



---

### Highlights :

- Increase performance and scalability of your IBM AIX and Linux and Windows™ systems
  - Maintain high data availability with virtually disruption-free expansion and migration
  - Improve efficiency through enterprisewide information sharing
  - Ease data management and enhance administrative control using information lifecycle management toolset
  - Simplify file system administration with extensible management and monitoring infrastructure
  - Gain greater return on investment and grow affordably with better hardware and infrastructure utilization
- 

# Intelligent storage management with IBM General Parallel File System

*Optimize storage management and boost your return on investment*

The explosive growth of data, transactions and digitally aware devices is straining IT infrastructure and operations, while storage budgets shrink and user expectations continue to multiply. As infrastructures support more users and become overburdened with storage, availability of data at the file level may degrade—and data management becomes more complicated.

In the past, organizations have addressed data management challenges by clustering servers or using network attached storage. But clustered servers can be limited in scalability and often require redundant copies of data. Traditional network attached storage solutions are restricted in performance, security and scalability. A single file server cannot scale, and even a roomful of file servers is not flexible enough to provide the dynamic, 24x7 data access required of a data-intensive computing environment. To overcome these issues, you need to look at a new, more effective approach to managing data.

The IBM General Parallel File System (IBM GPFS™) can help you move beyond simply adding storage to optimizing data management. GPFS is a high-performance, shared-disk file management solution that can provide faster, more reliable access to a common set of file data. Enabling a view of distributed data with a single global namespace across platforms, GPFS is also designed to provide:

- Online storage management
- Scalable data access through tightly integrated information lifecycle tools capable of managing petabytes of data and billions of files
- Centralized administration
- Shared access to file systems from remote GPFS clusters
- Improved storage use




## Handling the explosive growth of digital information more efficiently

GPFS provides several essential services to allow you to effectively manage growing quantities of unstructured data. GPFS leverages its cluster architecture to provide quicker access to your file data. File data is automatically spread across multiple storage devices, providing optimal use of your available storage to deliver high performance.

GPFS differs significantly from Network File System (NFS). With GPFS, there is no single-server bottleneck or protocol overhead for data transfer. GPFS takes file management beyond a single system by providing scalable access from multiple systems to a global namespace. GPFS interacts with applications like a local file system but is designed to deliver higher performance, scalability and fault tolerance by allowing access to the data from multiple systems directly and in parallel.

GPFS allows multiple applications or users to share access to a single file simultaneously while maintaining file-data integrity. For example, multiple animators or editors can work on different parts of a single video file from multiple workstations at the same time. This capability helps you reduce the amount of storage and management overhead required to maintain several copies of the source file. GPFS uses fine-grained locking based on a sophisticated scalable token (lock) management system to help ensure data consistency during concurrent access and prevent multiple applications or users from updating the same portion of a file at the same time.



Expanding beyond a storage area network (SAN), a single GPFS file system can be accessed by nodes using a TCP/IP or InfiniBand connection. Using this block-based network data access, GPFS can outperform network-based sharing technologies like NFS—and even local file systems such as EXT3 journaling file system for Linux® or Journaled File System.

## Providing greater data access to optimize performance for demanding applications

GPFS is designed for high-performance parallel workloads. Data and metadata flow from all the nodes to all the disks in parallel under control of a distributed lock manager. It has a very flexible cluster architecture which enables the design of a data storage solution that not only meets current needs but can also quickly be adapted to new requirements or technologies. GPFS configurations include direct-attached storage, network block input and output (I/O) or a combination of the two, and multisite operations with synchronous data mirroring.

Applications requiring the highest per node throughput often deploy the direct-attached model, where disks are physically attached to all nodes in the cluster that accesses the file system. The direct-attached model requires an operating system device interface to the storage, for example, using a Fibre Channel SAN or InfiniBand connection.

To optimize the cost of your storage infrastructure, deploy a large number of systems or seamlessly integrate multiple platforms and interconnect storage technologies, the network block I/O (also called network shared disk [NSD]) model can be used. Network block I/O is a software layer that transparently forwards block I/O requests from a GPFS client application node to an NSD server node to perform the disk I/O operation, and then passes the data back to the client. Using a network block I/O configuration can be more cost effective than a full-access SAN. And you can use this type of configuration to tie together systems across a wide area network (WAN). You can optimize the configuration of clustered systems to attached storage for your specific mix of workloads. You can also develop a data workflow within one location or across the WAN to help save time and the cost of managing file data.

To exploit disk parallelism when reading a large file from a single-threaded application whenever it can recognize a pattern, GPFS intelligently prefetches data into its buffer pool, issuing I/O requests in parallel to as many disks as necessary to achieve the peak bandwidth of the underlying storage-hardware infrastructure. GPFS recognizes multiple I/O patterns, including sequential, reverse sequential and various forms of striped access patterns. In addition, for high-bandwidth environments like digital media, GPFS can read or write large blocks of data in a single operation,

minimizing the overhead of I/O operations. GPFS has been proven in applications with aggregate data and bandwidth requirements that exceed the capacity of typical shared file systems.

### **Configuring file systems for high availability and reliability**

For optimal reliability, GPFS can be configured to eliminate single points of failure. The file system can be configured to remain available automatically in the event of a disk or server failure.

A GPFS file is designed to transparently fail over token (lock) operations and other GPFS cluster services, which can be distributed throughout the entire cluster to eliminate the need for dedicated metadata servers. GPFS can be configured to automatically recover from node, storage and other infrastructure failures.

GPFS provides this functionality by supporting data replication to increase availability in the event of a storage media failure; multiple paths to the data in the event of a communications or server failure; and file system activity logging, enabling consistent fast recovery after system failures. In addition, GPFS supports snapshots to provide a space-efficient image of a file system at a specified time, which allows online backup and can help protect against user error.

GPFS offers time-tested reliability and has been installed on thousands of nodes across industries, from weather research to broadcasting, retail, financial industry analytics and web service providers. GPFS also is the basis of many cloud storage offerings.

### **Simplifying data management with advanced tools**

GPFS provides an information lifecycle management (ILM) toolset that helps simplify data management by giving you more control over data placement. The toolset includes storage pooling and a high-performance, scalable, rule-based policy engine.

Storage pools enable you to transparently manage multiple tiers of storage based on performance or reliability. You can use storage pools to transparently provide the appropriate type of storage to multiple applications or different portions of a single application

within the same directory. For example, GPFS can be configured to use low latency disks for index operations and high capacity disks for data operations of a relational database. You can make these configurations even if all database files are created in the same directory. In addition, storage pools can simplify storage management by allowing you to transparently retask older storage systems, moving them from a tier 1 role to tier 2.

As files are created, policy-driven automation places file data in the appropriate storage pool based on business rules you define, regardless of where the file is placed in the directory structure. Once files exist, the policy engine can move files among storage pools or tape or delete them after a specified time period, or provide reports on the contents of the file system.

The high-performance policy engine leverages the parallel metadata infrastructure of GPFS to provide highly scalable policy processing, which allows you to actively manage billions of files across multiple storage tiers.



### **Easing file system administration**

GPFS provides a common management interface that simplifies administration even in very large environments. The administration model is easy to use and consistent with standard IBM AIX® and Linux® file system administration. Cluster operations can be managed from any node in the GPFS cluster. These commands support cluster management and other standard file system administration functions including user quotas, snapshots and storage management.

The ability to run multiple versions of GPFS in the same cluster means you can do rolling upgrades of operating systems and GPFS versions. Rolling upgrades can provide a near-zero

downtime method for upgrading cluster nodes, one node at a time, allowing application file data access throughout the upgrade process.

Heterogeneous servers and storage systems can be added to and removed from a GPFS cluster while the file system remains online, further simplifying management and helping enable 24x7 operations. When storage is added or removed, the data can be dynamically rebalanced to maintain optimal performance.

## Helping reduce the cost of managing your storage environment

GPFS can reduce data duplication and make more efficient use of storage components by combining isolated islands of information into a centralized, high-performance storage infrastructure. It also can help improve server hardware utilization by allowing dynamic storage access to all data from any node. By improving storage use, optimizing process workflow and simplifying storage administration, GPFS lets you take a multipronged approach to reducing storage costs.

## Summary

GPFS already is managing data volumes that most companies will not need to support for five years or more. You may not need a multipetabyte file system today, but with GPFS you know you will have room to expand as your data volume increases. GPFS can scale to meet the dynamic needs of your business, providing a solid application infrastructure with NFS file serving capabilities.

GPFS provides world-class performance, scalability and availability for your file data and is designed to optimize the use of storage, support scale-out applications and provide a highly available platform for data-intensive applications. A dynamic infrastructure that includes GPFS can help your company streamline data workflows, improve service, reduce costs and manage risk—delivering real business results today while positioning you for future growth.

## More information

To learn more about IBM GPFS, contact your IBM representative, IBM Business Partner or visit the GPFS website:

[ibm.com/systems/gpfs](http://ibm.com/systems/gpfs)



---

© Copyright IBM Corporation 2010

IBM Corporation  
Marketing Communications  
Systems Group  
Route 100  
Somers, NY 10589  
U.S.A.

Produced in the United States of America  
July 2010  
All Rights Reserved

IBM, the IBM logo, [ibm.com](http://ibm.com), AIX and GPFS are trademarks or registered trademarks of International Business Machines Corporation in the United States, other countries, or both. If these and other IBM trademarked terms are marked on their first occurrence in this information with a trademark symbol (® or ™), these symbols indicate U.S. registered or common law trademarks owned by IBM at the time this information was published. Such trademarks may also be registered or common law trademarks in other countries. A current list of IBM trademarks is available on the web at "Copyright and trademark information" at [ibm.com/legal/copytrade.shtml](http://ibm.com/legal/copytrade.shtml).

Linux is a registered trademark of Linus Torvalds in the United States, other countries, or both.

Microsoft, Windows, Windows NT, and the Windows logo are trademarks of Microsoft Corporation in the United States, other countries, or both.

InfiniBand is a trademark and/or service mark of the InfiniBand Trade Association.

Other company, product, and service names may be trademarks or service marks of others.



Please Recycle

---