

Multi-site Global File System

Combines File Shares Distributed Across Campuses and Hybrid Clouds for Easy Access, Optimal Capacity Use, Scalable Throughput, and High Availability

vFiLO BENEFITS

- Increased visibility and control of scattered data
- Complete hardware-independence and flexibility
- High operational efficiency
- Simplified file access, sharing, and collaboration
- Non-disruptive data migration between NAS/ filers and cloud/object storage
- High data availability
- Lower storage costs

The rapidly growing number and diversity of storage systems and locations make file sharing, collaboration, and data placement increasingly challenging. Plus, the recurring restructuring of file shares every time space runs low upsets users and disrupts applications.

DataCore vFiLO solves these problems by creating a scale-out global file system across distributed sites spanning on-premises and cloud-based file shares. Previously isolated folders spread across different systems roll up under a single global namespace for convenient access from any location via NFS and SMB protocols. The vFiLO software pools resources from discrete filers at each site, making optimal use of their capacity and horsepower.

vFiLO continually load balances, safeguards, and migrates files between active primary filers and secondary S3 object storage based on site-specific policies. The software-defined architecture offers the unique flexibility to incorporate existing NFS NAS and file servers into virtual pools. The filers can be non-disruptively expanded and replaced over time with new devices of your choosing based on cost, performance, and other preferences.

USE CASES: HOW vFiLO CAN HELP YOUR ORGANIZATION



Pool resources from distributed NAS and file servers for global access and optimal utilization



Offload inactive data from premium filers to lower-cost object/cloud storage



Scale out NAS to distribute loads, improve responsiveness, and ensure high availability



Collaborate between sites through selective file sharing and replication



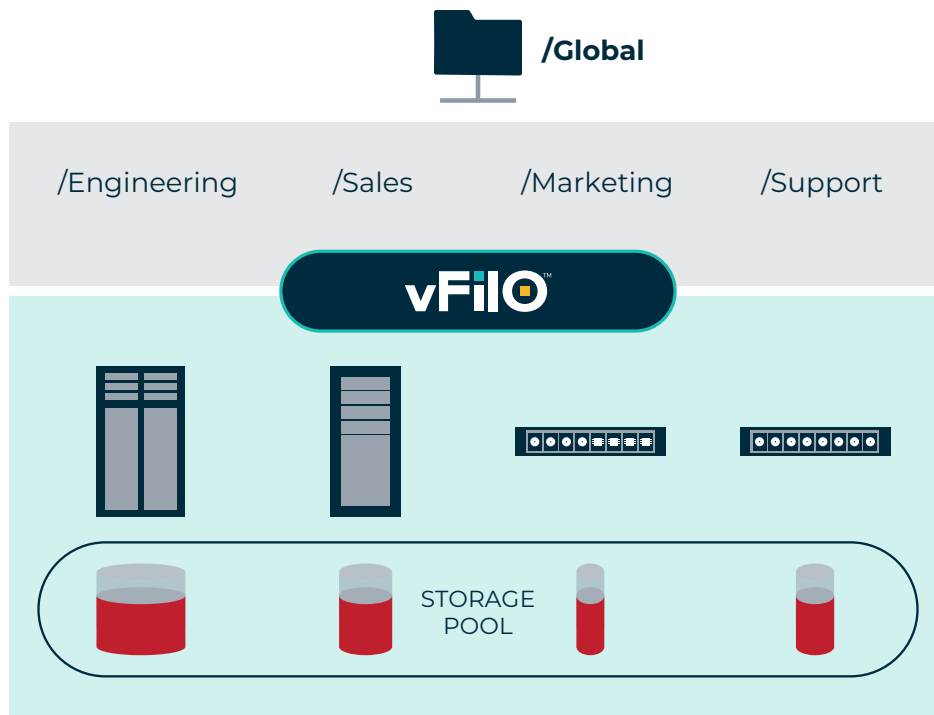
Replicate data between disparate filers to enhance BC/DR readiness and response



Switch workloads to remote sites during peak loads and planned downtime

POOL CAPACITY FROM ON-PREMISES NAS DEVICES AND FILE SERVERS

Rather than incurring long, stressful data migrations normally associated with technology transitions, vFiLO layers on top of existing NAS devices and file servers in the same campus, aggregating their resources under a single global file share. The combined capacity and horsepower of the virtual storage pool can then be used to balance the load, enhance data availability, and tier data placement.



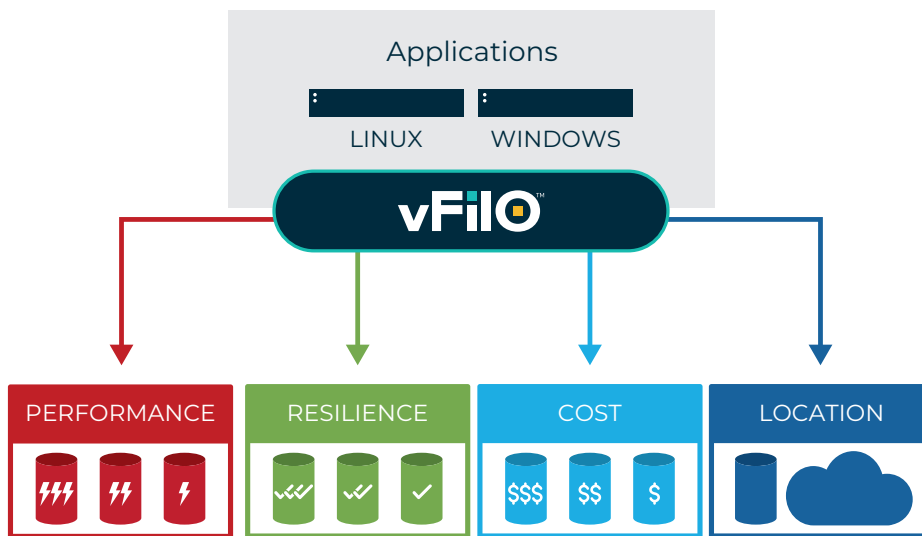
Pool storage capacity across diverse NAS devices and filers

The initial step called “assimilation” gathers the metadata describing folder hierarchies, ownership, permissions, and file locations from each filer to create a global catalog. That catalog is kept separate from the actual file contents, effectively de-coupling how data is organized from where it is stored. vFiLO then exports the original shares as subfolders under a global mountpoint (for example, /Global/Engineering). Clients momentarily disconnect from the filers and remount the globally-accessible shares from the vFiLO portal.

Applications and users continue accessing their data from the same familiar folder paths. Now, vFiLO is free to non-disruptively replicate and relocate files in the background as conditions and policies dictate. Normal business operations continue undisturbed even when adding new hardware or decommissioning legacy equipment. For utmost data protection and uninterrupted access, fully redundant configurations with multiple replicas of critical files may be configured.

AUTOMATIC DATA PLACEMENT

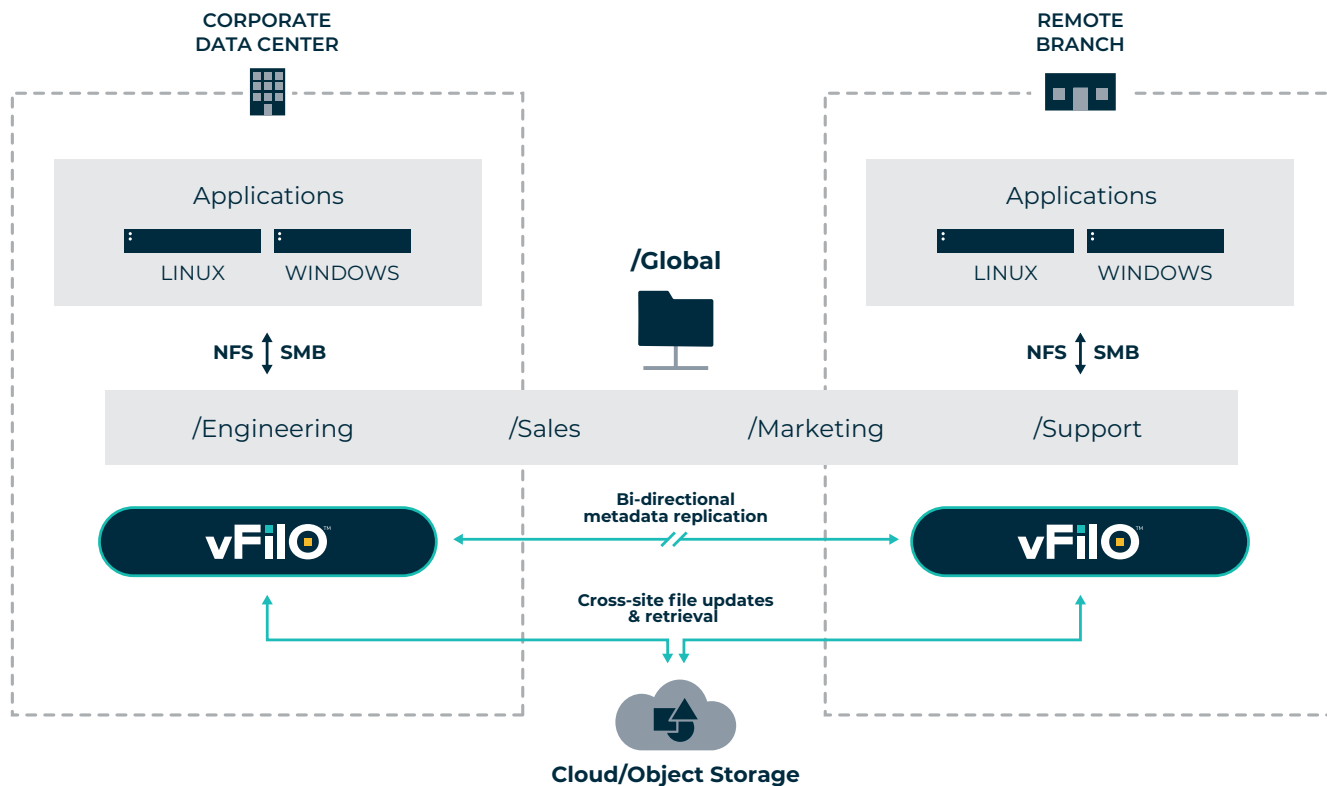
vFiLO automates the near-impossible task of juggling where files should be relocated to in order to satisfy business intentions, despite constantly changing conditions. The system administrator simply sets a few high-level objectives that guide how files meeting specific (or broad) criteria should be treated. Parameters including file type, aging, access frequency, ownership, and origin can drive data placement. These parameters help vFiLO choose between fast and cheap storage, on-premises and the cloud, or the number of copies necessary to meet resiliency, performance, and governance goals. Its AI/ML algorithms regularly sweep the metadata and assess the fluctuating state of the hardware in order to align files with those goals. Even when archiving files from active primary tiers to secondary S3-compatible cloud/object storage, the directory structure remains intact. There’s no need for IT intervention to retrieve them.



vFilo uses AI/ML to automate file placement based on administrator-defined business objectives

MULTI-SITE GLOBAL NAMESPACE

The vFilo global file system extends beyond a single campus. Once-isolated shares housed in different facilities are combined under a single namespace, making file sharing and collaboration extremely easy. Regularly-accessed files may be stored locally for the fastest response, whereas infrequently used ones appear local, but are retrieved behind the scenes from remote locations. This approach significantly reduces the capacity required. Files transferred between sites are also deduplicated and compressed in a cloud or object storage intermediary to reduce data transmission and space consumed. Only metadata is synchronized regularly to ensure that all users have the latest view of the global catalog.



Access files stored in multiple locations from a single namespace for effective inter-site collaboration

KEY FEATURES / DATA SERVICES

Several hands-free, file-granular services provided by vFileO dynamically govern data mobility, durability, and availability. They include synchronous mirroring between active filers on the same cluster, asynchronous replication between sites and to cloud/object storage, automatic data migration and rapid snapshots/clones. Archives placed on cloud/object storage are globally de-duplicated and compressed, yet remain accessible, as do recently deleted files.

CONSUMERS			
END USERS	APPLICATION & WEB SERVICES		DEVICES
ACCESS METHODS			
NFS		SMB	
OPERATION & INSIGHTS	DATA SERVICES		COMMAND & CONTROL
EXTENSIBLE METADATA	MULTI-SITE GLOBAL NAMESPACE	ACTIVE ARCHIVE	ACCESS CONTROLS
DATA MIGRATION	PARALLEL NFS	ASYNCHRONOUS REPLICATION	CLI
HISTORICAL / REAL-TIME CHARTS	POOLING, ASSIMILATION OF NAS/FILE SERVERS	AUTOMATED DATA PLACEMENT	CONSOLE
HEALTH & PERFORMANCE GRAPHS	SELF-SERVICE UNDELETE	DEDUPLICATION / COMPRESSION*	FILE GRANULARITY
ALERTS	SNAPSHOTS / CLONES	ENCRYPTION*	PLUG-INS
PROVISIONING	SYNCHRONOUS MIRRORING	LOAD BALANCING	
SUPPORTED STORAGE			
FILE	OBJECT	BLOCK	CLOUD

*FOR INACTIVE FILES PLACED ON OBJECT/CLOUD STORAGE

BENEFITS



Visibility & Control

- Organize widely scattered files under a single namespace for convenient access by apps and users
- Measurably improve productivity and collaboration across sites with easier file sharing
- Fulfill data governance obligations through policies for data placement and protection



Ultimate Flexibility

- Integrate new technologies into your infrastructure without expensive forklift upgrades
- Benefit from a location and device-independent approach to expansion and modernization
- Place data where it makes most sense, traversing sites and hardware architectures



Efficiency & Simplicity

- Fully exploit available storage assets by pooling their resources and load balancing across them
- Recover time lost from manually shuffling files and backing them up
- Avoid disruptive and time-consuming data migrations

Compelling Economics

- Save costs by transparently relocating inactive files, snapshots, and replicas on lower-cost object storage, where they are deduped and compressed
- Modernize while maximizing the value from existing equipment without having to rip out and replace perfectly good assets
- Defer buying more premium filers by fully utilizing existing resources on the active primary filers that deserve it
- Avoid consuming costly space and bandwidth by performing discrete functions on specific files rather than on volumes or entire folders / file shares

DEPLOYMENT

vFileO clusters at each site are comprised of three key components:



Existing NAS Devices and File Servers

Provide the original segregated file shares. They are mounted and accessed by vFileO over NFS.



Metadata Service Nodes (Anvil servers)

Control the cluster, its administrative interface, metadata operations, and assimilation. Anvil nodes are deployed as a pair for redundancy.



Data Service Nodes (DSX servers)

Responsible for all data services (synchronous mirroring, replication, data placement, snapshots/clones, etc.). They provide redundant portals for data access and also serve as on-premises file storage. The number of data service nodes can be scaled up or down based on business needs.

The clusters may be joined to an Active Directory domain for comprehensive access controls.

Scalability

- Scales from 50 TBs to multiple petabytes with billions of files in a single namespace
- Scales up and out to 40 data service nodes per site
- Up to eight sites can participate in the global file system

LICENSING

Benefit from simple, transparent, and flexible licensing based on managed capacity available to vFileO from active file servers, on-premises object storage, and public cloud S3 storage.

Pricing includes 24x7 Premier Support and software updates.

MINIMUM HARDWARE / SOFTWARE REQUIREMENTS

	METADATA SERVICE		DATA SERVICE	
	VIRTUAL MACHINE	BARE METAL	VIRTUAL MACHINE	BARE METAL
CPU Cores	8 vCPU	8+ cores Intel based 2.5+ GHz processor(s)	4+ vCPU	4+ cores Intel based 2.4+ GHz processor(s)
Memory	16+ GB (+ 1GB per vFIO share)	16+ GB (+ 1GB per vFIO share)	8+ GB	8+ GB
Boot Disk	200 GB	200 GB Hardware RAID recommended	100 GB	100 GB Hardware RAID recommended
Additional Disks	(2x) 400 GB backed by SDs Supports up to 20 million files	(2x) 400 GB NVMe recommended (Hardware RAID is not needed)	Only needed when Data Service instance is providing additional storage	Only needed when Data Service instance is providing additional storage
Network Adapter	(2x) 10 GbE (one link is dedicated to HA services)	(2x) 10 GbE (one link is dedicated to HA services)	(1x) 10 GbE	(1x) 10 GbE

0621



Discover the Ultimate Flexibility of DataCore Software

DataCore Software delivers the industry's most flexible, intelligent, and powerful software-defined storage solutions for block, file, and object storage, helping more than 10,000 customers worldwide modernize how they store, protect, and access data. With a comprehensive product suite, intellectual property portfolio, and unrivaled experience in storage virtualization and advanced data services, DataCore is The Authority on Software-Defined Storage. www.datacore.com

GET STARTED