

# Veritas Storage Foundation™ Cluster File System by Symantec

## Concurrent access to shared data

Veritas Storage Foundation Cluster File System enables concurrent file access from multiple servers to provide a flexible, high-performance, and highly available platform for shared data in a SAN environment. It's built for commercial transactional workloads and provides a mature single file system schema that is cache-coherent and POSIX-compliant. A cluster-wide volume and file system configuration allows for simplified management. Organizations also benefit from an integrated cluster volume manager that presents every node in the cluster with the same logical view of shared device configurations.

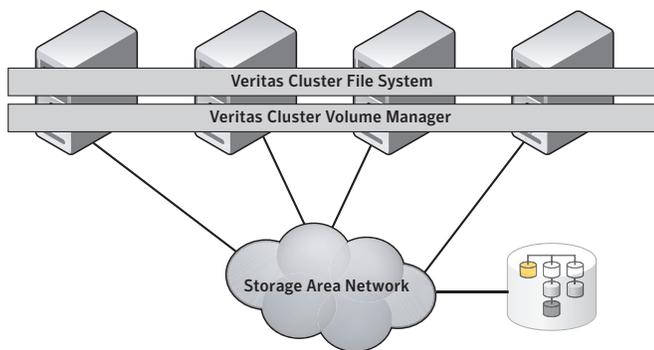


Figure 1. Veritas Storage Foundation Cluster File System Architecture

### Highlights

- **Benefit from storage consolidation**—Enable storage consolidation by allowing multiple applications access to a single storage pool
- **Improve application performance and scalability**—Add servers dynamically to meet scalability requirements without compromising performance

- **Minimize application downtime**—Utilize faster application failover to reduce planned and unplanned downtime
- **Ensure data consistency**—Maintain cache and storage coherency with atomic writes throughout the cluster
- **Simplify storage administration**—Manage multiple nodes as one with centralized management and advanced snapshot capabilities
- **Reduce storage costs**—Assign data to multiple tiers of storage based on predefined policies
- **Ensure data integrity**—Prevent data corruption in the event of a “split brain” condition with I/O fencing

### Benefit from storage consolidation

IT organizations that have information-sharing requirements across multiple applications can minimize the need to create redundant copies of information by implementing Veritas Storage Foundation Cluster File System. With the integration of a cluster file system and cluster volume manager, IT organizations can architect applications so that applications residing on multiple server nodes have concurrent access to the same information through shared storage (i.e., SAN-attached storage), thus minimizing the need to replicate, copy, or move information. With Veritas Storage Foundation Cluster File System, organizations can achieve near linear scalability by adding up to 32 nodes, all with concurrent access to a single storage pool, greatly increasing the benefits of physical storage consolidation via a storage area network.

### Improve application performance and scalability

IT organizations have two options when examining application scalability requirements: a scale-up (larger SMP servers) or a scale-out (multiple smaller SMP servers) approach. However, if an organization selects a scale-out approach, a single cluster file system across all nodes is essential to coordinate read/write access to the storage pool and ensure data integrity. Cluster File System does this and also ensures that performance isn't compromised, with its advanced file-locking capability providing parallel access to the same file from multiple nodes. A critical enabling technology is the advanced file-locking architecture, Multiple Transaction Server, that distributes file ownership across all nodes in the cluster (supporting up to 32 nodes), ensuring near linear scalability without requiring unnecessary metadata communication throughout the cluster.

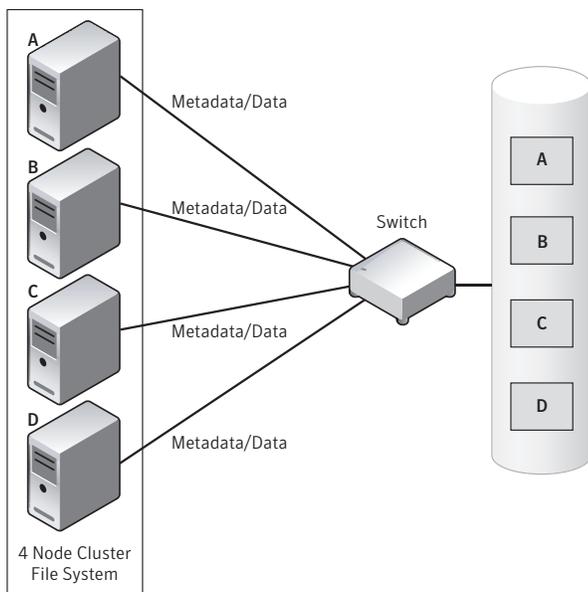


Figure 2. Multiple Transaction Server

### Minimize application downtime

Veritas Storage Foundation Cluster File System is tightly integrated with Veritas™ Cluster Server by Symantec, providing a comprehensive application failover solution minimizing both planned and unplanned downtime across all nodes in the cluster, up to 32. In the event of an application failure or a node failure, the application is dynamically migrated to an available node in the cluster without application or business interruption, and eliminates any single point of failure. Furthermore, cluster ownership is dynamically redirected in the event of the “primary node” failure. An additional benefit to a cluster file system architecture is that, as all nodes in the cluster have visibility to a shared storage pool, there is no need to manually mount storage to unique nodes in the event of application or node failures—minimizing application downtime, both planned and unplanned.

### Ensure data consistency

In an application architecture where multiple server nodes have access to the same storage pool, cache coherency is of the utmost importance, ensuring data consistency and integrity. There must be a single version of all files in a cluster that are visible to all nodes in the cluster; if each node has “its version” of the file (particularly during writes), it results in corrupt information. Veritas Storage Foundation Cluster File System uses the Veritas Global Lock Manager (GLM) to make writes throughout the cluster atomic. In doing so, even though application performance is maximized by having files cached at the local node level, GLM ensures that during any application write, there is only a single version of that file, preventing data corruption.

### Simplify storage administration

Veritas Storage Foundation Cluster File System minimizes the storage administration burden in many ways. First, as all nodes in the cluster have visibility to the same storage pool, the cluster can share a single set of configuration and data files. Therefore, administrators only need to “manage” a single node in the cluster and need not have the management burden scale with the number of nodes in the cluster. Second, the necessary processing overhead of backup/recovery operations can be isolated to a single node (since all nodes have access to all files). A dedicated node can perform all backup/recovery tasks (isolating CPU and I/O resources) for the entire cluster, leaving the application processing to be shared among other nodes in the cluster. Third, if business requirements warrant additional levels of data protection, FlashSnap™, a feature of Veritas Storage Foundation Cluster File System, can create point-in-time copies of production information (either at the volume or file system level). Volume-level point-in-time copies can be used for off-host processing without hindering the production cluster, or for disk-based recovery in the event of system failure/corruption. File system point-in-time copies can also provide checkpoints for disk-based recovery, and the file system can be “rolled back” to a consistent point in time. Just as backup/recovery operations can be dedicated to a single node, FlashSnap functions can be as well, freeing the other nodes in the cluster for application transactions.

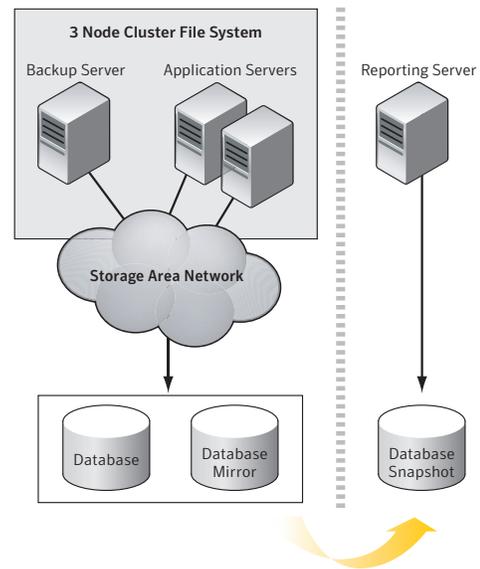


Figure 3. Veritas Storage Foundation Cluster File System Architecture

### Reduce storage costs

With the Dynamic Storage Tiering capabilities of Storage Foundation, unimportant or out-of-date files can be moved to less expensive storage devices without changing the way users or applications access those files. Policies can be created that will move files based on date created, last time accessed, owner, size, or name. These files will be dynamically moved without having to be taken offline. Because Cluster File System can span multiple volumes, applications/databases can be stored on multiple tiers of storage hardware. This move is completely transparent to the users and applications that own the files—they don’t need to know the files have moved. As data is moved to different storage tiers (leveraging declining disk prices and new technologies—e.g., SATA), the policies are centrally managed and dynamic, and support a heterogeneous server and storage infrastructure that require no application, database, or backup/recovery policy modifications.

### Ensure data integrity

When multiple systems/nodes have access to data via shared storage, the integrity of the data depends on internode communication ensuring that each node is aware when other nodes are writing data. When the coordination between the nodes fail, it results in a “split brain” condition—a situation in which two servers try to independently control the storage, resulting in application failure or data corruption. I/O fencing is the method of ensuring the integrity of critical information by preventing data corruption, and Veritas Storage Foundation Cluster File System has implemented I/O fencing using the industry-standard SCSI-3 persistent group reservation technology. This allows a set of systems to have temporary registrations with the disk and coordinate a write-exclusive reservation with the disk containing the data. With I/O fencing, Symantec ensures that errant nodes are “fenced” and do not have access to the shared storage, while the eligible node(s) continue to have access to the data.

---

### Product Highlights

- **Flexible management**—Offer intuitive and centralized Web, Java™, or command line interface options for local or remote management
- **Single file system schema**—Manage all applications as if they resided on a single server
- **Heterogeneous hardware support**—Support multiple server hardware platforms and most storage subsystems
- **CIO support**—Support Concurrent I/O (CIO) to provide significant performance improvement for applications that utilize CIO (e.g., DB2)

- **Range locking**—Allow parallel access to the same file from multiple nodes, improving I/O performance
- **Scale out**—Simplify application scalability as new nodes adopt cluster-wide configurations
- **Dynamic Multi-pathing**—Provide I/O load balancing and availability throughout the storage area network

---

### Supported operating systems

- IBM® AIX®
- HP-UX
- Sun™ Solaris™
- SUSE Linux
- Red Hat® Linux

---

### Related products

- Veritas Storage Foundation by Symantec
- Veritas Storage Foundation for Databases by Symantec
- Veritas Storage Foundation for Oracle® RAC by Symantec
- Veritas Cluster Server by Symantec



Data Sheet: Storage Management  
Veritas Storage Foundation Cluster File System by Symantec

**More information**

*Visit our Web site*

<http://enterprise.symantec.com>

*To speak with a Product Specialist in the U.S.*

Call toll-free 1 (800) 745 6054

*To speak with a Product Specialist outside the U.S.*

For specific country offices and contact numbers, please visit our Web site.

*About Symantec*

Symantec is the world leader in providing solutions to help individuals and enterprises assure the security, availability, and integrity of their information. Headquartered in Cupertino, Calif., Symantec has operations in more than 40 countries. More information is available at [www.symantec.com](http://www.symantec.com).

*Symantec World Headquarters*

20330 Stevens Creek Boulevard

Cupertino, CA 95014 USA

+1 (408) 517 8000

1 (800) 721 3934

[www.symantec.com](http://www.symantec.com)

