



The Storfirst Gateway Storage Operating System: Cloud-as-a-Tier ILM

A Technical Marketing Whitepaper – version 5.0

we will help you evolve⁷

Making the Hybrid Cloud Model Work

The proliferation of low cost/high performance storage infrastructure has emphasized the dependence that many of today's enterprises have on immediate access to company-wide data. The "feed-me-now" mentality of data access and availability has never been more prevalent and expectations will continue to soar. Examples of "infinite institutional memory" have wide reaching applications. Take for example healthcare, where digital imaging modalities create gigabytes of clinical data per patient; all of which is expected to be available to clinicians with minimal delay virtually anywhere on the network. Additionally, your typical Fortune 500 organization is responsible for managing and storing billions of files and email that are retained based on policies closely tied to litigation and regulatory requirements. Today we see these industries discovering new value in digitizing their assets to provide on-demand access to increasingly smaller and specialized audience segments.

This rapidly expanding class of data; fixed content which is accessed infrequently, is commonly referred to as "unstructured data" or "Big Data". This data presents unique challenges to today's IT infrastructure. Immediate access requirements preclude the application of traditional backup and "shelf management" approaches to old information. In fact most IT professionals realize that backup commonly translates into a "Write Once Read Never" approach to data management which no longer satisfies users' requirements for data availability. The size and relative infrequency of access also challenges traditional SAN fabrics which have been designed to support online transactional processes. Compounding all of the challenges caused by the enormous growth of fixed content is the inherent lack of file system support for multi-petabyte archives and imbedded retention, encryption and data protection policies required by this new generation of data center managers and CIOs.

While significant efforts have been made to address the hardware characteristics unique to long-term storage (i.e. CAS, object stores, cloud), little has been offered by traditional file systems to meet the functional requirements of the problem domain. Users are quickly discovering that "throw-in" CIFS or NFS presentation layers that have been quickly developed for access into general purpose storage platforms (re. cloud) lack the scalability, data security, and performance to cope with the influx of unstructured content. At Seven10 we have applied our collective experience with file systems and unstructured data to develop Storfirst Gateway; a multi-tiered archive for cloud and hybrid cloud storage environments. This paper describes the challenges we addressed during the design process and provides detailed information about our implementation of Storfirst Gateway in the enterprise.

Architectural Drivers of Storfirst Gateway™ Cloud Integrated Storage

Seven10's Storfirst Gateway was developed as a response to the many shortcomings of existing cloud gateway, virtualization, HSM and tiered storage systems. Its architecture was engineered to address those inadequacies, and to respond to the needs of application providers in the healthcare, financial and compliance driven markets, amongst others.

A consistent set of demands called for a long-term data management solution that could easily address:

- Scalability – Clients should not be concerned that their data or applications will outgrow the ability of the storage system where content resides
- Security – The solution should provide security at a level that matches the capability of the most secure OS while exceeding the features of almost all appliance based object stores
- Flexibility – Storage hardware technologies, and the associated configuration requirements, are many and varied and evolve over time. A storage system must be designed to eliminate storage obsolescence and allow flexibility of hardware configuration, both at installation, and throughout the lifespan of the organization

The Traditional Approach

Almost all of the competitive solutions available fall into two main camps: the Linux hosted storage appliance, or the Windows virtualization stub model.

The offerings in each camp use a similar set of technologies in their implementation, and therefore project a similar set of features and characteristics. These are:

LINUX APPLIANCES

These are either Linux or proprietary OS based implementations that present an NFS interface and a CIFS/SMB emulation for Windows clients. These appliances could be open systems, in that the O/S distribution is a known Linux distribution and subject to administrative view into their inner workings, or closed systems. A closed systems approach offers no view into the O/S whatsoever, preferring administration be accomplished exclusively through a vendor provider GUI. These vendors commonly use FUSE (Filesystem in Userspace) to mount local paths using a pluggable module. However FUSE offers no kernel support or more sophisticated kernel-based solution which would effectively stub a local path to some external storage devices.

For security, these appliances provide a POSIX security model where each user has an ID (UID) and belongs to only one group (GID).

Linux appliances generally rely on hard disk exclusively for the metadata and provide functionality to protect data by leveraging an existing back-up infrastructure.

WINDOWS SOFTWARE

Windows virtualization offerings are most often implemented as Windows file system filters drivers layered on top of an existing local drive formatted with well-known file systems such as NTFS. To export this file system to the outside world, CIFS/SMB shares or NFS exports are created using the Windows O/S. File system filters are primarily constrained by, and dependent upon, the underlying file system. While NTFS is a very capable file system, it is neither designed for, nor does it provide, tiered storage and massive scalability. Filter based virtualization solutions often rely on an NTFS volume to provide the metadata storage as well as the 'arena' for data access. They then control what data is present or absent on that primary volume through migration to and from it. The filter detects access to a reparse point, and intercedes to bring the file data back into the arena based on the location information stored in the reparse point structure.

Again, these systems either support hard disk exclusively, or provide functionality to backup to tape or provide backup tie-ins.

The Seven10 Advantage

Recognizing that the vast majority of enterprise IT infrastructure and associated applications (PACS, Document Imaging, Email Archiving, MS Office, etc.) runs on a Microsoft OS foundation, Storfirst Gateway is engineered to operate on a physical or virtual Windows platform, and is optimized for the majority of application environments.

Rather than emulating interface technologies and having to run the constant gauntlet of incompatibilities and out-of-date support that emulators inevitably face, Storfirst Gateway is embedded firmly in Windows as an installable file system. This gives the underlying file system of Storfirst Gateway an opportunity to leverage key standards and functionality directly from the Microsoft OS.

In addition to its massive scale and flexibility, Storfirst Gateway offers a kernel level platform enabler that uses multiple abstraction layers to allow the rapid support of entirely new classes and instances of storage hardware (i.e. Cloud: REST, OpenStack). In addition, the Storfirst enabler uses a combination of industry

standard and proprietary file systems such that each discrete LUN or device in the system is self contained for both data and metadata.

SCALE

Scale refers to a system's ability to store increasing numbers of entities (files and subdirectories) efficiently, so that the overall size of the system grows without constraint and the performance of the system remains constant.

The Traditional Approach – Almost all appliance-based competitive systems have the capability to store approximately 16 EB in a single file or a complete system. Most have a limit of 4 billion entities.

The performance of most systems degrades rapidly as file counts increase above 200 million.

The Seven10 Advantage – Storfirst Gateway has no practical limit on the number of files in a given directory or on the number of files/directories on a volume. It also has no practical limit on the number of volumes in the system, neither is there any practical limit to the total managed system storage capacity. The Windows file system stack supports individual file size and individual volume size up to 16 EB, or the maximum size of a 64-bit integer, and therefore Storfirst follows suite.

To store the directory structure, Storfirst Gateway uses a metadata database based on the highly acclaimed Oracle Berkeley DB. Together with Berkeley's huge scalability and industry leading performance, the highly optimized database design results in only a very gradual performance degradation as file system size increases into the low millions of files due to the limited amount of memory to cache metadata. After this point, performance maintains steady as scale increases into the hundreds of millions to billions of files.

Storfirst maintains all necessary metadata to make the storage environment self describing. The same metadata is also externalized by the Storfirst metadata service. This has significant implications for performance. Traditional file systems require that metadata on the volume be scanned for queries and searches. This imposes significant performance challenges in very large storage systems.

Searching such systems with a traditional file system can result in poor performance (imagine a full meta-data search on a 10 PB volume) or pathological behavior (disk thrashing). The highly indexed metadata service provides immediate response query operation with less than 0.1% storage overhead.

SECURITY

All storage systems must provide the ability to manage access to data. The Windows security model and the Microsoft security management tools built into the OS combine into an IT industry standard that provides a solid, flexible and unified interface for access control.

The Traditional Approach – The Linux implementations present an emulation of Windows security on top of the CIFS emulation, involving them in a constant game of catch-up. The result is a poor compromise of incomplete support for the constantly changing, undocumented SMB network protocol APIs balanced against 80/20 implementation goals.

The Seven10 Advantage – Storfirst Gateway implements the storage of Security Descriptors and associated ACLs in a standard way that allows the familiar Microsoft tools and capabilities to operate within Storfirst exactly as they do in NTFS. The complexities of the security system are opaque to Storfirst, and revisions and enhancements to functionality require no changes to the Storfirst Gateway codebase. Two stunning example of this was support for SMB2 and multi-tier

encryption which were developed without any need to reengineer the file system. Details of our encryption support are provided here:

Encryption at Rest (AES256)

In cryptography, the Advanced Encryption Standard (AES) is a symmetric-key encryption standard adopted by the U.S. government. The design and strength of all key lengths of the AES algorithm (i.e., 128, 192 and 256) are sufficient to protect classified information up to the SECRET level. TOP SECRET information will require use of either the 192 or 256 key lengths. Storfirst Gateway implements a 256-bit AES key. This key is stored in the local Storfirst Gateway server database and never leaves the customer's site. The key will be randomly generated by the Storfirst Gateway software and is not displayable on the screen. When file data is flushed to remote media, it is encrypted in line with that transfer and therefore is always encrypted once it leaves the Storfirst server.

Encryption in Transit (SSL)

SSL support is a method of both authenticating and encrypting the connection between the Storfirst Gateway server and the REST web services layer. The certificate file section is used for authentication of the client server, in this case Storfirst Gateway, with the storage layer. In other words, it is a way for the storage server to verify you, Storfirst Gateway, are who you say you are. This is usually a PKCS12 keystore file which contains both a public X509v3 certificate and a private key. The password section is a place for users to configure the passphrase for the client certificate. These fields are only required if you wish to have two-way, or bilateral, authentication. This means that both the client and the server must prove to each other who they are. The authorization file section is for a server certificate and is used for authentication of the storage server with the client machine, in this case Storfirst Gateway. In other words, it is a way for the Storfirst Gateway server to verify the storage server is in fact who they say they are. This field is recommended if you wish to have one-way, or unilateral, authentication and is required if you wish to have two-way, or bilateral, authentication.

Additionally, Storfirst was designed to protect original content from malicious or accidental file modifications through file locking, retention and write-once provisions.

FLEXIBILITY

Seven10 recognizes that storage evolution is constant and unstoppable. Throughout our existence, Storfirst Gateway has been at the forefront of the storage evolution with first of its kind support for Centera, first to support Cloud Storage via REST and Openstack, de-dupe appliances, object stores, SAN/NAS and tape or VTL. Flexibility is reflected in how a product allows you to configure whatever devices you need, and specify the handling of your data in a way that allows you to solve your archival problem whatever they may be.

The Traditional Approach – Early attempts at solving the tiered storage management issue have required the use of proprietary platforms to store and retrieve data from these “all-in-one” storage appliances.

Competitive products provide a limited range of capabilities that primarily focus on the provision of tiered disk storage at a block level with the capability for tape or optical to be used as backup. While providing some form of simple, global retention control combined with an optional read-only emulation feature.

Often these solutions are monolithic in their approach to storage support and only adopt the latest and greatest technologies in a shortsighted development model that is designed to “predict a storage platform winner”.

The Seven10 Advantage – Storfirst Gateway eliminates vendor lock-in with a multi-tiering management feature that unifies all types of storage subsystems under a single presentation layer.

With its support for Cloud Storage, object stores, fiber channel and SCSI attached disk, fiber channel and SCSI attached tape, and many other storage platforms, Storfirst Gateway is architected to allow for seamless adoption or transition to new devices and new classes of storage.

Data Migration

In July of 2011 Storfirst Gateway became the first (and only) file system to support feature-based, software driven, data migration from legacy storage to cloud or hybrid-cloud storage environments. Seven10's patent-pending Storfirst Migration seamlessly inventories content on legacy storage and synch's the data onto the new storage tier providing a complete audit trail throughout. Non-disruptive to the applications storing the data, Storfirst Migration provides the most transparent move possible for organizations moving from one storage system to another.

By making sure information is highly accessible and secure, the Storfirst Gateway hybrid cloud file system promotes the concept of an "always available" data set, irrespective of which storage platform the data resides on and regardless of how infrequently accessed said data becomes. This design philosophy was implemented to greatly reduce the dependency on traditional back-up applications. By focusing on intelligent data placement (instead of a stub-based data migration policy associated with virtualization), Storfirst Gateway proactively manages data from its initial read/write workspace across a landscape of indelible archival storage platforms.

Policies define levels of data retention and lifespan per tier, so that the archive can preserve data for definable periods (0 days – permanent), and make best use of the various characteristics of tiered storage.

Summary

Storfirst Gateway will provide application providers, partners and prospects with all of the essential elements to support long-term compliant storage without blurring the line between existing services and application level responsibilities. Simple to implement, secure, and proven - the Seven10 cloud integrated storage software eliminates complexity and many of the limitations offered by the competition. These attributes bring users all of the virtues of the cradle-to-grave data lifecycle concept, including storage vendor neutrality, policy driven data retention, lower administrative burden and data immutability and security.