

# data management

IBM.COM/DMMAGAZINE

KNOWLEDGE. PERFORMANCE. RESULTS.



**Tune  
it up,  
make it  
pay**

Maintenance moves  
that put the squeeze  
on costs

**PROTECT AND  
SERVE**

Inside InfoSphere  
Guardium

**DISTRIBUTED DBA**

Partition tables  
for performance

Building the engines of a Smarter Planet:


## Robust backup. Rapid recovery. Reliable return on your investment.

As the engines of a smarter planet, midsize businesses are busy fueling innovation and a global economy. And with so much information flowing across their organizations, the implications of lost data can be tremendous. They need a dynamic infrastructure that provides smarter ways to help protect and quickly recover data when anything goes wrong, no matter the cause. And they need a reliable, cost-effective solution. The IBM® Comprehensive Data Protection Solution Express™ can help protect your company from loss of data by providing automated backup and near-instant recovery. In other words, every time you make a change, it makes a backup—which helps to shorten backup gaps from days to minutes. Here's how to get started:

**1 Understand where you are.** Know your business risks and what's most important to you now. Evaluating your current recovery processes and prioritizing critical data help you see risk exposure and develop a backup schedule.



**2 Recover data quickly.** IBM Tivoli® Storage Manager FastBack automatically looks for data as it's created and recovers it when needed. Worry less with a backup schedule that you define. IBM technologies can recover data in just minutes.




**3 Protect your assets.** IBM System x® Express servers pull data and send it to the IBM System Storage® DS3200, which houses data in a reliable location.

The IBM Comprehensive Data Protection Solution Express includes IBM Tivoli Storage Manager FastBack, the IBM System x3250 M3 Server and the IBM System Storage DS3200 array, priced together starting at

**\$364** per month  
for 48 months<sup>1</sup>



**4 See the savings.** The data protection solution can pay for itself in less than a year by helping reduce the risk of revenue loss and the costs of manual data backup and recovery. Savings are estimated at \$15,000 annually.




### Midsize businesses are the engines of a Smarter Planet.

The IBM Express Advantage™ Concierge can connect you to the right IBM Business Partner. Call **877-IBM-ACCESS** or visit [ibm.com/engines/data2](http://ibm.com/engines/data2)



<sup>1</sup>Actual savings and costs will vary depending on individual customer configurations and environment. IBM Global Financing offerings are provided through IBM Credit LLC in the United States and other IBM subsidiaries and divisions worldwide to qualified commercial and government customers. Rates are based on a customer's credit rating, financing terms, offering type, equipment type and options, and may vary by country. Other restrictions may apply. Rates and offerings are subject to change, extension or withdrawal without notice. IBM, the IBM logo, ibm.com, Smarter Planet and the planet icon are trademarks of International Business Machines Corp., registered in many jurisdictions worldwide. Other product and service names might be trademarks of IBM or other companies. A current list of IBM trademarks is available on the Web at [www.ibm.com/legal/copytrade.shtml](http://www.ibm.com/legal/copytrade.shtml). © International Business Machines Corporation 2010.



# IDUG

The Worldwide  
DB2 User Community

## IDUG EMEA

8–12 November 2010

Hilton Vienna Hotel  
Vienna, Austria

### Can't make Vienna?

IDUG North America  
May 2-6, 2011  
Anaheim, California  
USA

## DB2 Users Working Together What's Upcoming? IDUG EMEA

Formerly IDUG Europe, IDUG EMEA is *the* conference for user-delivered, unbiased, trusted DB2 content in Europe, Middle East and Africa. IDUG EMEA delivers:

- User-driven technical sessions to sharpen your skills
- Best practices to drive efficiencies and maximize your investments
- Direction on new and existing database and related product issues
- Candid discussions with customers, IBM experts, developers and DB2 consultants
- A customer loyalty program that affords 80% in company savings for those qualified

### IDUG brings you the best of the best in DB2.

IDUG EMEA is just one way to tap in. Explore all the benefits of the #1 DB2 User Group—visit [www.IDUG.org](http://www.IDUG.org).

MENTORING

CERTIFICATION

EXPERTS

NETWORKING

EDUCATION



# data management

KNOWLEDGE. PERFORMANCE. RESULTS.

ISSUE 3, 2010 / VOLUME 15 NO. 3

## Features

### 18 InfoSphere Guardium

How this database security and monitoring software works its magic

### 25 Think Inside the Box

Strategies for using Informix as an embedded database



## Maintenance Makes Money

Need to trim database management costs? Look closely.

### 22 Industry Focus: Energy and Utilities

A solid data architecture sparks smart grid builders

## Departments

#### 4 Editor's Note

By Cameron Crotty

#### 6 NewsBytes

#### 10 IIUG User View

By Stuart Litel

#### 11 IDUG User View

By David Beulke

#### 28 Data Architect

You Want to Govern What?

By Robert Catterall

#### 31 Distributed DBA

Creating and Using Partitioned Tables

By Roger E. Sanders

#### 34 Programmers Only

Reducing

Conversations with DB2 for z/OS: Part 3

By Bonnie Baker

#### 38 Informix DBA

Building Fast

Data Warehouse

Schemas: Part 2

By Lester Knutsen

#### 40 Smarter is...

Delivering social services where they're needed most

By Crystal Bedell

### Ad Index

<b>DBI</b> .....	C3
www.DBISoftware.com	
<b>IBM</b> .....	C2
www.ibm.com	
<b>IBM Information On Demand</b> .....	5
www.ibm.com/events/informationondemand	
<b>Intel</b> .....	C4
www.intel.com	
<b>International DB2 Users Group</b> .....	1
www.idug.org	
<b>Melissa Data</b> .....	36
www.melissadata.com	
<b>Quest Software</b> .....	9
www.quest.com	
<b>Relational Architects International</b> .....	37
www.relarc.com	
<b>Responsive Systems</b> .....	3
www.responsivesystems.com	
<b>Safari Books Online</b> .....	30
www.safaribooksonline.com	

Optimize Both Performance & DB2 Memory Utilization With

# Buffer Pool Tool<sup>®</sup> 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

Tel: 732 972.1261

Web: [www.responsivesystems.com](http://www.responsivesystems.com)

Morganville, NJ 07751

Fax: 732 972.9416

**EDITOR**

Cameron Crotty  
editor@tdagroup.com

**MANAGING EDITOR**

Stephanie S. McLoughlin

**ART DIRECTOR**

Iva Frank

**DESIGNERS**

David Chan, Lalaine Gagni, Margie Preston

**CONTRIBUTING WRITERS**

Merv Adrian, Bonnie Baker, Crystal Bedell,  
David Beulke, Robert Catterall, Albert Cooley,  
Mirav Kapadia, Lester Knutsen, Sheryl M. Larsen,  
Stuart Litel, Roger E. Sanders

**EDITORIAL BOARD OF DIRECTORS**

Jeannie Cramer, Jeff Jones, Steve Miller, Jennifer  
Reese, Bernie Spang

**AD SALES**

Jill Thiry  
advertise@tdagroup.com

**AD COORDINATORS**

Katherine Hartlove  
Kim Johnsen

**SPECIAL THANKS TO**

Lea Anne Bantsari  
Matthew Golden

**SUBSCRIPTION SERVICES AND REPRINTS**

To subscribe to the print or digital version of *IBM Data Management* magazine, change your address, or make other updates to your information, please go to [ibm.com/dmmagazine](http://ibm.com/dmmagazine). For instant access to the *IBM Data Management* magazine digital edition, visit [ibmdmmagazinedigital.com](http://ibmdmmagazinedigital.com). For information about reprints, please send an e-mail to [customerservice@tdagroup.com](mailto: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 © 2010, 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.



we learn and we memorize, but technology keeps advancing. New tools provide better ways to accomplish our goals, and we have to adjust. Take driving a car: If you're of a certain age, you may remember learning to pump your brakes on a slick road to keep the wheels from locking up and sending you into a skid. But in these days of antilock braking systems (ABS), that's exactly what you're not supposed to do: when you hit an icy patch, the best way to stop is to stand on the brake pedal and let the ABS chip handle the rest.

When it comes to DBMSs... well, you can see where this is headed. I don't want to steal any thunder from Sheryl Larsen's fantastic cover story on how keeping your databases up-to-date can bring serious rewards, so go read it now! Practically every year, the development teams at IBM release new, more efficient, more effective ways to get the results that you are looking for. Do your databases and applications take advantage of those new features? Or are you still pumping the brake pedal on a car that has ABS, and paying the price in reduced performance?

Beyond Sheryl's hands-on advice, we have lots of new developments to explore. Earlier this year, IBM acquired Guardium, an organization that brings an innovative approach to real-time enterprise database monitoring and protection. In this issue, we pull back the curtain on the technical details of the Guardium solution—it's a story that you won't want to miss.

It's also turning out to be a very exciting year for Informix. You may have already heard the buzz, but skip on over to Stuart Litel's column for the full scoop. Need more Informix? Check out the feature story on preparing Informix for embedded deployments and join Lester Knutsen for Part 2 of his series on architecting data warehouses.

There's lots more to see in this issue, so I'll leave you with one last thought: If you haven't started thinking about attending the Information On Demand Global Conference this October in Las Vegas, what are you waiting for? From what I'm hearing as we plan the next issue, it's going to be worth the trip.

Thanks for reading,

Cameron Crotty  
Editor



# Register Today

## Information On Demand 2010

The Premier Conference for Information and Analytics

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

Drive innovation and business optimization with the best in technical and business education. More than 300 customers will share their real-world experiences about how they are transforming their businesses through information and analytics.

### Register Early and Save!

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


### Event Highlights

- More than 9,000 conference attendees
- 300 customer speakers
- 110 Industry-focused Business & IT Leadership sessions
- 600 Technical Skill Building sessions including 160 Cognos and SPSS sessions
- 128 Hands-on labs
- 3 free Information Management certification tests
- IBM's largest EXPO
- Networking opportunities

### Register Today!

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

[ibm.com/events/InformationOnDemand](http://ibm.com/events/InformationOnDemand)





## Workload-Optimized Systems Hit the Mark

IBM integrates hardware, software, and services to maximize performance

In these days of virtualization, off-the-shelf hardware, and flexible infrastructure, it can be easy to forget that different types of computing jobs place different demands on software and hardware. You wouldn't try to edit photographs on a netbook, and you'd hardly fire up your fastest workstation for a quick e-mail session.

IBM has long been a proponent of deploying targeted systems tuned to the task at hand, and the company recently advanced that argument by announcing a series of systems consisting of integrated servers, storage, software, and services—all configured and tuned to deliver maximum performance for specific workloads (see "IBM Smart Analytics

System family at a glance," at right). Two systems lead the charge: IBM pureScale Application System is tuned for transactional workloads, while IBM Smart Analytics System is optimized for analytic workloads.

IBM pureScale Application System combines IBM POWER7-based servers with WebSphere Application Server and DB2 pureScale software to handle heavy transactional workloads. The pureScale technology provides active-active services, creating a cluster-based, shared-disk architecture where each member has direct memory-based access to centralized locking and caching services. The pureScale Application System is designed to provide application transparency and continuous

availability, and grows with near-linear efficiency. As the amount of data increases, the system provides an economical and efficient approach to scaling capacity—without forcing organizations to overspend on excess hardware and software.

IBM Smart Analytics System is an integrated platform that provides broad analytics capabilities on a powerful warehouse foundation with IBM servers and storage. Deeply integrated and optimized for analytic workloads, IBM Smart Analytics System provides a single point of support for end-to-end analytics solutions. Plus, it can be deployed quickly: for most configurations, the system can be "analytics ready" in days.

 [MORE INFORMATION ibm.com/systems/smarter](http://ibm.com/systems/smarter)



## Sybase ASE 12.5 Compatibility Arrives for DB2

New features provide commonly used APIs and behaviors

Building on the support announced for Oracle database applications last year, IBM recently unveiled additional features for DB2 that are designed to make it easier for developers to move applications written for Sybase Adaptive Server Enterprise (ASE) to DB2.

The new DB2 SQL Skin for Sybase ASE makes it possible for applications using Sybase databases, T-SQL stored procedures, error handling, data types, and other capabilities unique to Sybase to run on DB2 with minimal application changes and without recompiling. In other words, DB2 can now behave as Sybase to applications written for Sybase.

The SQL Skin feature enables organizations to take advantage of the cost-saving aspects of DB2, like Deep Compression and autonomies, to yield savings that can then be invested in strategic initiatives such as analytics. Organizations such as BJC HealthCare and Pep Boys are already testing the feature to move their industry applications from Sybase ASE to DB2 to help reduce IT costs and achieve overall better performance.

This fully integrated capability was jointly developed with ANTs Software, an IBM Business Partner and provider of database migration solutions. The development began more than two years ago and early adopters have participated in a beta program since late 2009.

### MORE INFORMATION

[ibm.com/software/data/db2/linux-unix-windows/editions-features-sql-sybase.html](http://ibm.com/software/data/db2/linux-unix-windows/editions-features-sql-sybase.html)

## IBM Informix V11.50 Editions:

More choices,  
greater value

IBM recently announced a new suite of editions for Informix to deliver greater value to clients and partners. This suite includes a new Informix edition that is free to download, develop, and deploy, as well as a low-cost, feature-rich new edition for Mac OS X and Microsoft Windows environments.

The Informix Innovator-C Edition provides small-to-midsize businesses with common database functionality across a wide variety of platforms at no cost (redistribution requires a separate license). Comprehensive support is available as an optional purchase.

The Informix Choice Edition for Mac OS X and Windows is a feature-rich edition for developers, solution providers, and ISVs. With features such as High-availability Data Replication (HDR), Enterprise Replication, and clustering up to two nodes, the Choice Edition brings near-enterprise-class database capabilities to the midmarket, enabling customers and partners to develop solutions easily, quickly, and at low cost.

Read about the new editions of Informix and get links to download the free or trial editions at [ibm.com/informix](http://ibm.com/informix).

### IBM Smart Analytics System family at a glance

	5600	7600	9600
<b>Core software</b>	InfoSphere Warehouse Enterprise Edition; Cognos 8 BI	InfoSphere Warehouse Enterprise Edition; Cognos 8 BI	InfoSphere Warehouse Enterprise Edition; Cognos 8 BI
<b>Analytics options</b>	Query and reporting; multidimensional analysis (OLAP); data mining; text analytics	Query and reporting; multidimensional analysis (OLAP); data mining; text analytics	Query and reporting; multidimensional analysis (OLAP)
<b>Operating system</b>	SUSE Enterprise Linux	IBM AIX	IBM z/OS
<b>Server platform</b>	IBM System x3650 M2	IBM POWER 550	IBM System z10 technology
<b>Storage</b>	IBM System Storage DS3400; Fusion-io solid-state drive (optional)	IBM System Storage DS5300	IBM System Storage DS8700; Intelligence Disk Controller; solid-state drive (optional)
<b>Build, deploy, health check, and support service</b>	Yes	Yes	Yes



## From IBM Press

### *The Art of Enterprise Information Architecture: A Systems-Based Approach for Unlocking Business Insight*

Tomorrow's intelligent enterprises will bring together far more diverse sources of data, analyze it in more powerful ways, and deliver immediate insight to decision makers throughout the organization. Today, however, most companies fail to apply the information they already have and struggle with the complexity and costs of their existing information environments.

In this book, a team of IBM's leading information management experts guides you on a journey toward becoming an intelligent enterprise.

The authors present an information-centric approach to architecture and powerful new models that will benefit any organization. Using these strategies and models, companies can systematically unlock the business value of information by delivering actionable, real-time information in context to enable better decision making throughout the enterprise—from the shop floor to the top floor.

Buy now and get a special 35% discount on the book when using the coupon code IBM5717: [www.ibmpressbooks.com/bookstore/product.asp?isbn=0137035713](http://www.ibmpressbooks.com/bookstore/product.asp?isbn=0137035713)

## IDUG Unveils New Look and New Conference Focus

The International DB2 Users Group (IDUG) is taking several steps to reinforce its identity as the premier independent, DB2-focused, user-run organization driven by and for its members and volunteers. The first change is a new visual identity that celebrates the volunteers who make the group possible.

IDUG has also revisited its event lineup to ensure relevance to the greater DB2 community, and has renamed and expanded its upcoming IDUG Europe event to include attendees from more non-European countries. Under its new name—IDUG EMEA ([www.idug.org/emea](http://www.idug.org/emea))—the event will be held November 8–12, in Vienna, Austria, and will enable DB2 users in Europe, Africa, and the Middle East to come together as part of the wider community of DB2 users.

# Deploy Predictive Analytics Fast

IBM SPSS Decision Management software provides framework to enhance collaboration

Business users are clamoring for predictive analytics, but providing and supporting fast, effective, easy-to-use tools can be a serious challenge for data management professionals. IBM SPSS Decision Management is new software that enables business users to quickly build a predictive model within a configurable Web browser interface and to run simulations and what-if scenarios that compare and test the best business outcomes before the model is ever deployed.

The software also enhances collaboration between business users and professional analysts by providing one framework so they can work together seamlessly. Business users can easily design and build models based on their own business scenarios, and their expert

analysts can validate and refine those models to improve results.

Two industry-specific offerings are now available, and more are planned for future release. The first, IBM SPSS Decision Management for Claims, is designed for the insurance industry and helps identify fraudulent insurance claims, which account for approximately US\$30 billion in losses a year. The insurer can set up processes to identify claims that qualify for quick approval, as well as those that seem suspicious and require follow-up. This means insurers can resolve most legitimate claims in a single interaction and detect fraudulent cases at an earlier stage so that they can be routed for investigation.

The second, IBM SPSS Decision Management for Customer Interactions, is designed for organizations in all industries. It helps manage inbound interaction via call center, Web, point-of-sale, or e-mail, creating a view of the customer that can help a salesperson or service representative offer personalized recommendations or up-sell/cross-sell opportunities. For example, if a high-value retail banking customer dials the call center to complain, the new software may predict—based on the customer's data—that the individual is likely to leave. The information about the complaint, combined with the customer's history, can then be used to create a customized retention offer on the spot.

 **MORE INFORMATION**

[www.spss.com/software/deployment/decision-management](http://www.spss.com/software/deployment/decision-management)

## Coming to Grips with Big Data

IBM InfoSphere BigInsights provides tools for analyzing petabyte-class data volumes

Organizations and the Web are exploding with data, and business professionals want to analyze that data—both structured and unstructured, such as feeds, tweets, videos, and images—to gain actionable insight. But grappling with this “Big Data” can be an overwhelming prospect. To help organizations leverage and benefit from these massive sources of information, IBM created a new portfolio of services and solutions called IBM InfoSphere BigInsights.

Powered by Apache Hadoop, an open-source technology designed for analyzing large volumes of data, the BigInsights portfolio will help organizations analyze and visualize

petabyte-sized quantities of data. Consisting of specific analytics solutions that can be used by business professionals and be easily deployed in data center and cloud configurations, the portfolio includes the following:

- ▶ BigInsights Core is designed to help data professionals quickly get started with Big Data analytics, including design, installation, integration, and monitoring.
- ▶ BigSheets is designed to help business professionals extract, annotate, and visually uncover insights from vast amounts of information quickly and easily through a Web browser.

- ▶ Industry-specific solutions for the finance, risk management, and media and entertainment sectors assist with information analysis on a large scale.

Performing analytics on massive quantities of data enables people to spot patterns and uncover insights hidden in the information they have stored across their organization as well as data available on the Web. For example, using the BigInsights solutions, a financial analyst could combine historical market data trends with historical currency fluctuations and current news feeds to better predict risk exposure.

▶ **MORE INFORMATION**  
[ibm.com/software/data/infosphere/hadoop](http://ibm.com/software/data/infosphere/hadoop)

DATABASE MANAGEMENT

# Solving the DB2 Management Puzzle is Easy

## Quest Management Suite Simplifies DB2 Administration

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

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

Read our technical brief, *“The Top 10 Things a DBA Should Know About Toad® for DB2”* at [www.quest.com/puzzle](http://www.quest.com/puzzle)

 **QUEST SOFTWARE®**  
Smart Systems Management





**Stuart Litel**  
is president of the  
International  
Informix Users

Group (IIUG; [www.iiug.org/president](http://www.iiug.org/president)), CTO of Kazer Technologies ([www.kazer.com](http://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.

# Big News for Informix

## Exciting announcements unveiled at IOD EMEA 2010

Just as I thought, this year is turning out to be an exciting one for Informix. The five days I spent with other Informix users, partners, and IBM employees at the Information On Demand (IOD) Europe, Middle East, Africa (EMEA) Conference in Rome confirmed it.

IOD EMEA saw the announcement of some great news, starting with the most important pricing change to the Informix product line since IBM purchased the product line in 2001—in my opinion, of course. Now, a free Informix edition called Innovator-C is downloadable; not only is it for test and development use, but it also can be deployed for production use within your installation. To redistribute Innovator-C, you will need to obtain a license from IBM.

So what else is new? We now have two editions that replace the old Enterprise and Workgroup Editions of Informix Dynamic Server: Informix Ultimate Edition and Informix Growth Edition.

Back in the good old days (as my father used to say), Informix had a certain reputation: you buy it, you get everything. Well, the new Ultimate Edition is inspired by the good old days and includes everything (make that almost everything—compression being the one exception), with no restrictions.

How does Growth Edition stack up to the previous Workgroup version? Pretty darn well: Growth Edition includes more features at no extra charge. Plus, you can now deploy a three-node cluster of shared disk on the Growth Edition.

If you're an IBM Business Partner, the new pricing and packaging gives you the opportunity to go after the market in different ways, sell more, and offer a lower entry price to get new customers on board with Informix. For existing customers, you get the chance to start those additional projects or add to the ongoing ones in a simplified and less expensive way. Or maybe you can utilize features that cost extra before but are now included in the new packaging. Either way, the array of new versions opens the door for Informix and new opportunities.

If you want to download your no-charge copy or learn more about the new versions of Informix, you can do so at [ibm.com/informix](http://ibm.com/informix) or check out links for all the new packages at the IIUG Web site at [www.iiug.org/news/announcements/new\\_prices.php](http://www.iiug.org/news/announcements/new_prices.php).

Speaking of delighting the Informix community, the IBM Information On Demand 2010 Global Conference is scheduled for October 24–28 at Mandalay Bay Resort and Casino in Las Vegas, Nevada.

This event highlights the full breadth of the IBM Information Management portfolio, ranging from hundreds of technical sessions to more than 100 hands-on labs and much, much more, including the rumored announcement of the next release of Informix, code-named Panther.

IBM goes all out at this event, so try not to miss it! For more details about IOD, check out [ibm.com/software/data/conf](http://ibm.com/software/data/conf). See you all in Las Vegas, and be aware that a few Panthers will be running around Mandalay Bay.

I'm also proud to report that this past spring's IIUG Informix Conference was a major success. We surpassed the 2009 attendance and had more sponsors, more attendees (from more than 22 different countries), and more technical sessions than ever before. Among the presentations at the conference was the announcement of the 2010 IIUG Annual Director's Awards, which were presented to Thomas Girsch of Hilton Worldwide and Bruce Weed of IBM for their outstanding contributions to the Informix community. The event date and location are not yet set for the 2011 IIUG Informix Conference, but stay tuned to this column and the IIUG Web site ([www.iiug.org](http://www.iiug.org)) or the conference Web site ([www.iiug.org/conf](http://www.iiug.org/conf)) for more details. ✱

# Real-Life DB2 Superheroes

## IBM puts DB2 powers on display at the IDUG conference



**David Beulke** (*dave@davebeulke.com*) is president of Pragmatic Solutions, Inc. (PSI), a

training and consulting company that specializes in designing and improving SQL, application, and system performance on DB2 for Linux, UNIX, and Windows, and z/OS. He has experience in the architecture, design, and performance tuning of large data warehouses and OLTP solutions. He is also a former president of IDUG.

I am happy to report that the IDUG 2010 North America conference was a huge success, and I want to send a big thank-you to the leaders and volunteers of the organization for all their time and hard work. The conference featured a large number of DB2 users who represent organizations in vital industries: banks, stock markets, credit cards, and airlines, just to name a few. These companies continue to trust and build new applications on DB2 because of its performance, reliability, and ease of development—not because of a marketing campaign featuring a fantasy comic book character.

At the conference, DB2 users and IBM showcased the latest DB2 10 for z/OS technology, which helps reduce CPU usage by as much as 20 percent according to early IBM testing. With new features such as hash access, index include columns, an increased number of connections, security roles, enhancements to XML within the DB2 engine, and temporal tables, DB2 10 stands to not only save the DB2 community tremendous amounts of CPU processing but also development and overall support time.

For companies that have too many systems and not enough weekends to migrate all of them, DB2 10 offers the

opportunity to migrate directly from DB2 8 to DB2 10. This direct migration helps consolidate testing for all DB2 9 and DB2 10 features and can potentially shorten the overall migration time to DB2 10. The opportunity for skip version migration comes around only occasionally and can be a big help when you need to stay current with the latest DB2 capabilities.

DB2 9.7 for Linux, UNIX, and Windows (LUW) got some time in the spotlight, too, as users talked about the ease of migrating from Oracle to DB2, the data and index compression, and the high performance of the XML features. Companies are finding that they can easily migrate their systems from Oracle to DB2 in only days with the help of the IBM migration toolkit. Data and index compression are leading to elapsed-time savings within a company's application runtime and to dramatically shortened backup processes. And storage savings from compression are significant; from what I've heard working with clients and talking to others, the compression features can take databases from terabytes to gigabytes and can improve overall I/O application efficiency by as much as 80 percent.

XML functionality within the DB2 engine—in both DB2 9.7 for LUW and coming soon in DB2 10 for z/OS—provides great support for all types of unstructured data. Many companies are leveraging XML to interface with a variety of legacy and new service-oriented architecture (SOA) systems, and are using their existing IMS skills to build high-performance XML architectures. The DB2 engine excels at XML performance, and the ability to store anything in DB2 provides tremendous flexibility for supporting new mobile, spatial, Web, and other exotic applications.

Even if you didn't get to this year's conference, the great wealth of DB2 knowledge shared at the event (and now available on the Web) can answer your questions and help you address your company's challenges. Check out the IDUG Web site ([www.idug.org](http://www.idug.org)) to purchase proceedings from the 2010 conference as well as other North American and European conference information from 2009. There are a lot of great presentations, and when you subscribe to the IDUG Web site you get access to all of them. See you at the next event—IDUG EMEA 2010 in Vienna, Austria, November 8–12 (more conference information is available at [www.idug.org/emea](http://www.idug.org/emea)). \*

# Maintenance Makes Money



Five ways to cut database costs—just by paying attention

*By Sheryl M. Larsen*

**F**ew things do well without proper care. Houses, cars, pets, friends, and family all need our attention to stay happy and healthy, and database management systems are no different.

Without care and attention, a DBMS will break down, slow down, and eventually have to be shut down because it costs too much to operate. Keeping maintenance levels current to take advantage of changes in technology is one thing; applying the correct technology is another. This article points out five ways to reduce the cost of running DB2 by providing proper care to the system, using the right tool for the job at hand, and paying attention to detail.

## 1. Use the right SQL for the job—and make sure you know what it is

There are many ways to code a query that will yield the desired results; the trick is finding the method that plays to the strengths of your DBMS. That may seem obvious, but it's not as easy as it sounds because the tools and technology available to you change over time.

A few years ago, one of my clients inherited a customer relationship management (CRM) application that had been ported to its DB2 for z/OS Version 8 subsystem. One program was bringing the company's new IBM System z9 to its knees, consuming 25 percent of all CPU cycles on a daily basis. The program had 10 **CREATED** Global Temporary Tables, and most were embedded in recursive SQL statements. Six of the temporary tables could contain between 10,000 and 100,000 rows, depending on the depth of the sales chain, and were accessed repetitively.

But there is a second type of Global Temporary Table on DB2 for z/OS: **DECLARED**. And although **CREATED** temporary tables could not use any index technique until DB2 9, **DECLARED** tables could. So, I recommended converting six **CREATED** temporary tables to **DECLARED** temporary tables with clustering indexes. The remaining four **CREATED** temporary tables were appropriate for the job since they contained a small number of rows and were accessed only once. This change reduced the time required to complete each transaction from 1.78 seconds to 0.07 seconds, which doesn't look that impressive until you know that this program is executed



a minimum of 6,000 times per day. Therefore, the one small change saved the client 2.8 CPU hours every business day. The results were so dramatic that the senior DBA on site even joked about returning the new z9!

All of that occurred three years ago. Today, the client has DB2 9, which automatically generates sparse indexes on repetitively accessed work files, including temporary tables. But this sort of situation occurs all the time. SQL is a moving target: DB2 SQL had very humble beginnings as a query language when it started out in 1983—at that time it supported only inner joins, subqueries, the **GROUP BY** clause, the **HAVING** clause, the **ORDER BY** clause, and approximately 21 built-in functions. Now, 27 years later, the SQL available with DB2 9 for z/OS has exploded (see sidebar, “Look How You've Grown”).

Some SQL features are synonymous, but most provide unique functionality. And it is important to fully exploit new features as they become available, because many can increase performance and cut the cost of running your applications.

Searching static and dynamic SQL for keywords is an effective way to gauge how advanced your production SQL portfolio is. Figures 1 and 2 are SQL report cards prepared for clients by searching based on what was running in their current production environments. Client XYZ used very little SQL technology within its 3,000 financial reporting applications. This translates to almost all relational business rules and filtering performed in the application program, instead of taking advantage of high-performance SQL functions that could carry that load.

## LOOK HOW YOU'VE GROWN

**When DB2 SQL was first introduced, you could practically count the functions on one hand. These days, there's a lot more to choose from:**

TABLE EXPRESSIONS, COMPLEX CORRELATION, GLOBAL TEMPORARY TABLES, CASE, 100+ BUILT-IN FUNCTIONS, LIMITED FETCH, SCROLLABLE CURSORS, UNION EVERYWHERE, MIN/MAX SINGLE INDEX SUPPORT, SELF REFERENCING UPDATES WITH SUBQUERIES, SORT AVOIDANCE FOR ORDER BY, AND ROW EXPRESSIONS, 2M STATEMENT LENGTH, GROUP BY EXPRESSION, SEQUENCES, SCALAR FULLSELECT, MATERIALIZED QUERY TABLES, COMMON TABLE EXPRESSIONS, RECURSIVE SQL, CURRENT PACKAGE PATH, VOLATILE TABLE SUPPORT, STAR JOIN SPARSE INDEX, QUALIFIED COLUMN NAMES, MULTIPLE DISTINCT CLAUSES, IS NOT DISTINCT FROM, ON COMMIT DROP, TRANSPARENT ROWID COLUMN, GET DIAGNOSTICS, STAGE1 UNLIKE DATA TYPES, MULTI-ROW INSERT, MULTI-ROW FETCH, DYNAMIC SCROLLABLE CURSORS, MULTIPLE CCSIDS PER STATEMENT, ENHANCED UNICODE, AND PARALLEL SORT, TRUNCATE, DECIMAL FLOAT, VARBINARY, OPTIMISTIC LOCKING, FETCH CONTINUE, MERGE, CALL FROM TRIGGER, STATEMENT ISOLATION, FOR READ ONLY KEEP UPDATE LOCKS, SET CURRENT SCHEMA, CLIENT SPECIAL REGISTERS, LONG SQL OBJECT NAMES, SELECT FROM INSERT, UPDATE, DELETE, MERGE, INSTEAD OF TRIGGER, NATIVE SQL PROCEDURE LANGUAGE, BIGINT, FILE REFERENCE VARIABLES, XML, FETCH FIRST & ORDER BY IN SUBSELECT AND FULLSELECT, CASELESS COMPARISONS, INTERSECT, EXCEPT, NOT LOGGED TABLES

* Client XYZ's SQL report card	
Over 3,000 static financial reports	
V1	3 joins
V1	Zero subqueries
V4	Zero created global temp tables
V4	Zero table expressions
V5	Zero CASE expressions
V6	Zero declared temp tables
V7	Zero limited fetch

Figure 1: Client XYZ's SQL report card as of DB2 7 for z/OS.

* Client ABC's SQL report card	
V4	Zero created global temp tables
V4	3 useful table expressions
V5	158 CASE expressions
V6	Zero declared temp tables
V7	1613 limited fetch
V8	Zero SELECT INTO with ORDER BY
V8	26 common table expressions
V8	Zero recursive SQL
V8	Zero multi-row fetch

Figure 2: Client ABC's SQL report card as of DB2 8 for z/OS.


Client ABC's report card was much better but still lacked implementation of `SELECT INTO` with `ORDER BY` and multi-row fetch. These features typically reduce CPU demand by 30 percent and 50 percent, respectively. In thousands of occurrences and millions of executions, Client ABC was using obsolete single-row fetches instead of the current technology. Current SQL coding skills need to be cultivated in today's programming pool to transfer the technology into the production environment.

By moving to DB2 9, you get access to many new features, including:

- ▶ `TRUNCATE`
- ▶ `MERGE`
- ▶ `SELECT FROM UPDATE/DELETE/MERGE`
- ▶ `EXCEPT, INTERSECT`
- ▶ `RANK, DENSE_RANK, ROW_NUM`
- ▶ `ORDER BY / FETCH FIRST` in Subselect

Not taking advantage of these features, especially in a data warehousing environment, means that you need external tools to provide the functionality desired. This requirement drives up cost in terms of licensing fees and can have a significant impact on performance, since many of these tools require data to be transferred out of DB2 for external processing. (The searches used to gauge your SQL portfolio can also be used to determine whether a soon-to-be-purchased application package uses advanced features.)

## 2. Create optimal index design



Base table indexes are the first line of defense for performance-related issues. The most popular or critical transactions and services should have indexes to support all of the joining, filtering, and sequencing tasks required. Their omission can result in seriously delayed filtering or joining, all of which requires extra resources.

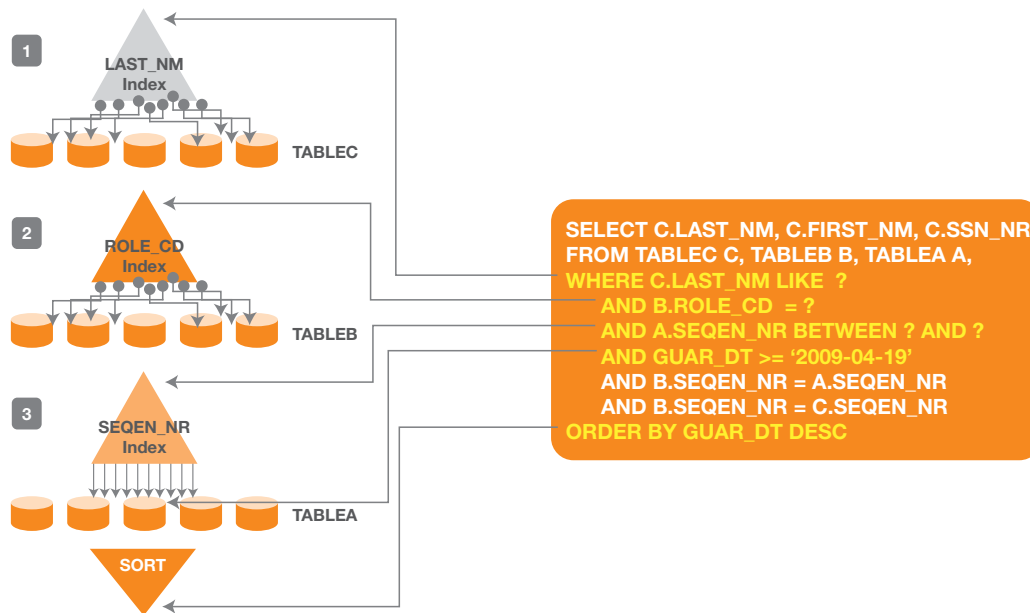
Consider the following scenario: A company had a Web page that its customers could access, and the company used the common tactic of executing a series of services when the customer logged in and storing the results so they could readily be displayed when the customer clicked on the associated tabs. This popular "P-Search" architecture let the customer quickly access commonly used functions, but such a strategy can be costly if the services being executed aren't indexed properly.

In this case, the P-Search service was consuming 9.9 CPU seconds of "in DB2" time on IBM System z10; during peak season, the customers waited between 10 and 40 seconds to see the Web page. The P-Search service contained a dynamic query that was responsible for the CPU consumption: a simple, three-table join on `TABLEA`, `TABLEB`, and `TABLEC`. All tables were partitioned and clustered by a single column `SEQEN_NR`; however, `SEQEN_NR` was not added to any other index. This forced the optimizer to choose between filtering and joining.

Figure 3 shows what happened:

**Step 1.** `TABLEC` is the first table accessed. DB2 Optimizer chooses to apply the local filter `LAST_NM`, but because `SEQEN_NR` is not included in the chosen index, DB2 Optimizer incurs random I/O to retrieve all `SEQEN_NRs` needed for the join to table `TABLEB`. Because the join is not performed using an





**Figure 3:** The first table accessed, **TABLEC**, applied the **LAST\_NM** filter on the non-clustering index, delaying **SEQEN\_NR** pickup. The second table accessed, **TABLEB**, applied the **ROLE\_CD** filter, delaying the join. The third table accessed, **TABLEA**, applied the join filter on the clustering index, delaying the **GUAR\_DT** filter.

index structure including **LAST\_NM** and **SEQEN\_NR**, it's delayed until data pages are retrieved. (See No. 1 in Figure 3.)

**Step 2.** The second table accessed, **TABLEB**, suffers from the same index design and applies the local filter **ROLE\_CD**, again incurring random I/O before applying the join predicate. Because the join was not performed using an index structure, it is also delayed until the data pages are retrieved. (See No. 2 in Figure 3.)

**Step 3.** **TABLEA**, the last table accessed, also suffers but this time uses the primary/clustering index for the join, uses sequential I/O for the data pages, applies the local filter **GUAR\_DT**, and then sorts for the **ORDER BY**. (See No. 3 in Figure 3.)

All three tables accessed had filtering or joining delayed until data pages were accessed. This is because DB2 Optimizer had to choose between joining or filtering, as no composite indexes satisfied both.

Solutions for this frequent, business-critical service consisted of the following:

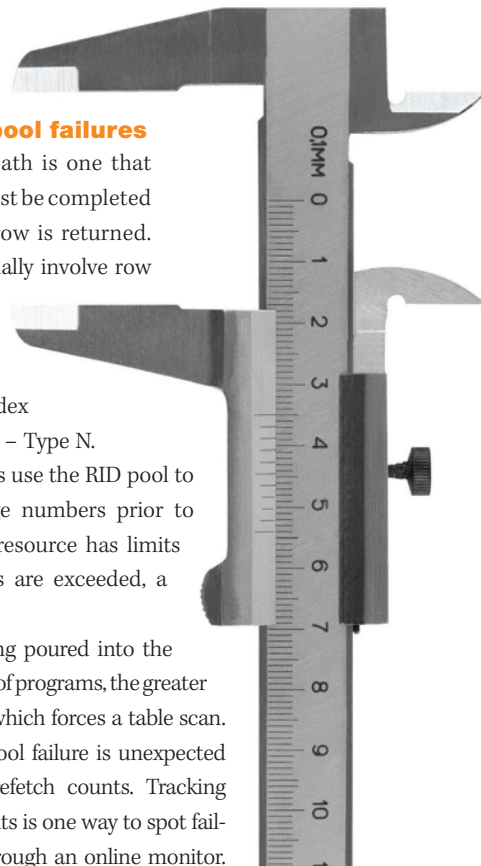
- ▶ Altering the **LAST\_NM.FIRST\_NM** index to add **SEQEN\_NR** and **SSN\_NR** to achieve index-only access to completely eliminate all random I/O to the data pages
- ▶ Altering the **ROLE\_CD** index to add **SEQEN\_NR** to combine the filtering and the joining
- ▶ Adding a fifth index on **TABLEA** of **GUAR\_DT.SEQEN\_NR** to combine the filtering and joining and eliminate the sort for the **ORDER BY**

The new index solutions brought the CPU time down to 0.02 seconds by moving all filtering, joining, and sequencing to composite index structures.

### 3. Monitor RID pool failures

An expensive access path is one that has many steps that must be completed before a single result row is returned. These access paths usually involve row identifier (RID) sorting to eliminate random I/O and are called List Prefetch, Multiple Index Access, and Hybrid Join – Type N. All of these access paths use the RID pool to pre-sequence data page numbers prior to access. This RID pool resource has limits and, when those limits are exceeded, a table scan occurs.

The more RIDs being poured into the RID pool by any number of programs, the greater the chance of a failure, which forces a table scan. One symptom of RID pool failure is unexpected spikes in sequential prefetch counts. Tracking sequential prefetch counts is one way to spot failures. A better way is through an online monitor.






Figure 4: A program with Hybrid Join – Type N access path gets 61 RID pool failures between 12 p.m. and 1 p.m.

Figure 4 shows the data from an online monitor of one program executing in a one-hour period using Hybrid Join – Type N that observed 61 RID pool failures.

One solution to the problem is to increase the size of the RID pool, but that may only temporarily suppress failures; worse, it may have no impact at all. A better solution is to identify queries that rely heavily on the RID pool by using DB2 traces or an online monitor. Each query needs to justify its RID pool use by fetching through to the end of the entire result set. Any queries not passing this test need to turn off heroic access paths that use the RID pool by adding either the `OPTIMIZE FOR n ROWS` or the `FETCH FIRST n ROWS ONLY` clause to the end of the query. DB2 10 will provide relief from these failures by utilizing external work files for the overflow, thus masking the need to analyze. However, it is best to find and resolve such issues before migrating to DB2 10, or you will have work file madness.

#### 4. Neutralize frozen or poorly performing SQL

Frozen SQL is SQL embedded in a purchased application that cannot be touched. The queries inside frozen SQL sometimes do not follow any performance guidelines and you are stuck with them. You need these queries to execute using as few resources as possible. For example, `WHERE DATE(col_TS) BETWEEN :date1 AND :date2` is a Stage 2/Residual condition. This predicate should be rewritten to `WHERE col_TS BETWEEN TIMESTAMP(:date1, '00:00:00') AND TIMESTAMP(:date2, '59:59:99')`, which is indexable Stage 1/Sargable. But you are forbidden to make that change.

Getting the work done ahead of time and storing the result on disc is one option. For DB2 8 sites, a solution is to create a materialized query table (MQT) that `SELECTS DATE(col_TS) as`

`NEW_COL FROM TABLE` and to create an index on `NEW_COL` column. Now all dynamic queries will be redirected to use the index. Another disc-based solution for DB2 9 sites is to use the feature called Index on Expression—`CREATE INDEX on DATE(col_TS)`. This will yield the same result with only one structure and can be used by both static and dynamic queries.

Both techniques are for queries that are ultra expensive to run, ultra complex to compute, or just plain ultra popular and may not be frozen. If your site is missing either type of structure, you are missing out on CPU demand reductions.

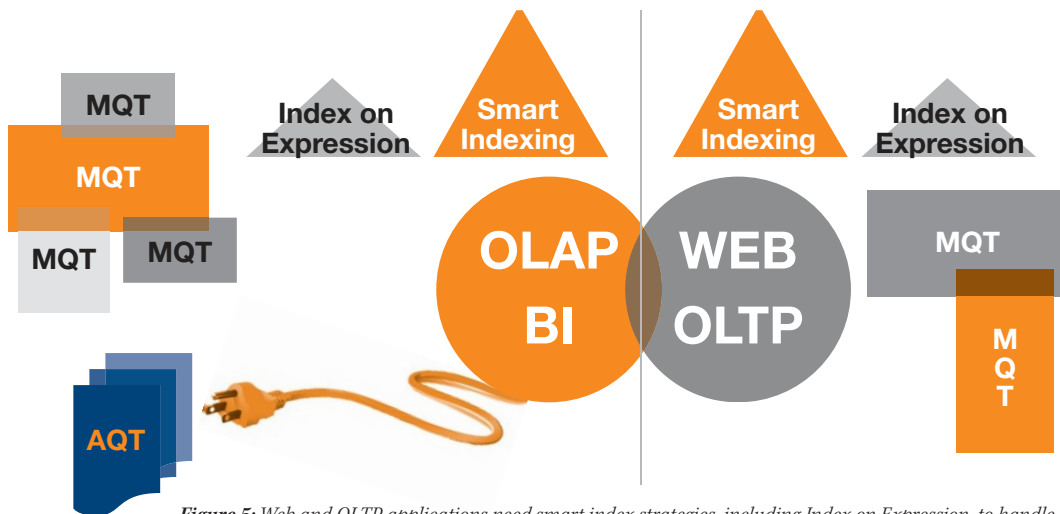
Another solution not yet available is memory based and can accelerate longer-running



queries by 5 to 10 times without requiring changes to the application, according to IBM lab

and customer beta experiences. Here's how:

Using IBM Data Studio (available at no charge), select a fact table and dimension tables to be pre-loaded in memory. This will copy and highly compress portions of your data warehouse into an all-memory, network-attached appliance, forming an accelerated query table (AQT). The appliance is called the IBM Smart Analytics Optimizer and is currently in beta. The DB2 9 Query Router used in redirecting queries from base tables/indexes to MQTs/indexes will be able to transparently redirect `INNER/LEFT JOIN` queries to use AQTs. This memory-based appliance will gain functionality over time to redirect a larger portion of queries for even more mind-bending query speed-up. Scalability requires a



**Figure 5:** Web and OLTP applications need smart index strategies, including Index on Expression, to handle large numbers of short-answer queries for probing and scrolling. OLAP and BI applications benefit from MQTs and AQTs to handle broad business questions that require vast amounts of data.

smart performance layer; MQTs, Index on Expression, and AQTs make up this performance layer.

## 5. Create the optimal schema design for your workload pattern

Scalability also requires a smart schema design that efficiently accommodates processing data in a workload pattern. For example, online transaction processing (OLTP) and intranet applications weather best with a third normal-form table design to mitigate update/delete anomalies and when heavily indexed for probing. This normalized design is well suited for static and focused business questions. Online analytical processing (OLAP) and business intelligence (BI) weather best with star schema and denormalization. The fact table and dimension table design are well suited for ad hoc and wide business questions that span vast amounts of data.

It makes sense that these environments would have disparate base table designs, index designs, MQTs, Indexes on Expressions, and in the future, AQTs, to take advantage of all DB2 has to offer for the workload pattern (see Figure 5).



### Show your database you care

Organizations need to be intimate with their DBMS of choice—attending to its proper care and paying attention to detail. To scale well, packaged DB2 applications also need to adhere to performance rules. By following simple rules—such as using the right SQL, indexes, access paths, MQTs, table design, and soon-to-come AQTs for the job—you can reduce the cost of data management. These rules require some thought and time beyond today's rapid application development cycle. But without this extra effort, your DBMS may start looking worse for wear. \*

*Sheryl M. Larsen* ([www.smlsql.com](http://www.smlsql.com)) is an internationally recognized researcher, consultant, and lecturer, specializing in DB2. She has more than 20 years' experience in DB2, is known for her extensive expertise in SQL, and performs detailed DB2 performance reviews for many clients. Larsen was voted an IBM Information Champion in 2009 and 2010 and the IDUG Speaker Hall of Fame in 2001, and has been a member of IBM's DB2 Gold Consultants program since 1994.

### RESOURCES

#### IBM Redbooks: DB2 9 for z/OS Technical Overview:

[www.redbooks.ibm.com/abstracts/SG247330.html?Open](http://www.redbooks.ibm.com/abstracts/SG247330.html?Open)

#### DB2 for z/OS technical resources:


[ibm.com/support/docview.wss?rs=64&uid=swg27011656#icc](http://ibm.com/support/docview.wss?rs=64&uid=swg27011656#icc)

#### DB2 9 for z/OS Information Center:

<http://publib.boulder.ibm.com/infocenter/dzichelp/v2r2/topic/com.ibm.db29.doc/db2prodhome.htm>

#### developerWorks: DB2 for z/OS Exchange:

[ibm.com/developerworks/exchange/dw\\_categoryView.jspa?categoryID=25](http://ibm.com/developerworks/exchange/dw_categoryView.jspa?categoryID=25)



How IBM's new database security and monitoring software helps protect sensitive information and reduce compliance costs

By Albert Cooley

**D**espite the noise you hear about data leakage through lost laptops, backup tapes, and unstructured data, databases are the primary target for external hackers and insider attacks. According to the *2009 Verizon Business Data Breach Investigations Report*, 75 percent of breached records originated in database servers; backup tapes, laptops, and workstations accounted for less than 1 percent of records breached.

It's easy to understand why: databases contain an organization's most valuable information, including customer records, payment card data, and financial results. Statistics show that hackers are skilled at using techniques such as cross-site scripting to penetrate perimeter defenses and reach the database. Existing security solutions, such as intrusion detection systems (IDSes), lack the knowledge of database protocols and structures required to detect inappropriate activities. Other solutions that rely on native DBMS logs, such as security information and event management (SIEM) systems, do not operate in real time, can be evaded by users with elevated privileges (which hackers often acquire), and introduce problematic overhead.

A growing number of mandates encompass this type of sensitive information as well, including various financial regulations (such as the Sarbanes-Oxley Act), industry-specific mandates (the Payment Card Industry Data Security Standard [PCI DSS]), and local data privacy laws. Each mandate has unique aspects, but they generally require organizations to detect, record, and remediate unauthorized access or changes to sensitive data, including those by privileged users, while providing a secure audit trail to validate compliance. Information security and database managers struggle to implement these types of controls, especially with respect to monitoring privileged users. Heightened focus on business-reputation risk and sensitive data protection is also driving closer internal scrutiny of controls. The result of all this is clear: providing effective database security and compliance has become anything but easy.

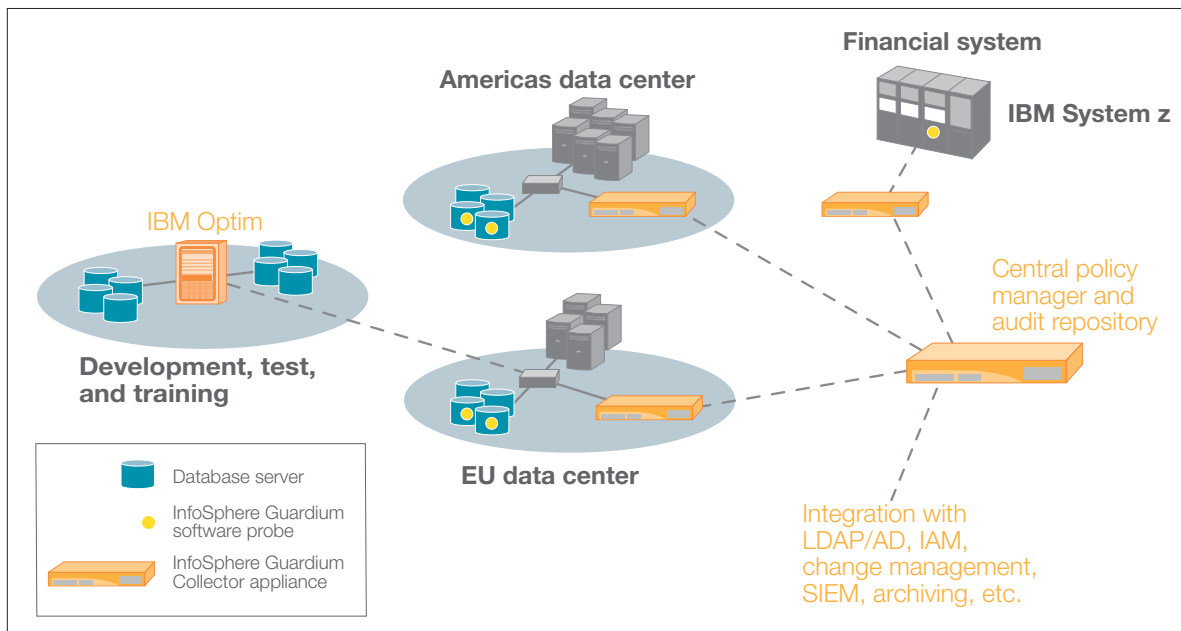
### The DIY approach

To date, meeting these needs typically involved developing custom solutions. The most common approach has been to enable the DBMS's audit facility to record transactions, then use UNIX scripts, Perl scripts, or C++ code to scrape and parse the logs and create a separate audit database or file. At regular intervals, individuals in the database administration organization must painstakingly review the data, identify policy violations, notify the parties responsible for investigating the violations, and record the investigation results.

This approach creates a variety of problems. First, developing, maintaining, and managing custom solutions often requires the equivalent of several highly skilled, full-time employees—the same personnel required for strategic business initiatives. Second, from a system-resource perspective, enabling the audit facility incurs overhead that is often unacceptable. And finally, auditors have begun to challenge the security of custom solutions: log files can be changed or deleted by hackers or insiders seeking to cover their tracks, and with a custom solution, security duties cannot be separated from administrative functions.

### Creating and managing granular policies

With the acquisition of Guardium, a leading supplier of database activity monitoring (DAM) and database protection solutions, IBM now offers an automated, effective, and efficient way to directly address growing database security and compliance challenges.



**Figure 1:** InfoSphere Guardium's scalable multi-tier architecture protects sensitive data through centralized aggregation of audit data and centralized, enterprise-wide management of security policies.

The IBM InfoSphere Guardium solution continuously monitors database transactions through lightweight software probes (see Figure 1) installed on the database servers. The probes monitor all database transactions, including those of privileged users, at the operating system kernel level without relying on database audit logs. The probes forward transactions to a hardened Collector appliance on the network, where they are compared to previously defined policies to detect violations. The system can respond with a variety of policy-based actions, including generating an alert and blocking the transaction in real time.

With InfoSphere Guardium, developing a policy to protect sensitive data is straightforward. For an example, consider the process of building a policy for payment card data; such data is frequently targeted by hackers, and it is subject to the very comprehensive and global PCI DSS standard.

Figure 2 illustrates a simple rule (see the sections outlined in red) specifying that any database transaction not from an “Authorized Client IP” address that attempts to access objects specified in our “Cardholder Objects” group will trigger an action to terminate the session. The Authorized Client IP group is created by entering the IP addresses of personnel with the business need to access cardholder data in a separate screen.



**Figure 2:** InfoSphere Guardium enables the development of granular policies to detect and block unauthorized database access or changes.

Rules can be made as detailed as necessary and can include a wide range of parameters—from user names and database type to command, time period, and counts. InfoSphere Guardium also provides integration with typical Lightweight Directory Access Protocol (LDAP) and directory services within your environment to automatically populate the groups in the security policy. As a result, you can enforce the access control portion of user management. InfoSphere Guardium allows you to develop the type of proactive controls specified by PCI DSS Requirements 10 (track and monitor access to cardholder data), 7 (restrict access to cardholder data), and 3 (protect stored cardholder data).

Note that in this example, the policy is database agnostic and can be applied uniformly across all databases supported, including IBM DB2, IBM Informix, Oracle, Microsoft SQL Server, Sybase, MySQL, and Teradata.

InfoSphere Guardium also supports IBM System z, a common platform in large enterprise environments. While IBM Resource Access Control Facility (RACF) is sometimes perceived as a sufficient control for mainframes, it does not capture a granular audit trail of what users did while accessing DB2 on z/OS, nor does it identify administrative users who abuse their privileges to view or change sensitive tables without a legitimate business need. IBM DB2 Audit Management Expert technology is used to capture mainframe database transactions and integrate them into a unified, enterprise-wide view via Collector appliances.

### How is it done?

Behind the scenes, the InfoSphere Guardium solution uses a series of technology strategies to support high levels of security while minimizing the impact on enterprise systems.

On the DBMS server, the solution must ensure that all traffic is captured—including direct connections to the DBMS server, such as SSH—and that the solution cannot be disabled without authorization. To meet this requirement, InfoSphere Guardium software probes are installed as a system account (root) and run as a single process on the operating system. The probes are installed on Microsoft Windows servers as a Windows service, and on UNIX servers as a daemon that, if killed, is restarted by the operating system. Also, because the probes are installed at the system software level, no changes to the database or applications are required, and the solution supports all major DBMS platforms.

Because databases are user-level programs that get services from the operating system, the probes can monitor all database transactions, whether they originate locally or remotely, by viewing these service requests. They can also support as many database instances as are installed on the server via a single probe process. This approach provides the separation of duties required by auditors, because only the security administrator will have the authority to disable the probe. As an additional security measure, there is a heartbeat between the appliance and the probe that will alert the administrator if the probe is disabled for any reason.

As the data is captured, it must be analyzed in real time to quickly identify policy violations. The probes transmit the data to the Collector appliances, which strip the data headers and analyze the content to determine how it relates to parameters specified by the policy. For example, the network layer will provide information on the client IP address, the OS user name will be extracted from the OS layer, the database user name will

be extracted from the database layer, and lastly, the payload will provide the actual SQL content. Relevant information is stored in a secure internal database so it can easily be used for enterprise-wide analytics, compliance reporting, and forensics.

Because transaction analysis is performed by the Collector and the server is only tasked with forwarding transactions, the performance impact on the server is very low—typically less than 5 percent.

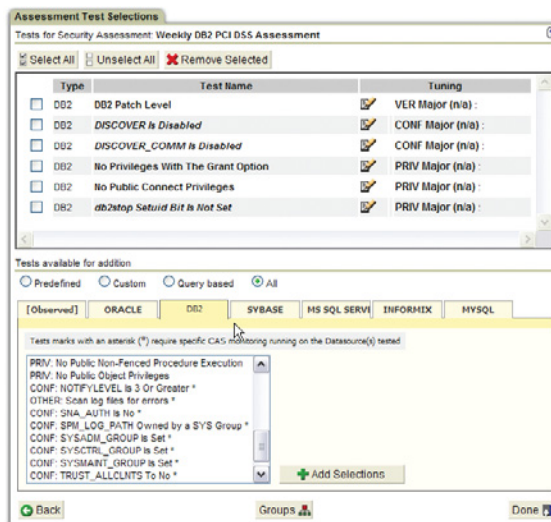
Data can be aggregated across Collectors, enabling centralized policy management, reporting and compliance workflow management—even across enterprises that have thousands of geographically dispersed databases. Clustering and high-availability architectures are often found in these types of large enterprise environments. To accommodate these architectures, probes can be configured to automatically fail over to other Collectors on the network, providing both high availability and load balancing. In clustered environments, probes are installed on all nodes, and each probe is configured as though that node is the primary node. All probes will be active, but only the primary node will forward monitored transactions.

### Beyond rules and policies

InfoSphere Guardium supports many other database security and compliance needs. Capabilities include discovering databases that may have been added outside normal processes by crawling the network; enforcing change control by matching actual database changes detected by the probes to authorized changes imported from a corporate change ticketing application; and detecting fraud by identifying actual end-user identities in transactions executed from enterprise applications such as SAP, Oracle E-Business Suite, and PeopleSoft that utilize pooled database connections.

The solution also provides vulnerability assessment and configuration-auditing capabilities that enable DBAs to directly evaluate the security of their databases. The former allows administrators to schedule tests—specific to their installed databases (see Figure 3)—to detect and prioritize (through a risk score) issues such as missing patches, weak passwords, misconfigured privileges, and inappropriate behavioral activity like credential sharing. The latter detects changes in critical system files and configuration objects. These are important capabilities called out in PCI DSS Requirements 2 (don't use vendor default passwords), 6 (maintain secure systems), and 11 (regularly test system security).

InfoSphere Guardium provides powerful monitoring, detection, and protection capabilities—but companies also need proof that all policy violations have been investigated and remediated. The Compliance Workflow Automation module satisfies that need, managing the regular distribution of compliance



**Figure 3:** The InfoSphere Guardium Vulnerability Assessment module scans the database infrastructure for missing patches, default passwords, misconfigured privileges, and other vulnerabilities.

reports to oversight teams, tracking the results of electronic sign-offs, and storing the results of this oversight process in the secure repository for review by auditors.

### Improve security, free up resources

Groups responsible for database administration should partner with other stakeholders—like the security team—to plan the InfoSphere Guardium implementation so they can help identify sensitive data, define appropriate protective policies, and design automated workflow to ensure incidents are investigated and closed in a manner that satisfies compliance requirements. Resources previously dedicated to the maintenance of custom solutions and the review of audit data can be redeployed to support the organization's strategic initiatives.

IBM has historically delivered strong security controls in its database solutions, including core authentication and authorization functionality, as well as advanced features such as Trusted Contexts. InfoSphere Guardium adds controls that will enable organizations using IBM database servers to significantly improve their security and compliance posture, while freeing up scarce technical resources. \*

*Albert Cooley is director of product marketing for Guardium, an IBM company.*

## RESOURCES

**DAM and InfoSphere Guardium:**  
[www.guardium.com/index.php/t1r](http://www.guardium.com/index.php/t1r)



# Utility PLAYERS

Data management technologies underpin energy manufacturers' smart grid efforts

*By Cameron Crotty*

Information and data will play a critical role in the future of the energy industry. Driven by demands for greater efficiency, more transparency, and reduced environmental impact, energy utilities are adopting new technologies and strategies faster than ever before, especially smart energy grids and the services they make possible: smart meters that monitor energy volume and time of use; monitoring applications that track and adjust electricity rates continuously based on availability and demand, or that alert customers when their usage spikes into a different rate level; and networked equipment, such as transformers and power lines, that automatically report their status and capacity.



All of those devices and systems generate a lot of data—and businesses in the industry need to figure out how to capture and manage it. “The energy and utilities industry has one of the greatest data challenges out there,” says Alan Webber, information technology analyst and partner at the Altimeter Group.

As a result, the standards that energy businesses use for communication are receiving some serious attention from data professionals. The International Electrotechnical Commission (IEC) Common Information Model (CIM) is the electric power industry standard that defines (among other things) the data architecture for information about electrical network configuration and status. And as the energy industry adapts to new conditions, that architecture is being poked, prodded, explored, and expanded.

#### Data architecture for an energy city

One group facing that challenge head-on is the IBM team in Mannheim, Germany, that created the data infrastructure to support the Model City Mannheim (MOMA), a project created by the German federal government as part of its E-Energy initiative to research the impact of information technology on the energy distribution grid.

The goal of the MOMA project was to create a smart grid that would provide a variety of services, including:

- ▶ Smart metering services for reading, validating, qualifying, and aggregating meter data to support centralized meter data management systems
- ▶ Energy management services for direct or indirect load control functions based on available load-balancing capabilities
- ▶ Forecasting services for predicting energy consumption
- ▶ Pricing services for calculating individual tariff-based prices for users

In addition to these new information-driven services, the MOMA smart grid needed to support many different participants. “Interoperability was always a critical part of our discussion,” says

Andreas Herdt, advisory IT architect for IBM. “The energy market is fragmenting into more and more businesses, from network operators to specialized energy generators, resellers, and others. They need to be able to exchange data.”

“The CIM data model moves pretty fast. It’s updated every 6 to 12 months, which is amazingly fast for the energy and utilities industry.”

—Andreas Herdt, *Advisory IT Architect, IBM*

The IBM team also wanted to create an architecture that would both mesh with the present while blazing a path for the future. “Our approach was to try not to build new models for the applications or the use cases that we needed to support,” says Christian Schiller, senior IT architect and managing consultant for IBM. “We wanted to work as much as possible with the CIM specifications that had already been accepted by the industry.”

#### To model or not to model?

The IBM data architects started planning by talking to all of the partners on the project and gathering as many different data concepts as possible. Some of the concepts didn’t always fit cleanly into the CIM data model definitions, and the team faced a constant temptation to solve the problem by defining a data concept in a specific application.

To help them make these critical decisions, the team laid down a set of criteria. “We had three rules of thumb,” says Herdt. “The first rule was to look at the model and find the concept. The second rule was to look carefully at the model and find the concept. And the third rule was to look *very carefully* at the model and find the concept!”

For the most part, the IBM team found objects in the model that could work with a little effort. “Sometimes we couldn’t achieve a perfect semantic match, especially at the attribute level, but we could usually get very close,” says Herdt. “It was



important to build on the existing CIM model as much as possible because we didn't want to re-create the wheel, and also because we want to build a foundation that other teams and organizations can work from."

### If a standard isn't defined, is it a standard?

However, the IBM team ran into challenges with the CIM, starting with documentation that was often ambiguous and sometimes nonexistent. Some of the model was documented at the object level and occasionally at the field level, but there was no documentation on specific use cases. In other instances, the documentation was wrong. "There was a situation where there was a master class for the model, and in a few occasions the documentation didn't correctly describe the object inheritance," says Schiller.

The fields or objects that the team was uncertain about were often central to the function of a given device. As a result, the team found itself embroiled in debates over the original intent of a given specification.

One of the most striking examples arose as the team attempted to build a data model for meter readings. The CIM data model offered two different ways to record meter readings: *reading* and *interval-reading*. The standard *reading* object appeared to have been designed for infrequent use and was fairly bulky. "Using the standard meter reading definition, you would create a container for every reading, which computes to 96 objects per day per meter if you read a meter every 15 minutes," says Herdt. Not a problem if you're recording one reading per month, but with smart meters, the team envisioned reading at intervals measured in minutes, which meant data volumes would mount quickly.

When the team dug into the actual model of the *interval-reading* object, they discovered that for it to work, the interval between every meter reading needed to be precisely the same. "We couldn't guarantee that the reading intervals would be identical because the underlying smart metering infrastructure did not offer readings exactly on the second," says Herdt. Instead, the team used the standard meter reading object, understanding that data management policies would be needed to help offset the impact of the high data volume.

### Searching but not finding

A few times, the IBM team couldn't find a concept that came close to matching their requirements, as in the case of weather forecasting. Weather dictates both how energy is consumed (air-conditioning in the summer, heating in the winter) and generated (such as wind farms, solar

facilities, and hydropower). However, the concept of weather was not even contemplated in the CIM model, so the team built their own data model to avoid misusing the semantic.

All of the work to address holes and document the model did not go to waste: the weather forecasting model (and the semantics supporting it) is one of many changes that the IBM team will propose back to the standards body that maintains the CIM. "The CIM data model moves pretty fast," says Herdt. "It's updated every 6 to 12 months, which is amazingly fast for the energy and utilities industry." That flexibility is a good thing, as the team can already see future needs. "The model covers smart metering now, but it doesn't cover all the different tariffs that will be needed," says Herdt. "We were able to create time- and volume-dependent tariffs, but eventually we'll need to make tariffs that depend on weather, energy availability, grid load, and other variables."

One of the team's ongoing challenges with the CIM involves finding ways to make the overall data model more specific. The model is not yet well-enough defined that separate organizations can simply follow it and feel confident that their data will match that of other organizations. "Right now, the model doesn't provide rules at the content level, which means that not everyone may use data containers in the same way," says Schiller. "For instance, a device serial number field could contain the manufacturer's serial number, an asset tag number issued by the organization that installs it, or something completely different."

### Lights, network, action!

The true test of the data model and semantics defined by the IBM team will come soon, as the partners involved in the E-Energy project begin to use the model for their own projects. "Over the next months, our model and code will face test data from real smart meters and the utilities' back-end systems," says Herdt. Later this year, Herdt and Schiller expect to see a full field trial that is completely based on the model they helped create. After that, who knows? But in any case, the energy expended now to establish the data infrastructure for smart grids will be a bright investment in the future—for data professionals and customers alike. \*

---

**Cameron Crotty** ([editor@tdagroup.com](mailto:editor@tdagroup.com)) is the editor of IBM Data Management *magazine*.

# Think Inside the Box

By Mirav Kapadia

Using Informix as an embedded database

Often, enterprise applications are separate from the databases that they rely on. But there is ongoing demand for database-driven, self-contained applications that work straight out of the box without additional configuration or administration.

In a server-based enterprise environment, professionals who have a deep understanding of database design and the dynamics of application behavior configure and maintain the databases. In many cases, however, this separation between application and database isn't desirable, or even possible. For example, ISVs must decide which versions of a DBMS their applications will support, knowing that each extra version adds development time, quality assurance cycles, and customer-support costs.

For some businesses, evaluating, maintaining, and supporting a stand-alone DBMS simply adds cost. What's the solution to these ills? Try an *embedded* database server. An embedded database server is integrated within the application, invisible to the end user, and requires either minimal or zero administration.

Working with a database server that will be part of an embedded deployment requires a different approach from a traditional enterprise deployment. In this article, we discuss some of the critical needs of an embedded environment and the Informix features that help developers meet those requirements. We also explain how to create a fully automated embedded Informix deployment using the new Deployment Utility available in Informix 11.50.xC6 and later versions.

## The unique challenges of embedded database deployments

Many of the characteristics that make a database server a good fit in a traditional enterprise environment also apply to an embedded database server. But the unique challenges of an embedded deployment change the priority—and sometimes even the meaning—of measures such as reliability, performance, and ease of deployment.

### Reliability and independence

The single most important quality of an embedded database is reliability: it cannot fail, ever. The database server is hidden; there is no database administrator, and no direct manual intervention is possible. Achieving this level of stability starts with robust database engine design (a hallmark of Informix), but that is usually just the beginning.

An embedded database server should also provide programmable and autonomic capabilities, so that the application can perform any necessary maintenance tasks. If a problem does arise, the application must be able to recover the database server and quickly return it to full operation.

Informix has an array of tools for enabling self-managing and self-tuning behavior. One of the most important—the Remote SQL Administration API—enables applications to programmatically issue SQL commands and perform many of the most common administration tasks necessary to keep the database running smoothly. For example, as the sizes of the databases supporting an application grow or shrink, the application can use the API to add storage space.

# Where is Informix?

With the help of more than 2,500 business partners, numerous Informix database server instances have been deployed in enterprises across many industries. Most of these deployments are “invisible”—the end user does not know that an Informix instance is running behind the scenes.

- If you have ever placed a cell phone call in North America, Europe, or Asia, it's likely that an application with an embedded Informix database validated your call information.
- Database applications in emergency call centers have relied on enterprise data replication features in Informix to ensure that critical systems are available 24x7.
- IP telephony call-processing systems use Informix as their embedded database server, as it can handle massive transaction volumes without sacrificing performance.
- It's highly likely that an Informix database verified your major credit card transactions or registered your purchase at a point-of-sale application in a retail store.
- Informix is the default embedded content store for IBM Cognos Express—an integrated business intelligence and planning solution for midsize companies.

Applications can use the Informix Scheduler API to automate stored procedures or user-defined routine executions at regular intervals or in response to specific triggers. Using the example of storage space, the embedded Informix instance could be configured to monitor disk usage and add more storage when any DBspace consumes 80 percent of its allocated space.

An embedded Informix database server can also provide the application with a comprehensive view of database activity through the sysmaster, sysadmin, syscdr, and sysusers system databases. The sysmaster database is particularly useful because its tables can show dynamic information about the state of the database server, such as monitoring buffer pool activity, sessions, and transactions. For example, an application can scan the sysmaster database for sessions that may indicate idle connections, terminate the sessions, and recover those processing resources.

## Performance and consistency

Performance is important for any database server, but in an embedded scenario, high performance must be accompanied by consistency and predictability. The database server should chug along continuously and evenly with the same performance characteristics as the organization's demands increase. The Informix database server can automatically adapt to changing workloads without degrading performance.

## Ease of deployment

Finally, deploying an embedded database server carries special challenges. Embedded database servers must often respect strict storage and memory limits. And just as the application it supports must be ready to operate out of the box, an embedded database server must require little or no setup. Any configuration must be fully automatic and completed by the application.

## Creating an automated embedded Informix deployment

Informix provides many features, such as Deployment Wizard and Deployment Utility (DU), to create and deploy a customizable footprint for optimizing storage requirements in embedded deployments. You can use Deployment Wizard to select the individual components that need to be installed on disk. DU was introduced in Informix 11.50.xC6; it is a new technology that can be used to rapidly automate deployments for preconfigured instances.

Consider this example: IBM Cognos Express uses DU to rapidly deploy a preconfigured and tuned Informix database server as its default content store. The end user is not aware of the Informix database server's existence, and the Cognos application handles the complete administration.

To create an embedded Informix deployment, you install and configure a source instance, take a snapshot of the product and database spaces, and finally, use DU to silently deploy the preconfigured instance on the target.

### Step 1

On the source system, install and configure the Informix database server instance. During the application development phase, you will use this database server instance to design, develop, and tune the databases.

To optimize the storage footprint, you can selectively install features from the Informix 11.50 product using the Deployment Wizard (see Figure 1). The application will dictate which features need to be installed for the product. For example, if the application doesn't intend to use On-Bar utilities, then that feature can be safely excluded from

the source instance. Selecting specific features through the Deployment Wizard is also supported during silent install.

The end result of this development phase: an application with an optimally tuned Informix instance that is configured for its life as an embedded database server. This instance will now serve as a template for deployment on the target.

### Step 2

On the source, compress the contents of \$INFORMIXDIR to an archive, such as `source_infxdir.tgz`. If you are working with an Informix installation on Microsoft Windows, then compress the contents of %INFORMIXDIR% to an archive, such as `source_infxdir.zip`. Also, compress the database spaces to another archive, such as `source_dbspaces.tgz` or `source_dbspaces.zip`.

To maximize reliability, we want to ensure that the database spaces are in a consistent state so their use can be resumed on the target. To achieve this, compress the database spaces after the instance has been blocked out or shut down. Currently, DU supports only database spaces that are configured as files in regular file systems (cooked files), which reside in a single directory.

### Step 3

On the target, copy the compressed archives, `ifxdeploy` executable, and `ifxdeploy.conf` from the source; extract `source_dbspaces.tgz`; update the configuration parameters in `ifxdeploy.conf`; and set the required environment variables.

The following parameters will need to be updated in `ifxdeploy.conf`:

- ▶ INFORMIXSERVER
- ▶ PROTOCOL1
- ▶ PORT1
- ▶ SERVERNUM
- ▶ INFORMIXDIR
- ▶ ONCONFIG
- ▶ SNAPSHOT
- ▶ INFORMIXPASSWORD
- ▶ ROOTPATH

The following environment variables will need to be set:

- ▶ INFORMIXDIR
- ▶ INFORMIXSERVER
- ▶ INFORMIXSQLHOSTS
- ▶ ONCONFIG
- ▶ PATH

### Step 4

On the target, execute the DU (`ifxdeploy`) to deploy the snapshot of the instance.

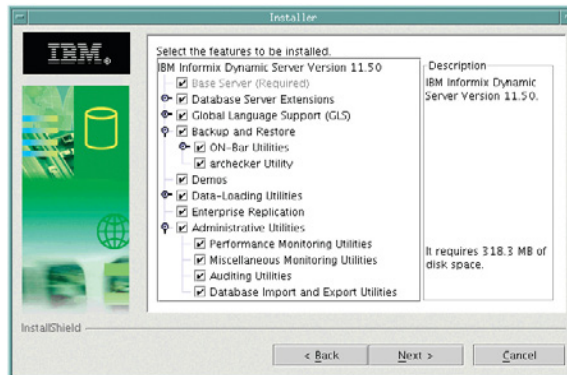


Figure 1: The Informix Deployment Wizard.

For example:

```
ifxdeploy -y -config ifxdeploy.conf
-f /mnt/hgfs/DEMO/source_infxdir.tgz
-rootpath /opt/IBM/informix/data/online_root
-relocate /data/IBM/informix/demo/server=/opt/IBM/Informix/data
```

You now should have a fully configured Informix database instance that is ready to run on the target.

### More to come

Informix is particularly well suited for use as an embedded database. Its proven technology qualifies it as an industrial-strength DBMS for an embedded database server on multiple platforms.

IBM continues to maintain a significant interest in enhancing the capabilities of Informix in various embedded environments, focusing on features such as deep embedding and integrated administration. Active research is in progress to make embedded deployments easier and to eliminate as much manual intervention as possible. \*

*Mirav Kapadia (mirav@us.ibm.com) is an advisory software engineer in the Informix Embed, Virtualization and Cloud Computing Team at IBM. He is an IBM Certified Database System Administrator for Informix V11.*

### RESOURCES

**Embedding Informix:** [ibm.com/informix/embed](http://ibm.com/informix/embed)

**Embedding Informix Dynamic Server: An introduction:**  
[www.redbooks.ibm.com/abstracts/sg247666.html](http://www.redbooks.ibm.com/abstracts/sg247666.html)

**Automate small footprint, embedded Informix Dynamic Server deployments:**  
[ibm.com/developerworks/data/tutorials/dm-1003idsemdatabilitytoolkit/index.html](http://ibm.com/developerworks/data/tutorials/dm-1003idsemdatabilitytoolkit/index.html)

**IBM virtual appliances:**  
[ibm.com/developerworks/wikis/display/im/ibm+virtual+appliances](http://ibm.com/developerworks/wikis/display/im/ibm+virtual+appliances)

# You Want to Govern What?



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

A database guy discovers there's more to data governance than meets the eye

I am a database guy (a DB2 guy, to be more precise). I do my work at the basement level of the IBM Information Management software stack. If the people who work with analytics, master data management, data integration, and other high-level software technologies are up on the bridge of the Ship of Data, peering through binoculars and wearing white jackets with gold epaulets, I'm down in the engine room. When the calls come in from the bridge, asking for more and more data to be beamed up with ever-greater speed, I call back with, "I canna' change the laws o' physics, Captain!" ... and then I find some way to get the job done. You folks up top keep talking about strategies and paradigms. I'll keep talking bits and bytes, speeds and feeds. Welcome to my world. Down here, we like to keep it real.

The latest in a long line of squishy-sounding concepts to come my way is "data governance." Excuse me, chief, but that sounds like a solution looking for a problem. I'll let you in on a

little secret: data is inert, which means it doesn't need to be "governed." It needs to be administered, and my friends and I have that situation well in hand, thank you. Go find something else to "govern."

What's that? You are, in fact, planning on governing something else when you talk about governing data? OK, I'm curious now. Tell me a little more. I'm listening.

## The focus is on people

Who am I listening to? None other than Steve Adler, director of Information Governance Solutions at IBM and chair of the IBM Information Governance Council. He's on the other end of the phone, telling me that the oxymoronic nature of the term "data governance" actually serves a useful purpose: it prompts people to ask questions. "What in the world do you mean by 'data governance'?" Boom. Conversation started, and a chance for Adler to say that the real aim is governing behavior.

People interact with databases: they direct data flows; they interpret and service data requests. People also make mistakes—often the result of shortcomings in the processes they are directed to follow and the application systems upon which they rely—that introduce errors in the information pipeline, creating downstream data quality problems. Data governance, says Adler, is largely about systematically identifying and addressing error entry points. The goal is to build data that decision makers can trust—and also provide them with evidence that the data can be trusted.

Now that's a pretty good pitch. But before joining your cause, I'd like to see Big Blue eat some of its own dog food. Adler asks if I know about IBM's product catalog. Sure I do—I only worked for the company for 17 years. Lots of information there—to the tune of 120 million records. Apparently, a high percentage of those records contained incorrect information: data that was wrong, missing, or published before it should have been. Out of 255 IBM product announcements chosen at random, only five were completely error-free. That's a motivating statistic.

To find out where the data quality problems originated, a team set "traps" at various points in the data flows that led to the product database. The traps helped the team discover several errors that could be expected to occur given a certain set of circumstances. With the causes identified, the team could then design and implement process- and technology-based solutions to eliminate the sources of information inaccuracy.

### Not my problem (or is it?)

I'm all for eliminating circumstances that lead to data quality problems, so you governance types go to it. I wish you success. No need to snoop around the databases that the DBAs and I are looking after—they're solid. How solid? At least 99.9 percent accurate. How many data records am I talking about? Well, a production database would probably hold north of a billion rows. Yeah, OK, even a very small error rate in a database of that size gives

# People interact with databases: they direct data flows; they interpret and service data requests. People also make mistakes that introduce errors in the information pipeline.

you a pretty big number. Point taken. Putting some of those traps in and around the database might be a good idea, after all.

But don't stop there, Adler says, because data quality problems are not just a matter of factual inaccuracies in data records. Sometimes, the pressing issue has to do with data categorization. Another story: An organization, partly due to acquisitions, ended up serving some of its larger corporate clients through multiple lines of business. One such client company complained that the same question, communicated to the service-providing organization through representatives of different lines of business, received different answers. It turned out that people employed in those different lines of business ascribed different meanings to the same terms. That's a data definition problem—the kind of problem that data governance aims to eliminate through effective master data management (MDM).

### How do you get from A to B?

I'll admit that I'm starting to see some value in this data governance stuff. It's not quite as squishy as I'd first thought. Still, seeing potential value and getting real value are two different things. How do you get some traction with a data governance effort? Where do you start? How do you keep things moving forward?

Adler told me that his favored approach involves six steps, to wit:

1. **Identify your goals.** Some goals will be of the sustainable variety (these will be ongoing), and others will be situational in nature (for example, dealing with a data quality problem). *Right—you need to know where you're going before you start on your way.*
2. **Identify what you're going to measure.** If you want to boost data quality, how will you know if you've made progress? Perhaps you'll examine some percentage of documents in a repository and record instances of factually incorrect or missing information. Deciding on a process and criteria for measurement is important for assessing the baseline situation (where's the pain?) as well as for tracking progress. *I get this. Talk is cheap: if you want to convince me that you're doing well, show me numbers.*
3. **Understand your organization's decision-making model.** Is it an autocratic model? Representative? Democratic? Whatever it is, is that the right model for your company? Does the data governance policy you're developing support the decision-making model? What would it mean if decision making were "better"? Would a larger volume of decisions be made? Or perhaps decisions would be made more quickly? *I suppose that as long as you're working on improving data quality, you might as well take the time*

to assess how data drives decisions in your environment. Maybe that'll be seen as something that needs improving, too.

4. **Communicate the data governance policy effectively.** How will information about the policy be communicated to stakeholders and other interested parties? Via e-mail? Through newsletters? *Probably wouldn't do to just rely on the ol' office grapevine to get the word out.*
5. **Measure results.** What has your data governance policy actually achieved? If you came up with a plan for measuring progress (item 2 in this list), cranking out the actual numbers shouldn't be too difficult. *Interpreting the numbers could be interesting.*
6. **Audit the whole thing.** Are processes being followed? Is technology being appropriately employed? Are controls effective? *The word "audit" kind of sets my teeth on edge, but a policy—whether focused on data governance or something else—isn't worth spit if it's ignored. I don't like being checked up on any more than the next guy, but I understand the need to do this sometimes.*

These are pretty high-level checkpoints, and a detailed project plan will have far more than six steps. But when faced with a complex task, framing the challenge in the right way can really help to focus

people's efforts. To me, Adler's approach looks to be a good framing of the problem space. It's something you could build on.

### OK, I'm in

A resident of one town once derisively said of another city, "There's no 'there' there." That's the way I used to think about data governance: interesting concept, but come on—where's the substance? Steve Adler makes a good case in arguing the business value of data governance. I do think now that there is some "there" there.

So if one of those high-up information management types visits you in the database engine room to talk data governance, give him or her a listen. Better yet, see what you can do to get involved in some way. When a grand plan is leavened with insights from people who have a rubber-meets-the-road perspective, the result is usually positive.

Gotta go. Maybe I'll see you around a coffee machine sometime. Who knows? We might even have a non-snarky talk about data governance. ✱

### RESOURCES

**IBM Information Governance:** [ibm.com/software/info/itsolutions/information-governance](http://ibm.com/software/info/itsolutions/information-governance)

# SIGN UP YOUR TEAM

FOR THE SAFARI BOOKS ONLINE OPEN HOUSE! – no strings attached!



**UNLIMITED ACCESS. UNLIMITED VALUE.**

Safari Books Online, Your Library in the Cloud.

Over 12,000 books, videos and certification guides <


Search titles from IBM Press, O'Reilly, Prentice Hall and more <

Read books before they are published <

**Safari**  
Books Online



# Creating and Using Partitioned Tables



## Roger E. Sanders

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

Special thanks to Liping Zhang, senior software engineer and technical lead for DB2 Data Partitioning Development at IBM, for providing feedback on the information presented in this article.

Improve query performance and achieve greater scalability

Table partitioning is a data organization scheme in which table data is divided across multiple storage objects called *data partitions*, based on values in one or more columns. A partitioned table can contain significantly more data than an ordinary table; yet, by taking advantage of a process known as *partition elimination*, queries run against partitioned tables typically execute faster and require less disk I/O than when they are run against nonpartitioned tables. (The DB2 Optimizer is data-partition aware, and only relevant data partitions are scanned during query execution.)

In this column, I'll show you how to create partitioned tables, and I'll show you

how a partitioned table's definition determines where individual records are stored.

### Range-partitioned tables

A data partition, also referred to as a *range* (currently DB2 supports only a range partitioning scheme), contains a subset of rows that are stored separately from other sets of rows in the table. Different data partitions can reside in different table spaces or in the same table space. The specifications provided in the `PARTITION BY` clause of the `CREATE TABLE` statement define how data in a table is partitioned. The syntax for this optional clause varies depending upon whether you want DB2 to generate each partition range evenly

within a specified table data range or you want to manually control the boundary for each range. The syntax for using automatic partition generation for a given table data range is:

```
PARTITION BY <RANGE>
  ([ColumnName] <NULLS LAST | NULLS FIRST> ,...)
  (
    STARTING <FROM> [Start | MINVALUE | MAXVALUE] |
    STARTING <FROM> ([Start | MINVALUE | MAXVALUE] ,...)
    <INCLUSIVE | EXCLUSIVE>
    ENDING <AT> [End | MINVALUE | MAXVALUE] |
    ENDING <AT> ([End | MINVALUE | MAXVALUE] ,...)
    <INCLUSIVE | EXCLUSIVE>
    EVERY <(>[Constant] <DurationLabel> <(>
    ,...)
```

The syntax for manually specifying partitions is:

```
PARTITION BY <RANGE>
  ([ColumnName] <NULLS LAST | NULLS FIRST> ,...)
  (
    <PARTITION [PartitionName]>
    STARTING <FROM> [Start | MINVALUE | MAXVALUE] |
    STARTING <FROM> ([Start | MINVALUE | MAXVALUE] ,...)
    <INCLUSIVE | EXCLUSIVE>
    ENDING <AT> [End | MINVALUE | MAXVALUE] |
    ENDING <AT> ([End | MINVALUE | MAXVALUE] ,...)
    <INCLUSIVE | EXCLUSIVE>
    <IN [TSName]>
    <INDEX IN [IndexTSName]>
    <LONG IN [LongTSName]>
    ,... )
```

where:

- ▶ **ColumnName:** Identifies one or more columns (up to 16) by name, whose values are used to determine which data partition a particular row is to be stored in. (The group of columns specified make up the partitioning key for the table—see sidebar, “Choosing a table partitioning key.”) No column with a data type that is a [LONG VARCHAR](#), [LONG VARGRAPHIC](#), [BLOB](#), [CLOB](#), [DBCLOB](#), [XML](#), a distinct type based on any of these data types, or a structured data type can be used as part of a data partitioning key.
- ▶ **PartitionName:** Identifies the unique name to be assigned to the data partition to be created.
- ▶ **Start:** Specifies the low end of the range for each data partition.
- ▶ **End:** Specifies the high end of the range for each data partition.
- ▶ **Constant:** Specifies the width of each data-partition range when the automatically generated form of the syntax is used. Data partitions

will be created starting at the [STARTING FROM](#) value and will contain this number of values in the range. This form of the syntax is supported only if the partitioning key is made up of a single column that has been assigned a numeric, date, time, or time stamp data type.

- ▶ **DurationLabel:** Identifies the duration that is associated with the *Constant* value specified if the partitioning key column has been assigned a date, time, or time stamp data type. The following values are valid for this parameter: [YEAR](#), [YEARS](#), [MONTH](#), [MONTHS](#), [DAY](#), [DAYS](#), [HOUR](#), [HOURS](#), [MINUTE](#), [MINUTES](#), [SECOND](#), [SECONDS](#), [MICROSECOND](#), and [MICROSECONDS](#).
- ▶ **TSName:** Identifies the table space in which each data partition is to be stored.
- ▶ **IndexTSName:** Identifies the table space in which the partitioned index for each data partition is to be stored.
- ▶ **LongTSName:** Identifies the table space in which the values of any long columns are to be stored.

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

Thus, if you wanted to create a table named SALES and partition it such that each quarter's data is stored in a different partition and each partition resides in a different table space, you could do so by executing a [CREATE TABLE](#) statement like this:

```
CREATE TABLE sales
  (sales_date    DATE,
   sales_amt     NUMERIC(5,2))
  IN tbsp0, tbsp1, tbsp2, tbsp3
  PARTITION BY RANGE (sales_date)
    (STARTING '1/1/2010' ENDING '12/31/2010'
     EVERY 3 MONTHS)
```

## CHOOSING A TABLE PARTITIONING KEY

Choosing an effective table partitioning key column is essential to taking advantage of the benefits of table partitioning. The most effective table partitioning key column(s) are columns that provide advantages in partition elimination. For example, if you normally query a table for records on the basis of date, partition the table on a date or time column.

Similarly, if you want to remove or archive portions of a table's data over time, partition the table according to the way you plan to archive records. For example, if you want to archive any data that is more than three years old, partition the table by week, month, or quarter so that you can roll out an old partition at the end of each week, month, or quarter, respectively.

Figure 1 illustrates how the resulting table would look.



**Figure 1:** Data stored in a simple partitioned table.

On the other hand, if you wanted to create a table named INVENTORY and partition it such that rows with item numbers that fall in the range of 1 to 100 are stored in one partition that resides in one table space, rows with numerical values that fall in the range of 101 to 200 are stored in another partition that resides in another table space, and so on, you could do so by executing an SQL statement that looks like this:

```
CREATE TABLE inventory
(item_no INT,
desc VARCHAR(20))
PARTITION BY (item_no NULLS FIRST)
(PARTITION PRODUCE STARTING MINVALUE ENDING 100 IN tbsp0,
PARTITION DAIRY STARTING 101 ENDING 200 IN tbsp1,
PARTITION BAKERY STARTING 201 ENDING 300 IN tbsp2,
PARTITION MEAT STARTING 301 ENDING 400 IN tbsp3)
```

In this case, rows with an ITEM\_NO value from 1 to 100 will be stored in a partition named PRODUCE (whose data is written to table space TBSP0), rows with an ITEM\_NO value from 101 to 200 will be stored in a partition named DAIRY (whose data is written to table space TBSP1), and so on; rows with a NULL value for ITEM\_NO will be stored in the PRODUCE partition.

It is important to note that when the NULLS FIRST option is specified, the first partition must start at MINVALUE. (Likewise, if the NULLS LAST option is used, the last partition must end with MAXVALUE.) Otherwise,

an attempt to insert a record with a NULL value for the partitioning key will result in a “data out of bounds” error. Additionally, each partition can have its data, index, and long data in separate table spaces. If you don’t specify an index or long table space, partitioned indexes and long column data will go in the same table space as the data.

By default, range boundaries are inclusive. To prevent specific records from being stored in a particular partition, you can create ranges using the EXCLUSIVE option. For example:

```
CREATE TABLE sales
(sales_date DATE,
sales_amt NUMERIC(5,2))
IN tbsp0, tbsp1, tbsp2, tbsp3
PARTITION BY RANGE (sales_date)
(STARTING '1/1/2010' ENDING '3/31/2010' EXCLUSIVE,
STARTING '3/31/2010' ENDING '6/30/2010' EXCLUSIVE,
STARTING '6/30/2010' ENDING '9/30/2010' EXCLUSIVE,
STARTING '9/30/2010' ENDING '12/31/2010')
```

In this example, a record with a sales date of 3/31/2010 won’t be stored in table space TBSP0. Instead, it will be stored in table space TBSP1.

When you insert a row into a partitioned table, it is automatically placed into the appropriate data partition based on its key value and the range it falls within. If the key value falls outside of all defined ranges for the table, the insert will fail and an error will be generated.

### Easy roll-in and roll-out of data

Another advantage of using partitioned tables is that new data can be easily added to a table (as another data partition) while old or obsolete data can be easily removed and archived. In my next column, I’ll walk you through the process of adding (rolling in) and removing (rolling out) data partitions. I’ll also show you how enhancements made in DB2 9.7 make the process of rolling data partitions in and out much faster and less disruptive. \*

## RESOURCES

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

**DB2 9 table partitioning:**  
[ibm.com/developerworks/edu/dm-dw-dm-0612read-i.html](http://ibm.com/developerworks/edu/dm-dw-dm-0612read-i.html)


**IBM Redbooks: Database partitioning, table partitioning, and MDC for DB2 9:** [www.redbooks.ibm.com/abstracts/SG247467.html?Open](http://www.redbooks.ibm.com/abstracts/SG247467.html?Open)

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

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

# Reducing Conversations with DB2 for z/OS: Part 3

Different ways of thinking about ordinary things



## Bonnie Baker

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

In my two previous columns, I wrote about unnecessary SQL and gave examples of SQL that is executed far too many times and SQL that should be replaced with newer, better-performing SQL. In this column, I will tell you about more situations where SQL can be deferred, combined, or avoided completely, all with the goal of reducing conversations with DB2.

### Three steps to remove unnecessary re-reads

Screen-display applications are notorious for unnecessary re-reads. We read to get an image to display; the user changes values on the screen and hits a function key to update; the program again takes control and re-reads the row (usually this time with a `CURSOR` using `FOR UPDATE OF`). After the `FETCH`, the program does “before and after” checking to see if anything on the row was changed by a different user while our user was staring at the screen display.

If the row is the same as the originally displayed row, our user's `UPDATE` is done. If any column on the row has changed, the appropriate information/error is provided to the user on the returned display.

The solution for this unnecessary SQL problem is to use an “optimistic locking” technique. We are confident that most, if not all, of the time, our row will not change during the

screen display. Our `UPDATE` glass is half full, not half empty. Therefore, we will do our `UPDATE` without re-reading the row. But we are neither foolish nor stupid; we will use a tried-and-true technique to ensure that our `UPDATE` is done to an unchanged row: we include in our `WHERE` clause one or more predicates that test for change. We have many options to use to test for change, some allowing more throughput than others. For example, we can:

1. Design our table to include a DB2-maintained `ROW UPDATE TIMESTAMP` (new in DB2 9), and include a predicate in our `WHERE` clause to check that column to see if its value is the same as the originally read column value.
2. Compare a user-maintained last-update timestamp to its original value (and reset that timestamp in our `SET` clause) to make sure that nothing on the row has changed since we received our image of the row.
3. Include `WHERE` clause predicates for every column on the row to see if each column value is the same as its original value.
4. Include `WHERE` clause predicates for only the columns that are important and relevant to us. We definitely want to make sure that the row still qualifies, and we may want to verify that columns we selected have not changed.

## Major changes to improve performance

Say you have an index on `TRANDATE`, `PROCESS-DATE`, and `CUSTNO`. You want to `FETCH` rows and write to a report all rows that will be deleted if they are: older than three months, inactive, and have a `CUSTNO` in a list found in your program's working storage.

You can:

```
Declare Cursor csr-delete-old-rows for
  Select ponbr, custno, trandate, process-date
  From bigtable
  Where trandate < current date - 3 months
  And status = :hv-inactive
```

```
Open csr-delete-old-rows
```

```
Fetch csr-delete-old-rows into
  :hvponbr, :hvcustno, :hvtrandate, :hvprocess-date
```

```
Write to a report
```

```
Delete from bigtable
  Where current of csr-delete-old-rows
```

If there are 3,000 index rows pointing to table rows that have a `TRANDATE` older than three months, but only 500 of those 3,000 table rows have a `STATUS` that is "inactive," then the worst-case scenario is that you will:

1. `CONNECT` to DB2 once for the `OPEN` to do `CURSOR` setup
2. `CONNECT` to DB2 for each `FETCH`. Then, for each of the 3,000 connects:
  - a. Do multiple `GET PAGE` requests for the pages in the multi-level index tree
  - b. Read an average of six index entries
  - c. For each of the six index `RIDs`, do a `GET PAGE` request to the table, followed by a potential synchronous wait on I/O
  - d. After the `GET PAGE` requests (and read I/Os) to the table, apply the `STATUS` predicate to the six table rows, reject five and accept one
3. Return a row to the program
4. Of the 500 qualified rows, only 25 of the `CUSTNO`s are in the working storage list (I call conditions that reject rows in your program after the return from DB2 "Stage 3 predicates"—and Stage 3 predicates are far worse than Stage 2 predicates)
5. `CONNECT` to DB2 again for each fully qualified row, this time to do the `DELETE` (25 connects)
6. Repeat 24 times for a grand total of 528 `CONNECT`s and way too many `GET PAGE`s to both the index and the table

Now what can we do to reduce the number of conversations in this situation? I suggest four improvements (you will remember the advice in step 2 from the last column: [www.ibmddmagazinedigital.com/dmmagazine/issue\\_2\\_2010/?folio=44#pg46](http://www.ibmddmagazinedigital.com/dmmagazine/issue_2_2010/?folio=44#pg46)):

1. If feasible, add `STATUS` to the index to make the qualification index-only
2. Change the `CURSOR` to use `ROWSET POSITIONING` to read more than one row at a time
3. Use `SELECT FROM DELETE` to see the rows
4. Get rid of the `STAGE 3` disqualification (the program check of `CUSTNO`) by `INSERT`ing the numbers into a `CREATED GLOBAL TEMPORARY TABLE` (e.g., `CTT_CUST`) and by using that temporary table in your SQL

```
Declare Cursor csr-select-and-delete-old-rows
  With rowset positioning
  For
  Select ponbr, custno, trandate, process-date
  From final_table
  (Delete from bigtable
  Where trandate < current date - 3 months
  And status = :hv-inactive
  And custno in (select custno from ctt_cust))
```

```
Open csr-select-and-delete-old-rows
(One CONNECT)
```

```
Fetch rowset from final_table
For 100 rows into
  :hvponbr-array, :hvcustno-array, :hvtrandate-array, :hvprocess-date-array
(One CONNECT with a +100 SQLCODE returned for row 26)
```

```
Close csr-select-and-delete-old-rows
(One CONNECT)
```

The new best-case scenario is that you will:

1. `CONNECT` to DB2 once for the `OPEN` to do `CURSOR` setup
2. `CONNECT` to DB2 for the one `FETCH` and
  - a. Do `GET PAGE` requests for the pages in the index tree
  - b. Read 25 index entries, fully qualifying the rows before reading the table
  - c. For each of the 25 `RIDs`, do a `GET PAGE` request to the table (if we are lucky and some of our rows share the same pages, we will do fewer `GET PAGE`s)
  - d. For each `GET PAGE` we may have to do a `READ I/O`

- e. As each row is found, **DELETE** it
  - f. Insert the **SELECT**ed columns of each row into final\_table
  - g. Return 25 rows to the program
3. Write to the report:

Close csr-select-and-delete-old-rows  
(One **CONNECT**)

Leveraging the preceeding four techniques reduced the conversations from 528 to 3 while lowering the number of table **GET PAGE** requests from 3,000 to at most 25.

### Eliminating Stage 3

I'd like to elaborate on the fourth suggested improvement (Get rid of the STAGE 3 disqualification). One of the most common coding missteps I see as a consultant is the practice of validating and rejecting rows in the program instead of in DB2.

These STAGE 3 predicates, especially those that reject a majority of the rows, are usually the result of old VSAM logic that did not utilize a **WHERE** clause. Back then, we used keyed reads and then applied program-coded **IF/THEN/ELSE** logic to the records. **IF/THEN/ELSE** logic

is fertile ground for figuring out how to use **WHERE** clause predicates so rows can be rejected sooner in DB2's address space.

### This is your conscience speaking

I wrote this three-part series to emphasize the importance of totally eliminating or deferring—until the very last moment—conversations with DB2. In the future, as you code each SQL statement, I hope a little Bonnie will be sitting on your shoulder, shaking my finger and asking, "Is this necessary? Can you do this read/maintenance to more than one row at a time? Can you defer this SQL to a later time, until you are absolutely sure that it needs to be done?" \*

### RESOURCES

**DB2 for z/OS:** [ibm.com/db2/zos](http://ibm.com/db2/zos)

**IBM Data Management magazine archive:**  
[ibm.com/developerworks/data/dmmag/archive.html](http://ibm.com/developerworks/data/dmmag/archive.html)

## Data Quality Tools for IBM



Ask about our 120-Day ROI Guarantee

Find the perfect data quality tools to fit your budget and business needs. Melissa Data offers easy-to-integrate APIs, Web services, and enterprise platforms to cleanse, verify, standardize, dedupe, and enrich your contact data.

- Verify addresses for over 240 countries
- Clean, validate and match data at point-of-entry or in batch
- Add lat/long coordinates with rooftop geocoding
- Work natively with Java, C, C++
- Power 32/64 AIX 5.2 or greater
- VisualAge or GCC 3.32

Request your free trial of one or more data quality tools at [MelissaData.com/myibm](http://MelissaData.com/myibm) or call 1-800-MELISSA (635-4772)

**MELISSA DATA**  
Your Partner in Data Quality

IBM, Power, AIX and VisualAge are trademarks or registered trademarks of International Business Machines. Other company, product and service names may be trademarks or service marks of others.

# **Will your z/OS batch jobs complete or capsize?**

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

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

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

**Restart made simple™**

**Download our White Paper:**

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

**For a free trial visit [www.relarc.com](http://www.relarc.com), or call +1 201 420 -0400**

**fai Relational  
Architects  
International**

DB2, WebSphere, z/OS and IMS are registered trademarks of IBM Corp.

# Building Fast Data Warehouse Schemas: Part 2

Optimize the design of your  
data warehouse dimension tables



## Lester Knutsen

([lester@advanceddatatools.com](mailto: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.

# M

My last column discussed why and how we use dimensional modeling to organize the data in a data warehouse. To briefly reiterate, the goal of a data warehouse is to give users fast access to data so they can make better business decisions, and to give them a separate area to run queries and to play with data that doesn't affect the operational database. Dimensional modeling helps us optimize our warehouse data for queries and analysis.

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

### Dimension tables: An introduction

A fact table contains things that you measure and record—typically numbers. A dimension table contains descriptive information about facts. A dimension table should contain no facts, only metadata. If a billing system fact table contains item quantity, amount paid, and amount due, its related dimension tables might include product information, customer information, and time or date information. Include as much descriptive information as possible! Useful dimensions have data that describes the who, when, where, what, and why of your facts.

### Dimension hierarchies

One of the most important design decisions you face is how to capture hierarchy information in a dimension table. Hierarchies are the logical structures that an organization creates to describe relationships among its data. For example, a product may belong to a product group, which then belongs to a brand, which then belongs to a division. A geographical hierarchy for a customer record may start at the street address and go up through the town, state, zip code, and country. Or, the dimension table may need to account for data from a geographic information system (GIS).

Business users often analyze data by comparing facts summarized at different levels of a hierarchy. Dimension table design is an essential part of enabling users to move easily through the data, drilling in and out to see what is going on at different levels.

Organizations often categorize their data in several ways, so you may need to account for multiple independent hierarchies. For example, if you have two ways of rolling up products, then you need to support both hierarchies in the dimension tables (see Table 1). I once worked with an organization that had independent offices that reported to regional offices. However, not every state



had regional offices. Because the organization needed to be able to analyze data by region and state, we had to support two hierarchies—states that had regional offices, and states that didn’t—in our dimension tables. There are no magic shortcuts to designing dimension hierarchies; it requires working with your business users and understanding their data. A whiteboard and a group meeting are the best tools for this process.

Not all dimensions will come from your source systems. You may purchase some dimension data, such as customer demographic data, from external suppliers, and some you may need to add manually. For example, most data warehouses have a time dimension, but this almost always must be created independently. For nearly every data warehouse system I build, I start with a spreadsheet and enter days for the next 10–20 years and the past 10–20 years, depending on how far back data will be loaded. I then add time periods such as the day of the week, holidays, fiscal periods, and seasons, building a hierarchy so people can slice and dice the data easily by time periods.

### Changing data in dimension tables

As you’re designing your dimension table, you need to plan for changing dimensional data. When a product name changes, do you show the name of the product when it was sold or the new product name? When a customer moves, do you show their address at the time they bought the product or their current address?

There are three established methods of dealing with changing dimensions. First, you can simply overwrite the data: your data will always be current, but you lose historical context.

The second method is to insert a new record for each change. This gives you a more accurate picture, but you must relate the old and the new records. One compromise method I have used is to insert a new record for a change only once a year or at specific time periods. Add a date field to both your dimension and fact tables when you do this, so you can see what they look like at the end of a chosen time period.

The third method is to add a field to the dimension table that captures old values. This method requires a schema alteration every time a record changes, because you’re going to be adding one new field for each change. To capture previous values, some sites add three or four fields when they build the dimension table. This field-adding method works best for data that’s only likely to change once or twice—such as a last name.

### Design tips and techniques

Here are a few things to remember when creating dimension tables:

- ▶ A dimension table needs a generalized key as the primary key data; this becomes the foreign key in the fact table. Informix serial numbers work really well as generalized keys, and I am a fan of using a serial number column on all dimension tables.
- ▶ You may have good descriptive information (such as notes or

* Product dimension fields	Comments
product_id	Primary key generated by data warehouse
product_code	Business users' primary key: product hierarchy level 1 and market hierarchy level 1
description	Description
manufacture	Metadata
version	Metadata
product_group	Product hierarchy level 2
brand	Product hierarchy level 3
division	Product hierarchy level 4
target_market	Market hierarchy level 2
market_sector	Market hierarchy level 3
market_sector_group	Market hierarchy level 4

**Table 1:** This dimension table defines two hierarchies: product and market.

comments) about a transaction that you do not want to lose, but that isn’t useful for analyzing the transaction data. Putting the descriptive information in the fact table will increase the table size and slow down queries. Instead, create a transaction dimension table with one record for each record in the fact table. That way, you can access the data and keep it out of the way.

- ▶ As you are building your data warehouse, think about which dimensions can be reused in other data warehouses. These are called conforming dimensions, and they are an early step to building an enterprise data warehouse.
- ▶ Generally speaking, a dimension table should be denormalized and flat, with as many columns as needed. Don’t worry about disk savings here: performance and ease of use are your main concerns.

Another design approach is a snowflake dimension, in which you normalize the dimensions to create dimension tables that link to other dimension tables. Most of the time, they just cause problems for people doing queries. One advantage of the snowflake dimension is the space savings, because you normalize your dimension tables. The disadvantage is that there are more tables to join in building a query, and users are less likely to use them. The simpler it is for end users to build a query, the more they will use your data warehouse, and the more business benefit it will provide.

### Next up: Avoid the data garbage dump

In my next column, we will look at summary tables. Summary tables are required in a data warehouse for fast query processing and to allow users to drill up and down into the facts. They can make the difference between a successful data warehouse and a “data garbage dump,” where gigabytes of data are stored but unusable. \*

# Smarter is...

Delivering social services  
where they're  
needed most



By Crystal Bedell

**M**uch of the time, data management work is all about schemas and SQL, or TCO and ROI. For caseworkers at California's Alameda County Social Services Agency (Alameda SSA), however, the ability to work with data effectively means more than just another percentage point on the revenue line. The agency delivers services and assistance that help people get back on their feet—a helping hand that's in demand now more than ever.

But with state and federal budgets pinched, qualifying for grant money is getting tougher. For example, under California's Deficit Reduction Act, the work participation rate (WPR) of welfare recipients must be at 50 percent; otherwise, the welfare agency may face stiff fines. And to qualify for a federal Title IV-E block grant, agencies must generate performance metrics on a daily basis to prove that the allocations are creating the desired effect.

Alameda SSA leadership knew that one of the best ways to support its caseworkers was with up-to-date, accurate, organized information—no simple task for an agency that has more than 2 million active records. “We needed to give our caseworkers direct access to information about their cases, at the individual case level. We needed to give the workers the right information, on the spot,” says Don R. Edwards, assistant agency director for Alameda SSA.

Because Alameda SSA caseworkers handle as many as 600 cases each, it was difficult for them to keep up with changes that needed immediate attention and action. Birthday notifications, for example, can be important: a child turning 18 could qualify his or her family for new benefits such as Medicaid or food stamps.

Working with IBM, Alameda SSA implemented a Social Services Integrated Reporting System (SSIRS). Based on IBM DB2, it combines IBM InfoSphere Identity Insight Solutions with IBM Cognos business intelligence (BI) software and an IBM InfoSphere Warehouse Enterprise Edition database delivered on a preconfigured IBM InfoSphere Balanced Warehouse D5000 to create a powerful data warehouse.

Alameda SSA chose DB2 in keeping with its philosophy of having all system components provided by one vendor—in this case, IBM. “This gives us better integration of software and services, and it reduces the complexity of the project,” says Edwards. “Plus, DB2 is robust software that enhances our backup and recovery processes.”



SSIRS enables Alameda SSA to generate the data it needs to demonstrate it is reaching its WPR goal and to secure additional funding. Before the project even began, the agency secured major funding from The Casey Foundation, a private philanthropic organization that was eager to see how effective a BI analytic data warehouse would be in improving the lives of the children in Alameda SSA's care.

The DB2-based system also helps reduce waste in welfare-to-work and child care programs; it can help Alameda SSA save up to US\$11 million and make it possible to direct funding and resources where they are needed most.

“[SSIRS is] helping us reach out to those who need us, making sure they don't fall through the cracks,” says Edwards. “It's about using our resources wisely and effectively to improve people's lives, better than we ever could before. This system is helping us in our fight against poverty in Alameda County.” \*

---

*Crystal Bedell is a freelance copywriter specializing in B2B technologies.*

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

Do you let the  
fox into your  
henhouse?

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

**Isn't that a conflict of interest?**

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

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



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



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



# Mission critical. Mission accomplished.

Introducing the new Intel® Xeon® processor 7500 series.  
Mainstream affordability meets mission-critical reliability.  
Goodbye server status quo. Learn more at [intel.com/inside](http://intel.com/inside).

Sponsors of Tomorrow.™ 

Copyright ©2010 Intel Corporation. Intel Xeon, Intel and the Intel logo are registered trademarks of Intel Corporation in the United States and other countries.