

**Taming the explosion of
information with
IBM storage solutions**



Making the most of more information

As the world becomes more interconnected, instrumented and intelligent, the amount of information generated by the people, technologies and objects that populate the planet is swelling to unprecedented proportions. Digital data is being generated in mind-boggling amounts—15 petabytes, or more than eight times the information contained in all U.S. libraries, every single day.¹ This growth shows no signs of stopping. Data volumes in general are doubling every 18 to 24 months.² At least one industry analyst expects annual worldwide information growth rates of 50 to 100 percent to be the norm as businesses and economies recover.³ And in just three years, the world's annual IP traffic is expected to total more than half a zettabyte.⁴ (That's a trillion gigabytes—or a one followed by 21 zeroes.)

Dealing with that amount of data would be challenging enough; what makes it even more daunting is that information isn't just growing, it's changing. Organizations are not only dealing with growth in structured data that is neatly stored in fields in databases. They are increasingly dealing with unstructured data, too. Unstructured data comes from sources like emails, documents, spreadsheets, blogs, web pages, images, audio files and video files—and it now constitutes up to 80 percent of data.⁵ Because it doesn't conform to traditional formats, it requires someone or something to interpret it in order for it to be transformed into actionable information that organizations can use.

Those who are charged with managing this data face a complex challenge—one that simultaneously raises issues of how to keep the cost of data management low while keeping the reliability of information access high. What's required is an information infrastructure that helps reduce complexity and cost while at the same time providing reliable access to trusted, up-to-date data. That isn't easy: According to a recent survey conducted by IBM, half of business leaders said they don't have sufficient information from across their organization to do their jobs; and one-third said they frequently make major decisions with information they don't trust, or with incomplete information.⁶ Only with the right infrastructure in place will the right information be reliably available to the people who need it to make the right business decisions.

“It doesn't matter if the economy is good or the economy is awful, data growth never abates. It is the one constant thing we deal with.”

—Steve Duplessie, Senior Analyst, Enterprise Strategy Group

Meeting the challenge with IBM Information Infrastructure solutions

IBM offers an innovative approach to help manage information growth more effectively, with an information infrastructure that efficiently, securely and cost-effectively stores and protects information and optimizes information access. IBM Information Infrastructure solutions can help in three key areas: managing the information explosion, moving and managing information, and adding capacity without introducing complexity.

Managing the information explosion

To better manage more—and more kinds of—information, IBM Information Infrastructure solutions include storage systems that are optimized for unstructured data and unpredictable workloads. IBM offers the top technologies that enable storage optimization, including solid-state storage optimization, virtualization, automated failover, and centralized management. IBM makes it easier to manage archiving and retention operations with integrated disk and tape archives, progressive incremental backup, and policy-based retention. IBM deduplication offers a high level of efficiency that can shrink backup time requirements significantly. And IBM deduplication, thin provisioning, virtualization, space-efficient point-in-time copy, compression, and storage resource management solutions can help organizations lower physical storage requirements by up to 50 percent.⁷

Intelligently moving and managing information

Today's smarter systems require continuous, reliable access to information, wherever it resides. To help ensure this, the IBM Information Infrastructure includes optimized solutions that enable the smart movement and management of information across storage tiers. These solutions are easy to configure and manage, allowing organizations to optimize storage price and performance without adding administrative complexity. When organizations are able to intelligently move data between storage solutions, they gain specific reliability and cost benefits:

- **Higher reliability:** The ability to shift workloads without disrupting users means fewer disruptions in service, and the ability to move active workloads to solid-state service optimizes performance.
- **Lower costs:** By migrating older information to more efficient, lower-cost storage systems, organizations can reduce storage costs. They can also control power and cooling requirements while expanding throughput and capacity, by moving active data to solid-state storage and inactive data to low-cost SATA or tape.

Capacity growth without added complexity

As volumes of data grow, it ultimately becomes too complex and costly to respond by continuing to add storage space and administration staff. IBM enables organizations to manage more capacity with less administration—even in environments that include storage from multiple vendors—with storage virtualization and storage management solutions that allow administrators to manage the infrastructure using familiar tools. IBM supports advanced capabilities for non-IBM storage, including online migration, point-in-time copy, and performance management. And IBM Scale Out Network Attached Storage (SONAS) delivers the capacity of multiple NAS devices with the management requirements of one. IBM solutions optimize storage so that organizations have less data to manage, higher disk utilization rates, and a smaller footprint for the storage and backup environment.



Smarter systems for a Smarter Planet

Meeting the need for storage optimization and information retention

Organizations seeking to upgrade their information infrastructures in phases will often focus on one of two broad categories: storage optimization, or information retention. Storage optimization initiatives help improve performance and service levels for production workloads. Information retention initiatives help manage long-term storage functions, such as backup and archive, more effectively.

Storage optimization: Reducing costs and improving performance

Storage optimization is essential for keeping costs down and performance up in primary data environments such as production databases, applications and file systems. IBM storage optimization solutions help meet growing throughput, scalability and reliability requirements with innovative technology, proven best practices and knowledgeable experts. These solutions enable organizations to:

- Improve storage price/performance by increasing disk utilization and leveraging tiered storage for performance optimization.
- Reduce operating expenses by consolidating data onto new, more efficient storage systems.
- Sustain service levels as the amount of information grows by automating critical storage management processes.

Information archiving and retention: Managing data throughout its useful life

The costs of managing and storing information can be lowered by simply eliminating duplicate, irrelevant or expired information. IBM can help organizations develop a smarter archiving strategy that enables them to retain only information that is relevant—based on business policy—and delete the rest.

By coupling smarter archiving with a smarter backup strategy, organizations can shorten the backup process, optimize the storage infrastructure and reduce storage management costs. IBM archiving and retention solutions specifically enable organizations to:

- Know what information to keep and what to delete.
- Store information safely and affordably for decades.
- Retrieve information after long archive periods, even if the logical format and physical media have changed.
- Prevent duplicate or inactive information from slowing down production systems.
- Move inactive or rarely used data out of costly top-tier storage.
- Keep audit information safe, unaltered and undeleted for the length of its required retention period.

“The combination of IBM’s storage technology, information management products, aggressive financing, and best-of-breed integrated services supported by world-class expertise and proven experience, provide the building blocks for the world’s strongest information infrastructure portfolio.”

—B. R. Allen Associates, “Creating a Dynamic Information Infrastructure,”
September 2009

Information infrastructure solutions from the leader in storage innovation

IBM is the world’s largest storage and data services organization, with experts around the world, a broad ecosystem of skilled IBM Business Partners, and service products to help speed time to value. IBM has practical experience providing storage for some of the world’s largest and most demanding environments, including its own information infrastructure, which supports 400,000 employees around the world, large outsourced data centers, and IBM cloud services. IBM leads the industry in a number of key areas.

Consolidating and moving data with storage virtualization

Storage virtualization provides the ability to pool storage systems into consolidated, shared capacity that can be managed from a central point of control. IBM is the industry leader in virtualization technology, even developing the technology to virtualize EMC and other disks.

With the IBM market-leading technology for storage area network (SAN) virtualization, IBM System Storage® SAN Volume Controller enables higher overall utilization than individual storage units can attain.

Virtualization helps decrease administration costs while increasing flexibility, disk utilization and throughput. It improves performance, makes storage easier to manage, and simplifies migration. Storage virtualization also complements server virtualization—and IBM can help organizations manage both, with a comprehensive portfolio that includes support for all types of storage virtualization, including tape, disk and file systems.

Reducing costs with tiered storage

IBM Information Infrastructure solutions can help organizations reduce costs by using policy-based data movement to get data out of costly Tier 1 storage and into less expensive Tier 2 storage. As data volumes grow, the financial pressure grows to

reserve Tier 1 for more active data, while migrating less active data to Tier 2—so that active data responds faster, and inactive data costs less to store.

IBM tiered disk solutions help optimize data decision making and migration processes, which can help organizations get their data onto the right tier with less effort. IBM can help create a tiered storage infrastructure that integrates all tiers of storage, from solid-state disk to tape, to provide more flexible choices for matching price and performance requirements.

Containing the information explosion with improved storage efficiency

IBM data reduction solutions include deduplication, space-efficient point-in-time copy, compression, and storage resource management solutions. IBM capabilities to improve disk utilization include thin provisioning and disk virtualization. Storage efficiency solutions can help organizations lower physical storage requirements by up to 50 percent.⁷

Data deduplication works by helping to eliminate redundant data, which lessens storage demands. IBM ProtecTIER® deduplication addresses the biggest source of duplicate data—backups—with efficient, scalable solutions that won't disrupt existing backup scripts and processes. IBM solutions can shrink backup windows significantly, speed most restores, reduce media costs and reduce data storage capacity demands.

IBM ProtecTIER deduplication solutions use patented algorithms that can specifically reduce the amount of space required for storage in the backup environment by up to a factor of 25, according to IBM tests, with enterprise-class performance of over 1,000 MB/sec. The capacity reduction that results from data deduplication is often expressed as a ratio—essentially the ratio of nominal data to the physical storage used. A 10:1 ratio, for example, means data is reduced by a factor of 10, and can be stored in 90 percent less space.

Expanding delivery choices with public and private cloud storage options

IBM offers multiple flexible delivery choices including remote managed services, cloud computing, and standard delivery options. Unlike many other vendors, IBM offers enterprise-ready storage cloud solutions for storage capacity on demand and archive. And IBM helps clients prepare for cloud computing with practical consulting, foundation technology, and integrated private cloud solutions.

Cloud computing has the potential to make an enormous financial impact on an organization by reducing IT labor costs for configuration, operations, management and monitoring, improving capital utilization, reducing provisioning cycle times, and reducing end-user IT support costs. Storage services such as capacity on demand, backup and archive are among the easiest to adapt to the cloud computing model. IT storage teams commonly provide storage as a service to other departments, so best practices and service level agreements have become somewhat standardized over time.

IBM helps organizations implement public, private and hybrid cloud storage services, with expert consulting and workload-optimized systems. IBM specifically offers three types of cloud solutions for storage and other services:

- IBM Smart Business on the IBM Cloud, with standardized services provided on a pay-per-use basis
- Smart Business Cloud services, or private cloud services behind an organization's firewall
- Smart Business Systems, which are purpose-built, integrated service delivery platform solutions, ready for rapid private cloud deployment

Improving resilience and availability with storage replication

Traditional backups take time, and restores take even longer—which can be a problem when continuous access to data is critical. IBM helps clients implement more resilient systems with advanced disk mirroring that can failover to alternate sites without disrupting user operations. MetroMirror and GlobalMirror features in IBM System Storage DS8000® and SAN Volume Controller can be used to ensure failover to an alternate data center within as little as three to five seconds, enabling near-continuous operations. IBM also offers IBM Tivoli® Storage Manager FastBack, an innovative solution that saves data to another system as the data is modified, rather than waiting for a traditional backup to be performed. Users therefore no longer have to wait for restores to complete, enabling near-immediate access to data in an emergency.

For more information

To learn more about how the IBM Information Infrastructure can help keep application reliability up and storage costs down, contact your IBM representative or IBM Business Partner, or visit: ibm.com/information_infrastructure

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



© Copyright IBM Corporation 2011

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

Produced in the United States of America
February 2011
All Rights Reserved

IBM, the IBM logo, and ibm.com are trademarks or registered trademarks of International Business Machines Corporation in the United States, other countries, or both. If these and other IBM trademarked terms are marked on their first occurrence in this information with a trademark symbol (® or ™), these symbols indicate U.S. registered or common law trademarks owned by IBM at the time this information was published. Such trademarks may also be registered or common law trademarks in other countries. A current list of IBM trademarks is available on the web at "Copyright and trademark information" at ibm.com/legal/copytrade.shtml

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

This document could include technical inaccuracies or typographical errors. IBM may not offer the products, services or features discussed in this document in other countries, and the product information may be subject to change without notice. Consult your local IBM business contact for information on the product or services available in your area. Any statements regarding IBM's future direction and intent are subject to change or withdrawal without notice, and represent goals and objectives only. The information contained in this document is current as of the initial date of publication only and is subject to change without notice. All performance information was determined in a controlled environment. Actual results may vary. Performance information is provided "AS IS" and no warranties or guarantees are expressed or implied by IBM. Information concerning non-IBM products was obtained from the suppliers of their products their published announcements or other publicly available sources. Questions on the capabilities of the non-IBM products should be addressed with the suppliers. IBM does not warrant that the information offered herein will meet your requirements or those of your distributors or customers. IBM provides this information "AS IS" without warranty. IBM disclaims all warranties, express or implied, including the implied warranties of noninfringement, merchantability and fitness for a particular purpose or noninfringement. IBM products are warranted according to the terms and conditions of the agreements under which they are provided.

- ¹ "2008 Annual Review: Intelligent Information," Thomson Reuters. <http://ar.thomsonreuters.com>
- ² "Data Volume is Becoming Unmanageable, Say Executives," Government Technology News Report, August 5, 2008. www.govtech.com/gt/articles/385068
- ³ Babineau, Brian, "IBM Information Infrastructure Initiative Tames the Information Explosion," Enterprise Strategy Group, April 2009. <ftp://ftp.software.ibm.com/common/ssi/sa/wh/n/oil03009usen/OIL03009USEN.PDF>
- ⁴ "Hyperconnectivity and the Approaching Zettabyte Era," Cisco, June 2009. www9.cisco.com/en/US/solutions/collateral/ns341/ns525/ns537/ns705/ns827/VNI_Hyperconnectivity_WP.html
- ⁵ IBM estimate.
- ⁶ LaValle, Steve, "Business analytics and optimization for the intelligent enterprise," IBM Institute for Business Value, December, 2009. <ftp://public.dhe.ibm.com/common/ssi/pm/xb/n/gbe03211usen/GBE03211USEN.PDF>
- ⁷ Hamano, John, "Introducing IBM N series – UA-SIS Deduplication solution," IOD Conference, October 2007, as one example.



Please Recycle