

EMC

Exam E20-020

Cloud Infrastructure Specialist Exam for Cloud Architects

Version: 7.0

[Total Questions: 64]

Question No : 1

A cloud architect is designing a private cloud for an organization. The organization has no existing backup infrastructure. They want to offer consumers the ability to backup virtual machine instances using image-based backups.

What should the cloud architect look for when selecting a backup application for this environment?

- A. Virtual machine hardware is on the backup application vendor's compatibility list
- B. Hypervisor servers' hardware is on the backup application vendor's compatibility list
- C. Backup application can be integrated with the selected CMP components
- D. Backup application supports a cloud gateway for accessing the cloud-based virtual machines

Answer: D

Explanation:

A cloud storage gateway provides basic protocol translation and simple connectivity to allow the incompatible technologies to communicate transparently. The gateway can make cloud storage appear to be a NAS filer, a block storage array, a backup target or even an extension of the application itself.

References: <http://searchcloudstorage.techtarget.com/definition/cloud-storage-gateway>

Question No : 2

When implementing QoS across a cloud network, how is storage traffic usually prioritized?

- A. Most important
- B. More important than tenant traffic but less important than management traffic
- C. More important than management traffic but less important than tenant traffic
- D. Least important

Answer: C

Question No : 3

Which set of general criteria should the cloud architect consider when selecting a cloud management platform?

- A. Multi-tenancy, API version, and cross-platform awareness

- B. Hypervisor, hybrid cloud capabilities, and authentication
- C. Geographic region, networked storage, and Internet access
- D. Network storage, operating system drivers, and update services

Answer: B

Explanation:

At a bare minimum, new clouds must be able to access all the physical and virtual infrastructure within an organization—Windows, AIX, or mainframe; hypervisors like vSphere, KVM, and Hyper-V; as well as public clouds including Microsoft Azure, Amazon Web Services (AWS), IBM SoftLayer, and more.

References: <http://www.bmc.com/blogs/3-fundamental-requirements-for-cloud-management-platforms/>

Question No : 4

You are designing consumer compute resources in an onsite private cloud. During an assessment, you discover that the organization's IT staff wants secure access to the underlying host OS. What should be included in the design to support this requirement?

- A. Host IDS configuration
Secure key infrastructure
Bridged management network
- B. Perimeter firewall configuration
VPN encryption
Separate management network
- C. Host OS firewall configuration
Central logging
Physically isolated management network
- D. Host OS firewall configuration
Secure key infrastructure
Separate management network

Answer: B

Question No : 5

A cloud architect is designing a distributed block storage solution that will support application HA. The solution will consist of 10 nodes with all SAS devices. There are five racks available in the data center. Each rack has a single top-of-rack access layer switch

with sufficient bandwidth to the aggregate layer. Each node will have multiple connections to the local top-of-rack switch.

The architect wants to provide a design that has the fewest number of storage pools maximizing the number of fault domains. How should the storage pools be designed?

- A. One pool with SSD nodes located in one rackOne pool with all SAS nodes located in a different rack
- B. One pool with 50% SSD and 50% SAS nodes distributed equally across five racks
- C. One pool with SSD nodes distributed equally across five racksOne pool with all SAS nodes distributed equally across five racks
- D. One pool with 50% SSD and 50% SAS nodes located in one rackOne pool with 50% SSD and 50% SAS nodes located in a different rack

Answer: C

Question No : 6

A cloud design contains multiple hosts that are running a hypervisor. Each host has only two physical 10 Gb network adapters that are aggregated for bandwidth and failover. The physical switch ports for each host are configured as access ports. Each host will support numerous virtual machines and network segments. These virtual machines will communicate with other virtual machines on the same network segments. Why would overlay networks be required for this design?

- A. Support for more than two virtual switches
- B. Support for more than one virtual segment
- C. Support for more than two virtual machines
- D. Support for more than one distributed virtual switch

Answer: B

Explanation:

An overlay network is a computer network that is built on top of another network. Nodes in the overlay network can be thought of as being connected by virtual or logical links, each of which corresponds to a path, perhaps through many physical links, in the underlying network.

Virtual overlay networks use tunneling protocols to extend isolated network segments between servers for multi-tenant data center networks.

References: <http://searchsdn.techtarget.com/tip/Virtual-overlay-networks-Tunneling-protocols-enable-multi-tenancy>

Question No : 7

What describes the storage categories represented by OpenStack Swift and EMC XtremIO requirements?

- A. Swift = Distributed Object Storage XtremIO = Central Storage
- B. Swift = Central Storage
XtremIO = Distributed File Storage
- C. Swift = Distributed Block Storage XtremIO = Distributed Object Storage
- D. Swift = Distributed File Storage XtremIO = Distributed Block Storage

Answer: A

Explanation:

OpenStack Swift is a globally-distributed object storage with a single namespace that's durable enough for the most demanding private clouds and now brought to you in an easy-to-deploy/scale/manage system.

XtremIO is a flash-based Storage Array.

References:

<https://www.swiftstack.com/#testimonial/2> <http://www.emc.com/collateral/white-papers/h11752-intro-to-XtremIO-array-wp.pdf>

Question No : 8

An organization plans to deploy a spine/leaf network topology to support a cloud design. Leaf switches will use layer-3 protocols to communicate with the spine switches. Hosts will each connect to two leaf switches using layer-2 protocols.

Which technology must be enabled between the host and leaf switches to provide the maximum throughput for a single data stream?

- A. Generic Network Virtualization Encapsulation
- B. Spanning Tree Protocol
- C. Equal-Cost Multi-Path Routing
- D. Multi-Chassis Aggregation

Answer: A

Explanation:

Generic Network Virtualization Encapsulation (Geneve) is the peacemaking protocol drafted to unify VXLAN, NVGRE, and whatever other tunneling protocols emerge for network virtualization. Geneve doesn't exactly replace VXLAN and other protocols. Rather,

it provides a common superset among them, so that outside software can provide hooks to Geneve rather than having to accommodate multiple encapsulation standards.

Note: NVGRE (Network Virtualization using Generic Routing Encapsulation) is a network virtualization technology that attempts to alleviate the scalability problems associated with large cloud computing deployments. It uses Generic Routing Encapsulation (GRE) to tunnel layer 2 packets over layer 3 networks. Its principal backer is Microsoft.

References: <https://www.sdxcentral.com/articles/news/intel-supports-geneve-unify-vxlan-nvgre/2014/09/>

Question No : 9

In addition to the operating system, what other components does the consumer manage in an IaaS cloud service model?

- A. Application, data, storage, and physical networking
- B. Data, middleware, application, and runtime
- C. Runtime, physical servers, application, and middleware
- D. Middleware, runtime, hypervisor, and application

Answer: B

Explanation:

In the case of IaaS the computing resource provided is specifically that of virtualised hardware, in other words, computing infrastructure.

IaaS clouds often offer additional resources such as a virtual-machine disk-image library, raw block storage, file or object storage, firewalls, load balancers, IP addresses, virtual local area networks (VLANs), and software bundles.

Figure: Cloud-computing layers accessible within a stack

Incorrect:

Not A: not Physical networking not C: Not physical servers. Not D: Not Hypervisors.

References: <http://www.interoute.com/what-iaas>

Question No : 10

An organization is implementing a backup solution for their private cloud. They are concerned that having the backup data stored onsite will expose them to loss in the event of a site-wide disaster. They are considering replicating the backup storage to an external site.