



## Highlights

### High Performance BeeGFS Parallel File System

A BeeGFS deployment at UCSB delivered over 13 GB/s with RAID 6 data protection. BeeGFS transparently spreads data across multiple servers, linearly scaling performance and capacity, from small clusters to supercomputer-class systems with thousands of nodes.

### UCSB Configuration

- 4x Object Storage Servers with 36x6TB NL SAS disks each
- 2x Metadata servers with Buddy Mirroring for high availability
- To meet performance and capacity requirements, Applied Data Systems deftly architected the appropriate drive capacity and number of Object Servers
- Optional features such as BeeOND deliver the performance of all-flash arrays as compute jobs require
- Affordable object server building blocks can be added in order to scale both performance and capacity

### 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

### 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

### White Glove Installation and Support

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

### BeeGFS Powers the Newest Computing Cluster at UC Santa Barbara

Researchers at UC Santa Barbara now have 300 teraFLOPs of computing power at their fingertips, thanks to the recent installation of the campus's new computing cluster "Pod", named for the term describing a group of whales. The \$1.1 million high-performance computing (HPC) cluster is hosted by the campus's Center for Scientific Computing (CSC) at the California NanoSystems Institute (CNSI) and Materials Research Laboratory (MRL), for University scientific and engineering research.

### Applied Data Systems ExtremeStor B BeeGFS Appliance Technical Details

UCSB's Pod cluster consists of 64 Applied Data Systems AgilityFlex compute nodes, each with 40 core and 192GB of RAM, plus 4 nodes with over 1TB of RAM (768GB of RAM + 300 GB Optane Memory Drive), and 3 GPU nodes with four NVIDIA V100/32GB GPUs with NVLINK. The cluster interconnect is a 100Gb Intel Omni-Path network.

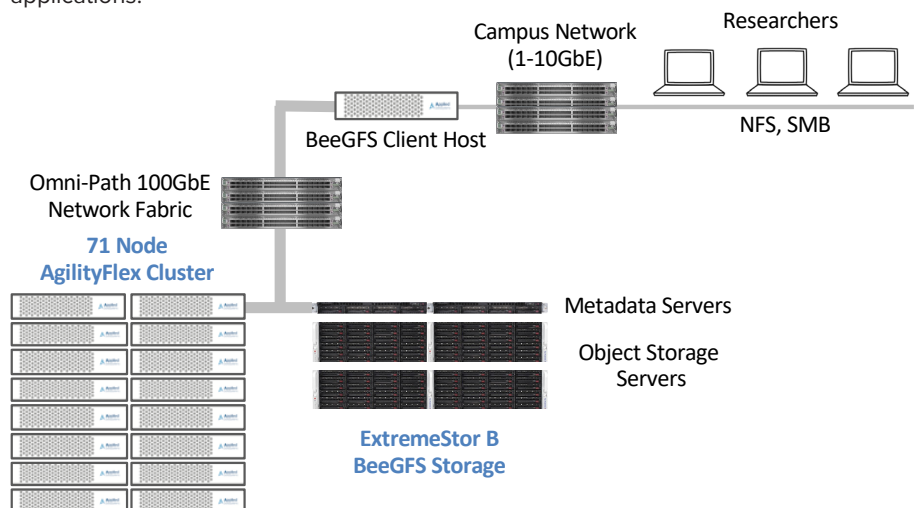
The 650TB BeeGFS parallel filesystem consists of two mirrored Metadata Servers and four Object Servers, with all BeeGFS servers attached to the 100Gb Omni-Path Network. Each Metadata server has four 240GB SSDs configured as RAID 10. Each of the four Object Storage servers include 36 x 6TB NL SAS hard drives in a RAID 6, 8+2 configuration. BeeGFS is open source and supported through the community. Optional professional support was included in the configuration.

### BeeGFS Performance Exceeded University Requirements

Applied Data Systems expertly architected a configuration that exceeded the University's performance, capacity and economic objectives. The Applied Data Systems deployment at UCSB delivered over 13 GB/s - beating the University's selection criteria by 30%, while providing robust RAID 6 data protection.

### BeeGFS Delivers Performance to Drive Machine Learning and AI Workloads

Engineering, chemistry and physics researchers are among the top users of the CSC-hosted computing cluster. In addition, Pod features GPU processor units to meet the emerging research needs of technology in the fields of machine learning and artificial intelligence, with BeeGFS providing the data throughput needed for these advanced applications.



### BeeGFS Powers "Pod", the Newest Computing Cluster at UC Santa Barbara