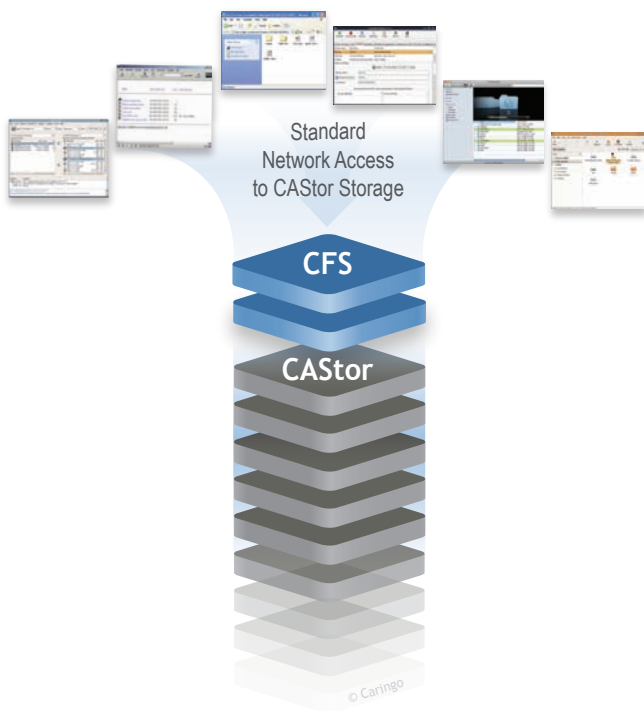


Product Brief

Flexible and Scalable File Storage

When the first file system hierarchy was implemented the sheer volume of file data being generated today was unforeseen. Up to 90% of new data being created is unstructured or file-based and today's demands require file systems to manage tens of millions of files and hundreds of terabytes of capacity. Organizations are challenged with file storage solutions that are becoming increasingly complex, more expensive to manage, difficult to backup and more challenging to maintain. Too much file data is stored in primary file stores and regulatory compliance means retaining this data for longer periods of time. A new approach to storage is needed. Companies need to store files more efficiently, more economically, and with less effort and guaranteed integrity.



CAStor™ Content File Server

File system interface to object-based storage. The file system metaphor and folder/sub-folder interface has a long history and is firmly entrenched with end users and in applications. Traditional file systems and NAS technologies built on block storage and i-node infrastructure are facing significant challenges for managing the massive volume of files and capacity required today. CAStor Content File Server (CFS) is unlike other file systems that must manage the complexity of multiple blocks for every file. CFS provides location independent access to object storage through a familiar file system interface. By separating the data from its' file system hierarchy, CFS leverages CAStor's object storage platform to store and protect your data in efficient and self-managing object-based storage clusters. With CAStor CFS, performance never decreases as more data is added and limits on file size and number of files are never a problem.

Supports standard file protocols. Using standard protocols, CAStor CFS interfaces to CAStor object-based storage clusters that deliver high performance, scalable and reliable storage. The CAStor and CAStor CFS solution is specifically designed to efficiently store active file-based data with all the functionality required for short or long-term archiving.

Business Benefits

Reduce or eliminate file system sprawl. Scale to billions of files and petabytes of capacity without adding new file systems.

Reduce administrative overhead and cost. No provisioning or configuration needed to add new capacity. Eliminates administrative effort to mitigate "hot spots" encountered by 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 cluster ensuring continuous data availability and minimizing backup and recovery time.

Regulatory compliance is built in. Set life cycle policies for retention and protection with WORM storage and file authentication.

CAStor™ Content File Server

Features

Standard File System Interface

Linux file system supports NFS, CIFS, FTP and WebDAV simultaneously

Massive Scale-out File Storage

Scale from millions to billions of files and terabytes to petabytes of capacity in a single file system

Seamless Scalability

Add capacity without having to provision, configure and present new storage to the file system

Cluster CFS Nodes

File system scales by clustering additional CFS servers for performance across common namespace and storage

Namespace

Maintains and maps file system metadata to stored objects

Timescapes (Continuous Snapshots)

Enables administrators to mount and view the contents of a file system at a particular point in time and restore content

File Level Permissions or ACL's (Access Control Lists)

Supports file-level read, write and execute permissions for users and groups

Manage and Define Metadata

Set metadata values at the mount point to automatically manage retention, replication and distribution

Automated Backup of CFS Namespace

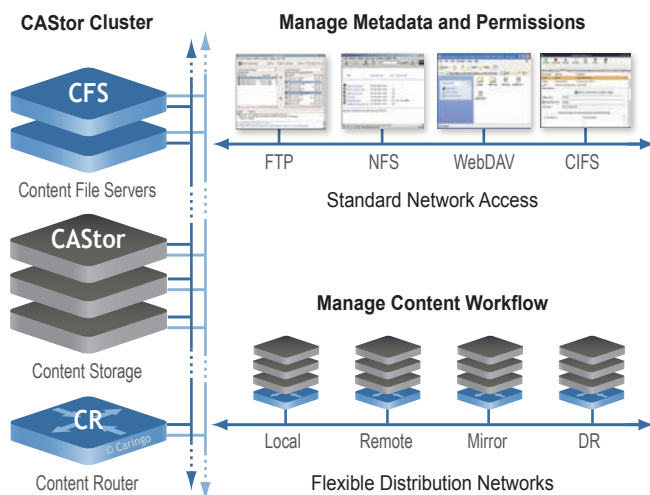
CFS namespace is automatically backed up to the CAStor cluster and replicated for recovery

Integration with External Namespace

Supported and defined API allows content to be ingested from other sources into a shared repository

Policy-Based Replication and Distribution

CAStor Content Router can be configured to automatically replicate and/or distribute files to one or more remote locations



Recommended Infrastructure

- ▶ Network Time Protocol (NTP) service
- ▶ Gigabit Ethernet connection speed

Required OS and Software

- ▶ 64-bit Ubuntu Linux – version 10.04 server
Coming soon: Red Hat Enterprise Linux
- ▶ CAStor Content Storage Software (server)

Recommended Hardware Configuration

CPU	64-bit, Intel Xeon, AMD Athlon 64, or equivalent
Hard Drive	SATA, SAS, or SCSI, with RAID configuration for failover
Other Drive	CD ROM to install the operating system
Memory	2GB minimum, additional mounts require additional RAM
Network Interface	Dual Gigabit Ethernet
CAStor Content Storage Cluster	Storage cluster of 3 or more server nodes

About CAStor

Object-based storage software enabling simple, scalable and reliable storage clusters for active and archive content on affordable, x86 server nodes. CAStor delivers the most complete set of features to address the explosion of digital content.

About CAStor Content Router

Intelligent and policy-based content distribution and replication across distributed sites is provided by filtering metadata to drive a robust rules-based workflow engine.

About CAStor Content File Server

Caringo's CAStor object-based storage software requires minimal administration and management. Due to the nature of file systems, CFS requires more administrative effort than CAStor including Linux expertise and experience with SAMBA and NFS file protocols. Contact us regarding requirements and the in-house expertise necessary to successfully implement CAStor CFS in your storage environment.

To learn more visit our web site or e-mail info@caringo.com.

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