



BeeGFS[®]

The Leading Parallel
Cluster File System



ABOUT BEEGFS

What is BeeGFS

BeeGFS (formerly FhGFS) is the leading parallel cluster file system, 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.



Why use BeeGFS

BeeGFS transparently spreads user data across multiple servers. By increasing the number of servers and disks in the system, you can simply scale performance and capacity of the file system to the level that you need, seamlessly from small clusters up to enterprise-class systems with thousands of nodes.

Get The Most Out Of Your Data

The flexibility, robustness, and outstanding performance of BeeGFS help our customers around the globe to increase productivity by delivering results faster and by enabling new data analysis methods that were not possible without the advantages of BeeGFS.

System Architecture



KEY ASPECTS

Maximum Scalability

BeeGFS offers maximum performance and scalability on various levels. It supports distributed file contents with flexible striping across storage servers on a per-file or per-directory basis as well as distributed metadata.

BeeGFS is optimized especially for use in environments where performance matters to provide:

- **Best** in class client throughput: 8 GB/s with only a single process streaming on a 100Gbit network, while a few streams can fully saturate the network.
- **Best** in class metadata performance: Linear scalability through dynamic metadata namespace partitioning.
- **Best** in class storage throughput: BeeGFS servers allow flexible choice of underlying file system to perfectly fit the given storage hardware.

BeeGFS Storage Pools make different types of storage devices available within the same namespace.


By having SSDs and HDDs in different pools, pinning of a user project to the flash pool enables all-flash storage performance for the current project while still providing the advantage of the cost-efficient high capacity of spinning disks for other data.

Maximum Flexibility

BeeGFS supports a wide range of Linux distributions such as RHEL/Fedora, SLES/OpenSuse or Debian/Ubuntu as well as a wide range of Linux kernels from ancient 2.6.18 up to the latest vanilla.

The storage services run on top of an existing local filesystem (such as xfs, zfs or others) using the normal POSIX interface and clients and servers can be added to an existing system without downtime.

BeeGFS supports multiple networks and dynamic failover in case one of the network connections is down.



BeeGFS client and server components can also run on the same physical machines. Thus, BeeGFS can turn a compute rack into a cost-efficient converged data processing and shared storage unit, eliminating the need for external storage resources and providing simplified management.

BEEGFS ON DEMAND

Maximum Usability

The BeeGFS server components are userspace daemons, while the client is a native kernel module that does not require any patches to the kernel itself. All BeeGFS components can be installed and updated without even rebooting the machine.

For installation and updates there are rpm/deb package repositories available; for the startup mechanism, easy-to-use system service scripts are provided.


BeeGFS was designed with easy administration in mind. The graphical administration and monitoring system enables dealing with typical management tasks in a simple and intuitive way, while everything is of course also available from a command line interface:

- Live load statistics, even for individual users
- Cluster installation
- Storage service management
- Health monitoring
- And more...

Excellent documentation helps to have the whole system up and running in one hour.

BeeOND

BeeOND (BeeGFS on demand) allows on the fly creation of a complete parallel file system instance on a given set of hosts with just one single command.




BeeOND was designed to integrate with cluster batch systems to create temporary parallel file system instances on a per-job basis on the internal SSDs of compute nodes, which are part of a compute job. Such BeeOND instances do not only provide a very fast and easy to use temporary buffer, but also can keep a lot of I/O load for temporary or random access files away from the global cluster storage.

BUDDY MIRRORING

Fault tolerance

BeeGFS storage servers are typically used with an underlying RAID to transparently handle disk errors. Using BeeGFS with shared storage is also possible to handle server failures. The built-in BeeGFS Buddy Mirroring approach goes even one step further by tolerating the loss of complete servers including all data on their RAID volumes - and that with commodity servers and shared-nothing hardware.



The built-in BeeGFS Buddy Mirroring automatically replicates data, handles storage server failures transparently for running applications and provides automatic self-healing when a server comes back online, efficiently resyncing only the files that have changed while the machine was offline.

BeeGFSv7: Newest features

- Storage pools
- Free space balancing when adding new hardware
- Metadata event logging enhancement
- fsck improvements
- Kernel 4.14, 4.15 support
- NIC handling during startup

More Features

We already talked about the BeeGFS key aspects scalability, flexibility and usability and what's behind them. But there are way more features in BeeGFS:


- Runs on various platforms, such as x86, OpenPOWER, ARM, Xeon Phi and more
- Re-export through Samba or NFSv4 possible
- Support for user/group quota and ACLs
- Fair I/O option on user level to prevent a single user with multiple requests from stalling requests of other users
- Automatic network failover, e.g. if InfiniBand is down, BeeGFS automatically switches to Ethernet and back later
- Online file system sanity check that can analyze and repair while the system is in use
- Built-in benchmarking tools to help with optimal tuning for specific hardware and evaluate hardware capabilities
- Support for cloud deployment on e.g. Amazon EC2 or Microsoft Azure

Not in the list? Just get in touch and we're happy to discuss all your questions.

BENCHMARKS


Metadata Operations

BeeGFS was designed for extreme scalability. In a testbed¹ with 20 servers and up to 640 client processes (32x the number of metadata servers), BeeGFS delivers a sustained file creation rate of more than 500,000 creates per second, making it possible to create one billion files in as little time as about 30 minutes.



Throughput Scalability

In the same testbed system with 20 servers, each equipped with a single node local performance of 1332 MB/s (write) and 1317 MB/s (read), and 160 client processes, BeeGFS demonstrates linear scaling to a sustained throughput of 25 GB/s - which is 94.7% of the maximum theoretical local write and 94.1% of the maximum theoretical local read throughput.



¹Benchmark System:

20 servers with 2x Intel Xeon X5660 @ 2.8 GHz, 48GB RAM running a Scientific Linux 6.3, Kernel 2.6.32-279. Each server is equipped with 4x Intel 510 Series SSD (RAID 0) running ext4 as well as QDR Infiniband. Tests performed using BeeGFS version 2012.10.

LICENSING MODEL

Free to use & Open source

BeeGFS is free to use for self-supporting end users and can be downloaded directly from www.beegfs.io

Professional Support

If you are already using BeeGFS but you don't have professional support yet, here are a few reasons to rethink this.

| | Free Edition | With Support |
|--|--------------|--------------|
| Community Mailing List | ● | ● |
| Free Updates | ● | ● |
| Ready-to-use binary packages & Source Code | ● | ● |
| Enterprise Features (High-availability, quota, ACLs) | | ● |
| Next Business Day Service Level Agreements for Support | | ● |
| Direct Contact to the File System Developers | | ● |
| Early Updates and Hotfixes | | ● |
| Customer Section: HowTos and more Documentation | | ● |

User Comments

"After many unplanned downtimes with our previous parallel filesystem, we moved to BeeGFS more than two years ago. Since then we didn't have any unplanned downtime for the filesystem anymore."

Michael Hengst, University of Halle, Germany

"We are extremely happy with our 3.1PB BeeGFS installation on 30 servers. It is rock-solid."

Rune Møllegaard Friborg,
University of Aarhus, Denmark

"Now under heavy load, our large BeeGFS system is performing well - bioinfo users are seeing >2x speedup in their apps, with no more hotspots when hitting common index files. Unlike the previous system, BeeGFS does not lock up under heavy load and even under heavy load, interactive use remains zippy. The only complaint is how long it's taking us to move data off the old system."

Harry Mangalam, UC Irvine, USA

"The network file access performance of BeeGFS is on the same level that we see when our applications are running on servers with a direct-attached RAID."

Genomatix, Germany

HAPPY USERS WORLD-WIDE

Scientific Computing

BeeGFS is widely popular among universities and the global research community, powering some of the fastest supercomputers in the world to help scientists analyze large amounts of data efficiently every day.



Life Sciences

BeeGFS is the parallel file system of choice in life sciences. The fast growing amount of genomics data to store and analyze quickly in fields like Precision Medicine make BeeGFS the first choice for our customers.



Finance, Oil & Gas, Media, Automotive, ...

BeeGFS is used in many different industries all around the globe to provide fast access to storage systems of all kinds and sizes, from small scale up to enterprise-class systems with thousands of hosts.



ThinkParQ GmbH

Corporate Headquarters

Trippstadter Str. 110
67663 Kaiserslautern, Germany

Global Sales & Consulting

Phone: +49 631 277 576 300
sales@thinkparq.com

www.thinkparq.com
www.beegfs.io



v1711