Question #1 *Topic 1*

A company has decided to scale its e-commerce application from its corporate datacenter to a commercial cloud provider to meet an anticipated increase in demand during an upcoming holiday.

The majority of the application load takes place on the application server under normal conditions. For this reason, the company decides to deploy additional application servers into a commercial cloud provider using the on-premises orchestration engine that installs and configures common software and network configurations.

The remote computing environment is connected to the on-premises datacenter via a site-to-site IPSec tunnel. The external DNS provider has been configured to use weighted round-robin routing to load balance connections from the Internet.

During testing, the company discovers that only 20% of connections completed successfully.

Instructions -

Review the network architecture and supporting documents and fulfill these requirements:

Part 1:

✑ Analyze the configuration of the following components: DNS, Firewall 1, Firewall 2, Router 1, Router 2, VPN and Orchestrator Server.

✑ Identify the problematic device(s).

Part 2:

✑ Identify the correct options to provide adequate configuration for hybrid cloud architecture.

If any time you would like to bring back the initial state of the simulation, please click the Reset All button.

Part 1 -

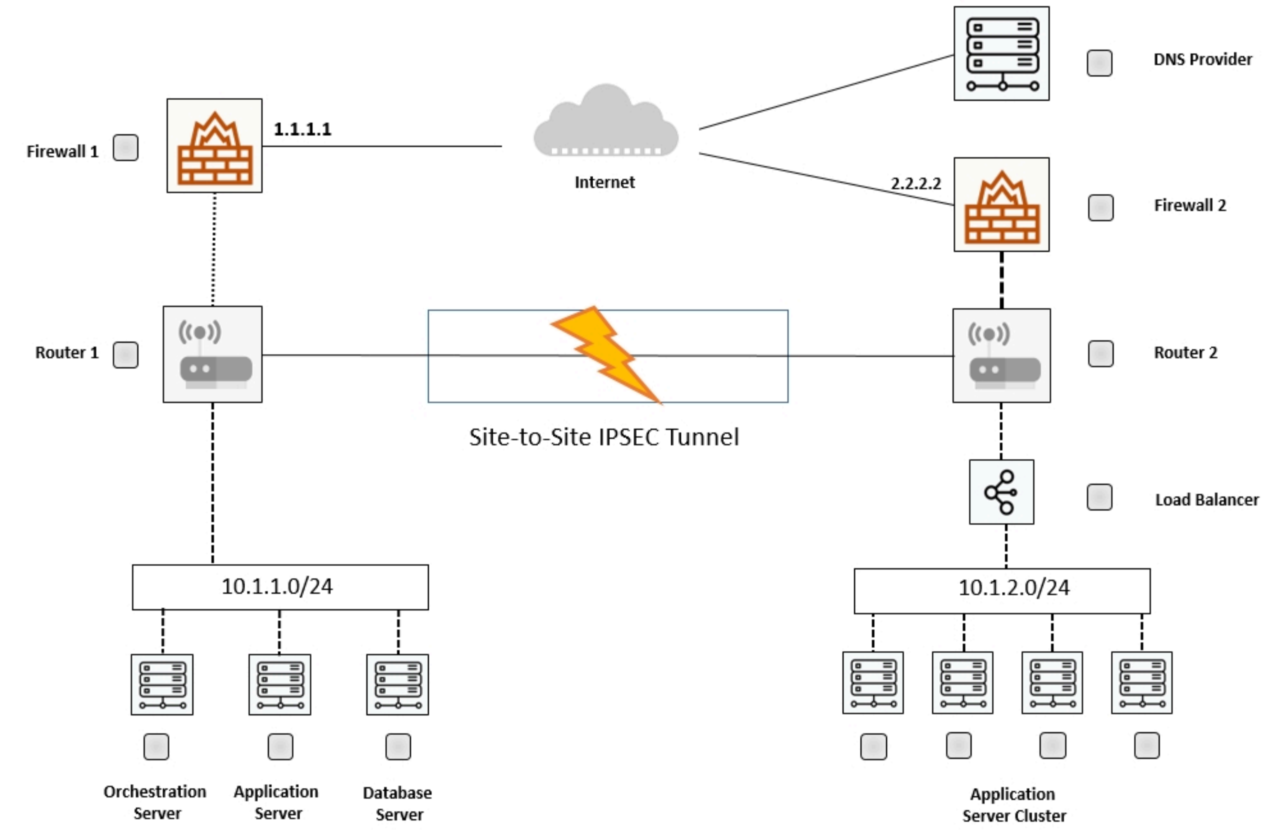

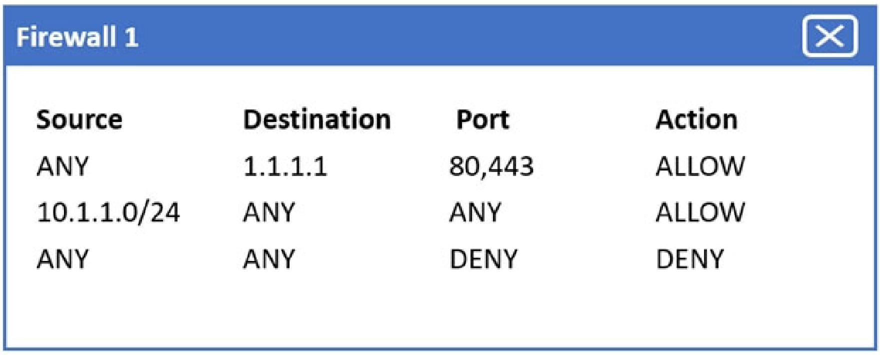
Firewall 1

Source	Destination	Port	Action
ANY	1.1.1.1	80,443	ALLOW
10.1.1.0/24	ANY	ANY	ALLOW
ANY	ANY	DENY	DENY

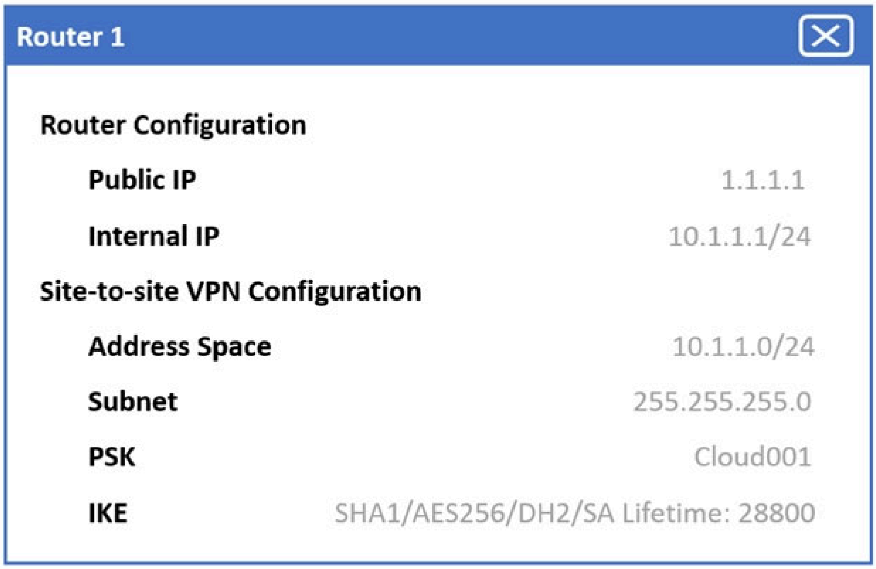

Router 1

Router Configuration

Public IP	1.1.1.1
Internal IP	10.1.1.1/24

Site-to-site VPN Configuration

Address Space	10.1.1.0/24
Subnet	255.255.255.0
PSK	Cloud001
IKE	SHA1/AES256/DH2/SA Lifetime: 28800

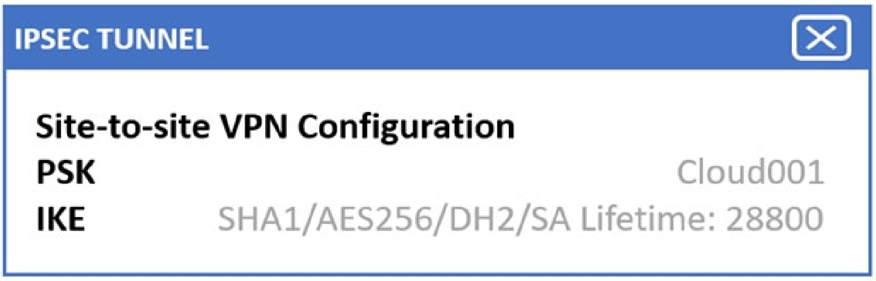

IPSEC TUNNEL

Site-to-site VPN Configuration

PSK	Cloud001
IKE	SHA1/AES256/DH2/SA Lifetime: 28800

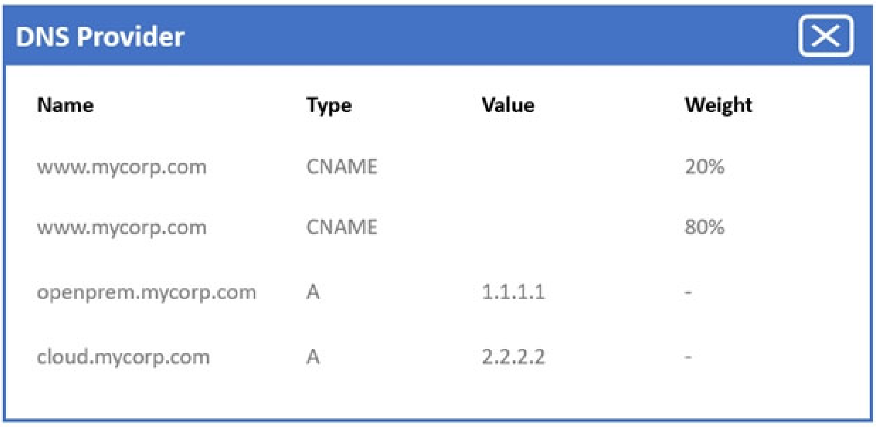

DNS Provider

Name	Type	Value	Weight
www.mycorp.com	CNAME		20%
www.mycorp.com	CNAME		80%
openprem.mycorp.com	A	1.1.1.1	-
cloud.mycorp.com	A	2.2.2.2	-

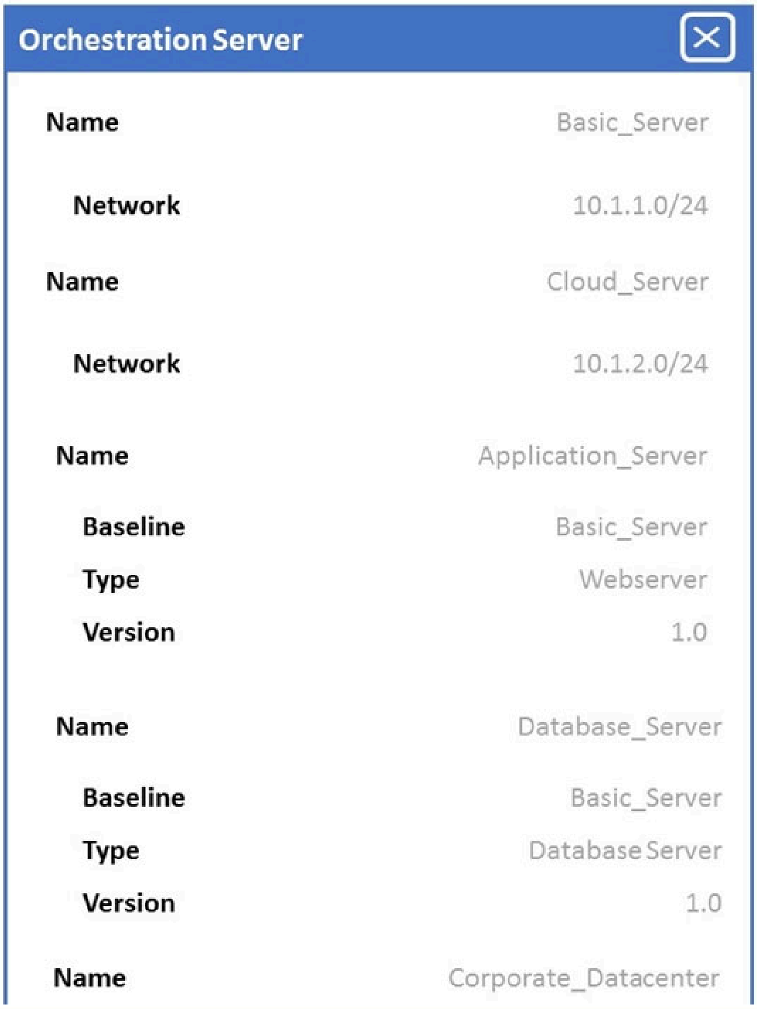

Orchestration Server

Name	Basic_Server
Network	10.1.1.0/24
Name	Cloud_Server
Network	10.1.2.0/24
Name	Application_Server
Baseline	Basic_Server
Type	Webserver
Version	1.0
Name	Database_Server
Baseline	Basic_Server
Type	Database Server
Version	1.0
Name	Corporate_Datacenter

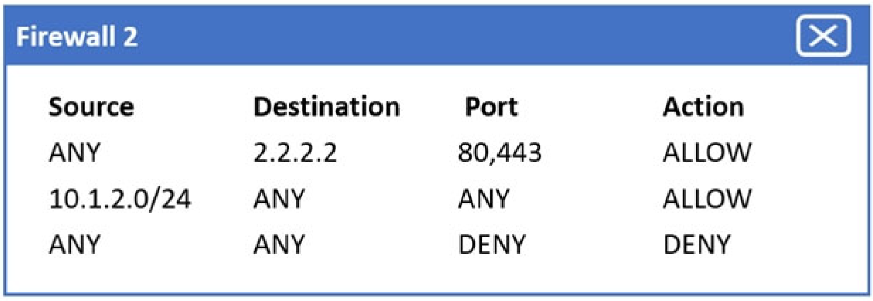

Firewall 2

Source	Destination	Port	Action
ANY	2.2.2.2	80,443	ALLOW
10.1.2.0/24	ANY	ANY	ALLOW
ANY	ANY	DENY	DENY

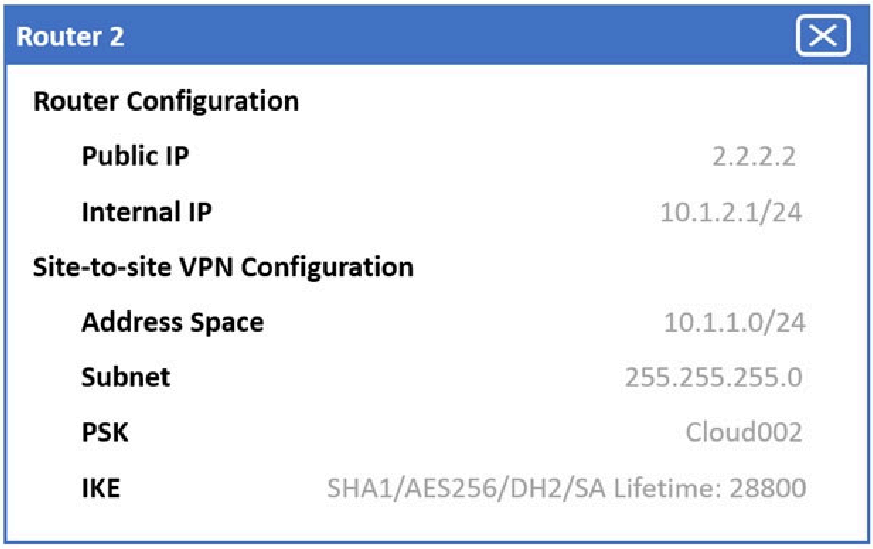

Router 2

Router Configuration

Public IP	2.2.2.2
Internal IP	10.1.2.1/24

Site-to-site VPN Configuration

Address Space	10.1.1.0/24
Subnet	255.255.255.0
PSK	Cloud002
IKE	SHA1/AES256/DH2/SA Lifetime: 28800

Part 2 -

Only select a maximum of TWO options from the multiple choice question. (Choose two.)

A. Update the PSK (Pre-shared key) in Router 2.

B. Update the A record on the DNS from 2.2.2.2 to 1.1.1.1.

C. Promote deny All to allow All in Firewall 1 and Firewall 2.

D. Change the Address Space on Router 2.

E. Change internal IP Address of Router 1.

F. Reverse the Weight property in the two CNAME records on the DNS.

G. Add the Application Server at on-premises to the Load Balancer.

Suggested Answer: *AD*

Community vote distribution

AD (63%) A (25%) 13%

Sweety_Certified7 3 months ago

Selected Answer: AD

Correct Options:

A. Update the PSK (Pre-shared key) in Router 2.

D. Change the Address Space on Router 2.

Explanation:

A. Updating the PSK in Router 2 ensures that it matches the configuration of the VPN endpoint on the commercial cloud provider's side, resolving any authentication issues.

D. Changing the Address Space on Router 2 might be necessary if there are IP address conflicts between the on-premises datacenter and the commercial cloud provider's network, ensuring proper routing of traffic through the VPN tunnel.

upvoted 2 times

Sweety_Certified7 3 months ago

Correct Options:

A. Update the PSK (Pre-shared key) in Router 2.

D. Change the Address Space on Router 2.

Explanation:

A. Updating the PSK in Router 2 ensures that it matches the configuration of the VPN endpoint on the commercial cloud provider's side, resolving any authentication issues.

D. Changing the Address Space on Router 2 might be necessary if there are IP address conflicts between the on-premises datacenter and the commercial cloud provider's network, ensuring proper routing of traffic through the VPN tunnel.

upvoted 1 times

Deeeeez_nuts 3 months, 1 week ago

there are 5 application servers. 1 is on premises, the other 4 are in a cluster. 20% of the connections are successful which probably means the only server accessible is the one on premises. The cloud architecture is not working, how do we make it work? We need access to those other 4 application servers so there must be a misconfiguration in the address space in router 2 and the PSK is wrong imo.

upvoted 3 times

FrancisDrake 3 months, 2 weeks ago

I took the test today. Passed it with the help of examtopics and some udemy. I got a 771. Not great but I didn't really want this cert anyway. Just got it to renew other certs. I still am not sure about the answer to this one. It was on the test. As has been hinted at by others it seems that you need to select one defective device (yes, it's router 2) but then the network needs to be configured for a hybrid architecture. So it's either the DNS or add application server. One mans opinion. Good luck on the test.

upvoted 2 times

kuzummjakk 4 months ago

Selected Answer: A

ADG. I believe it wrongly says select 2.

A: Ofc it's true

D: The ipsec tunnel is generally configured to be aware of both subnets

G: "Part 2" CLEARLY logically differentiates "fixing the connection" from "more appropriately configure for the cloud". Notice the more, insinuating that the current configuration is not "bugged" but "could be better". D is related to part 1 and would be classified as "necessary", not "preferable". D is also NOT classified as a "cloud-appropriate" configuration. It's an appropriate configuration, just not specifically a cloud one.

upvoted 1 times

⊟ **kuzummjakk** 4 months ago

Hold on, it says for "hybrid cloud architecture". G isn't necessarily hybrid cloud specific either. Better hybrid cloud architecture COULD just be fixing the "hybrid" part which is D?

upvoted 2 times

⊟ **FrankyD92** 4 months, 2 weeks ago

umm....yes. I can say I did not study enough on this portion and I REALLY wish I did

upvoted 1 times

⊟ **FrancisDrake** 5 months, 1 week ago

Selected Answer: A

I think the address space for the VPN configuration on router 2 is correct. How else are router 1 and 2 going to communicate? I think people are confusing router 2's internal address with the VPN configuration. The PSK is incorrect. What the second selection is I'm not sure. I think this question does not have all of the needed information.

upvoted 1 times

⊟ **SecPlus2022** 1 year ago

Selected Answer: AD

The PSK and address space on router 2 are clearly wrong.

upvoted 1 times

⊟ **ROCompTIA** 1 year, 1 month ago

Selected Answer: AD

You can see pretty clear ip class on the router 2

upvoted 1 times

⊟ **Trebor28** 1 year, 4 months ago

Selected Answer: AD

AD is the answer.

upvoted 1 times

⊟ **dryshell** 1 year, 4 months ago

Selected Answer: AG

i agree w sal

upvoted 1 times

⊟ **Sal** 1 year, 8 months ago

router 2

A. and H.

upvoted 1 times

⊟ **ryanzou** 1 year, 9 months ago

The problematic device should be route2.

upvoted 1 times

Question #2 *Topic 1*

An organization suffered a critical failure of its primary datacenter and made the decision to switch to the DR site. After one week of using the DR site, the primary datacenter is now ready to resume operations.
Which of the following is the MOST efficient way to bring the block storage in the primary datacenter up to date with the DR site?

A. Set up replication.

B. Copy the data across both sites.

C. Restore incremental backups.

D. Restore full backups.

Suggested Answer: *A*

Community vote distribution

A (100%)

TheFivePips 1 month, 2 weeks ago

Selected Answer: A

Replication involves continuously copying data changes from the primary datacenter to the DR site in near real-time. Since the DR site has been operational for a week, setting up replication would allow the primary datacenter's block storage to quickly catch up with the changes made at the DR site. This approach ensures that the primary datacenter becomes up-to-date with minimal data loss and downtime.

Option B (Copy the data across both sites) would likely be inefficient and time-consuming, especially if the amount of data is large, as it involves manually copying data across both sites, which may not capture changes made during the past week.

Option C (Restore incremental backups) and option D (Restore full backups) are not as efficient in this scenario because they involve restoring data from backups, which may not capture all changes made at the DR site since the failure occurred. Additionally, restoring backups could result in longer downtime and potential data loss compared to replication, which provides near real-time data synchronization.

upvoted 2 times

kuzummjakk 4 months ago

Selected Answer: A

C and D are close, but replication is more closely related to "multi-site" for "coldsite/hotsite" operations. D would take too long, and for C, incrementals just speed up the time for backing up, but you still need to restore the full backup before restoring all the incrementals, so C would also take too long for this case.

upvoted 1 times

Question #3 *Topic 1*

A cloud administrator is building a new VM for machine-learning training. The developer requesting the VM has stated that the machine will need a full GPU dedicated to it. Which of the following configuration options would BEST meet this requirement?

A. Virtual GPU

B. External GPU

C. Passthrough GPU

D. Shared GPU

Suggested Answer: *C*

Community vote distribution

C (100%)

Boats Highly Voted 11 months, 1 week ago

Selected Answer: C

GPU pass-through is used when one or more GPUs are dedicated to a single VM. Pass-through allows the VM to address the GPU directly without having to go through the hypervisor stack. This significantly improves performance for the application. Pass-through is used for very GPU-intensive workloads such as deep learning, artificial intelligence, and data analytics.

Vanderburg, Eric A.. CompTIA Cloud+ Certification All-in-One Exam Guide (Exam CV0-003) (p. 127). McGraw Hill LLC. Kindle Edition.

upvoted 11 times

Alizadeh Most Recent 4 months ago

Selected Answer: C

C. Passthrough GPU would be the best configuration option to meet the requirement of a full GPU dedicated to the machine-learning training VM.

Passthrough GPU, also known as GPU pass-through, allows a physical GPU to be directly assigned to a virtual machine, providing the VM with full access to the GPU's resources. This means that the VM will have a dedicated physical GPU, providing the performance required for machine-learning training.

upvoted 3 times

HeliosABC 10 months, 2 weeks ago

C is correct ans

upvoted 1 times

Question #4 *Topic 1*

Which of the following service models would be used for a database in the cloud?

A. PaaS

B. IaaS

C. CaaS

D. SaaS

Suggested Answer: *A*

Community vote distribution

A (100%)

Boats 3 months, 1 week ago

Selected Answer: A

I would go with PaaS.

Both AWS and Azure's primary Database solutions are PaaS. However you can run DBs in a VM making it also IaaS. I don't see where any providers solution has DBs classified as SaaS.

DBaaS: Diversity of Solutions
There's one more term we haven't addressed so far—Database as a service (DBaaS). It typically refers to databases offered as PaaS or SaaS.

All major cloud platforms now offer DBaaS solutions. Some of these are closer to SaaS; others are slightly closer to PaaS. Popular examples of DBaaS solutions include Amazon Relational Database Service (RDS), Azure SQL Database, MongoDB Atlas and Amazon DynamoDB.

upvoted 4 times

54a6b25 5 months, 3 weeks ago

A is right

upvoted 1 times

TheFivePips 7 months, 3 weeks ago

Selected Answer: A

A is the best answer but its annoying because you can use all the others to set up a database one way or another. PaaS just does all the dirty work for you.

upvoted 1 times

kuzummjakk 10 months ago

ig it's not SaaS because "in the cloud" seems to assume that it's from the perspective of "aws" or "azure" not "user who happens to be interacting with a cloud service over a web GUI"

upvoted 1 times

VVV4WIN 1 year, 1 month ago

Selected Answer: A

A - PaaS

From Microsoft: Azure SQL Database is a relational database-as-a-service (DBaaS) hosted in Azure that falls into the industry category of Platform-as-a-Service (PaaS).

upvoted 2 times

Alizadeh 1 year, 10 months ago

Selected Answer: A

A. PaaS (Platform as a Service) would be the service model used for a database in the cloud.

PaaS provides a platform for developers to build, deploy, and manage applications without the need to manage the underlying infrastructure. In the case of a database, a PaaS provider would offer a managed database service that would handle tasks such as software updates, data backups, and scaling, allowing developers to focus on building and maintaining their applications.

upvoted 4 times

HeliosABC 2 years, 4 months ago

Selected Answer: A

Paas is correct

upvoted 4 times

bx88 2 years, 6 months ago

PaaS is the most suitable service model would be used for a database in the cloud.

IaaS is of course can be used to bring up instances and then install database application on top of those instances. However, this approach does not utilize full advantages and benefits of cloud computing. The administrators have to manage and maintain all necessary infrastructure that can be offload by PaaS solution.

Example: Azure SQL Database is a fully managed platform as a service (PaaS) database engine that handles most of the database management functions such as upgrading, patching, backups, and monitoring without user involvement. SQL Server on Azure Virtual Machines (VMs) is an IaaS offer and allows you to run SQL Server inside a virtual machine in the cloud.

upvoted 4 times

SimplyDebonair 2 years, 9 months ago

The answer is "D." SaaS isn't just for on-demand software application usage. It can be utilized as a database and/or add database functionality to your application. Additionally, it comes as a "out-of-the-box" solution, it is simpler to set-up, and easier to train on. IaaS and PaaS can and do offer options to run databases, but these methods require more training and personnel to manage it. CaaS isn't a viable choice as it deals with communications.

https://rockset.com/blog/what-is-a-cloud-database-iaas-paas-saas-dbaas-explained/
https://www.ibm.com/cloud/learn/iaas-paas-saas

upvoted 4 times

maiathiago 1 year, 11 months ago

Agree!

upvoted 1 times

Question #5 *Topic 1*

A VDI administrator has received reports from the drafting department that rendering is slower than normal. Which of the following should the administrator check
FIRST to optimize the performance of the VDI infrastructure?

A. GPU

B. CPU

C. Storage

D. Memory

Suggested Answer: *A*

Community vote distribution

A (71%) B (25%) 4%

Pongsathorn Highly Voted 1 year, 3 months ago

Selected Answer: A

GPU
Depending on the workload, some processing can be offloaded to a graphics card, freeing the CPU to deal with tasks only it can manage. Not all workloads benefit from GPU offloading.

Workloads that benefit from GPU processing:
- Machine learning
- High-performance computing
- Graphics intensive
- Data analysis

Ref: The Official CompTIA Cloud+ Study Guide (Exam CV0-003)
upvoted 6 times

TheGinjaNinja Highly Voted 1 year, 11 months ago

Selected Answer: B

B. CPU should be the first thing the administrator checks to optimize the performance of the VDI infrastructure. High CPU usage can cause slow rendering and other performance issues. The administrator should monitor the CPU usage and make sure it is within normal limits, and take action if necessary (such as adding more CPU resources or optimizing the workloads) to reduce the usage and improve performance.
upvoted 5 times

54a6b25 Most Recent 5 months, 3 weeks ago

A is the right answer
upvoted 1 times

kuzummjakk 10 months ago

All of them, even storage to an extent, could be bottlenecks. "First" is rough though. Since it's not labeled as a GPU intensive task like "gaming" or "3d graphics" and it simply says "rendering", and since work software where you "make" tends to be more memory intensive than GPU intensive (any sort of "editing" software like for videos or images or even network diagrams), then D makes enough sense.
upvoted 1 times

> **kuzummjakk** 10 months ago
>
> TLDR: Work tasks are more memory intensive than GPU intensive, which can bottleneck rendering.
> upvoted 1 times

Sweety_Certified7 10 months, 1 week ago

Selected Answer: A

the VDI (Virtual Desktop Infrastructure) administrator should check to optimize the performance of the VDI infrastructure, especially for rendering tasks, is the GPU (Graphics Processing Unit) (Option A). Rendering tasks, particularly in drafting applications, heavily rely on GPU acceleration for efficient processing and display of graphics-intensive content.
upvoted 1 times

buckthesystem 1 year, 2 months ago

GPU ~ I think the key here is "drafting." My guess is that they would have a VDI solution optimized for work with CAD, thus the physical host would be equipped with a GPU. A lot of VDI solutions simply utilize the server CPU as it's a lot easier to deploy that way. So CPU isn't wrong IMO but good old CompTIA likes to give you obscure questions that in real life would require a lot more detail to properly answer.

upvoted 2 times

backdooranon 1 year, 2 months ago

Selected Answer: A

Rendering > Graphics > check GPU utilization first.

upvoted 2 times

Alizadeh 1 year, 8 months ago

Selected Answer: A

Given that the issue is related to rendering performance, the administrator should check the GPU (Graphics Processing Unit) first. Rendering tasks, particularly in drafting or design environments, rely heavily on the GPU's capabilities. By checking and optimizing the GPU performance, the administrator can improve the overall rendering speed and address the concerns of the drafting department.

upvoted 4 times

tutita 1 year, 8 months ago

Selected Answer: A

Should it be A? I'm confused, someone help, most likely you are using a lot of GPU power

upvoted 1 times

bettyboo 1 year, 8 months ago

Selected Answer: A

A. GPU

When it comes to rendering, the graphics processing unit (GPU) plays a crucial role in VDI performance. The GPU is responsible for accelerating the processing of graphics-intensive workloads such as rendering. Therefore, if the drafting department is reporting slow rendering, the VDI administrator should check the GPU usage and ensure that the appropriate GPU hardware and drivers are installed and configured correctly. Optimizing the GPU resources can significantly enhance the performance of VDI infrastructure for rendering workloads.

While the other options (CPU, storage, memory) are also important factors for VDI performance, they are typically not the primary bottleneck for rendering workloads. It's still important to ensure that they are also configured appropriately for the VDI environment, but the GPU is the first thing that the VDI administrator should check in this scenario.

upvoted 3 times

concepcionz 1 year, 9 months ago

Selected Answer: A

The drafting department likely uses applications that require graphic-intensive tasks, such as 3D rendering or CAD. Therefore, the first thing the administrator should check to optimize performance is the GPU ????

upvoted 3 times

Grayson2023 1 year, 10 months ago

According to AutoCAD, when rendering RAM is very important. Also noted, typically it uses only one CORE.
https://knowledge.autodesk.com/support/autocad/learn-explore/caas/sfdcarticles/sfdcarticles/Which-hardware-components-are-significant-for-the-use-of-AutoCAD.html

upvoted 1 times

Trebor28 1 year, 10 months ago

Selected Answer: D

D is correct.

upvoted 1 times

AustinKelleyNet 1 year, 11 months ago

Selected Answer: B

I would say B

upvoted 2 times

> **Sweety_Certified7** 9 months ago
>
> While the CPU (Central Processing Unit) is indeed a crucial component in any computing environment, including VDI infrastructures, it may not directly impact rendering performance as significantly as the GPU (Graphics Processing Unit), especially for graphics-intensive tasks such as rendering in drafting software.

upvoted 1 times

⊟ **AustinKelleyNet** 1 year, 11 months ago

Ignore previous comment. The answer provided is correct.

upvoted 1 times

Question #6 *Topic 1*

A Chief Information Security Officer (CISO) is evaluating the company's security management program. The CISO needs to locate all the assets with identified deviations and mitigation measures. Which of the following would help the CISO with these requirements?

A. An SLA document

B. A DR plan

C. SOC procedures

D. A risk register

Suggested Answer: *D*

Community vote distribution

D (100%)

SimplyDebonair Highly Voted 👍 2 years, 3 months ago

Selected Answer: D

The correct answer would be "D." A risk register would outline all the assets with identified deviations and their mitigation measures. It would it also outline the risks, the risk(s) descriptions, the impact of the risks, and their likelihood of occurring. SOC procedures is poorly worded and doesn't clarify on whether this for your actual SOC's documentation/procedure processes. Or if this is related to SOC reports (Type I, II, or III) when it comes to GRC or RMF functions.

upvoted 8 times

SimplyDebonair 2 years, 2 months ago

All-in-One CV0-003 pgs. 395-397

SOC Procedures – outlines the individual steps required to complete a task. Furthermore, security procedures ensure that those who follow the procedures will do the following:

- Perform the task consistently.
- Take the predictable amount of time to perform the task.
- Require the same resources each time the task is performed.

Risk Register – a document that tracks: the lists of risks, a description of the risk, the impact of the risk would have on the business if actualized, and the likelihood of the risk.

- Risk registers may document mitigating controls that reduce the risk. If they're mentioned, the register will then show what the residual risk is after the mitigating control is factored in.

upvoted 7 times

kuzummjakk Most Recent 4 months ago

Selected Answer: D

The question is asking for the effect of a vulnerability scanner (in a confusing way) so yeah D.

upvoted 1 times

yyCherubim 7 months, 3 weeks ago

Selected Answer: D

Although testpreplabs states this answer is C, a risk register makes more sense to me.

upvoted 1 times

Jt11 2 years ago

Selected Answer: D

Agree with you, it should be D Risk Register.

upvoted 4 times

Question #7 Topic 1

A cloud engineer is responsible for managing a public cloud environment. There is currently one virtual network that is used to host the servers in the cloud environment. The environment is rapidly growing, and the network does not have any more available IP addresses. Which of the following should the engineer do to accommodate additional servers in this environment?

A. Create a VPC and peer the networks.

B. Implement dynamic routing.

C. Enable DHCP on the networks.

D. Obtain a new IPAM subscription.

Suggested Answer: *A*

Community vote distribution

A (80%) D (15%) 5%

braveheart22 3 months ago

Option D is the right answer to this question.
Obtain a new IPAM (IP Address Management) subscription. This will allow the engineer to acquire additional IP addresses to accommodate the growing number of servers in the cloud environment. In the AWS Cloud Ecosystem IPAM is used by cloud engineers and administrators to request more IP addresses whenever there is IP address shortage and needs IP address scaling.

upvoted 1 times

Sweety_Certified7 9 months ago

Selected Answer: A

A. Create a VPC and peer the networks.
Explanation:

Creating a Virtual Private Cloud (VPC) allows for the segmentation and isolation of network resources within the cloud environment.
Peering the networks enables communication between multiple VPCs while keeping them logically separate.
By creating additional VPCs and peering them with the existing network, the engineer can effectively expand the available IP address space and accommodate the growing number of servers in the environment.

Option D (voted by 20%), obtaining a new IPAM (IP Address Management) subscription, may help manage IP addresses but does not inherently expand the IP address space.

upvoted 3 times

kuzummjakk 10 months ago

Selected Answer: A

A VPC lets you use "private space" aka non-routable IPs, meaning you can pretty much extensively use all the possible IPs bc you dont have to worry about not using the same IP as some random public device. With peering, you can set up a public machine that can access private space, so a connection can use that machine to jump from public to private, making private routable.

upvoted 2 times

FrancisDrake 11 months, 1 week ago

Selected Answer: A

I think the clue in the question is this "There is currently one virtual network..." and I'm not sure how IPAM helps if you're out of IP addresses.

upvoted 4 times

AllenTaylor 11 months, 3 weeks ago

Selected Answer: A

Create a VPC and Peer the Networks. This option directly addresses the issue by expanding the available IP address space through the creation of a new network (VPC) and then linking it with the existing network for seamless operation. This approach not only resolves the immediate problem but also provides a scalable structure for future expansion.

upvoted 2 times

Pongsathorn 1 year, 3 months ago

Selected Answer: A

Peering
Network connectivity that permits instances to communicate between two virtual private clouds using reserved IP addresses.

Peering connects two or more virtual networks. The virtual networks appear to consumers as a single network. In addition, fast connectivity is provided between the two networks, making data and resource access very efficient.

Peering is used in the hub-and-spoke model to connect the spoke networks with the hub network. Note that the spoke networks are not peered to each other in the hub-and-spoke model.

Ref: The Official CompTIA Cloud+ Study Guide (Exam CV0-003)

upvoted 3 times

SecPlus2022 1 year, 6 months ago

Selected Answer: A

Neither "B", "C" or "D" will provide addresses to the VPC in question that it doesn't already have. Creating a new VPC and peering the two will allow growth on the new VPC and both VPC's can communicate as if they're on the same network.

upvoted 1 times

PatrickH 1 year, 7 months ago

Its definatly NOT C. DHCP cannot give out addresses you just dont have. Its either A or D. I will do a bit more research and come back with my take.

upvoted 1 times

Zak11 1 year, 8 months ago

Selected Answer: C

DHCP is a protocol that automatically assigns IP addresses to devices on a network. By enabling DHCP on the network, the cloud engineer can automatically assign IP addresses to new servers as they are provisioned in the environment. This will allow the environment to accommodate additional servers without having to manually configure IP addresses.

upvoted 1 times

samuel186 1 year, 8 months ago

I think it's C. Enabling DHCP on the networks would be the most appropriate solution to accommodate additional servers in the public cloud environment. DHCP (Dynamic Host Configuration Protocol) is a network protocol that allows the automatic assignment of IP addresses and other network configuration settings to devices on a network. By enabling DHCP on the virtual network, the cloud engineer can automatically assign IP addresses to new servers as they are added to the network, without needing to manually assign IP addresses.

upvoted 1 times

betty_boop 1 year, 8 months ago

To accommodate additional servers in the environment when the current virtual network has no more available IP addresses, the cloud engineer should enable DHCP on the network.

Dynamic Host Configuration Protocol (DHCP) is a network protocol that automatically assigns IP addresses to devices on a network. By enabling DHCP on the network, the engineer can ensure that any new servers that are added to the environment will be automatically assigned an available IP address.

Creating a new Virtual Private Cloud (VPC) and peering the networks or implementing dynamic routing are other options that could help with network scalability. However, in this scenario, enabling DHCP on the existing network is likely the most straightforward and efficient solution.

Obtaining a new IP Address Management (IPAM) subscription is not necessary, as enabling DHCP on the existing network can handle the situation.

upvoted 1 times

concepcionz 1 year, 9 months ago

Selected Answer: A

Obtaining a new IPAM subscription may provide additional IP addresses, but it is not the most efficient or cost-effective solution. Creating a new network with a VPC and peering the networks is a better approach as it allows for more flexibility in managing and scaling the environment.

upvoted 1 times

TestDummies 1 year, 10 months ago

A is the correct answer, at first I thought IPAM as well, but it would not address the issue of running out of IP addresses, just managing them. So VPC would allow more.

upvoted 4 times

⊟ **AustinKelleyNet** 1 year, 11 months ago

Selected Answer: D

I think D for reasons already stated by TheGinjaNinja

upvoted 1 times

⊟ **TheGinjaNinja** 1 year, 11 months ago

Selected Answer: D

D. Obtain a new IPAM (IP Address Management) subscription. This will allow the engineer to acquire additional IP addresses to accommodate the growing number of servers in the cloud environment.

upvoted 2 times

⊟ **BaseliosG** 2 years, 4 months ago

Answer A

upvoted 4 times

Question #8 *Topic 1*

A system administrator is migrating a bare-metal server to the cloud. Which of the following types of migration should the systems administrator perform to accomplish this task?

A. V2V

B. V2P

C. P2P

D. P2V

Suggested Answer: *D*

Community vote distribution

D (100%)

BaseliosG Highly Voted 10 months, 1 week ago

Physical to virtual it is . D

upvoted 6 times

AustinKelleyNet Most Recent 5 months ago

Selected Answer: D

The answer provided is correct

upvoted 3 times

Question #9 *Topic 1*

A company is utilizing a private cloud solution that is hosted within its datacenter. The company wants to launch a new business application, which requires the resources below:

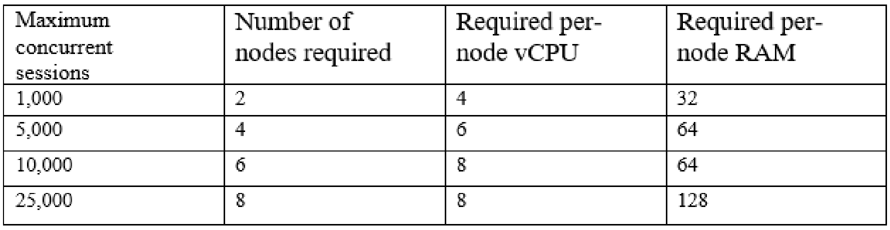

Maximum concurrent sessions	Number of nodes required	Required per-node vCPU	Required per-node RAM
1,000	2	4	32
5,000	4	6	64
10,000	6	8	64
25,000	8	8	128

The current private cloud has 30 vCPUs and 512GB RAM available. The company is looking for a quick solution to launch this application, with expected maximum sessions to be close to 24,000 at launch and an average of approximately 5,000 sessions. Which of the following solutions would help to company accommodate the new workload in the SHORTEST amount of time and with the maximum financial benefits?

A. Configure auto-scaling within the private cloud.

B. Set up cloud bursting for the additional resources.

C. Migrate all workloads to a public cloud provider.

D. Add more capacity to the private cloud.

Suggested Answer: *B*

Community vote distribution

B (80%) D (20%)

cbo Highly Voted 1 year, 11 months ago

Why not auto-scaling A?

upvoted 5 times

FrancisDrake 10 months, 3 weeks ago

According to the scenario there aren't enough internal resources.

upvoted 2 times

kuzummjakk 10 months ago

But the chart says it'd take less than "30 vcpu's" and "512GB" of memory.

I think it's moreso that while auto-scaling WOULD work, autoscaling is more appropriate for when the average load increases, not when there's a short period of a boost in activity. Cloud bursting is just "more relevant".

upvoted 1 times

TheFivePips 7 months, 3 weeks ago

if I'm understanding this correctly, then the current PC doesn't have enough resources for 24,000 sessions, but could handle the average 5000. It is a little unclear because you arnt given the exact numbers to work with but we know for sure we don't have enough for 25,000 sessions. So while auto scaling could work on a potentially lower load (say something like 20,000 maybe), I don't think it would work here. So really the only other feasible option is cloud bursting.

upvoted 1 times

TurtleManWeedler Most Recent 1 month, 1 week ago

Selected Answer: B

It's B. At 25,000 sessions it requires 8 nodes and 8 vCPUs per node. The Maximum number of vCPUs they have is 30. (8 Nodes x 8 vCPUS = 64 vCPUs.) Without even needing to look at the rest of the question we can see they don't have the requirements to support 25,000 sessions, though they state the max sessions to be 24,000 for them. That being said Cloud bursting would allow them to pay for the extra resources that they need ONLY when they need it. It's not auto-scaling because were talking about an On-Prem stiuation that would need to get it's additional resources from the clous and autoscaling typically refers to a cloud resource scaling up or out from within the cloud or at least isn't as defined as Cloud-Bursting is. Adding more capacity is certainly not cost effective considering they don't always have the maximum workload of 24,000 sessions. B is your best bet.

upvoted 1 times

Sweety_Certified7 10 months, 1 week ago

Selected Answer: B

Given the requirement for the shortest time to accommodate the new workload and maximum financial benefits, option B: Set up cloud bursting for additional resources appears to be the most suitable choice.
Cloud bursting leverages the existing private cloud infrastructure while providing additional resources from the public cloud as needed, offering scalability, quick deployment, and potential cost savings.

Answer from GPT

upvoted 2 times

⊟ **yyCherubim** 1 year, 1 month ago

Selected Answer: B

There are a whole lot of distracting numbers and words in this question; stick with simple.

upvoted 1 times

⊟ **Pongsathorn** 1 year, 3 months ago

Selected Answer: B

Cloud Bursting
Another approach to meeting service requirements is cloud bursting. Cloud bursting relies on a hybrid cloud model. If the private, on-premises cloud service reaches saturation, the workload can spill over into the public cloud. Of course, using public cloud resources incurs costs, but such costs are better than poor or no service due to a lack of resources.

upvoted 3 times

⊟ **betty_boop** 1 year, 8 months ago

Selected Answer: B

B. Cloud bursting

upvoted 1 times

⊟ **AustinKelleyNet** 1 year, 11 months ago

Selected Answer: B

I like B

upvoted 1 times

⊟ **Lenell** 2 years ago

Selected Answer: B

Cloud Bursting can be used for both compute and storage. This question is about compute capability. "Compute Bursting" unleashes the high-performance compute capabilities of the cloud for processing locally created datasets. (reference: https://www.ctera.com/it-initiatives/cloud-bursting/)

upvoted 1 times

⊟ **josernan** 2 years, 3 months ago

cloud bursting Concept of running an application on the organization's
internal computing resources or private cloud and "bursting" that application
into a public cloud on demand when the organization runs out of resources on
its internal private cloud.

upvoted 1 times

⊟ **Moosafat** 2 years, 3 months ago

Selected Answer: D

The quickest method would be to configure auto-scaling within the private cloud. There is no mention of the company having any public cloud solutions which is what cloud bursting consists of.

upvoted 3 times

⊟ **SimplyDebonair** 2 years, 9 months ago

Selected Answer: B

The correct answer is "B." Cloud Bursting will provide the quickest deployment for the private cloud by overflowing into the public cloud. This will save time and money in order to provide continued service without interruption.

https://azure.microsoft.com/en-us/overview/what-is-cloud-bursting/

upvoted 2 times

⊟ **SimplyDebonair** 2 years, 9 months ago

After some re-evaluations with my peers, I would like to change my answer to "D." For the following reasons:

For maximum financial benefit, the load has to be able to handle a maximum expected capacity without dropping sessions. The different scaling unfortunately makes this portion confusing... however, it is enough for a 5,000 maximum (with the expectation of a 5,000 average), there will be some that wouldn't be able to be accommodated for. For more clarity:

If 5,000 concurrent sessions are needed, the 4 nodes are required. Each node requires 6 vCPUs at the 5,000 mark. So if you're math is 4 nodes x 6 vCPU per node, you will equal out to 24 vCPUs. By doing the same thing for RAM, each node requires 64GB of RAM at the 5,000 mark. Then at that point, your math would be 4 nodes x 64GB of RAM, you will equal out to 256GB of RAM at that level. That is for the "AVERAGE" consumption.

So for the initial consumption, you wouldn't have enough resources.

upvoted 7 times

TheFivePips 7 months, 3 weeks ago

That option is by no means the cheapest or the fastest however, since it would require purchasing and standing up hardware in your own data center.

upvoted 1 times

JeanClaud 3 years, 1 month ago

Selected Answer: B

Correct answer is cloud bursting due to shortest amount of time and most financial benefit.

upvoted 1 times

TT 3 years, 3 months ago

I saw this same question in the CV1-003 and they had the answer as "Auto-scaling". B is correct. The shortest amount of time and provides the "financial benefits" as well.

upvoted 2 times

dvd21 3 years, 4 months ago

B. Shortest amount of time.

upvoted 4 times

Question #10 *Topic 1*

A cloud administrator is reviewing the authentication and authorization mechanism implemented within the cloud environment. Upon review, the administrator discovers the sales group is part of the finance group, and the sales team members can access the financial application. Single sign-on is also implemented, which makes access much easier. Which of the following access control rules should be changed?

A. Discretionary-based

B. Attribute-based

C. Mandatory-based

D. Role-based

Suggested Answer: *D*

Community vote distribution

D (100%)

cybertechb 4 months, 3 weeks ago

Attribute-based access control (ABAC) allows access decisions to be based on various attributes associated with users, resources, and environmental conditions. In this case, the administrator should review and potentially modify the attributes associated with the sales and finance groups to ensure that members of the sales group do not have access to the financial application unless explicitly authorized. This could involve adjusting the attributes that define group membership or refining the access policies based on specific attributes such as job role or department

upvoted 1 times

betty_boop 1 year, 2 months ago

Selected Answer: D

Role based

upvoted 1 times

AustinKelleyNet 1 year, 5 months ago

Selected Answer: D

The answer provided is correct

upvoted 4 times

Question #11 *Topic 1*

A company developed a product using a cloud provider's PaaS platform and many of the platform-based components within the application environment. Which of the following would the company MOST likely be concerned about when utilizing a multicloud strategy or migrating to another cloud provider?

A. Licensing

B. Authentication providers

C. Service-level agreement

D. Vendor lock-in

Suggested Answer: *D*

Community vote distribution

D (100%)

dvd21 Highly Voted 2 years, 10 months ago

D. Lock-in. Using a platform, apps can be tied to the vendor's implementation of the tech stack.

upvoted 7 times

TheFivePips Most Recent 1 month, 2 weeks ago

Selected Answer: D

Here's why the other options are less likely to be the primary concern:

A. Licensing: While licensing considerations may be important, especially when migrating to another cloud provider, they are generally more manageable compared to the challenges of vendor lock-in. Licensing issues can often be addressed through negotiations with vendors or by selecting cloud providers with favorable licensing terms.

B. Authentication providers: Authentication providers, such as identity and access management (IAM) services, are important for security and user management but are typically more adaptable to different cloud environments. While there may be some differences in authentication mechanisms between cloud providers, they are generally less restrictive than platform-specific components.

C. Service-level agreement (SLA): Service-level agreements define the terms of service availability, performance, and support provided by a cloud provider. While SLAs are important for ensuring reliability and performance, they are generally negotiable and can be adjusted to meet the company's requirements when migrating to another cloud provider. However, vendor lock-in may limit the flexibility to negotiate favorable SLAs with alternative providers.

upvoted 2 times

betty_boop 1 year, 2 months ago

Selected Answer: D

Vendor lock in

upvoted 1 times

AustinKelleyNet 1 year, 5 months ago

Selected Answer: D

D is correct

upvoted 1 times

Jt11 2 years ago

Selected Answer: D

Agree, should be D. Using all the custom features of PAAS risks vendor lock in, which would make multi cloud difficult.

upvoted 3 times

SimplyDebonair 2 years, 3 months ago

Selected Answer: D

"D" would be the correct answer due to the nature of the scenario, “A company developed a product using a cloud provider’s PaaS platform…” Using proprietary tools to develop something in one cloud, then moving to another cloud or multi-cloud strategy risks vendor lock-in.

upvoted 4 times

SimplyDebonair 2 years, 3 months ago

"D" would be the correct answer due to the nature of the scenario, “A company developed a product using a cloud provider’s PaaS platform…”
Using proprietary tools to develop something in one cloud, then moving to another cloud or multi-cloud strategy risks vendor lock-in.

upvoted 1 times

Greg01 2 years, 3 months ago

Selected Answer: D

D would be correct

upvoted 2 times

JeanClaud 2 years, 7 months ago

Selected Answer: D

I agree with DVD. Pretty clear the answer is D here.

upvoted 2 times

Question #12 *Topic 1*

A systems administrator is trying to establish an RDP session from a desktop to a server in the cloud. However, the connection appears to be refused even through the VM is responding to ICMP echo requests. Which of the following should the administrator check FIRST?

A. The firewall

B. The subnet

C. The gateway

D. The services

Suggested Answer: *A*

Community vote distribution

A (79%) D (21%)

ramrod1738 3 months, 1 week ago

The first thing the systems administrator should check when trying to establish an RDP session from a desktop to a server in the cloud and encountering a connection error is the network firewall configuration.

Firewall rules can block incoming RDP traffic, even if the server is responding to ICMP echo requests. The administrator should verify that the firewall is configured to allow incoming RDP traffic on the appropriate port (typically TCP port 3389). They should also check if the firewall is blocking incoming RDP traffic for the specific IP address or network range that the desktop is using.

If the firewall configuration is correct and the connection is still refused, the administrator should also check the network security groups or network access control lists in the cloud environment to ensure that they are configured to allow incoming RDP traffic.

Finally, the administrator should verify that the server has Remote Desktop Services enabled and that a remote desktop connection is allowed for the specific user account.

upvoted 1 times

Zak11 3 months, 1 week ago

Selected Answer: A

the first thing the administrator should check is the firewall.

Firewalls are designed to prevent unauthorized access to a network or system by blocking or allowing specific types of traffic. If the firewall is not configured to allow RDP traffic, then the connection will be refused. The administrator should verify that the firewall rules are correctly configured to allow RDP traffic to pass through.

While the subnet, gateway, and services can also impact the ability to establish an RDP connection, they are less likely to be the cause of the problem in this scenario.

upvoted 3 times

weaponxcel 3 months, 1 week ago

Selected Answer: A

A should be correct

When trying to establish a Remote Desktop Protocol (RDP) session and the connection is refused, one of the most common reasons is that the firewall is blocking the RDP port (typically port 3389). Firewalls are designed to control incoming and outgoing network traffic based on an organization's security policy. If the RDP port is not explicitly allowed, the connection will be refused.

D. The services (Wrong)

While it's essential to ensure that the RDP service is running on the server, the initial connection refusal is more likely due to a firewall rule blocking the RDP port rather than the service not running.

upvoted 2 times

> **kuzummjakk** 10 months ago
>
> "Services" here is in the context of cloud services so for aws, "the ec2 service"
>
> upvoted 1 times

E_Byte 5 months, 2 weeks ago

A firewall can refuse or blackhole the connection attempt. Blackholing is more secure, keeps the threat actor guessing. If a service is not bound to the port (TCP/3389) then the OS will usually refuse connections, A and D are both likely problems because firewalls are usually set to block all traffic except what is allowed. Obviously ICMP is allowed, is RDP, and from the tech's IP. Best practice is to limit where you can RDP from.
Bad Question... but I would likely go with A and probably get it wrong.

upvoted 1 times

⊟ **TheFivePips** 7 months, 3 weeks ago

Selected Answer: A

While ICMP echo requests (ping) being successful suggests that network connectivity is established and that the VM is reachable over the network, it does not necessarily mean that all types of traffic are allowed through the firewall.

Firewalls can be configured to allow or block specific types of traffic based on predefined rules. For example, ICMP echo requests (ping) may be allowed through the firewall while other types of traffic, such as RDP (Remote Desktop Protocol), may be blocked.

In this scenario, even though ICMP echo requests are successful, the RDP connection is being refused, indicating that the firewall may be blocking RDP traffic.

upvoted 1 times

⊟ **kuzummjakk** 10 months ago

Services here is in the context of "cloud services" like aws's EC2 or VPC services. In the cloud, you don't mess with firewall rules, you interact with the providers services. Now, say one of those services is a firewall, it would be more cloud agnostic to say "the services" since a "firewall service" isn't quite a staple in cloud networks. They could all present networking configurations like block/allow in a different manner.

upvoted 1 times

⊟ **kuzummjakk** 10 months ago

Although, "first" gives my explanation a hole. It insinuates that yes, a firewall is something you'd address directly in the cloud. This is probably lower on the list since it very well could vary what equates to a "firewall" cloud-to-cloud.

upvoted 1 times

⊟ **Chiaretta** 11 months, 1 week ago

Selected Answer: D

The given answer is correct. All others state networking problems.

upvoted 1 times

⊟ **Chiaretta** 11 months, 3 weeks ago

Selected Answer: D

Given that the server and desktop are inside the cloud and server is responding to ICMP echo requests but the RDP connection is being refused, the administrator should first check the SERVICE, all other question mention network problem.

upvoted 1 times

⊟ **FrancisDrake** 11 months, 1 week ago

You can ping a machine even if the firewall is blocking the port or service.

upvoted 4 times

⊟ **veliyath** 1 year ago

Given that the server in the cloud is responding to ICMP echo requests but the RDP connection is being refused, the administrator should first check A. The firewall.

The firewall settings might be configured to allow ICMP (ping) requests but could be blocking RDP (Remote Desktop Protocol) traffic. Checking the firewall rules and ensuring that the appropriate ports for RDP (TCP port 3389 by default) are open and allowed through the firewall would be the logical first step in troubleshooting the connectivity issue for establishing an RDP session.

upvoted 2 times

⊟ **yyCherubim** 1 year, 1 month ago

Selected Answer: D

The question states the two systems are successfully talking to each other via ICMP, so not primarily a firewall problem. If you FIRST verify the services are up and running, then check the firewall rules.

upvoted 1 times

⊟ **FrancisDrake** 11 months, 1 week ago

The connection is being refused. Sounds like a firewall. You can ping a machine even if you are blocked by a firewall.

upvoted 1 times

⊟ **mutatoo** 1 year, 6 months ago

RDP is layer 7- so therefore services is the answer as ICMP is going through (ICMP is on layer 3-Firewall). therefore that is clearly not it.

upvoted 1 times

maelo 1 year, 6 months ago

Selected Answer: A

I clearly would check firewall 1st, as this in handled by the CSP's frontend. It's faster checking, and a probable deny due to security reasons. Checking service availability without RDP established and thus without visual access to the instance, is anyways harder to achieve.

upvoted 3 times

betty_boop 1 year, 8 months ago

Selected Answer: A

Firewall

upvoted 2 times

concepcionz 1 year, 9 months ago

Selected Answer: A

The connection "refused" message typically indicates that there is an issue with the network or firewall settings.

upvoted 2 times

bagsik89 1 year, 10 months ago

Selected Answer: D

I would check the services on the server FIRST before even considering firewall rules. Don't overthink troubleshooting.

upvoted 1 times

FrancisDrake 11 months, 1 week ago

The key word is "refused". Something is not allowing the connection.

upvoted 1 times

AustinKelleyNet 1 year, 11 months ago

Selected Answer: A

It seems like a firewall rule is in place blocking the connection.

upvoted 1 times

TheGinjaNinja 1 year, 12 months ago

Selected Answer: A

Firewall

upvoted 2 times

Question #13 *Topic 1*

Which of the following would be the BEST option for discussion of what individuals should do in an incident response or disaster recovery scenario?

A. A business continuity plan

B. Incident response/disaster recovery documentation

C. A tabletop exercise

D. A root cause analysis

Suggested Answer: *C*

Community vote distribution

C (100%)

dvd21 Highly Voted 3 years, 4 months ago

C. Tabletop exercise. The question asks for discussion.

upvoted 6 times

TurtleManWeedler Most Recent 1 month, 1 week ago

Selected Answer: C

C since it asks for a discussion but there were so many more practical ways to ask this question. This question is terribly worded, and it is likely the type of wording you would see on the real exam too. CompTIA needs to learn how to write and stop wasting people's time and money.

upvoted 1 times

braveheart22 3 months ago

Well, I go with option A (Business Continuity Plan) The question is not asking for a discussion, rather it says "Which of the following would be the BEST option for discussion of WHAT INDIVIDUALS SHOULD DO in an incident response or disaster recovery scenario?" In other words, the key discussions during a Tabletop exercise shoud be the Business Continuity Plan (BCP) which is a blueprint of condensed strategies that will help the organization recover from a disaster and continue in business.

upvoted 1 times

Pongsathorn 1 year, 3 months ago

Selected Answer: C

The BEST option for discussing what individuals should do in an incident response or disaster recovery scenario is:

C. A tabletop exercise

A tabletop exercise is a practical and interactive approach to simulate various disaster scenarios, allowing individuals to discuss and practice their roles and responsibilities during an incident. It provides an opportunity for teams to work together, identify potential issues, and refine their incident response and disaster recovery plans. It is an effective way to ensure that everyone understands what they should do in a controlled and realistic environment. While items like business continuity plans, incident response/disaster recovery documentation, and root cause analyses are essential components of incident management, a tabletop exercise directly involves individuals in the discussion and practical application of their roles during an incident.

upvoted 4 times

betty_boop 1 year, 8 months ago

Selected Answer: C

Tabletop exercise

upvoted 1 times

AustinKelleyNet 1 year, 11 months ago

Selected Answer: C

C is correct

upvoted 1 times

scott5010 2 years, 1 month ago

Selected Answer: C

C is the best answer

upvoted 1 times

Admiral_Crunch 2 years, 6 months ago

C, A business continuity plan is not the place to have the discussion on what to do that's where you keep the ideas that have been decided upon during the tabletop exercise.

upvoted 1 times

SimplyDebonair 2 years, 9 months ago

Selected Answer: C

The correct answer would be "C." The scenario inquires what would be the "...BEST option for discussion..." for IR or DR scenario. A tabletop exercise would informally outline the IR/DR instance with team members, their roles, and their responses to that instance.

https://cybersecurity.wa.gov/tabletop-exercises
https://euclidsecurity.com/2020/12/17/what-is-a-cybersecurity-tabletop-exercise/

upvoted 3 times

SimplyDebonair 2 years, 9 months ago

o C would be the correct answer due to the scenario inquiring what would be the "...BEST option for discussion..." for IR or DR scenario. A tabletop exercise would informally outline the IR/DR instance with team members, their roles, and their responses to that instance.

upvoted 1 times

Hobbit 2 years, 11 months ago

Tabletop exercises are discussion-based sessions where team members meet in an informal, classroom setting to discuss their roles during an emergency and their responses to a particular emergency situation. A facilitator guides participants through a discussion of one or more scenarios

upvoted 3 times

DVD300 3 years, 1 month ago

it should be tabletop i did not get the question well

upvoted 3 times

DVD300 3 years, 1 month ago

A. the BCP

upvoted 1 times

> **Sal** 2 years, 2 months ago
>
> discussion is key to this question so its tabletop
>
> upvoted 1 times

Rodmas 3 years, 1 month ago

it should be A. the BCP

upvoted 1 times

Question #14 *Topic 1*

A systems administrator has migrated an internal application to a public cloud. The new web server is running under a TLS connection and has the same TLS certificate as the internal application that is deployed. However, the IT department reports that only internal users who are using new versions of the OSs are able to load the application home page. Which of the following is the MOST likely cause of the issue?

A. The local firewall from older OSs is not allowing outbound connections.

B. The local firewall from older OSs is not allowing inbound connections.

C. The cloud web server is using a self-signed certificate that is not supported by older browsers.

D. The cloud web server is using strong ciphers that are not supported by older browsers.

Suggested Answer: *D*

Community vote distribution

D (88%) 13%

braveheart22 3 months ago

After AI Review, I am changing my answer to "C" I spent some time reviewing my initial response to this question, and doing some google research as well. So I m finally leaning towards option "C"
Using a self-signed TLS certificate, common connection issues include: browser warning pop-ups alerting users about the untrusted certificate, potential for Man-in-the-Middle (MITM) attacks due to lack of trust validation, difficulty connecting to services that require trusted certificates, and user hesitation to proceed due to the security warnings displayed by their browser; essentially, any situation where a user's device cannot verify the authenticity of the certificate, leading to potential security concerns and disruption of the connection.

upvoted 1 times

braveheart22 3 months ago

After a careful examination of the question, I think option D is the correct answer.

upvoted 1 times

braveheart22 3 months ago

AI Review
Hey guys I have spent some time reviewing my initial response to this question, and doing some Google research as well. So I m finally leaning towards option hen using a self-signed TLS certificate, common connection issues include: browser warning pop-ups alerting users about the untrusted certificate, potential for Man-in-the-Middle (MITM) attacks due to lack of trust validation, difficulty connecting to services that require trusted certificates, and user hesitation to proceed due to the security warnings displayed by their browser; essentially, any situation where a user's device cannot verify the authenticity of the certificate, leading to potential security concerns and disruption of the connection.

upvoted 1 times

kuzummjakk 10 months ago

I understand why C is what it's looking for.
The question says that the only change was the server moving from internal to public. If it was a TLS version issue, they should've seen it when they were internal too. That crosses out D making C the only "possible" answer.

upvoted 1 times

TheFivePips 7 months, 3 weeks ago

Its not a TLS issue. Its an old OS issue. The question implies that there are new and old OSs internally, and only the old ones are having an issue. Now it is possible that the OS doesn't support self signed certificates, but do you really think that the older OS would have a more restrictive policy than the new one? self signed has been around for a while, I think its much more likely that it just doesn't support the new cryptographic protocols

upvoted 1 times

TheFivePips 7 months, 3 weeks ago

just to add on here, almost all browsers support self signed certificates, usually with a warning. It predates public key infrastructure, and was kind of the only game in town for a while. All that to say its more likely that the browser doesn't support the new cryptographic standards that come out every 3-5 years

upvoted 1 times

veliyath 1 year ago

The most likely cause of the issue based on the provided information is:

C. The cloud web server is using a self-signed certificate that is not supported by older browsers.

Older operating systems and browsers might not support or trust self-signed certificates, causing them to fail in establishing a secure connection with the web server. This issue would result in the inability of users on older OS versions to load the application home page despite having the same TLS certificate.

upvoted 1 times

nowaydude1 1 year, 2 months ago

Selected Answer: C

It's obviously C. So many people seem to have missed the memo on self-signed certificates. To host something on the web you need a PUBLIC CA to issue your certificate for it to be trusted. Since they were using the same certificate as when they had it running internally only... then its self-signed. The cipher thing is stupid. you really think everyone is running out of date browsers? come on, question the obvious. This is obviously a self-signed certificate issue.

upvoted 2 times

kuzummjakk 10 months ago

It specifically says that everyone with an old OS can't access it. The last sentence matters too.

upvoted 2 times

FrancisDrake 1 year ago

It does say internal users with the newer os. So I assume that internal users with the older os are also having issues. Also simply because it is self-signed does not necessarily mean that you cannot access the web server.

upvoted 3 times

Pongsathorn 1 year, 3 months ago

Selected Answer: D

The MOST likely cause of the issue where only internal users with newer OS versions can load the application home page in a TLS-encrypted connection is:

D. The cloud web server is using strong ciphers that are not supported by older browsers.

Here's why:

TLS (Transport Layer Security) connections involve encryption and decryption processes that rely on cryptographic ciphers. Older web browsers or OSs may not support the latest, most secure cryptographic ciphers due to security updates and compatibility issues.

upvoted 3 times

Pongsathorn 1 year, 3 months ago

When a web server is configured to use strong ciphers, older browsers or OSs that lack support for these ciphers will have difficulty establishing a secure connection. As a result, users with older systems may experience connection failures or loading issues when accessing the application.

To address this issue, the administrator should consider adjusting the server's cipher suite to support a wider range of clients, including those with older OSs and browsers. This can involve configuring the server to use a more backward-compatible cipher suite or enabling backward-compatible cipher suites in addition to strong ones. This approach ensures broader compatibility while still maintaining security.

upvoted 2 times

Zak11 1 year, 8 months ago

Selected Answer: D

The MOST likely cause of the issue is that the cloud web server is using strong ciphers that are not supported by older browsers. Strong ciphers provide a higher level of security but are often not supported by older browsers, which can lead to connection issues. To resolve the issue, the systems administrator can configure the cloud web server to support weaker ciphers that are compatible with older browsers.

upvoted 4 times

betty_boop 1 year, 8 months ago

Selected Answer: D

Strong ciphers

upvoted 1 times

Grayson2023 1 year, 10 months ago

I agree with C, the server is using the same TLS certificate both internally and externally. Regardless if the site was used internally or externally, the question states the same certificate was used.

upvoted 1 times

Grayson2023 1 year, 10 months ago

I recant this statement.

upvoted 2 times

AustinKelleyNet 1 year, 11 months ago

Selected Answer: D

This seems like an obvious answer.

upvoted 1 times

ramrod1738 1 year, 11 months ago

The MOST likely cause of the issue is that the cloud web server is using strong ciphers that are not supported by older browsers (Option D).

TLS is a protocol that encrypts data transmitted over the internet. When a client (such as a web browser) connects to a server (such as a web server), the two parties negotiate the encryption protocol and cipher that will be used for the connection. If the client and server do not support the same encryption protocols and ciphers, the connection will fail.

In this scenario, it is likely that the cloud web server is using strong ciphers that are not supported by older browsers. This would cause the connection to fail for internal users who are using older versions of the OSs.

upvoted 1 times

WeLikeSpamHere 1 year, 11 months ago

Selected Answer: D

Answer is D

upvoted 2 times

ryanzou 2 years, 3 months ago

Selected Answer: D

ANSWER MUST BE D, no doubts, met the same scenario before

upvoted 3 times

Admiral_Crunch 2 years, 6 months ago

D, I agree with bx88 statement.

upvoted 2 times

bx88 2 years, 6 months ago

Answer is D

The firewall policies should NOT have any impact in this scenario. If the policies block inbound or outbound traffics, users could not even access the application while it was hosted on-prem.

Since the new web server and has the same TLS certificate as the internal application that is deployed, certificate is less likely the cause of the issue.

If the web server is using strong ciphers that are not supported by older browsers (for example - the web server requires AES 256 bits while the old browser supports 128 and 192 bits only), the users using old browsers could not access the page. This is the most likely cause of the issue.

upvoted 2 times

dvd21 3 years, 1 month ago

Answer is D

upvoted 1 times

Question #15 *Topic 1*

A systems administrator is configuring RAID for a new server. This server will host files for users and replicate to an identical server. While redundancy is necessary, the most important need is to maximize storage. Which of the following RAID types should the administrator choose?

A. 5

B. 6

C. 10

D. 50

Suggested Answer: *A*

Community vote distribution

A (88%) 13%

⊟ **bx88** Highly Voted 👍 2 years, 6 months ago

Answer is A

Since the replication to another identical servers has taken care of high-availability responsibility already, we need to find out which RAID types provide maximum storage capacity.

There are N disks with capacity to store X amount of data

-- With RAID 5, the equivalent capacity of one drive is used for redundant information. Total storage capacity is (N-1) * X
-- With RAID 6, the equivalent capacity of two drives is used for redundant information. Total storage capacity is (N-2) * X
-- With RAID 10, the equivalent capacity of a half of total drives is used for redundant information. Total storage capacity is (N/2) * X
-- With RAID 50, the minimum equivalent capacity of two drives is used for redundant information. Maximum storage capacity is (N-2) * X

With this, RAID 5 is the best choice.

upvoted 7 times

⊟ **SecPlus2022** Highly Voted 👍 3 months, 1 week ago

Selected Answer: A

Since the minimum number of required disks for RAID 60 is 8, we'll use 8, 1TB drives for this example.

RAID 5 Usable storage = 7 TB
RAID 6 Usable Storage = 6 TB
RAID 50 Usable Storage = 6 TB
RAID 60 Usable Storage = 4 TB

This is not debatable, the best RAID to maximize storage is RAID 5. The question asks nothing about performance, only storage.

upvoted 5 times

⊟ **AllenTaylor** Most Recent 🕒 11 months, 3 weeks ago

Selected Answer: A

maximize storage = RAID 5

upvoted 1 times

⊟ **Zak11** 1 year, 8 months ago

Selected Answer: C

RAID 10 would be the best choice in this scenario. RAID 10 combines the benefits of RAID 1 (mirroring) and RAID 0 (striping) to provide both redundancy and improved performance. With RAID 10, data is mirrored across multiple drives for redundancy, and then the mirrored sets are striped for performance. This configuration provides good performance and fault tolerance, while maximizing the available storage space. RAID 5 and 6 are good choices for fault tolerance and performance, but they typically use more storage space for parity information. RAID 50 is a combination of RAID 5 and RAID 0, and would not be the best choice if maximizing storage is the most important need.

upvoted 1 times

⊟ **kuzummjakk** 10 months ago

Mirroring is heavier on size than the parity used in RAID 5.

upvoted 1 times

⊟ **bagsik89** 1 year, 10 months ago

In regards to storage capacity, RAID 5 will have more usable storage as RAID 6 stores more parity information than RAID 5.

Sorage Capacity = RAID 5

More Redundant = RAID 6(RAID 6 can withstand loss of at least two drives while RAID 5 can handle 1)

https://petri.com/raid-5-vs-raid-6/#What_are_the_best_use_cases_for_RAID_5

upvoted 1 times

⊟ **AustinKelleyNet** 1 year, 11 months ago

Selected Answer: A

Raid 5 has redundancy and gets more disk efficiency than Raid 6 and 10. Raid 0 has no redundancy.

upvoted 1 times

⊟ **ramrod1738** 1 year, 11 months ago

If the most important need is to maximize storage while ensuring data redundancy, the administrator should choose RAID 6 (Option B).

RAID 6 is a type of RAID that provides data redundancy and disk striping across multiple disks. In this setup, data is split across multiple disks and two sets of redundant data are stored. This means that if two disks fail, the data can still be reconstructed.

Compared to RAID 5, RAID 6 provides better data redundancy and can sustain the failure of two disks, but at the cost of some additional disk space being reserved for data redundancy.

If data redundancy is not a major concern and maximizing storage is the priority, RAID 0 or RAID 10 would be more appropriate. However, these RAID types do not provide data redundancy, so if a disk fails, data will be lost.

upvoted 1 times

⊟ **dvd21** 3 years, 1 month ago

Cleary answer is A.

upvoted 1 times

Question #16 Topic 1

A cloud architect is designing the VPCs for a new hybrid cloud deployment. The business requires the following:
✑ High availability
✑ Horizontal auto-scaling
✑ 60 nodes peak capacity per region
✑ Five reserved network IP addresses per subnet
✑ /24 range
Which of the following would BEST meet the above requirements?

A. Create two /25 subnets in different regions.

B. Create three /25 subnets in different regions.

C. Create two /26 subnets in different regions.

D. Create three /26 subnets in different regions.

E. Create two /27 subnets in different regions.

F. Create three /27 subnets in different regions.

Suggested Answer: *A*

Community vote distribution

A (50%) F (20%) D (20%) 10%

TheFivePips 1 month, 2 weeks ago

Selected Answer: B

2^(number of host bits) = # of addresses
ex. /24 means there are (32 -24) 8 host bits. 2^8 = 256

/26 subnets only have 64 addresses. With 5 reserved addresses that means they only have 59 usable addresses. This is insufficient for the question requirement of 60 nodes per region.
/27 subnets only have 32 addresses. similarly insufficient.
That leaves only /25 subnets

A. Create two /25 subnets in different regions.

Each subnet has 128 IP addresses, but with five reserved addresses, the usable IP addresses would be 123.
The total available capacity across both subnets would be 2 * 123 = 246.

B. Create three /25 subnets in different regions.

Similar to option A, each subnet has 128 IP addresses, but with five reserved addresses, the usable IP addresses would be 123.
The total available capacity across three subnets would be 3 * 123 = 369

The question does not concern itself with cost, overheard, complexity, or performance. Since it asks for what BEST fits the requirements, I will say option B provides slightly more capacity and fault tolerance with an additional subnet.
upvoted 2 times

TheFivePips 1 month, 2 weeks ago

I am changing this answer to A because I overlooked that you cannot create more than 2 subnets from a /24 network. To subnet a network, you borrow bits from the host portion to create smaller subnets. The number of bits you borrow determines the number of subnets you can create and the size of each subnet.

Here's how you can determine the number of subnets you can create:

Determine the Number of Subnet Bits: Subtract the original prefix length from the desired new prefix length. For example, if you want to subnet a /24 network into /26 subnets, you would subtract 24 from 26, resulting in 2 subnet bits.

Calculate the Number of Subnets: Raise 2 to the power of the number of subnet bits. This gives you the number of subnets you can create. For example, if you have 2 subnet bits, you can create 2^2 = 4 subnets.

this means that from a /24 network you can create only 2 /25 subnets(25-24 = 1, 2^1 = 2)
4 subnets of /26 (26-24 = 2, 2^2 = 4) (theses ones still fail to make enough nodes per region however)
8 subnets of /27 (27-24 = 3, 2^3= 8)
upvoted 1 times

kuzummjakk 4 months ago

Selected Answer: F

The "Create two/three" and the "HA" and "horizontal scaling" threw me off.

Requirements: 60 IPs total (per region), 5 reserved IPs (per subnet)
/27 gives you 30 IPs. Since the reserved IPs are taken PER subnet, it'll be 25 per subnet. If you make 3 of these, you have 75 IPs for the nodes.
upvoted 2 times

TheFivePips 1 month, 2 weeks ago

The requirement is for 60 per region, or 60 per subnet, if I understand the wording of the answers correctly.

This is a tricky question because
1) you cannot even make a subnet in a different regions from the original network as far as I know (making this entire question impossible). This might be possible on some CSP but if it is, I dont know about it.
2) you need to know how to calculate how many subnets you can make from a given network 3) you need to know how many hosts you can make for each subnet
upvoted 1 times

anonymonkey 7 months, 3 weeks ago

E.
Five reserve network IP addresses per subnet denotes /27 (/25 yields 2, /26 yields 4 & /27 yields 8 meaning a /27 is required for 5 network IP from the subnetted /24)
60 nodes peak capacity per region means you need two /27 with 30 usable each by giving 5 bits to the host field. (1 yields 2, 2 yields 4, 3 yields 8, 4 yields 16 and 5 yields 32. Minus two IPs for the network and broadcast leaving 30.) Two of these would be peak capacity of 60.
Answer is two /27 subnets in different regions.
upvoted 2 times

SecPlus2022 1 year ago

Selected Answer: A

You need 65 IPs per region. /27 will only give you 32 and /26 will only give you 64. You need a /25 network which will give you 128 per region. Having two regions will give you the requirement of high availability. Yes, 3 regions is better than 2, but it doesn't ask for the best high availability option, it just asks for high availability which 2 regions provides.
upvoted 3 times

Securityguy42 3 months ago

This. I was thinking C. But as its pointed out, need 5 reserved IP addresses. So A makes the most sense.
upvoted 1 times

TheFivePips 1 month, 2 weeks ago

This question is hard. You literally cannot make 3 /25 subnets from a /24 network. A is correct because B is impossible, not because A is enough. The question doesn't say anything about cost or performance or anything, so choosing A over B because A provides enough HA is bad reasoning. IF B were possible it would have more than enough HA to meet all the requirements, just the same as A
upvoted 1 times

Zak11 1 year, 2 months ago

Selected Answer: D

The requirement of 60 nodes peak capacity per region suggests that the solution should be designed for scalability. Additionally, the requirement of high availability indicates that the solution should be distributed across multiple regions.

Each subnet should have at least five reserved network IP addresses. This leaves 251 addresses per subnet in a /24 range.

To meet these requirements, the BEST solution would be to create three /26 subnets in different regions. This would provide a total of 753 IP addresses per region (3 subnets x 251 IP addresses per subnet) and allow for horizontal auto-scaling up to 60 nodes per region. Having three subnets per region would provide redundancy and high availability.

upvoted 2 times

AustinKelleyNet 1 year, 5 months ago

Selected Answer: A

A is correct

upvoted 2 times

ramrod1738 1 year, 5 months ago

To meet the above requirements, the best option would be to create two /26 subnets in different regions (Option C).

A /24 network range provides 256 IP addresses, which would be too large for the requirements of 60 nodes peak capacity per region and five reserved network IP addresses per subnet.

A /26 subnet provides 64 IP addresses, which is a better fit for the requirements. With two /26 subnets in different regions, the architect can ensure high availability through load balancing and auto-scaling by allocating the IP addresses dynamically. Additionally, two subnets provide ample space for five reserved IP addresses and 60 nodes peak capacity per region.

Options A, B, E, and F would provide either too few or too many IP addresses to meet the requirements, so they are not the best choices.

upvoted 1 times

concepcionz 1 year, 3 months ago

A /26 subnet give you 62 usable host, 64 - 2 (Network and Broadcast) but also it's asking for 5 reserved IP address in that case you'll need at least 65 usable IPs, therefore the answer is A because you want to stay in the /24 range.

Hope that makes sense

upvoted 4 times

kuzummjakk 4 months ago

although ig it also says "create two" so you'd actually have "124" usable addresses. For /25 it also says "create two" so you'd way overshoot it with "252" usable addresses

upvoted 1 times

bx88 2 years ago

Correct answer is A

upvoted 3 times

SimplyDebonair 2 years, 3 months ago

"A" would be the correct answer for this scenario on the premise: HA is guaranteed through the horizontal scaling, allowing different subnets in different regions. With 60 nodes (IP addresses), you can count out the /26 (64 available IPs, when you need 5 IPs in reserve) and /27 (32 available IPs aren't enough). And finally, it is physically impossible to create three /25 subnets due to IPs multiplying/dividing in increments of 2 (i.e., /25 = 128 and /24 = 256).

upvoted 4 times

u2637ps 2 years, 4 months ago

So if the question states ony having a /24 to work with it would have to A

upvoted 1 times

Question #17 *Topic 1*

A company recently experienced a power outage that lasted 30 minutes. During this time, a whole rack of servers was inaccessible, even though the servers did not lose power. Which of the following should be investigated FIRST?

A. Server power

B. Rack power

C. Switch power

D. SAN power

Suggested Answer: *C*

Community vote distribution

C (92%) 8%

SecPlus2022 Highly Voted 1 year, 1 month ago

Selected Answer: C

The question tells you that the servers did not lose power. If the servers in the rack did not lose power, the problem is not the rack power. The most likely reason the severs were inaccessible is due to connectivity, thus I would check the switch power.

upvoted 9 times

TestTurtle468 Most Recent 6 days, 5 hours ago

Selected Answer: A

The racks didn't lose power nor did the servers (which if physically checked if possible) they would be fine. Sounds like a connection issue.

upvoted 1 times

Francois1984 10 months, 2 weeks ago

Selected Answer: C

according to chatgpt: In this scenario, where a company experienced a power outage lasting 30 minutes and a rack of servers became inaccessible even though they did not lose power, the first thing to investigate should be the network connectivity and network infrastructure

upvoted 3 times

betty_boop 1 year, 2 months ago

Selected Answer: B

B. Rack Power

upvoted 1 times

kuzummjakk 4 months ago

would make the server lose power

upvoted 2 times

ramrod1738 1 year, 5 months ago

B. Rack power. The whole rack of servers was inaccessible during the power outage, indicating a problem with the rack power. The first step in resolving the issue would be to investigate the rack power.

upvoted 2 times

kuzummjakk 4 months ago

But the servers did not lose power.

upvoted 2 times

Question #18 *Topic 1*

A cloud provider wants to make sure consumers are utilizing its IaaS platform but prevent them from installing a hypervisor on the server. Which of the following will help the cloud provider secure the environment and limit consumers' activity?

A. Patch management

B. Hardening

C. Scaling

D. Log and event monitoring

Suggested Answer: *B*

Community vote distribution

B (100%)

SimplyDebonair Highly Voted 👍 2 years, 3 months ago

Selected Answer: B

o B would be the correct answer because hardening a system means tightening down access or limiting access on the principal of least privilege. Log and event monitoring are processes utilized after hardening a system. Patch management would fit in the same mold as Log/Event monitoring, but scaling has no purpose for this scenario.

upvoted 6 times

FrancisDrake Most Recent ⏱ 4 months, 3 weeks ago

Selected Answer: B

Going with B. But why would you want to block a customer from installing a hypervisor...?

upvoted 2 times

> **FrancisDrake** 4 months, 3 weeks ago
>
> um duh. my bad. unfortunately I can't delete this and must live with it forever.
>
> upvoted 4 times

betty_boop 1 year, 2 months ago

Selected Answer: B

Hardening

upvoted 1 times

bx88 2 years ago

correct is B

upvoted 3 times

JeanClaud 2 years, 7 months ago

Selected Answer: B

Agree with TT. Log and event monitoring would not prevent anything. Hardening the machine would.

upvoted 2 times

TT 2 years, 9 months ago

I'm going with B. They want to prevent unauthorized access to a server. That's what hardening is for. Log and Event Monitoring won't prevent unauthorized access. Scaling and Patch Management make so sense here.

upvoted 3 times

Question #19 *Topic 1*

A resource pool in a cloud tenant has 90 GB of memory and 120 cores. The cloud administrator needs to maintain a 30% buffer for resources for optimal performance of the hypervisor. Which of the following would allow for the maximum number of two-core machines with equal memory?

A. 30 VMs, 3GB of memory

B. 40 VMs, 1,5GB of memory

C. 45 VMs, 2 GB of memory

D. 60 VMs, 1 GB of memory

Suggested Answer: *B*

Community vote distribution

B (76%) 12% 12%

TT Highly Voted 3 years, 3 months ago

The answer is B. after subtracting 30% from the memory and the cores, we end up with 63GB and 84 Cores left. We can make 42 2 core vms. That eliminates C and D. Quick math with the remaining options: A) 3GB x 30= 90GB. we only have 63GBs to work with so A is wrong. B) 1.5GB x 40= 60GB. There are leftover cores and memory but this is the best choice.

upvoted 20 times

samuel186 Most Recent 3 months, 1 week ago

Selected Answer: C

To calculate the number of VMs that can be hosted in the given resource pool, we need to account for the hypervisor buffer of 30%.

So, we have 90 GB of memory and need to reserve 30% for the hypervisor, which leaves us with 63 GB of usable memory.

We also have 120 cores, but we need to reserve 30% for the hypervisor, which leaves us with 84 cores available.

Each VM needs 2 cores and an equal amount of memory. So, we can calculate the number of VMs as follows:

84 cores / 2 cores per VM = 42 VMs

63 GB / 2 GB per VM = 31.5 VMs (rounded down to 31 VMs)

Therefore, the maximum number of two-core machines with equal memory that can be hosted in the resource pool is 31 VMs, each with 2 GB of memory.

So, the correct answer is (C) 45 VMs, 2 GB of memory. None of the other options would fit within the available resources, even accounting for the hypervisor buffer.

upvoted 1 times

maelo 1 year, 6 months ago

The above calculation sets memory as priority, but only number of VMs was asked to be maximized, regardless of mem. So either B with keeping a certain buffer, or D with even higher exploit of buffer.

upvoted 1 times

TheFivePips 7 months, 3 weeks ago

Selected Answer: B

Given:

Total memory: 90 GB

Total cores: 120

30% buffer for resources

2 cores per machine

Calculate the buffer:

Memory buffer: 30% of 90 GB = 0.3 * 90 GB = 27 GB

Core buffer: 30% of 120 cores = 0.3 * 120 cores = 36 cores

Now, let's determine the available memory and cores after considering the buffers:
Available memory after buffer: Total memory - Memory buffer = 90 GB - 27 GB = 63 GB
Available cores after buffer: Total cores - Core buffer = 120 cores - 36 cores = 84 cores

We also need to determine how many VMs we can even use, if each VM uses 2 machines, and we factor in the 30% buffer:
Maximum number of two-core machines = 84 cores / 2 cores per machine = 42 machines

A) 30Vms x 3GB = 90. Enough machines(30) but too much memory(90)
B) 40Vms x 1.5 GB = 60. Enough machines(40) and this seems within our memory limitations(60)
C) 45VMs x 2 GB = 90. Too many machines(45) and too much memory(90)
D) 60Vms X 1GB = 60. Too many machines(60) but within our memory limitations(60)

Best answer is B
upvoted 1 times

⊟ **kuzummjakk** 10 months ago
Selected Answer: B
"Correct answer" is wrong unless by "resources" the question only meant "memory" and not "cores" which wouldn't make sense.
upvoted 1 times

⊟ **Chiaretta** 11 months, 3 weeks ago
Selected Answer: B
The math says B
upvoted 3 times

⊟ **yyCherubim** 1 year, 1 month ago
Selected Answer: D
Simple math: 120 total cores, you need to make 2-core VMs. There was only one answer w/60 VMs. Next: With 90 GB of RAM, each VM gets 1GB, leaving 30 GB (which is also 30% of 90) remaining for the "Buffer".
upvoted 2 times

⊟ **kuzummjakk** 10 months ago
"resources" includes cores.
upvoted 1 times

⊟ **backdooranon** 1 year, 2 months ago
Selected Answer: B
Not C because 30% of 120 is 36, which leaves 84 usable cores, translating into 42 VMs max. 45 VMs translates into 90 cores. 42 VMs max means it is either A or B, however A provides each VM with 3GB of memory, translating into 90GB of memory in total. The buffer requires 30% of 90GB memory so only 63GB will actually be usable. Hence B is the only logical answer.
upvoted 2 times

⊟ **Zak11** 1 year, 8 months ago
Selected Answer: C
To maintain a 30% buffer for resources, the total available memory and cores will be:
Memory: 90 GB x 0.7 = 63 GB
Cores: 120 cores x 0.7 = 84 cores
Each two-core machine will require 4 GB of memory (2GB/core x 2 cores). Therefore, the maximum number of two-core machines that can be deployed with equal memory will be:

Memory: 63 GB / 4 GB per machine = 15.75 machines
Cores: 84 cores / 2 cores per machine = 42 machines
The option with the maximum number of machines is C. Therefore, the answer is:
C. 45 VMs, 2 GB of memory
upvoted 1 times

⊟ **maelo** 1 year, 8 months ago
Although also favouring B, here just an assumption: A "buffer" is a flexible safety, that can be exploited in case of needs + D keeps 30% buffer for RAM, while CPU will suffer. At least D doesn't oversubscribe resources (not considering the HV).
upvoted 1 times

betty_boop 1 year, 8 months ago

Selected Answer: B

40 VMs, 1,5GB of memory

upvoted 1 times

AustinKelleyNet 1 year, 11 months ago

Selected Answer: B

B is correct.

upvoted 1 times

ramrod1738 1 year, 11 months ago

The cloud administrator can calculate the buffer by taking 30% of the total resources, which is (90 GB + 120 cores) * 30% = 54 GB + 36 cores.

Then, subtract the buffer from the total resources to get the available resources, which is (90 GB - 54 GB) + (120 cores - 36 cores) = 36 GB + 84 cores.

Finally, divide the available resources by the number of cores and memory per machine to determine the maximum number of machines, which is (36 GB / 2 GB per machine) + (84 cores / 2 cores per machine) = 18 machines + 42 machines = 60 machines.

upvoted 1 times

Ty_G_S 2 years, 2 months ago

Selected Answer: B

120 - 30% (36) = 84 Cores for use

84/2 = 42 VMs with x2 cores

Only B is close enough to be the right answer based on the Price is Right math.

upvoted 2 times

Not_That_Guy 2 years, 2 months ago

Selected Answer: B

Only B allows for the 30% buffer.

upvoted 1 times

jiminycriminal 2 years, 3 months ago

Selected Answer: B

Given answer (D) is incorrect. Cannot have 60 2-core machines after subtracting the 30% buffer. Answer is B as it gives us the maximum amount of 2-core machines with equal ram.

upvoted 1 times

[Removed] 2 years, 6 months ago

@TT can you explain why D is incorrect?

I believe the question referenced Maximum which ideally should be D in this case.

upvoted 1 times

> **jiminycriminal** 2 years, 3 months ago
>
> We can only 84 cores to use after subtracting 30%. We cannot have 60 2-core machines. B is gives us the most 2-core machines with equal ram.
>
> upvoted 1 times
>
> **StudyBM** 2 years, 3 months ago
>
> D doesn't take into account a 30% buffer for the cores
>
> upvoted 1 times

Question #20 *Topic 1*

A company that utilizes an IaaS service provider has contracted with a vendor to perform a penetration test on its environment. The vendor is able to exploit the virtualization layer and obtain access to other instances within the cloud provider's environment that do not belong to the company. Which of the following BEST describes this attack?

A. VM escape

B. Directory traversal

C. Buffer overflow

D. Heap spraying

Suggested Answer: *A*

Community vote distribution

A (100%)

kuzummjakk 4 months ago

Selected Answer: A

lol infrastructure as a service service provider

upvoted 2 times

betty_boop 1 year, 2 months ago

Selected Answer: A

A. VM Escape

upvoted 2 times

AustinKelleyNet 1 year, 5 months ago

Selected Answer: A

A is correct

upvoted 2 times

Question #21 *Topic 1*

A systems administrator is troubleshooting network throughput issues following a deployment. The network is currently being overwhelmed by the amount of traffic between the database and the web servers in the environment. Which of the following should the administrator do to resolve this issue?

A. Set up affinity rules to keep web and database servers on the same hypervisor.

B. Enable jumbo frames on the gateway.

C. Move the web and database servers onto the same VXLAN.

D. Move the servers onto thick-provisioned storage.

Suggested Answer: *A*

Community vote distribution

A (57%) B (24%) C (19%)

TheGinjaNinja Highly Voted 1 year, 11 months ago

C. Move the web and database servers onto the same VXLAN.

VXLAN (Virtual Extensible LAN) is a network virtualization technology that allows for the creation of virtual networks on top of physical networks. By moving the web and database servers onto the same VXLAN, the systems administrator can create a logical segmentation of the network, reducing the amount of traffic between the servers, and thus reducing the load on the network.

A. Setting up affinity rules to keep the web and database servers on the same hypervisor will not help in resolving the issue as the traffic is between the servers not inside the hypervisor.

B. Enabling jumbo frames on the gateway would help to improve network throughput by allowing for larger packets to be sent over the network, but this would not address the issue of overwhelming traffic between the web and database servers.

D. Moving the servers onto thick-provisioned storage would not help in resolving the network throughput issues as it is a storage related concern not a network related one.

upvoted 14 times

reto1 3 months, 2 weeks ago

VXLAN (Virtual Extensible LAN): VXLAN is a network virtualization technology that encapsulates Layer 2 Ethernet frames in Layer 4 UDP packets. By moving the web and database servers onto the same VXLAN, you can effectively segment the traffic within a virtual network, which can help in optimizing and isolating traffic between these servers. This can improve performance and reduce network congestion.

upvoted 1 times

kuzummjakk 10 months ago

Your approach is reducing the amount of traffic, but the question says the traffic is "between" the web server and database, so as long as they're talking to each-other, it's the same amount of traffic. Even though the traffic doesn't "originate" from the hypervisor, it crosses the hypervisor since these are presumably virtual machines.

upvoted 1 times

Kobigasi 4 months, 1 week ago

except switched traffic is wayyyy faster than routed traffic. Since one of the answers says move them to the same VLAN, we can assume they are NOT on the same vlan. Putting them on the same vlan would make traffic flow much faster.

upvoted 1 times

Nordeen23 Most Recent 1 month, 2 weeks ago

C is the correct answer.

upvoted 1 times

Alvin_L 7 months ago

Selected Answer: C

Creates a logical network segment on top of a physical network. By placing the web servers and database servers on the same VXLAN, their communication can be encapsulated and tunneled within the VXLAN, effectively isolating their traffic from the rest of the network. This reduces congestion on the overall network and improves throughput.

upvoted 2 times

TheFivePips 7 months, 3 weeks ago

Selected Answer: C

Affinity rules ensure that specific virtual machines (VMs) run on the same physical hypervisor host. While this might improve communication between the web and database servers by reducing network latency, it does not address the underlying network throughput issue.

Jumbo frames increase the Maximum Transmission Unit (MTU) size, which can reduce overhead and improve network efficiency. However, enabling jumbo frames requires support from all network devices in the path, including switches and routers. Additionally, it may not directly address the specific traffic congestion between the database and web servers.

VXLAN (Virtual Extensible LAN) is a network virtualization technology that allows for the creation of logical networks overlaying existing Layer 3 infrastructure. Moving the web and database servers onto the same VXLAN can optimize network traffic by reducing hops and improving communication efficiency.

Of these options C is the most optimized solution. A and B help to reduce latency across the board. But only C is specific enough to optimize where it needs it most. D is completely irrelevant.

upvoted 2 times

TheFivePips 7 months, 3 weeks ago

This is a bad question BTW. Three of the answers provide basic optimization, but none actually really seem to provide more throughput or reduce the actual load. I hate it here

upvoted 3 times

kuzummjakk 10 months ago

Selected Answer: A

"amount of traffic between the database and the web servers in the environment". Amount of traffic from one server to another, in "the" environment. Insinuating they're in the same environment so no gateway so not B.

upvoted 1 times

kuzummjakk 10 months ago

A is more related to latency. Less latency might improve throughput, but the more relevant answer is B since it directly addresses throughput.

upvoted 1 times

kuzummjakk 10 months ago

i forgot to vote

upvoted 1 times

kuzummjakk 10 months ago

hold on, the question says they're troubleshooting throughput, then it says there's an issue with the "amount" of traffic overwhelming the network between the server and database. A sounds more correct then, since then that cuts out the network bottleneck.

upvoted 1 times

kuzummjakk 10 months ago

and jumbo frames "on the gateway" only addresses traffic leaving or entering ...this is traffic WITHIN the cloud provider since it's presumably two servers on the same cloud provider *facepalm* I recede my original answer

upvoted 1 times

utied 1 year ago

Selected Answer: A

The question is about solving the problem of too much traffic between database and web server.

A. Would solve this by grouping the database and webserver on same hypervisor.

B. Jumbo frames would solve this by reducing packets, but there is no indication that packet size is the problem.

C. vxLAN would solve this if there was 'other' traffic that was creating the congestion. vxLAN solves the problem of only 4096 possible network segments(mainly for CSPs). It also can create tunneling, so remote networks act as local and has traffic mirroring. None of these solutions directly correct this questions underlying problem.

D. thick provision solves availability of storage with increased pricing. Not a storage problem.

upvoted 2 times

Kobigasi 4 months, 1 week ago

Do you actually do this for a living? putting them on the same hypervisor solves nothing since the VMs themselves can still be in separate vlans. All that traffic is still hitting the router. You need to have them in the same vlan since switched frames are wayyyy faster than routed packets.

upvoted 1 times

yyCherubim 1 year, 1 month ago

Jumbo Frames Really!? That's the best CompTIA can come up with. Guess I'm going to get this question wrong, because I'm not going with stupid Jumbo Frames!

upvoted 2 times

> **FrancisDrake** 11 months, 1 week ago
>
> Why...?
>
> upvoted 1 times

backdooranon 1 year, 2 months ago

Selected Answer: B

In some circumstances, using jumbo frames can result in better performance, while in others it can lower performance.In controlled networks with high utilization, using jumbo frames can lead to improved network throughput due to reduced overhead.
https://www.techtarget.com/searchnetworking/definition/jumbo-frames

upvoted 1 times

Pongsathorn 1 year, 3 months ago

Selected Answer: B

The correct option to resolve network throughput issues caused by overwhelming traffic between the database and web servers is:

B. Enable jumbo frames on the gateway.

Enabling jumbo frames increases the Maximum Transmission Unit (MTU) size, allowing larger packets to be transmitted over the network. This can reduce the overhead associated with transmitting smaller packets and potentially improve network throughput, especially in situations with high volumes of traffic between servers like databases and web servers.

Options A, C, and D are not directly related to optimizing network throughput and may not address the specific issue of network congestion caused by the high traffic between the database and web servers.

upvoted 1 times

Tomtom11 1 year, 4 months ago

Selected Answer: B

I would choose Answer B
"The network is currently being overwhelmed by the amount of traffic between the database and the web servers in the environment."

It is often more efficient to use a larger Ethernet frame size than the standard Ethernet
MTU inside the data center to reduce networking overhead. Jumbo frames allow for higher
network performance by reducing the overhead in each Ethernet frame by using fewer but
larger frames. Jumbo frames also reduce the number of times that a CPU will be interrupted
to process Ethernet traffic since each jumbo frame can be up to six times as large as a standard
frame.

upvoted 1 times

SecPlus2022 1 year, 6 months ago

Selected Answer: A

It's not likely that the database and the web servers referenced are on different networks therefor none of the traffic would be traversing through a gateway.

upvoted 2 times

> **FrancisDrake** 11 months, 1 week ago
>
> A good point.
>
> upvoted 2 times
>
> > **FrancisDrake** 11 months, 1 week ago
> >
> > But the scenario reads as if they are separated by a gateway. Hmmm...
> >
> > upvoted 2 times
> >
> > > **kuzummjakk** 10 months ago
> > >
> > > "...the database and web server in the environment."
> > >
> > > upvoted 1 times

Slambang 1 year, 6 months ago

Selected Answer: B

Given the options provided, the most appropriate choice for resolving network throughput issues would be B. Enable jumbo frames on the gateway.

upvoted 2 times

Zak11 1 year, 8 months ago

Selected Answer: A

Affinity rules will ensure that the web and database servers are placed on the same physical host, reducing the network traffic and improving throughput by allowing for communication over the faster internal network.

upvoted 1 times

samuel186 1 year, 8 months ago

Selected Answer: A

Setting up affinity rules to keep web and database servers on the same hypervisor is the BEST option to resolve this issue. Affinity rules will ensure that the web and database servers are located on the same physical host, which will help reduce network latency and improve throughput. This will also help to improve performance as the network traffic will not have to traverse multiple hypervisors. Enabling jumbo frames on the gateway, moving the web and database servers onto the same VXLAN, and moving the servers onto thick-provisioned storage are not likely to resolve network throughput issues between web and database servers.

upvoted 1 times

betty_boop 1 year, 8 months ago

Selected Answer: C

C. Move the web and database servers onto the same VXLAN

upvoted 2 times

AustinKelleyNet 1 year, 11 months ago

Selected Answer: A

I agree with NOT_THAT_GUY

upvoted 1 times

Question #22 *Topic 1*

An update is being deployed to a web application, and a systems administrator notices the cloud SQL database has stopped running. The VM is responding to pings, and there were not any configuration changes scheduled for the VM. Which of the following should the administrator check NEXT?

A. Logs on the VM

B. Firewall on the VM

C. Memory on the VM

D. vGPU performance on the VM

Suggested Answer: *A*

Community vote distribution

A (100%)

CXSSP 1 month, 1 week ago

Selected Answer: A

Logs on the VM typically contain detailed information about system events, errors, and application behavior. By examining the logs, the administrator can gain valuable insights into why the cloud SQL database stopped running. The logs might reveal error messages related to the database service itself, resource limitations, or any unexpected events that caused the crash. Analyzing these logs should provide clues to diagnose the root cause of the database outage and help the administrator determine the appropriate recovery steps.

upvoted 3 times

TheFivePips 1 month, 2 weeks ago

Selected Answer: A

This question sucks because I doesn't feel like there is enough information to really give a decent answer. That being said, when we need more information we should probably check the logs. That's my best guest. Its probably not a firewall issue because it responds to pings (doesn't rule it out tho). Probably not a memory issue because that wouldn't stop the Db from running unless it was a huge memory issue which would probably also stop pings(doesn't rule it out tho). vGPU is also not likely for similar reasons (doesn't rule it out tho). When in doubt, check the logs

upvoted 4 times

kuzummjakk 4 months ago

While it COULD be C, B, or even D, the question insinuates it stopped working specifically while an update is being pushed that "shouldn't" have touched it. Since it shouldn't, and clearly is "somehow", more information is required.

upvoted 2 times

Question #23 *Topic 1*

A company is concerned about the security of its data repository that contains customer PII. A systems administrator is asked to deploy a security control that will prevent the exfiltration of such data. Which of the following should the systems administrator implement?

A. DLP

B. WAF

C. FIM

D. ADC

Suggested Answer: *A*

Community vote distribution

A (100%)

kuzummjakk 4 months ago

too obvious im scared

upvoted 3 times

Zak11 1 year, 2 months ago

Selected Answer: A

DLP can monitor and control sensitive data leaving the network, whether it's through email, file transfer, or other network traffic. It can also detect and block unauthorized access to sensitive data by users within the network.

upvoted 2 times

betty_boop 1 year, 2 months ago

Selected Answer: A

A. DLP

upvoted 1 times

AustinKelleyNet 1 year, 5 months ago

Selected Answer: A

A is correct

upvoted 1 times

Question #24 *Topic 1*

An engineer is responsible for configuring a new firewall solution that will be deployed in a new public cloud environment. All traffic must pass through the firewall.
The SLA for the firewall is 99.999%. Which of the following should be deployed?

A. Two load balancers behind a single firewall

B. Firewalls in a blue-green configuration

C. Two firewalls in a HA configuration

D. A web application firewall

Suggested Answer: *C*

Community vote distribution

C (100%)

54a6b25 5 months, 1 week ago

C. Two firewalls in a HA configuration

Explanation:
To achieve a high SLA of 99.999% (often referred to as "five nines" uptime), you need a highly available and redundant setup. Deploying two firewalls in a High Availability (HA) configuration ensures that if one firewall fails, the other can take over without disrupting traffic, thereby maintaining continuous service.

upvoted 3 times

kuzummjakk 10 months ago

Selected Answer: C

LOL two load balancers for one firewall. if you see that on a network, run. blue/green can help maintain HA but only during an update.

upvoted 2 times

Zak11 1 year, 8 months ago

Selected Answer: C

The other options, such as load balancers, blue-green configuration, or a web application firewall, do not provide the same level of redundancy and failover capabilities as a HA firewall configuration.

upvoted 2 times

betty_boop 1 year, 8 months ago

Selected Answer: C

C. Two firewalls in a HA configuration

upvoted 1 times

ramrod1738 1 year, 11 months ago

C. Two firewalls in a HA (High Availability) configuration is the best option to achieve the desired SLA of 99.999%. With a HA configuration, two firewalls are deployed and configured to work in tandem, so that if one firewall fails, the other will take over seamlessly and ensure continuous operation with minimal downtime. This helps to achieve high availability and improve reliability, which is critical for meeting the demanding SLA requirement.

upvoted 2 times

CapJackSparrow 1 year, 11 months ago

High availability (HA) allows you to place two firewalls in a group and synchronize their configuration. This prevents a single point of failure on your network. The two firewalls have a heartbeat connection, which ensures failover if one of the firewalls goes down

upvoted 4 times

Question #25 *Topic 1*

After a hardware upgrade on a private cloud system, the systems administrator notices a considerable drop in network performance. Which of the following is
MOST likely the cause?

A. The driver

B. The memory

C. The cluster placement

D. The CPU

Suggested Answer: *A*

Community vote distribution

A (65%) | C (30%) | 5%

54a6b25 5 months, 3 weeks ago

The network driver is the most likely culprit for a drop in network performance following a hardware upgrade. Ensuring that the correct driver is installed and properly configured is essential to restore optimal network performance.

A. The driver: A hardware upgrade often requires updated drivers to function correctly. If the network driver is outdated, incompatible, or improperly configured, it can lead to significant drops in network performance. Ensuring the driver is up-to-date and correctly configured is crucial for optimal hardware performance.

upvoted 1 times

TheFivePips 7 months, 3 weeks ago

Selected Answer: A

If the only thing being changed in this scenario is the hardware, then it stands to reason that the memory and CPU might be improved, and the cluster placement should remain the same if the configuration is the same. It makes no mention either way so we can assume no change. The only thing that could be detrimental in this scenario is the driver. The software that tells the computer how to use the new hardware. This is a common problem when getting new hardware, making it a very likely candidate.

upvoted 2 times

kuzummjakk 10 months ago

Selected Answer: C

In cloud environments, when the hardware your VM is on needs to get replaced, they don't just take down your VM until the hardware is replaced then put it back on, they migrate it to another server "seamlessly" when they decommission that hardware (imagine their availability rating yikes). Drivers CAN affect performance, just as the CPU and Memory CAN affect performance, but it's assuming a bit too much out of the information of "the cloud vendor upgraded the hardware".

upvoted 1 times

TheFivePips 7 months, 3 weeks ago

I think its assuming too much to say the cluster placement is changed, when nowhere in the question is that stated. Also typically when doing an upgrade you expect performance to increase in key areas such as CPU and Memory. If either of those are doing worse, its because they are either not installed correctly or THE DRIVER IS WRONG

upvoted 1 times

jwicky 10 months, 2 weeks ago

Answer C: Don't overthink the question. It doesn't mention anything about upgrading one device. It is upgrading the private cloud system. Cluster placements often effect
performance and is used to increase performance. If it is used to increase performance a change could decrease it as well.

upvoted 3 times

reto1 3 months, 2 weeks ago

where does it say it's in a cluster? It could be just one host

upvoted 1 times

AllenTaylor 11 months, 3 weeks ago

A is the answer.

upvoted 1 times

FrancisDrake 1 year ago

It's not clear what hardware was upgraded but it is common to update drivers with new hardware installations. I'm going with driver.

upvoted 2 times

Frogman1981 1 year ago

D. The CPU: The CPU is the brain of the system and can significantly impact network performance. If the new CPU is not compatible with the existing network infrastructure or if it is not configured properly, it can cause a considerable drop in network performance. This is the most likely cause of the considerable drop in network performance after a hardware upgrade.

In conclusion, option D (the CPU) is the most likely cause of the considerable drop in network performance after a hardware upgrade on a private cloud system.

upvoted 1 times

kuzummjakk 10 months ago

"Can", but so can memory, so can the driver, so can cluster placement. Any one of these things can be a bottleneck for the network.

upvoted 1 times

yyCherubim 1 year, 1 month ago

Selected Answer: A

There is not enough information in the question to conclude it's a "Cluster" problem. Did they upgrade a server in the cluster, or the sound card on the public affairs officer's workstation so they could listed blast music at the next employee social?

upvoted 1 times

SecPlus2022 1 year, 6 months ago

Selected Answer: D

After further thought, need to change my answer. At first I was positive the answer was the driver, but drivers are not part of the CV0-003 objectives. Cluster placement can affect performance, but the most likely cause is incorrect configuration of the new hardware. The number 1 cause of network performance issues is high cpu usage. Therefore, I'm going with "D".

upvoted 1 times

kuzummjakk 10 months ago

I think D assumes too much information out of the question imo

upvoted 1 times

SecPlus2022 1 year, 6 months ago

Selected Answer: A

Exactly what "concepcionz" states.

upvoted 2 times

betty_boop 1 year, 8 months ago

Selected Answer: A

A. The driver

upvoted 2 times

concepcionz 1 year, 9 months ago

Selected Answer: A

When hardware is upgraded, it is common for drivers to be updated to be compatible with the new hardware.

upvoted 4 times

ramrod1738 1 year, 11 months ago

A. The driver could be the cause of the drop in network performance after the hardware upgrade. Drivers are software components that enable communication between the operating system and the hardware devices in a computer. If the new hardware is not compatible with the current driver, it could lead to performance issues, including slow network performance. Upgrading or updating the driver to a version that is compatible with the new hardware could resolve the issue. The other options, such as memory, cluster placement, and CPU, could also potentially impact performance, but the issue with the driver is the most likely cause based on the description provided.

upvoted 2 times

beamage 1 year, 11 months ago

Selected Answer: C

Nobody said servers, private cloud hardware upgrade, that could be switch's, routers, firewalls and so on.

upvoted 2 times

⊟ **scott5010** 2 years, 1 month ago

Selected Answer: C

According to my study guide, AWS attempts to spread out new instances across hardware, if that is the case, a bad change would not have a huge impact unless the cluster placement is wrong or concentrated. I vote for cluster placement

upvoted 3 times

⊟ **Brianhealey136** 2 years ago

Considering they said private cloud comparing it to AWS (a public cloud) makes no sense.

Answer is A

upvoted 1 times

⊟ **JVen** 2 years, 1 month ago

Selected Answer: A

A makes the most sense

upvoted 1 times

⊟ **Agr321** 2 years, 3 months ago

They key word is hardware upgrade. Answer is Driver

upvoted 2 times

Question #26 *Topic 1*

A systems administrator is trying to reduce storage consumption. Which of the following file types would benefit the MOST from compression?

A. System files

B. User backups

C. Relational database

D. Mail database

Suggested Answer: *B*

Community vote distribution

B (60%) C (40%)

⊟ **54a6b25** 5 months, 3 weeks ago

B. User backups: User backups typically contain a variety of file types, many of which (such as text documents, spreadsheets, and other office files) can be highly compressible. Compressing these backups can lead to significant storage savings.
User backups, with their diverse and often highly compressible content, stand to gain the most in terms of reduced storage consumption through compression.

upvoted 2 times

⊟ **TheFivePips** 7 months, 3 weeks ago

Selected Answer: C

The effectiveness of compression depends on the nature of the data being compressed. Generally, file types that contain repetitive patterns or redundant information benefit the most from compression because they can be significantly reduced in size.

Out of all of the options only the relational database has data that is FOR SURE repeated (because its relational).

upvoted 1 times

⊟ **kuzummjakk** 10 months ago

Selected Answer: B

ChatGPT can get stuff wrong btw. Got my net+ study material wrong good thing i caught it.

backups add up. that stuff gets exponential <----------

upvoted 1 times

⊟ **utied** 1 year ago

Selected Answer: B

When you compress data, you trade physical bit size for compute. Small disk size, but more compute to read the data. Typically you only compress stuff that will be in storage and not read that much. A,C,D are all in use, so you would increase compute for the smaller size benefit.

upvoted 2 times

⊟ **TheFivePips** 7 months, 3 weeks ago

This has nothing to do with what the question is asking. What would benefit most from compression means what will compress the best. You're reading too much into the question

upvoted 1 times

⊟ **anonymonkey** 1 year, 1 month ago

C saves over 20-80%
B saves 20-70%

Ultimately both user backups and relational databases can benefit from compression, but the decision should be based on specific requirements, such as storage efficiency, performance considerations, and resource utilization in regards which method would yield the greatest overall return.

upvoted 2 times

⊟ **yyCherubim** 1 year, 1 month ago

Selected Answer: B

I first though D because there is TONS of dead space in email data. However, B also has tons of dead space, and it's not used.

upvoted 1 times

nmap_king_22 1 year, 3 months ago

Selected Answer: C

The file type that would benefit the MOST from compression among the options provided is:

C. Relational database

Relational databases often contain structured data with repetitive patterns, making them well-suited for compression. Compressing a relational database can significantly reduce storage consumption without sacrificing data integrity or performance.

upvoted 3 times

NRJ6425 1 year, 7 months ago

Selected Answer: B

User Backups

upvoted 1 times

betty_boop 1 year, 8 months ago

Selected Answer: B

B. User backups

upvoted 1 times

Question #27 *Topic 1*

A technician just received the lessons learned from some recent data that was lost due to an on-premises file-server crash. The action point is to change the backup strategy to minimize manual intervention. Which of the following is the BEST approach for the technician to implement?

A. Backup as a service

B. RAID 1

C. Long-term storage

D. New backup devices

Suggested Answer: *A*

Community vote distribution

A (100%)

54a6b25 5 months, 3 weeks ago

A. Backup as a service: This approach automates the backup process, reducing the need for manual intervention. It typically includes regular, scheduled backups, offsite storage, and sometimes additional features like disaster recovery. This ensures data is consistently backed up without relying on manual processes.

upvoted 1 times

Alvin_L 7 months ago

Selected Answer: A

This is from Ghat GPT lol

A. Backup as a service (BaaS):

Automation: BaaS solutions are designed to automate the backup process, reducing the need for manual intervention. This ensures regular and consistent backups without the technician needing to manage the process actively.

Off-site Storage: BaaS typically involves storing data in the cloud, which protects against on-premises disasters such as server crashes.

Scalability and Reliability: Cloud-based backup services offer scalable storage and high reliability, often with built-in redundancy and robust recovery options.

Management and Monitoring: These services provide management and monitoring tools that alert administrators to any issues, ensuring that backups are completed successfully.

upvoted 1 times

FrankyD92 10 months, 3 weeks ago

Is this ChatGPT answering these questions or something? A is the obvious answer in my head. not sure how "Long-term Storage" fixes a backup issue

upvoted 3 times

> **kuzummjakk** 10 months ago
>
> LOL i was thinking the same thing
>
> upvoted 1 times

AllenTaylor 11 months, 3 weeks ago

A makes the most sense.

upvoted 2 times

Pongsathorn 1 year, 3 months ago

Selected Answer: A

The BEST approach for the technician to implement in order to minimize manual intervention in the backup strategy is:

A. Backup as a service.

Backup as a service (BaaS) typically involves using cloud-based backup solutions that are managed by third-party providers. These services automate the backup process, provide redundancy, and often include features like regular snapshots, retention policies, and disaster recovery capabilities. BaaS can significantly reduce manual intervention and enhance data protection compared to traditional on-premises backup solutions.

Options B, C, and D do not directly address the need to minimize manual intervention in the backup process, and they may not provide the same level of automation and convenience as a BaaS solution. RAID 1 is a storage redundancy technique, long-term storage refers to archiving data, and getting new backup devices wouldn't necessarily reduce manual intervention unless the devices are part of a more automated backup strategy.

upvoted 3 times

betty_boop 1 year, 8 months ago

Selected Answer: A

A. Backup as a service

upvoted 2 times

AustinKelleyNet 1 year, 11 months ago

Selected Answer: A

A all the way

upvoted 2 times

Alizadeh 1 year, 11 months ago

Selected Answer: A

A. Backup as a Service

Backup as a Service (BaaS) is a cloud-based service that allows companies to remotely back up and store their data. It generally provides automatic and scheduled backups that are done remotely, minimizing the need for manual intervention. BaaS can also provide easy access to data restoration and disaster recovery, allowing the technician to quickly and easily restore data in case of a crash or other disaster.

upvoted 2 times

JVen 2 years, 1 month ago

Selected Answer: A

A is the only one that makes sense

upvoted 2 times

scott5010 2 years, 1 month ago

Selected Answer: A

agree, it looks like A

upvoted 2 times

JVen 2 years, 1 month ago

Selected Answer: A

Specifically states no manual intervention, has to be A

upvoted 2 times

Admiral_Crunch 2 years, 6 months ago

A, the best solution to reduce manual intervention is Backup as a Service, think of it as Drop box or One drive

upvoted 3 times

SimplyDebonair 2 years, 9 months ago

o A would be the correct answer due to the specifications outlined in the scenario: "…minimize manual intervention." Backup-as-a-Service (BaaS) provides offsite storage for files, folders, and/or the entire hard drive that is regularly backed up by a secure, remote cloud-based repository over a network connection. This minimizes the risk of loss by user error, hacking, or any other type of technological disaster. See the link provided: https://www.netapp.com/cloud-services/cloud-backup/what-is-backup-as-a-services-baas/

upvoted 3 times

martin451 2 years, 11 months ago

I agree the answer is A Backup as a Service (BaaS) offers businesses a flexible backup solution that reduces their time and effort spent on taking manual data backup

upvoted 4 times

dvd21 3 years, 1 month ago

Minimize manual intervention. Answer is A.

upvoted 2 times

Question #28 *Topic 1*

A marketing team is using a SaaS-based service to send emails to large groups of potential customers. The internally managed CRM system is configured to generate a list of target customers automatically on a weekly basis, and then use that list to send emails to each customer as part of a marketing campaign. Last week, the first email campaign sent emails successfully to 3,000 potential customers. This week, the email campaign attempted to send out 50,000 emails, but only 10,000 were sent. Which of the following is the MOST likely reason for not sending all the emails?

A. API request limit

B. Incorrect billing account

C. Misconfigured auto-scaling

D. Bandwidth limitation

Suggested Answer: *A*

Community vote distribution

A (100%)

54a6b25 5 months, 3 weeks ago

A. API request limit: SaaS-based email services often have limits on the number of API requests that can be made within a certain time frame. If the email campaign exceeded the API request limit, the service would stop processing additional requests until the limit resets. This is a common issue when scaling up the number of emails sent significantly.

upvoted 1 times

kuzummjakk 10 months ago

Selected Answer: A

Flat 10,000 is too coincidental for a bandwidth issue.

upvoted 1 times

Pongsathorn 1 year, 3 months ago

Selected Answer: A

The MOST likely reason for not sending all the emails in this scenario is:

A. API request limit.

When using a SaaS-based service to send emails in an automated manner, such services often have API rate limits in place to prevent abuse or overuse of their infrastructure. These rate limits restrict the number of API requests (in this case, sending emails) that can be made within a certain timeframe. If the CRM system attempted to send 50,000 emails in a short period and hit the API request limit, it would result in only a portion of the emails being sent, which aligns with the described issue.

Options B, C, and D are less likely to be the cause of the problem. An incorrect billing account typically doesn't affect the number of emails that can be sent. Misconfigured auto-scaling and bandwidth limitations are more related to infrastructure scaling and network performance, which might not be the primary issue when dealing with email API requests.

upvoted 1 times

betty_boop 1 year, 8 months ago

Selected Answer: A

A. API request limit

upvoted 1 times

rob88Silva 1 year, 11 months ago

Selected Answer: A

API request limit is the MOST likely reason for not sending all the emails. Many SaaS-based services have API request limits which determine the number of requests that can be made to their systems within a certain period of time. If the marketing team exceeded this limit, the service may have stopped processing the remaining requests, resulting in only a portion of the emails being sent. This can be confirmed by checking the documentation of the SaaS service or by contacting their customer support.

upvoted 3 times

Question #29 *Topic 1*

A developer is no longer able to access a public cloud API deployment, which was working ten minutes prior. Which of the following is MOST likely the cause?

A. API provider rate limiting

B. Invalid API token

C. Depleted network bandwidth

D. Invalid API request

Suggested Answer: *B*

Community vote distribution

B (94%) 6%

SimplyDebonair Highly Voted 2 years, 9 months ago

Selected Answer: B

The correct answer is “B.” An API token is necessary before the API request can be sent. If the token is no longer working, the API request wouldn’t be submitted.

upvoted 9 times

54a6b25 Most Recent 5 months, 3 weeks ago

B. Invalid API token: API tokens can expire or become invalid due to several reasons (e.g., token expiration, token revocation for security reasons, or misconfiguration). If the token expired or was revoked, the developer would lose access suddenly, which fits the scenario described.

upvoted 1 times

kuzummjakk 10 months ago

Selected Answer: B

Who's going out here assigning ChatGPT answers to these questions lol.
Cannot access "the api", so the entire api. It assumes too much to say "only one request they made was wrong".

upvoted 1 times

AllenTaylor 11 months, 3 weeks ago

Selected Answer: A

I would say rate limit. For example, with twitter you can get 50 tweets within a 15 minute time period, additional requests will be denied until the 15 minutes pass.

upvoted 1 times

Pongsathorn 1 year, 3 months ago

Selected Answer: B

The MOST likely cause, given the information provided, is:

B. Invalid API token.

If a developer was able to access a public cloud API and suddenly loses access, one common reason for this is an invalid API token. API tokens often have a limited duration of validity, and if the token expires or becomes invalid for some reason, the developer won't be able to access the API anymore.

While the other options (rate limiting, depleted network bandwidth, and invalid API requests) are also potential causes of API access issues, the sudden loss of access in this scenario is more likely related to an authentication issue, such as an expired or invalid API token.

upvoted 4 times

AustinKelleyNet 1 year, 11 months ago

Selected Answer: B

B all the way

upvoted 1 times

ramrod1738 1 year, 11 months ago

B. Invalid API token.

An API token is typically used to authenticate the client making API requests to the cloud API. If the token becomes invalid, either because it has expired or because it has been revoked, the client will no longer be able to access the API. This can result in the developer being unable to access the API deployment, even if it was working previously.

API provider rate limiting, depleted network bandwidth, and invalid API request can also cause issues with accessing an API deployment, but are less likely to be the cause if the API was working previously and suddenly stopped. These issues are more likely to result in a slow or inconsistent response from the API, rather than a complete inability to access it.

upvoted 1 times

Agr321 2 years, 3 months ago

I'm going to go with A: API limit
The question ask about a PUBLIC API, which I don't think a session token would be used to authenticate with.....

upvoted 1 times

Agr321 2 years, 3 months ago

Answer is B API token.

upvoted 1 times

ryanzou 2 years, 3 months ago

Selected Answer: B

It's B, token is refreshed

upvoted 1 times

Rob69420 2 years, 3 months ago

Rate limiting is a strategy for limiting network traffic. It puts a cap on how often someone can repeat an action within a certain timeframe – for instance, trying to log in to an account. Rate limiting can help stop certain kinds of malicious bot activity. It can also reduce strain on web servers.

Rate limiting protects an API by applying a hard limit on its access.

upvoted 2 times

Michell999 2 years, 6 months ago

Why not A?

upvoted 4 times

reto1 3 months, 2 weeks ago

Yeah GPT says A.

upvoted 1 times

Agr321 2 years, 3 months ago

I was originally thinking A because it said Public, thinking it was referring to a PUblic API. Which I dismissed the API token because public API didn't use authentication.

I re-read the question. It is a API provided from a cloud provider. Public cloud providers are AWS,Azure, etc. you would need to authenticate into the public cloud to use the API.

Hope you understand the logic.

upvoted 2 times

martin451 2 years, 11 months ago

answer should be B

upvoted 2 times

dvd21 3 years, 1 month ago

What a shit question. I lean B.

upvoted 2 times

Question #30 *Topic 1*

A support engineer wants to prevent users from running malware on several IaaS compute instances. Which of the following will BEST achieve this objective?

A. Encrypt all applications that users should not access.

B. Set the execute filesystem permissions on the desired applications only.

C. Implement an application whitelisting policy.

D. Disable file sharing on the instance.

Suggested Answer: *C*

Community vote distribution

C (100%)

Zak11 Highly Voted 1 year, 8 months ago

Selected Answer: C

Whitelisting policies only allow approved applications to run, which would effectively prevent any unauthorized application from executing, including malware.

Option A is not a good solution to prevent users from running malware, as encryption is not designed to prevent users from running malware but to protect data from unauthorized access.

Option B can help to limit which applications can run, but it won't prevent users from running malware if the malware is disguised as an approved application.

Option D won't be a good solution either as it only restricts access to shared files, not the execution of applications.

upvoted 6 times

54a6b25 Most Recent 5 months, 3 weeks ago

C. Implement an application whitelisting policy: Application whitelisting allows only approved and trusted applications to run on the compute instances. This is a proactive approach to security that prevents users from executing any software that is not explicitly permitted, effectively blocking malware from running.

upvoted 1 times

Question #31 *Topic 1*

A cloud administrator is reviewing the annual contracts for all hosted solutions. Upon review of the contract for the hosted mail solution, the administrator notes the monthly subscription rate has increased every year. The provider has been in place for ten years, and there is a large amount of data being hosted. Which of the following is a barrier to switching providers?

A. Service-level agreement

B. Vendor lock-in

C. Memorandum of understanding

D. Encrypted data

Suggested Answer: *B*

Community vote distribution

B (100%)

⊟ **54a6b25** 5 months, 3 weeks ago

B. Vendor lock-in: Vendor lock-in occurs when a customer becomes dependent on a single provider due to the high costs, complexity, or technical challenges associated with moving to a different provider. In this case, the large amount of data being hosted and the historical reliance on the provider make it difficult and potentially costly to switch to a new provider.

upvoted 3 times

⊟ **AustinKelleyNet** 1 year, 11 months ago

Selected Answer: B

This should be obvious

upvoted 2 times

Question #32

Topic 1

A systems administrator is creating a VM and wants to ensure disk space is not allocated to the VM until it is needed. Which of the following techniques should the administrator use to ensure this?

A. Deduplication

B. Thin provisioning

C. Software-defined storage

D. iSCSI storage

Suggested Answer: *B*

Community vote distribution

B (100%)

54a6b25 5 months, 3 weeks ago

B. Thin provisioning: Thin provisioning allows disk space to be allocated to a VM on-demand as data is written, rather than reserving all the specified disk space upfront. This ensures that disk space is only used as needed, optimizing storage utilization.

upvoted 1 times

Pongsathorn 1 year, 3 months ago

Thin provisioning: Specify an initial quantity of available storage capacity, and then set a maximum storage size. The storage grows dynamically up to that maximum size, consuming only what it needs.

Thick provisioning: Specify a single quantity of storage capacity that is then reserved for the instance, whether needed or not.

Cloud bursting: Configure supplementary public cloud storage for when the organization's private cloud storage is full.

upvoted 2 times

AustinKelleyNet 1 year, 11 months ago

Selected Answer: B

Thin provisioning= Using resource when needed

Thick provisioning = constantly reserving resource

upvoted 2 times

Question #33 Topic 1

After accidentally uploading a password for an IAM user in plain text, which of the following should a cloud administrator do FIRST? (Choose two.)

A. Identify the resources that are accessible to the affected IAM user.

B. Remove the published plain-text password.

C. Notify users that a data breach has occurred.

D. Change the affected IAM user's password.

E. Delete the affected IAM user.

Suggested Answer: *D*

Community vote distribution

D (91%) 9%

ironman_86 Highly Voted 1 year, 10 months ago

the answer is B & D

upvoted 9 times

Rob69420 Highly Voted 1 year, 10 months ago

Remove the password

Change the password

upvoted 6 times

PatrickH Most Recent 3 months, 3 weeks ago

Selected Answer: D

What to do FIRST. Removing the password is good but potentially the damage has been done. I would suggest A and D. Change the Password)id imagine everyone agrees on that one) and indentify what resources may be impacted by any breech, so proactive measures can be put in place.

upvoted 2 times

kuzummjakk 3 months, 4 weeks ago

Selected Answer: D

The question says FIRST. Think damage control. If you delete the password first, anyone who's seen it still has access to the account. The first step is denying access, not preventing more access.

upvoted 2 times

Locy333 4 months, 3 weeks ago

Selected Answer: D

Answer is D over B. If you do B first, anyone that saw it prior to deletion would have access to the account. If you do D first, the account is secure first, then you can do cleanup. D is more time sensitive.

upvoted 2 times

eacunha 6 months ago

Selected Answer: B

Após carregar acidentalmente uma senha para um usuário do IAM em texto simples, as ações que um administrador de nuvem deve realizar PRIMEIRO são:

1. **B. Remova a senha de texto simples publicada.**
- Isso ajuda a mitigar imediatamente a exposição da senha publicada inadvertidamente.

2. **D. Altere a senha do usuário IAM afetado.**
- Para garantir que a senha comprometida seja imediatamente invalidada e substituída por uma nova.

Portanto, as opções corretas são:
- **B. Remova a senha de texto simples publicada.**
- **D. Altere a senha do usuário IAM afetado.**

upvoted 1 times

sheilawu 11 months, 3 weeks ago

Selected Answer: D

this was on my test 7/2 and there was only one answer to choose, the right answer is D

upvoted 2 times

AustinKelleyNet 1 year, 5 months ago

Selected Answer: D

B & D are correct. I think they had a programming error with their buttons

upvoted 2 times

> **kuzummjakk** 3 months, 4 weeks ago
>
> The question says "first". You can't do both first this is deliberate.
>
> upvoted 1 times
>
> > **TheFivePips** 1 month, 1 week ago
> >
> > it also says to choose 2. this is deliberate
> >
> > upvoted 2 times

Question #34 *Topic 1*

A cloud administrator has deployed a new VM. The VM cannot access the Internet or the VMs on any other subnet. The administrator runs a network command and sees the following output:

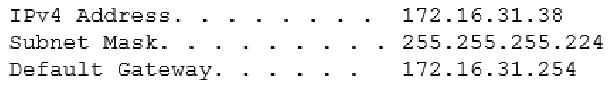

```
IPv4 Address. . . . . . . .  172.16.31.38
Subnet Mask. . . . . . . . . 255.255.255.224
Default Gateway. . . . . .   172.16.31.254
```

The new VM can access another VM at 172.16.31.39. The administrator has verified the IP address is correct. Which of the following is the MOST likely cause of the connectivity issue?

A. A missing static route

B. A duplicate IP on the network

C. Firewall issues

D. The wrong gateway

Suggested Answer: *D*

Community vote distribution

D (89%) 11%

i_bird Highly Voted 2 years, 3 months ago

I came to learn, not unlearn..how can it not be D?

upvoted 9 times

kuzummjakk 10 months ago

If the gateway is wrong, the traffic wouldn't leave the machine at all. Meaning it can't even talk to a machine on the same network like the question says. The gateway is the FIRST place the traffic hits.

upvoted 3 times

Kobigasi 4 months, 1 week ago

This is wrong. The question says ' VMs on other subnets'. The device can still talk on the same subnet, even if the gateway is wrong. It will use ARP, and the traffic will not hit the gateway first as you state.

upvoted 1 times

SimplyDebonair Highly Voted 2 years, 9 months ago

Selected Answer: D

The determining factor is the subnet mask. Plugging the IPv4 address into a IPv4/Subnet calculator, you come out with 30 usable hosts. The ranges for those hosts are 172.16.31.33 - 172.16.31.62. When you plug in the default gateway's IP, it also comes back with 30 usable hosts. However, the ranges for those hosts are 172.16.32.225 - 172.16.31.254, which is outside the range of the IPv4 address. Thus, the issue isn't a duplicated IP, but the gateway is wrong.

upvoted 8 times

kuzummjakk 10 months ago

technically the subnet mask can be wrong and traffic can still reach the gateway if the IP is correct, however if the gateway is wrong traffic won't leave the machine at all so talking to another machine on the same network aint happening like the question says. Duplicate IP still doesn't make sense though given the information in the question and how it can reach machines in its subnet just fine. Cloud vendors usually let you configure some sort of "route table" to control what traffic goes between what subnets and you're meant to configure the route table to send traffic that "falls through" as going to the "internet gateway" so by being the most related to "no internet access, but internal access", I think it's A.

upvoted 1 times

Kobigasi 4 months, 1 week ago

Actually traffic will leave the machine, even if the gateway is wrong. devices use ARP to talk to other devices in the same subnet and it doesn't hit the gateway at all.

That said I still believe the answer is D

upvoted 1 times

54a6b25 Most Recent 5 months, 3 weeks ago

D. The wrong gateway

Explanation:

Given the subnet mask of 255.255.255.224, the subnet range for the VM's IP address (172.16.31.38) would be 172.16.31.32 to 172.16.31.63, with 172.16.31.33 as the first usable address and 172.16.31.62 as the broadcast address. The default gateway for this subnet should be within this range, but typically, it would be one of the first or last usable addresses in the range, like 172.16.31.33 or 172.16.31.62.

The provided gateway, 172.16.31.254, does not fall within the 172.16.31.32/27 subnet range, which is why the VM cannot access the Internet or other subnets.

To resolve the issue, the correct gateway should be set to an address within the VM's subnet range. For example, if the standard practice in the network is to use the first usable address as the gateway, the gateway should be set to 172.16.31.33. If the last usable address is used, it should be 172.16.31.62.

upvoted 1 times

Alvin_L 7 months ago

Selected Answer: D

It's deinitely wrong gateway, this subnet is /27, should only has 32 devices, 254 is way out of range

upvoted 1 times

kuzummjakk 10 months ago

Selected Answer: A

If it can't reach the gateway, why can it reach that other machine? In the cloud, the network often has to be manually configured to allow access outside of the subnet to the internet so my bet's on A.

upvoted 2 times

AllenTaylor 11 months, 3 weeks ago

Selected Answer: D

It is D. Why is B shown as the answer?

upvoted 1 times

gijack88 1 year ago

Selected Answer: D

The Answer is D.

IPv4 address 172.16.31.38
with a subnet 255.255.255.224
has a network address of 172.16.31.32 and a broadcast of 172.16.31.63
this means the gateway can only fall between 172.16.31.33 - 172.16.31.62.

The gateway listed of 172.16.31.254 is not in the same subnet, which means it cannot reach the gateway necessary to reach other networks.

upvoted 1 times

yyCherubim 1 year, 1 month ago

Selected Answer: D

Answer "B" is more of CompTIA Stupid!
If you have a duplicate IP on the network, your not going to talk to anything.

upvoted 2 times

sheilawu 1 year, 5 months ago

Selected Answer: D

Was on my test 7/2 , answer should be D

upvoted 1 times

Zak11 1 year, 8 months ago

Selected Answer: D

The most likely cause of the connectivity issue is the wrong gateway. The solution is to correct the default gateway to the correct value, which should be within the same subnet as the VM's IPv4 address and subnet mask. In this case, the correct gateway address would be within the range of 172.16.31.33 - 172.16.31.62.

upvoted 1 times

scott5010 2 years, 1 month ago

BS question, default gateway has to be on the same subnet as the host, in this case it would be a /27 with 30 hosts.

upvoted 1 times

ryanzou 2 years, 3 months ago

Selected Answer: D

D with no doubt

upvoted 2 times

ironman_86 2 years, 3 months ago

D is the correct one

upvoted 3 times

DVD300 3 years, 1 month ago

D. wrong gateway coz the gateway is in a separate subnet

/27 goes in multiples of 32

upvoted 3 times

All_ultrex 3 years, 3 months ago

This has to be D, wrong gateway. The subnet is a /27 which allows for 30 usable hosts. The IP and gateway are not on the same subnet.

upvoted 4 times

Question #35 *Topic 1*

A company is switching from one cloud provider to another and needs to complete the migration as quickly as possible. Which of the following is the MOST important consideration to ensure a seamless migration?

A. The cost of the environment

B. The I/O of the storage

C. Feature compatibility

D. Network utilization

Suggested Answer: *C*

Community vote distribution

C (80%) B (20%)

SimplyDebonair Highly Voted 2 years, 9 months ago

Selected Answer: C

The answer would be "C." The cost of the environment would already be factored before agreeing to the cloud migration. I/O of storage would not factor here, but it is still important to know. Network utilization isn't clarified for its purpose unless it is trying to outline the speed of the migration. But feature compatibility would clarify if what is being migrated, will work in the new environment without issue (i.e., emphasis on seamless).

upvoted 7 times

Granddude Highly Voted 2 years, 6 months ago

Selected Answer: C

I am going with C based off of https://cloud.netapp.com/blog/data-migration-from-aws-to-azure-reasons-and-challenges

upvoted 5 times

54a6b25 Most Recent 5 months, 3 weeks ago

C. Feature compatibility: Ensuring that the features and services used in the current cloud environment are compatible with those in the new cloud environment is crucial for a seamless migration. This includes compatibility of APIs, data storage formats, networking configurations, and any other cloud services being utilized. If there are discrepancies or incompatibilities, they can lead to application failures, data corruption, or significant rework, which can severely hinder the migration process.

upvoted 2 times

Chiaretta 11 months, 3 weeks ago

Selected Answer: B

The key word in this question is as soon as possible. For me the right is I/O speed of memory

upvoted 3 times

dvd21 3 years, 1 month ago

Feature compatibility

upvoted 3 times

Question #36 *Topic 1*

A systems administrator disabled TLS 1.0 and 1.1, as well as RC4, 3DES, and AES-128 ciphers for TLS 1.2, on a web server. A client now reports being unable to access the web server, but the administrator verifies that the server is online, the web service is running, and other users can reach the server as well. Which of the following should the administrator recommend the user to do FIRST?

A. Disable antivirus/anti-malware software.

B. Turn off the software firewall.

C. Establish a VPN tunnel between the computer and the web server.

D. Update the web browser to the latest version.

Suggested Answer: *D*

Community vote distribution

D (100%)

54a6b25 5 months, 3 weeks ago

D. Update the web browser to the latest version.

Explanation:

The issue arises because the systems administrator disabled older TLS versions (1.0 and 1.1) and deprecated weaker ciphers (RC4, 3DES, AES-128) on the web server. This security enhancement is necessary to protect against vulnerabilities, but older web browsers may not support the newer TLS versions or stronger ciphers.

upvoted 2 times

AustinKelleyNet 1 year, 11 months ago

Selected Answer: D

D is correct. The problem is that the browser doesn't support the latest standards.

upvoted 2 times

Question #37 *Topic 1*

A cloud administrator has finished setting up an application that will use RDP to connect. During testing, users experience a connection timeout error. Which of the following will MOST likely solve the issue?

A. Checking user passwords

B. Configuring QoS rules

C. Enforcing TLS authentication

D. Opening TCP port 3389

Suggested Answer: *D*

Community vote distribution

D (100%)

54a6b25 5 months, 3 weeks ago

The most likely solution to solve the connection timeout error when users are attempting to connect via RDP (Remote Desktop Protocol) is:

D. Opening TCP port 3389.

Explanation:

RDP uses TCP port 3389 by default for communication between the client and the server. If users are experiencing connection timeout errors, it typically indicates that the port through which RDP traffic should flow is not open or accessible. By opening TCP port 3389 on the firewall or network security group (NSG) rules in the cloud environment, you allow incoming RDP connections to reach the server hosting the application.

upvoted 2 times

AustinKelleyNet 1 year, 11 months ago

Selected Answer: D

RDP is port 3389

upvoted 3 times

Question #38 *Topic 1*

A company just successfully completed a DR test and is ready to shut down its DR site and resume normal operations. Which of the following actions should the cloud administrator take FIRST?

A. Initiate a failover.

B. Restore backups.

C. Configure the network.

D. Perform a failback.

Suggested Answer: *D*

Community vote distribution

D (100%)

54a6b25 5 months, 3 weeks ago

D. Perform a failback: Failback is the process of returning operations from the DR site back to the primary site. This involves ensuring that all systems, applications, and services are transitioned back to the primary environment seamlessly. It is critical to perform failback in an orderly manner to minimize disruption and ensure business continuity.

upvoted 2 times

Pongsathorn 1 year, 3 months ago

Selected Answer: D

When a company has successfully completed a Disaster Recovery (DR) test and is ready to resume normal operations, the FIRST action the cloud administrator should take is typically to **perform a failback.**

Here's a brief explanation of each option:

1. **Failover**: Failover is the process of switching from the primary site (which is inoperative due to a disaster or testing) to the secondary or DR site to ensure business continuity. If the DR test was successful, you might have already performed the failover during the test. In a production environment, failover would have been triggered by a disaster.

2. **Restore Backups**: This step may be necessary during the failback process or as part of the overall recovery plan. However, restoring backups is generally done after initiating the failback because you want to ensure that your primary environment is operational before restoring data to it.

upvoted 3 times

Pongsathorn 1 year, 3 months ago

3. **Configure the Network**: Configuring the network is an important step in the overall recovery process, but it should also be done as part of the failback process. During the DR test, network configurations may have been altered to route traffic to the DR site. Configuring the network correctly during the failback ensures that traffic is routed back to the primary site.

4. **Perform a Failback**: A failback is the process of returning operations from the DR site (or backup environment) to the primary site once it is ready to resume normal operations. This step includes reversing any changes made during the failover or DR test and ensuring that the primary environment is fully operational.

In summary, initiating a failback is typically the FIRST action taken after a successful DR test to return operations to the primary site and resume normal business activities. The other actions may be necessary as part of the overall recovery plan but usually follow the failback process.

upvoted 4 times

Pongsathorn 1 year, 3 months ago

Selected Answer: D

A company just successfully completed a DR test and is ready to shut down its DR site and resume normal operations. Which of the following actions should the cloud administrator take FIRST?

A. Initiate a failover.

B. Restore backups.

C. Configure the network.

D. Perform a failback.

upvoted 2 times

sheilawu 1 year, 5 months ago

Selected Answer: D

was on my exam, I go for D

upvoted 3 times

Sal 2 years, 2 months ago

How is it not A???

upvoted 1 times

Princee450 1 year, 4 months ago

It already failover now it needs to failback

upvoted 2 times

SimplyDebonair 2 years, 9 months ago

Selected Answer: D

D would be the correct answer based on the criteria in the scenario. A failover is when you fall forward to mitigation operations for issues per your DR plan(s). Once those issues are over, you are to fall back on normal operations (i.e., failback).

upvoted 3 times

SimplyDebonair 2 years, 9 months ago

Further clarification:

Failover – uses a constant communication mechanism between two systems called a heartbeat. If this heartbeat continues uninterrupted, failover to the redundant system won't initiate. If the heartbeat between systems fails, the redundant system will take over the processing for the primary system.

Failback – when primary systems become operational again, the organization can initiate a failback. Failback is the process of restoring the processing back to the original node. If this were a failure situation, failback would revert processing to the node that had failed once it has been fixed.

upvoted 3 times

SimplyDebonair 2 years, 9 months ago

D would be the correct answer based on the criteria in the scenario. A failover is when you fall forward to mitigation operations for issues per your DR plan(s). Once those issues are over, you are to fall back on normal operations (i.e., failback).

upvoted 2 times

JeanClaud 3 years, 1 month ago

Selected Answer: D

Correct answer should be D. Perform a Failback.

upvoted 3 times

JeanClaud 3 years, 1 month ago

Correct answer should be D. Perform a Failback.

upvoted 3 times

Question #39 *Topic 1*

An IaaS application has a two-hour RTO and four-hour RPO. The application takes one hour to back up its data or restore from a local backup file. A systems administrator is tasked with configuring the backup policy. Which of the following should the administrator configure to achieve the application requirements with the LEAST cost?

A. Back up to long-term storage every night.

B. Back up to object storage every three hours.

C. Back up to long-term storage every four hours.

D. Back up to object storage every hour.

Suggested Answer: *B*

Community vote distribution

B (75%) 13% 13%

⊟ **SimplyDebonair** Highly Voted 👍 2 years, 8 months ago

Selected Answer: B

The correct answer is B.

Source: IT Pro Academy

Long-term storage is used to store the data accessed rarely or mostly saved for archive purposes. It should not be used to store backup data than must be restored in two hour since data in the long-term storage can take several hours to retrieve depending on the specified rehydration priority.

Back up to object storage every hour or three hours both accommodate the requirement of RTO and RPO. However, the option of backup every three hours is less expensive.

upvoted 10 times

⊟ **Kobigasi** Most Recent 🕘 4 months, 1 week ago

Selected Answer: B

It says for the LEAST cost. A and D are already out, and you can't use long term storage for a quick recovery of 4 hours. If your last backup is a MAX of 3 hours, you have an hour to restore and still meet the 4 hour RPO.

upvoted 1 times

⊟ **54a6b25** 5 months, 3 weeks ago

C. Back up to long-term storage every four hours: This option aligns with the RPO of four hours, ensuring that data is backed up at least every four hours. It also balances the cost by using long-term storage, which is typically cheaper compared to more frequent backups to object storage.

B. Back up to object storage every three hours: This option meets the RPO of four hours and provides more frequent backups, but it may not be necessary based on the application's RPO and RTO requirements, potentially increasing cost unnecessarily.

upvoted 1 times

⊟ **TheFivePips** 7 months, 2 weeks ago

This question should be a little more clear with what it wants. C is the cheapest option that fits the requirements that are stated in the question, but generally speaking you wouldn't put backup data into a long term storage. It is possible to retrieve your data within the time frame you want, but not always. B is the "more correct" answer but it is also more expensive. Annoying

upvoted 1 times

⊟ **SamiUlBaseer** 9 months, 3 weeks ago

Selected Answer: D

Given:

RTO (Recovery Time Objective): 2 hours
RPO (Recovery Point Objective): 4 hours
Backup duration: 1 hour
We need to choose an option that allows us to restore the application within 2 hours (RTO) and ensures that the data loss does not exceed 4

hours (RPO).

Option D. Back up to object storage every hour.

This option aligns with the RPO of 4 hours, as backups are taken every hour.
It also aligns with the RTO of 2 hours, as the backups can be restored from object storage within 2 hours.
Since the application takes one hour to back up its data or restore from a local backup file, this option ensures that backups are taken frequently enough to meet the RPO and RTO requirements effectively.
Therefore, the correct option is D. Back up to object storage every hour.

upvoted 2 times

SecPlus2022 1 year, 6 months ago

Selected Answer: B

Just as "SimplyDebonair" states.

upvoted 1 times

Zak11 1 year, 8 months ago

Selected Answer: C

Based on the application's RTO and RPO, the administrator needs to make sure that data can be restored within four hours, and the application can be brought back online within two hours. As the backup takes one hour, this leaves only one hour for the restoration process.

Option A, backing up to long-term storage every night, would have an RPO of 24 hours, which does not meet the application's requirements.

Option B, backing up to object storage every three hours, would have an RPO of three hours and would provide more frequent backups than necessary, increasing costs.

Option D, backing up to object storage every hour, would provide more frequent backups than necessary, increasing costs, and would not meet the RPO of four hours since data restoration takes one hour, and backups occur every hour.

Therefore, the best option would be to back up to long-term storage every four hours, with an RPO of four hours, which meets the application's requirements with the least cost. So, the correct answer is option C.

upvoted 2 times

dvd21 3 years, 1 month ago

Answer is likely B

upvoted 1 times

Question #40 *Topic 1*

A systems administrator is troubleshooting performance issues with a Windows VDI environment. Users have reported that VDI performance has been slow since the images were upgraded from Windows 7 to Windows 10. This VDI environment is used to run simple tasks, such as Microsoft Office. The administrator investigates the virtual machines and finds the following settings:
✑ 4 vCPU
✑ 16GB RAM
✑ 10Gb networking
✑ 256MB frame buffer
Which of the following MOST likely needs to be upgraded?

A. vRAM

B. vCPU

C. vGPU

D. vNIC

Suggested Answer: *C*

Community vote distribution

C (100%)

SimplyDebonair Highly Voted 👍 2 years, 8 months ago

Selected Answer: C

The correct answer is C.

Source: IT Pro Academy

The latest operating system from Microsoft Windows 10 is the most graphically intensive operating system and is designed to deliver improved user experience across PC and mobile devices. Desktop virtualization allows IT to more easily manage and deploy these new upgrades however new considerations must be weighed when selecting the right vGPU profile for your Windows 10 deployment.

According to NVDIA, the minimum frame buffer for a Windows 10 instance is 512MB. In reality, while 512MB will work for some Windows 10 workloads, there are several factors that will increase frame buffer usage above the 512MB threshold and require a 1GB profile to support.

Hence, vGPU MOST likely need to be upgraded.

FYI
- https://blogs.nvidia.com/blog/2016/11/29/vgpu-profile-for-windows-10/
- https://resources.nvidia.com/en-us-grid/vgpu-profile-sizing-guide

upvoted 11 times

54a6b25 Most Recent ⏲ 5 months, 3 weeks ago

Based on the information provided about the Windows VDI environment and the reported performance issues since upgrading from Windows 7 to Windows 10, the component that MOST likely needs to be upgraded is:

B. vCPU
vCPU (Virtual CPU): Windows 10 generally requires more processing power compared to Windows 7, especially for tasks involving newer versions of Microsoft Office and potentially other applications. The current configuration of 4 vCPUs may not be sufficient to handle the increased workload and demand from Windows 10 and modern Office applications.
vGPU (Virtual GPU): The environment described does not indicate a need for a vGPU unless there are specific graphics-intensive applications being used. Since the tasks mentioned are simple (Microsoft Office), it's unlikely that upgrading the vGPU would significantly improve performance.

upvoted 1 times

AllenTaylor 11 months, 3 weeks ago

Selected Answer: C

C is the answer.

upvoted 1 times

maelo 1 year, 6 months ago

Selected Answer: C

A (D) 10Gb connection should be way over any needs, but (C) 256MB frame buffer are low for Win10.

upvoted 2 times

WeLikeSpamHere 1 year, 11 months ago

Selected Answer: C

The correct answer is C.

upvoted 1 times

Question #41 *Topic 1*

An IaaS provider has numerous devices and services that are commissioned and decommissioned automatically on an ongoing basis. The cloud administrator needs to implement a solution that will help reduce administrative overhead. Which of the following will accomplish this task?

A. IPAM

B. NAC

C. NTP

D. DNS

Suggested Answer: *A*

Community vote distribution

A (100%)

⊟ **JohnMangley** Highly Voted 👍 2 years, 1 month ago

Selected Answer: A

from microsoft: IP Address Management (IPAM) is an integrated suite of tools to enable end-to-end planning, deploying, managing and monitoring of your IP address infrastructure, with a rich user experience. IPAM automatically discovers IP address infrastructure servers and Domain Name System (DNS) servers on your network and enables you to manage them from a central interface.

upvoted 7 times

⊟ **reto1** Most Recent 🕘 3 months, 2 weeks ago

IPAM provides a way to manage IP address allocation and DHCP services automatically, which is particularly useful in environments where devices and services are frequently commissioned and decommissioned.

upvoted 1 times

⊟ **54a6b25** 5 months, 3 weeks ago

The solution that will help reduce administrative overhead for an IaaS provider with numerous devices and services being commissioned and decommissioned automatically on an ongoing basis is:

A. IPAM (IP Address Management)

Explanation:

A. IPAM (IP Address Management): IPAM solutions automate the management of IP address spaces, including the allocation, tracking, and deallocation of IP addresses. In an IaaS environment where devices and services are frequently commissioned and decommissioned, IPAM can significantly reduce the administrative burden by automating these tasks, ensuring efficient IP address utilization, and preventing conflicts.

upvoted 1 times

Question #42 *Topic 1*

An organization's web server farm, which is hosted in the cloud with DNS load balancing, is experiencing a spike in network traffic. This has caused an outage of the organization's web server infrastructure. Which of the following should be implemented to prevent this in the future as a mitigation method?

A. Enable DLP.

B. Configure microsegmentation.

C. Enable DNSSEC.

D. Deploy a vADC appliance.

Suggested Answer: *D*

Community vote distribution

D (100%)

Rob69420 Highly Voted 2 years, 4 months ago

Is this D?

vADC is a virtual appliance that improves network application performance and safety. In addition to Layer 4 based load balancing, vADC provides balancing traffic based on Layer 7 HTTP application information like URL, HTTP Cookies or any other HTTP request information.

DNSSEC helps prevent DNS attacks like DNS cache poisoning and DNS spoofing. DNSSEC does not protect the entire server, it only protects the data exchanged between signed zones.

upvoted 12 times

54a6b25 Most Recent 5 months, 3 weeks ago

D. Deploy a vADC (virtual Application Delivery Controller) appliance: A vADC can intelligently distribute network traffic across multiple servers, optimizing resource use and preventing any single server from becoming overwhelmed. It can handle spikes in traffic more effectively by balancing the load and providing features such as caching, compression, and traffic shaping. This ensures higher availability and reliability of the web server farm.

Deploying a vADC appliance is the most effective solution to mitigate the risk of outages due to spikes in network traffic by efficiently managing and distributing the load across the web server infrastructure.

upvoted 1 times

AllenTaylor 11 months, 3 weeks ago

Selected Answer: D

D. Deploy a vADC Appliance. This solution directly addresses the need for better traffic management and infrastructure resilience in the face of high network traffic loads, which is the core issue in the described scenario.

upvoted 1 times

Pongsathorn 1 year, 3 months ago

Selected Answer: D

To mitigate the issue of a spike in network traffic causing an outage in the organization's web server farm hosted in the cloud with DNS load balancing, you should consider **deploying a vADC (Virtual Application Delivery Controller) appliance.**

Here's why:

1. **vADC Appliance**: A vADC is designed to manage and distribute traffic to your web servers efficiently. It can handle traffic spikes, distribute incoming requests evenly across your server farm, and provide various traffic management features like load balancing, SSL termination, caching, and more.

2. **Load Balancing**: vADCs are designed to intelligently distribute traffic to backend servers, ensuring that no single server becomes overwhelmed with traffic. They can use various load balancing algorithms to achieve this.

upvoted 3 times

Pongsathorn 1 year, 3 months ago

3. **Traffic Optimization**: A vADC can optimize traffic by offloading SSL processing, compressing content, and caching frequently accessed resources. This can improve the performance of your web servers.

4. **Health Monitoring**: vADCs can continuously monitor the health of your backend servers and route traffic away from servers that are experiencing issues, ensuring high availability.

5. **Scalability**: You can easily scale up or down your vADC capacity to accommodate changing traffic patterns. This flexibility is crucial during traffic spikes.

While the other options (DLP, microsegmentation, DNSSEC) may have their own importance in a security or network context, they do not directly address the issue of preventing an outage during traffic spikes. Deploying a vADC is a more appropriate solution to handle traffic management and prevent outages caused by sudden increases in network traffic.

upvoted 1 times

⊟ **ATill** 1 year, 8 months ago

Selected Answer: D

As Rob69420 said. D.

upvoted 1 times

⊟ **JVen** 2 years, 1 month ago

Selected Answer: D

D makes the most sense. See Rob69420's comment for why

upvoted 2 times

⊟ **scott5010** 2 years, 1 month ago

vADC is the answer, D.

upvoted 1 times

⊟ **Moosafat** 2 years, 3 months ago

Selected Answer: D

virtual Application Delivery Controller

upvoted 2 times

⊟ **ironman_86** 2 years, 3 months ago

VADC virtual Application Delivery Controller

upvoted 2 times

Question #43 *Topic 1*

A vendor is installing a new retail store management application for a customer. The application license ensures software costs are low when the application is not being used, but costs go up when use is higher. Which of the following licensing models is MOST likely being used?

A. Socket-based

B. Core-based

C. Subscription

D. Volume-based

Suggested Answer: *C*

Community vote distribution

C (75%) D (25%)

SimplyDebonair Highly Voted 2 years, 8 months ago

Selected Answer: C

The correct answer is C.

Source: IT Pro Academy

Volume licensing is a special type of software licensing setup that uses a single license key to authorize the software on multiple computers. For instance, a company may purchase a software license that allows up to 50 of their employees to use the software at the same time.

The socket based license means that a license for each physical socket is bought regardless of how many CPU cores it has. This license gets less expensive as the number of CPU cores per socket increase and load more VMs onto the system.

Core-based licensing requires all physical cores in the server to be licensed. Servers are licensed based on the number of processor cores in the physical server.

These 3 licensing models do not provide the required flexibility and elasticity: software costs are low when the application is not being used, but costs go up when use is higher.

The subscription model or the pay-as-you-grow model allows billing for only their use of resources.

upvoted 11 times

TheFivePips Highly Voted 7 months, 2 weeks ago

Selected Answer: C

This is from the comptia certmaster:

Volume based: One license that permits a specified number of installations, for example, installation of the software on up to 100 computers

Subscription: Periodic cost; usually includes at least basic technical support, maintenance, and possibly upgrades

Neither answer sound particularly good here, but we at least know subscription costs can be based on usage (tiers of use) and Volume based is basically just buying in bulk. Volume based don't scale at all.

If you were like me and guessed that volume based licensing would be based on the amount of usage, because it sure sounds like it might be, CompTIA would like to remind you that you're an idiot.

upvoted 5 times

54a6b25 Most Recent 5 months, 3 weeks ago

D. Volume-based: This model charges based on the amount of usage, such as the number of transactions, amount of data processed, or number of users accessing the application. Costs increase with higher usage and decrease with lower usage, making it the most likely model being described in the scenario.

Subscription: A subscription model charges a recurring fee, usually monthly or annually, for the use of the software. While this can sometimes be based on the level of service or number of users, it typically does not fluctuate directly with usage levels.

upvoted 1 times

FrancisDrake 11 months ago

Selected Answer: D

In my experience subscriptions are a fixed cost. I'm going with Volume based.

From Wikipedia: "In software licensing, volume licensing is the practice of using one license to authorize software on a large number of computers and/or for a large number of users. Customers of such licensing schemes are typically business, governmental or educational institutions, with PRICES for volume licensing varying depending on the type, QUANTITY and applicable subscription-term. For example, Microsoft software available through volume-licensing programs includes Microsoft Windows and Microsoft Office."

upvoted 2 times

Frogman1981 1 year ago

D.)
Volume-based Licensing:
What it is:
Customers pay for the software based on the volume of data they process, the number of transactions they perform, or any other quantifiable metric related to its usage.

Subscription-based Licensing:
What it is:
Customers pay a recurring fee, usually monthly or yearly, for access to the software, regardless of their usage level.

upvoted 4 times

Pongsathorn 1 year, 3 months ago

Selected Answer: C

Per user
One license for each user that consumes the software or service

Socket based
One license for each CPU that attaches to the socket of a motherboard, regardless of the number of cores the CPU might contain

Core based
One license for each core in a CPU in a server

Volume based
One license that permits a specified number of installations, for example, installation of the software on up to 100 computers

Perpetual
One-time fee for a license that may include additional support costs; however, the license is good for the life of the software

Subscription
Periodic cost; usually includes at least basic technical support, maintenance, and possibly upgrades

upvoted 2 times

sheilawu 1 year, 5 months ago

Selected Answer: D

I am go for D because the AZ-900 also has the question similar of this senario and it definded the volume based as the answer

upvoted 2 times

sheilawu 1 year, 6 months ago

Selected Answer: D

volume based is correct

upvoted 1 times

nate612 1 year, 8 months ago

Why not D, volume-based?

upvoted 1 times

maelo 1 year, 8 months ago

While A+B seem physical hardware based and static, C+D are elastic. Yet, subscription usually needs contractual interaction for (easy) up/down sizing, while volume licensing should derive costs right from the parameter of interest.

upvoted 1 times

beamage 1 year, 11 months ago

Selected Answer: C

Subscriptions may be increased and decreased as needed....

upvoted 2 times

TheGinjaNinja 1 year, 11 months ago

Selected Answer: D

D. Volume-based licensing is the most likely model being used in this scenario. This model charges customers based on the amount of usage, such as the number of transactions, number of users, or amount of data stored, rather than a one-time purchase or a recurring subscription fee. This aligns with the scenario where costs are low when the application is not being used, but increase when usage is higher.

upvoted 2 times

rob88Silva 1 year, 11 months ago

Selected Answer: D

D. Volume-based

A volume-based licensing model is most likely being used in this scenario, where the cost of the application is based on the level of usage. Under this model, the vendor charges the customer a set fee for a certain number of users or transactions, and additional fees for any usage above that limit. This licensing model would allow the vendor to offer lower costs when the application is not being used, but higher costs when usage is higher.

upvoted 2 times

JVen 2 years, 1 month ago

Selected Answer: C

It can only be C

upvoted 3 times

scott5010 2 years, 1 month ago

Selected Answer: C

gotta be C

upvoted 2 times

Moosafat 2 years, 3 months ago

Selected Answer: C

The subscription model or the pay-as-you-grow model allows billing for only their use of resources.

upvoted 2 times

martin451 3 years, 2 months ago

this should be A socket-base

upvoted 1 times

> **JeanClaud** 3 years, 1 month ago
>
> Hard disagree. Socket-based cost would remain the same unless amount of sockets changed.
>
> upvoted 5 times
>
> > **EyaT** 2 years, 9 months ago
> >
> > https://ccna7.org/a-vendor-is-installing-a-new-retail-store-management-application-for-a-customer-the-application-license-ensures-software-costs-are-low-when-the-application-is-not-being-used-but-costs-go-up-when-use/
> >
> > upvoted 1 times

Question #44 *Topic 1*

A systems administrator in a large enterprise needs to alter the configuration of one of the finance department's database servers. Which of the following should the administrator perform FIRST?

A. Capacity planning

B. Change management

C. Backups

D. Patching

Suggested Answer: *B*

Community vote distribution

B (100%)

Pongsathorn 3 months, 4 weeks ago

Selected Answer: B

Change management refers to preparing, supporting, and managing new or updated business processes or technology. There are many change management models, but in general, the steps look like this:

1. Propose and plan: Understand the business need and how the solution meets it.
2. Approve: Agree to provide resources (funding, time, and staff) for the proposed change.
3. Develop: Develop and test the new application, service, or other feature.
4. Deploy: Roll the change out for day-to-day use.
5. Close: The change is complete and moves to a maintenance phase.

Change management is a major undertaking in IT, and entire career paths can be dedicated to this field.

upvoted 3 times

Lenell 1 year ago

Selected Answer: B

The SA would do the other three regardless of the need to alter configurations. In this situation, the SA would have to present the change to the CCB in order to do the alteration.

upvoted 2 times

Agr321 1 year, 3 months ago

Change Management

A systems administrator in a large enterprise needs to alter the configuration of one of the finance department's database servers. Which of the following should the administrator perform FIRST?

upvoted 1 times

jiminycriminal 1 year, 3 months ago

Does an administrator "perform" change management? It says he needs to make a change, and I feel that implies the change was already authorized? What a vague question.

upvoted 2 times

> **Agr321** 1 year, 3 months ago
>
> Taking the perform out of the question and replace it with STEP.
> When approaching these questions when faced with a policy or physical change.
> It's the finance department SQL server, not some developers SQL server.
>
> If you approach questions like above, you will have a better time answering them correctly.
>
> upvoted 1 times

SimplyDebonair 1 year, 9 months ago

The correct answer would be B. There is no clarification on whether the change management process has been gone through. Any changes, regardless of how small or big, must go through the change management process. This allows proposals to be heard by end-users, management, and possibly stockholders. From there, it will be reviewed and either approved or denied, with reasons specified. From there, the administrator(s) can do whatever processes are necessary.

upvoted 4 times

⊟ **Granddude** 1 year, 6 months ago

I agree. Besides, wouldn't backups already be in place for database servers?

upvoted 1 times

Question #45 *Topic 1*

A database analyst reports it takes two hours to perform a scheduled job after onboarding 10,000 new users to the system. The analyst made no changes to the scheduled job before or after onboarding the users. The database is hosted in an IaaS instance on a cloud provider. Which of the following should the cloud administrator evaluate to troubleshoot the performance of the job?

A. The IaaS compute configurations, the capacity trend analysis reports, and the storage IOPS

B. The hypervisor logs, the memory utilization of the hypervisor host, and the network throughput of the hypervisor

C. The scheduled job logs for successes and failures, the time taken to execute the job, and the job schedule

D. Migrating from IaaS to on premises, the network traffic between on-premises users and the IaaS instance, and the CPU utilization of the hypervisor host

Suggested Answer: *A*

Community vote distribution

A (92%) 8%

TheGinjaNinja Highly Voted 1 year, 5 months ago

Selected Answer: A

A. The IaaS compute configurations, the capacity trend analysis reports, and the storage IOPS. These factors can affect the performance of the IaaS instance and thus the performance of the scheduled job. The cloud administrator should evaluate the compute configurations to ensure that the instance has the appropriate resources such as CPU, memory and storage to handle the increased workload due to the new users. The capacity trend analysis reports can provide insight into the performance of the instance over time. The storage IOPS can also affect the performance of the job because it relates to the speed at which the data can be read and written.

upvoted 10 times

Monkeyman1500 Most Recent 4 months, 3 weeks ago

Selected Answer: C

Ehhh C sounds right. Sure you onboarded a bunch of users, but it never stated the job was related to the users being onboarded or that the job was taking any longer than normal. You would investigate the job for errors, schedule, average time to complete, etc... before assuming there was a problem.

upvoted 1 times

kuzummjakk 3 months, 3 weeks ago

That'd make sense, but not as a test answer to the question. The question is very much saying "nothing about the job changed. The only change is onboarding 1,000 users *inserts cloud infrastructure information for cloud test*"

upvoted 1 times

Pongsathorn 9 months, 3 weeks ago

To troubleshoot the performance of the scheduled job on the IaaS instance in the cloud after onboarding 10,000 new users, the cloud administrator should evaluate the following:

A. The IaaS compute configurations, the capacity trend analysis reports, and the storage IOPS.

Here's why:

1. **IaaS Compute Configurations**: The cloud administrator should check if the compute resources allocated to the IaaS instance are sufficient to handle the increased load after onboarding new users. If the instance doesn't have enough CPU, RAM, or other resources, it can slow down job processing.

2. **Capacity Trend Analysis Reports**: Monitoring capacity trends over time can help identify if the instance's resource utilization has significantly changed after onboarding new users. This analysis can reveal resource bottlenecks.

upvoted 3 times

Pongsathorn 9 months, 3 weeks ago

3. **Storage IOPS**: Input/Output Operations Per Second (IOPS) on the storage system can impact database performance. If the storage system is not providing sufficient IOPS, it can cause delays in data retrieval or updates, affecting the job execution time.

Option B, which mentions hypervisor-related logs and metrics, might not be as relevant in this case, as the issue seems to be more about the IaaS instance's performance and resource allocation rather than the underlying hypervisor.

Option C, checking scheduled job logs, is important for understanding the job's execution time and any potential issues within the job itself. However, it doesn't directly address the performance problem caused by the onboarding of new users.

upvoted 3 times

Pongsathorn 9 months, 3 weeks ago

Option D, migrating from IaaS to on-premises, is a drastic step and should only be considered after thoroughly investigating and optimizing the current cloud environment. Additionally, migrating may not necessarily solve the performance issue if the on-premises infrastructure is not adequately configured.

Therefore, option A is the most appropriate initial step to evaluate and troubleshoot the performance of the scheduled job in the cloud environment.

upvoted 2 times

Zak11 1 year, 2 months ago

Selected Answer: A

@TheGinjaNinja is correct

upvoted 1 times

Question #46 *Topic 1*

A systems administrator is reviewing two CPU models for a cloud deployment. Both CPUs have the same number of cores/threads and run at the same clock speed. Which of the following will BEST identify the CPU with more computational power?

A. Simultaneous multithreading

B. Bus speed

C. L3 cache

D. Instructions per cycle

Suggested Answer: *D*

Community vote distribution

D (100%)

54a6b25 5 months, 3 weeks ago

D. Instructions per cycle (IPC): IPC is a measure of how many instructions a CPU can execute in one clock cycle. A higher IPC indicates a more efficient CPU that can perform more work per clock cycle, thus providing greater computational power.

upvoted 2 times

TheFivePips 7 months, 2 weeks ago

Selected Answer: D

Clock speed (measured in GHz) refers to how fast the CPU can process instructions per second. It's like the speed limit on a highway; a higher speed limit allows more cars to pass through in a given amount of time.

Instructions per cycle (IPC) refers to how many instructions a CPU can execute in one clock cycle. It's like how efficient each car is at carrying passengers; a car with more seats can carry more people in a single trip.

Relationship: Clock speed determines how many clock cycles occur in a second, essentially setting the pace of the CPU. IPC determines how many instructions the CPU can process within each of those clock cycles, indicating its efficiency.

upvoted 4 times

kuzummjakk 9 months, 4 weeks ago

Selected Answer: D

I second this

upvoted 1 times

Question #47 *Topic 1*

A systems administrator is building a new visualization cluster. The cluster consists of five virtual hosts, which each have flash and spinning disks. This storage is shared among all the virtual hosts, where a virtual machine running on one host may store data on another host. This is an example of:

A. a storage area network.

B. a network file system.

C. hyperconverged storage.

D. thick-provisioned disks.

Suggested Answer: *C*

Community vote distribution

C (86%) 14%

54a6b25 5 months, 3 weeks ago

C. Hyperconverged Storage: Hyperconverged storage is an approach where storage resources are integrated directly into the virtual hosts (compute nodes) and managed together with compute and networking in a single system. In hyperconverged infrastructure, storage is typically pooled and shared across the cluster, allowing virtual machines to use storage resources from any node in the cluster. This matches the scenario described, where the storage is shared among all virtual hosts, and a virtual machine on one host may store data on another host.

upvoted 3 times

FrancisDrake 12 months ago

Selected Answer: C

A SAN would mean separate storage from the hosts. This question indicates that each host has its own storage yet shares it with the other hosts.

upvoted 2 times

FrancisDrake 12 months ago

A SAN would mean separate storage from the hosts. This question indicates that each host has its own storage yet shares it with the other hosts.

upvoted 1 times

VVV4WIN 1 year, 1 month ago

Selected Answer: C

Definitely C

upvoted 1 times

backdooranon 1 year, 2 months ago

Selected Answer: C

Question specifically stated that each host in the cluster has its own storage, which is then shared with one another. SAN requires dedicated storage infrastructure.

upvoted 3 times

backdooranon 1 year, 2 months ago

See this VMWare definition of hyperconverged storage if you're unsure:

Hyperconverged storage is one facet of hyperconverged infrastructure (HCI), in which storage is bundled with compute and networking in a single virtualized system. With this software-defined approach, flexible pools of storage replace dedicated hardware. Each node includes a software layer that virtualizes the resources in the node and shares them across all the nodes in a cluster, creating one large storage pool.

https://www.vmware.com/topics/glossary/content/hyperconverged-storage.html

upvoted 3 times

maelo 1 year, 4 months ago

Selected Answer: A

Although I originally tended to "hyperconverged infrastructure", I finally chose A, "SAN", because hyperconverged infrastructure is defined as a mix of different resources types, such as hypervisor, processing, storage, network and others.

upvoted 1 times

kuzummjakk 9 months, 4 weeks ago

but it says "storage" not "infrastructure.

upvoted 1 times

Question #48 *Topic 1*

A VDI administrator has received reports of poor application performance. Which of the following should the administrator troubleshoot FIRST?

A. The network environment

B. Container resources

C. Client devices

D. Server resources

Suggested Answer: *D*

Community vote distribution

D (89%) 11%

AllenTaylor Highly Voted 5 months, 2 weeks ago

Selected Answer: D

Multiple reports would indicate the server as the issue.

upvoted 5 times

Jay987654 Most Recent 5 months ago

Selected Answer: D

The administrator should troubleshoot D. Server resources first.

In a Virtual Desktop Infrastructure (VDI) environment, poor application performance is often related to the resources on the server hosting the virtual desktops. This could include CPU, memory, storage, or network resources. If these resources are insufficient or not performing optimally, it can lead to poor performance for all users on the server.

upvoted 3 times

nmap_king_22 9 months, 1 week ago

Selected Answer: C

When a VDI (Virtual Desktop Infrastructure) administrator receives reports of poor application performance, the first step in troubleshooting should typically focus on the most likely causes of the issue. While a comprehensive investigation may be required, addressing the most probable root causes first can often resolve the problem more efficiently. In this scenario, you should prioritize checking:

C. Client devices

Client devices can often be a common source of poor VDI performance.

upvoted 1 times

Gwcan 8 months, 2 weeks ago

They're probably using thin clients. All the real computational resources are in the cloud so why would client devices be the answer?

upvoted 5 times

kuzummjakk 3 months, 3 weeks ago

ChatGPT type answer. Lookup VDI

upvoted 1 times

Question #49 *Topic 1*

A company has developed a cloud-ready application. Before deployment, an administrator needs to select a deployment technology that provides a high level of portability and is lightweight in terms of footprint and resource requirements. Which of the following solutions will be BEST to help the administrator achieve the requirements?

A. Containers

B. Infrastructure as a code

C. Desktop virtualization

D. Virtual machines

Suggested Answer: *A*

Community vote distribution

A (100%)

kuzummjakk 3 months, 3 weeks ago

keyword portable

upvoted 2 times

bsalama 8 months, 3 weeks ago

Selected Answer: A

A. Containers

Containers are a technology designed for high portability and efficiency. They provide a lightweight, isolated environment for applications, ensuring that the application and all its dependencies are bundled together. Containers are known for their low resource requirements, quick startup times, and ease of deployment across various cloud platforms and environments. They are highly portable and can run consistently on different infrastructure, which makes them a popular choice for cloud-native applications.

upvoted 3 times

Question #50 *Topic 1*

A systems administrator is deploying a VM and would like to minimize storage utilization by ensuring the VM uses only the storage if needs.
Which of the following will BEST achieve this goal?

A. Compression

B. Deduplication

C. RAID

D. Thin provisioning

Suggested Answer: *D*

Community vote distribution

D (100%)

bsalama Highly Voted 2 months, 3 weeks ago

Selected Answer: D

D. Thin provisioning

Thin provisioning allows you to allocate storage space to a VM on an as-needed basis. It means that when you create a VM, it's given a small initial allocation of storage, and additional storage space is allocated only as the VM consumes it. This approach optimizes storage utilization, as you avoid allocating a large amount of storage upfront that may go unused. It's an efficient way to make the most of your storage resources while still accommodating the VM's actual storage needs.

upvoted 5 times

Question #51 *Topic 1*

An organization requires the following to be achieved between the finance and marketing departments:
✑ Allow HTTPS/HTTP.
✑ Disable FTP and SMB traffic.
Which of the following is the MOST suitable method to meet the requirements?

A. Implement an ADC solution to load balance the VLAN traffic.

B. Configure an ACL between the VLANs.

C. Implement 802.1X in these VLANs.

D. Configure on-demand routing between the VLANs.

Suggested Answer: *B*

Community vote distribution

B (100%)

54a6b25 5 months, 3 weeks ago

B. Configure an ACL between the VLANs: Access Control Lists (ACLs) are used to permit or deny traffic based on specified criteria such as protocol type, source IP, and destination IP. By configuring ACLs between the VLANs, the administrator can allow HTTPS/HTTP traffic and deny FTP and SMB traffic, meeting the requirements effectively.

upvoted 2 times

nmap_king_22 1 year, 3 months ago

Selected Answer: B

B. Configure an ACL between the VLANs.

The most suitable method to meet the specified requirements of allowing HTTP/HTTPS traffic while disabling FTP and SMB traffic between the finance and marketing departments is to use Access Control Lists (ACLs).

upvoted 2 times

Question #52 *Topic 1*

A company wants to check its infrastructure and application for security issues regularly. Which of the following should the company implement?

A. Performance testing

B. Penetration testing

C. Vulnerability testing

D. Regression testing

Suggested Answer: *C*

Community vote distribution

C (86%) 14%

bsalama Highly Voted 1 year, 2 months ago

Selected Answer: C

C. Vulnerability testing

Vulnerability testing, or vulnerability assessment, is a security practice that helps identify weaknesses or vulnerabilities in an organization's infrastructure and applications. It involves scanning systems and applications for known vulnerabilities and potential security issues. This proactive approach allows the company to identify and address security weaknesses before they can be exploited by malicious actors.

The other options are not focused on security assessments:

A. Performance testing is about evaluating the performance characteristics of applications but doesn't primarily address security.
B. Penetration testing involves simulating cyberattacks to identify vulnerabilities but is typically done periodically or on an ad-hoc basis and may not cover regular security assessments.
D. Regression testing is a testing technique to ensure that new code changes do not negatively impact existing functionality but is not focused on security checks.

upvoted 6 times

54a6b25 Most Recent 5 months, 3 weeks ago

C. Vulnerability testing: Also known as vulnerability scanning, this process involves using automated tools to regularly scan systems and applications for known vulnerabilities. This helps in identifying security issues continuously and ensuring they are addressed promptly.

upvoted 1 times

maelo 1 year, 4 months ago

Selected Answer: B

An infrastructure and app scan is more a pentest than a vultest.

upvoted 1 times

kuzummjakk 9 months, 4 weeks ago

Good point. Vultests are generally associated with hosts; not always, but generally. However the keyword here is "security issues". Very vulnerability scan wording.

upvoted 1 times

Gwcan 1 year, 2 months ago

I think the keyword is "regularly". You probably don't simulate attacks regularly, but you do scan for vulnerabilities regularly. I'd go with C.

upvoted 3 times

Question #53 *Topic 1*

A systems administrator is analyzing a report of slow performance in a cloud application. This application is working behind a network load balancer with two VMs, and each VM has its own digital certificate configured. Currently, each VM is consuming 85% CPU on average. Due to cost restrictions, the administrator cannot scale vertically or horizontally in the environment. Which of the following actions should the administrator take to decrease the CPU utilization? (Choose two.)

A. Configure the communication between the load balancer and the VMs to use a VPN.

B. Move the digital certificate to the load balancer.

C. Configure the communication between the load balancer and the VMs to use HTTP.

D. Reissue digital certificates on the VMs.

E. Configure the communication between the load balancer and the VMs to use HTTPS.

F. Keep the digital certificates on the VMs.

Suggested Answer: *BE*

Community vote distribution

BE (51%) BC (49%)

TheGinjaNinja Highly Voted 1 year, 11 months ago

Selected Answer: BE

B. Move the digital certificate to the load balancer

E. Configure the communication between the load balancer and the VMs to use HTTPS

By moving the digital certificate to the load balancer, the system administrator is offloading the processing of the SSL/TLS encryption to the load balancer and not the VMs. This can help to decrease the CPU utilization on the VMs. Additionally, configuring the communication between the load balancer and the VMs to use HTTPS also eliminates the need to process encryption on the VMs, this will also decrease the CPU utilization on the VMs.

upvoted 12 times

reto1 3 months, 2 weeks ago

chatgpt approved

upvoted 1 times

FasterN8 10 months, 1 week ago

No. HTTPS does not eliminate the need to process encryption. The 'S' part of that is 'Secure' which uses TLS encryption. Frontload the decryption to the load balancer and go with a non-encrypted protocol on the backside to "... reduce CPU utilization" which is the focus of the question.

upvoted 5 times

K_J_ Most Recent 3 weeks, 1 day ago

Selected Answer: BC

HTTP requires less CPU processing compared to HTTPS.

upvoted 1 times

54a6b25 5 months, 3 weeks ago

B. Move the digital certificate to the load balancer.

C. Configure the communication between the load balancer and the VMs to use HTTP.

Explanation:

B. Move the digital certificate to the load balancer: Offloading SSL/TLS termination to the load balancer can significantly reduce CPU utilization on the VMs. The load balancer will handle the encryption/decryption process, freeing up CPU resources on the VMs.

C. Configure the communication between the load balancer and the VMs to use HTTP: After moving the digital certificate to the load balancer, the traffic between the load balancer and the VMs can be sent over HTTP, which requires less CPU processing compared to HTTPS. This further reduces the CPU load on the VMs.

upvoted 3 times

kuzummjakk 9 months, 4 weeks ago

Selected Answer: BC

"Moves cert to the load balancer so the VMs dont have to encrypt/decrypt"

"Makes the VM's communicate in HTTPs so they have to encrypt/encrypt"

what. cannot possibly be B and E.

upvoted 2 times

FasterN8 10 months, 1 week ago

Selected Answer: BC

No. HTTPS does not eliminate the need to process encryption. The 'S' part of that is 'Secure' which uses TLS encryption. Frontload the decryption to the load balancer and go with a non-encrypted protocol on the backside to "... reduce CPU utilization" which is the focus of the question.

upvoted 3 times

FrancisDrake 12 months ago

Selected Answer: BC

I am moving the certificates to the load balancer and configuring http communication between the VMs and the load balancer.

upvoted 2 times

FrancisDrake 12 months ago

This is an interesting question. If you are trying to reduce resource usage I would move the digital certificates to the load balancer and enable http communication between the VMs and the load balancer (HTTP I believe uses fewer resources).

upvoted 1 times

MJ06 1 year ago

Selected Answer: BC

People! If you remove the certificate from the VMs, how can you configure a secure connection between the load balancer and the VMs???

upvoted 2 times

> **FrancisDrake** 12 months ago
>
> I think because they are on the same internal network that is less of a concern.
>
> upvoted 3 times

SecPlus2022 1 year, 6 months ago

Selected Answer: BE

As explained by "TheGinjaNinja".

upvoted 2 times

Alizadeh 1 year, 8 months ago

Selected Answer: BC

B. Move the digital certificate to the load balancer.

C. Configure the communication between the load balancer and the VMs to use HTTP.

upvoted 3 times

LeDarius3762 1 year, 10 months ago

B) Move the digital certificate to the load balancer

E) Configure the communication between the load balancer and the VMs to use HTTPS

https://www.httpvshttps.com/

upvoted 3 times

rob88Silva 1 year, 11 months ago

Selected Answer: BE

Moving the digital certificate to the load balancer would offload the workload of encrypting and decrypting traffic from the VMs. This would reduce the CPU utilization on the VMs, allowing them to handle more requests.

Configuring the communication between the load balancer and the VMs to use HTTPS would also offload the workload of encrypting and decrypting traffic from the VMs. Additionally, using HTTPS would improve the security of the communication between the load balancer and the VMs.

upvoted 4 times

[Removed] 1 year, 11 months ago

B&C

Why not E? because no scope to scale out or scale up (as mentioned) and need to reduce CPU utilization. Please note, https is something recomended but it uses high CPU which is a problem here so going with http would reduce overhead which results lower cpu. :)

upvoted 1 times

scott5010 2 years, 1 month ago

Selected Answer: BC

here is a compelling article for bc, https://www.claudiokuenzler.com/blog/687/encrypted-http-connection-https-use-four-times-more-cpu-resources-load

upvoted 2 times

Not_That_Guy 2 years, 2 months ago

Selected Answer: BE

B obviously; move the overhead of dealing with certificates to the load balance. E most likely; data privacy standards require data leaving the VM to be encrypted (https).

upvoted 1 times

SimplyDebonair 2 years, 8 months ago

Selected Answer: BC

The correct answers are "B" and "C." As previously mentioned by dvd21: the overhead for HTTP compared to HTTPS is lower.

https://serverfault.com/questions/570387/https-overhead-compared-to-http

upvoted 4 times

> **i_bird** 2 years, 3 months ago
>
> HTTPS uses digital certificate, not HTTP, you still have to secure your connection right?
>
> upvoted 2 times
>
> > **ironman_86** 2 years, 3 months ago
> >
> > @i_bird, the load balancer is securing the communication as the cert has been move to it and only the communication between the load balancer and VMs will be not secured
> >
> > upvoted 1 times

dvd21 3 years, 1 month ago

answer is B and C. Http reduces CPU load.

upvoted 3 times

Question #54 *Topic 1*

A private IaaS administrator is receiving reports that all newly provisioned Linux VMs are running an earlier version of the OS than they should be. The administrator reviews the automation scripts to troubleshoot the issue and determines the scripts ran successfully. Which of the following is the MOST likely cause of the issue?

A. API version incompatibility

B. Misconfigured script account

C. Wrong template selection

D. Incorrect provisioning script indentation

Suggested Answer: *C*

Community vote distribution

C (100%)

kuzummjakk 3 months, 3 weeks ago

Selected Answer: C

For a bit I thought of the "automation script" as being separated from whatever deployed the machines. C makes sense since it's an automation script that just deployed new machines

upvoted 1 times

nmap_king_22 9 months, 1 week ago

Selected Answer: C

Template Selection: When provisioning virtual machines in an IaaS environment, administrators often use templates or images as a base. These templates contain a specific version and configuration of the operating system. If the wrong template is selected during the provisioning process, it will result in VMs being deployed with the earlier OS version specified in the selected template.

upvoted 4 times

Question #55 *Topic 1*

A cloud administrator is reviewing a new application implementation document. The administrator needs to make sure all the known bugs and fixes are applied, and unwanted ports and services are disabled. Which of the following techniques would BEST help the administrator assess these business requirements?

A. Performance testing

B. Usability testing

C. Vulnerability testing

D. Regression testing

Suggested Answer: *C*

Community vote distribution

C (100%)

kuzummjakk 3 months, 3 weeks ago

Selected Answer: C

Keyword is KNOWN bugs and fixes; ESPECIALLY ports and services. Although, the grammar of "all the known bugs and fixes are applied" is pretty bad; you don't apply known bugs.

upvoted 2 times

alittlesmarternow 5 months ago

Regression testing: This technique re-runs existing test cases against the new implementation to identify any unintended changes or regressions introduced during development. It directly addresses whether known bugs have been fixed and if the existing functionality remains intact.

upvoted 1 times

> **kuzummjakk** 3 months, 3 weeks ago
>
> Regression testing doesn't address known bugs, it addresses accounting for unknown bugs. It's specifically related to app development and the case where your in-production app gets unexpected app breaking bus and you need to "regress" to an older version of that app.
>
> upvoted 1 times
>
> > **kuzummjakk** 3 months, 3 weeks ago
> >
> > also it's specifically related to "regressing" in the case of bugs, not fixing bugs.
> >
> > upvoted 1 times

Question #56 *Topic 1*

A DevOps administrator is automating an existing software development workflow. The administrator wants to ensure that prior to any new code going into production, tests confirm the new code does not negatively impact existing automation activities. Which of the following testing techniques would be BEST to use?

A. Usability testing

B. Regression testing

C. Vulnerability testing

D. Penetration testing

Suggested Answer: *B*

Community vote distribution

B (100%)

nmap_king_22 3 months, 1 week ago

Selected Answer: B

Regression Testing: Regression testing is specifically designed to check whether new changes (new code in this case) have introduced any unintended side effects or defects in the existing functionality of the software. It ensures that the new code changes haven't negatively impacted the existing automation activities and that everything still works as expected.

upvoted 2 times

Alizadeh 8 months, 4 weeks ago

Selected Answer: B

B. Regression testing

Regression testing would be the best technique to use in this scenario. Regression testing is conducted to ensure that new code changes do not negatively impact the existing functionality or automation activities. It involves re-running test cases from a previously verified test suite to confirm that the new code has not introduced any bugs or issues in the previously working functionalities.

upvoted 3 times

Question #57 *Topic 1*

Some VMs that are hosted on a dedicated host server have each been allocated with 32GB of memory. Some of VMs are not utilizing more than 30% of the allocation. Which of the following should be enabled to optimize the memory utilization?

A. Auto-scaling of compute

B. Oversubscription

C. Dynamic memory allocations on guests

D. Affinity rules in the hypervisor

Suggested Answer: *C*

Community vote distribution

C (100%)

AustinKelleyNet 5 months ago

Selected Answer: C

C is correct

upvoted 2 times

martin451 1 year, 8 months ago

answer should be c Dynamic memory allocations on guests

upvoted 3 times

Question #58 *Topic 1*

A systems administrator would like to reduce the network delay between two servers. Which of the following will reduce the network delay without taxing other system resources?

A. Decrease the MTU size on both servers.

B. Adjust the CPU resources on both servers.

C. Enable compression between the servers.

D. Configure a VPN tunnel between the servers.

Suggested Answer: *A*

Community vote distribution

A (65%) C (35%)

BeauChateau Highly Voted 1 year, 8 months ago

Selected Answer: C

C. Enable compression between the servers.

Enabling compression between the servers can help to reduce the network delay without taxing other system resources. Compression reduces the amount of data that needs to be transmitted over the network, thereby reducing the overall network delay. Compression can be particularly effective for data that is highly redundant or compressible, such as text files or log files.

upvoted 5 times

reto1 3 months, 2 weeks ago

C. ChatGPT approved.

Enabling compression will help reduce the amount of data that needs to be transmitted over the network, thereby potentially reducing the time it takes for data to travel between the two servers. This approach minimizes network delay without putting additional strain on CPU or other resources, as it primarily optimizes the data transfer itself.

upvoted 1 times

Kobigasi 4 months, 1 week ago

It's A, NOT C. In what world does compression decrease latency?? Reducing MTU size means faster packet transmission.

upvoted 1 times

TheFivePips Most Recent 7 months, 2 weeks ago

Selected Answer: A

Given the options and considering the potential impact on system resources, if I had to choose one, I would lean towards Option A: Decrease the MTU size on both servers.

While both options have their pros and cons, reducing the MTU size addresses potential delays caused by packet fragmentation, which can occur when packets are too large to traverse certain network segments without being fragmented. By reducing the MTU size, we aim to minimize the need for packet fragmentation, thus potentially reducing delays associated with retransmission of fragmented packets.

Additionally, adjusting the MTU size is a network-level configuration that doesn't directly impose additional computational overhead on the servers, unlike enabling compression. While decreasing the MTU size may lead to slightly more packets being sent, it generally doesn't require significant additional computational resources compared to compression.

However, it's important to note that the effectiveness of this approach depends on the specific network environment and the underlying causes of network delay that the question simply does not give us.

upvoted 3 times

kuzummjakk 9 months, 4 weeks ago

Selected Answer: A

I was gonna say C, but i realized it doesn't say "other systems", it says "other system RESOURCES".

upvoted 1 times

cobbs 1 year, 3 months ago

Selected Answer: A

Increasing the MTU can improve performance, but decreasing the MTU can resolve packet loss and fragmentation problems when it is too high (and will not tax system resources like "enable compression" would).

upvoted 2 times

SecPlus2022 1 year, 7 months ago

Selected Answer: A

As Maelo states, compression and decompression takes CPU resources to process. Reducing the MTU has ZERO affect on system resources (which is a requirement as per the question).

upvoted 2 times

maelo 1 year, 7 months ago

Selected Answer: A

Compressed transmission needs compression and decompression processes before using the data. Reducing MTU size scales headers transmissions up, reduces transmission efficiency, but speeds up packet send/receive actions, not waiting for more data to fill the MTU. Extinguish a fire better with small buckets or big buckets? Throughput of reasonably small buckets is better.

upvoted 3 times

scott5010 2 years, 1 month ago

Selected Answer: A

C is the best solution but with question is looking for a solution that doesn't effect other components MTU size is the answer

upvoted 2 times

nate612 2 years, 2 months ago

Any insight on D for the answer?

upvoted 1 times

> **Not_That_Guy** 2 years, 2 months ago
>
> A VPN tunnel would provide an opportunity for encryption and a virtual point-to-point connection for easier IP addressing, but the actual traffic would still flow through the same intermittent switches/routers. The additional overhead from the encryption and encapsulation would increase delay if anything.
>
> upvoted 1 times

ryanzou 2 years, 3 months ago

Selected Answer: C

Correct answer is C. decrease the MTU will decrease the packet size in the channel, the transmission time will increase.

upvoted 2 times

ironman_86 2 years, 3 months ago

I think the answer is C

upvoted 1 times

u2637ps 2 years, 10 months ago

depending on the application reducing the MTU will actually increase delay

upvoted 1 times

Question #59 *Topic 1*

An SQL injection vulnerability was reported on a web application, and the cloud platform team needs to mitigate the vulnerability while it is corrected by the development team. Which of the following controls will BEST mitigate the risk of exploitation?

A. DLP

B. HIDS

C. NAC

D. WAF

Suggested Answer: *D*

Community vote distribution

D (92%) 8%

uzey 4 months, 2 weeks ago

Selected Answer: D

A Web Application Firewall (WAF) is the most effective immediate control to mitigate an SQL injection vulnerability. It can inspect and filter incoming web traffic, detecting and blocking malicious SQL injection attempt

upvoted 1 times

yyCherubim 1 year, 1 month ago

HIDS huh? Because it did such a good job at detecting the first SQL Injection, that we should rely on it this time?

upvoted 2 times

reto1 3 months, 2 weeks ago

HIDS (Host Intrusion Detection System): Monitors a single host for malicious activity but may not effectively prevent SQL injection attacks on a web application level.

upvoted 1 times

maelo 1 year, 7 months ago

Selected Answer: B

WAF = web application FW. Wikipedia: "specific form of application firewall that filters, monitors, and blocks HTTP traffic to and from a web service". This doesn't match SQL type.

upvoted 1 times

kuzummjakk 9 months, 4 weeks ago

The SQL vulnerability was detected on a WEB APPLICATION. SQL's in the backend, but it's interacted with via HTTP.

a HIDS only DETECTS intrusions, it does nothing about it and seeing as the SQL injection is made through interacting with the web application and not intruding the host, it makes even less sense.

upvoted 1 times

BeauChateau 1 year, 8 months ago

Selected Answer: D

D. WAF (Web Application Firewall)

A WAF is the best control to mitigate the risk of SQL injection vulnerabilities while the development team fixes the issue. A WAF can identify and block SQL injection attacks by analyzing the traffic between the application and the user. It can also help to protect against other types of attacks that exploit web application vulnerabilities.

upvoted 4 times

bagsik89 1 year, 10 months ago

Selected Answer: D

WAF is the best technical control against SQL Injection

upvoted 1 times

JVen 2 years, 1 month ago

Selected Answer: D

This should be D

upvoted 3 times

Not_That_Guy 2 years, 2 months ago

Selected Answer: D

Clearly WAF

upvoted 2 times

ryanzou 2 years, 3 months ago

Selected Answer: D

DO DOUBT, it is WAF.

upvoted 1 times

i_bird 2 years, 3 months ago

Even the nugget of information given with the answer point to WAF...

upvoted 4 times

achow26 2 years, 3 months ago

Answer should be D.

upvoted 2 times

ironman_86 2 years, 3 months ago

Why not D? HIDS will only detect and will not prevent the exploitation.

upvoted 2 times

> **maelo** 1 year, 7 months ago
>
> HIDS = host-based IDS/IPS - Intrusion Detection/*Prevention* System
>
> upvoted 1 times
>
> > **Locy333** 10 months, 3 weeks ago
> >
> > HIDS and HIPS are different systems. HIDS detects only, if the answer stated HIPS then it would be a viable prevention option.
> >
> > upvoted 1 times

Question #60 *Topic 1*

To save on licensing costs, the on-premises, IaaS-hosted databases need to be migrated to a public DBaaS solution. Which of the following would be the BEST technique?

A. Live migration

B. Physical-to-virtual

C. Storage-level mirroring

D. Database replication

Suggested Answer: *D*

Community vote distribution

D (86%) 14%

⊟ **Agr321** Highly Voted 👍 2 years, 3 months ago

Selected Answer: D

You are migrating the database to save on licensing, Not migrating a Server. Migrating to vendor host database.

upvoted 6 times

⊟ **reto1** 3 months, 2 weeks ago

ChatGPT approved

upvoted 1 times

⊟ **BeauChateau** Highly Voted 👍 1 year, 8 months ago

Selected Answer: D

D. Database replication

Database replication would be the best technique to migrate on-premises IaaS-hosted databases to a public DBaaS solution while saving on licensing costs. Database replication involves copying a database from one location to another, which can be done in near real-time. This allows the data to be continuously synchronized between the on-premises database and the public DBaaS solution, without disrupting operations or requiring a complete migration all at once. Additionally, database replication can be done with minimal downtime and minimal impact to the performance of the database.

upvoted 5 times

⊟ **kuzummjakk** Most Recent 🕘 9 months, 4 weeks ago

Selected Answer: D

Whoever said an on-premise IaaS DB is physical need to go back to study.

upvoted 1 times

⊟ **Pongsathorn** 1 year, 3 months ago

Selected Answer: D

To save on licensing costs and migrate on-premises databases to a public Database as a Service (DBaaS) solution, the BEST technique would be:

D. Database replication.

Here's why:

A. Live migration typically involves moving a virtual machine or application from one host or environment to another with minimal downtime. It might not be directly applicable when migrating to a DBaaS, as DBaaS solutions often require a different approach.

B. Physical-to-virtual (P2V) migration is used when converting physical servers to virtual machines. This is not relevant when migrating to a DBaaS, as DBaaS typically abstracts the underlying hardware.

upvoted 1 times

⊟ **Pongsathorn** 1 year, 3 months ago

C. Storage-level mirroring is a technique for redundancy and disaster recovery, not a method for migrating databases to a DBaaS platform.

D. Database replication involves creating a copy of the database in the target environment, and it's a common technique for database migrations. Depending on the DBaaS provider and the database technology being used, various replication methods (e.g., snapshot replication, transactional replication) can be employed to ensure minimal downtime during migration.

Database replication allows for a smooth transition from on-premises to DBaaS while minimizing downtime and ensuring data consistency, making it the most suitable option for this scenario.

upvoted 1 times

Zak11 1 year, 8 months ago

Selected Answer: D

Live migration and physical-to-virtual are techniques used for migrating VMs, not databases. Storage-level mirroring is a data protection technique and is not used for database migration.

upvoted 2 times

Lenell 2 years ago

Selected Answer: B

The key principle is "on-premises migration to DBaas." DBaas is less expensive than any IaaS whether the "Infrastructure" is in the cloud or on-premise. On the use of "Iaas" in the question is a distraction. The use of a DB doesn't change but the infrastructure does; i.e., physical-to-virtual. Comptia Cloud+ mindset....

upvoted 3 times

kuzummjakk 9 months, 4 weeks ago

on-premise IaaS means private cloud

upvoted 2 times

Agr321 2 years, 2 months ago

Selected Answer: D

https://aws.amazon.com/getting-started/hands-on/move-to-managed/migrate-sql-server-to-amazon-rds/

upvoted 3 times

jiminycriminal 2 years, 3 months ago

on-prem and IaaS is kinda contradictory aint it?

upvoted 3 times

kuzummjakk 9 months, 4 weeks ago

dont forget about private clouds

upvoted 2 times

Lenell 2 years ago

IaaS is a distraction. The key is that it is on-premise. By nature any "on-premise" network is IaaS to the organizational users.

upvoted 1 times

brickcity86 2 years ago

That's what im saying what kind of word salad is this question?

upvoted 2 times

achow26 2 years, 3 months ago

The current servers are IaaS so it cannot be P2V. Answer should be D. Provision the instance on DBaaS and enable replication.

upvoted 5 times

i_bird 2 years, 3 months ago

Physical hardware underlays every IAAS system, whether on-prem provided or CSP provided.

upvoted 1 times

maelo 1 year, 8 months ago

Any *aaS has HW underneath. Service provider abstracts HW even in IaaS. No physical asset is directly touched in the outlined move.

upvoted 1 times

jiminycriminal 2 years, 3 months ago

Database Replication is about redundancy, not migration.

upvoted 2 times

brickcity86 2 years ago

DB replication can copy the functioning DB over to the public environment in real-time while keeping the source DB operational

upvoted 2 times

Question #61 *Topic 1*

A systems administrator swapped a failed hard drive on a server with a RAID 5 array. During the RAID resynchronization, a second hard drive failed. Which of the following actions will make the server fully operational?

A. Restart the RAID resynchronization process.

B. Perform a P2V migration of the server.

C. Swap the failed hard drive with a fresh one.

D. Restore the server from backup.

Suggested Answer: *D*

Community vote distribution

D (97%)

JVen Highly Voted 👍 1 year, 7 months ago

Selected Answer: D

Yeah this has to be D. Cant have two drives fail, even if it was rebuilding at the time. Rebuild has to finish before another can fail in RAID 5.

upvoted 9 times

BeauChateau Highly Voted 👍 1 year, 2 months ago

Selected Answer: D

D. Restore the server from backup.

With two hard drives failed on a RAID 5 array, the server will not be fully operational until the data is restored from a backup. RAID 5 arrays are designed to tolerate a single drive failure, but if a second drive fails during the resynchronization process, the array will be unable to rebuild and the data will be lost.

upvoted 6 times

TheFivePips Most Recent 1 month, 1 week ago

Selected Answer: D

RAID 0 provides striping without parity or mirroring. Minimum Drives: 2. Data is distributed across multiple drives for improved performance, but there is no redundancy. If one drive fails, all data in the array is lost.

RAID 1 provides mirroring without striping or parity. Minimum Drives: 2. Data is duplicated across multiple drives, providing redundancy. If one drive fails, data remains accessible from the mirrored drive.

RAID 5 uses block-level striping with distributed parity. Minimum Drives: 3. Data and parity information are distributed across all drives in the array. RAID 5 can tolerate the failure of one drive without data loss.

RAID 10 combines mirroring and striping. Minimum Drives: 4. Data is mirrored across pairs of drives, and then the mirrored sets are striped for performance.

RAID 6 uses dual parity for fault tolerance. Minimum Drives: 4. It can withstand the failure of up to two drives without data loss. RAID 6 provides greater fault tolerance than RAID 5 but has slightly lower performance and higher storage overhead due to the additional parity information.

upvoted 3 times

TheFivePips 1 month, 1 week ago

this was just a reminder for people like me who constantly forget the RAID stuff

upvoted 2 times

Sweety_Certified7 3 months ago

Selected Answer: D

D. Restore the server from backup.

While it's true that RAID 5 arrays can tolerate a single drive failure, they become vulnerable to data loss if a second drive fails before the array has finished rebuilding after the first failure.

With two hard drives failed in a RAID 5 array during the resynchronization process, the array would be unable to rebuild, resulting in data loss. Restoring the server from backup is necessary to recover the lost data and make the server fully operational again.

upvoted 2 times

kuzummjakk 3 months, 3 weeks ago

why does it say "swapped a failed hard drive WITH a RAID 5 array". does it mean ON a RAID 5 array?

upvoted 1 times

yyCherubim 7 months, 3 weeks ago

Selected Answer: D

I think they are confusing RAID-5 with RAID-6.

I personally know resync will not work on RAID-5 with two failed drives ....it happened on my home RAID-5 NAS. The difference is I had no backups, and yes, it sucks to be me.

upvoted 1 times

SecPlus2022 1 year, 1 month ago

Selected Answer: D

As long as the RAID5 array is in the rebuild process it will be in a degraded state. The failure of a drive in said array while it is in a degraded state will render the data lost. D is the correct answer.

upvoted 1 times

bagsik89 1 year, 4 months ago

Selected Answer: C

C. Swap the failed hard drive with a fresh one.

Raid 5 can only survive one failed drive.

I don't think it's D as it doesn't change the fact that you have a failed drive.

upvoted 1 times

scott5010 1 year, 7 months ago

Selected Answer: D

Answer is D

upvoted 3 times

Agr321 1 year, 8 months ago

Selected Answer: D

Other comments are correct.

D is the answer. Restore

upvoted 3 times

Not_That_Guy 1 year, 8 months ago

Selected Answer: D

RAID 5 can only tolerate 1 drive failure. Best answer id D.

upvoted 2 times

Rob69420 1 year, 9 months ago

If a second disk in a RAID level 5 disk array fails, you must replace the failed disks, then delete and recreate the disk array. You must then recreate the file systems on the disk array and copy data to the restored disk array from your backup media.

Regardless of how many drives are in use, a RAID 5 array only allows for recovery in the event that just one disk at a time fails.

upvoted 1 times

ironman_86 1 year, 10 months ago

How is the answer is A?! I think it's D.

upvoted 3 times

> **jiminycriminal** 1 year, 9 months ago
>
> Yeah if two drives fail in RAID5, aren't you done for? People go for RAID6 now partly because of this.
>
> upvoted 1 times

Question #62 *Topic 1*

A cloud administrator recently deployed an update to the network drivers of several servers. Following the update, one of the servers no longer responds to remote login requests. The cloud administrator investigates the issue and gathers the following information:
✑ The cloud management console shows the VM is running and the CPU and memory utilization is at or near 0%.
✑ The cloud management console does not show an IP address for that server.
✑ A DNS lookup shows the hostname resolves to an IP address.
✑ The server is a member of the same security group as the others.
✑ The cloud administrator is able to log in remotely to the other servers without issue.
Which of the following is the MOST likely cause of the server being unavailable?

A. The network driver updates did not apply successfully, and the interface is in a down state.

B. The ACL policy for the server was updated as part of the server reboot, preventing login access.

C. The server was assigned a new IP address, and DNS entry for the server name was not updated.

D. The update caused an increase in the output to the logs, and the server is too busy to respond.

Suggested Answer: *A*

Community vote distribution

A (100%)

⊟ **KairKnows** Highly Voted 👍 2 years, 6 months ago

Answer should be C. It is clearly hinted at in the question, and why would only 1 server be affected if it was a network driver issue.

upvoted 6 times

⊟ **reto1** 3 months, 2 weeks ago

The server was assigned a new IP address: If this were the case, you would typically see the server listed with a new IP in the management console, and the DNS would be expected to update, which is not indicated here.

upvoted 1 times

⊟ **kuzummjakk** 9 months, 4 weeks ago

Re-read bullet point 2

upvoted 1 times

⊟ **jiminycriminal** 2 years, 3 months ago

Absolutely not. If the IP was changed, the management console would show it regardless. The network state is down. Sometimes a specific server will fail at something in a batch process for various reasons.

upvoted 2 times

⊟ **TheFivePips** Most Recent ⊙ 7 months, 2 weeks ago

Selected Answer: A

A. Server unresponsiveness to remote login requests and absence of the server's IP address in the cloud management console—are consistent with a network interface issue. If the network driver updates failed to apply correctly, it could result in the network interface being in a down state, effectively rendering the server unreachable over the network.

C. The server was assigned a new IP address, and the DNS entry for the server name was not updated.

While plausible, the fact that DNS resolution correctly resolves the hostname to an IP address suggests that DNS is functioning correctly. If the server's IP address had changed, DNS resolution would likely fail. Additionally, DHCP typically handles DNS updates automatically, reducing the likelihood of a misconfigured DNS entry.

upvoted 1 times

⊟ **BeauChateau** 1 year, 8 months ago

Selected Answer: A

A. The network driver updates did not apply successfully, and the interface is in a down state.

Based on the given information, the most likely cause of the server being unavailable is that the network driver updates did not apply successfully, and the interface is in a down state. The fact that the cloud management console does not show an IP address for the server suggests that the network interface is not functioning correctly. The fact that a DNS lookup shows the hostname resolving to an IP address indicates that DNS is working correctly, but the server is not responding to network requests.

upvoted 3 times

⊟ **achow26** 2 years, 3 months ago

Answer should be A. Question clearly states "The cloud management console does not show an IP address for that server."

upvoted 1 times

⊟ **Granddude** 2 years, 6 months ago

Selected Answer: A

I am going with A because the network drivers were just updated. I think the update failed on this server alone. Updating the drivers shouldn't cause DNS to assign a new IP address.

Most if not all VMs I have seen in real-world have their IP addresses statically assigned.

upvoted 3 times

⊟ **nate612** 2 years, 2 months ago

Yeah static addressing is best practice.

upvoted 1 times

Question #63 *Topic 1*

A company is planning to migrate applications to a public cloud, and the Chief Information Officer (CIO) would like to know the cost per business unit for the applications in the cloud. Before the migration, which of the following should the administrator implement FIRST to assist with reporting the cost for each business unit?

A. An SLA report

B. Tagging

C. Quotas

D. Showback

Suggested Answer: *B*

Community vote distribution

B (70%) D (30%)

☐ **uzey** 4 months, 2 weeks ago

Selected Answer: B

To accurately allocate costs to specific business units in a cloud environment, it's essential to implement a robust tagging system. By applying relevant tags to cloud resources, such as applications, servers, and storage, the administrator can track and allocate costs based on business unit, department, or project

upvoted 3 times

☐ **reto1** 3 months, 2 weeks ago

Tagging allows resources in the cloud to be labeled with specific identifiers that can include details such as business unit, project name, or environment type. By applying tags to cloud resources, the company can effectively track usage and costs associated with each business unit.

upvoted 2 times

☐ **Me_Me_Me** 11 months, 1 week ago

Selected Answer: B

Tags are metadata or labels that can be attached to cloud resources, such as VMs, storage, or networks. Tags can help organize, identify, and manage cloud resources, as well as track their usage and costs. Tags can also be used to implement chargeback or showback policies, which are methods of allocating the cloud services bill to different departments or consumers based on their consumption of resources .

Tags are needed to produce showbacks.

upvoted 1 times

☐ **FrancisDrake** 12 months ago

FIRST is in all caps. These questions are tricky.

upvoted 1 times

☐ **yyCherubim** 1 year, 1 month ago

Selected Answer: D

Lookup what "Showback" is and you will not be confused why it is the answer.

upvoted 3 times

☐ **Princee450** 1 year, 4 months ago

"to assist with reporting the cost for each business unit" would mean its a chargeback and not a showback. You need tagging to correctly use a chargeback so you can identify each department cost

Answer is tagging

upvoted 1 times

☐ **backdooranon** 1 year, 2 months ago

showback also does that, the difference is chargeback requires the unit to actually pay for the costs while showback doesnt

upvoted 4 times

☐ **Tomtom11** 1 year, 5 months ago

Showbacks offer departmental visibility into IT resource usage without charging departments for their use.

When a company implements an IT chargeback policy, they bill each department for the number of technology resources used for a given period

The Answer is Tagging

upvoted 2 times

⊟ **sheilawu** 1 year, 5 months ago

I go for B because it make more sence if it defines the "unit"

upvoted 1 times

⊟ **SecPlus2022** 1 year, 6 months ago

Selected Answer: B

Need to change my answer. Showback is a process of reporting cost and usage information to departments and users. You implement tagging whereas showback is a process. Showback provides usage and costs and they're looking to obtain this information and you need tagging in order to do so.

upvoted 2 times

⊟ **SecPlus2022** 1 year, 6 months ago

Selected Answer: D

The applications have not yet moved to the cloud. You can't add tagging to a cloud environment that doesn't exist yet. To learn the costs BEFORE the migration.

upvoted 2 times

⊟ **BeauChateau** 1 year, 8 months ago

Selected Answer: B

B. Tagging

Implementing tagging would be the first step to assist with reporting the cost for each business unit. Tagging resources with labels that indicate which business unit or project they belong to will enable the cloud administrator to easily separate the cost of the resources used by each unit or project.

An SLA (Service Level Agreement) report would not be useful for reporting the cost for each business unit, as SLA reports focus on availability and performance metrics.

Quotas are used to set limits on the amount of resources that can be consumed by a particular user or group, but they would not provide information about the cost of the resources used.

Showback is a reporting mechanism that provides visibility into the cost of cloud resources used by each business unit or department, but it requires accurate cost data, which can only be obtained through the implementation of tagging or another cost allocation mechanism.

upvoted 3 times

⊟ **concepcionz** 1 year, 9 months ago

Selected Answer: B

I'll go with B

"Tagging helps you manage your resources more effectively by enabling you to categorize them in different ways, such as by purpose, owner, environment, or other criteria. By tagging resources with business unit information, you can track costs across business units and report on usage and costs by business unit"

upvoted 1 times

⊟ **Trebor28** 1 year, 10 months ago

Selected Answer: D

D. before migration

upvoted 1 times

⊟ **beamage** 1 year, 11 months ago

Before the migration, Sorry sticking with D

upvoted 2 times

> ⊟ **LeDarius3762** 1 year, 10 months ago
>
> Sticking with D too because is before migration and he wants to know how much money is per business unit, suitable for showback totally. Tagging is good too just if you're already on the cloud
>
> upvoted 1 times

⊟ **TheGinjaNinja** 1 year, 11 months ago

Selected Answer: B

B. Tagging is the most suitable option to implement first as it will allow the administrator to assign specific metadata to resources such as VMs, storage, and network interfaces, so that they can be grouped and tracked by business unit. This will help in generating the cost per business unit report that the CIO is requesting. Tagging can be done before the migration, so the administrator can ensure that the cost can be tracked and reported correctly.

upvoted 4 times

Question #64 *Topic 1*

A cloud administrator would like to deploy a cloud solution to its provider using automation techniques. Which of the following must be used? (Choose two.)

A. Auto-scaling

B. Tagging

C. Playbook

D. Templates

E. Containers

F. Serverless

Suggested Answer: *CD*

Community vote distribution

CD (100%)

⊟ **jiminycriminal** Highly Voted 👍 2 years, 3 months ago

C and D. Playbooks and Templates. Auto-scaling is not a technique for automated deployment of a service, it's just a feature.

https://aws.amazon.com/blogs/infrastructure-and-automation/automate-ansible-playbook-deployment-amazon-ec2-github/

"Ansible is an open-source automation tool that uses playbooks to enable you to make deployments faster and scale to various environments"

upvoted 12 times

⊟ **reto1** 3 months, 2 weeks ago

C. Playbook: In the context of automation, a playbook typically refers to a configuration management tool (like Ansible) that automates tasks and deployments. It defines the steps needed to configure and manage cloud resources.

D. Templates: Templates (such as those used in Infrastructure as Code frameworks) allow for the predefined configuration of cloud resources. They enable quick and consistent deployments, which are essential for automation.

upvoted 1 times

⊟ **kuzummjakk** Most Recent ⏲ 9 months, 4 weeks ago

Selected Answer: CD

some quirky AI putting B

upvoted 1 times

⊟ **Zak11** 1 year, 8 months ago

Selected Answer: CD

Templates provide a pre-defined and reusable framework that can be quickly deployed to create an environment for the cloud solution. Playbooks are a set of procedures that automate the deployment of the cloud solution by defining the tasks to be executed in a particular sequence.

upvoted 4 times

⊟ **BeauChateau** 1 year, 8 months ago

Selected Answer: CD

D. Templates and C. Playbook

To deploy a cloud solution to a provider using automation techniques, templates and playbooks must be used. Templates are pre-configured blueprints that define the infrastructure and resources needed for an application or environment, and can be used to automate the deployment of cloud resources. Playbooks are used to automate the configuration of software and services on cloud resources, and can be used to ensure consistency and reduce manual effort.

upvoted 1 times

⊟ **Alizadeh** 1 year, 8 months ago

Selected Answer: CD

C. Playbook

D. Templates

upvoted 2 times

⊟ **achow26** 2 years, 3 months ago

It should be C and D. Had the options been Images then yes Images and Templates but option A is Auto-scaling which is not appropriate for this situation.

upvoted 2 times

⊟ **MarcusWG86** 2 years, 4 months ago

A & D

Templates and Images:

If you need redundant VMs with a particular configuration or specific software, you can create your own custom reusable template. Just configure a VM the way you need it and then use that image to create your other VMs. Also, custom images go hand in hand with autoscaling. Create a custom template image with your application already installed, and as the workload increases, autoscaling can provision new, preconfigured VMs from your template.

CompTIA Cloud+ Study Guide—Exam CV0-003, Third Edition

ISBN: 9781119810865

upvoted 4 times

⊟ **Agr321** 2 years, 2 months ago

They may have changed the available answers as IMAGES is not available choice now.

C&D are correct

A. Auto-scaling

B. Tagging

C. Playbook

D. Templates

E. Containers

F. Serverless

upvoted 3 times

Question #65 *Topic 1*

A systems administrator needs to configure monitoring for a private cloud environment. The administrator has decided to use SNMP for this task. Which of the following ports should the administrator open on the monitoring server's firewall?

A. 53

B. 123

C. 139

D. 161

Suggested Answer: *D*

Community vote distribution

D (100%)

TheFivePips 1 month, 1 week ago

Selected Answer: D

Port 53 is used for the Domain Name System (DNS) service. DNS is responsible for translating domain names (e.g., www.example.com) into IP addresses (e.g., 192.0.2.1) and vice versa.

Port 123 is used for the Network Time Protocol (NTP) service. NTP is used for synchronizing the time of networked devices, ensuring accurate timekeeping across systems. It helps maintain synchronization between servers, workstations, routers, switches, and other networked devices.

Port 139 is used for the NetBIOS Session Service, which is part of the Server Message Block (SMB) protocol suite. It is commonly associated with file and printer sharing in Windows environments. Port 139 facilitates communication between devices for accessing shared files, directories, and printers over a network.

Port 161 is used for the Simple Network Management Protocol (SNMP) service. SNMP is used for monitoring and managing network devices, such as routers, switches, servers, and printers. It allows network administrators to collect information about device status, performance, and configuration through SNMP agents running on managed devices.

upvoted 2 times

kuzummjakk 3 months, 3 weeks ago

wait, we're gonna get port questions?

upvoted 1 times

Pongsathorn 9 months, 3 weeks ago

Selected Answer: D

To enable SNMP monitoring in a private cloud environment, you should open **port 161** on the firewall of the monitoring server. SNMP (Simple Network Management Protocol) uses this port for communication between the monitored devices (in this case, your private cloud resources) and the SNMP manager (the monitoring server).

So, the correct answer is **D. 161**.

upvoted 2 times

Question #66 *Topic 1*

A cloud administrator is switching hosting companies and using the same script that was previously used to deploy VMs in the new cloud. The script is returning errors that the command was not found. Which of the following is the MOST likely cause of the script failure?

A. Account mismatches

B. IP address changes

C. API version incompatibility

D. Server name changes

Suggested Answer: *C*

Community vote distribution

C (100%)

dvd21 Highly Voted 3 years, 1 month ago

it's the API compatibility

upvoted 8 times

reto1 3 months, 2 weeks ago

C

When switching hosting companies, the new cloud provider may have a different API version or set of commands compared to the previous provider. If the script was written using commands specific to the old provider's API, it might not be compatible with the new provider's API, leading to errors such as "command not found."

upvoted 1 times

kuzummjakk Most Recent 9 months, 4 weeks ago

dear exam topics,

who is closing their eyes and selecting random answers as the "correct answer". Please tell them, keep it up. im on my toes

upvoted 1 times

eacunha 12 months ago

Selected Answer: C

C. Incompatibilidade de versão da API

Se o administrador de nuvem estiver usando um script que interage com a API da nuvem para realizar operações, mudanças na versão da API entre a antiga e a nova empresa de hospedagem podem resultar em comandos não encontrados ou erros relacionados à incompatibilidade de versão. Recomenda-se verificar a documentação da API da nova empresa de hospedagem para garantir que o script esteja alinhado com a versão correta da API.

upvoted 1 times

sheilawu 1 year, 5 months ago

Selected Answer: C

should be C not D

upvoted 2 times

SecPlus2022 1 year, 6 months ago

Selected Answer: C

It states "command not found" not "server not found" so the answer is "C", not "D".

upvoted 1 times

Alizadeh 1 year, 8 months ago

Selected Answer: C

C. API version incompatibility

upvoted 1 times

Frikandel 1 year, 9 months ago

Has to be C. it says clearly that "the command was not found"

upvoted 1 times

TheGinjaNinja 1 year, 11 months ago

Selected Answer: C

API Version

upvoted 2 times

achow26 2 years, 3 months ago

Answer should be A API Incompatibility. It says server name and not naming convention.

upvoted 3 times

MarcusWG86 2 years, 4 months ago

It is D, here is an example from GCP

Passing a Windows startup script from a local file to a new VM

Create a VM and pass the contents of a local file to be used as the startup script by using the gcloud compute instances create command with the --metadata-from-file flag:

```
gcloud compute instances create VM_NAME \
--image-project=windows-cloud \
--image-family=windows-2019-core \
--metadata-from-file=windows-startup-script-ps1=FILE_PATH
```

Replace the following:
VM_NAME: the name of the VM
FILE_PATH: the relative path to the startup script file

https://cloud.google.com/compute/docs/instances/startup-scripts/windows

upvoted 2 times

u2637ps 2 years, 10 months ago

Cloud providers have specific naming conventions

upvoted 1 times

u2637ps 2 years, 10 months ago

if the name is the same in a new cloud why would it fail?

upvoted 1 times

Question #67 *Topic 1*

A systems administrator is configuring network management but is concerned about confidentiality. Which of the following should the administrator configure to address this concern?

A. SNMPv3

B. Community strings

C. IPSec tunnels

D. ACLs

Suggested Answer: *A*

Community vote distribution

A (95%) 5%

jiminycriminal Highly Voted 2 years, 3 months ago

Selected Answer: A

I think it should be A - SNMP v3. Keywords: network management. SNMP v3 is encrypted. IPSec is too general of an answer.

upvoted 9 times

reto1 3 months, 2 weeks ago

SNMPv3 (Simple Network Management Protocol version 3) includes built-in security features such as authentication and encryption, which help ensure confidentiality and integrity of the management data being transmitted. This is critical for protecting sensitive information in network management.

upvoted 1 times

Rob69420 Highly Voted 2 years, 3 months ago

What are the three important services of SNMPv3 security model?

SNMPv3 protects against these threats by providing the security services of Data Integrity, Authentication, Privacy (Confidentiality), and Message Timeliness .

What does IPsec protect against?

IPsec is used for protecting sensitive data, such as financial transactions, medical records and corporate communications, as it's transmitted across the network. It's also used to secure virtual private networks (VPNs), where IPsec tunneling encrypts all data sent between two endpoints.

Pick your poison?

upvoted 5 times

kuzummjakk Most Recent 9 months, 4 weeks ago

Selected Answer: A

ah yeah keyphrase "network management"

upvoted 1 times

utied 1 year ago

Selected Answer: A

The question is about configuring network management but is concerned about confidentiality.

A. SNMPv3: Simple Network Management Protocol. v3 is encrypted. Gathers data(statistics) from agents running on other devices.

B. Community strings: Is a means of SNMP accessing statistics stored within a router or other device.

C. IPSec tunnels: encrypt all traffic in tunnel.

D. ACL: permissions to control access to resources.

upvoted 3 times

irazak 1 year, 2 months ago

SNMPv3 focuses on securing SNMP communication specifically, which may not cover other network communication within the same device. IPsec secures network communication at the IP layer, protecting all traffic between endpoints, not just specific protocols.

upvoted 1 times

Pongsathorn 1 year, 3 months ago

Selected Answer: A

To address concerns about confidentiality in network management, the administrator should configure **SNMPv3 (Simple Network Management Protocol version 3)**. SNMPv3 provides authentication and encryption, ensuring the confidentiality and integrity of SNMP messages. It allows you to set up user authentication and access control, making it the most secure option among the ones listed.

So, the correct answer is **A. SNMPv3**.

upvoted 1 times

concepcionz 1 year, 9 months ago

Selected Answer: A

"... network management."

upvoted 2 times

rob88Silva 1 year, 11 months ago

Selected Answer: C

IPSec (Internet Protocol Security) is a set of protocols that provide security for IP communications by authenticating and encrypting each IP packet in a communication session. By using IPSec tunnels, the administrator can secure the SNMP traffic between the management station and the managed devices, ensuring that the information being exchanged cannot be intercepted or tampered with by unauthorized parties.
SNMPv3 is a version of SNMP that provides more robust security features than previous versions, such as authentication and encryption, but it's not enough to secure the communication.

upvoted 1 times

TheFivePips 7 months, 2 weeks ago

SNMPv3 (Simple Network Management Protocol version 3) offers comprehensive security features, including authentication, encryption, and access control. It supports message integrity through HMAC (Hash-based Message Authentication Code) and data encryption using protocols like AES (Advanced Encryption Standard). SNMPv3 is the most secure version of SNMP and is recommended for protecting the confidentiality of management data.
It is also the only option that has anything to do with network management.

upvoted 1 times

Brianhealey136 1 year, 12 months ago

Selected Answer: A

Could easily be both A and C. I'm going with A. Dumb question.

upvoted 2 times

TheFivePips 7 months, 2 weeks ago

only if you miss the network management part

upvoted 1 times

Question #68 Topic 1

Which of the following will provide a systems administrator with the MOST information about potential attacks on a cloud IaaS instance?

A. Network flows

B. FIM

C. Software firewall

D. HIDS

Suggested Answer: *A*

Community vote distribution

D (50%) A (50%)

i_bird Highly Voted 2 years, 3 months ago

Answer is HIDS
cloud IaaS instance aka a host or node or vm.

upvoted 15 times

reto1 3 months, 2 weeks ago

HIDS monitors the host for suspicious activities and provides detailed insights into potential attacks, including unauthorized access attempts, changes to system files, and other indicators of compromise. It analyzes logs and system behavior, offering alerts and reports on potential threats specific to that instance.

upvoted 1 times

KOINU7 Highly Voted 1 year, 11 months ago

I am going with A because it says potential attacks. HIDS is detection of current attack.

upvoted 5 times

uzey Most Recent 4 months, 2 weeks ago

Selected Answer: A

Network flows provide the most comprehensive and detailed information about potential attacks on a cloud IaaS instance. By analyzing network traffic patterns, anomalies, and deviations from normal behavior, a system administrator can identify potential threats such as port scans, DDoS attacks, and unauthorized access attempts.

upvoted 1 times

Sweety_Certified7 9 months, 1 week ago

Selected Answer: D

HIDS monitors activity on individual hosts, including file system changes, log entries, and system calls, to detect and respond to potential security threats. It can provide valuable insights into attacks targeting specific hosts within the IaaS environment.

And, the question does not explicitly refer to network-related information, HIDS would be the correct answer. Not A: Network flows.

upvoted 2 times

kuzummjakk 9 months, 4 weeks ago

Selected Answer: D

Absolutely HIDS. You CAN use network flows in some cases, but that's not it's purpose. It's a HIDS's purpose.

upvoted 1 times

germancano14 1 year, 3 months ago

Selected Answer: D

HIDS.
Chatgpt says network flow but HIDS is specifically for the host it is installed and the question is talking about 1 specific instance.
Option D, Host-Based Intrusion Detection System (HIDS), is indeed an important security tool, but it primarily focuses on monitoring activities and events at the host or individual instance level. While it can provide valuable information about potential attacks or anomalies occurring on a specific system, it may not offer as comprehensive information about attacks on a cloud IaaS instance as network flow analysis does.

upvoted 2 times

Francois1984 1 year, 4 months ago

Selected Answer: D

"Intrusion Detection/Intrusion Prevention System (IDS/IPS)" would provide a systems administrator with the MOST information about potential attacks on a cloud Infrastructure as a Service (IaaS) instance. (Copied from ChatGPT)

upvoted 3 times

Sunshine_boy38 1 year, 5 months ago

My answer goes for HIDS (D) - HIDS is designed to monitor and analyze activities and events occurring on a specific host or server. To provide valuable information about potential attacks by analyzing system logs, detecting unauthorized access attempts, identifying unusual or suspicious behavior, and raising alerts for known attack patterns or signatures. It can help the administrator identify and respond to security incidents promptly, protecting the cloud IaaS instance and its resources.

However, choice for network flows only focuses on monitoring network traffic and analyzing network flows for anomalies. While network flow monitoring can provide insights into network activity, it may not provide comprehensive information about potential attacks targeting the IaaS instance itself.

upvoted 1 times

BeauChateau 1 year, 8 months ago

Selected Answer: A

A. Network flows.

Network flows provide the most information about potential attacks on a cloud IaaS instance because they capture all traffic between the instance and the network, including both inbound and outbound traffic. Network flows provide insights into network behavior, such as communication patterns, traffic volume, and protocols used. This information can be used to detect anomalies, such as a sudden increase in traffic or traffic from unexpected sources, which could indicate a potential attack.

upvoted 5 times

TheFivePips 7 months, 2 weeks ago

HIDS is literally designed to monitor and analyze activity on individual host systems, looking for signs of suspicious behavior or known attack patterns. It would give you much much more information that is RELEVANT to security.

upvoted 1 times

kuzummjakk 9 months, 4 weeks ago

give us your comptia answer, not your chatgpt answer

upvoted 1 times

Stella02 1 year, 11 months ago

Selected Answer: D

It says IAAS instance. HIDS is the correct answer.

upvoted 1 times

TheGinjaNinja 1 year, 11 months ago

Selected Answer: A

Answer is HIDS

upvoted 1 times

Not_That_Guy 2 years, 2 months ago

Selected Answer: A

HIDS is more PaaS and just alerts to potential threats, Network Flows gives the most information for IaaS.

upvoted 2 times

Question #69 *Topic 1*

A cloud administrator is designing a multiregion network within an IaaS provider. The business requirements for configuring the network are as follows:

✑ Use private networking in and between the multisites for data replication.

✑ Use low latency to avoid performance issues.

Which of the following solutions should the network administrator use within the IaaS provider to connect multiregions?

A. Peering

B. Gateways

C. VPN

D. Hub and spoke

Suggested Answer: *A*

Community vote distribution

A (74%) D (17%) 4%

kuzummjakk 9 months, 4 weeks ago

Selected Answer: D

It's A or D. NOT VPN because the question says "use private networking" so you don't need to VPN into it (generally used if you're going over a public network). Peering is probably more point A to point B and less multi-site. But you can argue that you can just peer all the sites together, but the "comptia" answer is "you can argue anything, but what's more relevant". So I'm picking D

upvoted 1 times

kuzummjakk 9 months, 4 weeks ago

Peering is more cloud-related though. Cloud peering is less "point to point" centric and more "open connection" centric.

upvoted 1 times

FrancisDrake 12 months ago

Selected Answer: A

I believe peering provides a direct physical connection. And can be configured for privacy. So you get performance and security. VPN would seem to have more overhead attached to it.

upvoted 4 times

reto1 3 months, 2 weeks ago

A is correct.

Peering allows for direct, private connections between different regions within an IaaS environment. This method typically offers lower latency compared to other options and facilitates efficient data replication without traversing the public internet.

upvoted 1 times

FrancisDrake 12 months ago

BUT I'm not sure that peering would be used for multiple sites. Site to site seems to be more the norm. Hmmm...

upvoted 1 times

VVV4WIN 1 year, 1 month ago

Selected Answer: D

Hub and spoke for multiregion and low latency

upvoted 3 times

Pongsathorn 1 year, 3 months ago

Selected Answer: A

To meet the business requirements of using private networking for data replication between multiple regions and ensuring low latency, the network administrator should use **Peering** within the IaaS provider.

Peering allows direct and private connectivity between virtual networks or regions within the same cloud provider's infrastructure. It provides low-latency communication between different regions or sites while keeping the traffic within the provider's network, ensuring data privacy.

So, the correct answer is **A. Peering**.

upvoted 1 times

Francois1984 1 year, 4 months ago

Selected Answer: C

Based on the provided business requirements of using private networking for data replication between multisites and ensuring low latency, the appropriate solution within an IaaS (Infrastructure as a Service) provider would be to use a Virtual Private Network (VPN) with Direct Connect or Dedicated Interconnect.

upvoted 1 times

> **kuzummjakk** 9 months, 4 weeks ago
>
> Hey ChatGPT, they said it's over a private network so a VPN is unnecessary.
>
> upvoted 1 times

ROCompTIA 1 year, 6 months ago

Selected Answer: A

By using peering, you can achieve private networking and potentially lower latency compared to routing traffic over the public internet.

upvoted 4 times

> **ROCompTIA** 1 year, 6 months ago
>
> Hub and spoke is the right answer
> By using a hub and spoke architecture, the administrator can accomplish the following:
>
> Private Networking: The hub and spoke model allows for the use of private networking within and between the multisites. This means that the data replication between the sites can occur over private connections, enhancing security and isolation.
>
> Low Latency: The hub and spoke model enables the administrator to optimize the network for low latency. By centralizing the connectivity through a hub, the administrator can establish direct connections between the hub and each spoke. This minimizes the latency compared to routing traffic through multiple hops or relying on public internet connections.
>
> upvoted 4 times

BeauChateau 1 year, 8 months ago

Selected Answer: B

B. Gateways.

To meet the business requirements of using private networking in and between the multisites for data replication and using low latency to avoid performance issues, the cloud administrator should use gateways within the IaaS provider to connect multiregions.

Gateways are used to establish dedicated connections between regions or data centers within an IaaS provider. They allow private network connectivity between regions or data centers, providing a low latency connection for data replication. They also ensure secure and reliable connectivity between the different regions, without the need for additional VPN or peering configurations.

Peering is a method of connecting networks through a direct connection between the two networks. While it can be used to connect multiple regions, it does not provide the dedicated, low-latency connection required for data replication.

upvoted 1 times

AustinKelleyNet 1 year, 11 months ago

Selected Answer: A

I believe the answer is A

upvoted 1 times

Brianhealey136 1 year, 12 months ago

Selected Answer: A

Peering is connecting multiple VPCs (Virtual Private Clouds) together. A VPC is just a VPN except in the cloud. I'll go with A. Could also be C

upvoted 3 times

> **kuzummjakk** 9 months, 4 weeks ago
>
> A VPC is NOT a VPN in the cloud. That is absolutely false.
>
> upvoted 1 times

jiminycriminal 2 years, 3 months ago

Selected Answer: A

VPN is not generally a low latency solution. I'm going with Peering.

upvoted 4 times

⊟ **jiminycriminal** 2 years, 3 months ago

I am wrong. VPC Peering can only take place within the same regions.

https://awsinsider.net/articles/2017/02/13/vpc-peering-with-aws-part-1.aspx#:~:text=Functionally%2C%20VPC%20peering%20is%20similar,no%20VPN%20connection%20is%20required.

Question specifically says multiregion. So the answer is likely VPN.

upvoted 5 times

⊟ **jiminycriminal** 2 years, 3 months ago

Nevermind, AWS introduced inter-region VPC peering in 2017 lol.

https://aws.amazon.com/about-aws/whats-new/2017/11/announcing-support-for-inter-region-vpc-peering/

So I think I'm going back to A. Peering.

upvoted 4 times

⊟ **Rob69420** 2 years, 4 months ago

Azure compute services, namely virtual machines (IaaS) and cloud services (PaaS), that are deployed within a virtual network can be connected through the private peering domain. The private peering domain is considered to be a trusted extension of your core network into Microsoft Azure. You can set up bi-directional connectivity between your core network and Azure virtual networks (VNets). This peering lets you connect to virtual machines and cloud services directly on their private IP addresses.

You can connect more than one virtual network to the private peering domain. Review the FAQ page for information on limits and limitations. You can visit the Azure Subscription and Service Limits, Quotas, and Constraints page for up-to-date information on limits. Refer to the Routing page for detailed information on routing configuration.

upvoted 1 times

⊟ **u2637ps** 2 years, 10 months ago

https://docs.aws.amazon.com/devicefarm/latest/developerguide/amazon-vpc-cross-region.html

If the addresses are different you can Peer

upvoted 1 times

Question #70 *Topic 1*

An organization is hosting a cloud-based web server infrastructure that provides web-hosting solutions. Sudden continuous bursts of traffic have caused the web servers to saturate CPU and network utilizations. Which of the following should be implemented to prevent such disruptive traffic from reaching the web servers?

A. Solutions to perform NAC and DLP

B. DDoS protection

C. QoS on the network

D. A solution to achieve microsegmentation

Suggested Answer: *B*

reto1 3 months, 2 weeks ago

B

DDoS protection is specifically designed to detect and mitigate Distributed Denial of Service (DDoS) attacks, which can overwhelm web servers with excessive traffic. Implementing DDoS protection helps to filter out malicious traffic before it reaches the web servers, ensuring they remain operational during traffic spikes.

upvoted 1 times

54a6b25 5 months, 3 weeks ago

B. DDoS protection: Distributed Denial of Service (DDoS) protection is specifically designed to detect and mitigate large volumes of traffic that can overwhelm web servers. DDoS protection solutions can filter out malicious traffic, allowing legitimate traffic to reach the servers and ensuring continued availability and performance.

upvoted 2 times

kuzummjakk 9 months, 4 weeks ago

B is the most related, even though a "continuous burst of traffic" doesn't automatically mean DDoS. QoS is related to prioritizing specific internal traffic (phone calls), NAC worries about WHO's allowed in, DLP worries about what's allowed out, and microsegmentation is more of a zero-trust idea.

upvoted 2 times

alittlesmarternow 11 months, 1 week ago

DDoS Protection: This specifically targets distributed denial-of-service attacks, which are characterized by sudden and continuous bursts of traffic aimed at overwhelming web servers. It filters and mitigates malicious traffic before it reaches the servers,

upvoted 1 times

Question #71 *Topic 1*

A systems administrator needs to configure a set of policies to protect the data to comply with mandatory regulations. Which of the following should the administrator implement to ensure DLP efficiency prevents the exposure of sensitive data in a cloud environment?

A. Integrity

B. Versioning

C. Classification

D. Segmentation

Suggested Answer: *C*

Community vote distribution

C (100%)

54a6b25 5 months, 3 weeks ago

C. Classification: Data classification involves identifying and categorizing data based on its sensitivity and importance. This is fundamental for DLP, as it allows the system to apply appropriate policies and controls to protect sensitive data, ensuring it is handled according to regulatory requirements.

upvoted 3 times

Pongsathorn 1 year, 3 months ago

Selected Answer: C

To ensure Data Loss Prevention (DLP) efficiency and prevent the exposure of sensitive data in a cloud environment, the administrator should implement **Classification**.

Classification involves categorizing data based on its sensitivity and importance. This categorization allows for the enforcement of policies that dictate how different types of data should be handled, shared, or stored. With proper classification, sensitive data can be identified, and DLP policies can be configured to ensure its protection, including measures like encryption, access controls, and monitoring.

So, the correct answer is **C. Classification**.

upvoted 2 times

ROCompTIA 1 year, 6 months ago

Selected Answer: C

Classification is the process of categorizing data based on its sensitivity and importance. By implementing data classification policies, the administrator can assign appropriate labels or tags to different types of data, indicating their level of sensitivity.

upvoted 3 times

BeauChateau 1 year, 8 months ago

Selected Answer: C

C. Classification.

To ensure DLP (Data Loss Prevention) efficiency in preventing the exposure of sensitive data in a cloud environment and comply with mandatory regulations, a systems administrator should implement a classification policy.

Classification policy involves identifying sensitive data and classifying it based on its sensitivity level. Once the data is classified, appropriate controls can be put in place to ensure that it is protected according to its classification level. This can include encryption, access controls, and monitoring.

upvoted 2 times

packbd 2 years, 3 months ago

Key word in the question also mentions "mandatory regulations". This implies Classification should be used.

upvoted 1 times

Granddude 2 years, 6 months ago

I am going with C based off of https://satoricyber.com/data-classification/data-classification/

upvoted 2 times

⊟ **tsmgunny** 2 years, 11 months ago

I was leaning towards C also but https://digitalguardian.com/blog/what-data-integrity-data-protection-101makes a strong case for A

upvoted 1 times

⊟ **jiminycriminal** 2 years, 3 months ago

No it doesn't.

"That said, data integrity is a desired result of data security, but the term data integrity refers only to the validity and accuracy of data rather than the act of protecting data."

upvoted 1 times

⊟ **martin451** 3 years, 2 months ago

I belive this should be c

upvoted 2 times

Question #72 *Topic 1*

Users of an enterprise application, which is configured to use SSO, are experiencing slow connection times. Which of the following should be done to troubleshoot the issue?

A. â€¢ Perform a memory dump of the OS. â€¢ Analyze the memory dump. â€¢ Upgrade the host CPU to a higher clock speed CPU.

B. â€¢ Perform a packet capture during authentication. â€¢ Validate the load-balancing configuration. â€¢ Analyze the network throughput of the load balancer.

C. â€¢ Analyze the storage system IOPS. â€¢ Increase the storage system capacity. â€¢ Replace the storage system disks to SSD.

D. â€¢ Evaluate the OS ACLs. â€¢ Upgrade the router firmware. â€¢ Increase the memory on the router.

Suggested Answer: *B*

Community vote distribution

B (100%)

Rob69420 Highly Voted 2 years, 4 months ago

Selected Answer: B

B. SSO == Auth

upvoted 5 times

reto1 3 months, 2 weeks ago

Packet Capture: This allows for detailed analysis of the authentication process, helping to identify if there are delays occurring during the SSO handshake or authentication requests.
Load-Balancing Configuration: Verifying the load-balancing setup ensures that requests are being distributed evenly among servers, which can help prevent bottlenecks that may lead to slow response times.
Network Throughput: Analyzing the network throughput of the load balancer can identify any potential network issues that may be contributing to the slow connections.

upvoted 1 times

uzey Most Recent 4 months, 2 weeks ago

Selected Answer: B

This option focuses on the network and authentication aspects of the SSO process, which are most likely to be causing the slow connection times.Packet capture will help identify network bottlenecks or errors during the authentication process.

upvoted 1 times

54a6b25 5 months, 3 weeks ago

B.Perform a packet capture during authentication.
Validate the load-balancing configuration.
Analyze the network throughput of the load balance

upvoted 1 times

yyCherubim 1 year, 1 month ago

Selected Answer: B

What in the <bleep> does storage have to do with slow SSO and Authentication?

upvoted 2 times

Pongsathorn 1 year, 3 months ago

Selected Answer: B

To troubleshoot slow connection times for users experiencing issues with Single Sign-On (SSO) in an enterprise application, the most relevant actions would be related to network and authentication. Option B is the best choice:

B. Perform a packet capture during authentication. Validate the load-balancing configuration. Analyze the network throughput of the load balancer.

Here's why:

1. **Packet Capture**: Capturing network packets during the authentication process can help identify any anomalies or delays in the communication between the user, the authentication server, and the application. It can provide insights into where the slowdown is occurring.

upvoted 1 times

Pongsathorn 1 year, 3 months ago

2. **Load-Balancing Configuration**: SSO systems often rely on load balancers to distribute traffic. Incorrect or misconfigured load balancing can lead to delays. Validating the load-balancing configuration ensures that requests are evenly distributed.

3. **Network Throughput Analysis**: Analyzing the network throughput of the load balancer can reveal if it's handling the traffic efficiently. If it's overwhelmed or has performance issues, it could be a bottleneck causing slow authentication.

Options A, C, and D don't directly address the specific issue of slow connection times during SSO authentication. They focus on different aspects of the infrastructure that are less likely to be the root cause of authentication delays.

upvoted 1 times

Sunshine_boy38 1 year, 5 months ago

B is the answer

upvoted 1 times

BeauChateau 1 year, 8 months ago

Selected Answer: B

B. Perform a packet capture during authentication. Validate the load-balancing configuration. Analyze the network throughput of the load balancer.

To troubleshoot the slow connection times for users of an enterprise application configured with SSO, a systems administrator should perform the following steps:

Perform a packet capture during authentication to determine if there are any delays or errors during the authentication process.

Validate the load-balancing configuration to ensure that the load balancer is configured correctly and distributing the load evenly across the servers.

Analyze the network throughput of the load balancer to identify any network-related issues, such as congestion or network latency.

upvoted 1 times

Agr321 2 years, 3 months ago

I'm leaning to B: Perform a packet capture, etc..
The keyword in the question is troubleshoot! not fix, upgrade, etc....
Really don't understand how capturing encrypted frames of SSO authentication actually help.

upvoted 2 times

packbd 2 years, 3 months ago

Moderators, please fix this. Answer is clearly B. Nothing here related to storage.

upvoted 2 times

jiminycriminal 2 years, 3 months ago

Selected Answer: B

lol we are having connection issues. Nothing to do with storage speed. I also agree that B seems to be correct.

upvoted 1 times

achow26 2 years, 3 months ago

B seems to be the correct answer

upvoted 2 times

Granddude 2 years, 6 months ago

It cant be C. It would be very expensive to upgrade storage to SSD, wouldn't it? Plus, it is troubleshooting...B is the quickest and most cost-effective. To buy or upgrade anything for troubleshooting purposes would not go over well without approvals and change management. I am going with B.

upvoted 2 times

emaney 2 years, 7 months ago

can answer B go under review?

upvoted 2 times

dvd21 3 years, 1 month ago

B. SSO == Auth

upvoted 3 times

Question #73 *Topic 1*

A systems administrator recently upgraded the processors in a web application host. Upon the next login, the administrator sees a new alert regarding the license being out of compliance. Which of the following licensing models is the application MOST likely using?

A. Per device

B. Per user

C. Core-based

D. Volume-based

Suggested Answer: *C*

Community vote distribution

C (100%)

alittlesmarternow 5 months ago

Selected Answer: C

Processor upgrade triggered the alert: This suggests the license is tied to the processing power of the host, not the device itself, users, storage, or network usage.

upvoted 2 times

Zak11 1 year, 2 months ago

Selected Answer: C

the recent processor upgrade may have caused the license to become out of compliance.

upvoted 1 times

Question #74 *Topic 1*

A company is currently running a website on site. However, because of a business requirement to reduce current RTO from 12 hours to one hour, and the RPO from one day to eight hours, the company is considering operating in a hybrid environment. The website uses mostly static files and a small relational database.
Which of the following should the cloud architect implement to achieve the objective at the LOWEST cost possible?

A. Implement a load-balanced environment in the cloud that is equivalent to the current on-premises setup and use DNS to shift the load from on premises to cloud.

B. Implement backups to cloud storage and infrastructure as code to provision the environment automatically when the on-premises site is down. Restore the data from the backups.

C. Implement a website replica in the cloud with auto-scaling using the smallest possible footprint. Use DNS to shift the load from on premises to the cloud.

D. Implement a CDN that caches all requests with a higher TTL and deploy the IaaS instances manually in case of disaster. Upload the backup on demand to the cloud to restore on the new instances.

Suggested Answer: *C*

Community vote distribution

B (50%) C (50%)

reto1 3 months, 2 weeks ago

B. Implement backups to cloud storage and infrastructure as code to provision the environment automatically when the on-premises site is down. Restore the data from the backups.

Explanation:
Cost Efficiency: This approach avoids the expense of running a continuously active cloud environment by utilizing cloud storage for backups, which is generally cheaper than maintaining a fully replicated environment.
Automated Provisioning: Using infrastructure as code allows for quick and efficient deployment of resources when needed, reducing the RTO to the required one hour by automating the recovery process.
Backup Restoration: By performing regular backups to cloud storage, the RPO can be met at eight hours by ensuring that data is up to date within that timeframe.
upvoted 2 times

uzey 4 months, 2 weeks ago

Selected Answer: C

This option provides a balance of cost-effectiveness, performance, and disaster recovery capabilities.

While other options might address some aspects of the requirements, they either incur higher costs or do not fully meet the RTO and RPO objectives
upvoted 1 times

kuzummjakk 9 months, 3 weeks ago

Selected Answer: B

C is more expensive than B. With C, you're paying for two operational websites 100% of the time.
upvoted 3 times

FasterN8 10 months, 1 week ago

Selected Answer: B

Option B is the only one that deals with recovery from a failure of the application (RPO/RTO). Auto-scaling and load balancing and CDNs are about performance and availability, not recovery.
upvoted 2 times

alittlesmarternow 11 months, 1 week ago

Selected Answer: C

Option A: Creates a full cloud replica, duplicating on-premises costs.
Option B: Incurs regular backup costs and requires manual infrastructure provisioning, increasing complexity and potential human error.

Option D: While CDNs reduce load on the origin server, manually deploying instances and uploading backups on demand can be labor-intensive and expensive.

upvoted 1 times

⊟ **Sunshine_boy38** 1 year, 5 months ago

C. Implement a website replica in the cloud with auto-scaling using the smallest possible footprint. Use DNS to shift the load from on-premises to the cloud.

upvoted 1 times

⊟ **Alizadeh** 1 year, 8 months ago

Selected Answer: C

C: Implement a website replica in the cloud with auto-scaling using the smallest possible footprint.

upvoted 3 times

⊟ **maelo** 1 year, 8 months ago

Key is static files and small DB, can be easily backed up very frequently. Infrastructure as code is used to efficiently raise new instances for getting back to service, thus reducing the RTO and RPO.

upvoted 2 times

Question #75 *Topic 1*

After analyzing a web server's log, a systems administrator sees that users are connecting to the company's application through HTTP instead of HTTPS. The administrator then configures a redirect from HTTP to HTTPS on the web server, and the application responds with a connection time-out message. Which of the following should the administrator verify NEXT?

A. The TLS certificate

B. The firewall rules

C. The concurrent connection limit

D. The folder permissions

Suggested Answer: *A*

Community vote distribution

A (59%) B (41%)

dvd21 Highly Voted 3 years, 1 month ago

B. Firewall rules. Users are seeing a connection time, not cert based errors.

upvoted 17 times

reto1 3 months, 2 weeks ago

A. The TLS certificate.

Explanation:

TLS Certificate: It's essential to ensure that the TLS (Transport Layer Security) certificate is correctly installed and valid. If there are issues with the certificate, such as it being expired, misconfigured, or not trusted, clients may not be able to establish a secure connection, leading to time-out errors.

upvoted 1 times

reto1 3 months, 2 weeks ago

If the firewall were blocking HTTPS traffic, users wouldn't be able to connect at all, but the specific issue here is the redirect causing a time-out.

upvoted 1 times

ROCompTIA 1 year, 6 months ago

When configuring a redirect from HTTP to HTTPS, it's essential to have a valid and properly installed Transport Layer Security (TLS) certificate on the web server. The TLS certificate is required to establish a secure connection over HTTPS.This usually happens also on a mail server on that webhost when the certificate has problems !!!

Firewall rules are less likely to directly cause a connection time-out after a redirect from HTTP to HTTPS.

The error here will be Connection Refused, or Destination Unreachable. Therefore, verifying the TLS certificate is the most relevant next step for the administrator.

upvoted 3 times

kuzummjakk 9 months, 3 weeks ago

"it's important to have a TLS cert for HTTPS" yes obviously.

If the connection is refused by a firewall, they typically simply don't respond on behalf of the device, leading a timeout. This is especially true in the cloud. One of the reasons is to deny port scanners the clear answer of "this port is closed".

upvoted 1 times

WereAllinThisTogether Most Recent 1 month, 2 weeks ago

Selected Answer: B

If there was something wrong with the Certificate itself it would propagate a different message that would more specific to the issue it would not generate a time-out. Since accessibility is confirmed utilizing port 80 but time-outs occur when utilizing port 443 when redirected it would be presumed this is a firewall issue blocking the connection. 100% the answer is B. Firewall Rules. Surprisingly i find this to be one of the easiest questions in the question bank but somehow the most debated?

upvoted 1 times

uzey 4 months, 2 weeks ago

Selected Answer: A

Since the redirect to HTTPS is successful but users are encountering a connection timeout,

the most likely culprit is an issue with the TLS certificate. This could be due to several reasons:

Certificate expiration: The certificate might have expired.
Certificate chain issues: There might be problems with the certificate chain, preventing browsers from trusting it.
Certificate configuration errors: The certificate might not be configured correctly on the web server.

By verifying the TLS certificate, the administrator can determine if it is the root cause of the connection timeouts.

upvoted 2 times

TheFivePips 7 months, 2 weeks ago

Selected Answer: B

Both A or B could be correct. I would expect a TLS issue to have a more specific error, and I would definitely think to check the ports first since it is pretty easy. Id go with B but the question is bad.

upvoted 1 times

kuzummjakk 9 months, 3 weeks ago

Selected Answer: B

Absolutely B. If there was a TLS issue, it'd be a TLS error.
There is no shortage of TLS errors for TLS issues so it wouldn't simply not give a TLS error.

upvoted 2 times

salthehash 11 months, 3 weeks ago

Selected Answer: B

The firewall rules: It's possible that the firewall is blocking HTTPS traffic, preventing the successful redirection. The administrator should check the firewall rules to ensure that traffic on the HTTPS port (typically port 443) is allowed

upvoted 4 times

FrancisDrake 12 months ago

Yes, tricky. But the question has NEXT in all caps. Firewall is the easiest thing to check NEXT.

upvoted 1 times

VVV4WIN 1 year, 1 month ago

Trick question, but application server responded, meaning there was connectivity to it, thus firewall not blocking connections. First I was convinced it was the firewall, but Not_That_Guy answer below convinced me otherwise....

upvoted 1 times

> **TheFivePips** 7 months, 2 weeks ago
>
> The connection time-out error is generated by the client-side network stack or web browser, not by the server itself. It occurs when the client's request to establish a connection with the server goes unanswered for too long.
> A connection time-out error does not necessarily indicate whether the server is working or not. It simply means that the client was unable to establish a connection with the server within the allotted time. The server may be unavailable due to various reasons, such as network issues, server downtime, firewall restrictions, or overloaded server resources.
>
> upvoted 1 times

yyCherubim 1 year, 1 month ago

I can see the case for both A & B. It's the 50/50/90 rule of test-taking: You have a 50/50 change of picking the correct answer; 90% of the time you will pick the wrong one. :]

upvoted 1 times

backdooranon 1 year, 2 months ago

Selected Answer: B

B is correct. A is possible if error was certificate related (expired or not trusted).

upvoted 1 times

ROCompTIA 1 year, 6 months ago

Selected Answer: A

I explained why below

upvoted 1 times

SecPlus2022 1 year, 7 months ago

Selected Answer: A

You can have timeouts caused by TLS issues and timeouts caused by TCP issues (both separately or together). The question tells you that the company's web application RESPONDED to requests after it was switched to https. This proves 443 is already open.

upvoted 2 times

⊟ **TheFivePips** 7 months, 2 weeks ago

You got a timeout error. That is not the same thing as the server responding to requests. The errors originate from the client side stack or web browser.

upvoted 1 times

⊟ **maelo** 1 year, 8 months ago

Selected Answer: B

B. Connection time-out message is not cert related.

upvoted 1 times

⊟ **davidsvida** 1 year, 10 months ago

An SSL certificate error occurs when the browser cannot verify the SSL certificates returned by the server. When the error happens, the browser blocks the website and warns the user that the website cannot be trusted as shown below.

upvoted 1 times

⊟ **TheGinjaNinja** 1 year, 11 months ago

Selected Answer: A

A. The TLS certificate

When the administrator configures a redirect from HTTP to HTTPS, the browser will try to establish a secure connection with the server. The browser will check the authenticity of the server's certificate before establishing the connection. If the certificate is invalid, expired or not trusted, the browser will not establish the connection, and the user will get the connection time-out message. Therefore, the administrator should verify the validity, expiration and trust of the server's certificate.

upvoted 2 times

⊟ **CapJackSparrow** 1 year, 11 months ago

also notice it says "users" not "some users" meaning the site has always been using port 80... so port 80 is open. My guess port 443 has not been open in the firewall the whole time, until the admin noticed the log. Im going firewall..

upvoted 1 times

⊟ **CapJackSparrow** 1 year, 11 months ago

Causes of Connection Timed Out Error

There can be various causes that can turn out to be the potential reasons for the connection timed out error. Some of these causes are mentioned below:

Slow Internet Connection: Slow Internet connection is one of the possible reasons for the error. If the system does not have a fast and reliable Internet connection, then there is a possibility that the server accepted the request, but due to slow Internet, the system took too long to respond.

Invalid URL: If the user requests an invalid URL to access the data, then the server automatically denies the request.

Server Delay or Error: It is not mandatory that there might be a problem at the user end only, it is also possible that there might be server delay and before the server could release, data timeout occurred.

Incorrect Settings: There is a possibility that the Windows Firewall did not allow data packets to get transferred into the system because of which the browser displays this error.

upvoted 1 times

Question #76 *Topic 1*

A company needs to access the cloud administration console using its corporate identity. Which of the following actions would MOST likely meet the requirements?

A. Implement SSH key-based authentication.

B. Implement cloud authentication with local LDAP.

C. Implement multifactor authentication.

D. Implement client-based certificate authentication.

Suggested Answer: *B*

Community vote distribution

B (87%) 13%

ryanzou Highly Voted 2 years, 3 months ago

Selected Answer: B

B is correct

upvoted 8 times

reto1 3 months, 2 weeks ago

B. Implement cloud authentication with local LDAP.

Explanation:
Cloud Authentication with Local LDAP: This approach allows the company to integrate its existing corporate identity management system (using LDAP) with the cloud provider's authentication mechanism. This enables users to log in to the cloud administration console using their corporate credentials, which meets the requirement for accessing the console with the corporate identity.

upvoted 2 times

uzey Most Recent 4 months, 2 weeks ago

Selected Answer: B

This approach allows users to use their existing corporate credentials for accessing the cloud console,
simplifying the authentication process and enhancing security.

upvoted 1 times

TheFivePips 7 months, 2 weeks ago

Selected Answer: B

SSH key-based authentication is typically used for accessing remote servers via SSH (Secure Shell) protocol. While SSH key-based authentication provides secure access to servers, it is not directly applicable to cloud administration consoles, which typically use web-based interfaces rather than SSH connections.

Integrating the cloud administration console with a local LDAP (Lightweight Directory Access Protocol) server allows users to authenticate using their corporate identity credentials stored in the LDAP directory. LDAP integration provides centralized user management and authentication, enabling users to log in to the cloud administration console using their existing corporate usernames and passwords.

While MFA enhances security, it does not directly address the requirement to use the company's corporate identity for accessing the cloud administration console.

Client-based certificate authentication involves issuing digital certificates to users, which they present to authenticate themselves. While client-based certificate authentication provides a secure method of authentication, it doesnt meet the requirement to use a corporate identity.

upvoted 2 times

Sunshine_boy38 1 year, 5 months ago

Correct Answer: D

upvoted 1 times

SecPlus2022 1 year, 6 months ago

Selected Answer: B

Third and final comment. Changing my answer to "B". I'm reading too much into the question. It's not about proving a corporate identity, it's simply about accessing the cloud console. The company most likely has domain environment so use LDAP for corporate directory resources.

upvoted 2 times

SecPlus2022 1 year, 7 months ago

I'm in need of adding an addendum to my comment. Yes, in addition to LDAP reading and modifying AD information, it can provide authentication via the bind operation. However, I'm still sticking with my answer ("D"), because a corporate identify isn't going to established via a username / password authentication process, it's going to be proven via a certificate issued by a CA (who has verified identity during the certificate issuing process).

upvoted 2 times

SecPlus2022 1 year, 7 months ago

Selected Answer: D

LDAP is a protocol that allows reading and modifying user information stored within Active Directory, it does not prove identity. The answer is "D".

upvoted 2 times

maelo 1 year, 8 months ago

I understand "corporate identity" as not user specific, but company-global. This can be provided by a respective certificate, to be imported/trusted at clients, or by trusting a root cert.

upvoted 2 times

ironman_86 2 years, 3 months ago

i think it should be B

upvoted 2 times

achow26 2 years, 3 months ago

Answer should be B.

upvoted 2 times

Question #77 *Topic 1*

A cloud administrator is planning to migrate a globally accessed application to the cloud. Which of the following should the cloud administrator implement to BEST reduce latency for all users?

A. Regions

B. Auto-scaling

C. Clustering

D. Cloud bursting

Suggested Answer: *A*

Community vote distribution

A (100%)

⊟ **alittlesmarternow** 5 months ago

Selected Answer: A

Regions: By deploying the application in multiple regions closer to users worldwide, the distance data needs to travel is reduced, leading to lower latency. This allows users to access the application faster regardless of their location.

upvoted 3 times

Question #78 *Topic 1*

A company needs a solution to find content in images. Which of the following technologies, when used in conjunction with cloud services, would facilitate the
BEST solution?

A. Internet of Things

B. Digital transformation

C. Artificial intelligence

D. DNS over TLS

Suggested Answer: *C*

Community vote distribution

C (100%)

alittlesmarternow 11 months, 1 week ago

Selected Answer: C

AI for image content search:
Techniques like computer vision and image recognition can analyze and understand the content of images, enabling users to search for specific objects, scenes, or attributes within them.
These techniques require massive datasets and significant computational power for training and inference, making cloud services ideal due to their scalability, flexibility, and access to high-performance computing resources.

upvoted 3 times

reto1 3 months, 2 weeks ago

C. Artificial intelligence.

Explanation:
Artificial Intelligence (AI): AI technologies, particularly those involving machine learning and image recognition, can analyze images to identify and categorize content. This capability is essential for tasks such as image classification, object detection, and facial recognition, making AI the most suitable choice for extracting information from images.

upvoted 1 times

Question #79 *Topic 1*

An organization is developing a new solution for hosting an external website. The systems administrator needs the ability to manage the OS. Which of the following methods would be MOST suitable to achieve this objective?

A. Deploy web servers into an IaaS provider.

B. Implement a cloud-based VDI solution.

C. Provision web servers in a container environment.

D. Use PaaS components in the cloud to implement the product.

Suggested Answer: *A*

Community vote distribution

A (100%)

ryanzou Highly Voted 2 years, 3 months ago

Selected Answer: A

Only IaaS can let you manage the OS

upvoted 18 times

reto1 3 months, 2 weeks ago

A. Deploy web servers into an IaaS provider.

Explanation:

IaaS (Infrastructure as a Service): This model provides the systems administrator with full control over the virtual machines, including the ability to manage the operating system, install software, and configure server settings as needed. This level of control is essential for managing the OS effectively.

upvoted 1 times

StudyBM Highly Voted 2 years, 3 months ago

Selected Answer: A

You don't manage an OS with Containers

upvoted 7 times

jiminycriminal 2 years, 3 months ago

This. While you can have containers in a compute instance, CSPs allow containers environments outside of that where you don't manage any OS.

upvoted 1 times

kuzummjakk Most Recent 9 months, 3 weeks ago

Selected Answer: A

"technically" there's a tiny OS in the container, but the question is very clearly talking about the VM's OS. Don't read too far into it.

upvoted 1 times

Robenger 11 months, 2 weeks ago

C

Yes, you can manage an operating system (OS) with a container. A container is a lightweight and portable software package that includes everything needed to run an application, including the code, runtime, libraries, and system tools. Containers can be used to package and deploy applications, as well as to manage the underlying OS.

For example, Google's Container-Optimized OS is an operating system image that is optimized for running Docker containers. With Container-Optimized OS, you can bring up your Docker containers on Google Cloud Platform quickly, efficiently, and securely.

upvoted 1 times

Sunshine_boy38 1 year, 5 months ago

A. Deploy web servers into an IaaS provider.

upvoted 1 times

BeauChateau 1 year, 8 months ago

Selected Answer: A

A. Deploy web servers into an IaaS provider.

To achieve the objective of having the ability to manage the OS for hosting an external website, the most suitable method would be to deploy web servers into an IaaS (Infrastructure-as-a-Service) provider.

upvoted 2 times

ironman_86 2 years, 3 months ago

after rethink the question, i see C as the correct answer.

upvoted 1 times

ironman_86 2 years, 3 months ago

why not A?

upvoted 1 times

Question #80 *Topic 1*

A systems administrator is provisioning VMs in a cloud environment and has been told to select an OS build with the furthest end-of-life date. Which of the following OS builds would be BEST for the systems administrator to use?

A. Open-source

B. LTS

C. Canary

D. Beta

E. Stable

Suggested Answer: *B*

Community vote distribution

B (100%)

reto1 3 months, 2 weeks ago

B. LTS (Long Term Support)

LTS versions are designed to receive updates and support for an extended period, typically several years, making them a reliable choice for long-term use in production environments. In contrast, options like Canary and Beta are often experimental and not suitable for production, while Stable builds may not offer the same extended support duration as LTS versions.

upvoted 1 times

sheilawu 1 year, 5 months ago

was on my exam 7/2

upvoted 2 times

maelo 1 year, 8 months ago

LTS doesn't take actual date into account, "LTS" label is not revoked. Means there are old/deprecated LTS versions in SW space, or an aging LTS version can be replaced by next LTS version of same product in short term. No good answer available. Simply bad question wording, might better say "... to look out for FIRST - and consider dates".

upvoted 3 times

BeauChateau 1 year, 8 months ago

Selected Answer: B

B. LTS (Long-Term Support)

When provisioning VMs in a cloud environment and selecting an OS build with the furthest end-of-life date, the best option for the systems administrator to use would be an OS build with Long-Term Support (LTS).

upvoted 2 times

SimplyDebonair 2 years, 9 months ago

Selected Answer: B

The correct answer is "B." A stable build typically only has limited support, usually only for a year at most. A Long-Term Support (LTS) however, will have gone through the motions of being in a operating capacity. It will also have evidence of a strong, "stable" backbone to work off of which is why it typically has support for three to five years.

-CompTIA Cloud+ Certification All-in-One Exam Guide (Exam CV0-003)

upvoted 4 times

> **SimplyDebonair** 2 years, 8 months ago
>
> For further clarification. I under this isn't related wholly to Cloud, but the premise remains the same between LTS and Stable.
>
> https://stackoverflow.com/questions/34829167/what-is-the-difference-between-the-lts-version-and-the-stable-version-of-node-js
>
> upvoted 1 times

dvd21 3 years, 1 month ago

LTS - Long Term Support

upvoted 2 times

Question #81 *Topic 1*

A cloud administrator set up a link between the private and public cloud through a VPN tunnel. As part of the migration, a large set of files will be copied. Which of the following network ports are required from a security perspective?

A. 22, 53, 445

B. 22, 443, 445

C. 25, 123, 443

D. 137, 139, 445

Suggested Answer: *B*

Community vote distribution

B (86%) 14%

Khairulhak 4 months, 1 week ago

Selected Answer: B

22 is for SSH
443 is for HTTPS
445 is for SMB

upvoted 3 times

Pongsathorn 1 year, 3 months ago

Selected Answer: B

From a security perspective, you want to ensure that your VPN tunnel is set up to allow secure and encrypted communication between your private and public clouds. Common VPN protocols include SSL/TLS (used on port 443) and IPsec (which can use various ports, but UDP 500 for IKE negotiation is common). SMB (Server Message Block) ports (137-139 and 445) are typically associated with file sharing and may not be needed for a VPN tunnel.

So, the **best option** for secure VPN communication in this scenario would be:

B. 22, 443, 445

- Port 22 is for SSH, which can be used for secure remote administration.
- Port 443 is commonly used for SSL/TLS encrypted communication, which is secure.
- Port 445 is associated with SMB, which is less likely to be needed for a VPN tunnel and is often used for file sharing.

upvoted 4 times

Sunshine_boy38 1 year, 5 months ago

B. 22, 443, 445

upvoted 1 times

SecPlus2022 1 year, 7 months ago

Selected Answer: B

Given the VPN is taking care of the secure part of the equation, I guess you could technically get away with "D", but NetBIOS is deprecated. SFTP is better suited for moving large files compared to SMB. There's no mention of overhead concern here so securely transferring your files (SFTP) over a secure connection (VPN) isn't something I would see as being problematic (can never be too secure). SMB may need to rely on NetBios to communicate, but ONLY if the device in question is old and is unable to support direct hosting of SMB over TCP/IP.

upvoted 1 times

craigbharrell 1 year, 9 months ago

Selected Answer: B

B, because SFTP/SSH, HTTPS, and SMB are more secure than netbios and can be used for file transfers. Also, netbios isn't on exam objectives and SFTP is (not like that always matters)

upvoted 3 times

AustinKelleyNet 1 year, 11 months ago

Selected Answer: B

I think B

upvoted 1 times

⊟ **maiathiago** 1 year, 11 months ago

Selected Answer: D

137 - NetBIOS - Allow windows machines to "talk" on the network.
139,445 - SMB - Protocol for sharing data.

Usually, you need the conjunction of them to have SMB working.

https://www.varonis.com/pt-br/blog/smb-port

upvoted 2 times

Question #82 *Topic 1*

A company is preparing a hypervisor environment to implement a database cluster. One of the requirements is to share the disks between the nodes of the cluster to access the same LUN. Which of the following protocols should the company use? (Choose two.)

A. CIFS

B. FTP

C. iSCSI

D. RAID 10

E. NFS

F. FC

Suggested Answer: *CF*

Community vote distribution

CF (100%)

i_bird Highly Voted 2 years, 3 months ago

Is RAID10 a protocol?

Leaning towards
FC and iSCSI

upvoted 7 times

reto1 3 months, 2 weeks ago

Both iSCSI and Fibre Channel are used for block storage access and are capable of sharing disks between multiple servers, allowing them to access the same LUN effectively.

upvoted 1 times

jiminycriminal 2 years, 3 months ago

https://docs.netapp.com/us-en/ontap/san-admin/lun-san-environments-concept.html

"The FC protocol and iSCSI protocol both provision storage through the use of LUNs"

I think you're correct.

upvoted 8 times

uzey Most Recent 4 months, 2 weeks ago

Selected Answer: CF

Both iSCSI and FC are designed for block-level data transfer, which is essential for sharing disks between nodes in a database cluster. They provide the necessary performance and reliability for such environments.

upvoted 1 times

c703728 5 months, 2 weeks ago

C. iSCSI - Internet Small Computer Systems Interface (iSCSI) is a protocol that allows clients (called initiators) to send SCSI commands to SCSI storage devices (targets) on remote servers. It is used to facilitate data transfers over intranets and to manage storage over long distances.

F. FC - Fibre Channel is a high-speed network technology primarily used to connect computer data storage. Fibre Channel is especially suited for connecting servers to shared storage devices and is used in storage area networks (SANs).

upvoted 1 times

kuzummjakk 9 months, 3 weeks ago

RAID10 has nothing to do with sharing disks. Sure it's nice and import, but not relevant to the question. ChatGPT type answer

upvoted 1 times

Sunshine_boy38 1 year, 5 months ago

The answer is C and E. iSCSI and NFS.

iSCSI is a network protocol that allows block-level storage to be shared over a network. This makes it a good choice for sharing disks between the nodes of a database cluster.

NFS is a file-level sharing protocol that allows files to be shared over a network. This makes it a good choice for sharing files between the nodes of a database cluster.

The other protocols listed are not suitable for sharing disks between the nodes of a database cluster.

CIFS and FTP are file-level sharing protocols that are not designed for sharing block-level storage.
RAID 10 is a type of RAID that provides fault tolerance, but it does not allow disks to be shared between nodes.
FC is a Fibre Channel protocol that is designed for high-performance storage, but it is not as widely supported as iSCSI or NFS.
upvoted 1 times

Sunshine_boy38 1 year, 5 months ago
Correction: i reverse back my answer to C. iSCSI & F. FC (Fibre Channel)
FC (Fibre Channel): FC is a high-speed storage area network (SAN) protocol that provides block-level access to storage devices. It allows multiple servers to share the same storage LUN over a dedicated and high-performance Fibre Channel fabric. FC is often used in enterprise-level environments where high performance and low latency are critical, such as database clusters.

NFS (Network File System): NFS is a file-level protocol used for sharing files over a network. Like CIFS, it does not offer the block-level access required for shared disks in a database cluster.
upvoted 3 times

NRJ6425 1 year, 7 months ago
Selected Answer: CF
C and F
upvoted 2 times

BeauChateau 1 year, 8 months ago
Selected Answer: CF
C. iSCSI
F. FC

To share disks between the nodes of a database cluster to access the same LUN, the company should use either iSCSI or FC protocols. Both of these protocols provide block-level access to shared storage, allowing multiple nodes to access the same LUN simultaneously.
upvoted 3 times

bagsik89 1 year, 10 months ago
C&F ISCSI and Fibre Channel
upvoted 1 times

AustinKelleyNet 1 year, 11 months ago
Selected Answer: CF
Raid 10 = Redundant Array of Independent Disks. NOT a protocol.
upvoted 4 times

beamage 1 year, 11 months ago
Selected Answer: CF
No Raid 10 is not a protocol!!!
upvoted 4 times

Rob69420 2 years, 3 months ago
A LUN is used by a transport protocol associated with an SCSI, iSCSI, Fibre Channel (FC) or similar interface. LUNs are central to the management of block storage arrays shared over a SAN.
upvoted 4 times

Question #83 *Topic 1*

A cloud administrator is working in a secure government environment. The administrator needs to implement corrective action due to recently identified security issue on the OS of a VM that is running a facility-management application in a cloud environment. The administrator needs to consult the application vendor, so it might take some time to resolve the issue. Which of the following is the FIRST action the administrator should take while working on the resolution?

A. Shut down the server.

B. Upgrade the OS

C. Update the risk register.

D. Raise a problem ticket.

Suggested Answer: *D*

Community vote distribution

D (61%) C (39%)

⊟ **Securityguy42** Highly Voted 8 months, 4 weeks ago

Selected Answer: D

For me, it just makes sense to create a problem ticket first, then update the ticket stating you udated the risk register.

D. Raise a problem ticket.

Raising a problem ticket should be the first action taken by the administrator. This allows the issue to be formally documented and tracked, ensuring that it is properly addressed within the organization's processes and procedures. Additionally, it provides a means for communication with relevant stakeholders, including the application vendor, and helps prioritize the resolution of the security issue. Shutting down the server or upgrading the OS might be necessary steps in resolving the issue but should come after the problem ticket has been raised and the appropriate actions planned. Updating the risk register is important but not the first action to take in this scenario, as the immediate focus should be on addressing the security issue and initiating the resolution process.

upvoted 5 times

⊟ **uzey** Most Recent 4 months, 2 weeks ago

Selected Answer: C

Updating the risk register

is the first and most appropriate action in this scenario.

Here's why:

Immediate action: While consulting the vendor is necessary, it's a longer-term solution. Updating the risk register immediately documents the identified security issue, its potential impact, and the ongoing mitigation efforts.

Risk management: This step aligns with the principle of proactive risk management, which is essential in a secure government environment.

Communication: Updating the risk register ensures that all relevant stakeholders are aware of the issue and its potential consequences.

Once the risk register is updated, the administrator can proceed with other actions, such as isolating the VM or applying temporary mitigations while waiting for a permanent solution.

upvoted 1 times

⊟ **reto1** 3 months, 2 weeks ago

Updating the risk register helps document the identified security issue, assesses its impact, and tracks any actions taken in response. This is crucial in a secure government environment, as it maintains a clear record of potential risks and the steps being taken to address them while the administrator works on consulting the application vendor.

upvoted 1 times

⊟ **kuzummjakk** 9 months, 3 weeks ago

Selected Answer: C

"first action WHILE working on a resolution" so something secondary to communicating with the vendor, essentially insinuating that they already have by this point.

Since none of the answers says hot fix, the next possible logical step would be to figure out all the affected devices by updating the risk register.

upvoted 2 times

Chiaretta 11 months, 3 weeks ago

Selected Answer: C

The first action is update the risk register!!!

upvoted 2 times

Zak11 1 year, 8 months ago

Selected Answer: C

In a secure government environment, it is important to follow proper procedures and protocols when dealing with security issues. The risk register should be updated first to document the security issue and ensure that all necessary steps are taken to mitigate the risk. Shutting down the server or upgrading the OS should only be done after proper risk assessment and consultation with the application vendor. Raising a problem ticket may be necessary, but it should not be the first action taken in this situation.

upvoted 2 times

BeauChateau 1 year, 8 months ago

Selected Answer: D

D. Raise a problem ticket.

The FIRST action that the cloud administrator should take while working on resolving the recently identified security issue on the OS of a VM that is running a facility-management application in a cloud environment is to raise a problem ticket.

A problem ticket will allow the administrator to track the progress of the issue and ensure that it is being addressed in a timely manner. It will also allow for proper documentation of the issue and any actions taken to resolve it.

upvoted 3 times

Alizadeh 1 year, 8 months ago

Selected Answer: D

D. Raise a problem ticket.

The first action the administrator should take while working on the resolution is to raise a problem ticket. This will ensure that the issue is properly documented, tracked, and communicated with the appropriate stakeholders. It also helps in prioritizing and managing the issue until it is resolved. In the meantime, the administrator can consult the application vendor for guidance on resolving the security issue.

upvoted 3 times

[Removed] 1 year, 10 months ago

I think it's correct with D. It emphasizes FIRST, and the first thing that is normally done in an IT department is create a ticket.

When the Log4j vulnerability happened, our department created a ticket for each machine it was found on and assigned it to the VM owner.

upvoted 1 times

namangel 2 years, 2 months ago

C. The risk register is updated with information on new risks as an output of this process.

upvoted 2 times

jiminycriminal 2 years, 3 months ago

If it's an OS security issue, why are they contacting the application vendor? I'm not fully understanding the question.

upvoted 1 times

FrancisDrake 12 months ago

Right. CompTIA trying to confuse us.

upvoted 2 times

Question #84 *Topic 1*

A DevOps administrator is designing a new machine-learning platform. The application needs to be portable between public and private clouds and should be kept as small as possible. Which of the following approaches would BEST meet these requirements?

A. Virtual machines

B. Software as a service

C. Serverless computing

D. Containers

Suggested Answer: *D*

Community vote distribution

D (100%)

Pisces225 3 months ago

Selected Answer: D

Gonna agree with the default on this one

upvoted 4 times

Question #85 *Topic 1*

A cloud administrator wants to have a central repository for all the logs in the company's private cloud. Which of the following should be implemented to BEST meet this requirement?

A. SNMP

B. Log scrubbing

C. CMDB

D. A syslog server

Suggested Answer: *D*

Community vote distribution

D (100%)

BeauChateau 2 months, 2 weeks ago

Selected Answer: D

D. A syslog server

To have a central repository for all the logs in the company's private cloud, a syslog server should be implemented.

A syslog server is a centralized logging solution that collects logs from different sources and stores them in a central location. It provides a convenient way for administrators to search, analyze, and manage logs from different systems and applications in the cloud environment.

upvoted 2 times

Alizadeh 2 months, 3 weeks ago

Selected Answer: D

D. A syslog server

upvoted 2 times

Question #86 *Topic 1*

A company needs to rehost its ERP system to complete a datacenter migration to the public cloud. The company has already migrated other systems and configured VPN connections. Which of the following MOST likely needs to be analyzed before rehosting the ERP?

A. Software

B. Licensing

C. Right-sizing

D. The network

Suggested Answer: *B*

Community vote distribution

B (56%) D (44%)

uzey 4 months, 2 weeks ago

Selected Answer: B

Key considerations include:

Licensing models: Understanding the ERP's licensing model (per user, per processor, etc.) is crucial to determine the costs and requirements in the cloud environment.
License compatibility: Ensuring that the ERP licenses are compatible with the cloud platform is essential to avoid licensing issues and additional costs.
Licensing optimization: Identifying opportunities to optimize licensing costs in the cloud environment can lead to significant savings.

By thoroughly analyzing the licensing aspect, the company can ensure a smooth migration and avoid unexpected costs or compliance issues.

upvoted 2 times

reto1 3 months, 2 weeks ago

Licensing is crucial in this scenario because it determines whether the existing licenses for the ERP software can be transferred to the cloud environment, whether additional licenses are needed, and any compliance issues that may arise during the migration. Ensuring that licensing is in order helps avoid legal and financial complications after the migration.

upvoted 1 times

TheFivePips 7 months, 2 weeks ago

Selected Answer: B

Given that VPN connections are already configured, the primary network considerations (connectivity, latency, bandwidth) are likely addressed. Ensuring that network requirements are met is important, but licensing typically presents more immediate and potentially blocking issues.

Right-sizing is important to optimize resource utilization and cost in the cloud, but this typically follows after ensuring that the software can be legally and compliantly rehosted in the new environment. Licensing issues need to be resolved before detailed performance and resource planning can be effectively done.

ERP systems typically come with complex licensing agreements that may include specific terms for deployment on different types of environments, such as on-premises versus cloud. Many software vendors have distinct licensing models for cloud deployments which might differ in cost and terms from on-premises licenses. Ensuring that the licenses are compliant with the new deployment in the public cloud is crucial to avoid potential legal and financial penalties. Analyzing the licensing terms can reveal potential cost differences, which is essential for budget planning and cost management in the cloud environment.

upvoted 1 times

kuzummjakk 9 months, 3 weeks ago

I'm gonna say C.
B is a pretty solid choice, but it might be drawing out too much extra context to assume that they need to license their ERP solution. However, the cloud SUPER hits hard on right-sizing for migrations which lines up pretty well.

upvoted 1 times

Robenger 11 months, 2 weeks ago

D
"The company has already migrated other systems and configured VPN connections."
This means to me that the network is already situated as you couldn't set up a VPN otherwise.
upvoted 1 times

FrancisDrake 12 months ago
Or could be IaaS. Either way it appears that they would be responsible for licensing when moving to a public cloud.

https://docs.flexera.com/flexera/EN/ITAssets/Lic-BYOSL4cloud.htm
upvoted 1 times

FrancisDrake 12 months ago
Selected Answer: B
It sounds like this is from the perspective of the company providing ERP so licensing would be a consideration for them. Sounds like a PaaS model.
upvoted 3 times

utied 1 year ago
Selected Answer: B
ERP: enterprise resource planning. Allowing HR, Sales, Marketing, Accounting... to access the same database for approved information.
So basically ERP is a SaaS solution to enterprise data held in a database format. It allows the management of who can access what of company data as a cloud solution.

This website: https://www.netsuite.com/portal/resource/articles/erp/cloud-erp.shtml
show subscription licensing to use the software. The licensing needs to match the type of hosting(on-prem, cloud-based-by-them, remote cloud hosted). It looks like the anser is B. Licensing.
upvoted 3 times

ROCompTIA 1 year, 6 months ago
Selected Answer: D
Before rehosting the ERP system to complete the datacenter migration to the public cloud, it is essential to analyze the network requirements and considerations. The network analysis is crucial because the ERP system relies on network connectivity to function properly and communicate with other systems and services.
upvoted 4 times

kuzummjakk 9 months, 3 weeks ago
Dude, EVERYTHING relies on network connectivity to function. What makes it the better option? ChatGPT type answer.
upvoted 3 times

Zak11 1 year, 8 months ago
Selected Answer: D
D. The network. Before rehosting the ERP system, the network needs to be analyzed to ensure it can support the ERP's requirements and to identify any potential issues that could affect the migration. This includes reviewing the current network topology, bandwidth requirements, and latency between the on-premises datacenter and the public cloud. Once the network analysis is complete, the other factors such as software, licensing, and right-sizing can be addressed.
upvoted 3 times

Question #87 Topic 1

An administrator is securing a private cloud environment and wants to ensure only approved systems can connect to switches. Which of the following would be
MOST useful to accomplish this task?

A. VLAN

B. NIPS

C. WAF

D. NAC

Suggested Answer: *D*

Community vote distribution

D (100%)

reto1 3 months, 2 weeks ago

NAC solutions enforce security policies on devices trying to connect to the network, allowing only authorized systems to access network resources. This is crucial for maintaining security in a private cloud by ensuring that only compliant and approved devices are allowed access to switches and other network components.

upvoted 1 times

BeauChateau 1 year, 8 months ago

Selected Answer: D

D. NAC

NAC (Network Access Control) is the most useful option to ensure only approved systems can connect to switches.

NAC is a security solution that helps to enforce network security policies by controlling access to network resources. It ensures that only authorized devices can access the network by validating their identity and checking their compliance with security policies.

upvoted 3 times

Question #88 *Topic 1*

An organization is currently deploying a private cloud model. All devices should receive the time from the local environment with the least administrative effort.
Which of the following ports needs to be opened to fulfill this requirement?

A. 53

B. 67

C. 123

D. 161

Suggested Answer: *C*

Community vote distribution

C (100%)

BeauChateau Highly Voted 2 months, 2 weeks ago

Selected Answer: C

C. 123

To ensure all devices receive the time from the local environment with the least administrative effort in a private cloud model, the organization needs to use the Network Time Protocol (NTP). NTP uses port 123, so the organization needs to open port 123 to allow NTP traffic to flow through the network.

Port 53 is used for DNS traffic, port 67 is used for DHCP traffic, and port 161 is used for SNMP traffic, but they are not related to the time synchronization requirement.

upvoted 6 times

Question #89 *Topic 1*

A cloud security analyst needs to ensure the web servers in the public subnet allow only secure communications and must remediate any possible issue. The stateful configuration for the public web servers is as follows:

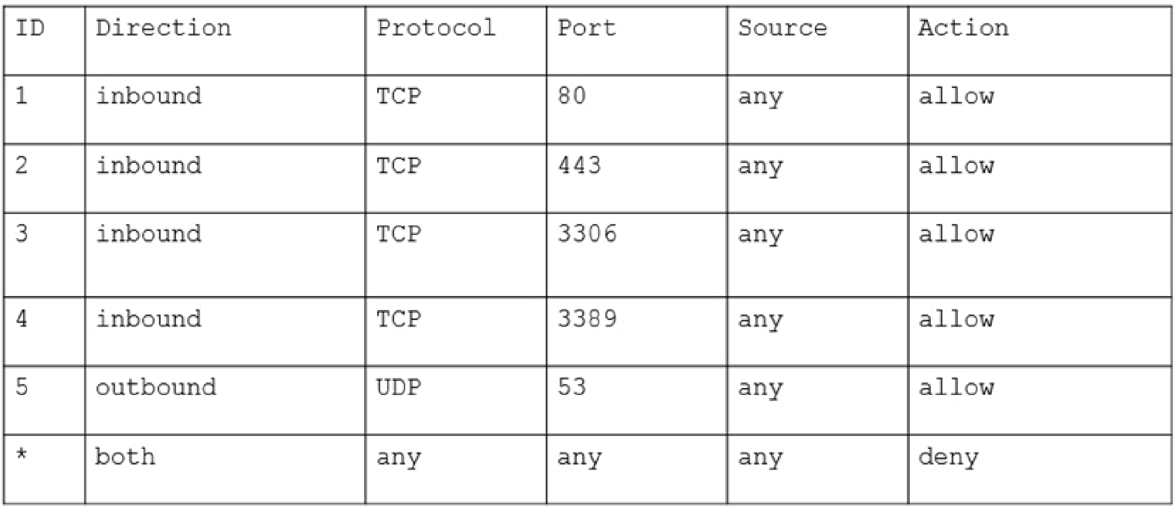

ID	Direction	Protocol	Port	Source	Action
1	inbound	TCP	80	any	allow
2	inbound	TCP	443	any	allow
3	inbound	TCP	3306	any	allow
4	inbound	TCP	3389	any	allow
5	outbound	UDP	53	any	allow
*	both	any	any	any	deny

Which of the following actions should the analyst take to accomplish the objective?

A. Remove rules 1, 2, and 5.

B. Remove rules 1, 3, and 4.

C. Remove rules 2, 3, and 4.

D. Remove rules 3, 4, and 5.

Suggested Answer: *B*

Community vote distribution

B (100%)

TheFivePips 1 month, 1 week ago

Selected Answer: B

Just a reminder for people like me who forget the less common ports

Port 3306: Used by MySQL for database connections.
Port 3389: Used by Microsoft RDP for remote desktop connections.

upvoted 2 times

lilegg 3 weeks, 3 days ago

I don't understand this question :(

upvoted 2 times

AustinKelleyNet 1 year, 5 months ago

Selected Answer: B

Must be B

upvoted 2 times

Question #90 *Topic 1*

Which of the following definitions of serverless computing BEST explains how it is different from using VMs?

A. Serverless computing is a cloud-hosting service that utilizes infrastructure that is fully managed by the CSP.

B. Serverless computing uses predictable billing and offers lower costs than VM compute services.

C. Serverless computing is a scalable, highly available cloud service that uses SDN technologies.

D. Serverless computing allows developers to focus on writing code and organizations to focus on business.

Suggested Answer: *D*

Community vote distribution

D (100%)

cuzzindavid Highly Voted 2 years, 3 months ago

the answer is "D"

"...AWS offers technologies for running code, managing data, and integrating applications, all without managing servers. Serverless technologies feature automatic scaling, built-in high availability, and a pay-for-use billing model to increase agility and optimize costs. These technologies also eliminate infrastructure management tasks like capacity provisioning and patching, so you can focus on writing code that serves your customers...."

upvoted 9 times

reto1 3 months, 2 weeks ago

Serverless computing allows developers to focus on writing code and organizations to focus on business.

This highlights the key advantage of serverless computing: it abstracts infrastructure management, enabling developers to concentrate on application development rather than server management, which is a primary distinction from traditional VM-based approaches where infrastructure management is still a concern.

upvoted 1 times

jiminycriminal Highly Voted 2 years, 3 months ago

Selected Answer: D

C is wrong as it applies to VMs just as well. The answer is D. It is a more concise answer that applies directly to serverless computing.

upvoted 5 times

uzey Most Recent 4 months, 2 weeks ago

Selected Answer: D

Serverless computing
abstracts away the underlying infrastructure, allowing developers to concentrate solely on writing code without worrying about server management, scaling, or provisioning.

upvoted 1 times

TheFivePips 7 months, 2 weeks ago

Selected Answer: D

A. Incomplete Explanation: While it is true that the infrastructure is managed by the CSP, this definition does not fully capture the core benefit of serverless computing, which is the focus on code rather than infrastructure management.

B. Not Always True: While serverless computing can be cost-effective, billing can be less predictable due to its event-driven nature. Costs depend on usage patterns, which can fluctuate.

C. Too General: This definition highlights scalability and availability, which are also characteristics of VMs and other cloud services. It does not emphasize the main distinction of focusing on code over infrastructure.

D. Serverless computing allows developers to focus on writing code and organizations to focus on business best captures the essence of serverless computing, highlighting the key difference from VMs: removing the burden of infrastructure management and allowing more focus on application development and business goals.

upvoted 1 times

Pongsathorn 1 year, 3 months ago

Selected Answer: D

D. Serverless computing allows developers to focus on writing code and organizations to focus on business.

This definition best explains how serverless computing differs from using virtual machines (VMs). In serverless computing, developers don't need to manage or provision servers or VMs. Instead, they can focus solely on writing the application code, and the cloud provider takes care of the underlying infrastructure, including server provisioning, scaling, and maintenance. This allows organizations to concentrate on their core business logic rather than infrastructure management, which is typically required when using VMs. Serverless computing abstracts away the concept of servers or VMs, making it a more abstract and code-centric approach to cloud computing.

upvoted 3 times

Tomtom11 1 year, 5 months ago

Selected Answer: D

Compute may refer to one of two things: IaaS virtual machines, or so-called
serverless computing.
IaaS Compute may refer to an IaaS service that lets you provision virtual machines,
storage, and networking resources in the cloud.
Serverless/FaaS Compute can also refer to what the marketers call serverless computing
and what the technophiles call function-as-a-service (FaaS). In this model, the cloud provider
hands you a slick interface into which you can upload your own application code
written in a variety of programming languages, and the cloud provider executes it on
compute infrastructure that they fully manage. This model obviates the need to provision
virtual machines. Instead, the cloud provider handles the compute infrastructure, so all
you have to do is deal with the application code. FaaS is a type of PaaS offering.

upvoted 1 times

Sunshine_boy38 1 year, 5 months ago

Selected Answer: D

upvoted 1 times

ROCompTIA 1 year, 6 months ago

Selected Answer: D

The best definition that explains how serverless computing is different from using virtual machines (VMs) is serverless computing allows developers to focus on writing code and organizations to focus on their business, rather than managing and provisioning servers.

upvoted 3 times

BeauChateau 1 year, 8 months ago

Selected Answer: D

D. Serverless computing allows developers to focus on writing code and organizations to focus on business.

Serverless computing is a cloud computing model where the cloud provider manages the infrastructure and automatically allocates resources as needed to run and scale applications. Unlike VMs, which require the user to manage the entire operating system and associated software, serverless computing allows developers to focus on writing code without worrying about infrastructure management. With serverless computing, the user only pays for the computing resources used by their application, rather than paying for a fixed amount of resources as with VMs.

upvoted 2 times

Markedexam 1 year, 10 months ago

Voting C.
Google "SDN Technologies" & you get "Software-Defined Networking (SDN) is an approach to networking that uses software-based controllers or APIs to communicate with underlying hardware infrastructure and direct traffic on a network" That's what the questions is focused on - "how definitions of serverless computing BEST explains how it is different from using VMs" C offers "Serverless computing is a scalable, highly available cloud service that uses SDN technologies." Accessing VM's on the Cloud entails management of the underlying tech - whereas the serverless model means code *(where ever it is) can be relied on for hosting the required service. Using a VM seems somewhat static in the face of distributed "software-defined" networking (Cloud.) I'd vote C on this. Not confident this question is focused on what the business models demand - but more on how the tech differs.

upvoted 1 times

beamage 1 year, 11 months ago

Selected Answer: D

Quoted Right out of the Book

upvoted 3 times

TheGinjaNinja 1 year, 11 months ago

Selected Answer: D

Serverless computing is a method of providing backend services on an as-used basis, rather than provisioning and maintaining servers. In this model, the cloud provider (such as AWS Lambda or Azure Functions) manages the servers and underlying infrastructure, while the developer is responsible for writing the code that runs on those servers. This allows the developer to focus on writing code and logic to solve business problems, while the organization can focus on its core business, rather than managing and scaling infrastructure. This can also offer a cost advantage as the user only pays for the resources and execution time needed, instead of having to pay for idle servers.

upvoted 3 times

namangel 2 years, 2 months ago

The answer is A. Infrastructure is fully managed by the CSP.

upvoted 3 times

Rob69420 2 years, 3 months ago

Serverless computing is a method of providing backend services on an as-used basis. A serverless provider allows users to write and deploy code without the hassle of worrying about the underlying infrastructure.

upvoted 1 times

ironman_86 2 years, 3 months ago

can someone explain that?

upvoted 1 times

Question #91 *Topic 1*

A system administrator has provisioned a new web server. Which of the following, in combination, form the best practice to secure the server's OS? (Choose three.)

A. Install TLS certificates on the server.

B. Forward port 80 traffic to port 443.

C. Disable TLS 1.0/1.1 and SSL.

D. Disable password authentication.

E. Enable SSH key access only.

F. Provision the server in a separate VPC.

G. Disable the superuser/administrator account.

H. Restrict access on port 22 to the IP address of the administrator's workstation.

Suggested Answer: *EGH*

Community vote distribution

EGH (50%) 11% 11% Other

TheGinjaNinja Highly Voted 1 year, 5 months ago

Selected Answer: EGH

Disregard my previous comment, I believe it is EGH

E. Enable SSH key access only: SSH key-based authentication is more secure than password-based authentication because a private key is much harder to crack than a password. Additionally, SSH keys can be configured to require a passphrase, which adds an extra layer of security.

G. Disable the superuser/administrator account: Disabling the superuser/administrator account and using a non-privileged account for daily tasks can help prevent privilege escalation attacks. It is also a good practice to use a different account for administrative tasks and to avoid logging in as the root account.

H. Restrict access on port 22 to the IP address of the administrator's workstation: By restricting access to the server's SSH port to a specific IP address, the administrator can ensure that only authorized users are able to access the server. This can help prevent unauthorized access and potential attacks.

upvoted 9 times

Sweety_Certified7 Most Recent 3 months, 1 week ago

Selected Answer: CEH

C. Disable TLS 1.0/1.1 and SSL: Disabling outdated and vulnerable encryption protocols enhances the security of the server's OS by ensuring that only modern and secure protocols are used for communication.

E. Enable SSH key access only: This option enhances security by allowing access to the server only through SSH keys, which are generally more secure than passwords and provide stronger authentication.

H. Restrict access on port 22 to the IP address of the administrator's workstation: By limiting SSH access to specific IP addresses, particularly the administrator's workstation, this option adds an extra layer of security to prevent unauthorized access to the server.

upvoted 1 times

kuzummjakk 3 months, 3 weeks ago

Selected Answer: DEH

Read carefully. "when in combination" and "to secure the server's OS".
The only 3 options that both relate to the OS's security AND relate to each other is DEH. They all relate to SSH (when in combination) and locking down protocols relates to OS security.

mic drop

upvoted 1 times

⊟ **Deeeeez_nuts** 3 months, 3 weeks ago

when in doubt, chatgpt it out

upvoted 1 times

⊟ **badgerino** 4 months, 2 weeks ago

Selected Answer: CEH

CEH makes the most sense.

C. Disable TLS 1.0/1.1 and SSL.
E. Enable SSH key access only. Most Voted
H. Restrict access on port 22 to the IP address of the administrator's workstation. Most Voted

- TLS 1.0 / 1.1 is insecure should be disabled

- Enabling SSH key access only negates the need to disable password auth

- Disabling administrator account is not realistic in a business environment you'll still need IT admins to have access.

- Restricting access to port and IP address of the admin workstation helps secure it the best

upvoted 1 times

⊟ **Monkeyman1500** 4 months, 1 week ago

It also says to disable SSL. Otherwise it would be right

upvoted 1 times

⊟ **FrancisDrake** 6 months ago

Selected Answer: DEG

Disabling passwords scares me but I think that it is correct. Along with disabling admin account (standard admin). SSH key access on the other hand seems like a no brainer.

upvoted 1 times

⊟ **FrancisDrake** 5 months ago

Disabling passwords goes hand in hand with SSH key access.

upvoted 1 times

⊟ **Zak11** 1 year, 2 months ago

Selected Answer: ADE

The best practices to secure the server's OS in combination are:

A. Install TLS certificates on the server.
D. Disable password authentication.
E. Enable SSH key access only.

These three measures help to secure the web server by implementing encryption and securing the authentication process. By disabling password authentication and enabling SSH key access only, the server is less vulnerable to brute-force attacks. Installing TLS certificates on the server helps to encrypt communications, preventing data interception and tampering.

upvoted 2 times

⊟ **BeauChateau** 1 year, 2 months ago

Selected Answer: CDE

The best practices to secure a server's OS are:

C. Disable TLS 1.0/1.1 and SSL. This is because these protocols have known vulnerabilities and should not be used in a production environment.
D. Disable password authentication. This makes it more difficult for attackers to guess or crack user passwords.
E. Enable SSH key access only. This provides a more secure way of authenticating users and prevents password-based attacks.

Option A is not relevant to securing the OS but is instead related to securing the web application running on the server. Option B is not a security best practice, but rather a way to redirect HTTP traffic to HTTPS. Option F and G are not related to securing the server's OS but rather related to network and user management. Option H is a good security practice, but it is not sufficient on its own to secure the OS.

Therefore, the correct options are: C, D, and E.

upvoted 1 times

concepcionz 1 year, 3 months ago

Selected Answer: CDE

I'll go with the followings

"C. Disable TLS 1.0/1.1 and SSL: TLS 1.0/1.1 and SSL have known vulnerabilities, so it's recommended to disable them to ensure secure communication.

D. Disable password authentication: Passwords can be guessed or stolen, so it's recommended to disable password authentication and use public key authentication instead.

E. Enable SSH key access only: Restrict access to the server to only those who possess the private key, which is much more secure than password authentication."

upvoted 1 times

erreyesarroyo 1 year, 5 months ago

just to add more to clusterf...

C. Disable TLS 1.0/1.1 and SSL.
E. Enable SSH key access only.
G. Disable the superuser/administrator account.

Disabling TLS 1.0 and 1.1 and SSL will remove the vulnerabilities in older encryption protocols and ensures that the data is transmitted securely. Enabling SSH key access only will provide a secure method of access to the server and reduces the risk of brute-force attacks on the server. Disabling the superuser/administrator account will prevent the attacker from using the default credentials to gain access to the server.

Other options such as installing TLS certificates, forwarding port 80 traffic to port 443, disabling password authentication, provisioning the server in a separate VPC, and restricting access on port 22 to the IP address of the administrator's workstation, can also be considered as a best practice, but they are not as critical as disabling deprecated protocols, enabling key access, and disabling the superuser/administrator account.

upvoted 4 times

TheGinjaNinja 1 year, 5 months ago

Selected Answer: AEG

I think AEG

upvoted 1 times

ironman_86 1 year, 10 months ago

For me, its A,C,E

upvoted 3 times

Question #92 *Topic 1*

A technician needs to deploy two virtual machines in preparation for the configuration of a financial application next week. Which of the following cloud deployment models should the technician use?

A. XaaS

B. IaaS

C. PaaS

D. SaaS

Suggested Answer: *B*

Community vote distribution

B (100%)

JVen Highly Voted 1 year, 7 months ago

Selected Answer: B

This has to be B

upvoted 8 times

Robenger Most Recent 5 months, 1 week ago

C

Yes, you can deploy virtual machines (VMs) using a Platform as a Service (PaaS) cloud service model. PaaS provides a platform for developers to build and deploy applications without worrying about the underlying infrastructure. PaaS providers manage the infrastructure, including the servers, storage, and networking, while developers focus on building and deploying their applications.

Microsoft Azure provides a PaaS offering called Azure App Service, which allows you to deploy web apps, mobile app backends, and RESTful APIs in a fully managed environment. You can also use Azure Virtual Machines to deploy and manage VMs in the cloud.

upvoted 1 times

kuzummjakk 3 months, 3 weeks ago

not according to Cloud+

upvoted 2 times

FrancisDrake 6 months ago

Selected Answer: B

IaaS. You deploy simple VMs using this.

upvoted 1 times

sheilawu 11 months, 2 weeks ago

Selected Answer: B

If it only state with launch an application, I will go for Paas, since the question is asking for VM, it must be go for B Iaas

upvoted 2 times

BeauChateau 1 year, 2 months ago

Selected Answer: B

The technician should use the IaaS (Infrastructure-as-a-Service) cloud deployment model to provision the virtual machines. IaaS provides virtualized computing resources, such as servers, storage, and networking, on which the user can deploy and manage their own applications and operating systems. The technician can then install the financial application on the virtual machines and configure them as required.

upvoted 3 times

JMorrison 1 year, 8 months ago

Answer is B,

PaaS solutions are often aimed at developers and database administrators (DBAs). These individuals use the provided platform to develop whatever applications or database services are needed by the organization without having to first build the platforms. PaaS solutions also scale quickly and easily, providing consistent development platforms as needs change.

upvoted 3 times

ryanzou 1 year, 9 months ago

PaaS, no doubts

upvoted 1 times

⊟ **jiminycriminal** 1 year, 9 months ago

Agreed, with PaaS you don't "deploy a VM". You need to borrow infrastructure to do that. B should be the answer.

upvoted 3 times

⊟ **achow26** 1 year, 10 months ago

Deploying VMs can be done with IaaS and not PaaS. So Answer should be B.

upvoted 4 times

Question #93 *Topic 1*

A system administrator supports an application in the cloud, which includes a restful API that receives an encrypted message that is passed to a calculator system. The administrator needs to ensure the proper function of the API using a new automation tool. Which of the following techniques would be BEST for the administrator to use to accomplish this requirement?

A. Functional testing

B. Performance testing

C. Integration testing

D. Unit testing

Suggested Answer: *A*

Community vote distribution

A (64%) C (36%)

concepcionz Highly Voted 1 year, 3 months ago

Selected Answer: A

FUNCTIONAL TESTING is a type of software testing that validates the software system against the functional requirements/specifications. The purpose of Functional tests is to test each function of the software application, by providing appropriate input, verifying the output against the Functional requirements.

upvoted 7 times

TheGinjaNinja Highly Voted 1 year, 5 months ago

Selected Answer: C

C. Integration testing would be the best option for the system administrator in this scenario, as it involves testing the interactions between different components of the system to ensure they work together as expected. In this case, the administrator would want to test the interactions between the restful API, the encryption system, and the calculator system to ensure that the encrypted message is properly received and processed by the calculator system through the API.

upvoted 6 times

Securityguy42 Most Recent 2 months, 3 weeks ago

Selected Answer: A

Where are people getting Integration Testing? lol Not in the exam objectives.

• Testing techniques

- Vulnerability testing

- Penetration testing

- Performance testing

- Regression testing

- Functional testing

- Usability testing

upvoted 2 times

Deeeeez_nuts 3 months ago

A is the answer.

Functional testing is a quality assurance test that evaluates whether a system or application meets its specifications—does it do what it's supposed to do?

Integration testing isn't mentioned once in comptia certmaster.

upvoted 2 times

samCarson 3 months, 4 weeks ago

Selected Answer: C

Integration testing would be the most appropriate technique for ensuring the proper function of the API in conjunction with the calculator system.

upvoted 1 times

powpao 8 months ago

Selected Answer: C

Agree with C

upvoted 1 times

⊟ **BeauChateau** 1 year, 2 months ago

Selected Answer: A

A. Functional testing would be the BEST technique for the administrator to use in order to ensure the proper function of the API using a new automation tool. Functional testing is a testing technique that involves testing the functionality of an application to ensure that it performs the intended tasks correctly. In this case, the API is expected to receive an encrypted message and pass it to a calculator system, so functional testing can help ensure that the API is working properly and is passing the message correctly to the calculator system. Other types of testing, such as performance testing, integration testing, and unit testing, may also be useful in some contexts, but functional testing is the most appropriate option for ensuring that the API is functioning properly.

upvoted 2 times

⊟ **craigbharrell** 1 year, 3 months ago

Selected Answer: A

It literally says Function in the question, he is testing the function (also an exam objective btw, integration is not)

upvoted 3 times

⊟ **FrancisDrake** 6 months ago

That would be the 'tell' which CompTIA is so fond of.

upvoted 1 times

Question #94 *Topic 1*

A cloud solutions architect needs to determine the best strategy to deploy an application environment in production, given the following requirements:
* No downtime
* Instant switch to a new version using traffic control for all users
Which of the following deployment strategies would be the BEST solution?

A. Hot site

B. Blue-green

C. Canary

D. Rolling

Suggested Answer: *B*

Community vote distribution

B (100%)

BeauChateau Highly Voted 1 year, 2 months ago

Selected Answer: B

The BEST deployment strategy to meet the given requirements is the "Blue-green" deployment.

In a blue-green deployment, two identical environments are maintained, one for production (blue) and the other for the next version or release (green). All production traffic is initially directed to the blue environment, while the green environment is prepared and tested thoroughly without affecting the production environment. Once the green environment is ready, the traffic is switched instantly from the blue to the green environment using traffic control, with no downtime or impact on users.

upvoted 5 times

Stonetales987 Most Recent 1 month, 3 weeks ago

Selected Answer: B

• Blue-green deployment These deployments always have an active system and one that is used for testing. When testing is complete, the testing system becomes active, and the former production system is available for testing. One system is labeled “blue” and the other “green.”

upvoted 2 times

Sal 1 year, 8 months ago

Answer is blue-green

https://www.seaflux.tech/blogs/blue-green-deployment-in-azure-cloud

upvoted 1 times

Question #95 *Topic 1*

A cloud security analyst is implementing a vulnerability scan of the web server in the DMZ, which is running in an IaaS compute instance. The default inbound firewall settings are as follows:

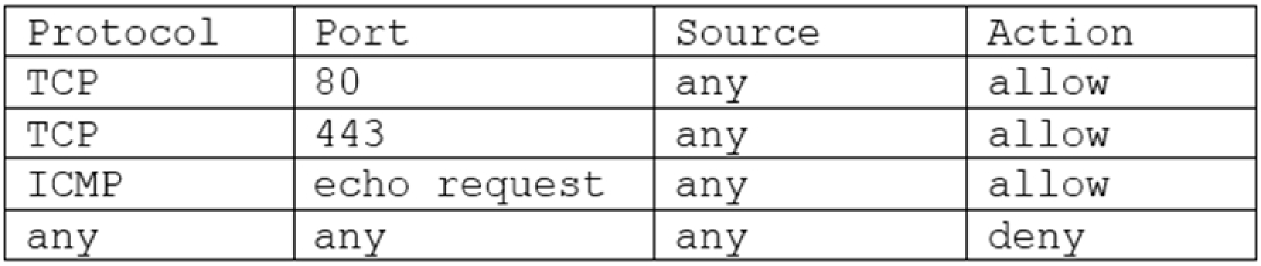

Protocol	Port	Source	Action
TCP	80	any	allow
TCP	443	any	allow
ICMP	echo request	any	allow
any	any	any	deny

Which of the following will provide the analyst with the MOST accurate report?

A. An agent-based scan

B. A network vulnerability scan

C. A default and common credentialed scan

D. A network credentialed vulnerability scan

Suggested Answer: *D*

Community vote distribution

D (67%) B (33%)

Jay987654 5 months ago

Selected Answer: D

The analyst should use Option D: A network credentialed vulnerability scan.

A network credentialed vulnerability scan provides the most accurate report because it uses valid credentials to log in to the scanned systems and gather detailed information about the operating system and installed software, including configuration issues and missing security patches. This type of scan can identify vulnerabilities that may not be visible during an unauthenticated scan.

upvoted 3 times

Jhonattan0032 5 months, 1 week ago

Selected Answer: D

D. A network credentialed vulnerability scan

This option allows the vulnerability scanner to perform an authenticated scan of the server.

upvoted 3 times

Pongsathorn 9 months, 1 week ago

Selected Answer: B

For performing a vulnerability scan on the web server in the DMZ running in an IaaS compute instance with the provided firewall settings, the most appropriate choice is **B. A network vulnerability scan**.

Here's why:

- **Agent-based scan (A)** typically involves installing an agent on the target system. In this case, since you're dealing with a DMZ web server in an IaaS compute instance, it may not be feasible to install an agent directly on the instance, especially if it's externally facing. Additionally, agent-based scans are more commonly used for endpoints and servers within your organization's network.

upvoted 3 times

Pongsathorn 9 months, 1 week ago

- **Network vulnerability scan (B)** is the best choice for this scenario. Network vulnerability scanning tools, like Nessus or OpenVAS, can scan the target system over the network without requiring an agent. Given the provided firewall settings, the web server allows traffic on ports 80 (HTTP), 443 (HTTPS), and ICMP echo requests. Network vulnerability scanners can assess the vulnerabilities of the web server based on these open ports and the services running behind them.

- **Default and common credentialed scan (C)** usually involves using default or common credentials to check for vulnerabilities on the target system. This might not be suitable for an externally facing web server, as it could be a security risk to use credentials that are not specifically configured for this purpose.

upvoted 1 times

Pongsathorn 9 months, 1 week ago

- **Network credentialed vulnerability scan (D)** typically implies using credentials to authenticate with the target system and perform a vulnerability scan. While this can provide more in-depth results, it may not be suitable for an externally facing web server in a DMZ due to security concerns. Moreover, the provided firewall settings might not allow the necessary ports and protocols for credential-based scans.

In summary, a network vulnerability scan is the most accurate and appropriate choice for assessing the security of the web server in the DMZ with the given firewall settings.

upvoted 2 times

Question #96 *Topic 1*

A systems administrator needs to configure SSO authentication in a hybrid cloud environment. Which of the following is the BEST technique to use?

A. Access controls

B. Federation

C. Multifactor authentication

D. Certificate authentication

Suggested Answer: *B*

Community vote distribution

B (100%)

AustinKelleyNet 5 months ago

Selected Answer: B

This is correct

upvoted 3 times

Question #97 *Topic 1*

A systems administrator wants to verify the word "qwerty" has not been used as a password on any of the administrative web consoles in a network. Which of the following will achieve this goal?

A. A service availability scan

B. An agent-based vulnerability scan

C. A default and common credentialed scan

D. A network port scan

Suggested Answer: *C*

Community vote distribution

C (100%)

Jay987654 5 months ago

Selected Answer: C

C: A default and common credentialed scan.

upvoted 1 times

Jay987654 5 months ago

Option C: A default and common credentialed scan.

A default and common credentialed scan checks for commonly used or default passwords on network devices. This type of scan can help identify if weak passwords like “qwerty” are being used, which could potentially compromise the security of the network.

upvoted 2 times

Question #98 *Topic 1*

An administrator has been informed that some requests are taking a longer time to respond than other requests of the same type. The cloud consumer is using multiple network service providers and is performing link load balancing for bandwidth aggregation. Which of the following commands will help the administrator understand the possible latency issues?

A. ping

B. ipconfig

C. traceroute

D. netstat

Correct Answer: *C*

Community vote distribution

C (100%)

eacunha 6 months ago

Selected Answer: C

C. traceroute:

O comando `traceroute` é uma ferramenta que rastreia a rota que os pacotes levam de um ponto para outro em uma rede. Ele exibe a lista de saltos (hops) que um pacote de dados faz para chegar ao destino. Isso pode ajudar o administrador a identificar onde ocorrem atrasos ou latências na rede.

Ao usar o `traceroute`, o administrador pode analisar cada salto e avaliar o tempo de resposta em cada ponto intermediário. Isso fornece insights sobre onde a latência pode estar ocorrendo, permitindo que o administrador localize e aborde possíveis problemas de desempenho na rede.

upvoted 1 times

AustinKelleyNet 1 year, 5 months ago

Selected Answer: C

The given answer is correct. This will show you the route and not just how much latency, which ping will do

upvoted 2 times

Question #99 *Topic 1*

A company has an in-house-developed application. The administrator wants to utilize cloud services for additional peak usage workloads. The application has a very unique stack of dependencies. Which of the following cloud service subscription types would BEST meet these requirements?

A. PaaS

B. SaaS

C. DBaaS

D. IaaS

Suggested Answer: *D*

Community vote distribution

D (100%)

AustinKelleyNet Highly Voted 5 months ago

Selected Answer: D

SaaS- Provider takes care of everything

PaaS- You take care of data and the application

IaaS- You take care of software they take care of hardware

In this case I think they want to take care of all the software, so it is IaaS

upvoted 6 times

Question #100 *Topic 1*

A systems administrator notices that a piece of networking equipment is about to reach its end of support. Which of the following actions should the administrator recommend?

A. Update the firmware.

B. Migrate the equipment to the cloud.

C. Update the OS.

D. Replace the equipment.

Correct Answer: *D*

⊟ **FrancisDrake** 6 months ago

It says 'end of support' NOT 'end of life'. I suppose replace the equipment is the BEST answer but in the real world it seems you would call the vendor to extend the support.

upvoted 1 times

⊟ **Securityguy42** 2 months, 3 weeks ago

Well, calling the vendor to extend support is not an option. This is not the real world. Its CompTIA's world.

upvoted 1 times

Question #101 *Topic 1*

An organization will be deploying a web application in a public cloud with two web servers, two database servers, and a load balancer that is accessible over a single public IP. Taking into account the gateway for this subnet and the potential to add two more web servers, which of the following will meet the minimum IP requirement?

A. 192.168.1.0/26

B. 192.168.1.0/27

C. 192.168.1.0/28

D. 192.168.1.0/29

Correct Answer: *C*

Community vote distribution

C (100%)

Pisces225 3 months ago

Selected Answer: C

Some of you guys are going to fail your Net+ or anything else network related. Here's the proper breakout. C is indeed correct, but some of these comments are just terrible and no one has laid it out, so here it is:

1 - gateway (this will be our public IP but will have internal as well as all gateways always do!)
2 - load balancer
3 - web server #1
4 - web server #2
5 - database server
6 - database server
7 - new webserver #1
8 - new web server #2
9 - broadcast IP

Our /28 is going to provide a maximum of 16 IPs which will cover the 9 addresses needed for current and future expansion.

upvoted 3 times

Stonetales987 1 month, 3 weeks ago

10 - Network Address - The first address of each subnet is also unusable and is used to identify the subnet itself and cannot be assigned to a host.

upvoted 1 times

db93ae3 3 months, 4 weeks ago

The load balancer will have a public IP. Doesn't that exclude it from the private network, leading to /29 being sufficient for 6 hosts

upvoted 1 times

Not_That_Guy 1 year, 8 months ago

Selected Answer: C

7 addresses are needed, /28 provides 8 addresses but only 6 of those are available after discounting the network and broadcast addresses. So /29 is the smallest subnet that meets requirements.

upvoted 1 times

Not_That_Guy 1 year, 8 months ago

Oops, I reversed those. /29 provides 8 addresses; /28 is the smallest subnet that meets requirements.

upvoted 2 times

ryanzou 1 year, 9 months ago

28 is correct

upvoted 2 times

achow26 1 year, 10 months ago

Total IP requirement is 7 including future growth so a /29 can provide 8 IPs. So why a /28?

upvoted 1 times

concepcionz 1 year, 3 months ago

192.168.1.0/29

192.168.1.8/29

192.168.1.16/29

192.168.1.24/29

and so on

So the Network is 192.168.1.0 and the Broadcast 192.168.1.7, that leaves 6 usable address (1,2,3,4,5,6)

upvoted 3 times

i_bird 1 year, 9 months ago

/29 provides 6 not 8..

correct ans: /28 = 14 Ips

upvoted 5 times

Question #102 *Topic 1*

All of a company's servers are currently hosted in one cloud MSP. The company created a new cloud environment with a different MSP. A cloud engineer is now tasked with preparing for server migrations and establishing connectivity between clouds. Which of the following should the engineer perform FIRST?

A. Peer all the networks from each cloud environment.

B. Migrate the servers.

C. Create a VPN tunnel.

D. Configure network access control lists.

Suggested Answer: *C*

Community vote distribution

C (100%)

Stonetales987 1 month, 3 weeks ago

Selected Answer: C

VPN first (Secure the connection) ,then peer. The VPN is needed because you are preparing for server migrations and establishing connectivity between TWO cloud environments hosted by DIFFERENT Managed Service Providers (MSPs).

upvoted 1 times

FrancisDrake 5 months, 4 weeks ago

If the answer is VPN then that migration is going to take a long time.

upvoted 1 times

utied 6 months, 2 weeks ago

Selected Answer: C

Peering: Establishes direct communication between VPCs without traversing the public internet, minimizing latency and enhancing performance and security. Appear to customer as single network.

I think the key is "The company created a new cloud environment with a different MSP. A cloud engineer is now tasked with preparing for server migrations and establishing connectivity between clouds." Using public internet is not peering.

C. Create a VPN tunnel.

upvoted 1 times

sham_ba_lam 1 year, 4 months ago

Selected Answer: C

They are all good steps, but FIRST would be to establish a VPN connection that the data migration process is secure.

upvoted 1 times

CapJackSparrow 1 year, 5 months ago

google always seems to be the wildcard... try googling something like "peer aws and azure VPC" the only-thing I can find is VPN tunnels.

upvoted 3 times

Francois1984 10 months, 1 week ago

i recently started using chatgpt. this is what it says:

Before migrating servers and workloads, establishing connectivity between the two cloud environments is crucial. You'll need to set up networking components to ensure communication between servers in both environments. This generally involves creating Virtual Private Clouds (VPCs), Virtual Networks, VPN connections, or Direct Connect links, depending on the cloud providers you are working with.

upvoted 1 times

buckthesystem 7 months, 3 weeks ago

You should never use ChatGPT to find factual information. Due to the underlying technology it can be entirely inaccurate and/or make up information that sounds convincingly correct. There are many example with even the latest models that have been "opened up" to the internet are getting things scarily wrong.

upvoted 1 times

Not_That_Guy 1 year, 8 months ago

I think peering works within the same cloud provider, because this question refers to different providers then VPN Tunnel is probably correct.

upvoted 4 times

⊟ **tonytonyyyyy** 1 year, 8 months ago

I think it's A.

https://cloud.google.com/vpc/docs/vpc-peering

Google Cloud VPC Network Peering allows internal IP address connectivity across two Virtual Private Cloud (VPC) networks regardless of whether they belong to the same project or the same organization.

VPC Network Peering enables you to connect VPC networks so that workloads in different VPC networks can communicate internally. Traffic stays within Google's network and doesn't traverse the public internet.

upvoted 2 times

Question #103 *Topic 1*

A web server has been deployed in a public IaaS provider and has been assigned the public IP address of 72.135.10.100. Users are now reporting that when they browse to the website, they receive a message indicating the service is unavailable. The cloud administrator logs into the server, runs a netstat command, and notices the following relevant output:

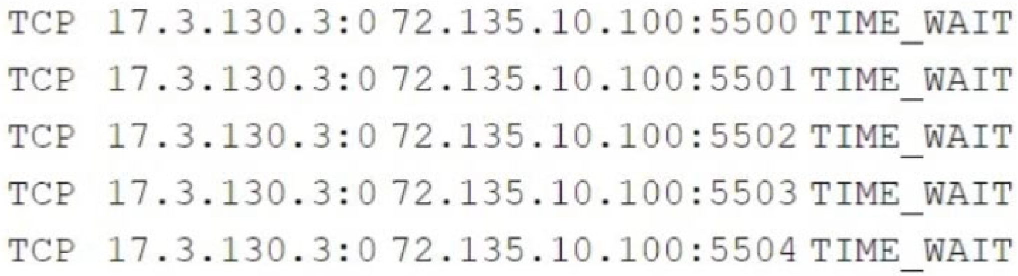

```
TCP 17.3.130.3:0 72.135.10.100:5500 TIME_WAIT
TCP 17.3.130.3:0 72.135.10.100:5501 TIME_WAIT
TCP 17.3.130.3:0 72.135.10.100:5502 TIME_WAIT
TCP 17.3.130.3:0 72.135.10.100:5503 TIME_WAIT
TCP 17.3.130.3:0 72.135.10.100:5504 TIME_WAIT
```

Which of the following actions should the cloud administrator take to resolve the issue?

A. Assign a new IP address of 192.168.100.10 to the web server.

B. Modify the firewall on 72.135.10.100 to allow only UDP.

C. Configure the WAF to filter requests from 17.3.130.3.

D. Update the gateway on the web server to use 72.135.10.1.

Correct Answer: *C*

Community vote distribution

C (80%) D (20%)

FrancisDrake 4 months, 2 weeks ago

Selected Answer: C

of all the suggested answers C is most likely

upvoted 1 times

FrancisDrake 4 months, 2 weeks ago

Hard to imagine D as the answer. 72.135.10.100 is a public ip address. For D to be the correct (gateway address of 72.135.10.1) the subnet would have to be 255.255.255.0 and that makes no sense.

upvoted 1 times

cobbs 10 months ago

Selected Answer: C

Denial of service attack. Since the attacker's public IP is taking up all the available network sockets (not all shown), legitimate customers are unable to connect. Block this with the firewall (WAF). Port 0 will show up if there’s fragmented IP traffic, like a DNS response which exceeds the historic maximum size of 512 bytes.

upvoted 2 times

concepcionz 1 year, 3 months ago

Selected Answer: C

Im going with C

upvoted 1 times

mattygster 1 year, 4 months ago

i dont see why D would be the correct choice, NETSTAT "TIME_WAIT" indicates normal traffic and waiting for any more packets after the syn/ack/syn, so there must be connections to the webserver from client machines. It seems like these answer choices are not related to the actual problem. BUT.....if i had to choose one, maybe the number of connections coming from a single IP could be signs of a resource exhaustion attack. In theory, if an attacker was hitting the server hard, it could cause an apache/IIS server to crash thus resulting in a "service not available".

I wish there was another choice pointing to the listening port needing to be reconfigured.

upvoted 1 times

beamage 1 year, 5 months ago

Selected Answer: D

Period.....

upvoted 1 times

⊟ **FrancisDrake** 5 months, 4 weeks ago

There's not enough information to make that determination. You should try explaining your reasoning. That would be more helpful.

upvoted 3 times

⊟ **CapJackSparrow** 1 year, 5 months ago

I'm looking at C pretty hard here, possible attack?

upvoted 1 times

⊟ **Not_That_Guy** 1 year, 8 months ago

This is one of those "Best Answer" scenarios; there is no "Good" answer. A, B, and C can be eliminated leaving D as the best answer even though D isn't a clear answer with the information provided. A - private IP on a public web server, NO. B - using only UDP, NO. C - block the user with issues vs resolving the issues, NO.

upvoted 2 times

⊟ **ryanzou** 1 year, 9 months ago

sometimes, WAF will filter out the SYNC packet during the TCP handshake. C may be the answer.

upvoted 1 times

⊟ **jiminycriminal** 1 year, 9 months ago

None of these makes sense. The rightmost IP address should be the foreign address, not the public IP of the server. Who made this?

upvoted 1 times

⊟ **jiminycriminal** 1 year, 9 months ago

I'm assuming the application gateway has been assigned the same IP as the server, so that's why the answer is D. It's calling out to the wrong IP? Man I really have no idea.

upvoted 1 times

⊟ **i_bird** 1 year, 9 months ago

TCP TIME_WAIT is a normal TCP protocol operation, it means after delivering the last FIN-ACK, client side will wait for double maximum segment life (MSL) Time to pass to be sure the remote TCP received the acknowledgement of its connection termination request. By default, MSL is 2 minutes. For the maximum, it can stay in TIME_WAIT for 4 minutes known as two MSL

From Network perspective, TCP TIME_WAIT status is just a normal behavior that after closing the session, TCP stack will hold the high port for little more time to ensure the other side receive the last FIN-ACK packet and no more data will be received in this conversation. TIME_WAIT is not the problem. I would suggest you focus on RPC error for this issue.

https://docs.microsoft.com/en-us/answers/questions/230227/time-wait-from-netstat.html

upvoted 1 times

⊟ **i_bird** 1 year, 9 months ago

Can someone explain why it is D?

upvoted 1 times

Question #104 *Topic 1*

A cloud administrator has been using a custom VM deployment script. After three months of use, the script no longer joins the LDAP domain. The cloud administrator verifies the account has the correct permissions. Which of the following is the MOST likely cause of the failure?

A. Incorrect encryption ciphers

B. Broken trust relationship

C. Invalid certificates

D. Expired password

Suggested Answer: *D*

Community vote distribution

D (60%) B (40%)

eastwouldd 5 months, 3 weeks ago

Constantly run into this issue during os migrations. Its almost always has been expired passwords

upvoted 1 times

Kailer 1 year, 3 months ago

I literally have this issue all the time at my current role and most of the times is B. Broken trust relationship. Also, chatgpt agrees on B so imma go with

upvoted 1 times

FrancisDrake 12 months ago

What is the cause of the broken trust relationship?

upvoted 2 times

Pongsathorn 1 year, 3 months ago

Selected Answer: D

The MOST likely cause of the failure in the custom VM deployment script to join the LDAP domain, even when the account has the correct permissions, is:

D. Expired password.

Here's why:

LDAP authentication typically relies on username and password credentials to establish trust between a system and the LDAP server. If the password for the account being used in the deployment script has expired, it will prevent successful authentication and the ability to join the LDAP domain.

To resolve this issue, the cloud administrator should verify the password's expiration status and update it if necessary. Additionally, it's essential to ensure that the script is correctly handling password authentication, including any password change requirements.

upvoted 3 times

ROCompTIA 1 year, 6 months ago

Selected Answer: B

Because Ldap

upvoted 2 times

Question #105 *Topic 1*

A cloud administrator is managing an organization's infrastructure in a public cloud. All servers are currently located in a single virtual network with a single firewall that all traffic must pass through. Per security requirements, production, QA, and development servers should not be able to communicate directly with each other.

Which of the following should an administrator perform to comply with the security requirement?

A.

✑ Create separate virtual networks for production, QA, and development servers.

✑ Move the servers to the appropriate virtual network.

✑ Apply a network security group to each virtual network that denies all traffic except for the firewall.

B.

✑ Create separate network security groups for production, QA, and development servers.

✑ Apply the network security groups on the appropriate production, QA, and development servers.

✑ Peer the networks together.

C.

✑ Create separate virtual networks for production, QA, and development servers.

✑ Move the servers to the appropriate virtual network.

✑ Peer the networks together.

D.

✑ Create separate network security groups for production, QA, and development servers.

✑ Peer the networks together.

✑ Create static routes for each network to the firewall.

Correct Answer: *B*

achow26 Highly Voted 1 year, 10 months ago

Answer should be A. If the peering is in place between networks, all the traffic is not passing thru the single firewall.

upvoted 18 times

brickcity86 1 year, 6 months ago

Agreed, A is the only option mentioning restricting traffic through the firewall

upvoted 1 times

[Removed] 1 year, 5 months ago

I dont think, it is correct answer....

Simply deploying FW does not make any impact on traffic unless you route the traffic using UDR. It will keeep on communicating directly.....

even NSG also says, vnet to vnet traffic is allowed by default....

upvoted 1 times

Pisces225 Most Recent 3 months ago

Going with A, odd there's no voting comments available.

upvoted 1 times

FrancisDrake 5 months, 4 weeks ago

I would think that the answer is A. The scenario stipulates no direct connection between networks.

upvoted 2 times

backdooranon 8 months, 2 weeks ago

If you create multiple VLANs and then peer them together it does not fulfill "no direct connection between networks" condition

upvoted 1 times

Pongsathorn 9 months, 1 week ago

The correct option to comply with the security requirement of ensuring that production, QA, and development servers should not be able to communicate directly with each other in a public cloud environment is:

A.

- Create separate virtual networks for production, QA, and development servers.
- Move the servers to the appropriate virtual network.
- Apply a network security group to each virtual network that denies all traffic except for the firewall.

upvoted 1 times

Pongsathorn 9 months, 1 week ago

Here's why:

- Creating separate virtual networks for each server group isolates them from each other, meeting the requirement for no direct communication.
- Moving servers to their respective virtual networks ensures they are in the correct network segment.
- Applying network security groups (NSGs) to each virtual network to deny all traffic except for the firewall enforces the desired isolation while allowing traffic to pass through the firewall for necessary communication.

Option B suggests peering the networks together, which would enable communication between them, violating the security requirement. Option C, while suggesting separate virtual networks, also suggests peering them, which again would allow communication between them, not meeting the requirement. Option D suggests peering networks together and creating static routes, which is not necessary and doesn't guarantee isolation as required.

upvoted 1 times

maelo 9 months, 3 weeks ago

Answer should be A. Candidate B suggests network security groups + peering. I see no multiple networks created, just security groups. A allows strict traffic management.

upvoted 1 times

AustinKelleyNet 1 year, 5 months ago

Gotta be A

upvoted 1 times

[Removed] 1 year, 5 months ago

Answer is D
Separate Networks for Prod, Dev n QA.
Peer all these network.
Defining route will override system route and will force traffic to move via FW.
#HubNspoke connctivity.. :)

upvoted 1 times

CapJackSparrow 1 year, 5 months ago

Hub-and-spoke VPC design, with separate Production, Development, and Research spoke VPCs connected to a central "hub" VPC.

Peering
Peering connects two or more virtual networks. The virtual networks appear to consumers as a single network. In addition, fast connectivity is provided between the two networks, making data and resource access very efficient.

Peering is used in the hub-and-spoke model to connect the spoke networks with the hub network. Note that the spoke networks are not peered to each other in the hub-and-spoke model.

upvoted 1 times

scott5010 1 year, 7 months ago

answer is B, all traffic MUST pass through the firewall and peering to create a hub spoke network

upvoted 2 times

ryanzou 1 year, 9 months ago

A is the answer

upvoted 1 times

Question #106 *Topic 1*

A cloud administrator is upgrading a cloud environment and needs to update the automation script to use a new feature from the cloud provider. After executing the script, the deployment fails. Which of the following is the MOST likely cause?

A. API incompatibility

B. Location changes

C. Account permissions

D. Network failure

Suggested Answer: *A*

Community vote distribution

A (88%) 13%

⊟ **Not_That_Guy** Highly Voted 👍 2 years, 2 months ago

A is if you run a script written for one API against a different API (i.e. different cloud providers). More likely the admin is trying to access a feature via API which they have not subscribed to yet.

upvoted 7 times

⊟ **CapJackSparrow** 1 year, 11 months ago

doesn't say anything about using another provider.

upvoted 2 times

⊟ **rhnorwoodjr** Most Recent 🕘 4 months, 3 weeks ago

Selected Answer: C

My gut goes toward A for many of the reasons stated. But in this situation the assumption could be made that new permissions are required to use the new "feature" and thus C could be just as valid. I have run into this situation a number of times working in three different AWS environments. It wasn't an API issue but an IAM permission that needed to be added.

upvoted 1 times

⊟ **Pongsathorn** 1 year, 3 months ago

Selected Answer: A

The MOST likely cause of a deployment failure after updating an automation script to use a new feature from the cloud provider is:

A. API incompatibility

Here's why:

- API (Application Programming Interface) changes, including the introduction of new features, can lead to compatibility issues with existing scripts. If the automation script relies on specific API calls or features that have changed or are no longer supported, it can result in a deployment failure.

The other options (B, C, and D) could also potentially cause deployment issues, but given the scenario of updating the script to use a new cloud provider feature, API incompatibility is the most probable reason for the failure.

upvoted 3 times

⊟ **SecPlus2022** 1 year, 6 months ago

Selected Answer: A

As stated by "Not_That_Guy"

upvoted 2 times

⊟ **concepcionz** 1 year, 9 months ago

Selected Answer: A

API incompatibility

upvoted 1 times

⊟ **sham_ba_lam** 1 year, 11 months ago

Selected Answer: A

The cloud administrator may need to modify the script to ensure that it is compatible with the new APIs and can successfully use the new feature.

Location changes, account permissions, and network failures can also cause deployment failures, but they are less likely to be the cause in this scenario, as these issues typically have a broader impact on the entire cloud environment, not just a single deployment.

upvoted 1 times

ironman_86 2 years, 3 months ago

i think the answer is A

upvoted 3 times

Question #107 *Topic 1*

A systems administrator has been asked to restore a VM from backup without changing the current VM's operating state. Which of the following restoration methods would BEST fit this scenario?

A. Alternate location

B. Rolling

C. Storage live migration

D. In-place

Suggested Answer: *D*

Community vote distribution

D (67%) A (27%) 7%

⊟ **TheGinjaNinja** Highly Voted 👍 11 months, 2 weeks ago

Selected Answer: D

An In-place restore is when the data is restored to the original location and overwrites the existing data. This method is best when the administrator wants to restore the VM without changing the current VM's operating state or configuration.

An alternate location restore is when the data is restored to a different location, rolling restore is when the data is restored in a rolling fashion, one node at a time. Storage live migration restore is when the data is restored by migrating the storage to the new location.

upvoted 7 times

⊟ **Pongsathorn** Most Recent ⊙ 3 months, 3 weeks ago

UNDERSTAND RESTORATION METHODS

You may specify different restore locations, depending on your needs.

In-Place/Overwrite Restoral

In the case of an actual incident, you may recover data back to its original location, overwriting whatever data currently resides there. For example, you may be rolling data back to a known point in time or setting a server back to an original configuration.

Alternate Location Restoral

In the restore program, specify an alternate location path to define where to restore the data. This is a great option when testing to ensure that the backup jobs work as intended and that you know the recovery procedure. In a test scenario, the production server is still up and in use, so you need to restore data to another location. VMs are often great targets for alternate location restorals for testing.

upvoted 1 times

⊟ **maelo** 4 months, 3 weeks ago

Selected Answer: A

Only alternate location will not interfere with existing VM.

upvoted 2 times

⊟ **SecPlus2022** 6 months, 2 weeks ago

Selected Answer: D

Given In-Place VM restoration is an Azure feature, I'm going with "D". Yes, technically, an alternate location restoration would not affect the state of the VM in questions, but "D" is the most logical answer. CompTia, I'm sure, wants to make sure you're aware of this cloud capability.

https://azure.microsoft.com/en-us/blog/an-easy-way-to-bring-back-your-azure-vm-with-in-place-restore/

upvoted 3 times

⊟ **Maged_nader12** 7 months ago

just don't know why I paid for this site :D , three different answers with the same point at the end !!!!!

upvoted 4 times

⊟ **BeauChateau** 8 months, 3 weeks ago

Selected Answer: A

A. Alternate location restoration method would be the BEST fit for this scenario. With the alternate location restoration method, the backup of the VM is restored to a different location without changing the current state of the VM. The administrator can then verify the backup and perform additional tasks, such as testing or updating the restored VM, before switching over to the restored VM.

upvoted 2 times

⊟ **sham_ba_lam** 11 months ago

Selected Answer: C

Storage Live Migration involves moving the virtual disk files associated with a VM from one storage location to another, without disrupting the running state of the VM. This allows the systems administrator to restore a VM from backup while keeping the current operating state intact.

An in-place restoration involves overwriting the existing virtual disk files, which would result in a change to the current operating state of the VM.

upvoted 1 times

⊟ **beamage** 11 months, 2 weeks ago

What are you talking about? In place restore while change the state of the current VM....

upvoted 1 times

Question #108 *Topic 1*

Due to a policy change, a few of a customer's application VMs have been migrated to synchronously replicated storage. The customer now reports that performance is lower. The systems administrator checks the resource usage and discovers CPU utilization is at 60% and available memory is at 30%. Which of the following is the MOST likely cause?

A. There is not enough vCPU assigned.

B. The application is not compatible with the new settings.

C. The new configuration is adding latency.

D. The memory of the VM is underallocated.

Correct Answer: *C*

Community vote distribution

C (100%)

⊟ **TheGinjaNinja** Highly Voted 1 year, 5 months ago

Selected Answer: C

I believe it is C

upvoted 8 times

⊟ **Stonetales987** Most Recent 1 month, 3 weeks ago

Selected Answer: C

Synchronously Replicated Storage would add latency because each write operation must be confirmed by the secondary site before it can be considered complete.

upvoted 1 times

⊟ **FrancisDrake** 5 months ago

Selected Answer: C

Because synchronous replication adds latency.

https://blog.purestorage.com/purely-informational/synchronous-replication-vs-aynchronous-replication/

upvoted 2 times

⊟ **ROCompTIA** 1 year ago

Selected Answer: C

Most likely cause of the performance degradation is the added latency from the new storage configuration

upvoted 1 times

⊟ **maelo** 1 year, 2 months ago

Selected Answer: C

Application is running in VM. VM was moved. Storage profile should be transparent to the application and its compatibility, but overall storage behaviour adds complexity, so latency. Synchronity means waiting for all storage parties to finish a task. Wait = Latency.
Example: Move VM from single SDD to HDD-RAID-1. App will work, but write slower.

This caught my attention: "... a FEW of a customer's application VMs have been migrated ...". Does this change anything?

upvoted 1 times

⊟ **CapJackSparrow** 1 year, 5 months ago

OR maybe...

define the problem (The new configuration is adding latency)
establish a theory of probable cause (The application is not compatible with the new settings)

upvoted 1 times

⊟ **CapJackSparrow** 1 year, 5 months ago

If it was not compatible, wouldn't it not work at all?

upvoted 2 times

Question #109 *Topic 1*

A systems administrator wants to ensure two VMs remain together on the same host. Which of the following must be set up to enable this functionality?

A. Affinity

B. Zones

C. Regions

D. A cluster

Correct Answer: *A*

Community vote distribution

A (80%) D (20%)

⊟ **Binkyboobah** 2 months, 2 weeks ago

Selected Answer: A

It's A

upvoted 2 times

⊟ **nmap_king_22** 3 months, 2 weeks ago

Selected Answer: D

So, to achieve the desired functionality of keeping two VMs together on the same host, configuring them to run within the same cluster is the appropriate approach. This ensures that the VMs will share the same physical server resources within the cluster, assuming the cluster configuration allows for it.

upvoted 1 times

⊟ **AustinKelleyNet** 11 months, 1 week ago

Selected Answer: A

https://vkinfotek.com/azureqa/what-is-affinity-group-azure.html#:~:text=By%20creating%20an%20affinity%20group%2C%20one%20can%20group,network%20latency%20and%20increases%20the%20performa

upvoted 2 times

Question #110 *Topic 1*

An organization is implementing a new requirement to facilitate users with faster downloads of corporate application content. At the same time, the organization is also expanding cloud regions. Which of the following would be suitable to optimize the network for this requirement?

A. Implement CDN for overall cloud application.

B. Implement auto-scaling of the compute resources.

C. Implement SR-IOV on the server instances.

D. Implement an application container solution.

Correct Answer: *A*

Community vote distribution

A (100%)

ROCompTIA 6 months, 3 weeks ago

Selected Answer: A

Content Delivery Network (CDN) for the overall cloud application would be the most suitable solution

upvoted 1 times

beamage 11 months, 2 weeks ago

CDN Content Delivery Network means putting data closer to the customer regionally

upvoted 4 times

beamage 11 months, 2 weeks ago

Selected Answer: A

Right out of the Book

upvoted 3 times

Sal 1 year, 2 months ago

Could it be C?

https://access.redhat.com/documentation/en-us/red_hat_openstack_platform/13/html/network_functions_virtualization_planning_and_configuration_guide/part-sriov-nfv-configuration

upvoted 1 times

TheGinjaNinja 11 months, 4 weeks ago

Regionally, it would be best to implement CDN

upvoted 4 times

Question #111 *Topic 1*

After a few new web servers were deployed, the storage team began receiving incidents in their queue about the web servers. The storage administrator wants to verify the incident tickets that should have gone to the web server team. Which of the following is the MOST likely cause of the issue?

A. Incorrect assignment group in service management

B. Incorrect IP address configuration

C. Incorrect syslog configuration on the web servers

D. Incorrect SNMP settings

Correct Answer: *A*

Community vote distribution

A (100%)

CXSSP 4 weeks, 1 day ago

Selected Answer: A

By reviewing the service management assignment groups, the storage administrator can identify and rectify the issue of web server incidents being routed incorrectly. This will ensure that incidents are directed to the appropriate team for faster resolution.

upvoted 3 times

Question #112 *Topic 1*

A systems administrator is deploying a solution that includes multiple network I/O-intensive VMs. The solution design requires that vNICs of the VMs provide low- latency, near-native performance of a physical NIC and data protection between the VMs. Which of the following would BEST satisfy these requirements?

A. SR-IOV

B. GENEVE

C. SDN

D. VLAN

Suggested Answer: *A*

Community vote distribution

A (100%)

Brianhealey136 Highly Voted 1 year ago

The single root I/O virtualization (SR-IOV) interface is an extension to the PCI Express (PCIe) specification. SR-IOV allows a device, such as a network adapter, to separate access to its resources among various PCIe hardware functions.

upvoted 8 times

nmap_king_22 Most Recent 3 months, 2 weeks ago

Selected Answer: A

To satisfy the requirements of low-latency, near-native performance of a physical NIC, and data protection between the VMs for a solution with multiple network I/O-intensive VMs, the BEST option would be:

A. SR-IOV (Single Root I/O Virtualization).

upvoted 2 times

Pongsathorn 3 months, 3 weeks ago

Single Root Input/Output Virtualization

Server motherboards include the PCIe bus, and one common component on the bus is a network interface card (or more than one). With traditional virtualization, the hypervisor controlled access between VMs and the bus. This control added a layer of complexity and latency to the structure.

The PCIe bus has been extended to address this concern. The single root input/output virtualization (SR-IOV) functionality permits direct access between the VM and the PCIe bus and, by extension, the NICs on the bus. This permits much faster communication between the physical network connection and the VMs on the host server.

The SR-IOV is particularly attractive to private cloud administrators that may be trying to gain as much efficiency as possible with the hardware deployments in their private data centers.

upvoted 3 times

Question #113 *Topic 1*

A global web-hosting company is concerned about the availability of its platform during an upcoming event. Web traffic is forecasted to increase substantially during the next week. The site contains mainly static content. Which of the following solutions will assist with the increased workload?

A. DoH

B. WAF

C. IPS

D. CDN

Suggested Answer: *D*

nmap_king_22 Highly Voted 3 months, 2 weeks ago

To assist with the increased workload for a global web-hosting company with mainly static content during an upcoming event, the most suitable solution would be:

D. CDN (Content Delivery Network).

upvoted 5 times

Question #114 *Topic 1*

A company wants to implement business continuity, and the cloud solution architect needs to design the correct solution. Which of the following will provide the data to measure business continuity? (Choose two.)

A. A service-level agreement

B. Automation scripts

C. Playbooks

D. A network diagram

E. A backup and restore

F. A recovery time objective

Suggested Answer: *AF*

Community vote distribution

AF (100%)

Not_That_Guy Highly Voted 1 year, 8 months ago

Selected Answer: AF

Only A and F provide the target parameters needed to measure success/failure.

upvoted 8 times

AustinKelleyNet Highly Voted 1 year, 5 months ago

Selected Answer: AF

If this isn't obvious to you, then I recommend that you reschedule your test and study some more.

upvoted 6 times

FrancisDrake 5 months, 4 weeks ago

Why don't you explain why the answer is AF?

upvoted 4 times

sham_ba_lam Most Recent 1 year, 4 months ago

Both the SLA and RTO are critical components of measuring business continuity and ensuring that the solution is effectively providing the desired level of service and performance. Other components such as automation scripts, playbooks, network diagrams, and back-up and restore processes are also important in ensuring business continuity, but they are not used specifically to measure the effectiveness of the solution.

upvoted 5 times

ironman_86 1 year, 9 months ago

i think it shoud be A & F

upvoted 3 times

jiminycriminal 1 year, 9 months ago

Yeah the only things that "measure" anything is the SLA and RTO.

upvoted 5 times

Question #115 *Topic 1*

A systems administrator is about to deploy a new VM to a cloud environment. Which of the following will the administrator MOST likely use to select an address for the VM?

A. CDN

B. DNS

C. NTP

D. IPAM

Suggested Answer: *D*

Community vote distribution

D (100%)

eacunha 6 months ago

Selected Answer: D

D. IPAM

O IPAM (Gerenciamento de Endereços IP, na sigla em inglês) é uma ferramenta usada para planejar, administrar e rastrear endereços IP em uma rede. Ao implantar uma nova VM, o administrador de sistemas provavelmente usará o IPAM para selecionar um endereço IP disponível para a VM na rede.

upvoted 1 times

AustinKelleyNet 1 year, 5 months ago

Selected Answer: D

IPAM = IP Address Management.

upvoted 2 times

Question #116 *Topic 1*

Which of the following is relevant to capacity planning in a SaaS environment?

A. Licensing

B. A hypervisor

C. Clustering

D. Scalability

Suggested Answer: *A*

Community vote distribution

A (57%) D (43%)

⊟ **Starz146** Highly Voted 👍 1 year, 9 months ago

Thoughts on licensing here?

upvoted 10 times

⊟ **VVV4WIN** Highly Voted 👍 1 year, 1 month ago

I would go with licensing here. Use O365/M365 as an example and compare it with BeauChateau's response, how will you be able to contact and discuss Microsoft's scalability in anyway shape or form? One will just assume they have scalability, just like with most other SaaS providers. You also have no control over this whatsoever. All one can control (using O365/M365 as the example again) is your licensing and comparing the staff intake with the current licenses to see if the company will be able to cover any new staff members and for how long.

A very rough example, but all that makes sense for me here is licensing.

upvoted 7 times

⊟ **FrancisDrake** 12 months ago

A good point.

upvoted 2 times

⊟ **braveheart22** Most Recent 🕘 2 months, 2 weeks ago

D is the correct option hands down.

Scalability is more directly relevant to capacity planning in a SaaS environment. Scalability directly impacts the ability to accommodate growth in users and usage patterns.

Effective capacity planning focuses on ensuring that the infrastructure can handle increasing loads, adapt to fluctuating demands, and maintain performance. This means designing the system to scale efficiently—whether that's adding more resources, optimizing performance, or managing load balancing.

Licensing, on the other hand, is typically a separate concern that deals with the legal and financial aspects of software usage rather than the technical capacity to support user demand. Licensing considerations are also important for compliance and revenue management.

From the above argument, it is more than evident that Scalability is relevant for Papacity Planning, and not Licensing.

upvoted 1 times

⊟ **nedeajob12** 4 months, 2 weeks ago

Selected Answer: A

think of office 365. its a saas application which costs go up per license

upvoted 1 times

⊟ **PatrickH** 9 months, 3 weeks ago

Selected Answer: A

Tough question because its not clear enough. Can definatly make an argument for BOTH A and D. Having started with D as the correct answer Im coming around to A, Because of the choices between Per User License, Socket Based, Core Based etc...

upvoted 1 times

⊟ **FrancisDrake** 10 months, 2 weeks ago

Selected Answer: A

Licensing. I will not repeat VVV4WIN's argument for him, I will only give him credit and direct you to it.

upvoted 1 times

⊟ **VVV4WIN** 1 year, 1 month ago

Selected Answer: A

Answer A - licensing, please see my explanation below

upvoted 1 times

BeauChateau 1 year, 8 months ago

Selected Answer: D

D. Scalability is relevant to capacity planning in a SaaS (Software-as-a-Service) environment. Capacity planning is the process of determining the IT infrastructure resources required to meet future demand for a service or application. In a SaaS environment, the provider must ensure that the infrastructure can scale to meet the demands of all customers who use the service simultaneously. This requires careful planning of resources such as CPU, memory, storage, and network bandwidth to ensure that they are adequate for current and future needs.

upvoted 3 times

FrancisDrake 11 months ago

You make a good argument. The problem with this question is that it is not made clear whether it is referring to the provider of the SaaS or consumer of the SaaS.

upvoted 3 times

Question #117 *Topic 1*

An organization is hosting a DNS domain with private and public IP ranges. Which of the following should be implemented to achieve ease of management?

A. Network peering

B. A CDN solution

C. A SDN solution

D. An IPAM solution

Suggested Answer: *D*

Community vote distribution

D (100%)

BeauChateau 2 months, 2 weeks ago

Selected Answer: D

To achieve ease of management for a DNS domain with private and public IP ranges, an IP Address Management (IPAM) solution should be implemented. IPAM allows the organization to manage IP addresses and related data centrally, and automate IP address assignment and management tasks, reducing the complexity of managing IP addresses.

upvoted 4 times

Question #118 *Topic 1*

A systems administrator is performing upgrades to all the hypervisors in the environment. Which of the following components of the hypervisors should be upgraded? (Choose two.)

A. The fabric interconnects

B. The virtual appliances

C. The firmware

D. The virtual machines

E. The baselines

F. The operating system

Suggested Answer: *CF*

Community vote distribution

CF (75%) CE (25%)

⊟ **Jay987654** 5 months ago

Selected Answer: CF

The two components of the hypervisors that should be upgraded are C. The firmware and F. The operating system.

The firmware and the operating system are integral parts of a hypervisor, and keeping them up-to-date is crucial for security, performance, and compatibility reasons.

upvoted 1 times

⊟ **utied** 6 months, 2 weeks ago

Selected Answer: CE

I choose C,E.
Reason: Hypervisors Type 1 does not have OS. It manages the VMs and allocates hardware directly from the kernel.
Once you update the hypervisor, you need to create a new baseline depending on how the updates affected the hypervisor.

upvoted 1 times

⊟ **Pongsathorn** 9 months, 1 week ago

Selected Answer: CF

When performing upgrades to hypervisors in an environment, the following components should typically be upgraded:

C. The firmware: Hypervisor hosts often have firmware, including BIOS or UEFI firmware, that may need updates to ensure compatibility with new hardware, security patches, or improvements in performance and stability.

F. The operating system: Hypervisor hosts run on an underlying operating system. Updating the OS ensures that it's running the latest security patches, drivers, and hypervisor-specific software to support virtualization.

upvoted 2 times

⊟ **Pongsathorn** 9 months, 1 week ago

The other options mentioned may also be relevant but are generally not considered as components of the hypervisors themselves:

- **A. The fabric interconnects:** This is more relevant in Cisco Unified Computing System (UCS) environments and deals with network connectivity. It's not directly related to hypervisor upgrades.

- **B. The virtual appliances:** Virtual appliances are software applications packaged to run on virtualized infrastructure. Upgrading them is important but not the same as upgrading the hypervisor itself.

upvoted 1 times

⊟ **Pongsathorn** 9 months, 1 week ago

- **D. The virtual machines:** VMs themselves don't need to be upgraded in the same way that hypervisors do. VMs can run various operating systems and software, and upgrades to their software should be managed independently.

- **E. The baselines:** Baselines are typically used in configuration management or compliance management and may not directly relate to hypervisor upgrades. However, ensuring that your environment complies with security baselines is important for overall security.

So, while these other components may require attention during maintenance, they are not the core components of the hypervisor software itself.

upvoted 1 times

Question #119 *Topic 1*

A SAN that holds VM files is running out of storage space. Which of the following will BEST increase the amount of effective storage on the SAN?

A. Enable encryption.

B. Increase IOPS.

C. Convert the SAN from RAID 50 to RAID 60.

D. Configure deduplication.

Suggested Answer: *D*

Community vote distribution

D (80%) C (20%)

8c4769c 2 months, 2 weeks ago

Selected Answer: D

Deduplication

upvoted 1 times

FrancisDrake 4 months, 2 weeks ago

I select C if I want to destroy all of my data.

upvoted 2 times

iliketacos 6 months ago

Selected Answer: D

converting from RAID 50 to RAID 60 would decrease the amount of usable storage space, given you have the same number of disks. Dedupe would free up space

upvoted 1 times

nmap_king_22 9 months, 2 weeks ago

Selected Answer: C

The option that would BEST increase the amount of effective storage on the SAN in the given scenario is:

C. Convert the SAN from RAID 50 to RAID 60.

RAID (Redundant Array of Independent Disks) configurations can significantly impact storage capacity and performance. In this case, converting the SAN from RAID 50 to RAID 60 would likely provide the greatest increase in effective storage while maintaining data redundancy and performance.

upvoted 1 times

Tomtom11 11 months ago

Selected Answer: D

Deduplication works the same way, except instead of replacing information only in a
single file, the storage system looks for redundant information across every bit of data it
stores. In the case of a SAN performing deduplication, it would seek out and deduplicate
redundant information stored in every LUN. This allows for tremendous efficiency when you
have hundreds of virtual disks all containing the same OS. Assuming the OS alone consumes
8 GB, deduplication could potentially free up terabytes of space!

upvoted 2 times

Question #120 *Topic 1*

Which of the following actions should a systems administrator perform during the containment phase of a security incident in the cloud?

A. Deploy a new instance using a known-good base image.

B. Configure a firewall rule to block the traffic on the affected instance.

C. Perform a forensic analysis of the affected instance.

D. Conduct a tabletop exercise involving developers and systems administrators.

Suggested Answer: *B*

Community vote distribution

B (100%)

8c4769c 2 months, 2 weeks ago

Selected Answer: B

B. seems like the only answer that fits for containing an incident.

upvoted 4 times

Question #121 *Topic 1*

A systems administrator has migrated a web application to the cloud with a synchronous uplink speed of 100Mbps. After the migration, the administrator receives reports of slow connectivity to the web application. The administrator logs into the firewall and notices the WAN port is transmitting at a constant 12.5MBps. Which of the following BEST explains the reason for the issue?

A. Misconfigured subnetting

B. Insufficient compute

C. Firewall issues

D. Not enough upload bandwidth

Suggested Answer: *D*

Community vote distribution

D (100%)

jiminycriminal Highly Voted 1 year, 9 months ago

I think the given answer is correct. 100Mbps = 12.5MBps. If we're transmitting a constant 12.5MBps out, people are fighting over bandwidth.

upvoted 7 times

Tomtom11 Most Recent 11 months ago

1 Mbps = 0.125 MB/s

upvoted 1 times

AustinKelleyNet 1 year, 5 months ago

Selected Answer: D

This should be obvious

upvoted 2 times

FrancisDrake 5 months, 4 weeks ago

Why is it obvious? Can you explain your thinking?

upvoted 3 times

Question #122 *Topic 1*

A user reports a poor-quality remote VDI session. Which of the following should the help desk technician do FIRST to troubleshoot the issue?

A. Check the FAQ section of the vendor's documentation.

B. Ask the user if the client device or access location has changed.

C. Reboot the user's virtual desktop.

D. Request permission to log in to the device remotely.

Suggested Answer: *B*

Community vote distribution

B (100%)

Chiaretta 5 months, 2 weeks ago

Selected Answer: B

B is the first step of CompTIA model.

upvoted 3 times

FrancisDrake 5 months, 4 weeks ago

In real life you do the reboot. For the test you ask if the device or access location has changed.

upvoted 1 times

BeauChateau 1 year, 2 months ago

Selected Answer: B

B. Ask the user if the client device or access location has changed.

Before starting any troubleshooting process, it is always essential to gather information about the problem. Asking the user if anything has changed with the client device or access location is a quick way to identify if the issue is related to the user's connection or hardware. This step may eliminate unnecessary troubleshooting and provide insight into the root cause of the problem.

upvoted 4 times

Alizadeh 1 year, 2 months ago

Selected Answer: B

B. Ask the user if the client device or access location has changed.

upvoted 3 times

jiminycriminal 1 year, 9 months ago

Wow this is the epitome of tech. "Have you tried turning it off and on again?" I'm not saying that's the correct solution, but I don't think I've seen that answer before in any test lol.

Anyway, I think it's B. Help desk should ask relevant questions. A new thin client or poor quality connection from a different location are relevant imo.

upvoted 1 times

i_bird 1 year, 9 months ago

First Step of TS model

Question the user and identify user changes to computer and perform backups before making changes:
A system if creating a problem, it is needed to be solved immediately. But before solving the problem, the very problem is needed to be identified. Once sorted out, and ultimately that is to be fixed. But before the problem is fixed, the system should be brought to action. For that, one is needed to have sor run. In this article, there will be some tips regarding the temporary solution of the problems in system.

https://www.examcollection.com/certification-training/a-plus-explaining-troubleshooting-theory.html#:~:text=Establish%20a%20theory%20of%20probable%20cause%20(question%20the%20obvious)&text=The%20questions%20are%20regarding'

upvoted 1 times

ironman_86 1 year, 9 months ago

shouldn't the help desk check and confirm first before doing the reboot?
upvoted 1 times

⊟ **LeDarius3762** 1 year, 4 months ago

yes and according to the provided answers that's the closest to B
upvoted 1 times

Question #123 *Topic 1*

A systems administrator is examining a managed hosting agreement and wants to determine how much data would be lost if a server had to be restored from backups. To which of the following metrics should the administrator refer?

A. RTO

B. MTBF

C. RPO

D. MTTR

Suggested Answer: *C*

Community vote distribution

C (100%)

⊟ **BeauChateau** Highly Voted 👍 1 year, 2 months ago

Selected Answer: C

The metric to refer to determine how much data would be lost if a server had to be restored from backups is RPO (Recovery Point Objective). RPO is the maximum amount of data loss that is acceptable, measured in time. It indicates the point in time to which data must be restored after an outage, and therefore, how frequently backups should be taken to ensure that data loss is kept within acceptable limits.

upvoted 5 times

⊟ **8c4769c** Most Recent 🕓 2 months, 2 weeks ago

Selected Answer: C

RPO is correct.

upvoted 2 times

Question #124 *Topic 1*

A systems administrator for an e-commerce company will be migrating the company's main website to a cloud provider. The principal requirement is that the website must be highly available. Which of the following will BEST address this requirement?

A. Vertical scaling

B. A server cluster

C. Redundant switches

D. A next-generation firewall

Suggested Answer: *B*

Community vote distribution

B (100%)

nmap_king_22 3 months, 1 week ago

Selected Answer: B

To achieve high availability for an e-commerce website when migrating to a cloud provider, the BEST option among the provided choices is:

B. A server cluster

A server cluster involves distributing the website's workload across multiple servers or instances. This setup ensures that if one server fails, another can take over, minimizing downtime and maintaining high availability. Server clusters are designed to handle traffic spikes and provide fault tolerance, which is crucial for e-commerce websites to ensure continuous operation and prevent revenue loss during outages.

upvoted 3 times

Question #125 *Topic 1*

A development team recently completed testing changes to a company's web-based CMS in the sandbox environment. The cloud administrator deployed these
CMS application changes to the staging environment as part of the next phase in the release life cycle. The deployment was successful, but after deploying the
CMS application, the web page displays an error message stating the application is unavailable. After reviewing the application logs, the administrator sees an error message that the CMS is unable to connect to the database. Which of the following is the BEST action for the cloud administrator to perform to resolve the issue?

A. Modify the deployment script to delete and recreate the database whenever the CMS application is deployed.

B. Modify the ACL to allow the staging environment to access the database in the sandbox environment.

C. Modify the CMS application deployment to use the previous version and redeploy the application.

D. Modify the configuration settings of the CMS application to connect to the database in the current environment.

Suggested Answer: *D*

Community vote distribution

D (100%)

Jay987654 5 months ago

Selected Answer: D

D. Modify the configuration settings of the CMS application to connect to the database in the current environment.

The error message indicates that the CMS application is unable to connect to the database. This suggests that the application's configuration settings may be pointing to the wrong database environment (for example, it might still be pointing to the sandbox environment instead of the staging environment). By updating the configuration settings to point to the correct database in the current environment, the administrator can resolve the connection issue.

upvoted 4 times

Question #126 *Topic 1*

A cloud administrator is responsible for managing a cloud-based content management solution. According to the security policy, any data that is hosted in the cloud must be protected against data exfiltration. Which of the following solutions should the administrator implement?

A. HIDS

B. FIM

C. DLP

D. WAF

Suggested Answer: *C*

Community vote distribution

C (100%)

Tomtom11 5 months ago

Data loss prevention is a security solution that identifies and helps prevent unsafe or inappropriate sharing, transfer, or use of sensitive data. It can help your organization monitor and protect sensitive information across on-premises systems, cloud-based locations, and endpoint devices.
https://www.microsoft.com/en-ie/security/business/security-101/what-is-data-loss-prevention-dlp#:~:text=Data%20loss%20prevention%20is%20a,based%20locations%2C%20and%20endpoint%20devices.

upvoted 1 times

BeauChateau 8 months, 3 weeks ago

Selected Answer: C

To protect against data exfiltration, a cloud administrator should implement a Data Loss Prevention (DLP) solution. DLP solutions can identify, monitor, and protect sensitive data by enforcing policies that prevent unauthorized access, sharing, or copying of data. By monitoring outbound traffic, DLP solutions can also detect and prevent data exfiltration attempts, ensuring that data remains secure while hosted in the cloud. HIDS (Host-based Intrusion Detection System) can detect unauthorized access or suspicious activity on the server. FIM (File Integrity Monitoring) can detect changes to files or configurations, helping to identify any unauthorized modifications. WAF (Web Application Firewall) can protect web applications against common attacks such as SQL injection, cross-site scripting, and cross-site request forgery. However, WAF is not specifically designed to protect against data exfiltration.

upvoted 4 times

bagsik89 10 months, 1 week ago

Selected Answer: C

The answer is C. A WAF will not prevent a data exfil.

upvoted 2 times

AustinKelleyNet 11 months, 1 week ago

Selected Answer: C

DLP = Data Loss Prevention

upvoted 2 times

WeLikeSpamHere 11 months, 3 weeks ago

Selected Answer: C

The answer is C. DLP.

upvoted 4 times

PM_PM 1 year, 1 month ago

Selected Answer: C

DLP sounds correct to me

upvoted 3 times

Not_That_Guy 1 year, 2 months ago

Selected Answer: C

Data exfiltration = Data Loss Prevention - C is definitely the best answer.

upvoted 4 times

ryanzou 1 year, 3 months ago

Answer is C

upvoted 2 times

⊟ **ironman_86** 1 year, 3 months ago

i think the answer is C

upvoted 3 times

Question #127 *Topic 1*

Lateral-moving malware has infected the server infrastructure. Which of the following network changes would MOST effectively prevent lateral movement in the future?

A. Implement DNSSEC in all DNS servers.

B. Segment the physical network using a VLAN.

C. Implement microsegmentation on the network.

D. Implement 802.1X in the network infrastructure.

Suggested Answer: *C*

Community vote distribution

C (69%) B (31%)

braveheart22 2 months, 3 weeks ago

Selected Answer: B

VLAN is the correct answer. Preparing for Sec+ and CySA+ made this clear.

upvoted 1 times

Pongsathorn 1 year, 3 months ago

Selected Answer: C

C. Implement microsegmentation on the network.

Microsegmentation is a security technique that divides a network into smaller, isolated segments or microsegments. Each microsegment can have its own security policies and controls. This is an effective approach to prevent lateral movement by limiting communication between different parts of the network. If lateral-moving malware infects one segment, it won't be able to easily propagate to other segments because communication is restricted.

upvoted 4 times

Pongsathorn 1 year, 3 months ago

While the other options (A, B, and D) have their own benefits and security implications, they may not be as effective as microsegmentation in preventing lateral movement. DNSSEC (option A) enhances DNS security but doesn't directly prevent lateral movement. Segmenting the network using VLANs (option B) can help, but it may not provide the same level of granular control and isolation as microsegmentation. Implementing 802.1X (option D) is important for network access control but doesn't directly address lateral movement within the network.

upvoted 2 times

SecPlus2022 1 year, 6 months ago

Selected Answer: C

"Another characteristic of APTs is that they move laterally by exploiting open ports and gaps in firewall rules. This lateral movement can be contained by micro-segmenting the network and applying intent-based security policies". Source: https://colortokens.com/blog/advanced-persistent-threats-apt/

upvoted 2 times

bagsik89 1 year, 10 months ago

The answer is C. The keyword is "most effective". B is a useful control but C is more effective. Microsegmentation is a control for (APT)Advanced Persistent Threats.

upvoted 2 times

davidsvida 1 year, 10 months ago

Network segmentation breaks the network into zones that typically consist of multiple devices and the applications that they host. Micro-segmentation takes this a step further, placing each device or even each application within its own segment.

upvoted 1 times

beamage 1 year, 11 months ago

Selected Answer: C

Sorry Cloud Micro-segmentation software uses network virtualization technology to create increasingly granular secure zones in data centers and cloud deployments.

https://www.vmware.com/topics/glossary/content/micro-segmentation.html
upvoted 3 times

⊟ **Daymeyon** 1 year, 10 months ago
great link... from that same page:

Micro-segmentation helps in networking by creating “demilitarized zones” for security within one data center and across multiple data centers. By tying fine-grained security policies to individual workloads, micro-segmentation software limits an attacker’s ability to move laterally through a data center, even after infiltrating the perimeter defenses.
upvoted 1 times

⊟ **beamage** 1 year, 11 months ago
Selected Answer: B
Micro segmentation means a switch, every switchport is a segment that's not right. VLANS would separate the servers
upvoted 1 times

⊟ **JohnMangley** 1 year, 12 months ago
C sounds correct based on some read ups
https://www.paloaltonetworks.com/cyberpedia/what-is-microsegmentation
upvoted 1 times

⊟ **Zettke** 2 years, 1 month ago
Selected Answer: B
I think B is the answer here
upvoted 2 times

Question #128 *Topic 1*

A systems administrator is creating a playbook to run tasks against a server on a set schedule. Which of the following authentication techniques should the systems administrator use within the playbook?

A. Use the server's root credentials.

B. Hard-code the password within the playbook.

C. Create a service account on the server.

D. Use the administrator's SSO credentials.

Correct Answer: *C*

Currently there are no comments in this discussion, be the first to comment!

Question #129 *Topic 1*

Which of the following should be considered for capacity planning?

A. Requirements, licensing, and trend analysis

B. Laws and regulations

C. Regions, clusters, and containers

D. Hypervisors and scalability

Suggested Answer: *A*

Community vote distribution

A (100%)

nmap_king_22 3 months, 1 week ago

Selected Answer: A

A. Requirements, licensing, and trend analysis

Capacity planning involves ensuring that an organization's IT infrastructure can meet its current and future performance and resource requirements. To achieve effective capacity planning, the following factors should be considered:

upvoted 4 times

Question #130 *Topic 1*

A cloud administrator is setting up a new coworker for API access to a public cloud environment. The administrator creates a new user and gives the coworker access to a collection of automation scripts. When the coworker attempts to use a deployment script, a 403 error is returned. Which of the following is the MOST likely cause of the error?

A. Connectivity to the public cloud is down.

B. User permissions are not correct.

C. The script has a configuration error.

D. Oversubscription limits have been exceeded.

Suggested Answer: *B*

Brianhealey136 Highly Voted 👍 6 months ago

HTTP 403 is an HTTP status code meaning access to the requested resource is forbidden. The server understood the request, but will not fulfill it.

upvoted 6 times

Question #131 *Topic 1*

A systems administrator is troubleshooting a performance issue with a virtual database server. The administrator has identified the issue as being disk related and believes the cause is a lack of IOPS on the existing spinning disk storage. Which of the following should the administrator do NEXT to resolve this issue?

A. Upgrade the virtual database server.

B. Move the virtual machine to flash storage and test again.

C. Check if other machines on the same storage are having issues.

D. Document the findings and place them in a shared knowledge base.

Suggested Answer: *B*

Community vote distribution

C (50%) B (50%)

Dejo1990 1 month, 1 week ago

Selected Answer: C

sys admins BELIEVES. He must test it to be sure

upvoted 1 times

Chiaretta 11 months, 3 weeks ago

Selected Answer: C

The key word is NEXT.

upvoted 2 times

FrancisDrake 12 months ago

This looks a troubleshooting question in disguise. C is going backward in the troubleshooting scheme. B would be testing the theory that the admin has formulated. I would select B.

upvoted 3 times

buckthesystem 1 year, 1 month ago

Selected Answer: C

Classic CompTIA poorly written question. Verifying that it's not a config issue first would potentially avoid unnecessary expenses involved with swapping over to SSDs. But who knows what they're going for here. Both B & C could be argued.

upvoted 3 times

BeauChateau 1 year, 8 months ago

Selected Answer: B

B. Move the virtual machine to flash storage and test again.

Explanation:

If a performance issue has been identified as disk related and the cause is determined to be a lack of IOPS on the existing spinning disk storage, the next step to resolve the issue should be to move the virtual machine to flash storage and test again. Flash storage is known to have a much higher IOPS capacity than spinning disk storage, which should help to resolve the performance issue. Upgrading the virtual database server or checking if other machines on the same storage are having issues may not necessarily resolve the disk-related performance issue. Documenting the findings and placing them in a shared knowledge base is important, but it should not be the next step in resolving the issue.

upvoted 1 times

bagsik89 1 year, 10 months ago

Selected Answer: C

C.

He was diagnosing an issue of lack of IOPS from the server. He needs to check the other VMs running on that disk to rule out that it's not a faulty disk.

upvoted 2 times

LeDarius3762 1 year, 10 months ago

Selected Answer: B

Answer C would help to see if it's a bad config/bad disk issue, but it won't solve the issue right away, the answer that would solve the issue if it's a IOPS problem is answer B

upvoted 1 times

AustinKelleyNet 1 year, 11 months ago

Selected Answer: B

He has established the theory. Now he needs to test it.

upvoted 3 times

beamage 1 year, 11 months ago

sorry could be a bad config on the one device which is why you should check other machines

upvoted 2 times

beamage 1 year, 11 months ago

Selected Answer: C

Could be a bad disk which is why you should check other machines

upvoted 2 times

TheGinjaNinja 1 year, 11 months ago

Selected Answer: B

B. Move the virtual machine to flash storage and test again.

The administrator has identified that the issue is disk-related and that the cause is a lack of IOPS on the existing spinning disk storage. Moving the virtual machine to flash storage would improve the IOPS and would help to resolve the performance issue.
The other options would not directly resolve the issue.
It is always good to document the findings and place them in a shared knowledge base, but this should be done after resolving the issue.

upvoted 4 times

Question #132 *Topic 1*

In an existing IaaS instance, it is required to deploy a single application that has different versions. Which of the following should be recommended to meet this requirement?

A. Deploy using containers.

B. Install a Type 2 hypervisor.

C. Enable SR-IOV on the host.

D. Create snapshots.

Suggested Answer: *A*

Community vote distribution

A (100%)

nmap_king_22 3 months, 1 week ago

Selected Answer: A

A. Deploy using containers.

To deploy a single application with different versions, using containers is the recommended approach

upvoted 2 times

Question #133 *Topic 1*

An organization is using multiple SaaS-based business applications, and the systems administrator is unable to monitor and control the use of these subscriptions.

The administrator needs to implement a solution that will help the organization apply security policies and monitor each individual SaaS subscription. Which of the following should be deployed to achieve these requirements?

A. DLP

B. CASB

C. IPS

D. HIDS

Suggested Answer: *B*

Community vote distribution

B (100%)

⊟ **BeauChateau** 2 months, 2 weeks ago

Selected Answer: B

To monitor and control the use of SaaS subscriptions, the best solution to deploy is a Cloud Access Security Broker (CASB). A CASB provides visibility and control over data that resides in cloud applications, including SaaS-based applications. It can help an organization apply security policies and monitor each individual SaaS subscription by providing features such as identity and access management, data loss prevention, threat protection, and compliance reporting. Therefore, option B is the correct answer.

upvoted 4 times

Question #134 *Topic 1*

A cloud administrator needs to reduce the cost of cloud services by using the company's off-peak period. Which of the following would be the BEST way to achieve this with minimal effort?

A. Create a separate subscription.

B. Create tags.

C. Create an auto-shutdown group.

D. Create an auto-scaling group.

Suggested Answer: *D*

Community vote distribution

D (60%) C (40%)

PatrickH 3 months, 2 weeks ago

Selected Answer: C

C seems correct:
Azure has an integrated feature that allows you to schedule the shutdown of your virtual machine (VM) for a specific time and time zone. To access this feature, use the path below for an Azure virtual machine resource: Operations -> Auto-shutdown.

You can also specify a Webhook URL and an email address to receive alert notifications before auto-shutdown.

upvoted 1 times

Chiaretta 5 months, 2 weeks ago

Selected Answer: D

Auto-scaling group is the right aswer, the auto shutdown is referred about domain computer left turned on by users.

upvoted 2 times

utied 6 months, 2 weeks ago

Selected Answer: C

I think C is the best answer, but Cloud+ exam objectives do not mention auto-shutdown groups. They only talk about auto-scaling groups. So who knows. Either one could be correct.

upvoted 2 times

userguy890 7 months, 3 weeks ago

Selected Answer: C

auto shutdown saves the most money since nothing is running during off peak times

upvoted 1 times

dcdc1000 10 months ago

Selected Answer: D

Create an auto-scaling group. A reduced workload (scaling down) during off-peak hours results in less heat generated and less electricity consumed. Which results in less cost. Not to mention, most CSP charge based on actual consumption.

upvoted 4 times

Question #135 *Topic 1*

A media company has made the decision to migrate a physical, internal file server to the cloud and use a web-based interface to access and manage the files.
The users must be able to use their current corporate logins. Which of the following is the MOST efficient way to achieve this goal?

A. Deploy a VM in a cloud, attach storage, and copy the files across.

B. Use a SaaS service with a directory service federation.

C. Deploy a fileshare in a public cloud and copy the files across.

D. Copy the files to the object storage location in a public cloud.

Correct Answer: *B*

Currently there are no comments in this discussion, be the first to comment!

Question #136 *Topic 1*

A systems administrator is configuring a storage array. Which of the following should the administrator configure to set up mirroring on this array?

A. RAID 0

B. RAID 1

C. RAID 5

D. RAID 6

Suggested Answer: *B*

Community vote distribution

B (100%)

⊟ **nmap_king_22** 3 months, 1 week ago

Selected Answer: B

B. RAID 1

To set up mirroring on a storage array, you should configure RAID 1. RAID 1, also known as mirroring, duplicates data across multiple drives in the array. In a RAID 1 setup, every write operation is mirrored to two or more drives simultaneously. This redundancy ensures that if one drive fails, the data is still accessible from the mirrored drive(s), providing data protection and fault tolerance.

upvoted 3 times

⊟ **BeauChateau** 8 months, 3 weeks ago

Selected Answer: B

To set up mirroring on a storage array, a systems administrator should configure RAID 1. RAID 1, also known as disk mirroring, uses two identical disks to store the same data. If one disk fails, the other disk can be used to recover the data. RAID 0 is not a mirroring RAID level but a striping RAID level that spreads data across multiple disks to enhance performance. RAID 5 and RAID 6 are parity-based RAID levels that provide fault tolerance and can recover data in the event of a single disk failure. However, they are not designed specifically for mirroring.

upvoted 4 times

Question #137 *Topic 1*

A systems administrator is working in a globally distributed cloud environment. After a file server VM was moved to another region, all users began reporting slowness when saving files. Which of the following is the FIRST thing the administrator should check while troubleshooting?

A. Network latency

B. Network connectivity

C. Network switch

D. Network peering

Correct Answer: *A*

Currently there are no comments in this discussion, be the first to comment!

Question #138 *Topic 1*

A systems administrator adds servers to a round-robin, load-balanced pool, and then starts receiving reports of the website being intermittently unavailable. Which of the following is the MOST likely cause of the issue?

A. The network is being saturated.

B. The load balancer is being overwhelmed.

C. New web nodes are not operational.

D. The API version is incompatible.

E. There are time synchronization issues.

Suggested Answer: *C*

Community vote distribution

C (71%) B (29%)

beamage Highly Voted 1 year, 5 months ago

Selected Answer: C

Obviously there is a node or nodes that are not on.....

upvoted 6 times

ROCompTIA Most Recent 1 year ago

Selected Answer: C

When new servers are added to a load balancer pool, there is a possibility they are not fully functioning or properly configured. If the load balancer is routing traffic to these non-operational servers, it would result in the website becoming unavailable.

upvoted 4 times

Sweety_Certified7 4 months, 1 week ago

If the new servers were not operational, the issue would likely result in consistent unavailability rather than intermittent unavailability. So answer is B.

upvoted 1 times

TheGinjaNinja 1 year, 5 months ago

Selected Answer: B

When a load balancer is overwhelmed, it can't distribute incoming traffic effectively, which can lead to intermittent unavailability of the website. This is the most likely cause of the issue as the administrator has recently added servers to the load-balanced pool, and the website is being reported as being intermittently unavailable.

Other options like network saturation, new web nodes not being operational, API version incompatibility, and time synchronization issues can also cause similar symptoms, but as per the information provided, the load balancer being overwhelmed is the most likely cause.

upvoted 4 times

Daymeyon 1 year, 5 months ago

Only disagree because of the verbiage “and then starts” as opposed to “and continues”. Sounds like something is wrong with the newly added servers. Going with C here

upvoted 1 times

Question #139 *Topic 1*

A company has a cloud infrastructure service, and the cloud architect needs to set up a DR site. Which of the following should be configured in between the cloud environment and the DR site?

A. Fallback

B. Playbook

C. Zoning

D. Replication

Suggested Answer: *D*

Community vote distribution

D (100%)

Not_That_Guy Highly Voted 1 year, 2 months ago

Selected Answer: D

Without replication, you wouldn't have a viable system to failover to in the first place.

upvoted 9 times

i_bird Highly Voted 1 year, 3 months ago

How is it fallback or failback??

I should be some kinda of replication??

upvoted 7 times

ROCompTIA Most Recent 6 months, 3 weeks ago

Selected Answer: D

Replication involves copying data between the primary cloud environment and the DR site to ensure the DR site has up-to-date information in the event of a failover. This is a key part of establishing a functioning DR solution.

upvoted 2 times

bagsik89 10 months, 1 week ago

Selected Answer: D

D. Replication

upvoted 2 times

TheGinjaNinja 1 year ago

Selected Answer: D

Replication

upvoted 3 times

s6i3 1 year, 1 month ago

I think replication is a right choice

upvoted 2 times

jiminycriminal 1 year, 3 months ago

cant really fallback and be operational without configuring replication "in-between". So I'm going with replication.

upvoted 4 times

Question #140 *Topic 1*

A cloud administrator is building a new VM for a network security appliance. The security appliance installer says the CPU clock speed does not meet the requirements. Which of the following will MOST likely solve the issue?

A. Move the VM to a host with a faster CPU.

B. Add more vCPUs to the VM.

C. Enable CPU masking on the VM.

D. Enable hyperthreading on the virtual host.

Suggested Answer: *A*

Community vote distribution

A (100%)

ROCompTIA 6 months, 3 weeks ago

Selected Answer: A

If the CPU on the current host is not fast enough for the appliance, moving it to a host with a faster CPU is the most direct way to resolve the issue

upvoted 1 times

strale 1 year, 2 months ago

Can anyone explain? Why not D?

upvoted 1 times

ROCompTIA 6 months, 3 weeks ago

Enabling hyperthreading may provide more logical CPU cores for the VM to use but does not increase the base CPU clock speed.

upvoted 1 times

JVen 1 year, 1 month ago

A is the only one that affects the actual IPC of the processor being used for the VM(clock speed helps improve IPC). Adding more cores doesn't. IPC can also be referenced as single core performance, this is separate from multi-core performance which scales based on the number of vCPU/cores being used.

upvoted 1 times

Question #141 *Topic 1*

A technician is working with an American company that is using cloud services to provide video-based training for its customers. Recently, due to a surge in demand, customers in Europe are experiencing latency. Which of the following services should the technician deploy to eliminate the latency issue?

A. Auto-scaling

B. Cloud bursting

C. A content delivery network

D. A new cloud provider

Suggested Answer: *C*

Community vote distribution

C (100%)

sheilawu 5 months, 2 weeks ago

Selected Answer: C

AZ900 is C so I go for C

upvoted 1 times

BeauChateau 8 months, 3 weeks ago

Selected Answer: C

C. A content delivery network (CDN) would be the most appropriate solution to eliminate the latency issue. A CDN is a network of geographically distributed servers that cache and deliver content to users based on their geographic location, resulting in reduced latency and faster content delivery. By using a CDN, the company can store and distribute the training videos from multiple locations, closer to the end-users, reducing the distance the data must travel, and therefore, reducing latency.

upvoted 2 times

DocHacker 1 year ago

Selected Answer: C

https://www.cloudflare.com/learning/cdn/what-is-a-cdn/

"A content delivery network (CDN) refers to a geographically distributed group of servers which work together to provide fast delivery of Internet content."

upvoted 1 times

Sal 1 year, 2 months ago

How is this not auto-scaling?

upvoted 1 times

aznstylewalk 1 year ago

Auto scaling wouldn't help with latency. Auto scaling would be adding more resource, content Delivery Network allows you to store information in another area/region and have users pull from the closest node decreasing latency

upvoted 4 times

Question #142 *Topic 1*

A systems administrator has received an email from the virtualized environment's alarms indicating the memory was reaching full utilization. When logging in, the administrator notices that one out of a five-host cluster has a utilization of 500GB out of 512GB of RAM. The baseline utilization has been 300GB for that host.
Which of the following should the administrator check NEXT?

A. Storage array

B. Running applications

C. VM integrity

D. Allocated guest resources

Suggested Answer: *B*

Community vote distribution

B (71%) C (29%)

Jay987654 5 months ago

Selected Answer: B

B. Running applications

Given that the memory utilization on one host in the cluster has increased significantly from a baseline of 300GB to 500GB, the next step for the administrator should be to check the running applications (Option B). Identifying the specific applications or processes that are consuming the additional memory can provide insights into the cause of the increased utilization.

upvoted 2 times

Robenger 5 months, 1 week ago

Selected Answer: B

In this scenario, the administrator has identified a sudden increase in memory utilization on one host in a virtualized environment. The next step to investigate would typically involve checking the following:

B. Running applications

Checking the running applications on the host is crucial to identify any processes or applications that might be consuming excessive memory resources. An abnormal increase in memory usage could be attributed to a specific application or process misbehaving.

upvoted 3 times

utied 6 months, 2 weeks ago

Selected Answer: C

I maybe overthinking the problem.
Cluster: two or more servers that have the same data for failover. treated as a single entity. network ip for both is single ip address.
Node: each cluster member is called a node. all nodes will be configured the same way and share the same ip.

So five nodes, all should be identical. One is using 200GB ram more than the other four. VM integrity has been compromised.
I think the answer is C.

upvoted 2 times

Question #143 *Topic 1*

Company A has acquired Company B and is in the process of integrating their cloud resources. Company B needs access to Company A's cloud resources while retaining its 1AM solution. Which of the following should be implemented?

A. Multifactor authentication

B. Single sign-on

C. Identity federation

D. Directory service

Suggested Answer: *C*

Community vote distribution

C (100%)

Pongsathorn 3 months, 2 weeks ago

Selected Answer: C

C. Identity federation

Identity federation allows users from one organization to access resources or services in another organization without the need to duplicate user accounts or credentials. In this scenario, Company B can retain its IAM (Identity and Access Management) solution while using identity federation to enable its users to access Company A's cloud resources. This helps maintain security and access control while allowing for the integration of resources between the two companies.

upvoted 3 times

PatrickH 7 months, 2 weeks ago

Just FYI thats a Typo. Its IAM not 1AM. Answer is correct

upvoted 4 times

Question #144 *Topic 1*

A systems administrator wants to have near-real-time information on the volume of data being exchanged between an application server and its clients on the
Internet. Which of the following should the systems administrator implement to achieve this objective?

A. A stateful firewall

B. DLP

C. DNSSEC

D. Network flows

Suggested Answer: *D*

Alizadeh Highly Voted 👍 5 months, 1 week ago

Network flows (Option D) is a technology that allows you to collect and analyze information about the flow of network traffic, including the volume of data being exchanged between the application server and its clients. This information can be used to identify patterns, trends, and anomalies in network traffic, and can be used to troubleshoot network issues or identify potential security threats. Network flow tools can provide near-real-time information on the volume of data being exchanged, making it the best option to achieve the systems administrator's objective.

upvoted 6 times

CapJackSparrow Most Recent 🕘 5 months, 2 weeks ago

What? CompTIA always playing games..

upvoted 1 times

Question #145

Topic 1

A company had a system compromise, and the engineering team resolved the issue after 12 hours. Which of the following information will MOST likely be requested by the Chief Information Officer (CIO) to understand the issue and its resolution?

A. A root cause analysis

B. Application documentation

C. Acquired evidence

D. Application logs

Correct Answer: *A*

Currently there are no comments in this discussion, be the first to comment!

Question #146 *Topic 1*

A company needs to migrate the storage system and batch jobs from the local storage system to a public cloud provider. Which of the following accounts will
MOST likely be created to run the batch processes?

A. User

B. LDAP

C. Role-based

D. Service

Suggested Answer: *D*

Tomtom11 5 months ago

Why not Answer A
User accounts are created for each and every user who needs to gain access to the cloud objects and resources. A user account is usually associated with an individual person but could be expanded to include other entities such as servers or applications that need to access cloud objects. (In data center parlance, these accounts are called service accounts.) More generally, user and service accounts are called identities. Once you create an identity,

upvoted 1 times

Francois1984 4 months, 1 week ago

i asked chatgpt......

When migrating batch jobs from a local storage system to a public cloud provider, the Service Account will MOST likely be created to run the batch processes.

A Service Account is a special type of account created for the purpose of running automated processes, scripts, or batch jobs. It's distinct from regular user accounts and is often used to ensure that the automated processes run with the appropriate permissions and access rights, without requiring direct user interaction.

upvoted 2 times

Question #147 *Topic 1*

A cloud administrator has finished setting up an application that will use RDP to connect. During testing, users experience a connection timeout error. Which of the following will MOST likely solve the issue?

A. Checking user passwords

B. Configuring QoS rules

C. Enforcing TLS authentication

D. Opening TCP port 3389

Suggested Answer: *D*

Community vote distribution

D (100%)

Pongsathorn 3 months, 2 weeks ago

Selected Answer: D

D. Opening TCP port 3389

The issue described suggests that the RDP (Remote Desktop Protocol) connection is likely being blocked due to a closed or restricted port. TCP port 3389 is the default port used for RDP connections, so opening this port in the network's firewall or security settings should resolve the connection timeout error and allow users to establish RDP connections to the application.

upvoted 3 times

AustinKelleyNet 11 months, 1 week ago

Selected Answer: D

rDp uses port 3389.

upvoted 4 times

beamage 11 months, 2 weeks ago

Selected Answer: D

No not QOS Firewall..... https://docs.microsoft.com/en-us/troubleshoot/windows-server/remote/troubleshoot-remote-desktop-disconnected-errors

upvoted 2 times

JVen 1 year, 1 month ago

Selected Answer: D

I agree with the others, D is the answer

upvoted 3 times

Not_That_Guy 1 year, 2 months ago

Selected Answer: D

QoS rules would be localized to the client's LAN/switch topology, which the cloud admin may not even have access to. D is the best answer.

upvoted 3 times

achow26 1 year, 3 months ago

D should be the answer

upvoted 4 times

ironman_86 1 year, 3 months ago

i think answer should be D

upvoted 3 times

Question #148 *Topic 1*

A storage array that is used exclusively for datastores is being decommissioned, and a new array has been installed. Now the private cloud administrator needs to migrate the data. Which of the following migration methods would be the BEST to use?

A. Conduct a V2V migration.

B. Perform a storage live migration.

C. Rsync the data between arrays.

D. Use a storage vendor migration appliance.

Suggested Answer: *D*

Community vote distribution

D (53%) A (27%) 13% 7%

BeauChateau Highly Voted 1 year, 2 months ago

Selected Answer: D

Since the storage array is used exclusively for datastores, the private cloud administrator should use a storage vendor migration appliance to migrate the data to the new array. This will ensure that the migration is completed with minimal downtime and data loss. V2V migration is used to migrate virtual machines between hypervisors. Storage live migration is used to move virtual disks between datastores. Rsync is a command-line tool used for copying and syncing files.

upvoted 5 times

FasterN8 Most Recent 4 months, 1 week ago

Selected Answer: D

This is a lifecycle process that the vendor has planned for. That application is purpose-built for this event (and new onboarding). Also, what Beau said.

upvoted 1 times

Granddude 9 months, 3 weeks ago

Selected Answer: C

on second thought, I am going with C. Because it's the fastest to IMPLEMENT than the other choices.

upvoted 1 times

SecPlus2022 1 year ago

Selected Answer: D

BeauChateau explains it best.

upvoted 1 times

concepcionz 1 year, 3 months ago

Selected Answer: D

We got A, B, C

upvoted 1 times

AniMon 1 year, 4 months ago

Rsync is the easiest way for storage migration

upvoted 1 times

TheGinjaNinja 1 year, 5 months ago

Selected Answer: A

Had this; I think it's A actually.

A V2V (Virtual Machine to Virtual Machine) migration is the best method for migrating data from one storage array to another in a private cloud environment. This method allows the administrator to move the entire virtual machine, including its configuration, virtual disk, and memory state, from one location to another while the virtual machine is still running. This ensures minimal disruption to the users and applications accessing the data and allows the administrator to test the migration before making it permanent.

upvoted 4 times

TheGinjaNinja 1 year, 6 months ago

Selected Answer: B

I think it is B for this one.

upvoted 2 times

CapJackSparrow 1 year, 5 months ago

googling "storage live migration" really ticks all the boxes, BUT, resync the data seems pretty good too. I'm REALLY not liking this Cloud+ exam, comptia plays too many games.

upvoted 3 times

Question #149 *Topic 1*

A cloud administrator checked out the deployment scripts used to deploy the sandbox environment to a public cloud provider. The administrator modified the script to add an application load balancer in front of the web-based front-end application. The administrator next used the script to recreate a new sandbox environment successfully, and the application was then using the new load balancer.

The following week, a new update was required to add more front-end servers to the sandbox environment. A second administrator made the necessary changes and checked out the deployment scripts. The second administrator then ran the script, but the application load balancer was missing from the new deployment.

Which of the following is the MOST likely reason for this issue?

A. The license limit on the number of server deployments allowed per month was exceeded.

B. The deployment script changes made by the first administrator were not checked in and committed.

C. The new server images were incompatible with the application load-balancer configuration.

D. The application load balancer exceeded the maximum number of servers it could use.

Suggested Answer: *B*

i_bird 3 months, 2 weeks ago

How is this B??

I'm thinking C

upvoted 1 times

jiminycriminal 3 months, 2 weeks ago

Incompatibility does not really make sense and will not stop the deployment of an application load balancer. Search on youtube "aws application load balancer", it's not even part of the virtual sandbox environment, you have to set it up separately and target the applications. The only thing here that makes any sense is B, the modified script was not saved. No other reason the load balancer would not be deployed.

upvoted 4 times

Question #150 *Topic 1*

A systems administrator is configuring updates on a system. Which of the following update branches should the administrator choose to ensure the system receives updates that are maintained for at least four years?

A. LTS

B. Canary

C. Beta

D. Stable

Suggested Answer: *A*

Community vote distribution

A (100%)

AustinKelleyNet 5 months ago

Selected Answer: A

LTS stands for long term support

upvoted 3 times

Question #151 *Topic 1*

A systems administrator is performing an OS upgrade on a production VM. Which of the following actions should the administrator take before the upgrade to ensure the FASTEST recovery of the system in case the upgrade fails in an unrecoverable way?

A. Submit the upgrade to the CAB.

B. Perform a full backup.

C. Take a snapshot of the system.

D. Test the upgrade in a preproduction environment.

Suggested Answer: *C*

Community vote distribution

C (100%)

Sweety_Certified7 3 months, 1 week ago

Selected Answer: C

Obviously it is C

upvoted 2 times

Question #152 *Topic 1*

Users are experiencing slow response times from an intranet website that is hosted on a cloud platform. There is a site-to-site VPN connection to the cloud provider over a link of 100Mbps. Which of the following solutions will resolve the issue the FASTEST?

A. Change the connection to point-to-site VPN.

B. Order a direct link to the provider.

C. Enable quality of service.

D. Upgrade the link to 200Mbps.

Suggested Answer: *C*

Community vote distribution

C (79%) 11% 11%

ironman_86 Highly Voted 1 year, 9 months ago

i think it's C, not B.

upvoted 5 times

Rjimbo Most Recent 1 month ago

Selected Answer: D

Upgrading the link to 200Mbps (option D) is the fastest solution as it directly addresses the issue of slow response times by providing more bandwidth for data transfer between the users and the cloud platform. This solution can be implemented relatively quickly compared to other options that involve configuration changes or ordering new network connections.

While ordering a direct link can provide a more reliable and potentially faster connection, it typically involves longer lead times for installation and configuration. This solution would not resolve the issue the fastest.

upvoted 2 times

Sweety_Certified7 3 months, 1 week ago

Selected Answer: B

Option C, enabling Quality of Service (QoS), can indeed help improve performance by prioritizing traffic and managing bandwidth effectively. However, implementing QoS typically involves configuration and testing, which may take time to implement and optimize. Additionally, while QoS can help improve the overall performance of the network by prioritizing critical traffic, it may not directly address the underlying issue of slow response times from the intranet website.

Ordering a direct link to the provider (Option B) would likely provide the fastest resolution because it bypasses the limitations and potential congestion of the site-to-site VPN connection altogether. It offers a dedicated and higher bandwidth connection between the user's network and the cloud platform, resulting in lower latency and faster response times for the intranet website. So, option B is correct.

upvoted 1 times

Robenger 5 months, 1 week ago

Selected Answer: B

Order a direct link to the provider (Option B): This option involves establishing a direct connection to the cloud provider, bypassing the VPN, and potentially providing faster and more reliable connectivity. Direct connections are often preferred for critical applications that require low-latency access.

Enable quality of service (Option C): QoS is important for optimizing network performance, but it might not be the fastest solution in this context. It addresses traffic prioritization but doesn't necessarily increase the overall bandwidth.

upvoted 1 times

Granddude 9 months, 3 weeks ago

I am not sure which way I would vote on this question. Wouldn't it depend on the root cause of the slow response time?

upvoted 1 times

BeauChateau 1 year, 2 months ago

Selected Answer: C

The fastest solution to the slow response times from the intranet website hosted on a cloud platform would be to enable quality of service (QoS). QoS allows network traffic to be prioritized to ensure that important traffic, such as the intranet website traffic, is given priority over less important traffic. This can improve response times and ensure that users can access the intranet website quickly, even when other network traffic is present. Changing the connection to point-to-site VPN, ordering a direct link to the provider, or upgrading the link to 200Mbps may also help to improve response times, but they would likely take longer to implement and may involve additional costs.

upvoted 4 times

bagsik89 1 year, 4 months ago

Selected Answer: C

C would be the quickest.

upvoted 3 times

Trebor28 1 year, 4 months ago

Selected Answer: C

it is c.

upvoted 3 times

beamage 1 year, 5 months ago

Selected Answer: C

The fact that it's a point to point there would much unneeded traffic you could filter immediately with QOS

upvoted 2 times

beamage 1 year, 5 months ago

sorry Site to Site is what I meant

upvoted 1 times

CapJackSparrow 1 year, 5 months ago

Im thinking C here..

upvoted 1 times

TheGinjaNinja 1 year, 6 months ago

Selected Answer: C

It has to be C.

upvoted 3 times

Question #153 Topic 1

A company with a worldwide presence wants to improve the user experience for its website. Which of the following can a systems administrator implement to improve download speeds and latency for the end users?

A. A CDN solution

B. An MPLS connection between datacenters

C. A DNS round robin

D. A site-to-site VPN between datacenters

Correct Answer: *A*

Currently there are no comments in this discussion, be the first to comment!

Question #154 *Topic 1*

An administrator receives a ticket indicating the accounting application is not working. Which of the following should the administrator check FIRST?

A. DHCP

B. Service-level agreement

C. DNS

D. System logs

Suggested Answer: *D*

Community vote distribution

D (86%) 14%

⊟ **i_bird** Highly Voted 👍 1 year, 9 months ago

Selected Answer: D

System Logs are used for troubleshooting, not SLAs

upvoted 6 times

⊟ **Sweety_Certified7** Most Recent ⊙ 3 months, 1 week ago

Selected Answer: D

D. System logs.

Explanation:

System logs provide valuable information about system events, errors, and warnings that could help diagnose issues with applications. By examining the system logs, the administrator can quickly identify any errors or issues related to the accounting application. This includes errors that may indicate why the application is not functioning properly, such as service crashes, configuration errors, or resource limitations.

upvoted 1 times

⊟ **mitchbitch** 7 months ago

Selected Answer: D

in my work environment an admin or engineer isnt checking an SLA, we might go to our PM to check but thats not something we are tasks in doing

upvoted 1 times

⊟ **dcdc1000** 10 months ago

Selected Answer: B

Per the CompTIA Cloud+ book I'm reading. The answer would be SLA. An example of a simple internal SLA might focus on the help disk. The SLA defines how the help desk prioritizes tickets, provides timely responses, and communicates with the rest of the organization. I think this could apply to cloud as well.

upvoted 1 times

⊟ **SecPlus2022** 1 year ago

Selected Answer: B

"where does it say anything about cloud?"...this is a Cloud+ exam, it probably goes without saying. If you're the customer and your hosted application isn't working, you're going to want to check your SLA with the CSP wouldn't you? In addition, as Landoski notes, system logs are not the same as application logs.

upvoted 2 times

⊟ **FrancisDrake** 5 months, 3 weeks ago

the scenario doesn't say customer. nor does it say hosted application.

upvoted 1 times

⊟ **Sweety_Certified7** 3 months, 1 week ago

While system logs primarily capture information about system-level events, errors, and warnings, they can still provide valuable insights into the functioning of applications running on the system. For example, system logs might record errors related to application dependencies, resource allocation issues, or system-level configurations that could impact the performance or functionality of the accounting application.

Although checking application logs would be ideal for diagnosing application-specific issues, if application logs are not available or not specified as an option, system logs would be the next best choice.

upvoted 1 times

Sweety_Certified7 3 months, 1 week ago

while service-level agreements are important for determining support expectations, they are not directly related to troubleshooting the specific issue with the accounting application.

upvoted 1 times

BeauChateau 1 year, 2 months ago

Selected Answer: D

The correct answer is D. System logs.

Before looking at the network components, an administrator should check the system logs for any related error messages that could give insight into what is causing the issue. The logs may reveal issues with the application or the system on which it is running, helping the administrator pinpoint the problem quickly and efficiently. Once the problem has been identified, the administrator can then move on to other troubleshooting steps, such as checking the network configuration or consulting the service-level agreement.

upvoted 2 times

bagsik89 1 year, 4 months ago

Selected Answer: D

D. System Logs

upvoted 2 times

Landoski 1 year, 7 months ago

System logs is not the same as application logs. This about a cloud based application (SaaS)- SLA is the answer.

upvoted 4 times

CapJackSparrow 1 year, 5 months ago

where does it say anything about cloud?

upvoted 1 times

JVen 1 year, 7 months ago

Selected Answer: D

I agree with the others, D has to be the answer

upvoted 2 times

Not_That_Guy 1 year, 8 months ago

Selected Answer: D

Nothing the SLA says is going to help fix the issue.

upvoted 4 times

Question #155 *Topic 1*

A company recently adjusted its load-balancer encryption policies to support only TLSv1.3. Soon after the change was made, several customers began reporting they could not access their website. Which of the following is the MOST likely cause of the issue?

A. The certificate is expired.

B. There is a mismatch between the key and the certificate.

C. The customers are using an unsupported OS.

D. The load balancer was misconfigured.

Suggested Answer: *C*

Community vote distribution

C (100%)

BeauChateau Highly Voted 2 months, 2 weeks ago

Selected Answer: C

The most likely cause of the issue is that the customers' browsers or operating systems do not support TLSv1.3 encryption protocol. While TLSv1.3 is a more secure protocol than earlier versions, it is not yet widely supported by all web browsers and operating systems. The load balancer was configured correctly, as it was set to support only TLSv1.3, but this caused compatibility issues with some customers. The other options listed are less likely to be the cause of the issue in this scenario.

upvoted 5 times

Question #156 *Topic 1*

A cloud architect is reviewing the design for a new cloud-based ERP solution. The solution consists of eight servers with a single network interface. The allocated

IP range is 172.16.0.0/28. One of the requirements of the solution is that it must be able to handle the potential addition of 16 new servers to the environment.

Because of the complexity of the firewall and related ACL requirements, these new servers will need to be in the same network range. Which of the following changes would allow for the potential server addition?

A. Change the IP address range to use a 10.0.0.0 address.

B. Change the server template to add network interfaces.

C. Change the subnet mask to use a 255.255.255.128 range.

D. Change the server scaling configuration to increase the maximum limit.

Suggested Answer: *C*

Community vote distribution

C (100%)

Pisces225 3 months ago

Selected Answer: C

Yes, it's overkill because the next subnet up at /27 would accommodate at 32 IPs (30 usable), but of the choices listed C does indeed work.

upvoted 2 times

Question #157 *Topic 1*

An organization is required to set a custom registry key on the guest operating system. Which of the following should the organization implement to facilitate this requirement?

A. A configuration management solution

B. A log and event monitoring solution

C. A file integrity check solution

D. An operating system ACL

Suggested Answer: *A*

Community vote distribution

A (100%)

Sweety_Certified7 3 months, 1 week ago

Selected Answer: A

answer A: A configuration management solution, such as Puppet, Chef, or Ansible, is designed to automate the configuration and management of systems. These tools allow organizations to define and enforce desired configurations across their infrastructure, including setting custom registry keys on guest operating systems. By using a configuration management solution, the organization can centrally manage and enforce the desired registry key settings, ensuring consistency and compliance across all systems.

upvoted 3 times

Pongsathorn 9 months, 2 weeks ago

Selected Answer: A

To set a custom registry key on the guest operating system, the organization should implement **a configuration management solution**. Configuration management tools, such as Puppet, Chef, Ansible, or Microsoft Group Policy (for Windows systems), are designed to automate and manage configurations across large numbers of servers and operating systems. These tools can be used to set custom registry keys, manage software installations, enforce security policies, and more.

While the other options (log and event monitoring, file integrity checks, and operating system ACLs) can be useful for different purposes, they are not specifically designed for managing and enforcing configurations on a large scale as configuration management tools are.

upvoted 2 times

Question #158 *Topic 1*

A systems administrator wants the VMs on the hypervisor to share CPU resources on the same core when feasible. Which of the following will BEST achieve this goal?

A. Configure CPU passthrough.

B. Oversubscribe CPU resources.

C. Switch from a Type 1 to a Type 2 hypervisor.

D. Increase instructions per cycle.

E. Enable simultaneous multithreading.

Suggested Answer: *E*

Community vote distribution

E (50%) B (40%) 10%

Sweety_Certified7 3 months, 1 week ago

Selected Answer: E

Correct Answer: E. Enable simultaneous multithreading.

Explanation:

Simultaneous Multithreading (SMT), also known as Hyper-Threading in Intel processors, allows multiple threads to run on each CPU core simultaneously. By enabling SMT, the hypervisor can better utilize CPU resources by scheduling multiple virtual CPUs (vCPUs) from different VMs to run on the same physical core when feasible. This enables better CPU resource sharing and improves overall CPU utilization efficiency.

upvoted 2 times

FrancisDrake 4 months ago

I don't think any of these options are the answer. Some kind of affinity rule would seem to be in order.

upvoted 3 times

Pongsathorn 9 months, 3 weeks ago

Selected Answer: E

Simultaneous Multi-Threading

Modern processors can manage more than one processing thread. Doing so may improve performance, depending on the situation. Intel CPUs support the use of a form of simultaneous multi-threading (SMT) named Hyper-Threading.

upvoted 1 times

dcdc1000 10 months ago

Selected Answer: B

Answer is Oversubscription. The term refers to the practice of allocating more resources to the VMs than the physical server actually has. Admins might do this based on the anticipated workload for a given set of VMs. i.e. not all VM's are expected to consume their fully allocated resources.

upvoted 1 times

Tomtom11 10 months, 2 weeks ago

Selected Answer: E

Simultaneous Multithreading (SMT) allows multiple execution threads to be executed on a single physical CPU core. The technology is known by a number of different names, such as Hyper-Threading, but operate along similar principles

upvoted 1 times

BeauChateau 1 year, 2 months ago

Selected Answer: E

The option that would best achieve the goal of allowing VMs on the hypervisor to share CPU resources on the same core when feasible is: E. Enable simultaneous multithreading.

Simultaneous multithreading (SMT), also known as hyper-threading, is a technology that allows multiple threads to run concurrently on a single core, providing better CPU utilization and performance. By enabling SMT on the hypervisor, the VMs will be able to share CPU resources more efficiently and effectively, leading to better overall performance.

upvoted 1 times

CapJackSparrow 1 year, 5 months ago

I have to go with D on this one. Oversubcription would be more for maximization of hardware utilization, and I don't get that from this question.. Again, CompTIA playing games.

upvoted 1 times

TheGinjaNinja 1 year, 5 months ago

Not D bro.

upvoted 2 times

TheGinjaNinja 1 year, 6 months ago

Selected Answer: B

This is not recommended in the field (for various reasons), but technically- this would be the solution.

upvoted 3 times

LeDarius3762 1 year, 4 months ago

I agree, answer B

upvoted 1 times

tonytonyyyyy 1 year, 8 months ago

Selected Answer: D

The term multithreading is ambiguous, because not only can multiple threads be executed simultaneously on one CPU core, but also multiple tasks (with different page tables, different task state segments, different protection rings, different I/O permissions, etc.). Although running on the same core, they are completely separated from each other. Multithreading is similar in concept to preemptive multitasking but is implemented at the thread level of execution in modern superscalar processors. I think the answer is correct. D.

upvoted 1 times

TheGinjaNinja 1 year, 6 months ago

How would you "increase" IPC? This is set by the CPU architecture; you would have to replace the CPU

upvoted 1 times

Sal 1 year, 8 months ago

I think this should be B.

upvoted 3 times

Question #159 *Topic 1*

A systems administrator needs to implement a security control that will prevent unknown malware from infecting a system in case the antivirus solution fails. Which of the following should the administrator implement?

A. A software whitelist

B. File integrity monitoring

C. A host-based IDS

D. Hardened baselines

Suggested Answer: *A*

Community vote distribution

A (71%) B (29%)

⊟ **i_bird** Highly Voted 1 year, 9 months ago

i'm thinking A

Application whitelisting is the practice of specifying an index of approved software applications or executable files that are permitted to be present and active on a computer system.

The goal of whitelisting is to protect computers and networks from potentially harmful applications.

upvoted 10 times

⊟ **FrancisDrake** Most Recent 5 months ago

Selected Answer: B

If the anti-virus solution has failed what good is the whitelist? The system has been infected. I'm going with file integrity monitoring.

upvoted 2 times

⊟ **Pongsathorn** 9 months, 3 weeks ago

Selected Answer: A

Software Whitelist (Option A): A software whitelist, also known as application whitelisting, is a security approach where only authorized and known applications are allowed to run on a system. This prevents the execution of any unauthorized or unknown applications, including malware. Even if the antivirus solution fails to detect a new malware variant, it won't execute on the system because it's not on the whitelist.

The other options are also valuable security controls, but they may not directly address the prevention of unknown malware:

upvoted 2 times

⊟ **Big_Gabe** 10 months, 1 week ago

Selected Answer: A

agree with A

upvoted 1 times

⊟ **Big_Gabe** 10 months, 1 week ago

Agree with A

upvoted 1 times

⊟ **BeauChateau** 1 year, 2 months ago

Selected Answer: A

A. A software whitelist is a security control that allows only known and trusted software to run on a system, while blocking all others. By implementing a whitelist, the system will only execute approved applications, preventing unknown malware from executing even if the antivirus solution fails. This will provide an additional layer of security and protect the system from unauthorized code. Therefore, A is the correct answer.

upvoted 1 times

⊟ **TheGinjaNinja** 1 year, 6 months ago

Selected Answer: A

Agree with i_bird on this one

upvoted 1 times

⊟ **Rob69420** 1 year, 9 months ago

A hardening process establishes a baseline of system functionality and security. The goal of hardening a system is to remove any unnecessary functionality and to configure what is left in a secure manner.

upvoted 2 times

⊟ **ironman_86** 1 year, 9 months ago

A is the correct answer

upvoted 3 times

Question #160 *Topic 1*

A cloud administrator is setting up a DR site on a different zone of the same CSP. The application servers are replicated using the VM replication, and the database replication is set up using log shipping. Upon testing the DR site, the application servers are unable to access the database servers. The administrator has verified the systems are running and are accessible from the CSP portal. Which of the following should the administrator do to fix this issue?

A. Change the database application IP.

B. Create a database cluster between the primary site and the DR site.

C. Update the connection string.

D. Edit the DNS record at the DR site for the application servers.

Suggested Answer: *C*

Community vote distribution

C (100%)

Pongsathorn 3 months, 3 weeks ago

Selected Answer: C

Update the Connection String (Option C): When you set up a DR site, it's common for the connection information to change. This includes the IP addresses or hostnames of the database servers in the DR site. The connection string used by the application servers must be updated to point to the correct database servers in the DR site. This ensures that the application servers can establish a connection to the database servers in the new location.

upvoted 3 times

BeauChateau 8 months, 3 weeks ago

Selected Answer: C

The administrator should update the connection string to fix the issue. Since the application servers are unable to access the database servers, it is likely that the connection string is pointing to the wrong IP address or hostname. By updating the connection string to the correct IP address or hostname of the database servers, the application servers will be able to communicate with the database servers and access the necessary data. Changing the database application IP, creating a database cluster, or editing the DNS record may not resolve the issue, as the root cause is likely related to the connection string being incorrect.

upvoted 3 times

ironman_86 1 year, 3 months ago

can someone explain why C?

upvoted 1 times

jiminycriminal 1 year, 3 months ago

The only thing I can think of is replicated VMs are probably trying to access the original database. A new connection string is needed for the DR VMs to connect to the DR database? Maybe?

upvoted 2 times

Question #161 *Topic 1*

A company recently subscribed to a SaaS collaboration service for its business users. The company also has an on-premises collaboration solution and would like users to have a seamless experience regardless of the collaboration solution being used. Which of the following should the administrator implement?

A. LDAP

B. WAF

C. VDI

D. SSO

Suggested Answer: *D*

Community vote distribution

D (100%)

BeauChateau 2 months, 2 weeks ago

Selected Answer: D

To provide users with a seamless experience, the administrator should implement Single Sign-On (SSO). SSO allows users to authenticate once and have access to multiple systems and services without having to enter their credentials repeatedly. This would allow the users to have a consistent login experience between the SaaS-based collaboration service and the on-premises collaboration solution.

upvoted 3 times

Question #162 *Topic 1*

A systems administrator has finished installing monthly updates to servers in a cloud environment. The administrator notices certain portions of the playbooks are no longer functioning. Executing the playbook commands manually on a server does not work as well. There are no other reports of issues. Which of the following is the MOST likely cause of this issue?

A. Change management failure

B. Service overload

C. Patching failure

D. Job validation issues

E. Deprecated features

Suggested Answer: *E*

Community vote distribution

E (69%) A (31%)

BeauChateau Highly Voted 8 months, 3 weeks ago

Selected Answer: E

E. Deprecated features are the most likely cause of the issue. It is possible that the monthly updates included changes that removed support for certain features or commands that were previously used in the playbook. When executing the playbook commands manually on a server, the administrator might have encountered errors or unexpected behavior due to the removal of those features or commands. It is important to review the release notes of updates before applying them to identify any deprecated features that could impact existing configurations.

upvoted 6 times

Pongsathorn Most Recent 3 months, 3 weeks ago

Selected Answer: E

Deprecated Features (Option E): Updates, especially in cloud environments, can deprecate or change features, APIs, or functions. If the playbooks were relying on deprecated features or functionality that has changed, they may no longer work as expected after the updates.

Change Management Failure (Option A): Change management typically involves tracking and coordinating changes to prevent issues. However, if the issue is related to deprecated features or changes in functionality caused by updates, it may not necessarily be a change management failure.

upvoted 2 times

SecPlus2022 6 months, 3 weeks ago

Selected Answer: E

Reason as sated by BeauChateau.

upvoted 1 times

TheGinjaNinja 1 year ago

Selected Answer: A

I think it might be A. As it is referring to playbooks

upvoted 4 times

tonytonyyyyy 1 year, 2 months ago

Why is it not C patching failures?

upvoted 2 times

Question #163 *Topic 1*

A systems administrator notices several VMs are constantly ballooning, while the memory usage of several other VMs is significantly lower than their resource allocation. Which of the following will MOST likely solve the issue?

A. Right-sizing

B. Bandwidth increase

C. Cluster placement

D. Storage tiers

Suggested Answer: *A*

Community vote distribution

A (100%)

Jay987654 5 months ago

Selected Answer: A

The most likely solution to this issue is A. Right-sizing.

Right-sizing involves adjusting the resources allocated to each VM based on their actual usage. This can help ensure that resources are distributed more evenly and efficiently, reducing the likelihood of some VMs constantly ballooning while others are underutilized. It's a common practice in virtualization management to optimize resource utilization.

upvoted 3 times

Question #164 *Topic 1*

The security team for a large corporation is investigating a data breach. The team members are all trying to do the same tasks but are interfering with each other's work. Which of the following did the team MOST likely forget to implement?

A. Incident type categories

B. A calling tree

C. Change management

D. Roles and responsibilities

Suggested Answer: *D*

Community vote distribution

D (100%)

Pongsathorn 3 months, 1 week ago

Selected Answer: D

The issue described, where team members are interfering with each other's work during a data breach investigation, suggests a lack of well-defined roles and responsibilities within the incident response team. This is often a crucial component of an effective incident response plan.

Answer: D. Roles and Responsibilities

Having clearly defined roles and responsibilities for team members in an incident response plan is essential to ensure that each person knows what they are responsible for and what tasks they should focus on during an incident. This helps in preventing duplication of efforts, streamlining the response process, and ensuring that critical tasks are not overlooked.

upvoted 3 times

Pongsathorn 3 months, 1 week ago

Roles and responsibilities within an incident response team may include designating incident coordinators, forensic analysts, communication liaisons, legal advisors, and other specialized roles. Each role should have a specific set of responsibilities and tasks assigned to it, and team members should work within the boundaries of their roles to maximize the efficiency and effectiveness of the response effort.

In addition to roles and responsibilities, incident response plans often include incident type categories (Answer A) to help categorize and prioritize incidents, but these alone may not prevent interference among team members. Calling trees (Answer B) are used for communication during incidents but don't address the issue of overlapping responsibilities. Change management (Answer C) is important for controlling changes to the IT environment but may not directly address the problem of team members interfering with each other's work during an ongoing incident.

upvoted 2 times

Pongsathorn 3 months, 3 weeks ago

Selected Answer: D

In a large corporation, when dealing with a data breach or any incident, it's crucial to have clearly defined roles and responsibilities for each team member. This ensures that everyone knows their specific tasks and areas of focus during the incident response process. Without well-defined roles and responsibilities, team members can interfere with each other's work, leading to confusion, duplication of efforts, and potentially making the incident response less effective.

upvoted 1 times

Question #165 *Topic 1*

A cloud engineer is responsible for managing two cloud environments from different MSPs. The security department would like to inspect all traffic from the two cloud environments. Which of the following network topology solutions should the cloud engineer implement to reduce long-term maintenance?

A. Chain

B. Star

C. Mesh

D. Hub and spoke

Suggested Answer: *D*

Community vote distribution

D (100%)

Sweety_Certified7 3 months, 1 week ago

Selected Answer: D

Hub and spoke topology is a centralized networking model where all traffic flows through a central hub. In this setup, the two cloud environments (spokes) connect to a central hub. The hub can be a network security device or a virtual network appliance capable of inspecting and monitoring traffic between the spokes. This topology simplifies management by centralizing control and inspection points, making it easier to implement security policies, monitor traffic, and perform maintenance tasks.

By implementing a hub and spoke topology, the cloud engineer can efficiently inspect all traffic from the two cloud environments without the complexity of a full mesh topology or the limitations of a chain topology. This reduces long-term maintenance efforts while ensuring comprehensive security monitoring across the network.

Therefore, option D. Hub and spoke is the most suitable choice.

upvoted 2 times