



ExtremeStor™ B

BeeGFS® Parallel File System Appliance With Expert Architectural Design, Implementation and Support

Highlights

High Performance Parallel File System

BeeGFS transparently spreads data across multiple servers, linearly scaling performance and capacity, from small clusters to supercomputer-class systems with thousands of nodes

Easy to Manage

Graphical administration and monitoring with a command line interface (with few commands) provide easy management, avoiding the notorious complexity of legacy open source parallel file systems

No Charge Open Source Software

Basic BeeGFS is free of charge open source software, combined with vendor agnostic industry standard hardware make ExtremeStor B economically compelling

Additional Enterprise Data Services

High-availability, quotas, ACLs, and with ZFS as the underlying file system: RAID Z, data compression, snapshots, and powerful management tools are available under additional support contract

Network and Protocol Support

- Standard TCP/IP and RDMA over Converged Ethernet (RoCE), InfiniBand and OmniPath
- POSIX Client, NFS and SMB export

Scalable Distributed Metadata

BeeGFS uses multiple dedicated metadata servers to manage global metadata in order to deliver best in class metadata performance and linear scalability

Expert Architectural Solution Design, Tuning, and Implementation

Applied Data Systems performs extensive analysis of existing and future needs, comprehensive solution architecture and validated hardware and software build that ships fully integrated

World's Fastest Growing Parallel File System

ExtremeStor B uses BeeGFS, the fastest growing open source parallel file system developed with a strong focus on performance and designed for easy installation and management. ExtremStor B is ideal for demanding, high-performance, high-throughput workloads found in technical computing for simulation, life sciences, oil & gas, deep learning, analytics, media, defense and financial services.

ExtremeStor B Delivers Extreme Performance and Scale

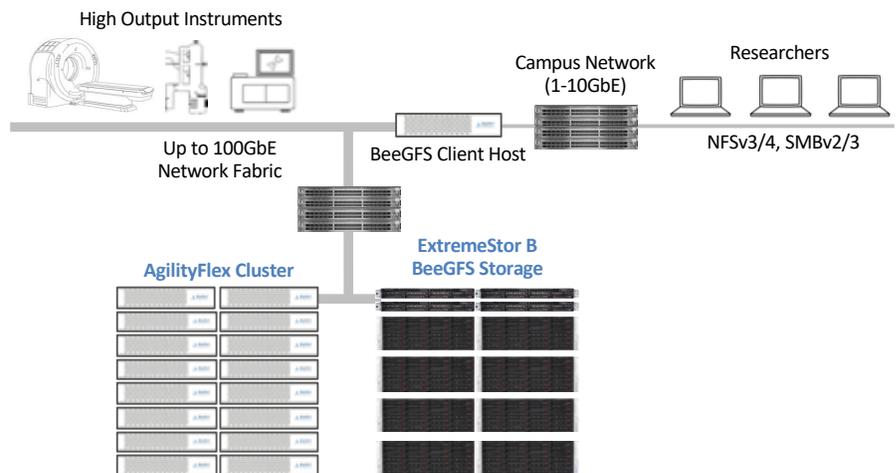
ExtremeStor B delivers maximum performance and scalability on fully integrated, top quality, industry standard hardware for maximum performance, reliability and data protection across a wide range of technical applications. ExtremeStor B with BeeGFS distributes files in parallel across multiple storage servers with dedicated, distributed metadata processing. ExtremeStor B delivers up to 8 GB/s client throughput with a single process streaming on a 100Gb network, with only a few streams capable of fully saturating the network. ExtremeStor B delivers best in class metadata performance with linear scalability through dynamic metadata name-space partitioning and distributing metadata operations per directory and subdirectory across metadata nodes in a simplified manner.

Under the Hood, What Makes BeeGFS so Fast

In contrast to other parallel file systems, BeeGFS uses all available RAM on its storage servers to quickly write bursts of data into the server RAM cache and to quickly read data from it. BeeGFS also serves data direct from the cache if it has already been recently requested by another client. BeeGFS also aggregates small I/O requests into larger blocks before writing them to disk. A single large file is distributed across multiple storage targets for high throughput.

Easy Appliance Deployment, Setup and Administration

In addition to performance, BeeGFS was designed for easy deployment and administration. The graphical administration and monitoring system facilitates simple and intuitive management including cluster installation, load statistics, storages service management and health monitoring. ExtremeStor B is delivered as an integrated hardware software appliance from Applied Data Systems.



ExtremeStor B High Performance, Single Namespace Architecture

Storage Pools Combine the Performance of Flash with the Economics of Disk

BeeGFS Storage Pools make different types of storage devices available within the same namespace. Economic high capacity disks can be accessed in parallel for high throughput and capacity, combined with a high performance flash tier.

Dynamic Network Fail-Over

BeeGFS supports multiple networks and dynamic fail-over in case one of the network connections is down

BeeOND (BeeGFS On Demand)

BeeOND allows on the fly creation of temporary parallel file system instances on the internal SSDs of compute nodes on a per-job basis for burst-buffering

Built-in High Availability By Replication

BeeGFS includes a replication HA mechanism called Buddy Mirroring that is fully integrated and does not rely on special hardware

Leadership Reliability

ExtremeStor B is delivered fully integrated, with top quality, industry standard hardware for maximum performance, reliability and data protection

Flexible Building Blocks

ExtremeStor B is delivered as a modular, repeatable, and highly supportable solution consisting of best of breed industry standard components

White Glove Installation and Support

ExtremeStor B is expertly installed and supported by Applied Data Systems who is the single point of contact for all support issues

Exclusively Focused on Technical Computing and Data Management

Applied Data Systems is 100% focused on technical computing, specializing in CPU and GPU computing, low latency networking and high-performance storage

Engineered for High Availability, Data Protection and Fault Tolerance

ExtremeStor BeeGFS storage servers come with underlying RAID (either RAID-6 or RAID-Z2) to transparently handle disk errors. BeeGFS includes a HA mechanism that is fully integrated and which does not rely on special hardware. This approach is called Buddy Mirroring, based on the concept of pairs of servers (the so-called buddies) that internally replicate each other and that help each other in case one of them has a problem. The built-in BeeGFS Buddy Mirroring approach can tolerate the loss of complete servers including all data on their RAID volumes - on commodity servers and shared-nothing hardware. Buddy Mirroring can also be used to put buddies in different failure domains, different racks or different server rooms.

Compelling Open Source Economics and Simplicity

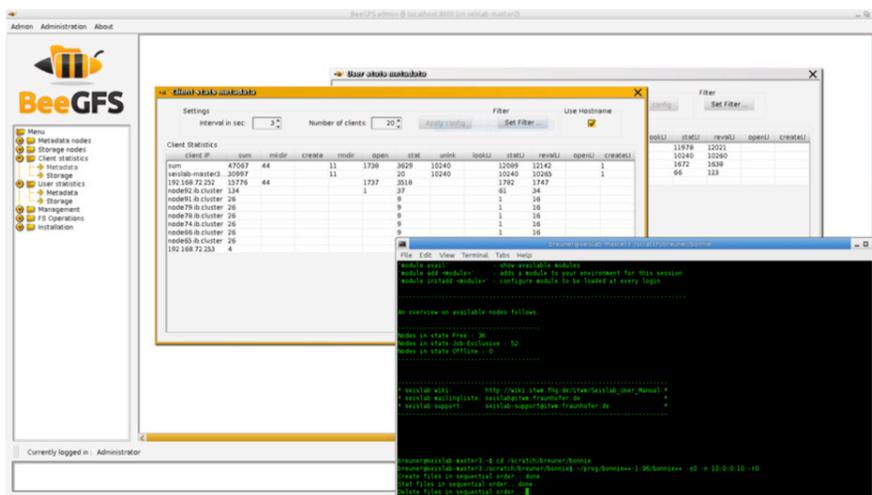
With basic BeeGFS file system software being available free of charge, ExtremeStor B reduces the expense of commercially licensed offerings and the notorious complexity of legacy open source parallel file systems. ExtremeStor B with BeeGFS is delivered fully integrated with expert solution design and support from Applied Data Systems.

Storage Pools Combine the Performance of Flash with the Economics of Disk

All-flash storage is still too expensive for systems that require large capacity. ExtremeStor B BeeGFS Storage Pools make different types of storage devices available within the same namespace. By having SSDs and HDDs in different pools, an application can enjoy all-flash performance for a project while still providing the advantage of cost-efficient high capacity spinning disks for other data. The placement of the data is fully transparent to applications. Data stays inside the same directory when moved to a different pool and files can be accessed directly without any implicit movement, no matter which pool the data is currently assigned to – all within the same namespace.

Technical Applications Can Take Immediate Advantage of BeeGFS Performance

BeeGFS is POSIX based software defined storage, so applications do not have to be rewritten or modified to take advantage of BeeGFS. BeeGFS clients accessing the data inside the file system communicate with the storage servers via network, via any TCP/IP based connection or via RDMA-capable networks like InfiniBand (IB), Omni-Path (OPA) and RDMA over Converged Ethernet (RoCE).



BeeGFS Administration and Monitoring GUI



For more information from Applied Data Systems:
844.371.4949 | info@applieddatasystems.com | www.applieddatasystems.com