PoINT Storage **Manager**

Your data is the key, so don't lose it.

Contents

A valuable good ____

Data volumes rising

The extent to which data volumes are changing varies between companies. But the direction is certainly the same – it is steadily going up. New technologies are capturing more and more unstructured data at ever higher levels of quality and detail. Among the main drivers of massive data growth are new analytics capabilities that make the mass and permanent storage of data from a wide variety of sources lucrative and relevant to new business models.

IT departments face the challenge of providing a storage infrastructure suitable for this purpose. Permanently expanding primary storage is technically impractical and economically unfeasible without higher IT budgets. Important insights and a basis for decision-making for a sustainable storage strategy can be gained by analyzing the file systems. Analyses of unstructured data show that more than 70 percent of the data is usually inactive, i.e. it has not been used for a long time. While the active part of the data is located on the fast primary storage, the inactive part can be moved to other storage through a data management concept that meets the requirements of the

respective lifecycle. In addition to technical advantages, the combination of several storage technologies achieves high cost savings.

Unstructured data

● **Passive**   ● **Active**

The higher the volume of data on the primary storage, the higher the volume of data backup. Consequently, the costs for the backup increase. If data can no longer be backed up within an acceptable backup time window, approaches are required that offload the primary storage and thus reduce the backup volume.

Secure archiving

If data contains valuable information for the company or information that is potentially valuable in the future, the interest in long-term storage increases. In addition, companies are confronted with an increasing number of laws and regulations, e.g. DSGVO, GoBD, Basel or SOX, which require long-term archiving. Legal retention periods vary depending on the type of data and are often ten and in some cases even 30 or more years.

Archived files must be protected from any changes, be it accidental deletion by the user, malicious manipulation or malware infection. The danger of ransomware in particular, which encrypts files and demands ransom payments from victims, has increased enormously in recent years.

At the same time, it must be ensured that data can be located and accessed in the future. In addition, there is the need to selectively delete information such as personal data from the archive if necessary.

Long-term strategy required

If files are to be stored for years or even decades, the question arises which data storage system and which format are best suited for this purpose. The lifetime of the data storage used is determined, among other things, by environmental influences, functional capability and the factor of whether continued operation of the storage is economical. Especially the latter point often occurs within a few years with the end of regular product support. Migrations of archived data to new storage technologies are therefore inevitable. In order to change technology, a software-based archiving approach that is independent of storage manufacturers is required. Due to dependency, proprietary formats are unsuitable for long-term storage.

No Vendor-Lock-In

Solution approach

As Independent Software Vendor, in short ISV, PoINT Software & Systems offers with PoINT Storage Manager a file-based tiering and archiving solution which has been developed independently from storage vendors. PoINT Storage Manager offers customers the possibility to transfer unstructured data into a multi-tiered storage architecture by means of a comprehensive set of policies. The tiering of inactive data offloads cost-intensive primary storage and thus reduces the storage volume.

In addition to active tiering from existing primary storage systems, PoINT Storage Manager also provides an archive file system which users and applications can use to store the data to be archived.

By supporting a wide range of storage technologies PoINT Storage Manager integrates seamlessly into existing infrastructures. Multiple access options ensure easy access to archived files in any use case. The software architecture as well as the integrated retention management enable long-term and secure archiving.

For storing, PoINT Storage Manager uses standardized formats and offers a comfortable as well as easy-to-use background migration, so that customers can smoothly switch to future storage platforms.

File Tiering vs. Block Tiering

A comparison

Storage tiering is a proven method of data and storage management that moves data between different storage classes and storage tiers. The objective is to make optimal use of storage systems. Most companies have a large inventory of so-called cold data:

- This cold data is rarely (or never) accessed.
- Nevertheless, this data must be retained, for example due to legal requirements.

Tiering of this data is automated on the basis of predefined policies. For example, data that has not been accessed for a certain period of time is moved to a lower-performance, but less expensive storage system.

Storage tiering brings cost and time savings as well as efficiency improvements:

- Primary storage systems are offloaded from inactive data.
- Backup times are significantly shortened and backup volumes are reduced.

Tiering can be performed according to different methods – analog to the different storage methods. A distinction is made between file-based and block-based tiering. File tiering belongs to the field of data management. This method is based on file attributes (e.g. size, age or last access to the file). Tiering is flexibly controlled by software, independent of the storage systems used. PoINT Storage Manager is a software which performs vendor-independent, file-based tiering. The software moves files within a multi-level storage architecture to the storage level corresponding to their age or usage. The policies for this tiering are set in advance by the user.

Block tiering, on the other hand, is a vendor-specific storage management feature. Block tiering involves moving individual blocks of files between storage tiers. The basic factor here is the frequency of access to the respective block.

Tiering and workflows

The choice of the tiering method has a decisive influence on the users' workflows.

The following table compares the two methods.

	File Tiering (PoINT Storage Manager)	Block Tiering (Storage System)
File integrity	The files are retained as such. They are stored on the primary or secondary storage according to the tiering policies and remain identifiable by attributes.	Individual fragments of a file (blocks) are stored on different storage tiers. A moved block cannot be identified.
Tiering policies	The tiering policies can be specified individually. Tiering can also be triggered manually.	The policies for tiering are predefined and cannot be adjusted to individual requirements (no manual tiering, no user-defined policies).
File access on secondary storage	The user can access the files seamlessly and independently from the primary storage after they have been moved to the secondary storage.	File access is only possible through the primary storage system.
Read access	Restoring to primary storage is not necessary for read-only access. This maintains performance and saves storage space.	Moved blocks must be restored to the primary storage (e.g. for a backup).

Investment protection and future proofness

	File Tiering (PoINT Storage Manager)	Block Tiering (Storage System)
Vendor independence	The tiering software moves files between storage tiers – regardless of the storage systems used.	Block tiering is a functionality of the storage system itself and is therefore vendor-specific.
Backup	The backup application identifies and skips moved files. It stores only the links or stubs that refer to the archived files. This provides significant savings in terms of storage space and backup time.	The backup application cannot identify or skip moved blocks.

|	File Tiering (PoINT Storage Manager)	Block Tiering (Storage System)
Storage migration	Primary and secondary storage can be migrated independently without restoring moved files	During the migration process, a full restore of moved blocks on the primary storage system is required. Migration of the secondary storage only is not supported.
Read access	The user can access the files seamlessly and independently from the primary storage after they have been moved to the secondary storage.	File access is possible only through the primary storage system.
Legally compliant archiving	Within the scope of file tiering, methods such as WORM, retention management and versioning are supported.	Typical functions that enable legally compliant archiving (WORM, retention management, versioning) are not supported by block tiering.

Conclusion

Block tiering is a method performed by the operating system of the storage system. Transparency about the location of a file is lost because a file is fragmented and only the NAS operating system knows where which blocks of a file are stored. In addition, migrations become very difficult because blocks of a file that have first been moved out have to be moved back in on the primary storage. In many cases, this is no longer even possible due to space limitations.

File tiering as performed by PoINT Storage Manager offers enormous savings potential and independence from the storage manufacturer or storage service provider. On the basis of automatically or manually executable policies, files which have not been accessed for a long time, for example, are moved as a whole to more cost-effective storage classes. In addition, file tiering can meet compliance and archiving requirements.

Block tiering is a storage system specific method and leads to vendor lock-in. Software-based file tiering with PoINT Storage Manager, on the other hand, represents independent data management.

PoINT Storage Manager ———

PoINT Storage
Manager

Archive Tier

Capacity Tier

Primary Storage

PoINT Storage Manager is a file-based tiering and archiving solution. The core element of the software architecture is a so-called Storage Vault. The definition of a Storage Vault includes the source system or the archive file system, one or more target storage systems, a set of policies for tiering and archiving as well as further configuration parameters. Up to 64 Storage Vaults can be set up per server instance. The maximum number of managed files and directories is currently two billion per Storage Vault. In large distributed environments, multiple PoINT Storage Manager instances can be monitored centrally via the Status Monitor.

Active und passive approach

Basically, two different approaches must be considered for an archiving system. In the active approach, the solution collects the data to be archived from the source system itself and stores it on the configured archive storage (HSM/ILM). In contrast, in the passive approach, the solution provides an interface that accepts data from users or applications (Archive File System). Thus, it has to be differentiated by which component in the archiving process the

data is transferred to the archive system. PoINT Storage Manager supports both methods of operation.

HSM and ILM

PoINT Storage Manager provides hierarchical storage management (HSM) and information lifecycle management (ILM) for unstructured data to move them to the most appropriate storage tech-nology according to their access behavior or their value and usage. By means of the comprehensive set of policies and the support of multiple source and target systems, customers can realize automated file tiering and archiving.

In PoINT Storage Manager's architecture, the existing primary storage represents the top level of the hierarchy which is directly accessed by users and applications. Keeping this access point, PoINT Storage Manager integrates a second and optionally a third storage level into the infrastructure. Via interfaces of the primary storage PoINT Storage Manager scans the file system for files which comply with the defined set of policies.

Comprehensive ruleset

Within a Storage Vault, policies are used to define the conditions a file has to fulfill in order for PoINT Storage Manager to perform defined actions. A set of policies can consist of one or more policies, which are processed sequentially, from top to bottom. This way, workloads can be mapped and specific files can be archived.

The possible conditions of a policy include:
- Name of the file (e.g. a specific file extension)
- Status of the file (new/modified, archived, or replaced by placeholder)
- Age of the file (creation or modification date)
- Last access (access date)
- Attributes of the file
- File size

Among others, the following actions can be applied to files:
- Archive the file
- Replace the file with a placeholder (stub or link)
- Deleting the file

In addition to these conditions and actions for archiving, PoINT Storage Manager provides recovery policies to recreate accidentally or maliciously deleted placeholders or to copy archived files back to primary storage. Execution is controlled by a schedule or triggered manually.

Archive File System

If an application itself has the possibility to write files into a predefined directory (file system) or if users are supposed to copy files manually into the archive, then the integration of the primary storage is not necessary and PoINT Storage Manager instead provides a file system into which files to be archived can be stored.

As with the active approach (HSM/ILM), also with the passive approach (Archive File System) customers benefit from the WORM functionality and the Retention Management of PoINT Storage Manager. Retention Management allows users to define a retention period for archived files. During this period files cannot be deleted from the archive.

Furthermore, the Archive File System of PoINT Storage Manager can be used as gateway to provide a file system based access to an object storage or to a public cloud provider.

Vendor independence

PoINT Storage Manager supports numerous source and target systems of different manufacturers. This provides high flexibility in the choice of secondary/archive storage and allows an easy change of storage technology. Vendor dependency and vendor lock-in are thus avoided.

As target storage PoINT Storage Manager supports cloud and object stores, NAS systems, tape systems and optical systems.

PoINT Storage Manager can write to up to four storage systems of the same or different technologies in parallel in one Storage Vault. Basically, any NAS systems with CIFS/NFS interface can be used as target storage. Likewise, appliances providing e.g. additional WORM functionalities are supported.

By using multiple target storages simultaneously, companies increase data protection through a technological and geographical break. By means of a multi-cloud approach, dependency on a cloud provider is reduced.

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			
Scality Ring			
SUSE Enterprise Storage			

Accessing the Archive ‗‗‗‗


Many ways possible

PoINT Storage Manager provides several options for archive access to meet different requirements. During the actual archiving process PoINT Storage Manager copies files which correspond to the defined conditions in the set of policies to the configured archive storage system. In the same step or subsequently, the original files are replaced by a placeholder, which is only a few KByte in size, in the file system to offload the primary storage.

Stubs

For NetApp FAS, Dell EMC Unity and Windows-based source systems PoINT Storage Manager supports the so-called stubbing method. Hereby, a stub is left in the source system which does not differ in its behavior from the original file in order to provide maximum transparency for users and applications.

The properties of the original file, such as the icon, the file size or the name including the file extension, are retained. Additionally, depending on the operating system version used, a small "x" is added to the icon to indicate that this element lies in the archive. A stub has a small size of only a few KByte, so that storage space is freed on the primary storage. When you call up the file properties, this effect becomes clear because the „size" is taken from the original file, but the occupied „size on disk" is minimal.

Stubs can be opened by users and applications in the usual way. In case of a read-only operation, the "Pass Through on Read" functionality of PoINT Storage Manager takes effect to provide requested blocks of the file without restoring the file. If

a file is edited by the user or an application, this modified file is stored once again on the source system. As soon as this new version complies with the defined set of policies, PoINT Storage Manager versioning is applied during archiving and another version of this file is created in the archive. Files in the archive are generally not modified.

Symbolic Links

For NAS systems with CIFS shares which do not support stubbing, e.g. Dell EMC PowerScale (Isilon), symbolic links are supported. A symbolic link acts as a reference to the file archived by PoINT Storage Manager and thus provides access to archived data.

Web Links

For any NAS system PoINT Storage Manager can replace the original file by so-called web links. Besides being independent of the source system, this access method is also independent of the client and its operating system. URL files for Windows-only en-vironments or HTML files, which are also

supported by Linux or MacOS clients, can be selected as web links. When replacing, the Web Link receives the permissions of the original file.

When the user opens a Web Link, a TLS-secured download of the archived file is provided via the configured default browser. Support for integrated authentication eliminates the need for an additional login process (Single Sign-On).

Data Browser

The administrator can open the Data Browser integrated in PoINT Storage Manager directly from the user interface. Using a File Explorer-like view, selected file versions can be copied from the archive. Numerous filter options allow to quickly limit the search to find the files you are looking for.

In addition, the Data Browser displays the retention period and other parameters of the archived file. The Data Browser can also be used to verify the data integrity of files.

Web Client

If no direct relationship to the archived files is to be configured for the file system of the primary storage, PoINT Storage Manager offers the so-called Web Client for access. This option can also be an addition to Stubs, Symbolic Links or Web Links.

The platform-independent web portal shows authorized users a directory tree of the archive. Users can selectively download required versions of archived files. If several files or an entire folder are selected, they are provided as a ZIP file. Alternatively, the files are offered via a temporary network share, which allows convenient copying to the destination via the File Explorer.

Archiving features ———

User Controlled Archiving

In addition to automatic archiving, PoINT Storage Manager offers the functionality "User Controlled Archiving" based on the same set of policies. Here, users are provided with a Windows Explorer extension to perform specific archiving actions for self-selected documents.

The selectable commands that the user sees in the software component on his workstation are determined by the administrator. A command, whose name is freely definable, comprises a set of policies to be executed in sequence. Possible scenarios include, for example, the archiving of invoices by the accounting department or specialist departments that can independently archive folders relating to completed projects and restore them as required.

WORM and Retention Management

The WORM function of PoINT Storage Manager makes sure that once archived files can only be read, but not modified. In case of a modification, a new version is created, leaving the original file in its original state.

Retention Management allows customers to specify a retention period for archived files. During this period, files cannot be deleted from the archive. This provides the basis for legal compliance for retention requirements in many industries. Both an absolute date and a relative time period from the time the respective file is archived are supported. If a file has to be deleted before the specified retention period has expired, e.g. due to legal requirements, the so-called privileged deletion is available. These deletions are logged by PoINT Storage Manager for proof.

Archive Volumes

During an archiving cycle PoINT Storage Manager collects the files which meet the conditions of the defined set of policies. In the temporary image directory PoINT Storage Manager creates so-called Archive Volumes (UDF containers) from the files to be archived. These Archive Volumes are then written to the archive storage. The size of the Archive Volumes can be adjusted so that the value is optimally adapted to the properties of the archive storage used and the customer-specific workflow.

The used Universal Disc Format (UDF) is a standardized and platform independent File System Format. The Archive Volumes of PoINT Storage Manager are based on this standardized format.

Large Archive Volumes improve the writing speed and accelerate a later migration to a new archive storage. Optionally, a file can be split over several Archive Volumes.

Even without PoINT Storage Manager companies can access archived files, because the created Archive Volumes are readable with on-board tools of the respective operating system (Windows, Linux, MacOS). Additionally, file lists can be exported as CSV files for allocation by means of the Data Browser.

Optionally, Archive Volumes can be encrypted. PoINT Storage Manager uses an AES256 and CBC based encryption method on block level.
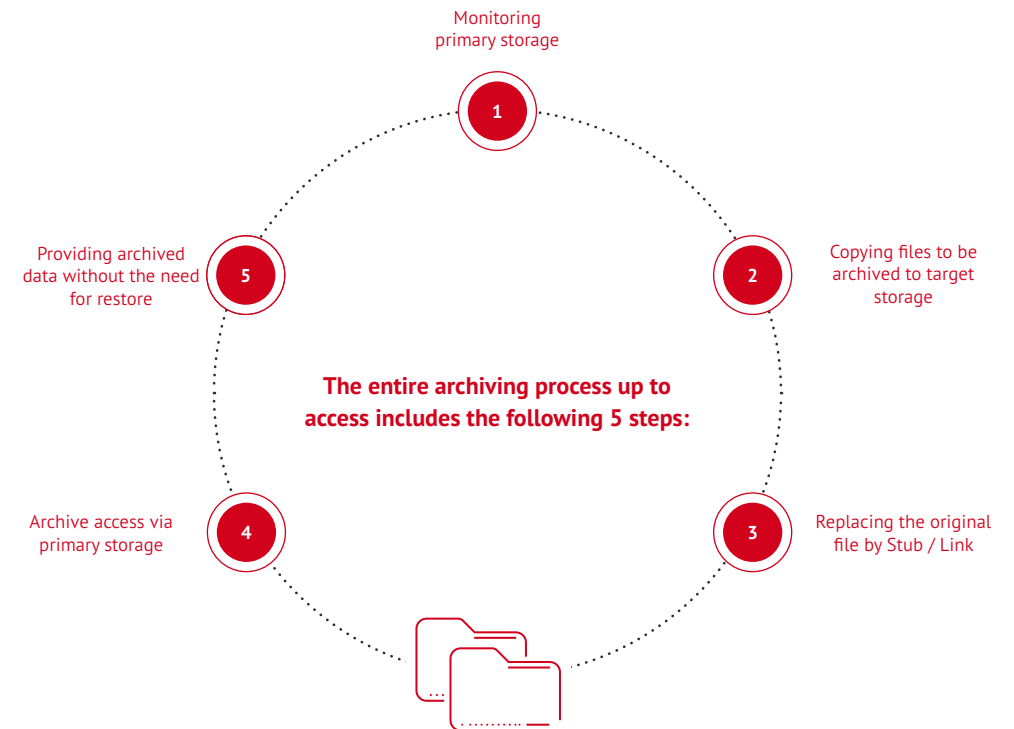
Archive migration

If data is stored over a long period of time, multiple changes of storage technology are necessary. PoINT Storage Manager offers a comfortable and interruption-free migration of the archive which can be set up in just a few steps. The actual transfer process takes place in the background so that no business interruption is necessary during the migration.

As PoINT Storage Manager works with large archive volumes, the transfer rate is drastically accelerated compared to a copy process of single files. After the migration, the IT administrator receives a log file with a protocol of the migration.

Transparent File Tiering and Archiving

Monitoring primary storage

1

Providing archived data without the need for restore
5

2
Copying files to be archived to target storage

The entire archiving process up to access includes the following 5 steps:

Archive access via primary storage
4

3
Replacing the original file by Stub / Link

Files to archive

Creation of Archive Volumes (UDF containers) in the image directory

Writing the finalized Archive Volumes to the archive device

Use Cases

PoINT Storage Manager is used by many customers worldwide, including small businesses with a low data volume of a few TB as well as large enterprises with several PBs in distributed structures. The most common use cases are listed below.

Compliance through archiving

Reasons for long-term archiving of data can vary. In addition to commercial interest, there are in particular legal requirements that make it necessary to retain data for years or decades. The most important technical points to take into account are data protection, migration to new storage technologies and support for all major storage systems and technologies.

—— CHALLENGES

- Archiving of inactive data and data to be archived
- Automatic and/or user-controlled archiving
- Modification protection (WORM) for archived data
- Retention management at archive storage level
- Archive storage migration without access interruption

—— SOLUTION 1 - AUTOMATIC ARCHIVING

- Realization of a two-tier HSM architecture
- Automated and policy-based archiving
- Transparent file access to archived data via the primary storage file system

—— SOLUTION 2 - USER CONTROLLED ARCHIVING

- Standardized file system (CIFS/NFS) for all applications
- Support for different storage technologies and systems
- Instant archiving to configured storage systems

—— BENEFITS

- Efficient use of primary storage through automatic archiving of inactive data
- Holistic approach through Archive File System
- Permanent availability of all data
- Fulfillment of compliance requirements through the Archive Tier
- Future-proof concept through integrated migration functions
- High investment protection through storage of archived data in standardized format
- Cost and time savings through reduced backup data volume on primary storage

PoINT Storage Manager

Archive Storage

Primary Storage

Centralization through Private/Public Cloud

Difficulties in companies with multiple sites and a decentralized storage infrastructure are caused by the high management effort and the associated rising costs. In order to downsize locations that have grown on the IT side, cold data should be moved from primary storage and stored centrally. A private cloud with high-availability object storage is very well suited for this.

SOLUTION

- Policy-based and automatic file tiering to a central private or public cloud
- Transparent access to migrated and archived files without restore
- Flexible storage architecture
- Integrated replication and migration functionality

CHALLENGES

- Integration of a private or public cloud
- Reduction of the required primary storage capacities per site
- Reduction of administrative effort
- Prevention of vendor lock-in

BENEFITS

- Efficient use of primary storage systems
- Fulfillment of compliance and archiving requirements
- Minimization of the backup window and reduction of backup storage
- Independence from public and private cloud providers

Public / Private Cloud

CIFS / NFS

Distributed sites

PoINT Storage Manager

Optimizing infrastructure using ILM and tiering

On the one hand, primary storage is designed for high performance requirements and is highly available thanks to methods such as synchronous data mirroring. On the other hand, however, it is precisely these two points that increase the costs per TB. In addition, data back-up is becoming increasingly difficult due to the massive increase in data volume. By migrating inactive data, the primary storage is significantly relieved and data backup becomes easier again.

—— CHALLENGES

- Optimized use of existing resources
- Unchanged workflow for users and applications
- Simplification of infrastructure
- Reduction of backup data volume

—— SOLUTION

- Policy-based file tiering and archiving of cold data
- Multi-level storage architecture
- Transparent read access without restore
- Consolidation of primary storage systems

—— BENEFITS

- Optimal use of primary storage for productive data
- Reduced data volume on the primary storage system
- Cost and time savings through reduced backup data volume
- High investment protection due to vendor independence

Old infrastructure

Secondary Storage

Primary Storage

Primary Storage

PoINT Storage Manager

New infrastructure
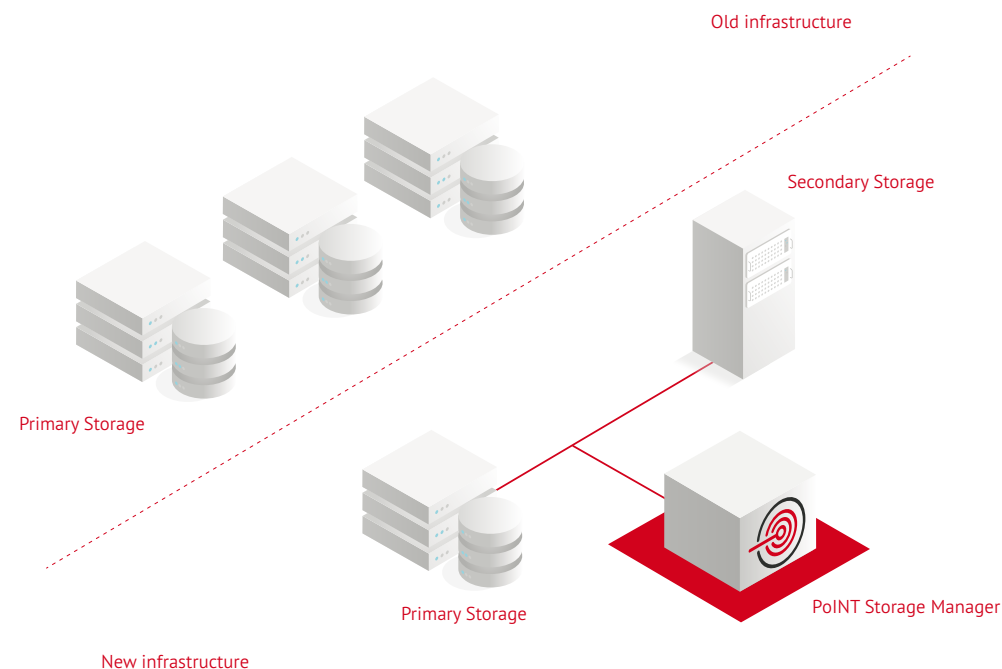
Hybrid infrastructure integrating on- and off-premises storage

Cloud offers for storing data are very attractive for many companies. The advantages include fast, flexible and on-demand provision of services, monthly billing for the resources used and the avoidance of maintenance work or upgrades by in-house IT staff. The advantages of on- and off-premises solutions can be combined through a hybrid concept.

—— CHALLENGES

- Integration of on- and off-premises solutions
- Unchanged workflows for applications or users
- Independence from cloud solution providers

—— SOLUTION

- Multi-tier and hybrid storage architecture
- Homogeneous integration of on-prem and off-prem solutions
- Flexible individual set of policies to represent the workflow
- Transparent read access without re-storage
- Data protection through encryption

—— BENEFITS

- Native support of different on-prem and off-prem solutions
- No modifications to existing workflows
- Fulfillment of compliance requirements
- Vendor independence through migration functions
- High reliability through synchronous replication

Off-Premises

On-Premises

PoINT Storage Manager

Secondary / Archive Storage

Primary Storage

Data Archiving

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.