

## Solution Brief

# AI on E-Series and BeeGFS

Deploy and scale out your AI environment with ease

### Key Business Benefits

#### Increase Flexibility

- Scale easily to meet capacity or computational needs by using BeeGFS
- With a building block architecture, sizing a system to meet performance expectations is much simpler

#### Control Costs

- NetApp® EF570 and E5700 systems are proven to be some of the most price performant arrays on the market
- BeeGFS is open source and free to get started

#### Simplicity

- E-Series systems are easy to deploy and simple to manage
- BeeGFS runs in user space and requires no kernel changes, making it the easiest parallel file system to set up and configure

### The Challenge

Modern workloads such as artificial intelligence, machine learning, and deep learning (AI, ML, and DL) require innovative solutions. As datasets and computational needs continue to grow, expanding the infrastructure to support them is increasingly difficult. The requirement for accelerators like graphics processing units (GPUs) to do the parallel processing necessary for deep learning further complicates the ability to size systems to meet customers' needs.

### The Solution

Combining the BeeGFS parallel file system and NetApp EF570 and E5700 systems gives you a powerful set of tools to easily add compute and storage to meet performance requirements. The NetApp E5700 hybrid array makes an ideal storage target for parallel file systems by using high-capacity, cost-effective NL-SAS hard drives. The midrange NetApp EF570 all-flash array is a perfect target for the metadata requirements of a parallel file system, requiring just 2U of rack space to deliver extreme IOPs, sub 100-microsecond response times, and up to 21GBps of bandwidth. Both E5700 and EF570 systems include these enterprise-proven availability features:

- Redundant components with automated failover
- Intuitive storage management with comprehensive tuning functions
- Advanced monitoring and diagnostics with proactive repairs

A BeeGFS parallel file system offers the following benefits:

- Compared to other parallel file systems, it's easy to install and configure
- Open source and free to use (basic features) with very lightweight metadata and management services
- Can handle large and small files simultaneously with little impact on performance
- Provides graphical administration and monitoring, unlike complex legacy open-source parallel file systems

### Increase Flexibility

By using BeeGFS with NetApp E-Series systems and implementing a building-block architecture, you can scale out to meet your performance requirements, making sizing a system much simpler. For example, if you have the necessary GPUs and want to train at up to 90,000 images per second, you may find that acquiring just one of the systems is enough to support that task. A parallel file system such as BeeGFS also enables a user or a group of users to run more than one job simultaneously by separating BeeGFS clients to run different jobs.

## Control Costs

BeeGFS is free to use with its base features. Of course, there are caveats; Some features, as well as BeeGFS support, are not free. Some of the pay-to-use features include:

- **BeeOND.** A parallel file system that is spun up instantly across clients' local SSDs, creating a burst buffer to accelerate the training of an AI system.
- **Buddy mirroring.** An HA solution created by BeeGFS to ensure the safety of data.
- **Storage pools.** A manual tiering solution that enables the user to place data that needs faster access on media such as SSDs.

Storagereview.com had this to say about the NetApp EF570: "The EF570 may carry a midmarket price, but it offers an enterprise-grade performance profile that blows away everything we've seen in the lab. When it comes to demanding block storage workloads like analytics and emerging AI-style use cases where latency and availability mean dollars, it's important to have a storage tool that's designed for the job. The NetApp EF570 is a perfect solution for these workloads, bludgeoning those demanding databases into submission to deliver intelligence in the fastest way we've seen to date."<sup>1</sup>

As well as being highly performant, E-Series systems are extremely reliable. With redundant components and automated failover, E-Series systems can claim 5 nines (99.999%) of availability.

## Simplicity

Simplicity is among the most attractive features of both NetApp E-Series systems and the BeeGFS parallel file system. Because BeeGFS requires no kernel changes and runs in user space, you can set up your BeeGFS system in minutes. And you can expand the BeeGFS environment with ease, requiring only the IP of your management node and a service to be spun up (metadata, storage, or client service). Like BeeGFS, E-Series systems are very easy to set up and configure. On-box NetApp SANtricity® management makes deployment easy.

## Conclusion

NetApp E-Series and BeeGFS mesh very well to create a cost-effective and easy-to-scale solution. Like NetApp, BeeGFS recognizes customers' preference for simplicity and has created a very usable and easily deployable solution.. Customers can take advantage of E-Series and BeeGFS to get the job done.

For more information about deploying BeeGFS with NetApp E-Series, go to <https://www.netapp.com/us/media/tr-4755.pdf>.

## Solution Components

- Tensorflow-gpu
- OpenMPI NCCL
- HOROVOD
- CuDNN
- CUDA

## BeeGFS with GPU servers

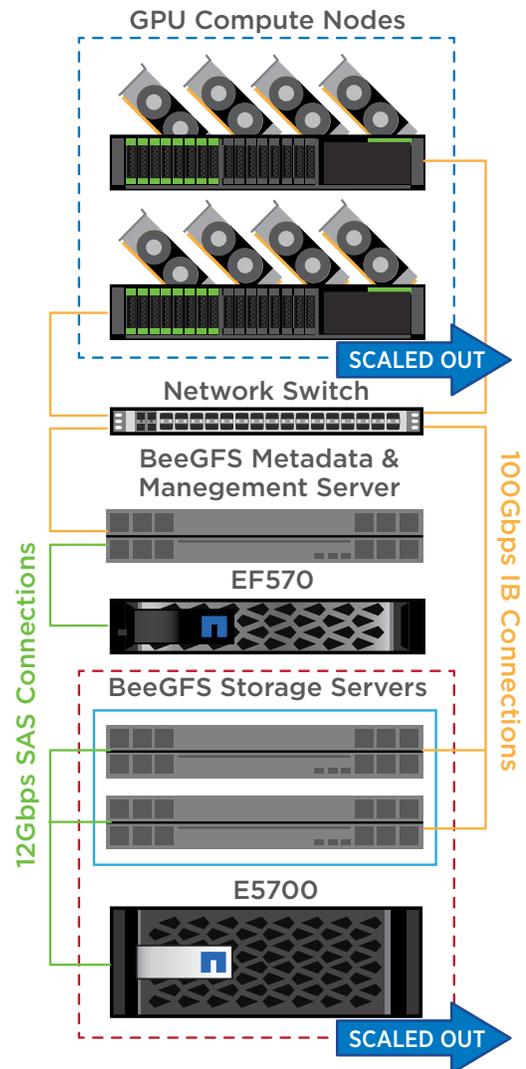


Figure 1

## About NetApp

NetApp is the data authority for hybrid cloud. We provide a full range of hybrid cloud data services that simplify management of applications and data across cloud and on-premises environments to accelerate digital transformation. Together with our partners, we empower global organizations to unleash the full potential of their data to expand customer touchpoints, foster greater innovation, and optimize their operations. For more information, visit [www.netapp.com](http://www.netapp.com). #DataDriven

<sup>1</sup>Armstrong, Adam. "NetApp EF570 All-Flash Array Review." *Storagereview.com*, October 2, 2018. [https://www.storagereview.com/netapp\\_ef570\\_allflash\\_array\\_review](https://www.storagereview.com/netapp_ef570_allflash_array_review)