

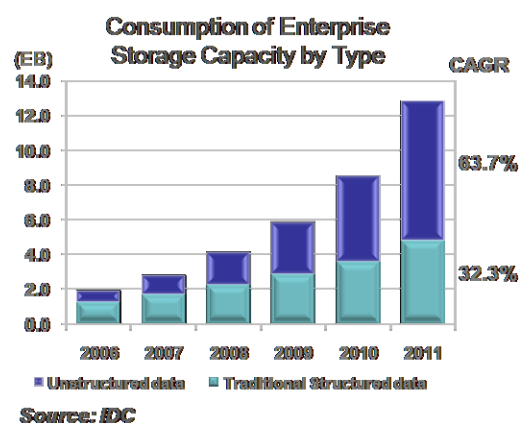


Tiered Storage by PoINT Storage Manager

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

Growth of Unstructured Data

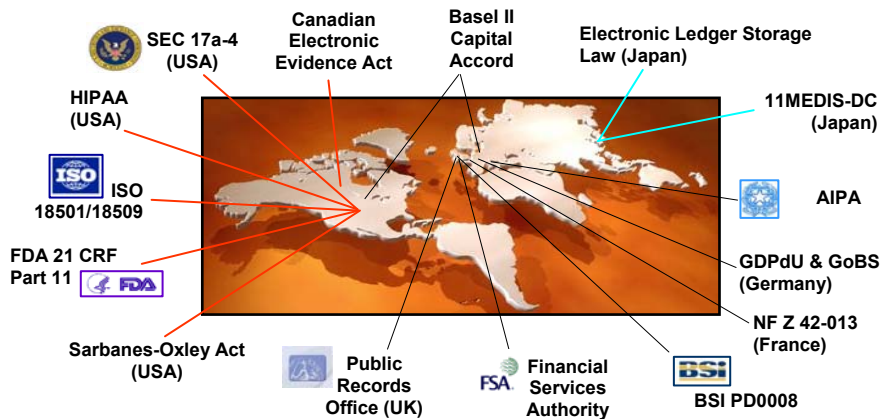
Unstructured data is growing at a greater rate than any other type of enterprise data. Market analysis and research (e.g. from IDC, see figure below) clearly prove this trend. Whereas structured data, like the content of a database, has a moderate growth, unstructured data, like emails and their attachments,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



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 to migrate 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 with 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 migration between tiers and within a 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 with standard NTFS interface 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 with 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 (see section below).

According to the policies which are specified by the administrator, PoINT Storage Manager automatically migrates 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 migrated 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 restore of migrated data into Tier 1, e.g. in cases where migrated 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

Tier 2 is realized by one or more cost optimized hard disk systems and/or optical libraries. The following systems and connection types are supported:

- hard disks (DAS/SAN or NAS),
- optical libraries (DAS).

PoINT Storage Manager supports in this tier all hard disk systems which offer a standard file system interface, e.g. FAT32, NTFS, CIFS/NFS. Optical libraries are supported as direct attached storage (DAS). Files migrated 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.

PoINT Storage Manager allows multiple devices for usage within Tier 2. The software is able to move files from one Tier 2 device to another or to generate copies on multiple Tier 2 devices, whereas the file system view presented by Tier 1 is not affected and the original file structure remains unchanged. This way PoINT Storage Manager realizes file virtualization with transparent file system access through Tier 1.

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 migrated to this tier and protected against modifications (“WORM Mode”) under control of the retention management module of PoINT Storage Manager (see section below).

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/or 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. 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 restored into Tier 1/2 to make them again available for transparent file system access.

PoINT Storage Manager provides integrated support for most tape and optical storage devices available in the market. This means that no third party software products are required for their integration.

Transparent Migration

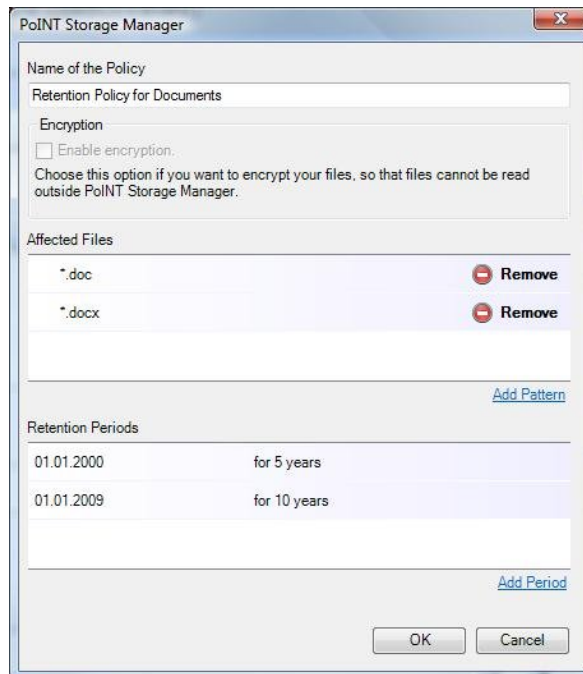
PoINT Storage Manager implements transparent migration of data stored in Tier 1 by the following two mechanisms:

- **PoINT Tiered Archive File System (PoINT TAFS)**
In Tier 1 storage which consists of e.g. performance optimized hard disk systems with Windows NTFS file systems, the PoINT TAFS realizes transparent migration by replacing files migrated 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 is also provided for NetApp FAS systems by the PoINT NetApp FAS Agent. This module performs transparent migration of files according to configurable rules – without any need for changes in the application. Thereby, files migrated 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 migrate data between devices configured within Tier 2 and 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.

Retention Management

The retention management module of PoINT Storage Manager monitors data stored in Tier 3 (“Archive Tier”) based on a set of rules called “Retention Policies”. A Retention Policy defines an arbitrary number of retention periods, every period associated with a validity period. In addition every Retention Policy specifies the file types, monitored directories, etc. it is used for. A Retention period can be specified as fixed period, but also by a fixed end date (including “infinite”).



The screenshot shows the 'PoINT Storage Manager' dialog box for configuring a retention policy. The title bar reads 'PoINT Storage Manager'. The main content area is divided into several sections:

- Name of the Policy:** A text field containing 'Retention Policy for Documents'.
- Encryption:** A section with a checkbox labeled 'Enable encryption.' and a note: 'Choose this option if you want to encrypt your files, so that files cannot be read outside PoINT Storage Manager.'
- Affected Files:** A list of file patterns with 'Remove' buttons. The list contains:*.doc (Remove) and *.docx (Remove). Below the list is an 'Add Pattern' link.
- Retention Periods:** A table with two rows:

Start Date	Duration
01.01.2000	for 5 years
01.01.2009	for 10 years

Below the table is an 'Add Period' link.

At the bottom of the dialog are 'OK' and 'Cancel' buttons.

WORM Mode

A file under control of a Retention Policy is set to WORM (write-once read many) mode, which means that it cannot be modified nor deleted. This is also supported for corresponding links in Tier 1 which maintain a reference to an archived file. If the File Versioning feature of PoINT Storage Manager is enabled, modifications of an archived file result in a new version of that file. At any time the originally archived file is protected against changes.

For every archived file with monitored retention exactly one Retention Policy is permanently assigned. Assignment can be done automatically, e.g. based on the file type or the monitored directory, or manually.

File Deletion

After expiration of the retention period an archived file is released for deletion. Depending on the used storage technology in Tier 3 deletion of a file is performed in one or two steps. First and immediately at the time of deletion the file is logically deleted in file system structure and therefore no longer accessible or visible by applications. If the storage technology allows physical deletion (e.g. hard disk systems) the corresponding files are also physically deleted.

File Encryption

Retention management can be combined with data encryption. This means each archived file is encrypted and the corresponding decryption keys are maintained in the database of PoINT Storage Manager. In this case deletion of a released file does also result in the deletion of the corresponding decryption keys. This means that an encrypted archived file cannot be decrypted and therefore it is

no longer readable. Obviously, file encryption makes much sense for (physical) WORM devices, e.g. LTO WORM tape or Blu-ray Disc, if configured as Tier 3.

Conclusion

POINT Storage Manager helps you to reduce the need for primary storage (“Performance Tier”) by migrating inactive or “dead” information off the primary storage tier to lower-cost “Capacity and Archive Tiers”. When your 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 your business risk 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 BD 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 migrating 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.

Tiered Storage by POINT Storage Manager.doc
Version 1.6 / CS 13.04.2010

©2010 POINT Software & Systems GmbH

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>

