



CDW Documentation

nvidia_bcm-installation

[Back to NVIDIA page](#)

[Back to BCM page](#)

[NVIDIA Base Command Manager Installation Manual](#)

Topic: Installation Overview and Graphical Installer

What This Unit Covers

- The main BCM installation approaches
- Why the graphical installer is the preferred path for most standard deployments
- What information should be prepared before installation begins
- What each major installer screen is doing
- Which installation choices have the biggest operational impact later
- What happens immediately after the installer finishes

What BCM Is Doing During Installation

- NVIDIA Base Command Manager is cluster management software used to streamline cluster provisioning, workload management, and infrastructure monitoring.
- During installation, the BCM installer deploys both the selected Linux distribution and the BCM software on the head node.
- At the same time, BCM creates a default disk image for compute nodes that is derived from that same Linux distribution.
- In practice, this means the installer is not just putting an application on the system. It is laying down the base operating environment for the cluster manager and preparing the foundation that regular nodes will later use for provisioning.
- This is one reason installation planning matters so much. Early installer choices influence how the head node is built, how nodes boot, how networks are defined, and how the cluster is brought up afterward.

BCM Installation Methods

NVIDIA documents more than one way to install BCM, but they are not all equal in simplicity or risk.

1. Bare-Metal Installation

- This is the standard installation flow where the target head node is booted directly from BCM installation media.
- The installer lays down the operating system and BCM onto the head node.
- This is the most straightforward option for a fresh deployment because the starting point is controlled and predictable.

2. Graphical Installer

- The graphical installer is the recommended method for a standard bare-metal installation.
- It provides a wizard-based workflow that walks the administrator through the required configuration in sequence.
- The graphical installer is generally preferred because it is easier to validate choices screen by screen and easier to use when reviewing hardware, networks, interfaces, and partitioning before deployment begins.

3. Text Installer

- The text installer is a minimal ncurses-based version of the GUI installer.
- It is useful when a graphical environment is not practical, but the graphical installer remains the recommended option in standard cases.

4. Add-On Installer

- NVIDIA also documents an add-on installation approach using Ansible.
- This method installs BCM onto a machine that is already running a supported Linux distribution.
- The add-on approach is not recommended for inexperienced cluster administrators because an existing OS can introduce conflicts, extra repository requirements, dependency issues, deprecated package issues, and custom-configuration troubleshooting. * In other words, add-on installation gives flexibility, but it also reduces the predictability that makes bare-metal installs cleaner.

Core Installation Principle

- The installer is essentially a guided configuration and deployment workflow.
- It gathers information about the cluster, the networks, the node naming model, the storage layout, and the management interfaces.
- Those values are then used to build the head node and establish the initial cluster design.
- This is why the installer should not be treated like a simple “next, next, finish” wizard.
- Each screen is capturing deployment assumptions that affect provisioning, connectivity, manageability, and later expansion.

Why Planning Comes First

- NVIDIA’s deployment guidance emphasizes that physical installation and network switch configuration should be completed before BCM deployment begins.
- NVIDIA also states that information about the intended deployment should be recorded in a site survey before installation.
- The site survey is important because it provides the installer inputs for items such as hostnames, IP ranges, DNS, time servers, topology choices, BMC details, and interface assignments.
- A good way to think about it is this: the installer is only as accurate as the design information you prepared ahead of time.

Before You Start Installation

- Confirm that the hardware is physically installed and ready.
- Confirm that the switch configuration and required network preparation are already done.
- Record deployment values in a site survey before booting the installer.
- Make sure the BIOS of the target head node is configured in UEFI mode.
- Make sure the boot order points to the BCM installation media.
- If using DGX-based guidance, verify expected network port designations and interface expectations ahead of time.
- For the most common Type 1 topology, NVIDIA documents at least two Ethernet NICs on the head node.
- For storage, NVIDIA deployment guidance recommends using a redundant storage device for the head node OS, such as hardware or software RAID, rather than relying on a single non-redundant device.

Boot Media and Starting the Installer

- The BCM installer can be started from installation media such as DVD or USB.
- In vendor deployment workflows, the media may also be mounted through BMC virtual media, depending on platform support.
- When the system boots from the BCM media, the ISO boot menu appears first. * The boot menu includes a countdown that will boot from the hard drive unless interrupted.
- To install BCM, you must interrupt that default behavior and explicitly choose **Start Base Command Manager Graphical Installer**.
- The text installer is also available from the boot menu, but the graphical path is the preferred one for this unit.
- The boot menu also allows kernel parameters to be passed to the installer before launch.

Important Advanced Note: Remote Installation

- NVIDIA documents a custom kernel parameter named **netconf**.
- When configured, netconf can set login and network values for the installer and launch SSH and VNC services as the installer welcome screen starts.
- This allows the cluster administrator to perform a remote installation instead of remaining physically at the console.
- This applies to both the GUI and text installer paths.
- This is not always needed in basic installs, but it is an important capability to know exists.

Graphical Installer Walkthrough

Step 1: Start Installation

- After choosing the graphical installer from the boot menu, the installer splash screen appears.
- The next action is to select **Start installation**.
- At this point, the installer environment is loaded and the guided configuration process begins.

Step 2: EULA Acceptance

- The installer presents more than one agreement.
- You must accept the NVIDIA EULA.
- You must also accept the EULA for the base operating system included in the ISO.
- This matters because BCM installation is deploying both NVIDIA software and the selected Linux distribution onto the head node.

Step 3: Kernel Modules

- The Kernel Modules screen shows the modules recommended for loading based on a hardware probe.
- In normal cases, the defaults are usually sufficient. * However, this screen is more capable than it may first appear.
- Modules can be reordered, removed, added manually, or blacklisted. * Module parameters can also be supplied.
- This screen becomes important if expected storage, NICs, or other hardware are not being recognized correctly.
- In most installations, the safest path is to leave defaults alone unless there is a clear detection issue or NVIDIA instructs otherwise.

Step 4: Hardware Information

- The Hardware Info screen shows what the installer detected after the hardware probe and kernel module loading.
- This is a verification checkpoint, not just an informational screen.
- You should confirm that expected storage devices are visible.
- You should confirm that expected network interfaces are visible.
- NVIDIA deployment guidance specifically calls out validating that the expected Ethernet-mode interfaces are detected and that storage devices look consistent with the known hardware.
- If expected hardware is missing here, the correct move is not to continue blindly.
- The right move is to go back, revisit kernel modules, and correct detection before proceeding.

Step 5: Installation Source

- The Installation Source screen identifies the available installation media and requires the administrator to choose the correct source.
- NVIDIA also documents an optional media integrity check here.
- If selected, the installer verifies the media before continuing.
- This integrity step is worth paying attention to because corrupted installation media can create subtle issues that may not be obvious until later.
- NVIDIA also notes in the installation manual that copied USB media should have its checksum validated so corruption does not cause hard-to-trace downstream problems.

Step 6: General Cluster Settings

The General Cluster Settings screen captures foundational cluster information.

- NVIDIA documents the following items here:
 - Cluster name
 - Administrator email
 - Time zone
 - Time servers
 - Nameservers
- Search domains Environment modules configuration
- Time settings matter more than they may seem at first glance.
- NVIDIA explicitly notes that time servers are recommended to avoid problems caused by time discrepancies between nodes.
- Accurate nameserver and domain information also matters because it supports predictable name resolution and cleaner post-install management.
- NVIDIA deployment guidance notes that some of these values can be changed later, but it is still best to enter them correctly during installation.

Step 7: Workload Manager Selection

- The Workload Manager screen allows a supported workload manager to be selected during installation.
- NVIDIA states that a workload management system is highly recommended for running multiple compute jobs on a cluster.
- If a workload manager is not selected during installation, it can be installed later after the base cluster deployment is complete.
- NVIDIA also notes that the head node can optionally be selected to run jobs as an additional compute resource in smaller clusters, if it has resources to spare.
- The practical takeaway is that workload management is important, but it is not a blocker for getting the base BCM environment installed.

Step 8: Network Topology Selection

This is one of the most important screens in the installer because later network definitions are based on this choice.

NVIDIA documents three supported network topologies.

Type 1 Topology

- Type 1 is the default and most common topology.
- The regular nodes sit on a private internal network.
- The head node connects both to that internal cluster network and to an external network.
- Traffic from regular nodes to outside networks is routed and forwarded through the head node.
- In this model, the head node provides DHCP and PXE services to an isolated worker-node network during provisioning.
- This keeps cluster traffic segregated and is one reason Type 1 is considered the simplest and most common deployment pattern.
- A limitation is that access to worker nodes from outside the cluster generally requires routing or proxying.

Type 2 Topology

- In Type 2, nodes are connected through a router to a public network.
- Regular nodes do not need to go through the head node to reach outside networks.
- There is no externalnet in this topology.
- Because the cluster may be sharing or touching an existing broader network, NVIDIA warns that care should be taken to avoid DHCP conflicts between the head node and any existing DHCP service.
- Type 2 is useful when worker nodes must remain directly reachable on the broader data plane, such as in service-hosting scenarios.

Type 3 Topology

- In Type 3, the head node and regular nodes are on different routed networks.
- The compute-node side is on **internalnet** by default.
- The head-node side is on **managementnet** by default.
- Communication between head node and regular nodes crosses Layer 3 boundaries.
- Because DHCP is a Layer 2 mechanism, DHCP/PXE traffic is not directly supported across that routed separation.
- NVIDIA therefore notes that a DHCP relay agent is typically required in this topology.
- Type 3 can also require relay considerations in some Type 2 scenarios where nodes are spread across multiple subnets.
- This is a good example of why topology is not just an installer preference. It affects boot mechanics, routing expectations, and network dependencies.

Step 9: Head Node Settings

The Head Node Settings screen is where the administrator sets:

- Hostname
- Administrator or root password
- Hardware manufacturer

This is simple on the surface, but it is still a meaningful checkpoint because the hostname becomes part of the cluster's management identity and the manufacturer selection can influence platform-specific behavior or defaults.

Step 10: Compute Node Settings

The Compute Node Settings screen defines the initial model for the regular nodes.

NVIDIA documents configuration of:

- Number of racks
- Number of nodes
- Node naming format
- Hardware manufacturer

The naming format includes:

- Base node name
- Start number
- Number of digits

By default, this produces names such as node001, node002, and so on. This matters because consistent naming pays off later in provisioning, monitoring, troubleshooting, and role assignment.

Step 11: BMC Configuration

The BMC Configuration screen handles out-of-band management for platforms using compatible management controllers.

NVIDIA documents support for BMCs compatible with:

- IPMI
- iDRAC
- iLO
- CIMC
- Redfish

These BMCs can be configured for head nodes and compute nodes. If BMC use is enabled, BMC-related network options appear.

NVIDIA notes that a new Layer 3 subnet can be created specifically for BMC interfaces. NVIDIA also notes that a BMC interface can share a physical interface with another network, but a dedicated physical BMC interface is recommended because shared arrangements can create problems during early BIOS checks.

Another important note is that if a BMC is configured, the BMC password is set to a random value. Operationally, this means BMC planning is not just about power control. It is about secure out-of-band manageability, cleaner troubleshooting, and better hardware visibility.

Step 12: Network Configuration

The Networks screen displays predefined networks based on the earlier topology and BMC choices.

NVIDIA documents that:

- Type 1 defines an external network and an internal network
- Type 2 defines an internal network but no external network
- Type 3 defines an internal network and a management network
- Additional BMC networks may also appear depending on configuration
- Additional custom networks can be added here.
- The installer validates the network settings before allowing you to continue.

NVIDIA makes an important distinction here:

validation confirms that settings are syntactically or logically acceptable, but that does not guarantee they are the intended values for the real environment. In other words, a valid network configuration can still be the wrong network configuration. This is why matching the site survey and coordinating with the network team matters.

Step 13: Head Node Interfaces

- The Head Node Interfaces screen lets you review the proposed head node interface assignments and IP settings.
- Interfaces can be created or removed. * Proposed values can be changed using dropdown selections.
- Network interfaces can even be swapped if the detected assignment is not what you want.
- If a BMC network is being shared with a regular network, an alias interface may appear.
- This screen is where design meets real hardware naming, so it is worth slowing down and verifying carefully.

Step 14: Compute Node Interfaces

- The Compute Node Interfaces screen lets you review the proposed compute-node interface configuration.
- NVIDIA documents that the **BOOTIF** interface is used by a node to pick up its image through node provisioning.
- By default, BCM assigns IP addresses sequentially, starting after the base network address.
- The IP offset setting changes where that sequential assignment begins.
- This is operationally useful because it allows address space to be reserved at the low end of a subnet.
- NVIDIA deployment guidance gives practical reasons for doing this, such as reserving addresses for gateways, VRRP, or separating BOOTIF and IPMI ranges so they do not overlap confusingly on the same network.
- The main point is that the offset is not cosmetic. It is there to support cleaner address planning.

Step 15: Disk Layout

- The Disk Layout screen is where the administrator selects the head node drive on which BCM will be installed.
- This is a major decision because it determines the physical target of the cluster manager installation.
- NVIDIA deployment guidance recommends choosing a redundant storage device for the installation target.
- For production-minded setups, this is a strong reliability consideration, not just a preference.

Step 16: Disk Layout Settings

After selecting the target disk, the installer moves to Disk Layout Settings.

This screen is used to define partitioning layouts for both the head node and compute nodes.

NVIDIA documents separate options for:

- Head node disk layout
- Compute nodes disk layout
- Default behavior depends on drive size and node type.

NVIDIA documents that the head node defaults to one big partition when the drive is smaller than about 500 GB, and to several partitions when the drive is about 500 GB or larger. Compute nodes default to several partitions using the standard default layout.

NVIDIA also documents additional layout possibilities including RAID, failover, STIG-compliant, and LUKS-based schemes. Custom layouts can be imported or edited through the layout XML configuration. This means the disk layout stage is not just “pick a disk.” It is where the storage design for head and regular nodes becomes explicit.

Step 17: Additional Software

- The Additional Software screen shows software that can be added if it was bundled into the installer ISO when the ISO was generated.
- NVIDIA notes that CUDA is one possible example of additional software that may appear here.
- NVIDIA also states that if NVIDIA AI Enterprise is going to be used, these additional software options must not be selected.
- This is an easy screen to overlook, but it can matter depending on how the ISO was built and what downstream platform expectations exist.

Step 18: Summary

The Summary screen provides a consolidated view of installation settings chosen in earlier steps. This is your last meaningful checkpoint before deployment starts. NVIDIA explicitly notes that you can still go back and correct values at this stage. This screen should be treated as a deliberate review point, especially for topology, interfaces, addressing, BMC choices, and disk layout.

Step 19: Deployment Execution

- Once installation begins, the Deployment screen shows progress.
- NVIDIA notes that you can no longer navigate back to previous screens after deployment has started.
- The install log can be viewed in detail from this screen.
- The system can be rebooted manually after deployment, or configured to reboot automatically when complete.
- During reboot, you may need to remove installation media or adjust BIOS boot order so the system boots from the newly installed disk instead of looping back into installer media.

Step 20: Completion and First Login

- After reboot, the head node presents a login prompt.
- NVIDIA states that the administrator can log in as root using the password set during installation.
- The installation manual then points administrators toward updating the system with the latest packages and continuing with licensing and configuration.
- In other words, the graphical installer gets the head node built, but it is the start of cluster

bring-up, not the end of cluster administration.

Important Operational Takeaways

- The graphical installer is the preferred and most approachable installation path for a standard BCM deployment.
- BCM installation is doing more than installing an application.
- It is establishing the operating environment for the cluster manager and the initial compute-node image model.
- A site survey should be completed before installation so installer values are driven by planned design rather than guesswork.
- Hardware verification during the installer matters. If expected NICs or storage do not appear, fix that before proceeding.
- Topology selection is one of the most important architectural choices because it changes how networks are predefined and how boot, routing, and access behave.
- BMC design deserves careful attention. A dedicated BMC interface is preferred over a shared one when possible.
- Network validation in the installer is only a sanity check. It does not replace design validation.
- IP offsets are useful for reserved address space and for keeping provisioning and management ranges organized.
- Disk selection and partition design should be treated as deliberate infrastructure decisions, not just installer clicks.
- The Summary screen is the last chance to review choices before they become a live deployment. * After the installer finishes, the work continues with login, updates, licensing, and post-install configuration.

What Actually Matters Most

If you are thinking like an administrator instead of just memorizing screens, the highest-value checks during this installer are:

- Is the hardware detection correct?
- Did I choose the correct network topology for the real environment?
- Do the site survey values match what I entered?
- Are the correct interfaces mapped to the correct networks?
- Is the BMC design clean and intentional?
- Am I installing onto the correct redundant storage target?

If those are right, the rest of the installer tends to go much more smoothly.

One-Sentence Summary The BCM graphical installer is a guided bare-metal deployment workflow that installs the base OS and BCM on the head node, prepares the cluster's initial network and provisioning design, and sets the foundation for the rest of the cluster bring-up process.

[Back to NVIDIA page](#)

[Back to BCM page](#)