

HPE DEPLOYMENT AND INTEGRATION SERVICES FOR NVIDIA GPU CLOUD

Advisory and Professional Services

SERVICE OVERVIEW

Next-generation capabilities in artificial intelligence (AI) and high-performance computing (HPC) are pushing the limits of what is possible to achieve with machines, from cancer research to autonomous driving. NVIDIA GPU Cloud (NGC) is a container registry that delivers GPU-accelerated cloud platform software stacks, designed for deep-learning and scientific-learning AI workloads.

The HPE Deployment and Integration Services for NGC integrates the HPE server infrastructure with NVIDIA® GPU accelerators and container repository to deliver a comprehensive deployment stack to utilize the GPU accelerator capabilities. During the delivery of these services, Hewlett Packard Enterprise provides the installation, configuration, and validation of your new GPU-enabled, NGC-ready HPE infrastructure. This service is delivered both remotely and on-site.

The service is available in three options with choice of remote or on-site delivery.

TABLE 1. Service options

	Core	Cluster	Core-edge
HPE infrastructure	HPE ProLiant DL380, HPE Apollo servers	HPE ProLiant DL380, HPE Apollo servers	HPE Apollo, HPE Edgeline servers
Container platform installation and configuration	✓	✓	✓
Container scheduler installation and configuration	✗	✓	✗
Cross-environment integration	✗	✗	✓

SERVICE BENEFITS

- Augment your team—the service complements your IT resources with globally available Advisory and Professional Services assistance from HPE Pointnext Services
- Time to production—NGC allows you to quickly spin up containers to accelerate your development cycle
- Certified hardware—the service is delivered to complement HPE hardware that is certified as NGC-ready
- Early return on investment—leverage qualified architectures to help you obtain the best from your investment early on
- Lightweight investment—you do not need to put significant effort to get going with both the infrastructure and environment
- Easy bundling—the service can be easily attached to hardware offerings and delivered alongside your new infrastructure
- Operational readiness—the service will validate your new HPE infrastructure is fully commissioned and ready to use

SERVICE FEATURE HIGHLIGHTS

- Service planning
- HPE GPU-enabled infrastructure configuration and integration
- NGC configuration, integration, and testing
- Knowledge transfer

TABLE 2. Service features

Feature	Delivery specifications
Service planning	<p>Service planning is delivered remotely. It is designed to provide the Customer with a detailed plan and schedule for the service delivery of the HPE infrastructure. During service planning, HPE conducts service planning meetings with the Customer's designated stakeholders to communicate the list of service activities and confirm that the predeployment prerequisites have been addressed. During these meetings, HPE will:</p> <ul style="list-style-type: none"> • Work with the Customer to review the service details, deliverables, and estimated timelines • Review Customer-provided documents, which include but are not limited to current network architecture and principles to effectively plan for the service delivery • Identify any additional changes to the existing network required to prepare for the integration of new HPE server infrastructure products and documents them in the delivery plan and schedule • Detail any product licensing and prerequisites that must be fulfilled for HPE to successfully perform the services • Identify any additional information required to allow HPE to deploy resources and begin service delivery • Ascertain possible dates for deployment and integration completion • Create a mutually agreeable plan and schedule, as well as determine the appropriate resources necessary to implement the service <p>HPE will work remotely with the Customer to coordinate the necessary activities, including identifying any service prerequisites, scheduling, and the coordination of services.</p>
HPE GPU-enabled infrastructure configuration and integration	<p>To enable cloud-like agility in workloads, scalability, and simplicity in the utilization of GPU hardware, a container infrastructure is highly recommended. The solution integrates HPE server infrastructure, with NVIDIA's GPU accelerators, drivers, and container repository to deliver a comprehensive stack to utilize the GPU accelerator's capabilities within those containers. Organizations can then deploy containers from the NGC repository as a starting point to build their own solutions. During delivery of this service, the HPE specialist assists the Customer by scheduling remotely delivered collaboration meetings that are aimed at helping ensure the service deployment and integration runs smoothly. Depending on the Customer requirements, HPE works remotely or on-site during service delivery. The initial infrastructure configuration and integration phase includes:</p> <ul style="list-style-type: none"> • Reviewing key elements of the purchased solution and the associated high-level design and existing core network as applicable • Evaluating the recommended network connectivity between the HPE server infrastructure nodes, as well as the larger core network; HPE also advises the Customer on the different connectivity strategies • Reviewing with the Customer the predefined choices of deployment components such as underlying operating system as well as any credentials or connectivity information that maybe required • Assessing server BIOS options, server configuration options such as RAID and storage configurations • Appraising configuration of the base operating system, package management system, and subscription services (where required) • Installation and configuration of accelerator drivers and libraries, container engine and runtime, and plugins • Installation and configuration of container scheduler (Kubernetes), package manager, and plugins (cluster option) • Documenting the configuration for the installed nodes
NGC configuration, integration, and testing	<p>During delivery of this service, the HPE specialist assists the Customer with the configuration and integration of the NGC service and any Customer-specific keys and service subscriptions. HPE may work remotely and/or on-site depending on the Customer requirements. This service includes configuration of the NGC service, deployment of a predefined set of containers, and demonstration and documentation of the container utilizing the GPU accelerator. Steps included in this phase are:</p> <ul style="list-style-type: none"> • Integration of the NGC service • Deployment of repository keys and subscriptions where necessary • Deployment and running of predefined container packages • Documentation of correct GPU utilization by the running container packages
Knowledge transfer	<p>During delivery of this service, the HPE specialist works remotely to present the Customer with final documentation and a knowledge transfer session to the appropriate Customer stakeholders. The activities associated with this service include:</p> <ul style="list-style-type: none"> • Performing a remote or on-site knowledge transfer session to facilitate a transfer of operational control to the Customer, including: <ul style="list-style-type: none"> – Providing an overview of the completed implementation, with the broad steps taken to achieve the implementation together with the functionality of the deployment – Offering a demonstration of the NGC implementation, with the steps required to deploy a sample container from the NGC repository check that it is running • Delivering operational tips specific to the NGC service from NVIDIA such as pointers to user guides, containers, and best practices • Describing examples of expansion possibilities or additional advantages



COVERAGE

This service is available on regular HPE workdays (excluding weekend days and HPE holidays) during country-specific HPE standard business hours.

CUSTOMER RESPONSIBILITIES

The Customer will:

- Assign a primary stakeholder to participate in the service planning meeting and follow-on service activity.
- Ensure that a primary stakeholder or designated staff person is assigned and who, on behalf of the Customer, will grant all approvals, provide information, and confirm that the hardware, firmware, and software needed to deliver this service are available. They will also make sure that software products are properly licensed and otherwise be available to assist HPE in facilitating the delivery of this service. The designated primary contact will be:
 - Responsible for all Customer aspects of the assigned work efforts
 - Authorized to make all decisions relative to the project, including identification and assignment of Customer resources
 - Available and able to interface with the HPE assigned resources on day-to-day issues throughout the project
 - Authorized to sign status reports and approve project changes
 - Able to coordinate all work efforts and meeting schedules
- Make sure that all service prerequisites—including, but not limited to, those identified during service planning—have been met prior to delivery of remote and/or on-site services configuration and integration services.
- Provide suitable virtual private network (VPN) or other suitable connectivity for delivery of the remotely delivered services.
- Allow HPE full and unrestricted network access to all locations where the service is to be performed on-site.
- Be responsible for all current-state and future-state architectures, designs, and integration projects, within the environment.
- Give access to network services such as network time protocol (NTP), domain name system (DNS), default gateways, routes, internet proxies and remediation, if needed.
- Ensure that all products associated with the tasks to be performed by HPE are ordered and available on-site prior to the start of the services and/or the arrival of the HPE specialist.
- Verify that all products associated with the tasks to be performed by HPE are installed, cabled with the necessary networks configured (HPE iLO/management and LAN), and powered on within the data center prior to the start of the service and/or the arrival of the HPE specialist.
- Provide HPE with the current architecture, standards, and detailed design documentation that may include, but is not limited to:
 - Project plans and schedules
 - Network topology diagrams
 - Rack placement diagrams
 - Cable maps (device and end station cable numbers, patch panel designation and port numbers, device port numbers, and VLAN information)
 - IP address maps
 - Network environment administrative and management parameters and variables (hostname, administration, users, authorization, administrative IP address, management passwords, SNMP, NTP, DNS server addresses, DHCP, and logging)
- Make modifications to the existing network including validation of connectivity to all endpoints that are required and identified during the planning stages of this service, prior to HPE performing configuration and integration tasks.
- Advise HPE of any special security, health, and safety matters applicable to the Customer site where the service is to be provided.
- Provide to HPE, on request, any information that we may request about the execution of the service.



- Coordinate all required internal/third-party participation and cooperation.
- Assign or make available experienced subject matter and technical experts, upon request or as needed.
- Provide HPE with the necessary access to the Customer building facilities and computer room facilities, as well as access credentials for logging into all servers and services for the service planning, as required.
- Purchase or provide all hardware, software, licenses, staff, current maintenance contracts, and environments necessary for HPE to deliver the service.
- Be responsible for any notifications to operations and any change control documentation that must be completed.
- Take responsibility for formal delivery of the configured infrastructure with HPE's assistance.
- Provide a suitable work and meeting area commensurate with the number of on-site HPE consultants and Customer subject matter experts assigned to the analysis, including desks, chairs, telephones, and internet/HPE network access through a VPN.

SERVICE LIMITATIONS

- HPE does not modify any configurations of any equipment that is not part of HPE Deployment and Integration Services for NGC.
- Operating system licensing/support contracts are not included within the price of the service, should they be required (such as Red Hat® subscriptions), the Customer should ensure these are procured ahead of the service delivery.
- NGC and NGC content licensing/support contracts are not included within the price of the service, should they be required, the Customer should ensure these are procured.
- Customer acknowledges and agrees that HPE may use resources outside the country of purchase for delivery of these services.
- On-site service assistance is provided at one physical location in the country where the service is sold. The Customer should check with their local HPE authorized representative to find out whether a specific location is eligible for this service.

GENERAL PROVISIONS/OTHER EXCLUSIONS

- To the extent HPE processes personal data on the Customer's behalf in the course of providing services, the HPE Data Privacy and Security Agreement Schedule—HPE Support and Professional Services found at hpe.com/info/Customer-privacy.html shall apply. HPE Deployment and Integration Services for NGC are governed by HPE standard terms for Professional Services as part of these services.
- Our ability to deliver this service is dependent upon the Customer's full and timely cooperation with HPE, as well as the accuracy and completeness of any information and data the Customer provides to HPE.
- Upon receipt of an acceptable order, HPE will contact the Customer within seven business days to organize a service delivery date. Service delivery dates are subject to resource availability and may be scheduled up to 30 days from the order acceptance date. The service eligibility expires at the end of 120 days from the date of purchase if not used. Under no circumstances shall the Customer be entitled to a credit or refund of any unused services.
- Services are either performed remotely or at the Customer's site, depending on the ordered product number. Travel for on-site service delivery occurs during the first morning and last workday of each week (or as jointly determined by both parties).
- HPE assumes that all information provided by the Customer is accurate. HPE collaborates with the Customer to determine acceptable estimates for any information that is not available.
- Service hours are inclusive of on-site and off-site service, and travel time. Any request for flexible work hours must be agreed to in advance by HPE and may impact the calculation of hours expended as detailed here. The on-site/off-site schedule of the HPE specialist is mutually agreed to prior to the commencement of services.
- Any after-hours work must be requested a minimum of one week in advance and must be approved by HPE.



ORDERING INFORMATION

Availability of service features and service levels may vary according to local resources and may be restricted to eligible products and geographic locations. To obtain further information or to order HPE Deployment and Integration services for NGC, contact a local HPE sales representative and reference the following product number(s):

- HU7R1A1—HPE Deployment and Integration SVC for NGC Core—remote delivery
- HU7R2A1—HPE Deployment and Integration SVC for NGC Core—on-site delivery
- HU7R3A1—HPE Deployment and Integration SVC for NGC Cluster—remote delivery
- HU7R4A1—HPE Deployment and Integration SVC for NGC Cluster—on-site delivery
- HU7R5A1—HPE Deployment and Integration SVC for NGC Core-edge—remote delivery
- HU7R6A1—HPE Deployment and Integration SVC for NGC Core-edge—on-site delivery
- H9SX2A1—HPE Big Data Analytics SOW Consulting SVC

LEARN MORE AT

hpe.com/services/ai-data

hpe.com/pointnext

Make the right purchase decision.
Contact our presales specialists.



Chat



Email



Call



HPE support



Get updates