An Analytics and Business Intelligence Solution Using IBM Smart Analytics System on eX5 Servers

- A comprehensive and advanced suite of business analytics tools
- Powerful, reliable, and scalable IBM eX5
- Based on the Intel Xeon E7 processor family

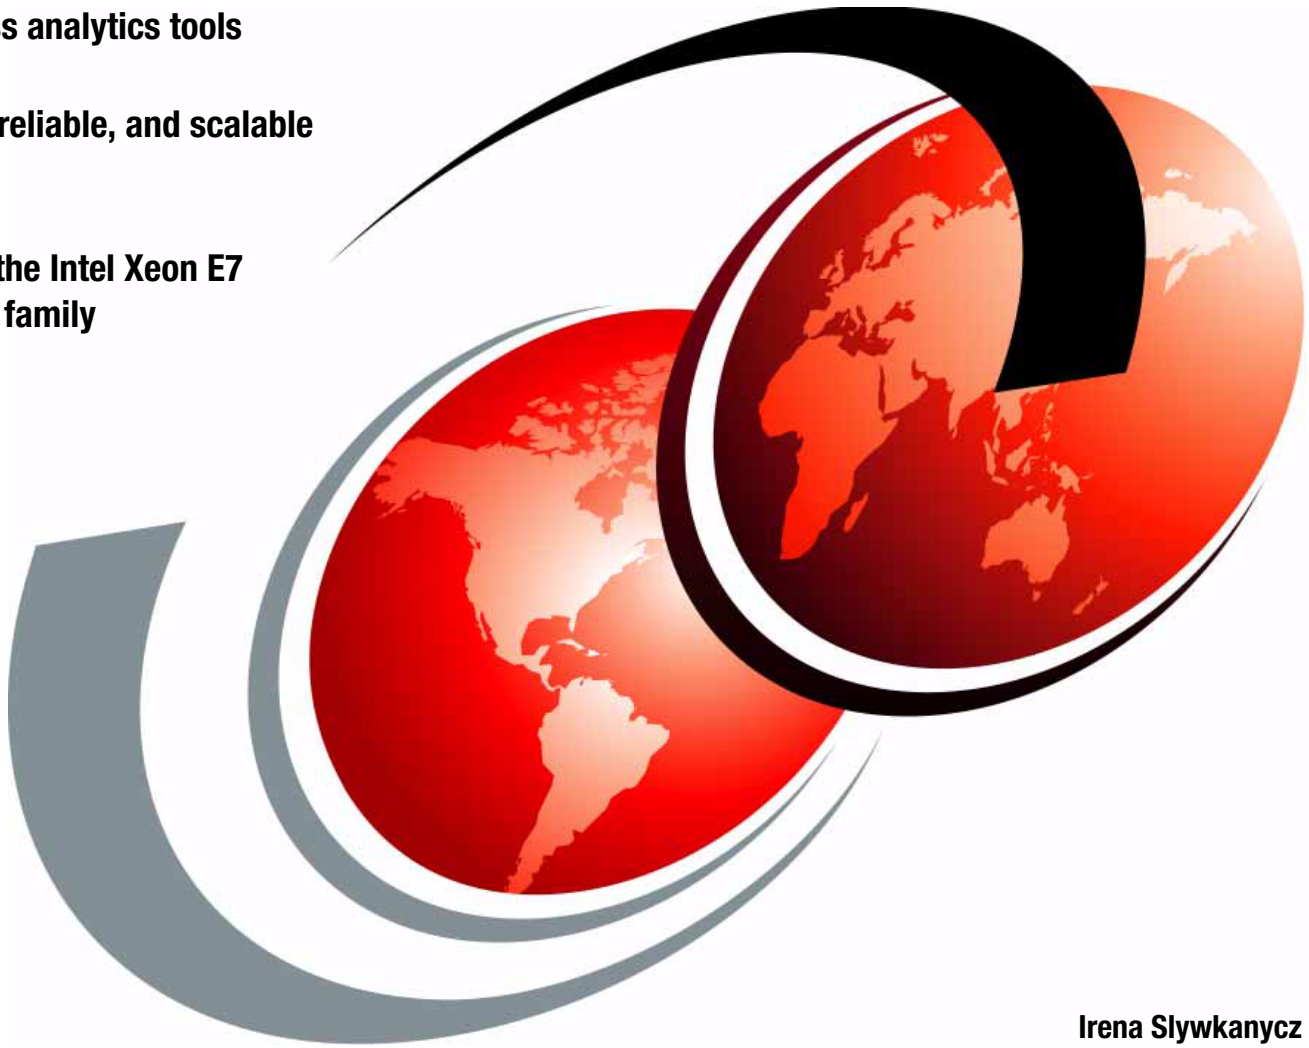

Irena Slywkanycz

Redpaper

ibm.com/redbooks

International Technical Support Organization

An Analytics and Business Intelligence Solution Using IBM Smart Analytics System on eX5 Servers

April 2011

REDP-4743-00

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

First Edition (April 2011)

This edition applies to IBM System x eX5 servers and IBM Smart Analytics System.

Contents

Notices

This information was developed for products and services offered in the U.S.A.

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 give you any license to these patents. You can send license inquiries, in writing, to:
IBM Director of Licensing, IBM Corporation, North Castle Drive, Armonk, NY 10504-1785 U.S.A.

The following paragraph does not apply to the United Kingdom or any other country where such provisions are inconsistent with local law: 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 states 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 supply in any way it believes appropriate without incurring any obligation to you.

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.

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 the names and addresses used by an actual business enterprise 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.

Trademarks

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. These and other IBM trademarked terms are marked on their first occurrence in this information with the appropriate symbol (® or ™), indicating US 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 http://www.ibm.com/legal/copytrade.shtml

The following terms are trademarks of the International Business Machines Corporation in the United States, other countries, or both:

Cognos®
Dynamic Infrastructure®
Express Advantage®
Express®
IBM Systems Director Active Energy Manager™
IBM®
InfoSphere™
OnForever™
Redbooks®
Redpaper™
Redbooks (logo)  ®
Smarter Planet™
System Storage®
System x®
X-Architecture®

The following terms are trademarks of other companies:

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

Intel Xeon, Intel, 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.

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

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

Preface

In today's business growth, data is being generated at an incredible rate. At the same time, the move toward real-time computing requires faster and more reliable data access, especially when databases are used to drive client-facing applications. Some data is structured, residing in relational databases that are distributed throughout the business, but far more data is unstructured, residing in emails, call center logs, documents, web pages, and multimedia files, to name only a few repositories.

Most businesses understand the tremendous value locked within these varied data repositories. Businesses use data warehouses, data marts and analytics tools to aggregate and analyze important segments of their data. When implemented successfully, these solutions deliver substantial and measurable business value, helping business leaders better understand their clients, competitors and operations, and also helping them identify opportunities, risks and inefficiencies.

In this IBM® Redpaper™ publication, we discuss the specific benefits, advantages, and features of the IBM Smart Analytics System working with IBM System x® eX5 servers. Using sample configurations, we illustrate how these systems together can provide you with advanced business analytics that make them a valuable, powerful, and cost-effective solution for IT organizations.

The team who wrote this paper

This paper was produced by the International Technical Support Organization (ITSO).

Irena Slywkanycz is a technical writer at the International Technical Support Organization, Poughkeepsie Center. A documentation specialist with 20 years of experience, she has worked in various organizations designing and creating user content such as online help, intranets, web content, newsletters, field news, release notes, management memos, press releases, application guides, and manuals.

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

Lydia Parziale
Mike Ebbers
International Technical Support Organization, Poughkeepsie Center

Aman Lalla
IBM Toronto Lab

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 client 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 papers to be as helpful as possible. Send us your comments about this paper 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

Dealing with business growth

We live in a time of rapid data growth. Business growth can accelerate this increase, as can mergers, acquisitions, new application deployments, and application upgrades that increase data processing workloads. Database solutions must be able to absorb increasing data volumes and expanding workloads while sustaining performance and avoiding disruptive capacity upgrades.

1.1 The problem of rapid data growth

There are some things that you do not mind waiting for. Business intelligence is not one of those things. Every day, your organization collects more data about its clients, partners, and markets. BI and analytics give you the ability to make sense of that ocean of data by crystallizing the patterns, trends, and knowledge that help top executives and front-line employees make smart, fast decisions; find opportunities; and keep your business ahead of the pack.

But if organizations cannot afford to wait for BI, neither can they afford to write a blank check to create it. IT budgets are already stretched to the breaking point. Racks of data center servers absorb scarce floor space, power, and cooling capacity. Meanwhile, IT organizations are under tremendous pressure to extract maximum value out of every dollar spent. But with IT data volumes expected to grow by leaps and bounds over the next five years, senior staffers are watching maintenance as troubleshooting tasks crowd strategic projects off their calendars.

When it comes to business and technology, though, one size does not fit all. Mid-sized businesses simply do not operate in the same way that giant conglomerates do. You do not have massive IT teams and multimillion dollar budgets, but you still need to leverage your data and information to make smart business decisions. You need flexible, out-of-the-box technology solutions that are easy to implement and provide immediate value. However, you should still be able to rely upon best-of-breed solutions with unrivaled performance and reliability.

1.2 IBM Smart Analytics on System ex5

IBM has amassed extensive expertise and established best-practices libraries based on real world experience. It has a comprehensive and advanced suite of business analytics software. Intel® brings another critical piece to the puzzle: the computing power needed to handle the enormous workloads of a business analytics solution on a flexible, affordable, yet highly resilient hardware infrastructure. The latest IBM System eX5 servers with the Intel Xeon® processor E7 family deliver exceptional performance, scalability, energy efficiency and resilience for complex, data-intensive analytics workloads. IBM has used these processors as the foundation for a new generation of enterprise-class servers that are well-suited to demanding and mission-critical analytics workloads.

IBM Smart Analytics System 1050 and IBM Smart Analytics System 2050 are cost-effective solutions for the midmarket and for departments looking to quickly deploy analytics and business intelligence capabilities on a powerful warehousing foundation.

Rich in functionality and delivered analytics-ready, the Smart Analytics System is a ready-to-go solution consisting of a preconfigured integration of analytics software, powerful data warehouse, hardware and services. We have done the integration, testing, and configuration work for you. Before a Smart Analytics System is shipped, IBM laboratories configure and stress test technology. IBM System x with DS3500 Storage is non-proprietary hardware that can easily be scaled within an organization in accordance with changing business needs, and all support is from a single vendor to allow for ease of optimization, troubleshooting, and future upgrades.

1.2.1 Smart Analytics System benefits

IBM Smart Analytics System 1050 and IBM Smart Analytics System 2050 are compact, appliance-like solutions. They are single systems delivering the right balance of power and simplicity. Preintegrated and optimized with industry-leading IBM software and hardware based on years of experience, you can gain results faster, often in days instead of months. Available on both Windows® and Linux® platforms to maximize your choice, IBM Smart Analytics System 1050 and IBM Smart Analytics System 2050 offer a range of preconfigured capacities from 330 GB to 13.2 TB of user space or 600 GB to 24 TB of total available storage.

These IBM workload optimized systems offer the following benefits:

- Reduced risk
  - Preconfigured and certified for optimal performance
  - Based on years of best practices experience
  - Support from a single proven vendor for ease of optimization, troubleshooting, and future upgrades
- Faster time to value
  - Delivered analytics-ready
  - Can accelerate your business intelligence deployment with optionally installed Cognos® 8 Business Intelligence software
- Simplicity
  - Predefined configurations for reduced complexity
- Flexibility for growth
  - Reliable, nonproprietary

Smart Analytics System 1050 and Smart Analytics System 2050 help achieve greater access to information and analytics, faster time to value, and reduced cost of ownership. They are delivered on energy-smart servers using Intel Xeon and, for the first time, IBM eX5 processors. In addition, the core data warehousing software is based on IBM InfoSphere™ Warehouse V9.7. Innovations in InfoSphere Warehouse V9.7 can help reduce the time to value of warehouse implementations, improve the security, availability, and scalability of warehouses, and increase the ease of implementation and delivery of business analytics.

1.2.2 System eX5 benefits

IBM efficient, scalable, and smart System eX5 servers give you the power to maximize benefit from your data and allow you the flexibility to add capability and features as needed without replacing your system.

System x hardware is easily tailored to your workload demands

The IBM System eX5 offers simple-to-deploy, preconfigured systems that are optimized for specific workloads. The eX5 database optimized systems offer the right balance of processing power, memory, and eXFlash high-performance storage technology for your database workloads. The eX5 virtualization optimized systems with embedded VMware vSphere Hypervisor come standard with MAX5 for unmatched memory capacity, thereby allowing you to increase the size and quantity of your virtual machines.

Do more with your investments

Four- to ten-core processor options and expanded memory capacity let you manage memory-intensive applications with higher utilization and greater throughput and bandwidth. You can consolidate older servers on one eX5 server for a lower total cost, and you can optimize your investments with more and larger virtual machines per server and lower licensing costs, while achieving greater business productivity.

Consolidate with confidence

A long-time leader in reliability, many of the features IBM pioneered with earlier generations of X-Architecture® have now become industry standards. The x3850 X5 continues to build on this leadership, enabling you to consolidate workloads without sacrificing availability or your peace of mind. OnForever™ features such as redundant power and cooling help minimize the potential for catastrophic business disruption. Other features include enhanced IBM Memory ProteXion for greater data integrity, memory mirroring and Predictive Failure Analysis for increased system uptime.

Own and operate with ease

Consolidated single points of management and an energy-smart design help reduce the time and money needed to run your system. Advanced light path diagnostics and around-the-clock remote access management facilitate ease of use and proactive problem-solving. In addition, IBM Systems Director Active Energy Manager™ helps monitor, measure, and manage power consumption to lower the energy draw and costs.

Select configurations of the x3850 X5 are part of the IBM Express® Advantage® Portfolio designed to meet the needs of midsize businesses. Easy to manage, Express models and configurations vary by country.

2

Overview of IBM Smart Analytics System on eX5

Business analytics has become one of the most powerful tools available for making a company more competitive. To keep pace with growing needs, businesses need solutions that provide industry-leading analytics capabilities that also minimize cost, complexity and implementation risk. The IBM Smart Analytics System deployed on Intel Xeon processor-based IBM System x servers delivers fully integrated business analytics on a powerful, cost-effective hardware infrastructure. Additionally, as their need grow, companies can add software functionality and hardware capacity in increments and without replacing the entire system.

2.1 A comprehensive integrated and optimized solution

Unlike competing solutions, the IBM Smart Analytics System, together with Intel Xeon processors, provides a comprehensive, integrated system with features that can be easily tailored to meet the needs of mid-size to large organizations. The key system features include the following:

- A scalable, high-performance data warehouse for aggregating data from multiple sources throughout the business
- Advanced analytics and business intelligence tools for analyzing the data and for distributing information and insights to applications and users
- Development tools and a management framework for extending the solution and maintaining high performance and availability as needs evolve and workloads grow and change

The IBM Smart Analytics System, running on System x servers based on a flexible and affordable Intel Xeon processor based-hardware infrastructure, allows companies to quickly and efficiently perform key functions, for example:

- Establish a single, trusted version of the truth by eliminating information silos and establishing a common framework for storing, understanding, and using information across the business
- Provide powerful, self-service analytics tools tailored to the needs and abilities of a wide range of individual users, from expert data analysts to top executives to front-line employees
- Channel real-time, actionable information and insights back into business operations, so that users and applications have access to necessary information, in context and in the most usable format
- Deliver these capabilities in a single service-oriented architecture (SOA) that enables efficient development, reuse, and management of information and analytics to promote continual improvement
- Create a flexible and scalable hardware and software foundation for business analytics that can grow as the complexity of business information grows and as new tools and strategies become available

In the Smart Analytics system using System x servers, IBM provides a a complete set of hardware, software, services, and support in a single package. This package also gives you a single price, a single point of service, and support that will help minimize your project risk.

2.2 Product solution highlights

The x3850 X5 provides the ultimate in configuration choice, availability and customized performance including the following:

- Highly scalable memory capacity and a range of processor core and hard disk drive options that provide for ultimate flexibility
- Built on the latest IBM X-Architecture with fifth-generation eX5 chipset design enhancements
- Easy upgrade path and pay-as-you-grow expansion offer greater investment protection
- Greater performance with higher utilization, throughput and bandwidth for virtualization and enterprise applications

- Advanced service and remote management capabilities for greater ease of use and improved productivity
- Power-optimized, energy-smart design for enhanced performance per watt

In the IBM Smart Analytics System, IBM has assembled the most comprehensive portfolio of information management, hardware, software, and services capabilities on the market today. The IBM Smart Analytics System provides an ideal solution for rapid deployment of these capabilities and accelerates delivery of new analytic innovations, including those from IBM Research.

The base building blocks of the solution are the x3850 X5 server and the MAX5 memory expansion drawer. The x3850 X5 is a 4U system with four processor sockets and up to 64 DIMM sockets. The MAX5 memory expansion drawer is a 1U device that adds 32 DIMM sockets to the server.

The IBM x3850 X5 and x3950 X5 are the follow-on products to the eX4-based x3850 M2, and like their predecessor, are four-socket systems. The x3950 X5 models are optimized for specific workloads such as virtualization and database workloads.

The MAX5 memory expansion unit is a 1U device, which you connect to the x3850 X5 or x3950 X5, and provides the server with an additional 32 DIMM sockets. It is ideal for applications that can take advantage of as much memory as is available.

2.3 IBM Smart Analytics System 1050 and System 2050 features

IBM Smart Analytics System 1050 and IBM Smart Analytics System 2050 are cost-effective solutions for the midmarket and for departments looking to quickly deploy analytics and business intelligence capabilities on a powerful warehousing foundation. These compact, appliance-like solutions are single systems that deliver the right balance of power and simplicity and arrive preintegrated and optimized with industry-leading IBM software and hardware. This means you can gain results faster, often in days instead of months.

Available on both Windows and Linux platforms, IBM Smart Analytics System 1050 and IBM Smart Analytics System 2050 offer a range of preconfigured capacities that will effectively meet your enterprise needs.

The IBM 1050 and 2050 Systems are workload optimized to offer the following benefits:

- Well-tested and well-configured solution on specific hardware reduces risk.
- Preconfigured and certified for optimal performance.
- Based on years of IBM experience and best practices expertise.
- A single proven vendor provides support, which translates into ease of optimization, troubleshooting, and smoother future upgrades.
- Faster time to value.
- Systems are delivered analytics ready.
- Systems can accelerate your business intelligence deployment with optionally installed Cognos 8 Business Intelligence software.
- Design simplicity.
- Predefined configurations for reduced complexity.

- Flexibility for growth.
- Reliable and nonproprietary.

Smart Analytics Systems 1050 and System 2050 give you greater access to information and analytics, faster time to value, and reduced cost of ownership. They are delivered on energy-smart servers using Intel Xeon and, for the first time, IBM eX5 processors. In addition, the core data warehousing software is based on IBM InfoSphere Warehouse V9.7.

Innovations in InfoSphere Warehouse V9.7 offer many advantages. For example, InfoSphere Warehouse V9.7 can:

- Reduce the time to value of warehouse implementations.
- Improve the security, availability, and scalability of warehouses.
- increase the ease of implementation and delivery of business analytics.

IBM Smart Analytics System 2050 is built upon the IBM System x3850 X5, providing enhanced power and unrivaled reliability and confidence in your IT solution deployments. The x3850 X5 takes performance, efficiency, and reliability to the next level with a balanced departmental server built on the next generation of IBM X-Architecture. It features an unmatched combination of x86 performance and scalability with a balanced design.

The x3850 X5 scales to four processors and 1 TB of RAM, and Intel Nehalem processors give twice the performance over the previous x86 technology, saving you money because the server memory technology uses less power and it utilizes new, high-efficiency power supplies.

Smart Analytics System 1050 and System 2050 integrate select InfoSphere and Cognos software with System x servers and storage systems, and are available on both Windows and Linux platforms.

2.4 IBM Smart Analytics Systems at a glance

Table 2-1 lists the major components and features of the Smart Analytics System.

Table 2-1 IBM Smart Analytics System at a glance

Category	Model 1050	Model 2050
Data Warehousing and Cubing Services	InfoSphere Warehouse Departmental Base Edition or InfoSphere Warehouse Departmental Edition	InfoSphere Warehouse Departmental Base Edition or InfoSphere Warehouse Departmental Edition
Data Mining and Text Analytics	InfoSphere Warehouse Departmental Edition	InfoSphere Warehouse Departmental Edition
Business intelligence	Cognos 8 BI Reporting and Query Studio	Cognos 8 BI Reporting and Query Studio
Operating system	Novell SUSE Linux 11; Windows Server 2008	Novell SUSE Linux 11; Windows Server 2008
Hardware	IBM System x 3500 M3	IBM System x 3850 X5
Storage	IBM System Storage® DS3500	IBM System Storage DS3500
User space capacity	330 GB; 1650 GB; 3.3 TB	3.3 TB; 6.6 TB; 13.2 TB
Total usable storage capacity	Up to 6 TB	Up to 24 TB
Architecture	Single server (tower or rack)	Single server (rack)

Category	Model 1050	Model 2050
Pricing model	Authorized user (20-user minimum for Departmental Base or 30-user minimum for Departmental) or by server	By server

2.5 IBM System 1050 and System 2050 configurations

IBM Smart Analytics System 1050 (Table 2-2) and IBM Smart Analytics System 2050 (Table 2-3) are delivered with all hardware and software preinstalled, integrated, and configured.

Table 2-2 Preset configurations for the IBM Smart Analytics System 1050

User space size	IBM System x Server	Form factor	External storage
330 GB	x3500 M3, 2 x Quad-Core Intel Xeon Processor E5640 series, 16 GB RAM, 4 x 300 BG 10k SAS 2.5-inch Slim-HS HDD	Tower	NA
1650 GB	x3500 M3, 2 x Quad-Core Intel Xeon Processor E5640 Series, 16 GB RAM, 2 x 300 GB 10K SAS 2.5-inch Slim-HS HDD	Rack	1 x DS3500 12 x 300 GB 10K SAS HDD
3300 GB	x3500 M3, 2 x Quad-Core Intel Xeon Processor E5640 Series, 24 GB RAM, 2 x 300 GB 10K SAS 2.5-inch Slim-HS HDD	Rack	1 x DS3500 24 x 300 GB 10k SAS HDD

Determine which system best suits your needs and base your selection on the expected database size, or the user space size.

Table 2-3 Preset configurations for the IBM Smart Analytics System 2050

User space size	IBM System x Rackable System	External storage
3300 GB (Small)	x3850 X5, 2 x Quad-Core Intel Xeon Processor E7520 Series, 16 GB RAM, 4 x 300 GB 10K SAS HDD	1 x DS3500 24 x 300 GB 10K SAS HDD
6600 GB (Medium)	x3850 X5, 4 x Quad-Core Intel Xeon Processor HDD E7520 Series, 32 GB RAM, 4 x 300 GB 10K SAS HDD	2 x DS3500, 48 x 300 GB 10K SAS
13200 GB (Large)	x3850 X5, 4 x Quad-Core Intel Xeon Processor HDD E7520 Series, 64 GB RAM, 4 x 300 GB 10K SAS HDD	4 x DS3500, 96 x 4 x DS3500, 96 x 300 GB 10K SAS

3

IBM System x eX5 servers

With rising costs, increasing workloads and demands for 24×7 availability, there is no room in your organization for underutilized or unreliable servers. Built on the latest next generation of enterprise IBM X-Architecture eX5 technology and Intel Xeon processors, the IBM Smart Analytics on System x offers superior performance, unbeatable reliability and great flexibility within an energy- and wallet-friendly design.

3.1 A new generation of powerful, enterprise-class servers

IBM eX5 servers based on the Intel Xeon processor E7 product family represent a major leap forward in supporting mission-critical workloads on industry-standard servers. With on-demand scalability, flexible partitioning, exceptional memory capacity, and advanced reliability features, the eX5 servers provide powerful data center platforms that can be deployed affordably and then scaled easily to meet changing requirements.

The Intel Xeon processor E7 product family plays an essential role in delivering the performance, scalability, and reliability to support demanding applications on large, enterprise-class systems. Servers based on these processors can help IT organizations scale performance as workloads grow and maintain high availability, all while preserving the compelling cost advantages of Intel Xeon processor-based solutions.

3.2 Why choose the eX5 platform

The IBM Smart Analytics server x3850 X5 allows freedom of choice with extremely flexible configurations plus memory expansion capabilities. A modular building block design lets you customize your system for current needs while providing the ability to react to changing workloads. You can expand your 4-socket, 64-DIMM x3850 X5 to 4 sockets and 96 DIMMs, or up to 8 sockets and 128 DIMMs. You can also reallocate resources as your environment changes. The x3850 X5 meets your needs today, while providing an easy, cost-effective upgrade path to change your environment when you are ready.

Note that x3850 X5 and x3950 X5 represent the follow-on product to the eX4-based x3850 M2, and like that predecessor, they are four-socket systems. The x3950 X5 models are optimized for specific workloads such as virtualization and database workloads. The MAX5 memory expansion unit is a 1U device, which you connect to the x3850 X5 or x3950 X5. It provides the server with an additional 32 DIMM sockets. It is ideal for applications that can take advantage of as much memory as is available.

The IBM System x3850 X5 and x3950 X5 servers offer the following benefits:

- Increase performance on a smaller IT budget.
- Increase database and virtualization performance without adding more CPUs, especially when software is licensed on a per-socket basis.
- Add memory capacity on top of existing processing power, so that overall performance goes up while software licensing costs remain static.
- Achieve the desired memory capacity with smaller capacity DIMMs.
- Pay for the system you need today, and grow both memory capacity and processing power as needed.

Compelling reasons to choose the eX5 platform include:

- Twice the memory capacity, cores, and hard disk drives of previous generations for ultimate flexibility and choice.
- Built on the latest IBM X-Architecture with fifth-generation eX5 chipset design enhancements.
- Greater performance with higher utilization, throughput and bandwidth for enterprise applications.

- Advanced service and remote management capabilities for greater ease of use and improved productivity.
- Power-optimized, energy-smart design for enhanced performance per watt.

3.3 The IBM eX5 family

The IBM eX5 product portfolio represents the fifth generation of servers built upon Enterprise X-Architecture. Enterprise X-Architecture is the culmination of bringing IBM technology, often derived from our experience in high-end enterprise servers, to the x86 server market. Now with eX5, IBM scalable systems technology for Intel processor-based servers has been delivered to mid-sized and high-end x86 server systems. These servers can be expanded on demand, and configured by using a building block approach that optimizes system design for your workload requirements.

As a part of the IBM Smarter Planet™ initiative, our Dynamic Infrastructure® charter guides us to provide servers that improve service, reduce cost, and manage risk. These servers scale to more CPU cores, memory, and I/O than previous systems, enabling them to handle greater workloads than the systems they supersede. Power efficiency and machine density are optimized, making them affordable to own and operate.

The ability to modify the memory capacity independently of the processors, and the high speed local storage options, such as eXFlash, means these system can be highly utilized, yielding the best return from your application investment. eXFlash is a dedicated SSD-based storage array for eX5 systems deployed for high performance database applications and workloads. eXFlash solutions incorporate a unique IBM controller that can support up to 1.6 TB of capacity, and leverage enhanced company I/O technologies to eliminate system bottlenecks. These systems allow your enterprise to grow in processing, I/O, and memory dimensions, so you can provision what you need now, and expand the system to meet future requirements. System redundancy and availability technologies are more advanced than previously available in the x86 systems

IBM System x3850 X5

The x3850 X5 server is the fifth generation of the Enterprise X-Architecture, delivering innovation with enhanced reliability and availability features to enable optimal performance for databases, enterprise applications, and virtualized environments.

The IBM System x3850 X5 server is a unique and compelling high-end modular server targeted at enterprise clients looking for increased consolidation opportunities with expanded memory capacity. It is a 4U four-socket Intel 7500-based (Nehalem-EX) platform with 64 DIMM sockets and can be scaled up to eight processor sockets and 128 DIMM sockets by connecting a second server to form a single system image and maximize performance, reliability, and scalability. This server is a modular building-block server that can be configured with one or two chassis to build a scale-up server with up to 160 logical threads and 3 TB memory. The server is shown in Figure 3-1.

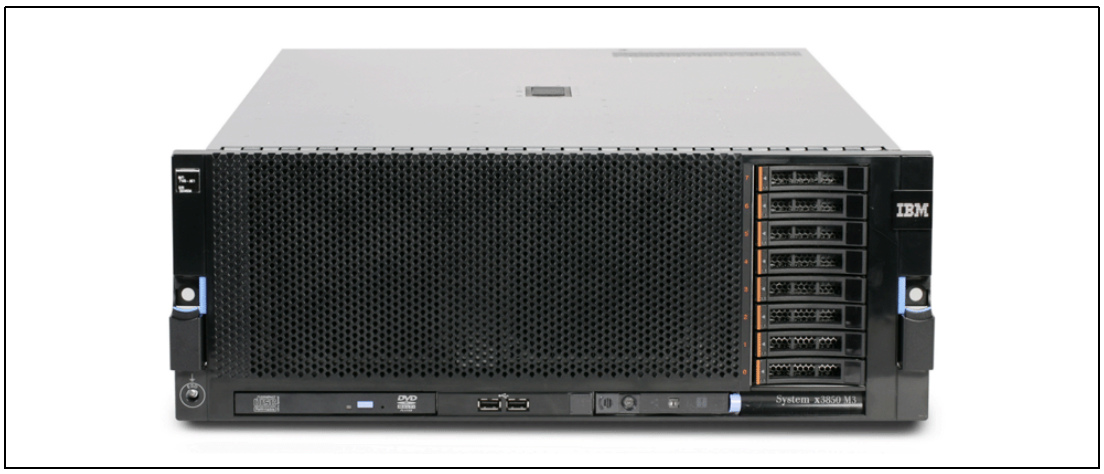

Figure 3-1 Front view of the x3850 X5

The x3850 X5 provides the ultimate in configuration choice, availability and customized performance including the following:

- Highly scalable memory capacity and a range of processor core and hard disk drive options that provide for ultimate flexibility.
- Built on the latest IBM X-Architecture with fifth-generation eX5 chipset design enhancements.
- Easy upgrade path and pay-as-you-grow expansion offer greater investment protection.
- Greater performance with higher utilization, throughput and bandwidth for virtualization and enterprise applications.
- Advanced service and remote management capabilities for greater ease of use and improved productivity.
- Power-optimized, energy-smart design for enhanced performance per watt.

IBM MAX5 memory expansion unit

The IBM MAX5 for System x (MAX5) memory expansion unit has 32 DDR3 DIMM sockets, one or two 675-watt power supplies, and five 40 mm hot-swap speed-controlled fans. It provides added memory and multinode scaling support for the x3850 X5 server. A view of the server is shown in Figure 3-2.

The MAX5 expansion module is based on eX5, the next generation of Enterprise X-Architecture. The MAX5 expansion module is designed for performance, expandability, and scalability. Its fans and power supplies use hot-swap technology for easy replacement without requiring the expansion module to be turned off.

With the addition of the MAX5 memory expansion unit, the x3850 X5 gains an additional 32 DIMM sockets for a total of 96 DIMM sockets. Using 16 GB DIMMs means that a total of 1.5 TB of RAM can be installed.

All DIMM sockets in the MAX5 are accessible, regardless of the number of processors that are installed on the host system.

Figure 3-2 x3850 X5 with the MAX5 memory expansion unit attached

IBM System x3950

For certain enterprise workloads, IBM offers preconfigured models known as the x3950 X5 models. They do not differ from standard models in terms of machine type or the options used to configure them. However, because they are configured with components that make them optimized for specific workloads, they are differentiated by this naming convention. No model of x3850 X5 or x3950 X5 requires a scalability key for eight-socket operation (as was the case with the x3950 M2). Also note that because the x3850 X5 and x3950 X5 use the same machine type, they can be scaled together into an eight-socket solution, assuming that they each use identical CPUs, and each has four CPUs.

IBM offers the x3950 X5, optimized for database workloads and virtualization workloads. Virtualization-optimized models of the x3950 X5 include a MAX5 as standard. Database-optimized models include eXFlash as standard.

3.4 Intel Xeon processor E7 family

The x3850 X5 is powered by the highly scalable Intel Xeon processor E7 family. Built to handle the most demanding applications, the Intel Xeon processor E7 family delivers a quantum leap in enterprise computing performance. The Intel Xeon processor E7 family features Intel Advanced Reliability Technology, which provides automatic detection and correction of errors, dynamic reassignment of workloads across processors, interconnect error detection and recovery, and individual virtual machine recovery in virtualized environments.

Figure 3-3 shows some of the features and benefits of this processor family, followed by a discussion of such benefits as performance, scalability, energy efficiency, and reliability.

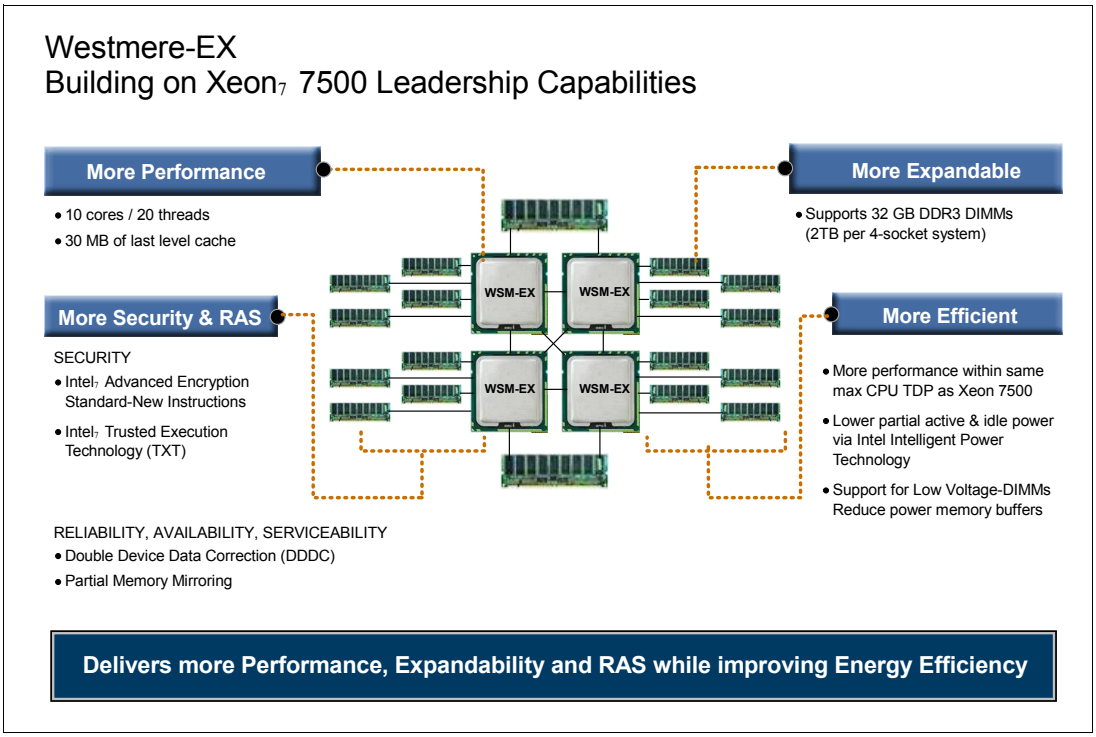

Figure 3-3 The Intel Xeon processor E7 product family

Performance

The Intel Xeon processor E7 family offers outstanding performance for heavy, data-intensive workloads. Built on a new 32nm micro architecture, this product family can provide up to 10 cores, 20 threads, and 30 MB of shared cache per processor. Four advanced, high-bandwidth interconnect links allow multiple processors to be directly connected to each other to increase performance and reduce latency.

These new capabilities deliver tangible performance results. Using the Intel Xeon E7 product family, organizations can accelerate database transactions by up to 40% compared with previous-generation four-socket processors.[1] As a result, organizations can generate results faster or process larger data volumes without having to increase the server footprint.

Scalability

The Intel Xeon processor E7 family provides the scalability required for adding or expanding workloads to accommodate new business opportunities. Intel QuickPath Interconnects allow scaling of processors from two to eight sockets. The Quad-Channel Integrated Memory Controller supports up to 16 memory slots per processor socket, enabling administrators to deploy up to 2 TB of memory in a 4-socket system to handle peak demands and leave headroom for database growth. The IBM MAX5 memory expansion drawer adds 32 DIMM slots for up to a total of 3 TB of memory in a 4-socket system. This large-scale memory capacity enables organizations to place an entire database in memory and run queries in real time, speeding the path from data to decision.

[1] This claim is based on performance on three industry-standard, common enterprise benchmarks (TPC Benchmark* C, SPECint*_rate_base2006, and SPECfp*_rate_base2006) comparing the best published/submitted results on a 4-socket server equipped with the Intel Xeon E7-4870 processor compared with a 4-socket server using the Intel Xeon processor X7460 and 4-socket server using the Intel Xeon processor X7560, as of March 26, 2010. Performance tests and ratings are measured using specific computer systems and/or components and reflect the approximate performance of Intel products as measured by those tests. Any difference in system hardware or software design or configuration may affect actual performance. Buyers should consult other sources of information to evaluate the performance of systems or components they are considering purchasing.

We also have node scalability. One x3850 node can scale to two, with the option of EXA scaling using MAX5 for FlexNode dynamic partitioning and failover capability. This gives these systems the ability to dynamically change to dual nodes and back to single node.

Energy efficiency

The x3850 has a power-optimized, energy-smart design for enhanced performance per watt. IBM exclusive eXFlash SSD technology is designed to improve application and database performance and to reduce storage and energy costs.

- New lower power memory DIMMs at 1.35V (versus 1.5V) provide additional energy efficiencies.
- Integrated power gates.
- Memory power efficiency.

Active Energy Manager

IBM Systems Director Active Energy Manager measures, monitors, and manages the energy components of the eX5 systems. Monitoring functions include power trending, thermal trending, PDU support, and support facility providers. Management functions include power capping and power savings mode.

This application helps clients monitor energy consumption to allow better utilization of available energy resources. The application software enables clients to trend actual energy consumption and corresponding thermal loading of IBM Systems running in their environment with their applications. Active Energy Manager can help with the following tasks:

- Allocate less power and cooling infrastructure to IBM servers
- Lower the power usage on select IBM servers
- Plan for the future by viewing trends of power usage over time
- Determine power usage for all components of a rack
- Retrieve temperature and power information through wireless sensors
- Collect alerts and events from facility providers related to power and cooling equipment
- Better understand energy usage across your data center by performing the following tasks:
  - Identify energy usage
  - Measure cooling costs accurately
  - Monitor IT

Reliability

The Intel Xeon processor E7 family, shown in Figure 3-3 on page 16, includes more than 20 features designed to help organizations protect data and meet service level agreements while refocusing resources on innovation.

- Self-healing capabilities enable continued operation even in the event of component failures. Machine Check Architecture-recovery (MCA-r) technology works with the operating system to recover from memory errors without crashes.
- The Intel Xeon processor E7 family adds support for Double Data Device Correction (DDDC), which facilitates recovery from two DRAM device failures.
- Support for Advanced Encryption Standard–New Instructions (AES-NI) enhances security by significantly reducing the performance penalties usually experienced with pervasive encryption.

The Intel Xeon processor E7 family includes more than 20 reliability, availability, and serviceability (RAS) features, including Machine Check Architecture Recovery, which enables synchronized error handling and automated recovery from many types of previously uncorrectable errors. Figure 3-4 shows some of these RAS features.

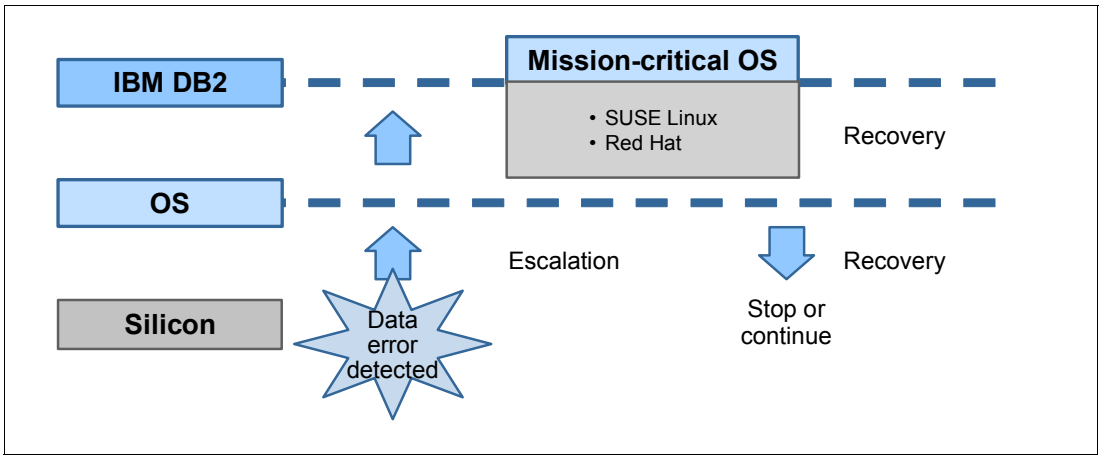

Figure 3-4 RAS features of the Intel Xeon processor

IBM eX5 servers take advantage of these processors to provide exceptional performance and scalability for mission-critical databases. This fifth generation of the IBM X-Architecture platform represents the biggest IBM investment ever in Intel Xeon processor–based servers. Systems are available in two-socket and four-socket configurations. Organizations can scale processor, memory, networking I/O, and storage I/O resources independently in each system. They can also connect a pair of four-socket servers with a high-performance link to create a single server with as many as eight sockets, 80 cores, 160 threads, and 4 TB (6 TB with MAX5) of memory.[2] A fully configured IBM eX5 server provides ample horsepower for hosting heavy, data-intensive workloads, and organizations have considerable flexibility for tailoring their servers to match their workloads.

IBM eX5 servers offer more than just scalable performance. They also deliver major enhancements to improve data integrity and system resilience. IBM supports and extends the silicon-based reliability, availability, and serviceability (RAS) features of the Intel Xeon processor E7 product family with features such as advanced error management, dynamic node partitioning, and automatic failover to improve utilization and uptime for mission-critical workloads.

Servers based on this architecture also include features such as integrated management modules and advanced light path diagnostics to provide sophisticated monitoring and proactive problem resolution. Redundant power supplies help eliminate a key source of potential downtime. Easily accessible components help speed repair time if issues arise.

[2] An IBM eX5 server can be configured with up to 2 TB of memory without an optional external IBM MAX5 memory expansion drawer, and up to 3 TB with such an expansion drawer.

4

Sample configurations

This chapter shows you typical configurations for the IBM Smart Analytics System and System 1050 and 2050. The IBM System x 1050 and 2050 are delivered with all hardware and software preinstalled, integrated, and configured.

4.1 IBM Smart Analytics at a glance

Table 4-1 lists details of the IBM Smart Analytics System.

Table 4-1 IBM Smart Analytics System at a glance

Category	Model 1050	Model 2050
Data Warehousing and Cubing Services	InfoSphere Warehouse Departmental Base Edition or InfoSphere Warehouse Departmental Edition	InfoSphere Warehouse Departmental Base Edition or InfoSphere Warehouse Departmental Edition
Data Mining and Text Analytics	InfoSphere Warehouse Departmental Edition	InfoSphere Warehouse Departmental Edition
Business Intelligence	Cognos 8 BI Reporting and Query Studio	Cognos 8 BI Reporting and Query Studio
Operating System	Novell SUSE Linux 11; Windows Server 2008	Novell SUSE Linux 11; Windows Server 2008
Hardware	IBM System x 3500 M3	IBM System x 3850 X5
Storage	IBM System Storage DS3500	IBM System Storage DS3500
User space capacity	330 GB; 1650 GB; 3.3 TB	3.3 TB; 6.6 TB; 13.2 TB
Total usable storage capacity	Up to 6 TB	Up to 24 TB
Architecture	Single server (tower or rack)	Single server (rack)
Pricing model	Authorized user (20-user minimum for Departmental Base or 30-user minimum for Departmental) or by server	By server

4.2 IBM Smart Analytics System 1050

IBM System 1050 is delivered with three preset capacity configurations, as listed in Table 4-2.

Table 4-2 IBM Smart Analytics System 1050

User space size	IBM System x Server	Form factor	External storage
330 GB	x3500 M3, 2 x Quad-Core Intel Xeon Processor E5640 series, 16 GB RAM, 4 x 300BG 10k SAS 2.5-inch Slim-HS HDD	Tower	NA
1650 GB	x3500 M3, 2 x Quad-Core Intel Xeon Processor E5640 Series, 16 GB RAM, 2 x 300 GB 10K SAS 2.5-inch Slim-HS HDD	Rack	1 x DS3500 12 x 300 GB 10K SAS HDD
3300 GB	x3500 M3, 2 x Quad-Core Intel Xeon Processor E5640 Series, 24 GB RAM, 2 x 300 GB 10K SAS 2.5-inch Slim-HS HDD	Rack	1 x DS3500 24 x 300 GB 10k SAS HDD

4.3 IBM Smart Analytics System 2050

IBM Smart Analytics System 2050 is also delivered with preconfigured settings, as listed in Table 4-3.

Table 4-3 Preset configurations for the IBM Smart Analytics System 2050

User space size	IBM System x Rackable System	External storage
3300 GB (Small)	x3850 X5, 2 x Quad-Core Intel Xeon Processor E7520 Series, 16 GB RAM, 4 x 300 GB 10K SAS HDD	1 x DS3500 24 x 300 GB 10K SAS HDD
6600 GB (Medium)	x3850 X5, 4 x Quad-Core Intel Xeon Processor HDD E7520 Series, 32 GB RAM, 4 x 300 GB 10K SAS HDD	2 x DS3500, 48 x 300 GB 10K SAS
13200 GB (Large)	x3850 X5, 4 x Quad-Core Intel Xeon Processor HDD E7520 Series, 64 GB RAM, 4 x 300 GB 10K SAS HDD	4 x DS3500, 96 x 4 x DS3500, 96 x 300 GB 10K SAS

5

Value of IBM eX5 and the IBM Smart Analytics System

In today's IT-driven business world, data is being generated at an incredible rate, both internally and from an enormous and growing array of external sources. This data provides businesses with a wealth of potential information and insight about their clients, their marketplace and their operations. But turning this data into actionable, trusted information in real time can be a complex and costly undertaking. The IBM Smart Analytics System, running on powerful and cost-effective Intel Xeon processor-based IBM System eX5 servers, provides a solution.

These fully integrated, highly optimized systems deliver advanced business analytics capabilities in comprehensive solutions that can be deployed quickly and with confidence, and can then be scaled and adapted as business needs grow. The latest Intel Xeon processors provide the robust performance needed to support demanding analytics workloads on an affordable hardware infrastructure. IBM integrates these processors into complete hardware and software solutions that help dramatically simplify deployment, integration, maintenance and growth. Businesses of all sizes now have a high-value, low-risk pathway for extracting the immense value locked in their data, thereby enabling them to make better, faster business decisions at every level of the organization.

The Smart Analytics System combines the high performance of the latest Intel Xeon processors with industry-leading IBM software and system architectures to form a single system that is preconfigured, tuned, and tested to deliver remarkable analytics and business intelligence (BI) performance and time-to-value straight out of the box. With workload-optimized systems, you can cut through the complexity of BI and implement an integrated, ready-to-use analytics solution designed to accelerate your business today and grow along with it tomorrow, a valuable balance of performance and power.

Related publications

The publications listed in this section are considered particularly suitable for a more detailed discussion of the topics covered in this paper.

IBM Redbooks

The following IBM Redbooks publications provide additional information about the topic in this document. Note that some publications referenced here may be available in softcopy only.

- *IBM Smart Analytics System*, SG24-4908

  http://www.redbooks.ibm.com/abstracts/SG24908.html

- *IBM eX5 Portfolio Overview: IBM System x3850 X5, x3950 X5, x3690 X5, and BladeCenter HX5*, REDP-4650

  http://www.redbooks.ibm.com/abstracts/redp4650.html

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 website for System eX5

  http://www-03.ibm.com/systems/info/x86servers/ex5/

- IBM Smart Analytics System 2050 website

  http://www-01.ibm.com/software/data/infosphere/smart-analytics-system/2050

Help from IBM

- IBM Support and downloads

  ibm.com/support

- IBM Global Services

  ibm.com/services

An Analytics and Business Intelligence Solution Using IBM Smart Analytics System on eX5 Servers

Redpaper™

A comprehensive and advanced suite of business analytics tools

Powerful, reliable, and scalable IBM eX5

Based on the Intel Xeon E7 processor family

In today's business growth, data is being generated at an incredible rate. At the same time, the move toward real-time computing requires faster and more reliable data access, especially when databases are used to drive client-facing applications. Some data is structured, residing in relational databases that are distributed throughout the business, but far more data is unstructured, residing in emails, call center logs, documents, web pages, and multimedia files, to name only a few repositories.

Most businesses understand the tremendous value locked within these varied data repositories. Businesses use data warehouses, data marts and analytics tools to aggregate and analyze important segments of their data. When implemented successfully, these solutions deliver substantial and measurable business value, helping business leaders better understand their clients, competitors and operations, and also helping them identify opportunities, risks and inefficiencies.

In this IBM Redpaper publication, we discuss the specific benefits, advantages and features of the IBM Smart Analytics System working with IBM System x eX5 servers. Using sample configurations, we illustrate how these systems together can provide you with advanced business analytics that make them a valuable, powerful, and cost-effective solution for IT organizations.

REDP-4743-00

INTERNATIONAL TECHNICAL SUPPORT ORGANIZATION

BUILDING TECHNICAL INFORMATION BASED ON PRACTICAL EXPERIENCE

IBM Redbooks are developed by the IBM International Technical Support Organization. Experts from IBM, Customers and Partners from around the world create timely technical information based on realistic scenarios. Specific recommendations are provided to help you implement IT solutions more effectively in your environment.

For more information:
ibm.com/redbooks