QSTAR TECHNOLOGIES

SECURING THE FUTURE OF YOUR DATA®

QStar White Paper

# Archive Storage Management

Table of Contents


Introduction 1
The Data Storage Dilemma 1
The Solution 2
The QStar Solution 3
The QStar Product Line 4
How QStar Software Works 6
- QStar Integral Volumes 7
- Hierarchical Directory Structure 8

Media Management 9
- The Hard Disk 9
- Media Sets 9
  - Ease of administration
  - Ease of handling
  - Information security
  - Transportability
  - Duplication
- Sparse Mounting 10
- Device Management 11

QStar Implementation 12
- QStar QSCSI Driver
- QStar Standard Data Format (SDF) File System
- QStar CD/DVD-ROM File System Manager
- QStar Universal Disk Format (UDF) File System
- QStar Library Driver
- QStar Volume Librarian
- QStar Magnetic Cache File System

QStar's Capabilities 13
- Automatic Data Migration
- Dynamic Storage Allocation
- Removable Media Management
- On-Disk Media Format
- Disaster Prevention and Recovery

QStar HSM 15
- Migration Management 15
- The Migration Cycle 15
  - Primary status
  - Archived status
  - Replicated status
- Magnetic Disk Utilization 15
- Utilization Levels 15
  - Total Capacity
  - High and Low Primary Capacity
  - Replication Reserve
- Archiving 15
  - Periodic Archiving
  - Demand Archiving
  - User Initiated Archiving
  - All Delayed Events
  - Low Primary Capacity
  - Agetime

- **File Migration Modes** 19
- Per-File Migration Control
  - Full File Migration
  - Archive Never
  - Keep In Cache
- Media Management
- **Media Types** 20
- **Media Format** 21
  - Transportability
  - Flexibility
  - Disaster Recovery

**Library Volume Management** 22

- **Volume Catalog**
- Electronic Labeling
- Location Tracking
- Dynamic Allocation
- Media Set Quotas

**QStar Data Director** 24

- Data Mirroring
- WAN and LAN Support
- Hot Standby
- Improved Disaster Recovery
- Library Load Balancing
- File system Resynchronization

**QStar Simulated Disk and Simulated Tape Jukebox** 26

**QStar API** 27

- Low-Level Device Inquires
- Library Control
- Request and Error Trapping
- Seamless Integration
- Expert Help

**Advanced Features of the QStar Software** 28

- **Compaction** 28
  - Optimized Media Usage
  - Absolute Deletes
- **Digital Signature** 28
- **Privileged Users** 29
- **Automatic Copydisk/Copytape** 29
- **Email Notification of Pending Requests and Errors** 29
- **Grace Period** 30
- **Automounter** 30
- **Fibre Channel Support – SAN Solution** 30

**QStar Software and Open Systems** 31

  - Transportable
  - Networkable
  - Scalable
  - Flexible

**Configuring for Individual Applications** 33

- **Interactive Model** 33
- **Session-Oriented Model** 34

# Introduction

## The Data Storage Dilemma

IT organizations today are facing relentless growth of data volumes. Applications are creating more and larger files than ever before. Industry organizations and regulators are mandating the type of data that must be retained and for how long it must be retained. The length of time data is to be stored can be as much as 50-years or more. However, much of this data will never be accessed again, until its time to delete if from the system. The price of hard disk storage has come down significantly, and as such, the solution to the problem has been to buy more disk storage solutions, however, disk storage is not designed to manage data continuously for many years. Many organizations are now finding that data stored on disk based archive systems are failing in such a way to create higher maintenance costs than expected with unacceptable data loss. They are also finding that as time goes on the problem is only getting greater as more data is stored on larger systems, with more and larger hard drives.

In developing a solid archive strategy, one must consider the goal. The first things to consider regarding the archive solution are the problems you're trying to address. What are the key issues that need to be overcome?

Data reduction on primary storage:

- Is there too much data to manage effectively on the current system?
- Is there too much data residing on costly high performance RAID or NAS?

- What is in the file system and who owns it?
  - How much of the file system is redundant?
  - How much of the file system is garbage or unauthorized files?
- Can and how much data should be stored off-line?
- What will be the tracking method if data is stored off-line?
- How often will archive data need to be accessed, and what will the network traffic look like?
- Are there automated policies in place to migrate data to archives?
- Are there key data storage compliance issues that govern your application?

By addressing these questions, organizations can begin to streamline data and enhance disaster recovery strategies while consolidating hardware and reducing operational overhead.
By identifying stagnant data in the file system and moving it to low cost archive media, IT departments will shrink backup windows and have the ability more efficiently manage hard disk stores while reducing the need for new hardware purchases, thus gaining an immediate return on investment, with both short and long-term benefits.

## The Solution

Many IT managers express the problem of having too much data to manage on the file system which creates a data protection problem. Hard disk systems need to be protected from disk failure; however, backup systems are not always fast enough to move a snap shot of the system to the backup device within a reasonable time frame. Data deletion is a possible solution, however, the IT manager is rarely the owner of the data, and therefore the IT manager is usually very reluctant to delete any files on the system, unless they have been deemed unauthorized, such as non-work related material. Vendors have thrown technology at the problem such as Disk based backup systems, which solve some of the transfer speed issues; however, they do not solve the capacity issue.

It is also useful to know what is in your file system. This can give IT managers a very clear picture of the types of files they have in their system, who owns the data, when the data was created, and how much of the data is redundant. By knowing how much data is redundant, allows IT managers to shrink their data set without deleting any files completely, just the redundant ones.

Once its been determined what is in the file system, a strategy can be formed to decide what get moved, what can stay and for how long. Regulations generally fall into two categories: how long must I retain a file, and must it be revision secure? When it has been determined what files fall under regulatory requirements, then polices can decide where the file is moved, how long it must remain there, and on what type of media is should reside based on time and security requirements. While QStar supports virtually all types of archive technologies, not all types of archive technologies are a fit for all data. The media chosen for the archive and if the media is to be kept on or off-line should be carefully considered.

Cost and data access requirements are typically the deciding factors if data is to be stored on-line, near-line or off-line. It is usually less costly to keep data off-line if it does not need to be accessed. If, however, the data needs to be accessed, even once, then off-line data will usually be more costly to load back into the system than if it had been kept in a near-line solution if the data is kept off-site. If stored off-line, an automated tracking system should be put into place along with a data security strategy, such as encryption in the event that the data should become lost or stolen.

Data access requirements should be a consideration as to what archive media type is to be used. Archive data that is likely to be accessed during its archive life should be stored on rotational media, where data that has little change of being recalled for any reason, might be better on a tape solution. In either case storage libraries can be used to keep the data near-line within a secure network for.

Data retention is another criterion for considering the media type. All media have a shelf life and a useable life. Data retention regulations will be the driving factor for how long data will reside in the archive. A media should be selected with a longer usable life that the regulation.
For example, most SATA hard disks have a manufactures specification between 3 and 5 years. Data should not be stored on SATA disks that are expected to be retained for longer than 3 to 5 years. Most tape media is expected to have a useable life of about 10 to 30 years and optical media is specified from 10 to 50 years.

Automation is the key to a successful archive. The ability to automatically send data to the archive and retrieve data back from the archive without changing the user experience in an active manner will determine the success of the system. If this process is manual, then the archive is likely to be ineffective. The user should not have to look for the data once it's been archived. User definable policies should automatically manage the data movement to and from the archive media. These policies should be set within the scope of any regulations that apply to the data being archived

Once an active archive management system is running on the network the migration of data to an archive platform, reduces the data set and thus shortens the backup window, and can help retain data to meet the many business regulations.

## The QStar Solution

QStar Software consolidates all of the physical mass storage subsystems on a network – hard disk, optical, or tape – into a single, hierarchically managed resource, or integral volume, treated as one logical entity *(Figure 6)*.

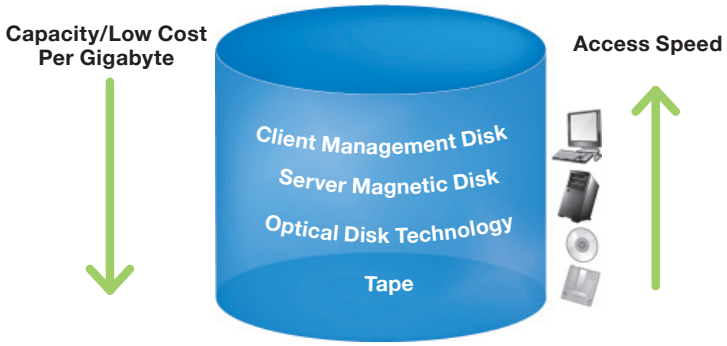


Using this concept, the most recently used data is available from hard disk, providing fast access for users. Less recently used data is stored on optical disks or other archival media and is automatically moved to the hard disk if a user accesses it. The movement of data between the hard disk and the archival media is completely transparent to the end-users and applications accessing it.

The software's automatic storage management provides the benefits of virtually unlimited storage capacity without sacrificing access time to critical data. The result is a network storage environment that is productive, cost-efficient, and manageable in helping users meet the business goals of their enterprise.

## The QStar Product Line

The QStar storage management product line includes the following software packages:

**QStar HSM**

**QStar Data Director**

**QStar MultiStor**

**QStar Network Migrator**

**QStar Simulated Disk Jukebox**

**QStar Simulated Tape Library**

**QStar CD/DVD Axxess**

**QStar CD/DVD Master**

**QStar API**

Each QStar product is designed to address specific application requirements *(Figure 8)*.

- QStar HSM supports more complex client/server systems with full-featured hierarchic storage management data migration to archive media.
- QStar Data Director expands the capability of HSM with the ability to provide high level mirroring and disaster recovery functionality for archive data protection.
- QStar's simulated tape and disk modules add the ability to manage a NAS or RAID system like a tape or optical library. With SDJ/STJ users can mirror NAS or RAID space to a storage library.
- QStar MulitStor and Network Migrator provide mulit-tiered based data migration between layers of storage, as well as the ability to manage data retention and deletion policies based on user definable policies. Both MultiStor and Network Migrator are fully compatible with the rest of the QStar product line.
- QStar CD/DVD Axxess adds the capability to automatically manage multiple compact disks and DVD-ROMs in varying formats.
- QStar CD/DVD Master enhances the Axxess product with a CD/DVD-R writing module.
- The QStar API provides the ability to seamlessly integrate between QStar's software and an ISV's application by allowing more control of the storage device through specialized API calls.

Each feature of QStar's data storage management products is described in more detail in following sections of this document.

| Features | HSM | Data Director |
|---|---|---|
| Application Compatibility | ✓ | ✓ |
| Data Accessibility | ✓ | ✓ |
| Dynamic Onode Allocation | ✓ | ✓ |
| Platter Interchangeability | ✓ | ✓ |
| Automatic Scheduling | ✓ | ✓ |
| Transaction Logging Cache | ✓ | ✓ |
| Automatic Data Migration | ✓ | ✓ |
| On-Disk Format | ✓ | ✓ |
| Dynamic Space Allocation | ✓ | ✓ |
| Disaster Recovery | ✓ | ✓ |
| Compaction | ✓ | ✓ |
| Advanced Error Logging | ✓ | ✓ |
| Remote Administration | ✓ | ✓ |
| Stand-alone Drive Support | ✓ | ✓ |
| Integrates with MultiStor | ✓ | ✓ |
| Integrates with Network Migrator | ✓ | ✓ |
| Data Mirroring | ✓ | |
| Hot Standby | ✓ | |
| Library Load Balancing | ✓ | |
| Automatic Failover | ✓ | |
| Filesystem Resynchronization | ✓ | |
| Digital Signature | ✓ | |

*Figure 8 – QStar Storage Management Product Line*

# How QStar's Software Works

### Solving the Data Management Problem

Traditionally, data in computer systems has been stored on hard disk, which provides quick response time, but a limited amount of storage space at a high cost. When the hard disk fills up, the information systems manager has to decide whether to buy more hard disks or archive the information off-line.

Generally, users don't need all of their data on-line all of the time, because only a small subset of their data is required at any given time. On the average, about 20% of the data available accounts for 80% of data use *(Figure 9)*.

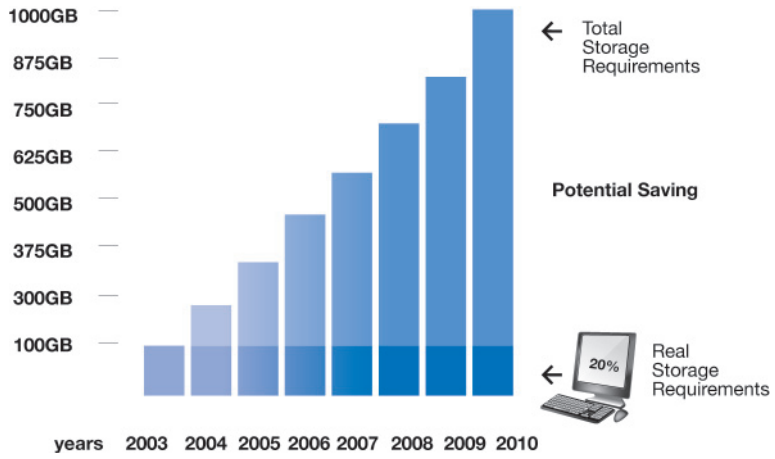


Even though users may access only a small number of their files on a regular basis, they perceive their use to be more frequent than it is, and are reluctant to store any of their data off-line. IT departments often set aside large amounts of free space to prevent capacity failures. Analysts report storage utilization rates from 20 to 50 percent. On top of this, consider that 60 to 80 percent of the data in these systems has often not been accessed for months or even years.

These architectures waste system resources, require more attention from administrators, increase backup costs and take longer to restore from an outage.

The dichotomy between these two views is the founding principle of the QStar's core software. The needs for both information systems managers and their end-users are satisfied by software that consolidates and controls all of the physical mass storage subsystems on a network. Then, users can easily access data, and systems administrators can control a single, logical entity.

## QStar Integral Volumes

QStar capitalizes on the best features of all of the media to provide users the benefit of increasing storage capacity without sacrificing access time to critical, valuable, or frequently used data. The speed of hard disk makes it ideal for creating, updating, and accessing data. The reliability, storage capacity, shelf life, and low cost of removable secondary storage make an ideal archival medium. Users' most frequently accessed data is stored on hard disk for rapid access, while less frequently used data is automatically and transparently migrated to less expensive removable media.

To further increase the on-line capacity of the system, QStar has the ability to manage a range of archive devices such as Content Addressable Storage (CAS) hard drive systems; optical disk, CD, DVD, and tape libraries.

A CAS device is a hard drive based RAID device usually with management software that minimizes the data redundancy stored on the CAS archive. In some cases it may have the ability to manage data retention within the storage device itself.

A library is a robotic device that varies in size and is capable of storing and accessing multiple removable media cartridges. A robotic arm moves the media to and from storage shelf positions to a read/write drive for near-line access. The storage capacity of current models ranges from 500-Gigabytes to over 100-Terabytes. Multiple libraries may also be physically connected to the same server to create an even larger pool of available resources.

All of the combined storage media that QStar manages is referred to as an integral volume *(Figure 10)*. An integral volume is the sum total of all of the items in the storage hierarchy.


To users, the integral volume seems like a standard hard disk with unlimited storage capacity. It is completely compatible with existing applications and on UNIX or Linux based platform, it may be NFS (Network File System) mounted by the systems administrator.

Integral volumes can also be completely off-line, such as volumes that contain infrequently accessed data. QStar manages all of the media in the integral volume as a single unit and knows the location of all of the files and media, including those that are off-line, at all times.

## Hierarchical Directory Structure

Like a standard hard disk, an integral volume has a hierarchical directory structure, and it can be navigated by using normal system commands. The difference is that at any point in time, a file on an integral volume can actually reside on different types of storage media. For example, the entire file can be on hard disk only or on secondary storage only; the entire file can be on secondary storage with a copy on hard disk; or part of the file can be on hard disk, while the entire file is on secondary storage. As depicted in a tree hierarchy *(Figure 11)*, the files controlled by QStar are composed of independent data segments. The data segments that make up the file are migrated according to user access patterns.

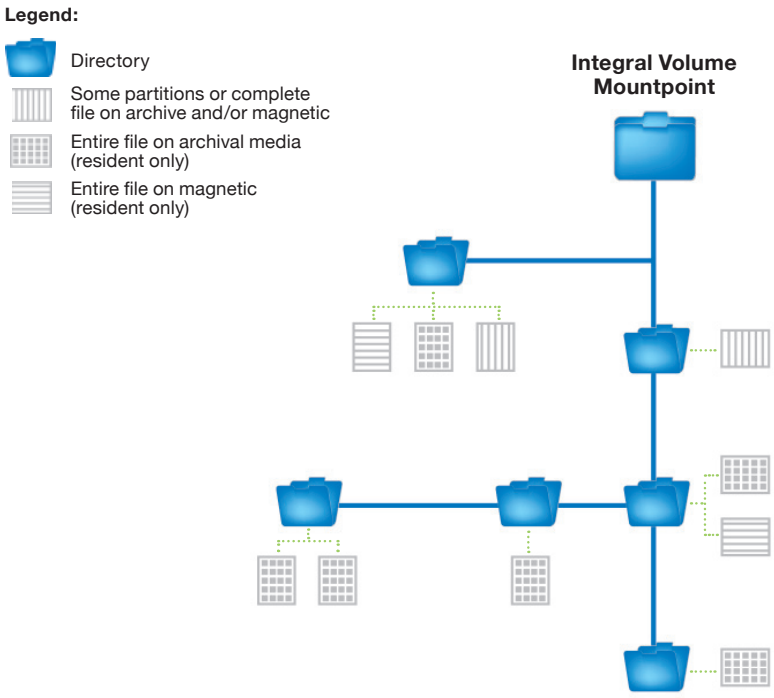


A copy of directory trees and file attributes always resides on the hard disk. This allows for faster file lookup, open, and directory listing operations without losing performance. In addition, a copy of this directory and file information is maintained on archival media, using the QStar volume format feature, to prevent catastrophic loss in case of total hard disk failure.

# Media Management

## The Hard Disk

In a QStar integral volume, files are always first written to and read from the hard disk; namely, users normally access the faster magnetic medium directly, rather than the slower archival media. As the hard disk fills up, the least recently used data is moved automatically to the archival media, where it can be easily retrieved later, if needed. The hard disk portion of the integral volume should be large enough to accommodate general usage, based on the average file size and the average number of users who simultaneously access the integral volume.

## Media Sets

A media set is a group of archival media (hard disk, CAS, UDO, Blu-ray, MO, DVD, CDs or tapes) in an integral volume. The media set consists of the sequenced surfaces of the physical archival media. Once the integral volume is created, the media set is treated as a single entity, even though it may consist of several individual media surfaces, which can be on-line or near-line (resident in the library) or off-line (stored outside the library).

By viewing the archival media as groups, or media sets, rather than single, autonomous units, QStar lets the application access large amounts of stored data without having to track the physical location of the file. The advantages of using media sets include:

- **Ease of administration:** Administrators can export a single file system rather than cope with multiple file systems. Migration across operating systems is also possible.
- **Ease of handling:** Large data files that require more than one surface are easily accommodated.
- **Information security:** Media sets can be organized and controlled by the type of information stored in them (e.g., accounting, personnel, and engineering), to which administrators can assign different levels of access security.
- **Transportability:** Complete media sets can be exported from a secondary storage device and then imported into another secondary storage device. The destination storage device can reside on the same system or any other platform that QStar supports, as the file system is transparent between operating systems.
- **Duplication:** Administrators can make fast copies of the media sets. This feature enables them to create baseline copies of critical sets and then store these media volumes off-line for disaster recovery.

QStar implements media sets according to each user's needs. Some users may want to have one single data repository, while others may prefer to departmentalize their data onto several sets. Administrators have complete control over the data, including; who has access it, how it is accessed, and when they may access it. Administrators can configure media sets to match the business needs in any combination of the following:

- A single media set per storage system, or across all available devices, or across the entire media library (on-line, near-line or off-line). Single media sets give the administrator greater ease of administration, handling, and transportability, and let them set common controls for entire sets.

- Multiple media sets per storage system, or across all devices, or across the entire media library (on-line, near-line or off-line). Multiple sets facilitate off-line storage and transportability. They also let administrators conveniently segregate data by type and set as many levels of security as required.

## Sparse Mounting

To use a media set, it first must be mounted. The mount process is a sparse mount, using information about the archival media resident on the hard disk instead of physically reading all archival media in the set *(Figure 12)*. Media sets can be mounted even if most of the media are off-line, although the then-current write medium must be on-line for the mount command to work.

The mounted media set appears to the host system as an ordinary mounted file system. Media sets support all normal file system characteristics including security and network access.

As media volumes are accessed, additional mount information is supplied to the hard disk. This effectively reduces robotic movement and brings the media set on-line very quickly.

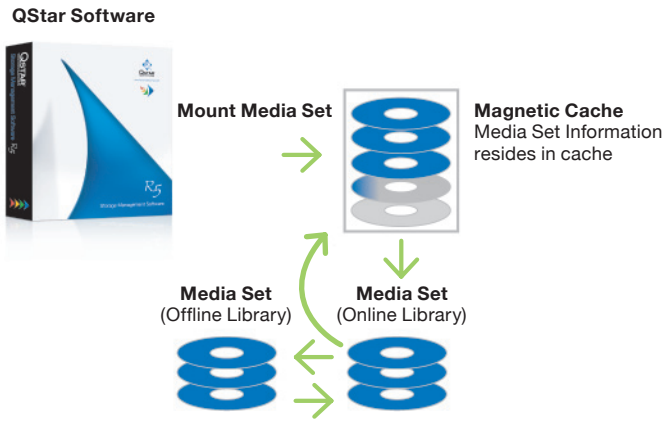

## Device Management

To manage all user requests as efficiently as possible, QStar views all devices, including configured multiple devices, as part of the same managed pool *(Figure 13)*. If a data request is made for a medium associated with an on-line device, it is resolved without human intervention. If a request is made for data that is not available on-line or near-line, the operator is prompted via email notification that a platter, disc or tape must be brought in from an off-line location. QStar uses any compatible device that is available to resolve the request. This feature gives the QStar Software the ability to maximize usage of all devices under its management.


Any archival medium in a set (except the current read or write medium), as well as complete media sets, can be exported from a library and then imported into another library. The software automatically senses that the medium belongs to an unknown set, and with two commands, the administrator can make this set and the data contained within quickly accessible to their users.

The software uses the aggregate set of devices, thus eliminating a single point of failure. If a drive in a library experiences a hardware error, the error is diagnosed, and the software automatically takes the faulty drive off-line, and continues to process data requests by using other available drives in the library. Administrators can also make duplicates of critical media sets to create baseline copies and then store these media sets off-line for disaster recovery.

## QStar Implementation

The QStar Software is composed of several integrated, functional modules *(Figure 14)*.
Each module has been tested and proven in operational environments.
The QStar modules are:

- **QStar QSCSI Driver**
  QSCSI, a prerequisite for installing all other modules and consists of high level SCSI UNIX or Windows NT device drivers for optical, tape, CD, and DVD drives and low level library queries.

- **QStar Standard Data Format (SDF) File System**
  Available as a faster alternative to the UDF file system, SDF provides support across disk, optical, UDO, Blu-ray, DVD, and tape media types.

- **QStar CD/DVD-ROM File System Manager**
  Utilizing an intelligent caching mechanism, this module provides an easy to use interface to see many different CD/DVD-ROMs as a single file system. With support for the ISO 9660 standard and UDF, grouping CDs or DVDs together under one mount point is easily configured.

- **QStar Universal Disk Format (UDF) File System**
  The UDF file system is provided as a world recognized non-proprietary file format. Available on either DVD, MO, UDO, or Blu-ray, the UDF ISO standard is accessible via most machine types without any special software.

- **QStar Library Driver**
  Library management software provides a complete application transparent interface to libraries, scheduling the insertion of a particular disk or tape into a drive based on demand for that volume, and then allowing access as in a stand alone drive environment.

- **QStar Volume Librarian**
  Automatically manages and electronically labels all removable media to guard against errors and automatically documents the location of each media.

- **QStar Magnetic Cache File System**
  MCFS, the migration manager, manages all the data on the hard disk media. It's goal is to achieve optimal performance for the user by tracking all the data in an integral volume and making real time decisions to either keep data on the hard disk or archive it to the storage media.

## QStar's Capabilities

QStar's software provides an automated, tiered based storage management technique by applying the concepts of virtual memory to network storage management. The software brings together the reliability and cost-effectiveness of multiple storage technologies by integrating all storage media types on the network into one, centrally managed environment. It does this while retaining the high speeds and performance levels of a hard disk, thereby making the most productive use of data storage while reducing overall costs.

By keeping all files easily accessible, whether they are stored on hard disk or on archival media, QStar's software satisfies both the information systems manager's need for centralized cost-efficient management of data storage and the end-user's desire for prompt, easy accessibility. Transparently migrating files between hard disk and archival storage media, based on frequency of use ensures that the hard disks will never fill up. Thus, the end-user has seemingly unlimited on-line storage capacity across the enterprise as needed. QStar's on-disk format also makes archival media easily transportable between cross-platform file servers and a wide variety of secondary storage devices.

QStar's tiered storage management of all storage media as one integral volume provides the following features and attendant benefits:

- **Automatic Data Migration**
  Data within the network volume is automatically migrated between the different levels of the storage hierarchy. Most recently used data is automatically promoted to the fastest available medium (hard disk) for instant accessibility. Less recently used data is automatically migrated to the most abundant, and least expensive, medium *(Figure 7)*.

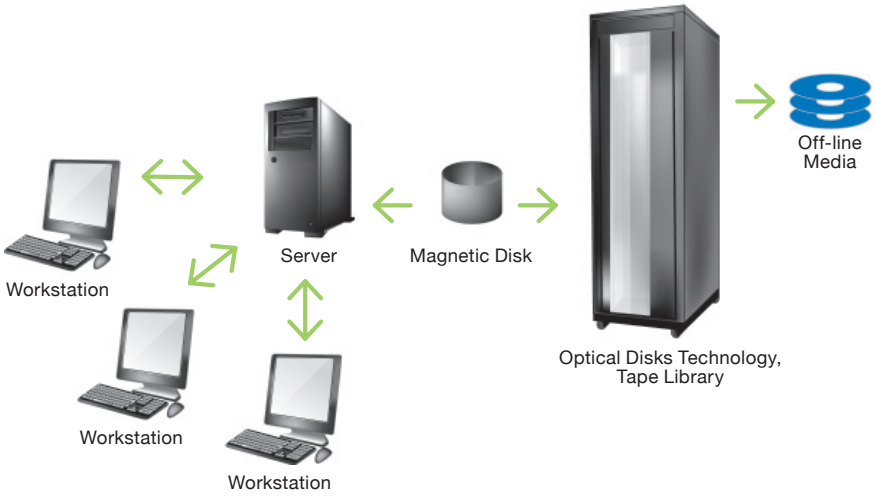


- **Dynamic Storage Allocation**
  The capacity of an integral volume automatically and transparently increases as required, providing users with virtually unlimited capacity for on-line data storage. This is done without intervention from the administrator.

### Removable Media Management

On-line and off-line media are internally and externally labeled, documented, and automatically managed to guard against operator errors. QStar's software handles all tracking and processing of removable media, storing their location in the media volume database so that media not currently in the drive can be identified and accessed. A special interface lets the customer transport data storage media between libraries on different hosts, even if they are running on a different operating system.

### On-Disk Media Format

All archival media are self-contained, with file directory information, data, and indexes on the same disk, or tape. This feature provides transportability between libraries and file servers.

### Disaster Prevention and Recovery

Because of the on-disk format, any hard disk in the storage hierarchy can be completely rebuilt from the underlying archival media, thus preventing catastrophic data loss.

# QStar HSM

## Migration Management

Migration is the movement of data between the hard disk partition and the archival media that make up the integral volume. The migration manager allocates the data to the appropriate storage medium to achieve optimal performance while maintaining low cost per megabyte for on-line or near-line storage. QStar MCFS is the migration manager module in the QStar Software. Written to use a transaction-logging cache, increased speeds of writes and advanced error recovery are standards.

## The Migration Cycle

All newly created data is first written to hard disk, called the source medium. Over time, new data is copied, or archived, from the hard disk file system to the archival media, called the target media. When a user requests access to data that is no longer on the hard disk, the data is copied, or replicated, from its archival medium back to the hard disk. The migration cycle continues when the replicated data is updated and later archived to target media *(Figure 15)*.

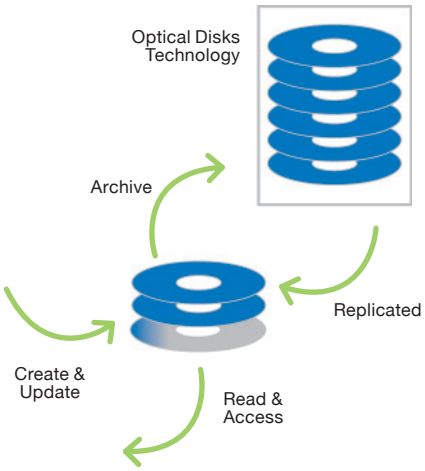


The reverse cycle occurs when a user accesses data from an integral volume and a copy of the archived data segments is not located on the source medium. The requested files are then replicated from the target medium onto the source medium. Changes to this file are made on the replicated file, and are archived to the target media.

If a file has not been accessed for a specified period, and the space occupied by that file is needed for other data, the file is archived and then removed from the hard disk. Directory information for the entire volume set is always maintained on the hard disk; facilitating magnetic media speed for routine file system operations, such as directory listings and file searches.

The migration manager tracks data by assigning a status to all data files:

- **Primary status:** means that the file on the hard disk is new or updated data that has not yet been archived to a secondary storage device.

- **Archived status:** means that the file has been written to the archival medium, but remains on the hard disk until an I/O request needs that space.

- **Replicated status:** means that the file has been archived to the target medium, but has been requested to return to the hard disk for I/O purposes. Replicated data on the hard disk is merely a copy of the same data that exists on the archival medium.

## Magnetic Disk Utilization

Ideally, the hard disk should be fully utilized up to its optimal capacity. Any room on the disk that is not being used for primary data is used for replicated and archived data. This increases the probability that a user's request for a file on the integral volume results in quick access to the cached data and does not require time-consuming access from the archival medium.

The migration manager tracks all of the data in an integral volume and makes real-time decisions about data placement. These decisions are based on:

- How recently the data has been accessed

- What proportion of the hard disk the system administrator has assigned for primary and for non-primary data

- How certain parameters are configured and certain commands implemented

The overall ratio of primary to replicated and archived data on the hard disk changes over time *(Figure 16)*. As new primary data is added, least recently used replicated or archived data is assigned non-cache status and deleted from the hard disk. However, any room on the disk that is not used for primary data is given to the most recently used replicated data, increasing the probability that a request for a file results in quick access.

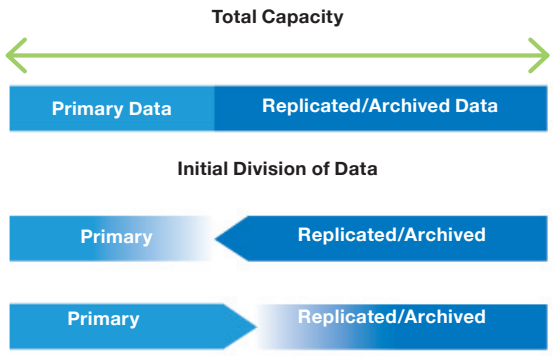

## Utilization Levels

The utilization levels set the boundary conditions for using the hard disk for primary, replicated or archived data.

- **Total Capacity**: This is the level at which the hard disk is fully used by primary, replicated or archived data. Total capacity is a configurable parameter. The migration manager ensures that the amount of data on the hard disk never exceeds the level set as the total capacity.

- **High and Low Primary Capacity**: These are the levels of capacity on the hard disk that are used to store primary data. When the primary data levels equal or exceed high primary capacity, the migration manager begins automatic archiving. The migration manager archives the primary data until all the data has been archived or until the low primary capacity level is reached.

- **Replication Reserve**: This is the space between total capacity and high primary capacity. This is the amount of hard disk space set aside only for replicated data. This reserve is useful when the system's administrator wants to achieve a balance in response levels between read and write access for users. The larger the replication reserve, the less likely that a user's write activity will impact another user's reading activity.

The various parameters of disk utilization *(Figure 17)* determine what portions of the hard disk are used for primary data and what portions are used for replicated and archived data.

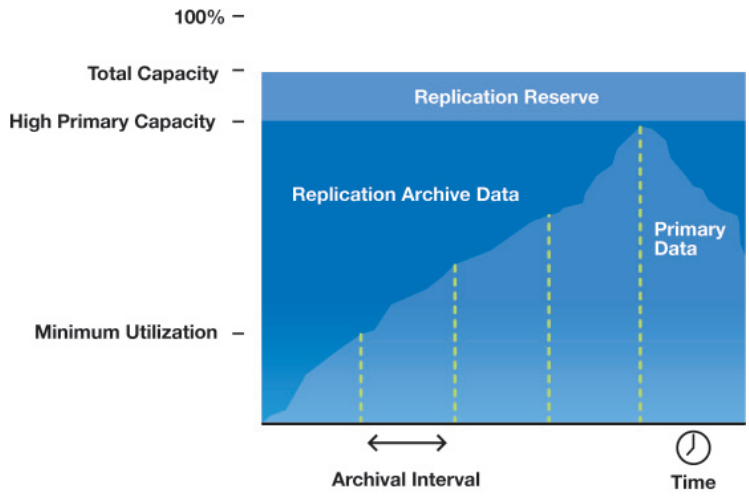

## Archiving

Archiving is the copying of primary data from the hard disk to the archival media. QStar supports three start archiving policies:

- **Periodic Archiving**: Here the systems administrator specifies a discrete time interval when the migration manager begins archiving. When periodic archiving is started, the migration manager begins processing delayed events until a stop archiving policy is satisfied. A copy of the data remains on the hard disk as archived data so it can be quickly accessed as needed without querying the slower archival media.

- **Demand Archiving**: The migration manager automatically begins archiving when the amount of primary data on the magnetic cache equals or exceeds the high primary capacity.

- **User Initiated Archiving**: The administrator initiates archiving on a selected integral volume. This method lets user archive critical files as soon as they are written.

Any archiving method, or all three, can occur in an integral volume. However, a well-configured integral volume uses periodic archiving as its primary method of controlling the migration cycle.

QStar also supports stop archiving policies. These rules determine when the archiving cycle should cease. The supported policies are:

- **All Delayed Events:** All primary data and all delayed events are archived. This policy ensures that all files as well as all high level changes such as changing ownership or permissions are processed on the archive media. This is the default stop archiving policy.

- **Low Primary Capacity**: The migration manager stops archiving once a certain amount of data has been archived. This amount is determined by the value of the tunable parameter, Low Primary Capacity. This parameter is useful in controlling the archiving to a defined data amount decreasing the length of the archiving cycle.

- **Agetime**: Only files written prior to the agetime are archived. This user definable value maintains recently written data in primary form. This limit is highly useful for customers who mainly access recently written files or those that wish to keep data on magnetic cache a certain number of days before archiving. These files will remain in the cache until they exceed the agetime value.

The system administrator sets the primary agetime and low primary capacity parameters based on integral volume usage characteristics.

## File Migration Modes

QStar's migration manager allows migration from the hard disk cache to the archive media to occur in two modes: data segment mode and full file mode.

In data segment mode, a file is viewed as a collection of data segments of discrete size. Each segment can be migrated individually or collectively with other segments. In full file mode, the migration manager always migrates the complete file as a single entity *(Figure 18)*.

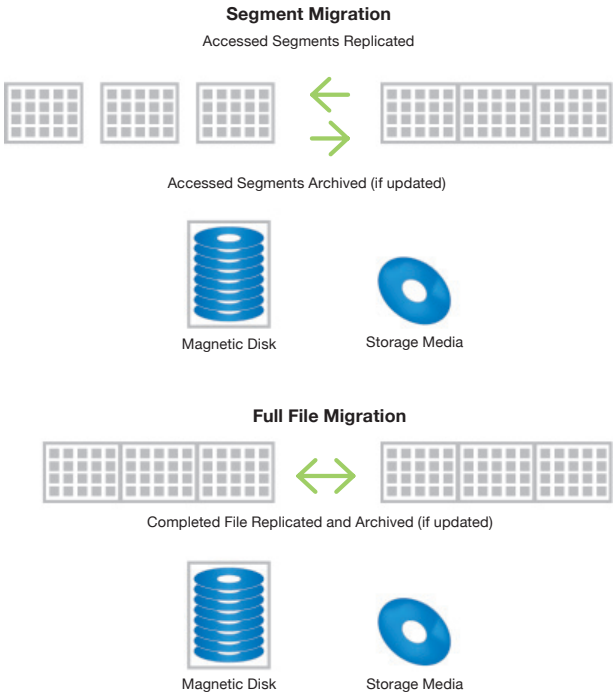


With large files, segment migration is more efficient than full file migration. For example, imaging applications typically store a multi-page document as a single file, often as large as 100 Megabytes or more. In segment migration mode, only the pages the customer are accessing or modifying are migrated. This will result in a substantial time savings: it's more efficient to migrate, modify, and archive two Megabytes (or less) than do the same with 100 Megabytes of data. Hard disk and archival media resources are used most efficiently, leaving more space on the hard disk for other data. Also, segment migration mode makes it possible to have files partially resident on magnetic and partially on the archival media.

For smaller files, full file migration mode is more efficient. The systems administrator defines a full file migration parameter, called the full-file limit, for each integral volume. For files larger than the full-file count, the migration manager performs segment mode migration. For files less than or equal to the full-file count, the migration manager performs full file mode migration. QStar's software can provide the flexibility of both alternatives.

### Per-File Migration Control

Both the systems administrator and privileged users can set a number of migration control functions and attributes for each file. Each of these functions override standard migration control behaviors specified for an integral volume. When used correctly, these functions can greatly increase individual user performance. Examples include:

- **Full File Migration**: The migration manager performs full file migration on the selected file. This attribute overrides the normal full-file migration limit for the integral volume where the file resides.

- **Archive Never**: The migration manager does not archive any primary data associated with this attribute. This is useful for temporary working files that have no permanent value.

- **Keep In Cache:** The migration manager always maintains a replicated copy of the file on the hard disk. This is useful for files for which immediate user access is required.

### Media Management

QStar's software can handle massive amounts of data with minimal user interaction. This is made possible by combining the latest in media technology with the QStar media formatting components: QStar Optical File System, QStar SDF File System, QStar CD/DVD File System Manager, QStar UDF File System and QStar QSCSI drivers.

## Media Types

The type and quantity of data stored by users determines the media they use to store it. QStar supports a wide variety of archival media, including CAS, optical platters, Blu-ray, DVD and CD disks and digital tape, as well as optical, DVD, CD or tape libraries.

Optical disks, or platters, are an external storage medium and come in many sizes, the most common being 51⁄4 inches such as Magneto Optical (MO) and Ultra Density Optical (UDO) or 120mm such as DVD and Blu-ray. Optical disks are portable and inexpensive to purchase and are most often used for archive applications. Optical is a random access media however the access time to data is slower than that of hard disk. Users may choose from a wide variety of single drives and libraries, varying from 600 Megabytes to multi-Terabytes of near-line storage capacity.

Tape is an external storage medium, usually both readable and writable, consisting of a loop of flexible celluloid-like material that can store data in the form of tiny magnetic fields that can be written, read and erased (if not a WORM media). The magnetic tape is housed in a plastic cartridge similar to that of an audio or video cassette. Because the tapes, which are recorded by a device called a tape drive, are portable and inexpensive to purchase, tape is often used for backing up or archiving data. A drawback of tape is that it can only be accessed by starting at the beginning and rolling through the tape until the desired data is located. QStar supports a number of tape formats including DLT, SDLT, LTO, AIT and SAIT.

The software supports write-once and rewritable media, as well as multi-function devices (capable of handling both rewritable and write-once media). Write-once media allows the ability

to write directly to a given location on disk only once; the user cannot edit or modify the data once written. Write-once is the most reliable form of data storage security, because the data cannot be altered once written. It is ideal for data that requires an audit trail, or invaluable data that has cost a great deal and whose loss or alteration the customer cannot risk. Rewritable media allows the ability to recycle space on the disk as if it were hard disk.

CAS technology is a hard drive based write-once solution where by special firmware is imbedded in the hard drive array that is WORM-like. Depending upon manufacturer, the CAS management system may also provide for single instant storage, encryption of data, self healing functionality, and internal replication of data. QStar supports several CAS hardware vendors.

The QStar media format managers likewise accommodate various media types and platter sizes, or a combination of them.

## Media Format

QStar's media format managers store information about the data on the surface of the medium, thus providing the template for the data written to the archival medium. All archival media are self-contained, with file directory information, data, and indexes on the same disk, platter or tape. The volume format optimizes the layout of data on an archival medium for maximum performance and transportability between the archive devices and file servers. All of the file components can reside on an archival storage device *(Figure 19)*.

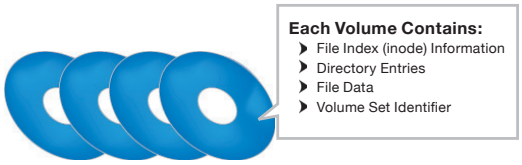


The benefits of media format include:

- **Transportability:** Media format lets the user easily move media from one host system to another, regardless of manufacturer. Thus, QStar's media format managers protect the enterprise investment in current hardware and allow access to critical data using other manufacturers' products.

- **Flexibility:** Media format describes the contents of a single archival medium allowing the medium to be used as part of a logical group, or media set. As part of a set, file and directory information is not restricted to a single platter or tape; it may span several platters or tapes, giving the system contiguous space for large files.

- **Disaster Recovery:** Making all archival media self-contained provides the customer with the means for disaster prevention and recovery. Any hard disk in the storage hierarchy can be completely rebuilt from the archival media, thus preventing catastrophic data loss.

# Library Volume Management

The storage media in a system can number up to thousands of tapes, DVD disks, CDs or optical platters in libraries. QStar's Volume Librarian module is responsible for managing all archival media known to the software, referred to as the media library. The Volume Librarian knows the location (physical and logical) and status of every archival medium, whether available on-line in a library or near-line in a storage facility. Its capabilities eliminate any errors that may occur in manual volume management, while increasing productivity and performance. With on-line media, QStar automatically manages all optical platter, DVD or CD disk, or tape allocation and movement without operator intervention. For requests that require off-line archival media access, the operator is prompted via the QStar logging feature as well as email notification to retrieve the optical platter, DVD or CD disk, or tape from its storage location and insert it into the managed device.

## Volume Catalog

To ensure the integrity of large numbers of archival media in the media library, QStar uses a volume catalog to index the location and status of each disk or tape *(Figure 20)*.

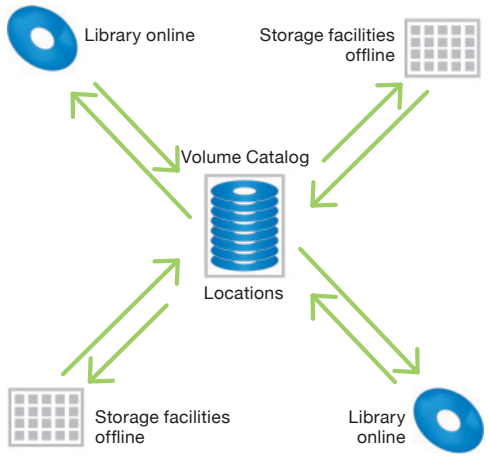


## Electronic Labeling

The key to each archival medium is a unique electronic ID, applied directly to each disk or tape. This label is used to record the surface, the physical medium, and any additional information the software needs. Using the electronic label, QStar can quickly locate any medium under its control.

## Location Tracking

As platters or tapes are moved from one location to another, the volume catalog is updated with the new information. For example, if information is residing on a platter outside a library, the volume librarian prompts the user to "retrieve the correct volume from the storage facility and place it in the library". Once inserted, the catalog changes the new location for volume to the library.

## Dynamic Allocation

As media sets become filled with data, it may be necessary to add additional media. QStar's software can dynamically allocate additional media to a set *(Figure 21)*. The user benefits, because there is no time spent resolving a request for more space. QStar simply appends additional media to the set, and the user's request is satisfied.

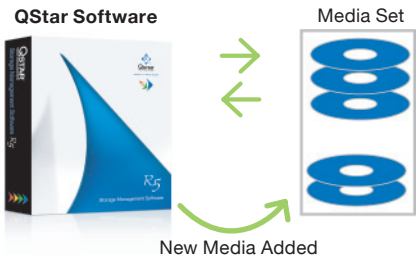


## Media Set Quotas

To prevent excessive usage of media, the software allows administrators to pre-define the set quota, or maximum number of archival media in a set, at set creation time. This feature gives the administrator additional control over the use of the media sets, particularly against runaway processes that consume more space than expected. Quotas can be implemented during the creation of a media set or after it is established.

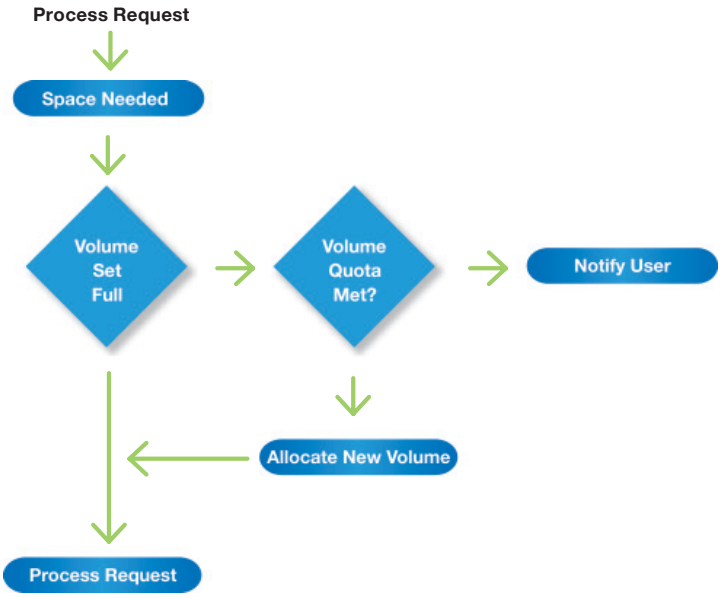


If a set requests additional media, the volume librarian first checks if the set quota has been reached and, if so, notifies the administrator via email *(Figure 22)*. If the quota is within limits, the volume librarian attempts to acquire an additional medium from a pool of blank platters, DVDs, CDs or tapes and append it to the media set. If a blank platter or tape is not available, the administrator is prompted to insert an additional one into the managed device. The administrator may also manually add media to a set as needed. The additional media is appended to the on-line media set without interrupting current processes using the set.

# QStar Data Director

The power and flexibility of QStar HSM can be enhanced by the addition of specialized disaster recovery software. Data Director adds the ability to do real-time mirroring, automatic failover and resynchronization of data.

QStar Data Director introduces data storage resilience to local and wide area networks in the form of data mirroring across multiple archive mediums. This software is a high performance and economical method of maintaining hot standby data storage while minimizing the vulnerability of data to loss. With mirroring, the software provides an efficient means of balancing the read-access loading of the storage sub-system. QStar Data Director offers the following features:

- **Data Mirroring**

  QStar Data Director provides an innovative way of solving the problem of backing up and restoring massive data storage repositories. By utilizing QStar's unique event-based mirroring technology, it is no longer necessary to schedule hours of downtime for backups.

  The software mirrors each file system event, which occurs to one or more distinct storage devices in real time. This means that the customer has exact duplicates of the data at any instance in time without the utilization of manpower and time-expensive backup processes. Even more important, the mirroring takes place automatically and without the need for user intervention.

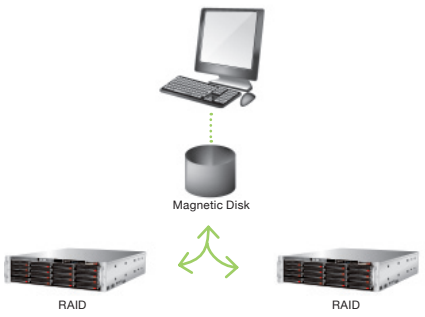


- *RAID - RAID*


- *RAID - Library*


- *Media Set - Media Set*

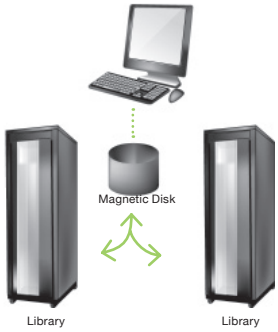


- *Library - Library*
- *Media Set - Media Set*

Data Director is completely independent of the archive device. By using combinations of hard disk space, RAID, optical, DVD, and tape libraries, different configurations can be obtained depending on the customer's needs *(Figure 23)*.

### WAN and LAN Support

Mirrored sets are not limited to the local server. Other systems on the local LAN or across a broad WAN regardless of their platform can operate as a secondary mirror site. This allows for duplicate sets to be secured by separating the libraries to other rooms or buildings, or adding another layer of security.

### Hot Standby

The hot standby feature of QStar Data Director allows the user, in the event of a device failure, to automatically switch over accessing data on one of the mirror devices without shutting down or reconfiguring the system. Once the failed device has been repaired, it can be brought back on-line as part of the mirror set with a simple software command.

### Improved Disaster Recovery

By making all the mirrored devices self-contained, Data Director provides the means for disaster prevention and recovery. Any device in the storage hierarchy can be completely rebuilt from the mirror device, thus preventing catastrophic data loss.

### Library Load Balancing

Because the data is mirrored, this allows the read-access loading to be balanced between archival storage devices, resulting in faster data access and system response during peak access times.

### File system Resynchronization

Also known as "Hot Sync", this feature will resynchronize the mirrored file systems should one be brought off-line because of hardware or other failures. This feature can also be used to take an established QStar HSM set and later add QStar Data Director to the software package to mirror this set.

## QStar Simulated Disk and Simulated Tape Jukebox

The advantage of QStar Simulated Disk and Simulated Tape Jukebox is the ability to configure a disk based system to look like a storage library on the network, then either copy that data to a physical storage library or with QStar Data Director, mirror between the simulated library and a physical storage library. By utilizing this technology, users are able to overcome some inherent operating system limitations and gain the ability With mirroring now available between disk based simulated devices and actual removable media devices, and by taking the security of a transaction logging cache, a system independent disk format with added event based mirroring, QStar can offer hybrid archive systems that have hard disk performance with mirrored data protection and low-cost redundant tape or optical backup functionality.

QStar offers two types of Simulated Jukeboxes within a Hard Drive configuration, **Simulated Disk Jukebox (SDJ)** and **Simulated Tapes Jukebox (STJ). SDJ/STJ** can group together multiple native file systems as a single large storage repository. This allows multiple devices to be managed as a single device or volume, overcoming operating system limitations on the size of a single logical volume. Also with SDJ/STJ there is no limitation to the number of files that can be stored in the QStar Standard File Format volume nor is there any file size limitation.

SDJ/STJ allows administrators to take advantage of features from other QStar HSM and Data Director previously only available for removable media technologies, such as, make copies of data from simulated volumes to removable media volumes automatically with QStar's Copydisk function. This allows users to make backup or distribution copies of data when needed. SDJ/STJ also allows users to create disk images in a standard exportable format such as ISO or UDF, which can later be burned onto optical disk or written to tape. In addition users can create a data set that can be mirrored at a later point in time. Volumes can be configured to logical views that make sense for a specific application. Users can either select the block-size and number of blocks for each simulated media or they can use a predefined value for the medium size.

SDJ/STJ can manage physical disk space available on a hard drive (Local Disk, RAID, DAS, NAS) or hard drive based write-once appliance (CAS) as if it were media loaded into an actual library as an Integral Volume set. The hard drive or write-once appliance can be partitioned to simulate virtual shelves allowing full use of library commands and behaviors. SDJ/STJ allows users to further protect archive data on a CAS device by writing a second copy either to another CAS device or to an optical or tape storage library. Should the CAS array fail for any reason, data is still accessible from the secondary device. This extends the level of protection beyond a single technology for a more robust archive disaster prevention plan.

## QStar API

As integrators, VARs and end-users plan and implement storage solutions, either locally or for a customer, the need to provide a complete package is becoming common place.
APIs, Application Program Interface, facilitate a path to integrate front-end applications with QStar Software, whether trapping error messages or issuing low-level library calls, QStar API provides the catalyst between the user application and QStar storage management software.

- **Low-Level Device Inquires**
  With QStar API, companies can make low-level SCSI queries to the library controller and/or the drives. All standard SCSI-2 commands can be passed through the QStar QSCSI driver directly to the device.

- **Library Control**
  Optical platters, CDs, DVDs and tapes can be manipulated in their library by passing through the QStar library management module. This allows for greater control of the managed devices as well as the ability to, in some cases, right out native file systems to optical platters.

- **Request and Error Trapping**
  An application, utilizing QStar API, can trap error codes and requests for human intervention. Once obtained, the application can then decide what is the best course of action

- **Seamless Integration**
  By utilizing the API, the QStar Software is invisible to the customer. All the commands can be linked within an application that looks as if it is running the library.

- **Expert Help**
  API customers get the best help offered, right from the developers of the QStar Software. Let QStar help with the integration.

# Advanced Features of the QStar Software

## Compaction

When working with optical, DVD, or tape media and the QStar filesystem SDF, the need may arise to delete many old documents. This leaves behind “dead space”, deleted files and old file versions, where these documents once resided. QStar provides a way to reclaim that space and provide assurance that the files are removed, by using the compaction utility. This specialized utility ensures that the media set uses the available space to its fullest extent.

- **Optimized Media Usage**: By removing “dead space” by moving the valid data to the end of the set, media is used to it’s fullest. Recovered media can then be erased, if applicable, and either re-added to the set or allocated to another *(Figure 25)*.

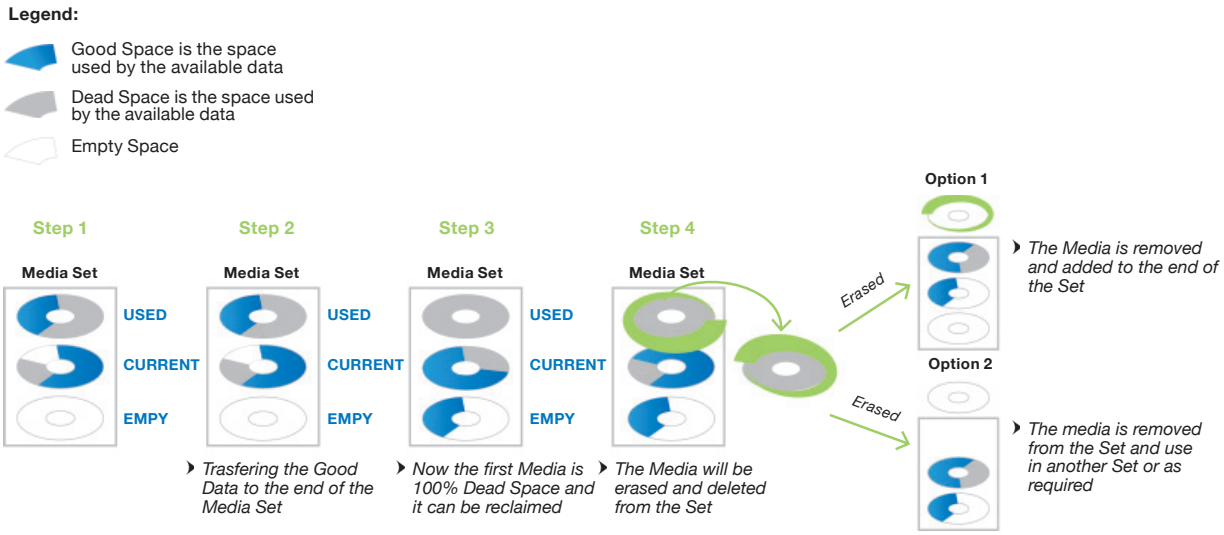


- **Absolute Deletes**: Customers that must comply with government policies or internal company procedures that mandate absolute deletion of files after a given amount of time, can rest assured that files are removed permanently from the storage medium. Removing the “dead space” during compaction and then erasing that piece of media, if applicable, guarantees that these standards are met.

## Digital Signature

Adding the ability to provide digital signatures to stored data and manage the digital signature process provides another level of information security within the QStar product line to protect customer’s data. Digital signatures can be used with any kind of document, whether it has been encrypted or not. The digital certificate contains the digital signature of the certificate-issuing authority so that anyone can verify that the certificate is real, and thus the document is authentic. This is a key component for a number of highly regulated market segments such as Drug Development, Medical and Life Sciences. The signature is there to ensure that any changes made to the data have been signed and cannot go undetected.

QStar Technologies, Inc. has teamed up with IPM to create a file management digital signature solution that complies with the data storage and data management regulations in the medical market. HIPAA, the regulation committee in the US for medical records, defines a digital signature as an electronic signature that cannot be forged. It is a computed digest of the text that is encrypted and sent with the text message. The recipient decrypts the signature and re-computes the digest from the received text. If the digests match, the message is authenticated and proved intact from the sender.

A digital signature ensures that the document originated with the person signing it and that it was not tampered with after the signature was applied. However, this is not enough to ensure that the sender is not an imposter. To verify that the message was in fact sent by the person claiming to have send it requires a digital certificate (digital ID) which is issued by a certification authority.

The system uses a one-way hash function to compute a small digest of the text message. Using a private key, the sender encrypts the digest, turning it into a digital signature. The signature and the message are then encrypted using the recipient's public key and transmitted. The recipient uses a private key to decrypt the text and derive the still-encrypted signature. Using a public key, the receiver then decrypts the signature back into the sender's digest and then re-computes a new digest from the text message. If the digests match, the message is proved to be authentic.

## Privileged Users

In some instances it is necessary to provide certain trusted users with the ability to run QStar commands as needed. These could include the ability to mount an integral volume, add media to a set, create a new set, gather proprietary statistical information, etc. QStar provides a super user owned file that can be easily modified to give individual users these abilities. With the simple addition of their names, the added users will have the ability to run all QStar commands as if they were super user.

## Automatic Copydisk/Copytape

The QStar software allows administrators to make exact duplicates of the media in an integral volume with an included duplication utility. This utility can be ran at the administrator's discretion or can be set to automatically duplicate the media on a set-by-set basis. One option is to have the copydisk/copytape utility begin making the duplicate as soon as the media is full. Another is to run on an incremental schedule, copying the data each day. This option is useful when the application may take more than just a day to write a full media. Each night the current write surface can just write the newly created data for that day. This is referred to as "delta" copying. Once the destination media copying is completed, the duplicate can be removed from the storage device and placed at a secure off-site location. If anything catastrophic should happen to the original set or library, the duplicate media could be loaded into another library on any platform supported by QStar to a fully accessible state.

## Email Notification of Pending Requests and Errors

The QStar Software can be configured to notify the system administrator when problems arise or a request is pending. As errors are sent to the QStar log file, a copy can be forwarded to the system administrator so that immediate action can take place. Pending jobs would include requests for off-line platters and adding media to a library when no erased media is available.

## Grace Period

Grace period is a feature that can be used on a set-by-set basis. This feature effectively turns a set, even at the magnetic disk cache level, into a WORM volume. Grace period is defined by an amount of time. This time dictates how long the cache will wait after a file has been written before it makes the file non-removable and non-alterable. This feature overrides permissions, even of the administrator, and ensures that the files can not be altered. This is an effective feature in compliance rich industries. Even if the company is not mandated by a regulatory compliance mandate, it is a useful to tool to protect the company from any internal foul play.

## Automounter

As used with the QStar CD/DVD Axxess or Master software, the QStar automounter can be used to place many single or multi media integral volume sets under one mount point.
Each surface will look as a subdirectory of the mount point. To minimize load on the system when a large number of single media sets are being managed by the automounter, for example, the software will mount and dismount the requested media as needed. This feature ensures that only sets with data that has been requested by a user are mounted and occupying the library drives. By limiting the amount of mounted media, valuable system resources are not consumed by inactive volumes.

## Fibre Channel Support – SAN Solution

With the growing need for SAN solutions today, QStar has added Fibre Channel support to all of the supported platforms. When utilizing the benefits of Fibre Channel, the server can remain in its current location while the library is stored in a separate room designed specifically for storage devices. This ensures that the media sets are secure in case of a major disaster.

Another use of the SAN solution is to free up the application server from its duties as the file server. Placing heavy strain on this server, by servicing both needs, many times requests are not serviced as quickly as they would be if the server were only providing applications. Having a separate file server allows for several application servers to all access and benefit from the library storage device. Having only the responsibility of provided I/O, this server can facilitate faster request times to all machines accessing it.

## QStar Software and Open Systems

Generally, hierarchical storage management requires that file systems know about the storage hierarchy and be able to account for the location of data within the hierarchy. UNIX, Linux or Windows do not provide this type of information. Storage management software vendors usually solve the problem in one of two ways:

- Supplying a proprietary file system with hierarchical awareness. However, this file system may not accommodate the existing file system data, the native operating system, or the native system's enhancements.

- Extending the native operating system and file system by adding a layer of migration support code to the native system kernel.

QStar Software is completely independent of application, kernel, and file systems, and does not affect them in any way. The software operates by inserting a library pseudo-driver one level under the file system *(Figure 26)*. A QStar managed storage system is, consequently, as open as the system's operating system.

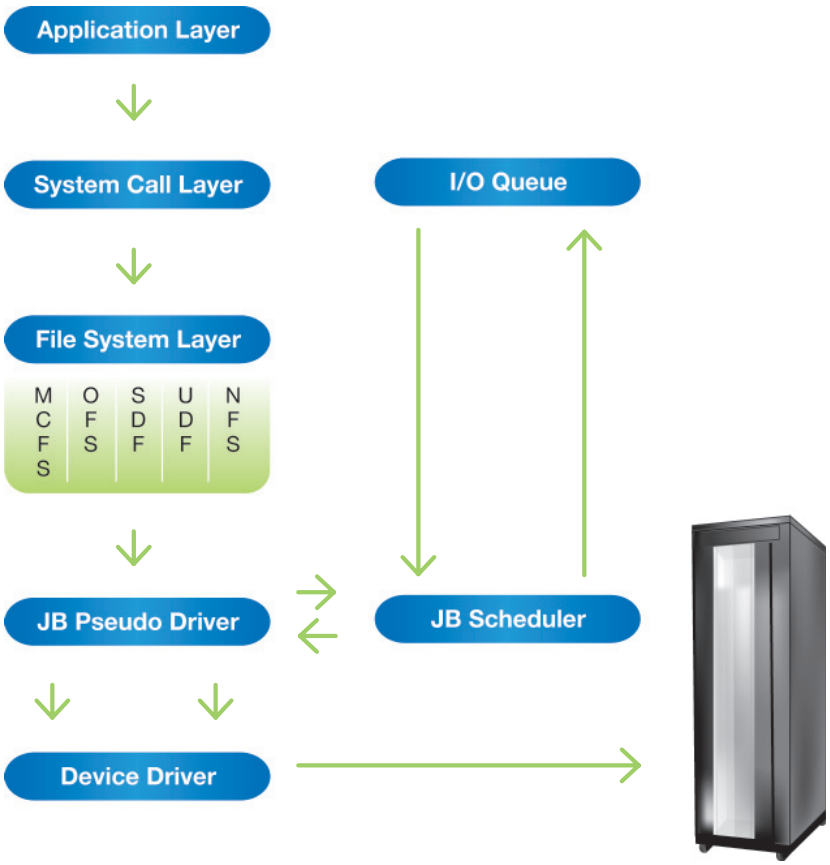

Designed to conform to, and go a step further than, current industry standards for open systems, QStar Software is:

- **Transpoble**
  QStar's software is media and application independent. Supported on the top UNIX operating systems (including Sun Solaris, Compaq Tru64, HP/UX, IBM AIX, LINUX and SGI IRIX) as well as Windows. QStar Software can be installed on any industry-standard client/server network.

- **Networkable**
  On each supported platform, QStar supports each of the standard networking protocols, including TCP/IP, NFS, Novell, NetBEUI, Netware, AppleTalk and LanManager.

- **Scalable**
  System configurations can range in size from one Gigabyte personal archive systems to multi-Terabyte enterprise storage configurations for servers, workstations, and desktop platforms.

- **Flexible**
  Media transportability, automatic device failure recovery, and automatic cache recovery allow for rapid deployment and redeployment of mission-critical systems.

QStar's software line is the storage solution for the client/server environment. With automated, hierarchical storage management, optical, DVD, CD and tape libraries, and file recovery; client/server networks can now enjoy the same level of data reliability, availability, security, and integrity as provided in the mainframe environment.

## Configuring for Individual Applications

QStar's software is shipped with a set of defaults that, from QStar's experience, are appropriate for most applications. However, today's applications vary widely in how they use mass storage. For this reason, QStar provides parameters for tuning performance on selected files or for an integral volume. The following sections describe how the administrator can set different integral volume parameters to achieve maximum performance in two selected scenarios.

### Interactive Model

In this model, we assume that an integral volume is configured to run as a network storage disk for general purpose user files. The amount of file read activity is approximately equal to the amount of file write activity. Also, the majority of the data that the users work with daily fits on the hard disk *(Figure 27)*.

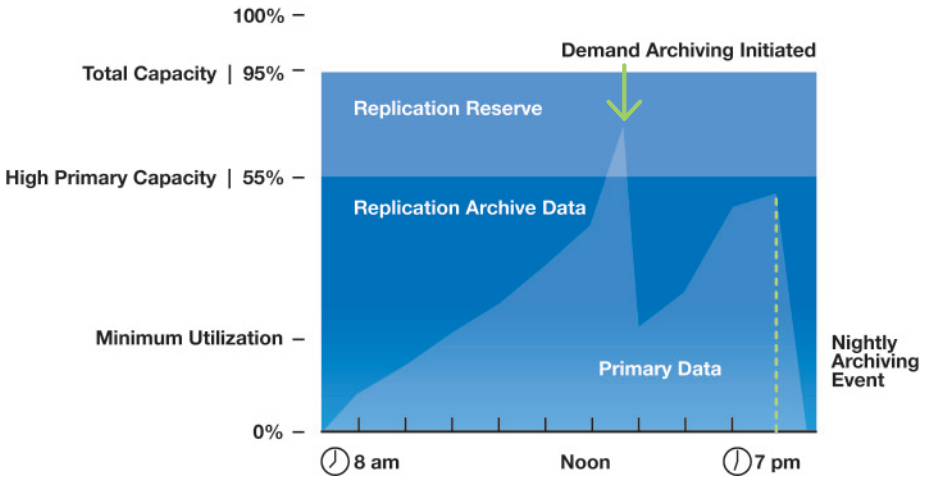


High primary capacity has been set to 55% of optimal capacity. This decreases the possibility that during peak write times, i.e. just before lunch when users save their files, read access users do not suffer a performance penalty due to over-utilization of the hard disk by primary data. Additionally, the archive interval is set to 24 hours, with initiation time after close of business. This means all primary data created during the workday is archived to permanent storage after hours. Thus, archiving does not affect user performance during the working hours.

## Session-Oriented Model

This session-oriented model assumes that an integral volume is used by different users over discrete time intervals. In this example, a high-end CAD/CAM workstation is shared by several engineers. In this model, a user sits at a workstation and begins a design session to access, create, and update a number of drawings.

The systems administrator sets the high primary capacity to include most of the hard disk *(Figure 28)*. Thus, performance is not affected by archiving events. At the end of the session, the engineer initiates an “archive now” function so that the primary data is archived to permanent storage. The workstation is now ready for uninterrupted use by the next user.

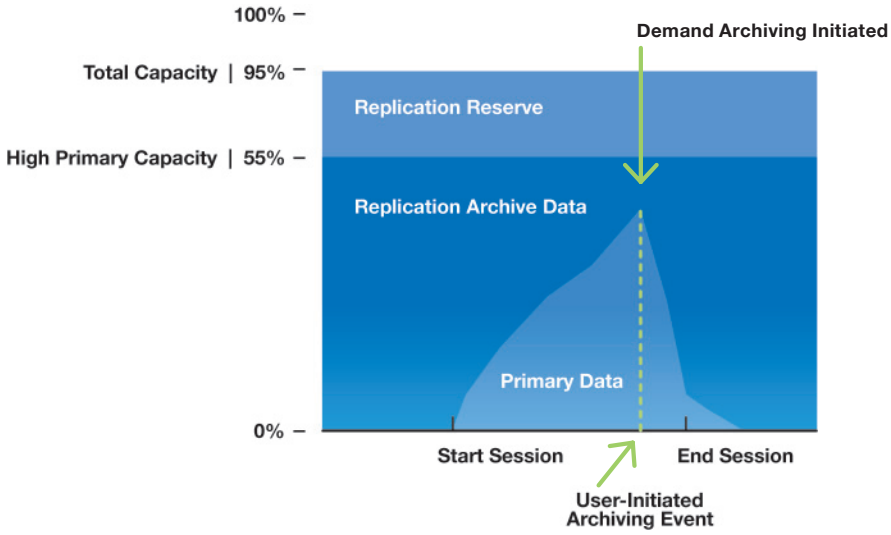


SECURING THE FUTURE OF YOUR DATA®

info@qstar.com
info@qstar.it

**www.qstar.com**