MagFS:
The File System for the Cloud

Table of Contents

New technologies and social trends are tearing down geographic barriers and changing the way people interact and collaborate, both at work and at play. Within the enterprise, these shifts are fundamentally changing how people consume data. Traditional storage players have yet to address the challenges posed by the new user-centric consumption models, and newer market entrants focus only on pieces of the picture. Yet this is where most enterprise data is going. While traditional solutions will continue to define how data is stored and managed within the datacenter, Maginatics developed MagFS to align storage technology with the pressing needs of a highly-distributed and mobile enterprise and to address a variety of “hot-button” requirements such as NAS consolidation, in-cloud scale-out workloads driven by high-performance computing (HPC) and other parallel applications, and secure, scalable, enterprise-grade file sharing and collaboration.

Introduction

Enterprise data can be divided into two main categories based on the way it is consumed. One category, comprised mainly of structured data, lives and is consumed within the datacenter. For this data, the most important storage requirements are very low latency and high throughput.

The remaining enterprise data, all of it unstructured, is stored centrally but consumed by geographically-distributed users and applications. Unstructured data accounts for most enterprise data today and is growing must faster than structured data. Its exponential growth is driving the need for storage systems that are scalable, elastic and cost-effective. At the same time, these systems must enable easy access for an increasingly distributed and mobile user base. In short, the key storage requirements for this data category are capacity and accessibility.

Traditional enterprise storage solutions were never designed to meet these requirements. Instead, they were optimized to deliver outstanding speed, latency and throughput for local users and applications that share a connection to the same corporate LAN. This is not a failing of traditional storage, but a case where rapidly changing access patterns have created a different set of needs.

New storage vendors have sprung up to try to meet these needs for capacity and accessibility. While their numbers are large, they can be divided into two groups: online file-sharing solutions and cloud gateways. MagFS represents a third category that combines the advantages of the first two while overcoming their inherent limitations.

We will first consider how other solutions attempt to solve the problem and then explore more deeply the underpinnings of the Maginatics approach.

Traditional Network File Systems

While filers are not a “new” solution, it is useful to consider their place within the datacenter to set the context. Traditional network attached storage is optimized to give LAN-connected users reliable access to local data, not to provide efficient connectivity for users in remote offices or field locations, let alone from mobile devices. Third-party WAN acceleration or thin-client solutions can provide workarounds, but at considerable added cost and complexity. Moreover, even though scale-out NAS exists, the majority of filers were never meant to scale to very large data sets; most hit the wall at a few hundred million files.

Cloud Gateways

Cloud gateways focus on solving the capacity problem for IT organizations by providing cloud connectivity and, in some cases, file system semantics for use with the cloud. Cloud gateways are typically deployed as appliances that sit in-line with traffic in branch offices and present a monolithic cache that creates the perception of LAN-speed access to the cloud for local users.

While solving the problem of capacity for headquarters and branch office locations, cloud gateways suffer from some of the same accessibility limitations as traditional enterprise storage. Because it exposes a standard CIFS/SMB or NAS interface, a cloud gateway is no better than a standard filer at providing efficient access to data from outside the office where the gateway is deployed. Users logging in from home or the field will experience very poor performance. Moreover, like traditional storage, cloud gateways lack native support for mobile access. Some vendors provide a measure of mobile support

by bundling a Web app that connects mobile devices to the CIFS interface exposed by the gateway. Even when they add data compression and authentication, however, this approach cannot deliver the performance or data security demanded by enterprise IT.

Gateways share another limitation with traditional filers: they essentially swap one point of management (the array at the branch) for another point of management (the gateway itself). For large companies with very distributed operations, the cost of deploying and maintaining a gateway at each branch can offset the efficiency gains offered by the cloud.

Finally, gateways break the security model of large enterprises, whose IT organizations typically insist that the encryption keys must be hosted exclusively in the datacenter. Gateways require distributing encryption keys to all branches, which creates security concerns since few branch offices benefit from the same level of physical security as the datacenter.

Online File Sharing

Many online file sharing solutions are services external to the enterprise, and all are sync-based. Rather than provide a general-purpose storage system, sync solutions focus on allowing users to access shared data stored in the cloud from anywhere using any device. Sync was introduced and popularized as a consumer technology that gives individuals personal access to their own data across multiple devices. It has been applied to business and enterprise customers in the last several years, and many vendors have introduced what they term “enterprise-grade” solutions based on sync.

Sync is ideal for file sharing and collaboration as long as the size of the team and shared data set remains small. The reason is that the sync process itself is inherently inefficient and unscalable. Sync blindly pushes every change to all devices of all users sharing the

same namespace, regardless of whether any of given recipient actually needs the changes.

With large teams or data sets, sync's blind-push approach consumes vast amounts of network bandwidth and client disk space. Even worse, it can take a days or even weeks for the sync operation to complete when teams or active data sets are large. For sufficiently large data sets, the operation may *never* complete. Either outcome defeats the very purpose of sync, which is to enable collaboration.

Maginatics MagFS

MagFS combines the best aspects of traditional NAS, cloud gateways and sync tools while avoiding their limitations.

Like traditional NAS, MagFS provides standard file system semantics, strong data consistency and application compatibility. Like cloud gateways, MagFS is a general-purpose storage system that allows IT to securely consume cloud storage. Like sync solutions, MagFS enables access to enterprise data from remote locations and mobile devices and provides the workforce with simple file sharing and collaboration workflows. Yet MagFS offers a capability that none of the others can match; namely, near real-time collaboration with user mobility at enterprise scale (and without requiring appliances at the edge).

Key Advantages of Maginatics Architecture

A large percentage of today's enterprise workforce is mobile, and virtual team members can find themselves on opposite ends of the earth. Even with broadband becoming ubiquitous worldwide, users accessing data over distance will continue to suffer a performance penalty regardless of the quality of their wide area network – a penalty that impacts their productivity and ability to collaborate effectively.

Data sets are also becoming larger. This is not only true of user documents but also applies to electronic medical records, field survey data, security camera feeds and myriad other kinds of data produced and consumed far from where it is stored.

As more and more data is stored centrally and consumed over distance by distributed and mobile teams, demand grows for storage systems that can deliver cost-effective capacity on the back end and anywhere, any-device accessibility on the front end. Traditional storage vendors are not yet addressing this distributed consumption model. Maginatics fills this market gap by extending the file system all the way to endpoint devices, creating an enterprise-wide storage web that connects users and devices with centralized data under the firm control of IT.

Maginatics offers the following advantages over other approaches:

- **Highest security.** Files are broken into chunks, each of which is independently encrypted using its own key. Data is always protected with AES-256 encryption in-flight and at rest, and every transmission occurs over a secure SSL/TLS connection.

- **On-Premises Model.** Just as important as the *level* of data security is the security management *model*. Where are the enterprise encryption keys stored? How are authentication and authorization handled? What are the workflows for credentials management, access control and system administration? With MagFS, all of these follow the traditional enterprise model.

- **Efficient access to data without boundaries.** Whether data is consumed in or out of the office, Maginatics provides efficient data access via multi-layer adaptive caching, in-line de-duplication, compounded operations and other optimizations.

- **Scalable architecture.** A split data/metadata architecture allows MagFS to scale linearly with the number of active users. By separating the data plane from the control plane, the server component performs only lightweight metadata operations, removing any risk of a bottleneck. On the client side, each endpoint device accesses data independently from all others, again eliminating the pinch points that accompany any monolithic appliance approach.

- **Near real-time file sharing.** MagFS is a full-fledged distributed file system. File system semantics make all changes instantaneously available to all users, who never have to wait for files to sync. Not only can files be shared with other members of one's organization, they can also be shared with partners and other third-parties using the MagFS External Sharing feature.

- **Global inline de-duplication.** Every Maginatics share benefits from global per share, location-independent, chunk-based de-duplication. If multiple clients are writing from different parts of the world (not only from within the same office), a given chunk will be moved only once.

- **Strong consistency.** Full Windows file system semantics eliminate revision conflicts and help support multi-writer and multi-reader workloads.

- **Granular access control.** Access control can be set at the file, folder and share level and is tied directly to the enterprise Active Directory or other identity management system.

- **Leverage existing infrastructure.** MagFS does not require dedicated appliances or caches at the branches, though its architecture allows for the addition of a cache where it can be useful, such as in very large branches.

- **Reduced points of management.** Deploying appliances (physical or virtual) or storage arrays in branch offices creates additional points of management. MagFS requires no points of management outside the data center.

The Three MagFS Design Pillars

The advantages of MagFS described above are built on three main design pillars that together differentiate MagFS from other solutions.

Scalable, split control-plane/data-plane architecture

MagFS features a distributed client-server architecture in which there is a clean separation, both physical and logical, between metadata and data. The server component of the system, which resides in the enterprise datacenter, functions as the gatekeeper to data but hosts only metadata. Clients communicate with the server over the “control plane” to request permission to perform data operations. Once permission is granted, clients communicate directly with the object store to perform those operations using metadata provided by the server. This complete separation enables system scalability as well as enhanced security and control.

On-premises security and management model

The MagFS Server, which is the equivalent of a traditional NAS head, resides in the datacenter and benefits from the same physical security as any other critical IT asset. An integrated key management system allows the MagFS Server to manage the enterprise encryption keys, which are never persistently stored outside the datacenter. The MagFS Server also integrates natively with Active Directory and other identity management systems to deliver airtight access control based on existing user profiles. In general, MagFS uses the same security and management model, workflows and toolsets as any traditional storage system regardless of whether data is stored on-premises or in a public cloud.

End-to-End Acceleration

To support user mobility and cloud connectivity, MagFS provides sophisticated optimizations including inline de-duplication, compounded operations and multi-layer adaptive caching to deliver the perception of locality even when data is far away.

Architecture

Overview

MagFS is based on a highly-distributed client-server architecture. The server is packaged as virtual appliances deployed in the data-center, protected by the hypervisor's High Availability feature and managed by IT. A client component runs on every end-user device and it is managed by the end user or IT. Storage is provisioned via "shares," each providing a single, independent namespace. Administrators may create, and users may be given access to, an unlimited number of shares. MagFS scales linearly with the number of shares, as each is served by one or more Metadata Servers (MDS). Clients connect to shares via native apps available for all supported platforms, which include mobile devices, PC and Mac laptops and desktops, server and virtual machines.

As illustrated below, there are two deployment models for MagFS: on-premises and in-cloud. In each case, the MagFS system and all of its components remain the same – only the location of the components changes.

PROVISION OF TRADITIONAL ENTERPRISE FILE SERVICES

In this deployment model, the MagFS Server is located on-premises and the object store can be either on-premises or in the cloud (or both).The clients can be anywhere. This model allows the enterprise to centralize capacity in a cost-effective commodity pool and to provide NAS-like services to users and applications that are co-located with the capacity or geographically distributed. Regardless of where data is stored, metadata (user identities and encryption keys) are stored exclusively on-premises. Solutions enabled by this deployment option include Enterprise File Sharing and Collaboration, Branch Office Storage Consolidation and Distribution of Large Files over Slow Connections.

Deployment Model for Provision of Enterprise File Services

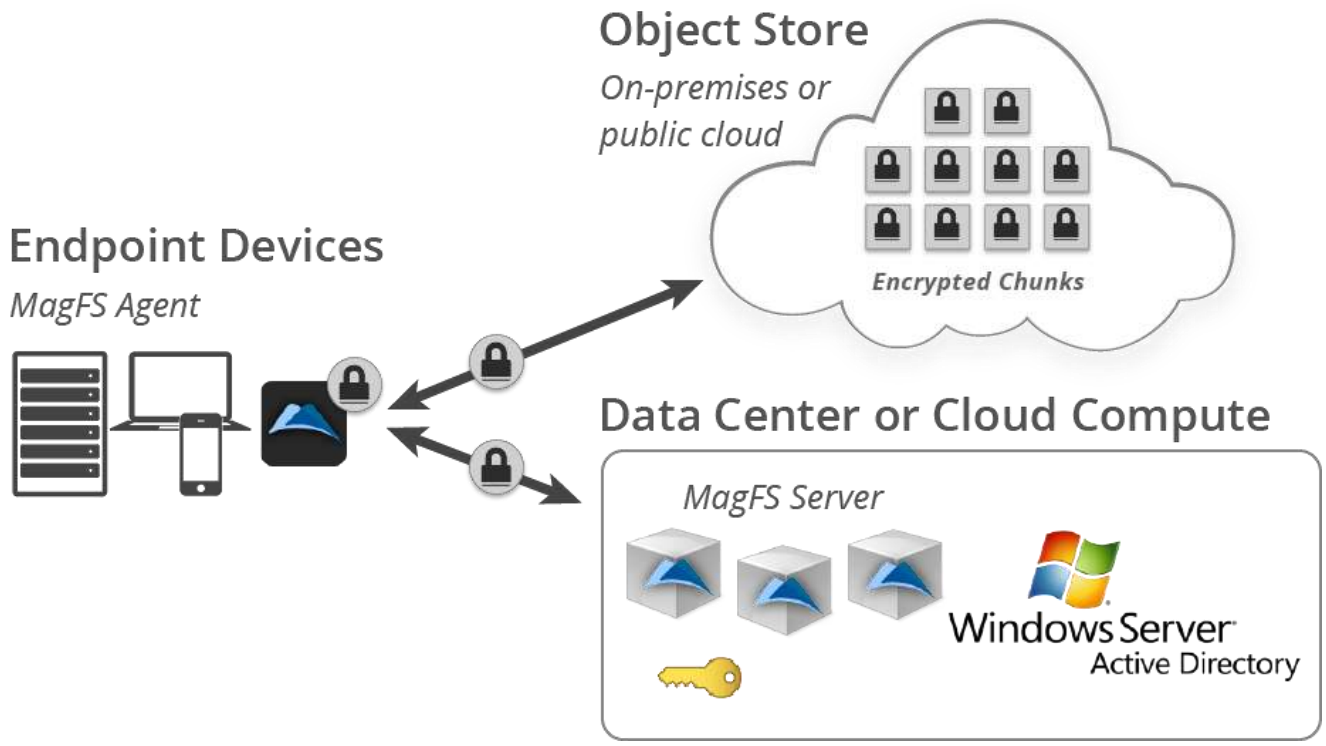

SCALE-OUT NAS IN THE CLOUD

In this deployment model, the MagFS Server, the object store and the clients are all co-located, either on-premises or in the cloud. This model allows users to build elastic, distributed NAS-like functionality on an object store. Applications that require NAS can now run against an object store without modification, whether in the data center or in the cloud. Parallel applications and workloads can benefit most from this deployment option and include those in the Life, Media and Entertainment (e.g., video rendering) and Large-scale Web server farms.

Deployment Model for Scale-Out NAS in the Cloud

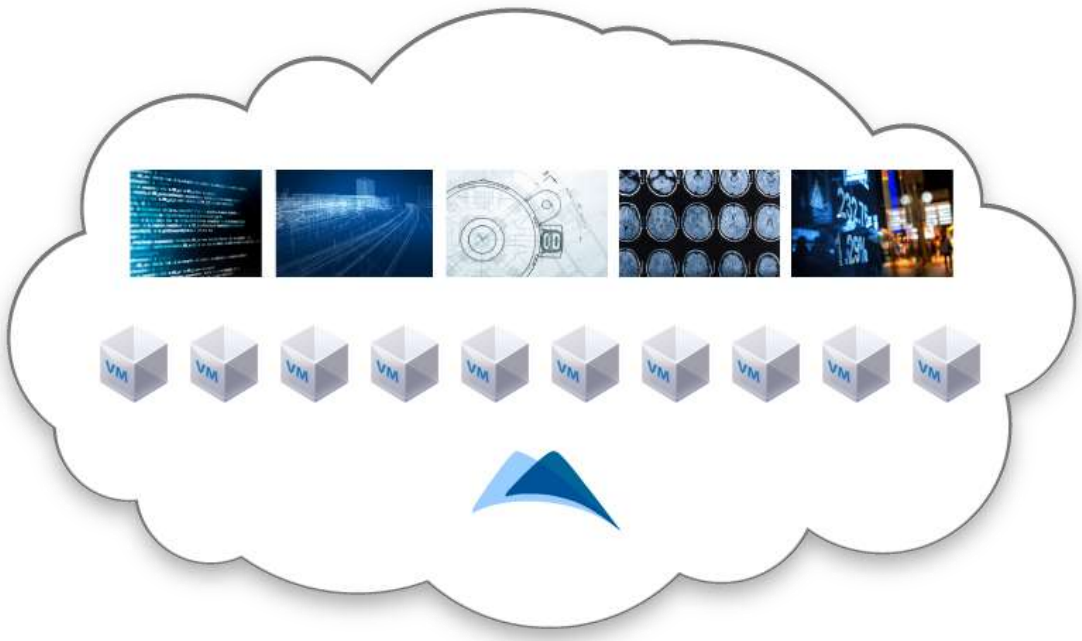

MagFS Components

Management Console

The management server provides a modern, Web-based console for managing a MagFS deployment. IT administrators can use the UI to provision, manage and monitor virtual filers and shares.

Metadata Server (MagFS Server)

In traditional storage systems, a single server (the “head” of the filer) typically becomes the bottleneck as the number of clients and/or storage capacity increases. Both data and metadata flow through the filer, increasing resource contention and requiring costly and complex hardware to provide good performance.

The main component of the MagFS Server is the MDS, which acts

as the “brain” for a given share. No file data ever flows through the MDS, allowing it to provide superior resource utilization and scalability. The MDS is responsible, among other things, for mediating all metadata operations as well as access control, data integrity, cache consistency and storing de-duplication chunk maps. Customers can add more capacity simply by augmenting a share via a new license key or by provisioning new virtual filers and shares.

Native Client (MagFS Client)

MagFS provides native clients for popular desktop and mobile platforms. The client User Interface is minimalist and non-intrusive – its goal is to remain largely out of the user’s way while providing tight integration with the underlying platform. A very large percentage of the client code base is common among all operating system platforms, enabling Maginatics to quickly add support for new platforms and OS versions.

To end users, a MagFS share is indistinguishable from a regular network share in appearance and functionality. Like any network share, MagFS is application-compatible, meaning existing applications can run against the share without modification. The MagFS

Client provides feature parity on all platforms while preserving platform-specific semantics -- for example, providing POSIX-style permissions on Mac and Linux but exposing Windows-style ACLs on Windows. A multi-layered, adaptive, encrypted local cache is used to accelerate performance and hide latency and bandwidth issues over distance. Clients use an efficient, bi-directional, firewall-friendly transport layer to communicate with the MDS. A single client can connect to multiple virtual filers and mount multiple MagFS shares.

Maginatics supports both an IT-centric consumption model and a user-centric model. In the former, policies are distributed directly from an Active Directory script so that users, upon login, will automatically receive access to their assigned portion of the MagFS

namespace, without any knowledge that Maginatics is operating behind the scenes.

In the end user-centric consumption model, IT provides a service to which the end user connects via a MagFS client tray app. This model is most relevant to organizations that have elected to support the BYOD (bring your own device) model.

Cloud-optimized Protocol

The MagFS Client communicates with the MDS using a cloud-optimized protocol that was designed from the ground up to leverage the MagFS split control-plane/data-plane architecture. It uses state-of-the-art network file system techniques for efficient metadata operations and leverages sophisticated WAN-acceleration techniques for transferring data to and from the object store.

Key Benefits

Enterprise Security

MagFS provides native integration with the existing enterprise Active Directory or other LDAP system. There is no need to create new profiles or to manually keep the user directory in-sync. IT uses its existing centralized identity management system with MagFS as if the latter were a traditional filer.

To enhance security, every file is broken into many small chunks (clients determine chunk boundaries), and each chunk is independently encrypted with its own AES-256 key before being sent to the data store. Files resident in the cache of an end-user device are also stored as a set of individually-encrypted chunks. In-flight, an additional layer of encryption is applied to both the chunks and metadata using SSL/TLS. Unencrypted data is never transmitted or persistently stored anywhere in the system, including on end-user devices. Critically, all keys are managed by the MDS itself and are never stored outside the datacenter.

MagFS ensures data integrity and guards against data loss by using a strong cache-consistency protocol. Metadata replication is done via snapshots which, from an administrator's standpoint, translate to full namespace snapshots.

Performance and Scalability

The MagFS system is linearly scalable from end-to-end. On the back end, MagFS leverages the scalability of the underlying object-based storage, which is known to scale linearly with the number of objects managed. MagFS pairs this back end with a front end that scales linearly with the number of active user sessions (not just concurrent users). With its split control-plane/data-plane architecture, the MDS is never a bottleneck for data transfers.

The MDS guarantees strong data consistency by using a lease-based mechanism to arbitrate caching capabilities. Lease-based, strong cache consistency means that clients can request the right to read from or write to their local cache for both file and namespace operations. Smart, adaptive caching at each client is enabled by an in-memory cache that holds hot data. Colder data is migrated to the encrypted on-disk cache; this staging is workload-aware. End users can configure the size of their local cache.

Inline de-duplication and garbage collection improve efficiency and reduce storage costs. Since the MDS maintains the file-to-chunk mapping, it is able to de-duplicate identical chunks across all files. Chunks that are not referenced by any file in the system are garbage-collected in the background to ensure that customers only pay for data they actually use.

Ease of Use

MagFS offers click-through provisioning. While IT projects generally can take weeks to complete and require the involvement of complex planning and assistance from professional services, implementing MagFS is typically very straightforward, with setup to startup requiring only minutes.

MagFS employs a modern UI that delivers a responsive management experience. The Web-based management console provides a single point of control for IT administrators with robust features including system dashboard, performance monitoring and upgrade support.

From the user's perspective, MagFS looks, feels and functions like any network folder and can be transparent if IT chooses to drive mount operations centrally through an Active Directory script. Existing applications can run against MagFS without modification. Users also benefit from a minimalist and user-friendly client UI, including tray apps on Mac and Windows.

Finally, native mobile apps allow users to open, preview and transfer (upload or download) files to and from other applications on the same device to support near real-time mobile collaboration.

Business Continuity

MagFS metadata snapshots can be securely stored off-site, either in the cloud or at a secondary site. In the event of a primary site disaster, operations can be rapidly resumed from any other location simply by spinning up replacement VMs and downloading the latest snapshot.

Also, once data is stored in the public cloud using MagFS, conventional backup is no longer required thanks to the inherent immutability of data stored in the cloud.

Use Cases

MagFS is well suited to any workflow that requires centralized storage of large amounts of unstructured data and/or de-centralized access to such data. The following are specific use cases for the solution.

- **Enterprise File Sharing and Collaboration.** Unlike any other solution in the market today, MagFS enables a genuinely enterprise-grade file sharing and collaboration platform by meeting all of the following requirements:
    - Complete security from end to end, including on end-user devices;
    - Near real-time collaboration regardless of team size or data set size;
    - Anywhere, any-device access; and
    - As easy to use as opening a network share or folder.

  Geographically distributed and mobile teams share a consistent view of shared content of any kind, from productivity documents to rich media files, CAD drawings, sensor data, medical records and other files of all sizes. In-line de-duplication ensures that only unique data is ever transferred from the cloud to a device or vice-versa. MagFS uses an “intelligent pull” system in which file changes are not moved to an end-user device until they are needed; when they are moved, only unique chunks are transferred, and all transfers are accelerated. This is in stark contrast to sync solutions, whose “blind push” approach propagates all file changes to all devices regardless of user need. Because MagFS does not sync data, users never have to wait for syncing to finish. As soon as an operation completes on one device, the new content is immediately available to all other users sharing the same name space.

- **NAS Consolidation.** MagFS enables enterprise IT organizations to consolidate data while preserving branch-office user experience and enhancing or adding remote and mobile access. With a solution powered by MagFS, IT can deliver efficient data access for users located in branch offices without having to replace one point of management in those branches (a local filer) with another (gateway appliance).
- **Migrating Applications to the Cloud.** MagFS overlays cloud storage with a distributed file system that looks, feels and functions like a traditional file system to applications and end users. As a result, migrating unmodified applications with a file system interface to the cloud is transparent to the application (and to end users). Scale-out workloads have the most to benefit from this capability. Whether in the Life Sciences (e.g., multi-step genomic analysis), media and entertainment (e.g., video rendering and other image manipulation) or large-scale Web environments (e.g., sites with highly variable traffic served by content management systems such as Drupal), these workloads require a multitude of compute nodes to have ready access to a shared capacity pool. In this case, traditional NAS falls short because all I/O must flow through the NAS appliance, which quickly becomes a bottleneck; and other distributed file system solutions – while they leverage the benefits of cloud compute – fail to expose the economics or agility of cloud storage. Pairing these application with the MagFS distributed file system eliminates bottlenecks and delivers the full benefits of the cloud.
- **Field Operations.** MagFS is ideal for provisioning data services in support of field operations, especially in situations where it would be very difficult or very costly to deploy a physical appliance (think of an oil rig or any remote outpost).
- **Hosted Online File Sharing.** MagFS is ideal for service providers to develop secure managed file sharing solutions for their customers.

Conclusion

MagFS expands NAS beyond its traditional boundaries by transforming it into a highly distributed system that connects the datacenter (where control resides), the geographically distributed end-user devices (where data is produced and consumed) and the cloud (where data is stored).

More broadly, MagFS offers NAS-like functionality in place of traditional NAS or where NAS is unavailable or not an option: on–premises to support distributed enterprises and mobile environments, and in the cloud to support scale-out workloads, enable disaster recovery solutions and extend the full breadth of cloud benefits to enterprises without causing disruption.

MagFS serves multiple objectives: to boost end-user productivity by enabling near real-time file sharing and collaboration, NAS consolidation, optimized in-cloud parallel workloads and other forms of NAS-like functionality for the highly-distributed enterprise; to enhance IT flexibility and economics by enabling enterprises to take full advantage of the cloud model; and to return firm control over data to IT.

Maginatics, Inc.
info@maginatics.com
(800) 360-1620 or (650) 265-1659
www.maginatics.com