IBM Storwize V7000 Unified Disk System

The most powerful and easy-to-use innovative disk system in the storage marketplace

Highlights

- Delivers sophisticated enterprise-class storage function for businesses of all sizes
- Supports your growing business requirements while controlling costs
- Provides up to 200 percent performance improvement with automatic migration to high-performing SSDs[1]
- Consolidates block and file storage in a single system for simplicity and greater efficiency
- Enables near-continuous availability of applications through dynamic migration
- Supports faster and more efficient data copies for online backup, testing or data mining
- Offers fiexible server and storage management

In storage management today, breaking the cycle of increased complexity and explosive data growth can be a big challenge. The old ways of buying and managing storage have become less effective. Due to resource constraints—both physical storage resources and human resources—IT organizations must act quickly to optimize and simplify their infrastructure. Unchecked complexity and growth will only become bigger problems over time.

Small and midsize organizations may also suffer from a range of challenges:

- Disruptive migrations
- Difficulty deploying tiered storage
- Inability to share storage among servers
- Reduced productivity and increased cost caused by isolated server and storage management tools
- Inability to use virtualized storage, like virtual servers, as a tool for optimizing expenditures, resources and capabilities

To stand up to these challenges and allow businesses to respond to a rapidly changing marketplace, IBM® Storwize® V7000 Unified is a virtualized storage system to complement virtualized server environments that provides unmatched performance, availability, advanced functions

and highly scalable capacity never seen before in midrange disk systems. Storwize V7000 Unified is a powerful midrange disk system that has been designed to be easy to use and enable rapid deployment without additional resources. Storwize V7000 Unified consolidates block and file workloads into a single storage system for simplicity of management and reduced cost. Storwize V7000 Unified is virtual storage that offers greater efficiency and flexibility through built-in solid state drive (SSD) optimization and thin provisioning technologies. Storwize V7000 Unified advanced functions also enable nondisruptive migration of data from existing storage, simplifying implementation and minimizing disruption to users. Storwize V7000 Unified also enables you to virtualize and reuse existing disk systems, supporting a greater potential return on investment (ROI).

Managing the information infrastructure

The need to increase storage efficiency has led many IT organizations to turn to consolidation, virtualization and automated tiering to reduce capital and operational expenses. IBM offers solutions today that can become part of your highly efficient, highly capable, next-generation information infrastructure, whether your storage environment supports a small or midsize organization or a large, complex data center.

Consolidation

Optimizing resources through consolidation can reduce costs and improve productivity. Consolidation can also lead to more efficient maintenance and management of your information infrastructure. By enabling you to scale storage efficiently, consolidation can deliver the capacity you need within the budget you have for the performance you want.

Virtualization

Virtualizing your storage infrastructure can optimize your expenditures, resources and capabilities. It allows you to scale system capacity and performance more easily to meet your growing information infrastructure needs, reduce the complexity of management and reduce the risk to your business of system failure. In server environments, virtualization technologies are often used to improve server utilization, reduce complexity, speed provisioning, consolidate application migration and provide improved flexibility in disaster recovery plans. Storage virtualization is designed to provide similar advantages for your storage environment. Combining storage and server virtualization can build a more powerful virtualized infrastructure for your business and provide greater benefits than either virtualization solution deployed alone.

Support for VMware vStorage APIs enables Storwize V7000 Unified to take on some storage-related tasks that were previously performed by VMware, which helps improve efficiency and frees up server resources for other more mission-critical tasks.

Tiering

Tiering optimizes storage by enabling data to be located in a way that can improve system performance, reduce costs and simplify information management. Tiering can enhance performance and reduce operating expenses by automating data movement. And tiering allows you to scale storage performance based upon your business needs. Using IBM System Storage® Easy Tier™ technology, you can deploy SSDs confidently, effectively and economically by automatically and dynamically moving only the appropriate data to the SSDs in the system, based on ongoing performance monitoring. Such effective storage tiering enables users to enjoy the performance benefits of SSDs without requiring administrators to create and manage storage tier policies. It does so without the excessive costs associated with placing too much of the wrong data on these relatively expensive drives. With a wide range of disk drives and SSDs, Storwize V7000 Unified provides the ability to optimize a tiered storage configuration to meet diverse application requirements.

Introducing Storwize V7000 Unified

Storwize V7000 Unified is a powerful storage system that combines hardware and software components to provide a single point of control to help support improved storage efficiency. By enabling virtualization, consolidation and tiering in business of all sizes, it is designed to improve application availability and resource utilization. The system offers easy-to-use, efficient and cost-effective management capabilities for both new and existing storage resources in your IT infrastructure.

Storwize V7000 Unified combines together a variety of IBM technologies including thin provisioning, automated tiering, storage virtualization, clustering, replication, multi-protocol support and a next-generation graphical user interface (GUI). Together, these technologies are designed to enable Storwize V7000 Unified to deliver extraordinary levels of storage efficiency.

Enhancing access with Easy Tier

Easy Tier provides automatic migration of frequently accessed data to high performing SSDs, enhancing usage efficiencies. Operating at a fine-grained granularity, the Easy Tier function automatically repositions pieces of the data to the appropriate class of drives based on I/O patterns and drive characteristics with no further administrative interaction.

Easy Tier makes it easy and economical to deploy SSDs in your environment. A hybrid pool of storage capacity is created containing two tiers: SSD and hard disk drive (HDD).

- The busiest portions of volumes are identified and automatically relocated to high-performance SSDs.
- Remaining data can take advantage of higher capacity, price-optimized drives for the best customer value.

Volumes in an SSD- or HDD-managed disk group are monitored and can be managed automatically or manually by moving hot extents to SSD and cold extents to HDD.

With an online database workload, Easy Tier improved throughput up to 200 percent and reduced transaction response time by up to 30 percent compared to a configuration using only HDD.

Using thin provisioning to optimize efficiency

Using thin provisioning, applications consume only the space they are actually using, not the total space that has been allocated to them. Designed to keep business overhead low, thin provisioning optimizes efficiency by allocating disk storage space in a flexible manner among multiple users, based on the minimum space required by each user at any given time. This reduces use of storage hardware but also can save electrical energy use, lower heat generation and reduce hardware space requirements.

For example, a database might be planned to grow to 100 TB but is only 10 TB today. Using thin provisioning, a storage administrator can allocate 100 TB of virtual capacity to meet expected future requirements while consuming only 10 TB of physical capacity. As the database grows, Storwize V7000 Unified allocates additional physical capacity as required. This approach minimizes the amount of physical capacity consumed and helps enable a more efficient approach to storage purchases while also minimizing future configuration changes as the database grows.

Avoiding disruptions with dynamic migration

Storwize V7000 Unified uses virtualization technology to help insulate host applications from physical storage changes. This ability can help enable applications to run without disruption while you make changes to your storage infrastructure. Your applications keep running so you can stay open for business.

Moving data is one of the most common causes of planned downtime. Storwize V7000 Unified includes a dynamic data migration function that is designed to move data from existing storage into the new system or between arrays in a Storwize V7000 Unified system, while maintaining access to the data. The data migration function might be used, for example, when replacing older storage with newer storage, as part of load balancing work or when moving data in a tiered storage infrastructure.

Using the Storwize V7000 Unified dynamic migration capabilities can provide efficiency and business value. Dynamic migration can speed time-to-value from weeks or months to days, minimize downtime for migration, eliminate the cost of add-on migration tools, and may help avoid penalties and additional maintenance charges for lease extensions. The result can be real cost savings to your business.

Consolidating data for efficiency and simplicity

Many users have deployed storage area network (SAN) attached storage for their applications requiring the highest levels of performance while separately deploying network attached storage (NAS) for its ease of use and lower cost networking. This divided approach adds complexity by introducing multiple management points and also creates islands of storage that reduce efficiency.

Storwize V7000 Unified provides the ability to combine both block and file storage into a single system. By consolidating storage systems, multiple management points can be eliminated and storage capacity can be shared across both types of access, helping to improve overall storage utilization. Storwize V7000 Unified also presents a single, easy-to-use management interface that supports both block and file, helping to simplify administration further.

Storwize V7000 Unified builds on the high-performance design of the Storwize V7000 system and integrates proven IBM software capabilities to deliver new levels of efficiency. The system includes IBM Active Cloud Engine™, which is designed to deliver policy-based management of files to reduce costs through use of tiered storage and improve data governance. For example, Active Cloud Engine policies can be used to move less frequently used data to lower cost tiers of storage, including tape in an IBM Tivoli® Storage Manager system. Active Cloud Engine policies can also be used to delete unwanted or expired data automatically.

Storwize V7000 Unified integration with McAfee and Symantec Antivirus is designed to provide the ability to isolate or delete compromised files, and leverage the most commonly deployed ISV antivirus applications.

Foundation for cloud deployments

Improving efficiency and delivering a flexible, responsive IT infrastructure are essential requirements for any cloud deployment. Key technologies for delivering this infrastructure include virtualization, consolidation and automation.

With its virtualized storage design and tight affinity with technologies such as IBM PowerVM™ and VMware, Storwize V7000 Unified is the ideal complement for virtualized servers that are at the heart of cloud deployments. Storwize V7000 Unified helps enable consolidation of multiple storage systems for greater efficiency and clustered systems drive the value of consolidation even further. Automated tiering technologies such as Easy Tier, Active Cloud Engine and Tivoli software help make the best use of the storage resources available.

Protecting data with replication services

Storwize V7000 Unified includes a rich IBM FlashCopy® function that is designed to create an almost instant copy of active data, which can be used for backup purposes or for parallel processing activities. Up to 256 copies of each volume may be created.

When combined with Storwize V7000 Unified thin provisioning, you can create copies using only a fraction of the amount of storage needed for a full physical copy. This type of copy, called a “snapshot,” is designed to help improve overall storage utilization and reduce the amount of capacity required for copies.

IBM Tivoli Storage FlashCopy Manager is designed for today's business world, where application servers are operational 24 hours a day—yet data must remain fully protected. If you have a 24×7 environment, you can't afford to lose any data. You also can't afford to stop critical systems for hours so you can protect the data adequately. FlashCopy Manager exploits the Storwize V7000 Unified snapshot capabilities to provide high-speed, backup and restore functionality with low impact to applications. Automated policy-based management of multiple snapshot backup versions, combined with a simple and guided installation and configuration process provide an easy-to-use and quick-to-deploy data protection solution that enables the most stringent database recovery time requirements to be met. FlashCopy Manager can help deliver the highest levels of protection for mission critical IBM DB2®, SAP, Oracle, Microsoft Exchange and Microsoft SQL Server applications through integrated application-aware snapshot backup and restore capabilities. Custom application support offers the ability to extend FlashCopy Manager capabilities to any application on IBM AIX®, Linux and Solaris.

The Metro Mirror and Global Mirror functions operate between Storwize V7000 Unified systems at different locations to help create copies of data for use in the event of a catastrophic event at a data center. For even greater flexibility, Metro Mirror and Global Mirror also support replication between Storwize V7000 Unified systems and a wide variety of storage systems supported by System Storage SAN Volume Controller. Metro Mirror is designed to maintain a fully synchronized copy at "metropolitan" distances (up to 300 km) whereas Global Mirror is designed to operate asynchronously and so helps maintain a copy at much greater distances (up to 8000 km). Both functions are designed to support VMware vCenter Site Recovery Manager to help speed disaster recovery. Enhancements to Global Mirror are designed to provide new options to help administrators balance network bandwidth requirements and recovery point objectives for applications, helping to reduce operating costs for disaster recovery solutions.

IBM Tivoli Storage Manager FastBack® provides an additional complementary capability to replicate highly-efficient deduplicated snapshots over TCP/IP connections to a remote FastBack Disaster Recovery hub, efficiently storing the disaster recovery DR snapshots on a Storwize V7000 Unified. This option helps provide a lower cost approach to delivering effective disaster recovery capability.

For NAS data, Storwize V7000 Unified offers data protection through a space-efficient file system and fileset-level snapshots (up to 256 per file system). Snapshots of a fileset provide a way to partition the namespace into smaller, more manageable units. File snapshots protect against accidental deletion or modification of files and that enable you to restore at the file level. The system also provides asynchronous replication for disaster recovery and business continuity. In addition, asynchronous replication offers encrypted file replication over extended distances between two sites. This function is integrated with IBM Active Cloud Engine, which can provide a high-speed scan of the source file system to determine files and directories that have been created, modified or deleted. The replication is performed by an "rsync" tool created by IBM that can move only the changed portions of files to a destination, offering network savings.

Storwize V7000 Unified has specific exploitation and integration with IBM Tivoli Storage Manager to provide efficient and extremely fast backup and restore processes, and the movement of files to external disk or tape. In addition, Storwize V7000 Unified provides support for the Network Data Management Protocol (NDMP) to provide full and incremental backup of files as well as restoring of these files and related file system data. Support for NDMP allows for backing up Storwize V7000 Unified with third-party backup applications over the LAN.

Leverage proven ISV solutions

IBM is committed to continuous improvement and seamless application integration to optimize your business results and minimize time-to-value. Our commitment is visible through ongoing work and enduring partnerships with ISVs such as Microsoft, Oracle, SAP, Symantec and VMware.

Combining Storwize V7000 Unified with leading ISV applications can provide increased flexibility and deliver a more robust information infrastructure for your business. Solutions have been qualified for Storwize V7000 Unified for select applications that focus on key solution areas, including backup/restore, disaster recovery, clustering, server virtualization and database and performance optimization. IBM is also committed to certifications with key ISVs aligned with various industries including healthcare, financial services, telecommunications and the public sector.

Integrated management

This solution provides a tiered approach to management designed to meet the diverse needs of different organizations. The Storwize V7000 Unified management interface is designed to give administrators intuitive control of the system and provides a single integrated approach for managing both block and file storage requirements in the same system. For organizations looking to manage both physical and virtual server infrastructures and the storage they consume (including provisioning and monitoring for higher availability, operational efficiency and infrastructure planning) Storwize V7000 Unified is integrated with IBM Systems Director Storage Control. A single administrator can manage and operate IBM servers (IBM System x®, IBM Power Systems™ and IBM BladeCenter®) along with networking infrastructure and IBM storage (including Storwize V7000 Unified) from a single management screen. For organizations looking to improve the operational efficiency of storage specialists, Tivoli Storage Productivity Center is designed to provide a SAN-wide perspective of storage health, I/O path performance analytics and capacity use for Storwize V7000 Unified and the surrounding storage infrastructure. Plug-ins to support Storwize V7000 Unified with Microsoft System Center Operations Manager (SCOM) and VMware vCenter help enable more efficient consolidated management in these environments.

A new performance dashboard provides at-a-glance access to key high-level real-time system performance information, which helps monitor and optimize the virtualized environment. IBM Tivoli Storage Productivity Center provides access to—and analysis of—historical performance data.

High-performance SSD support

For applications that demand high disk speed and quick access to data, IBM provides support for SSDs in 200, 300 or 400 GB 2.5-inch E-MLC (enterprise-grade multilevel cell) SSDs, or up to 96 TB of physical capacity in a single system enabling scale-out high performance SSD support.

External storage virtualization

External storage virtualization is the ability of the Storwize V7000 Unified system to manage capacity in other disk systems. When Storwize V7000 Unified virtualizes a disk system, its capacity becomes part of the Storwize V7000 Unified system and is managed in the same way as capacity on internal drives within Storwize V7000 Unified. Capacity in external disk systems inherits all the functional richness and ease-of-use of the Storwize V7000 Unified system including advanced replication, thin provisioning and Easy Tier. When virtualizing external storage, virtualizing external storage helps improve administrator productivity and boost storage utilization while also enhancing and extending the value of an existing storage asset.

Storwize V7000 Unified system description

The Storwize V7000 Unified storage system is packaged in 2U rack-mountable enclosures that house up to twenty-four 2.5-inch drives or up to twelve 3.5-inch drives. Control enclosures contain drives, redundant dual-active intelligent RAID controllers and dual power supplies, batteries and cooling components. Expansion enclosures contain drives, switches, power supplies and cooling components. You can attach up to nine expansion enclosures to a control enclosure supporting up to 240 drives. Two control enclosures (each with up to nine expansion enclosures) may be clustered together in a single system for even greater capacity and performance growth potential. Other components and characteristics of the system include:

- Internal storage capacity: Up to 36 TB of physical storage per enclosure using twelve 3 TB near-line SAS disk drive modules or up to 14.4 TB of physical storage per enclosure using twenty-four 2.5-inch 600 GB SAS disk drive modules
- Disk drives: SAS disk drives, near-line SAS disk drives and SSDs. Intermix of these drive types in the Storwize V7000 Unified controller and expansion enclosures add flexibility
- Cache memory: Sixteen GB cache memory (8 GB per internal RAID controller) as a base feature—designed to improve performance and availability
- Ports per control enclosure: Eight 8 Gbps Fibre Channel host ports (four 8 Gbps FC ports per RAID controller), four 1 Gbps and optionally four 10 Gbps iSCSI host ports (two 1 Gbps and optionally two 10 Gbps iSCSI host ports per RAID controller)
- Ports per File Module: Two 10Gbps Ethernet ports for server attachment, two 8 Gbps FC ports for attachment to Storwize V7000 control enclosures

Power and cooling (typical environments)

	Power consumption	Cooling
12-bay control enclosure	380 W	1300 BTU/hr
24-bay control enclosure	410 W	1400 BTU/hr
12-bay expansion enclosure	175 W	600 BTU/hr
24-bay expansion enclosure	205 W	700 BTU/hr
File Module	150 W	520 BTU/hr

Storwize V7000 Unified control and expansion enclosures are each available in models that support twelve 3.5-inch disk drive bays or twenty-four 2.5-inch disk drive bays. The system supports intermixing 12-bay or 24-bay enclosures in a single system. The expansion enclosures connect to the control enclosure using four SAS 6 Gbps disk expansion ports.

- Control enclosure: Supporting attachment of up to nine expansion enclosures with configurations up to 360 TB physical storage capacities.
- Expansion enclosure: Packaged in a 2U rack-mountable enclosure that houses twenty-four 2.5-inch drive bays or twelve 3.5-inch drive bays and dual power supplies with cooling components. Physical storage capacity of up to 36 TB per storage expansion enclosure using twelve 3.5-inch 3 TB near-line serial attached SCSI (SAS) disk drive modules and up to 14.4 TB per storage expansion enclosure using twenty-four 2.5-inch 600 GB SAS disk drive modules.
- File module: Packaged in a 2U rack-mountable enclosure, a Storwize V7000 file module provides attachment to 10 Gbps NAS environments. File modules are always deployed in pairs for redundancy.
- Clustered systems: Two control enclosures (each with up to nine expansion enclosures) may be clustered together in a single system for even greater capacity and performance growth potential. Clustered systems support up to 480 disk drives and up to 720 TB capacity.

Electrical power

- 12-bay and 24-bay control enclosures: 120 - 240 V ac, 3.8 - 9.0 A, 50/60 Hz
- 12-bay and 24-bay expansion enclosures: 100 - 240 V ac, 3.2 - 8.0 A, 50/60 Hz
- File modules: 100 - 240 V ac, 3.8 - 7.8 A, 50/60 Hz

Environment: all systems

- Temperature (operating)
  - 10° to 35°C (50° to 95°F) at 0 to 914 m (0 to 3,000 ft)
  - 10° to 32°C (50° to 90°F) at 914 to 2,133 m (3,000 to 7,000 ft)
- Temperature (powered off):
  - 10° to 43°C (50° to 109°F)
- Temperature (storage):
  - 1° to 60°C (34° to 140°F) at 0 to 2,133 m (0 to 7,000 ft)
- Temperature (shipping):
  - 20° to 60°C (-4° to 140°F) at 0 to 10,668 m (0 to 35,000 ft)
- Relative humidity (operating and powered off): 8 percent to 80 percent
- Relative humidity (storage): 5 percent to 80 percent
- Relative humidity (shipping): 5 percent to 100 percent (including condensation but excluding rain)
- Wet bulb
  - Wet bulb (operating temp): 23°C
  - Wet bulb (powered off temp): 27°C
  - Wet bulb (storage and shipping temp): 29°C
- Noise level: 6.5 decibels LwAd—when operating in a 2146 system rack

Note: The noise emission level stated is the declared (upper limit) sound power level, in decibels, for a random sample of machines. All measurements are made in accordance with ISO 7779 and reported in conformance with ISO 9296.

Host interface	SAN-attached 8 Gbps Fiber Channel (FC), 1 Gbps iSCSI and optional 10 Gbps iSCSI NAS-attached 10 Gbps Ethernet
User interface	Graphical user interface (GUI)
Supported drives	3.5-inch disk drives: • 2 TB, 3 TB 3.5 in. 7.2k Near-Line SAS disk 2.5-inch disk drives: • 146 GB, 300 GB 2.5 in. 15k SAS disk • 300 GB, 450 GB, 600 GB 2.5 in.10k SAS disk • 200 GB, 300 GB, 400 GB 2.5 in. E-MLC (enterprise-grade multilevel cell) solid state drive (SSD) • 1 TB 2.5 in. 7.2k Near-Line SAS disk
RAID levels	RAID 0, 1, 5, 6 and 10
Maximum drives supported	240 per control enclosure; 480 per clustered system
Fans and power supplies	Fully redundant, hot swappable
Rack support	Standard 19 inch
Management software	Storwize V7000 Unified software
Cache per controller/control enclosure/clustered system	8 GB/16 GB/32 GB
Advanced features included with each system	System Storage Easy Tier, FlashCopy, thin provisioning
Additional available advanced features	Remote mirroring, external virtualization unified storage, IBM® FlashCopy Manager, IBM Tivoli® Storage Productivity Center Midrange Edition, IBM Tivoli Storage Manager, IBM Tivoli Storage Manager FastBack®, IBM Systems Director, IBM Active Cloud Engine™
Warranty	Hardware: • 3 year limited warranty • Customer replaceable units • On-site service • Next Business Day 9×5 • Service upgrades available Software: • Software Maintenance Agreement available

Replication services	FlashCopy, FlashCopy Manager, Metro Mirror (Synchronous), Global Mirror (Asynchronous), local and asynchronous remote file-based replication
Dimensions	Control and expansion enclosures • Width: 483 mm (19.0 in.) • Depth: 630 mm (24.8 in.) • Height: 87.9 mm (3.46 in.) File Modules • Width: 482 mm (19.0 in.) • Depth: 729 mm (28.7 in.) • Height: 85.2 mm (3.45 in.)
Weight	12-bay enclosures: • Drive-ready (without drive modules installed): 17.7 kg (37.6 lb) • Fully configured (12 drive modules installed): 27.2 kg (59.8 lb) 24-bay enclosures: • Drive-ready (without drive modules installed): 17.7 kg (37.6 lb) • Fully configured (24 drive modules installed): 25.2 kg (55.4 lb)
Supported systems	For a list of currently supported servers, operating systems, host bus adapters, clustering applications and SAN switches and directors, refer to the System Storage Interoperation Centre.
ISV solutions	For a list of high quality solutions with our partner **Independent Software Vendors (ISVs), including access to solution briefs and white papers, refer to the ISV Solutions Resource Library.**

Why IBM?

The performance and availability of your storage environment can either enhance or hamper your business processes. That's where IBM comes in. As a market leader in the storage industry, we can help you handle the challenges, whether you are a small to midsize company or a large enterprise.

Innovative technology, open standards, excellent performance, a broad portfolio of proven storage software, hardware and solutions offerings—all backed by IBM with its recognized industry leadership—are just a few of the reasons you should consider storage solutions from IBM, including Storwize V7000 Unified.

With IBM, you get some of the best storage products, technologies, services and solutions in the industry without the complexity of dealing with different hardware and software vendors and system integrators.

For more information

To learn more about IBM Storwize V7000 Unified, please contact your IBM sales representative or IBM Business Partner, or visit the following website: **ibm.com**/storage/storwizev7000

For a list of currently supported servers, operating systems, host bus adapters, clustering applications and SAN switches and directors, refer to the System Storage Interoperation Center at: **ibm.com**/systems/support/storage/config/ssic

For a list of high quality solutions with our partner Independent Software Vendors (ISVs), including access to solution briefs and white papers, refer to: **ibm.com**/systems/storage/solutions/isv

Additionally, financing solutions from IBM Global Financing can enable effective cash management, protection from technology obsolescence, improved total cost of ownership and return on investment. Also, our Global Asset Recovery Services help address environmental concerns with new, more energy-efficient solutions. For more information on IBM Global Financing, visit: **ibm.com**/financing

IBM Systems and Technology Group
Route 100
Somers, NY 10589
U.S.A.

Produced in the United States of America
October 2011

[1] IBM lab measurements - August 2010

Please Recycle

TSD03111-USEN-05

Storwize V7000 Unified overview

Contents

Storwize V7000 Unified overview

The IBM® Storwize® V7000 Unified system is a virtualizing RAID storage system that supports both block protocols and file protocols. This unified system includes the IBM Storwize V7000 File Module and the Storwize V7000 storage system.

Storwize V7000 Unified hardware

The file module is a clustered system comprised of two units that provide file systems for use by network-attached storage. The file module uses the Storwize V7000 storage system to provide the file module with volumes. Other volumes, which are block volumes, are provided on the SAN to be presented to hosts.

The Storwize V7000 storage system consists of a drive enclosure called the *control enclosure*. Both regular and solid-state drives (SSDs) are supported. The control enclosure contains disk drives and two *node canisters* that are managed as a single clustered system. *Expansion enclosures* contain drives and are attached to the control enclosure. *Expansion canisters* include the serial-attached SCSI (SAS) interface hardware that enables the node hardware to use the drives of the expansion enclosures.

Figure 1 shows the Storwize V7000 system as a traditional RAID storage system. The internal drives are configured into arrays, and volumes are created from those arrays.

Figure 1. Storwize V7000 system as a RAID storage system

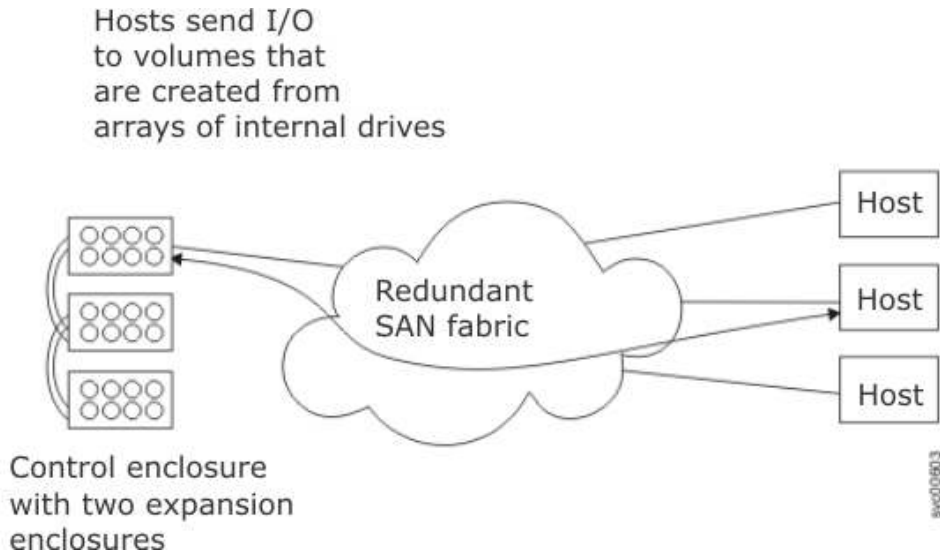

The Storwize V7000 system can also be used to virtualize other storage systems as shown in Figure 2.

Figure 2. Storwize V7000 system virtualizing other storage systems

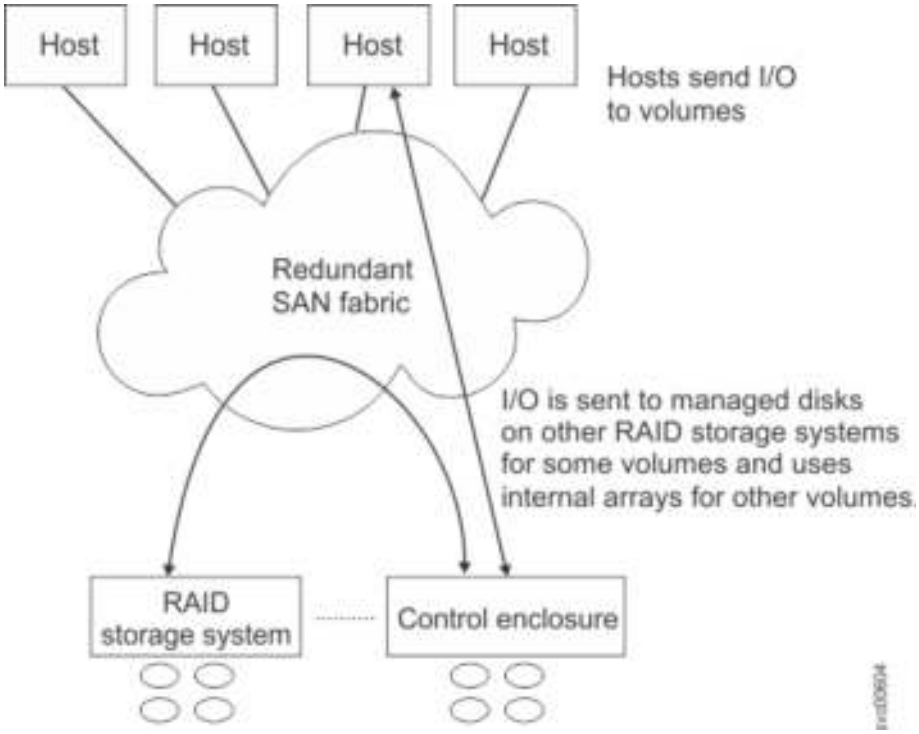

The two node canisters in each control enclosure are arranged into pairs known as *I/O groups*. A single pair is responsible for serving I/O on a given volume. Because a volume is served by two node canisters, there is no loss of availability if one node canister fails or is taken offline.

Figure 3 shows how a Storwize V7000 Unified system can be configured.

Figure 3. Storwize V7000 Unified configuration

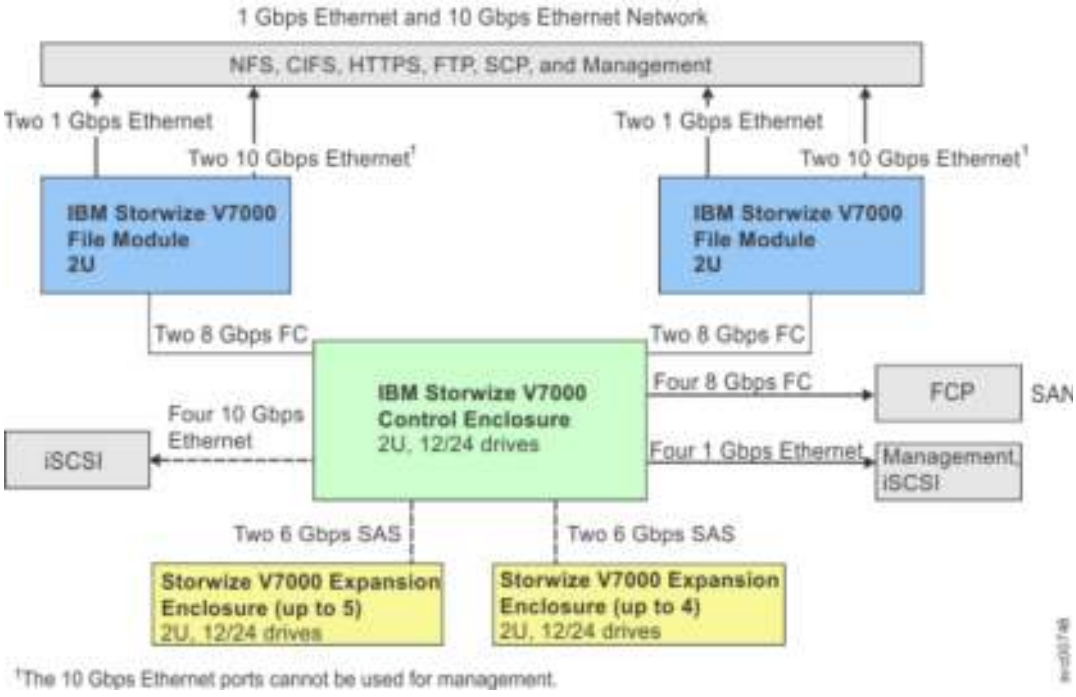

Storwize V7000 Unified software

The IBM Storwize V7000 File Module Software within the Storwize V7000 Unified system contains the interface node and management node functions. The management node function provides for configuring, administering, and monitoring a system. The interface node function connects a system to an Internet Protocol (IP) network. You can connect by using any of the following protocols:

- Common Internet File System (CIFS)
- Network File System (NFS)
- File Transfer Protocol (FTP)
- Hypertext Transfer Protocol Secure (HTTPS)

- Secure Copy Protocol (SCP)

The Storwize V7000 Unified software also performs the following functions for the host systems that attach to the Storwize V7000 Unified system by using block protocols:

- Creates a single pool of storage
- Provides logical unit virtualization
- Manages logical volumes
- Mirrors logical volumes
- Large scalable cache
- Copy Services
  - IBM FlashCopy® (point-in-time copy) function, including thin-provisioned FlashCopy to make multiple targets affordable
  - Metro Mirror (synchronous copy)
  - Global Mirror (asynchronous copy)
  - Data migration
- Space management
  - IBM System Storage® Easy Tier™ to migrate the most frequently used data to higher performing storage
  - Metering of service quality when combined with IBM Tivoli® Storage Productivity Center
  - Thin-provisioned logical volumes

Figure 4 shows an example of a management GUI view of the various components in a Storwize V7000 Unified system.

Figure 4. Storwize V7000 Unified components

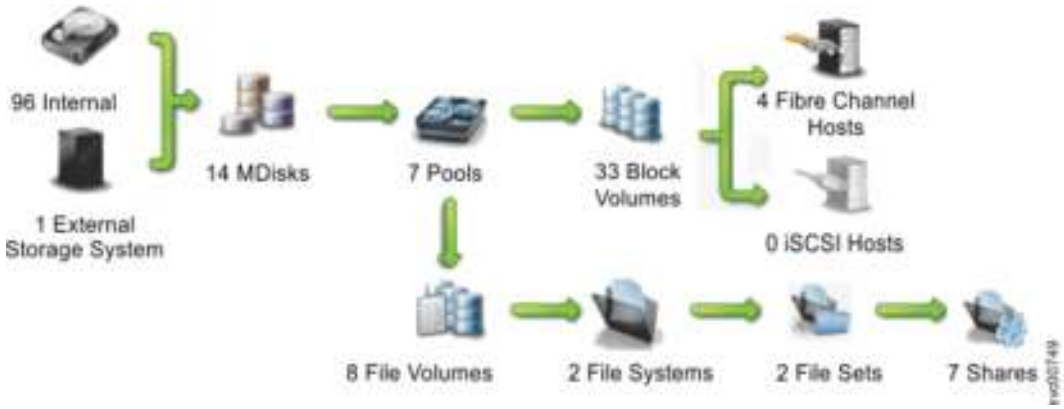

System management

The management node function within the file module presents a single point of control for system management and service. You use the command line interface (CLI) or the graphical user interface (GUI) for configuring, administering, and monitoring the system. The management node function also provides a CLI interface and web interface for performing hardware service actions.

Fabric types

I/O operations between hosts and Storwize V7000 nodes and between Storwize V7000 nodes and RAID storage systems are performed by using the SCSI standard. The Storwize V7000 nodes communicate with each other by using private SCSI commands.

Storwize V7000 uses the SCSI commands over the Fibre Channel SAN and either 1 Gbps Ethernet or 10 Gbps Ethernet. Table 1 shows the fabric type that can be used for communicating between hosts, nodes, and RAID storage systems. These fabric types can be used at the same time.

Table 1. Storwize V7000 communications types

Communications type	Host to Storwize V7000	Storwize V7000 to storage system	Storwize V7000 to Storwize V7000
Fibre Channel SAN	Yes	Yes	Yes

Communications type	Host to Storwize V7000	Storwize V7000 to storage system	Storwize V7000 to Storwize V7000
iSCSI (1 Gbps Ethernet or 10 Gbps Ethernet)	Yes	No	No

Note: To reach a Storwize V7000 Unified system using external clients, a public IP network must be configured and available. This includes all network interfaces used to provide NAS services to external clients in the customer network. These NAS services network interfaces are configured exclusively on the file modules.

Parent topic: Product overview

IBM Tivoli Storage Manager, Version 6.3

Deduplicating data

Data deduplication is a method for eliminating redundant data. Only one instance of the data is retained on storage media, such as disk or tape. Other instances of the same data are replaced with a pointer to the retained instance.

Deduplicated data must be stored in sequential-access disk (FILE) primary, copy, and active-data storage pools. that you enable for data deduplication. Because you can store more data on disk than on tape, data deduplication can reduce the overall amount of time that is required to retrieve data.

Restriction: When a client backs up or archives a file, the data is written to the primary storage pool specified by the copy group of the management class that is bound to the data. To deduplicate the client data, the primary storage pool must be a sequential-access disk (FILE) storage pool that is enabled for data deduplication.

consumption occurs on the client system during the backup. With server-side data deduplication, the processor consumption occurs on the server during the duplicate identification process. Consider factors such as processor usage, network bandwidth, restore performance, and compression when you decide to use data deduplication.

Data deduplication and data compatibility
With client-side data deduplication, data is compressed extent by extent, which means data compatibility is possible between the client and the server and between different versions of Tivoli Storage Manager. Also, data extents that are created by different operations are compatible, which can mean that a greater proportion can be deduplicated.

Data deduplication and disaster recovery management
Use the following scenarios to help you consider factors that can affect disaster recovery management when you use deduplicated primary or copy storage pools.

Parent topic: Managing storage pools and volumes

Feedback

IBM Storwize V7000 Version 6.2 delivers support for VMware VAAI, real-time performance monitoring, and 10 Gigabit iSCSI connectivity

Table of contents

At a glance

The IBM® Storwize V7000 disk system utilizes IBM System Storage® SAN Volume Controller (SVC) technology to deliver high performance, advanced function, high availability, and modular and scalable storage capacity:

- Supports RAID 0, 1, 5, 6, and 10
- Provides SAN-attached 8 Gbps Fibre Channel (FC) host connectivity, and 1 Gigabit Ethernet (GbE) and 10 GbE iSCSI host connectivity
- Accommodates up to twenty-four 2.5-inch disk drives or twelve 3.5-inch disk drives installed within the IBM Storwize V7000 Control Enclosure, with attachment support for up to nine IBM Storwize V7000 Expansion Enclosures, providing modular and highly scalable storage solutions that range up to 240 TB physical storage capacity, and 480 TB physical storage capacity in a clustered system
- Supports intermixing of SAS drives, Nearline SAS drives, and solid-state drives within the IBM Storwize V7000 Control Enclosures and the IBM Storwize V7000 Expansion Enclosures
- Includes IBM Easy Tier technology for automatically moving heavily used data extents onto high-performance storage devices
- Supports attachment of other storage devices via the FC interface just as the SVC
- Supports a complete set of SVC functions, including FlashCopy®, RemoteCopy, Volume Mirroring, and thin provisioning

Overview

IBM Storwize V7000 V6.2.0 delivers improved tracking with the introduction of real-time performance monitoring. Immediate performance information can be received, including CPU utilization and I/O rates, to monitor environmental changes, troubleshoot, and when pairing up the information with historical detailed data from Tivoli® Storage Productivity Center, can be better positioned to develop the best performance solutions.

VMware virtual environments can be improved with Storwize V7000 by using the vStorage API for Array Integration (VAAI). This new API delegates certain VMware functions to the Storwize V7000 for enhanced performance. In the vSphere 4.1 release, this offload capability to Storwize V7000 supports full copy, block zeroing, and hardware-assisted locking. The introduction of 10 GbE hardware for the Storwize V7000 allows continued focus on cost efficiency with higher network performance by offering 10 Gigabit iSCSI host attachment. The Storwize V7000

also provides up to twice the capacity and twice the throughput to support business growth by providing clustered systems that include up to two Storwize V7000 control enclosures. With external storage virtualization, the Storwize V7000 is scalable to manage up to 32 PB of storage by allowing managed disks to be as large as 256 TB on select storage systems.

Storwize V7000 interoperability now supports managing additional storage products, including HP StorageWorks P9500 Disk Array, Hitachi Data Systems Virtual Storage Platform, Texas Memory Systems RamSan-620, and EMC VNX models.

Key prerequisites

Software requirements

IBM Storwize V7000 Software V6.2 (5639-VM1, 5639-EV1, and 5639-RM1) has the same software prerequisites as IBM Storwize V7000 V6.1. Refer to the Description section and the announcement letters listed in the Reference information section for more information.

Hardware requirements

This program requires at least one IBM Storwize V7000 control enclosure (2076-112, 2076-124, 2076-312, or 2076-324) for installation. Current support summaries, including specific software, hardware, and firmware levels supported, are maintained at the IBM Storwize V7000 Support website

http://www.ibm.com/storage/support/storwize/v7000

For information on the IBM Storwize V7000 hardware, see the announcements listed in the Reference information section.

Planned availability date

- June 10, 2011: IBM Storwize V7000 Software V6.2.0 AAS new order shipments begin
- June 30, 2011: Availability of the Plug-In for VMware vCenter

Description

IBM Storwize V7000 V6.2 highlights include support for real-time performance monitoring, VMware VAAI, and 10 Gigabit iSCSI connectivity. The following functions, features, and support options help maximize the benefits of Storwize V7000.

Real-time performance monitoring

IBM Storwize V7000 V6.2.0 provides improved tracking by introducing real-time performance monitoring. Real-time performance monitoring gathers the following system level performance statistics:

- CPU utilization
- Port utilization and I/O rates
- Volume and MDisk I/O rates
- Bandwidth
- Latency

Real-time performance monitoring can immediately track environmental changes and for performance troubleshooting. The sampling rate can be set as low as five seconds. Historical data is kept for about five minutes, therefore Tivoli Storage Productivity Center can capture more detailed performance information and

to analyze trends in order to get a complete picture when developing the best performance solutions.

VMware vStorage API for Array Integration (VAAI)

VMware virtual environments can be improved with Storwize V7000 by using VAAI. This new API delegates certain VMware functions to Storwize V7000 to enhance performance. In the vSphere 4.1 release, this offload capability to Storwize V7000 supports:

- Full copy: Enables the Storwize V7000 to make full copies of data without requiring the ESX server to read and write the data
- Block zeroing: Enables Storwize V7000 to zero out large numbers of blocks, shortening the time needed to provision virtual machines
- Hardware-assisted locking: Improves the scalability of large ESX server farms sharing volumes on the Storwize V7000 by providing an alternate means to protect the metadata for VMFS cluster-file systems

Plug-In for VMware vCenter

The Storwize V7000 Plug-In for VMware vCenter allows VMware administrators to monitor and manage Storwize V7000 volumes from the VMware software. This plug-in enables a common set of Storwize V7000 tasks to be performed from VMware vCenter without needing to switch between user interfaces. This plug-in will initially be certified for use with Storwize V7000 V6.1.

10 Gigabit iSCSI

The new 10 Gigabit Ethernet hardware for the Storwize V7000 have the ability to attach hosts to the Storwize V7000 using 10 Gigabit iSCSI.

Clustered systems

Storwize V7000 now allows clustered systems with up to two Storwize V7000 control enclosures. Each control enclosure in a clustered system has the advantage of processing unique host system I/O workloads while configuration and management of the storage is performed at the system level. This provides up to twice the capacity and twice the throughput to support business growth, and it enables more consolidation, including external virtualized storage, for greater efficiency.

256 TB managed disks and 32 PB managed storage

Storwize V7000 now allows virtualization of managed disks up to 256 TB in capacity for select storage systems, improved from 2 TB. With the use of 8 GB extent sizes, one Storwize V7000 system can now effectively manage up to 32 PB of storage, up from the 8 PB maximum. For a list of the storage systems supporting managed disks up to 256 TB, visit

http://www.ibm.com/storage/support/2145

FlashCopy target as Remote Copy source

Using a FlashCopy target volume as the source volume in a remote copy relationship is now available in SVC V6.2. This complements the current capability for remote copy target volumes to be source volumes in a FlashCopy relationship, and allows administrators greater flexibility in setting up disaster recovery configurations.

Revised licensing rules for Remote Mirroring Software (5639-RM1)

Rather than purchasing a license for IBM Storwize V7000 Remote Mirroring Software (5639-RM1) for every licensed enclosure (internal and external) managed by the IBM Storwize V7000 Disk System at both the primary and the secondary systems as first released, some customers will benefit from the following modifications to the 5639-RM1 licensing rules:

- For primary and secondary systems with differing numbers of enclosures using Remote Mirroring, the number of 5639-RM1 licenses needed for each system is the number of enclosures on the smaller of the two systems (see scenarios 1 and 2 below)
- When multiple production systems replicate to a single disaster recovery system, the number of 5639-RM1 licenses at the disaster recover system must equal the sum of the 5639-RM1 licenses at the production systems (see scenario 3 below)

The following are examples of how to license Remote Mirroring Software (5639-RM1) under the modified licensing rules:

Scenario 1: The primary system is a four-enclosure Storwize V7000 with nothing externally virtualized, therefore it has four 5639-VM1 licenses. The secondary system is a two-enclosure Storwize V7000 with nothing externally virtualized, therefore it has two 5639-VM1 licenses. The Storwize V7000 Remote Mirroring licensing would be as follows: two 5639-RM1 licenses for the primary system plus two 5639-RM1 licenses for the secondary system, for a total of four 5639-RM1 licenses required.

Scenario 2: The primary system is a one-enclosure Storwize V7000 managing a DS5020 consisting of three enclosures, therefore it has one 5639-VM1 license plus three 5639-EV1 licenses. The secondary system is a three-enclosure Storwize V7000 with nothing externally virtualized, therefore it has three 5639-VM1 licenses. The Storwize V7000 Remote Mirroring licensing would be as follows: three 5639-RM1 licenses for the primary system plus three 5639-RM1 licenses for the secondary system, for a total of six 5639-RM1 licenses required.

Scenario 3: There are three primary systems replicating to a central disaster recovery system. The three primary systems are as follows:

- A two-enclosure Storwize V7000 with nothing externally virtualized, therefore it has two 5639-VM1 licenses
- A one-enclosure Storwize V7000 managing a DS5020 consisting of three enclosures, therefore it has one 5639-VM1 license plus three 5639-EV1 licenses
- A one-enclosure Storwize V7000 with nothing externally virtualized, therefore it has one 5639-VM1 license.

The central disaster recovery system is a nine-enclosure Storwize V7000 with nothing externally virtualized, therefore it has nine 5639-VM1 licenses. The Storwize V7000 Remote Mirroring licensing would be as follows: a sum of seven 5639-RM1 licenses for the primary systems plus seven 5639-RM1 licenses for the central disaster recovery system, for a total of fourteen 5639-RM1 licenses required.

Added Interoperability

Storwize V7000 now supports many more heterogeneous data center environments with the addition of interoperability support for:

- EMC VNX models
- HP StorageWorks P9500 Disk Array
- Hitachi Data Systems Virtual Storage Platform
- Texas Memory Systems RamSan-620

For all the specific models and host environments supported, visit

http://www.ibm.com/storage/support/storwize/V7000

ISV Solutions

In support of the Storwize V7000 V6.2 release, a number of Storwize V7000 solutions have been tested and documented which leverage applications of major Independent Software Vendors (ISVs) including Microsoft®, VMware, Oracle, Symantec, and SAP. Key applications that are featured as part of this announcement include, but are not limited to VMware View 4.5 Reference Architecture, vSphere

Virtualized Solution, vStorage API for Array Integration (VAAI) features/benefits, Microsoft Exchange Solution Reviewed Program (ESRP) with up to 32K and 60K Mailboxes/Users, Oracle Database with Global Mirroring, and Symantec NetBackup.

In addition to the certifications and solutions indicated, IBM performs certifications with ISVs aligned with the various vertical industries including Healthcare, Financial Services, and Telco. Additional certifications and solutions are updated regularly. Please refer to the latest ISV white papers at:

http://www.ibm.com/systems/storage/solutions/isv/index.html

Auto renewal registration of software maintenance for Storwize V7000 5639-VM1 (Feature number 9000)

Software maintenance renewals in the AAS ordering system for Storwize V7000 5639-VM1 software are now available via a services contract. Selecting feature number 9000 with the software maintenance registration ordered with 5639-VM1 indicates that the customer wishes to continue their software maintenance on a services contract at the end of the initial software maintenance term. Feature number 9000 is available on the following Storwize V7000 base software maintenance registration product IDs:

- 5639-SM1 - 1 Year Software Maintenance Registration
- 5639-SM6 - 2 Year Software Maintenance Registration
- 5639-SM3 - 3 Year Software Maintenance Registration
- 5639-SM7 - 4 Year Software Maintenance Registration
- 5639-SM8 - 5 Year Software Maintenance Registration

Note: Auto renewal services are not available in all countries or for all initial software maintenance contract terms. Refer to the services announcement found in the Reference information section for more details.

45-Day migration allowance

With the benefit of external virtualization, the Storwize V7000 allows customers to bring a system into their storage environment and very quickly and easily migrate data from existing storage systems to the Storwize V7000. In order to facilitate this migration, IBM allows customers 45 days from the date of purchase of the Storwize V7000 system to use the external virtualization function for the purpose of migrating data from an existing storage system to the Storwize V7000 system. Any use thereafter, and ongoing use of the external virtualization function of Storwize V7000, requires the purchase of a Storwize V7000 External Virtualization license (5639-EV1) at a quantity equal to the capacity managed under the Storwize V7000 system. Migrations performed at later points in time that are to completely replace other storage systems with the Storwize V7000, thereby requiring temporary virtualization of that external storage system to perform that replacement activity, are granted a 45-day period for use of external virtualization without having to purchase a license to complete such a migration effort. Customers must make their sales representative aware of their intent and when they will be starting this migration so that an end date can be tracked. It is the customer's responsibility to ensure that they are properly licensed for all external storage managed by Storwize V7000 after those 45 days.

Reference information

For information on Storwize V7000 hardware machine type 2076, refer to Hardware Announcement AG11-0090, dated May 09, 2011, and Hardware Announcement AG10-0229, dated October 07, 2010.

For information on System Storage Productivity Center, refer to Hardware Announcement AG10-0266, dated November 23, 2010.

For information on the previous IBM Storwize V7000 Software V6.1.0, refer to Software Announcement AP10-0332, dated October 07, 2010.

For information on the software maintenance auto renewal registration services offering, refer to Services Announcement AS11-0010, dated April 05, 2011.

Availability of national languages

Planned and supported geography translations will be available for web download by June 24, 2011, at

http://www.ibm.com/storage/support/storwize/v7000

Program number

Program number	VRM	Program name
5639-VM1	6.2.0	IBM Storwize V7000 Software V6
5639-EV1	6.2.0	IBM Storwize V7000 External Virtualization Software V6
5639-RM1	6.2.0	IBM Storwize V7000 Remote Mirroring Software V6

Product identification number

PID numbers for IBM Storwize V7000 Software V6 are:

Program PID number	Maintenance 1 year PID number	Maintenance 3 year PID number
5639-VM1	5639-SM1 reg/ren 5639-SM2 afl	5639-SM3 reg 5639-SM4 renewal 5639-SM5 afl

Maintenance 2 year PID number	Maintenance 4 year PID number	Maintenance 5 year PID number
5639-SM6 reg	5639-SM7 reg	5639-SM8 reg

PID numbers for IBM Storwize V7000 External Virtualization Software V6:

Program PID number	Maintenance 1 year PID number	Maintenance 3 year PID number
5639-EV1	5639-VVA reg	5639-VVC reg

Maintenance 2 year PID number	Maintenance 4 year PID number	Maintenance 5 year PID number
5639-VVF reg	5639-VVG reg	5639-VVH reg

PID numbers for IBM Storwize V7000 Remote Mirroring Software V6:

Program PID number	Maintenance 1 year PID number	Maintenance 3 year PID number
5639-RM1	5639-MMA reg	5639-MMC reg

Maintenance 2 year PID number	Maintenance 4 year PID number	Maintenance 5 year PID number

5639-MMF reg 5639-MMG reg 5639-MMH reg

Publications

The following publications can be downloaded in PDF format from the websites that follow this list:

Title	Order number
IBM Storwize V7000 CIM Agent Developer's Guide	GC27-2292-01
IBM Storwize V7000 Read First	GC27-2293-00
IBM Storwize V7000 Quick Installation Guide	GC27-2290-01
IBM Storwize V7000 Troubleshooting, Recovery, and Maintenance Guide	GC27-2291-01
IBM System Storage SAN Volume Controller and Storwize V7000 Command-Line Interface User's Guide	GC27-2287-01
IBM System Storage Multipath Subsystem Device Driver User's Guide	GC52-1309
IBM Storwize V7000 Safety Notices	GC27-3924-00

Additional copies of the above listed publications will be available on June 10, 2011.

Use the following options to order or download additional copies:

- To view online information and link to a download site, visit the IBM Storwize V7000 Information Center

  http://publib.boulder.ibm.com/infocenter/storwize/ic/index.jsp
- To order additional selected hardcopy publications for a fee or download free PDF format publications, visit the IBM Publications Center website

  http://www.ibm.com/shop/publications/order
- To directly download publications in PDF format, visit the IBM product support website

  http://www.ibm.com/storage/support/storwize/v7000
- Contact your IBM representative.

The IBM Publications Center

Hardcopies of selected IBM Storwize V7000 publications are available for a fee from the IBM Publications Center, as well as free softcopy publications. The Publications Center is a worldwide central repository for IBM product publications and marketing material with a catalog of 70,000 items. Extensive search facilities are provided. Payment options for orders are via credit card or customer number. A large number of publications are available online in various file formats, and they can all be downloaded free of charge.

Publications may also be obtained on DVD via use of media feature 5809 (on 5639-VM1).

- IBM Storwize V7000 Customer Documentation DVD, SK5T-8813-01

Ordering information

Charge metric

Program name	Part number or PID number	Charge metric
IBM Storwize V7000 Base Software V6	5639-VM1	Per Storage Device
IBM Storwize V7000 External Virtualization Software V6	5639-EV1	Per Storage Device
IBM Storwize V7000 Remote Mirroring Software V6	5639-RM1	Per Storage Device

Product group: IBM Systems and Technology Group, System Storage Division

Product Identifier Description	(PID)
IBM Storwize V7000 Base Software V6 per Storage Device	5639-VM1
IBM Storwize V7000 External Virtualization Software V6 per Storage Device	5639-EV1
IBM Storwize V7000 Remote Mirroring Software V6 per Storage Device	5639-RM1

Software Maintenance

Program name: IBM Storwize V7000 Base Software Maintenance

One-year Registration and Renewal

PID: 5639-SM1

One-year No Charge Registration:

Description	Program number	Feature number
SWMA Auto Renewal Registration	5639-SM1	9000

Program name: IBM Storwize V7000 Base Software Maintenance

Three-year Registration

PID: 5639-SM3

Description	Program number	Feature number
SWMA Auto Renewal Registration	5639-SM3	9000

Program name: IBM Storwize V7000
Base
Software Maintenance

Two-year Registration

PID: 5639-SM6

Description	Program number	Feature number
SWMA Auto Renewal Registration	5639-SM6	9000

Program name: IBM Storwize V7000
Base
Software Maintenance

Four-year Registration

PID: 5639-SM7

Description	Program number	Feature number
SWMA Auto Renewal Registration	5639-SM7	9000

Program name: IBM Storwize V7000
Base
Software Maintenance

Five-year Registration

PID: 5639-SM8

Description	Program number	Feature number
SWMA Auto Renewal Registration	5639-SM8	9000

Terms and conditions

The information provided in this announcement letter is for reference and convenience purposes only. The terms and conditions that govern any transaction with IBM are contained in the applicable contract documents such as the IBM International Program License Agreement, IBM International Passport Advantage® Agreement, and the IBM Agreement for Acquisition of Software Maintenance.

Licensing

IBM International Program License Agreement including the License Information document and Proof of Entitlement (PoE) govern your use of the program. PoEs are required for all authorized use.

This software license includes Software Subscription and Support (also referred to as Software Maintenance).

Agreement for Acquisition of Software Maintenance

The IBM Agreement for Acquisition of Software Maintenance (Z125-6011) agreement applies for Subscription and Support (also referred to as Software Maintenance) and does not require customer signatures.

License Information form number

LC23-5059 License Information for IBM System Storage SAN Volume Controller 6.2.0 and IBM Storwize V7000 6.2.0

LC23-5098 License information for IBM System Storage SAN Volume Controller and IBM Storwize V7000 Information Center 6.2.0

On or near the planned availability date, the current License Information documents will be available for review on the IBM Software License Agreement website

http://www.ibm.com/software/sla/sladb.nsf

Limited warranty applies

Yes

Limited warranty

IBM warrants that when the program is used in the specified operating environment, it will conform to its specifications. The warranty applies only to the unmodified portion of the program. IBM does not warrant uninterrupted or error-free operation of the program or that IBM will correct all program defects. You are responsible for the results obtained from the use of the program.

IBM provides you with access to IBM databases containing information on known program defects, defect corrections, restrictions, and bypasses at no additional charge. For further information, consult the IBM Software Support Handbook found at

http://www.ibm.com/support/handbook

IBM will maintain this information for at least one year after the original licensee acquires the program (warranty period).

Money-back guarantee

If for any reason you are dissatisfied with the program and you are the original licensee, you may obtain a refund of the amount you paid for it, if within 30 days of your invoice date you return the program and its PoE to the party from whom you obtained it. If you downloaded the program, you may contact the party from whom you acquired it for instructions on how to obtain the refund.

For clarification, note that for programs acquired under any of IBM's On/Off Capacity on Demand (On/Off CoD) software offerings, this term does not apply since these offerings apply to programs already acquired and in use by you.

Other terms

Volume orders (IVO)

Yes. Contact your IBM representative.

Passport Advantage applies

No

Software Subscription and Support (Software Maintenance) applies

Yes. All distributed software licenses include Software Subscription and Support (also referred to as Software Maintenance) for a period of 12 months from the date of acquisition, providing a streamlined way to acquire IBM software and assure technical support coverage for all licenses. Extending coverage for a total of three years from date of acquisition may be elected.

While your Software Subscription and Support (Software Maintenance) is in effect, IBM provides you assistance for your routine, short duration installation and usage (how-to) questions, and code-related questions. IBM provides assistance via telephone and, if available, electronic access, only to your information systems (IS) technical support personnel during the normal business hours (published prime shift hours) of your IBM support center. (This assistance is not available to your end users.) IBM provides Severity 1 assistance 24 hours a day, every day of the year. For additional details, consult your IBM Software Support Handbook at

http://www.ibm.com/support/handbook

Software Subscription and Support (Software Maintenance) does not include assistance for the design and development of applications, your use of programs in other than their specified operating environment, or failures caused by products for which IBM is not responsible under this agreement.

For more information about the Passport Advantage Agreement, visit the Passport Advantage website at

http://www.ibm.com/software/passportadvantage

System i Software Maintenance applies

No

Variable charges apply

No

Educational allowance available

15% to qualified educational institution customers.

Prices

For all local charges, contact your IBM representative.

Passport Advantage

For Passport Advantage information and charges, contact your IBM representative or authorized IBM Business Partner for Software ValueNet®. Additional information is also available at

http://www.ibm.com/software/passportadvantage

AP distribution

Country/region	Announced
AP IOT	
ASEAN*	Yes
India/South Asia**	Yes
Australia	Yes
People's Republic of China	Yes
Hong Kong S.A.R of the PRC	Yes

```
 Macao S.A.R of the PRC      Yes
 Taiwan                      Yes
 Korea                       Yes
 New Zealand                 Yes
Japan IOT
 Japan                       Yes

 * Brunei Darussalam, Indonesia, Cambodia, Lao People's Democratic
   Republic, Malaysia, Philippines, Singapore, Thailand, and Vietnam
 **Bangladesh, Bhutan, India, Sri Lanka, Maldives, Nepal, and
   Afghanistan
```

Trademarks

IBM, System Storage, FlashCopy, Tivoli, Passport Advantage and ValueNet are registered trademarks of IBM Corporation in the United States, other countries, or both.

Microsoft is a registered trademark of Microsoft Corporation in the United States, other countries, or both.

Other company, product, and service names may be trademarks or service marks of others.

Terms of use

IBM products and services which are announced and available in your country can be ordered under the applicable standard agreements, terms, conditions, and prices in effect at the time. IBM reserves the right to modify or withdraw this announcement at any time without notice. This announcement is provided for your information only. Reference to other products in this announcement does not necessarily imply those products are announced, or intend to be announced, in your country. Additional terms of use are located at:

http://www.ibm.com/legal/us/en/

For the most current information regarding IBM products, consult your IBM representative or reseller, or visit the IBM worldwide contacts page

http://www.ibm.com/planetwide/

Using System Storage:
An Aussie Storage Blog
with Anthony Vandewerdt

VAAI and vCenter plugin for Storwize V7000 and SVC

anthonyv | May 10 | Tags: storwize v7000 svc | Comments (0) | Visits (2,183) Ratings 0

0

*** Updated 25/07/2011: The VAAI plugin can be downloaded from here: http://www-933.ibm.com/support/fixcentral/swg/selectFixes?parent=ibm~Storage_Disk&product=ibm/Storage_Disk/IBM+Storwize+V7000+(2076)&release=6.2&platform=All&function=all ***

The May 9 announcement that SVC and Storwize V7000 will support VAAI is very welcome news. The fundamental point is that the SVC and Storwize V7000 virtualise external storage. This means that the mountains of DS3000, DS4000, DS5000, AMS1000s, CX3s, etc, that are currently being virtualized behind these products, will inherit VAAI as soon as the virtualization layer supports it. This is yet another feature to add to the list of functions that IBM Storage virtualization can provide, such as: EasyTier; Thin Provisioning; multiple consistency groups; snapshots; remote mirroring; dynamic data relocation... the list goes on.

In addition we are releasing a plug-in for vCenter that enables VMware administrators to manage their SVC or Storwize V7000 from within the VMware management environment

Functions will include:

- Volume provisioning and resizing
- Displaying information about volumes
- Viewing general information about Storwize V7000 and SVC systems
- Receiving events and alerts for Storwize V7000 systems and SVC attached to vSphere
- The Storwize V7000 and SVC plug-in for vCenter will also supports virtualized external disk systems

The plug-in will be available at no charge on June 30 (for Version 6.1 software) and July 31 (Version 6.2). Here is a sneak peak of what it will look like:

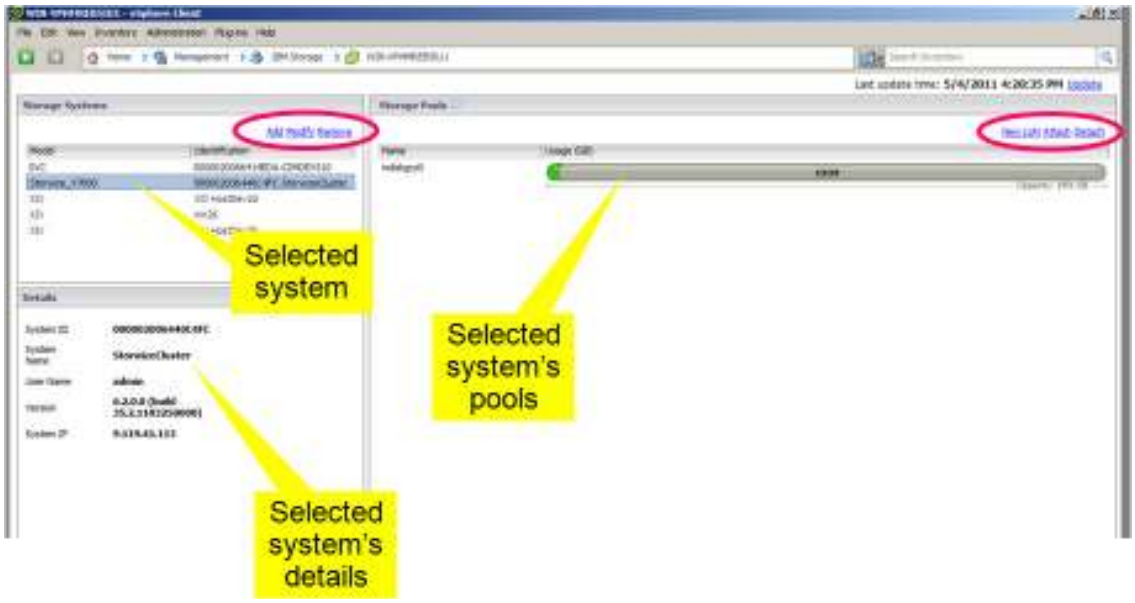

And to get an independent viewpoint have a read of Stephen Fosketts blog entry here:

http://blog.fosketts.net/2011/05/09/ibm-adds-vaai-support-xiv-svc/trackback/

DISCOVER RELATED RESOURCES TAGGED [storwize v7000 svc]

This Blog	All Blogs	Reddit	Del.icio.us	Twitter
storwize	storwize	storwize	storwize	storwize
v7000	v7000	v7000	v7000	v7000
svc	svc	svc	svc	svc

Add a Comment Recommend this Entry More Actions

COMMENTS [0]
There are no comments to display

Add a Comment

Feed for Blog Entries | Feed for Blog Comments | Feed for Comments for this Entry Previous | Main | Next

Product overviews >

IBM System Storage Archive Manager, Version 6.3

Compliance with data retention regulations

System Storage® Archive Manager is a software solution that provides long-term, retention protection when you archive business records, files, or data. The product can help you to comply with regulations for data retention.

In the following circumstances, consider System Storage Archive Manager for part of your storage management solution:

- You must meet regulatory compliance requirements to protect certain data from destruction for a specified length of time.
- You have internal controls that require that certain data must be archived for extended periods of time
- You must obtain approval before certain data can be destroyed.
- Your existing backup infrastructure is being overtaxed by long-term archival mandates and must be expanded or supplemented.

Archiving data to meet regulatory compliance requirements requires additional safeguards or protections, called *data retention protection*. These safeguards help guarantee that data is not deleted prematurely, either inadvertently or maliciously. To meet compliance requirements, System Storage Archive Manager provides additional protections for data archived by Tivoli® Storage Manager.

System Storage Archive Manager provides storage management services that permits users to archive files from their workstations or file servers to archive retention protected storage. Archived copies of files can be retrieved to local workstations. System Storage Archive Manager also includes an application programming interface (API) client program that you can use to enhance a content-management application with storage management services. When an application is registered with a server as a client node, the application can archive and retrieve objects from archive retention-protected storage.

System Storage Archive Manager uses chronological and event-based retention policies. *Chronological retention* is a calendar-based policy in which the final expiration countdown begins when an object is sent to System Storage Archive Manager storage. *Event-based retention* requires a predefined activation event to occur before the final expiration countdown starts. System Storage Archive Manager provides the ability to override prescribed retention policies using the deletion hold and release events.

Parent topic: Product overviews

Related concepts
Data retention protection
Deletion hold
Event-based policy retention protection

Related tasks
Setting up archive data retention protection
Suspending and allowing deletion and expiration of archive objects

Related reference
SET ARCHIVERETENTIONPROTECTION server command
Enablearchiveretentionprotection client option
Set Event client command

Feedback

Windows

IBM Tivoli Storage Manager, Version 6.3

Methods for backing up and recovering data on NAS file servers when using CIFS

Tivoli® Storage Manager can process network attached storage (NAS) file-server data that is accessed using the Common Internet File System (CIFS).

You can use the following methods to back up and recover data on NAS devices:

- Use a Tivoli Storage Manager backup-archive client to back up and restore data, by using CIFS or NFS to access files from the backup-archive client. Data can be stored on the Tivoli Storage Manager server with file-level granularity, by using the progressive-incremental backup method. The data is stored in the Tivoli Storage Manager storage hierarchy and can be migrated, reclaimed, and backed up to a copy storage pool.

  This method increases processor usage when the Tivoli Storage Manager client accesses individual files. The method requires that data flow through the Tivoli Storage Manager client. This method also requires that the data flows through the Tivoli Storage Manager server unless a LAN-free configuration is used.

- Use the snapdiff option to mitigate the performance problems of CIFS backup. This option stores data with file-level granularity using progressive incremental backup for NFS or CIFS.
- Use a Tivoli Storage Manager backup-archive client running on the NAS device, provided that you can use external programs with the NAS operating system.

  This method decreases processor usage of CIFS or NFS. Data can be stored on the Tivoli Storage Manager server with file-level granularity using progressive-incremental backup. The data is stored in the Tivoli Storage Manager storage hierarchy and can be migrated, reclaimed, and backed up to a copy storage pool. This method requires that data flow through the Tivoli Storage Manager client. This method also requires that the data flows over a network and through the Tivoli Storage Manager server unless a LAN-free configuration is used.

- Use NDMP with Tivoli Storage Manager. File systems are backed up as full images (all files) or differential images (all files that have changed since the last full backup). Backed up images are stored on a tape device that is accessed by the NAS file server. This method provides high performance because there is no data flow through a Tivoli Storage Manager client or server. Data that is backed up to Tivoli Storage Manager using NDMP cannot be migrated, reclaimed, or backed up to a copy storage pool.

These are some limitations of NAS file server data when using CIFS:

- Performance degradation occurs because data is being accessed remotely.
- The mapped drives appear to Tivoli Storage Manager as NTFS file systems, but they might not have the full functionality of NTFS functionality. For example, the encryption attribute of a file is set, but when the client backs up the file, the backup fails because the volume-level encryption setting indicates that encryption cannot be used for the volume.

  Tip: Use NDMP with Tivoli Storage Manager on a NAS file server to back up and restore volumes instead of backing up and restoring the volumes using remote mapped drives.

Parent topic: Mac OS X | AIX | HP-UX | Linux | Solaris | Windows Back up NAS file systems using Network Data Management Protocol

Related reference
Snapdiff

IBM System Storage Archive Manager Version 6.3 for Clients

Downloadable files

Abstract

This document describes how to download IBM® System Storage™ Archive Manager Version 6.3 for Clients, using the Passport Advantage® Online website.

Download Description

This document describes how to download IBM System Storage Archive Manager Version 6.3 client installation images using the Passport Advantage Online website.

IBM System Storage Archive Manager provides a client/server policy-based managed storage system to help meet regulatory and archive requirements for data retention and expiration. Archive and retrieval of retention managed data is through a System Storage Archive Manager client or the API client.

More information

To access a wide variety of technical resources for this product, see the Tivoli Storage Manager Information Center.

Note: There can be multiple clients in one download image.

Space requirements for AIX, 64-bit

Backup-archive client

- 170 MB of disk space to store the downloadable part.
- 200 MB of disk space to unpack the parts.

  NOTE: The AIX Runtime Library consumes approximately 1 MB when the part is unpacked. Installation of the Version Runtime Library is optional depending on whether the minimum required library level is already installed.

- 500 MB of disk space for the installed product.

HSM JFS2

- 420 MB of disk space to store the downloadable part.

Document information

IBM System Storage Archive Manager

Software version:
6.3

Operating system(s):
AIX, HP-UX, Linux, Linux on Power systems, Linux zSeries, Linux/x86, Mac OS, Solaris, Windows

Software edition:
All Editions

Reference #:
4030532

Modified date:
2011-10-21

- 460 MB of disk space to unpack the parts.

  NOTE: The AIX Runtime Library consumes approximately 1 MB when the part is unpacked. Installation of the Version Runtime Library is optional depending on whether the minimum required library level is already installed.

- 750 MB of disk space for the installed product.

HSM GPFS

- 450 MB of disk space to store the downloadable part.
- 490 MB of disk space to unpack the parts.

  NOTE: The AIX Runtime Library consumes approximately 1 MB when the part is unpacked. Installation of the Version Runtime Library is optional depending on whether the minimum required library level is already installed.

- 800 MB of disk space for the installed product.

Space requirements for Apple Macintosh

200 MB of disk space to store the downloadable part.
200 MB of disk space to unpack the parts.
400 MB of disk space for the installed product.

Space requirements for HP-UX Itanium®

350 MB of disk space to store the full set of downloadable parts.
550 MB of disk space to unpack the parts.
1000 MB of disk space for the installed product.

Space requirements for Linux®

x86_64 backup-archive client

- 200 MB of disk space to store the full set of downloadable parts.
- 200 MB of disk space to unpack the parts.
- 500 MB of disk space for the installed product.

x86_64 HSM GPFS client

- 300 MB of disk space to store the full set of downloadable parts.
- 300 MB of disk space to unpack the parts.
- 700 MB of disk space for the installed product.

System z®

- 80 MB of disk space to store the full set of downloadable parts.
- 80 MB of disk space to unpack the parts.
- 300 MB of disk space for the installed product.

on Power systems™

- 80 MB of disk space to store the full set of downloadable parts.
- 80 MB of disk space to unpack the parts.
- 300 MB of disk space for the installed product.

Space requirements for Solaris

Solaris SPARC®

- 200 MB of disk space to store the full set of downloadable parts.
- 300 MB of disk space to unpack the parts.

- 500 MB of disk space for the installed product.

Space requirements for Windows®

Note: These numbers are for the base client only and assume that all client components are installed.
Windows X32

- 175 MB of disk space to store the downloadable part.
- 370 MB of disk space to unpack the parts.
- 370 MB of disk space for the installed product.
- 370 MB of space in the tmp directory to run the installation program.

Windows X64

- 235 MB of disk space to store the downloadable part.
- 530 MB of disk space to unpack the parts.
- 530 MB of disk space for the installed product.
- 530 MB of space in the tmp directory to run the installation program.

Downloading and assembling the product

1. Download the part that you need. The parts are described by platform in the tabbed sections of this document.
2. Unpack each downloaded part into a single temporary directory on your system.
3. Follow the instructions in the IBM System Storage Archive Manager Quick Start Guide to install the product.

- AIX/Linux/UNIX Server
- Windows Server
- Client

Supported platforms

- AIX
- Apple Macintosh
- HP-UX
- Linux
- Solaris
- Windows

Downloadable part

Part number	Description
CI3KSML.tar.gz	IBM Tivoli Storage Manager V6.3, AIX
CI3KVML.tar.gz	IBM Tivoli Storage Manager V6.3, Apple Macintosh
CI3KTML.tar.gz	IBM Tivoli Storage Manager V6.3, HP-UX Itanium
CI3KUML.tar.gz	IBM Tivoli Storage Manager V6.3, Linux
CI3KWML.tar.gz	IBM Tivoli Storage Manager V6.3, Solaris
CI3KYML.exe	IBM Tivoli Storage Manager V6.3, Windows

Prerequisites

For system requirements, see IBM System Storage Archive Manager Supported Platforms.

Installation Instructions

Retrieve a copy of the System Storage Archive Manager Quick Start Guide.
The Quick Start Guide is available with the product and included on the Quick

Start DVD. The Quick Start Guide is also available for download in PDF format in all supported languages at the website listed below. Follow the instructions in the System Storage Archive Manager Quick Start Guide to install the product.

URL	LANGUAGE	SIZE(Bytes)
System Storage Archive Manager Quick Start Guide	English	147903

Download package

Download	RELEASE DATE	LANGUAGE	SIZE(Bytes)	Download Options
Log on to Passport Advantage Online	21 Oct 2011	English	21451	HTTP

Rate this page:

Rated by 0 users

Copyright and trademark information

IBM

IBM System Storage Archive Manager Version 6.3 for AIX/Linux/UNIX Servers

Downloadable files

Abstract

This document describes how to download IBM® System Storage™ Archive Manager Version 6.3 AIX®, Linux®, and UNIX® servers, using the Passport Advantage® Online website.

Download Description

This document describes how to download IBM System Storage Archive Manager Version 6.3 server installation images using the Passport Advantage Online website.

IBM System Storage Archive Manager provides a client/server policy-based managed storage system to help meet regulatory and archive requirements for data retention and expiration. Archive and retrieval of retention managed data is through a System Storage Archive Manager client or the API client.

More information

To access a wide variety of technical resources for this product, see the Tivoli Storage Manager Information Center.

Space requirements for AIX

- 2009 MB of disk space to store the downloadable part.
- 2776 MB of disk space to unpack the part.

  NOTE: The AIX Version 9.0.0.8 C Runtime Library consumes approximately 78 MB when the part is unpacked. Installation of the Version 9.0.0.8 Runtime Library is optional depending on whether the minimum required library level is already installed.

- 2156 MB of disk space for the installed product, which comprises 1838 MB in /opt, 297 MB in /usr, and 21 MB in /var.

Space requirements for HP-UX Itanium®

- 2168 MB of disk space to store the downloadable part.
- 2925 MB of disk space to unpack the part.
- 5510 MB of disk space for the installed product, which comprises 4569 MB in /opt, 929 MB in /usr, and 12 MB in /var.

Space requirements for Linux

Document information

IBM System Storage Archive Manager

Software version:
6.3

Operating system(s):
AIX, HP-UX, Linux, Linux on Power systems, Linux zSeries, Linux/x86, Solaris

Software edition:
All Editions

Reference #:
4030530

Modified date:
2011-10-21

- x86_64
  - 1656 MB of disk space to store the downloadable part.
  - 1656 MB of disk space to unpack the part.
  - 1684 MB disk space for the installed product, which comprises 1494 MB in /opt, 185 MB in /usr, and 5 MB in /var.
- on Power systems™
  - 1563 MB of disk space to store the downloadable part.
  - 1563 MB of disk space to unpack the part.
  - 1653 MB of disk space for the installed product, which comprises 1444 MB in /opt, 203 MB in /usr, and 6 MB in /var.
- System z®
  - 1486 MB of disk space to store the downloadable part.
  - 1730 MB of disk space to unpack the part.
  - 1356 MB disk space for the installed product, which comprises 1168 MB in /opt, 182 MB in /usr, and 6 MB in /var.

Space requirements for Solaris

- Solaris SPARC®
  - 1689 MB of disk space to store the downloadable part.
  - 2602 MB of disk space to unpack the part.
  - 1803 MB of disk space for the installed product, which comprises 1544 MB in /opt, 254 MB in /usr, and 5 MB in /var.

Downloading and assembling the product

Download the part that you need. The parts are described by platform in the tabbed sections of this document.
Unpack each downloaded part into a single temporary directory on your system.
Follow the instructions in the IBM System Storage Archive Manager Quick Start Guide to install the product.

- AIX/Linux/UNIX Server
- Windows Server
- Client

Supported platforms

- AIX
- HP-UX
- Linux
- Solaris

Downloadable part

Part number	Description
CI3K6ML.bin	IBM Tivoli Storage Manager V6.3, AIX
CI3K7ML.bin	IBM Tivoli Storage Manager V6.3, HP-UX
CI3KAML.bin	IBM Tivoli Storage Manager V6.3, Linux x86_64
CI3KBML.bin	IBM Tivoli Storage Manager V6.3, Linux on Power Systems
CI3KCML.bin	IBM Tivoli Storage Manager V6.3, Linux System z
CI3K8ML.bin	IBM Tivoli Storage Manager V6.3, Solaris SPARC

Optionable parts

Optionally, you can download any of the following parts:

Part number	Description
CI3KFML.bin	IBM Tivoli Storage Manager V6.3 Administration Center, AIX

CI3KHML.bin	IBM Tivoli Storage Manager V6.3 Administration Center, Linux x86
CI3KIML.bin	IBM Tivoli Storage Manager V6.3 Administration Center, Linux System z
CI3KGML.bin	IBM Tivoli Storage Manager V6.3 Administration Center, Solaris SPARC
CI3ZGML.bin	IBM Tivoli Monitoring for Tivoli Storage Manager V6.3, AIX
CI3KNML.tar	IBM Tivoli Monitoring for Tivoli Storage Manager V6.3, Linux x86
CI3KPML.bin	IBM Tivoli Monitoring for Tivoli Storage Manager V6.3, Linux x86_64

Prerequisites

For system requirements, see IBM System Storage Archive Manager Supported Platforms.

Installation Instructions

Retrieve a copy of the System Storage Archive Manager Quick Start Guide. The Quick Start Guide is available with the product and included on the Quick Start DVD. The Quick Start Guide is also available for download in PDF format in all supported languages at the website listed below. Follow the instructions in the System Storage Archive Manager Quick Start Guide to install the product.

URL	LANGUAGE	SIZE(Bytes)
System Storage Archive Manager Quick Start Guide	English	147903

Download package

Download	RELEASE DATE	LANGUAGE	SIZE(Bytes)	Download Options
Log on to Passport Advantage Online	21 Oct 2011	English	21451	HTTP

Rate this page:

Rated by 0 users

Copyright and trademark information

IBM

IBM System Storage Archive Manager Version 6.3 for Windows Servers

Downloadable files

Abstract

This document describes how to download IBM® System Storage™ Archive Manager Version 6.3 using the Passport Advantage® Online website.

Download Description

This document describes how to download IBM System Storage Archive Manager Version 6.3 Windows® server installation images using the Passport Advantage Online website.

IBM System Storage Archive Manager provides a client/server policy-based managed storage system to help meet regulatory and archive requirements for data retention and expiration. Archive and retrieval of retention managed data is through a System Storage Archive Manager client or the API client.

More information

To access a wide variety of technical resources for this product, see the Tivoli Storage Manager Information Center.

Space requirements for Windows, 64-bit

- 1288 MB of disk space to store the full set of downloadable parts.
- 1711 MB of disk space to unpack the parts.
- 1687 MB of disk space for the installed product.

Downloading and assembling the product

1. Download all of the parts that you need, both required and optional. The parts are described by platform in the tabbed sections of this document.
2. Unpack each of the parts into a single temporary directory on your system.
3. Use short path names for the directory used to extract the packages, for example: C:\tsm_images. Path names inside the package are long and extracting the package to a lengthy directory name can cause problems with the 255 character file path limit on Windows.

   Note: Ensure that the path does not contain non-ASCII characters.

4. Follow the instructions in the IBM System Storage Archive Manager Quick Start Guide to install the product.

Document information

IBM System Storage Archive Manager

Software version:
6.3

Operating system(s):
Windows, Windows 2008 server

Software edition:
All Editions

Reference #:
4030531

Modified date:
2011-10-21

- AIX/Linux/UNIX Server
- Windows Server
- Client

Supported platform

- Windows Server

Downloadable parts

There is one required downloadable part for Windows. Refer to the tables below for details about what each part contains. At a minimum, you must download the first part. The other part is optional.

Required parts

To install IBM Tivoli Storage Manager V6.3 on the Windows platform, you need to download the following part:

Part number	Description
CI3KDML.exe	IBM Tivoli Storage Manager V6.3, Windows, 64-bit

Optionable parts

Optionally, you can download any of the following parts:

Part number	Description
CI3KJML.exe	IBM Tivoli Storage Manager V6.3 Administration Center, 32-bit
CI3KKML.exe	IBM Tivoli Storage Manager V6.3 Administration Center, 64-bit
CI3KQML.exe	IBM Tivoli Monitoring for Tivoli Storage Manager V6.3, 32-bit
CI3KRML.exe	IBM Tivoli Monitoring for Tivoli Storage Manager V6.3, 64-bit

Prerequisites

For system requirements, see IBM System Storage Archive Manager Supported Platforms.

Installation Instructions

Retrieve a copy of the System Storage Archive Manager Quick Start Guide. The Quick Start Guide is available with the product and included on the Quick Start DVD. The Quick Start Guide is also available for download in PDF format in all supported languages at the website listed below. Follow the instructions in the System Storage Archive Manager Quick Start Guide to install the product.

URL	LANGUAGE	SIZE(Bytes)
System Storage Archive Manager Quick Start Guide	English	147903

Download package

Download	RELEASE DATE	LANGUAGE	SIZE(Bytes)	Download Options
Log on to Passport Advantage Online	21 Oct 2011	English	21451	HTTP

Rate this page:

Rated by 0 users

Copyright and trademark information

close window

IBMSystems magazine

Print

Easy Archiving

IBM has myriad solutions to make archiving and retention easier than expected

August 2008 | by Jim Utsler

Due to increasing regulatory pressures and other business considerations, more and more organizations are being pushed to archive and retain specific documents. Unfortunately, many of them do so using typical backup procedures, thinking that this is as good as any method to make sure they maintain important business content.

But dedicated archiving and retention solutions are much more effective, allowing users to keep only those documents and other information that specifically pertain to whatever regulations or business-specific requirements to which they must adhere. Rather than doing wholesale backups, organizations can cherry pick and store only what they need.

Not only does this method help ensure that companies are in compliance with regulations and other business-driven reasons for archiving and retention, but it also helps in other ways, including improving system performance, with those need-to-keep documents now being offloaded to low-cost disk and/or tape.

This may sound like a no-brainer, but some organizations aren't sure how to best meet their archiving and retention needs, including how long they should keep the documents or whether the documents could be kept secure enough to meet the toughest mandates. Now, however, thanks to some hardware and software solutions, those needs can be easily met, helping organizations automatically retain those key documents, apply whatever measures they need to meet regulatory compliance, and track the data throughout its lifespan.

A Huge Blow

Businesses need to keep documents and records handy in order to conduct business. But there are reasons why some organizations should "retain" specific documents. "Some do it to comply with industry or government regulatory mandates, others because of risk of litigation and yet others because they want to maintainÑand be able to refer back toÑcritical documents for future product or business development reasons," remarks Funda Eceral, archive solutions marketing manager, IBM System Storage*.

In some cases, organizations using archiving and retention technologies to meet regulatory mandates also use that information to support their legal departments. If a company gets sued, the company might have to provide certain legally binding documents. Not having them could result in the loss of a suit or other legal action taken against a company or even company officers.

"In the financial sector, the SEC demands that broker-dealers retain account-record information for six years. The six-year period begins either at the time the account is closed or

when the information is replaced or updated," Eceral says. "This includes even e-mail correspondences. If a company gets audited and can't produce those documents, they can be penalized for it. Usually, though, compliance officers are well-versed in the regulations and retention requirements particular to any given industry."

Another reason for some archiving activity is industry-specific data-retention requirements. For example, insurance companies have to keep signed, binding documents for some time, depending on the type of insurance that's being offered. In the case of life-insurance policies, customer-signed documents have to be kept at least until they're paid out and then for several years afterward.

In other cases, organizations may want to keep documents for reference. Auto manufacturers or aeronautical companies, for instance, often keep old blueprints and designs that might be useful for future product development and competitiveness. If the data isn't retained, it may well be lost to the ages, never to be seenÑor usedÑagain. "This can be a huge blow to continued development efforts," Eceral says.

The medical industry is another example of how important archiving and retention are. Documents such as MRIs and X-rays should be kept for a specified period of time. This allows medical professionals to easily access files as needed, especially if they're stored on a network-capable archiving and retention system, letting them share information with other medical service providers. National libraries worldwide are trying to preserve any and every document possible as they pertain to country's national heritage.

> *"As a leader in storage and enterprise-content management, IBM can help [companies] along the way so that they can manage their information growth, improve productivity and gain operational efficiencies." ÑFunda Eceral, archive solutions marketing manager, IBM System Storage*

A Keystroke Away

These are some business-driven examples of why organizations should archive documents. Operational aspects come into play, including the explosion in the amount of data being captured by companies. According to IDC, storage capacity shipments are growing at a compound annual growth rate of 54 percent from 2008-2012 (IDC, "Worldwide Disk Storage Systems 2008-2012 Forecast: Content-Centric Customers Reshaping Market Demand," Doc #212177, May 2008), with traditional structured data growing by 32 percent, unstructured data by 64 percent and, perhaps most importantly, replicated data by 44 percent (IDC, "Storage Infrastructure Innovations for the Future Data Center," Doc #DR2008_1RV, March 2008).

The results of these growth rates have an impact on everyday operational issues. Documents that should be archived often remain on production servers, slowing server and application performance and increasing backup windows, factors that sometimes make it impossible for organizations to meet service-level agreements.

That's in part why organizations must differentiate between regular backups, which are used for continuity purposes, and archiving and retention. In the former case, all data is copied,

typically to tape, so companies can recover accidentally deleted documents or to recover from some form of catastrophic disaster, whether because of a system failure or natural calamity.

In the latter case, only certain documents are moved to either disk or tapeÑor a combination of both. Although regularly scheduled backups should continue to take place, some of the information contained on backed-up servers can be permanently or temporarily offloaded to media other than production servers, thereby addressing some of the concerns related to the explosion of the IDC-noted data growth. As Eceral remarks, “Some organizations keep so much data on their main servers that everything slows considerably, including application performance and server and query response times.”

In most of the aforementioned business-related examples, only one copy of a document needs to be retained. If companies don’t take archiving-specific actions, and instead depend on backups, they maintain redundant copies of documents that must be stored only once, contributing to that 49-percent growth in replicated data.

Cost is a factor that should be considered when companies contemplate archiving and retention. Rather than keeping data on an expensive disk system, archived content can be moved to less expensive storage media, including nearline (SATA) disk systems and tape. A related benefit is that more expensive production system-disk space is now freed up for everyday processing requirements, with the timeliest documents and data being stored locally; only a keystroke away.

A Blended Solution

No matter how data is archived and retained, certain specifications must be considered. These include whether the data should be encrypted or be classified as non-erasable or non-rewritable. Again, this depends on the industry. In the finance world, for example, the SEC Rule 17a-4(f) dictates that records should be preserved exclusively in a non-erasable, non-rewritable format, without the possibility of anyone altering, erasing or overwriting the contents.

Solutions are available that make archiving and retention easier than expected. This is true given the intelligent nature of some of the software that underpins archiving and retention efforts. These include such offerings as, on the data-content layer, IBM CommonStore, IBM Content Manager, IBM FileNet*, IBM Optim and IBM Grid Medical Archive Solution; and, on the policy-management layer, Tivoli* Storage Manager, System Storage Archive Manager (SSAM) and Data Facility Storage Management System (DFSMS*). SSAM, a specialized version of Tivoli Storage Manager, provides storage controls to keep data in non-erasable, non-rewritable form, and it offers automated migration services, as well as chronological, event-based retention, deletion hold and policy-based data movement through storage tiers.

Many capable solutions exist on the hardware end. These include the System Storage DR550 and N Series with SnapLock providing non-erasable, non-rewritable storage repositories; IBM disk and tape systems; and write once, read many (WORM) 3996 Optical Libraries. The award-winning DR550 is particularly well suited to archiving and retention, including for both structured and unstructured data. In addition to being tightly integrated with SSAM, it also supports disk and tape encryption options.

As this list indicates, disk and tape are viable options for archiving and retention. However, a combination of the two media is seen as the most viable solution for several reasons. When documents are put to disk, they’re more accessible than they would be on tape. Conversely, tape is cheaper and lasts longer than disk. Disk might be used for documents accessed more

often and tape for those that aren't.

Organizations can move documents from disk to tape after a specified period of time, with this blended solution can help them realize, according to an internal IBM study ("IBM TCO study for System Storage DR550," 2007Ñ10-year example, which assumes 250 TB storage at 25-percent growth per year), up to a 50-percent lower total cost of ownership than by going with disk alone. Given that most industries require that documents be kept for the long-term, blended solutions make sense.

To support businesses with their archiving needs, IBM established a first in the industryÑthe Global Archive Solutions Center, in Guadalajara, Mexico, which was developed to present and demo archiving and retention solutions. The solutions available there will include e-mail archiving, database archiving, file/image archiving with IBM Enterprise Content Management, Tivoli Storage Manager, Optim and DR550 storage systems.

Primary Objectives and By-Products

Whatever the reasons organizations must archive and retain data, it's critical that they find the right solutions that meet their particular requirementsÑand not rely on scheduled backups. Once they do, they can rest assured that they're meeting their obligations, whether to meet regulatory guidelines, trend to industry standards or retain critical business information. As a by-product, they'll also improve system and application performance, optimize their storage systems and improve their cost postures.

As Eceral puts it, "When companies accept the reality that they may need to archive and retain business information, they'll find many, many tools and services that will help them do it. As a leader in storage and enterprise-content management, IBM can help them along the way so that they can manage their information growth, improve productivity and gain operational efficiencies."

White Paper

IBM's Information Archive

Addressing Long Term Information Retention Requirements

By Brian Babineau

January, 2010

Contents

Introduction

Organizations retain information for two reasons: because they want to (it has some inherent business value, business intelligence, analytics, reporting, etc.) or because they have to (due to regulatory requirements, legal discovery, etc.). These issues are exacerbated by data retention due to a lack of formal ways to determine what is or is not needed in addition to a lack of enforcement of consistent deletion policies. As a result, it is hard to find a business that doesn't have a reason to save information. On the other hand, it is easy to find one experiencing difficulties with reigning in information management costs.

The process of retaining information, more commonly referred to as "archiving," is challenging for several reasons:

- Companies continue to generate new information, at least a subset of which the business wants or has to save.
- Information is generated by multiple applications in multiple formats. E-mail, instant messages, productivity files, database records, SharePoint sites, wikis, and videos are just a few examples of information types that may need to be archived.
- Because of the importance of the information being archived, organizations have to treat it as "mission critical" and put proper data protection and business continuity processes in place to prevent downtime or data loss.
- Companies have to secure archives in accordance with industry- or government-specific privacy regulations.
- Archived information is kept for a long time. ESG research suggests that more organizations are saving archived information between six and ten years than for any other period of time (see Figure 1).[1] With most IT departments viewing the world of technology in three year increments, cost effectively storing and managing archived information requires an entirely new perspective.

Figure 1. Average Archive Retention Period By Content Type (E-mail, Database, File)

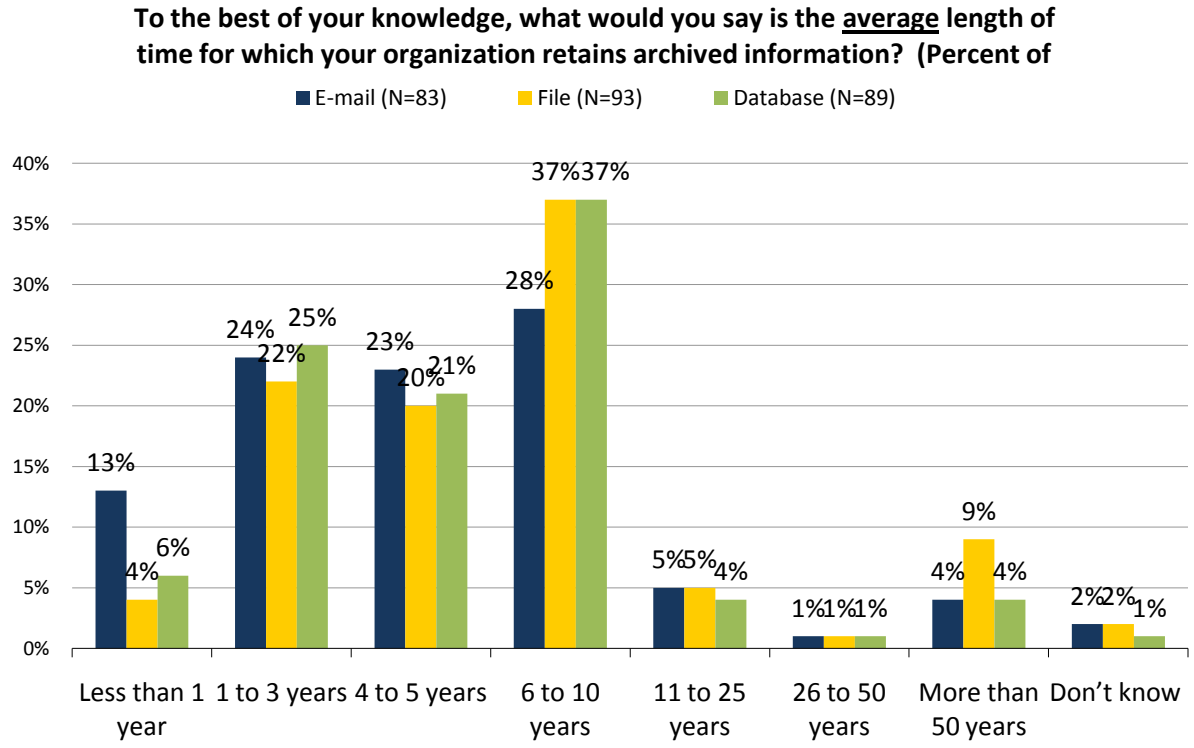

Source: Enterprise Strategy Group, 2007.

[1] Source: ESG Research Report, *E-mail, Database, and File Archiving Surveys*, November 2007.

Organizations can tackle a few of these issues immediately by deploying information-specific archiving solutions with the storage of their choosing. However, cost and complexity arise when another information type has to be archived and, due to legal preservation reasons, needs to be saved on immutable storage for a period that differs from the original information type. These individual archive stovepipes may be supported by a number of technology vendors, which can lead to management complexities, higher cost, and uncontrollable risk. A short-term perspective on information archiving can easily create unnecessary expenses and operational inefficiencies.

IBM's Smart Archive strategy, bolstered by a series of hardware, software, and services offerings, is designed to help customers transition from short-term archiving decisions and their associated risks to a longer-term, simpler approach to information retention. The offerings enable organizations to unify several aspects of the archive process (including the collection, analysis, management, retention, storage, and access of information) while providing a variety of technology consumption models inclusive of integrated appliances, managed services, and cloud-based services to help archiving fit into existing business and IT operating procedures—and not the other way around.

Archiving strategies such as IBM's allow businesses to implement individual products, safe in the knowledge that there are supporting and complementary offerings available for future deployments. Because archiving involves saving information for extended periods of time, the storage infrastructure is a logical place to start: the right storage infrastructure will help reduce costs resulting from the challenges outlined above. This paper details these costs, as well as how IBM's Information Archive—an integrated archive storage appliance and a cornerstone of the Smart Archive strategy—addresses them.

Information Retention Realities

Multiple Business Drivers

To understand storage-specific archive costs, organizations must first understand why they are saving information. ESG groups archiving drivers into four major categories:

- ***Compliance***. Almost every business, regardless of industry or geographic location, is subject to record retention regulations. Record retention processes must transition to incorporate digital information, but doing so also introduces ambiguity. Most rules mandate which records must be saved and for how long, but they do not specify how to do it or which technology to use to meet requirements—yet failure to comply with record retention requirements can result in extensive fines, severe damage to the business reputation, and business dissolution.
- ***Corporate governance***. Corporate boards and executives often define policies and practices separate from legal compliance to ensure accountability, proper reporting, and fulfillment of their responsibilities to shareholders and stakeholders. Information retention processes are frequently included in corporate governance initiatives to demonstrate diligence and fiduciary responsibility.
- ***Electronic discovery***. There are two retention catalysts when it comes to saving information for legal purposes. The first occurs when information is relevant to a specific legal or regulatory matter and must be retained and preserved until the matter is resolved. The second involves corporate policies that prevent the deletion of certain information in the event that it is needed to support or defend a claim. These mandates often lead to everything being saved, forever.
- ***Business reference.*** Organizations recognize the value of institutional knowledge and strive to retain information about products, processes, customers, programs, etc. for future use. To be useful for business reference, information also needs to be appropriately cataloged and classified to facilitate easy searching and retrieval. The more information available to employees, the better their decisions will be. In this situation, many companies are simply trying to figure out how to keep more information accessible at lower costs. One of the easiest ways to do so is by archiving or moving the content from the primary application infrastructure to a lower cost one. Companies are also trying to catalog or index as well as classify information to facilitate simple search and retrieval.

Dealing with Unique Circumstances

Due to various business drivers, information retention requirements—and thus storage needs—can differ widely. As an example, companies often have to archive several different information types from a variety of sources. This occurs regularly in electronic discovery: companies are asked to identify, preserve, and produce any information relevant to a specific topic or date range. Potential evidence may take the form of e-mail, video, or a database (see Figure 2).[2] An archive storage system must be flexible enough to support all these content types: structured (database, unstructured (files), and semi-structured (e-mail).

Figure 2. Frequently Requested Record Types During Electronic Discovery Processes

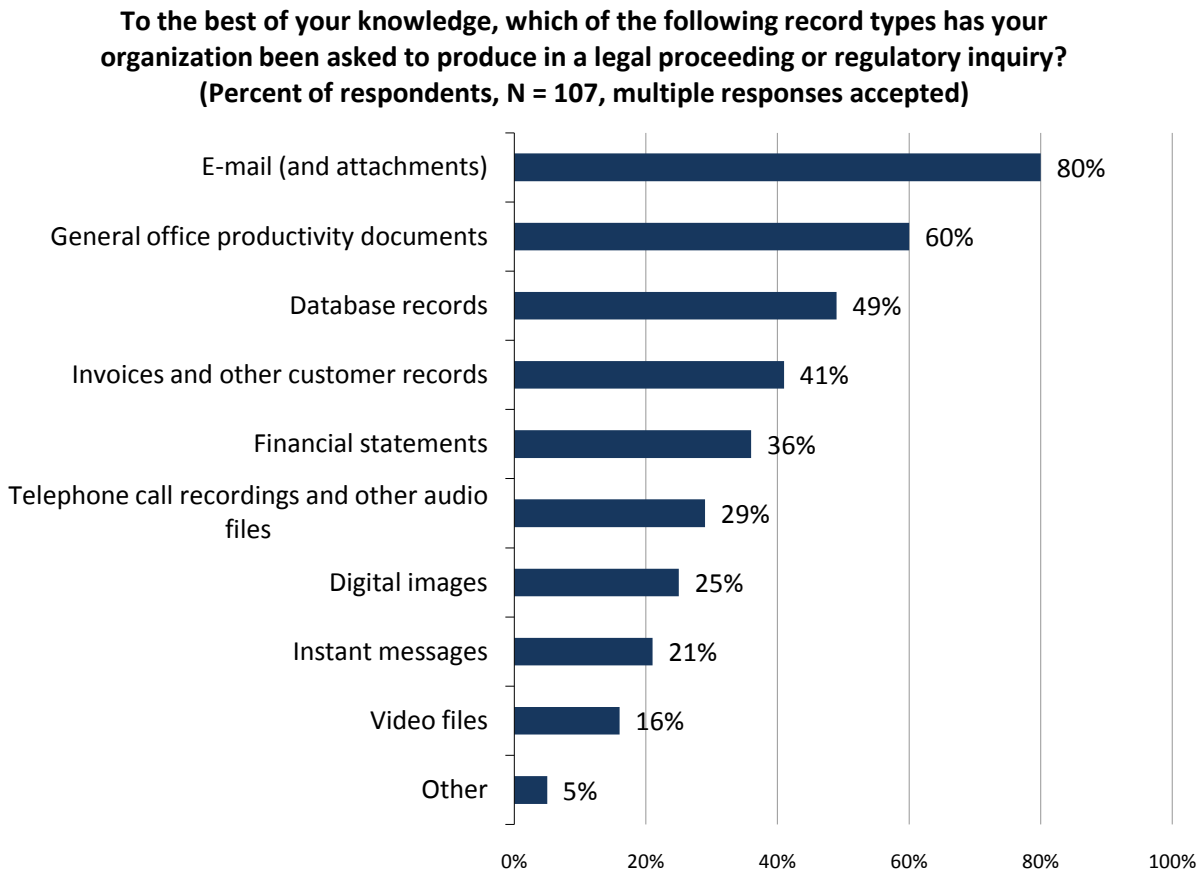

Source: Enterprise Strategy Group, 2007.

Organizations must also cope with a myriad of retention policies and parameters, including government- and industry-specific rules requiring different records be kept for varying lengths of time. When dealing with multiple electronic discovery matters, there is no formal timetable for legal holds, further complicating retention policies. After figuring out what information to keep and for how long, organizations must then decide what to do when a retention policy does expire: delete the information permanently or keep it?

Some archived information must be stored in such a way that allows its authenticity and integrity to be proven. This occurs frequently when information is archived for compliance and electronic discovery purposes and can be facilitated by a storage system that supports encryption and immutability (Write Once, Read Many data formatting) configurations. Other content types may not have such strict requirements, but companies choose to leverage these formatting and security options in the interest of data protection. There are, of course, many instances where archived data just needs to be kept somewhere.

[2] Source: ESG Research Report, *Electronic Discovery Requirements Escalate*, November 2007.

The biggest potential obstacle organizations face when archiving information is balancing accessibility requirements with storage costs. The conundrum exists because certain situations, such as electronic discovery, require archived data to be more accessible than business records that are in their 20th year of a 30 year retention policy. From a compliance and corporate governance perspective, companies may choose to keep recent records more accessible than those created a few years ago to support internal and external audit processes. More commonly, data saved for business reference is usually kept reasonably accessible. The issue is that the more accessible the storage media, the more it costs. From an archive perspective, "accessibility" usually leads to a comparison between disk and tape. Disk is typically faster, especially when application access is random, but tape is less expensive. Companies need to figure out which information is best for each media type and how long it should be kept there.

Costs Abound

Companies planning to save more information for longer periods of time have to budget for more storage capacity. However, acquisition costs are not the only expenses involved in implementing an archive. Because they contain valuable information—in some cases, vital business records and evidence—archives need to be protected and require backup and recovery systems. If archives support mission critical processes, organizations should also account for disaster recovery solutions to mitigate the risk of downtime.

There are several operational expenses associated with an archive storage environment, including:

- **Power to run and cool the storage systems**. Depending on a data center's power constraints, this could be a fairly sizeable expense. Adding new devices can be disruptive if a company must procure additional energy or reorganize the facility to optimize airflow.
- **Management (labor).** IT has to find ways to set up, provision, configure, and operate the archive storage infrastructure, including the data protection and disaster recovery systems. In archive environments, resources are often needed to constantly add more capacity due to data growth, customize interfaces so that a system can support a specific type of content that has to be saved, and update configurations to address various retention period requirements.
- **Migrations**. Most organizations include migration expenses in their labor costs, but archive environments involve larger (due to the amount of data) and more frequent migrations. These migrations are driven by retention periods that extend beyond the useful life (or the warranty) of the underlying storage system. Additionally, some IT departments proactively move data between systems—migrating older data from disk to tape, freeing capacity for fresh content. Any form of data movement can be extremely resource intensive as IT has to set up the target and source systems and then monitor the progress to ensure that no data is lost during the migration. Losing data creates unnecessary legal and compliance risks.

IBM's Information Archive Creates a Robust Foundation

Overview

IBM's Information Archive is a new unified and secure storage repository for archived information. It is a critical component of the IBM Smart Archive, which is a comprehensive, unified, integrated, and information-aware *archiving strategy* from IBM. IBM Information Archive is a factory integrated storage appliance capable of storing up to 304 TB of raw data on disk. With disk drive capacities constantly increasing, ESG expects this number to increase substantially over the next 12 to 18 months. In addition to the disk, Information Archive also supports hundreds of IBM and non-IBM tape systems, which increases the solution's logical capacity into the petabyte range. Customers can start small and add capacity—disk, tape, or both—as their archive storage needs dictate. This scalable approach enables customers to consolidate their archive storage rather than constantly purchase new devices as capacity needs increase.

With NFS and System Storage Archive Manager (SSAM, a variant of Tivoli Storage Manager) interfaces, Information Archive can connect to a variety of data sources simultaneously, including primary applications, purpose-built archive applications including IBM's Content Collector and Optim Data Growth solutions, and general purpose

production file shares. Flexible interfaces, specifically NFS, allow customers to archive multiple content types in a single system, eliminating the need for a different archive storage solution for every data source. The Information Archive's versatile architecture, leveraging IBM's General Parallel File System (GPFS) technology, leads ESG to believe that IBM will be able to support other interfaces such as CIFS and HTTP, which will only extend the solution's reach in the future.

"Collections" Concept Extend Flexibility

To date, the market has been wary of large-scale, consolidated storage archives because of the various configuration parameters needed to support different business drivers and content types. While any company could fill up a consolidated archive, not all the data needs to be saved and managed the same way. For example, a portion of archived information may need to be kept on non-erasable, non-rewritable media while other data simply needs to be stored for a few years in a reasonably accessible format for auditing purposes. As a result, companies preferred to implement several purpose-built archive storage systems with different configurations as well as a multitude of performance and availability characteristics to address each of their information retention requirements (compliance, electronic discovery, governance, etc.) and content types. This stovepipe archive storage strategy creates more issues—including poor utilization and management complexity—than it ultimately solves.

To address the stovepipe issue, Information Archive leverages "collections," which are virtual repositories within a single appliance. In its first release, it supports up to three collections per appliance and each collection can be customized to address a specific archive need. Customers may elect to set up collections based on archive business drivers to handle different content types and data sources or some combination thereof. There is little risk in filling up a collection as each of these virtual repositories can handle an estimated one billion files or "objects" (an object is the combination of a file and its metadata). With the Information Archive, customers get multiple "virtual" archives while managing one storage system.

Three Retention Policy Enforcement Choices

After setting up the collections, customers can select one of three data protection levels to best address their archive business requirements:

- **Basic** enables applications and users to delete documents before the retention period or retention hold expires. Users can also increase or decrease a retention policy and modify the collection's protection level if warranted. This option may be used when information is archived for general business reference purposes or to meet broad corporate governance requirements.
- **Intermediate** allows users to increase or decrease retention periods, but prevents deletion before the expiration of the retention period or the retention hold requirements. Users can increase the protection level to "Maximum," but cannot decrease it to "Basic." Customers looking to save important project information or enforce strict corporate governance requirements are likely to use this protection level.
- **Maximum** only allows users to increase retention periods and prevents information from being deleted until the retention period or retention hold has expired. This protection level cannot be altered and is ideal for companies that have to archive data to meet strict regulatory compliance requirements. This protection level may be compared to storage systems configured in WORM format.

Multiple Retention Options

With each collection having its own protection level, customers must then determine what data goes in the respective virtual repositories and how long to keep that information. This is determined by retention policies, which are also configured on a per collection basis. A retention policy can be a retention period determining how long a piece of information is to be kept. When customers know how long they want to save information (which is often the case when archiving for compliance reasons), they can leverage a "time-based" retention period. For example, a company may be forced to save all client communications including e-mail and instant messages for three years from the date of their creation. In situations where the ending time period is unknown, Information Archive customers can leverage "event-based" retention periods where information is kept until a particular occurrence. For example, employee contract information might be kept only as long as the employee remains with the company.

Another example of a retention policy is a retention hold, which prevents any information from being expired even if its retention period has ended. This option is most commonly used in electronic discovery situations where content must be preserved until the matter is resolved, even if the information is no longer needed to satisfy compliance, governance, or business reference requirements.

When configuring a retention policy within a collection, a customer needs to establish criteria used to identify how content should be managed and the actual retention period or hold requirement. This is based on file metadata properties such as the owner, date of creation, and format. As an example, a company may want to retain all Excel spreadsheets created by the finance organization for one year, any images generated by marketing for two years, and put executive employment contracts on retention hold because it is part of an ongoing legal matter.

As content is moved into a collection, it is analyzed and assigned a retention policy. Fine-grained retention policy management enables customers to set retention policies on a per file basis while saving the data in a single virtual repository. Information Archive also accepts and executes retention policies that are set by applications such as e-mail archiving, records management, and enterprise content management systems connected to it.

Controlling Cost and Risk

Data Reduction

One of the easiest ways to reduce archive related storage costs is to save less information. While this may sound counter-intuitive given the amount of information most companies have and want to save, it is entirely possible via Information Archive thanks to the appliance's deduplication and compression capabilities. Deduplication ensures that the same bytes of information are stored only once (reference pointers are used to track any redundant bytes sent to the system) while compression shrinks the size of a given set of bytes. Together, they pare down the amount of physical capacity consumed by archive data, reducing storage acquisition and operating costs.

Deduplication and compression capabilities are also configured on a per collection basis, providing customers a choice in how to maximize storage utilization. This is extremely useful in situations where one of the information reduction options doesn't align with the data type being archived. Some databases compress very well, but do not contain a significant amount of redundant data. Alternatively, attorneys may take a conservative approach in saving all information, choosing to leverage compression and not deduplication.

Multiple Storage Tiers

Information Archive also controls storage costs by supporting disk and tape within the same system and automating the movement of data between the two. Within the policy engine, customers establish criteria determining when they want to move content from disk to tape or vice versa. Criteria can include data growth triggers: when a collection reaches a certain size, older data is moved to tape freeing up capacity on disk for new information. Some customers may simply move all data based on when it was last accessed.

Moving older data to tape means that customers defer disk purchases, reducing archive storage acquisition costs. Policy-based migration also minimizes the burden on IT resources that often have to copy archives between storage media options to optimize storage resources.

Security is Imperative

Information Archive includes role-based security access, audit logs, encryption, and a locking cabinet for physical security. Its Enhanced Tamper Protection capability provides an additional level of security for data that absolutely cannot be altered or deleted during its retention period. When this feature is enabled, neither the customer nor IBM has root login authority. Common administrative and support operations are pre-programmed to remove the requirement for root access. Once turned on, the Enhanced Tamper Protection capability cannot be disabled.

Together, these security features mitigate the chance of unauthorized access and data tampering while audit logs track configuration information as well as any actions taken within the system. Customers can leverage the audit log data to prove that no one intentionally or maliciously tried to alter data without the appropriate permissions. They can also be used to track the success of collections and retention policy enforcement, which is imperative when companies need to demonstrate the authenticity and integrity of archived information often needed in regulatory and discovery situations.

The combination of the security features with "Maximum" Data Protection level make Information Archive an ideal target for compliance data that cannot be deleted during a retention policy. The system stores data in WORM format, prevents unauthorized access and tracks any attempts by users to change or delete information (which, of course, would be unsuccessful given how data is stored). To further demonstrate these capabilities, Information Archive was recently audited by a third party, Cohasset Associates, a consulting firm specializing in records management. The findings of the audit determined that Information Archive capabilities were satisfactory to address SEC Rule 17a-4(f) requirements—a regulation widely known as having the most stringent storage media requirements pertaining to the retention of business records. In short, this regulation requires broker/dealers to store requisite business records in non-erasable, non-rewritable format for specified (in SEC Rule 17a-3) retention period. And, several interpretations of this regulation have stated that electronic storage media as well as combination of electronic storage media and software that store information in WORM format can be used to meet the requirement.

Note that the ultimate determination of compliance falls within the customer's responsibility to notify its regulatory examining authority when they change the storage media which stores their business records. The examining authority may choose to question and inspect the media or not respond at all—the latter of which is usually interpreted as an "acceptance" of the solution. If the customer has to produce any of the business records stored on the system, they will need to prove and defend the means in which these records are stored and preserved on non-erasable, non-rewritable storage media.

Availability is Paramount

Similar to other storage systems, Information Archive includes redundant components and RAID protection to ensure high availability and data protection within the system. To prevent data loss due to corruption, customers can back up the archived data using traditional file system methods including NDMP as well as copying the data directly to tape using the Tivoli Storage Manager component. Additionally, the Information Archive database—which contains all of the metadata and file mapping information—can be protected with standard onsite and offsite backup policies.

For disaster recovery, Information Archive supports enhanced remote mirroring, enabling two real-time synchronous or asynchronous data copies to be maintained on separate Information Archive devices. Like many other Information Archive capabilities, data replication capabilities are configured on a per collection basis. If one collection contains critical data used to support compliance or discovery processes, customers do not have to replicate the entire system, saving storage costs and consuming less bandwidth.

Simplified Management and Operations

Customers may experience the greatest cost savings through Information Archive's management efficiencies. The first set of management efficiencies comes from ease of use and set up of the appliance itself. Wizards guide the user to create and configure collections as well as establish protection levels and retention and data management policies. Monitoring and troubleshooting tasks are integrated in a single interface, enabling IT to easily track system health.

The second set of management efficiencies is the direct result of Information Archive's ability to consolidate archive storage into a single system while addressing a myriad of archiving business requirements. When customers think about protecting their archives from downtime due to a disaster, they only have to set up replication on one system. Retention policies are set up in a central location as opposed to several devices with different protection levels. Floor space consumption is kept to a minimum and power draw can be minimized when customers deploy Information Archive with tape.

A Good Start to a Smart Archive

A New Approach

Organizations can continue to save information the way they always have: save everything forever, keep it on tape, and hope that no one ever has to access it. While this approach may sound simple, it is extremely complex and cumbersome as information retention requirements change. IBM's Smart Archive strategy is designed to help companies maintain a simple approach to information archiving while maintaining the flexibility to address whatever information archive challenges exist today and what may arise in the future.

In order to keep things simple, IBM's Smart Archive strategy includes products (hardware and software), services (consulting, implementation, etc.), and consumption models (appliances, cloud, etc.) so that customers can figure out what combination is best suited to meet their information retention needs. As an example, customers may turn to any number of IBM's Information Management software offerings to unify the identification, analysis, classification, and management of several information sources. In turn, these software solutions can send data into the Information Archive for long term, efficient retention.

There are several other examples of how IBM's different products and services can be used together to create a smart archive. However, customers should not lose sight of the fact that the whole point of IBM's Smart Archive is to do things differently in a more cost effective and more risk adverse manner than they are today.

Moving Forward Faster

Information Archive is an ideal starting place for a customer taking a new approach to archiving as it encompasses many of the facets of IBM's Smart Archive strategy. When a customer invests in an Information Archive, IBM Global Technology Services assist with installation and set up, accelerating the "time to archive." Often, archive projects get bogged down when customers realize they do not have the resources to properly connect data sources to the underlying storage system or spend the time to create retention policies. IBM helps Information Archive customers overcome these objectives and, for those organizations with existing archives, the services engagement can be extended to include data migration services.

IBM's services portfolio also includes managed services options for companies lacking the internal staff to run the archive environment. In some cases, customers may not be able to handle the explosive storage growth and associated data protection functions needed to properly manage an archive environment. Others may want to turn management of the entire archive infrastructure, including an Information Archive system and the applications which connect to it, over to IBM so they can focus limited resources on primary IT operations.

Enabling Access and Expiration

Saving the right—not all—information is a critical component of IBM's Smart Archive Strategy, as is expiring data when it is no longer needed for compliance, governance, legal, and business reference purposes. Determining what data to keep is a process handled by archive, content management, and other information management software applications. Information Archive helps customers expire content in a consistent manner through one its three protection levels assigned to a given collection. When customers begin to delete information, the storage savings exponentially increase because capacity is suddenly freed for new data.

Before data is expired, Information Archive facilitates secure information access by maintaining a clustered index of all the data stored within it. Applications and users can execute searches against the index, with results being displayed based on the user or application's security privileges. Those privileges may dictate if a user or application can view results across the entire Information Archive appliance or a specific collection. The index is always stored on disk (even if the data is moved to tape by the Information Archive), ensuring fast query response times even as archive storage capacity increases.

Heading for the Cloud

There is no bigger topic in IT right now than "cloud." While companies try to figure out their strategies in leveraging "public" or external cloud-based services, many are observing cloud-provider IT architectures using the insight to build internal or "private clouds." When examining what characteristics define a cloud architecture—flexibility, ease of provisioning, and real-time scalability—it is evident that IBM's Information Archive can be the cornerstone of a private archive cloud. Customers can connect most data sources to it, there are multiple "provisioning" options including setting up a new collection to address a new business requirement dictating the retention of information, and the appliance scales by adding capacity to the existing system or integrating another tier of and seamlessly migrating data to it.

When a customer is ready to leverage a public cloud, Information Archive can be deployed at a trusted third party service provider including IBM. A customer may do this to shift the day-to-day archive storage management to the service provider. In some cases, this transition may be triggered by shifting the software connected to the Information Archive from an on-premise deployment to a cloud or Software as a Service model. As an example, IBM's Content Collector software is already cloud-enabled which gives customers a choice in how they want to implement and consume archiving software. By having an architecture that supports cloud implementation characteristics, Information Archive provides customers with flexible deployment options.

The Bigger Truth

Archiving used to involve saving information for extended periods of time. Now, there are many more variables involved in archiving processes as companies have to deal with multiple information types, different access requirements, evolving retention policies, and lengthening retention periods. All of these impact archive storage needs and costs.

Companies could continue saving everything forever on the same storage media or buying separate storage systems designed to address individual archive requirements. However, these alternatives can get very expensive over the long term. With its "collections"-based architecture, IBM's Information Archive is flexible enough to address different requirements in a single system solving the stovepipe archive storage issue. And, whether you measure scalability by capacity (which Information Archive optimizes through information reduction features), resource management (which Information Archive has covered by supporting over 300 terabytes of disk and petabytes of tape), or number of objects (which Information Archive can handle up to three billion of in its first release), it is clear that customers can grow their archives with ease. Information Archive is also optimized to protect and secure archived information over the long term with capabilities such as Enhanced Tamper Protection.

As part of IBM's Smart Archive strategy, IBM's Information Archive provides customers with a variety of ways to solve evolving information retention challenges—an approach that has the potential to dramatically reduce capital and operating costs even as archive capacity explodes.

Enterprise Strategy Group | **Getting to the bigger truth.**

20 Asylum Street | Milford, MA 01757 | Tel:508.482.0188 Fax: 508.482.0218 | www.enterprisestrategygroup.com

TSW03051-USEN-00