Vendor: Nutanix

Exam Code: NCP-AII

Exam Name: AI Infrastructure

Version: DEMO

QUESTION 1
An administrator is configuring node categories in BCM for a DGX BasePOD cluster. They need to group all NVIDIA DGX H200 nodes under a dedicated category for GPU-accelerated workloads. Which approach aligns with NVIDIA's recommended BCM practices?

A. Avoid categories and configure each DGX node individually via CLI.
B. Create a new "dgx-h200" category, assign all DGX H200 nodes to it.
C. Use the existing "dgxnodes" category without modification, as it is preconfigured for all DGX systems.
D. Assign nodes to the "login" category to simplify Slurm integration.

Answer: B
Explanation:
In BCM, node categories are used to apply common configuration consistently across groups of similar systems. Creating a dedicated category for DGX H200 nodes and assigning those nodes to it supports repeatable management of GPU-accelerated workload settings across the BasePOD cluster.

QUESTION 2
To validate bisectional bandwidth across two racks in a Spectrum-X Ethernet fabric, which NCCL test configuration isolates East-West traffic?

A. Run without splits and analyze per-rack averages
B. NCCL_TESTS_SPLIT= "MOD 2" ./all_reduce_perf -g 8
C. NCCL_TESTS_SPLIT= "DIV 8" ./all_reduce_perf -g 1
D. NCCL_TESTS_SPLIT= "OR 0x7" ./all_reduce_perf -g 8

Answer: B
Explanation:
NCCL_TESTS_SPLIT="MOD 2" separates ranks into alternating groups, which is commonly used to force traffic across rack boundaries when ranks are arranged by rack. Running all_reduce_perf with all GPUs participates in the collective while isolating East-West paths needed to validate bisectional bandwidth between the two racks.

QUESTION 3
A network engineer is tasked with configuring the management, storage, and compute networks for a new DGX BasePOD deployment. Which statement best describes the network segmentation required for optimal operation?

A. A single VLAN for all types of network traffic.
B. Two networks: one for management and one for compute.
C. Four networks: compute, storage, out-of-band, and management.

Answer: C
Explanation:
DGX BasePOD deployments use separate network segments for distinct traffic classes so that management access, out-of-band hardware management, storage traffic, and high-performance compute communication do not interfere with each other. This segmentation improves performance, isolation, troubleshooting, and operational reliability.

QUESTION 4

During cluster deployment, the UFM Cable Validation Tool reports “Wrong-neighbor” errors on multiple InfiniBand links. What is the most efficient way to resolve this issue?

A. Verify LLDP data against topology files and remediate.
B. Reboot all leaf switches to force LLDP rediscovery.
C. Disable FEC on all switches to bypass neighbor validation.
D. Replace all affected cables with higher-grade OM5 fiber optics.

Answer: A
Explanation:
“Wrong-neighbor” errors indicate that the discovered cable peer does not match the expected topology. The efficient resolution is to compare LLDP-discovered neighbor data with the planned topology files, identify the mismatched links, and correct the cabling or topology definition accordingly.

QUESTION 5
During HPL execution on a DGX cluster, the benchmark fails with “not enough memory” errors despite sufficient physical RAM. Which HPL.dat parameter adjustment is most effective?

A. Increase block size to 6144 to maximize GPU utilization.
B. Reduce the problem size while maintaining the same block size.
C. Set PMAP to 1 to enable process mapping.
D. Disable double-buffering via BCAST parameter.

Answer: B
Explanation:
The HPL problem size directly controls the memory footprint of the benchmark. Reducing the problem size lowers the amount of memory required while keeping the block size unchanged, which is the most effective adjustment when HPL reports memory allocation failures despite the system having sufficient physical RAM.

QUESTION 6
Your company is planning to expand its AI capabilities significantly over the next five years. To future-proof your storage infrastructure, you need a solution that can scale in both capacity and performance. Which of the following strategies best ensures that your storage infrastructure remains adaptable to future AI demands?

A. Implement on-premises block storage system with periodic hardware upgrades.
B. Implement single-tier cloud storage solution to leverage cloud scalability.
C. Use a hybrid cloud model combining scalable cloud resources with on-premises infrastructure.
D. Deploy an all-flash array and remove data tiering to reduce latency.

Answer: C
Explanation:
A hybrid cloud model provides the flexibility to scale capacity and performance over time by combining on-premises infrastructure for predictable, high-performance workloads with cloud resources for elastic growth. This approach keeps the storage architecture adaptable as AI data volumes, training workloads, and performance requirements increase.

QUESTION 7
A system administrator needs to install a container toolkit and successfully run the following

commands:

```
sudo apt-get update
sudo apt-get install -y nvidia-container-toolkit
sudo nvidia-ctk runtime configure --runtime=docker
```

What step should be taken next to finish the installation?

A. systemctl restart docker
B. apt-get remove nvidia-container-toolkit for
C. dpkg –i doca-host-repo-ubuntu<version>_am64.deb</version>
D. apt-get install cuda-drivers

Answer: A
Explanation:
After configuring the NVIDIA Container Toolkit runtime for Docker with nvidia-ctk runtime configure --runtime=docker, Docker must be restarted so it reloads the updated runtime configuration and can run GPU-enabled containers correctly.

QUESTION 8
If two ports must be connected, but one is SFP and one is QSFP, for example, to connect a 25 GbE HOST CHANNEL ADAPTER to a QSFP port capable of both 100 GbE and 25 GbE, which of the following solutions would best meet this requirement?

A. QSA Adapter
B. SFP to 1G BASE-T (RJ45) adapter
C. SFP Connectors

Answer: A
Explanation:
A QSA adapter is used to connect an SFP/SFP28 transceiver or cable to a QSFP/QSFP28 port when the QSFP port supports the lower-speed mode. This fits the requirement of connecting a 25 GbE SFP-based host adapter to a QSFP port capable of operating at 25 GbE.

QUESTION 9
A system administrator is installing a GPU into a server and needs to avoid damaging the device. What item should be used?

A. Anti-ESD strap
B. Gloves
C. Electric screwdriver
D. Protective film

Answer: A
Explanation:
An anti-ESD strap should be used when installing a GPU to prevent electrostatic discharge from damaging sensitive electronic components on the device.

QUESTION 10
After configuring NGC CLI with ngc config set, a user receives “Authenticated failed” errors when

pulling containers. What step was most likely omitted?

A. Installing the CLI with apt-get instead of manual extraction.
B. Entering the API key during ngc config set or storing it in ~/.ngc/config.
C. Running sudo systemctl restart docker after configuration.
D. Setting –-format_type=json to enable API interactions.

Answer: B
Explanation:
NGC authentication requires a valid API key to be provided during ngc config set or stored in the user's ~/.ngc/config file. Without the API key, container pulls from NGC fail because the CLI cannot authenticate to the registry.

QUESTION 11
When updating the firmware on an NVLink switch transceiver, how can an engineer apply new firmware without interrupting the network?

A. flint -d -lid 27 -linkx -link_auto_update -activate
B. Physically disconnect and reconnect the transceiver.
C. mlxfwreset -d lid 27 reset --yes to reset the transceiver.
D. nv action reboot system to force immediate activation.

Answer: A
Explanation:
flint with the LinkX auto-update and activate options is used to update and activate NVLink switch transceiver firmware in a non-disruptive way when supported. This applies the firmware without requiring a physical reseat or disruptive device/system reset.

QUESTION 12
You are a network administrator responsible for configuring an East-West (E/W) Spectrum-X fabric using SuperNIC. The Bluefield 3 devices in your network should be a set to NIC mode with RoCE enabled to optimize data flow between servers. You have access to the Spectrum-X management tools and the necessary documentation. You need to use specific configuration commands to achieve this setup. Which of the following steps and commands are necessary to configure the Bluefield 3 devices in NIC mode for the E/W Spectrum-xX fabric using SuperNIC? (Choose two.)

A. Use the command sudo mixconfig -d /dev/mst/<device> set DISABLE_SPECTRUM_X=1 to reduce overhead.
B. Use the command sudo mixconfig -d /dev/mst/<device> set DPU_MODE=1 to set up the Bluefield 3 devices in DPU (Data Processing Unit) mode.
C. Use the command sudo mixconfig -d /dev/mst/<device> set INTERNAL_CPU_OFFLOAD_ENGINE=1 to configure the SuperNIC to operate in NIC mode.
D. Use the command sudo mixconfig -d /dev/mst/<device> set LINK_TYPE_P1=2 to enable Ethernet on the Bluefield 3 devices.

Answer: CD
Explanation:
For an East-West Spectrum-X fabric using SuperNIC, the BlueField-3 device must operate in NIC/SuperNIC mode and use Ethernet links for RoCE traffic. Setting the internal CPU offload engine enables the NIC-mode behavior required for SuperNIC operation, and setting the port link type to Ethernet enables the correct network mode for Spectrum-X RoCE connectivity.

QUESTION 13
A 24-hour HPL burn-in fails with “illegal value” errors during the first iteration. Which initial troubleshooting step resolves this without compromising burn-in validity?

A. Reduce test duration to 12 hours.
B. Switch from FP64 to FP32 precision.
C. Verify the matrix size is divisible by block size.
D. Disable GPU affinity.

Answer: C
Explanation:
HPL “illegal value” errors at the first iteration commonly indicate an invalid HPL.dat configuration. Verifying that the matrix size is compatible with the block size corrects the parameter issue while preserving the intended burn-in workload duration and precision.

QUESTION 14
An administrator needs to verify HA functionality after configuring BCM. Which command confirms the active head node and failover readiness?

A. systemctl restart cmdaemon to force a failover test.
B. nvsm show health to validate GPU status on both head nodes.
C. ping <secondary-head-node-ip> to test basic connectivity.
D. cmsh status to check HA status and active/standby roles.

Answer: D
Explanation:
cmsh status is the appropriate BCM command to verify HA state, including which head node is currently active and whether the standby node is ready for failover. It provides the cluster management service status needed to confirm HA functionality after configuration.

QUESTION 15
Refer to the output:

```
~$ sudo nvsm show healthinfo
--Timestamp: Sat Dec 16 16:26:32 2017 -0800
Version: 17.12-5
Checks------BIOS Revision [5.11].....
DGX Serial Number [YSY72800016].....
Verify installed DIMM memory sticks.....Healthy
...[output truncated]
Verify Ethernet controllers.....Healthy
Verify installed GPU's.....Unhealthy
Checking output of 'lspci' for expected GPU's
Missing GPU at PCI address "07:00.0"
Verify installed InfiniBand controllers.....Healthy
Verify PCIe switches.....Healthy
...[output truncated]
```

What insights can a system administrator gain regarding the DGX system’s health?

A. The system has passed the hardware health check successfully.
B. A GPU driver upgrade has failed.
C. A GPU is missing on the DGX system.
D. A GPU tray upgrade failed.

Answer: C
Explanation:
The health output shows the GPU check is unhealthy and reports a missing GPU at PCI address 07:00.0. This indicates the DGX system is not detecting one of the expected GPUs.

QUESTION 16
During East-West fabric validation on a 64-GPU cluster, an engineer runs all_reduce_perf and observes an algorithm bandwidth of 350 GB/s and bus bandwidth of 656 GB/s. What does this indicate about the fabric performance?

A. Critical failure; bus bandwidth exceeds hardware capabilities
B. Suboptimal performance; algorithm bandwidth should match bus bandwidth
C. Inconclusive; rerun with point-to-point tests
D. Optimal performance; bus bandwidth near theoretical peak for NDR InfiniBand

Answer: D
Explanation:
An all_reduce_perf result showing around 350 GB/s algorithm bandwidth and 656 GB/s bus bandwidth indicates strong collective communication performance for a 64-GPU NDR InfiniBand fabric. The bus bandwidth being near the expected peak confirms the East-West fabric is performing optimally and GPUDirect RDMA paths are healthy.

QUESTION 17
An engineer is tasked with configuring Out-of-Band (OOB) management for a DGX BasePOD deployment. Which network design will best ensure secure and reliable OOB management operations?

A. Place all BMC and management interfaces on an isolated OOB network with access restricted by firewall rules.
B. Use a single VLAN for both OOB management and compute fabric to simplify network design.
C. Configure OOB management interfaces to be accessible from any subnet within the data center for maximum flexibility.
D. Connect OOB management ports to the same switch as user traffic for easier troubleshooting.

Answer: A
Explanation:
OOB management should be isolated from production and compute traffic because BMC interfaces provide privileged hardware-level access. Placing BMC and management interfaces on a dedicated OOB network with firewall-restricted access improves security, reliability, and operational control for DGX BasePOD management.

Thank You for Trying Our Product

Lead2pass Certification Exam Features:

★ More than **99,900** Satisfied Customers Worldwide.

★ Average **99.9%** Success Rate.

★ **Free Update** to match latest and real exam scenarios.

★ **Instant Download** Access! No Setup required.

★ Questions & Answers are downloadable in **PDF** format and **VCE** test engine format.

★ Multi-Platform capabilities - **Windows, Laptop, Mac, Android, iPhone, iPod, iPad**.

★ **100%** Guaranteed Success or **100%** Money Back Guarantee.

★ **Fast**, helpful support **24x7**.

View list of all certification exams: http://www.lead2pass.com/all-products.html

10% Discount Coupon Code: **ASTR14**