



Mercury Network Storage System

BlueArc's Next Generation Platform for Consolidation

HIGHLIGHTS

- **Massive computing parallelism delivers the performance to support multiple, varied applications**
- **Scales up to 2PB usable capacity, a 256TB file system, and thousands of concurrent users**
- **Clusters up to eight nodes with Cluster Namespace to enable horizontal scalability***
- **Concurrent support for iSCSI, NFS and CIFS eliminates storage silos**
- **Dynamic, policy-based data migration and caching simplify management of infrequently accessed data**
- **Advanced Virtualization Framework delivers Thin Provisioning and Virtual Servers**
- **Integration with many legacy third-party network attached storage devices protects prior investment**
- **Support for multiple, best-in-class storage options for varied workloads**

The BlueArc Mercury™ Server is a next generation network storage platform that consolidates capacity across multiple applications and simplifies storage management for businesses with mid-range storage requirements—without compromising performance and scalability. The Mercury platform deploys BlueArc's hardware accelerated file system, and its open architecture excels across a variety of application environments, including general purpose file systems, database, messaging and online fixed content. When combined with BlueArc's advanced data management services and an integrated Storage Ecosystem, Mercury provides for an optimized storage solution that:

- Consolidates application storage requirements, leveraging advanced file system capabilities and virtualization to reduce the TCO of storage infrastructures
- Enables the highest performance and scalability of any product in its class through a hardware accelerated file system
- Simplifies management of large, dynamic data sets with Intelligent Tiered Storage to reduce the cost of keeping more data online
- Supports best-in-class storage devices and complementary third-party technologies through an integrated Storage Ecosystem to eliminate storage silos

Using Mercury, organizations can replace a number of dedicated legacy systems with far fewer Mercury systems, supporting multiple applications and a larger number of concurrent users. As a result, they can eliminate isolated storage silos, reduce infrastructure and energy costs, and simplify management—reducing overall total cost of ownership.

Hybrid-Core Architecture Delivers Performance

Mercury's Hybrid-Core Architecture delivers the highest performance available in mid-range storage systems. This architecture employs field programmable gate arrays (FPGAs) and traditional multi-core processors to optimize and separate data movement and management processes that normally compete for system resources. Rather than using shared buses and shared memory that require arbitration and can cause significant performance fluctuations and reductions, data is transferred between logical blocks in a point-to-point fashion over dedicated data transfer paths. As a result, the Hybrid-Core Architecture supports advanced file system functions while preventing conflicts or bottlenecks and ensuring consistently high performance.

*Initial software supports clustering 2 nodes, 8 node clusters will be supported in a subsequent software upgrade.

Hardware Accelerated File System Simplifies Management

Mercury's unique Hardware Accelerated File system delivers no-compromise system performance and scalability while multiple levels of virtualization overcome the complexities of large scale file system management and data availability. BlueArc's Cluster Namespace is a single, global namespace that integrates all elements of the storage system to provide a single logical view of the data regardless of where it resides in physical storage. The Cluster Namespace feature also makes it easy to assign and reassign file systems to virtual servers and physical nodes as usage and performance requirements change, without impacting user access to files and data. Generous capacity limitations of 2PB usable storage and 16 million objects per directory provide room for consolidation and future growth.

Also, comprehensive virtualization tools simplify the administration of file system functions and optimize resource utilization. Each physical Mercury node or cluster can be partitioned into 64 virtual servers with their own IP addresses, management policies and file systems. These virtual servers can be moved across the cluster for load balancing or to manage capacity and downtime. A maximum file system size of 256TB reduces the need to manage multiple small file systems.

Mercury also logically organizes RAID storage into shared, virtualized storage pools from which multiple file systems, physical servers and virtual servers can be provisioned. Administrators can expand file systems by using unallocated free space from this pool on demand or automatically using policy settings without downtime. Dynamic Storage Balancing stripes the data set across the maximum number of available drive spindles, optimizing performance and resource utilization while eliminating hot spots. As capacity is added, file blocks are redistributed across the new spindles to immediately improve performance while maintaining a high level of resource utilization.

Intelligent Tiered Storage Supports Multiple Applications

Mercury's Intelligent Tiered Storage allows administrators to reduce infrastructure costs by automatically placing data on the most cost-effective or highest utility tier of storage without compromising accessibility. Leveraging Data Migration, administrators can easily set up policies based on rules, parameters or triggers to automatically migrate data from a file system or volumes to the appropriate tier of storage—optimizing cost and performance. As data and copies move transparently between storage tiers, the file system remains intact so users can continue to find their files in the same location.

For environments with large, dynamic unstructured data sets and unpredictable demand, Dynamic Read Caching capabilities instantly copy active files from lower performance storage tiers to a high performance Fibre Channel or SSD cache storage tier for use across physical or virtual servers. This aggregates bandwidth and improves response time to prevent quality of service issues during spikes in demand.

A graphical, rules-based policy engine simplifies data management by providing "what if" analysis tools, policy templates, a scheduler for recurring policies and integrated data protection safeguards. As they create policies, administrators can select from variables including last access, file name, path and user identity to create quotas and thresholds that govern automated data movements.

Storage Ecosystem Eliminates Storage Silos

The Mercury Storage Ecosystem enables organizations to deploy a wide variety of storage arrays and media as a single entity to create an optimal tiered storage environment, including high-performance online, moderate-performance nearline, and infrequently accessed archival data. The Mercury Storage Ecosystem supports best-in-class storage arrays from multiple vendors that include SSD, FC, SAS and SATA drives, as well as leading tape libraries. Third party storage devices such as other network storage appliances and complementary technologies such as deduplication or archiving appliances are easily integrated into the Storage Ecosystem as another tier of storage within the Namespace.

For entry level configurations, Mercury servers and Mercury 2-node clusters can be directly connected to up to 2 storage arrays using built-in Fibre Channel SAN ports, reducing cost and installation complexity. In more advanced configurations, a switched Fibre Channel SAN creates a networking layer between the Mercury servers and the storage arrays to enable tiered storage. This SAN enables LAN free data movement between integrated tiers of storage with full data path redundancy. This architecture also ensures that a solution can scale up or out independently.

Lower TCO

The Mercury Network Storage System enables organizations with midrange networked storage requirements to consolidate storage infrastructure and simplify data management of large unstructured data sets—with high performance and scalability. Learn how these capabilities will help lower storage total cost of ownership, reduce administrative complexity, and accelerate revenue.



BlueArc Corporation
Corporate Headquarters
50 Rio Robles Drive
San Jose, CA 95134
t 408 576 6600
f 408 576 6601
www.bluearc.com

BlueArc UK Ltd.
European Headquarters
Queensgate House
Cookham Road
Bracknell RG12 1RB, United Kingdom
t +44 (0) 1344 408 200
f +44 (0) 1344 408 202