ISSUE 4, 2010 / VOLUME 15 NO. 4

IBM

data management

IBM.COM/DMMAGAZINE KNOWLEDGE. PERFORMANCE. RESULTS.

Migrate from_____to DB2 in weeks.*

(ANSWER ON PAGE 20)

ALSO INSIDE:
Hard Numbers on Informix TCO
Page 38

* Hey, strategy meetings take forever.

Smarter technology for a Smarter Planet:

It's time to ask smarter questions.

What exactly does a benchmark mean? For the last five years, IBM DB2® on Power Systems™ has ranked first on three of the industry's leading performance benchmarks, longer than Oracle and Microsoft combined.[1] But shouldn't we be asking our IT to deliver more than just raw performance? What really matters isn't some abstract measure of performance, it's what companies actually do with it. For instance, Coca-Cola Bottling Company is using DB2 on Power to reduce licensing, maintenance and storage fees by $350,000. EuResist is using an integrated analytics solution to predict the most effective drug combinations for individuals with HIV, with 78% accuracy. And the Dubai Gold & Commodities Exchange is working with IBM Security Services to achieve system uptime of over 99.9%. On a smarter planet, these are the benchmarks that matter.

A smarter business is built on smarter software, systems and services.
Let's build a smarter planet. ibm.com/questions

1. Based on number of days of performance leadership for the TPC-C, TPC-H 10TB and SAP 3-Tier SD benchmarks between June 1, 2005 and June 1, 2010. For more information, see http://www.tpc.org and http://www.sap.com/solutions/benchmark. TPC, TPC-C and TPC-H are trademarks of the TPC. IBM, the IBM logo, ibm.com, DB2, Power Systems, Smarter Planet and the planet icon are trademarks of International Business Machines Corp., registered in many jurisdictions worldwide. Other product and service names might be trademarks of IBM or other companies. A current list of IBM trademarks is available on the Web at www.ibm.com/legal/copytrade.shtml. © International Business Machines Corporation 2010.

IBM.COM/DMMAGAZINE

IBM data management

KNOWLEDGE. PERFORMANCE. RESULTS.

ISSUE 4, 2010 / VOLUME 15 NO. 4

Features

In-Depth Report:

Migrate to DB2

20 **Whether you're coming from Oracle or Sybase,** DB2 9.7 makes migration faster and easier than you previously thought possible

28 **What's New: DB2 10 for z/OS**
The future is now for CPU savings, performance gains, and new tools for time-based data

34 **Development Notebook: The Evolution of the IBM Smart Analytics System**
The inside story from the team behind the business analytics blockbuster

38

Informix TCO: The Facts Are In

Analyst numbers tell a simple story: Informix costs less to run

Departments

4 **Editor's Note**
By Cameron Crotty

6 **NewsBytes**

10 **IIUG User View**
By Stuart Litel

12 **IDUG User View**
By David Beulke

13 **Data Manager**
You Catch More Flies with Honey
By Merv Adrian

45 **Data Architect**
DB2 Indexes and Query Performance: Part 1
By Robert Catterall

47 **Distributed DBA**
Adding and Removing Data Partitions
By Roger E. Sanders

50 **Programmers Only**
Does the Order of SQL Predicates Matter?
By Bonnie Baker

53 **Informix DBA**
Building Fast Data Warehouse Schemas: Part 3
By Lester Knutsen

56 **Smarter is...**
Breaking Data and Disease Down to Size
By David Geer

Ad Index

Applied Analytix, Inc. **37**
www.applied-analytix.com

DBI **C3**
www.DBIsoftware.com

Fourth Millennium Technologies **23**
www.fmtusa.com

IBM **C2**
www.ibm.com

IBM Client Reference Program **51**
custref@us.ibm.com

IBM Information On Demand Virtual **27**
bit.ly/ioddm4

IBM Server and Storage **32-33, C4**
www.ibmdatamanagement.com

Intel **16-19, 41-44**
www.intel.com

International DB2 Users Group **5**
www.idug.org

Melissa Data **15**
www.melissadata.com

Quest Software **9**
www.quest.com

Relational Architects International **31**
www.relarc.com

Responsive Systems **3**
www.responsivesystems.com

Safari Books Online **11**
www.safaribooksonline.com

IBM Information On Demand **55**
www.ibm.com/events/informationondemand

IOD 2010 Conference Sponsors

Intel
SAP
Netezza
Compuware Corporation
CONSIST
CSC
Dassault Systèmes
Deloitte
enChoice, Inc.
Esri
Fourth Millennium Technologies
HCL
Infosys Technologies Ltd
JCB Partners
Jeskell, Inc.
LPA Systems
Mellmo
NetApp
Niteo Partners
Perficient
Pronto Software
QueBIT Consulting
Renewtek Pty Ltd
Sogeti
SYSCOM, Inc.
Thunderhead
TriTek Solutions

Optimize Both Performance & DB2 Memory Utilization With

Buffer Pool Tool® For DB2

Reduce Your Processing Costs, Save $$ Now!

RESPONSIVE SYSTEMS

281 Hwy 79 Morganville, NJ 07751
Tel: 732 972.1261 Fax: 732 972.9416
Web: www.responsivesystems.com

IBM data management
KNOWLEDGE. PERFORMANCE. RESULTS.

EDITOR
Cameron Crotty
editor@tdagroup.com

MANAGING EDITOR
Stephanie S. McLoughlin

ART DIRECTOR
Iva Frank

DESIGNERS
David Chan, Lalaine Gagni, Margie Preston

CONTRIBUTING WRITERS
Merv Adrian, Bonnie Baker, David Beulke, Ives Brant, Robert Catterall, David Geer, Terrie Jacopi, Lester Knutsen, Stuart Litel, Roger E. Sanders, Greg Thomas, Chris Young

EDITORIAL BOARD OF DIRECTORS
Jeannie Cramer, Jeff Jones, Steve Miller, Jennifer Reese, Bernie Spang

AD SALES
Jill Thiry
advertise@tdagroup.com

AD COORDINATORS
Katherine Hartlove
Kim Johnsen

SPECIAL THANKS TO
Lea Anne Bantsari
Matthew Golden

SUBSCRIPTION SERVICES AND REPRINTS
To subscribe to the print or digital version of *IBM Data Management* magazine, change your address, or make other updates to your information, please go to **ibm.com**/dmmagazine. For instant access to the *IBM Data Management* magazine digital edition, visit ibmdmmagazinedigital.com. For information about reprints, please send an e-mail to customerservice@tdagroup.com.

tda GROUP

PRESIDENT
Paul Gustafson

VICE PRESIDENT, STRATEGY AND PROGRAMS
Nicole Sommerfeld

VICE PRESIDENT, EDITORIAL DIRECTOR
Debra McDonald

VICE PRESIDENT, CONTENT SERVICES
Paul Carlstrom

Printed in the U.S.A.

How long does it take to migrate from Oracle or Sybase to DB2? Not too long ago, the answer to that question was "months." These days, well—the answer is on the cover of this issue. Things get even more interesting when you realize that "weeks" includes the entire project from beginning to end, from the first kickoff meeting to the last phases of testing. The actual migration of code and data can often be accomplished in *days*.

You've probably already heard that DB2 9.7 includes a set of compatibility features that are nothing short of astounding. According to IBM migration engineers, the typical organization migrating from Oracle to DB2 needs to modify only 2 percent of its code. The rest just works.

What's almost as amazing is the ongoing behind-the-scenes effort to whittle down that 2 percent. IBM engineers are out working with organizations every day to enable their applications for DB2, and are then bringing back that experience to plug into software updates. Make no mistake: the path to DB2 is smooth and getting smoother. More details are in our featured report on migration, and you owe it to yourself—and your organization—to take a look.

Migration isn't the only big news in this issue. Two blockbuster studies have been released in quick succession that document the serious total cost of ownership advantages delivered by Informix. Both are based on deep research and hard metrics; together, they add up to a business argument that is extremely compelling.

We also have an interview with two of the IBM team members behind the IBM Smart Analytics System. This all-in-one system, which combines software and hardware into an optimized package, is the fruit of years of field experience. Check out our question-and-answer session for an inside look at how the system came together—and what it means for data professionals.

Finally, if you're reading this at the IBM Information On Demand 2010 Global Conference, come see us on the Expo floor!

Thanks for reading,

Cameron Crotty

IDUG
The Worldwide DB2 User Community

IDUG Mentor Program

Bring your new trainee to a conference with you and start them on the path to a bright future.

IDUG Mentors are eligible for an **80% discount** to bring their Mentee to any of IDUG's 2011 conferences!

Get all the details at www.IDUG.org/Mentor

What's on the Horizon?

Two of the biggest events of 2011 are approaching. Through these independent events, attendees get user-delivered, unbiased DB2 content featuring:

- User-driven technical sessions to sharpen skills
- Best practices to drive efficiencies and optimize investments
- Direction on new and existing database and related product issues
- Candid discussions with customers, IBM experts, developers and DB2 consultants

IDUG Conferences

North America | 2 - 6 May, 2011
Anaheim, California

EMEA | 14 - 18 November, 2011
Prague, Czech Republic

MENTORING | CERTIFICATION | DB2 EXPERTS | NETWORKING | EDUCATION

Be a Part of the IDUG Community

Strength in numbers

more than **16,000** members representing **100** countries globally
thousands of DB2 user experiences
more than a **handful** of in-person events year-round

IDUG is an independent, user-driven community that helps DB2 users enhance their professional skills and network, drive efficiencies in their business, and optimize their DB2 technologies.

Experience IDUG year-round.
Sign-up is FREE at www.IDUG.org

Informix 11.70: Panther Pulls Ahead

Latest Informix release bolsters embeddability, enhances application development and security

With a strong reputation for reducing total cost of ownership (TCO) and providing high business value, the new release of Informix continues to build on its core strengths: high availability, performance, ease of use, and enterprise-class embeddability. The Informix 11.70 release (code name "Panther") asserts its dominance in these key areas and extends support for data warehousing, application development, and security. Informix 11.70 is expected to be generally available in October 2010.

High availability

The new release simplifies setup and maintenance for complex replication grids. With Informix 11.70, it's easy to define servers in a replication grid, add or remove servers as necessary, perform DDL operations, and manage tables. The improved Connection Manager streamlines workload management and transparent application failover across the grid. Informix 11.70 also supports backing up to and restoring from cloud storage.

Easy embeddability

Informix 11.70 makes embedding Informix in an application simpler than in previous verisons. It has a deployment utility to customize preconfigured instances for mass implementations, and an embeddability toolkit to help ease embedding and deployment. The self-maintaining functionality in 11.70 includes automatic storage provisioning and autonomic DB Scheduler procedures. *[Editor's note: For more information about embedding Informix, see "Think Inside the Box" from* IBM Data Management, *Issue 3, 2010 at* ***ibm.com****/dmmagazine.]*

Expanded Informix warehousing support

Informix 11.70 addresses the business need for a single database that is equally suited to running online transaction processing (OLTP) and warehousing workloads. Enhancements include automatic fragmenting of date-based data and enhanced handling of data fragments, making it easier to manage data growth. Informix 11.70 also features improved performance of complex queries against star and snowflake schemas, and light scans for much faster sequential scanning of large tables.

Cincom Synchrony Ready for DB2

Customer experience management solution available on DB2 9.7

Software and services provider Cincom Systems, an IBM Premier Business Partner, announced that its Synchrony customer experience management (CEM) solution has been validated with DB2 database software. Synchrony guides employees, step by step, through even the most complex customer-service and up-selling situations. It harnesses silos of knowledge, data, and resources into a single guidance center that instantly presents employees with insights about each customer. Synchrony users can benefit from the combined technologies to connect and interact at a personal level with customers, while gaining the performance, reliability, security, and ease of use provided by DB2.

"The validation of Cincom Synchrony on DB2 was quick and easy, and our engineers were able to quickly migrate our application from Oracle Database," says Randy Saunders, program director of Cincom Customer Experience Management. "DB2 provides a number of advantages to our customers, especially for larger enterprise accounts where performance is critical."

MORE INFORMATION
www.cincom.com

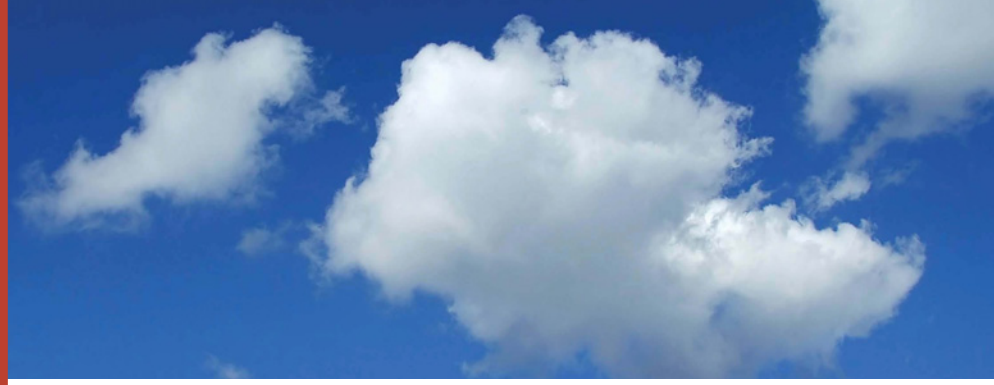

Sky's the Limit: Canonical Brings DB2 to the Cloud

Canonical, the company behind Ubuntu, has launched a virtual appliance with IBM DB2 Express-C software running on the Ubuntu cloud-computing platform in private and public cloud configurations. Canonical also announced that IBM has completed validation of the full version of DB2 software on Ubuntu 10.04 Long Term Support Server Edition.

This announcement is the latest in the growing ecosystem for Ubuntu 10.04, which launched in April with declarations of support from more than 80 organizations. Ubuntu is one of the most popular guest infrastructure layers on cloud services such as Rackspace and Amazon EC2. It is also increasingly deployed as the host cloud infrastructure layer (as Ubuntu Enterprise Cloud) by private organizations and ISPs. DB2 Express-C software will be available however Ubuntu is deployed on a cloud.

DB2 Express-C is a no-charge community edition of DB2 software. It is well suited for small businesses and multi-branch companies, as well as developers and business partners who serve these organizations. DB2 Express-C can be set up quickly, is easy to use, and includes self-managing capabilities. It also embodies all of the core features of more scalable DB2 editions, including pureXML technology for powering Web 2.0 and SOA-based solutions.

MORE INFORMATION
www.canonical.com
ibm.com/developerworks/downloads/im/udbexp/cloud.html

Enhanced application development and security management

Informix 11.70 includes a long list of features that make it easier to develop applications and enhance security:

- Improved compatibility with many popular open-source offerings, such as Drupal, Hibernate, Geronimo, iBATIS, MediaWiki, Tomcat, and XWiki
- Support for new SQL syntax and stored procedure debugging with IBM Data Studio and Microsoft Visual Studio
- Selective row-level auditing to audit only specific row activities for specific tables. DBAs can minimize the performance impact of auditing while simplifying the management of audit records.
- User mapping that enables DBAs to control which externally authenticated users are allowed to connect to Informix and their privileges. Version 11.70 also adds "trusted context" support for multitier Web applications that require secure and auditable database access.

Two recently released studies from analyst firms Forrester and ITG put hard numbers behind the Informix TCO and return on investment (ROI) story. See "Informix TCO: The Facts Are In," on page 38 for more about the results of the studies.

MORE INFORMATION
ibm.com/informix

Sensing Sensitivity: Dataguise Adds DB2 Support

The free trial version of DgDiscover from Dataguise now supports DB2 environments. With DgDiscover, DBAs can audit their database environment and quickly determine where sensitive information resides throughout the enterprise. The company, Dataguise, offers automated and advanced database security solutions to help organizations ensure regulatory compliance and protect against data theft.

MORE INFORMATION
www.dataguise.com

IBM Smart Analytics Optimizer:

BI Gets Fast

In July, IBM announced the IBM Smart Analytics Optimizer for System z. This new product is designed to accelerate long-running and complex analytical queries that scan a large amount of data within DB2 for z/OS.

IBM Smart Analytics Optimizer delivers information faster to decision makers by incorporating breakthrough technologies such as row and columnar stores with highly compressed data and an in-memory, massively parallel architecture. It helps eliminate the costs of traditional performance tuning, such as indices, materialized query tables (MQTs), and query plan tuning that required extensive time and expensive resources. It's also optimized for full table scans (one of the most common—and most time-consuming—business intelligence query access patterns). IBM Smart Analytics Optimizer helps accelerate the delivery of BI results and can even make it possible to execute queries that were previously too resource-intensive for systems to support.

A workload-optimized, deeply integrated DB2 for z/OS component, IBM Smart Analytics Optimizer will be delivered in an appliance-like package and attached to a System z via TCP/IP. It is based on new blade hardware, announced as IBM zEnterprise BladeCenter Extension. The deep integration allows a high degree of transparency for the exploiting applications, because the application programming interface (API) to DB2 for z/OS is unchanged. The integration of both worlds—DB2 for z/OS for online transactional processing (OLTP) and DB2 for z/OS for analytics—allows the building of new applications that integrate BI within operational processes.

MORE INFORMATION
ibm.com/software/data/infosphere/smart-analytics-optimizer-z

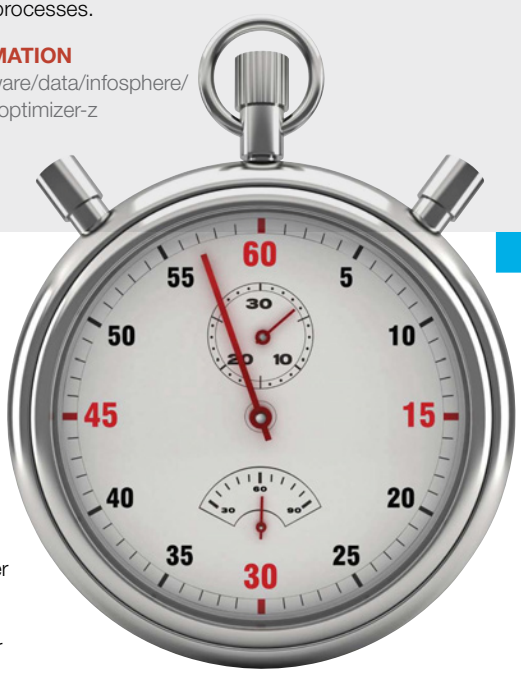

IBM Hits 10 Million Transactions Per Minute

IBM Power System with DB2 and IBM System Storage tops the competition in performance, energy, and cost

In August, an IBM POWER7-based system with DB2 database software and IBM System Storage broke all previous records and topped the 10 million transactions per minute (tpm) mark using the industry-standard TPC performance benchmark, besting all results previously achieved by competitors such as HP and Oracle.

IBM achieved the industry's highest-ever TPC-C (transaction processing) benchmark result using a Power System configuration with DB2, hitting 10,366,254 tpmC.[1] The IBM result delivered the following breakthrough results:

- Performance more than 2.5 times better than HP's best result,[2] 69 percent greater performance per core, and 2.1 times better price/performance
- Performance more than 35 percent better than Oracle's best performance result,[3] 2.7 times better performance per core, and 41 percent better price/performance

In addition, IBM calculations on the IBM configuration show it requires 35 percent less energy per transaction compared to the energy usage data published by Oracle on the Oracle configuration.[1]

The TPC-C benchmark with DB2 9.7 was performed with a cluster of three IBM Power 780 servers, each with 8 processors, 64 cores, and 256 threads, achieving a throughput of 10,366,254 tpmC at US$1.38/tpmC with an availability date of October 13, 2010. Full details of this and other TPC results are available at www.tpc.org.

[1] IBM POWER7 benchmark result: IBM Power 780: 10,366,254 tpmC at US$1.38/tpmC, available October 13, 2010, running on three nodes with a total of 24 processors, 192 cores, and 768 threads. TPC-C results available at www.tpc.org. Energy requirements generated using IBM energy estimate based on IBM calculations using customer-available energy estimation tools for IBM servers, storage energy estimation reports available from IBM Techline services, and published component active power consumption specifications. Energy estimates are not related to, and should not be compared to, official TPC-Energy results.

[2] HP benchmark result: HP Integrity Superdome: 4,092,799 tpmC at US$2.93/tpmC, available August 6, 2007, running on one node with a total of 64 processors, 128 cores, and 256 threads.

[3] Oracle Sun benchmark result: Sun SPARC Enterprise T5440: 7,646,486 tpmC at US$2.36/tpmC, available March 19, 2010, running on 12 nodes with a total of 48 processors, 384 cores, and 3,072 threads. Energy requirements taken from an Oracle-commissioned report located at www.oracle.com/features/strategic-focus-report.pdf. Energy estimates are not related to, and should not be compared to, official TPC-Energy results.

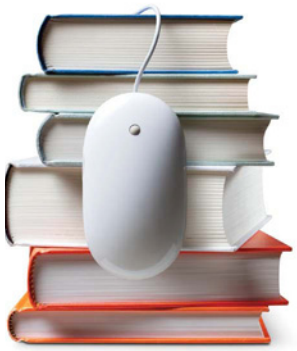

New Course:
DB2 9 for LUW Advanced Database Administration for Experts

This five-day course is available in public classroom, private onsite, and instructor-led online (ILO) formats. It is designed to help students fully use the advanced technical functions and features of DB2 for LUW.

MORE INFORMATION
ibm.com/services/learning

IBM DB2 Workshop Helps Oracle Professionals Extend Their Skills to DB2

IBM is offering a technical workshop to extend database skills to DB2 in a way that is fast, easy, and available at no charge for a limited time. Designed especially for Oracle and other database practitioners, this two-day session covers the new compatibility features in DB2 9.7 and includes certification testing at no charge.

New PL/SQL compatibility features in DB2 make it very easy to extend existing database and PL/SQL knowledge and skills to DB2. After completing this training, you will be better able to support an expanded set of platforms, migrate applications more easily, and take advantage of DB2's automation, scalability, and reliability features.

Workshop dates are coming up in October, November, and December in cities worldwide. For the full agenda and to register, visit **ibm.com**/db2workshop. If your city is not currently listed, use the sign-up link to be notified when a local date is available.

Read All About It:
Save on IBM Press Books

Save 40 percent off the price of IBM Press books and get free ground shipping in the United States by entering coupon code IBMTRAINING during checkout. Select from critical topics facing today's business and IT professionals.

MORE INFORMATION
www.ibmpressbooks.com

Stuart Litel *is president of the International Informix Users Group (IIUG; www.iiug.org/president), CTO of Kazer Technologies (www.kazer.com), an IBM Gold Consultant, member of the IBM Data Champion Inaugural 2008 class, and recipient of the 2008 IBM Data Professional of the Year award.*

Why Do I Use Informix?

The answer isn't what you think

A few weeks ago, an old friend who used to be an Informix user called me on the phone. His company's current database system had crashed a few times recently, and he asked me, "What is the appeal of Informix? Why do you continue to use it after all these years?"

At first I just sat silently, trying to figure out how to answer the question properly. I think he was waiting for me to start naming the numerous features added to the product in recent years. But after a good 30 seconds I replied, "It's the same reason that I've had for the last 20 years. I SLEEP WELL."

It is that simple. Since my friend's company left Informix—their choice, not his—they have used Sybase; DB2 for Linux, UNIX, and Windows; and Oracle. I won't tell you which they currently use, but according to my friend, none of them were as reliable as Informix was in 2001. He has spent many nights and weekends answering his cell phone and participating in numerous conference calls over the past nine years, fighting fires that he never fought during his Informix days.

About two years ago, IBM created a video titled *The Truth Mobile Investigates Informix* (available at www.iiug.tv). An actor goes looking for Informix, finds it is everywhere, and learns that the most important job of an Informix DBA is to leave work early and head to the golf course. Well, we all might not be able to leave work early and head to the golf course, but when was the last time your Informix database crashed and it was the fault of the database software? When was the last time you had to reboot the Informix database on the server?

> "I sleep well using Informix, because it just runs with no need for even planned downtime."

So yes, my old friend. We don't need to reboot the database weekly, monthly, or even annually. It is just always on. I sleep well using Informix, because it just runs with no need for even planned downtime.

After we got off the telephone, I sent my friend a quick follow-up e-mail that included the link to *The Truth Mobile Investigates Informix* video. I also included a few more links, such as how he can read all about moving that Oracle database to Informix (www.redbooks.ibm.com/abstracts/sg247730.html) and of course migrating from Microsoft SQL Server to Informix (www.redbooks.ibm.com/Redbooks.nsf/RedbookAbstracts/sg247847.html), so he can sleep as well as I do and have his database system incur almost no planned or unplanned downtime.

If you want to learn how to sleep well, leave work early, and find out just where else Informix is running, the date has been set for the 2011 IIUG Informix Conference. It will kick off on May 15, 2011, at the Marriott Overland Park in Kansas City, Kansas. Why do we stay at the same location year after year? Because we can be close to the Informix development and support staff, and holding the event at the same place keeps the costs down. For most attendees from the U.S., their total airfare, hotel, and conference registration is less than the cost of just registering for other major conferences. So mark your calendar, check back in this column, or visit www.iiug.org/conf for more details.

13,000+ titles from IBM Press, McGraw-Hill, O'Reilly, Packt Publishing, and more!

YOUR LIBRARY IN THE CLOUD.

More than 15 million IT professionals, developers, web designers, creative professionals and academic learning institutions are using Safari Books Online as their primary resource for research, problem solving, just-in-time learning and certification training.

Try it now for **FREE!**

safaribooksonline.com/openhouse/ibmmag

Safari Books Online

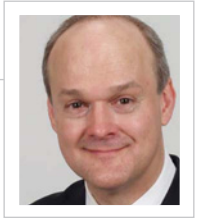

David Beulke *(dave@davebeulke.com) is president of Pragmatic Solutions, Inc. (PSI), a training and consulting company that specializes in designing and improving SQL, application, and system performance on DB2 for Linux, UNIX, and Windows, and z/OS. He has experience in the architecture, design, and performance tuning of large data warehouses and OLTP solutions. He is also a former president of IDUG.*

Is There Anything It Can't Do?

IBM zEnterprise runs your business—any system, any application

IBM recently announced the zEnterprise computer, and it was great to see the mainframe take another giant leap. In my opinion, the zEnterprise is far ahead of any other configuration available, providing more flexibility and capabilities. The new zEnterprise runs mainframe workloads and—with the newly integrated IBM POWER7 and IBM System x blade environments—can run any of the Linux, UNIX, and Windows (LUW) systems, allowing the consolidation of diverse application workloads. The IDUG community in particular will appreciate these capabilities, as they provide the opportunity to improve data integration and decrease data latency between platforms. The zEnterprise can serve as a central location for all administration and performance management of different types of enterprise data or applications.

Just as many organizations have consolidated UNIX systems into the mainframe environment, now they can use the zEnterprise—with the integration of blade environments and the new z/VM 6.1 virtualization hypervisor—to consolidate myriad Windows or x86-based systems. According to IBM, the new capabilities support 100,000 virtual images, providing plenty of room to deliver the mainframe's reliability, availability, and security to other platforms that must receive regular security patches or that struggle with inadequate management-support tools. These consolidation capabilities can help lower overall computing and energy costs. The architecture also helps reduce data copies and remote islands of data by putting it on this central hub within your company, helping to eliminate duplicate data-backup and disaster-recovery planning and procedures.

The zEnterprise system is built for performance and security. It has 96 processors running at 5.2 GHz, optimization facilities for better overall availability, and EAL5 security certification. The new z/OS 1.12 operating system for the zEnterprise also helps improve existing application performance, and more instructions are burned into the chips for faster execution. According to IBM tests, WebSphere and other Java workload performance can be improved by 66–88 percent by fully exploiting the more than 70 new Java hardware instructions. In addition, C++, COBOL, and PL/I have another 73 instructions burned directly into the hardware to boost performance. Simply moving applications to the integrated zEnterprise platform should result in performance improvements.

So the zEnterprise is the future because it runs any application or system regardless of the code base, operating system, or platform. What is also very interesting is that IBM has decided to sell a small, preconfigured, prepackaged version for less than US$1 million. These smaller systems provide a great, low-cost advantage over other distributed platforms. The zEnterprise can quickly deliver return on investment (ROI), making it a strong competitor against other hardware platforms and database environments. Companies can also build their own private cloud-computing environment with a time-tested and reliable zEnterprise system to run it all.

More information will be coming out at the next IDUG conferences: IDUG EMEA (November 8–12, 2010; Vienna, Austria), and the IDUG North America conference (May 2–6, 2011; Anaheim, California). And don't forget to check out www.idug.org and see all the great DB2 regional conferences coming to a city near you.

You Catch More Flies with Honey

At IBM, the success of a massive BI centralization project reveals ways to win the hearts and minds of business execs

Merv Adrian
is principal of IT Market Strategy, a research consultancy that analyzes software trends and advises leading IT firms on market strategy, the competitive landscape, and go-to-market execution issues.

How do you make sure you get the most value from business analytics? Not by empowering the tool users, but by enabling their internal clients, the consumers of the information. But organizational structure, costs, onboarding challenges, and lack of coordination get in the way—and the problems multiply with the size of the organization.

IBM is no different; it's one of the largest organizations in the world, and its business intelligence (BI) processes and practices had become more and more complex, redundant, and expensive over time. In 2008 it set out to change that with a System z–hosted, cloud-based initiative named Blue Insight, and the company is now well on its way toward its goal of saving US$20 million over five years. How it got there is a lesson in effective management of technology, culture, resources—and restraint.

The limits of local control

Like other giant companies, IBM was awash in tools in 2008: Hyperion, Actuate, Arcplan, Brio, and others were in wide use. Lawrence Yarter, chief architect of the IBM Worldwide Business Analytics Center of Competence (BACC), notes that analytics investments were driven by business needs but were not coordinated. "We had a lot of BI strategies; it was essentially 'roll your own' within business unit within geography," he says.

Every internal organization—from product development to sales and marketing, finance, and beyond—produced reports. Each organization had its own BI tools and infrastructure, each was funded with its own budget, and each was driven by its own metrics and priorities. From all appearances, the situation was only going to deteriorate. IBM's many internal CIOs had BI investments on their radar, and most involved new hardware and software acquisitions. As many as 50 new installations—sizable ones—were already on the drawing board when Yarter and his team kicked off the Blue Insight project.

Tip 1: Define clear objectives

Yarter and his team began a six-month process to create a three-year plan with three watchwords: *consolidate, virtualize, automate.* It began with a simple set of objectives—centralized operational support and onboarding processes, common hardware and software,

shared data, and a BACC to provide and promote the value of a service provider model to internal clients.

The overarching objective was to rationalize costs and improve effectiveness yet preserve the individual businesses' ownership of their content, says Yarter. "It's essential in any plan to respect the way business units want to function. They don't want to lose the autonomy that lets them be effective at what they are being measured on: creating revenue."

The planning process set the stage, and Yarter was ultimately able to use the big-picture perspective to identify opportunities for early successes that would help motivate the team and the clients.

> "We agreed to maintain the hardware and put security around it. After that, we made a clear statement: 'Data strategy for process owners is outside our purview. You decide that the data is trustworthy; then we'll connect to it and use it.'"
>
> —**Lawrence Yarter,** *Chief Architect of the IBM Worldwide Business Analytics Center of Competence*

Tip 2: Do your homework

Early results from the analysis revealed complexity that extended beyond product duplication. In six weeks, the portfolio analysis uncovered more than 50 multiproduct, departmental deployments, each with over 100 users, and running across more than 60 data sources—200,000 global named users in all.

Legacy portals had been deployed with custom authorization and authentication methods, each with custom code that needed to be updated and maintained. The resulting redundancy, lack of standards, absence of sharing, and opacity of costs were leading to inefficient systems that lacked agility and relied on multiple nonstrategic skill sets.

These metrics showed the scope of the problem, and they created a basis for documenting improvement as the project progressed. By paying attention and staying flexible, Yarter's team was able to identify and incorporate opportunities to improve on the original plan. For example, the team revised the project goals to include a services delivery model that would simplify deployment to new audiences. This would help reduce costs and lower barriers to entry, which would likely make the needed changes an easier sell.

Tip 3: Support teams, don't dismantle them

Business unit autonomy also guided decisions on staffing issues. If business managers were to have confidence in the outcome, they would want to keep their specialists on their own teams, not surrender them to a centralized authority. "We had many talented people with different skill sets—whatever transformation program emerged, I could not assume that I could move people around," says Yarter. Since the title "analytics IT specialist" did not mean the same thing in every department, the BACC's mandate to be a service provider, training center, and evangelist was crucial.

At every turn, Yarter made it a point to think about how the project would be perceived by local executives. One of the important socialization tactics that Yarter adopted was to be careful not to appear to provide a "solution." Executives are responsible for their own success—they want to piece together their own answers. "If they feel they maintain autonomy," says Yarter, "nobody fights you."

Tip 4: Empower constituents to speed acceptance

In fact, defining standard services—creating a "data deli" that the business units could access—became the point of the spear for Yarter. This facilitated even more recognition of business unit autonomy within the final structure. "We agreed to maintain the hardware and put security around it," says Yarter. "After that, we made a clear statement: 'Data strategy for process owners is outside our purview. You decide that the data is trustworthy; then we'll connect to it and use it.' "

This was a clear signal that Blue Insight was not yet about centralizing governance. Recognizing that the participants "owned" their data was not only useful politically, it also sped up the adoption process by eliminating lengthy negotiations and the creation of data "ownership standards"—at least for the time being. "We will facilitate that at some point," says Yarter. "I can analyze the replication and redundancy, and as more participants join, we can add governance later. Right now, it would affect our ability to go fast."

Tip 5: Reduce costs to create a powerful selling point

A few months later, the project was well under way. A dedicated infrastructure was in place, the cross-functional BACC was staffed, and the converged BI infrastructure—which includes IBM WebSphere Application Server, IBM HTTP Server, DB2, Java Runtime Environment 32-bit and IBM SDK for Java, and IBM Cognos 8 BI—had been built. Defined, standardized processes for onboarding participants were in place. The team set about migrating the identified business units into Blue Insight.

The trade-off for participants is clear: participation drives value. "The named user model, where participants sign up for services, has driven down costs for all," says Yarter. Participants eliminated the need to buy new hardware and software from their

own budgets and, where they were already in place, eliminated maintenance costs for both.

Even with latest upgrades, Blue Insight's cost is about US$10 per head per year. This gets past the difficult part of agreeing on standards and general availability. Business units simply cannot build and deploy their own comparably equipped environment for less than what Blue Insight costs.

Thus far, the project has exceeded expectations. The target was 55,000 users in the first year—a high bar that Yarter and his team beat by 31 percent, getting to 72,000 by the end of 2009. The new target is 120,000 users by the end of 2010 and 200,000 in 2011, at which point growth is expected to level off.

As of midyear 2010, Yarter reports that they have onboarded 113,000 users, a giant piece of the year's 120,000-user goal. "Adopting executives represent all of our geographies and business process areas. We have more than 50 adopters in production, and user communities range from fewer than 50 users to more than 50,000 users," Yarter says. He also expects to meet Blue Insight's 2010 target allocation of the projected US$20 million savings in short order. "More than 80 percent of the projected savings for 2010 has already been booked," says Yarter.

> "It's essential in any plan to respect the way business units want to function. They don't want to lose the autonomy that lets them be effective at what they are being measured on: creating revenue."
>
> —**Lawrence Yarter,** *Chief Architect of the IBM Worldwide Business Analytics Center of Competence*

The softer goals are being met as well. Within IBM, adding users to the existing z-based system is now 13 times faster than creating new deployments on prior IBM System p–based stand-alone platforms. Because of the elasticity of the environment, Yarter says, "I don't have to be concerned with whether resources are there."

Blue Insight is now ubiquitous, always on, and well defined. Because it's positioned as a service, not a solution, data preparation may need to occur—new constituents may want to add content and deliver it via a portal with a specific lens. Blue Insight need not become involved with that; new users understand the services offered. As a result, the time to value is much shorter. In this case, if they build it, they have already come. Lesson learned.

Start Faster, Reach Further

IBM Smart Analytics System and Powerful Intel Xeon Processor–Based IBM System x Servers Deliver Business Insights

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

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

IBM and Intel have been working for years to deliver optimized computing solutions and to bring enterprise-class technologies into industry-standard systems.

The IBM® Smart Analytics System family combines hardware, software, and services into systems that are integrated and optimized, can be installed in days instead of months, and can be quickly scaled out to support changing business needs. Available in a range of configurations and capabilities, including several models powered by the latest, most powerful Intel® Xeon® processor–based servers, the IBM Smart Analytics System family is designed to deliver cost-effective, enterprise-class analytics—without the wait.

Smarter Systems Support Smart Decision Making

IBM Smart Analytics System integrates leading data warehouse, business reporting, and analytics capabilities, enabling you to rapidly deploy comprehensive business analytics. It is built on a powerful IBM server, storage, and data management platform that is optimized for analytic workloads.

The Smart Analytics System is designed to deliver insights for capturing new opportunities and gaining competitive advantage. As a preconfigured solution, it facilitates rapid deployment, enabling organizations to get to work in days, not months. It can help cut costs by reducing the need for extra staff and in-depth expertise for implementation, tuning, and maintenance. Working with a single vendor also helps simplify acquisition, implementation, and support processes.

The Smart Analytics System provides a comprehensive, fully integrated BI solution, with a scalable warehouse foundation and broad analytic capabilities from the industry-leading IBM analytics and BI software portfolio, including multidimensional Cubing Services, data mining and text analytics, dashboards, and reporting (see Figure 1). Part of the IBM smarter systems for a smarter planet approach to optimizing hardware, software, and services, the Smart Analytics System is based on IBM InfoSphere™ Warehouse and is powered by IBM DB2® database software. It takes advantage of DB2's advanced performance features, such as:

- Deep data compression, which helps reduce the cost and footprint of the data warehouse
- Workload management, which enables you to manage, enforce, and monitor workloads and their service levels for users across your warehouse

The Smart Analytics System also integrates easily with existing data sources to provide a single version of the truth and to deliver powerful insights across your business.

The Power of Intel Xeon Processors

Today's markets may shift on a moment's notice, so your business must be agile enough to respond quickly—and accurately—if an opportunity presents itself. That means you need systems that respond fast and deliver insight and analysis on demand.

IBM Smart Analytics Systems draw on the power of the latest Intel Xeon processors to deliver high levels of enterprise computing capability.

IBM Smart Analytics System Integrated Stack

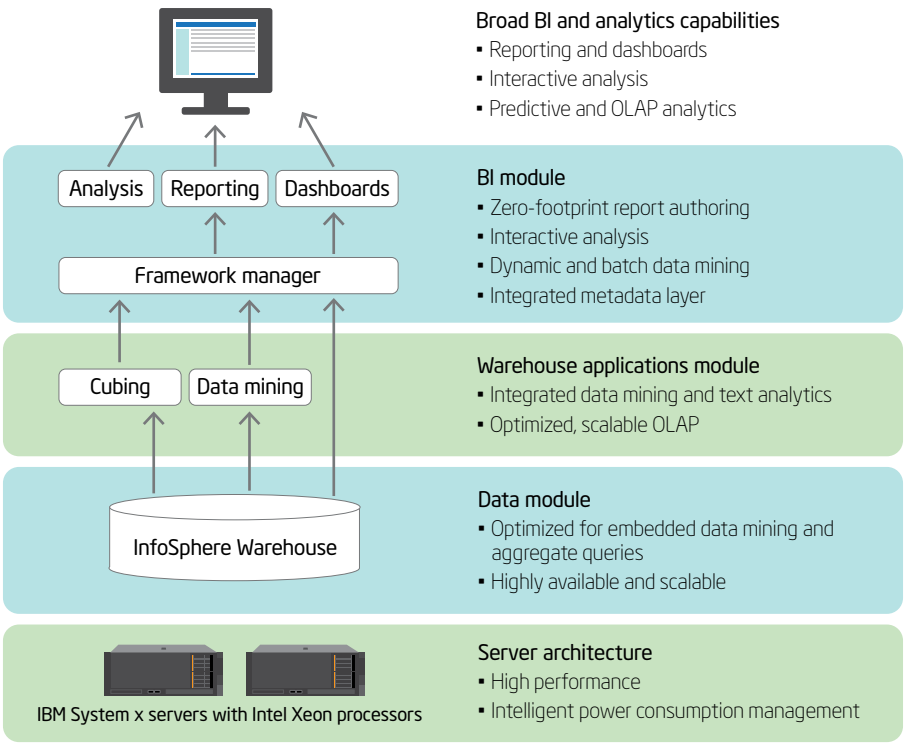

Figure 1: The IBM Smart Analytics System includes BI and warehousing capabilities, built on a high-performance server architecture.

Intel Xeon Processor 7500 Series

The Intel Xeon processor 7500 series delivers the largest performance leap ever for an Intel processor generation (see Figure 2), with an average of 3x higher performance across a wide range of industry benchmarks[1] compared with its predecessor. It provides up to 20x better performance per server versus single-core servers,[2] plus massive increases in scalability and more than 20 new, mainframe-inspired reliability, availability, and serviceability (RAS) features to help protect data and improve system resilience.

Large on-die cache (up to 24 MB per socket), Intel® Hyper-Threading Technology, and Intel® Turbo Boost Technology[3] deliver exceptional performance for complex BI workloads. Further, Intel® Advanced Reliability technology helps ensure data integrity and high availability to keep critical BI systems up and running. Intel and IBM have demonstrated scalable performance and consistent query response times for data sets as large as 10 TB on Intel Xeon processor 7500 series–based servers.[4]

Intel Xeon Processor 5600 Series

The Intel Xeon processor 5600 series for two-socket servers automatically regulates power consumption and intelligently adjusts server performance according to application needs. It delivers up to 15x performance per server over single-core servers,[5] enabling up to 95 percent lower energy costs[6] and helping you maximize performance while making the most of limited data center space, power, and cooling resources.

IBM Smart Analytics System Family

IBM Smart Analytics System 1050

Cost-effective support for smaller deployments

- IBM System x®3500 M3
- Intel Xeon Processor 5600 series
- Novell* SUSE Linux* Enterprise 11 and Microsoft* Windows Server* 2008
- IBM System Storage® DS3500
- IBM InfoSphere Warehouse Departmental Edition or InfoSphere Warehouse Departmental Base Edition
- IBM Cognos® Reporting and Query

IBM Smart Analytics System 2050

Scalability for departmental enterprise deployments and midsize businesses

- IBM System x3850 X5
- Intel Xeon processor 7500 series
- Novell SUSE Linux Enterprise 11 and Microsoft Windows Server 2008
- IBM System Storage DS3500
- IBM InfoSphere Warehouse Departmental Edition or InfoSphere Warehouse Departmental Base Edition
- IBM Cognos Reporting and Query

IBM Smart Analytics System 5600

Modular system for companies that need powerful analytics capabilities and growth flexibility

- IBM System x3650 M3
- Intel Xeon processor 5600 series
- Linux
- IBM System Storage DS3500; Fusion-io SSD (optional)
- IBM InfoSphere Warehouse Enterprise Edition
- IBM Cognos 8 BI

Intel Xeon Processors

Intel Xeon Processor 7500 Series

Exceptional scalable performance with advanced reliability

- Up to eight cores, 16 threads, 24 MB of on-die cache
- Intel Advanced Reliability technology

Intel Xeon Processor 5600 Series

The next generation of intelligent server processors

- Automatically regulates power consumption
- Adjusts server performance according to application needs

IBM Server Family Architectures: A Solid Foundation

IBM Smart Analytics Systems harness the power of the Intel Xeon processors through two IBM server architectures, the M3 and eX5 servers. The IBM eX5 server architecture is ideal for midmarket and developmental workloads, delivering performance and per-core energy efficiency for smaller deployments, while the IBM M3 server architecture delivers modular scalability with value and investment protection for enterprise-class environments.

For even more complex workloads, you can add optional modules with Fusion-io* solid-state device (SSD) technology to help improve response time and boost efficiency. The advanced design of the Smart Analytics System puts this high-performance technology to work where it is the most cost-effective and will have the greatest positive impact on system performance.

In an analytics system, for example, writing data to spinning disks can cause big performance bottlenecks. When equipped with solid-state storage, the Smart Analytics System captures highly disruptive temporary write I/Os—which are common in analytic workloads—to the solid-state storage first. Moving even a small percentage of data writes from spinning disks to solid-state storage can significantly increase I/O and cut query response, giving you the capability to expand hardware capacity on existing systems as your BI needs grow.

IBM Smart Analytics Systems based on the eX5 architecture offer flexible configurations with affordable starting points and easy expansion to handle the most demanding analytics and BI initiatives. You can start with a two-socket or four-socket server with just a single processor, and then scale that system as needed.

Reduce Time to Value through Integration

The modular, easily expandable nature of the Smart Analytics System family showcases the extensive integration work done by IBM and Intel. Every Smart Analytics System is a turnkey analytics solution that combines hardware, software, and services. All of the components have been configured and optimized to work smoothly together as a seamless unit, and the entire system has been tuned to maximize performance for analytic workloads.

As integrated solutions, Smart Analytics Systems are application-ready and data-ready almost from the moment they arrive. Organizations typically have Smart Analytics Systems up, running, and delivering analysis in a few days, compared to the weeks or months necessary for solutions built from scratch. That simplicity helps reduce the staff time and expertise that you need to reserve for implementation, which means that your senior IT staff can focus on delivering analytics to every level of your organization instead of wading through time-consuming setup and configuration.

The Biggest Performance Leap Ever for an Intel Processor Generation

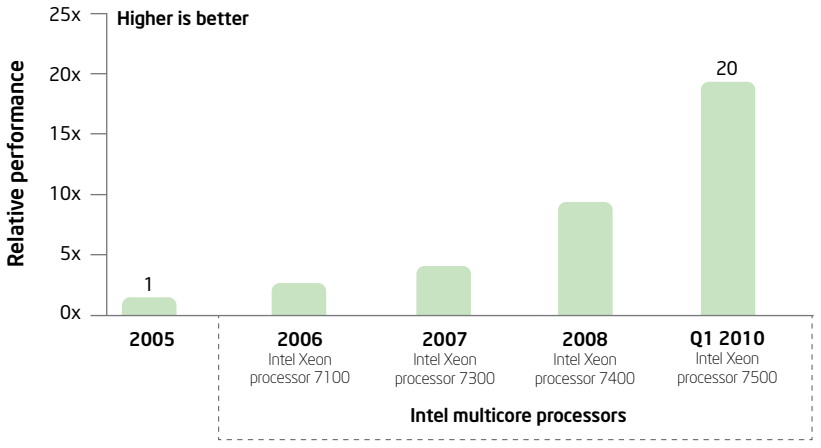

Source: Intel Internal Measurements, January 15, 2010. Results calculated based on geometric mean of five enterprise benchmark scores, including server-side Java*, integer throughput, floating-point throughput, ERP, and OLTP.

Figure 2: The Intel Xeon processor 7500 series delivers the biggest leap ever in generation-to-generation performance for an Intel processor.

Once the systems are in place, the extensive integration helps minimize conflicts and confusion. You can rely on a unified, certified solution stack and enjoy the benefits of a single point of service and support—no more chasing problems through multiple vendors.

Get Results Now

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

Learn More

More about smarter systems from IBM: www.ibm.com/systems/smarter

More about IBM Smart Analytics System: www.ibm.com/smart-analytics-system

More about the Intel Xeon processor 7500 series: www.intel.com/itcenter/products/xeon/7500/index.htm

More about the Intel Xeon processor 5600 series: www.intel.com/itcenter/products/xeon/5600/index.htm

---

[1] Average of 3x performance claim based on geometric mean of four industry-standard, common enterprise benchmarks (SPECjbb*2005, SPECint*_rate_base2006, SPECfp*_rate_base2006, and TPC Benchmark* E) comparing best published/submitted results on four-socket (4S) Intel® Xeon® processor X7560–based server platform to best published 4S Intel Xeon processor X7460–based server platform as of March 26, 2010.

[2] Claim: "Up to 20x performance per server" Disclaimer: Intel performance comparison using SPECjbb2005* business operations per second between five-year-old single-core Intel® Xeon® processor 3.33GHz–based servers and one new Intel Xeon processor X7560–based server. Performance tests and ratings are measured using specific computer systems and/or components and reflect the approximate performance of Intel products as measured by those tests. Any difference in system hardware or software design or configuration may affect actual performance. Buyers should consult other sources of information to evaluate the performance of systems or components they are considering purchasing. For more information, visit www.intel.com/performance/server.

[3] Intel® Hyper-Threading Technology requires a computer system with a processor supporting HT Technology and an HT Technology–enabled chipset, BIOS, and operating system. Performance will vary depending on the specific hardware and software you use. For more information, including details on which processors support HT Technology, see www.intel.com/technology/platform-technology/hyper-threading.

Intel® Turbo Boost Technology requires a platform with a processor with Intel Turbo Boost Technology capability. Intel Turbo Boost Technology performance varies depending on hardware, software, and overall system configuration. Check with your platform manufacturer on whether your system delivers Intel Turbo Boost Technology. For more information, see www.intel.com/technology/turboboost.

[4] Source: IBM internal measurements as of March 2010. All of the benchmark results presented are from a database with schema extended from the TPoX 2.0 schema. The benchmark was run with a series of 16 complex analytical SQL/XML queries. The resulting database size was built from approximately 10 TB data. Platform: three servers each with four Intel® Xeon® processor X7500, 2.27 GHz; 128 GB, running DB2 9.7 on Linux RHEL 5.4 64-bit with DB2 Compression and STMM. DS8700 configured with 384 300 GB 15K RPM Fibre Channel drives.

[5] Claim: "Up to 15x performance per server" Disclaimer: Intel performance comparison using SPECjbb2005* business operations per second between four-year-old single-core Intel® Xeon® processor 3.8 GHz with 2M cache–based servers and one new Intel Xeon processor X5670–based server. Performance tests and ratings are measured using specific computer systems and/or components and reflect the approximate performance of Intel products as measured by those tests. Any difference in system hardware or software design or configuration may affect actual performance. Buyers should consult other sources of information to evaluate the performance of systems or components they are considering purchasing. For more information, visit www.intel.com/performance/server.

- Baseline platform: Intel server platform with two 64-bit Intel Xeon processor 3.80 GHz with 2M L2 Cache, 800 FSB, 8x1 GB DDR2-400 memory, 1 hard drive, 1 power supply, Microsoft* Windows Server* 2003 Ent. SP1, Oracle* JRockit* build P27.4.0-windows-x86_64 run with 2 JVM instances.
- New platform: Intel server platform with two six-core Intel Xeon processor X5670, 2.93 GHz, 12 MB L3 cache, 6.4QPI, 12 GB memory (6x2 GB DDR3-1333), 1 hard drive, 1 power supply, Microsoft Windows Server 2008 64 bit SP2, Oracle JRockit build P28.0.0-29 run with 2 JVM instances.

[6] Claim: "Up to 95 percent lower energy costs" Disclaimer: Intel comparison replacing 15 four-year-old single-core Intel® Xeon® processor 3.8 GHz with 2M cache–based servers with one new Intel Xeon processor X5670–based server. ROI claims and costs have been estimated based on internal Intel analysis and are provided for informational purposes only. Performance tests and ratings are measured using specific computer systems and/or components and reflect the approximate performance of Intel products as measured by those tests. Any difference in system hardware or software design or configuration may affect actual performance. Buyers should consult other sources of information to evaluate the performance of systems or components they are considering purchasing. For more information on performance tests and on the performance of Intel products, visit Intel Performance Benchmark Limitations.

- Baseline platform: Intel server platform with two 64-bit Intel Xeon processor 3.80 GHz with 2M L2 Cache, 800 FSB, 8x1 GB DDR2-400 memory, 1 hard drive, 1 power supply, Microsoft Windows Server 2003 Ent. SP1, Oracle JRockit build P27.4.0-windows-x86_64 run with 2 JVM instances.
- New platform: Intel server platform with two six-core Intel Xeon processor X5670, 2.93 GHz, 12 MB L3 cache, 6.4QPI, 12 GB memory (6x2 GB DDR3-1333), 1 hard drive, 1 power supply, Microsoft Windows Server 2008 64 bit SP2, Oracle JRockit build P28.0.0-29 run with 2 JVM instances.

Any performance data contained herein was determined in a controlled environment. Therefore, the results obtained in other operating environments might vary significantly. Users of this document should verify the applicable data for their specific environment.

Migrate from Oracle or Sybase

(INSERT VENDOR NAME HERE)

Rewriting code for a new platform? That's old news. Time to reset your preconceptions about migrating from Oracle or Sybase to DB2 9.7. We think you'll be pleasantly surprised.

By Greg Thomas and Jeff Jones

to DB2 in weeks.

22 Migrating from Oracle to DB2
25 Oracle and DB2 Terminology
26 Migrating from Sybase ASE to DB2
26 More Migration Resources

Ask most DBAs how they feel about migrating applications, and you'll probably learn some new and colorful adjectives. After all, what could be more fun than opening up a business-critical app and rewriting legacy code so that it works with a new platform? You could follow up with a few joy-filled months of reworking complex, long-standing schemas so that they make sense to a new RDBMS. And then, how about a nice, relaxing trip to the dentist for a root canal?

But wait just a second. If you're talking about bringing applications from Oracle or Sybase environments to IBM DB2 for Linux, UNIX, and Windows (LUW), it's time to revise your expectations. DB2 9.7 is compatible with applications and databases developed for Oracle and Sybase at a level that you might not have thought possible.

The release of DB2 9.7 and the introduction of a new DB2 SQL Skin (available in DB2 9.7 Fix Pack 2) for applications compatible with Sybase Adaptive Server Enterprise (ASE) have greatly simplified migration to DB2 from both Oracle and Sybase environments. Migrations can now be accomplished faster, more cost-effectively, and often with little or no application rewriting. It wouldn't be fair to say that migrating to DB2 from these two platforms is effortless—but it is fair to say that most Oracle and Sybase applications can now be moved to DB2 9.7 with only minor modifications.

Skeptical? Sure you are. So let's take a closer look at the tools, techniques, and timing of migrating from Oracle and Sybase environments to DB2 9.7. We'll discuss the key steps in planning and conducting a migration, and then we'll examine potential hurdles and ways to eliminate them.

Migrating from Oracle to DB2

Previously, moving from Oracle to DB2 meant that the application needed to be ported, or rewritten to account for proprietary SQL, different locking mechanisms, and client interfaces that differed in both semantics and syntax. In DB2 9.7, IBM and EnterpriseDB (an IBM Business Partner) combined their expertise and developed a different approach, offering native support for many commonly used Oracle features, including:

- **Expanded SQL support** that incorporates popular keywords and semantics from the Oracle SQL dialect
- **Extended data-type support**, including support for the most common nonstandard data types used by Oracle Database
- **Added PL/SQL support** to help avoid the translation of procedural language (PL) code
- **Support for many built-in Oracle Database packages** to help simplify application migration; DB2 provides many of these packages, including DBMS_OUTPUT, UTL_FILE, DBMS_ALERT, DBMS_PIPE, DBMS_JOB, DBMS_LOB, DBMS_SQL, DBMS_UTILITY, UTL_MAIL, and UTL_SMTP
- **Scripting support** offers a SQL*Plus-compatible command-line processor called CLP Plus to help connect databases as well as to define, edit, and run statements, scripts, and commands
- **Enhanced concurrency model** enables administrators to adopt the same concurrency behavior as Oracle Database

These capabilities can dramatically reduce the need for manual modifications of application code. To date, IBM has worked with dozens of organizations to migrate millions of lines of code, and IBM DB2 for LUW SQL Architect Serge Rielau is one of the folks who spends much of his time on the migration front lines. "As we work with customers, we are finding that usually between 90 and 99 percent of SQL and PL/SQL statements need no changes," says Rielau. "On average, the typical organization needs to modify only 2 percent of its code."

Oracle migration step 1: Planning and assessment

An Oracle migration usually follows a five-step process: migration assessment, database object migration, application migration, solution deployment, and skills transfer and DB2 ramp-up. For now, let's focus on the first three steps.

The first step is to assess your environment and applications to identify potential roadblocks. If you're looking for hard data on how well your Oracle databases and applications will play with DB2, you'll want to download the free IBM Migration Enablement Evaluation Tool (MEET). The tool analyzes Oracle database objects and procedures, and quickly identifies those that use features not supported in DB2 9.7. The tool delivers an HTML report that identifies unsupported code, lists details and source code line numbers, and provides summary statistics (see Figure 1).

Evaluating and assessing your infrastructure and database requirements will help you understand the scope of your migration project. In step 3, we'll look at some of the most common challenges reported by organizations that have already completed migrations.

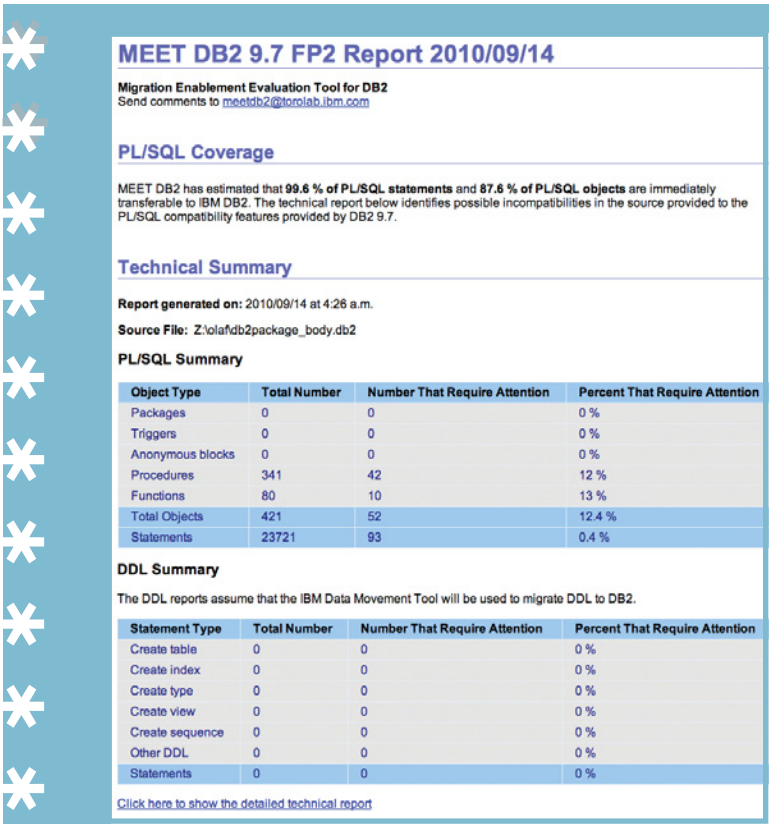

MEET DB2 9.7 FP2 Report 2010/09/14

Migration Enablement Evaluation Tool for DB2
Send comments to meetdb2@torolab.ibm.com

PL/SQL Coverage

MEET DB2 has estimated that **99.6 % of PL/SQL statements** and **87.6 % of PL/SQL objects** are immediately transferable to IBM DB2. The technical report below identifies possible incompatibilities in the source provided to the PL/SQL compatibility features provided by DB2 9.7.

Technical Summary

Report generated on: 2010/09/14 at 4:26 a.m.

Source File: Z:\olaf\db2package_body.db2

PL/SQL Summary

Object Type	Total Number	Number That Require Attention	Percent That Require Attention
Packages	0	0	0 %
Triggers	0	0	0 %
Anonymous blocks	0	0	0 %
Procedures	341	42	12 %
Functions	80	10	13 %
Total Objects	421	52	12.4 %
Statements	23721	93	0.4 %

DDL Summary

The DDL reports assume that the IBM Data Movement Tool will be used to migrate DDL to DB2.

Statement Type	Total Number	Number That Require Attention	Percent That Require Attention
Create table	0	0	0 %
Create index	0	0	0 %
Create type	0	0	0 %
Create view	0	0	0 %
Create sequence	0	0	0 %
Other DDL	0	0	0 %
Statements	0	0	0 %

Click here to show the detailed technical report

Figure 1: *IBM MEET scans and identifies Oracle database objects and statements that will run on DB2 without modification, noting any possible incompatibilities.*

Oracle migration step 2: Database object migration

The next step is to migrate Oracle database objects to DB2. The native support for Oracle PL/SQL and Oracle SQL provided by DB2 9.7 greatly simplifies this process, as does another resource now freely available: the IBM Data Movement Tool. This tool automatically copies Oracle database objects—including tables, packages, or entire schemas—to DB2.

To use the IBM Data Movement Tool, start it up and connect it to both your Oracle and DB2 databases. Once you are fully connected, you can extract the Data Definition Language (DDL) only, or both the DDL and the data. You'll likely want to use the tool's interactive deploy mode, which displays a navigation tree with all the objects extracted from the Oracle database. As the tool copies objects to DB2, it records its progress. If an object does not deploy correctly, the tool highlights the object in the navigation tree; click on the object, and the tool reveals the DDL and the error DB2 encountered, so that you can fix the definition and redeploy. With the tool, data migration is a relatively straightforward and automatic process that should be nearly problem-free.

Oracle migration step 3: Application migration

Generally speaking, DB2 9.7 significantly reduces the need for manual modifications of application code. However, you can't quite kick back with a cup of coffee and let DB2 do all the work. Here are a few of the places where you'll probably need to intervene, starting with the most simple and working up to the more complex.

PHP/Perl

Migrating PHP or Perl applications from Oracle to DB2 involves just one alteration: changing the library call from an Oracle library to a DB2 library. Modifying that call should require only a global replacement of the call names in the code using a text editor. The SQL inside the PHP remains unchanged.

Java

Converting Java code is similarly straightforward. The application programming interface (API) itself is well defined and database independent—the database connection logic is encapsulated in standard J2EE DataSource objects. The Oracle- or DB2-specific terms, such as user name or database name, are then configured declaratively within the application.

Oracle partitioning	DB2 data organization	Oracle 10g syntax	DB2 9 syntax
No equivalent	Round robin	None	Default: occurs automatically on single partition database
Range partitioning	Table partitioning	PARTITION BY RANGE	PARTITION BY RANGE
Hash partitioning	Database partitioning	PARTITION BY HASH	DISTRIBUTE BY HASH
List partitioning	Table partitioning with generated column	PARTITION BY LIST	PARTITION BY RANGE
Composite partitioning: hash-range hash-list	Combination of: database partitioning, table partitioning, multi-dimensional clustering	PARTITION BY RANGE, SUBPARTITION BY HASH, SUBPARTITION BY LIST	DISTRIBUTE BY HASH, PARTITION BY RANGE, ORGANIZE BY DIMENSIONS
No equivalent	Multidimensional clustering	None	ORGANIZE BY DIMENSIONS

Figure 2: *Configuring partitioning in DB2 may require code changes to accommodate differences in syntax.*

Converting code requires changing only the Java source code to the appropriate API driver (JDBC or SQLJ), the database connect string, and any incompatible SQL statement. DB2 9.7 also enables you to use Hibernate (an open-source persistence and query service for Java), which is now as easy to use with DB2 as it is with Oracle.[1]

Oracle Calling Interface

Oracle Calling Interface (OCI) is one of the many programming interfaces used by C/C++ developers to interact with an Oracle database. DB2 9.7 Fix Pack 1 introduced the DB2 Call Interface (DB2CI), which provides compatibility for the OCI—developers will have a familiar interface for both environments.

Oracle Forms

Oracle Forms is a legacy software product used to create data-entry systems for the database. Some organizations have hundreds of Oracle Forms screens, which constitute all or part of an application.

IBM has partnered with Realease to offer Oracle Forms–to–Java conversion capabilities. Realease provides tooling that translates Oracle Forms to Java, preserving the look and feel of the original GUI. In many cases, the translation work can be accomplished in one week rather than months.

Triggers

DB2 does not (yet) allow you to perform updates to tables from within a BEFORE trigger. In most cases, you can use AFTER triggers to perform these actions. Also, DB2 does not yet allow trigger actions to be combined. So if you have a PL/SQL multi-action trigger, you'll need to copy it into separate DB2 SQL PL triggers, using a Boolean variable for the predicates.

Partition handling

DB2 can be configured to organize data in several ways, including table partitioning, database partitioning, multi-dimensional clustering, or a combination of these organization schemes. If you used a form of partitioning with Oracle, you may need to update your code to accommodate the differences in syntax (see Figure 2).[2]

Third-party dependencies

Some applications have third-party software dependencies written into their code. Those dependencies could be difficult to identify if existing systems have been functioning successfully for years.

The migration assessment process can help you uncover these third-party dependencies and enable you to determine which dependencies need to be altered or eliminated. For some organizations, recoding might be required if the third-party software is no longer available.

[1] For more information on using Hibernate with DB2, read *Using Hibernate to Persist Your Java Objects to IBM DB2 Universal Database,* **ibm.com**/developerworks/data/library/techarticle/0306bhogal/0306bhogal.html.

[2] For specific syntax changes, see *Oracle to DB2 Conversion Guide: Compatibility Made Easy* at www.redbooks.ibm.com/abstracts/sg247736.html

Oracle and DB2 Terminology

Many commonly used concepts and much of the terminology varies between the Oracle and DB2 worlds. We asked Ted Alexander of The Praxium Group to share a few of his hot-button terms, as found in the IBM Redbook *Oracle to DB2 Conversion Guide: Compatibility Made Easy*.

Oracle term	DB2 term	Description
Oracle Parallel	DB2 Enterprise DPF	Support server partitioning
Oracle Gateway	DB2 Connect	DRDA access to hosts
Instance	Instance	Processes and shared memory; in DB2, it also includes a permanent directory structure. An instance is usually created at install time (or can be later) and must exist before a database can be created. A DB2 instance is also known as the *database manager* (DBM). A DB2 instance can have multiple databases, but an Oracle instance can have only one database.
Database	Database	Physical structure containing data: in Oracle, multiple instances can use the same database, and an instance can connect to one and only one database; in DB2, multiple databases can be created and used concurrently in the same instance.
Control files and .ora files	DBM and database configuration files	In Oracle, these are files that name the locations of files making up the database and provide configuration values. In DB2, each instance/DBM and each database has its own set of configuration parameters stored in a binary file. There are also other internal files and directories; none are manually edited.
Table spaces	Table spaces	Contains actual database data
Data files	Containers	Entities inside the table spaces
Segments	Objects	Entities inside the containers/data files
Extents	Extents	Entities inside the objects/segments
Data blocks	Pages	Smallest storage entity in the storage model
Database link	Federated system	In Oracle, this is an object that describes a path from one database to another. In DB2, a federated system is used. One database is chosen as the federated database and within it wrappers, servers, nicknames, and other optional objects are created to define how to access the other databases (including Oracle databases) and objects in them. Once an application is connected to the federated database, it can access all authorized objects in the federated system.
Clusters	N/A	A data structure that allows related data to be stored together on disk. This data can be table or hash clusters. The closest facility to this in DB2 is a clustering index, which causes rows inserted into a table to be placed physically close to the rows for which the key values of this index are in the same range.
Data dictionary	System catalog	Metadata of the database
N/A	SMS	System-managed table space

Migrating from Sybase ASE to DB2

Spurred by frustration with increasing maintenance and administration costs and a lack of a solid road map from Sybase, many Sybase ASE users are looking for an alternative.

Until recently, proprietary syntax and functionality in Sybase ASE made migrations difficult; you had to invest significant time and money to rewrite application code and then run test cases for the application. With the introduction of DB2 SQL Skin for applications compatible with Sybase ASE, DBAs now can accomplish this migration while minimizing—and in some cases completely avoiding—application coding changes as well as subsequent test-case changes.

Developed jointly by IBM and IBM Business Partner ANTs Software, DB2 SQL Skin allows Transact SQL (T-SQL) code—including queries, functions, triggers, and stored procedures—from Sybase ASE to interact transparently with DB2 with little or no rewriting, recompiling, or re-linking. DB2 SQL Skin provides the features, functions, and data formats required for code from Sybase ASE to run natively on DB2. The only necessary change is resetting connection parameters so the application connects to the DB2 server instead of the Sybase server.

Clearing the Path to DB2: More Migration Resources

Maybe you can't drop off your fleet of Oracle and Sybase applications on the DB2 server and have them magically get themselves up and running while you make coffee (well, not yet, anyway). But the native compatibility features built into DB2 9.7 really do handle most of the heavy lifting, and lots of tools and resources are available to smooth out any remaining rough patches.

For information about IBM DB2 SQL Skin feature for applications compatible with Sybase ASE, visit **ibm.com**/software/data/db2/linux-unix-windows/editions-features-sql-sybase.html or www.ants.com.

For more information about the steps for migrating from Oracle to DB2, read the IBM Redbook *Oracle to DB2 Conversion Guide: Compatibility Made Easy* at www.redbooks.ibm.com/abstracts/sg247736.html.

Check out the following sites for more on some of the IBM migration tools:

IBM Data Movement Tool: **ibm.com**/developerworks/data/library/techarticle/dm-0906datamovement

IBM Migration Enablement Evaluation Tool (MEET): https://www14.software.ibm.com/webapp/iwm/web/preLogin.do?lang=en_US&source=swg-meetdb2

Realease: www.realeasellc.com

Reducing costs through database migration: **ibm.com**/db2/lowerdatabasecosts/migrate-to-db2.html

For more database migration answers from IBM, contact:
askdata@ca.ibm.com or db2mig@us.ibm.com

Planning

As with an Oracle migration project, you should begin the migration process by assessing your environment. The assessments will help you understand the work involved and help you decide whether assistance from IBM or an IBM Business Partner would enable you to complete the migration easier, faster, and with better efficiency. Assistance will be particularly useful for organizations that have limited IT resources or numerous applications whose interactions must be coordinated.

During the assessment process, you may opt to modify applications so they run natively on DB2 or decide to use DB2 SQL Skin to run Sybase applications as-is. Because DB2 SQL Skin can simplify the migration process, you could use it during the initial migration and then rewrite some applications later, when your developers become more comfortable with the DB2 SQL dialect.

Execution

Sybase applications can use the same APIs on DB2. With DB2 SQL Skin, many of the behaviors that Sybase applications expect are now provided by DB2. As a result, the application can still run and return the data in the existing format, but it will work with DB2 even though the format is different within the database. All of this data format handling and T-SQL function handling is transparent to the application, which still thinks it is talking to Sybase ASE.

Greg Thomas is a technology and business writer in the San Francisco Bay Area. Over the past 18 years, he has contributed to numerous publications for leading high-tech companies.

Jeff Jones has IT analyst relations responsibility for the Data Management area of IBM Information Management Software and assists in the production of IBM Data Management *magazine, having been with the magazine and its predecessors since 1997.*

What's New: DB2 10 for z/OS

Dramatic CPU savings and data availability enhancements, check. But have you seen the new temporal data support?

By Terrie Jacopi

The latest release of DB2 for z/OS delivers the most aggressive package of CPU savings and performance improvements from IBM in more than 20 years. Not enough? Let's add a tall stack of new ways to keep your data continuously available, no matter what. Still want more? We've also got a big set of enhanced application-development capabilities.

Just listing every improvement in DB2 10 for z/OS would take more pages than we have. So in this article, we'll focus on the capabilities that tie to your top concerns: saving money and strengthening your data infrastructure. We'll also dig into a sharp new feature that will make it possible for you to manipulate time-based data faster and easier than ever before.

Out-of-the-box savings

IBM internal testing and reports from the beta program are showing some impressive numbers: Simply migrating to DB2 10 and rebinding will deliver significant savings for transactions, queries, and batch processes. Most customers can achieve CPU savings of 5 to 10 percent for traditional transaction workloads and up to 20 percent savings for nontraditional workloads, such as those using SQL stored procedures and some distributed relational database architecture (DRDA)–based workloads. DB2 10 reduces CPU usage by optimizing processor times and memory access, and by taking greater advantage of the latest processor improvements, larger amounts of memory, solid-state drives, and z/OS enhancements.

Of course, your results may vary. For example, customers that currently have virtual storage constraints or latching issues are likely to see greater improvements. Concurrent sequential single inserts that have many indexes can be reduced by 5 to 40 percent. Customers moving from DB2 9 can expect a small (up to 7 percent) reduction in CPU usage for utilities, while customers moving from version 8 will see reductions in CPU usage as high

as 20 percent. DB2 10 also includes a number of new database design capabilities, such as inline large objects (LOBs), hash access, and RELEASE(DEALLOCATE), that can improve performance and cut CPU usage.

More scalable and more available

DB2 10 sports an array of features designed to take data scalability and availability to new heights.

Improved scalability

DB2 10 substantially increases the amount of virtual storage that can be addressed by moving most DB2 working memory from 32-bit to 64-bit. One of the benefits of this increase is that a single DB2 10 subsystem supports 5 to 10 times more concurrent users than previous releases supported—as many as 20,000 concurrent threads. This greatly improves the vertical scalability of your DB2 subsystem.

Continuous availability enhancements

Online schema enhancements allow you to make changes to database objects (indexes and table spaces) while maximizing the availability of the altered objects. ALTER statement enhancements let you change indexes and table spaces without needing to unload the data, drop and re-create the objects, regenerate all of the security authorizations, re-create the views, and reload the data. The changes materialize when the altered objects are reorganized.

DB2 10 can now automatically reorganize disjoint partition ranges of a partitioned table space. This new feature, along with improvements to SWITCH phase performance and diagnostics, increases the usability and performance of online reorganization. DB2 10 also removes restrictions on the online reorganization of base table spaces that use LOB columns.

The DB2 catalog has been restructured to reduce lock contention by removing all links in the catalog and directory. New row-level locking functionality improves the lock avoidance techniques of DB2, and improves concurrency by holding acquired locks for less time and preventing writers from blocking the readers of data.

DB2 10 lets you access currently committed data to dramatically minimize transaction suspension. Now, a read transaction can access the currently committed and consistent image of rows that are incompatibly locked by write transactions without being blocked. This concurrency control can greatly reduce timeout situations between readers and writers accessing the same data row. DB2 10 also provides flexibility and increased performance to applications that only require available and committed data to be returned from DB2 tables.

Getting to DB2 10

DB2 10 supports direct migration from DB2 8 and DB2 9. Customers can choose which version is best for their business needs and plan their migration accordingly. The end of service for DB2 8 has been set for April 2012. Customers still on version 8 need to build a plan to move.

If you're building a migration plan and need more information about DB2 8 migration steps, go to: **ibm.com**/support/docview.wss?uid=swg27005641

Application programming improvements: Temporal tables and versioning

Raw performance and availability improvements will warm the heart of any DBA. But DB2 10 also brings a number of new application programming features to the table. One of the most exciting is built-in support for temporal data.

You need temporal support anytime you want to ask a question with a time-based element. For example, what level of insurance coverage did a client have six months ago when they had an accident? What was the medical condition of a patient at a specific time? What changes were made to a client's financial account during the last five years?

Previously, answering questions like these required developers to hardcode complex logic into their applications. Now, developers can instruct DB2 10 to automatically maintain a history of database changes and track effective business dates. This new capability uses simple SQL statements and provides a consistent approach to tracking time-related events, and managing and maintaining versioned data.

DB2 10 introduces two new concepts—system time and business time. System time tracks when changes are made to the state of the table, such as when an insurance policy was modified. Business time tracks effective dates of certain business conditions, such as interest rates. Bitemporal tables allow you to track both system and business time in a single table, and time periods can be added to indicate start and end points.

Using system time and business time

Defining a table with a system time period instructs DB2 10 to automatically capture changes made to the state of the table and to save the "old" rows in a history table. Simple SQL

Policy						
ID	**VIN**	**Annual_mileage**	**Rent**	**Coverage_amt**	**Sys_start**	**Sys_end**
111	A11	10000	Y	750000	01-31-2011	12-31-9999
141	B09	14000	N	750000	11-15-2010	12-31-9999
Policy_History						
ID	**VIN**	**Annual_mileage**	**Rent**	**Coverage_amt**	**Sys_start**	**Sys_end**
111	A11	10000	Y	500000	11-15-2010	01-31-2011

Figure 1: *When a record is updated, DB2 10 automatically adds a copy of the old row to the* POLICY_HISTORY *table.*

Policy						
ID	**VIN**	**Annual_mileage**	**Rent**	**Coverage_amt**	**Bus_start**	**Bus_end**
444	A44	40000	N	600000	01-01-2011	12-31-2011

Figure 2: *The original insurance policy includes mileage, coverage, and start/end dates.*

Policy_History						
ID	**VIN**	**Annual_mileage**	**Rent**	**Coverage_amt**	**Bus_start**	**Bus_end**
444	A44	40000	N	600000	01-01-2011	06-01-2011
444	A44	40000	N	750000	06-01-2011	09-01-2011
444	A44	40000	N	600000	09-01-2011	12-31-2011

Figure 3: *The SQL command requires DB2 to split the rows for the policy, adjusting for the new coverage amount and date information.*

queries that reference the current table but also need data in the history table will cause DB2 to transparently access the history table as needed, providing easy access to historical data without complex WHERE clauses with various timestamp and join conditions. For inserts, DB2 generates the appropriate values for system and transaction start times without having to reference them in the INSERT statement. When updating current data, DB2 automatically maintains old versions of the data in the appropriate history table—transparently.

Figure 1 shows what happens when an automobile insurance policy is updated with a new coverage amount by using the following:

```
UPDATE policy
SET coverage_amt = 750000
WHERE id = 111
```

DB2 updates the row in the current table and moves a copy of the old row to the history table, recording the system time start and end values. Subsequent updates are handled similarly. A simple DELETE command causes DB2 to remove the data from the current table and maintain an old version in the history table, including the end time of the deleted data. Users can access this data with simple SQL containing a time period specification.

To query information about the coverage amount recorded in the database for policy 111 for December 1, 2010, you simply use SELECT coverage_amt FROM policy FOR SYSTEM_TIME AS OF '12-01-2010' WHERE id = 111. DB2 transparently accesses the history table to retrieve the correct information.

You can also update, delete, and query data using a FOR PORTION OF BUSINESS_TIME clause. The following example shows how DB2 splits the appropriate rows for an automobile insurance policy by adjusting effective dates and coverage information through very simple SQL.

The customer's original policy ID number is 444 and the coverage amount is US$600,000 (see Figure 2).

The customer makes a request to increase the coverage amount to US$750,000 during a three-month trip, This is done with a SQL command that causes DB2 to split the row in two and insert a new row indicating the new coverage amount and the time period during which it is effective (see Figure 3):

```
UPDATE policy
FOR PORTION OF BUSINESS_TIME FROM '06-01-2011' TO '09-01-2011'
SET coverage_amt = 750000
WHERE id = 444
```

Bitemporal tables allow you to manage data with both system and business time simultaneously, combining the benefits of both concepts. Bitemporal data can also be used in history tables.

DB2 and System z: A perfect match

DB2 10 delivers significant, out-of-the-box CPU savings for many workloads, and customers can see the benefits by simply migrating and rebinding. Application development has never been easier for DB2 for z/OS. With DB2 10's temporal data features, application developers and system administrators can support time-based data more easily than ever before. This is one release you don't want to miss.

Terrie Jacopi *is program manager for DB2 for z/OS.*

RESOURCES

DB2 for z/OS news and updates:
ibm.com/software/data/db2/zos/db2-10

White paper: IBM DB2 10 for z/OS—Reduce costs with improved performance:
www14.software.ibm.com/webapp/iwm/web/signup.do?source=sw-infomgt&S_PKG=db2_zos_reduce_costs

DB2 for z/OS Overview presentation:
ftp://ftp.software.ibm.com/software/data/db2/zos/presentations/overview/db2-10-overview-idug-na-2010-josten.pdf

DB2 and z/OS Leverage

Develop tools, applications, even MVS system exits with the speed, ease and versatility of REXX and REXX Langauge Xtensions (RLX). Prototype any DB2 or z/OS requirement.

- *REXX interfaces to SQL and pureXML*
- *REXX/VSAM interface*
- *Native REXX support for DB2 commands, IFI requests and RRSAF/CAF services*
- *Extensive set of Software Development Kit functions for REXX*
- *Supports both static and dynamic SQL syntax*
- *Compatible with DSNREXX and extends the IBM interface with support for pureXML, Large Objects and multi-row FETCH*
- *Load SQL query results directly into ISPF tables, then display and process those results on scrollable ISPF panels*
- *Interpretive REXX compiler*
- *CLIST to REXX translator*
- *Immediate feedback, extensive diagnostics and content sensitive help*

Full DB2 10 for z/OS support

Unsurpassed productivity with REXX and RLX

Let us show you how. Call *800 776-0771 / Tel: +1 201 420-0400 / www.relarc.com*

Relational Architects International

Server and Storage Focus

IBM Systems Deliver for Data Managers

When it comes to servers and storage, what keeps data managers awake at night? Welcome to the first installment of Server and Storage Focus, where IBM experts Bob Zuber, worldwide marketing manager for System x Enterprise Servers, and Bina Hallman, director of entry and midrange disk storage, explain the features of IBM servers and storage systems that help give data managers a good night's rest.

Performance-to-Cost Ratio

Q: Performance and capacity demands from internal and external customers are on the rise. But most database managers have had their capital and expense budgets either frozen or reduced, which is enough to give anyone insomnia. What can database managers do?

Zuber: IBM has been in the x86 server business forever. We understand the pain that data managers are going through, and we were thinking of them when we created our latest eX5-based servers. We changed the fundamental engineering at the heart of the system to unclog the bottlenecks of even the most demanding x86 environments. This new line of enterprise-class servers helps lower costs, improve utilization, and simplify deployments.

The new designs are ideal for organizations that need additional memory and enterprise-class performance in a two-socket system. The IBM® System x® x3690 X5 is a new class of server, and it uses the new Intel® Xeon® processors to provide improved memory scaling, which is especially critical in virtualization initiatives. This new design point provides twice the amount of memory than current two-socket designs, and it can be further enhanced by attaching a MAX5 unit to the system. The MAX5 is a memory drawer that provides an additional 32 DIMMs, doubling the number of DIMM slots from 32 to 64, while remaining in a 3U form factor. Now you have the flexibility of adding more memory without adding more processors, enhancing performance while realizing significant cost savings.

Hallman: We have a strong prescription to help database managers get some much-deserved rest: systems with the performance and capacity necessary to meet today's workloads at affordable prices and the ability to grow to protect your investment. The IBM System Storage™ DS5000 series can be easily scaled up, and also supports intermixing drive types and host interfaces, for investment protection and cost-effective tiered storage. All DS series systems can be centrally managed, maximizing utilization and lowering costs, and are designed to support high availability with key hot-swappable components and many non-disruptive firmware upgrades.

The leading manufacturers in the storage industry came together about 10 years ago to form the Storage Performance Council (SPC), a nonprofit corporation, to define, standardize, and promote storage subsystem benchmarks as well as to disseminate objective, verifiable performance data to the computer industry and its customers. IBM is a member of SPC, and the IBM DS5000 series of disk systems has consistently achieved some of the best results in the industry.

The SPC benchmarks provide a way for storage systems of differing sizes from different manufacturers to be examined in an "apples to apples" comparison. The tests deliver "value points," expressed in either dollars-per-IOPS or dollars-per-MBPS. The SPC-1 test employs a maximum latency ceiling (30ms) to ensure realistic performance at all workload levels. The price of the tested system is divided by the IOPS achieved and a Random Value Point (dollars-per-IOPS) is displayed. The SPC-2 test examines three different types of sequential performance in order to fully reveal a system's ability to handle various workloads. The results are averaged to reach a Sequential Value Point (dollars-per-MBPS).

Performance to cost is a good measurement, but other design requirements must be considered as well, such as performance under mixed workloads. IBM worked closely

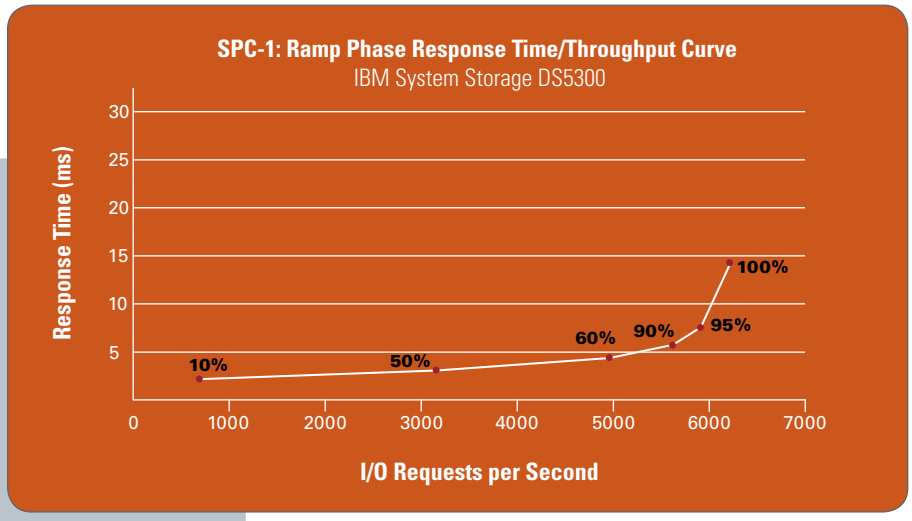

As these SPC-1 and SPC-2 test results show, the IBM DS5300 is capable of handling both highly random workloads along with highly sequential workloads, such as those typically encountered on a daily basis in many organizations. The SPC-1 response time/throughput curve illustrates the average response time and I/O request throughput at 100%, 95%, 90%, 80%, 50%, and 10% of the workload level used to generate the SPC-1 IOPS™ metric.

SPC-2 Reported Data

IBM System Storage DS5300 (8 GB, RAID-5)

SPC-2 MBPS™	SPC-2 Price/Peformance	ASU Capacity (GB)	Total Price	Data Protection Level
$5,634.17	$74.13	16,383.186	$417,648	RAID-5

The above SPC-2 MBPS value represents the aggregate data rate of all three SPC-2 workloads: Large File Processing (LFP), Large Database Query (LDQ), and Video on Demand (VOD)

with VMware, Microsoft, and Enterprise Strategy Group (ESG) to design a repeatable test methodology for demonstrating the real-world performance of storage systems in a virtual (both VMware® vSphere™ and Microsoft® Hyper-V™) mixed-workload environment. The tests show that the DS5000 systems deliver outstanding performance and scalability in a mixed-workload environment. The report of the test plan and the results have been published for more than two years and—to our knowledge—no other storage provider has published their product mixed-workload results.

You can find the SPC test results and mixed-workload test results at www.ibmdatamanagement.com.

Transparent Scalability

Q: With information volumes exploding, it can be difficult to stay focused on managing the data while worrying about running out of server bandwidth or storage disk space. Do you have something that will help data managers get their beauty sleep?

Hallman: Our engineers design features to make the data manager's life a little easier, by giving them a system that can easily grow over time.

Choice, flexibility, and no disruptions are all core design principles of the DS5000 family. The DS5000 series storage systems can scale from 5 to 480 drives, for more than 960 terabytes of raw storage capacity. We also provide a wide range (SSD, standard and self-encrypting Fibre Channel, large-capacity SATA) of hot-swap drive choices so that DB2 managers can adapt and tier the DS5000 system to their specific requirements.

Also, IBM engineers have built the DS5000 series to support online expansion, which allows you to non-disruptively add drives and expansion trays to the system. You can also non-disruptively add drives to existing groups to obtain greater performance; the DS5000 allocates the additional storage capacity in the background, without affecting the systems, applications, or end users.

Zuber: In the x86 world, the weakest link in delivering application performance has always been the amount of memory that could be installed in a server. If you needed more memory, you had to add processors. This, of course, drove higher levels of performance, but it also drove up cost. And often, many of those additional processors were going unused.

To solve this problem, IBM has delivered an x86 technical milestone. The MAX5 memory expander allows you to maximize memory without having to purchase additional costly processors and software licenses. MAX5 breaks a technical barrier of memory capacity that has hindered x86 computing for three decades.

Of course, the eX5 architecture also makes it easy to add processing power when necessary. For instance, the System x3850 X5 is a four-socket 4U server, providing great performance for databases and virtualization. But once that workload grows beyond the confines of the system for either memory or CPUs, the x3850 can be scaled to an eight-socket configuration by adding a node.

Learn more

IBM Server and Storage Systems: www.ibmdatamanagement.com

Watch for more installments of Server and Storage Focus in upcoming issues of *IBM Data Management* magazine.

Development Notebook:

The Evolution of the IBM Smart Analytics System

Optimization, tuning, and the art of fitting 20 years of warehousing and analytics experience in a single package

By Ives Brant

The IBM Smart Analytics System is another step forward in integrated appliances: highly optimized and powerful, with options for built-in analytics and some attractive features aimed at DBAs. We had a chance to sit down with two key development team members—Haider Rizvi, senior technical staff member, IBM Smart Analytics System architecture, and Nancy Kopp, program director, competitive and analyst strategy, IBM Information Management Software—and get some of the background on the system and the long history behind its development.

Q: What sparked IBM to start building analytic appliances?

Rizvi: Market demand. Data warehousing and analytics were continuing to spread everywhere, and they often became multivendor environments. We had customers run into problems with this. In one situation, the system simply stopped running, and it

wasn't clear where the problem originated. After 15 months of frustration, it was discovered that the servers were undersized for the application and workload, and the I/O on the storage unit was bottlenecking operations.

Q: So this kind of situation was the catalyst?

Rizvi: Not many cases are that extreme, of course, but refereeing between several vendors to determine the root cause of a problem was a serious customer pain point.

Kopp: We were also seeing customers getting stalled with the integration, trying to bring analytics online—and such delays become less tolerable all the time.

Q: And that's less tolerable given the demand for real-time data availability and query responses, and high availability, right?

Kopp: Of course. And you don't get high availability on mixed, high-stress workloads without a lot of design experience.

Q: But why build a complete, integrated system? Isn't every situation still different?

Rizvi: Good point. Yes, certain software features will vary. But our hardware building blocks are standard, fully integrated pieces that we can clip together. They are very heavily tested and optimized—we know how they'll behave in heavy usage, and they have been used extensively in the real world.

Q: How long have they been in the marketplace?

Kopp: These units go back almost to when IBM started building data warehouses in the early 1990s. We developed a methodology for putting together and optimizing building blocks of hardware, software, and OS [operating system] to handle warehousing and, later, analytics.

Q: And this led to the IBM Smart Analytics System?

Rizvi: The IBM Smart Analytics System is effectively in its fifth generation, but yes, it's been a long-term, incremental development path to get here. Originally we named it the BCU methodology, for Balanced Configuration Unit. Over the years, BCU became the building blocks of the Smart Analytics System. Today everything in the building blocks reflects what is now market-standard and proven with today's workloads. They are DBA-friendly.

Q: Because the integration and load testing has been done?

Rizvi: Yes, but more than that. We did several things with the DBA in mind, including the control features. But for starters, over nearly 20 years, we have built up skills and experience at optimizing integrated, ready-to-run systems with balanced operations, so every component from processor to memory to I/O and operating system and the database can handle a specified workload without getting bottlenecked, with an attractive total cost of ownership.

Q: With that approach, do you risk creating a result that's too generic for some customers?

Rizvi: Usually not. Doing this over time, we learned which integration tasks we are better equipped to handle than most customers, with a combined hardware-software solution, and the kind of support they require.

Q: Is building an analytics system really that complex? Lots of companies have created their own.

Kopp: It is a major task. Many of those companies you mention struggled with the implementation and needed up to a year to get the system working smoothly.

Rizvi: It's useful to understand two things. First, creating a high-performance analytics system is a delicate balancing act. All of the components, from the memory and processor up through the BI [business intelligence] tools, need to be tuned to work together and balanced against each other. IBM patented the methodology we use to do this. We're not talking about applying a system tuning checklist, or a couple of "a-ha!" moments.

Q: So even the patented methodology is not a story of a single nova-like breakthrough.

Rizvi: That's right. It's a story of persisting over two decades with nonstop, incremental improvements. That led to the strong, experience-based methodology for integrating and optimizing appliances—which was granted the patent. That brings us to the second point: we're talking about creating a system that wrings every drop of performance out of the hardware, and one that can be unboxed, loaded with data, and working in a few days or weeks. This is not something that just anyone can do.

Q: It can take months.

Rizvi: It's fair to say that customers who take on the task of buying and integrating their own mix of servers, memory, I/O, operating system, database, warehouse, analytic applications, and ETL [extract, transform, load]—and backup—expect to spend months on the task.

Q: What was the biggest challenge that the team faced?

Rizvi: [laughs] I'm not sure that I can pick one out! But we definitely stay up late thinking about how to keep I/O at maximum efficiency. Many pieces contribute to efficient I/O. At one point, when we saw the stresses that analytic workloads applied, we tried offloading temp space to a solid-state device. The performance gain was exceptional.

Q: That must have felt like a breakthrough moment.

Rizvi: It did—and this feature became a permanent part of the Smart Analytics System 5600 and 7700. The systems also leverage compression; it plays an effective role in driving down the cost of handling a specific workload.

Q: You mentioned that your integrated appliances are DBA-friendly. What else relates to that?

Kopp: Development of the IBM Smart Analytics System brought together our hardware and software groups within IBM, which assisted collaboration. Unifying the two teams led to features—such as the system management console—that span software and hardware, helping DBAs to take on more responsibility on both sides and to have a more strategic role.

Q: What is the system management console?

Rizvi: Every IBM Smart Analytics System 5600 and 7700 will be equipped—within the next several months—with a management console that gives the DBA or other administrator command-line control over the entire system: hardware, software, operating system, driver, firmware, and other components.

> "We're talking about creating a system that wrings every drop of performance out of the hardware, and one that can be unboxed, loaded with data, and working in a few days or weeks. This is not something that just anyone can do."
>
> — **Haider Rizvi,** *Senior Technical Staff Member, IBM Smart Analytics System Architecture*

Q: How is it different from typical software administration screens?

Rizvi: This console gives DBAs or other administrators the ability to maintain the whole cluster for not just software, but firmware and hardware as well. The administrator can see OS and firmware upgrades at a glance, and can orchestrate an update across an entire cluster that has dozens of data nodes.

Q: How does the administrator operate the management console?

Rizvi: It's command-line control oriented. The DBA can address the tooling for all DB2 software, IBM InfoSphere, the OS, and firmware. When conducting updates, at the right point it will even instruct the administrator to reboot so the changes take effect.

Kopp: The console eliminates certain surprises. You won't have a situation where the system crashes and that's how you learn a colleague updated the new OS version without telling you.

Q: So, how does that lead to a more strategic role for DBAs?

Rizvi: Database administrators have been operating beyond the database for years; for example, by learning the OS and determining how to lay data down. In so doing, they've become co-architects in the overall data warehouse environment.

Q: How does IBM ensure these units work right out of the box with the customer's data?

Rizvi: We're taking our direction from real-world deployments. We do extensive testing with the Smart Analytics System—it's as if a car manufacturer put your family and luggage in the vehicle and towed your boat repeatedly to wherever you vacation, before delivering it to you. When possible, we develop and test the Smart Analytics System using real customers' data.

Q: Can you be more specific about the testing process?

Rizvi: IBM runs three types of query scenarios: very complex and demanding, a group of medium complexity, and finally large numbers of small queries. Our workloads are in two categories: one for performance benchmarking, and the other for customer stress tests. Both mimic large numbers of users and put the I/O system to the test. We benchmark each unit—this is important to mention—and then we use the same benchmark tests after the unit is installed with the customer's data, to make sure no problems have entered the picture.

Q: What is the upcoming development path for the IBM Smart Analytics System?

Rizvi: We'll be studying and making sure we can handle growing data warehousing workloads. Pretty much what we've done for 20 years: ongoing, incremental improvement and optimization.

Q: How do you see the DBA's role evolving, and does IBM play a role in that evolution?

Rizvi: Database administrators' roles are changing in two major ways: first, spanning software and hardware both, and second, helping business users tap more of the potential of warehouses, analytics, and other new applications. That, as it happens, can occur when they are less often forced into a troubleshoot-and-repair mode.

We support their more strategic role with integrated appliances that give them broader control, give them fewer issues to troubleshoot because of our integration and testing, and contain important new applications that will take star roles in the years ahead. ✱

Ives Brant *has experience in the database and analytics industries, and was the first employee and editor-in-chief at* Tornado Insider *magazine, a European competitor to* Red Herring. *These days, he writes technical content for a wide range of companies.*

Advanced Analytics: The Newest Tools on the Farm

Based in Bakersfield, California, Sun World International grows and markets fresh fruit and vegetables. Although this description conjures a sentimental image that resonates with most people, the reality for Sun World is that the company must compete in a globalized commercial farming industry where rational analysis trumps sentimentality and timely, accurate decision support is more important than ever. According to Steve Greenwood, Director of Budgets and Reporting at Sun World, "Our aim is to produce quality products when they're currently not available. It's about hitting the market window when nobody else can."

Sun World has had an edge for several years now. The company put a critical building block in place when Greenwood brought in IBM Business Partner Applied Analytix, Inc., to redesign Sun World's financial statement analysis and reporting capability—replacing the manual input of data into electronic spreadsheets. Once the IBM Cognos TM1 software was implemented, the Finance Department had instant access to the financial information it needed to make fast business decisions.

Recently, Sun World decided to expand its analysis capability to the field level, so the company could better determine what crops go into the ground. Again Greenwood called on the expertise of Applied Analytix to create a Cognos TM1 system that addressed farming costs and harvesting efficiency on the company's 16,000 acres of farmland. Fed by Sun World's ERP system, the analytics solution can now provide a granular breakdown of sales, costs, and profitability by product type, variety, region, and individual ranch. For example, because Sun World can compare the costs of high-efficiency drip irrigation methods to traditional surface irrigation, it's been able to make better judgments about where to apply drip irrigation, resulting in lower water usage and better nourished produce. When growing table grapes, for example, targeted drip irrigation has resulted in a 5 percent reduction in harvesting costs, a 20 percent reduction in fuel usage, and a 50 percent increase in yield over the past five years.

To complete the efficiency circle, the Sun World sales staff has up-to-the-minute price data as they negotiate with buyers. Sun World displays this data on the selling floor via a ticker display, drawn from a near-real-time feed from Cognos TM1. "This gives our salespeople the ability to continually negotiate better prices," explains Greenwood. "A penny here, a penny there matters when you're shipping 11 million boxes per year."

applied analytix
Delivering Informed Decisions

For an in-depth view of how Steve Greenwood and his team are providing the decision support that has helped elevate Sun World—a midsize company—from a regional commodity producer to a world-class breeder and provider of specialty branded produce, look for his customer presentation at the Information On Demand 2010 Global Conference in Las Vegas in October. After IOD, download his presentation at ***www.applied-analytix.com/iod****.*

APPLIED ANALYTIX

Applied Analytix, Inc. has 15 years of experience bringing the power of advanced analytics to all types of businesses. Visit us at www.applied-analytix.com.

Informix TCO:

The Facts Are In

Research studies put hard numbers behind the business case for Informix

By Chris Young

Why would a company choose Informix database software to manage its data infrastructure? Many organizations easily answer that question: performance, reliability, stability, ease of use, and extremely low administration and maintenance costs.

Now, two well-researched studies help paint a more specific picture of how, why, and where Informix reduces total cost of ownership (TCO) and maximizes return on investment (ROI). The International Technology Group (ITG) recently completed a study that directly compared Informix with database software from another well-known global vendor and confirmed a variety of Informix benefits. Both that study and a recently completed "Total Economic Impact" commissioned study conducted by Forrester Consulting on behalf of IBM based on a large U.S. retailer demonstrate that these benefits can directly impact the bottom line.

32.7%
Reduction in three-year TCO for Informix compared to database software from another vendor

2 years
100% payback period for a representative Informix deployment

US$5.4 million
Total three-year benefit resulting from a representative Informix deployment

Studies reveal a compelling cost picture

For organizations seeking an economical database solution, the cost picture presented by the two studies is compelling. The ITG study is built around two sets of research data. First, a survey of 62 midsize Informix users in North America and Europe documented DBA staffing levels and confirmed that these levels are significantly lower for Informix than for equivalent competitive systems. The survey also documented other benefits of Informix use.

The users in this survey had all deployed Informix-based solutions on Windows x86 or, in a few cases, Intel Itanium–based servers, and had at least one year of experience operating them. Users reported between 6,000 and 7,000 employees, and sales of between US$1 million and US$1.6 billion. The survey population spanned a wide range of industries, including manufacturing, telecommunications, retail, financial services, healthcare, transportation, government, media, IT services, real estate, agribusiness, engineering, and construction.

As a second step, ITG compared configurations, software, and DBA staffing among these Informix users and equivalent users of database software from another vendor (see sidebar, "Who was that masked database server?") and calculated the TCO of using the different platforms. Three-year TCO was significantly lower for Informix, averaging 32.7 percent less than costs for the other database software (see Figure 1).

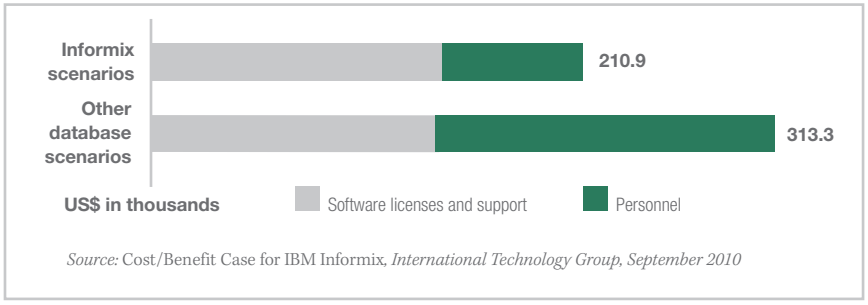

Source: Cost/Benefit Case for IBM Informix, *International Technology Group, September 2010*

Figure 1: *Three-year costs for use of Informix and database software from a competing vendor; averages for all installations*

Retailer sees favorable ROI

While the ITG study looked at the experience of dozens of different users, the Forrester study was a single-company analysis focused on a global retail organization. The retailer sought to expand the use of Informix for applications that have high transactional volumes—for example, a pharmacy application that provides access to prescription information across stores as well as back to the central distribution facility.

The objective of the Forrester study was to calculate the ROI or payback for this expansion project, which would also serve as an example of the ROI that other enterprises could expect from the Informix platform. Unlike the ITG study, which focused solely on the cost of software licensing, support, and personnel, the Forrester study looked at the total economic impact. This included not only software licensing, personnel, and annual maintenance, but also the cost of upfront implementation and hardware.

Forrester compared these costs to the value of the benefits provided by Informix in the retailer's environment. Specific benefits included improved administration efficiency, reduced impact of planned and unplanned downtime for applications, and improved server utilization of existing server assets, as well as lower overall support costs resulting from fewer system incidents. Forrester determined the value of these benefits to be more than US$5 million over three years.

Based on the benefits and costs, Forrester calculated an ROI of 32 percent with a break-even point of 2.0 years after deploying Informix (see Figure 2). ROI is calculated by dividing net benefits (benefits minus costs) by costs.

ROI	Payback period	Total benefits (PV)	Total costs (PV)	Net present value
32%	2.0	US$5,293,535	US$4,133,258	US$1,160,277

Source: The Total Economic Impact Of IBM Informix Database Server, *a commissioned study conducted by Forrester Consulting on behalf of IBM, September 2010*

Figure 2: *Three-year adjusted return on investment*

With the Informix expansion calculated to pay for itself in just two years, the retail organization analyzed by Forrester decided to proceed with the project. The Forrester study showed that Informix would help reduce the need for distributed, store-level staff to manage the pharmacy application—a point of particular interest to the retailer. It also showed that Informix would be able to deliver high levels of performance in an environment of rapid growth and could scale quickly depending on store demand.

Ease of administration cited as top benefit

Both studies validated the Informix reputation for ease of use. Improved management efficiency of the database environment was a key Informix benefit for the retailer in Forrester's study, and 73 percent of users in the ITG study said that a primary advantage of using Informix was that databases could be easily administered.

Among organizations of all sizes, users reported that very little time was spent on routine database maintenance tasks and that little monitoring was required. ITG survey respondents described Informix in such terms as "very simple…very easy to administer…hardly any administration…almost no DBA time…virtually no tuning…near-zero

Who was that masked database server?

To see these studies in their entirety, and to reveal the true identities of the database software that ITG compared to Informix, visit **ibm.com**/informix.

- *Cost/Benefit Case for IBM Informix*, International Technology Group, September 2010
- *The Total Economic Impact Of IBM Informix Database Server*, a commissioned study conducted by Forrester Consulting on behalf of IBM, September 2010

73%
Percentage of survey respondents citing ease of administration as the No. 1 advantage of Informix

US$96,000
Retailer's early reduction in annual hardware spending with Informix (fewer additional servers needed as Informix is scaled)

maintenance…runs unattended most of the time." Users praised specific Informix features as well: for example, DB Scheduler, which provides a framework to automatically schedule and monitor database activities, was described as "excellent…very easy to use."

A number of the companies studied by ITG used other databases from well-known vendors alongside Informix, and there was general agreement that DBA overhead was significantly lower for the Informix platform. In several cases, the number of DBAs required for administration fell from three or four to one after moving to an Informix database. Overall, DBA staffing for Informix was reported to be between two and four times less than for other databases.

Some users also cited the comparative ease with which non-DBAs could be trained to handle Informix administration tasks. For example, in an organization that has few staff members, a second individual could be trained to back up the primary DBA if the latter were unavailable.

Many unable to recall any Informix outages

High reliability and availability of the Informix platform were confirmed as well. Indeed, unplanned outages were sufficiently rare that many ITG survey respondents could not recall a single one. Several organizations reported that they had experienced outages due to other causes, such as hardware failures, power outages, and network disruptions. In these cases, Informix databases were rapidly recovered without data loss or corruption. Many users had not experienced either planned or unplanned outages over long periods, except for version upgrades and major application changes.

Respondents also mentioned the comparative simplicity of replication as a factor in choosing Informix. One organization that needed to support 24x7 business operations, for example, had reviewed Informix alongside databases from two other well-known vendors. It noted that Informix provided sophisticated failover and recovery features as standard, and that the use of clustering for the other platforms would have required additional, separately priced tools and a significant amount of custom consulting assistance.

Strong performance reduces hardware spend

While ease of use and reliability were top reasons for liking Informix, performance was also a big factor. ITG survey respondents reported strong performance not only for low-end, comparatively inexpensive installs by IT services companies, but also for other installations supporting larger-scale workloads. These included enterprise resource planning (ERP), query-intensive, image-intensive, and Internet workloads, as well as industry-specific systems in fields ranging from banking to distribution and retail.

In the Forrester study, improved performance through Informix allowed the retailer to more effectively utilize server resources. The organization in the study noted that, with Informix, it could delay or postpone the upgrade of existing servers to maintain performance compared with other platforms. By avoiding the need to purchase additional hardware server assets as it scaled its Informix footprint, the organization calculated that it could reduce its hardware spend by US$96,000 per year.

Making stretched resources do more

Traditionally, the core appeal of Informix has been that it enables organizations to achieve their goals while maintaining an essentially simple IT environment. These new studies confirm that this appeal translates into low administrative overhead and high levels of availability, which can significantly reduce costs and enhance ROI. Informix ownership costs were shown to be lower—by wide margins—than those of competing installations. All of these factors contribute to making Informix a logical choice for businesses looking for ways to make already stretched resources go further. ✱

Chris Young *is a technology writer based in the Pacific Northwest.*

Give Your Business the Fuel to Grow

IBM DB2 pureScale and IBM System x Servers Based on Intel Xeon Processors Drive Enterprise-Class Transaction Processing

In almost every industry you can name—retail, manufacturing, telecom, government, finance—transactions are the fuel that drives the business. They're credit-card payments at the mall, sell orders on the stock exchange, hotel bookings, and mobile phone call connections.

More transactions means more money coming in, more orders going out, more communication with customers and suppliers, more activity, and more business. In an increasingly instrumented, interconnected, and intelligent world, transaction volumes will only grow: estimates of the amount of data created each year have been rising exponentially, and that trend shows no sign of abating in the years ahead.

Organizations that master online transaction processing (OLTP) and develop the ability to easily handle ever-larger transaction volumes have the opportunity to differentiate themselves from the competition and stay a step ahead of the market. For businesses ready to take this next step, their choice of an OLTP platform is critical. It must not only provide outstanding performance; it must deliver that performance every minute of every day, regardless of fluctuating loads or maintenance needs. It must also stand ready to expand and adapt smoothly as the organization it supports grows and changes.

Intel and IBM have been working for years to create systems that change the economics of x86 computing. One of the results of that collaboration is an OLTP platform for clustered, distributed systems that delivers unprecedented flexibility and scalability.

IBM® DB2® pureScale™ provides near-limitless capacity and continuous availability running on IBM System x® servers powered by Intel® Xeon® processors.

With this new solution from IBM, organizations can build responsive, highly available OLTP systems that quickly scale in any direction needed, while remaining completely transparent to enterprise applications.

IBM DB2 pureScale

DB2 pureScale takes proven IBM DB2 for z/OS® mainframe technology and combines it with the latest distributed system techniques. This provides global caching and locking control, eliminating the challenges and limitations of hosting databases on distributed systems. In a DB2 pureScale cluster, one server provides centralized lock management services, a global cache for data pages, and other services. Members of the pureScale cluster can simultaneously access a shared database for read and write operations; they are also connected directly to the central server through remote direct memory access (RDMA) and DB2 pureScale technology, running over high-performance InfiniBand* network links (see Figure 1). The result is a cluster of DB2 systems that—to OLTP applications—looks and acts like a single, incredibly fast server.

DB2 pureScale technology raises the bar for availability and scalability in OLTP systems. Downtime can be virtually eliminated because if a cluster host fails, the system automatically fails over to an active host. DB2 provides full access to all pages of data that do not need to be recovered, and

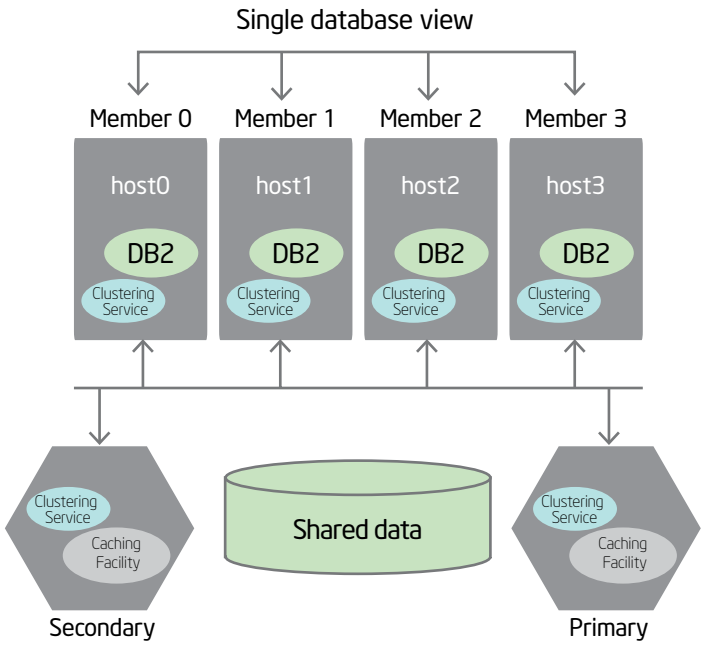

Figure 1: In a DB2 pureScale cluster, each member has remote direct memory access to the centralized locking and caching services of the DB2 pureScale server.

interrupted transactions are recovered or rolled back within seconds.

During the recovery, other members of the cluster can continue to process transactions—so the systems and customers that depend on the cluster can continue business as usual. Likewise, planned outages become a thing of the past, as IT can perform rolling maintenance one server at a time, without bringing the entire system down.

DB2 pureScale clusters can be easily and quickly scaled up—adding new servers requires just two commands. You can take a building-block approach to database expansion, beginning with two partially populated servers and adding capacity incrementally and without limit as requirements grow.

IBM has also built flexibility into DB2 pureScale licensing: clients can pay for the DB2 and DB2 pureScale they need, when they need it, and only for the period of time in use. That's great news for IT managers who are looking to cut costs.

The great news for developers is that no application changes are needed to take advantage of DB2 pureScale clusters. Applications do not need to be cluster-aware, so there is no need to alter or replace existing code. And, because DB2 offers compatibility technologies such as native support for commonly used syntax and the PL/SQL procedural language, organizations can run applications written for other database platforms—such as Oracle and Sybase—with few or no code alterations.

Hardware Performance Powered by Intel and IBM

Providing a complete, high-performance solution for OLTP applications, IBM System x servers powered by Intel Xeon processors are an ideal platform for DB2 pureScale clusters.

The Intel Xeon processor 7500 series delivers the largest performance leap ever for an Intel processor generation, with an average of 3x higher performance across a wide range of industry benchmarks[1] compared with its predecessor. It provides up to 20x better performance per server versus single-core servers.[2] Large on-die cache (up to 24 MB per socket), Intel® Hyper-Threading Technology, and Intel® Turbo Boost

Technology deliver exceptional performance for the thread- and memory-intensive workloads that are common to OLTP. Also, Intel® QuickPath Technology and two integrated memory controllers speed traffic between processor and I/O controllers, delivering up to 8x the memory bandwidth[3] for data-intensive transactional, database, and analytics applications.

The Intel Xeon processor 7500 series also features Intel® Advanced Reliability technology that provides automatic detection and correction of errors, dynamic reassignment of workloads across CPUs, interconnect error detection/recovery, and individual virtual machine recovery in virtualized environments. The result is high-end data integrity and availability for each individual host, providing an even more powerful and reliable foundation that helps keep transactions flowing reliably.

IBM eX5 servers take advantage of the Intel Xeon processor 7500 series to provide a powerful foundation for clustered database solutions. With eight-socket, four-socket, and two-socket configurations available, you can build DB2 pureScale clusters that fit your current needs and have plenty of room for future growth. Just as you can easily add servers to a DB2 pureScale cluster, you can easily scale individual servers to support massive workloads. For example, the eight-socket server configuration provides up to 64 processor cores, 128 threads, and 3 TB of memory, and there is no architectural limit to the number of hosts in a cluster. Also, with the Intel Xeon processor 7500 series, you won't sacrifice performance as the number of sockets increases (see Figure 2).

The architecture of the eX5 enterprise systems represents the largest investment that IBM has ever made in Intel processor-based servers, delivering major improvements in memory capacity, storage flexibility, system reliability, and deployment simplicity. For example, the IBM MAX5 external memory chassis enables memory that can be scaled independently of processors. The MAX5 increases the x3850 eX5 base memory capacity by 50 percent—from 1 TB to 1.5 TB—to support faster database performance.

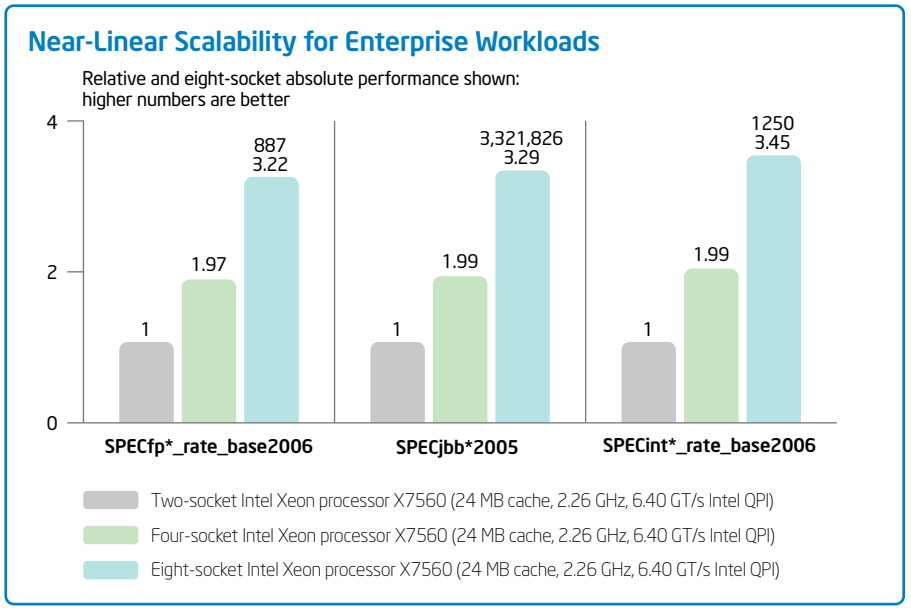

Source: www.intel.com/performance/server/xeon_mp/scalability.htm

Figure 2: *The Intel Xeon processor 7500 series delivers near-linear scalability in two-socket, four-socket, eight-socket, and larger system designs, providing an excellent platform for demanding enterprise workloads.*

Also, IBM eXFlash technology provides preconfigured solid-state storage media that can dramatically increase storage performance while reducing space and power requirements. Compared with traditional hard disk drive storage solutions, eXFlash storage delivers 99 percent greater performance per watt for database-type workloads running on local disks, while maximizing uptime with 64 times better reliability.[4]

For smaller OLTP workloads, IBM M3 servers based on the Intel Xeon processor 5600 series deliver exceptional performance and energy-efficiency per core. The Intel Xeon processor 5600 series automatically regulates power consumption and intelligently adjusts server performance according to application needs. It delivers up to 15x performance per server over single-core servers,[5] enabling up to 95 percent lower energy costs[6] and helping you to maximize performance in a resource-efficient and cost-effective way.

A Well-Oiled Business Engine

IBM DB2 pureScale brings mainframe concepts to the best of distributed computing technology to deliver practically unlimited capacity, continuous availability, and application transparency for transactional database applications. When combined with IBM System x servers powered by Intel Xeon processors, the result is a platform that can help reduce the risk and cost of business growth—and keep your organization's OLTP engine going strong.

Learn More

More about IBM DB2 pureScale: www.ibm.com/software/data/db2/linux-unix-windows/editions-features-purescale.html

More about the Intel Xeon processor 7500 series: www.intel.com/itcenter/products/xeon/7500/index.htm

More about smarter systems from IBM: www.ibm.com/systems/smarter

Sign up for a proof of concept by sending e-mail to go_db2@ca.ibm.com

---

[1] Average of 3x performance claim based on geometric mean of four industry-standard, common enterprise benchmarks (SPECjbb*2005, SPECint*_rate_base2006, SPECfp*_rate_base2006, and TPC Benchmark* E) comparing best published/submitted results on four-socket (4S) Intel® Xeon® processor X7560–based server platform to best published 4S Intel Xeon processor X7460–based server platform as of March 26, 2010.

[2] Claim: "Up to 20x performance per server" Disclaimer: Intel performance comparison using SPECjbb2005* business operations per second between five-year-old single-core Intel® Xeon® processor 3.33GHz–based servers and one new Intel Xeon processor X7560–based server. Performance tests and ratings are measured using specific computer systems and/or components and reflect the approximate performance of Intel products as measured by those tests. Any difference in system hardware or software design or configuration may affect actual performance. Buyers should consult other sources of information to evaluate the performance of systems or components they are considering purchasing. For more information, visit www.intel.com/performance/server.

[3] 8.2x memory bandwidth claim based on February 2010 Intel internal measurement using Intel internal memory bandwidth workload on comparable 4S Intel® Xeon® processor X7560 and 4S Intel Xeon processor X7460 servers.

[4] To operate a 240,000 IOPs database, a client would need 960 spinning disks (300 IOPs/disk). These would be deployed in 80 JBODs (12 disks each). These would require two entry servers (40 JBODs per server using a ServeRAID adapter with cascading feature).

[5] Claim: "Up to 15x performance per server" Disclaimer: Intel performance comparison using SPECjbb2005* business operations per second between four-year-old single-core Intel® Xeon® processor 3.8 GHz with 2M cache–based servers and one new Intel Xeon processor X5670–based server. Performance tests and ratings are measured using specific computer systems and/or components and reflect the approximate performance of Intel products as measured by those tests. Any difference in system hardware or software design or configuration may affect actual performance. Buyers should consult other sources of information to evaluate the performance of systems or components they are considering purchasing. For more information, visit www.intel.com/performance/server.

- Baseline platform: Intel server platform with two 64-bit Intel Xeon processor 3.80 GHz with 2M L2 Cache, 800 FSB, 8x1 GB DDR2-400 memory, 1 hard drive, 1 power supply, Microsoft* Windows Server* 2003 Ent. SP1, Oracle* JRockit* build P27.4.0-windows-x86_64 run with 2 JVM instances
- New platform: Intel server platform with two six-core Intel Xeon processor X5670, 2.93 GHz, 12 MB L3 cache, 6.4QPI, 12 GB memory (6x2 GB DDR3-1333), 1 hard drive, 1 power supply, Microsoft Windows Server 2008 64 bit SP2, Oracle JRockit build P28.0.0-29 run with 2 JVM instances

[6] Claim: "Up to 95 percent lower energy costs" Disclaimer: Intel comparison replacing 15 four-year-old single-core Intel® Xeon® processor 3.8 GHz with 2M cache–based servers with one new Intel Xeon processor X5670–based server. ROI claims and costs have been estimated based on internal Intel analysis and are provided for informational purposes only. Performance tests and ratings are measured using specific computer systems and/or components and reflect the approximate performance of Intel products as measured by those tests. Any difference in system hardware or software design or configuration may affect actual performance. Buyers should consult other sources of information to evaluate the performance of systems or components they are considering purchasing. For more information on performance tests and on the performance of Intel products, visit Intel Performance Benchmark Limitations.

- Baseline platform: Intel server platform with two 64-bit Intel Xeon processor 3.80GHz with 2M L2 Cache, 800 FSB, 8x1 GB DDR2-400 memory, 1 hard drive, 1 power supply, Microsoft Windows Server 2003 Ent. SP1, Oracle JRockit build P27.4.0-windows-x86_64 run with 2 JVM instances
- New platform: Intel server platform with two six-core Intel Xeon processor X5670, 2.93 GHz, 12 MB L3 cache, 6.4QPI, 12 GB memory (6x2 GB DDR3-1333), 1 hard drive, 1 power supply, Microsoft Windows Server 2008 64 bit SP2, Oracle JRockit build P28.0.0-29 run with 2 JVM instances

Any performance data contained herein was determined in a controlled environment. Therefore, the results obtained in other operating environments might vary significantly. Users of this document should verify the applicable data for their specific environment.

DB2 Indexes and Query Performance: Part 1

Weigh costs and benefits when considering additional indexes

Robert Catterall
(rfcatter@us.ibm.com) is an IBM DB2 specialist.

For plenty of you, indexes are the duct tape of DB2 performance fixes. Got a response-time issue? Add an index to the target table. In truth, reducing the elapsed time of a DB2-accessing query or batch job frequently involves creating a new index; however, while an additional index is very often a solution to a DB2 performance problem, it is not always the right solution. In this first part of a two-part series, I'll point out the costs that go along with the benefits of DB2 indexes, sketch out some rules for when to consider other solutions, and talk about ways to create more effective indexes. Next issue, I'll describe some query performance-tuning actions that do not involve new-index creation and that you can implement when a table is already heavily indexed and you are leery of adding more.

There's no free lunch

Indexes are definitely not without cost. Every additional index that you define on a table adds to the costs associated with your database. Start with CPU overhead: every index on a table multiplies the CPU cost of each insert or delete operation, or the CPU cost of any update that changes the value of an indexed column. Insert a row into a table with three indexes, and DB2 must add a corresponding entry to an index three times. Delete a row from a table with 10 indexes, and DB2 must remove all index entries associated with the deleted rows.

All of this index maintenance activity involves DB2 page requests (called GETPAGEs in a DB2 for z/OS system and *logical reads* in a DB2 for Linux, UNIX, and Windows [LUW] environment), and these tend to be the primary determinant of the CPU cost of SQL statement execution.

Extra indexes also add to the CPU consumption of certain DB2 utilities, such as the REORG utility. Indexes defined on keys that are not continuously ascending will naturally become less well-organized over time. New entries must be placed in particular index leaf pages (because index-key sequencing is strictly enforced, whereas

data-row clustering in a table is not), and if the page into which an entry must be placed is already full (not uncommon when the key is not continuously ascending), the page will be split and some of the entries will be relocated to a formerly empty page that might be a long way away in the index structure. REORG will restore a disorganized index to a well-organized state, at a cost of some CPU consumption.

Other utilities that get more expensive as more indexes are defined on a table include LOAD and RUNSTATS (index statistics in the DB2 catalog must be kept up-to-date), and the utilities that perform index-to-data consistency checking (that would be CHECK INDEX for DB2 for z/OS, and INSPECT with INDEXDATA for DB2 for LUW).

Indexes also have a disk space cost, although that doesn't need to be as brutal as it would have been a few years ago, thanks to the index compression features of DB2 9 for z/OS and DB2 9.7 for LUW. On both platforms, you can now get large space savings (often 50 percent or more) with little in the way of additional CPU overhead.

How many indexes?

As the number of indexes on a table grows, eventually the aggregate CPU cost of maintaining the indexes can outweigh whatever performance gains a new index might deliver. But where is that tipping point? The answer depends on the specific situation, but I generally don't like to see more than 4 or 5 indexes on a table in an online transaction processing (OLTP) environment, or more than 8 to 10 indexes on a table in a data warehouse system. I'm usually OK with a larger number of indexes per table in a data warehouse, because the emphasis there is typically on optimizing data retrieval performance, and query search arguments tend to be less predictable in a business intelligence setting.

Now, these index-limitation guidelines are more rules of thumb than rules to live by. Suppose a table in an OLTP database already has 6 or 7 indexes. Would I give a thumbs-up to adding another? Maybe, if the potential performance payoff is really big—and when I say big, I mean that the new index is expected to reduce the CPU cost of one (or several) of the more expensive queries in the system by 90 percent or more. Even if the proposed new index looks like a performance winner, before going ahead with implementation I might want to see if there are any good tuning alternatives that do not involve creating a new index (that will be the topic of Part 2 of this column).

If you're going to do it, do it right

When you do add an index to a table to cut the CPU cost and run time of a query, you want to get the most out of that new index as practically possible. That means, among other things, getting the order of the key columns right if it's a multi-index key. Why is that important? Because a query will run faster—sometimes much faster—if the query's predicates (its search arguments) match more columns of an index key; and when it comes to DB2's rules for matching predicates with index-key columns, order matters.

Here's a simplified (and therefore not all-inclusive) version of the rules for the most common predicate types, which include equal-type, range (for example, > or <=), and in-list (for example, COL1 IN ('DOG', 'CAT')):

1. For a multicolumn index key, predicate matching starts with the high-order column and proceeds as far as possible with other key columns, one at a time, in the order of their sequence within the index key. Put another way, matching starts at the far left of the index key and proceeds rightward.

2. Getting a match on the column in position *n+1* in the index key requires matching on the column in position *n* with an equal-type predicate or an in-list predicate. In other words, once a key column has been matched with a range predicate, there won't be any matching on columns to the right of that column.

3. Matching stops when a column in the key is skipped. That is, if you get a match on the column in position *n* of the key, and there is not a matching predicate for the column in position *n+1*, there can't be a match on the column in position *n+2*.

So, for example, if you have a query with predicates COL1 > 2 AND COL2 = 'BRICK', an index defined on COL2, COL1 will give you two predicate/key-column matches, while an index on COL1, COL2 will give you one match (refer to rule 2 in the preceding list: matching stops once a key column has been matched with a range predicate).

Use indexes, but choose wisely

Indexes can really turbocharge DB2 query performance, but if you put too many on a table, you could end up with response time and throughput going in the wrong direction. Start out conservatively, with just a few—maybe two or three—on each table in your database. That way, you'll have some reserve index-add capacity, and you'll probably be able to define additional indexes that will look really good from a cost-benefit perspective. Just don't overdo it, and don't waste precious index-add capital on indexes that deliver only marginal query performance benefits—you're looking for the big wins here.

In my next column, I'll show you that there are plenty of things you can do to improve query response times without creating new indexes. Combining those techniques with judicious use of additional indexes is the two-pronged performance-tuning approach that is most likely to get you where you want to go. ✱

RESOURCES

DB2 for z/OS family: ibm.com/software/data/db2/zos/family

Adding and Removing Data Partitions

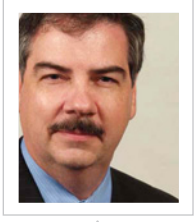

An effective way to migrate large blocks of data

Roger E. Sanders (roger_e_sanders@yahoo.com), a consultant corporate systems engineer at EMC Corporation, is the author of 19 books on DB2 for Linux, UNIX, and Windows and a recipient of the 2010 IBM Information Champion award. He is currently working on a new book titled From Idea to Print: How to Write a Technical Article or Book and Get It Published.

Special thanks to Mike Winer, senior technical staff member–DB2 kernel architect, and Liping Zhang, DB2 data partitioning development lead, for providing information that was used to develop this article.

In my last column, I explained how you can use table partitioning to organize data across multiple storage objects based on values in one or more columns. I also described two ways to use the PARTITION BY clause of the CREATE TABLE statement to create partitioned tables.

One advantage to using partitioned tables is that only relevant data partitions are accessed during query execution; because the DB2 Optimizer is data-partition aware, only relevant data partitions are scanned to resolve a query, resulting in fewer I/Os and higher query performance. Another advantage is that new partitions can easily be added to increase the table range or existing tables containing data can easily be attached (rolled in), while partitions containing old or obsolete data can be removed (rolled out) and archived, for example, to meet company policies or federal laws that mandate record retention.

In this column, I'll walk you through the process of adding and removing data partitions, and I'll show you how enhancements made in DB2 9.7 can make the process easier and faster.

Adding new partitions to a partitioned table

There are two ways to add new partitions to a partitioned table: you can add new ranges and increase the capacity of a table by adding one or more empty partitions, or you can add existing, populated tables by converting them into new partitions. As you might imagine, the process used depends on which approach you choose.

A new, empty partition can be added to a partitioned table by executing the ALTER TABLE statement with the ADD PARTITION option specified; the syntax used looks like this:

```
ALTER TABLE [TableName]
 ADD PARTITION <PartitionName>
 <STARTING <FROM> [Start | MINVALUE | MAXVALUE] |
  STARTING <FROM> ([Start | MINVALUE | MAXVALUE] ,...)
 <INCLUSIVE | EXCLUSIVE>>
 ENDING <AT> [End | MINVALUE | MAXVALUE] |
  ENDING <AT> ([End | MINVALUE | MAXVALUE] ,...)
 <INCLUSIVE | EXCLUSIVE>
 <IN [TSName]>
 <INDEX IN [IndexTSName]>
 <LONG IN [LongTSName]>
```

where:

- **TableName:** Identifies the partitioned table, by name, that a new partition is to be added to
- **PartitionName:** Identifies the unique name, if any, to be assigned to the partition to be added
- **Start:** Specifies the low end of the range for the data partition
- **End:** Specifies the high end of the range for the data partition
- **TSName:** Identifies the table space in which the new partition is to be stored
- **IndexTSName:** Identifies the table space in which any partitioned indexes for the new partition are to be stored
- **LongTSName:** Identifies the table space in which the values of any long columns are to be stored

Note: Parameters shown in angle brackets (< >) are optional; parameters or options shown in normal brackets ([ ]) are required and must be provided. A comma followed by ellipses (...) indicates that the preceding parameter can be repeated multiple times.

Thus, if you wanted to add an empty partition to a partitioned table named SALES, you could do so by executing an ALTER TABLE statement similar to this:

```
ALTER TABLE sales ADD PARTITION q4_sales
  STARTING '10/1/2010' ENDING '12/31/2010' IN tbsp3;
```

When such a statement is executed, a new table is created and logically made a part of the partitioned table. The system catalog table SYSCAT.DATAPARTITIONS is then updated to reflect the change.

Attaching a populated table to a partitioned table

Tables that already contain data can be added to a partitioned table as new partitions by executing the ALTER TABLE statement with the ATTACH PARTITION option specified. The syntax used for this type of operation looks like this:

```
ALTER TABLE [TableName]
  ATTACH PARTITION <PartitionName>
  <STARTING <FROM> [Start | MINVALUE | MAXVALUE] |
    STARTING <FROM> ([Start | MINVALUE | MAXVALUE] ,...)
  <INCLUSIVE | EXCLUSIVE>>
  ENDING <AT> [End | MINVALUE | MAXVALUE] |
    ENDING <AT> ([End | MINVALUE | MAXVALUE] ,...)
  <INCLUSIVE | EXCLUSIVE>
FROM [SourceTable]
<BUILD MISSING INDEXES | REQUIRE MATCHING INDEXES>
```

where *TableName*, *PartitionName*, *Start*, and *End* are the same as before and *SourceTable* identifies the base table that is to be used as the source of data for the new partition.

So, if you wanted to attach a base table named Q4_2010_SALES to an existing partitioned table named SALES as a new partition, you could do so by executing an ALTER TABLE statement that looks like this:

```
ALTER TABLE sales ATTACH PARTITION q4_sales
  STARTING '10/1/2010' ENDING '12/31/2010'
  FROM q4_2010_sales;
```

CONDITIONS FOR ATTACHING A TABLE AS A NEW PARTITION

Before an existing base table can be added (attached) to a partitioned table, the following conditions must be met:

- ☑ The source table must be a nonpartitioned table or a partitioned table that has only one partition and no attached or detached partitions.
- ☑ The source table cannot be a typed table or a range-clustered table.
- ☑ The source table must be droppable.
- ☑ The source and target table definitions must match:
  - ✓ The number of columns must be the same.
  - ✓ The data types of each column must be the same. (For columns with XML or LOB data types, INLINE LENGTH values must be the same.)
  - ✓ The nullability characteristics of each column must be the same.
  - ✓ Any default constraints defined for columns must be the same.
- ☑ Source and target table compression specifications must match.
- ☑ Characteristics (table space type, page size, extent size) of the table spaces used by the source and target tables must match.
- ☑ If the target table is distributed across database partitions, the source table must also be distributed using the same distribution method.
- ☑ If the target table is a multidimensional clustering (MDC) table, the source table must also be an MDC table.
- ☑ XML data formats used by the source and target tables must be the same.

After a base table is successfully attached to a partitioned table, the system catalog is updated to reflect the newly attached partition while information about the original base table is removed from the catalog. No data movement is involved. This means that the process of attaching a new partition takes place very quickly—typically within a few seconds. (It's important to note that to attach an existing base table to a partitioned table, the base table's design must be similar to that of the partitioned table; see sidebar, "Conditions for attaching a table as a new partition.")

Anytime a base table is attached to a partitioned table, the partitioned table is automatically placed in "Set Integrity Pending" state. This means that integrity checking must be performed immediately after a new partition is attached. To check integrity after adding a base table named Q4_2010_SALES to an existing partitioned table named SALES (as a new partition), you would need to execute a SET INTEGRITY statement that looks similar to this:

```
SET INTEGRITY FOR sales
  ALLOW WRITE ACCESS
  IMMEDIATE CHECKED;
```

Removing partitions

Just as you can add or attach new partitions to a partitioned table, you can also remove existing partitions. (Removed partitions become regular, stand-alone base tables.) Partitions can be removed by executing the ALTER TABLE statement with the DETACH PARTITION option specified; the syntax for this statement looks like this:

```
ALTER TABLE [TableName]
  DETACH PARTITION <PartitionName>
  INTO [TargetTable]
```

where:

- **TableName:** Identifies the partitioned table, by name, that a partition is to be removed from
- **PartitionName:** Identifies the partition, by name, that is to be removed
- **TargetTable:** Identifies the base table where the partition's data is to be stored

Thus, if you wanted to remove a partition named Q1_SALES from a partitioned table named SALES and store its data in a table named Q1_2010_SALES, you could do so by executing an ALTER TABLE statement that looks like this:

```
ALTER TABLE sales
  DETACH PARTITION q1_sales
  INTO q1_2010_sales;
```

Once the table Q1_2010_SALES becomes available, you can either drop it (to delete the data) or attach it to another partitioned table (for example, to archive it).

DB2 9.7 table partitioning enhancements

Prior to DB2 9.7, indexes for partitioned tables were nonpartitioned. This meant that whenever partitions were added or removed, corresponding records had to be added to or removed from any indexes that existed for the table.

With DB2 9.7, partitioned indexes are now supported, allowing each data partition to be indexed separately. This provides several benefits, including improved performance when adding or removing partitions and the ability to perform select maintenance operations at the partition level. Starting with FixPack1, you can perform REORG operations on a single data partition rather than on the entire partitioned table, resulting in faster REORG operations while increasing availability to other data partitions.

Other table partitioning enhancements introduced with DB2 9.7 include XML support (available with FixPack1 or greater) and higher availability during ALTER TABLE ... DETACH PARTITION operations—the removal of a partition no longer requires that all current transactions to the table be completed, allowing long-running queries to continue without impeding the partition removal process.

Partitioning for performance

Data partitioning provides many useful benefits, particularly when used in data warehouse and decision support environments. By taking advantage of these features, you can manage data more effectively, while delivering optimum database performance. *

RESOURCES

Unleash the power of table partitioning in your DB2 warehouse: **ibm.com**/developerworks/data/library/techarticle/dm-1006tablepartitioning/index.html

DB2 9 table partitioning: **ibm.com**/developerworks/edu/dm-dw-dm-0612read-i.html

IBM Redbooks: Database partitioning, table partitioning, and MDC for DB2 9: www.redbooks.ibm.com/abstracts/SG247467.html?Open

DB2 for Linux, UNIX, and Windows area on developerWorks: **ibm.com**/developerworks/data/products/db2luw/

IBM DB2 9.7 for Linux, UNIX, and Windows Information Center: http://publib.boulder.ibm.com/infocenter/db2luw/v9r7/index.jsp

- **Adding data partitions to partitioned tables** http://publib.boulder.ibm.com/infocenter/db2luw/v9r7/topic/com.ibm.db2.luw.admin.partition.doc/doc/t0021566.html
- **Attaching data partitions** http://publib.boulder.ibm.com/infocenter/db2luw/v9r7/topic/com.ibm.db2.luw.admin.partition.doc/doc/t0021577.html

Does the Order of SQL Predicates Matter?

Rearranging WHERE clause predicates may be an exercise in futility

Bonnie Baker
(bkbaker@bonniebaker.com) specializes in teaching on-site classes for corporations, agencies, and DB2 user groups. She is an IBM DB2 Gold Consultant, an IBM Information Champion, a five-time winner of the IDUG Best Speaker award, and a member of the IDUG Speakers' Hall of Fame. She is best known for her ability to demystify complex concepts through analogies and war stories.

In this column I like to answer questions that I am asked most often via e-mail, at conferences, and by my class and seminar attendees. One of the top 10 questions is: "Does the order in which I code my predicates matter?"

Sometimes this question is followed by the proclamation that the questioner is required to adhere to a "shop standard" that details the order in which predicates should be coded. Two examples of these standards are:

1. Code join predicates first, followed by local predicates (predicates on a single table) in the same order as the named tables appear in the FROM clause.

2. The most-filtering predicates should be coded before the least-filtering predicates. (This is the standard that I hear most frequently.)

My response to this top 10 question is: Rearranging WHERE clause predicates may be an exercise in futility. You may occasionally be able to fine-tune some (but certainly not all) SQL to improve performance and reduce CPU usage by rearranging the predicates. In these rare situations, the results will not be noticeable if the SQL is executed infrequently and addresses very few rows. But, if the SQL is popular and is executed millions of times a day, addressing hundreds of thousands of rows each time, the cumulative total could be significant.

In case you're already wondering, in some SQL neither of the standards dictated previously will result in performance gains. That said, there is a standard you can use to ensure that your predicates are coded in an order that will never be detrimental and may actually be beneficial to performance.

So, let's look at how DB2 feels about the order in which your predicates are coded. For the purposes of discussion, our chosen index is a three-column index called IX1 created on (COLA, COLB, COLC).

Our popular, important SQL with its 10 predicates is:

```
Select …
  from big_table
Where colc = :hvc
  and colb > :hvb
  and cola = :hva
  and cold = :hvd
  and cole > :hve
  and colf in (:hvf1, :hvf2, :hvf3)
  and colg between :hvgbegin and :hvgend
  and colh + coli = :hvtotal
  and colj = :hv1
  and colj like :hvj
```

Now, let's get down to business.

In what order should your predicates be coded?

Sometimes it just doesn't make a difference. Why? Because DB2 is going to rearrange your predicates before applying them, and that new order may be exactly right for your SQL.

Most of the time you can take advantage of any order that makes the SQL more readable and easier to maintain without affecting performance. An example of such a "readability" order is mentioned in shop standard example 1 (code join predicates before local predicates).

When does the order make a difference?

At static and dynamic BIND time, the DB2 Optimizer parses your SQL and—regardless of the order in which you coded them—rearranges the predicates in a predetermined order. This order is based upon the filter-factor-driven approach, which says that predicates that filter out (that is, eliminate) the most rows the soonest should be applied before those that filter out the least (that is, qualify the most) rows.

In our earlier SQL, the parsing would result in our predicates being applied in the following order:

```
Where cola = :hva
  and colb > :hvb
  and colc = :hvc
  and cold = :hvd
  and colj = :hv1
  and cole > :hve
  and colg between :hvgbegin and :hvgend
  and colf in (:hvf1, :hvf2, :hvf3)
  and colj like :hvj
  and colh + coli = :hvtotal
```

How did we end up with that arrangement? To start, DB2 always applies index predicates first, following the order in which the index is created. Our chosen index is on (COLA, COLB, COLC), so those predicates—the first three lines in the preceding example—will be applied in that order. If you want those three predicates applied in the exact order in which you code them, you must drop and re-create the index (or create a new index) on (COLA, COLC, COLB).

Second, regardless of the order in which they are coded, Stage 1 non-index predicates are applied in the following order:

1. Equal predicates
2. Range predicates
3. In-list and like predicates

DB2 took our Stage 1 non-index predicates and put the two equal predicates before our two range predicates (> and between), which are placed before our two in-list and like predicates.

Finally our Stage 2 predicate (and colh + coli = :hvtotal) is applied. Why last? Because now those higher-CPU, mathematical predicates will be applied to fewer rows; that is, only the rows that are left after applying the other nine predicates.

You may have noticed that in the parsed SQL, the two equal predicates will be applied in the order in which they were originally coded. Likewise, the two range predicates will be applied in their coded order, and the in-list and like predicates were left in the order originally coded.

Besides creating your index in a different order, what can you do?

The like-kind non-index predicates are applied in the order coded. Therefore, if you think the COLI predicate should be applied before the COLD predicate (to filter out more rows sooner), you must code the two predicates in COLI, COLD order. If you want the between predicate applied before the > (greater than) predicate, you must code the predicates in COLG, COLE order. And if you want the like predicate applied before the in-list predicate, your coding order must reflect that.

What I'm saying here is that for non-index predicates, DB2 does not consider the COLCARD or the COLUMN DISTRIBUTION statistics. Those statistics are used to aid in index selection, not for predicate rearrangement.

What standard should you use?

Based on all of this discussion, it would appear that the standard that advocates coding the most-filtering predicates before the least-filtering predicates is desirable. However, that standard may, in some situations, degrade performance for our popular SQL. What if we added the following non-index predicates to our SQL:

```
Where … our first 10 predicates
  and colm = :hvm
  and colw = :hvw
```

In this case, coding the most-filtering predicate (the one with the higher column cardinality) before the least-filtering predicate works just fine because we want to disqualify as many rows as possible as soon as possible. But what if our non-index predicates look like this:

```
Where … our first 10 predicates
  and ( colm = :hvm
    or  colw = :hvw )
```

In that example, coding the most-filtering predicate before the least-filtering predicate will work to our disadvantage. Some would think that with OR logic, the order of the predicates wouldn't make a difference since they both must be applied anyway. But the truth is that it does matter, because both do *not* always need to be applied.

You see, as soon as the row qualifies, predicate application within the OR list stops. The second predicate is applied only when the row is disqualified by the first predicate. Therefore, unlike AND logic, with OR logic we want to code the least-filtering predicate before the most-filtering predicate. To say it another way: we want to code the more-qualifying predicate before the least-qualifying predicate.

Is there a standard that is appropriate for all SQL?

Yes. Create your indexes in the order in which you want those index column predicates applied. Then code your SQL with AND predicates coded with most-filtering predicates before least-filtering predicates and OR predicates coded the opposite way, with least-filtering predicates coded before most-filtering predicates.

Keep something else in mind: if you really know your data and your search criteria, you may know that some predicates filter out more than others regardless of COLCARD, because the search values are highly distributed. With host variables or literals, DB2 does not consider this factor for non-index predicates; only you can take this fact into account and code accordingly.

Stay on top of your standards

Coding standards should be readdressed with each release of DB2 and with each learning curve. A new release or a newly learned fact can make you rethink those old rules.

With that in mind, my next column will be a rewrite of a very popular past column in which I explained how predicting the order of your result rows is not as simple it sounds. The latest releases of DB2 have altered and added to my opinions on this subject, and I want to share those thoughts with you.

RESOURCES

DB2 for z/OS: ibm.com/db2/zos

***IBM Data Management* magazine archive: ibm.com**/developerworks/data/dmmag/archive.html

Building Fast Data Warehouse Schemas: Part 3

Design speedy summary tables

Lester Knutsen *(lester@advancedatatools.com) is president of Advanced DataTools Corporation, an IBM Informix consulting and training partner specializing in data warehouse development, database design, performance tuning, and Informix training and support. He is president of the Washington, D.C., Area Informix User Group; a founding member of IIUG; an IBM Gold Consultant; and an IBM Data Champion.*

My last two columns discussed why and how we use dimensional modeling to organize the data in a data warehouse. To briefly reiterate, the goal of a data warehouse is to give users fast access to data so they can make better business decisions, and to give them a separate area to run queries and play with data so they don't affect the operational database. The goal of dimensional modeling is to optimize warehouse data for easy queries and analysis by business users.

There are three key parts of a dimensional model: fact tables, dimension tables, and summary tables. We have already covered fact and dimension tables, so in this column, we will look at how to design summary tables.

Summary tables are an important part of creating a high-performance data warehouse. A summary table stores data that has been aggregated in a way that answers a common (or resource-intensive) business query. Summary tables are all about speed. They're smaller than fact tables, which means they generally respond more quickly (fewer rows to query), and they deliver answers without calculating every result from scratch.

Most of the time the summary table contains the results of an SQL SUM function, but you can also use the COUNT, MAX, MIN, or AVERAGE functions as needed. The COUNT function is very useful to show at any level how many records make up each aggregated row in the summary table.

Ideally, the structure of the data warehouse offers a summary table for any question that business users want to ask. From the summary table, they can then pick the dimensions that they want to drill down into and see more detail. Users should be able to drill down from one summary table into another summary table, ultimately arriving at the fact table (if that's the level of detail they need).

A well-designed data warehouse has a pyramid of summary tables, each one containing more detail until you arrive at the fact table, at the bottom of the pyramid. The summary table at the top of the pyramid is small, so queries on it are very fast. In fact, the goal is to have 90 percent of your queries running against the summary tables so the user gets very fast responses. As the user drills down into more detail, each summary table contains more rows, until the user reaches the fact table (see Figure 1).

For example, when I built a financial data warehouse for a government agency many years ago, they had a fact table with millions of records containing every debit and credit transaction in their financial system for several years. A single payment could have had 6 to

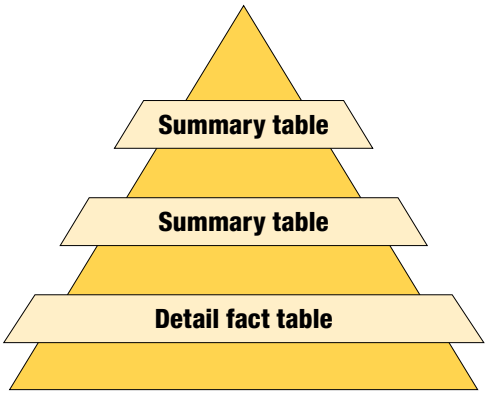

Figure 1: *Pyramid of summary tables*

12 debit and credit transactions. However, most users did not want to see this level of detail; they needed to see only the income and expense results of these transactions.

We built a summary table that aggregated the income and expense detail to show the end result for each transaction. This table still contained millions of records, so we built a summary table to aggregate the detail at a monthly level, a table to aggregate the detail at a year-to-date level, and a final summary table for the division level. This prestored, aggregated data allowed agency employees to see results very quickly at the level they needed. When they found issues that required more detail, they could then drill down into another summary table or the detail fact table.

As you might guess, summary tables are the starting point for most canned reports and dashboards. Summary tables are like data cubes, and a cube is basically a separate software product or technology that creates summary tables. Hence, you can do everything with summary tables that you do with cubes.

Building summary tables

One way to build summary tables is to put an SQL trigger on the fact table, so every time a record gets inserted into the fact table, the trigger updates the summary table. That way you never have to worry about building the summary table again and again. Another way is to have the extract, transform, and load (ETL) process update the summary table as it loads the fact table. A third method is to have the summary tables rebuilt after the ETL process using SQL SUM and aggregate functions. Higher-level summary tables in a pyramid of tables may be built off lower-level tables, rather than off of the fact table.

A key design decision is the grain of aggregate tables. When you start with a fact table that has billions and billions of records, you need to plan summary tables in a hierarchy that summarizes the detail with the goal of reducing the record count by a factor of 100,000 to 1,000,000. Your top-level summary table may have about 100,000 records; the next-level summary table will probably have millions of records. Finally, your detailed fact table will have billions of records.

Comparing across the grain

Another benefit of summary tables is that they let you compare results from two different fact tables (maybe even in separate data marts) that have different grains. Matching up grain levels when building comparisons is a critical concept in designing a data warehouse. One project I developed had financial data at the detail grain transaction level and program accomplishments at a day grain level. It was not possible to compare dollars spent with accomplishments without summarizing the dollars to the same grain as that of the program level. We built one summary table that contained dollars and program accomplishments at the same level, so users could see what it cost per program accomplishment and which offices were the most productive for the lowest costs.

This is how your users are going to make decisions—by comparing facts from different fact tables with different grains. No matter what data your users want to compare—budgeted expenses to actual, planned income to actual, projected sales to actual, or sales to inventory—the data must be summarized to the same grain so the results are meaningful. If you don't summarize data to the same level of detail, users run the risk of reporting on incorrect data or drawing incorrect conclusions from the data.

Designing a dimensional data warehouse is a complex subject, and I have just touched the surface of it in these three articles. I would like to share some other resources that offer more information (see sidebar, "Resources"). The first is from IBM Redbooks, and I am really excited to see it published: *Data Warehousing with the Informix Dynamic Server*. You can download it from the IBM Redbooks Web site.

A second resource I have found helpful is a manual titled *Designing and Implementing a Database,* available in the IBM Informix Dynamic Server 11.5 Information Center and as a separate download. The entire Section 4 is on designing dimensional databases, and the manual offers a very good overview of dimensional data modeling and implementing a dimensional database design with Informix.

I also did a series of Webcasts earlier this year on this subject, available on the Advanced DataTools Web site. Happy designing!

RESOURCES

Data Warehousing with the Informix Dynamic Server:
www.redbooks.ibm.com/abstracts/sg247788.html

Designing and Implementing a Database:
http://publib.boulder.ibm.com/infocenter/idshelp/v115/topic/com.ibm.ddi.doc/ddi.htm

Data warehouse design Webcasts:
www.advancedatatools.com/Solutions/DataWarehouseWebcast.html

Informix home page:
ibm.com/informix

Information On Demand 2010

The Premier Forum for Information and Analytics

IBM invites the readers of *IBM Data Management* magazine to visit the sponsors of this year's event.

Intel, a leader in silicon innovation and global manufacturing, develops technologies and initiatives to advance how people work and live. Intel® Xeon® processors deliver enhanced, energy-efficient performance for data-intensive business applications. Combined with the latest IBM® System x® servers, Intel Xeon processors provide outstanding performance for IBM Information Management software.
www.intel.com/xeon

SAP is a leading global provider of business software that enables companies of all sizes and in more than 25 industries to become best-run businesses.
www.sap.com

Netezza Corporation (NYSE: NZ) is the global leader in data warehouse, analytic and monitoring appliances that dramatically simplify high-performance analytics across an extended enterprise.
www.netezza.com

Compuware Corporation
www.compuware.com

CONSIST
www.consist.com

CSC
www.csc.com

Dassault Systèmes
www.3ds.com

Deloitte
www.deloitte.com

enChoice, Inc.
www.enchoice.com

Esri
www.esri.com

Fourth Millennium Technologies
www.fmtusa.com

HCL
www.hcltech.com

Infosys Technologies Ltd
www.infosys.com

JCB Partners
www.jcbpartners.com

Jeskell, Inc.
www.jeskell.com

LPA Systems
www.lpa.com

Mellmo
www.roambi.com

NetApp
www.netapp.com/ibmnseries

Niteo Partners
www.niteo.com

Perficient
www.perficient.com

Pronto Software
www.pronto-software.com

QueBIT Consulting
www.quebit.com

Renewtek Pty Ltd
www.renewtek.com

Sogeti
www.sogeti.com

SYSCOM, Inc.
www.syscom.com

Thunderhead
www.thunderhead.com

TriTek Solutions
www.triteksol.com

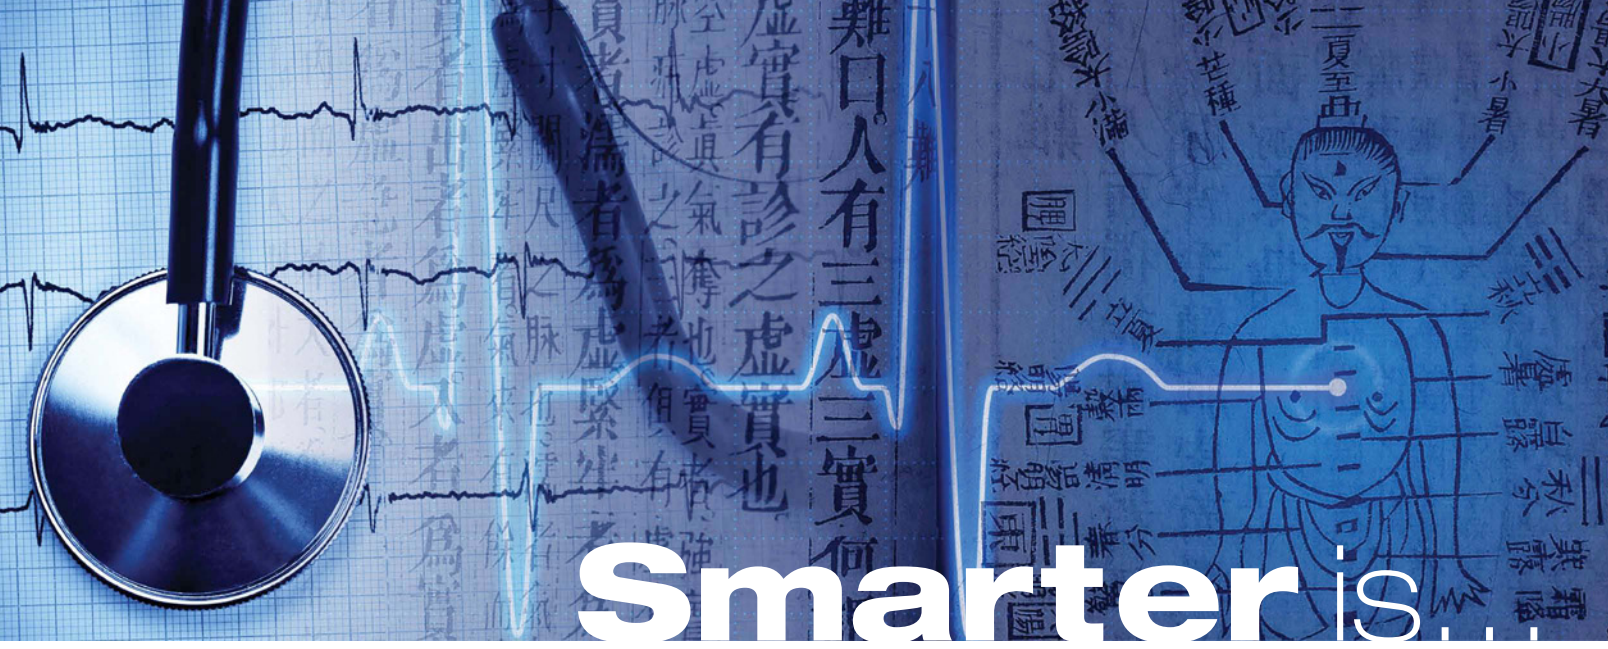

Breaking Data and Disease Down to Size

By David Geer

Chronic kidney disease (CKD) is the ninth leading cause of death in the United States, according to the Centers for Disease Control and Prevention.[1] Modern Western medicine has had little impact in staying the disease's deadly hand.

But Eastern doctors have found hope for patients with CKD in traditional Chinese medicine (TCM). Over the years TCM has proven effective in delaying the need for kidney dialysis. Guang Dong Hospital of TCM in Southern China treats 10,000 kidney patients a year; if doctors can establish the combinations of treatments that are consistently effective for different categories of CKD patients, they can apply evidence-based treatment plans with confidence, helping to save lives.

Although the Guang Dong Hospital of TCM has access to millions of electronic medical records, much of the patient data will, in the future, be captured in large XML documents that are difficult and time-consuming to sift through and analyze, according to Sarah Knoop, manager for Healthcare Systems Research at IBM Research–Almaden, in San Jose, California.

To address this problem, IBM and the Guang Dong Hospital teamed up to create Healthcare Information Warehouse for Analysis and Sharing (HIWAS), a prototype plug-in–based tool for IBM InfoSphere Warehouse. The hybrid solution extracts the existing, complicated XML data and enables it to be converted into more easily analyzed formats, says Knoop. "[These formats are] the most appropriate, efficient storage representation for the business problem at hand leveraging DB2 technology," she says.

"For clinicians to use this data in the clinical setting, they have to be able to extract it in order to analyze it," explains Knoop. "However, doing so can be difficult because these XML models are typically complex and more general."

Knoop also says that HIWAS helps extract valuable information for further analysis so it can be used by IBM Cognos analysis tools. HIWAS converts the excised data into warehouse schemas for hybrid relational-XML or pure relational databases, Knoop explains. Cognos is then able to analyze the data and create reports that highlight the significant relationships between the data. Doctors use those relationships to build a base of clinical evidence that reveals best practices around treatments for CKD, according to Knoop.

With HIWAS, doctors can analyze the data and categorize CKD patients according to those who became sick because of other diseases or in combination with other chronic diseases. This will help doctors understand which combinations of Western treatment plans and TCM approaches are most effective for certain categories of patients, says Lu Yu Bo, president of Guang Dong Hospital of TCM.

"As clinical doctors, we can use IBM technology to find more useful information from the patients' cases," says Mao Wei, professor of nephrology at Guang Dong Hospital of TCM. "This useful data can help us to improve our effectiveness in traditional Chinese medicine in the treatment of chronic disease."

David Geer *is an Ohio-based technology journalist.*

[1] Albright, Ann. CDC Congressional Testimony, July 1, 2010. "The Battle Against Diabetes: Progress Made, Challenges Unmet." www.cdc.gov/washington/testimony/2010/t20100701.htm

Hardware sales guys hate us. **Your bottom line will love us.**
Get Remarkable Breakthrough Results for DB2 LUW.

Do you let the fox into your henhouse?

That hardware company will tell you that you need yet another server to tune the database software they sold you.

Isn't that a conflict of interest?

DBI uses patented technologies to help DBAs uncover and isolate sources of excessive CPU utilization, I/O bottlenecks, and poor response times. Sophisticated proactive alerts further assure your business sustains peak performance at minimal operational costs.

Let us show you how easy effective database tuning can be with measurable results and instant ROI in just five mouse clicks.

dbi

Your Performance IS Our Business

DBI is the recognized leader in database performance tuning and optimization. Our unique products enable DBAs to adopt a proactive methodology - escaping the trap of simple "reactive" tools. Visit www.DBISoftware.com/henhouse or call toll-free 1-866-773-8789 (outside USA: +1-512-249-2324)

TDA GROUP
800 West El Camino Real, Suite 101
Mountain View, CA 94040
U.S.A.

POWERFUL SCALABLE RELIABLE & EFFICIENT

Come see us in **IOD Booth #702**

IBM System x servers and DS5000 storage systems for your database applications

Data center managers need infrastructure systems that meet today's needs as well as tomorrow's.

They need:

- *Integrated server and storage solutions that deliver high performance, strong value, and fast ROI*
- *Intelligent infrastructure options that scale transparently, on demand*
- *Highly utilized, power- and space-efficient systems that lower TCO*
- *Dependable, highly reliable solutions with outstanding support from a trusted partner*

IBM System x3690 X5

IBM System x3850 X5

IBM System Storage DS5300

IBM System Storage EXP5060

IBM System x Servers and DS5000 Storage Systems *Deliver*

IBM

For more information visit
www.ibmdatamanagement.com