

Information as a Service Unfolds

System z9 Mainframe/DB2 Premier Data & Information
Server for SOA

About this White Paper

Data (*both structured and unstructured*) is at the core of most enterprise businesses today, and must be leveraged much more strongly for competitive advantage. Today the now vast, still fast growing, but often scattered, enterprise data pools renders easy access/availability to critical data too difficult and costly.

Most enterprises need new/better information management (*IM*) solutions to help them to unify, integrate and transform this wealth of raw data into useable information and business insight. They need "Information on Demand", essential for driving new business strategies, to deliver just the information needed, when they need it, where they need it, and more quickly.

"**Information as a Service**" (*IAAS*) is IBM's strategy and vision to deliver Information on Demand (*IOD*). After \$3B, 15 acquisitions, and 3 years of development, IAAS has now unfolded and is being realized. IBM is now delivering a full portfolio of advanced information management service products that deliver IOD enterprise-wide under a Service-Oriented Architecture (*SOA*).

For over four decades, IBM mainframes have been leaders in data and transaction serving. Now IBM is extending and strengthening its System z9-DB2 mainframe offering as the premier enterprise data-serving platform, exploiting its legendary business values and power to serve as the enterprise hub for data and the new information services of IAAS. It has even introduced a new, dedicated information processing CPU on the mainframe, the zIIP, to underpin this central new role.

IAAS information services include data-serving, content management, master data management, information integration, information analysis and discovery, and the new area of entity analytics. Under IAAS, all these are to be delivered as services "on demand" through a unified, cohesive, integrated and open standards-based suite of advanced IBM SOA-compatible middleware software.

In this major White Paper, analysts Software Strategies evaluates how, why and where IBM's IAAS strategy and product-set will help enterprises resolve today's serious information management challenges. It also assesses how the new System z9-DB2 data-serving hub platform fits into this framework, and examines each of the above information services (*and the products that support them*), with special focus on System z9 mainframe environments.

Written by: Ian Bramley

Published: 1st Edition February 2006

Design & Layout: iok design – e-mail: info@iok-design.co.uk

© 2006 Software Strategies. This publication is property of Software Strategies and cannot be adapted, copied, or reproduced, by any means whatsoever, without prior written permission from Software Strategies.

Important Notice: The information available in this publication is given in good faith and is believed to be reliable. Software Strategies expressly excludes any representation or warranty (express or implied) about the suitability of materials in this publication for your purposes and excludes to the fullest extent possible any liability in contract, tort or howsoever, for implementation of, or reliance upon, the information contained in this publication. All rights reserved. This publication, or any part of it, may not be reproduced or adapted by any method whatsoever, without prior written Software Strategies consent. Software Strategies acknowledges all trademarks.



White Paper

Information as a Service Unfolds

System z9 Mainframe/DB2 Premier Data & Information Server for SOA

CONTENTS

| | |
|---|----|
| 1. Executive Summary | 5 |
| 2. Today's Information Management Challenges and SOA | 8 |
| 3. The SOA Application Generation, the Mainframe & Information Services | 11 |
| 4. Introducing Information as a Service – Crucial SOA Enabler | 16 |
| 5. System z9 – Premier Enterprise Data-Serving Platform – Major Advances | 24 |
| 6. System z9 zIIP Specialty Engine – Transforms z/OS DB2 Data-serving Economics | 34 |
| 7. System z9 Content Management and Content Integration Service | 38 |
| 8. Master Data Management as a Service – IBM Steps Up | 44 |
| 9. Enterprise Information Integration Services and System z9 | 49 |
| 10. Information Analysis & Discovery (IAD) Services | 54 |
| 11. Entity Analytic Solutions (EAS) – New IM Services for Identity Resolution | 60 |
| Appendix A: SOA & the zSeries/System z9 Mainframe – A Marriage Made in Heaven | 62 |
| Appendix B: The New System z9 109 Mainframe Generation – Analyst Overview | 64 |
| Appendix C: IBM WebSphere SOA Foundation Software – Overview | 68 |
| Appendix D: Impressive IBM Mainframe Storage Solutions Take Off | 69 |
| Related Software Strategies Research | 72 |
| Software Strategies | 73 |
| About the Author | 73 |

TABLE OF CHARTS

| | |
|--|----|
| Figure 1: Information Management – Turning Data to Information to Insight | 5 |
| Figure 2: What is Service Oriented Architecture? | 9 |
| Figure 3: Information Management Challenges | 10 |
| Figure 4: Service Oriented Architecture Takes Off – Some Indicators | 11 |
| Figure 5: Accelerating zSeries Mainframe Resurgence – Main Drivers | 12 |
| Figure 6: New IBM System z9 109 Mainframe Strategic Positioning – Collaborative Computing Hub | 13 |
| Figure 7: The IBM SOA Reference Architecture & Information Services | 15 |
| Figure 8: Information as a Service Overview | 18 |
| Figure 9: Information as a Service Needs an Ecosystem | 19 |
| Figure 10: IAAS – Information Services and Information Accelerators | 20 |
| Figure 11: Information Services within the SOA Lifecycle | 21 |
| Figure 12: Information Virtualization by IAAS | 22 |
| Figure 13: IBM Information Services – Market Leadership – Powering a SOA | 23 |
| Figure 14: System z9 Data-serving Roadmap 2006-2007 | 25 |
| Figure 15: System z9 Database Servers – Strong Current Release Advances | 27 |
| Figure 16: Extensive DB2 Tools Portfolio for z/OS | 28 |
| Figure 17: Comprehensive IMS Tools Portfolio (z/OS) | 29 |
| Figure 18: DB2 "Viper" Next-Generation Database – Evolving to a Multi-structure Database Model | 31 |
| Figure 19: The zIIP – System z9 Integrated Information Processor | 34 |
| Figure 20: What Content Management Needs to Provide | 38 |
| Figure 21: IBM Content Management Platform – Flexible, Modular, and Open-standards-based | 39 |
| Figure 22: The IBM Master Data Management Approach | 46 |
| Figure 23: IBM WebSphere Information Integration Portfolio | 50 |
| Figure 24: IBM WebSphere Information Integration Connectors | 52 |
| Figure 25: The IBM Information Server | 53 |
| Figure 26: Information Analysis and Discovery – Turning Information into Insight | 56 |
| Figure 27: IBM DB2 EAS Identity Resolution Platform | 61 |
| Figure B1: The New IBM System z9 109 High-end Mainframe | 64 |
| Figure C1: End-to-end Process Capabilities for SOA – New SOA Foundation Software | 68 |
| Figure D1: New Enterprise Storage Systems – Dec 2004 – Complement zSeries, System z9 | 70 |



1. Executive Summary

Our research findings from the body of this White Paper are summarized below:

1. **Enterprise Data Major Challenges:** Enterprises now face large challenges to manage/use their fast-growing, but highly distributed data, both structured and unstructured. This renders easy access to/availability of critical data/information too difficult and costly. (See Section 2.)
2. **Radically Improved Information Management Needed:** Organizations need better information management (IM) approaches that are able to turn their wealth of enterprise data into actionable information/insight. (See Figure 1.)

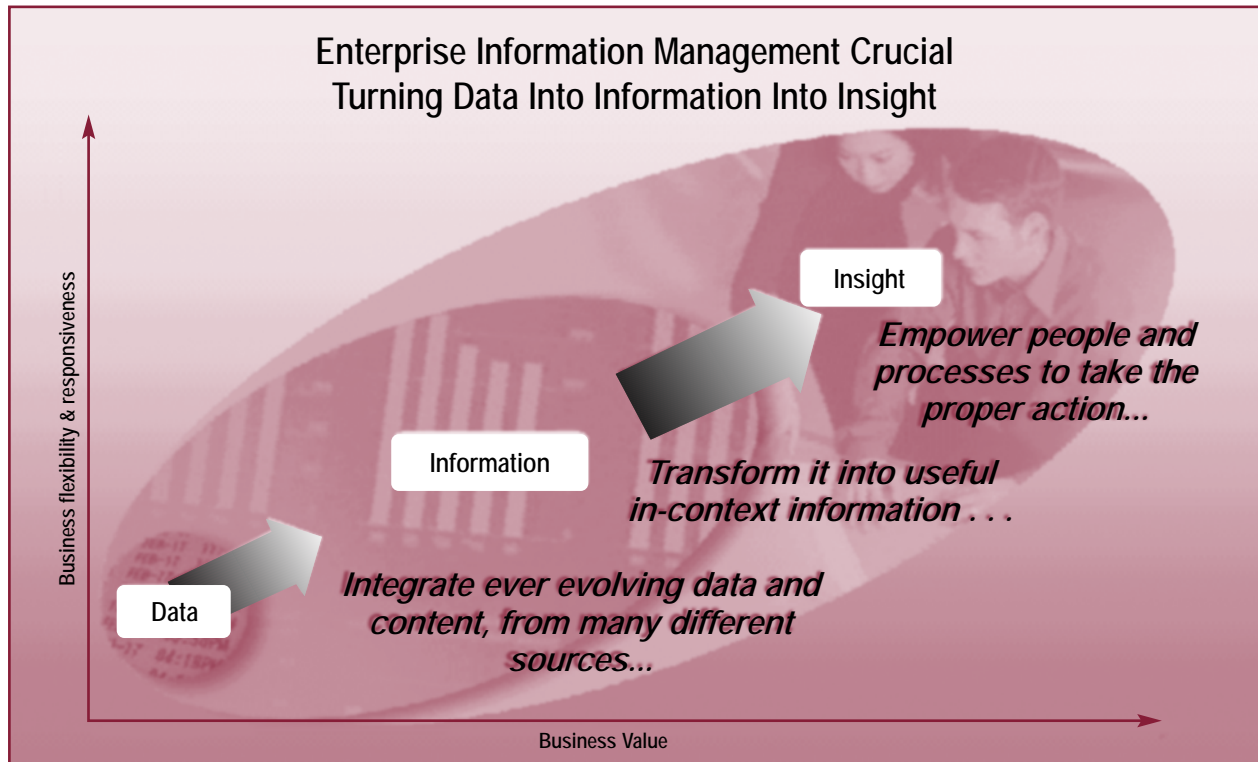


Figure 1: Information Management – Turning Data to Information to Insight

3. **Resurgent Mainframe Already Main Data Host:** The resurgent IBM mainframe (including the latest System z9) hosts/manages a majority of enterprise machine-readable data. Should we now consider a wider future role for the platform supporting these improved IM environments? What are its capabilities, and what benefits flow?
4. **SOA – Next-generation Applications Architecture:** SOA is the most transformative software technology/architecture to emerge since the Internet, and is also the hottest topic/fastest growing market in IT. New IM approaches must fit SOA, support the new SOA applications, and embrace their open standards.
5. **Mainframe Values, Strengths Ideal for SOA:** Distinctive business values, and major advances with the new System z9-109, post powerful claims for the mainframe to be an enterprise node platform for SOA, and the corporate hub for serving data as well as delivering information services. (See Appendices A and B.)
6. **IBM SOA Architecture Highlights Information Services:** IBM's SOA Reference Architecture provides a clear enterprise SOA solutions framework. "Information Services" are a central element. The firm suggests the IM challenges can best be met by treating "Information as a Service" (IAAS) across the enterprise. (See Section 3.)
7. **"Information as a Service" (IAAS) Unfolds:** IBM's IAAS vision/strategy, after \$3B and 3 years investment, is now being realized and delivered (see Section 4). It adds an enterprise-wide IM software services layer and solves most IM challenges. IAAS closely fits into and supports a SOA approach (although can be implemented independently), and is now being delivered in new IBM IM servers/software.
8. **IAAS – A Major Advance:** We assess IAAS to be a major advance, offering new ways to manage, understand and fully exploit today's dauntingly complex and diverse enterprise information resources much more effectively.

System z9-z/OS-DB2 Data Server Leader: This combination already smoothly runs most of the world's largest relational databases.

9. **Will Extend IBM's IM Market Leadership:** IBM is the share leader in the main IM software market segments today. The innovative IAAS advances, here or near, will bring IBM wider technology leadership over its competitors, and fuel further market share growth.

10. **System z9-z/OS-DB2 Data Server Leader:** This combination already smoothly runs most of the world's largest relational databases. With the new System z9 mainframe advances (*capacity doubled*) and recent/forthcoming DB2 advances, IBM now champions further, wider usage of the trio as the premier enterprise data-server at the heart of IAAS. (See Section 5.)

11. **Compelling Roadmap:** DB2 Version 8, z/OS 1.7 and the new System z9 have already delivered major advances. The combined roadmap for the trio is now packed with further significant advances that further strengthen its case.

12. **DB2 UDB for z/OS Version 9? To Raise Bar:** The next-generation DB2 UDB for z/OS Version 9? (*"Viper"*) will raise the RDBMS bar, adding breakthrough native XML support. Its new hybrid relational-XML model offers large advantage over IBM's database competitors, plus many other improvements, thus defining a new level of data server capability, and should debut late in 2006.

13. **New zIIP Specialty Engine:** New System z9 zIIP specialty processor engines (*600 MIPS each*) add a welcome boost, running offloaded qualified DB2 workloads with significant (*software*) cost savings, to encourage wider System z9-z/OS-DB2 data-serving usage. (See Section 6.)

14. **Content Management (CM) Advances:** Managing the exploding volumes of digital content is now a major IM challenge. CM services are also considered central to the IAAS strategy. Enhanced System z9-z/OS-DB2 CM products include:

- **DB2 Content Manager for z/OS Version 8.3:** Scalable advanced ECM solution that fully exploits the unique strengths of the host zSeries/System z9-z/OS-DB2 platform.
- **WebSphere Information Integrator Content Edition:** Provides comprehensive and advanced content integration services to feed the above DB2 CM platform, on the mainframe and elsewhere.
- **WebSphere Information Integrator OmniFind Edition:** Innovative text analytics-based enterprise search capability, IBM's first, able to mine unstructured content – not only keywords – that work across multiple content repositories from z/OS, and onto other platforms.

IBM is the clear CM market leader and these advances will fuel further mainframe-centered CM growth. (See Section 7.)

15. **Information Integration (II) Central for IAAS:** II is a core information service that is vital for new SOA composite applications, and for realizing IBM's full IAAS strategy.

16. **IBM II Offering Unifies EAI, ETL and EII:** IBM's WebSphere Information Integration (WII) suite now unifies EAI, ETL and EII capabilities into a single, cohesive, open framework for the first time. Recent advances across the suite have extended its clear market/technology lead today, adding much-extended capability, parts drawn from the major 2005 Ascential acquisition. (See Section 9.)

II for zSeries/System z9 mainframe environments is already well-supported...

17. **System z9 II Well-supported:** II for zSeries/System z9 mainframe environments is already well-supported, with seven WII products available on the platform today. These support the extended mainframe role as the central node for new enterprise SOA applications, and in the role as the central enterprise data-serving hub platform.

18. **Major Further II Advances Due in 2006:** Further leading-edge developments/innovations are well-advanced and will deliver large advances in this EII framework through spring 2006 deliveries of the ground-breaking, new:

- **The IBM Information Server:** High-function, information integration run-time services platform in a single, unified and comprehensive server. Will be a core SOA foundation, alongside WebSphere Application Server and WebSphere Process Server.
- **WebSphere Metadata Server:** Will provide unified metadata management as a service to all the other products and tools within the IBM WII suite platform.
- **WebSphere Information Analyzer:** Business-driven, end-to-end data and content profiling tool closely integrated into the enhanced WII suite.

that are slated to sit at the heart of the framework, and thus be core SOA server foundations.

19. UIMA Standard for Unstructured Content Takes Off: IBM researched/developed this new text analytics standards framework, with DARPA and leading universities, for 4 years. Now embedded in IBM IM products, UIMA was given to open source, has been adopted by 15 ISVs, and is now the de facto text analytics standard.

20. Master Data Management (MDM) – Key IM Service: This newer IM discipline is now an important, fast-growing middleware/services IM market segment. IBM's 1,000-strong MDM unit, assembled from acquisitions, now offers two powerful foundation MDM middleware products:

- **WebSphere Product Center:** For enterprise product master data/information management.
- **WebSphere Customer Center:** For transactional customer master data integration.

each leaders in their space, built on a WebSphere platform base. MDM is a vital information service in an IAAS SOA environment. We expect IBM to build out its MDM portfolio, solutions, accelerators and services fast. (*See Section 8*)

21. Analytics Demands Increase: Faster business cycles and soaring data volumes now need more real-time analytics, and more scalable, robust BI (*business intelligence*) data servers. Improved analytics are urgently required, to yield enhanced insight and understanding for better decisions. BI builds raw data into informational warehouse databases that support such analysis/reporting to find new knowledge, develop informed insight, and allow better decisions. (*Section 10*)

22. Analysis & Discovery Foundation & BI Delivery Tools

Needed: To deliver point 21 above, enterprises need two things: robust, powerful analysis and discovery foundations, and the right BI delivery tools. IBM (*Inventor of the RBMS core and data warehouse of BI*) is the leading provider of the former. Specialized ISV firms provide most of the BI delivery tools.

In the 1980s the dominant BI server was the IBM mainframe. UNIX and Windows/Intel became the BI platform fashion of the 1990s. In 2006, the System z9 mainframe... is again platform of choice for today's larger-scale, near-real-time BI workloads.

23. BI Delivery Tool Platforms Come Full Circle: In the 1980s the dominant BI server was the IBM mainframe. UNIX and Windows/Intel became the BI platform fashion of the 1990s. In 2006, the System z9 mainframe (*immensely more powerful and economic today*), is again platform of choice for today's larger-scale, near-real-time BI workloads. Likewise, BI tools have circled from host-based fat server thin-client (*1980s*), through fat PC client to thinner server (*1990s*), back to Web-browser, thin-client with fat-server (*2000s*) architectures, the latter ideally complementing powerful servers. (*Such as the mainframe.*)

24. Strong IBM System z9 Analysis and Discovery Foundation: IBM offers a robust, proven set of heavy-duty analysis and discovery foundation products for the mainframe, including:

- **DB2 UDB for z/OS database.**
- **QMF query/reporting.**
- **DB2 Data Warehouse Manager for z/OS.**
- **DB2 Intelligent Miner for z/OS data mining.**
- **DB2 Entity Analytics identity resolution.**
- **WebSphere Information Integration suite (z/OS) for BI integration.**

These, plus IBM's industry BI solutions, BI services, and partner BI delivery tools, provide complete enterprise analysis, discovery and BI solutions.

25. System z9 and zIIP – Ideal BI Server: Improved power and capability of System z9 mainframes, plus improved BI economics from the new zIIP specialty engine (*point 13 above*), make the mainframe today an ideal hub for large-scale BI, close to the core transactional systems also run on the platform.

26. Powerful DB2 Entity Analytics Suite: Powerful, new-technology IM entity analytics information services resolve crucial identity challenges widespread in intelligence, security, fraud prevention, and crime fighting, etc. The suite contains DB2 Identity Resolution, DB2 Relationship Resolution, and DB2 Anonymous Resolution. (*See Section 11*)

27. DB2 EA Suite DB2 for z/OS Support: DB2 Identity Resolution and DB2 Relationship Resolution now provide DB2 for z/OS Version 8 database support, bringing their powerful capabilities to bear directly on major mainframe-resident databases and applications under a SOA.

2. Today's Information Management Challenges and SOA

Introduction – Towards On Demand, Mainframe Central

With the faster pace of modern business, organizations seek to improve their core business processes, integrate their people, processes and information more closely together, and to connect their systems with their collaborative ecosystem of partners, suppliers and channels. This movement towards On Demand business is widespread, spanning across the globe, and across major industries, as organizations seek to increase their flexibility/responsiveness, whilst reducing costs and cutting wasteful activities. A central part of On Demand business is Information on Demand (*IOD*), which envisions the free flow of information that supports On Demand business processes across the enterprise.

This strategic direction requires major improvements in enterprise business applications portfolios and IT infrastructures to date ill equipped to support these challenging business-led demands.

For most large enterprises, the IBM mainframe platform continues to play a central and indeed increasing role at the heart of their IT infrastructures. These systems host/support many of the largest and most mission-critical enterprise applications, a host of new workload applications, and manage 80% of enterprise's machine-readable data. 2000-2005 saw a dramatic transformation, market re-evaluation and resurgence of the mainframe. This began with the zSeries in 2000, and further accelerated with the new-generation System z9 recently shipping from September 2005. We have assessed/evaluated the mainframe resurgence in several of our Papers. (See the "Related Software Strategies Research" Section on page 72, items 1, 2,3, 4, 6 and 9.)

CEO & CIO Priorities

After widespread CEO focus on cost reductions prevalent in earlier years of this decade, late 2004/2005 saw CEOs place top-line revenue growth as their foremost priority, whilst cost-containment continues to remain important. CEOs say their organizations must become far more responsive to rapidly changing markets, opportunities and threats. They consider attaining more effective integration of their people, core business processes, and of their wide sweep of information, are now a critical success factor in meeting the two top-priority goals above.

In the IT organization, these new CEO/boardroom priorities have also changed CIO/IT management team strategic priorities. These now are more closely aligning IT and business goals to enable and support new top-line revenue growth, whilst containing costs through more efficient core processes enabled by new-generation IT solutions. They also need to improve organizational responsiveness and agility through faster, more responsive IT solution development/delivery that mirrors rapid rates of change/advance that businesses now need. Finally, CIOs must deliver new-generations of IT solutions that empower and enable the organization's people, teams, partners, and broader ecosystem to be more effective, productive, and to collaborate more effectively.

In this new management climate, how are these demanding and challenging goals to be met? What will enable such a remodeling of business processes, applications systems and solutions, and IT infrastructure? What is the role in this of IM? SOA, discussed below, is the new application software technology that alone offers the transformative advances that can best help organizations meet these demanding new CEO and CIO objectives, but parallel advances are needed in IM.

Service Orientated Architecture – Transformative Software Technology

SOA is undoubtedly the most important software technology to gain industry-wide prominence since the Internet. The SOA-enabling Web services technology and standards have now matured, and today support broad SOA adoption by all businesses, large and small.

SOA promises new levels of speed, flexibility, integration and adaptability for business applications in every industry and in every size of organization, and has become IT's hottest topic/fastest growing market during 2005.

At its heart is the notion of composing new application solutions from a set of open-standards-based, reusable business software and systems "services" that can support new enterprise business process workflows, optimized to run faster and adapt more quickly to business needs, over the Internet. Figure 2 (on page 9) highlights the basics of SOA. This deceptively simple concept offers the most profound and far-reaching benefits. At its heart, SOA is centrally concerned with improving the level of integration of all three primary enterprise resources of people, processes and information.

IBM is now the clear industry thought and market leader in SOA, offering its comprehensive IBM SOA Foundation portfolio of middleware software, accelerators, solutions, methods, tools and services to support organization-wide SOA deployment. At the heart of this comprehensive and innovative offering stands the new WebSphere SOA Foundation set of middleware software infrastructure and development tool products, which provide a complete, flexible and incrementally deployable platform for organizations to implement SOA solutions.

What is Service Oriented Architecture?

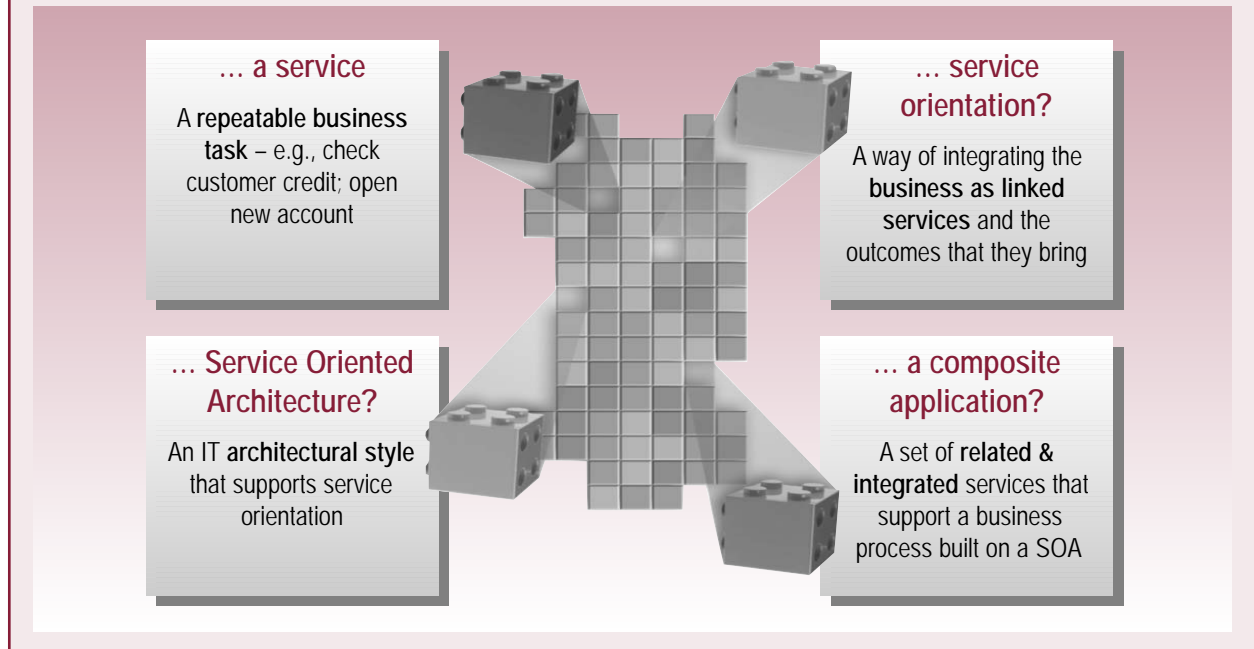


Figure 2: What is Service Oriented Architecture?

Our recent SOA White Paper (see the "Related Software Strategies Research" Section on page 72, item 1) fully evaluated and assessed SOA, the benefits it offers, its maturity and current adoption levels. The zSeries/System z9 mainframe, we found, fitted well into the new role of central platform node for enterprise SOA development/deployment. This Paper also provided this analyst's advice on how organizations should start on their enterprise SOA adoption roadmap journey, identifying the types of external resources, products and services that are required. It also closely examined and evaluated the IBM SOA Foundation, a comprehensive suite of SOA middleware software, tools, methods, accelerators, solutions and services designed to help enterprises rapidly but incrementally adopt and benefit quickly from SOA software technology. The Paper also included an extensive glossary of SOA terms and concepts, for those new to the topic. Our Executive Summary findings from this research are included here in Appendix A: this summarizes the case for SOA and for using the mainframe as a principal enterprise SOA node.

Enterprise Information Challenges Today Are Formidable

The information environment in most medium and larger enterprises has become large, complex and highly diverse, as the metrics in Figure 3 (on page 10) highlight.

A striking finding (top left), is that a full 60% of CEOs consider that their organizations need to do a better job of capturing/understanding information rapidly to make swifter, more fact-based business decisions. Equally striking, research showed only one third of CFOs considered their organization's current information was easy to use, tailored, cost-effective, or integrated. Both these findings indicate business executives seek information access and analysis improvements across their enterprises.

Volumes of machine-readable information have continued to grow exponentially. The average Fortune 500 company had 122 terabytes of disc storage on their data center floors in 2005. This had grown dramatically from the average of 7 terabytes of disc installed in 1996. This was a sustained compound 37% per annum disc growth, a rate expected to continue at least through to the year 2007, and probably beyond.

Usual perception of computerized information focuses on the conventional, structured, alphanumeric data associated with traditional IT applications. However, today 85% of these vast information volumes are of unstructured information. These now include documents, images, e-mail messages, databases, Web content, XML documents, drawings and designs, and the like. Throughout this White Paper, we use the term data to encompass both structured and unstructured information.

Organizations now spend a lot of their IT resources managing this flood of information. For example, on average 17% of total IT budgets went on storage, hardware, software, and people to manage it, in 2005. In application development, 30-50% of application design time was spent creating and managing copies of information.

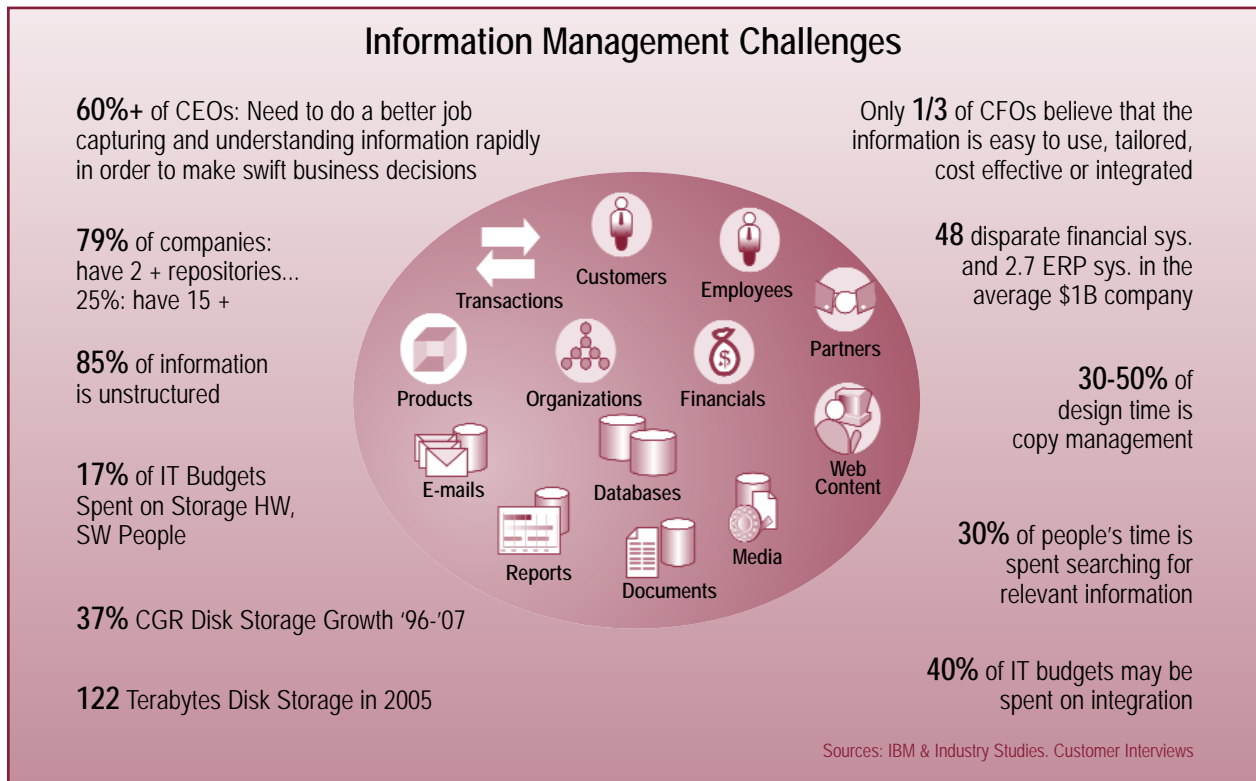


Figure 3: Information Management Challenges

Perhaps even more striking, across the entire organization, 30% of employees' time is spent searching for relevant information, which shows the scale of improvement in information accessibility needed. Research also showed that as much as 40% of total IT budgets are now typically spent on integrating applications, people and information.

Enterprise application portfolios grew diverse, complex and often overlapped, as new applications were layered on top of older systems through the last decade. For example, the average \$1B company was found to have 47 separate financial systems and 2.7 ERP systems in a recent study. This diversity and fragmentation of systems, and their associated information files and databases, is at the heart of today's giant enterprise information challenge and of the high costs associated with dealing with all this information.

Our Analysis

The above shows enterprises today wrestling with a real set of big challenges in managing, accessing, and in being able to fully use, their disparate and highly distributed pools of data strewn across the enterprise. This complexity is costing them a fortune in duplicated storage, wasted capacity, complex integration efforts to bridge the disparate information pools, in storage management, security, back-up, etc. Most importantly, real and easy access to, and availability of, critical data/information is still too difficult. This means substantial opportunities abound for new and better approaches to IM.

Urgently needed today are new and better ways of managing, accessing, distributing, aggregating and analyzing this deluge of data...

Urgently needed today are new and better ways of managing, accessing, distributing, aggregating and analyzing this deluge of data, both structured and unstructured. New technologies and solutions for turning this wealth of data into useful, accessible

information, when and where needed, for the new generation of SOA business applications that are being built, and for all groups of enterprise users that need to be better supported, are now essential.

The transformed and resurgent IBM mainframe (*zSeries and the new System z9 mainframe generation*) already plays a central role in hosting and managing up to 80% of enterprise machine-readable data, both structured and unstructured. It must therefore rationally play, and is well suited for, a much wider future role in this much-needed new IM environment.

SOA is the most important software technology/architecture to emerge to industry-wide prominence since the Internet. The SOA-enabling Web services technology and standards have now matured, and today support broad SOA adoption by all businesses, large and small. SOA promises new levels of speed, flexibility, integration and adaptability for business applications in every industry and in every size of organization, and has become IT's hottest topic/fastest growing market during 2005.

The vital, much-needed, new generation of IM must fit well within a SOA architecture, must support the new generation of SOA applications now rapidly proliferating, and must also support their open industry standards.

Indeed, many analysts, including Gartner Group, go rather further and state that, without a new level of IM that addresses and resolves the challenges mentioned above, efforts to implement a SOA architecture are likely to fail or fall short of their acknowledged high potential.

3. The SOA Application Generation, the Mainframe & Information Services

Service Oriented Architecture Takes Off

SOA has rapidly become the prime, unifying software architectural approach for the creation and deployment of business software applications on modern, distributed computing infrastructure, and is now supported across, and by, almost the entire, IT industry. SOA has risen rapidly to this now-central prominence within the last eighteen months to become the most important, fundamental change in how business application software can best be built, deployed and integrated and managed that we have seen in the fifty-year history of commercial computing.

The concept of service-oriented software architecture is not itself new. However, the recent developments of Web services technologies, and their accompanying open systems standards, have brought the major breakthrough that allows SOA to become such a transformative approach to business applications software architecture. In addition, leading vendors have now developed the robust and complete middleware software and development tooling, methods and best practices needed to support SOA development and deployment for mainstream customer use. These, in combination, now allow the SOA approach to bear centrally on, and to solve, many of the long-standing, and hitherto intractable, challenges of building, running, integrating and adapting flexible, responsive business applications that closely map to the transformed new business processes that we described above.

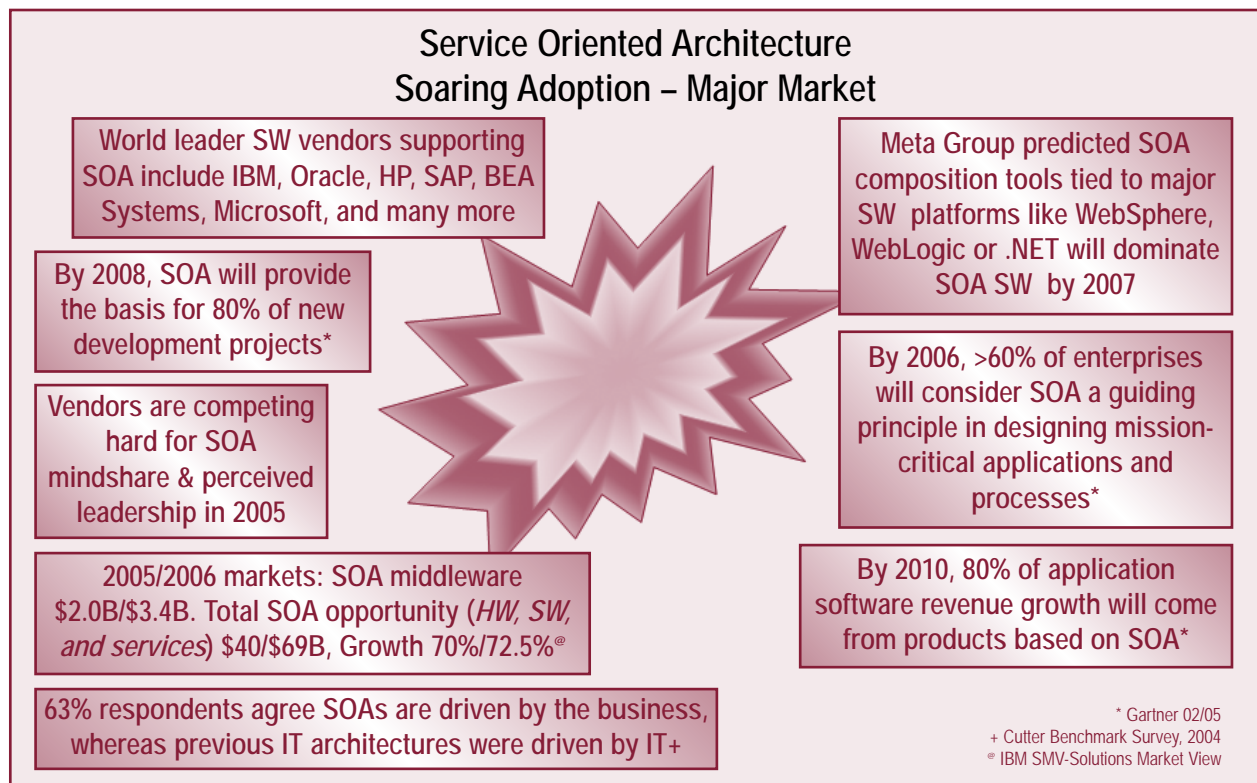


Figure 4: Service Oriented Architecture Takes Off – Some Indicators

Figure 4 shows a selection of important facts and assessments from other leading industry authorities. These highlight the large-scale importance and widespread adoption of SOA, and underline the rapid growth it is expected to undergo over the next several years. All the essentials for take-off are finally in place. We are now firmly in the early majority adoption phase of a decade-long industry migration to this next generation of SOA-based, business process-centric applications. This gives enterprises the ideal opportunity to re-evaluate what are the best platforms these should run on, and/or have their databases hosted upon. Analysts agree that enterprises should not replicate the overly complex, distributed, multi-tier infrastructures installed with client/server through the 1990s for SOA over the next decade.

Indeed, many have made major efforts to consolidate/simplify their IT infrastructures already, and this continues to be a high priority. Centralizing, simplifying, securing, virtualizing and automating the IT infrastructure are all essential to make it more scalable, flexible, resilient, and easier to operate/support, both for existing and for SOA applications.

SOA provides the software architecture to enable far better integration, easier migration to the optimal host, and presents a good natural opportunity for a substantial platform mix change. There is little question that this swings the balance of argument back in favor of a more centralized, virtualized hub platform with the right characteristics, such as mainframes, and we found the arguments for this were convincing in our recent SOA on zSeries/System z9 White Paper.

Mainframe Resurgence Positions System z9 for Major SOA Role

2004 saw the 40th anniversary of the IBM mainframe, the central business application-computing platform for thousands of larger firms for up to four decades. Today, mainframes host several trillion dollars worth of business applications developed over this period and process hundreds of billions of business transactions per year.

The IBM zSeries mainframe saw a dramatic resurgence in the market, and in industry-wide esteem and respect, between 2000 and 2005. It enjoyed a quite remarkable return to favor, gained 16 points of "over \$250,000 server" segment market share, and saw more mainframe MIPS sold in the last four years than in the previous 36; 75% of these devoted to running "new-to-mainframe" workloads. The platform had been transformed by deep changes, and several \$B of IBM investment, continuous innovation and development. We found the principal drivers of this striking resurgence of the mainframe platform were the factors shown in Figure 5.

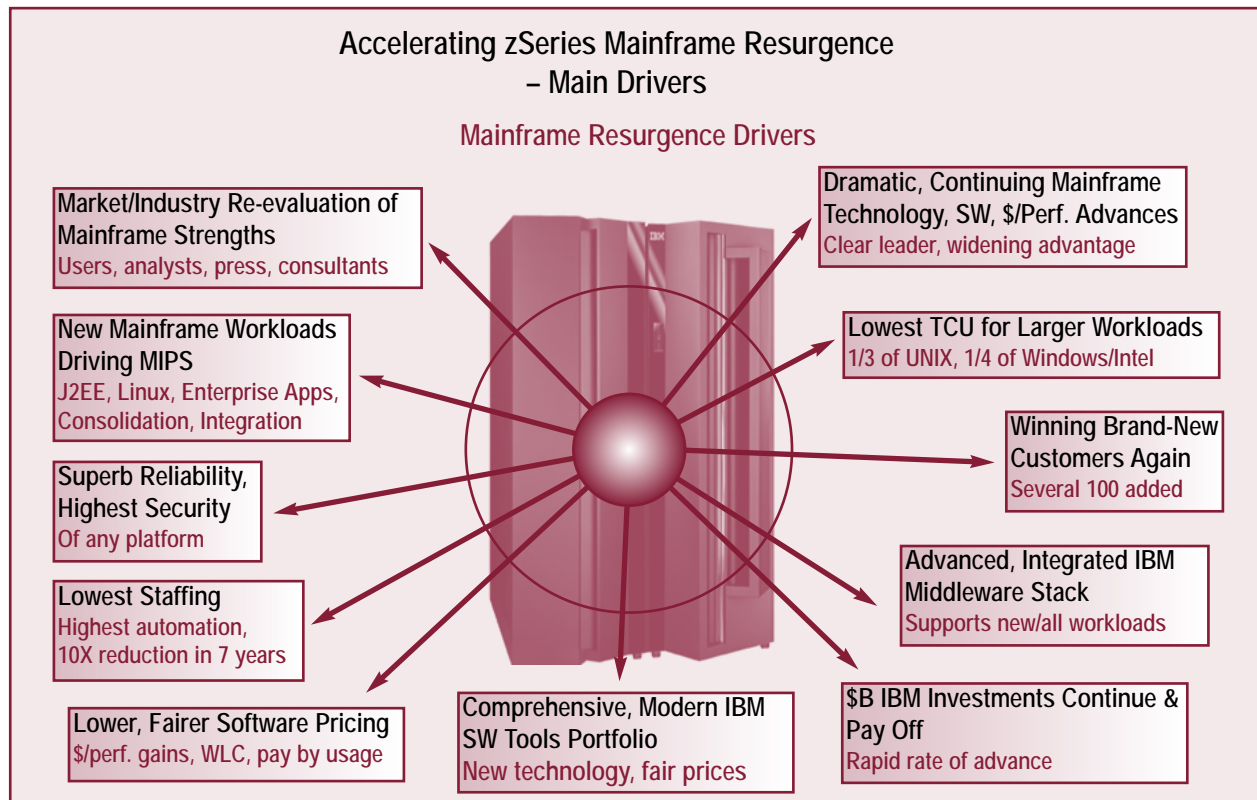


Figure 5: Accelerating zSeries Mainframe Resurgence – Main Drivers

This market revival began in 2000 with the introduction of the zSeries, and by 2004 it had strengthened to make the mainframe a dominant player in the market again. The zSeries mainframe had again become widely recognized as the "Five Star" enterprise computing platform, largely because of the unique business values and technology leadership it continued to offer to support business transformation to On Demand.

The newest mainframe generation, the System z9, which started shipping in September 2005, brought further major advances, with focus on supporting surging growth in collaborative computing, now enabled by extensive support for SOA, under which the next generation of business applications are now being built. We recently published an in-depth study of the new System z9 hardware and software (see the "Related Software Strategies Research" Section on page 72, item 2), and our main findings from this research are summarized in Appendix B.

SOA Virtualization of Applications – Platform Independence Strength, Not Platform Indifference

One of the most important characteristics of SOA application architecture is that it insulates the developer using a software service from the details of its implementation, and thus from the technologies of the platform that service is running on. This powerful “virtualization of applications” attribute of SOA allows application components from multiple platforms to inter-operate and collaborate in new business process workflow-supportive composite applications over the Internet much more easily, and without the developer requiring deep knowledge of its implementation. This platform independence and loosely-coupled integration model is a major advantage of SOA. It supports its flexibility, rapid development and incrementally implementable characteristics that are so important.

However, this should not at all be construed as meaning that SOA developers should be indifferent to the nature and attributes of the platform that their important SOA services run upon. Nor should it be mistaken as implying that under SOA all platforms suddenly become equally capable, or can all deliver the same qualities of service. This remains, as ever, far from being the case.

Major differences in the achievable qualities of service and the attributes, performance and behavior of different systems platforms remain, and indeed have actually widened substantially in recent years. The zSeries/System z9 mainframe platform, for example, is most highly differentiated by its high qualities of service in many areas, based on IBM's continuous innovation, development and investment over the past five years. We have analyzed this in several earlier Papers and Reports (*see the “Related Software Strategies Research” Section on page 72, items 1,2,3,4, 6 and 9).*

New System z9 109 Mainframe – Principal “Information as a Service” Node for Data and SOA

On July 26th 2005, IBM introduced “one of the most sophisticated computing systems ever”, its new IBM System z9 109 high-end mainframe (*pictured in Figure B1 on page 64*) first of a new generation of mainframe systems. IBM said that it had designed and optimized the System z9 109 to serve as the principal enterprise hub platform in the new era of collaborative computing. Specifically it intends the mainframe to serve as a principal enterprise hub for data/information and as a main enterprise node for SOA amongst other new roles. Our assessment of the strategic positioning of these recent new mainframe systems is shown in Figure 6.

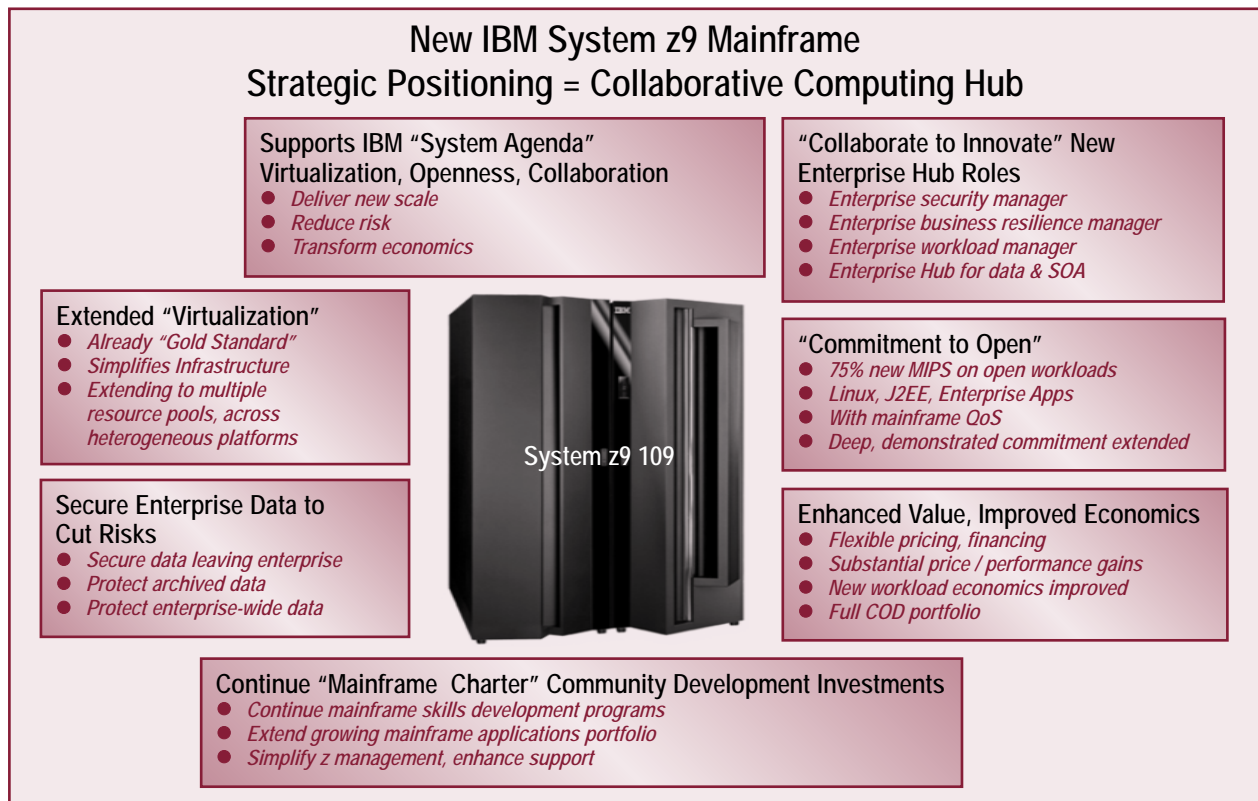


Figure 6: New IBM System z9 109 Mainframe Strategic Positioning – Collaborative Computing Hub

Providing twice the capacity and memory of its successful z990 predecessor (*often tagged as T-Rex and introduced in 2003*), the System z9 109 emphasizes new levels of security, enhanced reliability, and extensions to its leadership virtualization capabilities. The System z9 109 was the outcome of a three-year-long, \$1.2 billion development effort said to have included 5,000 global IBM system designers, engineers, security experts and software developers, continuing the sustained investment in mainframe innovation and technology IBM has consistently made since the late 1990s.

The new System z9 109 mainframes will not only fulfill their well-established, current roles at higher levels of scale, with reduced risks through extended security, and with further improvements in mainframe economics. IBM now envisions the mainframe assuming important new **“Principal Enterprise Node”** roles, notably in the context of this White Paper, as **the main enterprise node platform for SOA and as the hub for data/information**, sitting at the heart of cross-ecosystem, collaborative computing. To this role, the mainframe brings its outstanding virtualization, QoS, extreme security, and advanced workload management to control and optimize heterogeneous, multi-tier infrastructures, not only within the enterprise, but also across customers' ecosystems of collaborating partners.

Increased Use of the Mainframe Node – Route to Slash Soaring System Management Costs

...costs of system management and support for these systems has also rocketed, from some \$33B in 1996 to a projected \$140B in 2008 on current trends...

Complex, multi-tier, distributed IT environments are responsible for the soaring costs of enterprise IT infrastructure support and management. For example, the cumulative total number of servers installed worldwide has soared from just 6 million in 1996 to an astonishing projected 35 million in 2008, a 6-fold increase. The annual worldwide costs of system management

and support for these systems has also rocketed, from some \$33B in 1996 to a projected \$140B in 2008 on current trends, over three times the expected expenditure on new servers themselves.

A large reduction in these costs is the main way to release resources for new developments, and requires major systems technology innovations and advances to allow users to drive these costs down sharply. Whilst infrastructure consolidation and simplification has been a major theme for several years, much clearly remains to be done to bring this runaway, server profusion-driven systems management cost growth back under control, and drive them down to more acceptable levels.

There are therefore compelling customer economic and staff resource benefits to be realized when the powerful business value attributes of the latest mainframes extend their power, control, efficiency and security out across much more of wider multi-platform enterprise IT environments in such new “principal enterprise node” roles. The new System z9 109 represents a substantial step toward delivery of this “Grand Challenge” vision. To implement it, IBM has extended mainframe server hardware, operating systems, middleware run-time software engines, development tooling, systems management, virtualization fabric technologies, as well as the storage and networking capabilities, in concert. Such an integrated, holistic approach was essential, and has been successfully implemented. We found this to be of considerable direct benefit to mainframe customers.

These moves now extend acknowledged and unique mainframe QoS across the customer's wider existing IT infrastructure, to support large-scale collaborative computing under SOA application architectures. IBM claims the best way to do this is by creating an extended and simplified virtualization foundation that extends the mainframe's “Gold Standard” virtualization strengths out across the enterprise. Another essential is that this effort must be entirely based on open industry standards throughout, vital to enable other vendors, their systems and software, to participate; preserving and adding value to customers' existing investments. By providing robust, secure and scalable mainframe data management, data-serving, data storage support, informational processing, and management capabilities as services to other platforms, the whole infrastructure QoS can be improved, and management costs sharply reduced.

Mainframe – Main Enterprise SOA Node Platform and Data/Information Hub

IBM proposes that its latest generation of mainframes should become the main enterprise node platform for SOA and hub for data/information. The mainframe already has unrivalled capabilities to manage massive data volumes securely and, with new middleware software releases through 2005/early 2006, is now well equipped to support the next generation of composite applications integrated under a SOA.

Many enterprise applications (*and their data*) reside on the mainframe platform today. Many of these are already being restructured into new composite SOA applications. It is a natural extension to have the mainframe serve as central node for managing and maintaining the new heterogeneous SOA applications and data of collaborative computing out over other platforms. In this extended role, the System z9 109 will also further extend its established position as the principal business integration node for the enterprise, which its middleware product set has been supporting for several years. With the IBM System z9 109, IBM delivers significant technology advances that implement this “central SOA node and data hub” strategy and new role.

Another factor arguing for this new mainframe role is that more than a decade of information warehouse and datamart development/deployment has multiplied the transactional distributed systems complexity problem. These added numerous additional informational databases, and the multiple, complex ETL processes needed to create and update them. These warehouses and datamarts were intended to provide cleaned, consolidated management information, for BI, decision support, and for deep analysis and data mining. Many of these are scattered across distributed platforms as well. The much-increased data-serving capacity and power of the mainframe, combined with the most sophisticated storage, and with the industry-leading information analysis and discovery foundation software capabilities (*assessed in Section 10*) of the mainframe today makes a strong case. These factors, together with today's far lower mainframe costs, mean it now makes strong technical and economic sense to re-centralize many next-generation data warehouses on the mainframe enterprise hub. This can better provide an overall enterprise BI view, and can run BI at nearer to real-time speed in the faster moving On Demand business environment.

The IBM SOA Reference Architecture and Information Services

IBM has developed, refined and published its SOA Reference Architecture, which defines the primary services required for enterprise SOA deployment, and how these interact together. This Reference Architecture is shown in Figure 7.

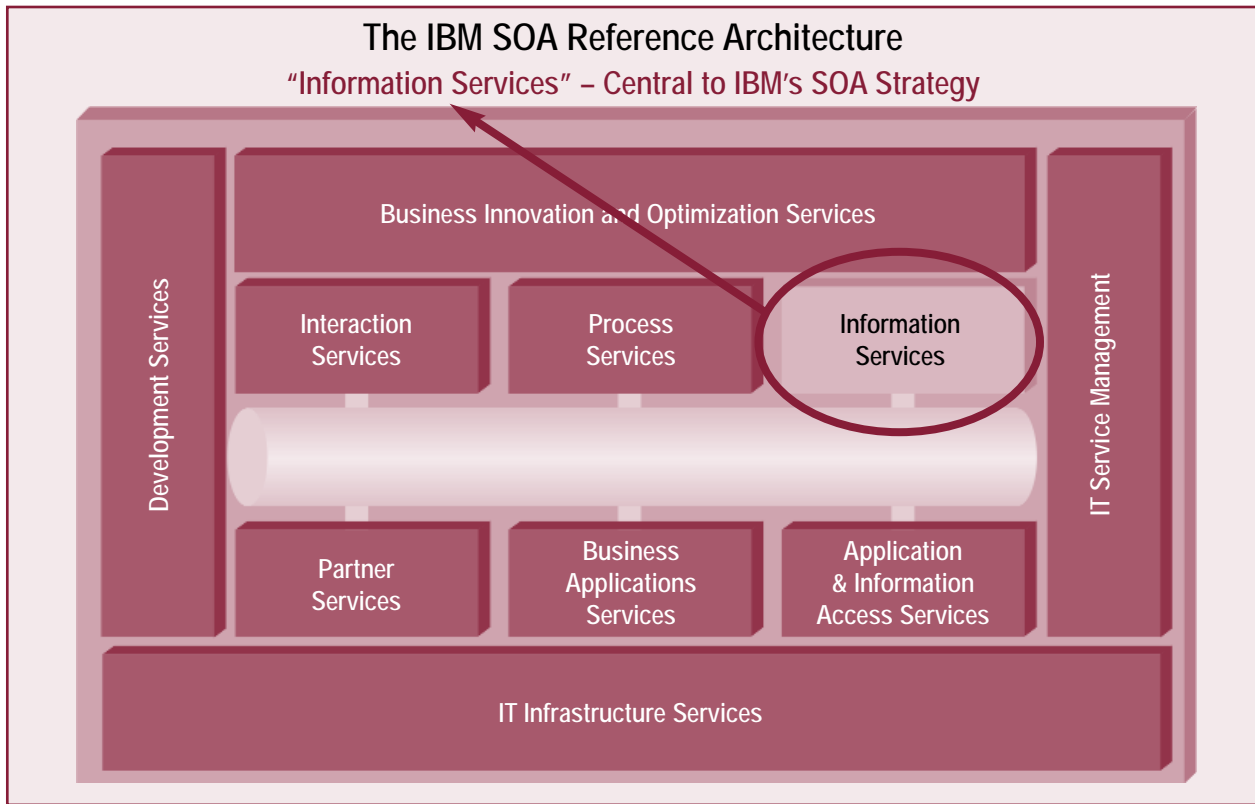


Figure 7: The IBM SOA Reference Architecture & Information Services

The IBM WebSphere SOA Foundation middleware/development tools suite that implements much of this architecture (*assessed in our White Paper – see the "Related Software Strategies Research" Section on page 72, item 1*) are all available now (*or are due out in 1H 2006*) for the zSeries/System z9 mainframe platform. In that Paper, we examined in some detail the Business Innovation and Optimization Services, the Process Services, the Interaction Services, the Development Services and the Business Application Services elements of the architecture, and the new software products that now support these services on the platform. Our summary assessment of the WebSphere SOA Foundation products that support this architecture on the System z9 is given in Appendix C.

An equally important element (*and the main subject of this current White Paper*) is the "Information Services" component of the architecture, highlighted by the ring in Figure 7. The notion here is that "Information" should be provided as enterprise services that can participate in, and support the SOA as a whole. We define and assess the concepts of IAAS in Section 4. Information Services are therefore a core part of IBM's overall SOA strategy. IAAS promises to deliver the major IM advances needed to meet the demanding challenges highlighted in Section 2. IAAS can also be adopted and implemented independently of, and ahead of SOA applications, with major business and systems benefits in its own right.

Our Analysis

As enterprises across the world move towards On Demand, seeking to become more responsive, flexible and adaptable to constantly changing markets and demands, they are integrating their operations internally and externally with customers and partners with more interactive or collaborative processing systems and applications.

Building enterprise business applications has always raised difficult, even intractable challenges, much-compounded by the spread of complex, distributed, heterogeneous IT environments, and by the much faster recent pace of business change which also demands a faster IT applications response. The wide gap between what business leaders needed and what IT organizations could provide could not continue.

Earlier decades saw many application development technology advances, but none had decisively or comprehensively resolved the challenges. A more transformative software architecture that could break free of the constraints and overcome the challenges was long needed.

SOA has now emerged to become, in our assessment and that of other analysts, that much-needed transformative business software architecture, the most significant such advance yet to emerge. Many indicators, including widespread software industry support, analyst support and uptake projections, the substantial scale and pace of customer SOA adoption, a growing number of

SOA has now emerged to become, in our assessment and that of other analysts, that much-needed transformative business software architecture, the most significant such advance yet to emerge.

success stories, and the size forecasts for SOA-related IT markets, all confirm the momentum and importance of SOA.

The unique business values and significant further advances offered by the new System z9 mainframe, stake a claim that the platform should now serve as the principal enterprise node for SOA, and as the hub for data and IAAS, at the center of customers' collaborative computing/SOA developments.

IBM's SOA Reference Architecture provides a clear framework and structure for enterprise SOA. One of the central services of the architecture is "Information Services". There are major benefits that directly address the deep and wide-ranging enterprise information challenges we described in Section 2 in treating IAAS.

4. Introducing Information as a Service – Crucial SOA Enabler

Information on Demand – An Introduction

IBM's On Demand strategy promised to enable businesses to closely integrate their people, processes and information across the enterprise, and across its ecosystem of partners, to speed up core business processes, reduce costs and enable more rapid "sense and response" to changing market and business conditions, opportunities and threats. A central component of this broad strategy is IBM's **Information on Demand (IOD)** strategy. The Information On Demand vision is about creating enhanced business value and reducing risk by capturing, creating, integrating, analyzing and optimizing all types and sources of information throughout their lifecycle, and thus enabling the organization to act more quickly and systematically upon that information. Information on Demand provides powerful business benefits across the entire organization, including:

- Managing risk and streamlining compliance efforts.
- Providing methods and tools to better understand and manage business risk.
- Ensuring employees make their decisions based on up-to-date, accurate and complete information.
- Enabling the organization to transition from running on estimations, towards running on accurate measurements.
- Enabling organizations to transform their business practices from reactive tactics to real-time monitoring and control.
- Opening the organization's information securely and efficiently to all the suppliers, partners and customers in the business ecosystem, enabling improved collaborative processes that benefit all.
- Gaining control of Master Data across the enterprise.
- Reducing overall IT costs by simplifying today's unduly complex data and information management tasks.

These are powerful, compelling and wide-ranging business benefits that are widely applicable, and are now practically realizable, by organizations in every industry and geography.

IAAS Investments Enable Breakthrough Information on Demand Realization

IBM has long been the market leader and prime innovator in information management middleware software, tools and services, best known for its database leadership with DB2, IMS, and for many other long-established and widely-used IM products. However, over the last three years, the giant has invested over \$3B, in new research-based in-house development and in a string of 15 complementary information management software acquisitions. These have now all been integrated and combined with IBM's own products to provide the market-leading, comprehensive and highly-integrated IAAS offering that is now emerging to deliver all the capabilities required to realize the Information on Demand vision.

With the early availability of recent major IAAS developments, IBM has unquestionably now designed the industry's first complete information infrastructure for delivering IAAS/Information on Demand, an infrastructure that will help businesses truly leverage all their information assets, throughout their life cycles, for strategic advantage. The full IAAS capability will be rolled out over 2006 and 2007, with major steps arriving in 2006.

This is a major industry breakthrough, providing the first comprehensive approach to unifying and exploiting information fully. It overcomes the long-standing information fragmentation, and is a principal enabler for the SOA approach to next-generation applications that is rapidly taking off. It can also provide powerful solutions to many of the information challenges in

As an example of the much-extended scope of Information as a Service, consider a customer history. The first view that springs to mind is usually of the standard transactional history of orders/business written with that customer that is held in a sales order processing database as conventional database table rows. However, the enterprise has much wider information available. For example, purchase orders from that customer, the customer contract or contracts, supply chain transactional data, order-specific e-mails, and historical purchasing patterns. With Information as a Service, all this new and historic content, combined into a unified and more comprehensive customer information set, can be delivered in the specific context a user or a business process requires, based upon security access constraints, the user role, and the task being executed.

By providing information within a specific context the user is provided with all the information they need to act within their role, but without wasteful searching of large amounts of potentially irrelevant data, or without missing crucial aspects.

The essence of Information as a Service is that neither the people nor the processes requesting information should need to care what format or system the information came from, so long as it is timely, complete, accurate, and in-context, providing a single version of the whole truth.

Additional business value and flexibility can then be achieved by applying analytic techniques to the base information that can enable predictions about what is likely to occur in future. Such further analytics can provide recommendations or prescriptions for actions the user, or the business processes, should take for best results. This analytics based insight, which enables appropriate and timely action to be taken, is the true power and the real business leverage of accurate, integrated information.

IBM's New "Information as a Service" Vision & Offerings

IBM's Information as a Service vision and offerings have been designed to realize and deliver the Information on Demand benefits above. It will allow information to be made more easily available in an open and easily accessible manner, regardless of its format or location, across the enterprise. **"Information as a Service" (IAAS)** also fits tightly and naturally into the services-based approach of SOA. IBM Information Services and Information Accelerators supply the core services and a set of proven ways of speeding deployment. The former cover the services needed to manage and access all forms of information in a SOA approach. Based entirely on the open industry standards that are emerging (*which IBM is driving*), IAAS will render enterprise information truly open (*via these industry standards*) and easily accessible (*in minutes not months*) to people, applications, and business processes. This improves systems and information agility, because information will no longer be closely tied to proprietary systems, formats, or technologies, avoiding vendor lock-in to any single database, operating system or server platform. IBM will help customers create information services by taking data (*structured and unstructured*) from its current passive, closed state, applying open standards technologies and industry know-how to turn them into active "services" that handle business processes. These can be used internally, externally and even be sold as live services.

Organizations are increasingly exploiting SOA to introduce much-increased flexibility into their business applications and processes by rearchitecting their new-generation applications into smaller, more accessible and reusable software services. In a SOA, business processes and the IT infrastructures that support them are treated as secure, standardized services that can easily be reused, managed and composed in new ways to rapidly address changing business needs and priorities. IAAS is therefore a central foundation for achieving more value from enterprise information assets. Gartner strongly advises that organizations should develop an IM strategy as part of their SOA architecture. It says organizations will waste their investment in SOA unless they have enterprise information that SOA can exploit.

Figure 8 shows a conceptual overview of the IAAS vision.

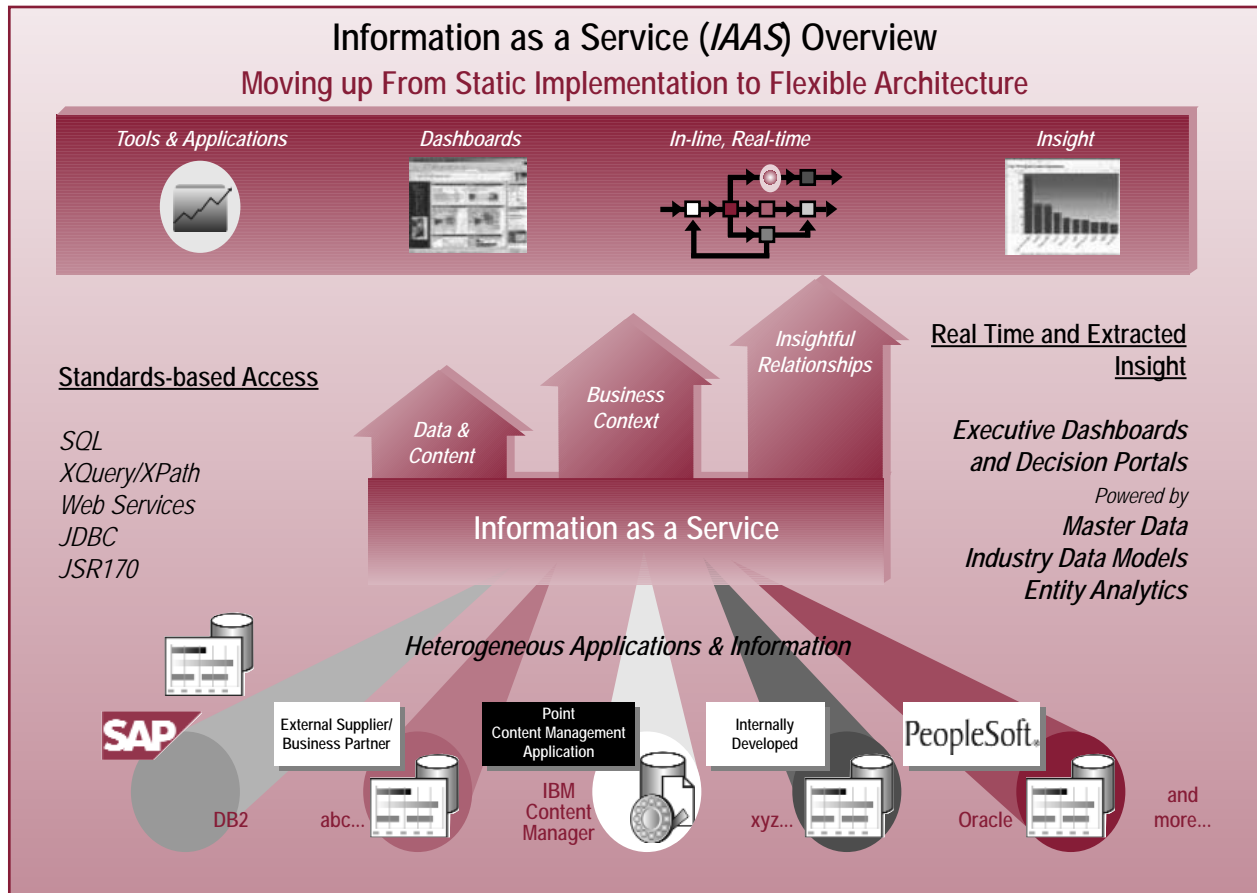


Figure 8: Information as a Service Overview

At the bottom of the chart are typical heterogeneous enterprise information sources and applications. In this example, these include enterprise packaged applications (such as SAP or PeopleSoft), internally developed applications and databases, external partner and supplier applications, content management applications, and DB2 or Oracle database applications that are the raw sources for information. In practice, many other combinations of sources may be present in individual enterprises.

The center of the chart shows three levels of value-add that IAAS brings. These are **access** to data and content at the most basic level, **business context** adding further value at the next level, and **insightful relationships** identified in the data by analytics at the highest level.

Crucial to the IAAS model, shown centre left, are the open information access and management standards that will make IAAS a truly open and inclusive approach, including SQL, XQuery/XPath, Web services, JDBC and JSR170. IBM is playing a leading role in driving the newer of these standards to adoption and maturity.

Shown at the top of the chart are the different forms into which IAAS will deliver enriched information. These include into new SOA applications, into graphical executive dashboards or decision portals, as on-line real-time analytic results, and more broadly as new and revealing insights into relationships revealed by the information and analytic techniques.

The broader IAAS goal is to help customers make money by finding new business opportunities that exploit hidden gems underneath today's exploding volumes of data, such as customer information, supply chain patterns, customer service and public opinion about product quality.

Information as a Service Ecosystem

IAAS requires a strong, supportive systems and software ecosystem, which IBM has been actively building out for many years, and an overview of this ecosystem is shown in Figure 9. In the center, a base layer of data assets and systems is the foundation, and this includes servers, storage systems, filers and networks. These run and support the other layers above. IBM has long held technology and market leadership in the provision of all the elements of this layer.

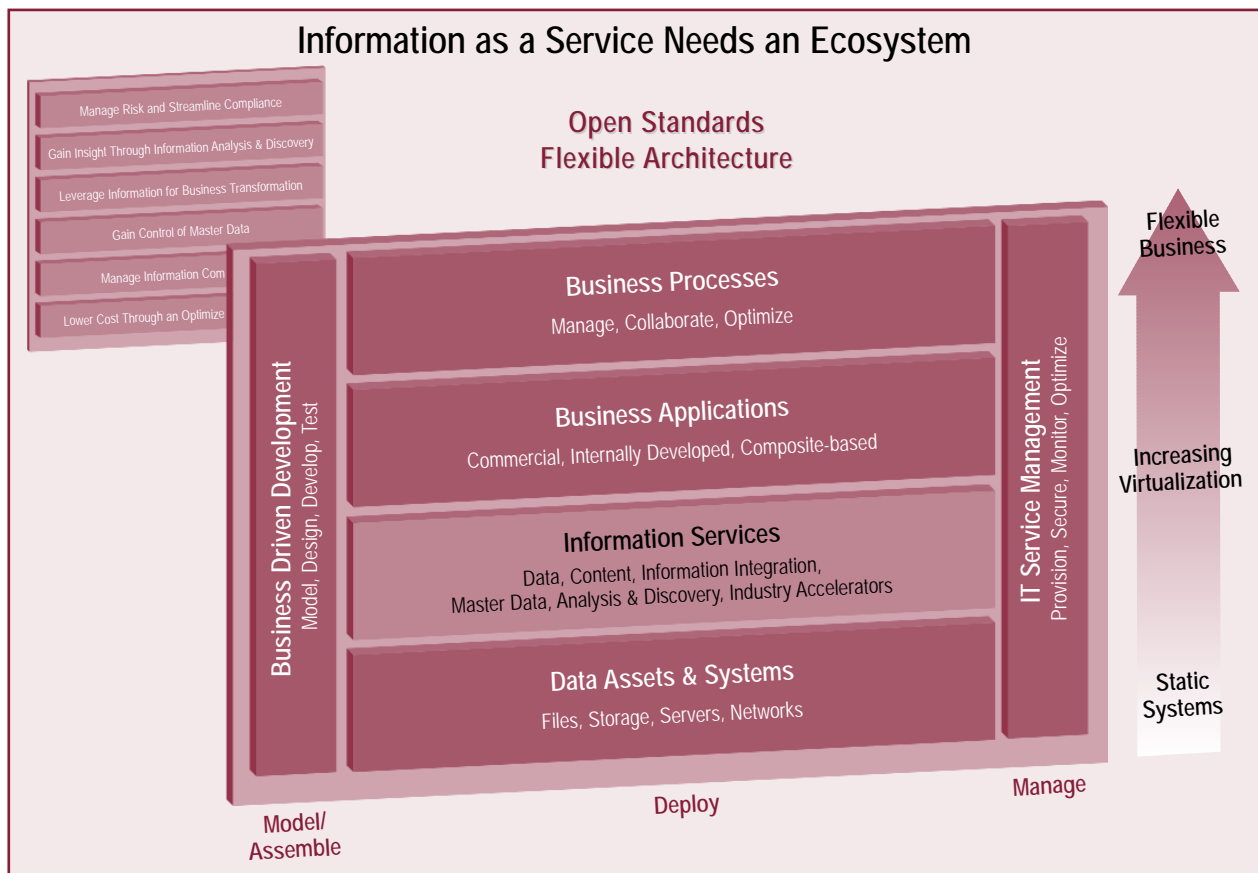


Figure 9: Information as a Service Needs an Ecosystem

The next layer of Information Services are the enablers that deliver Information as a Service, run on top of and access the data assets and systems in the base layer, and which also provide the information services to support the business applications in the layer above. These business applications themselves in turn support the core business processes that companies are seeking to modernize and transform. Left to right on the chart, we need to be able to handle the entire lifecycle of each of these layers, from their initial modeling/assembly/creation, into their deployment and use, and to manage them effectively over their lifecycle.

On the right, IT Service Management is essential to monitor, control and optimize all the layers in operation, to meet IT service level goals and objectives. The vertical arrow on the right also highlights how we have moved from static systems, through an intense and continuing process of resource virtualization (*of servers, storage, applications (SOA), and now information (IAAS)*). This virtualization advance is increasingly enabling the much higher level of business flexibility and adaptability needed for On Demand business.

On the left hand side of the chart, a new level of Business-Driven Development is needed to better align business and IT, to ensure that business processes are modernized and optimized effectively, and that the new business applications support and enable these processes effectively. IAAS plays a central role in delivering the right, current, accurate information to these improved business applications and processes, regardless of its underlying source or platform.

Information Services and Information Accelerators

To deliver IAAS, IBM has defined and developed a complete infrastructure of five broad categories of **“Information Service”** that combine to provide a comprehensive, enterprise-wide, standards-based set of services for accessing, managing and taking timely and proper action upon information effectively.

In addition, it has developed, packaged, and now offers, a set of best-practices-based **“Information Accelerators”**. These are industry-specific technology best practice models, maps, schemas, policies, etc that IBM or partners have developed and that make it easier, faster and cheaper for customer organizations to adopt and deploy IAAS strategies throughout their enterprise. These decrease the time to value, and reduce the risk associated with providing information in a SOA, providing best practices approaches that support innovative applications and processes, and that have been developed by IBM or its partner network.

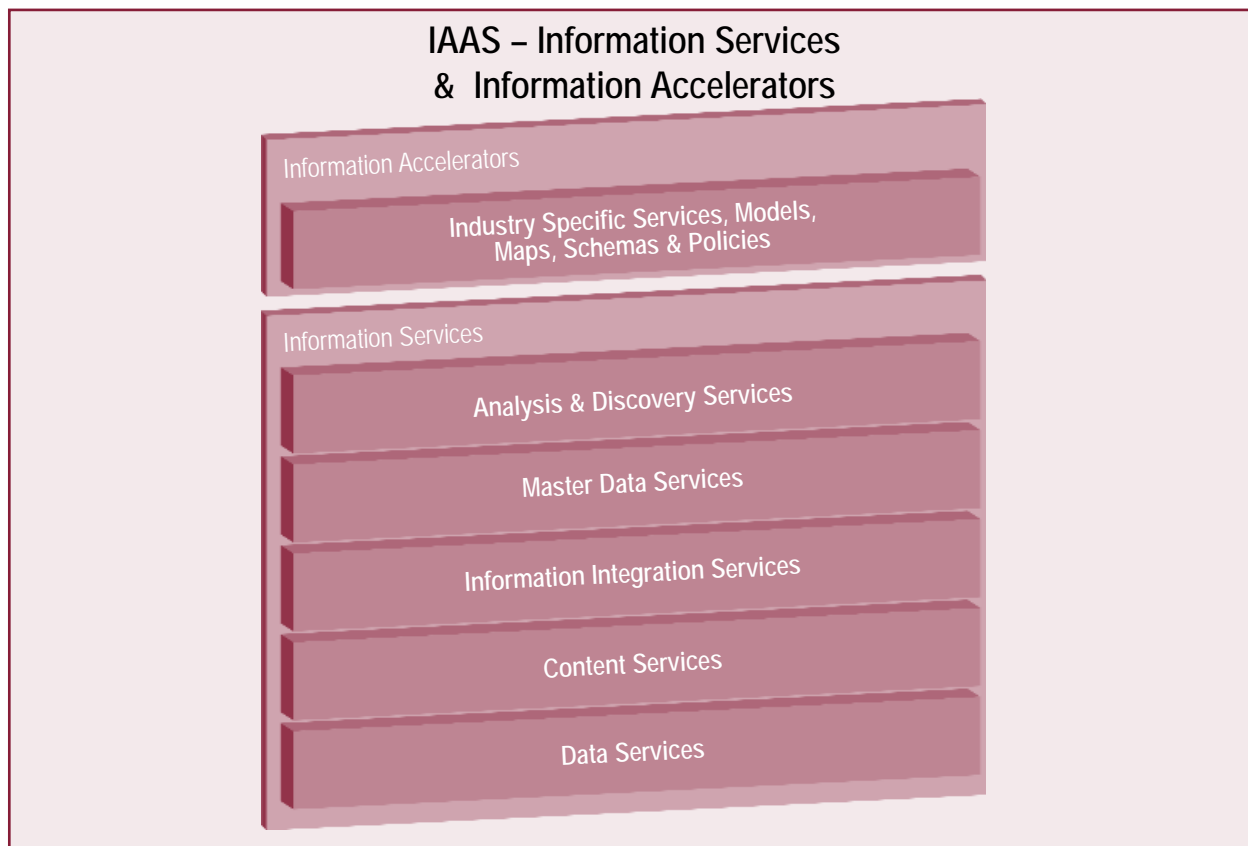


Figure 10: IAAS – Information Services and Information Accelerators

The five principal **Information Services** that provide the building blocks for a next-generation IM infrastructure are shown in Figure 10, and are discussed and more fully defined below:

- **Data Management Services:** These are the most familiar and longest-established information management services provided by IBM's leading and fast-evolving database servers which form the core foundation for all Information Services.
- **Content Management Services:** Information infrastructure services that include management of those documents, Web content, imaging, reports, digital assets, e-mail, and records, etc. declared to be critical assets, throughout their lifecycle. IBM boldly (*but accurately*) claims that these services enable its users to access, protect and deliver more forms of information than any other vendor.
- **Information Integration Services:** Services that deliver accurate, consistent, timely and coherent business information to people and processes enterprise-wide. These services allow real-time, integrated access to business information regardless of location or format, providing the broadest set of information integration capabilities to semantically align information across disparate sources.
- **Master Data Management Services:** Flexible framework middleware services to support enterprise structured and unstructured data and business services that are aligned with key business processes. Enables organizations to create and manage consolidated "master" data for customer and product information (*and other prime entities in future*) across heterogeneous environments.
- **Analysis & Discovery Services:** Includes information analysis and discovery services that help identify and analyze the information available. Also includes foundation services specifically designed to prepare and act on structured data for querying, reporting and analysis, enabling powerful analysis through leading BI tools (*third party*) from fully-integrated data to derive new insights for competitive advantage.

Information services are created and deployed by the information management professionals who understand the information best. They select the information that they wish to share and make available, specify the rules that govern that information, and publish flexible services that provide that information, often with a few simple mouse clicks.

Systems developers can then easily connect to, call and use these information services in their new SOA applications, without having to understand the specific APIs and semantics of the underlying source systems. Data processing consistency and quality is then automatically maintained by the services that provide the data themselves. This enables application developers to focus on the design of new business processes and applications, secure in the knowledge that they are using consistent, accurate data provided by the well-defined and approved set of enterprise information services that have been made available.

IAAS, using all these types of Information Services, will be a real advance over the current situation in most enterprises, where the daunting complexity and diversity of information resources is often chaotic. This complexity means that business-critical information is currently often disaggregated, disorganized and difficult to understand – rendering it difficult and costly to provide the consistent and timely information needed for effective business management.

Information in the SOA Lifecycle

Within an IM architecture supporting a SOA, all information needs to be managed across its entire lifecycle, and IBM now offers a full range of capabilities for deploying and managing information within a SOA that span the whole SOA and information lifecycle.

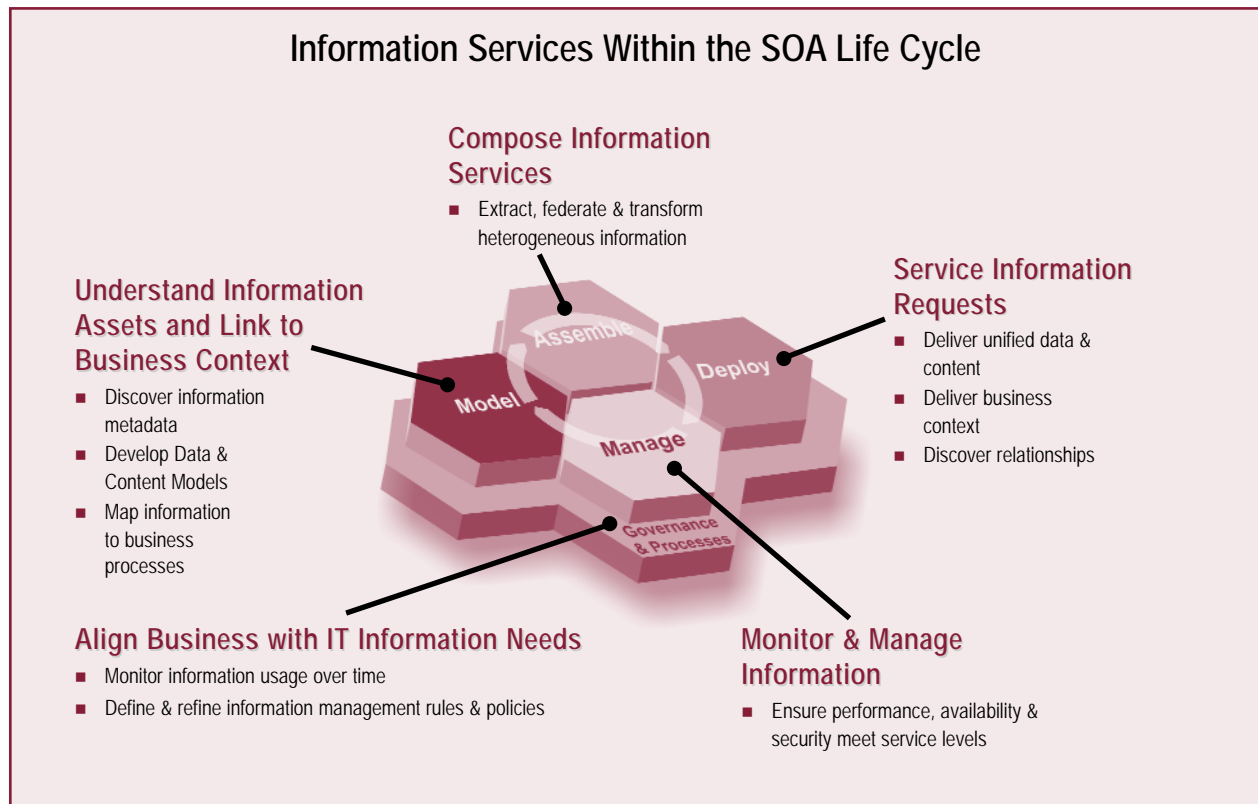


Figure 11: Information Services within the SOA Lifecycle

The SOA lifecycle runs **Model-Assemble-Deploy-Manage-Governance, Process**, and the supported information operations for each stage are shown against each heading on Figure 11.

The Virtualization of Information

IT professionals are now familiar with the widespread application of server virtualization to enable, amongst other valuable benefits, the now widespread de-coupling of server workloads/images from physical server hardware resources. Server virtualization enables far more efficient and flexible utilization of server hardware capacity, and the rapid provisioning of new server images via server partitioning. Indeed, the zSeries and System z9 mainframes offer the industry's "Gold Standard" in server virtualization, with all system resources (*CP, memory, I/O, and internal networking*) fully virtualized.

More recently, we have also seen widespread adoption of storage virtualization technologies, which enable disparate storage systems on a SAN to be managed as a single virtual storage pool, with similar major gains in improved storage utilization, ease of data migration, and considerable management effort and storage cost savings.

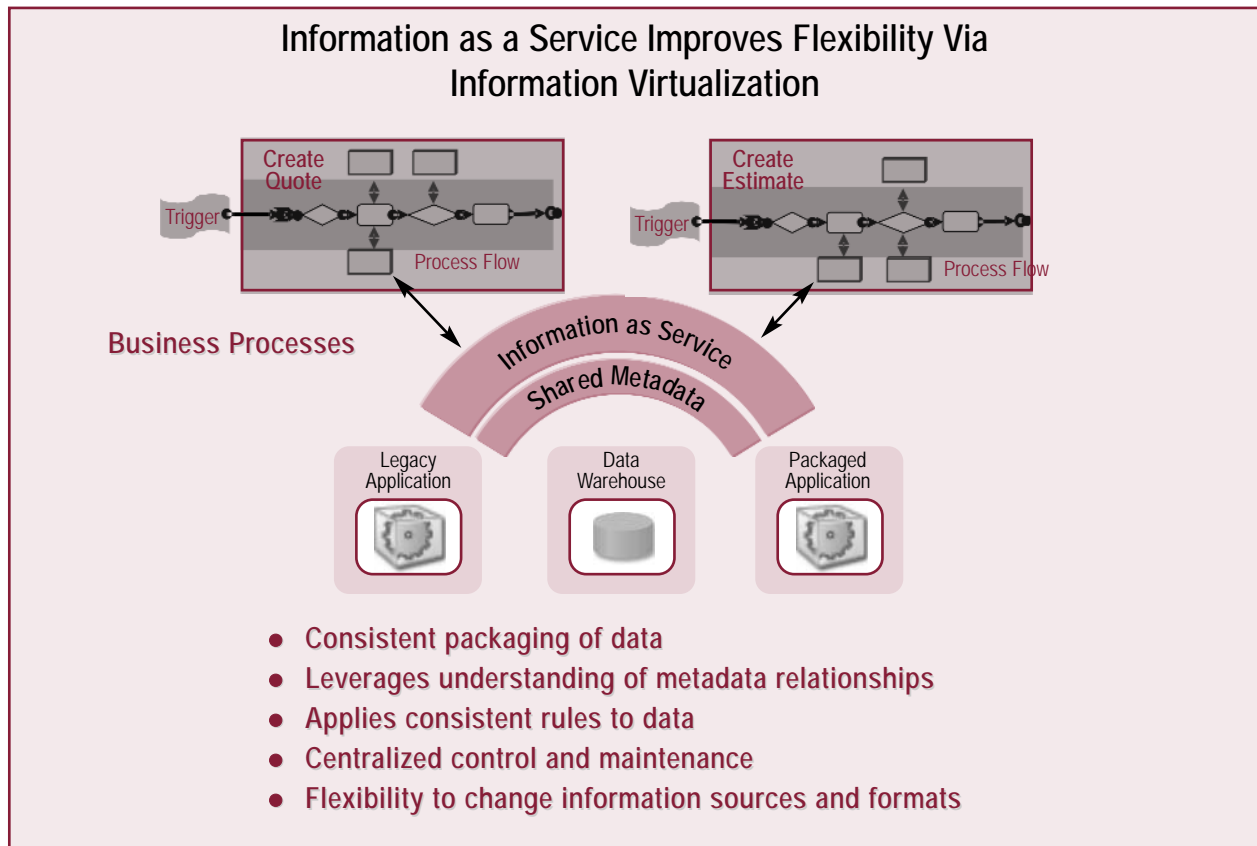


Figure 12: Information Virtualization by IAAS

IAAS essentially extends virtualization technology to the IM sphere. IAAS and its set of Information Services provide a higher-level abstraction layer, de-coupling services from their underlying source databases, files, applications, and/or APIs. This enables much simpler, more secure and transparent, open-standards-based access to, and use of, those Information Services, with the IAAS middleware infrastructure automatically providing the navigation, federation or other underlying operations required to fulfil the service requests. This major advance is illustrated in Figure 12. This shows the types of information sources (*center*) that are usually virtualized by the IAAS Information Services virtualization layer. It also shows how this layer supports and provides information services to the new generation of SOA composite applications, and business process workflows (*upper part of chart*). The considerable benefits of this information virtualization are also enumerated and highlighted (*bottom of chart*).

Information as a Service and the z Series/System z9 Mainframe Platform

IAAS, as presented above, is an enterprise-wide vision that is logically independent of any underlying hardware platform, and intrinsically needs to embrace and support the usual diversity of platforms in use within the typical enterprise today. Our focus in this White Paper, however, is primarily upon the zSeries/System z9 that is the premier enterprise data-serving platform in the industry today. In the remainder of this White Paper we examine and evaluate how IAAS will be supported with recent and planned major advances in the System z9 hardware, operating system, database engines, and with new or enhanced IM middleware offerings for the mainframe platform.

Our Analysis

IBM's new IAAS strategy is unfolding with the emergence of a market-leading, comprehensive and highly integrated offering that will deliver on the firm's Information on Demand vision. It represents a significant industry breakthrough, and heralds a more comprehensive approach to unifying and exploiting information fully. It adds a new, enterprise-wide information management software services layer that offers powerful capabilities to transform enterprise data from many disparate sources into accurate/complete, timely, in-context, and effectively managed over its lifecycle, information for action.

The IBM Information Services and Information Accelerators that comprise the offering supply the core services that are needed to manage and access all forms of information under a SOA approach and a number of complementary accelerators that speed implementation. These fit tightly and naturally into the services-based approach to applications development of SOA that has itself rapidly become the dominant applications development model for the next decade, but can also be used independently and in advance of SOA adoption.

Based on open industry standards, the promise of IAAS is to render enterprise information truly open (*via these industry standards*) and more easily accessible (*in minutes not months*) to people, applications, and business processes across the organization. The prime areas where the IAAS strategy is expected to deliver major benefits are broad. Making money by better understanding customers, interacting with data in new ways, handling mergers and acquisitions, ensuring regulatory compliance, keeping data secure, supporting the explosion of mobile devices, and exploiting the latest information management/storage technologies in new industry-specific solutions are amongst these.

IAAS, with its familiar and new types of Information Services, represents a major advance over the current situation in most enterprises. It offers new methods of managing, understanding and exploiting the daunting complexity and diversity of information resources that is often the reality today. It also unifies and provides a common approach that embraces and handles both structured and unstructured data under one single, cohesive framework.

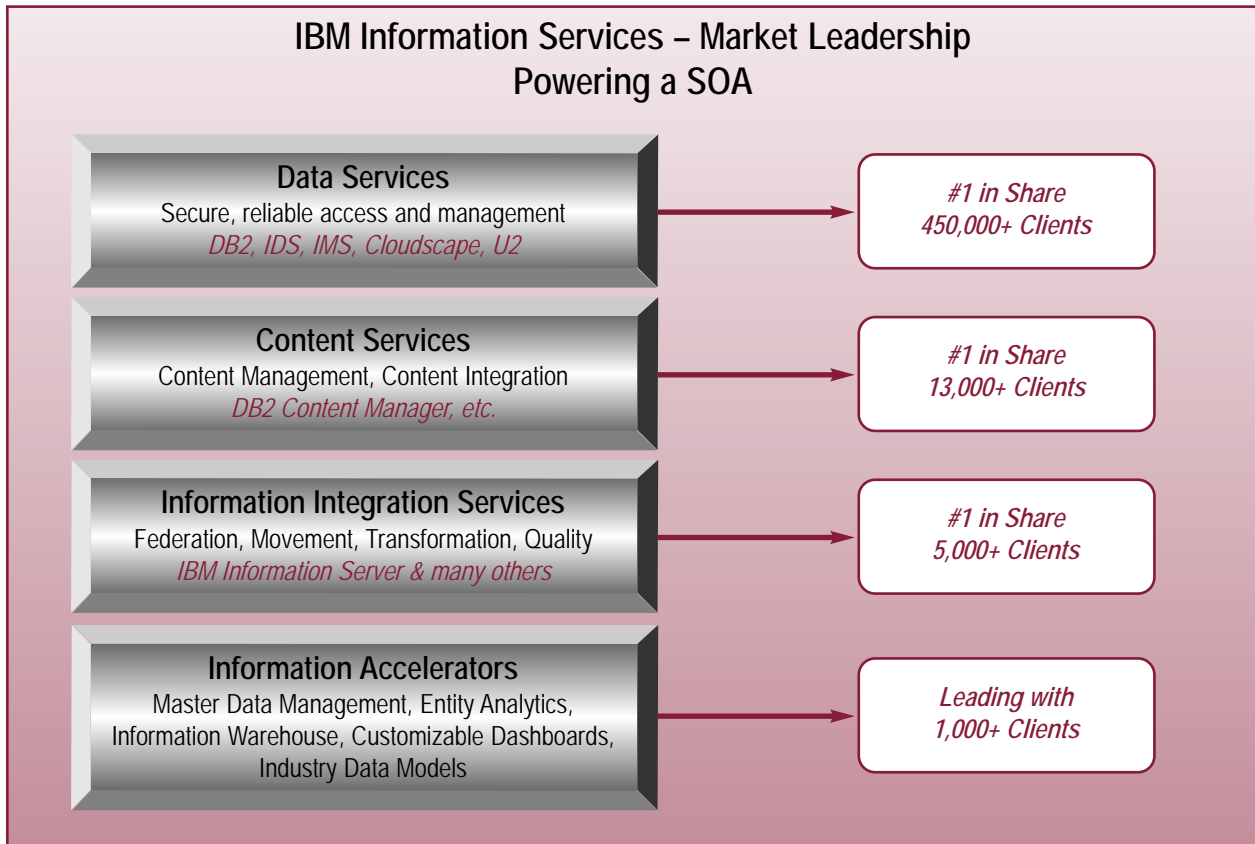


Figure 13: IBM Information Services – Market Leadership – Powering a SOA

IBM has already established a commanding market leadership position in most information management marketplaces, as shown by Figure 13. In the largest three main segments, the company already holds the market share leadership position, with massive customer bases using its information services software and solutions. In the much newer, but rapidly emerging MDM services segment, IBM has also made a formidable start and now has a major presence that we assess in Section 8.

In Section 5, we review in depth the data-serving capabilities of the System z9 mainframe, notably the powerful DB2 on z/OS combination, and assess the recently announced, extended vision for the platform to play a much larger role as the enterprise data server hub.

The unveiling of this compelling IAAS vision, and of the new and extended Information Services products and Accelerators that support it, is the outcome of one of IBM's largest and most significant software investment and development programs. This has run over the last three years, cost \$3B, and involved 15 software acquisitions. It has also deployed numerous sophisticated advances, algorithms and technologies pioneered by IBM Research, which are now included in the new offerings. Whilst current products are already a major landmark on this extensive roadmap, the effort is ongoing, and further large steps have already been announced for 2006 delivery.

This White Paper focuses on the System z9 mainframe. Our coverage thus centers on IAAS capabilities that are, or will be, available on the mainframe (*or that directly support it*), whilst also briefly covering some other important IAAS offerings currently offered on other platforms only to provide a more complete picture.

5. System z9 – Premier Enterprise Data-serving Platform – Major Advances

Introduction

On January 26th 2006, IBM launched a new vision, strategy and development roadmap to extend the already powerful strengths of System z9-z/OS-DB2 as the industry's "Gold Standard" enterprise data-serving platform, and to promote and grow its wider market adoption for the next decade. This announcement complements/supports the IBM IAAS strategy by strengthening this central enterprise data-serving hub.

This mainframe resurgence/market growth momentum was boosted by the Q4 2005 introduction of the next-generation System z9 109 high-end systems, which began shipping in mid-September 2005. For example, IBM's FY Q4 2005 financial results, released on January 17th 2006, stated that revenues from the zSeries/System z9 mainframe product increased 5% (*on an as reported basis*) or by 10% (*at constant currency*), compared with the year-ago period. Total delivery of zSeries computing power, which is measured in MIPS (*millions of instructions per second*), increased 28% in Q4 2005. Commenting on these strong zSeries and System z9 results, IBM CFO Mark Loughridge also reported:

"This was our largest quarter of MIPS shipments on record and our highest revenue since the fourth quarter of 1998. Today, over 60 percent of our (mainframe) revenue is driven by new workloads, such as Linux and Java, compared to only 15 percent at the end of 1998. This increased adoption of new workloads by our customers highlights the broader applicability of the platform and enhanced customer benefits of our latest System z9."

New-generation mainframe products typically ramp up to their peak adoption over several quarters, and we expect additional healthy mainframe revenue growth rates will therefore be posted for the Q1 and Q2 2006 periods. The fact that Q4 05 MIPS capacity shipment growth (*at 28%*) was 2.8 to 5.6 times higher than revenue growth (*5%/10% per basis used*), indicates the continuing sharp price/performance gains IBM delivered to mainframe customers in Q4 2005, further improving customer value.

IBM will build further on this market momentum with this new data-serving vision. It also announced the game-changing, complementary zIIP (*System z9 Integrated Information Processor*) specialty processor engine. This offloads selected System z9-z/OS-DB2 enterprise application, data warehousing and database maintenance workloads (*initially, others to follow*) to the dedicated, cost-effective zIIP specialty processors. We assess this significant development, which is to be delivered within 2006, in Section 6.

System z9 – Premier Enterprise Data-serving Platform

For over four decades, IBM mainframes, with their DB2 and IMS database software, have been the clear industry leader in high-end data and transaction serving, delivering levels of scale, performance, reliability, availability, security, manageability and recoverability unmatched by any other data-serving platform.

We found that the System z9-z/OS (*together with its SOA support from new WebSphere SOA Foundation middleware software – available on z/OS in Q1 2006*) now allows mainframe customers to unleash their accumulated mainframe data and applications equity by extending these to the new generation of SOA composite applications.

DB2 on z/OS allows users to combine and support many database-based solutions (*including transactional applications, data warehouses, BI, entity analytics, and support for ECM, etc.*) on this single, powerful, database-operating system combination, on the robust, reliable and scalable System z9 hardware platform. By combining many solutions in one, infrastructure can be simplified, and a single view of the data is easier to attain. The powerful partitioning, load balancing, and workload management of the System z9 easily enables it to simultaneously run many such DB2 data-serving workloads, each to assured QoS, on a single system.

Our research had isolated and ranked the nine foremost areas of unique business value that mainframe customers valued most highly about the platform. Several of these are widely understood across the industry (*Highest Resource Use Efficiency/Utilization, Highest Scalability & Capacity, Highest Performance & Quality of Service (QoS), and Low Risk via World-class IBM Support*). Others contradict widespread, but now obsolete, mainframe "myths" (*Lowest Total Cost of Ownership (TCO) & Cost/User, Much-improved System Costs*). These are highly desirable attributes for an enterprise data-serving platform destined to host much of the next generation of SOA-based enterprise applications. For example, to quantify just a few of these top business value attributes:

- The availability now obtainable from a single-site and multi-server System z9's exceeds 99.999% at the hardware level.
- The scalability of the platform as a whole now ranges from a single system of 54 processors SMP, 17,000 MIPS and 60 LPAR partition capacity, to a 32-System z9 Parallel Sysplex cluster of up to 32 times those figures.

- DB2 on z/OS manages many of the world's largest and most demanding databases, for example:
 - 23.1TB DB2 for z/OS database at the Land Registry for England and Wales, managing land registry records for every registered property and land tract for these countries.
 - The UPS shipping system that achieves peak workload rate of 1.1 Billion SQL statements per hour on DB2 for z/OS.

The trillion-dollar inventory of customers' application software, and the wealth of corporate data (*estimated to amount to 80% of machine-readable information*) accumulated over the last four decades on the mainframe platform, are also both major assets. These are now being widely mined and turned into reusable SOA software and information services. These in turn are being reused in the new SOA enterprise applications that customers are building. With so many such core business application and data assets residing on the mainframe, building and deploying next-generation SOA applications reusing many of these software and data assets on the mainframe platform offers substantial time-to-value and performance advantages.

With extensive support for open-industry standards, plus the most advanced virtualization capabilities of any platform, the mainframe is now well placed and equipped to extend its enterprise data-serving footprint. Its top levels of security, that assures the highest standards of data and privacy protection, is another argument for such deployment, as are the advanced, intelligent recovery capabilities of the platform that minimize business downtime and avoid data loss.

Roadmap for Further System z9-z/OS-DB2 Data-serving Advances

With the January 26th announcement, IBM set out a medium-term vision and roadmap charting numerous developments designed to further extend the leadership data-serving capabilities of the System z9-z/OS-DB2 combination, and this roadmap is shown in Figure 14.

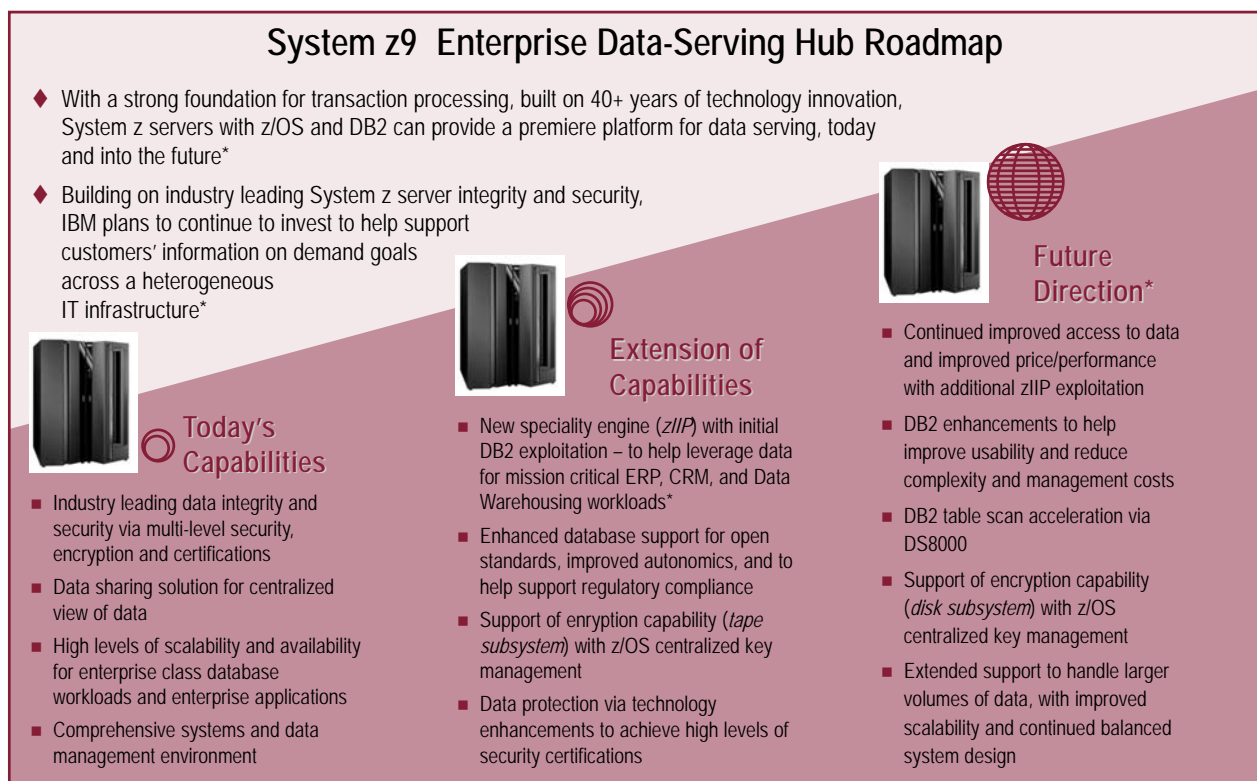


Figure 14: System z9 Data-serving Roadmap 2006-2007

The roadmap highlights (*upper part of chart*) the core strengths of the System z9-z/OS-DB2 platform. In the first column under the "Today's Capabilities" heading, it shows currently available elements which have themselves seen major advances in recent releases. The second column, under the "Extensions of Capabilities" heading, showcases the new zIIP speciality processor engine that supports the offloading of important qualifying DB2 workloads onto the zIIP engines. Also highlighted here are important new DB2 advances, and significant System z9 security advances (*including extended encryption for tape, high-level security certification for the platform, and a new DB2 encryption exit*) and we expect these to be delivered during 2006 via the deliverables reviewed below.

New z/OS improvements will further increase DB2 on z/OS performance and functionality. These will be delivered in the next major release of z/OS that we expect will be delivered around September 2006, likely to be named z/OS Version 1.8 (*on past z/OS release timing and naming practice*). Details of this release have not yet been disclosed, and are not expected much before Q2 2006.

Looking further ahead, under the "Future Directions" column of Figure 14, five further areas of advance, which we expect to arrive during 2007, are listed. These include:

- **Additional z9 offloads to zIIP engine:** This statement indicates other System z9 workloads are under consideration for offload to zIIP engines, which would further enhance the value-proposition of the zIIP solution. These are likely to be other information processing-intensive operations, but no specific workloads have yet been detailed, and these may be DB2 and/or non-DB2.
- **DB2 Enhancements:** New functionality for DB2 on z/OS will be delivered in the next-generation release of DB2 for z/OS. This new release (*code-named "Viper"*), is the first enterprise database management system to provide full native XML support, in a major database advance to a new hybrid relational/XML database model. It will also add numerous other DB2 for z/OS advances. This next-generation release has not yet been publicly named, and no z/OS version availability date has yet been declared. If IBM follows past DB2 release naming patterns, these would suggest a "DB Version 9" or similar naming. We use the form "DB2 Version 9?" here.
- **Exploitation of DB2-DS8000 synergy:** The advanced DS8000 enterprise storage systems IBM introduced in December 2004 have been successful in the market, and are the perfect complement to the System z9 mainframe and DB2. The statement here indicates DB2 table-scan acceleration, through exploitation/optimization of DS8000-DB2 synergy, will be made available. (*Our findings on these impressive, leadership enterprise storage systems are summarized in Appendix D, and our full research is highlighted in the "Related Software Strategies Research" Section on page 72, item 7*).
- **Continued DB2 security enhancements:** Support of encryption facility (*disc subsystem*) with System z9 centralized key management.

In addition, not shown on the chart, important new **Information Management software products for z/OS** are expected through 2006/2007 under the IAAS software strategy. These "missing pieces of the IAAS jigsaw puzzle" are known new IAAS products not yet currently announced for System z9-z/OS at time of writing. These could include:

- The new IAAS flagship **IBM Information Server** (*information integration services*) product. (*Assessed in Section 9*).
- The new **IBM Metadata Server** (*enterprise metadata services*) product. (*Assessed in Section 9*).
- The **IBM MDM products** (*WebSphere Customer Center and/or WebSphere Product Center for master data services*). (*Assessed in Section 8*).
- The sophisticated **DB2 Entity Analytics** products (*DB2 Identity Resolution, DB2 Relationship Resolution, and DB2 Anonymous Resolution*). (*Assessed in Section 11*).

...compelling roadmap of further major advances for the premier System z9-z/OS-DB2 data-serving platform...

We see compelling cases for several of these to be made available on IBM's flagship enterprise platform itself, and expect that one or several of these will be offered on System z9 by 2007.

Overall, this is a compelling roadmap of further major advances for the premier System z9-z/OS-DB2 data-serving platform

though 2007, particularly when the strengths of some of the individual items, reviewed below, are considered.

Strong System z9-z/OS-DB2 Data-serving Capability Advanced, Available Today

2004 and 2005 saw major new releases of the two mainframe database management systems, DB2 UDB for z/OS Version 8, and IBM Version 9, already widely adopted and well received by mainframe customers worldwide. These current versions provided extensive support for new-generation SOA applications and closer integration with the WebSphere middleware stack that supports SOA of System z9, and many other major advances.

Also introduced in September 2005 was the latest release of z/OS, Release/Version 1.7, which provided significant advances across the board. An important new System z9-z/OS Version 1.7 facility from a DB2 viewpoint was the MIDAW (*Modified Indirect Addressing Word*) capability, which considerably improved DB2 (*and other data-intense operations*) performance, and is reviewed below.

We provide our summary assessment of these current capabilities below:

- **DB2 Universal Database (UDB) for z/OS, Version 8**

DB2 for z/OS is IBM's flagship relational database management system, and is the foundation for its comprehensive set of data and information management services for the mainframe. (DB2 UDB is also supported on most other significant IT platforms). DB2 manages a high proportion of the world's largest databases on the zSeries/System z9 platform, and is equally suited to supporting large online transaction processing/databases, major enterprise applications packages (such as SAP and Siebel), and to handling enterprise BI and data warehousing applications.

DB2 is the global market share leader in the database market, with over one million server licenses deployed, and an estimated 40 million end users, according to IBM. DB2 Version 8, which became generally available on 03.26.2004, was rated the most significant advance of DB2 on the mainframe for many years. The middleware underwent a complete rewrite to extend scalability by fully exploiting the 64-bit z/Architecture of the mainframe, and to increase its commonality with the distributed platform implementations, now on a common code base. DB2-based information management solutions also include data warehousing, data mining, content management and information integration solutions, tightly coupled with this industrial-strength, multimedia, Web-enabled, high-performance, flexible and reliable/available database engine. These capabilities enable DB2 UDB Version 8 to support the most sophisticated and demanding enterprise-scale On Demand applications, and the core business processes of national and global enterprises.

DB2 UDB for z/OS Version 8, the twelfth version of the mainframe product, delivered more advances and improvements (100+) than any previous DB2 release. It substantially extended integration of DB2 with the hardware, other key IBM software elements, and with enterprise applications. Also delivered were **scalability advances**, to provide headroom for applications growth with flexibility, and **increased availability** by enabling more on-line changes to be made without requiring a shutdown in further moves towards continuous availability at the database level. DB2 UDB for z/OS Version 8 was the first IBM middleware product to break the 31-bit computing barrier and to fully exploit the massive scalability potential of the 64-bit mainframe z/Architecture, allowing customers to drive higher workloads through their systems and to "do more with less".

| Major z/OS | Key Feature/Function Advances in New Middleware Releases |
|---|--|
| DB2 UDB For z/OS Version 8 <i>Released 03.26.2004</i> | <ul style="list-style-type: none"> ● Enhanced SQL: More functions, better diagnostics, tighter DB2 UDB family compatibility, extending DB2's SQL beyond its current boundaries. ● Schema Evolution: "Alter a table and go" – no need to drop and redefine. This important change brings less system downtime and more data availability. ● 64-bit Virtual Storage Support: Simplifying main storage, increasing system availability and scalability. ● Longer Names for Tables & Columns, Longer SQL Statements: Enhancing DB2 family application portability and increasing database functionality. ● Enhanced Java™ and Unicode Support: Improving application support and reengineering for international business. ● Enhanced Utilities: Full utility support for the extensive changes in DB2 for z/OS, plus greater DB2 family compatibility. ● Multilevel Security at Row Level: Enables finer-grain database security at row level. ● DB2 Tools: Extensive suite of 32 complementary DB2 tools, all aligned for DB2 V8. |
| IMS Version 9 <i>Available since October 2004</i> | <p>IMS Version 9 Database Manager:</p> <ul style="list-style-type: none"> ● Broadened Access: Java™ & XML enhancements plus XML data storage in IMS databases. ● Integrated IMS Connect Function: Provides fast-to-install, easy-to-use, high-performance/high-volume, and secure transparent access to IMS apps./data from any app. environment, including Linux. Offers network management, workload-balancing commands for better resource utilization. ● Ease-of-use/Extended Access: Expanded, autonomic, user-friendly commands and interfaces that are accessible across environments. ● Easier to Handle & More Secure: Easier installation/system generation enhanced security and serviceability. ● Improved Availability and Recovery: Fully integrated on-line reorganization for HALDB, providing concurrent on-line updates and availability of data. ● Improved Performance and Capacity: For Virtual Storage Constraint Relief (VSCR), Database Recovery Control (DBRC), and Fast Path. <p>IMS Version 9 Transaction Manager:</p> <ul style="list-style-type: none"> ● Integrated IMS Connect Function: See above. ● Enhanced AD & Deployment: Broadened AD and execution tools, with XML and Java enhancements. ● Extended Autonomic Network Switchover Capability. ● Easier to Handle & More Secure: Easier installation and system generation, as well as enhanced security and serviceability. ● Improved Performance and Capacity: Improved system availability, performance, and capacity for VSCR, DBRC, and Fast Path message handling. <p>IMS Tools: Enhanced, extensive suite of 36 IMS tools, all aligned for IMS V9.</p> |

Figure 15: System z9 Database Servers – Strong Current Release Advances

DB2 UDB for z/OS Version 8 also provided comprehensive database support for building and deploying new **composite applications under SOAs**, both through closer integration with the WebSphere software platform, and through DB2 support for XML and Web services open standards.

In addition to these overall Version 8 improvements that broadly enhanced scalability, application porting, security, architecture, and continuous availability, our selection of specific new DB2 UDB for z/OS Version 8 features and advances are highlighted in Figure 15 (*on page 27*).

Also offered is a **comprehensive portfolio of modern DB2 tools** on z/OS that support high-productivity database administration, recovery and performance management, application management, and replication and utilities management. These tools were all aligned to support the DB2 UDB for z/OS Version 8 release, and provided support for building, optimizing and managing the whole DB2 infrastructure on zSeries cost-effectively. This extensive DB2 portfolio for z/OS adds hugely to the strength of the complete DB2 offering on the platform, as can be seen from Figure 16, which shows z/OS DB2 tools by category now available.



Figure 16: Extensive DB2 Tools Portfolio for z/OS

The DB2 UDB for z/OS Version 8 release's **performance enhancements** drove down the TCO of DB2 applications, making each transaction cheaper to execute, and allowing a given server to handle larger workloads. Many system **limitations within the product were removed** or extended, enabling much larger applications. For example, table name sizes were increased from eighteen to 128 characters, and the number of database partitions in a database was increased from 254 to 4,096. By removing barriers that prevented non-DB2 applications from being easily ported to DB2 on the zSeries/System z9 platform, Version 8 also facilitated server consolidation, and new "to the mainframe" database migration, projects.

In addition, the **increased automation and intelligence** built into the product and the supporting DB2 tools significantly increased productivity, allowing skilled DB2 administrators and developers to spend less time on routine tasks and to devote more of their efforts to adding value with better-written, more efficient applications.

Migration must be from DB2 UDB for z/OS Version 7. z/OS Version 1.3 executing in 64-bit addressing mode is the minimum prerequisite operating system level.

- **IMS Database and Transaction Servers, Version 9**

Although our main focus in this Section is on DB2, IMS remains a bedrock database engine on the System z9 mainframe platform, and also saw recent major advances. IMS (*Information Management System*) is z/Series' long-standing (37-years) and respected high-performance, high-volume database and transaction server, used for the most demanding, highest volume OLTP applications. IMS Version 9, the latest major release, became generally available on 10.29.2004. Version 9 was developed to deliver the highest QoS for enterprise transactions and data, and to extend these qualities to new composite applications under a SOA. IMS, almost alone, can reliably support tens of thousands of users with sub-second response times, and can process more than 100M transactions per day. Peak transaction rates exceeding 21,000 transactions per second have been posted on a single mainframe IMS system. It can also access multi-terabyte-sized databases, and run near-continuously with as little as 2-3 hours of planned and unplanned outage per year.

Thus it is no surprise that IMS usage has again been enjoying robust usage growth in recent years. The largest financial, transportation, utility and manufacturing firms worldwide, who have long used IMS, are experiencing e-business-driven growth. New customers, such as the banking sector in China, utilize this database server engine to support their surging business transaction growth. IBM reports that 95% of Fortune 1000 companies use IMS, that some 50B transactions per day are processed on the software, serving 200M end users, and supporting in excess of 15B GB of production data.

IMS Version 9 provided a database and transaction environment able to manage large and unpredictably growing volumes of data. It took further important steps towards offering continuous availability, and had been optimized to deliver consistent, rapid response times for critical transactions with customers, suppliers and partners. A comprehensive range of IMS tools provided IMS database administration, performance management, recovery management, application management, transaction management, utilities management, and replication services for IMS system optimization, which increased staff productivity to support the core servers and their applications more easily and at lower cost. OMEGAMON XE for IMS, for example, provided advanced performance monitoring for IMS environments and applications, a fruit of the 2004 IBM acquisition of Candle Corporation. The tool suite products were all aligned to support IMS Version 9 from day one, enabling customers to immediately exploit the significant advances included. Figure 17 shows this wealth of IMS tools by category available in 2006.



Figure 17: Comprehensive IMS Tools Portfolio (z/OS)

Version 9 of IMS also enhanced high-performance, transparent access to IMS transactions and data from any other application environment, including Linux. It provided **native XML-based interchange** that enabled all parts of the enterprise, partners and customers to interact and exchange data with core IMS data servers. It also provided **Web services standards support**, so that traditional IMS applications could be published and choreographed into new composite applications under a SOA. Figure 15 highlights our selection of enhanced features incorporated in IMS Version 9.

Java has become central for new enterprise applications, and so IMS Version 9 and the IMS Connector for Java offered **enhanced Java support**. They allowed customers take advantage of the latest tooling, as well as providing enhanced performance. Many IMS customers are extending/modernizing their existing software assets into new On Demand business application. These advances enabled IMS to now support the vital applications, information and operational integration needed. They helped provide a seamless end-to-end capability for IMS application use and data access, enabling applications to talk to others in new ways to support changing core business processes. The enhanced development tools support offered also made this extension and reuse of core IMS application assets faster and easier, improving developer productivity and project time-to-business value.

Many other improvements made IMS itself easier to install, easier to use and manage, extended environment management commands and tools, improved resource utilization and performance, as well as increasing operations productivity.

IMS Version 9 was fully enabled for use with IBM WebSphere Application Server, and with the WebSphere development tools, to rapidly transform static Web sites into sources of dynamic Web content. Version 9 allowed customers to transform IMS transactions into Web services that can be reused in composite applications for SOAs. In our assessment, these advances presented a convincing case for IMS users to migrate to IMS Version 9, and adoption has been brisk for IMS since the release became available.

- ***z/OS Flagship Version 1.7 – A Major Release/Advance***

The System z9 109's most advanced features are optimized under the flagship z/OS operating system that most fully exploits z/Architecture's 64-bit capabilities. z/OS provides a secure, scalable, high-performance base on which to deploy Internet and Java technology-enabled applications, providing a comprehensive and diverse application execution environment that meets the most demanding QoS requirements for enterprise transactions and data for large-scale, collaborative computing. Particular strengths of z/OS include autonomic capability, integration, virtualization, and openness. The newest production release, z/OS Version 1.7, which shipped in September 2005, enabled the transport of data in a security-rich environment between multiple computing platforms, and the virtualization of hundreds of applications, including Java-based applications.

z/OS Version 1.7 added facilities that extended the capabilities of the zSeries/System z9 109 platform to be able to serve as the **enterprise data hub and principal data server**, and as the **interconnect point for business logic under SOA**. These extensions to the role of the mainframe apply its core strengths to managing and serving enterprise data, and to the deployment, execution, management and control of business logic, not only on the zSeries/System z9 itself, but also for the **other connected hardware platforms** in the enterprise IT infrastructure.

Other business value enhancements in z/OS V1.7 included further scalability improvements, application integration enhancements, further security improvements, enhanced TCP/IP networking and applications availability, additional self-optimization features, and enhanced ease-of-use. Also included were enhancements to deliver increased availability of z/OS UNIX System Services (*z/OS UNIX*), to support new security standards, and to improve enterprise-wide workload management. Ease of-use was enhanced by the integration of the Health Checker within z/OS that provides checking and "best practices" for z/OS components, and which can be used by both ISV and middleware products. In addition, z/OS 1.7 supported new hardware features and functionalities of the new System z9 109 mainframe generation.

With z/OS V1.7, IBM also simplified z/OS to make the mainframe platform easier to manage by a new generation of IT professionals. In addition to continuing to extend its robust capabilities, IBM worked to automate, simplify and eliminate many z/OS management and maintenance tasks to free mainframe staff for other activities. A new graphical z/OS management console has subsequently been made available in support of these goals. Other new features in z/OS V1.7 focused on simplifying transaction encryption, balancing workloads across a cluster, and continuing to reduce both planned and unplanned outages.

z/OS now optimizes the performance of the IP networks across a cluster of servers with a new TCP/IP Sysplex Load Balancing Advisor that coordinates with switches, such as those from Cisco. This enables better load balancing decisions that protect busy servers and provided resources where the business needs them.

We assessed z/OS Version 1.7 (*and the previous V1.6*) in considerable detail in a 2005 White Paper, where we found it offered major advances across the spectrum, and the summary above was drawn from this research. (*See the "Related Software Strategies Research" Section on page 72, item 4.*)

- **DB2 Content Manager for z/OS Version 8.3**

ECM is one of the most important enterprise information services required in an IAAS environment. For the System z9 mainframe, DB2 Content Manager for z/OS, in its latest Version 8.3, provides a comprehensive ECM and content integration service solution. Since this is such an important part of the IAAS environment for the System z9 mainframe, we cover DB2 Content Manager for z/OS Version 8.3 in more depth in Section 7.

- **20% I/O Response Time reduction for DB2 using MIDAWs**

New with the System z9-109, the MIDAW (*Modified Indirect Access Word*) facility provides a more efficient replacement for certain categories of data-chained CCW I/O operations, particularly beneficial to database processing operations. The MIDAW facility improves FICON performance by reducing channel, director, and control unit overhead. MIDAW, which works on a single I/O block, makes it possible to efficiently send small control blocks embedded in a disk record to separate buffers than those used for larger data areas within the record. IBM reports substantial performance gains in a number of practically important DB2 tasks with MIDAW. VSAM, PDSE, HFS, zFS and other extended format datasets will also see performance benefit from MIDAW in many cases. For example, 20% I/O response time reductions have been measured in many DB2 applications using MIDAWs.

DB2 for z/OS Version 9? – The Next-generation DB2 for Mainframe Release

Preliminary information on the next generation of DB2 was disclosed in November 2005 at the XML Conference, at which the start of open testing/evaluation of the product by qualified customers, developers and partners was also announced. The new product is expected to become available for System z9-z/OS around the end of 2006. Code-named DB2 “Viper”, this next-generation release has been designed to help customers manage and access data across a SOA with higher levels of flexibility and speed.

DB2 “Viper” is the industry’s first database designed with both **native XML data management** and relational data capability. DB2 “Viper” is expected to be the only database product able to seamlessly manage both conventional relational data and XML data without requiring the XML data to be reformatted or placed into a large object within the database. This breakthrough will enable customers to increase the availability, speed and versatility of their information, while dramatically reducing administrative costs associated with existing data management techniques. It also will significantly reduce the complexity and time and effort developers will need to spend creating applications able to access both relational data and XML repositories.

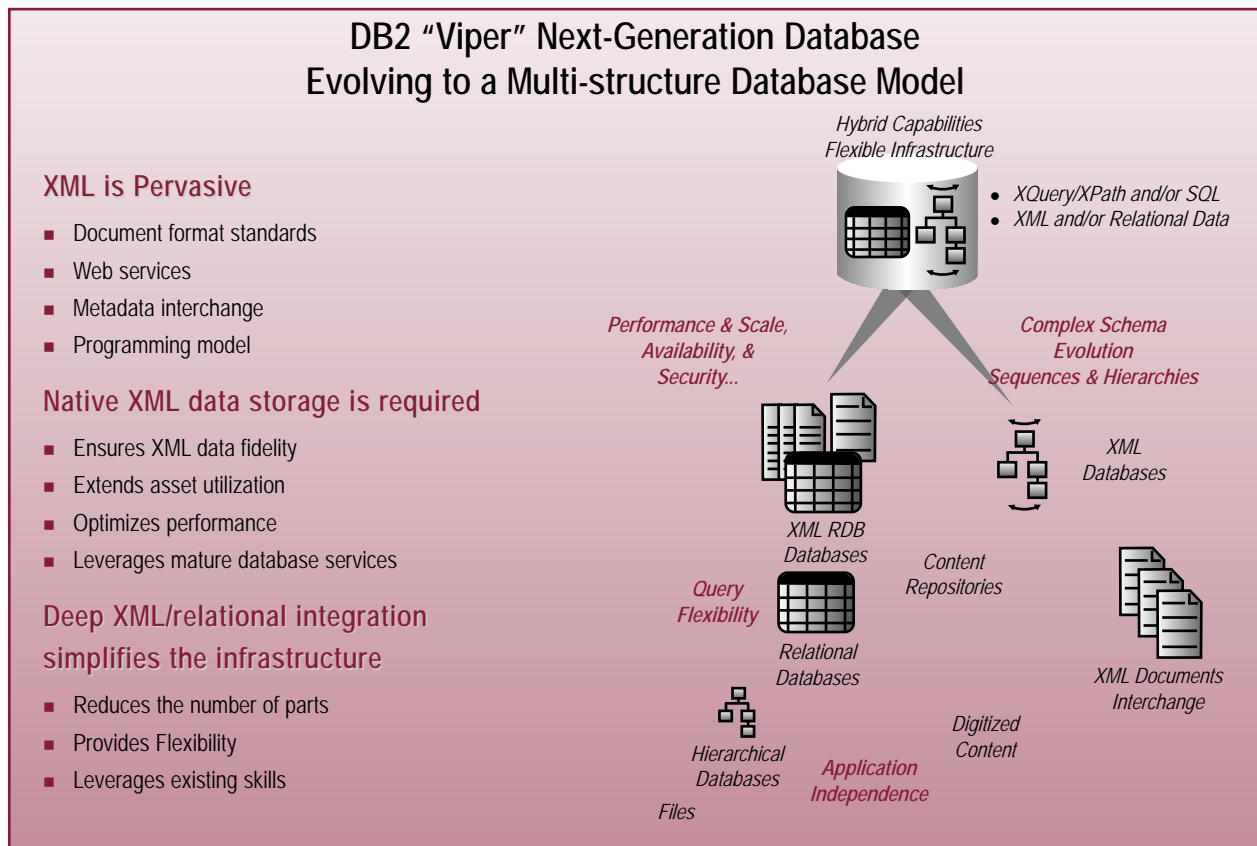


Figure 18: DB2 “Viper” Next-Generation Database – Evolving to a Multi-structure Database Model

The right-hand graphic part of Figure 18 above highlights the new DB2 "Viper" multi-structure database model that combines hierarchical management of XML content (*accessed with XQuery/XPath*) with traditional relational database management (*accessed with SQL*) within one unified, high-performance, highly-available database infrastructure. The bullet points on the left-hand side of Figure 18 highlight the pervasive role XML now plays in enterprise SOA architectures. They make the case that native XML database support is now definitively required for data fidelity, asset utilization, performance optimization, and to exploit the mature data management services found in an industrial-strength database management system like DB2. This new breed of hybrid, multi-structure database management system will simplify enterprise infrastructures, increase flexibility, and take advantage of the widespread existing DB2 skill-base in the market.

In his keynote address at the XML 2005 Conference (*where the product was first previewed*), Bob Picciano, IBM VP of Database Servers, commented:

"DB2 Viper will fundamentally change the rules of the database game. With Viper at the heart of their information infrastructure, customers will rapidly transition from conventional data management practices to unprecedented information management techniques that enable them to leverage information as a service."

These DB2 Viper native XML capabilities will benefit SOA implementations, which rely upon the ability to access a myriad of data stored across multiple formats. They free data from the static form it has been forced into by relational-only database products, enabling DB2 Viper to deliver IAAS that is readily accessible throughout SOA environments.

Currently, only two alternatives exist. One is to run a dedicated XML database (*capable of managing hierarchical XML data structures*) alongside the main RDBMS, which complicates the infrastructure. The other is to force XML data unnaturally into a relational database with XML "shred" technologies, or to hold XML documents in the RDBMS as large objects. The former loses much of the value and structure of XML message information, and the latter does not allow message content processing, so neither are real, long-term solutions. Highlighting the major advantages native XML database support provides, early DB2 Viper user Thore Thomassen, Senior Enterprise Architect for leading Norwegian insurance and financial services company Storebrand, commented in the announcement:

"Our development time using Viper's native XML store is a radical improvement over existing XML 'shred' technology. We are now able to make schema changes in minutes rather than days and this will dramatically improve our customer response time."

Viper's native XML technology will also provide XQuery support. XQuery is an emerging industry-standard language that extends XPath and is specially designed for processing XML data. Applications can use XQuery, standard SQL or both to retrieve documents from either, or both, underlying storage formats.

In addition to this groundbreaking native XML capability, DB2 "Viper" is also the first database that can support all three common methods of database partitioning simultaneously – a major innovation in improving data management and information availability. By simultaneously supporting range partitioning, multi-dimensional clustering and hashing, Viper will enable organizations to arrange and order their information in the way that best suits their individual business requirements and demands.

It will also support handling increasing regulatory compliance, adding enhanced auditing and privacy capabilities, support for trusted security contexts, and for defined database roles in the compliance environment. To reduce risks, the new release will also support native SQL stored procedures, will provide default databases and tablespaces, and will provide automatic unique indexes to support and define primary keys. To decrease complexity and reduce database administration costs, it will also provide fast table replacement, the ability to partition a database by growth, index compression, and a new Optimization Service Center.

We assess DB2 "Viper"/DB2 UDB for z/OS Version 9? To be a substantial industry advance. There are real advantages in its new hybrid relational/XML multi-structure database model. It allows developers to store and handle both relational tabular data and XML data within a single data management infrastructure much more easily, with equivalent access via SQL and XQuery, without any loss of the richness of XML content, and with strong performance. This major advance places IBM far ahead of its RDBMS competitors Oracle and Sybase who offer limited XML support, and still further ahead of Microsoft's SQL Server. IBM executives say that around 20% of enterprise data today resides in relational databases, but 35% is already in XML, so extending the database to be able to natively manage this XML data provides a huge potential expansion of database usage.

With the rapid, industry-wide adoption of SOA applications architectures, to which XML is intrinsic and central, DB2 "Viper" now stands alone with this breakthrough native XML support amongst the major RDBMS. It also advances DB2 "Viper" well beyond the existing definition and scope of a database, creating the new and broader wider category of an enterprise data server.

Our Assessment

The zSeries/System z9-z/OS-DB2 enterprise data-serving platform has a long-established record of accomplishment in handling many of the largest customer databases and applications in the industry. The significant hardware and system advances of the new System z9 mainframe generation now provide large extensions to the scalability, availability and security of the environment that now allows larger enterprise-wide and extra-enterprise database applications to be cost-effectively, securely and efficiently hosted on the z9 platform.

The zSeries/System z9-z/OS-DB2 enterprise data-serving platform has a long-established record of accomplishment...

It therefore makes perfect sense for IBM to extend and promote further and wider usage of the System z9-z/OS-DB2 platform for broader enterprise data-serving roles at the heart of the new enterprise infrastructures supporting SOA. We warmly welcome this announcement and the publication of this System z9-z/OS-DB2 data-serving roadmap.

It also supports the IBM IAAS strategy/vision by providing the high-end data/information services at enterprise hub level that are needed for next-generation, large-scale, collaborative SOA applications.

The new zIIP specialty engine is a welcome addition that provides encouragement and economic advantages for deploying additional remote application DB2 data-serving and data warehousing workloads on the platform, and we evaluate this fully in Section 6. By providing a significantly lower cost for these popular mainframe DB2 workloads, IBM has also dealt a powerful blow to its main high-end enterprise database competitor Oracle, whose Oracle RAC cluster solutions claim to deliver high-end database functionality and capacity using commodity server clusters. Whilst few, if any, Oracle RAC databases approach the size of the larger mainframe DB2 site databases, Oracle has typically made much market noise about their RAC offering as a mainframe alternative. With the advent of the zIIP, it is likely that any prior Oracle RAC cost advantage will be eliminated or much reduced, and the case for the more manageable, automated, reliable DB2-z/OS combination will become more unassailable. We must await details of the typical zIIP savings attainable to fully quantify this advance.

2004 and 2005 saw big advances in the main zSeries/System z9 data management and operating systems environments, with DB2 UDB for z/OS Version 8, IMS Version 9, and z/OS Version 1.7. These strengthened the mainframe data-serving platform with numerous advances, and all are available today.

Other related, parallel IAAS advances with current or future System z9 application are covered in Sections 7-11 and add further software weight to the case.

The most important additional advance for the System z9-z/OS-DB2 data-serving platform will be the availability of DB2 UDB for z/OS Version 9? (*DB2 "Viper"*), which brings the breakthrough and industry-leading native support for XML in a new hybrid relational-XML multi-structure database model that offers real and large advantages over both current alternative approaches, and over IBM's database competitors. This will be further supported and enhanced by the new 2006 release of the z/OS operating system, z/OS Version 1.8? (*Features of which have yet to be announced*) but which can be expected to focus on fully exploiting newer System z9 hardware capabilities, and on supporting the System z9 data-serving role.

Further ahead, the roadmap for 2007 highlights a number of important and intriguing further advances in DB2 security, hardware exploitation, additional zIIP workload offloads, and the debut of additional major IAAS software products on the z/OS platform.

We consider this a compelling, well-conceived roadmap that will give enterprise users considerable confidence and motivation to further extend their use of the System z9-z/OS-DB2 data-serving platform...

IBM has also indicated that it will, in the future, offer new DB2 functions that will support DB2 scan acceleration via the DS8000 enterprise storage system, exploiting closer System z9-DS8000 integration synergies, and offloading part of this host workload to these powerful storage systems. (*See Appendix D*)

We consider this a compelling, well-conceived roadmap that will give enterprise users considerable confidence and motivation to further extend their use of the System z9-z/OS-DB2 data-serving platform to support more of their enterprise data-serving workloads through 2006 and 2007.

6. System z9 zIIP Specialty Engine – Transforms z/OS DB2 Data-serving Economics

zIIP Introduction

For an IAAS IM strategy for SOA, data serving is the core foundation upon which all other enterprise information services are based, and is thus of central importance to all enterprise IT users.

On January 26th 2006, in a major strategic move designed to encourage, facilitate and support the extended enterprise data-serving role proposed for the System z9 mainframe, IBM announced an important new weapon.

The new System z9 Integrated Information Processor, snappily named the zIIP, is a new System z9 dedicated specialty processor. The zIIP automatically offloads remote SQL requests for DB2 database-serving submitted over TCP/IP, DB2 star schema data warehouse parallel query-processing workloads, and some DB2 tools index maintenance operations, from the System z9's CP general-purpose processors (*where these workloads have run to date*). It then automatically runs them on zIIP engine(s) designated on the System z9.

The zIIP will serve as a high-speed engine that better enables data to be centralized on the mainframe. The ease of implementation and economy of the zIIP solution helps eliminate the barriers between transactional data stores on the mainframe and BI, ERP and CRM applications that run on distributed computers. With the zIIP capability, the System z9 mainframe will minimize the need to maintain duplicate copies of the data and will provide better security between the applications and the data.

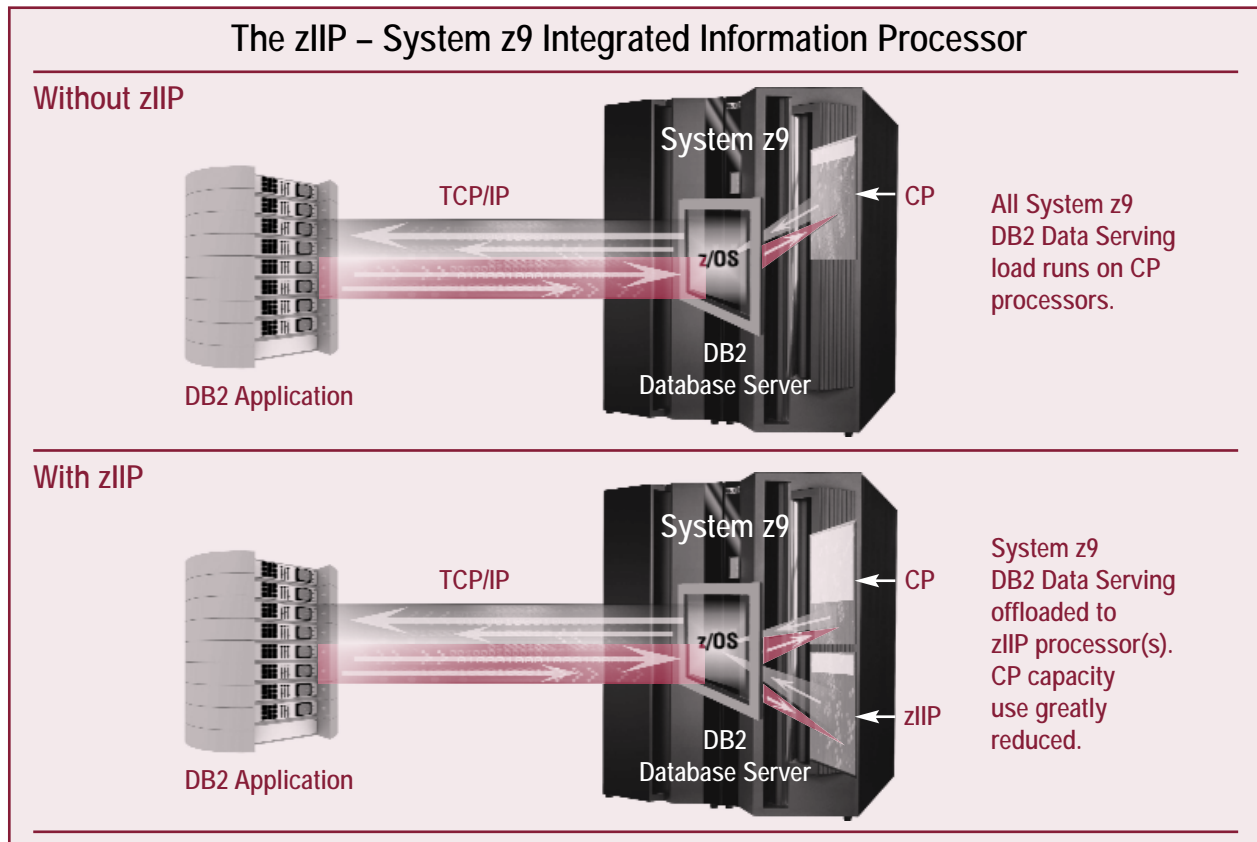


Figure 19: The zIIP – System z9 Integrated Information Processor

The operation of the zIIP for the first of these workloads is shown in Figure 19, which illustrates incoming remote SQL requests over TCP/IP being served by DB2 running on z/OS on a System z9 mainframe, with and without the zIIP. Without the zIIP, the System z9's general-purpose CP processors process all the data-serving work. With the zIIP, much of this work is offloaded to the zIIP specialty processors, reducing the load on the general-purpose CPs, which are thus freed up for other workloads.

Removing Barriers to System z9 DB2 Data-serving Workload Expansion

DB2 running under z/OS on the zSeries and System z9 mainframe platform has long been the industry's premier enterprise data-serving environment. It offers rock-solid, high-performance, secure, large-enterprise-scale relational database data management services for traditional mission-critical transactional workloads, for large-scale BI/data warehousing workloads, and increasingly for next-generation, SOA-based composite applications. DB2 on z/OS has long provided these data management services not only to the myriad applications that themselves run on the zSeries/System z9 platform, but also to the numerous requests which come into the mainframe database from applications hosted on other distributed platforms.

With its unique QoS, the zSeries/System z9-z/OS-DB2 platform is without real rivals in its globally proven ability to manage the largest and most demanding data-serving workloads in the industry, with outstanding manageability, recoverability, reliability and performance.

Indeed, the recent years of zSeries/System z9 mainframe market resurgence was driven by customers placing substantial new workloads on the platform, most of them exploiting the platform's data-serving powerhouse strengths. These include new SOA and J2EE database applications, ecosystem collaboration applications, and major enterprise packaged applications such as SAP and Siebel, which all require the powerful back-end data-serving support that DB2 on z/OS provides. Continuing reductions in the cost of zSeries mainframe hardware and software, seen consistently over this decade, have also facilitated and encouraged this clear trend.

Only one constraint limited wider adoption of this successful data-serving role for the mainframe. Whenever such a workload drove up System z9 general-purpose CP usage for data-serving that required a capacity increase, it triggered increased System z9 software costs for the LPAR partition concerned. The advent of Workload License Charge (*WLC*) pricing (*earlier this decade*) had improved this issue. WLC limited software license cost changes just to the CP capacity increase powering the specific partition concerned, rather than across the zSeries/System z9 system as a whole. None-the-less, this factor was undoubtedly the constraint limiting wider use of the System z9-DB2 data-serving combination. Enterprises using DB2 on z/OS have long been familiar with the strengths and advantages of the combination, and have established skills and experience in managing them. Many had asked IBM to make it easier and more affordable for them to extend the data-serving role of the combination to much wider workloads.

Now, in a dramatic, game-changing strategic move, IBM has taken up its sledgehammer and smashed this last obstacle to wider System z9 DB2 data-serving deployment with the zIIP.

The zIIP in Detail

The zIIP is a dedicated specialty System z9 processor specifically designed to help improve resource optimization and help lower the cost of ownership of the System z9 for eligible DB2 data-serving workloads. With the introduction of the zIIP, System z9 z/OS users can free up the corresponding capacity and optimize their resource utilization on their general-purpose System z9 CP processors, by having z/OS automatically direct eligible work to these new specialty engines. The zIIP will offload three important z/OS DB2 database serving workloads initially:

- **Host DB2 data-serving workloads generated by remote SQL calls**, submitted via over a TCP/IP connection, from existing or new single or multi-tiered DB2 applications on other platforms, that access a DB2 database residing on a System z9 mainframe. This will benefit a wide range of multi-tier ERP, CRM, BI and SOA application serving workloads that use a System z9 DB2 data server, and will encourage further and wider use of these configurations. The popular and widely used DB2 Connect software is one major example of IBM software that uses DRDA over TCP/IP and that would thus generate qualifying zIIP workloads of this type on the DB2 host system.
- **Star schema parallel query workloads** on a System z9 DB2 database, whether the applications submitting the query are resident on the System z9 itself, or are running on another platform and remotely submitting the queries to the System z9 DB2 database server. This creates new opportunities and encouragement to deploy more data warehousing and BI solutions that exploit star schema queries on System z9 DB2 database servers cost-effectively.
- **Select DB2 tools database maintenance operations**, internal DB2 utility functions used to maintain index maintenance structures. These operations include the DB2 Load, Reorg. and Rebuild Index utility functions that are often substantial workloads for larger or more complex DB2 databases. This is highly complementary to the other two zIIP workloads above, because these database maintenance operations are usually run overnight or at weekends, outside of prime online service hours. They will therefore exploit installed zIIP capacity when demand from the other two online workloads are likely to be at their lowest.

These DB2 workloads are those that use Service Request Block (SRB) mode for some or all of their processing, part of which can be automatically offloaded by z/OS to the zIIP. As with the other System z9 specialty engines, IBM anticipates no changes at all will be required to the DB2 applications concerned in order to take advantage of the zIIP capability, and workload diversion will be made completely automatically by the z/OS operating system. The zIIP benefits, where two or more qualifying types of workload co-exist in a single application, are additive. At the January 26, 2006 announcement, IBM also gave very clear hints that additional z/OS workloads will be enabled to exploit zIIP capacity in future.

As with the other existing specialty processors such as the zAPP (*Java workloads*), the zIIP specialty processors must be installed 1:1 with the number of general-purpose CP processors designated/installed. (*The 1:1 ratio applies to all types of specialty processors installed, not to a single type.*)

Adding zIIPs does not increase the MSU rating or model designation of a System z9 system, and so does not increase operating or middleware software license costs. So no additional IBM software charges will apply for the zIIP capacity, unlike for general CP processors, and this provides the new levels of price/performance the zIIP will offer for these DB2 data-serving workloads. Each zIIP engine on the System z9 provides 600 MIPS of capacity to execute the qualified offloaded DB2 data-serving workloads.

The full set of System z9 Capacity Upgrade on Demand options will also be offered for the zIIP, allowing customers to change their zIIP capacity permanently or temporarily as needed.

Adopting the zIIP requires a System z9 mainframe system, the z/OS 1.6 or z/OS 1.7 operating system releases, and DB2 Version 8 (*with enabling PTFs*), or above. A detailed announcement of the zIIP, with full pricing and availability details, will be made early in Q2, and actual availability is planned "within 2006". Substantial guidance on the proportion of work that customers can expect to see offloaded to the platform will then be provided. We expect IBM to try hard to deliver the zIIP for the Q4 prime mainframe sales quarter. However, a considerable amount of development, testing and measurement is needed to bring a new System z9 specialty engine to market, as these do take time to perfect and test thoroughly.

The other available System z9 specialty workload processors, the zAAP (*Java workloads*) and the IFL (*Linux workloads*) are each currently priced at a \$125,000 per-processor purchase cost. So, for example, over a typical seven-year system life-span, the initial processor hardware cost of one zAAP or IFL would average a fixed cost of just \$17,857 per year for the high Java™ or Linux throughput each specialty engine can deliver. On public announcement, IBM also set the pricing for the zIIP at the same fixed \$125,000 price at these other specialty engines. With the likely offload savings attainable, this will radically improve the economics for the mainframe deployment of these DB2 workloads.

Customers may concurrently use zIIP and zAAP engines to offload the qualifying DB2 workload and Java workload parts of the same enterprise application for additive benefits and savings.

As with the other System z9 specialty engines, all the zIIP engines installed in a System z9 may be managed overall as a single resource pool, providing flexibility, support and back up for many qualifying applications.

The zIIP will also be able to use the new weight management (*previously announced for general-purpose processors (CP) and zAAPs on the z9-109*). This allows separate processing weight specification for zAAPs in an LPAR, to allow flexible changes to the independent weight values within the zAAPs respective pools. The zIIP will also be able to take advantage of such weighting factors.

A "Playbook" offering, to provide a "zIIP equivalent" via through commercial T&C mechanisms, is planned for prospective customers with qualifying new DB2 applications who need to go live prior to actual availability of the zIIP facility, and/or of its enabling prerequisites at the customer site.

How Much DB2 Data-serving Workload can be Offloaded to the zIIP?

An obvious question for potential zIIP customers running, or planning to run, qualifying DB2 data serving workloads on their System z9(s) is what proportion of each such workload will be automatically offloaded to the installed zIIP(s) under consideration? IBM is currently conducting extensive tests and analyses to answer this question, will publish estimating guidelines and examples at the time of the early Q2 06 full announcement, and expects to provide more detailed planning and implementation materials by zIIP availability. However, preliminary studies of typical enterprise application data-serving workloads, such as SAP applications, have indicated a broad potential for an average 50% workload offload for this type of application.

Users will clearly have to wait for these guidelines and test results to determine the overall solution cost reductions the use of the zIIP will bring to these qualifying workloads, but these are certain to be substantial on all indicators seen to date.

Other Advantages of the zIIP

The advantages conferred by System z9 zIIP use are not only these substantial, general-purpose CP capacity releases and associated substantial cost savings. By enabling more multi-tier enterprise applications to enjoy the rock-solid, reliable and manageable System z9 z/OS DB2 data-serving capability, they will improve the performance, stability and manageability of this important class of multi-tier enterprise applications.

By enabling cost-effective operation of more large-scale star schema data warehousing and BI query-intense applications on the System z9-z/OS-DB2 platform, users will reap substantial gains from the improved performance, easier single-platform ETL processes, and better manageability and QoS the mainframe provides at much lower costs than hitherto.

By offloading DB2 database maintenance operations to the zIIP(s), customers may well also be able to reduce or squeeze overnight or weekend windows, or overcome existing batch schedule constraints.

The zIIP will also enable companies to recentralize their DB2 data on the more manageable, reliable and recoverable mainframe platform, and also to run their data warehouse databases more closely alongside their transactional applications on this robust platform, at a considerably more affordable cost,

These are valuable additional benefits that flow from deployment of the zIIP.

Our Analysis

In our assessment, the introduction of the zIIP is a highly significant milestone that will further accelerate the continuing resurgence of the mainframe platform. It will make it easier, and considerably more cost effective, for customers to support many more of the qualifying DB2 data-serving workloads on the rock-solid, "gold-standard" enterprise System z9-z/OS-DB2 data-serving platform.

The same technique, that of offering fixed-price and software-cost, impact-free, dedicated processors for Linux and Java™ J2EE™ workloads, has been successful for both IBM and for its customers. These have encouraged/enabled the now widespread exploitation of Linux, and the fast-growing wider deployment of enterprise Java applications, on zSeries/System z9. They have swung the economic balance more in favor of deployment on secure, manageable and reliable zSeries/System z9 mainframes (*as opposed to on complex distributed server farms*).

...the zIIP is a highly significant milestone that will further accelerate the continuing resurgence of the mainframe platform.

The advent of the zIIP will, in our assessment, have a similar effect, by swinging the economic balance in favor of running the qualifying DB2 database serving workloads on System z9 z/OS DB2. The QoS, performance and ease of management advantages of the platform, which DB2 on z/OS fully exploits, are already favorable, and these, combined with the zIIP economic advantages, make a strong case.

In the roadmap IBM announced plans for further extending/enhancing the data-serving strengths of the System z9-z/OS-DB2 platform (*discussed in Section 5*), it also indicated that other IBM subsystem workloads are under consideration for offload to the zIIP (*and or zAPP*) specialty processors in future. These moves, when realized, will further strengthen the appeal and value of the zIIP solution.

The main limitation of the initial zIIP offering is that it does not support offload of local, System z9-resident DB2 application data-serving to the zIIP specialty engines, except where the application and data server are running in separate LPAR partitions and communicating via HiperSockets (*which uses TCP/IP*). The former category comprises a significant proportion of overall System z9 DB2 workloads, and thus generates a similar proportion of IBM's mainframe hardware and software revenues. The firm has taken the view that these "native workloads" have already shown themselves to be cost-effective and business-justified in their current deployment on the System z9 platform. This is an understandable position, but we do expect IBM will receive many customer requests to widen the qualifying workloads span.

The motivation for introducing the zIIP, it must be understood, is to attract and generate additional, incremental and new workloads onto the System z9-z/OS-DB2 platform, and we consider it will strongly encourage such workload growth over the next several years.

The motivation for introducing the zIIP, it must be understood, is to attract and generate additional, incremental and new workloads onto the System z9-z/OS-DB2 platform, and we consider it will strongly encourage such workload growth over the next several years.

7. System z9 Content Management and Content Integration Service

Content Management Introduction

The second information service category supporting IAAS is content management (*with their associated content integration services*). Content services provide the ability to capture, integrate, manage and deliver all forms of digital content across an entire company to create real business value, and to deliver it as a service. Content integration services are concerned with capturing and integrating content from all its native sources, and placing it under common content management repository control.

Amongst the main current business challenges that enterprises are facing in content management are:

- Records and retention management, archiving and discovery.
- The rapidly accelerating move from paper to electronic forms of content.
- The need to manage content based on user roles, and to deliver industry expertise.
- A need to integrate content and business processes across complex heterogeneous environments.

To meet these challenges, an ECM solution needs to provide a considerable breadth and depth of support for every type of digital content, and a wide range of primary services to support their lifecycle management. These ECM solution requirements are shown in Figure 20.

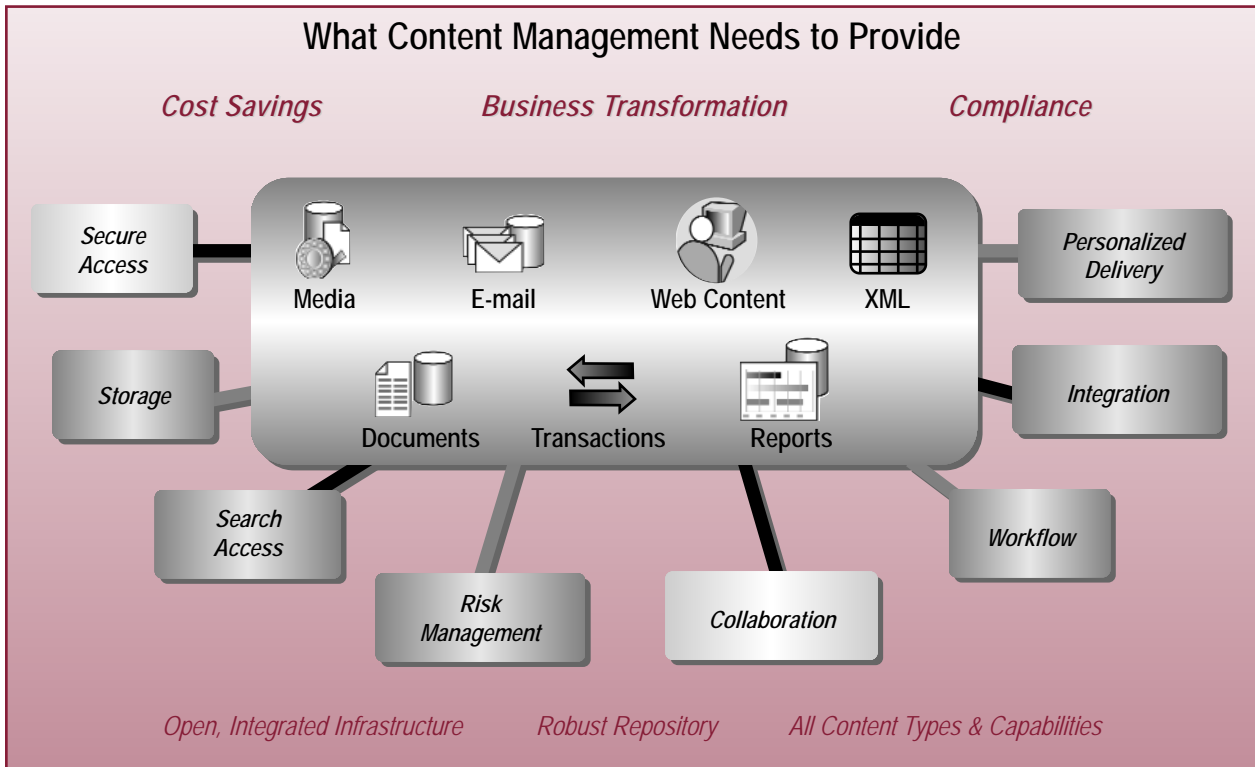


Figure 20: What Content Management Needs to Provide

In the center of the chart, today's many types of digital content are illustrated. These include rich media, e-mails, Web content, XML and other documents, transactions and reports. The ECM functionality required includes secure access to content, content storage, search access to find required content, risk management, support for collaboration and business process workflows, content integration, and personalized and/or role-based content delivery to users. Ideally, all these types of content, and all these primary content services/functions, should be embraced within, and be provided by, one unified ECM platform that is used enterprise-wide. There are clear advantages in this approach, as opposed to the more commonly found current patchworks of overlapping partial solutions that are duplicative, costly and complex to operate and support, and that often provide incomplete coverage.

Such a complete ECM platform provides the best basis for supporting business transformation, enabling regulatory compliance, and for achieving cost savings across the ECM environment. At the same time, most enterprises want and need to be able to implement their content management strategy incrementally, in line with business priorities, and taking advantage of their existing infrastructures and solutions, so a modular approach that embraces existing elements and supporting open standards, is called for.

IBM Content Management Product Portfolio

IBM has been the market leader in the provision of content management and integration middleware software, accelerators, and industry solutions for several years. After recent major software developments, releases, and the integration of several acquisitions, IBM now offers a fully integrated, comprehensive ECM product portfolio and a wide range of industry solutions. However, the portfolio is modular, can be implemented incrementally, is deeply based on open standards, and supports many existing environment components, as well as providing its own extensive content management and integration services. A variety of products, components/optional features, and packaging/pricing options provide this modularity, and are designed to suit different levels and scales of customer requirements, and the major hardware platforms including System z9-z/OS (*our focus here*) are supported.

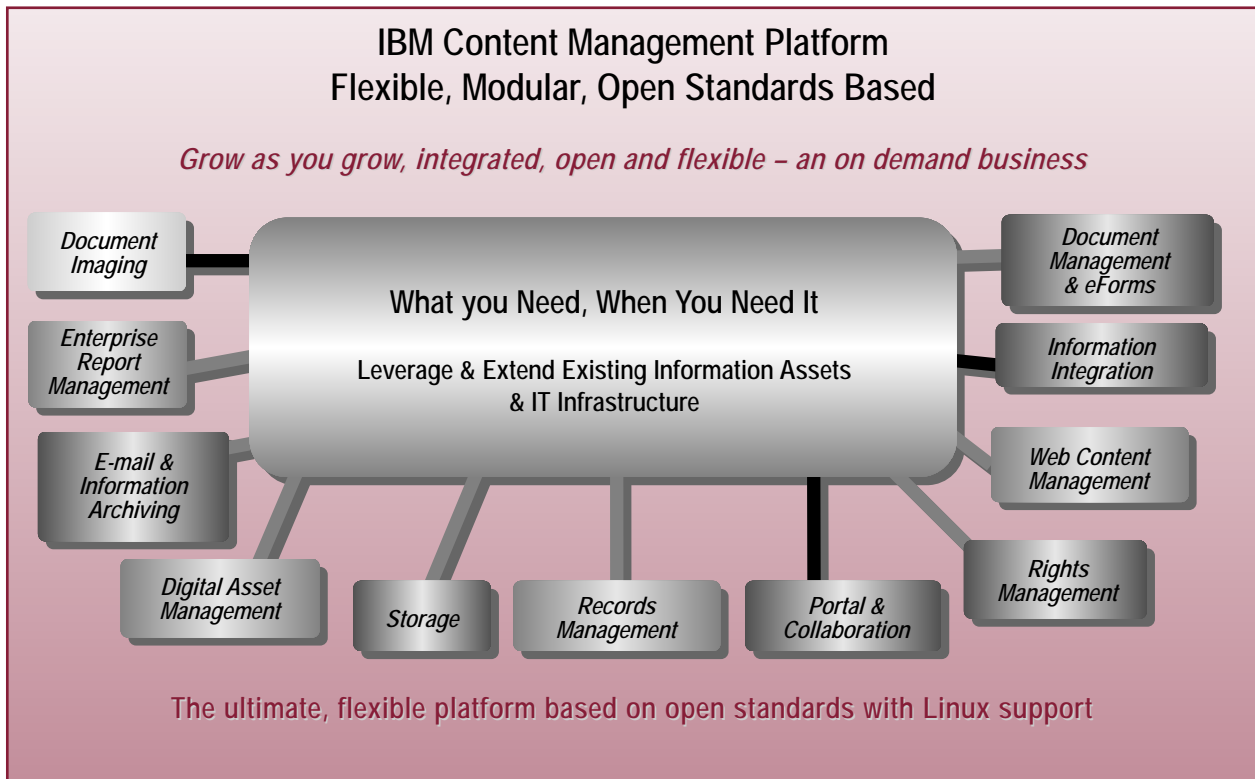


Figure 21: IBM Content Management Platform – Flexible, Modular, and Open-standards-based

The functionality included within the IBM solution, shown broadly in Figure 21, is extensive. It now includes the following content management services that can be used to support the broad range of business processes that use unstructured data:

- Document imaging support.
- Enterprise report management.
- E-mail, message and information archiving.
- Digital asset management.
- Content storage.
- Records management.
- Portal and collaboration.
- Rights management.
- Web content management.
- Information integration.
- Document management and e-forms.

The IBM content management product portfolio that delivers these common, integrated content management services now includes:

- **DB2 Content Manager:** Manages images, documents, Web content, multimedia assets, and others, with built-in workflow. Enterprise and Standard Editions are offered. (*zSeries/System z9, iSeries, and Windows/UNIX/Linux server platforms supported.*)
- **DB2 Document Manager:** Manages the complete lifecycle of business documents. (*Windows server and client platforms.*)
- **DB2 Content Manager OnDemand:** Provides output/report management and report distribution. (*zSeries/System z9, iSeries, and Windows/UNIX/Linux server platforms supported.*)
- **DB2 Records Manager:** Is an award-winning, certified records management solution. Federated Records Management is a multiple repository records management solution. (*Windows/UNIX server platforms.*)
- **DB2 CommonStore for SAP:** Archives SAP data. (*UNIX, Linux, Windows and OS/400 server platforms.*)
- **DB2 CommonStore for Lotus Domino:** (*Archives Domino e-mail and attachments. AIX and Windows server platforms.*)
- **DB2 CommonStore for Exchange Server:** (*Archives Exchange e-mail and attachments. Windows server platform.*)

This portfolio now provides a standards-based content repository framework and advanced content integration capabilities that bring information of all forms together. The spectrum of content management services offered are built on a portal framework that delivers information to knowledge workers, on demand. These content management facilities are fully integrated with the IBM Workplace client to enable role-based content management and collaboration for casual and power users.

DB2 Content Manager Architectural Integration with WebSphere Family

Architecturally, DB2 Content Manager acts as the foundation content repository first architectural layer, and delivers the event framework, document library, document process management, and content lifecycle management functionality of the platform. It is based on IBM DB2 UDB, which provides the content storage and data management services CM requires. DB2 Content Manager integrates upwards into the second architectural layer of WebSphere Application Server that provides run-time support for content services (*which include collaboration, workflow, personal information management and e-mail, dynamic learning, ECM and others*). An IBM Content Management SDK supports integration of the above with the third layer elements of IBM WebSphere Information Integrator (*which provides content integration, enterprise search and data integration services*), and with IBM WebSphere Business Integration (*which provides process management and application integration services*). These integrate into the fourth layer of IBM WebSphere Portal (*which provides application aggregation, personalization, process choreography, and pervasive computing support services*). This in turn delivers a fifth presentation layer of the architecture, via the IBM Workplace client, to the end user.

Latest IBM Content Management Portfolio Releases Provide Content Management as Services

2005 saw the announcement and availability of new releases of several main products within the IBM content management portfolio, the Version 8.3 releases. These delivered what IBM claimed is the industry's first and only content and information integration platform to leverage information across diverse repositories, and to provide a single view of business-critical information. It also delivered automated indexing and searching capabilities, providing better control over structured and unstructured data repositories, and making it easier for end users to store and manage content. These new releases included:

- **IBM DB2 Content Manager Version 8.3:** IBM's flagship content management product at the centre of the portfolio is IBM DB2 Content Manager, and the latest Version 8.3 release was announced on 03.09.2005. This extended the software's reach to enable business processes, with improved XML handling, simplified records control, and new industry solutions. The portfolio now uses Web services to easily enable all types of data to be captured, managed and searched as part of normal business processes, including support for Microsoft .NET environments. The workflow capabilities now help clients to automate business processes by providing graphical end-user tools. High-performance document routing integrates with these workflow capabilities to help streamline business processes. Complete automation of the capture and management of XML documents is enabled in a common content repository. To help streamline compliance efforts, IBM also added the first automated records control for content, through integration with DB2 Records Manager. This integration ensures automatic, consistent and accurate records declaration and classification across all forms of business content. Other enhancements included new, rich media streaming technology.

- **IBM DB2 Document Manager 8.3:** New records management integration provided simpler, more automated document lifecycle management. Other enhancements included expanded language support, improved usability/productivity via keyboard accessibility and search enhancements, and new security measures (*through single sign-on capabilities between Content Manager and Document Manager*). Together with **IBM DB2 Records Manager 4.1.1**, the software supports government and corporate industry standards, and is integrated with standard office and CAD applications, enabling end users to simply save documents as corporate records.
- **IBM DB2 CommonStore 8.3:** Provides comprehensive archive/retention management of e-mails, attachments and messaging system content. Records management integration allows end users to save e-mail messages and attachments as records, manually or automatically. Manages, retains and disposes of messaging system content to help comply with regulatory, legal and corporate policy requirements. Improved full text indexing, full text searching and retrieval capabilities provide access to all content, message body and attachments from either native client interfaces or archive repository interfaces. Single instance archiving ensures that only one copy of a document is kept in the archive, regardless of how many times the same document was archived by different users. Extended storage platform support helps to use the most appropriate storage devices according to lifecycle value of the content.

The DB2 Content Manager portfolio products were developed at the IBM Silicon Valley Laboratory and the Almaden Research Center sites in California.

IBM DB2 Content Manager for z/OS Version 8.3, Enterprise Edition

For zSeries/System z9 users, the main content management product offering is now IBM DB2 Content Manager for z/OS, Version 8.3, which became available on 05.06.2005. This release broadly incorporates the Content Manager advances described above to provide a completely integrated, reusable, and open strategic ECM platform that is secure, scalable, and extensible. It also exploits the unique QoS of the System z9-z/OS mainframe platform, and integrates with other mainframe software offerings. An Enterprise Edition provides the most complete package of content management capabilities for the mainframe platform.

Other specific advances in IBM DB2 Content Manager for z/OS, Version 8.3, included support for Tivoli Storage Manager (*TSM*) APIs, Java™ Connector APIs, document routing enhancements, query/search enhancements, and a new pricing structure. These enable exploitation of existing infrastructure software (*TSM*), enable access to Content Manager within customer business applications, and improve search effectiveness. Using TSM can increase maximum object size, offers higher concurrency (*than with OAM on DB2*), and offers the expanded choice of hardware storage and retention devices that TSM supports.

Document routing was considerably enhanced to include decision points, actions, action lists, parallel routing, and LOB access nodes. A new graphical builder (*in the system administration client*) enabled non-technical business users to define document routing processes. This supported process modeling of simple sequential routing and sophisticated routing with decision points, actions/action lists, parallel routing, and LOB access.

New Java host z/OS-based APIs were made available under UNIX Systems Service (*USS*), and allowed integration with z/OS-based Java applications running under USS, or with WebSphere Application Server Java applications on z/OS.

Administration data export/import facilities were enhanced to allow the Windows system administration client to export system administration data into an XML-readable file, and to import such data from an XML file, useful to copy administrative settings from one server to another. DB2 Content Manager data model instances can also be imported and exported as XML documents. A GUI mapping tool now maps user XML schemas to DB2 Content Manager data models.

The search and query functions in Version 8.3 were enhanced to include query on checked-out items, row-based view filtering in queries, to provide query result counts, to add a new IN operator, and to introduce an internal query optimizer that shortens generated SQL statements.

Accessibility (*usability*) and serviceability (*logging control*) improvements were also included.

Included in the IBM DB2 Content Manager for z/OS package license are limited-use licenses for DB2 UDB, WebSphere Application Server, Directory Server, WebSphere MQ and WebSphere MQ Workflow, and DB2 Information Integrator for Content, all for use in content management roles, providing the complete IBM software infrastructure needed.

The inclusive DB2 Information Integrator for Content capability provides a consistent content integration framework across all supported servers. A federated connector is provided as the common interface for federated DB2 Content Manager integration applications. This allows data and digital content sources to be accessed with common APIs and components. Applications can be written to this federated connector, as well as to the following native connector interfaces that are included: DB2 Content Manager V8.2 & V8.3, DB2 Content Manager OnDemand for Multiplatforms V7.1, V8.3, DB2 Content Manager V7, DB2 Content Manager for iSeries V5.1, V5.3, DB2 Content Manager OnDemand for iSeries V5.2, V5.3, DB2 Content Manager for OS/390, V2.3, DB2 Content Manager OnDemand for z/OS V7.1, DB2 Content Manager ImagePlus OS/390 FAF V3.1, ImagePlus/390 ODM V3.1, DB2 UDB 8.21, JDBC with JDK 1.4.21 and ODBC 3.01. These provide comprehensive content integration support for existing multi-platform content management and data environments.

The product fully exploits the power, scalability, security, and advanced storage support and software environment of the zSeries/System z9 z/OS platform...

The product fully exploits the power, scalability, security, and advanced storage support and software environment of the zSeries/System z9 z/OS platform to provide a truly enterprise scale and strength ECM platform.

Content Management as Services Well Advanced, Supports IAAS Strategy

As can be seen from the foregoing, DB2 Content Manager and the supporting products in the IBM content management portfolio/architecture have advanced in supporting the delivery of content management functionality as services, supporting the broader IAAS vision and strategy.

These advances are based on message-based communication between system components, now supported by the Web services interface introduced in Content Manager Version 8.3 and through CM services support for Web services-oriented tools (*Rational tools*). They are also supported through the new model of SOA inter-component communication provided by the new **WebSphere ESB product**. Another important advance has been the development and adoption of open language binding (*or API*) standards for content, with the new Java Content Repository (*JCR*) standard developed through the JSR170 working group, Version 1 of which was ratified in 2005, and which IBM strongly supported. IBM has already introduced JCR support in the previous release of WebSphere Portal Server V5.1, is using JCR in the Workplace collaboration family, and continues to invest via strong participation in continuing JCR standards efforts.

Other IBM Content Management Advances & Investments

IBM has also been active with further investments and other related developments in unstructured information and content management, and we briefly summarize some of these other recent contributory developments below:

- **Venetica Acquisition:** On 10.07.2004, IBM acquired Venetica, the leading pioneer in Enterprise Content Integration (*ECI*) software and services, for its unstructured content integration software (*VeniceBridge*). This enabled real-time federated access across disparate unstructured content sources, and provided the ability to work with the underlying content, and to make its information actionable. It added an additional abstraction layer that allowed users to incorporate content functionality directly into enterprise applications. The Venetica software supported a wide range of important content sources. This acquired technology has now been integrated/incorporated into the major new IBM ECI product described below.
- **WebSphere Information Integrator Content Edition, Version 8.3:** This new IBM product release, which became available on 11.30.05, provides a unified interface to work with content from any source, and provides bi-directional access to support the functionality of the underlying ECM systems supported. WIICE's flexible, highly scalable abstraction layer makes applications repository-independent. WIICE provides federation services for working across multiple content sources, development components for integrating content capabilities into business applications, and supports connectivity to a range of leading ECM systems. These include major imaging/document management, report management, and Web content/digital asset management packaged systems, customer applications, network file systems, and workflow/BPM systems. The wide range of out-of-the-box WIICE connectors unifies most major content sources and workflow systems, and slashes the cost, complexity and risk of custom programming otherwise needed to integrate these.

WIICE connectors available include: IBM DB2 Content Manager, IBM DB2 Content Manager OnDemand (*both distributed and z/OS*), IBM WebSphere MQ Workflow, IBM WebSphere Portal Document Manager (*read save*), IBM Lotus Domino Document Manager, Documentum Content Server, FileNet Content Services, FileNet Image Services, FileNet Image Services Resource Adapter, FileNet P8 Content Manager, FileNet P8 Business Process Manager, Open Text Livelink, Microsoft Index Server/NTFS, Stellent Content Server, Interwoven TeamSite, Hummingbird Enterprise DM. In addition, read-only relational database access is supported to DB2 UDB, Oracle, and any database accessible through the WebSphere Information Integrator federated data server. (*AIX, Linux, Solaris and Windows server platforms.*)

- **WebSphere Information Integrator OmniFind Edition, Version 8.3:** WIIOE is IBM's first commercial enterprise search and text analytics offering. It provides a broad platform of robust searching capabilities across the enterprise, not just for keywords but within content. The product is focused to meet enterprise requirements of security, support for a variety of enterprise content sources, understanding relevancy for content, scalability, performance and availability. The product employs an open architecture, supporting third-party analytics and partner solutions, and is highly complementary to IBM's portal, content management, collaboration and information integration offerings, discussed throughout this White Paper. It can be used to power Intranets, Extranets, and corporate public Web sites, to find the most relevant enterprise data for employees, partners, and customers, and also provides a rich platform for building high-value, text analytic solutions.

Out-of-the-box connectors are available for the following systems: IBM DB2 Content Manager Multiplatform V8.2, V8.3, IBM DB2 Universal Database for Linux, UNIX, Windows V8.1 and V8.2, IBM DB2 Universal Database for z/OS V7 & V8, IBM DB2 Universal Database for iSeries V5.3, EMC/Documentum 4.2x, .2.5, and 5.3, FileNet CS 5.3 and 5.4, FileNet P8 CM 3.0, 3.5, Hummingbird DM 5.1.0.5 with SR4, Informix Dynamic Server V9, IBM Lotus Domino Document Manager 6.5.1, IBM Lotus Notes/Domino Server 5.0.9a, or later, 6.0.2CF2, Lotus Domino QuickPlace 6.5.1, Microsoft Exchange Server 2000, 2003, Microsoft SQL Server 2000, Newsgroup (*NNTP*), Open Text Livelink Enterprise Server 9.2 and 9.5, Oracle 9i and 10g, Sybase 11.9.2, 12.0 and 12.5, UNIX file systems, Windows file systems, Web (*HTTP or HTTPS*), IBM WebSphere Portal 5.1 Web pages, WebSphere Portal 5.1.0.1 Document Manager, Workplace Web Content Management 2.5, 5.1, and mainframe sources accessible through IBM WebSphere Information Integrator Classic Federation for z/OS (*VSAM, IMS, CA-Datcom, Software AG Adabas*). This is an impressively comprehensive set of searchable sources. This product release was made available on 12.20.2005. (*AIX, Linux, Windows and Solaris Server platforms*.)

- **UIMA standard framework for text analytics:** Unstructured Information Management Architecture (*UIMA*) is an important emerging standard, developed by IBM Research. It provides an open software framework to support sophisticated text analytic applications that discover and understand latent meaning, relationship and relevant facts buried within unstructured information content. UIMA makes it easy to integrate text analytic software tools with end-to-end enterprise application, and provides tools to speed the creation of new, reusable analytic software components to handle unstructured information. UIMA was the result of over 4 years of IBM Research development, and received significant support from the Defense Advanced Research Projects Agency (*DARPA*), and working group inputs from a number of leading universities and companies.

On 08.08.2005, IBM announced plans to make UIMA available through open source. The technology has been presented to the Open Source Technology Group, with availability through SourceForge around the end of 2005. Widespread commercial adoption of UIMA by more than 15 software vendors was announced on the same date.

The UIMA framework has already been embedded in IBM products, including IBM WebSphere Information Integrator OmniFind Edition (*see above*), the first commercially available software platform for processing content based on the UIMA standard. IBM WebSphere Portal Server and Lotus Work Place also leverage UIMA for content processing.

- **PureEdge Acquisition:** On 08.05.2005, IBM completed its acquisition of PureEdge Solutions. IBM is integrating the acquired PureEdge e-forms functionality into its collaboration technology portfolio, including the IBM Workplace, WebSphere, and Lotus offerings. This acquisition also complements its fast-growing content management business. E-forms allow structured exchange of data (*such as inventory figures, customer data or pricing information*) between end users and back-end systems, more efficiently and more accurately than paper forms. This new e-form technology is based on XML, which provides a consistent format for sharing documents/transactions between applications. PureEdge had actively supported the XForms electronic form document open industry standard, a platform-independent markup language for data capture/validation in electronic forms. IBM now actively supports this open standard that provides customer flexibility and choice (*avoiding proprietary forms technology vendor lock-in*) and is participating heavily in its development.
- **iPhrase Systems Acquisition:** IBM acquired Phrase Systems, Inc. on 11.01.2005. iPhrase software improves e-commerce sales, online service and support, and call center productivity by allowing Web site users to more easily find answers, make purchases and solve problems without expert assistance, using insight into the intent of user requests. It allows deeper query understanding, adaptive results presentation, fuller business control of the user experience, and the delivery of in-context merchandising capabilities. iPhrase technology is compliant with UIMA, discussed above. iPhrase technologies are being integrated into IBM's information management software portfolio, broadening its enterprise search and content management capability, including its incorporation into a future WebSphere Information Integrator OmniFind Edition release, and elsewhere. iPhrase was the fifteenth software acquisition IBM has made since 2001 in support of its initiative to deliver IAAS.

These are a compelling set of unstructured content products, acquisitions and open standards developments from IBM. They highlight only a few of the more recent and significant of IBM's many IAAS advances.

Our Analysis

The content management product portfolio and content management services described above help organizations manage and integrate all forms of digital information, including scanned images, audio, and information from multiple software vendors' products that are located in various data sources.

For larger enterprises, DB2 Content Manager for z/OS Version 8.3 provides a scalable, technologically-advanced ECM solution. It fully exploits the power, scalability and capacity, availability and reliability of the zSeries/System z9 and DB2 for z/OS platform, and is the natural choice for the central, enterprise-wide, content management hub.

...DB2 Content Manager for z/OS Version 8.3 provides a scalable, technologically-advanced ECM solution.

This, and the other associated complementary content management portfolio products described above, now deliver the most extensive set of ECM and enterprise content integration services in the industry, realizing IBM's IAAS vision for content management.

Content management with the IBM content management portfolio centered on DB2 Content Manager has already helped many customers worldwide digitize and manage all forms of information, from documents to photographs to e-mails, and to leverage these contents to support new business applications that create competitive advantage. Many published references and success stories attest to the benefits.

Content consolidation and integration services needs are supported with the latest release of WebSphere Information Integrator Content Edition, as described above.

The widespread needs for more sophisticated enterprise search that can mine unstructured content itself across the multiple content repositories in use today in most enterprises is another critical service, now ably provided by the impressive WebSphere Information Integrator OmniFind Edition, with its innovative text analytics search capabilities.

Making such deeper use of content material in search, and in line of business applications, required new open standards. The UIMA framework that IBM researched with DARPA and others, and is now open source, has gained traction as the leading open-source text analytics standards framework that alone can facilitate further collaborative innovation. UIMA is already supported in IBM products, and 15 other software ISVs declared support at announcement.

IBM has also developed a wide range of content management industry solutions based on this infrastructure software, but space constraints preclude coverage here.

8. Master Data Management as a Service – IBM Steps Up

Introduction to Master Data Management

Many enterprises today operate multiple overlapping applications systems, each providing different and usually incomplete data about the main entities of the business, such as products, customers, suppliers, locations, and employees, etc. This is a rapidly increasing requirement for organizations to get a single view of these information types due to compliance, market efficiency, accuracy, etc. Each application defined and used its own versions (*and data attributes*) about one or several main entities, which resulted in mis-aligned, inconsistent, and often incomplete, master data being held, deeply embedded in these applications. Older transaction-focused systems held a narrow data set, restricted to the barest structured information attributes needed for the application's original purpose. Today, a much richer set of machine-readable and often unstructured information, is now also required. Items such as photographs, drawings, graphic images, compound documents, video clips, Podcasts, and other rich media elements are now a central part of Web, marketing, publishing and e-business systems.

In global enterprises, even the seemingly obvious need for globally unique/consistent product numbers/codes, or employee IDs, are missing. A patchwork of locally developed codes or identifiers, deeply embedded in older applications, preventing a global view or analysis being consistently made, is often found. This causes costly data redundancies, wasteful complexity, and ultimately misleading analytics.

MDM is a newer information management discipline that seeks to overcome this widespread proliferation, inconsistency and incompleteness of central entity data. MDM has been defined as the set of disciplines, technologies, and solutions used to create and maintain consistent, complete, contextual and accurate business data for all enterprise stakeholders (*within and outside the enterprise*). At its center, MDM is based on the concept of data objects that represent the main business entities of the enterprise. Principal master data objects include products, customers, organizations, locations, trading partners, employees, consumers, citizens, assets, accounts, and policies, etc. Others, including Oracle, use the term "data hub" to refer to the same concept and approach as MDM.

At its simplest, MDM seeks to give businesses a single, accurate view of customers, partners or suppliers, by drawing data from separate stores of data scattered across an enterprise into a unified, centralized, shared view.

MDM – Now an Important New Information Management Middleware Market

MDM has rapidly emerged in recent years to become an important and rapidly growing middleware market, supported with processes, methods and tools that address these challenges. This rapid MDM growth has been driven by the increasing efforts to integrate applications that enterprises have been undertaking as they transform business processes. In addition, the migration towards a new generation of SOA applications, and towards IAAS, has also made these data consistency and incompleteness issues a real roadblock hindering progress.

Market researchers describe MDM as a fast-growing, but currently highly fragmented, market segment for middleware software and services. Worth an estimated \$5.4B in 2005, the MDM market is expected to grow to \$10.4 billion by 2009, a compound growth rate of 13.8%, according to market research firm IDC.

Ad hoc, in-house and service provider service-based approaches predominated in the early growth of this market. Specialist new software vendors created the new class of MDM middleware software now growing rapidly in the segment.

IBM, and most analysts, expect that the MDM market will migrate away from the ad-hoc, custom solution building approaches that have predominated to date, towards one based on widespread use of standard MDM middleware software solutions. (*See below.*)

MDM is closely linked to, is dependent upon, and is crucial to the success of, the other primary types of information services in an IM architecture, including data management, content management, information integration and BI.

IBM Carves MDM Leadership Position

Recognizing its importance, IBM made a major entry/commitment to the MDM market. It acquired pioneering MDM middleware vendor Trigo Technologies on 04.06.2004. The acquired Trigo Product Center product became IBM's WebSphere Product Center offering described below. The more recent acquisition of DWL on 08.02.2005 brought in that firm's leading-edge customer data integration middleware, now IBM's WebSphere Customer Center offering. Three other smaller, privately held companies were also acquired in this space. IBM now has over one thousand employees dedicated to its Enterprise Master Data Solutions organization, focused on MDM middleware software, solutions and service provision, and has invested \$1.5B in its acquisitions and developments in MDM to date.

IBM's main MDM offerings are SOA-based middleware that provide organizations with flexible frameworks to support enterprise structured and unstructured data and business services, aligned with key business processes. IBM claims it now brings together all the components required for a successful enterprise MDM strategy: master data integration, content management, BI, and MDM for specific data objects – including product, customer, and supplier – and master data solutions for specific industries.

Other industry majors actively addressing MDM are the two giant enterprise application vendors, SAP and Oracle. Each of these also regards MDM as a crucial strategic area. Their traditional applications had essentially created in-built, de facto standards for master data, but these were traditionally embedded deeply within their applications, which created a form of vendor lock-in for customers needing to integrate these vendor enterprise applications with their other systems. Today, both vendors are rearchitecting and building next-generation enterprise applications based on SOA approaches, built on their standards-based middleware, and MDM is a crucial part of these efforts.

What is Driving MDM Growth Today?

Several important business drivers moved MDM from its long-standing "chronic problem" status, to become an acute and serious problem. Different initiatives demand that organizations gain control of comprehensive cross and inter-enterprise master information and processes, including:

- Industry initiatives, compliance mandates, legislation, liability, and other regulatory compliance-related demands.
- Ever more connected supply and demand chains, with decreasing business cycles, needing better-integrated systems.
- Outsourcing, mergers and acquisitions that create formidable data and systems integration challenges of ERP/application consolidation.
- On Demand Business needing improved data warehousing/BI and analytics, much more expensive and difficult to achieve without MDM in place.

These widespread drivers create a compelling business case, based around cost reduction, revenue enhancement, customer satisfaction, compliance, and reduced time to market, for enterprise MDM solutions.

Why Master Data in Middleware, Not Enterprise Applications?

To meet these challenges, master information must be decoupled and separated from the individual transaction applications in which it was traditionally embedded, and moved into common MDM middleware. These complete, accurate sources can then ensure consistent master information is shared and used, across not only all the transactional systems that require it, but also across the many analytical (*data warehousing/BI*) systems needed today, which have become problematic to keep aligned. This open and flexible middleware MDM approach offers solutions and benefits to the challenges above and can provide MDM services to all next-generation enterprise applications, both packaged and in-house built, in the customer environment.

Analysts, IBM, and both leading enterprise application vendors all now agree on this. However, the enterprise application vendors seek to extend existing lock-ins, by making their customers dependent on their own new MDM middleware tied to their applications.

IBM, alone amongst the big three, provides open MDM software and solutions that support all packaged, legacy and custom applications, and that will enable new, differentiating applications as ERP becomes a commodity, without such application lock-in.

The IBM MDM Approach

The overall approach IBM has taken to MDM is illustrated in Figure 22.

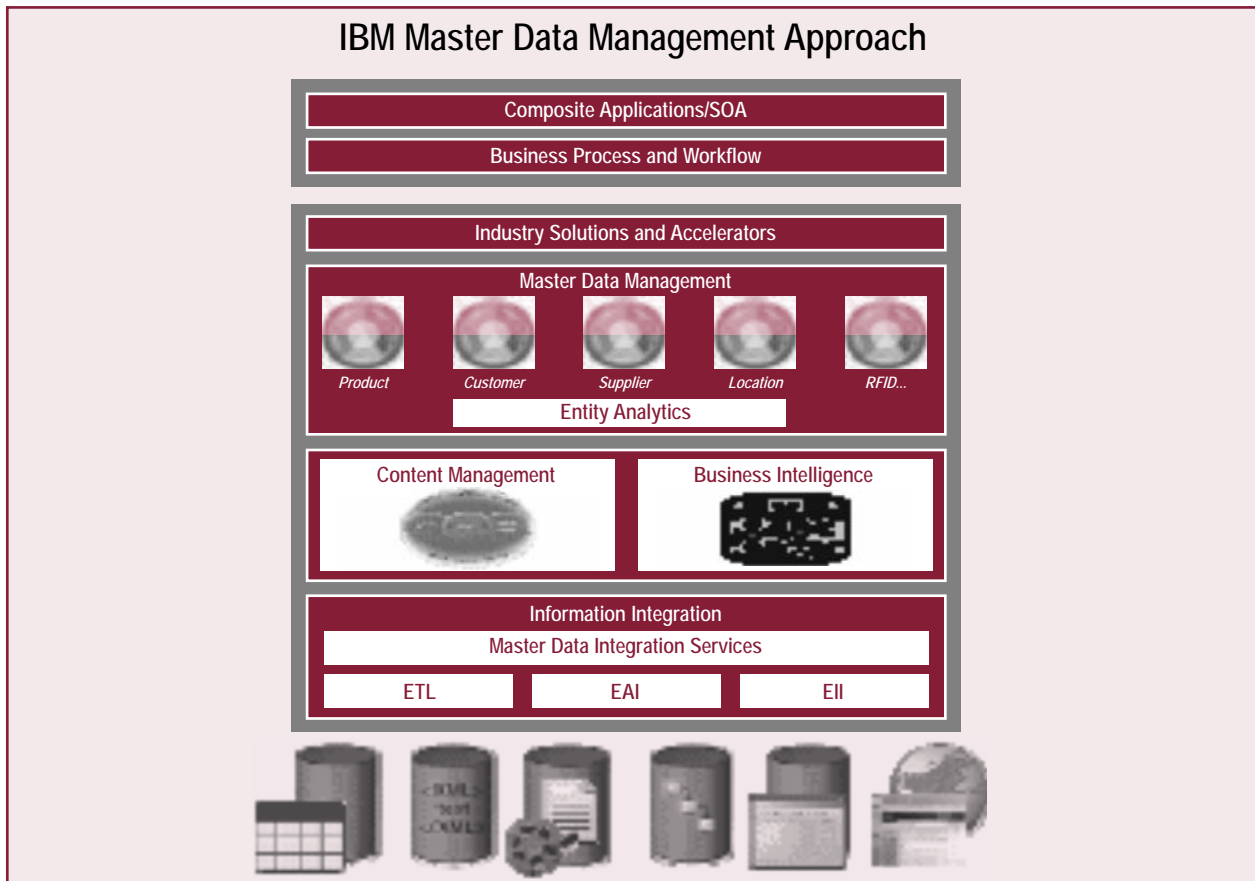


Figure 22: The IBM Master Data Management Approach

MDM systems for central data objects such as products and customers (*center*) support master data solutions that support and enable new business process workflows implemented in new SOA composite applications. (*Upper layers.*)

MDM systems draw heavily on the content management systems that hold unstructured data, and are primary feeds of metadata and master data to BI and data warehousing analytic applications.

Master data integration is a vital component, to access, cleanse, transform and load master data from the disparate source systems (*shown at the bottom of the chart*) into the MDM systems. This requires comprehensive ETL, EAI, and EAI capabilities that are well integrated and open to process the master data flows required. The same full suite of advanced IBM WebSphere Information Integration products that we assess in Section 9 are used for master data integration and for standard information integration applications.

One of the newest information management technologies, entity analytics, also fits into the center of this picture, and we discuss the important, newer IBM Entity Analytics solutions technologies in Section 11.

Product Information Management with WebSphere Product Center, Version 5.2

Many types of business, including retailing, distribution, manufacturing, consumer electronics, computer hardware, and others, are centered on the products/product lines they sell, distribute, and/or manufacture. For these businesses, the foremost MDM challenge is managing product master data effectively throughout the product lifecycle, across all the many enterprise systems that touch products.

Product information management is an enterprise MDM approach to collecting, authoring, rationalizing, and publishing or syndicating product information, across multiple internal systems and for exposing product information externally.

The Internet, e-commerce, and evolving industry standards have changed expectations about the accessibility and complexity of product information. IBM's WebSphere Product Center is a product information management solution that provides companies with a consistent, flexible, accurate and highly scalable central repository for managing and linking information about product, location, trading partner, organization, and terms of trade. It enables easier synchronization of this information internally with existing enterprise systems, and externally with trading partners. The product is based on open standards, and provides a middleware foundation for companies to address strategic product-centric initiatives, and to comply with industry standards. Broad capabilities include:

- Leverages core product information, including product attributes, price and location, increasing accuracy and making it more directly useable and valuable in business processes, which it thus helps to streamline.
- Helps companies efficiently deploy their product and service information across the many customer, partner and employee touch points involved.
- Provides a good foundation for integrated multi-channel sales and marketing initiatives that span many customer touch points, including e-commerce sites, retail stores, print, catalogs and more.
- Enables companies to implement a strategic approach to exchanging standards-based data with partners through Global Data Synchronization.
- Delivers rich product information to Web sites and e-commerce applications, printed documents and marketing collateral, kiosks and mobile devices, as well as directly to customers and trading partners through various access points.
- Links product information to product-related terms of trade information, such as pricing, establishing valuable linkages that can be leveraged.

WebSphere Product Center is a member of the IBM MDM family, and integrates with the WebSphere Information Integration and DB2 Entity Analytics products. The latest release, WebSphere Product Center Version 5.2, became available on 03.24.2005. (*ASX5L, Linux, HP-UX and Solaris server platforms.*)

Transactional Customer Data Integration (CDI) with WebSphere Customer Center, Version 6

Many categories of business are intrinsically customer-centric, and their success depends upon maximizing what is sold to each customer (*whether business or consumer*). Industry sectors where such a customer-centric view is central include retail and wholesale banking, insurance, credit card providers, other financial services, telecommunications, healthcare, and other service-centric businesses. In these, retaining customers, and successfully cross-selling multiple products to the customer base, both crucially depend on accurate and consistent customer information.

CDI is an MDM approach that combines the technology, processes and services needed to create and maintain a single, accurate, timely and complete view of the customer across multiple channels and business lines. Despite the significant time and resources they have already invested in CRM, many organizations still lack a true enterprise-wide view and update environment for their customers. By connecting these systems to an enterprise customer hub, WebSphere Customer Center (*formerly called DWL Customer*) provides a unified view of the customer across multiple business and product silos. It delivers this single view and update environment to all channels and CRM applications, enabling multi-channel integration and consistent customer service. Key features and capabilities include:

- Provides real-time, transactional CDI.
- This integrated customer data delivers a single version of customer "truth" to all customer-facing channels and front- and back-office systems through multiple interfaces.
- An industry-leading CDI solution for performance and scalability, with publicly available benchmark results available.
- Based on open standards, the product is designed for implementation within a Service-Oriented Architecture, and includes over 480 business services to manage and maintain customer data.
- It provides infrastructure foundations for companies to become more customer-centric, improving customer service and cross/up-sell execution.

Additional benefits include support for compliance with various government regulations such as Sarbanes-Oxley and Basel II, cost savings from enabling the recognition and processing of duplicate customer records, and the ability to track complex customer hierarchy and relationship information in B2C and B2B environments. WebSphere Customer Center is part of the IBM MDM family of products and integrates with the WebSphere Information Integration suite, and DB2 Entity Analytics products. WebSphere Customer Center became available on 11.14.2005. (*AIX5L server platform.*)

Master Data Management and the zSeries/System z9

zSeries and System z9 mainframes already run a large proportion of the most demanding, mission-critical enterprise applications, and provide much of the enterprise data-serving horsepower for database-intense workloads, for which they are well-suited. In our recent study (see the “Related Software Strategies Research” Section on page 72 item 1) we found the zSeries/System z9 mainframe was already playing a central role in hosting new-generation, enterprise-scale, SOA applications, and that the case for more of these to run on the platform was compelling.

MDM has a close relationship with these heavy mainframe workloads, and will have many touch-points, bi-directional data transfers, and integrations with these mainframe-resident applications.

MDM solutions should clearly be centrally located, and should be run as close to the IT central hub of the enterprise as possible. Here they can most easily integrate with, and provide services to, other enterprise applications including those running on the mainframe.

WebSphere Product Center and WebSphere Customer Center are both built on the same IBM WebSphere middleware infrastructure software and open standards that support IBM customer’s applications, use the same WebSphere Information Integration middleware for their master data integration, and support the same DB2 Entity Analytic solutions. This means they fit well into enterprise SOA architecture, and can provide their MDM capabilities as services across the whole enterprise IT infrastructure, over SOA backbone capabilities such as WebSphere ESB.

Some MDM products, such as WebSphere Product Center, are essentially managing data about data (*a form of metadata if you like*). Although strategically important, their master data volumes, system throughputs, and transaction rates are relatively modest compared to live enterprise applications, and arguably do not yet require the scale, resources, power and QoS that direct System z9 mainframe native deployment would provide.

WebSphere Product Center and WebSphere Customer Center are the IBM-named and branded successors to the third-party products acquired with their vendors by IBM (*Trigo Technologies and DWL*). Like most such newer ISV applications of recent years, both were originally developed for the UNIX mid-range platforms that are the home to the majority of ERP and CRM enterprise applications.

The fact these products are based on common IBM middleware means that they could potentially be ported to other platforms that support the same IBM middleware (*which includes the mainframe*) in future, should there be sufficient demand – although no such directions or plans have been announced.

IBM has also invested heavily in methodologies, MDM industry solutions, and other accelerators to create a complete offering around these central MDM middleware products, but space constraints preclude assessment here.

Our Analysis

MDM has emerged as a vital information discipline, and as a rapidly growing new middleware and services marketplace. In a short time, IBM has created a formidable presence, large MDM organization, and initial MDM products and solutions by acquisition that are now enjoying the full force of IBM global marketing promotion. These products are based on open standards and the IBM

In a short time, IBM has created a formidable presence, large MDM organization, and initial MDM products and solutions...

WebSphere middleware stack, and form a highly complementary MDM component of IBM’s IAAS strategy as a whole. They will play an increasingly important part in enterprise customer, next-generation SOA applications development, by resolving long-standing master data issues in product and customer-centric applications.

We expect IBM will add other core MDM solutions for other important business entities over the next several years, as it strives to move from early leadership in this new market space to the commanding share lead we expect it to achieve.

MDM systems have many touch points with zSeries/System z9 mainframe enterprise applications, its extended data-serving role and the new SOA applications that will require MDM services. The common WebSphere middleware infrastructure underpinnings both the latter share, and the power and breadth of IBM’s complementary information management offerings (*data management, content management, information integration, entity analytics and data warehousing/BI*) that span platforms, will facilitate this interoperation under a SOA approach.

9. Enterprise Information Integration Services and System z9

Introduction to Information Integration Today

Enterprise IT infrastructures grew in complexity over the last decade, and today comprise a variety of applications, technologies, platforms and databases. As a result, information integration between these many disparate systems and sources of information has become an overwhelming challenge as organizations seek to modernize their business processes supported by next-generation SOA applications. Today, information integration has become a wide-ranging information management service area that is crucial to this migration, and has now brought together a number of integration middleware technologies, standards and solutions.

Integration Middleware Streams Converge into Broad Information Integration Church

A number of generic, middleware-based integration technologies emerged over this period to help address these wide-ranging challenges. **Enterprise Application Integration (EAI)** which sought to integrate and bridge disparate enterprise applications was one of these. **Extract Transformation and Load (ETL)** middleware tools emerged, to support the growth of data warehouses, datamarts, and other mediated data conversions, by providing powerful extraction, data transformation and database loading functions that simplified and automated these processes. **Enterprise Information Integration (EII)** middleware technology emerged more recently, to link and connect the multitude of disparate data sources found in enterprises today, and to provide one single logical view and streamlined access to these.

Recently, it became clear all three approaches needed to be combined and melded with common metadata services, extensive connector libraries, and with full SOA open-standards support. Such “new-EII” platforms would provide a more unified, cohesive enterprise integration capability able to support all the transactional and BI integration needs of the enterprise within a single, common information integration framework.

IBM and Information Integration, Ascential Acquisition

IBM made an early entry, staked a substantial position in EAI with its CrossWorlds acquisition, and pioneered EII with its own WebSphere Information Integration offerings. Then in 2005 with its major acquisition of ETL leader Ascential, it gained a well-established suite of ETL technologies. After extensive integration of these three elements, and other substantial, innovative, research and development, IBM has emerged with the industry's first complete new EII middleware offering that unifies and integrates EAI, ETL and EII into one cohesive platform. This is the WebSphere Information Integration suite. A full range of integrated products implementing all the base elements of this vision are available today, covering a wide range of platforms, types of integration, data sources, and available in a number of different targeted editions to meet specific segment needs.

However, a large ongoing development effort is also well advanced to extend and further unify this comprehensive offering to a higher level, with a new class of central server, and advanced metadata services shared across the entire portfolio, both announced and well underway for delivery in 2006. Under the project code name “Serrano-Hawk”, these further advances combine the Hawk developments underway at Ascential with the Serrano developments IBM itself had begun. When delivered, these further developments will extend the sophistication, integration and power of what is already the industry's broadest information integration offering.

IBM already holds a number one market share position in this now more broadly defined EII market, with over 5,000 customers using its EII middleware and solutions today, and many ISVs and partners supporting the framework/platform.

WebSphere Information Integration now provides a broad strategic framework to help customers speed new application deployment and control IT costs. It gives companies real-time, integrated access to business information – structured and unstructured, public and private, mainframe and distributed. The cornerstones of the framework include data and content federation, replication, event publishing, and enterprise search (*with WebSphere Information Integrator OmniFind Edition*).

These capabilities provide EII as a set of services. These support IBM's IAAS vision/strategy, and can be deployed within a SOA to support the information integration needs of customers' next-generation, composite applications that are today being built and deployed on the WebSphere SOA Foundation application platform (*reviewed in Appendix D*).

WebSphere Information Integration Portfolio Overview Today

A simple introductory overview of the areas of functionality, and of the principal IBM middleware software products available today, in the IBM WebSphere Information Integration suite, is shown in Figure 23 on page 50.

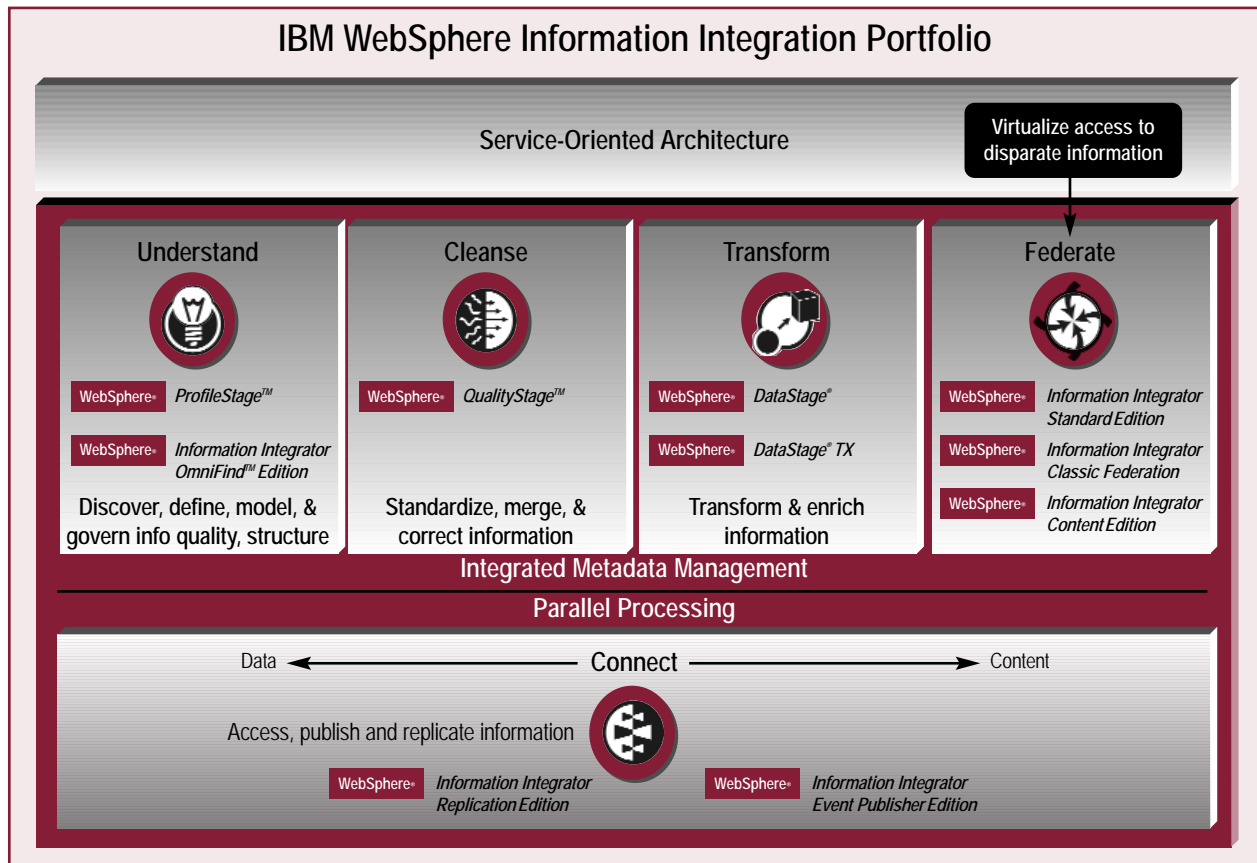


Figure 23: IBM WebSphere Information Integration Portfolio

The Figure shows that the whole offering operates under a SOA, and encompasses four broad areas of integration functionality at the next level, labeled as “Understand”, “Cleanse”, “Transform”, and “Federate”, each described by a short explanation, and with the currently available WebSphere Information Integration products shown within each area. These are supported by the next layers of integrated metadata management, and support for parallel processing, that sit above the base “Connect” area that spans from data to content, and covers access, publishing and replication of information.

The architecture embraces all types of data, including transactional, operational, and analytical information, structured and unstructured information, and mainframe or distributed platform-based information, private or public information, and can support both live, real-time access as well as batch processes.

The same information integration products (*services*) support new enterprise customer applications, as well as content integration and master data integration. They also help ensure data quality and consistency by applying the same advanced data quality and cleansing algorithms across ETL, EAI and EII integration operations. The products are also designed and architected for heavy-duty enterprise-scale use, with high performance and scalability attributes, exploiting parallel processing technologies to give high throughput where needed.

This open EII platform has been designed/developed to provide businesses with the maximum advantage from their current investments in data and content, business logic, and infrastructure that are essential for competitive advantage. We consider the offering provides a breadth of functionality that no other integration vendor offers, making it clearly the leading single-vendor enterprise integration solution that stands up to all the requirements of today’s complex enterprise architectures. It therefore offers a more complete solution, lower costs, less complexity, and lower risks than any patchwork of multi-vendor integration solutions could provide.

Main WebSphere Information Integration Suite Products

Below we briefly summarize the main currently available products that comprise the WebSphere Information Integration Suite today, to position and explain the specific offerings shown in Figure 23:

- **WebSphere ProfileStage:** (*Understand*). This product automates data discovery and source analysis, including column analysis, table analysis and cross-table analysis, and helps in any type of data migration/evolution project. It finds relationships in the data, filters data by business rules, and can generate DDL and ETL job definitions for the target. The product essentially reverse-engineers the metadata from the source data by providing a good understanding of the data, it reduces time-to-value and minimizes overall costs and resources for critical data integration projects. Enterprise and Extended Editions are offered. (*AIX5L, HP-UX, Solaris, Linux and Windows server platforms.*)

- **WebSphere Information Integrator OmniFind Edition:** (*Understand*). Enterprise search capability, described in Section 7 on page 38.
- **WebSphere QualityStage:** (*Cleanse*). This product provides information standardization and cleansing by resolving format and context data inconsistencies between systems. It uses a variety of investigation techniques (*parsing, lexical analysis, and context sensitivity*) in this cleansing process, and handles standardization, matching and survivorship. Bringing inconsistent data into clean and complete shape before integration is the vital role performed. Enterprise (*AIX5L, HP-UX, Solaris, Linux, Tru-64 UNIX and Windows server platforms*) and Enterprise MVS (*z/OS server platform*) Editions are offered.
- **WebSphere DataStage:** (*Transform*). Integrates data across multiple and varied types and high volumes of data sources and target applications – in the time available. Provides ETL capabilities for large amounts of data in real-time or bulk mode, and manages multiple integration processes. Handles all transformations, from the simplest to the most complex. Optional metadata management, high-performance parallel ETL processing, and on demand capabilities, are offered, depending on Edition. SOA, Extended, Extended SOA, Enterprise, Enterprise SOA, and Enterprise MVS Editions are offered, with differing capabilities/feature sets. (*Enterprise MVS Edition runs on z/OS server platform.*)
- **WebSphere DataStage TX:** (*Transform*). Integrates applications across platforms through a codeless, graphical approach, supporting enterprise-class transformation for mission-critical integration. The products can handle many-to-many and complex data integrations, supports Web services environments allowing data integration logic to be deployed within a SOA, and supports key open standards like Web services, JMS and EJB. Event-driven integration is also supported. High availability, high performance, and high reliability, with smooth request load balancing across WebSphere DataStage TX servers, are also supported. SOA, Extended and Extended SOA Editions are offered, with combinations of the above features. (*AIX, HP-UX, HP-Non-Stop, OS/390 and z/OS, Solaris and Windows server platforms.*)
- **WebSphere Information Integrator:** (*Federate*). IBM's flagship information integration framework brand that helps customers get new applications out faster, capture more value and insight from existing assets, and control IT costs. Meets a wide range of data integration requirements for BI and business integration. Multiple Editions provide a range of capabilities, including enterprise search, data federation, data transformation, data placement (*caching and replication*), and data event publishing. These add value to integration projects like building a corporate Intranet, augmenting a data warehouse with real-time data, building a unified view of customers or products, or managing data consistency, distribution, or synchronization across applications. Editions include OmniFind and Content (*described elsewhere*), Event Publisher, Replication, Standard, Advanced, Advanced Unlimited, and Developer. Version 8.2 is the current release. (*Server platforms in general are AIX, HP-UX, Linux, Solaris and Windows.*)
- **WebSphere Information Integrator Classic Federation for z/OS:** (*Federate*). This important z/OS product provides direct, real-time SQL access to mainframe databases and files without mainframe programming. JDBC or ODBC SQL statements are dynamically translated into native read/write APIs. The product can access data stored in VSAM, IMS CA-IDMS, CA-Datcom, Software AG Adabas, and DB2 UDB for z/OS databases, and SQL SELECT, INSERT, UPDATE, DELETE statements and stored procedure calls are supported. Driven by metadata that contains a mapping of physical databases and files to logical relational tables, all these mainframe data sources appear as one logical relational database, in a product that provides high-performing, scalable mainframe data integration. Windows and UNIX tools and applications are supported. Combined with WebSphere Information Integrator V8.2, it delivers enterprise-wide structured and unstructured data integration. Version 8.2 is the current release. (*Server platform z/OS.*)
- **WebSphere Information Integrator Content Edition:** (*Federate*). Enterprise content integration, described in Section 7, on page 38.
- **WebSphere Information Integrator Replication for z/OS:** (*Connect*). New offering that supports business continuity and workload distribution across DB2 UDB instances through high throughput, low latency queue-based replication. The product synchronizes mainframe DB2 UDB data with DB2 UDB, Oracle, Informix, Sybase, Microsoft SQL Server, and Teradata-based applications, through flexible SQL-based replication architecture. It manages data distribution and consolidation between headquarters and retail or branch sites, and can deliver more current data to data warehouses for tactical decision making. It also offers the ability to initiate business processes based on DB2 UDB transactions. Version 8.2 is the current release. (*Server platform z/OS 1.4+.*)
- **WebSphere Information Integrator Event Publisher for DB2 UDB for z/OS:** (*Connect*). This product captures changes made in DB2 UDB for z/OS and publishes these changes to WebSphere MQ. These data events can be used by an application, tool, or message broker to drive additional processing, such as data replication, application-to-application communication, and workflow initiation, etc. The changed data items are captured using DB2 UDB recovery log mechanisms. All captured data is reformatted into a consistent XML message format used by WebSphere Information Integrator products regardless of the type of source data. Changes to the applications that use the published data events can be made with almost no impact on the event-driven integration. Typical applications include pushing operational customer data updates to a CRM, delivering changed data to a data warehouse, or initiating a workflow action on a data event. Version 8.2 is the current release. (*Server platform z/OS 1.4+.*)

This broad range of information integration products, in multiple editions, available today meets almost every information integration scenario need under a common framework. In the System z9 context of this White Paper, seven of the major products above are available for the zSeries/System z9-z/OS environment, providing extensive coverage of the integration styles and data source coverage most needed in mainframe environments.

Extensive Range of Connectors

Across enterprise environments, scores of different data sources are deployed that may require connection in information integration projects. WebSphere Information Integration supports the widest range of standard connectors that enable out-of-the box integration with that source without development. Figure 24 shows this extensive collection, which needs no further comment.

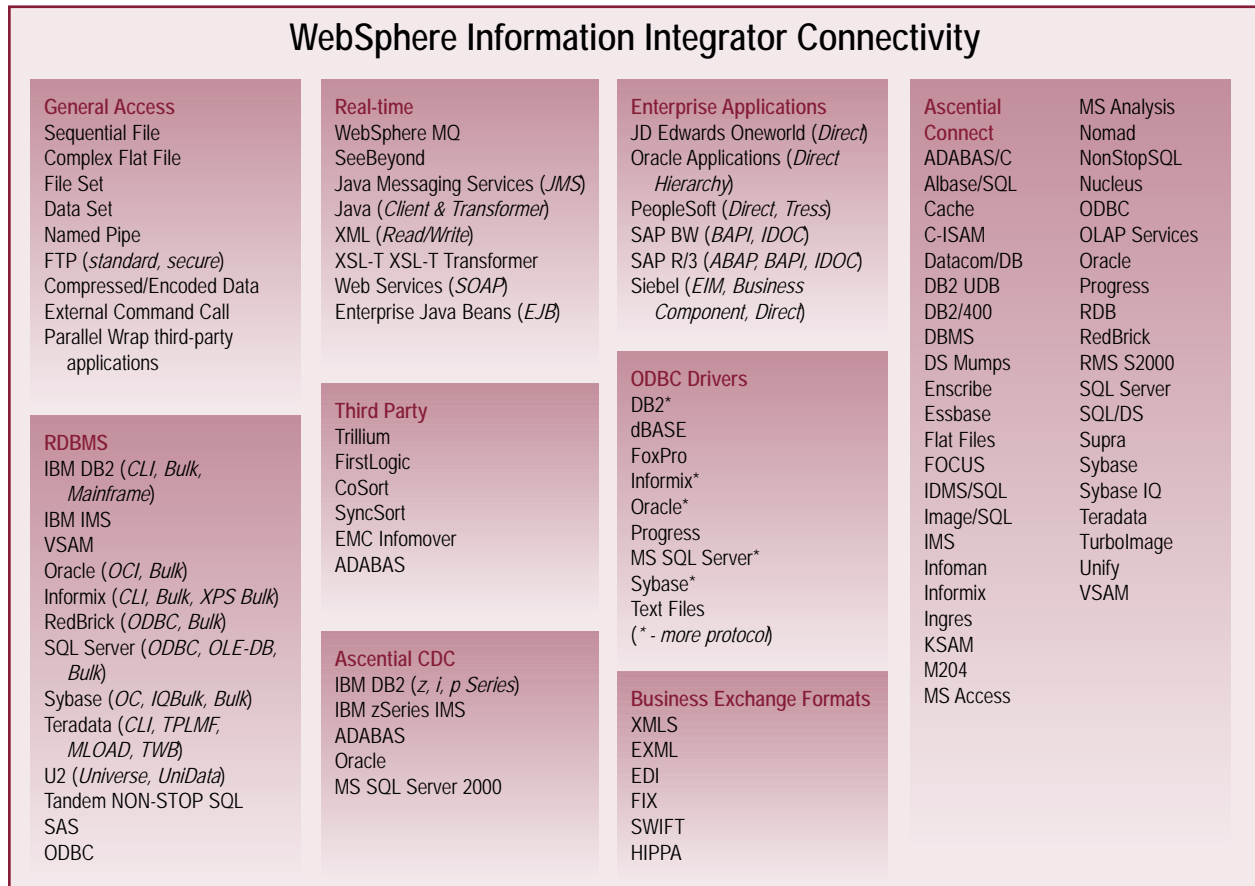


Figure 24: IBM WebSphere Information Integration Connectors

Major Further Developments Coming

As mentioned above, extensive further development has been in progress to extend and advance the WebSphere Information Integration suite to a new level of cross-suite integration, power and capability. Under the project code name "Serrano-Hawk", these major further advances have combined the Hawk developments that were underway at Ascential, with the Serrano developments IBM itself had begun. When delivered, these will considerably extend the sophistication, integration and power of what is already the industry's broadest information integration offering.

Several important new products for the suite resulting from these efforts have already been previewed, and are expected to become available during 2006. To complete this review of IBM's Information Integration capabilities, we briefly outline these major strategic additions below:

- The IBM Information Server:** This major new engine can best be described as doing for information integration what the application server did for J2EE, by providing a comprehensive, high-function information integration run-time services platform in a single, unified and comprehensive server. The centerpiece of IBM's IAAS strategy, WebSphere Information Server will take its place at the heart of SOA infrastructures, alongside WebSphere Application Server and WebSphere Process Server, as one of the three foundation IBM middleware engines for SOA. Figure 25 (*on page 53*) shows the broad positioning of WebSphere Information Server.

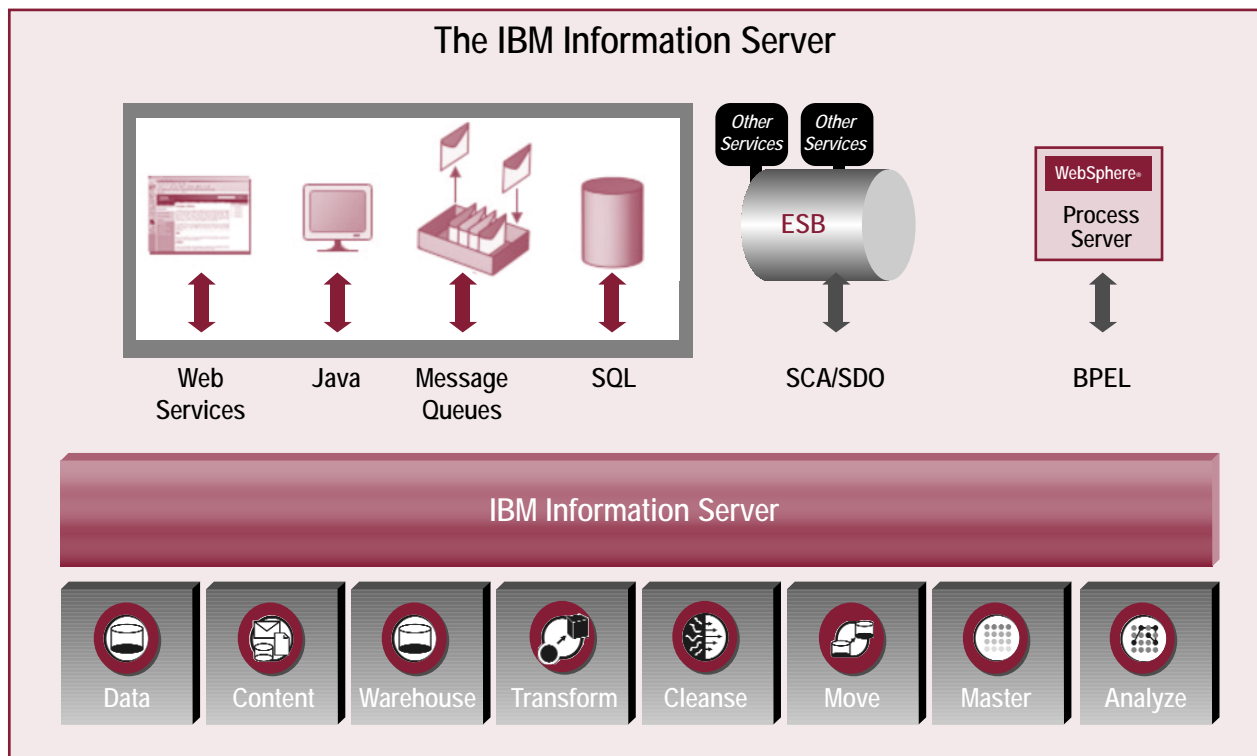


Figure 25: The IBM Information Server

Providing trusted information and control, the IBM Information Server is designed to provide rich capabilities for understanding, cleansing, transforming, moving and accessing diverse data sources. It also provides a single framework for Information Services by delivering shared metadata, logging, security, services registry, configuration, and installation; thereby providing more control over, and manageability of, data assets.

It essentially implements a full enterprise information virtualization approach (pictured in Figure 12 on page 22) that completes the de-coupling of a consuming application from its underlying data source, for much greater flexibility, and thus provides a unified view of all structured and unstructured data. This provides complete data transparency, where all data appears as one source regardless of how and where the data is actually stored, which means that applications will continue to work regardless of any changes in how data is actually stored. The approach is heterogeneous, and supports federated access to a wide range of diverse data sources of many types, including relational, other structured, XML, messages, and Web, etc. The IBM Information Server is also extensible, and a Wrapper Development Toolkit is provided so those customers may add support for non-standard data sources. The product is designed to be highly autonomous and non-disruptive to data sources, existing applications and systems.

The Server supports rapid deployment, enabling new information services (including WSDL, directory entries, and service artifacts in application server) to be deployed in minutes without any J2EE skills. It offers flexible bindings, so that services can be deployed as EJBs, JMS, or Web services, and as SCA in Q2 2006 to meet diverse project requirements. Metadata is linked into the common repository (see below), providing rich metadata visibility including lineage & impact analysis.

Support for the SOA foundation open-standards is naturally a vital aspect, and these include: Web services, SQL, JMS, EJB, SCA, II4BPEL, and CDC.

The server has also been architected for scalable high performance, with sophisticated optimization of distributed queries and advanced caching, and for fault tolerance for high availability.

This richly functional new server is a major advance, and we expect it to become available in late Q1 2006 for select multi-platforms, and on other major platforms such as z/OS from late Q3 2006.

- WebSphere Metadata Server:** Many important business initiatives like regulatory compliance, risk management and BI require organizations to define where information comes from, how it relates to other information, and when and how it is derived. This has raised the importance of metadata, the attributes that define the content, context, relationships, and structure of information. IBM's new metadata services infrastructure is designed to allow metadata to be more easily managed, accessed by those who need it, and shared across heterogeneous technologies through a service-oriented architecture.

To enable this capability, IBM announced the IBM WebSphere Metadata Server, which provides metadata management as a service to products within the IBM WebSphere Information Integration suite. IBM BI, MDM, and WebSphere portfolio products will also exploit the Metadata Server, in the future. The product is based on the open Eclipse Modeling Framework model. It provides the industry's first unified repository architecture for data integration that will provide dynamic, shared metadata across the complete WebSphere Information Integration suite. This means that metadata import/export between tools will no longer be required, because it will be made directly visible within the tools from the WebSphere Metadata Server. Part of IBM's previously announced "Project Hawk" initiative, this new product will be available, along with other IBM Information Integration products, in the spring of 2006. We assess this as a genuinely major advance that will make IBM the clear metadata leader.

Our Analysis

Information Integration is a central information service category that plays an ever-increasingly-important role in new SOA composite applications. IBM has built the industry's broadest information integration framework in WebSphere Information Integration, which combines the previously separate disciplines of EAI, ETL and EII into a single, cohesive and open framework.

IBM has built the industry's broadest information integration framework in WebSphere Information Integration, which combines the previously separate disciplines of EAI, ETL and EII...

The now integrated Ascential products provided a major advance in the information integration capabilities and products now offered and available from IBM today.

A number of these major products are supported on the System z9 mainframe. These provide information integration services well adapted to mainframe environments that are increasingly the hub of many large-scale SOA developments.

Further leading-edge developments and innovations are well-advanced to deliver another giant step in this framework, with the spring 2006 delivery of the ground-breaking IBM Information Server and WebSphere Metadata Server (*both assessed above*) that become the hub of the framework, and core SOA server foundations. These will be accompanied by further upgrades of several of the other existing WII products.

In our assessment, the IBM Information Server is a major industry "game changer", in the same way that enterprise application and process servers were when they were first introduced. It brings together, unifies, and integrates under one robust engine, all the primary information services functionality needed, and will hugely simplify infrastructure and operations.

IBM's large and sustained investments in Information Integration are now coming to fruition, with an outstanding line-up that can only further strengthen its already # 1 market share position.

10. Information Analysis & Discovery (IAD) Services

Introduction

Information analysis and discovery services are concerned with gathering, managing and analyzing enterprise data, to generate real information and new knowledge to allow better, informed insight and thus improved business decisions. Data warehouses/datamarts long provided merged, cleansed and repackaged operational and transactional data summaries in informational databases more appropriate for business intelligence (BI), decision support and analysis. These informational analysis processes are highly query-intensive and place substantial demands on the database systems supporting them and this on the server system's power/capacity. Numerous ETL processes must also be run, to update and refresh the data warehouses regularly from operational/transactional sources. These often-extensive ETL processes also require substantial batch processing and I/O computing power.

Faster Business Cycles, Soaring Data Volumes Drive-up BI Workloads

As today's faster business pace, IAD and BI also needs to speed up. Older weekly or monthly analysis cycles must be replaced by nearer to real-time, daily or even hourly, update cycles. Only these can allow rapidly changing customer, market and product demands to be identified and acted upon within remodeled business processes fast enough.

Soaring data volumes generated by e-business and the explosion of the Web, by the tens of millions of mobile and pervasive devices now in use, and by new tracking technologies such as RFID, are another big issue. These have multiplied transaction and data storage volumes several-fold in this decade alone, a trend we expect to continue. Such soaring transactional volumes directly multiply the volumes of ETL processing needed to create today's informational databases (*within data warehouses and datamarts*), which themselves grow correspondingly larger.

IAD Foundations and Business Intelligence Tools

In considering the enterprise software infrastructure required for this new AID/BI environment today, we make an important distinction between **IAD foundational infrastructure** capabilities, and the **BI delivery tools** that directly support end users. A close and symbiotic relationship exists between these two elements.

- **IAD foundational infrastructure:** This is IAD infrastructure information services software built with the scale, strength and robustness to underpin a successful enterprise IAD/BI strategy. Such an infrastructure includes:
 - Enterprise-class RDBMS, with advanced inbuilt support data management for high-performance BI databases.
 - Enterprise-class database query and reporting services tools closely coupled with the database management system.
 - Data mining and discovery services that can operate upon transactional and informational databases to discover/reveal new knowledge, patterns and relationships in the data.
 - Powerful, scalable ETL software services able to handle extraction, cleansing and transformation of large volumes of operational data, and its efficient loading into informational databases/ warehouses, are essential. In addition, robust and broad EII capabilities, to provide connectivity to diverse data sources needing to be integrated into data warehouses, are also vital.
 - Advanced analysis services providing specialized services to the delivery tools and to customer applications. OLAP and data cube generation facilities are two examples. The new field of entity analytics (*See Section 11*) is another.
 - Accelerators – offerings that make it faster/easier for a customer to deploy this IAD infrastructure – are invaluable. They include best practices, methodologies, industry-specific information warehouse models, and complete industry solutions.

IAD infrastructure suites are mainly offered by leading enterprise RDBMS software providers, notably clear leader IBM around DB2, Microsoft around SQL Server, and Oracle around the Oracle database. The investments required developing and supporting such RDBMS/BI foundation capabilities ensure this concentration will continue.

- **BI delivery tools:** These software tools/systems act upon (*and run on top of*) the IAD infrastructures above. They provide higher-level analysis, reporting, query, dashboard displays, business scorecards, plus other types of BI output. They also deliver/distribute these results to the groups of end users, ranging from executives, to managers, to knowledge workers and professionals, that are involved. Many terms have been used to categorize/describe such tools, including BI, decision support, executive information systems, enterprise reporting systems, 4GLs, and others. Today, a handful of third-party software vendors dominate this space, with varyingly comprehensive enterprise BI delivery tools/suites. SAS, Cognos, Business Objects, Information Builders and MicroStrategy are leading examples. Their suites integrate with the main relational database platforms through SQL and ODBC industry standards.

IBM IAD Foundation Infrastructure Leader

World IT provider leader IBM also led in IAD foundation software for nearly two decades. IBM invented both the relational database and the data warehouse concepts which underpin this domain. It has helped thousands of customers generate, manage and exploit enterprise and departmental data warehouses since the early 1980s. IBM combines its industry-leading server/storage hardware with its wide portfolio of information management software (*described in this White Paper*) from database software like DB2, to information integration, through MDM, and with advanced analytics services. These provide a strong IAD foundation for customer BI solutions ranging from department-level datamarts to massive enterprise data warehouses. IBM also delivers several complete industry-specific solutions, such as its Retail Data Warehouse, Banking Data Warehouse, Crime Information Warehouse and Telco Data Warehouse, all based on its BI foundation technologies. These can be implemented via its industry-knowledgeable service professionals. Our broad view of the whole IBM IAD foundation offering is shown in Figure 26 on page 56.

Informational Computing and the Mainframe – Full Circle?

From 1980 to the mid-1990s, a large majority of informational computing (*which grew rapidly over this period*) ran on the mainframe, then the dominant enterprise IT platform and thus the host for most operational/transaction data. The terms used at the time included end-user computing, 4GL, and Information Center. These became a significant part (*often 30-40%*) of total mainframe workloads by the end of this period. They supported thousands of new end users in the larger customer sites. Leading popular BI delivery tools on the mainframe then included FOCUS, SAS, RAMIS, NOMAD and IBM's AS. Most mainframe sites ran one or more such products (*as many still do with successors today*). These were mainframe, host-based, server-centric, interactive (*interpretive*) products, also able to run in batch mode. The mainframe proved adept at supporting these large-scale informational workloads, particularly after DB2 reached industrial strength. The trade-offs between directly accessing operational data, or the alternative use of staged information/data warehouse, were extensively debated. The data warehouse approach (*first proposed /championed by IBM*) emerged as most satisfactory.

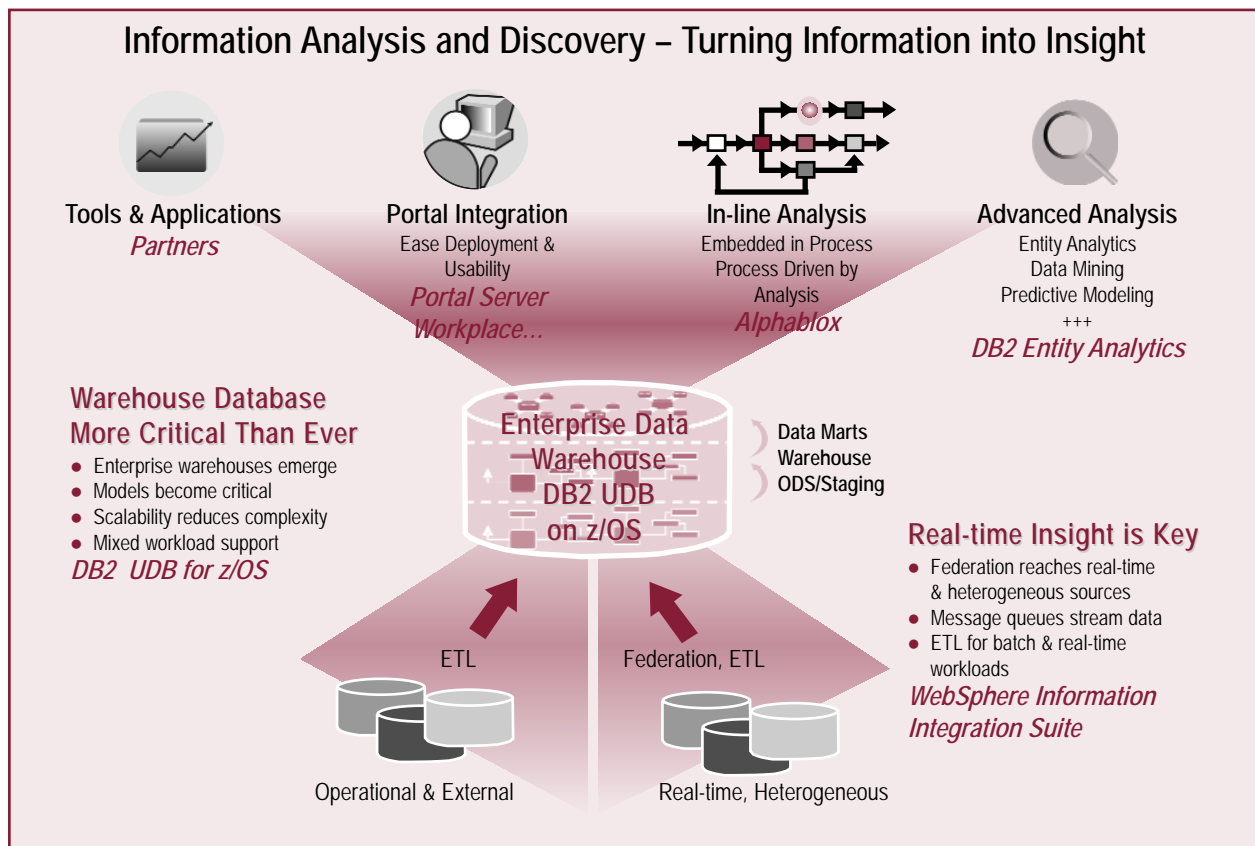


Figure 26: Information Analysis and Discovery – Turning Information into Insight

Only two challenges inhibited these big informational workloads on mainframes. One was a performance/capacity use challenge: these high query/report workloads, with their supporting substantial ETL loads, running on interpretative software tools, consumed much mainframe horsepower/capacity. On the single-image, few CPU-multiprocessor mainframes of the time, these demands sometimes conflicted with stringent OLTP workload service level needs. Running several mainframes with loads segregated was one answer (*"Information Center"*). The other challenge was these informational loads grew much faster than normal mainframe OLTP once business users discovered how easily they could unlock access to mainframe data. High BI workload growth drove up MIPS use/storage capacity demands rather faster than the steady OLTP growth. At mainframe costs, with the software licensing models of the time, platform costs could rise faster than expected.

When UNIX systems arrived, some businesses began running small-medium BI workloads on these distributed systems. The relational database became available on UNIX to support this. Rather later, Windows-Intel platforms became powerful enough to run small/medium BI workloads on Microsoft software (*SQL Server*) as well. This was fuelled by rapid advances in UNIX (*and Intel*) performance, capacity, and price/performance and by the crisis of confidence the mainframe underwent in the mid-1990s. By the late 1990s, conventional wisdom widely favored UNIX as the preferred platform for BI, mostly for cost reasons.

Today everything is different, and this conventional wisdom is now outdated, in our view. The mainframe has been completely transformed, has over 100 times the scale and capacity today, its costs have tumbled more than 70-fold in twelve years, and virtualization and partitioning enables one system to easily support multiple mixed workloads (*including transactional and informational*) safely, securely and to assured service levels. DB2 has gone from strength-to-strength, added many IAD/BI supportive features and capabilities, and, with the latest System z9 mainframe under z/OS, is undoubtedly the premier enterprise data-serving platform by all measures. Today's mainframes have a full 64-bit architecture, under a balanced design providing excellent large-scale informational/query processing (*based on their channel architecture, dedicated I/O processors, support for high-end, high-capacity, high-performance storage systems, automated tape, robust database management software, rock-solid security, recoverability, and resilience, etc.*). Moreover, they remain as pre-eminent as ever for the most demanding OLTP workloads that are becoming enterprise SOA applications today. Their automation, low staffing needs, more favorable software licensing, and much lower hardware costs, now typically make mainframe solutions much less expensive in TCU (*for larger workloads*) than distributed alternatives. The latter need multiple servers and armies of staff to support the same loads a single mainframe can easily deliver from its up to 17,000 MIPS power.

We therefore consider the mainframe has now come full circle, to reclaim its place again as often the best platform to host large-scale enterprise IAD/BI databases and solutions. Another key driver for this trend relates to the changing architecture of BI delivery software, which has also gone through a complete 360-degree circle (*see below*).

BI Delivery Tools Also Go Full Circle, Back to Host Server-centric Thin-client Model

As informational workloads originally grew on the mainframe during the 1980s, all the BI delivery software tools were host-centric, server-based products. These ran in interactive and batch modes entirely on the mainframe host, delivering to the thinnest of client devices, the 3270 terminal (*or light PC emulator thereof*) or the printer.

Through the 1990s distributed computing era of UNIX and Windows growth, the client/server-computing model (*of fat PC local client running heavy local functionality, supported by a server back end*) was the dominant BI delivery tool architectural model. Most BI leaders over this period used this client/server model, but customers found these solutions costly to distribute and support, and lacking the ability to scale well for larger user populations.

With the explosion of the Internet, by the start of this decade, the predominant computing architecture has become the thin Web browser client, communicating with a fat, host-based server that encompasses all the functionality of the application. BI delivery software has morphed again, and today leading BI vendor solutions such as Information Builders' WebFOCUS, SAS, Cognos and Business Objects, have all adopted this model. It slashes PC client costs and complexity, and scales much better to support larger user communities more easily. This BI delivery software architecture is also ideally suited to the mainframe, the most scalable of enterprise server platforms that is supported by the strongest data server in DB2 on z/OS.

Near Real-time Analytics and Data Mining – Newer IAD Workloads

On Demand business runs in faster cycle times; it thus needs nearer to real-time analytics to run at now-shorter business process cycle time, not at the traditional weekly/monthly batch summary/aggregation cycles of earlier BI. We consider this nearer to real-time IAD/BI analytic processing can run best close to the main SOA applications and their live data, increasingly centered on mainframe hubs. ETL workloads have also grown with data volumes to become substantial compute and I/O intensive jobs that also run well on powerful mainframes, by exploiting their balanced, high CPU and I/O capacities and parallel processing capabilities. More of these ELT processes are also being run more often, to support and feed near-real-time IAD/BI, which is easily handled fast on mainframes.

Data and compute-intensive data-mining applications are being used to detect patterns and correlations in the data for a gamut of other sophisticated uses. These also require close proximity to enterprise databases and data warehouses, and need considerable horsepower that the mainframe amply provides.

IBM's IAD Foundation Infrastructure Offerings

IBM's IAD foundation infrastructure software for the z Series/System z9 platform today includes:

- **DB2 UDB For z/OS Version 8:** Fully assessed in Section 5, and the premier enterprise data-serving platform, with extensive in-built support for BI and data warehousing. To be followed later by DB2 UDB for z/OS Version 9? (*Viper*) with further advances.
- **DB2 Tools:** Extensive set of supporting DB2 tools, a number contributing to BI/data warehousing solutions. (*Assessed elsewhere in the "Related Software Strategies Research" Section on page 72, items 6 and 9.*)
- **DB2 Data Warehouse Manager for z/OS Version 8.2:** DB2 Warehouse Manager for z/OS delivers a complete warehouse solution package for use with DB2 for z/OS. It combines tools to build, manage, and access DB2 data warehouses. It is based on proven technologies enhanced to offer additional function and to provide tighter integration between components, and includes:
 - DB2 Data Warehouse Center (*a GUI that integrates with the DB2 Control Center*).
 - Warehouse agent for z/OS (*executes for the Data Warehouse Center and enables point-to-point data movement*).
 - Warehouse transformers (*stored procedures/user-defined functions providing commonly needed DW-building transformations*).
 - The Information Catalog (*helps end users find, understand, and access relevant DW information*).
 - DB2 UDB V8.2 Enterprise Server Edition (*restricted use license*).

The current product release became available 12.03.2004.

- **Query Management Facility (QMF) Version 8.1:** QMF is the long-standing, and almost universally used, DB2 query and reporting tool suite that is a core complement to the database, and BI foundation. QMF has been much refined and extended to support BI tasks across the enterprise, and extends outwards to support varied individual needs. It supports drag-and-drop queries and analytics, provides executive dashboards (*with QMF Visionary Studio*), and supports DB2 V8 long names, etc. QMF is also enabled for WebSphere Application Server (*DB2 QMF for WebSphere*), provides extensive data visualization, and offers a database explorer for easy database browsing. QMF also boasts excellent performance and reliability for accessing relational databases, with native support for advanced features of current DB2 releases. On z/OS for DB2, three DB2 QMF Editions are offered – Enterprise, Distributed and Classic – with differing components. A Distributed Edition for Multiplatforms is also available.
- **DB2 Intelligent Miner for Data for z/OS Version 8.1:** This product enables knowledge workers to identify/extract high-value BI from enterprise data, and to communicate results to business users for analysis and decision making. It provides integrated, sophisticated analysis of large amounts of enterprise data, using proven data mining technology with best-of-breed scalability, and good usability/productivity for data mining analysts. It enables customers to gain new business insights and to harvest valuable BI from their enterprise data by mining even high-volume transaction data generated by point-of-sale, ATM, credit card, call center, or e-commerce activities. Typical uses include identifying market segment clusters, discovering product sales associations, identifying sequential patterns of buying or service call behavior, developing distinct classifications of customers, and to develop predictive scoring algorithms.

DB2 Intelligent Miner for Data Version 8.1, the current release, offers new analysis, data exploration, and scalability options that enhance core data mining strengths. It also brings data mining value to more BI users by widening access to mining functions/results at the analyst's desktop. Enhancements include integration with warehouse and BI tools, and improved data mining results communication. The release fully exploits DB2 Version 8.1 advances such as its index features and extended column names, and enables mining data from, and applying model results back to, DB2 UDB Version 8 databases. DB2 Intelligent Miner for Data is a client/server product that mines data directly from DB2 hosts (*as well as from other data sources via DB2 DataJoiner*). Data and metadata can also be exchanged with warehouses and BI tools and applications. The product also has an API, enabling development of customized, industry-specific mining applications by customers, IBM IGS consultants, and IBM Business Partners. Supported clients are Windows 2000/NT and AIX. (*Other supported server platforms are AIX, HP-UX, Linux, Solaris and Windows.*)

- **DB2 Entity Analytics:** New category of identity resolution IM tools – fully assessed in Section 11 following.
- **WebSphere Information Integration Suite:** Covered fully in Section 9, provides extensive and sophisticated EII and ETL information management services and capabilities, including z/OS implementations, that support data warehouse staging and loading, and enterprise information data integration.

These are a broad set of IAD foundation software capabilities on the System z9-z/OS platform. They enable it to offer premier enterprise IAD services and support large-scale BI. Many other advanced DB2 BI capabilities are also offered on other platforms, but are therefore outside the scope of this White Paper. (*Including advances in DB2 Data Warehouse Edition, DB2 Cube Views, DB2 Alphablox, and DB2 BCU, etc.*)

Industry Business Intelligence Solutions

Building on the above IAD infrastructure foundation software, and its extensive vertical industry business knowledge, IBM has now assembled a set of industry-specific, strategic BI solutions. These accelerators help customers in these sectors to implement major BI solutions to common industry needs much faster, and include:

- Banking Data Warehouse (*BDW*).
- Basel II Risk and Compliance Solution.
- Crime Information Warehouse (*CIW*).
- Information Framework Critical Business Process Models.
- Information Framework Object Modules.
- Insurance Information Warehouse (*IIW*).
- Retail Business Intelligence.
- Telecom Data Warehouse (*TDW*).

The associated IGS industry and technology consulting teams provide deep business technical skills and infrastructure experience expertise for the successful implementations of these DB2 Data Warehouse solutions.

Our Analysis

IAD/BI is a vital information service class that gathers, manages and analyzes enterprise data. It transforms raw operational data to cleaner, broader, more appropriately structured informational databases within data warehouses or datamarts. Query, reporting, advanced analytic, and data mining tools are then used off these informational databases to generate real information, new knowledge, and the informed insight that enables better business decisions.

The number of data source multiplied, as increasingly complex enterprise applications portfolios become the norm. IAS software infrastructure needs to combine these multiple data sources for meaningful analysis. Today's faster business pace now demands faster, nearer real-time IAD/BI cycle-times to match. IAD/BI infrastructure must also cope with exploding data volumes generated by e-business, the Web, mobiles, and RFID, etc. An enterprise-scale, robust, and powerful BI software/systems infrastructure, plus a powerful analysis tool portfolio, are essential to meet these tough needs.

IAD foundational infrastructure demands information services software with the scale, strength and robustness to provide the underpinnings of an effective enterprise IAD/BI strategy. IBM is the leader, and other major DBMS vendors the competition, for these IAD platforms. BI delivery tools are the information services software tools that provide analysis, reporting, query, dashboards, and scorecards, etc., that deliver BI results to all enterprise users groups. Specialist BI ISVs provide these BI delivery suites.

IBM invented both the RDBMS and data warehousing, and has been the leader in IAD foundational infrastructure for 20 years. It offers industry-leading server/storage hardware, leadership DB2 database software, information integration, MDM and advanced analytic service software. These, plus industry solutions, IBM services, and IBM partner BI delivery tools optimized for its IAD foundations, provide complete IBM IAD/BI solutions of all levels.

Through the 1980s, the mainframe was the dominant platform for informational computing, and such use grew dramatically. In the 1990s, much BI was added on UNIX (*then also on Windows/Intel*) platforms. Today, with the massive demands above, IAD/BI is returning full circle to the System z9 mainframe. These alone have the power, scale, capacity and robust and sophisticated IAD foundation software to drive the new enterprise-wide IAD now needed, and are fast becoming the main enterprise data hubs.

First-generation informational/BI delivery tools proliferated through the 1980s: these were interactive, mainframe server, host-based, thin-client tools delivering to dumb terminals. Through the 1990s new BI tools migrated to use the fat-client Windows PC-mainly, UNIX server client/server architectures in vogue at this time. These scaled very poorly, were expensive to buy, and costly to support. This decade, the Web changed everything, and today's leading BI tools have returned full circle to a new server-centric, Web-browser, thin-client model, ideal for powerful enterprise server platforms such as the mainframe. (*And similar to the 1980s model in today's Web environment.*)

IBM's robust, proven IAD foundation information services for the System z9 mainframe are delivered by the DB2 UDB for z/OS database, the QMF query/reporting software, the DB2 Data Warehouse Manager for z/OS, and the DB2 Intelligent Miner for z/OS data mining. They also include the newer, advanced DB2 Entity Analytics suite for identity resolution on z/OS DB2 data (*see Section 17*). These are supported, for information integration, by the extensive WebSphere Information Integration suite on z/OS and other platforms (*see Section 9*). These capabilities deliver IAD as services that can be embedded in SOA applications and reused across the enterprise, under IBM's IAAS strategy, and support open standards strongly. Built atop these strong product foundations, IBM also offers a portfolio of pre-built vertical industry BI solutions, backed by experienced services, for fast time-to-value in those niches.

The market resurgence of the mainframe, with the power, capacity and new economics of the new System z9 systems are well supported by the above IAD foundation software capabilities. These, aided by the introduction of the new zIIP z9 specialty processor (*that offloads important DB2 BI workloads*), have now positioned the mainframe to recapture some of its earlier dominance. It provides an ideal hub for large-scale IAD/ BI, close to the transactional systems that also run on the platform, that is secure, reliable, gives excellent performance, and now comes at more affordable costs.

11. Entity Analytic Solutions (EAS) – New IM Services for Identity Resolution

Introduction

The seemingly simple questions of “Who is Who?” and “Who Knows Who?” occur in many business and government problems today. IBM's DB2 Entity Analytic Solutions (EAS) portfolio provides solutions to these age-old problems with levels of speed and accuracy not previously possible. These DB2 EAS solutions strengthen IBM's information management portfolio by providing a multidimensional approach to analytics that extends the capabilities of identity-based applications. The combination provides value to customers, and to business partners who deliver BI and other applications, that might require a single customer view, fraud detection, or customer relationship management, across many industries such as government, banking, insurance and healthcare.

Addresses a Panoply of Common and Challenging Identity-related Issues

At a time when being able to accurately recognize a person's real identity has never been as important as it is today, Entity Analytics now provides a means to enhance and extend knowledge discovery in a diverse and important set of real-world problem areas that include:

- Anti-Money Laundering (AML).
- Border safety.
- Company insider threats.
- Employee screening.
- Multi-party data sharing.
- Customer loyalty.
- Clinical research.
- Utilities fraud.
- Vendor screening.
- Fraudulent slip and fall.
- International data exchange.
- Patriot Act compliance.
- Customer relationship management.
- Reputation risk.
- Executive self-payouts.
- OFAC compliance.
- Benefit management.
- Homeland security.
- Inventory shrinkage.
- Credit card bust out.

As can be seen above, the applications of this technology are wide-ranging, sensitive and all carry high stakes.

An Open EAS Approach

The DB2 EAS technology complements IBM's open computing strategy through its ability to work with both IBM and non-IBM data sources, including Oracle and Microsoft databases. This enables customers to integrate these solutions with their existing technology infrastructures. The technology is based on both SQL and the Java industry-standard J2EE, allowing developers to more easily develop, deploy and maintain analytical business applications needing identity resolution.

DB2 Entity Analytic Solutions – Unique Capabilities

IBM DB2 Entity Analytic Solutions (EAS) bring a higher accuracy, precision and fidelity to identity recognition. Enhancing privacy by answering both of these questions anonymously is the latest capability of the DB2 EAS portfolio. This advance encompasses non-obvious relationships and anonymous data sharing, dramatically extending the capabilities of identity-based applications, and offering a level of due diligence and security not possible before.

Other solutions do not resolve individual identities to the depth and degree of accuracy required to conclusively determine who is really who. Nor do they manage identities in real-time as new information is received, or include the impact of both obvious and non-obvious relationships.

IBM DB2 Entity Analytic Products

An overview of the DB2 EAS portfolio products is shown in Figure 27 (on page 61), and brief descriptions of each follow below.

- **DB2 Identity Resolution:** Identifies “Who is Who?” This product can resolve inconsistent, ambiguous identity and attribute information into a single resolved entity across multiple data sets. The product can overcome deliberate attempts to confuse or misrepresent individual identity. Version 3.9.1 is the current release, and was made available on 10.28.2005. Supported server platforms include AIX5L, Linux (32- and 64-bit), Solaris, Windows and HP-UX. Databases supported include IBM DB2 for z/OS V8, IBM DB2 UDB 8.2, Oracle 9R2, Oracle 10g and MS SQL Server 2000.

- **DB2 Relationship Resolution:** Recognizes “Who Knows Who?” This product extends the value of existing data assets by linking unique resolved identities to outside obvious and non-obvious entities, to establish relationships. These may be to individuals or organizations, and can uncover social, professional, value, and criminal networks. Version 3.9.1 is the current release, and was made available on 10.28.2005. Platforms and databases supported are as above.
- **DB2 Anonymous Resolution:** Who is Who and Who Knows Who... Anonymously? This, the newest IBM EAS product, enables multiple organizations to selectively share data and leverage proprietary data in a manner that never exposes sensitive information, while still identifying relationships and developing leads. This anonymization allows two or more organizations to compare and correlate their proprietary data, without exposing the organizations’ sensitive data values, to identify matches, as well as obvious and non-obvious relationships on demand. Version 3.8.1 is the current release, and was made available on 06.14.2005. (Fully supported server platforms are Windows, AIX, Solaris, with database support for DB2, DB2 and Oracle respectively, and limited support platforms are HP-UX, SGI, Linux and Windows Server.)

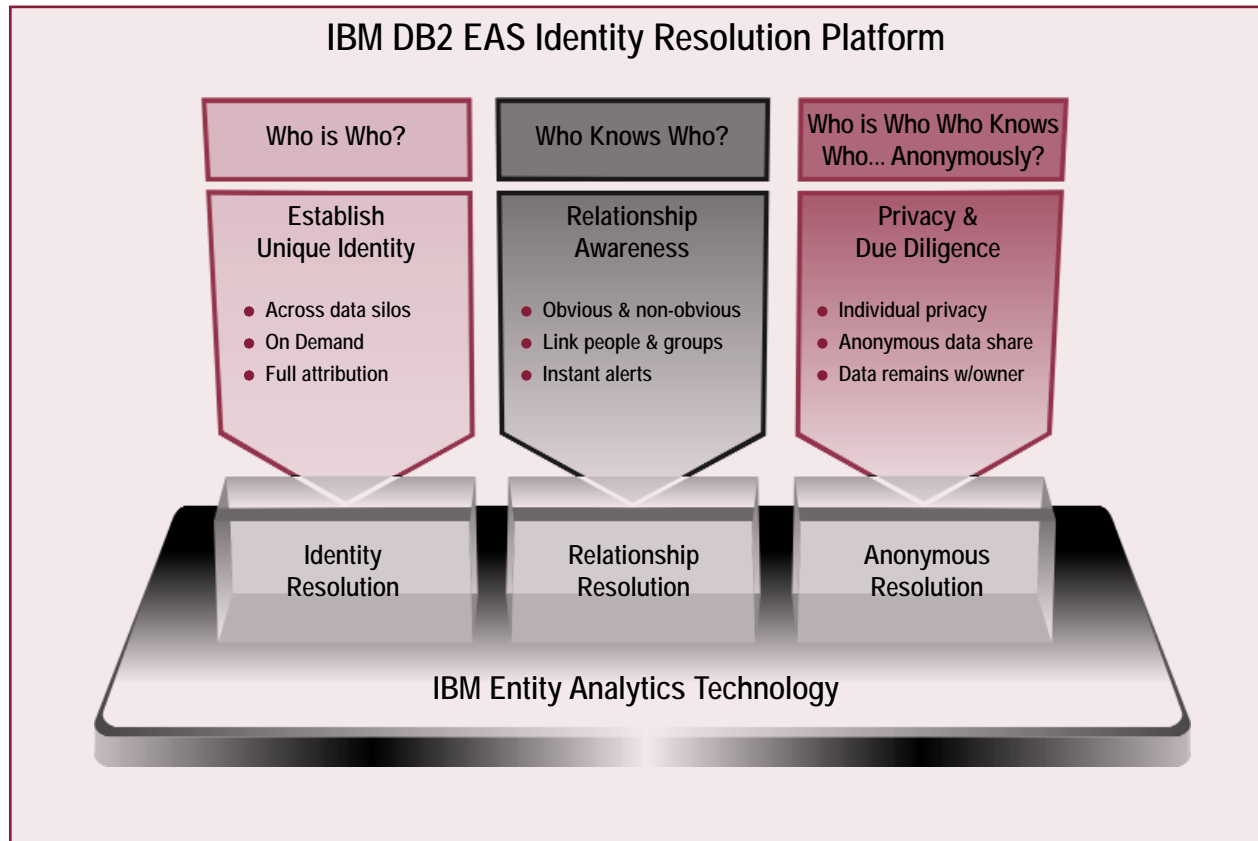


Figure 27: IBM DB2 EAS Identity Resolution Platform

Our Analysis

Entity Analytics are a relatively new technology for resolving crucial identity challenges that are now widespread in intelligence, security, fraud prevention, crime fighting, and many other high-risk, high-return areas highlighted above. The DB2 Entity Analytics portfolio provides an innovative set of identity resolution information services that add value to many types of applications and solutions addressing the challenges above.

With DB2 Identity Resolution and DB2 Relationship Resolution now providing database support for DB2 for z/OS Version 8, these powerful identity resolution technology tools can now be brought to bear directly on major mainframe-resident databases and applications under a SOA.

These DB2 EAS technology advances are amongst the numerous investments IBM made to extend its analytical information services solutions, which include DB2 Data Warehouse Edition, DB2 Cube Views and DB2 Alphablox. These underpin the firm's growing range of industry-specific BI solutions, such as the IBM Banking Data Warehouse for financial institutions and the IBM Crime Information Warehouse for law enforcement that IBM now offers.

Appendix A: SOA & the zSeries/System z9 Mainframe – A Marriage Made in Heaven

Our White Paper “SOA Takes Off – New WebSphere SOA Foundation Extends IBM’s Lead with New System z9 Mainframes as the Hub of the Enterprise” (*Software Strategies White Paper, September 2005*) deeply explored the case for adoption of Service Oriented Architecture. It also examined the central role the zSeries/System z9 mainframe should play in this new-generation of enterprise applications technology. Our main findings from this research are summarized in its Executive Summary, and reproduced here:

1. **Collaborative Integration Vital:** Organizations worldwide are moving towards becoming On Demand Businesses. They want to become more responsive, flexible and adaptable to faster-changing, constantly evolving business conditions and markets. To do this, they are integrating internal operations, and linking with customers and partners externally in collaborative processes and applications.
2. **Building Applications Challenging** Building enterprise IT applications has long been difficult. Complex, distributed, heterogeneous IT environments compounded this. Recent faster business changes and challenges required the IT applications to change/respond faster, and be more secure and resilient, needs that IT could often not meet.
3. **Prior Application Development (AD) Advances Fail to Resolve:** Many AD advances emerged over prior decades. However, none until now decisively solved the challenges. There remained a wide gap between what business leaders needed and what IT could provide. A breakthrough was desperately needed.

SOA is now the hottest topic and the fastest growing market for the IT industry in 2005/2006.

4. **Service Oriented Architecture (SOA) the Breakthrough:** SOA is the transformative business software architecture finally able to make this breakthrough. Other top analysts, and we, consider SOA to be the most significant, applications advance ever to emerge. It offers drastic improvements in how business applications are built, supported, deployed and adapted to support core business processes and enable rapid change.
5. **SOA Market Momentum:** SOA adoption has taken off on a large scale and at a rapid pace, as many indicators attest. SOA is now the hottest topic and the fastest growing market for the IT industry in 2005/2006.
6. **IBM Mainframe Central SOA Platform:** The unique business values of zSeries, and of the new System z9 with its further advances, position the mainframe to serve as the enterprise hub platform for SOA and data. This places it at the epicenter of customers’ collaborative computing and SOA developments.
7. **New, Redefined Distributed Computing Style:** SOA is an important, newer distributed computing style. It allows loosely coupled, independent, reusable business software services (*interoperable & technology-neutral*) to be choreographed into Internet-based composite applications, to support transformed business process workflows.
8. **Services Reuse Heart of SOA:** Business software services are reusable, so new applications can be built much faster. Services can be created from existing software assets and packaged applications, or be newly built in J2EE™ or .NET. Composite SOA applications are flexible, can easily be adapted to meet business changes, and make enterprises more responsive to new market opportunities and competitive threats, etc.
9. **Web Services, Standards Enablers:** SOA is not only based on, but extends, Web services technology, and exploits the now broad, maturing set of open Web services standards. These allow SOA applications to integrate and inter-operate across all hardware, OS and middleware platforms via the Web.
10. **Drastically Simplifies Business/System Integration:** Monolithic enterprise applications were hard and costly to integrate or change, and supporting both absorbed most IT capacity. For the reasons above, SOA applications are dramatically easier to integrate and change, freeing up scarce IT skills for new, high-business-value applications.
11. **Brings Tight Business-IT Alignment:** SOA extends the valuable advances of Web services and now allows tight alignment between business and the IT team to be achieved – resolving this long-standing and major issue.
12. **SOA Comes of Age – Start Now/Move Faster:** With universal and compelling SOA benefits, and maturing SOA solution offerings, we advise enterprises to now begin this important journey (*if not already begun*) or to redouble and accelerate their efforts (*if underway*).
13. **Mainframe Investment Brought Resurgence:** IBM invested over \$5B this decade to make today’s mainframe the most advanced, productive and cost-effective enterprise computing system for On Demand Business, we found in our recent System z9 study. Consequently, the platform has enjoyed a healthy resurgence in market share and esteem.

14. Compelling Case for SOA on Mainframe: Our research found compelling reasons why the mainframe should play a central role as the platform and the main node for enterprise customer SOA-enabled infrastructures, deployments of which are accelerating:

Our research found compelling reasons why the mainframe should play a central role as the platform and the main node for enterprise customer SOA-enabled infrastructures...

- A wealth of traditional enterprise application software and data assets on mainframe are easily converted to business services for SOA.
- Five years of intense deployment of new workloads on the mainframe means a large, modern software portfolio is now also running on the platform, which can be readily adapted to plug-and-play in enterprise SOA architectures.
- Unique mainframe business values are highly supportive of enterprise SOA business composite application needs, not easily (*or not at all*) delivered by other platforms.
- Advanced mainframe capabilities, such as extreme virtualization and intelligent self-optimizing workload management, underpin these values, are highly complementary to enterprise SOA composite business applications, and can often best provide the performance, behavior and characteristics needed.

15. Heterogeneous SOA infrastructures need central co-ordination of vital SOA-wide functions, and the mainframe now has well-advanced, specific capabilities and roadmaps to fulfill this enterprise SOA and data hub role well.

16. Approach SOA Business-focused: Begin SOA with a sharp focus on business strategies, with business process transformations prioritized, and with full top executive support. The business benefits of successful SOA deployment are most compelling, although the associated IT benefits are also valuable. Experienced industry business consultants can help uncover and rank the processes most in need of treatment, if this is not already clear.

17. Do a SOA Readiness Assessment: To clearly identify where your enterprise is positioned on the SOA capability-maturity curve, such an assessment is a highly-effective, rapid means of pinpointing the areas to be enhanced for fuller SOA adoption. Self-service SOA tools, and/or workshop services for this are available from some vendors.

18. Understand SOA Lifecycle/Roles: Understanding SOA applications lifecycles, and the business and IT roles involved, will help match with existing resources, and will pinpoint capability/capacity gaps that need to be addressed. Joint business-IT "SOA Program Governance" is recommended.

19. Identify New SOA Resources Needed: From the gaps identified above, the additional SOA resources (*staff and/or skills, SOA middleware software, SOA tools, education, knowledge and information, methods and best practices, business and SOA technology services, etc.*) needed to realize the enterprise SOA plan can be identified.

20. SOA Vendor Selection: A primary choice is between the few major vendors with more complete SOA "platforms", and offerings from small, "pure-play", newer vendors. Analysts expect the former to quickly dominate in the SOA market, because of the high investment needed and vendor risk perceptions. Wide differences in offering and platform strength exist between the major vendors, typical of an early market stage.

21. WebSphere SOA Foundation Software Centerpiece: An impressive new WebSphere SOA Foundation software suite stands at the heart of IBM's SOA offerings, comprising these new or enhanced products:

- WebSphere Business Modeler.
- WebSphere ESB.
- WebSphere Integration Developer.
- WebSphere Message Broker (*Advanced ESB*).
- WebSphere Process Server.
- WebSphere Business Monitoring.

We found each tool was of class-leadership level, represented major advances on their predecessors, and made this market's premier enterprise SOA software suite. All these core SOA engines and tools will be available of System z9-z/OS by spring 2006.

22. Mainframe Platform – SOA Support: Enhanced, SOA-supportive, major middleware engine releases, already available for z/OS, underpin and support the new WebSphere SOA Foundation (*available by Q1 2006 on zSeries/System z9*). These position the mainframe as a premier SOA host and hub.

23. IBM the Only SOA Choice for Mainframe Users: We found that where enterprises are substantially using IBM mainframes (*or mainframe plus diverse multi-platform environment*), the IBM SOA offering is not only the strongest overall, but is the only realistic SOA solution to adopt on the mainframe side.

24. **IBM SOA Leadership Acknowledged:** IBM has created a formidable strength of broad offerings, and a clear leadership position in SOA, both thought and market, as confirmed by recent analyst surveys, holding three times the SOA leadership recognition of the nearest competitor.
25. **IBM SOA Services, Support Unrivaled:** We found the services breadth, depth, scale and SOA experience of IBM Global Services (*IGS*) to far exceed any other SOA services vendors. Combined with the leadership WebSphere SOA Foundation software, a wealth of IBM SOA education, knowledge/information, tools/methods, and a fast growing SOA partner ecosystem in place, these strengths make IBM the natural partner of choice for enterprise SOA deployments.

Appendix B: The New System z9 109 Mainframe Generation – Analyst Overview

We recently published an in-depth, Enterprise Server Spotlight Report (see the “Related Software Strategies Research” Section on page 72, item 2) assessing/reviewing the new System z9 109 mainframe hardware, software and storage, highlighting their significance and value to enterprise IT users and vendors worldwide. This Appendix reproduces our Executive Summary findings, included to provide a concise picture of the new mainframe platform here:

1. **Spectacular New System z9 109 Mainframes – Designed for Collaborative Computing Era:** The new System z9 109 mainframes (shown in Figure B1) are positioned to support **collaborative computing**, which IBM sees as the next generation of enterprise computing. Enterprises today must collaborate with their partners more efficiently, with global collaborative processing tightly integrating people, processes, technology and data to their traditional transaction processing under SOA.

This calls for systems that can be more easily, more rapidly, and more intelligently linked to securely support On Demand Business. IBM claims that the new System z9 109 mainframes lead in collaborative computing technologies like virtualization, open standards and encryption that uniquely equip these systems to better share information and processes within and across collaborating businesses.



Figure B1: The New IBM System z9 109 High-end Mainframe

2. **System z9 Exemplifies New “Systems Agenda” Strategy:** IBM also unveiled a new “Systems Agenda” systems strategy. Its main thrusts are to “Virtualize Everything”, “Commit to Open”, and to “Collaborate to Innovate”, and to extend IBM’s deep strengths in these areas to better support the firm’s On Demand Business vision launched in 2002. The new IBM System z9 109-mainframe advances exemplify this new strategy.

3. **System z9 109 – New Enterprise Hub Roles:** The System z9 109 was designed/optimized to serve as the **enterprise hub** in this new **collaborative computing** era, offering new levels of security, enhanced reliability, extensions to its leadership virtualization capabilities, and much-increased scale/capacity. This vision defines four new mainframe **enterprise hub roles** in collaborative computing infrastructures, as well as its traditional, roles. These are:

The System z9 109 was designed/optimized to serve as the enterprise hub...

- Enterprise Security Manager.
- Enterprise Business Resilience Manager.
- Enterprise Workload Manager.
- Enterprise Hub for Data and Service Oriented Architecture (SOA).

With the IBM System z9 109, IBM delivers significant steps along the charted development roadmap of this “Hub” strategy, announced in October 2004, as well as extending support for traditional mainframe roles in numerous other advances.

4. **\$1.2B IBM Investment:** The System z9 109 took three years, cost \$1.2 billion, and used 5,000 global IBM developers and engineers to develop. This continued the multi-\$B mainframe investments IBM has consistently made since the late 1990s that have fuelled the market resurgence seen in recent years.

5. **New Branding:** Since October 2000, all the IBM server lines (*including the zSeries mainframe*) were branded with the “eServer” prefix. The brand evoked the sharing of IBM’s best technologies across its (*then four*) server lines, and supported its corporate “e-business” (*later “e-business on demand”*) strategies/messages. With the System z9 109, “eServer” is dropped, and the “System” brand returns (*a.k.a. System 360, System 38*). The next generation pSeries, iSeries and xSeries can be expected to follow suite. “eServer” worked well for five years, but became slightly discordant with the “On Demand Business” corporate message, so retirement was appropriate. IBM Systems’ marketers place much emphasis on branding, although analysts and customers are less moved.

Impressive technologies double the high-end capacity of the System z9 109 to 17,800 MIPS in the new top-end 54-way system.

6. **Doubles Capacity, Scalability:** Impressive technologies double the high-end capacity of the System z9 109 to 17,800 MIPS in the new top-end 54-way system. (*Up from 32-way on the z990*). The System z9 109 also supports doubled the memory at 512GB, and increases system bandwidth by 80%, in a balanced design that enables it to support much larger collaborative processing workloads. (*All compared to the z990*). IBM claims a new System z9 109 (*54-way*) can securely process 1B transactions per day, and that a maximum 32-system Parallel Sysplex cluster of System z9 109s can handle 25B transactions per day, ample for the largest business on the planet.

7. **Five-model Range:** The System z9 109 offers a five-model range – the Models S08, S18, S28, S38 and S54 – offering low entry points and fine-grained increments for exact system sizing. All five models shipped from September 16 2005, the 54-way S54 two months earlier than previously announced. Comprehensive upgrade paths, both within the new range and from most z990 and z900 systems, are offered as usual with IBM’s mainframes.

8. **More Powerful, Dual-core Processors, and Flexible Designation:** For the System z9 109, IBM built a new, near 600-MIPS Processor Unit (*PU*) (*35% up on the z990 PU*). Each PU may be designated and used as a CP (*general-purpose workload processor*), a System Assist Processor (*SAP – I/O processor*), an IFL (*specialty Linux processor*), an ICF (*specialty Sysplex coupling processor*), a zAAP (*specialty Java workload processor*), or as a spare processor. This allows customers considerable flexibility to configure an optimum System z9 109 system for their workload and budget. The IBM z9 processor chip is again a dual-core SOC that includes one or two active processor cores (*as on the z990*), L1 cache, and memory-bus interfaces. The z9 processor ships with a clock speed of 1.65GHz. (*Up from 1.2GHz. on the z990’s z8 PU*), and delivers 580-600 MIPS (*up from 450-480 MIPS on the z990, both depending on workload*).

9. **Now Just 2 Spare PUs Per System Needed:** In a welcome development, a minimum of only two PUs across the whole system need to be reserved for use as spares. Previously, the z990 required two spare PUs per book (*1-4 books per system*). This frees up 6-8 PUs for customer designation on higher-end System z9 109 models. A minimum of two SAP (*I/O*) processors per book must also be designated, to provide I/O processing horsepower.

- 10. Dense Physical Packaging:** The System z9 109 server occupies almost identical floor space, and uses similar electrical power, whilst providing twice the capacity/throughput of its z990 predecessor, an impressive packaging feat. A single two-frame chassis houses the entire system, with most of the space devoted to up to three I/O subsystem cages. The System z9 109 uses similar modular “book” construction to that of the z990, with one (*Model S08*) to four books (*Models S38 & S54*) per system. The z9-109 uses a dense, new, 3.5B transistor Multi-Chip Module (*MCM*) that mounts 16 chips in total, including eight dual-core z9 processor chips, and four Level 2 cache chips. Each book houses one MCM, and up to 128GB of shared main memory.
- 11. Secures Enterprise Data to Cut Risks With New Security Advances:** Substantial security and encryption advances will enable the proposed new “**Enterprise Security Manager**” hub role for the System z9 109, bringing valuable new protection options we consider customers will warmly welcome. These include securing data transported extra-enterprise; advanced new encryption; faster Secure Sockets Layer (*SSL*) transactions; easier Internet security implementation for mainframe workloads; and enhanced network-based security with Cisco. These, and the System z9 109’s centralized key management capabilities, add further strength to the mainframe platform’s already legendary security. This security capability will be of particular value to enterprises deploying broad SOA applications over multiple platforms.
- 12. Virtualization Virtuoso:** The System z9 109 extends mainframe virtuoso virtualization strengths, enabling customers to simplify their overly complex IT infrastructures. The System z9 109 now supports up to 60 LPARs (*double the 30 on the z990*), can support thousands of virtual servers on z/VM, enjoys/exploits the major IBM Virtualization Engine 2.0 platform advances, and adds important new storage virtualization developments/directions, extending its “virtualization fabric hub” leadership. For higher levels of SOA implementation, fullest virtualization of the IT infrastructure enables highest efficiency and automation, and the System z9 clearly holds industry-leading advantage here.
- 13. Extended Reliability and Availability:** System z9 109 extends already legendary zSeries mainframe reliability and availability with further significant advances, especially in further reducing planned outages. The System z9 109 offers new availability features, including specialty engine capacity back-up, enhanced book availability, higher memory availability, and more non-disruptive hardware driver maintenance. These all help drive down the number and duration of planned outages still further.
- 14. Outstanding Enterprise Workload Management:** Extension of the well-proven mainframe intelligent workload management strengths to other heterogeneous, IT platform-based applications is especially noteworthy, and addresses major user challenge. The IBM VE 2.0 Enterprise Workload Manager now enables the System z9 to optimize SOA application workload management and performance across heterogeneous environments.
- 15. Middleware Fully Supports Next-generation SOA Applications:** Collaborative computing demands new generation, SOA based, composite applications. For the System z9 109, in our assessment IBM now offers the industry’s most comprehensive, robust, and advanced SOA middleware and tools suite bar none. All the main z/OS middleware software engines, and their associated tools, had all been substantially upgraded though 2004 and 2005. These new servers included WebSphere Application Server for z/OS Version 6, WebSphere Portal for z/OS Version 5.1, WebSphere MQ for z/OS Version 6, and CICS Transaction Server Version 3.1/CICS Transaction Gateway Version 6, all of which have enhanced support for SOA, Web services, and J2EE™ 1.4. WebSphere Business Integration Server Foundation for z/OS Version 5.1 provides a business process server today. With the recent SOA Foundation announcement, WebSphere ESB, WebSphere Process Server, and a new version of WebSphere Message Broker (*Advanced ESB*) are also slated for Q1 2006 availability on z/OS. Recent DB2 for z/OS Version 8 and IMS Version 9 releases also offer enhanced and complementary database support. These engines provide full run-time support for deploying next-generation, collaborative computing applications under a SOA, implementing the **enterprise SOA node** role envisioned for the System z9.
- 16. SOA Tools Complement Platform:** The SOA Foundation announcement brought the much enhanced WebSphere Business Modeler process modeling tool for business analysts to design and model business processes. It also added, the new WebSphere Integration Developer for assembling SOA composite application, and the powerful WebSphere Business Monitoring tool, for monitoring live business process composite applications. When combined with the Rational Software Development Platform AD and team development tools for new software development, and the extensive suite of Tivoli management tools for z/OS, these provide a comprehensive suite of SOA development, monitoring and management tools for mainframe SOA solutions.
- 17. Double-digit Percentage Price/Performance Gains for Enhanced Value and Improved Economics:** The System z9 109 family continues the mainframe’s rapid price/performance advances, at double-digit percentage improvement per-year rates, for both hardware and software. Continuing previous aggressive moves, IBM has again improved price/performance substantially on many mainframe hardware and software cost items, to extend this sharp improvement in mainframe economics that had already made it the lowest TCU platform for enterprise-scale workloads.

- 18. Broad Operating Systems Support:** The System z9 109 supports z/OS, z/VM, VSE/ESA, z/VSE, TPF, z/TPF and Linux for System z9 (*31-bit and 64-bit distributions*), supporting ESA/390 and z/Architecture modes that ensures current applications will run unchanged. The flagship z/OS, in new Release 1.7, provides fullest support for the new systems and for their most advanced features.
- 19. Holistic Design/Extensive Collaboration:** IBM worked hard internally to design and optimize this sophisticated new mainframe system, its hardware, operating system, middleware and services, in a holistic approach to provide balanced, all-round performance, price/performance, security, reliability and manageability gains across all levels of its stack. It has also collaborated extensively with such select major partners as Cisco Systems, Network Appliance, enterprise management software vendors, and others. Under the Mainframe Charter, into which large investments continue to be made, extended support for and collaboration with ISVs has encouraged many new partners to offer their products for the mainframe software platforms.
- 20. Commitment to Open:** New open workloads drove 75% of new mainframe MIPS sold in the last four years. The success of Linux on the platform has been particularly noteworthy, with over 1,700 installations to date. The System z9 109 adds extended support for open-standards-based SOA application development, and for virtualization, and supports all other significant and established open industry standards, many IBM-driven. Today, the mainframe is now amongst the most open, accessible and interoperable of all-commercial systems.
- 21. Will Extend Mainframe Resurgence:** IBM shipped more mainframe MIPS since 2001 than in the previous 36 years, and gained 16% share points in the “over \$250,000 servers” segment (*in which it holds a commanding lead*), highlighting the scale of the mainframe resurgence. Sustained technology innovation, the core strengths and unique business values and QoS it offers, superior economics for large workloads, and a wholesale market rediscovery of current mainframe advantages, stand behind this resurgence. We assess the System z9 109 will accelerate further resurgence and growth, with its impressive portfolio of advances on all fronts.
- 22. Continues to Preserve Customers’ Mainframe Investments:** One of the most important strengths and advantages of the mainframe platform, carefully preserved and maintained by IBM, is the high level of customer investment protection that it offers. Upward hardware and software compatibility, high hardware residual values, stable core APIs, affordable hardware upgrade paths, multiple operating systems support, and compatibility for customers’ existing applications, are long term “givens”. Ensuring no disruptive discontinuities in the hardware and software roadmaps, together with highest-quality support and maintenance, are also established mainframe heritage strengths that are continued and extended with the System z9 109. These desirable attributes are difficult and costly for any system vendor to achieve. They are far from common in this industry, where even major vendors have often imposed costly and disruptive technology transitions on their customers that wrote off much of their users’ prior investments. Microsoft (*middleware software stack*) and Hewlett-Packard (*server lines consolidation 2003-2005*) are prime examples.
- 23. Coolest Computer System on the Planet:** IBM, in its usual conservative style, described the new mainframes as “One of the most sophisticated computing systems ever”. “Coolness” is just as an important a marketing and adoption factor in the IT industry as in fashion and popular culture. The “coolest” technologies in the industry today included dual core processors, SOC, virtualization, high-density, reliability, security, encryption, SOA, and several others. With its winning leadership on all of these technology “coolness” factors, the new System z9 109 can now justly be considered the “Coolest Computer System on the Planet”.
- 24. Next Generation of Mainframe Staff in Training:** Over 150 universities are now participating in the zSeries/System z9 Academic Initiative, and IBM/SHARE have pledged to add 20,000 new mainframe professionals to the global workforce by 2010. Each year thousands of new generations of young IT professionals are discovering the outstanding attributes and capabilities today’s mainframe offers with astonishment and interest. Many now graduate with these skills to join the established staff teams of mainframe users, who have long understood the mainframe’s unique advantages.
- 25. Attractive & Compelling Upgrade:** Our assessment found the System z9 109, first of a new generation of IBM mainframe systems, provides an attractive and compelling upgrade for both high-end z990 customers and mid-sized z900 and z990 users. With its massive increase in scale/capacity, superior virtualization, extended security, enterprise hub support, extensive SOA software stack, much-improved price/performance, and further enhanced availability, these new mainframes offer an attractive blend of benefits and advances. These advances, together with these system’s still lower per-user costs (*that run 3-4 times lower than UNIX or Windows systems on enterprise workloads*), render them ideal “enterprise hubs” for collaborative computing.

These advances, together with these system’s still lower per-user costs, render them ideal “enterprise hubs” for collaborative computing.

Appendix C: IBM WebSphere SOA Foundation Software – Overview

We recently reviewed the new IBM SOA Foundation middleware software suite in our White Paper entitled “SOA Takes Off – New WebSphere SOA Foundation Extends IBM’s Lead with New System z9 Mainframes as the Hub of the Enterprise” (*Software Strategies White Paper, September 2005*). Our summarized overview of the new WebSphere SOA Foundation software is reproduced below.

At the heart of the new IBM SOA Foundations is six new, or significantly enhanced, WebSphere SOA Foundation software products, shown and positioned on the SOA lifecycle in Figure C1.

- **WebSphere Business Modeler:** To model SOAs and business process, IBM offers a considerably enhanced new version of its WebSphere Business Modeler. This is an easy-to-use graphical modeling tool that enables both business and IT to work together to model and design the business process flow before deployment, and complements the SOA modeling capabilities of the Rational Software Architect tool for IT architects.

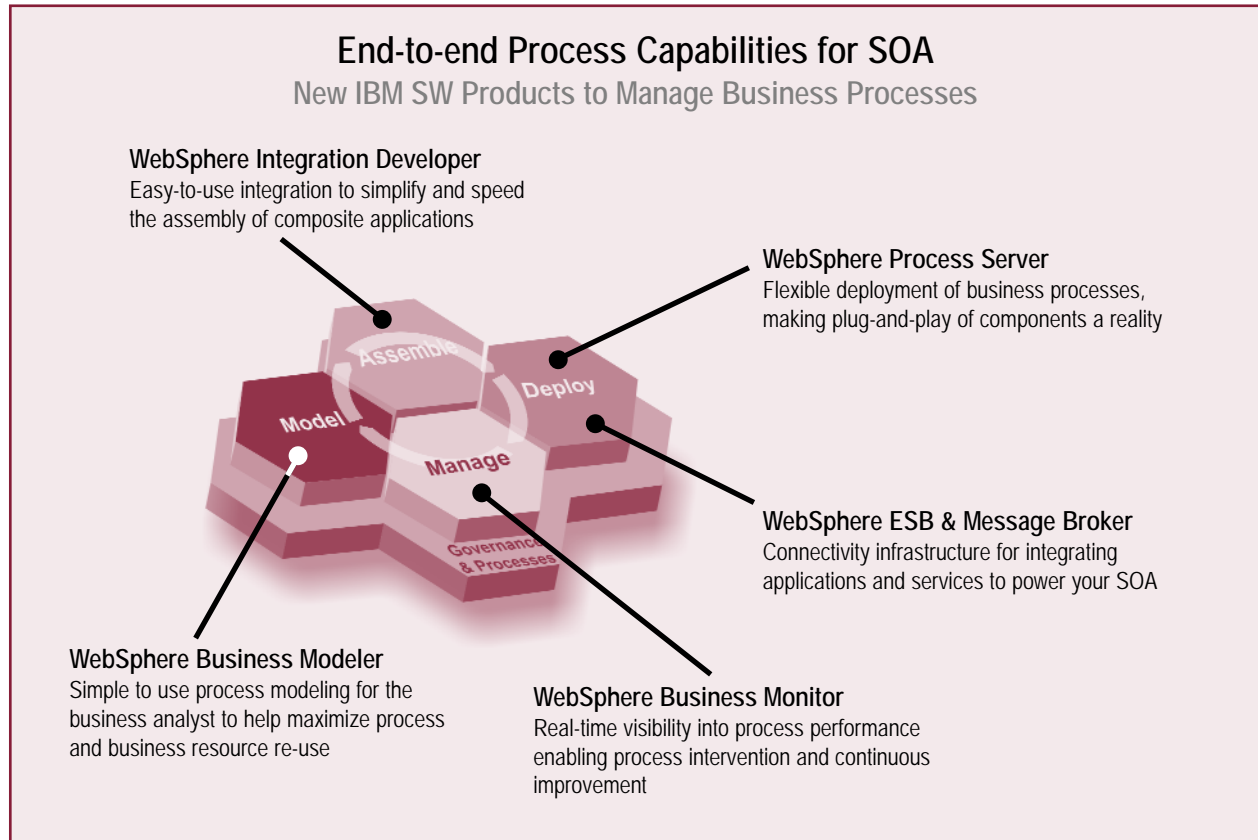


Figure C1: End-to-end Process Capabilities for SOA – New SOA Foundation Software

- **WebSphere Integration Developer:** To assemble (*choreograph or orchestrate*) new SOA applications, IBM announced WebSphere Integration Developer, an Eclipse-based application development tool to build and deploy business processes based on SOA. (*Eclipse is the leading open standards-based software platform for application development tools, supported by most tools leaders and players, that eases tool integration round the common Eclipse Platform.*) WebSphere Integration Developer allows developers of composite applications to view existing IT applications as services that can be easily wired together to compose full business processes. To further help customers assemble SOAs, IBM has also released a new version of its Rational Application Developer product, which is used to build new component services. (*WebSphere Integration Developer will supersede the current WebSphere Studio Application Developer – Integration Edition (WSAD-IE) product when available.*)
- **WebSphere Enterprise Service Bus:** To help clients deploy SOAs, IBM also announced its new WebSphere ESB product. Further strengthening its current ESB capabilities, the WebSphere ESB product provides connectivity and integration for Web services-focused applications and services. For advanced ESB functionality, IBM is to deliver a new version of WebSphere Message Broker, which has been a leader in providing universal connectivity and data transformation for applications, whether or not they comply with open standards. With this WebSphere ESB product, IBM has now bowed to the market inevitable, finally acknowledging that an ESB should be a specific software product and not just an architectural pattern as it had previously argued.

- **WebSphere Process Server:** Also new for SOA deployment is WebSphere Process Server, open-standards-based software powered by WebSphere ESB that helps simplify the integration of business processes that span people, systems, customers and business partners. WebSphere Process Server can help reduce the time, cost, and risk of integration projects, because it simplifies the movement of information between applications based on business rules.
- **WebSphere Business Monitor:** To manage SOAs, IBM also announced an enhanced version of its powerful WebSphere Business Monitor product. This software helps monitor business process performance and key performance indicators to guide and aid the overall management and control of the business processes running under the SOA.

The WebSphere SOA Foundation, in our assessment, now provides the industry's strongest SOA middleware and tools suite...

All of these new products are expected to ship on all supported platforms, including z/OS for mainframe, by the end of Q1 2006, and some considerably earlier. Exact general availability dates were not published for all these products at time of writing, and full pricing is likely to be disclosed with the individual availability announcements.

In our assessment, this was a cohesive announcement that strengthens and completes the core WebSphere SOA software stack in all key areas, and considerably clarifies and improves product naming and packaging, which some felt was previously unduly complex.

Our Analysis

The WebSphere SOA Foundation, in our assessment, now provides the industry's strongest SOA middleware and tools suite, extending the established middleware market leadership of the WebSphere platform into this important and fast-growing market. The culmination of several years of intense IBM development, several component acquisitions, much open standards support development, and considerable optimization and enrichment efforts on both the servers and the tools, stand behind this considerable achievement.

...we consider IBM is well positioned to strengthen its already clear overall SOA market leadership over the next two years.

Significant repackaging and much-improved product naming make the components of the Foundation more flexible and easier to understand.

When combined with the extensive SOA education, knowledge and information, professional service capability, methodologies and other supporting elements, we consider IBM is well positioned to strengthen its already clear overall SOA market leadership over the next two years.

The zSeries/System z9 mainframe should, in our assessment, play a key part in enterprise SOA deployment and centralized control. It is to be well equipped with this up-coming WebSphere SOA Foundation suite, due out on z/OS before end Q1 2006 or earlier and with existing equivalent products already available now for early starters.

Appendix D: Impressive IBM Mainframe Storage Solutions Take Off

Over the last two years, IBM has transformed its entire storage portfolio, with new, performance and leadership enterprise storage systems, successful enterprise storage virtualization solutions, new tape systems, and enhanced storage management software. These storage advances are highly complimentary to the storage needs and capabilities of the zSeries and new System z9 mainframe, and support its proposed new role as the enterprise hub for data and for the provision of enterprise-wide "IAAS" services.



Our Enterprise Storage Spotlight Report – **"New Power-driven, High-end and Modular Enterprise Storage Systems – Game-changing Server Technologies/Advances Supercharge IBM's Storage Market Leadership Bid"** (November 2004) – covered the key new DS8000 and DS6000 enterprise storage systems that are the centerpiece of this new generation of advanced IBM mainframe storage systems in depth. In this Appendix we include here our executive summary findings on these systems from this Report:

1. **On Demand Era Demands Storage Improvements:** On Demand's new, optimized business models/processes demand more flexible, responsive storage infrastructures that are better able to support variable demands. Customers are demanding simplified, more manageable, and more cost-effective enterprise storage infrastructures.
2. **Improving Economy, Storage \$ Growth at Last:** Global economies are growing in 2004, as is the storage market after a three-year decline. The \$5.5B high-end enterprise segment will post a low single-digit rise in 2004. The larger \$8.7B mid-range storage segment will grow at 8-9% for the next several years. With combined \$14.2B customer expenditure in these segments, these are rich prizes for the leading vendors.

3. **Storage Networks Too Complex:** As enterprises implemented storage networks to better share data, complexity crept in, and multiple SANs with poor or no compatibility were implemented for different server platforms. For On Demand, users now need ways to consolidate and integrate their enterprise storage networks more cost effectively.
4. **2 Major New Enterprise-class series of Storage Systems from IBM:** On October 12th 2004, IBM made its most important and dramatic enterprise storage announcement of this decade, by announcing not just the one expected, but two all-new enterprise-class storage server systems:
 - **DS6000 series:** The surprise was this ultra-compact, high-density 3U enterprise-class server that scales up to 67.2TB, and sets a new price point. It also offers complete interoperability with its DS8000 big brother, ESS 800 and ESS 750, redefining mid-range enterprise storage with “no compromises” performance, reliability, scalability and copy functionality in a small package.
 - **DS8000 series:** Expected, but impressive in its all-new specifications and advanced capabilities was the DS8000 high-end, enterprise-class storage server, in two models scaling up to a high 192TB, using IBM’s leadership POWER5 microprocessors and p570 server engines. The DS8000 introduced the industry’s first storage system LPARs, allowing two storage images to be run on one system. The DS8000 offers up to 6 times the scale, large performance increases in 20% less space, at significantly lower cost than IBM’s ESS (*Shark*) its predecessor.

The new systems became available from December 3rd 2004, with entry level list prices beginning at \$97,000 for the DS6000 and \$250,000 for the DS8000. Our highlight summary of these two systems is shown in Figure D1.

Radical New Enterprise Storage Systems – Dec 2004 – Complement zSeries, System z9

| | |
|---|--|
| <div style="display: flex; justify-content: space-between;"> <div style="width: 60%;"> <p>DS6000 Enterprise-Class Storage For All Sizes Of Customer</p> <ul style="list-style-type: none"> Full Enterprise Capabilities: At 50% of price Highly-Scalable: Up to 67 TB Dense Modular Packaging Excellent Performance: 2X EMC CX700 Enterprise-Class Resilience: Ultra-Competitive vs. EMC DMX800: 2X scalability, 50% of price Server-Inspired Technology: PowerPC based Easy To Install & Service: Customer install/repair Systems from: \$97,000 Complements: z890, System z9, Open Systems & ILM </div> <div style="width: 35%; text-align: center;">  </div> </div> | <div style="display: flex; justify-content: space-between;"> <div style="width: 60%;"> <p>DS8000 New Standard In High-End Enterprise Storage</p> <ul style="list-style-type: none"> All-New, HE ESS: Technology tour-de-force Dramatic Performance: 6X vs. ESS 800 Massive Scalability: <192TB today, >1PB SOD Exceptional Consolidation Capability: 1st With Storage System LPARs (2) Advanced server-based technology: POWER5 & p570 engines Enterprise Resiliency/BC: Near 6-nines Flexibility/Extensibility: Highly flexible, extensible ESS Much Lower Prices: 50/60% under ESS 800 Systems from: \$250,000 Complements: z990, System z9 & high-end Open systems </div> <div style="width: 35%; text-align: center;">  </div> </div> |
|---|--|

Common Features Both Support zSeries & Open Systems, Compatible/Interoperable Copy Services, Metro/Global Mirror & Copy, Common User interfaces, Common New Management Suite, Common Code Base, 1st 4-year warranties

Figure D1: New Enterprise Storage Systems – Dec 2004 – Complement zSeries, System z9

5. **Creates New Enterprise Storage Continuum:** To date, high-end and mid-range enterprise storage have been worlds apart, with separate product lines, copy service, and management interfaces adding complexity and cost for customers. The DS6000 and DS8000 break down these costly barriers to provide the first continuum of enterprise storage across the mid-range and high-end, with the same functionality, copy services and management software.
6. **Allow First Unified Enterprise Storage Infrastructures:** These systems now allow users to eliminate long-standing, costly, dual SAN environments, and deploy one smooth enterprise continuum to support both mainframes and Open Systems on one affordable, common, integrated SAN environment. Major savings from such consolidation are obtainable.

7. Systems Complete New DS Family Line: IBM named a new brand of IBM TotalStorage DS Family last month, adding new low-price point entry-level systems (*DS300 & DS400*) to its successful mid-range (*now DS4000 series*). With this October 2004 addition of the breakthrough DS6000 and DS8000 enterprise-class storage servers, the giant now claims the widest and newest single-vendor disk line-up on the market.

8. Customer Needs Foremost: In designing the new DS6000 and DS8000, IBM looked hard at, and listened carefully to, what enterprise customers most needed; to better protect data, simplify storage infrastructures, and to manage data efficiently throughout its lifecycle, in the faster On Demand Era. Four main needs were clear:

- Higher upward scalability with more granular horizontal growth.
- Unified solution to support heterogeneous servers.
- Needed range of Quality of Service (*QoS*) options to suit different tiers.
- Need improved disaster recovery processes.

With the DS6000 and DS8000, in our Report's detailed assessment, the firm has hit these needs squarely on the nose, and raised the competitive game substantially.

9. Balance Tips to Server-based Technology: Enterprise storage competitors like EMC long used proprietary, specialized storage systems, custom-built end-to-end for the role in hardware and software, and sold at high prices. IBM's new products are its fifth generation of general server-based enterprise storage systems, and heavily leverage IBM's milestone PowerPC/POWER5 microprocessor and server advances. Now these server-based systems have pulled ahead in scalability, performance and price, tipping the balance in IBM's favor. The strategic inflexion point in favor of server-based storage solutions is now truly here!

10. Mold-breaking DS6000 series Looks a Winner: With its dramatically smaller form, mid-range prices, but full enterprise-class capabilities, the DS6000 series, will be attractive to customers in at least five common situations, including:

- Mid-sized enterprises without enterprise-class storage today.
- Users with mainframes and Open Systems on separate SANs.
- Users of small-medium IBM ESS and other older competitors' systems.
- z890 mainframe customers.
- Customers needing lower-cost Disaster Recovery/Business Continuance (*DR/BC*) options.

These are widespread, the DS6000 completely outclasses competition, and we assess is certain to become a major success and sector leader.

11. DS8000 series Goes for the High Ground: The DS8000's impressive gains in performance, 6-fold higher scalability, 20% smaller footprint, and competitive pricing (*than the ESS 800*) are market game-changing, and present a formidable IBM challenge to leader EMC and other main competitor HDS. Prime opportunity "sweet-spots" for IBM with the DS8000 are:

- ESS replacements.
- Booming zSeries market.
- New-name open bid wins from EMC or HDS.
- Storage consolidation for multiple server environments
- High-end pSeries and iSeries.
- IBM Global Services (*IGS*) services-led storage improvement projects.

Most early DS8000 successes will come from these scenarios. Medium-term, IBM expects to engage full battle with HDS and EMC at the high-end, and win, with customers benefiting from this new market option.

12. Shared Copy Services Leadership Will Help: IBM's TotalStorage Resiliency Family products provide common copy services that enable affordable DR/BC and Information Lifecycle Management (*ILM*) solutions across the new DS8000, DS6000 and ESS enterprise storage systems continuum. This unique commonality allows CIOs to extend continuity and ILM strategies more broadly, easily and affordably. This strength will aid IBM's battle towards higher share and ultimate storage leadership.

13. **IBM's Storage Leadership Push Accelerates:** With these two powerful enterprise-class storage systems, within a new overall disk range, IBM makes no secret of its ambition to recapture the number-one position in both enterprise and overall storage markets. With these new systems, it has formidable product weapons, and extended channels, that are certain to see it gaining share and momentum through 2005 on the first of these goals. The competitors are tough and entrenched, but the IBM offerings are so radical that chances are good. It is likely to take three to four years for IBM to attain its target, and customers will benefit throughout from these innovations and stronger competition.
14. **Challenges Ahead:** IBM faces entrenched and capable competitors in EMC and HDS, who will not easily relinquish their current positions. EMC's customers are familiar with its copy services and software, and are generally well supported. HDS struck first, by announcing their new TagmaStore USP HE ECSS and SAN virtualization offering a timely month before IBM's announcement, and will, no doubt, be clearing inventory of Lightning 9900 V series at favorable prices. IBM needs to seize the bridgeheads of opportunity that we identify in 10 and 11 above rapidly to gain momentum, and to push the envelope of technology advance briskly. It must also ramp up its Business Partner channels to fully address the mid-market opportunities. These are demanding challenges, and will require consistent execution.
15. **Larger DS8000 Models in the Wings:** Adding gasoline to the fire it already lit, IBM gave a Statement of Direction (*SoD*) for an even larger, more powerful DS8000 model. Scaling up to 384TB, the SoD 8-way systems will keep IBM at the top of the enterprise storage bragging scale when the top segment of the market requires systems of this power. Word is that even larger 12-way models are also on IBM's roadmap.
16. **Storage Software Portfolio Supports Aims:** IBM extended its TotalStorage Open Software Family, including the TotalStorage Productivity Center Suite, which will be a good management option for DS6000 and DS8000 customers. These integrated storage management tools now include automated server and storage provisioning that cuts storage staff effort and time to service. These tools ensure customers get the best from their new DS6000 and DS8000 systems, and merit closer review and full consideration.

These are interesting times in the enterprise storage business! With this announcement, IBM has put a stake in the ground with two new and attractive major offerings that redefine enterprise storage, and will be welcomed warmly by its many customers. Established competitors will be forced to respond – to the benefit of customers.

Related Software Strategies Research

1. **“SOA Takes Off – New WebSphere SOA Foundation Extends IBM's Lead with New System z9 Mainframes as the Hub of the Enterprise.”** Software Strategies White Paper, 2nd Edition, November 2005, 56 p.p., 20 charts and tables. (*In-depth assessment of the case for adopting SOA, evaluation of the IBM WebSphere SOA Foundation software, methods and services, and the role of the System z9/zSeries mainframe as the enterprise SOA hub.*)
2. **“Spectacular System z9 Mainframes Leap Ahead with Doubled Power, Enterprise Hub Roles – Virtualization, Security, Availability, SOA & Value Advances.”** Software Strategies Enterprise Server Spotlight Report, September 2005, 72 p.p., 28 charts and tables. (*In-depth assessment of new IBM System z9 109 mainframe platform hardware, operating systems, middleware, storage and virtualization capability, first major new system to emerge under IBM Systems' new “Systems Agenda” strategy for the emergent era of collaborative computing.*)
3. **“New IBM Systems Agenda Strategy Revealed – Announces New High-end System z9 Mainframes, Virtualization Engine Version 2.0, & new Blade.org.”** Software Strategies Analyst Announcement Assessment, August 2005, 48 p.p., 22 charts and tables. (*IBM's most important Systems announcements of 2005 (July 26th) debuted its new “Systems Agenda” strategy for the emergent era of collaborative computing. This Assessment provides a comprehensive, independent review/analysis of these industry-significant developments.*)
4. **“Maximizing Business Value With Resurgent zSeries Mainframes – Platform Readiness Key in 2005.”** Software Strategies White Paper, September 2004, 56 p.p., 26 charts and tables. (*Focuses on key business values delivered by resurgent zSeries mainframe platforms, and assesses main advances in the 2005 mainframe software stack that combine to provide compelling advantages of adopting platform-readiness strategies exploiting these mainframe advances fully.*)
5. **“Enterprise Integration Challenge – Vendor Middleware-based Integration Solutions Offer Major Advantages Over Custom In-house Solutions.”** Software Strategies White Paper, 2nd Edition, April 2005, 36 p.p., 20 charts and tables. (*Reviews the actual methods enterprises have used to date for enterprise/application integration, and evaluates the case for using vendor enterprise integration middleware. Also reviews IBM WebSphere MQ Version 6, the latest release of this industry-leading integration MOM middleware suite.*)

6. **"Enterprise Transformation, Modernization & Integration Top Priority Today – Resurgent zSeries Mainframe Stakes Powerful Claim for Expanded Role."** Software Strategies White Paper, 3rd Edition, April 2004, 46 p.p., 20 charts and tables. (*Enterprise transformation, modernization and integration with resurgent zSeries. Software stack-centered, in-depth assessment of the mainframe as an enterprise platform as at fall 2004.*)
7. **"New Power-driven, High-end and Modular Enterprise Storage Systems – Game-changing Server Technologies/Advances Supercharge IBM's Storage Market Leadership Bid."** Software Strategies White Paper, 2nd Edition, November 2004, 42 p.p., 21 charts and tables. (*In-depth technology assessment of IBM's new DS6000 and DS8000 enterprise storage systems.*)
8. **"IBM WebSphere Business Integration Leads the Way with Fullest J2EE™ and Eclipse Support – Business Integration Software Swings to Open Standards."** Software Strategies White Paper, 2nd Edition, September 2004, 24 p.p., 16 charts and tables. (*Comparative evaluation of WebSphere Business Integration Suite & open standards support.*)
9. **"zSeries Mainframe Resurgence Beyond Question – Software/Costs See Major Advances: IBM Tools Break USV Grip."** Software Strategies White Paper, January 2004, 50 p.p., 18 charts and tables. (*zSeries resurgence and software stack assessment.*)

Software Strategies

Software Strategies is a specialist analyst firm focused on e-Infrastructure platform strategies and issues. Since 1997 we have published reports, white papers and newsletters, hosted and spoken at numerous successful industry events, and have worked closely with industry leaders, including: IBM; Microsoft; Intel; ICL; Unisys; Computer Associates; BMC; Stratus Computers; Misys Plc; Notability Solutions; and many others.

Specialist expertise on mainframes, servers, storage, processors, Microsoft technology, Intel environments and enterprise middleware has been a common thread. Tens of thousands of enterprise IT users have benefited from our authoritative events, presentations, conferences, newsletters, white papers, and reports.

About the Author



Ian Bramley
Managing Director – Software Strategies
ian.bramley@software-strategies.net

This White Paper was researched and written by Ian Bramley, Managing Director of Software Strategies, and published in February 2006. The views expressed are those of Software Strategies, and are based on our proprietary research. Bramley founded Software Strategies in 1997, is an experienced enterprise infrastructure analyst, a keynote speaker at many industry events, and has published many reports and papers. Before this, he served as Director of Enterprise Platforms at Butler Group and was Founder/Chairman of the Enterprise NT Management Forum from 1998 to 2001. Previously, he held a variety of executive positions with 4 international software and services vendors over a 30-year industry career.

