

serget12 2 years, 5 months ago

Leasing space at a data center, means you are paying to have the DC supply Internet/Cooling/Power, you get to CONTROL everything else. You want Citrix and VMware, go ahead, you your on SAN controller, network gear. You get all the control. Cost will go up since you must procure all the hardware/licensing/admin/management. Don't think this is a good question.

upvoted 1 times

akg001 2 years, 10 months ago

Selected Answer: B

B. Control

upvoted 2 times

kepalon 3 years ago

CORRECT- NEGATIVE: Control & POSIVITE: Costs

upvoted 3 times

kollmekay 1 year, 1 month ago

Positive rather should be advantage = Control, Negative is disadvantage = Cost

Question has to be looked once again

upvoted 1 times

Question #81 *Topic 1*

Which aspect of archiving must be tested regularly for the duration of retention requirements?

A. Availability

B. Recoverability

C. Auditability

D. Portability

Suggested Answer: *B*

In order for any archiving system to be deemed useful and compliant, regular tests must be performed to ensure the data can still be recovered and accessible, should it ever be needed, for the duration of the retention requirements.

Community vote distribution

B (100%)

Pika26 4 months, 1 week ago

Selected Answer: B

B. Recoverability

upvoted 1 times

akg001 1 year, 4 months ago

Selected Answer: B

B. Recoverability

upvoted 1 times

NastyNutsu 1 year, 6 months ago

things that are available might not always be recoverable (aka corrupted files)

upvoted 1 times

Ahbey_911 2 years, 8 months ago

It's what I thought too, but the explanation made some sense too. I'll argue recoverability is a subset of availability. This is a case of "examiner decides the correct answer".

upvoted 1 times

reckert001 2 years, 8 months ago

What makes recoverability the correct response as opposed to availability?

upvoted 1 times

Sa007788 2 years, 8 months ago

That data may be available but we can have an issue when we cover it.Of course if date is not available we can't cover it.So recoverability is most important to test.

upvoted 2 times

nidoz 2 years, 8 months ago

Archived Data always tested for recoverability purposes to make sure that in event of Disaster, it should be recovered as intended.

upvoted 7 times

Question #82 *Topic 1*

Which of the following represents a minimum guaranteed resource within a cloud environment for the cloud customer?

A. Reservation

B. Share

C. Limit

D. Provision

Suggested Answer: *A*

A reservation is a minimum resource that is guaranteed to a customer within a cloud environment. Within a cloud, a reservation can pertain to the two main aspects of computing: memory and processor. With a reservation in place, the cloud provider guarantees that a cloud customer will always have at minimum the necessary resources available to power on and operate any of their services.

Community vote distribution

A (100%)

akg001 4 months, 1 week ago

Selected Answer: A

A. Reservation

upvoted 2 times

Question #83 *Topic 1*

When is a virtual machine susceptible to attacks while a physical server in the same state would not be?

A. When it is behind a WAF

B. When it is behind an IPS

C. When it is not patched

D. When it is powered off

Suggested Answer: *D*

A virtual machine is ultimately an image file residing a file system. Because of this, even when a virtual machine is "powered off," it is still susceptible to attacks and modification. A physical server that is powered off would not be susceptible to attacks.

Community vote distribution

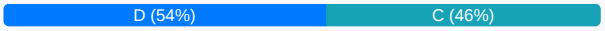

Masia767 Highly Voted 2 years, 8 months ago

Should be D, since vm is still vulnerable even when powered off, because if hypervisor or host machine gets compromised while powered ON, offline VM can be the target still. Physical machine when offline is not accessible .

upvoted 12 times

TheFivePips Most Recent 9 months ago

Selected Answer: D

I apologize for any confusion earlier. Let's address the question directly:

When is a virtual machine susceptible to attacks while a physical server in the same state would not be?

The correct answer is:

D. When it is powered off

Explanation:

Powered-off Virtual Machine (VM): Even when a virtual machine is powered off, its virtual disks and configurations are still stored on the hypervisor's storage system. While access to these resources is typically secured, vulnerabilities could potentially be exploited through administrative interfaces or if snapshots are not properly managed or encrypted.

Powered-off Physical Server: A powered-off physical server, on the other hand, typically requires physical access to the server hardware to compromise its data or security. Physical security measures such as locked server rooms or cabinets provide significant protection against unauthorized access.

upvoted 3 times

globy118 9 months, 3 weeks ago

Selected Answer: C

When a VM is not patched, it remains susceptible to attacks even when powered off. In contrast, a powered-off physical server is less exposed.

upvoted 2 times

Kneebee 1 year ago

When a virtual machine (VM) is powered off, it's still susceptible to certain types of attacks; whereas a physical server in the same state would not be.

upvoted 1 times

Krishna2637 1 year, 7 months ago

Selected Answer: D

Patch issues exists in both cases.

upvoted 4 times

Pika26 1 year, 10 months ago

Selected Answer: C

C. When it is not patched

upvoted 2 times

akg001 2 years, 10 months ago

Selected Answer: C

C. When it is not patched

upvoted 2 times

skis4u 3 years, 4 months ago

why not option C ? un-patched VM's can be susceptible to attacks.

upvoted 1 times

EdwardLeeBurtle 3 years, 3 months ago

I believe that the stipulation of the question that a physical server would not be susceptible to the same attack would disqualify C as both VMs and Physical servers are susceptible if unpatched.

upvoted 5 times

Question #84 *Topic 1*

Which of the following threat types involves an application developer leaving references to internal information and configurations in code that is exposed to the client?

A. Sensitive data exposure

B. Security misconfiguration

C. Insecure direct object references

D. Unvalidated redirect and forwards

Suggested Answer: *C*

An insecure direct object reference occurs when a developer has in their code a reference to something on the application side, such as a database key, the directory structure of the application, configuration information about the hosting system, or any other information that pertains to the workings of the application that should not be exposed to users or the network.
Unvalidated redirects and forwards occur when an application has functions to forward users to other sites, and these functions are not properly secured to validate the data and redirect requests, allowing spoofing for malware of phishing attacks. Sensitive data exposure occurs when an application does not use sufficient encryption and other security controls to protect sensitive application data. Security misconfigurations occur when applications and systems are not properly configured or maintained in a secure manner.

Community vote distribution

MaciekMT 2 months, 1 week ago

Selected Answer: C

from AI: Insecure direct object references (IDOR) is a threat type that occurs when an application developer leaves references to internal information in code that is exposed to the client.
Explanation
IDOR
A vulnerability that occurs when an application provides direct access to objects based on user-supplied input. This can happen when an application uses an identifier to access an object in a database without checking for access control or authentication.
Attackers
Attackers can use IDOR to bypass authorization and access resources in the system directly, such as database records or files.
Causes
IDOR can occur due to missing access control checks, which fail to verify whether a user should be allowed to access specific data.

upvoted 1 times

Pika26 4 months, 1 week ago

Selected Answer: C

C. Insecure direct object references

upvoted 1 times

xroxro 1 year, 1 month ago

Question not precise enough to clearly choose between A and C.
A if data is generic data (for example, version of the internal database)
C if used by a backend application (for example internal authentication token)

upvoted 3 times

akg001 1 year, 4 months ago

Selected Answer: C

C. Insecure direct object references

upvoted 2 times

certifiedgeek 1 year, 4 months ago

This question can also lean forward with "sensitive information disclosure".

upvoted 2 times

DarkHorse99 1 year, 2 months ago

true but no need to add confusion

upvoted 1 times

Question #85 *Topic 1*

Which of the following is the biggest concern or challenge with using encryption?

A. Dependence on keys

B. Cipher strength

C. Efficiency

D. Protocol standards

Suggested Answer: *A*

No matter what kind of application, system, or hosting model used, encryption is 100 percent dependent on encryption keys. Properly securing the keys and the exchange of them is the biggest and most important challenge of encryption systems.

Community vote distribution

A (100%)

Purespace Highly Voted 5 months, 2 weeks ago

Exam tip here: avoid bringing in "what if's" when evaluating the question. For example, when I first read this I thought that maybe the cipher strength would be the biggest concern "if" someone were to use a weak cipher that was easy to defeat. But I used an "if" in my logic. The question doesn't state anything about the cipher used, so we can't assume that a weak cipher would be the biggest concern when protection of encryption keys is critical across ALL encryption ciphers.

upvoted 5 times

akg001 Most Recent 1 year, 4 months ago

Selected Answer: A

A. Dependence on keys

upvoted 3 times

Question #86 *Topic 1*

Which of the following would NOT be considered part of resource pooling with an Infrastructure as a Service implementation?

A. Storage

B. Application

C. Mamory

D. CPU

Suggested Answer: *B*

Infrastructure as a Service pools the compute resources for platforms and applications to build upon, including CPU, memory, and storage. Applications are not part of an IaaS offering from the cloud provider.

Community vote distribution

B (100%)

Murakh Highly Voted 3 years, 10 months ago

doesn't change the answer but memory spelling is wrong

upvoted 20 times

nelombg Most Recent 4 months ago

B. Please correct the memory spelling.

upvoted 2 times

akg001 1 year, 4 months ago

Selected Answer: B

B. Application

upvoted 2 times

Question #87 *Topic 1*

Which technology is NOT commonly used for security with data in transit?

A. DNSSEC

B. IPsec

C. VPN

D. HTTPS

Suggested Answer: *A*

DNSSEC relates to the integrity of DNS resolutions and the prevention of spoofing or redirection, and does not pertain to the actual security of transmissions or the protection of data.

Community vote distribution

A (100%)

akg001 4 months, 1 week ago

Selected Answer: A

A. DNSSEC

upvoted 1 times

Sa007788 1 year, 8 months ago

if the question was about confidentiality, i think answer will be more accurate because DNSSEC refer to protect integrity and this is one of the CIA Security tirad also.

upvoted 1 times

Guest4768 2 years, 5 months ago

DNSSEC provides protection for the DNS intngrity, where integrity is of course one of the key security primitive, thus this answer explanation is too considerless. Still, the answer is correct.

upvoted 2 times

Question #88 *Topic 1*

Which of the following roles is responsible for gathering metrics on cloud services and managing cloud deployments and the deployment processes?

A. Cloud service business manager

B. Cloud service operations manager

C. Cloud service manager

D. Cloud service deployment manager

Suggested Answer: *D*
The cloud service deployment manager is responsible for gathering metrics on cloud services, managing cloud deployments and the deployment process, and defining the environments and processes.

ichnos Highly Voted 3 years ago
Correct is D:
Cloud service deployment manager: Gathers metrics on cloud services,
manages deployment steps and processes, defines the environment and
processes
upvoted 6 times

Kanthie 2 years, 11 months ago
D is correct based on ISC2
upvoted 2 times

MaciekMT Most Recent 2 months, 1 week ago
Selected Answer: D
A cloud service deployment manager is responsible for defining the environments and processes, gathering metrics on cloud services, managing cloud deployments, and overseeing the deployment process.
upvoted 1 times

DA95 9 months, 3 weeks ago
The correct answer is C. Cloud service manager. A cloud service manager is responsible for overseeing the deployment, operations, and management of cloud services. This typically involves gathering metrics on cloud service performance, managing cloud deployments and deployment processes, and ensuring that the cloud services are meeting the needs of the organization. Other roles that may be involved in cloud service management include cloud service business managers, who focus on the financial aspects of cloud services, and cloud service operations managers, who focus on the day-to-day operations and maintenance of the cloud services.
upvoted 2 times

kepalon 1 year, 6 months ago
I agree, this one is not clear. Actually there is a previous question which mentioned that in the CBK does not exist this role.
I will consider this to be CORRECT though!!!
upvoted 2 times

EdwardLeeBurtle 1 year, 9 months ago
Frustrating, I have 4 different study books and none of them even mention this role and it seems crazy to reference a source that isn't available for general access or reference.
upvoted 4 times

samir45 7 months, 3 weeks ago
It is answered in page 24 CCSP Certified Cloud Security Professional All-in-One Exam Guide 2022 edition.
upvoted 1 times

samir45 7 months, 3 weeks ago
The answer is D.
Cloud service deployment manager : Gathers metrics on cloud services, manages deployment steps and processes, and defines the environment and processes.
upvoted 1 times

KCjoe 3 years, 1 month ago

It should be C.

upvoted 1 times

SueHam2 3 years, 2 months ago

This seems close to another question, which the answer was: Cloud service manager...which is correct in this case?

upvoted 3 times

drop_table 3 years, 2 months ago

This answer looks correct to me, the Cloud Service Deployment Manager is defined in ISO 17789:2014, section 8.3.1.2. It's confusing because this is an ISO role, it's not defined in the ISC2 CBK.

upvoted 3 times

Question #89 *Topic 1*

Which of the following is considered an external redundancy for a data center?

A. Power feeds to rack

B. Generators

C. Power distribution units

D. Storage systems

Suggested Answer: *B*
Generators are considered an external redundancy to a data center. Power distribution units (PDUs), storage systems, and power feeds to racks are all internal to a data center, and as such they are considered internal redundancies.

MaciekMT 4 weeks, 1 day ago

Selected Answer: A

Power feeds to rack come from the external power grid or utility provider, making them an external redundancy.
They provide primary power to the data center but are not controlled internally.
If external power fails, the data center must rely on internal redundancy (e.g., generators, UPS systems).
B. Generators → Internal redundancy → Located inside the data center and provide backup power.
C. Power Distribution Units (PDUs) → Internal redundancy → Distributes power within the data center but does not generate or supply it.
D. Storage Systems → Internal redundancy → Ensures data availability within the data center, but is not an external backup mechanism.

upvoted 1 times

akg001 4 months, 1 week ago

B. Generators

upvoted 1 times

Question #90 *Topic 1*

Which of the following is the optimal humidity level for a data center, per the guidelines established by the America Society of Heating, Refrigeration, and Air
Conditioning Engineers (ASHRAE)?

A. 30-50 percent relative humidity

B. 50-75 percent relative humidity

C. 20-40 percent relative humidity

D. 40-60 percent relative humidity

Suggested Answer: *D*
The guidelines from ASHRAE establish 40-60 percent relative humidity as optimal for a data center.

MaciekMT 2 months, 1 week ago

Selected Answer: D

The American Society of Heating, Refrigerating, and Air-Conditioning Engineers (ASHRAE) recommends that data centers maintain a relative humidity (rH) of 40–60%

upvoted 1 times

muchsnow 4 months, 2 weeks ago

But what is the actual reference???

upvoted 1 times

akg001 2 years, 10 months ago

D. 40-60 percent relative humidity

upvoted 2 times

Question #91 *Topic 1*

What is the first stage of the cloud data lifecycle where security controls can be implemented?

A. Use

B. Store

C. Share

D. Create

Suggested Answer: *B*

The "store" phase of the cloud data lifecycle, which typically occurs simultaneously with the "create" phase, or immediately thereafter, is the first phase where security controls can be implemented. In most case, the manner in which the data is stored will be based on its classification.

Community vote distribution

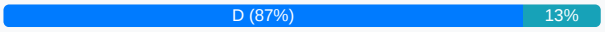

brandV Highly Voted 3 years, 2 months ago

Selected Answer: D

While security controls are implemented in the create phase in the form of SSL/TLS, this only protects data in transit and not data at rest. The store phase is the first phase in which security controls are implemented to protect data at rest.

upvoted 9 times

[Removed] Highly Voted 3 years, 2 months ago

Create.

as per CBK, Data classification is foundational security control. page 44, cbk 3rd edition.

upvoted 7 times

krauzo Most Recent 1 month, 1 week ago

Selected Answer: D

create - when data is created it can be properly classified

upvoted 1 times

Vee_Wang 3 months, 1 week ago

Selected Answer: B

create stage, you can define the classification . While security control is implemented in store stage.

upvoted 1 times

CBO2025 4 months, 2 weeks ago

Selected Answer: B

When your create data, you're just typing or modifying however, when you submit it, you'll trigger the store phase, hence the secuirt starts at the store phase

upvoted 1 times

Kneebee 11 months, 3 weeks ago

B is the correct choice - the CCSP official study indicates this is the first stage where security controls can be implemented to protect data at rest.

upvoted 2 times

FranklinG 1 year ago

This is a tricky question to try and trip you up. It says "the first stage" making it sound like the first phase in the data life cycle world, which would be "Create."

However, "Store" is the right answer, because in the "Create" phase the data owner is defined, then data is categorized, classified, labeled, tagged and marked. And if created remotely, data should be encrypted, and connections secured (VPN) and secure key management practices should be practiced. Now, in the "Store" phase which occurs almost concurrently with the "Create" phase is where it's immediately important to employ:

The use of backup methods on top of security controls to prevent data loss.

Additional encryption for data at rest. DLP and IRM technologies are used to ensure that data security is enforced during the Use and Share phases of the cloud data lifecycle.

upvoted 2 times

JBvino 1 year, 5 months ago

While security controls are implemented in the create phase in the form of SSL/TLS, this only protects data in transit and not data at rest. The store phase is the first phase in which security controls are implemented to protect data at rest.

upvoted 2 times

kollmekay 1 year, 1 month ago

But the question doesnt mention data in transit

upvoted 1 times

MartijnBdV 1 year, 8 months ago

Selected Answer: D

security controls can be initially implemented at the create phase as well, specifically in the form of technologies such as SSL/TLS with data that is inputted or imported.

upvoted 2 times

escaprix 1 year, 9 months ago

Selected Answer: B

Store no hay duda

upvoted 2 times

Purespace 1 year, 11 months ago

Selected Answer: D

Create - as has been said throughout the comments, data classification and labeling is most certainly a "security control" as defined in NIST SP 800-53, ISO 27001, HITRUST, etc. (look up "information handling" in the control sets).

upvoted 1 times

secisfun 2 years, 3 months ago

Selected Answer: D

Create

upvoted 1 times

kepalon 3 years ago

In Create is the 1st phase where labels can be assign;
In Store is the 1st phase where controls can be implemented.

upvoted 5 times

xaccan 3 years, 11 months ago

Store
The Store phase often happens in tandem with (or immediately after) the Create phase. During this phase, the created or modified data is saved to some sort of digital repository within the application or system. Storage can be in the form of saved files on a filesystem, rows and columns saved to a database, or objects saved in a cloud storage system.
During the Store phase, the classification level assigned during creation is used to assign and implement appropriate security controls. Controls like encryption (at rest), Access Control Lists (ACLs), logging, and monitoring are important during this phase. In addition, this phase is when you should consider how to appropriately back up your data to maintain redundancy and availability

upvoted 4 times

phanil1 3 years, 11 months ago

In Create, you can only define the controls like classification, you cannot apply until you store them. store is correct answer.

upvoted 5 times

asldavid 4 years ago

Should be "D". Data classification is done during the create phase.

upvoted 4 times

Rangakarthik 4 years, 2 months ago

Yes I do agree here

upvoted 1 times

Question #92 *Topic 1*

What controls the formatting and security settings of a volume storage system within a cloud environment?

A. Management plane

B. SAN host controller

C. Hypervisor

D. Operating system of the host

Suggested Answer: *D*
Once a storage LUN is allocated to a virtual machine, the operating system of that virtual machine will format, manage, and control the file system and security of the data on that LUN.

Community vote distribution
A (100%)

bessonf Highly Voted 3 years, 10 months ago
A . this is cloud management plan which allow to define secury level a storage object. The formatting dring by the type of storage object selected
upvoted 5 times

deegadaze1 3 years ago
A is the Correct Answer
upvoted 1 times

lolanczos Most Recent 3 months, 1 week ago
Selected Answer: A
It's A. Direct from the study guide.

(ISC)² CCSP Official Study Guide, Third Edition:

Chapter on Cloud Architecture and Design:
Describes the roles of different cloud planes (management, control, and data planes) and emphasizes that the management plane is responsible for the configuration and management of cloud resources, including storage.
upvoted 1 times

Mo22 5 months, 2 weeks ago
Selected Answer: A
To me A makes the most sense here but I agree very dump question and very poorly formated
upvoted 1 times

usuari000 11 months ago
the host is the hypervisor, the guest is the VM. If the storage is mounted inside the VM, it is not the host who formats that storage.
upvoted 1 times

kepalon 2 years, 6 months ago
I vote D. As well, as the VOLUME is assigned to a VM-OS, and it´s the one that controls the file system in the volume.
upvoted 2 times

[Removed] 2 years, 8 months ago
D.
B & C are ruled out as they are out of control for CC.
You cant format the disk from Management plane,
D is the only reasonable answer.
upvoted 2 times

NobleGiantz 3 years, 7 months ago
C is correct here. For volume storage, the hypervisor takes a chunk of the physical storage infrastructure and virtually assigns it to a VM.
upvoted 2 times

Mobi333 3 years, 9 months ago

Operating system of the host is correct, bcoz volumes assigned to OS and OS is the one which formats the volumes/Lun's

upvoted 2 times

Guivent 3 years, 10 months ago

Yes I think it should be the management plane

upvoted 1 times

Question #93 *Topic 1*

What does SDN stand for within a cloud environment?

A. Software-dynamic networking

B. Software-defined networking

C. Software-dependent networking

D. System-dynamic nodes

Suggested Answer: *B*

Software-defined networking separates the administration of network filtering and network forwarding to allow for distributed administration.

Community vote distribution

B (100%)

MaciekMT 2 months, 1 week ago

Selected Answer: B

Software defined networking (SDN) provides a way to centrally configure and manage networks and network services such as switching, routing, and load balancing in your data center. You can use SDN to dynamically create, secure, and connect your network to meet the evolving needs of your apps.

https://learn.microsoft.com/en-us/azure/azure-local/concepts/software-defined-networking

upvoted 1 times

akg001 4 months, 1 week ago

Selected Answer: B

B. Software-defined networking

upvoted 2 times

Question #94 *Topic 1*

From a legal perspective, what is the most important first step after an eDiscovery order has been received by the cloud provider?

A. Notification

B. Key identification

C. Data collection

D. Virtual image snapshots

Suggested Answer: *A*

The contract should include requirements for notification by the cloud provider to the cloud customer upon the receipt of such an order. This serves a few important purposes. First, it keeps communication and trust open between the cloud provider and cloud customers. Second, and more importantly, it allows the cloud customer to potentially challenge the order if they feel they have the grounds or desire to do so.

Community vote distribution

A (100%)

stack120566 5 months, 3 weeks ago

The phrase " from a legal perspective" indicates that there would be some specific law or regulation in place. if the quest reead ' form the customer acceptance perspective of the cloud service , then te answer would be multitennacy. For me . legal perspective implies regualtrion.

upvoted 1 times

akg001 1 year, 10 months ago

Selected Answer: A

A. Notification

upvoted 1 times

Question #95 *Topic 1*

Which of the following would make it more likely that a cloud provider would be unwilling to satisfy specific certification requirements?

A. Resource pooling

B. Virtualization

C. Multitenancy

D. Regulation

Suggested Answer: *C*
With cloud providers hosting a number of different customers, it would be impractical for them to pursue additional certifications based on the needs of a specific customer. Cloud environments are built to a common denominator to serve the greatest number of customers, and especially within a public cloud model, it is not possible or practical for a cloud provider to alter their services for specific customer demands.

kjjcraigskel Highly Voted 3 years, 11 months ago
Regulations make providers willing, not unwilling. This answer is Multitenancy, as they are unwilling to impose on all customers
upvoted 9 times

Xindydo Highly Voted 3 years, 1 month ago
I am wondering if the question was supposed to read "Which of the following would make it more likely that a cloud provider would be unwilling to satisfy specific CUSTOMER requirements?"

If so, that would make the answer of "multi-tenancy" make complete sense. Otherwise, it does not.
upvoted 6 times

globy118 Most Recent 3 months, 2 weeks ago
Copilot said D.Regulation.
C. Multitenancy: While multitenancy is a key aspect of cloud environments, it doesn't directly impact a cloud provider's willingness to comply with specific certification requirements. Multitenancy refers to multiple customers sharing the same infrastructure, but it doesn't inherently hinder certification efforts.
D. Regulation: This factor tends to have a more substantial influence. Compliance with varying legal requirements and regulations can significantly impact a cloud provider's ability to meet specific certifications. It often involves complex processes and adjustments.
Given this context, the more relevant answer is D. Regulation.
upvoted 1 times

Chungies 1 year, 3 months ago
The wording of the question is not right.
upvoted 1 times

Pegasus_orb 2 years, 9 months ago
I agree with you @ Xindydo
upvoted 2 times

ArizonaClassics 4 years, 8 months ago
REGULATION is the right answer
upvoted 4 times

kjjcraigskel 3 years, 11 months ago
Regulations would not ever make a provide unwilling.
upvoted 3 times

Question #96 *Topic 1*

Which of the following pertains to fire safety standards within a data center, specifically with their enormous electrical consumption?

A. NFPA

B. BICSI

C. IDCA

D. Uptime Institute

Suggested Answer: *A*

The standards put out by the National Fire Protection Association (NFPA) cover general fire protection best practices for any type of facility, but also specific publications pertaining to IT equipment and data centers.

Community vote distribution

A (100%)

MaciekMT 1 month, 1 week ago

Selected Answer: A

The National Fire Protection Association (NFPA) sets fire safety standards for various industries, including data centers, which consume enormous amounts of electricity and pose fire risks. The NFPA 75 and NFPA 76 standards specifically address fire protection and suppression in data centers and telecommunications facilities.

upvoted 1 times

akg001 4 months, 1 week ago

Selected Answer: A

A. NFPA

upvoted 1 times

Question #97 *Topic 1*

Which of the following roles involves the connection and integration of existing systems and services to a cloud environment?

A. Cloud service business manager

B. Cloud service user

C. Cloud service administrator

D. Cloud service integrator

Suggested Answer: *D*

The cloud service integrator is the official role that involves connecting and integrating existing systems and services with a cloud environment. This may involve moving services into a cloud environment, or connecting to external cloud services and capabilities from traditional data center-hosted services.

Community vote distribution

D (100%)

MaciekMT 1 month, 1 week ago

Selected Answer: D

A Cloud Service Integrator is responsible for connecting and integrating existing systems, applications, and services into a cloud environment. This role ensures seamless interoperability between on-premises infrastructure and cloud-based solutions.

Cloud Service Business Manager (A) → Handles financials, contracts, and customer relationships, not technical integration.
Cloud Service User (B) → Simply consumes the cloud services; they don't manage integration.
Cloud Service Administrator (C) → Manages cloud resources and configurations but doesn't necessarily perform system integration.

upvoted 1 times

akg001 2 years, 10 months ago

Selected Answer: D

D. Cloud service integrator

upvoted 2 times

Sa007788 4 years, 2 months ago

This role is not described in ISC2

upvoted 2 times

468fa2a 9 months, 1 week ago

Yep, neither in the 2022 or 2023 version. If they are asking questions not in their own material, but from referenced material, I am on the fence if that is fair. You cant read all the external references, and some you have to pay for too,

upvoted 1 times

Ahbey_911 4 years, 1 month ago

Check "ISO 17789" - Cloud Computing Reference Architecture

upvoted 6 times

Question #98 *Topic 1*

Which technique involves replacing values within a specific data field to protect sensitive data?

A. Anonymization

B. Masking

C. Tokenization

D. Obfuscation

Suggested Answer: *B*

Masking involves replacing specific data within a data set with new values. For example, with credit card fields, as most who have ever purchased anything online can attest, nearly the entire credit card number is masked with a character such as an asterisk, with the last four digits left visible for identification and confirmation.

Community vote distribution

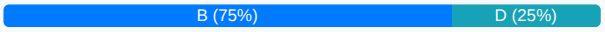

evilwizardington Highly Voted 3 years, 7 months ago

The CCSP Official Guide states: "The term obfuscation refers to the application of any of these techniques in order to make the data less meaningful, detailed, or readable in order to protect the data or the subject of the data."
The guide indicates that obfuscation can be achieved either by masking or anonymizing. And masking has several techniques.

From that explanation: obfuscation does not always mean replacing values (it can be anonymized the information, meaning, removing the information). And tokenization may be an option, but a more general answer is masking.

That's why the most accurate answer is masking.
upvoted 14 times

Ahbey_911 Highly Voted 3 years, 8 months ago

This is another case where the examiner decides what they want as the answer, all of the answers can be correct:
Data masking or Obfuscation is a process of hiding, replacing, or omitting sensitive
information e.g. PII, PHI, PCI
Data Anonymization is a technique for information sanitization with an intent to protect
privacy.
Tokenization is substituting sensitive information with non-sensitive information
upvoted 8 times

Ramye Most Recent 1 week, 5 days ago

Selected Answer: C

Ans is C Tokenization. Other options including masking does not replace data, masking obscure the data so it's not replaced so this can't be a answer.
upvoted 1 times

MaciekMT 1 month, 1 week ago

Selected Answer: B

Masking involves replacing values within a specific data field to protect sensitive data while keeping it usable for testing or display. It ensures that sensitive data, such as credit card numbers or social security numbers, is hidden from unauthorized users.

Anonymization (A) → Irreversibly removes personally identifiable information (PII) to prevent re-identification.
Tokenization (C) → Replaces sensitive data with a non-sensitive equivalent (token) that can be mapped back to the original data through a secure token vault.
Obfuscation (D) → Scrambles or alters data to make it difficult to understand but does not follow structured masking techniques.
upvoted 1 times

globy118 3 months, 2 weeks ago

Selected Answer: B

Obfuscation involves intentionally making code or data more complex or unclear. While it can provide some protection, it's not specifically designed for sensitive data.

upvoted 1 times

Mo22 6 months ago

Selected Answer: B

While "D. Obfuscation" is also a technique used to protect sensitive data, it generally refers to making data unclear or difficult to understand through various methods, rather than specifically replacing values within a data field. Obfuscation can include techniques like encryption, scrambling, or using dummy data.

upvoted 1 times

FranklinG 7 months ago

B. The answer is Masking.

Data obfuscation is the blanket term for transforming data into a different form to protect it. There are three main types of data obfuscation: data masking, tokenization, and encryption.

Data masking creates a substitute version of a dataset. The data values are changed, but the format remains the same. Remember we are looking for the BEST answer.

upvoted 1 times

Kneebee 11 months, 1 week ago

The correct answer is B (see AWS definition: https://aws.amazon.com/what-is/data-masking/)

upvoted 1 times

BuckLee 11 months, 3 weeks ago

Selected Answer: D

Token the field

upvoted 1 times

Pika26 1 year, 4 months ago

Selected Answer: B

B. Masking

upvoted 1 times

xroxro 2 years, 1 month ago

I think we have to stick to CCSP classification :

Data security strategies :

• Obfusaction

○ Data anonymization

○ Data masking

- Substitution

- Scrambling

- Deletion or nulling

•tokenization

•homomorphic encryption

• bit splitting

In our case, "replacing value with specific data" correspond to "data masking with substitution"

Answer B

upvoted 1 times

GregP 2 years, 2 months ago

I'm gonna say B because it says "within a specific data field" which makes masking the best answer.

upvoted 3 times

kepalon 2 years, 6 months ago

Masking is the best answer. IN this cases you are suppose to provide the best answer.

I agree though that they are similar concepts, but each one has its best definition.

upvoted 2 times

Warriors 2 years, 11 months ago

Tokenization is the most appropriate answer as there are more than one ways to mask or obfuscate data that includes tokenization. However, tokenization is the process of substituting a sensitive data with a non-sensitive data (token).

upvoted 3 times

Sa007788 3 years, 8 months ago

Masking is same as Obfuscation. This is question is less accurate

upvoted 1 times

Irivera 3 years, 8 months ago

answer is B masking

the other three are specific types of masking

upvoted 2 times

HCL 3 years, 10 months ago

I mean why Obfuscation is not the answer?

Obfuscation is the same as Masking, isn't it?

upvoted 3 times

Question #99 *Topic 1*

What expectation of data custodians is made much more challenging by a cloud implementation, especially with PaaS or SaaS?

A. Data classification

B. Knowledge of systems

C. Access to data

D. Encryption requirements

Suggested Answer: *B*
Under the Federal Rules of Civil Procedure, data custodians are assumed and expected to have full and comprehensive knowledge of the internal design and architecture of their systems. In a cloud environment, especially with PaaS and SaaS, it is impossible for the data custodian to have this knowledge because those systems are controlled by the cloud provider and protected as proprietary knowledge.

Community vote distribution

guest999 Highly Voted 4 years, 3 months ago
The term Data Custodian, is being used with different meaning by different folks. In general the Cloud Customer (the consumer of Cloud services) is the data owner in all types IaaS, PaaS, SaaS. They are responsible for the data. The Cloud provider is the Data Custodian. In this question the Data Custodian seems to refer to the Data Owners, the Cloud Customer!.
upvoted 9 times

xaccan 2 years, 11 months ago
Data custodian could be Database administrator, so the given answer is very correct.
upvoted 3 times

HCL 3 years, 10 months ago
Agree, so it is very confusing.
upvoted 2 times

evilwizardington 3 years, 7 months ago
I think you are seeing in the wrong perspective this. Data Owner is the main person responsible and who utilizes the information. And Data Custodian is the one which gives maintenance, administration, and set controls over the data. In this case, the custodian would never be the cloud vendor.

For example: for employees data, the Owner is HR, the Custodian is the security team or IT, and the cloud vendor acts as processor/subprocessor.
upvoted 12 times

MaciekMT Most Recent 1 month, 1 week ago
Selected Answer: B
Data custodians are typically expected to have detailed knowledge of the systems that store and process their data—this includes how the data is hosted, secured, and maintained. In a PaaS or SaaS model, however, much of the underlying infrastructure, configuration, and even some application components are completely opaque to the customer.
upvoted 1 times

MaciekMT 2 months, 1 week ago
Selected Answer: B
Knowledge of systems

Explanation:
In a cloud environment, particularly with PaaS or SaaS, the underlying infrastructure and systems are managed by the cloud provider. This makes it difficult for data custodians to gain detailed knowledge of the systems hosting their data. They are often abstracted from the hardware, software configurations, and other operational details, which complicates tasks such as:

Assessing security risks.
Verifying compliance.

Understanding system-specific behaviors.
This abstraction layer poses a challenge for custodians to fully understand and manage the environment compared to traditional on-premises systems.
upvoted 1 times

cloudenthusiast 7 months ago
Selected Answer: B
Cloud providers do not provide the information of their system to clients.
upvoted 1 times

Pika26 1 year, 4 months ago
Selected Answer: A
A: Data classification
upvoted 1 times

secisfun 1 year, 9 months ago
C. Access to data
upvoted 3 times

DA95 1 year, 9 months ago
C. Access to data is made more challenging by a cloud implementation, especially with PaaS or SaaS. In these scenarios, the cloud provider typically manages access to data and data storage, which can make it more difficult for data custodians to control who has access to the data and how it is used.
upvoted 2 times

akg001 2 years, 4 months ago
Selected Answer: B
B. Knowledge of systems
upvoted 2 times

NobleGiantz 3 years, 7 months ago
A data custodian or data processor processes the data on behalf of the data owner. The data custodian is responsible for adhering to the data owner's established requirements for using and securing the data, and must process the data in accordance with the data owner's established purposes.
upvoted 2 times

Question #100 *Topic 1*

What type of PII is controlled based on laws and carries legal penalties for noncompliance with requirements?

A. Contractual

B. Regulated

C. Specific

D. Jurisdictional

Suggested Answer: *B*

Regulated PII involves those requirements put forth by specific laws or regulations, and unlike contractual PII, where a violation can lead to contractual penalties, a violation of regulated PII can lead to fines or even criminal charges in some jurisdictions. PII regulations can depend on either the jurisdiction that applies to the hosting location or application or specific legislation based on the industry or type of data used.

Community vote distribution

MaciekMT 1 month, 1 week ago

Selected Answer: B

"Regulated PII" refers to personally identifiable information that is protected by specific laws (e.g., HIPAA, GDPR, FERPA) and imposes legal penalties for noncompliance. While all PII needs some level of protection, regulated PII must adhere to specific statutory or regulatory requirements, making it subject to legal enforcement.

upvoted 1 times

akg001 4 months, 1 week ago

Selected Answer: B

B. Regulated

upvoted 1 times

kns20 1 year, 1 month ago

Difference Between Contractual and Regulated Private Data
Contractual and regulated data may coexist within a single complementary context; a contract may be formulated to enforce the adherence to a regulation or set of regulations and a regulation may define the need to have contractual relationships between provider and consumer. The intent of a contract is to provide for a legally binding instrument that governs the acts, expectations, and behaviors between two or more parties. A regulation is typically confined to a specific industry or process that involves a provider and consumer (although the term regulation is used occasionally tied to laws).

upvoted 3 times

Question #101 *Topic 1*

Which if the following is NOT one of the three components of a federated identity system transaction?

A. Relying party

B. Identity provider

C. User

D. Proxy relay

Suggested Answer: *D*

Community vote distribution

D (100%)

NobleGiantz Highly Voted 2 years, 7 months ago

D is correct

upvoted 7 times

MaciekMT Most Recent 1 month, 1 week ago

Selected Answer: D

Typical federated identity transactions include three main players:

Identity Provider (IdP) – Manages the identity information and authenticates the user.
Relying Party (RP) – The service or application that depends on the IdP for authentication.
User – The individual who needs to access the relying party's service.
A proxy relay isn't one of these standard three components in a typical federated identity scenario.

upvoted 1 times

bp339 5 months, 2 weeks ago

Selected Answer: D

Proxy relay is not one of the three components of a federated identity system transaction. The three components of a federated identity system transaction are the user, the identity provider, and the relying party. The user is the person who is trying to access a resource, the identity provider is the entity that authenticates the user and provides identity information, and the relying party is the entity that is requesting the identity information and making the access decision. A proxy relay is not typically considered a component of a federated identity system transaction, although it may be used in some implementations to help facilitate the transaction.

upvoted 2 times

Marcelop 2 years, 10 months ago

If the federation solution has to be exposed to internet, then a proxy is needed.

upvoted 2 times

Question #102 *Topic 1*

Which value refers to the amount of time it takes to recover operations in a BCDR situation to meet management's objectives?

A. RSL

B. RPO

C. SRE

D. RTO

Suggested Answer: *D*

The recovery time objective (RTO) is a measure of the amount of time it would take to recover operations in the event of a disaster to the point where management's objectives are met for BCDR.

Community vote distribution

D (100%)

bp339 5 months, 2 weeks ago

Selected Answer: D

Recovery Time Objective (RTO)

upvoted 1 times

akg001 1 year, 4 months ago

Selected Answer: D

D. RTO

upvoted 1 times

Question #103 *Topic 1*

Which of the cloud deployment models requires the cloud customer to be part of a specific group or organization in order to host cloud services within it?

A. Community

B. Hybrid

C. Private

D. Public

Suggested Answer: *A*

A community cloud model is where customers that share a certain common bond or group membership come together to offer cloud services to their members, focused on common goals and interests.

Community vote distribution

A (100%)

akg001 4 months, 1 week ago

Selected Answer: A

A. Community

upvoted 1 times

Question #104 *Topic 1*

What provides the information to an application to make decisions about the authorization level appropriate when granting access?

A. User

B. Relying party

C. Federation

D. Identity Provider

Suggested Answer: *D*

Upon successful user authentication, the identity provider gives information about the user to the relying party that it needs to make authorization decisions for granting access as well as the level of access needed.

Community vote distribution

D (100%)

keresh Highly Voted 2 years, 3 months ago

Application is the relaying party in the context of federation. The Identity Provider passes the information to the relaying party, which is the application.

That's why D is correct

upvoted 5 times

Brittle 1 year, 1 month ago

Thanks

upvoted 1 times

MaciekMT Most Recent 1 month, 1 week ago

Selected Answer: D

In a federated identity scenario, the Identity Provider (IdP) supplies an assertion or token containing user identity data (often called claims). The application (relying party) then uses this information to make authorization decisions (i.e., what the user is allowed to do).

upvoted 1 times

hogancl42004 6 months ago

My thought is B, "The relying party is any member of the federation that shares resources based on authenticated identities. Relying parties then handle authorization based on their policies. This allows a relying party to determine their level of trust in third-party IdPs and to map permissions on their own rather than relay on the IdP to provide both authentication and authorization. " From the CCSP Official Study Guide, Third Edition pg 181

upvoted 2 times

GH1982 1 year, 11 months ago

A relaying party may authorize a user's request based on authorization attributes fetched from an IdP. Examples of authorization attributes include permissions/privileges assigned to the user or the user's role. IdP provides the attributes, and answer is D.

upvoted 1 times

kepalon 2 years ago

Selected Answer: D

it is the right one

upvoted 3 times

kepalon 2 years ago

Really confusing, as the Authorization is done by the "Relying Party" who is the one that needs to provide the Authrorization.
The identity provider, passes the identification+Authentication.
I understand your point, and I like the idea of thinking the Application = Relying Party, that way it is easier to point to the Identity Provider as the right answer.

upvoted 2 times

sans1241 2 years, 11 months ago

The question talking about which system gives the tokens to be consumed by relying party/application.

upvoted 2 times

pooppants 3 years ago

I would have said B. The Identity Provider doesnt touch the application. The information comes from the relying party?

upvoted 1 times

Question #105 *Topic 1*

What is a standard configuration and policy set that is applied to systems and virtual machines called?

A. Standardization

B. Baseline

C. Hardening

D. Redline

Suggested Answer: *B*

The most common and efficient manner of securing operating systems is through the use of baselines. A baseline is a standardized and understood set of base configurations and settings. When a new system is built or a new virtual machine is established, baselines will be applied to a new image to ensure the base configuration meets organizational policy and regulatory requirements.

Community vote distribution

B (100%)

gayan237 9 months, 2 weeks ago

baseline is the 'minimum' standard. this question needs proper explanation.

upvoted 1 times

kepalon 2 years ago

Selected Answer: B

baseline is correct

upvoted 2 times

Question #106 *Topic 1*

Which entity requires all collection and storing of data on their citizens to be done on hardware that resides within their borders?

A. Russia

B. France

C. Germany

D. United States

Suggested Answer: *A*

Signed into law and effective starting on September 1, 2015, Russian Law 526-FZ establishes that any collecting, storing, or processing of personal information or data on Russian citizens must be done from systems and databases that are physically located with the Russian Federation.

Community vote distribution

A (100%)

Ashishtv Highly Voted 2 years, 8 months ago

Which domain is this mentioned in

upvoted 5 times

MaciekMT Most Recent 1 month, 1 week ago

Selected Answer: A

Russia's data residency laws (Federal Law No. 242-FZ) require that all collection and storage of personal data about Russian citizens must be done on servers located within Russia's borders.

upvoted 1 times

sweetykaur 5 months, 2 weeks ago

Germany. Germany has strict data protection laws that require all collection and storage of data on its citizens to be done on hardware within its borders.

upvoted 1 times

kepalon 3 years ago

Selected Answer: A

Russia is the correct answer

upvoted 2 times

Banzaaai 3 years, 6 months ago

A, Russia.

upvoted 2 times

Question #107 *Topic 1*

Which of the cloud cross-cutting aspects relates to the ability to easily move services and applications between different cloud providers?

A. Reversibility

B. Availability

C. Portability

D. Interoperability

Suggested Answer: *C*

Portability is the ease with which a service or application can be moved between different cloud providers. Maintaining portability gives an organization great flexibility between cloud providers and the ability to shop for better deals or offerings.

Community vote distribution

C (100%)

BuckLee 5 months, 3 weeks ago

Selected Answer: C

Similarly, organizations should strive to avoid vendor lock-in whenever possible. Portability is a design principle that says workloads should be designed so that they don't leverage vendor-specific features and may be more easily shifted between cloud providers. This isn't always possible, but it is a good design practice.

(ISC)2 CCSP Certified Cloud Security Professional Official Study Guide (p. 22).

upvoted 1 times

nelombg 11 months, 2 weeks ago

The answer is C interoperability

upvoted 1 times

serget12 1 year, 5 months ago

Interoperability defines how easy it is to move and reuse application components regard□less of the provider, platform, OS, infrastructure, location, storage, format of data or APIs,
how well applications work together, and how well new applications work with other
solutions present in the business, organization, or provider’s existing architecture.

upvoted 2 times

akg001 1 year, 10 months ago

Selected Answer: C

C. Portability

upvoted 2 times

certifiedgeek 1 year, 10 months ago

The question is about "moving between cloud providers" (portability) and not retrieve&delete data from a cloud provider (reversibility) nor the ability for apps to work seamlessly with providers (interoperability).

upvoted 2 times

xav1er 1 year, 11 months ago

Its Portability. From CCSP Students Guide: Portability defines the ease with which application components are
moved and reused elsewhere regardless of the provider, platform,
OS, infrastructure, location, storage, format of data, or APIs.
Portability is a key aspect to consider when selecting cloud
providers, since it can both help prevent vendor lock-in and deliver
business benefits by allowing identical cloud deployments to occur
in different cloud provider solutions, either for the purposes of
disaster recovery or for the global deployment of a distributed single
solution.

upvoted 1 times

kepalon 2 years ago

C is the right Answer.

Reversibility refers to the possibility to remove ALL your data from the cloud; after moving it to another location first, of course.

There is a thinner line between Reversibility and Interoperability though, the second one refers to the possibility to work with more than one cloud solution.

upvoted 1 times

[Removed] 2 years, 2 months ago

Per CBK 3rd, page 18.

Portability refers to data or Architecture portability.

Reversibility is a measure of the extent your cloud services can be moved from one cloud to another.

A is the answer

upvoted 2 times

zaqwsx 1 year, 11 months ago

I will also go with A

I found

Reversibility is the extent to which cloud-based applications are designed so that they can be moved to other cloud providers or environments.

upvoted 1 times

Question #108 *Topic 1*

Which type of audit report is considered a "restricted use" report for its intended audience?

A. SAS-70

B. SSAE-16

C. SOC Type 1

D. SOC Type 2

Suggested Answer: *C*
SOC Type 1 reports are considered "restricted use" reports. They are intended for management and stakeholders of an organization, clients of the service organization, and auditors of the organization. They are not intended for release beyond those audiences.

Community vote distribution

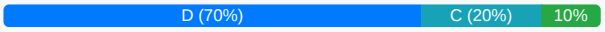

ragar123 Highly Voted 4 years, 3 months ago
the question doesn't seem correct The SOC 2 Type 1 is
not extremely useful for determining the security and trust of an organization. The SOC 2
Type 1 only reviews the design of controls, not how they are implemented and maintained,
or their function. The SOC 2 Type 2 report, however, does just that. This is why the SOC 2
Type 2 is the sort of report that is extremely useful for getting a true assessment of an organization's security posture.
upvoted 21 times

MartinRB Most Recent 1 month, 2 weeks ago
Selected Answer: C
A SOC Type 1 report (Service Organization Control Type 1) is considered a restricted use report because it is intended for a specific audience, typically management, auditors, and regulators. It focuses on the design and implementation of internal controls at a point in time, rather than ongoing operational effectiveness.
upvoted 1 times

Mo22 5 months, 2 weeks ago
Selected Answer: C
SOC Type 1 reports are often intended for restricted use, meaning they are designed for specific, intended users, such as management or those charged with governance, not for the general public or broader external use. They evaluate the design of a service organization's controls at a specific point in time.
upvoted 2 times

cloudenthusiast 7 months ago
Selected Answer: D
SOC2 TYPE 2
upvoted 1 times

ccKane 9 months, 2 weeks ago
Selected Answer: B
SSAE-16 includes SOC1 and SOC2. Both are restricted. "SSAE-16, which stands for Statement on Standards for Attestation Engagements No. 16, was introduced by the AICPA as a replacement for SAS-70. SSAE-16 introduced several changes and improvements to the auditing and reporting process for service organizations, particularly for those providing services that could impact their clients' financial reporting. SSAE-16 is part of a broader framework for attestation engagements, including SOC (Service Organization Control) reports (SOC1 and SOC2)."
upvoted 1 times

JohnnyBG 7 months, 4 weeks ago
SSAE is an auditing standard, not a report by itself. My take is SOC 2 is then answer.
upvoted 3 times

nirlion 1 year, 4 months ago

You may get type 2 reports, but never type 1 report (soc 1 or 2 does not matter). Type 1 reports are always classified with no exceptions in real life since it pertains to a “specific time” as against type 2. Ask any auditor friend, they will tell.
upvoted 1 times

Pika26 1 year, 5 months ago

Selected Answer: D

SOC Type 2 reports include a description of the service organization's system, a detailed testing of the design and operating effectiveness of controls, and an opinion provided by an independent auditor.
upvoted 1 times

LearnsNow 1 year, 8 months ago

Selected Answer: D

SOC 2 Type 2 is the correct answer.
upvoted 1 times

Lenell 1 year, 9 months ago

Selected Answer: C

Type 1 report just provides a report of procedures / controls an organization has put in place as of a point in time (no required audit so no outside audience; i.e., more restrictive). A Type 2 report has an audit period and provides evidence of how an organization operated its controls over a period of time (required audit so outside audience; i.e., less restrictive). Restrictive is observed from the perspective of the data owner's view.
upvoted 1 times

DA95 1 year, 9 months ago

A SOC Type 2 audit report is considered a "restricted use" report for its intended audience. SOC, or Service Organization Controls, is a set of auditing standards and guidelines developed by the American Institute of Certified Public Accountants (AICPA) to help service organizations demonstrate the effectiveness of their internal controls and processes. A SOC Type 2 audit report is a detailed assessment of a service organization's controls over a specific period of time, typically six to nine months. Because this report contains sensitive information about the organization's internal controls and processes, it is considered a "restricted use" report and is only intended for the organization's management, board of directors, and other stakeholders who have a need to know the information contained in the report.
upvoted 4 times

DERCHEF2009 1 year, 11 months ago

Selected Answer: D

Yes it is D
upvoted 4 times

certifiedgeek 2 years, 4 months ago

Request to update the choices as both "SOC Type 1" and "SOC Type 2" (whether SOC1 or SOC2) are both restricted to their intended users. Also SOC3 (which does not have any type) are for public use.
upvoted 3 times

kepalon 2 years, 6 months ago

Note that it does not mention SOC 1 or SOC 2, but Type 1 & Type 2.
There is something worng with this question.
Type 2 is in a period of time - 6 months
Type 1 is in a specific time, when the control/design was checked.
upvoted 1 times

keresh 2 years, 10 months ago

both are restricted but SOC1 is more restrictive

SOC 1 - Use of these reports is restricted to the management of the service organization, user entities, and user auditors.

SOC 2 - Use of these reports are restricted.

Taken from https://us.aicpa.org/interestareas/frc/assuranceadvisoryservices/serviceorganization-smanagement
upvoted 2 times

Warriors 2 years, 11 months ago

SOC 2 is the right answer
upvoted 1 times

serget12 1 year, 11 months ago

Correct, but either type 1 or 2 would fall under SOC 2. So the answer could be either one.

upvoted 1 times

xaccan 2 years, 11 months ago

They mean SoC 1 which is true, which is a control report that focuses strictly on an organization's financial statements and a service organization's controls that can impact a customer's financial statements

upvoted 3 times

Ahbey_911 3 years, 7 months ago

The options include SOC Type I & II, not SOC 2 Type II.
SOC Type I - provides a description of the controls provided by the audited organization and the auditor opinion based on the description, BUT... does not involve actual testing of controls. SOC Type 1 reports are intended for restricted use, only to be seen by the actual service organization, its current clients, or its auditors. These reports are not intended for wider or public distribution.
So, the answer is correct folks.

upvoted 2 times

evilwizardington 3 years, 7 months ago

First of all, there's no SOC Type I, and SOC Type 2. SOC 1 does not have both versions. Only SOC 2.

Under such premise, SOC 2 in any of its forms is intended only for restricted use. The only one for a wider audience its the SOC 3 report.

So I agree, the question or answers are incorrect.

upvoted 2 times

Question #109 *Topic 1*

What is the concept of segregating information or processes, within the same system or application, for security reasons?

A. fencing

B. Sandboxing

C. Cellblocking

D. Pooling

Suggested Answer: *B*

Sandboxing involves segregating and isolating information or processes from others within the same system or application, typically for security concerns. This is generally used for data isolation (for example, keeping different communities and populations of users isolated from other similar data).

Community vote distribution

B (100%)

cmarcos97 Highly Voted 2 years, 6 months ago

Disagree, Sandboxing can mean many things today, but in the realm of cloud computing, sandboxing refers to the concept of a protected area being utilized for testing untested or untrusted code or to better understand if an application is working the way it was intended to work. These sandboxes are usually protected areas in memory that will not allow processes of any kind to run outside the environment or allow access inside from any other application or process.

upvoted 6 times

Zeezee2 2 years, 4 months ago

What you see as sandbox is in the ccsp world considered an on-premise sandbox. In the cloud world, sandboxing simply means logically separating one or more environments based on eg. data classification in order to ensure different levels of security at each level.

upvoted 4 times

nelombg Most Recent 11 months, 3 weeks ago

Weird question.

upvoted 4 times

Kilani 4 months, 1 week ago

Really weird, sandboxing is nowadays referenced with advanced threat protection where you create an environment isolated to measure the damage that this threat has to it.

upvoted 1 times

akg001 1 year, 10 months ago

Selected Answer: B

B. Sandboxing

upvoted 1 times

Question #110 *Topic 1*

The European Union passed the first major regulation declaring data privacy to be a human right. In what year did it go into effect?

A. 2010

B. 2000

C. 1995

D. 1990

Suggested Answer: *C*

Adopted in 1995, Directive 95/46 EC establishes strong data protection and policy requirements, including the declaring of data privacy to be a human right. It establishes that an individual has the right to be notified when their personal data is being access or processed, that it only will ever be accessed for legitimate purposes, and that data will only be accessed to the exact extent it needs to be for the particular process or request.

Community vote distribution

C (100%)

MaciekMT 1 month, 1 week ago

Selected Answer: C

The European Union's Directive 95/46/EC (often referred to as the Data Protection Directive) went into effect in 1995 and famously recognized privacy as a fundamental human right. It laid the groundwork for data protection principles later incorporated into the General Data Protection Regulation (GDPR).

upvoted 1 times

Lenell 9 months ago

Selected Answer: C

The directive was repealed to the regulation in 2018 (see https://gdpr-text.com/read/article-94/). But all that is irrelevant. In the context of the ISC2 question, their answer is rightfully 1995.

upvoted 2 times

pwxfuchxaubgkfcqay 1 year, 3 months ago

Selected Answer: C

GDPR is a regulation.. 95/46 was only a directive... but doesn't matter in this context: C is correct

upvoted 2 times

Lenell 9 months ago

The directive was repealed to the regulation in 2018 (see https://gdpr-text.com/read/article-94/). But all that is irrelevant. In the context of the ISC2 question, their answer is rightfully 1995.

upvoted 1 times

akg001 1 year, 4 months ago

Selected Answer: C

C. 1995

upvoted 1 times

Question #111 *Topic 1*

Which of the following is NOT a key area for performance monitoring as far as an SLA is concerned?

A. CPU

B. Users

C. Memory

D. Network

Suggested Answer: *B*

An SLA requires performance monitoring of CPU, memory, storage, and networking. The number of users active on a system would not be part of an SLA specifically, other than in regard to the impact on the other four variables.

Community vote distribution

B (100%)

akg001 4 months, 1 week ago

Selected Answer: B

B. Users

upvoted 2 times

Question #112 *Topic 1*

Which of the following is the MOST important requirement and guidance for testing during an audit?

A. Stakeholders

B. Shareholders

C. Management

D. Regulations

Suggested Answer: *D*

During any audit, regulations are the most important factor and guidelines for what must be tested. Although the requirements from management, stakeholders, and shareholders are also important, regulations are not negotiable and pose the biggest risk to any organization for compliance failure.

Community vote distribution

D (100%)

MaciekMT 4 weeks, 1 day ago

Selected Answer: D

During an audit, testing must adhere to the most important requirement and guidance, which comes from applicable regulations and compliance frameworks. Regulations dictate what must be tested, how it should be conducted, and the standards that must be met.

While stakeholders, shareholders, and management are important, they do not provide direct testing requirements—they may set priorities or expectations, but regulatory compliance is the primary driver of audit testing.

upvoted 1 times

certifiedgeek 4 months, 2 weeks ago

Regulations requirements are mandatory if compliance is a must. These can get the organization fined or lose it compliance/certifications. Requirements from management, stakeholders, and shareholders can have different levels of scope which might or might not include regulations requirements.

upvoted 3 times

kepalon 6 months, 1 week ago

Selected Answer: D

REGULATIONS

upvoted 2 times

S_h_a_h 1 year, 6 months ago

Answer D is correct as listed

upvoted 1 times

Sa007788 1 year, 8 months ago

For me answer is C because Management will include also Regualtory requirement and may have more internal requirement

upvoted 2 times

xaccan 11 months, 2 weeks ago

read again the quetion.

upvoted 2 times

Question #113 *Topic 1*

Which value refers to the amount of data an organization would need to recover in the event of a BCDR situation in order to reach an acceptable level of operations?

A. SRE

B. RTO

C. RPO

D. RSL

Suggested Answer: *C*

The recovery point objective (RPO) is defined as the amount of data a company would need to maintain and recover in order to function at a level acceptable to management. This may or may not be a restoration to full operating capacity, depending on what management deems as crucial and essential.

Community vote distribution

MaciekMT 4 weeks, 1 day ago

Selected Answer: C

Recovery Point Objective (RPO) refers to the amount of data an organization would need to recover in the event of a Business Continuity and Disaster Recovery (BCDR) situation to reach an acceptable level of operations. It represents the maximum tolerable amount of data loss measured in time.

For example, if an organization has an RPO of 4 hours, it means that in the event of a failure, they must be able to restore data from backups no older than 4 hours to meet their operational needs.

Other options:
A. SRE (Site Reliability Engineering) – This is a discipline focused on improving system reliability, not a metric for data recovery.
B. RTO (Recovery Time Objective) – This defines the maximum downtime allowed before operations must be restored, but it does not specify how much data needs to be recovered.
D. RSL (Recovery Service Level) – This is not a commonly used term in BCDR planning.

upvoted 1 times

Mo22 5 months ago

Selected Answer: C

Yeah, updating my answer RPO is the correct answer

upvoted 1 times

Mo22 6 months ago

Selected Answer: D

D. RSL (Recovery Service Level)

In the context of Business Continuity and Disaster Recovery (BCDR), the value that refers to the amount of data an organization would need to recover in order to reach an acceptable level of operations is known as the Recovery Service Level (RSL).

upvoted 2 times

inmelet2 1 year, 4 months ago

Selected Answer: C

A. SRE Site Reliability Engineer
B. RTO Recovery Time Objective
C. RPO Recovery Point Objective
D. RSL Recovery Service Level

Answer C

upvoted 1 times

akg001 2 years, 4 months ago

Selected Answer: C

C. RPO

upvoted 3 times

Question #114 *Topic 1*

What must SOAP rely on for security?

A. Encryption

B. Tokenization

C. TLS

D. SSL

Suggested Answer: *A*
Simple Object Access Protocol (SOAP) uses Extensible Markup Language (XML) for passing data, and it must rely on the encryption of those data packages for security.

Community vote distribution
A (57%) C (43%)

MaciekMT 1 month, 1 week ago

Selected Answer: C
SOAP (Simple Object Access Protocol) does not include built-in security features, so it typically relies on transport-layer protection—most commonly TLS (Transport Layer Security)—to safeguard data in transit. While SSL was used in older implementations, it's largely deprecated, and TLS is the modern standard for securing SOAP communications.
upvoted 1 times

HW4301 4 months, 1 week ago
A. SOAP uses SSL/TLS protocols to secure communication channels between clients and servers.
upvoted 1 times

sweetykaur 5 months, 2 weeks ago
SSL (Secure Sockets Layer). SOAP (Simple Object Access Protocol) relies on SSL for ensuring secure communication and data transfer.
upvoted 2 times

globy118 9 months, 3 weeks ago

Selected Answer: C
SOAP itself does not inherently rely solely on encryption; it can work over both encrypted and unencrypted channels.
upvoted 2 times

Mo22 1 year ago

Selected Answer: C
C. TLS (Transport Layer Security)

for security. TLS is a protocol that provides encryption and secure communication over a network. While SOAP can use other security measures like encryption and tokenization, it typically relies on TLS (or its predecessor, SSL) to ensure secure communication between client and server.
upvoted 3 times

kepalon 3 years ago

Selected Answer: A
encryption is right.
SOAP does not provide encryption, so it needs the extra security
upvoted 4 times

Question #115 *Topic 1*

Which of the following is a commonly used tool for maintaining system configurations?

A. Maestro

B. Orchestrator

C. Puppet

D. Conductor

Suggested Answer: *C*
Puppet is a commonly used tool for maintaining system configurations based on policies, and done so from a centralized authority.

Community vote distribution
C (100%)

ichnos Highly Voted 2 years, 6 months ago
C is correct

Two very popular tools for maintaining system configurations and versioning of software are Puppet (https://puppet.com/) and Chef (https://www.chef.io/chef/).
upvoted 5 times

ggx Most Recent 6 months ago
Maestro, Puppet, Chef, Conductor are all orchestrators
upvoted 2 times

akg001 10 months, 2 weeks ago
Selected Answer: C
C. Puppet
upvoted 1 times

NileshGavali 2 years, 8 months ago
isnt this a specific questions to a Specific operating system? and not in general!
upvoted 4 times

Question #116 *Topic 1*

What type of data does data rights management (DRM) protect?

A. Consumer

B. PII

C. Financial

D. Healthcare

Suggested Answer: *A*

DRM applies to the protection of consumer media, such as music, publications, video, movies, and soon.

Community vote distribution

A (100%)

Ola4real Highly Voted 11 months, 2 weeks ago

DRM is Digital Right Management and it's designed limit distribution and unauthorized copying of content. So A is correct but it's not Data Right Management

upvoted 8 times

akg001 Most Recent 4 months, 1 week ago

Selected Answer: A

A. Consumer

upvoted 1 times

Question #117 *Topic 1*

Which type of testing uses the same strategies and toolsets that hackers would use?

A. Penetration

B. Dynamic

C. Static

D. Malicious

Suggested Answer: *A*

Penetration testing involves using the same strategies and toolsets that hackers would use against a system to discovery potential vulnerabilities.

Community vote distribution

A (100%)

akg001 4 months, 1 week ago

Selected Answer: A

A. Penetration

upvoted 2 times

Question #118 *Topic 1*

From a security perspective, which of the following is a major concern when evaluating possible BCDR solutions?

A. Access provisioning

B. Auditing

C. Jurisdictions

D. Authorization

Suggested Answer: *C*

When a security professional is considering cloud solutions for BCDR, a top concern is the jurisdiction where the cloud systems are hosted. If the jurisdiction is different from where the production systems are hosted, they may be subjected to different regulations and controls, which would make a seamless BCDR solution far more difficult.

Community vote distribution

C (100%)

kepalon 6 months, 1 week ago

Selected Answer: C

Jurisdiction is the right one!!! C:

For Pegasus - this is the type of confusion on the ISC2 questions...

upvoted 3 times

akg001 9 months, 1 week ago

I would like to vote : C

upvoted 1 times

Pegasus_orb 9 months, 1 week ago

Instead of saying "BCDR solutions" why not say something like "BCDR site location"

upvoted 4 times

Question #119 *Topic 1*

Which of the following is NOT a focus or consideration of an internal audit?

A. Certification

B. Design

C. Costs

D. Operational efficiency

Suggested Answer: *A*

In order to obtain and comply with certifications, independent external audits must be performed and satisfied. Although some testing of certification controls can be part of an internal audit, they will not satisfy requirements.

Community vote distribution

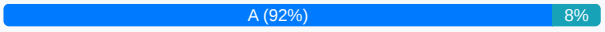

kjjcraigskel Highly Voted 4 years, 5 months ago

Costs is a focus of audit? disagree.

upvoted 13 times

axman832005 3 years ago

I'm with ya, internal audit is ur own department.. why would cost matter

upvoted 2 times

CptSweatbread 2 years, 7 months ago

A financial audit performed by the internal audit team with look at costs/budget line items

upvoted 1 times

Treebeard88 1 year, 2 months ago

Costs are the number one limitation for any company. They won't waste money just for the fun of it whether it's internal or external.

If it's an internal audit you aren't attempting to achieve any accreditation or certification so for me it would be certification.

upvoted 2 times

Zeezee2 Highly Voted 3 years, 4 months ago

Selected Answer: A

Speaking from experience doing security audits at clients, internal audits do NOT result in an official certification of any kind. See it as a way of 1) finding gaps and fixing them by a certain date by which an external audit will actually perform the audit and provide certification.

upvoted 7 times

Zeezee2 3 years, 4 months ago

And to add to this, everything has costs to it. Internal audit = blocking personnel from doing their day to day to answer to questions, to go through documentation, have people involved internally who actually do the audit, sometimes aided by consulting firms who work from an internal audit perspective as well... All these things cost $$€€

upvoted 6 times

MaciekMT Most Recent 1 month, 1 week ago

Selected Answer: A

An internal audit typically focuses on verifying the design and effectiveness of internal controls, operational efficiency, and sometimes cost considerations (e.g., financial compliance). Certification, on the other hand, is usually part of an external or third-party audit process, aiming to validate that the organization meets a recognized external standard (e.g., ISO, SOC, etc.).

upvoted 1 times

globy118 9 months, 1 week ago

Selected Answer: D

in my work, i don't think Internet audit focus on operation efficiency, our SOP will increase after their audit every time, our workload increases as well

upvoted 1 times

Lenell 2 years, 3 months ago

Selected Answer: A

Because external audits does not take on the role of a "trusted advisor" but more of a regulator with punitive capability, I see "A" as the BEST answer.

upvoted 2 times

Eric0223 2 years, 5 months ago

Selected Answer: A

The CCSP offical guide mentioned that internal audit covers cost, design, performance...while external covers some sort of the efficiency of control implementation.

upvoted 2 times

kepalon 3 years ago

Selected Answer: A

Certification,

upvoted 1 times

carls233 3 years, 5 months ago

Answer looks correct to me, note the word 'internal' got to be A

upvoted 2 times

Banzaaai 3 years, 6 months ago

C, costs out of audit scope

upvoted 2 times

CISSP_Wannabe 4 years ago

Agree that Costs should be the correct answer. Certification I think is a red-herring (distractor).

upvoted 2 times

Sa007788 4 years, 2 months ago

internal audit may be the first phase of certification process, before going with external auditor for certification.Internal audit is recomanded.Costs should be best answer

upvoted 2 times

Q2 4 years, 4 months ago

Me too. Anyone else want to chime in?

upvoted 4 times

Question #120 *Topic 1*

Which of the following is the sole responsibility of the cloud customer, regardless of which cloud model is used?

A. Infrastructure

B. Platform

C. Application

D. Data

Suggested Answer: *D*

Regardless of which cloud-hosting model is used, the cloud customer always has sole responsibility for the data and its security.

Community vote distribution

D (100%)

kepalon 6 months, 1 week ago

Selected Answer: D

DATA is always a cloud customer responsibility, with the Governance as well.

upvoted 4 times

Bill_nye_russian_guy 10 months, 2 weeks ago

Data is correct. Although you can transfer through insurance, This ultimately falls on the owner.

upvoted 4 times

Zeezee2 10 months, 1 week ago

Indeed ultimate liability is not transferred. Only financial ramifications of mishandling data that was brought to light by regulatory audit, contract breach or other may be subject to payback by insurance as long as the reason is acceptable following the contract with the insurer.

upvoted 3 times

Question #121 *Topic 1*

What process is used within a clustered system to provide high availability and load balancing?

A. Dynamic balancing

B. Dynamic clustering

C. Dynamic optimization

D. Dynamic resource scheduling

Suggested Answer: *D*

Dynamic resource scheduling (DRS) is used within all clustering systems as the method for clusters to provide high availability, scaling, management, and workload distribution and balancing of jobs and processes. From a physical infrastructure perspective, DRS is used to balance compute loads between physical hosts in a cloud to maintain the desired thresholds and limits on the physical hosts.

Community vote distribution

D (100%)

ichnos Highly Voted 3 years ago

D is correct but is not dynamic. DRS stands for Distributed Resource Scheduling (DRS)

upvoted 13 times

Sa007788 Highly Voted 2 years, 8 months ago

DRS=

Distributed Resource Scheduling

upvoted 5 times

Pika26 Most Recent 4 months, 1 week ago

Selected Answer: D

D. Dynamic resource scheduling

upvoted 1 times

ikamalbhatt 4 months, 4 weeks ago

Distributed resource scheduling, which focuses on providing resources to virtual machines to ensure they can meet performance service- level requirements. It also allows migrations of systems to other underlying infrastructure during maintenance and includes monitoring and management capabilities. In short, distributed resource scheduling is the ability to manage resources across a cluster or environment in a way that optimizes reliable and consistent service delivery.

upvoted 2 times

Question #122 *Topic 1*

Which of the following is NOT a function performed by the handshake protocol of TLS?

A. Key exchange

B. Encryption

C. Negotiation of connection

D. Establish session ID

Suggested Answer: *B*

The handshake protocol negotiates and establishes the connection as well as handles the key exchange and establishes the session ID. It does not perform the actual encryption of data packets.

Community vote distribution

B (100%)

kepalon 6 months, 1 week ago

Selected Answer: B

encryption - of data

upvoted 2 times

David_S 1 year, 4 months ago

The answer B. would be better presented as "Data encryption" or "Session data encryption." The TLS handshake does employ encryption, for secrets and keys, using asymmetric encryption. It is true, however, that the data of the TLS session is not encrypted in the handshake.

upvoted 4 times

Question #123 *Topic 1*

Unlike SOC Type 1 reports, which are based on a specific point in time, SOC Type 2 reports are done over a period of time. What is the minimum span of time for a SOC Type 2 report?

A. Six months

B. One month

C. One year

D. One week

Suggested Answer: *A*

SOC Type 2 reports are focused on the same policies and procedures, as well as their effectiveness, as SOC Type 1 reports, but are evaluated over a period of at least six consecutive months, rather than a finite point in time.

Community vote distribution

A (100%)

EdwardLeeBurtle Highly Voted 2 years, 2 months ago

They need to fix the language on these. SOC 1, SOC 2 Type 1 and Type 2 etc. Use the correct format or it makes studying more difficult.

upvoted 5 times

MaciekMT Most Recent 1 month, 1 week ago

Selected Answer: A

Although the AICPA does not mandate a strict minimum for the coverage period of a SOC 2 Type 2 report, the generally accepted industry practice (and what most clients and auditors expect) is a six-month coverage period at a minimum. This is because a Type 2 report is meant to verify the effectiveness of controls over time, rather than just at a single point.

upvoted 1 times

ra1paul 1 month, 3 weeks ago

Selected Answer: A

12 months is a long time without any control changes.

upvoted 1 times

qpodian 5 months, 3 weeks ago

It is 1 year. Just google it

upvoted 1 times

DA95 1 year, 3 months ago

C. One year

upvoted 1 times

kepalon 2 years ago

Selected Answer: A

6 months

upvoted 3 times

Question #124 *Topic 1*

What changes are necessary to application code in order to implement DNSSEC?

A. Adding encryption modules

B. Implementing certificate validations

C. Additional DNS lookups

D. No changes are needed.

Suggested Answer: *D*
To implement DNSSEC, no additional changes are needed to applications or their code because the integrity checks are all performed at the system level.

Community vote distribution
D (100%)

MaciekMT 1 month, 1 week ago

Selected Answer: D

DNSSEC (Domain Name System Security Extensions) is handled at the DNS infrastructure and resolver level, rather than at the application code level. As long as the underlying DNS resolvers and infrastructure support DNSSEC, applications typically don't require any additional changes to start benefiting from the secure DNS lookups.

upvoted 1 times

ArashV 2 months ago

Selected Answer: C

It's very unlikely this question is asked just to be pointless and with no action needed. I think as per specified by contributor DA95, option C is correct.

upvoted 1 times

nzboy123 3 months ago

Selected Answer: C

To facilitate signature validation, DNSSEC adds a few new DNS record types:

RRSIG - Contains a cryptographic signature
DNSKEY - Contains a public signing key
DS - Contains the hash of a DNSKEY record
NSEC and NSEC3 - For explicit denial-of-existence of a DNS record
CDNSKEY and CDS - For a child zone requesting updates to DS record(s) in the parent zone.

upvoted 1 times

DA95 3 months, 2 weeks ago

In order to implement DNSSEC, some changes to application code may be necessary. DNSSEC is a security extension to the Domain Name System (DNS) that provides authentication for DNS lookups. To implement DNSSEC, application code may need to be updated to perform additional DNS lookups in order to verify the authenticity of DNS records. This may involve adding code to perform cryptographic operations in order to validate DNSSEC signatures. Therefore, the correct answer is option C, "Additional DNS lookups."

upvoted 3 times

akg001 10 months, 2 weeks ago

Selected Answer: D

D. No changes are needed.

upvoted 2 times

Question #125 *Topic 1*

Which type of controls are the SOC Type 1 reports specifically focused on?

A. Integrity

B. PII

C. Financial

D. Privacy

Suggested Answer: *C*

SOC Type 1 reports are focused specifically on internal controls as they relate to financial reporting.

Community vote distribution

C (100%)

AliHamza Highly Voted 2 years, 9 months ago

It is SOC 1, SOC 2 - Type 1 & Type 2, SOC 3

upvoted 11 times

gjjw Highly Voted 2 years, 10 months ago

to be precise it's SOC 1, not SOC Type 1

upvoted 8 times

Sa007788 2 years, 2 months ago

Yes, agree

upvoted 2 times

serget12 Most Recent 5 months, 1 week ago

SOC 1 reports focus solely on controls at a service provider that are likely to be relevant to an audit of a subscriber’s financial statements.

upvoted 1 times

akg001 10 months, 2 weeks ago

Selected Answer: C

C. Financial

upvoted 1 times

kepalon 1 year ago

The question should read SOC 1; Not TYPE 1

upvoted 1 times

Question #126 *Topic 1*

Which security concept is based on preventing unauthorized access to data while also ensuring that it is accessible to those authorized to use it?

A. Integrity

B. Availability

C. Confidentiality

D. Nonrepudiation

Suggested Answer: *C*

The main goal of confidentiality is to ensure that sensitive information is not made available or leaked to parties that should not have access to it, while at the same time ensuring that those with appropriate need and authorization to access it can do so in a manner commensurate with their needs and confidentiality requirements.

Community vote distribution

xroxro Highly Voted 2 years, 1 month ago

what a messy question

preventing unauthorized access to data --> C confidentiality
ensuring that it is accessible to those authorized to use it --> B Availability

looks like we have to guess between B and C.
As CSSP is more security oriented, i am for C

upvoted 6 times

zxccvbnm Highly Voted 2 years, 1 month ago

Selected Answer: C

Availability cant prevent unauthorized access

upvoted 5 times

MaciekMT Most Recent 1 month, 1 week ago

Selected Answer: C

Confidentiality is about ensuring that only authorized individuals have access to data, while unauthorized individuals are barred from it. By definition, this security principle means preventing disclosure to those who shouldn't see the data, and implicitly allowing access to those who are properly authorized.

Availability focuses on ensuring the data/systems are online and reachable for authorized users when needed (but does not address unauthorized access).

Integrity focuses on preventing unauthorized modification of data and ensuring its accuracy and reliability.

Nonrepudiation ensures a party cannot deny (repudiate) having taken an action (e.g., sending a message or signing a document).

upvoted 1 times

TraceSplice 6 months, 1 week ago

Selected Answer: C

C confidentiality

upvoted 1 times

SameerDutta 6 months, 2 weeks ago

Selected Answer: C

C is the answer

upvoted 1 times

DA95 1 year, 9 months ago

The security concept that is based on preventing unauthorized access to data while also ensuring that it is accessible to those authorized to use it is availability (option B). This concept is an important part of information security and focuses on ensuring that data and systems are available to users when they need them, without being compromised or disrupted by unauthorized access or other security threats. Integrity (option A) is related to the accuracy and completeness of data, confidentiality (option C) is focused on protecting data from unauthorized disclosure, and nonrepudiation (option D) is a concept that is used to prevent someone from denying that they performed a certain action.

upvoted 1 times

Eric0223 1 year, 11 months ago

Selected Answer: C

CCSP book states that encryption is designed to prevent unauthorized access

upvoted 2 times

quagga 2 years, 2 months ago

Selected Answer: B

I'd say B. Availability. https://www.isc2.org/Certifications/CISSP/CISSP-Student-Glossary#:~:text=Availability

Availability: Ensuring timely and reliable access to and use of information by authorized users.

upvoted 2 times

CptSweatbread 2 years, 1 month ago

I thought B as well, but I think the key words here are "prevent unauthorized access". Availability on it's own does not prevent unauthorized access.

upvoted 3 times

akg001 2 years, 4 months ago

Selected Answer: C

C. Confidentiality

upvoted 3 times

Question #127 *Topic 1*

Which of the following is NOT a domain of the Cloud Controls Matrix (CCM)?

A. Data center security

B. Human resources

C. Mobile security

D. Budgetary and cost controls

Suggested Answer: *D*

Budgetary and cost controls is not one of the domains outlined in the CCM.

Community vote distribution

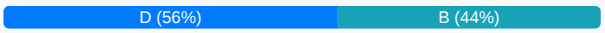

melvint Highly Voted 8 months ago

Selected Answer: D

https://cloudsecurityalliance.org/research/cloud-controls-matrix/

Answer is D

upvoted 5 times

MaciekMT Most Recent 4 weeks, 1 day ago

Selected Answer: D

The Cloud Controls Matrix (CCM) is a cybersecurity control framework developed by the Cloud Security Alliance (CSA). It includes domains that address various security aspects of cloud computing, but budgetary and cost controls are not part of its focus. The CCM is designed to ensure cloud security, not financial management.

Why the other options are part of CCM:

A. Data center security → This falls under Infrastructure & Virtualization Security (IVS) and Facility Security (FS) domains in CCM.

B. Human resources → Covered under Human Resources Security (HRS), which deals with security policies for employees.

C. Mobile security → Addressed under Mobile Security (MOS), ensuring security for mobile and endpoint devices.

upvoted 1 times

lolanczos 3 months, 1 week ago

Selected Answer: D

It's D.

Budgetary and cost controls are not included in the Cloud Controls Matrix (CCM), which focuses on security, compliance, and risk management rather than financial aspects.

upvoted 1 times

Pika26 10 months, 1 week ago

Selected Answer: B

B: Human resources

upvoted 1 times

Loveguitar 10 months, 1 week ago

D is the correct answer.

upvoted 4 times

kbnk 11 months ago

Selected Answer: B

The domain that is NOT a part of the Cloud Controls Matrix (CCM) is B. Human resources.

The Cloud Controls Matrix (CCM) is a set of controls designed to provide fundamental security principles to guide cloud vendors and to assist customers in assessing the overall security risk of a cloud provider. The domains of CCM include:

A. Data Center Security: Covers the physical and environmental security of the cloud provider's data center.

C. Mobile Security: Deals with security considerations for mobile devices, such as smartphones and tablets.

D. Budgetary and Cost Controls: Addresses the budgetary and cost controls of the cloud provider, such as pricing, billing, and chargeback.

Therefore, option B, Human resources, is not a domain of the Cloud Controls Matrix (CCM
upvoted 2 times

lolanczos 3 months, 1 week ago
You are 100% wrong. It's D that isn't part of the CCM. A, B, and C are there.
upvoted 1 times

bradseth 1 year, 1 month ago
Selected Answer: B
come on
upvoted 1 times

Banzaaai 2 years, 6 months ago
why not Human Resources?
upvoted 1 times

xaccan 2 years, 5 months ago
The domains covered in the new Cloud Controls Matrix (CCM) v4 are:

Application & Interface Security
Audit and Assurance
Business Continuity Mgmt & Op Resilience
Change Control & Configuration Management
Data Security and Privacy - DSP (old DSI)
Datacenter Security
Cryptography, Encryption and Key Management
Governance, Risk Management and Compliance
Human Resources Security
Identity & Access Management
Infrastructure & Access Management
Infrastructure & Virtualization
Interoperability & Portability
• Universal EndPoint Management
Security Incident Management, E-Discovery & Cloud Forensics
Supply Chain Management, Transparency & Accountability
Threat & Vulnerability Management
Logging and Monitoring
upvoted 12 times

serget12 1 year, 5 months ago
I'd double check the list provided, I did not see Logging and Monitoring although that seems like it should be part of the Matrix.
upvoted 2 times

Question #128 *Topic 1*

Which security concept, if implemented correctly, will protect the data on a system, even if a malicious actor gains access to the actual system?

A. Sandboxing

B. Encryption

C. Firewalls

D. Access control

Suggested Answer: *B*

In any environment, data encryption is incredibly important to prevent unauthorized exposure of data either internally or externally. If a system is compromised by an attack, having the data encrypted on the system will prevent its unauthorized exposure or export, even with the system itself being exposed.

Community vote distribution

B (100%)

kepalon 6 months, 1 week ago

Selected Answer: B

ENCRYPTION is the correct answer, if data is encrypted, then it will be protected even if disclosed/stolen during encryption.

upvoted 4 times

Question #129 *Topic 1*

Which of the following is the sole responsibility of the cloud provider, regardless of which cloud model is used?

A. Platform

B. Data

C. Physical environment

D. Infrastructure

Suggested Answer: *C*
Regardless of which cloud-hosting model is used, the cloud provider always has sole responsibility for the physical environment.

Community vote distribution
C (100%)

kepalon Highly Voted 1 year, 6 months ago

Selected Answer: C

Physical environment is always the responsibility of the cloud provider.
In IaaS the Infrastructure has a shared responsibility.
upvoted 5 times

ra1paul Most Recent 1 month, 3 weeks ago

Selected Answer: C

Customer will never have access to any physical environment.
upvoted 1 times

Rollizo 10 months, 3 weeks ago

Physical environment following ISC2 guide
upvoted 2 times

smiley 1 year, 9 months ago

Why not D? Who is always responsible for the infrastructure in the cloud?
upvoted 1 times

Joadeika 8 months, 3 weeks ago

Infrastructure and platform is what service provider can still provide but they can be provided by IaaS and PaaS providers. But the physical environment could could not be provided by any other person other than the service provider
upvoted 2 times

Question #130 *Topic 1*

Which of the following is NOT a factor that is part of a firewall configuration?

A. Encryption

B. Port

C. Protocol

D. Source IP

Suggested Answer: *A*

Firewalls take into account source IP, destination IP, the port the traffic is using, as well as the network protocol (UDP/TCP). Whether or not the traffic is encrypted is not something a firewall is concerned with.

Community vote distribution

A (100%)

kepalon Highly Voted 6 months, 1 week ago

Selected Answer: A

welcome to the ISC2 world.. questions are like this!!!

upvoted 7 times

ra1paul Most Recent 1 month, 3 weeks ago

Selected Answer: A

Firewalls do not handle encryption directly. they filter and control n/w traffic based on ports, protocols, IP.

upvoted 1 times

Cyber_Yeti 2 years, 7 months ago

This question needs revising. The stem does not provide enough context for the reader to understand the full picture.

upvoted 4 times

Zeezee2 10 months, 1 week ago

I think it is very clear that encryption is the right answer here, although the question is then very simple.

upvoted 2 times

Question #131 *Topic 1*

Which of the cloud deployment models involves spanning multiple cloud environments or a mix of cloud hosting models?

A. Community

B. Public

C. Hybrid

D. Private

Suggested Answer: *C*

A hybrid cloud model involves the use of more than one type of cloud hosting models, typically the mix of private and public cloud hosting models.

Community vote distribution

C (100%)

MaciekMT 1 month, 1 week ago

Selected Answer: C

Apparently (ChatGPT), A hybrid cloud deployment spans multiple cloud environments and can include a mix of:

Public cloud (e.g., AWS, Azure, Google Cloud)
Private cloud (on-premises or dedicated infrastructure)
Community cloud (shared infrastructure for a specific group)

upvoted 1 times

SameerDutta 6 months, 2 weeks ago

Selected Answer: C

hybrid

upvoted 1 times

hanyahmed 1 year, 9 months ago

Selected Answer: C

Hybrid is correct answer

upvoted 1 times

BossUK 2 years ago

Although Hybrid would often indicate a mixture of on-prem and cloud. It still can also be between different cloud environments

upvoted 3 times

CralP0ppa 2 years, 6 months ago

Wouldnät this be could multi-cloud ? Hybrid in my mind is a mixture of on-prem and cloud, not a mixture of cloud

upvoted 3 times

Temanny 2 years, 4 months ago

I thought the same too. The question is very confusing

upvoted 1 times

Question #132 *Topic 1*

Which of the following is NOT one of five principles of SOC Type 2 audits?

A. Privacy

B. Processing integrity

C. Financial

D. Security

Suggested Answer: *C*
The SOC Type 2 audits include five principles: security, privacy, processing integrity, availability, and confidentiality.

MaciekMT 4 weeks, 1 day ago

Selected Answer: C

SOC 2 (Service Organization Control Type 2) audits are based on the Trust Services Criteria (TSC), which consist of five key principles:

Security – Protection of systems from unauthorized access.
Availability – Ensuring systems are available for operation and use.
Processing Integrity – Ensuring system processing is complete, valid, accurate, timely, and authorized.
Confidentiality – Protection of sensitive data from unauthorized disclosure.
Privacy – Proper handling of personal information.
Financial controls are NOT part of SOC 2 audits. Financial reporting is covered under SOC 1, which is focused on internal controls over financial reporting (ICFR), not security and operational controls.

upvoted 1 times

kepalon 6 months, 1 week ago

Question should read: SOC 2 audits, or even SOC 2 type audits. BUT not SOC TYPE 2.

upvoted 3 times

xaccan 11 months, 4 weeks ago

SOC 2 defines criteria for managing customer data based on five “trust service principles”—security, availability, processing integrity, confidentiality and privacy.

upvoted 2 times

deegadaze1 1 year, 6 months ago

No, It is About SOC1 Type 1/2 [Precisely SOC1 type 2 (type II)]

upvoted 1 times

Sa007788 1 year, 8 months ago

question is about SOC2 type2

upvoted 3 times

evilwizardington 1 year, 7 months ago

Technically, it is about SOC 2, either Type 1 and Type 2.

upvoted 1 times

Ahbey_911 1 year, 7 months ago

I agree

upvoted 1 times

Question #133 *Topic 1*

Which aspect of cloud computing makes data classification even more vital than in a traditional data center?

A. Interoperability

B. Virtualization

C. Multitenancy

D. Portability

Suggested Answer: *C*

With multiple tenants within the same hosting environment, any failure to properly classify data may lead to potential exposure to other customers and applications within the same environment.

Community vote distribution

C (100%)

akg001 4 months, 1 week ago

Selected Answer: C

C. Multitenancy

upvoted 2 times

Question #134 *Topic 1*

What concept does the "T" represent in the STRIDE threat model?

A. TLS

B. Testing

C. Tampering with data

D. Transport

Suggested Answer: *C*

Explanation -

Any application that sends data to the user will face the potential that the user could manipulate or alter the data, whether it resides in cookies, GET or POST commands, or headers, or manipulates client-side validations. If the user receives data from the application, it is crucial that the application validate and verify any data that is received back from the user.

Community vote distribution

C (100%)

ikamalbhatt 4 months, 4 weeks ago

Selected Answer: C

STRIDE stands for spoofing, tampering, repudiation, information disclosure, denial of service, and elevation of privileges.

upvoted 1 times

akg001 1 year, 4 months ago

Selected Answer: C

C. Tampering with data

upvoted 2 times

Question #135 *Topic 1*

Which of the following would be a reason to undertake a BCDR test?

A. Functional change of the application

B. Change in staff

C. User interface overhaul of the application

D. Change in regulations

Suggested Answer: *A*

Any time a major functional change of an application occurs, a new BCDR test should be done to ensure the overall strategy and process are still applicable and appropriate.

Community vote distribution

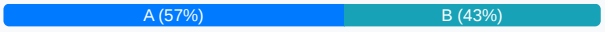

bdfb8cf 3 weeks ago

Selected Answer: B

Imagine rolling out functions on a weekly / monthly basis, would you test BCDR at this frequency?
But if you change personel, should test the playbook in case of a crises

upvoted 1 times

BlackListRapa 5 months, 1 week ago

Selected Answer: A

Major configuration changes within an application should entail new BCDR testing. Any major configuration change or update represents a significant shift in an environment, and, as such, proper testing is needed to ensure that all BCDR implementations and procedures are both still valid and still work as intended. the changes mentioned in the other answer choices are either minor or personnel changes that do not require new comprehensive testing.

upvoted 2 times

CptSweatbread 1 year, 1 month ago

I would agree more with A as an answer if the scenario stated it was a major functional change. Since majority functional changes could introduce dependencies or single points of failure which wasn't included in the initial BCDR plan. Just making a functional change is insufficient to retest the BCDR plan IMHO.

upvoted 1 times

kepalon 1 year, 6 months ago

Selected Answer: A

A is the correct answer to this question

upvoted 2 times

Ukashek 1 year, 10 months ago

From experience - you do not test BCDR after every functional change, but if regulator requires you to test particular environment semi annually instead of annually (for highly regulated institution) you must comply

upvoted 2 times

Zeezee2 1 year, 10 months ago

Selected Answer: B

Personally I think B is the most appropriate answer here, I disagree with C & D and I don't think any simple functional change should trigger BCP retesting. It should perhaps trigger BCP plan changes based on changes to that application's BIA if the new functionality makes the application more or less critical to the organisation however.

upvoted 3 times

CloudTrip 2 years, 1 month ago

Disagree with this answer. From practical exp, if it's a monolithic application and IaaS/PaaS then any application change will not require BCDR test as the infra layer or DR would have been the same. SImilarly containers and SaaS app change/update/config change also should not require BCDR test. Regulation change can be valid reason for BCDR test. Somebody commented here about RTO sarcastically but those RPO/RTO changes could be internal rules/policy with BCDR and definitely may require "next" DR test to test /confirm the new targets. More regulations can be state or Federal Govt restrictions for e.g. about failover from Govt DC to

happen only to other DCs within certain geographical perimeter. Like many other ISC question, language of this question also rather vague. ISC may expect answer for this as Application change but that's not something ideally trigger this.

upvoted 2 times

David_S 2 years, 4 months ago

This question has a few potentially correct answers. Even a change in staff, if those are the staff who need to perform the BC/DR procedures, would trigger the need for testing, with the new staff involved.

upvoted 3 times

phanil1 2 years, 5 months ago

Each time Application changes the design is changed, it can be called for a BCDR Test.

upvoted 3 times

Sa007788 2 years, 8 months ago

For me it's D, change in regulation may invole major change regarding metrics of BCDR

upvoted 3 times

evilwizardington 2 years, 7 months ago

And the changes in metrics of a BCDR do not mean you need to test it. =P

For example, assume that your RTO is lower. So what? the infrastructure didnt change, the steps didnt change, etc. So you dont need to test the BCDR to know the estimated time.

upvoted 3 times

Question #136 *Topic 1*

What is the biggest challenge to data discovery in a cloud environment?

A. Format

B. Ownership

C. Location

D. Multitenancy

Suggested Answer: *C*
With the distributed nature of cloud environments, the foremost challenge for data discovery is awareness of the location of data and keeping track of it during the constant motion of cloud storage systems.

Community vote distribution
C (100%)

phanil1 Highly Voted 2 years, 5 months ago
When the data is multi tenant, you cannot have access to the each tenant and get the data , it is a big concern, when the question says whats the biggest concern, it should be multi tenancy.
upvoted 7 times

MaciekMT Most Recent 4 weeks, 1 day ago
Selected Answer: C
The biggest challenge to data discovery in a cloud environment is location, because cloud data is often distributed across multiple regions, data centers, and even different cloud providers. This creates difficulties in:

Tracking where data resides → Cloud providers often store data in multiple locations for redundancy and performance optimization, making it hard to pinpoint its exact location.
Legal & Compliance Issues → Different countries have data sovereignty laws (e.g., GDPR, CCPA) that regulate where data can be stored, making discovery even more complex.
Dynamic Environments → Cloud workloads can scale, migrate, or get replicated across various regions, affecting data visibility.
upvoted 1 times

ikamalbhatt 4 months, 4 weeks ago
Selected Answer: C
Laws and regulations may limit the types or methods of discovery you can engage in, or what you can do with the data, as well as where and how you can store it. The EU's General Data Protection Regulation (GDPR) has requirements that impact where and how you store data and thus may have an impact on discovery, and other regulations may also apply, depending on your organization's compliance requirements.
upvoted 3 times

Pika26 5 months, 1 week ago
Selected Answer: C
Answer is C.
upvoted 1 times

AlanJP 2 years, 2 months ago
The ISC2 CBK lists the first challenge as "Identifying where your data is" so I guess Location is the 'right' answer despite what we think
upvoted 4 times

asldavid 2 years, 6 months ago
according to Official ISC2 doc, ■■ Identifying where your data
upvoted 2 times

babusartop17 2 years, 3 months ago
And? Did you forget to add rest of the information?
upvoted 2 times

Rangakarthik 2 years, 8 months ago

I agree

upvoted 1 times

yesj 2 years, 11 months ago

Multi-Tenancy is also there Why "Multi Tenancy can't be answer

upvoted 4 times

Sa007788 2 years, 8 months ago

Multitenancy impact more Confidentiality or Jurisdiction requirement but Discovery still same

upvoted 2 times

Question #137 *Topic 1*

Which crucial aspect of cloud computing can be most threatened by insecure APIs?

A. Automation

B. Redundancy

C. Resource pooling

D. Elasticity

Suggested Answer: *A*
Cloud environments depend heavily on API calls for management and automation. Any vulnerability with the APIs can cause significant risk and exposure to all tenants of the cloud environment.

MaciekMT 1 month, 1 week ago

Selected Answer: A

APIs are the backbone of cloud automation—they enable the rapid provisioning, configuration, and management of resources. If the APIs are insecure, attackers can manipulate or disrupt those automated workflows, impacting the entire orchestration of the cloud environment. While elasticity, resource pooling, and redundancy all rely on automation, compromised APIs most directly threaten the automation layer that ties everything together.

upvoted 1 times

sweetykaur 5 months, 2 weeks ago

D. Elasticity. Insecure APIs can compromise the ability to scale resources dynamically, which is a key aspect of cloud computing's elasticity.

upvoted 1 times

akg001 2 years, 10 months ago

A. Automation

upvoted 1 times

Question #138 *Topic 1*

Which of the following should NOT be part of the requirement analysis phase of the software development lifecycle?

A. Functionality

B. Programming languages

C. Software platform

D. Security requirements

Suggested Answer: *D*
Security requirements should be incorporated into the software development lifecycle (SDLC) from the earliest requirement gathering stage and should be incorporated prior to the requirement analysis phase.

Community vote distribution
B (71%) D (29%)

MahesswarShriMohanty Highly Voted 4 years, 2 months ago
Answer is correct.

The Secure Software Development Lifecycle (SDLC) Process phases are
1. Requirement Gathering and Feasibility
2. Requirement Analysis
3. Design
4. Development/Coding
5. Testing
6. Maintenance

Note: It is essential that security be included in discussions and the SDLC process from the very initial stages.

Carter, Daniel. CCSP Certified Cloud Security Professional All-in-One Exam Guide, Second Edition (p. 169). McGraw-Hill Education. Kindle Edition.
upvoted 15 times

Nord Highly Voted 3 years, 7 months ago
Should be B (Programming languages)
upvoted 8 times

NobleGiantz 3 years, 7 months ago
Security should be discussed at each stage of the SDCL
upvoted 6 times

Monchel Most Recent 4 weeks ago
Selected Answer: B
Programming languages: This relates to implementation details—how the software will be built—not what it should achieve. This is typically decided in the design or development phase, not requirement analysis.
upvoted 1 times

MaciekMT 4 weeks, 1 day ago
Selected Answer: B
During the requirement analysis phase of the Software Development Lifecycle (SDLC), the focus is on understanding what the software needs to do rather than how it will be implemented. The key aspects include:

A. Functionality → Defines what the software should do.
C. Software platform → Determines where the software will run (e.g., cloud, on-prem, Windows, Linux).
D. Security requirements → Identifies necessary security controls (e.g., encryption, access control).
However, B. Programming languages are typically chosen later, during the design or implementation phase, based on system architecture and project constraints.
upvoted 1 times

TraceSplice 6 months, 1 week ago

Selected Answer: B

programming languages should not be part of the requirement analysis phase

upvoted 2 times

Kneebee 11 months, 1 week ago

The correct answer is "B";

The requirement analysis phase of the software development lifecycle focuses on gathering and documenting the functional and non-functional requirements of the software system. Programming languages, on the other hand, are not a requirement but a technical decision made later in the development process. Programming languages are chosen based on the project's needs and the expertise of the development team. Therefore, they should not be part of the requirement analysis phase.

upvoted 2 times

escaprix 1 year, 3 months ago

Selected Answer: B

The programming languages should not be part of the requirement analysis phase of the software development lifecycle. The requirement analysis phase focuses on gathering, analyzing, and documenting the functional and non-functional requirements of the software. It involves understanding the desired functionality, user needs, business processes, and system constraints

upvoted 1 times

Pika26 1 year, 4 months ago

Selected Answer: B

B: Programming languages

upvoted 1 times

Pika26 1 year, 5 months ago

Selected Answer: B

Answer is B

upvoted 1 times

bp339 1 year, 5 months ago

Selected Answer: B

Security requirements should be part of the requirement analysis phase of the software development lifecycle, as security is an essential aspect of software development. During the requirement analysis phase, the software's security requirements should be identified and documented to ensure that the software is developed with the necessary security controls in place to protect against potential threats.

upvoted 1 times

Lenell 1 year, 9 months ago

Selected Answer: D

I comprehend why ISC2 wants "D" as the answer. Their point is that "security requirements" should be included at the EARLIEST possible point in the cycle. This would be the "Gathering and Feasibility" phase. Thus the question should be worded differently. The key is to comprehend that all of these have a role in this phase. Only one is firm before going into (prior to) this phase.

upvoted 5 times

hanyahmed 1 year, 9 months ago

Selected Answer: B

the correct answer is B Programming Languages should be defined in Desing phase.

upvoted 1 times

DA95 1 year, 9 months ago

The correct answer is B. Programming languages. The requirement analysis phase is the first phase of the software development lifecycle, and it involves gathering and documenting the functional and non-functional requirements for the software. This includes things like the desired functionality of the software, the software platform, and the security requirements. Programming languages are not typically part of this phase, as they are typically selected later in the development process based on the requirements and other factors.

upvoted 1 times

Eric0223 1 year, 11 months ago

Selected Answer: B

B, coz security would cover entire lifecycle of SDLC, and programming language should be discussing once finalizing this function or platform

upvoted 2 times

zxccvbnm 2 years, 1 month ago

Selected Answer: B

B (Programming languages)

upvoted 1 times

F34 2 years, 4 months ago

Selected Answer: B

B is correct

upvoted 1 times

kepalon 2 years, 6 months ago

Selected Answer: B

D is wrong, in the gathering requirements we already introduce the Security.

I will go for B: The specific programming language.

upvoted 1 times

Question #139 *Topic 1*

Which of the cloud cross-cutting aspects relates to the assigning of jobs, tasks, and roles, as well as to ensuring they are successful and properly performed?

A. Service-level agreements

B. Governance

C. Regulatory requirements

D. Auditability

Suggested Answer: *B*
Governance at its core is the idea of assigning jobs, takes, roles, and responsibilities and ensuring they are satisfactory performed.

Community vote distribution
B (100%)

Seke Highly Voted 2 years, 1 month ago
Governance is the system by which the provisioning and usage of cloud services are directed and controlled. Governance defines actions, assigns responsibilities, and verifies performance. Governance includes the policy, process, and internal controls that comprise how an organization is run. Everything from the structures and policies to the leadership and other mechanisms for management.
upvoted 7 times

Pika26 Most Recent 5 months, 1 week ago
Selected Answer: B
Answer is B.
upvoted 1 times

kepalon 1 year, 6 months ago
Selected Answer: B
GOVERNANCE!!!
upvoted 3 times

axman832005 1 year, 7 months ago
Wouldn't auditing be the one here? the question says "ensuring.." Governance doesn't seem correction at all..
upvoted 1 times

AlanJP 2 years, 2 months ago
According to the ISC2 CBK "The term 'governance' relating to processes and decisions looks to define actions, assign responsibilities, and verify performance."
upvoted 4 times

yesj 3 years ago
Goverance should be the one who keeps tracks of all the jobs should be completed
upvoted 3 times

bark101 3 years, 1 month ago
The answer should be SLA
upvoted 1 times

kjjcraigskel 2 years, 11 months ago
SLA is not one of the defined cross cutting aspects.
upvoted 2 times

Ahbey_911 2 years, 7 months ago
SLA is one of the defined cross cutting aspects though (ISO 17789)
Cross-cutting aspects include:
• auditability (clause 8.5.2);
• availability (clause 8.5.3);
• governance (clause 8.5.4);

- interoperability (clause 8.5.5);
- maintenance and versioning (clause 8.5.6);
- performance (clause 8.5.7);
- portability (clause 8.5.8);
- protection of personally identifiable information (clause 8.5.9);
- regulatory;
- resiliency (clause 8.5.10);
- reversibility (clause 8.5.11);
- security (clause 8.5.12);
- service levels and service level agreement (clause 8.5.13).

upvoted 5 times

Xindydo 2 years ago

SLA is performance-based...it has nothing to do with assigning roles or responsibilities. The SLA documents the agreed-upon performance levels within the cloud.

upvoted 2 times

Question #140 *Topic 1*

Which regulatory system pertains to the protection of healthcare data?

A. HIPAA

B. HAS

C. HITECH

D. HFCA

Suggested Answer: *A*

The Health Insurance Portability and Accountability Act (HIPAA) sets stringent requirements in the United States for the protection of healthcare records.

Community vote distribution

A (100%)

RalphLee 7 months, 4 weeks ago

HITECH also protect PHI.

upvoted 2 times

akg001 1 year, 4 months ago

Selected Answer: A

A. HIPAA

upvoted 1 times

Question #141 *Topic 1*

Which aspect of cloud computing makes it very difficult to perform repeat audits over time to track changes and compliance?

A. Virtualization

B. Multitenancy

C. Resource pooling

D. Dynamic optimization

Suggested Answer: *A*

Cloud environments will regularly change virtual machines as patching and versions are changed. Unlike a physical environment, there is little continuity from one period of time to another. It is very unlikely that the same virtual machines would be in use during a repeat audit.

Community vote distribution

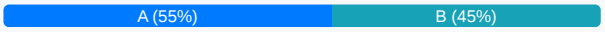

MaciekMT 3 weeks, 6 days ago

Selected Answer: D

Dynamic optimization in cloud computing refers to the continuous and automatic adjustment of resources (e.g., scaling, load balancing, and instance migrations) to optimize performance and efficiency. This makes it difficult to perform repeat audits over time because:

Constant Changes → Cloud resources are provisioned, modified, and decommissioned dynamically, making it hard to track what has changed between audits.

Ephemeral Workloads → Instances and containers may exist only for a short time, meaning traditional audit logs might miss key events.

Compliance Challenges → Regulatory frameworks often require evidence of consistent security configurations, but dynamic optimization can introduce unexpected changes.

Why the other options are incorrect:

upvoted 1 times

lolanczos 3 months, 1 week ago

Selected Answer: D

Guys, it's D. Dynamic optimization refers to the cloud's ability to automatically adjust and reallocate resources as needed, often without human intervention. This can make it challenging to track and audit specific changes over time, as configurations and resource allocations may change dynamically.

A. Virtualization: While virtualization is a fundamental technology in cloud computing, it does not inherently complicate repeat audits. It provides a consistent platform for tracking resources.

B. Multitenancy: Multitenancy involves sharing resources among multiple customers, which poses security and isolation challenges, but it does not specifically hinder repeat audits.

C. Resource pooling: Resource pooling refers to combining resources for efficiency. It doesn't directly affect the ability to perform audits, as pooled resources are typically managed and tracked systematically.

upvoted 1 times

Sivath 4 months, 1 week ago

Dynamic optimization : specifically refers to the frequent and automated changes that pose unique challenges to auditing and compliance tracking.

upvoted 1 times

sweetykaur 5 months, 2 weeks ago

Dynamic optimization. This aspect continuously adjusts resources and configurations, complicating repeat audits and tracking changes over time.

upvoted 1 times

Kneebee 11 months, 2 weeks ago

Answer A is correct. The virtual machines in the cloud environment regularly change due to patching and version updates.

upvoted 1 times

Kneebee 1 year, 5 months ago

The correct answer is "B".

Multitenancy in cloud computing refers to the practice of multiple customers (tenants) sharing the same physical resources within a cloud environment. While multitenancy can lead to cost efficiencies and resource optimization, it can also make it more challenging to perform repeat audits over time to track changes and compliance. This is because multiple tenants may be making changes and utilizing the same shared resources, making it difficult to isolate and track specific changes and ensure compliance for each individual tenant.

upvoted 2 times

escaprix 1 year, 9 months ago

Selected Answer: B

When multiple tenants share the same underlying infrastructure, it becomes complex to maintain a granular level of visibility and control over the data and activities of each tenant.

upvoted 1 times

Pika26 1 year, 11 months ago

Selected Answer: B

Answer is B. Multi-tenancy.

upvoted 1 times

bp339 1 year, 11 months ago

Selected Answer: B

it is a valid point that cloud environments are dynamic and constantly changing, with virtual machines being updated, patched, and replaced over time. This can make it challenging to maintain continuity and consistency between audits, as the environment may look different each time.

However, this is primarily a result of the virtualization and dynamic resource allocation that are common in cloud environments, rather than the multitenancy aspect specifically. Multitenancy refers to the sharing of resources and infrastructure among multiple tenants, which can complicate audit tracking and consistency in a different way.

upvoted 1 times

amarongas 2 years, 1 month ago

Selected Answer: A

Virtualization - A

upvoted 1 times

serget12 2 years, 5 months ago

Virtualization is cloud, especially with templates and best practices to automate. Which would make it very easy to track and repeat logging and auditing. Bad question

upvoted 1 times

zxccvbnm 2 years, 7 months ago

Selected Answer: A

virtualization as i think Multitenancy shouldnt impact audit

upvoted 2 times

quagga 2 years, 8 months ago

Selected Answer: B

B. Multitenancy

upvoted 2 times

kepalon 3 years ago

Selected Answer: A

multitenancy is correct, good explanation provided.

upvoted 3 times

styro 2 years, 7 months ago

I dont understand comments like these. The given answer is A. You select A, then state multitenancy is correct (which is B), and refer to the explanation provided being good, which is an explanation for virtualization being the answer. Why is this upvoted two times?

upvoted 11 times

Question #142 *Topic 1*

Which security concept would business continuity and disaster recovery fall under?

A. Confidentiality

B. Availability

C. Fault tolerance

D. Integrity

Suggested Answer: *B*

Disaster recovery and business continuity are vital concerns with availability. If data is destroyed or compromised, having regular backup systems in place as well as being able to perform disaster recovery in the event of a major or widespread problem allows operations to continue with an acceptable loss of time and data to management. This also ensures that sensitive data is protected and persisted in the event of the loss or corruption of data systems or physical storage systems.

Community vote distribution

B (100%)

akg001 4 months, 1 week ago

Selected Answer: B

B. Availability

upvoted 2 times

Question #143 *Topic 1*

Which of the following is NOT an application or utility to apply and enforce baselines on a system?

A. Chef

B. GitHub

C. Puppet

D. Active Directory

Suggested Answer: *B*
GitHub is an application for code collaboration, including versioning and branching of code trees. It is not used for applying or maintaining system configurations.

Community vote distribution
B (100%)

MaciekMT 3 weeks, 6 days ago

Selected Answer: B
GitHub is a code repository and version control platform, primarily used for storing, managing, and collaborating on code. While it can store configuration files or automation scripts, it does not directly apply or enforce baselines on a system.
Why the other options are correct:
A. Chef → A configuration management tool that automates system configuration and enforces baselines.
C. Puppet → Another configuration management tool used to apply system baselines automatically.
D. Active Directory → Used for Group Policy enforcement, which helps enforce security and system configuration baselines.
upvoted 1 times

kepalon 6 months, 1 week ago

Selected Answer: B
github is a a sw repository and collaboration tool. Not a control SW tool, the other ones are SW control tools.
And yes, you can enforce baselines with AD. GPO- Group Policy Objects & templates.
upvoted 2 times

AWSPro24 9 months, 1 week ago
I think gitHub is the right answer but there are certainly methods were IaC lives in github and it is therefore used as a baseline for building out infrastructure.
upvoted 1 times

NYF 1 year, 9 months ago
Microsoft is releasing security baselines for on-premises Active Directory connected devices using group policies. These are used by many organizations around the globe for decades. Using these security settings, administrators can control the state of the corporate devices and maintain the standards.

https://www.rebeladmin.com/2019/08/step-step-guide-apply-security-baselines-windows-10-devices-using-microsoft-intune/
upvoted 1 times

Guivent 1 year, 10 months ago
Can we enforce baselines on Active Directory?
upvoted 2 times

duracell 1 year, 4 months ago
GPOs (Group Policy Objects)
upvoted 5 times

Question #144 *Topic 1*

Which of the cloud cross-cutting aspects relates to the ability for a cloud customer to easily remove their applications and data from a cloud environment?

A. Reversibility

B. Availability

C. Portability

D. Interoperability

Suggested Answer: *A*
Reversibility is the ability for a cloud customer to easily remove their applications or data from a cloud environment, as well as to ensure that all traces of their applications or data have been securely removed per a predefined agreement with the cloud provider.

MaciekMT 1 month, 1 week ago

Selected Answer: A

Reversibility refers to the ability of a cloud customer to easily exit the cloud service and remove their applications and data from the environment. While portability deals with moving data between providers, reversibility specifically addresses the ease and completeness of withdrawal from a cloud environment.

upvoted 1 times

nelombg 4 months ago

correct Answer is "Reversibility"

upvoted 1 times

serget12 11 months, 3 weeks ago

Is Reversibility part of the cross-cutting aspects?

upvoted 3 times

zaqwsx 1 year, 5 months ago

it looks correct

"Reversibility is like the Undo button for cloud. It separates organizations that can handle changes at scale from organizations that experience downtime whenever a need arises for compliance with regulatory requirements, flexibility, or agility in strategy."

upvoted 3 times

Question #145 *Topic 1*

Which of the following is NOT a function performed by the record protocol of TLS?

A. Encryption

B. Acceleration

C. Authentication

D. Compression

Suggested Answer: *B*

The record protocol of TLS performs the authentication and encryption of data packets, and in some cases compression as well. It does not perform any acceleration functions.

Community vote distribution

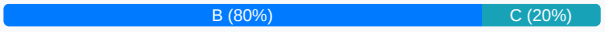

MaciekMT 1 month, 1 week ago

Selected Answer: B

The TLS record protocol is responsible for ensuring confidentiality (via encryption), data integrity (through authentication mechanisms like MACs), and it may also compress data before encryption. Acceleration is not a function of the TLS record protocol.

upvoted 1 times

rkumar16d 5 months ago

Answer is C. Authentication

upvoted 1 times

JohnnyBG 1 year ago

Selected Answer: C

Auth is not performed by record protocol for reasons mentioned bellow

upvoted 1 times

kepalon 3 years ago

Selected Answer: B

ACCELERATION IS CORRECT.

TLS Record protocol provides the actual secure communication method for transmitting data; encryption + authentication & in some cases also compression.

Acceleration is not a component of TLS, it might be provided by another external component

upvoted 4 times

[Removed] 3 years, 2 months ago

TLS Record Protocol manages the following:

Dividing outgoing messages into manageable blocks, and reassembling incoming messages.

Compressing outgoing blocks and decompressing incoming blocks (optional).

Applying a Message Authentication Code (MAC) to outgoing messages, and verifying incoming messages using the MAC.

Encrypting outgoing messages and decrypting incoming messages.

upvoted 2 times

AlanJP 3 years, 8 months ago

Authentication is actually done by the Handshake protocol. The Message Authentication Code is applied at the Record layer, but the authentication is not done here. TLS Accelerators are hardware encryption engines and are used in the Record layer.

upvoted 3 times

Question #146 *Topic 1*

What concept does the "R" represent with the DREAD model?

A. Reproducibility

B. Repudiation

C. Risk

D. Residual

Suggested Answer: *A*
Reproducibility is the measure of how easy it is to reproduce and successful use an exploit. Scoring within the DREAD model ranges from 0, signifying a nearly impossibly exploit, up to 10, which signifies something that anyone from a simple function call could exploit, such as a URL.

Zeezee2 Highly Voted 4 months, 1 week ago
The categories are:

Damage – how bad would an attack be?
Reproducibility – how easy is it to reproduce the attack?
Exploitability – how much work is it to launch the attack?
Affected users – how many people will be impacted?
Discoverability – how easy is it to discover the threat?
upvoted 11 times

Question #147 *Topic 1*

The SOC Type 2 reports are divided into five principles.
Which of the five principles must also be included when auditing any of the other four principles?

A. Confidentiality

B. Privacy

C. Security

D. Availability

Suggested Answer: *C*
Under the SOC guidelines, when any of the four principles other than security are being audited, which includes availability, confidentiality, processing integrity, and privacy, the security principle must also be included with the audit.

Community vote distribution
C (100%)

MaciekMT 1 month, 1 week ago

Selected Answer: C

The SOC 2 framework is built around five trust services principles: Security, Availability, Processing Integrity, Confidentiality, and Privacy. When performing an audit for any of the latter four principles, the Security principle is always a required baseline. It acts as the foundational control that underpins the effectiveness of the other areas.

upvoted 1 times

akg001 4 months, 1 week ago

Selected Answer: C

C. Security

upvoted 2 times

Question #148 *Topic 1*

How many additional DNS queries are needed when DNSSEC integrity checks are added?

A. Three

B. Zero

C. One

D. Two

Suggested Answer: *B*

DNSSEC does not require any additional DNS queries to be performed. The DNSSEC integrity checks and validations are all performed as part of the single DNS lookup resolution.

Community vote distribution

B (100%)

DA95 Highly Voted 2 years, 3 months ago

When DNSSEC integrity checks are added, an additional two DNS queries are needed. DNSSEC (Domain Name System Security Extensions) is a set of security extensions to the Domain Name System (DNS) that provide authentication and integrity for DNS data. When DNSSEC is used, additional DNS queries are needed in order to verify the authenticity and integrity of DNS records. This involves querying additional DNS resource records, such as the DNSKEY and DS records, which are used to verify the digital signatures on DNS data. As a result, two additional DNS queries are typically needed when DNSSEC is used, in addition to the initial query for the DNS data itself.

upvoted 5 times

MaciekMT Most Recent 1 month, 1 week ago

Selected Answer: B

DNSSEC is designed so that the extra security-related records (such as RRSIG, DNSKEY, and DS) are returned alongside the standard DNS responses. This integration means that the resolver does not need to issue any additional DNS queries beyond the original request.

upvoted 1 times

sweetykaur 5 months, 2 weeks ago

One. DNSSEC adds an extra DNS query to fetch the digital signatures necessary for verifying the authenticity and integrity of the DNS data.

upvoted 1 times

Pika26 1 year, 11 months ago

Answer is C. One.

upvoted 1 times

xroxro 2 years, 7 months ago

correct me if i'm wrong but with DNSsec the recursive DNS has to query public keys to verify signature, right ?
So, zero from a client point of view (only one query to his recursive DNS server) but many if count all needed queries

upvoted 2 times

akg001 2 years, 10 months ago

Selected Answer: B

B. Zero

upvoted 1 times

Question #149 *Topic 1*

Which of the following is the sole responsibility of the cloud customer, regardless of which cloud model is used?

A. Platform

B. Infrastructure

C. Governance

D. Application

Suggested Answer: *C*

Regardless of which cloud-hosting model is used, the cloud customer always has sole responsibility for the governance of systems and data.

Community vote distribution

C (100%)

Pillartech 6 months, 3 weeks ago

Selected Answer: C

governance

upvoted 1 times

akg001 2 years, 10 months ago

Selected Answer: C

C. Governance

upvoted 2 times

Question #150 *Topic 1*

Which of the following service categories entails the least amount of support needed on the part of the cloud customer?

A. SaaS

B. IaaS

C. DaaS

D. PaaS

Suggested Answer: *A*

With SaaS providing a fully functioning application that is managed and maintained by the cloud provider, cloud customers incur the least amount of support responsibilities themselves of any service category.

Community vote distribution

A (100%)

akg001 4 months, 1 week ago

Selected Answer: A

A. SaaS

upvoted 1 times

Question #151 *Topic 1*

Which of the following would NOT be a reason to activate a BCDR strategy?

A. Staffing loss

B. Terrorism attack

C. Utility disruptions

D. Natural disaster

Suggested Answer: *A*

The loss of staffing would not be a reason to declare a BCDR situation because it does not impact production operations or equipment, and the same staff would be needed for a BCDR situation.

Community vote distribution

A (100%)

DA95 3 months, 2 weeks ago

B. Terrorism attack

upvoted 1 times

akg001 10 months, 2 weeks ago

Selected Answer: A

A. Staffing loss

upvoted 3 times

Question #152 *Topic 1*

Which of the cloud cross-cutting aspects relates to the oversight of processes and systems, as well as to ensuring their compliance with specific policies and regulations?

A. Governance

B. Regulatory requirements

C. Service-level agreements

D. Auditability

Suggested Answer: *D*

Auditing involves reports and evidence that show user activity, compliance with controls and regulations, the systems and processes that run and what they do, as well as information and data access and modification records. A cloud environment adds additional complexity to traditional audits because the cloud customer will not have the same level of access to systems and data as they would in a traditional data center.

Community vote distribution

A (100%)

TraceSplice 6 months, 1 week ago

Selected Answer: A

Answer is A

upvoted 1 times

Fosca 8 months, 1 week ago

The OSG says it's governance on page 20. Auditability is a subset of governance. But like with many of these situations, the "right" answer is not necessarily the "correct" answer.

upvoted 2 times

Darln 8 months, 1 week ago

The answer should be governance as stated by the definition

upvoted 1 times

TheGinjaNinja 1 year, 5 months ago

Selected Answer: A

Answer is A

upvoted 1 times

Pika26 1 year, 5 months ago

Selected Answer: A

Answer is A. Governance.

upvoted 1 times

Brittle 1 year, 7 months ago

i think its A

upvoted 1 times

DERCHEF2009 1 year, 11 months ago

Selected Answer: A

Its A!

upvoted 3 times

serget12 1 year, 11 months ago

Regulatory Compliance:

Regulatory compliance is an organization’s requirement to adhere to relevant laws, regu□lations, guidelines, and specifications relevant to its business, specifically dictated by

the nature, operations, and functions it provides or utilizes to its customers.

upvoted 1 times

samsom 2 years, 2 months ago

Governance role is to set the directions in form of policies, standards, and guidelines. Management role is delivery based on processes & procedures that follow the policies, standards, and guidelines. Audit role is to review & inspect that delivery conforms to the policies, standards, and guidelines. Thus, answer is D & not A.
upvoted 1 times

axman832005 2 years, 7 months ago
It says oversight.. which is governance. Auditing is not oversight.. but then it says “ensuring their compliance..” that would be audit. I’d still pick D
upvoted 1 times

axman832005 2 years, 7 months ago
It says oversight.. which is governance. Auditing is not oversight.. but when the
upvoted 4 times

Seke 3 years ago
Auditability allows for users and the organization to access, report, and obtain evidence of actions, controls, and processes that were performed or run by a specified user.
upvoted 3 times

Vertho 3 years, 9 months ago
D is correct:
Cloud Cross-Cutting Aspects: These cross-cutting aspects include security, interoperability, portability, reversibility, privacy, availability, governance, performance, service levels, service level agreements, auditability and regulatory aspects
upvoted 2 times

HCL 3 years, 10 months ago
Why B - Regulatory Requirements is not the answer?
Regulatory requirements are also related to regulation, policy and compliance.
upvoted 2 times

kjjcraigskel 3 years, 11 months ago
Auditability isn't one of the cross cutting aspects. They are:

Security, Interoperability, Portability, Reversibility, Privacy, Availability and Governance
upvoted 1 times

Ahbey_911 3 years, 7 months ago
Yes, it is. It is at the top of the list. Check ISO 17789. I know the books don't list it.
Cross-cutting aspects include:
- auditability (clause 8.5.2);
- availability (clause 8.5.3);
- governance (clause 8.5.4);
- interoperability (clause 8.5.5);
- maintenance and versioning (clause 8.5.6);
- performance (clause 8.5.7);
- portability (clause 8.5.8);
- protection of personally identifiable information (clause 8.5.9);
- regulatory;
- resiliency (clause 8.5.10);
- reversibility (clause 8.5.11);
- security (clause 8.5.12);
- service levels and service level agreement (clause 8.5.13).

upvoted 8 times

ichnos 4 years ago
Correct answer: D

Most leading cloud providers supply their customers with a good deal of auditing, including reports and evidence that show user activity, compliance with controls and regulations, systems and processes that run and an explanation of what they do, as well as information, data access, and modification records.
upvoted 2 times

cthd 4 years ago

Why not A - Governance?

upvoted 1 times

CL888 4 years ago

Audits are done to make sure controls and processes do what they are supposed to do, and that the Org is compliant with regulations.

upvoted 1 times

Question #153 *Topic 1*

Which of the cloud cross-cutting aspects relates to the ability to reuse or move components of an application or service?

A. Availability

B. Interoperability

C. Reversibility

D. Portability

Suggested Answer: *B*

Interoperability is the ease with which one can move or reuse components of an application or service. This is maximized when services are designed without specific dependencies on underlying platforms, operating systems, locations, or cloud providers.

Community vote distribution

D (100%)

NYF Highly Voted 4 years, 3 months ago

B does not seem to be the right answer.

Portability seems like a better answer.

Interoperability is the ability of two or more systems or applications to exchange information and to mutually use the information that has been exchanged. Cloud interoperability is the ability of a customer's system to interact with a cloud service or the ability for one cloud service to interact with other cloud services by exchanging information according to a prescribed method to obtain predictable results.

Portability, on the other hand, is moving the data and/or applications from one system to another and having it remain useable or executable. Cloud data portability is the ability to easily move data from one cloud service to another without needing to re-enter the data. Cloud application portability is the ability to migrate an application from one cloud service to another or between a customer's environment and a cloud service.

---

upvoted 12 times

Trax 4 years, 3 months ago

Fully agreed. Should be D.

And it's important to understand difference between Interoperability & Portability....

upvoted 3 times

NobleGiantz 4 years, 1 month ago

Portability is the ease with which a party can move or reuse application or service components. Portability means that the service provider, underlying platform, operating system, API structure, format of data, or other factors do not present obstacles to seamlessly moving services from one solution to another.

upvoted 3 times

topcat Highly Voted 3 years, 10 months ago

Answer is given:

Interoperability

Move or reuse components of an application on the same time

Portability

Easily and seamlessly move between different cloud providers.

Key word is components

upvoted 11 times

MaciekMT Most Recent 3 weeks, 6 days ago

Selected Answer: D

Portability in cloud computing refers to the ability to reuse or move components of an application or service across different cloud environments with minimal modification. This includes:
Migrating applications between cloud providers (e.g., AWS → Azure).
Moving workloads between on-premises and cloud environments (hybrid cloud).
Ensuring data, configurations, and dependencies can function in different environments.
upvoted 1 times

sweetykaur 5 months, 2 weeks ago
Interoperability. It focuses on the ability to reuse or move components of an application or service across different systems and platforms.
upvoted 1 times

TraceSplice 1 year ago
Selected Answer: D
Should be D.
upvoted 1 times

Lee_Lah 1 year, 1 month ago
Selected Answer: D
Answer is D
upvoted 1 times

zxccvbnm 1 year, 10 months ago
Selected Answer: D
Should be D
upvoted 1 times

Pika26 1 year, 11 months ago
Selected Answer: D
Answer is D. Portability.
upvoted 1 times

Brittle 2 years ago
D for me
upvoted 1 times

xav1er 2 years, 11 months ago
Selected Answer: D
right answer is: Portability (D)

Definition of portability from ISC2 materials:
Portability
Portability defines the ease with which application components are moved and reused elsewhere regardless of the provider, platform, OS, infrastructure, location, storage, format of data, or APIs. Portability is a key aspect to consider when selecting cloud providers, since it can both help prevent vendor lock-in and deliver business benefits by allowing identical cloud deployments to occur in different cloud provider solutions, either for the purposes of disaster recovery or for the global deployment of a distributed single solution. Again, think of car components. Light bulbs, brakes, and other standard components could be switched out, yet the car would continue to function.
upvoted 5 times

sedr 2 years, 8 months ago
D is Correct.Page 47 CCSP Official Student Guide 4th edition
upvoted 2 times

Odenkyem 3 years, 9 months ago
Interoperability can be defined as a measure of the degree to which diverse systems or components can work together successfully. More formally, IEEE and ISO define interoperability as the ability for two or more systems or applications to exchange information and mutually use the information that has been exchanged. In the context of cloud computing, interoperability should be viewed as the capability of

public cloud services, private cloud services, and other diverse systems within the enterprise to understand each other’s application and service interfaces, configuration, forms of authentication and authorization, data formats, etc. in order to work with each other

upvoted 4 times

nidoz 4 years, 2 months ago

Given Answer is correct.

Interoperability defines how easy it is to move and reuse application components regardless of the provider, platform, OS, infrastructure, location, storage, format of data or APIs, how well applications work together, and how well new applications work with other solutions present in the business, organization, or provider’s existing architecture

upvoted 4 times

Question #154 *Topic 1*

Which of the following is a restriction that can be enforced by information rights management (IRM) that is not possible for traditional file system controls?

A. Delete

B. Modify

C. Read

D. Print

Suggested Answer: *D*

IRM allows an organization to control who can print a set of information. This is not be possible under traditional file system controls, where if a user can read a file, they are able to print it as well.

Community vote distribution

D (100%)

kepalon 6 months, 1 week ago

Selected Answer: D

D- Print is correct.

With file system controls you can limit the rights to only read, write or modify, once you can read you will be able to print. Therefore PRINT is the correct answer.

upvoted 4 times

leaf37 2 years, 1 month ago

All of them can be enforced by IRM that are not possible with traditional file management systems controls.

upvoted 1 times

Question #155 *Topic 1*

What strategy involves hiding data in a data set to prevent someone from identifying specific individuals based on other data fields present?

A. Anonymization

B. Tokenization

C. Masking

D. Obfuscation

Suggested Answer: *A*

With data anonymization, data is manipulated in such a way so as to prevent the identification of an individual through various data objects, and is often used in conjunction with other concepts such as masking.

Community vote distribution

A (100%)

MaciekMT 1 month, 1 week ago

Selected Answer: A

Anonymization involves modifying or removing personally identifiable information from a data set, so that individuals cannot be identified—even when other data fields are available that might otherwise allow for re-identification.

upvoted 1 times

akg001 4 months, 1 week ago

Selected Answer: A

A. Anonymization

upvoted 1 times

Question #156 *Topic 1*

What type of security threat is DNSSEC designed to prevent?

A. Account hijacking

B. Snooping

C. Spoofing

D. Injection

Suggested Answer: *C*

DNSSEC is designed to prevent the spoofing and redirection of DNS resolutions to rogue sites.

Community vote distribution

C (100%)

akg001 4 months, 1 week ago

Selected Answer: C

C. Spoofing

upvoted 1 times

Question #157 *Topic 1*

Which European Union directive pertains to personal data privacy and an individual's control over their personal data?

A. 99/9/EC

B. 95/46/EC

C. 2000/1/EC

D. 2013/27001/EC

Suggested Answer: *B*

Directive 95/46/EC is titled "On the protection of individuals with regard to the processing of personal data and on the free movement of such data."

Community vote distribution

B (100%)

ichnos Highly Voted 2 years ago

This question is outdated

Directive 95/46/EC is repealed with effect from 25 May 2018.
The GDPR supersedes the Data Protection Directive and will fully phase out the DPD and become national law for all EU Member States by May 25, 2018.

upvoted 12 times

akg001 Most Recent 4 months, 1 week ago

Selected Answer: B

B. 95/46/EC

upvoted 1 times

Question #158 *Topic 1*

Which of the cloud cross-cutting aspects relates to the requirements placed on a system or application by law, policy, or requirements from standards?

A. regulatory requirements

B. Auditability

C. Service-level agreements

D. Governance

Suggested Answer: *A*

Regulatory requirements are those imposed upon businesses and their operations either by law, regulation, policy, or standards and guidelines. These requirements are specific either to the locality in which the company or application is based or to the specific nature of the data and transactions conducted.

Community vote distribution

hanyahmed Highly Voted 1 year, 9 months ago

Selected Answer: A

A is right answer, Compliance is related to laws and standards it is no Governance

upvoted 5 times

MaciekMT Most Recent 1 month, 1 week ago

Selected Answer: A

This aspect deals with the legal, policy, and standards-driven constraints placed on systems and applications. It ensures that the system complies with external mandates, such as data protection laws or industry-specific regulations, which is distinct from internal policies like governance or operational requirements like SLAs and auditability.

upvoted 1 times

TraceSplice 6 months, 1 week ago

Selected Answer: A

should be A . Compliance is related to laws, regulation and standards

upvoted 1 times

Pika26 1 year, 5 months ago

Selected Answer: A

Answer is A. Clue: Law

upvoted 2 times

samsom 2 years, 2 months ago

Governance is the process by which decisions get made, implemented and enforced--it is internal to the pools. In contrast, regulation is how governments review and change the decisions of pools--it is external to the pools.

upvoted 1 times

kepalon 2 years, 6 months ago

Selected Answer: D

I vote for Governance

Let me share this from Official Study book:

It is important to not confuse the concepts of “governance” and “corporate governance.” Governance refers to the legal and regulatory mandates of regions and countries. Corporate governance is the relationship between shareholders and other stakeholders in the organization versus the senior management of the corporation.

In this case Governance could include both types of Governance!!!

upvoted 2 times

akg001 2 years, 4 months ago

agree Confirmed D

upvoted 1 times

pooppants 3 years, 5 months ago

Confirmed D

https://vceguide.com/which-of-the-cloud-cross-cutting-aspects-relates-to-the-oversight-of-processes-and-systems-as-well-as-to-ensuring-their-compliance-with-specific-policies-and-regulations/

upvoted 1 times

topcat 3 years, 4 months ago

that link is not the same question

upvoted 4 times

pooppants 3 years, 5 months ago

I thought that regulations were specifically to do with law and not standards

upvoted 1 times

Question #159 *Topic 1*

Which data point that auditors always desire is very difficult to provide within a cloud environment?

A. Access policy

B. Systems architecture

C. Baselines

D. Privacy statement

Suggested Answer: *B*

Cloud environments are constantly changing and often span multiple physical locations. A cloud customer is also very unlikely to have knowledge and insight into the underlying systems architecture in a cloud environment. Both of these realities make it very difficult, if not impossible, for an organization to provide a comprehensive systems design document.

Community vote distribution

B (100%)

JimmyHarckins 2 weeks, 6 days ago

Selected Answer: C

Dynamic nature of cloud makes difficult to provide baselines

upvoted 1 times

MaciekMT 1 month, 1 week ago

Selected Answer: B

Auditors typically desire a comprehensive view of a system's architecture to understand all the security controls and data flows. However, in a cloud environment, the detailed underlying infrastructure is often abstracted away or considered proprietary, making it very difficult for auditors to obtain this level of detailed information.

upvoted 1 times

akg001 4 months, 1 week ago

Selected Answer: B

B. Systems architecture

upvoted 1 times

Question #160 *Topic 1*

What type of host is exposed to the public Internet for a specific reason and hardened to perform only that function for authorized users?

A. Proxy

B. Bastion

C. Honeypot

D. WAF

Suggested Answer: *B*

A bastion host is a server that is fully exposed to the public Internet, but is extremely hardened to prevent attacks and is usually dedicated for a specific application or usage; it is not something that will serve multiple purposes. This singular focus allows for much more stringent security hardening and monitoring.

Community vote distribution

B (100%)

akg001 4 months, 1 week ago

Selected Answer: B

B. Bastion

upvoted 1 times

Question #161 *Topic 1*

Which security concept is focused on the trustworthiness of data?

A. Integrity

B. Availability

C. Nonrepudiation

D. Confidentiality

Suggested Answer: *A*

Integrity is focused on the trustworthiness of data as well as the prevention of unauthorized modification or tampering of it. A prime consideration for maintaining integrity is an emphasis on the change management and configuration management aspects of operations, so that all modifications are predictable, tracked, logged, and verified, whether they are performed by actual human users or systems processes and scripts.

Community vote distribution

Kneebee 5 months, 2 weeks ago

Integrity is directly related to the trustworthiness of data - it refers to the accuracy of the data. It ensures the data isn't corrupted or tampered with during storage, retrieval, or processing.

upvoted 2 times

ikamalbhatt 1 year, 4 months ago

Selected Answer: A

Integrity ensures that data is accurate, complete, and unaltered during transmission, storage, or processing. It involves maintaining the consistency, reliability, and validity of data, and protecting it from unauthorized modifications or deletions.

upvoted 1 times

Pika26 1 year, 5 months ago

Selected Answer: A

Answer is A. Integrity.

upvoted 1 times

luckflying 1 year, 6 months ago

Selected Answer: C

The purpose of Nonrepudiation is to ensure the data can be trusted because data has been hashed and no issue of integrity.

upvoted 1 times

akg001 2 years, 4 months ago

Selected Answer: A

A. Integrity

upvoted 1 times

Question #162 *Topic 1*

Which OSI layer does IPsec operate at?

A. Network

B. transport

C. Application

D. Presentation

Suggested Answer: *A*

A major difference between IPsec and other protocols such as TLS is that IPsec operates at the Internet network layer rather than the application layer, allowing for complete end-to-end encryption of all communications and traffic.

Community vote distribution

A (100%)

MaciekMT 1 month, 1 week ago

Selected Answer: B

Hold it. I'm confused now as Most IPsec implementations use UDP as the primary transport protocol, which makes it Transport layer and not Network

upvoted 1 times

MaciekMT 1 month, 1 week ago

Selected Answer: A

IPsec operates at the Network layer (Layer 3) of the OSI model. It secures IP communications by authenticating and encrypting each IP packet, making it a critical component for protecting data in transit across IP networks.

upvoted 1 times

akg001 4 months, 1 week ago

Selected Answer: A

A. Network

upvoted 1 times

Question #163 *Topic 1*

Which of the cloud cross-cutting aspects relates to the requirements placed on the cloud provider by the cloud customer for minimum performance standards and requirements that must be met?

A. Regulatory requirements

B. SLAs

C. Auditability

D. Governance

Suggested Answer: *B*

Whereas a contract spells out general terms and costs for services, the SLA is where the real meat of the business relationship and concrete requirements come into play. The SLA spells out in clear terms the minimum requirements for uptime, availability, processes, customer service and support, security controls and requirements, auditing and reporting, and potentially many other areas that define the business relationship and the success of it.

Community vote distribution

B (100%)

akg001 4 months, 1 week ago

Selected Answer: B

B. SLAs

upvoted 1 times

Question #164 *Topic 1*

Which of the following service capabilities gives the cloud customer the most control over resources and configurations?

A. Desktop

B. Platform

C. Infrastructure

D. Software

Suggested Answer: *C*

The infrastructure service capability gives the cloud customer substantial control in provisioning and configuring resources, including processing, storage, and network resources.

Community vote distribution

C (100%)

akg001 4 months, 1 week ago

Selected Answer: C

C. Infrastructure

upvoted 2 times

Question #165 *Topic 1*

What concept does the "I" represent with the STRIDE threat model?

A. Integrity

B. Information disclosure

C. IT security

D. Insider threat

Suggested Answer: *B*

Perhaps the biggest concern for any user is having their personal and sensitive information disclosed by an application. There are many aspects of an application to consider with security and protecting this information, and it is very difficult for any application to fully ensure security from start to finish. The obvious focus is on security within the application itself, as well as protecting and storing the data.

Community vote distribution

B (100%)

zaqwsx Highly Voted 5 months, 2 weeks ago

correct

Spoofing
Tampering
Repudiation
Information disclosure (privacy breach or data leak)
Denial of service
Elevation of privilege[4]

upvoted 5 times

akg001 Most Recent 4 months, 1 week ago

Selected Answer: B

B. Information disclosure

upvoted 2 times

Question #166 *Topic 1*

At which stage of the BCDR plan creation phase should security be included in discussions?

A. Define scope

B. Analyze

C. Assess risk

D. Gather requirements

Suggested Answer: *A*

Security should be included in discussions from the very first phase when defining the scope. Adding security later is likely to incur additional costs in time and money, or will result in an incomplete or inadequate plan.

Community vote distribution

A (100%)

kepalon Highly Voted 1 year, 6 months ago

Selected Answer: A

I will go for A: Fist stage is the Scope, and as usual we should include security from the beginning

upvoted 7 times

MaciekMT Most Recent 1 month, 1 week ago

Selected Answer: D

Security considerations should be integrated right from the beginning of the BCDR planning process. During the requirements gathering phase, stakeholders—including security teams—identify the controls, compliance needs, and protective measures necessary for the organization's critical systems and data. Addressing security at this stage ensures that the subsequent analysis, risk assessment, and planning phases are built upon a foundation that properly reflects security requirements, rather than treating security as an afterthought.

upvoted 1 times

Pika26 5 months, 1 week ago

Selected Answer: A

Answer is A.

upvoted 1 times

DA95 9 months, 3 weeks ago

Security should be included in discussions during the assess risk stage of the BCDR plan creation phase. This is because at this stage, the organization is identifying and analyzing potential risks to the organization and its assets, and determining how those risks should be mitigated. Therefore, the correct answer is C. Assess risk.

upvoted 2 times

acomsa 1 year, 9 months ago

I would say that A is the correct answer. You first need to see if BCDR is in the scope and what type is needed depending on the business requirements, and only after that you start to gather the requirements

upvoted 3 times

xaccan 1 year, 11 months ago

D is correct

upvoted 4 times

moutaz1983 1 year, 11 months ago

I think it is (D) … Gathering requirement is the first phase should be

upvoted 3 times

Banzaaai 2 years ago

why not D ?

upvoted 2 times

Question #167 *Topic 1*

Which approach is typically the most efficient method to use for data discovery?

A. Metadata

B. Content analysis

C. Labels

D. ACLs

Suggested Answer: *A*
Metadata is data about data. It contains information about the type of data, how it is stored and organized, or information about its creation and use.

Community vote distribution
A (100%)

MaciekMT 1 month, 1 week ago

Selected Answer: A
Metadata provides descriptive information about data—such as file type, creation date, owner, and tags—making it much more efficient for quickly locating and categorizing data compared to the more resource-intensive process of content analysis.

Labels can indeed be useful for data discovery, but they typically need to be applied manually or through less consistent processes. Metadata, on the other hand, is automatically generated and includes a wide range of details (like file size, creation date, file type, and sometimes even labels). This breadth of information makes metadata a more efficient and comprehensive method for automated data discovery, as it doesn't rely solely on the presence of human-assigned labels.
upvoted 1 times

rkumar16d 5 months ago
Labels
upvoted 1 times

rkumar16d 5 months ago
Should be Labels
upvoted 1 times

Dasccsp 6 months, 1 week ago
Should be Labels.
upvoted 2 times

jimmyraj18 11 months ago
C.Labels
upvoted 2 times

akg001 2 years, 10 months ago
Selected Answer: A
A. Metadata
upvoted 2 times

Question #168 *Topic 1*

Which of the following features is a main benefit of PaaS over IaaS?

A. Location independence

B. High-availability

C. Physical security requirements

D. Auto-scaling

Suggested Answer: *D*

With PaaS providing a fully configured and managed framework, auto-scaling can be implemented to programmatically adjust resources based on the current demands of the environment.

Community vote distribution

D (100%)

MaciekMT 1 month, 1 week ago

Selected Answer: D

PaaS platforms are designed to streamline application development and management, and one of their standout features is built-in auto-scaling. This means your application can automatically adjust resource allocation in response to demand, without you having to manually configure scaling rules as you would with IaaS. This abstraction lets developers focus on code and business logic, leaving the heavy lifting of resource management to the platform.

upvoted 1 times

gayan237 9 months ago

Location Independence is something you will definitely not get with IaaS. you can auto-scale in IaaS based on how you setup the systems, In PaaS coders only have to focus on building the application rather than thinking about the location cz the Could provider will take care of the underlining infrastructure.

upvoted 1 times

Pika26 10 months, 1 week ago

Selected Answer: D

D: Auto-scaling

upvoted 1 times

DA95 1 year, 3 months ago

A main benefit of PaaS over IaaS is the ability to focus on building and deploying applications, rather than managing the underlying infrastructure. This is often referred to as "location independence" (A. Location independence), because it allows developers to focus on their code and not on the underlying infrastructure, such as servers, storage, and networking. This can save time and resources and enable developers to be more productive.

upvoted 2 times

akg001 1 year, 10 months ago

Selected Answer: D

D. Auto-scaling

upvoted 1 times

kepalon 2 years ago

Selected Answer: D

auto-scaling is CORRECT

upvoted 2 times

bessonf 3 years, 3 months ago

B: HA is very often native on PaaS where it should implemented in IaaS.

upvoted 2 times

eliotn183 3 years, 3 months ago

you need to the CBK. D is correct.

upvoted 3 times

Question #169 *Topic 1*

Which audit type has been largely replaced by newer approaches since 2011?

A. SOC Type 1

B. SSAE-16

C. SAS-70

D. SOC Type 2

Suggested Answer: *C*

SAS-70 reports were replaced in 2011 with the SSAE-16 reports throughout the industry.

Community vote distribution

C (100%)

MaciekMT 3 weeks, 6 days ago

Selected Answer: C

SAS-70 (Statement on Auditing Standards No. 70) was an older standard used to audit service organizations' internal controls. However, it was replaced in 2011 by the SSAE-16 (Statements on Standards for Attestation Engagements No. 16), which later evolved into SSAE-18.

Why the other options are incorrect:
A. SOC Type 1 → Still in use today; focuses on a snapshot of controls at a specific point in time.
B. SSAE-16 → Introduced in 2011 to replace SAS-70, later updated to SSAE-18 in 2017.
D. SOC Type 2 → Still widely used; focuses on ongoing operational effectiveness of controls over a period of time.

upvoted 1 times

xav1er 5 months ago

Selected Answer: C

correct but a pointless one.

upvoted 3 times

AlanJP 1 year, 2 months ago

What a pointless question. Why do we need to care about an audit type that was superseded in 2011??

upvoted 2 times

ichnos 2 years ago

Correct: Answer C

The SAS 70 was a report used in the past primarily for financial reporting and was
oftentimes misused in the service provider context. The SSAE 18 standard and subsequent
SOC reports are its successors.

upvoted 4 times

guest999 2 years, 3 months ago

SSAE-16 has also been replaced by SSAE-18

upvoted 2 times

Question #170 *Topic 1*

Which of the following can be useful for protecting cloud customers from a denial-of-service (DoS) attack against another customer hosted in the same cloud?

A. Reservations

B. Measured service

C. Limits

D. Shares

Suggested Answer: *A*
Reservations ensure that a minimum level of resources will always be available to a cloud customer for them to start and operate their services. In the event of a
DoS attack against one customer, they can guarantee that the other customers will still be able to operate.

Community vote distribution
C (100%)

ichnos Highly Voted 4 years, 6 months ago
Correct Answer: A

A reservation will
typically guarantee to a cloud customer that they will have the minimum
required resources necessary to power on and operate their services within the
environment. Reservations also offer insurance against denial-of-service (DoS)
attacks or other customers using such large amounts of resources that the cloud
customer cannot operate their services.
upvoted 13 times

gayan237 1 year, 9 months ago
how would the cloud platform guarantee the "reservation" when DoS attack running wild in that environment ? instead, if you "limit" each tenant to upper threshold level, it will make sure, cloud platform will have enough available resources to allocate to other tenants.
upvoted 4 times

guest999 Highly Voted 4 years, 9 months ago
Limit is the appropriate answer, if the DoS here is referring to the over usage resources by one customer, and there by denying resource to another. If this referring to real DoS then again Reservation cannot stop it. So either way Reservation is incorrect.
upvoted 8 times

eliotn183 4 years, 3 months ago
Nope, Reservations will protect against DDOS.
upvoted 1 times

CL888 4 years, 6 months ago
I dont think so because, a Limit would will be the total amount of resources to be exhausted by the DoS, but, what guarantees that another customer will have the resources needed to run?, the reservation, so the reservation is what protects the other customers.
upvoted 5 times

GeneralGM 1 year, 12 months ago
You and I are reading the question the same... if the attack is coming from one source (current cloud customer or not), a limit will eventually block that DOS attack. The question is poorly worded and should say:
Which of the following can be useful for protecting cloud customers WHEN another customer faces a denial-of-service (DoS) attack hosted in the same cloud?
- Answer: Reservations
upvoted 2 times

GeneralGM 1 year, 12 months ago

You and I are reading the question the same... if the attack is coming from one source (current cloud customer or not), a limit will eventually block that DOS attack. The question is poorly worded and should say:
Which of the following can be useful for protecting cloud customers from a denial-of-service (DoS) attack WHEN another customer IS hosted in the same cloud?
- Answer: Reservations
upvoted 1 times

MaciekMT Most Recent 1 month, 1 week ago

Selected Answer: C

Setting limits on resource usage can help prevent one cloud customer from monopolizing resources, which might inadvertently lead to a denial-of-service situation for other tenants sharing the same infrastructure. Limits cap the maximum amount of CPU, memory, or bandwidth any single customer can use, providing a safeguard against runaway resource consumption—whether from a misbehaving application or a targeted DoS attack.
upvoted 1 times

Kneebee 6 months, 3 weeks ago

Correct Answer : C
In Cloud computing limits pertain to the constraints on resource usage for individual customers. By setting limits on the amount of resources (such as CPU, memory, bandwidth, etc.) each customer can consume, cloud providers can prevent any single customer from dominating resources and potentially affecting the performance of other customers.
upvoted 1 times

ezpzls 1 year, 7 months ago

Limits all the way
upvoted 1 times

Pika26 1 year, 10 months ago

Selected Answer: C

C: Limits
upvoted 2 times

ikamalbhatt 1 year, 11 months ago

Selected Answer: C

C Limits is the correct answer, Limits are put in order to throttle the services , hence preventing DOS attacks.
upvoted 1 times

serget12 2 years, 5 months ago

The concept of shares is used to arbitrate the issues associated with compute resource contention situations. Resource contention implies the existence of too many requests for resources based on the actual available resources currently in the system. If resource contention takes place, share values are used to prioritize compute resource access for all guests assigned a certain number of shares.
upvoted 1 times

certifiedgeek 2 years, 10 months ago

If a costumer has "limits" beyond the available "shares", they can consume resources intended to other costumers which can cause resource exhaustion. "Reservations" should prevent such exhaustion for other costumers though their resource requirements can no longer expand until this issue is fixed.
upvoted 2 times

gjjw 4 years, 10 months ago

Limits seem to be better...
upvoted 3 times

Guest4768 4 years, 10 months ago

Limits should be on the attacked customer. This question asks for the way to protect another customer. In general, you cannot limit other customers to protect a certain customer.
upvoted 3 times

Question #171 *Topic 1*

Which of the following service capabilities gives the cloud customer the least amount of control over configurations and deployments?

A. Platform

B. Infrastructure

C. Software

D. Desktop

Suggested Answer: *C*

The software service capability gives the cloud customer a fully established application, where only minimal user configuration options are allowed.

Community vote distribution

C (100%)

akg001 4 months, 1 week ago

Selected Answer: C

C. Software

upvoted 1 times

Question #172 *Topic 1*

What does the "SOC" acronym refer to with audit reports?

A. Service Origin Confidentiality

B. System Organization Confidentiality

C. Service Organizational Control

D. System Organization Control

Suggested Answer: *C*

Community vote distribution

dkd123 Highly Voted 4 years, 1 month ago

The right answer is: System and Organizational Controls.

upvoted 16 times

Sh0wMan 3 years, 11 months ago

https://www.aicpa.org/interestareas/frc/assuranceadvisoryservices/sorhome.html

upvoted 3 times

topcat Highly Voted 3 years, 10 months ago

The answer is Service Organization Controls

upvoted 8 times

ACNgo Most Recent 2 weeks, 2 days ago

Selected Answer: C

SOC in Audit

SOC stands for Service Organization Control.

It refers to a set of audit reports designed to evaluate the internal controls of service organizations, ensuring they meet specific standards for security, availability, processing integrity, confidentiality, and privacy.

Types of SOC Reports:

SOC 1: Focuses on controls related to financial reporting.

SOC 2: Focuses on controls related to security, availability, processing integrity, confidentiality, and privacy.

SOC 3: A summarized version of SOC 2, intended for general public use.

upvoted 1 times

MaciekMT 1 month, 1 week ago

Selected Answer: D

The term "SOC" in audit reports stands for System Organization Control. These reports, such as SOC 1, SOC 2, and SOC 3, evaluate a service organization's internal controls relevant to security, availability, processing integrity, confidentiality, or privacy.

upvoted 2 times

eb63e5a 1 month, 2 weeks ago

Selected Answer: C

Service Organizational Control (SOC)

upvoted 1 times

Kneebee 6 months, 3 weeks ago

Correct answer is D. System Organization Controls

System and Organizational Controls (SOC) reports are standards designed by the American Institute of Certified Public Accountants (AICPA). The AICPA guides reference the SOC acronym as System and Organizational Controls.

upvoted 1 times

e10045d 8 months, 2 weeks ago

Selected Answer: C

Service Organizational Control. Straight from their website (Domain 6.3).
https://www.isc2.org/certifications/ccsp/ccsp-certification-exam-outline
upvoted 2 times

FranklinG 1 year ago
The answer is C. In the ISC2 CCSP book it says SOC stands for Service Organization Control
upvoted 1 times

mdmdmd 1 year ago
The same book says this: System and organization controls (formerly known as service organization control) audit type and report-making D the legible answer
upvoted 2 times

FranklinG 1 year ago
Can you please point out the exact page on this book where it says it's formerly known as "Service Organization Control"
upvoted 1 times

cloudenthusiast 1 year, 1 month ago
Selected Answer: C
Service Org Control based on the book
upvoted 1 times

sdf23r23 1 year, 1 month ago
Selected Answer: C
googled
upvoted 1 times

Kneebee 1 year, 5 months ago
The CISSP Study Guide v2 & v3 (edition 3 was released in 2023) indicates SOC stands for System and Organization Controls).
upvoted 2 times

Schuiram 1 year, 10 months ago
Selected Answer: C
Ok, thats a mean one. If I look at wikipedia (or the ICS CCSP learning guidance) its "System AND Organizational Controls (see: https://en.wikipedia.org/wiki/System_and_Organization_Controls). But as seen in the wiki article its also referred to as "service organizations controls". Now actually D is not showing the AND - so techically its wrong. But it is a dirty question - like so many in CCSP.
upvoted 1 times

ikamalbhatt 1 year, 11 months ago
Selected Answer: D
should b System not Service
upvoted 1 times

Pika26 1 year, 11 months ago
Selected Answer: C
Answer is Service Org. Controls
upvoted 3 times

kylesam2017 2 years, 1 month ago
The correct answer is on page 253 of CCSP CBK Reference book (3rd Edition) as follows: "Service Organizations Controls 1 (SOC 1): These reports deal mainly with financial controls and are intended to be used primarily by CPAs that audit an entity's financial statements."
upvoted 1 times

bradseth 2 years, 1 month ago
none of the questions in CCSP appeared in the exam. THE ONE IN HERE IS NOT A DUMP BUT JUST A PRACTICE QUESTIONS
upvoted 2 times

bradseth 2 years, 1 month ago
Selected Answer: C
C crazy people
upvoted 1 times

Question #173 *Topic 1*

What does the REST API use to protect data transmissions?

A. NetBIOS

B. VPN

C. Encapsulation

D. TLS

Suggested Answer: *D*

Representational State Transfer (REST) uses TLS for communication over secured channels. Although REST also supports SSL, at this point SSL has been phased out due to vulnerabilities and has been replaced by TLS.

Community vote distribution

D (100%)

akg001 4 months, 1 week ago

Selected Answer: D

D. TLS

upvoted 1 times

zaqwsx 5 months, 2 weeks ago

it looks correct

upvoted 2 times

Question #174 *Topic 1*

What strategy involves replacing sensitive data with opaque values, usually with a means of mapping it back to the original value?

A. Masking

B. Anonymization

C. Tokenization

D. Obfuscation

Suggested Answer: *C*

Tokenization is the practice of utilizing a random and opaque "token" value in data to replace what otherwise would be a sensitive or protected data object. The token value is usually generated by the application with a means to map it back to the actual real value, and then the token value is placed in the data set with the same formatting and requirements of the actual real value so that the application can continue to function without different modifications or code changes.

Community vote distribution

kns20 Highly Voted 3 years, 1 month ago

Masking protects data in use

Tokenization typically use to protect data at rest

Obfuscation scrambles sensitive data

Anonymization permanently replaces sensitive data with a substitute value

upvoted 8 times

Ahbey_911 Highly Voted 3 years, 8 months ago

Tokenization is correct.

upvoted 5 times

TraceSplice Most Recent 6 months, 1 week ago

Selected Answer: C

C is correct - Tokenization, when applied to data security, is the process of substituting a sensitive data element with a non-sensitive equivalent, referred to as a token, that has no intrinsic or exploitable meaning or value. The token is a reference (i.e. identifier) that maps back to the sensitive data through a tokenization system.

upvoted 1 times

Lee_Lah 7 months, 3 weeks ago

Selected Answer: C

Tokenization, when applied to data security, is the process of substituting a sensitive data element with a non-sensitive equivalent, referred to as a token, that has no intrinsic or exploitable meaning or value. The token is a reference that maps back to the sensitive data through a tokenization system.

upvoted 2 times

JohnnyBG 7 months, 3 weeks ago

Selected Answer: A

Masking is the answer. Tokenisation mean the value is changed for a token that you can convert to the original value in another DB.

upvoted 1 times

akg001 2 years, 4 months ago

Selected Answer: C

C. Tokenization

upvoted 2 times

EdwardLeeBurtle 2 years, 9 months ago

Poorly worded question.

upvoted 2 times

HCL 3 years, 10 months ago

Why Masking or Obfuscation are not the answers?

Algorithmic substitution of Data Masking allows re-generation of real data too.

upvoted 1 times

evilwizardington 3 years, 7 months ago

Data masking not always is looking for back regeneration, only in few cases (as algorithmic substitution).

In contrast, tokenization is intended always for mapping back. Here the key word is: mapping.

upvoted 7 times

Question #175 *Topic 1*

With software-defined networking, what aspect of networking is abstracted from the forwarding of traffic?

A. Routing

B. Session

C. Filtering

D. Firewalling

Suggested Answer: *C*

With software-defined networking (SDN), the filtering of network traffic is separated from the forwarding of network traffic so that it can be independently administered.

Community vote distribution

A (100%)

kepalon Highly Voted 2 years, 6 months ago

Selected Answer: A

D is incorrect, A: Routing is correct,

Routing is taken over the management plane!!!

SDN contains a Management Plane & a Forwarding Plane, in the second one goes as well the Filtering

The concept of Software Defined Networking (SDN) is about taking the routing control away from the individual network elements, and putti

Nevion VideoIPath orchestration and SDN control system.

https://techex.co.uk/ip-production/software-defined-network-sdn-routing#:~:text=Software%20Defined%20Network%20%28SDN%29%20routing%20The%20concept%20of,example%20Nevion%20VideoI

upvoted 6 times

kepalon 2 years, 6 months ago

I assume, I think correctly, that ROUTING is a Networking aspect, as stated in the question.

upvoted 1 times

bdfb8cf Most Recent 3 weeks ago

Selected Answer: C

Filtering is abstracted from forwarding because SDN separates the control plane from the data plane.

upvoted 1 times

MaciekMT 3 weeks, 6 days ago

Selected Answer: A

In Software-Defined Networking (SDN), the control plane (which makes decisions about where traffic should be sent) is separated from the data plane (which forwards traffic based on those decisions). This means that routing decisions are abstracted from the actual forwarding of traffic.

In traditional networking, routing is performed by each individual network device. In SDN, however, a centralized SDN controller manages the routing logic and pushes forwarding rules to network devices dynamically.

Why the other options are incorrect:

B. Session → SDN doesn't directly manage sessions; session management occurs at higher layers.

C. Filtering → Filtering (e.g., ACLs) can be centrally controlled in SDN but is still enforced at the data plane.

D. Firewalling → Firewalls still inspect and block traffic as needed, but SDN enables more dynamic policy enforcement.

upvoted 1 times

chunmuga 3 months ago

Selected Answer: C

So the correct interpretation for the question could be that SDN abstracts filtering from the traffic forwarding process.

The controller decides which traffic should be allowed, blocked, or redirected (filtering), and the network devices apply these instructions while forwarding the traffic.

upvoted 1 times

TraceSplice 6 months, 1 week ago

Selected Answer: A

should be A

upvoted 1 times

escaprix 1 year, 3 months ago

Selected Answer: A

n traditional networking, routing decisions are typically made within network devices such as routers, which determine the path for traffic to reach its destination based on various routing protocols and algorithms. In SDN, the control plane, managed by a centralized controller, abstracts the routing functionality. It is responsible for making routing decisions and programming the forwarding behavior of network devices.

upvoted 1 times

ikamalbhatt 1 year, 4 months ago

Selected Answer: A

should be A

upvoted 1 times

Pika26 1 year, 5 months ago

Selected Answer: A

Answer is A:Routing.

upvoted 1 times

nelombg 1 year, 5 months ago

filtering traffic C

upvoted 1 times

samsom 2 years, 2 months ago

networking is abstracted from the forwarding of traffic.

Keyword is abstracted here. Thus, the answer will be filtering.

upvoted 2 times

FrankHot 2 years, 9 months ago

Should be routing

upvoted 2 times

Bill_nye_russian_guy 2 years, 10 months ago

At first I thought this was routing, but we are not routing traffic. We are forwarding so we are filtering this to a destination or port.

upvoted 2 times

Question #176 *Topic 1*

Which of the following does NOT fall under the "IT" aspect of quality of service (QoS)?

A. Applications

B. Key performance indicators (KPIs)

C. Services

D. Security

Suggested Answer: *B*
KPIs fall under the "business" aspect of QoS, along with monitoring and measuring of events and business processes. Services, security, and applications are all core components and concepts of the "IT" aspect of QoS.

MaciekMT 1 month, 1 week ago

Selected Answer: B

Key performance indicators are metrics used to measure the performance and effectiveness of IT services rather than being an inherent component of the IT aspects of quality of service. In the context of QoS, aspects such as applications, services, and security are considered part of the IT environment, while KPIs are the measures used to evaluate how well those aspects are performing.

upvoted 1 times

spencerryy 6 months ago

B - Key Performance Indicator

upvoted 1 times

Question #177 *Topic 1*

What does dynamic application security testing (DAST) NOT entail?

A. Scanning

B. Probing

C. Discovery

D. Knowledge of the system

Suggested Answer: *D*

Dynamic application security testing (DAST) is considered "black box" testing and begins with no inside knowledge of the application or its configurations.

Everything about the application must be discovered during the testing.

Community vote distribution

D (100%)

akg001 4 months, 1 week ago

Selected Answer: D

D. Knowledge of the system

upvoted 1 times

Question #178 *Topic 1*

Where is an XML firewall most commonly deployed in the environment?

A. Between the application and data layers

B. Between the IPS and firewall

C. Between the presentation and application layers

D. Between the firewall and application server

Suggested Answer: *D*

XML firewalls are most commonly deployed in line between the firewall and application server to validate XML code before it reaches the application.

Community vote distribution

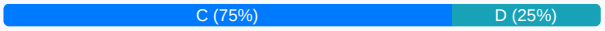

Lenell Highly Voted 2 years, 2 months ago

Selected Answer: D

XML Firewall is an appliance so A and C are eliminated. XML is an application of interest. D places the appliance in the most applicable location in the topology.

upvoted 5 times

Monchel Most Recent 3 weeks, 6 days ago

Selected Answer: D

The suggestion to put an XML firewall “between the presentation and application layers” (Option C) is not correct. An XML firewall is a network security device, not something that sits between internal software layers of an application. The CCSP guidance explicitly notes that placing an XML firewall at other points (such as between the presentation and application layers, or between an IPS and the firewall, etc.) “would not serve the intended purpose” of intercepting XML attacks. The proper location is in front of the application server (behind the external firewall), so that all XML traffic can be filtered before it hits the application. This aligns with best practices for web services security and is the placement recommended by the CCSP curriculum and other authoritative sources.

upvoted 1 times

MaciekMT 1 month, 1 week ago

Selected Answer: D

XML firewalls are specialized appliances that inspect and filter XML-based traffic—such as SOAP messages—in order to detect and block malicious XML content and attacks. They are most commonly deployed as a reverse proxy between the external firewall and the application server. This placement allows them to examine inbound XML traffic before it reaches the application server, thereby providing an additional layer of protection for XML-based web services.

upvoted 1 times

Sivath 4 months, 1 week ago

Between the firewall and application server.

upvoted 1 times

sweetykaur 5 months, 2 weeks ago

Between the firewall and application server. This placement helps in filtering and securing XML-based communications before they reach the application server.

upvoted 3 times

Mo22 11 months, 1 week ago

Selected Answer: C

An XML firewall is designed to manage and secure XML traffic, which is commonly associated with web services and API interactions. The most appropriate deployment for an XML firewall is C

upvoted 1 times

Pika26 1 year, 11 months ago

Selected Answer: C

C. Between the presentation and application layers

upvoted 2 times

Pika26 1 year, 10 months ago

REMOVE.

upvoted 1 times

DA95 2 years, 3 months ago

Selected Answer: C

A web application firewall (WAF - protected behind XML and SQL injection) is typically placed between the application and the presentation layers of the OSI (Open Systems Interconnection) model. The OSI model is a framework that is used to describe how data is transmitted over a network. It is divided into seven layers, each of which performs a specific set of functions to enable communication between devices.

upvoted 3 times

Zeezee2 3 years, 4 months ago

For reference, an XML firewall is a specialized device used to protect applications exposed through XML based interfaces like REST and scan XML traffic coming into and going out from an organization. Typically deployed in a DMZ environment an XML Firewall is often used to validate XML traffic, control access to XML based resources, filter XML content and rate limit requests to back-end applications exposed through XML based interfaces.

upvoted 3 times

Zeezee2 3 years, 4 months ago

You can't put a solution like XML firewall in the theoretical space between layer 6 & 7 of OSI model, it doesn't make sense

upvoted 4 times

Banzaaai 3 years, 6 months ago

why not C

upvoted 1 times

Question #179 *Topic 1*

What type of masking strategy involves replacing data on a system while it passes between the data and application layers?

A. Dynamic

B. Static

C. Replication

D. Duplication

Suggested Answer: *A*

With dynamic masking, production environments are protected with the masking process being implemented between the application and data layers of the application. This allows for a masking translation to take place live in the system and during normal application processing of data.

Community vote distribution

A (100%)

MaciekMT 3 weeks, 6 days ago

Selected Answer: A

Dynamic Data Masking (DDM) is a masking strategy that replaces or obfuscates data in real-time as it passes between the data and application layers. It does not alter the data at rest; instead, it modifies the data only when it is accessed by specific users or applications, ensuring that sensitive information remains hidden from unauthorized users.

Why the other options are incorrect:

B. Static → Static Data Masking (SDM) permanently replaces sensitive data at rest in a database, making it irreversible.

C. Replication → Data replication refers to copying data from one system to another for redundancy, not masking.

D. Duplication → Similar to replication, duplication means creating exact copies of data, but it does not involve masking.

upvoted 1 times

akg001 4 months, 1 week ago

Selected Answer: A

A. Dynamic

upvoted 1 times

Question #180 *Topic 1*

Which of the following is a widely used tool for code development, branching, and collaboration?

A. GitHub

B. Maestro

C. Orchestrator

D. Conductor

Suggested Answer: *A*

GitHub is an open source tool that developers leverage for code collaboration, branching, and versioning.

Community vote distribution

A (100%)

akg001 4 months, 1 week ago

Selected Answer: A

A. GitHub

upvoted 1 times

Question #181 *Topic 1*

Which aspect of security is DNSSEC designed to ensure?

A. Integrity

B. Authentication

C. Availability

D. Confidentiality

Suggested Answer: *A*

DNSSEC is a security extension to the regular DNS protocol and services that allows for the validation of the integrity of DNS lookups. It does not address confidentiality or availability at all. It allows for a DNS client to perform DNS lookups and validate both their origin and authority via the cryptographic signature that accompanies the DNS response.

Community vote distribution

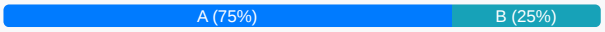

sweetykaur 5 months, 2 weeks ago

Authentication. DNSSEC (Domain Name System Security Extensions) is designed to ensure the authenticity of DNS data, preventing DNS spoofing attacks by verifying that the responses come from the correct source.

upvoted 1 times

Pika26 1 year, 11 months ago

Selected Answer: A

A. Integrity

upvoted 1 times

Rollizo 2 years, 4 months ago

Selected Answer: A

Integrity using authentication

DNSSEC definition

upvoted 2 times

xav1er 2 years, 11 months ago

Selected Answer: B

U would Say B - authentication.

DNS Security (DNSSEC)

RFC 5011 references a DNS Security DNSSEC specification that automates the trust anchor process of validating the thousands of possible DNS systems that may exist in a resolver's DNS hierarchy. The purpose of DNSSEC is to validate zone transfers with a digital signature and to thwart attackers that falsify responses to DNS queries, giving them the ability to redirect end users to bogus internet sites. On September 27, 2017, the Internet Corporation for Assigned Names and Numbers (ICANN) announced that in the first quarter of 2018, it planned to roll out a new key signing key (KSK) to support global DNSSEC. The rollout KSK-2017 was successfully completed on October 11, 2018. This represented the first change to the KSK since 2010.

upvoted 2 times

kepalon 3 years ago

Selected Answer: A

INTEGRITY!!

DNSSEC ensure the resolution you are receiving is the correct one!!!

upvoted 3 times

RadhaLaado 3 years, 6 months ago

The answer should be Authentication since DNSSEC authenticates responses to DNS Domain Name Look Ups. Authentication implies Integrity, but Integrity does not imply authentication.

upvoted 4 times

kepalon 3 years ago

DNSSEC does not provide authentication, but authenticity of the information, ensuring what you recive is the correct resolution, therefore it is INTEGRITY!!!

upvoted 1 times

Timsykcir91 3 years, 9 months ago

Is the question complete?

upvoted 2 times

Question #182 *Topic 1*

Which process serves to prove the identity and credentials of a user requesting access to an application or data?

A. Repudiation

B. Authentication

C. Identification

D. Authorization

Suggested Answer: *B*

Authentication is the process of proving whether the identity presented by a user is true and valid. This can be done through common mechanisms such as user

ID and password combinations or with more secure methods such as multifactor authentication.

Community vote distribution

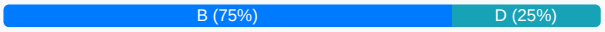

MaciekMT 1 month, 1 week ago

Selected Answer: B

Authentication is the process that verifies a user's identity by checking the credentials they provide (like a password, token, or certificate) against a trusted source. While identification involves stating who the user is, authentication is what proves that claim, ensuring that the access request is legitimate.

upvoted 1 times

Lee_Lah 7 months, 3 weeks ago

Selected Answer: B

Authentication PROVES the identity. Authorization confirms you have the right to access what you're trying to access. Answer is B - Authentication.

upvoted 1 times

BNike 9 months, 3 weeks ago

Selected Answer: D

I think Authorization is more likely the correct answer. You can be authenticated in a system, but you need authorization to access the data/application, which is another layer of security.

upvoted 1 times

BNike 9 months, 3 weeks ago

I think Authorization is more likely the correct answer. You can be authenticated in a system, but you need authorization to access the data/application, which is another layer of security.

upvoted 1 times

akg001 2 years, 4 months ago

Selected Answer: B

B. Authentication

upvoted 2 times

Question #183 *Topic 1*

Who would be responsible for implementing IPsec to secure communications for an application?

A. Developers

B. Systems staff

C. Auditors

D. Cloud customer

Suggested Answer: *B*

Because IPsec is implemented at the system or network level, it is the responsibility of the systems staff. IPsec removes the responsibility from developers, whereas other technologies such as TLS would be implemented by developers.

Community vote distribution

B (100%)

MaciekMT 1 month, 1 week ago

Selected Answer: B

Implementing IPsec involves configuring network-level security policies and ensuring proper encryption, authentication, and key management for data in transit. This task generally falls under the responsibilities of systems or network administrators, not developers, auditors, or the cloud customer.

upvoted 1 times

Pika26 5 months, 1 week ago

Selected Answer: B

B. Systems staff

upvoted 1 times

kepalon 1 year, 6 months ago

Selected Answer: B

administrators = systems staff.

IPSec, as far as I know, is provided at the system level, so it needs to be done by an admin.

upvoted 3 times

smiley 1 year, 9 months ago

What is system staff? Are they part of Cloud Customers?

upvoted 4 times

kepalon 1 year, 6 months ago

I agree that the naming is confusing, but I guess this is what happens in the real exam!!!

Sometimes you can get the answer but elimination, but for that you need to have a clear knowledge of the other 3 options.

upvoted 1 times

Question #184 *Topic 1*

What is the minimum regularity for testing a BCDR plan to meet best practices?

A. Once year

B. Once a month

C. Every six months

D. When the budget allows it

Suggested Answer: *A*

Best practices and industry standards dictate that a BCDR solution should be tested at least once a year, though specific regulatory requirements may dictate more regular testing. The BCDR plan should also be tested whenever a major modification to a system occurs.

Community vote distribution

MaciekMT 3 weeks, 6 days ago

Selected Answer: C

Best practices recommend testing a Business Continuity and Disaster Recovery (BCDR) plan at least every six months to ensure that:

Critical systems and processes remain functional in case of disruption.

New infrastructure, applications, or policies are accounted for.

Staff are trained and aware of their roles in a disaster scenario.

Why the other options are incorrect:

A. Once a year → While some organizations do annual tests, this may not be sufficient for rapidly changing IT environments.

B. Once a month → Monthly testing is too frequent and impractical for most organizations due to cost and resource constraints.

D. When the budget allows it → BCDR testing should be proactive, not dependent on budget fluctuations.

upvoted 1 times

ikamalbhatt 4 months, 3 weeks ago

Selected Answer: A

To meet best practices, the minimum regularity for testing a Business Continuity and Disaster Recovery (BCDR) plan typically involves conducting tests at least once a year.

upvoted 1 times

akg001 1 year, 4 months ago

Selected Answer: A

A. Once year

upvoted 1 times

Question #185 *Topic 1*

Other than cost savings realized due to measured service, what is another facet of cloud computing that will typically save substantial costs in time and money for an organization in the event of a disaster?

A. Broad network access

B. Interoperability

C. Resource pooling

D. Portability

Suggested Answer: *A*

With a typical BCDR solution, an organization would need some number of staff to quickly travel to the location of the BCDR site to configure systems and applications for recovery. With a cloud environment, everything is done over broad network access, with no need (or even possibility) to travel to a remote site at any time.

Community vote distribution

MaciekMT 1 month, 1 week ago

Selected Answer: C

Broad network access is one of the core characteristics of cloud computing—it ensures that services are accessible from anywhere via standard network connections. However, while it improves availability and flexibility, it doesn't directly translate into the cost savings during a disaster recovery scenario.

Resource pooling, on the other hand, is about sharing computing resources across many customers and dynamically reallocating them as needed. In the event of a disaster, the provider's ability to quickly reassign pooled resources to affected services can significantly reduce downtime and the costs associated with maintaining redundant infrastructure. This dynamic and efficient utilization of resources is what typically results in substantial time and cost savings during disaster recovery.

upvoted 1 times

Sivath 4 months, 1 week ago

Resource pooling in cloud computing allows multiple customers (tenants) to share a common pool of computing resources (e.g., storage, processing, memory, network bandwidth) dynamically allocated based on demand.

In the event of a disaster, resource pooling enables organizations to quickly access and scale up resources in the cloud without needing to invest in additional hardware or maintain expensive on-premises infrastructure. This elasticity and rapid provisioning save significant time and money during disaster recovery.

upvoted 1 times

jampol 9 months, 1 week ago

Selected Answer: C

In the event of a disaster, resource pooling enables rapid and flexible reallocation of resources to maintain business continuity, leading to substantial cost savings in both time and money.

upvoted 1 times

SteroidalRED 1 year, 3 months ago

Selected Answer: A

A: In the event of a disaster, you can continue to access your systems from anywhere.

upvoted 2 times

Omega06g 1 year, 5 months ago

The ability to work from anywhere that cloud brings

upvoted 2 times

stack120566 1 year, 8 months ago

Disaster is not defined. Broad Network access would be more valuable in situations where a disaster impacts the office space and data center, forcing users to disperse .

upvoted 1 times

Pika26 1 year, 11 months ago

Selected Answer: C

C. Resource pooling

upvoted 1 times

Zeezee2 3 years, 4 months ago

Broad network access, it eliminates the need for having either a second place of operations or having to lease it from a third party when there's a disaster. Both are very costly and are eliminated when people can work from anywhere.

upvoted 3 times

AlanJP 3 years, 8 months ago

It depends on the type of disaster, but 'broad network access' is probably the best answer here. If your premises has to shut down (e.g. Covid!) then you can still access your services from a backup location (e.g. from my bed!).

upvoted 3 times

deegadaze1 4 years ago

I would have chosen Rapid elasticity, but the nearest choice is Resource pooling.

Resource pooling-- is an IT term used in cloud computing environments to describe a situation in which providers serve multiple clients, customers or "tenants" with provisional and scalable services. These services can be adjusted to suit each client's needs without any changes being apparent to the client or end-user

upvoted 2 times

deegadaze1 4 years ago

Broad network access is something that, in a way, goes against the idea of a private cloud. However, as more employees use smartphones, tablets and other devices with online connectivity, they want to access company resources and continue to work from these devices.

upvoted 2 times

Question #186 *Topic 1*

Which of the following is NOT part of a retention policy?

A. Format

B. Costs

C. Accessibility

D. Duration

Suggested Answer: *B*

The data retention policy covers the duration, format, technologies, protection, and accessibility of archives, but does not address the specific costs of its implementation and maintenance.

Community vote distribution

B (100%)

akg001 4 months, 1 week ago

Selected Answer: B

B. Costs

upvoted 1 times

Question #187 *Topic 1*

Which aspect of cloud computing would make the use of a cloud the most attractive as a BCDR solution?

A. Interoperability

B. Resource pooling

C. Portability

D. Measured service

Suggested Answer: *D*

Measured service means that costs are only incurred when a cloud customer is actually using cloud services. This is ideal for a business continuity and disaster recovery (BCDR) solution because it negates the need to keep hardware or resources on standby in case of a disaster. Services can be initiated when needed and without costs unless needed.

Community vote distribution

D (100%)

MaciekMT 1 month, 1 week ago

Selected Answer: B

Resource pooling is a key cloud characteristic that makes it particularly attractive for BCDR solutions. It enables cloud providers to dynamically allocate and reassign resources from a shared pool, allowing organizations to quickly recover by leveraging redundant infrastructure without the need for significant capital investment. This rapid, flexible access to resources during a disaster can dramatically reduce downtime and recovery costs.

In contrast, while interoperability, portability, and measured service have their benefits, they do not directly provide the dynamic and rapid resource reallocation that is critical in a disaster recovery scenario.

upvoted 1 times

Pika26 1 year, 5 months ago

B. Resource pooling

upvoted 2 times

Pika26 1 year, 4 months ago

REMOVE

upvoted 1 times

akg001 2 years, 4 months ago

Selected Answer: D

D. Measured service

upvoted 2 times

Zeezee2 2 years, 10 months ago

Very hard to reply to, many answers are correct. Let's just trust this answer is what the exam is looking for...

upvoted 1 times

Zeezee2 2 years, 10 months ago

Looking into other questions, Portability is the second priority after the actual cost savings through measured services.

upvoted 2 times

cookiewanwan 9 months, 2 weeks ago

Agree - this one is tricky.

upvoted 1 times

Question #188 *Topic 1*

Which of the cloud deployment models offers the easiest initial setup and access for the cloud customer?

A. Hybrid

B. Community

C. Private

D. Public

Suggested Answer: *D*

Because the public cloud model is available to everyone, in most instances all a customer will need to do to gain access is set up an account and provide a credit card number through the service's web portal. No additional contract negotiations, agreements, or specific group memberships are typically needed to get started.

Community vote distribution

D (100%)

akg001 4 months, 1 week ago

Selected Answer: D

D. Public

upvoted 1 times

nelombg 1 year, 3 months ago

The answer is Public cloud

C

upvoted 3 times

Sa007788 1 year, 8 months ago

for entreprise it's more Hybrid thant public, Public answer means compagny will transfer all business to the cloud which is not so easy in the begining

upvoted 2 times

xaccan 11 months, 4 weeks ago

Read the question, the easiest is public cloud, Hybrid can also contain private cloud which is not easy to deploy resources from, it can take many days compare to public cloud which can take few minutes.

upvoted 3 times

Question #189 *Topic 1*

Which of the following is NOT something that an HIDS will monitor?

A. Configurations

B. User logins

C. Critical system files

D. Network traffic

Suggested Answer: *B*

A host intrusion detection system (HIDS) monitors network traffic as well as critical system files and configurations.

Community vote distribution

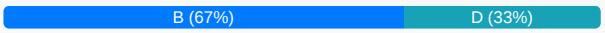

kjjcraigskel Highly Voted 4 years, 5 months ago

NIDS monitors network traffic! Not HIDS.

upvoted 23 times

kepalon 3 years ago

HIDS will check the network traffic as well, but only related to the host

upvoted 4 times

Seke Highly Voted 3 years, 3 months ago

Selected Answer: B

» Host IDS (HIDS): This type of IDS operates on a single host and monitors only *network traffic* that flows into and out of that host. In addition to monitoring a host's network traffic, HIDS are often able to monitor *critical configurations* and *files* on a host and can be configured to alert on suspicious modifications. Similar to other host-based security controls, HIDS are prone to compromise if an attacker gains root-level access on that host. To combat this, HIDS logs should immediately be sent a remote system (like your centrally managed SIEM), and HIDS configurations and settings should be locked down and managed on a remote system.

Consider installing a HIDS on your baseline images for your highly sensitive systems. Configure the HIDS to communicate with your SIEM or other centrally managed alerting dashboard. You can then deploy and manage those distributed HIDS in one fell swoop.

upvoted 10 times

JimmyHarckins Most Recent 2 weeks, 5 days ago

Selected Answer: D

HIDS: no network traffic monitoring

upvoted 1 times

bdfb8cf 3 weeks ago

Selected Answer: D

HIDS does not monitor network traffic—that is the role of a Network-Based Intrusion Detection System (NIDS).

upvoted 1 times

Monchel 3 weeks, 6 days ago

Selected Answer: D

Monitoring network packets and flows is generally the job of a Network-Based IDS (NIDS). In the CCSP Common Body of Knowledge, network traffic analysis is attributed to NIDS, not HIDS.

upvoted 1 times

MaciekMT 1 month, 1 week ago

Selected Answer: D

A Host Intrusion Detection System (HIDS) monitors the internal state of a host by tracking activities such as changes to configurations, user logins, and modifications to critical system files. Monitoring network traffic is generally handled by a Network Intrusion Detection System (NIDS), not a HIDS.

upvoted 1 times

sweetykaur 5 months, 2 weeks ago

Network traffic. Host-based Intrusion Detection Systems (HIDS) monitor configurations, user logins, and critical system files, but they don't typically monitor network traffic—that's the realm of Network-based Intrusion Detection Systems (NIDS).

upvoted 1 times

Mo22 11 months, 2 weeks ago

Selected Answer: D

NIDS dose monitor the network traffic not the HIDS

upvoted 1 times

Mo22 1 year ago

Selected Answer: D

An HIDS (Host-based Intrusion Detection System) monitors activities on a specific host or device, such as configurations, user logins, and critical system files. It does not typically monitor network traffic, which is the role of a Network-based Intrusion Detection System (NIDS).

upvoted 2 times

JohnnyBG 1 year, 1 month ago

Selected Answer: D

Host-based Intrusion Detection System [HIDS] focuses on monitoring & protecting individual hosts or devices within a network. Network-based Intrusion Detection System [NIDS] concentrates on monitoring network traffic to identify suspicious patterns & potential threats across the entire network

upvoted 1 times

escaprix 1 year, 9 months ago

Selected Answer: D

An HIDS primarily focuses on monitoring and analyzing activities occurring within the host or system itself. This includes monitoring configurations, user logins, critical system files, file integrity, process activity, and other host-specific events. The purpose of an HIDS is to detect suspicious or unauthorized activities on the host and raise alerts or take action accordingly.

While network traffic is crucial for overall security monitoring, it falls under the purview of network-based monitoring systems rather than host-based systems like HIDS

upvoted 1 times

Pika26 1 year, 11 months ago

Selected Answer: D

D. Network traffic

An HIDS (Host-based Intrusion Detection System) is designed to monitor and protect individual systems within a network by analyzing activities and events occurring on the host itself. It typically monitors configurations, user logins, and critical system files, among other things, to detect potential security threats or unauthorized activities.

upvoted 2 times

Pika26 1 year, 11 months ago

Selected Answer: D

D. Network traffic

An HIDS (Host-based Intrusion Detection System) is designed to monitor and protect individual systems within a network by analyzing activities and events occurring on the host itself. It typically monitors configurations, user logins, and critical system files, among other things, to detect potential security threats or unauthorized activities.

upvoted 1 times

infosecdummy 2 years, 4 months ago

Selected Answer: B

AIO CCSP states all but User Logins will be monitored.

upvoted 2 times

kepalon 3 years ago

Selected Answer: B

B is correct!!! HIDS will monitor the inbout/outbound traffic of the host and the rest options as well. So the only remaining one is USERS

upvoted 2 times

skis4u 3 years, 4 months ago

SHOULD BE D -- FOR NETWORK TRAFFIC WE HAVE NIDS

upvoted 1 times

AlanJP 3 years, 8 months ago

B is sort of correct - HIDS doesn't monitor all network traffic but it monitors inbound and outbound packets for the device only

upvoted 2 times

Question #190 *Topic 1*

Which of the following technologies is used to monitor network traffic and notify if any potential threats or attacks are noticed?

A. IPS

B. WAF

C. Firewall

D. IDS

Suggested Answer: *D*
An intrusion detection system (IDS) is designed to analyze network packets, compare their contents or characteristics against a set of configurations or signatures, and alert personnel if anything is detected that could constitute a threat or is otherwise designated for alerting.

Community vote distribution
D (100%)

MaciekMT 1 month, 1 week ago

Selected Answer: D
An Intrusion Detection System (IDS) is designed to monitor network traffic and alert administrators when it detects suspicious activities or potential threats. While an IPS (Intrusion Prevention System) also monitors traffic, it takes the additional step of blocking the threats, whereas an IDS is primarily focused on detection and notification.
upvoted 1 times

Pika26 5 months, 1 week ago

Selected Answer: D
D. IDS
upvoted 1 times

skis4u 1 year, 10 months ago
YES ITS ASKING JUST ABOUT MONITOR -- SO IDS WOULD BE THE BEST ANSWER
upvoted 1 times

yesj 3 years ago
Is this real question in exam. Looks like a joke. IPS and Firewalls also can "Monitor" and prevent the network traffic
upvoted 4 times

Ahbey_911 2 years, 8 months ago
Keyword is 'notify'. IDS is the only one that will notify in real time.
IPS will go ahead and prevent the payload and/or reconfigure device.
You will need to monitor the SIEM to see logs of the others.
upvoted 5 times

Sa007788 2 years, 8 months ago
Question is not about Prevent, we need to pay attention to each word in the question
upvoted 5 times

Question #191 *Topic 1*

What concept does the "A" represent in the DREAD model?

A. Affected users

B. Authentication

C. Affinity

D. Authorization

Suggested Answer: *A*

Affected users refers to the percentage of users who would be impacted by a successful exploit. Scoring ranges from 0, which means no users are impacted, to
10, which means all users are impacted.

Community vote distribution

A (100%)

budjones Highly Voted 8 months ago

Selected Answer: A

Damage – how bad would an attack be?
Reproducibility – how easy is it to reproduce the attack?
Exploitability – how much work is it to launch the attack?
Affected users – how many people will be impacted?
Discoverability – how easy is it to discover the threat?

upvoted 6 times

kepalon Most Recent 6 months, 1 week ago

Selected Answer: A

Affected Users

upvoted 4 times

Question #192 *Topic 1*

Which attribute of data poses the biggest challenge for data discovery?

A. Labels

B. Quality

C. Volume

D. Format

Suggested Answer: *B*

The main problem when it comes to data discovery is the quality of the data that analysis is being performed against. Data that is malformed, incorrectly stored or labeled, or incomplete makes it very difficult to use analytical tools against.

Community vote distribution

C (100%)

AlanJP Highly Voted 3 years, 8 months ago

I would have thought that Format is a bigger challenge - if you can't read the data it doesn't matter how good the quality is - but what do I know

upvoted 9 times

Zeezee2 3 years, 4 months ago

Indeed, but even with bad formatted data you could still have readable metadata and/or labels applied so it might not be a big issue. And Quality can mean literally f*** all in this context so let's just roll with it.

upvoted 5 times

sweetykaur Most Recent 5 months, 2 weeks ago

Volume. With the sheer amount of data generated today, the vast volumes can make data discovery challenging. It becomes difficult to sift through and manage the massive amounts of information.

upvoted 1 times

Mo22 1 year ago

Selected Answer: C

In the context of the CCSP official study materials, the biggest challenge for data discovery is identified as C. Volume. While data quality is certainly an important factor in data discovery and can pose its own challenges, the volume of data is specifically highlighted as the biggest challenge in the materials I reviewed.

upvoted 1 times

Pika26 1 year, 11 months ago

Selected Answer: C

C. Volume

upvoted 1 times

AWSPro24 3 years, 3 months ago

I think one issue here might be confusing eDiscovery - which is almost always in a legal context and "data discovery" which seems to imply multi-purpose.

upvoted 3 times

AWSPro24 3 years, 3 months ago

Study guide contains references to both but one in a section of "challenges"

pg. 122 "Big data: On big data projects, data discovery is more important and more challenging. Not only is the volume of data that must be efficiently processed for

discovery larger, but the diversity of sources and formats presents challenges that

make many traditional methods of data discovery fail."

pg 123 "Data Discovery Issues "Poor data quality: Data visualization tools are only as good as the information that is inputted. If organizations lack an enterprise-wide data governance policy, they could be relying on inaccurate or incomplete information to create their

charts and dashboards."

upvoted 3 times

Question #193 *Topic 1*

What does static application security testing (SAST) offer as a tool to the testers?

A. Production system scanning

B. Injection attempts

C. Source code access

D. Live testing

Suggested Answer: *C*
Static application security testing (SAST) is conducted with knowledge of the system, including source code, and is done against offline systems.

Community vote distribution
C (100%)

xav1er 5 months ago

Selected Answer: C
Answer C

Static Application Security Testing (SAST): This test is able to do a static
analysis
of source code. Source code is available for internally developed
software systems. Static testing will not find all vulnerabilities. SAST is a good
initial test to eliminate common vulnerabilities that can be found in this manner.
As the code is known, this is a form of white-box
testing
upvoted 2 times

zaqwsx 5 months, 2 weeks ago
correct

It is a type of white box testing method meaning they require access to source code to function. It finds all security vulnerabilities including software flaws and weaknesses such as SQL injection and others by examining code before it is deployed.
upvoted 2 times

Question #194 *Topic 1*

Which of the following service capabilities gives the cloud customer an established and maintained framework to deploy code and applications?

A. Software

B. Desktop

C. Platform

D. Infrastructure

Suggested Answer: *C*

The platform service capability provides programming languages and libraries from the cloud provider, where the customer can deploy their own code and applications into a managed and controlled framework.

Community vote distribution

C (100%)

akg001 4 months, 1 week ago

Selected Answer: C

C. Platform

upvoted 1 times

Question #195 *Topic 1*

What process is used within a cloud environment to maintain resource balancing and ensure that resources are available where and when needed?

A. Dynamic clustering

B. Dynamic balancing

C. Dynamic resource scheduling

D. Dynamic optimization

Suggested Answer: *D*

Dynamic optimization is the process through which the cloud environment is constantly maintained to ensure resources are available when and where needed, and that physical nodes do not become overloaded or near capacity, while others are underutilized.

Community vote distribution

D (100%)

kepalon Highly Voted 3 years ago

Selected Answer: D

Dynamic Optimization is a Microsoft concept for DRS- Distributed Resource Scheduler.

upvoted 5 times

KCjoe Highly Voted 4 years, 7 months ago

It is C. Dynamic resource scheduling. The answer is wrong.

upvoted 5 times

MaciekMT Most Recent 3 weeks, 6 days ago

Selected Answer: D

Dynamic optimization in cloud computing refers to the process of automatically adjusting and reallocating resources to maintain load balancing, availability, and performance. It ensures that computing resources (CPU, memory, storage, and networking) are distributed efficiently based on workload demands.

Why the other options are incorrect:

A. Dynamic clustering → Clustering refers to grouping multiple servers or nodes together for high availability and failover, but it does not actively balance resources in real time.

B. Dynamic balancing → Not a widely recognized term in cloud computing; resource balancing is part of dynamic optimization.

C. Dynamic resource scheduling → While similar, scheduling focuses on pre-allocating resources for workloads rather than real-time optimization.

upvoted 1 times

sweetykaur 5 months, 2 weeks ago

Dynamic resource scheduling. This process ensures resource balancing and availability in a cloud environment by dynamically allocating resources based on current demands.

upvoted 1 times

Pika26 1 year, 11 months ago

Selected Answer: D

D. Dynamic optimization

upvoted 1 times

likeahoss 2 years, 1 month ago

Distributed Resource Scheduling applies to clusters only whereas Dynamic Optimization spans the entire cloud

upvoted 3 times

AWSPro24 3 years, 3 months ago

I found this reference but not sure what its reference is.

https://ccsp.alukos.com/concepts/cloud#dynamic-optimization

upvoted 1 times

Ahbey_911 4 years, 1 month ago

Pay attention guys!
The option says "Dynamic Resource Scheduling"....there is no such thing.
DRS is Distributed Resource Scheduling. So, Dynamic Optimization is clearly correct.
upvoted 4 times

Zeezee2 3 years, 4 months ago
Another question seemed to refer to Dynamic Resource Scheduling though, with following explanation:
Dynamic resource scheduling (DRS) is used within all clustering systems as the method for clusters to provide high availability, scaling, management, and workload distribution and balancing of jobs and processes. From a physical infrastructure perspective, DRS is used to balance compute loads between physical hosts in a cloud to maintain the desired thresholds and limits on the physical hosts.
upvoted 7 times

xroxro 2 years, 7 months ago
Question #121Topic 1
What process is used within a clustered system to provide high availability and load balancing?
A. Dynamic balancing
B. Dynamic clustering
C. Dynamic optimization
D. Dynamic resource scheduling

Answer is D
upvoted 2 times

NYF 4 years, 3 months ago
D is the answer.
With increasing business unit IT spending on cloud services, IT leaders must prevent new risks, sprawl, cost overruns and missed SLAs.

Dynamic optimisation technology can help balance the benefits of agility with required governance controls for cloud services and virtualised infrastructure.
https://whitepapers.theregister.com/paper/view/5554/dynamic-optimization-technology-for-infrastructure-resources-and-cloud-services
upvoted 2 times

ichnos 4 years, 6 months ago
Correct answer: D

Dynamic optimization is the continual and automatic process within a cloud
environment of shifting resources and virtual machines between physical hosts
and resources to ensure a proper balance is maintained.
upvoted 3 times

CL888 4 years, 6 months ago
dynamic optimization: constantly maintaining that resources are available
dynamic resource scheduling: balance compute loads between hosts to maintain thresholds
upvoted 4 times

RVA1189 4 years, 7 months ago
How does this differ to DRS ???
upvoted 3 times

Question #196 *Topic 1*

Which value refers to the percentage of production level restoration needed to meet BCDR objectives?

A. RPO

B. RTO

C. RSL

D. SRE

Suggested Answer: *C*

The recovery service level (RSL) is a percentage measure of the total typical production service level that needs to be restored to meet BCDR objectives in the case of a failure.

Community vote distribution

C (100%)

MaciekMT 1 month, 1 week ago

Selected Answer: C

RSL stands for Recovery Service Level. It defines the percentage of production restoration needed to meet the organization's Business Continuity and Disaster Recovery (BCDR) objectives. For instance, an RSL of 95% means that during a recovery scenario, 95% of production-level performance or capacity must be restored to consider the recovery effort successful.

RPO (Recovery Point Objective) refers to the maximum acceptable amount of data loss measured in time.
RTO (Recovery Time Objective) is the target time within which systems should be restored after a disruption.
SRE is not typically used in the context of BCDR metrics.

upvoted 1 times

Pika26 5 months, 1 week ago

Selected Answer: C

C. RSL

upvoted 1 times

Niitetteh 10 months, 1 week ago

Answer: C
What is RTO RPO RSL?
RPO = Within the past 24 hours. RTO = 1 hour. MAD = 8 hours. MTTR = 41 minutes. RSL = 80% or 0.8

upvoted 4 times

akg001 1 year, 4 months ago

Selected Answer: C

C. RSL

upvoted 2 times

Pegasus_orb 1 year, 8 months ago

I don't think it's RPO, because RPO is measured in time, not a percentage

upvoted 1 times

Trax 2 years, 8 months ago

Answer A : RPO !

upvoted 2 times

Trax 2 years, 8 months ago

Forget my remark! RSL seemes correct... like question talk about a pourcentage !

upvoted 3 times

Joadeika 8 months, 2 weeks ago

https://community.spiceworks.com/topic/249491-what-is-rto-rpo-and-rsl

upvoted 1 times

Joadeika 8 months, 2 weeks ago

https://blog.triparagon.com/risk-assessment-for-disaster-recovery-and-business-interruption-planning-0-0

RSL measure what and how much of computing power is needed in an emergency condition
upvoted 1 times

nelombg 2 years, 3 months ago
I don't think so.RSL is the answer
upvoted 4 times

Ahbey_911 2 years, 8 months ago
RSL - percentage of the total production service level that has to be restored to meet the BCDR objectives (correct)
RPO - determines the maximum acceptable amount of data loss measured in time (Incorrect)
upvoted 6 times

Question #197 *Topic 1*

Over time, what is a primary concern for data archiving?

A. Size of archives

B. Format of archives

C. Recoverability

D. Regulatory changes

Suggested Answer: *C*

Over time, maintaining the ability to restore and read archives is a primary concern for data archiving. As technologies change and new systems are brought in, it is imperative for an organization to ensure they are still able to restore and access archives for the duration of the required retention period.

Community vote distribution

C (100%)

Yarly 8 months, 3 weeks ago

C. Recoverability

upvoted 1 times

akg001 1 year, 10 months ago

Selected Answer: C

C. Recoverability

upvoted 3 times

Question #198 *Topic 1*

What is an often overlooked concept that is essential to protecting the confidentiality of data?

A. Strong password

B. Training

C. Security controls

D. Policies

Suggested Answer: *B*

While the main focus of confidentiality revolves around technological requirements or particular security methods, an important and often overlooked aspect of safeguarding data confidentiality is appropriate and comprehensive training for those with access to it. Training should be focused on the safe handling of sensitive information overall, including best practices for network activities as well as physical security of the devices or workstations used to access the application.

Community vote distribution

B (100%)

akg001 4 months, 1 week ago

Selected Answer: B

B. Training

upvoted 2 times

Question #199 *Topic 1*

Which of the cloud deployment models offers the most control and input to the cloud customer as to how the overall cloud environment is implemented and configured?

A. Public

B. Community

C. Hybrid

D. Private

Suggested Answer: *D*

A private cloud model, and the specific contractual relationships involved, will give a cloud customer the most level of input and control over how the overall cloud environment is designed and implemented. This would be even more so in cases where the private cloud is owned and operated by the same organization that is hosting services within it.

Community vote distribution

D (100%)

akg001 4 months, 1 week ago

Selected Answer: D

D. Private

upvoted 1 times

Question #200 *Topic 1*

What concept does the "D" represent with the STRIDE threat model?

A. Data loss

B. Denial of service

C. Data breach

D. Distributed

Suggested Answer: *B*

Any application can be a possible target of denial-of-service (DoS) attacks. From the application side, the developers should minimize how many operations are performed for non-authenticated users. This will keep the application running as quickly as possible and using the least amount of system resources to help minimize the impact of any such attacks.

Community vote distribution

B (100%)

Zeezee2 Highly Voted 10 months, 1 week ago

Threat: Desired property

Spoofing: Authenticity

Tampering: Integrity

Repudiation: Non-repudiability

Information disclosure: Confidentiality

Denial of Service: Availability

Elevation of Privilege: Authorization

upvoted 9 times

akg001 Most Recent 4 months, 1 week ago

Selected Answer: B

B. Denial of service

upvoted 2 times

Question #201 *Topic 1*

Your boss has tasked your team with getting your legacy systems and applications connected with new cloud-based services that management has decided are crucial to customer service and offerings.
Which role would you be assuming under this directive?

A. Cloud service administrator

B. Cloud service user

C. Cloud service integrator

D. Cloud service business manager

Suggested Answer: *C*
The cloud service integrator role is responsible for connecting and integrating existing services and applications with cloud-based services.A cloud service administrator is responsible for testing, monitoring, and securing cloud services, as well as providing usage reporting and dealing with service problems. The cloud service user is someone who consumes cloud services. The cloud service business manager is responsible for overseeing the billing, auditing, and purchasing of cloud services.

Community vote distribution
C (100%)

MaciekMT 1 month, 1 week ago

Selected Answer: C
The role of a Cloud Service Integrator is all about connecting legacy systems with new cloud-based services, ensuring seamless interoperability so that the organization can offer enhanced customer service and support. This integration involves technical bridging between old and new environments—a perfect fit for the integrator role.
upvoted 1 times

akg001 4 months, 1 week ago

Selected Answer: C
C. Cloud service integrator
upvoted 1 times

Question #202 *Topic 1*

One of the main components of system audits is the ability to track changes over time and to match these changes with continued compliance and internal processes.
Which aspect of cloud computing makes this particular component more challenging than in a traditional data center?

A. Portability

B. Virtualization

C. Elasticity

D. Resource pooling

Suggested Answer: *B*
Cloud services make exclusive use of virtualization, and systems change over time, including the addition, subtraction, and reimaging of virtual machines. It is extremely unlikely that the exact same virtual machines and images used in a previous audit would still be in use or even available for a later audit, making the tracking of changes over time extremely difficult, or even impossible. Elasticity refers to the ability to add and remove resources from a system or service to meet current demand, and although it plays a factor in making the tracking of virtual machines very difficult over time, it is not the best answer in this case. Resource pooling pertains to a cloud environment sharing a large amount of resources between different customers and services. Portability refers to the ability to move systems or services easily between different cloud providers.

Community vote distribution
B (100%)

MaciekMT 1 month, 1 week ago

Selected Answer: C
While virtualization and resource pooling are foundational technologies that enable cloud computing, they don't inherently create the rapid, transient changes that complicate audit trails. Elasticity is the characteristic that allows resources to scale up or down quickly based on demand. This dynamic provisioning and deprovisioning of resources means that the environment is continuously shifting, which makes it harder to track changes over time and ensure ongoing compliance.

Virtualization abstracts hardware, but the virtual instances often remain relatively stable once provisioned.
Resource pooling means sharing resources among multiple customers, but it doesn't necessarily lead to rapid changes in individual system configurations.
Thus, it's the elastic nature of cloud resources that poses the biggest challenge for auditing over time.
upvoted 1 times

cloudenthusiast 7 months ago

Selected Answer: B
Virtualization as it is difficult for the audit trail
upvoted 1 times

gayan237 1 year, 2 months ago

so, if its virtualization only on cloud-computing, is that mean you don't have virtualization in on-prem ? resource pooling makes it tough in cloud because it leads to other clients also using the same resource (hardware CPU, RAM, space) as you which makes it tough to audit.
upvoted 3 times

Pika26 1 year, 4 months ago

Selected Answer: B
B: Virtualization
upvoted 1 times

Zeezee2 2 years, 10 months ago

Virtualization as the creation and destruction of virtualized resources happens all the time so meaningful compliance reporting becomes very difficult, even if you use a baseline for all resources.
upvoted 4 times

pooppants 3 years, 6 months ago

wow that is a tough question. I can see valid justifications for all of them!

upvoted 3 times

Question #203 *Topic 1*

In the wake of many scandals with major corporations involving fraud and the deception of investors and regulators, which of the following laws was passed to govern accounting and financial records and disclosures?

A. GLBA

B. Safe Harbor

C. HIPAA

D. SOX

Suggested Answer: *D*

The Sarbanes-Oxley Act (SOX) regulates the financial and accounting practices used by organizations in order to protect shareholders from improper practices and accounting errors.The Health Insurance Portability and Accountability Act (HIPAA) pertains to the protection of patient medical records and privacy. The
Gramm-Leach-Bliley Act (GLBA) focuses on the use of PII within financial institutions. The Safe Harbor program was designed by the US government as a way for
American companies to comply with European Union privacy laws.

Community vote distribution

MaciekMT 1 month, 1 week ago

Selected Answer: D

The Sarbanes-Oxley Act (SOX) was enacted in response to major corporate scandals, such as Enron and WorldCom, to ensure greater transparency and accountability in corporate financial reporting and disclosures. SOX imposes strict reforms to improve the accuracy and reliability of corporate disclosures and to protect investors from fraudulent practices.

upvoted 1 times

akg001 4 months, 1 week ago

Selected Answer: D

D. SOX

upvoted 1 times

Question #204 *Topic 1*

Which one of the following threat types to applications and services involves the sending of requests that are invalid and manipulated through a user's client to execute commands on the application under the user's own credentials?

A. Injection

B. Missing function-level access control

C. Cross-site scripting

D. Cross-site request forgery

Suggested Answer: *D*

A cross-site request forgery (CSRF) attack forces a client that a user has used to authenticate to an application to send forged requests under the user's own credentials to execute commands and requests that the application thinks are coming from a trusted client and user. Although this type of attack cannot be used to steal data directly because the attacker has no way of seeing the results of the commands, it does open other ways to compromise an application. Missing function-level access control exists where an application only checks for authorization during the initial login process and does not further validate with each function call. Cross-site scripting occurs when an attacker is able to send untrusted data to a user's browser without going through validation processes. An injection attack is where a malicious actor sends commands or other arbitrary data through input and data fields with the intent of having the application or system execute the code as part of its normal processing and queries.

Community vote distribution

MaciekMT 1 month, 1 week ago

Selected Answer: D

Cross-site request forgery (CSRF) involves tricking an authenticated user's browser into sending manipulated, unauthorized requests to an application. These requests are made using the user's own credentials, effectively causing the application to execute commands as if they were legitimately initiated by the user. This is distinct from injection (which manipulates input to exploit vulnerabilities) or cross-site scripting (which injects malicious scripts into web pages)

upvoted 1 times

akg001 4 months, 1 week ago

Selected Answer: D

D. Cross-site request forgery

upvoted 1 times

Question #205 *Topic 1*

Which cloud service category would be most ideal for a cloud customer that is developing software to test its applications among multiple hosting providers to determine the best option for its needs?

A. DaaS

B. PaaS

C. IaaS

D. SaaS

Suggested Answer: *B*

Platform as a Service would allow software developers to quickly and easily deploy their applications among different hosting providers for testing and validation in order to determine the best option. Although IaaS would also be appropriate for hosting applications, it would require too much configuration of application servers and libraries in order to test code. Conversely, PaaS would provide a ready-to-use environment from the onset. DaaS would not be appropriate in any way for software developers to use to deploy applications. IaaS would not be appropriate in this scenario because it would require the developers to also deploy and maintain the operating system images or to contract with another firm to do so. SaaS, being a fully functional software platform, would not be appropriate for deploying applications into.

Community vote distribution

B (100%)

akg001 4 months, 1 week ago

Selected Answer: B

B. PaaS

upvoted 2 times

Question #206 *Topic 1*

You just hired an outside developer to modernize some applications with new web services and functionality. In order to implement a comprehensive test platform for validation, the developer needs a data set that resembles a production data set in both size and composition.
In order to accomplish this, what type of masking would you use?

A. Development

B. Replicated

C. Static

D. Dynamic

Suggested Answer: *C*
Static masking takes a data set and produces a copy of it, but with sensitive data fields masked. This allows for a full data set from production for testing purposes, but without any sensitive data. Dynamic masking works with a live system and is not used to produce a distinct copy. The terms "replicated" and "development" are not types of masking.

Community vote distribution
C (100%)

MaciekMT 1 month, 1 week ago

Selected Answer: C
Static data masking involves creating a copy of the production data set where sensitive information is permanently replaced or altered. This produces a data set that closely resembles production in size and composition but is safe for use in non-production environments such as a comprehensive test platform.

Dynamic masking, on the other hand, applies changes in real time on live data and isn't as suited for creating a separate, static test copy. Development or replicated masking are not standard terms in this context.
upvoted 1 times

cloudenthusiast 7 months ago

Selected Answer: C
Static Application Security Testing (SAST) as it is not interrupting the workflow.
upvoted 1 times

DA95 1 year, 9 months ago

D. Dynamic
upvoted 2 times

akg001 2 years, 4 months ago

Selected Answer: C
C. Static
upvoted 3 times

Question #207 *Topic 1*

In order to prevent cloud customers from potentially consuming enormous amounts of resources within a cloud environment and thus having a negative impact on other customers, what concept is commonly used by a cloud provider?

A. Limit

B. Cap

C. Throttle

D. Reservation

Suggested Answer: *A*
A limit puts a maximum value on the amount of resources that may be consumed by either a system, a service, or a cloud customer. It is commonly used to prevent one entity from consuming enormous amounts of resources and having an operational impact on other tenants within the same cloud system. Limits can either be hard or somewhat flexible, meaning a customer can borrow from other customers while still having their actual limit preserved. A reservation is a guarantee to a cloud customer that a certain level of resources will always be available to them, regardless of what operational demands are currently placed on the cloud environment. Both cap and throttle are terms that sound similar to limit, but they are not the correct terms in this case.

pooppants Highly Voted 4 years ago
feels like a trick question to me...
upvoted 6 times

MaciekMT Most Recent 1 month, 1 week ago
Selected Answer: A
Cloud providers commonly set limits on resource usage to prevent any single customer from consuming excessive resources that could negatively affect others. While terms like "cap" or "throttle" are sometimes used in similar contexts, "limit" is the standard concept that directly sets a maximum allowable consumption for a given resource.
upvoted 1 times

sweetykaur 5 months, 2 weeks ago
Throttle. By throttling, cloud providers can limit the amount of resources any single customer can consume at a given time, ensuring fair resource distribution and minimizing the impact on other customers.
upvoted 1 times

Question #208 *Topic 1*

Where is a DLP solution generally installed when utilized for monitoring data at rest?

A. Network firewall

B. Host system

C. Application server

D. Database server

Suggested Answer: *B*

To monitor data at rest appropriately, the DLP solution would be installed on the host system where the data resides. A database server, in some situations, may be an appropriate answer, but the host system is the best answer because a database server is only one example of where data could reside. An application server processes data and typically sits between the data and presentation zones, and as such, does not store data at rest. A network firewall would be more appropriate for data in transit because it is not a place where data would reside.

Community vote distribution

B (100%)

MaciekMT 1 month, 1 week ago

Selected Answer: B

DLP solutions that monitor data at rest are generally deployed on host systems. This allows them to scan local file systems and storage to identify and protect sensitive data before it can be leaked. While application or database servers may store data, the DLP agent typically operates at the host level for comprehensive coverage.

upvoted 1 times

akg001 4 months, 1 week ago

Selected Answer: B

B. Host system

upvoted 2 times

Question #209 *Topic 1*

Which of the following aspects of security is solely the responsibility of the cloud provider?

A. Regulatory compliance

B. Physical security

C. Operating system auditing

D. Personal security of developers

Suggested Answer: *B*

Regardless of the particular cloud service used, physical security of hardware and facilities is always the sole responsibility of the cloud provider. The cloud provider may release information about their physical security policies and procedures to ensure any particular requirements of potential customers will meet their regulatory obligations. Personal security of developers and regulatory compliance are always the responsibility of the cloud customer. Responsibility for operating systems, and the auditing of them, will differ based on the cloud service category used.

Community vote distribution

B (100%)

akg001 4 months, 1 week ago

Selected Answer: B

B. Physical security

upvoted 1 times

Question #210 *Topic 1*

Humidity levels for a data center are a prime concern for maintaining electrical and computing resources properly as well as ensuring that conditions are optimal for top performance.
Which of the following is the optimal humidity level, as established by ASHRAE?

A. 20 to 40 percent relative humidity

B. 50 to 75 percent relative humidity

C. 40 to 60 percent relative humidity

D. 30 to 50 percent relative humidity

Suggested Answer: *C*
The American Society of Heating, Refrigeration, and Air Conditioning Engineers (ASHRAE) recommends 40 to 60 percent relatively humidity for data centers.
None of these options is the recommendation from ASHRAE.

Community vote distribution
C (100%)

akg001 4 months, 1 week ago

Selected Answer: C

C. 40 to 60 percent relative humidity

upvoted 1 times

Question #211 *Topic 1*

Within a SaaS environment, what is the responsibility on the part of the cloud customer in regard to procuring the software used?

A. Maintenance

B. Licensing

C. Development

D. Purchasing

Suggested Answer: *B*

Within a SaaS implementation, the cloud customer licenses the use of the software from the cloud provider because SaaS delivers a fully functional application to the customer. With SaaS, the cloud provider is responsible for the entire software application and any necessary infrastructure to develop, run, and maintain it.

The purchasing, development, and maintenance are fully the responsibility of the cloud provider.

Community vote distribution

bessonf Highly Voted 3 years, 9 months ago

in SaaS model, the licence is generally included and the customer is purshasing to the cloud provider the right to use the software. Cloud provider is bundling the licence into its model. for me: D

upvoted 11 times

NYF Highly Voted 3 years, 9 months ago

The answer should be Purchasing the License.

upvoted 8 times

MaciekMT Most Recent 1 month, 1 week ago

Selected Answer: D

In a SaaS (Software as a Service) environment, the cloud provider is responsible for procuring, maintaining, and licensing the software used. The customer simply subscribes to and uses the service. This means that the cloud customer does not have to handle the purchasing (or other procurement-related activities) for the software.

So, while the answer choice "Purchasing" might seem like an activity, in the context of SaaS, it's the provider—not the customer—who is responsible for procuring the software.

upvoted 1 times

Kneebee 5 months, 1 week ago

B. Licensing is the correct answer. In a traditional sense, software isn't purchased in a SaaS model. The customer pays a subscription fee; the software remains the property of the CSP.

upvoted 1 times

Pika26 1 year, 5 months ago

Selected Answer: B

B. Licensing

upvoted 2 times

Lenell 1 year, 8 months ago

Selected Answer: B

You are SUBSCRIBING to the license which gives you access to usage of the software. The CSP purchases the software to make it available to its customers. Therefore, the CSP is responsible for its availability (maintenance, development, patches, updates, etc.)

upvoted 2 times

Sven007 1 year, 10 months ago

SaaS is a temporary contract - it's not that you buy something

upvoted 1 times

xav1er 2 years, 5 months ago

Selected Answer: B

Proper answer is LICENSING
after ISC2 materials:
Cloud Service Capabilities
Application and software licensing: Customers no longer need to purchase licenses, support, and associated costs, as licensing is “leased” and is relevant only when in use (covered by the provider). Additionally, purchasing of bulk licensing and the associated CapEx is removed and replaced by a pay-per-use licensing model.
upvoted 4 times

kepalon 2 years, 6 months ago

Selected Answer: B

You do not need to purchase any sw, you pay for a licence, example per user to use the SW provided by the SaaS
upvoted 3 times

serget12 1 year, 11 months ago

So purchasing?
upvoted 1 times

Sven007 1 year, 10 months ago

no, you are licensing the / leasing the software. When you purchase something, you own it. In SaaS, you don't own the software, you just have a non-exclusive, temporary right to use a service
upvoted 3 times

AWSPro24 2 years, 9 months ago

Selected Answer: D

SaaS products are rarely licenses and almost always "subscriptions". I have to go with D here.
upvoted 4 times

xaccan 2 years, 11 months ago

purchasing is the only correct answer
upvoted 1 times

evilwizardington 3 years, 7 months ago

From CBK: Application and software licensing: Customers no longer need to purchase licenses, support, and associated costs because licensing is leased and is relevant only when in use (covered by the provider). Additionally, purchasing of bulk licensing and the associated CapEx is removed and replaced by a pay-per-use licensing model.

I think the question and answers are not correctly done. But I see why it is licensing here.
upvoted 4 times

Ahbey_911 3 years, 8 months ago

The CSC pay (i.e purchase) the CSP to be licensed to use the SaaS. "Purchasing" is the only reasonable answer here.
upvoted 2 times

Question #212 *Topic 1*

Implementing baselines on systems would take an enormous amount of time and resources if the staff had to apply them to each server, and over time, it would be almost impossible to keep all the systems in sync on an ongoing basis.
Which of the following is NOT a package that can be used for implementing and maintaining baselines across an enterprise?

A. Puppet

B. SCCM

C. Chef

D. GitHub

Suggested Answer: *D*
GitHub is a software development platform that serves as a code repository and versioning system. It is solely used for software development and would not be appropriate for applying baselines to systems. Puppet is an open-source configuration management tool that runs on many platforms and can be used to apply and maintain baselines. The Software Center Configuration Manager (SCCM) was developed by Microsoft for managing systems across large groups of servers.
Chef is also a system for maintaining large groups of systems throughout an enterprise.

Community vote distribution
D (100%)

akg001 4 months, 1 week ago
Selected Answer: D
D. GitHub
upvoted 1 times

whiteyeti9898 1 year, 7 months ago
so funny how they continuously use GitHub as a throwaway answer
upvoted 4 times

Zeezee2 10 months, 1 week ago
Someone wasn't happy that Microsoft bought it :D
upvoted 2 times

NileshGavali 2 years, 2 months ago
repeat question
upvoted 2 times

Question #213 *Topic 1*

From the perspective of compliance, what is the most important consideration when it comes to data center location?

A. Natural disasters

B. Utility access

C. Jurisdiction

D. Personnel access

Suggested Answer: *C*

Jurisdiction will dictate much of the compliance and audit requirements for a data center. Although all the aspects listed are very important to security, from a strict compliance perspective, jurisdiction is the most important. Personnel access, natural disasters, and utility access are all important operational considerations for selecting a data center location, but they are not related to compliance issues like jurisdiction is.

Community vote distribution

C (100%)

Sivath 4 months ago

From a compliance perspective, jurisdiction is the most critical consideration for data center location because it determines the laws and regulations governing the data stored and processed within that region.

upvoted 1 times

akg001 2 years, 10 months ago

Selected Answer: C

C. Jurisdiction

upvoted 3 times

Question #214 *Topic 1*

Different certifications and standards take different approaches to data center design and operations. Although many traditional approaches use a tiered methodology, which of the following utilizes a macro-level approach to data center design?

A. IDCA

B. BICSI

C. Uptime Institute

D. NFPA

Suggested Answer: *A*

The Infinity Paradigm of the International Data Center Authority (IDCA) takes a macro-level approach to data center design. The IDCA does not use a specific, focused approach on specific components to achieve tier status. Building Industry Consulting Services International (BICSI) issues certifications for data center cabling. The National Fire Protection Association (NFPA) publishes a broad range of fire safety and design standards for many different types of facilities. The
Uptime Institute publishes the most widely known and used standard for data center topologies and tiers.

Community vote distribution

A (100%)

MaciekMT 1 month, 1 week ago

Selected Answer: A

The International Data Center Authority (IDCA) takes a macro-level approach to data center design. Unlike traditional tiered methodologies, which primarily focus on the technical and physical aspects of data center infrastructure, the IDCA framework emphasizes the overall business impact, operational efficiency, and holistic design principles that support a data center's performance and resiliency.

The Uptime Institute is well-known for its tiered approach to data center design. Their tier system (Tier I through Tier IV) categorizes data centers based on redundancy and fault tolerance, helping organizations understand the expected level of uptime and resiliency.

upvoted 1 times

Pika26 5 months, 1 week ago

Selected Answer: A

A. IDCA

upvoted 1 times

DA95 9 months, 3 weeks ago

C. Uptime Institute

upvoted 2 times

akg001 1 year, 4 months ago

Selected Answer: A

A. IDCA

upvoted 1 times

Question #215 *Topic 1*

The European Union is often considered the world leader in regard to the privacy of personal data and has declared privacy to be a "human right."
In what year did the EU first assert this principle?

A. 1995

B. 2000

C. 2010

D. 1999

Suggested Answer: *A*
The EU passed Directive 95/46 EC in 1995, which established data privacy as a human right. The other years listed are incorrect.

Community vote distribution
A (100%)

Pika26 5 months, 1 week ago

Selected Answer: A

A. 1995

upvoted 1 times

akg001 1 year, 4 months ago

Selected Answer: A

A. 1995

upvoted 2 times

certifiedgeek 1 year, 4 months ago

Year asserted by this directive is no longer relevant but being familiar where "privacy is declared as a human right" is essential.

upvoted 1 times

Question #216 *Topic 1*

A DLP solution/implementation has three main components.
Which of the following is NOT one of the three main components?

A. Monitoring

B. Enforcement

C. Auditing

D. Discovery and classification

Suggested Answer: *C*
Auditing, which can be supported to varying degrees by DLP solutions, is not a core component of them. Data loss prevention (DLP) solutions have core components of discovery and classification, enforcement, and monitoring. Discovery and classification are concerned with determining which data should be applied to the DLP policies, and then determining its classification level. Monitoring is concerned with the actual watching of data and how it's used through its various stages. Enforcement is the actual application of policies determined from the discovery stage and then triggered during the monitoring stage.

Community vote distribution
C (100%)

[Removed] 8 months, 3 weeks ago
Correct as per page 58 CBK 3rd edition.
upvoted 3 times

AWSPro24 9 months ago
Selected Answer: C
This is a pretty lame question. Most modern DLP does support some auditing so you can see how you are doing against compliance targets. Obviously it can't track false negatives. But it can track positives so you can see if your training is paying off, for example. And here's a quote supporting this

"DLP also provides reporting to meet compliance and auditing requirements and identify areas of weakness and anomalies for forensics and incident response."

https://digitalguardian.com/blog/what-data-loss-prevention-dlp-definition-data-loss-prevention

I'm picking C as well, but I don't like it and would like to see a much more wrong answer.
upvoted 2 times

Question #217 *Topic 1*

What type of storage structure does object storage employ to maintain files?

A. Directory

B. Hierarchical

C. tree

D. Flat

Suggested Answer: *D*

Object storage uses a flat file system to hold storage objects; it assigns files a key value that is then used to access them, rather than relying on directories or descriptive filenames. Typical storage layouts such as tree, directory, and hierarchical structures are used within volume storage, whereas object storage maintains a flat structure with key values.

Community vote distribution

D (100%)

MaciekMT 1 month, 1 week ago

Selected Answer: D

Object storage uses a flat structure where each file is stored as an individual object, accompanied by metadata and a unique identifier. This design contrasts with hierarchical file systems that use directories and subdirectories to organize files.

upvoted 1 times

akg001 4 months, 1 week ago

Selected Answer: D

D. Flat

upvoted 3 times

Question #218 *Topic 1*

Which cloud storage type requires special consideration on the part of the cloud customer to ensure they do not program themselves into a vendor lock-in situation?

A. Unstructured

B. Object

C. Volume

D. Structured

Suggested Answer: *D*

Structured storage is designed, maintained, and implemented by a cloud service provider as part of a PaaS offering. It is specific to that cloud provider and the way they have opted to implement systems, so special care is required to ensure that applications are not designed in a way that will lock the cloud customer into a specific cloud provider with that dependency. Unstructured storage for auxiliary files would not lock a customer into a specific provider. With volume and object storage, because the cloud customer maintains their own systems with IaaS, moving and replicating to a different cloud provider would be very easy.

Community vote distribution

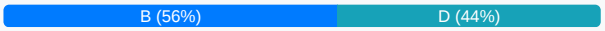

MaciekMT 1 month, 1 week ago

Selected Answer: B

Object storage systems are commonly accessed via cloud-specific APIs (e.g., Amazon S3), which can be proprietary. If a cloud customer builds their application tightly around a particular object's storage API, they risk becoming locked in to that vendor's ecosystem. Special consideration is needed to design in a way that minimizes this risk, such as abstracting the storage layer to allow for easier migration between vendors.

upvoted 1 times

Pika26 5 months, 1 week ago

Selected Answer: B

B. Object

Object storage in a cloud environment requires special consideration on the part of the cloud customer to ensure they do not program themselves into a vendor lock-in situation. Object storage systems use a unique approach to store data, which involves the use of flat address spaces and metadata associated with each object. Different cloud providers implement their own proprietary APIs and storage protocols for object storage, making it more challenging to move data between providers without significant effort or re-architecting.

upvoted 2 times

vavofa5697 7 months, 3 weeks ago

Selected Answer: B

It should be B. Object.

This is because object storage uses a flat addressing structure for files, which means that the file paths and folder structures are not stored as metadata with the files, making it difficult to switch to another vendor without significant reconfiguration of the data

upvoted 3 times

xroxro 1 year, 1 month ago

Unstructured storage for auxiliary files would not lock a customer into a specific provider

Why unstructured is limited to auxiliary files ? could be all files...

upvoted 1 times

AWSPro24 1 year, 9 months ago

Selected Answer: D

Agree with the answer but the explanation's last sentence of "With volume and object storage, because the cloud customer maintains their own systems with IaaS, moving and replicating to a different cloud provider would be very easy." is a little off. cloud storage is almost entirely PaaS not IaaS except for ephemeral drives on VMs. You are almost always connecting to a PaaS cloud service even for volumes. They are PaaS storage they are just presented or formatted as volumes.

upvoted 4 times

Question #219 *Topic 1*

Which cloud deployment model would be ideal for a group of universities looking to work together, where each university can gain benefits according to its specific needs?

A. Private

B. Public

C. Hybrid

D. Community

Suggested Answer: *D*

A community cloud is owned and maintained by similar organizations working toward a common goal. In this case, the universities would all have very similar needs and calendar requirements, and they would not be financial competitors of each other. Therefore, this would be an ideal group for working together within a community cloud. A public cloud model would not work in this scenario because it is designed to serve the largest number of customers, would not likely be targeted toward specific requirements for individual customers, and would not be willing to make changes for them. A private cloud could accommodate such needs, but would not meet the criteria for a group working together, and a hybrid cloud spanning multiple cloud providers would not fit the specifics of the question.

Community vote distribution

Lee_Lah 7 months, 3 weeks ago

Selected Answer: D

D - Community

upvoted 1 times

akg001 2 years, 4 months ago

Answer C: Hybrid? - using combination of public and private : they can collaborate and with their own private can gain benefits according to it's specific needs.

upvoted 1 times

[Removed] 2 years, 8 months ago

Correct Answer B.

Question says "where each university can gain benefits according to its specific needs?"

Need is not common/shared its specific to university, community cloud doesn't fit the into the need.

They can still collaborate using public cloud.

upvoted 1 times

zaqwsx 2 years, 5 months ago

yes but they ask about "group of universities looking to work together" if you choose public cloud you also don't meet requirements about "specific needs"

upvoted 1 times

Question #220 *Topic 1*

Data centers have enormous power resources that are distributed and consumed throughout the entire facility.
Which of the following standards pertains to the proper fire safety standards within that scope?

A. IDCA

B. BICSI

C. NFPA

D. Uptime Institute

Suggested Answer: *C*
The National Fire Protection Association (NFPA) publishes a broad range of fire safety and design standards for many different types of facilities. Building Industry
Consulting Services International (BICSI) issues certifications for data center cabling. The Uptime Institute publishes the most widely known and used standard for data center topologies and tiers. The International Data Center Authority (IDCA) offers the Infinity Paradigm, which takes a macro-level approach to data center design.

Community vote distribution
C (100%)

akg001 4 months, 1 week ago

Selected Answer: C

C. NFPA

upvoted 1 times

Question #221 *Topic 1*

Which of the following threat types involves an application that does not validate authorization for portions of itself beyond when the user first enters it?

A. Cross-site request forgery

B. Missing function-level access control

C. Injection

D. Cross-site scripting

Suggested Answer: *B*

It is imperative that applications do checks when each function or portion of the application is accessed to ensure that the user is properly authorized. Without continual checks each time a function is accessed, an attacker could forge requests to access portions of the application where authorization has not been granted. An injection attack is where a malicious actor sends commands or other arbitrary data through input and data fields with the intent of having the application or system execute the code as part of its normal processing and queries. Cross-site scripting occurs when an attacker is able to send untrusted data to a user's browser without going through validation processes. Cross-site request forgery occurs when an attack forces an authenticated user to send forged requests to an application running under their own access and credentials.

Community vote distribution

B (100%)

Morello 6 months, 3 weeks ago

Is this question still valid as Missing function-level access control is no longer listed in the OWASP top 10?

upvoted 1 times

akg001 2 years, 10 months ago

Selected Answer: B

B. Missing function-level access control

upvoted 1 times

Question #222 *Topic 1*

Clustered systems can be used to ensure high availability and load balancing across individual systems through a variety of methodologies.
What process is used within a clustered system to ensure proper load balancing and to maintain the health of the overall system to provide high availability?

A. Distributed clustering

B. Distributed balancing

C. Distributed optimization

D. Distributed resource scheduling

Suggested Answer: *D*
Distributed resource scheduling (DRS) is used within all clustered systems as the method for providing high availability, scaling, management, workload distribution, and the balancing of jobs and processes. None of the other choices is the correct term in this case.

Community vote distribution
D (100%)

Pika26 5 months, 1 week ago

Selected Answer: D

D. Distributed resource scheduling
upvoted 1 times

certifiedgeek 1 year, 4 months ago

There is no such definition for "Distributed optimization" in all CCSP references. There is however for "Dynamic" optimization.
upvoted 1 times

Banzaaai 2 years ago

why not C
upvoted 2 times

Zeezee2 1 year, 10 months ago

Dynamic optimization shifts resources and VM's between physical hosts to ensure a proper balance is maintained.
Dynamic resource scheduling will ensure that for any available resources/VMs (as provided by the dynamic optimization), the load is balanced evenly across them to ensure no over or under utilization.
upvoted 4 times

Question #223 *Topic 1*

Although the REST API supports a wide variety of data formats for communications and exchange, which data formats are the most commonly used?

A. SAML and HTML

B. XML and SAML

C. XML and JSON

D. JSON and SAML

Suggested Answer: *C*

JavaScript Object Notation (JSON) and Extensible Markup Language (XML) are the most commonly used data formats for the Representational State Transfer
(REST) API and are typically implemented with caching for increased scalability and performance. Extensible Markup Language (XML) and Security Assertion
Markup Language (SAML) are both standards for exchanging encoded data between two parties, with XML being for more general use and SAML focused on authentication and authorization data. HTML is used for authoring web pages for consumption by web browsers

Community vote distribution

C (100%)

akg001 4 months, 1 week ago

Selected Answer: C

C. XML and JSON

upvoted 1 times

Question #224 *Topic 1*

The share phase of the cloud data lifecycle involves allowing data to leave the application, to be shared with external systems, services, or even other vendors/ contractors.
What technology would be useful for protecting data at this point?

A. IDS

B. DLP

C. IPS

D. WAF

Suggested Answer: *B*
Data loss prevention (DLP) solutions allow for control of data outside of the application or original system. They can enforce granular control such as printing, copying, and being read by others, as well as forcing expiration of access. Intrusion detection system (IDS) and intrusion prevention system (IPS) solutions are used for detecting and blocking suspicious and malicious traffic, respectively, whereas a web application firewall (WAF) is used for enforcing security or other controls on web-based applications.

Community vote distribution
B (100%)

akg001 4 months, 1 week ago

Selected Answer: B

B. DLP

upvoted 1 times

zaqwsx 5 months, 1 week ago

correct

upvoted 1 times

Question #225 *Topic 1*

When an API is being leveraged, it will encapsulate its data for transmission back to the requesting party or service.
What is the data encapsulation used with the SOAP protocol referred to as?

A. Packet

B. Payload

C. Object

D. Envelope

Suggested Answer: *D*
Simple Object Access Protocol (SOAP) encapsulates its information in what is known as a SOAP envelope. It then leverages common communications protocols for transmission. Object is a type of cloud storage, but also a commonly used term with certain types of programming languages. Packet and payload are terms that sound similar to envelope but are not correct in this case.

Community vote distribution
D (100%)

akg001 4 months, 1 week ago

Selected Answer: D

D. Envelope

upvoted 1 times

Question #226 *Topic 1*

From a security perspective, what component of a cloud computing infrastructure represents the biggest concern?

A. Hypervisor

B. Management plane

C. Object storage

D. Encryption

Suggested Answer: *B*

The management plane will have broad administrative access to all host systems throughout an environment; as such, it represents the most pressing security concerns. A compromise of the management plane can directly lead to compromises of any other systems within the environment. Although hypervisors represent a significant security concern to an environment because their compromise would expose any virtual systems hosted within them, the management plane is a better choice in this case because it controls multiple hypervisors. Encryption and object storage both represent lower-level security concerns.

Community vote distribution

B (100%)

[Removed] Highly Voted 1 year, 8 months ago

Question says "cloud computing infrastructure " I would think it is either Hypervisor (multi tenancy, data bleed, VM escape) or the Encryption. storing keys etc

Management plane is more of a concern from insider threat perspective, can be easily secured using MFA and location/IP based controls to prevent external threats.

upvoted 5 times

MaciekMT Most Recent 1 month ago

Selected Answer: B

The management plane is the central control point for orchestrating and administering the entire cloud environment. If it's compromised, an attacker gains full administrative access to resources and services, making it a prime target from a security perspective. While components like the hypervisor are critical too, cloud providers often invest heavily in their hardening, whereas vulnerabilities in the management plane can have broader and more devastating consequences.

upvoted 1 times

gauravsapra 5 months ago

A. Hypervisor

upvoted 1 times

Pika26 5 months, 1 week ago

Selected Answer: B

B. Management plane

upvoted 1 times

Awraith 6 months, 2 weeks ago

I answered B as well, but Hypervisor & Encryption could definitely be answers too ...

Lame question.

upvoted 1 times

serget12 11 months, 2 weeks ago

Management plane is usually accessed from a jump box/bastion. Encryption requires usually a few components, making it susceptible to greater threats.

upvoted 1 times

akg001 1 year, 4 months ago

Selected Answer: B

B. Management plane

upvoted 1 times

Zeezee2 1 year, 10 months ago

This is the major internal threat at side of the cloud provider, but impacting all cloud customers.

upvoted 1 times

AWSPro24 1 year, 9 months ago

it's a pretty big threat to a single CSC too. If someone pops your management plane you're toast.

upvoted 2 times

Question #227 *Topic 1*

Which of the following is NOT one of the main intended goals of a DLP solution?

A. Showing due diligence

B. Preventing malicious insiders

C. Regulatory compliance

D. Managing and minimizing risk

Suggested Answer: *B*

Data loss prevention (DLP) extends the capabilities for data protection beyond the standard and traditional security controls that are offered by operating systems, application containers, and network devices. DLP is not specifically implemented to counter malicious insiders, and would not be particularly effective in doing so, because a malicious insider with legitimate access would have other ways to obtain data. DLP is a set of practices and controls to manage and minimize risk, comply with regulatory requirements, and show due diligence with the protection of data.

Community vote distribution

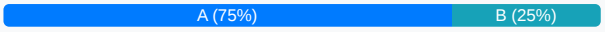

dhiru Highly Voted 3 years, 5 months ago

But DLP is an effective solution to prevent malicious insider/user from sending sensitive data out of the network.

upvoted 8 times

kjjcraigskel 2 years, 11 months ago

DLP doesn't prevent malicious insiders. It hinders them.

upvoted 6 times

Guest4768 3 years, 4 months ago

It is difficult to cover ALL data fraud cases by insiders with DLP. B is partially correct, and others are fully correct, so B is the answer.

upvoted 8 times

xaccan 1 year, 11 months ago

Malicious insider does not explicitly mean leaking sensitive data outside the company, it is a general term.

upvoted 1 times

akg001 Highly Voted 1 year, 4 months ago

Selected Answer: A

A. Showing due diligence

upvoted 5 times

MaciekMT Most Recent 1 month ago

Selected Answer: A

While a DLP solution does help demonstrate that an organization is taking proactive steps toward data protection (which can be useful in showing due diligence), its primary goals are to prevent unauthorized data exfiltration (including threats from malicious insiders), ensure regulatory compliance, and manage and minimize risk associated with data loss. "Showing due diligence" is more of a legal or reputational benefit rather than a direct technical or operational objective of the DLP solution

upvoted 1 times

Loveguitar 4 months, 1 week ago

The answer is correct (Insider threat prevention is not among the goals of a DLP system). The goals of a DLP strategy for an organization are to manage and minimize risk, maintain compliance with regulatory requirements, and show due diligence on the part of the application and data owner.

Carter, Daniel. CCSP Certified Cloud Security Professional All-in-One Exam Guide, Third Edition (p. 101). McGraw Hill LLC. Kindle Edition.

upvoted 3 times

Pika26 5 months, 1 week ago

Selected Answer: B

B. Preventing malicious insiders

upvoted 2 times

AJ2021 11 months, 1 week ago

Selected Answer: A

Not a very clear question, you could argue for both A & B to be correct, in my opinion A is correct

upvoted 1 times

serget12 11 months, 2 weeks ago

I believe the number 1 threat/ security issue has to do with internal risks.

upvoted 1 times

Biden 1 year, 7 months ago

Question is "NOT one of the main intended goals of a DLP solution?" !!

Whats does "Showing Due Diligence" have anything to do with DLP? Shudnt this be the correct answer ?

upvoted 5 times

akg001 1 year, 4 months ago

A. Showing due diligence

upvoted 2 times

[Removed] 1 year, 8 months ago

"Showing due diligence" to what ?

Question/answer is incorrect.

DLP is heavily used to protect from insider disclosures.

CCSP Official guide says "DLP can protect from malicious disclosure" which would equate malicious insider.

DLP goals

Additional security

Policy enforcement

Enhanced monitoring

Regulatory compliance

upvoted 1 times

funktribe 2 years, 3 months ago

DLP is not a breach detection technology hence B is correct

upvoted 1 times

Question #228 *Topic 1*

Data center and operations design traditionally takes a tiered, topological approach.
Which of the following standards is focused on that approach and is prevalently used throughout the industry?

A. IDCA

B. NFPA

C. BICSI

D. Uptime Institute

Suggested Answer: *D*
The Uptime Institute publishes the most widely known and used standard for data center topologies and tiers. The National Fire Protection Association (NFPA) publishes a broad range of fire safety and design standards for many different types of facilities. Building Industry Consulting Services International (BICSI) issues certifications for data center cabling. The International Data Center Authority (IDCA) offers the Infinity Paradigm, which takes a macro-level approach to data center design.

Community vote distribution
D (100%)

akg001 4 months, 1 week ago

Selected Answer: D

D. Uptime Institute
upvoted 2 times

Question #229 *Topic 1*

Jurisdictions have a broad range of privacy requirements pertaining to the handling of personal data and information.
Which jurisdiction requires all storage and processing of data that pertains to its citizens to be done on hardware that is physically located within its borders?

A. Japan

B. United States

C. European Union

D. Russia

Suggested Answer: *D*
The Russian government requires all data and processing of information about its citizens to be done solely on systems and applications that reside within the physical borders of the country. The United States, European Union, and Japan focus their data privacy laws on requirements and methods for the protection of data, rather than where the data physically resides.

Community vote distribution
D (100%)

Pika26 5 months, 1 week ago
Selected Answer: D
D. Russia
upvoted 1 times

akg001 1 year, 4 months ago
Selected Answer: D
D. Russia
upvoted 1 times

DA95 9 months, 3 weeks ago
Dear, in what Jurisdictions does the RODO apply that regulates this requirements?
upvoted 1 times

Syd 3 years, 1 month ago
Repeat question
upvoted 4 times

Zeezee2 1 year, 10 months ago
Different question nonetheless.
upvoted 1 times

Question #230 *Topic 1*

The management plane is used to administer a cloud environment and perform administrative tasks across a variety of systems, but most specifically it's used with the hypervisors.
What does the management plane typically leverage for this orchestration?

A. APIs

B. Scripts

C. TLS

D. XML

Suggested Answer: *A*
The management plane uses APIs to execute remote calls across the cloud environment to various management systems, especially hypervisors. This allows a centralized administrative interface, often a web portal, to orchestrate tasks throughout an enterprise. Scripts may be utilized to execute API calls, but they are not used directly to interact with systems. XML is used for data encoding and transmission, but not for executing remote calls. TLS is used to encrypt communications and may be used with API calls, but it is not the actual process for executing commands.

Community vote distribution
A (100%)

akg001 4 months, 1 week ago

Selected Answer: A

A. APIs

upvoted 2 times

Question #231 *Topic 1*

When dealing with PII, which category pertains to those requirements that can carry legal sanctions or penalties for failure to adequately safeguard the data and address compliance requirements?

A. Contractual

B. Jurisdictional

C. Regulated

D. Legal

Suggested Answer: *C*

Regulated PII pertains to data that is outlined in law and regulations. Violations of the requirements for the protection of regulated PII can carry legal sanctions or penalties. Contractual PII involves required data protection that is determined by the actual service contract between the cloud provider and cloud customer, rather than outlined by law. Violations of the provisions of contractual PII carry potential financial or contractual implications, but not legal sanctions. Legal and jurisdictional are similar terms to regulated, but neither is the official term used.

Community vote distribution

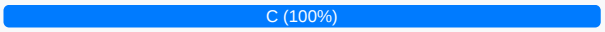

akg001 4 months, 1 week ago

Selected Answer: C

C. Regulated

upvoted 2 times

Question #232 *Topic 1*

Although the United States does not have a single, comprehensive privacy and regulatory framework, a number of specific regulations pertain to types of data or populations.
Which of the following is NOT a regulatory system from the United States federal government?

A. HIPAA

B. SOX

C. FISMA

D. PCI DSS

Suggested Answer: *D*
The Payment Card Industry Data Security Standard (PCI DSS) pertains to organizations that handle credit card transactions and is an industry-regulatory standard, not a governmental one. The Sarbanes-Oxley Act (SOX) was passed in 2002 and pertains to financial records and reporting, as well as transparency requirements for shareholders and other stakeholders. The Health Insurance Portability and Accountability Act (HIPAA) was passed in 1996 and pertains to data privacy and security for medical records. FISMA refers to the Federal Information Security Management Act of 2002 and pertains to the protection of all US federal government IT systems, with the exception of national security systems.

Community vote distribution

MaciekMT 1 month ago

Selected Answer: D

PCI DSS (Payment Card Industry Data Security Standard) is not a U.S. federal government regulation. It's a set of security standards developed by major credit card companies to protect cardholder data, and while it is widely enforced by the payment card industry, it is not a law enacted by the U.S. government.

Here's why the others are not correct:

HIPAA (Health Insurance Portability and Accountability Act): This is a federal law that establishes standards for protecting sensitive patient health information.
SOX (Sarbanes-Oxley Act): This is a federal law that sets requirements for financial reporting and corporate governance, particularly in response to corporate fraud.
FISMA (Federal Information Security Management Act): This is a federal law that requires federal agencies and contractors to implement information security programs.

upvoted 1 times

DA95 3 months, 2 weeks ago

SOX (Sarbanes-Oxley Act) is not a regulatory system from the United States federal government. SOX is a corporate governance law that was enacted in 2002 in response to a number of accounting scandals at major companies. The law sets out certain requirements for public companies, including the need for independent audits and improved financial reporting. SOX does not pertain to privacy or data regulation specifically.

The other three options (HIPAA, FISMA, and PCI DSS) are all regulatory systems from the United States federal government. HIPAA (Health Insurance Portability and Accountability Act) is a law that protects the privacy of individuals' health information. FISMA (Federal Information Security Management Act) is a law that establishes a framework for securing federal government information systems. PCI DSS (Payment Card Industry Data Security Standard) is a set of security standards for organizations that handle credit card information.

upvoted 1 times

akg001 10 months, 2 weeks ago

Selected Answer: D

D. PCI DSS

upvoted 1 times

funktribe 1 year, 9 months ago

The Federal Information Security Management Act of 2002 is a United States federal law enacted in 2002 as Title III of the E-Government Act of 2002. The act recognized the importance of information security to the economic and national security interests of the United States. PCI is rules created by the credit card companies hence global.

upvoted 1 times

NobleGiantz 2 years, 1 month ago

D is correct. Explanation is on point.

upvoted 2 times

Sa007788 2 years, 2 months ago

Explain is not true :
https://en.wikipedia.org/wiki/Federal_Information_Security_Management_Act_of_2002
according to this link :" According to FISMA, the head of each agency shall develop and maintain an inventory of major information systems (including major national security systems)"

upvoted 1 times

Zeezee2 1 year, 4 months ago

The act recognized the importance of information security to the economic and national security interests of the United States. The act requires each federal agency to develop, document, and implement an agency-wide program to provide information security for the information and information systems that support the operations and assets of the agency, including those provided or managed by another agency, contractor, or other source.

upvoted 1 times

Question #233 *Topic 1*

The president of your company has tasked you with implementing cloud services as the most efficient way of obtaining a robust disaster recovery configuration for your production services.
Which of the cloud deployment models would you MOST likely be exploring?

A. Hybrid

B. Private

C. Community

D. Public

Suggested Answer: *A*
A hybrid cloud model spans two more different hosting configurations or cloud providers. This would enable an organization to continue using its current hosting configuration, while adding additional cloud services to enable disaster recovery capabilities. The other cloud deployment models--public, private, and community-- would not be applicable for seeking a disaster recovery configuration where cloud services are to be leveraged for that purpose rather than production service hosting.

Community vote distribution
A (63%) D (38%)

pete1981 Highly Voted 1 year, 8 months ago
Is the company currently operating on prem or in cloud? If it's currently in cloud, which model? The answer incorrectly assumes that we know the information.
upvoted 5 times

MaciekMT Most Recent 1 month ago
Selected Answer: D
Public cloud services are often the go-to choice for robust disaster recovery (DR) configurations. They offer:

Scalability & Flexibility: Resources can be provisioned on-demand to match your DR needs.
Cost Efficiency: You typically pay only for what you use, avoiding significant upfront capital expenditure.
Geographical Redundancy: Public cloud providers have data centers in multiple locations, enhancing resiliency and failover capabilities.
In contrast, hybrid, private, or community models either involve more complex management or don't inherently offer the broad, on-demand resource availability that public clouds do for DR purposes.
upvoted 1 times

Pika26 5 months, 1 week ago
Selected Answer: A
A. Hybrid
upvoted 1 times

Lenell 8 months, 3 weeks ago
Selected Answer: A
The question implies that you DO NOT have a clod solution. DR context implies that your current location is no longer usable for ops so the DR location must be in a different location than your current location, Cloud would allow you to recover faster and with less cost. Hybrid is best cloud solution for this.
upvoted 1 times

hanyahmed 9 months, 1 week ago
Selected Answer: A
A is the correct answer, it is Hybrid DR model
upvoted 1 times

AJ2021 11 months, 1 week ago
Selected Answer: A
Another tricky question, where you could argue for both A & D, but in my opinion, the way it is written, I would go for A
upvoted 1 times

serget12 11 months, 2 weeks ago

Disaster takes place, it would make sense that the area the region that is impact would impact your On-prem. Public cloud would allow you to set-up your DR location in a totally different region giving you the best option. Not to mention Cost savings.

upvoted 1 times

nigthwish 1 year ago

A: is the correct answer. No other one would even closely come to benefit DR as well as hybrid would.

upvoted 2 times

quagga 1 year, 1 month ago

Selected Answer: A

A: Hybrid

upvoted 1 times

zaqwsx 1 year, 5 months ago

Selected Answer: D

I will go public is most efficent way IMO

upvoted 3 times

AWSPro24 1 year, 9 months ago

pg 25 of the study guide ""Cloud bursting" and disaster recovery can be enhanced by hybrid cloud deployments; "cloud bursting" allows for public cloud resources to be utilized when a
private cloud workload has reached maximum capacity."

Also, if you are "implementing" cloud services I am going to say Hybrid as you can't just go from on-prem to cloud unless you're only doing IaaS / VMs.

upvoted 2 times

cmarcos97 2 years ago

Public(AWS or Azure) makes more sense to me. Hybrid if needing more shared security responsibility.

upvoted 2 times

phanil1 2 years, 5 months ago

I think its public, question doesnt talk about PII or SPI/PHI, it doesn't say you have a mix of clouds. robust DR is only three in public.

upvoted 1 times

HCL 2 years, 10 months ago

The question doesn't mention the production environment is a private cloud environment. According to NIST's definition, it cannot be a Hybrid Cloud if it does not involves two of more distinct cloud infrastructures (private, community, or public).

upvoted 2 times

Ahbey_911 2 years, 8 months ago

I don't think the question need to specify that the production environment need to be private cloud for the answer to be hybrid. It's hybrid because sensitive and PII data are usually backed up on private cloud, while other data can be backed up on public cloud. So, private+public=hybrid in this case.

upvoted 3 times

Zeezee2 1 year, 10 months ago

In this specific case, nothing says that production would move from onprem to cloud so we should assume only bcp/dr elements may be put to cloud meaning either simply backup and archiving and perhaps also the replication of on-prem resources to their cloud counterparts whenever it is needed. In any case, this is a hybrid setup which may require some degree of interoperability.
Public cloud would be more suitable if they specify least cost and least effort without specific security needs, but it's not mentioned.

upvoted 4 times

Question #234 *Topic 1*

If you are running an application that has strict legal requirements that the data cannot reside on systems that contain other applications or systems, which aspect of cloud computing would be prohibitive in this case?

A. Multitenancy

B. Broad network access

C. Portability

D. Elasticity

Suggested Answer: *A*

Multitenancy is the aspect of cloud computing that involves having multiple customers and applications running within the same system and sharing the same resources. Although considerable mechanisms are in place to ensure isolation and separation, the data and applications are ultimately using shared resources.

Broad network access refers to the ability to access cloud services from any location or client. Portability refers to the ability to easily move cloud services between different cloud providers, whereas elasticity refers to the capabilities of a cloud environment to add or remove services, as needed, to meet current demand.

Community vote distribution

A (100%)

akg001 4 months, 1 week ago

Selected Answer: A

A. Multitenancy

upvoted 1 times

Ajaygilly 1 year, 1 month ago

Answer is correct

upvoted 1 times

Question #235 *Topic 1*

The REST API is a widely used standard for communications of web-based services between clients and the servers hosting them.
Which protocol does the REST API depend on?

A. HTTP

B. SSH

C. SAML

D. XML

Suggested Answer: *A*
Representational State Transfer (REST) is a software architectural scheme that applies the components, connectors, and data conduits for many web applications used on the Internet. It uses and relies on the HTTP protocol and supports a variety of data formats. Extensible Markup Language (XML) and Security Assertion
Markup Language (SAML) are both standards for exchanging encoded data between two parties, with XML being for more general use and SAML focused on authentication and authorization data. Secure Shell client (SSH) is a secure method for allowing remote login to systems over a network.

Community vote distribution
A (100%)

cloudenthusiast 7 months ago

Selected Answer: A

Rest API relies on HTTP, based on study material.

upvoted 1 times

Question #236 *Topic 1*

Which of the following actions will NOT make data part of the create phase of the cloud data lifecycle?

A. Modify data

B. Modify metadata

C. New data

D. Import data

Suggested Answer: *B*

Modifying the metadata does not change the actual data. Although this initial phase is called "create," it can also refer to modification. In essence, any time data is considered "new," it is in the create phase. This can come from data that is newly created, data that is imported into a system and is new to that system, or data that is already present and is modified into a new form or value.

Community vote distribution

B (100%)

akg001 4 months, 1 week ago

Selected Answer: B

B. Modify metadata

repeated.

upvoted 2 times

zhengdeshuo 1 year, 6 months ago

Answer is B. It's repeated question

upvoted 4 times

Question #237 *Topic 1*

Most APIs will support a variety of different data formats or structures.
However, the SOAP API will only support which one of the following data formats?

A. XML

B. XSLT

C. JSON

D. SAML

Suggested Answer: *A*
The Simple Object Access Protocol (SOAP) protocol only supports the Extensible Markup Language (XML) data format.
Although the other options are all data formats or data structures, they are not supported by SOAP.

Community vote distribution
A (100%)

akg001 4 months, 1 week ago

Selected Answer: A

A. XML

upvoted 2 times

Question #238 *Topic 1*

Which cloud storage type is typically used to house virtual machine images that are used throughout the environment?

A. Structured

B. Unstructured

C. Volume

D. Object

Suggested Answer: *D*

Object storage is typically used to house virtual machine images because it is independent from other systems and is focused solely on storage. It is also the most appropriate for handling large individual files. Volume storage, because it is allocated to a specific host, would not be appropriate for the storing of virtual images.

Structured and unstructured are storage types specific to PaaS and would not be used for storing items used throughout a cloud environment.

Community vote distribution

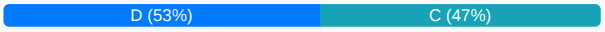

dnd1000 Highly Voted 2 years, 10 months ago

Selected Answer: D

The question is asking for storing VM images. These images are stored in Object Storage. You use volume storage to attach to the VMs as a form of hard drive, not to store the VM image.

upvoted 9 times

akg001 Highly Voted 2 years, 10 months ago

Selected Answer: C

C. Volume

upvoted 7 times

MaciekMT Most Recent 1 month ago

Selected Answer: D

Virtual machine images are typically stored in object storage. Object storage is designed for handling large amounts of unstructured data—like VM images—and offers scalability, durability, and a flat namespace that makes managing these images straightforward. In many cloud environments, VM images (such as Amazon Machine Images in AWS) are stored in an object storage system (e.g., Amazon S3) until they are deployed. This contrasts with volume (block) storage, which is used for the running instances rather than for storing the image itself.

upvoted 1 times

sweetykaur 5 months, 2 weeks ago

The typical cloud storage type used to house virtual machine images is C. Volume storage. Volume storage provides block-level storage that is well-suited for storing and managing virtual machine images.

upvoted 1 times

el3ctronick 7 months ago

Selected Answer: D

object is cheaper and better option for data that doesn't alter occasionally. i will go for object.

upvoted 1 times

JDav1 1 year, 2 months ago

D=Object. It is about storing image not running the VM

upvoted 1 times

joeee7 1 year, 8 months ago

volume

upvoted 2 times

Pika26 1 year, 11 months ago

Selected Answer: C

C. Volume

upvoted 1 times

[Removed] 3 years, 2 months ago

Volume storage is the correct Answer, S3 is Object, cant store VM files there. Object is flat structured storage.

upvoted 2 times

Question #239 *Topic 1*

With an API, various features and optimizations are highly desirable to scalability, reliability, and security.
What does the REST API support that the SOAP API does NOT support?

A. Acceleration

B. Caching

C. Redundancy

D. Encryption

Suggested Answer: *B*
The Simple Object Access Protocol (SOAP) does not support caching, whereas the Representational State Transfer (REST) API does. The other options are all capabilities that are either not supported by SOAP or not supported by any API and must be provided by external features.

Community vote distribution
B (100%)

akg001 4 months, 1 week ago

Selected Answer: B

B. Caching

upvoted 2 times

Question #240 *Topic 1*

Although much of the attention given to data security is focused on keeping data private and only accessible by authorized individuals, of equal importance is the trustworthiness of the data.
Which concept encapsulates this?

A. Validity

B. Integrity

C. Accessibility

D. Confidentiality

Suggested Answer: *B*
Integrity refers to the trustworthiness of data and whether its format and values are true and have not been corrupted or otherwise altered through unauthorized means. Confidentiality refers to keeping data from being access or viewed by unauthorized parties. Accessibility means that data is available and ready when needed by a user or service. Validity can mean a variety of things that are somewhat similar to integrity, but it's not the most appropriate answer in this case.

Community vote distribution
B (100%)

JDav1 8 months ago
Trustworthiness of the data is the key. The answer is B=Integrity.
upvoted 1 times

Voldamort 2 years, 8 months ago
Selected Answer: B
It has to be integrity since confidentiality is already covered. Consider 'of equal importance' when thinking about this question.
upvoted 3 times

Banzaaai 3 years ago
its D , confidentiality
upvoted 1 times

xaccan 2 years, 11 months ago
The question has two parts, the first is describing the confidentiality (allow only authorized people to access data) and the second part is asking about trustworthiness of the data which is obviously the integrity
upvoted 9 times

Question #241 *Topic 1*

Three central concepts define what type of data and information an organization is responsible for pertaining to eDiscovery. Which of the following are the three components that comprise required disclosure?

A. Possession, ownership, control

B. Ownership, use, creation

C. Control, custody, use

D. Possession, custody, control

Suggested Answer: *D*

Data that falls under the purview of an eDiscovery request is that which is in the possession, custody, or control of the organization. Although this is an easy concept in a traditional data center, it can be difficult to distinguish who actually possesses and controls the data in a cloud environment due to multitenancy and resource pooling. Although these options provide similar-sounding terms, they are ultimately incorrect.

Community vote distribution

MaciekMT 1 month ago

Selected Answer: D

In the context of eDiscovery, an organization is typically required to disclose data and information that it has in its "possession, custody, or control." This legal standard determines what information must be preserved and produced during litigation or investigations, making it the key framework for identifying relevant data.

upvoted 1 times

Aa7411111 7 months, 3 weeks ago

Selected Answer: D

Should be D.

upvoted 1 times

akg001 1 year, 10 months ago

Selected Answer: D

D. Possession, custody, control

upvoted 1 times

Question #242 *Topic 1*

Which of the following threat types involves the sending of commands or arbitrary data through input fields in an application in an attempt to get that code executed as part of normal processing?

A. Cross-site scripting

B. Missing function-level access control

C. Injection

D. Cross-site forgery

Suggested Answer: *C*

An injection attack is where a malicious actor will send commands or other arbitrary data through input and data fields with the intent of having the application or system execute the code as part of its normal processing and queries. This can trick an application into exposing data that is not intended or authorized to be exposed, or it could potentially allow an attacker to gain insight into configurations or security controls. Missing function-level access control exists where an application only checks for authorization during the initial login process and does not further validate with each function call. Cross-site request forgery occurs when an attack forces an authenticated user to send forged requests to an application running under their own access and credentials. Cross-site scripting occurs when an attacker is able to send untrusted data to a user's browser without going through validation processes.

Community vote distribution

MaciekMT 1 month ago

Selected Answer: C

Injection attacks involve sending malicious commands or arbitrary data through input fields, tricking the application into executing unintended commands as part of its normal processing. This contrasts with Cross-site scripting, which targets client-side code execution, and cross-site request forgery, which leverages authenticated sessions to force unwanted actions. Missing function-level access control is about inadequate permission checks rather than injecting code.

upvoted 1 times

zxccvbnm 7 months, 2 weeks ago

Selected Answer: C

C. Injection

upvoted 2 times

Pravinkarthik 8 months, 2 weeks ago

Selected Answer: C

C. Injection

upvoted 1 times

akg001 10 months, 2 weeks ago

Selected Answer: D

D. Possession, custody, control

upvoted 2 times

Question #243 *Topic 1*

With a cloud service category where the cloud customer is responsible for deploying all services, systems, and components needed for their applications, which of the following storage types are MOST likely to be available to them?

A. Structured and hierarchical

B. Volume and object

C. Volume and database

D. Structured and unstructured

Suggested Answer: *B*

The question is describing the Infrastructure as a Service (IaaS) cloud offering, and as such, the volume and object storage types will be available to the customer.

Structured and unstructured are storage types associated with PaaS, and although the other answers present similar-sounding storage types, they are a mix of real and fake names.

Community vote distribution

B (100%)

akg001 4 months, 1 week ago

Selected Answer: B

B. Volume and object

upvoted 2 times

Question #244 *Topic 1*

Which of the following roles would be responsible for managing memberships in federations and the use and integration of federated services?

A. Inter-cloud provider

B. Cloud service business manager

C. Cloud service administrator

D. Cloud service integrator

Suggested Answer: *A*

The inter-cloud provider is responsible for peering with other cloud services and providers, as well as overseeing and managing federations and federated services. A cloud service administrator is responsible for testing, monitoring, and securing cloud services, as well as providing usage reporting and dealing with service problems. The cloud service integrator is responsible for connecting existing systems and services with a cloud. The cloud service business manager is responsible for overseeing the billing, auditing, and purchasing of cloud services.

Community vote distribution

MaciekMT 3 weeks, 5 days ago

Selected Answer: A

In cloud computing, an Inter-cloud provider is responsible for managing interactions between different cloud environments, including:

Managing memberships in federations
Overseeing the use and integration of federated services
This role ensures seamless interoperability and integration across multiple cloud platforms.

Why Not the Others?

B. Cloud service business manager: Focuses on business aspects like contracts and pricing, not technical integration of federated services.

C. Cloud service administrator: Manages day-to-day operations within a single cloud environment, not cross-cloud federations.

D. Cloud service integrator: Integrates various services within a cloud but doesn't specifically manage federations between different cloud providers.

upvoted 1 times

Pillartech 6 months, 2 weeks ago

Selected Answer: A

Inter-cloud provider

upvoted 1 times

Pika26 1 year, 11 months ago

Selected Answer: D

D. Cloud service integrator

upvoted 2 times

hanyahmed 2 years, 3 months ago

Selected Answer: A

This role related to Cloud Service Provider "CSP"
A is the correct answer Inter-cloud manage the federation

upvoted 1 times

akg001 2 years, 10 months ago

Selected Answer: A

A. Inter-cloud provider

upvoted 1 times

abhiw13 3 years, 1 month ago

(ISC)2 Official Student Guide has this under ISO/IEC 17789 CSP, CSC, CSN Roles

upvoted 1 times

[Removed] 3 years, 2 months ago

Federation is not necessarily inter cloud. D could be the correct answer. but question itself looks incorrect.

upvoted 3 times

AWSPro24 3 years, 3 months ago

Can anyone provide a reference for this from anywhere?

upvoted 1 times

certifiedgeek 2 years, 10 months ago

Check this URL https://cloudgal42.com/cloud-computing-activities/ which summarizes all cloud roles/sub-roles based on the ISO/IEC 17789 document.

upvoted 1 times

abhiw13 3 years, 1 month ago

Answer is A: Inter-cloud Provider as per ISO/IEC 17789 CSP, CSC, CSN Roles

upvoted 4 times

Question #245 *Topic 1*

Which data state would be most likely to use TLS as a protection mechanism?

A. Data in use

B. Data at rest

C. Archived

D. Data in transit

Suggested Answer: *D*

TLS would be used with data in transit, when packets are exchanged between clients or services and sent across a network. During the data-in-use state, the data is already protected via a technology such as TLS as it is exchanged over the network and then relies on other technologies such as digital signatures for protection while being used. The data-at-rest state primarily uses encryption for stored file objects. Archived data would be the same as data at rest.

Community vote distribution

akg001 4 months, 1 week ago

Selected Answer: D

D. Data in transit

upvoted 1 times

Question #246 *Topic 1*

You are working for a cloud service provider and receive an eDiscovery order pertaining to one of your customers.
Which of the following would be the most appropriate action to take first?

A. Take a shapshot of the virtual machines

B. Escrow the encryption keys

C. Copy the data

D. Notify the customer

Suggested Answer: *D*
When a cloud service provider receives an eDiscovery order pertaining to one of their customers, the first action they must take is to notify the customer. This allows the customer to be aware of what was received, as well as to conduct a review to determine if any challenges are necessary or warranted. Taking snapshots of virtual machines, copying data, and escrowing encryption keys are all processes involved in the actual collection of data and should not be performed until the customer has been notified of the request.

Community vote distribution
D (100%)

akg001 4 months, 1 week ago

Selected Answer: D

D. Notify the customer
upvoted 1 times

Question #247 *Topic 1*

If a cloud computing customer wishes to guarantee that a minimum level of resources will always be available, which of the following set of services would compromise the reservation?

A. Memory and networking

B. CPU and software

C. CPU and storage

D. CPU and memory

Suggested Answer: *D*

A reservation guarantees to a cloud customer that they will have access to a minimal level of resources to run their systems, which will help mitigate against DoS attacks or systems that consume high levels of resources. A reservation pertains to memory and CPU resources. Under the concept of a reservation, memory and
CPU are the guaranteed resources, but storage and networking are not included even though they are core components of cloud computing. Software would be out of scope for a guarantee and doesn't really pertain to the concept.

Community vote distribution

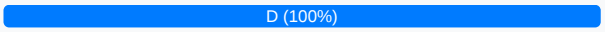

gjjw Highly Voted 3 years, 4 months ago

Should be ‘comprise’ instead of ‘compromise’ in the question

upvoted 15 times

Pika26 Most Recent 5 months, 1 week ago

Selected Answer: D

D. CPU and memory

upvoted 1 times

Smiles1963 1 year, 2 months ago

Please adjust the question. Change Compromise to comprise

upvoted 1 times

akg001 1 year, 4 months ago

Selected Answer: D

D. CPU and memory

upvoted 1 times

Zeezee2 1 year, 10 months ago

Network, software and storage are not typically the areas where a bottleneck of resources would occur (all are widely available and cheap in comparison to CPU/Memory which may be more limited in availability).

upvoted 1 times

Sa007788 2 years, 8 months ago

How cloud customer can work without networking.For me it's A

upvoted 1 times

Sa007788 2 years, 8 months ago

Sorry, i made confusion it's D correct

upvoted 2 times