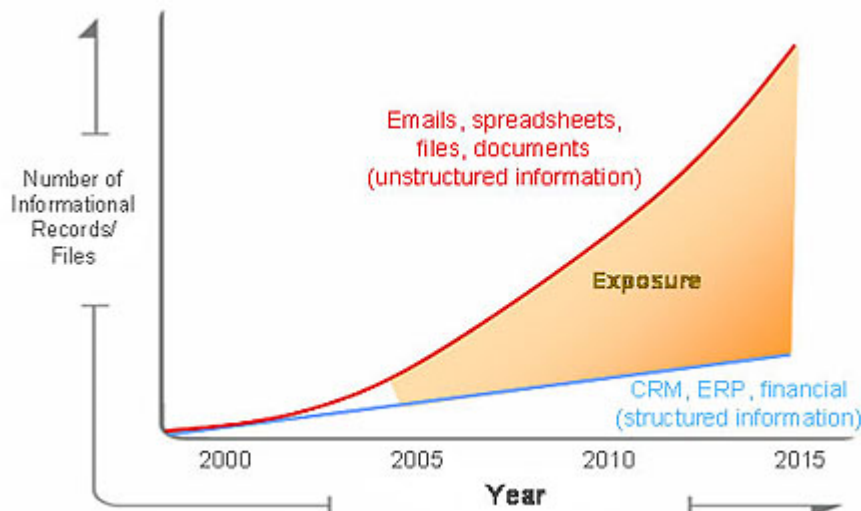


Automated Tiered Storage by PoINT Storage Manager

Optimizing the storage infrastructure concerning cost, efficiency and long-term availability

Growth of Unstructured Data

Unstructured data is growing at a greater rate than any other type of enterprise data. Market analysis and research (e.g. from IDC) clearly prove this trend. Whereas structured data, like the content of a database, has a moderate growth, unstructured data, like emails and their attachments, is causing a growth of enterprise data of typically 50% to 100% per year.



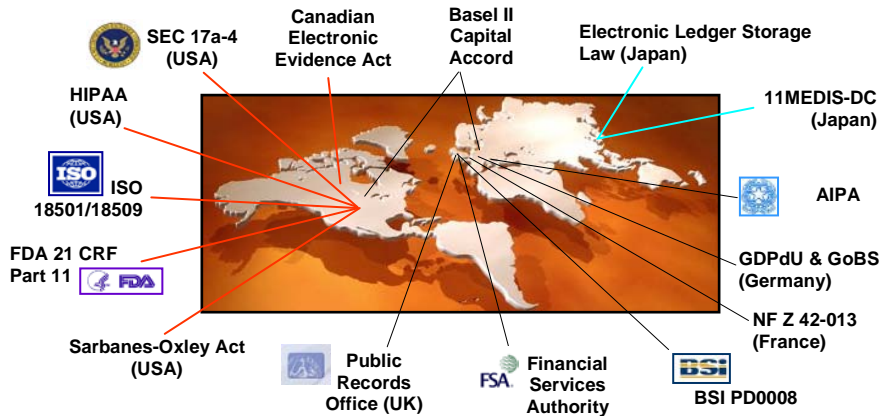
Source: Rational Retention Europe Limited

When considering the current situation in the industry with non-growing or even decreasing IT budgets, this leads to the demand for an intelligent data management concept addressing this

development, finally requiring a multi-tier storage approach.

Compliance Requirements


Another aspect of increasing importance in today's enterprise storage infrastructures is the requirement to fulfil compliance regulations. World-wide more than 25.000 regulations exist, the figure below shows some of them.



Compliance regulations and risk management require retaining and protection of business information from premature deletion or malicious modification. An intelligent storage management solution can help to reduce this business risk by providing as part of a tiered storage concept a storage layer which prevents alteration or deletion of stored data before the compliance retention period has elapsed.

Data Lifecycle

An analysis of a data object in a storage system is likely to reveal that initially it is accessed and updated very frequently. When the data object is growing older, its access frequency decreases to level which is almost negligible, as illustrated in the following figure:




According to this situation data can be considered to be in one of the following stages: active, less active or almost static, ready to be archived. At the same time as data is less frequently used the corresponding data volume is growing. This leads to the requirement that during its lifetime data should be moved to different physical storage locations because it would be too expensive and would not make sense to store all data on high performance storage systems (e.g. SAS RAID or SSD). Depending on the location of data during its lifecycle, it needs to be relocated to the most

appropriate storage device in a multi tier storage architecture. This requirement can be fulfilled by a storage management software which automatically migrates data to an appropriate storage tier depending on its status.

Typical Storage and Backup Situation

Today, most companies use only one single type of hard disk system, typically SCSI, FC or SAS RAID, as their primary storage and a tape based system for backup purposes. To fulfil the growing storage capacity requirements primary storage is usually upgraded by additional expensive hard disks although these disks are filled with more and more static data as a consequence of the lifecycle of data as described above. Additionally, the time to perform backup increases and resulting also in an increasing number of tape cartridges.



Storage Technologies

A tiered storage concept can solve this dilemma and can fulfil the requirements as described in the sections above by combining the advantages of different storage technologies into an optimal storage mix. The following list presents a short overview about the major technologies existing today as well as their typical characteristics.

- **SSD (Solid State Disk)**
Solid State Disks provide highest performance at high price. In a multi tier storage architecture this technology would typically be used for active data.
- **SAS/FC Disks**
RAID systems consisting of SAS or FC disks have lower performance figures compared to SSD but today this technology is still used as primary storage. Costs of SAS/FC based systems are still relatively high.
- **SATA Disks**
The performance of SATA based hard disk systems is lower compared to SAS/FC disks, however this technology has evolved as the most cost-effective drive technology over the last years.
- **LTO Tape**
LTO tape technology shows the best costs per TB ratio, however this technology does not provide random access capabilities which makes it inappropriate for use as primary storage. In

a tiered storage concept this technology is well suited for replication and (to some degree) for archiving.

- **Blu-ray Disc (BD), Optical**


Optical, in particular Blu-ray Disc (BD) technology perfectly fulfils all requirements for long-term archiving. Compared to tape this technology provides random access capabilities with adequate performance figures.

Tiered Storage Concept

According to market research and analyses, e.g. IDC, May 2009, a tiered storage concept incorporating the capabilities of the above described storage technologies provides enormous advantages regarding all critical aspects which are performance, capacity, compliance and cost. As illustrated in the figure below an approach consisting of three storage tiers fulfils all described requirements, where each tier implements a dedicated functionality:

- Tier 1 – “Performance Tier”
- Tier 2 – “Capacity Tier”
- Tier 3 – “Archive Tier”

Tier 1 consists of storage with highest performance. Tier 2 offers high capacity ensuring that stored data remain always accessible. Tier 3 is optimized for long-term archival and offers additional properties, in particular demanded for compliant archiving.



In the scope of this model Tier 1 typically represents the existing primary storage consisting of high performance, highly reliable hard disk systems (e.g. FC/SAS RAID) which are typically expensive.

Tier 2 implements a storage layer consisting of cost optimized storage systems which provide high capacity at lower performance (e.g. SATA RAID). This tier is typically used for the migration of data – in the following called “purging” (to distinguish between storage migration and data migration) – of data from Tier 1 which becomes less active.

Tier 3 implements the storage layer for long-term archiving, e.g. realized by optical or tape, and supports also offline media. Data stored in this tier is under control of a retention management and can only be altered or deleted under specific conditions, e.g. not before the corresponding retention

period has elapsed.

The described approach requires an intelligent software solution taking care of all aspects described above - PoINT Storage Manager.

Realization by PoINT Storage Manager

PoINT Storage Manager is a software solution which implements a multi tier storage architecture by means of a policy-based data management with automated purging and migrating between different tiers and within one tier. Independent of the physical location of stored data PoINT Storage Manager provides transparent file system access through Tier 1 ("Performance Tier") for applications.

Tier 1 – Performance Tier

In the tiered storage concept implemented by PoINT Storage Manager already existing hard disk systems typically used as primary storage are considered as Tier 1. This tier is supported by PoINT Storage Manager for the following hard disks systems and interface technologies:

- local NTFS hard disk partitions connected as direct attached storage (DAS) which are under control of the server-based PoINT TAFS Agent (see section below),
- SAN LUNs formatted as NTFS controlled by the server-based PoINT TAFS Agent,
- remote NTFS hard disk volumes connected as network attached storage (NAS) controlled by the client-based PoINT TAFS Agent,
- NetApp FAS Volumes controlled by the PoINT NetApp FAS Agent .

According to the policies which are specified by the administrator, PoINT Storage Manager automatically purges data to "lower" tiers ("Capacity Tier", "Archive Tier") but preserves transparent file system access through Tier 1. The "Pass Through" mechanism of PoINT Storage Manager ensures that purged data is delivered directly to applications from all tiers, which means that files are not re-stored again in Tier 1. The policies do also provide automatic retrieval of purged data into Tier 1, e.g. in cases where purged data is accessed again more frequently.

In addition to the interface technologies described above PoINT Storage Manager is also able to monitor standard NAS systems which provide a CIFS/NFS interface. For these types of systems PoINT Storage Manager provides policy based file copying (replication) to the Capacity and/or Archive Tier. Optionally source files can be deleted after successful copying.

Tier 2 – Capacity Tier

PoINT Storage Manager supports in this tier all hard disk systems (DAS, NAS, SAN), which offer a standard file system interface, e.g. FAT32, NTFS, CIFS/NFS. Files purged from Tier 1 are stored in their original representation as files. This means files and directories keep their original names and structure which allows read-only access to files stored in Tier 2 in case the Tier 1 storage fails.

Tier 3 – Archive Tier

As described, Tier 3 implements the archive storage layer for long-term archiving in the tiered storage concept of PoINT Storage Manager. Based on administrator defined policies data is purged to this tier, stored and protected against modifications ("WORM Mode") .

The following devices and connection types are supported:

- optical devices (single drives, autoloaders and libraries) direct or network attached (DAS/NAS),

- tape devices (single drives, autoloaders and libraries) direct attached (DAS),
- NAS devices accessible through an UNC path, e.g. NetApp FAS incl. SnapLock™.

Archived data is stored in UDF formatted images files (“Container Files”) generated in the scope of an automatically scheduled PoINT Storage Manager archiving job. A file which is archived in Tier 3 can also exist as copy and link in Tier 1 and/or Tier 2 and – depending on the settings – files accessed through Tier 1 can also be protected in these tiers against unauthorized modifications. If the option „File Versioning“ of PoINT Storage Manager is activated, changes of the archived data result in new versions. Thus the original file is protected against modifications at any time. Depending on the policies it is possible to delete files in Tier 1/2, e.g. in case of missing data access for a specific period of time. Either manually or policy-driven these files can be recovered into Tier 1/2 to make them again available for transparent file system access.

PoINT Storage Manager provides integrated support for LTO tape and for most optical formats (Blu-ray disc, DVD/CD, ...) available in the market. This means that no third party software products are required for their integration.

For administrating changed media like tapes and optical discs mostly used by libraries, PoINT Storage Manager offers an off-line media management. Thus a cost-effective and seamless method for storage enhancement is provided. Furthermore comprehensive functions for data search on off-line media is available. Additionally energy cost can be saved as no energy is used by storing media off-line.

Automated Purging

PoINT Storage Manager implements automated and transparent purging of data stored in Tier 1 by two mechanisms:

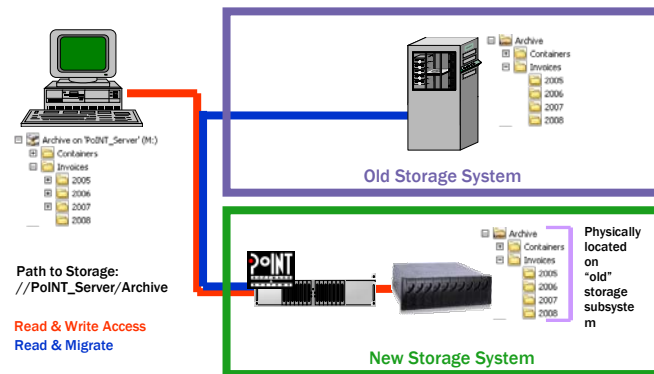
- **PoINT Tiered Archive File System (PoINT TAFS) Agent**
In Tier 1 storage which consists of e.g. performance optimized hard disk systems with Windows NTFS file systems, the PoINT TAFS Agent realizes transparent purging by replacing files purged to Tier 2 and/or Tier 3 by links (“Windows reparse points”). Unchanged access is retained for applications and users to all files originally stored in Tier 1 with unchanged path names, even if they are physically located on “lower” tiers.
- **PoINT NetApp FAS Agent**
The same functionality as offered by PoINT TAFS Agent is also provided for NetApp FAS systems by the PoINT NetApp FAS Agent. This module performs transparent purging of files according to configurable rules by NetApp ONTAP file system – without any need for changes in the application. Thereby, files purged from the NetApp FAS systems are replaced by so called “Stubs” (references) and stored safely by PoINT Storage Manager on Tier 2 and/or Tier 3.

In addition to transparent data migration from Tier 1 to “lower” tiers PoINT Storage Manager is also able to purge data between devices configured within Tier 3. This functionality allows migration of data from an “old” device (intended to be replaced) to a “new” device without interruption of the Tier 1 file system operability.

Storage Migration

A main feature of PoINT Storage Managers is the replacement of „legacy“ storage systems by a transparent migration to a new storage architecture. Hence an existing data base of a storage

systems which has to be replaced can be migrated without any business interruption. The data access remains even during the migration.



Data Protection

Data administrated by PoINT Storage Manager are secured and protected by versatile methods. The automated data replication (transparent for applications) provides the redundant storing of data on the configured storage systems. Therefore a system failure cannot cause a complete data loss. The simultaneous use of different storage technologies (e.g. hard disk and tape) with their specific advantages is an important module for secure data storage. Especially the usage of removable media (e.g. tape and optical) enable the off-site storage of data. Of course also a recovery is supported, for example if a storage environment has to be set up right from the beginning.

Furthermore backups of SQL and Microsoft Exchange data bases can be carried out automatically by PoINT Storage Manager¹.

Data Security

PoINT Storage Manager offers a comprehensive block-level encrypting mechanism based on AES256 and CBC. Storage Vaults can be encrypted by individual „Encryption Keys“ which offers a higher data security.

Besides data encryption PoINT Storage Manager supports also functions for data authentication. The implementation is based on SHA-1 standard and enables a data authentication check for the identification of possible modifications. The so called “Authentication Check Jobs“ supply a detailed check protocol.

Cloud Storage Integration

PoINT Storage Manager does not only integrate local storage technologies into a tiered storage architecture, but also integrates Cloud Storage² reasonable into the storage infrastructure. This is realized by integrating these services homogenously into Tier 3 („Archive Tier“). In such a configuration it is recommended that the storage cloud should not be used as the only storage instance, but as a supplement for the “on-site” storage. PoINT Storage Manager provides the necessary configuration possibilities to fulfill this.

¹ The support of SQL and Microsoft Exchange database backup is planned for Q2/2011.

² The support of Cloud Storage services by PoINT Storage Manager is planned for Q2/2011.

Conclusion

PoINT Storage Manager helps to reduce the need for primary storage (“Performance Tier”) by purging inactive or “dead” information off the primary storage tier to lower-cost “Capacity and Archive Tiers” with transparent access to the data. When the primary storage tier holds less data, applications can run faster and backup and recovery operations can be completed more rapidly.

By this approach, PoINT Storage Manager does not only reduce acquisition costs for new storage hardware but also optimizes the backup process. As migrated data has not to be protected once again – it is already stored redundantly on configured Tier 2 or Tier 3 storage – both, storage capacity (e.g. less number of tapes) and required time for backup are saved.

In addition, PoINT Storage Manager can help to reduce business risks by providing a storage repository (“Archive Tier”) in which stored information cannot be altered nor deleted until the compliance retention period has elapsed.

In a multi-tier storage architecture supported by PoINT Storage Manager the entire system may consist of a combination of different storage technologies (e.g. Tier 1: SAS RAID or NetApp FAS, Tier 2: SATA RAID, Tier 3: LTO tape or Blu-ray optical) to make use of technology-specific advantages and to ensure safety by redundancies.

The tiered storage concept of PoINT Storage Manger provides the following advantages and benefits:

- efficient usage of expensive primary storage by purging inactive data to less expensive Tier 2 and Tier 3 storage,
- permanent availability of active data through Tier 1,
- reduced management effort by high degree of automation (“policy-based management”),
- uninterrupted operation of applications by file virtualization realized by “transparent” data migrations,
- cost savings by reduced backup data volume,
- fulfilment of compliance requirements by Tier 3 - the “Archive Tier”,
- cost savings because no third party software products are required for connecting tape and optical systems.

PoINT Software & Systems GmbH believes the information in this publication is accurate as of its publication date. Such information is subject to change without notice. PoINT Software & Systems GmbH is not responsible for any inadvertent errors. The PoINT logo is a registered trademark of PoINT Software & Systems GmbH. All other trademarks belong to their respective owners.

Automated Tiered Storage by PoINT Storage Manager.doc

Version 1.7 / CK 250211

PoINT Software & Systems GmbH

Eiserfelder Str. 316
57080 Siegen, Germany
tel: +49 271 3841 0
fax: +49 271 3841 151
email: info@point.de
<http://www.point.de>



Glossary

Autoloader (Optical)

Autoloader for optical discs (CD/DVD, Blu-ray Disc) are devices, which record and print automatically optical media. These machines exist of one or multiple recorder, a robotic for moving the media and a printer for the label-printing.

Autoloader (Tape)

Autoloader for tapes are devices, which exist of one or multiple tape drives and tapes, which are inserted automatically into the tape driver.

Compliance

The term compliance describes in general the conformity with and the fulfillment of legal and regulatory requirements..

Legacy Storage

Legacy storage describes storage systems with a mostly strategic meaning, which often have been used for a long period. These storage systems are mostly not state of the technology and have to be replaced. The migration of these systems is very often critical, as the business operations should not be interrupted and the access to the stored data must be possible.

Multi Tier Storage Architecture

A Multi Tier Storage Architecture consists of multiple layers (mostly of three tiers), where a storage systems with specific quality is assigned to. Therefore the advantages of all storage systems can be combined within the complete system. The management and migration of data between and within a tier is carried out by the storage management software (PoINT Storage Manager).

Purging

In relation with a tiered storage solution, purging means the migration (verschieben oder auslagern?) of files stored in a higher storage tier to a lower on. The original file is replaced by a reference file (called 'link' or 'stub'). This reference file shows the same file properties as the original file but has no file content. By accessing to a reference file, the management software (PoINT Storage Manager) provides that the file content is delivered from the purged original file.

Retention Management

The survey and control of retention periods for archived files is called Retention Management. This function is carried out by a storage management software (PoINT Storage Manager), which manages the classification of retention periods to the data and furthermore prevents the modifications of archived data within the stated periods.

Retrieval

While accessing purged data, it can be reasonable to copy file content back to the original storage location. This operation, which is carried out background by PoINT Storage Manager, is named as immediate recovery.

SAS RAID

SAS is the abbreviation for 'Serial Attached SCSI' and RAID for 'Redundant Array of Independent Disks'. A SAS RAID System is a combination of hard disks, with SAS interface topology. A special characteristic of SAS RAID is the higher performance and reliability in comparison with SATA RAID ('Serial ATA RAID'). However the SAS RAID systems are considerably more expensive.

SSD

SSD means Solid State Disk and is a storage system without mechanical parts. The data are stored on electronic modules, which are able to store data permanently even without energy supply. The transfer and access rate of SSD in comparison to hard disks is considerably higher.

Structured Data

Structured data are data, whose contents are clearly defined and discernible (e.g. data in databases, ERP systems, accounting programs)

Transparent File Access

In combination with a tiered storage solution, transparent file access means that the access to a file in general, is possible by the file system interface. This access is independent from the storage device where the file content is physically stored. While accessing the file, PoINT Storage Manager provides that the data are delivered from the storage system, where the data are stored at this time. The storage location can be changed according to the defined policies, but the access interface is unaffected.

Unstructured Data

Unstructured data consist of optional information objects, whose content cannot be made directly accessible and where no separation from content and meta data exists (text documents, audio, pictures, videos, emails, presentations,...).