



PoINT Storage **Manager**



Contents

- 4 — **A valuable good**
 - Data volumes rising
 - Secure archiving
 - Long-term strategy required
 - Solution approach

- 8 — **PoINT Storage Manager**
 - Active and passive approach
 - HSM and ILM
 - Comprehensive ruleset
 - User-controlled archiving
 - Transparent file tiering and archiving
 - Virtual file system

- 14 — **Accessing the archive**
 - Stubs
 - Web links
 - Web Client
 - Data Browser
 - Archive volumes
 - Archive migration
 - Migration process
 - Supported storage technologies
 - Target storage system

- 20 — **Use Cases**
 - Compliance through archiving
 - Centralizing data by consolidating archive storage
 - Optimizing infrastructure using ILM and tiering
 - Hybrid infrastructure integrating on- and off-premises storage

A valuable good


Data volumes rising

The degree to which volumes of data are changing varies between businesses, but the direction of travel is universal: the numbers are constantly rising. New technologies mean that more and more **unstructured data** is being captured, in increasing levels of quality and detail. Key factors in this massive increase in data include new options for analysis, which are creating financial incentives to **permanently store large volumes** of data to serve new business models.

IT departments must now face the challenge of supplying a dedicated storage infrastructure for this purpose. From a technical perspective, constantly expanding the primary storage system is not a reasonable approach, nor is it financially possible without increasing IT budgets. Analysing file systems can deliver critical insights and a basis for making decisions on a long-term storage strategy. **Analyses of unstructured data** usually show that **more than 70% of data is inactive** – in other words, that it has not been used for a long time. While the active data remains on the high-performance primary storage, a data management plan can be used to

move the inactive portion to another storage system which suits its current life cycle stage. Aside from the technical benefits, combining multiple storage technologies can result in significant cost savings.

Unstructured data



The more data is kept on the primary storage system, the greater the scope of the data security plan must also be. As a result, **backup costs rise as well**. If data can no longer be backed up within an acceptable timeframe, an approach is needed that reduces the load on the primary storage system and therefore reduces backup volumes as well.

Secure archiving


If the data held by a company contains valuable information, or information that may potentially be valuable in the future, the question of long-term storage becomes more interesting. Meanwhile, companies face an increasing number of **laws and regulations**, including GDPR, GoBD, Basel and SOX, which require long-term retention of some data. Legal retention periods vary depending on the type of document, but they are often around 10 years, sometimes rising to 30 years or more.

Archived files need to be protected against any changes, be they accidental deletion, malicious manipulation, or malware attacks. In particular, the danger of **ransomware**, which encrypts data and demands payment from its targets to release files again, has increased dramatically in recent years.


Businesses also need to ensure that data will be findable and accessible in future, and they may also need to delete certain information (such as personal data) from the archive.

Long-term strategy required

If the intent is to store data for years or even decades, the question then becomes which storage systems and which format(s) best meet this requirement. The lifespan of the storage system used is determined by environmental factors, the availability of devices to read it, and the question of whether it makes financial sense to keep operating a given system. The latter point in particular often becomes relevant within a few years of the end of regular product support. As a result, it is often necessary to **migrate archived data multiple times** from one storage system to another. A software-based approach to archiving, independent of any storage system manufacturer, is necessary in order to successfully **migrate from technology to another**. Proprietary formats are not suitable for long-term storage, since they depend on a single developer.



Vendor-Lock-In



Solution approach

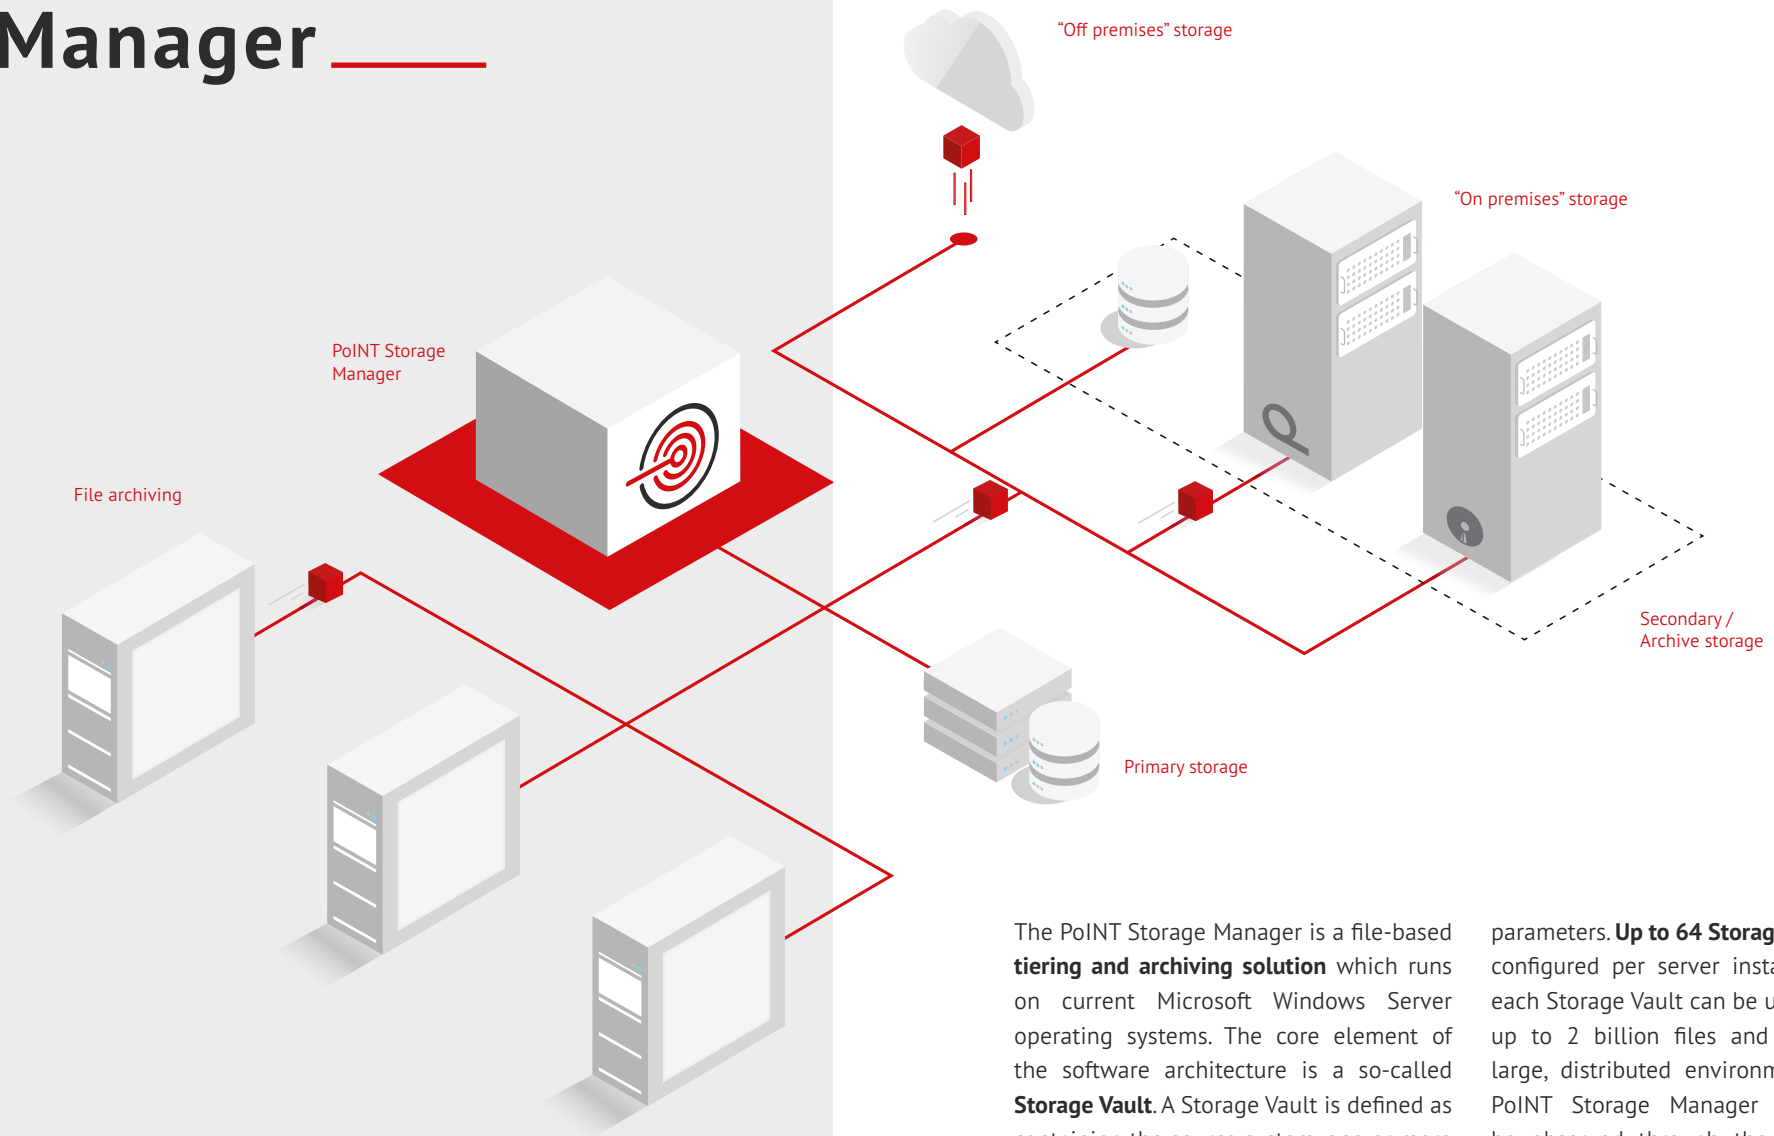
As an independent software vendor (ISV), PoINT Software & Systems offers the **PoINT Storage Manager** as a file-based **tiering and archiving solution** developed independently of any storage system manufacturer. The PoINT Storage Manager enables customers to convert unstructured data into a **multi-tier storage architecture** using a comprehensive ruleset. By relocating inactive data, the high-cost primary storage system sees its load reduced, which in turn reduces the scope of data protection operations.

Because it supports a wide spectrum of storage technologies, the PoINT Storage

Manager integrates seamlessly into existing infrastructures. Multiple access options ensure easy access to archived files, whatever the use case. The software architecture and integrated retention management system enable **secure long-term archiving**.

The PoINT Storage Manager uses a standardized storage format and offers a convenient, straightforward background migration process which allows customers to smoothly switch over to future storage platforms.

PoINT Storage Manager



The PoINT Storage Manager is a file-based **tiering and archiving solution** which runs on current Microsoft Windows Server operating systems. The core element of the software architecture is a so-called **Storage Vault**. A Storage Vault is defined as containing the source system, one or more target storage systems, a ruleset for tiering and archiving, plus additional configuration

parameters. **Up to 64 Storage Vaults** can be configured per server instance. Currently, each Storage Vault can be used to manage up to 2 billion files and directories. In large, distributed environments, multiple PoINT Storage Manager instances can be observed through the central status monitor.

Active and passive approach

Fundamentally, two different approaches need to be considered for an archiving system. In one approach (the active approach), the solution itself gathers archive data from the source system and saves it to whichever archive storage system is in use. A passive approach, by contrast, sees the solution supplying an interface which receives data from users or applications. The distinction, therefore, is in terms of which components in the archiving process are used to transfer data to the archive system. The PoINT Storage Manager supports **both approaches**.

HSM and ILM

The PoINT Storage Manager offers hierarchical storage management (HSM) and information lifecycle management (ILM) for unstructured data, in order to save this data on the appropriate storage system based on its **access history** as well as its **value and usage**. The extensive ruleset, and support for a wide range of source and target systems, mean that customers can realise automated file tiering and archiving.

Within the architecture of the PoINT Storage Manager, the existing primary storage system represents the top tier in the hierarchy, which users and applications access directly. The PoINT Storage Manager

integrates a second (and optionally a third) tier into the infrastructure while **retaining this same access point**. Using interfaces within the primary storage, the PoINT Storage Manager scans the file system for files meeting specified criteria.

Comprehensive ruleset

Rules within a Storage Vault are used to define the conditions a file must meet in order for the PoINT Storage Manager to carry out a given set of actions. A ruleset can consist of one or several rules, which are processed in order from top to bottom. This is how workloads are represented and files are archived in a highly targeted manner.

Possible **conditions** for a rule include:


- File name (e.g. a specific file extension)
- File status (new/edited, archived or replaced with a placeholder)
- File age (creation or edit date)
- Last access (access date)
- File attribute
- File size

The following **actions**, among others, can be applied to files:

- Archive file
- Replace file with placeholder (stub or web link)
- Delete file
- Create log entry
- Skip next rule



“The Solution for Information Lifecycle Management”



In addition to these conditions and actions, the PoINT Storage Manager also has restoration rules for re-creating accidentally or maliciously deleted placeholders, or to copy archived data back to the primary storage system. This is either done to a specified schedule or can be triggered manually.

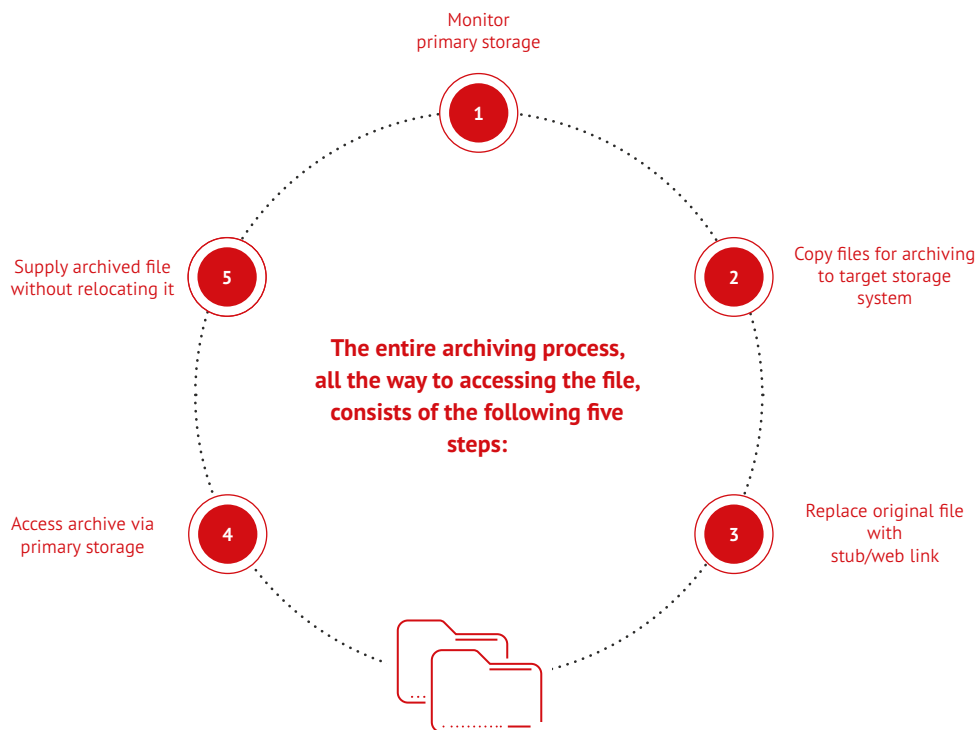
User-controlled archiving

In addition to its automatic archiving tools and based on the same ruleset, the PoINT Storage Manager supports **user-controlled archiving**. To this end, users are given a Windows Explorer add-on which allows them to carry out specific actions for selected documents.

The commands available to the user in the software on their workstation are set by

the administrator. A command, whose name can be freely customised, can consist of the complete set of available rules. Possible scenarios include archiving invoices in the accounting department or specialist offices, archiving folders on their own once projects are completed, and restoring them as needed.

Transparent file tiering and archiving



Virtual file system

If an application itself is able to write the files it produces to a pre-specified directory, or if users need to manually copy files to the archive, there is no need to involve the primary storage system, and the PoINT Storage Manager will instead supply a virtual file system in order to archive the files in question.


As with the active approach, customers also benefit from the PoINT Storage Manager's **WORM functionality** and **retention management** features when taking the passive approach. Once archived, files can be read, but they cannot be edited any more. If they are changed, a new version is created and the original file is kept in its original condition.

Retention management lets customers specify a **retention period** for archived files. During this period, files cannot be deleted

from the archive. In many industries, this represents the basis for legal compliance with retention requirements. Both absolute dates and relative timespans from the point of archiving are supported. If a file needs to be deleted before the retention period expires, e.g. in order to meet legal requirements, a so-called **privileged deletion** option is available. This process is logged for evidential purposes in the PoINT Storage Manager.


The PoINT Storage Manager can also function as a **gateway**, supplying file system-based access to an object storage system or public cloud provider.

Accessing the archive



The PoINT Storage Manager provides a range of options for accessing the archive, designed to meet a range of possible needs. During the actual archiving process, the PoINT Storage Manager copies files that meet the conditions in the ruleset to the specified archive device. In the same step, or afterwards, the original files are replaced within the file system by a placeholder just a few kilobytes in size, in order to relieve the load on the primary storage system.

if the file's properties are inspected, since the "size" of the original file will be used, but the "**size on disk**" attribute will be seen to be minimal.



Stubs


For NetApp FAS, Dell EMC Unity and Windows-based source systems, the PoINT Storage Manager supports "stubbing". Here, a "stub" is left in the source system which is almost indistinguishable from the original file in terms of its appearance and behaviour, in order to provide maximum transparency for users and applications.

The properties of the original file, including its icon, size and name (including extension) remain the same. Depending on the operating system in use, the icon will also show a small "x" to signal that it is actually located within the archive. The small size of each stub, just a few kilobytes, results in more free space on the primary storage system. This effect becomes clear

Stubs can be opened by users and applications in the normal way. If the files are only being read, the PoINT Storage Manager's "**pass through on read**" functionality is used to show the requested elements of the file without restoring it. If a user or application edits a file, the modified file will be saved back to the source system. As soon as this new version meets the specified ruleset, the PoINT Storage Manager's **versioning** functionality will take effect and a new version of the file will be placed in the archive. In principle, files in the archive are never edited.

Web Links

On any NAS system, the PoINT Storage Manager can replace an original file with PoINT Web Links. As well as being **independent of the source system**, this access method is also **independent of the client** (operating system). The Web Link can be a URL file (for Windows-only environments) or a HTML file, which is also compatible with Linux and MacOS clients. When it is put in place, the Web Link is given the same permissions as the original file.




When the user opens a Web Link, a TLS-secured download of the archived file will be supplied via the user's default browser. Support for integrated authentication means there is no need for an additional login (single sign-on).

Web Client

If there is no need to directly connect the primary storage file system to the archived files, the PoINT Storage Manager offers a Web Client. This option can also complement stubs or Web Links.

The **platform-independent web portal** shows authorised users a directory tree of



the archive. The user can download given versions of archived files. If multiple files are selected, or an entire folder, they will be provided in ZIP format. Alternatively, the files can be provided via temporary network access, allowing the user to simply copy them to the desired location with their file explorer.

Data Browser

The IT administrator can access the integrated Data Browser directly from the PoINT Storage Manager user interface. Using a view similar to that of a file explorer, selected versions of files can be copied from the archive. Numerous filtering options allow the user to quickly restrict what is shown, in order to find the files they need.

In addition, the Data Browser shows both the retention period and archive volume (see below) where the file is stored. The Data Browser can also be used to verify the integrity of files using checksums.

Archive Volumes


The Universal Disc Format (UDF) is a standardized, platform-independent file system format. The PoINT Storage Manager's Archive Volumes are based on the standard format described in version 2.01.

While archiving files, the PoINT Storage Manager collects files that meet the conditions of its defined ruleset. The PoINT Storage Manager generates so-called Archive Volumes (UDF containers) in the temporary image directory using the files to be archived, which are then written to the archive storage system. The size of the container can be specified, optimising the value in line with the properties of the archive storage system used and the customer's specific workflow.

Large archive volumes improve write speed and make later migration to a new archive storage system run faster. Optionally, a file can be split across multiple Archive Volumes (UDF containers).

Even without the PoINT Storage Manager, businesses can access archived files, since the UDF containers generated are readable using Windows, Linux or MacOS tools. File lists can also be exported for categorisation in CSV format using the Data Browser.

Archive Volumes can optionally be encrypted. In this case, the PoINT Storage Manager uses an AES256- and CBC-based encryption process at the block level.



Archive migration

Whenever data is stored for long periods of time, it will need to migrate several times across multiple storage technologies. The PoINT Storage Manager enables convenient, uninterrupted archive migration which can be configured in just a few steps. The actual transfer process runs in the background.

Migration process

- Set up new archive storage system
- Start background migration from source system to target
- Delete old archive storage

Since the PoINT Storage Manager works with large archive volumes, the transfer rate is dramatically faster than copying individual files. Once migration is complete, the IT administrator receives a log file with a record of the archive volumes transferred.

Supported storage technologies

As it is independent of any vendor, the PoINT Storage Manager supports a wide range of source and target systems from the widest possible range of manufacturers. This enables a high degree of flexibility when choosing an archive storage system and makes it easy to migrate across storage technologies.

Primary storage

The stubbing process is currently supported on the following primary storage systems:

- Dell EMC Unity (VNX)
- NetApp FAS (7-mode and cluster-mode)
- Microsoft Windows NTFS/ReFS based storage

The Web Link process is also supported on all NAS systems that the PoINT Storage Manager server can access using the SMB/ NFS protocol. Examples of NAS systems include:

- Dell EMC Isilon
- Hitachi NAS
- Huawei OceanStor 9000
- Qumulo



Old archive storage

New archive storage

Target storage system

The supported target storage systems are cloud and object stores, NAS systems, tape systems and optical systems. These are integrated using so-called connectors. The PoINT Storage Manager can write to up to four devices, of the same or different types, in parallel to one Storage Vault. In principle, any NAS system with an NFS/SMB interface can be used as the target storage system.


Certain appliances providing additional WORM functionality are also supported.

By using multiple target storage systems simultaneously, businesses can improve data security by putting technological and physical distance between their data. Using a multi-cloud approach, dependency on any one cloud provider is reduced.

The following is an incomplete list of supported target storage systems:

Cloud/object Stores	NAS systems appliances	Tape systems	Optical systems
AWS S3	NAS with SMB access	HPE	DACAL Jumbox
Microsoft Azure	Dell EMC Data Domain	Fujitsu	DISC ArXtor
Oracle Cloud	NetApp FAS (SnapLock)	IBM	HIT HMS / HDL
Wasabi	FAST LTA Silent Bricks	Oracle (ACSLs)	Sony PetaSite ODA
Cloudian HyperStore	FAST LTA Silent Cubes	Overland	
Dell EMC ECS	Quantum Artico	Quantum	
Hitachi Vantara HCP		Qualstar	
IBM COS		Spectra Logic	
NetApp StorageGRID Webscale			
Scality Ring			
SUSE Enterprise Storage			

Use Cases ---



The PoINT Storage Manager is used worldwide by many customers, including small businesses with low data volumes (a few TB) and large enterprises handling multiple petabytes of data across distributed structures. The most common use cases are listed below.

Compliance through archiving

The reasons for archiving data for a long period of time can vary. As well as commercial interests, there are especially numerous legal requirements which mandate years- or decades-long storage. The most important technical points to take into account are data security, migration to new storage technologies, and support for all critical storage systems and technologies.

— CHALLENGES


- Archive inactive data and data for archiving
- Protect archived data against modification (WORM)
- Retention management at the archive storage level

— SOLUTION

- Realise a two-tier HSM architecture
- Automated, rules-based archiving
- Transparent access to archived files via the primary storage's file system

— BENEFITS

- Use the primary storage system efficiently by archiving inactive data
- Permanent data availability
- Fulfil compliance requirements with WORM and retention management
- Future-proof design thanks to integrated migration functions
- Protect investments by storing archived data in a standard format
- Safe time and money by reducing backup data volumes on the primary storage system



Centralizing data by consolidating archive storage

One difficulty for businesses based in multiple locations with a decentralised storage infrastructure is the large amount of management work this causes, and the constantly rising costs that result. In order to streamline locations whose IT departments have grown, cold data should be transferred off the primary storage system and stored in a central location. A private cloud with a highly available object storage system is particularly well suited to this task.

— CHALLENGES


- Integrate a private or public cloud
- Reduce primary storage capacity required at each site
- Reduce administrative work
- Prevent vendor lock-in

— SOLUTION

- Automatic, rules-based file archiving in a central private or public cloud
- Transparent access to files moved to the archive without returning them to primary storage
- Flexible storage architecture
- Integrated replication and migration functions

— BENEFITS

- Efficient use of primary storage systems
- Meet compliance and archiving requirements
- Minimize the duration of the backup process and reduce backup size
- Independence from public/private cloud providers



Optimizing infrastructure using ILM and tiering

On one hand, primary storage systems are designed for high performance and high availability thanks to processes like synchronous data mirroring. On the other hand, these two factors also increase the cost per terabyte. Meanwhile, the rapid rise in data volumes makes data protection an increasingly complex task as well. By transferring inactive data away from primary storage, the primary system sees a significant reduction in capacity usage, making data protection simpler as well.

— CHALLENGES


- Optimize usage of available resources
- Unchanged workflows for users and applications
- Simplify infrastructure
- Reduce backup data volumes

— SOLUTION

- Rule-based file tiering and archiving cold data
- Multi-tier storage architecture
- Transparent read access without transferring data back to primary storage
- Consolidate primary storage systems

— BENEFITS

- Optimal usage of primary storage for productive data
- Reduced data volumes on the primary storage system
- Save time and money by reducing backup data volumes
- Protect investments by avoiding dependency on any one vendor



Hybrid infrastructure integrating on- and off- premises storage

For many businesses, cloud products are very attractive options for storing data. The benefits include fast, flexible access to services on demand, monthly billing based on resource usage, and no need for a business' own IT personnel to complete maintenance work or upgrades. A hybrid concept allows businesses to combine the benefits of both on- and off-premises solutions.

— CHALLENGES


- Integrate on- and off-premises solutions
- Workflows remain unchanged for applications and users
- Independent of any cloud solution provider

— SOLUTION


- Multi-tier hybrid storage architecture
- Homogeneously integrate on- and off-premises solutions
- Flexible, individual ruleset for mapping the workflow
- Transparent read access without transferring data back to primary storage
- Data protection thanks to encryption

— BENEFITS

- Native support for on- and off-premises solutions
- No need to adapt existing workflows
- Fulfil compliance requirements
- Vendor independence thanks to migration tools
- High system availability thanks to synchronous replication



PoINT Software & Systems



PoINT Software & Systems is specialized in the development and distribution of software products for storage, management and long-term archiving of data using all available mass storage technologies like hard disks/flash, magnetic tapes, optical media, object store and cloud storage. We work jointly together with leading manufacturers of storage systems. Thus, we can offer an early support of innovative storage technologies. Furthermore, we plan entire storage solutions and provide consultancy with our long-term and versatile expertise.

PoINT products allow efficient usage of storage systems and help to reduce costs and issues caused by data growth. The software solutions fulfil compliance and archiving requirements and provide independence from storage technologies and vendors. PoINT products are distributed by our partners world-wide and have been proven in more than two million installations. Our customers include many well-known companies from different industries, who comply with our solutions their complex demands by providing the necessary reliability and perfection.

PoINT Software & Systems GmbH
Eiserfelder Straße 316
57080 Siegen, Germany

P +49 271 3841-0
M info@point.de
W www.point.de