

Simple file system access to massively scalable storage is what customers demand to meet the challenge of rampant growth that is stretching the capabilities of existing file systems.

When the first file system hierarchy was implemented the sheer volume of file-based data being generated today was unforeseen. Organizations are seeking a way to simplify and ease the cost of managing greater amounts of file system data that delivers performance, scalability and protection for both active and archive files.

The solution is here.

## The Business Benefits of CASTor

### Simple integration with legacy applications

Applications that use common internet file systems (CIFS) or network file system (NFS) protocols can leverage the advantages of CASTor without application or business process modifications. Implementing CASTor requires minimal changes in your existing IT environments.

### File system scalability beyond compare

CASTor FSG presents a file system view on the front end with a massively scalable storage cluster and single-address space on the back end allowing it to scale with business demands.

### Reduce administrative overhead and cost

Zero provisioning or configuration is needed to add new storage capacity to the file system and also eliminates management of "hot spot" issues encountered by other file systems.

### Seamlessly upgrade without impact

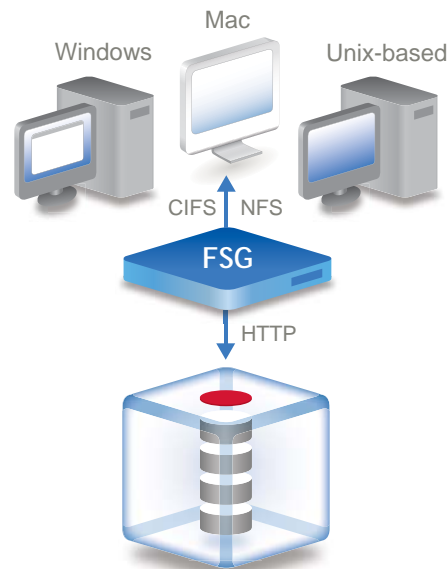
Upgrade storage hardware without application downtime or the need to migrate files.

### Reduce backup time or eliminate altogether

Two or more replicas of all files are stored in the CASTor cluster ensuring continuous data availability and minimizing backup time and recovery.

### Regulatory compliance

Set lifecycle policies for retention and protection with WORM capability and content authentication. Realize a single tier of storage for active and archive files.



## CASTor™ FSG

### Massively Scalable File System Gateway

The file system has a long history and is firmly entrenched in applications that use them to store unstructured, non-database data, or simply files. Besides a database, the file system is the next primary means for working with disk storage. Traditional block storage creates significant challenges for file systems that must manage millions to billions of files, and Terabytes to Petabytes of capacity. Unlike other file systems that must manage multiple individual blocks for every file, FSG's relational database only needs to manage a single unique identifier (UUID).

CASTor FSG was designed to provide applications the familiar file system interface to the CASTor content storage cluster. Classic file system mapping provides a virtual file system front end to CASTor that users can access through standard network file systems and protocols. The CASTor cluster delivers high performance, scalable and reliable storage that works with files as content objects instead of blocks.

simplicity  
performance  
scalability  
integrity  
value



## CAStor FSG Features

### Policy-based File System

Policies can be set on a per share basis, giving the user control over retention, number of replicas and phases of the file lifecycle.

### CIFS and NFS Support

The scalability and throughput of CAStor become accessible to applications without the need to change code.

### Self Managing

CAStor automatically manages all replication, policies, integrity and storage from the time a file is written to storage.

### Massive Scalability

There is no hard limit to the number of files or space capacity. Multiple CAStor FSG systems can run in parallel on top of the same CAStor Content Storage Cluster.

### High Performance File System

Multi-threaded with connection pooling for scalability and performance, it automatically takes advantage of the multiple contact points in a CAStor storage cluster as it scales up. Connections to each node are cached and multiple threads are spun on each connection.

### Enhanced Metadata Management

The database in CAStor FSG contains standard file metadata, mapping between the file name and Universally Unique Identifier (UUID) in CAStor as well as additional useful information.

### Multiple Virtual Views

CAStor FSG can surface multiple virtual views into a single content store. Content stored through FSG can be accessed via applications and UUIDs interchangeably.

### Authentication

CAStor FSG can operate in a standalone fashion or may be integrated with your existing user directories with standard authentication protocols such as NTLMSSP, SPNEGO, NTLMv2, Kerberos and Active Directory Passthrough at the group and user levels.

## Recommended Hardware Configuration

CAStor FSG will run on a variety of vendor hardware with the following recommended configuration:

CPU	Intel Xeon, AMD Athlon64, or equivalent
Internal HDD	Standard SATA or SCSI with RAID configuration for redundancy
Memory	2GB RAM*
Network Interface	Gigabit Ethernet
CD Drive	Required for installing FSG software
Operating System	Ubuntu Linux Version 7.10 (Recommended) or DebianLinux 4.0
Database	MySQL 5, Oracle 10g Enterprise, or Oracle 10g Express Edition, either on-board or in separate server
CAStor™ Content Storage Cluster	Storage cluster of 3 or more server nodes

\* Number of file shares implemented may require additional memory.

E-mail [info@caringo.com](mailto:info@caringo.com) or visit [www.caringo.com](http://www.caringo.com) for additional information on our company and our products.

Caringo, Inc.  
901 South MoPac Expressway  
Building II, Suite 410  
Austin, Texas 78746

