InveStore v4
Product Overview

Pegasus Disk Technologies, Inc.

Contents

INVESTORE V4 3

MMC GRAPHICAL USER INTERFACE 3

PRODUCT FEATURES 4

INVESTORE ENTERPRISE 5

ACCESS TO THE STORAGE RESOURCE 7

DEVELOPERS INTERFACE 8

VIM OPTIONS 9

ERROR RECOVERY AND MESSAGING 11

INTERNAL FILE SYSTEM FEATURES & FUNCTIONALITY 12

VOLUME SPANNING 14

STANDARD FEATURES 17

FOR MORE PRODUCT INFORMATION 18

InveStore v4

InveStore Storage Management Software v4 has been designed to meet the market demands for seamless integration within the Windows operating system that incorporates a user interface architecture based on Microsoft's MMC and including the use of COM/DCOM objects.

- InveStore v4 provides Explorer like look and feel, while maintaining compatibility and support for the historical media and storage devices that are used for storage management.

- InveStore v4 provides the functionality to expand single-surface limitations, spanning multiple volumes for increased data archival capability.

- InveStore v4 provides greater functionality and Enterprise level control of the data storage system and its individual components.

- InveStore v4 provides an easy-to-implement software solution for connecting your storage device (MO, WORM, CD-ROM, DVD-ROM, DVD-RAM, and UDO hardware) under Windows 2000/XP/2003. It provides superior capability for the developer using our product to integrate Data management and storage functions, directly into their applications.

- InveStore provides access to re-moveable storage devices such as standalone drives, libraries, auto-changers and towers and presents extended storage as a single, hierarchical file system. Through a single drive letter, applications can utilize terabytes of storage as easily as using a server hard drive. Data and files are accessed using the normal operating system commands.

- Application software independent. InveStore handles all commands to the storage device drives and robotics, allowing the application to access the storage media transparently in the operating system.

MMC Graphical User Interface

Using Microsoft Management Console (MMC) modules, the InveStore v4 interface allows greater functionality within the console and easy "snap-in" of your MMC compliant modules. Using the MMC as the basis for our management console (GUI) eliminates the need for a proprietary console. Users are presented with a standard look and feel for system management. Manage multiple InveStore storage servers, import and export media, rename platters, reserve disk space, adjust performance parameters, view drive & volume details, format media, create volume sets and more.

The need for extensive training is eliminated since the user is operating within the context of a familiar Windows framework. Because MMC has been released as a standard part of Microsoft Windows operating system, it provides a common console framework for server and network management applications.

Pegasus Disk Technologies, Inc. 18 Crow Canyon Court, San Ramon, CA 94583
Telephone: (925) 314-1800 • Fax: (925) 314-1801 • Http://www.pegasus-afs.com • sales@pegasus-afs.com

InveStore
Storage Management Software

Product Features

Feature	***Benefit***
Volume Spanning	Allows the user to specify a number of media to be included in a volume set(s). This volume set is addressed as one target for write and read operations.
MMC Console interface	Updated interface based on Microsoft Management Console (MMC). Right-Click functionality.
COM/DCOM API	Programming Interface architecture based on Microsoft's MMC and including the use of COM objects.
Runs as an Windows Service	InveStore automatically re-starts and re-establishes shares when server is started.
Fully independent application software	No additional protocols or third-party software products needed.
Supports any combination of write once, Re-writable, CD or DVD-ROM, DVD-RAM, UDO media under the same driver letter	Multiple SCSI storage devices appear as a single resource, accessible through a single drive letter. Each storage volume is represented as a first level subdirectory.
Supports multiple file formats	ISO 9660 (CD-ROM), High Sierra and the Pegasus high performance file systems.
Full 32 bit kernel	Full multitasking capabilities provides access for Up to 1000 users with only one server.
Extensive Volume Management	Automatically tracks files on-line (in the storage drive), near-line (in the jukebox), or even off-line (at the warehouse).
Multiprocessing on the server	Add, remove, or modify storage drive parameters and log system events and errors without disrupting the file server.

Pegasus Disk Technologies, Inc. 18 Crow Canyon Court, San Ramon, CA 94583
Telephone: (925) 314-1800 • Fax: (925) 314-1801 • Http://www.pegasus-afs.com • sales@pegasus-afs.com

InveStore Enterprise

The InveStore v4 console presents to the storage system Administrator four nodes or levels for the control of and access of the storage system. The primary node is that of the ***InveStore Enterprise***. From the ***InveStore Enterprise*** node is accessed the various levels leading to the physical components of your storage system.

The InveStore console presents the different levels of the storage system using an ***Explorer*** like interface. This interface displays all of the connected ***Servers***, ***Libraries***, ***Volumes*** and ***Volume Sets*** under the entire Storage Enterprise on your LAN.

Each node has a specific ***Action*** menu and ***Toolbar*** with tool Icons that provide the various functions related to the current level. This allows the InveStore v4 Console to provide greater detail, for ease of control and administration.

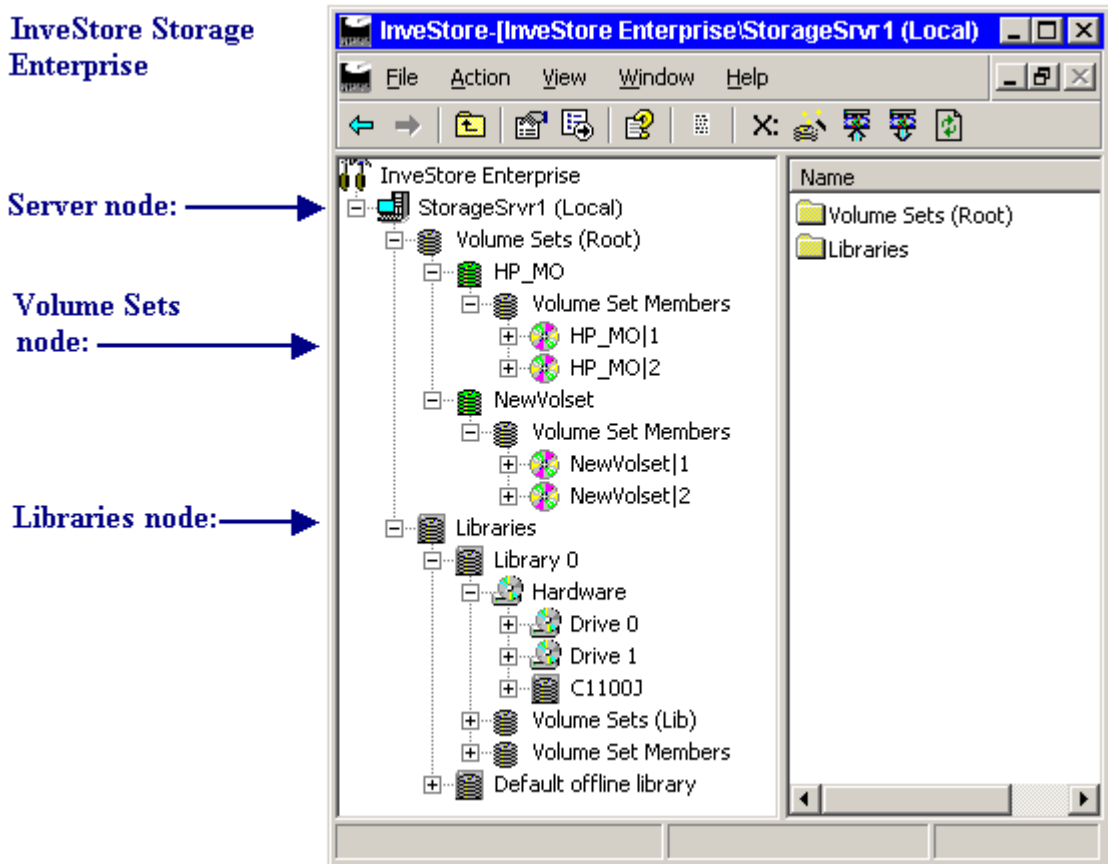

InveStore allows the operator to double-click to expand the various views in the MMC console to drag and drop where appropriate, and to resize windows. The menu items are context (node) sensitive and individual menu items will display information or functions specific to the context of whatever the current view may be.

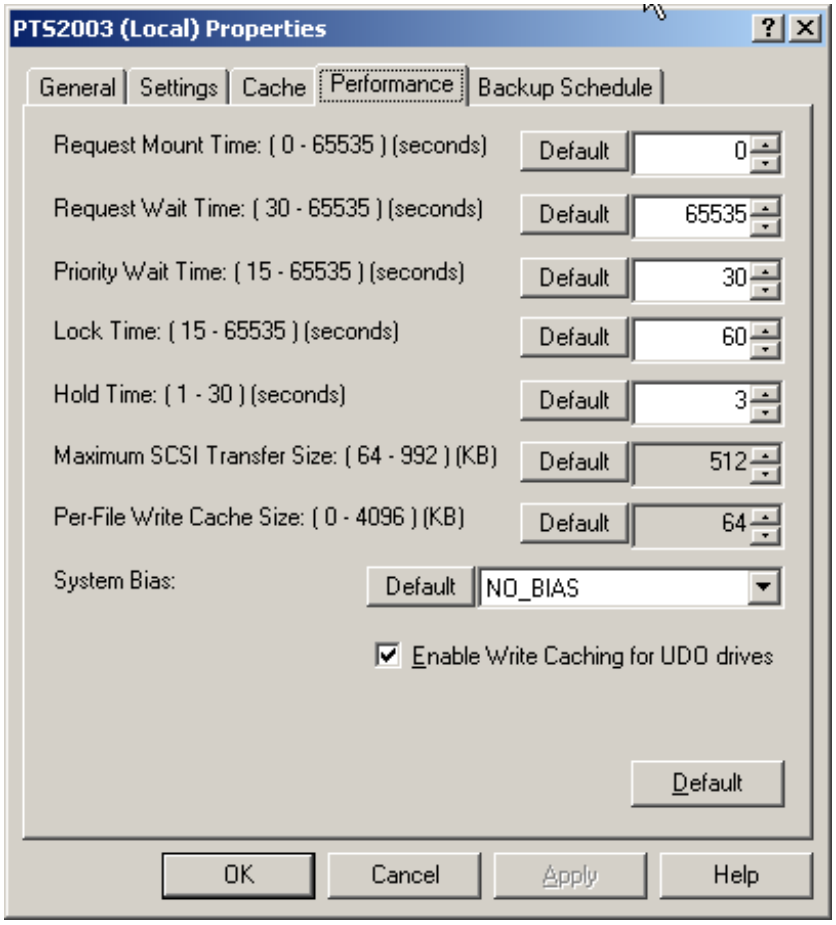

Access to the Storage Resource

Access to all storage library resources is made available through a single drive letter interface or via API integration. InveStore automatically tracks on-line data (in the storage drive), near-line data (in the storage library), or even off-line data (in a secondary storage location).

InveStore treats each storage volume set as a first level subdirectory. The following is a typical directory view of an InveStore controlled storage device as viewed from ***Windows Explorer***.

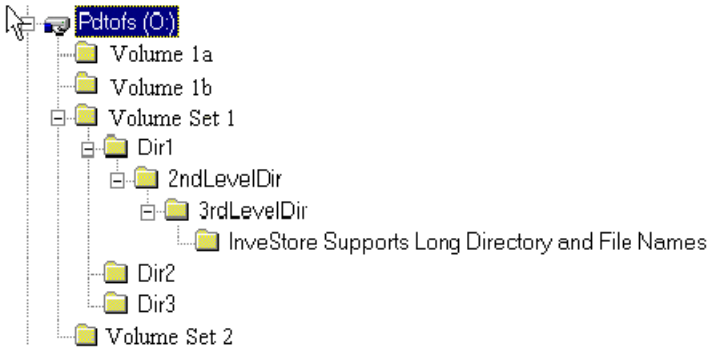

InveStore provides support for virtually all of the storage hardware from the industry's leading manufacturers. It handles combinations of drives and storage libraries, allowing daisy chaining up to seven devices on a single SCSI bus. Multiple SCSI adapters are simultaneously supported as well. InveStore provides an easy upgrade path from single standalone storage drives to high-end systems, as storage needs increase. Data is fully transportable from device to device and operating system to operating system.

Developers Interface

COM

Component Object Model (COM) is a Microsoft architecture that provides the framework for creating and using component-based applications. The InveStore console consists of client-side ActiveX COM modules, which interface with the server-side COM objects.

Using ActiveX COM modules allows the InveStore v4 console to incorporate and provide access to the widest choice in services, tools, languages and applications to developers. This interface makes it more accessible to VB programmers and lends itself better for creating a unique interface for a specific application.

Clients can quickly develop custom controls within their own application to control the various features and functions of the storage management software. From mounting and un-mounting cartridges, formatting new volume sets, backing up volume sets to error management can all be handled remotely with the Pegasus storage management COM API. The range of applications that can now be developed has been greatly expanded.

COM Internet Services (CIS) also introduces support for a new Distributed COM (DCOM) transport protocol known as Tunneling Transmission Control Protocol (TCP) that allows DCOM to operate over TCP port 80. This allows a client and a server to communicate in the presence of most proxy servers and firewalls, thereby enabling a new class of COM-based Internet scenarios.

DCOM

The Distributed Component Object Model (DCOM) is a protocol that enables software components to communicate directly over a network in a reliable, secure, and efficient manner. Previously called "Network OLE," DCOM is designed for use across multiple network transports, including Internet protocols such as HTTP. DCOM is based on the Open Software Foundation's DCE-RPC spec and will work with both Java applets and ActiveX components through its use of the Component Object Model (COM).

DCOM allows the InveStore API and the InveStore storage management server to be available from anywhere on the Internet. Standard web browsers can incorporate ActiveX components that integrate the management features of Pegasus directly into the web browser. The potential for such capabilities are unlimited.

By providing our management console as a plug-in to MMC we have eliminated the need for a proprietary management console. Our management features now coexist side by side with other consoles such as SQL Server, Information Server, etc. In addition our client can place their own management features within the MMC, providing the end user with a single, simplified management console for all applications.

Network Access

Developers are not limited to working on the server and can access data from any workstation at a remote location. COM works with multiple network transports such as TPI/IP, IPX, SPX and HTTP.

Security

COM provides flexible security options such as NT4 Security, SSL/Certificates, Kerberos and IPSEC.

VIM OPTIONS

Shown here is an example of the Pegasus VIM call options available under the library classification:

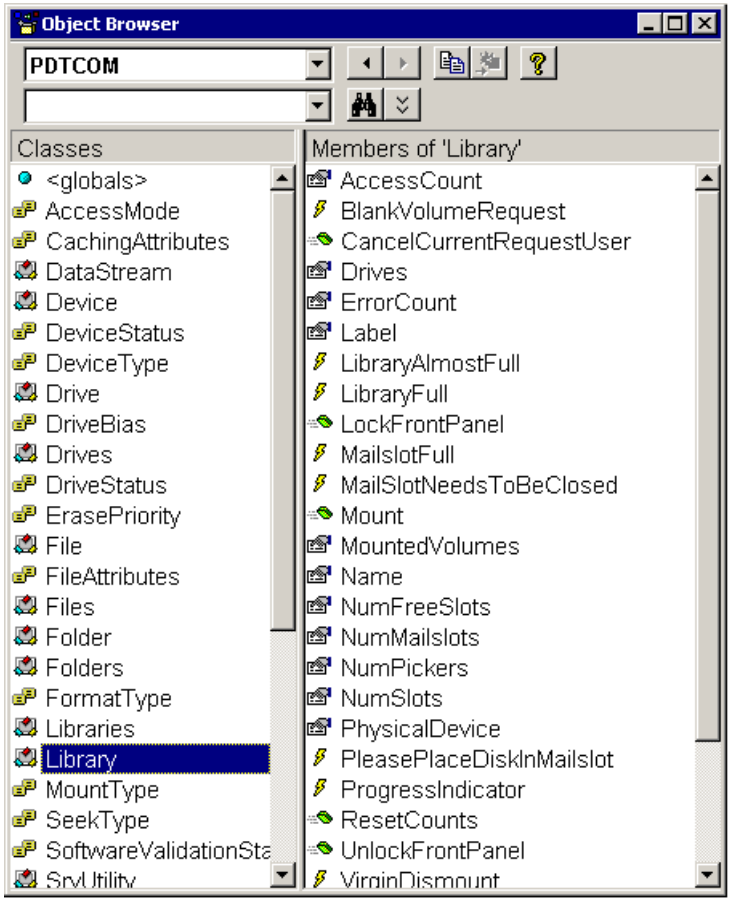

Shown here is an example of the Pegasus VIM call options available under the device classification:

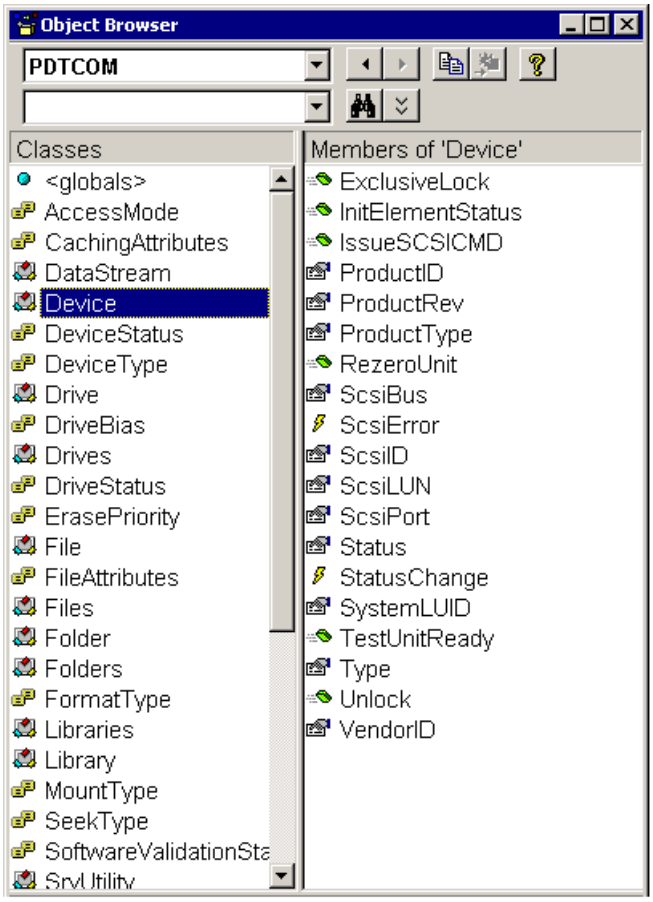

Shown here is an example of the Pegasus VIM call options available under the VIM classification:

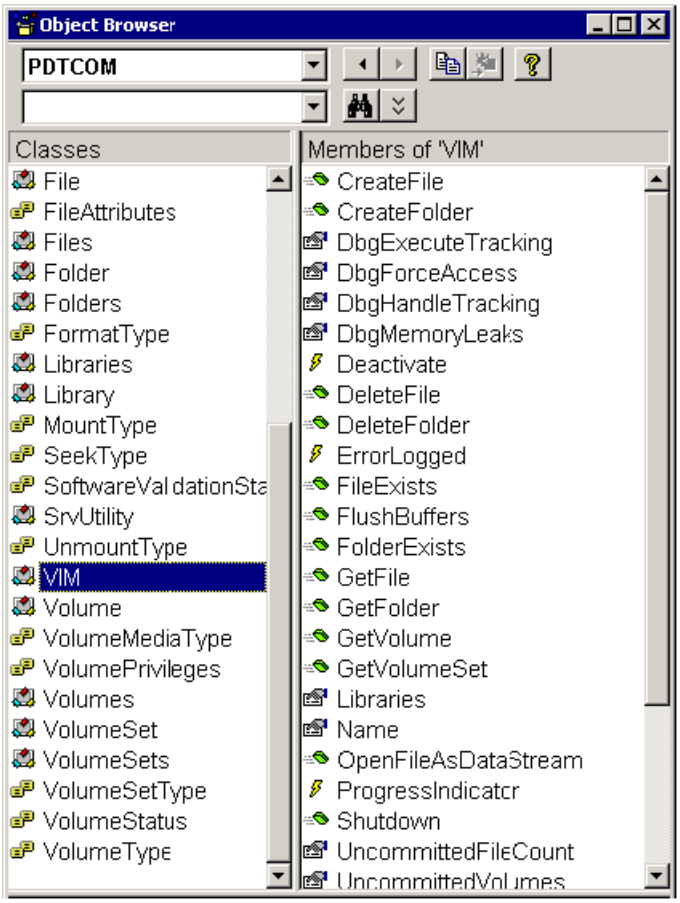

Error Recovery and Messaging

Error Recovery/Processing

Error processing and recovery is a major part of any good storage management system. Pegasus has spent an enormous amount of time enabling the InveStore product to diagnose and recover from any type of error. Error processing is built into all levels of the system. The design goal has been, and continues to be, that all errors can be easily diagnosed and/or corrected, either automatically or by viewing the log files that are stored to the storage server hard drive.

The error processing system alerts the user to error conditions. All error messages are recorded to the Windows event log as well as the InveStore event log file. The InveStore device drivers contain extensive error processing code for trapping SCSI or operating system related failures. Each device has its own custom error processor that will correct whatever errors it can or pass the failure back up to the console when it cannot. These types of error messages are trapped in detail within the InveStore log event file.

Messaging

Log messages and Administrator action items are displayed with pop-up message windows. Messaging is based on COM events. These "events" have a high degree of configurability. The following are some of the events that are covered by this feature:

- Startup Successful
- Start Failure
- Mount Volume Set (please insert disk, please insert next disk)
- Mount Volume (please insert disk)
- Unmount Volume Set (please remove disk)
- Unmount volume (please remove disk)
- Prescheduled Backup launches
- Backup Completions
- Format Volume Set (progress pop-ups)
- Format (please insert disk + progress indicator)
- Offline Volume Mount Request
- Please Shut Mail-Slot Door
- Severe Error Warnings
- Hardware Errors
- Activation

Internal File System Features & Functionality

Integrated 32 Bit Kernel

For a storage management system to be truly useful it must be tightly integrated into the operating system. Pegasus has integrated the InveStore product into Microsoft Windows environment using the File System Driver (FSD) approach. This is Microsoft's only approved way of adding foreign file systems to Windows. The FSD provides Pegasus with the maximum transparency and integration. This allows InveStore to be shared using the Native Windows server software.

InveStore also incorporates the 32 bit multi-tasking functionality of Windows for improved performance and system throughput. While the InveStore software has been designed to be operating system independent, the low level modules which provide the required abstract operating systems calls, such as memory allocation, events, semaphores, etc., are fully customized for the particular operating system and take advantage of what that OS has to offer.

Access and Security

The InveStore console design provides remote access and administration from one to many InveStore storage management system(s). In doing so the InveStore console eliminates the current one to one console to storage server limitation. A single console can connect to and administer any number of InveStore storage servers that exist on the network.

Redundant File System for WORM

The Pegasus File System for Archive WORM Media has been designed with reliable long-term archival storage as its primary goal. The on-disk directory structure uses twin directory entries that are doubly linked. If one entry is damaged the system can use the links to traverse around the damaged entry and retrieve the duplicate. In addition, the duplicate entries are stored in different areas of the disk, greatly reducing the chance that both entries will be damaged, making our file system fault tolerant.

Multiple File System Support

The InveStore software has been designed to be independent of any file system, including the Pegasus WORM file system. A major design goal from the beginning was to enable support for multiple formats with a minimum amount of work and accessibility through a single common interface. Pegasus currently supports two versions of WORM file system, ISO 9660, OSTA UDF through Section 3 and Joliet.

The ability to support multiple formats has also been beneficial to OEM clients of Pegasus who have legacy media written using their own proprietary format. This allows these OEM's to concentrate on developing value added features to their main product instead of wasting resources with legacy optical support.

The InveStore Storage Management Software is made up of several distinct components that handle the functions necessary to read, write and provide transparency to removable storage libraries under the Windows operating systems. One component, the Physical Format Manager (PFM), is responsible for understanding and manipulating the physical file system structure recorded on disk. This abstraction layer allows a single Logical File System (LFS) layer to manage multiple physical media formats. The LFS implements all of the high-level file functionality such as open file, create file and make directory, while the PFM layer deals with more primitive functions such as formatting a disk, allocating and freeing space within a file system partition, and creating and updating on-disk meta-data structures.

This design greatly simplifies what the PFM must do. File system semantics common to all file systems such as Create, Open, Close, Read, Write and Change File Position, are handled by the LFS layer and the PFM concentrates on simplistic meta-data handling routines. This approach allows additional file system PFM's to be added to InveStore with a minimum amount of coding. Each PFM that is added can take advantage of InveStore features including Hierarchical Cache Management (HCM), transparent operating system access and a huge array of storage device drivers for CD, DVD, Optical or Tape based media.

Volume Spanning

Spanning

One of the significant features developed in the InveStore v4 Enterprise Storage Management Software is that of Volume Spanning. This specifically applies to MO and WORM storage devices, as well as tape in a future release. Volume Spanning is the process of combining individual disk cartridges within a storage device and making them appear to be a single contiguous volume. InveStore v4 employs an algorithm that allows files and directories to span disk surfaces. The volume set is typically a group of disk cartridges that are grouped together as a single logical volume.

A volume set is composed of at least one disk cartridge, with each side of the disk incorporated into the volume set. The first surface of the first cartridge is typically referred to as the root volume, and all other surfaces are referred to as member volumes. The InveStore v4 data allocation algorithm allocates space from the first member volume until it is completely full, and then continues writing to other members of that set.

Below is a screen showing properties on a volume set:

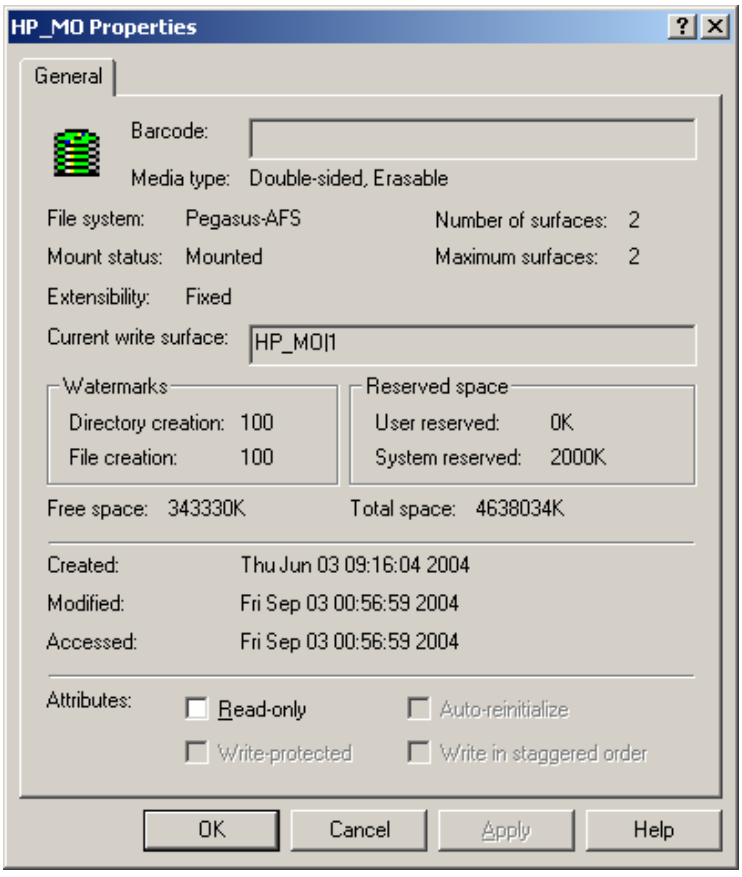

Intelligent Spanning

InveStore v4 Enterprise Storage Management Software includes volume sets that use our ***Intelligent Spanning*** capabilities. This provides application control of the volume set size. At format time, the application can choose one of the following options:

1. The volume set size is permanently set which prevents it from ever being grown.
2. The volume set size is programmed to grow automatically. When the volume set runs out of space, InveStore will automatically take any blank disk within the jukebox and add it to the volume set. An upper limit to such a volume set can be specified to prevent it from becoming too unwieldy.

To prevent catastrophic data loss, ***Intelligent Spanning*** does not store the entire directory on the root volume. Each surface within the volume set is self-describing. Should any surface within that volume become damaged, only the data on that surface is lost. All other members of the volume set are fully accessible allowing any volume set to be partially mounted and any number of cartridges can be offline for the volume set and the remaining surfaces will be fully accessible.

The order in which the disk surfaces of the volume set are written is fully configurable by the application.

1. Contiguous order writing instructs the system to write to side A followed by side B of each disk. This works well with single drive jukeboxes since only one surface can be accessed at a time.
2. Staggered order instructs the system to write to side A of each disk in the volume set followed by side B. This approach works best for multi-drive jukeboxes, as more data can be accessible at any given time since only one side of a disk can be accessed in the drive.

Intelligent Spanning automatically clusters all of the data within a subdirectory to a single surface whenever possible. The application can specify that the system reserve a given amount of space for each subdirectory created. If this space is not available, the directory and all its files are written to the next available surface.

Intelligent Spanning reserves a configurable amount of space for file updates on each surface. This space is not used for new files. This feature helps prevent files from being spread across multiple surfaces and greatly reduces the severe performance degradation that volume sets incur over time.

Volume Tracking

Volume management is an indexing system designed specifically for removable media. It allows users to track volumes (and their contents) whether they are stored in a drive, jukebox or on the shelf.

In volume management, each volume set is defined as a separate storage area for data or images. Since optical disks are removable from either a standalone drive or jukebox, the operating system addresses each volume set independently for "clustering" files, allowing a related group of files (or cartridge) to be removed without affecting the files remaining in a jukebox.

This capability is invaluable for document imaging or multi-media applications since documents typically consist of many pages of stored files.

InveStore maintains copies of each volume set's directory within a cache file on a magnetic disk. These cache files are copies of the up-to-date directory on the volume(s). The user (or application program) can browse through the directory structure without mounting the cartridge on an actual drive. Locating files or documents is quick and easy. This method significantly improves the search performance and significantly decreases system contention.

One of the most critical pieces of volume management is to insert and remove disks in an optimized fashion. The goal is to have the appropriate disk within a drive whenever it is possible and to minimize the user wait time whenever it is not. This approach to Volume Tracking helps to prevent a condition known as "thrashing" whereby disks are swapped excessively and needlessly.

InveStore
Storage Management Software

Standard Features

- Remote Administration
- Long term archival storage
- High speed access to near line storage
- Full tracking and directory access to off-line volumes
- Operating system Interchangeability
- Contention resolution (multiple users accessing multiple volumes with limited drive resources)
- Operating system independent
- File system independence
- Hardware independence
- Provides highest level of data integrity
- Hardware abstraction (you address files and directories, not drives, jukeboxes and disks)
- Volume set support that minimizes file and directory spanning
- Integrated caching for write operation

Pegasus Disk Technologies, Inc. 18 Crow Canyon Court, San Ramon, CA 94583
Telephone: (925) 314-1800 • Fax: (925) 314-1801 • Http://www.pegasus-afs.com • sales@pegasus-afs.com

InveStore
Storage Management Software

For More Product Information
Please Contact

Pegasus Disk Technologies Inc
2333 San Ramon Valley Blvd
San Ramon, CA 94583
(925) 314-1800

Or, you can receive product information by visiting our web page at: **www.pegasus-afs.com**