ISSUE 2, 2010 / VOLUME 15 NO. 2

IBM

data management

IBM.COM/DMMAGAZINE

KNOWLEDGE. PERFORMANCE. RESULTS.

QUALITY CONTROL
Why data governance matters

INFORMIX DBA
Optimizing warehouse schemas

Champions of Finance

Meet two data professionals committed to community

Smarter technology for a Smarter Planet:

Finding meaning in the noise.

An unprecedented amount of information flows through companies every day. But to what effect? A recent study found that 52% of managers have no confidence in the information they rely on to do their job. And 42% of them actually use the wrong information at least once a week. Without the right approach to business intelligence, companies struggle to turn all that information into sound decisions.

IBM business intelligence and performance management solutions give you the smarter tools you need to access the right information, making it available to the right people when and how they need it. Today IBM is helping over 20,000 companies spot trends, mitigate risk and make better decisions, faster. In fact, we helped a major retail supplier achieve this by cutting their average financial reporting time by almost 50%.

A smarter business needs smarter software, systems and services.
Let's build a smarter planet. ibm.com/intelligence

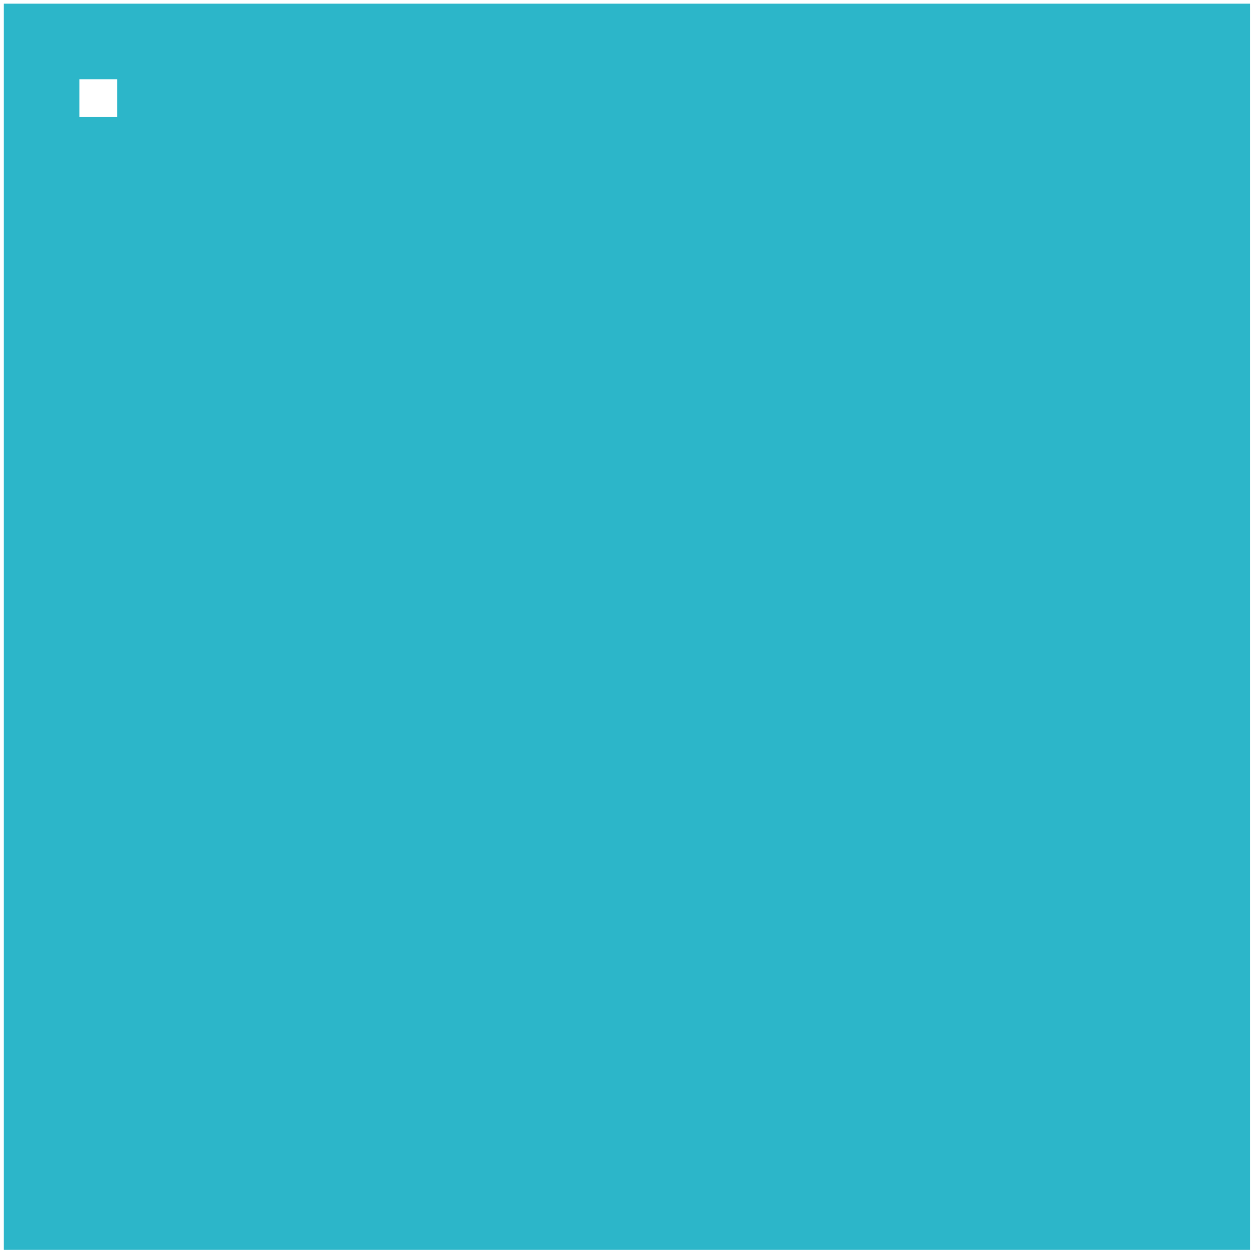

IBM
IBM.COM/DMMAGAZINE
IBM data management
KNOWLEDGE. PERFORMANCE. RESULTS.
ISSUE 2, 2010 / VOLUME 15 NO. 2

Features

Photos of Manuel Gómez Burriel by Denis Doyle. Photos of Banu Ekiz by Ali Kabas.

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

Departments

Ad Index

Register Today

for the IBM Information Management Global Conference!

InformationOnDemand2010

The Premier Conference for Information and Analytics

October 24–28, 2010 Mandalay Bay | Las Vegas, Nevada

IBM Information On Demand 2010 is the Premier IBM Information and Analytics Event. Drive innovation and business optimization with the best in technical and business education. More than 300 customers will share their real-world experiences about how they are transforming their businesses through information and analytics.

Register Early and Save!

May 3 - August 31	$1,895 ($500 savings)
September 1 - October 23	$2,195
October 24 - October 28	$2,395

2010 Event Highlights

- More than 9000 conference attendees
- 300 customer speakers
- 110 Business Leadership sessions
- 440 Technical Skill Building sessions
- 3 free Information Management certification tests
- Huge EXPO: 200 Business Partners and 120 IBM Exhibits

Register Today for IBM Information On Demand 2010

- Gain deep technical education and the best strategic insight and analysis.
- Explore the latest advances in IBM Hardware, Software and Services solutions – more usability sandboxes and hands-on labs.
- Hear from industry leaders who are leveraging information and analytics to realize and build competitive advantage in today's economic environment.
- Make an immediate impact on your organization by learning how to take the next step in transforming your business through information and analytics.
- Interact with your peers, industry experts, Business Partners and IBMers who share your interests and can provide support and insights.

ibm.com/events/InformationOnDemand

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, Howard Baldwin, Bonnie Baker, David Beulke, Robert Catterall, Elizabeth Dial, Lester Knutsen, Stuart Litel, Kent Milligan, Paul Pendle, Roger E. Sanders, Chris Young

EDITORIAL BOARD OF DIRECTORS
Debra J. Black, Cathy Elliott, Kim Evans, Jeff Jones, Kimberly C. Madia, Nancy Miller, Kathryn Zeidenstein

AD SALES
advertise@tdagroup.com

AD COORDINATOR
Katherine Hartlove

SPECIAL THANKS TO
Lea Anne Bantsari

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.

Here in the Silicon Valley offices of *IBM Data Management* magazine, we are no strangers to the IBM Information Champions program. With only one exception (who I'll come back to in a moment), all of our regular columnists are long-standing members: Roger E. Sanders, Stuart Litel, Lester Knutsen, Robert Catterall, David Beulke, and Bonnie Baker.

Information Champions do more than make their day-to-day living as data management professionals; they lead the way for the rest of us. They teach; give seminars; lead user groups; moderate forums; advise their colleagues; and write books, blogs, columns, and just about anything else you can think of that would share their knowledge. These are leaders in the community and some of the brightest lights in the IBM Information Management sky.

That's why we're pleased to be able to present this issue's cover story, featuring two Information Champion honorees. Banu Ekiz and Manuel Gómez Burriel work on opposite ends of the European continent, but their geographic separation fades beside the similarities in their work—both are in the financial industry—and their efforts to cultivate and expand the DBA and data management community.

As I mentioned earlier, only one of our regular columnists isn't an IBM Information Champion. But Merv Adrian is a data management leader in his own right, and he uses his decades of experience to look at data and information from a more strategic perspective. In this issue, he grills a panel of IBM experts on data governance with the goal of attaching some real-world meaning to a slippery, abstract concept.

Finally, if you haven't visited the digital edition of *IBM Data Management* magazine recently, you're in for a treat. We've rolled out a redesigned interface that streamlines and simplifies the navigation, making it easier for you to explore the issue and share articles. Check it out at www.ibmdmmagazinedigital.com and let us know what you think!

Thanks for reading,

Cameron Crotty
Editor

EMC²
where information lives®

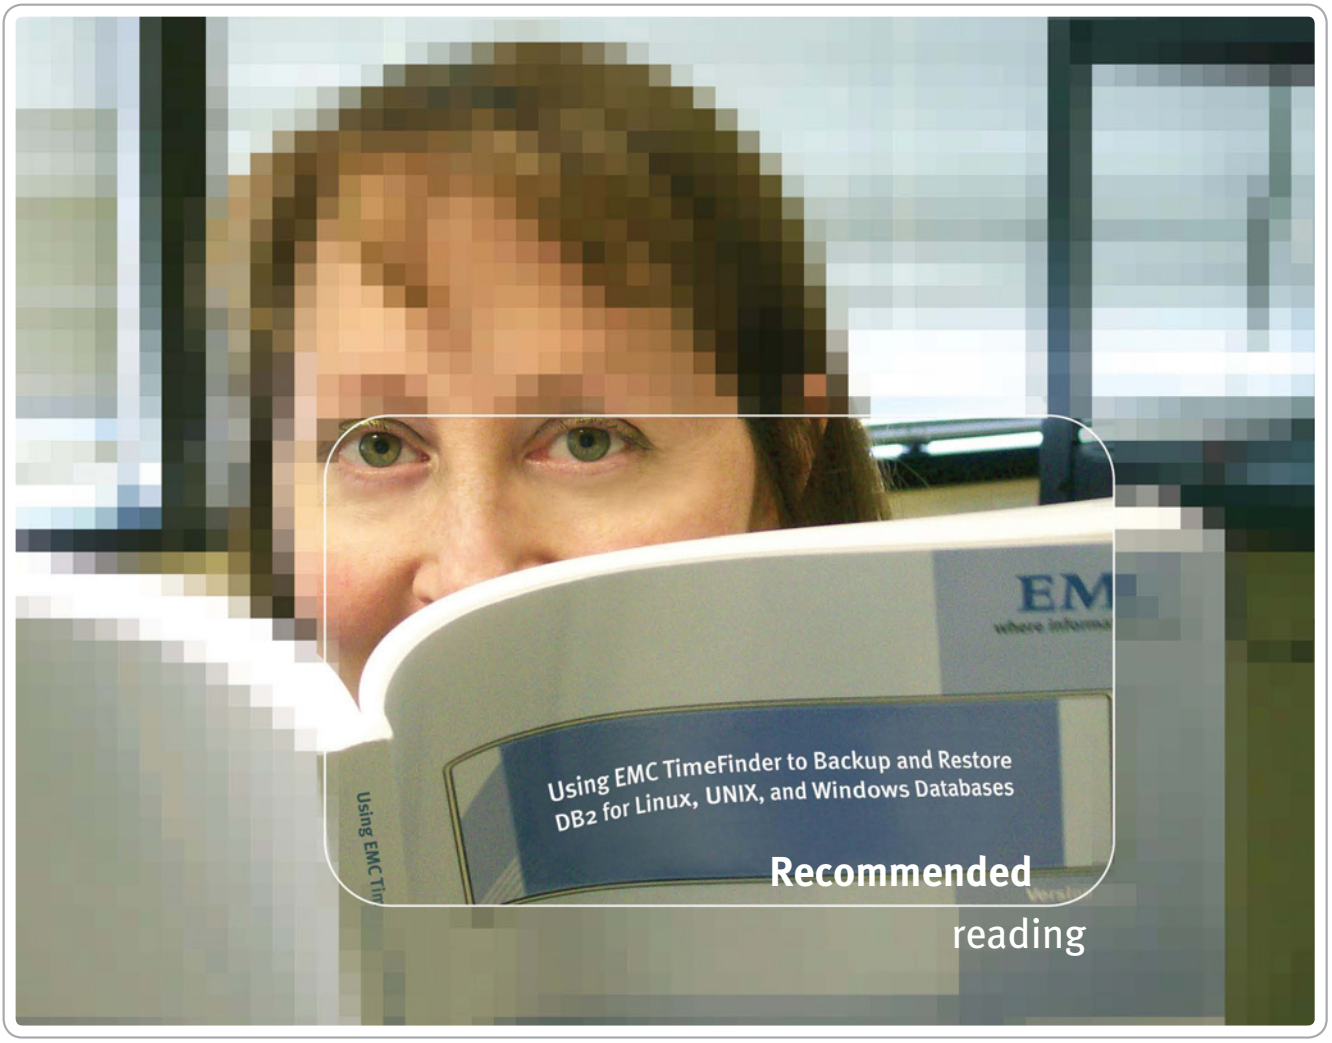

EMC and IBM join forces to help you back up mission-critical DB2 databases.

If you're dealing with mission-critical databases, your first mission is to secure a copy of this must-read EMC® TechBook. Co-authored by IBM engineer Dale McInnis and EMC engineers Roger E. Sanders, Paul Pendle, and Michele Backstrom, this highly technical exploration will show you why EMC solutions are ideal for backing up and restoring mission-critical databases. You'll get the best practices you need to back up and restore DB2 databases on EMC storage efficiently and effectively—direct from four of the brightest engineering minds around. It's information so critical to your mission, you'll want to read it again and again.

Learn more at www.EMC.com or www.vervante.com.

Discover these other EMC TechBooks:
DB2 for Linux, UNIX, and Windows on EMC Symmetrix Storage Systems
DB2 for z/OS Using EMC Symmetrix Storage Systems
Deploying DB2 for Linux, UNIX, and Windows on EMC Symmetrix Arrays

EMC²
TECHBOOKS

DB2 10 for z/OS Beta Program Begins

New version designed for improved performance, scalability, and cost savings

Building on the formidable capabilities of IBM DB2 9 for z/OS and the IBM System z platform, the next version of DB2 for z/OS has moved one step closer to release. Beginning in March, IBM made DB2 10 for z/OS available to selected organizations in a closed beta program. The new version includes features that can improve performance, increase the resiliency of business-critical data, and accelerate application and warehouse deployment.

Internal testing at IBM indicates that customers may realize significant out-of-the-box CPU savings with DB2 10, depending on the type of workload. Additional savings and higher performance can be achieved by taking advantage of features including hash access, index include columns, in-line large objects, parallel index updates, faster single-row retrieval, in-memory work files, index list prefetch, and a 1 MB page size for System z10 buffer pools.

The new version of DB2 can also support up to 10 times as many users per data-sharing group member than was possible with previous versions of DB2. Depending on the configuration and situation, this can mean that fewer data-sharing members are needed, greatly increasing the number of users that can be supported on reduced memory and CPU resources.

Maintaining the resiliency of business-critical information is another area of emphasis in this release. More data access, data management, and data definition functions can be handled concurrently, reducing overall downtime. Schema evolution, or data definition on demand, enables developers and DBAs to change the structure of the database in certain ways without requiring DB2 to come down—and fewer planned outages result in higher availability. DB2 10 also has several features aimed at enhancing security and simplifying regulatory compliance, including data masking, more granular administrative privileges, and new audit capabilities.

Enhancements for rapid application and warehouse deployment to support business growth round out the feature set for DB2 10, including support for temporal or versioned data, which facilitates a more automated process for rolling old data into archives over time. A number of SQL enhancements make the DB2 10 SQL implementation more consistent with other members of the DB2 family and make it easier to adapt SQL applications originally written for other database platforms. Examples include a 64-bit ODBC driver, Currently Committed locking, implicit casting or loose typing, time stamp with time zone, and moving sum and average.

DB2 10 also features enhancements to DB2 pureXML that focus on improving performance and usability of XML data applications. These include schema validation in the engine, a binary XML exchange format, multiversioning, easy update of subparts of XML documents, stored procedures, user-defined functions and triggers, and XML index matching with date/time stamp.

Query Management Facility (QMF), the DB2 query and reporting facility, also contains many updates for DB2 10, including more than 140 new analytical functions and support for HTML, PDF, and Flash report output.

Migration to DB2 10 is supported from both DB2 9 and from DB2 for z/OS V8.

MORE INFORMATION
ibm.com/software/data/db2/zos/db2-10

Power Systems with POWER7

IBM recently unveiled Power Systems running on new POWER7 processors. Data management products integrated with POWER7 technology leverage its performance, virtualization, and data integrity to deliver optimum workload performance.

The POWER7-based systems can reduce memory costs and energy needs and enable physical memory to be logically expanded by up to 100 percent on some workloads, including SAP. A single processor system can support up to 1,000 virtual machines for Linux, AIX, and IBM i. The POWER7-based systems help DBAs quickly and easily migrate applications to eliminate system outages, and they offer a set of new automated tools that can simplify system management.

MORE INFORMATION
ibm.com/power7

IBM Information On Demand 2010

The Countdown to IOD EMEA Is On!

Plan to attend the Information On Demand (IOD) Europe, Middle East, Africa (EMEA) Conference 2010, to be held May 19–21, in Rome, Italy. IOD EMEA promises to be the premier Information Management event in the EMEA region, offering leading technical and business leadership programs, an unrivalled forum to share best practices, opportunities to network with peers, and the chance to learn more about data management products and features. Register now to reserve your spot at the conference: **ibm.com**/software/uk/data/conf

IBM Optim Performance Manager Extended Edition **Speeds Analysis**

IBM recently announced Optim Performance Manager Extended Edition, which helps DBAs optimize the performance and availability of DB2 and DB2 applications with a proactive, end-to-end performance management approach.

Optim Performance Manager Extended Edition for DB2 on Linux, UNIX, and Windows (LUW), formerly known as DB2 Performance Expert, boasts a new architecture that makes problem isolation and analysis available anywhere, anytime, and faster than ever. The new browser-based interface includes performance overview displays with associated health indicators to quickly detect problems both in the overall environment and within a specific database. Intelligent diagnostic dashboards provide metrics that help DBAs focus on a particular problem area and quickly see the details needed to determine root cause.

This release expands upon the capabilities of DB2 Performance Expert by monitoring end-to-end database transactions to isolate response-time problems and ensure that service-level agreements are met.

In addition to supporting Java applications, Optim Performance Manager Extended Edition V4.1 now supports DB2 call-level interface (CLI) applications, and it provides predefined application views for IBM WebSphere Application Server, SAP, IBM Cognos, IBM InfoSphere DataStage, and IBM InfoSphere SQL Warehouse tasks.

The Extended Edition also features administrative and management tooling for DB2 Workload Manager (WLM), which provides the ability to quickly and proactively administer, manage, and monitor the workload in context with workload management settings. New integration between Optim Performance Manager Extended Edition and IBM Tivoli Composite Application Manager enables a consolidated view of the business transactions across the enterprise, with the comprehensive detail necessary to diagnose database-specific activities.

MORE INFORMATION
ibm.com/software/data/optim/performance-manager-extended-edition

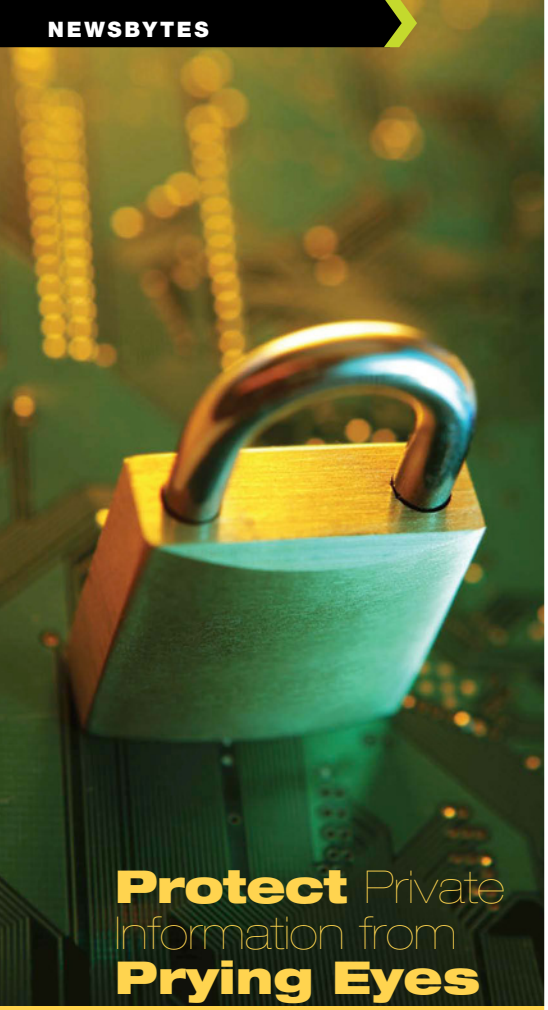

Protect Private Information from Prying Eyes

Almost every organization has sensitive information that must be kept confidential. But what happens when that data appears on forms and documents that must be distributed as part of normal business processes? No need to pull out the black marker and strike through private information: a new tool, IBM Optim Data Redaction, can help organizations protect sensitive data by automatically recognizing and removing sensitive content from documents and forms.

For example, a customer's credit scores in a loan document could be hidden from an office clerk, while still being visible to a loan officer.

Optim Data Redaction can help organizations reduce the cost of compliance, balance automated extraction with human review, prevent unintentional data disclosure, and control the data viewed by each user. The new tool, part of the IBM Optim Integrated Data Management portfolio, is currently available.

MORE INFORMATION
ibm.com/press/us/en/pressrelease/29316.wss

IMS Tools Solution Packs Help Get the Job Done

On February 9, IBM announced four new integrated IMS Tools Solution Packs, providing end-to-end solutions designed to help organizations save time and money. The IMS Tools Solution Packs can deliver faster return on investment by cutting CPU consumption, DBA labor costs, and application downtime.

IMS Database Solution Pack for z/OS combines features for managing IMS full-function databases, including High Availability Large Databases (HALDBs). It provides high-performance load, unload, index build, reorganize, backup, verify, and report capabilities.

IMS Fast Path Solution Pack for z/OS delivers an extensive set of utilities to help database administrators analyze, maintain, and tune IMS Fast Path databases. With this solution pack, many key functions can be performed without taking the IMS database offline.

IMS Performance Solution Pack for z/OS provides a comprehensive portfolio of database performance management tools. This solution pack tightly integrates the capabilities of IMS Connect Extensions, IMS Performance Analyzer, and IMS Problem Investigator, simplifying and accelerating end-to-end analysis of IMS transactions.

IMS Recovery Solution Pack for z/OS provides features, functions, and processes to implement best-practices backup and recovery scenarios. The solution pack combines IMS Recovery Expert and IMS Database Recovery Facility features and functions, and includes high-performance image copy, change accumulation, and index-building capabilities.

To support the solution packs, IMS Tools Base provides a consolidated installation package for the entire collection of IMS tools.

MORE INFORMATION
ibm.com/software/data/db2imstools/products/ims-tools.html

Learn Data Management on Your Schedule

Study when and where you want with self-paced virtual classes

Need to brush up on the constantly growing IBM Data Management portfolio? Now there's an option that delivers the same content, interactive exercises, and hands-on labs as IBM classroom training while allowing you to set your own schedule.

Self-paced virtual classes are 30-hour to 80-hour online courses that involve Web-based reading assignments, support from instructors and subject-matter experts, and multimedia presentations based on the same content used in IBM classroom courses, plus interactive exercises and hands-on lab assignments using the actual products.

New courses that cover data warehousing, Informix, DB2, and Optim include:

- DB2 9 Administration Workshop for Windows
- DB2 9 for LUW Quickstart for Experienced Relational DBAs
- DB2 9.7 for LUW Quickstart for Experienced Relational DBAs
- Fast Path to DB2 9 for Experienced Relational DBAs
- Informix Dynamic Server 11 System Administration
- Developing Applications Using Informix 4GL
- Implementing Web Services Using Informix 4GL
- InfoSphere Warehouse 9 Components

MORE INFORMATION
ibm.com/software/data/education/elearning.html

Information Governance: A New Perspective

Information governance is a quality-control discipline for assessing, managing, using, improving, monitoring, maintaining, and protecting organizational information. Effectively leveraging information requires governance to guarantee that the information is reliable, accurate, and authentic. Recently, IBM laid out its unique approach to this topic, looking at information governance from an information supply chain perspective. The IBM Information Governance initiative aids businesses in sustaining profitable, measurable growth by ensuring the viability of their information. This strategy aims to maximize the value of information by identifying potential issues before they happen.

As part of this initiative, IBM announced two new data protection software products: IBM InfoSphere Business Information Monitor and IBM Optim Data Redaction.

InfoSphere Business Information Monitor, based on Guardium technology, tracks the flow of information throughout the enterprise and alerts companies of potential flaws in information. Optim Data Redaction protects information by automatically detecting sensitive content and removing it from documents.

In addition, IBM Global Services detailed a number of expanded service offerings for information governance and formally announced its Information Governance Center of Excellence within the Business Analytics and Optimization area.

› MORE INFORMATION
ibm.com/software/info/itsolutions/information-governance

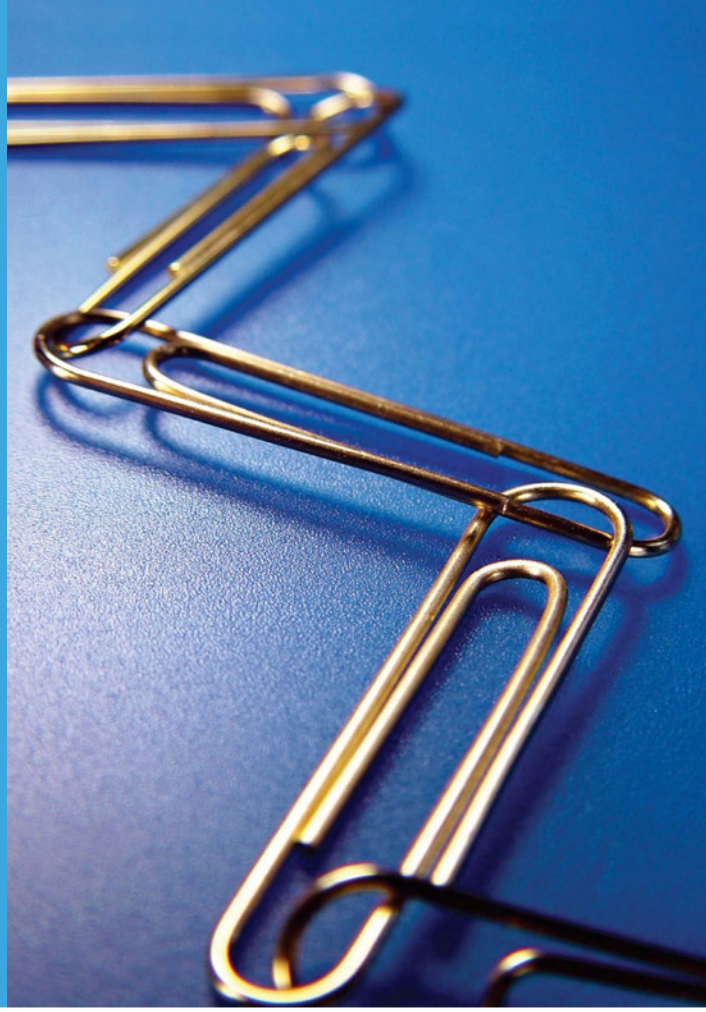

Solving the DB2 Management Puzzle is Easy

Quest Management Suite Simplifies DB2 Administration

Managing your complex DB2 environment just became easier, thanks to Quest® Management Suite for DB2 LUW. It combines all of the functionality you need for DB2 on Linux, Unix, and Windows, so you can manage multiple aspects of DB2 administration and performance with a single solution. Whether it's administration powered by Toad®, SQL tuning, workload analysis and trending or real-time diagnostics, you've got the tools you need to quickly and easily manage and control your day-to-day DB2 activities.

Watch the pieces of DB2 management fall into place with Quest.

Read our technical brief, *"The Top 10 Things a DBA Should Know About Toad® for DB2"* at **www.quest.com/puzzle**

QUEST SOFTWARE®
Smart Systems Management

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.*

Here Comes **Informix**

Get ready to feel the love as IBM puts Informix in a featured role

It has been about 30 years since Informix was founded in California, and about 9 years since IBM added Informix to its portfolio of data management offerings in April 2001. For those 9 years, Informix has shared a single spotlight with its brethren IBM database servers.

That will change in 2010. IBM Information Management is giving Informix its own well-deserved spotlight to harness the true potential of this great database product.

I've been sworn to secrecy on the specifics, but here's the basic plan:

- Delight the install base
- Build strategic differentiation
- Create a proactive sales culture
- Build new, profitable revenue streams that capitalize on the growing footprint of Informix

Each of those four items involves its own focused plan that is already under way. To me, that's a sign that this is sure to be an exciting year for Informix. You can get a taste of things to come at **ibm.com**/software/data/informix/discover-informix/.

One of the great things about seeing the new plan is that I know what a strong platform Informix is and that it's worthy of this level of attention. I have always said: let customers do an old-fashioned Coke vs. Pepsi database challenge, and most of them will choose Informix—and love the choice they made, even years later. Want an example? On the IT vendor satisfaction tracking site VendorRate (www.vendorrate.com), Informix has placed first or second in customer satisfaction each quarter since the third quarter of 2008, and first overall for 2009.

> IBM Information Management is giving Informix its own well-deserved spotlight.

The reason? It always works! I have told real stories about this characteristic for years in this magazine and its predecessors (www.dbmag.intelligententerprise.com/columns/iiuguserviews.jhtml). Remember my friend the Informix DBA who does not need a work pager or cell phone because his system is always up? How about the country that runs its entire phone system on Informix totally unattended, except for the person who puts a backup tape in the system once a week?

That's not all that's going on. IBM will be producing and committing to a variety of IBM Software Group products that will work seamlessly with Informix. I wish I could give you the product names or even some hints about their identities, but let's just say that the list I was given includes more than 10 products from within the IBM portfolio, plus numerous well-known open-source products.

I've been known to have a few choice words to say about Informix getting the support it deserves, and that's why I'm excited—it looks like we're going to be seeing a lot more from IBM very soon. (And oh yes, rumor has it the name will be INFORMIX, not IDS.) Want to learn more? Mark your calendar for the IBM Information On Demand 2010 Global Conference at Mandalay Bay Resort and Casino in Las Vegas this October. In the meantime, check out the IIUG Web site at www.iiug.org to get your Informix fix.

DB2 10 for z/OS

More details emerge about new features that can slash CPU costs

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.*

Have you heard about DB2 10 for z/OS yet? The announcement of the beta program and presentations from the IBM System z Summit road shows highlight many of the new functions and features, and the upcoming International DB2 Users Group (IDUG) conference should provide further details. All of this demonstrates that IBM continues to listen to its customers: DB2 10 for z/OS addresses performance, scalability, availability, security, and data warehousing with features that can improve operations and overall performance, and—most important—reduce overall CPU consumption and total cost of ownership.

Many of these improvements, especially in application integration, SQL, and XML, can slash CPU costs for existing applications. Specifically, the DB2 10 optimizer does more analysis when choosing between two indexes, and it compares access-path costs more closely when determining the best access for SQL statements with OR and IN predicates. DB2 10 also improves the handling of SQL that overflows the relative identifier (RID) pool limits by using the DB2 work file resources and avoiding a costly tablespace scan. Additional CPU-reduction features include improved parallelism, LOB/XML streaming capabilities, and better handling of SQL Stage 2 predicates. These improvements will immediately benefit almost every application, without requiring any application changes or redevelopment.

DB2 10 also introduces a completely new access method called Hash Access, which uses a new hash space that provides direct access to a data row. In some cases, this direct access reduces data access to a single I/O, dramatically decreasing CPU use and speeding up application response time. There are trade-offs for using Hash Access, though: parallelism is not available and traditional clustering keys are not allowed against the hash data. Nevertheless, Hash Access will be great for product or customer identity data, and for other information where unique direct keys are already used.

As DB2 databases continue to grow in size and transaction volume, the amount of administration time required for database changes becomes more of a challenge. A number of DB2 10 utility improvements can further minimize downtime during normal operational and database-change activities. One of the best new features enhances operations by replacing the old DROP/RECREATE and REBUILD indexes method of adding new columns to a table to ALTER the new columns into the table and then performing an online reorganization for the changes to take place. The new method of enhancing tables eliminates downtime and improves availability for mission-critical very large databases (VLDBs).

With DB2 10, DBAs will also be able to create or rebuild a non-unique index without any application impact or locking downtime. This will especially help newly installed applications by making it possible to quickly define an index that can improve SQL access and accelerate problem resolution. This enhancement alone almost instantaneously improves performance for any installed application.

I'm looking forward to the beta program, road shows, and IDUG 2010 North America conference to bring out more details about all of the DB2 10 features. It's great to see how IBM keeps enhancing DB2 so that its customers can continue to drive some of the biggest, most complex, and most robust database, data warehousing, and transaction systems in the world.

Data Pays

By Howard Baldwin

At either end of Europe, IBM Information Champions build financial foundations and data management communities

Dividends

At first glance, you might think that Istanbul and Madrid, two cities at Europe's eastern and western reaches, have nothing in common, but you'd be wrong. They're both large cities at 40 degrees latitude, both trace their roots back to Roman times, and both share Moorish influence. And they are both home to IBM Information Champions.

Like their home cities, Information Champions Banu Ekiz of Istanbul and Manuel Gómez Burriel of Madrid share certain characteristics as well. Along with their extensive database experience, both are battle-tested veterans of their country's banking industry. Ekiz has worked for two Turkish banks since receiving her master's degree in 1996; Gómez has worked for Confederacion Española de Cajas de Ahorros (CECA) [Spanish Confederation of Savings Banks], a consortium of 45 Spanish banks, since 1986.

They also share a passion for database technology and its ability to provide the crucial foundation for an industry where, now more than ever, success is clearly linked to the world's economic viability. In these profiles, Ekiz and Gómez talk about their work with the latest database technologies, and how it continues to inspire them for the future.

The IBM Information Champion Program (formerly the IBM Data Champion Program) recognizes exceptional individuals who contribute to the advancement of information management at corporate and governmental entities around the world. In 2009, just 152 professionals were awarded this designation, based on their uncompensated contributions to user groups, community Web sites, conferences, and online forums. For more information, visit **ibm.com**/software/data/champion/.

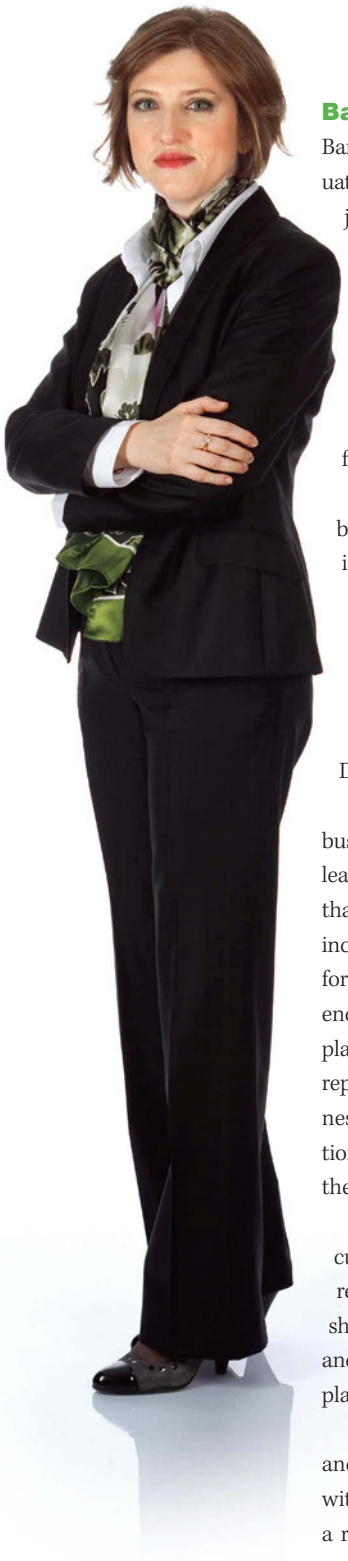

Banu Ekiz, Akbank

Banu Ekiz became interested in databases while an undergraduate at Bosporus University in Istanbul. During one summer job, she worked on computer security systems for a bank. "It was then I realized that database management systems had many capabilities. They were not limited to storing data, but their capabilities extended to any use of that data," she explains. She continued in a master's program at the same university, and for her master's thesis, she applied her DBMS knowledge to create a database and algorithms for studying and discerning healthy versus cancerous cells.

Her thesis showed the power of databases in healthcare, but in Turkey, the best computer-science students are enticed into two high-profile industries: finance and telecommunications. That's how Ekiz ended up in banking, working first for a leading bank in Turkey as a database administrator, then manager of its DBA group. In 2007, she moved to Akbank, one of Turkey's leading banks with headquarters in Istanbul and operations in the Netherlands, Germany, Dubai, and Malta.

Ekiz has put her database schooling to good use. She is the business intelligence coordination center manager at Akbank, leading a four-person group with extensive responsibilities that are crucial for the bank's efficiency. Her team (which includes outsourced development assistance) is responsible for the bank's enterprise information architecture, which encompasses database systems architecture, the hardware platform, the application development tools, the front-end reporting tools, metadata, enterprise data quality, and business intelligence (BI). Their current project: the implementation of a real-time operational reporting system, which will be the basis for a data warehouse running at near–real time.

For Ekiz, this is an exciting undertaking. She believes her current work is vital to the bank's success. "Most companies realize that they need a business intelligence strategy today," she says. "My group's work in leading the strategy definition and the architectural decisions about the business intelligence platform is the most interesting thing about my work."

For instance, she notes, the ability to access meaningful and clean data was just a dream a few years ago. However, with new technological changes, that dream can become a reality, she says. "Putting all the policies and procedures together while working with the business and accommodating its needs, however, is a big challenge." The bank's BI efforts require a well-defined strategy. "Without putting standards and rules in place, it's a real burden to keep up with the business's needs for a BI system. We realized that we had to define both the strategy and the architecture for our business intelligence solutions."

Crafting this strategy requires a high level of collaboration and communication, Ekiz believes. "We need to understand the business needs correctly. But the business has to trust IT about the technical decisions. That trust, along with a strong partnership, comes out of IT responding to the business's requests in a swift manner."

Her group is managing everything from data governance to performance. "Data governance is a new focus for us with the implementation of the real-time data warehouse. It's helping us provide good quality data," Ekiz says, noting that for high-level business intelligence efforts to be successful, they require a strong architectural foundation. At Akbank, the real-time data warehouse will use an IBM Cognos front end, and DB2 will serve as the source database. "We will be using replication techniques to move data from DB2 to our operational data store. For the replication, we will be taking advantage of the advanced monitoring features of DB2 so that the performance of our source systems will not be affected," she explains.

This kind of tuning is easier with modern systems, says Ekiz. "I started my career with DB2 version 2, and with every version, the robustness, availability, and business continuity functions have been far ahead compared to the other database systems." In her early days with databases, she remembers, "you had to develop your own monitoring and alert mechanisms" as well as tinker with the system to make it work to your specifications.

But just as databases have advanced, Ekiz believes that database administrators must evolve their skills as well. "DBAs should have great communication skills and the ability to work closely with their application development groups," she says. "They should also have teaching skills, so that they can explain new technological changes to others in their company." Ultimately, she believes that good DBAs should be able to combine their knowledge of databases with their knowledge of the needs and strategies of the business, and help the company move forward using those insights.

That's why she became involved with IBM database user groups, starting in 2003 as a speaker at the International DB2 Users Group (IDUG) Europe conferences. The following year, she joined the IDUG Conference Planning Committee. "I first met Banu in 2007 in Brussels for an IDUG conference planning meeting," says Surekha Parekh, worldwide DB2 for z/OS market manager at IBM. "My first impression was that she was easy to talk to and fun."

By 2008, Ekiz was chairing the IDUG Europe conference. "What I have learned in IDUG has helped me throughout my career," says Ekiz. "Listening to so many DB2 professionals come and explain their experiences at IDUG helped me realize how many things you can do with DB2."

Photos of Banu Ekiz by Ali Kabas

Now, she is delighted to be part of the IBM Information Champion Program, because "you get first-hand information about IBM products and strategies before other people, and this gives you a head start to make plans for new technologies."

Manuel Gómez Burriel, CECA

Manuel Gómez Burriel's database career started in the 1980s as an application programmer on batch-oriented mainframe applications. "It was very common for application programmers to want to move into system programming, and I took advantage of a chance I had to move into a database administration group," Gómez says. The database involved? IBM IMS version 1.3.

He's still working on IMS databases, and proudly so. Gómez is a DBA manager, supervising a staff of eight in the IT department of Madrid-based CECA, which supports 45 Spanish savings banks. Founded in 1928, the company's services include providing management statistics and analysis, regulation and financial analysis, fiscal advisory services, legal advice, and other general advisory services. Most important for Gómez, it offers technology services and support for its members.

An entity that well established has a reputation to maintain, especially in the financial services industry. Indeed, Gómez takes the reliability and performance of his databases very seriously; the challenge of ensuring both of those things is what he enjoys most about his job, almost as if the database were a mountain to be conquered (he also loves trekking).

"Our customers demand that they have access to their data on a 24/7 basis, and we focus on providing that service level because that's how we're judged in our job," Gómez says. For the past five years, he says proudly, his group has been striving to provide the highest level of uptime possible. "We are measuring annual downtime in hours, not days. Just this morning, we had a meeting to discuss a shutdown for maintenance, which we will take care of in less than two hours, between 4 a.m. and 6 a.m."

Gómez shrugs off having to occasionally work such unusual hours; it's part of the job. "Being a DBA, especially in banking, requires a lot of work and dedication, and frequently nighttime hours. The requests we get from business are getting more demanding. Not everyone wants to take this kind of job, but if they do, I can promise them a work life that will never be boring." Even after more than 20 years, he still finds his work "exciting and attractive."

Part of that excitement comes from watching databases evolve within the corporation. "When I started working as a DBA, the database was only for storing the information in a way that it would be available and compatible to everybody," he explains. "But now, while this definition still applies in those legacy applications, databases are also used for business intelligence, for data warehouses, and as the basis for new applications. They have become more important for more people everywhere and all the time; they're utilized as if they were office tools."

Gómez also believes that a key part of his job is to share both his excitement about his work and his database know-how. He is highly active in several database user groups, including IDUG and the Spanish regional user groups for IMS and DB2 that are part of Guide Share Europe. "Like Banu, Manuel is passionate, committed, and always very supportive," says IBM's Parekh. "Last year, Manuel set up the first Spanish DB2 Regional User Group. This was extremely successful, as we had over 100 attendees—a major achievement for a RUG meeting."

"Since the beginning of my life in IT, I have always believed in the concept of community as an easy and friendly way to share knowledge and experience," says Gómez, noting that the Web makes that easier than ever before. "We maintain forums on our Web site for topics, and to post opinions and messages. It helps to keep our community connected."

Certainly the IBM Information Champion Program offers that capacity as well, and Gómez sees it as a way to both share information and learn new skills. "It's really important for me to be part of this select and privileged group. I like having direct access to other colleagues, and I like sharing documents and recommendations. I'm very grateful to the program because of the opportunities it's offered me."

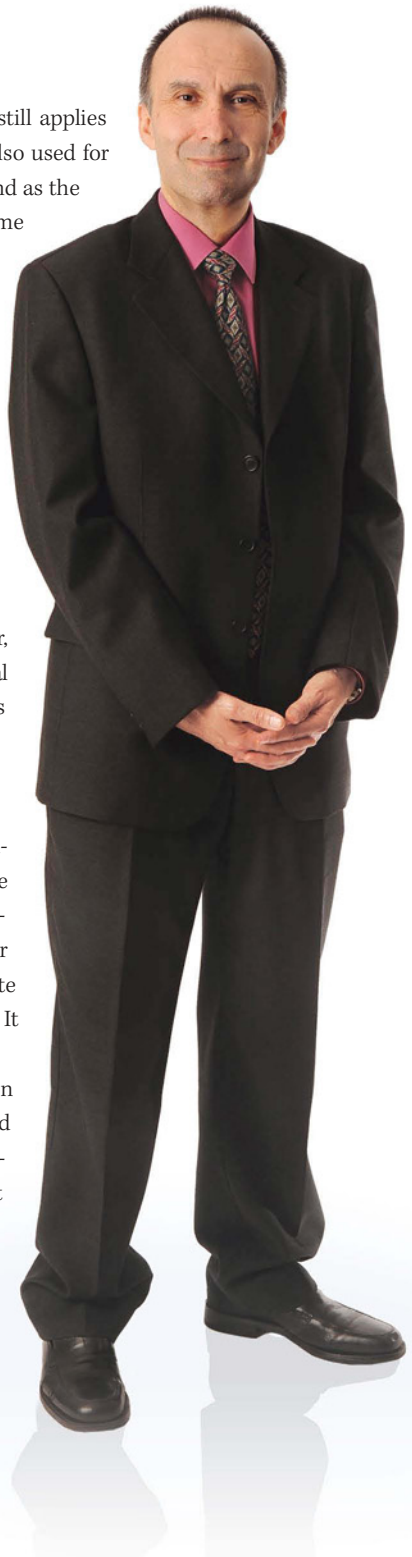

Photos of Manuel Gómez Burriel by Denis Doyle

Howard Baldwin is a Silicon Valley–based freelancer who writes about business and technology issues.

RESOURCES

IDUG: www.idug.org

IDUG Europe conference page: www.idug.org/idug-europe/index.html

Guide Share Europe: www.gse.org

Spanish DB2 Regional User Group: www.spdug.org

Why Governance Matters

By Merv Adrian

IBM experts debate the promise and perils of data governance and why data quality should be a top concern

Like painting the proverbial bridge, organizations are discovering that cleaning and maintaining enterprise data is a task that never ends. What's worse, it's possible to spend a sizable chunk of time—and cash—to clear out errors, omissions, and gaps in years of accumulated data, only to see the same problems reappear over time in enterprise data stores.

The challenge of maintaining data quality has become so large and intense that local fixes and one-time solutions are insufficient. Even information management projects and tools like master data management are not enough to govern data quality without a corporate commitment and culture-building effort. But how do you get that commitment? For that matter, where do you start when defining the problem?

Five years ago, IBM's Steven Adler convened a group of customers to tackle the question of governance, forming the IBM Data Governance Council (**ibm.com**/ibm/servicemanagement/data-governance.html) to investigate and implement best practices—and to advise IBM on how to proceed. At a recent conference, I invited Adler to join Brett Gow from IBM Global Business Services and Lise Neely from the IBM InfoSphere team to discuss the state of the issue and the market.

Steven Adler
Director, IBM Data Governance Solutions

Lise Neely
Market Segment Manager, IBM InfoSphere Master Data Management

Brett Gow
Associate Partner, IBM Global Business Services

How did we get to "data governance"?

ADRIAN: Steven, when you started the council several years ago, were a lot of people concerned about data governance?

ADLER: There were a lot of concerned people; they just didn't know what it was called. They were concerned about data quality and security. It had dawned on them that they might be spending too much money to protect worthless data and not enough on valuable data, and they couldn't tell the difference.

People were governing the use of information—they just weren't doing it with forethought, or consistency, or with a plan. Everybody makes governance decisions. Every decision is a policy—but how do you know whether it's the right policy, whether your outcomes are consistent with your goals? The challenge today isn't just to govern. It's to govern well.

GOW: Often, people don't understand the complexity [of governance]. Many think it's project-based or focused on technology. Others think it's an overlay to existing processes and procedures—activities, roles, and responsibilities different from what they're doing today. They are surprised when they understand the dynamics, the politics.

ADRIAN: For data managers, what are the symptoms that there's a data governance problem?

NEELY: If you are facing the same fire drill over and over, there is a process breakdown somewhere, and it's time to look at the root cause and address it there.

ADRIAN: Is it challenging to convert the abstract vision to something somebody in the organization might be willing to fund, and staff, and organize and drive?

GOW: Absolutely. There may be a visionary, but that's rare. More often it's regulatory mandates, audit issues that indicate the organization needs to take action. Sometimes governance is wrapped around specific initiatives: data warehouse development, master data management, or data quality issues.

ADLER: Following the credit crisis, regulators stepped in to do "stress tests." They discovered firms lacked systemic ways of reporting simple things; different business processes existed; people didn't want to share information. Regulation has had a catalytic impact.

How do I get started?

ADRIAN: So, where should we begin? Can I get started without the corporate vision statement? Start with business value, at the data manager's desk?

NEELY: Think about key initiatives driving the business. Even compliance is a business issue. If you begin by assessing your current state relative to your key business and IT initiatives, you'll see the biggest and quickest return.

ADLER: You can start anywhere. Executive sponsorship without a crisis is difficult to obtain. Often, programs just start. Nobody's going to build a council; nobody's sure if they're called a steward. They want to build systemic decision making, or control for glossary terms, or metadata, or master data management.

Define your sustainable goals for the program. Do you want to clean up metadata? Drive more revenue? Decide what to measure to demonstrate why you're not achieving your goals today.

When data problems pop up, people think of policy enforcement because policies have already been created. That's way too late; start with the premise that maybe your policies aren't effective. Collect evidence, real facts. Communicate, measure and audit results, and compare them to your goals.

> "Many clients think they can initiate information governance as a self-service exercise. But because of politics and their involvement in the environment, they don't have the true perspective."
>
> —Brett Gow
> *Associate Partner*
> *IBM Global Business Services*

The business end of data quality

It can be tempting to focus on technical fixes to data quality problems: better filters, different definitions, additional metadata. But the best tools and processes in the world can't—and won't—displace the people involved. During the panel discussion, Steven Adler brought up a great example of an intractable data quality problem that didn't get fixed until the solution considered the human element.

ADLER: People [weren't completing] information in new accounts because they got commissioned just on opening the accounts, not on what kind of information they put in. Name and the address and key information got the account open, but other information for cross-selling was lacking. Training didn't help.

[IBM Data Governance] Council members said, "You need a data quality cleansing tool or a new architecture." We threw theoretical ideas out. One after the other, [the customer would] knock them down: "Tried that; it didn't work."

I asked another member. He said, "We solved it. We invited sales, finance, and HR to a new governance council and decided that there was no technical solution to this problem: there was a compensation solution. We changed the model so branch operatives were also compensated for repeat business to the same customer." That changed behavior; people realized that without demographic information, they couldn't sell repeat business.

> "I've got to understand the purpose, what the business is doing, translate what I'm doing technically into what it does for the business. That's what gets traction. That's what gets attention. That's what gets budget."
>
> —Lise Neely
> *Market Segment Manager*
> *IBM InfoSphere Master Data Management*

ADRIAN: Lise, how do we collect this data and justify the ongoing effort?

NEELY: Say I need to understand the sources of customer information and how the organization uses them to get insight and create products and services that meet customers' needs. I have to understand the information supply chain, work with the business to figure out available metrics, and improve that information.

ADRIAN: So I don't need a definition for every data element in the organization. I need to know about the ones that relate to this problem and this goal.

NEELY: Right—something I can measure now, with a baseline. As data quality improves, measure it in terms of both data quality and impact on the business.

ADRIAN: Now I engage stakeholders with "I understand why we're having some of these problems that you've been calling me about. Let's talk about how we can fix them."

GOW: Absolutely. Identify root causes and potential solutions. Rally around the fastest or cheapest to fix, and determine the appropriate solution based on business need. Address the bulk of the problem or most of the core issues. Deliver that in a timely and cost-effective way.

Growing governance beyond the project

ADRIAN: Once we've identified a problem, articulated the value of its solution, and found some stakeholder support, how do we go to the more general principle of data governance?

ADLER: You've described a project-oriented approach. But people don't know what they don't know. They aren't aware that there are stewards or a chief privacy officer focusing on governance topics, and so they fix their own problem. There's no coordination. Governance isn't just something you do once in a while. It's a systemic approach to solving problems over time. You won't achieve 100 percent data quality overnight; you must learn from your successes and mistakes over time.

Teach people that even without a formal governance program, you're still governing: you're just not doing it with a systemic process. Monitoring and auditing are your most important tools because they tell you what you know and don't know.

NEELY: What bad thing would happen, and to whom? The policy decision becomes what we can put in place to prevent it, and how we can provide insight so we can take corrective action if something is breaking down. If you apply that approach to an initiative that people really care about—growing revenue, reducing cost, or controlling risk—you'll get a lot of attention.

ADRIAN: Sounds great when we're sitting around talking. Do people really get beyond the project, or are other steps needed to institutionalize this?

ADLER: There are six steps: set goals, define metrics, choose a decision-making model, communicate policy, measure outcomes, [and] audit. But even if it's one person solving one problem, it needs to be recorded. Every decision articulates policy.

Get experts in the room to make a decision you can't make on your own. Everybody gets that epiphany. You know the value of governance in your organization. Nobody has to tell you again.

ADRIAN: Data managers might say, "Sounds great, but I don't talk to those people." How do we shorten the cycle, get them into a room to go after this problem?

NEELY: Start small and build upon success. For example, understanding the banking customer is relevant to withholding requirements. As a result of data governance initiatives that one client took, they were better able to understand the risk profile and holdings across a particular group of customers, and to reduce reserve requirements because they understood their exposure. When word of this success got out, other groups raised their hands, so the project built upon its success and had a very positive effect.

ADRIAN: What kind of skills do data managers need for this?

NEELY: Communication. I've got to understand the purpose, what the business is doing, translate what I'm doing technically into what it does for the business. That's what gets traction. That's what gets attention. That's what gets budget.

ADRIAN: Is this the kind of thing that an outside party can help facilitate?

GOW: It's often critical. Many clients think they can initiate information governance as a self-service exercise. But because of politics and their involvement in the environment, they don't have the true perspective. Consultants can help people understand that they're not alone; competitors have the same problems, and tried-and-true approaches can yield some benefits. What freezes people is thinking they're not going to be able to deliver and they're going to be culpable.

ADLER: When we began the council, we didn't have enough experience, so we built a maturity model—a benchmarking tool. Many engagements and countless meetings around the world later, we know how to do it right. There is a science to doing governance well, and there are disciplines that need to be learned. People like Brett and Lise have the experience and know-how to help customers. Working with a partner will help you get it right. Without that experience and those disciplines, it's harder.

ADRIAN: Steven, let me close with you. What do we need for the next step?

> "When data problems pop up, people think of policy enforcement because policies have already been created. That's way too late."
>
> —Steven Adler
> *Director*
> *IBM Data Governance Solutions*

ADLER: We must provide more automation because governance is difficult. You can set goals in the abstract, but today there's no system for keeping them somewhere to be compared to outcomes. "Define your metrics" is nice to say, but there's no automation for collecting this information and deciding which facts are relevant and which ones require policy changes. I can talk about decision-making models, but most firms don't have tools to decide which model to use and when: when to be transparent and open and involve lots of people, and when to use the autocratic model.

Those processes are done today without any knowledge. Until we can provide automation that helps firms in a much more systemic way, we're not going to be able to move this market forward. ✱

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.

RESOURCES

IBM information governance solutions: ibm.com/
software/info/itsolutions/information-governance

IDUG
The Worldwide DB2 User Community

There's only one IDUG.
Your only source for independent, unbiased and trusted DB2 information.

Join IDUG to receive:

- Year-round resources tailored to DB2 users, by DB2 users
- Videos of the best DB2 technical presentations in the industry
- An archive of past conference presentations
- Access to hundreds of technical articles
- The answers to your toughest technical questions on the DB2-L listserv
- The opportunity to attend global conferences and regional events
- A calendar of Regional User Group meetings
- A special 45% discount off of books from IBM Press
- And we're just getting started…

Join IDUG, the International DB2 User Group today!

What are you waiting for?

IDUG Mentors can receive up to 80% off to bring a new member to a global conference!

IDUG 2010 Europe

Hilton Vienna Hotel , 8-12 November 2010

Information at: **www.IDUG.org**

Taking the First Steps Toward Data Quality

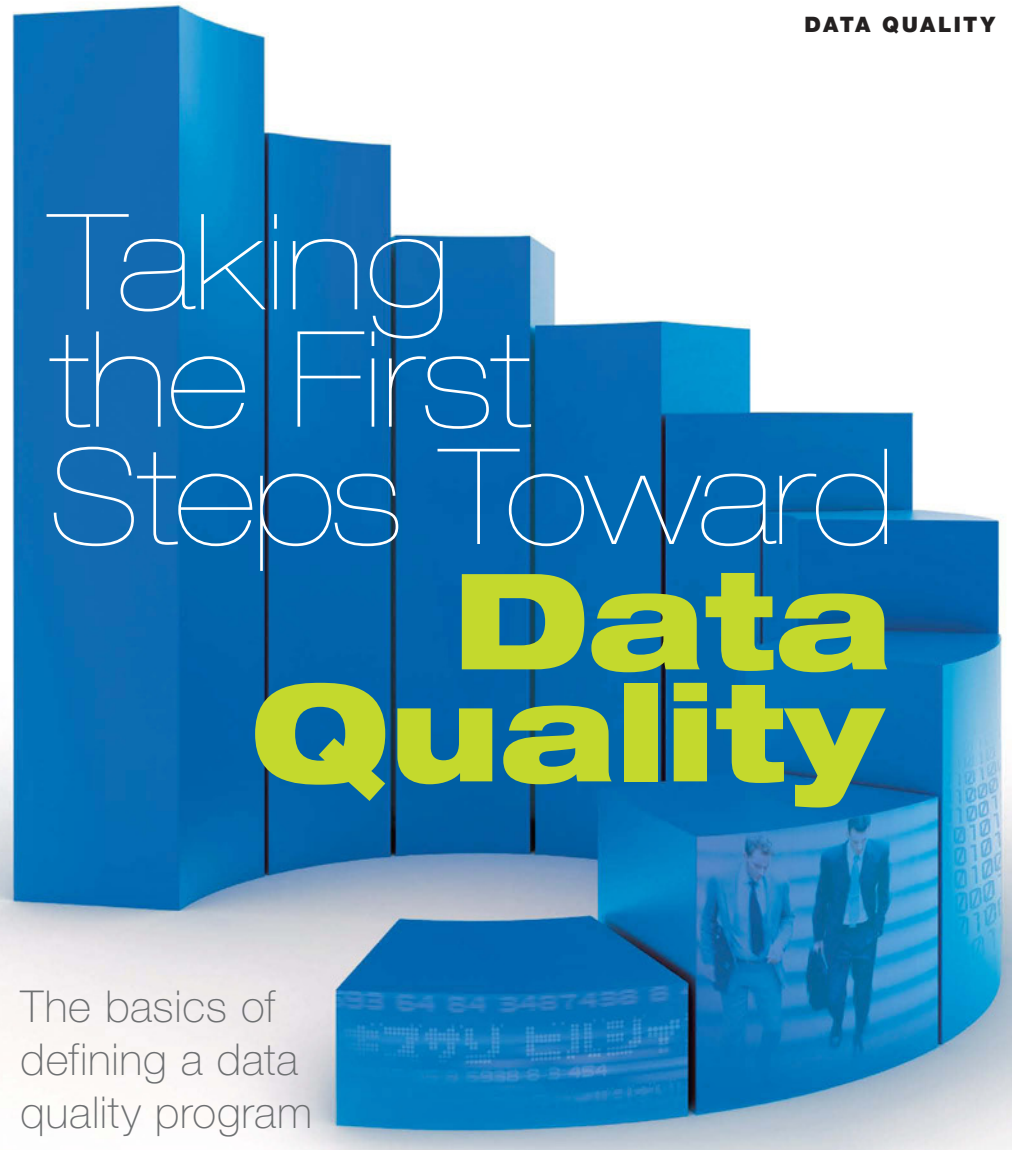

The basics of defining a data quality program

By Elizabeth Dial and Cameron Crotty

It would be difficult to find someone who thinks that data quality isn't important. Certainly, the effects of poor data quality are painfully clear: organizations depend on data to make strategic management decisions, provide customer service, and develop processes and timelines. If that data is obsolete, inconsistent, incoherent, or just plain wrong, it can cost a company time, customers, and revenue. Additionally, demonstrating data quality is often a requirement for regulatory compliance.

Trying to develop an overarching program to maintain and improve data quality can feel like chasing ghosts. In this article, we will present important concepts essential to a successful data quality program. We will also outline a plan for initiating a data quality program through a project tied to a specific business initiative.

SMART data quality

The SMART mnemonic is a staple of project-management theory, but it's also extremely applicable in the realm of data quality. Here are some SMART things that you should do in your data quality program:

S

Specific: Define data quality at a low enough level for it to be meaningful. Everyone wants good data quality; the question is what "good" means for the user, for the entity, and for the attribute.

M

Measurable: After defining data quality, measure and monitor the data.

A

Actionable: Reports should be at a level such that actions can be chosen in order to improve. A data quality program should also have guidelines as to how to take actions.

R

Realistic: Data quality does not improve overnight. Creating a plan that promises huge gains and benefits is sure to fail.

T

Time-driven: Break down the program into implementable milestones with achievable dates.

What is data quality?

The first step to creating a successful data quality program is to understand what data quality means in the context of a particular organization. Broadly, quality data is "fit for use": it can be trusted and it is suitable for its intended purpose. Assessing whether a specific set of data meets the criteria requires answering several questions: What data is being used, who is using it, how are they using it, when are they using it, and why? This becomes more complex as organizations begin sharing data across lines of business, departments, and other entities. It quickly becomes clear that to measure data quality effectively, it must be defined at the entity or even at the attribute level.

Data quality can be measured in many dimensions, including accuracy, reliability, timeliness, relevance, completeness, and consistency. Of course, different organizations will have different priorities. However, it's important to recognize that there are technical and business views of data quality, and both are important. Data that meets technical quality standards (such as consistent, correctly formatted, well-defined) but that is not perceived by users as reliable, accurate, or useful will have little impact on the organization. In short, ensuring data quality requires an awareness of both technical and business requirements.

Strategy and setting goals

One of the best ways to build a data quality program is to tie it to a strategic business project. Data quality isn't the ultimate goal—it's the means to the goal, which supports, extends, or enhances the business in some way. For example, a company that sets the goal of increasing sales at retail stores by 20 percent over the coming year might want to create a data quality program that ensures that information delivered to store managers about sales trends of high-value products is accurate, timely, and precise.

The charter, objectives, and plan for a successful data quality project should follow the well-known project management SMART mnemonic: Specific, Measurable, Actionable, Realistic, Time-driven (see sidebar, "SMART data quality"). This is also the time to address high-level organizational issues (such as who will own the program and who will be the major stakeholders) and technical issues (such as the tools to be used and the environment for data analysis).

Scope and definition

Once the goals of the data quality project are established, the next stage is discovery and assessment, starting with identifying the data that is within the scope of the project. With the owners of the data entities established, the

business and IT teams can move on to defining the data entities and their attributes. For every entity, there should be a business definition (such as what the data is and why it is meaningful), a technical definition (field sizes, types, relationships, and hierarchies; expected data patterns or formats; and so on), and a quality definition that includes expected and acceptable values along with business rules and formatting rules.

A tool such as IBM InfoSphere Business Glossary can be helpful at this stage, by providing a repository for data definitions and a simple user interface for entering, searching, and exploring vocabulary and definitions. An enterprise glossary helps ensure that definitions are consistent across projects, supports collaboration between business and IT and across lines of business, and helps build the common vocabulary and understanding of data.

Assessment and profiling

Assessing the actual data based on the criteria established by the business and technical teams is the next step. Here, software such as IBM InfoSphere Information Analyzer is used to profile the data. During profiling, the data is checked at the column, table, and cross-table level to assess its completeness, validity, and conformity to known or expected usage. If the business definitions for the data have been clearly established, the rules can be entered into InfoSphere Information Analyzer, which will use them to validate the data.

InfoSphere Information Analyzer also provides a central business rule repository, promoting reuse and consistency across different projects and implementations, and shares a metadata repository with InfoSphere Business Glossary, which simplifies data sharing and implementation. Other data quality tools make it possible to perform sophisticated, automated analysis on data depending on the data quality and validation needs (see sidebar, "Resources").

After the assessment, results should be reviewed by both technical and business teams to develop a complete understanding of the data. The next step is deciding what action to take based on the reports. Sometimes, the action will be technical, such as changing a data model or a user interface. Other times, the action will involve a business process or policy change, such as altering who is responsible for gathering and entering the data.

From assessment to program

At this point in the process, the organization should understand what at least part of its data environment looks like and know what its business objectives are. The next step is to create a data quality process that will move the organization from the current state to the desired state.

Building this program is beyond the scope of this article, but a data quality program has three essential elements. First, it continually uses the structure defined for the data quality assessment to regularly measure data quality. Second, it assigns stewards to continually monitor data quality. Finally, it provides a process for developing action plans for dealing with data quality issues identified during ongoing monitoring.

Today, many organizations discover data quality issues only when they impact the business—usually with a negative result. By actively assessing and monitoring data quality, organizations can graduate to identifying data issues and addressing them before they cause problems. By creating repeatable processes and reusable assets, organizations can ground the abstract concept of data quality in a real-world project, and use it to minimize risk and generate business value. ✱

Elizabeth Dial *is a technical solution architect in the IBM Information Agenda Architecture Group.*

Cameron Crotty *is the editor of* IBM Data Management *magazine.*

RESOURCES

IBM InfoSphere: ibm.com/software/data/infosphere

IBM InfoSphere Information Analyzer: ibm.com/software/data/infosphere/information-analyzer

IBM InfoSphere Business Glossary: ibm.com/software/data/infosphere/business-glossary

IBM InfoSphere QualityStage: ibm.com/software/data/infosphere/qualitystage

IBM InfoSphere Discovery: ibm.com/software/data/infosphere/discovery

RIGHT DATA RIGHT PLACE RIGHT TIME

By Paul Pendle

Storage tiering for the DB2 database administrator

To date, DBAs have had little say in what type of storage their databases lived on: they simply asked for a given amount of space and it was provided by the storage administrator.

Now, however, it's important for DBAs to think about the types of storage that they use, because storage isn't "just storage" anymore. There are high speeds, low speeds, and in-between speeds, and each is best suited for different uses.

The challenge lies in deciding where to put data to achieve the best return on investment (ROI). This article will explore how to create a tiered approach to storage with IBM DB2 for z/OS and achieve the ultimate goals of getting the right data in the right place at the right time to increase performance optimization and reduce total cost of ownership (TCO).

Disk storage tiers: Choices, choices

Never before has a DB2 DBA faced so many storage options with so many price and performance variables. At the high end, solid-state drives (SSDs) that can comfortably support random I/O access patterns at a rate of 5,000 I/O operations per second (IOPS) are an attractive choice for demanding, I/O-intensive applications. However, they can be relatively expensive.

On the other end of the scale, Serial ATA (SATA) drives can be as large as 2 TB but typically deliver only around 75 IOPS. They perform reasonably well for sequential workloads but do not do so well with random I/O access patterns.

Between these extremes sit the commonly used Fibre Channel (FC) drives that have been the backbone of enterprise storage arrays for many years. FC drives deliver 180 IOPS with acceptable response time (less than 5 ms). FC drive capacities are growing very quickly but appear to have reached a speed/response plateau.

However, individual devices are just part of the equation. Drives of all types are arranged in enterprise storage arrays in ways that deliver different price/performance ratios. For example, RAID-10 (striped mirrors) performs better but is more costly than RAID-5. Similarly, RAID-5 (striped single parity) has some performance advantages over RAID-6 (striped double parity), yet RAID-6 provides higher protection against data loss.

You may have some or all of these types of storage available to you, or you may have other variations. Whatever your environment, start your storage tiering plan by defining your tiers. High-cost, high-performance storage becomes your gold tier, followed by less-expensive and less-responsive silver and bronze tiers.

The challenge of choosing tiers

If defining storage tiers is relatively straightforward, placing the right data on them often isn't. The decision starts with the broad parameters of the project: Is the database critical to the business? What are the performance requirements? What are the availability requirements?

From there, the qualifications become more complicated. Within a single DB2 system, not all data is the same; it might make sense to have data for one database stored, in part, on each tier. Even within production data, some tables are rarely accessed while others are accessed continually. In a final twist, access profiles can vary over time (see Figure 1).

You also need to consider characteristics of the I/O that makes up your storage workload. Is it random or sequential? Read or write? What are the ratios of reads to writes? Are the I/Os small or large? Understanding your workload will allow you to make appropriate storage tiering decisions.

Characterizing the workload

Workloads have different intensities when viewed in time and when viewed from the storage perspective. In addition to the intensity, the type of workload can have a bearing on which storage tier is optimal for deployment. All I/Os are not equal!

Consider a write operation: Most enterprise storage arrays cache an inbound write and return a channel end/device end acknowledgment immediately to the host. This process is usually referred to as a direct access storage device (DASD) fast write. At a later time, the write is "destaged" to physical storage. In this case, does it really matter what the underlying storage is? To further muddy the waters, writes from the DB2 buffer pool are usually written asynchronously with respect to the transactional workload and are therefore a lower priority than, say, synchronous reads.

Sequential read workloads, such as those generated during a tablespace scan, are also interesting. When DB2 processes a plan that requires a tablespace scan, it performs prefetch activities, which are asynchronous read-ahead I/Os. When the storage array detects this sequential activity, it also performs prefetch by reading the table data into cache on the array. This puts the disk into a kind of streaming mode, notwithstanding interruptions from other I/O requests to the same physical storage media.

In streaming mode, spinning disks are actually as good, if not better than, solid-state disks. The gold tier might not be the best choice for this kind of DB2 workload. In fact, moving this data to the gold tier will not only fail to improve performance, it may take up valuable space that could be used for other I/O activity.

While selecting the correct tier for certain kinds of I/O activity may not be intuitive, an empirical measurement of the workload can lead to a more optimal tiering deployment.

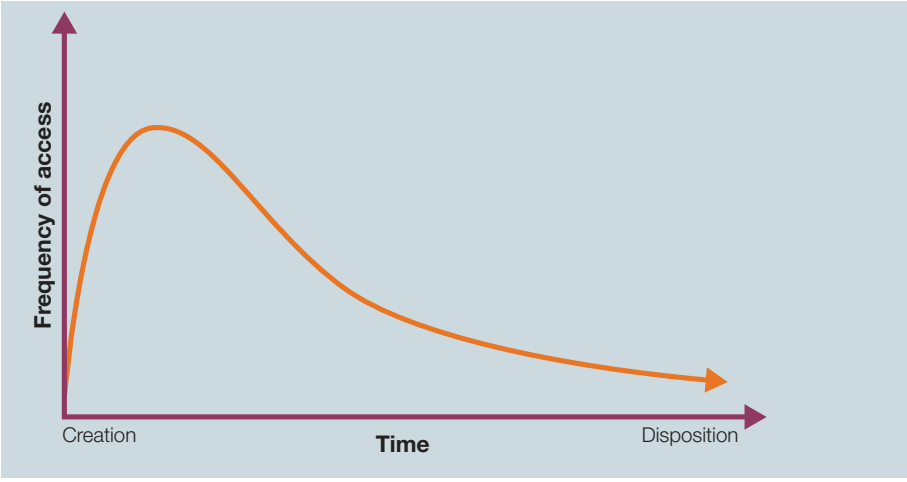

Figure 1: *The pattern of access to data, from creation to disposition, is usually characterized as the information life cycle.*

Selecting the best workload for each tier

The best way to determine which data belongs on which tier is by analyzing Resource Management Facility (RMF) data for the tablespaces to determine I/O characterization. For example, when looking for data that might be suitable for SSDs, it is instructive to examine System Management Facility (SMF) 42 subtype 6 records, which show the tablespaces with a high DISCONNECT time. The high DISC time is usually an indicator that the page to be read was not in the storage controller cache for the DB2 synchronous read. Data sets with high average DISC times in the interval would generally be likely candidates for moving to SSDs.

But even if you identify data sets with a high DISC time (that is, that are getting storage cache "misses"), data sets that have the highest miss percentages are usually the largest data sets because they cannot be held in storage cache easily. What you really need is to understand the "miss density" of your data sets. The miss density is how many cache misses you get per gigabyte of tablespace storage. You can compute this value with some artful spreadsheet work using DCOLLECT and the SMF records.

Other metrics from the SMF reports can indicate data sets with minimal I/O activity. Those data sets could be candidates for the bronze tier of storage, especially if they are large.

Although it is not possible to go into all the details of this type of analysis in this article, some general rules can be derived:

- Random read I/O activity is the best workload for the SSDs, especially the activity that results in a lot of storage cache misses.
- Highly sequential activity is best for the FC or SATA drives.

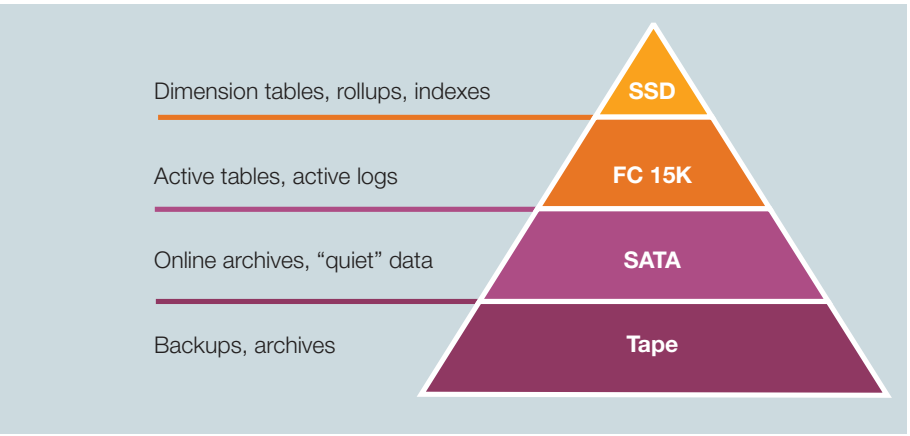

Figure 2: *Static tiering using SMS storage groups matches storage tiers to data characteristics.*

- High write activity is best for FC drives without using parity protection. If the storage controller is cache-constrained, the high write activity could be directed at SSDs since the writes can be destaged to disk faster than with FC or SATA drives.
- Low I/O activity is best for SATA drives.

Using DFSMS to create a manual storage tiering strategy

Ideally, storage mechanisms and databases would automatically work together to identify the best tier for each data, and then shift the data to that tier—but we're not there yet.

In the meantime, you can use IBM Data Facility Storage Management Subsystem (DFSMS) to deploy a rudimentary, manual tiering methodology. By constructing storage groups that contain volumes of similar performance characteristics, you can then use Automatic Class Selection (ACS) in combination with storage class settings to direct DB2 tablespaces to the most appropriate tier.

This is rudimentary for three reasons: First, the control of data set placement is at allocation time (when the data set is created). This does not take into account the data life cycle. What happens when the data is no longer accessed frequently and becomes a candidate for a lower-performing tier? Second, it assumes that you know tablespace performance requirements in advance. What if those requirements are unknown and the tablespace is inadvertently placed on the wrong tier? How do you move it? Third, maybe only some of the tablespace is "hot" and the rest is cold, so it should exist on multiple tiers.

A storage tiering solution

When tablespaces are known to have persistent, deterministic performance requirements, a static storage-tiering technique can be applied, perhaps using DFSMS. The overall goal is to match the storage tiers to the performance and I/O characteristics of the data. Figure 2 depicts a simple static tiering solution that could be implemented using Storage Management Subsystem (SMS) storage groups. In reality, the tiering model is rarely this simple due to the volatility of data access patterns.

Data movement between tiers

In any tiering model, there are a number of ways to move data from tier to tier transparently, without affecting database availability:

- **DB2 Reorg utility:** Using the online reorg to move a tablespace from one disk to another is simple and effective. Controlling where it goes is not so easy, but can be achieved using guaranteed space or by changing ACS routines.
- **DB2 partitioning:** Prudent choice of a partitioning key can allow the DBA to roll DB2 partitions for a tablespace from one tier of storage to another using the data set movement tools listed below.
- **Third-party tools:** These include IBM Softek TDMF, IBM Softek zDMF, EMC z/OS Migrator, EMC Virtual LUN Migrator, EMC FAST, Hitachi Tiered Storage Manager, FDRPAS, and FDRMOVE. Note that when using these tools, you must validate which ones you need based on whether you need to move the data at the data set level or the volume level, and whether the DB2 tablespace or volume is actively being used.

Storage tiering: Heading toward automation

Storage tiering for DB2 for z/OS can certainly improve ROI and increase performance optimization, although it can be a burden for DBAs to manage the whole strategy. DBAs need intelligent storage systems that can automatically act on database environments and place data at the tablespace level (or at an even lower level of granularity) on the optimal storage tier based on user-provided policy settings. In the meantime, it's good practice—and potentially lucrative—to familiarize yourself with the characteristics of the storage that you request or that you're assigned, and think about the best ways to use it.

Paul Pendle *is a consulting systems engineer for EMC Corporation. He has more than 30 years of experience in databases, hardware, software, and operating systems both from a database administrator's perspective and from a systems programming perspective. He has worked with DB2 for z/OS since version 1.2.*

Will your z/OS batch jobs complete or capsize?

Smart/RESTART lets your applications restart from near the point of failure — after abends, recompiles, even system IPLs. Your applications can run restartably, often without source changes!

Smart/RESTART guarantees that your program's sequential file and cursor position, working storage and VSAM updates stay in sync with changes to DB2, MQ, IMS and other RRS compliant resources. So you can restart fast with assured integrity.

Smart/RESTART is a robust, reliable and proven solution used by Global 2000 organizations worldwide to run their mission-critical z/OS batch applications. It's the standard for z/OS batch restart.

Restart made simple™

Download our White Paper:

"Beyond Restart and Concurrency: z/OS System Extensions for Restartable Batch Applications"

For a free trial visit www.relarc.com, or call +1 201 420 - 0400

Relational Architects International

What's New:

DB2 for i 7.1

Adaptive performance tuning and new SQL features give data professionals a smooth ride

By Kent Milligan

Since its early days on the AS/400 platform, IBM DB2 for i has been known for its ease of use and low administration requirements. With the newest release, which hit the streets in late April, the DB2 for i team focused on ways to support developers through self-managing technology and enhancements that simplify common tasks. At the top of the list of feature enhancements: simplified performance tuning, a new real-time self-tuning engine, additional SQL features, and a method to integrate advanced processes such as data encryption into existing applications without additional code.

Fast, simple performance tuning

The DB2 for i 7.1 release greatly simplifies performance tuning in several ways, starting with enhanced management tooling in IBM i Navigator, the central management console for DB2 for i. The latest release includes performance boosters in the DB2 engine, such as enabling solid-state drive and in-memory searches of DB2 databases. IBM DB2 Web Query has also been enhanced, making it easier for DBAs to meet the response-time requirements for both online transaction processing (OLTP) and business intelligence solutions.

But perhaps the most exciting new performance optimization feature in DB2 for i 7.1 is Adaptive Query Processing (AQP). With this feature, the DB2 for i query optimizer can make real-time plan adjustments—such as changing the join order or utilizing a new index—while the SQL request is running and without disrupting the application. Imagine a car engine that can tune itself as you cruise down the road instead of making you stop to adjust the cylinder timing or tighten the belts; that's what AQP can do.

AQP builds on the self-learning query optimization capabilities introduced in the prior release of DB2 for i (version 6.1). But where self-learning optimization technology must wait for a slow-running query to complete before it can learn from past executions and adjust the plan for future executions of the query, AQP can act in real time, as the query is running.

AQP relies on intelligent monitor agents to coordinate real-time performance adjustments, automatically assigning agents to queries that are expected to run longer than a few seconds. During query execution, the monitor agent periodically compares the runtime execution metrics with the optimizer's estimated costs. Whenever an agent detects a significant deviation between the runtime and estimated costs, the query optimizer is invoked to re-examine the query plan and look for a more efficient option.

If a more efficient plan is found (such as a new join order or utilization of a new index), the query is altered to use the new plan. The new plan and restart of the query are completely transparent to the application and the user—the only noticeable effect is improved performance.

AQP is available only to those SQL statements processed by the SQL Query Engine (SQE). But the 7.1 release adds support within SQE for logical file references on FROM clauses. As a result, even more applications will be able to leverage SQE and take advantage of the numerous cases where SQE delivers performance that can be magnitudes faster than the heritage query engine.

New SQL and field procedures

The DB2 for i 7.1 release also contains many new SQL features that make it easier and faster for developers to both extend the functionality of their existing applications and deliver new solutions. A new XML data type and set of XML publishing functions allow IBM i applications to easily interoperate with XML documents. Simple integration of stored procedures returning result sets is now available to RPG and COBOL applications using embedded SQL and to SQL routines. In addition, support for global variables, array types, Merge statements, Currently Committed concurrent access resolution, and the Or Replace clause also lower the barriers associated with porting applications based on Oracle and other database management system products.

Data privacy is a top concern for companies around the world, and new support for a field procedure exit routine, known as Fieldproc, enables developers to transparently deliver column-level data encryption with minimal coding changes. DB2 Fieldproc support enables developers to register an exit program at the field level that DB2 automatically calls each time that a row (record) is written or read. Registering the program at the field (or column) level ensures that the Fieldproc program will be called each time that a database read and write is performed from any interface or application—regardless of the programming language used. On write operations, DB2 will call the Fieldproc program to get the encoded value of the input data and then place the encoded value into the corresponding DB2 column.

The Fieldproc program can perform any type of encoding—including data compression—but it's expected that encryption algorithms such as Advanced Encryption Standard (AES) will be used most often for sensitive data. Correspondingly, any read operation would cause the DB2 engine to invoke the Fieldproc program to return a decoded value of the stored data, which DB2 then passes back to the application. The Fieldproc program could call system APIs to retrieve user or environment information to have conditional encoding and decoding behaviors. For example, it may be a requirement that only users belonging to the administrator group are allowed to see the complete credit card number, while all other users have access to only the last four digits (see Figure 1).

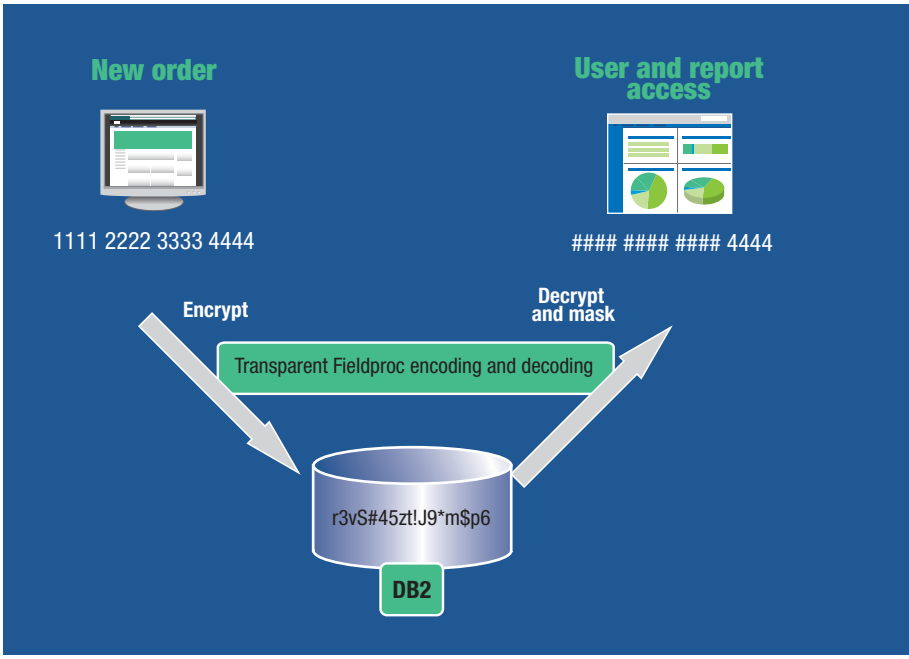

Figure 1: *With the field procedure exit routine, developers can easily add column-level encryption and data masking to existing applications.*

Using Fieldproc requires no changes to existing applications. With prior releases, adding column-level masking or encryption to applications using the native, non-SQL interfaces required developers to add encryption and decryption logic to every program that performed I/O on the table containing sensitive data. This is not a trivial task! Relying on Fieldproc support to centralize and integrate the encryption routines eliminates the need for extensive programming changes, dramatically simplifying the development and delivery of column-level encryption solutions.

The new release of DB2 for i adds up for both administrators and developers. By enabling better performance, simplifying tuning, and putting advanced technology to work, version 7.1 keeps you—and your end users—on Easy Street. ✱

Kent Milligan is a DB2 for i senior certified IT specialist in IBM ISV Solutions Enablement for the IBM i platform. He spent the first eight years of his IBM career as a member of the DB2 development team in Rochester, New York, and he speaks and writes regularly on DB2 for i relational database topics.

RESOURCES

IBM DB2 for i: ibm.com/systems/i/db2

IBM DB2 for i online publications:
ibm.com/systems/i/db2/books.html

IBM DB2 for i developerWorks forum:
ibm.com/developerworks/forums/forum.jspa?forumID=292

Intel and IBM Redefine the Limits of Data Warehousing

"The leadership performance of DB2 on new Intel Xeon processor 7500 series–based systems is the direct result of the deep collaboration between IBM and Intel. Together, we have achieved up to double the performance of previous-generation systems through full and efficient utilization of all available cores, threads, and memory to deliver even more value for our customers."

—Berni Schiefer
Distinguished Engineer,
Information Management Performance and Benchmarks, IBM

A performance-optimized data warehouse requires the right balance of database, server, and storage resources. The complexity involved in achieving this simple realization is the driving force behind a multi-year collaboration between IBM and Intel, who share the objective of delivering exceptional value—faster—to customers in need of business intelligence solutions. This is the underlying principle of the IBM® InfoSphere™ Balanced Warehouse™ portfolio, and a core requirement for next-generation analytics systems.

With more organizations choosing XML as their preferred transaction format, database administrators and warehouse architects are being challenged to integrate massive volumes of business-critical XML data into more sophisticated decision-support systems. Common design considerations for many of today's largest high-performance data warehouses already include the flexibility of the IBM DB2® pureXML® capability and the robustness of IBM System Storage™ solutions, but may not account for the latest innovations available from servers based on Intel architecture.

To this end, IBM and Intel have collaborated over multiple generations of products to optimize and measure the ever-increasing XML transaction processing capability of DB2 software on Intel® Xeon® processors, beginning nearly five years ago with single-core CPUs and humble 50 GB XML data sets. The latest performance tests reveal that multi-core Intel Xeon processors can indeed handle 1 TB XML repositories with ease.

But beyond transaction processing performance, measuring data warehouse performance is more about proving the solution's capacity to scale the performance of complex queries in tandem with the increasing data sizes those queries address. Can an Intel-based solution support a growing, multi-terabyte data warehouse without sacrificing query response times? To answer this question, IBM and Intel constructed a decision-support system based on high-performance components available in the market today, and extended the XML data size to 10 TB.

IBM DB2 9.7: The Heart of IBM InfoSphere Warehouse

IBM DB2 9.7 software ("DB2 9.7") represents the next generation of relational database servers, offering design innovations that make it well suited for demanding, mission-critical applications. For example, applications developed for service-oriented architectures should be able to scale with ease, be continuously available, and natively manage XML; DB2 software is engineered to handle all three requirements.

Sophisticated XML Data Management with IBM DB2 pureXML

Built on the high-performance XML data management capabilities of DB2 9.7, IBM InfoSphere Warehouse 9.7 software ("InfoSphere Warehouse 9.7") is optimized for data manipulation, query and retrieval, and data storage in a highly scalable, highly available, and secure architectural framework. Database design options for XML data—such

as hash partitioning, range partitioning, and multidimensional clustering—can help improve scalability and query performance, assist developers in exploiting parallel-processing environments, and simplify the addition and removal of time-sensitive data.

To support business collaboration and enhance application programmer and administrator productivity, DB2 9.7 offers IBM DB2 pureXML data schemas and scripts for major industry-specific XML messages, including Financial Information Exchange Markup Language (FIXML), Health Level Seven (HL7), Association for Cooperative Operations Research and Development (ACORD), News Markup Language (NewsML), and Human Resources XML (HR-XML).

IBM System Storage DS8700: Resiliency, Performance, Scalability, Value

The IBM System Storage DS8700 model ("DS8700") is the most advanced model in the IBM DS8000® lineup, offering a new level of performance for the company's flagship enterprise disk platform. Enabling up to 2.5 times the performance of the previous model, the DS8700 is designed to support demanding business applications with its unparalleled data throughput and resiliency features and five-nines availability (99.999 percent).[1] Moreover, with its tremendous scalability, new storage tier optimization, and broad server support, the DS8700 can help simplify the storage environment by consolidating disparate storage platforms onto a single system. At the same time, it provides the availability and performance levels required by your most important business applications.

The DS8700 model is a testament to IBM's reputation for outstanding quality and world-class engineering. And in today's challenging economic climate, having an enterprise disk platform that combines high levels of system and application availability with superior performance, flexibility, and total cost of ownership is essential.

Intel Xeon Processor 7500 Series: Powering Performance

Built to handle the most demanding applications, the Intel Xeon processor 7500 series delivers a quantum leap in enterprise computing performance. With up to 8 cores and 16 threads per processor, and a new high-bandwidth interconnect system that enables extreme scalability, this new platform is a top choice for supporting almost any high-performance data warehousing environment.

Green IT organizations can also take advantage of industry-leading virtualization performance and a server consolidation ratio of 20:1 to get more done with fewer servers, reduced power consumption, and less overhead.[2]

Mission-critical data warehouses will benefit from Intel Advanced Reliability Technology, a new feature in the Intel Xeon processor 7500 series that provides automatic detection and correction of errors, dynamic reassignment of workloads across CPUs, and even individual virtual machine recovery in consolidated environments—everything you need to help maintain data integrity, minimize downtime, and maximize productivity.

IBM DB2 and Intel Together: Doubling Performance

IBM and Intel have worked together for years to deliver multiple XML benchmarks featuring DB2 software on Intel Xeon processors, utilizing the industry-supported Transaction Processing over XML (TPoX) workload.

Figure 1 demonstrates that customers with DB2 9.7 deployed on previous-generation Intel Xeon processor–based servers can realize double the performance by upgrading to systems powered by the Intel Xeon processor 7500 series, such as the IBM System x® 3850 X5 server.[3]

Figure 1 also shows that the Intel Xeon processor 7500 series offers better virtualization efficiency. The amount of

"The new Intel Xeon processor 7500 series delivers to our mutual clients a new level of value for their business intelligence solutions. In collaboration with Intel, IBM has demonstrated InfoSphere Warehouse scaling linearly up to 10 TB of XML data on IBM System Storage DS8700—with similar query response times. This kind of efficient scalability improves the economics of solving complex business problems with IBM analytics solutions on Intel architecture."

—Salvatore Vella
Vice President of Development, Database Servers and Data Warehousing, IBM

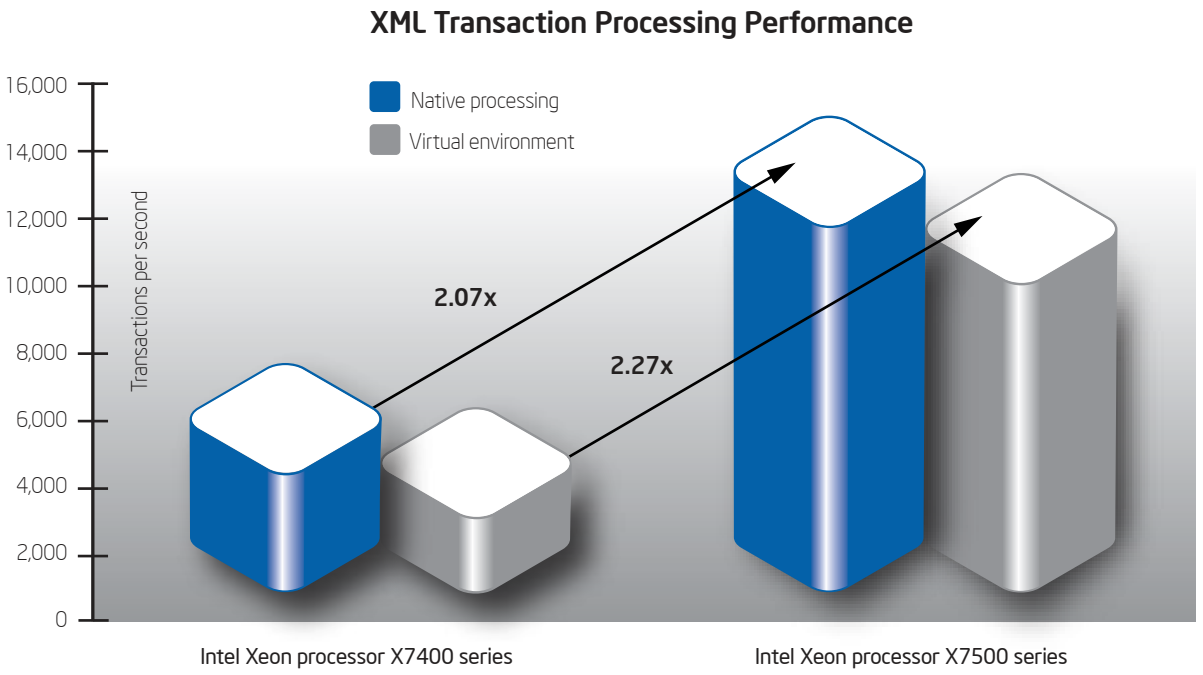

Figure 1: Compared to the previous-generation Intel Xeon processor 7400 series, XML transaction processing is doubled for both native and virtual environments on the Intel Xeon processor 7500 series.

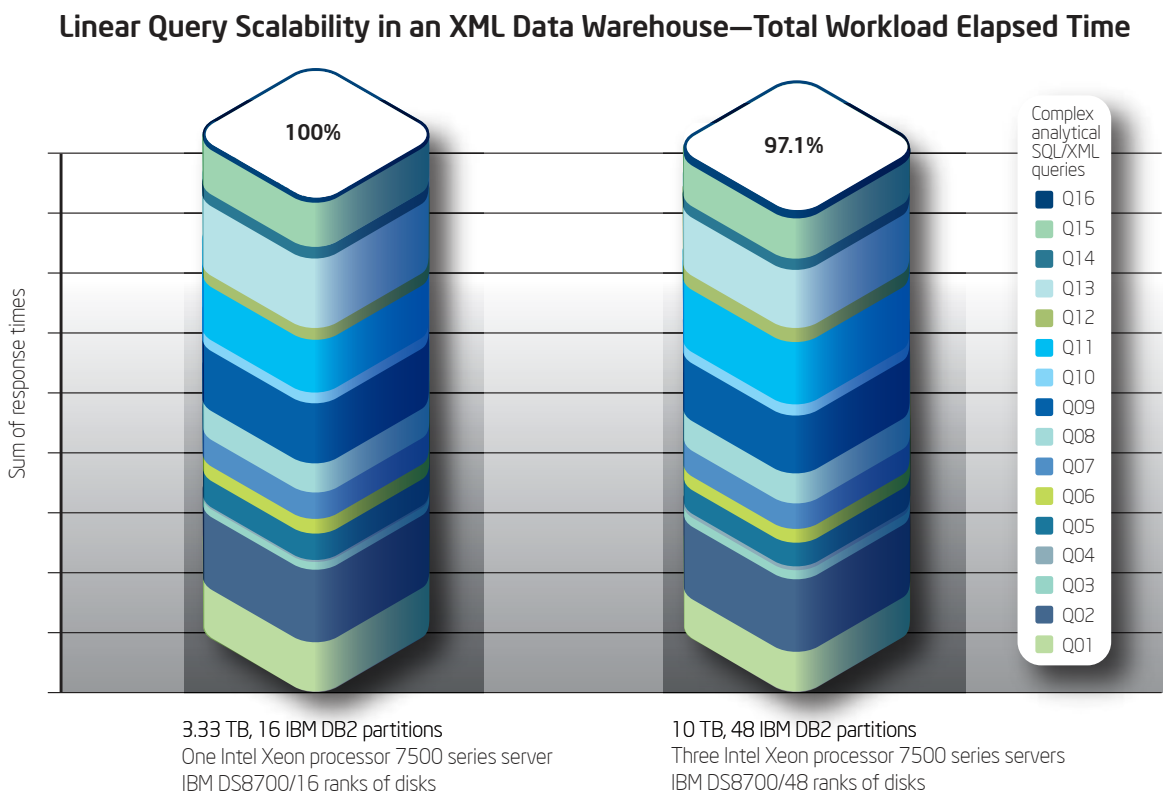

Figure 2: With the IBM-Intel solution, complex queries maintain consistent response times as both the size of the database and the number of servers triple.

processing overhead consumed by VMware on the Intel Xeon processor 7500 series is lower than the amount consumed on the previous Intel Xeon processor 7400 series–based hardware.[4]

These benchmarks are strong proof points of the value of the collaborative engineering efforts by IBM and Intel to maximize the innovations in DB2 9.7 and the Intel Xeon processor 7500 series.

Large XML Data Warehouses Become a Reality with Intel, IBM InfoSphere Warehouse, and IBM Storage

To determine if an Intel-based solution could sufficiently support a growing, multi-terabyte data warehouse, Intel and IBM paired InfoSphere Warehouse 9.7 with servers based on the Intel Xeon processor 7500 series and IBM System Storage DS8700. With this powerful reference configuration, the team studied response times for 16 complex analytical SQL/XML queries while growing the XML warehouse from 3.33 TB to 10 TB of raw data, representing up to 5.5 billion XML documents.

Figure 2 shows query response times at 10 TB that are nearly identical to those at 3.33 TB. In fact, IBM and Intel measured a difference in query performance of less than 3 percent between the two scale factors of the reference configuration.[5]

These initial tests demonstrate that IBM InfoSphere Warehouse 9.7 paired with the new Intel Xeon processor 7500 series and IBM System Storage DS8700 can provide linear query scalability as XML data warehouse size and complexity increase. Together, IBM and Intel have shown that as Intel compute resources are scaled in tandem with the data volume, the performance of the decision-support workload remains constant and predictable—a pivotal consideration for any high-performance data warehouse deployment.

Learn More

More about IBM DB2 9.7: www.ibm.com/db2

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

More about IBM System Storage DS8700: www.ibm.com/systems/storage/disk/ds8000/index.html

More about the TPoX 2.0 benchmark: http://tpox.sourceforge.net

---

[1] Source: IBM DS8000 performance team internal testing results. www-03.ibm.com/systems/storage/disk/ds8000/index.html

[2] Claim: "up to 20:1 server consolidation" based on comparison between 4S Intel® Xeon® MP CPU 3.3 GHz (single core with HT, 1 MB L2, 8 MB L3, Potomac) and 4S Intel® Xeon® X7560 2.26 GHz (8 core)–based servers. Calculation includes analysis based on performance, power, cooling, electricity rates, operating system annual license costs, and estimated server costs. This assumes 42U racks, US$0.10 per kWh, cooling costs are 2x the server power consumption costs, operating system license cost of US$900/year per server, per server cost of US$36,000 based on estimated list prices, and estimated server utilization rates. All US dollar figures are approximate. SPECint_rate_base2006* performance and power results are measured for X7560 and Xeon 3.3 GHz–based servers. Platform power was measured during the steady state window of the benchmark run and at idle. Performance gain compared to baseline was 20x.

–Baseline platform (measured score of 33.8): Intel server with four Intel Xeon MP CPU 3.3 GHz (single core with HT, 1 MB L2, 8 MB L3) processors, 16 GB memory (8x2 GB DDR2-400), two hard drives, one power supply, using Red Hat EL 5.3 x86_64 operating system.

–New platform (measured score of 705): Intel internal reference server with four Intel Xeon processor X7560 (24M cache, 2.26 GHz, 6.40 GT/s Intel QPI, Intel Hyper-Threading Technology, Intel Turbo Boost Technology), 256 GB memory (64x4 GB QR DDR3-1333), one hard drive, two power supplies, using SuSE* LINUX 11, cpu2006.1.1.ic11.1.linux64.binaries.nov242009 .tar.bz2 binaries.

[3] Source: Intel internal measurements as of February 2010. All of the benchmark results presented are derived using the TPoX 2.0 mixed benchmark simulating an OLTP workload with users running a mixture of read-only queries, inserts, deletes, and updates over a two-hour period. The resulting database size was approximately 1 TB. Platform: server with Intel Xeon processor X7500, 2.27 GHz; 256 GB versus Intel Xeon processor X7400, 2.67 GHz; 64 GB. Both platforms were running IBM DB2 9.7 on Linux SLES10-SP2 64-bit with DB2 Compression and STMM.

[4] Same workload as footnote 3, except that IBM DB2 9.7 on Linux SLES10-SP2 64-bit was run with VMware* ESX* 4.0 Update 1. Platform: server with Intel Xeon processor X7500, 2.27 GHz; 256 GB versus Intel Xeon processor X7400, 2.67 GHz; 128 GB.

[5] 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 IBM 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.

When Data Wears Out Its Welcome

Are cold data records gumming up the performance of your servers?

Robert Catterall
(rcatterall@catterallconsulting.com) is president of Catterall Consulting, a provider of DB2 consulting and training services.

A growing database can be a good news/bad news thing. On the one hand, a database that's getting bigger is something to smile about if it reflects growth in a company's business: more sales, more customers, more prospects. On the other hand, bigger can be bad when the increase in data volume occurs largely as a result of records going into the database and never coming out.

There comes a time when old rows in a table—particularly in tables that record business events such as sales transactions or insurance claims or stock trades—are just taking up space. They are rarely, if ever, requested by applications, so they provide little in the way of value. What you need to do is get them out of the database (without necessarily getting rid of them entirely; more on that momentarily). If you leave them where they are, they'll cost you—big time.

The high cost of old and cold data

I sometimes think of records in a database as having a temperature. Those that are frequently retrieved or updated by users or applications are hot. Those that are never or rarely retrieved or changed are cold. Generally speaking, the older a record gets, the colder it gets (so-called code or reference tables, which tend to be very static in nature, are an exception to this rule). A buildup of cold data in your database can cost you in several ways:

- **More expensive SQL (part 1):** The top determinant of an IBM DB2 application workload's CPU cost is the number of pages that DB2 must access in processing the workload (these accesses are known as GETPAGEs in a DB2 for z/OS environment and as logical reads in a DB2 for Linux, UNIX, and Windows environment). If data in a table is *not* clustered by date—if, for example, it is clustered by customer ID or order number—the percentage of rows in each page that are relatively cold will increase over time if old rows are not periodically removed from the database. That, in turn, can drive up page accesses (and, therefore, the CPU cost of SQL statement execution), because now DB2 must look in more pages to retrieve the same number of hot rows (and remember, these are the rows that are most frequently requested by users).

- **More expensive SQL (part 2):** You might be thinking, "That more-expensive SQL problem is not one that I face, because my business-event–recording tables are clustered by date." It's true that such a clustering scheme will concentrate cold rows at one end of the table, while pages at the other end of the table are dense with hot rows; however, there's the little matter of index access to consider. As you probably know, a DB2 index is logically structured like an upside-down tree, with a root page at the top, leaf pages with index key values and corresponding row IDs at the bottom and, unless the table is rather small, some number of levels (usually one or two) of non-leaf pages in between. When a table grows because cold rows are never deleted, eventually an n-level index on the table will become an n+1-level index. When a three-level index goes to four levels, the number of page accesses required to retrieve a row by way of an index goes up by 25 percent, from four (index root page, non-leaf page, leaf page, and table page) to five (an additional non-leaf index page).
- **More I/Os, longer run times:** As table pages become ever more full of cold rows (and index leaf pages with cold key values), the effectiveness of the buffer pool configuration goes down unless buffer pools are enlarged accordingly—another cost. Fewer buffer pool "hits" mean more I/O operations, and that increases SQL statement run times.
- **More expensive utilities:** Utility processes that are executed on a regular basis, such as those that back up and reorganize database objects and update catalog statistics used for query optimization, consume more CPU time as tables and indexes grow.
- **More disk space:** Obviously, the larger the database, the greater the amount of disk space required to hold the database objects. The cost per megabyte of disk storage is less than it used to be, but no one's giving the stuff away.

So rows in your database that are never, or hardly ever, retrieved or acted upon by programs increase the cost of your data-serving system on multiple fronts. Getting those cold rows out of the database would be a good thing, but don't take an ax to the problem. You'll need to think it through carefully.

Wouldn't it be nice if the same data archive could be used by different source database systems running on multiple hardware/OS platforms? Infrastructure simplification, anyone?

From the database to where? And when? And for how long?

A well-thought-out data archive and purge strategy can be a big help in maintaining the cost efficiency of your operational, run-the-business database systems. But cost control is not the whole story—there are also implications for regulatory compliance. The bottom line is that data archiving and purging can deliver a multitude of benefits if you do it right, and doing it right means effectively addressing several key questions, including:

- **When should data records be removed from operational databases?** Sometimes, this is a relatively easy question to answer, such as when the age of a record is the sole criterion for deletion from a table. For example, you might have an application that provides users with detailed account activity information over the past 12 or 24 months. In other cases, purge criteria can be more complex. It may be that a row recording a transaction cannot be deleted from a table if a certain related business event—perhaps recorded in another table—has or has not occurred. Establishing proper purge criteria requires a thorough understanding of the business processes associated with the data in question.
- **One-step or two-step removal?** It may be that an old, cold record is taken right out of the database at the appropriate time. Alternatively, a record may be moved, but not deleted, when it goes from hot to warm (accessed occasionally, but not frequently). Suppose 90 percent of requests for data in a table are for rows that have been inserted within the last 3 months, but you must support an application requirement for 12-month history retrieval. Given such a scenario, an organization may opt to keep the most recent 3 months of data in one table and data for the succeeding 9 months in another table—an arrangement that could help to optimize the performance of the "90 percent" data requests.
- **Into the bit bucket, or into an archive?** When you remove a record from an operational database, can it truly be thrown away, or must it be retained for legal reasons? If the latter, how long must it be retained?
- **What about the physical location?** Typically, access to archived data will be infrequent. That being the case, does it make sense to store archived data on expensive, high-performance disk systems? A better fit would likely be a higher-density storage device that provides online accessibility with reasonable performance. And while you're at it, wouldn't it be nice if the same data archive could be used by different source database systems running on multiple hardware/OS platforms? Infrastructure simplification, anyone?
- **What about security?** Your organization has probably taken pains to ensure that

data in your operational databases is properly secured, protecting it from unauthorized access. Will protection be similarly robust for archived data? It had better be, or good luck passing that audit.

- **How about usability?** OK, so you have a place where a record removed from the operational database can be economically and securely stored for a long time. Suppose someone needs to retrieve that record? Can it be quickly and easily located? Can the information in the record be presented in a usable format? Is there any metadata (data about the data) associated with the record that could make it more readily usable?

That's a pretty good list, there. How are you going to tackle it?

Build or buy?

It wasn't so long ago that do-it-yourself was a common organizational approach to data archive and purge (if, indeed, an organization did anything at all about it). Then came a truckload of legislated data retention and protection rules, a lot of mergers and acquisitions, more-complex criteria concerning the deletion of data from operational databases, and initiatives around simplifying and rationalizing IT infrastructure. You can still go the roll-your-own route, but that order has become a lot taller of late.

Want to explore vendor offerings? The demand for robust and flexible data archive and purge solutions has increased considerably in recent years, and a number of companies compete for that business (IBM is a big player in the market, with its Optim line of information lifecycle management products). An investment in a packaged archive and purge solution might offer an attractive payback for your organization.

Get your data right-placed

Is your database like a hotel where records check in, but never check out? That situation can be pretty sweet for the people who sell you servers and storage, but it sure doesn't help your company's bottom line. Get a handle on your data, figure out what doesn't belong in your operational databases, and implement a process that will move the cold data to a safe, secure, and economical archive. You've taken care of those old records for a long time. It's time they got a place of their own.

RESOURCES

IBM Optim Integrated Data Management solutions: **ibm.com**/software/data/optim

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

Will you be on the **short list**?

IBM
data management
KNOWLEDGE. PERFORMANCE. RESULTS.

Subscribe
ibm.com/dmmagazine

Advertise
650.919.1200
dm_adsales@tdagroup.com

- When readers turn to *IBM Data Management* magazine, they're actively searching for solutions to their most important problems. You can help them, but will they be able to find you before they start their next project?
- When you advertise in *IBM Data Management* magazine, you connect with 75,000 of IBM's most important customers and prospects. Each quarter, a whole new set of these customers embark on solving a whole new set of problems. Be there when they need you.
- Contact us today at dm_adsales@tdagroup.com to learn more about advertising opportunities and options for helping IBM customers.

IBM.COM/DMMAGAZINE

Changes to the Cursor Stability Isolation Level: Part 2

Understanding Currently Committed behavior

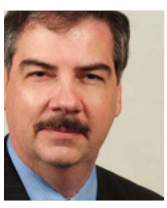

Roger E. Sanders
(roger_e_sanders@yahoo.com) is a consultant corporate systems engineer at EMC Corporation. He is the author of 18 books on DB2 for Linux, UNIX, and Windows and a recipient of the 2008 IBM Data Champion Inaugural award. He is currently working on a new book that outlines how to write technical magazine articles and books and get them published.

Special thanks to Senior Technical Staff Member–DB2 Kernel Architect Mike Winer for providing information used to develop this article.

In my last column, I explained how isolation levels play a key role in preventing databases from becoming inconsistent in multi-user environments. I also showed how the Cursor Stability (CS) isolation level worked both before and after lock avoidance techniques were implemented in IBM DB2 9.5. In this column, I'll introduce you to the Currently Committed (CC) semantics that were included in DB2 9.7, and I'll show you how these semantics provide faster data access and increased data concurrency for SQL statements running under the CS isolation level.

Currently Committed semantics

DB2 9.7 introduces a new implementation of the CS isolation level that incorporates CC semantics to further prevent writers from blocking readers when this isolation level is used. The intent is to provide a CS isolation level that avoids lock waits without violating CS semantics (like the use of the DB2_SKIPDELETED and DB2_EVALUNCOMMITTED registry variables does—see sidebar, "Registry variables used to delay or avoid acquiring locks in some circumstances").

Using full lock avoidance techniques (as described in my last column), a read-only SQL statement operating under CC semantics will not acquire a lock as long as DB2 can determine that the data needed has been committed. However, if DB2 is unable to tell whether the data needed has been committed, the transaction executing the statement will try to acquire a lock on the row in question.

If a lock can be acquired, processing will continue using traditional CS isolation level behavior. If a lock cannot be acquired (because another transaction holds an Exclusive lock on the row), DB2 will examine the lock that is held by the other transaction to obtain information about the row. The lock will contain one (and only one) of the following:

- **No information:** Indicates that the row is locked but nothing has been done to it (that is, no uncommitted changes are in flight)
- **An Uncommitted Insert identifier:** Indicates that the row is a newly inserted row that has not yet been committed
- **Log information:** Indicates that the row contains uncommitted data; in this case, the log information identifies the log record that corresponds to the first time the row was modified by the transaction that currently holds the lock on the row

If the lock contains no information, the row is treated as if the desired lock was acquired. If the lock contains an Uncommitted Insert identifier, the row is skipped since this identifier represents a row that has not yet been committed. If the lock contains log information, this information will be used to return the CC version of the row (that is, the row as it

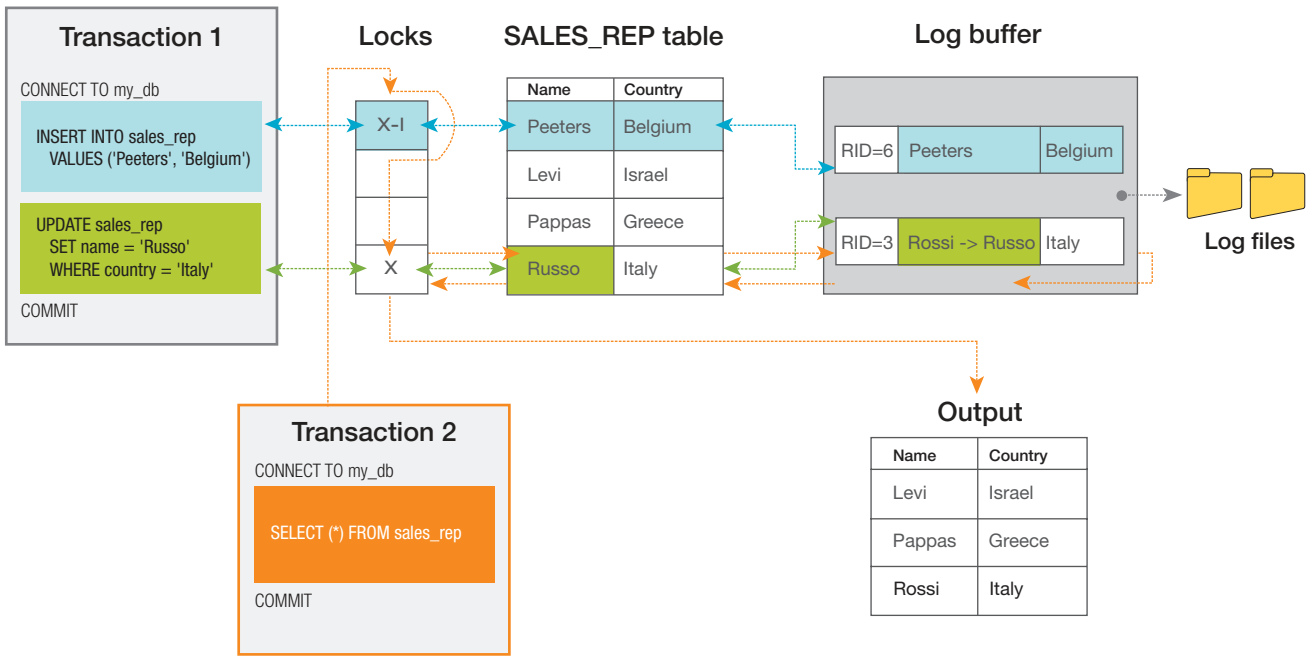

Figure 1: *In this example, Transaction 1 executed two DML statements, which caused log information and an Uncommitted Insert identifier to be written to the lock list for the SALES_REP table. When Transaction 2 queried the SALES_REP table, CC semantics allowed data for locked rows to be read from log records for previously committed transactions; the record for the Uncommitted Insert was not returned.*

existed before changes were initiated) from a log record stored in either the log buffer or from a transaction log file. DB2 uses the Log Sequence Number (LSN) to directly access the appropriate log record (see sidebar, "How DB2 can tell if data is committed").

Figure 1 illustrates how a SELECT statement using the CS isolation level with CC semantics will retrieve records. It is important to note that CC semantics can apply to SQL statements that are executed under both Read Stability (RS) and CS isolation levels. However, under the RS isolation level, CC provides only DB2_SKIPINSERTED behavior, which is the capability to no longer incur lock waits for uncommitted inserted rows.

Enabling Currently Committed semantics behavior

By default, CC semantics are turned on for new databases created with DB2 9.7. Existing databases that are upgraded to DB2 9.7 can take advantage of CC semantics by assigning either the value ON or the value AVAILABLE to the *cur_commit* database configuration parameter of the database that has been converted. If the *cur_commit* database configuration parameter is set to ON, CC semantics are applied database-wide for both the CS and RS isolation levels. If the *cur_commit* database configuration parameter is set to AVAILABLE, DB2 will store the appropriate information in locks and perform the extra logging overhead needed (to ensure that the logged data contains the full uncommitted version of the row being changed) to support CC semantics, but CC semantics behavior will have to be enabled on an application-by-application basis. This is done by either binding an embedded SQL application to the database using the CONCURRENTACCESSRESOLUTION USE CURRENTLY COMMITTED option or by specifying the SQL_ATTR_CONCURRENT_ACCESS_RESOLUTION connection attribute with CLI/ODBC and Java applications.

Note that the use of CC semantics requires an increase in log space for update operations to tables that are defined as DATA CAPTURE NONE. This additional space is used to log the first update of a data row by a current transaction; this data is used to retrieve the currently committed image of the row.

REGISTRY VARIABLES USED TO DELAY OR AVOID ACQUIRING LOCKS IN SOME CIRCUMSTANCES

When transactions are run concurrently, the following phenomena can occur:

- **DB2_SKIPINSERTED:** Allows CS/RS scans to skip uncommitted inserted rows
- **DB2_SKIPDELETED:** Allows CS/RS scans to skip uncommitted deleted rows and index keys
- **DB2_EVALUNCOMMITTED:** Allows CS/RS scans to apply and perform query predicate evaluation on uncommitted data; also allows the scans to skip uncommitted deleted rows. In effect, it treats the scan as an Uncommitted Read until it finds a qualifying row, at which time it may need to lock to ensure only committed data is processed/returned.

HOW DB2 CAN TELL IF DATA IS COMMITTED

All data row and index entries have a "flags" byte that contains a "Possibly UNCommitted" (PUNC) bit. If the PUNC bit is not set, the data row/index entry is guaranteed to be committed; otherwise the commit status is unknown.

Pages contain a "pageLSN" that identifies the LSN of the log record that corresponds to the last modification made to the page. If the pageLSN is older than the database's commitLSN or a table's readLSN, then the row/key is guaranteed to be committed; otherwise the commit status is unknown.

Conclusion

Maintaining data consistency and integrity, while allowing multiple transactions to access the same data at the same time, can be challenging, particularly in multi-user environments. Transactions, isolation levels, and locks are mechanisms to help maintain data consistency, and the CS isolation level tends to be used more often because it provides the greatest amount of concurrency while preventing dirty reads (which occur when a transaction reads data that hasn't yet been committed). Prior to DB2 9.7, use of this isolation level could result in delayed response times because of lock waits.

CC semantics provide a whole new implementation of the CS isolation level that does not violate ANSI standards for CS isolation levels. With CC semantics, transactions performing read and write operations avoid lock waits on uncommitted inserts, and transactions performing read-only operations end up trading a lock wait for a log read when they encounter uncommitted updates/deletes from concurrent transactions. This behavior practically guarantees that application response times will no longer be affected by locking.

Coming Attractions

Don't miss the next issue of *IBM Data Management* magazine, where you'll find:

- **Ideas, strategies, tactics, and hands-on projects** to help you cut costs and work more efficiently
- **Hard-won lessons from the energy and utilities industry** on long-term data archiving and regulatory compliance
- **Advice and tips** from our expert columnists, including more on data warehouse schemas in Informix DBA
- **A technical tour** of the next phase of database protection and monitoring with IBM Guardium

IBM data management
KNOWLEDGE. PERFORMANCE. RESULTS.

Subscribe
ibm.com/dmmagazine

Advertise
650.919.1200
advertise@tdagroup.com

IBM.COM/DMMAGAZINE

Reducing Conversations with DB2 for z/OS: Part 2

More tried-and-true basics and a few new bells and whistles

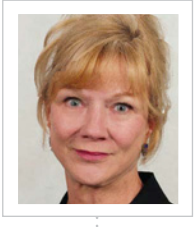

Bonnie Baker
(bkbaker@bonniebaker.com) specializes in teaching on-site classes for corporations, agencies, and DB2 user groups. She is an IBM DB2 Gold Consultant, an IBM Information Champion, a five-time winner of the IDUG Best Speaker award, and a member of the IDUG Speakers' Hall of Fame. She is best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

In my previous column, I wrote about unnecessary SQL. The goal was to make programmers stop and think about each statement and make sure that it was both absolutely necessary and absolutely necessary at that specific time. I gave examples of SQL that could be eliminated completely and SQL that could be deferred until a later time, until the last possible moment, and then executed only if still necessary.

In this column, I will continue with examples of SQL that may be executed far too many times and some that should be replaced with newer, better-performing SQL. So, welcome to part two of a three-part series on eliminating or reducing conversations with IBM DB2.

Tried-and-true techniques

Control break logic

One of the fundamental rules for performance that we should obey when writing programs, especially batch programs, is to use control break logic whenever possible. This is not exclusively a DB2 concept—it's just good programming practice. Checking for a control break, or a change in a value, reduces connections to DB2 because you look up a value in a table only if the value you are worried about is different from the last value that you looked up.

Optimal control break logic requires that the input data be in a specific order; for example, ITEM within STORE within REGION. But even if the input data is not sorted in a specific order, break logic can still be used, ensuring that any SQL that is required for a REGION is performed only once for each REGION; for a STORE, only once for each STORE within that REGION; and for an ITEM, only once for each ITEM within that STORE.

As a side note, DB2-enforced referential integrity does not use control break logic. Foreign-key values are checked on every INSERT, UPDATE, DELETE, and MERGE, even if the value of the column at issue is the same as the value in the row INSERTed (and so forth) before. With program-enforced referential integrity, we have the option of doing the validation only when the value of the column at issue is different from the value in the prior row; that is, on a control break.

Preloading small reference tables into program memory

Another fundamental rule for performance is to preload small reference tables into working storage tables (when it is smart to do so) to avoid connecting to DB2 an inordinate number of times.

A less-than-amiable person once asked me, "What kind of crummy product makes it necessary to load values into working storage to avoid connecting to the product?" I thought for a minute and said, "The kind of product that recognizes that a programmer has choices and a brain and understands that avoiding unnecessary calls, connects, and get pages is a good idea."

Just because DB2 is fast and has buffer pools to reduce actual I/O does not mean that connecting to DB2 to read a 10-row table a million times is smarter than connecting to DB2 once, reading 10 rows into a working storage table, and then addressing program memory a million times.

Reduce repetitive executions of subselects

Consider this SQL:

```
Select col1, col2, col3
From big_table where item = :hv-item-just-read
        And big_table.deptno in (select deptno from little_dept-table
                                 Where division = Eastern)
```

For each item on an input sequential file, the program logic requires a lookup in a big table to get the associated values for col1, col2, and col3. Hopefully, as recommended earlier, the input data set has been sorted by item number *and* the program is checking to see if this item number is different from the prior number before this lookup is done.

But look more closely at the SQL. The subselect is creating a list of departments that are in the Eastern division and then making sure that the department at issue is in that list. I'm sure that, with a wee bit of thought, we could create that "departments in the Eastern division" list once and then do our lookup without connecting to DB2. For example, we can use a V8 rowset-positioned cursor to connect to DB2 one time and fetch our list; then we change the subject SQL to a hard-coded in-list built from our array:

```
Declare CursorDept with rowset positioning for
    Select deptno from little_dept-table
    Where division = Eastern)
Fetch next rowset from CursorDept
    For 20 rows
Into :hvarray-deptno
(code to ensure that +100 was reached)
Select col1, col2, col3
 From big_table where item = :hv-item-just-read
    And big_table.deptno in (:hvdept1, :hvdept2, :hvdept3....:hvdept20)
```

Remember, if there are fewer than 20 department numbers (say, 15) and you do not want to code dynamic SQL, you can always perpetuate the last value in host variables 16 through 20.

Reading a row before update/delete—fetch/update/fetch/update

More-experienced programmers must overcome some of their entrenched practices. In the "old days," we had to read a record to update or delete it. With DB2, if you have no need to establish a before image of the row, you do not need to read the row to update or delete it. I often see batch programs that use a cursor to fetch, update, fetch, update, fetch, update..., when they could just do a reasonably sized relational SET UPDATE using the same WHERE clause that is in the DECLARE CURSOR. Instead of:

```
Declare CursorUpd for
    Select col1, col2, ...
        From tableA
        Where jobcode = :hvj for update of salary

Fetch CursorUpd into :hvcol1, :hvcol2, ...
Update salary set salary = salary + 1000.00 where current of CursorUpd
```

and repeating the fetch, update, fetch, update again and again, why not just:

```
Update tableA
    Set salary = salary + 1000.00
        Where jobcode = :hvj
```

SET processing is one of the huge strengths of relational design. You must, of course, ensure that the SET is a size that does not create an unacceptable unit of recovery or a problem with lock escalation, or cause locking issues (such as timeouts) for other concurrently running programs.

I sometimes see the preceding scenario (fetch, update) compounded by a subsequent singleton SELECT of the updated row—read it before, update it, look at it after. With newer programming techniques, as well as the realization that we do not need to see the before image, we can DECLARE a row-positioned cursor to SELECT from our SET UPDATE and then FETCH a rowset of reasonable size (say 100 rows at a time) to see the result of the maintenance.

Single-row fetches are becoming rare. Why connect 100 times when you can connect once and see the same 100 rows?

Stay tuned for part 3

In the next issue, I will continue with the final installment of this topic, and we will look at even more (and newer) techniques for avoiding or reducing connects to DB2. ✱

RESOURCES

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

The Mystery of DB2 Sorts:
www.dbmag.intelligententerprise.com/story/showArticle.jhtml?articleID=211300275

Programmers Only archive:
www.dbmag.intelligententerprise.com/columns/programmersonly.jhtml

Building Fast Data Warehouse Schemas

Optimize your data warehouse design, starting with the fact 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.

Why build a dimensional data model or star schema for a data warehouse? Why not use your existing transaction system? In this and the next couple of articles, we will answer those questions by taking a look at how to design a data warehouse schema.

The goal of a data warehouse is to give users fast access to data so they can make better business decisions. The data needs to be correct, and users need to be able to query and analyze it. In fact, the more users are able to analyze data (I call this playing with data), the better they learn and gain new insight.

These goals are not easy to accomplish with a regular transactional operational database, where the priority is to get records entered fast. You don't want users burning up CPU cycles running queries and analysis on an operational database, and you don't want those queries locking records that need to be open for data entry. On one of my very first data warehouse jobs, we decided to test our query tool against the operational system. Within minutes we got a call from the DBA asking why we were locking so many records in the database!

Data warehouses give users a separate area where they can do queries and play with data without affecting the operational database. And now that we've created a separate copy of data for the sole purpose of analysis, we can look at other ways to optimize it for the intended function, starting with how it's organized.

Just the facts

Dimensional modeling is the design methodology used to organize the data in the data warehouse. There are three key parts of a dimensional model: the fact table, the dimension tables, and the summary tables.

Facts are things that you measure and record. Typically, a fact is a number. When I look at a new schema, the first thing I do is highlight all the columns that are decimal or number. That is my starting point for determining what the facts are in a schema.

Dimensions are reference information—they provide descriptive information about the facts. Let's take a very simple example, as seen in Figure 1. This is a billing system with a table called bill, a table called payment, and a table called apply payment. In the bill table, we have information about what was on the bill and the bill amount and the quantity. In the payment table, we have information about the payment that came in and a payment amount. In the apply payment table, we have information about which payment was applied to which bill. There are various combinations of

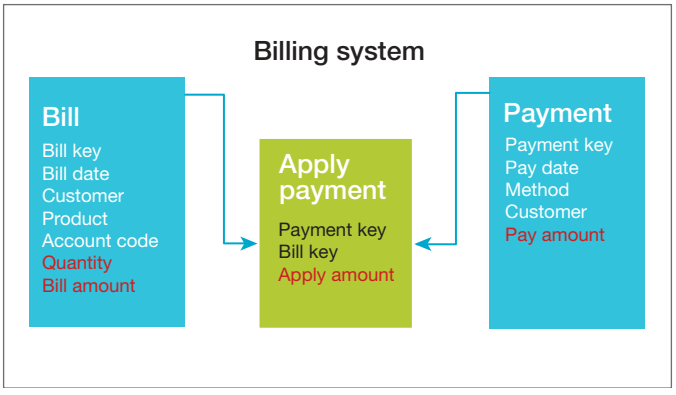

Figure 1: *Transactional schema*

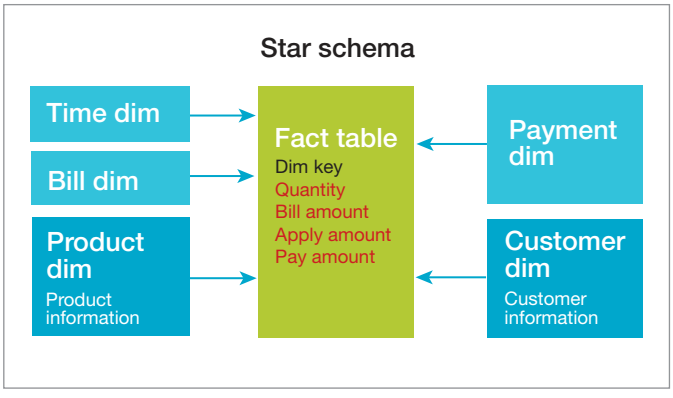

Figure 2: *Facts and dimensions*

Date	Bill no.	Customer	Product	GL codes	Quantity	Bill amount	Pay amount	Apply amount
1/1/2010	101	Smith	10	4000	1	100		
1/1/2010	102	Smith	52	4000	2	60		
1/5/2010		Smith		2000			160	
1/5/2010	101	Smith	10	2000				100
1/5/2010	102	Smith	52	2000				60
Totals						160	160	160

Figure 3: *The user view—a big spreadsheet*

payments for paying bills: one payment could pay multiple bills or pay only a part of a bill, or three payments could pay one bill.

This schema has four fields that are facts: quantity, bill amount, apply amount, and pay amount. All other fields are dimensional information describing the transactions.

Figure 2 shows what this example might look like in a data warehouse schema. There is a fact table that contains the facts. The payment dimension table describes things like the check number and the bank into which the check was deposited. A customer dimension table describes the customer, a product dimension table describes the product that was there, the bill dimension table describes the bill (bill date, terms, and so forth), and almost every data warehouse has a time dimension.

Basically what you're out to do in dimensional modeling is transform the schema from a transactional system (Figure 1), to something like Figure 2, where you have a dimensional model in place.

But the end users who are analyzing data do not want to see a schema. They want to see what I call "the big spreadsheet in the sky." They don't care about dimensions or about star schemas. What they want to see is a huge spreadsheet with all the data adding up correctly at the bottom.

This is the main reason we use a dimensional model: it is easier to create this spreadsheet in one query. If you do a `SELECT` from every table and every dimension in your dimensional model, you would get the big spreadsheet in the sky (see Figure 3).

The best way to describe the result of a dimensional model to people who are new to it is to view it as a big spreadsheet that has every dimension field and every fact field pre-joined together. It is then possible to work backward from that view to figure out which elements are facts and which are dimensions.

Fact table design

A fact table contains the numbers. It is central to the data warehouse. In addition to numbers, a fact table has a key to all the dimensional tables, so you can join from the fact table to the dimension tables.

Facts are number fields that can be added up most of the time. Some facts are different every time you measure them, such as temperature. I worked on an interesting data warehouse for a company that was manufacturing an item very sensitive to room temperature. The company needed to anticipate three days in advance what the temperature would be at the time of production, or the product would be defective. So the fact they were most interested in was a projection of what the temperature would be in three days, based on temperature changes over the last three days.

A key concept in building the fact table is to decide on the "grain," or granularity, of the fact. Grain is the level of detail that each record in the fact table will have, and there are a couple ways of looking at it. One type is a snapshot grain. A snapshot grain is usually something like an end-of-day snapshot, a monthly snapshot, or a weekly snapshot. The data is accurate to a point in time and static after that point in time.

The second type of grain is a transactional grain, which shows every little transaction that goes on. While a snapshot grain might show the state of a customer at the end of the day, a transactional grain will show every debit and credit transaction that took place with the customer over time.

The goal is to get the lowest possible grain, so always capture records at the lowest transactional level. It is much easier and faster to summarize than it is to reload a lower level of detail. You cannot drill down into detail when you don't have the data, so it is far better to put the lowest level of detail in the data warehouse from the beginning.

The starting point for a data warehouse design is the fact table, and you should now have an overview of what to look for and include in a fact table. In the next issue, we will look at dimension and summary tables.

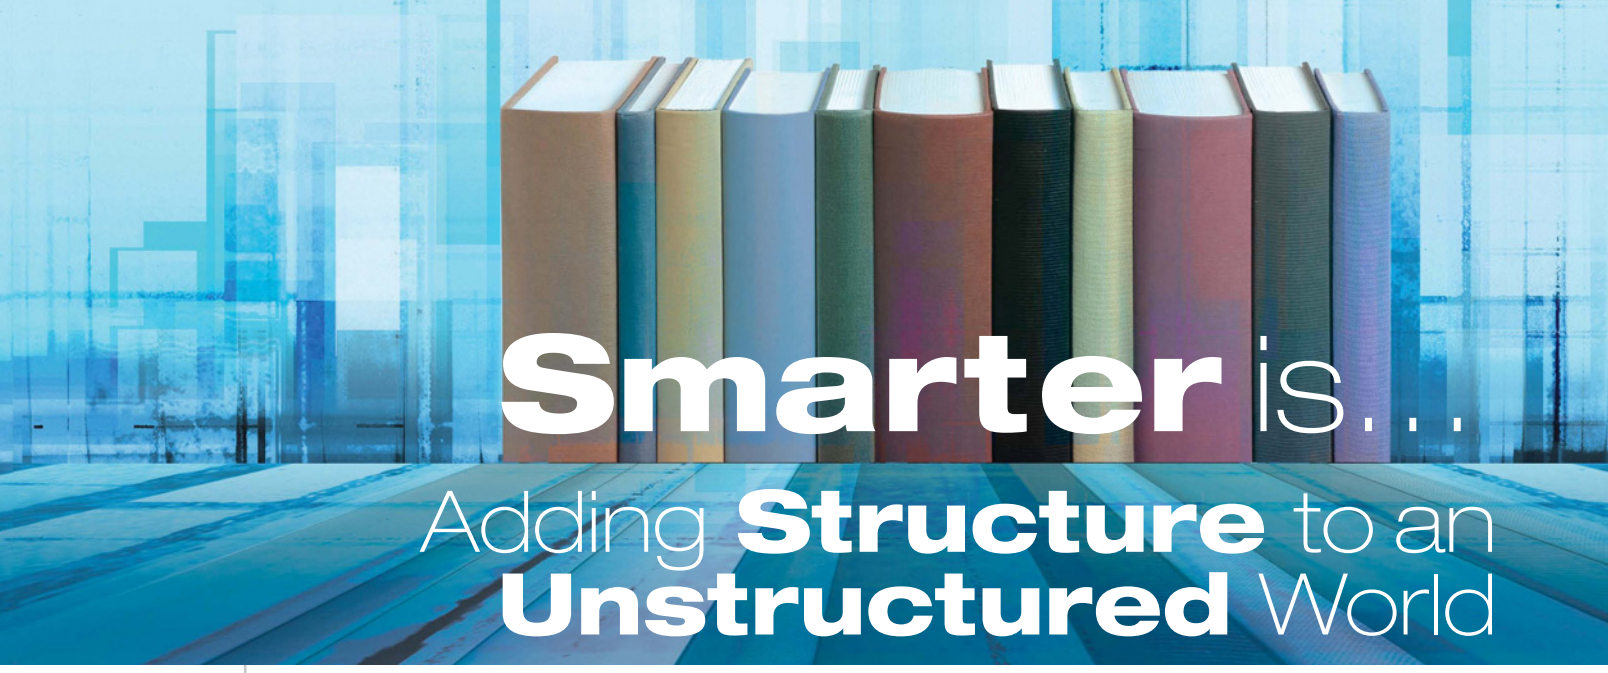

Smarter is...

Adding Structure to an Unstructured World

IBM BigSheets helps organizations extract big value from unstructured data

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

Data managers have created immense value from structured information. Now the challenge is to pull in data from the unstructured world—and mix it with internal data stores to gain fresh perspectives. And there's no problem finding unstructured data to analyze: David Boloker, chief technology officer for emerging technologies at IBM, estimates that of the 15 petabytes of data created around the world each day, about 80 percent comes from unstructured sources.

"The most daunting part of the challenge isn't collecting unstructured data, it's getting value from it," says Boloker. "Take the example of a pharmaceutical company that has a drug in clinical trials. Much of clinical data is unstructured, handwritten in patient records and then digitized. If there was a way to quickly and easily refine that data into a more structured form, the company might confirm benefits of the drug much earlier in the process or spot subtle problems that might otherwise be missed."

The British Library was grappling with just such a challenge. Tasked with archiving information from across the published spectrum, the staff needed a way to turn massive amounts of data from Web sites and other unstructured sources into a viable resource. Working with IBM, the library successfully implemented a prototype analytics technology called IBM BigSheets (**ibm.com**/software/ebusiness/jstart/bigsheets).

With IBM BigSheets software, users are able to access vast archives of data, submit queries to easily research the information, analyze it in a format that is organized like a spreadsheet, and explore it in other familiar visual contexts. For example, users can see search results in a pie chart and look at the data in a tag cloud. "As a data manager, my question is, 'How can I make all the unstructured information coming at me useful to my organization?' Now I have an answer," says Boloker.

Under the hood, BigSheets is built on the Apache Hadoop open-source framework for parallel processing large data sets on compute clusters, and it uses the Hadoop Distributed File System (HDFS) for high-throughput access to application data. The BigSheets software collects information from a variety of source applications, extracts the data, annotates it with tags, and enriches it for display.

BigSheets is already enabling the British Library to extract big value from unstrucutured data, but Boloker expects the technology to have an even larger impact in science, academia, and the private sector. "A business could do things like matching unstructured data from a given zip code to internal sales data and see what's causing an up or down trend," he explains. "We're now able to use information that was lost in the unstructured world and compare or contrast it with what we already have. It's really a new day for data managers and their clients."

TDA GROUP
800 West El Camino Real, Suite 380
Mountain View, CA 94040
U.S.A.