



Storage Partner Considerations for NVIDIA DGX BasePOD and NVIDIA DGX SuperPOD with DGX B200, DGX H200, or DGX H100 Systems

Storage Program Overview

Document History

Version	Date	Authors	Description of Change
01	May 2024	NVIDIA	Initial release
02	June 2024	NVIDIA	Branding update

Table of Contents

- Introduction4
- Program Qualification Criteria5
- Validation and Engagement.....7
 - Validation Process7
 - Continued Engagement Process.....7
- Go-To-Market and Promotion8
 - Go-To-Market Expectations.....8
 - Promotional Expectations8

Introduction

NVIDIA DGX B200, DGX H200, and DGX H100 systems can be deployed in ad hoc clusters or as part of NVIDIA DGX BasePOD or NVIDIA DGX SuperPOD AI infrastructure solutions, depending on the scale and complexity of the AI workloads. Both DGX BasePOD and DGX SuperPOD require high-performance storage solutions that can keep up with the massive data throughput and I/O demands of DGX systems.

NVIDIA works with selected storage partners to qualify, validate, and promote their storage solutions for DGX BasePOD and DGX SuperPOD. This document summarizes the key considerations for storage partners that are interested in applying to the DGX BasePOD or DGX SuperPOD with DGX B200, DGX H200, or DGX H100 systems storage programs.

Program Qualification Criteria

The DGX BasePOD reference architecture allows storage providers to create integrated AI appliance solutions that reduce the complexity of deploying AI infrastructure.

To apply for the DGX BasePOD storage program, prospective storage partner must meet the following minimum criteria:

- Storage partners are any of the following: storage solution providers, OEM partners, or integrated solution partners.
- Storage partner has AI customers and an NVIDIA partner business manager (PBM) assigned to them.
- Storage partner has a business case for DGX BasePOD approved by NVIDIA.
- Partner solution comprises a distributed file system product sold as an appliance. The solution is developed by Partner or by Partner with its providers.
- Storage partner has DGX systems in their lab and sells AI solutions with DGX systems.
- Storage partner's storage appliance provides a rich set of enterprise storage and data management features.

NVIDIA DGX SuperPOD is world-class, turnkey infrastructure designed to tackle the most demanding AI challenges. To achieve this, storage solutions used in DGX SuperPOD must meet the highest levels of performance and reliability.

To participate in the DGX SuperPOD storage program prospective storage partner must meet the following minimum criteria:

- Storage partner is an existing partner in the DGX BasePOD storage program described above.
- Storage partner delivers storage appliances that they develop with their own storage software, integrate, support, and service.
- Storage partner has a business case for DGX SuperPOD approved by NVIDIA.
- Program participation approval obtained by storage solution partner cannot be transferred to or reused by other organizations.
- The storage appliance is a distributed file system product that includes all applicable hardware, software, and firmware. Software-only storage solutions are not eligible.
- Storage partner has dedicated professional services and support teams for DGX SuperPOD deployments.
- Customer demand for integration with storage partner's appliance with DGX SuperPOD.

Validation and Engagement

Validation Process

Storage partners that meet the program qualification criteria must undergo a validation and approval process to ensure that their storage solutions are compatible and performant with (i) DGX BasePOD or (ii) DGX SuperPOD (with DGX B200, DGX H200, or DGX H100) systems. The validation process is as follows:

- For DGX BasePOD, storage partners must deploy a representative NVIDIA Base Command Manager cluster with DGX B200, DGX H200, or DGX H100 systems in their lab to execute NVIDIA-provided tests and submit results to NVIDIA.
- For DGX SuperPOD, storage partners must supply their storage solution, including all necessary hardware and software, to NVIDIA for a minimum of 6 to 12 months. The storage solution will be rigorously tested at NVIDIA using a comprehensive, NVIDIA-internal DGX SuperPOD test suite.
- In all cases, NVIDIA will make the final determination on acceptance of partners into the programs.

Continued Engagement Process

To be part of the storage program, storage partners who successfully complete the validation process must maintain a continued engagement process with NVIDIA to ensure that their storage solutions are up to date and aligned with the DGX BasePOD and DGX SuperPOD with DGX B200, DGX H200, or DGX H100 systems requirements. The continued engagement process is as follows:

- For DGX BasePOD, storage partners must periodically run DGX BasePOD tests to ensure compatibility with every major Base Command Manager release. The NVIDIA DGX team will provide Base Command Manager pre-release software to the storage partner for testing purposes, upon prior written request. The storage partners must also conduct business reviews every 3 to 6 months, joint sales and channel training, and joint professional services and support training with NVIDIA.
- For DGX SuperPOD, storage partners must follow the same continued engagement process as DGX BasePOD, plus they must also integrate their storage solutions with Base Command Manager and commit to deploy their storage solution within NVIDIA Professional Services (NVPS) timelines at each customer location.

Go-To-Market and Promotion

Go-To-Market Expectations

To be part of the NVIDIA DGX BasePOD and NVIDIA DGX SuperPOD with DGX B200, DGX H200, or DGX H100 systems storage programs, storage partners must also meet the minimum go-to-market (GTM) expectations of NVIDIA that are described below.

- For DGX BasePOD, storage partners must collaborate with and develop and sell AI solutions based on DGX BasePOD. The GTM and sales strategy is jointly agreed upon by the storage partner and executed by both parties. Storage partners are responsible for creating and executing a marketing and asset plan for their DGX BasePOD-based solutions.
- For DGX SuperPOD, storage partners must develop DGX SuperPOD AI solutions with NVIDIA. The GTM and sales strategy is primarily defined by NVIDIA.

Promotional Expectations

Storage partners who are validated to offer the DGX BasePOD and DGX SuperPOD with DGX B200, DGX H200, or DGX H100 systems storage programs must also meet the promotional expectations of NVIDIA. The promotional expectations are as follows:

- All outbound materials that include DGX BasePOD, DGX SuperPOD, DGX B200, DGX H200, DGX H100, DGX systems, or references NVIDIA in any way must be submitted to and approved by NVIDIA prior to release or publication.
- For DGX BasePOD, storage partners must co-develop promotional materials as part of their GTM plan highlighting DGX BasePOD and NVIDIA's AI software solutions. These materials must be reviewed by NVIDIA in advance of any use or publication.
- DGX SuperPOD storage partners must co-develop promotional materials highlighting DGX SuperPOD and NVIDIA's AI and HPC software solutions. These materials must be reviewed by NVIDIA in advance of any use or publication.
- Storage partners must leverage NVIDIA-hosted events to showcase their storage solutions for DGX BasePOD and DGX SuperPOD.

Notice

This document is provided for information purposes only and shall not be regarded as a warranty of a certain functionality, condition, or quality of a product. NVIDIA Corporation ("NVIDIA") makes no representations or warranties, expressed or implied, as to the accuracy or completeness of the information contained in this document and assumes no responsibility for any errors contained herein. NVIDIA shall have no liability for the consequences or use of such information or for any infringement of patents or other rights of third parties that may result from its use. This document is not a commitment to develop, release, or deliver any Materials (defined below), code, or functionality.

NVIDIA reserves the right to make corrections, modifications, enhancements, improvements, and any other changes to this document, at any time without notice.

Customer should obtain the latest relevant information before placing orders and should verify that such information is current and complete.

NVIDIA products are sold subject to the NVIDIA standard terms and conditions of sale supplied at the time of order acknowledgement, unless otherwise agreed in an individual sales agreement signed by authorized representatives of NVIDIA and customer ("Terms of Sale"). NVIDIA hereby expressly objects to applying any customer general terms and conditions with regards to the purchase of the NVIDIA product referenced in this document. No contractual obligations are formed either directly or indirectly by this document.

No license, either expressed or implied, is granted under any NVIDIA patent right, copyright, or other NVIDIA intellectual property right under this document. Information published by NVIDIA regarding third-party products or services does not constitute a license from NVIDIA to use such products or services or a warranty or endorsement thereof. Use of such information may require a license from a third party under the patents or other intellectual property rights of the third party, or a license from NVIDIA under the patents or other intellectual property rights of NVIDIA.

Reproduction of information in this document is permissible only if approved in advance by NVIDIA in writing, reproduced without alteration and in full compliance with all applicable export laws and regulations, and accompanied by all associated conditions, limitations, and notices.

THIS DOCUMENT AND ALL NVIDIA MATERIALS, ARE BEING PROVIDED "AS IS." NVIDIA MAKES NO WARRANTIES, EXPRESSED, IMPLIED, STATUTORY, OR OTHERWISE WITH RESPECT TO THE MATERIALS, AND EXPRESSLY DISCLAIMS ALL IMPLIED WARRANTIES OF NONINFRINGEMENT, MERCHANTABILITY, AND FITNESS FOR A PARTICULAR PURPOSE. TO THE EXTENT NOT PROHIBITED BY LAW, IN NO EVENT WILL NVIDIA BE LIABLE FOR ANY DAMAGES, INCLUDING WITHOUT LIMITATION ANY DIRECT, INDIRECT, SPECIAL, INCIDENTAL, PUNITIVE, OR CONSEQUENTIAL DAMAGES, HOWEVER CAUSED AND REGARDLESS OF THE THEORY OF LIABILITY, ARISING OUT OF ANY USE OF THIS DOCUMENT, EVEN IF NVIDIA HAS BEEN ADVISED OF THE POSSIBILITY OF SUCH DAMAGES. Notwithstanding any damages that customer might incur for any reason whatsoever, NVIDIA's aggregate and cumulative liability towards customer for the products described herein shall be limited in accordance with the Terms of Sale for the product.

Copyright

© 2024 NVIDIA Corporation. All rights reserved.

