Redbooks
ibm.com/redbooks

IBM System Storage Solutions Handbook

Ezgi Coskun
Mikael Lindström
Maciej Olejniczak
Oliver Stark
Megan Gilge

Cloud

Storage

IBM

Redbooks

International Technical Support Organization

IBM System Storage Solutions Handbook

July 2016

SG24-5250-11

Note: Before using this information and the product it supports, read the information in “Notices” on page vii.

Eleventh Edition (July 2016)

This edition applies to the IBM System Storage products as of April 2016.

Contents

Notices

This information was developed for products and services offered in the US. This material might be available from IBM in other languages. However, you may be required to own a copy of the product or product version in that language in order to access it.

IBM may not offer the products, services, or features discussed in this document in other countries. Consult your local IBM representative for information on the products and services currently available in your area. Any reference to an IBM product, program, or service is not intended to state or imply that only that IBM product, program, or service may be used. Any functionally equivalent product, program, or service that does not infringe any IBM intellectual property right may be used instead. However, it is the user's responsibility to evaluate and verify the operation of any non-IBM product, program, or service.

IBM may have patents or pending patent applications covering subject matter described in this document. The furnishing of this document does not grant you any license to these patents. You can send license inquiries, in writing, to:
IBM Director of Licensing, IBM Corporation, North Castle Drive, MD-NC119, Armonk, NY 10504-1785, US

INTERNATIONAL BUSINESS MACHINES CORPORATION PROVIDES THIS PUBLICATION “AS IS” WITHOUT WARRANTY OF ANY KIND, EITHER EXPRESS OR IMPLIED, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF NON-INFRINGEMENT, MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE. Some jurisdictions do not allow disclaimer of express or implied warranties in certain transactions, therefore, this statement may not apply to you.

This information could include technical inaccuracies or typographical errors. Changes are periodically made to the information herein; these changes will be incorporated in new editions of the publication. IBM may make improvements and/or changes in the product(s) and/or the program(s) described in this publication at any time without notice.

Any references in this information to non-IBM websites are provided for convenience only and do not in any manner serve as an endorsement of those websites. The materials at those websites are not part of the materials for this IBM product and use of those websites is at your own risk.

IBM may use or distribute any of the information you provide in any way it believes appropriate without incurring any obligation to you.

The performance data and client examples cited are presented for illustrative purposes only. Actual performance results may vary depending on specific configurations and operating conditions.

Information concerning non-IBM products was obtained from the suppliers of those products, their published announcements or other publicly available sources. IBM has not tested those products and cannot confirm the accuracy of performance, compatibility or any other claims related to non-IBM products. Questions on the capabilities of non-IBM products should be addressed to the suppliers of those products.

Statements regarding IBM's future direction or intent are subject to change or withdrawal without notice, and represent goals and objectives only.

This information contains examples of data and reports used in daily business operations. To illustrate them as completely as possible, the examples include the names of individuals, companies, brands, and products. All of these names are fictitious and any similarity to actual people or business enterprises is entirely coincidental.

COPYRIGHT LICENSE:

This information contains sample application programs in source language, which illustrate programming techniques on various operating platforms. You may copy, modify, and distribute these sample programs in any form without payment to IBM, for the purposes of developing, using, marketing or distributing application programs conforming to the application programming interface for the operating platform for which the sample programs are written. These examples have not been thoroughly tested under all conditions. IBM, therefore, cannot guarantee or imply reliability, serviceability, or function of these programs. The sample programs are provided “AS IS”, without warranty of any kind. IBM shall not be liable for any damages arising out of your use of the sample programs.

Trademarks

IBM, the IBM logo, and ibm.com are trademarks or registered trademarks of International Business Machines Corporation, registered in many jurisdictions worldwide. Other product and service names might be trademarks of IBM or other companies. A current list of IBM trademarks is available on the web at “Copyright and trademark information” at http://www.ibm.com/legal/copytrade.shtml

The following terms are trademarks or registered trademarks of International Business Machines Corporation, and might also be trademarks or registered trademarks in other countries.

Active Cloud Engine®
AIX®
AIX 5L™
Bluemix®
Cloudant®
Cognos®
DB2®
developerWorks®
Domino®
DS8000®
Easy Tier®
FICON®
FlashCopy®
GDPS®
Geographically Dispersed Parallel Sysplex™
Global Technology Services®
GPFS™
HyperFactor®
HyperSwap®
IBM®
IBM Cloud Managed Services™
IBM Elastic Storage™
IBM FlashCore™
IBM FlashSystem®
IBM Flex System®
IBM SmartCloud®
IBM Spectrum™
IBM Spectrum Accelerate™
IBM Spectrum Archive™
IBM Spectrum Control™
IBM Spectrum Protect™
IBM Spectrum Scale™
IBM Spectrum Storage™
IBM Spectrum Virtualize™
IBM Watson™
IBM z Systems™
IBM z13™
Insight™
Linear Tape File System™
Lotus®
MicroLatency®
OS/400®
Parallel Sysplex®
POWER®
Power Systems™
POWER7®
POWER7+™
POWER8®
PowerHA®
PowerVM®
ProtecTIER®
Real-time Compression™
Redbooks®
Redpaper™
Redbooks (logo) ®
Storwize®
System i®
System i5®
System p5®
System Storage®
SystemMirror®
Tivoli®
Tivoli Storage Manager FastBack®
Variable Stripe RAID™
WebSphere®
XIV®
z Systems™
z/OS®
z/VM®
z/VSE®
z13™
zEnterprise®

The following terms are trademarks of other companies:

SoftLayer, and SoftLayer device are trademarks or registered trademarks of SoftLayer, Inc., an IBM Company.

Intel, Intel Xeon, Intel logo, Intel Inside logo, and Intel Centrino logo are trademarks or registered trademarks of Intel Corporation or its subsidiaries in the United States and other countries.

ITIL is a Registered Trade Mark of AXELOS Limited.

Linux is a trademark of Linus Torvalds in the United States, other countries, or both.

Linear Tape-Open, LTO, Ultrium, the LTO Logo and the Ultrium logo are trademarks of HP, IBM Corp. and Quantum in the U.S. and other countries.

Microsoft, Windows, and the Windows logo are trademarks of Microsoft Corporation in the United States, other countries, or both.

Java, and all Java-based trademarks and logos are trademarks or registered trademarks of Oracle and/or its affiliates.

UNIX is a registered trademark of The Open Group in the United States and other countries.

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

Find and read thousands of IBM Redbooks publications

- Search, bookmark, save and organize favorites
- Get personalized notifications of new content
- Link to the latest Redbooks blogs and videos

Get the latest version of the Redbooks Mobile App

iOS **Download Now** Android

Promote your business in an IBM Redbooks publication

Place a Sponsorship Promotion in an IBM® Redbooks® publication, featuring your business or solution with a link to your web site.

Qualified IBM Business Partners may place a full page promotion in the most popular Redbooks publications. Imagine the power of being seen by users who download millions of Redbooks publications each year!

ibm.com/Redbooks
About Redbooks → Business Partner Programs

THIS PAGE INTENTIONALLY LEFT BLANK

Preface

The IBM® System Storage® Solutions Handbook helps you solve your current and future data storage business requirements. It helps you achieve enhanced storage efficiency by design to allow managed cost, capacity of growth, greater mobility, and stronger control over storage performance and management. It describes the most current IBM storage products, including the IBM Spectrum™ family, IBM FlashSystem®, disk, and tape, as well as virtualized solutions such IBM Storage Cloud.

This IBM Redbooks® publication provides overviews and information about the most current IBM System Storage products. It shows how IBM delivers the right mix of products for nearly every aspect of business continuance and business efficiency. IBM storage products can help you store, safeguard, retrieve, and share your data.

This book is intended as a reference for basic and comprehensive information about the IBM Storage products portfolio. It provides a starting point for establishing your own enterprise storage environment.

This book describes the IBM Storage products as of March, 2016.

Authors

This book was produced by a team of specialists from around the world working at the International Technical Support Organization, Tucson Center.

Ezgi Coskun was a Techline Storage Specialist in Turkey. She provided presales support for entitled business partners, distributors, and IBM sales teams in Middle East and Africa, and Europe. She has over 4 years of IT experience. She holds a bachelors degree in Control Engineering from Istanbul Technical University.

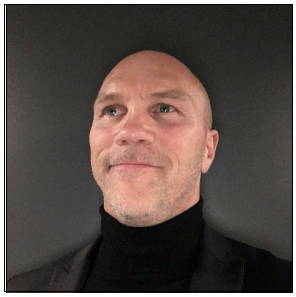

Mikael Lindström is an IBM Certified Consulting IT Specialist, working as a Resilience Software Defined Storage Architect in IBM Global Technology Services® and Technology, Innovation & Automation organization. He has 14 years of IT experience and has worked with IBM since 2006 in various technical and leadership roles. Mikael has authored several IBM Redbooks publications and published several Spectrum Protect white papers at IBM developerWorks®.

Maciej Olejniczak is an IBM Certified IT Expert within Actualizing IT Solutions and is The Open Group Master IT Specialist. He has been working at IBM based in Warsaw, Poland for 8 years. His areas of expertise include IT Architecture, IT Security, Storage, IBM Power Systems™, HPC, and middleware solutions. Maciej currently provides technical pre-sales and post-sales support for customers and sales teams in EMEA region with a specific focus on security products. He has broad experience with IBM Spectrum software.

Oliver Stark has a worldwide responsibility for competitive assessments in the area of Data Protection & Retention solutions. As an expert in physical and virtual tape solutions for open systems, he knows the marketplace very well. Since he joined IBM in 2000, he built his personal storage history in several storage-related positions. Oliver is based close to Mainz in Germany. His professional background is international business and sales. He has managed the IBM Storage Top Gun classes in Europe for years.

Megan Gilge is a Project Leader at the IBM International Technical Support Organization. Before joining the ITSO, she was an Information Developer in the IBM Semiconductor Solutions and User Technologies areas.

Thanks to the following people for their contributions to this project:

Larry Coyne
Bertrand Dufrasne
Karen Orlando
Jon Tate
International Technical Support Organization, Poughkeepsie Center

Liam Dowds
IBM Global Technology Services

Thanks to the authors of the previous editions of this book.

- Authors of the tenth edition, IBM System Storage Solutions Handbook, published in June 2016, were:

  Ezgi Coskun
  Rucel F. Javier
  Pauli Rämö
  Renato Santos
  Megan Gilge

Now you can become a published author, too!

Here's an opportunity to spotlight your skills, grow your career, and become a published author—all at the same time! Join an ITSO residency project and help write a book in your area of expertise, while honing your experience using leading-edge technologies. Your efforts will help to increase product acceptance and customer satisfaction, as you expand your network of technical contacts and relationships. Residencies run from two to six weeks in length, and you can participate either in person or as a remote resident working from your home base.

Find out more about the residency program, browse the residency index, and apply online at:

ibm.com/redbooks/residencies.html

Comments welcome

Your comments are important to us!

We want our books to be as helpful as possible. Send us your comments about this book or other IBM Redbooks publications in one of the following ways:

- Use the online **Contact us** review Redbooks form found at:

  ibm.com/redbooks
- Send your comments in an email to:

  redbooks@us.ibm.com
- Mail your comments to:

  IBM Corporation, International Technical Support Organization
  Dept. HYTD Mail Station P099
  2455 South Road
  Poughkeepsie, NY 12601-5400

Stay connected to IBM Redbooks

- Find us on Facebook:

  http://www.facebook.com/IBMRedbooks
- Follow us on Twitter:

  http://twitter.com/ibmredbooks
- Look for us on LinkedIn:

  http://www.linkedin.com/groups?home=&gid=2130806
- Explore new Redbooks publications, residencies, and workshops with the IBM Redbooks weekly newsletter:

  https://www.redbooks.ibm.com/Redbooks.nsf/subscribe?OpenForm
- Stay current on recent Redbooks publications with RSS Feeds:

  http://www.redbooks.ibm.com/rss.html

1

Introduction

This chapter introduces IBM storage systems and the IBM software-defined architecture.

As the planet becomes more integrated, the rate of data growth is increasing exponentially. This data explosion is making commonly accepted practices of data management inadequate. Social, mobile, cloud, big data, and analytics are driving an explosion in data volumes and causing new challenges for data protection, disaster recovery, regulatory requirements, and standards for compliance. Data management has become one of the highest priorities of organizations today.

This chapter provides an overview of the IBM storage offerings. It includes the following sections:

- IBM Spectrum family and software defined systems
- IBM storage systems
- IBM storage services

1.1 IBM Spectrum family and software defined systems

IBM is a proven leader in collaborating with clients worldwide to solve business problems through the implementation of world class, smart information technology solutions. From hyper-efficient storage arrays and massively scalable to smarter infrastructure management software to comprehensive implementation and operational services capabilities, the IBM portfolio of smart storage cloud capabilities is unmatched in the industry. This section provides an overview of IBM software-defined storage (SDS) products offering with a focus on the IBM Spectrum Storage™ family of products and their capabilities and benefits.

This section describes the business context for SDS within a larger IT software-defined infrastructure (SDI). Fundamental benefits to IT infrastructures and their ultimately supported enterprises are described at a strategic level in support of these key initiative areas:

- Cloud
- Analytics
- Mobile
- Social
- Security
- Traditional computing

1.1.1 Introduction to software-defined architecture

Fundamentally, software-defined architecture enables automation of infrastructure configuration that supports rapid deployment and is aligned to real-time application requirements. These are long standing goals of IT open systems optimization. Software-defined (SD) architecture is an evolving technology that has been made feasible by the abstraction of infrastructure component interfaces delivered through the virtualization of server, storage, and network infrastructures. The current state of SD is similar to where cloud architecture was several years ago. In fact, a key driver of SD development and deployment has been cloud configuration automation requirements. Although SD is finding widespread application in cloud implementations, it can also provide substantial agility and utilization improvements across IT environments, particularly those with rapidly changing application infrastructure support requirements.

SD architecture targets new business models by using tighter interactions with customers like Big Data, Analytics, Social Business, and Mobile. In addition, it targets traditional IT business workloads like enterprise resource planning (ERP), human resources (HR), and customer relationship management (CRM) systems that continue to be important within an integrated infrastructure. Customer-interaction-oriented models are often referred to as Systems of Engagement, whereas traditional back end structures are called Systems of Record.

Together, Systems of Record and Systems of Engagement form the IBM Systems of Insight™ as shown in Figure 1-1.

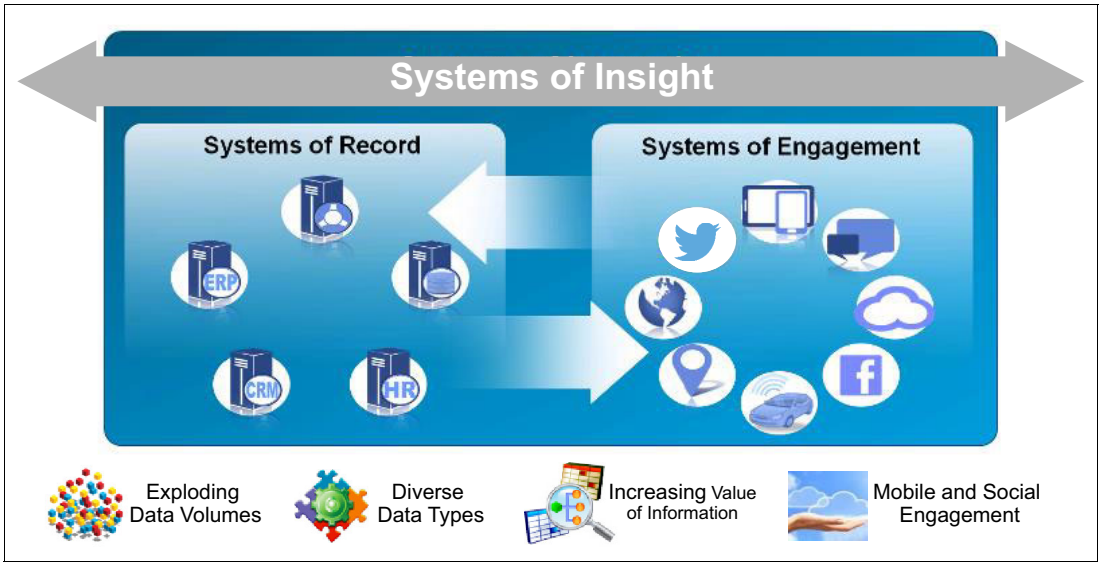

Figure 1-1 Systems of Insight

The requirement to support Systems of Insight workloads is driving IT to reshape itself to accommodate these new business needs while integrating them with traditional applications and infrastructure. Enabling IT's capability to facilitate these increasingly challenging business requirements is the primary goal of software-defined architecture.

SDS is a key component that supports the SDI framework along with software-defined compute and software-defined network constructs as shown in Figure 1-2 on page 4. Although each of these constructs can be used separately, substantial synergy and value results from an integrated approach, which most organizations should adopt. Figure 1-2 shows the SDI framework.

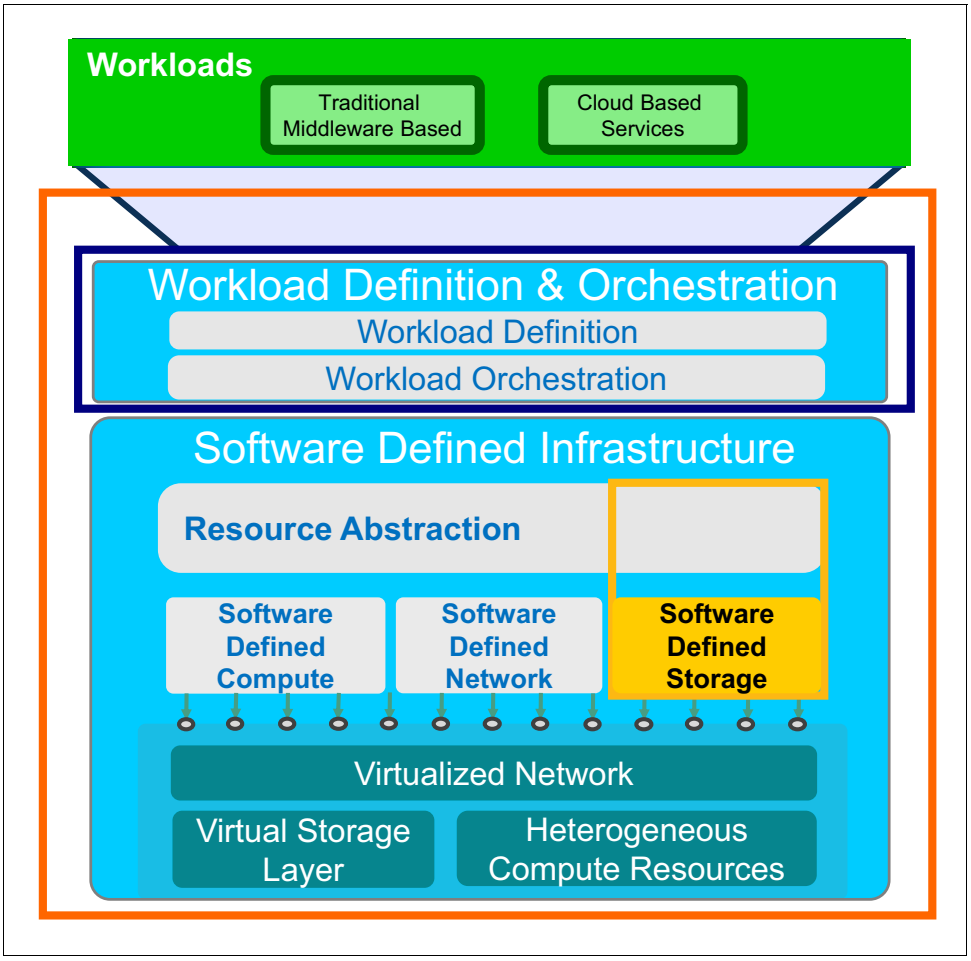

Figure 1-2 SDI framework

SDI drives efficiency by optimizing the connections between workloads and resources based on the needs of the business. Workloads are defined and orchestrated based on usage patterns, and resources are managed and deployed according to business rules and policies. IT Enterprises employing an SDI approach have several core advantages over processes that have traditionally been handled manually.

The fundamental idea of SDI consists of automatically orchestrating infrastructure resources (compute, storage, networks, and management) to programmatically (that is, by using software programming) meet workloads requirements in real time.

1.1.2 Introduction to software-defined storage

SDS, as the name implies, can easily be misinterpreted as storage that is defined in software. Which it is, but this definition lacks several essential important aspects of the currently intended meaning. Most IT storage arrays today are already based in either software, microcode, or both. SDS in today's business context refers to IT storage that goes beyond

typical array interfaces (for example, command line and graphic user) to operate within a larger architectural construct. SDS supports overall IT architectural definition, configuration, and operations (often referred to as SDI or the more restrictive software-defined data center (SDDC)). Although this approach is possible (and probably easier) in a homogeneous, single vendor implementation, its greatest value and versatility are as standardized programming interfaces applied across a heterogeneous multivendor IT infrastructure.

A fundamental inhibitor to fully realized IT optimization has been achieving integrated systems automation, which is often referred to as autonomic computing. This goal has been difficult primarily due to the lack of adequately standardized component interfaces to support a unified infrastructure automation approach. The high levels of customization required involved labor intensive and time consuming work, resulting in higher costs and affecting the ability of IT to rapidly react to new and changing business application requirements.

The primary business reasons for implementing SDS fall into these broad categories, which have interrelated characteristics and benefits:

- Facilitate IT automation to improve business and IT agility at lower cost
- Optimized performance tuning and easier capacity planning
- Optimized systems administration and control to allow effective and efficient resource utilization that lowers cost and supports business requirements
- Simplified architecture to reduce specialized components and skills requirements
- Virtually limitless elastic data scaling
- Support for block, file, and object data types
- Performance tuning
- Capacity planning
- Enable advanced application deployment by employing Systems of Engagement and Systems of Insight by using Systems of Record:
  - Cloud
  - Analytics
  - Mobile
  - Social
  - Security

1.1.3 SDS reference architecture

SDS is one the three main components of the new SDI architecture. The main characteristic of the SDS reference architecture, similar to the software-defined networking (SDN) environment, is the separation of the storage functions into two main layers:

- SDS Control Plane: The control plane is a software layer that manages the virtualized storage resources. It provides all of the high-level functions needed by the customer to run the business workload and enable optimized, flexible, scalable, and rapid provisioning storage infrastructure capacity. These capabilities span functions like storage virtualization, policies automation, analytics and optimization, backup and copy management, security, and integration with the API services, including other cloud provider services.
- SDS Data Plane: The data plane encompasses the infrastructure where data is processed. It consists of all basic storage management functions such as virtualization, RAID protection, tiering, copy services (remote, local, synchronous, asynchronous, and

point-in-time), encryption, and data deduplication that can be started and managed by the control plane. The data plane is the interface to the hardware infrastructure where the data is stored. It provides a complete range of data access possibilities, spanning traditional access methods like block I/O (for example, iSCSI) or File I/O (POSIX compliant), to object-storage or Hadoop Distributed File System (HDFS).

Figure 1-3 shows the SDS capabilities of the control plane and data plane.

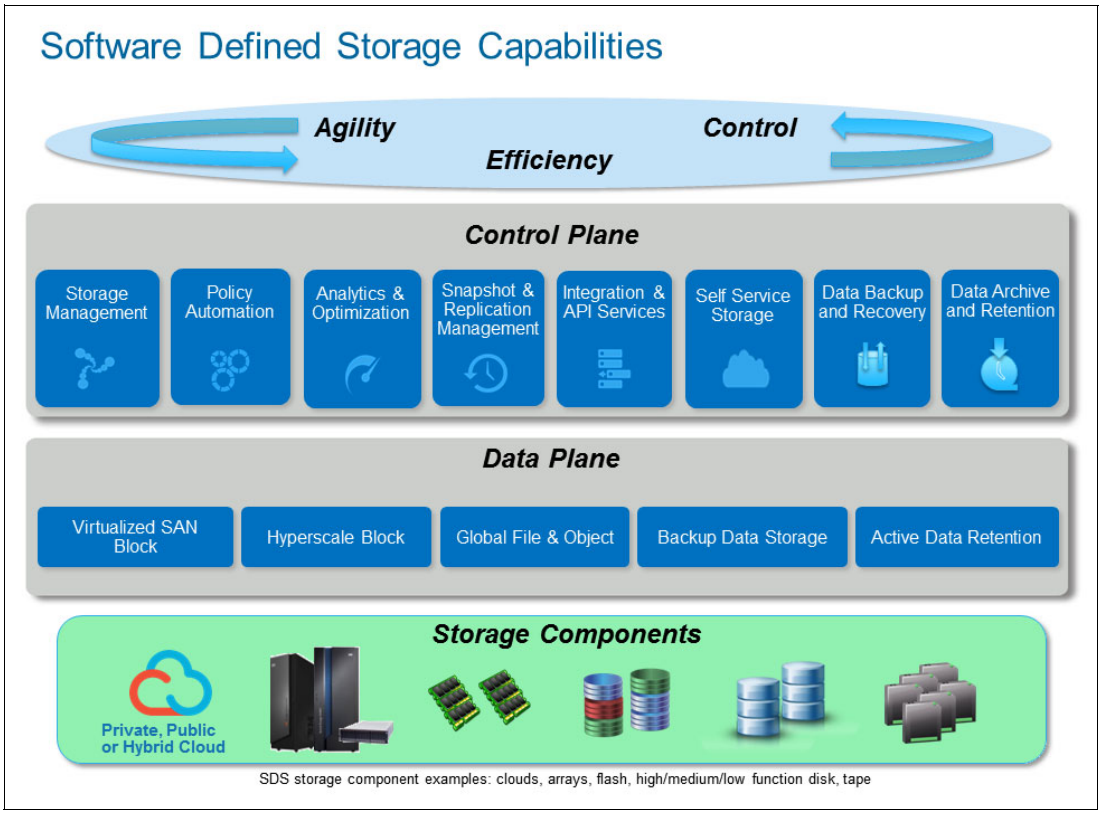

Figure 1-3 SDS capabilities through the control plane and data plane

SDS provides the agility, control, and efficiency that are needed to meet rapidly changing business requirements by dynamically optimizing infrastructure capabilities to application service level requirements. New business requirements are the driving force for the emergence of this new IT storage infrastructure architecture. SDS is built with standardized software application programming interfaces to provide organizations with the underlying capabilities to support applications aligned to the digital consumption model of the digital economy. It also provides a compelling value proposition for optimizing traditional workloads within this consolidated architectural construct.

For more information about the IBM software-defined storage architecture, see Chapter 2, “IBM Spectrum Storage family” on page 13.

1.2 IBM storage systems

In this era of high storage demand, the amount of information that most organizations deal with on a daily basis doubles as often as every two years. At the same time, big data projects can quickly double or triple this volume of information. Information is becoming more valuable, with the ability to analyze it counting as a competitive advantage. To stay ahead,

organizations need faster access to more data at the same time. Inflexible infrastructures with increasing complexity can create significant obstacles for organizations that are trying to take advantage of these tremendous opportunities.

IBM disk, hybrid disk, and flash memory systems provide storage efficiency solutions such as Real-time Compression, automated tiering, virtualization, and thin provisioning. These storage solutions increase the data storage optimization opportunities for organizations of all sizes to boost system performance and lower IT costs. These growing demands require a new approach to storage. They require a more intelligent, more efficient, and more automated approach to storage that fundamentally changes the way you think about it. Storage agility and efficiency for today's applications is intended for cloud, analytics, security, mobile, and social applications.

IBM storage products can help deliver the solutions that are required to implement a dynamic infrastructure. Such solutions include the following capabilities:

- Smarter information infrastructure
- Storage virtualization
- Space efficiency and reduced storage footprint in data centers
- Energy efficiency
- Data protection, information security, and resiliency
- Service and information management
- Strategic outsourcing (see 1.3, "IBM storage services" on page 12)

This following subsections provide an overview of the IBM storage systems and links to more information:

- IBM SAN Volume Controller and IBM Storwize family
- IBM FlashSystem family
- IBM DS8880 Storage System
- IBM XIV Storage System
- IBM Elastic Storage Server
- IBM tape storage
- IBM storage area network products

1.2.1 IBM SAN Volume Controller and IBM Storwize family

IBM SAN Volume Controller (SVC) and the IBM Storwize® family are designed to deliver the benefits of storage virtualization to large enterprises, small businesses, and midmarket companies.

IBM SAN Volume Controller

IBM Spectrum Virtualize™ functions benefit all virtualized storage. For example, IBM Easy Tier® optimizes use of flash storage. IBM Real-time Compression™ enhances efficiency even further by enabling the storage of more active primary data in the same physical disk space. Encryption helps improve data security, and high-performance thin provisioning helps automate provisioning. These benefits can help extend the useful life of existing storage assets, reducing costs.

Integrating these functions into SVC also means that they are designed to operate smoothly together, reducing management effort.

SVC can help you achieve the following benefits:

- Enhance data storage functions, economics, and flexibility with sophisticated virtualization
- Use hardware-accelerated data compression for efficiency and performance

- Use encryption to help improve security for data on existing storage systems
- Move data among virtualized storage systems without disruptions
- Optimize tiered storage, including flash storage, automatically with IBM Easy Tier
- Improve network utilization for remote mirroring with innovative replication technology
- Implement multi-site configurations for high availability and data mobility between data centers

For more information about IBM Spectrum Virtualize, see 2.3.1, “IBM Spectrum Virtualize” on page 29.

For more information about SVC, see 4.1, “IBM SAN Volume Controller” on page 82.

IBM Storwize

Built with IBM Spectrum Virtualize software, the IBM Storwize family provides hybrid solutions with common functionality, management, and flexibility. It includes built-in functions such as Real-time Compression and Easy Tier technology optimizing flash and hard disk drives to deliver extraordinary levels of efficiency and high performance. Available in a wide range of storage systems, IBM Storwize family delivers sophisticated capabilities that are easy to deploy, and help control costs for growing businesses.

This book describes the following IBM Storwize products:

- IBM Storwize V7000 Unified and IBM Storwize V7000. These virtualized systems can simplify storage by consolidating workloads into a single storage system.
- IBM Storwize V5000. This highly flexible, easy-to-use, virtualized storage system provides high performance and advanced functionality for midsized organizations.
- IBM Storwize V3700. This affordable, easy-to-use, self-optimizing disk system is designed for small and midsized organizations.

For more information about IBM Spectrum Virtualize, see 2.3.1, “IBM Spectrum Virtualize” on page 29.

For more information about the IBM Storwize family, see Chapter 4, “IBM SAN Volume Controller and IBM Storwize family” on page 81.

1.2.2 IBM FlashSystem family

IBM FlashSystem products are all-flash storage arrays that offer scalable performance, agile integration, and enduring economics. This book describes the following products:

- IBM FlashSystem V9000

  These all-flash storage systems deliver the full capabilities of FlashCore technology coupled with a rich set of the features found in the most advanced software-defined storage solutions. These features include IBM Real-time Compression, dynamic tiering, thin provisioning, snapshots, cloning, replication, data copy services, and high-availability configurations.

- IBM FlashSystem 900

  Easy to deploy and manage, IBM FlashSystem 900 is designed to accelerate the applications that drive business. Powered by FlashCore technology, IBM FlashSystem 900 delivers the extreme performance, IBM MicroLatency®, enterprise reliability, and operational efficiencies required for gaining competitive advantage in today’s dynamic marketplace.

For more information, see Chapter 5, “IBM FlashSystem family” on page 127.

1.2.3 IBM DS8880 Storage System

IBM DS8880 is a new family of enterprise hard disk drive (HDD) and flash data systems that delivers mission-critical acceleration, uncompromising availability, unparalleled integration with IBM z Systems™ and IBM Power Systems, and transformational efficiency through industry-leading capabilities and the highest levels of security.

For more information, see Chapter 6, “IBM DS8800 Storage System” on page 141.

1.2.4 IBM XIV Storage System

IBM XIV® Storage System is a high-end, grid-architecture storage system designed to support the diverse and dynamic workloads of open systems with simplicity, resiliency, and predictability. These characteristics make it ideal for cloud and big-data analytics environments. An edition for cloud providers also delivers outstanding flexibility and cloud economics.

For more information, see Chapter 7, “IBM XIV Storage System” on page 157.

1.2.5 IBM Elastic Storage Server

IBM Elastic Storage™ Server is an integrated network-attached storage (NAS) solution that provides the building blocks for IBM Spectrum Scale™ clusters. It significantly lowers the cost of storage by using IBM Spectrum Scale RAID, which allows the use of commodity disk shelves instead of storage controllers for physical storage. IBM Elastic Storage Server can scale out from 50 terabytes to several petabytes and accelerate business results by providing file, object, and Hadoop interfaces to a common storage pool.

For more information, see 2.3.3, “IBM Spectrum Scale” on page 38 and the following website:

http://www.ibm.com/systems/storage/spectrum/scale/

1.2.6 IBM tape storage

IBM offers a full range of tape storage solutions, including drives, autoloaders, libraries, virtual tape systems, and Spectrum Archive software that can make tape as easy to use as disk. IBM tape storage provides the following benefits:

- Provide high-capacity storage backup for IT infrastructures and low total cost of ownership
- Store removable cartridges to protect them from viruses, sabotage, and corruption
- Support scalability by simply adding more tape cartridges, versus drives
- Add data protection by moving portable media offsite
- Shorten backup windows and enable rapid data restores with virtual tape
- Address security compliance requirements by using encryption and Write Once Read Many (WORM) technologies

IBM tape drives

IBM enterprise-class tape storage products are designed to offer the high performance, availability, reliability, and capacity that are needed to provide enterprise-class solutions for mass storage, data archiving, enterprise backup, and disaster recovery.

This book describes the following enterprise-class tape storage products:

- IBM TS1150 tape drive
- IBM TS1140 tape drive
- IBM TS1130 tape drive
- IBM TS4500 tape library
- IBM TS3500 tape library
- IBM Tape Controller for IBM z Systems
- IBM TS7650G ProtecTIER® Deduplication Gateway
- IBM TS7740 and TS7720

IBM midrange tape storage products are designed to provide reliable and flexible data backup, archiving, and management, today and into the future.

This publication describes the following midrange tape storage products:

- IBM TS7620 ProtecTIER Deduplication Appliance
- IBM TS4500 Tape library
- IBM TS3500 Tape library
- IBM TS3310 Tape library

IBM entry-level tape products provide reliable, affordable data backup and protection.

This publication describes the following entry-level tape drives:

- IBM TS2270 tape drive
- IBM TS2360 tape drive
- IBM TS2260 tape drive
- IBM TS2250 tape drive
- IBM TS2240 tape drive
- IBM TS2900 tape autoloader
- IBM TS3200 tape library
- IBM TS3100 tape library

For more information, see 8.1, “IBM tape drives” on page 178.

IBM tape storage media

Today’s data centers must satisfy increasing demands for data retention, security, and availability while contending with reduced budgets. IBM tape storage media solutions have been developed to help you address these challenges.

This publication describes the following media:

- IBM 3592 tape cartridges
- IBM Linear Tape-Open (LTO) Ultrium tape cartridges
- IBM DAT/DDS tape cartridges

For more information, see 8.1, “IBM tape drives” on page 178.

Tape automation

IBM System Storage offers a full range of tape libraries and autoloaders that feature high performance and capacity for entry, midrange, and enterprise system environments. Use the advanced technology to handle your backup, save, and restore, and archival data storage needs.

This book describes the following products:

- “IBM TS2900 tape autoloader” on page 193
- “IBM TS3100 tape library” on page 195

- “IBM TS3200 tape library” on page 197
- “IBM TS3310 tape library” on page 199
- “IBM TS3500 tape library” on page 202
- “IBM TS4500 tape library” on page 207
- “TS4500 product description” on page 209

Virtual tape

A virtual tape library provides high performance backup and restore by using disk arrays and virtualization software. A virtual tape library is a unique blend of several storage tiers. The lifecycle of data from its creation at the server level migrates by backup software to a virtual tape library. A virtual tape library is a powerful combination of high performance storage area network (SAN)-attached disk and high performance servers that emulate a tape storage device. At this level, you have many options for the data. For example, the data can remain on the virtual tape library indefinitely if enough space is available, or it can be migrated to tape for offsite storage, archive, or both.

Virtual tape libraries fill a void in the backup infrastructure for data that needs to be restored at a certain moment. Many restore requests happen within six weeks of the data being backed up. Backup software can be configured to back up data to a virtual tape library and then create a virtual tape-to-tape copy for offsite deployment. It is no longer necessary to call the tapes back from an offsite location unless data is required from years past.

This publication describes the following products:

- “IBM TS7600 ProtecTIER systems” on page 216
- “TS7650G ProtecTIER Deduplication Gateway” on page 217
- “TS7620 ProtecTIER Deduplication Appliance Express” on page 219
- “ProtecTIER Manager” on page 220
- “ProtecTIER native replication” on page 220
- “IBM TS7700” on page 222
- “IBM Data Facility Storage Management Subsystem” on page 229

1.2.7 IBM storage area network products

The IBM SAN products provide the following benefits:

- Choose from a comprehensive portfolio of SAN switches, storage, software, and solutions
- Capitalize on enterprise-wide data sharing and collaboration with Fibre Channel connectivity
- Connect servers and storage systems in local, campus, metro, and global infrastructures by using intelligent SAN switches, directors, and routers
- Help drive innovation and create a dynamic, scalable infrastructure
- Access information cost-effectively

IBM provides the following types of SAN products:

- Enterprise SAN directors provide metro and global connectivity in addition to high availability, scalability, and performance between sites over different networking protocols. This publication describes the following products:
  - IBM System Storage SAN768B-2 and SAN384B-2
  - Cisco MDS 9710 Multilayer Director for IBM System Networking
  - Cisco MDS 9706 Multilayer Director for IBM System Storage
  - Cisco MDS 9500 Series Multilayer Directors for IBM System Storage

- Midrange SAN switches can scale to provide solutions for small and midsized business up to large enterprises. This publication describes the following products:
  - IBM System Networking SAN96B-5
  - Cisco MDS 9396S 16G Multilayer Fabric Switch for IBM System Storage
  - Cisco MDS 9148S 16G Multilayer Fabric Switch for IBM System Storage
- Entry SAN switches are designed to provide easy-to-use, affordable solutions for small and midsized businesses. This publication describes the following products:
  - IBM System Networking SAN24B-5
  - IBM System Storage SAN24B-4 Express
- SAN specialty switches deliver capabilities that range from connecting heterogeneous SAN fabrics, to enabling distance extension using Fibre Channel over IP, to converging SAN and Ethernet traffic on a single platform. This publication describes the following products:
  - IBM System Storage SAN42B-R
  - Cisco MDS 9250i Multiservice Fabric Switch for IBM System Storage
  - IBM System Storage SAN06B-R extension switch

The IBM SAN products are described in Chapter 9, “Storage networking products” on page 231.

For more information, see the following website:

http://www.ibm.com/systems/storage/san/

1.3 IBM storage services

IBM delivers deep technological expertise and rich industry insight to help you align your IT and business objectives. IBM offers a broad range of cloud and managed services.

For the current offerings, see the following links:

- IBM Global Technology Services

  http://www.ibm.com/services/us/en/outsourcing/index.html
- IBM Storage

  http://www.ibm.com/systems/storage
- IBM System Storage interactive product guide

  http://www.ibm.com/common/ssi/cgi-bin/ssialias?htmlfid=TSO00364USEN

2

IBM Spectrum Storage family

This chapter provides an overview of IBM software-defined storage (SDS) product offerings with a focus on the IBM Spectrum Storage family of products and their capabilities and benefits.

The products are organized by their functions within the SDS control plane or SDS data plane. The control plane is the software layer that manages administrative functions such as configuration, monitoring, replication, policy automation, and provisioning, for software-defined storage resources. Data is processed and stored in the data plane.

This chapter includes the following sections:

- SDS architecture
- SDS control plane
- SDS data plane

2.1 SDS architecture

Figure 2-1 shows the IBM SDS architecture with a mapping of the Spectrum Storage family of products across the SDS control plane and data plane.

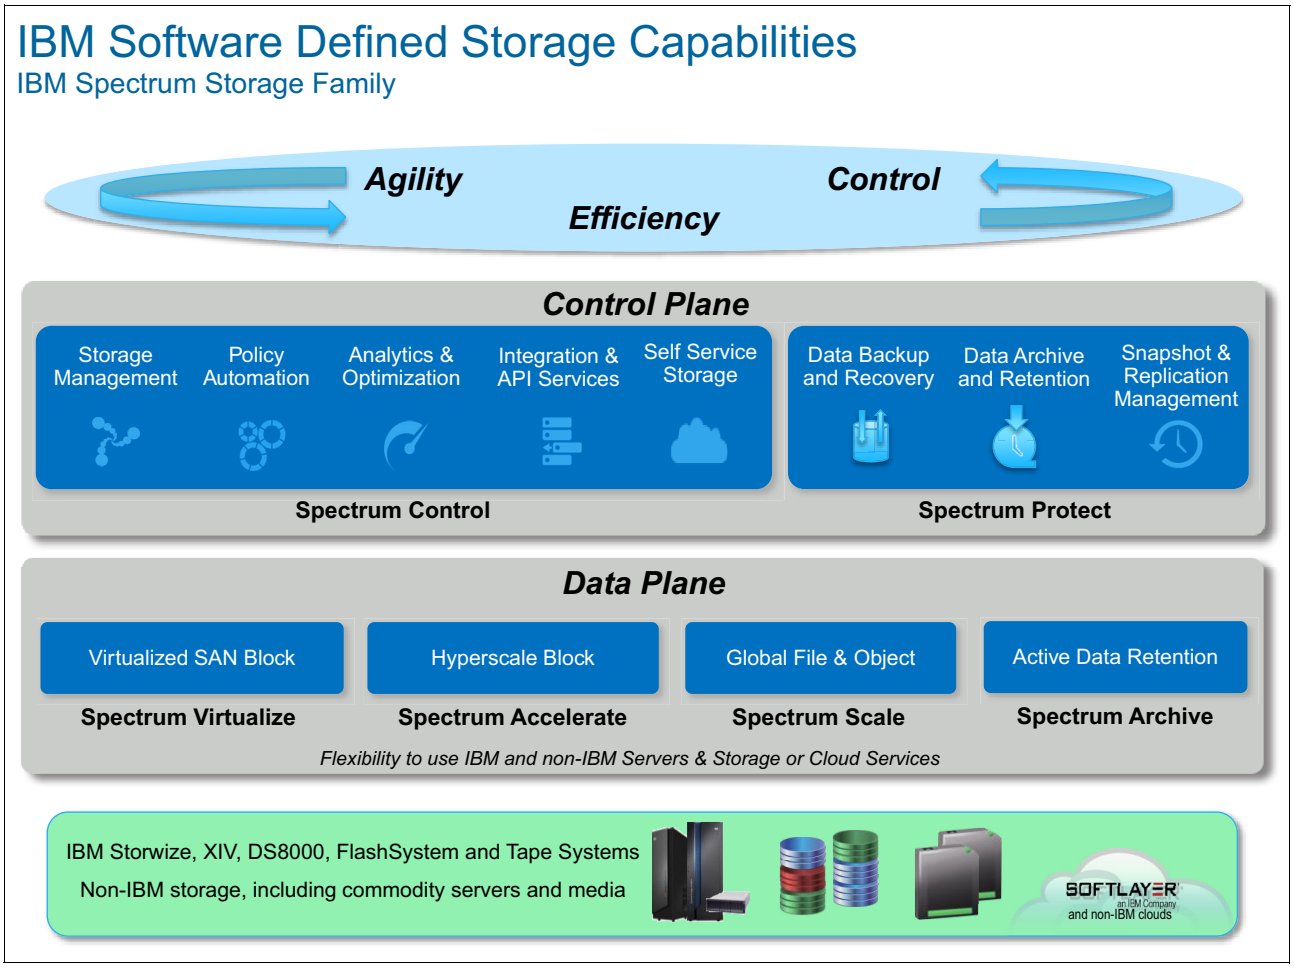

Figure 2-1 IBM Spectrum Storage family mapped to the SDS control plane and data plane

Table 2-1 is a high-level overview of the IBM Spectrum Storage family. It pairs descriptions with the products that provide the functions.

Table 2-1 IBM Spectrum Storage Family descriptions

	Spectrum Storage family member	Description	Former name
SDS Control Plane			
	Spectrum Control	Automated control and optimization of storage and data infrastructure	IBM Tivoli® Storage Productivity Center, management layer of IBM Virtual Storage Center, and IBM Spectrum Control™ Base Edition (IBM Storage Integration Server)
	IBM Spectrum Protect™	Optimized data protection for client data through backup and restore capabilities	IBM Tivoli Storage Manager Suite for Unified Recovery IBM Spectrum Protect Snapshot
SDS Data Plane			
	Spectrum Virtualize	Core SAN Volume Controller function is virtualization that frees client data from IT boundaries	IBM SAN Volume Controller
	Spectrum Accelerate	Enterprise storage for cloud that is deployed in minutes instead of months	IBM XIV as software
	Spectrum Scale	Storage scalability to yottabytes and across geographical boundaries	IBM General Parallel File System (GPFS™)
	Spectrum Archive	Enables long-term storage of low activity data	IBM Linear Tape File System™ Enterprise Edition, Library Edition, and Single Drive Edition

2.2 SDS control plane

The control plane is a software layer that manages the virtualized storage resources. It provides all the high-level functions that are needed by the customer to run the business workload and enable optimized, flexible, scalable, and rapid provisioning storage infrastructure capacity. These capabilities span functions like storage virtualization, policy automation, analytics and optimization, backup and copy management, security, and integration with the API services, including other cloud provider services.

This section describes the IBM software product offerings that provide the building blocks for the SDS control plane:

- IBM Spectrum Control
- IBM Spectrum Protect

2.2.1 IBM Spectrum Control

IBM Spectrum Control provides efficient infrastructure management for virtualized, cloud, and software-defined storage to simplify and automate storage provisioning, capacity management, availability monitoring, and reporting.

Key capabilities

Spectrum Control helps organizations transition to new workloads and updated storage infrastructures by providing these tools:

- A single console for managing all types of data on disk, flash, file, and object storage systems
- Simplified visual administration tools that include an advanced web-based user interface, a VMware vCenter plug-in, and IBM Cognos® Business Intelligence with pre-designed reports
- Analytics-driven tiered storage optimization that automatically moves data to the most cost-effective tier to help support business required service levels
- Near-instant data protection and recovery capabilities that protect data without affecting application performance
- Choice of license plans with three deployment options that can fit any size organization, including mid-sized organizations.

Benefits

The Spectrum Control benefits address the two most significant storage costs: Physical capacity and storage administration. Here are several highlights of those broad based benefits:

- Volume-level, cross-platform automated storage tiering that reduces users' cost of storage by up to 50%[1]
- Cloud and software defined storage that provides a service catalog, automated provisioning, and optimized utilization of pools of storage
- Intelligent Performance Management with application and department views of the storage infrastructure

[1] Based on IBM experience using storage analysis

Figure 2-2 shows the Spectrum Control dashboard window where all the managed resources in your data server are presented in an aggregated view.

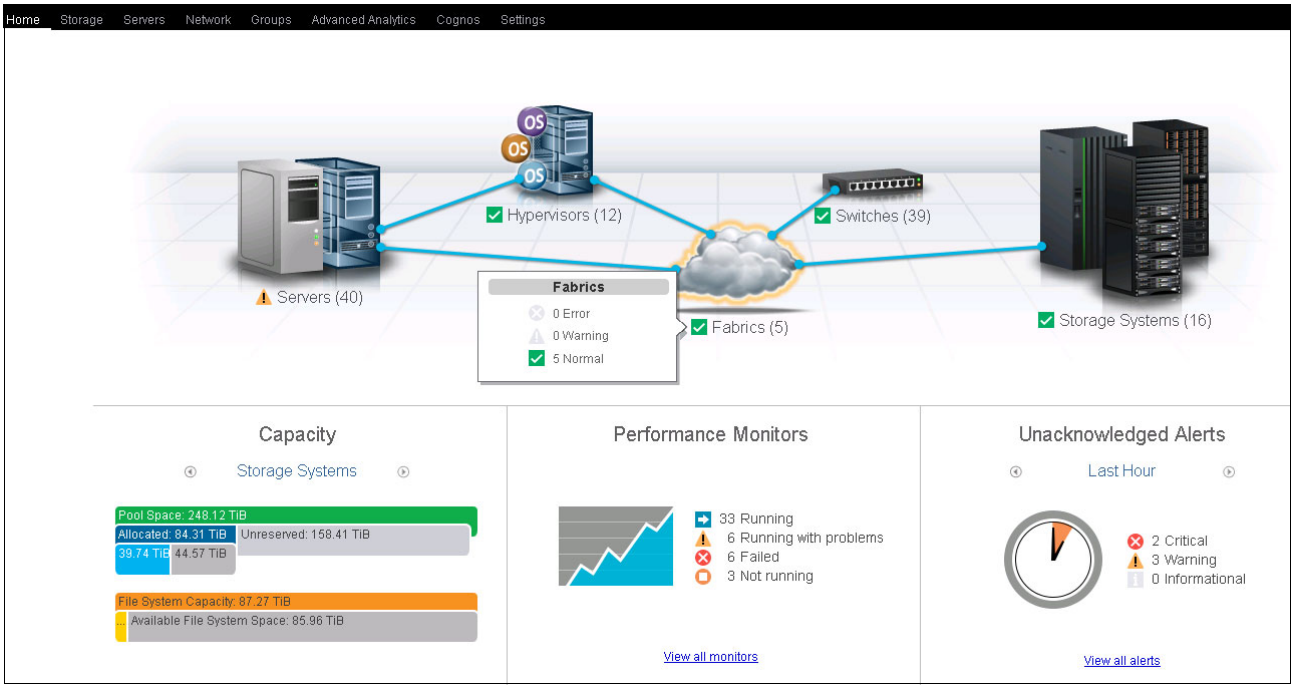

Figure 2-2 Single dashboard for monitoring all storage components

IBM Spectrum Control data management and storage management solutions provide comprehensive monitoring, automation, and analytics. They address the two most significant storage costs: Physical capacity and storage administration. You can get important insights into cloud, virtualized, and software defined storage environments.

IBM data management and storage management solutions can help with these tasks:

- Identify and categorize storage assets for file, block, and object data
- Generate departmental and application views of storage
- Optimize data placement within the infrastructure
- Identify unused and wasted storage space

IBM Spectrum Control is a comprehensive, end-to-end data and storage management solution that monitors, automates, and analyzes multivendor storage environments.

IBM Spectrum Control provides a single point of control that helps administrators manage every aspect of the storage infrastructure, between the hosts, through the fabric, and down to the physical disks, across multi-site storage environments. IBM Spectrum Control helps consolidate management of file, object, SAN, server-based, and software- defined storage and enables you to manage storage your way. Optimized interfaces are included for storage specialists, VMware vCenter users, and remote management with open application programming interfaces (APIs) such as OpenStack and IBM Cloud Orchestrator.

IBM Spectrum Control delivers a comprehensive solution that can significantly improve monitoring, automation, and analytics capabilities in multivendor storage environments:

- Heterogeneous storage systems, ports, and switches

  IBM Spectrum Control provides device-level, integrated storage infrastructure management capabilities for managing both IBM and non-IBM storage systems that are

managed with IBM SAN Volume Controller. Agent and agentless options are available. Heterogeneous storage support is offered by using the Storage Networking Industry Association (SNIA) Storage Management Initiative Specification (SMI-S) v1.0.2, v1.1, and v1.5 or later.

- Device management

  IBM Spectrum Control helps monitor and manage performance and measure service levels by storing received performance statistics in database tables for future use. IBM Spectrum Control sets performance thresholds for devices based on selected performance metrics, generating alerts when those thresholds are exceeded. These capabilities help the solution simplify the complex monitoring of multiple SAN-attached storage devices.

  IBM Spectrum Control includes these device management capabilities:

  - Proactive performance management from a single, integrated console for monitoring storage devices
  - Monitoring of metrics, such as I/O rates and cache utilization, and support for storage optimization through the identification of the best logical unit numbers (LUNs) across multiple storage subsystems
  - The measuring and tracking of service levels by storing historical performance statistics for analysis and report generation
  - Generation of timely alerts that enable event action by setting performance thresholds based on business policies

- SAN fabric management

  IBM Spectrum Control supports multivendor SANs and includes automatic resource and topology discovery, monitoring and alerts, zone control, and SAN error-prediction capabilities. IBM Spectrum Control diagnostic capabilities show which resources are impacted by an availability or performance issue in the SAN.

  IBM Spectrum Control includes these SAN fabric monitoring capabilities:

  - Automated device discovery, topology rendering, and zone control for SANs
  - Multiple SAN views, including physical, logical, and zone
  - Diagnostic capabilities that pinpoint resources affected by an availability or performance issue
  - Performance monitoring at the port and switch level
  - Enterprise scalability that enables upgrading from SAN islands to enterprise SANs

- Storage provisioning

  IBM Spectrum Control includes optimized interfaces for storage specialists, VMware vCenter users, IBM Cloud Orchestrator environments, and open APIs such as Storage Management API for Cloud (SMAC). It can help simplify and standardize storage provisioning, so users can quickly get the storage that they need to do their jobs.

- Advanced alert management

  Advanced alert management in IBM Spectrum Control can help administrators identify and resolve problems fast. Policy-based automation enables automated responses, based on business policies. Event processing can reduce the complexity of managing alerts in large storage environments.

- Scalable copy services management

  IBM Spectrum Control includes copy services management capabilities that help simplify the data replication process. It can automate complex replication tasks without scripts, and

can scale to support hundreds of replication sessions across thousands of volumes. IBM Spectrum Control supports the following session types: IBM FlashCopy®, IBM XIV snapshots, IBM Basic HyperSwap®, Metro Mirror, Global Mirror, Metro/Global Mirror, and Global Copy.

For copy services management, IBM Spectrum Control can use IBM DS8000®, XIV, IBM SAN Volume Controller (SVC), IBM Storwize V7000 Unified, IBM Storwize V7000, IBM Storwize V3700, and IBM Storwize V3500.

IBM Spectrum Control integrates with IBM Tivoli System Automation to enable a single point of control for both multi-tiered business applications and the corresponding copy services between two sites. This integration offers the following capabilities:

- Integrated monitoring and control of storage hardware replication sessions
- Automatic start, stop, and move of tiered applications
- Lower-priority business application shutdown to keep higher-priority business applications running
- End-to-end application recovery

- Tiered storage analysis and optimization

  IBM Spectrum Control Tiered Storage Optimizer helps organizations optimize price and performance without manual tuning. Tiered Storage Optimizer uses proprietary analytics from IBM Research that can help keep data on the correct storage tier over time. Storage tier decisions are based on data usage patterns, rather than on initial projections or guesswork. Tiered Storage Optimizer moves data volumes between storage pools and systems.

- SAN Error Predictor

  SAN Error Predictor is an autonomic computing capability that is designed to help predict SAN network outages before they occur. This functionality incorporates predictive failure analysis into the storage network environment that allows administrators to be more proactive in ensuring SAN availability. SAN Error Predictor enables network errors to be resolved before they affect data and application availability.

- Performance analysis and management

  IBM Spectrum Control integrates built-in, context- sensitive performance management that is accessible to any administrator. Performance management is server-centric, so administrators can more easily visualize storage performance problems from the user perspective. Administrators can quickly generate performance graphs from several perspectives and align them with a simple mouse click. Performance scenarios can be cloned, so administrators can look back to better understand how problems first developed

IBM Spectrum Control is available in a range of options to fit practically any size environment and budget. For example, IBM Spectrum Control Storage Insights is a cloud-based solution that deploys in minutes and includes advanced analytics for storage tier optimization. On-premises solutions include standard and advanced editions, available in capacity or per-enclosure license options. IBM Virtual Storage Center bundles IBM Spectrum Control Advanced Edition with multi-brand storage virtualization and snapshot management. IBM Spectrum Control Base consolidates VMware integration for IBM storage systems into a single download.

IBM Spectrum Control Standard Edition

IBM Spectrum Control Standard Edition is available in capacity or per-enclosure license options. The IBM Spectrum Control Standard Edition license is priced by the capacity of the storage systems that it manages. The IBM Spectrum Control Standard Select Edition license is priced by the number of storage enclosures that it manages.

These licenses provide the following features:

- Single dashboard view of the storage environment, which enables you to manage storage systems, hypervisors, servers, and networks
- Capacity and usage monitoring in addition to management and planning for the heterogeneous storage environment
- Performance monitoring for storage systems and networks
- Health and alerting for storage systems, hypervisors, servers, and networks
- Modeling of departments and applications to understand the consumer view of the environment
- Administrative tools, which includes the following:
  - VMware vCenter plug-in, which enables you to manage and provision virtual storage in a vSphere environment, and to view reports about storage systems that are monitored by IBM Spectrum Control
  - Reports from IBM Cognos Business Intelligence, which provides reporting and analytics information

This license includes a license for IBM Copy Services Manager (formerly known as IBM Tivoli Storage Productivity Center for Replication). Copy Services Manager manages 2-site replication, 3-site replication, and advanced copy services.

For more information about Spectrum Control Standard, see *IBM Spectrum Family: IBM Spectrum Control Standard Edition*, SG24-8321.

IBM Spectrum Control Advanced Edition

IBM Spectrum Control Advanced Edition is available in capacity or per-enclosure license options like Standard Edition. The IBM Spectrum Control Advanced Edition license is priced by the capacity of the storage systems that it manages. The IBM Spectrum Control Advanced Select Edition is priced by the number of storage enclosures that it manages.

These licenses contain everything that is in the IBM Spectrum Control Standard Edition license, and the following additional features:

- Analytics-driven, tiered-storage optimization that automatically moves data to the most cost-effective tier
- Pool balancing
- Block storage provisioning
- Application-aware snapshot-based protection that is offered by IBM Spectrum Protect Snapshot (formerly known as IBM FlashCopy Manager)

For more information, see the IBM Knowledge Center:

http://www.ibm.com/support/knowledgecenter/SS5R93_5.2.9/com.ibm.spectrum.sc.doc/tpc_kc_homepage.html

The IBM Tivoli Storage Productivity Center license is now IBM Spectrum Control Standard Edition.

IBM Spectrum Control Advanced Edition is a new software-only product that offers the same functions as the Storage Analytics Engine product that was part of IBM Virtual Storage Center.

For a complete list of devices that can be used with IBM Spectrum Control, see the following website:

http://www.ibm.com/support/docview.wss?uid=swg21386446

IBM Spectrum Control Base Edition

IBM Spectrum Control Base Edition is a centralized server that consolidates a range of IBM storage provisioning, automation, and monitoring solutions through a unified server platform. As shown in Figure 2-3, it provides a single-server back-end location and enables centralized management of IBM storage resources for the use of independent software vendor (ISV) platforms and frameworks. These frameworks currently include VMware vCenter Server, VMware vSphere Web Client (vWC), and VMware vSphere Storage APIs for Storage Awareness (VASA), VMware vCenter Operations Manager (vCOps), VMware vCenter Orchestrator (vCO). Spectrum Control Base Edition is available for no extra fee to storage-licensed clients.

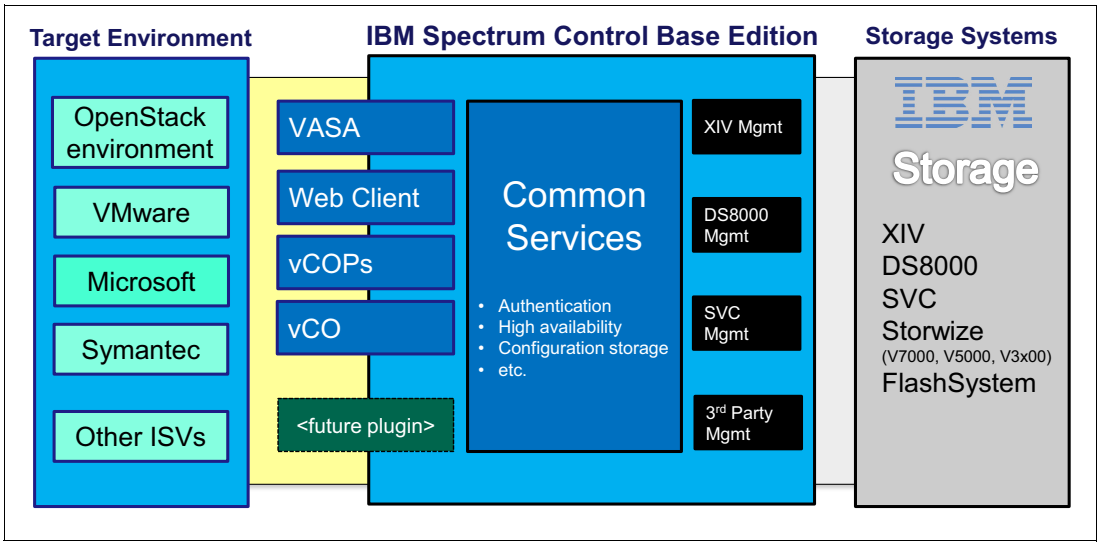

Figure 2-3 Spectrum Control Base Edition

Control plane: IBM Spectrum Control Base Edition is not in the data path. It runs in the control plane as shown in Figure 2-1 on page 14.

Spectrum Control Base Edition integrates between environments like OpenStack components, VMware, and SNIA SMI-S driven environments, and the storage subsystems. Spectrum Control Base supports requests from APIs like VASA and Next Generation Client (NGC).

Spectrum Control Base Edition provides common services such as authentication, high availability, and storage configuration for all IBM Storage in homogeneous and heterogeneous multiple target environments. Spectrum Control Base Edition manages IBM XIV Storage System, IBM DS8000 series, IBM SAN Volume Controller, the IBM Storwize family, and third-party storage subsystems.

Support for VMware virtual volumes

IBM Spectrum Control Base Edition delivers comprehensive storage virtualization support, by using VMware Virtual Volume (VVOL) technology.

Note: The virtual volume functionality is supported by the IBM FlashSystem V9000 (7.6 or later), IBM XIV (11.5.1 or later), and IBM Spectrum Virtualize/Storwize (7.6 or later) storage systems.

The VVOL architecture, introduced in VMware VASA 2.0, preserves the concept of a traditional datastore, maintaining familiarity and compatibility with previous data storage implementations. However, the virtual disks in a VVOL datastore can use different storage attributes (services), which can include thin/thick provisioning, snapshot support, encryption, and so on. Moreover, each VVOL can be managed independently. This implementation improves system scalability, ensures granular management, uses hardware features, and improves performance of storage systems at the VM level, providing a complete end-to-end cloud solution. Another entity, a storage space, includes one or more services, and can be assigned to different storage customers.

For more information, see the IBM Spectrum Control Base Edition information in the IBM Knowledge Center:

http://www.ibm.com/support/knowledgecenter/STWMS9/landing/IBM_Spectrum_Control_Base_Edition_welcome_page.html

For more information about VVOL architecture implementation in XIV, see *IBM Spectrum Control Base: Enabling VMware Virtual Volumes with IBM XIV Storage System*, REDP-5183.

For more information about VVOL implementation with IBM Spectrum Virtualize, see *Quick-start Guide to Configuring VMware Virtual Volumes for Systems Powered by IBM Spectrum Virtualize*, REDP-5321.

Spectrum Control in an OpenStack environment

The Spectrum Control OpenStack Cinder driver enables your OpenStack powered cloud environment to use your Spectrum Control installation for block storage provisioning. Spectrum Control provides block storage provisioning capabilities that a storage administrator can use to define the properties and characteristics of storage volumes within a particular service class. For example, a block storage service class can define RAID levels, tiers of storage, and various other storage characteristics.

IBM Storage Insights

IBM Storage Insights is an analytics-driven, storage resource management solution that is delivered over the cloud. The solution uses cloud technology to provide visibility into on-premises storage to help you optimize you storage environments. This SaaS solution, which runs on SoftLayer®, can deploy in as little as 5 minutes and show actionable insights in 30 minutes.

Storage Insights is a cloud data and storage management service that is deployed in a secure and reliable cloud infrastructure that provides the following features:

- Accurately identify and categorize storage assets
- Monitor capacity and performance from the storage consumer's view, including server, application, and department-level views
- Increase capacity forecasting precision by using historical growth metrics
- Reclaim unused storage to delay future purchases and improve utilization
- Optimize data placement based on historical usage patterns that can help lower the cost of storage by 50% per GB

Figure 2-4 shows an example of the Storage Insights dashboard.

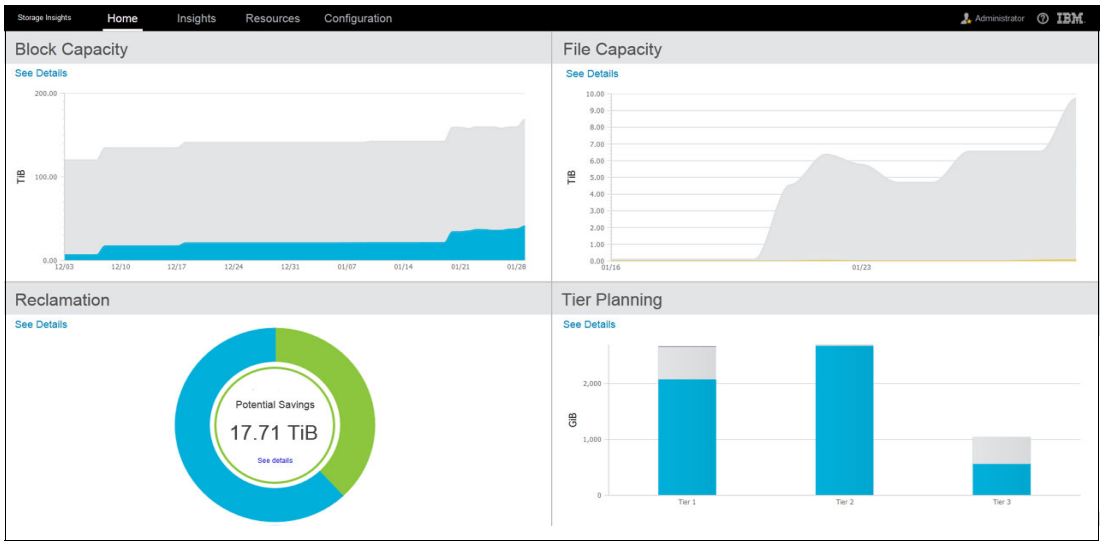

Figure 2-4 Storage Insights dashboard

For more information, see the following websites:

- IBM Spectrum Control Storage Insights product page

  http://www.ibm.com/systems/storage/spectrum/insights
- IBM Spectrum Control Storage Insights

  http://www.ibm.com/marketplace/cloud/analytics-driven-data-management/us/en-us

2.2.2 IBM Spectrum Protect

IBM Spectrum Protect is an integrated, industry-leading, data protection solution that covers your entire enterprise from small systems to virtual environments to enterprise applications. It provides a set of product features that allow you to design adaptive and comprehensive data protection solutions. It is a comprehensive data protection, recovery, and resiliency solution for software-defined, virtual, physical, and cloud environments.

Key capabilities

IBM Spectrum Protect has the following key capabilities:

- Protects virtual, physical, and cloud data with one solution.
- Native object storage support.
- Simplified user experience that simplifies backups by consolidating administration tasks under a single pane of glass.
- Reduction of backup and recovery infrastructure costs with inline data deduplication.
- Incremental forever backups with progressive incremental methods that back up only new or changed versions of files, which can greatly reduce data redundancy, network bandwidth, and storage pool consumption as compared to traditional methods.

- Supports many file systems and platforms, including IBM DB2®, Oracle, SAP, SAP HANA, Microsoft SQL Server, Microsoft Exchange, and IBM Lotus® Domino®. Support for extra platforms is available from IBM Business Partners through the partner ecosystem.
- Delivers full capabilities at significant savings for small and mid-sized organizations.

For an updated list of the available products and support in the IBM Spectrum Protect family, see the following website:

http://www.ibm.com/software/products/en/spectrum-protect

Benefits

IBM Spectrum Protect provides a single point of control and administration for storage management needs.

IBM Spectrum Protect data deduplication

Inline data deduplication improves storage utilization, and eliminates the need for a dedicated data deduplication appliance. It provides the following benefits:

- Significant scalability increase with up to 100 TB daily deduplication ingest throughput including ingest, deduplication, replication, and maintenance.[1]
- Improved efficiency with inline data deduplication reduces storage requirements for server-side data deduplication. New, efficient, server-side inline data deduplication in IBM Spectrum Protect 7.1.3 and later allows the use of inexpensive backup storage pool disks.
- No need for identify and reclamation processing and a reduced need for offline reorganization.
- Deduplicated storage pools can natively use on-premises object storage or off-premises cloud-based object storage.

When the IBM Spectrum Protect server receives data from a client, the server identifies duplicate data extents and stores unique instances of the data extents in a directory-container storage pool. If the same byte pattern occurs many times, data deduplication greatly reduces the amount of data that must be stored or transferred. In addition to whole files, IBM Spectrum Protect can also deduplicate parts of files that are common with parts of other files.

IBM Spectrum Protect provides the following types of data deduplication:

- Server-side data deduplication

  The server identifies duplicate data extents and moves the data to a directory-container storage pool. The server-side process uses inline data deduplication, where data is deduplicated at the same time that the data is written to a directory-container storage pool. Deduplicated data can also be stored in other types of storage pools. With inline data deduplication on the server, the need for reclamation is eliminated.

- Client-side data deduplication

  With this method, processing is distributed between the server and the client during a backup process. The client and the server identify and remove duplicate data to save storage space on the server. In client-side data deduplication, only compressed, deduplicated data is sent to the server. The server stores the data in the compressed format that is provided by the client. Client-side data deduplication can reduce the amount of data that is sent over the local area network (LAN). It can also reduce processing power and time that is required to remove duplicate data on the server. It can also improve database performance because the client-side data deduplication is also inline. You can combine both client-side and server-side data deduplication in the same production

[1] IBM internal tests. Includes ingest, deduplication, replication, and maintenance.

environment. The ability to deduplicate data on either the client or the server provides flexibility in terms of resource utilization, policy management, and data protection.

- Progressive incremental backup

  In a progressive incremental backup process, the IBM Spectrum Protect server monitors client activity and backs up any files that have changed since the initial full backup. Entire files are backed up so that the server does not need to reference base versions of the files. This backup technique eliminates the need for multiple full backups of client data, thus saving network resources and storage space.

Server-side compression

With IBM Spectrum Protect 7.1.5, you can increase the amount of available space in a storage pool by enabling inline compression. When you enable inline compression, data is compressed as it is written to a storage pool. Data compression is available for directory-container and cloud-container storage pools.

Inline compression works on non-encrypted data in both cloud and directory container pools. It has negligible impact on performance.

IBM Spectrum Protect Operations Center

IBM Spectrum Protect Operation Center is a light-weight management application that offers the daily dashboard and management interface for the IBM Spectrum Protect servers. It provides overview, visibility, automation, and control over the environment, and enables administrators to manage multiple IBM Spectrum Protect servers, including server diagnostics and email reporting on a single window as shown in Figure 2-5.

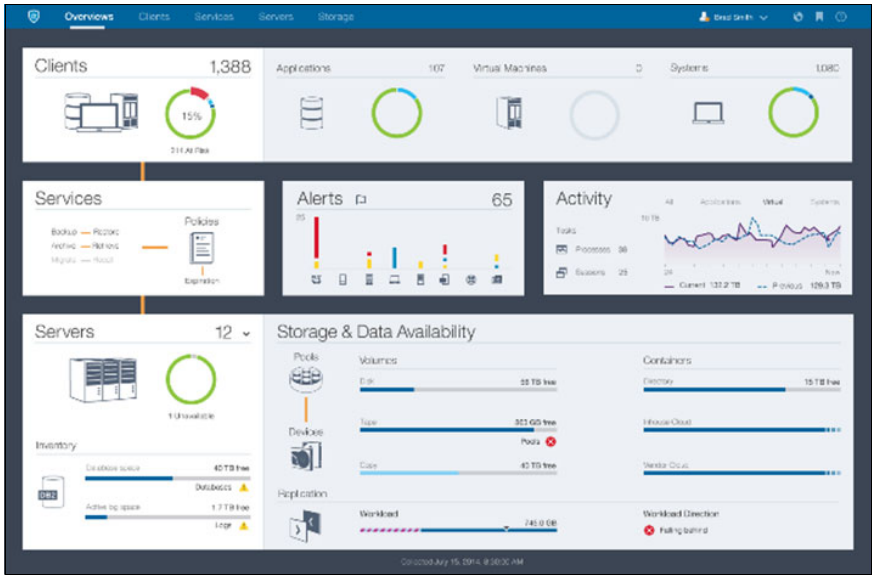

Figure 2-5 Operation Center

The Operation Center follows a hub-spoke model, with the first IBM Spectrum Protect server connected to it acting as the hub server. This system in turn connects to the spoke servers to query information and run commands. There is a correlation between the version of the Operation Center and the version of the hub server as new functions are introduced over time, requiring updates on the IBM Spectrum Protect servers.

The IBM Spectrum Protect Operation Center can be deployed on the system that hosts the primary IBM Spectrum Protect server or on, for example, a virtual machine.

For information about compatibility, see the following website:

http://www.ibm.com/support/docview.wss?uid=swg21640917

Also, explore the IBM Spectrum Protect Operation Center live demonstration environment at the following website:

https://demo.tsm.ibmserviceengage.com:11090/TSMLiveDemo

For information about Operation Center hardware and software requirements, see the following website:

http://www.ibm.com/support/docview.wss?uid=swg21653418

IBM Spectrum Protect for Virtual Environments

IBM Spectrum Protect for Virtual Environments (also known as Data Protection for VMware) protects VMware and Microsoft Hyper-V virtual machines.

It can be quickly and easily installed on virtual machines or physical hosts. The web-based portal provides self-service file recovery. No software is required on the user's system and the user does not need to know about IBM Spectrum Protect.

IBM Spectrum Protect provides the following benefits:

- Improves efficiency with deduplication, incremental "forever" backup, and other advanced IBM technology to help reduce costs
- Simplifies backups and restores for VMware with a practical interface that you can access from within VMware vCenter or a self-service restore portal
- Integrates with VMware snapshots and OpenStack clouds for faster, more frequent backups of your most critical virtual machines
- Flexible recovery and copy options from image-level backups give you the ability to perform recovery at the file, mailbox, database object, volume, or VM image level by using a single backup of a VMware image
- Eliminates overhead with optimized virtual machine backup by supporting VMware vStorage APIs for Data Protection and Microsoft Hyper-V technology, which simplifies and optimizes data protection.
- Instant access and restore that enables the IBM Spectrum Protect server to act as a virtual datastore
- Object-level recovery from image-level backups (MS Exchange, MS SQL Server, and files)
- The ability to use all IBM Spectrum Protect enterprise-class features including Operations Center

Figure 2-6 shows the file-level recovery capabilities of IBM Spectrum Protect for Virtual Environments.

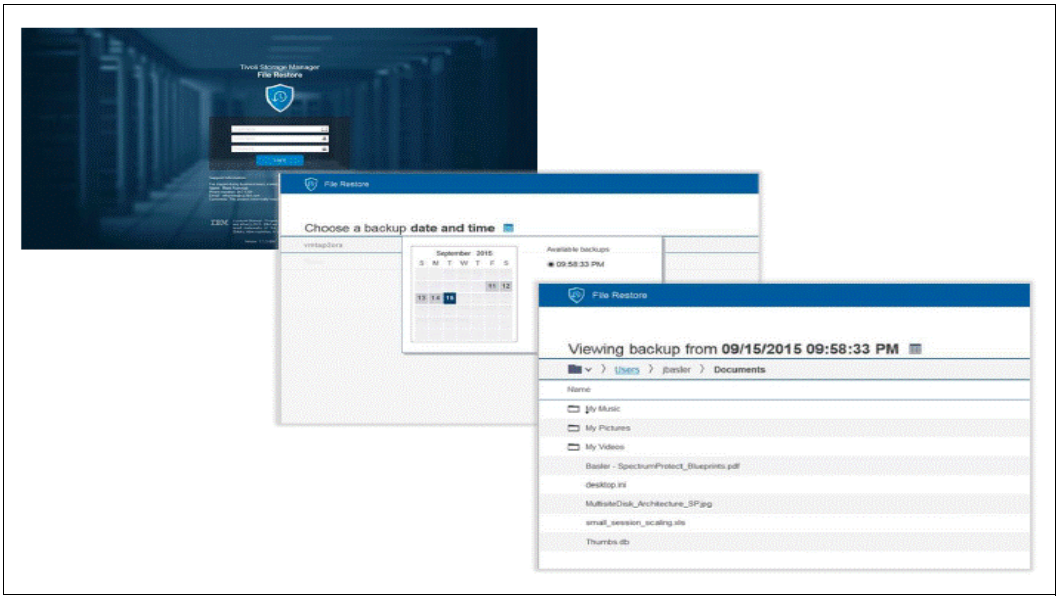

Figure 2-6 Web-based portal for self-service file recovery

For more information, see the following website:

http://www.ibm.com/software/products/en/spectrum-protect-for-virtual-environments

Support for IBM Cleversafe as a target

With IBM Spectrum Protect Version 7.1.5, you can configure cloud-container storage pools to use Cleversafe, an IBM object storage system, to back up data.

The Cleversafe object storage system uses the Simple Storage Service (S3) protocol to communicate with a cloud server. The Cleversafe object storage system can help you simplify storage management, improve server performance, and secure data by using encryption.

For more information, see the version 7.1.5 information in the IBM Knowledge Center:

http://www.ibm.com/support/knowledgecenter/SSGSG7_7.1.5/tsm/welcome.html

IBM Spectrum Protect Snapshot

IBM Spectrum Protect Snapshot (formerly Tivoli Storage FlashCopy Manager) delivers high levels of protection for key applications and databases by using advanced integrated application snapshot backup and restore capabilities.

IBM Spectrum Protect Snapshot provides support for IBM Spectrum Scale, IBM Storwize family, IBM DS8000, IBM XIV Storage systems, storage managed with IBM SAN Volume Controller and Spectrum Virtualize, and various virtual environments. It also manages application-aware snapshots on EMC, Hitachi, NetApp, and IBM and VSS-compatible Windows storage.

Traditional backups of files, databases, and similar objects that are sent across networks (LAN or SAN) to backup servers provide a degree of recoverability. However, the speed of backup and recovery is constrained by resource capacity (server, network, storage, and so on). The ability to make point-in-time snapshots with limited impact to the application while rapidly recovering large data sets by reverting to a snapshot greatly reduces the recovery time for these data sets.

The snapshots that are captured by IBM Spectrum Protect Snapshot can be retained as backups on local disk. With optional integration with IBM Spectrum Protect, you can use the full range of advanced data protection and data reduction capabilities such as data deduplication, progressive incremental backup, hierarchical storage management, and centrally managed policy-based administration as shown in Figure 2-7.

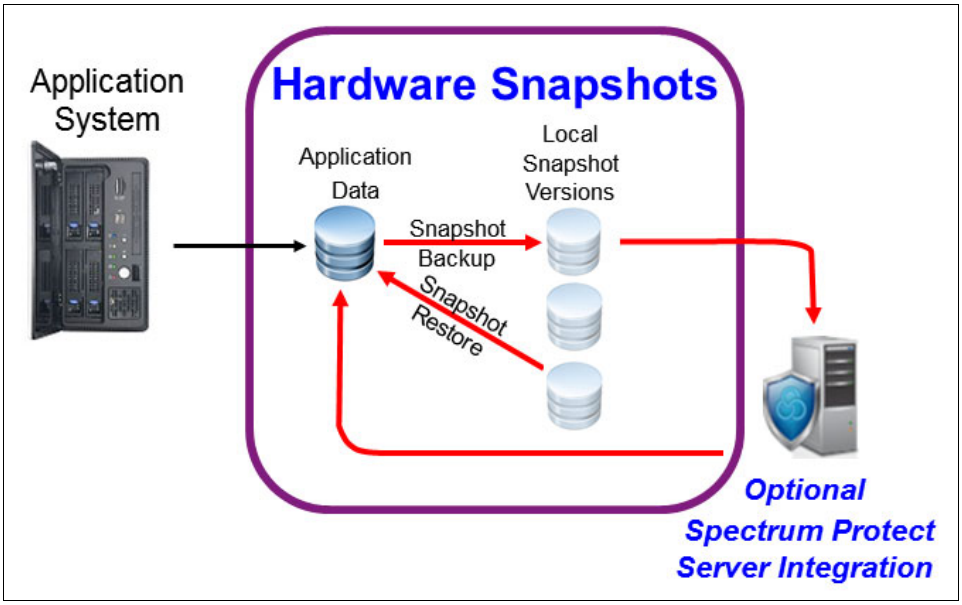

Figure 2-7 IBM Spectrum Protect Snapshot storage snapshot capabilities

Because a snapshot operation typically takes much less time than the time for a tape backup, the window during which the application must be aware of a backup can be reduced. This reduction facilitates more frequent backups, which reduces the time that is spent performing forward recovery through transaction logs, increases the flexibility of backup scheduling, and eases administration.

For more information, see the following website:

`http://www.ibm.com/software/products/en/spectrum-protect-snapshot`

For hardware and software requirements for IBM Spectrum Protect Snapshot, see the following website:

`http://www.ibm.com/support/docview.wss?uid=swg21427692`

IBM Spectrum Protect Blueprints

Blueprints can be thought of as the instructions for how to build a physical appliance. A blueprint consists of a prescriptive “cookbook” document that describes up to three reference architectures in detail (including IBM hardware model numbers and configuration requirements) and scripts to help with deployment.

Blueprints help improve time to value with faster server deployment. They can reduce implementation time from hours to as little as 10 minutes with systematic design and build instructions that align with software-defined data protection.

The following platforms are supported:

- Linux and Microsoft Windows
- IBM AIX® on POWER®
- Linux on POWER, big endian

The following storage is supported:

- IBM Storwize V7000, V5000, and V3700
- IBM Elastic Storage Server GL4, GL2 (powered by IBM Spectrum Scale)

Blueprints are available for small (up to 6 TB), medium (6 - 20 TB), and large (20 - 100 TB) environments.

For more information, see the following website:

http://ibm.co/1MnTQoD

2.3 SDS data plane

The data plane encompasses the infrastructure where data is processed. It consists of all basic storage management functions that can be started and managed by the control plane, which include:

- Virtualization
- RAID protection
- Tiering
- Copy services (remote, local, synchronous, asynchronous, and point-in-time)
- Encryption
- Data deduplication

The data plane is the interface to the hardware infrastructure where the data is stored. It provides a complete range of data access possibilities, spanning traditional access methods like block I/O (for example, iSCSI) and File I/O (POSIX compliant), to object-storage and Hadoop Distributed File System (HDFS).

This section describes the IBM software product offerings that provide the building blocks for the SDS data plane:

- IBM Spectrum Virtualize
- IBM Spectrum Accelerate™
- IBM Spectrum Scale
- IBM Spectrum Archive™

2.3.1 IBM Spectrum Virtualize

IBM Spectrum Virtualize provides industry-leading storage virtualization that enhances existing storage to improve resource utilization and productivity to achieve a simpler, more scalable, and cost efficient IT infrastructure. IBM SAN Volume Controller delivers the functions of IBM Spectrum Virtualize.

SAN Volume Controller, part of the IBM Spectrum Storage family, has been improving infrastructure flexibility and data economics for more than 10 years. SAN Volume Controller's innovative data virtualization capabilities provide the foundation for the entire IBM Storwize family. When using SAN Volume Controller integrated into the IBM Virtual Storage Center, you have SDS that spans both the control plane and the data plane.

For more information, see Chapter 4, "IBM SAN Volume Controller and IBM Storwize family" on page 81.

Key capabilities

The highly scalable storage virtualization system provides common functions, management, and mobility across heterogeneous storage types. Spectrum Virtualize has the following key capabilities:

- Pools storage from multiple systems
- Compresses data with high performance IBM Real-time Compression for extraordinary efficiency
- Enables nondisruptive movement of data among storage systems
- Supports high availability, multi-site configurations

Benefits

The sophisticated virtualization, management, and functions of Spectrum Virtualize provide these storage benefits, among others:

- Improves storage utilization up to 100%
- Supports up to 5x as much data in the same physical space
- Simplifies management of heterogeneous storage systems
- Enables rapid deployment of new storage technologies for greater return on investment (ROI)
- Improves application availability with virtually zero storage-related outages

The SAN Volume Controller combines software and hardware into a comprehensive, modular appliance that uses symmetric virtualization.

Symmetric virtualization is achieved by creating a pool of managed disks (MDisks) from the attached storage systems. Those storage systems are then mapped to a set of volumes for use by the attached host systems. System administrators can view and access a common pool of storage on the SAN. This function helps administrators to use storage resources more efficiently and provides a common base for advanced functions. Figure 2-8 shows the Spectrum Virtualize functions.

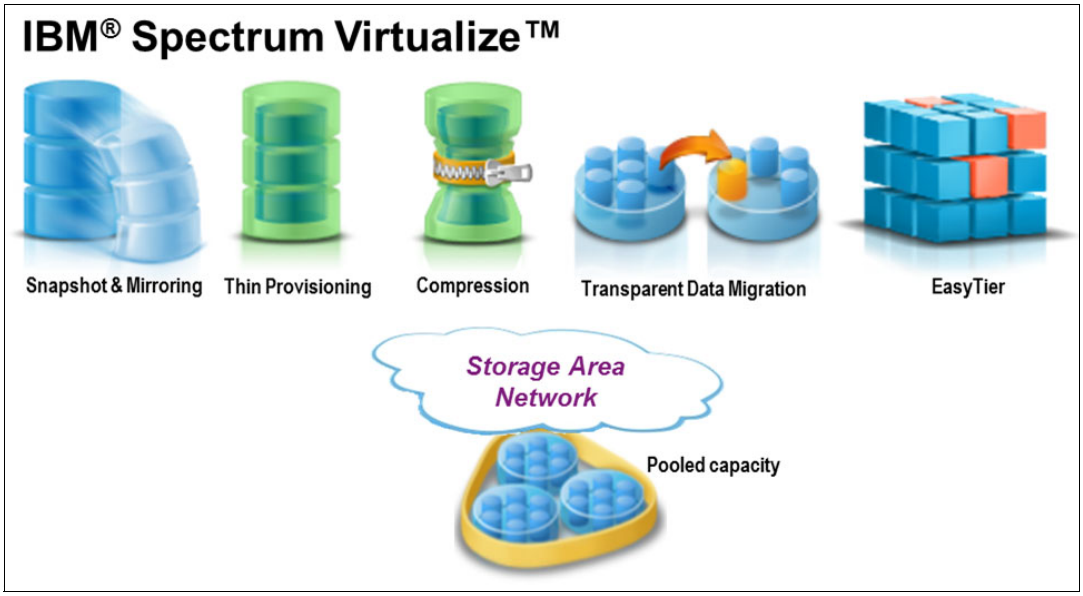

Figure 2-8 IBM Spectrum Virtualize functions

Table 2-2 describes the Spectrum Virtualize features and benefits.

Table 2-2 Spectrum Virtualize features and benefits

Feature	Benefits
Single point of control for storage resources	► Designed to increase management efficiency ► Designed to help support business application availability
Pools the storage capacity of multiple storage systems on a SAN	► Helps you manage storage as a resource to meet business requirements, not just as a set of boxes ► Helps administrators better deploy storage as required beyond traditional “SAN islands” ► Can help increase utilization of storage assets ► Insulates applications from physical changes to the storage infrastructure
Clustered pairs of IBM SAN Volume Controller data engines	► Highly reliable hardware foundation ► Designed to avoid single points of hardware failure
IBM Real-time Compression	► Increases effective capacity of storage systems up to five times, helping to lower costs, and floor-space, power, and cooling requirements ► Can be used with a wide range of data, including active primary data, for dramatic savings ► Hardware compression acceleration helps transform the economics of data storage
Innovative and tightly integrated support for flash memory	► Designed to deliver ultra-high performance capability for critical application data ► Move data to and from flash memory without disruption; make copies of data onto hard disk drives (HDDs)
Support for IBM FlashSystem	Enables high performance for critical applications with IBM MicroLatency, coupled with sophisticated functions
Easy-to-use IBM Spectrum Storage family management interface	► Single interface for storage configuration, management, and service tasks regardless of storage vendor ► Helps administrators use their existing storage assets more efficiently
IBM Storage Mobile Dashboard	Provides basic monitoring capabilities to securely check system health and performance
Dynamic data migration	► Migrate data among devices without taking applications that use that data offline ► Manage and scale storage capacity without disrupting applications
Manage tiered storage	Helps balance performance needs against infrastructure costs in a tiered storage environment
Advanced network-based copy services	► Copy data across multiple storage systems with IBM Spectrum Protect Snapshot ► Copy data across metropolitan and global distances as needed to create high-availability storage solutions
Integrated Bridgeworks SANrockIT technology for IP replication	► Optimize use of network bandwidth ► Reduce network costs or speed replication cycles, improving the accuracy of remote data

Feature	Benefits
IBM HyperSwap	► Provide highly available, concurrent access to a single copy of data from data centers up to 300 km (186 miles) apart ► Enable nondisruptive storage and virtual machine mobility between data centers
Thin provisioning and snapshot replication	► Dramatically reduce physical storage requirements by using physical storage only when data changes ► Improve storage administrator productivity through automated on-demand storage provisioning
IBM Spectrum Protect Snapshot application-aware snapshots	► Performs near-instant application-aware snapshot backups, with minimal performance impact for IBM DB2, Oracle, SAP, VMware, Microsoft SQL Server, and Microsoft Exchange ► Provides advanced, granular restoration of Microsoft Exchange data

Virtualizing storage with SAN Volume Controller helps make new and existing heterogeneous storage arrays more effective by including many functions that are traditionally deployed within disk array systems. By including these features in a virtualization system, SAN Volume Controller standardizes functions across virtualized storage for greater flexibility and potentially lower costs.

Figure 2-9 shows how SAN Volume Controller stretch virtual volume standardizes heterogeneous storage across data centers.

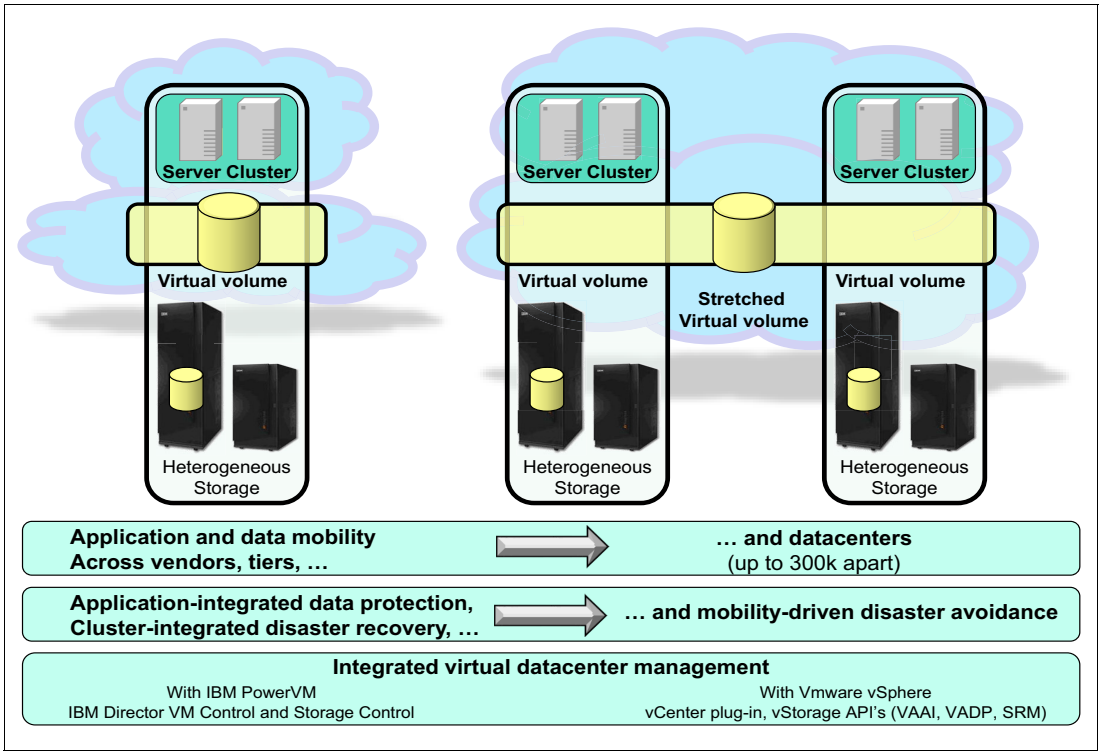

Figure 2-9 SAN Volume Controller stretches virtual volume across heterogeneous storage

IBM Spectrum Virtualize and SAN Volume Controller support attachment to servers through Fibre Channel (FC) protocols and Internet Small Computer System Interface (iSCSI) protocols over IP networks at 1 Gbps and 10 Gbps speeds. These configurations can help

reduce costs and simplify server configuration. Spectrum Virtualize also supports Fibre Channel over Ethernet (FCoE) protocol.

IBM Spectrum Virtualize and SAN Volume Controller combines hardware and software in an integrated, modular solution that is highly scalable. An I/O group is formed by combining a redundant pair of storage engines that are based on IBM System x server technology. Highly available I/O groups are the basic configuration element of a Spectrum Virtualize cluster.

The configuration flexibility means that your implementation can start small and grow with your business to manage large storage environments. The scalable architecture and tight integration enable your business to take advantage of the high throughput of flash drives. This configuration supports high performance for critical applications.

IBM Spectrum Virtualize also includes the IBM System Storage Easy Tier function, which helps improve performance at lower cost through more efficient use of flash drives. The Easy Tier function automatically identifies highly active data within volumes and moves only the active data to a flash drive. It targets flash use to the data that benefits the most, which delivers the maximum benefit even from small amounts of flash capacity, and helps move critical data to and from flash drives as needed without disruption to applications.

IBM Spectrum Virtualize helps increase the amount of storage capacity that is available to host applications by pooling the capacity from multiple disk systems within the SAN. In addition, it combines various IBM technologies that include thin provisioning, automated tiering, storage virtualization, IBM Real-time Compression, clustering, replication, multiprotocol support, and an updated graphical user interface (GUI). Together, these technologies enable a IBM Spectrum Virtualize and SAN Volume Controller to deliver exceptional storage efficiency.

IBM Spectrum Virtualize and SAN Volume Controller includes a dynamic data migration function to move data from one storage system to another while still maintaining access to that data. The Volume Mirroring function stores two copies of a volume on different storage systems. This configuration helps improve application availability during a failure or disruptive maintenance to an array or disk system. The controller's stretched cluster configuration automatically uses whichever copy of the data remains available.

With the stretched cluster, administrators can apply a single set of advanced network-based replication services that operate in a consistent manner. This set is applied regardless of the type of storage that is being used. The Metro Mirror and Global Mirror functions operate between systems at different locations. They help create copies of data for use during a catastrophic event at a data center. For even greater flexibility, Metro Mirror and Global Mirror also support replication between SAN Volume Controller systems and IBM Storwize V7000 Unified systems.

The IBM FlashCopy function stretched cluster quickly creates a copy of active data that can be used for backup purposes or for parallel processing activities. This capability enables disk IBM Spectrum Virtualize and SAN Volume Controller Enhanced Stretched Cluster with VMware backup copies to be used to recover almost instantly from corrupted data, which significantly speeds application recovery.

What is new with IBM Spectrum Virtualize Software V7.6

IBM Spectrum Virtualize software delivers a powerful solution for storage efficiency in environments where IBM is the primary storage. When used in combination with SVC, Storwize V7000 (Unified), Storwize V5000, Storwize V3700, and Storwize V3500, IBM Spectrum Virtualize provides the benefits of storage virtualization and advanced storage capabilities, including HyperSwap, VVOL support, FlashCopy, replication, Real-time Compression, Easy Tier, and Encryption.

Spectrum Virtualize 7.6 includes the following new functionalities:

- Encryption

  For improved data security and confidentiality, encryption is now available for both internal capacities of SVC, Storwize V7000 and FlashSystem V9000 and externally virtualized capacities of Storwize V7000, SVC, and FlashSystem V9000. Encryption can be applied to virtualized storage arrays, even if the virtualized array does not have encryption capabilities. Encrypted volumes are transparent to applications, easing implementation and operation.

  For supported externally virtualized storage configurations, data encryption can be enabled by the Storwize V7000 or SVC controllers. The data is encrypted and sent over the storage network to the virtualized storage devices. Encryption by Spectrum Virtualize is transparent to hosts and applications, and the data is encrypted and protected across the network between the Storwize and SVC controllers and the storage devices.

- Distributed RAID

  For higher availability, Distributed RAID is now available for internal capacities, offering up to 10 times faster RAID array rebuild times. As drives grow in capacity, rebuild times increase, increasing the possibility of a second drive failure and resulting in a loss of data. Distributed RAID addresses this issue by distributing the data across a larger number of physical drives, reducing the load on each individual drive from rebuild activity. Distributed RAID can also deliver increased performance because data can be read from/written to more drives for each I/O. Distributed RAID volumes are transparent to applications for simplified implementation and operation.

- IP quorum base support

  Current Stretched Cluster (SVC) and HyperSwap (SVC, Storwize V7000, and Storwize V5000) implementations require FC storage on a third site to cope with tie-break situations during an inter-site link failing. IP quorum base support enables the use of a lower-cost IP network-attached host as quorum disk for simplified implementation and operation.

- HyperSwap GUI support

  Current local HyperSwap is functionally complete, but requires complex configuration through the existing command-line interface. This new release of HyperSwap offers a new command-line interface that simplifies the set-up process to a handful of commands, and also adds the ability to configure and manage local HyperSwap through the GUI.

- VMWare VVOL integration

  Migration of space efficient volumes between storage containers maintains the space efficiency of volumes and delivers resiliency for VMs in case the volumes start running out of space. Cloning a virtual machine will achieve a full independent set of virtual volumes.

- Increase maximum quantity of iSCSI hosts

  For more scalable iSCSI host support, version 7.6 increases the maximum number of iSCSI host sessions from 256 to 1024 per node, offering more fanout for each SVC I/O group.

- Integrated Comprestimator

  Real-time Compression is a key differentiator of the Storwize Family, and the Comprestimator is its key sizing tool to estimate how much capacity savings the customer can expect. The Comprestimator can analyze the patterns of the actual customer data, and estimate the compressibility of data per volume. Currently, storage administrators have to get server administrator collaboration and authorization to install Comprestimator on servers, potentially on hundreds of servers. The integration of Comprestimator within Spectrum Virtualize code eases the process of estimating capacity savings by having this sizing tool run natively in SVC, Storwize V7000, and FlashSystem V9000, avoiding the

need to install Compresti mator on hundreds of servers and enabling estimates of Real-time Compression effectiveness from a central console.

IBM Spectrum Virtualize and VersaStack solution by Cisco

VersaStack, an IBM and Cisco integrated infrastructure solution, combines computing, networking, and storage into a single integrated system. It combines the Cisco Unified Computing System (Cisco UCS) Integrated Infrastructure with IBM Spectrum Virtualize, which includes IBM FlashSystem V9000 and IBM Storwize storage offerings. This combination allows quick deployment and rapid time to value for the implementation of modern infrastructures.

For more information about VersaStack solutions, visit this webpage:

http://www.ibm.com/systems/storage/disk/cisco-versastack.html

2.3.2 IBM Spectrum Accelerate

A key member of the IBM Spectrum Storage portfolio, IBM Spectrum Accelerate is a highly flexible, software-defined storage solution that enables rapid deployment of block data storage services for new and traditional workloads, both on and off premises. Spectrum Accelerate allows you to run hotspot-free, grid-scale software that deploys on-premises on x86 commodity servers and on the optimized XIV Storage System, and off-premises as a public cloud service on SoftLayer. It offers proven grid- scale technology, mature features, and ease of use, and is already deployed on over 100,000 servers worldwide.

IBM Spectrum Accelerate delivers a design that was made for the cloud, ensuring hotspot and tuning-free service levels across mixed workloads to help meet the unpredictability of today's cloud world. It delivers a single management experience across software defined infrastructure with the Hyper-Scale Manager GUI. It also helps cut costs through reduced administration effort and training, reduced procurement, standardization of data center storage hardware operations and services, and licensing flexibility that enables cost-efficient cloud building.

Figure 2-10 shows how straightforward scaling is by building a storage grid with Spectrum Accelerate.

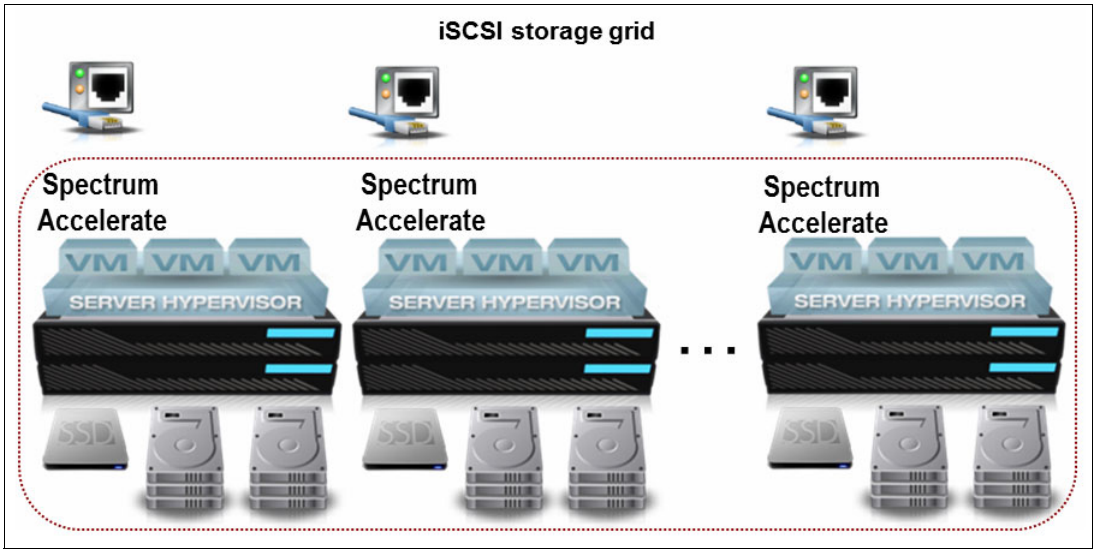

Figure 2-10 Spectrum Accelerate iSCSI storage grid

Key capabilities

Spectrum Accelerate gives organizations these capabilities:

- Enterprise cloud storage in minutes, using commodity hardware
- Hotspot-free performance and QoS without any manual or background tuning needed
- Advanced remote replication, role-based security, and multi-tenancy
- Deploy on-premises or on the cloud
- Management of dozens of petabytes
- Best in class VMware and OpenStack integrations

Spectrum Accelerate runs as a virtual machine on vSphere ESXi hypervisor. This configuration enables you to build a server-based SAN from commodity hardware that includes x86 servers, Ethernet switches, solid-state drives (SSDs), and direct-attached, high-density disks. Spectrum Accelerate essentially acts as an operating system for your self-built SAN storage, grouping virtual nodes and spreading the data across the entire grid.

IBM Spectrum Accelerate release 11.5.3 manages up to 15 nodes in a grid. It provides a single point of management of up to 144 grids of 15 nodes connected through Hyper-Scale Manager, for a total of 2160 nodes. It supports 6 TB HDDs for 50% more capacity in the same footprint than 4 TB drives and support for 600 GB 10K rpm drives.

IBM Hyper-Scale mobility over iSCSI networks provides the following capabilities:

- Delivers online, transparent, nondisruptive data mobility between IBM Spectrum Accelerate storage clusters
- Helps maximize utilization and effective performance balancing across multiple clusters and facilitate aggressive thin-provisioning

Spectrum Accelerate allows you to deploy storage services flexibly across different delivery models, including customer-choice hardware, existing data center infrastructure, and IBM storage systems.

Benefits

Spectrum Accelerate provides the following benefits:

- Cost reduction by delivering hotspot-free storage to different deployment models, on and off premises, paying less for the same capacity, acquiring less hardware, reducing maintenance outlay, and minimizing administrative effort
- Increased operational agility through easy cloud building, faster provisioning, small capacity increments, and licensing flexibility that enables you to easily repurpose servers
- Rapid response through enterprise-class storage availability, data protection, and security for new and traditional workload needs in the data center and other sites while flexibly balancing capital and operational expenses
- Simplified management with seamless management across 144 virtual systems, the XIV Storage System, or the infrastructure of your off- premises services provider
- An administrator can deploy the IBM Spectrum Accelerate storage solution on existing virtualized servers in minutes, enabling a tuning-free storage environment to be built quickly with outstanding ease of management and agility

Table 2-3 describes Spectrum Accelerate features with their associated benefits.

Table 2-3 Spectrum Accelerate features and benefits

Feature	Benefit
Performance	► Ensures even data distribution through massive parallelism and automatic load balancing, including capacity add ► Distributed cache
Reliability and Availability	► Grid redundancy maintains two copies of each 1-MB data partition with each copy on a different virtual machine (VM), proactive diagnostics, fast and automatic rebuilds, event externalization ► Advanced monitoring; network monitoring; disk performance tracking/reporting; data center monitoring; shared monitoring for some components; data and graphical reports on I/O, usage, and trends ► Self-healing, which minimizes the rebuild process by rebuilding only actual data ► Automated load balancing across components with minimized risk of disk failure due to rapid return to redundancy
Management	► Intuitive GUI: Scales to over 40 PB with IBM Hyper-Scale Manager; extensive CLI; RESTful API; mobile app support with push notifications; multi-tenancy with quality of service by domain, pool, or host
Cloud automation and self-service	► OpenStack; VMware vRealize Orchestrator through IBM Spectrum Control Base
Snapshot management	► Space efficient snapshots: Writable, snapshot of snapshot, restore from snapshot, snapshots for consistency groups, mirroring
Thin provisioning and space reclamation	► Thin provisioning per pool, thick-to-thin migration; VMware, Microsoft, Symantec space reclamation support
Mirroring	► Synchronous/asynchronous; volumes and consistency groups, recovery point objective (RPO) of seconds; online/offline initialization; failover/failback; mirroring across platforms, including with XIV Storage System
Security	► Role-based access management, multi-tenancy, iSCSI Challenge Handshake Authentication Protocol (CHAP), and auditing; integrates with Lightweight Directory Access Protocol (LDAP) and Microsoft Active Directory servers

For more information about IBM Spectrum Accelerate, see the following resources:

- IBM Spectrum Accelerate

  http://www.ibm.com/systems/storage/spectrum/accelerate/
- *IBM Spectrum Accelerate Deployment, Usage, and Maintenance*, SG24-8267
- *Deploying IBM Spectrum Accelerate on Cloud*, REDP-5261

2.3.3 IBM Spectrum Scale

IBM Spectrum Scale is a proven, scalable, high-performance file management solution that is based on IBM GPFS. IBM Spectrum Scale provides world-class storage management with extreme scalability, flash accelerated performance, and automatic policy-based storage tiering from flash to disk, then to tape. IBM Spectrum Scale reduces storage costs up to 90% while improving security and management efficiency in cloud, big data, and analytics environments.

First introduced in 1998, this mature technology enables a maximum volume size of 8 YB, a maximum file size of 8 EB, and up to 18.4 quintillion (two to the 64th power) files per file system. IBM Spectrum Scale provides simplified data management and integrated information lifecycle tools as software-defined storage for cloud, big data, and analytics. It introduces enhanced security, flash accelerated performance, and improved usability. It also provides capacity quotas, access control lists (ACLs), and a powerful snapshot function.

Key capabilities

IBM Spectrum Scale adds elasticity with the following capabilities:

- Global namespace with high performance access scales from departmental to global
- Automated tiering, data lifecycle management from flash (6x acceleration) to tape (10x savings)
- Enterprise ready: Data security (encryption), availability, reliability, large scale proven
- Open: POSIX compliant, integrated with OpenStack components and Hadoop

Benefits

IBM Spectrum Scale provides these benefits:

- Improves performance by removing data-related bottlenecks
- Automated tiering, data lifecycle management from flash (acceleration) to tape (savings)
- Lowers cost by eliminating duplicate data
- Enables sharing of data across multiple applications
- Reduces cost per performance by placing data on most applicable storage (flash to tape)

IBM Spectrum Scale is part of the IBM market-leading software defined storage family and has the following characteristics:

- As a software-only solution: Runs on virtually any hardware platform and supports almost any block storage device. IBM Spectrum Scale runs on Linux (including Linux on IBM z Systems), AIX, and systems based on Windows.
- As an integrated IBM Elastic Storage Server solution: A bundled hardware, software, and services offering that includes installation and ease of management with a graphical user interface. Elastic Storage Server provides unsurpassed end-to-end data availability, reliability, and integrity with unique technologies that include IBM Spectrum Scale RAID.
- As a cloud service: IBM Spectrum Scale delivered as a service provides high performance, scalable storage, and integrated data governance for managing large amounts of data and files in the IBM SoftLayer cloud.

IBM Spectrum Scale features enhanced security with native encryption and secure erase. It can increase performance by using server-side flash cache to increase I/O performance up to six times. IBM Spectrum Scale provides improved usability through data replication capabilities, data migration capabilities, Active File Management (AFM), File Placement Optimizer (FPO), and IBM Spectrum Scale Native RAID.

Figure 2-11 shows an example of the Spectrum Scale architecture.

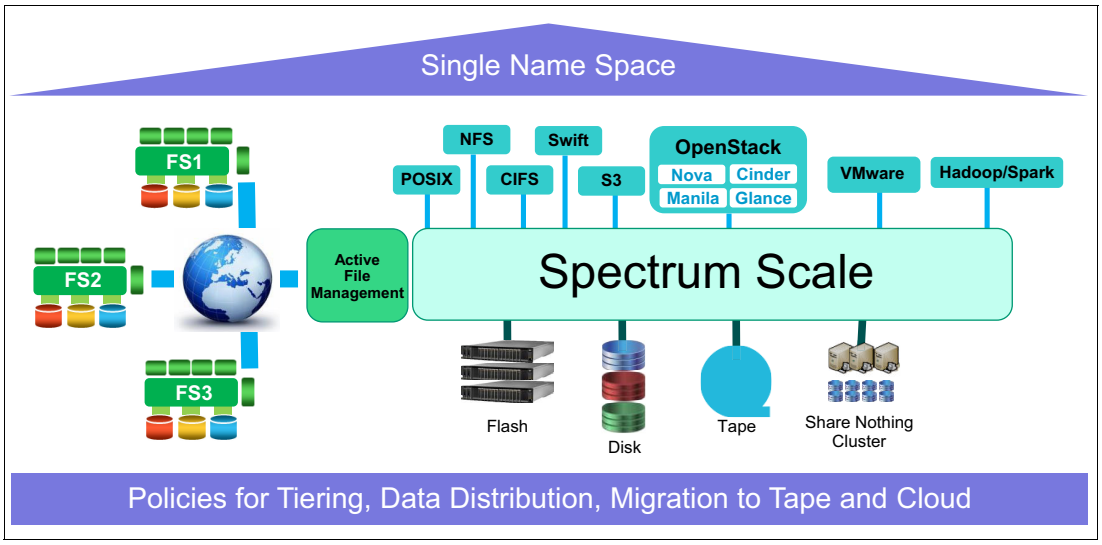

Figure 2-11 Spectrum Scale architecture

IBM Spectrum Scale is based around the following concepts:

- Storage pools
- File sets
- Policy engine
- Mirroring, replication, and migration capabilities
- Active File Management
- File Placement Optimizer
- IBM Spectrum Scale Native RAID
- Licensing

Storage pools

A storage pool is a collection of disks or arrays with similar attributes. It is an organizational structure that allows the combination of multiple storage locations that are of identical characteristics. There are three different types of storage pools:

System Pool The system pool contains file system metadata and can be used to store data. One system pool is needed per file system.

Data Pool A data pool is used to store file data. A data pool is optional.

External Pool An external pool is used to attach auxiliary storage, such as tape, to IBM Spectrum Scale. An external pool is optional.

File sets

IBM Spectrum Scale creates a single name space. Tools can provide a fine grained management of the directory structure. A file set acts like a partition of a file system, a subdirectory tree. File sets can be used for operations such as quotas or used in management policies. It is a directory tree that behaves like a “file system” within a file system and has the following characteristics:

- Part of the global namespace
- Can be linked and unlinked (like mount/unmount)

- Policy scan can be restricted to only scan file sets. This feature can be helpful when the file system has billions of files
- A file set can be assigned to a storage pool

There are two kinds of file sets:

Dependent file sets A dependent file set allows for a finer granularity of administration. It shares the inode space with another file set.

Independent file set An independent file set has a distinct inode space. An independent file set allows file set level snapshots, independent file scans, and enables advanced features like Active File Management.

Policy engine

The policy engine uses an SQL style syntax to query or operate on files based on file attributes. Policies can be used to migrate all data that has not been accessed in 6 months to less expensive storage or just used to query the contents of a file system. Management policies support advanced query capabilities, though what makes the policy engine most useful is the performance. The policy engine is capable of scanning billions of objects as shown in Table 2-4.

Table 2-4 Speed comparison for GPFS policy engine

Search through 1000000000 (1 billion) files	
`find`	~ 47 hours
GPFS policy engine	~ 5 hours

Table 2-4 shows the power of the GPFS policy engine. Although an average **`find`** across 1 billion files took approximately 47 hours, the GPFS policy engine could satisfy the request within five hours. The GPFS policy engine can also create a candidate list for backup applications to use to achieve a massive reduction in candidate identification time.

IBM Spectrum Scale has next generation availability with features that include rolling software and hardware upgrades. You can add and remove servers to adapt the performance and capacity of the system to changing needs. Storage can be added or replaced online, and you can control how data is balanced after storage is assessed.

Mirroring, replication, and migration capabilities

In IBM Spectrum Scale, you can replicate a single file, a set of files, or the entire file system. You can also change the replication status of a file at any time by using a policy or command. Using these capabilities, you can achieve a replication factor of two, which equals mirroring, or a replication factor of three.

A replication factor of two in Spectrum Scale means that each block of a replicated file is in at least two failure groups. A failure group is defined by the administrator and contains one or more disks. Each storage pool in a file system contains one or more failure groups. Failure groups are defined by the administrator and can be changed at any time. So when a file system is fully replicated, any single failure group can fail and the data remains online.

For migration, IBM Spectrum Scale provides the capability to add storage to the file system, migrate the existing data to the new storage, and remove the old storage from the file system. All of this can be done online without disruption to your business.

Active File Management

AFM enables the sharing of data across unreliable or high latency networks. With AFM, you can create associations between IBM Spectrum Scale clusters and define the location and flow of file data. AFM allows you to implement a single name space view across clusters, between buildings, and around the world.

AFM operates at the file set level. This means that you can create hundreds of AFM relationships in each file system. AFM is a caching technology though inode, and file data in a cache file set is the same as an inode and file data in any IBM Spectrum Scale file system. It is a “real” file that is stored on disk. The job of the cache is to keep the data in the file consistent with the data on the other side of the relationship.

AFM can be implemented in these modes:

- Read-Only (ro): Create a flash cache for heavily read data
- Local-Update (lu): Provide a development copy of data
- Single-Writer (sw): Collect data at a remote location
- Independent Writer (iw): Create a global interactive name space
- Asynchronous DR: Create asynchronous copies of file data

File Placement Optimizer

FPO is optimized for Hadoop in a shared-nothing deployment. FPO can use the local storage of the individual cluster members to form a high performance, resilient storage environment. Using policies, you can control the placement of data and replicas within the cluster to optimize for your access patterns.

IBM Spectrum Scale Native RAID

IBM Spectrum Scale Native RAID provides next generation performance and data security. IBM Spectrum Scale native RAID just a bunch of disks (JBOD) are directly attached to the systems running IBM Spectrum Scale software. This technology uses declustered RAID to minimize performance degradation during RAID rebuilds and provides extreme data integrity by using end-to-end checksums and version numbers to detect, locate, and correct silent disk corruption. An advanced disk hospital function automatically addresses storage errors and slow performing drives so your workload is not affected.

IBM Spectrum Scale native RAID is available with the IBM Power8 architecture in the IBM Elastic Storage Server.

Licensing

IBM Spectrum Scale V4.2 offers different editions so you only pay for the functions that you need:

- Express Edition contains the base IBM Spectrum Scale data access function
- Standard Edition includes the base function plus information lifecycle management (ILM), AFM, and either integrated multiprotocol support, which includes Ganesha NFS or traditional clustered NFS
- Advanced Edition includes encryption of data at rest, secure erase, asynchronous multisite disaster recovery, and all the features of Standard Edition

For each of these editions, you can choose an IBM Spectrum Scale license for Server, Client, and FPO.

For more information, see the following resources:

- IBM Spectrum Scale (IBM developerWorks)

  http://www.ibm.com/developerworks/servicemanagement/tc/gpfs/index.html
- IBM Spectrum Scale

  http://www.ibm.com/systems/storage/spectrum/scale/index.html
- IBM Spectrum Scale (IBM Knowledge Center)

  http://www.ibm.com/support/knowledgecenter/STXKQY/ibmspectrumscale_welcome.html
- IBM Spectrum Scale Wiki

  https://ibm.biz/BdFPR2
- IBM Elastic Storage Server

  http://www.ibm.com/systems/storage/spectrum/ess

2.3.4 IBM Spectrum Archive

IBM Spectrum Archive, a member of the IBM Spectrum Storage family, enables direct, intuitive, and graphical access to data stored in IBM tape drives and libraries by incorporating the Linear Tape File System (LTFS) format standard for reading, writing, and exchanging descriptive metadata on formatted tape cartridges. Spectrum Archive eliminates the need for additional tape management and software to access data. Spectrum Archive offers three software solutions for managing your digital files with the LTFS format: Single Drive Edition (SDE), Library Edition (LE), and Enterprise Edition (EE). With Spectrum Archive Enterprise Edition and Spectrum Scale, tape can now add savings as a low-cost storage tape tier.

Network attached unstructured data storage with native tape support using LTFS delivers the best mix of performance and lowest cost storage.

Key capabilities

Spectrum Archive options can support small, medium, and enterprise businesses with these advantages:

- Seamless virtualization of storage tiers
- Policy-based placement of data
- Single universal namespace for all file data
- Security and protection of assets
- Open, non-proprietary, cross platform interchange
- Integrated functionality with IBM Spectrum Scale

Benefits

IBM Spectrum Archive enables direct, intuitive, and graphical access to data stored in IBM tape drives and libraries. It does so by incorporating the LTFS format standard for reading, writing, and exchanging descriptive metadata on formatted tape cartridges. Spectrum Archive eliminates the need for additional tape management and software to access data.

Spectrum Archive takes advantage of the low cost of tape storage while making it easy to use. Spectrum Archive provides these benefits, among others:

- Access and manage all data in stand-alone tape environments as easily as though it were on disk
- Enable easy-as-disk access to single or multiple cartridges in a tape library

- Improve efficiency and reduce costs for long-term, tiered storage
- Optimize data placement for cost and performance
- Enable data file sharing without proprietary software
- Scalable and low cost

IBM Linear Tape File System

In 1997, IBM invented a new file system for files stored on tape cartridges, which was initially called *IBM Long Term File System*. This invention allowed it to use a tape drive as if it were a hard disk drive or a removable flash drive to store and access data. Because of the nature of the tape technology, the file access is still sequential.

In 2010, this technology was adopted by the Linear Tape-Open (LTO) consortium that defines the standard for all LTO technology specifications. For more information, see 8.1.1, “LTO technology” on page 178). The LTFS format specification and file system implementation were released on April 12, 2010 with the support of IBM, HP, Quantum, and the LTO Consortium. From this time, the file system was called LTFS.

LTFS is the first file system that works with LTO generation 7, 6, and 5 tape technology (or IBM TS1150 and TS1140 tape drives). It sets a new standard for ease of use and portability for open systems tape storage. With this application, accessing data that is stored on an IBM tape cartridge is as easy and intuitive as using a USB flash drive. Tapes are self-describing, and you can quickly recall any file from a tape without having to read the whole tape.

Further, any LTFS-capable system can read a tape that is created by any other LTFS-capable system regardless of the operating system and platform. Any LTFS-capable system can identify and retrieve the files that are stored on it. LTFS-capable systems have the following characteristics:

- Files and directories are displayed as a directory tree listing.
- More intuitive searches of cartridge and library content are now possible due to the addition of file tagging.
- Files can be moved to and from LTFS tape by using the familiar drag-and-drop metaphor common to many operating systems.
- Many applications that were written to use files on disk can now use files on tape without any modification.
- All standard File Open, Write, Read, Append, Delete, and Close functions are supported.

Spectrum Archive editions

As shown in Figure 2-12, Spectrum Archive is available in different editions that support small, medium, and enterprise businesses.

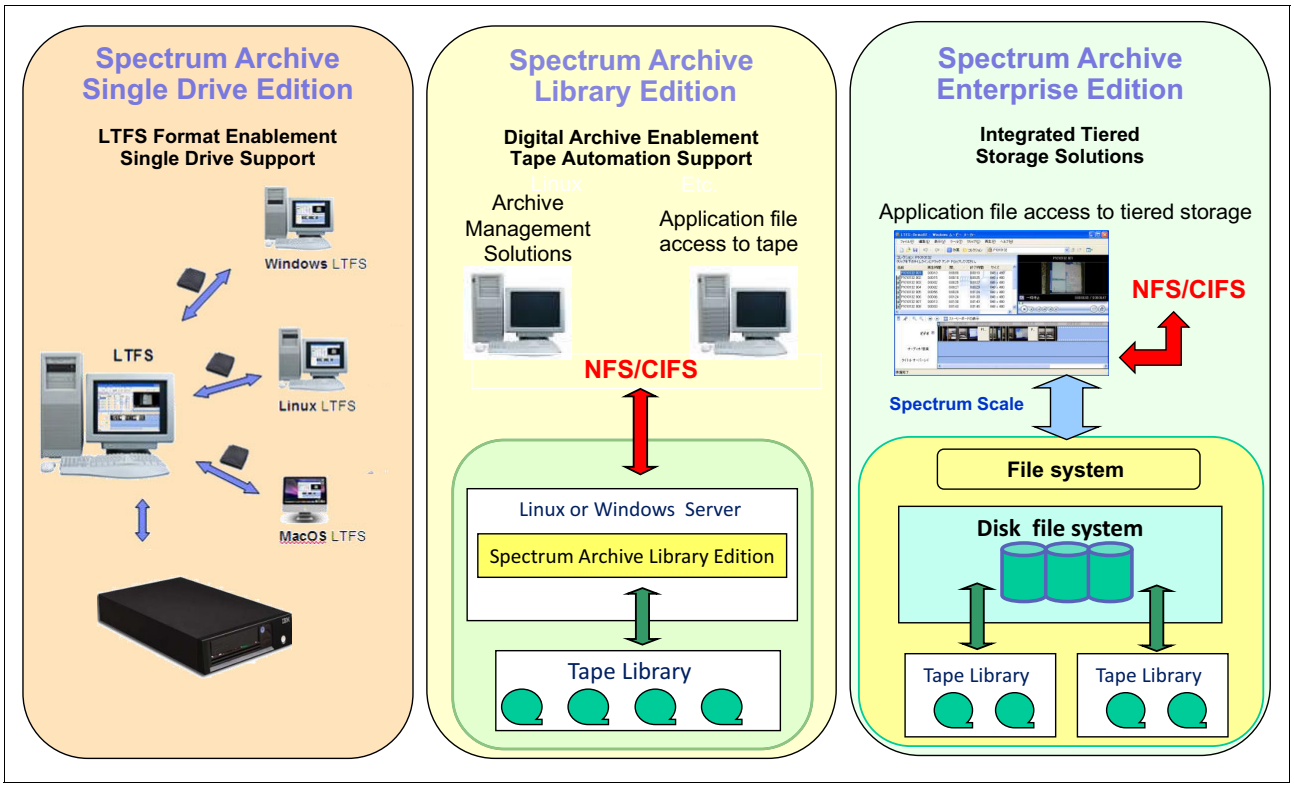

Figure 2-12 IBM Spectrum Archive Singe Drive Edition, Library Edition, and Enterprise Edition implementations

IBM Spectrum Archive Enterprise Edition

IBM Spectrum Archive EE gives organizations an easy way to use cost-effective IBM tape drives and libraries within a tiered storage infrastructure. By using tape libraries instead of disks for Tier 2 and Tier 3 data storage (data that is stored for long-term retention), organizations can improve efficiency and reduce costs. In addition, Spectrum Archive EE seamlessly integrates with the scalability, manageability, and performance of IBM Spectrum Scale, which is an IBM enterprise file management platform that enables organizations to move from simply adding storage to optimizing data management.

Here are some of the Spectrum Archive Enterprise Edition highlights:

- Simplify tape storage with the IBM LTFS format, combined with the scalability, manageability, and performance of IBM Spectrum Scale
- Help reduce IT expenses by replacing tiered disk storage (Tier 2 and Tier 3) with IBM tape libraries
- Expand archive capacity by simply adding and provisioning media without affecting the availability of data already in the pool
- Add extensive capacity to IBM Spectrum Scale installations with lower media, floor space, and power costs
- Support for attaching up to two tape libraries to a single Spectrum Scale cluster

Spectrum Archive EE for the IBM TS4500, IBM TS3500, and IBM TS3310 tape libraries provides seamless integration of Spectrum Archive with Spectrum Scale by creating an LTFS

tape tier. You can run any application that is designed for disk files on tape by using Spectrum Archive EE. Spectrum Archive EE can play a major role in reducing the cost of storage for data that does not need the access performance of primary disk. This configuration improves efficiency and reduces costs for long-term, tiered storage.

With Spectrum Archive EE, you can enable the use of LTFS for the policy management of tape as a storage tier in a Spectrum Scale environment and use 7, 6, and 5 tape drives as a critical tier in the storage environment. Spectrum Archive EE supports IBM LTO Ultrium 6 and 5 tape drives, IBM System Storage TS1150, and TS1140 tape drives that are installed in TS4500, and TS3500 tape libraries, or LTO Ultrium 7, 6, and 5 tape drives that are installed in the TS3310 tape libraries.

The use of Spectrum Archive EE to replace disks with tape in Tier 2 and Tier 3 storage can improve data access over other storage solutions because it improves efficiency and streamlines management for files on tape. Spectrum Archive EE simplifies the use of tape by making it transparent to the user and manageable by the administrator under a single infrastructure. Figure 2-13 shows the integration of Spectrum Archive EE archive solution with Spectrum Scale.

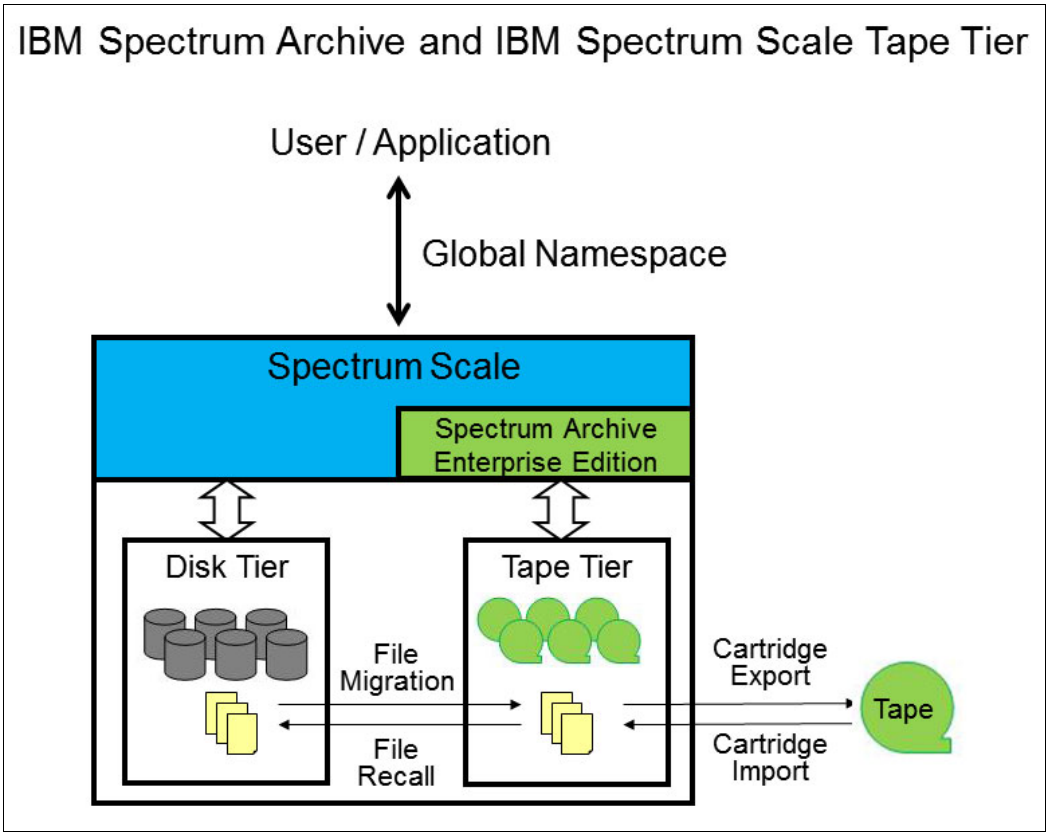

Figure 2-13 Integration of IBM Spectrum Scale and Spectrum Archive Enterprise Edition

The seamless integration offers transparent file access in a continuous name space. It provides file level write and read caching with disk staging area, policy-based movement from disk to tape, creation of multiple data copies on different tapes, load balancing, and high availability in multi-node clusters. It also offers data exchange on LTFS tape by using import and export functions, fast import of file name space from LTFS tapes without reading data, built-in tape reclamation and reconciliation, and simple administration and management.

For more information, see IBM Spectrum Archive Enterprise Edition at the following link:

http://www.ibm.com/systems/storage/tape/ltfs-ee/

IBM Spectrum Archive Library Edition

IBM Spectrum Archive LE extends the file management capability of the IBM Spectrum Archive SDE. Spectrum Archive LE is introduced with Version 2.0 of LTFS. It enables easy-as-disk access to single or multiple cartridges in a tape library.

LTFS is the first file system that works with IBM System Storage tape technology to optimize ease of use and portability for open-systems tape storage. It manages the automation and provides operating system-level access to the contents of the library. Spectrum Archive LE is based on the LTFS format specification, enabling tape library cartridges to be interchangeable with cartridges that are written with the open source SDE version of Spectrum Archive. IBM Spectrum Archive LE supports most IBM tape libraries:

- TS2900 tape autoloader
- TS3100 tape library
- TS3200 tape library
- TS3310 tape library
- TS3500 tape library
- TS4500 tape library

IBM TS1150 and IBM TS1140 tape drives are supported on IBM TS4500 and IBM TS3500 tape libraries only.

Spectrum Archive LE enables the reading, writing, searching, and indexing of user data on tape and access to user metadata. *Metadata* is the descriptive information about user data that is stored on a cartridge. Metadata enables searching and accessing of files through the GUI of the operating system. Spectrum Archive LE supports both Linux and Windows.

Spectrum Archive LE provides the following product features:

- Direct access and management of data on tape libraries with LTO Ultrium 7, 6, and 5 tape drives, and TS1150 and TS1140 tape drives
- Tagging of files with any text, allowing more intuitive searches of cartridge and library content
- Use of the partitioning of the media in LTO-5 tape format standard
- One-to-one mapping of tape cartridges in tape libraries to file folders
- Capability to create a single file system mount point for a logical library that is managed by a single instance of LTFS and runs on a single computer system
- Capability to cache tape indexes, and to search, query, and display tape content within an IBM tape library without having to mount tape cartridges

The IBM Spectrum Archive LE offers the same basic capabilities as the IBM Spectrum Archive SDE with additional support of tape libraries. Each LTFS tape cartridge in the library appears as an individual folder within the file space. The user or application can go to these folders to access the files that are stored on each tape. The Spectrum Archive LE software automatically controls the tape library robotics to load and unload the necessary LTFS Volumes to provide access to the stored files.

For more information, see IBM Spectrum Archive Library Edition at the following link:

http://www.ibm.com/systems/storage/tape/ltfs-le/

IBM Spectrum Archive Single Drive Edition

The IBM Spectrum Archive SDE implements the LTFS Format and allows tapes to be formatted as LTFS Volumes. These LTFS Volumes can then be mounted by using LTFS to allow users and applications direct access to files and directories that are stored on the tape. No integration with tape libraries exists in this edition. You can access and manage all data in stand-alone tape environments as simply as though it were on disk. The limitations of having access to tape as a sequential medium apply.

Use cases for IBM Spectrum Archive SDE are local data import/export; data exchange between locations; all applications that deal with large files like video surveillance, video editing, and streaming; image archiving; and many more.

For more information, see IBM Spectrum Archive Single Drive Edition at the following link:

http://www.ibm.com/systems/storage/tape/ltfs/overview.html

3

IBM Storage Cloud

This chapter explains the relationship of cloud computing to storage and the larger IT landscape. Clouds can be characterized in terms of ownership, that is, public, private, hybrid, and community clouds. They can also be categorized by the type of services that they provide.

Before focusing specifically on storage clouds, it is useful to describe the larger IT landscape for a general understanding of cloud computing concepts.

This chapter includes the following sections:

- Why IBM Cloud?
- Cloud computing overview
- Storage cloud
- IBM storage hardware cloud offerings
- IBM storage cloud software offering
- Security and data protection
- Take the next step

3.1 Why IBM Cloud?

In today's digital economy, the boundaries between virtual and physical experiences are collapsing. Leading businesses understand that customer experience has become the most important source of competitive advantage. With this in mind, IBM has created a collection of fully integrated services to help your business use data across all digital channels to understand its customers and anticipate their needs. It begins with a four-tiered approach.

3.1.1 Why cloud?

Cloud computing, often referred to as simply "the cloud," is the delivery of on-demand computing resources, everything from applications to data centers, over the Internet on a pay-for-use basis. Cloud computing offers these advantages:

- Elastic resources: You can scale up or down quickly and easily to meet demand
- Metered service: You only pay for what you use
- Self service: All the IT resources you need with self-service access

3.1.2 IBM Cloud software as a service

Cloud-based applications, or software as a service (SaaS), run on distant computers "in the cloud" that are owned and operated by others and that connect to users' computers through the Internet and, usually, a web browser

SaaS provides the following benefits:

- You can sign up and rapidly start using innovative business apps
- Apps and data are accessible from any connected computer
- No data is lost if your computer breaks because your data is in the cloud
- The service is able to dynamically scale to usage needs

3.1.3 IBM Cloud platform as a service

Platform as a service (PaaS) provides a cloud-based environment with everything required to support the complete lifecycle of building and delivering web-based (cloud) applications. It does so without the cost and complexity of buying and managing the underlying hardware, software, provisioning, and hosting.

PaaS provides the following benefits:

- Develop applications and get to market faster
- Deploy new web applications to the cloud in minutes
- Reduce complexity with middleware as a service

3.1.4 IBM Cloud infrastructure as a service

Infrastructure as a service (IaaS) provides companies with computing resources that include servers, networking, storage, and data center space on a pay-per-use basis.

IaaS provides the following benefits:

- No need to invest in your own hardware
- Infrastructure scales on demand to support dynamic workloads
- Flexible, innovative services available on demand

3.1.5 Public cloud (SoftLayer)

Public clouds are owned and operated by companies that offer rapid access over a public network to affordable computing resources. With public cloud services, users do not need to purchase hardware, software or supporting infrastructure, which is owned and managed by providers.

Public cloud has the following key aspects:

- Innovative SaaS business apps for applications that range from customer relationship management (CRM) to transaction management and data analytics
- Flexible, scalable IaaS for storage and compute services on a moment's notice
- Powerful PaaS for cloud-based application development and deployment environments

3.1.6 IBM infrastructure for private cloud

A private cloud is infrastructure that is operated solely for a single organization, whether managed internally or by a third party, and hosted either internally or externally. Private clouds can take advantage of cloud's efficiencies while providing more control of resources and steering clear of multi-tenancy.

Private cloud has the following key aspects:

- A self-service interface controls services, allowing IT staff to quickly provision, allocate, and deliver on-demand IT resources
- Highly automated management of resource pools for everything from compute capability to storage, analytics, and middleware
- Sophisticated security and governance designed for a company's specific requirements

3.1.7 IBM hybrid cloud

A hybrid cloud uses a private cloud foundation combined with the strategic integration and use of public cloud services. The reality is that a private cloud cannot exist in isolation from the rest of a company's IT resources and the public cloud. Most companies with private clouds will evolve to manage workloads across data centers, private clouds, and public clouds, which creates hybrid clouds.

Hybrid clouds have the following key aspects:

- Allow companies to keep the critical applications and sensitive data in a traditional data center environment or private cloud
- Enable taking advantage of public cloud resources like SaaS, for the latest applications, and IaaS, for elastic virtual resources
- Facilitate portability of data, applications, and services
- Provide more choices for deployment models

3.1.8 Storage cloud

Companies are increasingly turning to service providers to address their cloud infrastructure needs. To meet demand, service providers must ensure that their cloud infrastructure will grow in lockstep with their clients' needs. They must be able to deploy and scale rapidly while keeping IT service management simple, and without compromising performance or service

levels. Similarly, service providers must also minimize costs and prevent disruption to their own businesses even as they expand IT infrastructure and enhance the delivery of services. Along with establishing a robust networking and server foundation, service providers need to deploy data storage that is suited for a cloud environment to ensure their success.

Robust and agile IBM storage solutions help cloud deliver the following benefits:

- Simplicity and economics to drive speed and efficiency,
- Reliability and interoperability to deliver continuous operations
- Predictability and speed to guarantee service levels

For more information about IBM System Storage for Cloud, see the following link:

http://www.ibm.com/systems/storage/solutions/cloud-storage.html

3.2 Cloud computing overview

The National Institute of Standards and Technology (NIST) provides the following definition[1] for cloud computing:

> "Cloud computing is a model for enabling convenient, on-demand network access to a shared pool of configurable computing resources (for example, networks, servers, storage, applications, and services) that can be rapidly provisioned and released with minimal management effort or service provider interaction."

Users interact with cloud computing environments with the services that the cloud environment provides. The following examples are of services that are provided by a cloud (cloud services):

- Virtual servers
- Database services
- Email applications
- Storage

A company can use cloud services that are provided by third parties, or it can build its own cloud. The company can then provide services from the cloud to internal company users, to selected business partners or customers, or to the world at large.

For a service to be considered a cloud service, it needs to exhibit the following characteristics:

- Support self-service provisioning
- Be accessible through the Internet or corporate intranet
- Provide resources from a resource pool without the user needing knowledge of the pool
- Provide simple and fast resource elasticity, as users demand changes
- Provide the ability to monitor resources with a dashboard view on cloud health status
- Support a metering capability, which enables a dynamic chargeback model

To provide these characteristics, the infrastructure that enables the cloud services takes advantage of two key enablers:

- Virtualization: Allows computing resources to be pooled and allocated on demand. It also enables pay-per-use billing to be implemented.
- Automation: Allows for the elastic use of available resources, and for workloads to be moved to where resources are available. It also supports provisioning and deprovisioning of service instances to support scalability.

[1] See NIST Special Publication (SP) 800-145, *A NIST Definition of Cloud Computing*: http://nvlpubs.nist.gov/nistpubs/Legacy/SP/nistspecialpublication800-145.pdf

As the user of any service, whether it is being provisioned as a cloud service might be immaterial. However, you are likely to be *using* a cloud service when the service that you are using exhibits the characteristics that are listed previously.

3.2.1 Cloud service models

When discussing cloud services, a helpful approach is to organize service capabilities into groups. NIST formally describes a standard for grouping cloud services, referring to them as *service models*. These service models are sometimes referred to as *delivery models* because they describe the services that are delivered by the cloud model. The following sections describe the NIST service models.

Infrastructure as a service

The IaaS model is the simplest for cloud service providers to provision. It can include the following elements:

- Processing
- Storage
- Network

As an IaaS user, you can deploy and run your chosen software, including operating systems and applications. You do not need to manage or control the underlying cloud infrastructure, but you have control over the operating systems, storage, and deployed applications.

Examples of commercial implementations of IaaS include IBM SoftLayer, IBM Cloud Managed Services™, IBM Cloud Managed Backup, Amazon Elastic Compute Cloud (EC2), and Rackspace.

Platform as a service

The PaaS model includes services that build on IaaS services. They add value to the IaaS services by providing a platform on which the cloud users can provision their own applications, or conduct application development activities. The user does not need to manage the underlying cloud infrastructure (network, storage, operating systems), but can control configuration of the provisioned platform services. The following services are provisioned in PaaS models:

- Middleware
- Application servers
- Database servers
- Portal servers
- Development runtime environments

Examples of commercial implementations of PaaS environments include IBM Bluemix®, IBM Cloudant®, Amazon Relational Database Service, and Microsoft Azure.

Software as a service

The SaaS model provides software services that are complete applications that are ready to use. The cloud user simply connects to the application, which is running at a remote location. The user might not know where the system is located. The cloud service provider is responsible for managing the cloud infrastructure, the system on which the application is running, and the application itself.

SaaS is sometimes referred to as *applications as a service*. SaaS also includes content services (for example, video on demand) and higher value network services (for example

Voice over Internet Protocol (VoIP)) as typically encountered in communication service provider scenarios.

Examples of commercial implementations of SaaS environments include IBM Watson™ Analytics, IBM API Management on Cloud, IBM Payment Systems, SalesForce, and NetSuite.

Business process as a service

In recognition of the IT industry direction of provisioning business process as a service (BPaaS) from within a cloud environment, IBM developed a definition of the BPaaS model.

The BPaaS model combines software and workflow elements to deliver end-to-end business processes as a service. Many business processes have the potential to be delivered through this model: Both horizontal applications (such as payroll, technical support, and billing) and vertical markets (such as healthcare and insurance).

Examples of commercial implementations of BPaaS include IBM Source to Pay on Cloud, IBM Customer Experience on Cloud, IBM Watson Business Solutions, ADP HR, and Google Adsense.

Cloud service model layering

Figure 3-1illustrates how the service models described previously can be layered. It also contrasts the level of effort required of the service provider with that of the service user through the service model layers. As you travel up the service model layers, the service provider is responsible for providing more effort as the level of functionality increases. By contrast, as you travel down the service layers, the service user must provide more effort in terms of environment customization.

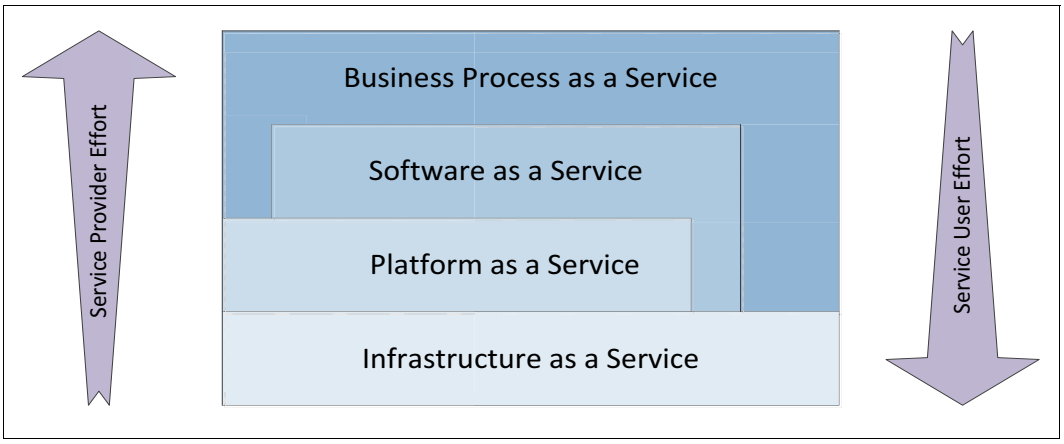

Figure 3-1 Cloud service models

Table 3-1 lists the functions that are provided by the cloud service provider and the cloud service user for each service model.

Table 3-1 Cloud service provider and service user responsibilities by service model

Service model	Cloud service provider-delivered functions	Cloud user-delivered functions
Business process as a service	Business process ↓	Business process configuration
Software as a service	Applications ↓	Application configuration ↑
Platform as a service	Languages Libraries Tools Middleware Application servers Database servers ↓	Applications ↑
Infrastructure as a service	Processing Storage Network	Languages Libraries Tools Middleware Application servers Database servers ↑

3.2.2 Cloud delivery models

Cloud delivery models refer to how a cloud solution is used by an organization, where the data is located, and who operates the cloud solution.

The following cloud delivery models are available:

- Public cloud
- Private cloud
- Hybrid cloud
- Community cloud

Figure 3-2 illustrates these cloud delivery models, and identifies some of their characteristics in terms of roles, users, and accessibility.

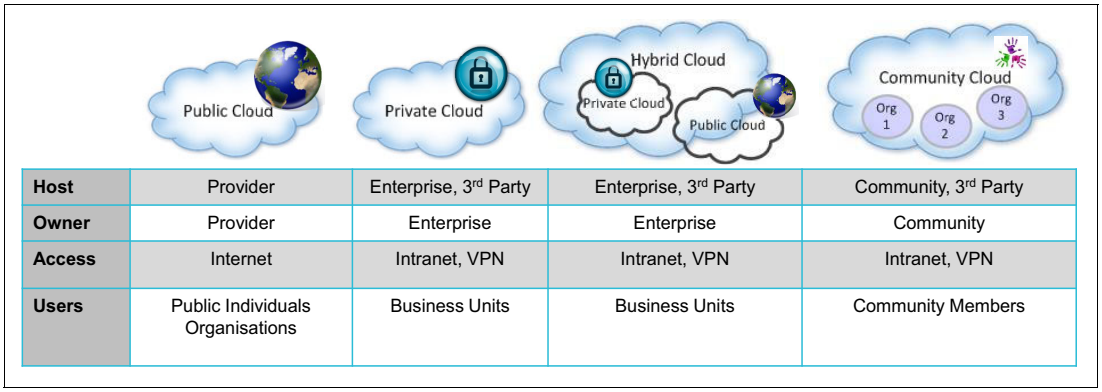

Host	Provider	Enterprise, 3rd Party	Enterprise, 3rd Party	Community, 3rd Party
Owner	Provider	Enterprise	Enterprise	Community
Access	Internet	Intranet, VPN	Intranet, VPN	Intranet, VPN
Users	Public Individuals Organisations	Business Units	Business Units	Community Members

Figure 3-2 Cloud delivery models

Public clouds

A public cloud is one in which the cloud infrastructure is made available to the general public or a large industry group over the Internet. The infrastructure is not owned by the user, but by an organization that provides cloud services. Services can be provided either at no cost, as a subscription, or as a pay-as-you-go model.

Examples of public clouds include IBM SoftLayer, Amazon EC2, Google AppEngine, and Microsoft Azure App Service.

Private clouds

A private cloud refers to a cloud solution where the infrastructure is provisioned for the exclusive use of a single organization. By consolidating and centralizing services into a cloud, the organization benefits from centralized service management and economies of scale.

A private cloud provides an organization with some advantages over a public cloud. Private clouds are ideal when the type of work being done is not practical for a public cloud because of network latency, security, or regulatory concerns.

A private cloud can be owned, managed, and operated by the organization, a third party, or a combination. The private cloud infrastructure is usually provisioned on the organization's premises, but it can also be hosted in a data center that is owned by a third party.

Hybrid clouds

A hybrid cloud is a combination of various cloud types (public, private, and community). Each cloud in the hybrid mix remains a unique entity, but is bound to the mix by technology that enables data and application portability.

The hybrid approach allows a business to take advantage of the scalability and cost-effectiveness of a public cloud without exposing applications and data beyond the corporate intranet. A well-constructed hybrid cloud can service secure processes, such as receiving customer payments (a private cloud service), and secondary processes such as employee payroll processing (a public cloud service).

The challenge for a hybrid cloud is the difficulty in effectively creating and governing such a solution. Services from various sources must be obtained and provisioned as though they originated from a single location, and interactions between private and public components make the implementation even more complicated.

Community clouds

A community cloud shares the cloud infrastructure across several organizations in support of a specific community that has common concerns, such as mission, security requirements, policy, or compliance considerations. The primary goal of a community cloud is to have participating organizations realize the benefits of a public cloud, such as shared infrastructure costs and a pay-as-you-go billing structure, with the added level of privacy, security, and policy compliance that is associated with a private cloud.

The community cloud infrastructure can be provided on-premises or at a third party's data center, and can be managed by the participating organizations or a third party.

For more up-to-date information, see the following website:

`http://www.ibm.com/cloud-computing/`

3.3 Storage cloud

Cloud data storage is a critical component in the cloud computing model. Without cloud storage, there can be no cloud service. This section explores how these challenges can be addressed in the various storage cloud models that are aligned to cloud computing constructs.

A storage cloud provides SaaS to storage consumers. It might be delivered in any of the previously described cloud delivery models (public, private, hybrid, community). A storage cloud can be used to support a diverse range of storage needs, including mass data stores, file shares, backup, and archive. Implementations range from public user data stores to large private storage area networks (SAN) or network-attached storage (NAS), hosted in-house or at third-party managed facilities. The following examples are publicly available storage clouds:

- IBM SmartCloud® offers various storage options, including archive, backup, and object storage.
- Skydrive from Microsoft allows the public to store and share nominated files on the Microsoft public storage cloud service.
- Email services, such as Hotmail, Gmail, and Yahoo, store user email and attachments in their respective storage clouds.
- Facebook and YouTube allow users to store and share photos and videos.
- Dropbox allows users to store and access files anywhere, in the web, on a mobile device, on a hard disk drive (HDD), or on the go.

Storage cloud capability can also be offered in the form of *storage as a service*, where you pay based on the amount of storage space used.

Figure 3-3 shows how various electronic or portable devices can access storage through the Internet without necessarily knowing the explicit details of the type or location of storage that is used underneath.

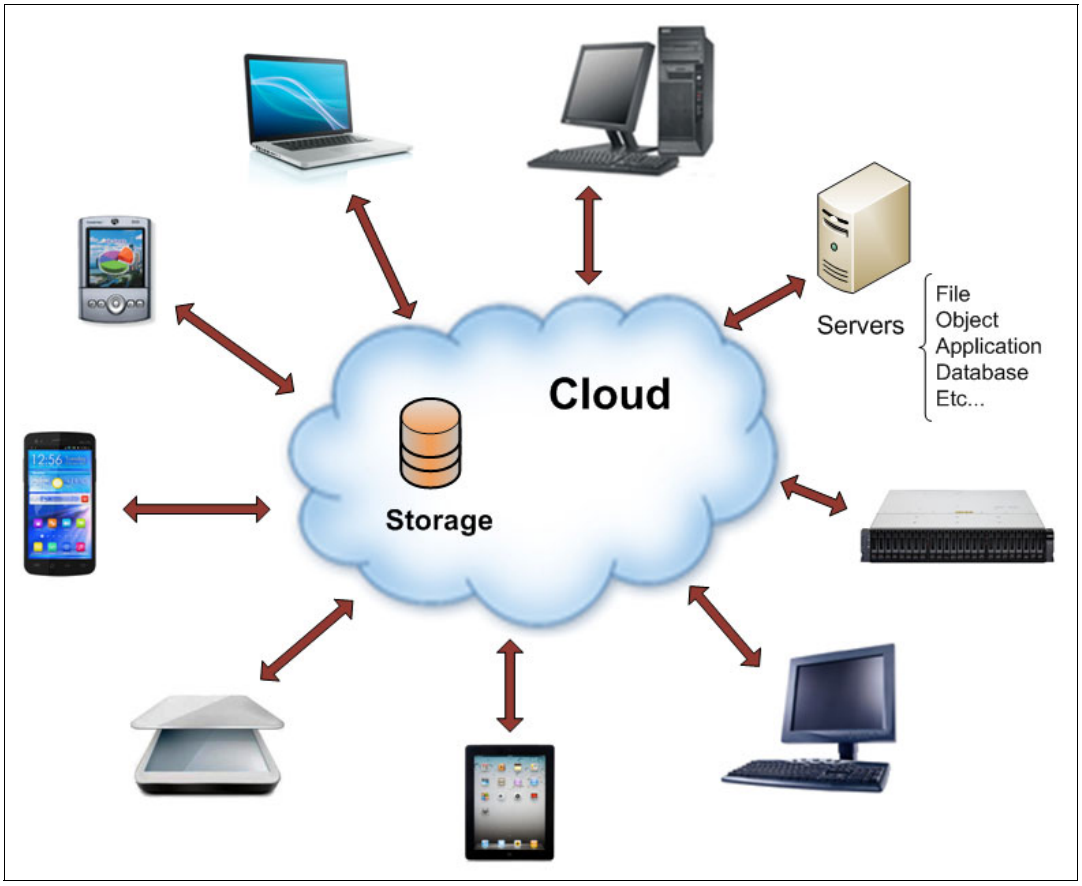

Figure 3-3 Overview of a storage cloud

3.3.1 Storage usage differences within a cloud computing infrastructure

Within a cloud computing infrastructure, a useful distinction can be made between how storage capacity is used, which is similar to the difference that exists in traditional IT between system data (files, libraries, utilities, and so on), and application data and user files. This distinction becomes important for storage allocation in virtual server implementations.

Storage cloud

Storage cloud is the storage capacity service that is provided for client data. A storage cloud exhibits the characteristics that are essential to any cloud service (self-service provisioning, Internet and intranet accessibility, pooled resources, elastic, and metered). It is a cloud environment on which the offered services allow you to store and retrieve data on behalf of computing processes that are not part of the storage cloud service. Storage in a storage cloud can be categorized in the following ways:

- Hosted storage

  This category is primary storage for block, file, or object data that can be written and read on demand. This storage is provisioned as higher performance and availability storage.

- Reference storage

  This category is fixed content storage to which blocks or files are typically written to once, and read from many times. The types of data that are on reference storage typically include multimedia, archival data, medical imaging, surveillance data, and log files.

Storage for cloud

Storage for cloud is a general name that is applied to the type of storage environment. It is implemented in cloud computing that is required to provision cloud computing services. For example, when a virtual server machine is created, storage capacity is required. This storage is provisioned as part of the virtual machine (VM) creation process to support the operating system and runtime environment for the instance. It is not delivered by a storage cloud. The types of storage that are provisioned for a cloud service can be categorized in the following ways:

- Ephemeral storage

  This storage is required only while a VM is running. It is freed from use and made available to the storage pool when the VM is shut down. Examples of this category of storage include boot volumes, page files, and other temporary data.

- Persistent storage

  This storage is required across VM reboots. It is retained even when a VM is shut down. It includes “gold” (master template) images, systems customization, and user data.

Figure 3-4 illustrates the categories of storage found in cloud computing.

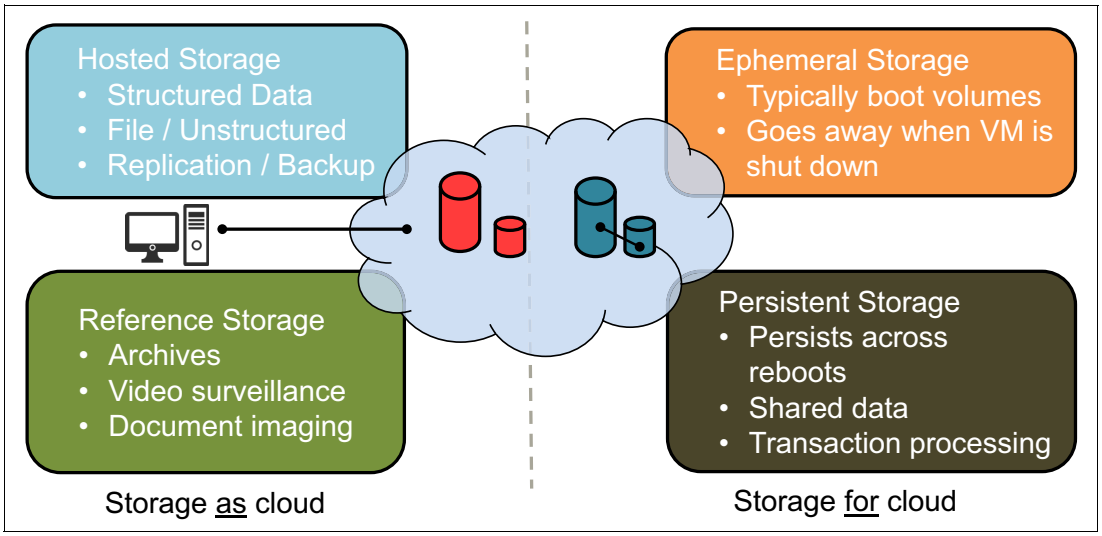

Figure 3-4 Storage categories used in cloud

3.3.2 Traditional storage versus storage cloud

This section compares the various challenges of traditional and cloud storage, outlines the advantages of cloud storage, and explains key implementation considerations for potential storage cloud infrastructure deployments.

Challenges of traditional storage

Before exploring the advantages and benefits of storage cloud, this section lists several limitations of current IT infrastructure that businesses deal with daily. This categorization is from a high level. Challenges in one category can sometimes be applicable to other categories.

Constrained business agility

The time that is required to provision storage capacity for new projects or unexpectedly rapid growth affects an organization's ability to quickly react to changing business conditions. The following constraints are examples:

- Time that is required to deploy new or upgraded business function
- Downtime that is required for data migration and technology refresh
- Unplanned storage capacity acquisitions
- Staffing limitations

Often substantial reserve capacity is required to support growth that requires planning and investment far in advance of the actual need to store data. The following additional issues can affect business agility:

- Inability to meet demand for data availability
- Inability to access the correct data at the correct time to make better business decisions
- Inability to support unplanned acquisitions and staffing limitations

Suboptimal utilization of IT resources

The variation in workloads and the difficulty in determining future requirements typically results in IT storage capacity inefficiencies:

- Difficulty in predicting future capacity and service-level needs
- Peaks and valleys in resource requirements
- Over and under provisioning of IT resources

Extensive capacity planning effort is needed to plan for varying future storage capacity and service level requirements. Capacity is often underutilized because the storage infrastructure requires reserve capacity for unpredictable future growth requirements and therefore cannot be easily scaled up or down.

Organizational constraints

Another barrier to efficient use of resources can be traced to artificial resource acquisition, ownership, and operational practices:

- Project-oriented infrastructure funding
- Constrained operational budgets
- Difficulty implementing resource sharing
- No chargeback or showback mechanism as incentive for IT resource conservation

The limited ability to share data across the enterprise, especially in the context of interdepartmental sharing, can degrade overall use of IT resources including storage capacity. Parallel performance requirements in existing storage systems result in one node supporting one disk, leading to multiplication of nodes and servers.

IT resource management

Efficient IT support is based on cost-effective infrastructure and service-level management to address business needs:

- Rapid capacity growth
- Cost control
- Service-level monitoring and support (performance, availability, capacity, security, retention, and so on)
- Architectural open standardization

The continued growth of resource management complexity in the storage infrastructure is often based on a lack of standardization and high levels of configuration customization. For

example, adjusting storage performance through multiple RAID settings and manual tuning the distribution of I/O loads across various storage arrays consumes valuable staff resources.

Duplicate data that exists in the form of copies across organizational islands within the enterprise leads to higher costs for data storage and backup infrastructure. Compounding this problem are ever-shrinking operational and project budgets, and lack of dynamic chargeback or showback models as incentives for IT resource conservation.

Advantages of a storage cloud

Storage cloud has redefined the way storage consumers can do business, especially those who have seasonal or unpredictable capacity requirements, and those requiring rapid deployment or contraction of storage capacity. Storage cloud can help them focus more on their core business and worry less about supporting a storage infrastructure for their data. Storage cloud provides the advantages:

- Facilitates rapid capacity provisioning, which supports business agility
- Improves storage utilization by avoiding unused capacity
- Supports storage consolidation and storage virtualization functionality
- Chargeback and showback accounting for usage as incentive to conserve resources

Improvement in quality of service (QoS), by automating provisioning and management of underlying complex storage infrastructure, helps improve the overall efficiency of IT storage. Cloud features such as deduplication, compression, automatic tiering, and data migration capabilities are generally built-in options and also support the optimizing of storage costs by implementing tiered storage.

Often the growth in file-based systems is restricted to approximately a few terabytes (TB). This restriction can be easily overcome with storage cloud. Ubiquitous access to data over the Internet, intranet, or both, provides location-independent access. It can also provide a single management platform to manage hundreds of nodes, with data flowing from all the nodes to all the storage arrays.

Capital expenditure can be reduced with a cloud operational-based, pay-as-you-go model. Storage clouds can be tailored or services acquired to support key storage operations such as backup and recovery, remote site disaster recovery, archive, or development and test operations.

Figure 3-5 shows layers that provide unique benefits in the storage cloud.

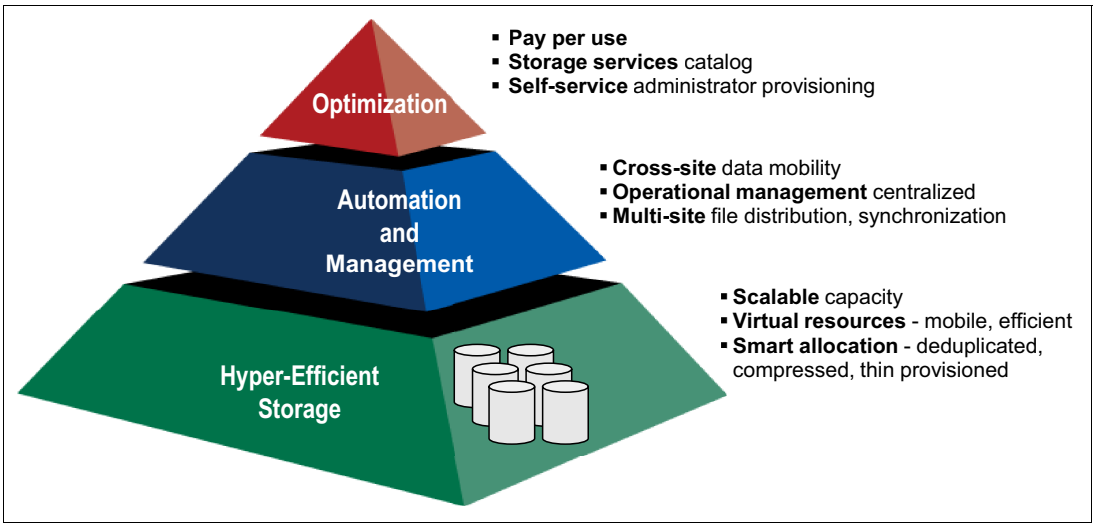

Figure 3-5 Storage cloud characteristics

Implementation considerations for storage cloud

Storage cloud offers many advantages. However, you need to be aware of some of the challenges:

- Have a reliable and robust network infrastructure for remote data access. Because the storage is accessed over the Internet or intranet, a good network connection is essential.
- Security is an important factor. Consider storage device-level encryption for sensitive data.
- Maintain security and control of data that is stored off-site, especially in third-party locations.
- Ensure that regulatory compliance is preserved for various standards.
- Beware of vendor lock-in because of proprietary protocols for access of data by separate storage cloud providers.
- Know the overall reliability of the cloud storage provider. Are service level agreements (SLAs) required, and will providers offer adequate assurance of service delivery? Will the provider remain viable in the future?
- Multitenancy (isolation) is critical. Data needs to be protected from other clients, security threats, viruses, and so on, because data is stored on a common shared storage infrastructure.
- Difficulty in applying policies across many independent filers in an enterprise can cause operational problems.
- Determine whether the cloud storage provider can scale to your capacity and maintain required performance service levels.
- Be able to manage complexity of separate hardware from multiple vendors. Standardization can simplify management for heterogeneous storage devices. Storage virtualization across SAN arrays, such as SAN virtualization with SAN Volume Controller, or Global Namespace solutions, such as IBM Scale Out Network Attached Storage, can provide solutions to this issue.

Benefits and features of storage cloud

The overall benefits of storage cloud vary significantly based on the underlying storage infrastructure. Figure 3-6 identifies basic differences between the traditional IT model and a storage cloud model.

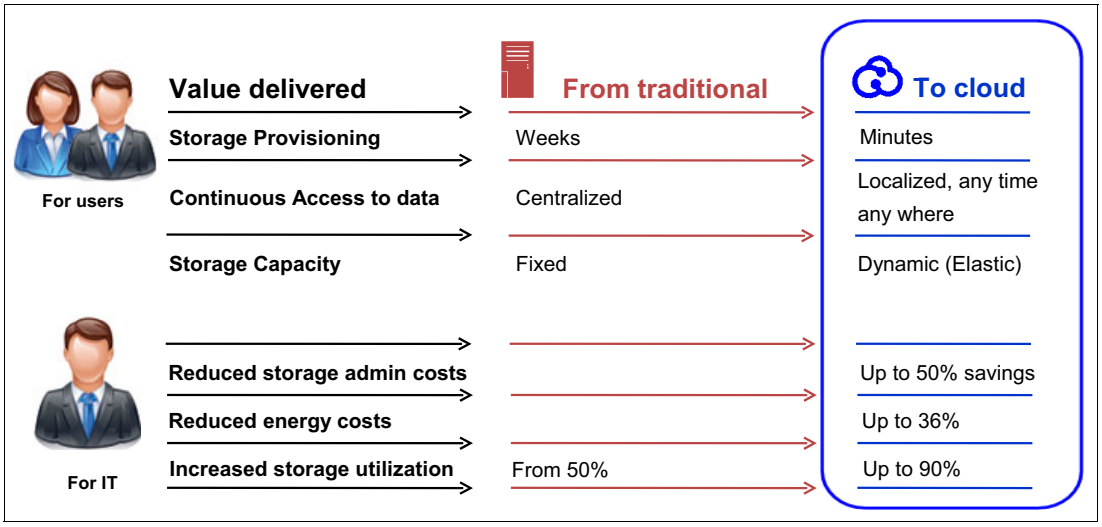

Figure 3-6 Benefits of moving to storage cloud from traditional IT infrastructure

Storage clod provides these benefits and features:

- Dynamic scaling and provisioning (elasticity)

  One of the key advantages of storage cloud is dynamic scaling, also known as *elasticity*. Elasticity means that storage resources can be dynamically allocated (scaled up) or released (scaled down) based on business needs. A storage cloud can start small and grow incrementally with business requirements, or even shrink in size to lower costs if appropriate to capacity demands.

- Faster deployment of storage resources

  New enterprise storage resources can be provisioned and deployed in minutes compared to less optimized traditional IT, which typically takes more time, sometimes days or even months.

- Reduction in total cost of ownership and better return on investment

  Enterprise storage virtualization and consolidation lower infrastructure total cost of ownership (TCO) significantly, with centralized storage capacity and management driving improved usage and efficiency, generally providing a higher return on investment (ROI) through storage capacity cost avoidance.

- Reduce cost of managing storage

  Virtualization helps in consolidating storage capacity and helps achieve much higher utilization, significantly reducing the capital expenditure on storage and its management.

- Greener data centers

  By consolidating geographically dispersed storage into fewer data centers, you achieve a smaller footprint in terms of rackspace. You can save on energy (electrical power) and charges for infrastructure space.

- Dynamic, flexible chargeback model (pay-per-use)

  By implementing storage cloud, an organization pays only for the amount of storage that is used rather than paying for an incremental spare capacity, which remains idle until

needed. Savings can also be realized from hardware and software licensing for functionality such as replication and point-in-time copy.

- Multiuser file sharing

  By centralizing the storage infrastructure, all users can have parallel and simultaneous access to all the data across the enterprise.

- Self-service user portal

  A self-service user portal that is based on a service catalog empowers clients to automatically provision based on predefined templates. You can manage IT infrastructure that is based on the users needs.

- Integrated storage and service management

  The storage cloud infrastructure usually includes integrated management software, which helps to manage the complete storage infrastructure from a single console.

- Improved efficiency of data management

  Consolidation and standardization of storage resources facilitates less infrastructure complexity. Consistent policies and process with integrated management tools support geographically diverse infrastructure requirements.

- Faster time to market

  Automation, self-service portals, rapid deployment, dynamic scaling, and centralized storage management enhance business agility. Businesses can focus on building their core products and competencies instead of worrying about the management of their IT infrastructure.

For more detailed information about Cloud Storage Solutions, see *IBM Private, Public, and Hybrid Cloud Storage Solutions*, REDP-4873.

3.4 IBM storage hardware cloud offerings

IBM is a proven leader in working with clients worldwide to collaborate in solving business problems through the implementation of world class, smart information technology solutions. From hyper-efficient storage arrays and massively scalable virtualization engines to smarter infrastructure management software to comprehensive implementation and operational services capabilities, the IBM portfolio of smart storage cloud capabilities is unmatched in the industry.

If you are building your own storage cloud solution or want a complete solution immediately, the IBM storage cloud portfolio can fulfill that need. Consider the following offerings:

- IBM XIV Storage System. The IBM XIV Storage System is a high-end fully scalable general-purpose disk storage system with a design that is an ideal match with cloud delivery models. It offers an outstanding virtualized grid design, which allows massive parallelism that always allocates system resources evenly. It scales performance with capacity, transparently providing elasticity, which is essential for cloud implementations.

  For more information about the IBM XIV Storage System, see Chapter 7, “IBM XIV Storage System” on page 157.

- IBM Storwize V7000 and V7000 Unified Storage System. IBM Storwize V7000 Unified is a virtualized storage system that is designed to consolidate block and file workloads into a single storage system for simplicity of management, reduced cost, highly scalable capacity, performance, and high availability.

  IBM Storwize V7000 Unified storage also offers improved efficiency and flexibility through built-in solid-state drive (SSD) optimization, thin provisioning, and nondisruptive migration of data from existing storage. The system can virtualize and reuse existing disk systems to provide greater potential ROI. With Integrated IBM Active Cloud Engine®, you can use all those features to build your storage cloud.

  For more information about IBM Storwize V7000 Unified, see Chapter 4, “IBM SAN Volume Controller and IBM Storwize family” on page 81.

- IBM SAN Volume Controller. SAN Volume Controller is a storage virtualization system that enables a single point of control for storage resources to help support improved business application availability and greater resource utilization. The objective is to manage storage resources in your IT infrastructure. You can ensure that they are used to the advantage of your business, and used quickly, efficiently, and in real time, while preventing increases in administrative costs.

  For more information about IBM SAN Volume Controller, see Chapter 4, “IBM SAN Volume Controller and IBM Storwize family” on page 81.

- IBM FlashSystem storage. IBM Flash System delivers extreme performance to derive measurable economic value across the data architecture: Servers, software, applications, and storage.

  For information about IBM FlashSystem storage, see Chapter 5, “IBM FlashSystem family” on page 127.

- IBM Storage DS8880. The IBM DS8880 is the IBM flagship enterprise class Tier 1 block storage device.

  For more information about IBM DS8880, see Chapter 6, “IBM DS8800 Storage System” on page 141.

3.5 IBM storage cloud software offering

IBM storage cloud is a family of integrated enterprise-class cloud computing technologies and services for securely building and using private and public clouds. Built upon open standards and combined with deep IBM expertise and preferred practices, storage cloud is designed for complex high-performance computing environments. Storage cloud goes beyond securely delivering new cloud efficiencies and cost savings to drive fundamental innovation for lasting marketplace advantage.

3.5.1 IBM software defined storage

This section gives an overview of IBM software-defined storage (SDS) products offering with a focus on the IBM Spectrum Storage family of products and their capabilities and benefits.

The products are organized by their functions within the SDS control plane or SDS data plane. The control plane is the software layer that manages administrative functions (for example, configuration, monitoring, replication, policy automation, and provisioning) for SDS resources.

New IBM Spectrum Storage family contains these products:

- Spectrum Control

  Automated control and optimization of storage and data infrastructure (formerly IBM Tivoli Storage Productivity Center, Management layer of IBM Virtual Storage Center, and IBM Spectrum Control Base Edition - IBM Storage Integration Server)

- Spectrum Protect

  Optimized data protection for client data through backup and restore capabilities. (formerly IBM Tivoli Storage Manager Suite for Unified Recovery IBM Tivoli FlashCopy Manager)

- Spectrum Virtualize

  Core SVC functionality is virtualization that frees client data from IT boundaries (formerly IBM SAN Volume Controller)

- Spectrum Accelerate

  Enterprise storage for cloud deployed in minutes instead of months (formerly IBM XIV as Software)

- Spectrum Scale

  Storage scalability to yottabytes and across geographical boundaries (formerly IBM General Parallel File System (GPFS))

- Spectrum Archive

  Enables long-term storage of low activity data (formerly IBM Linear Tape File System Enterprise Edition, Library Edition, and Single Drive Edition)

For more information about IBM SDS, see Chapter 2, “IBM Spectrum Storage family” on page 13.

3.5.2 OpenStack cloud software in cloud computing

Pre-built, pre-tested cloud workflow IaaS offerings are popular for organizations that need to change quickly to stay competitive. By adopting a proven IaaS solution, an IT organization can obtain and implement a reliable template and toolset for implementing true cloud capabilities within the IT organization.

The OpenStack architecture goal is to provide an open source cloud operating system IaaS platform for creating and managing large groups of virtual private servers in a cloud computing environment. OpenStack cloud software is an open source IaaS cloud operating system, which was released under the terms of the Apache 2.0 license.

By adopting and using offerings such as OpenStack cloud software, the IT organization is able to organize, develop skill sets, and deploy cloud computing around proven offerings that already implement industry cloud computing best practices.

OpenStack architecture is one implementation of a best practices cloud workflow. Regardless of the cloud operating system environment that is used, the following key summary points apply:

- Cloud operating systems provide the necessary technology workflow to provide truly elastic, pay-per-use cloud services.
- OpenStack cloud software provides a vibrant open source cloud operating system that is growing quickly.

- OpenStack has the following storage components:
  - The Cinder component provides support, provisioning, and control of block storage.
  - The Swift component provides support, provisioning, and control of object storage.
  - The Manila component provides support, provisioning, and control of file storage.

Figure 3-7 shows the OpenStack architecture.

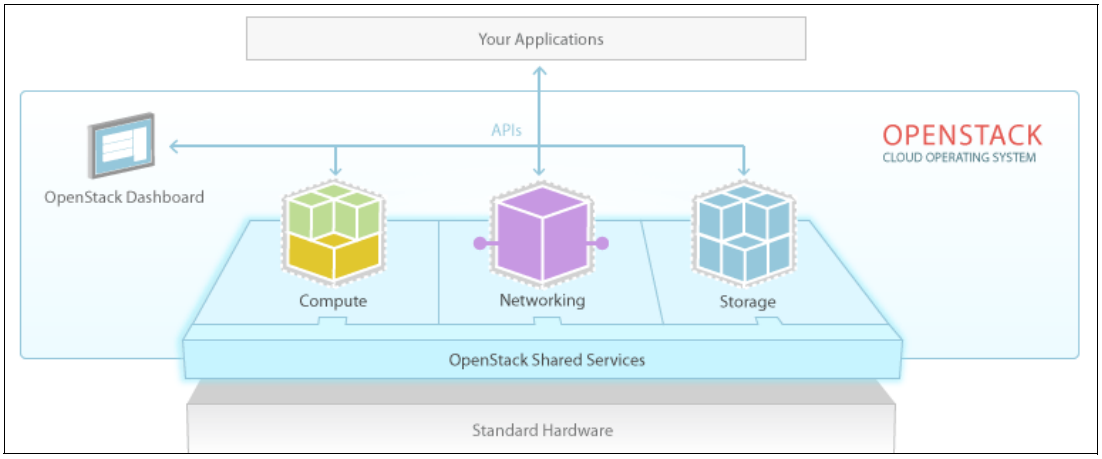

Figure 3-7 OpenStack high architecture high-level overview

OpenStack storage components

Within the overall cloud workflow, the following OpenStack components support storage:

- Cinder (block storage)
- Swift (object storage)
- Manila (file storage)

IBM, a platinum member of the OpenStack Foundation, is an active contributor to OpenStack cloud software and to the Cinder component (and other components of the OpenStack architecture).

OpenStack Cinder

Cinder is the portion of OpenStack cloud software that is responsible for the following functions, among others:

- Creating, modifying, extending, and deleting block storage LUNs
- Attaching/detaching these LUNs to server images created in the Nova component
- Performing block storage functions such as snapshots (create, delete, list)
- Provides support for backup and restore of volumes (volume from snapshot/image)
- Protection functions such as Volume Replications and Consistency Groups

OpenStack software contributors write storage drivers that allow specific storage systems to be supported, provisioned, and managed by Cinder.

IBM has written and contributed OpenStack drivers for the following IBM block storage systems to be supported by Cinder in an OpenStack environment:

- IBM Spectrum Scale (Havana release)
- IBM XIV Storage System (Folsom release)
- IBM Storwize family (Icehouse release)
- IBM System Storage DS8880 (Havana release)
- IBM FlashSystem (Kilo release)

To install the Spectrum Control Cinder driver, make sure that you have access to Virtual Storage Center license with IBM Spectrum Control version 5.2.2, OpenStack (Havana release). The OpenStack Cinder node must run either Red Hat Enterprise 6.4 or higher, or Ubuntu 12.04.

For more information about the IBM storage drivers and functions that are supported in the various OpenStack releases, see:

https://wiki.openstack.org/wiki/CinderSupportMatrix

OpenStack Swift

The OpenStack Swift (Object Store) component provides object storage and allows users to store or retrieve files in a blob mode. Object store is a great fit for data that does not change much such as backups, archives, video, audio, VM images, and Internet-scale repositories of data.

The following are some of the object storage capabilities:

- OpenStack cloud software provides redundant, scalable object storage by using clusters of standardized servers that are capable of storing petabytes of data.
- Object Storage is not a traditional file system, but rather a distributed storage system for static data such as virtual machine images, photo storage, email storage, backups, and archives. Having no central point of control provides greater scalability, redundancy, and durability.
- Objects and files are written to multiple disk drives spread throughout servers in the data center, with the OpenStack software responsible for ensuring data replication and integrity across the cluster.
- Storage clusters scale horizontally simply by adding new servers. If a server or hard disk drive fails, OpenStack software replicates its content from other active nodes to new locations in the cluster. Because OpenStack technology uses software logic to ensure data replication and distribution across different devices, inexpensive commodity hard disk drives and servers can be used in lieu of more expensive equipment.

One thing to note about OpenStack Swift (Object Store) is the eventual consistency. To achieve massive scaling and grow horizontally, the object storage does not have strong consistency, which is typically needed for databases with transactions, where all reads need to be guaranteed to return the most recent data.

IBM Spectrum Scale (as of OpenStack Juno release) combines the benefits of Spectrum Scale with the most widely used open source object store today, OpenStack Swift. Spectrum Scale provides enterprise information lifecycle management (ILM) features. To ensure compatibility with the Swift packages over time, no code changes are required to either Spectrum Scale or Swift to build the solution.

For more information, see *A Deployment Guide for IBM Spectrum Scale Object*, REDP-5113.

OpenStack Manila

The OpenStack Manila (File) component provides file storage that allows coordinated access to shared or distributed file systems. Although the primary consumption of shares would be OpenStack compute instances, the service is also intended to be accessed independently, based on the modular design established by OpenStack services.

The following capabilities are available in Manila:

- Shared file system services for VMs
- Vendor-neutral API for NFS/CIFS and other network file systems

- IBM Spectrum Scale Manila (in Kilo):
  - Extends Spectrum Scale data plane into VM
  - Supports both kNFS and Ganesha 2.0
  - Create/list/delete Shared and Snapshots
  - Allow/deny access to a share based on IP address
  - Multi-tenancy

OpenStack architecture Summary

The rapid growth, adoption, and open source nature of OpenStack cloud software, along with its well-designed support of block, object, and file storage, make it an attractive option for implementing a ready, proven cloud best practices workflow.

For more information about OpenStack technology, see the following site:

http://www.openstack.org

3.5.3 IBM Cloud Orchestrator

IBM Cloud Orchestrator provides an open and extensible cloud management platform for managing heterogeneous hybrid environments. With IBM Cloud Orchestrator, you can manage your cloud infrastructure.

IBM Cloud Orchestrator helps you with end-to-end service deployment across infrastructure and platform layers. It also provides integrated IT workflow capabilities for process automation and IT governance, resource monitoring, and cost management. The product offers an extensible approach to integration with existing environments, such as network management tools. It facilitates integration with customer-specific service management processes, such as processes that are defined in the Information Technology Infrastructure Library (ITIL).

IBM Cloud Orchestrator provides the following features and benefits:

- Quickly deploy and scale on-premises and off-premises cloud services
- Standardization and automation of cloud services through a self-service portal
- Reusable workload patterns to enable dynamic cloud service delivery
- Reduce the administrator workload and error-prone manual IT tasks.
- Use open technologies such as OpenStack cloud software to build an interoperable infrastructure foundation to provision workloads, provide multi-tenancy, and enable administration.
- Deliver services with IBM SoftLayer, existing OpenStack platforms, IBM PowerVM®, IBM z systems, VMware, or Amazon EC2.

Figure 3-8 shows the dashboard of the IBM Cloud Orchestrator that has a “one pane of glass” design that shows all the required information that a cloud administrator needs to effectively manage the cloud resources.

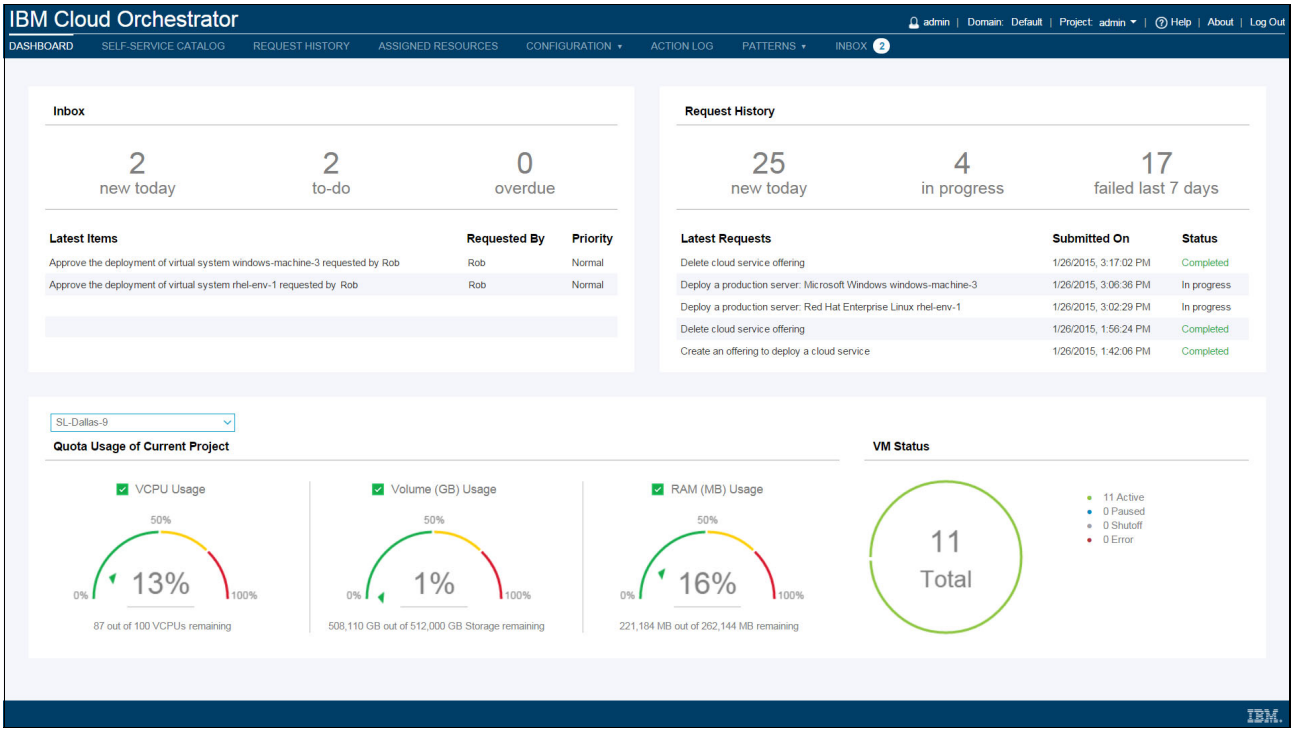

Figure 3-8 IBM Cloud Orchestrator dashboard

Standardization and automation of cloud services help provide the following benefits:

- Allow cloud administrators to expose cloud services in a simpler self-service portal for user consumption
- Control image sprawl, manage image drift, and reduce security vulnerability through analytics, image versioning, and federated image library features
- Include monitoring and capacity analytics capabilities to help consolidate and balance workloads
- Allow you to measure the cost of cloud services with metering and charge-back capabilities

Dynamic cloud service delivery benefits from the following functions of reusable workload patterns:

- Enable rapid deployment and integration of cloud capabilities by using reusable workload patterns for simple infrastructure services and complex multitier applications.
- Allow workload patterns to be created by using templates, or from scratch. After a pattern is created, it can be reused to create multiple identical instances in the cloud.
- Can integrate with middleware components and infrastructure resources to help optimize components for a particular type of application workload.

When different cloud services are used, every service provider has its own management and provisioning tools. Hybrid clouds can refer to a combination of public cloud service and private cloud on-premises. However, hybrid clouds can also consist of two public clouds provided by different vendors or even a combination of a traditional IT system with a cloud. A

setup where existing traditional IT systems are combined with a public cloud service is the most frequent use case of a hybrid cloud.

Because different cloud services can be provided by different vendors, there is no complete picture available showing the total number of deployed instances. An orchestration layer can be a solution for the problem. This layer uses a single interface for all cloud-related tasks that the orchestration layer will communicate with the different cloud services through application programming interfaces. Standardized APIs are being introduced in the industry under the OpenStack initiative. Because IBM Cloud Orchestrator has its infrastructure built on OpenStack, it can provide a unifying view across hybrid cloud services.

Monitoring and metering

IBM Cloud Orchestrator Enterprise enables monitoring of workloads and instances through IBM Cloud Monitoring (ICM). This application includes monitoring virtual environments. IBM Cloud Monitoring can also do capacity planning for correct workload placement, which includes support for both VMware and PowerVM.

Metering capability on a cloud enterprise can be enabled with IBM SmartCloud Cost Management (SCCM). IBM SmartCloud Cost Management can meter the cloud services based on allocation and utilization.

IBM SmartCloud Cost Management requires a data source so that data collectors can collect data from a database or web service. SmartCloud Cost Management supports these four data source types:

- Database
- Message broker
- Server
- Web service

To administer and manage the SmartCloud Cost Management, you must define rate codes, rate groups, and a calendar, and set some configuration options.

Offerings allow you to easily manage the rate groups and rates by logically categorizing related rate groups and rates within a rate table. Rate groups contain rates that can be optionally categorized into offerings.

Rate groups allow all rate codes to be assigned into a group. Creating and using these rate groups let you create rate subtotals in reports, graphs, and spreadsheets. You can group rates such as Windows Charges, UNIX Charges, and Mainframe Charges and all of them will be summarized in a way that is meaningful.

Through SmartCloud Cost Management, a storage administrator can also define Pricing Models to use different pricing models for the same data metering. Differential Pricing uses a rate table for differential pricing and Rate Group/Rate Reordering to show and reorder the rates that are displayed in SmartCloud Cost Management.

For more information about IBM Cloud Orchestrator, see the following link:

https://www.ibmserviceengage.com/cloud-orchestration/learn

For more information about IBM Cloud Orchestrator metering and billing, see:

https://www.ibm.com/support/knowledgecenter/SS4KMC/welcome

3.5.4 Cleversafe solution

On November 6, 2015, IBM completed the acquisition of Cleversafe (now referred to as Cleversafe, an IBM Company). Cleversafe is a web-scale object storage leader with over 350 patents that delivers a breakthrough private cloud software platform that solves petabyte-and-beyond storage challenges. Supporting some of the world's largest data repositories, the Cleversafe solution can significantly reduce storage costs while providing carrier-grade security, high reliability, and simplified storage management. Cleversafe has several customers who use its object storage solution at over 100PBs

Cleversafe's software platform delivers the scalability, security, and cost efficiency needed to harness the power of web-scale data.

Cleversafe provides IBM Cloud with core technology to drive IBM SoftLayer Object Storage service, creating the scalable and reliable storage platform upon which many services can be built. Services like analytics and mobile solutions will be able to use the SoftLayer storage platform to create dynamic and innovative applications with the Cleversafe technology at its foundation. Cleversafe technology also provides IBM clients and service providers with a leading object storage solution that helps the storage administrator manage large amounts of unstructured data. In addition to being implemented in SoftLayer, Cleversafe's solutions will also be implemented in the storage division's Spectrum Storage family of software defined storage.

Increasing revenue and lowering costs requires a holistic data strategy for next generation mobile, social, cloud, and Internet of Things (IoT) applications. The combination of IBM and Cleversafe addresses a broad set of workloads with an expansive set of object-based solutions for unstructured data, such as images, video, and audio files. Cleversafe provides scalability and security without complexity and in many instances at less cost than traditional methods. IBM clients benefit from consistent, scalable, cost-efficient hybrid deployment options, both on-premise when control and security of data is important and in the IBM cloud, making it easier for clients to scale, respond to, and right-size their data deployments.

For more information about Cleversafe, see the following website:

https://www.ibm.com/marketplace/cloud/object-storage/us/en-us

3.6 Security and data protection

Organizations today demand that their data be protected from corruption and loss, whether by accident or intent. This section highlights the security and data protection mechanisms available within storage clouds, and their relevance to providing the data integrity, which businesses have come to expect.

3.6.1 Backup and disaster recovery

Despite rapid data growth, data protection and retention systems are expected to maintain service levels and data governance policies. Data has become integral to business decision-making and basic operations, from production to sales and customer management. Data protection and retention are core capabilities for their role in risk mitigation and for the amount of data involved.

The storage environment offers three functions that improve efficiency and effectiveness of data protection and retention:

- Backup and recovery: Provides cost-effective and efficient backup and restore capabilities, improving the performance, reliability, and recovery of data in respect to SLAs. Backups protect current data, and are unlikely to be accessed unless data is lost.
- Archiving: Retains data that has long-term data retention requirements, either for compliance or business purposes, by providing secure and cost effective solutions with automated process for retention policies and data migration to different storage media.
- Continuous data availability: Ensures uninterrupted access to data for critical business systems, reducing the risk of downtime and providing capability to fail over transparently and as instantaneously as possible to an active copy of the data. The total mirroring strategy needs to be automated to ensure automated failover and then an appropriate automated fail-back.

IBM Spectrum Protect

IBM Spectrum Protect is a family of products that helps manage and control the “information explosion” by delivering a single point of control and administration for storage management needs. It provides a wide range of data protection, recovery management, movement, retention, reporting, and monitoring capabilities by using policy-based automation.

IBM Spectrum Protect enables data protection from failures and other errors by storing backup, archive, space management, and restore data as well as compliance and disaster-recovery data in a hierarchy of auxiliary storage. IBM Spectrum Protect can help protect computers that run various operating systems, on various hardware platforms and connected together through the Internet, wide area networks (WANs), local area networks (LANs), or SANs.

The progressive incremental methodology that is used by IBM Spectrum Protect backs up only new or changed versions of files. This process greatly reduces data redundancy, network bandwidth, and storage pool consumption as compared to traditional methodologies.

For the complete description of IBM Spectrum Protect products family, see Chapter 2, “IBM Spectrum Storage family” on page 13.

Table 3-2 lists the main features, advantages, and benefits that are offered by IBM Spectrum Protect.

Table 3-2 Main features, advantages, and benefits of IBM Spectrum Protect

Features	Advantages	Benefits
Backup and recovery management	Intelligent backups and restores using a progressive incremental backup and restore strategy, where only new and used files are backed up	Centralized protection based on smart-move and smart-store technology, leading to faster backups and restores with fewer network and storage resources needed
Hierarchical storage management	Policy-based management of file backup and archiving	Ability to automate critical processes that relate to the media on which data is stored while reducing storage media and administrative costs associated with managing data

Features	Advantages	Benefits
Archive management	Managed archives	Ability to easily protect and manage documents that need to be kept for a certain period
Advanced data reduction	Combines incremental backup, source and target data deduplication, compression, and tape management to provide data reduction	Reduces the costs of data storage, environmental requirements, and administration

For the most current list of IBM Spectrum Protect products, see the following website:

http://www.ibm.com/systems/storage/spectrum/

3.6.2 Multitenancy

Multitenancy is an architecture in which a single instance of a software application serves multiple customers. Each customer is called a tenant. Tenants might be given the ability to customize some parts of the application, such as color of the user interface (UI) or business rules, but they cannot customize the application's code.

Multitenancy in cloud environments

The term *multitenancy* refers to an architecture that is typically used in cloud environments. Instead of providing each cloud service consumer (tenant) a separate, dedicated infrastructure (single-tenancy architecture), all consumers share one common environment. Shared layers must behave as though they were set up in a dedicated fashion in terms of customization, isolation, and so on.

In general, multi-tenancy implies multiple non-related consumers or customers of a set of services. Within a single organization, this could be multiple business units with resources and data that must remain separate for legal or compliance reasons. Most hosting companies require multi-tenancy as a core attribute of their business model. This requirement might include a dedicated physical infrastructure for each hosted customer or logical segmentation of a shared infrastructure by using software-defined technologies.

A cloud environment has two primary technology stacks where multitenancy is relevant:

- The management environment (cloud management stack)
- The managed environment (infrastructure, platform, or application that is provided as a service)

Depending on the service model, the level and degree of shared infrastructure varies as illustrated in Figure 3-9. For IaaS, typically hypervisors are installed on the managed hardware infrastructure. For PaaS, there is a multitenancy-enabled middleware platform, and for SaaS, the multitenancy-enabled software application is divided into virtual partitions.

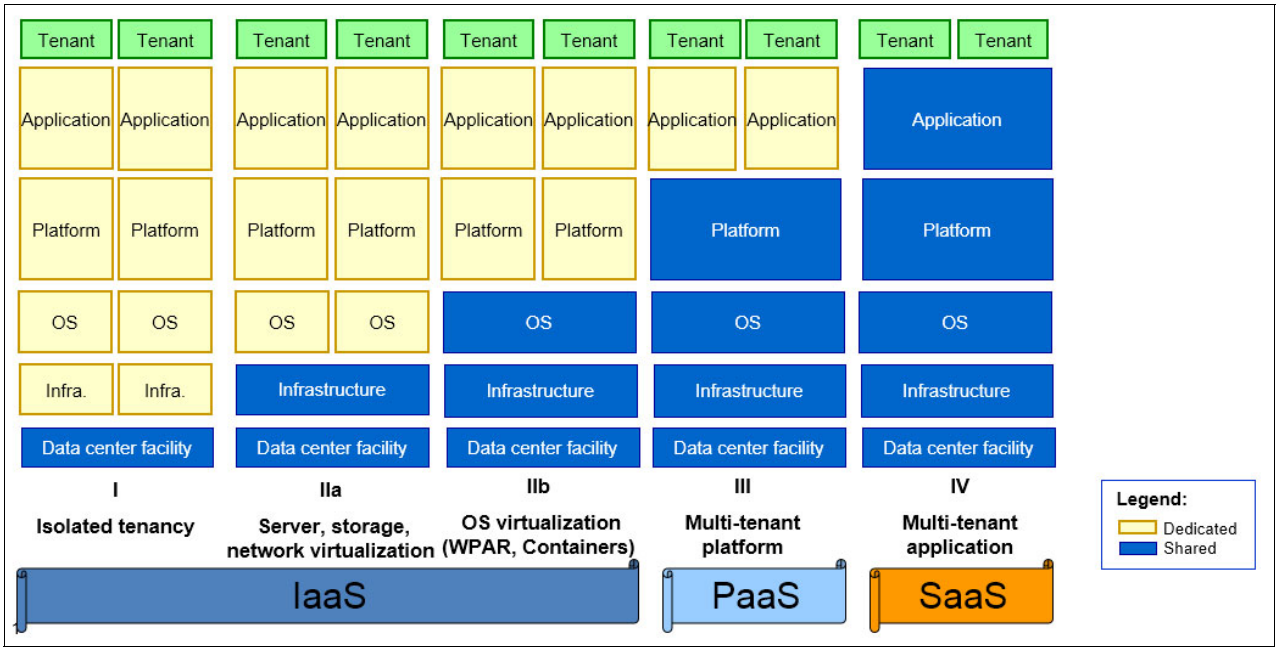

Figure 3-9 Multitenancy in cloud environments

Multitenancy offers several main benefits:

- Can quickly scale to more tenants
- Is cost-effective because the infrastructure is shared by all tenants
- Requires less management effort than a virtualized or mediated approach
- Requires less storage

IBM Tivoli Service Automation Manager

Tivoli Service Automation Manager (see 3.5.3, “IBM Cloud Orchestrator” on page 69) allows service providers to create cloud environments that can be used by multiple customers (multitenancy). Its main strength is data segregation. Teams of users are assigned to a customer and, although a single cloud can support multiple customers, each user sees only the objects that are associated with the customer that that user is assigned to. Customers can share a set of servers, storage, and network resources, or be assigned to different physical resources.

Furthermore, the cloud service provider can assign quotas to each cloud service consumer to define limits on the usage of specific resource pools. These limits define the amount of resources that can be requested by an individual cloud service consumer, such as the amount of storage, memory, physical CPU, and disk.

IBM Spectrum Scale and IBM Spectrum Virtualize

IBM Spectrum Scale and IBM Spectrum Virtualize provide several features that can be used to support scenarios where multiple tenants, can share the large amount of file type storage:

- Security
  - Authentication: IBM Spectrum Scale and IBM Spectrum Virtualize connect to a single, trusted authentication source that can be either Active Directory or Lightweight Directory Access Protocol (LDAP). This configuration allows different tenants such as different departments to use their own, isolated user directories.
  - Authorization: The access control lists (ACLs) that are used to store common permission attributes for NFS and CIFS are the baseline for consistent mappings across all platforms.
  - Audit logging: IBM Spectrum Scale and IBM Spectrum Virtualize support audit logging for all command-line interface (CLI) and graphical user interface (GUI) commands.
- Separation of data
  - Hardware separation: Physical storage is organized in storage pools that allow a separation on the hardware level.
  - Logical separation: From a flexibility perspective, logical separations are far superior to physical separation. Dynamic allocation from a common physical hardware pool to logical entities allows quota-based over-provisioning, and also small step-expansion.
- Utilization and chargeback
  - Accounting: IBM Spectrum Scale and IBM Spectrum Virtualize allow reporting on these levels: User, group, share (if created as file set), and file system.
  - SLA restrictions: Dedicating physical components to individual users can be used for mitigation. For example, assigning a limited number of physical ports to certain servers will throttle the bandwidth that becomes available to those servers. Creating dedicated file systems and file sets on dedicated storage pools can restrict the available disk performance.

Antivirus

NAS systems are designed to serve many users by connecting to them using various file-based protocols, such as NFS or CIFS. The integrity of data that is created or accessed by using these protocols is potentially vulnerable to threats such as viruses, worms, Trojan horses, and other forms of malware. An antivirus solution that works directly with your NAS system helps you to effectively protect your data against those threats.

Encryption

The primary security controls for restricting access to sensitive information that is stored on user devices are encryption and authentication. The appropriate encryption solution for a particular situation depends primarily on the type of storage, the amount of information that needs to be protected, the environments where the storage is located, and the threats that need to be mitigated.

Encryption is a technique that is used to encode data with an encryption key so that the information content of the data can be decoded only with knowledge of the decryption key. Data that is encrypted is referred to as *ciphertext*. Data that is encrypted into ciphertext is considered secure from anyone who does not have possession of the decryption key.

DS8000 disk encryption

Self-encrypting drives protect and secure enterprise information from external and internal threats. The DS8000 (DS8700, DS8800, DS8870) disk system supports data encryption with

the IBM Full Disk Encryption (FDE) drives. All disks in the DS8000 must be FDE drives. No intermix is allowed.

IBM FlashSystem encryption

IBM FlashSystem offerings (840, V840, 900, V9000) are used to make applications and data centers faster and more efficiently. Encryption is provided optionally to protect data at rest against the potential exposure of sensitive user data and user metadata that are stored on discarded or stolen flash modules.

IBM Spectrum Virtualize Gen2 encryption

The IBM Storwize family (Gen2 systems) built with IBM Spectrum Virtualize software also provides an optional encryption for data at rest by using the Protection Enablement Process (PEP). PEP transitions the system from a state of not protection-enabled to protection-enabled.

IBM Spectrum Accelerate Gen3 encryption

The IBM XIV Storage System offers a data-at-rest encryption solution that uses self-encrypting disks (SEDs) and flexible key manager software. When encryption is enabled, the optional solid-state drive (SSD) disks used as flash cache are also encrypted by using software-based AES 256-bit encryption. Uniquely among IBM products, the XIV Storage System offers hot encryption. When encryption is enabled on the XIV system, all data that resides on it is encrypted within minutes, with no performance impact..

IBM XIV

Although most of the storage products offer encryption at extra cost, XIV encryption is available at no additional fee. However, an external key server implemented by IBM Security Lifecycle Manager (SKLM) is required. The IBM XIV solution offers also a simple solution for securely erasing (cryptographic erasure) any disk drive that is being retired or repurposed.

IBM Spectrum Protect encryption

IBM Spectrum Protect client data encryption is the ability to encrypt the actual data file payload for a backup, archive, restore, or retrieve session. The encryption and decryption is done on the client, and all data that is transferred between client and IBM Spectrum Protect server is kept in encrypted cypher text for added protection. For the IBM Spectrum Protect client, transparent encryption is also available. When you use this option, the encryption key is sent to the Spectrum Protect server (the encryption key is itself encrypted) and is stored on the Spectrum Protect server.

IBM tape and virtual libraries

Tape continues to be the most cost-effective, flexible, and scalable medium for high-capacity storage for backup and archiving. Tape can help address compliance requirements with encryption and Write Once Read Many (WORM) solutions. IBM high capacity, high performance tape, and virtual tape offerings can reduce back up windows and allow consolidation of more data onto fewer cartridges.

IBM Security Key Lifecycle Manager

IBM Security Key Lifecycle Manager helps meet regulations and standards such as the Payment Card Industry Data Security Standard (PCI DSS), Sarbanes-Oxley, and the Health Insurance Portability and Accountability Act (HIPAA). It supports the OASIS Key Management Interoperability Protocol (KMIP) standard.

Version V2.5 delivers these advantages:

- An improved and simplified user interface
- Easier installation by using the IBM Installation Manager
- Quicker, more silent installation that lasts approximately 15 minutes
- Seamless integration of all installation components such as IBM WebSphere® Application Server, IBM DB2, and Security Lifecycle Manager application installations.
- Centralized logs (all logs in one place)
- Significantly improved scalability to support larger numbers of keys
- Also includes these advantages:
  - All required installation information is gathered and available before the installation commences
  - Free-flowing installation with minimal intervention from the user

IBM Security Services

IBM Security Services helps you protect your cloud data through a wide array of software and hardware data protection and encryption technologies. The following are the benefits of using IBM Security Services:

- Protect sensitive data that is accessed, stored, and transmitted on your endpoint devices
- Enforce device-usage policy and support regulatory compliance
- Reduce the need for in-house security specialists
- Promote security-rich information exchange by monitoring sensitive data

3.7 Take the next step

You can focus on how you are going to take the next step. Do you have sufficient resources to take the next step on your own? Do you have sufficient skills to navigate the options? Will a technology partner make a cost-effective contribution?

Cloud storage, as with any other emerging technology, is experiencing growing pains. Some facets are immature, fragmented, and lack standardization. Vendors are promoting their own particular technology as the emerging standard. Although standards are lacking, IBM believes that a set of web services API-based capabilities, accessed through non-persistent connections on public or private networks, provides the fundamental frame of reference for accessing storage cloud services. This definition allows for both public service offerings and private use, and provides a basis for expansion of solutions and offerings.

As a leader in cloud computing, IBM has the resources and experience to help businesses implement and use cloud services, including storage cloud. IBM offers hardware and software technologies and key services to help you take advantage of cloud computing. IBM can assist you in planning, designing, building, deploying, and even managing and maintaining a storage cloud environment.

Whether on your premises or someone else's, IBM can make the journey move more quickly, and in many cases deliver value to your business much more rapidly, ultimately saving you money.

Clients that have implemented an IBM Smart Business Storage Cloud solution are projecting savings as follows:

- A large client with 1.5 PB of usable unstructured file system capacity projects savings of over $7.1 million (USD) over the course of five years in hardware acquisition and maintenance, and environmental and administration costs.
- A medium client with 400 TB of usable unstructured file system capacity projects savings of over $2.2 million in hardware acquisition and maintenance, and environmental and administration costs.
- A small client with 200 TB of usable unstructured file system capacity projects savings of over $460,000 in hardware acquisition and maintenance, and environmental and administration costs.

The latest information related to IBM cloud offerings is available at the following website:

http://www.ibm.com/cloud-computing/us/en/index.html

IBM personnel can assist you by developing a high-level architecture and implementation plan. This plan includes a supporting business case to justify investment based on a compelling return on investment, with improved service levels and lowered costs for your cloud infrastructure. IBM consultants use the unique Cloud Computing Reference Architecture (CCRA) and the IBM Cloud Workload Analysis Tool to help you analyze your existing environment. They can then determine the cloud computing model that is best suited for your business. They help you identify the business areas and workloads that, when changed to a cloud computing model, can enable you to reduce costs and improve service delivery t in line with your business priorities.

For more information about Cloud Storage Solutions, see *IBM Private, Public, and Hybrid Cloud Storage Solutions*, REDP-4873.

4

IBM SAN Volume Controller and IBM Storwize family

The IBM Storwize family provides hybrid solutions with common functionality, management, and flexibility. It includes built-in functions such as Real-time Compression and Easy Tier technology optimizing flash and hard disk drives to deliver extraordinary levels of efficiency and high performance. Available in a wide range of storage systems, IBM Storwize family delivers sophisticated capabilities that are easy to deploy, and help control costs for growing businesses.

The all-flash enterprise-class V5000 and V7000F solutions deliver high performance.

This chapter includes the following sections:

- IBM SAN Volume Controller
- IBM Storwize V3700
- IBM Storwize V5000 Gen2
- IBM Storwize V7000 Unified, IBM Storwize V7000, and IBM Storwize V7000F

4.1 IBM SAN Volume Controller

Built with IBM Spectrum Virtualize software, part of the IBM Spectrum Storage family, IBM SAN Volume Controller (SVC) is a dependable system that improves data value, security, and simplicity for new and existing storage infrastructure. Proven over 12 years in thousands of deployments, its innovative data virtualization capabilities help organizations achieve better data economics by supporting new workloads that are critical to their success. SVC systems can handle the massive volumes of data from mobile and social applications, enable rapid and flexible cloud services deployments, and deliver the performance and scalability that are needed to gain insights from the latest analytics technologies.

IBM SAN Volume Controller enhances storage capabilities with sophisticated virtualization, management, functionality, and provides the following benefits:

- Enhance data storage functions, economics, and flexibility with sophisticated virtualization
- Use hardware-accelerated data compression for efficiency and performance
- Use encryption to help improve security for data on existing storage systems
- Optimize tiered storage, including flash storage, automatically with IBM Easy Tier
- Improve network utilization for remote mirroring with innovative replication technology
- Implement multi-site configurations for high availability and data mobility between data centers
- Move data among virtualized storage systems without disruptions

4.1.1 SAN Volume Controller Software

IBM SAN Volume Controller is built with IBM Spectrum Virtualize software, which is part of the IBM Spectrum Storage family.

The software provides these functions for the host systems that attach to SAN Volume Controller:

- Creates a single pool of storage
- Provides logical unit virtualization
- Manages logical volumes
- Mirrors logical volumes

The SAN Volume Controller system also provides these functions:

- Large scalable cache
- Copy Services
  - IBM FlashCopy (point-in-time copy) function, including thin-provisioned FlashCopy to make multiple targets affordable
  - IBM HyperSwap (active-active copy) function
  - Metro Mirror (synchronous copy)
  - Global Mirror (asynchronous copy)
  - Data migration
- Space management
  - IBM Easy Tier function to migrate the most frequently used data to higher-performance storage
  - Metering of service quality when combined with IBM Spectrum Control

– Thin-provisioned logical volumes
– Compressed volumes to consolidate storage

For more information about IBM Spectrum Virtualize software, see 2.3.1, “IBM Spectrum Virtualize” on page 29.

4.1.2 IBM SAN Volume Controller architecture

The SAN Volume Controller combines software and hardware into a comprehensive, modular appliance that uses symmetric virtualization.

Symmetric virtualization is achieved by creating a pool of managed disks (MDisks) from the attached storage systems. Those storage systems are then mapped to a set of volumes for use by attached host systems. System administrators can view and access a common pool of storage on the storage area network (SAN). This functionality helps administrators to use storage resources more efficiently and provides a common base for advanced functions.

A SAN is a high-speed Fibre Channel network that connects host systems and storage devices. In a SAN, a host system can be connected to a storage device across the network. The connections are made through units such as routers and switches. The area of the network that contains these units is known as the *fabric* of the network.

Figure 4-1 shows only one SAN fabric and two zones: Host and storage. In a real environment, it is a preferred practice to use two redundant SAN fabrics. The SVC can be connected to up to four fabrics.

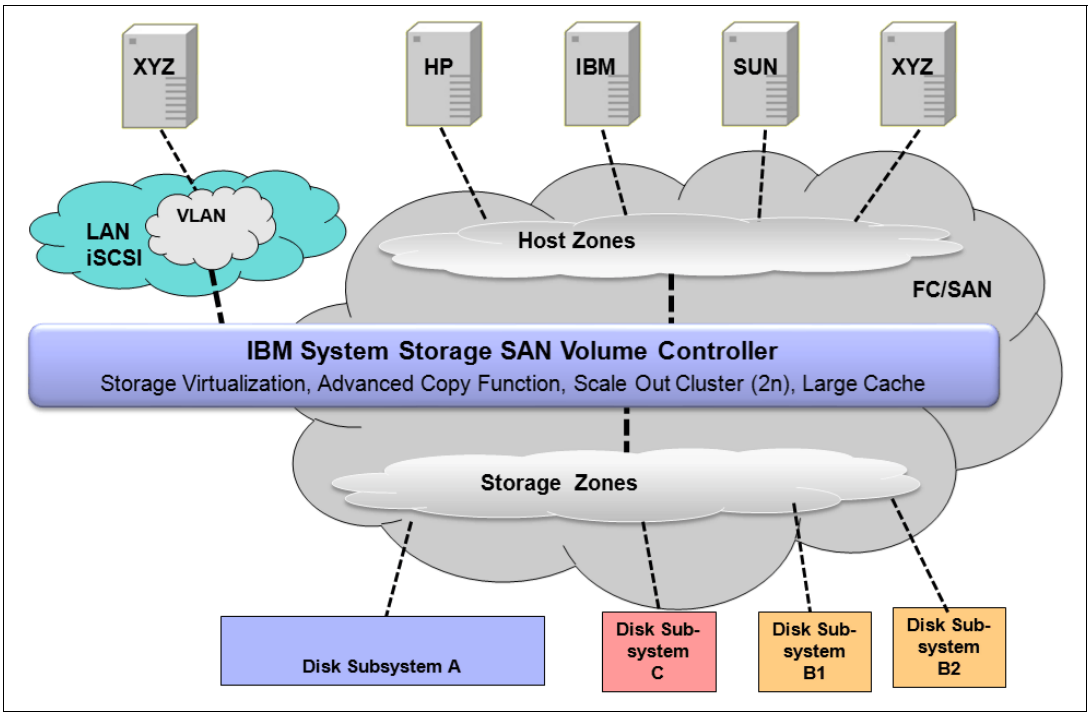

Figure 4-1 SVC conceptual and topology overview

Volumes

A system of SAN Volume Controller nodes presents volumes to the hosts. Most of the advanced functions that SAN Volume Controller provides are defined on volumes. These volumes are created from MDisks that are presented by the RAID storage systems or by

RAID arrays that are provided by flash drives in an expansion enclosure such as SAN Volume Controller 2145-24F. All data transfer occurs through the SAN Volume Controller node, which is described as symmetric virtualization.

Figure 4-2 shows the data flow across the fabric.

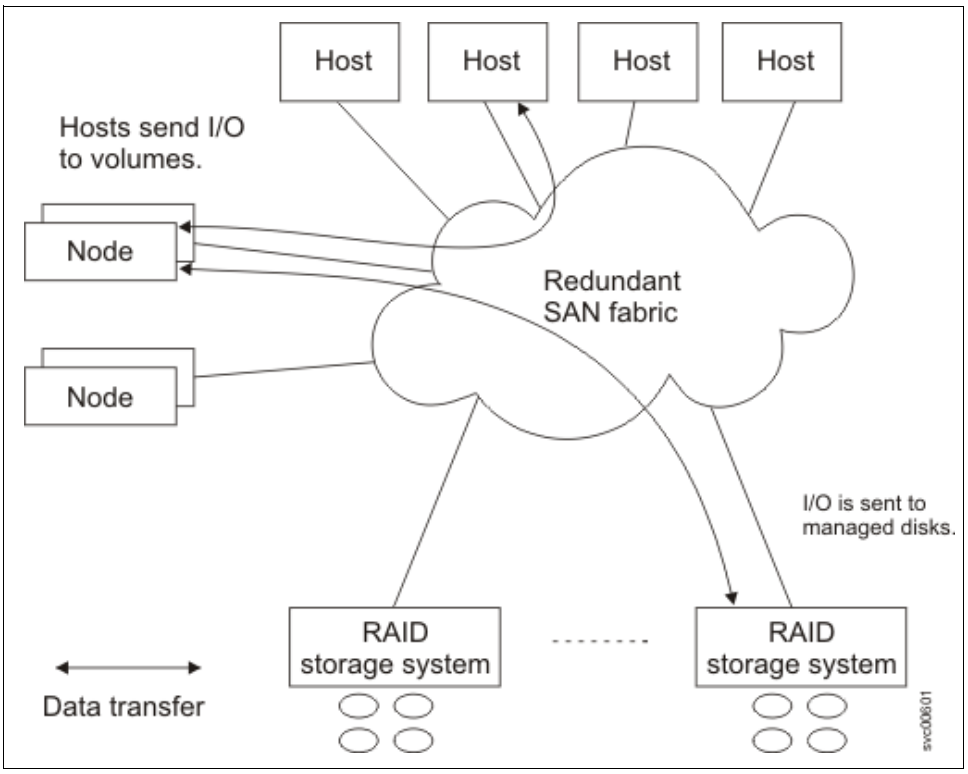

Figure 4-2 Data flow in a SAN Volume Controller system

The nodes in a system are arranged into pairs that are known as *I/O groups*. A single pair is responsible for serving I/O on a volume. Because a volume is served by two nodes, no loss of availability occurs if one node fails or is taken offline.

Volume types

You can create the following types of volumes on the system.

- *Basic volumes*, where a single copy of the volume is cached in one I/O group. Basic volumes can be established in any system topology. Figure 4-3 shows a standard system topology.

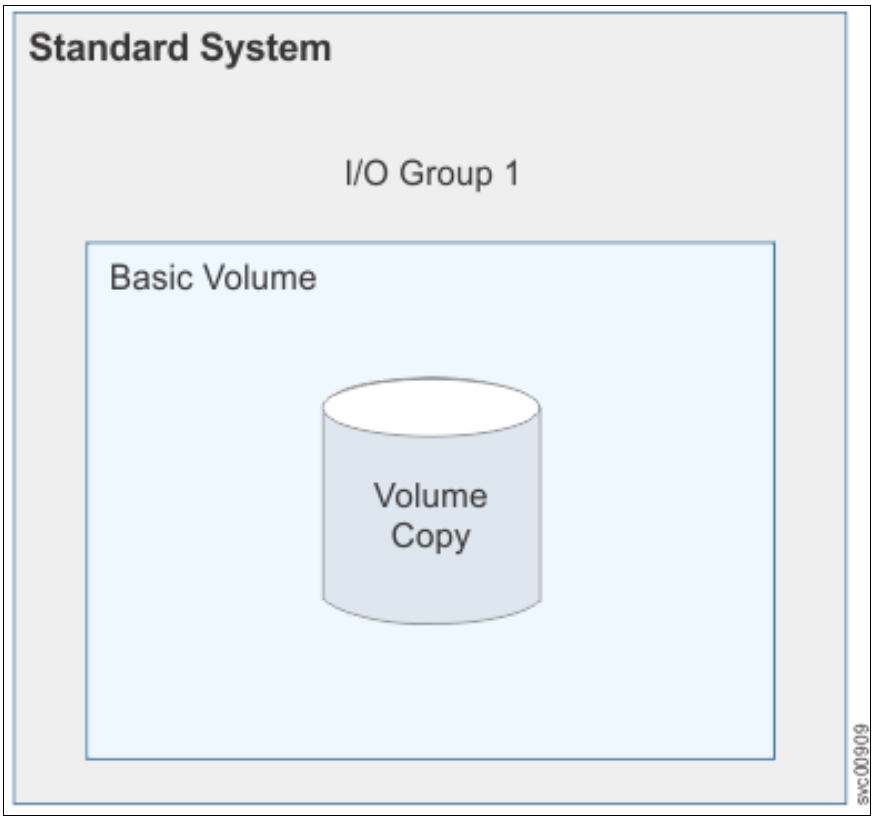

Figure 4-3 Example of a basic volume

- *Mirrored volumes*, where copies of the volume can either be in the same storage pool or in different storage pools. The volume is cached in a single I/O group. Typically, mirrored volumes are established in a standard system topology. Figure 4-4 shows an example of mirrored volumes.

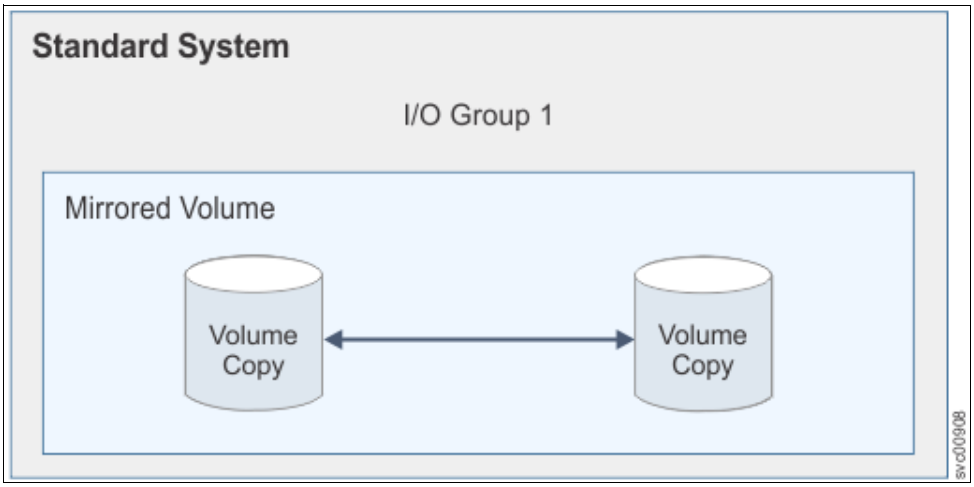

Figure 4-4 Example of mirrored volumes

- *Stretched volumes*, where copies of a single volume are in different storage pools at different sites. The volume is cached in one I/O group. Stretched volumes are only available in stretched topology systems. Figure 4-5 shows an example of a stretched volume.

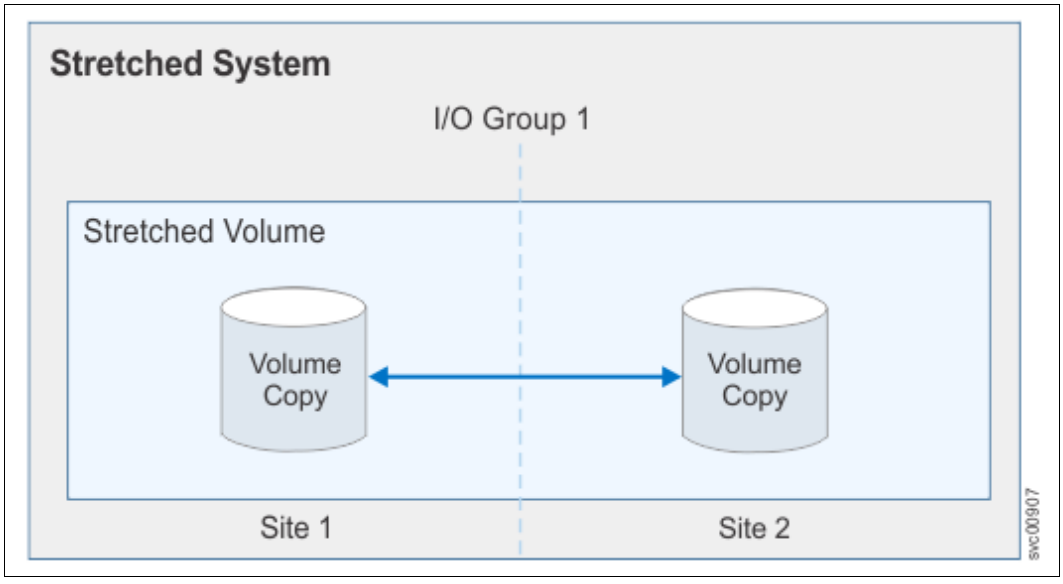

Figure 4-5 Example of stretched volumes

- *HyperSwap volumes*, where copies of a single volume are in different storage pools that are on different sites. The volume is cached in two I/O groups that are on different sites. These volumes can be created only when the system topology is HyperSwap. Figure 4-6 shows an example of a HyperSwap volume.

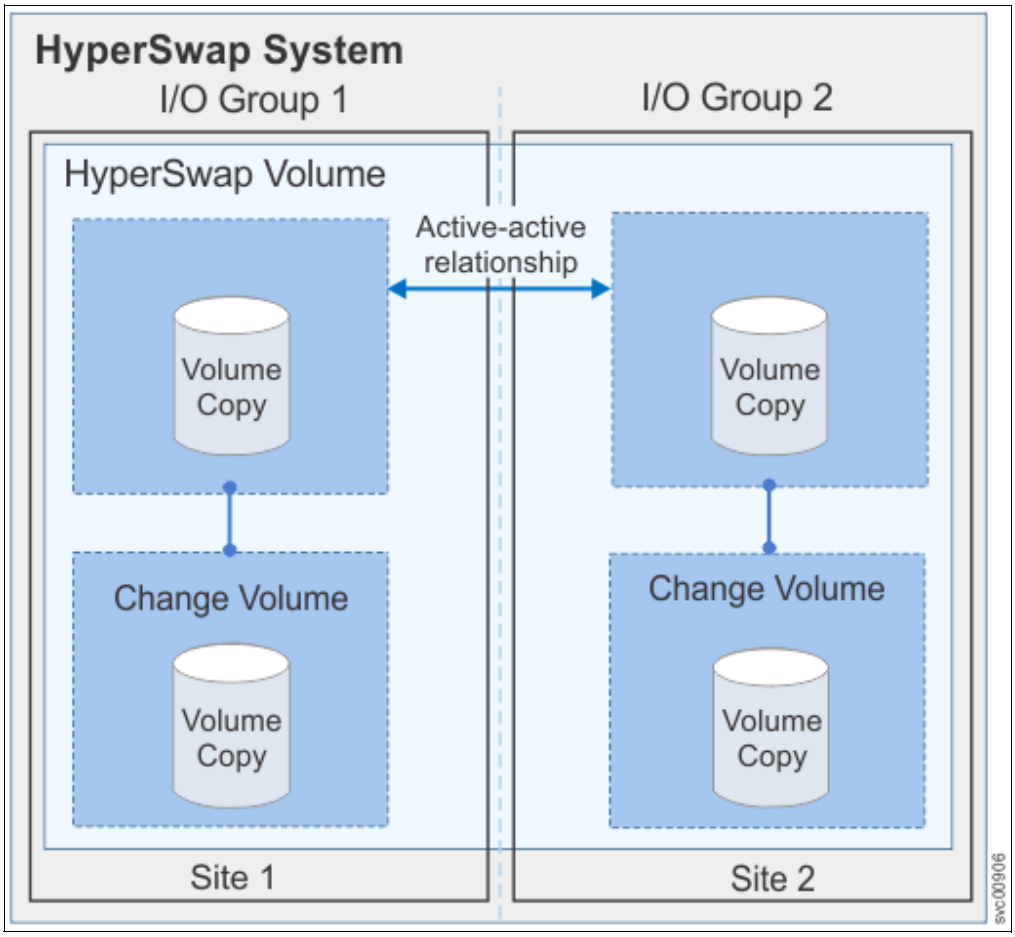

Figure 4-6 Example of HyperSwap volumes

System topology

The topology property of a SAN Volume Controller system can be set to one of the following states:

- The standard topology where all nodes are at the same site or where each node of an I/O group is at a different site. Higher availability can be achieved by using Global Mirror or Metro Mirror to maintain a copy of a volume on a different system at a remote site. The copy at the remote site can be used for disaster recovery. Figure 4-7 shows an example of a standard system topology.

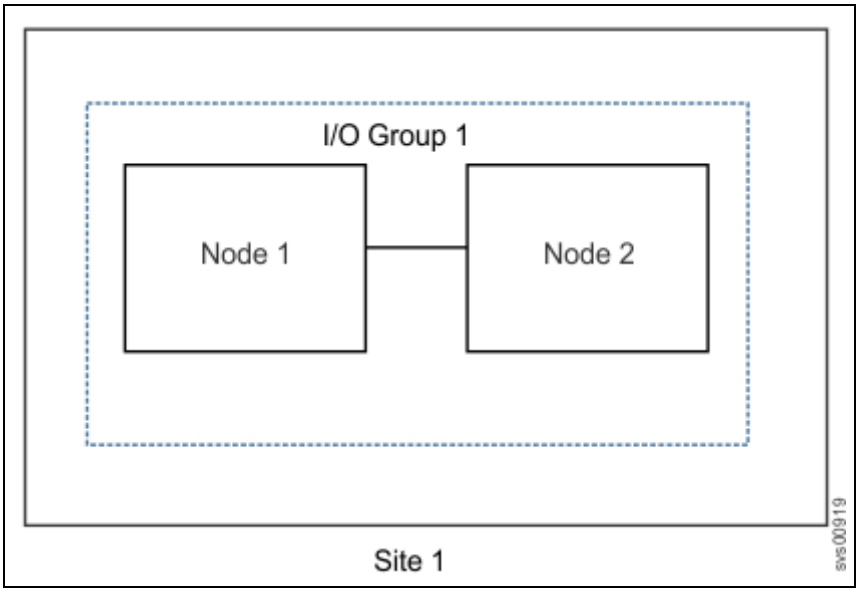

Figure 4-7 Example of a standard system topology

- The stretched topology where each node of an I/O group is at a different site. Access to a volume can continue when one site is not available but with reduced performance. Figure 4-8 shows an example of a stretched system topology.

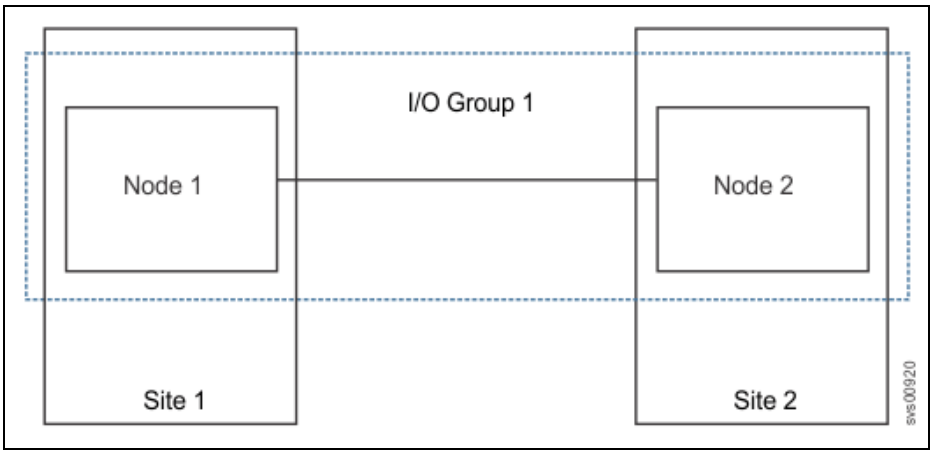

Figure 4-8 Example of a stretched system topology

- The HyperSwap topology where each I/O group is at a different site. A volume can be active on two I/O groups so that it can immediately be accessed through the other site when a site is not available. Figure 4-9 shows an example of a HyperSwap system topology.

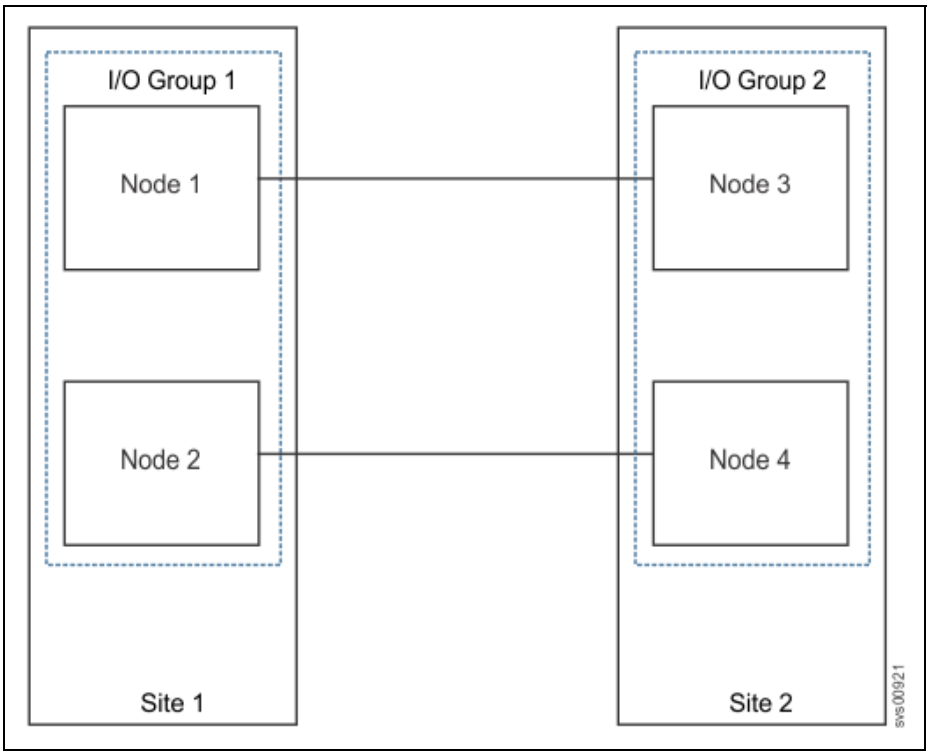

Figure 4-9 Example of a HyperSwap system topology

Note: You cannot mix I/O groups of different topologies in the same system.

Table 4-1 summarizes the types of volumes that can be associated with each system topology.

Table 4-1 Summary of system topology and volumes

	Volume Type				
Topology	**Basic**	**Mirrored**	**Stretched**	**HyperSwap**	**Custom**
Standard	X	X			X
Stretched	X		X		X
HyperSwap	X			X	X

System management

The SAN Volume Controller nodes in a clustered system operate as a single system and present a single point of control for system management and service. System management and error reporting are provided through an Ethernet interface to one of the nodes in the system, which is called the *configuration node*. The configuration node runs a web server and provides a command-line interface (CLI). The configuration node is a role that any node can take. If the current configuration node fails, a new configuration node is selected from the remaining nodes. Each node also provides a command-line interface and web interface for initiating hardware service actions.

Fabric types

I/O operations between hosts and SAN Volume Controller nodes, and between SAN Volume Controller nodes and RAID storage systems, use the SCSI standard. The SAN Volume Controller nodes communicate with each other through private SCSI commands.

Fibre Channel over Ethernet connectivity is supported on SAN Volume Controller models 2145-DH8 and 2145-CG8.

Table 4-2 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 4-2 SAN Volume Controller communications types

Communications type	Host to SAN Volume Controller	SAN Volume Controller to storage system	SAN Volume Controller to SAN Volume Controller
Fibre Channel SAN	Yes	Yes	Yes
iSCSI (1 Gbps Ethernet or 10 Gbps Ethernet)	Yes	No	No
Fibre Channel Over Ethernet SAN (10 Gbps Ethernet)	Yes	Yes	Yes

Flash drives

Some SAN Volume Controller nodes contain flash drives or are attached to expansion enclosures that contain flash drives. These flash drives can be used to create RAID MDisks that in turn can be used to create volumes. In SAN Volume Controller 2145-DH8 nodes, the flash drives are in an expansion enclosure that is connected to both sides of an I/O group.

Flash drives provide host servers with a pool of high-performance storage for critical applications. Figure 4-10 shows this configuration. MDisks on flash drives can also be placed in a storage pool with MDisks from regular RAID storage systems. IBM Easy Tier performs automatic data placement within that storage pool by moving high-activity data onto better performing storage.

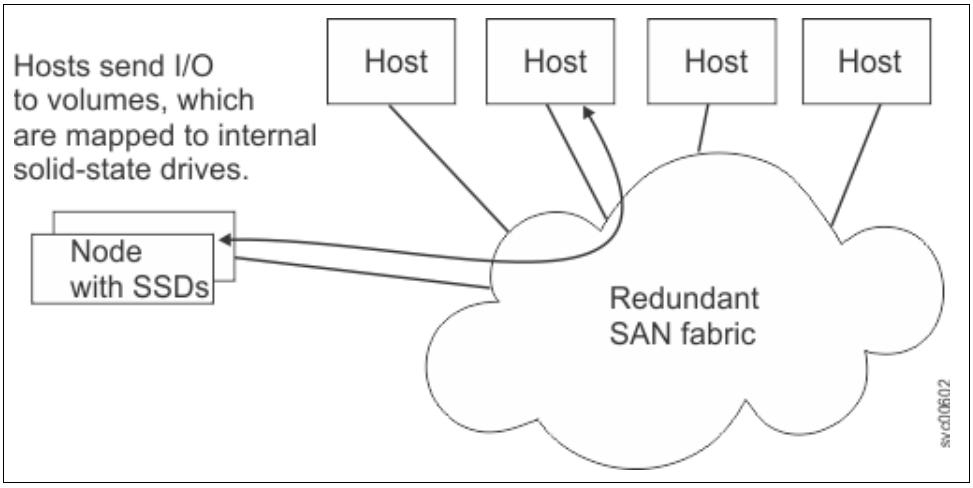

Figure 4-10 SAN Volume Controller nodes with internal Flash drives

Easy Tier

SAN Volume Controller includes IBM Easy Tier, a function that responds to the presence of drives in a storage pool that also contains hard disk drives (HDDs). The system automatically and nondisruptively moves frequently accessed data from HDD MDisks to flash drive MDisks, thus placing such data in a faster tier of storage.

Easy Tier eliminates manual intervention when you assign highly active data on volumes to faster responding storage. In this dynamically tiered environment, data movement is seamless to the host application regardless of the storage tier in which the data belongs. Manual controls exist so that you can change the default behavior, for example, such as turning off Easy Tier on pools that have any combinations of the three types of MDisks.

The system supports these tiers:

Flash tier

The flash tier exists when flash MDisks are in the pool. The flash MDisks provide greater performance than enterprise or nearline MDisks.

Enterprise tier

The enterprise tier exists when enterprise-class MDisks are in the pool, such as those built from serial-attached SCSI (SAS) drives.

Nearline tier

The nearline tier exists when nearline-class MDisks are used in the pool, such as those drives built from nearline SAS drives.

All MDisks belong to one of the tiers, which includes MDisks that are not yet part of a pool.

Encryption

The SAN Volume Controller 2145-DH8 system provides optional encryption of data at rest. This encryption protects against the potential exposure of sensitive user data and user metadata that is stored on discarded, lost, or stolen storage devices. Encryption of system data and system metadata is not required, so system data and metadata are not encrypted. Encryption is performed by the SVC Storage Engine for data stored within the SVC Expansion Enclosures and externally virtualized storage subsystems. Encryption is enabled on the SVC Storage Engine through the acquisition of the Encryption Enablement feature on each SVC DH8 Storage Engine.

The Encryption USB Flash Drives (Pair) feature is required when the Encryption Enablement feature is acquired. This feature provides two USB flash drives for storing the encryption master access key.

Encryption by Spectrum Virtualize is transparent to hosts and applications, and the data is encrypted and protected across the network between the Storwize and SVC controllers and one of more storage devices.

IBM SAN Volume Controller hardware

Each SAN Volume Controller node is an individual server in a SAN Volume Controller clustered system on which the SAN Volume Controller software runs.

The nodes are always installed in pairs, with a minimum of one and a maximum of four pairs of nodes constituting a system. Each pair of nodes is known as an I/O group.

I/O groups take the storage that is presented to the SAN by the storage systems as MDisks and translates the storage into logical disks (volumes) that are used by applications on the hosts. A node is in only one I/O group and provides access to the volumes in that I/O group.

Figure 4-11 shows the front-side view of the SVC 2145-DH8 node.

Figure 4-11 SVC 2145-DH8 storage engine

If you have a SAN Volume Controller 2145-DH8 model, you can add expansion enclosures to expand the available capacity of the system. Each system can have a maximum of four I/O groups with two expansion enclosure attached to each I/O group. Other system models do not support expansion enclosures.

An expansion enclosure houses the following additional hardware: Power supply units (PSUs), canisters, and drives. Enclosure objects report the connectivity of the enclosure. Each expansion enclosure is assigned a unique ID. It is possible to change the enclosure ID later.

Note: A maximum of one expansion enclosure is allowed for each SAS chain. SAN Volume Controller 2145-DH8 nodes support a maximum of two SAS chains.

The SAN Volume Controller 2145-DH8 node has the following features:

- A 19-inch rack-mounted enclosure
- At least one Fibre Channel adapter or one 10 Gbps Ethernet adapter
- Optional second, third, and fourth Fibre Channel adapters
- 32 GB memory per processor
- One or two, eight-core processors
- Dual redundant power supplies
- Dual redundant batteries for better reliability, availability, and serviceability than for a SAN Volume Controller 2145-CG8 with an uninterruptible power supply
- Up to two SAN Volume Controller 2145-24F expansion enclosures to house up to 24 flash drives each
- iSCSI host attachment (1 Gbps Ethernet and optional 10 Gbps Ethernet)
- Supports optional IBM Real-time Compression
- A dedicated technician port for local access to the initialization tool or the service assistant interface.

Model DH8 includes three 1 Gb Ethernet ports standard for iSCSI connectivity. It can be configured with up to four I/O adapter features providing up to sixteen 16 Gb FC ports, up to sixteen 8 Gb FC ports, or up to four 10 Gb Ethernet (iSCSI/Fibre Channel over Ethernet (FCoE)) ports.

Using compression reduces the amount of physical storage across your environment. Compressed volumes are a special type of volume where data is compressed as it is written to disk, saving more space. To use the compression function, you must obtain the IBM Real-time Compression license. In addition, the hardware level for both nodes within the I/O group must be either SAN Volume Controller 2145-DH8, 2145-CG8, or 2145-CF8 for that I/O group to support compression. SAN Volume Controller 2145-DH8 nodes must have two CPUs and at least one compression acceleration adapter installed to use compression. Each compression accelerator increases the speed of I/O transfers between nodes and compressed volumes. Enabling compression on SAN Volume Controller 2145-DH8 nodes does not affect non-compressed host to disk I/O performance.

4.1.3 Copy Services functions

IBM SAN Volume Controller provides Copy Services functions that can be used to improve availability and support disaster recovery. The following Copy Services functions are available for all supported hosts that are connected to SAN Volume Controller:

- IBM FlashCopy function makes an instant, point-in-time copy from a source volume to a target volume.
- Metro Mirror provides a consistent copy of a source volume on a target volume. Data is written to the target volume synchronously after it is written to the source volume so that the copy is continuously updated.
- Global Mirror provides a consistent copy of a source volume on a target volume. The data is written to the target volume asynchronously and the copy is continuously updated. However, the copy might not contain the most recent updates when a disaster recovery operation is performed.
- Global Mirror with change volumes provides support for Global Mirror with higher recovery point objectives (RPOs) by the use of change volumes. This function is for environments that have less bandwidth available between the sites than the update rate of the replicated workload.
- IBM HyperSwap function provides dual-site, active-active access to a volume. Data is written to the copies on both sites synchronously so that both copies are continuously updated and can be used immediately if the other copy becomes unavailable.

Metro Mirror, Global Mirror, Global Mirror with change volumes, and HyperSwap functions are types of remote copy or remote replication.

FlashCopy function

The FlashCopy function creates a point-in-time copy of data stored on a source volume to a target volume.

In its basic mode, the IBM FlashCopy function copies the contents of a source volume to a target volume. Any data that existed on the target volume is lost and is replaced by the copied data. After the copy operation has completed, the target volumes contain the contents of the source volumes as they existed at a single point in time, unless target writes have been processed. The FlashCopy function is sometimes described as an instance of a time-zero copy (T 0) or point-in-time copy technology. Although the FlashCopy operation takes some time to complete, the resulting data on the target volume is presented so that the copy appears to have occurred immediately.

Although it is difficult to make a consistent copy of a data set that is constantly updated, point-in-time copy techniques help solve this problem. If a copy of a data set is created using a technology that does not provide point-in-time techniques and the data set changes during the copy operation, the resulting copy might contain data that is not consistent. For example, if a reference to an object is copied earlier than the object itself and the object is moved before it is copied, the copy contains the referenced object at its new location, but the copied reference still points to the previous location.

More advanced FlashCopy functions allow operations to occur on multiple source and target volumes. FlashCopy management operations are coordinated to provide a common, single point-in-time for copying target volumes from their respective source volumes. This process creates a consistent copy of data that spans multiple volumes. The FlashCopy function also allows multiple target volumes to be copied from each source volume. This function can be used to create images from different points in time for each source volume.

HyperSwap, Metro Mirror, and Global Mirror

Metro Mirror and Global Mirror are two types of remote-copy operations. You can use these functions to set up a relationship between two volumes, where updates made to one volume are mirrored on the other volume. For Metro Mirror and Global Mirror operations, the volumes can be on two different systems (intersystem). Metro Mirror can also support volumes that reside on the same system (intrasystem).

Although data is only written to a single volume, the system maintains two copies of the data. If the copies are separated by a significant distance, the Metro Mirror and Global Mirror copies can be used as a backup for disaster recovery. A prerequisite for Metro Mirror and Global Mirror operations between systems over Fibre Channel connections is that the SAN fabric to which they are attached provides adequate bandwidth between the systems. SAN fabrics are not required for IP-only connections.

For both Metro Mirror and Global Mirror copy types, one volume is designated as the primary and the other volume is designated as the secondary. Host applications write data to the primary volume, and updates to the primary volume are copied to the secondary volume. Normally, host applications do not run I/O operations to the secondary volume.

Metro Mirror is a type of remote copy that creates a *synchronous copy* of data from a primary volume to a secondary volume. A secondary volume can either be on the same system or on another system. With synchronous copies, host applications write to the primary volume but do not receive confirmation that the write operation has completed until the data is written to the secondary volume. This process ensures that both the volumes have identical data when the copy operation completes. After the initial copy operation completes, the Metro Mirror function always maintains a fully synchronized copy of the source data at the target site.

The Global Mirror function provides an *asynchronous copy* process. When a host writes to the primary volume, confirmation of I/O completion is received before the write operation completes for the copy on the secondary volume.

When Global Mirror operates without cycling, write operations are applied to the secondary volume as soon as possible after they are applied to the primary volume. The secondary volume is generally less than 1 second behind the primary volume, which minimizes the amount of data that must be recovered if a failover occurs. However, a high-bandwidth link must be provisioned between the two sites.

The system supports the following types of relationships and consistency groups:

- Active-active (for HyperSwap volumes)
- Metro Mirror
- Global Mirror without change volumes (cycling mode set to None)
- Global Mirror with change volumes (cycling mode set to Multiple)

If necessary, you can change the copy type of a Metro Mirror or Global Mirror remote-copy relationship or consistency group without re-creating the relationship or consistency group with the different type. For example, if the latency of the long-distance link affects host performance, you can change the copy type to Global Mirror to improve host performance over high latency links. For Global Mirror relationships with multiple cycling mode, changes are tracked and copied to intermediate change volumes. The changes are transmitted to the secondary site periodically to lower bandwidth requirements.

Note: You cannot change the type of an active-active relationship or consistency group.

The Metro Mirror and Global Mirror functions support the following operations:

- Intersystem copying of a volume, in which one volume belongs to a system and the other volume belongs to a different system.

  Note: A system can participate in active Metro Mirror and Global Mirror relationships with itself and up to three other systems.

- Intersystem and intrasystem Metro Mirror relationships can be used concurrently on the same system.
- Bidirectional links are supported for intersystem relationships. This configuration means that data can be copied from system A to system B for one pair of volumes while data is copied from system B to system A for a different pair of volumes.
- The copy direction can be reversed for a consistent relationship.
- You can change the copy type for relationships and consistency groups between Metro Mirror and Global Mirror with or without change volumes.
- Consistency groups are supported to manage a group of relationships that must be kept synchronized for the same application. This configuration also simplifies administration because a command that is issued to the consistency group is applied to all the relationships in that group.
- The system supports a maximum of 8192 Metro Mirror and Global Mirror relationships per system.

For more information about the HyperSwap function, see the IBM Knowledge Center:

http://www.ibm.com/support/knowledgecenter/STPVGU_7.6.0/com.ibm.storage.svc.console.760.doc/svc_copyservicesovr_21p99u.html?lang=en

4.1.4 More information

For more information, see the following IBM Redbooks publications:

- *Implementing the IBM System Storage SAN Volume Controller with IBM Spectrum Virtualize V7.6*, SG24-7933
- *IBM SAN Volume Controller 2145-DH8 Introduction and Implementation*, SG24-8229
- *IBM System Storage SAN Volume Controller and Storwize V7000 Best Practices and Performance Guidelines*, SG24-7521
- *IBM SAN Volume Controller and Storwize Family Native IP Replication*, REDP-5103
- *IBM Spectrum Virtualize and SAN Volume Controller Enhanced Stretched Cluster with VMware*, SG24-8211

Also, see these websites:

- SAN Volume Controller product page

  http://www.ibm.com/systems/storage/software/virtualization/svc/

- SAN Volume Controller V7.6 Recommended Software Levels

  http://www.ibm.com/support/docview.wss?uid=ssg1S1005424

- SAN Volume Controller information in the IBM Knowledge Center

  http://www.ibm.com/support/knowledgecenter/STPVGU/welcome

- SAN Volume Controller Support

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

4.2 IBM Storwize V3700

IBM Storwize V3700 is a powerful, easy to use, self-optimizing, and affordable Storwize system, bringing Storwize V7000 value to entry clients.

IBM Storwize V3700 supports up to two redundant RAID controllers with either 12 or 24 drive configurations. There are two types of expansion enclosures, each with either 12 or 24 drives. IBM Storwize V3700 can scale up to a maximum of 10 enclosures with up to 240 drives, depending on the configuration. Built on the proven Storwize V7000 technology, IBM Storwize V3700 provides the following features:

- Great consolidation flexibility through support for the following host attachment protocols: iSCSI, FC, FCoE, and SAS
- The same management interface as Storwize V7000, providing a familiar experience
- Advanced virtualization features that were previously only in midrange systems, such as IBM Storwize V7000
- Sophisticated data protection that uses FlashCopy technology with optional upgrades
- Six Gbps SAS drive attachment technology
- Multiple internal drive options: Flash drives, 15k r SAS, 10k r SAS, and nearline SAS

Figure 4-12 shows the 12-slot model.

Figure 4-12 IBM Storwize V3700 12-slot control or expansion enclosure

Figure 4-13 shows the 24-slot model.

Figure 4-13 IBM Storwize V3700 24-slot control or expansion enclosure

Table 4-3 lists the IBM Storwize V3700 specifications.

Table 4-3 IBM Storwize V3700 specifications

Component	Description
Model	2072-12C - V3700 Control Enclosure 12 x 3.5” drives 2072-24C - V3700 Control Enclosure 24 x 2.5” drives 2072-12E - V3700 Expansion Enclosure 12 x 3.5” drives 2072-24E - V3700 Expansion Enclosure 24 x 2.5” drives 2072-2DC- V3700 SFF DC Dual Control Enclosure 2072-2DE- V3700 SFF DC Dual Expansion Enclosure
RAID controller	Dual active, hot-swappable controllers
Cache per controller	4 GB cache per controller with 8 GB upgrade (battery-backed)
Host interface	Six 6 Gb SAS and four 1 Gb Ethernet ports that are standard for SAS and iSCSI connectivity. In addition, select from these options (per one node canister): ► Four 8 Gbps Fibre Channel ports ► Four 1 Gbps iSCSI ports ► Two 10 Gbps iSCSI ports ► Four 6 Gbps SAS ports
Supported drives	6 Gbps SAS 3.5-inch drives: ► 2 TB, 3 TB, 4 TB, 6 TB, and 8 TB 7,2k r nearline ► 900 GB, 1.2 TB, and 1.8 TB 10k r ► 300 GB, and 600 GB 15k r 6 Gbps SAS 2.5-inch drives: ► 300 GB 15k r ► 600 GB 15k r ► 600 GB 10k r ► 900 GB 10k r ► 1 TB 7.2k r nearline ► 2 TB 7.2k r nearline ► 1.2 TB 10k r ► 1.8 TB 10k r 2.5-inch Flash drives: ► 200 GB, 400 GB, 800 GB, and 1.6 TB
RAID	RAID 0, 1, 5, 6, and 10
Maximum drives supported	Up to 240 drives (high-performance, enterprise class disk drives, high-capacity, archival-class nearline disk drives, and Flash drives)
Fans and power supplies	Fully redundant, hot-swappable
Rack support	Standard 19-inch rack-mount enclosure
Warranty	3-year limited warranty, customer-replaceable units, onsite service, next-business-day 9x5, service upgrades available

For the latest specification information about the IBM Storwize V3700, see this website:

http://www.ibm.com/systems/storage/disk/storwize_v3700/index.html

4.2.1 IBM Storwize V3700 components

IBM Storwize V3700 is a 2U rack mountable enclosure with up to two node or expansion canisters, 12 or 24 drives, and two power supplies that contain fans.

Node canisters contain 4 GB of cache, which can be upgraded to 8 GB that provide up to 16 GB of cache for the system. Each node canister contains a battery to provide cache backup during a power failure. Six Gbps SAS ports are available for expansion connectivity. Node canisters can be upgraded with host interface cards (HICs) to accommodate your host environment:

- Four port 8 Gbps Fibre Channel HIC
- Four port 1 Gbps iSCSI HIC
- Four port 6 GB SAS HIC
- Two port 10 Gbps iSCSI/FCoE HIC

Attachment: IBM Storwize V3700 allows directly attached FC hosts. A maximum of four redundant FC hosts can attach directly.

An expansion canister houses two 6 Gbps SAS ports that are used for connection to the control enclosure and additional expansions.

IBM Storwize V3700 uses new mini SAS high-density (HD) connectors to connect the expansions to the control enclosure. Figure 4-14 shows the difference between mini SAS cables, which are used by IBM Storwize V7000; and mini SAS HD cables that are used by IBM Storwize V3700.

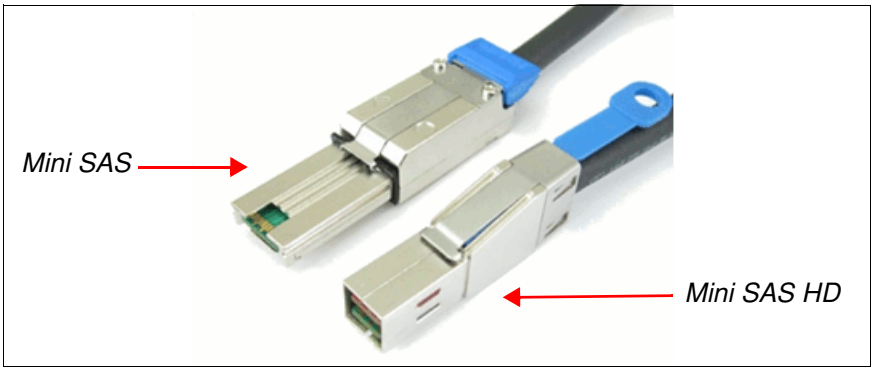

Figure 4-14 Comparison of mini SAS and mini SAS HD cables

IBM Storwize V3700 has a single expansion SAS chain that supports up to nine expansion enclosures. Node canisters use SAS port 4 for the expansion connection.

Figure 4-15 shows the IBM Storwize V3700 expansion cabling topology.

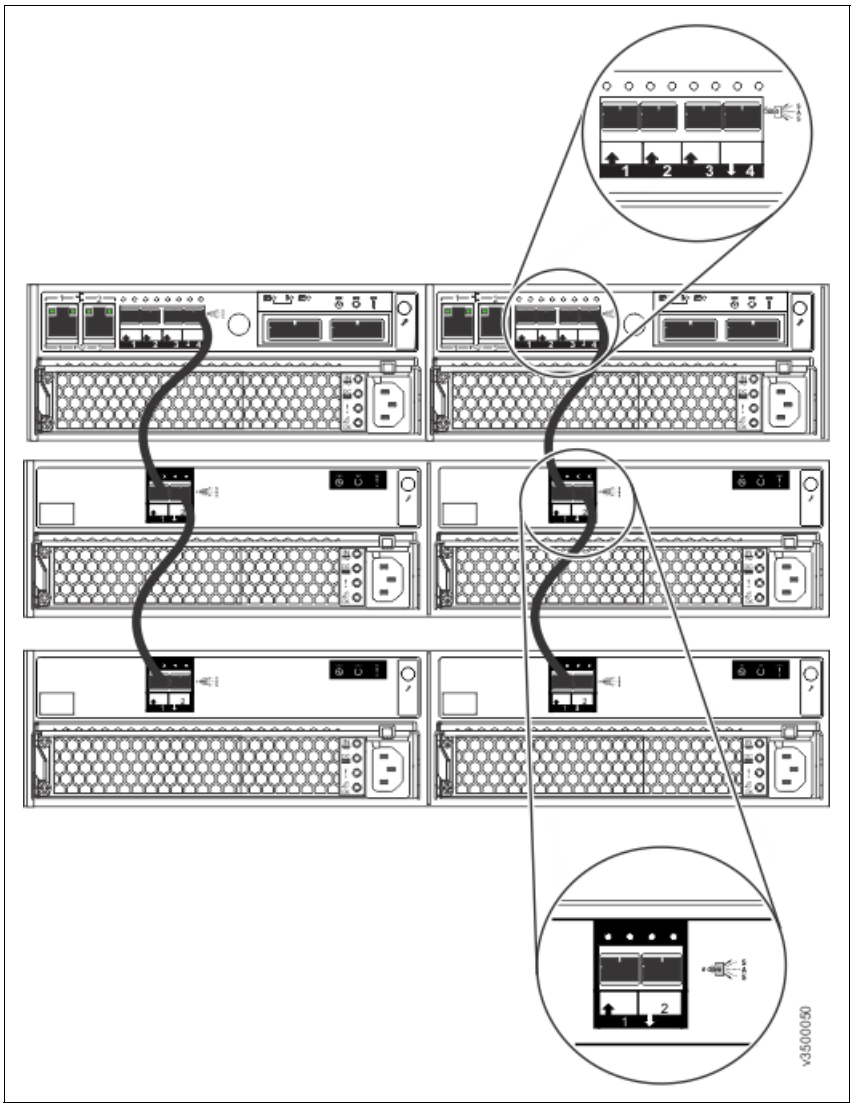

Figure 4-15 IBM Storwize V3700 expansion cabling topology

Most of the V3700 components are customer-replaceable units (CRUs), except for the enclosure midplane:

- Power supplies (with fans)
- Node and expansion canisters
- Battery Backup Unit (BBU)
- Cache memory
- HIC
- All disk drives

4.2.2 IBM Storwize V3700 software and configuration

IBM Storwize V3700 software is based on the SAN Volume Controller and V7000 software stack. However, unlike the IBM SAN Volume Controller and V7000, V3700 software is distributed as Licensed Machine Code (LMC). Therefore, users do not have to purchase and install a separate license to use the system.

The following functions are included with every Storwize V3700 system:

- RAID levels 0, 1, 5, 6, and 10 provide the flexibility to choose the level of data protection that is required.
- Virtualization of internal storage enables rapid, flexible provisioning and simple configuration changes.
- Thin provisioning optimizes efficiency by allocating disk storage space in a flexible manner among multiple users, based on the minimum space that is required by each user at any time. With thin provisioning, applications use only the space they need, not the total space that is allocated to them.
- Data migration enables easy and nondisruptive moves of volumes from another storage system onto the Storwize V3700 system with Fibre Channel (FC) or SAS connectivity. You can easily migrate data from DS3200/DS3500 to Storwize V3700. For more information, see “Migrating” in the IBM Knowledge Center:

  http://www.ibm.com/support/knowledgecenter/STLM5A_7.4.0/com.ibm.storwize.v3700.740.doc/svc_icmigrating_42ki9p.html?lang=en
- FlashCopy creates copies of data for backup, parallel processing, testing, and development. The copies are available almost immediately. Storwize V3700 supports up to 64 FlashCopy targets per system.

Several licensed functions have 90-day trial licenses available. You can use the trial licenses to determine whether the function works as expected, create a business justification for a licensed function, or test the benefit in your environment. If you use a trial license, the system warns you when the trial is about to expire at regular intervals. If a license is not purchased and activated before the trial period expires, the system automatically suspends any configuration that is associated with the function.

The following additional licenses can be purchased to expand the capabilities of your system.

Easy Tier

Easy Tier responds to pools with a mixture of flash, enterprise, and nearline storage in any combination. The system automatically and nondisruptively moves data between the MDisks to optimize performance. For example, Easy Tier moves frequently accessed data from enterprise or nearline MDisks to flash drive MDisks, placing the data in a faster tier of storage. A 90-day trial version of this function is available.

Remote copy

With the remote copy function, you can set up a relationship between two volumes so that updates that are made by an application to one volume are mirrored on the other volume. The volumes can be in the same system or on two systems. The license settings apply only to the system on which you are configuring license settings. For remote copy partnerships, a license is also required on any remote systems that are in the partnership. A 90-day trial version of this function is available.

FlashCopy upgrade

The FlashCopy upgrade extends the base FlashCopy function that is included with the product. The base version of FlashCopy limits the system to 64 target volumes. With the FlashCopy upgrade license activated on the system, an unlimited number of FlashCopy mappings are allowed. If you reach the limit that is imposed by the base function before you activate the upgrade license, you cannot create any more FlashCopy mappings.

Turbo Performance

The Turbo Performance function provides enhanced performance for the system. A 90-day trial version of this function is available.

An IBM Storwize family system is in one of two layers: the replication layer or the storage layer. The system layer affects how the system interacts with the SAN Volume Controller and external Storwize family systems. The SVC can be in the replication layer only. The IBM Storwize V3700 and V7000 can be in the storage layer, which means that they can be put behind another SAN Volume Controller, V3700, or V7000. Also, they can be in the replication layer, which means that they can virtualize another V3700 or V7000, which is used as the storage layer. The Turbo Performance function provides the opportunity for flexible organization and virtualization configurations.

Important: Only image mode disks can be virtualized by IBM Storwize V3700.

The IBM Storwize V3700 uses basic storage units that are called *managed disks* and collects them into one or more storage pools. These storage pools then provide the physical capacity to create volumes for use by hosts.

The Storwize V3700 supports hot-spare drives. Spare drives are global spares, which means that any spare that is at least as large as the drive that is being replaced can be used in an array. The spare system prefers the best possible match based on the following factors: technology (SAS, Flash, or nearline SAS), speed, capacity, and location.

4.2.3 More information

For more information, see *Implementing the IBM Storwize V3700*, SG24-8107.

Also, see the IBM Storwize V3700 information in the IBM Knowledge Center:

http://www.ibm.com/support/knowledgecenter/STLM5A/landing/V3700_welcome.html

4.3 IBM Storwize V5000 Gen2

Designed for software-defined environments and built with IBM Spectrum Virtualize software, the IBM Storwize family is an industry-leading solution for storage virtualization. It includes technologies to complement and enhance virtual environments, delivering a simpler, more scalable, and cost-efficient IT infrastructure.

IBM Storwize IBM Storwize V5000 is a highly flexible, easy to use, virtualized hybrid storage system that is designed to enable midsized organizations to overcome their storage challenges with advanced functionality. IBM Storwize V5000 second-generation models offer improved performance and efficiency, enhanced security, and increased scalability with a set of three models to deliver a range of performance, scalability, and functional capabilities.

4.3.1 IBM Storwize V5000 Gen2 models

IBM Storwize V5000 second-generation models deliver a range of performance, scalability, and functional capabilities.

Table 4-4 provides an overview of IBM Storwize V5000.

Table 4-4 IBM Storwize V5000 model overview

Model	Processor	Expansion enclosures supported	Additional available advanced functions[a]
IBM Storwize V5030	Two 6-core processors and up to 64 GB of cache	Up to 20 IBM Storwize V5000 expansion enclosures	► External storage virtualization ► IBM Real-time Compression ► Encryption of data at rest
IBM Storwize V5020	Two 2-core, 4-thread processors and up to 32 GB of cache	Up to 10 IBM Storwize V5000 expansion enclosures	Encryption of data at rest
IBM Storwize V5010	Two 2-core, 2-thread processors and 16 GB of cache	Up to 10 IBM Storwize 5000 expansion enclosures	

a. In addition to IBM Spectrum Virtualize functions that are supported by the entire family. See , “More detailed information” on page 106.

All IBM Storwize V5000 second-generation models include these features:

- I/O connectivity options for 16 Gb FC, 12 Gb SAS, 10 Gb iSCSI / FCoE, and 1 Gb iSCSI
- Twelve 3.5-inch large form factor (LFF) or twenty-four 2.5-inch small form factor (SFF) drive slots within the enclosure
- Support for the attachment of second-generation IBM Storwize V5000 12 Gb SAS expansion enclosures
- Support for IBM Spectrum Virtualize functions, including thin provisioning, Easy Tier, FlashCopy, and remote mirroring
- A 2U, 19-inch rack mount enclosure with either AC or DC power
- A one- or three-year warranty

For a detailed list of specifications, see the following website:

`http://www.ibm.com/systems/storage/disk/storwize_v5000/specifications.html`

For more information, see *Implementing the IBM Storwize V5000*, SG24-8162

IBM Storwize V5030

IBM Storwize V5030 control enclosure models offer the highest levels of performance, scalability, and functionality with these features:

- Two 6-core processors and up to 64 GB of cache
- Support for 504 drives per system with the attachment of 20 IBM Storwize V5000 expansion enclosures and 1,008 drives with a two-way clustered configuration
- External virtualization to consolidate and provide IBM Storwize V5000 capabilities to existing storage infrastructures
- Real-time Compression for improved storage efficiency
- Encryption of data at rest stored within the IBM Storwize V5000 system and externally virtualized storage systems

- An all-flash enterprise class solution option
- Connectivity:
  - Standard: 10 Gb iSCSI, 1 Gb iSCSI
  - Optional: 16 Gb Fibre Channel,12 Gb SAS, 10 Gb iSCSI / Fibre Channel over Ethernet (FCoE), 1 Gb iSCSI

Storwize V5030 systems can be clustered to help deliver even greater performance, bandwidth, and storage capacity. With two-way system clustering, the maximum size of a Storwize V5030 system increases to 1,008 drives. A Storwize V5030 system can be added into existing Storwize V5000 clustered systems that include previous generation Storwize V5000 systems.

Further scalability of a Storwize V5030 system can be achieved with virtualization of external storage. When Storwize V5030 virtualizes an external disk system, capacity in the external system inherits the functional richness and ease of use of Storwize V5030.

IBM Storwize V5020

IBM Storwize V5020 control enclosure models offer mid-level performance, scalability, and functionality with these features:

- Two 2-core, 4-thread processors and up to 32 GB of cache
- Support for 264 drives per system with the attachment of 10 IBM Storwize V5000 expansion enclosures
- Encryption of data at rest stored within the IBM Storwize V5000 system
- Connectivity
  - Standard: 12 Gb SAS, 1 Gb iSCSI
  - Optional: 16 Gb Fibre Channel, 12 Gb SAS, 10 Gb iSCSI / FCoE, 1 Gb iSCSI

IBM Storwize V5010

Storwize V5010 control enclosure models offer entry-level performance, scalability, and functionality with:

- Two 2-core, 2-thread processors and 16 GB of cache
- Support for 264 drives per system with the attachment of 10 Storwize V5000 expansion enclosures
- Connectivity
  - Standard: 1 Gb iSCSI
  - Optional: 16 Gb Fibre Channel, 12 Gb SAS, 10 Gb iSCSI / FCoE, 1 Gb iSCSI

4.3.2 Components

This section describes the IBM Storwize V5000 Gen2 hardware.

Enclosures

An enclosure is the rack-mounted hardware that contains all the main components of the system: Canisters, drives, and power supplies. The term *enclosure* is also used to describe the hardware and other parts that are plugged into the enclosure.

The system has two different types of enclosures: *Control enclosures* and *expansion enclosures*. A control enclosure manages your storage systems, communicates with the host,

and manages interfaces. In addition, each control enclosure can have multiple attached expansion enclosures, which expand the available capacity of the existing control enclosure. The system supports up to two control enclosures with up to nine expansion enclosures per control enclosure.

An expansion enclosure houses the following extra hardware: PSUs, canisters, and drives. Enclosure objects report the connectivity of the enclosure.

Storwize V5000 drives can only be used in Storwize V5000 enclosures. Storwize V7000, Flex System V7000 Storage Node, Storwize V3500, and Storwize V3700 drives cannot be used in a Storwize V5000 enclosure.

Note: IBM Storwize V5010 and IBM Storwize V5020 control enclosures systems can support up to 10 expansion enclosures on one chain. Storwize V5030 systems can support two chains, and each chain can support up to 10 expansion enclosures.

A Storwize V5000 expansion enclosure can only be used with a Storwize V5000 control enclosure. A Storwize V5000 control enclosure can only manage a Storwize V5000 expansion enclosure.

For more information about enclosures, including ports, see the IBM Storwize V5000 Gen2 hardware components information in the IBM Knowledge Center:

http://www.ibm.com/support/knowledgecenter/STHGUJ_7.6.1/com.ibm.storwize.tb5.761.doc/tbrd_hwinfo_6npwsu.html?lang=en

For more information about supported configurations, see the Storwize V5000 supported environment information in the IBM Knowledge Center:

http://www.ibm.com/support/knowledgecenter/STHGUJ_7.6.1/com.ibm.storwize.tb5.761.doc/svc_siteaudit_210blf.html?lang=en

Supported drives

The IBM Storwize V5000 Gen2 enclosures support a range of enterprise-class, nearline-class, and flash drives as shown in Table 4-5.

Table 4-5 Drive types

Drive type		Speed	Size
2.5-inch form factor	SSD	N/A	400, 800 GB, and 1.6, 3.2 TB
2.5-inch form factor	SAS	15,000 rpm	300 GB, 600 GB
2.5-inch form factor	SAS	10,000 rpm	900 GB, 1.2 TB, 1.8 TB
2.5-inch form factor	Nearline SAS	7,200 rpm	2 TB
3.5-inch form factor	SAS	15,000 rpm	300 GB, 600 GB[a]
3.5-inch form factor	SAS	10,000 rpm	900 GB, 1.2 TB, 1.8 TB[b]
3.5-inch form factor	Nearline SAS	7,200 rpm	4 TB, 6 TB, 8 TB

a. 2.5-inch drive in a 3.5-inch drive carrier
b. 2.5-inch drive in a 3.5-inch drive carrier

More detailed information

For more information about the hardware, including enclosures, host interface adapter ports and indicators, and other components, see the IBM Storwize V5000 Gen2 hardware components information in the IBM Knowledge Center:

http://www.ibm.com/support/knowledgecenter/STHGUJ_7.6.1/com.ibm.storwize.tb5.761.doc/tbrd_hwinfo_6npwsu.html?lang=en

4.3.3 Functional capabilities

Storwize V5000 uses IBM Spectrum Virtualize V6.1 or later software to deliver innovative, proven, and comprehensive capabilities. The following functions are supported on all Storwize V5000 second-generation models, except as noted:

- Virtualization of internal storage: Enables rapid, flexible provisioning and simple configuration changes
- External storage virtualization: Virtualizes existing IBM and non-IBM storage to make it part of the IBM Storwize V5000 system where it inherits the advantages of the IBM Storwize V5000 (IBM Storwize V5030 only)
- Thin provisioning: Supports business applications that need to grow dynamically, while consuming only the space that is actually used
- Encryption of data at rest: Helps protect data from unauthorized access when disks are physically removed from the IBM Storwize system or from externally virtualized storage systems (IBM Storwize V5030 and IBM Storwize V5020 only)
- Real-time Compression: Improves efficiency by storing more active primary data in the same physical disk space (IBM Storwize V5030 only)
- FlashCopy: Enables copies of data to be created with minimal effect to production environments
- Easy Tier: Enables automatic and intelligent migration of frequently accessed data elements to high-performing drives or flash storage (an additional license is required)
- Remote mirroring: Enables synchronous or asynchronous data replication between IBM Storwize V5000 systems or with any other IBM Storwize family system, including IBM Storwize V7000 and IBM SAN Volume Controller
- Copy Services: Improves availability and support disaster recovery (additional licenses are required):
  - Metro Mirror provides a consistent copy of a source volume on a target volume. Data is written to the target volume synchronously after it is written to the source volume so that the copy is continuously updated.
  - Global Mirror provides a consistent copy of a source volume on a target volume. The data is written to the target volume asynchronously and the copy is continuously updated. However, the copy might not contain the most recent updates when a disaster recovery operation is performed.
- HyperSwap: Delivers high-availability configurations for resilient virtualized environments (IBM Storwize V5030 only)
- Dynamic migration: Delivers efficiency and business value through a nondisruptive migration function. A Fibre Channel connection is required to import data from an existing storage controller.

- IBM Storage Mobile Dashboard: Offers basic monitoring capabilities to securely check the health and performance of IBM Storwize family systems
- Graphical user interface: Delivers intuitive data management designed with point-and-click system management capabilities

For more information about Storwize V5000 functional capabilities and software, see the IBM Storwize V5000 Gen2 overview in the IBM Knowledge Center:

http://www.ibm.com/support/knowledgecenter/STHGUJ_7.6.1/com.ibm.storwize.tb5.761.doc/tbrd4_tbrd4ovr.html?lang=en

4.3.4 VersaStack Solution for IBM Storwize V5000

The VersaStack solution by Cisco and IBM can help you accelerate the deployment of your data centers. It reduces costs by more efficiently managing information and resources while still maintaining your ability to adapt to business change.

The VersaStack solution is based on Cisco UCS Mini. It integrates Cisco Unified Computing System (Cisco UCS) Mini blade servers, IBM Storwize V5000 disk system, and UCS Central software as a single system, providing versatility, efficiency, and simplicity for data centers.

For more information, see the following website:

http://www.ibm.com/systems/storage/disk/cisco-versastack.html

4.3.5 More information

For more information, see the following links:

- IBM Storwize V5000 information in the IBM Knowledge Center

  http://www.ibm.com/support/knowledgecenter/STHGUJ
- IBM Storwize V5000 product page

  http://www.ibm.com/systems/storage/disk/storwize_v5000/

4.4 IBM Storwize V7000 Unified, IBM Storwize V7000, and IBM Storwize V7000F

IBM Storwize V7000 Unified and IBM Storwize V7000 are virtualized, enterprise-class hybrid storage systems that provide the foundation for implementing an effective storage infrastructure and transforming the economics of data storage. With industry-first hardware accelerated Real-time Compression, they can reduce the cost of storage by up to 80 percent while maintaining application performance.

IBM Storwize V7000 supports *block* workloads. IBM Storwize V7000 Unified consolidates *block and file* workloads into a single system and can be combined with Storwize V7000 File Modules to create a Storwize V7000 Unified solution that consolidates block and file workloads into a single storage system for greater simplicity and efficiency.

Built with IBM Spectrum Virtualize software, IBM Storwize V7000 Unified and Storwize V7000 provide the latest storage technologies for unlocking the business value of stored data, including virtualization and Real-time Compression. In addition, the systems include a powerful hardware platform that can support the massive volumes of data created by today's

demanding cloud, analytics, and traditional applications. They are designed to deliver outstanding efficiency, ease of use, and dependability for organizations of all sizes.

IBM Storwize V7000 Unified and IBM Storwize V7000 provide the following key benefits:

- Hardware accelerated Real-time Compression reduces storage acquisition costs by using up to 80% less disk and flash capacity
- File support consolidates block and file storage in a system for simplicity and greater efficiency (Storwize V7000 Unified only)
- Active File Management enables highly efficient policy-based management of files to reduce costs through use of tiered storage (Storwize V7000 Unified only)
- Ability for block systems to both scale up and out for performance and capacity with clustered systems
- IP replication reduces remote mirroring costs with innovative network optimization
- Available IBM performance and capacity guarantees help you focus on your business, not your storage, and deploy with confidence
- Automated storage tiering with IBM Easy Tier provides advanced technology for automatically migrating data between storage tiers based on real-time usage analysis patterns
- New generation GUI provides easy-to-use data management designed with a graphical user interface and point-and-click system management capabilities
- It is part of VersaStack, which is an integrated infrastructure solution jointly developed by IBM and Cisco to provide faster application and workload delivery, and IT agility

4.4.1 IBM Storwize V7000 Unified

The IBM Storwize V7000 Unified Disk System integrates the serving of storage and file-related services, such as file sharing and file transfer capabilities, in one system. The Storwize V7000 Unified system can provide storage system virtualization as well, using the mature virtualization capabilities of the IBM SAN Volume Controller. It is an integrated storage server, storage virtualization, and file server appliance.

To serve logical volumes and files, the hardware and software to provide these services are integrated into one product. Viewed from its clients, one part of the Storwize V7000 Unified Disk System is a storage server and the other part is a file server. Therefore, it is called *Unified*.

IBM Storwize V7000 Unified system is a single, integrated storage infrastructure with unified central management that simultaneously supports Fibre Channel, IP Storage Area Networks (iSCSI), and network-attached storage (NAS) data formats, and is centrally managed.

Storwize V7000 Unified storage subsystem: The Storwize V7000

The V7000 Unified system uses internal storage to generate and provide logical volumes to storage clients, so it is a storage system. It can manage the virtualization of external storage systems as well.

The storage subsystem consists of the hardware and software of the IBM Storwize V7000 storage system (Storwize V7000). At the time of writing, it runs the IBM Spectrum Virtualize 7.6 software.

The storage subsystem is used for these functions:

- Provision of logical volumes to external storage clients
- Provision of logical volumes to the internal storage clients, the file modules

Storwize V7000 Unified system file server subsystem: The file modules

In addition to providing logical volumes, the Storwize V7000 Unified Disk System provides access to file system space and to files in those file systems. It uses file sharing protocol, file access protocols, and file transfer or file copy protocols, so it acts as a file server.

The file server subsystem of V7000 Unified system consists of two IBM Storwize V7000 file modules. The file module is a volume storage system that is composed of two units that provide file systems for use by NAS. 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 file system is based on IBM General Parallel File System (GPFS) technology.

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 allows you to configure, administer, and monitor 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 Disk System consists these components:

- A single IBM Storwize V7000 control enclosure
- 0 - 20 Storwize V7000 expansion enclosures (storage server subsystem)
- Two file modules (file server subsystem)
- The connecting cables

You can add up to three control enclosures for more I/O groups. Each additional control enclosure, together with the associated expansion enclosures (up to twenty per control enclosure), provides a new volume I/O group. The file modules remain directly connected to the original control enclosure, which presents I/O group 0. That is the default configuration. Plan to add only block volumes in the new I/O groups. File volumes that are created for you when a new file system is created must continue to be in I/O group 0.

Note: Adding more file modules to the system is not supported.

Storwize V7000 Unified system file server subsystem configuration

The file server subsystem of the V7000 Unified system consists of two file modules. They are IBM System x servers x3650 M3.

The following details are for one file module:

- Form factor: 2U
- Processor: Single four core Intel Xeon C3539 2.13 GHz, 8 G L3 cache (or similar)
- Cache: 72 GB
- Storage: Two 600 GB 10 K SAS drives, RAID 1
- Power supply units: Two (redundant), 675 W

The following interfaces are on one file module:

- Four 1 GbE ports
  - Two ports for external connectivity to file clients and file level remote copy
  - Two ports for the management network between the file modules for unified system clustering
- Two 10 GbE ports for external connectivity to file clients and file level remote copy
- Two 8 Gb FC ports, one of which is internally connected to each Storwize V7000 node canister

The internal and external interfaces of the V7000 Unified system, including the optional 10 GbE interfaces on the Storwize V7000 subsystem, are shown in Figure 4-16.

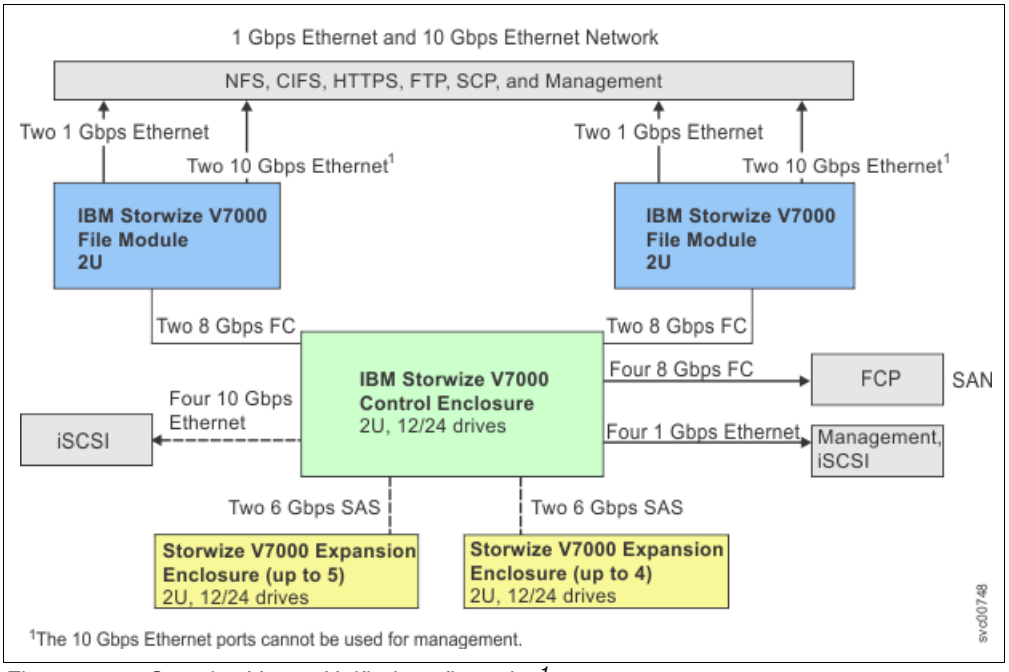

Figure 4-16 Storwize V7000 Unified configuration[1]

For more information about IBM Storwize V7000 Unified, see *Implementing the IBM Storwize V7000 Unified Disk System*, SG24-8010.

4.4.2 IBM Storwize V7000

IBM Storwize V7000 is a storage-system-based in-band block virtualization process, in which intelligent functionality, including advanced storage functions, is available for internal storage and any virtualized storage device.

IBM Storwize V7000 next-generation models offer increased performance and connectivity, integrated compression acceleration, and additional scalability with the following features:

- Two node canisters, each with an eight-core processor and integrated hardware-assisted compression acceleration
- 64 GB cache (per system) standard, with optional 128 GB cache (per system) for Real-time Compression workloads

[1] The 10gbpsEthernet ports cannot be used for management

- 16 Gb FC, 8 Gb FC, 10 Gb iSCSI / FCoE, and 1 Gb iSCSI connectivity options
- 12 Gb SAS expansion enclosures supporting twelve 3.5-inch LFF or twenty-four 2.5-inch SFF drives
- Scaling for up to 504 drives per system with the attachment of 20 Storwize V7000 expansion enclosures and up to 1,056 drives in a four-way clustered configuration
- The ability to be added into existing clustered systems with previous generation Storwize V7000 systems
- Compatibility with IBM Storwize V7000 Unified File Modules for unified storage capability

Storwize V7000 can improve the use of your storage resources, simplify your storage management, and improve the availability of your applications. It can also reduce the number of separate environments that need to be managed down to a single environment. It also provides a single interface for storage management. After the initial configuration of the storage subsystems that you are going to virtualize, all of the day-to-day storage management operations are performed from Storwize V7000.

IBM Storwize V7000 architecture

Storwize V7000 is a SAN block aggregation virtualization appliance that is designed for attachment to various host computer systems.

Two major approaches currently exist that you should consider for the implementation of block-level aggregation and virtualization:

- Symmetric: In-band appliance

  The device is a SAN appliance that sits in the data path, and all I/O flows through the device. This implementation is referred to as *symmetric virtualization* or *in-band*.

  The device is both target and initiator. It is the target of I/O requests from the host perspective, and the initiator of I/O requests from the storage perspective. The redirection is performed by issuing new I/O requests to the storage. Storwize V7000 uses symmetric virtualization.

- Asymmetric: Out-of-band or controller-based

  The device is usually a storage controller that provides an internal switch for external storage attachment. In this approach, the storage controller intercepts and redirects I/O requests to the external storage as it does for internal storage. The actual I/O requests are themselves redirected. This implementation is referred to as *asymmetric virtualization* or *out-of-band*.

Figure 4-17 shows variations of the two virtualization approaches.

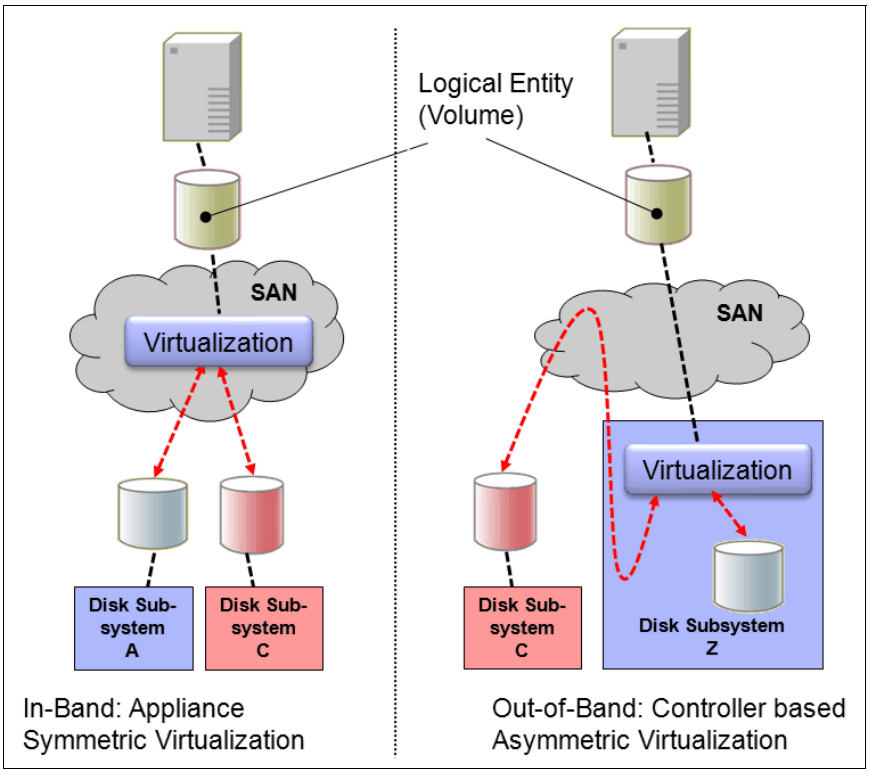

Figure 4-17 Overview of block-level virtualization architectures

The IBM Storwize V7000 solution provides a modular storage system that includes the capability to virtualize both external SAN-attached storage and its own internal storage. This solution is built upon the IBM SAN Volume Controller technology base, and uses technology from the IBM System Storage DS8000 family.

IBM Storwize V7000 system provides several configuration options that are aimed at simplifying the implementation process. It also provides automated wizards, called *Directed Maintenance Procedures* (DMP), to help resolve any events that might occur. A Storwize V7000 system is a midrange, clustered, scalable, and external virtualization device.

Included with a Storwize V7000 system is a graphical user interface (GUI) that enables storage to be deployed quickly and efficiently. The GUI runs on the Storwize V7000 system, so a separate console is not needed. The management GUI contains a series of preestablished configuration options that are called *presets*, and that use common settings to quickly configure objects on the system. Presets are available for creating volumes and FlashCopy mappings, and for setting up a RAID configuration.

The Storwize V7000 solution provides a choice of up to 1056 SAS drives for the internal storage in a clustered system. It uses SAS cables and connectors to attach to the optional expansion enclosures. In a clustered system, the Storwize V7000 can provide about 4 pebibytes (PiB) of internal raw capacity.

The Storwize V7000 solution consists of 1 - 4 control enclosures and, optionally, up to 80 expansion enclosures. It also supports the intermixing of the different expansion enclosures. Within each enclosure are two canisters. Control enclosures contain two node canisters, and expansion enclosures contain two expansion canisters.

4.4.3 IBM Storwize V7000 components

Storwize V7000 provides the following benefits:

- Brings enterprise technology to midrange storage.
- Specialty administrators are not required.
- Client setup and service can be done easily.
- The system can grow incrementally as storage capacity and performance needs change.
- Multiple storage tiers are in a single system with nondisruptive migration between them.
- Simple integration can be done into the server environment.

The Storwize V7000 subsystem consists of a set of drive enclosures. Control enclosures contain disk drives and two nodes (an I/O Group), which are attached to the SAN fabric. Expansion enclosures contain drives, and are attached to control enclosures.

The simplest use of Storwize V7000 is as a traditional RAID subsystem. The internal drives are configured into RAID arrays and virtual disks created from those arrays. IBM Storwize V7000 can also be used to virtualize other storage controllers.

IBM Storwize V7000 supports regular and solid-state drives (SSDs) and uses IBM System Storage Easy Tier to automatically place volume hot spots on better-performing storage. This section briefly explains the basic architecture components of Storwize V7000.

Nodes

Each IBM Storwize V7000 hardware controller is called a *node* or *node canister*. The node provides the virtualization for a set of volumes, cache, and copy services functions. Nodes are deployed in pairs, and multiple pairs make up a *clustered system* or *system*. A system can consist of 1 - 4 Storwize V7000 node pairs.

One of the nodes within the system is known as the *configuration node*. The configuration node manages the configuration activity for the system. If this node fails, the system chooses a new node to become the configuration node.

Because the nodes are installed in pairs, each node provides a failover function to its partner node during a node failure.

I/O Groups

IBM Storwize V7000 can have 1 - 4 pairs of node canisters known as *I/O Groups*. Storwize V7000 supports eight node canisters in the clustered system, which provides four I/O Groups.

When a host server performs I/O to one of its volumes, all of the I/Os for a specific volume are directed to the I/O Group. Also, under normal conditions, the I/Os for that specific volume are always processed by the same node within the I/O Group.

Both nodes of the I/O Group act as preferred nodes for their own specific subset of the total number of volumes that the I/O Group presents to the host servers (a maximum of 2048 volumes per I/O Group). However, each node also acts as a failover node for its partner node within the I/O Group. Therefore, a node takes over the I/O workload from its partner node, if required, with no effect to the server's application.

In an IBM Storwize V7000 environment that is using an active/active architecture, the I/O handling for a volume can be managed by both nodes of the I/O Group. Therefore, it is mandatory for servers that are connected through FC connectors to use multipath device drivers to be able to handle this capability.

The IBM Storwize V7000 I/O Groups are connected to the SAN so that all application servers that access volumes from the I/O Group have access to them. Up to 2048 host server objects can be defined in four I/O Groups.

If required, host servers can be mapped to more than one I/O Group in the IBM Storwize V7000 system. Therefore, they can access volumes from separate I/O Groups. You can move volumes between I/O Groups to redistribute the load between the I/O Groups.

However, moving volumes between I/O Groups cannot always be done concurrently with host I/O, and requires in some cases a brief interruption to remap the host. On the following website, you can check the compatibility of IBM Storwize V7000 Non Disruptive Volume Move (NDVM) function with your hosts:

http://www.ibm.com/support/docview.wss?uid=ssg1S1004622

Important: The active/active architecture provides availability to process I/Os for both controller nodes, and enables the application to continue running smoothly, even if the server has only one access route or path to the storage controller. This type of architecture eliminates the path and LUN thrashing typical of an active/passive architecture.

System

The system or clustered system consists of 1 - 4 I/O Groups. Certain configuration limitations are then set for the individual system. For example, the maximum number of volumes that are supported per system is 8192 (having a maximum of 2048 volumes per I/O Group), or the maximum managed disk supported is 32 petabytes (PB) per system.

All configuration, monitoring, and service tasks are performed at the system level. Configuration settings are replicated to all nodes in the system. To facilitate these tasks, a management Internet Protocol (IP) address is set for the system.

A process is provided to back up the system configuration data onto disk so that it can be restored during a disaster. Note that this method does *not* back up application data. Only IBM Storwize V7000 system configuration information is backed up. For the purposes of remote data mirroring, two or more systems must form a *partnership* before creating relationships between mirrored volumes.

System configuration backup: After backing up the system configuration, save the backup data on your hard disk (or at the least outside of the SAN). If you are unable to access IBM Storwize V7000, you do not have access to the backup data if it is on the SAN.

For details about the maximum configurations that are applicable to the system, I/O Group, and nodes, see the following link:

http://www.ibm.com/support/docview.wss?uid=ssg1S1004628

RAID

The IBM Storwize V7000 setup contains several internal drive objects, but these drives cannot be directly added to storage pools. The drives need to be included in a RAID to provide protection against the failure of individual drives.

These drives are referred to as members of the array. Each array has a RAID level. RAID levels provide various degrees of redundancy and performance, and have various restrictions regarding the number of members in the array.

IBM Storwize V7000 supports hot spare drives. When an array member drive fails, the system automatically replaces the failed member with a hot spare drive and rebuilds the array to

restore its redundancy. Candidate and spare drives can be manually exchanged with array members.

Each array has a set of goals that describe the location and performance of each array. A sequence of drive failures and hot spare takeovers can leave an array unbalanced (with members that do not match these goals). The system automatically rebalances such arrays when the appropriate drives are available.

The following RAID levels are available:

- RAID 0 (striping, no redundancy)
- RAID 1 (mirroring between two drives)
- RAID 5 (striping, can survive one drive fault)
- RAID 6 (striping, can survive two drive faults)
- RAID 10 (RAID 0 on top of RAID 1)

MDisks

An MDisk is the unit of storage that IBM Storwize V7000 virtualizes. This unit could be a logical volume on an external storage array presented to IBM Storwize V7000, or a RAID array that consists of internal drives. IBM Storwize V7000 can then allocate these MDisks into various storage pools. An MDisk is not visible to a host system on the SAN because it is internal or zoned only to the IBM Storwize V7000 system.

The MDisks are placed into storage pools where they are divided into several extents, which can be 16 - 8192 megabytes (MB), as defined by the storage administrator. See the following link for an overview of the total storage capacity that is manageable per system regarding the selection of extents:

http://www.ibm.com/support/docview.wss?uid=ssg1S1005250#_Extents

A volume is host-accessible storage that has been provisioned out of one *storage pool* or, if it is a mirrored volume, out of two storage pools. The maximum size of an MDisk is 1 PB. IBM Storwize V7000 system supports up to 4096 MDisks (including internal RAID arrays).

At any point in time, an MDisk is in one of the following four modes:

- Array

  Array mode MDisks are constructed from drives by using the RAID function. Array MDisks are always associated with storage pools.

- Unmanaged MDisk

  An MDisk is reported as *unmanaged* when it is not a member of any storage pool. An unmanaged MDisk is not associated with any volumes, and has no metadata stored on it. Storwize V7000 does not write to an MDisk that is in unmanaged mode, except when it attempts to change the mode of the MDisk to one of the other modes. Storwize V7000 can see the resource, but the resource is not assigned to a storage pool.

- Managed MDisk

  Managed mode MDisks are always members of a storage pool, and they contribute extents to the storage pool. Volumes (if not operated in image mode) are created from these extents. MDisks operating in managed mode might have metadata extents allocated on them, and can be used as *quorum disks*. This mode is the most common and normal mode for an MDisk.

- Image mode MDisk

  Image mode provides a direct block-for-block translation from the MDisk to the volume by using virtualization. This mode is provided to satisfy three major usage scenarios:

  – Image mode enables the virtualization of MDisks already containing data that was written directly, and not through IBM Storwize V7000. Rather, it was created by a direct-connected host. This mode enables a client to insert Storwize V7000 into the data path of an existing storage volume or LUN with minimal downtime. The image mode is typically used for data migration from old storage systems to new.

  – Image mode enables a volume that is managed by IBM Storwize V7000 to be used with the native copy services function provided by the underlying RAID controller. To avoid the loss of data integrity when IBM Storwize V7000 is used in this way, it is important that you disable the IBM Storwize V7000 cache for the volume.

  – IBM Storwize V7000 provides the ability to migrate to image mode, which enables IBM Storwize V7000 to export volumes and access them directly from a host without the IBM Storwize V7000 in the path.

  Each MDisk presented from an external disk controller has an online path count that is the number of nodes that have access to that MDisk. The *maximum count* is the maximum number of paths detected at any point in time by the system. The *current count* is what the system sees currently. A current value less than the maximum can indicate that SAN fabric paths have been lost.

  SSDs (flash drives) that are in IBM Storwize V7000 are presented to the cluster as MDisks. To determine whether the selected MDisk is a flash drive, click the link on the MDisk name to display the Viewing MDisk Details pane. The Viewing MDisk Details pane displays values for the Node ID, Node Name, and Node Location attributes.

Quorum disk

A *quorum disk* is an MDisk that contains a reserved area for use exclusively by the system. The system uses quorum disks to break a tie when exactly half the nodes in the system remain after a SAN failure. This situation is referred to as *split brain*. Quorum functionality is not supported on flash drives in IBM Storwize V7000. There are three candidate quorum disks. However, only one quorum disk is active at any time.

Disk tier

It is likely that the MDisks (LUNs) presented to the IBM Storwize V7000 system have various performance attributes due to the type of disk or RAID on which they reside. The MDisks can be on 15,000 disk revolutions per minute (RPMs) Fibre Channel or SAS disks, Nearline SAS or Serial Advanced Technology Attachment (SATA) disks, or even flash drives.

Therefore, a storage tier attribute is assigned to each MDisk, with the default being **`enterprise`**. A tier 0 (zero)-level disk attribute (**`ssd`**) is available for flash drives, and a tier 2-level disk attribute (**`nearline`**) is available for nl-sas.

Storage pool

A *storage pool* is a collection of up to 128 MDisks that provides the pool of storage from which volumes are provisioned. A single system can manage up to 128 storage pools. The size of these pools can be changed (expanded or shrunk) at run time by adding or removing MDisks, without taking the storage pool or the volumes offline.

At any point in time, an MDisk can only be a member in one storage pool, except for image mode volumes.

Each MDisk in the storage pool is divided into several extents. The size of the extent is selected by the administrator when the storage pool is created, and cannot be changed later. The size of the extent can be 16 - 8192 MB.

It is a leading practice to use the same extent size for all storage pools in a system. This approach is a prerequisite for supporting volume migration between two storage pools. If the storage pool extent sizes are not the same, you must use volume mirroring.

Figure 4-18 shows an overview of a Storwize clustered system with an I/O Group.

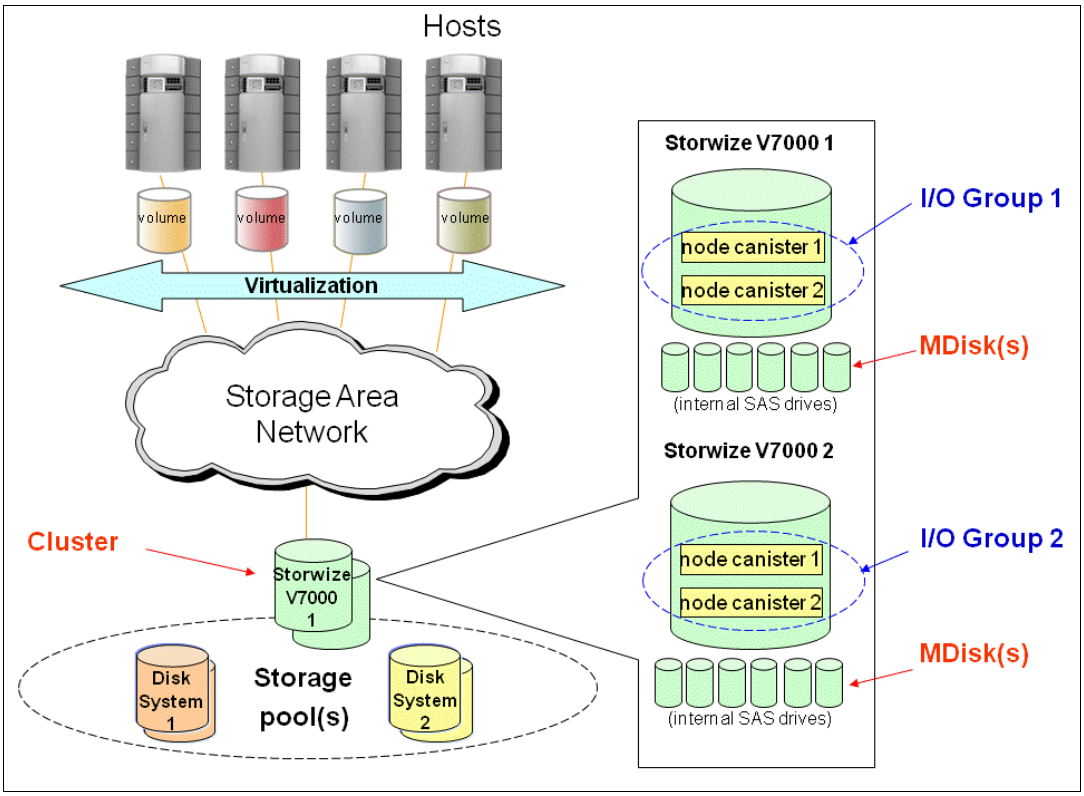

Figure 4-18 IBM Storwize V7000 clustered system

IBM Storwize V7000 limits the number of extents in a system to 2^{22} = ~4 million. Because the number of addressable extents is limited, the total capacity of a Storwize V7000 system depends on the extent size that is chosen by the Storwize V7000 administrator. The capacity numbers that are specified in Table 4-6 for a Storwize V7000 system assume that all of the defined storage pools have been created with the same extent size.

Table 4-6 Extent size-to-addressability matrix

Extent size maximum	System capacity	Extent size maximum	System capacity
16 MB	64 terabytes (TB)	512 MB	2 PB
32 MB	128 TB	1024 MB	4 PB
64 MB	256 TB	2048 MB	8 PB
128 MB	512 TB	4096 MB	16 PB
256 MB	1 PB	8192 MB	32 PB

For most systems, a capacity of 1 - 2 PB is sufficient. In a Storwize V7000 environment, generally use the default size of 1 gigabyte (GB) as the standard extent size.

Volumes

Volumes are logical disks that are presented to the host or application servers by Storwize V7000. The hosts cannot see the MDisks. They can only see the logical volumes created from combining extents from a storage pool.

There are three types of volumes: Striped, sequential, and image. These types are determined by how the extents are allocated from the storage pool:

- A volume created in striped mode has extents allocated from each MDisk in the storage pool in a round-robin fashion.
- With a sequential mode volume, extents are allocated sequentially from an MDisk.
- Image mode is a one-to-one mapped extent mode volume.

Using striped mode is the best method to use for most cases. However, sequential extent allocation mode can slightly increase the sequential performance for certain workloads.

Figure 4-19 shows the striped volume mode and sequential volume mode, and illustrates how the extent allocation from the storage pool differs.

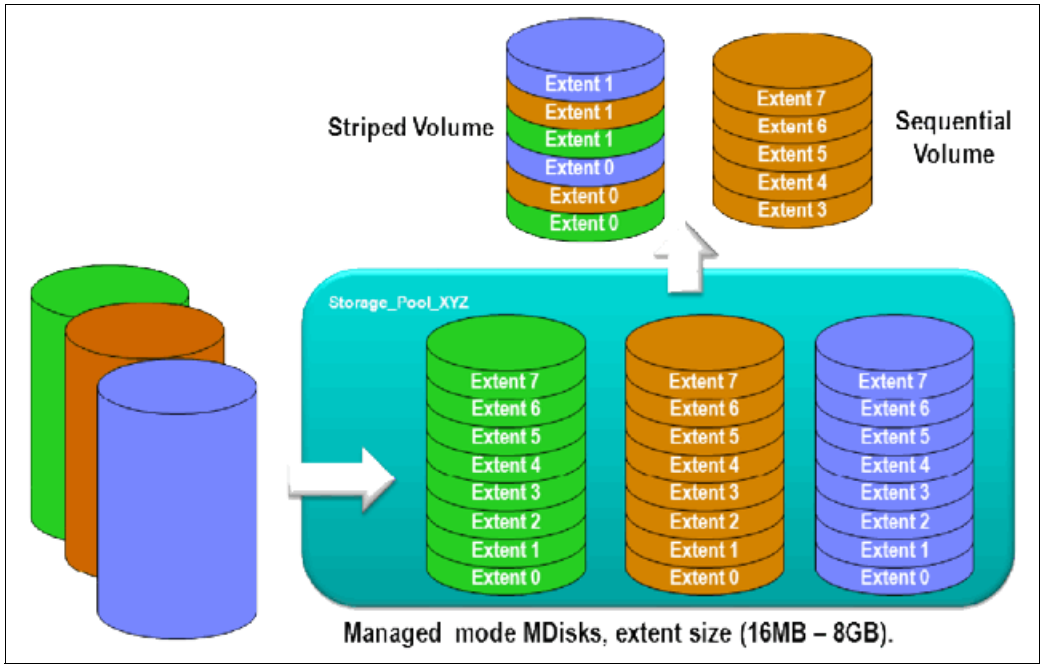

Figure 4-19 Storage pool extents overview

You can allocate the extents for a volume in many ways. The process is under full user control at volume creation time, and can be changed at any time by migrating single extents of a volume to another MDisk within the storage pool.

Hosts

Volumes can be mapped to a *host* to allow access for a specific server to a set of volumes. A host in Storwize V7000 is a collection of host bus adapter (HBA) worldwide port names (WWPNs) or iSCSI qualified names (IQNs) defined on the specific server. Note that iSCSI

names are internally identified by “fake” WWPNs, or WWPNs that are generated by the IBM Storwize V7000.

Volumes can be mapped to multiple hosts, for example, a volume that is accessed by multiple hosts of a server system. iSCSI is an alternative means of attaching hosts. However, all communication with back-end storage subsystems, and with other Storwize V7000 systems, is still through FC.

Node failover can be handled without installing a multipath driver on the iSCSI server. An iSCSI-attached server can simply reconnect after a node failover to the original target IP address, which is now presented by the partner node. To protect the server against link failures during network or HBA failures, using a multipath driver is mandatory.

Volumes are LUN-masked to the host’s HBA WWPNs by a process called *host mapping*. Mapping a volume to the host makes it accessible to the WWPNs or IQNs that are configured on the host object. For a SCSI over Ethernet connection, the IQN identifies the iSCSI target (destination) adapter. Host objects can have both IQNs and WWPNs.

Easy Tier

Easy Tier is a performance function that automatically migrates or moves extents off a volume to, or from, one MDisk storage tier to another MDisk storage tier. In Storwize V7000, the Easy Tier automatically moves extents between highly used and less-used MDisks within the same storage tier. This function is called *Storage Pool Balancing*, and it is enabled by default without any need for licensing.

Balancing cannot be disabled by the user. Easy Tier monitors the host I/O activity and latency on the extents of all volumes with the Easy Tier function turned on in a multitier storage pool, over a 24-hour period.

New in Storwize family software V7.3: Easy Tier V3 integrates the automatic functionality to balance the workloads between highly used and less-used MDisks within the same tier. It is enabled by default, cannot be disabled by the user, and does not need an Easy Tier license.

Next, it creates an extent migration plan based on this activity, and then dynamically moves high-activity (or *hot*) extents to a higher disk tier in the storage pool. It also moves extents whose activity has dropped off (or *cooled*) from the high-tier MDisks back to a lower-tiered MDisk.

Easy Tier: The Easy Tier function can be turned on or off at the storage pool level and the volume level. It supports any combination of three tiers within the system. Flash drives are always marked as Tier 0. Turning off Easy Tier does not disable Storage Pool Balancing.

To experience the potential benefits of using Easy Tier in your environment before installing expensive flash drives, turn on the Easy Tier function for a *single-level* storage pool. Next, turn on the Easy Tier function for the volumes within that pool. Easy Tier then starts monitoring activity on the volume extents in the pool.

Easy Tier creates a report every 24 hours, providing information about how Easy Tier behaves if the tier were a multitiered storage pool. So, even though Easy Tier extent migration is not possible within a single-tiered pool, the Easy Tier statistical measurement function is available.

The Easy Tier function can make it more appropriate to use smaller storage pool extent sizes. The usage statistics file can be off-loaded from Storwize V7000. You can then use the IBM Storage Tier Advisor Tool to create a summary report.

Thin provisioning

Volumes can be configured to be either *thin-provisioned* or *fully allocated*. A thin-provisioned volume behaves regarding application reads and writes as though they were fully allocated. When creating a thin-provisioned volume, the user specifies two capacities:

- The real physical capacity allocated to the volume from the storage pool
- The virtual capacity available to the host

In a fully allocated volume, these two values are the same.

Therefore, the real capacity determines the quantity of MDisk extents that is initially allocated to the volume. The virtual capacity is the capacity of the volume reported to all other Storwize V7000 components (for example, FlashCopy, Cache, and remote copy), and to the host servers. The real capacity is used to store both the user data and the metadata for the thin-provisioned volume. The real capacity can be specified as an absolute value, or a percentage of the virtual capacity.

Thin-provisioned volumes can be used as volumes assigned to the host, by FlashCopy to implement thin-provisioned FlashCopy targets, and also with the mirrored volumes feature. When a thin-provisioned volume is initially created, a small amount of the real capacity is used for initial metadata.

Write I/Os to grains of the thin volume that *were not* previously written to cause grains of the real capacity to be used to store metadata and the actual user data. Write I/Os to grains that *were* previously written to update the grain where data was previously written. The grain size is defined when the volume is created, and can be 32 kilobytes (KB), 64 KB, 128 KB, or 256 KB. The default grain size is 256 KB, which is the strongly suggested option. If you select 32 KB for the grain size, the volume size cannot exceed 260,000 GB.

The grain size cannot be changed after the thin-provisioned volume has been created. Generally, smaller grain sizes save space, but require more metadata access, which can adversely affect performance. If you are not going to use the thin-provisioned volume as a FlashCopy source or target volume, use 256 KB to maximize performance. If you are going to use the thin-provisioned volume as a FlashCopy source or target volume, specify the same grain size for the volume and for the FlashCopy function.

Figure 4-20 illustrates the thin-provisioning concept.

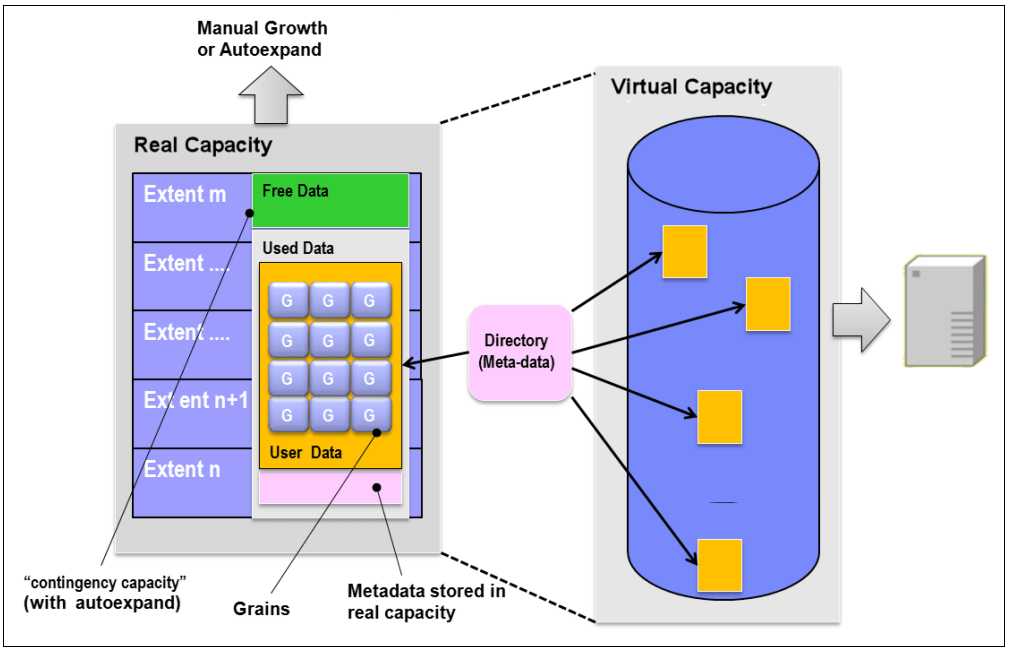

Figure 4-20 Conceptual diagram of a thin-provisioned volume

Thin-provisioned volumes store both user data and metadata. Each grain of data requires metadata to be stored. Therefore, the I/O rates that are obtained from thin-provisioned volumes are less than the I/O rates that are obtained from fully allocated volumes.

The metadata storage fixed use is never greater than 0.1% of the user data. The fixed resource use is independent of the virtual capacity of the volume. If you are using thin-provisioned volumes in a FlashCopy map, for the best performance, use the same grain size as the map grain size. If you are using the thin-provisioned volume directly with a host system, use a small grain size.

The real capacity of a thin volume can be changed if the volume is not in image mode. Increasing the real capacity enables a larger amount of data and metadata to be stored on the volume. Thin-provisioned volumes use the real capacity that is provided in ascending order as new data is written to the volume. If the user initially assigns too much real capacity to the volume, the real capacity can be reduced to free storage for other uses.

A thin-provisioned volume can be configured to *autoexpand*. This feature causes Storwize V7000 to automatically add a fixed amount of extra real capacity to the thin volume as required. Autoexpand attempts to maintain a fixed amount of unused real capacity for the volume, which is known as the *contingency capacity*.

The contingency capacity is initially set to the real capacity that is assigned when the volume is created. If the user modifies the real capacity, the contingency capacity is reset to be the difference between the used capacity and the real capacity. A volume that is created without the autoexpand feature, and therefore has a zero contingency capacity, goes offline as soon as the real capacity is fully used and needs to expand.

Autoexpand does not cause the real capacity to grow much beyond the virtual capacity. The real capacity can be manually expanded to more than the maximum that is required by the current virtual capacity, and the contingency capacity is recalculated.

To support the auto expansion of thin-provisioned volumes, the storage pools from which they are allocated have a configurable capacity warning. When the used capacity of the pool exceeds the warning capacity, a warning event is logged. For example, if a warning of 80% has been specified, the event is logged when 20% of the free capacity remains.

A thin-provisioned volume can be converted nondisruptively to a fully allocated volume, or vice versa, by using the volume mirroring function. For example, you can add a thin-provisioned copy to a fully allocated primary volume, and then remove the fully allocated copy from the volume after they are synchronized.

The fully allocated to thin-provisioned migration procedure uses a zero-detection algorithm so that grains containing all zeros do not cause any real capacity to be used.

Real-time Compression

Compressed volumes are a special type of volume where data is compressed as it is written to disk, saving additional space. To use the compression function, you must obtain the IBM Real-time Compression license. With the IBM Storwize V7000 model (2076-524), you already have one compression acceleration adapter included in the base product, and you can add one more.

It is also suggested to upgrade your memory to 64 GB for best use of Real-time Compression. Enabling compression on Storwize V7000 nodes does not affect non-compressed host-to-disk I/O performance. Like thin-provisioned volumes, compressed volumes have virtual, real, and used capacities. Use the following guidelines before working with compressed volumes:

- Real capacity is the extent space that is allocated from the storage pool. The real capacity is also set when the volume is created and, like thin-provisioned volumes, can be expanded or shrunk down to the used capacity.
- Virtual capacity is available to hosts. The virtual capacity is set when the volume is created, and can be expanded or shrunk afterward.
- Used capacity is the amount of real capacity that is used to store client data and metadata after compression.
- Capacity before compression is the amount of client data that has been written to the volume and then compressed. The capacity before compression does not include regions where zero data is written to unallocated space.
- An I/O Group can contain a maximum of 200 compressed volumes and compressed volume mirrors.
- You can also monitor information about compression usage to determine the savings to your storage capacity when volumes are compressed. To monitor system-wide compression savings and capacity, select **Monitoring** → **System** and either select the system name or **Compression View**. You can compare the amount of capacity that is used before compression is applied to the capacity that is used for all compressed volumes.

  In addition, you can view the total percentage of capacity savings when compression is used on the system. Furthermore, you can also monitor compression savings across individual pools and volumes. For volumes, you can use these compression values to determine which volumes have achieved the highest compression savings.

Cache

The primary benefit of storage cache is to improve I/O response time. Reads and writes to a magnetic disk drive suffer from both seek and latency time at the drive level, which can result in from 1 millisecond (ms) - 10 ms of response time (for an enterprise-class disk).

The IBM Storwize V7000 nodes combined with IBM Spectrum Virtualize software V7.6 provide 32 GB memory per node, or 64 GB (128 GB) per I/O Group, or 256 GB (512 GB) per system. The Storwize V7000 provides a semi-flexible cache model, and the node's memory can be used as read or write cache, either one as an I/O workload cache. The size of the write cache is maximally 12 GB of the node's memory. The remaining part of the memory is split between read cache allocation and compression allocation.

When data is written by the host, the preferred node saves the data in its cache. Before the cache returns completion to the host, the write must be mirrored to the partner node, or copied into the cache of its partner node, for availability reasons. After having a copy of the written data, the cache returns completion to the host. A volume that has not received a write update during the last two minutes will automatically have all modified data destaged to disk.

Note: Optional cache upgrade of 32 GB on Storwize V7000 is reserved for RtC, and it is not used when RtC is disabled.

Starting with Storwize V7000, the concept of the cache architecture has been changed. Storwize V7000 now distinguishes between upper and lower cache that enables the system to be more scalable:

- Required for support beyond 8192 volumes
- Required for support beyond 8 node clusters
- Required for 64-bit addressing beyond 28 GB
- Required for larger memory in nodes
- Required for more processor cores
- Required for improved performance and stability

The architectural overview is shown in Figure 4-21.

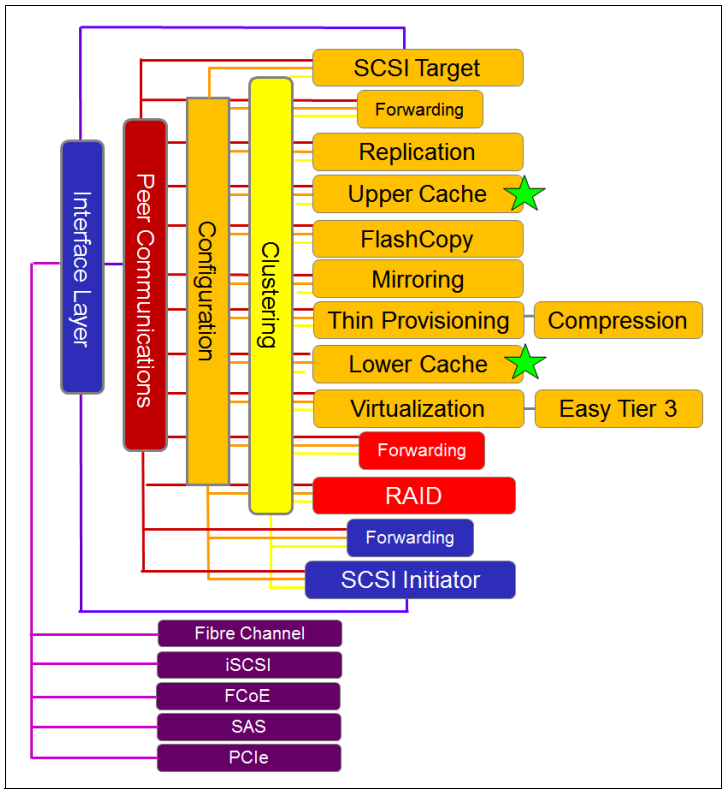

Figure 4-21 New cache architecture

If one node of an I/O Group is missing, due to a restart or a hardware failure, the remaining node empties all of its write cache and proceeds in operation mode. This is referred to as *write-through mode*. A node operating in write-through mode writes data directly to the disk subsystem before sending an `I/O complete` status message back to the host. Running in this mode can degrade the performance of the specific I/O Group.

Write cache is partitioned by storage pool. This feature restricts the maximum amount of write cache that a single storage pool can allocate in a system. Table 4-7 shows the upper limit of write-cache data that a single storage pool in a system can occupy.

Table 4-7 Upper limit of write cache per storage pool

One storage pool	Two storage pools	Three storage pools	Four storage pools	More than four storage pools
100%	66%	40%	33%	25%

Storwize V7000 will treat part of its physical memory as non-volatile. *Non-volatile* means that its contents are preserved across power losses and resets. Bitmaps for FlashCopy and Remote Mirroring relationships, the virtualization table, and the write cache are items in the non-volatile memory.

During a disruption or external power loss, the physical memory is copied to a file in the file system on the node's internal disk drive so that the contents can be recovered when external power is restored. The functionality of uninterruptible power supply units is provided by internal batteries, which are delivered with each node's hardware. The batteries ensure that there is sufficient internal power to keep a node operational to perform this dump when the external power is removed. After dumping the content of the non-volatile part of the memory to disk, Storwize V7000 shuts down.

IBM Storwize V7000F

IBM Storwize V7000F is an all-flash, highly scalable, virtualized, enterprise-class storage system that is designed to consolidate workloads into a single system for ease of management, reduced costs, superior performance, and high availability.

IBM Storwize all-flash storage systems

Today's organizations need to meet ever-changing demands for storage while also improving data economics. IT staff must deliver more services faster and more efficiently, enable real-time insight, and support more customer interaction. All-flash storage can meet these requirements for fast, reliable, and consistent access to data. The technology is accessible to more businesses than ever.

IBM Storwize family offers the new IBM Storwize V7000F and IBM Storwize V5000F systems as all-flash, virtualized, enterprise- class storage systems designed to deliver the high performance needed to derive real- time insights from business data combined with advanced management functions. These all-f lash solutions address the need to handle the massive volumes of data created by today's demanding cloud, analytics, and traditional applications.

Built with IBM Spectrum Virtualize software, Storwize V7000F and Storwize V5000F provide the latest storage technologies for unlocking the business value of stored data. These technologies include virtualization, IBM Real-time Compression, thin provisioning, snapshots, cloning, replication, and high- availability configurations. They are designed to deliver outstanding performance, efficiency, ease of use, and dependability for organizations of all sizes.

The IBM Storwize all-flash storage systems have the following highlights:

- High performance

  The all-flash Storwize solutions deliver both high performance and dramatically improved data economics. They help accelerate multiple workloads for faster decision making with best-of-class performance. The solutions are most often deployed to support high-performance enterprise applications, such as databases like IBM DB2, Oracle, and SAP, as well as virtual desktop infrastructure (VDI) and server virtualization.

- Efficiency

  IBM Spectrum Virtualize software is designed to deliver extraordinary levels of efficiency, helping to revolutionize data economics and drive down costs for cloud, analytics, virtual server, and other enterprise-class deployments. As a result, organizations no longer have to choose between performance and efficiency. They can get both within a single storage solution.

  IBM Real-time Compression with hardware acceleration, available in Storwize V7000F, enables the system to deliver higher performance for compressed data than traditional systems offer for uncompressed data. This performance enables its use for practically all data types. It is designed to enable storing up to five times as much data in the same physical drive space by compressing data as much as 80 percent. The benefits include reduced acquisition cost (because less hardware is required), reduced rack space, and lower power and cooling costs throughout the lifetime of the system. And, when combined with external data virtualization, Real-time Compression can significantly enhance the usable capacity of existing storage systems, extending their useful life even further.

  When replicating block data for business continuity, Storwize V7000F and Storwize V5000F can use IP network connections for simplicity and lower cost. Integrated Bridgeworks SANrockIT technology helps improve network utilization up to three times compared with traditional approaches, which can help reduce networking costs and accelerate replication cycles.

- Data virtualization

  IBM Spectrum Virtualize data virtualization technology helps insulate applications from physical storage. This characteristic enables applications to run without disruption, even when changes are made to the storage infrastructure.

  Storwize V7000F and Storwize V5000F also extend data virtualization to other storage systems. When virtualized, data in a disk system becomes part of the Storwize system, and it can be managed in the same way as internal drives. Data in external disk systems inherits all the Storwize functional richness and ease-of-use features, including advanced replication, high performance thin provisioning, data migration, and Real-time Compression. Virtualizing external storage helps improve administrator productivity and boost storage utilization while also enhancing and extending the value of existing storage investments.

  Moving data is one of the most common causes of planned downtime. Data virtualization enables moving data from existing storage into the new system or between arrays, while maintaining access to the data. This function can be used when replacing older storage with newer storage, as part of load-balancing work, or when moving data in a tiered storage infrastructure from disk drives to flash.

  Data virtualization can improve efficiency and business value. Nondisruptive 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 can help avoid penalties and additional maintenance charges for lease extensions. The result can be real cost savings for your business.

- Ease of use

  Storwize V7000F and Storwize V5000F are easy to use from the start. An intuitive management interface enables administrators to easily manage the solution. IBM Spectrum Control, based on IBM Tivoli Storage Productivity Center, can also provide organizations with an end-to-end view of storage health, long-term performance analytics, and capacity statistics for Storwize V7000F, Storwize V5000F, and the surrounding storage infrastructure.

  What's more, IBM Spectrum Virtualize technologies, including Real-time Compression and IP replication with Bridgeworks SANrockIT technology, operate automatically and require little or no customization.

- Dependability

  Storwize V7000F and Storwize V5000F are part of the proven IBM Storwize family, with more than 90,000 enclosures and 2 exabytes of capacity deployed in organizations worldwide.

  With their virtualized storage design and tight affinity with IBM PowerVM, OpenStack, Microsoft ODX, VMware vSphere v6, and VMware vSphere Virtual Volumes (VVOL), Storwize V7000F and Storwize V5000F are an ideal complement for virtualized servers that are at the heart of cloud deployments.

4.4.4 VersaStack Solution for IBM Storwize V7000 and V7000 Unified

The VersaStack solution combines the performance and innovation of Cisco UCS Integrated Infrastructure, which includes the Cisco Unified Computing System (Cisco UCS), Cisco Nexus and Cisco MDS 9000 Family switches, and Cisco UCS Director, with the performance and efficiency of the IBM Storwize storage system. The IBM Storwize V7000 includes technologies that both complement and enhance virtual environments with built-in functions such as IBM Data Virtualization, Real-time Compression, and Easy Tier that deliver extraordinary levels of performance and efficiency.

For more information, see this website:

http://www.ibm.com/systems/storage/disk/cisco-versastack.html

4.4.5 More information

For more information about the IBM Storwize family, see the following materials

- *Implementing the IBM Storwize V7000 and IBM Spectrum Virtualize V7.6*, SG24-7938
- *Implementing the IBM Storwize V7000 Gen2*, SG24-8244
- *IBM Real-time Compression in IBM SAN Volume Controller and IBM Storwize V7000*, REDP-4859
- *IBM System Storage SAN Volume Controller and Storwize V7000 Best Practices and Performance Guidelines*, SG24-7521
- *IBM SAN Volume Controller and Storwize Family Native IP Replication*, REDP-5103
- *Implementing the IBM Storwize V7000 Unified Disk System*, SG24-8010
- *IBM Storwize V7000, Spectrum Virtualize, HyperSwap, and VMware Implementation*, SG24-8317

5

IBM FlashSystem family

The IBM FlashSystem products are designed to help companies improve performance and reduce server sprawl, power consumption, cooling, and floor space requirements. All of this in turn can help clients save money, improve performance, and invest more in innovation.

The IBM FlashSystem family can help you take control of your storage, harness existing data, and manage your IT infrastructure in a much more cloud-like fashion.

This chapter describes the IBM FlashSystem storage systems, which can either complement or replace traditional hard disk arrays.

This chapter includes the following sections:

- Overview
- Why flash matters
- IBM FlashSystem 900 storage
- IBM FlashSystem V9000 storage

5.1 Overview

The IBM FlashSystem portfolio of flash storage delivers high performance, efficiency, and reliability for shared enterprise storage environments, helping clients around the world address performance issues with their most important applications and infrastructure.

These storage systems can either complement or replace traditional hard disk drive (HDD) arrays, even arrays that incorporate solid-state devices (SSDs) or other flash technology, in many applications. Applications include online transaction processing (OLTP), business intelligence (BI), online analytical processing (OLAP), virtual desktop infrastructures, high-performance computing, and content delivery solutions (such as cloud storage and video on demand).

5.2 Why flash matters

Flash technology is a critical component of the new IBM approach to the design and deployment of storage infrastructures. It is part of a holistic approach that uses solid-state storage technology with disk and tape technologies to solve complex problems. Flash delivers performance and value to boost the applications that run your business.

Today, global organizations depend on the ability to unlock business insights from massive volumes of data, and they need to be able to do it faster than the competition. Flash technology changed the economics of today's data center by eliminating storage bottlenecks. Flash technology redefined the performance, efficiency, and reliability of data-intensive applications. In fact, the IBM FlashSystem storage offerings provide extreme input/output operations per second (IOPS) performance in a power-efficient footprint.

The IBM FlashSystem storage family delivers high performance and efficiency in an easy-to-integrate offering so that businesses can more readily compete in the high velocity marketplace. Extreme performance, IBM MicroLatency, macro efficiency, and enterprise-grade reliability make the IBM FlashSystem a powerful and cost-effective tool for accelerating OLAP systems and gaining competitive advantage. More importantly, the capabilities and capacity of the FlashSystem arrays enable commercial and governmental enterprises to address multiple compute challenges in current 24x7x365 operational environments while empowering growth and innovation in the future.

5.3 IBM FlashSystem 900 storage

Today's global organizations depend on the ability to unlock business insights from massive volumes of data. With IBM FlashSystem 900, they can make faster decisions based on real-time insights and unleash the power of the most demanding applications, including online transaction processing (OLTP) and analytics databases, virtual desktop infrastructures (VDIs), technical computing applications, and cloud environments.

Easy to deploy and manage, the IBM FlashSystem 900 is designed to accelerate the applications that drive business. Powered by IBM FlashCore™ Technology, IBM FlashSystem provides the following characteristics:

- Extreme performance with IBM MicroLatency delivers 90 µs write latency and 155 µs read latency
- Macro efficiency lowers operating costs and increases efficiency of IT infrastructure by using much less power and space than traditional HDD and SSD solutions
- Enterprise reliability protects critical assists with two-dimensional RAID protection, IBM Variable Stripe RAID™, redundant and hot-swappable components, and concurrent code loads

FlashSystem 900 can be configured in capacity points as low as 2.4 TB usable after RAID 5 protection. Coupled with 10 Gbps iSCSI, FlashSystem is positioned to bring extreme performance to small and medium businesses (SMBs) and growth markets. Figure 5-1 shows the FlashSystem 900.

Figure 5-1 IBM FlashSystem 900

FlashSystem V9000 is available with either a one-year or a three-year warranty, so you can select the warranty period that best addresses your business and financial needs.

FlashSystem 900 has these characteristics:

- FlashSystem 900 is configurable with 2.4 - 57 TB of usable capacity for scalability and flexibility.
- FlashSystem 900 provides flexible interface types that include Fibre Channel (FC), Fibre Channel over Ethernet (FCoE), iSCSI, and InfiniBand to easily integrate into existing SAN environments.
- FlashSystem 900 offers hardware-based AES 256 data-at-rest encryption with no performance impact.

5.3.1 FlashCore technology

At the heart of FlashSystem 900 is IBM FlashCore Technology, which consists of the following three key elements:

- Hardware-accelerated architecture that is engineered for flash, with a hardware-only data path
- IBM MicroLatency modules that are designed for low latency, density, and reliability
- IBM Advanced Flash Management that improves MLC flash endurance 9x over standard implementations without sacrificing latency

Figure 5-2 shows IBM FlashCore technology.

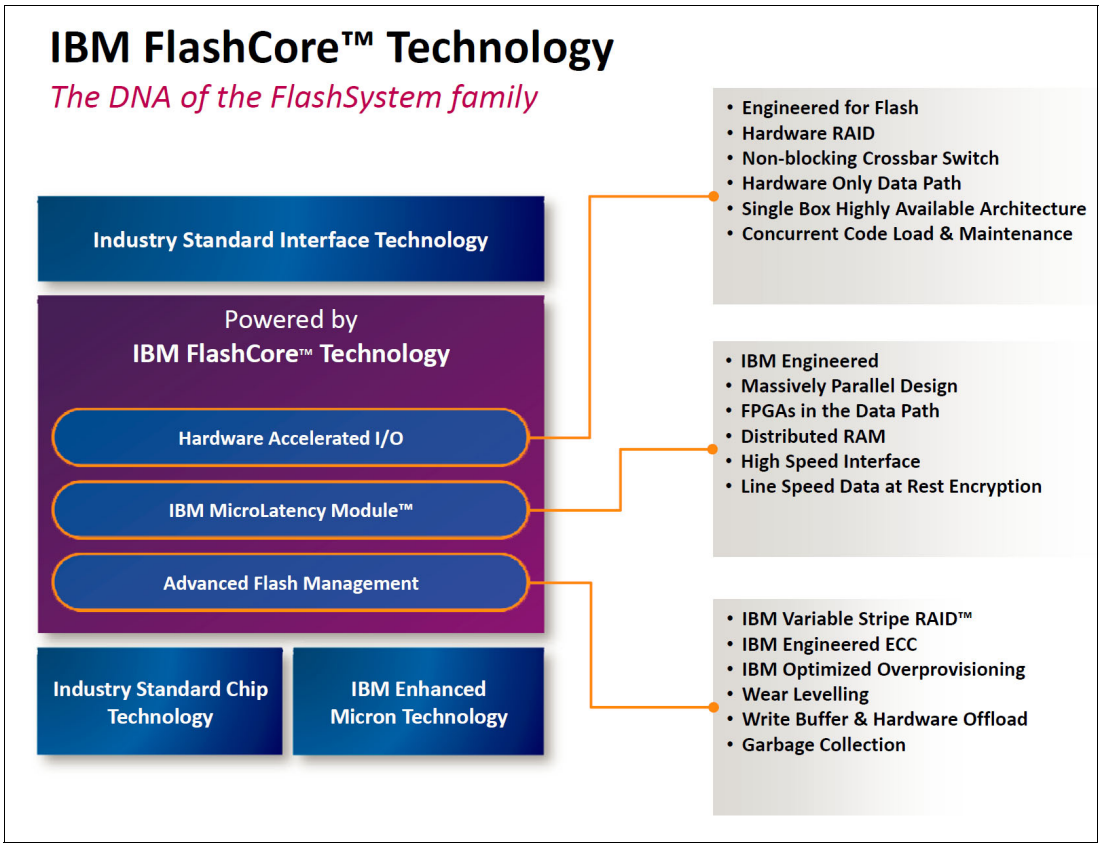

Figure 5-2 IBM FlashCore Technology

5.3.2 Key features

FlashSystem 900 provides extreme performance with IBM MicroLatency, macro efficiency, and enterprise reliability.

Scalability and performance

FlashSystem 900 has the following scalability and performance features:

- IBM FlashCore technology for consistently high performance at lower cost
  - 90us/155us read/write latency
  - Up to 1.1 million random read 4 K IOPS
  - Up to 10 GBps read bandwidth
- IBM-enhanced Micron MLC technology for higher storage density and improved endurance
- Flexible interface types of Fibre Channel, FCoE, iSCSI, and InfiniBand
  - Up to 16 ports of 8 Gbps or eight ports of 16 Gbps Fibre Channel
  - Up to 16 ports of 10 Gbps FCoE
  - Up to 16 ports of 10 Gbps iSCSI
  - Up to eight ports of 40 Gbps quad data rate (QDR) InfiniBand

- Slots for up to 12 hot-swappable IBM MicroLatency storage modules (1.2 TB, 2.9 TB, or 5.7 TB)
- Configurable 2.2- 57 TB of capacity for increased flexibility

Enterprise-class reliability features

FlashSystem 900 delivers the following enterprise-class reliability features:

- Concurrent code load enables customer applications to remain online during firmware upgrades to all components, including the IBM MicroLatency modules.
- Hot-swappable IBM MicroLatency modules by way of tool-less front panel access. If a MicroLatency module failure occurs, critical customer applications can remain online while the defective module is replaced.
- All active components are redundant and hot-swappable, including IBM MicroLatency modules, RAID controllers, management modules, interface, batteries, fans, and power supplies.
- All components are easily accessible through the front or rear of the enclosure. IBM FlashSystem 900 does not need to be moved in the rack and no top access panels or cables need to be extended.
- Two-dimensional (2D) Flash RAID, which consists of IBM Variable Stripe RAID and system-wide RAID 5. Variable Stripe RAID technology helps reduce downtime and maintain performance and capacity during partial or full flash chip failures. System-wide RAID 5, with easily accessed hot swappable flash modules, helps prevent data loss and promote availability.
- IBM Advanced Flash Management improves flash endurance 9x over standard implementations with IBM engineered ECC, advanced wear leveling, and proprietary garbage collection, relocation, and block picking algorithms.

Energy and space efficiency

FlashSystem 900 has the following energy and space efficiency characteristics:

- 40% greater density than previous generation
- 4x greater capacity than competitive all-flash products
- Reduced power, space, and cooling with only 625 watts typical operating power and 2U form factor

Manageability and security

FlashSystem 900 offers the following manageability and security features:

- Hardware-accelerated AES-XTS 256-bit data at rest encryption including Hot Encryption Activation and Hot Encryption Re-key
- Intuitive graphical user interface (GUI) and command-line interface (CLI) which is available in any supported browser
- Simple Network Management Protocol (SNMP)
- Email alerts
- Syslog redirect to send system log messages to another host

Figure 5-3 shows IBM FlashSystem 900 Management GUI (front view of the system).

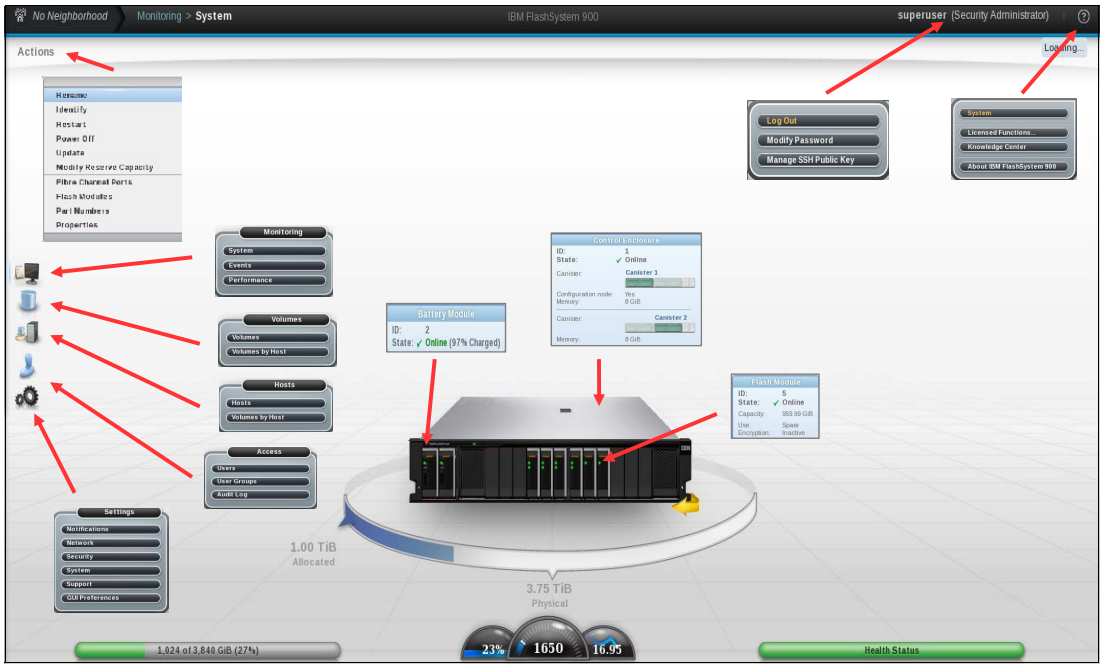

Figure 5-3 IBM FlashSystem 900 Management GUI (front view of the system)

5.3.3 More information

For more information, see the following links:

- *FlashSystem 900 Product Guide*, TIPS1261
- *Implementing IBM FlashSystem 900*, SG24-8271
- IBM FlashSystem 900 product page

  http://www.ibm.com/systems/storage/flash/900/

5.4 IBM FlashSystem V9000 storage

IBM FlashSystem V9000 is a comprehensive all-flash enterprise storage solution that delivers the full capabilities of IBM FlashCore technology plus a rich set of software-defined storage features. These features include IBM Real-time Compression, dynamic tiering, thin provisioning, snapshots, cloning, replication, data copy services, and HyperSwap for high availability. With the release of FlashSystem V9000 Software V7.6, extra functions are available including VMWare vSphere Virtual Volumes (VVOL) support, Encryption support for externally virtualized storage, and an integrated compretimator utility. Support also exists for VMware 6 that will enhance and improve scalability in a VMware environment.

IBM FlashSystem V9000 Software version 7.6 replaces version 7.5 and is available to all IBM FlashSystem V9000 customers with current warranty or software maintenance agreements.

Figure 5-4 shows the IBM FlashSystem V9000.

Figure 5-4 IBM FlashSystem V9000

FlashSystem V9000 has these characteristics:

- FlashSystem V9000 with IBM FlashCore Technology drives real-time analytical insights with up to 50x faster performance than enterprise disk systems.
- FlashSystem V9000 scales up to 2.5 million IOPS and over 2 PB effective capacity with IBM MicroLatency, under a single, fully integrated management interface.
- FlashSystem V9000 now supports encryption for externally virtualized storage, even if the supported virtualized array does not have encryption capabilities.
- FlashSystem V9000 with VVOL support, allows more efficient operations and control of external storage resources.
- FlashSystem V9000 with HyperSwap offers simplified setup and management through a GUI.
- FlashSystem V9000 now incorporates Integrated Comprestimator, a key sizing tool to estimate how much capacity savings the customer can expect with IBM Real-time Compression.
- FlashSystem V9000 now includes IP quorum base support for lower-cost IP attached hosts as a quorum disk.
- FlashSystem V9000 supports a new four port 16 Gb Fibre Channel host adapter (feature code #AF44).
- FlashSystem V9000 support with Virtual Storage Center (VSC) includes performance statistics and metrics of monitored storage systems and switches. These reports can be viewed in the VSC web-based GUI or stand-alone GUI. For more information, see the VSC Knowledge Center:

  http://www.ibm.com/support/knowledgecenter/SS5R93_5.2.9/com.ibm.spectrum.sc.doc/fqz0_t_starting_tpc_web_based.html

5.4.1 Key features

FlashSystem V9000 is a rack-mount shared flash memory device that is based on IBM enhanced multi-level cell (MLC) flash technology. It provides macro efficiency with up to 57 TB of protected capacity in a 6U form factor, enterprise reliability through IBM Variable Stripe RAID and two-dimensional flash RAID, and extreme performance with MicroLatency. FlashSystem V9000 provides advanced data services, including business continuity with replication services (Metro Mirror and Global Mirror), data protection with FlashCopy

services, and higher storage efficiency with thin provisioning, Real-time Compression, IBM Easy Tier 3rd Generation, external virtualization, space-efficient copies, and HyperSwap capability. The FlashSystem V9000 baseline configuration is composed of the following components:

- Two FlashSystem V9000 control enclosures

  FlashSystem V9000 Control Enclosure provides advanced data services for flash memory. It is a 2U, 19-inch rack-mount enclosure. FlashSystem V9000 includes two control enclosures, also called controller nodes, for redundant host access. The control enclosures include up to 16 ports that are used for connectivity, with options for 8 Gb Fibre Channel, 16 Gb Fibre Channel, and 10 Gb Ethernet FCoE and iSCSI.

- One FlashSystem V9000 storage enclosure

  FlashSystem V9000 Storage Enclosure is a purpose-built, all-flash storage shelf. It is a 2U 19-inch rack-mount enclosure with 12 slots for flash modules. The storage enclosure provides eight 16 Gb Fibre Channel ports to connect to FlashSystem V9000 control enclosures, either directly or through dedicated internal switches (with the scalable building block). Flash modules within any individual building block are available in 1.2 TB, 2.9 TB, or 5.7 TB capacity. Capacities cannot be intermixed. FlashSystem V9000 is available with either one-year or three-year warranties, so clients can select the warranty period that best addresses their business and financial needs.

Easy to deploy and manage, FlashSystem V9000 is designed to accelerate the applications that drive business. Powered by IBM FlashCore Technology, IBM FlashSystem V9000 provides three dimensions of value. Figure 5-5 shows scalable performance, enduring economics, and agile integration.

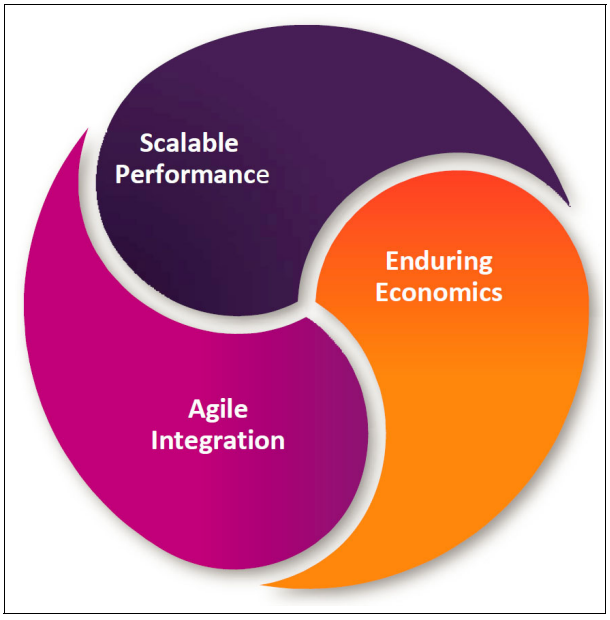

Figure 5-5 IBM FlashSystem V9000 dimensions of value

For more information about FlashCore technology, see 5.3.1, “FlashCore technology” on page 129 and the following webpage:

http://www.ibm.com/systems/storage/flash/900/technology.html

Scalability and performance

FlashSystem V9000 has the following scalability and performance features:

- FlashSystem V9000 eliminates I/O bottlenecks while generating higher levels of application efficiency (improved performance)
- Up to 57 TB usable and 285 TB effective capacity in only 6U
- Up to 456 TB usable and 2.28 PB effective capacity in only 36U
- Up to 2.5 million IOPS and 19.2 GBps bandwidth
- IBM MicroLatency

Reliability, availability, and serviceability

FlashSystem V9000 delivers the following enterprise-class reliability features:

- Concurrent code load enables customer applications to remain online during firmware upgrades to all components, including the flash modules.
- Redundant hot-swappable components: FlashSystem V9000 Storage Enclosure has two clustered, hot-swappable canisters that each contain two hot-swappable fan modules, two management controllers, two management Ethernet ports, and a USB port. The batteries, fans, and power supplies are all redundant and hot-swappable. If a flash module failure occurs, critical customer applications can remain online while the defective module is replaced. IBM Variable Stripe RAID, which is a patented IBM technology that provides an intra-module RAID stripe within each flash module. Variable Stripe RAID technology helps reduce downtime and maintain performance and capacity during partial or full flash chip failures.
- Two-dimensional (2D) flash RAID that consists of IBM Variable Stripe RAID and system-wide RAID 5.
- FlashSystem V9000 control enclosures are an active-active pair and support concurrent code load.

Flash for less than the cost of disk

By offering extreme performance and MicroLatency, FlashSystem V9000 can take the place of multiple racks of HDDs in a small footprint. FlashSystem V9000 can decrease your total cost of ownership (TCO) by lowering your power, space, and cooling costs, and reducing your software licensing expenses.

With Real-time Compression, FlashSystem V9000 can increase the effective capacity of your flash memory up to 5x, decreasing the cost for effective capacity up to 80%. Real-time Compression supports active data, unlike other data reduction solutions. The Model AC2 Control Enclosure offers several features for Real-time Compression workloads. These features include a second Xeon E5 v2 Series eight-core processor with 32 GB of memory and up to two optional compression accelerator cards for hardware-assisted compression.

Comprestimator is the key sizing tool for FlashSystem V9000 to estimate how much capacity savings the customer can expect. Comprestimator can recognize the patterns of the actual customer data, and estimate the compressibility of data per volume. The integration of Comprestimator in FlashSystem V9000 software eases the process of estimating capacity savings by having this sizing tool natively in FlashSystem V9000. This configuration avoids the need to install Comprestimator, and allows estimates of Real-time Compression effectiveness from a central console.

Software-defined services

FlashSystem V9000 merges IBM software-defined storage with the scalable performance of IBM FlashSystem technology to help you more easily manage your entire storage environment while preserving your investments in storage. Software-defined storage services enable you to use the following features across all of your storage:

- Thin provisioning: Enables dynamic growth so you can purchase only the storage you need when you need it.
- Easy Tier flash storage management: Optimizes performance at lower overall cost.
- High availability configurations: Enables near-continuous data availability.
- Copy Services: Enables space efficient backups.
- Disaster recovery techniques: Allows you to practice disaster plans and validate business continuity plans.
- Simple GUI: Allows storage to be quickly deployed and efficiently managed.
- HyperSwap capability: Allows each volume to be presented by two I/O groups. The configuration can tolerate combinations of node and site failures.

Note: IBM Storage Mobile Dashboard, version 1.5.4, supports the IBM FlashSystem V9000 GUI. You can download the dashboard at no cost from iTunes:

https://itunes.apple.com/us/app/ibm-storage-mobile-dashboard/id677826483?mt=8

Deep application integration

IBM FlashSystem V9000 Software V7.6 includes the following features, which enable tight integration with VMware:

- vCenter plug-in: Enables monitoring and self-service provisioning of the system from within VMware vCenter.
- vStorage API for Array Integration (VAAI) support: Supports hardware accelerated virtual machine (VM) copy/migration, hardware accelerated VM initiation, and accelerates VMware virtual machine file system (VMFS).
- SRM for Site Recovery Manager: Supports automated storage and host failover, failover testing, and failback.
- VVOLs integration for better usability: The migration of space efficient volumes between storage containers will maintain the space efficiency of volumes. Cloning a virtual machine will achieve a full independent set of virtual volumes, and resiliency has been improved for VMs in case volumes start running out of space.

VVOL

Before the availability of VVOLs, a VM in a VMware environment would be presented a disk in the form of a file called a virtual machine disk (VMDK). This file represented a physical disk to the VM. The disk could then be accessed by the operating system installed on the VM in the same way a physical volume on a regular server was. The VMDK file was then placed onto a file system called VMFS hosted by a standard volume (LUN), for example implemented on external storage system such as FlashSystem V9000. With the availability of the vSphere Virtual Volume technology, each VM disk can now be mapped to an external storage volume (for example, a FlashSystem V9000 volume).

With VVOL, FlashSystem V9000 becomes aware of individual VMDK files, and data operations such as snapshot and replication can be performed directly by FlashSystem V9000, at the VMDK level rather than the entire VMFS datastore.

Note: The integration of VVOL with FlashSystem V9000 is based on the VMware Storage APIs for Storage Awareness (VASA). The IBM support for VASA is delivered as part of IBM Spectrum Control. VASA version 2 is required to use VVOL capability.

IBM Spectrum Control Base Edition

FlashSystem V9000 currently supports Integration of VASA and VAAI by using IBM Spectrum Control Base Edition 2.1. This is a centralized server system that consolidates a range of IBM storage provisioning, virtualization, cloud, automation, and monitoring solutions through a unified server platform. It provides insight and awareness to VMware and vSphere about the configurations, capabilities and storage health, and events of a storage system. With this capability, VMware administrators can independently and centrally manage their storage resources on IBM storage systems.

IP quorum base support

For lower implementation and operation costs for a high availability solution, IP quorum base support enables the use of lower-cost IP network attached hosts as a quorum disk.

Scale up and scale out

FlashSystem V9000 has a scalable architecture that allows flash capacity to be added (scaled up) to support multiple applications. The virtualized system can also be expanded (scaled out) to support higher IOPS and bandwidth, or the solution can be simultaneously scaled up and out to improve capacity, IOPS, and bandwidth while maintaining MicroLatency. As a result, your organization can gain a competitive advantage through a more flexible, responsive, and efficient storage environment. FlashSystem V9000 has the following scalability features per building block:

- Slots for up to 12 hot-swappable flash memory modules (1.2 TB, 2.9 TB, or 5.7 TB modules)
- Configurable 2.4 - 57 TB of capacity for increased flexibility per storage enclosure
- FlashSystem V9000 has the following flexible scalability configuration options:
  - Base configuration
  - Scale up: Add capacity
  - Scale out: Add controllers and capacity

Figure 5-6 illustrates the increments in the scalable capacity of FlashSystem V9000. It also shows that additional storage enclosures can be added to a single building block, and also two, three, or four building blocks.

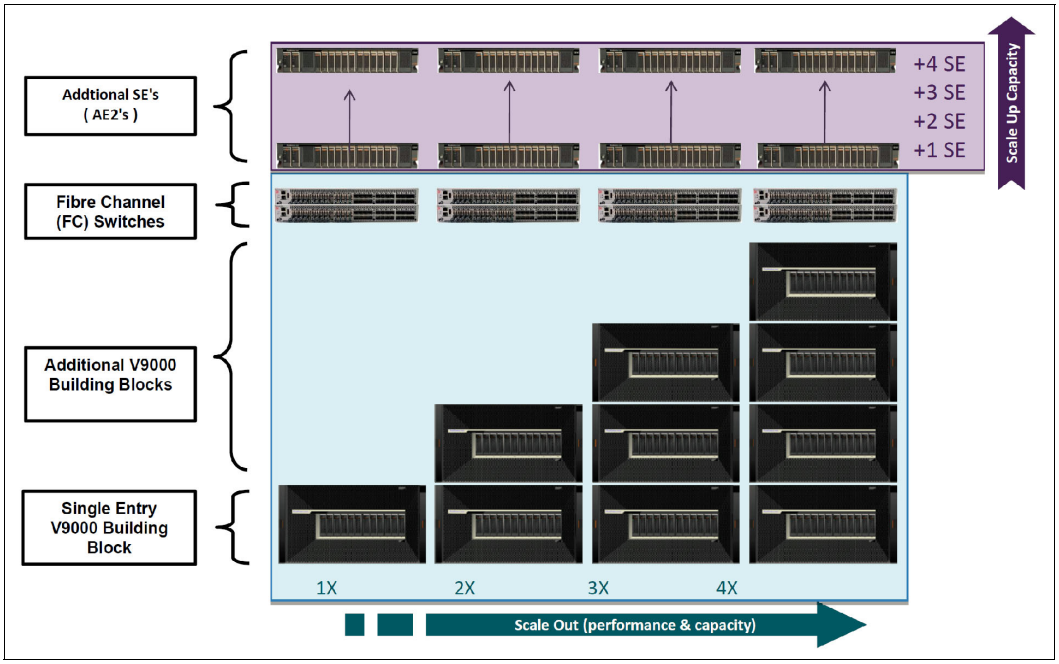

Figure 5-6 FlashSystem V9000 scalable capacity

Manageability and security

FlashSystem V9000 offers the following manageability and security features:

- Advanced security for data at rest with hardware-accelerated AES-XTS 256 encryption.
- A GUI to manage the FlashSystem V9000 Control Enclosure and the FlashSystem V9000 Storage Enclosure. The GUI is available in any supported browser. Also included is the FlashSystem V9000 CLI, which is a collection of commands that you can use to manage the FlashSystem V9000.
- Email alerts.
- Syslog redirect to send system log messages to another host.

5.4.2 Software and licensing

FlashSystem V9000 uses IBM Storwize software-defined storage features. FlashSystem V9000 data services are provided through FlashSystem V9000 Software. FlashSystem V9000 has both base and optional software licenses.

The following functions are provided with the FlashSystem V9000 base software license:

- Thin provisioning
- Data migration
- Simple GUI
- EasyTier
- Automatic restriping of data across storage pools

The following functions are included with the FlashSystem V9000 base software license only for internal storage:

- FlashCopy
- Real-time Compression
- Remote Mirroring

The following optional licensed features are offered with the FlashSystem V9000 Software for external storage:

- External storage virtualization
- Real-time Compression
- FlashCopy
- Remote Mirroring

For more information about software and licensing, see *IBM FlashSystem V9000 Product Guide*, REDP-5317.

5.4.3 IBM FlashSystem V9000 in a VersaStack environment

VersaStack, an IBM and Cisco integrated infrastructure solution, combines computing, networking, and storage into a single integrated system. It combines the Cisco Unified Computing System (Cisco UCS) Integrated Infrastructure with IBM Spectrum Virtualize, which includes IBM FlashSystem V9000 and IBM Storwize storage offerings, for quick deployment and rapid time to value for the implementation of modern infrastructures. With comprehensive reference architectures that include Cisco Validated Designs (CVDs), IBM Redbooks publications, sizing guidelines, and single-call support, the solution sets a benchmark to accelerate data center infrastructure deployment and to help manage information and resources efficiently amid business change.

Figure 5-7 illustrates the VersaStack solution with IBM FlashSystem V9000 at a glance.

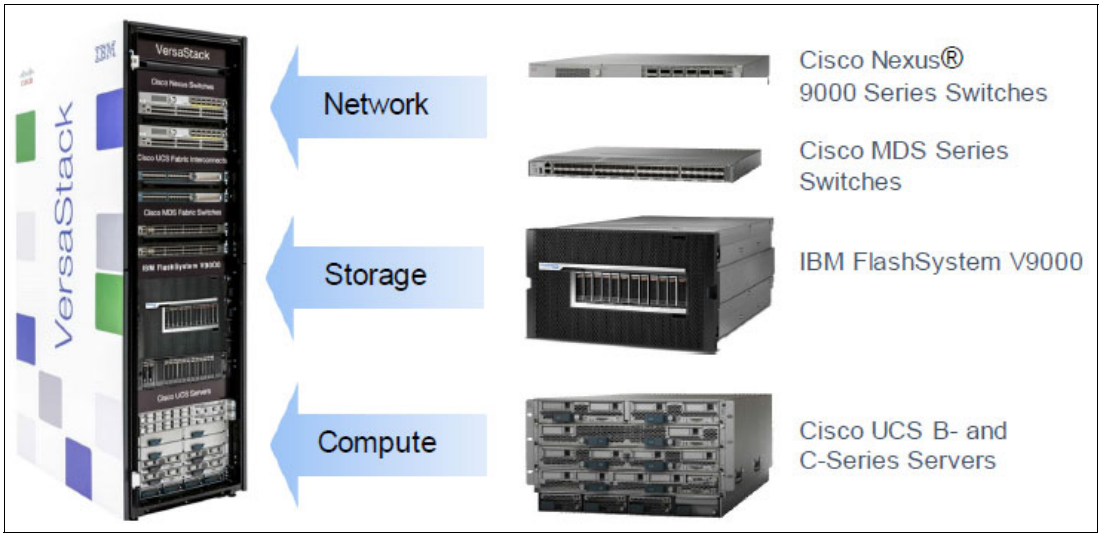

Figure 5-7 VersaStack at a glance

For more information about IBM FlashSystem V9000 in a VersaStack environment, see *IBM FlashSystem V9000 in a VersaStack Environment*, REDP-5264.

5.4.4 More information

For more information and detailed specifications, see the following website:

http://www.ibm.com/systems/storage/flash/v9000/overview.html

Also, see the following IBM Redbooks publications:

- *IBM FlashSystem V9000 Product Guide*, REDP-5317
- *Introducing and Implementing IBM FlashSystem V9000*, SG24-8273
- *IBM FlashSystem V9000 in a VersaStack Environment*, REDP-5264

6

IBM DS8800 Storage System

As an easy-to-deploy, nondisruptive addition to DS8880, this flash enclosure puts optimized flash storage capacity within easy reach of both new and legacy deployments of DS8880 storage infrastructures.

Designed to deliver greater than “six nines” (99.9999 percent) availability, DS8880 is an enterprise-class storage system that is ideally suited to addressing heavy workload requirements. IBM brings together the latest in high-performance flash storage with the proven reliability and extreme high availability of IBM enterprise-class storage solutions. Organizations can access business-transforming insights instantly, keep e-commerce sites up and running, and do it all confidently with a proven, highly resilient system.

The family of IBM Enterprise Disk Systems offers a broad range of scalable solutions to address various enterprise storage needs. By using leading IBM technology, the enterprise disk systems provide choice in function, performance, and resiliency.

Industry-leading IBM flash solutions reach across the IBM storage portfolio, and now include IBM High-Performance Flash Enclosure, designed for DS8880 storage systems.

This chapter includes the following sections:

- DS8880 overview
- Product highlights
- Models
- Related information

6.1 DS8880 overview

This chapter describes the features and major characteristics of the IBM DS8880 storage family products.

The new IBM DS8880 family introduces business-critical, hybrid data systems that span a wide range of price points. It includes two high-performance models:

- DS8884 Model 980 with its associated DS8884 Expansion Unit Model 98B
- DS8886 Model 981 with its associated DS8886 Expansion Unit Model 98E

Each model represents the most recent in this series of high-performance, high-capacity, flexible, and resilient storage systems. These systems are aimed at addressing the needs of the most demanding clients and enabling unique and powerful storage integration with IBM z Systems.

6.2 Product highlights

This section provides an overview of the features and functions that are offered by the IBM DS8880 models with firmware release 8.0.

Like its predecessors, the IBM DS8880 architecture is server-based. Powerful IBM POWER8® processor-based servers manage the cache to streamline disk input/outputs (I/Os), maximizing performance and throughput. These capabilities are further enhanced with the availability of High-Performance Flash Enclosures (HPFEs).

The DS8880 family includes the following models:

- IBM DS8884: DS8884 (model 980, with model 98B expansion frames) is the successor of the DS8870 Business Class model. It enables organizations to overcome storage challenges with advanced, easy-to-use functionality for running critical workloads on mainframes, POWER servers or distributed systems, either as a dedicated platform for consolidated systems or for multiple platforms, delivered within an affordable, flexible, and space-saving package.
- IBM DS8886: The DS8886 (model 981, with 98E expansion frames) is the successor of the DS8870 Enterprise Class model. It helps accelerate mission-critical applications with up to two times better performance, backed by 24 x 7 availability and superior functionality for multi-site replication, and deep z Systems, Power, or distributed systems integration. All of these features are provided in a dense yet expandable model.

The IBM DS8886 and DS8884 excel in supporting the IBM zEnterprise® server and IBM Power server environments. Both models can be ordered with a one-year, two-year, three-year, or four-year support period.

Figure 6-1 shows the frontal view of the DS8884 and DS8886 base frames. The racks can be slightly different in height: 40U for the DS8884, and up to 46U for the DS8886.

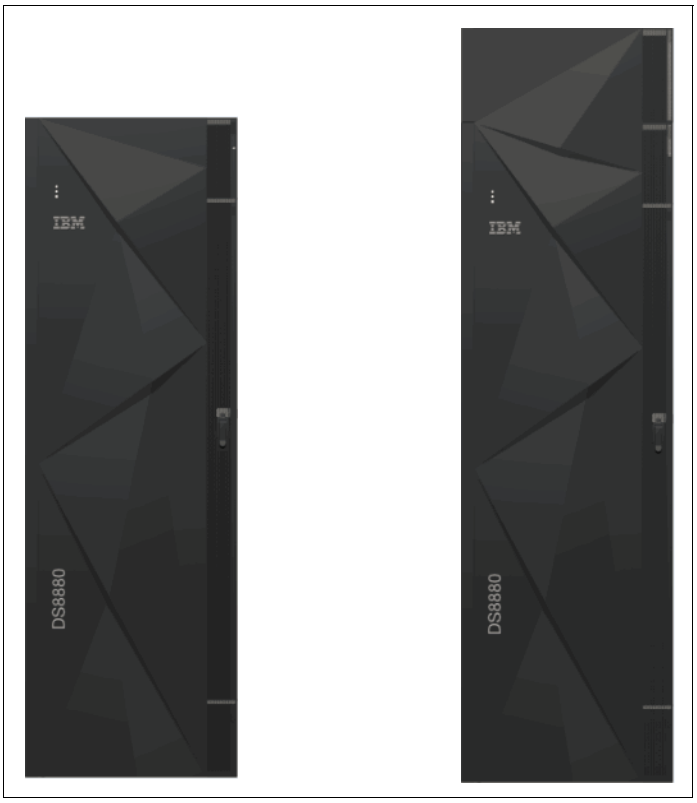

Figure 6-1 IBM DS8884 and DS8886 front views, base frame

6.2.1 Features

The IBM DS8880 family has these key features, among others:

- The processor complexes of the DS8880 family are based on IBM Power System servers like the S824 and S822 POWER8 models that are also independently available as stand-alone servers running AIX and Linux operating systems.
- The IBM DS8880 models feature an HPFE that provides outstanding flash performance. The HPFE is directly attached to the Peripheral Component Interconnect Express fabric, enabling increased bandwidth and transaction processing capability. This 1U storage enclosure contains a pair of powerful flash RAID adapters and up to thirty 400 gigabyte (GB) encryption-capable flash cards in each enclosure, for a maximum of 240 flash cards in the IBM DS8886 model.
- The IBM Storage Driver for OpenStack is a software component that integrates with the OpenStack cloud environment and enables usage of storage resources provided by supported IBM storage systems. After the driver is installed on the OpenStack Cinder (OpenStack Block Storage) nodes, storage volumes can be allocated by the Cinder nodes to the Nova-compute nodes. Virtual machines on the Nova-compute nodes can then use these storage resources.
- IBM Easy Tier, a well-proven feature of the DS8000 series, is available at no charge as part of the Base license package. It dynamically optimizes performance for multi-tiered systems. It can also rebalance data within a single tier to help maintain optimal performance. Now in its eighth generation, Easy Tier provides full support of the HPFE

and associated flash cards, while maintaining extra features like Easy Tier Heat Map Transfer or Easy Tier Application.

- With the frame width reduced from 33 inches to 19 inches, the DS8880 comes with a 42% narrower footprint than the previous DS8870 model.

6.2.2 Scalability and performance

The IBM DS8880 models feature POWER8 server technology to help support higher performance:

- The DS8880 family is available with different processor options, which range from dual 24-cores in the DS8886 to a single 6-core systems in the DS8884. The simultaneous multithreading (SMT) mode enables the POWER8 processor to maximize the throughput of the processor core and offers more threads than competing processor technologies.
- The first generation of DS8870 (IBM POWER7®) came with 1.2 billion processors per core. The second DS8870 generation (IBM POWER7+™) came with 2.1 billion transistors. The DS8880 models come with 4.2 billion transistors per core.
- Cache configurations are available that range from 64 GB up to 2 terabytes (TB) of system cache. The storage server architecture of the IBM DS8880 with its POWER8 processors makes it possible to manage large caches with small cache segments of 4 KB (kilobytes), and hence large segment tables, without the need to partition the cache.
- Peripheral Component Interconnect Express (PCIe Generation 3) improves sequential read/write throughput and input/output operations per second (IOPS). The I/O enclosures are directly connected to the DS8880 storage servers with point-to-point PCI Express Gen3 cables.
- The HPFE is populated with 1.8-inch encryption capable 400 GB flash cards on a dedicated architecture, which outperform SSDs (solid-state drives). Each HPFE can contain up to thirty 400 GB flash cards.

  Using the HPFE, the internal design of the IBM DS8880 family features a true flash architecture. By bypassing the classical device adapters (DAs), the PCIe-based architecture of the HPFE enables it to meet extreme IOPS and response time requirements, in ultra-dense packaging.

- The IBM DS8880 models have an improved 19-inch high-density frame design. As shown in Figure 6-2, it can support 1536 drives, plus 240 flash cards, in a small, high-density footprint (base frame and three expansion frames). This smaller footprint helps to preserve valuable raised floor space in data center environments.

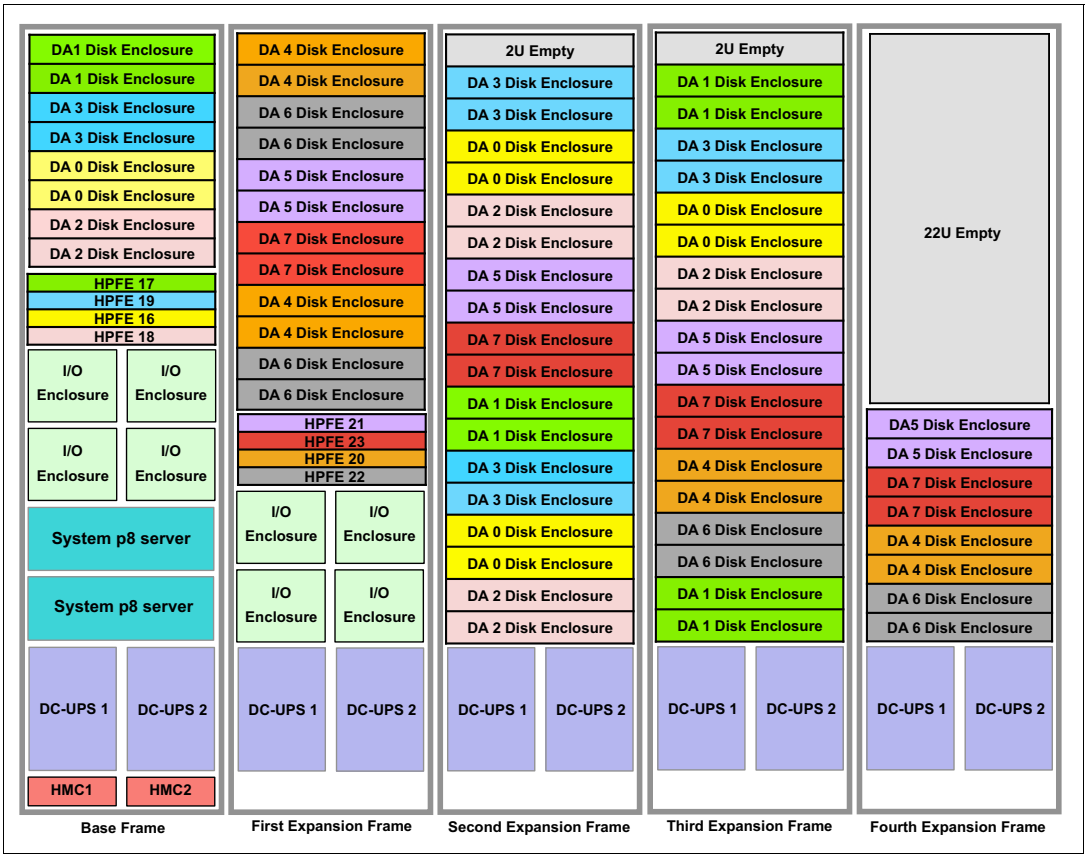

Figure 6-2 Front view of fully configured DS8886

- The IBM DS8886 scales to more than 3 petabytes (PB) of raw drive capacity, and supports multiple storage tiers. The IBM DS8880 family supports a broad range of drives that use different technologies. Nine types of drives are managed in up to three different tiers (flash cards and flash drives as tier 0, fibre-attached serial-attached SCSI (SAS) drives with 15,000 and 10,000 revolutions per minute (RPM) as tier 1, and near-line SAS drives as tier 2).
- The DS8886 model provides up to 32 Fibre Channel host adapters, supplying a total of up to 128 ports (16 Gbps) in a single-machine footprint.
- The smaller IBM DS8884 model option supports up to 768 small form-factor (SFF) drives or flash drives, plus up to 120 high-performance flash cards.
- Write data is always protected by maintaining a copy of modified data in nonvolatile storage (NVS) until the data is destaged to hardened storage.
- The Adaptive Multi-stream Prefetching (AMP) caching algorithm can dramatically improve sequential performance, reducing times for backup, processing for business intelligence, and processing for streaming media. Sequential Adaptive Replacement Cache is a caching algorithm that allows you to run different workloads, such as sequential and random workloads, without negatively affecting each other.
- IBM Easy Tier offers capabilities such as manual volume capacity rebalance, auto performance rebalancing in both homogeneous and hybrid pools, hot-spot management,

rank depopulation, manual volume migration, and thin provisioning support. It is available at no extra charge. Easy Tier also includes the following capabilities:

  - Easy Tier Application enables clients to assign distinct application volumes to a particular tier in the Easy Tier pool, disregarding the Easy Tier advanced data migration function.
  - Easy Tier Heat Map Transfer is able to provide whatever the data placement algorithm is on the Metro Mirror/Global Copy/Global Mirror (MM/GC/GM) primary site. It then reapplies it on the MM/GC/GM secondary site through the Easy Tier Heat Map Transfer utility when failover occurs.
  - Easy Tier includes additional reporting improvements, such as workload skew curve, workload categorization, and a data movement daily report. In addition, it has intra-tiering and micro-tiering support for storage tier with more than one drive technology (mixed high performance flash cards and flash drives, or mixed Enterprise 15,000 and 10,000 RPM drives).

- Storage Pool Striping (rotate extents) provides a mechanism to distribute a volume's or logical unit number's (LUNs) data across many RAID arrays.
- Storage Tier Advisor Tool (STAT) is a tool that is used with the Easy Tier feature to help clients understand their current storage system workloads.
- I/O Priority Manager is a feature that provides application-level quality of service (QoS) for workloads that share a storage pool. This feature provides a way to manage QoS for I/O operations associated with critical workloads, and gives them priority over other I/O operations associated with non-critical workloads. For IBM z/OS®, the I/O Priority Manager allows increased interaction with the host side.
- 8 Gbps FC/IBM FICON® host adapters in the IBM DS8880 offer enhanced connectivity, with 4-port and 8-port host adapters in the I/O enclosures that are directly connected to the processor complexes. The 16 Gbps adapters are available as 4-port only.
- The host adapters support FICON attachment to IBM z Systems servers. You can configure each port to operate as a Fibre Channel Protocol (FCP) port, or FICON port.
- High-Performance FICON for IBM z Systems (zHPF) is z/OS's new I/O architecture, and already comes with several generations of enhancements. zHPF is included in the DS8880 z feature code package. Enhancements to zHPF include Extended Distance capability, zHPF List Pre-fetch support for IBM DB2 and utility operations, and zHPF support for sequential access methods. All DB2 I/O is now zHPF-capable.
- The IBM DS8880 models provide large volume support for LUN sizes up to 16 TiB. This configuration simplifies storage management tasks. In a z/OS environment, extended address volumes (EAVs) with sizes up to 1 TB are supported.
- Traditional performance features for IBM z Systems environments are parallel access volumes (PAVs), HyperPAV, Multiple Allegiance, I/O priority queuing, I/O Priority Manager, and zHPF.
- SAN Fabric I/O Priority Management: Starting with newer z Systems like IBM z13™, QoS functions have also been extended into the SAN fabric. Like for the I/O Priority Manager with z/OS, there is cooperation between the z Systems Workload Manager (WLM) and that SAN QoS function. The WLM is capable to send the priority information of each I/O to the SAN fabric and to the DS8000, completing the management of the entire end-to-end flow of an I/O.

- FICON Dynamic Routing (FIDR) is another performance-relevant function that is available with newer FICON cards like the z13 offers. When considering many paths in a SAN, for instance with a larger number of inter-switch links (ISLs), the traditional static routing often led to imbalanced ISLs with not all available bandwidth used. FIDR leads to an optimally balanced SAN, what translates into a bigger effective regain of SAN ISL bandwidth.
- Additionally, the system offers end-to-end I/O priorities, cooperative caching, long busy wait host tolerance, or automatic port queues.

6.2.3 Availability and serviceability

The DS8880 family is designed and implemented with component redundancy to avoid potential single points of failure. The DS8880 models offer high availability and multiplatform support, including z Systems.

- Up to 128 host adapter ports can be configured to operate as an Fibre Channel (FC) port or a FICON ports. The host bus adapter (HBA) cards can operate (auto-negotiate) up to two speed factors down: For the 16 Gbps HBAs, down to 4 Gbps, and for the 8 Gbps HBAs, down to 2 Gbps if needed. There is a full-duplex data transfer, over either longwave or shortwave fibre links.
- A nondisruptive upgrade path with each DS8880 model allows configuration upgrades for processors, memory, and storage capacity enhancement to be performed concurrently. The IBM DS8886 supports a scalable upgrade path from the smallest to the largest processor configuration.

  An upgrade from DS8884 to DS8886 is not possible.
- The IBM DS8880 family features Smart Rebuild, a function that is designed to help reduce the possibility of secondary failures and data loss in RAID arrays. The IBM DS8000 series supports RAID 5, RAID 6, and RAID 10.
- The Dynamic Volume Expansion simplifies management by enabling easier, online volume expansion (for Open Systems and z Systems) to support application data growth, and to support data center migration and consolidation to larger volumes to ease addressing constraints.
- For data protection and availability, the IBM DS8000 series supports advanced disaster recovery (DR) and business continuity solutions, such as Copy Services functions, available for thin-provisioned volumes as well.
- Resiliency highlights include the following capabilities:
  - FlashCopy handling of volume reservations provides the detection of SCSI reserves that exist for devices in a Metro Mirror or Global Mirror environments. It resets the reserve when performing a FlashCopy after it has been identified that it is not a valid reserve for a running server.
  - z/OS Soft Fence prevents any system from accessing data from the former remote mirror primary site when an unplanned HyperSwap occurs.
- Active Volume Protection is a feature that prevents the deletion of volumes still in use.
- Support for T10 Data Integrity Field is standard. The Data Integrity Field standard of SCSI T10 enables end-to-end data protection, from the application or host HBA down to the storage drives.
- Lightweight Directory Access Protocol (LDAP) authentication support, which allows single sign-on (SSO) functionality, can simplify user management by allowing the IBM DS8000 to rely on a centralized LDAP directory rather than a local user repository. For more information, see *LDAP Authentication for IBM DS8000 Storage*, REDP-4505.

- The IBM DS8880 models are certified as meeting the requirements of the IPv6 Ready Logo program, which indicates its implementation of IPv6 mandatory core protocols and the ability to interoperate with other IPv6 implementations. They can be configured in native IPv6 environments.

6.2.4 Security and encryption

Combined with the world-class business resiliency and encryption features, the IBM DS8880 family provides a unique combination of HA, performance, and security.

To counteract the growing threat of security breaches, the IBM DS8880 has self-encrypting drives as a standard feature, and has security capabilities, such as broad-based access-control and tamper-proof audit logging, in addition to other security features:

- Self-encrypted drives are a standard feature. You just need to activate encryption on drives with IBM Security Key Lifecycle Manager (SKLM). For more information, see *IBM DS8870 Disk Encryption*, REDP-4500.
- Full Disk Encryption (FDE) can protect business-sensitive data by providing drive-based hardware encryption that is combined with a sophisticated key management software. FDE is available for all drive types, including flash cards and flash drives (solid-state drives (SSDs)). Because encryption is done by the drive, it is transparent to host systems, and can be used in any environment, including z/OS.
- Security improvements in the IBM DS8880 family enable customers to become compliant with the Special Publication (SP) number 800-131a, which is a National Institute of Standards and Technology (NIST) directive that provides guidance for protecting sensitive data by using cryptographic algorithms that have key strengths of 112 bits. For more information, see *IBM DS8870 and NIST SP 800-131a Compliance*, REDP-5069.
- The following specific features of encryption key management help address Payment Card Industry Data Security Standard (PCI DSS) requirements:
  - Encryption deadlock recovery key option. When enabled, this option enables you to restore access to an IBM DS8880 when the encryption key for the storage is unavailable because of an encryption deadlock scenario.
  - Dual-platform key server support. This support is important if key servers on z/OS share keys with key servers on open systems. The DS8880 family requires one isolated key server in encryption configurations. Dual-platform key server support allows two server platforms to host the key manager, with either platform operating in either clear key or secure key mode.
  - Recovery Key Enabling/Disabling and Rekey data key option for the FDE feature. Both of these enhancements can help clients satisfy PCI security standards.

6.2.5 Advanced functions

The IBM DS8880 family has the following advanced functions:

- The IBM DS8880 provides DS Storage Manager (the graphical user interface (GUI)) and DS command-line interface (CLI) management interfaces to configure the system or query status information.
- Thin Provisioning enables the creation of over-provisioned devices for more efficient use of the storage capacity for Open Systems.
- Quick Initialization provides fast volume initialization for Open Systems LUNs and count key data (CKD) volumes, and therefore allows the creation of devices, making them available as soon as the command completes.

- The following list includes advanced Copy Services features in the IBM DS8000:
  - IBM FlashCopy is a feature that allows the creation of volume copies (and data set copies for z/OS) nearly instantaneously. Different options are available to create full copies, incremental copies, or copy-on-write copies.
  - The DS8880 models and code provide the same remote-mirroring options as previous models of the DS8000 family. Synchronous remote mirroring (Metro Mirror) is supported up to 300 km. Asynchronous copy (Global Mirror) is supported for unlimited distances.
  - Metro Mirror, Global Copy, Global Mirror, Metro/Global Mirror, z/OS Global Mirror, and z/OS Metro/Global Mirror business continuity solutions are designed to provide the advanced functionality for almost any recovery point or recovery time objective.
  - The Copy Services can be managed and automated with the IBM Copy Services Manager (CSM). For z/OS environments, IBM Geographically Dispersed Parallel Sysplex™ (IBM GDPS®) provides an automated DR solution.
  - With AIX operating systems, the DS8880 family supports Open IBM HyperSwap replication. The Open HyperSwap is a special Metro Mirror replication method designed to automatically fail over I/O from the primary logical devices to the secondary logical devices in the event of a primary storage system failure.
  - In co-operation with z/OS's Data Mover, another option is available for z/OS: Global Mirror for z/OS. Another important feature for z/OS Global Mirror (2-site) and z/OS Metro/Global Mirror (3-site) is Extended Distance FICON.
  - Remote-Pair FlashCopy enables you to establish a FlashCopy relationship where the target is a remote mirror Metro Mirror primary volume, keeping the pair in the full-duplex state.
  - The Easy Tier Heat Map Transfer function is also integrated with IBM Copy Services Manager or with newer GDPS versions, and all of the functions are available through the IBM Copy Services Manager.
  - The Resource Groups feature is a policy-based resource scope limiting function that enables the secure use of Copy Services functions by multiple users on a DS8000 series storage system.
- The IBM DS8880 provides support for VMware vStorage application programming interfaces (APIs) for Array Integration (VAAI). VAAI enables certain storage tasks to be offloaded from the server hardware to the storage array. Support is included for the Atomic Test and Set (ATS) primitive, the Cloning Blocks primitive, and the Zeroing Blocks primitive.
- The DS8880 family supports VASA 2.0, and the RESTful API.
- The IBM DS8880 models also support the IBM Storage Management Console for VMware vCenter. The IBM Storage Management Console for VMware vCenter is a software plug-in that integrates into the VMware vCenter server platform and enables VMware administrators to independently and centrally manage their storage resources on IBM storage systems.

6.3 Models

The DS8880 delivers extensive scalability as compared to previous DS8870 systems:

- IBM DS8886 Model 981 configuration

  With 8-, 16- or 24-core processor complexes, the DS8886 Model 981 supports up to 192 2.5-inch disk drives and 120 flash cards for a maximum capacity of up to 355.2 TB. It also supports up to 2 TB of processor memory with up to 64 Fibre Channel/FICON ports. With an optional expansion unit (DS8886 Model 98E), it scales as follows:

  - With one DS8886 Model 98E Expansion Unit, the DS8886 Model 981 (16- or 24-core) supports up to 128 Fibre Channel/FICON ports, 480 disk drives and 240 flash cards. This configuration provides up to 864 TB of physical capacity with 1.6 TB SSD drives, or up to 240 drives for up to 1.06 PB of physical capacity with 4 TB nearline SAS disk drives.
  - With two DS8886 Model 98E Expansion Units, the DS8886 Model 981 (16- or 24-core) supports up 912 disk drives. This configuration provides up to 1.46 PB of physical capacity with 1.6 TB SSD drives, or up to 456 drives for up to 1.82 PB of physical capacity with 4 TB nearline SAS disk drives.
  - With a third DS8886 Model 98E Expansion Unit, the DS8886 Model 981 (16- or 24-core) supports up 1 344 disk drives. This configuration provides up to 2.15 PB of physical capacity with 1.6 TB SSD drives, or up to 672 drives for up to 2.69 PB of physical capacity with 4 TB nearline SAS disk drives.
  - With a forth DS8886 Model 98E Expansion Unit, the DS8886 Model 981 (16- or 24-core) supports up 1 536 disk drives. This configuration provides up to 2.46 PB of physical capacity with 1.6 TB SSD drives, or up to 768 drives for up to 3.1 PB of physical capacity with 4 TB nearline SAS disk drives.

- IBM DS8884 Model 980 configuration

  With a 6-core processor complex, the DS8884 Model 980 supports up to 240 2.5-inch disk drives and up to 60 flash cards for a maximum capacity of up to 408 TB. It also supports up to 256 GB of processor memory with up to 32 Fibre Channel/FICON ports. With an optional expansion unit (DS8884 Model 98B), it scales as follows:

  - With one DS8884 Model 98B Expansion Unit, the DS8884 Model 980 (128 or 256 GB memory) supports up to 64 Fibre Channel/FICON ports, 528 disk drives and 120 flash cards. This configuration provides up to 892.8 TB of physical capacity with 1.6 TB SSD drives, or up to 264 drives for up to 1.1 PB of physical capacity with 4 TB nearline SAS disk drives.
  - With two DS8884 Model 98B Expansion Units, the DS8884 Model 980 (127 or 256 GB memory) supports up to 768 disk drives, providing up to 1.23 PB of physical capacity with 1.6 TB SSD drives, or up to 384 drives for up to 1.54 PB of physical capacity with 4 TB nearline SAS disk drives.

To meet performance and configuration needs, the new DS8880 family features a complex of dual POWER8 servers. Each server is equipped with either single processors (6-, and 8-core), or dual processors (16-, and 24-core).

Table 6-1 provides more information about drive and adapter counts with the two available configuration options.

Table 6-1 DS8880 models and characteristics

Model	Cores per Processor Complex	Physical capacity	Disk drives	Processor memory	Host adapters	9xE attach	Flash enclosures
DS8886 configuration:							
981	8-core	355 TB	192	256 GB	16	0	4
981	16-core	2.55 PB	1536	512 GB	32	0 to 4	8
981	24-core	2.55 PB	1536	2048 GB	32	0 to 4	8
DS8884 configuration:							
980	6-core	408 TB	240	64 GB	8	0	2
980	6-core	1.28 TB	768	256 GB	16	0 to 2	4
Model	**Cores per Processor Complex**	**Physical capacity**	**Disk drives**	**Processor memory**	**Host adapters**	**9xE attach**	**Flash enclosures**
DS8886 configuration first expansion:							
98E	N/A	508.8 TB	288	N/A	16	981	4
DS8886 configuration second/third expansions:							
98E	N/A	681.2 TB	432	N/A	0	981	0
DS8886 configuration fourth expansion:							
98E	N/A	307.2	192	N/A	0	981	0
DS8884 first expansion configuration first expansion:							
98B	N/A	460.8	288	N/A	8	980	2
DS8884 first expansion configuration second expansion:							
98B	N/A	384 TB	240	N/A	0	980	0

Figure 6-3 on page 152 shows the front view of a DS8886, with one expansion frame attached.

- The base frame shows the two POWER8 server-based processor complexes. Processors and system memory can be upgraded concurrently.
- Four HPFEs are shown in the base frame, and four in the first expansion frame.
- Four I/O bay enclosures are shown in the base frame, and four in the first expansion frame. The adapters that are contained in the I/O enclosures can be either DAs or host adapters, or you can have the connection to the HPFEs.
- Primary (and optional secondary) Hardware Management Console (HMCs) are at the bottom of the base frame.

The power subsystem is based on direct current uninterruptible power supply (DC-UPS) with integrated battery sets. Each frame contains two DC-UPS power supplies in a fully redundant configuration.

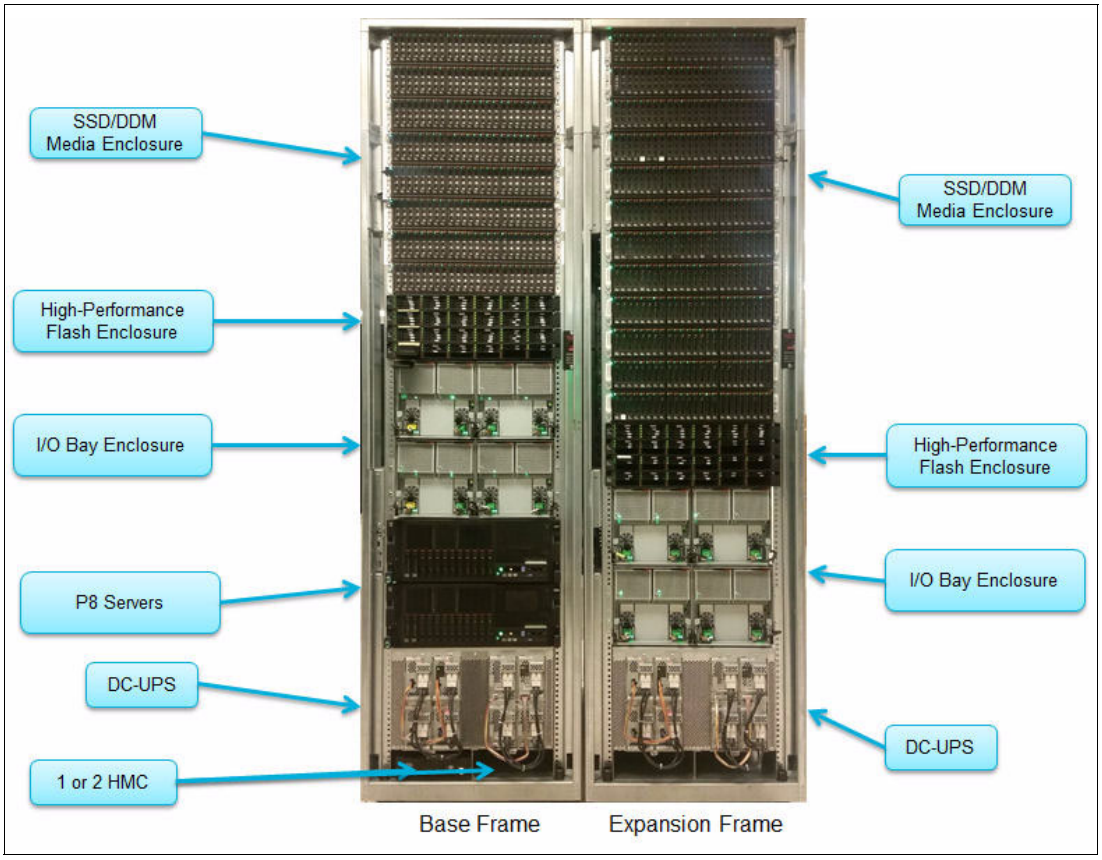

Figure 6-3 DS8886 front view, with one expansion frame shown

For more information about the expansion frames and specifications, see the *IBM DS8880 Introduction and Planning Guide*, GC27-8525.

6.3.1 Other configuration features

This section highlights that main features, upgrades, and options that can be requested with new DS8880 orders or installed later.

- A secondary HMC can be ordered as a redundant HMC to provide continuous availability to the HMC functions. The use of a second HMC can be useful in storage environments where Encryption capability or Copy Services functions have been activated.
- The extended power line disturbance (ePLD) option is designed to protect your storage unit for 40 seconds, rather than only 4 seconds, from a power-line disturbance. If no ePLD option is ordered, one Battery Set Module (BSM) set per DC-UPS (for both the main and expansion racks) is needed. If the ePLD option is ordered, two BSM sets per DC-UPS (for both the main and expansion racks) are needed.
- An optional overhead cabling feature is available. The overhead cabling option includes a top-exit bracket for fiber cables.

6.3.2 Warranty information and upgrades

DS8880 offers the Enterprise Choice warranty of 1, 2, 3 or 4 years on both the hardware *and* the advanced function software in the following way:

- 4 years on type 2834 models
- 3 years on type 2833 models
- 2 years on type 2832 models
- 1 year on type 2831 models

6.3.3 Scalable upgrades

With the DS8884 model, it is possible to start with a single-frame 6-core configuration with storage enclosures for 48 DDMs, and grow to a full-scale, 768-drive, 120 Flash Card, three-frame configuration.

With the DS8886 model, it is possible to start with a single-frame 8-core configuration, and grow to a 1536-drive, 240 Flash Card, five-frame configuration. All frame and capacity upgrades are concurrent, regardless of configuration type.

6.3.4 Licensed functions

Many of the software optional licensed functions of the DS8880 need a License Key to be activated. In fact, licensed functions are enabled through a licensed function indicator feature plus a licensed function authorization feature number.

The following three licensed functions are available:

- Base Function: The Base Function license is required for each DS8880 storage system.
- Copy Services: Copy Services features help you implement storage solutions to keep your business running 24 hours a day, 7 days a week by providing data duplication, data migration, and disaster recovery functions.
- z-synergy Services: The z-synergy Services include z/OS licensed features that are supported on the storage system.

Licensed functions are activated and enforced within a defined *License Scope*. License scope refers to the type of storage, and therefore the type of servers, that the function can be used with: Fixed Block (FB), Count Key Data (CKD) or both FB and CKD. Orders come in increments of a specific capacity.

The Base Function license is ordered to cover the full physical capacity of the storage system. It includes the following features:

- Operating Environment License (OEL)
- Database protection
- Encryption Authorization
- Easy Tier
- I/O Priority Manager
- Thin Provisioning.

The Copy Services license includes the following features (including Multi-Target PPRC relations):

- Metro Mirror
- Global Mirror
- Metro/Global Mirror

- Point-in-Time Copy / FlashCopy
- z/OS Global Mirror
- z/OS Metro/Global Mirror Incremental Resync (RMZ Resync).

The z-synergy Services license is available for CKD scope and includes the following features:

- Parallel Access Volumes (PAV)
- HyperPAV
- High-Performance FICON for z Systems (zHPF)
- z/OS Distributed Database Backup.

For more information, see the *IBM DS8880 Introduction and Planning Guide*, GC27-8525.

6.3.5 Supported environments

The DS8000 offers connectivity support across a broad range of server environments, including IBM Power Systems, z Systems, servers from Oracle and Hewlett-Packard, non-IBM Intel, and AMD-based servers.

The DS8000 supports over 90 platforms. For the most current list of supported platforms, see the DS8000 System Storage Interoperation Center (SSIC) at this website:

http://www.ibm.com/systems/support/storage/config/ssic/

For information about supported environments, see *IBM System Storage DS8000: Host Attachment and Interoperability*, SG24-8887.

For information about VMware, see *IBM DS8870 and VMware Synergy*, REDP-4915.

This rich support of heterogeneous environments and attachments, along with the flexibility to easily partition the DS8880 storage capacity among the attached environments, can help support storage consolidation requirements and dynamic environments.

Disk Magic

Disk Magic is a Microsoft Windows-based performance modeling tool that is used by IBM and IBM Business Partners to model storage system performance. It supports storage systems from multiple vendors and offers detailed support for IBM storage systems.

Important: IBM Disk Magic for Windows is a product of IntelliMagic, which is licensed to IBM and IBM Business Partners to model disk storage system performance. Contact your IBM Representative or IBM Business Partner to discuss a Disk Magic study.

6.4 Related information

For more information, see the following publications and links:

- *IBM DS8880 Architecture and Implementation (Release 8)*, SG24-8323
- *IBM DS8880 Introduction and Planning Guide*, GC27-8525
- IBM System Storage DS8884 and DS8886 (Machine type 2831) Models 980 and 981 with one-year warranty - Product Announcement (20 October 2015):

  http://www.ibm.com/common/ssi/ShowDoc.wss?docURL=/common/ssi/rep_ca/9/897/ENUS115-139/index.html

- IBM System Storage DS8884 and DS8886 (Machine type 2832) Models 980 and 981 with two-year warranty - Product Announcement (20 October 2015):

  http://www.ibm.com/common/ssi/ShowDoc.wss?docURL=/common/ssi/rep_ca/2/897/ENUS115-142/index.html

- IBM System Storage DS8884 and DS8886 (Machine type 2833) Models 980 and 981 with three-year warranty - Product Announcement (20 October 2015):

  http://www.ibm.com/common/ssi/ShowDoc.wss?docURL=/common/ssi/rep_ca/0/897/ENUS115-140/index.html

- IBM System Storage DS8884 and DS8886 (Machine type 2834) Models 980 and 981 with four-year warranty - Product Announcement (20 October 2015):

  http://www.ibm.com/common/ssi/ShowDoc.wss?docURL=/common/ssi/rep_ca/1/897/ENUS115-141/index.html

- IBM DS8880 (Machine type 283x) high-performance flagship Function Authorizations - Product Announcement (20 October 2015):

  http://www.ibm.com/common/ssi/ShowDoc.wss?docURL=/common/ssi/rep_ca/8/897/ENUS115-138/index.html

- *IBM DS8870 Copy Services for Open Systems*, SG24-6788
- *IBM DS8870 Copy Services for IBM z Systems,* SG24-6787
- *IBM System Storage DS8000 Copy Services Scope Management and Resource Groups*, REDP-4758
- IBM DS8870 Multiple Target Peer-to-Peer Remote Copy, REDP-5151
- *IBM DS8000 Easy Tier*, REDP-4667
- *DS8870 Easy Tier Application*, REDP-5014
- *IBM DS8870 Easy Tier Heat Map Transfer,* REDP-5015
- *DS8000 Thin Provisioning*, REDP-4554
- *DS8000 I/O Priority Manager*, REDP-4760
- *LDAP Authentication for IBM DS8000 Storage*, REDP-4505
- *IBM DS8870 Disk Encryption*, REDP-4500
- *Introduction to IBM Assist On-site Software for Storage*, REDP-4889
- IBM Assist On-Site:

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

- IBM Tivoli Assist On-site Remote Support Utility. User's Guide:

  https://aos.us.ihost.com/AssistOnSiteAdmin/docs/AOS_Utility_User_Guide.pdf

- *IBM System Storage DS8000: Host Attachment and Interoperability*, SG24-8887:
- *DS8000 Performance Monitoring and Tuning,* SG24-8318
- *IBM DS8870 and VMware Synergy,* REDP-4915
- *IBM DS8000 Series Command-Line Interface User's Guide*, SC27-8526
- IBM Knowledge Center / IBM System Storage DS8000 Information Center:

  http://www.ibm.com/support/knowledgecenter/ST5GLJ/product_welcome/ds8000_kcwelcome.html

- DS8880 support:

  https://www.ibm.com/support/entry/portal/product/system_storage/disk_systems/enterprise_storage_servers/ds8880
- IBM System Storage Interoperation Center (SSIC):

  http://www.ibm.com/systems/support/storage/ssic
- IBM drive storage feature activation (DSFA):

  http://www.ibm.com/storage/dsfa
- *Using IBM DS8870 in an OpenStack Environment,* REDP-5220

7

IBM XIV Storage System

The IBM XIV Storage System is a high-end, flash-optimized, fully scalable enterprise disk storage system that is based on a grid of standard, off-the-shelf hardware components. The enterprise-class business enabler is designed around a self-healing, data protected cloud-centric foundation ideal for public, private, or hybrid clouds.

XIV cloud economics provide world-class business continuity for applications that require zero tolerance downtime. Its predictable, consistent, high performance design allows organizations to take control of their storage, and gain business insights from their data. The architecture of the system is designed to deliver performance, scalability, and ease of management while harnessing the high capacity and cost benefits of serial-attached SCSI (SAS) drives.

Optional IBM Real-time Compression technology helps reduce the amount of physical storage that is required in your environment. You can reuse free space in the existing storage without archiving or deleting data.

This chapter includes the following sections:

- Product overview
- IBM Real-time Compression on the IBM XIV Storage System
- Architecture and key components
- Related information

7.1 Product overview

IBM XIV is a high-end, grid-scale storage system that stands out for tuning-free high performance, extreme ease of use, and excellent data economics including inline, field-proven IBM Real-time Compression. XIV is ideal for cloud. It offers high service levels for dynamic workloads, easy hyper-scaling including in multi-tenant environments, and flexible consumption models, as well as robust cloud automation and orchestration through OpenStack, RESTful API, and VMware. It offers security and data protection through advanced mirroring, hot encryption and self-healing, and XIV storage abstraction through vSphere Storage APIs for Storage Awareness (VASA) 2.0 support. Industry benchmarks underscore stellar XIV performance and cost benefits.

Built with IBM Spectrum Accelerate, IBM XIV offers unified management and operational agility across your data center and hybrid clouds. For more information about IBM Spectrum Accelerate, see 2.3.2, “IBM Spectrum Accelerate” on page 35.

7.1.1 XIV models

The XIV Gen3 Storage System family consists of two machine types and two models:

- XIV Storage System Gen3 Model 214
- XIV Storage System Gen3 Model 2810/ 2812- 314

XIV Model 214

The XIV Storage System Gen3 Model 214 shares all of the concepts and software architecture of its predecessor Model 114. The primary hardware component upgrades to the previous model are in the following areas:

- Interface and data modules:
  - Processor
  - Memory
  - Power Supply Unit
  - 1 TB, 2 TB, 3 TB, 4 TB, and 6 TB SAS self-encrypting drives (SEDs)
  - 400 GB flash drive cache option or 800 GB flash drive with 4 TB configurations
- Interface modules:
  - Two-port 10 GbE adapter
- Software:
  - IBM XIV software version 11.5

XIV Model 314

IBM XIV Storage System Model 314 is the latest addition to the XIV family. It is a hardware-enhanced XIV Gen3 storage system that is targeted for customers requesting high utilization of XIV IBM Real-time Compression without performance degradation. It has these improvements:

- Compared to Model 214, its compressed capacity is now up to five times larger.
- On compressed workloads, performance sees an improvement of 50% or more (with IBM XIV software version 11.6.1). It can also compress smaller volumes.
- The Real-time Compression feature is included in the IBM XIV license, and the feature is enabled by default.

7.1.2 Did you know?

IBM XIV provides the following advantages:

- IBM XIV has one of the fastest rebuild times in the storage industry.
- IBM XIV has all functions and features included and supporting software at no extra cost, such as remote mirroring, snapshots, encryption, and VMware integration.
- IBM XIV Gen3 offers industry-standard data at rest encryption while avoiding the performance impact of SED. Encryption can be turned on non-disruptively at any time.
- IBM XIV self-optimizes automatically upon changes in the hardware configuration, such as the addition of modules or replacement of modules upon failure.
- Capacity on Demand (CoD) configurations of XIV Storage Systems are available to provide a method to quickly provision new physical storage and can be ordered at any of the valid partially populated capacities
- The IBM XIV Gen3 Model 314 has the following features:
  - A compression feature is enabled by default, with no separate licensing required
  - Dedicates 48 GB of RAM to IBM Real-time Compression
  - The user configurable max_soft_size can be up to 2 PB
  - The total number of compressed volumes and snapshots is increased to 1024

The IBM XIV Storage System is designed to be a scalable enterprise storage system based on a grid array of hardware components.

It is an innovative, high-end, open disk system series that is designed to support business requirements for a highly available information infrastructure. XIV architecture is a grid of standard Intel and Linux components connected in any-to-any topology. This architecture is designed to support enterprise-class reliability, performance, scalability, and energy efficiency.

Figure 7-1 shows the XIV GUI with the front view of the XIV Gen3 system.

Figure 7-1 BM XIV Management GUI System view

7.1.3 Core features

The XIV design has the following features:

- Workload balancing. The workload is evenly distributed over all hardware components at all times. All disks and modules are used equally, regardless of access patterns. Pseudo-random distribution ensures consistent load-balancing even after adding, deleting, or resizing volumes, as well as adding or removing hardware. This balancing of all data on all system components eliminates the possibility of a hotspot being created.
- True virtualization. Unlike other system architectures, storage virtualization is inherent to the basic principles of the XIV Storage System design. Physical drives and their locations are hidden from the user, which dramatically simplifies storage configuration. The automatic layout maximizes the system's performance by using system resources for each volume, regardless of the user's access patterns. With the XIV, you no longer decide or plan which type of data protection to use for storing data or plan how many drives you need to dedicate for a specific application. You just have to allocate the needed storage capacity, and the system does the rest. Storage administrators can respond to growing storage needs in minutes instead of the hours or days required with traditional storage systems.
- Ease of management. Ease of management that goes below the surface, with a virtualized architecture, highly intuitive Management GUI, and clever implementation of

role-based management all add up to create a streamlined user experience with less administration.

- Fast rebuild. The XIV rebuild times are fast because it keeps track of what blocks contain data and only rebuilds these blocks. This approach results in rebuild times of less than an hour for a 6 TB drive failure, when the XIV is 100% utilized.

7.1.4 Storage efficiency

The following XIV characteristics contribute to storage efficiency

- Real-time Compression. The Real-time Compression engine used in the XIV has been field-proven for several years. For Model 214, a separate license is required for Real-time Compression. For Model 314, Real-time Compression is enabled by default and does not require an additional license.

  The XIV Storage software V11.6 and higher includes a comprestimator utility that provides specific compression predictions for workloads. Using compression with XIV results in much higher data density, reducing the storage system total cost of ownership (TCO). The compression feature can be applied with a single click to an existing volume if the system predicts a good compression ratio by using its built-in estimator. There is no need to schedule compression or to schedule conversions between compressed/uncompressed states.
- Efficient capacity usage. Capacity utilization functionality, including space-reclamation for virtualized environments, space-efficient snapshots, and advanced thin provisioning, helps support savings in capacity and cost with fewer administration requirements.
- High-density storage. Enterprise-class, high-density SAS disk drives create increased physical space efficiencies. Using 6 TB drives, XIV Gen3 provides up to 485 TB usable capacity in a single rack. The ability to store so much data in one system with fewer, large capacity drives, and the use of multi-core processors, can help to reduce power and cooling expenses for a more energy efficient solution.

7.1.5 Why XIV is ideal for cloud

IBM XIV Cloud Storage for Service Providers (or client hosted cloud solutions) require elastic, rapid response, and generate dynamic demand that is met with IBM XIV's 485 TB linear scalability and immediate and automatic redistribution, without interruption, of added capacity. The grid design stripes data across all modules and disks, incorporating data redundancy.

IBM XIV is prepared for the cloud with VMware, OpenStack, and IBM SoftLayer Cloud. It has these characteristics:

- Capacity: Up to 2 PB of effective capacity per system and 288 PB with hyper- scale management of 144 systems. XIV dramatically reduces total cost of ownership by increasing the effective capacity while providing cost savings on physical space, power consumption, and related administrative costs.
- Predictability: Grid architecture that supports SLAs through optimal resource sharing and high I/O predictability without the need for complex analysis or performance tuning.
- Simplicity: Autonomic data distribution, tuning-free high performance, easy capacity planning, and multi-tenant administration; acquisition flexibility; no hidden costs
- Resiliency: Grid redundancy; fast disk rebuilds; uninterrupted data access; data-at-rest encryption
- IBM XIV integrates with VMware vRealize Suite. This functionality is available to all IBM XIV releases through the IBM Spectrum Control Base Edition version 2.1. It is available for

no additional charge and enables VASA, VMware vSphere Web Client, vCenter Orchestrator (vCO), vCenter Operations Manager (vCOps), and VMware vCloud Automation Center (vCAC). VMware Virtual Volumes (VVOLs) is designed to enable upcoming VMware vSphere environments to easily automate IBM XIV provisioning, offload snapshots and cloning, and instantly reclaim space. The IBM Spectrum Control Base Edition version 2.1 delivers a range of IBM storage integration services in cloud-based architectures that provide provisioning, automation, and monitoring. For more information, see *XIV Storage System in a VMware Environment*, REDP-4965.

- IBM XIV can automate storage discovery and provisioning in OpenStack environments by using OpenStack Cinder integration. To manage the cloud more easily, deploy the IBM Cloud Manager. For more information, see *Using XIV in OpenStack Environments*, REDP-4971.
- IBM XIV can automate custom storage operation by using the IBM XIV RESTful API. For more information, see *Exploring RESTful API Support with XIV*, REDP-5064.
- IBM XIV multi-tenancy allows secure isolation into domains of IBM XIV storage resources among numerous tenants, with the ability of setting different quality of service (QoS) levels per domain. It enables the division of storage system administration tasks into logical domains by using role-based permissions. It also enables rapid deployments while also minimizing the need for extensive planning, tuning, or field-upgrades.
- IBM XIV scales storage capacity and performance linearly and can manage up to 144 systems by using IBM Hyper-Scale Manager.
- The IBM XIV Cloud Storage for Service Providers product further empowers users with the flexibility to combine a robust set of base functions with license-per-need features. The ability to pay for functionality when it is required enables cloud providers to tailor services to customer needs even more cost effectively.
- Together with IBM Spectrum Accelerate, IBM XIV can be configured to support disaster recovery in Cloud. See *Deploying IBM Spectrum Accelerate on Cloud*, REDP-5261.

7.1.6 Security and encryption

IBM XIV storage system addresses the data-at-rest encryption demands.

- *S*elf-encrypting drives: IBM XIV Gen3 Model 314 offers industry-standard data at rest encryption at no extra cost, while avoiding performance impact with SEDs. IBM Security Key Lifecycle Manager is used to manage the security keys. Encryption can be turned on non-disruptively with data initially already on disks. For more information about encryption with IBM XIV, see *IBM XIV Security with Data-at-Rest Encryption*, REDP-5047.
- Payment Card Industry Data Security Standard (PCI DSS)*:* Security enhancements such as auditing of user actions and user interface locking after a predefined period of idle time, allow IBM XIV to meet the requirements that are imposed by PCI DSS.

7.1.7 Management

The highly intuitive IBM XIV Management GUI and built-in management tools make administrative tasks easy and efficient, with little training or expertise required, from provisioning volumes to monitoring multiple systems.

Figure 7-2 is an example of a system view with a module pullout in the IBM XIV Management GUI.

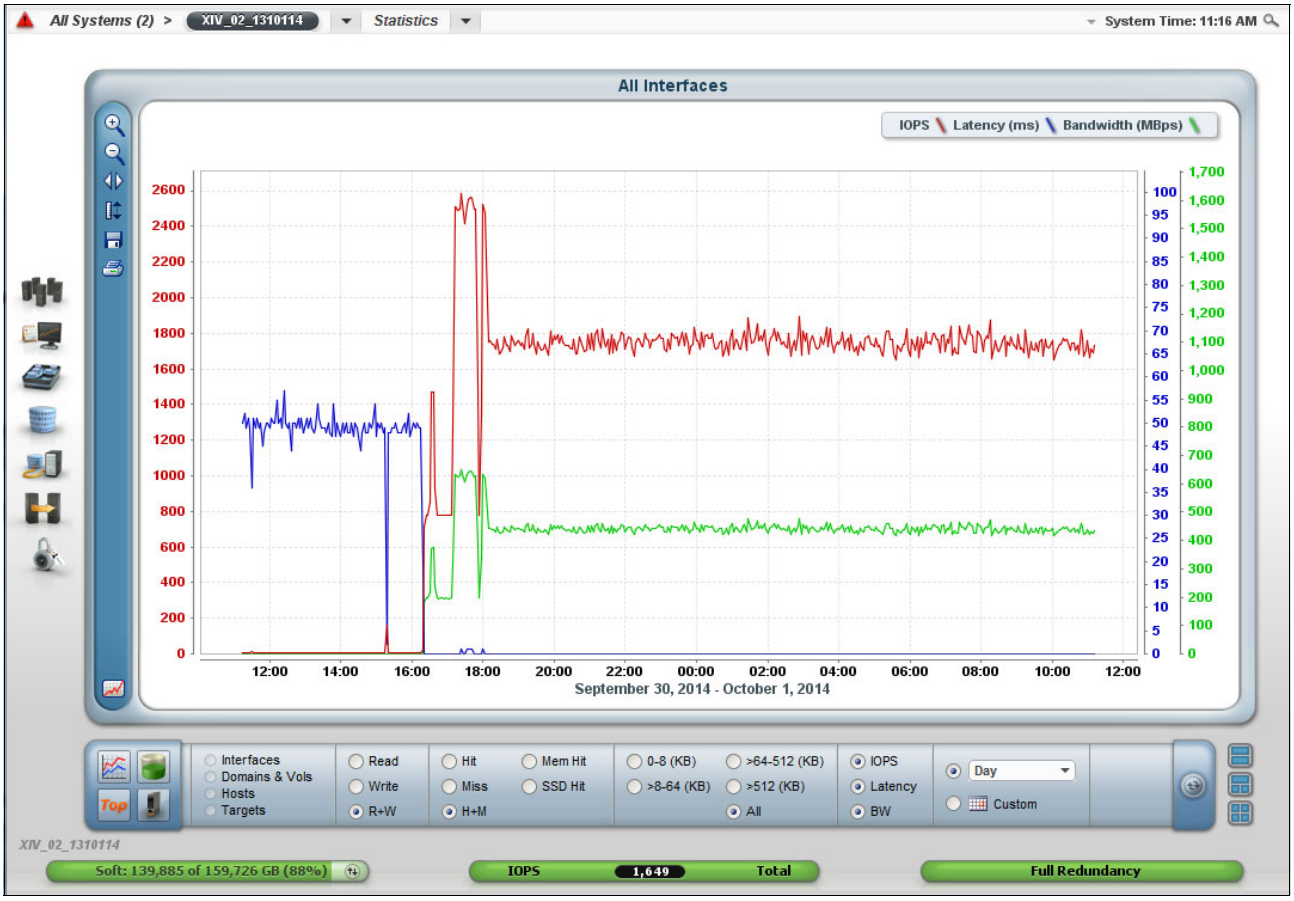

Figure 7-2 Statistics monitor view in IBM XIV Management GUI

XIV includes the following features to aid with management:

- The IBM XIV Management GUI acts as the management console for the storage system. A simple and intuitive GUI enables storage administrators to manage and monitor all system aspects easily, with almost no learning curve.
- The IBM XIV TOP application allows the user to view and monitor performance information for defined volumes and hosts in real time. It can be launched independently or from within the IBM XIV Management GUI.
- The XCLI is a comprehensive command-line interface that allows the user to configure and monitor the IBM XIV Storage System. All the functions available in the IBM XIV Management GUI are also available in the XCLI. It can be used in a shell environment for interactive commands or as part of a script to automate lengthy or complex tasks.
- Performance statistics can be monitored through the IBM XIV Management GUI and XCLI at any time. Monitoring through the IBM XIV Management GUI is easily done by selecting filters.

7.1.8 Monitoring with XIV Storage Management GUI and XCLI

The XIV Storage Management graphical user interface (GUI) and XIV Storage System Command-Line Interface (XCLI) include many features that allow you to monitor the system. These features include monitoring for alerts, events, and failed components.

You can show the health state, alerts, and events for multiple IBM XIV systems from the **Systems** menu. By selecting the Systems view, you can work with a particular storage system or show **All Systems Alerts** or **All Systems Events** for every XIV Storage System that is defined to the GUI. The **Systems** menu does not offer multiple selections when only one XIV Storage System is defined to the GUI.

The XIV Storage Management GUI allows users to work with alerts across multiple systems, regardless of which XIV Storage System is selected. You can access the alert information regardless of which system you are working with, or which window is displayed.

You can access the **Alerts**, **Events**, **Statistics**, and **QoS Performance Class** for the currently selected machine from the **Monitor** menu.

Extensive information and many events are logged by the XIV Storage System. The system captures entries about problems with various levels of severity, including warnings and other informational messages. These messages include information about logins, configuration changes, and the status of attached hosts and paths. All of the collected data can be reviewed in the Events window.

References: For information about the extensive XIV performance monitoring features, and the no extra licensing costs “all included” concept, see I*BM XIV Storage System Architecture and Implementation*, SG24-7659.

7.1.9 IBM Hyper-Scale

IBM Hyper-Scale is a family of growing technologies designed around an innovative approach to storage scalability. IBM Hyper-Scale includes the following features:

- Hyper-Scale Manager, which reduces operational complexity and enhances capacity planning through integrated management for large and multi-site IBM XIV deployments. It enables the IBM XIV Management GUI to access and operate on multiple IBM XIV systems concurrently. It also provides support for the RESTful application programming interface (API). The Hyper-Scale Manager runs on a single instance of a virtual machine server or on several servers.

Figure 7-3 illustrates how the Hyper-Scale Manager allows a storage administrator to work with volumes from multiple systems, in this case to create cross-system snapshots.

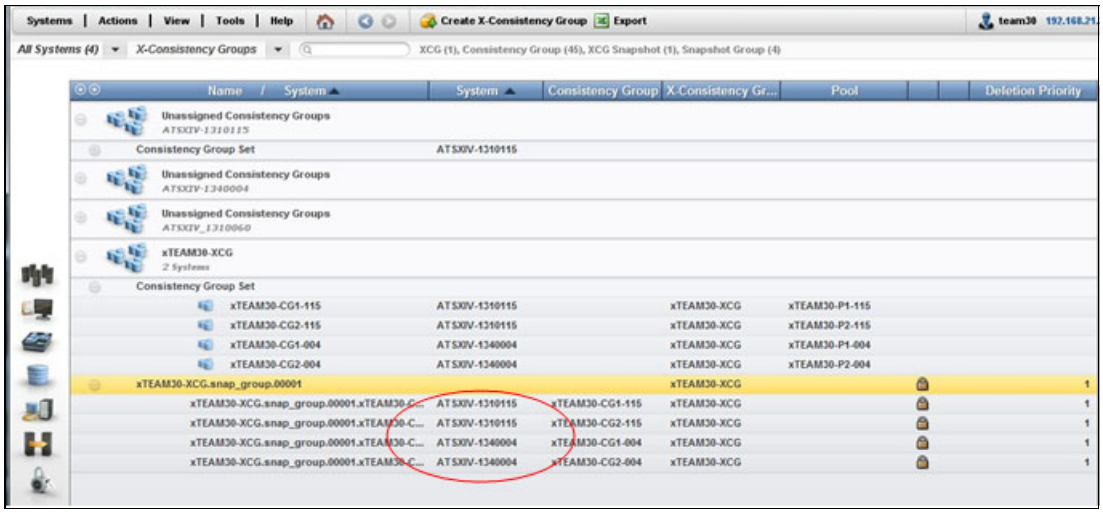

Figure 7-3 Hyper-Scale Manager

- Hyper-Scale Mobility, which is a powerful function that allows the user to move volumes between storage systems transparently, with no disruption to host applications. It is only available with Gen3 systems.
- Hyper-Scale Consistency, which with cross system consistency (or snapshot) groups enables a coordinated creation of snapshots for inter-dependent consistency groups on multiple systems. This feature is available only through the IBM Hyper-Scale Manager.

7.1.10 Business continuity

The IBM XIV Storage System provides a rich set of copy services functions suited for various data protection scenarios, and enables clients to enhance their business continuance and data migration capabilities:

- Synchronous mirroring is a data replication solution to use between two storage systems that achieves a recovery point objective (RPO) of zero with a distance of less than 100 km. In synchronous mirroring, a host write operation is completed on both the local and remote sites before an acknowledgement is returned to the host. Synchronous mirroring ensures local and remote always have the same copy of data for a zero RPO.
- Asynchronous mirroring is designed for a non-zero RPO at greater distances. Consistent sets of data are copied to the remote location at predefined intervals, whereas the host writes are acknowledged after they are written on the local site.
- Three-site (or 3-way) mirroring is a star topology multi-target disaster recovery solution. It utilizes proven, efficient IBM XIV technology with synchronous and asynchronous mirroring. This option is only available with Gen3 systems.
- IBM XIV data migration is a seamless data transfer tool for migrating data from another source system by simulating host behavior. It synchronizes data between the two storage systems by using transparent copying to the IBM XIV Storage System as a background process. It requires only a short outage on the host side to switch LUN ownership to the IBM XIV and begin the migration process.

7.1.11 Scalability and performance

IBM XIV Storage System is a scalable enterprise storage system based on a grid array of hardware components. The architecture offers the highest performance through maximized and balanced utilization of all disks, distributed cache implementation, and exceptional performance characteristics. It provides these advantages:

- Massive parallelism: With the grid architecture, the system ensures full usage of all system components. All volumes are spread across all spindles in the system. The system harnesses all storage capacity and all internal bandwidth. It takes advantage of all available processing power for host-initiated I/O activity and system-initiated activity, such as rebuild processes and snapshot generation.
- Processing power: The IBM XIV Storage System open architecture uses the latest processor technologies and is more scalable than solutions that are based on a closed architecture. The IBM XIV Storage System avoids sacrificing the performance of one volume over another, and therefore requires little to no tuning.
- Innovative cache memory: Up to 720 GB of total system cache with flexible and powerful cache implementation allows the IBM XIV system to use large slots for reads and pre-fetched data while enabling it to manage a smaller slot size for improved cache management and better performance.
- Flash Cache (optional): For ultra-high performance needs, IBM XIV Gen3 optionally offers up to 12 TB of management-free Flash Caching available to all system data. The flash caching option can be installed non-disruptively and can provide up to a 90% reduction in I/O latency for random read workloads with block size up to 64 KB. Operating with advanced flash algorithms, the latest IBM XIV Gen3 components help meet requirements for extremely high performance workloads.
- Enhanced connectivity: IBM XIV Gen3 offers multiple active/active IO interfaces with improved host connectivity with up to twenty-four 8 Gb FC ports and up to twelve 10 Gb Ethernet ports or twenty-two 1 Gb Ethernet ports for connecting to iSCSI-attached hosts.

7.1.12 Availability and serviceability

The IBM XIV Storage System maximizes continuous operation and minimizes the performance degradation that is associated with non-disruptive planned and unplanned events, while also providing the capability to preserve the data during a disaster. It provides these features:

- Self-healing and self-tuning allow disk failure to be taken care of by an efficient rebuild process that brings the system back to full redundancy in minutes. In addition, the IBM XIV Storage System extends the self-healing concept, resuming redundancy even after failures in components other than disks, or upon changes in the hardware configuration, such as the addition of modules.
- The rapid restoration of redundant data across all available drives and modules in the system during hardware failures and the automatic redistribution of data across all newly installed hardware are fundamental characteristics of the IBM XIV Storage System architecture.
- IBM XIV Storage System dynamically maintains the pseudo-random distribution of data across all modules and disks while ensuring that two copies of data always exist when the system reports *Full Redundancy*. When a disk drive or a module fails, the IBM XIV data distribution algorithms automatically identify the non-redundant partitions. The system then begins to make copies of just those non-redundant partitions and places this data in reserved areas of the disk drives that exist for just this circumstance.

- Non-Disruptive Code Load (NDCL) enables upgrades to the IBM XIV Storage System software from a current version to a later version without disrupting the application service.
- The provision within the IBM XIV Storage System to efficiently and flexibly create snapshots, coupled with the ability to define consistency groups of volumes, are integral elements of the data preservation strategy. In addition, the IBM XIV Storage System data mirroring functionality facilitates excellent recovery point and recovery time objectives as a central element of the full disaster recovery plan.

7.2 IBM Real-time Compression on the IBM XIV Storage System

Continued data growth and economic pressures are driving the rapid adoption of data reduction technologies. Although much of the attention regarding data reduction has focused on backup, many businesses are applying data reduction throughout the entire data lifecycle.

Traditionally, data deduplication was generally used with highly redundant data sets found in backup-to-disk, virtual tape, and physical tape library applications. Data deduplication can provide an acceptable solution for sequential access workloads that are less sensitive to performance. However, it has limited use cases, and delivers savings in areas such as virtual desktop infrastructure (VDI) and backup.

Other data that is not repetitive by nature does not typically benefit from deduplication savings. Therefore, it cannot meet the demanding performance requirements of primary storage for random-access, high transaction volumes, and high throughput.

Compression is suitable for most enterprise workloads and use-cases. It is the most common form of data reduction technology for backups and backups to disk, and might be the most common data reduction technology in disk-based systems.

IBM Real-time Compression Software that is embedded in IBM XIV Storage System addresses all the requirements of primary storage data reduction, including performance. It does so by using a purpose-built technology called IBM Real-time Compression.

Companies need to reduce data and maintain performance, which can be challenging with traditional compression methods.

Using IBM Real-time Compression provides an innovative approach that is designed to overcome these challenges.

7.2.1 IBM Real-time Compression with XIV

IBM Real-time Compression with the IBM XIV Storage System Gen3 effectively and efficiently answers a key requirement that typically challenges traditional approaches to capacity compression. It reduces capacity while maintaining high performance of the storage system. Using IBM Real-time Compression reduces the amount of physical storage that is required in your environment. You can reuse free space in the existing storage without archiving or deleting data.

IBM Real-time Compression provides the following benefits:

- It delivers substantial savings across a versatile range of enterprise workloads.
- It uses the IBM Random Access Compression Engine (RACE) technology, which was purpose-built for real-life primary application workloads. IBM Real-time Compression takes advantage of data temporal locality to maximize data savings and system performance.

It is related to the effect of reading in the same (or similar) order to the write order, and is one of the unique RACE technology core pillars (for a detailed explanation of data temporal locality). RACE enables a larger number of read I/Os to be served from cache, potentially resulting in faster application response time.

- It is easy to use. An administrator simply selects a **Compressed** check box when creating a new compressed volume. For an existing volume, the system can display an accurate estimation of potential compression savings in the XIV GUI.
- Being able to easily assess potential savings before compressing existing data facilitates deployment of IBM Real-time Compression in existing environments. Furthermore, nondisruptive compression, or alternatively, conversion of uncompressed volumes to compressed volumes (and vice versa) for existing volumes can provide an easy way to reclaim capacity and accelerate return on investment (ROI).
- It benefits from the XIV architecture because evenly distributed compression load across system resources increases performance and efficiency.
- It benefits from the XIV scaling of performance and the amount of system resources that can be allocated to IBM Real-time Compression.
- It works with primary active data. Due to the system's ability to preserve high performance consistency with compression, IBM Real-time Compression can be used with active primary data. Therefore, it supports workloads that are not candidates for compression in other solutions, and also supports compression of existing data.

Various configuration items affect the performance of compression on the system. To attain high compression ratios and performance, consider the following guidelines:

- If you have only a small number (10 - 20) of compressed volumes, configure them on one I/O group and do not split compressed volumes between different I/O groups.
- For larger numbers of compressed volumes on systems with more than one I/O group, distribute compressed volumes across I/O groups to ensure that access to these volumes is evenly distributed among the I/O groups.
- Identify and use compressible data only. Different data types have different compression ratios, and it is important to determine the ratio of the data that is currently on your system. You can use tools that estimate the compressible data or use commonly known ratios for common applications and data types. Storing these data types on compressed volumes saves disk capacity and improves the benefit of using compression on your system. Table 7-1 shows the compression ratio for common applications and data types.

Table 7-1 Compression ratios for different data types

Data Types/Applications	Compression Ratios
Databases	Up to 80%
Server or Desktop Virtualization	Up to 75%
Engineering Data	Up to 70%
Email	Up to 80%

IBM Real-time Compression uses the reliable, field-proven, and patented IBM RACE technology to achieve a valuable combination of high performance and compression efficiencies:

- It can lower the effective capacity requirements of a volume to 1/5 of the uncompressed capacity.
- No additional hardware is required to use IBM Real-time Compression.

- It can reduce cost for software that is licensed by capacity because less physical storage is required for compressed data.
- It can provide operational expense (OPEX) benefits because it requires no changes to the existing storage environment. It is fully integrated into XIV Gen3 version 11.6, and allows for nondisruptive compression of volumes.
- It can be enabled without changing your existing storage environment (applications, hosts, networks, fabrics, or external storage systems). The solution is not apparent to hosts, so users and applications continue to work as-is.
- Compression occurs within the XIV system itself. The conversion from non-compressed to compressed is inline and does not require downtime.
- Compressed volumes provide an equivalent level of availability as regular volumes. Compression can be implemented into an existing environment without an impact to service. The exception is mirroring, which must be temporarily stopped while volumes are converted. Existing data can be compressed transparently while still being accessed by users and applications.
- Compressed volumes can be mirrored, minimizing requirements for replication bandwidth and capacity requirements on the target system, and correspondingly maximizing system performance due to the reduction of data to transfer.

Remote volume copies are always compressed if the source is compressed. This process not only reduces storage requirements, but also uses less bandwidth because the data is transferred compressed. Mirroring and Hyper-Scale Mobility are faster and require less bandwidth because less data is transferred.

7.2.2 Common use cases

Compression savings are different on every data set. Some data sets can be compressed more than others, resulting in higher capacity savings. For example, image and video data, except for lossless types, are usually already compressed. Therefore IBM Real-time Compression might not provide more savings for these data types.

This section addresses the most common use cases for implementing compression:

- General-purpose volumes
- Databases
- Virtualized infrastructures

Understanding compression rates, ratios, and savings: To clarify the meaning of the terms *compression ratio*, *compression savings rate*, and *compression savings*, consider a use case where the original data physical capacity before compression was 100 TB, and the physical data capacity after compression is 40 TB.

The following values help to clarify these terms:

- Compression rate = 60%
- Compression savings rate = 60%
- Compression savings = 60 TB
- Compression ratio = original size (100 TB) divided by the size on disk after compression (40 TB) = 2.5:1

When you consider savings, it is easiest to use the compression rate.

The compression ratio helps in understanding how much effective data you can store on your system. So, when you have a 2.5:1 compression ratio, you will be able to store 250 TB of data on 100 TB of physical capacity.

Figure 7-4 shows typical capacity savings with IBM Real-time Compression.

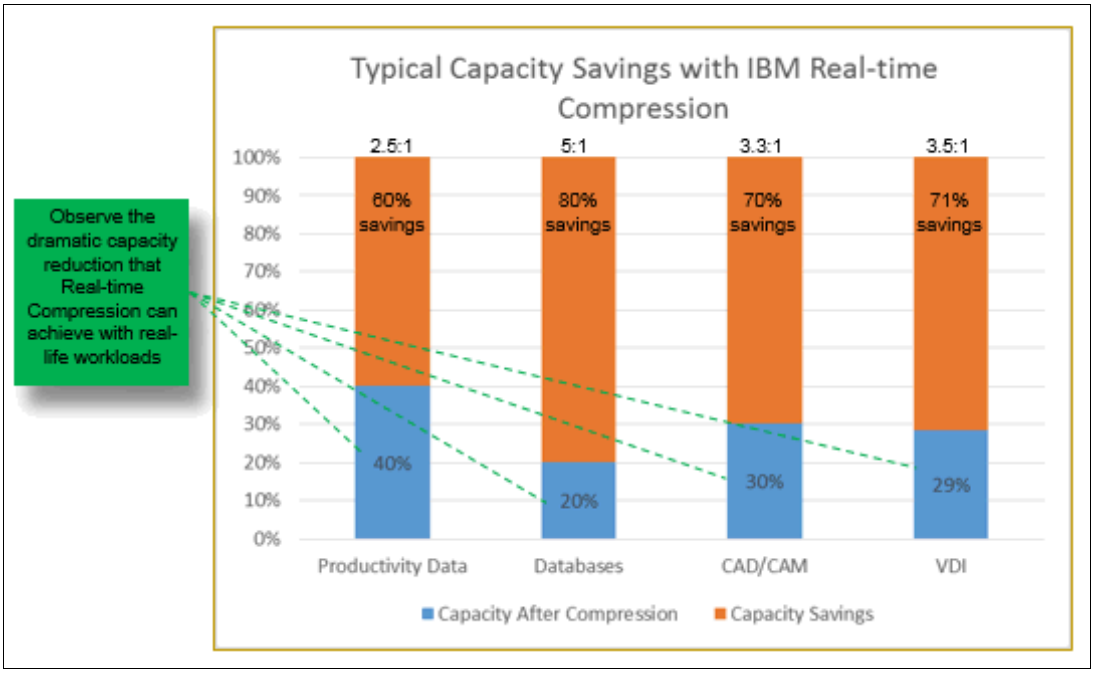

Figure 7-4 Typical capacity savings with IBM Real-time Compression

These websites are also relevant as further information sources:

- IBM XIV Storage System documentation

  http://www.ibm.com/support/knowledgecenter/STJTAG/com.ibm.help.xivgen3.doc/xiv_kcwelcomepage.html

- IBM XIV Real-time Compression Evaluation User Guide

  http://public.dhe.ibm.com/common/ssi/ecm/ts/en/tsj03645usen/TSJ03645USEN.PDF

Also, see the following publications:

- *Implementing IBM Real-time Compression on the IBM XIV Storage System*, REDP-5216
- *IBM Real Time Compression with IBM XIV Storage System Model 314*, REDP-5306

7.3 Architecture and key components

The IBM XIV Storage System Gen3 provides several architectural design factors that contribute to its unique operational capabilities. The IBM XIV Storage System is composed of these physical components (numbers correspond to those in Figure 7-5 on page 172):

- IBM XIV storage system is delivered in a standard IBM T42 rack, fully cabled for 15 modules, even if a partially populated rack is ordered.
- The hardware of the Interface Modules **[1]** and Data Modules **[2]** in the IBM XIV Storage SystemGen3 is based on an Intel server platform optimized for data storage
- A patch panel **[3]** at the rear of the rack is used for connecting the IBM XIV Storage System to the customer's network. The Fibre Channel ports on the Interface Modules are connected to the patch panel by using 50-micron cables. All external connections must be made through the patch panel. In addition to the host connections and to the network connections, more ports are available on the patch panel for service connections. **[4]**
- An Automatic Transfer Switch (ATS) **[5]** allows the IBM XIV Storage System to attach to two independent power sources to protect against loss of customer utility power. During an external power failure, the ATS automatically transfers the load to the redundant power supply.
- The IBM XIV Storage System has three internal Uninterruptible Power Supplies (UPSs) **[6]** that are used as cache batteries, so the data in cache always will be written down to disk during a power outage. Two of these UPSs are enough to maintain cache battery power availability. This redundant design is n+1 because the IBM XIV Storage System has one redundant UPS. The UPS complex has enough battery reserves to sustain two emergency power shutdowns.
- The 1U Maintenance Module **[7]** and the modem, which are installed in the middle of the rack, are used for IBM XIV Storage System support and for the IBM personnel to maintain and repair the system. This device is only used to gain remote access to the IBM XIV System through the modem for support personnel.
- The IBM XIV Storage System utilizes redundant InfiniBand switches **[8]** for communications between the interface and data modules. Each InfiniBand switch contains 36 ports that have 40-Gbps full bidirectional bandwidth per port. The switches are powered by redundant power supplies and fan modules to eliminate any single point of failure.

Figure 7-5 illustrates these physical components.

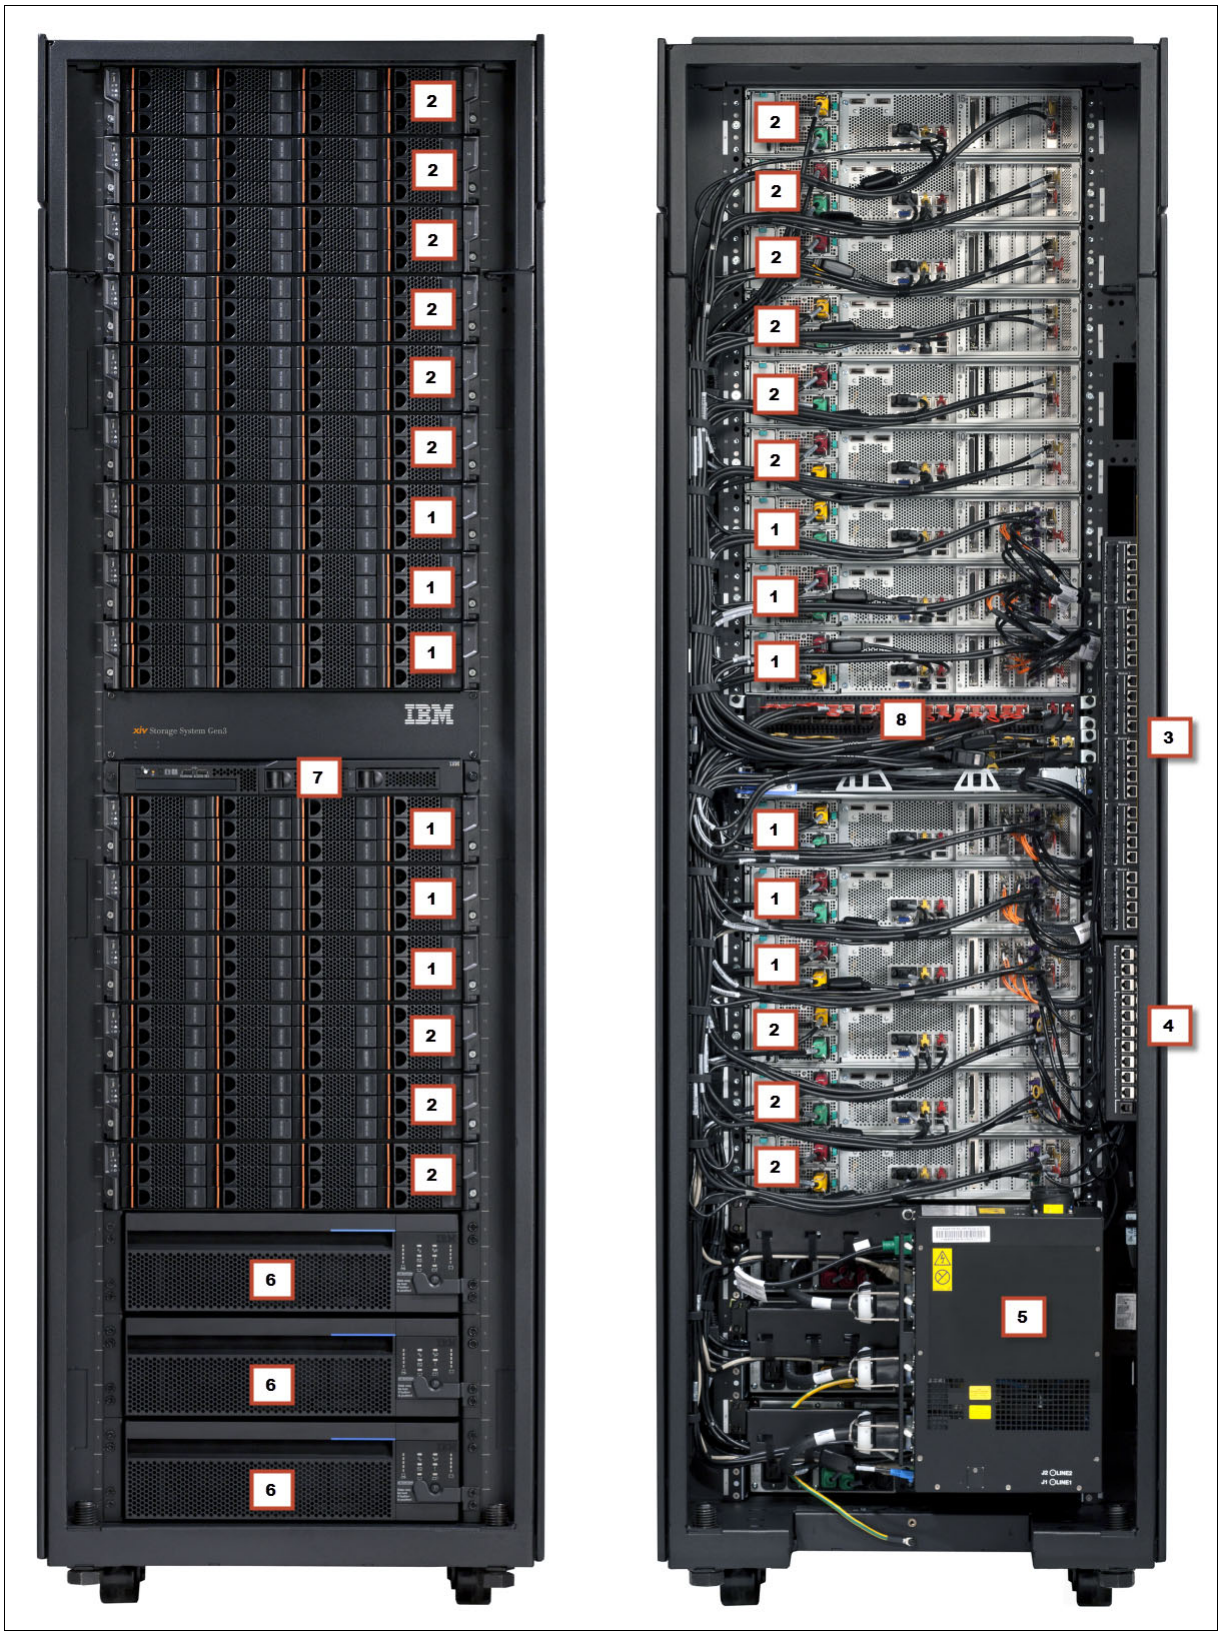

Figure 7-5 The physical components of an IBM XIV Storage System

7.3.1 Configuration

The rack specifications for capacity, connectivity, system resources, and power usage can be found in Figure 7-6. Note, the specifications differ based on the number of modules in the system, and the size of the disk drives it contains.

Rack Configuration

	Model 214 Only							
Total number of modules (Configuration type)	6 Partial	9 partial	10 partial	11 partial	12 partial	13 partial	14 partial	15 full
Total number of data modules	3	3	4	5	6	7	8	9
Total number of interface modules	3	6	6	6	6	6	6	6
Number of active interface modules	2	4	4	5	5	6	6	6
Interface module 9 state		Disabled	Disabled	Enabled	Enabled	Enabled	Enabled	Enabled
Interface module 8 state		Enabled	Enabled	Enabled	Enabled	Enabled	Enabled	Enabled
Interface module 7 state		Enabled	Enabled	Enabled	Enabled	Enabled	Enabled	Enabled
Interface module 6 state	Disabled	Disabled	Disabled	Disabled	Disabled	Enabled	Enabled	Enabled
Interface module 5 state	Enabled	Enabled	Enabled	Enabled	Enabled	Enabled	Enabled	Enabled
Interface module 4 state	Enabled	Enabled	Enabled	Enabled	Enabled	Enabled	Enabled	Enabled
FC ports	8	16	16	20	20	24	24	24
iSCSI ports (1 Gbps – mod 114)	6	14	14	18	18	22	22	22
iSCSI ports (10 Gbps – mod 214)	4	8	8	10	10	12	12	12
Number of disks	72	108	120	132	144	156	168	180
Usable capacity (1 / 2 / 3 TB) (Mod 214 only)	28 TB 55 TB 84 TB	44 TB 88 TB 132 TB	51 TB 102 TB 154 TB	56 TB 111 TB 168 TB	63 TB 125 TB 190 TB	67 TB 134 TB 203 TB	75 TB 149 TB 225 TB	81 TB 161 TB 243 TB
Usable capacity (4 / 6 TB)	112 TB 169 TB	177 TB 267 TB	207 TB 311 TB	225 TB 338 TB	254 TB 382 TB	272 TB 409 TB	301 TB 453 TB	325 TB 485 TB
# of CPUs (one per Module) Mod 214 # of CPUs (two per Module) **Mod 314**	6 n/a	9 **18**	10 **20**	11 **22**	12 **24**	13 **26**	14 **28**	15 **30**
Memory (24 GB per module w 1/2/3 TB) Memory (48 GB per module w 4/6 TB) **Mod 314: Memory (96 GB per module w 4/6 TB)**	144 GB 288 GB **n/a**	216 GB 432 GB **864 GB**	240 GB 480 GB **960 GB**	264 GB 528 GB **1056 GB**	288 GB 576 GB **1152 GB**	312 GB 624 GB **1248 GB**	336 GB 672 GB **1344 GB**	360 GB 720 GB **1440 GB**
{Optional for 1, 2 ,3, 4, 6 TB XIVs} 400 GB Flash Cache {Optional for 4, 6 TB XIVs} 800 GB Flash Cache	2.4 TB 4.8 TB	3.6 TB 7.2 TB	4.0 TB 8.0 TB	4.4 TB 8.8 TB	4.8 TB 9.2 TB	5.2 TB 10.4 TB	5.6 TB 11.2 TB	6.0 TB 12.0 TB

Figure 7-6 IBM XIV Gen3 rack configuration

Some of the configurations that are shown in Figure 7-6 include the following items:

- Partially populated configurations of 6, 9, 10, 11, 12, 13, or 14 modules, allowing for more granularity of capacity options. Modules can be added as capacity needs increase up to 15 modules. For these configurations, some of the interface modules are not available. Those are labeled as *disabled* in Figure 7-6.

  Note: The first column of configuration in Figure 7-6 is applicable to Model 214, not to Model 314. Model 314 starts with a nine module configuration.

- A fully populated rack contains nine data modules and six interface modules for a total of 15 modules. The total usable capacity is 161 TB for a complete system when equipped with 2 TB drives, 243 TB of usable capacity when fully populated with 3 TB drives, 325 TB of usable capacity when fully populated with 4 TB drives, and 485 TB of usable capacity when fully populated with 6 TB.
- Fully and partially populated systems can be configured with either 1 or 10 Gbps iSCSI ports. The number of ports available for the 1 Gbps option is 6 - 22. The number of 10 Gbps ports ranges is 4-12 per system.
- For Model 214, the memory option for fully and partially populated systems is 24 GB per module for systems configured with 1, 2, or 3 TB drives, and 48 GB per module for 4 or 6

TB drives. For Model 314, the memory for fully and partially populated systems is 96 GB per module for systems configured with 4 or 6 TB drives, and 48 GB of the 96 GB RAM is dedicated to Real-time Compression. Figure 7-7 specifies XIV Storage Systems detailed configuration.

IBM XIV Storage System (Model 214 and Model 314) — System specifications	
General properties	
Capacity per drive (nearline SAS)	1 TB,* 2 TB, 3 TB, 4 TB or 6 TB self-encrypting hard drives (SEDs). **Model 314** only supports 4 TB and 6TB hard drives, and is 8 TB and 10 TB hard drives ready.
Number of disk drives (min/max)	Model 214: 72/180, **Model 314**: 108/180
Encryption	All disk drives are SEDs. Encryption requires an external key management solution, such as IBM Security Key
Lifecycle Manager	
Performance features	
Maximum number of CPUs	15, **Model 314**: 30
Maximum number of CPU cores	90, **Model 314**:180 physical (360 logical cores with Intel Hyper-Threading technology)
Maximum memory	Up to 360 GB (24 GB of RAM per module for hard disks of 1 TB, 2 TB, 3 TB capacity) Up to 720 GB (48 GB of RAM per module for hard disks of 4 TB, and 6 TB capacity) **Model 314**: Up to 1440 GB (96 GB of RAM for hard disks of 4 TB and 6 TB capacity)
Maximum cache-to-disk bandwidth	480 Gbps
Flash caching (optional)	Up to 6 TB (400 GB flash caching per module) – all hard disk capacity points Up to 12 TB (800 GB flash caching per module) – 4 TB or 6 TB hard disks **Model 314**: Up to 12 TB (800 GB flash caching per module) – 4 TB or 6 TB hard disks
Note: When XIV encryption is enabled, the data on the flash drives is also encrypted.	
Connectivity	
Maximum number of ports—Fibre Channel	24 (8 Gbps ports)
Maximum number of ports—iSCSI over gigabit Ethernet	22 (1 Gbps ports) or 12 (10 Gbps ports)

Physical features	
Temperature	10°C – 35°C (50°F – 95°F)
Altitude (max)	2,134 m/7,000 ft
Humidity	25% – 80% noncondensing
Dimensions (height × width × depth)	202 cm × 66 cm × 120 cm (79.53 in. × 25.98 in. × 47.24 in.)
Maximum weight	1,041.5 kg (2,296 lb)
Clearance front/rear	120 cm/120 cm (47.24 in./47.24 in.)
Redundant power feed	Yes
Input voltage	180 – 264 V ac at 60 A or 30 A (±10%)
Power usage	Less than 10% increase over comparable Model 214 configuraiton
Host connectivity	
Fibre Channel rates	8 Gbps
iSCSI rates	1 Gbps or 10 Gbps
Capacity-on-demand configurations	Yes
Warranty	1 and 3 year limited warranty, onsite service, same day 24×7

Figure 7-7 Detailed Model 214 and Model 314 specification

7.3.2 Configuration options

The IBM XIV Storage System can be purchased with different options. Here are some highlights:

- Flash Cache is an optional feature available for IBM XIV Gen3 systems. The system can be equipped with 400 GB of flash cache for every module. In addition, the IBM XIV Gen3 systems with 4 TB or 6 TB drives can be equipped with 400 GB or 800 GB flash cache for every module. For Model 314, 800 GB flash cache is used regardless of the capacity of hard disks used.
- IBM XIV Cloud Storage for Service Providers offers an innovative pay-per-need scalability matrix. The offer empowers cloud providers to cater to tenant requirements with extreme ease, and efficiency through flexible software licensing. See announcement letter ENUS214-173 at the following link:

  http://www.ibm.com/common/ssi/cgi-bin/ssialias?infotype=AN&subtype=CA&appname=gpateam&supplier=897&letternum=ENUS214-173&pdf=yes
- Advanced system placement is a pay-as-you-go program that lets an organization purchase the IBM XIV Storage System for a fraction of the price upon installation. When the system reaches a predetermined capacity threshold, the client will be charged for the full balance of the system and will be able to acquire the next system for $1. This program is designed for rapidly growing environments.
- Capacity on demand is a program that is designed to help organizations with less aggressive growth projections. A Gen3 system can be ordered with a certain amount of authorized storage capacity along with extra capacity that is not intended for initial use. The extra capacity will be purchased as needed when allocated to a storage pool.

7.3.3 Warranty info and upgrades

IBM XIV Storage System Model 314 Device Type 2812 and Device Type 2810 offer a flexible warranty choice. Device Type 2812 supports a 3-year warranty to complement the 1-year warranty offered by the functionally equivalent Device Type 2810. The warranties offer these benefits:

- Technical Advisors (TAs) are available with the initial warranty and for an extra fee after the warranty is up. The TA will support the initial installation and updates of the IBM XIV Storage System. The TA will also enhance end to end support by being a client advocate to proactively manage problem resolution through the support process for hardware.
- Call home and remote support
  - Call home is the capability of the system, when enabled, to send event notifications to the IBM XIV Remote Support Center. This capability enables proactive and failure notifications to be sent directly to IBM for analysis.
  - Remote support enables trained IBM service personnel to connect to the IBM XIV Storage System to analyze a problem, repair it remotely if possible, and assist the onsite IBM service support representative (SSR).

7.4 Related information

For more information, see the following resources:

- *IBM XIV Storage System Architecture and Implementation*, SG24-7659
- *IBM XIV Storage System Business Continuity Functions*, SG24-7759

- *IBM XIV Gen3 with IBM System Storage SAN Volume Controller and Storwize V7000*, REDP-5063
- *IBM XIV Storage System: Host Attachment and Interoperability*, SG24-7904
- *Solid-State Drive Caching in the IBM XIV Storage System*, REDP-4842
- *IBM XIV and VMware Synergy with IBM Spectrum Control Base Edition*, REDP-5131
- *XIV Storage System in a VMware Environment*, REDP-4965
- *IBM XIV Security with Data-at-Rest Encryption*, REDP-5047
- *IBM Hyper-Scale in XIV Storage*, REDP-5053
- *IBM XIV Storage System: IBM Hyper-Scale Mobility Overview and Usage*, REDP-5007
- *IBM XIV Storage System Thin Provisioning and Space Reclamation*, REDP-5001
- *Using XIV in OpenStack Environments*, REDP-4971
- *Exploring RESTful API Support with XIV*, REDP-5064
- *IBM XIV Storage System Multi-Site Mirroring*, REDP-5129
- *Space Reclamation in IBM XIV with Windows 2012*, TIPS1011
- IBM Hyper-Scale: A powerful new approach to scaling storage management (XIV)

  http://www.ibm.com/common/ssi/cgi-bin/ssialias?infotype=SA&subtype=WH&htmlfid=TSW03281USEN#loaded
- IBM Offering Information page (announcement letters and sales manuals)

  http://www.ibm.com/common/ssi/index.wss?request_locale=en

8

IBM tape products

This chapter describes the full range of IBM storage tape drives, tape automation products, and tape virtualization products.

The IBM storage tape products portfolio covers the needs of a wide spectrum of implementations, from entry-level to large enterprise.

This chapter includes the following sections:

- IBM tape drives
- IBM tape automation products
- IBM tape virtualization products

8.1 IBM tape drives

This section describes the IBM tape drives.

It also describes Linear Tape-Open (LTO) technology and the related Ultrium specifications, including the latest IBM LTO-7 tape drive and cartridges. It provides information about the TS1100 (formerly 3592) tape technology. It also briefly describes the IBM midrange and enterprise storage media that is used in LTO and TS1150 tape drives.

Finally, it provides comprehensive information about the strategic IBM tape systems and briefly summarizes the next steps in the development of tape technology.

8.1.1 LTO technology

LTO is an open format technology that is defined through a *tape cartridge standard*. This standard means that users can have multiple sources of products and media. Compatibility of all compliance-verified LTO Ultrium manufacturers is guaranteed. The LTO technology established a new open format specification for high-capacity, high-performance storage products and addresses a growing client need for improved data interchange across platforms.

LTO technology was developed jointly by IBM, Hewlett Packard (HP), and Seagate in 1997 to provide a clear and viable choice in an increasingly complex array of tape storage options. The current technology provider companies are IBM, HP, and Quantum. In 2000, the first generation of LTO became a product called LTO Ultrium 1. The current generation is LTO Ultrium 7.

During the generations, additional features and functions were agreed by the consortium:

- Write Once, Read Many (WORM) cartridges have been certified since LTO Ultrium 3. They are designed for applications that require an audit trail, such as specified in the Sarbanes-Oxley Act, Health Insurance Portability and Accountability Act (HIPAA), and many other legal regulations in different countries. See Figure 8-1.

Figure 8-1 WORM cartridge

- Data encryption has provided information security with support for encryption and key management since LTO Ultrium 4.
- Cartridge partitioning enables the use of the Linear Tape File System (LTFS) format, which is supported in LTO 5 and later. For more information, see 2.3.4, “IBM Spectrum Archive” on page 42.

- Changes to the compression buffer have allowed higher data compression since LTO Ultrium 6 (from 2:1 to 2.5:1)
- An IBM exclusive feature called Statistical Analysis and Reporting System (SARS) that provides a high level of preventive diagnostic reporting to help isolate failures between media and hardware.

The LTO technology offers investment protection due to compatibility with earlier versions for read and write. This compatibility is defined by reading two generations backward and reading / writing one generation backwards. For example, LTO Ultrium 7 drives can read and write LTO Ultrium 6 cartridges, and read LTO Ultrium 6 and Ultrium 5 cartridges.
LTO Ultrium 4 cartridges are not supported in LTO Ultrium 7 drives)

The uncompressed capacity per cartridge has improved from 100 GB (LTO Ultrium 1) to 6 TB per cartridge at LTO Ultrium 7.

Figure 8-2 shows a document published by the LTO consortium that displays the history and future of the LTO technology.

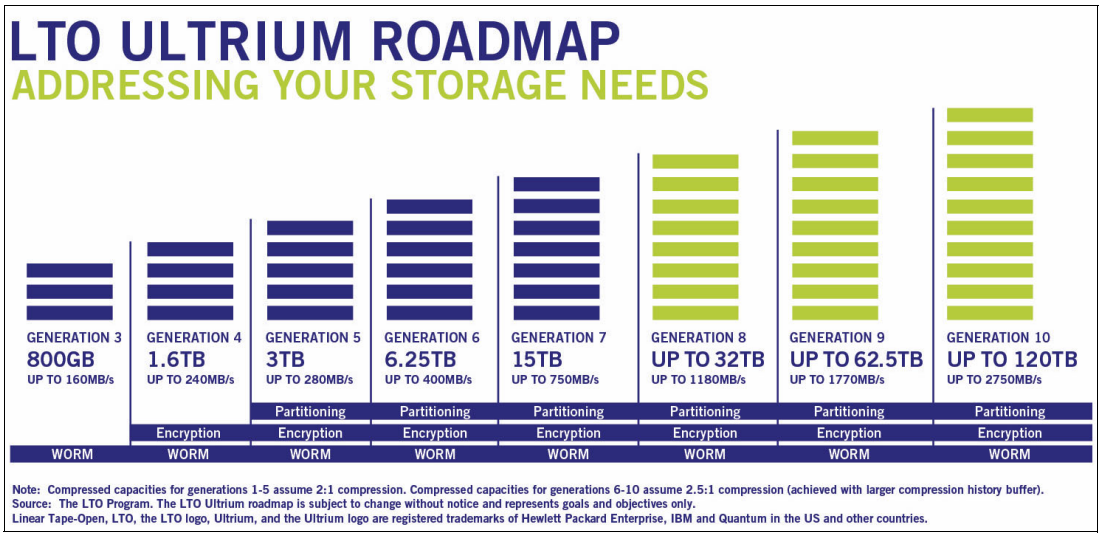

Figure 8-2 LTO Ultrium roadmap as published by the consortium (IBM, HP, Quantum)

For more information about the IBM exclusive LTO features, see the following website:

`http://www.ibm.com/systems/storage/media/lto/`

See more general specifications and consortium details on the official LTO website:

`http://www.lto.org`

8.1.2 IBM LTO tape drives

The IBM LTO Ultrium family of tape drives consists of the following products:

- IBM TS2250 tape drive
- IBM TS2260 tape drive
- IBM TS2270 tape drive
- IBM TS2350 tape drive
- IBM TS2360 tape drive

IBM LTO half-height tape drives

The IBM half-height form-factor LTO tape drives are designed for rack environments where space is a premium requirement. With its enclosure design that allows two units to be mounted side by side in a 19-inch IBM server rack, these drives offer lower-cost implementation.

IBM half-height LTO drives offer the same high capacity and the same data transfer rate as full-height LTO drives. The stand-alone or rack mountable single drives match current backup and restore or archiving requirements in small or medium environments. These drives are often used for departmental backup or data import and export.

Table 8-1 shows a comparison of the current available IBM LTO drives with a half-height form factor.

Table 8-1 Comparison of IBM LTO half-height drives

	TS2250	**TS2260**	**TS2270**
Model	3580 H5S	3580 H6S	3580 H7S
Physical capacity	1.5 TB native 3 TB with 2:1 compression	2.5 TB native 6.25 TB with 2.5:1 compression	6 TB native 15 TB with 2.5:1 compression
Data transfer rate	Up to 140 MBps native	Up to 160 MBps native	Up to 300 MBps native
Server attachment	6 Gbps SAS	6 Gbps SAS	6 Gbps SAS
Encryption and WORM support	Yes		Yes
Partitioning / LTFS support	Yes / Spectrum Archive single drive edition	Yes / Spectrum Archive single drive edition	Yes / Spectrum Archive single drive edition
Rack mountable	Yes. Use the IBM 7216 Multimedia Enclosure. http://www.ibm.com/systems/storage/tape/7216/		
Web links	http://www.ibm.com/systems/storage/tape/ts2250/index.html	http://www.ibm.com/systems/storage/tape/ts2260/index.html	http://www.ibm.com/systems/storage/tape/ts2270/index.html

Note: All IBM LTO tape drives offer compatibility with earlier versions for cartridge investment protection. This is defined are reading LTO cartridges two generations backward, and reading and writing LTO cartridges one generation backward.

Figure 8-3 shows a TS2270 LTO half height tape drive for reference. The TS2250 and TS2260 have the same size and form factor.

Figure 8-3 TS2270 tape drive

IBM LTO full-height tape drives

The IBM full-height form-factor LTO tape drives are designed for stand-alone or desk top usage. IBM full-height LTO drives match the requirements of today's backup and restore or archiving requirements in small or medium environments. These drives are often used for departmental backup or data import and export.

Table 8-2 shows a comparison of the current available IBM LTO drives with a full-height form factor.

Table 8-2 Comparison of IBM LTO full-height drives

	TS2350	**TS2360**
Model	3580 S53	3580 S63
Physical capacity	1.5 TB native 3 TB with 2:1 compression	2.5 TB native 6.25 TB with 2.5:1 compression
Data transfer rate	Up to 140 MBps native	Up to 160 MBps native
Server attachment	6 Gbps SAS	6 Gbps SAS
Encryption and WORM support	Yes	
Partitioning/LTFS support	Yes/Spectrum Archive single drive edition	Yes/Spectrum Archive single drive edition
Rack mountable	No	
Web links	http://www.ibm.com/systems/storage/media/lto/	http://www.ibm.com/systems/storage/tape/ts2360/index.html

Note: All IBM LTO tape drives offer compatibility with earlier versions for cartridge investment protection. This compatibility is defined by reading LTO cartridges two generations backward, and reading / writing LTO cartridges one generation backward

Figure 8-4 shows a TS2360 LTO full-height tape drive for reference. The TS2350 has the same size and form factor.

Figure 8-4 TS2360 tape drive

8.1.3 IBM TS1100 tape drive family

In 2006, IBM demonstrated technology leadership by recording 8 TB to a half-inch format tape. In 2010, IBM demonstrated recording 35 TB of uncompressed data to a half-inch format tape. In 2015, IBM scientists demonstrated an areal recording density of 123 billion bits of uncompressed data per square inch on low cost, particulate magnetic tape. This breakthrough represents the equivalent of a 220-terabyte tape cartridge that can fit in the palm of a hand.

IBM used the evolutionary progression of technology building blocks that was set in place over the preceding years to make these advances. An enterprise tape drive roadmap (Figure 8-2 on page 179) was laid out to ultimately reach and far exceed the 1 TB in native cartridge capacity within the 3592 tape drive generation. The 3592 model J1A became the first tape drive generation of the Enterprise Tape family. It enabled the storage of 300 GB of data on a cartridge (900 GB with 3:1 compressible data).

The same cartridges were to be reused by the second generation of 3592 tape drives. By using the JA media, 500 GB can be stored. With the high-capacity JB cartridge, 700 GB without compression was able to be stored. With the third generation of IBM 3592 tape drive and the JA media, it was possible to store 640 GB and 1 TB on the JB media. The first milestone was achieved: IBM delivered a 1 TB tape technology to the market.

With the fourth generation of the IBM 3592, the TS1140 model E07, IBM again takes tape capacity to a new level. The TS1140 can store 1.6 TB of data on the JB cartridge, and 4 TB of data on the new advanced JC cartridge. IBM kept its promise for IBM 3592 tape drives to

provide a 4 TB tape drive. The TS1140 tape drive is not only a new drive, but proof of the commitment of IBM to further tape technology development.

The new TS1150 tape drive gives organizations an easy way to deliver fast access to data, improve security, and provide long-term retention with far less expense than disk solutions. The TS1150 offers high-performance, flexible data storage with support for data encryption. This fifth-generation drive can help protect investments in tape automation by offering compatibility with existing automation. In addition, an upgrade model is available for existing TS1140 tape drives. Also, the TS1150 supports the LTFS format for direct, intuitive, and graphical access to data with IBM Spectrum Archive.

For more information, see the following website:

http://www.ibm.com/systems/storage/tape/index.html

Nomenclature

Because of many common considerations for the TS1120, TS1130, TS1140, and TS1150 tape drives, and to make reading more convenient, the following names are used for tape drives in this chapter:

- 3592 is used to refer to all models
- 3592-J1A is used for the first 3592 generation
- TS1120 is used for the 3592 E05 tape drive
- TS1130 is used for the 3592 E06 tape drive
- TS1140 is used for the 3592 E07 tape drive
- TS1150 is used for the 3592 E08 tape drive

IBM TS1130 tape drive overview

The IBM TS1130 tape drive, which is also referred to as the IBM 3592 model E06 and EU6, is the third tape drive generation of the IBM 3592 tape family. This generation provides higher capacity and performance when compared to its predecessor 3592 model E05.

The TS1130 tape drive features storage capability to help you establish easy access to data, better security, long-term retention, and data governance and regulatory compliance. The TS1130 tape drive offers high-performance flexible data storage with support for data encryption. The TS1130 tape drive can help you protect your investments in tape automation by offering compatibility with existing automation. To further protect your investment, an upgrade model is available for your existing TS1120 tape drives. To support a heterogeneous server environment, the TS1130 offers multi-platform support.

The TS1130 tape drive supports the TS3400 and TS3500 tape libraries, IBM 3494 tape libraries, IBM TS7700, IBM racks that enable stand-alone installation, and IBM 3952 tape frames model C20 (3952C20 frame) when they are attached to a Sun 9310 library.

Use the tape drive performance to help meet your business needs. The TS1130 tape drive supports a native data transfer rate of up to 160 MBps to deliver information at the speed that your business demands. Help reduce backup and recovery times or require fewer resources to support the environment by using the TS1130 tape drive to transfer data.

Figure 8-5 shows the TS1130 tape drive.

Figure 8-5 TS1130 tape drive

Highlights

The TS1130 tape drive has the following highlights:

- Provides information security with support for encryption and key management
- Optimizes information retention with support for existing IBM tape automation
- Supports WORM cartridges to help satisfy compliance requirements
- Offers high performance and high capacity for storage consolidation
- Provides virtual backhitch and high-resolution directory to help improve access to data and reduce wear and tear on media

Key features of the TS1130 tape drive

The TS1130 tape drive has the following key features:

- Support for 320 MBps in open system environments where data typically compresses at 2:1
- High performance with data transfer rate of up to 350 MBps with 3:1 compression
- Flexible media, including short and long length cartridges and rewritable and WORM formats
- Small form factor and high capacity
- Support for encryption and key management
- Compatibility with existing IBM tape automation installations
- Support for IBM System i5®, IBM System p5®, z Systems, and System x

Hardware summary

The TS1130 tape drive includes the following features:

- Linear serpentine recording
- 1 TB capacity with JB/JX media
- 640 GB capacity with JA/JW media
- 128 GB capacity with JJ/JR media
- 400 MBps burst data rate
- Compact 3.8 inch x 7.8 inch x 18.4 inch dimensions

For more information, see the following link:

`http://www.ibm.com/systems/storage/tape/ts1130/index.html`

IBM TS1140 tape drive overview

Business demands require applications that rely on high capacity, fast access to data, and long-term data retention. The TS1140 tape drive features storage capabilities to help you establish easy, rapid access to data, improve security, provide long-term retention, and help maintain data governance and regulatory compliance. The TS1140 tape drive offers high-performance, flexible data storage with data encryption support. The drive can help you protect your investments in tape automation by offering compatibility with existing automation.

To further protect your investment, an upgrade model is available for your existing TS1130 tape drives. For flexible storage in heterogeneous server infrastructures, the TS1140 tape drive offers multi-platform support. It can scale from midrange to enterprise environments. The TS1140 tape drive, which is also referred to as the *3592 model E07*, is the latest and fourth tape drive generation of the IBM 3592 tape family. The TS1140 tape drive provides higher levels of performance, reliability, and cartridge capacity than the TS1130 tape drive model E06.

Figure 8-6 shows the TS1140.

Figure 8-6 TS1140 tape drive

The TS1140 tape drive supports the TS3500 and TS4500 tape libraries in addition to IBM racks that enable stand-alone installation. The TS1140 features three options for type C media. The IBM 3592 Advanced data tape cartridge, JC, provides up to 4 TB native capacity, and up to 4 TB are provided by the IBM 3592 Advanced WORM cartridge, JY. With a limited capacity of up to 500 GB, the IBM 3592 Economy cartridge, JK, offers fast access to data. The TS1140 can also read and write on earlier type B (JB and JX) media, and read only on type A media (JA, JW, JJ, and JR).

To help optimize drive utilization and reduce infrastructure requirements, the TS1140 tape drive can be shared among supported open-system hosts on a SAN or between IBM FICON mainframe hosts when they are attached to the IBM tape controller model C07 for IBM z Systems.

The TS1140 tape drive maintains the same features and technology enhancements that are introduced with the TS1120 tape drive and extended by the TS1130 tape drive. In addition, the TS1140 tape drive offers several enhancements over the predecessor models.

Media reuse of the TS1140 tape drive

The following features help you reuse media with the TS1140 tape drive:

- Support high-performance, high-capacity tape processing for storage consolidation
- Help enhance information security with support for encryption and key management
- Improve information retention with support for existing IBM tape automation
- Support WORM cartridges to help satisfy compliance requirements

Key features of the TS1140 tape drive

Key features of the TS1140 tape drive are described:

- High performance with a data transfer rate of up to 650 MBps with compression.
- Flexible media, including short and long length cartridges, rewritable and WORM formats, and media partitioning.
- Small form factor and high capacity.
- Support for encryption and key management.
- Compatibility with existing IBM tape automation installations.
- Support for IBM Power Systems, IBM System i®, IBM System p, IBM z Systems, and IBM System x.
- Support for IBM LTFS helps reduce the cost of long-term data retention by providing access to tape storage without the need for proprietary applications. Fast access to archived data with drag-and-drop simplicity. For more information, see 2.3.4, “IBM Spectrum Archive” on page 42.

Hardware summary of the TS1140 tape drive

The hardware is summarized for the TS1140 tape drive:

- Linear serpentine recording
- 4 TB capacity with JC/JY media
- 1.6 TB capacity with JB/JX media
- 500 GB capacity with JK media
- 800 MBps burst data rate
- Compact 3.8 inch x 7.8 inch x 18.4 inch dimensions

For more information, see the following website:

`http://www.ibm.com/systems/storage/tape/ts1140/index.html`

IBM TS1150 tape drive overview

Reduce the costs of long-term data protection in big data, cloud, mobile, and social environments. In today’s data-hungry world, cost-efficient storage is more important than ever. Business analytics increase the value of data that is stored for long periods of time, while user demands for data accessibility keep escalating because of the growing use of cloud, mobile, and social media technologies. Yet storing massive volumes of static, rarely accessed data on expensive storage media can put a huge strain on shrinking IT budgets.

The TS1150 tape drive gives organizations an easy way to deliver fast access to data, improve security, and provide long-term retention for less expense than disk solutions. The TS1150 tape drive offers high-performance, flexible data storage with support for data encryption. This fifth generation drive can help protect investments in tape automation by

offering compatibility with existing automation. In addition, an upgrade model is available for existing TS1140 tape drives. The TS1150 tape drive supports the LTFS format for direct, intuitive, and graphical access to data.

Figure 8-7 shows the TS1150 tape drive.

Figure 8-7 TS1150 tape drive

The TS1150 tape drive supports the TS4500 and TS3500 tape libraries in addition to IBM racks that enable stand-alone installation. The TS1150 tape drive features three options for type D media. The IBM 3592 Advanced data tape cartridge, JD, provides up to 10 TB native capacity. Up to 10 TB are provided by the IBM 3592 Advanced WORM cartridge, JZ. A type D3592 Economy cartridge, JL, is also available that stores 2 TB of data and provides fast access to data. The TS1150 tape drive can also read and write on the previous type C media (JC, JY, and JK).

Key features of the TS1150 tape drive

The TS1150 tape drive has the following key features:

- Up to 7 TB native capacity on an IBM 3592 JC/JY cartridge.
- High capacity helps reduce the need for more tape drives and cartridges.
- Fewer necessary drives reduce the complexity of your tape infrastructure.
- Data transfer rate up to 700 MBps with compression.
- Improved performance of your tape operations to help maintain business performance.
- Media options include short and long lengths, rewritable and WORM formats, and media partitioning.
- Support for a single encryption key management approach, which can help reduce audit and compliance costs.
- Host-based encryption of data or the use of specialized encryption appliances avoided.

- IBM LTFS support. Long-term data retention costs reduction potential due to access to tape storage without the need for proprietary applications. Fast access to archived data with drag-and-drop simplicity. For more information, 2.3.4, “IBM Spectrum Archive” on page 42.
- Media partitioning to update tape partitions independently.
- Small form-factor improves space efficiency of your tape infrastructure.
- Compatibility with existing IBM tape automation installations. Protects your existing tape investments while it improves performance with next-generation technology.
- Support for IBM Power Systems, IBM System i, IBM System p, and IBM System x.

Hardware summary of the TS1150 tape drive

The TS1150 tape drive hardware includes the following features:

- Recording technique: Linear serpentine 32-channel concurrent recording.
- Native capacity (uncompressed) 10 TB (using JD/JZ media), 7 TB (using JC/JY media), or 2 TB (using JL media).
- Native sustained data rate (uncompressed) 360 MBps (type D media) or 300 MBps (type C media reformatted).
- Adaptive instantaneous data rates:
  - Twelve speeds, 112 - 365 MBps, for 3592 JD/JL/JZ cartridges that are initialized in Gen. 5 format
  - Twelve speeds, 99 - 303 MBps, for 3592 JC/JK/JY cartridges that are initialized in Gen. 5 format
  - Twelve speeds, 90 - 252 MBps for 3592 JC/JY in Gen. 4 format
- Burst data rate 800 MBps.
- High-speed search (maximum) 12.4 m/s end-to-end search rate.
- Operating environment:
  - Temperature with media 16°C – 32°C (60°F – 90°F)
  - Relative humidity 20% – 80% non-condensing (limited by media)
  - Maximum continuous operating power dissipation (drive and integrated blower): 46 watts
- Platform support:
  - IBM Power Systems (IBM AIX, IBM i, and Linux)
  - IBM System p (AIX and Linux)
  - IBM System i (IBM i and IBM OS/400®)
  - IBM System x (Open systems support)
  - Hewlett-Packard UNIX (HP-UX)
  - Oracle Solaris
  - Servers with Intel or AMD processors and SUSE Linux
  - Red Hat
  - Microsoft Windows
- Encryption support (IBM i, AIX, HP, Sun, Linux, and Windows)

For more information, see the following website:

http://www.ibm.com/systems/storage/tape/ts1150/index.html

8.1.4 IBM midrange and enterprise storage media

Dedicated tape media exist for each type of tape drive: LTO, 3592, and 3590. Each drive type supports its own media only. For example, you cannot use an LTO cartridge in a 3592 tape drive. When an IBM tape library is ordered, it is supplied with one Ultrium data cartridge or one 3592 data cartridge and one cleaning cartridge at no additional charge. Each member of the TS3*xxx* and TS4*xxx* tape library family has unique media features and rules that apply when you place the order. Make sure that the corresponding numbering scheme for bar code labels is ordered with the tape media. Use this information when you order more media.

To order media for your Ultrium, 3590, or 3592 tape drives, go to this website:

http://www.ibm.com/systems/storage/media/tape_compatibility.html

Select your geographic location and your country, and a list of authorized distributors is presented.

For a tape media compatibility matrix, see the following website:

http://www.ibm.com/systems/storage/media/tape_compatibility.html

LTO media

IBM delivers high-capacity media that provides clients with the ability to meet the demanding growth in data storage. The increase in cartridge capacity reduces the amount of equipment, space, and human intervention that are required for daily tape operations. In addition, it reduces the number of cartridges that are needed for backup and restore operations, and helps to lower operational costs throughout the enterprise.

For more information about IBM LTO media, which is referred to as *midrange tape media*, see the following website:

http://www.ibm.com/systems/storage/media/

The Ultrium cartridge is a single-reel cartridge, which means that the whole tape is wrapped around a single reel when the cartridge is not loaded in a drive. During the loading process, the threader of the drive catches the leader pin of the tape and threads it through the drive and the machine reel. During the read/write process, the tape is stored on the machine reel and the cartridge.

Ultrium tape formats

Figure 8-8 shows the IBM Ultrium cartridges, which are distinguished by color.

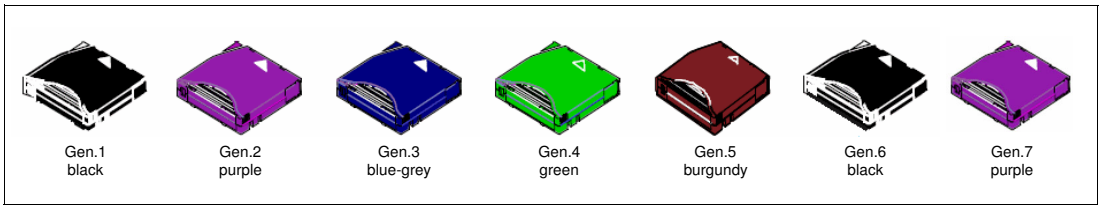

Figure 8-8 IBM LTO cartridges by generation and color

The IBM LTO Ultrium 7 cartridge is purple with a silk screen label on top that specifies "Ultrium 7- 60000 GB." The IBM LTO Ultrium 6 cartridge is black with a silk screen label on top

that specifies “Ultrium 6 - 2500 GB.” The IBM LTO Ultrium 5 cartridge is burgundy with a silk screen label on the top that specifies “Ultrium 5 - 1500 GB.” The IBM LTO Ultrium 4 cartridge is green with a silk screen label on the top that specifies “Ultrium 4 - 800 GB.” The IBM LTO Ultrium 3 cartridge is blue-gray. The IBM LTO Ultrium 2 cartridge is purple, and the LTO Ultrium 1 cartridge is black.

WORM data cartridges are two tones to distinguish them from other data cartridges. Each WORM cartridge is the color as the same generation of data cartridge on the top, but it is gray on the bottom. All generations of cartridges contain 1/2-inch, dual-coat, metal-particle tape.

LTO colors: The colors of LTO Ultrium cartridge shells are standardized. Nevertheless, certain manufacturers do not completely follow that convention. To prevent a confusing assortment of conflicting colors, use only vendors who use the same scheme.

TS1100 media (3592 Advanced Data Cartridge)

IBM 3592 tape cartridges are available in three lengths (limited capacity, standard capacity, and extended capacity) and in either rewritable or WORM formats. The cartridges can be used in IBM TS1150, IBM TS1140, IBM TS1130, IBM TS1120, and 3592 J1A tape drives.

With enterprise tape libraries that hold thousands of cartridges, the cost of keeping the technology current can be significantly impacted by the cost of upgrading the tape media. IBM 3592 tape cartridges support media reuse by enabling the drive to reformat and upgrade prior generation media cartridges. That means existing media cartridges achieve both improved performance and capacity, which help you mitigate the impact of a tape subsystem upgrade.

Highlights

The following features are included (as shown in Figure 8-9):

- Designed for IBM TS1150, IBM TS1140, IBM TS1130, and IBM TS1120 tape drives and IBM enterprise tape drive 3592
- Available in rewritable and WORM cartridge models
- Available in three formats: Fast access, standard capacity, and extended capacity
- Memory chip tracks cartridge and tape drive usage
- JD/JZ cartridges provide best-in-class performance of up to 360 MBps
- JL and JK cartridges support fast access to data
- JD/JZ and JC/JY cartridges support high capacity requirements

Figure 8-9 IBM 3592 tape media

For more information about the IBM 3592 media, see this website:

http://www.ibm.com/systems/storage/media/3592/index.html

8.2 IBM tape automation products

IBM offers tape storage, backup, and recovery solutions for data protection and retention.

IBM invented the concept of magnetic tape storage in 1952. Since then, IBM has delivered many innovations in tape storage and that innovation continues today.

This section describes the IBM midrange and enterprise tape automation products, which provide the highest levels of performance and reliability of any IBM tape subsystem. All IBM Tape automation products are designed for autonomic operations in backup / recovery and archiving and retrieval environments.

The following tape libraries are covered:

- IBM TS2900
- IBM TS3100
- IBM TS3200
- IBM TS3310
- IBM TS3500
- IBM TS4500

8.2.1 Important IBM tape library features

With the continued development of IBM tape products, device drivers, and management tools, IBM is introducing various new features and components. The following order shows their importance:

- Advanced Library Management System (ALMS)
- High Density (HD)
- Multipath architecture
- Control Path Failover and Data Path Failover
- High availability features
- Tape System Reporter (TSR)

Advanced Library Management System

The ALMS provides enhanced flexibility and capabilities for partitioning the TS3500 and TS4500 tape libraries. ALMS virtualizes the Small Computer System Interface (SCSI) element addresses while also maintaining the approach of the multipath architecture and the use of the SCSI-3 Medium Changer commands. Without ALMS, everything is based on the SCSI element address (location-centric), and partitioning is based on real cartridge slots and drive slots. With ALMS, no affinity exists between a real slot address and a SCSI element address that is reported to the server and used by the server. ALMS provides an affinity with the volume serial number (VOLSER) on the bar code label of the cartridge.

ALMS eliminates downtime when capacity on demand storage is added, logical libraries are added or removed, or logical library storage allocation is altered. ALMS also reduces downtime when adding expansion frames, adding or removing tape drives, or changing logical drive allocation.

ALMS offers the following new capabilities on the TS3500 and TS4500 tape libraries:

- Dynamic partitioning:
  - Storage slot pooling
  - Flexible drive assignment
- Adding and removing physical storage capacity transparently to any host application

- Configuring drives or L*xx* storage capacity without taking the library offline
- Virtualizing I/O slots to automatically manage the movement of cartridges between I/O station slots and storage slots

The TS3500 and TS4500 tape libraries comply with the SCSI Medium Changer standard regardless of whether ALMS is enabled. When it is enabled, ALMS is transparent to the application. The SCSI Medium Changer can be thought of as a “location-centric” interface.

The application that controls a SCSI Medium Changer device specifies a source and destination location for each request to move a cartridge. The traditional SCSI library does not control the cartridge locations. Instead, the SCSI library acts on behalf of the server.

For more information, see *IBM Tape Library Guide for Open Systems*, SG24-5946 or *IBM TS4500 R2 Tape Library Guide*, SG24-8235.

IBM High Density feature

The IBM HD feature of specific IBM tape libraries and expansion frames is IBM patented technology that allows multiple cartridges to be stored in a tiered architecture. The depth of a cartridge location in an HD slot is known as a *tier*. HD slots are designed to contain multiple cartridges in up to five tiers, depending on the type and model of the tape library or expansion module.

Each column has a spring-loaded mechanism that pushes a tape cartridge into Tier 1 when it is the only cartridge in that column. A storage element address is assigned to each cartridge at the time that the cartridge is inserted into the library.

The IBM HD feature significantly reduces the overall tape library footprint within the data centers and increases the total capacity of the tape solution for a reasonable price.

Control Path Failover and Data Path Failover

Control Path Failover configures multiple physical control paths to the same logical library within the device driver and provides automatic failover to an alternate control path when a permanent error occurs on one path. It is transparent to the running application. Control Path Failover is available for AIX, Linux, Solaris, Hewlett-Packard UNIX (HP-UX), and Microsoft Windows hosts. For example, consider a simple multipath architecture connection that consists of two host bus adapters (HBAs) (in a host) that are connected to a library with two or more drives. The control ports are enabled in two drives. The two HBAs connect to the first and second control port drives.

This simple configuration provides two physical control paths to the library for redundancy if one path from an HBA to the library fails. When the server starts, each HBA detects a control port to the library, and two medium changer devices (in AIX, smc0 and smc1) are configured. Each logical device is a physical path to the same library. However, an application can open and use only one logical device at a time, either smc0 or smc1.

When the alternative pathing support is enabled on both smc0 and smc1, the device driver configures them internally as a single device with multiple paths. The application can still open and use only one logical device at a time (either smc0 or smc1). If an application opens smc0 and a permanent path error occurs, the current operation continues on the alternative path without interrupting the application.

Requirements: Control Path Failover can work only with IBM tape libraries that contain at least two tape drives. The license key needs to be ordered separately and the appropriate host device driver is crucial.

Data Path Failover provides a failover mechanism in the IBM device driver so that you can configure multiple redundant paths in a storage area network (SAN) environment. If a path or component fails, the failover mechanism is designed to provide automatic error recovery to retry the current operation using an alternative, preconfigured path without stopping the current job. This function improves flexibility in SAN configuration, availability, and management.

When you access a tape drive device that is configured with alternative pathing across multiple host ports, the IBM device driver automatically selects a path through the HBA that has the fewest open tape devices and assigns that path to the application. This autonomic, self-optimizing capability is called *load balancing*.

The dynamic load balancing support is designed to optimize resources for devices that have physical connections to multiple HBAs in the same machine. The device driver is designed to dynamically track the usage of each HBA as applications open and close devices, and balance the number of applications that use each HBA in the machine. It can help optimize HBA resources and improve overall performance. Furthermore, the Data Path Failover provides autonomic self-healing capabilities that are similar to Control Path Failover, with transparent failover to an alternate data path in a failure in the primary host-side path.

Data Path Failover is enabled automatically when Control Path Failover license keys are installed within the supported IBM tape libraries and the appropriate device drivers are correctly configured on the host system.

Tape System Reporter

The IBM Tape System Reporter application is a monitoring server that is based on Java with an optional Windows-based graphical user interface (GUI) so that you can monitor and gather data for multiple libraries. You can generate general and specific data reports for the multiple tape cartridges, tape drives, and frames that you are monitoring.

The IBM Tape System Reporter application enables operators and administrators of the TS3500 and TS4500 tape libraries to monitor and report on storage devices from any location in an enterprise environment. This application communicates directly with your library to collect and store pertinent data so that you can generate and view performance trends. The IBM Tape System Reporter application is bundled with your ALMS purchase.

Data is available from IBM 3592 tape drives (models J1A, E05, E06, EU6, and E07) and from LTO Ultrium 2, Ultrium 3, Ultrium 4, Ultrium 5, and Ultrium 6 tape drives.

For more information, see *IBM Tape Library Guide for Open Systems*, SG24-5946 or *IBM TS4500 R2 Tape Library Guide*, SG24-8235.

8.2.2 Overview of available tape libraries

IBM provides a complete spectrum of tape libraries, highlighting high performance and capacity for entry, midrange, and enterprise systems. The IBM tape libraries handle backups, save and restore, and archival data storage needs with minimal or even no manual intervention of onsite personnel.

8.2.3 IBM TS2900 tape autoloader

The TS2900 tape autoloader is a single-drive, low profile automated tape solution and the first entry automation solution for small to midmarket tape environments. The TS2900 tape

autoloader is an external 1U stand-alone or rack-mountable unit that incorporates a single Ultrium Half High serial-attached SCSI (SAS) tape drive and nine tape cartridge capacity.

The TS2900 tape autoloader (Machine Type 3572) provides compact high capacity, low-cost solutions for simple unattended data backup. The library has a compact 1U form factor with easy access to tape cartridges by a removable magazine. The TS2900 tape autoloader is an external stand-alone or rack-mountable unit that incorporates an IBM Ultrium 4, Ultrium 5, Ultrium 6, or Ultrium 7 Half-High tape drive. It is equipped with a SAS host adapter attachment that has a data transfer rate of up to 6 Gbps. Figure 8-10 shows the compact TS2900 tape autoloader.

Figure 8-10 IBM TS2900 tape autoloader

The TS 2900 tape autoloader has a removable cartridge magazine that provides nine data cartridge slots, including a configurable two slot I/O station. The TS 2900 tape autoloader is an entry point for IBM LTO tape automation. This autoloader uses the IBM patented HD slot technology.

These cartridges are supported in the TS2900:

- IBM LTO tape cartridge Ultrium 7 (6 TB native physical capacity)
- IBM LTO tape cartridge Ultrium 6 (2.5 TB native physical capacity)
- IBM LTO tape cartridge Ultrium 5 (1.5 TB native physical capacity)
- IBM LTO tape cartridge Ultrium 4 (800 GB native physical capacity)
- Write Once Read Many (WORM) cartridges

Cartridge support depends on the tape drive that is installed in the TS2900 tape autoloader.

The library media capacity can be further increased by using hardware compression (a 2:1 compression factor for Ultrium 5 and lower tape drives, and a 2.5:1 compression factor for Ultrium 6 and Ultrium 7 tape drives).

The TS2900 tape autoloader has a capacity of a maximum of nine tape cartridges, which provides a media capacity of up to 54 TB (135 TB with a 2.5:1 compression) data storage.

The TS 2900 tape autoloader has a 6 Gbps single-port small form factor (SFF)-8088 SAS connector. Designed for tape automation, the TS2900 tape autoloader can be attached to Open Systems Servers only. To determine the latest supported servers, see the IBM System Storage Interoperation Center (SSIC) website:

http://www.ibm.com/systems/support/storage/config/ssic/

Table 8-3 summarizes the capabilities of the IBM TS2900 tape library.

Table 8-3 TS2900 tape autoloader

Model	TS2900 tape autoloader
Drive options	IBM LTO Ultrium 7 Half High: 6 Gbps SAS IBM LTO Ultrium 6 Half High: 6 Gbps SAS IBM LTO Ultrium 5 Half High: 6 Gbps SAS IBM LTO Ultrium 4 Half High: 6 Gbps SAS
Optional feature codes	Desk-side cover kit (FC 7010) Additional tape magazine (FC 8111) Rack-mount kit (FC 7006)
Tape drive type	IBM LTO Ultrium 7 Half Height IBM LTO Ultrium 6 Half Height IBM LTO Ultrium 5 Half Height IBM LTO Ultrium 4 Half Height
Number of drives	One
Number of tape cartridges	Nine
Number of mail slots	One
Physical capacity	IBM LTO Ultrium 7: Up to 54 TB (135 TB with 2.5:1 compression) IBM LTO Ultrium 6: Up to 22.5 TB (46.25 TB with 2.5:1 compression) IBM LTO Ultrium 5: Up to 13.5 TB (27 TB with 2:1 compression) IBM LTO Ultrium 4: Up to 7.2 TB (14.4 TB with 2:1 compression)
Data transfer rate	IBM LTO Ultrium 7: Up to 300 MBps IBM LTO Ultrium 6: Up to 160 MBps IBM LTO Ultrium 5: Up to 140 MBps IBM LTO Ultrium 4: Up to 120 MBps
Attachment and systems support	The TS2900 supports 6 Gbps attachment available for selected IBM System servers and other Linux and Microsoft Windows open-system servers.

For more information about the TS2900 tape autoloader, see the*IBM Tape Library Guide for Open Systems*, SG24-5946.

For specific details, configuration information, and operations help, see this website:

`http://www.ibm.com/systems/storage/tape/ts2900/index.html`

8.2.4 IBM TS3100 tape library

The TS3100 tape library (Machine type 3573) handles backup, save and restore, and archival data-storage needs for small to medium-size environments. The TS3100 tape library is an external 4U stand-alone or rack-mountable unit. It incorporates a single Ultrium Full High or two Ultrium Half High tape drives, and has a total capacity of 24 tape cartridges. It is a compact, high-capacity, low-cost solution for simple, unattended data backup and archive.

The TS3100 tape library is configured with two removable cartridge magazines: One cartridge magazine on the left side (12 data cartridge slots) and one cartridge magazine on the right side (12 data cartridge slots). Additionally, the left magazine includes a single I/O station slot to help support continuous library operation while it imports and exports media. The tape library can be configured in sequential or random access mode and has a bar code reader installed by default. The operator front control panel includes the power button, front

panel LEDs, control keys, and the operator control panel display. To manage the library, an easy-to-use web interface is available.

Figure 8-11 shows the physical TS3100 library.

Figure 8-11 TS3100 library

The TS3100 tape library is a driveless product that is adaptable to your business and backup solution needs. Table 8-4 lists the available configurations.

Table 8-4 TS3100 in detail

Model	**TS3100**
Drive options	► Single Ultrium 7 Full Height tape drive - 8 Gbps FC ► Up to two Ultrium 7 Half Height tape drives - 6 Gbps SAS or 8 Gbps FC ► Single Ultrium 6 Full Height tape drive - 8 Gbps FC ► Up to two Ultrium 6 Half Height tape drives - 6 Gbps SAS or 8 Gbps FC ► Single Ultrium 5 Full Height tape drive - 6 Gbps SAS or 8 Gbps FC ► Up to two Ultrium 5 Half Height tape drives - 6 Gbps SAS or 8 Gbps FC ► Single Ultrium 4 Full Height tape drive - LVD SCSI, 3 Gbps SAS, or 4 Gbps FC ► Up to two Ultrium 4 Half Height tape drives - 3 Gbps SAS or 4 Gbps FC
Optional feature codes	► Path Failover feature FC 1682 ► Transparent LTO encryption FC 5900 ► Rack mount FC 7002 ► Right-side magazine set FC 8106 ► Left-side magazine FC 8109
Tape drive type	► LTO Ultrium 7 half-height and full-height ► LTO Ultrium 6 half-height and full-height ► LTO Ultrium 5 half-height and full-height ► LTO Ultrium 4 half-height and full-height
Number of drives	1 - 2
Number of tape cartridges	Up to 24
Number of mail slots	1
Physical capacity	IBM LTO Ultrium 7: 144 TB native (up to 360 TB compressed 2.5:1) IBM LTO Ultrium 6: 60 TB native (up to 150 TB compressed 2.5:1) IBM LTO Ultrium 5: 36 TB native (up to 72 TB compressed 2:1) IBM LTO Ultrium 4: 19.2 TB native (up to 38.4 TB compressed 2:1)
Data transfer rate / drive	IBM LTO Ultrium 7: Up to 300 MBps IBM LTO Ultrium 6: Up to 160 MBps IBM LTO Ultrium 5: Up to 140 MBps IBM LTO Ultrium 4: Up to 120 MBps

Model	TS3100
Attachment and systems support	8 Gbps Fibre Channel 6 Gbps SAS interfaces LVD SCSI, 4 Gbps Fibre Channel, and 3 Gbps SAS interfaces (LTO Ultrium 4 full-height) Native device driver support is available for selected IBM System servers and other Linux and Microsoft Windows open-system servers.

Avoid a combination of two Ultrium Half Height tapes drives of different generations because it might cause unlikely complications during their configuration on host systems and in the backup applications. Also, avoid mixing the Ultrium media generations within the same TS3100 tape library.

The library supports the IBM Spectrum Archive Library Edition featuring IBM Linear Tape File System technology if LTO Ultrium 5 (or higher) drives are used. For more information about Spectrum Archive and LTFS, see 2.3.4, “IBM Spectrum Archive” on page 42.

For up-to-date and detailed operating system and attachment requirements, see the following website:

http://www.ibm.com/systems/support/storage/config/ssic

For more information about the TS3100 tape library, see the *IBM Tape Library Guide for Open Systems*, SG24-5946.

For more information, see the following website:

http://www.ibm.com/systems/storage/tape/ts3100/index.html

8.2.5 IBM TS3200 tape library

The TS3200 tape library is designed for backup, save and restore, and archival data-storage needs for small to medium-size environments. The TS3200 tape library is an external 4U stand-alone or rack-mountable unit that incorporates up to two Ultrium Full High or four Ultrium Half High tape drives and 48 tape cartridges, in addition to three I/O slots.

Like the TS3100 tape library, the TS3200 tape library (Figure 8-12) provides compact, high-capacity, low-cost solutions for simple, unattended data backup and archive. It offers high capacity and performance technology for midrange environments. The TS3200 tape library is an external 4U stand-alone or rack-mountable unit that incorporates up to two Full Height LTO IBM Ultrium 7, Ultrium 6, Ultrium 5, or Ultrium 4 tape drives or up to four Half-Height tape drives of the Ultrium 7, Ultrium 6, Ultrium 5, and Ultrium 4 generations.

Figure 8-12 TS3200 Library

The TS3200 tape library is an excellent tape storage solution for organizations with existing digital linear tape or organizations that require high-performance automated tape backup. The

TS3200 tape library is also designed for organizations with limited physical space in their IT environments.

Operating in a rack environment allows organizations to take the advantage of placing the TS3200 tape library in a standard 19-inch rack. This configuration provides up to 288 TB of compressed tape storage in only a 4U space. Remote management and a bar code reader are standard components of the library, and it can run in sequential or random access mode. Optional available features are a rack mount kit, an extra power supply, Control Path Failover, and Data Path Failover.

The TS3200 tape library has IBM patented Multipath Architecture for sharing the library robotics. It allows a library with at least two drives to be partitioned into two logical libraries, for sharing between servers and applications.

The TS3200 tape library (model 4UL) has the convenience of ordering it driveless, which is adaptable to clients' business and backup solution needs. Table 8-5 lists configurations that are available for purchase.

Table 8-5 TS3200 in detail

Model	**TS3100**
Drive options	► Up to two Ultrium 7 Full Height tape drive - 8 Gbps FC ► Up to four Ultrium 7 Half Height tape drives - 6 Gbps SAS or 8 Gbps FC ► Up to two Ultrium 6 Full Height tape drive - 8 Gbps FC ► Up to four Ultrium 6 Half Height tape drives - 6 Gbps SAS or 8 Gbps FC ► Up to two Ultrium 5 Full Height tape drive - 6 Gbps SAS or 8 Gbps FC ► Up to four Ultrium 5 Half Height tape drives - 6 Gbps SAS or 8 Gbps FC ► Up to two Ultrium 4 Full Height tape drive - LVD SCSI, 3 Gbps SAS or 4 Gbps FC ► Up to four Ultrium 4 Half Height tape drives - 3 Gbps SAS or 4 Gbps FC
Optional feature codes	► Path Failover feature FC 1682 ► Transparent LTO encryption FC 5900 ► Rack mount FC 7002 ► Right-side magazine set FC 8106 ► Left-side magazine, upper FC 8107 ► Left-side magazine, lower FC 8108
Tape drive type	► LTO Ultrium 7 half-height and full-height ► LTO Ultrium 6 half-height and full-height ► LTO Ultrium 5 half-height and full-height ► LTO Ultrium 4 half-height and full-height
Number of drives	1 - 4
Number of tape cartridges	Up to 48
Number of mail slots	3
Physical capacity	IBM LTO Ultrium 7: 288 TB native (up to 720 TB compressed 2.5:1) IBM LTO Ultrium 6: 120 TB native (up to 300 TB compressed 2.5:1) IBM LTO Ultrium 5: 72 TB native (up to 144 TB compressed 2:1) IBM LTO Ultrium 4: 38.4 TB native (up to 76.8TB compressed 2:1)

Model	TS3100
Data transfer rate	IBM LTO Ultrium 7: Up to 300 MBps IBM LTO Ultrium 6: Up to 160 MBps IBM LTO Ultrium 5: Up to 140 MBps IBM LTO Ultrium 4: Up to 120 MBps
Attachment and systems support	8 Gbps Fibre Channel 6 Gbps SAS interfaces LVD SCSI, 4 Gbps Fibre Channel, and 3 Gbps SAS interfaces (LTO Ultrium 4 full-height) Native device driver support is available for selected IBM System servers and other Linux and Microsoft Windows open-system servers.

The TS3200 library supports the IBM Spectrum Archive Library Edition featuring IBM Linear Tape File System technology if LTO Ultrium 5 (or higher) drives are used. For more information about Spectrum Archive and LTFS, see 2.3.4, “IBM Spectrum Archive” on page 42.

For up-to-date and detailed operating system and attachment requirements, see the following website:

http://www.ibm.com/systems/support/storage/config/ssic

For more information about the TS3200 tape library, see the *IBM Tape Library Guide for Open Systems*, SG24-5946.

For more information, see the following website:

http://www.ibm.com/systems/storage/tape/ts3200/index.html

8.2.6 IBM TS3310 tape library

The TS3310 tape library is a highly expandable Ultrium LTO library. You can start small with a 5U base unit (model L5B), which is available in a desktop or rack-mounted configuration single five Electronic Industry Alliance (EIA) rack unit. The base unit contains the library control module, a fixed tape cartridge storage of 35 slots, a configurable I/O station of six slots, a touchscreen display, and cartridge handling robotics. The base unit also includes up to two LTO Ultrium 7, Ultrium 6, Ultrium 5, or Ultrium 4 tape drives.

Over time, as your need for tape backup expands, you can add another 9U (model E9U) expansion module. Each E9U expansion module can accommodate up to four LTO Ultrium 7, Ultrium 6, Ultrium 5, or Ultrium 4 tape drives; additional storage slots; a configurable I/O station of 12 slots; and redundant power supply. The entire system grows vertically. Currently, available configurations include the 5U base library module alone or with up to four 9U modules.

The TS3310 tape library offers a broad range of configuration possibilities. The smallest configuration includes a base unit with one to two LTO Ultrium 7 trough LTO Ultrium 4 tape drives, up to 246 TB (LTO Ultrium 7) of native tape storage (35 slots), and six I/O slots. This configuration is upgradeable to a fully configured rack-mounted library 41U high with up to 18 LTO Ultrium 7 through LTO Ultrium 4 tape drives with over 2.45 PB of native tape storage (LTO Ultrium 7) and up to 54 I/O slots. If two or more expansion modules are installed, the whole library must be rack-mounted.

The IBM Ultrium LTO tape drives that are available in the TS3310 tape library configuration are encryption capable. They support application-managed encryption (AME), system-managed encryption (SME), and library-managed encryption (LME).

Figure 8-13 shows the TS3310 tape library base L5B unit with one expansion module. On the right side of each module, you can see the I/O station door (six slots in the base L5B or 12 slots in the E9U).

Figure 8-13 TS3310 Library base module L5B plus one expansion module E9U

For organizations that are unsure of their short-term or long-term tape capacity needs, the TS3310 tape library Capacity on Demand (CoD) capability allows the system to scale as your needs grow. In the initial configuration, an E9U has half of its storage cells enabled. As your business grows, you can purchase a CoD key to enable the second half of the model E9U storage cells.

This library supports LTO Ultrium 7, LTO Ultrium 6, or LTO Ultrium 5 native switched fabric FC attachment and LTO Ultrium 4 tape drives with either SAS or native switched fabric FC attachment for connection to a wide spectrum of open system servers.

The library supports the IBM Spectrum Archive Library Edition featuring IBM LTFS technology if LTO Ultrium 5 (or higher) drives are used. For more information, see 2.3.4, “IBM Spectrum Archive” on page 42.

Figure 8-14 shows the possible expansion scenarios. Each expansion model can host drives, cartridge slots, or a combination of both.

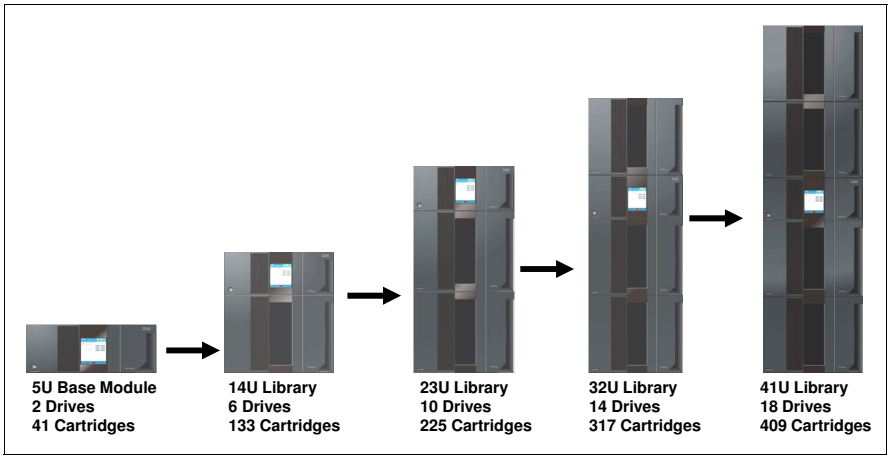

Figure 8-14 TS3310 expansion options

The specifications for the TS3310 tape library are listed in Table 8-6.

Table 8-6 TS3310 tape library specifications

Type and model	**3576 model L5B**	**3576 model L5B and four E9U models**
Configuration	Base library	Base library and four expansion units
LTO storage slots (maximum)	41	409
LTO input/output slots (minimum/maximum)	6	54
Maximum tape drives	2	18
Total physical capacity 2.5:1 compression	615 TB	6.13 PB
Capacity on Demand increments	n/a	46 cartridges per license key
Maximum logical libraries	2	18
Hot-swap components	Power supplies and tape drives	
Systems management	Data encryption and SMI-S enabled	
Operating systems and servers supported	OS/400, V5R2, V5R3, or later IBM AIX 5L™ V5.1, V5.2, V5.3, V6.1, or later Sun Solaris 8 and 9 Microsoft Windows 2003 all editions Microsoft Windows 2008 all editions HP-UX 11.0, HP-UX 11.11, and HP-UX 11.231 Linux distributions: Red Hat Enterprise Linux version 4 and SUSE Linux Enterprise Server 9	

For the latest features, functions, and specifications about the TS3310, see this website:

http://www.ibm.com/systems/storage/tape/ts3310/index.html

For more technical information, library setup, and configuration guidance, see this website:

http://www.ibm.com/support/docview.wss?uid=ssg1S7003238

For supported systems and platforms, see the SSIC website:

http://www.ibm.com/systems/support/storage/config/ssic

Also, see *IBM Tape Library Guide for Open Systems*, SG24-5946.

8.2.7 IBM TS3500 tape library

The TS3500 tape library combines IBM automation and drive technology to provide a highly scalable, automated tape library for z Systems, and open systems backup and archive in midrange to enterprise environments.

With the TS3500 tape library, you can use both the LTO and Enterprise 3592 drive technologies within the same library. The TS3500 tape library was previously known and marketed as the IBM 3584 tape library, machine type 3584.

The TS3500 tape library provides tape storage solutions for the large, unattended storage requirements for today's midrange and enterprise (z/OS and open systems) environments. Information is provided about the TS3500 tape library attachment in an open systems environment only.

For information about the TS3500 tape library attachment to a z Systems environment, see *IBM TS3500 Tape Library with System z Attachment A Practical Guide to Enterprise Tape Drives and TS3500 Tape Automation*, SG24-6789.

Combining reliable, automated tape handling and storage with reliable, high-performance IBM LTO Ultrium tape and 3592 drives, the TS3500 tape library offers outstanding retrieval performance with typical cartridge move times of less than three seconds.

The TS3500 tape library can be partitioned into multiple logical libraries. This feature makes it an excellent choice for consolidating tape workloads from multiple heterogeneous open systems servers.

The TS3500 tape library provides outstanding reliability and redundancy through the provision of redundant power supplies in each frame, an optional second cartridge accessor, Control Path Failover and Data Path Failover, dual grippers within each cartridge accessor, and the Shuttle Complex system to connect parallel TS3500 tape library strings. Both library and drive firmware can now be upgraded nondisruptively, that is, without interrupting the normal operations of the library.

In physically constrained data center environments, the option to connect multiple library strings to create a *shuttle complex* (see "TS3500 tape library shuttle complex (model SC1)" on page 206) greatly increases opportunities for growth, in addition to the maximum cartridge capacity without the need for more tape drives. With the addition of high-density frames, it also greatly increases capacity without requiring more floor space.

The TS3500 tape library supports tape encryption on the following tape drives:

- TS1070 (3588-F7A) tape drive
- TS1060 (3588-F6A) tape drive
- TS1050 (3588-F5A) tape drive
- TS1040 (3588-F4A) tape drive
- IBM 3592 model E*xx* tape drives

The IBM 3592 model E*xx* tape drives include the TS1140 tape drive, TS1130 tape drive, and TS1120 tape drive. All three encryption methods are supported: AME, SME, and LME.

The TS3500 tape library model HA1 allows two robotic accessories to operate simultaneously in two 16-frame configurations. Figure 8-15 shows the minimum and maximum 16 frame TS3500 tape library string configurations.

Figure 8-15 TS3500 tape library minimum and maximum configurations

The libraries are built of the following modules:

- Every library requires a base frame (model L*xx*) to which optional expansion frames (model D*xx or* model *Sxx*) can be added. Only one base frame is permitted in each library configuration.
- Base and expansion frames that support either of the following types:
  - LTO 1, 2, 3, 4, 5, 6 and 7 tape drives (model *x*53, x52, and x32 frames)
  - IBM 3592 tape drive models E*xx* and IBM 3592 tape drive model J1A (model *x*23 and x22 frames)
- An optional second accessor is available through the addition of the model HA1 frames.

Tip: The L32, D32, L22, D22, L52, and D52 frames were withdrawn from marketing, but they are mentioned for reference because they can be upgraded to L53 or D53 models. New L and D frame models can be attached to the withdrawn models.

Model L*xx* and D*xx* frames can be intermixed with each other and installed frame models. The minimum requirement is one base frame in each library. The S24 can be attached to the tape library with installed 3592 drives. The S54 can be attached to the tape library with installed LTO drives.

IBM HD frames

The TS3500 tape library storage-only frame includes the model S24 frame and the model S54 frame. The TS3500 tape library model S24 and S54 frames are high-density storage-only expansion frames that are compatible with existing TS3500 tape libraries and frames. Each model S24 frame supports up to 1,000 IBM 3592 cartridge slots. Each model S54 frame supports up to 1,320 LTO cartridge slots.

For the use of HD frames, the ALMS feature must be enabled. For more information about ALMS, see “Advanced Library Management System” on page 191.

Virtual I/O slots

When ALMS is enabled, you can enable virtual I/O slots in the TS3500 tape library so that the library automatically queues all cartridge moves between the I/O station and the storage slots. In this way, the library enhances its import and export capabilities. This capability makes the process of adding and removing cartridges much easier and faster.

For more information about ALMS and virtual I/O slots, see *IBM System Storage TS3500 Tape Library Introduction and Planning Guide IBM 3584 Tape Library* at this website:

http://www.ibm.com/support/docview.wss?uid=ssg1S7001738

TS3000 System Console

The TS3000 System Console (TSSC) is a service tool that provides many service and support functions from a centralized location. These functions include error-initiated problem reporting and remote services, including remote monitoring support and wellness checking, that are beneficial to both the user and the IBM service support representative (SSR).

TSSC offers these main benefits:

- Remote support enablement
- Reduced telephone line charges
- Faster data offload by broadband Call Home through Electronic Customer Care (ECC)
- Reduced repair costs
- Improved serviceability
- Better proactive maintenance

Tape System Library Manager (TSLM)

The Tape System library Manager (TSLM) software provides consolidation and simplification benefits in an open systems TS3500 tape library environment.

TSLM provides a resource management layer between applications, such as Tivoli Storage Manager, and the tape library hardware. Essentially, TSLM decouples tape resources from applications. This decoupling simplifies both the aggregation and the sharing of tape resources.

TSLM can combine the capacity of multiple TS3500 tape libraries into a single reservoir of tape storage that can be managed from a single point. This single point of control offers more effective management, monitoring, and reporting of the use of tape storage resources for new and existing TS3500 tape library systems.

TSLM is software that meets today’s tape storage management challenges for a wide variety of clients who are involved in high-performance computing (HPC) environments, such as oil and gas exploration and genomic analysis. In HPC environments, dozens of tape libraries and hundreds of tape drives must be shared and managed to back up and archive petabytes of data at the lowest cost possible. The TS3500 systems with a shuttle connector, TSLM, and Tivoli Storage Manager are the ideal solutions for these clients.

IBM TS3500 tape library specifications

The specifications for TS3500 tape library are listed in Table 8-7.

Table 8-7 TS3500 tape library specifications

Model	IBM TS3500 tape library
Frame definition	L23 – base frame for TS1150, TS1140, TS1130, TS1120, or IBM 3592 tape drive; D23 – drive-capable expansion frame for TS1150, TS1140, TS1130, TS1120, or IBM 3592 tape drive, or S24 – storage-only expansion frame for IBM 3592 cartridges L53 – base frame for LTO, TS1070 tape drive, TS1060 tape drive, TS1050 tape drive, or TS1040 tape drive; D53 – drive-capable expansion frame for LTO, TS1070 tape drive, TS1060 tape drive, TS1050 tape drive, or TS1040 tape drive; S54 – storage-only expansion frame for LTO cartridges HA1 – High Availability service bay frame for use with the dual accessor feature
Tape drive types	TS1150, TS1140, TS1130, TS1120 or IBM 3592 tape drives, or IBM LTO Ultrium 7 and lower generations
Number of frames	One base frame, up to 15 expansion frames The TS3500 tape library model HA1 installation provides one of the two extra frames that are required as service bays in a dual accessor library.
Number of drives	Up to 12 per frame (up to 192 per library)
Number of tape cartridges	L23 – up to 260; D23 – up to 400; S24 – up to 1,000; Total supported: 15,000 L53 – up to 287; D53 – up to 440; S54 – up to 1,320; Total supported: 20,000
Number of logical libraries	Maximum of 192 (up to the number of installed drives)
Number of 3953 systems	Maximum of four per TS3500 tape library subsystem
Number of TS7700s	Maximum of eight per TS3500 tape library subsystem
Attachment and systems support	The TS3500 tape library can attach to IBM Power Systems, System p, System i, and System x servers, non-IBM servers, and workstations Attach to z Systems by using the IBM 3953 Tape System (3953 tape system) or the IBM TS7740
Operating systems support	Device driver support is available for IBM AIX; IBM OS/400; IBM i; Windows 2000; Windows Server 2003; Linux; Sun Solaris; and HP-UX

For more information about the TS3500 tape library, see the *IBM System Storage TS3500 Tape Library Introduction and Planning Guide IBM 3584 Tape Library* at the following website:

`http://www.ibm.com/support/docview.wss?uid=ssg1S7001738`

For the latest specifications, features, and functions, see this website:

`http://www.ibm.com/systems/storage/tape/ts3500/index.html`

TS3500 tape library shuttle complex (model SC1)

The TS3500 tape library shuttle complex enables extreme scalability of over 300,000 LTO cartridges (or over 2.25 EB of uncompressed TS1150 tape drive data) in a single library image by supporting the transport of cartridges from one TS3500 tape library string to another TS3500 tape library string. Application software that supports this new capability can move tape cartridges directly from their home logical library to the destination logical library. Shuttle connections span HD TS3500 tape library S24 or S54 frames from different TS3500 tape library strings. The TS3500 tape library shuttle complex supports new and existing TS3500 tape library installations and is suited for High Performance Storage System (HPSS) environments.

Tip: HPSS and TSLM are the library management solutions that support the TS3500 tape library shuttle complex.

The Shuttle Complex consists of three main components:

- TS3500 tape libraries with model SC1
- Library Management Software (HPSS or TSLM) with Shuttle Call System (SCS)
- Total Storage System Console (TSSC)

Table 8-8 lists the capacities of the supported combinations of frames, tape drives, and their capacities, including the extreme scalability of the TS3500 tape library shuttle complex.

Table 8-8 TS3500 tape library configurations, including the shuttle complex

Supported configurations			
A library models	**Drives**	**Number of cartridges[a,b]**	**Maximum native capacity[c]**
L5x, D5x, and S54	Ultrium tape drives	> 20,000	120 PB
L32, D32, and S54	Ultrium tape drives	> 20,000	120 PB
L2x, D2x, and S24	3592 tape drives	> 15,000	60 PB
L5x, D5x, S54, and SC1	Ultrium tape drives	> 300,000	1.8 EB
L32, D32, S54, and SC1	Ultrium tape drives	> 300,000	1.8 EB
L2x, D2x, S24, and SC1	3592 tape drives	> 225,000	2.25 EB

a. All node cards in the library must be *xx*3-equivalent node cards. For *xx*2 models, you must use an *xx*3 model conversion or the enhanced node card feature (FC1700 or FC1701).

b. To increase the maximum number of cartridges to more than 6,887, or to support a shuttle complex, logical libraries must use LTO Ultrium 4, 3592-E05, or later tape drives as control path drives. To support more than 6,887 cartridges, and for logical libraries with shuttle stations assigned, check the minimum code levels that are required.

c. The maximum key features native capacity figures are based on library configurations with HD frames and all Ultrium 5 or 3592-E07 tape drives. In addition, the figures for libraries with model SC1 (shuttle connection) are based on the maximum shuttle complex configuration.

Figure 8-16 illustrates the TS3500 tape library shuttle complex and shows the capability of moving tapes from one library string to another by bypassing the intermediate library strings in comparison to a traditional pass-through method. The TS3500 tape library transports tape cartridges in shuttle cars that pass *over* the libraries. This method of transporting cartridges is called *direct flight*. With the direct flight capability, if no drive is available in the home logical library, the cartridge is moved across a *shuttle connection* to a logical library with an available drive. This configuration of interconnected parallel library strings is called a *shuttle complex.*

Figure 8-16 TS3500 library in a 3 library shuttle complex

To review and understand the physical requirements for the TS3500 tape library shuttle complex, see the *IBM System Storage TS3500 Tape Library with ALMS Introduction and Planning Guide*, GA32-0593.

8.2.8 IBM TS4500 tape library

The IBM TS4500 tape library is a next-generation storage solution that is designed to help midsize and large enterprises respond to storage challenges in open system environments:

- High data volume
- Growth in data centers
- Increasing cost of data center storage footprints
- Difficulty migrating data across vendor platforms
- Increased complexity of IT training and management as staff resources shrink

The TS4500 tape library combines reliable, automated tape handling and storage with reliable, high performance in an open systems environment. Incorporating IBM LTO Ultrium tape and 3592 drives, the TS4500 tape library offers outstanding retrieval performance with typical cartridge move times of less than 3 seconds.

The TS4500 tape library can be deployed as a single frame library and upgraded to a total of 18 frames, with a combination of either LTO and 3592 frames. This configuration can be partitioned into multiple logical libraries. This feature makes it an excellent choice for consolidating tape workloads from multiple heterogeneous open systems servers.

The TS4500 library protects investment by providing for redeployment of S24 and S54 frames from the TS3500 onto the TS4500.

The TS4500, installed with the High Availability (HA) option, provides dual active accessors that provide accessor redundancy, and can double the robot performance during tape move operations.

The library provides outstanding reliability and redundancy through provision of redundant power supplies in each drive frame, control and data path failover, and dual grippers within

the cartridge accessor. Both library and drive firmware can be upgraded nondisruptively, without interrupting normal operations. Encryption is supported on the following tape drives:

- TS1070 tape drive (Model 3588 F7C)
- TS1060 tape drive (Model 3588 F6C)
- TS1050 tape drive (Model 3588 F5C)
- TS1150 tape drive (Model 3592 EH8)
- TS1140 tape drive (Model 3592 EH7)

The TS4500 supports the encryption methods AME and LME.

The following features are supported in TS4500 Release 2 and later:

- Automatic media verification
- Flexible remote authentication
- Primary control system failover
- Mixed media types within the same TS4500 library
- Scalability to 18 frames
- Up to 128 tape drives
- SNMP query configuration
- Redeployment of S24 and S54 frames from TS3500 to TS4500

The following features are supported in TS4500 Release 3 and later:

- High availability with dual active accessors
- Support for external TSSC/IMC
- Flexible growth options with new flex track design

Note: From Release 3, the TS4500 is supported on z/OS systems that have TS7700 attached. This support is provided on TS7700 subsystems that have Release 3.4 installed. To provide this support, install two integrated 16 Gb FC switches into the bottom of a L25 or D25 frame, or use an external TSSC. The TS4500 Management Interface has preset TS7700 logical library support.

Figure 8-17 shows a three-frame version of the TS4500 tape library. An individual library can consist of one L frame and up to 17 expansion frames, and can include up to 128 tape drives with more than 23,000 tape cartridges as shown in Table 8-9 on page 210.

Figure 8-17 TS4500 in a 3-frame configuration

The TS4500 tape library provides the following capabilities:

- High Availability dual active accessor option with integrated service bays to reduce inactive service space by 40%. The Elastic Capacity option can be used to completely eliminate inactive service space.
- All of the frames include HD slot technology.
- Additional HD2 frame models can be placed in any active position, so the library can grow from both the right side and the left side of the first L frame.
- HD generation 1 frames, from existing TS3500 library, can be redeployed into a TS4500. These frames must be installed to the right of the Lx5 frame, and must have FC 1742 applied to them before they can exist in a TS4500 library string.
- New Single Deep Cell technology.
- Integrated management console (IMC).
- New user interface for improved usability.
- Updated control system.
- Input/output (I/O) magazine to allow individual cartridge handling to be performed independently of the library.
- Top-rack space to house extra tape solution components within the library footprint.
- Support for HD2 compatible models of the TS1150 (3592 EH8), TS1140 (3592 EH7), LTO Ultrium 5 (3588 F5C), LTO Ultrium 6 (3588 F6C), and LTO Ultrium 7 (3588 F7C) tape drives.

The TS4500 tape library is available with several tape drives, frame models, and feature options to meet your specific needs. Additional features of the TS4500 tape library are highlighted in the following list:

- Advanced Library Management System (ALMS)
- Ability to attach multiple simultaneous heterogeneous servers
- Remote management with the TS4500 management GUI or the TS4500 command-line interface (CLI)
- Remote monitoring by using Simple Network Management Protocol (SNMP), email, or syslog
- SNMP query configuration
- Media health verification
- Multipath architecture
- Drive and media exception reporting
- In-depth reporting by using the TSR
- Host-based path failover
- Up to 288 I/O slots (36 I/O slots standard for LTO libraries and 32 I/O slots standard for 3592 libraries with extra I/O slots available as a feature add-on for all D25 and D55 frames)

8.2.9 TS4500 product description

The IBM TS4500 tape library (Machine Type 3584) is a modular tape library that consists of a high-density base frame and up to three high-density expansion frames. The frames join side by side and can grow to the left or to the right of the base frame. All frames are supported by a single cartridge accessor. You can install a single-frame base library (Figure 8-17 on

page 208) and grow it to four frames, tailoring the library to match your system capacity requirements.

Table 8-9 lists the supported combinations of frames, tape drives, and their capabilities.

Table 8-9 TS4500 tape library capabilities

Models	Drives in frames	Maximum cartridges	Maximum native capacity
L25, D25, S25, and S24.	3592 tape drives	17 550	175 PB
L55, D55, S55, and S54	LTO tape drives	23 170	139 PB
The maximum native capacity figures are based on library configurations with one base frame with all LTO-7 or TS1150 tape drives, and 17 storage-only HD frames.			

Note: Customer-printed bar code labels are not currently supported on the TS4500. Some bar code labels that are previously successfully scanned by the TS3500 are not readable by the TS4500 at general availability (GA). Issues with reading media labels are most often those labels that were printed by the customer.

Figure 8-18 shows the TS4500 with a maximum configuration of 18 frames.

Figure 8-18 Maximum TS4500 tape library configuration with 18 frames

Eight types of frames are supported in the current TS4500 tape library range. Each frame is identified by a three-character model number (L25, D25, L55, D55, S25, S55, S24, and D24) that describes the nature of the frame.

The TS4500 tape library can contain a mix of 3592 and LTO frames. The TS4500 also supports the addition of a top rack frame. The top rack model, 3584 Model TR1, provides an extra 10U of rack space on any frame in a library without requiring more floor space.

The new XIV style GUI makes configuration and administration of the library intuitive. Figure 8-19 shows the TS4500 GUI with the general library view. You can hover over the icons on the left to expand them.

Figure 8-19 TS4500 library GUI

For a detailed description of the IBM TS4500 tape library, see the *IBM TS4500 R2 Tape Library Guide*, SG24-8235.

8.3 IBM tape virtualization products

This section describes tape virtualization, data deduplication, and associated IBM tape virtualization products and solutions:

- IBM ProtecTIER products:
  - IBM TS7650G ProtecTIER Deduplication Gateway
  - IBM TS7620 ProtecTIER Deduplication Appliance Express

    The IBM TS7620 Expansion Drawer provides added repository capacities for TS7620 systems.
- IBM TS7700:
  - IBM model TS7740
  - IBM model TS7720
  - IBM model TS7720T
- DFSMSrmm (formerly known as Enterprise Removable Media Manager (eRMM))

8.3.1 Introduction to tape virtualization

A virtual tape library provides high performance backup and restore by using disk arrays and virtualization software. A virtual tape library is a unique blend of several storage tiers. The lifecycle of data from its creation at the server level migrates by backup software to a virtual tape library. There is no better place for the data to be because the virtual tape library is a combination of high performance SAN-attached disk and high performance servers that emulate a tape storage device. At this level, you have many options for the data. For example, the data can remain on the virtual tape library indefinitely if there is enough space, or it can be migrated to tape for offsite storage, archive, or both.

Virtual tape libraries fill a void in the backup infrastructure for data that needs to be restored at a certain moment. Many restore requests happen within six weeks of the data being backed up. Backup software can be configured to back up data to a virtual tape library and then create a virtual tape-to-tape copy for offsite deployment. It is no longer necessary to call the tapes back from an offsite location unless data is required from years past.

One of the biggest challenges with backup planning today is that the amount of data that is being backed up is growing, but the time that is allotted for a backup (the *backup window*) is shrinking or remaining static. Applications need to be up and operational nearly 24 hours a day. To manage the need for increased data capacity and data protection, clients must find ways to shrink their backup windows and recover as quickly as possible.

You can expect the following main benefits from tape virtualization:

- Brings efficiency to the tape operation environment
- Reduces the batch window
- Provides high availability and disaster recovery configurations
- Provides fast access to data through caching on disk
- Provides use of current tape drive, tape media, and tape automation technology
- Provides the capability of filling high capacity media to 100%
- Provides many tape drives or concurrent use
- Provides data consolidation, protection, and sharing
- Requires no additional software
- Reduces total cost of ownership (TCO)

In traditional backup environments, as shown in Figure 8-20, the client data is backed up in two ways:

- Local area network (LAN) clients write their backup data by the backup server to tape.
- LAN-free clients use the SAN for direct backup to the tape devices.

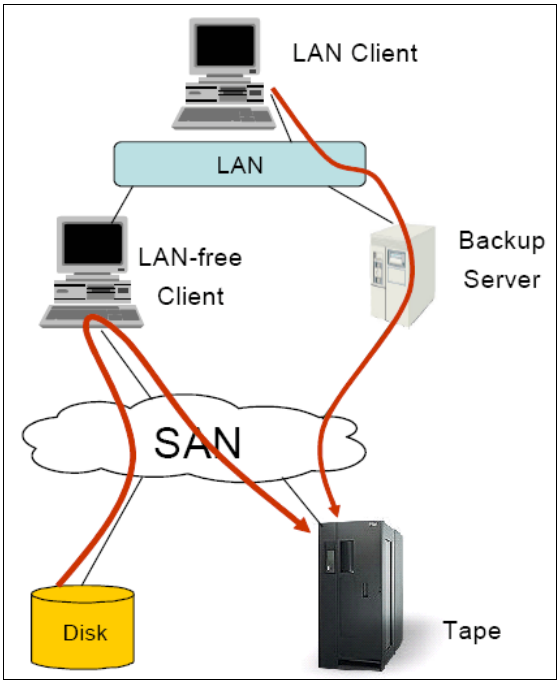

Figure 8-20 Traditional backup across the LAN, SAN to tape, or both

The following problems are encountered with these traditional backups:

- LAN backups cannot benefit from fast tape drives because the network transfer rate is the limiting factor. Data cannot be delivered fast enough to keep the tape drives streaming.
- LAN-free clients require dedicated tape drives for every backup session. In this case, many tape drives are needed during the backup window.

A solution for these problems is to use disk-to-disk (D2D) or disk-to-disk-to-tape (D2D2T) backups.

LAN backups can be performed by using a disk buffer, as shown in Figure 8-21. The client writes the backup data by LAN first to a disk buffer on the backup server. From there, it can be recalled rapidly if needed. After the backup on disk is completed, the backup server can migrate the data from the disk buffer to a physical tape with no impact on the client. Furthermore, the disk buffer works as a cache. The most recently backed up data is in the fast disk buffer until it is overwritten by newer data.

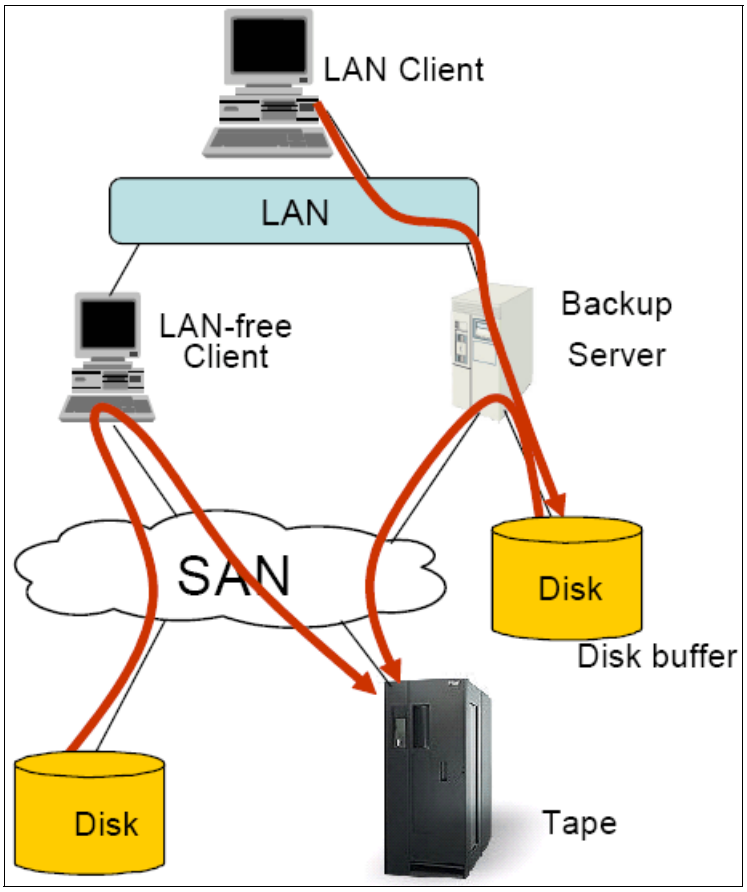

Figure 8-21 LAN Backup to a disk buffer (second step to tape) and LAN-free backup to tape

A LAN backup is a good solution if the backup software supports it. IBM Spectrum Protect supports a hierarchy of storage pools, and it can automatically migrate from disk to tape. Other backup applications might not automatically migrate from disk to tape.

An alternative solution to optimizing backup is to introduce a virtualization and emulation layer. This layer appears to both the backup server and the LAN-free clients as a tape drive.

The combination of hardware (servers, cache controllers, and cache modules), and software (tape drive and robotic emulation) is called a virtual tape library (VTL) as shown in Figure 8-22.

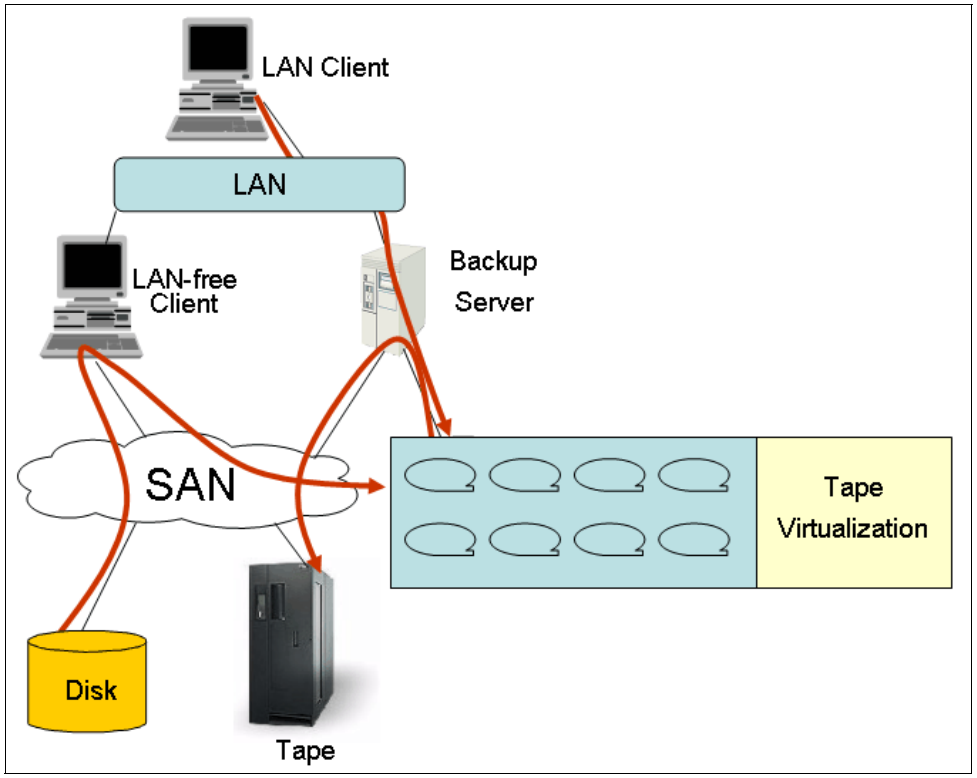

Figure 8-22 Tape virtualization where traditional disk looks like a tape library to the backup host

8.3.2 Concept of data deduplication

Data deduplication is a technology that is used to reduce the amount of space that is required to store data on disk. It is achieved by storing a single copy of data that is backed up repetitively. The principle of data deduplication is shown in Figure 8-23.

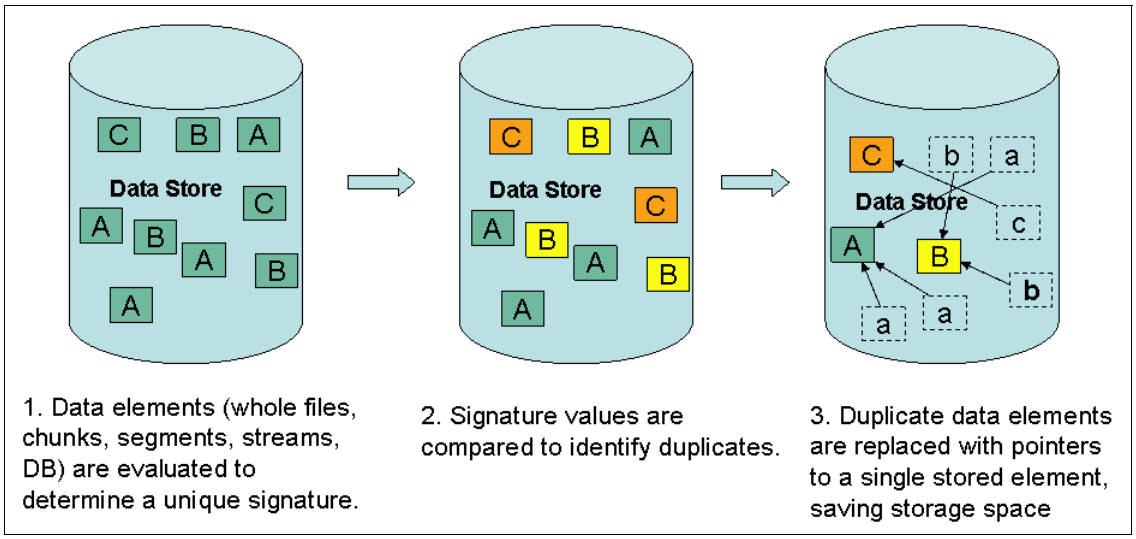

Figure 8-23 The basic concept of data deduplication

With data deduplication, data is read by the data deduplication product while it looks for duplicate data. Different data deduplication products use different methods of breaking up the data into elements, but each product uses a technique to create a signature or identifier for each data element. As shown in item 1 in Figure 8-23 on page 215, the data store contains three unique data elements, A, B, and C, each with a distinct signature. Whether using real-time or post processing deduplication, data element signature values are compared to identify duplicate data. In item 2, all of the duplicate data items are identified (shown by color coding). After the duplicate data is identified, one copy of each element is retained, pointers are created for the duplicate items, and the duplicate items are not stored.

IBM HyperFactor deduplication technology

IBM HyperFactor® is a patented technology that is used in IBM ProtecTIER Enterprise Edition V2.5 and later software. All available IBM ProtecTIER deduplication products use inline (real-time) deduplication processing and are designed to provide 100% data integrity by avoiding the risk of hash collisions.

HyperFactor technology, as shown in Figure 8-24, uses a pattern algorithm that can reduce the amount of required space for storage in the backup environment by up to a factor of 25, based on evidence from existing implementations. The capacity expansion that results from data deduplication is often expressed as a ratio, essentially the ratio of nominal data to the physical storage that is used.

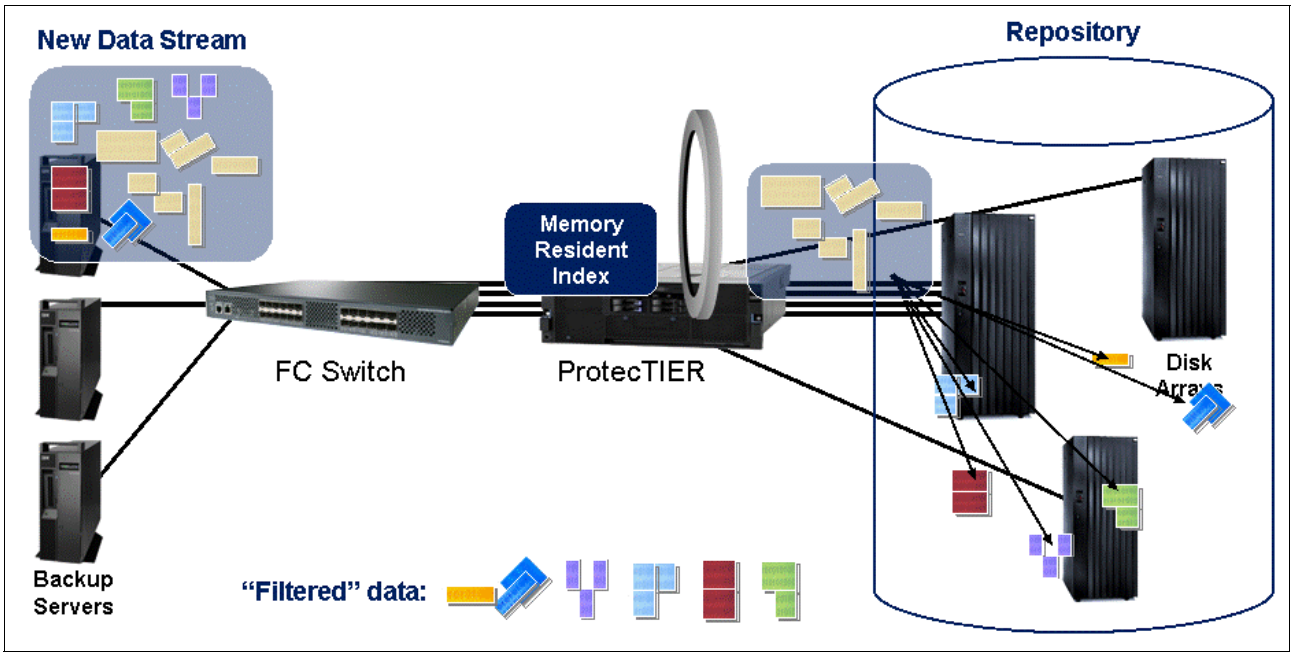

Figure 8-24 IBM HyperFactor technology

8.3.3 IBM TS7600 ProtecTIER systems

The IBM TS7600 ProtecTIER Deduplication solution meets the disk-based data protection needs of the enterprise data center while enabling significant infrastructure cost reductions. The solution offers high-performance inline data deduplication, optional 2-node clustering, and up to 1 PB of physical storage capacity per system. Combined with IBM storage, the ProtecTIER Gateway solution provides an optimal disk-based target for backup and archive data.

With ProtecTIER version 3.4, the File System Interface Network File Interface (FSI-NFS), Common Internet File System (FSI-CIFS), and VTL configurations are supported.

The IBM TS7600 ProtecTIER Deduplication solution can scale to meet the needs of small businesses to the needs of an enterprise-class data protection platform. This solution is designed to quickly and safely protect business information while reducing the amount of required space for storage. Clients can experience an up to 25:1 storage capacity reduction. The TS7600 solution is designed to provide the following benefits:

- Improve backup and recovery operations without changing existing backup processes, procedures, or policies.
- Shorten the backup window and simplify data management.
- Consolidate heterogeneous tapes and gain control of decentralized open systems data.

ProtecTIER models for open systems

The ProtecTIER TS7600 family consists of the following models:

- IBM TS7650G ProtecTIER Deduplication Gateway
- TS7620 ProtecTIER Deduplication Appliance Express

The IBM TS7620 Expansion Drawer provides additional repository capacities for TS7620 systems.

8.3.4 TS7650G ProtecTIER Deduplication Gateway

The IBM TS7650G ProtecTIER Deduplication Gateway is designed to meet the disk-based data protection needs of the enterprise data center while enabling significant infrastructure cost reductions. The solution offers industry-leading inline deduplication performance and scalability up to 1 PB of physical storage capacity per system, and can provide up to 25 PB or more backup storage capacity. Combined with IBM or third-party storage, the TS7650G ProtecTIER Deduplication Gateway provides a powerful disk-based repository to improve the performance, retention, and availability of backup and archive data.

TS7650G ProtecTIER Deduplication Gateway (3958-DD6)

The latest TS7650G gateway model is the 3958 DD6, which is a two rack unit device. It does not include a disk storage repository, which allows you to choose from various supported disk storage options. The DD6 is available in single node, dual node independent, and dual node cluster options. Dual power and cooling are included in the chassis. Call-home and support functions are contained within the device so there is no longer a requirement for a TSSC.

Figure 8-25 shows the latest IBM TS7650G Gateway, model 3958 DD6.

Figure 8-25 IBM TS7650G deduplication gateway

The latest TS7650G DD6 gateway offers an 80% reduced rack hardware footprint over the previous TS7650G DD5 gateway model, as illustrated in Figure 8-26.

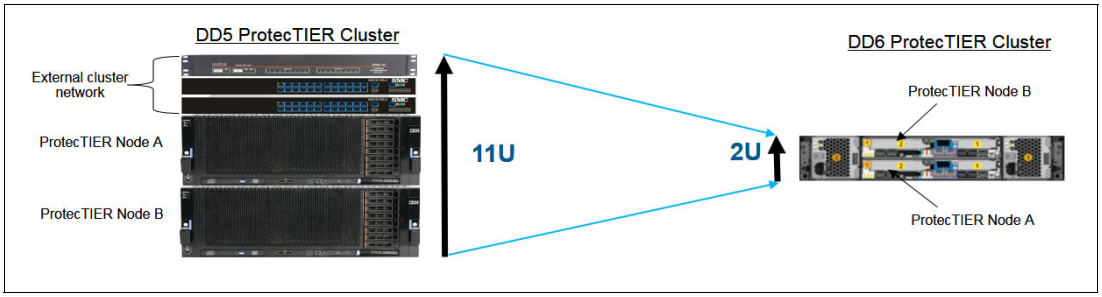

Figure 8-26 TS7650G Reduced Rack Hardware Footprint

Note: The previous TS7650G DD5 model was available in May 2012 and was withdrawn on 18 June 2016. It was based on the x7145 model type. It required a TSSC for call-home and support purposes.

The TS7650G, which consists of the 3958-DD6 hardware that is combined with the IBM ProtecTIER Enterprise Edition software, is designed to address the data protection needs of enterprise data centers. The solution offers high performance, high capacity, scalability, and a choice of disk-based targets for backup and archive data. The TS7650G offers these benefits:

- Inline data deduplication that is powered by HyperFactor technology.
- Multi-core virtualization and deduplication engine.
- Clustering support for higher performance and high availability.
- Fibre Channel (FC) ports for host and server connectivity.
- Performance up to 2500 MBps or more sustained inline deduplication backup (2-node clusters, depending on the configuration of attached disks (VTL configuration)).
- Up to 3200 MBps or more VTL sustained recovery performance.
- Virtual tape emulation of up to 16 VTLs per single node or 2-node cluster configuration, and up to 512 virtual tape drives per 2-node cluster or 256 virtual tape drives per single node.
- FSI-NFS support for NAS backup environments based on Linux.
- FSI-CIFS support for NAS backup environments based on Microsoft Windows.
- Emulation of the TS3500 tape library with IBM LTO Ultrium 3 tape drives
- Scales up to 1 PB of physical storage over 25 PB of user data (depending on data deduplication ratio).

Disk controllers and expansion units

You can select the disk controller and expansion units for use with the TS7650G. A list of compatible controllers and their expansion units is available at the following IBM Tape Systems Resource Library website:

http://www.ibm.com/systems/storage/tape/library.html

Verify compatibility in the TS7650G ISV and the interoperability matrix document that you find on the website.

ProtecTIER does not require more compression to be effective. It performs compression, by default, after the data deduplication process.

For more information, see the TS7650G ProtecTIER Deduplication Gateway product page:

http://www.ibm.com/systems/storage/tape/ts7650g

8.3.5 TS7620 ProtecTIER Deduplication Appliance Express

The IBM TS7620 Deduplication Appliance Express is designed for midsize companies that need a more holistic approach to data protection. The IBM TS7620 ProtecTIER Deduplication Appliance Express uses unique data deduplication technology to help solve critical backup and recovery challenges. The TS7620 ProtecTIER Deduplication Appliance Express is the machine type 3959 SM2.

Available in two configuration options, the TS7620 ProtecTIER Deduplication Appliance Express is an integrated server and storage hardware platform that includes preinstalled IBM ProtecTIER deduplication software. This solution has a preconfigured repository and can be configured with either a VTL or an FSI interface.

Figure 8-27 shows the IBM TS7620 appliance.

Figure 8-27 IBM TS7620 appliance

TS7620 features and benefits

The TS7620 Appliance Express provides the following features and benefits:

- Single 3U integrated appliance in 6 TB and 12 TB physical capacities, upgradeable to 35 TB
- Plug-and-play installation
- VTL or FSI CIFS/NFS support
- Improves backup and recovery performance with high-speed, disk-based data protection
- Accelerates recovery of critical data
- Optimizes storage infrastructure and helps reduce TCO
- Helps achieve business resilience objectives without changing existing backup procedures and practices
- Uses data deduplication technology to help solve backup and recovery challenges in midsize environments

For more information, see the TS7620 ProtecTIER Deduplication Appliance Express product page:

http://www.ibm.com/systems/storage/tape/ts7620

8.3.6 ProtecTIER Manager

The ProtecTIER Manager (PT Manager) application is a GUI that is required for service or maintenance on ProtecTIER systems. It must be installed on a client-provided workstation. You use it to view, configure, manage, upgrade, and troubleshoot the operation of the TS7600 family of products.

ProtecTIER Manager can support more than one ProtecTIER server. If the ProtecTIER Manager Workstation is connected to more than one ProtecTIER server, they can be connected to a workstation by a switch or by a terminal server. The workstation and switch must be provided by the client.

8.3.7 ProtecTIER native replication

ProtecTIER native replication is a logical feature that enables clients to replicate, or move any or all of their virtual tape cartridges from their main site to a remote disaster recovery location and vice versa. The TS7620 Appliance, the TS7650 Appliance, and the TS7650G Gateway models employ an IP-based replication design. Data is replicated between sites from one 2-node cluster to another. Tape archiving capabilities outside of the production environment are also provided with this feature.

In a disaster, the TS7620 or TS7650G at the disaster recovery site can become the production site (failover) until the main site comes back online. At that point, the client can replicate or move the newly created tapes back to the main production site (failback). ProtecTier and Native Replication enable more frequent testing of disaster recovery plans.

Types of replication

Two types of replication are available:

- Many-to-One replication: One hub can receive up to 12 spokes for replication.
- Many-to-many replication: In a 4-way group of systems with ProtecTIER, you can replicate from one to each of the three others, bidirectionally. This type of replication is available with ProtecTIER version 3.1 software and later. For FSI, the many-to-many topology is supported for up to eight bidirectional servers in a grid with up to three target destinations for each source.

Figure 8-28 shows the concept of ProtecTIER native replication in a VTL environment with an IBM ProtecTIER Gateway as a hub at the central (disaster) site and four spokes at branch offices (remote sites). As shown in the figure, each hub supports up to 12 spokes, depending on the specific model and deduplication solution.

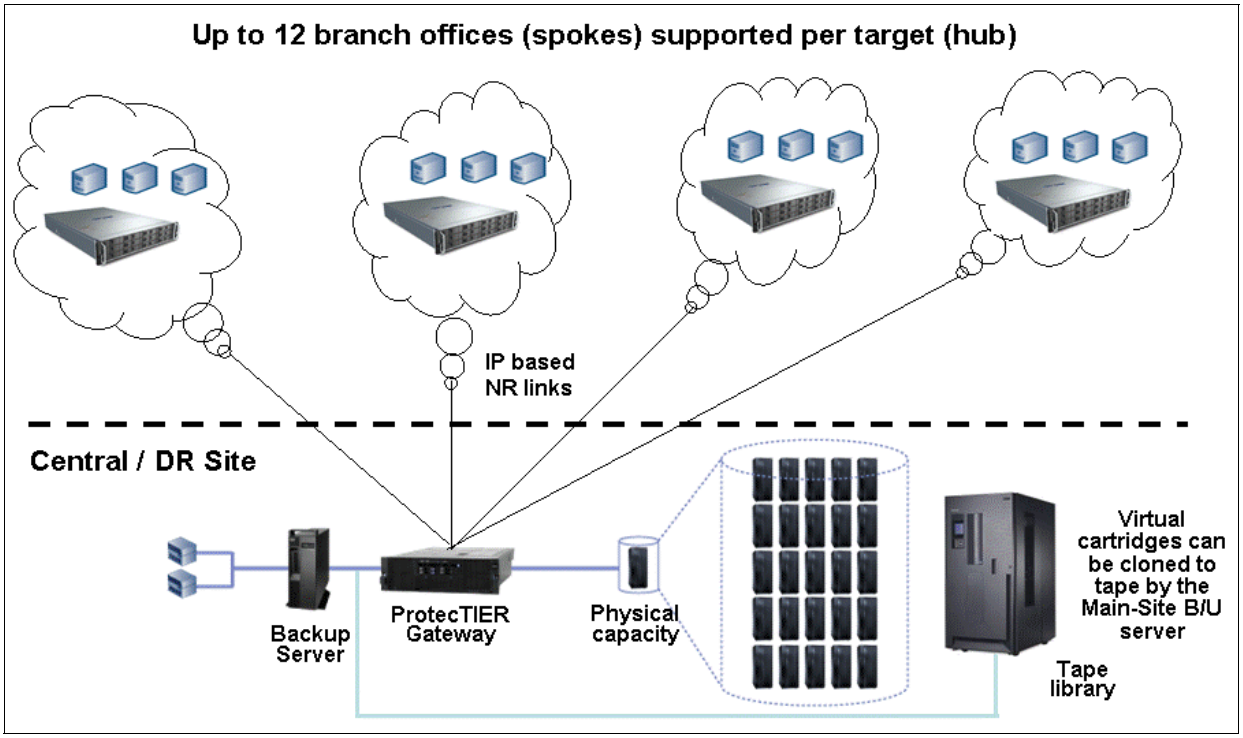

Figure 8-28 Native Replication that uses an IBM ProtecTIER solution in a VTL environment

Benefits of ProtecTIER native replication

Because ProtecTIER deduplicates data before storing it, only the changes of the data are transferred to the disaster recovery site over the replication link. This method translates into substantial savings in the bandwidth that is needed for the replication link. Replication can run in parallel to backup and restore activity.

ProtecTIER Native Replication eliminates the risks that are associated with human intervention and physical tape transportation. Replication occurs automatically (by policy) and simultaneously as data is deduplicated and backed up.

ProtecTIER Native Replication enables clients to improve their disaster recovery capabilities in the following ways:

- Accelerated recovery time: With ProtecTIER, all data is stored on high-speed disk, making recovery fast and efficient, and enabling systems to be brought online at a fraction of the time of a tape-based recovery environment.
- Expanded disaster recovery coverage: IBM ProtecTIER enables more production applications and associated data to be protected with replication, which was previously reserved only for Tier 1 applications and mission-critical data.
- Automated Operations: Native replication is a tightly integrated feature of ProtecTIER that enables IT organizations to automate the process of getting backup data to an offsite location quickly and reliably, avoiding the hassles and security risks that are associated with the transport of physical tapes.

For more information about best practices for ProtecTIER offerings, see *IBM ProtecTIER Implementation and Best Practices Guide*, SG24-8025.

8.3.8 IBM TS7700

To effectively protect data in z System environments, deliver uninterrupted data access, address regulatory requirements, and archive business data, organizations need a solution that can manage both primary and backup data that is active, inactive, or even archived.

IBM TS7700 is a family of mainframe virtual tape solutions that optimize data protection and business continuance for z Systems data. By using virtualization and disk cache, the IBM TS7700 family operates at disk speeds while maintaining compatibility with existing tape operations. The TS7700 fully integrated tiered storage hierarchy takes advantage of both disk and tape technologies to deliver performance for active data and the best economics for inactive and archive data.

The IBM TS7700 family of z Systems tape virtualization is offered in three models:

- The IBM TS7720 features encryption-capable high-capacity cache that uses 3 TB SAS disk drives with RAID 6 to scale to large capacities with the highest level of data protection.
- The IBM TS7740 features encryption-capable 600 GB SAS drives with RAID 6 protection.
- The IBM TS7720T (tape-attached) combines the benefits of the TS7720 and TS7740. It also provides Disk Cache Partition, which provides better control of how workloads use the disk cache, and Delay Premigration, which provides the ability to delay movement to tape.

The IBM TS7740 attaches a TS3500 tape library to the 3592 family of tape drives. The IBM TS7720 uses disk cache to store the data of the attached host system.

Both models write data by policy to physical tape through attachment to high-capacity, high-performance TS1150 tape drive models and earlier IBM 3592 tape drive models, which are installed in TS3500 libraries. Physical tape support is optional on the IBM TS7720.

Figure 8-29 shows a TS7720 with an attached expansion frame for higher disk cache capacity.

Figure 8-29 TS7720 with an optional expansion frame for higher disk capacity

Features and benefits

Table 8-10 lists the features and benefits of the IBM TS7700.

Table 8-10 IBM TS7700 features and benefits

Feature	Benefits
IBM TS7700 offers tiered hierarchy of disk and tape storage	► Helps improve performance and capacity to help achieve a low TCO for tape processing ► Helps improve application availability and shorten batch-processing windows
4,000,000 virtual logical volumes regardless of the number of clusters in the grid	► Helps reduce or eliminate bottlenecks in tape environments ► Supports the rereference of volumes without the physical delays that are typical of tape I/O ► Helps increase the performance of tape processes
Advanced policy management (subset)	► Enables cache management for volume retention and deletion ► Volume pooling allows grouping of logical volumes on physically separate cartridges or cartridge pools ► Copies export logical volumes from stand-alone and logical WORM ► Cross-site replication creates a copy of a logical volume at different sites

Feature	Benefits
Up to six IBM TS7700 systems can be configured to participate in a grid environment	► Helps to support high-availability requirements ► Helps enhance availability during planned maintenance and service or system upgrades ► Helps avoid transporting cartridges in a disaster
Support for both disk and tape drive-based encryption	► Helps secure sensitive at-rest data
Synchronous mode copy	► Provides true synchronous copy through real-time duplexing (simultaneous local mount and remote mount) ► Copy consistency at two locations up to any implicit or explicit sync point to provide a zero recovery point objective (RPO)
IBM FlashCopy support for disaster-recovery testing	► Performs disaster-recovery tests without interfering with production runs
Remote-mount IP link failover	► Improves remote read/write operations by using alternative links in a path failure ► If a failure occurs, the failover path searches another IP link to continue the job, without stopping or affecting the host or client

Key features

The IBM TS7700 has the following key features:

- Can be attached to supported IBM z Systems servers
- Support for a maximum of 4,000,000 virtual volumes
- Maximum of 1,536 virtual drives in a 6-way grid configuration
- Web-based management tools
- Advanced policy management
- TS3000 System Console
- Synchronous mode copy
- Remote-mount IP link failover
- Attach up to 96 TS1100 or IBM 3592 tape drives
- Up to 48 Fibre Channel connection (FICON) channels based on 6-cluster grid configurations

IBM TS7700 highlights

The following list provides key features of the IBM TS7700:

- Gain innovative data protection and business continuance for IBM z Systems servers
- Keep data secure, continuously available, and easy to manage
- Simplify z Systems connected tape operations and improve batch window performance
- Support disk and physical tap
- Disk cache refresh (CC9/CS9 and CS9/XS9)
- Disk cache encryption for data at rest
- Four million logical volumes
- Native Lightweight Directory Access Protocol (LDAP)
- Simplified Internet Protocol Security (IPsec)/IPv6 (on/off only - no configuration choices)

- TSSC remote access security (also known as LDAP)
- Reliability, availability, and serviceability (RAS) enhancements
- z Systems FICON attachment (up to four adapters per cluster)
- IBM TS7740 with up to 28.61 TB cache (by using 600 GB SAS disks)
- IBM TS7720 disk-only with up to 623.86 TB cache (by using 3 TB SAS disks)
- Support for 2-cluster, 3-cluster, 4-cluster, 5-cluster, or 6-cluster Grid
- Grid support with mixed IBM TS7720 and IBM TS7740
- 3592 drives, including native TS1120/TS1130/TS1140 support (up to 16)
- TS3500 library attachment
- Up to 1536 virtual drive addresses
- Logical WORM support
- Out-of-band tape encryption support
- Broadband Call Home (ECC)
- Code installation improvement
- Unified GUI in XIV style for intuitive administration and monitoring as shown in Figure 8-30

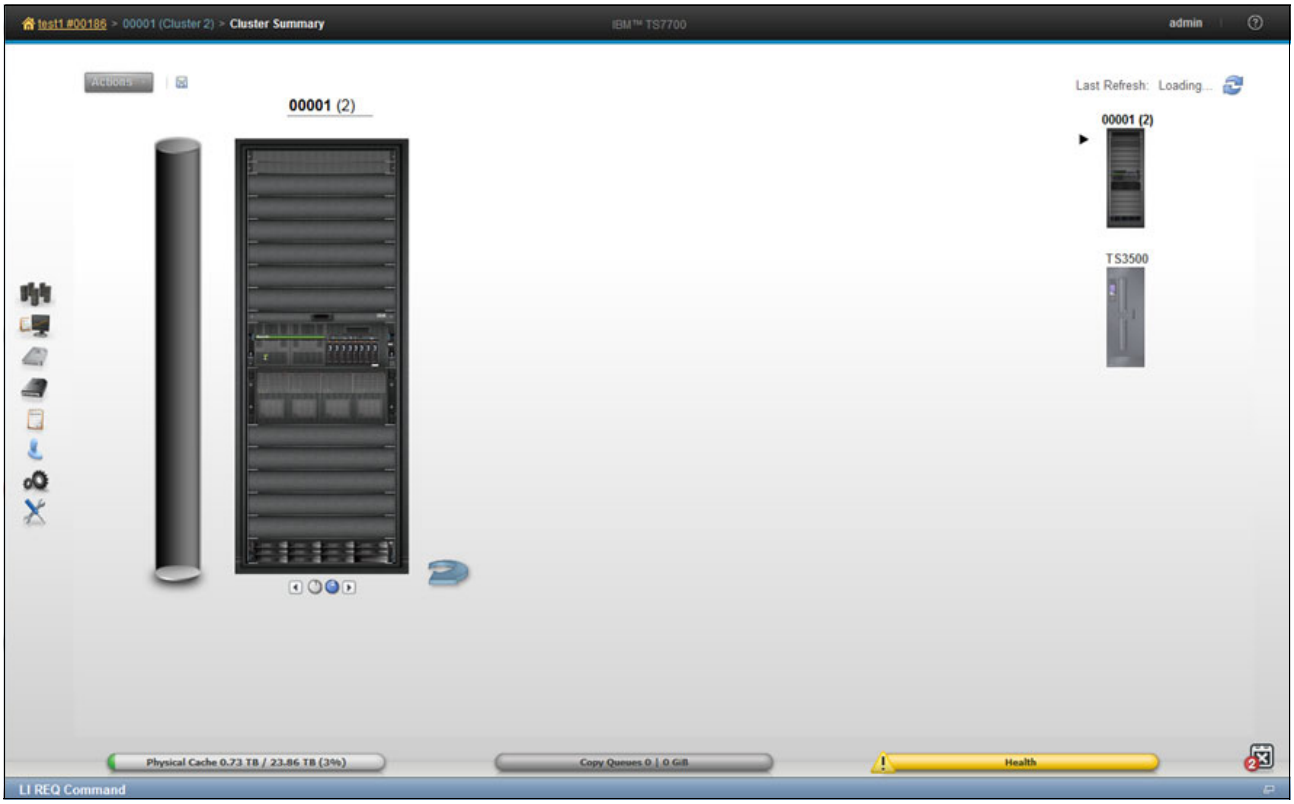

Figure 8-30 TS700 GUI for intuitive administration and monitoring

These capabilities are designed to improve performance and capacity to help lower the TCO for tape processing and avoid human errors. An IBM TS7700 can help improve the efficiency of mainframe tape operations by efficiently using disk storage, tape capacity, and tape speed, and by providing many tape addresses.

The TS7700 uses outboard policy management functions to manage cache and volume pools, and to control selective dual copy, dual copy across a grid network, copy mode control, encryption, and Copy Export. It also includes a new Unified GUI and enhanced statistical reporting.

Multiple TS7700s can form a clustered grid to enhance fault tolerance, increase availability, and eliminate downtime. Figure 8-31 shows a solution where multiple TS7700 engines with different models are connected.

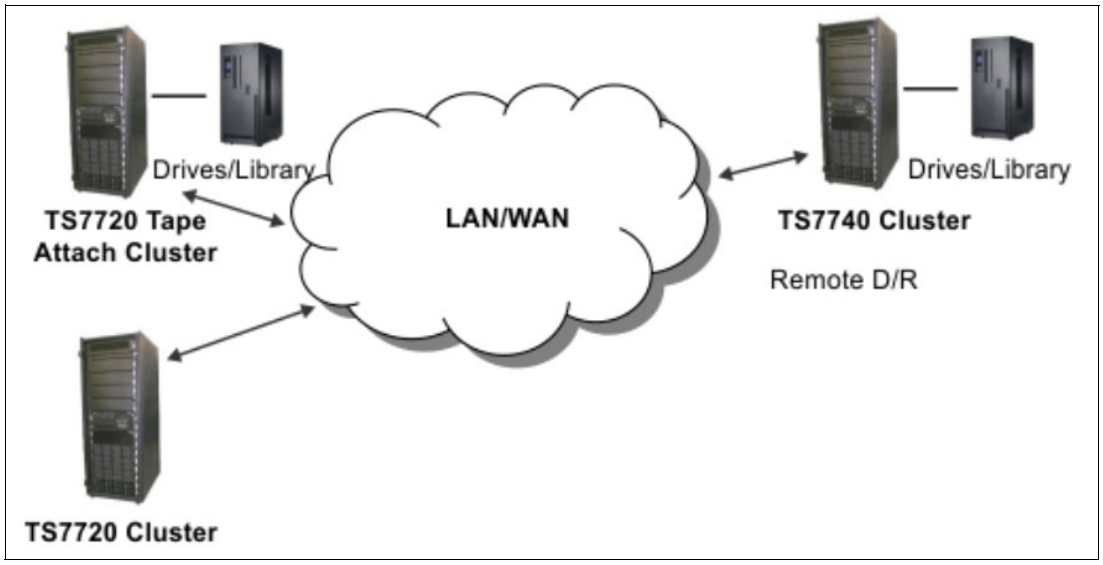

Figure 8-31 Multiple TS770 engines can be connected

IBM TS7700 release 3.3

Release 3.3 is enhanced to include support for the latest TS1150 tape drives and offering optimization in migration and disk encryption.

The following are the highlights of the actual release 3.3:

- Grid to Grid Migration (GGM)
- TS1150 tape drive support
- External managed disk encryption: External key management stores the key in an IBM Security Key Lifecycle Manager (SKLM) server (or ISKLM on z/OS)
- Management interface (MI) updates: Wider browser support, update for TS1150, encryption functionalities, GGM, and more.
- Other enhancements and new feature codes

Note: SKLM for external managed disk encryption is a separate software offering.

IBM TS7700 release 3.2

Release 3.2 is enhanced to include support for more virtual tape drives.

The IBM TS7700 includes 256 logical devices (or virtual tape drives) in 16 subsystems (least recently used (LRUs), LIBPORT-IDs, control units (CUs)) per cluster. Release 3.2 increases the number of supported virtual tape drives to a maximum of 496 devices. These additional devices can be purchased with Feature Code (FC) 5275, Additional Virtual Devices. Each instance of this feature supports one logical subsystem with 16 logical tape drives. A maximum of 15 Additional Virtual Device features can be installed on a single cluster.

8 Gb FICON adapters (FC 3438 or FC 3439) must be installed in a cluster before these additional devices can be defined. Existing configurations with 4 Gb FICON adapters do not support these additional devices.

All clusters in a grid must run release 3.2 or later before this feature can be operational. The same quantity of Additional Virtual Device features does not need to exist on every cluster. For example, this configuration means that one cluster can have the standard 256 virtual devices and another cluster can have 384 devices.

Specifications of the TS7700 models

Table 8-11 lists the specifications for the IBM TS7700.

Table 8-11 IBM TS7740 and IBM TS7720 at a glance

IBM TS7740 at a glance						
Specifications	**Single node**	**Two-cluster grid**	**Three-cluster grid**	**Four-cluster grid**	**Five-cluster grid**	**Six-cluster grid**
Usable RAID 6 array cache (TB)	Up to 1,008	Up to 2,016	Up to 3,024	Up to 4,032	Up to 5,040	Up to 6,048
Virtual drives	256	512	768	1,024	1,280	1,536
TS1100 or 3592 tape drives	4 - 16	8 - 32	12 - 48	16 - 64	20 - 80	24 - 96
Virtual volumes	4,000,000	4,000,000	4,000,000	4,000,000	4,000,000	4,000,000
Eight Gbps FICON channels	8	16	24	32	40	48
Maximum logical paths	4,096	8,192	12,288	16,384	20,480	24,576
Warranty	One-year onsite repair	One-year onsite repair	One-year onsite repair	One-year onsite repair	One-year onsite repair	One-year onsite repair
IBM TS7720 at a glance						
Specification	**Single node**	**Two-cluster grid**	**Three-cluster grid**	**Four-cluster grid**	**Five-cluster grid**	**Six-cluster grid**
Usable RAID 6 array cache (TB)	Up to 28	Up to 56	Up to 85	Up to 113	Up to 141	Up to 169
Virtual drives	256	512	768	1,024	1,280	1,536
TS1100 or 3592 tape drives	4 - 16	8 - 32	12 - 48	16 - 64	20 - 80	24 - 96
Virtual volumes	4,000,000	4,000,000	4,000,000	4,000,000	4,000,000	4,000,000

IBM TS7740 at a glance						
Specifications	**Single node**	**Two-cluster grid**	**Three-cluster grid**	**Four-cluster grid**	**Five-cluster grid**	**Six-cluster grid**
Eight Gbps FICON channels	8	16	24	32	40	48
Maximum logical paths	4,096	8,192	12,288	16,384	20,480	24,576
Warranty	One-year onsite repair	One-year onsite repair	One-year onsite repair	One-year onsite repair	One-year onsite repair	One-year onsite repair

Physical information

The following physical information applies to the IBM TS7700:

- Width: 644 mm (25.4 in.)
- Depth: 1,102 mm (43.4 in.)
- Height: 1,804 mm (71.0 in.)
- Weight: 481.9 kg (1,062 lb.)

The IBM TS7700 supports the following environments:

- IBM z/OS
- IBM z/VM®
- IBM z/VSE®
- IBM z/TPF

Figure 8-32 shows a TS7700 single frame configuration.

Figure 8-32 TS7700 single frame configuration

Characteristics of the IBM TS7700 components

The characteristics of these virtualization engines are summarized:

- The IBM TS7700 Server provides host connections for up to four FICON channels for attachment to IBM z Systems servers with the appropriate levels of z Systems software.
- Attached to a TS3500 tape library, the IBM TS7740 supports from 4 - 16 drives per cluster, with up to 96 in a 6-cluster grid.
- The IBM TS7720 does not attach to any tape library, but it can attach to the IBM Virtualization Engine TS7740, which can use the TS3500 tape library.
- Each IBM TS7700 supports up to a maximum of 256 virtual drives per cluster with up to 1,526 in a 6-way grid.
- The IBM TS7700 Single or Multi-Cluster Grid supports up to 4,000,000 logical volumes. Each logical volume has a maximum capacity of 1.2 GB - 18 GB (assuming 3:1 compression and 400 MB - 6000 MB volume sizes).
- Supported systems include IBM z/OS v1.9, z/VM v5.4, IBM z/VSE V4.1, and z/TPF 4.1

For more information, see the following link:

http://www.ibm.com/systems/storage/tape/ts7700/index.html

8.3.9 IBM Data Facility Storage Management Subsystem

IBM Data Facility Storage Management Subsystem (DFSMS) is a software suite that automatically manages data from creation to expiration. DFSMS provides the following functions:

- Allocation control for availability and performance
- Backup/recovery and disaster recovery services
- Space management
- Tape management
- Reporting and simulation for performance and configuration tuning

It offers the following benefits:

- Integrated, comprehensive data management solution
- Managed, predictable service to meet service level agreements (SLAs) for storage
- Up to 50% increased effective storage capacity with reduced overall storage cost

IBM Data Facility Storage Management Subsystem Removable Media Manager for z/OS environments

IBM Data Facility Storage Management Subsystem Removable Media Manager (DFSMSrmm) is the functional component of DFSMS. DFSMSrmm manages all removable media resources, including automatic, manual, and virtual tape libraries (such as the IBM TS7700 series), for the z/OS environment.

The following key functions are provided by DFSMSrmm:

- Tape data set recording: DFSMSrmm automatically tracks and records metadata for data that is written to all types of tape volumes.
- Inventory management: DFSMSrmm automatically records metadata for tape volumes that are used on the system. With DFSMSrmm, you can define any tape volume, enabling your complete tape inventory to be monitored and managed. DFSMSrmm supports tape volumes that are used anywhere in your enterprise, including z/VM, and the volumes that are managed by the IBM Integrated Removable Media Manager for the Enterprise on

z Systems (IRMM). You can easily partition your tape inventory and tape libraries and use DFSMSrmm to control which system can use the volumes.

- Lifecycle management: DFSMSrmm provides policies to enable tape data sets and tape volumes to be managed throughout their life. You define your tape data service levels and tape technology retention criteria, and DFSMSrmm automatically manages the lifecycle.
- Scratch pool management: DFSMSrmm supports either a global scratch pool or you can define one or more specific scratch pools for non-system managed tape volumes. DFSMSrmm manages the selection of the pool, and the return of volumes back to the scratch pool after the tape data expires.

DFSMSrmm highlights

DFSMSrmm has the following highlights:

- DFSMSrmm is included as part of DFSMS with the z/OS operating system. DFSMSrmm cooperates with and uses many system components that are required for the complete security, management, and well-being of your tape library.
- DFSMSrmm volume retention and movement policies can optionally be specified interactively with Interactive System Productivity Facility (ISPF) panels. It allows authorized application owners to alter existing values without contacting the tape librarian.
- DFSMSrmm includes the report generator, which is an ISPF-based tool for easily building and customizing many storage management reports. Sample report types and reports enable reporting that is based on DCOLLECT data, Data Facility Hierarchical Storage Manager (DFSMShsm) functional statistics, and DFSMSrmm-created data.
- In addition to the regular release updates, the DFSMSrmm enhancements provide many improvements, including support for the following products:
  - IRMM for the enterprise on z Systems
  - D/T3592 model E06 and larger tape volumes
  - IBM TS7700 release 1.6 and Copy Export
  - Migration health checks for V1R10 and V1R11
- DFSMSrmm is a full-function removable media management system for your enterprise. DFSMSrmm supports tape volumes, which are used anywhere in your enterprise, including z/VM, and volumes that are managed by the IRMM for the enterprise on z Systems.
- DFSMSrmm is functionally compatible with existing tape management systems and runs in parallel during conversion. The easy-to-use conversion tools, related documentation, and IBM Redbooks publications simplify the migration of your existing tape environment to management by DFSMSrmm.
- The special expiration date formats of 99*nnn* and 98*nnn* are supported by DFSMSrmm.

For DFSMSrmm enhancements, past and present, see the following website:

http://www.ibm.com/support/docview.wss?rs=0&q1=T1010391&uid=isg3T1010391&loc=en_US&cs=utf-8&cc=us&lang=en

To learn more about DFSMS, see the *DFSMS V1.10 and EAV Technical Guide*, SG24-7617.

9

Storage networking products

This chapter introduces storage networking options that can be employed to build an optimal storage environment. It describes both Ethernet and Fibre Channel options.

This chapter includes the following sections:

- Overview
- SAN and related technologies
- Storage area networking products
- Selecting the best alternative
- More information

9.1 Overview

Data access requirements of the many applications in use across an organization are different, and therefore cannot be solved with a single storage networking approach. There are trade-offs in cost, ease-of-management, performance, distance, and maturity.

The two popular models for attaching the storage to the network are storage area network (SAN) and network-attached storage (NAS). These methods help to remove direct attachments between storage and server, providing more flexibility in storage provisioning and the maximization of valuable storage and host resource.

The SAN can be viewed as an extension of the storage bus concept that enables storage devices and servers to interconnect by using elements such as routers, switches, directors, and gateways. Storage I/O, command, and control communications are performed on this dedicated network, providing an any-to-any connection for processors and storage on that network.

The most common transport media is Fibre Channel (FC). Fibre Channel Protocol (FCP) is the I/O protocol for open systems. Fibre Connectivity (FICON) and FCP are available for mainframe environments to communicate. These protocols offer high data rates and low latency communication, which are a requirement for today's data intensive applications.

Internet SCSI (iSCSI), Fibre Channel over IP (FCIP), Internet Fibre Channel Protocol (iFCP), Fibre Channel over Ethernet (FCoE), and Fibre Channel over Converged Enhanced Ethernet (FCoCEE) all provide the ability to implement storage connectivity by using Ethernet based solutions to extend the distance of an existing FC environment or maximize the Ethernet-based resources that are available.

9.2 SAN and related technologies

This section provides an overview of some key SAN technologies.

9.2.1 Storage area network

A SAN is a dedicated network for storage devices and the systems that access those devices. Figure 9-1 on page 233 represents the typical installation of a SAN. Storage networking today typically uses the FC technology, but the concept of a SAN is independent of the underlying type of network. I/O requests to disk storage on a SAN are called *block I/Os* because, similar to the direct-attached disk, the read and write I/O commands identify a specific device (disk drive or tape drive) and, for disks, specific block (sector) locations on the disk.

Figure 9-1 shows a typical SAN environment.

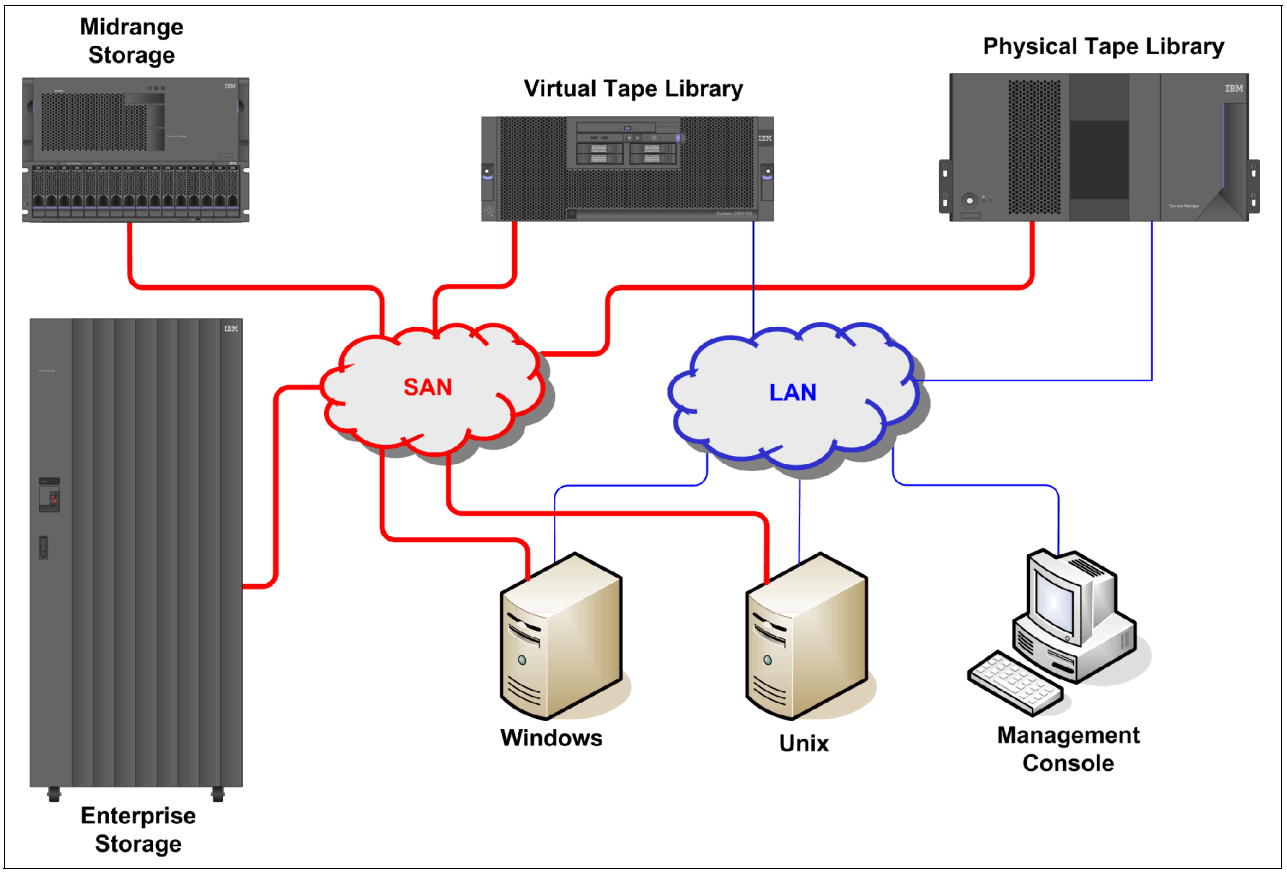

Figure 9-1 Typical SAN environment

SANs are used to connect shared storage arrays and tape libraries to multiple servers, and are used by clustered servers for failover. Through correct architecture and configuration, both availability and performance are increased.

A SAN facilitates direct, high-speed data transfers between servers and storage devices, potentially done in any of the following ways:

- Server to storage: This capability is the traditional model of interaction with storage devices. The advantage is that the same storage device can be accessed serially or concurrently by multiple servers.
- Server to server: With this capability, the SAN can be used for high-speed, high-volume communications between servers. In addition, the SAN is used as a communication path between cluster nodes in clustered systems, such as IBM PowerHA® SystemMirror® for AIX.
- Storage to storage: This outboard data movement capability enables data to be moved without server intervention, freeing up server processor cycles for other activities such as application processing. Examples include a disk device backing up its data to a tape device without server intervention, or remote device mirroring across the SAN.

SANs allow applications that move data to perform better, for example, by sending the data directly from the source to the target device with minimal server intervention. SANs also enable network architectures where multiple hosts access multiple storage devices that connect to the same network.

SAN usage can potentially offer the following benefits:

- Improvements to application availability: Storage is independent of applications and accessible through multiple data paths for better reliability, availability, and serviceability.
- Higher application performance: Storage processing is off-loaded from servers and moved onto a separate network.
- Centralized and consolidated storage: Management, scalability, flexibility, and availability are simpler.
- Data transfer and vaulting to remote sites: Remote copy of data is enabled for protection from disasters and malicious attacks.
- Simplified centralized management: Having a single image of storage media simplifies management.
- Storage cloud services: Storage virtualization and dynamic storage provisioning require the reliable SAN layer.

Fibre Channel

FC is well-established in the open systems environment as the underlying architecture of the SAN. FC is a technology standard that allows data to be transferred from one network node to another at a high speed. The interconnections between nodes are not necessarily based on *fiber* optics, but can also be based on copper cables. FC is ideal for moving large volumes of data across long distances quickly and reliably. In current implementations, the FC standard speed is generally available from 2 Gbps up to 32 Gbps. However, older 4 Gbps and 8 Gbps equipment is being replaced by faster connections.

FC is structured with independent layers, as are other networking protocols. Five layers exist, where 0 is the lowest layer. The physical layers are 0 - 2. These layers carry the physical attributes of the network and transport the data that is created by the higher-level protocols such as:

- Small Computer System Interface (SCSI)
- Transmission Control Protocol/Internet Protocol (TCP/IP)
- Fibre Channel connection (FICON)

Converged network

SAN and LAN are often implemented as separate network infrastructures by using Ethernet network interface cards (NICs), and Fibre Channel by using FC host bus adapters (HBAs). With today's high-performance networks, many servers are only able to use a small fraction of the available capacity of either network infrastructure.

Each of these networks requires its own adapters, fabrics, cables, tools, switches, management, and skills. If all of these components are combined, or converged, the potential for reducing cables, adapters, switches, and the skills that are required is clear. Replacing multiple networks with one in sometimes is an available option.

At a minimum, a converged network requires an adapter at the server that can carry the storage and IP networking workload. The access layer of the network (network switches or routers) must support multi-protocol traffic that is a combination of FC-based and IP-based data transport.

In contrast to the traditional server model, Figure 9-2 shows how the Converged Network Adapter (CNA) in a server, which is connected to the Enhanced Ethernet, has the potential to reduce the components that are required.

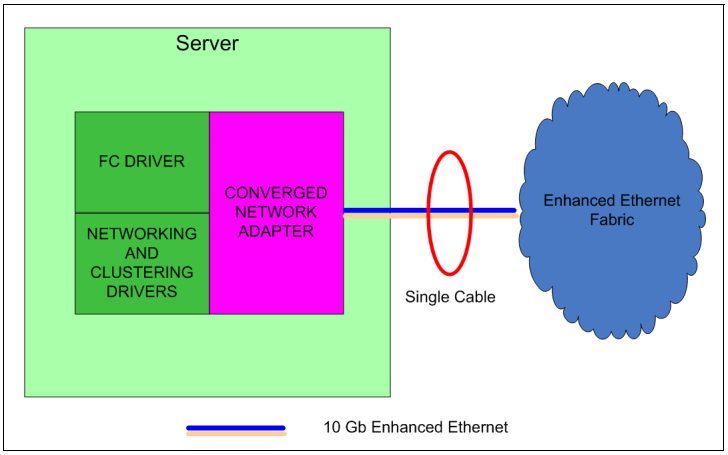

Figure 9-2 Serve with CNA installed

By using an FC driver, the CNA functionally represents a traditional FC HBA to the server's operating system. By using NIC or clustering drivers, the CNA functionally represents a traditional networking or clustering device to the server's operating system. The FC traffic is encapsulated into FCoE frames and these FCoE frames are converged with networking or clustering traffic.

The fabric contains FCoE-capable switches that can pass FC traffic to the attached SANs and Ethernet traffic to the attached Ethernet network. These switches must be able to support Enhanced Ethernet.

9.2.2 IP storage networking technologies

In contrast with the converged network, where the access layer transports the storage and IP-based workloads on a single medium (single cable and adapter), IP storage networking uses the IP-based networks only for storage workloads. The LAN or WAN cable that connects to the server's NIC transports only storage data and does not serve IP-based traffic.

With the continuous development of the IP-based networking technologies and increasing nominal capacities (10 Gbps and 40 Gbps), the potential for implementation of the IP storage networking solutions is growing. However, Internet Small Computer System Interface (iSCSI), Fibre Channel over IP (FCIP), and iFCP are SAN extension technologies that are ideal for connecting smaller departmental and less I/O-intensive systems into a SAN.

SAN deployment and its resulting benefits are primarily focused on mission-critical islands of application servers within individual data centers. The difficulty and cost that are associated with migrating the large number of data center midrange servers to FC made it impractical for IT managers to extend the benefits of SAN to midrange applications.

A basic diagram of these technologies is shown in Figure 9-3.

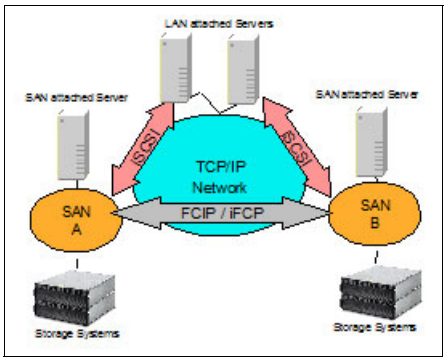

Figure 9-3 IP Storage Technology

Internet Small Computer System Interface

iSCSI allows storage to be accessed over a Internet Protocol network as though it was locally attached. The server pushes out the SCSI commands through the Ethernet NIC. As the SCSI commands exit out through the server, they are encapsulated within an IP packet and forwarded across the LAN to a LAN/SAN gateway interface. These iSCSI packets are then translated onto the FC SAN through the TCP/IP transport and conversion protocol.

The iSCSI is an ideal point-to-multipoint solution for connecting dispersed SAN islands. By using the iSCSI, midrange servers can gain access to consolidated storage while they retain their existing IP infrastructure. This configuration allows a cost-effective extension of SAN benefits to midrange applications on servers within the data center and on departmental servers throughout the enterprise.

The iSCSI uses TCP/IP for reliable data transmission over potentially unreliable networks. The iSCSI layer interfaces to the operating system's standard SCSI set. It includes encapsulated SCSI commands, and data and status reporting capability. When, for example, the operating system or application requires a data write operation, the SCSI Command Descriptor Block (CDB) must be encapsulated for transport over a serial gigabit link and delivered to the target.

Fibre Channel over IP

FCIP uses a tunneling protocol to transport FC frames over an existing IP infrastructure. It is therefore better suited for point-to-point solutions (Figure 9-4). By using FCIP, clients can use their current WAN infrastructure for connecting remote SAN islands over long distances.

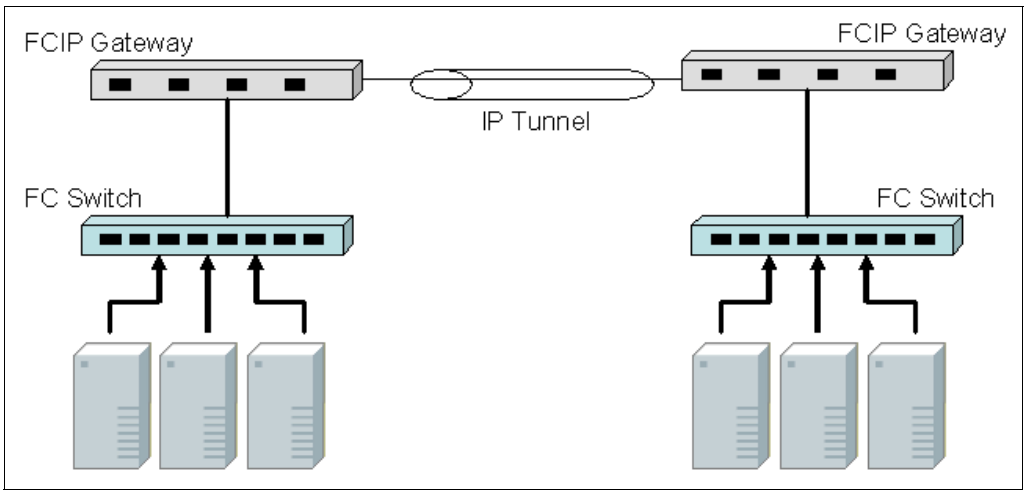

Figure 9-4 FCIP with IP Tunneling example

FCIP is an ideal combination of technologies to address the dual requirements of storage networking and networking over distance. FC is a mature technology that is optimized for storage data movement within the campus and data center.

FC represents a major investment in software compatibility, interoperability, and applications for campus-based storage networking. Likewise, IP is a mature technology that is optimized for data movement across WAN distances. It represents a major investment in software compatibility, equipment interoperability, and applications for WAN-based data networking.

FCIP solutions encapsulate FC and transport it over a TCP socket. As in all IP networks, performance can vary based on these factors:

- The types of switches and routers
- The number of hops that the packets must traverse
- The level of congestion in the network

Internet Fibre Channel Protocol

iFCP is TCP/IP-based method for interconnecting FC SANs, SAN devices, and IP networks. iFCP technology provides multipoint access to FC devices. iFCP capitalizes on Internet Protocol network services while using the performance and interoperability capabilities of the FC network. With iFCP, existing SCSI and FC networks can be interconnected into the existing Internet Protocol environment. iFCP can be used with FC switching and routing protocols, or it can completely replace them.

iFCP is designed for clients that might have a wide range of FC devices (that is, HBAs, subsystems, hubs, switches, and so on) and want the flexibility to interconnect these devices with an IP network. iFCP can interconnect FC SANs with IP and allow clients the freedom to use Internet Protocol (IP) networks in place of FC networks for the SAN itself. Through the implementation of iFCP as a gateway-to-gateway protocol, these clients can maintain the benefit of their FC devices while using a highly scalable, manageable, and flexible enterprise IP network as the transport medium.

iFCP enables FC device-to-device communication over an IP network, providing more flexibility compared to enabling only SAN-to-SAN communication.

A drawback of the FCIP protocol, which makes iFCP more attractive, is that FCIP is a protocol that uses tunneling to encapsulate FC data packets for forwarding over the IP network. Therefore, FCIP works only within an FC infrastructure. However, iFCP can handle both iSCSI and FCIP traffic. Applications that are developed for FC SAN environments are supported over iFCP.

The benefits of iFCP for storage networks are that scalability, distance, and connectivity issues are virtually eliminated. The existing Ethernet protocol structure allows the rapid deployment of applications and solutions that already use the TCP/IP protocol layers. With one less routing protocol to support, network complexity and management are also potentially reduced. More importantly, the lower cost of InfiniBand network switches as opposed to FC switches enables a lower total cost of ownership (TCO) of the enterprise SAN.

Fibre Channel over Ethernet

FCoE is an enhancement that expands FC into the Ethernet by combining two leading-edge technologies (FC and the Ethernet). It is the transport, or mapping, of encapsulated FC frames over the Ethernet. The Ethernet provides the physical interface and FC provides the transport protocol, the combination of which provides an FC frame that is delivered in an Ethernet frame.

Mixing FCoE and existing FC and Ethernet networks is easy and expected. The most probable implementation is “from the edge.” This process involves adding FCoE with new equipment while keeping existing Ethernet and FC hardware and cabling in place until it makes sense to replace it with FCoE.

FCoE has these highlights:

- FCoE uses CNA.
- CNAs transfer either Ethernet NIC traffic, FC traffic, or both.
- Enhanced Ethernet protocol supports FC traffic.
- Enhancements add lossless data transmission and more management functions.
- FCoE is also called FCoCEE.
- CNAs use 10 Gb physical Ethernet ports.

Each port can run all NIC, all FC, or mixed NIC and FC traffic.

For more information about FCoE and iSCSI, see *Storage and Network Convergence Using FCoE and iSCSI*, SG24-7986.

9.2.3 SAN topologies

FC-based networks support three types of topologies:

- Point-to-point (FC-P2P)

  Point-to-point is the simplest topology for an FC SAN. It allows the host and storage to connect directly.

- Loop (arbitrated) (FC-AL)

  Arbitrated loop, which is also known as FC-AL, is an FC topology in which devices are connected in a one-way loop fashion in a ring topology.

- Switched (FC-SW)

  Switched fabric is a computer network topology where many storage devices connect to each other through switches.

These topologies can be implemented separately or interconnected to form a *fabric*. The fabric can also be extended to cover even greater distances.

More definitions of SAN topologies exist, but they are always variations or combinations of these three basic concepts, such as these topologies:

- Single switch topology
- Cascaded and ring topology
- Mesh topology
- Core-Edge topology
- Edge-Core-Edge topology

Point-to-point topology

The point-to-point topology is the easiest FC configuration to implement, and it is also the easiest to administer. This simple link can be used to provide a high-speed interconnection between two nodes, as shown in Figure 9-5. A node is any device with one or more FC ports. Because connectivity is limited to two nodes, the exploitation of point-to-point in tape environments is limited. However, the distance between nodes can be up to 10 km (6.2 miles), which enables a tape library to be at another site.

When greater connectivity and performance are required, each device can be connected to a fabric without incurring any additional expense beyond the cost of the fabric.

Figure 9-5 shows a point-to-point topology.

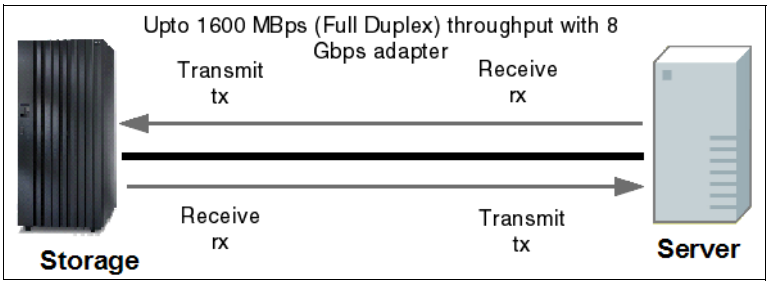

Figure 9-5 Point-to-point topology

Switched topology

FC switches provide increased bandwidth, scalable performance, an increased number of devices, and, in certain cases, increased redundancy. FC switches vary in the number of ports and media types that they support.

Multiple switches can be connected to form a switched fabric that can support many host servers and storage subsystems, as shown in Figure 9-6 on page 240. When switches are connected, all configurations are merged into one fabric configuration that is shared throughout the SAN.

The switches can be interconnected in a number of topology configurations to provide better performance, more available ports, and improved redundancy.

Traffic can be routed in many ways. For technical, security, or other reasons, various levels of zoning or other mechanisms can be used to restrict the any-to-any access. Performance

monitoring and configuration changes or upgrades that are needed to keep the network performing adequately are more complex.

Figure 9-6 shows a switched topology.

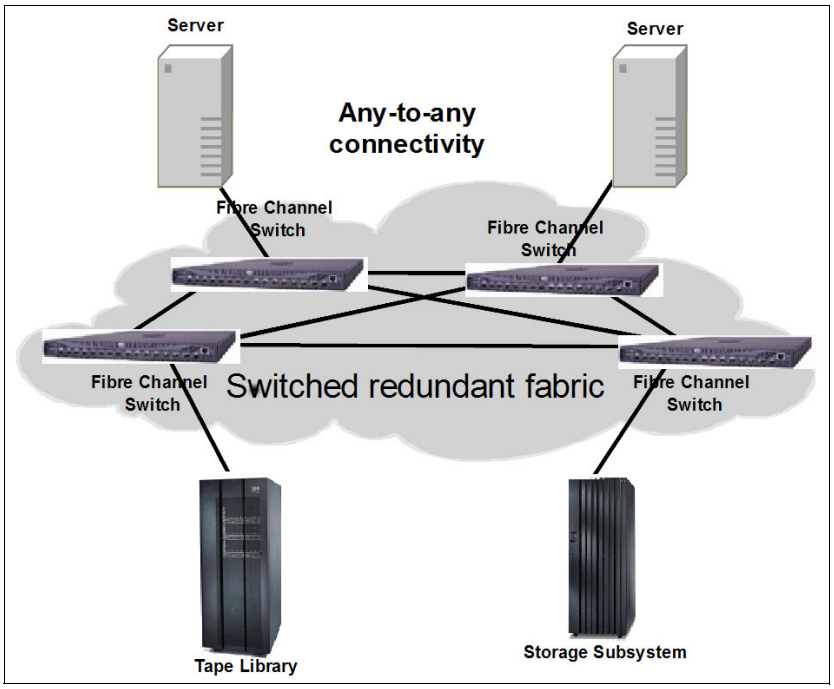

Figure 9-6 Example of a switched fabric

Loop (arbitrated) topology

Note: Loop technology is now rare in production environments. It is included briefly here because some existing tape technology still in production today uses it.

FC-AL offers relatively high bandwidth and connectivity at a low cost. For a node to transfer data, it must first arbitrate to win control of the loop. When the node has control, it is now free to establish a virtual point-to-point connection with another node on the loop. After this point-to-point (virtual) connection is established, the two nodes use all of the loop's bandwidth until the data transfer operation is complete. After the transfer is complete, any node on the loop can now arbitrate to win control of the loop.

Figure 9-7 shows a loop (arbitrated) topology.

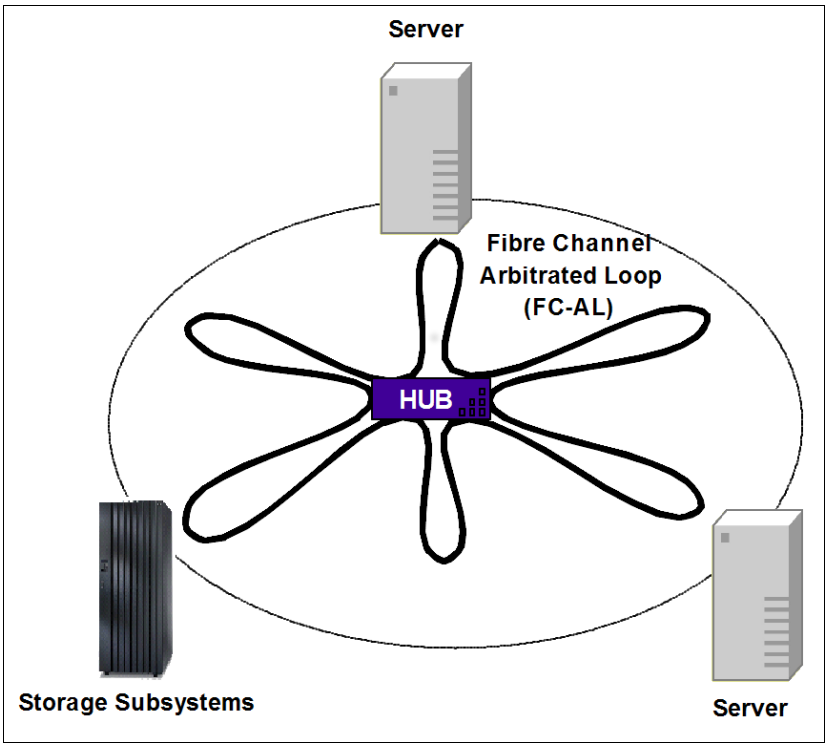

Figure 9-7 Arbitrated loop topology

9.2.4 Physical components of SAN infrastructure

SAN infrastructure is composed of a number of physical components. HBAs, cables, transceivers, switching, routing, and storage subsystem host ports all combine to create the end to end communication path.

Cables and connectors

FC connectors are available mainly in two types as LC and LC connectors. LC connectors are now the standard. See Figure 9-8.

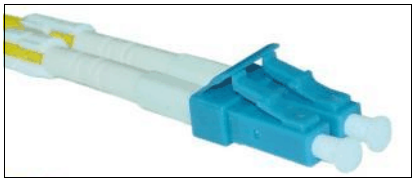

Figure 9-8 Standard LC connector

Transceivers

Transceivers are plugged on switches or directors on each port and used to convert the internal communication transport to gigabit transport. 2 Gbps, 4 Gbps, 8 Gbps, 10 Gbps, 16 Gbps, and 20 Gbps transceivers are called Small Form-Factor Pluggable (SFP) media. Figure 9-9 shows an SFP.

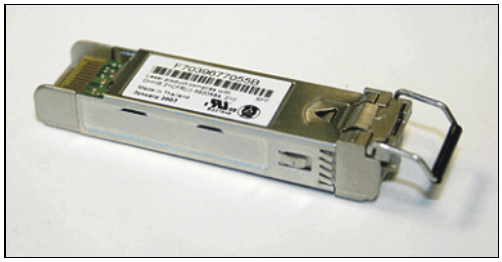

Figure 9-9 Fiber Optic SFP

Host bus adapters

The HBA connects to the bus of the host or storage system. It can connect to the cable that leads to the SAN. The function of the HBA is to convert the parallel electrical signals from the bus into a signal that can be transported to the storage or host.

HBAs are available for Fibre Channel, iSCSI, and FCoE, among others. Figure 9-10 shows FC, iSCSI, and FCoE HBAs.

HBAs with two or four FC ports are typically available.

Figure 9-10 Available HBAs

SAN switches, routers, and directors

SAN switches help servers, tape drives, libraries, and SAN storage to connect in a switched fabric. It provides connectivity and easy access to the data that is stored on your SAN storage devices, SAN switches provide high-speed reliable access.

Figure 9-11 shows a SAN384B-2 director class switch.

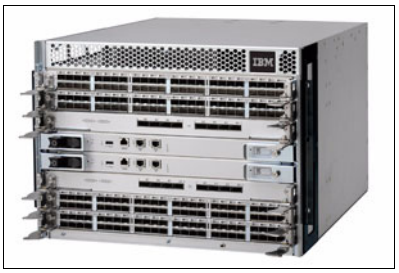

Figure 9-11 IBM Storage Networking SAN384B-2

Fibre Channel connection

FICON is a high-speed I/O interface for mainframe computer connections to storage devices. A FICON channel is a high-bandwidth connection between the processor and a storage device.

FICON channels increase I/O capacity through a combination of architecture and faster physical link rates. This combination makes them more efficient than IBM Enterprise Systems Connection (ESCON), the previous IBM fiber optic channel standard.

FICON is based on FC and runs on an FC infrastructure, but it requires separate FICON support from FC switches.

For more information about FICON, see the *FICON Native Implementation and Reference Guide*, SG24-6266.

9.2.5 Network-attached storage

NAS is a device on a network that can be shared with non-storage traffic. Currently, the network is usually an Ethernet LAN, but it can be any network that supports the IP-based protocols that NAS uses.

Figure 9-12 on page 244 shows a diagram of an NAS appliance. In contrast to the “block I/O” that is used by SANs, NAS I/O requests are called “file I/Os”. File I/O is a higher-level type of request that specifies the file to be accessed, an offset into the file (as though the file was a set of contiguous bytes), and a number of bytes to read or write beginning at that offset. File I/O requests are mainly Common Internet File System (CIFS), Network File System (NFS), or specialized I/O protocols for file access and file sharing. Unlike block I/O, no awareness exists of a disk volume or disk sectors in a file I/O request. Inside the NAS product, an operating system tracks where files are on disk and issues a block I/O request to the disks to fulfill file I/O read and write requests.

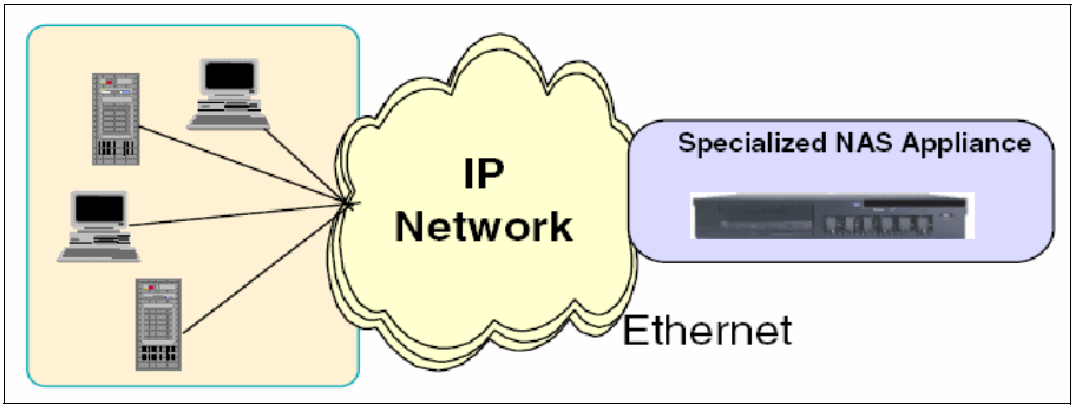

Figure 9-12 NAS network diagram

In contrast to SAN devices that can usually also be direct-attached (for example, by point-to-point FC) and network-attached by SAN hubs and switches, an NAS device is generally only an NAS device and attaches only to processors over a local area network (LAN) or wide area network (WAN).

NAS benefits

NAS products provide a wide range of network attachment capabilities to a broad range of host and client systems.

An NAS device offers the following benefits:

- Ease of installation

  NAS is generally easier to install and manage than a SAN. An NAS appliance can usually be installed on an existing LAN/WAN network. NAS manufacturers often cite “up and running” times of 30 minutes or less. NAS can be accessed by any operating system.

  Hosts can potentially start to access NAS storage quickly, without needing disk volume definitions or special device drivers. In contrast, SANs take more planning, including the design of an FC network and the selection and installation of SAN management software.

- Resource pooling

  NAS allows capacity within the appliance to be pooled. That is, the NAS device is configured as one or more file systems, each on a specified set of disk volumes. All users that access the same file system are assigned space within it on demand. That approach is more efficient than buying each user their own disk volumes or DAS, which often results in certain users having too much capacity and others too little. Therefore, NAS pooling can minimize the need to manually reassign capacity among users.

  However, NAS pooling resides within an NAS appliance, and little if any sharing of resources occurs across multiple appliances. It raises costs and management complexity as the number of NAS nodes increases. In contrast, an advantage of a SAN is that all devices (multiple disk and tape systems) on a SAN can be pooled. As total capacity grows, a SAN can become easier to manage and more cost-effective.

- File sharing

  NAS provides file sharing by using the NFS and CIFS protocols.

NAS gateways

An NAS gateway provides the function of a conventional NAS appliance, but without integrated disk storage. The disk storage is attached externally to the gateway, possibly sold separately, and can also be a stand-alone offering for direct or SAN attachment. The gateway

accepts a file I/O request (for example, by using the NFS or CIFS protocols). It then translates that request to a SCSI block I/O request to access the external attached disk storage. In the simplest terms, an NAS gateway is an NAS system in which the storage is externally attached rather than captively attached.

The gateway approach to file sharing offers the benefits of a conventional NAS appliance, with more potential advantages:

- Increased disk capacity scalability (compared to the capacity limits of an integrated NAS appliance).
- Ability to offer file sharing and block I/O on the same disk system. Disk capacity in the SAN can be shared (reassigned) among gateway and non-gateway use. Therefore, a gateway can be viewed as an NAS/SAN hybrid that increases flexibility, and potentially lowers costs.

9.3 Storage area networking products

SAN switches interconnect multiple host servers with storage servers and devices to create a SAN. The SAN switches can be used either as stand-alone devices to build a simple SAN fabric, or they can be interconnected with other switches to build a larger SAN fabric.

With the continuously increasing bandwidth of data centers' Ethernet, clients want unified storage and IP-based networking. IBM offers the solution for converged networking where both types of traffic use the same network infrastructure.

This section describes the IBM SAN product portfolio, which includes FC switches and directors with FC and FICON ports. It describes the following entry, midrange, and enterprise-level switch products:

- IBM SAN b-type switches and directors
- Cisco switches and directors

For the latest SAN products and information, see this website:

`http://www.ibm.com/systems/storage/san/`

IBM SAN switch offerings can be positioned into four broad product groups:

- Entry SAN switches
- Midrange SAN switches
- Enterprise SAN switches
- SAN specialty switches

9.3.1 Common characteristics

All IBM SAN products offer the following capabilities:

- Industry standard performance with 2, 4, 8, 10, and, in many cases, 16 Gigabits per second (Gbps) throughput
- Intelligent Fabric Services Architecture that provides switch interoperability
- Enterprise-level scalability with fault-tolerant core-to-edge SAN fabrics that contain thousands of devices
- Open fabric management that allows support for the widest range of solutions, from small workgroup SANs up to large enterprise SAN fabrics with thousands of devices

- Flexible management options, including Tivoli Ready certification for the centralized management of large enterprise SAN fabrics
- Manageable by IBM Tivoli Storage Productivity Center for Fabric
- Common enterprise SAN fabric that simplifies deployment, management, and network growth
- Flexible FC connectivity provides connectivity to a host of IBM and other vendors' servers and storage products
- Pay-as-you-grow scalability provides scalable network growth in a modular, cost-effective, and non-disruptive manner

9.3.2 Features specific to the IBM b-type SAN switches and directors

The following features are specific to the IBM b-type SAN switches and directors:

- Common firmware

  The common firmware enables the introduction of new switch technologies while protecting earlier switch investments, in addition to mixing and matching the SAN components, as required.

- IBM Network Advisor

  IBM Network Advisor is an intuitive graphical user interface (GUI) that allows network managers to monitor and manage SAN fabrics that consist of switches with a Java capable web browser from standard desktop workstations.

- Advanced Security

  Advanced Security significantly reduces the security holes that are left by traditional SAN implementations. It also greatly improves the ability to minimize SAN-specific vulnerabilities by providing a comprehensive, policy-based security system for IBM SAN switch fabrics.

- Advanced Zoning

  Advanced Zoning provides data exchange between devices in the same zone and prohibits exchange to any device that is not in the same zone. The zoning is enforced in application-specific integrated circuit (ASIC) hardware, preventing unauthorized devices from accessing the fabric.

- Extended Fabric

  The Extended Fabric feature provides extensions within the internal switch buffers. This feature maintains performance for distances greater than 10 km (6.2 miles).

- Fabric Manager

  Fabric Manager provides a Java based application that can simplify the management of a multiple switch fabric. It administers, configures, and maintains fabric switches and SANs with host-based software.

- Inter-switch link (ISL) Trunking

  The ISL Trunking feature allows ISLs between the same pair of switches to be grouped and to act as a single, high-speed "pipe" or trunk. Up to eight ISLs can be combined into a single logical ISL with a total bandwidth of 128 Gbps that can support any number of devices. ISL Trunking is designed to significantly reduce traffic congestion in storage networks.

- Virtual Fabrics

  The IBM b-type switches support the American National Standards Institute implementation of Virtual Fabrics (VFs). VFs add the capability for physical switches to be partitioned into independently managed logical switches, each with its own data, control, and management paths.

9.3.3 Fabric Vision

Fabric Vision with IBM SAN b-type Gen 5 Fibre Channel technology offers enhanced capabilities for managing your storage network that include simplified monitoring, increased availability, and reduced costs.

Fabric Vision simplifies and improves network monitoring with comprehensive visibility into network health and performance:

- A customizable health and performance dashboard simplifies management with critical information provided together on one window.
- Predefined policies, rules, and actions simplify the deployment of monitoring, but allow customization to match the aggressiveness of alerts and actions that you want.
- The solution's ClearLink diagnostics capability enables visibility to proactively discover issues with links, such as cabling or optics, to avoid problems that are caused by signal degradation.
- The ability to increase instrumentation and granularity helps identify latency, congestion, cyclic redundancy check (CRC) errors, timeouts, and other issues in the fabric. These help storage administrators visualize the network's health and performance.

For more information about the IBM b-type SAN switches and directors, see *IBM b-type Gen 5 16 Gbps Switches and Network Advisor*, SG24-8186.

9.3.4 Features specific to the Cisco SAN switches and directors

The following features are specific to the Cisco SAN switches and directors:

- Common firmware

  The common firmware enables the introduction of new switch technologies while protecting earlier switch investments, in addition to mixing and matching the SAN components, as required.

- Multiprotocol intelligence

  The Cisco SAN switches and directors transparently support FC, IBM FICON, and TCP/IP-based storage protocols.

- Virtual SAN (VSAN)

  The integrated hardware-based VSAN technology enables the physical switch infrastructure to be divided into several separate logical SAN fabrics, each with their own fabric services. This feature enhances the reliability of the SAN fabrics and helps in implementing interoperability between different SAN products.

- Inter-VSAN routing

  The integrated hardware-based inter-VSAN routing enables the routing of traffic at a line rate between VSANs, without requiring external hardware.

- Fabric-wide quality of service (QoS)

  Fabric-wide QoS helps manage bandwidth and control latency to prioritize critical traffic and enables migration from a SAN island to enterprise-wide storage networks.

- Virtual machine-aware SAN deployment

  Provides end-to-end visibility all the way from the VM down to storage, with resource allocation, performance measurements, and predictions that are available on a per-VM basis to enable rapid troubleshooting in mission-critical virtualized environments.

- Advanced traffic management

  Advanced traffic management capabilities are designed to simplify deployment and optimization of large-scale fabrics:

  - Virtual output queuing: Helps ensure line-rate performance on each port, independent of traffic patterns, by eliminating head-of-line blocking.
  - Up to 4,095 buffer-to-buffer credits: Can be assigned to an individual port for optimal bandwidth usage across distances.
  - Fibre Channel shortest path first (FSPF)-based multipathing: Provides the intelligence to load balance across up to 16 equal cost paths and, during a switch failure, dynamically reroute traffic.

- Port Channels

  The Port Channels feature allows the aggregation of up to 16 physical ISLs into a single logical bundle, providing optimized bandwidth usage across all links. The bundle can consist of any speed-matched ports from any module in the chassis, helping to ensure that the bundle can remain active even in a module failure.

- Cisco Data Center Network Manager (DCNM)

  Cisco DCNM (formerly Cisco Fabric Manager) is an easy-to-use application that simplifies management across multiple switches and converged fabrics.

For more information about the Cisco SAN switches and directors, see *IBM and Cisco: Together for a World Class Data Center*, SG24-8105.

9.3.5 Entry-level SAN switches

The IBM SAN switch entry-level products are designed specifically to address the needs of small to medium-size SAN environments. They can be used to create a wide range of high-performance SAN solutions, from simple single-switch configurations to larger multi-switch configurations that support fabric connectivity and advanced business continuity capabilities.

Infrastructure simplification solutions for IBM Power Systems include storage consolidation and high-availability server clustering with IBM disk storage arrays. Business continuity solutions include data protection with IBM tape libraries and devices and IBM Tivoli Storage Manager data protection software. IBM entry fabric switches provide up to twenty-four 16 Gbps ports for fully non-blocking performance, and advanced intelligence features.

This section introduces the ideal IBM SAN switch solutions for entry-level applications.

The following entry-level products are included:

- IBM System Storage SAN24B-4 Express
- IBM System Networking SAN24B-5

IBM System Storage SAN24B-4 Express

The SAN24B-4 Express (model 2498-24E) is a high-performance scalable switch that provides a 24-port fabric with a Ports on Demand feature, so you can scale your network as you grow. You can get the switch with an 8-port license and then can upgrade it conveniently up to 24 ports by enabling the license. With auto-sensing link speeds at 1, 2, 4, and 8 Gbps and a flexible design that allows you to configure this switch as a fabric switch or an Access Gateway, it is suitable for small to midsize businesses.

A single SAN24B-4 Express switch can serve as the cornerstone of a SAN when you want to obtain the benefits of storage consolidation and implement FC storage systems. An entry-level 8-port storage consolidation solution can support up to seven servers with a single path to either disk or tape. The Ports on Demand (POD) feature is designed to enable a base switch to grow to 16 - 24 ports to support more servers and more storage devices without taking the switch offline. Figure 9-13 shows the SAN24B-4 Express fabric switch.

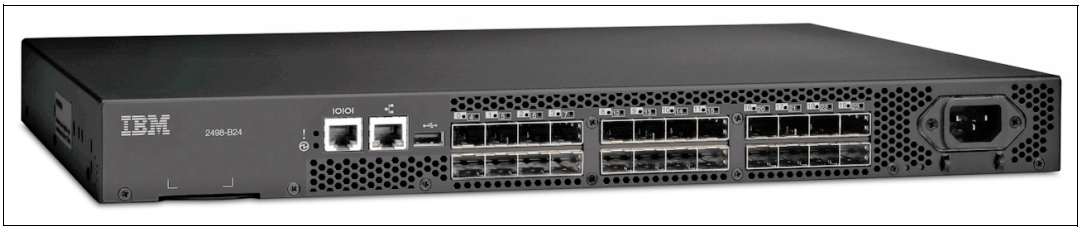

Figure 9-13 IBM System Storage SAN24B-4 Express

The SAN24B-4 fabric switch requires Fabric OS V6.1.0 or later. The switch offers easy to use Web Tools, 8-Gb FC, Long-Distance Support, Advanced Zoning, Full-fabric Support, Fiber Watch, Advanced Performance Monitoring, Fabric Vision Technology, Enhanced Group Management, and ISL Trunking. The base switch also offers eight default ports. The POD licenses are available in 8-port increments. With flexible architecture that is based on GoldenEye2 ASIC, the switch supports F, FL, E, and M ports at 8 Gbps. The switch also has USB port support for firmware download, configuration upload and download, and supportsave. The SAN24B-4 has a 1U form factor and is a single field-replaceable unit (FRU) with no field-replaceable parts. The switch has one power supply and three integrated fans.

Support: The Access Gateway mode is supported only in 24-port configurations, and *only* 2 GB Brocade branded USB drives are supported on the USB port. The 4 Gbps and 8 Gbps link speeds are supported only with Brocade-branded SFPs.

For more information, see this website:

`http://www.ibm.com/systems/storage/san/b-type/san24b-4/`

IBM System Networking SAN24B-5

The IBM System Networking SAN24B-5 switch is designed to provide outstanding price and performance value, combining flexibility, simplicity, 16 Gbps FC technology, and enterprise-class functions in an entry-level switch. The SAN24B-5 is configurable in 12 or 24 ports and supports 2, 4, 8 or 16 Gbps speeds in an efficiently designed 1U form factor. The base unit includes one (2498-24G/2498-X24) or two (2498-F24) integrated power supplies and fans. A second power supply provides more redundancy for increased resiliency.

The IBM System Networking SAN24B-5 switch is shown in Figure 9-14.

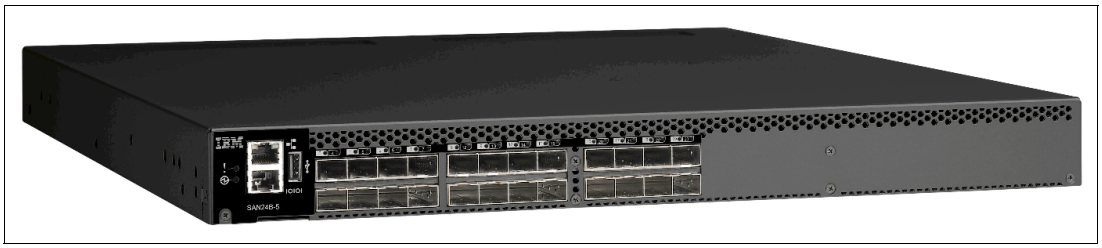

Figure 9-14 The IBM System Networking SAN24B-5 switch

The SAN24B-5 fabric switch requires Fabric OS V7.0.1 or later. The switch offers easy to use Web Tools, 16-Gb FC, Long-Distance Support, Advanced Zoning, Full-fabric Support, Fiber Watch, Advanced Performance Monitoring, Fabric Vision Technology, Enhanced Group Management, and ISL Trunking. The base switch also offers 12 default ports. The POD license is available for the remaining 12 ports. With flexible architecture that is based on Condor3 ASIC, the switch supports F, E, EX, D, and M ports at 16 Gbps. The FL port is not available. The switch also has USB port support for firmware download, configuration upload and download, and supportsave. The SAN24B-5 has a 1U form factor and is a single FRU with one hot-swappable power supply as a replaceable part. The secondary redundant power supply and fan assembly is an optional feature.

For more information, see *IBM System Networking SAN24B-5 Switch*, TIPS1128.

Also, see the product website:

http://www.ibm.com/systems/storage/san/b-type/san24b-5/

9.3.6 Midrange SAN switches

The IBM midrange SAN solutions provide more capability, features, and benefits beyond the simple entry solutions. They provide 1, 2, 4, 8, and 16 Gbps port-to-port non-blocking throughput with auto-sensing capability for connecting to older 1 Gbps host servers, storage, and switches. They are available in 48 - 96 ports with the POD scalability models.

The IBM SAN switches can be used to create dedicated, reliable, and high-performance networks for storage products, such as disk subsystems, tape drives, and tape drive libraries. In addition, all of these models are fully interoperable with the previous IBM SAN switches. They can be added to existing fabrics, enabling transition from existing FC storage networks to the faster technology.

The IBM SAN switch solutions are ideal for midrange applications. The following midrange SAN switches are included:

- IBM System Storage SAN48B-5
- IBM System Networking SAN96B-5
- Cisco MDS 9148S 16G Multilayer Fabric Switch for IBM System Storage
- Cisco MDS 9396S 16G Multilayer Fabric Switch for IBM System Storage

IBM System Storage SAN48B-5

The IBM System Storage SAN48B-5 SAN switch (model 2498-F48) is designed to meet the demands of hyper-scale, private cloud storage environments by delivering 16 Gbps FC technology and capabilities that support highly virtualized environments. To enable greater flexibility and investment protection, the SAN48B-5 is configurable in 24, 36, or 48 ports and supports 2, 4, 8, 10, or 16 Gbps speeds in an efficiently designed 1U package. The switch is shown in Figure 9-15.

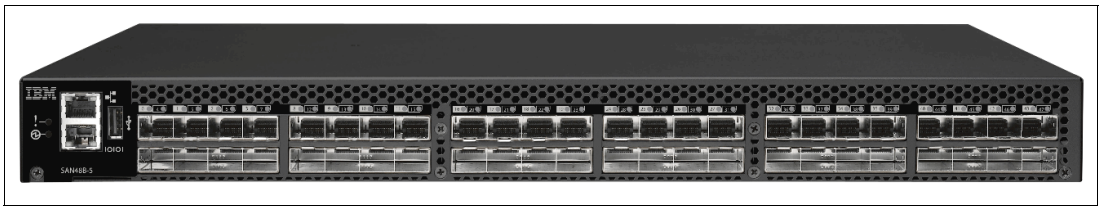

Figure 9-15 Front view of IBM System Storage SAN48B-5 switch

The SAN48B-5 fabric switch requires Fabric OS V7.0.1 or later. The base model of the switch has 24 ports that are enabled, and the POD licenses are available in 12-port increments. The switch can simplify server virtualization and virtual desktop infrastructure (VDI) management. The switch meets the high-throughput demand of enterprise data center switch capabilities, such as native 10 Gbps support for dense wavelength division multiplexing (DWDM) or in-flight data compression and encryption.

The SAN48B-5 can be deployed as a full-fabric switch or as an Access Gateway, which simplifies fabric topologies and heterogeneous fabric connectivity. The default mode setting is a switch. Access Gateway mode uses N_Port ID Virtualization (NPIV) switch standards to present physical and virtual servers directly to the core of SAN fabrics. This configuration makes it transparent to the SAN fabric, which greatly reduces the management of the network edge. The SAN48B-5 in Access Gateway mode can connect servers to NPIV-enabled b-type and m-type SAN directors, switches, and routers.

For more information, see *IBM System Storage SAN48B-5*, TIPS1125.

Also, see the product website:

`http://www.ibm.com/systems/storage/san/b-type/san48b-5/`

IBM System Networking SAN96B-5

The IBM System Networking SAN96B-5 SAN switch (models 2498-F96 and 2498-N96) is designed to meet the demands of hyper-scale, private cloud storage environments by delivering 16 Gbps FC technology and capabilities that support highly virtualized environments. To enable greater flexibility and investment protection, the SAN96B-5 is configurable in 48, 72, or 96 ports, and supports 2, 4, 8, 10, or 16 Gbps speeds in an efficiently designed 2U package.

The IBM System Networking SAN96B-5 is shown in Figure 9-16.

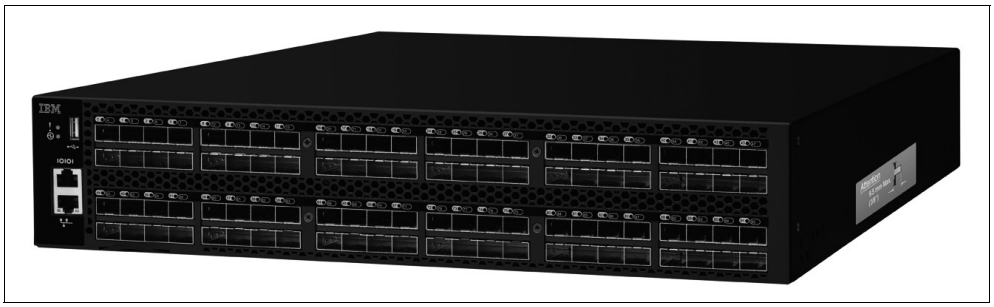

Figure 9-16 IBM System Networking SAN96B-5

High availability features make the IBM System Networking SAN96B-5 suitable for use as a core switch in midrange environments or as an edge-switch in enterprise environments where a wide range of SAN infrastructure simplification and business continuity configurations are possible. IBM Power Systems, z Systems, and many other vendor servers, disk, and tape devices are supported in many common operating system (OS) environments.

Optional features provide specialized distance extension, dynamic routing between separate or heterogeneous fabrics, link trunking, FICON, performance monitoring, and advanced security capabilities.

A single physical chassis can be subdivided into two or more logical switches that create a logical fabric with other switches. Integrated Routing is a licensed feature that is supported on every port of the switch and requires the POD license for all 96 ports.

The ports on the switch are grouped in 8-port groups that match the trunk group, and with ISL Trunking, speeds of up to 128 Gbps can be achieved per trunk. Dynamic Path Selection can be used for optimizing the performance and load balancing, and the switch can be managed by using Web Tools.

The IBM System Networking SAN96B-5 switch is available in two models to accommodate different rack airflow requirements:

- 2498-F96 with airflow from non-port side to port side
- 2498-N96 with airflow from port side to non-port side

For more information, see *IBM System Networking SAN96B-5*, TIPS1103.

Also, see the product website:

`http://www.ibm.com/systems/storage/san/b-type/san96b-5/`

Cisco MDS 9148S 16G Multilayer Fabric Switch for IBM System Storage

The Cisco MDS 9148S 16G Multilayer Fabric Switch for IBM System Storage is the latest generation of the highly reliable, flexible, and low-cost Cisco MDS 9100 Series switches. It combines high performance with exceptional flexibility and cost effectiveness. This powerful, compact one rack-unit (1RU) switch scales from 12 to 48 line-rate 16 Gbps FC ports.

The Cisco MDS 9148S is excellent for the following purposes:

- A stand-alone SAN in small departmental storage environments
- A top-of-the-rack switch in medium-sized redundant fabrics
- An edge switch in enterprise data center core-edge topologies

Figure 9-17 shows the Cisco MDS 9148S switch.

Figure 9-17 Cisco MDS 9148S switch

For more information, see *Cisco MDS 9148S 16G Multilayer Fabric Switch for IBM System Storage*, TIPS1255.

Also, see the product website:

http://www.ibm.com/systems/storage/san/ctype/9148S/

Cisco MDS 9396S 16G Multilayer Fabric Switch for IBM System Storage

MDS 9396S provides up to 96 autosensing Fibre Channel ports, capable of speeds of 2, 4, 8 and 16 Gbps, which results in 16 Gbps of dedicated bandwidth for each port. The switch comes with 48 enabled ports, and it can be scaled up by adding one or more 12-port MDS 9396S On-Demand Port Activation licenses.

MDS 9396S scales from 48 to 96 high-performance Fibre Channel ports in a 2RU form factor. It offers more buffer-to-buffer credits than previous-generation fabric switches. In addition, it supports 32 virtual SANs (VSANs), making it an excellent choice for stand-alone small and mid sized business Fibre Channel networks

Figure 9-18 shows the Cisco MDS 9396S 16G Multilayer Fabric Switch for IBM System Storage

Figure 9-18 Cisco MDS 9396S 16G Multilayer Fabric Switch for IBM System Storage

For more information, see the product website:

http://www.ibm.com/systems/storage/san/ctype/9396S

9.3.7 Enterprise SAN switches

The IBM enterprise SAN switches and directors address enterprise SAN client requirements for infrastructure simplification and improved business continuity.

The SAN directors are designed to operate with other members of the IBM System Networking Switch family. You can configure a wide range of highly scalable solutions that address demands for integrated, heterogeneous mainframe and open server enterprise SANs.

The following enterprise products are included:

- IBM System Storage SAN384B-2 and SAN768B-2
- Cisco MDS 9500 Series Multilayer Directors
- Cisco MDS 9706 Multilayer Director for IBM System Storage
- Cisco MDS 9710 Multilayer Director

This section introduces IBM System Networking Switch solutions ideal for enterprise environments. Advanced security with comprehensive, policy-based security capabilities can improve availability and simplify operation.

IBM System Storage SAN384B-2 and SAN768B-2

The IBM System Storage SAN384B-2 (model 2499-416) and IBM System Storage SAN768B-2 are fabric backbone switches. They are the premier platform for the consolidation of your data center connectivity, providing high-performance and highly available data networking with industry-leading 16 Gbps FC technology. They are also among the members of the IBM b-type SAN family that are designed to use the new IBM and Brocade data center fabric architecture.

The SAN384B-2 and SAN768B-2 interoperate with other members of the IBM b-type SAN family, in addition to other fabrics:

- They can be configured with a wide range of connectivity options, including 16, 10, 8, 4, and 2 Gbps FC, up to 4 Gbps FICON, and FCIP over 1 gigabit per second Ethernet (GbE).
- They are designed to enable support for emerging high-performance and high-function network protocols, including FCoCEE.
- They are designed to serve as the basis for transforming existing networks into a unified, high-performance data center fabric, connecting applications with their data and virtual servers with virtual storage.

As members of the IBM family of b-type SAN products, the SAN384B-2 and SAN768B-2 are designed to participate in fabrics that contain other b-type SAN devices that are manufactured by Brocade. Their versatile hardware can serve as a new top tier (or *backbone*) in a complex fabric and provide connections to other b-type SAN directors, switches, and routers.

The IBM System Storage SAN384B-2 and SAN768B-2 are shown in Figure 9-19.

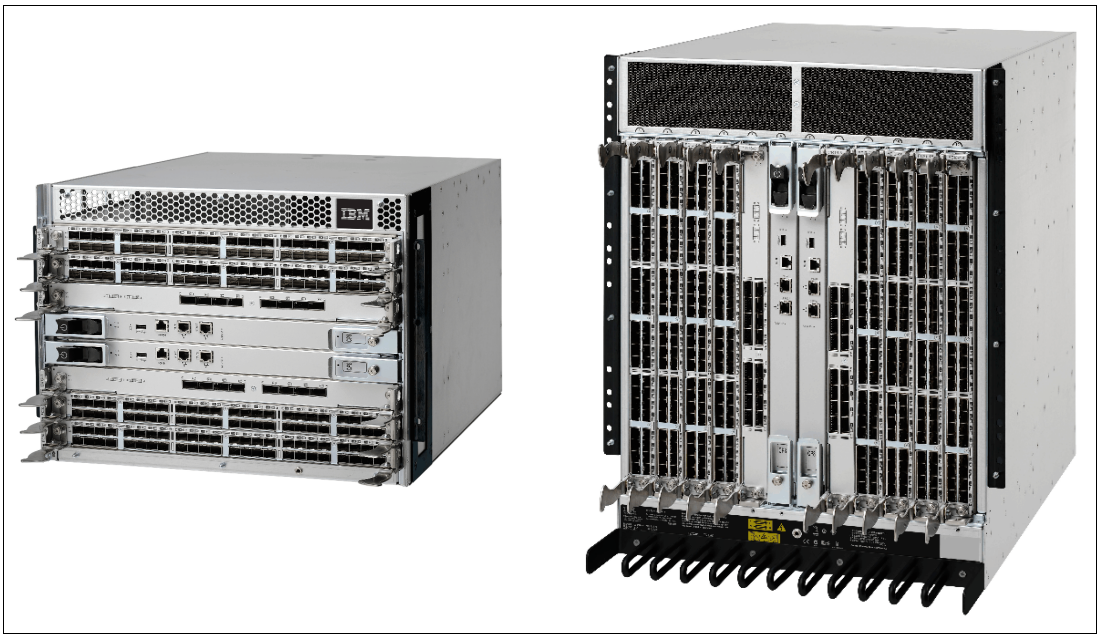

Figure 9-19 The IBM System Storage SAN384B-2 (left) and SAN768B-2 (right)

The IBM System Storage SAN384B-2 and SAN768B-2 offer the following hardware features:

- The SAN384B-2 chassis includes two control processor blades and two core module blades. It also has space for up to four port blades, dual power supplies, and two fan modules in an 8U rack height (plus 1U exhaust shelf).
- The SAN768B-2 chassis includes two control processor blades and two core module blades. It also has space for up to eight port blades and dual power supplies with an option for two more power supplies, and three fan modules in a 14U rack height.
- Up to ten SAN384B-2 or SAN768B-2 switches can be connected with high-speed UltraScale optical inter-chassis links.
- The following blades are available:
  - 16 Gbps 32-port and 48-port FC blades, capable of speeds of 2, 4, 8, 10, and 16 Gbps
  - 16 Gbps 64-port FC blade, with 16 Quad Small Form-factor Pluggable (QSFP) ports and Multiple-Fiber Push-On/Pull-off (MTP/MPO) connectivity
  - 8 Gbps 64-port FC blade, with special mini optical transceiver (mSFP) optics included
  - 8 Gbps Enhanced Extension blade that contains 22 ports that are capable of 8 and 4 Gbps FC connection or 1 Gbps copper Ethernet connection, and two 10 GbE ports
  - 8 Gbps 16-port FC Encryption blade
  - 10 Gbps 24-port FCoE blade (SAN768B-2 only)
- The following transceivers are available:
  - 16 Gbps, 10 Gbps, and 8 Gbps shortwave SFPs
  - 16 Gbps, 10 Gbps, and 8 Gbps 10 km (6.2 miles) longwave SFPs
  - 16 Gbps and 8 Gbps 25 km (15.5 miles) extended distance longwave SFPs

- In addition, the following transceivers are available for the Enhanced Extension blade and the 24-port FCoE blade:
  - 10 Gbps Ethernet SR and LR SFP+
- In addition, the following transceivers are available for the Enhanced Extension blade only:
  - 4 Gbps shortwave and 10 km longwave SFP
  - 1 Gbps copper SFP
- Full Fabric and Universal Port operation (E_Port, F_Port, and FL_Port)

For more information, see *IBM System Storage SAN768B-2 and SAN384B-2 Fabric Backbones*, TIPS1127.

Also, see the product website:

http://www.ibm.com/systems/storage/san/b-type/san768b-2/

Cisco MDS 9500 Series Multilayer Directors

The Cisco MDS9500 Series Multilayer Directors include these models:

- Cisco MDS9506 - model 2054-E04 with up to 192 ports
- Cisco MDS9513 - model 2054-E11 with up to 528 ports

Note: As of 19 April 2014, Cisco MDS9509 - model 2054-E07 was withdrawn from marketing and replaced by the new Cisco MDS9710, which is model 9710-E08.

The Cisco MDS 9500 Series Multilayer Directors support 1, 2, 4, 8, and 10 Gbps FC switch connectivity and intelligent network services. These services help improve the security, performance, and manageability that are required to consolidate geographically dispersed storage devices into a large enterprise SAN.

Both models of Cisco MDS9500 Series Multilayer Directors use two Supervisor-2A Modules that are designed for high availability and performance. The Supervisor-2A Module combines an intelligent control module and a high-performance crossbar switch fabric in a single unit. It uses FSPF multipath routing, which provides intelligence to load balance across a maximum of 16 equal-cost paths and to dynamically reroute traffic if a switch fails.

Each Supervisor-2A Module provides the necessary crossbar bandwidth to deliver full system performance in the Cisco MDS 9500 directors with more Advanced Fibre Channel Switching modules or Fabric3 Switching modules. It is designed to eliminate the impact on system performance of the loss or removal of a single crossbar module.

The Cisco MDS9500 MDS 9500 Series Multilayer Directors support the deployment of FCoE by using 10 Gbps Ethernet-SR X2 transceivers.

The Cisco MDS9500 Series Multilayer Directors are shown in Figure 9-20.

Figure 9-20 Cisco MDS 9500 Series Multilayer Directors: MDS9513 (left) and MDS9506 (right)

For more information, see the product website:

http://www.ibm.com/systems/storage/san/ctype/9500/

Cisco MDS 9706 Multilayer Director for IBM System Storage

Cisco MDS 9706 Multilayer Director for IBM System Storage is a director-class SAN switch that is designed for deployment in small to mid sized storage networks that can support enterprise clouds and business transformation. It layers a comprehensive set of intelligent features onto a high-performance, protocol-independent switch fabric

Using Cisco MDS 9700 Family switching modules, MDS 9706 supports up to 192 ports in a 6-slot modular chassis, with up to 768 ports in a single rack. Ports can be configured in these ways:

- Fibre Channel (2/4/8 Gbps, 4/8/16 Gbps, or 10 Gbps)
- FCoE (10 Gbps)
- A mix of both Fibre Channel and FCoE

MDS 9706 supports the same Fibre Channel and FCoE switching modules as Cisco MDS 9710 Director for IBM System Networking for a high degree of system commonality.

In addition to meeting the basic requirements of non-disruptive software upgrades and redundancy of all critical hardware components, MDS 9706 software architecture offers outstanding availability. MDS 9706 provides redundancy on all major hardware components, including the supervisor and fabric modules and the power supplies. Cisco MDS 9700 Series Supervisor Module automatically restarts failed processes, making MDS 9706 exceptionally robust.

Figure 9-21 shows the Cisco MDS 9706 Multilayer Director for IBM System Storage.

Figure 9-21 Cisco MDS 9706 Multilayer Director for IBM System Storage

For more information about the Cisco MDS 9706 Multilayer Director for IBM System Storage, see the following website:

`http://www.ibm.com/systems/storage/san/ctype/9706`

Cisco MDS 9710 Multilayer Director

The Cisco MDS 9710 Multilayer Director offers an enterprise-class solution for large virtualized data centers or cloud installations that require high-performance, scalable, and reliable low-latency SANs.

The Cisco MDS 9710 Multilayer Director supports 2, 4, 8, 10, and 16 Gbps FC switch connectivity and intelligent network services. These features help improve the security, performance, and manageability that are required to consolidate geographically dispersed storage devices into a large enterprise SAN. It also supports 10 Gbps FCoE connectivity in parallel with FC, offering a seamless transition to IP-based storage networking.

The Cisco MDS 9710 Multilayer Director uses two new hot-swappable, redundant Supervisor-1 Modules that are designed for high availability and performance. In addition, the director uses separate hot-swappable, redundant crossbar modules.

The Cisco MDS 9710 Multilayer Director is shown in Figure 9-22.

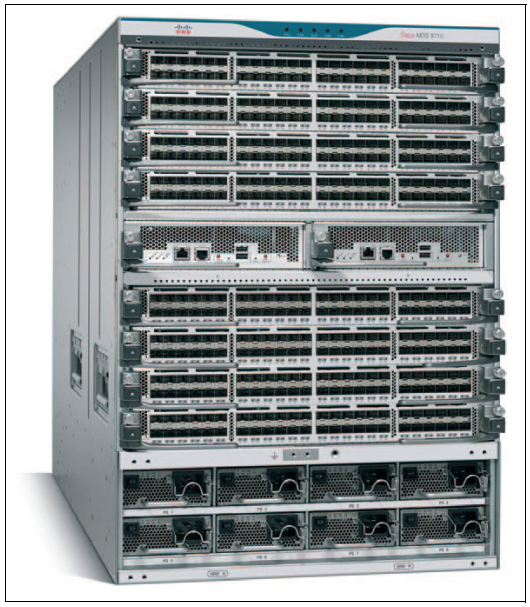

Figure 9-22 Cisco MDS 9710 Multilayer Director

For more information, see *Cisco MDS 9710 Multilayer Director for IBM System Networking*, TIPS1046, or the product website:

http://www.ibm.com/systems/storage/san/ctype/9700/

9.3.8 SAN specialty switches

The SAN switches offer special functions that relate to multi-protocol routing.

The following enterprise products are described:

- IBM System Storage SAN06B-R Extension Switch
- IBM System Storage SAN42B-R Extension Switch
- Cisco MDS 9250i Multiservice Fabric Switch for IBM System Storage

IBM System Storage SAN06B-R Extension Switch

A wide range of IBM System Storage midrange and enterprise SAN infrastructure simplification and business continuity solutions can be created with the IBM System Networking SAN06B-R Extension Switch. Infrastructure simplification solutions for the IBM Power Systems and System x families include disaster tolerance over metropolitan and global IP networks with IBM disk arrays, tape libraries, and IBM Tivoli Storage Manager data protection software. Separate SAN islands can also be consolidated by using FC routing. Support for z Systems servers is provided by the optional High-Performance Extension and FICON Control Unit Port (CUP) Activation features.

Local site infrastructure simplification solutions can be extended to one or more remote sites for enhanced data protection and disaster tolerance. The IBM System Storage SAN06B-R Extension Switch provides FCIP and FCIP Tunneling Service for distance extension, which can enable cost-effective and manageable metropolitan and global business continuity solutions. This extended distance connectivity can help create consolidated remote tape vaulting data protection, plus Metro Mirror and Global Mirror disk-based, disaster-tolerant solutions.

Since the introduction of SANs, clients built multiple SAN networks (or islands) for various applications, often with fabric switch components from various manufacturers. Certain islands were built by various departments within a company, whereas other islands resulted from mergers, acquisitions, or reorganizations. These islands, which were often created to isolate important applications or to implement equipment with various capabilities, constrain opportunities for enhanced infrastructure simplification and vital business continuity solutions.

The IBM System Networking SAN06B-R Extension Switch (model 2498-R06) provides the FC-FC Routing Service. This service allows the interconnection of multiple SAN islands without requiring that the separate fabrics are merged into a single large SAN. This capability can help create a tiered or extended enterprise SAN infrastructure without needing to redesign or reconfigure the entire environment.

The IBM System Networking SAN06B-R Extension Switch is shown in Figure 9-23.

Figure 9-23 The IBM System Networking SAN06B-R Extension Switch

For more information, see *IBM System Storage SAN06B-R Extension Switch*, TIPS1126, or the product website:

http://www.ibm.com/systems/storage/san/b-type/san06b-r/

IBM System Storage SAN42B-R Extension Switch

The IBM System Storage SAN42B-R switch is intended as a platform for FCIP. This configuration enables the transmission of FC data over long distances by way of IP networks by wrapping FC frames in IP packets. Each end of the FCIP communication path must be a compatible FCIP device. The switch can operate independently or in a fabric that contains multiple extension switches.

The SAN42B-R provides the following software features:

- Multiple logical FCIP tunnels with a maximum tunnel bandwidth up to 20 Gbps allow for scalable connectivity between sites.
- The FCIP Trunking feature allows multiple IP source and destination address pairs. These pairs are defined as *FCIP circuits,* through multiple 1/10 GbE or 40 GbE interfaces to provide a high-bandwidth FCIP tunnel and lossless failover resiliency. In addition, each FCIP circuit supports four QoS classes (Class-F, High, Medium, and Low Priority), each as a Transmission Control Protocol (TCP) connection.
- Hardware-based compression delivers the ability to maximize throughput over lower bandwidth links in the WAN, optimizing the cost efficiencies of FCIP. The SAN42B-R compresses FC frames before they are encapsulated into FCIP packets.
- Key protocol features are enabled in the FCIP implementation to optimize the performance of Extension over IP networks, including WAN Optimized TCP, 9 K jumbo frames, and end-to-end path maximum transmission unit (MTU) auto discovery.
- Hardware-based IPsec supports a mix of secure and non-secure tunnels on the same Ethernet port, jumbo frames, and VLAN-tagged connections. The IPsec function can support both IPv4 and IPv6.
- FastWrite, Open Systems Tape Pipelining, and Advanced Accelerator for FICON mitigate the latency effect of a long-distance FCIP distance connection over an IP WAN.
- Built-in WAN link tester generates traffic over an IP connection to test for maximum throughput, congestion, loss percentage, out of order deliver, latency, and other network conditions. It helps determine the health of a WAN link before you deploy it for use.
- Fabric Vision advanced monitoring provides the following functions:
  - Policy-based monitoring monitors FCIP connectivity and WAN anomalies by using multi-layer metrics.
  - Flow monitoring reports I/O operations per second (IOPS) and the data rate of individual I/O flows of inter-data center (DC) replication and tape backup operations.

- Flow generator generates FC frames for a defined flow with a default or custom size and pattern, and sends them across an FCIP tunnel to help validate end-to-end network setup and configuration.

- Fabric OS delivers distributed intelligence throughout the network and enables a wide range of applications.

The IBM System Storage SAN42B-R switch is available in three configurations, as shown in Table 9-1.

Table 9-1 IBM System Storage 42B-R configurations

Product configuration	FC ports	Ethernet ports	WAN rate limiting	Approximate application throughput
Base configuration	24 x 16 Gbps	16 x 1/10 Gbps	5 Gbps	15 Gbps
Medium configuration (Base plus WAN Rate Upgrade 1)	24 x 16 Gbps	16 x 1/10 Gbps	10 Gbps	30 Gbps
Maximum configuration (Base plus WAN Rate Upgrade 1 plus WAN Rate Upgrade 2)	24 x 16 Gbps	16 x 1/10 Gbps 2 x 40 Gbps	Unlimited	80 Gbps

Note: The approximate application throughput values in Table 9-1 assume that a degree of data compression is occurring. However, IBM makes no promises, guarantees, or any indication that a level of compression is possible for client-specific data. Certain data is highly compressible and other data cannot be compressed. The amount of application throughput varies depending on data compressibility and the selected compression mode.

The IBM System Storage SAN42B-R Extension Switch is shown in Figure 9-24.

Figure 9-24 IBM System Storage SAN42B-R Extension Switch

The IBM System Networking SAN06B-R Extension Switch offers the following highlights:

- Rack-mountable 2U chassis
- Base configuration includes a comprehensive set of advanced services, including Fabric Vision, Extension Trunking, Adaptive Rate Limiting, ISL Trunking, IPsec, Compression, and Extended Fabric
- Optional licenses are available for Integrated Routing, FICON CUP, and Advanced FICON Accelerator

- Built-in encryption and compression of storage data flows over long-distance links with no performance penalty
- Extension Trunking increases WAN use and protects against WAN link failures

For more information, see *IBM System Storage SAN42B-R Extension Switch*, TIPS1209, or the product website:

http://www.ibm.com/systems/storage/san/b-type/san42b-r/

Cisco MDS 9250i Multiservice Fabric Switch for IBM System Storage

The Cisco MDS 9250i Multiservice Fabric Switch, the next generation of the Cisco MDS 9200 Series Multiservice Switches, is an optimized platform for deploying high-performance SAN extension solutions, distributed intelligent fabric services, and cost-effective multiprotocol connectivity for both open systems and mainframe environments.

With a compact form factor and advanced capabilities that are typically available only on director-class switches, the Cisco MDS 9250i Multiservice Fabric Switch is an ideal solution for departmental and remote branch-office SANs and in large-scale SANs with the Cisco MDS 9710 Multilayer Director.

The Cisco MDS 9250i Multiservice Fabric Switch offers up to 40 16-Gbps FC ports, two 1/10 Gigabit Ethernet IP storage services ports, and eight 10-Gigabit Ethernet FCoE ports in a fixed 2RU form factor.

The Cisco MDS 9250i Multiservice Fabric Switch is shown in Figure 9-25.

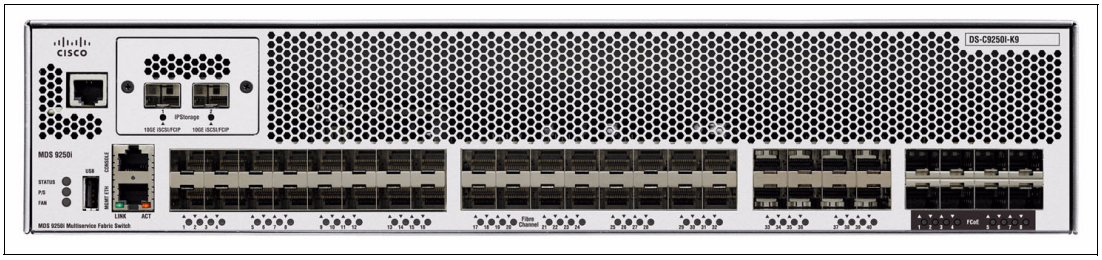

Figure 9-25 Cisco MDS 9250i Multiservice Fabric Switch

For more information, see *Cisco MDS 9250i Multiservice Fabric Switch for IBM System Storage*, TIPS1167, or the product website:

http://www.ibm.com/systems/storage/san/ctype/9250i/

9.3.9 SAN solutions

IBM SAN solutions integrate IBM and other vendor server, storage, SAN switches, and software components into solution templates that are extensively tested to provide high availability, scalability, security, and simplicity of management. IBM SAN solutions are offered with worldwide IBM service and end-to-end solution support. These solutions templates might be customized by IBM Business Partners or IBM Global Services to address individual customer requirements.

For more information, see the following website:

http://www.ibm.com/systems/storage/san/index.html

9.4 Selecting the best alternative

Which storage networking alternative is the best for a specific organization might be obvious based on organizational objectives, current storage infrastructure, and what the alternatives provide. Or it might be an open question. Storage technology has clearly become more varied and sophisticated, and the decisions are more complex than ever. Choice means flexibility, and that is good, but which choice to make is not always clear:

- If a group of individual users with PCs needs to share disk storage capacity and also share files in that storage, an NAS might be easiest to install and manage.
- If application servers need to share disk storage and are each accessing independent (block I/O) databases, an FC-based SAN might be appropriate.
- For a few servers where no SAN exists, iSCSI might be less expensive and less complex.
- For many servers and greater distances, FCoE can be used over existing Ethernet infrastructures. However, it can add increased complexity.
- If required, FCIP or iFCP provides a cost-effective way to achieve business protection, enabling solutions, such as remote tape archiving.

9.5 More information

For more information about all of the SAN-related topics and the products that are described in this chapter, see the following publications:

- *Introduction to Storage Area Networks*, SG24-5470
- *IBM b-type Gen 5 16 Gbps Switches and Network Advisor*, SG24-8186
- *IBM and Cisco: Together for a World Class Data Center*, SG24-8105
- *SAN Storage Performance Management Using Tivoli Storage Productivity Center*, SG24-7364
- *Implementing an IBM b-type SAN with 8 Gbps Directors and Switches*, SG24-6116
- *IBM/Cisco Multiprotocol Routing: An Introduction and Implementation*, SG24-7543
- *Implementing an IBM/Cisco SAN*, SG24-7545

For more information, see the following website:

http://www.ibm.com/systems/storage/san/index.html

A

Withdrawn products

This appendix lists the most recently withdrawn products in the IBM storage family.

Table A-1 lists IBM Storage product withdrawals for the two years before the writing of this book.

Table A-1 Storage products that are withdrawn

Description	Type	Model	Withdrawal date	Announcement number
Disk				
IBM Storwize V7000	2076	124, 224, 324	Jan 13, 2015	915-001
IBM Storwize V7000	2076	124, 224, 324	Jan 13, 2015	915-001
IBM Flex System® V7000 (AAS)	4939	A29. A49	Jan 17, 2015	914-185
IBM Flex System V7000 (AAS)	4939	A29. A49	Jan 17, 2015	914-185
IBM Scale Out Network Attached Storage	2851	RXA, RXB, RXC, SS2, DR2, DE2	March 10, 2015	915-058
IBM Scale Out Network Attached Storage	2851	RXA, RXB, RXC, SS2, DR2, DE2	March 10, 2015	915-058
FlashSystem				
FlashSystem 840, FlashSystem V840, and select features	9840, 9843 9846, 9848	AE1 AC1, AE1	May 19, 2015	915-131

Description	Type	Model	Withdrawal date	Announcement number
Tape				
Crossroads ReadVerify Appliance	3222	RV1	Jul. 14, 2015	915-074
Tape Controller Model C07	3592	C07	Nov 13, 2015	915-168
TS1130 tape drive	3592	E06	Apr 29, 2016	915-150
TS1040 Ultrium 4 tape drive	3588	F4A	Apr 29, 2016	915-150
Ultrium 4 Fibre Channel tape drive	3576	L5B, E9U	Apr 29, 2016	915-150
Ultrium 4 LVD SCSI Drive	3573	L2U, L4U	Apr 29, 2016	915-150
Ultrium 4 Fibre Drive	3573	L2U, L4U	Apr 29, 2016	915-150
Ultrium 4 SAS Drive	3573	L2U, L4U	Apr 29, 2016	915-150
Storage area networks				
Cisco MDS 9222i for IBM System Storage	2054	E01	Jan 15, 2015	915-022
Cisco MDS 9148 for IBM System Storage	2417	C48	Jun 30, 2015	915-138
IBM RackSwitch G7028	7120	24L	June 30, 2015	915-079
Storage software				
IBM Tivoli Storage Manager FastBack® IBM Tivoli Storage Manager FastBack for Microsoft Exchange IBM Tivoli Storage Manager FastBack for Bare Machine Recovery	5641	A01, A02, A03	Apr 14, 2015	915-114

For more information about any product withdrawal, see the Withdrawal Announcement Letter. For each product withdrawal, Table A-1 on page 265 provides the announcement number in the last column. To access the announcement letter by the number, see the Offering Information website:

http://www.ibm.com/common/ssi/index.wss

Insert the number in the “Keywords” field in the simple search view or directly in the “Letter number” field in the advanced search view.

Related publications

The publications listed in this section are considered particularly suitable for a more detailed discussion of the topics covered in this book.

IBM Redbooks

The following IBM Redbooks publications provide additional information about the topic in this document. Note that some publications referenced in this list might be available in softcopy only.

- *A Deployment Guide for IBM Spectrum Scale Object*, REDP-5113
- *Deploying IBM Spectrum Accelerate on Cloud*, REDP-5261
- *DS8000 Performance Monitoring and Tuning,* SG24-8318
- *FlashSystem 900 Product Guide*, TIPS1261
- *IBM DS8880 Architecture and Implementation (Release 8)*, SG24-8323
- *IBM DS8870 Copy Services for IBM z Systems,* SG24-6787
- *IBM DS8870 Copy Services for Open Systems*, SG24-6788
- *IBM DS8870 Disk Encryption*, REDP-4500
- *IBM FlashSystem V9000 Product Guide*, REDP-5317
- *IBM FlashSystem V9000 in a VersaStack Environment*, REDP-5264
- *IBM Hyper-Scale in XIV Storage*, REDP-5053
- *IBM Private, Public, and Hybrid Cloud Storage Solutions*, REDP-4873
- *IBM Spectrum Accelerate Deployment, Usage, and Maintenance*, SG24-8267
- *IBM System Storage DS8000: Host Attachment and Interoperability*, SG24-8887
- *IBM Software-Defined Storage Guide*, REDP-5121
- *IBM Tape Library Guide for Open Systems*, SG24-5946
- *IBM XIV and VMware Synergy with IBM Spectrum Control Base Edition*, REDP-5131
- *IBM XIV Gen3 with IBM System Storage SAN Volume Controller and Storwize V7000*, REDP-5063
- *IBM XIV Storage System Gen3 Architecture, Implementation, and Usage*, SG24-7659
- *IBM XIV Storage System Business Continuity Functions*, SG24-7759
- *IBM XIV Storage System: Host Attachment and Interoperability*, SG24-7904
- *Implementing IBM FlashSystem 900*, SG24-8271
- *Implementing the IBM Storwize V5000*, SG24-8162
- *Introducing and Implementing IBM FlashSystem V9000*, SG24-8273
- *IBM Assist On-site for Storage Overview*, REDP-4889
- *Solid-State Drive Caching in the IBM XIV Storage System*, REDP-4842
- *Using XIV in OpenStack Environments*, REDP-4971
- *XIV Storage System in a VMware Environment*, REDP-4965

You can search for, view, download or order these documents and other Redbooks, Redpapers, Web Docs, draft and additional materials, at the following website:

ibm.com/redbooks

Online resources

These websites are also relevant as further information sources:

- IBM Assist On-Site:

  http://www.ibm.com/support/assistonsite/
- IBM Global Technology Services:

  http://www.ibm.com/services/us/en/outsourcing/index.html
- IBM Storage:

  http://www.ibm.com/systems/storage
- IBM System Storage interactive product guide:

  http://www.ibm.com/common/ssi/cgi-bin/ssialias?htmlfid=TSO00364USEN
- IBM System Storage Interoperation Center (SSIC):

  http://www.ibm.com/systems/support/storage/config/ssic

Help from IBM

IBM Support and downloads

ibm.com/support

IBM Global Services

ibm.com/services

Redbooks

IBM System Storage Solutions Handbook

SG24-5250-11

ISBN 0738441740

IBM

(0.5” spine)
0.475”<->0.873”
250 <-> 459 pages

SG24-5250-11

ISBN 0738441740

Printed in U.S.A.