



NetApp Verified Architecture

NVIDIA DGX SuperPOD with NetApp

NVA design

Ryan Rodine, NetApp
June 2022 | NVA-1167-DESIGN

Abstract

This NetApp Verified Architecture describes the design of the NVIDIA DGX SuperPOD with NetApp® BeeGFS building blocks. This solution is a full-stack data center platform that is validated on a dedicated acceptance cluster at NVIDIA.

TABLE OF CONTENTS

Executive summary	3
Program summary	3
Solution overview	3
Solution technology	3
Use case summary	5
Technology requirements	5
Hardware requirements	5
Power requirements	6
Software requirements	6
Solution verification	6
Conclusion	6
Where to find additional information	6
Version history	6

LIST OF TABLES

Table 1) Hardware requirements	5
Table 2) Software requirements	6

LIST OF FIGURES

Figure 1) A high-level view of the solution.	4
Figure 2) Front view of a default sized SuperPOD.	4
Figure 3) NetApp BeeGFS building block.	5

Executive summary

Although AI enhances consumers' lives and helps organizations in all industries worldwide to innovate and to grow their businesses, it is a disrupter for IT. To support the business, IT departments are scrambling to deploy high-performance computing (HPC) solutions that can meet the extreme demands of AI workloads. As the race to win with AI intensifies, the need for an easy-to-deploy, easy-to-scale, and easy-to-manage solution becomes increasingly urgent.

The NVIDIA DGX SuperPOD makes supercomputing infrastructure easily accessible for your organization and delivers the extreme computational power that you need to solve even the most complex AI problems. To help you deploy at scale today, this NVIDIA and NetApp turnkey solution removes the complexity and guesswork from infrastructure design and delivers a complete, validated solution including best-in-class compute, networking, storage, and software.

Program summary

NVIDIA DGX SuperPOD with NVIDIA Base Command Manager brings together a design-optimized combination of AI computing, network fabric, storage, software, and support. Every implementation is validated on a dedicated acceptance cluster at NVIDIA. The recommended design is replicated beforehand, and a suite of performance results is produced. So, when DGX SuperPOD is deployed on site, it runs exactly as it was intended.

Solution overview

NVIDIA DGX SuperPOD is an AI data center infrastructure platform delivered as a turnkey solution for IT to support the most complex AI workloads facing today's enterprises. It simplifies deployment and management while delivering virtually limitless scalability for performance and capacity. In other words, DGX SuperPOD lets you focus on insights instead of infrastructure.

With NetApp EF600 all-flash arrays at the foundation of your NVIDIA DGX SuperPOD, you get an agile AI solution that scales easily and seamlessly. The flexibility and scalability of the solution enable it to support and adapt to evolving workloads, making it a strong foundation to meet your future storage requirements. Modular storage building blocks give you a granular approach to growth. You can scale seamlessly from terabytes to petabytes. By increasing the number of storage building blocks, you can scale up the performance and capacity of the file system, enabling your solution to manage the most extreme workloads with ease.

Solution technology

- NVIDIA DGX SuperPOD is a validated deployment of 20 x 140 DGX A100 systems with validated externally attached shared storage:
 - Each DGX A100 SuperPOD scalable unit (SU) consists of 20 DGX A100 systems and is capable of 48 AI PFLOPS. It usually contains at least two NetApp BeeGFS building blocks depending on the performance and capacity requirements for a particular installation.
- NetApp BeeGFS building blocks consists of two NetApp EF600 arrays and two x86 servers:
 - With NetApp EF600 all-flash arrays at the foundation of your NVIDIA DGX SuperPOD, you get a reliable storage foundation backed by six 9s of uptime.
 - The file system layer between the NetApp EF600 and the NVIDIA DGX A100 systems is the BeeGFS parallel file system. BeeGFS was created by the Fraunhofer Center for High-Performance Computing in Germany to solve the pain points of legacy parallel file systems. The result is a file system with a modern, user space architecture that is now developed and delivered by ThinkParQ and used by many supercomputing environments.

- NetApp support for BeeGFS places the assistance of NetApp's excellent support organization at your fingertips. You get access to superior support resources, early access to BeeGFS releases, and access to select BeeGFS enterprise features such as quota enforcement and high availability (HA).
- The combination of NVIDIA SuperPOD SUs and NetApp BeeGFS building blocks provides an agile AI solution in which compute or storage scales easily and seamlessly.

Figure 1 shows a high-level view of the NVIDIA DGX SuperPOD with NetApp solution.

Figure 1) A high-level view of the solution.

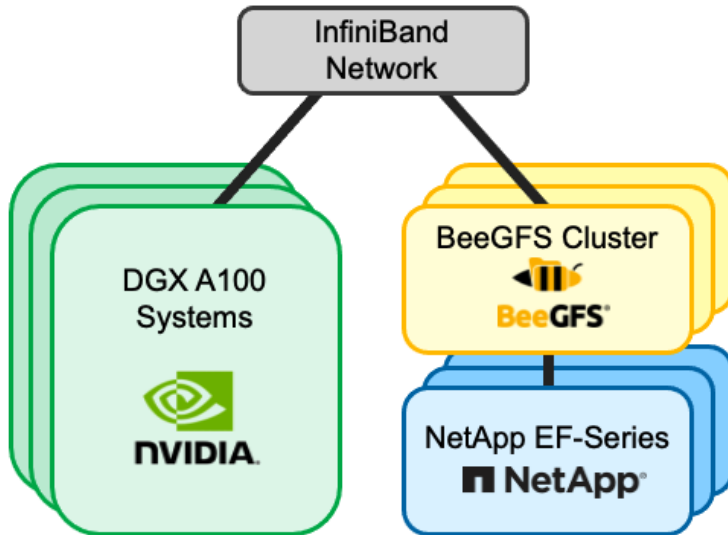


Figure 2 shows the frontal view of a default sized NVIDIA DGX SuperPOD system.

Figure 2) Front view of a default sized SuperPOD.





Figure 3 shows an overview of a NetApp BeeGFS building block.

Figure 3) NetApp BeeGFS building block.



Use case summary

This solution applies to the following use cases:

- Artificial Intelligence (AI) including machine learning (ML), deep learning (DL), natural language processing (NLP), and natural language understanding (NLU)
- Medium to large scale AI training
- Computer vision, speech, audio, and language models
- HPC including applications accelerated by message passing interface (MPI) and other distributed computing techniques
- Application workloads characterized by the following:
 - Reading or writing to files larger than 1GB
 - Reading or writing to the same file by multiple clients (10s, 100s, and 1000s)
- Multiterabyte or multipetabyte datasets
- Environments that need a single storage namespace optimizable for a mix of large and small files

Technology requirements

This section covers the technology requirements for the NVIDIA DGX SuperPOD with NetApp solution.

Hardware requirements

Table 1 lists the hardware components that are required to implement the solution for a single SU. The solution sizing starts with 20 NVIDIA DGX A100 systems and two or three NetApp BeeGFS building blocks.

A single NetApp BeeGFS building block consists of two NetApp EF600 arrays and two x86 servers. You can add additional building blocks as the deployment size increases. For more information, see the [NVIDIA DGX A100 SuperPOD reference architecture](#) and [NVA-1164-DESIGN: BeeGFS on NetApp NVA Design](#).

Table 1) Hardware requirements.

Hardware	Quantity
NVIDIA DGX A100	20
NVIDIA Quantum QM8700 switches	12 leaf, 7 spine
NetApp BeeGFS building blocks	2-3

Power requirements

Power requirements depend on the performance and capacity requirements for a particular installation. A default starting configuration 20-node SuperPOD with the recommended BeeGFS storage blocks require up to approximately 143kW maximum power. See the NVIDIA DGX SuperPOD Data Center Design Reference Guide for a full compute breakdown of thermal and power requirements.

Software requirements

Table 2 lists the software components required to implement the solution. The software components that are used in any particular implementation of the solution might vary based on customer requirements.

Table 2) Software requirements.

Software
NVIDIA DGX software stack
NVIDIA Base Command Manager
ThinkParQ BeeGFS parallel file system

Solution verification

NVIDIA DGX SuperPOD with NetApp was validated on a dedicated acceptance cluster at NVIDIA by using NetApp BeeGFS building blocks. Acceptance criteria was based on a series of application, performance, and stress tests performed by NVIDIA. For more information, see the [NVIDIA DGX SuperPOD: NetApp EF600 and BeeGFS Reference Architecture](#).

Conclusion

NetApp and NVIDIA have a long history of collaboration to deliver a portfolio of AI solutions to market. NVIDIA DGX SuperPOD with the NetApp EF600 all-flash array is a proven, validated solution that you can deploy with confidence. This fully integrated, turnkey architecture takes the risk out of deployment and puts you on the path to winning the race to AI leadership.

Where to find additional information

To learn more about the information that is described in this document, review the following documents and/or websites:

- NVA-1164-DESIGN: BeeGFS on NetApp NVA Design
<https://www.netapp.com/media/71123-nva-1164-design.pdf>
- NVA-1164-DEPLOY: BeeGFS on NetApp NVA Deployment
<https://www.netapp.com/media/71124-nva-1164-deploy.pdf>
- NVIDIA DGX SuperPOD Reference Architecture
<https://www.nvidia.com/en-us/data-center/resources/nvidia-dgx-superpod-reference-architecture/>

Version history

Version	Date	Document version history
Version 1.0	June 2022	Initial release.

Refer to the [Interoperability Matrix Tool \(IMT\)](#) on the NetApp Support site to validate that the exact product and feature versions described in this document are supported for your specific environment. The NetApp IMT defines the product components and versions that can be used to construct configurations that are supported by NetApp. Specific results depend on each customer's installation in accordance with published specifications.

Copyright information

Copyright © 2022 NetApp, Inc. All Rights Reserved. Printed in the U.S. No part of this document covered by copyright may be reproduced in any form or by any means—graphic, electronic, or mechanical, including photocopying, recording, taping, or storage in an electronic retrieval system—without prior written permission of the copyright owner.

Software derived from copyrighted NetApp material is subject to the following license and disclaimer:

THIS SOFTWARE IS PROVIDED BY NETAPP "AS IS" AND WITHOUT ANY EXPRESS OR IMPLIED WARRANTIES, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE, WHICH ARE HEREBY DISCLAIMED. IN NO EVENT SHALL NETAPP BE LIABLE FOR ANY DIRECT, INDIRECT, INCIDENTAL, SPECIAL, EXEMPLARY, OR CONSEQUENTIAL DAMAGES (INCLUDING, BUT NOT LIMITED TO, PROCUREMENT OF SUBSTITUTE GOODS OR SERVICES; LOSS OF USE, DATA, OR PROFITS; OR BUSINESS INTERRUPTION) HOWEVER CAUSED AND ON ANY THEORY OF LIABILITY, WHETHER IN CONTRACT, STRICT LIABILITY, OR TORT (INCLUDING NEGLIGENCE OR OTHERWISE) ARISING IN ANY WAY OUT OF THE USE OF THIS SOFTWARE, EVEN IF ADVISED OF THE POSSIBILITY OF SUCH DAMAGE.

NetApp reserves the right to change any products described herein at any time, and without notice. NetApp assumes no responsibility or liability arising from the use of products described herein, except as expressly agreed to in writing by NetApp. The use or purchase of this product does not convey a license under any patent rights, trademark rights, or any other intellectual property rights of NetApp.

The product described in this manual may be protected by one or more U.S. patents, foreign patents, or pending applications.

Data contained herein pertains to a commercial item (as defined in FAR 2.101) and is proprietary to NetApp, Inc. The U.S. Government has a non-exclusive, non-transferrable, non-sublicensable, worldwide, limited irrevocable license to use the Data only in connection with and in support of the U.S. Government contract under which the Data was delivered. Except as provided herein, the Data may not be used, disclosed, reproduced, modified, performed, or displayed without the prior written approval of NetApp, Inc. United States Government license rights for the Department of Defense are limited to those rights identified in DFARS clause 252.227-7015(b).

Trademark information

NETAPP, the NETAPP logo, and the marks listed at <http://www.netapp.com/TM> are trademarks of NetApp, Inc. Other company and product names may be trademarks of their respective owners.

NVA-1167-DESIGN-0622