

data management

IBM.COM/DMMAGAZINE

KNOWLEDGE. PERFORMANCE. RESULTS.



Save Time, Save Money

DB2 9.7: TUNED FOR ACTION

Tighter Compression,
Enhanced Automation,
and SQL Compatibility

CONSOLIDATED GAINS

How Informix Optimizes
Storage for Data Warehousing

HEALTHCARE'S XML HEARTBEAT

SOA and pureXML Ease
Medical Data Management

Dive Deeper

With *IBM Data Management* magazine, you get more than a surface glance: more detail, more information, more context, and more viewpoints on the technologies and strategies you deal with every day.

IBM
data
management
KNOWLEDGE. PERFORMANCE. RESULTS.

▶ **IBM Data Management digital edition**

Sign up for easy online access and quick content searches across the entire magazine.

▶ **IBM Data Management community**

Explore current and previous issues, and post your views to the magazine's blog.

▶ **My developerWorks and developerWorks DBA Central**

Customize developerWorks to deliver your information, your way—news, blogs, downloads, feeds, forums, and more.

Dive deeper.
Get to the bottom of things at:

IBM.COM/DMMAGAZINE



data management

KNOWLEDGE. PERFORMANCE. RESULTS.

ISSUE 2, 2009 / VOLUME 14 NO. 2

Features

DB2 9.7: It All Adds Up

Save time and money with the enhanced compression, virtualization, and management features in IBM DB2 9.7 for Linux, UNIX, and Windows

14

19 SQL Compatibility in DB2 9.7

IBM DB2 9.7 is a developer's dream. Try out some of the new Oracle PL/SQL compatibility features in our hands-on demonstration

22 SQL Tuning: Not Just for Hardcore DBAs Anymore

IBM Optim Query Tuner helps developers create efficient queries and build tuning skills



Pure ³⁰ Concentrated Data

Informix Dynamic Server adds speed and ease-of-use features to improve performance in data warehouse and business intelligence deployments



Industry Focus:
Healthcare

Healthcare's XML
Heartbeat

26

The healthcare industry is using XML—and DB2 pureXML technology—to keep up with vast amounts of rapidly changing data while protecting patients' security and privacy

Smarter Is...

40

**Predictive, Preventive
Healthcare**

How clinical data warehouses are turning information into smarter treatment plans



DB2 pureXML anchors the patient-oriented document system at UCLA Health System

Departments

4 **Editor's Note**

By Cameron Crotty

6 **NewsBytes**

9 **IDUG User View**

By David Beulke

10 **IIUG User View**

By Stuart Litel

12 **Data Manager**

The Performance and Cost Optimization Ripple Effect
By Scott Hayes

32 **Data Architect**

DB2 Data Warehouse Performance, Part 2
By Robert Catterall

34 **Distributed DBA**

Reducing Downtime with HADR
By Roger E. Sanders

37 **Informix DBA**

Cool New Stuff for Informix
By Lester Knutsen

39 **Skills Zone**

How to Get Certified for DB2
By Howard Fosdick

Ad Index

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

Relational Architects International..... **5**
www.relarc.com

International Informix Users Group..... **C3**
www.iiug.org

Optimize Both Performance & DB2 Memory Utilization With **Buffer Pool Tool[®] For DB2** Reduce Your Processing Costs, Save \$\$ Now!

PRODUCT FUNCTIONS	BUFFER POOL TOOL	BMC POOL ADVISOR	IBM POOL ANALYZER	UBS HAINER BPA4DB2
Collects trace data for the critical performance period	✓		✓	✓
Easily processes up to 16GB of trace data	✓			
Huge performance trace must be down-loaded to the Workstation			✓	✓
Simulates/predicts pool & object I/O performance	✓			
Simulates/predicts object partition usage & I/O rate	✓			
Simulates I/O for moving objects to other (& new) pools	✓			
Simulates/predicts Group Buffer Pool performance	✓			
Proven I/O rate prediction accuracy, ask our clients	✓			
Provides a <i>proven methodology</i> for pool tuning	✓			
Proven low CPU overhead for data collections	✓			
Continual CPU overhead on your processor		✓	✓	✓
Reports memory usage & paging of DB2 & entire LPAR	✓			
Shows you how to reduce transaction elapsed times	✓			
Windows based graphic performance analysis	✓		✓	✓
Performance Tuning Wizard	✓			
Reports CPU costs by object for I/O and Scan activity	✓			
Easily installs in less than half an hour	✓			
Client references that will discuss their savings	✓			

**RESPONSIVE
SYSTEMS**

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

IBM data management

KNOWLEDGE. PERFORMANCE. RESULTS.

EDITOR

Cameron Crotty
editor@tdagroup.com

MANAGING EDITOR

Stephanie S. McLoughlin

ART DIRECTOR

Iva Frank

DESIGNERS

David Chan, Lalaine Gagni, Margie Preston

CONTRIBUTING WRITERS

Merv Adrian, Eric Ahrendt, Bonnie Baker, David Beulke, Robert Catterall, Howard Fosdick, Scott Hayes, Lester Knutsen, Cliff Leung, Tony Leung, Stuart Litel, Anil Mahadev, Ken North, Roger E. Sanders, Bob Scheier, Ray Willoughby, Kathryn Zeidenstein

EDITORIAL BOARD OF DIRECTORS

Rick Myllenbeck (Chair), Scott M. Bisang, Felicity Carson, Jeff Jones, Kimberly C. Madia, Nancy Miller, Bob Sawyer, Kathryn Zeidenstein

AD SALES EXECUTIVE

Randy Byers
advertise@tdagroup.com

AD COORDINATOR

Katherine Hartlove

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/dmmagazine. For information about reprints, please send an e-mail to customerservice@tdagroup.com.

IBM and the IBM logo are registered trademarks of the International Business Machines Corporation and are used by TDA Group under license.

Material published in *IBM Data Management* magazine copyright © 2009, International Business Machines. Reproduction of material appearing in *IBM Data Management* magazine is forbidden without prior written permission from the editor.



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.



Welcome to the premier issue of *IBM Data Management* magazine!

You may already know us as *IBM Database Magazine*, a dependable resource that began as *IBM DB2 Magazine* in 1996. But the role of the data management professional is evolving—as it always has—so we're growing and changing to help you make the most of new opportunities.

Businesses run on data and information. Your organization depends on data flowing smoothly, and your colleagues depend on information being accurate. The challenge has always been to efficiently manage that data, but today that means more than just tuning databases or developing applications. It also means creating information architectures and finding new ways to make data work to your organization's best advantage.

Our goal is to help you meet those challenges. We're committed to delivering the same in-depth technical content, practical advice, and expert commentary that you've come to expect, with a new design that's more accessible and easier to navigate. Plus, we've added coverage on how data is shaping the way the world does business, and how organizations can use their information as an engine to drive business success.

But don't take my word for it; check out the stories that we've gathered for this issue. We have several angles on the big news this year so far: the announcement of DB2 9.7 for Linux, UNIX, and Windows. Start with our cover story to explore what's new and improved, then move on to "SQL Compatibility in DB2 9.7," and dig deeper with IBM Data Champion Anil Mahadev for a close look at one of the most talked-about features. Want more? Regular columnist and former IDUG president David Beulke checks in with a personal take on the new version.

David is just one of many familiar faces that longtime readers will recognize: Lester Knutsen, Robert Catterall, and Roger E. Sanders all return with hands-on technical advice and commentary. Next to these well-known favorites, we've added a new column, "Data Manager," that takes a broader business view. This month, DBI Software CEO Scott Hayes joins us to explain what every business needs to know about performance tuning—and how you can reap the benefits.

Read on and discover how XML is transforming the healthcare industry, what's new with Informix, and why everyone from DBAs to developers will be excited about the powerful new IBM Optim Query Tuner.

We hope you enjoy our new look and find the magazine more useful, relevant, and insightful than ever before. We invite you to tell us about how we can continue to improve by sending an e-mail to editor@tdagroup.com. And be sure to check out the magazine's new online home at ibm.com/dmmagazine.

Thanks for reading,

Cameron Crotty
Editor



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 Fortune 1000 organizations worldwide to run their mission-critical z/OS batch applications. It's the standard for z/OS batch restart.

Restart made simple™

DB2, Websphere MQ and IMS are registered trademarks of IBM Corp.

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

rai **Relational
Architects
International**

From Snapshots to Stream Computing

IBM System S software delivers real-time analysis of data flows

» IBM UPDATES INFOSPHERE DATA ARCHITECT

Version 7.5.2 simplifies process of estimating storage requirements

IBM InfoSphere Data Architect Version 7.5.2 features several improvements, including simplified storage needs estimating. Volumetric information is used to approximate storage requirements and forecast growth rates. Other new enhancements include copy/paste to invoke privacy attributes, mapping across multiple data projects, and toggling between tree navigation and flat navigation in the explorer to allow faster navigation between data sources.

» **MORE INFORMATION**
ibm.com/software/data/studio/data-architect

Today, the result of most data analysis resembles a single frame from a motion picture: a snapshot of one point in time. But now, after seven years of development, IBM has made its stream computing technology, evolving from its System S project in IBM Research, available commercially as IBM InfoSphere Streams.

The software uses a streaming architecture and algorithms developed by the IBM Research team to process thousands of simultaneous data streams, providing nearly instantaneous analysis that is continually refined as new data becomes available. System S software can help organizations respond to changing conditions in real time by seeing and responding to events across complex systems.

The System S architecture represents a significant advancement in the relationship between queries and data as well as in computing system organization and capability. Traditional data analysis applies changing queries to static data; System S evaluates constantly changing data against a set of defined queries derived from user expressions of information needs and interests.

User requirements are translated into specifications of how potentially available raw data and existing information can be transformed to satisfy organizational objectives. The run-time environment accepts these specifications, determines how it might reorganize itself to best meet those requirements, and automatically makes the

required changes. The run-time environment can now continually adapt to the state and utilization of its computing resources, as well as the information needs expressed by the users and the availability of data to meet those needs.

IBM is making System S trial code available at no cost. The code includes Streams Studio, the development environment, and Streams Runtime, to test applications. Toolkits and adapters are available as well. In addition, IBM announced the opening of the IBM European Stream Computing Center in Dublin, Ireland. The center will serve as a hub for research, advanced testing, and support for European clients interested in adopting stream computing technology.

» **MORE INFORMATION**
ibm.com/software/data/infosphere/streams

Data Studio Joins Optim Family

Enhanced toolset offers increased integration with extended support

IBM has expanded and extended the Optim portfolio. The IBM Optim line now includes Optim Development Studio; Optim pureQuery Runtime; Optim Database Administrator for DB2 for Linux, UNIX, and Windows; and Optim Query Tuner. Formerly encompassed under the Data Studio name, the Optim family also offers new performance-tuning capabilities; a common environment for developing Oracle, DB2, and Informix applications; and enhanced portfolio integration.

Optim Development Studio V2.2 and Optim pureQuery Runtime for Linux, UNIX, and Windows V2.2 both have extended support to Oracle databases, including native support for PL/SQL stored procedures. As a result, organizations doing development for multiple environments can use a single database development environment across Oracle, DB2, and Informix. Additionally, these releases support IBM DB2 9.7 and its new SQL compatibility features (see "SQL Compatibility in DB2 9.7," this issue).

Optim pureQuery Runtime improves SQL best practices by providing conversion and consolidation of literals. Use of literals instead of variables or parameters is generally discouraged within SQL because it makes each statement unique, which drives down cache hit ratios. Now, multiple statements that had each been unique

due to the use of literals will be consolidated into a single statement using parameters to improve application performance. In addition, Optim pureQuery Runtime V2.2 introduces a more lifecycle-driven means for capturing application SQL, simplifying the use of pureQuery over time. This release also introduces the ability to configure data source parameters on the application server online, increasing the flexibility of managing a pureQuery deployment.

Optim Database Administrator for DB2 for Linux, UNIX, and Windows simplifies DB2 database administration by providing an easy-to-use interface and built-in analysis to automate complex database changes and migrations with confidence. This release extends support for large database environments and disaster recovery.

MORE INFORMATION
ibm.com/software/data/data-management

IBM Acquires Exeros

Data discovery technology helps transform data into a strategic business asset



IBM announced that it acquired data discovery software vendor Exeros in May. The acquisition will help organizations using the IBM Information Management software portfolio to better manage their data.

Exeros technology automatically uncovers hidden relationships between databases, helping users make sense of disparate sources without the labor-intensive process of manually mapping

data. This capability helps generate new levels of intelligence that can identify emerging market trends and help predict business outcomes.

"All organizations today are faced with the daunting challenge of turning massive amounts of information into insights to guide their businesses, but many are held back by the complexity of corporate data sources," said Ambuj Goyal, general manager,

IBM Information Management. "The combination of IBM and Exeros will enable companies to more intelligently manage their data across all formats and computing platforms, creating a smarter enterprise."

Over the next several months, IBM intends to integrate Exeros with the IBM Information Management software portfolio. The acquisition will also further the IBM Information On Demand

initiative, which helps customers use information as a competitive and strategic business asset.

In addition, the acquisition will support the recently announced IBM Business Analytics and Optimization consulting practice to meet growing demand for advanced analytics capabilities to create new, more intelligent business systems.

MORE INFORMATION
www.exeros.com

» GUARDIUM EXPANDS SUPPORT FOR IBM DATABASES

Database security platform reaches more IBM platforms—and more customers

Database security company Guardium announced expanded support for IBM database software and operating systems. The Guardium database security platform uses real-time, policy-based monitoring to protect against data leakage, ensure data governance, discover sensitive data for regulatory compliance, and enhance database security postures. Paired with IBM software, Guardium database security offerings help customers better manage and protect their enterprise data. In fact, three of the top four global banks, several U.S. government agencies, and several major healthcare providers are among the organizations that depend on Guardium and IBM.

Guardium has now added support for some of the most popular IBM database platforms, enabling organizations to protect sensitive databases while reducing IT costs. Guardium has added support for IBM DB2 9 for z/OS; IBM DB2 for i; IBM DB2 9.5 for Linux, UNIX, and Windows; IBM Cognos 8; IBM Informix Dynamic Server 11.5; Red Hat Enterprise Linux for IBM System z; and SUSE Linux Enterprise Server for IBM System z.

» MORE INFORMATION
www.guardium.com



Oracle Migrations Made Easier and Less Expensive with DB2 9.7

IBM builds on EnterpriseDB technology to deliver Oracle compatibility

The latest version of DB2 for Linux, UNIX, and Windows boasts a number of benefits, including business analytics capabilities for both relational and XML data and reduced storage and energy consumption. But one of the biggest pieces of news in the DB2 9.7 release is its SQL compatibility features, including native support for Oracle PL/SQL.

To deliver these features, IBM augmented an in-house development initiative with technology licensed from EnterpriseDB.

As a result, organizations can now run applications written for Oracle Database on DB2 9.7 while reducing the cost, time, and risk of migrating applications to PL/SQL.

DB2 9.7 now supports Oracle Database data types, SQL syntax and semantics, functions, and built-in packages; the Oracle PL/SQL stored procedure language (including packages); and the Oracle Database call interface for application developers.

DB2 9.7 customers can now leverage the existing skills of

Oracle Database professionals; there is no need to retrain or hire specialized personnel to take advantage of new DB2 9.7 features. DBAs will also benefit from the ability to deploy new Oracle Database-compatible applications. They can deploy instances of DB2 9.7 and Oracle Database that interoperate seamlessly, and they can migrate existing Oracle Database applications to DB2 9.7.

» MORE INFORMATION
ibm.com/db2/9

IBM IMS Education and Updates

Three online teleconferences help you learn about the new features of IBM Information Management System (IMS) 10. Also, don't miss the video showing how E. Sun Bank in Taiwan is using IMS to create a highly reliable and scalable service-oriented architecture (SOA) solution.

Teleconference: IMS 10 Connectivity Enhancements with IMS Connect API

IMS Connect, the TCP/IP gateway to IMS transactions, operations, and now data, has been significantly enhanced. Learn how the new features, together with the IMS Connect API, position IMS Connect as the foundation for all future IMS connectivity.

Speakers: Jack Yuan, David Cameron, and Kevin Flanigan

ibm.com/software/os/systemz/telecon/may12

Teleconference: What's New in IMS Database Recovery

IMS 10 includes many new enhancements for Database Recovery Control (DBRC) that simplify DBRC access by vendor tools and ease version-to-version upgrades for customers using tools that interface with the RECON data set.

Speakers: Sally Gehring and Allen Hall

ibm.com/software/os/systemz/telecon/jun23

Teleconference: Data Sharing Management with IMS Sysplex Manager

This teleconference focuses on how IMS Sysplex Manager can help IMS users manage workload balancing and data sharing issues in a Sysplex environment.

Speaker: Andy Nguyen

ibm.com/software/os/systemz/telecon/jun30

Video: IMS Delivers Service at E. Sun Bank

E. Sun Bank in Taiwan enables quick business integration and exemplary service with an SOA solution based on IMS, IBM WebSphere, IBM Rational, and IBM System z.

ibm.com/software/data/ims/video-esunbank.html

First Look at DB2 9.7



David Beulke is former president of IDUG (www.idug.org) and a DB2 Data Champion. He has more than 22 years experience in architecture, design, and development of high-performance DB2 data warehouse and OLTP systems across all platforms.

New features improve performance and compatibility and are easy to implement

The new release of IBM DB2 9.7 for Linux, UNIX, and Windows offers many advantages for better performance, better maintenance, and lower total cost of ownership, especially for large and robust databases. But after exploring the new version myself, what strikes me most about the features is how easy they are to implement.

The first feature that makes this release really stand out is index compression, which reduces both I/O and the size of the index structure. Through this new compression technology, DB2 is capable of achieving compression ratios as high as 50 percent or more. This index size reduction has big implications for dramatically reducing disk requirements for terabyte data warehouse implementations and for reducing overall I/O and response time for transaction systems. With this feature, DB2 is keeping well ahead of the competition.

The next new feature within DB2 is an improvement to its transaction concurrency model. This concurrency enhancement provides more granularity for the transactions. Within some high-performance systems, one transaction writing data and another transaction reading data can be blocked or delayed before the writer transaction is committed. This delay or blocking situation can cause some high-performance

systems issues, so the DB2 team developed a new isolation level called Currently Committed that is now the default within the system.

The new Currently Committed isolation level allows writer transactions to continue processing while the other reader transactions within the system get their data. It works by retrieving the data row

Your application developers will like the new scalar functions, especially the date and time functions that will provide performance boosts for processing.

information that the reader transaction needs from the DB2 logs—usually from the log's memory cache, which is very fast. As a result, the reader transaction can continue processing without waiting for the writer transaction to be committed. This functionality provides more robust concurrency and implements it in a way that does not cause extra overhead for the system or transaction processing.

New releases of DB2 almost always include SQL improvements, and version 9.7 has many SQL-related features that improve

functionality, compatibility, and overall performance. Your application developers will like the new scalar functions, especially the date and time functions that will provide performance boosts for processing. There are too many to mention here, but your development group should immediately evaluate the extensive number of new SQL functionalities.

Next, a number of new phrases were added to the DB2 lexicon to make it more compatible with Oracle SQL syntax and PL/SQL. The improvements should enable you to migrate Oracle applications to DB2 in days rather than months. Plus, the ability to use the same SQL syntax will allow your development team to quickly shift to the DB2 platform with its self-tuning and automated memory management.

These are only some of the DB2 9.7 features; more information can be found at ibm.com/db2/9. Many of these features were hot topics at the International DB2 User Group conference in Denver, held in May 2009. The conference included many great presentations by beta customers, consultants, and IBMers, talking through the internals as well as implementation and performance details. If you weren't able to attend, look at the presentation grid at www.idug.org/conference/CD/NA09/grid.html and buy the conference CD. ✪



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.

Insider Info on the IIUG Informix Conference

The reasons why Kansas is key to the annual gathering's ongoing success

During the 2009 IIUG Informix Conference in April, I was repeatedly asked two questions: 1) Why is the conference so successful, even in these tough economic times; and 2) Why is it in the Kansas City area two years in a row—and where will it be next year?

Okay, that's three questions, but the answers are related. The conference is so successful for one major reason: It is in Kansas City (actually Overland Park, Kansas), the hometown of the largest group of developers and support staff from the IBM Informix development center located in Lenexa, Kansas. That means the conference can draw on those folks as speakers and invite them to socialize with users and other event attendees. We have all heard the saying "direct from the horse's mouth"; well, because the conference is in Overland Park, the users and attendees can meet and talk directly with the actual Informix developers and support staff and vice versa.

See, a few years ago I was at the Lenexa development lab for an afternoon pizza party sponsored by the IIUG for the staff when they started asking me questions. The first question was "How do our users utilize Informix?" The second question was "What are our users like?" That was when I realized that bringing the developers and users together would be invaluable for both sides. We all know

that Informix users run 8 of the 10 largest retailers, 20 of the 25 top supermarkets, the 911 telephone system, and the famous soccer ball. But the IBM developers have never seen these applications, nor have they met the users who created them. The IIUG conference allows the two groups to meet and create a better, more useful, and truly user-driven product.

Let me tell you another secret about why we have gone to Overland Park two years in a row and will more than likely return in 2010: economics! It is much less expensive to hold an event in Overland Park than it is in Boston, New York, Las Vegas, or even San Jose, and it is a less expensive place for attendees to visit. Combine that with the proximity to the Lenexa development center, and the IIUG conference becomes an easy trip for your boss to approve: You get to meet the developers and support staff of the product that your company uses every day, and the trip doesn't cost too much.

So if you are asked why the IIUG conference is in Overland Park, Kansas, now you know: It is both technically and financially beneficial to everyone, and it is why—even this past year, when other user groups, events, and conferences were struggling—the IIUG 2009 conference was successful.

This year's conference attendees came from more than 30 different countries. Some

came from as far away as Cameroon and Croatia; we had attendees from Europe, Asia, and South America. Whenever I asked, "Will you come back next year?" the answer was always, "As long as the boss lets me!"

Special thanks to the entire IIUG conference team; to Dr. Arvind Krishna and Dr. Anant Jhingran, for their outstanding keynotes; and to all the sponsors for making this such a great event. And congratulations to IIUG 2009 Directors Award winners Alan Grady and Peter Perez.

If you missed the 2009 conference and would like to get a copy of many of the presentations, check out the IIUG Web site member area at www.iiug.org/members and log on to download the presentations available to any IIUG member (remember, membership is free). In the meantime, get ready to mark your calendar: we will announce next year's conference location and dates in the next few weeks, and we would love to have as many user speakers and attendees as possible.

What's new in the Informix world? Did you know that Informix Dynamic Server 11.50 now features compression capability? Or have you heard about the new Informix Warehouse product? Or the ability to check out Informix Dynamic Server on the Amazon Cloud? If you missed any of these new features, you're not reading the IIUG Newsletter at www.iiug.org/insider. *

Coming Attractions

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

IBM
data
management
KNOWLEDGE. PERFORMANCE. RESULTS.

- ▶ **Smarter, more effective strategies** for data warehousing and business intelligence
- ▶ **An in-depth look** at how data management is changing the way state and local governments operate
- ▶ **Advice and tips** from our expert columnists, including Programmers Only with Bonnie Baker
- ▶ **Best practices** for data scoring in IBM InfoSphere Balanced Warehouse environments

Subscribe
ibm.com/dmmagazine

Advertise
Randy Byers
650.919.1200 x144
advertise@tdagroup.com



The Performance and Cost Optimization Ripple Effect

What every manager needs to know to make database optimization pay off



Scott Hayes

(Scott.Hayes@dbisoftware.com) is president and CEO of DBI Software. He is an IBM Data Champion and Gold Consultant; a frequent speaker at IDUG, IBM, ISACA, and ISSA conferences; a published author; and a regular blogger on DB2 for Linux, UNIX, and Windows performance.

Your database is probably the last place you would look for IT cost savings. After all, what's most important about a database is that it delivers the correct information on time, every time. But the beauty of database optimization is that it both cuts costs and improves performance. Optimizing your organization's databases reduces unnecessary and wasteful CPU consumption, generating a ripple effect of benefits: lower CPU consumption, lower energy costs for hardware and cooling, improved server consolidation and virtualization, lower hardware costs, lower software licensing costs, and improved organizational productivity through better, more reliable, and more predictable response times.

But to get the most out of your database optimization efforts, there are two things that you need to know. The first thing is: *The physical design of your database is critically important, especially the indexes.* Indexes are important to databases for much the same reasons that they're important to people: they make it possible to find things faster with less effort.

Imagine you have a bowl of blueberries in your refrigerator and you want to use them in a recipe. You get out your favorite cookbook to find recipes that use blueberries, but it's pretty big: 592 pages. You could scan through every page of the cookbook and scan ingredients lists looking for blueberries, but scans are costly. They take time, energy, and might even make your eyes hurt. More likely, your cookbook has a good index in the

back and you can look up "blueberries," find the recipes, and go straight to those pages to pick the one you want.

Databases work the same way. Your 592-page cookbook wouldn't make for a very large database table—almost certainly small enough to fit on a 32 MB USB memory stick. But if you asked a cookbook application to search for recipes with blueberries in the ingredient list, the application would send a query to the database, and the database would scan recipe entries. The scan might take 592 CPU instructions or more if the cookbook is not already in database memory, and it might take five seconds to complete. Now, if a useful index on **INGREDIENT** is available, the same search could be completed using only six CPU instructions and would be finished in a fraction of a second.

Your business system applications probably don't do frequent searches for blueberries, but the same principles apply. Just about every business system database and data warehouse has a large number of frequently accessed small tables that would fit on USB jump drives. Ask your DBAs how many tables exist in your databases that are smaller than 64 MB—the answer may surprise you.

Small database tables are important because of the second thing you need to know: *Maximum performance optimization is possible only when you consider the entire system.*

DBAs use a utility provided by database vendors called **EXPLAIN** to review the

database's access strategy to satisfy queries. The **EXPLAIN** utility reports the anticipated computing resource cost of executing the query. DBAs normally review the anticipated costs before putting queries into production. If the anticipated cost is high, they tune the statement and the physical design of the database to reduce its anticipated cost, often by adding an index.

If the anticipated cost of a query is relatively low—say a five-second query about blueberries—a DBA will often pay no attention to it. The problem is that **EXPLAIN** provides estimates for only a single execution of a given query. If your application frequently accesses smaller tables without proper indexes, the database will load them into memory each time and scan them.

Processing waste like this is prolific in at least 9 out of 10 production databases. At a large bank I worked with recently, a multi-terabyte

OLTP banking application spent 34 percent of all its CPU time running queries against a table with only 32 rows. At a large retailer, adding a missing index to a small, 2,000-row table in an 8 TB warehouse reduced the elapsed time of a critical decision support query from three hours to two minutes. At another online retailer, 97 percent of the CPU cost on a system with eight CPUs was attributable to a single query accessing a relatively small table. As we say in Texas, "I'm going to shoot you straight"—wasteful, high-cost statements like these exist in almost every database, but they are often hidden from plain view because they access smaller tables.

Databases with self-tuning memory capabilities will detect costly queries and attempt to keep small tables in memory, mitigating I/O costs and reducing response time somewhat. But these queries are still running at a CPU cost that is 99 percent higher than it

should be. Self-tuning memory simply reacts to observed performance characteristics and does its best to compensate for physical design flaws.

In the current economic climate, the immediate need for and value of database performance cost optimization and tuning cannot be overstated. Unless your organization is flush with cash and has no regard for energy and the environment, it is past time to give your databases a proper tune-up. Ignoring this call to action makes no more sense than driving your car around with underinflated tires and a trunk full of rocks. Consider: If you could cut your number of database servers in half, your license costs in half, your energy costs in half, and simultaneously double or triple performance and improve reliability, would you? And what would you be willing to invest to obtain these savings and improvements? *

FORMERLY IBM DATABASE MAGAZINE



Visit ibm.com/dmmagazine and sign up for your complimentary digital subscription.



Join us online

- ▶ **Quick content searches** throughout the entire issue
- ▶ **Direct links** to related community resources
- ▶ **Easy information access**, sharing, and printing

You'll get the same in-depth technical content, practical advice, and hands-on tips about how to improve productivity and performance as the print edition—now including real-world commentary about how your peers are using data and information architectures to reduce costs and improve business results.



DB2

Automation, compression,
and management
improvements in the new
version of DB2 save valuable
time—and money

9.7:

It All Adds Up

By Eric Ahrendt

TIME IS MONEY. BUT IN THE DATA MANAGEMENT BUSINESS, THIS saying might be better phrased as “Time costs money...and so does everything else.” Servers, storage, power, administrative time, development time, and support time all show up on the balance sheet these days, and pretty much everyone from the data center team to the CEO is trying to figure out how to spend less.

This environment is both the backdrop and the driver for IBM DB2 9.7 for Linux, UNIX, and Windows (LUW). Announced in April, the latest version takes direct aim at reducing costs across the board with a broad range of features and improvements. Some of the headliners include new Deep Compression technology to help organizations reduce the amount of storage needed, as well as new and improved tools to streamline administration and workload management, accelerate development, and automate repetitive tasks.



Compression supports performance gains

The database might not be the first place you would look to help reduce hardware costs, but IBM has been working on that line of thinking with DB2 for some time. For example, the long-term performance advantage maintained by DB2 can ultimately translate into a reduced need for hardware purchases, because you aren't forced to constantly add hardware to keep up with performance demands.

Data compression can have the same beneficial effects on an organization's need for storage, and DB2 9.7 offers improvements to the Deep Compression technology that first appeared in DB2 9.1. Whereas version 9.5 added the capability to automatically start compressing data once there's enough data to create a meaningful dictionary, DB2 9.7 adds the ability to compress more than just row data, expanding to include indexes, log files, temporary tables, inline XML data, and large objects. These enhancements will be of particular interest to organizations with large business intelligence (BI) applications and data warehouses, which depend on indexes and temporary tables.



Unlocking XML data

By enabling DB2 to store XML data in its native format, the pureXML engine has given organizations that depend on XML tremendous advantages in performance and flexibility. A major enhancement in DB2 9.7 is the ease of moving back and forth between SQL data and XML data—developers don't need to know which is which and can develop with whatever tools they prefer.

More organizations are capturing XML data but aren't sure what to do with it. For example, suppliers may communicate via XML with an organization's EDI system, so that a single system can handle the interaction with all suppliers. Until now, the organization's transactional system has had to convert XML data to relational data or simply delete it from the warehouse. But DB2 9.7 can store XML in a warehouse and can scale to accommodate high volumes of data—even a terabyte or more. By taking advantage of this capability, organizations can now use their BI tools directly against XML to discover business insight previously locked in their XML data. The enhancements to Deep Compression for XML have led to great results, with compression rates over 65 percent and performance acceleration of more than 1.5 times.



The additional compression also leads to better performance with indexes and temp tables. Companies involved in the beta testing program for the most recent version of DB2 are consistently reporting compression rates of 70 percent or more, which translates into storage savings of up to 50 percent with no performance penalties.

DB2 9.7 compresses more than just row data; it can also compress indexes, log files, temporary tables, inline XML data, and large objects.

Virtualization on the march

Most databases continue to command dedicated physical servers, but virtualization is on the rise—so IBM has announced support for virtualization in all editions of DB2 9.7 for LUW, from IBM DB2 Express to IBM DB2 Enterprise, as well as the IBM InfoSphere Warehouse for DB2 editions. IBM supports a broad array of virtualization environments for x86 and x64 architectures, including both full virtualization (VMware ESX, Red Hat Enterprise Linux, and SUSE Linux Enterprise Server), and OS virtualization (Solaris Zones, Parallels Virtuozzo Containers). For a complete list of virtualization support, visit ibm.com/developerworks/wikis/display/im/DB2+Virtualization+Support.

IBM has also optimized DB2 9.7 for VMware. As workloads or virtual partitions change, DB2 will react and dynamically allocate memory. You can also take advantage of the VMware vMotion feature to move a virtual machine from one physical server to another with no downtime. IBM is certified by VMware on this capability.

To help lower costs, IBM offers flexible virtualization license costs, where you pay only for the number of virtual resources that you have deployed. Processor Value Unit (PVU) sub-capacity licensing lets you license DB2 for less than the full capacity of your server or group of servers. It provides the licensing granularity needed to leverage various multi-core chip and virtualization technologies. Also, IBM has provided additional usage options by announcing DB2 availability through the Amazon Elastic Compute Cloud (EC2) infrastructure service (aws.amazon.com/ibm).



DB2 stays solid on security

Making sure that the right people—and only the right people—have access to specific data is a critical challenge for data management professionals. IBM tackles this issue in the new version of DB2 with enhanced security and audit features.

DB2 9.7 for LUW increases access control granularity, making it possible to give DBAs full control over a database without access to the actual data. This enables DBAs to do everything to the database that they would

ordinarily do—perform regular administration tasks, load data, use utilities—but not see or alter the data. Separating the ability to administer the database from the ability to access the data it carries makes it easier for organizations to establish and maintain security policies without interfering with necessary maintenance and administration functions.

Auditing was also recently tweaked for DB2. The audit facility was redesigned to improve performance and to provide

fine-grained auditing (FGA). The audit facility now produces all audit records based on the audit configuration specified, controlling critical information about who is accessing DB2, when, and from where. The audit facility still provides the ability to audit at both the instance and the individual database level, independently recording all instance- and database-level activities in separate logs. The improvements help data managers track connections and authorizations,

statement text, application IDs, and the originating request's IP, along with timestamps for important events.

Rounding out the security improvements in DB2 9.7 is Encryption Expert, which allows users to encrypt data at rest, onsite, and offsite, and to store passwords at a central security server. In-transit data is also secured by passwords so that even if a tape is lost during delivery, the data on it is not accessible to unauthorized users.

How to control time

Saving DBA time and effort is another big part of the DB2 9.7 feature set, which expands on the familiar DB2 controls and automated tools for configuring, optimizing, and protecting the database. One of the most important new administrative features expands on the workload management capabilities of DB2, which gives database administrators the ability to prioritize workloads and be sure that their databases are processing the right job at the right time.

In DB2 9.5, users gained the ability to prioritize workloads and assign rankings to different users, roles, groups, application names—any combination of factors. In DB2 9.7, this capability is enhanced with time-based functionality. Integrated tooling from IBM Data Studio Base, in the IBM Optim Performance Manager for DB2 for Linux, UNIX, and Windows tool (previously IBM Data Studio and IBM DB2 Performance Expert; see sidebar, “Optim development environment grows”), allows users to automatically boost job priorities at specific times of day to meet deadlines. Users can also assign a high priority to batch jobs during off-peak hours so they can be completed within designated windows.

A waterfall workload management ability also enables DBAs to automatically lower the importance of some workloads when certain thresholds, such as CPU used or rows read, are hit. This helps ensure that rogue queries do not

Optim development environment grows



IBM continues to build on its vision of Integrated Data Management in parallel with the development of DB2 9.7. Key product announcements to support the strategy were first made in October 2008. In June 2009, IBM is announcing extensions to its Optim product line in the areas of query and query workload performance tuning; a common environment for developing applications with DB2, IBM Informix, and Oracle database management systems; and enhanced integration of products across the Optim, IBM InfoSphere, and IBM Data Studio portfolios.

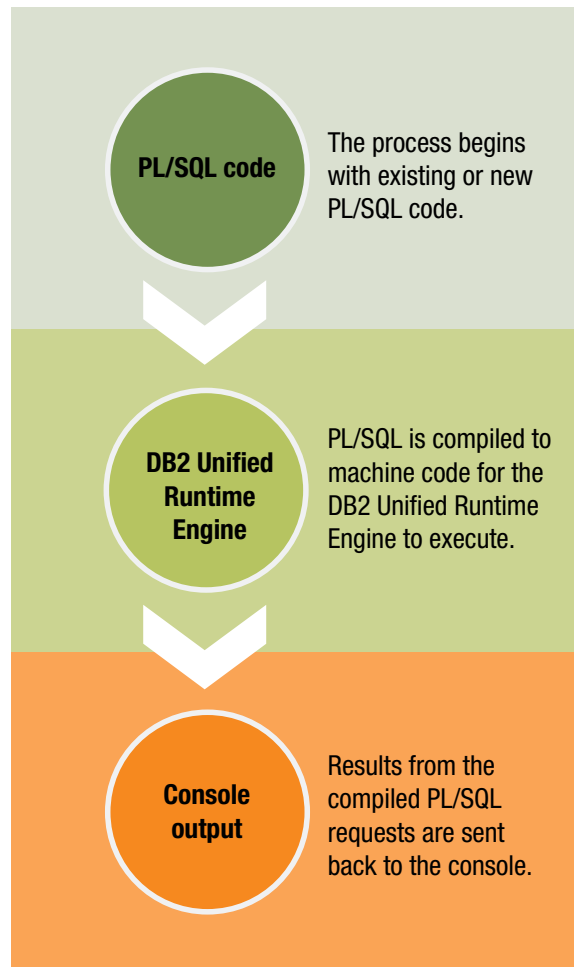
IBM is introducing Optim Development Studio, Optim pureQuery Runtime, and Optim Database Administrator (formerly Data Studio Developer, Data Studio pureQuery Runtime, and Data Studio Administrator). These products have been enhanced and renamed under the Optim umbrella. In addition, IBM is announcing the new Optim Query Tuner product.

SQL compatibility features make DB2 9.7 for Linux, UNIX, and Windows easier for application developers to adopt.

take over a database. DBAs can use the same technology to increase the priority of certain workloads at key thresholds, such as time, to make sure important queries that must meet service-level agreements get resource priority.

The new version of DB2 9.7 for LUW also makes it easy to manage mixed workloads on your system, such as a high-volume transactional system with reporting or a BI system with occasional trickle feeds. A newly added “currently committed” locking syntax eliminates the conflict between

Figure 1: The PL/SQL compilation process with IBM DB2 9.7 for Linux, UNIX, and Windows has three basic steps.



read and write workloads so the database can deliver a true point-in-time response.

In addition to new features for managing database time, DB2 9.7 also helps administrators minimize downtime by making it possible to move database schemas without taking the system down. Changing tables—even changing column type—can be done with no downtime, as can moving a table from one table space to another. In fact, you can transfer an entire schema, including all the tables and associated objects, from development to QA to performance-testing and incur no downtime.

DB2 9.7: Simplifying life for both DBAs and developers

With advances in compression, development, administration, virtualization support, and development, DB2 9.7 is poised to help database managers use server and storage resources more efficiently, automate workload management, and simplify application development. Application developers and ISVs familiar with the Oracle DBMS can more rapidly get applications and tools running on DB2 with new PL/SQL support and flexible concurrency model and data-typing technology (see Figure 1). The new DB2 9.7 release also offers organizations more options in security management and high availability.

With businesses keeping a close eye on both time and money these days, the new features in IBM DB2 9.7 for LUW will help get the most out of both. *

Eric Ahrendt writes on technology issues for a range of Fortune 500 companies.

RESOURCES

DB2 9.7 for Linux, UNIX, and Windows:

ibm.com/breakfree

ibm.com/db2/9

IBM Optim, IBM Data Studio, and IBM Integrated Data Management:

ibm.com/software/data/data-management

Industry analyst thoughts on DB2 9.7:

<http://mervadrian.wordpress.com/2009/04/22/db2-runs-plsql-say-what>

<http://mervadrian.wordpress.com/2009/04/23/db2-97-focuses-on-costs-simpler-management>

www.gartner.com/DisplayDocument?doc_cd=167786&ref=g_homelink

SQL

By Anil Mahadev

Compatibility in DB2 9.7

From triggers to object enhancements: A demonstration of the new Oracle compatibility features

CALLING ALL DEVELOPERS! DO YOU HAVE Oracle database skills? Do you wish you could reuse your PL/SQL knowledge and code on another database platform?

IBM DB2 for Linux, UNIX, and Windows (LUW) has had a makeover. Driven by present and future customer requirements for cost-saving enhancements, the latest release, DB2 9.7, provides a host of new technologies in familiar areas like autonomics (self-management) and Deep Compression (see “DB2 9.7: It All Adds Up” in this issue). But the icing on the cake is a collection of capabilities focusing on SQL compatibility. These features make DB2 easier to adopt as your database software so you can take advantage of its cost-saving and high-performance characteristics, even if you are currently most familiar with competing databases.

In this article I will show examples of several of the new SQL compatibility features in DB2 9.7. We’ll start with triggers and new data types, move into support for Oracle SQL and PL/SQL, and wrap up with a discussion of object and concurrency enhancements. There is something for everyone as DB2 “gets compatible.”

To get started, connect to the Command Line Processor Plus (CLP Plus) window (see Figure 1), replacing “db2admin” with your chosen login credentials:

```
clpplus db2admin@localhost:50000/sample
```

SQL compatibility

The first step in using the SQL compatibility features in DB2 9.7 is to enable the `DB2_COMPATIBILITY_VECTOR` parameter. Setting `db2set DB2_COMPATIBILITY_VECTOR=ORA` enables the SQL compatibility features related to Oracle for all subsequent database creation. For an explanation of some of the most important parameters used by `DB2_COMPATIBILITY_VECTOR`, see Table 1.

To enable this feature, you must have administrative privileges. Enter the following commands at the DB2 command window:

```
db2set DB2_COMPATIBILITY_VECTOR=ORA
db2stop
db2start
```

To create a simple database, enter: `db2 create database [dbname]`.

You can then connect to the database by entering:

```
db2 connect to [dbname];
```

The next step is to create tables using Oracle data types. Enter the following `CREATE TABLE` statement using either the CLP Plus window or IBM Data Studio:

```
Create table employee (EMPNO NUMBER(5), ENAME VARCHAR2(50),
DEPTNO NUMBER(5))
/
```

```
Administrator: CLP Plus
CLPPlus: Version 1.0
Copyright (c) 2008, IBM CORPORATION. All rights reserved.
SQL> help index
Type 'HELP [topic]' for command line help.
DLL          CHANGL      IMPORT      EXECUTE
ACCEPT       EXIT        DESCRIBE   HOST
BTITLE       SET         DISCONNECT SPOOL
GLT_DB_CFG   GET_DBM_CFG CURRENT_SCHEMA CLLAR_BREAKS
UNDEFINE     REMARK      UPDATE_DBM_CFG  QUIT
START        BREAK       @           UPDATE_DB_CFG
LIST        TABLE     SAVED      LIST
CLEAR_COMPUTES SERVEROUTPUT INDEX       GET
COMPUTE      CLEAR      HELP       RESET_DBM_CFG
DEFINE       APPEND     CONNECT    RESET_DB_CFG
SHOW        PROMPT    COLUMN     CLPPLUS
PRINT       JCCLOGMODE LOAD        PASSWORD
LIST_TABLES LOGMODE    PAUSE      VARTABLE
INPUT       EXPORT
```

Figure 1: In DB2 9.7, the CLP Plus interface allows you to try out the SQL compatibility features.

Bit position	Feature	Explanation
(0x01)	ROWNUM	Enables the use of ROWNUM as a synonym for ROW_NUMBER() OVER() and permits ROWNUM to appear in the WHERE clause of SQL statements.
(0x02)	DUAL	Enables the DUAL dummy table.
(0x04)	Outer join operator	Enables support for the outer join operator (+).
(0x08)	Hierarchical queries	Enables support for hierarchical queries using the CONNECT BY clause.
(0x10)	NUMBER data type	Enables the NUMBER data type and associated numeric processing.
(0x20)	VARCHAR2 data type	Enables the VARCHAR2 data type and associated character string processing.
(0x40)	DATE data type	A combined date and time value.
(0x80)	TRUNCATE TABLE	Enables alternate semantics for the TRUNCATE statement under which IMMEDIATE is an optional keyword that is assumed to be the default if not specified. An implicit commit operation is performed before the TRUNCATE statement executes if the TRUNCATE statement is not the first statement in the logical unit of work.
(0x100)	Character literals	Enables the assignment of the CHAR or GRAPHIC data type (instead of the VARCHAR or VARGRAPHIC data type) to character and graphic string constants whose byte length is less than or equal to 254.
(0x200)	Collection methods	Enables the use of methods to perform operations on arrays, such as first, last, next, and previous. Also enables the use of parentheses in place of square brackets in references to specific elements in an array; for example, array1(i) refers to element i of array1.
(0x400)	Data dictionary-compatible views	Enables the creation of data dictionary-compatible views.
(0x800)	PL/SQL compilation	Enables the compilation and execution of PL/SQL statements and language elements.

PL/SQL features

We will now explore the features that allow you to create PL/SQL procedures. DB2 9.7 supports the following:

- ▶ User-defined packages
- ▶ PL/SQL procedures/functions/anonymous blocks
- ▶ Built-in packages (**DBMS_OUT.PUT_LINE**, **DBMS_PIPE**, **DBMS_UTL**, and so on)
- ▶ Associative (**INDEX BY**) arrays
- ▶ Varrays
- ▶ Triggers
- ▶ **%ROWTYPE**
- ▶ **%TYPE**
- ▶ **EXCEPTIONS**
- ▶ Ref-Cursors

Note that PL/SQL support is available only in the Enterprise Server Edition and Workgroup Edition of DB2 9.7 for LUW.

PL/SQL is not supported in DB2 Express-C, DB2 Express, or DB2 Personal Edition.

PL/SQL procedures in DB2

DB2 supports much of the PL/SQL language. Let's see how this works with some sample code. Once it is compiled and executed, this PL/SQL procedure displays a welcome message to the user. To see the output on the console, first issue the command **SET SERVEROUTPUT ON**.

```
CREATE OR REPLACE PROCEDURE message_proc (myname varchar2)
IS
BEGIN
    DBMS_OUTPUT.PUT_LINE('Welcome to DB2 9.7 (! ' || myname
    || ' . Now supports for ORACLE SQL and PL/SQL as well :-)!');
END message_proc;
```

To call this PL/SQL procedure in DB2, issue the following commands:

```
EXEC proc_name
EXEC message_proc('Anil');
```

Here is an example that uses a **FOR** loop:

```
BEGIN
    FOR a IN 10 .. 20 LOOP
        DBMS_OUTPUT.PUT_LINE('Counter ' || a);
    END LOOP;
END;
```

This loop will produce the following output:

```
Counter 10
Counter 20
```

The TRUNCATE TABLE command

Those of you with an Oracle database background have probably used the command **TRUNCATE TABLE** at some point in your careers. DB2 now supports this command as well. The following code creates a simple table, adds rows to it, and then uses the **TRUNCATE** command to remove all the rows from the table:

```
CREATE TABLE CLIENTS
(CLIENTID NUMBER(5) PRIMARY KEY NOT NULL, CLIENTNAME
VARCHAR2(50));
```

Insert some rows:

```
INSERT INTO CLIENTS (CLIENTID, CLIENTNAME) VALUES (1,'IBM');
INSERT INTO CLIENTS VALUES (2,'MSFT');
INSERT INTO CLIENTS VALUES (3,'EDB');
```

Let's look at the rows:

```
SELECT * FROM CLIENTS;
```

Now issue the following command to truncate these rows:

```
TRUNCATE TABLE CLIENTS;
```

Let's prove that they're gone:

```
SELECT * FROM CLIENTS;
```

You will get zero rows displayed.

Object and concurrency enhancements

I'd like to wrap up by mentioning two important improvements to DB2 9.7 that can save you time and increase performance: first-use validation, and a new default locking state.

First-use validation

Whenever underlying database objects (tables, views, procedures, functions, and so on) are altered, the next time an object is retrieved with its dependent objects, it will be automatically revalidated. This is definitely a time-saver for application developers. The automatic dependency check is done whenever an object is altered.

Read currently committed (default in DB2)

In previous releases of DB2, it was not possible to have one user trying to read and one trying to write to the same row at the same time. This is based on the principle that readers

wait for the outcome of pending writes. DB2 9.7 now uses the default-locking state so that readers will not wait for the outcome of a pending write. Instead they will read the currently committed version of the row. Hence, readers don't block writers and writers don't block readers.

I am very impressed by these new SQL compatibility features. I have only been able to scratch the surface in this article, So I encourage you to take DB2 9.7 for a spin. You won't be disappointed! *

Anil Mahadev is an IBM Data Champion and a member of the Board of Directors for the IDUG India Forum. He is currently working as an independent consultant in Bengaluru, India, and exploring database implementation, Windows clustering, high availability, and strategies for multi-database platforms and applications.

RESOURCES

General information on IBM DB2: ibm.com/db2/9

DB2 product documentation: ibm.com/support/docview.wss?rs=71&uid=swg27009474

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

Upcoming DB2 technologies:
ibm.com/db2/technology-sandbox

Dive Deeper

With *IBM Data Management* magazine, you get more than a surface glance: more detail, more information, more context, and more viewpoints on the technologies and strategies you deal with every day.

IBM
data
management
KNOWLEDGE. PERFORMANCE. RESULTS.

Dive deeper.
Get to the bottom of things at:

IBM.COM/DMMAGAZINE

▶ IBM Data Management digital edition


Sign up for easy online access and quick content searches across the entire magazine.

▶ IBM Data Management community

Explore current and previous issues, and post your views to the magazine's blog.

▶ My developerWorks and developerWorks DBA Central

Customize developerWorks to deliver your information, your way—news, blogs, downloads, feeds, forums, and more.



IBM Optim Query Tuner for Linux, UNIX, and Windows helps developers create efficient queries and build tuning skills

By Cliff Leung,
Tony Leung,
Ray Willoughby, and
Kathryn Zeidenstein

SQL TUNING:

Not *just*
for **hardcore**
DBAs *anymore*

AS GOOD AS IBM DB2 IS AT OPTIMIZING QUERY PERFORMANCE, TUNING SQL STATEMENTS sometimes requires skilled human intervention. There are those who have built careers around their deep knowledge of how to tune SQL to achieve optimized performance. And no wonder: a poorly written SQL statement can wreak havoc, impacting service-level agreements, consuming precious resources, or even causing application outages—all of which can cost businesses time and money.

Recently, the need for highly optimized SQL crashed headlong into new advances in application programming. Frameworks that automatically handle database operations, such as Hibernate and Java Persistence Architecture (JPA), are becoming more popular with Java developers. Development frameworks can make developers more productive, but these frameworks and the DBAs who use them are unable to see or control the generated SQL. In fact, the frameworks may generate SQL “under the covers” that would make experienced DBAs tear out their hair.

Add to this the fact that these SQL statements are dynamic, and it becomes even more difficult to track down the originating code that produced the SQL. The result: SQL performance issues that may not be detected until the application is under a full load in the production environment.

What if it were easier for developers to do some SQL tuning during development? Could tuning become a routine task that every developer does easily and effectively before deploying the application? With IBM Optim Development Studio (formerly Data Studio Developer), the answer can be “yes.”

With Optim Development Studio, DBAs and developers now have the ability to *see and manipulate* the actual SQL that is being used by the Java program in the context of the actual line of source code that produced the SQL. In addition, Optim Development Studio provides visualization capabilities that allow developers to see which SQL statements cost the most and, maybe even more important, how often they are run. Finally, the new query tuning offerings provide tuning advice designed to help developers and DBAs collaborate effectively to produce fast, less costly enterprise-ready database applications.

Introducing IBM Optim Query Tuner

It’s understandable why developers don’t make query tuning a high priority. First of all, developers tend to focus on the results a query returns, asking only, “Am I getting the data I need?” Second, developers are often under tight deadlines and don’t have time to truly examine how queries perform—and they may not be able to unit test under a significant load to uncover problems. Finally, understanding how the `EXPLAIN` tables relate and how to interpret the information contained in them can be intimidating and cryptic.

Even if a developer spots a SQL issue, there is still a gap between knowing you have a problem and knowing what to do about it. IBM DB2 Visual Explain can provide a picture of what the database is doing, but that picture can be pretty mysterious to a developer, or even to a new DBA. Plus, DB2 Visual Explain only describes what the database is doing; it

doesn’t provide insight into the performance ramifications of that choice. Experienced DBAs who may be able to solve the problem are often so busy with other high-value activities that they may not be able to help until the problem lands in their lap on the production system.

Optim Query Tuner can help both novice and expert database developers write better-performing SQL. Its easy-to-use advisors and query visualizations can help cut costs and improve performance by providing expert advice on writing high-quality queries and improving database design.

As the name suggests, Optim Query Tuner provides single-query tuning assistance. You can invoke Optim Query Tuner for SQL statements from:

- ▶ Packages, SQL stored procedures, triggers, and user-defined functions (UDFs)—all from within the Data Source Explorer in the Eclipse environment
- ▶ The integrated query editor and routine editor within IBM Data Studio or other Optim development offerings
- ▶ The SQL Outline and the Java editor within Optim Development Studio
- ▶ The DB2 catalog, for database packages and stored procedures

This single-query tuning capability is announced for IBM Optim Query Tuner 2.2. You can get combined single query and workload tuning today for IBM DB2 for z/OS with DB2 Optimization Expert for z/OS (which will soon be renamed to IBM Optim Query Workload Tuner for DB2 for z/OS).

Visualization and advisor options

Once you select a statement for tuning, there are several visualizations and advisors available to help you.

Developers with less SQL experience may rely heavily on the *Query Advisor*. This advisor makes recommendations on how to rewrite a query to enhance efficiency based on best-practice rules. When you select a recommendation, the line of SQL is highlighted, and you are presented with a description of how to rewrite the query, as well as an explanation of the recommendation (see Figure 1).

Note that recommendations for changing queries are based on the assumption that the DB2 catalog statistics are accurate. With *Query Annotation*, you can format the SQL statement for improved readability and navigation and see which statistics are being used at the table and column level. If there are no statistics, then you’ll need to update the DB2 catalog statistics (the Query Tuner Statistics Advisor can help you do this).

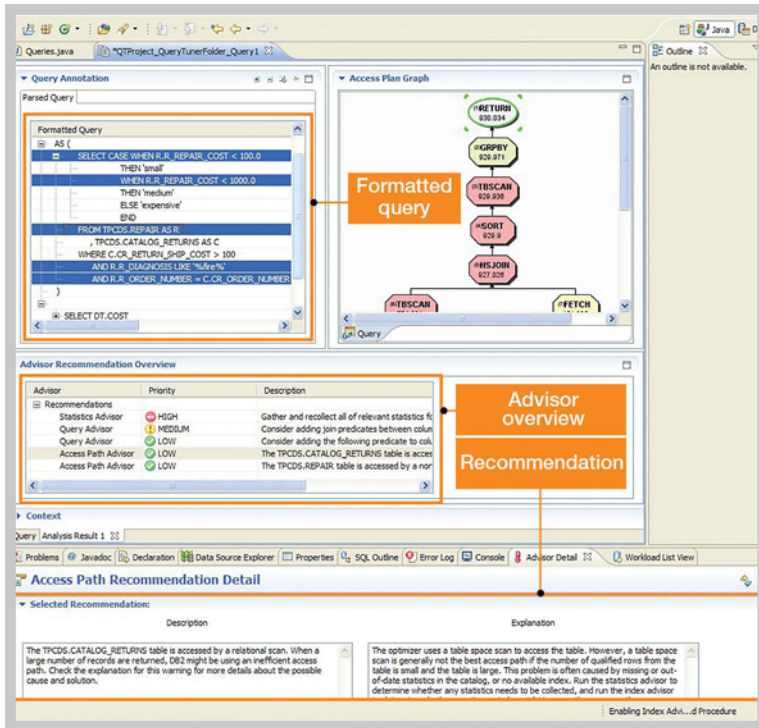


Figure 1: Optim Query Tuner includes a query formatter and advisors.

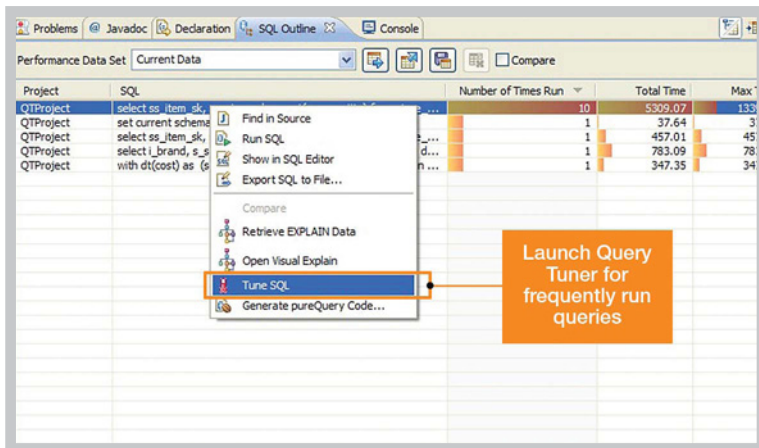


Figure 2: The SQL outline view allows you to tune queries based on frequency or elapsed time.

With the formatting capability provided by Query Annotation, you can expand and collapse sections of the SQL query. If you click on a table in the FROM clause, it will highlight related columns in the SELECT and WHERE clauses. Conversely, when you select a column in the SELECT or WHERE clause, it will highlight the table the column is related to. Additionally, if your statement references a view, Query Annotation provides the ability to drill down into the SQL query in the view.

Those with more experience will find the Access Plan Graph and Access Path Advisor helpful. The Access Plan

Graph is the same graph as DB2 Visual Explain but provides additional statistical or costing information. The Access Path Advisor identifies potential access path issues and provides an explanation and recommendations for dealing with those issues.

If you are a DBA, or if you are the virtual DBA for your own development database, you will also be interested in the Statistics Advisor and Index Advisor. As mentioned earlier, the DB2 optimizer is only as good as the catalog statistics are current. The Statistics Advisor will let you know when it's time to update those catalog statistics and generate the command to do so—all you have to do is click a button. Even if you are running auto-stats, the Statistics Advisor can provide additional information about column distribution.

The Index Advisor will analyze the query and let you know if the use of indexes on the columns used in the query may help improve the performance of the query. And yes, it will generate the index Data Definition Language (DDL) for you, so all you have to do is click a button to create the index.

When using the Index Advisor within Optim Query Tuner, remember that the advice is only relevant to the single SQL statement at hand; be sure to consider how this index will perform within the overall workload. The IBM DB2 Design Advisor can provide that analysis.

Using Optim Query Tuner with Optim Development Studio

Because Query Tuner can be installed into the same Eclipse instance as Optim Development Studio and there are integration points between these two products, Java developers now have natural points in their work where they can invoke the query tuning capability and get advice on writing better queries without slowing down the development process. There is even a way for DBAs or developers who don't have access to source code to use Optim Development Studio to replace poorly performing queries with ones that have been tuned using Optim Query Tuner. The following three scenarios illustrate these points.

Identifying and tuning performance hotspots during Java development

Optim Development Studio lets you capture SQL statements and related performance metrics for any new or existing Java application, including those created using a framework such as Hibernate (see "Optimize your existing JDBC applications using pureQuery" in the "Resources" sidebar for more detail). These captured statements and metrics are displayed in the Optim Development Studio "SQL outline" view (previously known as the "pureQuery outline").

As seen in Figure 2, the SQL outline view shows you how frequently SQL statements are run and how long they run. It provides a visual display of elapsed time for statements, including total, minimum, maximum, and average time. In addition, it tells you how many times each statement runs. You can easily identify the SQL hot spots and launch Optim Query Tuner directly from this view.

From the SQL outline view, you can also generate EXPLAIN information, then sort by cost to quickly find the most costly queries. Simply right-click on the query and launch Optim Query Tuner to tune that statement. With the combination of cost, elapsed time, and number of times executed, you can more easily decide which queries to focus your tuning efforts on.

Optimizing queries during new Java development

Let's assume you're developing a new Java application using the Optim Development Studio Java editor. Not only do you get SQL assistance to help you formulate queries correctly, you can now right-click on any SQL statement in the Java editor to be launched directly into Optim Query Tuner (see Figure 3). You will see your query formatted and annotated, as well as the access plan and advice on how to fix potential problems.

Optimizing queries for existing Java database applications

This scenario is for DBAs or developers who are working without the ability to change the underlying source code or who need to make a quick fix for a critical performance problem before a source code change can be tested and rolled out. In this scenario, you run the application to capture the SQL statements, invoke Optim Query Tuner directly from the captured SQL file, and use the Optim Development Studio capture file editor to *replace* the poorly performing statement with one that you have tuned—all without changing source code. When you set a driver property for the application, it will pick up that improved SQL statement the next time it runs. For more information about SQL replacement, see the "What's new and exciting in IBM Data Studio Developer 2.1" article in the "Resources" sidebar.

Improving the development process

We hope that this article has given you some things to think about in terms of roles and responsibilities, skill building, and overall development process improvements. If you are a DBA who manages DB2 for Linux, UNIX, and Windows systems, you may also be interested to know that IBM is

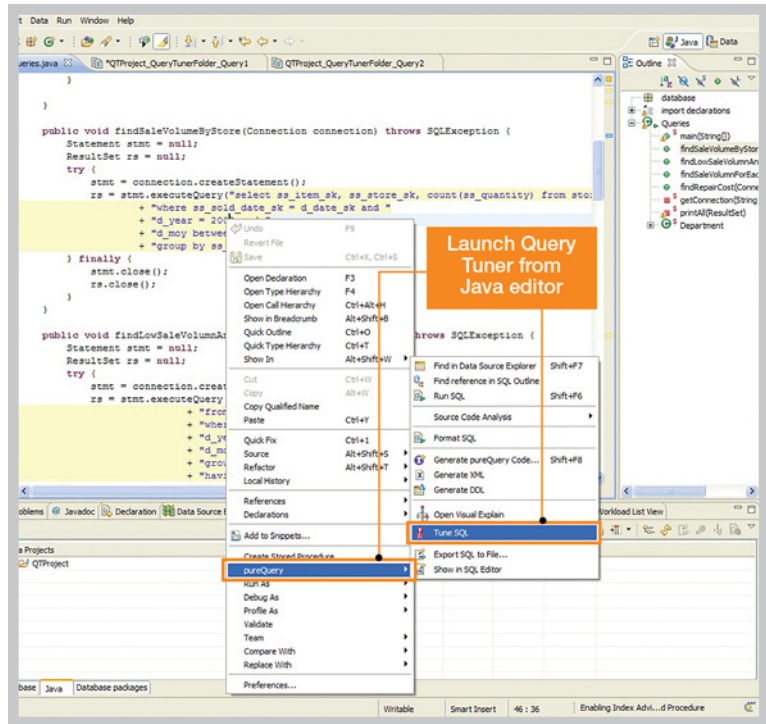


Figure 3: You can tune queries directly from the Optim Development Studio Java editor by right-clicking on any SQL statement.

actively working on another capability for these systems that exists today for DB2 for z/OS: query workload tuning and advisory capabilities to help identify, isolate, and tune active queries in a production workload. ✱

The authors are part of the IBM Integrated Data Management development team based in the Silicon Valley lab in San Jose, California. Cliff Leung and Tony Leung are software architects. Ray Willoughby is the technical product manager for query tuning offerings. Kathryn Zeidenstein is a member of the technical enablement team, focused on community efforts and communication.

RESOURCES

Optim Query Tuner:

ibm.com/software/data/optim/query-tuner

Optim Integrated Data Management portfolio:

ibm.com/software/data/optim

Optim Development Studio:

ibm.com/software/data/optim/development-studio

Optimize your existing JDBC applications using pureQuery:

ibm.com/developerworks/edu/dm-dw-dm-0808titzler-i.html

What's new and exciting in IBM Data Studio Developer 2.1:

ibm.com/developerworks/data/library/techarticle/dm-0812surange/index.html



Healthcare's XML Heartbeat

By Ken North

DB2 pureXML anchors the patient-oriented document system of a multi-hospital healthcare provider

F

OR A DECADE, ACADEMIA, GOVERNMENT, AND INDUSTRY groups have used XML for document storage and data integration. Database and software providers, including IBM, recognized the potential of XML early on: they helped define standards and industry-specific markup languages based on XML, while working to integrate XML with SQL in their database platforms. By first releasing the IBM DB2 XML Extender, Informix Web DataBlade, and Informix XSLT DataBlade, then more deeply integrating XML support into the RDBMS servers, IBM and Informix advanced the integration of XML and SQL database processing.

XML is an important enabling technology for multiple industries, particularly industries with complex integration, archiving, and compliance requirements. In particular, the healthcare industry can exploit service-oriented architectures (SOAs) and SQL/XML-powered databases for building new applications while addressing compliance and standards issues—and for supporting electronic medical records systems.

The rise of XML in healthcare

The rise of XML in the healthcare industry has been driven partly by legislation intended to protect patients' security and privacy, including the Health Insurance Portability and Accountability Act (HIPAA). Enacted by the U.S. Congress

to protect insurance coverage, HIPAA includes standards for electronic transactions and provisions for the privacy and security of data and applies to claim, payment, benefit inquiry, claim status, and other transactions. HIPAA also requires the U.S. Department of Health and Human Services to define rules for the dissemination of healthcare information.

Translating those requirements into usable standards is often the work of standards-development organizations; one of the most prominent in the healthcare industry is Health Level Seven (HL7). HL7 produces standards for operations involving the exchange of administrative and clinical data in healthcare domains, including claims processing, imaging, and pharmacy (see sidebar, "HIPAA and HL7"). The healthcare industry has also developed specifications for an operational data model and study data tabulation model (Clinical Data Interchange Standards Consortium). These standards apply to communication between internal systems and external entities, such as the U.S. Food and Drug Administration.

New standards in various industries, including healthcare, have prompted development of XML-enabled applications. This new wave of technology means that we can build applications—often composite applications using SOA plumbing—to access medical data with a combination of interoperable services and rich database support for XML.



HIPAA and HL7

But efficiently managing large amounts of XML data can be a challenge. With DB2 9, IBM introduced the pureXML solution, which permits storage, indexing, and querying of documents in their native XML format. Several leading institutions have taken advantage of the native XML capabilities of IBM DB2 to build systems that not only exploit healthcare industry standards, but also improve data access and performance.

The challenge of electronic medical records

With sizeable amounts of rapidly and frequently changing data, large healthcare organizations rely heavily on XML. The UCLA Health System is one such organization: a multi-hospital healthcare provider with a diversity of clinical and healthcare applications. It includes the Ronald Reagan UCLA Medical Center, Santa Monica UCLA Medical Center and Orthopaedic Hospital, Mattel Children's Hospital UCLA, and Resnick Neuropsychiatric Hospital at UCLA, as well as the UCLA Medical Group of primary-care and specialty-care offices. The staff of more than 2,000 physicians handles more than 1 million clinic visits and 80,000 hospital visits per year.

A healthcare system of that size must process large amounts of data on a daily basis, including medical records updates; lab results; MRI, computed tomography (CT),

HL7 has been developing standards designed to address HIPAA legislation since 1996, when it formed a Claims Attachment working group to standardize the information needed to process insurance claims. In that same year, HL7 began actively working with XML through its SGML/XML special interest group. The initial deliverable was six recommended attachments for claims processing.

Since then, HL7 has developed a specification for messaging, a Clinical Document Architecture (CDA), and a Reference Information Model (RIM). HL7 has also worked on a standard for electronic submission for CDA Public Health Case Reporting (PHCR) to state and local public health departments.

As the standards matured, XML became an increasingly important technical component. For example, the first version of the CDA defined an XML architecture for exchange of clinical documents based on XML Document Type Definitions (DTDs) included in the specification with semantics defined using the HL7 RIM and HL7 registered coded vocabularies. The upcoming version 3 release of the CDA is expected to use only XML encoding.

"IBM DB2 native support of XML allows it to store content in the healthcare industry-standard HL7 CDA format," says Karla Norsworthy, vice president, software standards at IBM. "IBM is committed to healthcare interoperability and innovation through open standards. We have seen the benefits of flexibility, time to market, and innovation that come from widely adopted, open standards such as Java, XML, and healthcare standards developed in organizations such as HL7."



and electron beam CT angiography images; admission, discharge, and transfer data; and pharmacy orders. All of this data must be stored securely and reliably (and eventually archived intelligently) while remaining accessible on demand. In addition, the data must be easily searched, transmitted, and organized by a wide variety of employees spread across multiple locations. Add the need to easily and quickly input and update information, and you're looking at a complex piece of IT infrastructure.

The UCLA Health System met those demands with its patient-oriented documentation system (PODS), an electronic medical records (EMR) repository that provides storage and retrieval capabilities for more than 20 million documents. In the bigger picture, PODS is a critical source of patient information for the UCLA document management system extended SOA (xSOA). Along with PODS, the xSOA provides viewer interfaces to a GE BDM pharmacy information system, a CliniComp Essentris acute care system, an Orion Soprano clinical data system, and a forms portal. The xSOA Central Document Bus connects to a GE picture archiving and communication system (PACS), clinical applications and images, and the PODS repository. The Image Bus provides access to patient diagnostic images, whereas a Forms Bus handles the large variety of electronic forms used by UCLA Health Services. An HL7 Message Bus provides HL7-compliant communications.

Combining an SOA with DB2 databases, PODS provides 2,000 doctors and 3,000 nurses with access to patient records. The system supports more than 400 electronic forms for data entry; these forms replace the 1,000 paper forms used previously, helping to eliminate errors due to handwriting misinterpretation or blank fields. The database holds information for about 2 million patients and grows by 12,000 documents per day as new test results, doctor's notes, and other patient-related data are added.

PODS includes a document repository and metadata repository, pairing medical record files with a DB2 database. PODS stores medical record image files on file servers; it uses an IBM DB2 database to store corresponding metadata and a network attached storage array for the image files, including PDFs and text.

When a document enters the system, it is stored in the file server and indexed in the DB2 database. Documents are not deleted because the PODS repository also serves as an archive of patient data. To ensure survivability and high availability for 24x7 operations, the PODS architecture includes redundant servers and databases, with data replication to synchronize between database servers.

The DB2 database metadata store currently contains 30 million rows of information. The metadata is stored as XML, using the DB2 9 native XML engine. According to Dr. Charles Wang, architect manager at UCLA Medical Center Computing Services, the more than 400 schemas used for PODS comply with the W3C XML Schema language. The PODS software maps those different schemas into a single, virtual schema for the entire system. The system uses a four-key composition to create a unique identifier for a document paired with its metadata.

To safeguard patient privacy, the PODS design uses a multilevel security model. Besides built-in DB2 security capabilities, the software architecture includes a document and metadata handler that is integrated with a security service. It also features role-based security and a single sign-on capability. Security, concurrency control, parallel processing, and versioning are but a few of the advantages of using DB2 for storage of XML schemas and documents, as opposed to dealing with those problems on an ad hoc basis when using a file system to manage XML documents and schemas.

UCLA Health System builds on SOA

The PODS implementation is a good example of how an SOA enables disparate applications to use essential services—in this case, services for accessing patient information. Documents enter the system via a document service interface and are placed in queues. The PODS architecture

includes IBM WebSphere MQ for asynchronous messaging and queuing. For HL7 messaging, the UCLA Health System uses Sun SeeBeyond eGate Integrator which, according to Dr. Wang, is “the enterprise-wide standard for all application interfaces.”

The PODS queue manager operates with an input queue, exception queue, and replication queue. The standard services for managing metadata and image files provide application programming interfaces (APIs) for uploading, downloading, querying, and updating documents. The system supports auditing by generating a report of all activity against the database except uploads.

The evolution of PODS

The UCLA Health System PODS implementation supports access to patient documentation with DB2 pureXML capabilities for loading, querying, and updating data. It provides a set of Web services interfaces that enable clinical systems to upload and query data.

The most recent versions of the system are PODS3 and PODS4. Both were built on an SOA, but use different DB2 capabilities for XML processing. For XML messaging in the form of SOAP-based Web services, PODS3 and PODS4 use a combination of Systinet and IBM WebSphere software.

However, the PODS4 implementation marked a transition from the DB2 XML Extender to the pureXML capabilities of DB2 9. For example, DB2 9 introduced support for a feature defined by the SQL:2003 standard, an XML column type treated as a first-class data type. You can use the XML type in Data Definition Language (DDL) statements, functions, and stored procedures.

Another benefit of the move to DB2 9 is the hybrid storage engine and a query optimizer that “understands” XML (mapping to relational algebra for queries involving XML). The migration to PODS4 did not change the PODS functional requirements, but pureXML processing simplified metadata processing while meeting the system’s response time and scalability goals.

The upload process illustrates the difference between PODS3 and PODS4. When there was an upload of a PDF document, for example, the PODS3 upload process stored the document and an XML metadata file on the EMC file server. It decomposed the metadata for use by the DB2 SQL storage engine as an XCollection, a type implemented by the DB2 XML Extender. The upload also validated the schema using the IBM WebSphere Application Server parser and indexed the path to the PDF and XML files in the DB2 database. The DB2 transaction associated with an upload included generating a unique document ID, logging the upload in the activity

history table, and executing a `SQL INSERT` into 18 tables.

The PODS4 upload process treats the XML document metadata differently. Instead of storing metadata across 18 tables, the PODS4 upload process saves it using DB2 columns of type XML and does an `INSERT` query into four tables.

Simplified processing and standards compliance

Moving from PODS3 to PODS4 with DB2 9 greatly simplified the UCLA Health System’s database administration and replication tasks. The PODS3 architecture used DB2 8 with the XML Extender, and a database consisting of 28 tables. The PODS4 database design, using XML columns, required only 10 tables and eliminated 20 stored procedures used by PODS3.

Because the UCLA Health System uses XML for patient metadata, supporting a new electronic form in the PODS collection involves creating the data definition or schema for form content. DBAs dealing with a constant stream of new forms want efficient solutions for defining new types of data, such as test results. Moving to DB2 9 reduced the amount of time required to add new forms and schemas to the system: supporting a new form in PODS4 typically takes two hours, compared with two weeks in PODS3.

In any industry, standards are essential for interoperability and efficient data interchange. XML has become a powerful tool for healthcare providers in part because it offers an effective tool for markup and for defining vocabularies for data interchange and archiving. However, robust applications require a reliable data management infrastructure. PODS illustrates how one healthcare provider is addressing the challenge of creating, storing, and exchanging electronic medical records. With PODS, UCLA Health System—like other healthcare institutions—has embraced XML technology and started down a path that leads to sophisticated electronic medical records, compliance with HIPAA guidelines and HL7 standards, and increased productivity. *

*Ken North is a consultant, author, industry analyst, and editor of www.SQLSummit.com. He teaches Expert Series seminars and has chaired the XML DevCon 200x conference series, *LinkedData Planet*, and *DataServices World* conferences.*

RESOURCES

IBM DB2 pureXML: ibm.com/db2/xml

DB2 pureXML enablement wiki: ibm.com/developerworks/wikis/display/db2xml/Home

Pure Concentrated Data

By Bob Scheier

Informix Dynamic Server compresses, consolidates data to boost data warehousing



MENTION INFORMIX DYNAMIC SERVER (IDS) AND MOST DBAS THINK OF A FAST, reliable, and low-maintenance platform for online transaction processing (OLTP). But those transactions hold valuable insights into business trends, leading many organizations to also use IDS as a data warehouse against which they can run business intelligence (BI) queries.

To meet those needs, Informix has unveiled a series of enhancements that bring its long-standing speed and ease-of-use benefits to data warehousing and BI. These include tools to help customers model, schedule, and execute the data transformation and data flows required to create data warehouses. Most recently, IBM announced new data compression and consolidation features that reduce the cost (and boost the performance) of both BI and OLTP applications on IDS.

The key new technology is the IDS Storage Optimization Feature. It reduces the size of not only OLTP data stores but also data warehouses that, if they grow too large, can make business analysis overly complicated and expensive.

Introducing Storage Optimization

Based on technology used in IBM DB2, Storage Optimization compresses and consolidates the data within IDS. Early customer trials show that Storage Optimization reduces the amount of space required to store data either in memory or on disk by an average of 50 percent. This can cut the time required to process queries by as much as 20 percent because more data can be kept in memory, reducing the number of I/O operations to slower mechanical disk drives. Cutting the database size can also, of course, delay or even eliminate the need to upgrade disk storage.

The first of the three Storage Optimization components is compression, which examines each row in the database for recurring patterns of data. It stores the individual recurring patterns in a dictionary, replacing those patterns with shorter strings of symbols. Unlike other compression techniques that scan only a portion of each row for repeating patterns, Storage Optimization scans the entire row, regardless of how many columns it intersects. By scanning a larger area than other compression techniques, it can find and compress more repeating patterns, and thus achieve very high compression ratios (see Figure 1).

Repack (or Coalesce), the second component, consolidates the free space created within each partition, while the final capability, Shrink, removes the unused portion of the partition and returns it for reuse by IDS. These larger, contiguous spaces are much easier for IDS to reuse than smaller, isolated free spaces. It is this compression and consolidation that speeds query performance, while holding down the amount of physical disk space required for data warehousing.

Compression on autopilot

Database compression isn't new, but Storage Optimization automates key tasks to make it easier to use. For example, it builds the compression dictionary automatically, instead of requiring the DBA to manually specify the patterns to be compressed.

Informix also provides guidelines (see sidebar, "Resources") for when optimization would be least useful, such as for small tables that fit into

memory without compression and for temporary tables and system catalogues. But even in those cases, the CPU cycles needed to compress the data and the table scans required to repack and shrink data have little impact on overall performance.

While compression and decompression do require CPU cycles, "We're finding the price you pay in CPU overhead is outweighed by the benefit you get from having fewer pages and fewer I/Os," says IDS Chief Architect Kevin Brown. "If before, your database couldn't fit in memory, and now it can, you're going to see a really big performance boost."

In cases where compression might compromise application performance, administrators can specify which fragments to compress. "If a table has 10 fragments, for example, you might want to leave the most recently created fragments uncompressed for performance purposes, and compress the older fragments, since they're less likely to be accessed," says Brown. A recompress option allows administrators to rebuild the compression dictionary to ensure optimal compression of the tables, which is useful if the data within them has changed significantly over time.

DBAs can use any or all of the Storage Optimization capabilities while the database is running, thus preventing interruption to either transactions or queries. They can also manage Storage Optimization from within the same IDS OpenAdmin Tool (OAT) they use to manage the size of volumes and to monitor performance. Graphical tools show administrators or business users how much space they can save by running any of the functions within Storage Optimization. Storage Optimization can also be controlled through a command-line interface, which is useful for the many ISVs that bundle IDS within their own applications.

The benefits of the bundle

Storage Optimization is available in a new data warehouse platform package called IBM Informix Warehouse, which includes IDS, its new Storage Optimization Feature, and the IBM Informix Warehouse Feature V11.50. Informix Warehouse




Figure 1: Saving repeating patterns as shorter character strings allows Storage Optimization to achieve very high compression ratios.

Feature V11.50 includes the SQL Warehouse (SQW) Client with Design Studio for data modeling, schema design, data transformation design, and data flow design; SQW Warehouse Server, with an administration console to schedule and manage data flows; and SQL Warehouse runtime to perform data transformation within the IDS data server.

The Informix Warehouse infrastructure also supports integrated external tooling for BI analytics from IBM Cognos, for managing data growth with IBM Optim, and for data transformation and cleansing with IBM InfoSphere DataStage and InfoSphere QualityStage.

With the new IDS data warehousing tools, DBAs can perform business analysis on their operational data stores to ensure access to the most current data. They can also create separate data warehouses within IDS to prevent analytic queries from slowing production applications or run a "shared disk" configuration in which separate servers access the same data for OLTP and business analysis.

The bottom line: businesses that have been using IDS for data warehousing and BI now have better, more targeted tools to achieve their goals. "We give organizations the choice so they can make the best use of their data and infrastructure based on their technical and business requirements," says Brown. *

Bob Scheier is a veteran IT writer based in Swampscott, Massachusetts.

RESOURCES

Data warehousing with Informix: ibm.com/informix/warehouse

IDS Storage Optimization: <http://download.boulder.ibm.com/ibmdl/pub/software/data/sw-library/informix/whitepapers/compression.pdf>

IBM Informix Warehouse: ibm.com/common/ssi/rep_ca/2/897/ENUS209-142/ENUS209-142.PDF

IBM Informix Warehouse Feature: ibm.com/common/ssi/rep_ca/1/897/ENUS209-041/ENUS209-041.PDF

DB2

Data Warehouse Performance

Part 2

Find problem SQL queries and tune them for quicker response



Robert Catterall

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

Part 1 of my column, in the previous issue of *IBM Database Magazine*, covered IBM DB2 data warehouse performance management from a system and database perspective. In this follow-on, I'll provide recommendations for SQL statement-level performance tuning.

Priority one: Target selection

When it comes to query tuning, it's important for you to do not only good work, but the right work. Your first impulse might be to work on queries with long run times, but these aren't necessarily the "problem" queries. A query that runs for 10 seconds can cause a lot of dissatisfaction if it is executed frequently and is part of a process that users expect to complete in a couple of seconds.

Often, the best indicator of where tuning work is most needed is "the voice of the user." Absent specific complaints, you'll want to spend your time on the queries that stand out in terms of both run time (and/or CPU time) and frequency of execution.

There are several tools on the market that can facilitate query tuning target selection (IBM's offerings include DB2 Query Monitor for z/OS and DB2 Performance Expert for Linux, UNIX, and Windows). However, you can get useful query target selection information from DB2 itself. Mainframers should look at the `STMTCACHE ALL` option of the `EXPLAIN` statement (introduced via DB2 for z/OS V9, and delivered for V8 by way of the fix for APAR PQ88073). `EXPLAIN STMTCACHE ALL` will insert

into the `DSN_STATEMENT_CACHE_TABLE` a row for each SQL statement in the dynamic statement cache. Among the 40-plus columns in the `DSN_STATEMENT_CACHE_TABLE` are those in which a query's statement text, accumulated elapsed time, accumulated CPU time, and number of executions are recorded. That should help you to zero in on statements that are likely to provide a good return on your investment of performance-tuning time.

Over on the Linux, UNIX, and Windows (LUW) side, DB2 folks should take advantage of the administrative views introduced with DB2 9 (the high-level qualifier for these views is `SYSIBMADM`; you can read about them in the *DB2 9 for LUW System Monitoring Guide and Reference*). One of the views that can be particularly useful is `SYSIBMADM.TOP_DYNAMIC_SQL`, which includes columns named `STMT_TEXT`, `AVERAGE_EXECUTION_TIME_S`, and `NUM_EXECUTIONS`.

Access path selection: What did DB2 decide, and why?

Statement-level tuning usually involves finding a better way to access the data needed for generating a query result set. Will a table be scanned in its entirety? Will an index be used? If tables are to be joined, what join method should be used? Answering these questions generates the query's access path; your tuning aim is to get a better path that will reduce statement execution time.

But first, you need to understand the path selected by the DB2 optimizer. That information is obtained via the DB2 `EXPLAIN` statement (`EXPLAIN` output for a statement can be examined in text form by querying `EXPLAIN` tables, or in a graphical form using the Optimization Service Center for DB2 for z/OS or the Visual Explain functionality of DB2 for LUW). Once you have the `EXPLAIN` output, look at how the query's result set is generated, and then consider alternative access path possibilities.

To properly assess these access path alternatives, you'll need information from the DB2 catalog about the tables accessed by the query and the indexes defined on those tables. Keep in mind that this catalog data is what the optimizer uses to make access path decisions. If the catalog statistics are inaccurate, get them up to date with `RUNSTATS`. That simple step could lead to a new access path that might dramatically improve query performance.

When you've looked at the access path selected by DB2 for your target query, ask yourself this question: Should DB2 have selected a different access path for the query? If a path other than the one selected by DB2 looks better to you, check to see if the tables targeted by the query—and their associated indexes—are well organized. If the catalog statistics are up to date and DB2 knows that a particular index is poorly organized, DB2 may decide not to use it. In that case, reorganizing the index might lead to its selection as part of a better-performing query access path.

A more likely explanation for a difference of opinion between you and DB2 is that you know something that DB2 doesn't. Perhaps you know that data values in a certain column are not evenly distributed. Without column data-value distribution information in the catalog, DB2 will assume an even distribution of values, which may lead the optimizer to a sub-optimal access path.

If a significant data value skew exists for a column referenced in a query predicate, make sure that DB2 knows this, preferably by generating column-value histogram statistics using `RUNSTATS`. (Histogram statistics, available

in DB2 for LUW for quite some time, can be generated with DB2 for z/OS V9 `RUNSTATS`. If you are using DB2 for z/OS V8, the next best thing is the `FREQVAL` option of `RUNSTATS`.) Note that advice for enriching DB2 for z/OS catalog statistics for the benefit of a particular query can be obtained from the Optimization Service Center for DB2 for z/OS—a free version of which can be downloaded from ibm.com/db2/zos/downloads/osc.html.

When you and DB2 agree

If you and DB2 agree on the best available access path but that path does not deliver satisfactory performance, you need to provide DB2 with a new access path choice. There are several ways to do this, including adjusting your indexes, using a materialized query table, and table re-clustering.

Index-related actions can reduce the number of pages that DB2 has to examine to generate the query's result set. You might define an index on a predicate-referenced column that is not currently indexed. You might define an index on a set of columns that is already indexed, but with a different column order specified to increase the number of index key columns on which DB2 can perform predicate-related value matches. (For example: given a compound predicate of the form `COL_A > y AND COL_B = x`, and an existing index on `COL_A | COL_B`, a new index on `COL_B | COL_A` would enable DB2 to match on both columns of the index key instead of just matching on the first key column.) You might add a column to an existing index to get index-only access for a portion of the query's processing.

Another way to provide a new and better access path for a query is through the use of DB2's materialized query table (MQT) functionality. An MQT is defined by way of a `SELECT` statement—usually a result set that would otherwise have to be dynamically built when a query executes. Here are three great things about an MQT:

1. DB2 can automatically rewrite a query to take advantage of an MQT.

2. Because the result set in the MQT is already there, the run time for a query that is rewritten by DB2 to access the MQT can be reduced dramatically, saving time that would otherwise be spent dynamically building that result set at query execution time.
3. An MQT can be indexed to provide an even greater boost to query performance.

Table re-clustering is another means of providing DB2 with a better access path by reducing the number of pages that DB2 has to examine in executing the query. Data clustering can have an especially large impact on data warehouse query performance, because rows are often retrieved in large numbers. If a table was originally defined to be clustered on an account number column, and it subsequently turned out that users were retrieving rows from the table on a date-range basis, it might be worth considering a change in the table's clustering sequence to get retrieved rows into close proximity to each other in the table. Also, the performance of join operations may be improved significantly if joined tables are clustered in join-column sequence (e.g., by `CUSTOMER_ID` when the join predicate is `TABLE_A.CUSTOMER_ID = TABLE_B.CUSTOMER_ID`).

If you do decide to re-cluster a table, keep in mind that DB2 for z/OS V9 makes this much easier through the new `CLUSTER` and `NOT CLUSTER` options of the `ALTER INDEX` statement. If re-clustering is on the table, make sure that you give some thought to exploiting clustering-related enhancements delivered in the last couple of DB2 versions, such as multi-dimensional clustering for DB2 for LUW, and the ability to partition a table on one key and cluster within partitions on another key in a DB2 for z/OS environment.

Here's hoping that your data warehouse query-tuning efforts will deliver breakthrough results. Business intelligence is as hot an application area as you can find these days, and getting into it (if you're not already there) will give you all kinds of opportunities to deliver real value to your organization. Tune on! ✨

Reducing Downtime with HADR

A step-by-step guide to minimizing the impact of failures on your database environment

In a perfect world, mission-critical databases would never be subjected to issues like hardware, network, or software malfunctions, nor would they be exposed to natural disasters like fire and flood. But in reality, any database can be impacted by one of these events and the consequences can vary—downtime for recovery, loss of critical data, or the need to completely rebuild the entire database infrastructure. If your database environment consists of single-partition DB2 databases, you can minimize the consequences of such events by taking advantage of an IBM DB2 for Linux, UNIX, and Windows feature known as High Availability Disaster Recovery (HADR).

What is HADR?


HADR is a DB2 database replication feature that is designed to minimize the impact to a database system when a partial (hardware, network, or software malfunction) or a complete site (fire, flood, and so forth) failure occurs. HADR protects against data loss by replicating data changes from a source database, called the *primary*, to a target database, called the *standby*. Synchronization with the standby database occurs by rolling forward transaction log records that were generated for the primary database and have been shipped to the standby database (see sidebar, “What gets replicated?”).

Without HADR, a partial site failure requires restarting a server, one or more

instances of DB2, and the databases under each instance’s control. If the transaction load was heavy at the time the partial site failure occurred, it can take several minutes to restart the instances and the databases affected. With HADR, a failover occurs and the standby database takes over, replacing the failed primary database in seconds. Furthermore, you can automatically redirect any clients that were using the primary database to the standby database either by using Automatic Client Reroute or by adding retry logic in each application that interacts with the database.

Because HADR uses TCP/IP to communicate between the primary and the standby database, each database can reside in a different location. For example, the primary database might be located at one data center in one city, while the standby database might be located at another data center in another city. If a complete site failure occurs at the primary site, data availability is maintained by having the remote standby database take over as the primary database.

After the emergency is over and the failed original primary server/database has been repaired, it can rejoin the HADR pair as the standby database—provided both copies of the database can be made consistent. And once the original primary database is reintegrated into the HADR pair as the standby database, you can switch the roles so that the original primary database once again functions as the primary database. This is known as a *failback* operation.



Roger E. Sanders

(roger_e_sanders@yahoo.com)
is consultant corporate systems engineer at EMC Corporation. He is the author of 18 books on DB2 for Linux, UNIX, and Windows and teaches classes at many DB2 conferences. 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 Dale McInnis, senior technical staff member (STSM) and DB2 availability architect at the IBM Toronto Lab, for reviewing the material presented in this article.

Requirements for HADR environments

To take advantage of HADR, you must have two dedicated servers available: one for the primary database and one for the standby database. Furthermore, both servers must adhere to a set of very specific requirements (see sidebar, “Server setup checklist”).

Setting up HADR

Setting up an HADR environment is fairly straightforward. First, ensure that the systems to be used as primary and standby servers are identical and that a TCP/IP connection exists between them. Then follow these steps:

1. **Determine the host name, host IP address, and the service name or port number** for both the primary and the standby database server. If a server has multiple network interfaces, ensure that the HADR host name or IP address maps to the intended interface. You will need to allocate separate

HADR ports for each protected database—these cannot be the same as the ports that have been allocated to the DB2 instance. The host name can map to only one IP address.

2. **Create the standby database** by restoring a backup image or initializing a split mirror copy of the database that will serve as the primary database on the standby server. When using the `RESTORE DATABASE` command to create the standby database from a backup image, IBM recommends that the `REPLACE HISTORY FILE` option be specified. Avoid use of the `TABLESPACE`, `INTO`, `REDIRECT`, and `WITHOUT ROLLING FORWARD` options.

When using the `db2inidb` utility to create the standby database from a split mirror copy, use the `AS STANDBY` option (to initialize the split mirror image as a clone); the `AS SNAPSHOT` and `AS MIRROR` options should be avoided. The `RELOCATE USING` option can be specified to change one or

more of these configuration attributes: instance name, log path, and database path. However, you must not change the database name or the table space container paths.

Once you have created the standby database, do not use the `ROLLFORWARD DATABASE` command as a recovery option. The results of performing a roll-forward recovery operation could differ slightly from replaying the logs on the standby database using HADR. If the primary and standby databases are not identical when HADR is started, an error will occur.

3. **Assign appropriate values to the HADR configuration parameters** for both the primary and the standby database. After the standby database has been created, but before HADR is started, set the HADR-specific configuration parameters shown in Table 1 for both the primary and the standby database. You can choose your level of protection from

WHAT GETS REPLICATED?

Operations replicated under HADR:

- ▶ Execution of Data Definition Language (DDL) statements (`CREATE`, `ALTER`, `DROP`)
- ▶ Execution of Data Manipulation Language (DML) statements (`INSERT`, `UPDATE`, `DELETE`)
- ▶ Buffer pool operations
- ▶ Table space operations, as well as storage-related operations performed on automatic storage databases (`ALTER DATABASE`)
- ▶ Online reorganization
- ▶ Offline reorganization
- ▶ Changes to metadata (system catalog information) for stored procedures and user-defined functions (UDFs)

Operations not replicated under HADR:

- ▶ Stored procedure and UDF object and library files
- ▶ Non-logged operations, such as changes to database configuration parameters and to the recovery history file

SERVER SETUP CHECKLIST

Here's a checklist of HADR setup items that shouldn't be missed:

Hardware

Server: The servers hosting the primary and standby databases should be as similar as possible, from the vendor to the architecture

Memory: Ideally, identical amounts of memory should be available to support replayed buffer pool operations

Network: A TCP/IP interface is required between the HADR host systems, connected over a high-speed, high-capacity network

Software

Operating system: Identical versions on both servers, including patch level

DB2 software: Identical bit size (32 or 64) and version—during rolling upgrades, the modification level of the database system for the standby database can be higher than that of the primary database for a short while; always apply fix packs to the standby database system first

Database configuration

- ▶ Single-partition databases with identical names
- ▶ Identical storage space allocated for transaction log files; raw devices are not supported
- ▶ Archival logging on current primary database only
- ▶ Table space properties, such as table space name, table space type (DMS, SMS, or Automatic Storage), page size, size, container path, container size, and container type (raw device, file, or directory), must be identical on both databases
- ▶ Before adding containers to existing table spaces by executing the `ALTER TABLESPACE` statement against the primary database, table space containers involved with such statements must exist on both systems
- ▶ Automatic storage databases are supported; any storage paths specified with `ALTER DATABASE` statements must exist on both the primary and the standby server

potential loss of data by specifying one of three synchronization modes:

- ▶ **Synchronous:** A database change made on the primary database won't be committed until the associated log records have been written to the transaction log file on the standby DB2 instance.
- ▶ **Near synchronous:** Same as synchronous mode, except that the change made to the primary database can be committed when the associated log records have been written to memory on the standby DB2 instance (rather than written to the standby instance's transaction log).
- ▶ **Asynchronous:** Committing a database change on the primary database is not held up in any way by the communication of log records to the standby database.

4. **Connect to the standby database** and start HADR on the standby server by executing the `START HADR` command.

5. **Connect to the primary database** and start HADR on the primary server with the `START HADR` command.

You can also set up an HADR environment using the Set Up HADR Databases Wizard, which can be activated by selecting the High Availability Disaster Recovery action from the Databases menu found in the Control Center.

Configurations for Automatic Client Reroute and HADR

Automatic Client Reroute is a DB2 feature that redirects client applications from the primary database to the standby database immediately after a failover operation is performed. If you set up HADR using the Set Up HADR Databases Wizard, Automatic Client Reroute is enabled by default. If you set up HADR manually, you can enable the Automatic Client Reroute feature by executing the `UPDATE ALTERNATE SERVER FOR DATABASE` command;

Parameter	Value range/default	Description
<code>hadr_db_role</code>	N/A	Read-only; indicates the current role of the database, if it is part of a HADR environment. Valid values are <code>STANDARD</code> , <code>PRIMARY</code> , or <code>STANDBY</code> .
<code>hadr_local_host</code>	Any valid character string Default: <code>NULL</code>	Specifies the local host for HADR TCP communication; either a host name or an IP address can be used.
<code>hadr_local_svc</code>	Any valid character string Default: <code>NULL</code>	Specifies the TCP service name or port number for which the local HADR process accepts connections.
<code>hadr_remote_host</code>	Any valid character string Default: <code>NULL</code>	Specifies the TCP/IP host name or IP address of the remote HADR node.
<code>hadr_remote_inst</code>	Any valid character string Default: <code>NULL</code>	Specifies the instance name of the remote server. Administration tools, such as the Control Center, use this parameter to contact the remote server. HADR also checks whether a remote database requesting a connection belongs to the declared remote instance.
<code>hadr_remote_svc</code>	Any valid character string Default: <code>NULL</code>	Specifies the TCP service name or port number that will be used by the remote HADR node.
<code>hadr_syncmode</code>	<code>SYNC</code> , <code>NEARSYNC</code> , <code>ASYNC</code> Default: <code>NEARSYNC</code>	Specifies the synchronization mode to use for HADR. This determines how primary log writes are synchronized with the standby database when the systems are in peer state. Valid values for this configuration parameter are <code>SYNC</code> (Synchronous), <code>NEARSYNC</code> (Near Synchronous), <code>ASYNC</code> (Asynchronous).
<code>hadr_timeout</code>	1 – 4,294,967,295 Default: 120 seconds	Specifies the time (in seconds) that the HADR process waits before considering a communication attempt to have failed. (The value assigned to this configuration parameter must be the same for both the primary and the standby database.)
<code>hadr_peer_window</code>	0 – 4,294,967,295 Default: 0 seconds	Specifies the time (in seconds) that a HADR primary-standby database pair continues to behave as though still in peer state, if the primary database loses connection with the standby database. This helps ensure data consistency.

Automatic Client Reroute does not use the values stored in the `hadr_remote_host` and `hadr_remote_svc` database configuration parameters. For example, if a database named SALES, located at SVR1, is the primary database in an HADR environment and its corresponding standby database (also named SALES) resides on host SVR2 and listens on port number 456, you could set up Automatic Client Reroute by executing a command that looks like this:

```
UPDATE ALTERNATE SERVER FOR DATABASE sales
USING HOSTNAME svr2 PORT 456
```

Once this command is executed, clients connect to host SVR1 to obtain the alternate server information. If a communication error occurs between a client and the SALES database at host SVR1, the client will first attempt

to reconnect to the SALES database on host SVR1. If this fails, it will then attempt to establish a connection with the standby SALES database located on host SVR2.

A helping hand for disaster recovery

HADR provides an excellent solution for minimizing the impact on a database environment when hardware, network, or software malfunctions or natural disasters occur. In addition to reducing data loss and downtime, HADR enables an organization to keep a disaster recovery (DR) copy of a database available and up-to-date, even when the DR site is hundreds of miles away. And if Automatic Client Reroute is used in conjunction with HADR, client applications can automatically be rerouted to the DR database copy or site in the event of a failover.✱

Cool **New** Stuff for Informix

Informix Warehouse 11.50, Informix 11.50.UC4 release with data compression, and an exciting user conference

The first half of this year has been an exciting time. The new Informix Data Warehouse functionality was announced by IBM, a new release of Informix was announced, and the International Informix User Group Informix Conference was held in Kansas City, April 26 to 29. I am sure others have reported on the conference details so I will focus on the contest we held to honor the Fastest Informix DBA. And I will share with you some of my findings with IBM Informix Warehouse and data compression in the IBM Informix Dynamic Server 11.50.UC4 release.

Fastest Informix DBA contest

My company, Advanced DataTools, sponsored a contest to find the Fastest Informix DBA. I took a simple customer billing process, added some bad SQL, and added a default `ONCONFIG` file with a couple of bad configuration options. The benchmark unchanged took about 30 minutes to run and I challenged participants to make it run faster. I knew this was going to be fun because when I tried this out on my staff the week before the conference, they could not stop running the benchmark. It became an exciting game to make it all go faster—even I got hooked into finding different ways to get it to run faster! We gave each participant about 45 minutes to:

1. Tune and optimize the `ONCONFIG` file
2. Tune and optimize the dbspace and table layout
3. Optimize the SQL script
4. Run the benchmark script


The grand prize was an Apple iPod touch for the Fastest User DBA, who had to be a paid attendee at the conference (or conference speaker or volunteer, at the discretion of the judges), not a consultant, not an IBM employee, and employed by a company using Informix internally. We also had 10 categories for many other winners (see sidebar, “Fastest Informix DBA 2009 Winners”).

Congratulations to everyone who participated—this was a blast and I think everyone had fun and did a great job. They all took an `ONCONFIG` file and SQL script and tuned it so that it would run in less than four minutes. We are planning to do a podcast or webcast with the winners to find out how they did it. I will also use the analysis from how the winners did it as material for future articles in this magazine. And we are looking at ways to do this again, so stay tuned!

Informix Data Warehouse function release

I have been building data warehouses using Informix for over 12 years, as it is lighting-fast for loads and queries. I have built terabyte-sized systems with thousands of users, and it works great. In April, IBM announced the Informix Data Warehouse Feature and Informix Warehouse V11.50. I was in the beta program for this function and I am really excited about the new features.

The key feature is the SQL Warehouse (SQW) Design Studio, which allows you to build data connections to source and target databases, reverse engineer the source



Lester Knutsen

(lester@advanceddatatools.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.

physical data models, and design the target data warehouse data model. It then enables you to build the extract and load process, using SQL data flows or your own command and control functions.

I have a large library of scripts, programs, and functions that I use for loading data, and the ability from within Design Studio to call external programs or scripts from a command box made it very fast and easy to convert some of the load programs I had already developed. You can also use the command box to call Informix High Performance Loader scripts from Design Studio; once you have the process and data transformation rules defined and tested, there is a separate Warehouse Administration Console you can use to manage, schedule, and run your load process.

Data compression in Informix 11.50.UC4

Another key announcement was the release of fix pack UC4 for Informix 11.50, which includes data compression. I was also in the beta for this, and besides saving space, compression may speed up your system. The standard benchmark I use (and the Fastest Informix DBA contest was based on this benchmark) showed about a 25 percent performance improvement when using compression. With compression, fewer pages need to be read into memory from disk, making the overall system perform better. The larger the system, the better this seems to work.

To enable compression on a server, run the following SQL command to execute a function in the sysadmin database:

```
EXECUTE FUNCTION sysadmin.task ("enable compression");
```

Next you can compress and uncompress specific tables. To compress a table, run the following SQL command to execute the function in the sysadmin database:

```
EXECUTE FUNCTION sysadmin.task ("table compress", "table_name", "database_name", "table_owner");
```

Table 1: Common commands to compress, pack, and shrink tables

Turn on compression in a server	EXECUTE FUNCTION sysadmin.task ("enable compression");
Estimate how much space may be saved by compressing a table	EXECUTE FUNCTION sysadmin.task ("table estimate_compression", "table", "database", "owner");
Create a compression dictionary without compressing a table	EXECUTE FUNCTION sysadmin.task ("table create_dictionary", "table", "database", "owner");
Compress a table (automatically creates a dictionary)	EXECUTE FUNCTION sysadmin.task ("table compress", "table", "database", "owner");
Compress a table and return the freed table space to the pool of free space	EXECUTE FUNCTION sysadmin.task ("table compress repack shrink", "table", "database", "owner");
Uncompress a table	EXECUTE FUNCTION sysadmin.task ("table uncompress", "table", "database", "owner");
Compress a fragment of a partitioned table	EXECUTE FUNCTION sysadmin.task ("fragment compress", "partnum");
Consolidate space in a table without compressing it	EXECUTE FUNCTION sysadmin.task ("table repack", "table", "database", "owner");
Consolidate space in a table, locking and taking the table offline, without compressing it	EXECUTE FUNCTION sysadmin.task ("table repack_offline", "table", "database", "owner");
Consolidate space in a fragment, locking and taking the table offline, without compressing it	EXECUTE FUNCTION sysadmin.task ("fragment repack_offline", "partnum");
Shrink a table and return space to the pool of free space	EXECUTE FUNCTION sysadmin.task ("table shrink", "table", "database", "owner");
Shrink a fragment and return space to the pool of free space	EXECUTE FUNCTION sysadmin.task ("fragment shrink", "partnum");

Replace `table_name` with the name of your table, `database_name` with the database, and `table_owner` with the name of the owner. These are the only two commands I had to add to my benchmark script to see the performance gains with compression. There are other commands that can be used to uncompress tables, control locking while compressing, and to display space savings. IBM also has an IDS Compression Estimator program that you can download from ibm.com/Informix to see the space savings.

Conference kudos

This spring was a great time for Informix with all the new features released, and the IIUG Informix Conference in Kansas City at the end of April topped it off. The conference could not have taken place without the hard work of a core group of volunteer users on the steering committee. I would really like to thank the conference planning committee for all their work that resulted in such a fine user conference. ✱

FASTEST INFORMIX DBA 2009 WINNERS

- Grand Prize: Fastest User DBA—Andrew Ford
- Runner-Up and Fastest Mid-Aged DBA (30–50) and Domestic DBA—Tom Girsch
- Fastest IBM Developer—Spokey Wheeler (best overall time)
- Fastest Old-Timer (longest experience with Informix)—David Link
- Fastest Senior DBA (over 50)—John Fahey
- Fastest Professional Consultant DBA—Warren Donovan
- Fastest Youngest, Newest, and Fastest International DBA—Jean-Michel Ramseyer



How to Get Certified for DB2

What you need to know about sign-up details, process information, and exam formats



*Howard Fosdick is a certified DBA for DB2 and Oracle. His book *Rexx Programmer's Reference* covers everything Rexx—including database programming—and is available at www.amazon.com/rexx.*

With the exciting new design of *IBM Data Management* magazine, it's appropriate to revisit IBM DB2 certification as well.

Why certify? Consider both the tangible and intangible benefits. On the practical side, the *Certification Magazine* annual salary survey consistently shows that certified IT professionals make more than those who are not certified.

But there's more to the decision than salary. Certification rounds out your DB2 knowledge and ensures that you've been exposed to the product's many features. You'll also have the satisfaction of increasing your on-the-job productivity.

Most database vendors have certification programs, but the DB2 certification program is unique in several respects. First, you obtain a DB2 certification for each and every exam you pass. This sharply contrasts to certifications for competing database products, where you must pass more than one exam to obtain a single cert. With DB2 you can pursue certification, stop after any exam you pass, and still have something valuable to show for your efforts.

Second, IBM provides tons of free information and study materials to help you learn about and prepare for the tests. I've listed some of these in the Resources section at the end of this column. Leverage these freebies and you'll save money—while studying with materials prepared by those who know DB2 best.

Exam Number	Certification	Platforms
730	IBM Certified DB2 Database Associate	DB2 9
731	IBM Certified Database Administrator	Linux, UNIX, and Windows (LUW)
732	IBM Certified Database Administrator	IBM z/OS
733	IBM Certified Application Developer	LUW and z/OS

All DB2 certification candidates start by taking the same test, #730, called "DB2 9 Family Fundamentals." Once you pass this exam you have the option to take any of several exams to achieve higher-level certification (see Table 1).

You sign up online to take any of the IBM certification exams through Prometric's IBM test site (www.prometric.com/IBM/default.htm). Prometric runs testing centers in cities and towns across the nation.

The exams cost US\$200 each, but IBM runs frequent promotions. When you visit the IBM Information Management training and certification Web page, you'll see any current promotion prominently displayed. You can also certify for free while attending IBM's Information On Demand conference or the International DB2 Users Group conference.

Each exam contains about 60 to 70 multiple-choice questions. You get credit for every question you answer correctly and must answer roughly 60 percent correctly to pass. You have 90 minutes to complete the exam. (These numbers vary slightly by exam.)

Next steps

I've listed some of the key IBM certification Web sites in the "Resources" sidebar. Peruse them and spend some quality time there; I've also written an article on IBM developerWorks called "DB2 9 Certification: Everything You Need to Know." It covers all the bases and gives you a good overview of the entire process. ✱

RESOURCES

IBM Information Management certification home page: ibm.com/software/data/education/certification.html

"DB2 9 Certification: Everything You Need to Know":

ibm.com/developerworks/data/library/techarticle/dm-0401fosdick

DB2 Information Management online practice tests:

ibm.com/software/data/education/cert/assessment.html

Information Management books: ibm.com/developerworks/dwbooks

DB2 certification on Facebook: www.facebook.com/pages/IBM-DB2-Training-and-Certification/14382814766?ref=mf



Smarter is...

Predictive, Preventive Healthcare

Clinical data warehouses turn information into smarter treatments

Johnny's doctor is stumped. In this small-town hospital, she's never seen these symptoms before, and the nearest specialist is 200 miles away. There's no immediate danger, but how can she hone in on the problem, choose and conduct the right tests, and get a treatment plan—especially one the insurance payer will approve?

Getting a handle on healthcare costs is one of society's greatest challenges. In 2007, US\$2.2 trillion was spent on healthcare,¹ and that figure is expected to grow steadily while quality of care continues to vary dramatically.

Applying information technology to the problem has been a difficult task, even though some of the purely administrative benefits are easy to imagine. But what if we use healthcare data to improve decision making across the entire provider organization?

IBM partner Convergence CT, a clinical data warehouse provider, is one of the leaders in making this "what if" a reality. Working together with IBM, they have combined IBM InfoSphere and IBM Cognos technologies with healthcare expertise to create IBM InfoSphere Clinical Analytics, a pre-integrated clinical data warehouse. With preloaded data definitions for industry standards, predefined performance indicators, and prebuilt analytics, the warehouse gives providers a sophisticated tool that would take IT and medical specialists enormous amounts of time to create independently.

Seventy-five person-years of expertise and construction have gone into building a heuristic model of clinical analytics. The warehouse captures data from multiple sources, cleanses it, and stores it, performing more than 900 data validations to ensure that the data can be trusted.

IBM InfoSphere Clinical Analytics also addresses an important requirement of any healthcare IT project or solution: helping medical staff focus on healthcare duties instead of on data management. For example, predesigned dashboards and data cubes help staff track surgical outcomes and analyze benchmarks for practice management and clinical performance. Meanwhile, reports that automatically comply with—and document compliance to—privacy regulations help take the compliance burden off medical practitioners. Instead of forcing a nurse with patient responsibilities to maintain data quality and anonymity standards, the reports and applications take care of it.


The equations are simple. For the hospital, research facility, or medical practice:

**server + data management software +
clinical models = ready-to-use tools**

And for the patient:

**predictive + prescriptive = better health
outcomes**

Imagine that. ✱



Merv Adrian, an analyst and consultant, founded IT Market Strategy after three decades in the IT industry. During his tenure as senior vice president at Forrester Research, he was responsible for all of Forrester's technology research, covered the software industry, and launched Forrester's well-regarded analyst relations practice.

¹ Source: U.S. Department of Health and Human Services, www.cms.hhs.gov/NationalHealthExpendData/25_NHE_Fact_Sheet.asp

RESOURCES

IBM InfoSphere Clinical Analytics:
ibm.com/software/data/infosphere/clinical-analytics

Convergence CT: www.convergencect.com



When it comes to
performance

Rely on

Informix[®]

When it comes to
knowledge

Trust



**International
Informix
Users Group**

Join us, it's free.
Visit <http://www.iiug.org>

Whether it's with our conference, with our web TV—IIUG.tv, with Cheetah—IDS v11.5, or our forums, Insider, or web site, our passion is in sharing knowledge.

Come and join IIUG.

We speak Informix.

Join us online

- ▶ **Quick content searches** throughout the entire issue
- ▶ **Direct links** to related community resources
- ▶ **Easy information access**, sharing, and printing




**data
management**
KNOWLEDGE. PERFORMANCE. RESULTS.

Visit ibm.com/dmmagazine and sign up for your complimentary digital subscription.

You'll get the same in-depth technical content, practical advice, and hands-on tips about how to improve productivity and performance as the print edition—now including real-world commentary about how your peers are using data and information architectures to reduce costs and improve business results.

IBM.COM/DMMAGAZINE