



- Expert Verified, Online, **Free**.



CERTIFICATION TEST

- CertificationTest.net - Cheap & Quality Resources With Best Support

A system engineer needs to set the vGPU scheduling behavior for all GPUs to share the scheduling equally with the default time slice length. What command should be used?

- A. esxcli system module parameters set -m nvidia -p "NVreg_RegistryDwords=RmPVMRL=0x00"
- B. esxcli system module parameters set -m nvidia -p "NVreg_RegistryDwords=RmPVMRL=0x@1"
- C. esxcli graphics module parameters set -m nvidia -p "NVreg_RegistryDwords=RmPVMRL=0x01"
- D. esxcli system module parameters set -m nvidia -p "NVreg_RegistryDwords=FRL=@x01"

Suggested Answer: A

Community vote distribution

B (100%)

🗉 👤 **Rajivjain** 2 weeks, 3 days ago

Selected Answer: B

Setting RmPVMRL=0x01 via esxcli system module parameters enables Equal Share scheduling with the default time slice.
upvoted 1 times

During a multi-day NeMo burn-in, intermittent "GPU fell off bus" errors occur. Which diagnostic approach isolates hardware faults?

- A. Run DCGM diagnostics alongside burn-in to monitor GPU health metrics
- B. Switch from BERT to GPT models for simpler computations
- C. Enable HPL_USE_NVSHMEM for alternative memory sharing
- D. Reduce blocksize to 500MB to lower memory pressure

Suggested Answer: A

Community vote distribution

A (100%)

🗨️ 👤 **Rajivjain** 2 weeks, 3 days ago

Selected Answer: A

DCGM diagnostics run concurrently to catch Xid errors and hardware-level GPU faults that cause bus drop-offs.
upvoted 1 times

An engineer needs to verify the current firmware versions of all components (ATF, BSP, NIC, UEFI) on a BlueField-3 DPU's BMC. Which Redfish API command provides this information?

- A. `mstflint -d <PCI_ID> query full`
- B. `curl -k -u root:<password> -X GET https://<DPU-BMC-IP>/redfish/v1/UpdateService/FirmwareList`
- C. `curl -k -u root:<password> -X GET https://<DPU-BMC-IP>/redfish/v1/UpdateService/FirmwareInventory`
- D. `mlxconfig -d <dev> q`

Suggested Answer: C

Community vote distribution

C (100%)

🗨️ 👤 **Rajivjain** 2 weeks, 3 days ago

Selected Answer: C

The Redfish FirmwareInventory endpoint lists versions of all BMC-managed components (ATF, BSP, NIC, UEFI).
upvoted 1 times

An engineer needs to verify NVLink isolation on a single node with 8 GPUs.

Which NCCL test configuration stresses switch bisection bandwidth?

- A. Use NCCL_TESTS_SPLIT= "DIV 8" with point-to-point tests
- B. Use all_reduce_perf -b 8 -e 16G -f2 -g 8 with NCCL_TESTS_SPLIT= "AND 0x1"
- C. Use all_reduce_perf -b 8 -e 16G -f2 -g 8 without splits
- D. Use reduce_scatter_perf -b 8 -e 16G -f2 -g 4

Suggested Answer: B

Community vote distribution

B (100%)

🗨️ 👤 **Rajivjain** 2 weeks, 3 days ago

Selected Answer: B

The AND 0x1 split partitions GPUs into subgroups that force traffic across the NVSwitch fabric, stressing bisection bandwidth.
upvoted 1 times

You are validating the environment of an NVIDIA GPU-accelerated data center during post-deployment checks.

Which one action is essential to confirm that power and cooling are sufficient for the stable operation of NVIDIA DGX H100 systems?

- A. Use NVSM to disable unused PCIe devices to reduce overall system heat output.
- B. Review the system BIOS to ensure GPU overclocking is enabled for maximum performance.
- C. Verify that each DGX system is connected to redundant, properly rated PDUs and that all power supplies are reporting nominal input.
- D. Confirm the system fans are running at 100% under all workloads to prevent overheating.

Suggested Answer: C

Community vote distribution

C (100%)

🗨️ 👤 **Rajivjain** 2 weeks ago

Selected Answer: C

Verifying redundant, properly rated PDUs and nominal PSU input confirms the power/cooling infrastructure meets DGX H100 requirements for stable operation.

upvoted 1 times

An engineer needs to validate 400G DAC cable signal integrity in a DGX cluster.
Which CVT metric best identifies marginal cables needing replacement?

- A. Effective BER > $1.5E-254$ during a ≤ 6 -hour monitoring window.
- B. Transceiver model matching QSFP-DD specifications.
- C. Temperature fluctuations > 5°C during validation.
- D. Lane power variance < 3dB across all transceivers.

Suggested Answer: A

Community vote distribution

A (100%)

 **Rajivjain** 2 weeks ago

Selected Answer: A

An effective BER threshold above $1.5E-254$ within a short monitoring window flags marginal DAC cables for replacement.
upvoted 1 times

A user encounters "permission denied" errors when running GPU-accelerated containers on a Secure Boot-enabled system. What resolves this?

- A. Disable SELinux to relax unnecessary security policies.
- B. Reinstall Docker without the NVIDIA runtime.
- C. Enroll the MOK and sign NVIDIA kernel modules.
- D. Run Docker with sudo for elevated privileges.

Suggested Answer: C

Community vote distribution

C (100%)

🗨️ 👤 **Rajivjain** 2 weeks ago

Selected Answer: C

With Secure Boot enabled, NVIDIA kernel modules must be signed and the MOK enrolled for the kernel to load them, resolving the permission errors.
upvoted 1 times

An InfiniBand administrator needs to run performance benchmarks on new devices added to the fabric.
What tool should be used to check the latency?

- A. tcpdump
- B. ib_write_lat
- C. perfmon
- D. ibdiagnet

Suggested Answer: B

Community vote distribution

B (100%)

🗨️ 👤 **Rajivjain** 2 weeks ago

Selected Answer: B

ib_write_lat measures RDMA write latency between InfiniBand devices for performance benchmarking.
upvoted 1 times

A cluster administrator is preparing to update the firmware on a DGX H100 system, including the GPU tray (baseboard).

What is the correct sequence of steps to perform a safe and successful firmware upgrade?

- A. Update the BMC and skip the GPU tray and motherboard tray updates if the system appears healthy.
- B. Update the GPU tray first, then the motherboard tray, and reboot the BMC after all updates are complete.
- C. Stop all GPU activity, update and reboot the BMC, update motherboard and tray components, perform a cold reset, and verify completion.
- D. Perform a cold reset, stop all GPU activity, update and reboot the BMC, update motherboard and tray components, and verify completion.

Suggested Answer: *C*

Currently there are no comments in this discussion, be the first to comment!

One of the nodes in a cluster is not running as fast as the others and the system administrator needs to check the status of the GPUs on that system.

What command should be used?

- A. iblinkinfo
- B. lspci | grep NVIDIA
- C. nvidia-smi
- D. nvidia-gpu-status

Suggested Answer: C

Currently there are no comments in this discussion, be the first to comment!

A systems engineer is updating firmware across a large DGX cluster using automation.

What is the best practice for minimizing risk and ensuring cluster health during and after the process?

- A. To save time, simultaneously update all nodes in the cluster without draining or diagnostics.
- B. Drain nodes from the scheduler, run pre-update diagnostics, update firmware in batches, and verify health post-update before scaling to the next batch.
- C. Drain nodes from the scheduler, update firmware in batches, skip diagnostics and verify health post-update before scaling to the next batch.
- D. Update nodes that have reported faults, leaving others on older firmware.

Currently there are no comments in this discussion, be the first to comment!

A customer has just completed the first boot of their DGX system and is prompted to create an administrative user.

What is the correct approach for setting up this user to ensure secure BMC and GRUB access?

- A. Create a unique, strong, lower-case username and password that will be used for both BMC and GRUB access, avoiding default or weak credentials.
- B. Skip the creation of a new user and retain the default admin account for BMC and GRUB access.
- C. Create separate usernames for BMC and GRUB to maximize flexibility.
- D. Use "sysadmin" as the username and a simple password for ease of management.

Suggested Answer: A

Currently there are no comments in this discussion, be the first to comment!

As the infrastructure lead for an NVIDIA AI Factory deployment, you have just uploaded the latest supported firmware packages to your DGX system. It is now critical to ensure all hardware components run the new firmware and the DGX returns to full operational capability. Which sequence best guarantees that all relevant components are correctly running updated firmware according to NVIDIA's documentation and recommended operational steps?

- A. Execute a single AC power cycle on the DGX after the update process, then reset the software stack and verify status using diagnostic commands on each node for confirmation of all component updates.
- B. Initiate a cold power cycle on the system to activate firmware for components, reset the BMC using the recommended command, and perform an AC power cycle to ensure EROT and CPLD firmware is activated.
- C. Perform a software-driven restart on the operating system of every compute node, then use advanced tools to check firmware status, and reissue update commands if any firmware appears inactive afterward.
- D. Initiate a cold power cycle on all node trays to activate firmware, follow with a DGX reboot procedure, and use the management interface to finish activating CPLD firmware on the host.

Suggested Answer: *B*

Currently there are no comments in this discussion, be the first to comment!

You are tasked with updating both NVIDIA GPU drivers and DOCA drivers on a set of servers used for AI workloads. The environment previously had an older driver stack and custom kernel modules.

What is the most important step to successfully upgrade the drivers without causing conflicts?

- A. Update the GPU driver leaving the DOCA and OFED drivers unchanged as long as they are detecting the hardware properly.
- B. Uninstall all existing GPU and DOCA-related drivers and associated kernel modules before the new install
- C. Keep the older driver running alongside the new version in case you need to roll back the upgrade.
- D. Validate the driver version post-install since the fresh install will overwrite the legacy drivers.

Suggested Answer: *B*

Currently there are no comments in this discussion, be the first to comment!

During multi-node HPL burn-in, GPUs show uneven utilization.
Which configuration ensures balanced workload distribution?

- A. HPL_RUN_GEMM_TESTS to skip validation
- B. Set --gpu-affinity and --cpu-affinity to align GPU and NUMA nodes
- C. HPL_OOC_TILE_M to 8192 for larger blocks
- D. Enable HPL_USE_NVSHMEM=1 for shared memory acceleration

Suggested Answer: *B*

Currently there are no comments in this discussion, be the first to comment!

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.

Suggested Answer: *B*

Currently there are no comments in this discussion, be the first to comment!

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`

Suggested Answer: *B*

Currently there are no comments in this discussion, be the first to comment!

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.

Suggested Answer: C

Currently there are no comments in this discussion, be the first to comment!

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.

Suggested Answer: A

Currently there are no comments in this discussion, be the first to comment!

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.

Suggested Answer: *B*

Currently there are no comments in this discussion, be the first to comment!

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.

Suggested Answer: *C*

Currently there are no comments in this discussion, be the first to comment!

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

Suggested Answer: A

Currently there are no comments in this discussion, be the first to comment!

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

Suggested Answer: A

Currently there are no comments in this discussion, be the first to comment!

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

Suggested Answer: A

Currently there are no comments in this discussion, be the first to comment!

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.

Suggested Answer: *B*

Currently there are no comments in this discussion, be the first to comment!

A system administrator needs to validate a GPU-based server and ensure that no errors occur under load.

What command should be used?

- A. stress-test --usage
- B. nvsm stress-test
- C. nvsm dump health
- D. nvsm show health

Suggested Answer: *B*

Currently there are no comments in this discussion, be the first to comment!

After configuring HA, the administrator runs `cmsh status` and notices the secondary head node reports `mysql [FAIL]`. What is the most likely cause?

- A. The secondary head node lacks NVIDIA GPU drivers.
- B. The BCM license expired after HA configuration.
- C. The cluster nodes are powered on during the HA configuration.
- D. Network connectivity issues between the primary and secondary head nodes.

Suggested Answer: *D*

Currently there are no comments in this discussion, be the first to comment!

Why is it important to provide a large and high-performance local cache (using SSDs configured as RAID-0) for deep learning workloads on DGX systems?

- A. Using local SSD cache in RAID-0 enables direct GPU access to files without host CPU involvement, further boosting performance.
- B. Local SSD cache allows users to increase the number of NFS threads on the server without impacting storage reliability.
- C. A local SSD cache in RAID-0 ensures that most training data is read only once from the network, significantly reducing NFS traffic.
- D. Local SSD cache in RAID-0 is necessary to provide redundancy in case one of the drives fails during long training runs.

Suggested Answer: *C*

Currently there are no comments in this discussion, be the first to comment!

What information does the “ibnodes” command display?

- A. All host & server names
- B. All hosts & switches
- C. All server names
- D. All channel adapters

Suggested Answer: *B*

Currently there are no comments in this discussion, be the first to comment!

A system administrator needs to install a GPU/DPU in a server. The server has a free PCI-e slot, there are enough free PCI-e lanes, and there is enough room for the card.

Which procedure should be followed?

- A. Ensure the server has enough power. Make sure the server is up and running with attached cables. Wear an ESD bracelet.
- B. Ensure the server has enough power. Verify compatibility of cables with server's platform. Make sure the server is down to remove cables safely. Do not wear an ESD bracelet.
- C. Ensure the server has enough power. Verify compatibility of cables with server's platform. Make sure the server is down to remove cables safely. Wear an ESD bracelet.
- D. Ensure the server has enough power. Make sure the server is down to remove cables safely. Wear an ESD bracelet.

Suggested Answer: C

Currently there are no comments in this discussion, be the first to comment!

A cluster administrator needs to validate transceiver firmware versions across 200 ports using UFM.

Which GUI-based method provides a consolidated view?

- A. Use "Topology" view to visually inspect cable icons.
- B. Navigate to "Devices" > select a switch > "Cables" tab to see ASIC firmware and transceiver versions.
- C. Export all switch logs and grep for "FW Version"
- D. Run `mlxlink -d lid-<LID> -m` on each port manually

Suggested Answer: B

Currently there are no comments in this discussion, be the first to comment!

What command sequence is used to identify the exact name of the server that runs as the master SM in a multi-node fabric?

- A. ibstat, then sminfo
- B. sminfo, then smpquery ND
- C. ibis, then ibsim
- D. sminfo, then smpquery NI

Suggested Answer: *B*

Currently there are no comments in this discussion, be the first to comment!

An InfiniBand server stops working, and a system administrator runs the "ibstat" command that provides the following output:

```
CA 'mix5_1'
  CA type: MT4115
  Number of ports: 2
  Firmware version: 10.20.1010
  Hardware version: 0
  Node GUID: 0x0002c90300002f78
  System image GUID: 0x0002c90300002f7b
  Port 1:
    State: Initializing
    Physical state: LinkUp
    Rate: 100
    Base lid: 0
    LMC: 0
    SM lid: 0
    Capability mask: 0x0251086a
    Port GUID: 0x0002c90300002f79
    Link layer: InfiniBand
```

What is the cause of the issue?

- A. The HCA port is faulty.
- B. There is no running SM in the fabric.
- C. The cable is disconnected.
- D. The neighboring switch port is faulty.

Suggested Answer: B

Currently there are no comments in this discussion, be the first to comment!

After NCCL burn-in reports “transport retry count exceeded”, which corrective action addresses the underlying fabric issue?

- A. Switch from Ring to Tree algorithms via NCCL_ALGO=TREE
- B. Inspect InfiniBand link quality metrics (BER, symbol errors) and replace faulty cables
- C. Increase NCCL_IB_TIMEOUT to tolerate longer latencies
- D. Reduce message size to decrease network utilization

Suggested Answer: *B*

Currently there are no comments in this discussion, be the first to comment!

An infrastructure engineer in an AI factory has successfully replaced a power supply unit on an NVIDIA DGX H100. After installation, both the IN and OUT LEDs on the new power supply illuminate solid green.

Which NVSM CLI command should the engineer use to quickly verify the overall system status and ensure it is operating as expected?

- A. `nvsm show powermode`
- B. `nvsm show health`
- C. `nvsm show power`
- D. `nvsm show alerts`

Suggested Answer: *B*

Currently there are no comments in this discussion, be the first to comment!

After initial setup and health checks, the DGX H100 system administrator wants to verify that containers can access GPUs before running production workloads.

Which method is recommended for this validation?

- A. `sudo docker run --gpus all --rm nvcr.io/nvidia/cuda:12.1.1-base-ubuntu22.04 nvidia-smi`
- B. `sudo docker run --gpus all --rm nvcr.io/nvidia/cuda:12.1.1-base-ubuntu22.04 systemctl`
- C. `sudo docker run --gpus all --rm nvcr.io/nvidia/cuda:12.1.1-base-ubuntu22.04 ls-la`
- D. `sudo docker run --rm nvcr.io/nvidia/cuda:12.1.1-base-ubuntu22.04 nvidia-smi`

Suggested Answer: A

Currently there are no comments in this discussion, be the first to comment!

After ClusterKit reports "GPU-Host latency exceeds threshold", which NVIDIA diagnostic tool should be used to isolate hardware faults?

- A. `nvidia-smi topo -m` to inspect GPU topology connections
- B. Re-run ClusterKit with `--stress=gpu -Y 60` to extend test duration
- C. DCGM Diags `dcgmi diag -r 2`
- D. `ib_write_bw` to measure InfiniBand bandwidth between nodes

Suggested Answer: C

Currently there are no comments in this discussion, be the first to comment!

A media company is developing an AI platform for video content analysis that requires storing and processing large volumes of unstructured video data. The platform must support high throughput for data ingestion and provide efficient access for real-time analytics. Given these requirements, which storage strategy should the company implement?

- A. Block storage for low latency and high performance
- B. Object storage for scalability and metadata management
- C. Tape storage for its cost-effectiveness and archival capabilities
- D. File storage for hierarchical organization and easy navigation

Suggested Answer: *B*

Currently there are no comments in this discussion, be the first to comment!

A financial services firm is deploying an AI model for fraud detection that requires rapid inference and data retrieval across multiple sites. Which feature should their storage system prioritize?

- A. Tape backup systems.
- B. Low-cost HDD solutions.
- C. High capacity with moderate speed.
- D. Multi-protocol data access with low latency.

Suggested Answer: *D*

Currently there are no comments in this discussion, be the first to comment!

What command is needed to measure BER (Bit Error Rate)?

- A. `mlxlink -d <device> -c -e`
- B. `mlxconfig -d <device> q`
- C. `msflint -d <device> q full`
- D. `ethtool -S<device>`

Suggested Answer: A

Currently there are no comments in this discussion, be the first to comment!

A System Administrator needs to change the scheduling behavior of a single GPU to use a fixed share scheduler.

What command achieves this?

- A. `mixconfig -d /dev/mst/mt4123_pciconf0 set LINK_TYPE_P1=2`
- B. `esxcli -i 0 -mig 18`
- C. `nvidia-smi -i 0 -mig 1`
- D. `esxcli system module parameters set -m nvidia -p`

Suggested Answer: *D*

Currently there are no comments in this discussion, be the first to comment!

During cluster validation, the Cable Validation Tool (CVT) reports “Underperforming (BER)” for an InfiniBand link. Which BER thresholds indicate a critical signal quality issue requiring cable replacement?

- A. Raw BER > 1e-6 or Effective BER > 1.5E-254 for ≤6hr measurements
- B. Effective BER > 0 during the first 125 minutes of link operation
- C. Temperature > 85 °C on transceiver module
- D. Rx power variance > 3dB between lanes

Suggested Answer: A

Currently there are no comments in this discussion, be the first to comment!

During BCM cluster setup, an engineer must configure bonded network interfaces on DGX nodes for high availability.

Which cmsh command sequence properly configures a bond0 interface with two physical NICs?

- ```
device use dgx001
A. interfaces delete enp225s0f1np1
 interfaces delete enp97s0f1np1

 device use dgx001
B. interfaces set enp225s0f1np1 network internalnet
 interfaces set enp97s0f1np1 network internalnet

 device use dgx001
 interfaces add bond bond0
C. append interfaces enp225s0f1np1 enp97s0f1np1
 set mode 1
 set network internalnet

 device use dgx001
 interfaces add vlan vlan100
D. set parent bond0
 set mode 1
 set network internalnet
```

**Suggested Answer:** *C*

Currently there are no comments in this discussion, be the first to comment!

A company has a registered NGC account and their server has NCG CLI installed.

What step should be taken first to gain access to NGC?

- A. ngc config get
- B. ngc config update
- C. ngc config set
- D. ngc init

**Suggested Answer:** *C*

Currently there are no comments in this discussion, be the first to comment!

You are installing the operating system as part of the initial setup for a new NVIDIA Base Command Manager (BCM) cluster. Which two of the following actions are essential for a successful OS installation on the cluster's head node? (Choose two.)

- A. Start the head node OS installation process with the system BIOS set to legacy boot mode instead of UEFI.
- B. Configure network switches for PXE boot to all compute nodes before installing the OS on the head node.
- C. Set the desired time zone and configure NTP synchronization during the OS installation wizard.
- D. Download the latest BCM ISO and verify its integrity using the provided checksum, then start the installation.

**Suggested Answer:** *CD*

Currently there are no comments in this discussion, be the first to comment!

A user wants to restrict a Docker container to use only GPUs 0 and 2.

Which command achieves this?

- A. `docker run --device /dev/nvidia0./dev/nvidia2 nvidia/cuda:12.1-base nvidia-smi`
- B. `docker run --gpus "device=0.2" nvidia/cuda:12.1-base nvidia-smi`
- C. `docker run --gpus all nvidia/cuda:12.1-base nvidia-smi --id=0,2`
- D. `docker run --e NVIDIA_VISIBLE_DEVICES=0.2 nvidia/cuda:12.1-base nvidia-smi`

**Suggested Answer:** *B*

Currently there are no comments in this discussion, be the first to comment!

A team is validating a DGX BasePOD deployment. Using cmsh, they run a command to check GPU health across all nodes. What indicates that the system is ready for AI workloads?

- A. All GPUs report Status\_Health = OK and Health = OK for each device.
- B. Only the head node's GPUs need to be healthy.
- C. At least half of the GPUs report Status\_Health = OK
- D. The command output is ignored if the system powers on without errors.

**Suggested Answer:** A

Currently there are no comments in this discussion, be the first to comment!



A customer is designing an AI Factory for enterprise-scale deployments and wants to ensure redundancy and load balancing for the management and storage networks.

Which feature should be implemented on the Ethernet switches?

- A. Use only one switch for all management and storage traffic.
- B. MLAG for bonded interfaces across redundant switches.
- C. Implement redundant switches with spanning tree protocol.
- D. Disable VLANs and use unmanaged switches.

**Suggested Answer:** *B*

Currently there are no comments in this discussion, be the first to comment!

A team is installing the NVIDIA Run:ai control plane on a Kubernetes cluster.  
Which two options are most critical to validate before proceeding?

- A. All cluster nodes have NVIDIA GPUs installed.
- B. Helm is installed on the installer machine.
- C. Ensure Kubernetes is running on the cluster.
- D. NTP is disabled to simplify time synchronization.

**Suggested Answer:** *BC*

Currently there are no comments in this discussion, be the first to comment!

A leaf switch shows "FW Version Mismatch" alerts for transceivers after cluster expansion.

Which tool validates transceiver firmware against expected versions?

- A. ethtool
- B. iblinkinfo
- C. mlxconfig
- D. flint

**Suggested Answer:** *D*

Currently there are no comments in this discussion, be the first to comment!

During a 48-hour NeMo question-answering model burn-in test, GPU memory errors occur when processing large datasets. Which configuration strategy prevents Out-of-Memory (OOM) errors while maintaining processing efficiency?

- A. Increase files\_per\_partition to 1000 for larger batch processing.
- B. Switch from FP16 to FP32 precision for numerical stability.
- C. Disable add\_filename for Parquet files to reduce metadata.
- D. Set blocksize= "1GB" for data loading and enable RMM asynchronous allocation.

**Suggested Answer:** *D*

Currently there are no comments in this discussion, be the first to comment!

A system administrator noticed a failure on a DGX H100 server. After a reboot, only the BMC is available.  
What could be the reason for this behavior?

- A. There are more than two failed power supplies.
- B. A boot disk has failed.
- C. The network card has no link / connection.
- D. Multiple GPUs have failed.

**Suggested Answer:** A

Currently there are no comments in this discussion, be the first to comment!

You are evaluating the integration of NVIDIA BlueField DPUs into your data center's storage architecture to optimize AI workloads. The storage solution chosen has incorporated BlueField DPUs to enhance performance and efficiency.

Which of the following benefits directly results from this integration?

- A. Elimination of latency issues in data processing tasks.
- B. Enhanced I/O performance with NVMe storage access speeds.
- C. Reduced CPU load by offloading data processing tasks to DPUs.
- D. Unlimited scalability by adding more DPUs without architectural changes.

**Suggested Answer:** C

Currently there are no comments in this discussion, be the first to comment!

After running a 24-hour stress test on a DGX node, the administrator should verify which two key metrics to ensure system stability?

- A. Total energy consumption and NVLink bandwidth.
- B. SSD write endurance and RAM capacity.
- C. Average CPU usage  $\geq 80\%$  and Docker container uptime.
- D. No thermal throttling events and consistent GPU utilization  $\geq 95\%$  throughout the test.

**Suggested Answer:** *D*

Currently there are no comments in this discussion, be the first to comment!

After upgrading to HPL-AI 2.0 on a DGX A100 cluster, a 2x performance gain is observed. Which optimization is primarily responsible for this improvement?

- A. Reduction of problem size (N) to accelerate computation.
- B. Doubling of GPU clock speeds through firmware updates and relevant configuration.
- C. Automatic NVLink bandwidth doubling via driver updates.
- D. MPI-aware GPU communication that reduces CPU bottlenecks and GPU idle time.

**Suggested Answer:** *D*

Currently there are no comments in this discussion, be the first to comment!



For a 48-hour NCCL burn-in test, which parameters ensure sustained fabric stress while detecting silent data corruption?

- A. `all_reduce_perf -b 8G -e 32G -z 1 -G 1000`
- B. `broadcast_perf -b 4G -e 16G -w 100`
- C. `all_reduce_perf -b 8G -e 32G -c 1000 -z 1 -G 1000`
- D. `reduce_scatter_perf -f 2 -g 8`

**Suggested Answer:** *C*

Currently there are no comments in this discussion, be the first to comment!

An engineer needs to validate NVLink Switch functionality on a DGX H100 system with 8 GPUs.

Which NCCL command verifies intra-node NVLink bandwidth?

- A. `all_reduce_perf -b 8 -e 16G -f2 -g 4` with `NCCL_TESTS_SPLIT= "MOD 2"`
- B. `all_reduce_perf -b 8 -e 16G -f2 -g 8` with `NCCL_TESTS_SPLIT= "OR 0x7"`
- C. `broadcast_perf -b 8 -e 16G -f2 -g 8` without split configuration
- D. `all_reduce_perf -b 8 -e 16G -f2 -g 1` repeated 8 times

**Suggested Answer:** *B*

Currently there are no comments in this discussion, be the first to comment!

You are leading a project to enhance the energy efficiency of a data center that heavily relies on AI workloads. NVIDIA suggests moving beyond traditional metrics like Power Usage Effectiveness (PUE) to better capture the efficiency of modern data centers. Which strategy should you prioritize to develop more accurate energy-efficiency metrics?

- A. Focus on integrating kilowatt-hours into existing metrics to better reflect the actual energy used for productive work.
- B. Use watts-used as the primary measure of efficiency, as it accurately reflects the power input at any given time.
- C. Use Power Usage Effectiveness as the primary metric while supplementing it with additional measures of useful work done per unit of energy.
- D. Develop benchmarks tailored to specific workloads, such as MLPerf for AI applications, to better understand energy use in real-world scenarios.

**Suggested Answer:** *D*

Currently there are no comments in this discussion, be the first to comment!

A system administrator needs to configure a BlueField DPU and enable RShim on the baseboard management controller (BMC). Which command should be executed?

- A. `systemctl enable bmc-rshim.service`
- B. `ipmitool raw 0x32 0x6a 1`
- C. `systemctl restart rshim`
- D. `scp <path_to_bfb> root@<bmc_ip>:/dev/rshim0/boot</bmc_ip></path_to_bfb>`

**Suggested Answer:** *B*

Currently there are no comments in this discussion, be the first to comment!

An administrator installs NVIDIA GPU drivers on a DGX H100 system with UEFI Secure Boot enabled. After reboot, the drivers fail to load. What is the first action to resolve this issue?

- A. Reinstall drivers using `apt-get install nvidia-driver-550` without rebooting.
- B. Delete `/etc/X11/xorg.conf` to force driver reconfiguration.
- C. Enroll the Machine Owner Key (MOK) during system reboot and enter the recorded password.
- D. Disable Secure Boot permanently in BIOS/UEFI settings.

**Suggested Answer:** C

Currently there are no comments in this discussion, be the first to comment!

A system administrator receives an alert about a potential hardware fault on an NVIDIA DGX A100. The GPU performance seems degraded, and the system fans are operating loudly.

What step should be recommended to identify and troubleshoot the hardware fault?

- A. Run a deep learning workload to stress test the GPUs and check whether the issue persists.
- B. Check the NVIDIA System Management Interface (nvidia-smi) for GPU status and temperatures.
- C. Increase the fan speed to maximum and check whether the performance improves.
- D. Power drain then restart the DGX and check if the performance degradation resolves.

**Suggested Answer:** *B*

Currently there are no comments in this discussion, be the first to comment!

For an NVIDIA Enterprise AI Factory with 256 GPUs, which storage solution characteristic is most critical to validate during scaling tests?

- A. RAID rebuild times under disk failure.
- B. Consistent per-node throughput  $\geq 8$  GiB/s.
- C. Single-node write performance during idle clusters.
- D. Maximum 4K random read IOPS exceeding 1 million.

**Suggested Answer:** *B*

Currently there are no comments in this discussion, be the first to comment!

ClusterKit's NCCL bandwidth test shows 350 GB/s on a 400G InfiniBand fabric.  
How should this result be interpreted?

- A. Inconclusive; rerun with `--stress=cpu` to validate.
- B. Optimal performance, indicating healthy fabric and GPUDirect RDMA.
- C. Suboptimal performance; requires FEC tuning to reach 380+ GB/s.
- D. Critical failure; expected is  $\geq 390$  GB/s for HDR InfiniBand.

**Suggested Answer:** *B*

Currently there are no comments in this discussion, be the first to comment!



A systems administrator is preparing a new DGX server for deployment.

What is the most secure approach to configuring the BMC port during initial setup?

- A. Enable remote access to the BMC over the internet using the default admin credentials for initial troubleshooting.
- B. Leave the BMC port disconnected until after the operating system is fully configured and in production.
- C. Connect the BMC port to a dedicated and firewalled network and change the default admin credentials.
- D. Connect the BMC port directly to the production network and retain default admin credentials for convenience.

**Suggested Answer:** C

Currently there are no comments in this discussion, be the first to comment!

You are following the official steps to install the NVIDIA Container Toolkit using a package manager on Ubuntu. After importing the NVIDIA package repository and GPG key, what is the next action?

- A. Download the CUDA toolkit installer from NVIDIA's official website.
- B. Install the nvidia-container-toolkit package using your package manager.
- C. Reboot the host system to apply the repository changes and proceed.
- D. Format the disk to clear any existing NVIDIA-related dependencies first.

**Suggested Answer:** *B*

Currently there are no comments in this discussion, be the first to comment!

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.

**Suggested Answer:** A

Currently there are no comments in this discussion, be the first to comment!

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 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.

**Suggested Answer:** *CD*

Currently there are no comments in this discussion, be the first to comment!

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.

**Suggested Answer:** *C*

Currently there are no comments in this discussion, be the first to comment!

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.

**Suggested Answer:** *D*

Currently there are no comments in this discussion, be the first to comment!

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.

**Suggested Answer:** C

Currently there are no comments in this discussion, be the first to comment!

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

**Suggested Answer:** *D*

Currently there are no comments in this discussion, be the first to comment!



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.

**Suggested Answer:** A

Currently there are no comments in this discussion, be the first to comment!

Refer to the exhibit.



What is the maximum bandwidth supported by Direct Attach Copper (DAC) LinkX cables?

- A. 100 Gbps
- B. 200 Gbps
- C. 50 Gbps
- D. 400 Gbps

**Suggested Answer:** *D*

Currently there are no comments in this discussion, be the first to comment!

To verify correct installation of a DGX system, which command should a system administrator use to collect system configuration and diagnostic information?

- A. nvsm get health
- B. lspci | grep DPU
- C. nvsm stress-test
- D. nvsm show health

**Suggested Answer:** *D*

Currently there are no comments in this discussion, be the first to comment!

You are standing up a NVIDIA DGX system for enterprise production. Stakeholder teams require system reliability, performance consistency under load, and proper escalation processes before release. A recent system in another cluster experienced intermittent GPU failures attributed to missed early-stage validation.

Which deployment and validation sequence best addresses production readiness and mitigates the risk of avoidable downtime or performance loss?

- A. Complete hardware and cabling, power on the system, update firmware and drivers, run full hardware health checks and stress diagnostics (using NVSM), verify all GPU and system sensor logs, and validate GPU accessibility.
- B. Power on the system, install all AI frameworks, configure the CUDA and library stack, set up user environments, then plan stress tests and diagnostics as part of ongoing routine operations.
- C. Update network topology, assign static IPs and DNS entries, register the system with NVIDIA, then conduct basic OS-level checks and enable user access after log-in testing is successful.
- D. Install latest OS images and drivers, confirm OS and container functionality, invite users for a monitored production trial, and collect workload feedback to plan any further diagnostics or updates.

**Suggested Answer:** A

Currently there are no comments in this discussion, be the first to comment!

Which of the following steps are essential components of a recommended DGX cluster installation procedure? (Choose two.)

- A. Configure networking by validating all interfaces on each node – ensuring proper InfiniBand/Ethernet connectivity – prior to installing cluster software.
- B. Complete application containerization, run distributed jobs, validate node health or storage availability.
- C. Group nodes by function during initial setup and assign them to relevant categories in the cluster management tool.
- D. Install Slurm on the head node and then configure the compute nodes' default OS images.

**Suggested Answer:** AC

Currently there are no comments in this discussion, be the first to comment!

You are expanding a DGX-based deep learning cluster to train on large, high-resolution images that cannot fit into local cache. Multiple nodes will access this data concurrently and require high performance.

Which storage and networking solution best meets these requirements?

- A. Implement a standard NFS server on a 10GbE network since the cluster can access the export and job performance will not be impacted.
- B. Increase the SSD RAID-0 local cache size in each node so it can absorb most training data, making network storage type and speed less important for performance.
- C. Recommend general-purpose object storage for all training data as it's optimized for deep learning workloads and distributed data access at any scale
- D. Deploy a high-performance parallel filesystem across InfiniBand or 40/100GbE, ensuring at least 3GB/s per node and scalable aggregate bandwidth for all cluster workloads.

**Suggested Answer:** *D*

Currently there are no comments in this discussion, be the first to comment!

What is the purpose of using NCCL in verifying east/west fabric in an NVIDIA AI Factory? (Choose two.)

- A. To measure bandwidth between GPUs.
- B. To measure the latency between GPUs.
- C. To measure the power consumption of GPUs.
- D. To measure the storage network performance.

**Suggested Answer:** *AB*

Currently there are no comments in this discussion, be the first to comment!

While trying to bring-up the fabric, the administrator identifies one nonoperational InfiniBand server:

CA 'mix5\_1'

```
CA type: MT4115
Number of ports: 2
Firmware version: 10.20.1010
Hardware version: 0
Node GUID: 0x0002c90300002f7b
Port 1:
 State: Active
 Physical state: LinkUp
 Rate: 100
 Base lid: 0
 LMC: 0
 SM lid: 0
 Capability mask: 0x0251086a
 Port GUID: 0x0002c90300002f79
 Link layer: Ethernet
```

What is the cause based on the output?

- A. The HCA port is not part of the InfiniBand fabric.
- B. The HCA port is not assigned a LID number.
- C. The HCA port is faulty.
- D. OpenSM is not running on the fabric.

**Suggested Answer:** A

Currently there are no comments in this discussion, be the first to comment!



ClusterKit's NCCL bandwidth test shows 350 GB/s on a 400G InfiniBand fabric.  
How should this result be interpreted?

- A. Optimal performance, indicating healthy fabric and GPUDirect RDMA
- B. Critical failure; expected is  $\geq 390$  GB/s for HDR InfiniBand
- C. Suboptimal performance; requires FEC tuning to reach 380+ GB/s.
- D. Inconclusive; rerun with `--stress=cpu` to validate.

**Suggested Answer:** A

Currently there are no comments in this discussion, be the first to comment!

You are tasked with setting up High Availability (HA) for NVIDIA Base Command Manager (BCM) in a new GPU cluster. The cluster consists of a primary and secondary head node, as well as several compute nodes. The requirements are: automatic failover of BCM services, minimal disruption to workloads, and proper cluster health monitoring during and after installation.

During your BCM HA installation and configuration process, which two of the following actions are mandatory for ensuring a robust and verified HA cluster configuration? (Choose two.)

- A. After configuration is complete, simulate a failover by stopping BCM services on the active head node to verify that all services are running on the secondary node with no interruption.
- B. During configuration, explicitly synchronize both the configuration and state data directories from the primary to the secondary head node to ensure consistency.
- C. Compute nodes must be powered on and performing work to initiate the synchronization of the head nodes.
- D. Configure both head nodes to use independent static IP addresses for BCM services instead of relying on a shared virtual IP address.
- E. Assign a floating Virtual IP address that can automatically migrate between the primary and secondary head nodes during failover.

**Suggested Answer:** *AE*

Currently there are no comments in this discussion, be the first to comment!

After Spectrum-X fabric deployment, NCCL tests show intermittent latency spikes.  
Which network condition most severely impacts East-West bandwidth?

- A. Multiple transceiver firmware mismatches.
- B. 400G port utilization at 70% on several nodes during tests.
- C. Packet loss >0.001% causing NCCL pipeline stalls.
- D. Jitter below 5us with consistent latency.

**Suggested Answer:** *C*

Currently there are no comments in this discussion, be the first to comment!

What is the primary purpose of performing a NeMo Burn-in on a new AI infrastructure?

- A. To stress test the hardware and software stack with representative NeMo workloads, ensuring reliability.
- B. To benchmark production training speed and ensure all GPUs are running at identical clock speeds.
- C. To tune NeMo model hyperparameters for maximum accuracy on user datasets during cluster deployment.

**Suggested Answer:** A

Currently there are no comments in this discussion, be the first to comment!

Which of the following tests should be used to check for the lowest possible latency between two nodes in a fabric?

- A. ib\_read\_lat
- B. ib\_write\_lat
- C. ib\_read\_bw
- D. iperf

**Suggested Answer:** *B*

Currently there are no comments in this discussion, be the first to comment!

An engineer is reimaging a DGX system in a large cluster.

Which method ensures the most efficient and secure remote installation without physical access?

- A. Create a USB drive with the ISO and manually boot from it on the DGX system.
- B. Skip ISO verification and directly flash the OS to the disk via SSH.
- C. Use apt-get to upgrade the OS without rebooting the system.
- D. Build a software image on Base Command Manager and then reimage.

**Suggested Answer:** *D*

Currently there are no comments in this discussion, be the first to comment!