



Solution Brief

NetApp E-Series Storage with BeeGFS

Quickly deploy a cost-effective, easy-to-manage HPC solution, with support from NetApp

Key Benefits

Keep Pace with Your Extreme Workloads

- Accelerate workloads with a BeeGFS parallel cluster file system that runs on industry-leading high-performance, low-latency NetApp® E-Series storage.
- Reduce downtime by relying on E-Series storage, which features 99.9999% availability.

Simplify Deployment and Management

- Get up and running quickly with an as-delivered solution that is easy to install and requires minimal configuration.
- Use automation tools such as Ansible to deploy and to manage BeeGFS with E-Series at scale.

Respond Quickly to Business Needs

- Seamlessly adapt to evolving workloads with a building-block architecture that easily scales performance and capacity as needed.

Lower Your TCO

- Reduce operating costs with high-density drives and price/performance-optimized storage building blocks.

Reduce Risk and Enable Success

- Configure your solution right from the start and maintain peak performance and availability with the help of NetApp support experts.

The Challenge

As technologies such as the Internet of Things (IoT), artificial intelligence (AI), and 3D imaging evolve, the size and amount of data that your organization has to work with is growing exponentially. Keeping ahead of this data growth requires a high-performance computing (HPC) solution with a lightning-fast, highly reliable IT infrastructure to process, store, and analyze your data. However, the cost and complexity of deploying and operating an HPC infrastructure to manage this critical data can be daunting.

The Solution

Whether you're looking for the origins of the universe, the next big oil reserve, a foolproof way to predict financial markets, or a cure for cancer, ThinkParQ and NetApp can help. A BeeGFS parallel cluster file system backed by NetApp E-Series storage is a proven, integrated solution with a simple, reliable, scalable, and cost-effective HPC infrastructure that keeps pace with your most extreme workloads.

BeeGFS and E-Series storage systems have been developed with a strong focus on performance. Both are designed for easy installation and for simplified management and use. They're the optimal combination no matter what size your organization is, and whether you are an experienced HPC guru or are taking your first steps into HPC.

BeeGFS is an open source product. By purchasing a support contract from NetApp, you enhance the benefits that you receive from your NetApp and BeeGFS deployment. 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. When you choose NetApp as your BeeGFS support provider, storage experts are just a phone call away.

Keep Pace with Your Extreme Workloads

In the highly competitive world of business, speed is everything. However, even the fastest supercomputer can't meet expectations if it doesn't have equally fast storage to support it.

NetApp E-Series storage with BeeGFS gives you consistent, near-real-time access to your data. BeeGFS transparently spreads your data across multiple servers and their back-end storage to prevent bottlenecks and to enable continuous high performance. E-Series storage systems can support any number of workloads simultaneously and can deliver up to 2 million random read IOPS and 24GBps sustained (maximum burst) write bandwidth per scalable building block. Optimized for both flash and spinning media, E-Series systems include built-in technology that monitors your workloads and automatically adjusts configurations to maximize performance. With the 99.9999% reliability of E-Series systems, your HPC environment will not miss a beat.



A NetApp Solutions
Integrator Partner

Simplify Deployment and Management

Managing hundreds or thousands of compute nodes along with the network and storage components in an HPC environment is a daunting task. And when your environment needs to scale to meet growing demand, the complexities that are involved add to the challenge. In an environment where complexity rules, simplicity is critical.

Unlike other file systems that require multiple platforms and models to scale, the NetApp HPC and BeeGFS solution is easy to install and to support. The modular design enables your IT staff to nondisruptively add performance and capacity without complex deployments or migrations. Traditional scripting or automation platforms such as Ansible enable you to quickly configure both E-Series and BeeGFS components at any scale.

Respond Quickly to Business Needs

AI, IoT, and 3D imaging all have one thing in common: They generate huge amounts of data. To manage and process all that data in an HPC environment, you need a solution that can quickly respond to data growth.

NetApp and BeeGFS together gives you an agile enterprise storage solution that scales easily and seamlessly. The flexibility and scalability of the solution enable it to support and to adapt to evolving workloads, making it one of the strongest foundations for future storage requirements. Modular storage building blocks composed of E-Series hybrid flash arrays that are connected to BeeGFS-powered storage servers give you a granular, building-block approach to growth. By increasing the number of storage building blocks in your system, you can seamlessly scale up the performance and capacity of the file system, from small clusters up to enterprise-class systems with hundreds of nodes.

Lower Your TCO

When you're considering the cost of an HPC solution, your initial investment in hardware and software is just the beginning. To more accurately assess cost, you must look at the TCO for the solution. So, in addition to the deployment costs, you also must consider the costs of operating the solution after you have deployed it.

With the NetApp and BeeGFS solution, you get enterprise storage in price/performance-optimized building blocks that are built on the E-Series storage architecture and that make small to large configurations cost efficient. In addition, NetApp HDDs and SSDs meet strict quality, performance, and interoperability requirements before they are accepted. This extra scrutiny results in failure rates that are 4 times lower than for commodity drives. NetApp also leverages industry-leading density per terabyte, which lowers your power, cooling, and support costs. You also get the flexibility to support 100Gb InfiniBand (iSCSI Extensions for RDMA and SCSI RDMA Protocol), 100Gb NVMe over IB, 100Gb NVMe over RoCE, 32Gb NVMe over FC, 32Gb FC, 25Gb iSCSI, and 12Gb SAS connectivity. All these benefits add up to significantly lower TCO.

Reduce Risk and Enable Success

To enable your success, NetApp support for BeeGFS places the assistance of NetApp's excellent support organization at your fingertips. Let us help you by providing configuration advice, assisting with performance tuning, and troubleshooting any issues that might arise. As the storage experts, the NetApp support team can help assess and troubleshoot the overarching BeeGFS deployment, not just the file system. Our experts use proven best practices, tools, and methodologies to guide you to get optimal performance out of your NetApp and BeeGFS investment.

About BeeGFS

Delivered by ThinkParQ, BeeGFS is one of the leading parallel cluster file systems, developed with a strong focus on performance and designed for very easy installation and management. If I/O-intensive workloads are your problem, BeeGFS is the solution. For more information, visit www.beegfs.io.

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. #DataDriven