



# **EF-Series AI with NVIDIA**

## **NetApp Solutions**

NetApp  
July 31, 2024

# Table of Contents

- NetApp EF-Series AI with NVIDIA ..... 1
  - EF-Series AI with NVIDIA DGX A100 Systems and BeeGFS ..... 1
  - NVA-1156-DESIGN: NetApp EF-Series AI with NVIDIA DGX A100 systems and BeeGFS..... 4
  - NVA-1156-DEPLOY: NetApp EF-Series AI with NVIDIA DGX A100 systems and BeeGFS ..... 4

# NetApp EF-Series AI with NVIDIA

Overview of EF-Series AI converged infrastructure solutions from NetApp and NVIDIA.

## EF-Series AI with NVIDIA DGX A100 Systems and BeeGFS

- [Design Guide](#)
- [Deployment Guide](#)
- [BeeGFS Deployment Guide](#)

### NVIDIA DGX SuperPOD with NetApp - Design Guide



Amine Bennani, David Arnette and Sathish Thyagarajan, NetApp

#### 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 any organization and delivers the extreme computational power needed to solve even the most complex AI problems. To help customers 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 DGX H100 systems and NVIDIA Base Command brings together a design-optimized combination of AI computing, network fabric, storage, software, and support. The BeeGFS on NetApp architecture has been previously validated on a dedicated acceptance cluster at NVIDIA. The latest architecture extends that validation by maintaining the proven design while incorporating support for the latest hardware from NVIDIA.

#### 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 an NVIDIA DGX SuperPOD, customers 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 current and future storage requirements. Modular storage building blocks allow a granular approach to growth and scale seamlessly from terabytes to petabytes. By increasing the number of storage building blocks, customers can scale up the performance and capacity of the file system, enabling the solution to manage the most extreme workloads with ease.

## Solution technology

- NVIDIA DGX SuperPOD with NVIDIA DGX H100 systems leverates DGX H100 systems with validated externally attached shared storage:
  - Each DGX SuperPOD scalable unit (SU) consists of 32 DGX H100 systems and is capable of 640 petaFLOPS of AI performance at FP8 precision. It usually contains at least two NetApp BeeGFS building blocks depending on the performance and capacity requirements for a particular installation.

### *A high-level view of the solution*

image::EF\_SuperPOD\_HighLevel.png[]

- NetApp BeeGFS building blocks consists of two NetApp EF600 arrays and two x86 servers:
  - With NetApp EF600 all-flash arrays at the foundation of NVIDIA DGX SuperPOD, customers get a reliable storage foundation backed by six 9s of uptime.
  - The file system layer between the NetApp EF600 and the NVIDIA DGX H100 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 aligns NetApp's excellent support organization with customer requirements for performance and uptime. Customers 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.

### *NetApp BeeGFS building block*

image::EF\_SuperPOD\_buildingblock.png[]

## 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), natural language understanding (NLU) and generative AI (GenAI).
- 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 below lists the hardware components that are required to implement the solution for a single SU. The solution sizing starts with 32 NVIDIA DGX H100 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. Customers can add additional building blocks as the deployment size increases. For more information, see the [NVIDIA DGX H100 SuperPOD reference architecture](#) and [NVA-1164-DESIGN: BeeGFS on NetApp NVA Design](#).

**Table 1. Hardware requirements**

Hardware	Quantity
NVIDIA DGX H100	32
NVIDIA Quantum QM9700 switches	8 leaf, 4 spine
NetApp BeeGFS building blocks	3

### Software requirements

Table 2 below 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 customers can deploy with confidence. This fully integrated, turnkey architecture takes the risk out of deployment and puts anyone 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://docs.nvidia.com/dgx-superpod/reference-architecture-scalable-infrastructure-h100/latest/index.html#>

NVIDIA DGX SuperPOD Data Center Design Reference Guide

<https://docs.nvidia.com/nvidia-dgx-superpod-data-center-design-dgx-h100.pdf>

NVIDIA DGX SuperPOD: NetApp EF600 and BeeGFS

<https://nvidiagpugenius.highspot.com/viewer/62915e2ef093f1a97b2d1fe6?iid=62913b14052a903cff46d054&source=email.62915e2ef093f1a97b2d1fe7.4>

## **NVA-1156-DESIGN: NetApp EF-Series AI with NVIDIA DGX A100 systems and BeeGFS**

Abdel Sadek, Tim Chau, Joe McCormick and David Arnette, NetApp

NVA-1156-DESIGN describes a NetApp Verified Architecture for machine learning (ML) and artificial intelligence (AI) workloads using NetApp EF600 NVMe storage systems, the BeeGFS parallel file system, NVIDIA DGX A100 systems, and NVIDIA Mellanox Quantum QM8700 200Gbps IB switches. This design features 200Gbps InfiniBand (IB) for the storage and compute cluster interconnect fabric to provide customers with a completely IB-based architecture for high-performance workloads. This document also includes benchmark test results for the architecture as implemented.

[NVA-1156-DESIGN: NetApp EF-Series AI with NVIDIA DGX A100 systems and BeeGFS](#)

## **NVA-1156-DEPLOY: NetApp EF-Series AI with NVIDIA DGX A100 systems and BeeGFS**

Abdel Sadek, Tim Chau, Joe McCormick, and David Arnette, NetApp

This document describes a NetApp Verified Architecture for machine learning (ML) and artificial intelligence (AI) workloads using NetApp EF600 NVMe storage systems, the ThinkParQ BeeGFS parallel file system, NVIDIA DGX A100 systems, and NVIDIA Mellanox Quantum QM8700 200Gbps InfiniBand (IB) switches. This document also includes instructions for executing validation benchmark tests after the deployment is complete.

[NVA-1156-DEPLOY: NetApp EF-Series AI with NVIDIA DGX A100 systems and BeeGFS](#)

## Copyright information

Copyright © 2024 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.

LIMITED RIGHTS LEGEND: Use, duplication, or disclosure by the government is subject to restrictions as set forth in subparagraph (b)(3) of the Rights in Technical Data -Noncommercial Items at DFARS 252.227-7013 (FEB 2014) and FAR 52.227-19 (DEC 2007).

Data contained herein pertains to a commercial product and/or commercial service (as defined in FAR 2.101) and is proprietary to NetApp, Inc. All NetApp technical data and computer software provided under this Agreement is commercial in nature and developed solely at private expense. The U.S. Government has a non-exclusive, non-transferrable, nonsublicensable, 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) (FEB 2014).

## 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.