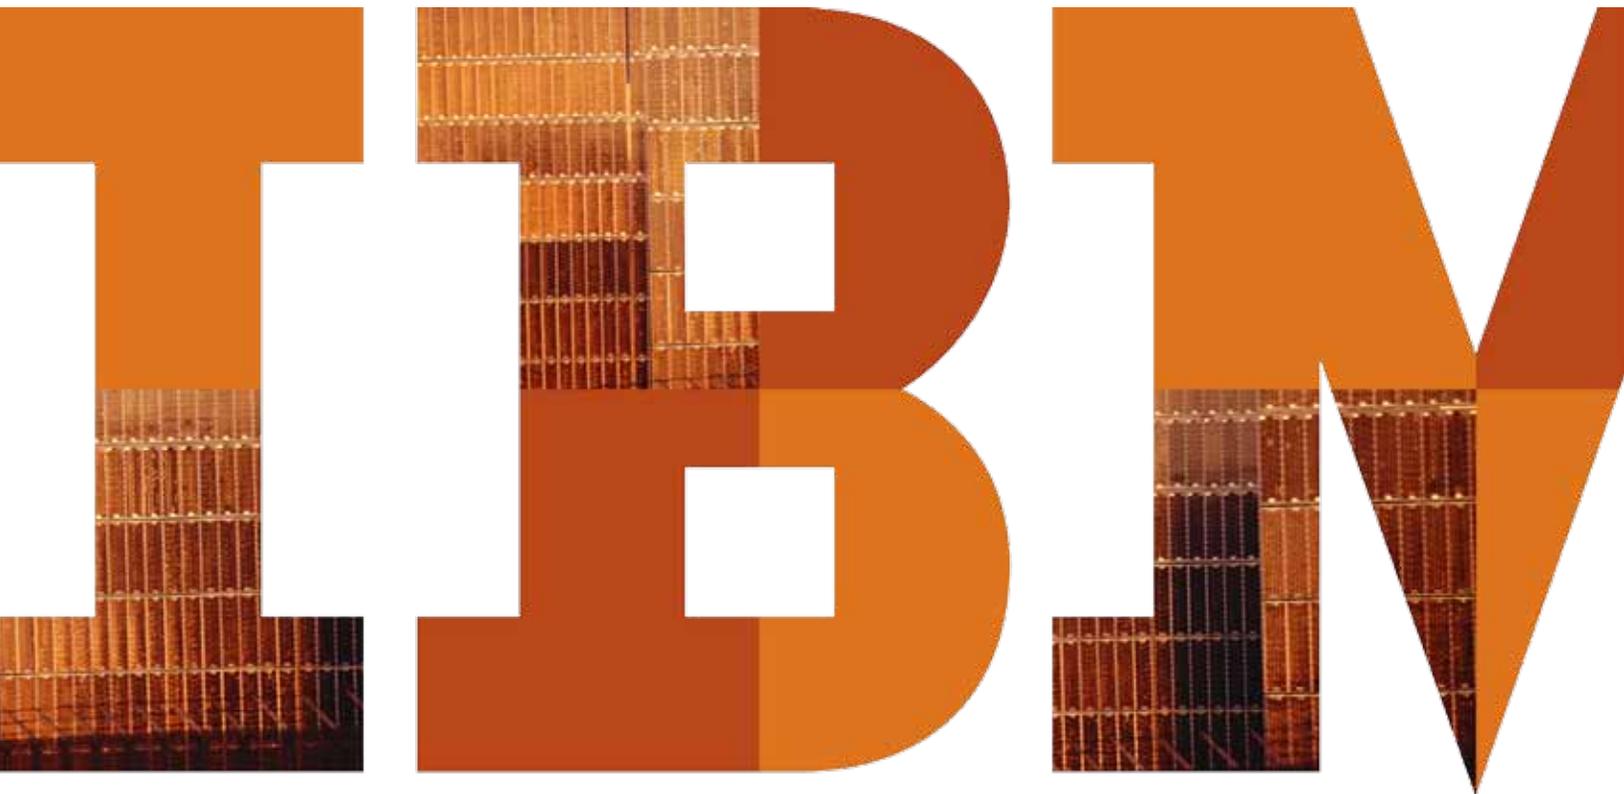


Integrated systems for operational analytics

*Discover insights and predict outcomes quickly with
integrated analytics, business intelligence and data
warehousing technology*



Contents

- 2 Executive summary
- 2 Rising volume of information creates a need for analytics
- 3 Criteria for a successful analytics solution
- 5 Building on a warehouse foundation
- 9 Business intelligence
- 9 IBM portfolio of integrated systems for operational analytics
- 11 Business benefits of integrated systems
- 11 Find actionable insights in mountains of data

Executive summary

IT infrastructures must evolve to cope with the exploding volume, variety and velocity of information available today. Through the smart use of analytics, organizations can capitalize on this data to uncover new opportunities, predict customer behaviors and business outcomes, and generate actionable, real-time insights to enhance business performance.

CIOs recognize that analytics and business intelligence (BI) can help their organizations meet emerging challenges and improve competitive positioning. But to maximize the benefits of analytics, they need solutions that will quickly deliver value, reduce administrative complexity and control the total cost of ownership (TCO). A proven approach is to select a pre-integrated solution that is optimized for analytics workloads.

Integrated systems for operational analytics from IBM are complete, ready-to-deploy, high-performance solutions that combine the powerful capabilities of analytics, BI and data warehousing. The fully integrated, workload-optimized design helps accelerate the time-to-value, enabling organizations to generate new insights rapidly. By providing a consolidated infrastructure and simplifying management, these systems also help lower the cost of deriving value from enterprise information. Available in multiple configurations and featuring modular architectures, these systems enable organizations to easily add capacity and adapt to changing business needs.

Rising volume of information creates a need for analytics

Data volumes are rising rapidly, and there are no signs that the flow of data will ebb any time soon. Organizations today collect data from every aspect of their operations, ranging from customer transactions to extended supply chain interactions. As these organizations continue to implement and access a greater number of instrumented, interconnected and intelligent systems, the volume, variety and velocity of data will grow.

Managers understand the importance of this data, but they need better tools to analyze it and distill its value. Implementing more powerful analytics tools can help managers address the information overload, find timely answers to specific questions and get ahead of the curve by predicting outcomes and developing timely, actionable insights. Adopting fully integrated solutions can help them realize those results rapidly.

The most valuable analytics solutions are those that empower individuals throughout the organization. These solutions transform analysis from a back-office activity for a handful of experts to an approach that can provide pervasive, predictive, real-time information for more decision makers at the point of impact. The focus of the entire organization moves from “sense and respond” to “predict and act.” The result: rapid, informed and confident decisions and actions that are based on consistent, trusted information.

Critical capabilities for supporting advanced analytics

- **Search and query:** To find and analyze relationships among a broad set of query workloads, data sizes and types, a solution must spread the workload over multiple systems with fast processors.
 - **Predictive analytics:** To predict the best answers after querying one or more enormous databases, a solution needs systems with large amounts of memory and the ability to perform complex calculations.
 - **Risk analysis:** To conduct compute-intensive risk analysis, a solution requires highly scalable accelerators optimized for the computational characteristics of the risk or hedge model being analyzed.
-

Analytics can benefit organizations across multiple industries. In government, analytics can help agencies deliver more convenient and reliable services to citizens. In banking, analytics can help guide business responses to consolidation and new business models. Telecommunications companies can use analytics to anticipate subscriber churn and improve the customer experience.

Although the benefits of analytics are clear, the technology path to reaching them is not. Many current analytics solutions leave customers frustrated or disappointed because they are too complex, cannot handle enterprise-level data volumes and take too long to deploy. Integrated systems for operational analytics from IBM help address these potential stumbling blocks, offering powerful, complete and integrated solutions that are optimized for analytics workloads.

Criteria for a successful analytics solution

What constitutes a successful analytics solution? In surveying analytics customers, IBM found several common criteria:

- **Consolidation of all necessary components into a single system.** The solution should be designed for enterprise analytics, offer pay-as-you-grow modular deployment and deliver world-class performance and reliability.
- **Fast time-to-value.** It must minimize start-up time by being pre-built for rapid deployment, offer high out-of-the-box performance and be pretested and preconfigured for customers' unique needs.
- **Reduced cost of ownership and streamlined system management.** The solution should require fewer resources in terms of storage, servers and floor space; enable IT to quickly add user and data capacity to meet growth requirements; and deliver reduced integration costs.
- **Immediate access to information and analytics.** Preconfigured BI modules should offer flexible functionality and delivery style; include advanced data discovery through data mining and text analytics; and deliver optimized and integrated online analytical processing (OLAP) capability.

IBM: A pioneer and leader in analytics solutions

IBM has long recognized the value of analytics and has a rich history in the database and warehouse technologies that serve as the foundation for analytics solutions. With each successive generation of analytics products, IBM has applied lessons learned to improve the ability of solutions to deliver actionable insights. Over the past decade, IBM also has added best-of-breed technologies to enhance its analytics portfolio by acquiring key companies and products, including Netezza, Ascential Software, Trigo, Cognos, SPSS, Vivisimo and others.

Integrated systems for operational analytics integrate a strong data warehouse foundation and broad analytics capabilities with IBM servers and storage to provide an end-to-end analytics solution. They are built with the latest IBM server and storage technologies, and enable organizations to select from a range of chipsets, storage options and performance levels. These solutions are part of a long-term IBM roadmap designed to continue producing solutions that respond to changing customer needs.

In creating these systems, IBM also incorporated the idea of Balanced Configuration Units (BCUs)—optimized allocations of hardware and software that can help speed deployment and reduce risks. The Integrated Systems for Operational Analytics feature a module-based format that allows organizations to incorporate functional add-ons of complementary software to customize the solution.

IBM® PureSystems™ offerings continue the IBM tradition of providing fully integrated solutions that draw on best practices developed from thousands of successful customer implementations. These expert integrated systems—which include the IBM PureData™ System for Operational Analytics—are designed to help organizations reduce complexity, accelerate time-to-value and control costs.

The IBM PureData System for Operational Analytics provides deployment patterns that enable repeatable, self-service deployment of a data warehouse within minutes. Drawing on deep expertise amassed by IBM and IBM Business Partners, deployment patterns are templates that define what software components are required and how they should best be configured to meet particular needs. Administrators can use the patterns to speed deployment of the data warehouse and help ensure consistency across multiple deployments. Organizations can choose from a range of existing patterns, easily customize those patterns or develop their own with tools available from IBM.

Building on a warehouse foundation

At the core of the integrated systems is IBM InfoSphere® Warehouse—a unified data warehouse that delivers access to structured and unstructured information in real time. InfoSphere Warehouse is powered by IBM DB2®, which is optimized to deliver industry-leading performance across multiple workloads while helping to lower administration, storage, development and server costs.

InfoSphere Warehouse offers a range of features and capabilities that help enhance the flexibility of the data warehouse and increase management efficiency:

- **Data partitioning:** InfoSphere Warehouse supports data partitioning, giving organizations multiple ways to distribute data across servers for large-scale parallelism and linear scalability. Because InfoSphere Warehouse can physically cluster data on multiple dimensions, order data by value range and limit input/output (I/O) to relevant data partitions, it can reduce the work needed to resolve many queries.
- **Compression:** Storage optimization technology can help organizations significantly reduce disk space requirements and improve query performance.
- **Workload management:** Organizations gain low-overhead, fine-grained control over the database execution environment. Administrators can assign processor and pre-fetch priorities to each workload, and execute workloads according to different business priorities. This level of control can help organizations meet service-level agreements (SLAs) by prioritizing execution of business-critical workloads.

Analytics in action: German travel and leisure organization

Business challenge

A German travel and leisure organization needed to build a faster, more stable information management platform for a new business model, and wanted to deliver web-based access to cross-system, standardized information. It also wanted to integrate BI into operational applications and better support the company's decision-making processes.

The organization selected an end-to-end IBM solution that includes a scalable hardware package, a dynamic warehousing concept, industry-leading BI and software that provides web-based access to all data.

IBM solution components

- IBM Cognos® BI
- IBM InfoSphere Balanced Warehouse
 - IBM InfoSphere Warehouse
 - IBM System Server
 - IBM System Storage®

Solution benefits

The organization is more agile, thanks to comprehensive and consolidated near-real-time information. It has shortened its product cycles, increased flexibility and lowered costs with a streamlined, integrated, company-wide system. Now the organization can give all employees the benefits of the insights gained through enhanced BI.

In addition, time-based controls help save money by automatically throttling low-priority work until it can execute without affecting high-priority work. Capabilities such as buffer pool I/O prioritization and integration with Linux workload management give administrators even greater control over the database execution environment.

- **Autonomics and self-tuning memory manager:** InfoSphere Warehouse simplifies memory configuration by automatically setting values for several configuration parameters. Database workloads seldom remain static, and changes can occur in seconds. A system tuned by even the most skilled administrator at one point in time may not be optimal at another time. Memory settings are especially vulnerable to workload changes—the wrong settings can affect response times and threaten SLAs.

The InfoSphere Warehouse self-tuning memory manager (STMM) includes an automatic tuner that dynamically distributes available memory resources among several memory consumers for the database. STMM responds to changes in workload characteristics, adjusting the values of memory configuration parameters and the sizes of buffer pools to optimize performance.

- **Performance management:** InfoSphere Warehouse provides a comprehensive and proactive performance management solution for database applications, while helping DBAs optimize the performance and availability of the database and resolve performance bottlenecks before they affect your business. The performance management capabilities:
 - Enable early and rapid problem detection to prevent impact on production systems.
 - Enable improved performance management to meet SLAs via optimization and tuning recommendations.
 - Support trend analysis across the performance warehouse for growth planning.

Multidimensional analysis (OLAP)

With cubing services, the integrated systems for operational analytics from IBM provide a new way of delivering information from the warehouse. Cubing services extend business insight through speed-of-thought access without moving or duplicating data. Unlike other OLAP tools, cubing services seamlessly access data within the warehouse, dynamically building cached cubes that are exposed through industry-standard OLAP interfaces such as Multidimensional Expressions (MDX), XML for Analysis (XMLA) and OLE DB for OLAP (ODBO). Cubing services are accessible from BI clients such as Cognos, Microsoft Excel, IBM DB2 AlphaBlox® and IBM DataQuant.

Analytics in action: Brazilian financial institution

Business challenge

A Brazilian financial institution needed to better predict client buying behaviors, identify cross-sell and up-sell opportunities and discern market trends.

IBM solution components

The organization deployed an IBM solution to consolidate data from more than 70 different sources into a data warehouse supporting customer relationship management (CRM) and risk-management applications. The components of the financial institution's solution are:

- IBM InfoSphere Warehouse
- IBM Power Systems™
- IBM System Storage
- IBM InfoSphere Information Server
- IBM ILOG® Optimization Decision Manager (ODM) for decision support for planning and scheduling
- IBM Industry Model

Solution benefits

The institution is now able to leverage database mining to understand clients and reduce operational risk. It has reduced its CRM model deployment cycle from eight months to a few weeks, and can now execute analytical models in near-real time to better predict customer buying behaviors and market trends.

Cubing services also provide optimization techniques to dramatically improve the performance of OLAP queries. In doing so, they simplify the delivery of business analytics and optimization results, and put more power into more decision makers' hands to analyze data and generate business insights.

The value of cubing services derives from their scalability and performance:

- **Scalability:** Cubing services can deliver performance with data derived from terabytes of source data, leveraging the scalability of DB2.
- **Performance:** Cubing services can provide speed-of-thought access through their advanced, dynamically constructed cubes. Cubing services build upon the Hybrid OLAP capability implemented through the OLAP accelerator tooling, taking metadata from the OLAP structures defined in InfoSphere Warehouse Design Studio and generating a dimensional cache within the cubing server for each imported cube definition.

Data mining and text analytics

Powerful yet simple to use, the text analytics and data mining capabilities available with the integrated systems for operational analytics from IBM enable integrated analysis of both structured and unstructured data. The system supports standard data mining models (clustering, associations, classification and prediction) that can be developed via drag-and-drop methods in InfoSphere Warehouse Design Studio. The data mining models can be executed in the production environment to provide real-time scoring of data records.

The integrated systems for operational analytics offer rich presentation components to enable visual analysis of data mining results. These models also may be imported in industry-standard Predictive Model Markup Language (PMML) format from third-party modeling tools. Plus, the models provide the ability to extend data mining workbench tools, such as IBM SPSS® Modeler, enabling data miners to post their mining models into the database for real-time execution.

The predictive analytics and pattern analysis capabilities can help users detect fraud, reduce customer churn, segment customers and simplify market-basket analysis. The in-database data mining capabilities integrate with existing systems to provide scalable, high-performing predictive and pattern analysis without moving data into proprietary data mining platforms.

Text analytics is a key IBM advantage. Many other solutions cannot access potentially valuable unstructured data captured across the organization such as call-center notes, customer feedback, free-form text fields, image documents and web pages. InfoSphere Warehouse supports the analysis of previously untapped unstructured data, helping to provide additional insights into customer and product issues.

Analytics in action: Large healthcare insurance provider

Business challenge

A large healthcare insurance provider needed to give employers, consumers and healthcare providers better decision support, more actionable information and insights from predictive analysis.

IBM solution components

- IBM Cognos
- IBM InfoSphere Warehouse
- IBM Power Systems
- IBM System Storage

Solution benefits

The insurance provider now has more and better information about treatment and medical-technology effectiveness, and greater insight into emerging trends in healthcare practice and delivery.

Business intelligence

BI modules for these integrated systems are based on IBM Cognos, which delivers a complete range of BI capabilities—reporting, analysis and dashboarding—on a single service-oriented architecture (SOA). The integrated systems for operational analytics and InfoSphere Warehouse from IBM provide universal access, building on open standards that support a multitude of third-party tools for discovery and analysis.

Integrated systems for operational analytics from IBM put industry-leading BI software into the hands of the business leaders who demand an easy-to-use solution. Users can access a consistent view of information from across the business and multiple data sources. They can uncover and share new insights, and make better decisions to drive the business forward.

Deep integration and optimization of BI capabilities enable organizations to hit the ground running and achieve a rapid return on investment. The flexibility of these integrated systems allows organizations to expand and adapt the BI environment as requirements change.

IBM portfolio of integrated systems for operational analytics

Integrated systems for operational analytics are available from IBM in multiple configuration models (see Table 1):

- **IBM Smart Analytics System 5710:** A compact, cost-efficient, single-server solution based on IBM System x® technology, this model is designed for midsize businesses and enterprise departments looking to quickly deploy analytics and BI capabilities.
- **IBM Smart Analytics System 5600:** A modular, enterprise-strength solution based on IBM System x technology and designed for business analytics workloads, this model offers an optional solid-state drive to reduce data latency.
- **IBM PureData System for Operational Analytics:** Based on the latest IBM Power Systems (POWER7®) technology, the PureData System for Operational Analytics takes integration, built-in expertise and a simplified user experience to a new level. Designed to handle 1,000 or more concurrent operational queries,¹ it is available in five configuration sizes and can scale up to a petabyte of data capacity.
- **IBM zEnterprise® Analytics System 9710:** An analytics platform based upon the new IBM zEnterprise 114 platform, this configuration model can support operational analytics, deep mining and analytical reporting with minimal data movement. It performs operational analytics with the quality of service of IBM System z® at an entry-level cost.
- **IBM zEnterprise Analytics System 9700:** This model offers an analytics platform that can support operational analytics, deep mining and analytical reporting with minimal data movement. It offers operational analytics with the same quality of service as an online transactional processing (OLTP) environment on System z.

Table 1: Integrated systems for operational analytics from IBM

	IBM Smart Analytics System 5710	IBM Smart Analytics System 5600	IBM PureData System for Operational Analytics	IBM zEnterprise Analytics System 9710	IBM zEnterprise Analytics System 9700
Core software	<ul style="list-style-type: none"> • InfoSphere Warehouse Departmental Edition • Cognos 	<ul style="list-style-type: none"> • InfoSphere Warehouse Enterprise Edition • Cognos 	<ul style="list-style-type: none"> • InfoSphere Warehouse Enterprise Edition • Cognos 	<ul style="list-style-type: none"> • Cognos for z/OS 	<ul style="list-style-type: none"> • InfoSphere Warehouse on System z • Cognos for Linux on System z or Cognos for z/OS • SPSS 14.2 Modeler
Analytical capabilities	<ul style="list-style-type: none"> • Query and reporting • Multidimensional analysis (OLAP) • Data mining • Text analytics 	<ul style="list-style-type: none"> • Query and reporting • Multidimensional analysis (OLAP) • Data mining • Text analytics 	<ul style="list-style-type: none"> • Query and reporting • Multidimensional analysis (OLAP) • Data mining • Text analytics 	<ul style="list-style-type: none"> • Query and reporting • Multidimensional analysis (OLAP) 	<ul style="list-style-type: none"> • Query and reporting • Multidimensional analysis (OLAP) • Data mining
Operating system	Linux	Linux	IBM AIX®	IBM z/OS	IBM z/OS with SUSE or Red Hat Enterprise Linux
Server platform	IBM System x	IBM System x	IBM POWER7	IBM System z	IBM System z
Storage platform	IBM System Storage (large configuration only)	IBM System Storage, solid-state drive (optional)	IBM System Storage, solid-state drive standard (optional expansion)	IBM System Storage, solid-state drive (optional)	IBM System Storage, solid-state drive (optional)
Build, deploy, health check and support service	Optional	Yes	Yes	Yes	Yes

Business benefits of integrated systems

Integrated systems for operational analytics from IBM offer a broad range of benefits for organizations seeking to analyze volumes of data rapidly:

- **Fast answers to business questions:** Quickly delivers information in context to decision makers, and provides an interactive, self-service environment for exploration and analysis.
- **Optimized business performance:** Enables decision makers to easily measure and monitor financial and operational business performance, analyze results and predict outcomes.
- **Better, faster decisions:** Provides decision makers and individual users across the organization with real-time information and new insights—where, when and how they need it.
- **New opportunities:** Helps business users discover opportunities to maximize customer and product profitability, minimize customer churn, detect fraud, improve patient care, increase campaign effectiveness and more.
- **Rapid time-to-value:** Provides fully integrated solutions that accelerate the time-to-value by enabling organizations to start loading data within hours.
- **Reduced TCO:** Enables organizations to reduce TCO by offering a single line of support and by minimizing integration and management costs.
- **Simplified management:** The IBM PureData System for Operational Analytics offers built-in management capabilities and integrated system upgrades to help simplify system administration and free up IT resources to work on strategic projects.

Find actionable insights in mountains of data

Smart organizations today are focused on collecting, synthesizing and making sense of the vast amounts of information available to them. To succeed at those tasks, they need a high-performance, adaptable, analytics-optimized system.

Integrated systems for operational analytics from IBM are complete, ready-to-deploy analytical solutions designed to reduce the time it takes to extract value from information and lower the cost of deriving value from enterprise information. The systems' modular architectures make it possible for organizations of all sizes to cost-effectively add capacity and capabilities as their analytics needs change over time.

Integrated systems for operational analytics give your organization the information it needs to work smarter by putting the right answers in the hands of your decision makers today while putting your business in the best position to quickly adapt and grow to answer the questions of tomorrow.

For more information

To learn more about the IBM PureData System for Operational Analytics, visit ibm.com/software/data/puredata/analytics/operational

To learn more about the IBM Smart Analytics System, please contact your IBM representative or IBM Business Partner, or visit ibm.com/software/data/infosphere/smart-analytics-system



© Copyright IBM Corporation 2012

IBM Corporation
Software Group
Route 100
Somers, NY 10589

Produced in the United States of America
October 2012

IBM, the IBM logo, ibm.com, AIX, AlphaBlox, DB2, Cognos, ILOG, InfoSphere, Power Systems, POWER7, PureData, PureSystems, SPSS, System Storage, System x, System z, zEnterprise and z/OS are trademarks of International Business Machines Corp., 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 ibm.com/legal/copytrade.shtml

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

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

This document is current as of the initial date of publication and may be changed by IBM at any time. Not all offerings are available in every country in which IBM operates.

The client examples cited are presented for illustrative purposes only. Actual performance results may vary depending on specific configurations and operating conditions. THE INFORMATION IN THIS DOCUMENT IS PROVIDED "AS IS" WITHOUT ANY WARRANTY, EXPRESS OR IMPLIED, INCLUDING WITHOUT ANY WARRANTIES OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE AND ANY WARRANTY OR CONDITION OF NON-INFRINGEMENT. IBM products are warranted according to the terms and conditions of the agreements under which they are provided. Actual available storage capacity may be reported for both uncompressed and compressed data and will vary and may be less than stated.

¹Based on IBM internal tests of prior-generation system, and on system design for normal operation under expected typical workload. Individual results may vary.



Please Recycle
