

Huawei

H12-893_V1.0 Exam

HCIP-Data Center Network V1.0

Questions & Answers Demo

Version: 4.0

Question: 1

Which of the following is not an advantage of link aggregation on CE series switches?

- A. Improved forwarding performance of switches
- B. Load balancing supported
- C. Increased bandwidth
- D. Improved reliability

Answer: A

Explanation:

Link aggregation, often implemented using Link Aggregation Control Protocol (LACP) on Huawei CloudEngine (CE) series switches, combines multiple physical links into a single logical link to enhance network performance and resilience. The primary advantages include:

Load Balancing Supported (B): Link aggregation distributes traffic across multiple links based on hashing algorithms (e.g., source/destination IP or MAC), improving load distribution and preventing any single link from becoming a bottleneck.

Increased Bandwidth (C): By aggregating multiple links (e.g., 1 Gbps ports into a 4 Gbps logical link), the total available bandwidth increases proportionally to the number of links.

Improved Reliability (D): If one link fails, traffic is automatically redistributed to the remaining links, ensuring continuous connectivity and high availability.

However, Improved Forwarding Performance of Switches (A) is not a direct advantage. Forwarding performance relates to the switch's internal packet processing capabilities (e.g., ASIC performance, forwarding table size), which link aggregation does not inherently enhance. While it optimizes link utilization, it doesn't improve the switch's intrinsic forwarding rate or reduce latency at the hardware level. This aligns with Huawei's CE series switch documentation, where link aggregation is described as enhancing bandwidth and reliability, not the switch's core forwarding engine.

Reference: Huawei CloudEngine Series Switch Configuration Guide – Link Aggregation Section.

Question: 2

In the DCN architecture, spine nodes connect various network devices to the VXLAN network.

- A. TRUE
- B. FALSE

Answer: A

Explanation:

In Huawei's Data Center Network (DCN) architecture, particularly with the CloudFabric solution, the spine-leaf topology is a common design for scalable and efficient data centers. VXLAN (Virtual Extensible LAN) is used to create overlay networks, enabling large-scale multi-tenancy and flexible workload placement.

Spine Nodes' Role: In this architecture, spine nodes act as the backbone, interconnecting leaf nodes (which connect to servers, storage, or other endpoints) and facilitating high-speed, non-blocking communication. Spine nodes typically handle Layer 3 routing and serve as VXLAN tunnel endpoints (VTEPs) or connect to devices that do, integrating the physical underlay with the VXLAN overlay network.

Connection to VXLAN: Spine nodes ensure that traffic from various network devices (via leaf nodes) is routed efficiently across the VXLAN fabric. They provide the high-bandwidth, low-latency backbone required for east-west traffic in modern data centers, supporting VXLAN encapsulation and decapsulation indirectly or directly depending on the deployment.

Thus, the statement is TRUE (A) because spine nodes play a critical role in connecting the underlay network (various devices via leaf nodes) to the VXLAN overlay, as per Huawei's DCN design principles. Reference: Huawei CloudFabric Data Center Network Solution White Paper; HCIP-Data Center Network Training Materials – VXLAN and Spine-Leaf Architecture.

Question: 3

Which of the following statements are true about common storage types used by enterprises?

- A. FTP servers are typically used for file storage.
- B. Object storage devices are typically disk arrays.
- C. Block storage applies to databases that require high I/O.
- D. Block storage typically applies to remote backup storage.

Answer: AC

Explanation:

Comprehensive and Detailed in Depth

A. FTP servers are typically used for file storage.

This is correct. FTP (File Transfer Protocol) servers are indeed a common way to store and share files. They are widely used for basic file storage and transfer needs.

B. Object storage devices are typically disk arrays.

This is incorrect. Object storage devices are not typically disk arrays in the traditional sense. Object storage is designed for massive amounts of unstructured data. While they use disks for persistence, they present data as objects with metadata, rather than as blocks or files. Object storage solutions often use distributed systems across many servers, not just a single array.

C. Block storage applies to databases that require high I/O.

This is correct. Block storage is ideal for applications that demand high I/O performance, such as databases. Block storage provides raw, unformatted data blocks, giving applications direct control and low latency.

D. Block storage typically applies to remote backup storage.

This is partially true, but not the typical primary use case. While block storage can be used for remote backups, it is generally considered less efficient and more expensive than object storage for this purpose. Object storage is better suited for large, unstructured backup datasets. Block storage is better for applications that need fast read/write speeds, such as databases and virtual machines.

Therefore, the correct answers are A and C.

Reference to Huawei Data Center Network documents:

Huawei storage product documentation detailing block storage (e.g., OceanStor Dorado), file storage, and object storage (e.g., OceanStor Pacific) characteristics and use cases.

Huawei white papers on data center storage architectures, which compare and contrast different storage types.

Huawei HCIP-Storage training materials, which will have very detailed information regarding each of the storage types, and their use cases.

Question: 4

Which of the following technologies are Layer 4 load balancing technologies? (Select All that Apply)

A. Nginx

B. PPP

C. LVS

D. HAProxy

Answer: A, C, D

Explanation:

Layer 4 load balancing operates at the transport layer (OSI Layer 4), using TCP/UDP protocols to distribute traffic based on information like IP addresses and port numbers, without inspecting the application-layer content (Layer 7). Let's evaluate each option:

A. Nginx: Nginx is a versatile web server and reverse proxy that supports both Layer 4 and Layer 7 load balancing. In its Layer 4 mode (e.g., with the stream module), it balances TCP/UDP traffic, making it a Layer 4 load balancing technology. This is widely used in Huawei's CloudFabric DCN solutions for traffic distribution. TRUE.

B. PPP (Point-to-Point Protocol): PPP is a Layer 2 protocol used for establishing direct connections between two nodes, typically in WAN scenarios (e.g., dial-up or VPNs). It does not perform load balancing at Layer 4 or any layer, as it's a point-to-point encapsulation protocol. FALSE.

C. LVS (Linux Virtual Server): LVS is a high-performance, open-source load balancing solution integrated into the Linux kernel. It operates at Layer 4, using techniques like NAT, IP tunneling, or direct routing to distribute TCP/UDP traffic across backend servers. It's a core Layer 4 technology in enterprise DCNs. TRUE.

D. HAProxy: HAProxy is a high-availability load balancer that supports both Layer 4 (TCP mode) and Layer 7 (HTTP mode). In TCP mode, it balances traffic based on Layer 4 attributes, making it a Layer 4 load balancing technology. It's commonly deployed in Huawei DCN environments. TRUE.

Thus, A (Nginx), C (LVS), and D (HAProxy) are Layer 4 load balancing technologies. PPP is not.

Reference: Huawei CloudFabric Data Center Network Solution – Load Balancing Section; HCIP-Data Center Network Training – Network Traffic Management.

Question: 5

Which of the following components is not required to provide necessary computing, storage, and network resources for VMs during VM creation?

- A. Nova
- B. Neutron
- C. Ceilometer
- D. Cinder

Answer: C

Explanation:

This question pertains to OpenStack, a common virtualization platform in Huawei's HCIP-Data Center Network curriculum, where components collaborate to create and manage virtual machines (VMs). Let's analyze each component's role in providing computing, storage, and network resources during VM creation:

A. Nova: Nova is the compute service in OpenStack, responsible for managing VM lifecycles, including provisioning CPU and memory resources. It's essential for providing computing resources during VM creation. Required.

B. Neutron: Neutron is the networking service, handling virtual network creation, IP allocation, and connectivity (e.g., VXLAN or VLAN) for VMs. It's critical for providing network resources during VM creation. Required.

C. Ceilometer: Ceilometer is the telemetry service, used for monitoring, metering, and collecting usage data (e.g., CPU utilization, disk I/O) of VMs. While useful for billing or optimization, it does not directly provide computing, storage, or network resources during VM creation. Not Required.

D. Cinder: Cinder is the block storage service, providing persistent storage volumes for VMs (e.g., for OS disks or data). It's essential for providing storage resources during VM creation if a volume is attached. Required.

Thus, C (Ceilometer) is not required to provision the core resources (computing, storage, network) for VM creation, as its role is monitoring, not resource allocation.

Reference: Huawei HCIP-Data Center Network Training – OpenStack Architecture; OpenStack Official Documentation – Service Overview.