DB2 Information Management Software

IBM

The Business Value of DB2 UDB for z/OS

- Learn about the synergy of DB2 with IBM eserver zSeries and z/OS
- Discover how you can benefit from the richness in function
- Be ready for On Demand Business

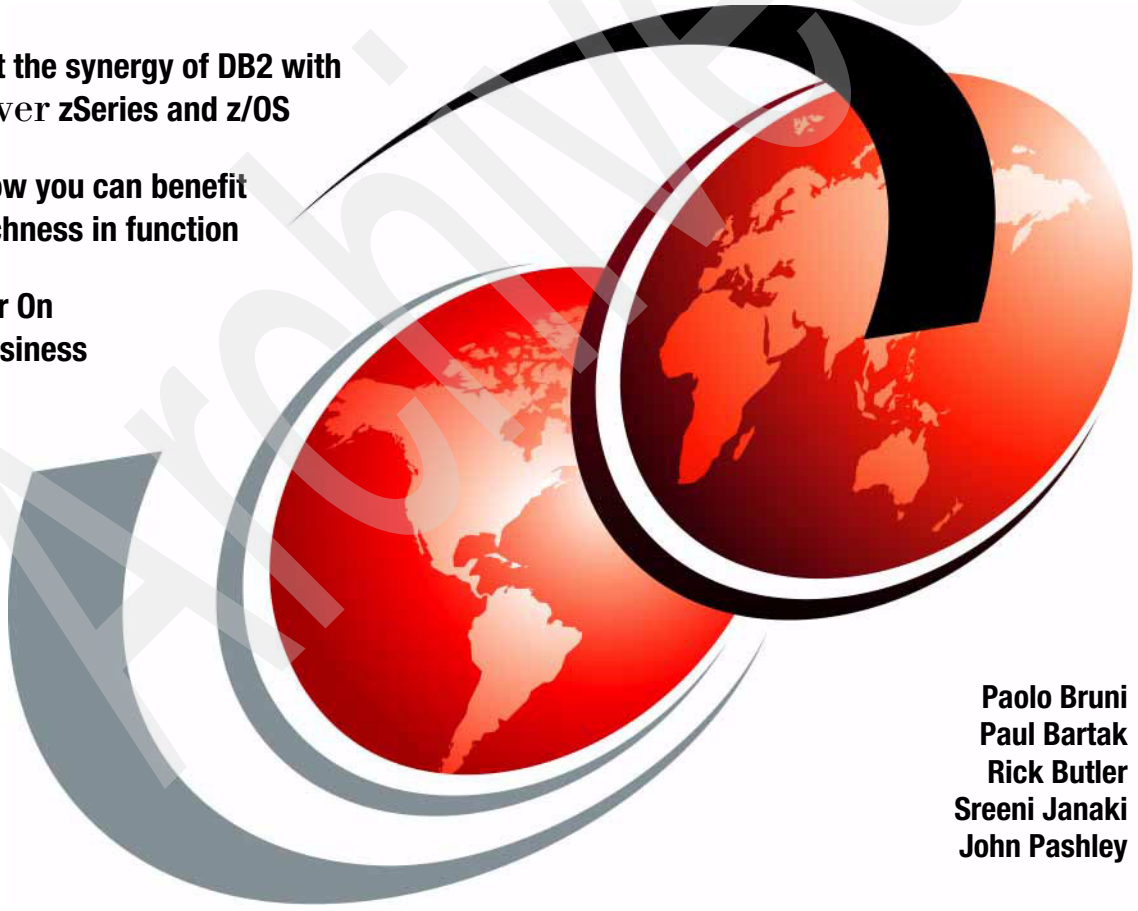

Paolo Bruni
Paul Bartak
Rick Butler
Sreeni Janaki
John Pashley

Redbooks

ibm.com/redbooks

International Technical Support Organization

The Business Value of DB2 UDB for z/OS

June 2005

SG24-6763-00

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

First Edition (June 2005)

This edition applies to IBM DB2 Universal Database for z/OS Version 8 (program number 5625-DB2).

Contents

Notices

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

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

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

The following paragraph does not apply to the United Kingdom or any other country where such provisions are inconsistent with local law: INTERNATIONAL BUSINESS MACHINES CORPORATION PROVIDES THIS PUBLICATION "AS IS" WITHOUT WARRANTY OF ANY KIND, EITHER EXPRESS OR IMPLIED, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF NON-INFRINGEMENT, MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE. Some states do not allow disclaimer of express or implied warranties in certain transactions, therefore, this statement may not apply to you.

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

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

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

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

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

COPYRIGHT LICENSE:
This information contains sample application programs in source language, which illustrates programming techniques on various operating platforms. You may copy, modify, and distribute these sample programs in any form without payment to IBM, for the purposes of developing, using, marketing or distributing application programs conforming to the application programming interface for the operating platform for which the sample programs are written. These examples have not been thoroughly tested under all conditions. IBM, therefore, cannot guarantee or imply reliability, serviceability, or function of these programs. You may copy, modify, and distribute these sample programs in any form without payment to IBM for the purposes of developing, using, marketing, or distributing application programs conforming to IBM's application programming interfaces.

Trademarks

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

@server®	DFSORT™	Parallel Sysplex®
@server®	DRDA®	PR/SM™
Redbooks (logo) ™	Enterprise Storage Server®	QMF™
eServer™	Enterprise System/3090™	Rational®
ibm.com®	ES/3090™	Redbooks™
iSeries™	ES/9000®	RACF®
pSeries®	ESCON®	RS/6000®
z/Architecture™	FlashCopy®	S/360™
z/OS®	FICON®	S/390 Parallel Enterprise
z/VM®	Geographically Dispersed	Server™
zSeries®	Parallel Sysplex™	S/390®
AIX®	GDPS®	Seascape®
CICS®	Hiperspace™	System/360™
CICSPlex®	HyperSwap™	System/370™
Database 2™	Informix®	System/390®
Distributed Relational Database	IBM®	SQL/DS™
Architecture™	IMS™	Tivoli®
Domino®	Lotus®	TotalStorage®
DB2 Connect™	Magstar®	VM/ESA®
DB2 Extenders™	MQSeries®	VSE/ESA™
DB2 Universal Database™	MVS™	WebSphere®
DB2®	MVS/ESA™	3090™
DFSMS/MVS®	NetView®	
DFSMSdss™	OS/390®	

The following terms are trademarks of other companies:

Java and all Java-based trademarks are trademarks of Sun Microsystems, Inc. in the United States, other countries, or both.

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

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

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

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

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

Preface

We are in a new phase of On Demand Business. Companies must be able to respond to fluctuating market conditions and provide products and services on demand to customers. DB2® for z/OS® and IBM @server® technology enable you to maintain your company's core competencies. They help you to manage global market and competition changes with consistent availability, security, and privacy, everywhere, all the time.

DB2 continues to deliver rich functionality for scalable and highly available data for enterprise-scale, on demand applications. Combine the power and capacity of IBM @server zSeries® and the high performance and high availability of z/OS with the new version of DB2. In doing so, you'll expand and extend your applications in the increasingly competitive on demand environment. DB2 Universal Database™ (UDB) for z/OS Version 8 delivers power, flexibility, and DB2 family compatibility through synergy with zSeries and z/OS.

This IBM® Redbook explains:

- How DB2 Database Servers are open systems that integrate with, and enable the sharing of information across multiple platforms.
- How the zSeries 990 server is equipped to deal with the explosion of data and handle it with incredible reliability and security: This in turn accelerates time-to-value while maximizing the return on investment of your digital information assets.
- How DB2 autonomic computing functions reduce the skills and staffing requirements to minimize cost outlays and reduce risk.

This IBM Redbook helps you to position DB2 UDB for z/OS as a springboard for the future. For a more technical description of the synergy between DB2 and the zSeries platform, and more details about DB2 functions, refer to the DB2 for z/OS Version 8 Web site at:

http://www.ibm.com/software/data/db2/zos/db2zosv8.html

The team that wrote this redbook

This redbook was produced by a team of specialists from around the world working at the International Technical Support Organization (ITSO), San Jose Center.

Paolo Bruni is a DB2 Information Management Project Leader at the ITSO, San Jose Center. He has authored several Redbooks™ about DB2 for z/OS and related tools, and has conducted workshops and seminars worldwide. During Paolo's many years with IBM, in development and in the field, his work has been mostly related to database systems.

Paul Bartak is an IBM DB2 for z/OS Specialist who is responsible for developing and delivering the DB2 for z/OS V8 Migration Planning Workshop. He has been working for 20 years on DB2 for MVS™, OS/390®, and z/OS (starting with Version 1 of the product), with roles that have included application design and development, database administration, DBA management, and Architect and Technical Leader. Since joining IBM in 1999, he has provided support to clients for all members of the DB2 family.

Rick Butler is a DBA Specialist and DBA Manager at Bank of Montreal (BMO) Financial Group, based in Toronto, Canada. He has 18 years of experience in designing and implementing DB2 databases for mainframe banking applications. Prior to this, he worked with network and hierarchical databases, and in London, as a programmer and designer on a foreign exchange application. He also has data warehousing experience with DB2 on AIX®. He has expertise in access path statistics and DRDA®. He holds a Bachelor of Science degree in operations research from Université de Montréal.

Sreeni Janaki is an IT consultant working for Covansys Corporation in the United States. He has more than 10 years of experience developing Business Intelligence systems, many of them implemented on DB2. His areas of interest include database design and database administration. He holds two Master of Science degrees in engineering from the University of Alaska at Fairbanks.

John Pashley is an IT Architect for IBM in the United Kingdom. He has more than 20 years of experience in the IT industry, the majority of which working on the IBM mainframe platform. He has experience in working across a diverse spectrum of industry sectors, with extensive involvement in finance and international payments.

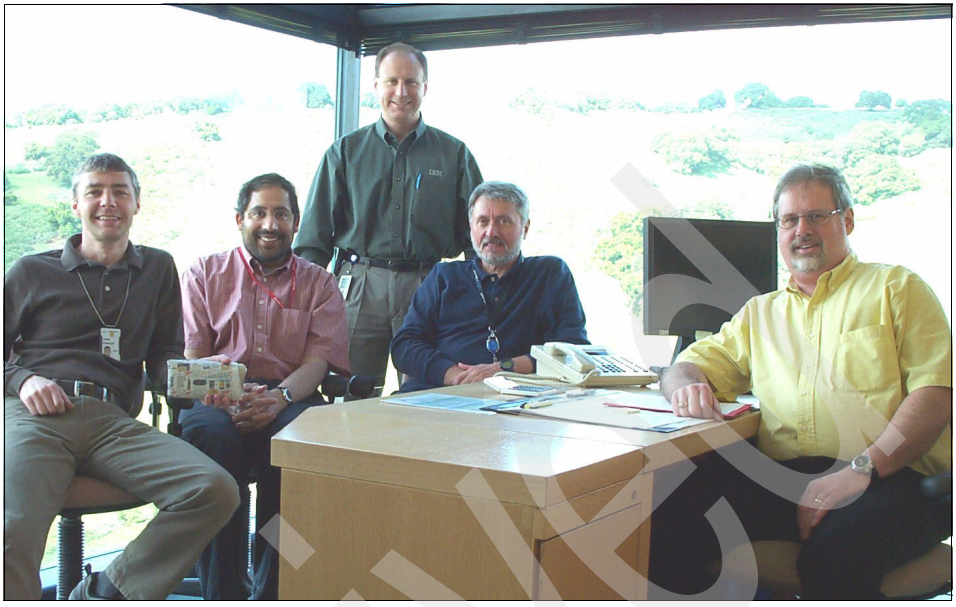

The team (left to right): John Pashley, Sreeni Janaki, Paul Bartak, Paolo Bruni, and Rick Butler

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

Frank Kyne
Leslie Parham
Sangam Racherla
ITSO

John Campbell
Curt Cotner
Gene Fuh
Terrie Jacopi
Jeff Josten
Glenn McGeoch
Roger Miller
Britta Ottersgard
Ron Williams
Maureen Zoric
IBM Silicon Valley Lab

Claudia Sweet
IBM Rochester

Paul DiMarzio
IBM Poughkeepsie

Become a published author

Join us for a two- to seven-week residency program! Help write an IBM Redbook dealing with specific products or solutions, while getting hands-on experience with leading-edge technologies. You'll team with IBM technical professionals, Business Partners, and/or customers.

Your efforts will help increase product acceptance and customer satisfaction. As a bonus, you'll develop a network of contacts in IBM development labs, and increase your productivity and marketability.

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

ibm.com/redbooks/residencies.html

Comments welcome

Your comments are important to us.

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

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

  ibm.com/redbooks

- Send your comments in an e-mail to:

  redbook@us.ibm.com

- Mail your comments to:

  IBM Corporation, International Technical Support Organization
  Dept. QXXE Building 80-E2
  650 Harry Road
  San Jose, California 95120-6099

1

Evolution of the mainframe

The mainframe as we know it today started with the introduction of the System/360™ (S/360™) in the mid-1960s. As companies caught on to the business benefits of computing, the workload volume and complexity started to increase. Even as demand for more processing power and throughput were met by revolutionary technology advancements, the mainframe never lost sight of its core principles.

Unrivalled reliability, availability, and serviceability (RAS) saw the mainframe become the focal point of enterprise IT infrastructure in most of the world's largest corporations. This is still true today where it is somewhat taken for granted as the dependable workhorse, in some respects suffering from its own success in striving for zero downtime.

Access to information and services entered the on demand era with the commercialization of the Internet and the global impact of e-business. The quiet dependability of the mainframe, married with ground-breaking technology, make this computing platform more attractive now than it is ever been. Together they allow evolutionary legacy application transformation as a means of enhancing competitiveness in the era of On Demand Business.

1.1 The evolution of z/Architecture

In the decades since the 1960s, mainframe computers have steadily grown to achieve enormous processing capabilities. Mainframes of today have an unrivaled ability to serve end users by the tens of thousands, to manage petabytes of data, and to reconfigure hardware and software resources to match changes in workload, all under centralized operational control. Currently, the zServer 900 and 990 (z900 and z990) family of mainframes handles such heavy computing needs.

1.1.1 Lower cost of entry to mainframes

In addition, IBM has lowered the point of entry into z/Architecture, for example, by introducing the zSeries 800 (z800) mainframe server in February 2002. Much less expensive than the z900, the z800 is minimally different in terms of features. At its introduction, clients could choose from 80 to more than 600 million instructions per second (MIPS).

IBM also introduced *z/OS.e*, which is a specially priced z/OS offering that can be attractive to clients as they battle to find a software and hardware configuration to handle their workloads. The flexibility of pricing is evident in its multiple available pricing options.

1.1.2 More options in pricing and performance

The more recent IBM offering of zSeries 890 (z890) shows the same IBM commitment to providing more alternatives to our clients based on their unique needs. The successor of the z800, the z890, was introduced in April 2004 and is based on a lower-capacity version of the z990 chip set. It offers several different performance levels and is upgradable from z800.

Performance options start from less than 30 MIPS and extend to approximately 1400 MIPS. This gives lower-end mainframe clients a great choice, with as many as nine options under 150 MIPS. z990 functions have also been brought into this market as the z890 is based on a five-processor multichip module that provides four instructions processors and one input/output (I/O) processor.

For details about the z890, visit:

http://www.ibm.com/servers/eserver/zseries/z890/

1.1.3 Leveraging investment in existing applications

In a fast-changing world, the IT department must embrace the latest industry trends if it is to keep the business current. Business-to-business trading platforms and customer expectations continue to evolve as On Demand Business becomes an intrinsic element of life. The enterprise IT infrastructure must adapt constantly to meet the requirements of the business as it battles increasing competition and industry pressures.

We have entered a new and exciting era as online business moves from acceptance to expectation. Web services move into the mainstream as the XML-based Web Services Description Language (WSDL) standard takes hold. Web services provide a foundation for Service-Oriented Architecture (SOA), a component-based interaction model that is set to advance from technical forum topic to business imperative.

However, this does not mean that clients no longer need the applications that run on their traditional systems. Many clients leverage their existing applications by enabling the newly developed business logic to invoke their existing transactions in real time, thus creating a composite application. These composite applications are a reality because the clients' investment in application programs, whenever they were built, has been preserved while innovation continued.

A logical partition (LPAR) on z/OS can run Information Management System (IMS™) as easily as it runs a zSeries Application Assist Processor (zAAP), which is useful in running Java™ workloads. zAAP is also known as an integrated facility for application (IFA) processor, and is available on z990 and z890 servers. As mentioned previously, zAAP is a processor unit that is used exclusively for running Java application workloads under z/OS. A zAAP works asynchronously with the central processor (CP) to execute Java programming under control of the IBM Java Virtual Machine (JVM). A CP is a processor unit that has the z/Architecture™ and ESA/390 instruction sets. The JVM processing cycles can be executed on the configured zAAP. Unlike the CP, zAAP can do nothing on its own, which is why IBM does not impose software charges on zAAP capacity.

There is opportunity for hybrid solutions using Linux® and z/OS. An example is when an application server (such as an SAP Application Server or IBM WebSphere® Application Server) runs on a Linux instance and a DB2 database server runs within a z/OS LPAR. Then you can use zSeries HyperSocket technology for communication between the servers. Linux can run on the "bare hardware" of the zSeries machine without any special software or support required.

Linux-z/OS hybrid solutions represent a formidable merging of technologies in a world where IT organizations look to maximize cost-effectiveness and performance for large, enterprise-class workloads. Linux provides a convenient

and cost-effective environment for migrating UNIX®-based and Microsoft® Windows® NT-based applications to the mainframe. z/OS provides the best environment for high-performance DB2 database server components in the solution.

Another example of a Linux-z/OS hybrid solution is the porting of a DB2 Connect™ implementation from a server into a Linux partition on the mainframe. This configuration potentially performs better and more flexibly, being closer to DB2. It also saves the cost and management of a server.

See the paper *Mainframe Linux Solutions Mature with mySAP.com for IBM eServer zSeries* by David Mastrobattista (2002) of the Giga Information Group, Inc. You can find this paper on the Web at:

http://www.ibm.com/servers/eserver/zseries/library/whitepapers/pdf/gigamar19mysaplinuxonz.pdf

1.2 Hardware

We now chart some of the significant events in the evolution of the IBM mainframe platform and call out important implications for the modern business environment. From this, you see how the strategy of steady innovation on a firm foundation has lead to the modern rock-steady platform on which an enterprise can build its business into the future.

The first general purpose automatic digital computer built by IBM dates back to 1944. It was an electromechanical machine developed in conjunction with Harvard University, known as the *Automatic Sequence Controlled Calculator*. It performed addition in one-third of a second and multiplication in six seconds. From then on, innovations continued, at times directly due to events such as the Korean war in the 1950s, resulting in faster and more versatile machines, and leading to the System/360.

1960s

On 7 April 1964, IBM introduced the System/360, a family of five increasingly powerful computers that ran the same operating system and used the same 44 peripheral devices.

It was a notable achievement. Before this point, computer makers manufactured unique computers for specific customers. Customers had to rewrite their software applications every time they upgraded to a larger system. None of their existing peripherals, such as printers, tape drives, or disk storage, would work on the new system without extensive modification. With the S/360, for the first time, clients could scale up without a complete reinvestment in software and peripherals.

The origins of z/OS also date back to the early 1960's time frame with the introduction of the IBM 360 family of mainframe computers. This shows that, for a long time, IBM has been building hardware with solid operating system software that exploits the hardware to the fullest. Read more about this in 1.3, "Operating systems" on page 10.

In 1968, IBM introduced the Customer Information Control System (CICS®) and IMS. They enable workplace personnel to enter, update, and retrieve data online. To date, they remain the industry's most popular transaction monitors.

1970s

In June 1970, IBM introduced the System/370™, which could also run System/360 programs, easing the upgrade burden on clients.

Note: This remains a core principle of IBM mainframe servers to this day. Significant attention is given to backwards compatibility throughout the evolution and technological advancement of the platform.

The System/370 was one of the first lines of computers to include "virtual memory" technology. This technique was developed to expand the capabilities of the computer by using space on the hard drive to accommodate the memory requirements of software.

In 1972, IBM announced the virtual machine (VM). VM has long been recognized as a robust computing platform, spanning the family of S/390® and zSeries servers. Today, z/VM® helps create an agile mainframe where resources can be quickly and effectively used in response to the dynamic requirements of the latest on demand solutions. Built on the solid VM/ESA® base, z/VM exploits the z/Architecture. It helps enterprises meet their growing demand for multisystem server solutions. And it offers a broad range of support for operating system environments such as z/OS, z/OS.e, OS/390, TPF, VSE/ESATM, CMS, and Linux for zSeries and S/390.

Also in 1972, SAP developed a revolutionary enterprise resource planning (ERP) system for the System/360. For the first time, companies could place orders and track inventory in real time. This helped to improve inventory control, delivery time, and customer relations.

Note: The special relationship between IBM and SAP remains today with specific attention given to the needs of SAP/R3 in the latest version of DB2 Universal Database (UDB) for z/OS.

In 1976, SAS software became available on the System/370. SAS helped to create a new competitive edge: business intelligence. This innovation transformed raw data into actionable intelligence that helped organizations develop more profitable customer and supplier relationships, and enable better, more accurate decisions.

Note: Do you think they knew 30 years ago how significant Business Intelligence (BI) would become?

Today DB2 UDB for z/OS and the zSeries mainframe represent a powerful force with their enviable performance in relation to storage management, data access, and complex query processing.

1980s

In 1980, the introduction of the 3081 processor offered a two-fold increase in internal performance over the previous mainframe processor. It also featured Thermal Conduction Modules (TCM) that significantly reduced space, cooling, and power requirements.

The early 1980s saw the commercialization of relational systems. Oracle entered the marketplace and IBM released several relational products with an eventual focus on Database 2™ (DB2).

Structured Query Language (SQL) gained recognition as the industry standard. DB2 became the IBM flagship database management system (DBMS) offering over IMS.

By 1988, IBM clients began deployment of DB2 beyond decision support systems (DSS) and into core transactional processing. This drove reductions in CPU costs and dramatic improvements in concurrency.

IBM introduced the Enterprise System/3090™ Model 600S. This gave clients 56% more processing power. It also lead to a new 10-model S series of advanced mainframe computers that took advantage of IBM enterprise-strength operating systems and data management software.

Note: Significant in today's On Demand Business environment is the partnership of DB2 and z/OS and their ability to process extreme numbers of transactions regardless of unpredictable workload peaks. Integrity of data and security of transaction is no less important.

1990s

In the 1990s, IBM introduced System/390® (S/390), the company's most comprehensive rollout of products, features and functions in more than a quarter of a century. Some industry pundits, however, did not think the mainframe would survive the early 1990s. They predicted that the rapid growth in personal computers and small servers would render large central systems obsolete.

IBM believes that serious, security-rich, industrial-strength computing will always be in demand. Therefore, IBM reinvented and redefined the mainframe from the inside, infusing it with an entirely new technology core and reducing its price.

Key elements of this process included the reduction of the cost of the processor by switching from bipolar technology to Complementary Metal Oxide Semiconductor (CMOS) chip technology. This allowed for the scalability and availability provided by a unique clustering approach for S/390 (the Parallel Sysplex®). It also supported open interfaces such as Ethernet, FDDI, ATM, TCP/IP over Systems Network Architecture (SNA), and the adoption of UNIX standard programming interfaces.

In 1994, IBM announced the first S/390 Parallel Sysplex offering. It encompasses the Coupling Facility, the S/390 Parallel Transaction Server, high-speed coupling links, and software enhancements. The Parallel Sysplex is designed to provide continuous availability of applications, reduce or eliminate planned application outages, and scale to virtually unlimited capacity.

Most of the transformation was completed by mid-1997.

In 1998, IBM announced the S/390 Parallel Enterprise Server™ Generation 5 Models. This consists of 16 new S/390 models, including fifteen general purpose models and one Coupling Facility model. The Coupling Facility model exploits the latest IBM CMOS technology yielding up to two times the performance of the corresponding previous G4 models.

Note: Parallel Sysplex is key to the phenomenal capabilities of z/Architecture to scale-out across several closely integrated high power mainframe servers to meet the needs of today's most demanding business critical implementations.

By the end of the decade, with the growth of the Internet and advent of a networked economy, the number of companies trying to make the transition to e-business started to explode. The demand for more mainframe power continued to grow.

In 1999, Capacity Upgrade on Demand (CUoD) debuted on the S/390. It provides extra processors installed as spare capacity that can be "turned on" as

dictated by business needs. It provides a critical tool that can help companies better manage spikes in demand and handle unpredictable changes.

Linux also appeared on the S/390, for the first time. It combined the flexibility of open source computing with the legendary scalability and reliability of the mainframe.

Note: Linux on the mainframe is one indication of IBM's commitment to open standards. The traditional strengths of the underlying hardware platform, coupled with support for the latest technology initiatives, make this a serious contender for on demand workload deployment.

2000s

IBM unveiled the z900, the first IBM mainframe built from the ground up with e-business as its primary function. The zSeries mainframe is built to handle unpredictable e-business, allowing thousands of servers to operate within one box. IBM also introduced z/OS, a new 64-bit operating system.

By 2002, with z/OS 1.4, the mainframe operating system set a new record for Secure Socket Layer (SSL) transactions per second. This allows companies to respond more rapidly to e-commerce customers in a highly secure environment.

IBM introduced the z800, a lower-priced, entry-class mainframe that fundamentally changes the economics of mainframe computing. With this move, IBM delivers for the first time advanced Parallel Sysplex clustering technology to entry-class mainframe clients.

Note: Smaller entry-level hardware, specialized operating system (OS) offering, new workload pricing models, Java and Linux support all point toward On Demand Business. They enable the mainframe to be presented as a viable server option, even for first-time clients.

In 2003, IBM introduced the Mainframe Charter. This document articulated the commitment to deliver innovative solutions, expand zSeries value, and foster a zSeries community to meet clients' On Demand Business requirements.

IBM also announced the z990, the world's most sophisticated server and the new flagship of the IBM @server family. It provides a balanced, highly secure foundation for dynamically balancing critical applications. It is the result of a four-year and more than USD $1 billion investment.

Important: True to its roots, the new server demonstrates a staggering increase in throughput rate on SSL transactions. As a key infrastructure component of On Demand Business, the support for high volume secure transaction processing is of major significance in choice of platform.

One hundred additional new Linux on zSeries applications were introduced. Linux on zSeries is the ideal foundation for on demand operating environments.

In 2004, the z890 acquired new options to help better manage software costs and innovative technologies designed to deliver a streamlined, more responsive infrastructure. The System/360's tradition of encompassing every user's need continues in the z890, with many advantages for medium-sized enterprises.

Further z990 enhancements and the zAAP were also announced in 2004.

Note: zAAP further demonstrates the modernization of the mainframe platform, providing significant cost benefits to Java application deployment in support of On Demand Business workloads.

Some of the strategic evolutionary enhancements traditionally seen across the mainframe lifeline include:

- More and faster processors
- More physical memory and greater memory addressing capability
- Dynamic capabilities for upgrading both hardware and software
- More sophisticated automated hardware error checking and recovery
- Enhanced devices for input/output (I/O) and more and faster paths (channels) between I/O devices and processors
- Sophisticated I/O attachments, such as local area network (LAN) adapters with extensive inboard processing
- Increased ability to divide the resources of one machine into multiple, logically independent and isolated systems, each running its own operating system
- Enhanced clustering technologies, such as Parallel Sysplex, and the ability to share data among multiple systems

Note: Despite the continual change, mainframes remain the most stable, secure and compatible of all computing platforms. The latest models can handle the most advanced and demanding customer workloads, yet continue to run many applications that were written in the 1970s or earlier.

1.3 Operating systems

In harmony with the mainframe hardware advances has been a steady evolution in operating system capabilities. The focus has been on maximizing the utilization of expensive technology and making provision for the throughput of ever increasing workloads across multiple applications.

Because the origin of z/OS started in the early 1960s and the topic of the book is the synergy between DB2 UDB and z/OS, we summarize the value of z/OS before we delve into the evolution of it to where it is right now.

1.3.1 Value of z/OS

To truly understand and appreciate the value of z/OS, you should be aware of its history and reason for being. As mentioned in the previous section, the origin of z/OS dates back to the early 1960s and the announcement of the IBM 360 family of mainframe computers. This family of computers, and the operating system software and middleware that run on it, were designed to support the important applications that drive business and are critical to business operation and generation of revenue. As such, reliability, availability, security, and serviceability were foremost on the minds of the developers.

Through the years that operating system has been constantly improved upon. Today in its current form of z/OS, developers still work toward the goal of a constantly available, high-performance operating environment. You can find a lot of the details, especially those which contrast UNIX with z/OS, on the Web at:

http://www.ibm.com/servers/eserver/zseries/zos/support/zos_availability.html

Serviceability

Architects and developers of z/OS work to reduce the occurrence of both planned and unplanned outages. Functional recovery routines (FRRs) are designed so that, in the event of a function failure, the recovery routine is capable of taking over and preventing the entire system from failing. Select IBM software components can also be serviced in a nondisruptive manner so that maintenance levels can be maintained without negative effect to business processes.

z/OS prides itself in having good first failure data capture (FFDC) techniques. They are designed to capture data and identify problems that occur during production without impacting system availability.

This functionality has been a core part of the operating system design since 1974. Because of its long history, sometimes we take this functionality for granted. In z/OS, if a task (comparable to a UNIX thread) fails, the system is

designed to quickly make a copy of the running programs, together with its data and system control blocks, to another location in memory. This data can then be dumped to disk or tape for later analysis. For user applications and many subsystems, the dump is designed to be created without impacting unrelated work on the system. Even for system-level problems, the focus is on collecting the data with the least impact on other work.

After a problem occurs and the dump is taken, z/OS and the subsystems typically either remain up or can be restarted without performing a re-initial program load (IPL). This capability is handled by recovery routines such as ESTAE and FRR being given control. These routines are designed to capture diagnostic data; identify and restore modified system structures, resources, and locks held by the failed user; and return control to the program at a specified point.

These serviceability features are available to be used by independent software vendors (ISV) and user applications. Contrast this with the UNIX environment. When there is a problem in the operating system code, subsystem, or even some application problems, the UNIX system typically takes a system (kernel) dump and then stops, requiring a re-IPL.

An important mechanism to help in problem determination is to use the system trace records. You may think of this as the z/OS flight recorder.

Various z/OS component and system level traces are active by default. Information about which module was entered and by which transaction is efficiently written to a trace table in memory where it can later be dumped and analyzed. In addition, flexibility is associated with the ability to turn on additional levels of tracing on a component-by-component basis without a restart for supplemental data gathering. The key here is that this can be highly efficient without significant impact on overall performance.

UNIX systems typically do not have tracing enabled by default. In most UNIX systems, tracing can be prohibitive from a performance perspective. Furthermore, tracing is typically activated by specifying a program variable and restarting the application. This is obviously disruptive and unacceptable to z/OS. All this makes UNIX problem determination more difficult, which can result in a less serviceable environment than z/OS.

System cleanup

The issue of system cleanup is also important. Processes and address spaces can hang for any number of reasons. One example is a deadlock situation between applications or in the latches obtained by the UNIX kernel (operating system) when two threads try to acquire system-level locks for resources held by the other.

Part of the fundamental z/OS recovery process is the cleanup of any resources that programs may have acquired. By cleaning up these resources, programs can free memory that was obtained, hardware and software locks can be released, and data sets (files) can be closed.

A program may encounter an error before it has the opportunity to clean up its resources on its own. Or the application may not invoke its own recovery routine. In either case, the system recovery termination manager (RTM) is designed to invoke the End of Task and End of Memory processing to release the locks, latches, and other system resources. In addition, z/OS has included extensive logic to restart canceled UNIX System Services processes and (non-USS) address spaces without requiring re-IPLs.

The UNIX operating system typically does not track or recover from deadlock conditions on the latches. If one or both of these threads are canceled, the problem is not resolved. UNIX still considers those latches and locks held because it does not go through the code to return the resources. All that remains is to reboot.

Security

Another key aspect of system availability is security. Worms and viruses have brought businesses and government computing systems to a halt. The security features of z/OS are designed to thwart unauthorized access. Access control and authentication are used to determine who is accessing the system, and what resources can be made available to them. In addition tools to help detect the possibility of possible attack, such as port scanning and flooding detection, are provided.

There are many functions that are critical to providing a security-rich environment. This includes the encryption of passwords as users sign on to the system. This provides:

- Encryption data being transmitted between servers
- Protection of the system and user files from unauthorized access
- Auditing capabilities
- Protection of the security manager database itself

Many of the these functions are available on UNIX systems by taking overt action. They may possibly require purchasing and installing ISV products or open source applications. As such, many clients may not know that they need to take these actions, possibly resulting in an insecure system. In addition, these may add unacceptable performance bottlenecks to the UNIX system. In the worst case, a security modification or addition can require rebuilding the kernel to install the packages which may invalidate any service contracts with the operating system provider.

z/OS security interfaces to an external security manager, such as RACF®, are designed, architected, and integrated into the product. They help to optimize performance and are generally much easier to set up and manage. In addition, since it is a centralized server, it is designed to be done only once.

Digital certificates can be a basic building block of a trusted infrastructure supporting secure transactions over the Internet. Public Key Infrastructure (PKI) services provide for the life-cycle management of digital certificates. These services include creation (or signing) of digital certificates and renewal after a period of time.

On other platforms, applications that use Digital Certificates may require the customer to go to third-party vendors to purchase and manage certificates which could be costly on a per certificate basis. This may be a significant expense for a client who wants to use digital certificates to establish secure identification of and communication with large numbers of its customers or business partners.

With the PKI services that are built into z/OS, clients have the capability to create and manage their own digital certificates in whatever numbers they choose. Cost effective z/OS PKI services can make digital certificate-based security an option where it may not have been an option.

On distributed systems, system traffic travels everywhere. Often, passwords can travel in the clear within a company's intranet, unless someone took the added effort to encrypt them as noted earlier. Using simple technology, a hacker can tap into a line to steal data and passwords. Hundreds of these distributed systems can be consolidated on a single zSeries server with data being passed to each other inside the box using virtual network technology such as HyperSockets, Guest LANs, and so on. Even between z/OS images on different servers within a Parallel Sysplex, data can be exchanged using the CF links. All this helps protect z/OS and zSeries from physical network attacks.

Part of the Communication Server, a base element of z/OS, is Intrusion Detection Services (IDS). IDS is built within the TCP stack and is designed to discover and defend against attacks such as scans, single packet attacks, and flooding. It automatically logs this activity for later analysis. z/OS and RACF are designed to protect resources by default. You can apply and manage this security across the system to handle access control lists. This applies equally to the standard z/OS files and to the z/OS hierarchical file systems.

Middleware

To help business to be conducted without interruption and according to the corporation's goals, z/OS and its key middleware components work as a team. CICS and IMS take advantage of functions of IBM mainframe hardware, such as Subsystem Storage Protection and Subspace Group Facility, to provide highly available transaction solutions.

DB2 is designed to provide a high performance, richly functional database management system for critical business data. To help clients meet their business goals, and to assign resources where it makes most sense, z/OS includes workload management functions. These functions are policy driven to grant important workloads priority access to key resources so that they can meet the performance goals set by the business. The resources can be CPU or storage that is to be dedicated to this task to meet its goal. Workload Manager (WLM) and Intelligent Resource Director (IRD) continuously monitor the system and adapt to changes in workload and configuration to meet the specified business goals.

Synergy with storage

Let us look more closely at how z/OS, in general, and the WLM component of z/OS exploit the storage capabilities. In this context, we look at the IBM TotalStorage® Enterprise Storage Server® (ESS). ESS can be used with other operating systems as well, but we focus on how z/OS works effectively with ESS.

Three main storage capabilities reduce the client's need for spending time in analyzing their workload and data access patterns:

- Multiple Allegiance (MA)

  This is the capability to support I/O requests from multiple systems, one per system. They are to be concurrently active against the same logical volume if they do not conflict with each other.

- Parallel Access Volume (PAV)

  The operating system allowed only one request at a time for each volume. With PAV, multiple unit addresses are associated with the same logical volume (aliases). These aliases are dynamically managed by WLM as a pool.

- I/O Request Priority (IORP)

  IORP is a means for associating a priority with each I/O request. Multiple applications within an operating system image may issue I/O requests concurrently. The work for each target logical volume is ordered by priority, with the number of requests matching the number of available PAV aliases.

Furthermore, there is no need for the client to spend time in ensuring that each individual volume is configured with enough resources to handle peak loads. You need to plan only for resources at an entire subsystem level.

You can find more information in 3.10, “Disk storage” on page 68. For details, see the article titled “z/OS support for the IBM TotalStorage Enterprise Storage Server” in the *IBM Systems Journal* Vol. 42, No. 2, 2003.

z/OS as a UNIX platform

Since 1993 (with MVS/ESA™ SP V4.3), z/OS has also been a UNIX platform. It has supported file systems (HFS and faster performing zFS), and has the same look and feel as other UNIX platforms from a user and application point of view. Yet, there are differences. UNIX applications can run on z/OS like they can run on other platforms. These applications can also enjoy the added qualities of service already mentioned that z/OS provides, as well as the z/OS Workload Manager and other features.

z/OS can provide availability benefits even for simple functions such as checking or recovering from a bad pointer. Although sigaction() or siglongjmp() UNIX validity checking calls can be of help to check for a bad address, these methods are rarely used. In z/OS, even if the subsystem is given bad data, the recovery routines are designed to protect the operating system, UNIX kernel, and subsystems. On a UNIX platform, a bad pointer may cause an error in the subsystem code which could in turn cause a system dump, requiring a re-IPL.

If an application should get hung, the only UNIX recourse is to end the process using the **`kill`** function. But sometimes this function does not work.

Since the UNIX implementation on z/OS uses z/OS services, you can take an additional recovery step, which is to use the z/OS UNIX **`Superkill`** command. **`Superkill`** invokes z/OS services designed to cancel the address space running the UNIX process. The z/OS End of Task and End of Memory services release system resources that were held. In addition, many UNIX servers running on z/OS are designed on z/OS to clean up after themselves so you can restart them without requiring a re-IPL.

UNIX systems and subsystems usually do not take advantage of specialized hardware because they are written to be general programs, running on standardized UNIX and Linux platforms. DB2 for z/OS is designed to be optimized and tightly integrated with the z/OS operating system and the zSeries platform. This helps provide better performance and the ability to take advantage of platform-specific availability options such as Parallel Sysplex clustering.

Horizontal scalability via a sysplex

The growth of processors within one operating environment is called *vertical scaling*. *Horizontal scaling* is leveraging multiple systems to work together on a common request in parallel. IBM mainframes running z/OS can be linked in a unique clustered environment called a *system complex*, or *sysplex* for short. A sysplex provides a tremendous potential for growth.

In most clusters, each system owns a specific piece of data or a portion of computational responsibility. The sysplex cluster is unique in that the data and computational responsibilities can be shared by all of the systems in the sysplex. This provides two key advantages:

- A near linear scale of compute power
- The ability to survive the outage of a software subsystem without losing the availability of the entire system

A sysplex may also be geographically dispersed. For example, if there is a catastrophic event in one location, the sysplex may be recovered in another location. This enables a business to continue despite a major outage at a given location.

You can find more information about the value of z/OS and Linux on zSeries in the IBM article “Positioning z/OS and Linux for zSeries”, from April 2004. This article is on the Web at:

`http://www.ibm.com/servers/eserver/zseries/library/literature/papers.html`

1.3.2 Evolution of z/OS

OS/360 was introduced in the 1960s. Also known as Principle Control Program (PCP), it was initially only capable of executing one task at a time. This changed with the introduction of multiprogramming with a fixed number of tasks (MFT) supporting initially four and later fifteen consecutive tasks. The fixed limitation was removed with the introduction of multiprogramming with a variable number of tasks (MVT).

In 1968, additional function was added to System/360 to enable two processors to share main storage and therefore work on the same job. This was the first significant advance towards application availability.

By 1972, we saw the introduction of System/370 and virtual storage support through dynamic address translation. Virtual storage gave the illusion that the processor had access to more memory than was really available by using disk storage and paging technology. IBM supported two variants of the operating system, OS/VS1 providing a single virtual storage (SVS) address space and OS/VS2, which was destined to allow multiple address spaces.

The actual support for multiple address spaces did not come until June 1974 with the release of OS/VS2 MVS, the first reference to Multiple Virtual Storage.

The year 1981 heralded the leap to 31-bit addressing and access to 2 GB of virtual storage address space. This came with the IBM 3081 processor and the introduction of Extended Architecture (XA).

This was followed in 1988 by Enterprise Systems Architecture (MVS/ESA) in support of the ES/3090™ processor. A new type of virtual storage space was introduced called the *data space*, which reflected its use for data only and not program code execution.

Next to come was *hiperspace*, another special use for virtual storage that exploited expanded memory. Hiperspace™ data provided a fast movement of data in and out of a regular address space and could not be processed directly. It gave rise to another storage technique called *hiperbatch*. This allowed sequential file data to be manipulated in hiperspace memory to improve batch file processing.

In 1990, IBM introduced the ESA/390 and ES/9000® family of processors and supporting operating systems. OS/390 would see support for open systems under Open Edition services which later became UNIX System Services (USS).

Today the latest zSeries servers support the fully functional 64-bit z/OS, which represents the premier parallel operating system for batch jobs, data sharing, and online transaction processing (OLTP) tasks through transaction processors such as CICS. However, the future expansion of MVS lies with its adaptation to support open systems and accommodate the unique characteristics of on demand workloads.

1.3.3 Other mainframe operating systems

Besides the traditional MVS, four other mainframe operating systems exist:

- z/Virtual Machine (z/VM)

  As a control program, z/VM is a hypervisor because it runs other operating systems in the virtual machines it creates. The control program artificially creates multiple virtual machines from the real hardware resources. To the user, it appears as though they dedicated use of the shared real resources.

- Virtual Storage Extended (VSE)

  VSE is popular with users of smaller mainframe computers. Many of these clients eventually migrate to z/OS when they grow beyond the capabilities of VSE. Compared to z/OS, the VSE operating system provides a smaller, less complex base for batch and transaction processing.

- Linux

  Linux on zSeries and S/390 is capable of inheriting the strengths and reliability features of the hardware while preserving the openness and stability of Linux.

- Transaction Processing Facility (TPF)

  The TPF operating system is a special purpose system. It is used by companies that require high volume transactions, such as credit card and airline companies.

1.4 Processor

When S/360 was first announced, the available hardware at the time was a single central processor. As the business world extended its use and reliance on computing, the demand for processing power outpaced the rate at which CP speed was progressing.

IBM introduced the ability to add a second processor providing more power and potentially more availability. You could conceptually continue processing even if one of your CPs failed. These machines were available either as Attached Processors (APs), where only one CP had an I/O subsystem, and Multiprocessors (MPs), where each CP had access to its own I/O subsystem. The DB2 Sort Assist was provided to favor internal sort operations.

In addition to providing more capacity, the processors, I/O channels, and storage could be physically “partitioned” to enable separate copies of the operating system to be run.

The next major hardware advance, in terms of flexibility for running multiple copies of the operating system, was Processor Resource/System Manager (PR/SM™) with the LPAR feature. PR/SM gave you the ability, even with one CP, to run up to four LPARs.

The latest zSeries servers continue to offer increased performance, resilience, and security with functions such as:

- 64-bit architecture: Platform compatibility and future proofing
- On-chip compression: Releasing computing power for other operations
- Improved chip storage organization: Improves cache hit ratios
- Integrated cryptographic coprocessors: Hardware encryption support
- Up to 30 LPARs each with 64-bit central memory addressability
- Up to 256 GB memory

Note: DB2 UDB for z/OS releases significant performance gains through synergy with the zSeries hardware capabilities.

1.5 Management

The mainframe has traditionally been seen as a resource-intensive platform that needs a large team of highly trained systems programmers to configure and manage the system and its peripherals. This situation has steadily changed over the years with improved management tools, interfaces, and the latest drive toward autonomics.

On the configuration and management front, the recent OS/390 Web-based assistants (wizards) have been joined by a new facility. z/OS Managed System Infrastructure for Setup (MSYS for Setup) employs the same easy interview style as the wizards for defining configuration values. Defaults and best practice values are suggested through a user friendly screen-based interface which simplifies subsystem configuration.

The capability to support multiple mixed workloads is a mainframe strength that is gaining recognition with the re-evaluation of server consolidation over large numbers of distributed dedicated servers. Effective management of the machine workload is critical to ensuring the cost effectiveness of such a strategy.

IBM recognized this issue and has consistently evolved the workload management components of the mainframe platform. With the latest WLM and IRD, the zSeries servers provide unparalleled dynamic prioritization of resources to ensure that business defined service levels are achieved.

Note: Enhanced automated systems management capabilities help to reduce total cost of ownership (TCO). Involvement in mundane activities is minimized enabling experienced staff to be better used on business value tasks. System assisted configuration means skill levels for certain system programming activities are reduced.

What's next?

The IBM autonomic computing initiative helps clients reduce the cost and complexity of their On Demand Business infrastructures still further. It aims to overcome the challenges of systems management using systems that self-optimize, self-configure, self-heal, and self-protect.

The zSeries plays a major role in this initiative since its ground-breaking capabilities function as a model for other IBM platforms. DB2 and Information

Management tools also provide functions which contribute to reducing cost and complexity (see “Autonomic” on page 88, and 8.4.1, “Autonomic computing” on page 195).

Enterprise Workload Manager (eWLM) extends the objectives of WLM to simplify, automate, and offer goal-oriented processing and monitoring, but it applies these goals to a cross-platform environment. End-to-end performance monitoring and management will be enabled across applications, servers, or logical partitions.

1.6 Preparing new mainframers

Executives in business and government may be wondering about satisfying future staffing needs on the mainframe as the current mainframe programmers and administrators reach retirement age. This is consistent with the Mainframe Charter announced by IBM on 22 August 2003, available at:

http://www.ibm.com/servers/eserver/zseries/announce/charter

IBM intends to provide innovation on the mainframe platform and provide skills and expertise for clients who are using this platform. For this to be fruitful, developing mainframe skills in fresh graduates from college is a necessity. This fostering of mainframe education in colleges and universities also resonates with the clients.

Marist College in Poughkeepsie, NY, and Clarkson University in Potsdam, NY, are examples of schools that impart IBM mainframe instruction as part of their curriculum. They have zSeries servers. Clarkson University won their mainframe in an IBM sponsored contest. Several students in these schools also receive focused training in zSeries through an IBM six-month IT Cooperative Education program in Poughkeepsie, where they work on large system projects. You can find more information about the IT co-op program at IBM Poughkeepsie on the Web at:

http://www.developer.ibm.com/us/en/university/scholars/products/pdfs/ITSkillsCo-opWhitePaper.pdf

IBM also has an IBM Scholars zSeries Program that offers colleges and universities educational resources to enhance the zSeries skilled resource base. It also helps students develop practical skills that enable them to find good jobs quickly upon graduation. While there is always room for doing more, this is one way IBM contributes to building budding mainframers.

1.7 2000 through 2005: The era of on demand

Mainframe technology continues to push the envelope of what's possible. IBM research makes the impossible real and eventually commonplace. The traditional batch and OLTP workloads that lay at the foundation of the mainframe success story have not gone away. The demands of the enterprise to support critical business objectives cross multiple application profiles. The technologies has never been greater.

In the visionary world of resource virtualization, the zSeries must quietly coexist and cooperate with any other server. On demand computing offers a new path for growth and IBM is working even more closely with its clients to help solve their problems.

One key message to take from this selective snapshot of mainframe history is that pure power is not everything and the ability to adapt is crucial. The technological world, in which global enterprises operate, demands reliability and security at its core. Competition will only increase and survival will depend on efficiency of operation and minimized cost base.

Note: The emphasis is on total cost per transaction, taking all factors into consideration, but underpinned by a secure, reliable technology framework.

2

The evolution of DB2 for z/OS

This chapter discusses the history of DB2 and z/OS and how they have evolved to reach new levels of availability, performance, and scalability. It covers the history of DB2, from the development of the relational model, to the birth of Structured Query Language (SQL), up to the current version of DB2 which is the premier enterprise relational database management systems (RDBMS).

In the 1960s, large enterprises demanded that their systems had high availability, performance, and scalability. The demand for these characteristics drove the development of database management systems (DBMSs) and operating systems. Databases and the mainframe operating system were developed to meet those needs during that time, and they have both advanced tremendously over the last 40 years.

In today's world, most businesses, be they large or small, have opened up their applications over the Web to business partners and customers. Their data needs to be available 24 hours a day from anywhere in the world, for many different clients, devices, and languages. Online customers expect a quicker response than they experience over a live telephone call. They also expect that the information used during this process remains absolutely secure. This is more than browsing information, when an order is created, it triggers transactions within multiple applications. If the response is not immediate, that customer is one click away from the competitor.

These requirements have again driven DBMSs and operating systems to evolve.

Another key requirement is driven by the need to reduce complexity of operations. That is that the DBMS, the operating system and the hardware must all be able to monitor themselves, be self healing, and correct problems when they happen, or even prevent them. As the amount of data grows exponentially, both for traditional and new data types, these autonomic features address the more mundane tasks. This enables database administrators (DBAs) to spend time on more complex work and to solve business problems.

2.1 The evolution of DB2 UDB for z/OS

In the early 1950s, the disk drive was developed in IBM San Jose, California. In the late 1960s, as large amounts of data were stored on disk, the research focus turned to data management systems.

At that time, network and hierarchical data management systems were in use. CODASYL and Information Management System (IMS) were state of the art technologies for automated banking, accounting, and order processing. This was enabled by the introduction of commercial mainframe computers. These early architectures mixed the physical manipulation of data with its logical manipulation. When the physical location of data changed, such as from one area of a disk to another, applications had to be updated to reference the new location.

Teams of programmers were needed to express queries to extract meaningful information. While such databases could be efficient in handling the specific data and queries for which they were designed, they were inflexible.

Respected IBM Fellow Dr. E.F. (Ted) Codd worked on a new kind of mathematical notation, but it took a while for it to receive the attention it deserved. Programs that would have been five pages long if represented in CODASYL were written by Codd in single lines. These would be queries such as “Find the employees who earn more than their managers.”

Some of the roots of DB2 are in IMS. However the paper *A relational model of data for large shared data banks*, written by Codd in June 1970, was the seminal starting point for SQL/DS™, DB2, and other RDBMSs. This landmark paper was published in the Communications of the Association for Computing Machinery. It proposed a new way to manage data stored in computers, replacing the hierarchical or navigational structure with simple tables containing rows and columns.

“Ted’s basic idea was that relationships between data items should be based on the item’s values, and not on separately specified linking or nesting. This notion greatly simplified the specification of queries and allowed unprecedented flexibility to exploit existing data sets in new ways,” said Don Chamberlin, an IBM research colleague of Codd’s.

The idea of relying only on value-based relationships was a radical concept at that time, and many people were skeptical. They did not believe that machine-made relational queries could perform as well as hand-tuned programs written by expert human navigators.

From Codd’s paper, a team in IBM research started building System R as part of a program of research in the relational model of data. The objective of this system was to handle ad hoc queries (one time) and “canned” queries (executed many times). For the relational model to be accepted, it had to be proven by an industrial-strength implementation. That was the goal of the System R project begun at IBM San Jose Research in 1973.

One team, called *Research Storage System* (RSS), worked on physical data management issues. Out of this team came some early ground breaking work that included developing a data management interface with support for indexes, logging, concurrency, and transactions.

The other, higher level team, was called *relational data system* (RDS). It took the relational calculus way of expressing a query and developed a syntax to express questions. The language was called *Specifying Queries as Relational Expressions* (SQUARE).

A similar and competing project started around this time at U.C. Berkeley. The teams were aware of each other’s intentions and a competition developed which spurred the progress of the relational approach.

Codd believed that computer users should be able to work at a more natural language level. They should not be concerned about the details of where or how the data was stored.

The SQUARE language was extended and enhanced to a more English keyword approach. This made it possible for nonprogrammers to interact with databases. The language became Structured English Query Language (SEQUEL). The original SEQUEL/DML (Data Manipulation Language) paper, written by Don Chamberlin and Ray Boyce, was made public at SIGMOD in May 1974. The paper described the syntax of this language and what it could do. This was put to good use by other companies who later brought other RDBMSs to market.

At the IBM laboratory, other concepts emerged:

- When querying a set of rows, lock the set.
- Instead of authorizing columns of a table, make it a view of the columns and authorize that.
- Transactions, like logical units, should be all or nothing.

If you compile, it will run faster. The idea was to compile rather than to interpret the language. You optimize once and then you package the whole thing. Then you invalidate the *code* if things change that impact the access strategy. It was a good idea, but experts didn't think it could perform. This was the major factor that influenced performance, and performance had to be proven for the relational concept to be accepted. It was cost-based optimization. Different costs were attributed to each possible access path, and the optimal one was chosen.

This was good for canned queries, but what about ad hoc? That challenge was met by providing both an interpreter and a compiler. The compiler handled canned queries and ad hoc queries. Ad hoc query formulation and execution allowed for rapid development and testing.

IMS kept records more organized, and because the relational model does not have a notion of sorting, it was OK to insert the records into the table and then sort them at the time you present the results. The concept of nonprocedural and lots of optimization took hold.

However, rather than making data available to nonprogrammers (eventually brought about via the development of the graphical user interface), System R and SEQUEL had achieved significant breakthroughs:

- They made it possible to embed queries in host languages.
- SEQUEL served as an interchange medium between different systems.
- The notion of data independence separated the physical representation of data from the logical representation presented to applications.

  Data could be moved or stored in a different format without causing applications to be rewritten. Application developers were freed from the tedious physical details of data manipulation and allowed underlying storage structures to change with time.

Some of the other benefits were:

- Nonprocedural requests and cost-based optimization
- Meaningful names for tables and columns stored in user visible system tables
- Authorization checking of all statements at compilation time, with only a single user to program an authorization check at execution time

- Dynamic database definition, allowing tables, users, and transactions to be added or dropped while the system is up
- Locking with multiple concurrent users accessing same database with predictable results

  This was part of the Algorithms for Recovery and Isolation Exploiting Semantics (ARIES) project.
- Transaction management, ability to commit, or a set of statements as a whole

The result was that for the Eagle Transaction 1 (ET1; a debit credit scenario now called TPC/A) benchmark, System R had a path length that was state of the art in performance. No database could run faster than System R on canned lightweight transactions, which is what you would think to be the weak link in a relational database.

The first installation of System R was 1977 at Pratt and Whitney for inventory control of their parts and supplies that they were using to build jet engines. Another joint venture with Upjohn was for the purpose of storing the results of clinical experiments that they were using in support of their drug applications to the Federal Drug Administration (FDA).

There were other users, mainly within IBM, but also many scientific centers all over the world. System R was not big in terms of code size. RDS had 38,000 lines of PL/I and 9000 lines of assembler. RSS was written in PLS and was about 35,000 lines of code. System R was a major influence on other DBMSs.

When the Santa Teresa laboratory was built, the development group was convened to that lab. It worked to build a product which would replace IMS and perform all System R functions. However the team did not want to generate machine code. They wanted to generate something slightly higher level and then interpret that.

One database, VSS, was to be developed for the MVS platform. There were a lot of challenges and discussions around VSS and it was stopped.

Another team was given the task of developing the SQL/DS database for the VM and DOS environments. The SQL/DS project proceeded and was delivered to market in the summer of 1980. SQL/DS was query-like and not really for larger databases. It contained a lot of System R code and is a current product.

Note: Around this time, SEQUEL was abbreviated to SQL because of a trademark legal challenge by the Hawker Sidley aircraft company in Great Britain. It became the industry standard language for describing what data should be retrieved.

When the MVS database project was revived, the code name had evolved many times, from VSS to DS/1 to Eagle (because of the eagle on the poster announcing the Santa Teresa lab). When the team discovered there were no eagles left at the Santa Teresa lab, the project was renamed Ampersand.

When the MVS project finally completed, the database was called DB2. It was not billed as a transaction system at the time. It was considered a decision support system, not ready to take on IMS. System R had a lot of influence on DB2.

DB2 history

Figure 2-1 shows a high level time line of the DB2 releases and some of the functionality delivered. Each release is explained in the list that follows.

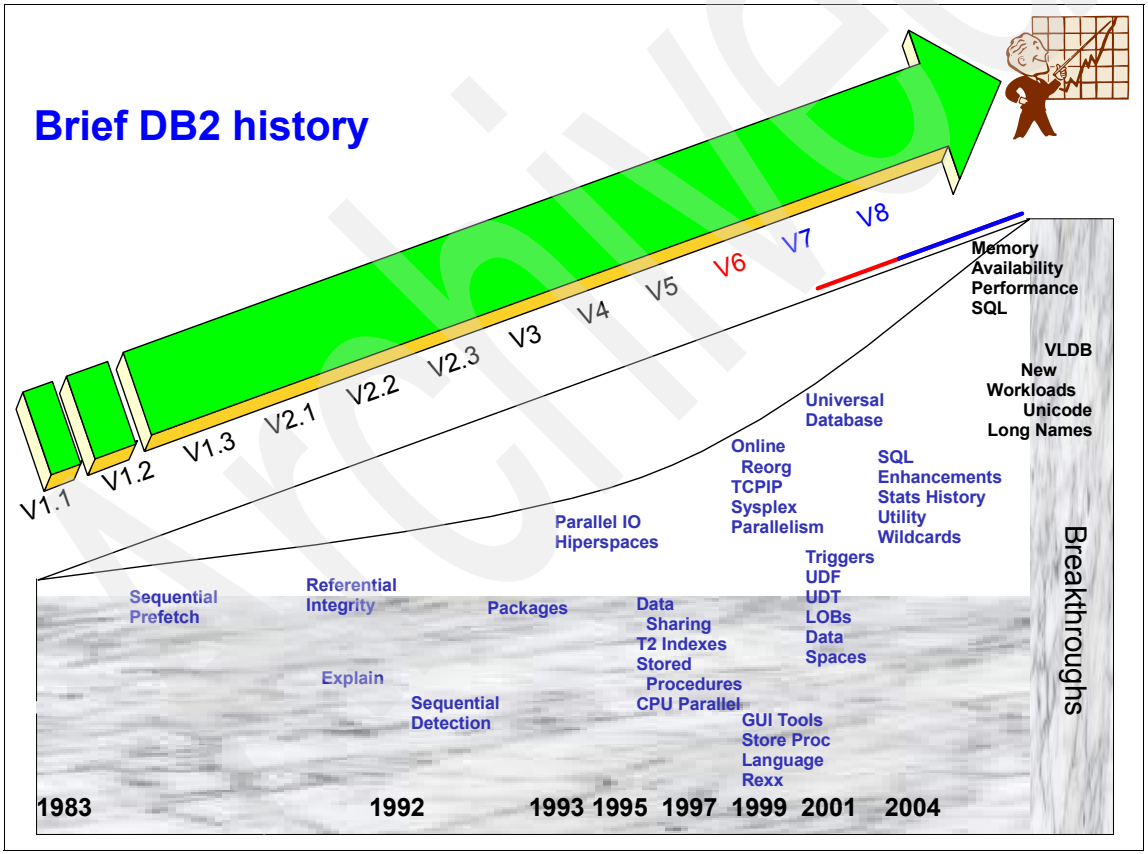

Figure 2-1 DB2 history time line

- Version 1 release 1 was announced on 07 June 1983.

  In 1985, DB2 development concentrated on the relational aspect of doing what a relational database needs to do to make clients successful.

- Version 2 Release 1 was a breakthrough in 1988 that included many performance enhancements and Referential Integrity. This integrity feature allowed the DBMS to enforce relationships based on business data values. This ensured, for example, that prior to adding an order row in a table, DB2 checked that the related customer existed in the parent table.

- Version 2 Release 2 improved query optimization and the private protocol for distributed database (MVS to MVS). This brought better performance and improved developer productivity with the ability to easily query a remote DB2 subsystem.

- Version 2 Release 3 supported the Distributed Relational Database Architecture™ (DRDA). It set up a standard for access to other relational products and DB2 on other platforms. DB2 sort was improved with assisted instructions, and batch run times were improved. IBM initially developed DRDA to allow different RDBMSs to talk to each other. This later became an industry standard.

- Version 3 was a strong release for synergy with S/390. It supported hardware assisted compression, hiper pools, and flexibility with buffer pools. This was a significant improvement over software compression and allowed more data to be stored in less space. Additional pools kept more data in memory. Lastly the enablement of I/O parallelism dramatically improved the elapsed time for some queries.

- Version 4 was the second breakthrough release, with locking improvements and new indexes. The effect was that more transactions could execute concurrently as locks were better managed. A new type of index brought more speed and no index locking.

  Data sharing provided a new dimension for scalability and availability. This tremendous advancement was unique and enabled upgrades without outages and the ability to add DB2 subsystems on the fly to a data sharing group. Version 4 also had major performance improvements for SQL and utilities as well as central processor (CP) parallelism. This was another step in better parallelism for long running queries.

- Version 5 was the major release for e-business, business intelligence, sysplex query parallelism.

  Tremendous scalability was achieved with DB2 and OS/390 as a robust solution for e-business. DB2 has evolved over the years by increasing the degree of parallelism in its database processing. This level enabled queries to run across the data sharing group, a major improvement to all query processing.

- Version 6 continued the theme of e-business, business intelligence, and enterprise applications

  Utility performance improved allowing you to load and reorganize your data faster and more efficiently. Scalability, large objects, and object-related function were provided. Online reorganization enabled you to reorganize your data while allowing business as usual access against your database.

  Triggers allowed you to encapsulate business rules within the database. This improved performance and ensured consistency as opposed to maintaining similar logic in multiple application programs.

  After general availability of Version 6, more key new functions were provided. These included identity columns, savepoints, declared temporary tables, and many performance enhancements.

  The star join method provided better performance for warehouse queries against your fact and dimension tables.

- Version 7 focused on utilities and provided new function, more availability, improved performance and improved usability.

  The new DB2 utilities deliver improved and powerful functions, enabling your business data to be better maintained and available. Several system parameters are changeable online, and some outages are eliminated.

  SQL improvements are substantial with scrollable cursors and union changes. Unicode opened new possibilities for international e-commerce. The Extensible Markup Language (XML) extender provided support for a new emerging standard and allowed you to store, retrieve, and search XML documents in a DB2 database.

- Version 8 is a breakthrough release, with fundamental changes in many areas. Much of DB2 was re-engineered to help clients grow beyond current limits and enable open systems On Demand Business through support for advanced Java and Unicode functionality.

  Online schema evolution allows structural changes to occur without drop and create, providing increased availability for your business. You can add or change columns on the fly, and you can add partitions in seconds.

  Version 8 offers partitioning improvements. This extends support for very large databases and increases the amount of data you can store. Improvements are made which benefit enterprise resource planning (ERP) applications such as SAP, PeopleSoft, and Siebel. One example is materialized query tables, that are automatically used without code changes. This can dramatically improve the performance of your complex decision support queries.

  Integration with IBM WebSphere and Java are key enhancements that allow easier porting of applications. There are major SQL enhancements including

support for larger table and columns names, and statements 2 MB long. Version 8 brings incremental scalability with 64-bit virtual addressing. It can use more storage more effectively, with less complexity. This reduces system constraints as more applications and users are added.

In addition, Version 8 provides the following enhancements:

- *Autonomic space management* of extents
- Significant indexing and optimization improvements
- New options for high security with *multilevel security* and *row level security*, a first in the family
- A solid foundation for your On Demand Business infrastructure

DB2 functions within an open and scalable architecture and continues to grow well. The growth comes from current applications that scale up, from the growth of many vendor applications, and from new customers.

Traditional OLTP has changed dramatically. Batch work has increased in volume and needs to be concurrent with other work. Applications have become more complex. OLTP has evolved into 24-hour *On Demand Business*.

DB2 had transitioned from databases acting as an abstract data management layer to being high performance, enabling high volume query processing across multiple engines.

The DB2 optimizer team has continued to deliver improvements that are more than a factor of 10 times faster for a significant fraction of the queries.

DB2 SQL performance improvements continue with stronger optimization. The optimization has been cost-based since the beginning, in System R. This cycle has improved use of indexes, reduced processing times, and increased use of parallelism. Star schema optimization, the ability to use multi-column and multiple indexes for access, materialized query tables, sparse indexes, improved buffer management, and more sophisticated query rewrite continue to deliver substantial improvements in every release.

In DB2, IBM has a state-of-the-art relational database management product. It has near scalable growth within the Parallel Sysplex clustering architecture.

For an introduction to future enhancements, see Chapter 8, “The future of DB2 for z/OS” on page 187.

2.2 DB2 and Parallel Sysplex industry leadership

This section describes significant features which are unique in the industry and offer tremendous benefit and cost savings.

2.2.1 Extreme availability

The highest levels of availability are obtained with DB2 and the Parallel Sysplex on the zSeries architecture. The Parallel Sysplex is built with failure prevention and availability as the foundational elements. It is designed from the ground up with recovery routines for every major component. This allows the system to recover and continue operation without interruption, even in the event of operating system function error.

The zSeries philosophy is to prevent failure. The mean time between failures is measured in decades. The more errors can be prevented at the lowest level (at the hardware and microcode levels), the less impact there is on applications and end users.

DB2 is also built for availability. Critical data sets and logs are duplexed. The granularity of utilities allows a portion of a table to be maintained while the rest of the data is available.

Online schema evolution is needed for business continuity and is a first in the industry. It means that you can make more and more structural table changes while the data is available for business as usual. For example, you can alter data type and length of an existing column. You can also add a partition dynamically while allowing access to the other partitions.

IBM realizes that your business does not stop and has adopted this goal of change on the fly. DB2 is aware of the new and prior versions of a table. This has considerable real benefits for the availability of your data. For example, planned outages are reduced. Changes which took hours to do in prior versions are handled in a matter of minutes. This enables a dramatic increase in availability of your business applications. For more information about this, see “Online schema evolution” on page 91.

When a server or DB2 member is unable to function, work is automatically and dynamically rerouted to other components. The failing components autonomically attempt recovery. See 3.2, “Parallel Sysplex and data sharing” on page 43.

DB2 and z/Architecture lead the industry in *autonomics*. For a description of this significant real-time characteristic, see 8.4.1, “Autonomic computing” on page 195.

Intelligent Resource Director and Workload Manager continuously track activity and direct work and resources in a manner which ensures that business service-level agreements are met. Workload throughput is maximized and “white space” is minimized.

This environment is mature and sophisticated. Extensive autonomic features of the operating system ensure that resources are continually monitored. This is about the server being up. It is about the application and the data and the connection to disk being up. This environment offers extreme availability. New high priority work can be added and immediately serviced adequately at the expense of lowest priority workloads. This is truly unique in the industry. It is the substance that ensures data is available on demand. You can learn more in 3.9, “WorkLoad Manager, Intelligent Resource Director, and VIPA” on page 63.

2.2.2 Near-linear scalability

DB2 in a Parallel Sysplex configuration offers tremendous near-linear scalability. You can scale vertically within a server or horizontally across the group. Within the zSeries architecture, the cost of adding a member to an existing group is minor. For more information about sysplex timers and the coupling facility, both key components of zSeries, see 3.2, “Parallel Sysplex and data sharing” on page 43, and 3.2.6, “Summary” on page 48.

2.2.3 Preservation of investment

IBM works hard a preserving the investment that you have in you application software and database structures. See 3.11, “zSeries 990” on page 75.

Backward compatibility

When you upgrade from one version to another, the design strategy is to ensure that existing code executes successfully without any changes. The few necessary incompatibilities that exist are highlighted early. IBM identifies and describes these in a timely manner allowing you to proactively research and adjust where necessary. Where practical, IBM provides automated scripts to locate incompatibilities.

Proactive notification

IBM has provided ample notification of upcoming changes which demand change in database structure or application coding. For example the deprecation of type 1 indexes was announced at least two versions before it took effect. Similarly there have been numerous announcements about the future demise of private protocol. In both cases, the alternatives, type 2 indexes and DRDA protocol, enabled a much higher level of usability and performance.

Regression testing

With DB2 Version 8, IBM significantly increased the amount of testing that is done. Stress testing time was almost doubled and included multiple concurrent mixed workloads. Duration of early Quality Partnership Program (QPP) testing was also doubled. This included using new functions and regression testing. QPP had cross industry/world-wide customer representation in manufacturing, banking, securities/investment, retail, distribution, and health care. Real customer catalogs were converted. Further stress testing was done on a 16-way zSeries 990 (z990) machine.

IBM understands the investment that businesses have made in DB2. IBM has raised the bar on testing and is committed to sustaining that investment level. IBM is serious about shipping new DB2 versions when they are ready.

2.2.4 Controlled enablement of new functionality on version upgrade

In the past, when first putting a new DB2 version in development, it was difficult to prevent application developers from using the new features during the transition phase. If such code slipped through and was delivered to production, it would have failed.

Now in Version 8, there is a formalization of what takes place. The first phase for Version 8 is the compatibility mode (CM). During this phase, no new function is possible. However, fallback to Version 7 is possible. When you decide to use new functionality, you must run an enabling-new-function mode (ENFM) job. When this job is started, there is no returning to compatibility mode. Then everything is available, and you switch to new-function mode (NFM).

This migration approach is tighter and gives you more control. IBM does not have to devote resources to complex fallback coding. This process protects you from surprises and ensures that new functionality is exercised when you are ready.

2.2.5 Rolling maintenance and new releases without an outage

In a sysplex configuration, DB2 maintenance or a version upgrade to a data sharing group can be done as a windowless install. You can temporarily redirect work away from a member while you upgrade it. Members of a data sharing group can coexist. For example, there can be three members at V8 compatibility mode and one member still at Version 7. You can do a rolling upgrade without impacting your workload. This is a considerable advantage and allows your business be continuously available during the upgrade.

Attention: For DB2 Version 8, to benefit from reduced global lock contention, you need a group shut down to implement Locking Protocol 2. This is a radical change that provides reduced locking contention and improved lock management. You can do this group down whenever you are ready to improve group buffer pool locking performance.

2.2.6 Ability to fallback

IBM has developed a controlled process for migrating from Version 7 to 8. As you progress through the compatibility, enabling-new-function, and new-function modes, you have control of the process. The door is open for you to fallback up until the start of the ENFM. Keep coexistence as short as possible and stay in ENFM for a short time. We recommend that the duration be short between turning on NFM in development and that you turn on NFM in production.

2.2.7 Conclusion

Nothing beats the availability of DB2 in a Parallel Sysplex configuration. The resultant availability and scalability are unmatched in the industry.

2.3 Summary for DB2

Dr. E.F. (Ted) Codd invented the relational model and made database management into a science. He described data in a meaningful way and worked persistently within IBM to gain acceptance of the relational model. This model became a compelling paradigm, which spurred the advancement of the database industry.

Among other awards, Codd was designated an IBM Fellow in 1976, the highest technical honor within IBM. In 1981, he received the Turing Award, the foremost technical recognition in the computing profession. His remarkable vision and intellectual genius ushered in a whole new realm of innovation that has shaped the world of technology today.

The use of DB2 has changed dramatically. Originally, DB2 focused on managing structured data, rows and columns of numbers and letters. The initial query and decision support applications were small and simple, but they have grown from megabytes to gigabytes to terabytes. These applications continue to increase in size and complexity.

Most of the innumerable data transactions we routinely make today, using bank accounts and credit cards, trading stock, making travel reservations and

participating in online auctions, use relational databases based on the abstract and sophisticated mathematical theory developed by Codd. Now able to support data on handhelds to supercomputers, the IBM DB2 family of databases handles billions of transactions per day.

Today the DB2 family of products (see Figure 2-2) offers an integrated end-to-end solution. It covers all platforms and includes existing data, from which clients can draw large benefits.

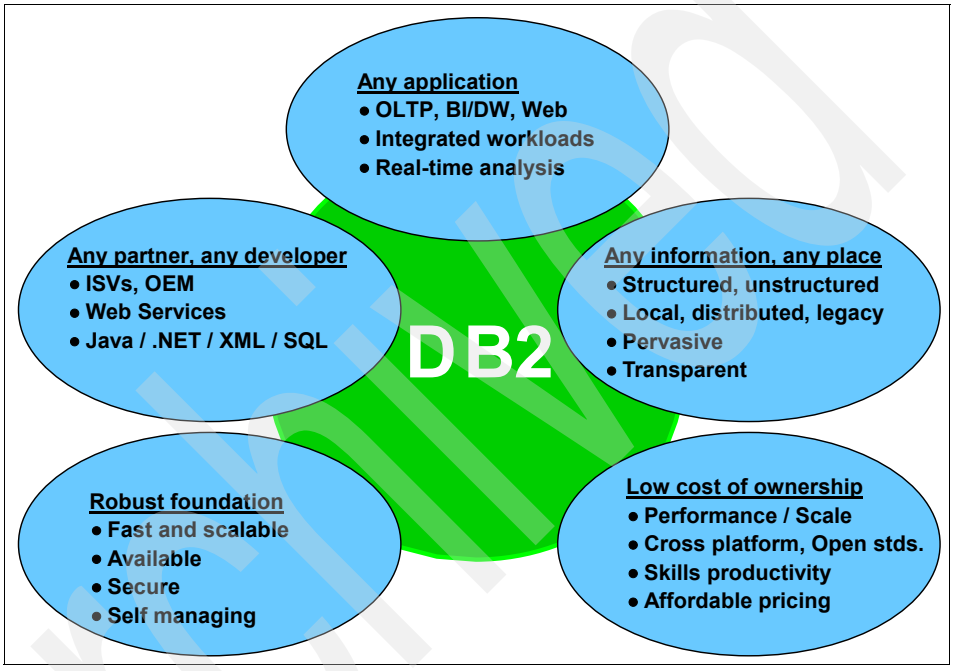

Figure 2-2 The DB2 vision

DB2 has been enhanced to handle different data types (extensibility) such as image, text, video, large objects and XML. It does work in parallel, handles distributed requests, and provides easy access to distributed heterogeneous data (federation via WebSphere Information Integrator).

IBM has partnered with many vendors that are growing rapidly, such as SAP, PeopleSoft and Siebel. These applications manage the enterprise and customer relationships. To succeed with enterprise applications, the DBMS must be able to provide solid support for transactions and for business intelligence.

DB2 provides quick, secure, and reliable access to terabytes of data. Meanwhile it supports the modern languages and interfaces which are part of how business

is done in this new era when the customers or partners need access to data from anywhere at anytime.

Figure 2-3 summarizes the family of DB2 database servers and shows the relative positioning.

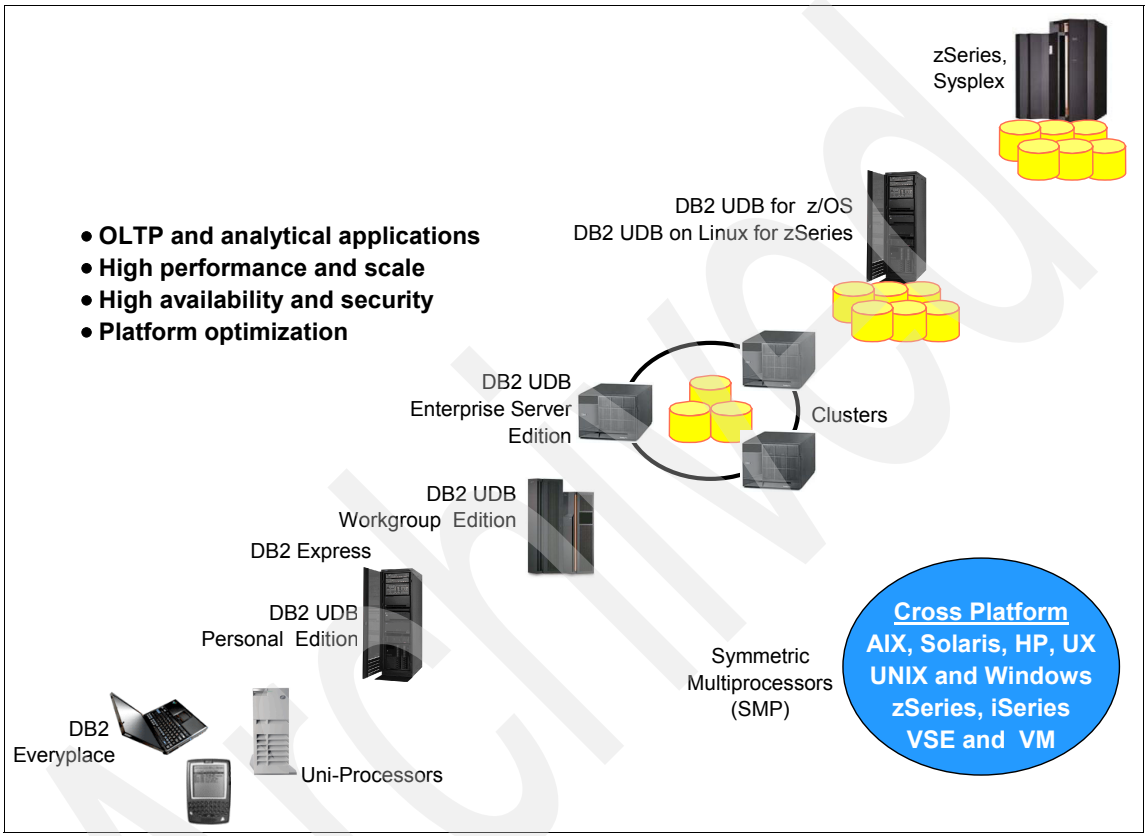

Figure 2-3 The family of DB2 database servers

Fundamental to the IBM database strategy is the ability to easily manage, extend, and exploit the database. This is provided by Information Management Tools. IBM tools exploit the features of new releases of DB2 immediately at general availability time. Their long-term goal is to provide expert advice and automatic management features to enable DB2 environments on all platforms to be easier to manage. The goal is also to require less administrative effort and less expertise to achieve outstanding performance and results.

Tools are a strategic initiative for IBM. The investments have increased every year since 2000. New tools, releases, versions, and features are added to keep

up with the demands of On Demand Business and mixed workloads while offering availability, scalability, and recoverability.

DB2 UDB for z/OS is the platform of choice when the client requires the highest level of availability, unparalleled scalability, consolidation of workload, integration with tools and other applications, and flexible growth. As clients process increasingly large amounts of data and more complex queries, DB2 UDB for z/OS delivers enhancements to keep pace with their needs. This can be for important applications of the top 100 companies worldwide or for the billing and accounting system of a small-to-medium-sized business.

DB2 and the Parallel Sysplex provide the industry's absolute highest levels of availability. This benchmark is made possible by 25 years of zSeries hardware and software evolution built upon a proactive approach to availability.

DB2 and the zSeries architecture are mature. DB2 DBAs on the mainframe can support three or four more times the data than other platforms.

DB2 delivers fast access to usable information in a secure manner. It can scale and continue to provide fast access. It has significant and extensive disaster recovery features. And, it works in conjunction with an operating system that has industry leading availability.

DB2 for z/OS manages access by many concurrent users to large volumes of data in a secure and extremely available manner. Its environment is built on open standards.

DB2 is available for On Demand Business. It is poised to meet the challenges ahead such as digitizing nontraditional data and the increasing demands for security and controls for audit and regulatory purposes.

The close relationship between DB2, z/OS, and the z990 platform is unique in the industry. It delivers world-class performance, scalability, and availability. This benefits all sizes of businesses including the top companies in the world today running their important applications. This tight partnership is a strategic ongoing IBM goal and a significant differentiator.

3

DB2 synergy with zSeries and z/OS

The zSeries environment maintains its benchmark position as a flexible, efficient and responsive platform for highly complex, integrated environments running a wide range of important workloads. This chapter discusses the zSeries hardware and software features which DB2 uses to achieve synergy with the platform as well as the z/Architecture. It describes the significance of this from a business point of view.

zSeries hardware, z/OS and DB2 have a unique partnership in the database industry, one which is key to IBM today and in the future. No other database management system (DBMS) vendor also owns the platform architecture and operating system. Nothing beats the availability of DB2 on zSeries hardware and software. This chapter also discusses how this partnership enables a business to respond far more dynamically to the demands of an always-on marketplace and economy.

3.1 The zSeries, z/OS, and DB2 partnership

If you need an industrial strength DBMS with superior quality of service (QoS) characteristics and truly on demand scalability and workload management, consider DB2 and the z/Architecture.

3.1.1 z/OS highlights

z/OS is the flagship IBM mainframe operating system. z/OS is the robust zSeries mainframe operating system designed to meet the demanding QoS requirements for On Demand Business. Here are some highlights:

- It provides a highly secure, scalable, high-performance base for on demand applications.
- It can simplify IT infrastructure by allowing the integration of applications in a single z/OS image.
- It takes advantage of the latest open software technologies to extend existing applications and add new on demand applications.
- It incorporates world-class optimization features, security services, distributed print services, storage management, and Parallel Sysplex availability.

z/OS Java execution environment

The zSeries Application Assist Processor (zAAP) is available on zSeries 890 (z890) and zSeries 990 (z990) servers. The zAAP is an attractively priced specialized processing unit. It provides an economical Java execution environment for clients who desire the traditional qualities of service and the integration advantages of z/OS.

When configured with general purpose processors with a z/OS partition, zAAPs may help increase general purpose processor productivity. They may also contribute to lowering the overall cost of computing for z/OS Java technology-based applications. For more information, see 3.14, “zSeries Application Assist Processor for Java” on page 80.

z/OS.e on the zSeries 800

z/OS.e is a specially priced offering of z/OS. It provides select functions for the zSeries 800 (z800) and z890 at an exceptional price.

z/OS.e is specifically configured for today’s hot applications, such as WebSphere, SAP, Lotus® Domino®, Siebel, Oracle, SAS, BEA or WebLogic. It is also configurable for any Web-based data transaction processing application using Java, Java 2 Platform Enterprise Edition (J2EE), or C/C++, giving these workloads the price/performance that clients expect. Also, z/OS.e is

complementary to those clients who want to run Linux applications. For example, z/OS.e can act as a database server for Linux application servers.

z/OS.e uses the same code base as z/OS. It invokes an operating environment that is comparable to z/OS in service, management, reporting, and reliability. Also, z/OS.e invokes the same z800 and z890 hardware functionality that z/OS does. No new z/OS skills and service procedures are required for z/OS.e. With few exceptions, z/OS.e has the same features and functionality as z/OS.

This offering does not support certain workloads and languages, such as traditional CICS, IMS, COBOL, and FORTRAN.

At only a fraction of the cost of z/OS, z/OS.e makes it easier to run new workloads on the mainframe.

On demand infrastructure: z/OS leadership

z/OS has optimization features to help provide the responsiveness features needed for on demand applications. The z/OS Workload Manager (WLM) is at the heart of z/OS optimization. It is designed to manage the priority of mixed workloads based on business policies, defined in service-level agreement (SLA) terms.

Intelligent Resource Director (IRD) extends the z/OS Workload Manager to work with PR/SM on zSeries servers with features to dynamically manage resources across multiple logical partitions (LPARs). Based on business goals, WLM is designed to adjust processor capacity, channel paths, and input/output (I/O) requests in real time across LPARs without human intervention.

Together with the zSeries servers, and with interoperability with Linux for zSeries and z/VM, z/OS can play a critical role in simplifying a client's infrastructure. Both z/OS and Linux for zSeries support much of the server differentiation that sets the zSeries apart from other servers. With zSeries servers, z/OS provides the base for the z/Architecture with support for 64-bit storage, IRD, and HyperSockets (for interpartition communications).

The high availability of z/OS is a key reason why so many clients rely on z/OS for their most critical applications. To continue to improve this availability, z/OS provides automation capabilities in a new element, *Managed System Infrastructure for Operations*. This element provides automation for single system and sysplex operations to help simplify operations and improve availability. MSYS for Operations plays an important role in outage avoidance.

z/OS is the foundation of the future of zSeries, an integral part of the z/Architecture designed and developed to quickly respond to the exigent QoS requirements for the on demand era.

3.1.2 zSeries architecture highlights

IBM provides world-class zSeries server technology to help today's enterprises respond to business conditions quickly and with flexibility. From automation to advanced virtualization technologies and open industry standards, zSeries servers help deliver competitive advantages for the On Demand Business.

The zSeries architecture provides innovative self healing features built into every zSeries server. The features are designed to manage and heal themselves and dynamically adapt to change. They help you to keep your system running at peak performance, reducing administration costs and letting you focus on running your business rather than managing your IT infrastructure.

The concept of servers that take care of themselves is modeled after the autonomic function of the human nervous system. Just as the human body automatically breathes, digests, and fends off viruses, a zSeries server is designed to manage, repair, and protect itself. Because the zSeries servers incorporate technologies that have been developed over decades of research and fine-tuning, they are leaders in autonomic computing abilities.

On/Off Capacity on Demand is designed to provide greater flexibility in managing the constantly changing workload demands of today's environment.

zSeries servers are among the most secure servers on the market, with mean time between failures (MTBF) measured in decades. In fact, the zSeries is designed for up to 99.999% availability with Parallel Sysplex clustering. The zSeries is designed to provide superior qualities of service to help support high volume, transaction-driven applications, and other critical processes. It supplies tremendous power and throughput for information-intensive computing requirements.

z/OS architecture is designed to support a failure in some of its components without affecting the others. This means that a zSeries server can survive a central processor (CP) failure without loss of availability of the system. Also a middleware software subsystem outage can be survived without loss of availability of the system.

3.1.3 DB2 UDB for z/OS

DB2 takes advantage of joint development and possible synergy of zSeries, z/OS, and DB2. DB2 is engineered from the ground up to use many features for high availability and continuous operations that are provided by z/OS and zSeries hardware.

How does DB2 take advantage of zSeries hardware and software? This is explained in the following sections. These sections review how DB2 for z/OS

uses the key components of the zSeries hardware and z/OS software to achieve higher performance and industry leading availability. This combination also implies lower people costs than other environments.

3.2 Parallel Sysplex and data sharing

Today's important and high performing applications need uninterrupted service to clients. Whether this is a previously existing billing application or a more recent round-the-clock Web store-front, a Parallel Sysplex and DB2 data sharing implementation offer DB2 applications and administrators the highest potential availability necessary to meet your business goals and your customers' expectations.

Within the z/OS environment, you can react to unexpected business demands by dynamically adding CPs to maintain performance and availability of your business applications. This is known as *vertical scaling*. Vertical scaling enables more efficient use of a machine's processing power.

A Parallel Sysplex provides the foundation for continuously available, clustered environments on the zSeries platform. Figure 3-1 shows two zSeries servers that are clustered together in the Parallel Sysplex. These can be LPARs within a single server or a combination of the two configurations. The two LPARs are members in the Parallel Sysplex, using the zSeries hardware and software for efficient system clustering.

One of the strengths of Parallel Sysplex technology is that every server in a Parallel Sysplex has access to all data resources and an application can run on any server. The Parallel Sysplex provides a "don't go down" environment versus "fast failover" with added multiple separate products on the distributed platform. It enables high levels of multi-programming versus specialized under-used servers. New work is routed to the most appropriate server in a cluster based on priorities and workload balancing, processor and I/O resources are moved to the work automatically and dynamically by IRD and WLM.

A Parallel Sysplex cluster is designed for continuous availability and performance. Its benefits include:

- A single system image, accessing multiple small or large databases
- Shared everything
- Dynamic workload management
- Sophisticated systems management
- Reduction of planned downtime, which represent 90% of outages

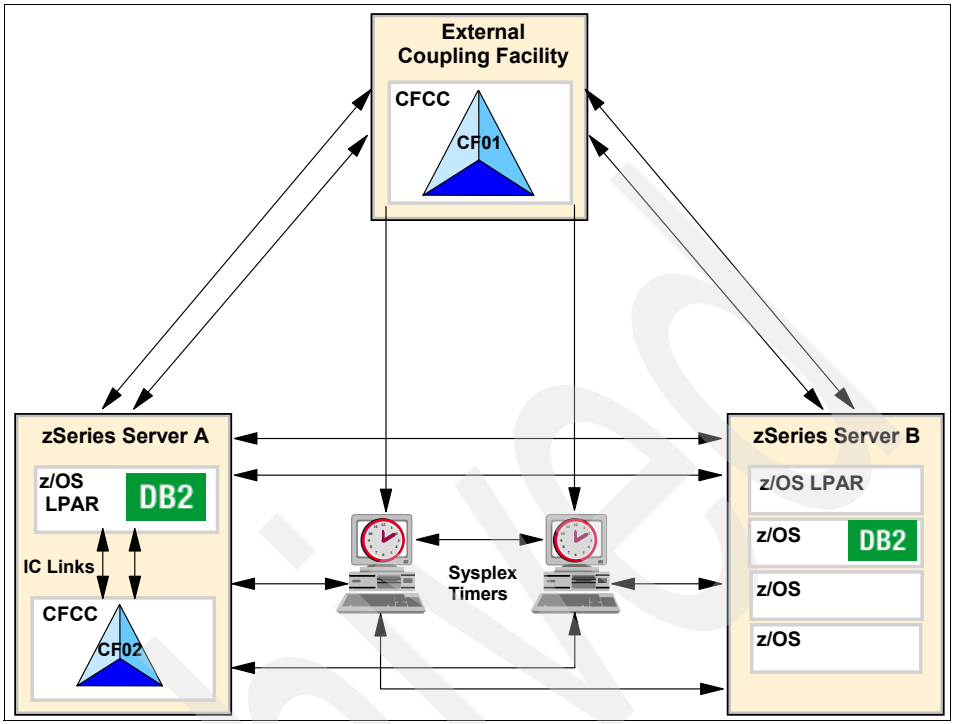

Figure 3-1 Parallel Sysplex components

There are several architectural choices when implementing a clustered solution to enhance database high availability and workload balancing. Some popular options are presented in the following sections.

3.2.1 Shared nothing

Several nodes or servers are clustered together by a communications layer. Each node has its own portion of the data, indexes, logs, and locks. Through the communications between the nodes of the database and a single database catalog, the user or application has a single-system view of the partitioned database. This configuration has these requirements:

- May require functions, data, or both to be shipped across nodes
- Scales well with high parallelism
- Requires careful workload management planning, with data accessible to only one node
- Requires attention to data placement to avoid uneven distribution
- Implies that work usually requires coordination across nodes
- Means that in a failover event, data is unavailable until an instance restart

3.2.2 Shared disk

The data is still viewed as a single object by the application or user. However, this time, the data can be accessed by any of the nodes. In this configuration, note the following factors:

- There is no shipping of functions or data between nodes.
- A piece of the workload can go to any node.
- Nodes can be added without redistribution of data.
- Locking and buffer invalidation require messaging between nodes which can grow exponentially when nodes are added. Scalability is not linear.

3.2.3 Shared data

A clustered DBMS architecture is essential for eliminating single points of failure and thus providing higher availability. And a good clustering architecture also allows for higher scalability by allowing access to the databases from multiple computers instead of just one. However, to gain these benefits, a performance overhead must be paid for the additional coordinating of activities between the nodes of the DBMS cluster.

IBM solved shared-disk scalability problems by implementing global locking and a shared-memory disk cache into dedicated hardware, its Sysplex Coupling Facility. It manages the locks with low overhead and high performance. Data sharing allows an application to run on one or more DB2 subsystems in a Parallel Sysplex environment. These DB2 subsystems are known as members of a data sharing group.

IBM also introduced high-speed links. These are memory-to-memory or physical links that connect LPARs and coupling facilities. They provide good performance with superior workload management and benchmark level fault tolerance.

The coupling facility, along with sysplex timers and high-speed links, enable industry leading levels of reliability, availability, and scalability (RAS) for your DB2 applications. We recommend that you duplex the coupling facility and sysplex timer. The three structures in the coupling facility (cache, List, and locking) are used by data sharing capable participants, such as DB2 for z/OS and OS/390.

Parallel Sysplex is a key example of the synergy of DB2 and zSeries. DB2 takes advantage of the zSeries Parallel Sysplex, with its superior processing capabilities. By allowing two or more processors to share the same data, you can:

- Maximize performance while minimizing cost
- Improve system availability and concurrency
- Expand system capacity
- Configure your system environment more flexibly

With data sharing, applications that run on more than one DB2 subsystem can read from and write to the same set of user databases concurrently and share the same catalog.

With a single subsystem, you can grow your system incrementally by adding additional central processors, power, and memory (vertical scaling). With data sharing, you can add processors, power, memory, and DB2 (horizontal scaling) to the data sharing group. You don't have to move part of the workload onto another system. This alleviates the need to manage copies of the data or to use distributed processing to access the data. You can configure your environment flexibly. For example, you can tailor each z/OS image to meet the requirements for the user set on that image. For processing that occurs during peak workload periods, you can bring up a dormant DB2 member to help process the work.

In the event that a DB2 image goes down, the other "surviving" members remain up and running. Work is automatically and dynamically routed away from the downed DB2 member. z/OS Automatic Restart manager can automatically restart failed DB2 members. You have near continuous availability.

3.2.4 Advantages of data sharing

Data sharing is a horizontal scaling option that has been an integral part of DB2 since Version 4. Up to 32 subsystems can be clustered together, sharing the same catalog, directory, and multiple user databases.

With the necessary links implemented, a structure can either be duplexed for immediate fail over or rebuilt via one of the connected LPARs. In the case of a rebuild, a momentary pause may, or may not, be detectable by an application or user. Because critical DB2 data sharing items, such as locks, are externalized to the coupling facility, data integrity is maintained in the event of member failures.

The coupling facility is also used as a big cache for DB2. This improves performance by avoiding large amounts of I/O.

One of the advantages of DB2 data sharing is to provide additional capacity without changing applications or copy data. If you exhaust the resources of a single DB2 subsystem, without data sharing, you need to split your data across multiple subsystems and provide logic to determine which subsystems to access for which data. With data sharing, you can add another DB2 member to the group and use the same applications and data. Even if you do not have capacity issues that preclude you from running your entire DB2 workload on one subsystem, consider DB2 data sharing for the availability improvements provided by a properly configured Parallel Sysplex environment.

Data sharing provides the following availability features that are not provided when running a single DB2 subsystem:

- Multiple DB2 subsystems, called members of a data sharing group, that can provide failover when one of the subsystems is unavailable because of an unplanned outage; this means increased availability
- Distribution of your workload over multiple members to reduce the CPU workload on any one member, without maintaining multiple copies of the data
- Application of maintenance to one of the members without an outage to your production applications, which can continue to access DB2 data from the other members

  Then the maintenance can be subsequently rolled across all members in a staggered manner.
- Migration to a new version of DB2 in a rolling fashion to maintain availability to your data and allowing coexistence of two versions
- Leaves application interfaces unchanged

  Your investment in people and skills is protected because existing SQL interfaces and attachments remain intact when sharing data.
- Can exploit sysplex query parallelism
- Data integrity maintained in the event of a DB2 member failure

The synergy with the z/OS operating system and the zSeries hardware provides for a clustered implementation. This achieves the key benefits of alternative shared nothing and shared disk approaches without suffering the pitfalls.

3.2.5 Cost benefit

Since the relationship between a DB2 data sharing member and the other members is via the coupling facility, adding members does not increase the overhead of inter-node communications on any other member. When implemented in a two-way configuration (two DB2 subsystems, as members in a Parallel Sysplex group), incorporating new members adds minimal overhead. There is approximately a 5% to 15% CPU overhead in moving into a two-way data sharing group.

Beyond two members, each additional member adds a minor incremental cost of usually between .5% and 1%. Other shared disks implementations continually add incremental costs similar to the initial overhead. There are no additional license costs when adding members to a data sharing group.

3.2.6 Summary

After decades of engineering in pursuit of high availability and continuous operations, the result is a robust and tightly integrated solution delivering high levels of availability that continue to be unmatched in the industry. If any part goes down, recovery is on the fly and fast.

One of the single biggest growth areas is applications that scale up to a certain point on the UNIX platforms before clients decide that they need to move to zSeries for better availability and scalability. The zSeries platform has a significant technological advantage, especially when you consider the Parallel Sysplex.

DB2 with zSeries and z/OS has fantastic scalability and availability. It relies on Parallel Sysplex technology and a coupling facility for global locking to provide near linear scalability with nominal overhead. You can add and remove DB2 members without changing configurations. DB2 adapts automatically to these capacity changes.

To meet your business demands and make the best use of your computing resources, you can scale vertically within a DB2 member or horizontally across DB2 members. Through the shared everything design, you have near-continuous availability for your business applications, which gives you a competitive advantage, allowing you to grow your business on demand.

3.3 Compression

Your data is growing exponentially. More and more data is required for historical analysis, data warehousing, and regulatory needs. You need to keep information for a prolonged period of time. Data is usually copied many times. It is difficult to plan direct access storage devices (DASD) growth, and disk storage is often the most expensive component in data center costs.

zSeries hardware processors offer even better performing hardware compression. You can reduce storage cost and archive even larger amounts of data.

DB2 Version 8 increases the number of partitions, allowing you to store much more data in a partitioned table, up to 128 TB. The upper limit of 4096 on the number of partitions means that for time series data, you can add a partition per day for more than 10 years. You don't have to pre-allocate these partitions. DB2 Version 8 lets you add partitions dynamically, as part of the *online schema evolution strategy*.

Compression combined with a higher number of possible partitions allows you to store more business data than ever before in a single table.

When you compress data, bit strings that occur frequently are replaced by shorter strings. Information about the mapping of bit strings to their replacements is stored in a compression dictionary. Computer processing is required to compress data before it is stored and to decompress the data when it is retrieved from storage.

Data compression allows you to:

- Store more data on disk and use less storage
- Save time retrieving the data since fewer disk I/O operations are required
- Store less data on the log
- Have higher in memory hits when accessing data because more compressed rows are in the buffer
- Reduce disk storage costs

Data compression can be implemented either via hardware or software. Because compression techniques have improved and become broadly accepted, most hardware vendors have implemented some of the widely used compression algorithms in hardware.

There is a price to pay for the CPU compression and decompression. However, for I/O bound access, the elapsed time improvement far outweighs the extra CPU cost. Since the buffer pools can be made more effective and avoid I/O, some CPU time and elapsed time can be saved. Figure 3-2 shows some sample compression ratios and sample effects of compression on workloads that are CPU intensive and I/O intensive.

On the zSeries platform, with the dedicated coprocessor unit, the z900 (G7) machine has significantly higher performance than the former G6 processor, reducing in the order of one third the number of cycles for compression. Increasing performance requirements for data compression was the main reason for choosing a coprocessor-based implementation.

The advantages of hardware compression are:

- Faster and less expensive than software-based algorithms
- Runs as a black box
- Executes in dedicated microprocessors, relieving the main processor for other tasks

It is a simple task for the DBA to determine if compression is suitable. IBM provides a DB2 stand-alone utility, DSN1COMP, which verifies the compression benefits. This can be done prior to activation and from a copy of the data and does not impact your production tables. z/OS compression allows customization

of the compression routine to maximize the compression of the specific data. This customization is automatically done by DB2, which samples the data to be compressed.

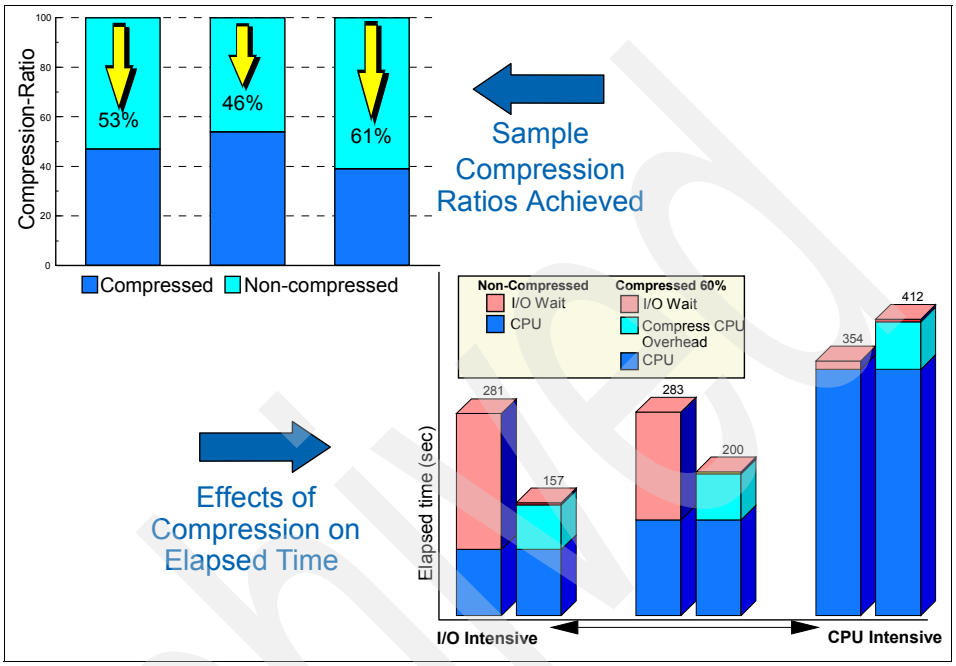

Figure 3-2 Effects of compression

Some data compresses very well, and some data is not well suited for compression. Generally data with a lot of repeating characters compresses well. An ideal candidate is a large table where rows are larger than 100 bytes, with a high probability of repetitive strings of character, accessed by either individual random SELECT or massive scans and periodically updated or added. Small rows do not compress well because of the maximum of 255 rows per page and the additional dictionary information stored with the table space.

If compression is suitable, the appropriate objects can be easily altered and reorganized to be compressed. Compression helps you store and manage more of your business data. With an ideal candidate table, accessed via random SELECTs, the CPU overhead for decoding is low especially compared to the I/O time. And for massive scans, the savings in disk space often outweigh the CPU overhead.

The zSeries server provides faster CPUs and even faster decompression instructions. It allows a better performing compression. Compression can now be considered to alleviate disk occupancy in a more widespread range of situations.

3.4 Security

These days, everyone is more aware of security. In today's business world, data is growing exponentially and needs to be accessible all of the time from anywhere. Along with this, there has been a general heightened awareness and interest in security and privacy. There are frequent news reports about security breaches and potential threats. Antivirus software is a necessity on client computers. Firewalls must be in place. New security legislation issues must be addressed.

More data is available to more people, through more channels. Data needs to be protected, regardless of whether it is internal company financial data or a customer's social security number or credit card information. Data should only be used for the purpose for which it was originally collected. In some cases, data from one part of the business needs to be isolated from other parts of the business.

3.4.1 zSeries security features

z/OS extends its robust mainframe security features to address the demands of on demand enterprises. Such technologies as Secure Sockets Layer (SSL), Kerberos V5, Public Key infrastructure, and multilevel security, as well as the exploitation of zSeries Cryptographic features are available in z/OS.

Integrated Cryptographic Service Facility (ICSF) is a part of z/OS. It provides cryptographic functions for data security, data integrity, personal identification, digital signatures, and the management of cryptographic keys. Together with cryptography features of the zSeries servers, z/OS provides high-performance SSL, which can benefit applications that use the z/OS HTTP Server and WebSphere, TN3270, and CICS Transaction Gateway server.

z/OS provides support for digital certificates, including the ability to provide full life-cycle management. With Public Key Services in z/OS, clients can create and manage digital certificates, leveraging their existing z/OS mainframe investments. This may provide significant cost savings over other digital certificate hosting options.

zSeries has attained Common Criteria Evaluation Assurance EAL5 Certification. This evaluation demonstrates that the zSeries can be an essential building block for server consolidation and for the integration of on demand applications and

traditional corporate workloads on a single server. You can run z/OS, z/VM, and Linux for zSeries in different LPARs on a single zSeries server. You can consolidate sensitive applications onto one system with highly secure isolation of workloads and data.

3.4.2 DB2 controls

DB2 has strong and granular access control. It controls access to its objects by a set of privileges. Default access is none. Until access is granted, nothing can be accessed. This is called *discretionary access control* (DAC).

DB2 has extensive auditing features. For example, you can answer such questions as, “Who is privileged to access what objects?” and “Who has actually accessed the data?”

The catalog tables describe the DB2 objects, such as tables, views, table spaces, packages, and plans. Other catalog tables hold records of every granted privilege or authority. Every catalog record of a grant contains information such as name of the object, type of privilege, IDs that receive the privilege, ID that grants the privilege, and time of the grant.

The audit trace records changes in authorization IDs, changes to the structure of data, changes to values (updates, deletes and inserts), access attempts by unauthorized IDs, results of GRANT and REVOKE statements, and other activities which are of interest to auditors.

You can use the Resource Access Control Facility (RACF) to:

- Control access to the DB2 environment
- Facilitate granting and revoking to groups of users
- Ease the implementation of multilevel security in DB2
- Fully control all access to data objects in DB2

DB2 defines sets of related privileges, called *administrative authorities*. You can effectively grant many privileges by granting one administrative authority.

Security related events and auditing records from RACF and DB2 can be loaded into DB2 databases for analysis. The DB2 Instrumentation Facility Component can also provide accounting and performance related data. This kind of data can be readily loaded into a standard set of DB2 tables (definitions provided). Security and auditing specialists can easily query this data to review all security events.

3.4.3 DB2 and multilevel security

A multilevel security (MLS) system is a security environment. It allows the protection of data based on both traditional discretionary access controls, and controls that check the sensitivity of the data itself through mandatory access controls.

These mandatory access controls are at the heart of an MLS environment. They prevent unauthorized users from accessing information at a classification to which they are not authorized. They also prevent users from changing the classification of information to which they do have access. These mandatory access controls provide a way to segregate users and their data from other users and their data regardless of the discretionary access they are given though access lists.

To create an MLS environment, you must have a combination of software and hardware components that enforce the security requirements needed for such a system. The security relevant portion of software and hardware components that make up this system is also known as the *Trusted Computing Base*.

Why multilevel security

The primary arena where MLS is valuable is government agencies that need a security environment that keeps information classified and compartmentalized between users. In addition to the fundamental identification and authentication of users, auditing and accountability of the actions by authenticated users on these systems is provided by the security environment.

In such highly secure environments, to manage the compartmentalization of information between users, each compartment is on its own system. This makes it difficult for classified information to spill from one system to another, since the connections between systems can be highly controlled. With MLS, these systems can be consolidated onto a single system, with each compartment independent of the other, so that no transfer of data can occur between compartments within that system. This takes advantage of the cost savings of not having to manage multiple systems, but only a few, or one system.

Commercial clients may also find some features of MLS useful, such as to separate sensitive customer information from the general populace or from other users. New government regulations, such as HIPAA (see the following list), or corporate mergers are examples where security of information based on the information itself is important in the commercial world.

MLS is implemented at the operating system level. DB2 Version 8 participates in this scheme and provides MLS security to the row level.

This additional mandatory access control feature helps your business to address the most common security issues. It does this along with the rock solid z/OS partition separation, certification for common criteria, improved DB2 encryption features, and the new DB2 and IMS encryption tool. Together these features also help your business to comply with existing and new regulations such as:

- Health Insurance Portability and Accountability Act of 1996 (HIPAA) in the U.S.A. for health care
- Gramm-Leach-Bliley Act of 1999 (GLBA) in the U.S.A. for financial services
- Sarbanes-Oxley, an act which aims to protect investors by improving the accuracy and reliability of corporate disclosures
- Personal Information and Electronic Documents (PIPEDA) Act in Canada
- United Kingdom Data Protection Act (Oct 1998)
- European Union Data Protection Directive 95/46/EC

Multilevel security concepts

Mandatory access control (MAC) imposes additional restrictions upon users. Now users access data based on a comparison of the classification of the user and the classification of the data as well as the standard *discretionary access control* (DAC) checking. This additional security check verifies that users can access only data and resources that their classification allows them to. This happens regardless of whether they have discretionary access to such data or resources. Mandatory means that subjects cannot control or bypass the access.

An MLS system is a security environment that allows the protection of data. It is based on both traditional discretionary access controls and controls that check the sensitivity of data itself through mandatory access controls. Using MLS allows you to classify objects and users with security labels that are based on hierarchical security levels and non-hierarchical security categories.

You can implement MAC via RACF at the operating system level. The key advantage is that MAC security can be integrated across the platform, with the same security for files, print, and DB2.

RACF, also known as the z/OS Security Server, has several options that can be turned on and off to manipulate different aspects of an MLS environment. It is possible to have some features of multilevel security on at one time (creating a partial MLS environment). Because of this, commercial customers may find this type of environment useful to meet these needs as opposed to running a complete multilevel security environment.

Using MLS, you can define security for DB2 objects. by assigning security labels to your DB2 objects. You can define a hierarchy between those objects. MLS then restricts access to the object based on the security level of that object.

Row level security as a subset of multilevel security

In today's complex world, organizations may have considerable needs to restrict access to data in their database applications. Privacy and data protection legislation, antitrust legislation, and considerations of national security are a few of the reasons why organizations need to ensure that people in one part of the business do not know and cannot determine what is happening in another part.

Companies may lose business if they cannot demonstrate to security-conscious potential customers that data relating to them is strictly protected from unauthorized access. Regulatory authorities may require that computer systems be separately maintained if it cannot be shown that the different divisions of a company are prevented from accessing details of the other divisions' operations.

In an increasingly interconnected era, organizations may want to offer limited access to their operational systems to the clients, suppliers and trading partners. They may note want to give those people freedom to roam through data pertaining to their competitors.

The requirement can be summarized as follows, using a customer order database as an example. It must include:

- A means of marking a customer as being in a set of protected customers
- A means of propagating such markings to related data after the customer marking is made
- A means of marking new data for such customers and transient data relating to such customers with the security markings of the customer
- The restriction of access to data based on such markings
- A means of ensuring that the security markings of the data are not changed other than by authorized users and processes

You can give a group of people access to a table and limit each user to only a subset of the data based on the particular individual's MLS definition. And you can do this without creating views or placing extra predicates in the SQL. Instead, you place a security label on the data row and then associate the security label with each of the users. The Data Manager layer knows how to compare these security labels and can see whether a particular user is allowed to access a row. It is flexible in terms of having one set of tables and many different ways of accessing the data or providing a subset of the data.

The concept of dominance is relevant in MAC. Consider a scheme where rows in a table are identified as either “Top secret”, “Sensitive” and “Unclassified”. A subject or user with a a security label of “Sensitive” sees all sensitive and unclassified rows, but is not aware of the presence or absence of top secret rows.

Figure 3-3 illustrates another security scheme using colors. With this hierarchy established in the security server, the system understands that users with authority to access RAINBOW can access anything.

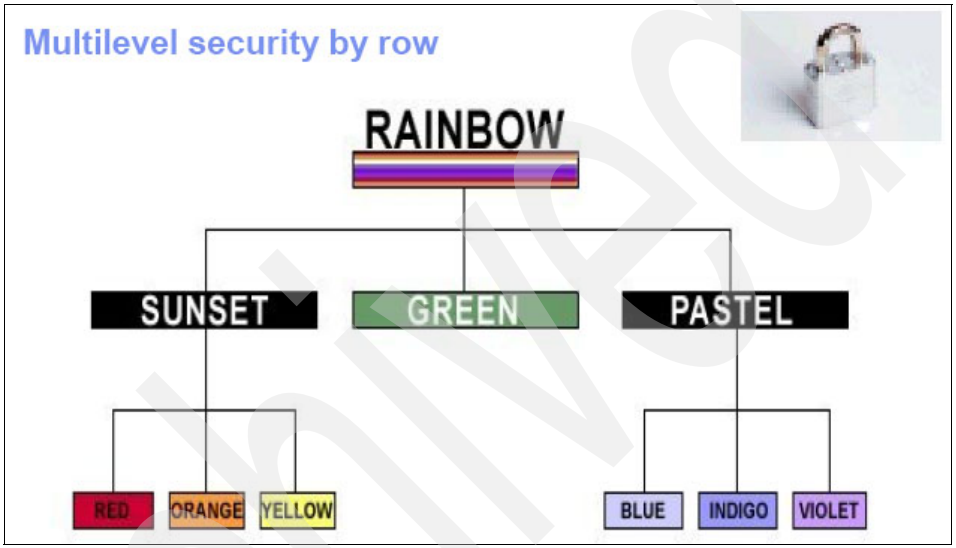

Figure 3-3 Rainbow MLS hierarchy

Someone with authority to access PASTEL information can access any row associated with BLUE, INDIGO, VIOLET, or PASTEL. Someone with SUNSET can access SUNSET, RED, ORANGE, YELLOW. This is a lot more powerful than having an exact match on a security label (such as having a user’s label match exactly the data’s label). It has the notion of “groups” that make security administration easier to manage. With this additional capability, you can implement this type of security scheme without requiring the application to access the data using special views or predicates.

DB2 provides row level MLS, which enables a high degree of flexibility and implementation of sophisticated security rules. DB2 leverages security functions in the operating systems and RACF to deliver leading edge security at low levels of granularity.

z/OS 1.5 together with DB2 UDB Version 8, provide a solution for multilevel security on the zSeries mainframe. This support provides row-level security labeling in DB2 and protection in z/OS designed to meet the stringent security

requirements of multi-agency access to data. This solution leverages zSeries leadership in scale, high availability, and self managing capabilities for highly secure single-database hosting.

3.4.4 Summary

Data is the key to running your business. DB2 and zSeries hardware and software give you the controls for safely and effectively administering it.

- DB2 leverages security functions in the operating system. With multilevel security, users can implement sophisticated security in their DB2 applications without writing their own code and be better positioned to obtain auditor certification.
- DB2 leverages cryptographic functions in the hardware.
- Both security and cryptographic functions enable delivery of leading-edge security at low levels of granularity, for example, individual rows and columns instead of tables.
- Common criteria Evaluation Assurance EAL5 Certification for zSeries and z/Os demonstrates that the zSeries can be an essential building block for server consolidation and the integration of on demand applications and traditional corporate workloads on a single server. This is desirable for reasons of economy, flexibility, security, or management.

The combination of DB2 Version 8 functionality, the z/OS security features, and the z/Architecture hardware encryption capability establish a highly secure environment for your DB2 data as it grows and as it is accessed from anywhere at anytime as part of On Demand Business.

Your data is secure to the row level and column level. This is enforced with the DB2 and zSeries state of the art hardware and software.

3.5 Encryption

This section introduces both DB2 encryption and the IBM Data Encryption Tool. It discusses some recent hardware enhancements that improve encryption performance.

Today, many organizations are paying much more attention to the security of their data. They may need to comply with various security regulations, such as HIPAA, Sarbanes-Oxley, or Gramm-Leach-Bliley Acts. Or they are prompted as a result of emerging new technologies, such as the emergence of the Internet, storage area networks (SANs), and more intelligent storage controllers. That is, do you trust your storage vendor not to look at your data?

Some governments require IT products to be certified against the *Common Criteria* for evaluation of IT security. Common Criteria is a contribution to the development of an international standard. It opens the way to worldwide mutual recognition of IT security evaluation results.

DB2 ships a number of built-in functions which allow you to encrypt and decrypt data. IBM offers an encryption tool called the *IBM Data Encryption for IMS and DB2 Databases*, program number 5799-GWD. This section introduces both DB2 encryption and the IBM Data Encryption tool. It also discusses recent hardware enhancements that improve encryption performance.

Data encryption has several challenges. These include making changes to your application to encrypt and decrypt the data, encryption key management, and the performance overhead of encryption.

DB2 encryption is available at the column level and at the row level.

Machines prior to the z990 have a Cryptographic Coprocessor Feature to improve the performance of encryption and decryption. However, only two of the CPUs could perform encryption tasks that require encryption. Running on other processors had to be dispatched to those CPUs. Performance could suffer with contention among tasks, such as parallel queries. Also dedicated LPARs may be unable to use the encryption hardware feature.

The z990 introduced a new hardware instruction, CP Assist for Cryptographic Function (CPACF), which can run on all CPUs. The CPACF feature is available only on the z990 hardware and later, not the older z900. The z990 also introduces a PCIXCC card, which is needed for the IBM Data Encryption tool, but not for the DB2 encryption function. The following sections briefly introduce these two encryption functions.

3.5.1 DB2 column level encryption

DB2 Version 8 ships a number of built-in functions which allow you to encrypt data at the *cell* level. These functions are ENCRYPT_TDES (or ENCRYPT), DECRYPT_BIN, DECRYPT_CHAR, and GETHINT.

The SET ENCRYPTION PASSWORD statement allows you to specify a password as a key to encryption. Because you can specify a different password for every row that you insert, you can encrypt data at the cell level in your tables. However, you are responsible for managing all these keys. Make sure you have a mechanism in place to manage the passwords that are used to encrypt the data. Without the password, there is no way to decrypt the data. These encryption functions use the Triple Data Encryption Standard (DES) to perform the encryption.

The DB2 built-in encryption functions require:

- DB2 Version 8
- Integrated Cryptographic Service Facility (ICSF)
- On z990, CPACF

> **Note:** The Peripheral Component Interconnect Extended Cryptographic Coprocessor (PCIXCC) card is not required, unless DRDA encryption is necessary.

- Pre-z990 cryptographic coprocessor

Each CP on the z990 has an assist processor on the chip in support of cryptography. This feature provides for hardware encryption and decryption support. PCIXCC provides a cryptographic environment with added function. To learn more about PCIXCC, refer to *IBM @server zSeries 990 (z990) Cryptography Implementation*, SG24-7070.

Applications that need to implement DB2 encryption must apply the DB2 encrypt and decrypt built-in functions to each column to be encrypted or decrypted. All encrypted columns must be declared “for bit data”. Unchanged read-applications see data in encrypted form. Applications may apply a different key for each column, but may also supply the key in a special register. We strongly recommend, for performance, that you specify the key in the special register.

The LOAD and UNLOAD utilities do not support DB2 encryption, but SQL-based programs such as DSNTIAUL do support encryption. Encryption of numeric fields is not supported. The length of encrypted columns must allow for an additional 24 bytes, rounded up to a double-word boundary. Space usage may be a concern if you plan to use DB2 to encrypt small columns.

Indexes are also encrypted. Predicates that depend on the collating sequence of encrypted columns (for example, range predicates) may produce wrong results (unless modified to use built-in functions correctly). For example, the following statement produces the wrong results:

```
SELECT COUNT(*) WHERE COL =:HV;
```

The following statement produces the correct results with almost no impact to performance.

```
SELECT COUNT(*) WHERE COL = ENCRYPT_TDES(:HV);
```

The following statement produces the produces wrong results.

```
SELECT COUNT(*) WHERE COL < ENCRYPT_TDES(:HV);
```

The next statement produces the correct results with a large impact on performance.

```
SELECT COUNT(*) WHERE DECRYPT_CHAR(COL) <:HV;
```

3.5.2 IBM Data Encryption for IMS and DB2 Databases

Are you the person responsible for protecting your company's sensitive IMS and DB2 data? Are you investigating how to comply with security legislation in such industries as health care and finance?

If so, IBM Data Encryption for IMS and DB2 Databases is the tool you need. It provides you with a data encryption tool for both IMS and DB2 for z/OS databases in a single product. It enables you to protect your sensitive and private data for IMS at the segment level and for DB2 at the table level

This tool performs row level encryption using EDITPROCs. Unlike the DB2 encryption functions shipped with DB2, the Data Encryption Tool uses different keys to encrypt different tables. The encryption keys can be either *clear*, such as the DB2 encryption functions, or *secure*. Plus they are managed through ICSF. Clear keys generally perform better. The tool also supports single, double, or triple DES. Again, refer to *IBM @server zSeries 990 (z990) Cryptography Implementation*, SG24-7070, to learn more about the clear and secure keys.

You can find more information about the IBM Data Encryption for IMS and DB2 Databases by visiting the Web at:

http://www.ibm.com/software/data/db2imstools/db2tools/ibmencrypt.html

The IBM Data Encryption for IMS and DB2 Databases tool supports all versions of DB2, and it encrypts only the whole row. No application changes are required. However you must modify the DDL to include the EDITPROC. The applications need not be aware of encryption keys.

3.5.3 Summary

Encryption enables you to leverage the power of SANs safely while complying with privacy and security regulations. The IBM Data Encryption for IMS and DB2 Databases tool is implemented via standard IMS and DB2 exits. The exit code invokes the zSeries crypto hardware to encrypt data for storage and decrypt data for application use. The tool can help you save the time and effort required to write and maintain your own encryption software for use with such exits or within your applications.

With the increased demand for data privacy and security, the need for data encryption has moved to the forefront of technology concerns. In today's on

demand environment, applications require authentication and confidentiality. IBM addresses this need.

With DB2 Version 8, z990 hardware improvements, and the new IBM Data Encryption for IMS and DB2 Databases tool, you have flexible options for complying with legislation and for protecting your sensitive data, whether internal company data or customer and partner data. zSeries crypto functions support industry standards. With the Crypto Coprocessor and Integrated Cryptographic Services Facility, you can implement high-performance cryptography. The functionality provided by these options can be easily understood by auditors.

3.6 Sort

DB2 uses two types of sort: DB2 sort is used for query-related sorts, such as GROUP BY and ORDER BY, while DFSORT™ is used for utility-related sorts.

DB2 sort has been using microcode-assisted sort instructions, available on all IBM processors, since 1992. The DB2 sort performance has taken advantage of the hardware improvement since then. The DB2 sort using the microcode-assisted sort instructions was referred to as ESA sort. With DB2 Version 4, the use of the ESA sort was extended to all DB2 sort operations. The enhancement resulted in reduced CPU time and elapsed time.

DFSORT *always* ships with z/OS, even though you may not be licensed for it. DB2 Version 8 provides a “special” license for DB2 to use DFSORT, without the user needing to acquire an actual DFSORT license.

With Version 8, the DB2 utilities benefit from using DFSORT. DFSORT prefers fewer, large sortwork data sets. Other products prefer more, smaller sortwork data sets. Other sort products are not always installed as re-entrant and the storage required limits the degree of parallelism. Sorting performance is improved with memory object sorting in DFSORT using 64-bit storage.

This enhancement allows DB2 to exploit particular functions and features of DFSORT. This means better sort performance and more robust DB2 utilities for your DB2 data.

3.7 z/OS Unicode Conversion Services

In this area of globalization, the ability for systems to handle data from around the world is paramount. However, workstations and servers may use different code pages, depending on the country in which they reside. In effect, they are speaking different languages, making communications difficult.

Unicode provides a consistent way to encode multilingual plain text and brings order to a chaotic state of affairs. DB2 understands Unicode and users do not have to convert existing data. DB2 can integrate newer Unicode data with existing data and handle the translations. The synergy between DB2 and z/OS Unicode Conversion Services helps this process to be high performing.

z/Architecture instructions exist that are designed just for Unicode conversions. There are significant Unicode functional and performance enhancements in z900, z/OS 1.4, z990, and DB2 Version 8.

3.8 UNIX System Services

The two major reasons for UNIX on the mainframe are the portability of skills and the portability of code. You can use the same skills on one platform and then run them on another. You can also move functions from one platform to another. z/OS UNIX System Services is tightly integrated into the operating system. You gain the best of both worlds: UNIX and z/OS. z/OS UNIX System Services are an important element of the IBM open and distributed computing strategy.

Here are some facts about UNIX System Services:

- z/OS UNIX is a certified UNIX system.
- z/OS UNIX is an integral element of z/OS.
- WebSphere Application Server, CICS, IMS, Java Runtime, Tuxedo, DB2, WebSphere MQ, SAP R/3, Lotus Domino, and Oracle Web Server all use z/OS UNIX.
- z/OS UNIX applications can communicate with DB2, CICS, IMS, and WebSphere MQ.
- z/OS UNIX is built for the enterprise where you can prioritize workloads for high performance when running with a mixed workload.
- A broad range of independent software vendor (ISV) applications is ported to z/OS UNIX.

- z/OS UNIX has a hierarchical file system familiar to UNIX users.
- Applications can work with data in both UNIX hierarchical file systems and traditional MVS data sets.

Here are some advantages of z/OS UNIX:

- Use of all enterprise services by z/OS applications
  - Through Workload Manager, the ability to manage the allocation of physical resources to maximize system responsiveness
  - High performance of database access
  - Access to 2-phase commit through RRS
  - Access to advanced security features of RACF
- Enterprise-class applications and middleware
- Extensive support in the marketplace

With UNIX System Services, you can manage and secure system resources from a single point, consistently run near 100% utilization, and benefit from Workload Manager allocation of resources. You can run your UNIX programs on the mainframe and be closer to the DB2 database. This avoids network hops and enables you to benefit from the high availability and scalability of the zSeries environment.

3.9 WorkLoad Manager, Intelligent Resource Director, and VIPA

Workloads have grown from business expansion or consolidation, and from developing and purchasing new applications. The competition for the mainframe resources is increasing and is continuous. A mixed application workload must be carefully managed to avoid conflicts and significant adverse impact on critical transactions. The complex constant management and adjustment required to optimize the workload throughput depend on efficient automated workload and resource software.

3.9.1 WorkLoad Manager

To this end, IBM introduced and evolved the zSeries WLM and IRD. Working in harmony, their unique capabilities ensure maximum usage of your infrastructure investment.

Dynamic resource allocation helps to allow for maximum transaction processing throughput, so your customers, suppliers, and employees can count on reliable service. Predefined business priorities allow the system to continually and dynamically direct resources to top-priority work, so your most important applications can have the power and bandwidth they need at the right time.

A typical z/OS configuration contains a variety of processor types and sizes, generally operating in LPAR mode. They support images that run batch, online transaction processing (OLTP), Web servers, application development, enterprise resource planning (ERP) applications, business intelligence (BI), and various other workloads. For larger enterprise installations, you can expect to add the complexity of multiple physical servers running in a Parallel Sysplex configuration. It is not practical to allow these applications conflict with each other or to maintain an expensive team of operations staff to monitor and adjust resource priorities.

Automated workload management is a necessity. It enables the control of executable units of work within and across systems to ensure the best utilization of system resources. It does this to meet the demands of DB2, CICS, batch, TSO, UNIX System Services, and Web servers.

When in goal mode system operation, WLM provides fewer, simpler, and consistent system externals that reflect goals for work expressed in terms commonly used in business objectives. WLM and Service Request Manager (SRM) match resources to meet those goals by constantly monitoring and adapting the system.

Figure 3-4 shows the difference between point-in-time reallocation of resources in the IBM @server pSeries® environment and the zSeries continuous reallocation of resources based on the user service level plan. This is a unique *autonomic* zSeries systems management feature, which is unmatched in the industry.

Workload Manager in goal mode has knowledge of the system utilization and workload goal achievement across all the systems in the sysplex. This cross-system knowledge and control has a much more significant impact in an IRD environment.

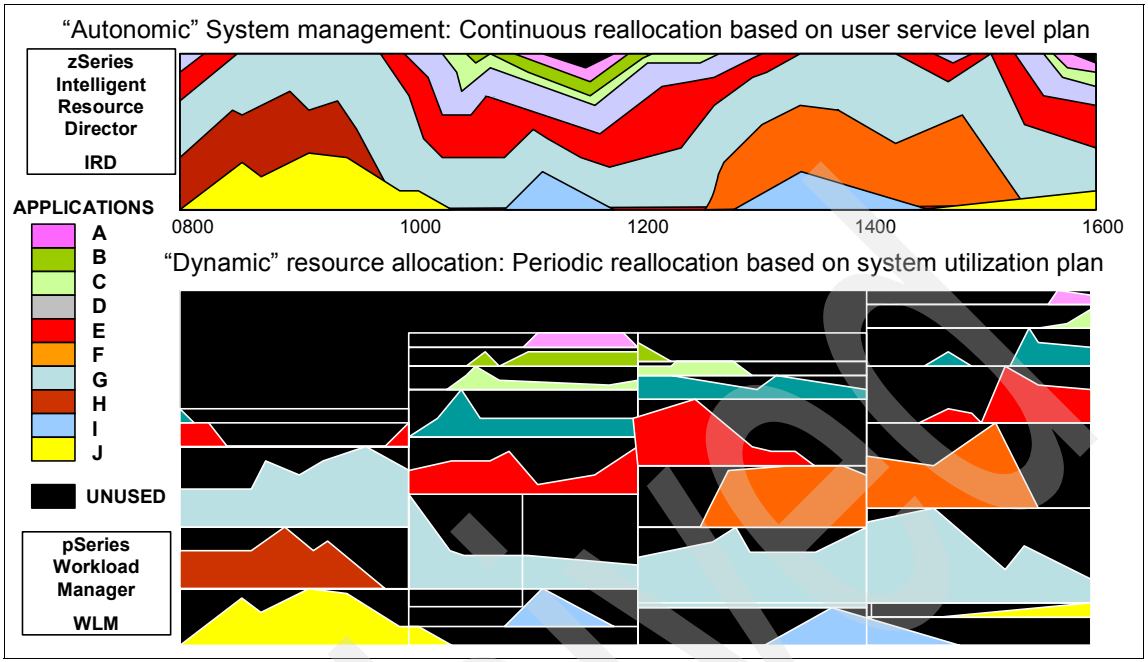

Figure 3-4 Autonomic system management

3.9.2 Intelligent Resource Director

Exclusive to the IBM z/Architecture is IRD. This function optimizes processor and channel resource utilization across LPARs based on workload priorities. IRD uses facilities in z/OS Workload Manager, Parallel Sysplex, and PR/SM to help you derive greater value from your zSeries investment.

Compared to other platforms, z/OS with WLM already provides benefits from the ability to drive a processor at 100% while providing acceptable response times for your critical applications. IRD amplifies this advantage by helping you ensure that all those resources are being used by the right workloads, even if the workloads exist in different LPARs. IRD gives you the ability to move the resource to where the workload is.

IRD is not a product or a system component. Rather it is three separate but mutually supportive functions:

- WLM LPAR CPU management

  Processor resources are distributed across an LPAR cluster by dynamically adjusting the LPAR weights in response to changes in workload requirements.

- Dynamic channel-path management

  Dynamic channel path management enables Workload Manager to dynamically move channel paths through the ESCON® director from one I/O control unit to another, in response to changes in workload management.

- Channel subsystem I/O priority queuing

  In an LPAR cluster, if important work is missing its goals due to I/O contention on channels shared with other work, it is given a higher channel subsystem I/O priority than the less important work.

Figure 3-5 shows the benefits of IRD with an example. To simplify, OLTP has importance 1, with BI 2 and batch 3. During the day, OLTP work has a higher priority over BI work. During this time, 75% of the processing power is devoted to LPAR 1, while BI work in partition 2 is processed with the remaining 25%. At night, this is automatically adjusted. The BI work is given 75% of the processing power, and the Batch work in partition 1 is done with the other 25%. Notice also the re-adjustment of the channel subsystems. During the prime shift, the OLTP/BATCH disks have more channels, where in the evening, there are more paths to the BI disks.

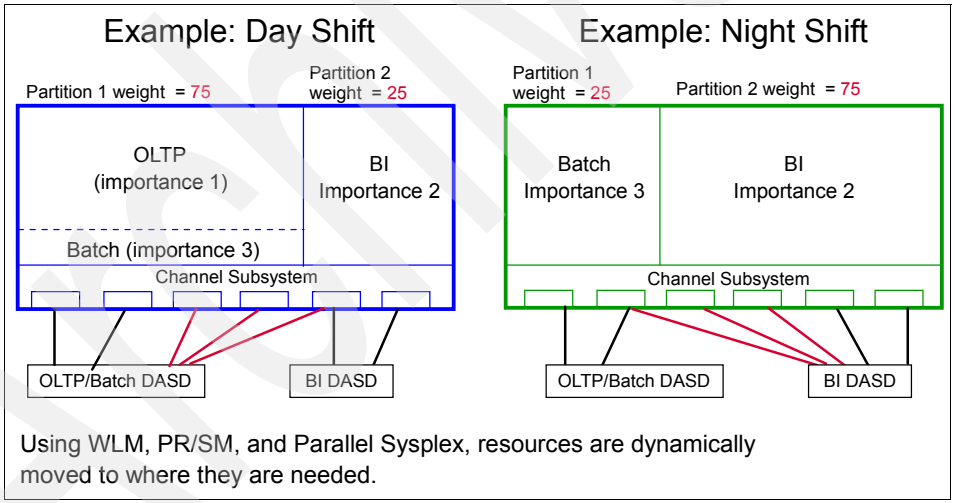

Figure 3-5 IRD benefits

3.9.3 Virtual IP addressing

Dynamic virtual IP addressing (VIPA) gives you the ability to assign a specific, virtual IP address to a data sharing group and to each member of the group. This address is independent of any specific TCP/IP stack within the Parallel Sysplex. Even if a member is moved to another z/OS system, as in the case of a failure or

maintenance, the member remains accessible. It retains the same virtual IP address, without any connection implications.

The Sysplex Distributor enables network-level load balancing for initial requests (to member with most capacity), followed by database-level load balancing for subsequent connections. The Sysplex Distributor monitors real-time information from Workload Manager, such as member status and QoS data.

With VIPA, the Sysplex Distributor handles fail over and ensures that connections are continuously routed to the best member to handle the work.

3.9.4 Summary

zSeries can dynamically reallocate resources to maintain “end-user service levels”, prioritized by business goals, using external performance measurements. Workload Manager provides a solution for managing workload distribution, workload balancing, and distributing resources to competing workloads. This results in higher utilization, since the system can schedule work on an ad hoc basis to sustain utilization levels approaching 100%.

zSeries capability for mixed workloads is built into the hardware, hypervisors, operating systems, and clustering. This is done efficiently and demonstrated over 25 years of production.

The zSeries server is more mature, more comprehensive, and more dynamic. There is no goal mode, or IRD for UNIX systems. UNIX WLMs tend to act more as governors. More utilities are needed in the UNIX environment.

With WLM and IRD, you can dynamically assign processors, I/O, memory resources, and TCP/IP connections within and across multiple z/OS images. This can be done by transaction type, user, and time of day. This provides end-to-end business management based on business priorities and SLAs.

Enabled by WLM and IRD, these powerful management capabilities across a complex enterprise infrastructure help a business protect its critical workloads by extending the prioritized dynamic distribution of resources across multiple servers. For example, you can be running at 100% utilization and add a high priority workload. This software accepts the new work and enables the desired throughput by slowing down some of the low priority work, effectively providing capacity on demand. This is unequalled in the industry.

z/OS unique features make SLAs possible under the most demanding and unpredictable environments.

3.10 Disk storage

DB2 continues to deliver synergy with FICON® (fiber connector) channels and disk storage features. It integrates with the storage management software.

DB2 uses Parallel Access Volume and Multiple Allegiance features of the IBM TotalStorage Enterprise Storage Server (ESS). ESS FlashCopy® is used for DB2 BACKUP and RESTORE SYSTEM utilities. FlashCopy V2 helps increase the availability of your data while running the CHECK INDEX SHRLEVEL(CHANGE) utility.

Larger control interval sizes help performance with table space scans. Striping is an effective way of increasing sequential throughput by spreading control intervals (CIs) across multiple devices.

3.10.1 ESS

The IBM TotalStorage ESS is the most powerful disk storage server from IBM, developed using IBM Seascape® architecture. The ESS provides unmatchable functions for all On Demand Business servers and for non-IBM (that is, Intel®-based and UNIX-based) families of servers. Across all of these environments, the ESS features unique capabilities that allow it to meet the most demanding requirements of performance, capacity, and data availability.

To meet the unique requirements of On Demand Business, where massive swings in the demands placed on your systems are common and continuous operation is imperative, you need high-performance and intelligent storage technologies. You also need systems that can support any server application in your business, today and into the future.

Since its initial availability with the ESS Models E10 and E20, and then with the succeeding F10 and F20 models, the ESS has been the storage server solution. It offers exceptional performance, extraordinary capacity, scalability, heterogeneous server connectivity, and an extensive suite of advanced functions to support users' important, high-availability, and multiplatform environments. The third-generation ESS Model 800 keeps up with the pace of users' needs by adding more sophisticated functions to the initial set. It enhances connectivity options and powers its performance features.

The IBM TotalStorage ESS Model 800 provides significantly improved levels of performance, throughput, and scalability. It continues to exploit the innovative features introduced with its preceding E and F models, such as Parallel Access Volumes, Multiple Allegiance, I/O Priority Queuing, the remote copy functions (synchronous, non-synchronous, and asynchronous), and the FlashCopy

point-in-time copy function. Also, the heterogeneous server support characteristics of previous models are continued with the ESS Model 800.

Figure 3-6 summarizes the ESS capabilities. The following sections provide a brief description of the most important ESS functions.

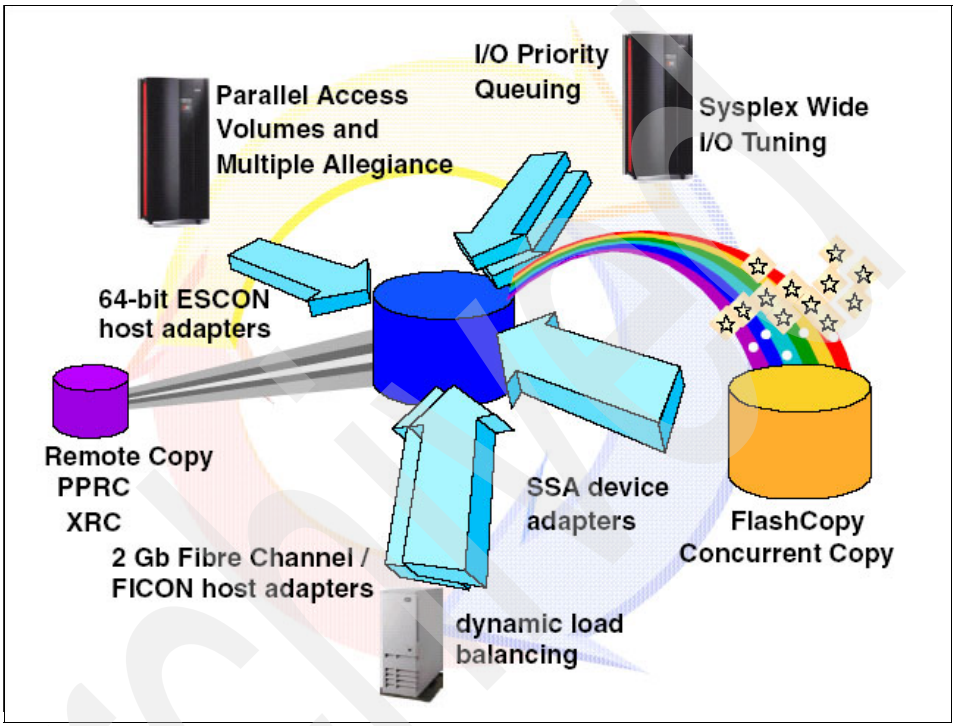

Figure 3-6 IBM TotalStorage Enterprise Storage Server capabilities

Multiple Allegiance, Parallel Access Volumes, and I/O Request Priority

In the early days, disk hardware was only capable of processing one I/O at a time, which limited OS/390 systems. They did not try to issue another I/O to a disk volume, represented in MVS by a Unit Control Block (UCB), while an I/O was already active for that device. In addition, the storage subsystems accepted only one I/O at a time from different system images to a shared disk volume, for the same reasons mentioned earlier.

The ESS has the capability to do more than one I/O to an emulated S/390 volume. The ESS introduces the concept of *alias addresses*. Instead of one UCB per logical volume, an OS/390 host can now use up to 256 UCBs for the same logical volume. Apart from the conventional base UCB, alias UCBs can be

defined and used by OS/390 to issue I/Os in parallel to the same logical volume. The function that allows parallel I/Os to a volume from one host is called *Parallel Access Volumes* (PAV). But I/Os are not limited to coming from one host in parallel. The ESS also accepts I/Os to a shared volume coming from different hosts in parallel. This capability is called *Multiple Allegiance*.

The ESS shows virtually no degradation when a DB2 parallel query scans one, two, or three partitions simultaneously on the same volume. The PAV function nearly eliminates the IOSQ time and removes the contention caused when multiple concurrent I/Os try to access the same volume. The same measurement on older devices would show substantial degradation in elapsed time. It would be two times worse for two partitions and over three times worse for three partitions on one volume.

Multiple Allegiance

Multiple Allegiance is the capability to support I/O requests from multiple systems, one per system, to be concurrently active against the same logical volume if they do not conflict with each other. Although Multiple Allegiance does not require software support, it is required for Parallel Access Volume.

Parallel Access Volume

Before PAV was available, the operating system allowed only one request at a time for each volume. Thus, when there was an active I/O request to a volume, its UCB was flagged as busy.

To support PAV, multiple unit addresses are associated with the same logical volume. Each such address is associated with a corresponding subchannel within the zSeries channel subsystem (CSS). These additional addresses of the volume are known as *PAV aliases* or simply *aliases*.

These aliases are dynamically managed by WLM for all releases of z/OS. With the WLM dynamic alias management function, the alias addresses are managed as a pool. The aliases in the pool are available for allocation to volumes, rather than being statically associated with specific volumes. The number of aliases associated with any particular volume is dynamically adjusted by WLM to achieve the workload goals or to improve overall efficiency when all workloads are achieving their goals. This eliminates work for the DBAs in relocating data sets to manage contention.

For database applications, the storage management tasks associated with obtaining high levels of parallelism are also simplified. Administrators no longer need to worry about this. They can let the z/OS Data Facility Storage Management Subsystem (DFSMS) facilities place DB2 partitions. Also WLM management of PAV can support high parallel query throughput, regardless of the data placement.

DFSMS is a set of components that are integrated within z/OS. It employs the concept of policy-based storage management. This involves defining policies which allow the system to take over many storage management tasks that were previously performed manually. All of these autonomic features are designed to improve productivity in your IT shops.

I/O Request Priority

I/O Request Priority (IORP) is a means for associating a priority with each I/O request. Multiple applications within an operating system image may issue I/O requests concurrently.

The I/O supervisor (IOS) component of z/OS queues the work for each target logical volume in order of priority. The number of requests started match the number of available PAV aliases. These requests are initiated by sending them to the CSS. The CSS orders them by priority on its internal work queue. When an available channel path is found, the I/O request is sent to it, where again it is started in priority order.

The I/O requests are transported through the I/O fabric, such as Fibre Channel, to the control unit where the higher priority work is allocated greater portions of the bandwidth available on the FICON channels. WLM also handles the dynamic management of I/O priorities. This complements the dynamic alias management that it does.

Although PAV increases parallelism for I/O requests, and reduces a major cause for I/O delays, contention for resources may exist at some level, either in the operating system, in CSS, or in ESS. WLM I/O priority management ensures that where there is contention, the work is processed so that the work with the highest business importance is executed first and has handled the processor and storage needs. z/OS with ESS provides end-to-end management of I/O priorities, coordinated across even the multiple systems of a sysplex.

FlashCopy

FlashCopy provides an instant, point-in-time copy of the data for application usage such as backup and recovery operations. It enables you to copy or dump data while applications are updating the data. Source and target volumes can reside on the same or different logical subsystem (LSS), also known as a *Logical Control Unit* (LCU). DFSMSdss™ automatically invokes FlashCopy when you issue the COPY FULL command on a subsystem that supports FlashCopy functions. This is a unique feature of the ESS.

DB2 for z/OS Version 8 takes advantage of the FlashCopy functions when running utilities.

3.10.2 FICON channels

FICON channels enable faster access to your DB2 data. They have many benefits over the traditional ESCON and parallel channels, including:

- Increased data transfer rates

  The FICON data link rate during data transfer is 100 MB/sec. With 2 Gb links, FICON Express can deliver greater sustained data rates in the range of 150 MB/sec to 170 MB/sec for large data transfers.

- Improved performance

  Native FICON attachment can offer much higher data transfer rates than ESCON attachments. For critical data-intensive applications, this can translate into significant improvements in overall elapsed time. For example, QSAM workloads (typical of batch) and large non-parallel DB2 queries (typical of data warehousing and data extracts) can run over two times faster with FICON than with ESCON.

- Reduced backup windows

  Working together, the FICON-attached ESS and 3590-A60 Magstar® tape system can slash elapsed times for backup operations by up to half, depending on data compressibility.

- Channel aggregation

  A single FICON channel can replace multiple ESCON channels. This means there are fewer channels, director ports, and control unit ports, or rather a simpler configuration to manage.

- Reduced infrastructure

  Infrastructure costs can be reduced by the aggregation of multiple ESCON channel fibers onto a single FICON fiber. This is most significant when connecting host systems to remote locations either across the street, across town, or even further.

- Increased addressability

  The current addressing limitations for ESCON are 1024 Unit Addresses (UA) per channel. With FICON, the addressability is increased to 16384 per FICON Channel. This increase can help clients who are currently constrained either by the number of CUs or number of devices attached to their systems.

- Increased distance

  FICON channels now support distances up to 20 km without repeaters and 100 km with repeaters without data rate droop.

FICON Express2, a new generation of FICON features, is designed to satisfy the requirements of bandwidth hungry applications. It offers increased connectivity,

doubling the number of channels in the same amount of physical space as FICON Express. It also provides up to 240 FICON Express2 channels on z990 and up to 80 channels on z890. Plus FICON Express2 offers an increased number of start I/Os over FICON Express and increased throughput. FICON Express2 is designed to deliver up to 270 MB/sec of throughput, depending upon application environment. FICON Express was designed to deliver up to 170 MB/sec of throughput.

As zSeries has increased its CPU performance capabilities, FICON has increased its I/O performance capabilities, to continue to deliver a balanced system design. In addition, FICON's greater distance and bandwidth capabilities over ESCON makes it an essential and cost effective component in many distance and disaster recovery solutions. FICON Express2, with double the connectivity in the same amount of physical space, and increased performance, can continue to satisfy the bandwidth demands of your applications.

FICON features offer the opportunity to substantially consolidate and simplify ESCON channel path configurations, particularly the complex configurations found in large z/OS environments. This is due to native FICON's increased attached volume addressability, multiplexing capability, and director and attachment availability. It is also due to the sharply increased throughput capability of FICON, FICON Express, and FICON Express2 as implemented on zSeries servers.

FICON directors enable channels and ports to be fiber shared. They decrease cross-site connections and reduce implementation cost for disaster recovery with Geographically Dispersed Parallel Sysplex™ (GDPS®) and remote copy. This translates into faster access to your DB2 data, faster backups of DB2 databases and allows your disk control units to be farther away from your host processor.

Laboratory measurements demonstrate the improvements that FICON channels can provide for various DB2 workloads when comparing to ESCON channels. These improvements can result in benefits for DB2 transaction workload, DB2 logging, DB2 queries, and DB2 utilities.

3.10.3 VSAM data striping

Striping is a technique to improve the performance of data sets which are processed sequentially. This is achieved by splitting the data set into segments or stripes and spreading those stripes across multiple volumes. This technique has been available for non-VSAM data sets since DFSMS/MVS® Version 1 Release 1. It became available for VSAM data sets with DFSMS in OS/390 V2R10. DFSMS implements VSAM data striping by spreading control intervals in a control area across multiple devices.

VSAM data striping has these characteristics:

- Data sets must be DFSMS managed and in extended format. You can use the AMS REPRO command to convert existing standard VSAM data sets to extended format or to revert back to standard VSAM.
- It is available for all VSAM data set organizations, including linear data sets. VSAM striping is limited to 16 stripes.

DB2 active log data sets are good candidates for VSAM data striping. In heavy update environments, the DB2 active log can become the limiting factor from a performance point of view. Laboratory measurements have shown that the DB2 log band width was 8.2 MB/sec on ESS model E20. It then became 11.6 MB/sec on ESS model F20. Striping has further increased throughput to 27 MB/sec.

3.10.4 Geographically Dispersed Parallel Sysplex

Geographically Dispersed Parallel Sysplex extends zSeries and DB2 availability through a service. This services combines Parallel Sysplex with different disk level redundancy techniques such as Peer to Peer Remote Copy (PPRC) and Extended Remote Copy (XRC).

A GDPS/PPRC implementation consists of a Parallel Sysplex spread across two sites separated by up to 100 kilometers of fiber with one or more z/OS systems at each site. By managing both planned and unplanned exception conditions, GDPS/PPRC can help maximize application availability, protect against metropolitan area disasters, and provide business continuity with near transparent disaster recovery.

The HyperSwap™ feature enables DASD failover to be transparent to DB2. See 3.10.5, “HyperSwap” on page 75.

GDPS/PPRC provides the Parallel Sysplex cluster continuous availability benefits. It significantly enhances the capability to recover from disasters and other failures, as well as manage planned actions. GDPS/PPRC is a near transparent disaster recovery solution. Its recovery time objective is less than one hour.

GDPS/XRC provides an automated solution to protect from regional disasters. Due to distances typically involved in disaster recovery solutions, it exploits XRC as a disk mirroring technique to avoid performance impacts on production systems. The Parallel Sysplex does not span the production (site 1) and recovery sites (site 2). The recovery site can be located at any distance from the production site. GDPS/XRC is a disaster recovery solution. Its recovery time objective is less than two hours.

3.10.5 HyperSwap

The HyperSwap function is a part of GDPS/PPRC (see 3.10.4, “Geographically Dispersed Parallel Sysplex” on page 74). It provides the ability to transparently switch the applications I/O activity to the secondary PPRC volume for both planned and unplanned reconfigurations. HyperSwap has been designed to provide the swap of large numbers of volumes quickly.

The important ability to re-synchronize incremental disk changes in both directions between primary and secondary PPRC disks is provided as part of this function.

This feature is remarkably significant because it removes disk subsystems as a single point of failure. DASD failure is transparent to DB2. The dynamic switch to secondary disk and redriving of the I/Os is done by GDPS and z/OS. Consider this for additional protection of DB2 catalog and critical business tables.

3.11 zSeries 990

The z990 is the IBM server for the on demand world. Clients of every size, and across every industry, are looking for ways to make their businesses more productive and more resilient in the face of change and uncertainty. They need the ability to react to rapidly changing market conditions, manage risk, outpace their competitors with new capabilities, and deliver clear returns on investments.

Welcome to the on demand era. To meet the challenges of constant change, companies must go beyond simply integrating business processes. They must have the flexibility to provide products and services to customers on demand. But success in this environment takes a new approach to business, and a new kind of IT infrastructure to support it.

The z990 delivers enriched functions for the on demand operating environment. zSeries is the enterprise-class server optimized for mixed workload integration, high transaction processing, and data serving for the on demand world.

In implementing the z/Architecture with new technology solutions, the zSeries is designed to facilitate IT and business transformation and reduce the stress of growing business to-business (B2B), business-to-customer (B2C), and business-to-employee (B2E) computing requirements. The zSeries represents an advanced generation of servers that feature enhanced performance, support for zSeries Parallel Sysplex clustering, and improved hardware management controls. They also feature open and industry standard technologies, such as Linux and Java, and innovative functions to address infrastructure simplification and On Demand Business operations.

The z990 is designed to provide a balanced system and improved price and performance. From processor storage to the system's I/O and network channels, end-to-end bandwidth is designed to deliver data where and when it is needed.

With the superscalar microprocessor and the CMOS 9S-SOI technology, the z990 is designed to further extend and integrate key platform strengths such as dynamic flexible partitioning and resource management for mixed and unpredictable workloads.

The advantage of these machines is more powerful and additional processors. In preliminary measurements of running DB2 UDB for z/OS Version 7 and Version 8 at Silicon Valley Lab, the range of improvements in processing for DB2 was expected to be in the range of 1.5 to 1.6 times faster than a z900 2C1.

The IBM On/Off Capacity on Demand capability of the z990 is designed to provide even greater flexibility. It allows you to turn on additional, temporary system resources at busy times of the year and then turn them off when they're no longer needed. This can give you exceptional control over costs while you meet your dynamic capacity needs.

The z990 can provide reserved emergency backup CPU capacity through its Capacity Backup Upgrade feature. This feature gives extra capacity to your operation in emergency situations where you have lost capacity in another part of your establishment and need to recover capacity on a designated z990 system.

The z990 is an open flexible server. It offers the flexibility to manage numerous operating systems on a single server, including z/OS, z/VM, VSE/ESA™, TPF and Linux on zSeries, and Linux on S/390. The z990 supports both ESA/390 and z/Architecture modes. Since both of these architectures are supersets of previous architectures, they are designed to provide seamless support for previous architectures dating back to the S/360. As such, they are designed to support existing application investments without change to realize the benefits of new technologies with the z990. z/OS and the z990 are designed to automatically adapt to support the relevant 24-bit, 31-bit and 64-bit addressing schemes.

The zSeries is designed to deliver the highest level of application availability required in today's on demand environment. It offers extremely high reliability. It is endowed with self-healing and self-managing features so your system can fine-tune itself to help provide the level of performance required for On Demand Business operations.

DB2 supports unique use of the z/Architecture instruction set. It supports several instructions that help to improve reliability, performance, and availability.

3.12 64-bit virtual storage

As application workloads and data volumes continue to grow, the demands on available memory space becomes an issue to rival those on processing power. Constraints on performance cannot be relieved by increasing CPU power alone. Time spent juggling objects in and out of central memory becomes a significant factor.

The introduction of 64-bit virtual addressing was a major change and enabled a new architectural level for the zSeries product line. DB2 is the first product on zSeries to make the corresponding move to the z/OS prerequisite and 64-bit addressing.

Ever increasing business pressures to compete in a fast moving on demand marketplace introduce significant demands on information management systems to cope with large volumes of data in real time. Web services, On Demand Business, and traditional OLTP workloads compete for access to processor resources. These resources must be supported by efficient and timely access to data with little or no delay.

The latest z/Architecture hardware and zSeries operating system deliver a significant benefit in the introduction of 64-bit addressability and enhanced resource management components. Virtual storage constraints, not uncommon in the large strategic and critical installations in which you find DB2 for z/OS, can be relieved in support of continued transaction growth and high volume business operation.

Faster hardware technology is inevitable in the race to keep up with increasing business demands driven by customer and internal trends towards instant information access. As the volume of data and transactions increases, it places additional demands on storage solutions and the speed with which they can deliver data into the system. Physical storage access continues to be the most costly in terms of performance. The closer the data is to the processor, the faster it can be served. The vast increases in virtual storage capacity, backed by increased real storage limits, enable more data to be cached in memory in readiness and anticipation of application demands.

Collection, interpretation, and manipulation of data has become a critical foundation for business decision making and knowledge management. More information is being harvested from core business operational processes to feed the appetite of data warehouse and datamart installations. Some of this data can be captured periodically with little imposition on core applications. And some has a currency aspect which requires it to be captured and made available in relative real time.

Business consolidation continues in popularity as enterprises seek rapid growth, complementary market potential, and competitive advantage. Such activity introduces additional workloads into the system, often with differing characteristics and demands to the existing applications. Increased transaction throughput, data volumes, and user base must be assimilated into the IT infrastructure framework as seamlessly as possible with minimal disruption to business process and services.

All this new business and technology evolution leads in only one direction: the need to process and manipulate larger volumes of data within the constraints of a real-time service-level agreement. The burden on the IT architecture to anticipate an application's data requests and maximize the potential for memory resident data access will only increase. The latest enhancements for DB2 UDB for z/OS underpinned by the advances in z/Architecture and z/OS are critical to the success of businesses as they strive to compete and retain a competitive edge.

With DB2 Version 8, now you can have much bigger buffer pools and a lot more central memory on your machine to back up those buffer pools. That means you can scale an individual DB2 system to much larger numbers. You don't expect data pages to move in and out of the buffer pool that much more rapidly. But you can have a lot more pages in the buffer pool, which lets you handle more concurrent users.

The z990 provides more central memory up to 256 GB. DB2 uses memory to allow scaling with large workloads. The additional memory and value for the money means that you can add the needed memory to improve scalability, availability, and ease of use. You have one large virtual space, rather than a number of hiperspaces and data spaces in addition to the address spaces. Several limits for storage use are lifted. While using more than 2 GB in one address space once meant an outage, it now only requires real storage tuning. You are unlikely to experience outages due to lack of virtual storage.

The 64-bit real storage support eliminates expanded storage, helps eliminate paging, and may allow you to consolidate your current systems into fewer LPARs or to a single native image. With 64-bit virtual storage, outages are reduced. Monitoring is simpler, so you can scale your DB2 applications farther using more memory, more effectively. This brings significant technological advancement to ease the burden and help the enterprise applications scale up to the challenges ahead.

3.13 HyperSockets

When applications go through Internet servers, application servers experience some latency from the additional hops through TCP/IP networks. If the application servers are in zSeries LPARs, then up to 16 high speed TCP/IP networks can be created via the memory bus, without using CPU cache. This does not affect other activity.

The benefits of HyperSockets include faster virtual server growth, savings of adapter or cage slot costs, and simpler configuration. These HyperSockets enable a “network in memory” at high speeds. They are also highly secure, available, and simple. Your applications are closer to DB2 and have a faster and safer connection to you DB2 data.

3.13.1 Specialized processor for Linux

The Integrated Facility for Linux (IFL) is a CP dedicated to Linux workloads. The attractively priced IFL processor enables you to purchase additional processing capacity exclusively for Linux workloads, without affecting the MSU rating or the zSeries model designation. This means that an IFL does not increase charges for zSeries software running on general purpose (standard) processors in the server.

The IFL hardware feature is isolated from general use. It is supported by z/VM, the Linux operating system, and Linux applications. It cannot run other IBM operating systems.

Running Linux on the zSeries mainframe is:

- An ideal platform for efficient and effective infrastructure simplification
- An ideal Linux platform to provide the security, flexibility, scalability, resiliency, and integration needed to respond to the demands of this dynamic marketplace

The IFL enables you to exploit the integration capabilities and core strengths of your zSeries server for Linux at a cost-effective price point. It provides a full capacity processor, independent of the capacity of the general purpose CPs. And IFLs are designed to operate asynchronously with general purpose CPs.

Linux applications accessing DB2 can greatly benefit from using specialized Linux processors and being located close to DB2. They can leverage legendary zSeries mainframe robustness and processes.

3.14 zSeries Application Assist Processor for Java

On demand is everywhere. Customer expectations continue to grow and business requirements are changing more rapidly than ever. The Web has become the tool of choice for commerce, and the Java language continues to take the programming world by storm. Considering this, organizations that want to remain competitive and responsive to their customers, partners, employees and stakeholders must develop and deploy new strategic business applications more quickly, efficiently and cost effectively than ever before. To develop, maintain and enhance these applications on a regular basis, sophisticated tools and a strategic approach are required.

Java adoption and application development continue to accelerate in the marketplace as an open, highly strategic, and productive programming model. Its design and richness of function make it the tool of choice for programmers to write robust, bug-free code that can run on almost any platform.

Applications developed using Java can be fairly unpredictable. They typically require significantly more processing and memory resources (often two to three times more) than traditional applications. This is due to Java's platform-independent architecture characteristics, which generate applications with high levels of abstraction, code generation, and reuse that result in longer path lengths, larger memory footprints, and other potential inefficiencies. IT budgets, unfortunately, are not growing exponentially. They are forcing customers to seek more cost effective and productive ways of deploying new Java technology-based applications.

The zSeries is already well known as one of the safest, most secure, most available, most scalable, and most reliable platforms on which to keep critical databases. By delivering the zAAP, IBM provides our clients with compelling and economic reasons to place their key Java applications there as well.

Leveraging the zAAP can be the smartest move your company makes. Why?

- zAAPs enable you to integrate and run On Demand Business Java workloads on the same server as your important database for high performance, reliability and availability. This also helps to simplify and reduce the infrastructure required for Web applications.
- zAAPs can be exploited in the security-rich IT operating environment that only z/OS can provide.
- zAAPs may help to increase system productivity by reducing the demands on general purpose processors, making capacity available for reallocation to other workloads.

- zAAPs allow you to run Java workloads, such as WebSphere for z/OS for a significantly lower total cost of ownership than previously possible.
- A Java workload is transparently executed on the zAAP processors. You don't have to change your applications at all.
- zAAPs are priced attractively at only USD $125000 per zAAP.
- IBM does not impose software charges on zAAP capacity.
- zAAPs are designed to operate in conjunction with the general central processors to execute Java programming under control of the IBM Java Virtual Machine (JVM).

IBM provides a tool that allows you to learn the potential for Java execution on zAAPs inherent in your existing applications. This tool gathers usage information about how much CPU time is spent executing Java code which can potentially execute on zAAPs.

Both zSeries CPs and zAAPs can take advantage of zSeries capacity on demand. During spikes in demand, this ability to scale up and out quickly can be the difference between flawless execution and slow response times or even system crashes. A single zSeries server can scale up to millions of transactions per day or scale out to manage tens to hundreds of virtual Linux servers.

In summary, the zAAP delivers a specialized z/OS Java execution environment with the integration advantages and traditional strengths of the zSeries platform. Its attractive price and infrastructure simplification advantages, combined with zSeries virtualization, applications availability and autonomic management capabilities, make the zSeries the ideal environment for On Demand Business computing. In addition, the traditional mainframe benefits that zSeries also offers, renowned performance, reliability, security and scalability, can help support on demand operations.

zAAP is an industry first. It is currently the only specialized processing unit for Java today.

3.15 Summary

The z990 can help almost any size company in any industry become more responsive, variable, focused, and resilient, the key components of an On Demand Business. With its balanced design that ensures high levels of utilization, its ability to automatically allocate resources where and when they are needed according to business priorities, its flexibility to adapt to change and, its incredible reliability, the z990 is the data center for the On Demand Business world. Consider that over time, in the mainframe environment, fewer technical staff are needed.

z/OS, the IBM flagship mainframe operating system, and DB2 are both designed to maintain high availability through rolling upgrades. This keeps your applications and data more available and allows business as usual.

Improved disaster recovery is built upon changes in z/OS, DFSMS, GDPS, and ESS.

Together, IRD and Parallel Sysplex are designed to provide the flexibility and responsiveness required by On Demand Business workloads. These capabilities can significantly help minimize the complexities inherent in broadly distributed architectures and can help lower your total cost of ownership. The z990 extends the power and versatility of z/OS's innovative Workload Manager from the edge of the network to the heart of your data.

zSeries servers are designed to handle sustained peak workloads utilization of 100% without service level degradation to high priority workloads. On a z990, individual component reliability, or MTBF, is designed for up to 50000 to 60000 hours. This equates to an MTBF measured in decades for zSeries servers. This is an extremely reliable machine.

Your data is growing at an unprecedented rate. You have a fundamental need to secure and protect data. Your workloads are increasing and sometimes unpredictable. This is true today and in the future. DB2 can handle massive amounts of data today and allows clients to continue to grow as their business grows in the future.

Let DB2 and zSeries hardware and software take the complexity out of managing and processing this data by providing you with 24x7 RAS and solid backup. The highest transaction processing workloads are achieved in this environment. zSeries and DB2 provide components that allow you to continually manage mixed workloads. You can maintain high availability in a dynamic business world, while using advanced hardware features to speed up encryption and compression.

The new zAAP, available on z890 and z990 servers, delivers a specialized z/OS Java execution environment. It allows you to extend the value of your existing zSeries investments by strategically integrating new Java technology-based workloads alongside your DB2 data and core business applications on the mainframe.

You can also consolidate existing Java application code onto the zSeries platform. In doing so, you benefit from the resiliency of the mainframe environment, without anticipated modifications to Java applications. The mainframe environment allows you to maximize people and system resources. It also enables you to maximize utilization of existing capacity.

DB2 is the premier exploiter of zSeries and z/OS. DB2 is vertically integrated with the z/OS operating system, zSeries hardware, DB2 tools, WebSphere, Workload Manager, IRD, and IBM storage products. Faster processors and larger memory can make your business more responsive. Security has improved and is more sophisticated.

DB2 leverages security functions in the operating systems and cryptographic functions in the hardware to deliver leading edge security and privacy at low levels of granularity. Security is a prerequisite for high availability. z/OS has the industry's most stringent processes, tools, and techniques for access control and asset protection. It maintains both the availability and the integrity of resources.

DB2 and zSeries hardware and software handle diverse workloads for important applications, while maximizing resource utilization. DB2 data sharing and the Parallel Sysplex architecture are engineered with availability as the core requirement. They allow you to open your door to the world and concentrate on growing your business.

It is important for companies adopting on demand strategies to team up with a trusted, experienced partner to help pave the way to success.

WLM and IRD provide for the continuous, automatic and dynamic adjustment of resources for your DB2 applications workloads. They direct resources to priority work and ensure that all your business goals are met. They also ensure that such system resources as processor and DASD channels are optimally used.

This DB2 and zSeries partnership does a great deal for you today. As market conditions shift, this unique combination allows your business to expand and scale either vertically within a server or horizontally within the Parallel Sysplex. DB2 has rich functionality, your applications are extremely available, and your data assets are secure. IBM will bring more autonomic features to allow systems to protect themselves and *heal* themselves. This enables you to extend the number of important decisions which can be made without thinking about them.

DB2 data sharing and the Parallel Sysplex infrastructure offer the foundation to enable the highest level of availability for your business data. They provide the redundancy to survive outages at various levels.

DB2, zSeries, and z/OS are the components that enable a business to have a tangible manifestation of its brand in a highly safe and secure IT infrastructure. This centralized environment is mature and has sophisticated management capabilities requiring less staff year over year.

This deep synergy between DB2, zSeries hardware, and z/OS has enduring value, and is an ongoing IBM strategic design objective. You have now learned the compelling reasons for using DB2 and zSeries to keep your business running. They provide the ultimate DBMS solution for Web-enabled enterprise and heterogeneous workloads to give you a competitive advantage.

4

Business resiliency

Don't you have enough to worry about? Today's IT shops are complicated. You likely support multiple types of hardware and software solutions. Although mean time between failure (MTBF) continues to improve, numerous components present the opportunity for a single point of failure.

DB2 for z/OS and the zSeries platform offer the foundation of resiliency for your environment, keeping your IT assets available to your internal and external customers. DB2 for z/OS and the zSeries platform provide the architecture to maintain maximum availability, when problems occur, and to adjust to changing and unpredictable workloads.

Several competitive products offer features that make for good evaluation checklist material. These items do not necessarily make your database infrastructure more resilient, but often provide for ease of use. DB2 for z/OS has, and is, incorporating features to enhance usability. This also makes it easier to migrate from competitive databases or to port applications to the zSeries platform. By making the migration, or port, easier, the ability to obtain enhanced application resiliency is an option for more applications from other platforms.

DB2 for z/OS Version 8 has broken through many limits and barriers of the past to deliver ground breaking benefits in all aspects of business resiliency. Business resiliency is expanded through the following topics:

- The DB2 business computing environment
- DB2 for z/OS and zSeries resiliency

4.1 The DB2 business computing environment

Some of the components that a CIO, CTO, architect or IT manager must concern themselves with on a daily basis include:

- Java
- Web services
- Distributed applications
- Open standards
- Security
- Unpredictable workloads
- Service-oriented architecture (SOA)
- Regulatory pressures
- Application development
- Time to market
- Critical applications
- Continuous availability
- Backup and recovery
- Autonomics

They must also face these questions:

- Are you meeting today's business demands?
- Are you positioned to meet those of tomorrow?
- Are your solutions for today and tomorrow two different strategies?
- How comfortable are you with the future of your current platform investment?

DB2 UDB for z/OS and the zSeries platform provide the deep synergy necessary to enable business resiliency for the concerns listed above, both for today and for tomorrow. See Figure 4-1.

Today, DB2 is universally accepted as the premier database system for the zSeries platform. Known for its high availability and rock-solid security, DB2 has been a core component of critical business applications over the last two decades.

The evolution of DB2 has maintained its status as the premier mainframe database, and has positioned it as *the* enterprise data serving platform. Whether you are supporting existing COBOL applications, or developing and running Web-based, customer interfaces, DB2 provides support of open standards for developing and accessing all of your information resources. At the same time, it offers rich features and unsurpassed availability and security.

The On Demand Business model requires an operating environment built on open standards to enable rapid and cost-effective innovation and reconfiguration. Among other things, the infrastructure to support an On Demand Business

environment must be available, scalable, and secure. DB2 is a core component of this infrastructure.

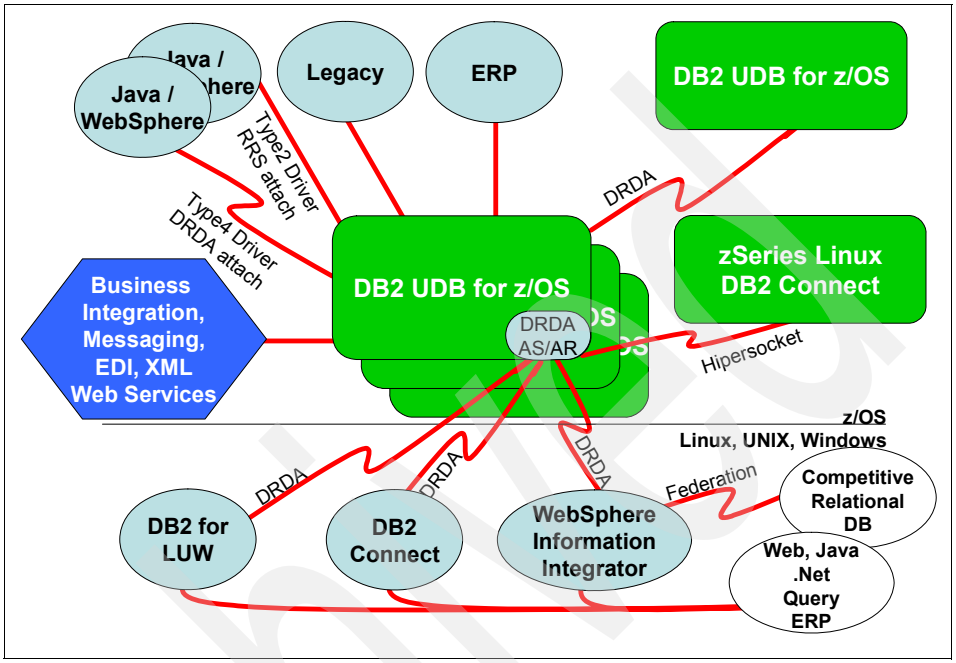

Figure 4-1 Business resiliency landscape

Integration

DB2 supports new Java and Microsoft workloads, and integrates with Web services and message queuing technologies. It manages structured and unstructured data. Heterogeneous, relational, and non-relational information is supported via WebSphere Information Integrator.

Open

IBM is deeply committed to Linux and standards for Java, Extensible Markup Language (XML), Web services, and distributed, multivendor database interoperability. DB2 UDB for z/OS provides a Java Universal Driver. It is the same code base across platforms, integrated XML support, object relational capabilities, support of Web services, and strong Distributed Relational Database Architecture (DRDA) support for serving and accessing data in a distributed environment.

When opening your environment, security becomes ever more important. DB2's integration with the zSeries platform provides solid access and privilege control.

It does this from connection authentication, to the request for data in a single row, with hardware assisted encryption for storage and transmission of information.

Virtualized

DB2's integration with DB2 Connect and WebSphere Information Integrator provides practical support for federated access to information throughout your enterprise. Superior workload management, 64-bit support, and clustered scalability provides for the expansion and utilization of a virtualized information environment. There are also specialized processors available for clustering, Linux, and Java.

Autonomic

Industry leading cost based optimization, integrated space management algorithms, and run-time data access learning all contribute to DB2 as a self-tuning and managing environment. This is an area exploding with opportunity and is discussed in 8.4.1, "Autonomic computing" on page 195.

4.2 DB2 for z/OS and zSeries resiliency

Before we review the key capabilities that support your DB2 UDB for z/OS selection as your enterprise data server, it is important to review a primary philosophy difference. Synergy with the z/OS operating system and the zSeries hardware positions DB2 to *keep your data available*. This is opposed to competing platforms and databases, which are focused on *recovering quickly*.

The ability to recovery quickly is an important attribute. However, your business goals are better achieved with a solution that focuses on eliminating outages, planned or unplanned. This core design difference resonates through to your business applications and maintenance processes, and therefore, to your ability to meet business goals and service-level agreements.

Key platform synergy items drive business resiliency to a level beyond the reach of competitive offerings.

- Hardware assistance for compression, encryption, sorting, and Unicode conversion
- Exploitation of 64-bit memory addressibility
- Exceptional workload management support, governing the priorities of processes running on z/OS, and balancing connections from distributed platforms
- Full utilization of the Parallel Sysplex architecture, enabling a scalable, clustered, incremental growth architecture

- Support for new workloads, using specialized processors for Java and Linux workloads

  Fast communication pipes, *HyperSockets*, are available for distributed connections between other logical partitions (LPARs) on the server. This includes Linux LPARs, which offer a cost saving consolidation option to host your applications and DB2 Connect servers.
- Integration with other core z/OS components, such as those for disk storage management, access control, UNIX system services, and failover
- Tooling and instrumentation for diagnosis and manageability

These qualities are covered in more detail in Chapter 3, "DB2 synergy with zSeries and z/OS" on page 39.

4.2.1 Reliability

In the commonly used abbreviation RAS, the R is for *reliability*. This is the union of the *availability* (A) and *scalability* (S) attributes, that work together to best use your IT investment and meet your business computing goals. A reliable system is available to respond to the requests of your internal and external customers. It performs to meet your business goals in the midst of unpredictable workloads.

In the mix of these attributes reside challenging technology opportunities that cannot be solved with a single product. DB2 UDB for z/OS, integrated with the z/OS operating system and the zSeries hardware platform, provides the solution.

To complete the RAS abbreviation, the topics of availability and scalability follow. The letter S is sometimes used to identify *serviceability*, another strong zSeries trait. This topic is addressed in 4.2.4, "Manageability and adaptability" on page 97.

4.2.2 Availability

To customers, planned and unplanned outages look the same. They cannot perform a task or complete a transaction. The customers may try to conduct a financial transaction with your company. Will they wait until your site is available, or will they go elsewhere? E-commerce makes the latter option a real possibility. Back office processing may have penalties associated with delays in meeting internal or external deadlines.

A significant component of having a highly available database is to run on a highly available platform. With an MTBF design point in excess of 30 years, the zSeries platform has demonstrated exceptional availability. DB2 UDB for z/OS exploits the zSeries technology, providing exceptional data availability.

DB2: Designed for availability

DB2 is written with high availability in mind. A DB2 system incorporates redundant components to survive localized failures that may occur. Such items as key control data sets and logs are duplexed to keep the engine running. Errors can be detected at a very granular level, so that a single data page can be recovered automatically.

Backup, recovery, and maintenance

When combined with the IBM DB2 for z/OS utilities, maintenance operations can also be accomplished online. The IBM DB2 utilities are a fee-based option for DB2. The utilities help you maintain the integrity of your data and performance of your applications. Third-party tools may also be available to fulfill these roles.

Backup options

You have numerous backup options to provide insurance, in case data needs to be recovered. A recovery can be required due to physical damage to the object such as with a disk failure, or logical damage which is possible with errant application logic.

Hot backups can be taken, where data read and write access is not interrupted. You can also choose from full object copies, incremental copies, or copies at a volume or subsystem level. Should a recovery be needed, the data and indexes can be recovered to current time, or to a point in time prior to the failure. Integrity checks within DB2 ensure to flag objects that have system defined relationships with the recovered objects, so that an administrator can be sure to validate relationship integrity.

Recovery optimizations are incorporated into the tools. They help to apply log records in a single pass and to know if entire log data sets can be skipped because they contain no changes for the object being recovered.

When combined with optional tools, additional backup and recovery options are possible. Changes to a table can be extracted from the log and applied to the last backup to produce a new backup. While post-processing is required to produce this backup, the object does not have to be copied again. Also recovery time can be reduced by having a more current back up of the data. The DB2 tools also allow for application recoveries, by analyzing the logs and reverse engineering the Structured Query Language (SQL) of the application.

Reorganization

Online reorganization and load utilities allow for maintenance with no, or a very limited, outage to the data. Reorganization of the data and indexes is important, over time, to ensure that the data is maintained according to your original design

criteria. This assists in maintaining top performance for both data retrieval and insert operations.

Off-site recovery

You can accomplish off-site recovery in many ways. Some options include:

- Backup copies and logs from the local site can be transported, physically or electronically, to a remote site for utility recovery purposes.
- A DB2 Tracker site can be set up to maintain a delayed copy of your primary site data. After an initial synchronization, logs from the primary site are transported to the remote site and applied to the remote site data. This can help to reduce data loss and the outage window.
- Replication offers a near-real time flow of data to the remote site, where it can be applied according to your timing requirements. The changes can be applied continuously or based on a configurable delay. Replication can further reduce data loss and outage time. Additionally, this approach allows for the remote data to be available for query workloads.
- Geographically Dispersed Parallel Sysplex (GDPS) allows for the extension of the data sharing environment (see “Data-sharing availability” on page 92) to a remote site, for failover purposes.

Online schema evolution

To support the fast-paced, changing business environment, it is necessary to modify data structures in support of new and changing application requirements. In the past, most schema modification operations were disruptive to data availability. This function enhances data availability, as well as benefits the management of your DB2 data.

Traditionally, schema changes have required the following steps:

- Unloading the data
- Generating data definition language (DDL) for the existing object and dependent objects, such as indexes, synonyms, views, referential integrity, check constraints, triggers, and security
- Dropping the existing object
- Recreating the object in the new format
- Possible reformatting and loading of the data
- Recreating dependent objects
- Re-establishing security

With DB2 Version 8, many of these operations can be accomplished online, without interrupting data access from the applications. DB2 tracks the version of

a data row and converts any rows on disk at an older version to the current schema at retrieval time. Therefore information is always seen in its most current state when retrieved. Dependent indexes and views are also changed as part of this operation. These alter operations are accomplished online.

The reorganization utility, such as the online reorganization mentioned earlier, can be used to reformat all rows on disk into the most current format. When the rows are reorganized, the row conversion processes are no longer necessary, until the next online schema change.

System parameters

There are hundreds of system parameters that govern the behavior of your DB2 subsystem. These parameters are called DSNZPARMs. They include the following functions:

- Identifying information for the DB2 subsystem
- Defining various internal memory areas
- Timing and behavior of internal processes, for example, for storage management and optimization choices
- Timers and controls that monitor thresholds, like locks
- Default choices for such items as index structures
- Environmental factors, such as language support

Changes to core definitions, such as these, often require that the database be restarted to enact the change. With DB2 for z/OS Version 8, the majority of new parameters can be changed online. About 75% of all parameters can now be changed on line, maintaining data availability. The remaining DSNZPARMs are mostly definition items that rarely, if ever, change.

Data-sharing availability

To achieve the highest data availability, DB2 exploits the Parallel Sysplex technology of the zSeries platform and z/OS. Up to 32 DB2 subsystems are clustered together to provide a shared data, horizontal growth option. The components of this implementation and comparison with other clustering techniques are discussed in 3.2, “Parallel Sysplex and data sharing” on page 43.

The benefits of data sharing are:

- Incremental growth and near-linear scalability

  This includes extensive parallelism, offering inter-node parallelism across all participating DB2 members, and intra-node parallelism within a single member. Data sharing scalability is covered further in 4.2.3, “Scalability” on page 93.

- Near continuous availability, with the option to keep your data available through maintenance and upgrade windows

  You can maintain data availability when performing such operations as migrating from Version 7 to Version 8.

- Intelligent workload management

  Workloads can be routed to the DB2 member with the best availability for the request.

- Fault tolerance, to survive potential software and hardware outages

  The architecture has the built-in redundancies to maintain data integrity and data availability during component failures. Working with z/OS, failing components can be restarted on surviving portions of your environment, while workload is shifted to maintain availability.

4.2.3 Scalability

Perhaps the most obvious support for scalability is the exploitation of 64-bit addressing. The largest z/Architecture address is 16 exabytes -1. DB2 now supports a 16-exabyte DBM1 address space. Virtual storage is now addressable by DB2, beyond the 2 GB bar. The new virtual storage area above the old 2 GB limit is referred to as *above the bar*. The 16-exabyte limit, potentially addressable by 64-bit, is referred to as the *beam*.

This support has expanded DB2's reach by placing many memory pools, of which buffer pools tend to be the largest participant, above the 2 GB bar and expanding each buffer pool maximum size to 1 TB. Other areas have also been moved above the bar. These includes compression dictionaries, global dynamic statement cache, and sorting pools, giving more room below the bar for thread related storage.

Scalability is also often associated with maximum table size, which is now 128 TB. However, features of DB2 have been available for years that promote the scalability of the zSeries platform.

Data sharing scalability

As mentioned in "Data-sharing availability" on page 92, DB2 data sharing exploits the Parallel Sysplex architecture, providing a clustered solution for superior availability. While this provides for unsurpassed availability, it also provides for near linear, cost effective scalability. DB2's use of Parallel Sysplex and a comparison of clustering techniques are further explained in *Achieving the Highest Levels of Parallel Sysplex Availability in a DB2 Environment*, REDP-3960, and *Achieving the Highest Levels of Parallel Sysplex Availability*, SG24-6061.

A single DB2 subsystem can scale vertically by adding processors. Multiple DB2 members can be clustered together to scale horizontally. Near linear horizontal scalability is achievable because of the use of the zSeries coupling facility (CF) and high speed coupling links. The critical coordination processes of global locking and buffer coherency are managed in the CF. Therefore as you add members above two, the additional overhead averages less than 1%.

The DB2 members communicate intelligently with the CF. CF communications are kept to a minimum until they are needed for coordinated access to the same table or index across members, at the partition level.

A shared data clustering solution that scales provides several benefits:

- Distribution of logging across multiple members
- Distribution of workload across multiple members
- Separate memory limits for each member

  Therefore some of the storage still below the 2 GB bar, such as thread storage, is distributed among the members.
- Incremental growth because members can be added as needed

  These members may be in the same LPAR, in a different LPAR, or on another server, as determined by your resource availability and goals. When adding members to an existing server footprint, you scale your investment as well. Because data sharing is an option, you do not have to purchase it separately for use with DB2.

DB2 data sharing was introduced with DB2 Version 4. A Parallel Sysplex contains, along with the processors that are running the DB2 subsystems, one or more processors that act as coupling facilities. Figure 4-2 shows the basic components of a DB2 data sharing environment.

The CFs run specialized control code that coordinates activity across the multiple members of the data sharing group. The CFs contain three major structures that are used by DB2. All three reside in one or more CFs.

- **Group buffer pools**: These are cache structures that contain data and index pages for DB2 objects that are being accessed across the data sharing group.
- **Lock structure**: The lock structure ensures that the same data is not updated by multiple members at the same time.
- **Shared communications area** (SCA): The SCA is used to coordinate recovery and startup across the group.

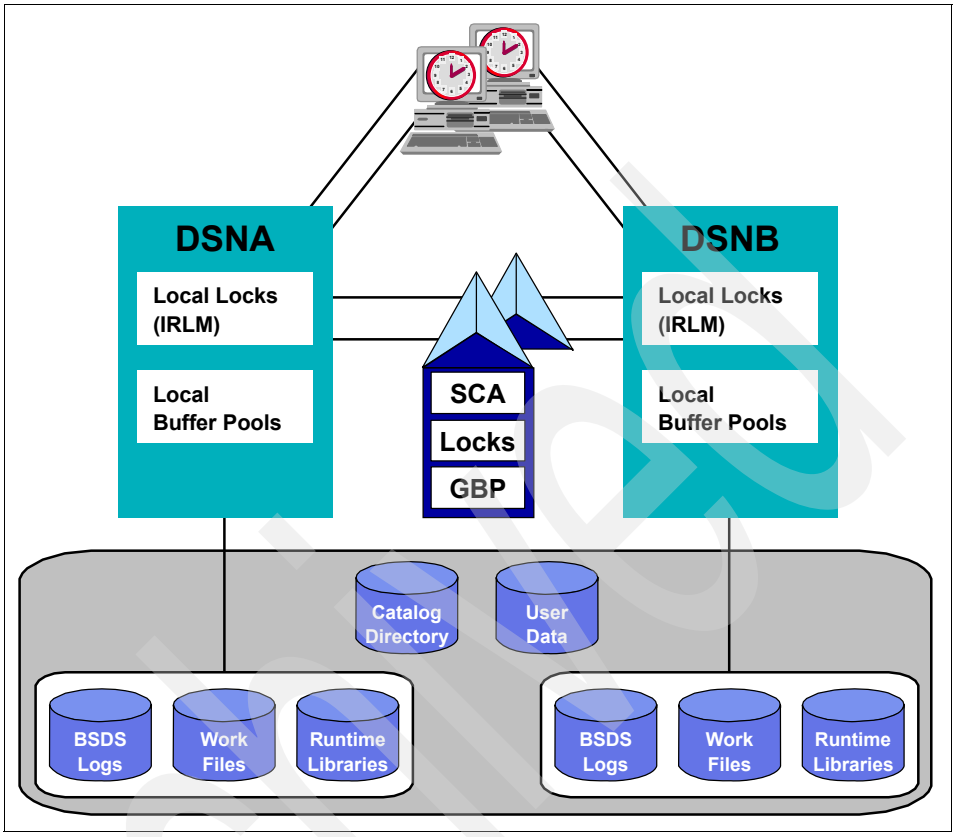

Figure 4-2 DB2 data sharing

Serving information to your enterprise

DB2, as a scaling data server, is a great choice for your existing and new workloads. The Workload Manager (WLM) for this platform is integrated with the zSeries Intelligent Resource Director (IRD), allowing for the continual adjustment of resource allocation based on your business priorities. DB2 and the zSeries allocate machine resources to meet your responsiveness goals. The period aging capability allows shorter tasks to obtain resources, by adjusting the resources used by longer running tasks.

DB2 also benefits by its relationship with WLM for the reasons explained in the following sections.

Workload routing

Distributed connections to data sharing groups are directed to the member with the most capacity. As a task progresses, it may be moved to other members for balancing or fault tolerance.

Stored procedures

Stored procedures are particularly useful for distributed data serving. A request can be made from a remote client, which then is executed on the mainframe, close to the database engine, returning the results to the client. DB2 can modify the number of tasks allocated to a stored procedure address space, or the number of address spaces, based on the workload.

Distributed connections

The Distributed Data Facility (DDF) of DB2 has grown in importance over the years, especially as the number of Web applications have expanded. DDF's role has become as important as a transaction manager. Its features have grown accordingly to effectively handle threads into the mainframe from distributed applications. Such items as connection pooling, two-phase commit, character translations, and advancements to support new SQL have enhanced the product to bring the reliability of DB2 for z/OS to your distributed applications.

Thread management is also assisted by the ability to mark a thread inactive. This reduces the memory footprint but does not destroy the connection. DB2 can then manage more total threads by switching the status between active and inactive. Each DB2 subsystem, or DB2 member, can actively serve 2000 threads, while managing a total of 150000 threads. This further complements the behavior of some client-based applications that may connect once but remain inactive for long periods of time.

DB2 supports open industry standards for communications. DRDA is a standards-based protocol for database interoperability, managed by The Open Group. The DDF address space manages DRDA communications to and from DB2 for z/OS. Either Systems Network Architecture (SNA), or more prominently, TCP/IP communications standards are supported.

Optimization

IBM has been refining cost-based optimization techniques from the beginning of DB2. These algorithms, considering statistics and environmental information, choose the least expensive access path to the data.

The optimizer is the core of what makes a relational database management system (RDBMS) non-navigational. Numerous index, join, sort, parallelism, and rewrite techniques are evaluated. Improvements in optimization techniques and statistics gathering continue to refine this decision making process. All ANSI standard isolation levels are supported. They allow you to choose the degree of locking necessary to ensure data integrity and concurrency for your application or individual statement.

DB2 can statically store these access paths. Bound *packages* make your application execution consistent and secure. Because the SQL in prepared in advance, the costs are not incurred at run time and a known access path is maintained. It is more secure because no table access privileges need be granted to the authentication ID of the application or user. Instead, you simply grant the ability to execute the plan. When binding a package, you can also choose to have the SQL re-evaluated at run time. In this case, a subsequent lower cost incremental bind is performed, and the query can be optimized for the specific data referenced in the SQL.

Dynamic SQL can also be managed at run time. Preparation of the SQL takes place at the time of execution. You can then save this access path in cache and reuse it if you run it again. There are also bind options to run SQL as dynamic, but use the table access privileges of the administrator binding the package. This eliminates the need to grant explicit privileges to the end user for that application.

Intelligent lock management

Working with WLM, DB2 can promote the priority of a task which has exclusive access to a resource needed by a waiting higher priority task.

4.2.4 Manageability and adaptability

Manageability and adaptability are a key focus area of DB2 for z/OS and the platform. As the features and options of this environment continue to expand, proper tooling is required to address the administration, tuning, and diagnostic tasks for the environment. These capabilities assist with the productivity of your existing skill base and building new administration skills.

Governors

Reactive and predictive governors can be configured to protect your environment from dynamic SQL that is considered too expensive. Predictive warning thresholds can be set, or reactive thresholds can stop queries that exceed their allocation of resources.

Space management

DB2 can manage VSAM extent allocations using a sliding scale algorithm. This takes some of the space management burden from your database administration staff. As your data grows, DB2 adjusts additional VSAM allocations so that you achieve optimal utilization of the storage.

Instrumentation

DB2 provides internal code points to flexibly externalize internal information for such activities as accounting, diagnosis, and tuning. This reporting is accomplished through trace classes and specific component identifiers. The output created for these categories is used by tooling and user-written programs to analyze DB2 usage, access, and performance.

- **Statistics**: This tool provides information reports about the usage of DB2 address spaces. Information such as this can be useful in capacity planning, performance, and trending analysis.
- **Accounting**: This tool reports application-related resource information. You can use it for application performance and charge back.
- **Audit**: This tool collects information about DB2 security controls. You can use the report to ensure appropriate data access.
- **Performance**: You can use the information from this tool to identify suspected performance issues and to perform tuning.
- **Monitor**: The monitor allows for the collection of information to be used by tools or user-written programs, such as the DB2 Performance Expert for z/OS.

Utilities

IBM and third parties provide utilities for DB2 that allow you to perform:

- Backup and recovery for data availability
- Reorganization and rebalancing to maintain application and database performance
- Loading for the introduction of data into your databases

- Integrity checks to ensure defined business rules are maintained as structures change and data is loaded
- Diagnosis and repair utilities to resolve problem situations

IBM is advancing in performance and autonomic capabilities by incorporating:

- Decision points, to determine if a utility is needed based on defined thresholds
- Wild carding to have a single utility encompass multiple objects and resistant to change as objects are added or removed
- More parallelism and performance
- Back up and recovery options at the subsystem level

Regardless of which company you choose to be your DB2 utility vendor, IBM delivers the IBM DB2 utilities with your DB2 for z/OS shipment. If you choose non-IBM utilities, IBM Support will provide you with temporary use of the IBM utilities if warranted by a support situation.

Tools

IBM and several third-party vendors offer DB2 tools to assist in the administration and operation of your DB2 environment. Tools provide learning opportunities as well, in that a new DB2 administrator may know what they want to accomplish but may still be learning how to perform the task. We mention the IBM DB2 tools and associate them with two categories: those delivered as DB2 features at no extra charge and those that are separately priced.

DB2 Management Client Package

The DB2 Administrative Client includes a set of graphical tools that provide a spectrum of basic capabilities (Figure 4-3).

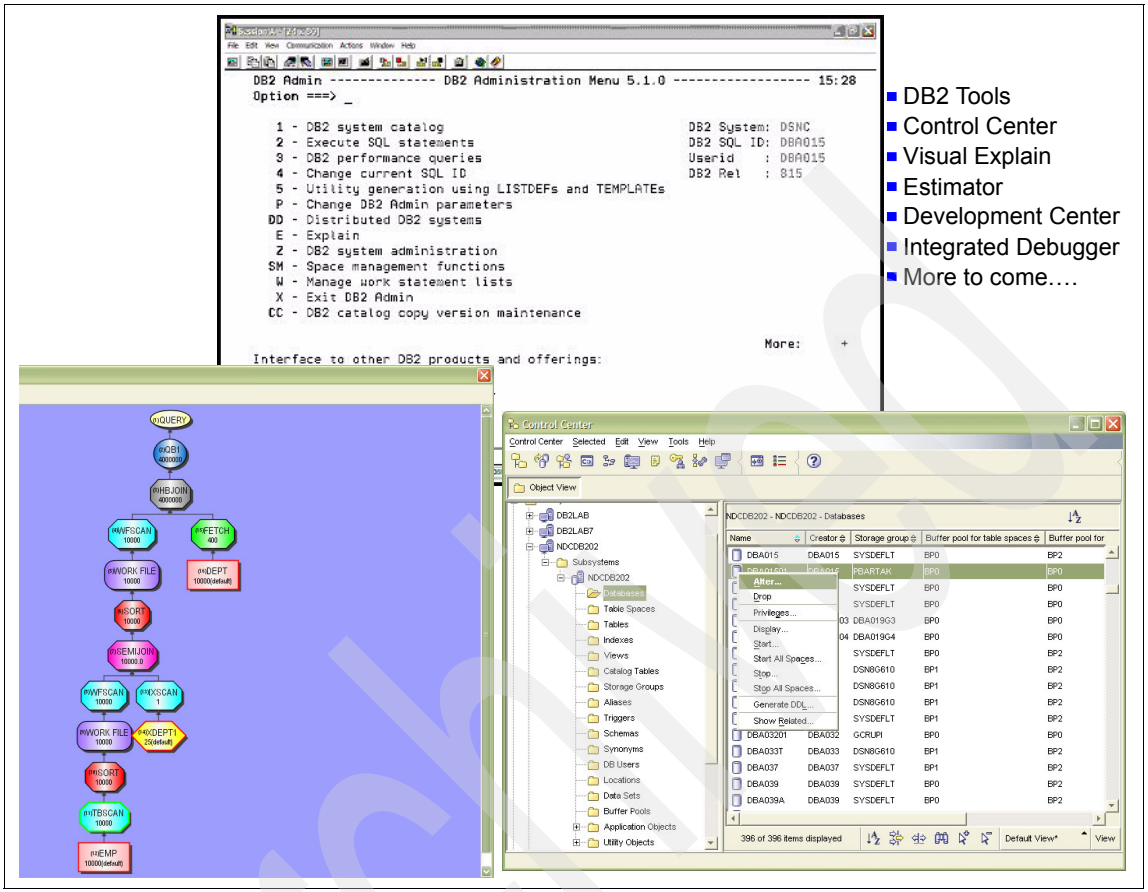

Figure 4-3 DB2 Control Center, DB2 Administration Tool, and Visual Explain

To use the Administrative Client, you must order and install the optional, no-charge IBM DB2 UDB for z/OS Management Client Package (FMIDs HDAS810, JDB881D) feature. To fully use the capabilities, there are some installation requirements for the z/OS platform.

The Administrative Client includes these items:

- **Control Center**: Allows for catalog viewing, some object maintenance and generation, data sampling, configuration, and utility operations
- **Replication Center**: Administers replication definitions, operations, and monitoring across platforms

- **Development Center**: Is used to develop Java and SQL stored procedures

  An integrated debugger assists in testing your SQL stored procedures in DB2 for z/OS.

- **Visual Explain**: Can be invoked from the control center or installed separately

  This tool offers a graphical display of access paths and assists with explains from dynamic statement cache. Statistics for the access path are available, and you can save results for later comparisons. It also produces an XML document of an access path which can be shared with IBM Support to assist with serviceability.

  DB2 Visual Explain Version 8 also works with DB2 UDB for OS/390 and z/OS Version 7. You can download it from:

  `http://www.ibm.com/software/data/db2/zos/osc/ve/index.html`

- **DB2 Estimator**: A separate tool that you can download from the Web

  This tool allows you estimate the resource requirements of an application by modeling workloads, database structures, and transaction volumes. DB2 Estimator Version 8, is available from the Web at:

  `http://www.ibm.com/software/data/db2/os390/estimate/`

IBM DB2 and IMS Tools

Smart tools together with an intelligent database engine, capable of autonomic self-healing and self-tuning, help to keep your business at optimal performance. By taking the administrative burden away from specialist staff, such tools enable their expertise to be used better in strategic activity in support of the business' bottom line.

This section presents only the current list of available DB2 tools in Figure 4-4. For details about functions and releases, refer to:

`http://www.ibm.com/software/data/db2imstools/`

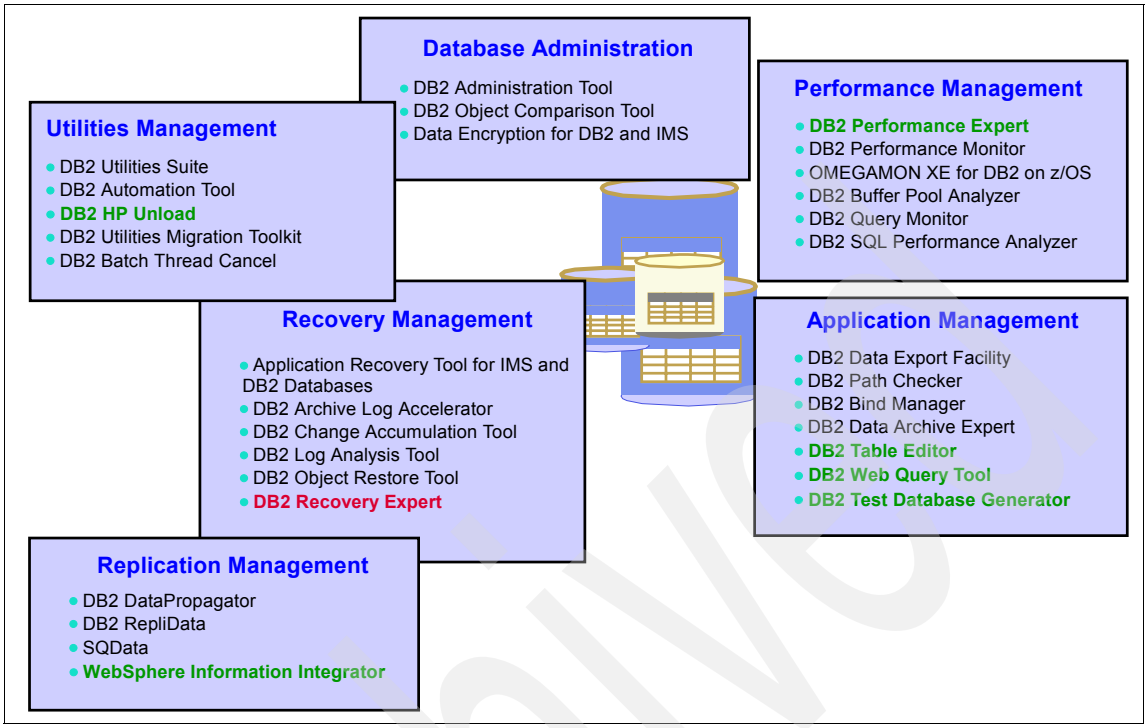

Figure 4-4 IBM DB2 Tools for z/OS and Multiplatform product portfolio

The optional tooling enhances productivity in the areas discussed in the following sections.

Administration

The administration tools assist with the creation and maintenance of DB2 objects, authorities, and operational parameters. These tools also assist with discovery, by providing the framework to mine the DB2 catalog information. See Figure 4-3 on page 100. Catalog comparison features allow for the movement of schema differences between environments.

Operational enhancement

Facilities are provided to extend DB2's capabilities. Examples include logging and recovery enhancements. These tools can also optimize your environment by determining when to run utility maintenance.

Replication

Data replication fulfills a variety of requirements from the population of reporting warehouses or data marts, to the maintenance of hot sites for high availability. Copying full data sources, or moving incremental changes, is possible across the DB2 family and the major relational competitors.

With the more recent introduction of WebSphere Information Integrator Replication for z/OS (Q Replication), DB2 data can be transported via WebSphere MQ Series. Q Replication enables the propagation of change information to a stored procedure or to XML for event publications. Q Replication enables lower latency data movement to better support high availability uses. It enables event publishing to facilitate business integration options with the support of XML publications.

Event publishing has been expanded to other non-relational z/OS data sources. WebSphere Information Integrator Classic Event Publisher captures, formats, and publishes XML messages from IMS, CA-IDMS, and VSAM events via CICS.

Application development and management

These tools assist with the creation, reporting, and optimization of access path. They can also assist with the generation of test databases, test data, and the pruning of production data.

Performance management

Monitoring and simulation software assist in gaining a view to what is happening inside your DB2 system and what may happen if you change parameters or workload.

Query tools

In addition to Query Management Facility, the traditional feature of DB2 mentioned with all its new components in 6.10, “DB2 Query Management Facility” on page 160, other tools provide easy access to your data through the building of a query and formatting of the results. Data reporting and analytics are supported by a variety of tools with varying capabilities that cater to different needs.

4.2.5 Security

DB2 works with the security server of the operating system to *authenticate* a connection request to the database. When the ID for the connection is authenticated, the ID must have *privileges* to access objects within the DB2 system. Privileges can be owned via the creation of an object or by having the privilege granted to the ID.

After you have completed authentication and have determined that you have the privilege to perform the operation requested against an object, multilevel security (MLS) with row level granularity can limit the rows to which you have access. This feature works with the security server to determine an ID's position within a security hierarchy. Access to rows, within an MLS-enabled table, is filtered so the ID is presented only information within the proper security hierarchy.

VSAM data sets on disk ultimately house the DB2 managed data. Therefore, to protect against access to your data outside of DB2, a few encryption options are available.

Encryption tool

IBM, as well as third parties, offer an encryption tool to secure data at the table level. This is implemented via a table exit and uses the same encryption technique for all rows.

Encryption functions

Encryption functions are DB2 built-in functions that allow for encryption at the column and row level. Each execution of the function allows for the specification of a different encryption key, making this a granular technique.

Network encryption

You can optionally enable encryption of the data being served by DB2 for a distributed request. This protects access to your data over the network.

Auditing

To ensure that your data access within DB2 is meeting your security policies, you can turn auditing on at the table level. This may be necessary for IDs that have higher level roles to maintain your system or out of convenience, but should not access the data.

Plans and packages

These objects store the access path for the SQL in an application that is prepared in advance. They provide the ability for you to use the privileges of the ID binding the plan or package, when checking for authority to perform the operation. The bind is typically accomplished via a controlled administrative ID. The privilege to execute the plan or package is given to the users who must run the application without them needing direct access to the data. Therefore this provides fast and consistent access path determination at run time and provides a layer of security.

4.2.6 Accessibility and extensibility

Driving DB2 family compatibility and portability have been key focus areas for the last several versions of DB2.

Extensibility

What makes DB2 a universal database (UDB)? DB2 is branded UDB when it incorporates key object relational features, such as stored procedures, triggers, user-defined functions, user-defines data types, and unstructured data management. Currently, all DB2s, except for VM/VSE, are branded UDB.

Encoding scheme

Support for ASCII tables was added in Version 5, and Unicode was added in Version 7. Version 8 completed the integration of multiple encoding schema support by enabling SQL access to EBCDIC, ASCII, and Unicode in the same SQL statement. The majority of the DB2 catalog tables has been converted to Unicode. Such key DB2 processes as program preparation and SQL parsing are done in Unicode.

New workload support

Type 2 and type 4 Java universal drivers are available for use with DB2 for z/OS, whether your applications run in a zSeries LPAR or off the platform. When your Java workloads run on the mainframe, the z/OS Application Connectivity to DB2 for z/OS (FMID HDDA210) provides the drivers for the z/OS platform.

The type 2 driver is best used when running in the same LPAR as DB2, using DB2's RRS attach facility. The type 4 driver is best used when in a different LPAR using DRDA to communicate with DB2. The Java drivers that IBM delivers off the mainframe are developed by the same team and are the same code base, providing cross-platform compatibility.

Improvements for Open Database Connectivity (ODBC) and .Net continue. DB2 Version 8 delivers extended support and a new SQL cancel capability. See also Chapter 6, "Readiness for new workloads" on page 135.

SQL enhancement

DB2 family compatibility and portability make it easier to move applications to DB2 for z/OS to obtain the business resiliency benefits highlighted in this chapter, without incurring significant application redesign. Figure 4-5 lists some of the recent SQL improvements. It shows the top porting requirements from major independent software vendor (ISV) partners. The top requirements were all incorporated into Version 8 to assist in building and porting applications on DB2 for z/OS.

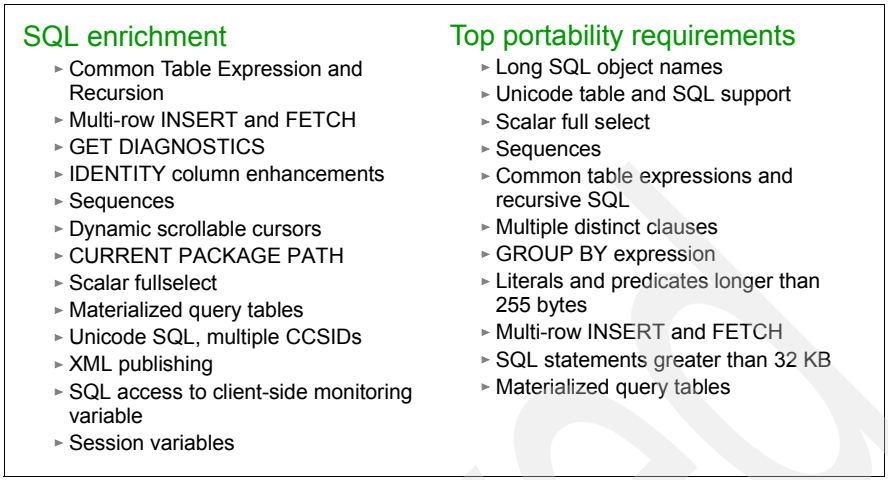

Figure 4-5 SQL enrichment and portability

In addition, the new *SQL Reference for Cross-Platform Development Version 2* guide has been published. You can download this document from the Web at:

`http://www.ibm.com/software/data/db2/zos/v8books.html`

This guide can assist you in avoiding the decreasing number of cross-platform differences, which may cause porting problems.

WebSphere enhancements have been included to make these applications more efficient. READ ONLY USING UPDATE LOCKS and IS NOT DISTINCT FROM have been added to provide enhanced read, update, and null checking.

Stored procedures

Internal and external stored procedures are important in the support of network computing. This encapsulated procedural logic, stored close to or in the database, allows for efficient data access. DB2 supports external stored procedures written in the languages that are available on the server. These include COBOL, PL/I, Assembler, C, C++, REXX, and Java. Internal stored procedures are written in the standard SQL stored procedures language, and are common with the SQL stored procedures on DB2 UDB for Linux, UNIX, and Windows. The DB2 Version 8 enhancement, increasing the SQL statement length to 2 MB, is particularly helpful for these modules.

The DB2 Development Center assists in the building of Java and SQL stored procedures in DB2 for z/OS. It also has an integrated debugger for SQL stored procedures. At run time, Workload Manager supports the adjustment of tasks for the stored procedure workload.

Network computing

Whether using stored procedures or making SQL calls to DB2 for z/OS, DB2 on this platform provides the data serving capabilities you need to support network computing. DB2 has incorporated support of the DRDA open standard for database interoperability. DB2 for z/OS can be a DRDA application requester (AR) or an application server (AS), working with other AS and AR components on or off the mainframe.

DB2 Connect is a common DRDA AR used off the mainframe to enable communications with DB2 for z/OS for data serving. DB2 Connect enables communications from Linux, UNIX, and Windows platforms. And it supports query tools, remote applications, remote Web servers, and information integration, to mention a few of the solutions.

DB2 Connect, working with DB2 for z/OS and the Sysplex Director, enables connection fault tolerance and workload balancing. Connection pooling and DB2's inactive thread management allow for a high performing network computing environment. If DB2 Connect is installed in a Linux LPAR on zSeries, you can use HyperSocket communications between the LPARs. This provides high-speed network communications without going over a network.

As mentioned earlier, DB2 for z/OS can also be an application requester. Communications can be configured with DB2 for Linux, UNIX, or Windows to be the application server. Through the use of WebSphere Information Integrator, competitive relational databases can also be the application server, even if they do not support the DRDA standard for database communications.

When configured, communications with tables in remote databases can be accomplished with SQL. Either a CONNECT statement is issued to direct the SQL, or a three-part table name is used in the SQL. The three-part table name is an attractive option because it further insulates the requester from the data location. When a three-part name is used in an SQL statement, DB2 uses the first component of the name to establish communications with the server configured in the DB2 communications database. An additional layer of insulation can be achieved by creating an *alias* over the three-part name. Then the remote table appears to be a regular local table to the application.

The DB2 load through cursor capability, *DB2 Cross Loader*, can also take advantage of this infrastructure. This feature allows you to flow the results of a cursor to the DB2 load utility. When a three-part name or alias is used, data from databases on distributed platforms can be loaded directly into your DB2 for z/OS tables.

Information integration

Several WebSphere Information Integrator-related products enable data integration. One form of integration, *federation*, involves accessing the data where it resides in real time. You can also achieve this to a limited degree with *DB2 Connect*. DB2 Connect can federate the DB2 family and Informix®.

The WebSphere Information Integrator products that provide relational federation capabilities do not run on the mainframe. These products are DRDA application servers. Therefore, they can work with DB2 for z/OS as an application requester. After configuration, requests from the mainframe can be made to distributed databases. This opens the possibility to have existing programs, such as COBOL programs, accessing competitive relational data from other platforms without any changes in the COBOL code. The benefits to query and reporting tools may be more evident. These tools that run on the mainframe can include distributed data from where it resides.

Replication

Another form of information integration is replication, or the copying of data from a source to a target. This technology supports:

- Complete copies or incremental data movement based on source changes
- Source to target mapping, for vertical and horizontal partitioning of the data
- Data transformations
- Distribution and consolidation topographies
- Scheduling options

IBM offers two primary replication architectures that fit different missions.

SQL Replication

SQL Replication moves data via SQL over DRDA. Changes are captured at a source and staged into DB2 tables. Apply processes, working to move data to a target platform, handle the mapping, transforming, partitioning, and timing of the data movement to the target.

SQL Replication enables heterogeneous replications by working with the federation capabilities of the Information Integrator products and the replication administrative tools to include non-DB2 sources and targets. Several competitive relational databases can be either the source or the target for data movement.

This architecture is best suited to distribution topographies (one to many), data transformation scenarios, and heterogeneous data movement.

Q Replication

For low latency, DB2-to-DB2 data movement, WebSphere Information Integrator Replication for z/OS provides another option. This architecture uses WebSphere MQ Series as the transport for the data, and employs parallel apply agent technology to deliver low latency data movement. Stored procedure targets and event publishing offer new options for information integration. The event publishing capability, which can be purchased separately, formats XML documents from the captured data and publishes this information via WebSphere MQ Series. This can drive intra- or inter-business integration.

See *WebSphere Information Integrator Q Replication: Fast Track Implementation Scenarios*, SG24-6487.

4.2.7 Summary

In today's complicated, multifacited environments, DB2 UDB for z/OS is a cornerstone. DB2 provides deep synergy with the zSeries platform and z/OS. This delivers a data serving platform that is adaptable to change and provides the highest levels of availability.

While it is important to deliver quick-recovery capability, DB2 for z/OS and the z/OS platform are designed to go beyond this and deliver continuous availability. This difference in philosophy places DB2 and the z/OS platform at a higher level of capability. As explained in this chapter, deployment choices of data sharing, online utilities, and autonomic definitions allow DB2 to deliver higher availability while allowing you to make required system and application changes.

The improvement of existing features, and the delivery of new capabilities, continue to provide superior results on a long history of business resiliency. This delivers support for both legacy and new workloads, optimized data serving to the enterprise, and reduced administrative costs.

5

The on demand environment

Pressure to compete in the on demand era requires an enterprise to be flexible and adaptable to unpredictable market forces. Now more than ever the IT infrastructure must support the business by being flexible and responsive. At the same time, it must maintain the highest standards of security, integrity, and availability of information.

Through z/Architecture, supporting open standard languages and technologies, you can fully use the strengths of the platform and provide a future-proof IT environment on which to run your business. Specifically the capabilities for Java and Linux can provide a cost-effective deployment option for new Web-based applications. Both of these technologies are gaining rapid acceptance and increased focus in the development community, both for in-house applications and by independent software vendors (ISVs) for packaged solutions.

This chapter discusses z/Architecture adaptation through evolution and the latest technologies. It also discusses enhanced application support and z/OS.e for on demand workloads.

5.1 z/Architecture adaptation through evolution

Mainframe technology has consistently adapted over the years in response to demands, on the businesses it supports. As the market landscape has changed to embrace the Internet and the World Wide Web, IT systems and infrastructure have evolved to accommodate new business process models.

One of the greatest challenges that came with the e-business revolution was for corporations to realize that their business model had to change to support continuous availability from multiple channels. Competitive pressures led companies to open their business to global real-time communication with partners, suppliers, and a new Web channel customer base.

From e-business, we have progressed into the On Demand Business era where IT as a business enabler becomes increasingly more flexible, responsive, and adaptable to market forces. IT and its business customer have never needed to be more united if they are to meet the new challenges together. With the latest z/Architecture, IBM provides a perfect platform on which to consolidate and regroup in readiness for the battle ahead.

Throughout this evolutionary journey, IBM has not lost sight of the immense business value and investment wrapped up in legacy applications. For instance:

- Of the world's business data, 70% is processed by COBOL.
- Over 200 trillion lines of COBOL code are in existence.
- Daily 30 billion COBOL transactions are processed.

Many of these applications still provide efficient business function that IBM can help organizations exploit through a cost-effective legacy transformation program. Enterprise COBOL V3.3 enables developers to leverage more than 30 years' worth of applications in new endeavors. It delivers function to continue the integration of COBOL and Web-oriented business processes.

Note: With ever decreasing hardware, software, and running costs, the mainframe environment becomes an increasingly attractive proposition for deployment of new business applications and consolidation of mixed application workloads.

To learn how zSeries has helped clients meet the challenges of making their business an On Demand Business, visit:

`http://www.ibm.com/servers/eserver/zseries/testimonials`

This section discusses some of the key considerations for a modern competitive IT infrastructure: one that is ready to compete in an on demand world.

5.1.1 Revenue protection and growth

This is undoubtedly the top priority for any Chief Executive Officer. With trade borders and boundaries becoming increasingly virtual, an enterprise must combat both national and international competition. Companies can turn this situation to their advantage by exploiting this expanded marketplace, but only if their business processes and supporting IT systems are up to the job.

In the fight for turf in an on demand world, your competitors range from traditional established companies to Web-based newcomers to the market. To a customer who is searching the best deal, it often doesn't matter. Reliable service at the right price is the governing factor.

In a marketplace where loyalty can evaporate as a result of a Web search, it is important to protect your existing revenue base while striving for new business. With its support for the latest technologies and industry methodologies, z/Architecture offers a number of benefits that keep you competitive while helping to satisfy your customers and business users.

Customer experience

Customers and partners will quickly seek alternatives in the event of unreliable or slow service. Built on proven technology, the zSeries family of servers and supporting operating environment provide unrivalled reliability and availability with performance to match.

- Consistent service

  Through the enablement of additional processors or the introduction of additional hardware to a Parallel Sysplex configuration, your z/Architecture infrastructure can scale-up or scale-out to meet projected workload demand.

- Consistent response

  With its distinct dynamic workload management capabilities, a zSeries server can constantly monitor and react to resource utilization. Through the business language specification of service-level agreements (SLA), the system balances its resources to ensure important workloads are prioritized accordingly.

Through constant monitoring, as opposed to periodic snapshots, it is possible to respond immediately to unpredictable workload peaks. Resources are automatically diverted to important tasks to maintain system throughput according to business priorities. Refer to 3.9, "WorkLoad Manager, Intelligent Resource Director, and VIPA" on page 63, to see the functions capable of providing the constant management and adjustment required to optimize the workload throughput.

5.1.2 Return on investment

Your IT infrastructure is expensive and should be used to its fullest capabilities.

System resource optimization

Mainframe technology has evolved over time from a single-processor architecture to multiple physical multiprocessor servers collaborating and cooperating as part of a Parallel Sysplex configuration. Each central processor (CP) can be configured with a number of logical partitions (LPAR) to ensure optimal use of the processor time.

Workload Manager (WLM) works together with Processor Resource/System Manager (PR/SM) technology to ensure that the CP is kept busy crunching the numbers. Idle processor time is a waste of precious resource. Once it is gone, you can never get it back.

With the latest addition of Intelligent Resource Director (IRD), the automated workload management capabilities are extended to spread the workload across multiple CPs, even where they exist in separate zSeries Parallel Sysplex servers. Significant processing capabilities can now be made available to direct the workload where it receives the best service.

Note: By constantly working to reallocate unused resources to workloads that need those resources, you can dramatically increase system utilization, leading to improved return on investment (ROI). Compare this efficiency with that on open systems where processors are often under utilized because there is usually only one application per server.

The data sharing capabilities of DB2 Universal Database (UDB) for z/OS enable applications to run on any server within a complex. The unique multisystem control capabilities of the zSeries coupling facility (CF) ensure that contention over data is managed, and integrity across the enterprise is maintained.

Virtualization

IBM mainframe servers have been investing in virtualization technologies for over 35 years. As a result of continuous evolution and innovation, zSeries servers offer advanced virtualization techniques that enable you to run a few applications on a single server. A properly configured zSeries server can support up to tens, hundreds, or even thousands of applications.

Note: Potential benefits include the ability to help lower costs by deploying hundreds of applications on virtual servers, rather than paying for hundreds of actual servers, their software, maintenance, and the people to manage them.

The zSeries can also help to reduce complexity within the IT environment by virtualizing many components in a single managed environment. It can increase the flexibility of the infrastructure with new virtual servers being added in minutes. Virtual servers that are no longer required can be de-activated and their resources reclaimed and redeployed automatically.

Efficient staff utilization

IBM continues to drive its *autonomic computing initiative* with more self-configuration, self-healing capabilities in the latest z/OS and DB2 UDB for z/OS offerings. Through autonomic activity, the system can take trivial tasks away from your expensive, highly skilled system programmers and database administrators (DBAs). In doing so, it enables them to perform more significant work that brings business value to the enterprise. See also 8.4.1, “Autonomic computing” on page 195.

5.1.3 Business resiliency

In today’s on demand environment, downtime is both unwelcome and costly. If your applications are not consistently available, your business suffers. The damage can extend well beyond the financial realm into key areas of customer loyalty, market competitiveness, and regulatory compliance. High in the list of critical business requirements today is the need to keep applications up and running in the event of planned or unplanned disruptions to your systems.

Note: The high availability of z/OS is a key reason why so many customers rely on it for their most critical applications. To continue to improve this availability, z/OS provides automation capabilities in a new element, the *Managed System Infrastructure for Operations*. This element provides automation for single system and sysplex operations to help simplify operations and improve availability. MSYS for Operations plays an important role in outage avoidance.

IBM leadership in delivering world-class application availability, and overall business resiliency has been 40 years in the making. No other computing platform has integrated availability-enhancing features throughout its architecture as long as or as thoroughly as zSeries. And no one is better able to bring availability technologies to your IT infrastructure.

zSeries business resiliency solutions help you to:

- Enable superior levels of application availability thanks to innovative hardware features and the recovery capabilities of the z/OS operating system
- Reduce downtime from hours to seconds

- Dynamically and automatically reallocate resources to meet the demands placed on your applications
- Provision new virtual servers in minutes to handle shifting workloads, with Linux running under the z/VM operating system
- Keep system failures invisible to employees, partners, and customers, with dynamic disk-swapping capabilities that keep applications and data available

Downtime costs will likely continue to rise and become more visible as enterprises increase their reliance upon IT solutions. Your customers often access products via the Internet, so the amount of time and effort it takes for them to switch loyalties to another supplier is minimal, to say the least. If your site is not up or responsive when your customers want it, they are more likely to go somewhere else in the time it takes to click a mouse. The fact is that frequent system outages in today's highly competitive marketplace can negatively impact your business.

See Chapter 4, "Business resiliency" on page 85, for more details.

5.1.4 Security

A major concern in today's business climate is security. Studies show that businesses consider security to be the number one inhibitor to doing on demand or network computing. Competing technology infrastructures simply do not provide the comfort level concerning security that an enterprise needs to exploit expanded access to business-critical data and resources. You want the most robust security solutions you can find. You already know that z/Architecture is considered more secure than other platforms and can meet the challenge.

The zSeries offers a comprehensive integrated set of support elements that address a wide range of security issues. From user identification and authentication, access control and auditing to distributed directory, networking security and security administration, zSeries hardware provides integrity, process isolation and cryptographic capability to address your specific requirements and concerns. On top of this solid hardware foundation, zSeries operating systems implement the various elements of security with the security server and communication server components while supporting customization.

The zSeries has extended these robust security solutions to meet the needs of business today with an eye toward the future. zSeries servers have implemented leading-edge technologies such as high-performance cryptography, large-scale digital certificate authority and management, improved Secure Sockets Layer (SSL) performance and advanced resource access control function. With Intrusion Detection Services, the zSeries has enhanced its ability to help resist network-based attacks while embodying industry and international standards.

Note: Version 8 of DB2 UDB for z/OS delivers several new and enhanced security features in support of traditional and on demand applications.

The most important security improvements provided by DB2 Version 8 are:

- New built-in functions for data encryption and decryption
- New data encryption tool
- Multilevel security with row-level granularity
- Easier identification of system users
- Session variables
- Improved encrypted security in distributed computing environments

Multilevel security is a security policy that lets you classify data and users based on a system of hierarchical security levels that is combined with a system of non hierarchical security categories. A description of DB2 security improvements is provided in 3.4, “Security” on page 51.

Version 8 provides four new special registers that you can use in your applications to more easily identify system users. This can be useful whenever you have a large server, such as WebSphere Application Server, connected to DB2 where all the threads for a connection otherwise show a single identity.

5.2 Advanced technologies

To remain competitive, a modern enterprise must be ready to adapt and react if it is to benefit from the advances made in the industry. In support of a responsive business, the IT infrastructure must also be flexible and adaptable. Support for the latest technology initiatives and methodologies is a key strength of IBM technology.

Note: Through its open standard support and accommodation of the latest enabling software, such as the WebSphere product suite, the z/Architecture makes a perfect platform on which to base a future-proof IT infrastructure.

5.2.1 Open standards

Compliance with open standards is a key commitment from IBM. It plays a part in all of IBM’s technology platforms.

There are two types of open standard:

- **Official**: Where a recognized standards organization has endorsed it
- **De facto**: Meaning the marketplace has made it a standard

The HTTP protocol used by Web browsers is an official industry standard, sanctioned by the World Wide Web Consortium. Windows and Java are de facto industry standards, sanctioned by success in the marketplace.

Connecting across department and company borders requires one of two things. Either everybody uses the same technology, or all technologies can connect and integrate. The latter relies on the adoption and adherence to a recognized standard.

Since the world embraced Internet protocols, open standards have been gaining widespread adoption, from Java and Extensible Markup Language (XML) to Web services to Linux. IBM has played a large part in this industry-wide adoption, spending over USD $1 billion to support and grow Linux. IBM is also one of the largest commercial investors in Java and is involved in several open standards initiatives, more than any other company in the industry.

Note: The latest business-driven technology initiatives, such as Web Services, rely on adherence to open standards from participating enterprises. Non-compliance or inability to adapt may become a significant barrier to trade.

Linux

Linux on zSeries can help transform your business into a true on demand operating environment. Built on open standards, Linux running on zSeries servers gives you freedom of choice in middleware. zSeries solutions incorporate workload management technology that can enable your Linux applications to share computing resources and to allocate them dynamically as needed. And with advanced virtualization capabilities, zSeries servers running Linux can empower you to accomplish more with fewer resources.

Note: Linux on zSeries can help lower your total cost of ownership (TCO) in a variety of ways

- Hardware savings

  Linux running on zSeries can give you the ability to rein in server sprawl through consolidation. By exploiting the virtualization capabilities of zSeries, you can consolidate hundreds of individual Linux servers onto a single server and deploy these virtual servers in a matter of minutes. Fewer servers can also mean additional savings in physical space and heating, air conditioning, and electricity costs.

- Software savings

  Because Linux is an open operating system, running your applications on a zSeries server with Linux helps save you from locking yourself into costly

licenses for proprietary software. If you've already invested money in certain software to run your core business applications, you may be able to implement that software on Linux on zSeries. This allows you to run diverse applications and middleware alongside Linux in multiple dynamic partitions. Consider also that Linux application software and middleware are usually priced on a per-engine basis. Since a single zSeries server can run multiple Linux applications on a single engine, your licensing costs can be reduced by consolidating distributed servers onto a zSeries server.

- Manpower resource savings

  Consolidating multiple servers onto a single zSeries server running multiple virtual Linux servers can mean less labor is required for system management and maintenance. The centralized system management and autonomic computing capabilities of Linux on zSeries can also help reduce the errors and workload-balancing tasks that can consume countless IT staff hours.

- Business continuity

  Linux on zSeries can give you the ability to scale up and out, on demand, to meet spikes in server activity, helping to minimize costly transaction delays and potentially devastating system crashes. With a suite of built-in self-healing, self-managing and self-optimizing features, zSeries servers can rapidly respond to, or even anticipate, threats to system health, helping to prevent budget-sapping downtime.

Java

In the on demand era, business requirements change more frequently than ever. To stay current and competitive, a business must develop new strategic Web-based applications. Java adoption continues to accelerate as an open programming model in support of these requirements. These applications typically require more IT resources than traditional applications due to levels of abstraction, code generation, and reuse. Unfortunately, IT budgets are not keeping pace with these needs, forcing customers to seek more cost-effective and productive ways of deploying new Java technology-based applications.

Note: The zSeries platform together with the latest versions of DB2 UDB for z/OS provide extensive support for the Java language. Enhancements in security and performance, together with massively scalable virtual server technology, make this an ideal environment for Java applications, both in-house and that support ISVs.

Enhancements in DB2 UDB for z/OS Version 8 remove roadblocks to performance and DB2 family compatibility. They do so by providing support for a common client and standardizing database connection protocols based on the

Open Group Technical Standard DRDA Version 3. Three components comprise the common client:

- Common client for Open Database Connectivity (ODBC)
- Java common client for Structured Query Language for Java (SQLJ) and Java Database Connectivity (JDBC)
- Administrative client

Together, these components provide consistent access to servers across DB2 environments through a common run-time environment that has a single access path for all applications. As a result, new function and new applications can go into production more quickly because you can write an application once and use it in any of the DB2 common client environments.

Note: DB2 UDB for z/OS Version 8 implements the Universal Driver which, through the implementation of a single protocol and pure Java code base, greatly enhances portability of Java applications.

DB2 Version 8 also adds DRDA XA protocol support. This implements the Java 2 Platform, Enterprise Edition (J2EE), Java Transaction Service (JTS), and Java Transaction API (JTA) specifications. By doing so, Version 8 provides compliance with the XA Open Standards specification for distributed transactions.

The J2EE specification is a standard for developing, assembling, deploying, running, and managing multi-tier server-centric enterprise applications. The specification applies to all aspects of the architecture and development of large-scale enterprise applications. J2EE encourages the adoption of a multi-tier architecture and a strict separation between business logic and presentation.

SQLJ

SQLJ was developed by The SQLJ Group, a consortium comprised of database vendors and Sun Microsystems. SQLJ is a multipart static SQL model for using SQL with Java. It complements the earlier dynamic SQL JDBC model, with:

- Embedded SQL in Java, which enhances the approach of embedding SQL statements in programs
- SQL routines using Java, providing the ability to invoke Java methods from SQL code
- SQL types using Java, defining SQL extensions for using Java classes as data types

Note: The DB2 Universal JDBC Driver is a single driver that includes JDBC type 2 and JDBC type 4 behavior, as well as SQLJ support.

SQLJ is the recommended choice for application programmers who are concerned with security, performance, and simplicity.

- Security

  With customized DB2 SQLJ, the SQL is executed statically. Therefore the SQL statements are executed using the package owner's privileges. Any other users who run SQLJ applications must be granted the EXECUTE privilege on its packages. That is, users who are authorized to run programs, do not necessarily have the right to use SELECT, UPDATE, DELETE, or INSERT against the same tables or views that the program queries or is modifying unless an appropriate privilege is explicitly granted.

- Performance

  SQLJ allows SQL statements to be embedded in a Java program. Application code is precompiled allowing the SQL statements to be optimized when the packages are stored in the database. SQLJ applications executed statically can lead to better and more consistent performance than JDBC.

- Simplicity

  With its simplified and standardized syntax, SQLJ modules are typically more concise and easier to write than JDBC modules. This can help reduce development and maintenance costs. Existing embedded SQL applications written in other languages can be migrated easily to Java using SQLJ.

Java 2 Enterprise Edition

Through compliance with J2EE, z/Architecture represents an ideal platform for On Demand Business. IBM has demonstrated its support for this strategy through the porting of its award winning WebSphere product portfolio to run on zSeries.

WebSphere Application Server is a comprehensive, sophisticated, J2EE and Web services technology-based application platform. It is specifically designed to leverage the qualities of service (QoS) inherent in the z/OS operating system.

Note: Enhanced WebSphere run-time support and development tools enable COBOL programs to be dynamically loaded into WebSphere Application Server for z/OS regions. You can access them from the J2EE enterprise beans running in those regions.

Java stored procedures

Stored procedures, a feature of DB2 UDB for z/OS and almost all relational databases, provide an execution environment physically and logically collocated with the relational database. Stored procedures can provide major benefits in the areas of application performance, code re-use, security, and integrity. DB2 has

offered improving support for developing and operating stored procedures. Enhancements related to tooling, language support, and system environment have opened new possibilities for secure, highly portable applications in line with the strategy of today's organizations.

Java workload cost reduction with zAAP

The zSeries 890 (z890) and 990 (z990) enterprise servers support a new type of processor called the zSeries Application Assist Processor (zAAP). zAAP is an attractively priced specialized processing unit that provides an economical Java execution environment for clients who desire the traditional QoS and the integration advantages of the zSeries platform.

Why run Java on scores of seemingly inexpensive servers that often operate at no more than 25% of capacity and do not integrate with other business applications?

Note: It may be far more efficient, and less costly, to run Java workloads on your existing z/Architecture platform (z890 or z990) using the new zAAP support. This is one more reason for consolidating mixed application workloads, including new Web-based Java development, onto the mainframe.

zAAPs enable you to:

- Simplify and reduce server infrastructures by integrating Java Web applications next to important data for high performance, reliability, availability, and security.
- Maximize the value of your zSeries investment through increased system productivity by reducing the demands and capacity requirements on general purpose processors which may then be available for reallocation to other zSeries workloads.
- Lower the overall cost of computing for WebSphere Application Server and other Java technology-based applications; through hardware, software and maintenance savings.

Important: With zAAP, you can consolidate workloads and eliminate whole tiers of Web application servers. Most companies don't even realize how costly these servers can be in the long run.

XML

For the past few years, XML has increasingly become the de facto data format on the Internet, on corporate intranets, and for data exchange. As part of its commitment to the open industry standard, IBM has expanded and evolved its XML support to address advanced e-business applications.

DB2 UDB for z/OS Version 8 provides you with a set of brand new built-in functions to help with XML publishing.

Note: DB2 built-in functions can reduce your application development effort by providing a convenient high-performance option for the provision of lightweight XML applications.

IBM developed an XML Toolkit for z/OS that can help improve the speed of deployment and reduce the complexities of information exchange within business integration. The XML Toolkit for z/OS is designed to provide a valuable infrastructure component to assist you in creating, integrating, and maintaining your business-to-business (B2B) solutions. The XML Toolkit fully supports the euro. It is based on cross-platform, open source code that is designed to be compliant with industry standards and can help to lower the TCO by capitalizing on the existing zSeries infrastructure.

UNIX

In 1994, in MVS SP4.3, IBM introduced UNIX, calling it OpenEdition. In OS/390 R2, OpenEdition received full UNIX95 branding, and OS/390 UNIX became a UNIX system according to the Open Systems Foundation.

Note: The two major reasons for UNIX on z/Architecture are portability of skills and portability of code. You can use the same skills on different platforms, and you can move function from one platform to another. ROI in expensive manpower is maximized.

z/OS UNIX System Services is tightly integrated into the operating system. You win the best of both worlds: UNIX and z/OS. z/OS UNIX System Services (USS) are a key element of the IBM open and distributed computing strategy.

Unicode

The Unicode Standard is the universal character encoding standard used for representation of text for computer processing. The Unicode Standard provides the capacity to encode all of the characters used for the written languages of the world. The z/OS support for Unicode implements these Unicode Version 3.0 standards: Character code page, case conversion, and normalization. Unicode

support also takes advantage of native z/Architecture Unicode HW instructions for faster processing.

Unicode is the official way to implement ISO/IEC 10646 and is required by modern standards such as XML, Java, Lightweight Directory Access Protocol (LDAP), and CORBA 3.0. It is supported in many operating systems, all modern browsers, and many other products. The Unicode standard, and the availability of tools supporting it, are among the most significant recent global software technology trends.

Note: Incorporating Unicode into client/server or multitiered applications and Web sites offers significant cost savings over the use of legacy character sets. Unicode enables a single software product, or a single Web site, to be targeted across multiple platforms, languages, and countries without re-engineering.

DB2 UDB for z/OS is increasingly being used as a part of large client server systems. In these environments, character representations vary on clients and servers across many different platforms and across different geographies. One area where this sort of environment exists is in the data centers of multinational companies. Another example is e-commerce. In both of these examples, a geographically diverse group of users interact with a central server, storing and retrieving data.

The traditional way of encoding characters requires hundreds of different encoding systems, because no single encoding scheme is adequate for all the letters, punctuations, and technical symbols in common use. These encoding systems also conflict with one another, because two encoding schemes can use the same code points for different characters. To eliminate these problems, the support of the Unicode encoding scheme was introduced in DB2 Version 7.

The Unicode support that was introduced in DB2 Version 7 allows you to store data in Unicode. However, the ability to store data in Unicode is not enough to solve all code page-related problems.

With DB2 Version 8, considerable new functionality is added to improve the use of the Unicode encoding scheme. Those improvements include:

- Unicode parsing of SQL and utility statements
- Unicode catalog
- Possibility of using multiple CCSIDs per SQL statement
- ODBC Unicode support

5.2.2 Web services

Web services is a technology that allows applications to communicate with each other in a platform and programming language-independent manner. A Web service is a software interface that describes a collection of operations that can be accessed over the network through standardized XML messaging.

Web services use protocols based on the XML language to describe an operation to execute or data to exchange with another Web service. A group of Web services interacting together in this manner defines a particular Web service application in a service-oriented architecture (SOA).

Note: Opening your operational systems to customers, suppliers, and business partners requires a comprehensive security infrastructure and policy framework. Building on the proven security partnership of DB2 and z/OS helps your business enter this new environment with a high degree of confidence.

5.2.3 Service-oriented architecture

An SOA is a component model that inter-relates the different functional units of an application, called *services*, through well-defined interfaces and contracts between these services. By describing application interfaces in an XML-based language, called Web Services Definition Language (WSDL), services have moved to a more dynamic and flexible interface system.

You can leverage the power of open and industry standards and the flexibility of an SOA by using IBM WebSphere product offerings designed for the z/OS operating environment. With these products, you can build and deploy composite applications that are designed to make it easy to reuse your existing infrastructure and adapt to change.

Note: Enterprises are beginning to recognize the significant value invested in their legacy applications. Web services and the SOA provide a possibility for future exploitation of these existing assets through new channels.

5.3 Enhanced application support

IBM software that is available for the z/OS operating system helps businesses collaborate effectively, maintain and access information efficiently, and conduct transactions reliably, scalably, and with excellent availability. The IBM software for z/OS allows you to extend and enhance the Web application programming environment. It also provides the base on which higher level application services for On Demand Business are built.

Note: Assisted GUI-based development reduces reliance on legacy mainframe skills with the associated costs in addition to increasing productivity and time to market.

Figure 5-1 shows the tool suite available for the zSeries platform support application development from inception and beyond deployment. Through the advanced autonomic capabilities of z/Architecture, the deployed applications benefit from a constantly monitored and regulated execution environment.

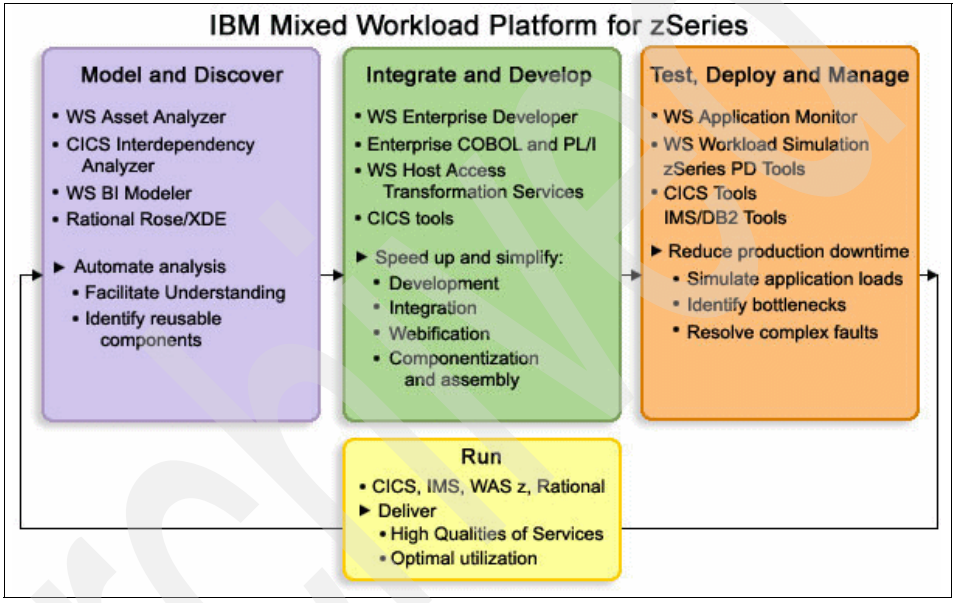

Figure 5-1 zSeries application toolset

5.3.1 WebSphere

zSeries provides a complete application development solution across the life cycle. This includes technologies that support the rapid discovery, development, deployment, and management of applications.

Development

This application life cycle is supported in a framework based on WebSphere Studio. It has a strong integration to compilers, performance tools, file management tools, and runtimes including CICS, IMS, and WebSphere.

WebSphere Studio provides an open, comprehensive development environment that tightly integrates with Ready for WebSphere Studio Partner plug-ins and

third-party Eclipse-based offerings. Founded on open technologies and built on Eclipse, WebSphere Studio provides a flexible, portal-like integration of multi-language, multi-platform and multi-device application development tools for building, testing, and deploying dynamic applications.

Eclipse technology provides a comprehensive set of functions useful to tool builders. Additional tools from IBM, IBM partners, Eclipse community members, and other tools providers extend the power and scope of the development environment.

Deployment

The application and transaction infrastructure delivers high-volume transaction processing for customers' important applications, through application environments with excellent security. It is built on the latest technologies such as Web services standards.

WebSphere Application Server

WebSphere Application Server for z/OS, V5, features the same deployment and administration capabilities as the WebSphere Application Server for multi-platforms. This allows you to reduce your training and education costs.

WebSphere Application Server for z/OS provides extensive integration with the operating system to help leverage the high reliability, strong security, and overall robust qualities of service of the zSeries hardware and z/OS.

Note: WebSphere is designed to take full advantage of the inherent strengths of the z/Architecture platform. It delivers exceptional self-optimizing and self-healing operational advantages to critical workloads for On Demand Business.

The self-optimizing capabilities of z/OS and WebSphere respond to periods of peak demand by managing the workloads to give high priority workloads processor resource. The self-healing capabilities of zSeries and z/OS include recovery from a processor failure and a system failure.

CICS Transaction Server

IBM CICS Transaction Server is an advanced transaction processing solution that provides powerful and flexible support for on demand applications.

CICS Transaction Server for z/OS V3.1 delivers the following capabilities:

- Increased ease of integration

  Web Services capabilities extend CICS applications to an SOA and support the industry-leading SSL security protocol.

- Enhanced application transformation

  This includes optimized CICS data exchange capabilities and the ability to leverage a single development tool (WebSphere Studio Enterprise Developer) for application transformation and integration.

- Improved performance and system management

  This provides simplified administration through the extension of the CICSPlex® System Manager Web User Interface, improved workload throughput, and enhanced C/C++ programs performance.

Note: CICS Transaction Server V3.1 provides new Web services capabilities to extend CICS applications to an SOA, enabling them to be exposed directly as Web Services.

IMS

IMS helps you with On Demand Business enablement, growth, availability, and systems management that the newer environments and cost measures require. The enhancements in IMS Database Manager (IMS DB) and IMS Transaction Manager (IMS TM) with Version 9 help you:

- Transform the way you do business with integrated information
- Build applications for On Demand Business that tolerate the rigors of doing business on the Internet
- Run a scalable, available, safe, and easily-managed environment
- Leverage your information to make more informed decisions

IMS Version 9 helps in addressing your On Demand Business needs through integration, openness, autonomic computing, and virtualization. It provides:

- Integration with other IBM products and with other products and platforms within the industry

  It supports the IBM commitment to open standards that benefit you. It enables the ability to develops supporting tools for application development and connectivity.

- Manageability through staging users to autonomic computing

  It eases use, eliminates and reduces outages, and minimized the education curve for users using IMS.

- Scalability through virtualization

  This assures flexibility for growth and expansion in a homogeneous environment, while using many heterogeneous data and application sources and new hardware and software facilities to optimize performance, capacity,

availability, and recovery. This includes new levels of enhanced availability for IMS High Availability Large Databases (HALDB), introduced in IMS Version 7, with fully integrated Online Reorganization support that provides concurrent online update and availability of data.

IMS Version 9 extends the reach of existing applications and data.

The integrated IMS Connect function, used with the WebSphere Development Tooling and the IMS Connector for Java, can significantly ease the development of solutions for On Demand Business that access IMS transactions. These solutions can be deployed in WebSphere Application Servers. They allow you to use Web applications, J2EE applications, or Web services to quickly transform static Web sites into sources of dynamic Web content, improving marketing effectiveness and customer service. They can also be used to transform C, COBOL, and MFS IMS transactions into Web services for SOA. This enables a quick response to new customer requirements, business operations, and competitive threads.

IMS SOAP Gateway support is an XML-based solution. It enables existing or new IMS applications to communicate outside of the IMS environment. It uses SOAP to provide and request services independently of platform, environment, application language, or programming model.

IMS Connector for Java is being enhanced to support using RACF key rings as SSL keystores. This support works in conjunction with the digital certificate support in RACF, allowing IMS Connector for Java to retrieve public and private keys from a RACF key ring. This eases security management for this environment. This support is being provided through the IMS Version 9 service process.

Business integration

The evolving technologies, business processes, applications, and people that make up your enterprise are all factors that influence the way you conduct business. Your ability to respond immediately to customer demands, support initiatives for On Demand Business, streamline mergers, and keep up with globalization requires a comprehensive business integration solution. It requires a solution that allows the disparate parts of your business to work together toward unified goals.

With a comprehensive enterprise integration solution, you have tools, technologies and best practices to create automated business processes that tie together these related entities and manage them as one. They enable you to respond proactively to changing business needs.

The comprehensive IBM WebSphere Business Integration portfolio provides five essential capabilities to help you respond rapidly to changing business requirements. It enables you to:

- Model and simulate business processes
- Integrate people, processes, and applications
- Connect customers and trading partners
- Monitor business processes end to end
- Manage and improve business effectiveness

5.3.2 DB2 development

The DB2 Development Center is an evolution of the Stored Procedure Builder. It is a rapid iterative development environment for building stored procedures, user-defined functions, structured data types, and more. This tool was rewritten from the ground up to allow for concurrent task execution, flexible docking, enhanced scalability, and higher productivity. WebSphere Studio Application Developer provides wizards to help you develop DB2 UDB applications.

Functionality includes:

- Support for the entire family of DB2 servers including DB2 UDB for z/OS, DB2 for IBM @server iSeries™, DB2 for UNIX, Linux, and Windows
- Enhanced z/OS support including specialized SQL IDs (package owner, build owner, and secondary SQL ID) and advanced build options
- Support for developing SQL and Java stored procedures, SQL scalar and table user defined functions, MQSeries®, OLE DB, and XML table functions, and structured data types for EJB methods and properties
- Support for viewing live database tables, views, triggers, stored procedures, and user-defined functions

The Development Center represents an evolution in DB2 application development tooling. It provides ease of use, scalability, and a rich set of functionality.

Note: With support for DB2 UDB for z/OS, your organization benefits from the development productivity gains of a Windows integrated development environment (IDE) and the runtime advantages of the mainframe environment.

5.3.3 IBM Rational

The IBM Rational® Software Development Platform (SDP) is a complete solution for developing software and software-based systems. It allows your organization to operate more efficiently and deliver innovative products and services that set you apart from your competitors.

The IBM Rational SDP gives you a common software development experience, spanning all the steps in a project, from inception through deployment. The result is a technology environment that maximizes your organization's collective strength across your business, operations, and development teams.

By defining, automating and integrating the strategic business process of software development in an open, proven and modular manner, a company can:

- Better respond to competitive and customer demands
- Maximize output
- Improve time-to-value
- Increase business agility

5.3.4 IBM Lotus

IBM Lotus Domino provides a multiplatform foundation for collaboration and On Demand Business. It drives solutions from corporate messaging to Web-based transactions and everything in between. This enterprise-class messaging and collaboration system is built to maximize human productivity by unleashing the experience and expertise of individuals, teams, and extended communities.

IBM Lotus Domino for z/OS helps zSeries clients build and deploy messaging and collaborative applications on the same system with enterprise applications and data. To reduce service delivery costs by eliminating server farms and consolidating to a single platform, IBM enterprise servers deliver the scalability to support thousands of users and applications with large storage requirements.

Note: The largest single system Domino deployment in production today (more than 35000 users on one hardware footprint) runs on z/OS.

Advancements in zSeries server technology and the z/OS operating system offer potential for even greater scalability. For unmatched reliability and availability for messaging and collaboration with virtually unlimited, nondisruptive growth, Domino for z/OS is the answer.

5.3.5 Microsoft .Net framework

For .NET Windows developers, IBM has recognized the strong demand for a tighter level of integration into that environment. IBM is strongly committed to the .NET platform, specifically Visual Studio .NET, where it has been part of the Microsoft Visual Studio Integrator Program (VSIP) since early 2002. IBM has delivered a native .NET managed provider and a tightly integrated set of tools for the Visual Studio 2002 and 2003 IDEs.

The IBM DB2 Development Add-In is a collection of features that integrates seamlessly into your Microsoft Visual Studio .NET development environment for working with DB2 servers and developing DB2 routines.

Note: Support for .Net and Visual Studio reflects IBM support for the latest development environments. GUI-based development aids the acceptance of mainframe technology by non-legacy staff, while retaining the significant benefits of zSeries runtime deployment.

5.3.6 Independent software vendors and packaged solutions

IBM is committed to supporting ISVs and making the zSeries platform an attractive platform for solution deployment. With the latest portability enhancements, modern language, and open systems support, there are even more reasons for choosing z/Architecture.

SAP

SAP is the recognized leader in providing collaborative business solutions for all types of industries and for every major market. SAP is the world's largest interenterprise software company and one of the world's largest independent software suppliers overall.

Note: The z/OS platform offers unmatched scalability, availability, performance, and manageability to reliably handle any information needed in an SAP R/3 solution using DB2 for z/OS as the database server.

Siebel

A comprehensive family of applications and services, Siebel applications for on demand business enable companies to draw from a consolidated source of customer information to service, as well as sell and market to customers across multiple channels. Using the software, companies are unifying and personalizing their interaction channels, offering Web self-service, and closely targeting their marketing campaigns, all in a quest for bottom-line leadership.

To successfully run Siebel applications for On Demand Business, you must have a powerful, scalable database that can provide data integrity, security, and ease of use.

PeopleSoft

PeopleSoft is a major provider of applications for on demand business that build a collaborative network of customers, suppliers, and employees. PeopleSoft helps organizations optimize every interaction with their customers, employees, and suppliers to create more loyal, collaborative, and profitable relationships.

There are several motivations for pursuing a PeopleSoft implementation based on zSeries Parallel Sysplex and DB2 data sharing. Many customers are deploying PeopleSoft applications in support of business-critical operations. Typical business drivers include a desire to run a single global PeopleSoft instance, customer and supplier access to Web-based PeopleSoft applications around the clock, and support of 24x7 manufacturing and distribution operations.

Those business drivers lead to the following IT requirements:

- Near-continuous system availability
- Central processor scalability through horizontal growth

5.3.7 DB2 utilities

Since the release of IBM DB2, Version 1, IBM has provided a set of core utilities to help DBAs manage their DB2 data. While these utilities have been continually enhanced in subsequent versions of DB2, DBAs have often opted to supplement them with stand-alone utilities and tools from ISVs.

The DB2 utilities match or exceed the best of the comparable products from the ISVs in terms of price/performance, functionality, ROI, and ease of maintenance.

Note: Consolidating on a single supplier toolset may not only reduce IT cost, but it can also simplify infrastructure management and administration overheads and simplify support or service with one vendor.

Reduced application development

Some of the latest database utilities offer functionality that historically would have been given to the IT department to solution through application development. You can save significant time, effort, and cost through the use of the utility functionality. It may not be viable to replace existing working applications. However with new requirements or heavy maintenance enhancements, you may want to see if there's a cheaper way to achieve the same results.

IBM utilities are optimized for the operating system and database. By deploying utility-based solutions, you automatically benefit from future performance enhancements.

- Replace insert programs with Online Load Resume
- Replace delete programs with Reorg Discard
- Transform data with Unload
- Transfer and transform with Cross Loader

Note: Utilities are not just for DBAs anymore. Programmers face tight schedules to deliver large and highly available application systems. Taking advantage of utility functions as part of the overall system design makes it easier to meet those challenges.

5.4 z/OS.e for workloads for On Demand Business

z/OS.e is a specially priced offering of z/OS. It provides select function for the zSeries 800 (z800) and z890 at an exceptional price.

z/OS.e is specifically configured for today's hot applications such as WebSphere, SAP, Lotus Domino, Siebel, Oracle, SAS, BEA WebLogic, or any Web-based data transaction processing application using Java, J2EE, or C/C++, giving these workloads a price/performance customers expect. Also, z/OS.e is complementary to those clients who want to run Linux applications. For example, z/OS.e can act as a database server for Linux application servers.

z/OS.e uses the same code base as z/OS and invokes an operating environment that is comparable to z/OS in service, management, reporting, and reliability. Also, z/OS.e invokes the same z800 and z890 hardware functionality that z/OS does. No new z/OS skills and service procedures are required for z/OS.e. With few exceptions, z/OS.e has the same features and functionality as z/OS. This offering does not support certain workloads and languages, such as traditional CICS, IMS, COBOL and FORTRAN.

Note: At a fraction of the cost of z/OS, z/OS.e makes it easier to run new workloads on the mainframe.

6

Readiness for new workloads

Businesses look at accommodating new workloads due to a variety of reasons, the primary reason being On Demand Business. Some of the reasons for introducing new workloads are:

- Development of new business transaction applications
- Enhancement of existing applications
- Purchase of new software packages
- Acquisition or merger of companies
- Consolidation
- Growth of customers
- Development or augmentation of Business Intelligence (BI) systems
- Requirement of storing new types of data such as video, audio, and XML

Of course, organizations do this to provide better customer service while reining in IT costs. This chapter first discusses the background information for situations for creating new workloads. Then it looks at the previously listed topics and shows how DB2 UDB for z/OS can be the ideal platform to handle these new workloads.

6.1 Background for new workload creation

Over the past decade, a number of companies have sprung up whose entire business model is based on providing customers, in a global marketplace, with the ability to conduct transactions through the Web. This made it necessary for new operational systems to be built embracing the use of graphical, easy-to-use front-end windows.

For existing traditional businesses, the business model also changed. Conducting business via the Web and remaining open for business 24 hours a day, seven days a week is now commonplace. The customer comes in through the Web and operates in a self-service fashion with mostly context-sensitive metadata to guide them. Often times, even with new Web applications, the data that is generated is the same as the data that was generated using mainframe green screens on dumb terminals before the advent of the Web application which wrote to a database on the mainframe.

Rewriting an entire online transaction processing (OLTP) application is expensive. But businesses want to develop newer enhancements to applications using a modern graphical user interface (GUI) to keep customers happy. Therefore, businesses also sometimes create a front-end, easy-to-use GUI, while passing user messages back to the same existing back-end programs and data repositories via software such as IBM MQSeries. As a result, software has to accommodate for such languages as XML, which is used for messaging.

Businesses also must handle more data resulting from acquisitions. The need for added BI resulting in the creation of data warehouses, operational data stores, and data marts has added a great amount of strain on the corporate IT department. Many businesses have tried other UNIX-based servers. They must maintain and pay software licenses for mainframe and non-mainframe systems and for hardware and storage.

There are business requirements of storing large objects (LOBs) of data such as pictures and big text documents as a record for business and legal purposes. Technology has not stood still for mainframe systems in the last 40 years. They have become adept at accommodating software that has been around for the last several decades such as IMS, DB2, and CICS. They have also become a host for running Web applications built in Java and accommodate the latest business requirements.

Several questions come to mind. For example, can this zSeries hardware with z/OS and DB2 running on it, rise up to all these challenges and be affordable at the same time? We discuss it in this chapter.

6.2 Development of new business transaction applications

When you develop new applications, chances are that you are building Web applications. You have already seen the remarkable features of z/OS with DB2 on it as a good foundation for your Web applications. This discussion carries you further into the other features that come from using this synergy of DB2 with z/OS for your new application, even if it is a Web application.

6.2.1 Prioritizing for new applications

New Web applications need database availability and performance. If the user does not receive a quick response, they will navigate away from the Web page. Workload Manager (WLM) on z/OS can give more priority to the user logging in from the Web and complete their transaction quickly. It can make a batch job wait for resources, while your Web customers or prospective customers use the system.

6.2.2 Security for new applications

Customers want to feel secure when they log into your new application. z/OS has code at the hardware level that can encrypt and decrypt data quickly. The cryptographic chips are duplexed for redundancy and offer a high rate of encrypt/decrypt operations.

zSeries 990 (z990) servers offer speed, with capabilities of greater than 11,000 Secure Sockets Layer (SSL) handshakes/second with z/OS Version 1.4. This was measured on a z990 server with 16 CPs and 6 Peripheral Component Interface Cryptographic Accelerator (PCICA) features. zSeries also has a PCI Cryptographic Coprocessor (PCICC) feature, which has a tamper proof design and supports symmetric encryption as well as public key encryption. This is widely used throughout the financial sector. The PCICC has earned the FIPS 140-1 Level 4 certification required by U.S. government agencies.

In addition to this, if you are using DB2 UDB for z/OS Version 8, row-level security is available. The user only sees the rows which are allowed based on his security label. See 3.4.3, “DB2 and multilevel security” on page 53, for more information.

6.2.3 Ability to accommodate growth

You may be building new applications, but worry about the ability of your hardware and software in handling this growth. Now business customers can have peace of mind in knowing that with Parallel Sysplex technology, more DB2 subsystems can be added and made to work in coordination with the existing DB2 subsystems in a data sharing fashion. This is *horizontal scalability.*

There is only one copy of your source data so there are no issues in dealing with replication. The subsystems which are data sharing members communicate with each other, taking the help of a coupling facility for managing the integrity of data. This internal communication is done in microseconds, so there is no performance overhead with the use of this technology.

Customers will also be pleased to learn that there is *vertical scalability.* Each z990 box can accommodate up to 32 central processors. Although adding further capacity can degrade performance as the processors compete for shared resources (the law of diminishing returns), there are many customers running with 16 processors and several other customers are moving beyond 16. The *Large Systems Performance Reference*, SC28-1187, is the source for IBM's assessment of relative processor capacity and can help in identifying the point beyond which horizontal scaling can come into play with the Parallel Sysplex.

Note: For more information, see also Chapter 3, "Why consolidate to S/390" in *Consolidating UNIX systems onto OS/390*, SG24-2090.

6.2.4 Rapid development of high performance applications

This section summarizes functions and tools for developing advanced applications.

Use of stored procedures on the server

You can choose from several development tools on the Windows platform. For example, you can use Development Center to develop stored procedures, save them on the mainframe server, and then test their execution on the mainframe. All your Web applications can share this stored procedure and execute it based on the business function needed. Then you receive predictable results which can be further based on your site standards.

Ability to use enhanced SQL features

The SQL that you can use in your programs has been greatly enriched. These features are described on the DB2 for z/OS Version 8 Web site at:

`http://www.ibm.com/software/data/db2/zos/db2zosv8.html`

The features are discussed here from a business perspective.

- Ability to fetch and insert multiple rows in one call

  Suppose you work for a mutual fund company and a customer wants to view her investment in her retirement plan over the past 90 days. In prior DB2 UDB for z/OS versions, you would declare a cursor and fetch one row at a time from the result set to the application program. With Version 8, you can use a multi-row fetch and fetch up to 32,767 rows into your application program at one time and display them quickly. Therefore, your customer does not have to wait as long.

  Similarly, you can insert multiple rows with one INSERT statement. There is a high reduction in CPU usage via the use of these features. This results in a savings in cost for both the company and the customer.

- Availability of a common table expression (CTE) that permits defining a result table with a table-identifier that can be specified as a table name in any FROM clause of the fullselect that follows

  This means, for example, that your development team does not need temporary tables created by a database administrator (DBA) to help in placing a query to help customers. Less time waiting for a DBA means more productivity from your development team with quicker development time and earlier implementation.

- Introduction of a powerful feature for recursive SQL

  CTEs provide the ability for you to perform recursive SQL. Recursive SQL enables you to concatenate result sets for each iteration in a result set table. It also lets you select from the consolidated result set table.

  This function, as delivered with DB2 Version 8, allows powerful SQL programming. It helps your development team to create and test user required data more quickly with lesser DBA support. And it enables DBAs to focus more on other day-to-day issues such as production issues, backup, and recovery.

- Capability of using a dynamic scrollable cursor which does not materialize the result table but allows you to scroll across the base table

  This reduces the use of memory while improving performance.

- Ability to store data in Unicode

  You do not have to worry about two users using two different Coded Character Set Identifiers (CCSID) to retrieve the same data. The DB2 catalog is now stored in Unicode. This helps in globalization of your business, because each character has exactly the same meaning across all languages on any coded character set page of Unicode.

- Usability of seven built in XML publishing functions in DB2 UDB for z/OS

  XML publishing is the creation of XML data from traditional relational databases. For the past few years, XML has increasingly become the de facto data format on the Internet, on corporate intranets, and for data exchange.

  IBM recognizes this and provides built-in functions with DB2. For example, it provides XML2CLOB, which is a cast function. It provides scalar functions such as XMLELEMENT, XMLATTRIBUTES, XMLFOREST, XMLCONCAT, and XMLNAMESPACES. And it provides XMLAGG, which is an aggregate function. Since these are built-in DB2 functions, they run within DB2 address spaces. This is unlike external user-defined functions (UDFs) that run in a Workload Manager managed address space outside of DB2.

  Furthermore, in previous versions of DB2, you had to create your own application that converted the DB2 data to the XML format or use a DB2 XML Extender. Now the use of these built-in functions into the DB2 engine improves performance.

- Use of INSERT within SELECT

  Surrogate keys are used in database design more so in recent years. They are artificially generated based on a business key. Therefore, they are not known at the time of insertion of a row. Sometimes you need the value so that you can use it in the insertion of a dependent table row. Now you can insert and select the value back in one statement, saving on network costs and improving performance.

- Introduction of *sequence* object

  This is a further improvement, in additional to the *identity column* that was introduced prior to DB2 Version 8. The value in the identity column generated by DB2 was at a table level but the same table could be created for several creators, owners, or qualifiers within a DB2 subsystem such as development, user acceptance, and so on.

  Sequences are at a subsystem level in a non-data sharing environment and across the entire data sharing group in a data sharing environment. They can reduce DBA time and enhance business productivity. The use of sequences also promotes portability from DB2 on other platforms and other relational databases.

There are many more functions such as star join processing enhancements, materialized query tables, and partitioning and indexing improvements. These are beneficial for the workloads resulting from using DB2 for z/OS in the development of BI applications. See also 6.7, “Development or augmentation of BI systems” on page 153.

Portability

Your development environment for new applications may be using data on DB2 UDB on Windows but the production implementation is on DB2 UDB for z/OS. In this case, there is less worry in regard to portability. This is because Version 8 of DB2 UDB for z/OS now supports longer names for tables, columns, views, stored procedures, and so on.

The name of tables and views for example has a length of 128 bytes and is common across DB2 for z/OS, Linux, UNIX, Windows and iSeries. Similarly the column name has a maximum length of 30 bytes, which is common across all members of the DB2 family. Sequences and the length of their names is 128 bytes, which is also common across all of these platforms. Common table expressions mentioned earlier in this chapter are also compatible across DB2 on all platforms.

Some differences still exist across the platforms for various limits such as for index names and schema names. These are all documented in Appendix A, “SQL limits” in *SQL Reference for Cross-Platform Development Version 2.* This IBM manual describes the rules and limits for preparing portable programs across all members of the DB2 family. You can download this book from:

`http://www.ibm.com/software/data/db2/zos/v8books.html`

Easier database administration

DBAs can now change more than 100 of DB2’s system configuration parameters while DB2 is active. Dynamic schema changes are allowed in several more situations as explained in 6.4.1, “Usability, availability, and scalability” on page 145. Backup and recovery of an entire subsystem are possible due to newer utilities as explained in the same section. The RUNSTATS utility now gathers distribution statistics for non indexed columns. Before DB2 Version 8, you had to use the DSTATS program for this.

6.3 Enhancement of existing applications

Much of the information relevant to new business application development is also pertinent to the enhancement of an existing application. If the DB2 UDB for z/OS is scalable and available for workloads that result from building new software applications, then it can accommodate the enhancement of existing applications.

While preparing your design specifications for enhancing your system, you can take advantage of more features being made available with each newer version of the z/OS operating system and DB2 UDB running on this operating system. Some of these features are mentioned earlier in 6.2, “Development of new business transaction applications” on page 137.

An additional point that is applicable for enhancement of your existing applications is the use of utilities instead of coding programs. This decreases development time and enhances productivity and performance. Let us examine this in the following sections.

6.3.1 Using utilities instead of developing programs

This section shows some examples of ready-made application solutions that are provided by DB2 utilities.

Using the LOAD utility instead of INSERT programs

The regular LOAD utility drains all access to a table space. Draining here means that DB2 waits for all applications to relinquish their claims. It blocks any new claims for the table space for INSERT, UPDATE, or DELETE while LOAD is running.

Development teams resorted to using INSERT programs in this situation. But INSERT programs cause contention with OLTP programs if they don't commit frequently. If you use LOAD RESUME YES SHRLEVEL CHANGE, then it operates like an INSERT program. It uses Data Manager INSERT instead of Data Manager LOAD. It also tries to maintain the cluster sequence which does not happen if regular LOAD utility is used.

IBM benchmarks have found that LOAD RESUME YES SHRLEVEL CHANGE is about 25% faster in elapsed time and uses 30% less CPU time than a comparable INSERT program. This was noted in the article "Programs vs. Utilities", by Jim Ruddy, in *DB2 Magazine* from Quarter 2, 2003.

Replacing DELETE programs with REORG DISCARD

Sometimes you may need either to delete a significant number of rows or archive and then delete a significant number of rows. In either situation, using SQL DELETEs of a large amount of data can be time consuming because they generate log records. This is also a performance inhibitor on the entire system because a "log latch" is taken when these log records are being written.

By using REORG TABLESPACE SHRLEVEL NONE or REFERENCE or CHANGE you can delete data, conditioning it with DISCARD FROM TABLE tablename WHEN condition. You can have DB2 either discard these rows or write them to a data set and generate a corresponding LOAD statement for loading them into an archive table. Specifying LOG NO on this REORG statement reduces logging and avoids system bottlenecks. The discarded rows are efficiently removed, and the table space is completely reorganized.

REORG SHRLEVEL NONE allows neither read nor write access to data. SHRLEVEL REFERENCE allows only read access. REORG SHRLEVEL CHANGE, also called online REORG, lets applications have read and write access to the data. At the same time it unloads, sorts, and reloads the data and then applies log records to catch up with the transaction activity.

Transformation of data with UNLOAD

You may want to change a data type from SMALLINT to INTEGER, strip away trailing blanks of a character column, or change the encoding scheme from EBCDIC to Unicode. In this case, you can use the UNLOAD utility instead of writing a program.

By using the UNLOAD utility, you can unload a SMALLINT column as an INTEGER, use STRIP to remove leading or trailing blanks, and use TRUNC to truncate a column to a specific length. You can convert character data to Unicode by specifying a particular CCSID. For quick analysis and testing, you can only unload a subset of rows by specifying a percentage in the SAMPLE parameter or specifying a number of rows using the LIMIT parameter. You can also use a WHEN condition to qualify your subset of rows.

Both the LOAD and UNLOAD utilities in DB2 UDB for z/OS Version 8 offer delimited data support when you add the keyword DELIMITED.

Using Cross Loader for transfer and transform

The DB2 Family Cross Loader is an option of the LOAD utility that is activated when you use the EXEC SQL statement above the LOAD statement in a load job. It combines the power of SQL functions, Distributed Relational Database Architecture (DRDA) communications, and the IBM LOAD utility.

To load data into a DB2 table from different sources such as Oracle, SQL Server, Informix, you can use an SQL statement in a LOAD job. You need DRDA communication and WebSphere Information Integrator to access non-DRDA compliant data or nonrelational data. By using these functions, the SQL statement in the LOAD utility job can access the source data and load it into the target DB2 table. You can also have the source and target tables as DB2.

You may have often faced the situation where you go through a painstaking process, such as to create your TIME dimension table with 50 years of data in a test subsystem, and have must use an SQL program to select and insert these rows into a production TIME dimension table. You can use the Cross Loader option to load the data.

You may also have needed to consolidate data from different data sources. In this case, you no longer have to unload the data from each source into flat files, move them separately onto the mainframe, place them into temporary tables, and then

use SQL. Now you can execute SQL statements directly on these sources and load data into your DB2 table on the mainframe. This saves a lot of time for your application development team, provides excellent performance, and enhances overall productivity. Recent enhancements also allow Cross Loader to directly load large objects (LOBs).

6.4 Purchasing new software packages

Over the past several years DB2, first on OS/390 and subsequently on z/OS, has been enhanced with features that are helpful in running packaged applications smoothly and efficiently. We review a few features that have been built into the DB2 mainframe database server over the past few versions. This gives you an idea of the cumulative improvements that are going into the database product, keeping in mind the software packages that you may buy.

For example, consider the DB2 for OS/390 Version 6. You could defer defining the underlying VSAM data sets for application tables that you were not going to use immediately. You could temporarily suspend logging while a backup of the logs and database was done and then resume logging.

Optimizer-related enhancements provided the ability to force the optimizer to choose the index even if the RUNSTATS utility indicated a smaller number of pages, which would have precluded the optimizer from using the index. This was the case to account for packaged application tables which are volatile and quickly grow in size. Packaged applications use a lot of dynamic SQL. The ability provided for larger environmental descriptor manager (EDM) pool space for the prepared-statement cache in which the dynamic SQL of your application is cached. This improved performance.

Then DB2 Version 7 for OS/390 and z/OS came along. It provided the ability to restart a DB2 data sharing member with minimal storage requirement in the event of a system failure. Enhanced warnings were created for long-running unit of recovery (UR) jobs. The number of log records that were written by the uncommitted UR jobs was taken into consideration for issuing warnings.

DB2 UDB for z/OS Version 8 has built upon this. It added features that make it an excellent choice as the database server for software packages that you buy for your organization. The following sections mention a few features of DB2 Version 8 and explain how they help packaged software applications.

6.4.1 Usability, availability, and scalability

Many features in DB2 UDB for z/OS Version 8 have enhanced usability, availability, and scalability. The following sections explain most important of these features.

DB2 universal driver for Java

Java Web applications need the services of a “driver” to access the database server. DB2 UDB for z/OS Version 8 provides a universal driver that uses the DRDA protocol to access data on any local or remote server. Prior to this version, there were different drivers for the DB2 runtime client, DB2 application development client, and for DB2 Connect. This affected portability. The new universal driver enhances usability and portability.

Unicode support

Architectural changes to DB2 Version 8 give you the ability to store data from around the world with each character having the same meaning across any language within the Unicode character set pages. The DB2 catalog stores information in Unicode. DB2 converts any SQL code to Unicode before parsing. This ensures proper results, meaning that one collating sequence is used, when dealing with global data. Tables belonging to multiple encoding schemes, such as EBCDIC and Unicode, can be joined within a single SQL statement. This promotes DB2 usability by packaged solutions.

Dynamic schema changes and the online schema evolution

Database availability is enhanced with the ability to perform schema changes while the database is online. Prior to Version 8, new tables could be added to a database and new columns could be added to an existing table at the end of all existing columns. However the change of the data type of a column was not possible without the dropping and recreation of the table.

DB2 UDB for z/OS Version 8 has new capabilities to handle changes dynamically. Some of these enhancements are:

- The ability to add a new partition to an existing partitioned table and to rebalance the data within this partitioned table
- Adding columns to indexes without having to drop and recreate the index
- Changing column data types and lengths in tables without dropping and recreating the tables
- Changing the cluster sequence of the table data
- The ability for most of DB2’s system configuration parameters (DSNZPARMS) to be changed while DB2 is active

- Automatic secondary space allocation

  Before Version 8, DBAs had to spend time manually resizing the amount of disk space suballocated for table spaces. This enhancement allows the DBA to instruct DB2 to manage all table secondary expansions autonomically.

64-bit architecture

Using DB2 UDB for z/OS Version 8, customers can now fully exploit the immense storage capabilities of the 64-bit memory model with no changes to their applications. This allows DB2 systems to cache significantly more data, allowing it to vertically scale. Since the length of many DB2 objects, such as table, view names, and column names, have increased, this additional storage could not have come at a better time. The customer can combine data sharing, which helps in scaling horizontally, and benefit from excellent availability, by being able to scale vertically. Customers can expect some increase in CPU requirements following the move to a 64-bit environment.

Backup and recovery

A quick backup and recovery implies more database availability. The tables of packaged software generally have application enforced Referential Integrity (RI) and rarely have database enforced RI. Therefore you cannot use the QUIESCE utility to establish a point where the set of table spaces related to the referential structures are in a consistent state.

To help the DBAs in such situations, DB2 Version 8 has introduced the BACKUP SYSTEM utility, which can backup an entire subsystem or data sharing group. Similarly you can recover the entire subsystem by using the RESTORE SYSTEM utility. The BACKUP SYSTEM utility when used along with FlashCopy makes the backup extremely fast and promotes database availability. You can learn about this entire backup and recovery process on the DB2 for z/OS Version 8 Web site at:

http://www.ibm.com/software/data/db2/zos/db2zosv8.html

6.4.2 Performance

Many features in DB2 UDB for z/OS Version 8 have enhanced performance. Some of these features are explained in the following sections.

SQL enhancements

These are mentioned from an application development point of view in 6.2.4, “Rapid development of high performance applications” on page 138. This section discusses those SQL improvements that are performance related and were not mentioned earlier.

- Volatile Table support

  Often in packaged applications, the size of the table can vary from zero to many rows at run time. By declaring the table as Volatile at create time, you force DB2 to use the index and save CPU time. Otherwise DB2 looks at the RUNSTATS statistics and uses a table space scan or a non-matching index scan to satisfy requests. It does this because it would be under the false impression that the table has few rows or none at all.

- More Stage 1 filtering of rows

  In DB2 Version 8, more predicates with mismatched data types and mismatched lengths can be evaluated by DB2 as Stage 1 predicates. Stage 1 predicates reduce processing costs because the Data Manager component itself can deal with the predicate. In previous versions, these were considered as Stage 2 predicates or residual predicates. They were passed on to the relational data system (RDS) component of DB2 for evaluation of the join predicate. This helps your Web applications coded in Java because Java does not have a fixed-length character string data type. Every string has a variable length.

- Multi-predicate access path enhancement

  The new access path in Version 8 can change the sequence of tables by participating in joins that involve multiple tables to take advantage of early filtering. This means fewer pages are accessed to execute the SQL statement resulting in better performance. Performance is further improved if you employ parallelism in such situations as multicolumn sort merge join. Since DB2 Version 8 uses more predicates than it did in Version 7 for mismatched data types, it can exploit this feature better.

- Other SQL performance improvements present in Version 8 such as Parallel Sort enhancement

  Other improvements include the usage of triggers, the ability to scan indexes backward, and the capability of using multiple DISTINCTs on different columns in the same SQL statement.

6.4.3 Tools and administration

You can use a variety of IBM tools, many with a GUI that can be used in working with DB2 family databases, for all your existing and packaged DB2 applications. The tools are varied such as BI tools, command line tools, development tools, general administration tools, information tools, monitoring tools, and set-up tools. Of these, you can use the General Administration GUI tool of DB2 Control Center to administer your DB2 database.

DB2 Control Center

You can use the DB2 Control Center to connect to the mainframe DB2 UDB for z/OS database server. It is an integrated part of DB2 UDB for Linux, UNIX, and Windows. The Control Center provides you a GUI to easily create tables within the schemas of your choice, views, aliases, triggers, schemas, indexes, table spaces and so on. Wizards are available that help you create databases, table spaces, tables, load tables and so on.

This product also ships with a set of stored procedures that you must install on each DB2 UDB for z/OS subsystem that you want to work with using the DB2 Control Center. These stored procedures help execute batch programs for database administration such as for object maintenance. The batch program DSNACCOB for example is used to run utilities, such as COPY, MODIFY, RECOVERY, REORG, RUNSTATS, and STATISTICS.

You can also use other stored procedures to delete a data set, check to see whether a data set exists, submit a JCL job, cancel a JCL job, and so on. All of these functions help to administer your DB2 database more easily. Other GUI tools, such as the DB2 Development Center, enable you to develop, for example, stored procedures of your own and run them on the mainframe.

Note: The bottom line is that GUI tools are available to help you perform a variety of tasks including administer your DB2 for UDB z/OS mainframe database server.

The features of DB2 that help enhance the BI capabilities for packaged applications, such as SAP Business Warehouse are covered in 6.7, “Development or augmentation of BI systems” on page 153.

6.5 Acquisition of new companies

As soon as a business acquisition occurs, there discussion about the business model of the merged companies. For example, you may have a company that sells sporting goods only through its shops in various locations in the United States. Customers can either come into the store and purchase what they desire or call in to place orders, but only during the normal business hours of 9:00 a.m. to 8:00 p.m. Now suppose that this company has purchased a Web-based sporting goods company for which there are no physical shop locations nor any sales people. Customers come in at any time to the Web store and place orders.

This means that the database availability requirement for online transactions has gone from 9:00 a.m. to 8:00 p.m. to 24 hours a day, 7 days a week. The database also needs to accept the initial consolidation of the data of both companies. In

addition, the database server needs to be scalable to accommodate the ongoing extra workload. The database server should be able to interact well with the Web application built possibly in Java. Tools should be available within the database software that allow for easy and excellent maintenance. All this should be done with low total cost of ownership (TCO) over the period of time in which the hardware and software supports you.

DB2 UDB for z/OS on zSeries servers can do all this. DB2 UDB for z/OS is reliable, available, and scalable. It has an excellent portfolio of utilities and tools. It provides all of this with high performance. And it safeguards your investment into the future by updating versions of the software, accounting for Unicode, 64-bit architecture and online schema evolution.

6.6 Consolidation

New workloads can result via new company acquisition, as explained in 6.5, “Acquisition of new companies” on page 148, and from server consolidation. The following sections explore both.

6.6.1 Consolidation via new business acquisition

For example, consider that you work for an insurance company, which has just acquired another company. With this acquisition, you need to make sure that online transactions have access to one consolidated set of products that are sold at a price based on one rules engine. You cannot continue to have two different systems that are operated by the company, where the same product is sold for two different prices, with such items as geographical location and coverage offered remaining the same. In addition, continuing to operate on two different systems duplicates the call centers where representatives help customers to secure policies and process claims.

First your online transaction system, the one you choose to keep for the present and future, has to be made capable of handling the functionality of the merged companies. Second, the data of both companies must be consolidated into one system. If a customer requests service, you need to pull the pertinent information from either company with relative ease. And third, the reporting and business analysis functions must also be consolidated.

6.6.2 Server consolidation

Consolidation is more than what is caused by a company acquisition. Consolidation can also be server consolidation. Server consolidation is also referred to as *IT optimization*. This is more than simply replacing smaller servers

with fewer, bigger and more powerful servers. It is about simplifying and optimizing your entire IT infrastructure including servers, databases, applications, networks and systems management processes. The goal is to reduce costs and complexities, while providing a more rational, stable foundation for growth and new solution deployment.

The major drivers for server consolidation are:

- Cost savings
- Manageability of staffing
- Security
- Reliability
- Process control

As smaller budgets are available to CIOs, server consolidation presents one of the biggest opportunities for realizing savings. Imagine that in addition to saving money, there is an increased opportunity of providing a more efficient and resilient on demand infrastructure. On demand has been mentioned several times in this book. This section relates it in the context of consolidation.

An On Demand Business is an enterprise whose business processes are integrated end-to-end across the company and with key partners, suppliers, and customers. It can respond with flexibility and speed to customer demand, market opportunity, or external threat. To realize an on demand operating environment, businesses need to examine:

- Virtualization

  This means to reduce the complexity of the management of hardware, software, and network while dynamically adjusting to increasing workloads.

- Automation

  This is the ability to reconfigure systems on the fly and recover from failing components. It is the ability to provide security and availability, among other things.

- Integration

  This is integrating people, processes, and information. Integrating people means allowing them to communicate and collaborate. Integrating processes means not allowing your business processes to be isolated islands, but to work together. The scope of integrating information is across and beyond the enterprise. It entails looking at any information that can provide you insight and looking at the data from different places as if it were connected and in one repository. WebSphere Information Integrator can help you achieve this.

How does server consolidation help you achieve an on demand environment, that is capable of better aligning technology with your business goals? Tactics for

server consolidation can be grouped into four categories. The best practices are to use the elements of all four of these approaches that best benefit your organization, in an ongoing effort.

- Centralization

  This is the placement of servers in fewer locations to reduce the number of physical sites. This saves on facility costs and in consolidating the operations of some data centers into others. At a minimum, the administration of the data centers can be centralized, even as the servers remain in physically different locations.

- Physical consolidation

  This involves replacing many small servers with fewer large servers. Traditional server farms employ discrete servers that can consume incremental expenses: hardware and maintenance, floor space, power and cooling, additional support staff, and per-server (engine) software fees. They can require miles of cables and can require days to deploy a new server. They rely on spares and reboots for high availability.

  Reducing the number of small servers means a simpler environment and can reduce your costs and improve your manageability. This can also reduce the personnel required to service multiple servers and possibly reduce dealing with multiple technologies.

- Application consolidation

  This refers to the shifting of operations of many applications onto a single server. This approach benefits companies that have many different applications running on different operating systems. An example is COBOL applications running on z/OS writing to DB2 UDB for z/OS, Java applications running on AIX and some .NET applications running on Windows, and some packaged applications running on Windows writing to an SQL server.

- Data consolidation

  By merging or migrating data from multiple locations into one single system, you consolidate your data. This means that if you consolidate your data onto DB2 UDB for z/OS, then you need only the administration personnel for this software. Furthermore, you do not have to deal with keeping multiple copies of data consistent. Also, your reporting and data warehousing can benefit from such consolidation. All this can mean reduced costs.

The question is, “Why is a zSeries system with z/OS and DB2UDB for z/OS the right choice for such a consolidation? The answer is in Chapter 2, “The evolution of DB2 for z/OS” on page 23, through Chapter 5, “The on demand environment” on page 111, as well as this chapter.

In addition to all this, let us tie the value of consolidation into creating an on demand operating environment:

- Virtualization is realized when you use such tools as Workload Manager (which adjusts to varying workloads), online schema evolution, dynamic virtual IP addressing (VIPA), and easier backup and recovery.
- System failures are kept invisible to users with dynamic disk-swapping capabilities. Such features as FlashCopy make it faster to make copies by providing more availability. Eliminating a single point of failure is the hallmark of zSeries and DB2 UDB for z/OS. Thus automation is realized.

If you do not integrate data, then in modern IT environments, often times you have to replicate data. This brings about its own share of problems in trying to keep the data synchronous. Data Integration or Federation has become possible with the introduction of IBM tools, such as WebSphere Information Integrator, which can reach any data within and beyond your enterprise, using DB2 technology among other things. Read more about WebSphere Information Integrator on the DB2 for z/OS Version 8 Web site at:

http://www.ibm.com/software/data/db2/zos/db2zosv8.html

In summary, server consolidation can:

- Reduce the number of your data centers
- Reduce the number of your mail servers
- Provide significant savings in capital, networking, and administrative costs
- Contribute to improved efficiency in server administration
- Provide increased ability to respond to new application needs

Consider the example of IBM. Since the early 1990s, it reduced its 155 data centers around the globe to 12 major centers worldwide by 2002. The IBM effort to consolidate its Global Notes Architecture in North America is another example of server consolidation. The Business Transformation and CIO Office of IBM built the business case by recommending all the four consolidation approaches discussed in this section. The approaches were centralization, physical consolidation, application consolidation, and data consolidation.

The most interesting is the consolidation at the Poughkeepsie site. Before this consolidation, IBM supported approximately 20000 users on RS/6000® SP nodes in the company's location in Southbury, Connecticut. It supported another 18000 users in Toronto, Canada. These users were consolidated into the Poughkeepsie data center and migrated to Domino running on zSeries servers. The Poughkeepsie site is scheduled to support approximately 100000 users with zSeries. This shows that IBM is consolidating to zSeries servers.

6.7 Development or augmentation of BI systems

Increasingly BI systems, which can be a central data warehouse (CDW), an operational data store (ODS), or application oriented data marts, are creating a great amount of strain on IT organizations. Managing the scope can be challenging in developing a CDW or data marts. There is complexity in knowing the proper relationships of existing data and portraying them logically in the form of an ER diagram. There is also complexity in integrating similar business data from one or more operational systems to provide a unified representation. In addition, there is the decision to find the right platform to physically implement the BI system.

Suppose you are considering to use DB2 UDB for z/OS with zSeries hardware as your choice. You need to see what improvements in the software and hardware can make this a worthy choice for accepting the increased workload of such a system.

Or you may already have your BI systems on DB2 UDB for z/OS and you wonder whether you should stay on such a platform. If you stay on this platform, you must determine how to reduce the job cycle of the growing BI system. This discussion is important for such customers. However, most of it is centered on which features are available on DB2 UDB for z/OS and which can handle such a new workload.

Another situation is to consolidate data of a BI system subsequent to an acquisition for example. Here are the important points to consider when dealing with such a new workload from a BI perspective, most of which are very large database (VLDB) considerations:

- Preliminary sizing estimate and capacity planning for growth
- Data movement into the BI system
- Availability
- Scalability
- Systems management
- Parallelism
- Query performance and throughput
- Database support skills

6.7.1 Preliminary sizing estimate and capacity planning for growth

You have to examine your source systems to see the volume of data that would be in the data warehouse, for example. You must also factor in the growth of operational data that would be piped into the data warehouse BI system over time. Based on the data storage requirements, you have to perform capacity planning for the hardware, software, and networks.

With all this in mind, DB2 UDB for z/OS has been architected to support 64-bit virtual storage. Large address spaces of up to 2^{64} bytes can now be defined for the database master address space (DBM1) and internal resource lock manager (IRLM) address spaces. The DBM1 address space uses memory for various operations in DB2. This includes buffer pools. Any time there is any update of data involved DB2 brings the data into these buffer pools. For that matter DB2 also reads data from buffer pools. Larger buffer pools enable I/O reduction for random access and enable larger page sizes, which benefit sequential access.

Each DB2 subsystem has its own IRLM. The IRLM controls access to the application data. It is used by DB2 as the lock manager to ensure the integrity of the data. More storage for IRLM enables locks to be at a smaller granularity, resulting in less lock contention. Now you can have from 1 to 4096 partitions of the table, with each partition on a separate table space that can grow as large as 4 GB.

6.7.2 Data movement into the BI system

Here are some key considerations for moving data to the BI system:

- Will the initial load times of very large volumes of data be fast enough?

  LOAD utilities can be quick in loading large volumes of data. With DB2 Version 8, the architectural limit of partitions was raised to 4096, and secondary indexes can be partitioned. This improves the parallelism, and the utility run time is dramatically reduced. Furthermore the loading process is simplified by the support of the delimiter format, allowing for input files coming directly from any other platform without intermediate preformatting.

  Benchmarks with DB2 Version 8 have shown up to 1.4 million rows (of variable size up to 4000 bytes) per second. Furthermore, two indexes were built for each row, and RUNSTATS was run concurrently.

- Will the periodic load and update of data in the data warehouse fit within my batch window?

  New SQL enhancements help you access and update the data faster.

- Can obsolete data be purged quickly?

  Such utilities as REORG DISCARD can help you accomplish this. You can delete a significant number of rows and archive them into a separate file if necessary and reorganize your database files in cluster key sequence without empty pages.

  You can also partition by one key, such as date, for a fast purge, and cluster by another key for processing with good performance.

- How much real-time updating can be done?

  You have already seen the ability to scale vertically or horizontally to have more processing power. In addition, storage improvements have included an increase in the size of buffer pools for updating data. The log files have increased from 31 in previous versions to 93 active logs in Version 8. And the 1000 archive logs in previous versions have increased to 10000 archive logs in Version 8.

6.7.3 Availability

As BI systems become more critical for businesses to compete against other businesses globally and to do well generally, the requirements for maintaining database availability across various time zones grows. You must consider data access to the BI system from Internet, intranet, and client/server environments.

Having Workload Manager when using DB2 UDB for z/OS allows you to use system resources for BI purposes. If more, higher priority online transactions need such resources, then they can use them without much waiting time.

6.7.4 Scalability

As your business users begin using the Data Warehouse, you will increasingly find yourself with more requests for newer reports and aggregate tables. This means more data volumes, additional users, and increased concurrency. If this happens within a budget cycle, then this increase should be met acceptably without an increase in capacity.

DB2 UDB for z/OS on zSeries servers can provide this scalability. It offers the possibility of vertically scaling up by adding more central processors or horizontally scaling by using the Parallel Sysplex data sharing environment.

6.7.5 Systems management

More of your users will begin to depend on your decision support systems for their decision making. As a result, it will be more important for you to provide the same systems management that you afforded to the day-to-day operational systems. If you have your OLTP data on DB2 UDB for z/OS on zSeries hardware and choose to have your BI on the same platform, then you can leverage your systems management for both purposes.

6.7.6 Parallelism

As you begin dealing with large tables in the data warehouse, you will want to know if you can:

- Perform parallel data loading into the data warehouse or data mart tables
- Conduct parallel reorganization or recovery processing
- Execute parallel query processing

Partitioning, allowing a table to be split up to 4096 parts, is available in DB2 UDB for z/OS to help you exploit parallelism. You can choose to partition by day, month, or geographical location, whichever makes analysis of your business quicker in a parallel way. Since each table partition is assigned a different table space, you can load each partition separately via LOAD utility jobs. You can choose not to log and take an image copy after the load jobs finish. This helps you in not having latch contention for the logs.

In DB2 UDB for z/OS, reorganization is at a table space level and you can reorganize different partitions of the table in different table spaces in parallel using the REORG TABLESPACE utility. Reorganization placed the data back in the cluster key sequence and reclaim space. You can also choose a compress option to compress the data when it is reloaded. Similarly you can recover table spaces in parallel using the RECOVER utility.

You should have table and index partitioning to enable good parallel query processing. You can do this with DB2 Version 8. We already looked at the ability to partition a table. As the table is being partitioned, the index can also be partitioned and parallelism can be applied to the usage of indexes.

You can have a partitioned partitioning index or a secondary index. A *secondary index* is any index that is not a partitioning index. A *partitioning index* has the same left-most columns in the same collating sequence as the columns which partition the table. In a partitioned index, each partition of the index corresponds to one data partition.

In DB2 Version 8, the data within a partition may be ordered differently than the partitioning key. Now an index can be created to match this ordering of underlying data which is called a *data partitioning secondary index* (DPSI). Therefore this index contains different columns than the partitioning columns. It must be a non-unique index. You can also choose to have a *nonpartitioned secondary index* (NPSI) which can be unique. In this index, you can choose to have columns different than the columns in the table partitioning key.

Query parallelism is possible with *merge join*. Merge join, also known as *sort merge join*, requires a predicate of the form table1.column = table2.column. This is called an *equality join predicate*. Merge join requires ordered input on the

joining columns, either through index access or by sorting. This sorting can be done in parallel and improve merge join performance.

Therefore, parallelism is primarily to reduce query elapsed time, which benefits your development team and your users.

6.7.7 Query performance and throughput

Businesses have built various decision support data repositories, such as data warehouses. They have also built front-end GUI interfaces using query and reporting tools such as Query Management Facility (QMF™). The GUI front end appeals to the business users as does the query builder facility on most of these tools. They help business users without much SQL experience to create a SQL query and obtain the data they need. In such a situation, more queries are placed against the data warehouse, scanning more amounts of data, with the users expecting quicker responses. In this business context, the following areas require extra importance:

- Consistent responses for short and medium queries
- Prioritization of critical queries, according to business priorities
- Prevention of large queries from monopolizing system resources while providing them with the ability to complete successfully
- Managing throughput in a mixed query environment

Query performance is improved via parallelism as mentioned in 6.7.6, “Parallelism” on page 156. It is also improved as a result of the compression of data. DB2’s compression implementation of Ziv-Lempel algorithms is hardware assisted. Storing compressed data in table spaces reduces storage cost. These savings are propagated on each backup copy that may have been taken for recovery purposes, including the relevant log records. See 3.3, “Compression” on page 48, for more information.

Query performance has also improved due to all the SQL enrichment that has happened to DB2 SQL with each newer version. Quite a few of these improvements are in 6.2.4, “Rapid development of high performance applications” on page 138, and in 6.4.2, “Performance” on page 146. Let us look at some more improvements from a BI perspective.

Materialized query tables

Any time you spend a lot of processing power and I/O on a query and you know that the end users want such a result set, then you can create a materialized query table (MQT) for them. If you already created a View for the users, then you can leverage the effort you put into the View by using the View to create the contents of your MQT.

The DB2 optimizer is aware of these MQTs. If you were to write queries against the base tables on which this MQT is based, then the optimizer sees if it can satisfy the request from the MQT. If it can, then it automatically rewrites the query to go to the MQT. You save CPU time and quicken query response time for end users by using MQTs.

In a real-world situation, you may use some views for reporting. The SQL of the view can be complex. The reason you have a view is that the user is interested in a snapshot of the information at each month end. They have no use in accumulating the information such as in a fact table. You don't need to request data from the view, requiring its result set to be generated several times during the month. Instead, you can have the MQT return the results more quickly, using fewer CPU resources.

Star join enhancements

BI systems can have a few data marts, which are built by organizations to quickly examine key business metrics. Such metrics include number of sales by the necessary context to understand the metric, such as by product, by time, and by geographical location. DB2 has created powerful enhancements related to the star join technique to facilitate end users in obtaining their results fast.

6.7.8 Database support skills

Fundamental database support skills are applicable across OLTP and data warehouse databases for a database management system such as DB2 UDB for z/OS. Organizations can leverage their OLTP DBAs for maintaining their data warehouses. However, as the data warehouse grows in size and complexity, businesses have to invest in training their DBAs on additional skills. These new areas include new skills, such as star and snowflake schema designs, conformed dimensions across datamarts, surrogate keys, and slowly changing dimensions.

6.7.9 Conclusions for BI systems on z/OS

When employing BI systems for z/OS, we know the following statements are true:

- Data Warehouse and data marts can coexist on DB2 UDB for z/OS, since a Workload Manager manages system resources effectively. And OLTP and Data Warehouse can coexist on the same system. Workload Manager can adjust the priorities for both throughout the day.
- Ease of moving data from the warehouse to the mart, with both being on the same platform, is a plus and an important consideration.
- Parallel processing of queries and DB2 utilities optimize performance.

- Centralizing processor resources allows for more effective sharing of the capacity, without creating isolated islands of capacity.
- Excellent new SQL enhancements reduce the elapsed time for end user queries.

6.8 Customer growth

Your company has an improved product offering and can launch effective marketing campaigns. All this has been made possible with BI systems that you have built. They have provided information about your customers and prospective customers and have matched the right products to your customers. This insight along with better customer service and globalization have caused a growth in your number of customers.

After reading this chapter, you can see that DB2 UDB for z/OS has unmatched features which can handle your added OLTP system workload and can satisfy the increased analytical demand of understanding the buying patterns of your customers. This can help your organization to retain your customers and continue to market to them the best products that they are likely to purchase.

6.9 Requirement for storing new data types

There is an increasing need to store large objects such as large documents, pictures, images of checks, as well as mixed media and their related data definition. DB2 lets you store all this as *large objects*.

The term large object refers to DB2 objects that you can use to store these large amounts of data. An LOB is a varying-length character string that can contain up to 2 GB minus 1 byte of data. There are three LOB data types.

- *Binary large object* (BLOB), which you use to store binary data such as pictures, voice, and mixed media
- *Character large object* (CLOB), which you use to store single byte character set (SBCS) or mixed character data, such as documents
- *Double-byte character large object* (DBCLOB), which you can use to store data that consists of only double byte character set (DBCS) data

DB2 supports XML. With XML, you have set up descriptive tags that identify the data universally. You can then apply a particular form to the data encapsulated in that form and transport your data anywhere. This support was first provided with the DB2 Extenders™. It is now being integrated in DB2 with a set of SQL built-in functions that allow applications to generate XML data from relational data.

These functions reduce application development efforts for generating XML data for data integration, information exchange, and Web services.

6.10 DB2 Query Management Facility

As information levels continue to grow, the need to leverage customer and operational data for business insights has become a strategic imperative. The DB2 QMF family helps you retrieve actionable intelligence from your data sources.

6.10.1 DB2 QMF components

The enhancements to DB2 QMF allow management teams to drive their businesses with insight and to respond with speed to changing market conditions. To accommodate today's diverse environments, DB2 QMF Enterprise Edition, a separately priced feature of DB2 for z/OS, consists of the following components. They allow users to easily apply their information strategies to single or multiple database and end-user platforms:

- DB2 QMF for TSO/CICS

  This component is for end users who require information access through TSO/CICS terminals. It includes powerful query building, automation, and data integration tools. Users from predominantly GUI workstation backgrounds can develop reports, as well as access and run saved QMF for TSO/CICS reports much faster and more easily.

- DB2 QMF High Performance Option

  This component enables scheduling, management, and control over end-user access to information by allowing DBAs to easily track and report end-user query activity. It includes monitoring resource usage, managing QMF objects, and compiling queries into COBOL programs for improved performance.

- DB2 QMF for Windows

  This component was designed for end users who need information access through Windows workstations. It allows application developers to reuse existing DB2 QMF for TSO/CICS objects (queries, forms, and procedures) or create new objects from the Windows environment, speeding the development process significantly. You can integrate query results with a variety of Windows desktop tools, such as spreadsheets, and DB2 for Windows or other desktop databases.

- DB2 QMF for WebSphere

  This is the DB2 QMF family's browser-based portal to business information on demand. DB2 QMF for WebSphere, which supports every IBM WebSphere platform, makes Web browsers zero-maintenance clients for the DB2 UDB family of products without plug-ins or special downloads. Users can connect easily via an intranet or the Internet to share many of the query, reporting, and OLAP functions found in DB2 QMF for Windows and DB2 QMF for TSO/CICS. It includes a completely redesigned and enhanced multidocument interface, allowing drag and drop grouping and formatting of query results.

- DB2 QMF Visionary Studio

  The new data visualization functionality of DB2 QMF allows developers to rapidly create visual information applications that offer rich interactive functionality specific to virtually any information need. In doing so, organizations can maximize the value of their data and take action based on facts and not guesswork.

DB2 QMF is also offered as a stand-alone product, called *DB2 QMF Distributed Edition for Multiplatforms*, which consists of DB2 QMF for Windows, DB2 QMF for WebSphere, and DB2 QMF Visionary Studio.

6.10.2 More on Visionary

In 6.7, "Development or augmentation of BI systems" on page 153, you learned how DB2 for z/OS can accept the BI workload gracefully with great query performance among other features. Nowhere else is this clearly visible than when the CEO wants a 360 degree view of the enterprise. This is an avenue where the Information Systems and Support group can show company executives their value to the enterprise.

DB2 QMF Visionary uses the data repositories that you have built in providing various operational metrics in the form of a dashboard. The Viewer can alter the thresholds in the visually rich, dynamic, and intuitive environment, which this tool fosters.

QMF Visionary consists of two major functions:

- QMF Visionary Studio

  This is a rapid application development authoring application, which you can use to develop and deploy your QMF Visionary world. A *Visionary world* is a graphic display of related data driven by dynamic SQL queries. A *world* consists of scenes, rather than traditional forms or reports. Therefore, as you navigate from scene to scene in a world, data displays can change according to context. Moreover, data is continually updated with what is current in the database.

- QMF Visionary WorldView

  This is a stand-alone end-user application for viewing and navigating the deployed Visionary worlds.

For a demonstration of this tool, visit:

http://www.ibm.com/software/data/qmf/demos.html

Let us see how this is useful in a business setting. Suppose you are the executive of an insurance company. You would want to see:

- The number of people who were reached this month by your marketing campaigns and the number that responded: You would want to know how many quotes were handed out and how many of these quotes became new business.
- The amount of the written premium this month: Compare that with last few months at the same time or with the same month last year.
- The number of claims that have come through for this month: Compare this number with the same time for the past few months.
- The resource utilization of your call center personnel and your claim service personnel

You can look at all this in a visually appealing environment. You can also view real-time graphs providing clear comparisons of business metrics over time. The business executives can change the time range if they want or vary thresholds when they look at their operational metrics using this Visionary tool. It allows insight into their business and pinpoints areas for improvement before these areas become full blown problems.

In addition to being useful for the commercial sector, DB2 QMF Visionary is equally important for various government organizations. Consider for example lawmakers in the legislature branch, at either the federal or state level, passing various laws such as for welfare reform. The government officials in the executive branch, such as the governor of a state, want to look on the effect of citizens at a county level.

Suppose the governor wants to know how many people are on welfare for various programs this month as of a given day as compared to last month at the same time. Then he simply needs to look at the QMF Visionary dashboard. In this case, the data repositories are built that can be queried from the Visionary world to provide this information to the governor, who can then obtain useful actionable information. The governor can also see the result of the laws that were passed during various months of the year for the citizens of the state.

Creation and deployment of a Visionary world are documented in the DB2 Information Management Software Information Center at:

http://publib.boulder.ibm.com/infocenter/dzichelp/index.jsp

Refer to the DB2 for z/OS Version 8 Web site, which covers more about DB2 QMF for Windows, DB2 QMF for TSO/CICS, DB2 QMF for WebSphere, DB2 QMF for High Performance Option, and DB2 QMF Visionary. You can find this site at:

http://www.ibm.com/software/data/db2/zos/db2zosv8.html

6.11 Summary

The focus of this chapter has been to identify the various situations in which your organization and your IT shop may have to tackle new workloads. For each situation, we described why DB2 UDB for z/OS on zSeries servers is the ideal platform for you to run your new workloads. A brief summary from a database perspective of this readiness for new workloads follows.

The business model of many enterprises has changed, requiring high database availability. DB2 benefits from the synergy with z/OS in using the Workload Manager to set priorities for tackling customer requests, while other batch jobs are completed in the background. z/OS has further synergy with storage systems, by which you can submit multiple I/O requests to the same volume from the same host, reducing contention. Your ability to back up and restore your database faster has also improved database availability. New software, such as Java, has come into play, giving businesses the ability to easily create GUI applications.

z/OS has evolved to let you handle Java workloads, using a ZAAP, which works with the CP in executing your Java workloads under control of the Java virtual machine (JVM). At the same time, businesses want to leverage existing investment made into existing systems by building composite applications whereby the new applications invoke existing code.

IBM recognizes the investment made by businesses in building applications on mainframe software, such as IMS and CICS, in the last 40 years. Therefore IBM makes sure that any advancement does not affect the existing software. The mainframe continues to support the historic code that still provides valuable capabilities for running the business.

The database extensibility of DB2 UDB for z/OS can handle newer data types such as big text documents, pictures, and audio. It offers newer performance features such as SQL advancements and parallelism. And it offers new availability and skill reducing options such as online schema evolution and easier

backup and recovery. These features help to improve employee productivity while helping you to meet the SLA with your customers.

Over the years, businesses have heavily invested in decision support systems. These systems have tried to understand more about customers and products, which has resulted in a huge development of BI systems. BI systems have different requirements for dealing with large sets of data compared with the traditional need of handling smaller read, insert, update, and delete transactions. DB2 UDB for z/OS is dealing well with both types of workloads. This is consistent with IBM Mainframe Charter, which calls for continuous innovation of mainframe zSeries to handle a wide range of workloads. Better portability of applications enables you, for example, to create a quick proof of concept using DB2 on Windows. It also enables you to leverage your work by porting the application to the mainframe and running it with fewer changes than before.

As more data comes in from new business acquisitions and server consolidation, horizontal scalability is available via sysplex and vertical scalability by adding processors to the same machine. This can help you create an on demand operating environment. Improved support is available for packaged applications, and a variety of tools, such as Development Center, help customers make productivity gains. You have seen improvements in getting the data into the database and have the ability to examine the data with clarity using QMF.

In Chapter 1, “Evolution of the mainframe” on page 1, we explain how IBM has lowered the cost of entry and brought the advanced mainframe technology to the mid-sized enterprise with its zSeries 890 (z890) mainframe offering. This, along with flexible pricing options that allow for On/Off Capacity on Demand, provide a great incentive to look at your mainframes as your partner in dealing with large and varied workloads as your business continues to grow. The mainframe continues to evolve to support you as you make your business agile and flexible to meet the demands of the customers now and for the years to come.

7

Evaluating the cost of your solution

IBM has added many enhancements in DB2 UDB for z/OS Version 8. Together with enhancements to the zSeries supporting platform, there are now considerable cost-based values for an enterprise that is deploying or considering this complementary partnership of technologies.

An enterprise IT infrastructure must maximize return on investment (ROI) if it is to continue to receive the support of the organization and be accepted as an effective differentiator in today's competitive marketplace. The dynamic transaction patterns for e-business and timely responsiveness required of an On Demand Business operating environment make the right choice of technology more important than ever before.

This chapter discusses the total cost of ownership (TCO) with zSeries and DB2 UDB for z/OS.

7.1 Total cost of ownership

Many and varied approaches exist for articulating the cost of IT and, in particular, the data center-related infrastructure components. Initial hardware purchase price and software license fees are only two elements in a complex cost structure that underpins the ROI in technology.

7.1.1 What is TCO?

The concept of TCO was developed by Gartner Group in the late 1980s and initially related to personnel computer deployment costs. Later, it was extended to cover other technology-related costs such as networks and elements of the data center. Over time, several tools and consultancy groups have materialized to support the modeling and assessment of TCO as it relates to IT.

Note: TCOnow! is an assessment tool from CIOview Corporation with which you can estimate the potential TCO savings for selected IBM Solutions compared with competitive alternatives. For details about TCOnow!, go to:

http://www.ibm.com/servers/solutions/serverconsolidation/tco/

The Gartner model categorizes costs as either direct or indirect.

Direct costs

These are visible IT costs that support related investments and expenses. The costs are related to:

- Hardware and software

  This covers the initial purchase or lease costs of hardware and the purchase or license costs of software. To these costs are added the related expenses for such items as:

  - Storage and network devices
  - Maintenance contracts and related expenses
  - Consumables

- Operations

  This covers employment and related office costs for support personnel including:

  - Systems programmers
  - Database administrators
  - Network specialists

- Security specialists
- Help desk support staff

- Administration

  This includes elements of costs for human resources, finance and other departments involved with the use of IT. Training costs also fall into this area.

Indirect costs

Indirect costs are intangible costs that are spread across the complete organization including business related IT expenses. They include:

- Unofficial support personnel

  As home PC ownership and utilization has evolved, most organizations have realized an increase in user IT awareness and computer literacy. With this growth in confidence comes the introduction of the user department *expert*. These individuals are often approached to investigate and help with departmental problems in favor of a perceived slow central support group. However the expenses relating to this person's time and any consequential costs are rarely recorded or accounted for.

- System downtime

  This refers to incidents relating to the IT infrastructure which cause a user to be disrupted and their productivity impacted, causing a hidden cost to the business.

7.1.2 Mainframe differentiators

Some of the more obvious mainframe capabilities that differentiate the mainframe from the competition and deliver TCO benefits are:

- The right confidence level for the enterprise in running important applications
- A basic capability to divide a single server into multiple partitions or virtual servers, each one running multiple applications simultaneously and securely
- The ability to self-manage the allocation of system resources based on business priorities among multiple workloads and virtual servers, allowing the mainframe to sustain high utilization rates (over 90%) while meeting IT service level objectives

  This compares to average utilization rates of less than 30% for UNIX and Microsoft Windows NT® environments.

- Industry leadership availability and scalability characteristics that allow for consolidation of large applications
- A high degree of systems management software and automation, coupled with mature IT management process providing higher people productivity

Note: The result is more efficient utilization of server capacity, higher people productivity, and lower environmental costs. This adds up to TCO advantages with reduced IT complexity.

The sections that follow explore several the cost considerations, both visible and most often overlooked, that contribute toward a longer term cost of ownership.

7.2 zSeries and DB2 UDB for z/OS

As stated previously, the unique characteristics of the IBM z/Architecture offerings make it a compelling proposition when considered in relation to TCO as opposed to simple acquisition-related expenses.

7.2.1 Mainframe Charter

IBM introduced Mainframe Charter to provide a framework for planned future investment and to highlight specific ways in which IBM intends to deliver ongoing value to zSeries clients. In August, 2003, IBM announced a comprehensive set of pricing initiatives designed to lower costs and encourage clients to accelerate their move to “on demand.”

A commitment to delivering value has many components. In essence IBM is delivering a continuous stream of technology and offerings that help lower the cost of mainframe computing on zSeries. IBM is delivering more performance for a lower unit price.

For further information about Mainframe Charter, go to:

http://www.ibm.com/servers/eserver/zseries/announce/charter/

Provide value

The charter outlines nine principles which fall into three action groupings.

- Innovation: Continuing innovation to maximize business value
- Value: Delivering value for an on demand world
- Community: Turning the On Demand Business vision into reality

Under the provision of a value group, IBM intends to:

- Enhance the value proposition and lower the cost of computing of zSeries solutions in a way that is compelling, clear, and consistent
- Extend the on demand characteristics of zSeries servers, highlighting its strengths as an environment for usage-based computing

- Increase the ability to account for allocation and use of zSeries resources in an on demand environment

Continuous advances in z/Architecture technology, married to the latest features of DB2 UDB for z/OS, make this platform ideally placed to run traditional and modern applications for On Demand Business in a mixed workload environment.

Note: Under a scenario running traditional applications and WebSphere Application Server applications for On Demand Business, IBM envisages a potential ongoing cost saving of up to 40% over three years.

Other scenarios suggest similarly impressive savings potential by upgrading to the latest servers, product offerings and pricing models.

Clients who are contemplating an Enterprise Application Solution (EAS) that uses Linux and DB2 should investigate projected cost benefits when evaluating their infrastructure options.

Persistent innovation

Under this commitment, IBM has a principle to improve the autonomic and self-management capabilities of the zSeries. We have already seen several enhancements in the z/Architecture platform that drive it in the direction of lower maintenance and support costs. This continues with the latest zSeries 990 (z990) hardware and the latest z/OS and DB2 UDB for z/OS Version 8 products, all of which incorporate unique autonomic capabilities.

Attention: The cost implications of a true autonomic platform are substantial with analyst reports quoting human error as the main cause of system downtime and instability. This can have serious consequences to the business, not all of which are immediately obvious as explained later.

A system that can manage itself also has the potential to return major benefits. Self optimization features and simplified administration tools help to keep the system running for more time while providing efficient utilization of resources.

Foster a community

Under the third charter category, IBM plans to revitalize the mainframe community to the benefit of both existing users and potential new clients. The need for expensive highly trained mainframe support staff has often been seen as a detractor. IBM has countered this with autonomic capabilities, as discussed earlier, together with a commitment to ensure the availability of design, development, and support skills for zSeries.

In addition to training its own staff to assist clients, IBM is working to increase the available zSeries skills base within the industry. The IBM Scholars zSeries Program offers colleges and universities educational resources to enhance the zSeries skilled resource base. It also helps students develop practical skills that enable them to find good jobs quickly upon graduation.

7.2.2 Hardware

You must consider several elements when evaluating mainframe costs. These have relevance to extending the current use, upgrading to the latest offerings, and considering the suitability of a mainframe-based solution.

Server

Currently four zSeries models are available that support a range of power and feature specifications.

- zSeries 800 (z800) and zSeries 890 (z890): Mainframes for the mid-sized enterprise

  zSeries functionality is offered at a lower entry capacity level than previously available on IBM z/Architecture hardware, with new granular growth options.

- z990: The server for the on demand world

  This server is optimized to help meet the requirements of today's on demand solutions. It can scale up to three times the capacity of the zSeries 900 (z900) Model 216. The ability to scale up physically and scale out virtually helps simplify IT infrastructure and integrate multiple workloads.

- z900: The platform for On Demand Business

  Balanced system design enables optimal use of system resources in dynamic On Demand Business environments.

- z/OS.e

  z/OS.e is a specially priced offering of z/OS, providing select z/OS function for the z890 and z800. This product offering is intended to help exploit the fast growing world of On Demand Business using mainframe technology and the attractively priced z/OS.e to deploy new application workloads. It will *not* support certain traditional workloads.

  Clients for whom the limitations of z/OS.e are not appropriate may instead want exploit the latest New Application License Charge (NALC) to achieve cost savings. The NALC provides clients with an affordable price for deploying qualifying new workloads under z/OS. This includes such application workloads as WebSphere, Domino, Siebel, and SAP.

Note: This combination of hardware and operating system gives the ability to run applications written exclusively in Java, Java 2 Platform, Enterprise Edition (J2EE), and C/C++ at a lower cost. The platform supports the zSeries Application Assist Processor (zAAP) for greater Java workload cost savings, as described later.

Memory

The cost of memory continues to drop while hardware memory capacity continually increases with each new server model. Today's application technologies for On Demand Business, including J2EE, Linux and 64-bit exploitation, take advantage of the latest data-in-memory application architectures. Large memory configurations can help improve performance of solutions based on these application technologies.

Note: DB2 UDB for OS/390 Version 6 began to use the larger, 64-bit real storage allowing to move buffer pools in data spaces. DB2 UDB for z/OS Version 8 took the next step, starting to exploit the 64-bit virtual architecture of the zSeries platform to enhance performance and free lower addressable memory to 31-bit applications.

Clients running varied application workload profiles may see additional benefit from larger memory configurations. As applications enable the higher memory area, a potential need to scale the enterprise infrastructure across an additional server may be removed or deferred.

Disk storage

Information is an increasingly valuable and costly organizational asset. Organizations seek to minimize risk, reduce costs, and increase flexibility by aligning IT investments according to information value. IBM provides clients with the industry's most advanced and complete portfolio of information management solutions. At the heart of these solutions is a simplified and resilient storage environment that helps efficiently manage information over its life cycle.

Through DB2 host data compression, savings can be made on storage costs. The consolidation of server applications onto a single server or server cluster can also realize savings through the de-duplication of data. Archiving and compressing less frequently used data to less expensive devices can also be beneficial. Tools are available that can be help in managing archival at row level.

Important: The trend for decentralized IT infrastructure leads to a proliferation of smaller server deployment and in an uncontrollable spread of local data copies. This increases the storage overhead across the wider organization and introduces unquantified security risk and unmanageable access control.

DB2 UDB for z/OS exploits many of the features of IBM TotalStorage Enterprise Storage Server (ESS) and zSeries technology to deliver added benefit to the business in terms of availability and performance. In particular, it supports these features:

- FlashCopy Version 2 for speedy database backup
- Parallel Access Volume and Multiple Allegiance for efficient concurrent access to shared data
- Advanced hardware compression for lower space utilization and increased data throughput
- On z890 and z990 servers, the use of latest device connectivity technology (FICON Express2) can deliver significant improvements in data transfer rates for storage area networks (SAN)

The new DB2 system level backup and restore utilities simplify and speed up these procedures while significantly reducing recovery times. See 3.10, “Disk storage” on page 68, for more information.

Maintenance

A simplified IT infrastructure means fewer components can go wrong. Combined with the proven reliability of mainframe technology, this makes for a robust environment on which to base an enterprise computing strategy.

Since the majority of system downtime is attributed to human error, the autonomic capabilities of zSeries and DB2 for self-configuration and self-healing bring further benefits.

Note: Capacity backup can also help reduce maintenance downtime as described later.

7.2.3 Mainframe pricing

IBM continues to drive down the cost of the mainframe platform. With its latest pricing initiatives, z/Architecture becomes a viable and cost-effective, long-term option for On Demand Business.

Capacity on Demand

IBM has introduced several initiatives to support its on demand ethos.

- IBM Capacity Upgrade on Demand (CUoD)

  This cornerstone of the IBM Capacity on Demand offering allows for the nondisruptive addition of one or more central processors (CPs), internal coupling facilities, and integrated facilities for Linux for predictable, permanent growth. CUoD can quickly add processors up to the maximum number of available inactive engines. This provides you with the capacity for much needed dynamic growth in an unpredictable business world.

- IBM On/Off Capacity on Demand (On/Off CoD)

  On/Off Capacity on Demand enables additional processor capacity to be turned on or off as necessary to satisfy business needs. zSeries servers can be provisioned with a number of additional processors configured as unavailable. As needed, these processors can be quickly brought online and are billed according to utilization.

  On/Off CoD is extended in z990 to include Integrated Facility for Linux (IFL). The concept of processors being available as needed has been extended to include the installation of additional dedicated Linux processors. IFL engines can be brought online as needed to satisfy demand peaks and are billed accordingly.

- Capacity BackUp (CBU)

  CBU for zSeries processors can provide reserved emergency backup processing capacity for situations in which you lose capacity in another part of your establishment. This helps you to recover by adding reserved capacity on a designated zSeries system. A CBU system normally operates with a “base” processing configuration and a pre-configured number of additional processing units reserved for activation in case of an emergency. The activation is done electronically. It quickly eliminates time that is otherwise spent waiting for a technician to arrive.

Workload License Charges

Workload License Charges (WLC) is a monthly license pricing metric designed to support today’s On Demand Business requirements. WLC enables customers to:

- Grow hardware capacity without necessarily increasing software charges
- Pay for key software, such as DB2, CICS, IMS, MQSeries, and z/OS with logical partition (LPAR)-level granularity
- Experience a low cost of incremental growth
- Manage software cost by managing workload utilization

WLC is subdivided into variable WLC (VWLC) and flat rate charges. VWLC apply to such products as z/OS, DB2, IMS, CICS, MQSeries, System Automation, NetView®, and Domino. Flat WLC apply to such legacy products as less current compilers and older MVS, VM, and VSE utilities.

Customers may choose to implement WLC in one of two ways:

- Full-capacity WLC

  Charges are based on the full zSeries server capacity where each VWLC product executes.

- Subcapacity WLC

  Charges are based on the utilization of the LPAR or LPARs where a VWLC product executes. A list of subcapacity eligible software products is available from the pricing Web site listed in the following Additional Information section.

Entry Workload License Charges

In support of an on demand operating environment, IBM offers Entry Workload License Charges (EWLC) for z800 and z890 clients. EWLC enables qualifying z800 and z890 clients to pay for subcapacity eligible IBM software based on usage of the LPARs where that product executes. This subcapacity pricing provides the potential to lower software charges on a stand-alone z800 or z890.

The zSeries Entry License Charge (zELC) is a monthly license charge pricing metric. It is designed to support the z800 server. IBM introduced zELC to deliver appropriate software price performance for a variety of z800 client needs.

Additional information

Software pricing in regard to workload planning is a complex operation. For further information, see:

`http://www.ibm.com/servers/eserver/zseries/swprice/`

Middleware

Enterprise application integration (EAI) software, such as IBM WebSphere MQ, provides enabling services that allow multiple processes running on one or more servers to interoperate across the enterprise. The more diverse and heterogeneous the infrastructure is, the more components and platform or package vendors are involved.

Consolidating your server base brings logistical benefits in middleware license management. It may also deliver significant cost savings.

Support software

The more complex your infrastructure is, the more support software and staff are required to administer and maintain it. In addition, the more heterogeneous your IT components are, the harder it is to standardize vendor offerings and impose a cost effective company software preference.

Simplifying your infrastructure can only bring benefit to the organization through rationalized processes and centralized management. IBM provides support software to help optimize the IT operation together with a growing framework of autonomic capabilities to reduce human intervention.

System management and monitoring

An extensive portfolio of products provides the industry's most comprehensive offering of monitoring and configuration capabilities for zSeries environments. IBM Tivoli® provides a common user interface with which to monitor the zSeries mainframe systems and proactively manage variable workloads found in complex enterprise environments. Infrastructure cost of ownership is reduced by an integrated view of the resources being used across the entire environment instead of relying on a patchwork of disparate architectures and solutions.

Providing a high-performance, zero-downtime environment involves multiple technologies, such as workload management and automation solutions, that help proactively adjust the infrastructure with shifts in demand. Systems management solutions ensure that applications operate optimally with capacity planning products that help predict future technology needs.

Note: Enhanced data management tools available with DB2 UDB for z/OS Version 8 make the corporate databases easier to manage and maintain. IBM provides earliest availability of compatible tools and utilities in reference to a version general availability date.

Future enhancements (see 8.4.1, "Autonomic computing" on page 195) will extend the expert advisers and automatic management features of DB2 to bring even greater value.

z/OS Managed System Infrastructure for Operations (MSYS for Operations) provides automation for single system and sysplex operations. It helps to simplify operations and improve availability and plays an important role in outage avoidance. MSYS for Operations provides functions that control and manage both hardware and software resources making fully automated solutions possible.

Simplified Configuration z/OS Managed System Infrastructure for Setup (MSYS for Setup) is the strategic solution for product installation, configuration, and

function enablement. MSYS for Setup allows usage of consistent interfaces with wizard-like configuration dialogs.

Database utility and tool software

IBM has responded to requirements for database tools that provide both price performance and functionality. With the latest suite of products, IBM provides functionality comparable to independent software vendor (ISV) offerings and is making them available at an affordable cost. IBM has been putting a new focus on tools that complement its premier database engines, DB2 UDB for z/OS and IMS.

It is now possible to select the necessary tools at an affordable price from a group of products whose depth and breadth demonstrates ongoing commitment to meet the challenges of today's environment for On Demand Business. By providing IBM developed, full function, competitively priced products, a client will not feel forced to invest in additional ISV tool sets.

Important: With a relatively small number of ISVs supporting the mainframe platform, there is a potential for prices to be artificially high. The latest IBM tools help address client fears by providing an element of choice and competitive pricing.

Network

Many data centers today have multitiered server-to-server networks. They include a a collection of middle-tier servers and gateways surrounding the mainframe database and transaction-serving resource. Connecting this multitude of servers together requires the complexity and cost of many networking components. The performance, availability, and response time of the business can depend on the stability of this data center server-to-server network. The more servers that are involved, the greater the number of connections, cost, and complexity is required to install, administer, and maintain.

HyperSockets act like a "network in a box" to provide high-speed communication between virtual servers that are consolidated on z/Architecture.

Security and cryptography

Such encryption software as RACF and Integrated Cryptographic Service Facility (ICSF), together with hardware-based crypto coprocessors, give a robust and cost-effective data security platform on zSeries servers. The ultra-fast built-in encryption facilities of the latest z990 server offer blistering performance to Secure Sockets Layer (SSL), virtual private networking (VPN), and data storing applications. They reduce the development overhead for many applications that

can use these functions. Web-based applications hosted on this platform benefit from the impressive SSL throughputs achievable with the latest server offerings.

IBM's own software applications, such as DB2 UDB for z/OS, make significant use of the z/Architecture encryption support to provide efficient and cost-effective security. The latest multi-level security (MLS) features in DB2 UDB for z/OS Version 8 provide further application-ready function that can drastically reduce development effort.

Java application cost savings

The zAAP, available on the z990 and z890 servers, is an attractively priced specialized processing unit. It provides an economical Java execution environment for z/OS. When configured with general purpose processors within logical LPARs running z/OS, zAAPs may help increase general purpose processor productivity. They may also contribute to lowering the overall cost of computing for z/OS Java technology-based applications.

zAAPs are designed to operate asynchronously with the general processors to execute Java programming under control of the IBM Java virtual machine (JVM). This can help reduce the demands and capacity requirements on general purpose processors, which may then be available for reallocation to other zSeries workloads.

In addition, zAAPs offer these benefits:

- Allow you to integrate and run Java workloads for On Demand Business on the same server as your database, helping to simplify and reduce the infrastructure required for Web applications
- Allow you to run Java workloads, such as WebSphere for z/OS, for a significantly lower TCO than previously possible
- Transparently execute Java workloads, without you having to change your applications

IBM does not impose software charges on zAAP capacity.

Linux software

Linux with a zSeries server can provide the best of both worlds. It can provide the flexibility and openness of Linux, as well as the reliability and robustness of mainframe technology.

Businesses often deploy tens and hundreds of servers to perform such functions as Web serving, mail serving, file/print serving, and File Transfer Protocol (FTP). It is not uncommon for these servers to have low utilization rates (10% to 20%) due to the need for redundancy and the requirement to handle peak workload

periods. The ability to consolidate or rehost these servers on the mainframe provides for much more efficient utilization of server resources, reduced overall complexity of IT infrastructure, and reduction in the people and physical facilities required to manage it.

Note: One z900, running multiple images of Linux, can do the job of hundreds of Sun, Hewlett Packard (HP), or Windows NT servers. Customer data suggests that a TCO advantage can be realized by consolidating as few as 40 Windows NT servers to the mainframe.

One central mainframe server can replace an entire server farm. Think of what that can save in terms of logistical simplification, floor space, cabling and network connectivity, and power bills! The list goes on.

Further cost benefits can be realized, where applicable, by employing the IFL, a CP option dedicated to Linux workloads. The attractively priced IFL processor enables the purchase of additional processing capacity exclusively for Linux workloads, without affecting the service unit price rating of the hardware. This means that an IFL does not increase charges for zSeries software running on general purpose (standard) processors in the server.

7.2.4 System utilization

The unique constant resource monitoring and adjustment of zSeries workload management offers near 100% system utilization. However, leaving smart system technology to manage your applications does not mean a lack of business control. Sophisticated coordination and prioritization of tasks, according to business specified policies, ensures critical requirements and service levels are met.

```
Lost utilization cost = % unused CPU x $ annual system cost
```

Important: This formula may be compounded by the fact that you are failing to fully use an expensive resource, and you are paying for a separate product to do the job.

Reluctance to deploy more than a single application to each UNIX server together with less effective workload management can lead to significant resource under utilization. Estimates of 70% capacity waste is not uncommon.

Workload consolidation

Consolidating multiple server applications on a single zSeries server can mean less labor is required for system management and maintenance. The centralized system management and autonomic computing capabilities of z/Architecture can

also help cut down on the errors and workload-balancing tasks that otherwise consume countless IT staff hours.

A single zSeries server is capable of running multiple diverse operating systems at one time. It extends the consolidation option to Linux applications alongside traditional MVS-based workloads.

7.2.5 Staff utilization

Keeping expensive staff resource fully utilized on tasks that add value to the business is a key element of overall IT return on investment. Both the company and staff members have a vested interest in removing mundane administrative activities.

Interaction with a company's IT infrastructure should be an efficient, effective, and user friendly experience. IBM has made significant progress in bringing the mainframe closer to other server platforms in a commitment to promote usability and virtualization of the company infrastructure. In the following sections we discuss a few of the ways in which different user groups gain from recent z/Architecture enhancements.

Development

Components across the complete IBM software development portfolio from WebSphere, Lotus, and Rational now support the zSeries environment. DB2 provides a comprehensive user friendly Windows and Linux tool suite for database application development with deployment under z/OS. And both the IBM WebSphere Studio and Microsoft .NET Studio integrated development environments (IDE) gain plug-ins for DB2 UDB for z/OS.

Note: Among many of its new and enhanced features, the latest DB2 Development Center provides wizards for fast and easy Java and SQL stored procedure building.

Open standards support through Java, XML, and Linux capabilities bring installations closer to the ideal of developing on any preferred platform for independent server deployment. With full compliance to open standards support, there are fewer reasons for ISVs to avoid the mainframe environment. IBM is partnered with a large number of ISVs with an increasing number of products being available for the platform.

For further detail about zSeries vendor support, refer to the following Web site:

http://www.ibm.com/servers/eserver/zseries/solutions/s390da/

Testing

Both the internal IT test organization and user department testers will benefit, such as the developers mentioned previously, from user friendly interfaces to IT systems. In addition, the consolidation of applications onto fewer platforms reduces training requirements and associated cost.

The ability to standardize on Linux and Windows as the user interface, Java and J2EE as the application environment, and DB2 UDB for z/OS as the enterprise data source, can greatly reduce complexity and usability experience.

Note: The latest advances in QMF for Windows and QMF for WebSphere are prime examples of IBM commitment to simplified information management. See 6.10, “DB2 Query Management Facility” on page 160.

Operational support

Fewer servers, in the data center and scattered across the enterprise, require less operational support staff. A myriad of diverse hardware and software implementations means careful management of such items as:

- License renewals
- Version upgrades
- Capacity upgrades (disk/memory)
- End of service replacement
- Hardware, software, and data auditing

The latest Tivoli Omegamon suite of products provides enterprise-wide systems management capabilities. This allows a view of application workloads across the entire company network.

System programmers and database administrators

The IT industry recognizes the need for less human intervention and greater infrastructure simplification. The autonomic ideal of a human-free data center is a ways off. z/Architecture is at the forefront of this initiative. In the meantime, such products as DB2 UDB for z/OS continue to simplify tasks with the introduction of smarter configuration assistance, online database administration, enhanced optimization, and improved utility products.

Network administrators

The addition of multiple channels and devices accessing corporate systems has increased the need for strategic network planning and continuous management. The need for network simplification, as part of an overall infrastructure simplification initiative, should return cost benefits through the reduction of person hours and potentially reduced cabling and component costs.

Security specialists

Regulatory interest in the use of customer data continues to increase pressure on business to place greater audit control over access to systems and data. Infrastructure simplification can help to provide tighter centralized control over this valuable but highly sensitive company asset.

A comprehensive security policy is essential. However, it can add significant expense through the need for specialist personnel and administration of a burgeoning array of software and hardware security components.

IBM mainframes are recognized as one of the most secure server platforms in terms of technology design and supporting software, such as RACF. Making this robust and proven security infrastructure available to applications can drastically simplify development effort and cut project time scales. Combine this with the inherent security capabilities within DB2 UDB for z/OS and you have a trustworthy security framework to support the growing needs of today's world of On Demand Business.

Note: As with the infrastructure virtualization initiative, there is a long-term objective to make the mainframe like "any other server" in regard to user interaction. Business users, support staff, and developers will be relieved from the burden of learning platform-specific skills.

Customer and user support

Infrastructure and application simplification can lead to reduced training costs and may mean better utilization of existing staff. Support staff costs encompass both external customer and internal user departments. The number of diverse applications and systems will affect the number of personnel needed to provide end user assistance.

Experienced and consistently trained help desk staff are an expensive resource, even when out sourced to a third party. This is recognized by the extensive use of agency resource and automated, self-help solutions. The increasing adoption of Web-based interfaces helps to standardize the look and feel of applications. However it does not remove platform-specific knowledge needed for technical issue resolution.

7.2.6 Quality of service

Quality of service (QoS) affects how business users and external customers perceive aspects and characteristics of the IT applications with which they interface. The cost impacts of bad QoS are indeterminate, from a customer's decision not to return, through bad press and a drop in share price.

The following sections list considerations that you should take into account when you attempt to qualify system experience and its impact on the business.

Customer experience

In regard to customer experience, weigh the following considerations.

- Poor performance: Response times

  You may supply the desired product or service. But consider how many users give up before reaching the checkout due to slow response or cumbersome application.

- Availability

  If your system is not there when your customers want it, day or night, your customers have plenty of alternatives available to them.

- Repeat business: Customer returns

  In the online world, there are always numerous results from a search for service or product. You may not be chosen a second time if the experience was not as expected, and expectations today are extremely high.

- Additional business: Customer referral

  A customer may discuss a bad experience of your system with friends, family, and colleagues but it won't achieve desirable benefits.

Business experience

In regard to the business experience, consider these items.

- Business Intelligence (BI) gathering

  The thriving industry in products and services around BI reflects the value being placed on the effective use of enterprise data. Many companies take the option to extract, maintain, and query operational data copies remote from the master source. With distributed systems and the lack of single source, the collection, collation, and audit of data may become complex and costly.

- IT department responsiveness

  Service-oriented architecture (SOA) advocates the re-assessment of the IT-business relationship with the IT department becoming a true service provider to the enterprise. The business takes greater ownership of IT requirements and starts to evaluate ROI of the service provided in respect to measurable business value.

  An effective and simplified IT infrastructure is imperative to provision a cost-effective operation.

- Time to market

  SOA and Web Services also have an impact in this area. With regard to packaged applications or externally provided pre-built services, today's competitive marketplace requires efficient and responsive IT systems that can respond rapidly to requirements.

  The reuse of existing business function and efficient access to data plays an important part in the effective use of IT systems in the future.

- Application user and transaction cost

  The mainframe has been associated with the provision of large transaction volume and high user access at low unit cost. This gets better with the latest technology platform. The impressive power and throughput of the zSeries servers is well supported by the significant data availability and performance enhancements in DB2 UDB for z/OS.

Lost revenue

A reliable, secure IT infrastructure provides a firm framework for the adoption of new initiatives and industry methodologies, such as Web Services and SOA.

- Missed opportunities

  Late entry to market or failure to satisfy the latest market expectations can impact reputation and result in lost revenue.

- Trade partners who go elsewhere

  The inability to support the latest industry initiatives or inflexible systems may cause potential or existing business partners to pursue alternatives.

System downtime

When the system becomes unavailable, your business performance and profitability are impacted. Reliance on technology has increased to such an extent that personnel find themselves at a loss when unable to interact with the system for their daily work activity.

```
Lost availability cost = % system unavailable x $ annual system cost
```

Important: This formula indicates the visible cost of system downtime. Add in the unquantifiable costs, such as bad press, dented reputation, and missed revenue, and an extended outage can prove expensive.

Here are some additional intangibles that you may want to keep in mind when considering the reliability factor of a technology platform for your business.

- Company share price
- Business reputation

- Business confidence
- Staff productivity

  You want to believe that your staff will find other jobs to keep them busy while their primary system is unavailable. But reality shows that people are easily distracted by unexpected events such as machine crashes. Productivity diminishes while the system is unavailable, but you're still paying salaries.

- Project delivery schedules

  Your development schedules may suffer from lost task time or a freeze on implementation activity may be required to reinstate business confidence. The third strike, and another hidden cost of system downtime, comes if you plough in additional staff resource or overtime to get schedules back on track.

Server performance expansion

You must also consider expanding server performance.

- Scale-up to add an additional processor

  This is included in the readiness of customer demand.

- Scale-out by using Parallel Sysplex

  This includes an external dedicated coupling facility (x2 for resilience) and an integrated CF processor.

7.3 Summary

Through several IBM consulting-led engagements, IBM has developed a methodology that can assess the cost and efficiency of the current IT infrastructure. It starts by producing a series of three server "scorecards" for mainframe, UNIX, and Windows NT servers. Each scorecard includes four financial or cost metrics: IT people efficiency, server utilization, application availability, and total cost. They can give an accurate cost and service-level snapshot of complex server infrastructures in any enterprise. You can use these metrics to model and build alternative future server investment cases.

Most importantly, these engagements have shown that there are major misperceptions of the relative cost of deploying new applications on mainframes, UNIX and Windows NT servers.

The conclusion of many of these cost studies is that, for a "typical" online application, the following cost ratios hold true, when all the people, hardware and software costs are included in a three-year business case.

For a zSeries solution, the incremental hardware, software and people costs are, for example, USD $1 million/year. This solution cost is dominated by software cost, mainly ISV software costs.

For a comparable UNIX solution, the incremental costs are typically $1.6 million to $2.4 million/year. These solutions are dominated by people cost.

For a Windows NT solution, the incremental costs are usually in the range of $2 million to $4 million/year and are dominated by people cost. In addition, unscheduled and scheduled outage costs can typically add $1 million to $2 million to the UNIX and Windows NT three-year cases.

TCO studies are showing more and more that Web application consolidation is financially viable and reduces the role of server farms while strongly re-establishing the role of the high-end server for the mainframe.

8

The future of DB2 for z/OS

DB2 for z/OS has a rich heritage of features that has expanded over the past 12 versions, as client requirements to satisfy new and more demanding workloads have increased. IT environments are faced with globalization, regulatory pressures, and on demand access to structured and unstructured data. Where there is demand for a reliable, available, and scalable information server to meet these requirements, DB2 for z/OS is found as a core part of the architecture. Significant features have been added to make DB2 for z/OS more accessible and manageable, positioning DB2 as *the* data server for all of your enterprise workloads.

The future for DB2 is bright. You can count on IBM being an innovative company, one of the top companies in terms of patents in their portfolio. DB2 development is advancing the engine to continue to deliver the features that clients need most. You can count on DB2 for z/OS to run your new workloads and to grow with existing workloads. New SQL capabilities will enable your future application needs, package application requirements, and porting of your data to DB2 for z/OS. Enhanced portability will enable you to gain higher levels of business resiliency for existing applications. Innovative security enhancements will offer flexibility in authorizations and role definitions.

In addition, industry surveys support DB2 for z/OS total cost of ownership (TCO) advantages when the assessment criteria is responsibly complete. Cost comparisons must include the higher utilization rates obtainable with zSeries and lower administrative costs achieved with centralized systems. Through

advancements in DB2's autonomic capabilities, the DB2 for z/OS financial advantage will continue to be compelling.

The long-term IBM vision for DB2 on zSeries is to continue to closely couple DB2 and the z/Architecture. IBM will package the database, operating system, and hardware together and roll in a "DB2 on Demand" black box. This will allow you to easily configure and manage this highly capable and powerful environment from a browser, without needing knowledge of such things as TSO and VSAM.

The intention is to dramatically reduce people costs. IBM will do this by eliminating the need for zSeries skills and reducing database administrator skills to an absolute minimum and then use an open, Web interface. Less specialized knowledge requirements imply lower people costs, with people as the principle cost going forward since hardware and software are becoming commodities.

This chapter covers some of the future advances in new workloads, business resiliency, DB2 for z/OS platform synergy, and TCO.

8.1 New workloads

DB2 for z/OS has been supporting enterprise applications for years. With the explosion of Web applications, DB2 and the zSeries platform have once again shown their resiliency to handle new z/OS and distributed applications while supporting unpredictable workloads. These features further DB2 for z/OS as a natural choice for new and existing applications in your environment.

8.1.1 Native SQL stored procedures

Stored procedures written in SQL procedure language enhance portability and ease of use when using DB2 for z/OS as your enterprise information source. This language is an ANSI standard language. It is similar to the proprietary stored procedure languages of several competitive databases, which assists in migrating and porting to DB2 for z/OS.

SQL stored procedures are supported by the DB2 Development Center tooling, providing an environment to code, test, and debug modules from your connected workstation. This language is currently converted to C when the CREATE PROCEDURE statement is executed. The C program is then automatically prepared, compiled, linked, and bound. The developer does not need to work with the C code.

SQL stored procedures code will be natively integrated into the DB2 engine, eliminating the conversion to C. Additionally, extensions to the bind command will allow for the promotion of the program and access paths between environments without needing to recreate the stored procedure.

8.1.2 Integrated XML

DB2 for z/OS extensibility has been implemented via *extenders*. Extenders for text, image, audio, video, and XML are delivered with DB2 and are optional installation components. They provide the tooling, user defined data types (UDT), user defined functions (UDF), and stored procedures to managed non-relational data. The XML extender provides for the storage and indexing of an XML document as a character large object (CLOB), or for the shredding of the XML document into relational columns for storage and query.

DB2 Version 8 expanded on XML support by implementing several XML publishing operations as built-in DB2 functions. This allows you to perform XML document composition from relational data with improved performance without the XML Extender.

There will be expanded support of XML in DB2 by integrating more features into the engine. This includes an XML data type, native storage of XML documents, integration of the XPath language, and catalog extensions to support definitions of XML schemas. Utilities will support creation and maintenance of XML data.

8.1.3 Portability, ERP support, and family compatibility

DB2 SQL and DDL features continue to improve. This benefits application vendors, custom developed systems, and DBAs by driving compatibility across the DB2 family and aiding in portability. Enhanced portability assists in the development of common code across platforms and the migration of applications to DB2 for z/OS. It allows your business to achieve higher levels of reliability, availability, and scalability without significant changes to the applications.

Truncate

Truncate Table provides for the rapid removal of rows from a table. You can use this function to delete the contents of a table before applying data from a DSN1COPY.

New data type support

Support will be made for:

- Decimal floating point numbers, similar to calculator mathematics and supporting the IEEE standard
- BIGINT support of double word (8 byte) integer values
- VARBINARY, providing better comparison support for binary strings

Unicode enhancements

Unicode improvements in DB2 Version 7 and Version 8 have made Unicode the native mode of DB2. These improvements will continue in the future, allowing character data to have enhanced comparisons and collation.

Large object improvements

Large objects (LOBs) were introduced in DB2 Version 6. Usage has increased substantially in the past few years, and major enhancements have been made in DB2 Version 8.

APARs on Version 8 deliver the ability to use utilities for loading and unloading large LOB data. File reference variables are used to let the large objects be accessed from data sets instead of from storage. The abilities to reorganize and to recover space are provided.

Future changes will help with improved function and usability, DB2 family compatibility, cost of ownership, performance, and scalability.

Merge

Online transaction processing (OLTP) workloads that need to place data for several rows into a table will benefit from merge SQL. This capability will allow for the updating of existing rows and creation of new rows through a single SQL statement.

Multi-row insert-type capability is extended, via the MERGE statement, to take an array of column values from a program and perform insert and update operations against a table. DB2 will use the matching criteria in the merge SQL to update existing rows and perform inserts for rows that do not exist, through a single SQL statement.

Select from DELETE, UPDATE, and MERGE

The object-relational capabilities of DB2 allow for the incorporation of business logic into the database. This extends the power of SQL. Sometimes the application needs to know the results of this logic, when applied to the SQL issued. A subsequent SELECT for the data adds complexity and execution time to the application.

The *insert within select* feature of DB2 for z/OS Version 8 has been expanded to include the retrieval of columns from rows that are modified via DELETE, UPDATE, and MERGE SQL. One SQL call to DB2 modifies the table contents and returns the resultant changes to the application program.

When used with DELETE, the application now has the option to code a destructive read from a table. This is particularly useful when a table is used as a data queue, as with many vendor packages.

Subquery improvements

Correlated and non-correlated subqueries will benefit from improved optimization. They will provide added flexibility with the support of ORDER BY and FETCH FIRST clauses.

Additional family compatibility

Additional compatibility will be added for:

- Default databases and table spaces
- Automatic unique indexes to support defined primary keys

8.2 Business resiliency

This section addresses the “abilities”, that together make too long of a title. Among these are reliability, availability, scalability, serviceability, accessibility, and manageability.

The zSeries has a long proven history of premium service in the spaces of reliability, availability, and scalability. The last three versions of DB2, since the addition of *Universal Database* (UDB) to the DB2 UDB Server for OS/390 title, have introduced significant function into the usability related “abilities”.

The following sections show this trend continuing, with expanded support for growing and new workloads.

8.2.1 Enhanced security

Today’s computing environment is subject to increasing regulatory pressures and potentially malicious attacks. The security of the information, to which you have been entrusted, has never been more critical. DB2 for z/OS already resides on one of the most secure platforms in the industry. It provides secure access to your data through internal or external authorities and privileges, encryption, and multilevel security with row-level granularity enforcement.

Security can be difficult to manage. The plan for DB2 is to provide you with additional capability while assisting you with security management.

Trusted security context

Today, you have the option to set a system parameter which indicates to DB2 that all connections are to be trusted. It is unlikely that all connection types, such as DRDA, RRS, TSO, and batch, from all sources will fit into this category. It is likely that only a subset of connection requests for any type and source may be trusted or that you want to restrict trusted connections to a specific server. More granular flexibility will allow for the definition of *trusted connection objects*.

Once defined, connections from specific users via defined attachments and source servers will allow trusted connections to DB2. The users defined in this context can also be defined to obtain a *database role*.

Database role

A database role is a virtual authorization ID that is assigned to the user via the context mentioned previously. DB2 privileges are assigned to the defined role. The role exists as an object independent of its creator, so creation of the role does not produce a dependency on its creator.

The role can be assigned and removed from individuals via the trusted authorization context as needed. This allows a DBA to perform object maintenance during a change control window on a Saturday night, for example. But when Monday arrives, they do not have the authority to do this same work. Auditing trails of the work completed during the maintenance window are available for verification by a security administrator or auditor.

8.2.2 Enhanced manageability

This section describes some of the functions that reduce the workload of data administrators.

Fast replacement of a table with another

This function will allow you to generate a copy of a current table with the same attributes and same data, in the same table space. It appears to an application or user as a very fast replacement of data within a table. Web-based applications, striving for maximum availability will benefit from the option of implementing a pair of tables that are clones of each others' structure. Copies for application testing and auditing can be easily created. These *clone* tables will have the unique ability to change names quickly. Applications can quickly and almost transparently switch between dual mirror tables.

The tables can be partitioned or non-partitioned, and are created with the CREATE TABLE syntax. The primary table's structure, including indexes, large objects (LOBs) and before triggers, will be copied. Information can then be inserted or loaded into the copy table, and the copy can have its own image copies. When the data in the copy table needs to become active to an application, an ALTER statement will switch the name, providing fast replacement of the original data.

Partitioned by growth

Partitioned tables have required key ranges to determine the target partition for row placement. When a table is partitioned, you gain the benefits of scaling objects to hold more data. You also benefit from more granular locking and parallel operations by spreading the data over more data sets.

The option to partition by growth will allow segmented tables to be partitioned as they grow, without needing key ranges. These segmented tables will gain increased table space limits and the SQL and utility parallelism, afforded for partitioned tables.

8.2.3 Table Append option

The Table Append option offers increased performance for inserting data into the end of a table. It reduces the instructions used to target the locations for new rows. Index maintenance techniques are also made more efficient when DB2 detects that entries are added to the beginning or ending of the index.

8.2.4 Index changes

A series of recent indexing enhancements have increased the scalability and manageability of DB2. The following additional enhancements are on their way:

- Support of larger index page sizes
- Indexing of expressions
- Index compression

8.2.5 Temporary storage architecture and use

Many recent improvements have enhanced the use of temporary storage when it is needed by the DB2 engine, for example, for interim materialized result sets. The different temporary pools are converged to a single source. Use of temporary storage is eliminated when processing small amounts of data.

8.3 DB2 for z/OS platform synergy

DB2's deep synergy with z/OS and zSeries enables the innovation which makes superior business resiliency possible. This history continues to advance in support of distributed and growing workloads.

8.3.1 Shared memory and distributed connections

Distributed connections to DB2 for z/OS will benefit from z/OS V1R7 changes that DB2 will exploit. DB2's distributed communication processes (the *distributed address space*) will access data directly from the database manager address space, instead of moving the data. The distributed address space will also exploit 64-bit addressing, as the database manager and lock manager address spaces do today with Version 8.

This internal change will benefit new and existing workloads, where distributed communications are configured with another logical partition (LPAR) or to an application running on the zSeries platform.

8.3.2 Index compression

The ability to compress data will be expanded to index structures. This will bring disk savings to large indexes and further use the synergy with zSeries.

8.4 Total cost of ownership

The informed IT investor knows that the zSeries platform and DB2 for z/OS combine to deliver significant value for the money spent. Providing the best TCO today is not enough when making a decision for your future. DB2 for z/OS improvements will continue to drive the maximum value for your IT investment.

8.4.1 Autonomic computing

Autonomic computing is a cross-IBM initiative with the mission to increase the self-managing capabilities of IT infrastructures. It will do this through the definition of business policies and service-level agreements (SLAs) that enable the system to automatically manage processes. This frees IT resources to concentrate on improving the business rather than running the system.

Autonomic computing is an essential attribute of the IBM on demand strategy. It is in the DNA of every product and solution that IBM offers. Our clients derive business value through increased automation which is delivered through autonomic computing. The goal of the autonomic computing initiative at IBM is to help clients build more automated IT infrastructures and to improve return on investment, deployment time, and quality of service.

Autonomic computing is a continuously evolving and dynamic state. It establishes the correct balance between that which is managed by a person and what is managed by the system. It lets you focus on your business, not your infrastructure. It makes the systems self-configuring, self-healing, self-optimizing and self-protecting. It means that the systems do the work, freeing IT professionals to focus on other critical business needs.

Self-configuring

Systems that are self-configuring have the ability to:

- Adapt dynamically to changing environments
- Add and remove components to and from the systems, changing the environment as necessary depending on workloads

Imagine a situation where you no longer need to worry about the interactions between components when you add a new application or capability. That's part of

what self-configuring is all about. The bottom line is IT agility, the ability to respond rapidly to changing demands.

Self-healing

Self-healing is all about business resiliency, the ability to discover, diagnose, prevent, and recover from disruptions, keeping the system going. With increased business dependence on IT and internal and external users demanding 24x7 availability, this is more important than ever.

Self-optimizing

Self-optimizing is about operational efficiency. It is about tuning the resources, balancing workloads, and making maximum use of the IT resources that are available. In today's world, this has to be done and redone continuously since the workloads are so variable.

In the past, when workloads were more predictable, it was a question of tuning the systems to support the workload and then letting them run. Today, the workloads are changing dynamically and dramatically. Systems need to continuously monitor and self-tune, while adapting and learning from the environment around them.

Self-protecting

Now that our IT systems are open to the public, so to speak, security is ever more important. The concept of "self-protecting" is about the ability of the system to secure information and resources. It must be able to anticipate, detect, identify, and protect from attacks of any kind.

Autonomic computing begins with the infrastructure that is already in place and evolves to become more and more self-managing. IBM already has several products with autonomic capabilities that can benefit our clients.

DB2 for z/OS offers several features, utilities, system parameters, and commands today. They span the spectrum from the *basic* to the *autonomic* levels as shown in Figure 8-1.

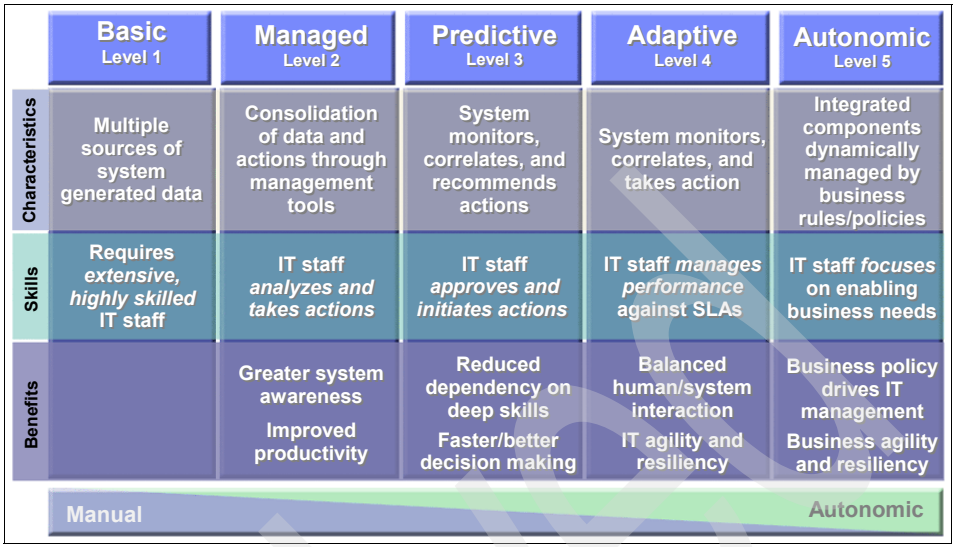

Figure 8-1 Autonomic evolution

Some examples of these features are:

- The cost-based optimizer

  Choosing from a myriad of access paths, it determines the least expensive access to your data, based on available information. Then during execution, the DB2 engine gathers heuristics about run-time behavior which can modify data access dynamically.

- A sliding scale secondary allocation for VSAM data sets

  It works to take best advantage of the disk space used by an object and reduces the administrative costs associated with tracking this activity.

- Automatic identification and invocation of recoveries at a page level if needed, without disrupting access to other data

- Detection and recovery from several component failure scenarios, such as duplexed objects in the coupling or logging facilities

- Real-time statistics that are collected during execution, which provide information about utilities that should be run

- Numerous internal memory guards and adjustments to reallocate resources, adjust sizing definitions, or govern internal processes

  Some examples include storage contraction, buffer pool monitoring, sequential detection, and data sharing inter-DB2 read/write interest.

- Sophisticated workload management, offering continuous resource adjustments to meet business goal responses

- Lock priority management, which can advance a lower priority workload that is holding a resource needed by a higher priority workload
- Identification of long running tasks that can impact later recoveries, should a recovery be needed

DB2 is just getting started. Additional autonomics will be incorporated in the engine and in the tools that support the engine. Meanwhile, watch for more design and configuration support, online goal-based tuning of thresholds, deeper synergies with zSeries and z/OS, and learning-based optimization.

8.4.2 Optimization

IBM invented and has been improving cost-based optimization from the beginning. This core component of DB2 needs good information to make good access path decisions. The ability to collect statistics about your data has been continually expanded.

In today's fast changing IT industry, application complexity grows rapidly and transaction volumes can fluctuate wildly. More enterprise resource planning (ERP) applications, such as SAP and Peoplesoft, as well as tools, generate tables, indexes, and queries. In such an environment, it is difficult for a DBA to determine the problem with a particular query or set of queries that are performing poorly. At the same time, the majority of the IT budget is allocated to people costs, as hardware and software costs reach commodity levels.

DB2's advanced work in the collection of multi-column and nonuniform distribution statistics, and reoptimization run-time techniques, positions DB2 to handle a mix of workloads in your production environment today. Further improvements will examine the gathering of more information during execution of the SQL. They will dynamically change the original access path decisions and learn information that can be made available for future executions. The tooling around these progressive optimization enhancements will provide an administrative environment to advance autonomic design choices and serviceability. These advanced features will provide execution and people efficiencies, making the most of your IT investment.

Optimization Service Center

The Optimizer Service Center (OSC) will initially capture meaningful information about your workload. It will provide helpful serviceability features by exposing insightful optimizer information.

In a following phase, the OSC will analyze what is running and make specific workload recommendations to reach optimal performance. It will suggest:

- Physical database design changes, such as indexes, materialized query tables (MQTs), and partitioning
- Statistics to be collected
- System configuration changes

At the next level, OSC can act on your behalf, if permitted, and implement the changes autonomically for you. More of the centers and advisors, shown in Figure 8-2, will be exposed for administrator use. This will free administrators to address more important tasks, such as delivering new features that further leverage your IT investment.

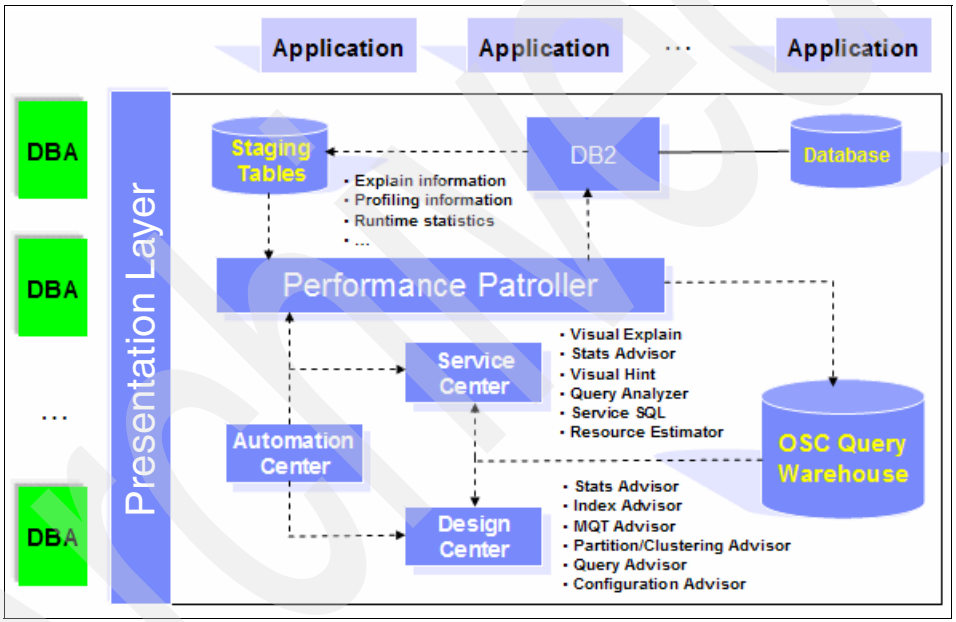

Figure 8-2 Optimization Service Center

IBM will stay ahead of the competition. It will continue to deliver advanced technology in DB2 and address the difficult computing issues of today and tomorrow. Optimization enhancements will position DB2 to deliver superior results, maximizing your investment in hardware, software, and technical skills.

8.5 Conclusion

Janet Perna, Information Management General Manage for IBM Software Group, best summarized the future of DB2 in her keynote address at the 2004 DB2 Technical Conference.

She stated, “From financial institutions managing investment portfolios to a pharmaceutical companies inventing medical cures, companies across all industries are facing a whirlwind of business pressures and regulatory legislation. Focused on improving customer service and reducing operational costs, companies are searching for new ways to leverage information to run their business with more insight.

At the same time, the computing infrastructure has become extremely heterogeneous, widely distributed and increasingly complex. Companies are challenged to extend traditional information management techniques to new forms of content from many different sources. As market leader, IBM DB2 technology, partnerships and solutions are evolving to meet these challenges.”

Abbreviations and acronyms

AC	autonomic computing
API	application programming interface
AR	application requester
AS	application server
B2B	business-to-business
BI	Business Intelligence
BLOB	binary large object
CBU	Capacity BackUp
CCSID	coded character set identifier
CDW	central data warehouse
CF	coupling facility
CICS	Customer Information Control System
CLOB	character large object
CMOS	complementary metal oxide semiconductor
CP	central processor
CPU	central processing unit
CTE	common table expression
CUoD	Capacity Upgrade on Demand
DAC	discretionary access control
DB	database
DB2	Database 2
DBA	database administrator
DBCLOB	double-byte character large object
DBCS	double-byte character set
DBM1	database master address space
DDF	distributed data facility
DDL	data definition language
DPSI	data partitioning secondary index
DRDA	Distributed Relational Data Architecture
DSNZPARMs	DB2's system configuration parameters
DSS	decision support systems
EAI	enterprise application integration
EAS	Enterprise Application Solution
EDM	environmental descriptor manager
EJB	Enterprise JavaBean
ENFM	enable-new-function mode
ERP	enterprise resource planning
ESS	Enterprise Storage Server
EWLC	Entry Workload License Charges
eWLM	Enterprise Workload Manager
GDPS	Geographically Dispersed Parallel Sysplex
GLBA	Gramm-Leach-Bliley Act of 1999
GUI	graphical user interface
HALDB	High Availability Large Databases
HTTP	Hypertext Transfer Protocol
HW	hardware
IBM	International Business Machines Corporation
ICSF	Integrated Cryptographic Service Facility
IDE	integrated development environments

IFL	Integrated Facility for Linux
IMS	Information Management System
IORP	I/O Request Priority
IRD	Intelligent Resource Director
IRLM	internal resource lock manager
ISV	independent software vendor
IT	information technology
ITSO	International Technical Support Organization
J2EE	Java 2 Enterprise Edition
JDBC	Java Database Connectivity
JTA	Java Transaction API
JTS	Java Transaction Service
JVM	Java Virtual Machine
LCU	Logical Control Unit
LDAP	Lightweight Directory Access Protocol
LPAR	logical partition
LSS	logical subsystem
MAC	mandatory access control
MBps	megabytes per second
MLS	multi-level security
MQT	materialized query table
MTBF	mean time between failures
MVS	Multiple Virtual Storage
NALC	New Application License Charge
NFM	new-function mode
NFS	Network File System
NPSI	nonpartitioned secondary index
ODBC	Open Database Connectivity
ODS	Operational Data Store
OLE	Object Link Embedded
OLTP	online transaction processing
OSC	optimizer service center
PAV	Parallel Access Volume
PCICA	Peripheral Component Interface Cryptographic Accelerator
PCICC	PCI Cryptographic Coprocessor
PPRC	Peer to Peer Remote Copy
PR/SM	Processor Resource/System Manager
QMF	Query Management Facility
QoS	Quality of Service
QPP	Quality Partnership Program
RACF	Resource Access Control Facility
RAS	reliability, availability and serviceability
RDBMS	relational database management system
RDS	relational data system
RI	Referential Integrity
ROI	return on investment
SAN	storage area networks
SBCS	store single byte character set
SDP	Software Development Platform
SLA	service-level agreement
SOA	service-oriented architecture
SOAP	Simple Object Access Protocol
SQL	Structured Query Language
SQLJ	Structured Query Language for Java
SRM	Service Request Manager
SSL	Secure Sockets Layer
TCO	total cost of ownership
TPF	Transaction Processing Facility

UA	Unit Addresses
UCB	Unit Control Block
UDB	Universal Database
UDF	user-defined functions
UDT	user-defined data types
UR	unit of recovery
USS	UNIX System Services
VIPA	Virtual IP Addressing
VLDB	very large database
VM	virtual machine
VSE	Virtual Storage Extended
VSIP	Visual Studio Integrator Program
VWLC	Variable Workload License Charges
wizards	Web-based assistants
WLC	Workload License Charges
WLM	Workload Manager
WSDL	Web Services Definition Language
XA	Extended Architecture
XML	Extensible Markup Language
z800	zSeries 800
z890	zSeries 890
z990	zSeries 990
zAAP	zSeries Application Assist Processor
zELC	zSeries Entry License Charge

Related publications

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

IBM Redbooks

For information about ordering these publications, see “How to get IBM Redbooks” on page 207. Note that some of the documents referenced here may be available in softcopy only.

- *Achieving the Highest Levels of Parallel Sysplex Availability in a DB2 Environment*, REDP-3960
- Consolidating UNIX systems onto OS/390, SG24-2090
- *Achieving the Highest Levels of Parallel Sysplex Availability*, SG24-6061
- *DB2 UDB for z/OS Version 8: Everything You Ever Wanted to Know, ... and More,* SG24-6079
- *Disaster Recovery with DB2 UDB for z/OS,* SG24-6370
- *zSeries Application Assist Processor (zAAP) Implementation,* SG24-6386
- *Multilevel Security and DB2 Row-Level Security Revealed,* SG24-6480
- *WebSphere Information Integrator Q Replication: Fast Track Implementation Scenarios*, SG24-6487

Other publications

These publications are also relevant as further information sources:

- *Large Systems Performance Reference for IBM zSeries and S/390*, SC28-1187
- Graziano Sloan, Susan. *The Official Introduction to DB2 for z/OS: Covers DB2 Universal Database for z/OS Version 8, 2nd Edition*, Prentice Hall PTR, 26 May 2004. ISBN 0-13-147750-1.
- “z/OS support for the IBM TotalStorage Enterprise Storage Server”, *IBM Systems Journal*, Vol. 42, No. 2, 2003.

- Ruddy, Jim. Articles “Programs vs. Utilities” in *DB2 Magazine* from Quarter 2, 2003 and “Programs vs. Utilities Revisited” in *DB2 Magazine* from Quarter 3, 2004.
- *SQL Reference for Cross-Platform Development Version 2*

  http://www.ibm.com/software/data/db2/zos/v8books.html
- *The IT co-op program at IBM Poughkeepsie* white paper

  http://www.developer.ibm.com/us/en/university/scholars/products/pdfs/ITSkillsCo-opWhitePaper.pdf
- Mastrobattista, David of the Giga Information Group, Inc. *Mainframe Linux Solutions Mature with mySAP.com for IBM eServer zSeries* paper, 19 March 2002

  http://www.ibm.com/servers/eserver/zseries/library/whitepapers/pdf/gigamar19mysaplinuxonz.pdf
- *The zSeries Server Consolidation and Application Integration* white paper, April 2004

  http://www.ibm.com/servers/eserver/zseries/library/whitepapers/pdf/gm130254.pdf
- *Building a business case: The role of server consolidation in IBM's transformation to an e-business* white paper, September 2002

  http://www.ibm.com/servers/solutions/serverconsolidation/pdf/cio.pdf
- *Positioning z/OS and Linux for zSeries* white paper, April 2004

  http://www.ibm.com/servers/eserver/zseries/library/literature/papers.html
- Stuhler, Julian of Triton Consulting. *DB2 for z/OS Version 8 - Driving Business Value* white paper

  http://www.ibm.com/support/docview.wss?rs=64&context=SSEPEK&uid=swg27005465

Online resources

These Web sites and URLs are also relevant as further information sources:

- DB2 UDB for z/OS Version 8

  http://www.ibm.com/software/data/db2/zos/db2zosv8.html
- The DB2 Information Management Software Information Center

  http://publib.boulder.ibm.com/infocenter/dzichelp/index.jsp
- “The Big Picture: IBM DB2 Information Management Software and DB2 Universal Database”, article by Jeff Jones from 09 January 2003

http://www.ibm.com/developerworks/db2/library/techarticle/0301jones/0301jones.html

- Testimonials on zSeries meeting the challenges of On Demand Business

  http://www.ibm.com/servers/eserver/zseries/testimonials

- The Mainframe Charter, announced by IBM on 22 August 2003

  http://www.ibm.com/servers/eserver/zseries/announce/charter

- zSeries ISV support and solutions

  http://www.ibm.com/servers/eserver/zseries/solutions/s390da/

- zSeries 890

  http://www.ibm.com/servers/eserver/zseries/z890/

- History of relational

  http://www.research.ibm.com/resources/news/20030423_edgarpassaway.shtml
  http://www.research.ibm.com/resources/awards.shtml
  http://www.intelligententerprise.com/031118/618strategic1_1.jhtml

- Business resiliency

  http://www.ibm.com/servers/eserver/zseries/resiliency/

How to get IBM Redbooks

You can search for, view, or download Redbooks, Redpapers, Hints and Tips, draft publications and Additional materials, as well as order hardcopy Redbooks or CD-ROMs, at this Web site:

ibm.com/redbooks

Help from IBM

IBM Support and downloads

ibm.com/support

IBM Global Services

ibm.com/services

Index

D

J

K

L

M

N

O

P

Q

R

S

T

U

V

W

X

Z

(0.2”spine)
0.17”<->0.473”
90<->249 pages

The Business Value of DB2 UDB for z/OS

Learn about the synergy of DB2 with IBM @server zSeries and z/OS

Discover how you can benefit from the richness in function

Be ready for On Demand Business

We are in a new phase of On Demand Business. Companies must respond to fluctuating market conditions and provide products and services on demand to customers. DB2 for z/OS and IBM @server technology enable you to maintain your company's core competencies. They help you to manage global market and competition changes with consistent availability, security, and privacy, everywhere, all the time.

DB2 continues to deliver rich functionality for scalable and highly available data for enterprise-scale, on demand applications. Combine the power and capacity of IBM @server zSeries and the high performance and high availability of z/OS with the new version of DB2. In doing so, you'll expand and extend your applications in the increasingly competitive on demand environment. DB2 Universal Database (UDB) for z/OS Version 8 delivers power, flexibility, and DB2 family compatibility through synergy with zSeries and z/OS.

This IBM Redbook helps you to position DB2 UDB for z/OS as a springboard for the future. It examines how DB2 Database Servers integrate with, and enable the sharing of information across, multiple platforms. It looks at how the zSeries 990 server manages the explosion of data with reliability and security. Plus it explores how DB2 autonomic computing functions reduce the skills and staffing requirements to minimize cost outlays and reduce risk.

INTERNATIONAL TECHNICAL SUPPORT ORGANIZATION

BUILDING TECHNICAL INFORMATION BASED ON PRACTICAL EXPERIENCE

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

For more information:
ibm.com/redbooks

SG24-6763-00 ISBN 0738490245