ISSUE 1, 2011 / VOLUME 16 NO. 1

IBM

data management

IBM.COM/DMMAGAZINE

KNOWLEDGE. PERFORMANCE. RESULTS.

INFORMIX 11.7
MORE FEATURES, MORE FASTER, MORE BETTER

DATABASE REVOLUTION
THE HISTORY OF A WORLD-CHANGING IDEA

DATA ARCHITECT
PERFORMANCE OPTIONS WITHOUT INDEXES

DOCTOR, DOCTOR:

Give Me The News

Data pros lead the way in healthcare

Smarter technology for a Smarter Planet:

What 3 million lines of code means to a piece of luggage.

It means Amsterdam Airport Schiphol will be able to accurately and efficiently move 70 million pieces of luggage per year—20 million more bags per year than they used to. The airport's automated baggage solution will allow them to increase their baggage handling capacity by 40%, so they can meet the growing demand placed on them as one of Europe's largest transport hubs. This system is built on IBM Rational® and Tivoli® software and runs on Power Systems™. A smarter business is built on smarter software, systems and services.

Let's build a smarter planet. ibm.com/luggage

A data visualization of the flow of baggage traffic at Amsterdam Airport Schiphol.

IBM.COM/DMMAGAZINE

IBM data management

KNOWLEDGE. PERFORMANCE. RESULTS.

ISSUE 1, 2011 / VOLUME 16 NO. 1

Features

16

Healthy Respect

Data professionals lead the way in healthcare

28

Organizing Principle

The database: A story of an idea that changed the world

22

Informix 11.7: The New Zoom

More features, more faster, more better

52

Smarter is...
Out of Chaos, Order

Departments

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

6 **NewsBytes**

10 **IIUG User View**
By Stuart Litel

11 **IDUG User View**
By David Beulke

12 **Data Manager**
DHL Mexico Follows the Money
By Merv Adrian

37 **Data Architect**
DB2 Indexes and Query Performance: Part 2
By Robert Catterall

41 **Distributed DBA**
How Well Do You Know the Rules?
By Roger E. Sanders

45 **Programmers Only**
New *ORDER BY* Information
By Bonnie Baker

49 **Informix DBA**
More Musings on Informix 11.7
By Lester Knutsen

52 **Smarter is...**
Out of Chaos, Order
By Lisa K. Stapleton

Sponsor Index

Applied Analytix, Inc. **51**
www.applied-analytix.com/installtm1toolbar

DBI **C3, 51**
www.dbisoftware.com/smarter

Fourth Millennium Technologies **9**
www.fmtusa.com

IBM **C2**
www.ibm.com/luggage

IBM Client Reference Program **40**
cusref@us.ibm.com

IBM Information On Demand **50**
www.ibm.com/events/InformationOnDemand

IBM Server and Storage **C4**
www.ibmdatamanagement.com

International DB2 Users Group (IDUG) **38**
www.idug.org

IIUG Informix Conference **48**
www.iiug.org/conf

Melissa Data **39, 50**
www.melissadata.com/myibm

Niteo Partners **51**
www.niteo.com/SmartPredictHC

PSS Systems **47**
www.pss-systems.com

Quest Software **7, 51**
www.quest.com/solve

Relational Architects International **21**
www.relarc.com

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

Safari Books Online **5**
www.safaribooksonline.com/ibmmag

Sogeti **51**
www.us.sogeti.com/bi

Sponsor Articles

Hardware Reliability Features That Maximize Uptime **14–15**
www.ibmdatamanagement.com

Migrating From Sybase ASE This Year? **27**
www.ibm.com/software/data/db2/linux-unix-windows/editions-features-sql-sybase.html

A Benchmark for the Real World **33–36**
www.ibm.com/software/data/db2/purescale

Optimize Both Performance & DB2 Memory Utilization With

Buffer Pool Tool® For DB2

Reduce Your Processing Costs, Save $$ Now!

RESPONSIVE SYSTEMS

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

IBM data management
KNOWLEDGE. PERFORMANCE. RESULTS.

EDITOR
Cameron Crotty
editor@tdagroup.com

MANAGING EDITOR
Stephanie S. McLoughlin

SENIOR EDITOR
Lisa K. Stapleton

ART DIRECTOR
Iva Frank

DESIGNERS
Lalaine Gagni, Margie Preston

CONTRIBUTING WRITERS
Merv Adrian, Bonnie Baker, David Beulke, Ives Brant, Robert Catterall, Art Kagel, Lester Knutsen, Stuart Litel, Carl Olofson, 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

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

tda GROUP

PRESIDENT
Paul Gustafson

VICE PRESIDENT, STRATEGY AND PROGRAMS
Nicole Sommerfeld

VICE PRESIDENT, EDITORIAL DIRECTOR
Debra McDonald

VICE PRESIDENT, CONTENT SERVICES
Paul Carlstrom

Printed in the U.S.A.

The weeks around the end of the old year and the start of the new year are a time to look back on what we've accomplished and think about what we want to achieve in the months ahead. In the spirit of the season, we spend this issue of the magazine looking forward and looking back.

Our cover story, "New World, New Work: Data Governance and the DBA," describes the evolution of data governance at BlueCross BlueShield of Tennessee. As companies realize that data and information are indispensible ingredients to a successful business, they're turning to data professionals for help in making both tactical and strategic decisions. In the years to come, data pros at all levels will find their roles changing, and this story provides a glimpse of what the future could bring.

At the other end of the spectrum, "The Database Revolution" is an absolutely fascinating piece by Carl Olofson that traces the history of database technology. In his day job, Carl is a research vice president at IDC who studies the database software markets; however, he's also a widely known and respected IT historian. We're thrilled at the opportunity to look into the past through his eyes, especially as IBM celebrates its centennial this year.

Those stories are just the start. Lester Knutsen and Art Kagel spent the latter part of 2010 tearing into Informix 11.7, and they've got a full report on what's new and how to use it. Merv Adrian has the latest from DHL Mexico and the company's successes with Informix Flexible Grid technology, and we also have some inside skinny on how new IBM acquisition Netezza might fit into your plans.

So while we might be a little late to wish you a "Happy New Year," here's hoping that the first few weeks of 2011 have already been exciting and prosperous. As always, we want to hear from you at editor@tdagroup.com.

Also, if you haven't seen the digital edition, go subscribe (for free, of course) at www.ibmdmmagazinedigital.com. It's a convenient and useful way to see the magazine.

Thanks for reading,

Cameron Crotty

LEARN TODAY. DEVELOP TODAY.
MOVE YOUR LIBRARY TO THE CLOUD

SIGN UP ›
FOR THE SAFARI BOOKS ONLINE OPEN HOUSE

Find all the latest and most relevant resources for IBM software developers and IT professionals at Safari Books Online.

LEARN MORE AT › safaribooksonline.com/ibmmag

UNLIMITED ACCESS. UNLIMITED VALUE.

Safari® Books Online

Sign up today for the **Safari Books Online Open House** and get your team or workgroup access to the world's most popular, fully searchable digital library. See why more than 15 million business and IT professionals, developers and web designers from corporations, government agencies and academic institutions access Safari Books Online for research, problem solving, just-in-time learning, professional development and certification training.

WHAT'S A NETEZZA?

Screamingly fast MPP and smart analytics boost enterprise BI

In September, IBM announced that it was acquiring Netezza for $1.7 billion as part of its strategy for giving customers as many options as possible to mine and exploit enterprise data. The speed and simplicity of deploying Netezza appliances such as Netezza's flagship TwinFin offering will help to accelerate the growth of IBM analytics among midsize businesses, while complementing and extending business intelligence (BI) capabilities in larger enterprises.

Because longtime IBM followers may still be somewhat unfamiliar with Netezza's data warehouse appliances, the following is a guide to Netezza's offerings, which include TwinFin. Designed to bring analytics directly into the hands of business users within every department of an organization, Netezza appliances are well suited to the needs of high-performance analytics. Built on IBM System x and often running alongside InfoSphere software, the TwinFin powers many analytics applications, such as Cognos and SPSS. The appliances require minimal administration and IT skills, so customers can usually run complex data queries within days of deploying the solution. The Netezza data warehouse and analytics appliance is valued by organizations in a wide variety of industries for its "Four S's"—speed, simplicity, scalability, and smart analytics.

Speed and simplicity

The TwinFin's massively parallel processing (MPP) architecture and huge memory capacity make it incredibly fast at analyzing huge volumes of data. TwinFin is built using IBM System x blades and IBM storage arrays—an industry-standard platform. The notable use of field-programmable gate array (FPGA) technology, plus the TwinFin hardware-assisted streaming architecture, deliver high price/performance ratios in an appliance platform.

More importantly, TwinFin is part of a recent trend toward more in-memory data processing. By keeping data in-memory and using sophisticated analytics—such as those written in Java, C++, Fortran, Python, MapReduce, Hadoop, SAS, and R—Netezza can attack business problems in time frames that used to be unthinkable.

TwinFin also comes with all its hardware, software, and storage integrated into one compact, powerful package. Many users report that they are able to install the appliance and finish loading and querying data in just a few days, leading to faster deployment and quicker ROI for advanced analytics.

Scalability

The TwinFin MPP architecture, combined with advanced, mature data warehouse software, enables TwinFin to run multiple enterprise applications in complex workload scenarios with predictable scalability. Rather than "throwing hardware at the problem" of scaling a system, TwinFin is designed in-balance from the ground up to provide scalable performance for data warehouse and analytics applications.

Smart analytics

Netezza i-Class software pushes advanced analytics deep into TwinFin's database. By supporting new data-analytic paradigms such as Hadoop and the R statistical programming language, TwinFin with i-Class makes high-performance, advanced analytics available to many more business users, without requiring them to have advanced degrees in statistics or computer science.

The smart money knows that data warehouse infrastructure is critical to a company maintaining a strategic edge in competitive marketplaces. Together, Netezza and IBM will continue to turbocharge enterprise data, helping deliver actionable insights to the entire organization.

MORE INFORMATION

www.netezza.com/data-warehouse-appliance-products/twinfin.aspx

www.netezza.com/areyouready

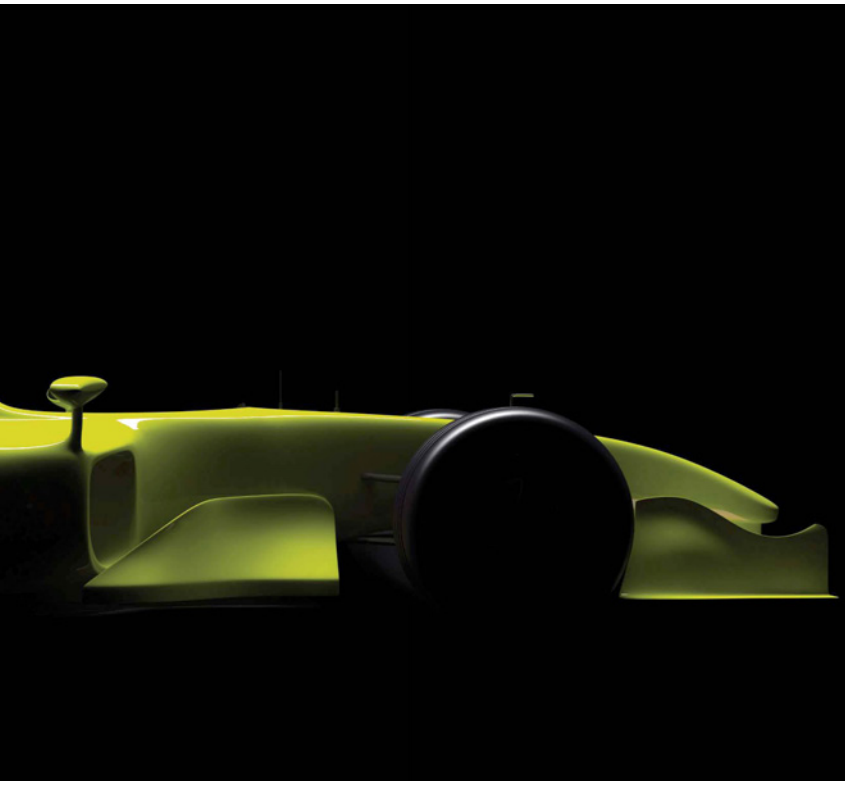

Partner News:

Triton Consulting and **DBI Software** Partner for DB2 Tools and Support

Triton Consulting, based in the U.K., and DBI Software, headquartered in Austin, Texas, are joining forces to help DB2 users on both sides of the Atlantic. The terms of the partnership cover three agreements: Triton will combine DBI's tools with its own Remote DBA service to enhance its proactive monitoring capabilities; DBI software will resell Triton's Remote DBA service in North America; and Triton Consulting will be authorized to sell the full range of DBI software solutions to the U.K. and EMEA markets. Observers say that combining DBI's tools for DB2 for Linux, UNIX, and Windows (LUW) with remote monitoring capabilities will help DBAs to get the useful information and alerts that DBI already provides, even from afar. Officials from both companies also emphasize that a wider distribution of each company's products will benefit the worldwide user community with better tools.

MORE INFORMATION
www.dbisoftware.com/press22.php

The New Era of Enterprise Business Intelligence: Using Analytics to Achieve a Global Competitive Advantage

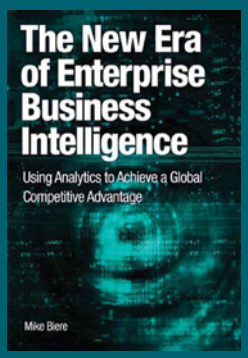

It's not easy to transform BI's potential into real business value. In *The New Era of Enterprise Business Intelligence*, top BI expert Mike Biere presents winning BI strategies and infrastructure, and describes how to maximize the value of information throughout the enterprise.

This product-independent guide brings together start-to-finish guidance and practical checklists for senior IT executives, planners, strategists, implementers, and business users. Drawing on thousands of hours working with enterprise customers, Biere helps decision makers choose among BI platform suites and appliances. He discusses how to evaluate BI solutions in the light of existing infrastructure, while avoiding technology biases and other hidden traps that can derail projects. Readers will also learn many skills critical to long-term success, including developing effective RFPs, cost-justifying the project, developing staff (and user) skills, and measuring the impact of the project on the business. Biere also draws on his extensive experience to examine coming developments in BI technology, giving readers a preview of what they're likely to see in the near future.

Use coupon code IBM5423 (case sensitive) and receive 35% off.

MORE INFORMATION
http://bit.ly/NewEraofEnterpriseBI

Support: Maintain Your Access to the IBM Electronic Support Portfolio

Staying active with subscription and support not only extends the value of your software with the latest product enhancements, it also ensures that you maintain access to our electronic support portfolio. Renew now and continue to get a customized support experience using our award-winning support portal.

MORE INFORMATION
http://bit.ly/SubscriptionAndSupport

Education: Complimentary Information Management Training Plans

One of the most important factors in determining an IT project's success is the team's skill level. As you plan and budget for 2011, make sure your team has the skills to drive your project's success.

IBM training plans help you budget accurately

To build an accurate budget, consider: Which team members need training? What kind of training? How soon? Once your needs are determined, IBM offers complimentary training plans to match.

Get started and request your complimentary training plan

Simply complete the online request form and Information Management Training experts will work with you to develop a custom training plan for your entire team that addresses the following:

- Training recommendations based on business roles, products, and project timelines
- Scheduling recommendations mapped to your project timelines
- Cost-saving strategies to stretch your training budget

MORE INFORMATION
ibm.com/training/plan

Test Drive a Self-Paced Virtual Class at No Charge

Self-paced virtual classes (SPVC) offer students maximum flexibility and control over their learning. The lack of fixed schedule or travel requirements means that students can study when and where they want. Plus, in today's business climate, self-paced virtual classes help organizations reduce their overall training costs.

Now students can experience this learning method through the SPVC Trials Offering. The following courses are available in a trial format:

- 1L610T – Advanced DB2 SAP for Administrators
- 2L250T – DB2 9 for LUW Multiple Partition Environment for Single Partition DBAs
- 2X101T – Developing Applications Using Informix 4GL
- DF040 – IBM FileNet P8 Prerequisite Skills 4.5
- 2X811T – Informix Dynamic Server 11 System Administration
- 1WA52T – InfoSphere Warehouse Components – SQL Warehousing Tool and Administration Console

Students can view detailed course content, instructor demos, and class exercises. Simply log in and experience an SPVC trial at no charge.

MORE INFORMATION
http://bit.ly/IBMSelfpacedVirtualClass

Analyst Reports:

Expert Analyses of IBM Smarter Systems, DB2, and Oracle Exadata

IT industry experts Richard Winter (Winter Corporation) and Philip Howard (Bloor Research) have delivered three in-depth studies of DB2, IBM Smart Analytics System, and IBM pureScale Application System. The Winter white paper looks at DB2 and data warehousing, while the Bloor studies show how IBM stacks up against Oracle.

Winter Corporation on warehousing with IBM

In the white paper *IBM DB2 9.7 and Smart Analytics System—Architecture and Key Capabilities for Data Warehousing*, Winter Corporation researchers examine how data warehouses are growing and changing; the architectural foundation of DB2 for Linux, UNIX, and Windows (LUW); DB2's capabilities for resource optimization and workload management; and IBM's ability to respond to data warehouse growth. The paper also describes IBM Smart Analytics System—a pre-integrated, modular, workload-optimized system for analytics—and concludes that "overall, DB2 9.7 and IBM's Smart Analytics Systems are among the most capable data warehouse products available on the market today."

MORE INFORMATION
http://bit.ly/IBMDB297andSmartAnalyticSystem

Bloor Research on IBM pureScale Application System and Oracle Exadata X2-2

Philip Howard of Bloor Research provides competitive analysis of these two OLTP systems for IT specialists. The report concludes that "IBM pureScale Application System out-competes Oracle Exadata in almost every area we have examined." Topics covered include scalability, performance, managing growth, administration and management, and costs.

MORE INFORMATION
http://bit.ly/hU9df0

Bloor Research on IBM Smart Analytics System and Oracle Exadata X2-2

Bloor Research also examines these two analytics systems, providing a comprehensive comparison for IT specialists. The report again finds that the IBM system out-competes Oracle in most areas. Areas examined are scale, performance, managing growth, administration, and costs.

MORE INFORMATION
http://bit.ly/ge4Lg9

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

Why Do You Love Informix?

Starting out the year with some tidbits, an introduction, and a request

The end of 2010 was a busy time for Informix, wasn't it? There were lots of new developments, including the release of a new version of Informix. Saying that might take some getting used to: IDS is no more; it's officially "Informix." In this case, it's Informix version 11.7, and it's packed with new features. You'll want to read more about this version at the International Informix Users Group (IIUG) site (www.iiug.org/news/announcements/ids_11.7.php). Also at the IIUG site is information on how IIUG members can download a free Developer Edition of the new release for many of the most popular platforms. [*Editor's note: See also the article "Hands-on First Look: Informix 11.7" in this issue.*]

I'd also like to introduce you to someone. Many in the Informix community already know Rob Thomas, the vice president of business development for IBM Information Management. He is personally overseeing all aspects of the Informix business around the world. Lots of old, incorrect rumors are still out there about Informix, and Rob is setting the world straight. You can read his updated letter to the Informix community on the IIUG Web site (www.iiug.org/news/letter/ibm_20oct10.html).

And, it's time for me to tell you about the upcoming fourth annual 2011 IIUG Informix Conference at the Overland Park Marriott (Kansas City area) May 15–18, 2011. This year's conference will feature over 80 technical sessions and more than eight half-day tutorials on the Sunday prior to the conference.

This is the largest gathering of Informix professionals anywhere, and vendors from all over the world will be there. At last year's event, we had attendees from four continents and 22 countries, including the famous Brazilian contingent (they definitely helped out the Kansas City economy). Don't wait until the last minute—each of the last two years we have nearly closed the registration due to space limitations. With a much better economy, we strongly anticipate outgrowing the hotel this year! For full information about the conference, visit http://iiug.org/conf.

One last thing: I am gearing up to write my future columns, and I need your help. See, the magazine editor always has to send me a friendly note every issue asking me where my column is. So I have an idea! Next issue I want to take up a theme that got started by accident a few weeks ago when I sent a worldwide e-mail blast to the IIUG members on "Why I use Informix." Rather than tell you all why I use it, let me ask you to tell me why you use Informix. If you read this column, please take two minutes, visit www.iiug.org/president, and fill in the form to send it to me. I promise not to use any names or company names unless you tell me I can. Please help me out and give me your thoughts, so I can get my column in on time! ✱

IDS is no more; it's officially "Informix." In this case, it's Informix version 11.7, and it's packed with new features.

Business Driving Technology or Vice Versa

Businesses are implementing DB2 10 quickly. Here's why.

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 the International DB2 Users Group (IDUG).*

Businesses that depend on mainframe systems don't make platform decisions lightly. So after talking with several people at the recent IBM Information On Demand 2010 and International DB2 Users Group (IDUG) conferences and hearing about not just interest in DB2 10, but also lots of early implementations, I can say that the tech and biz folks agree: DB2 10 is a winner.

The reasons I heard were common and clear, starting with the tough economic times. Businesses need to reduce costs now, not two years from now. DB2 10 can cut CPU and overall processing costs immediately, and businesses are pushing their tech staffs to implement it as soon as possible. In fact, some are skipping directly from DB2 8 to DB2 10.

Security and compliance are also prominent on every business's agenda. Executives are asking their security teams to control all data access by everyone in the organization—including the systems and DBA staff—down to the row and column level. For an executive suite that has security on its mind, the new role-based security model and label-based data access security layer in DB2 10 look compelling. DB2 10 provides comprehensive granular access controls, and the business can lock down its data and applications so only the proper personnel have access.

The new masking, encryption, and archiving facilities within DB2 10 help simplify regulatory compliance. Since they are built directly into the DB2 database engine, their performance beats any outside process, and DB2 10 provides standard procedures for implementing these new facilities on any old or new application database.

For plenty of companies, business intelligence is high on the priority list. The new DB2 10 temporal and history tables provide new data structures that deliver business or system time against any business question or data. These tables make it easier for a business to quickly understand its financial balance, order status, or business situation against any past, present, or future date. Many financial, insurance, and retail companies are defining their migration plans to leverage temporal data and gain competitive advantage.

All of this new DB2 10 information was just a small part of the IDUG EMEA 2010 conference in Vienna. The valuable user presentations, hallway conversations about business applications, and the keynote session with Jeff Jonas provided tremendous value to everyone. IBM has threaded the needle with DB2 10 cutting-edge technology that reduces CPU demand while providing better regulatory and business intelligence capabilities. Lots of companies think it's the perfect way to drive business demands faster. ✱

> The tech and biz folks agree: DB2 10 is a winner.

DHL Mexico Follows the Money

Informix Flexible Grid helps the global logistics company consolidate its IT infrastructure on its own terms

Merv Adrian, *now a research vice president with Gartner, completed this column while principal of IT Market Strategy, a research consultancy that analyzes software trends and advises leading IT firms on market strategy, the competitive landscape, and go-to-market execution issues.*

In the summer of 2011, IBM will celebrate the 10th anniversary of its acquisition of Informix, one of the early mega-investments the company made to expand its software portfolio. Skepticism abounded when the strategy appeared, but the ensuing years have seen customer retention, increasingly well-articulated market positioning, and continued development of innovative product features. Some, such as partitioning techniques, have made their way into other IBM database offerings. Others remain unique, and DHL Mexico, a longtime Informix customer, cites Flexible Grid as an example of IBM's continuing commitment to meet its needs.

Providing logistics services around the world is no small feat. For DHL Mexico, a local $370 million business that contributed $40 million in profit to its global enterprise parent during its past year, the required infrastructure is a complex, difficult-to-maintain patchwork. Tackling the issue of managing information technology cost-effectively is a priority. Felipe A. Villegas Moran, CIO of DHL Mexico, says the hardware and software combinations he must manage are numerous, and 10 percent to 15 percent of his annual budget goes to the upgrades required to keep everything interoperating smoothly. When as much as $10 million per year is dedicated to maintenance and the refreshment of systems, every financial advantage Villegas can find is crucial.

"We have huge needs. We are looking at doubling our capabilities not only in logistics, but also in information infrastructure. We move more than 75,000 shipments per day; that means more than 500,000 different checkpoints of information that we use to track every single shipment," Villegas says. Villegas is streamlining everything possible in the DHL Mexico IT infrastructure: simplifying applications, eliminating software platforms, and reducing the number of servers. DHL is already reducing its servers from 109 down to 80 in Mexico. Nonetheless, it has added enormous processing capacity because of the effectiveness of its new hardware and software. But there is much more to do—and software that is tightly wedded to its platform can be an impediment, limiting the operational and investment flexibility that is needed for optimization of the infrastructure.

"Many transnational companies are discovering this problem," Villegas says. "Five years ago, all of our units were much freer to do what they wanted for local autonomy; now we see the value of being global as well as local." DHL has a worldwide collection of applications—large and small—running on virtualized systems that include IBM, HP, and EMC hardware with IBM and VMware virtualization environments atop multiple operating systems—including Windows, Linux, and UNIX—that have varied programming environments as well. Above all this runs the application layer, and for DHL Mexico, much of it runs on IBM's Informix database.

Informix has long been a top choice for globally distributed enterprises that need to install it and leave it alone. All the DHL companies run Informix; the daily core applications, which operate across the company's 66 million checkpoints from Mexico to Malaysia, are based on Informix. DHL Mexico has many running copies of Informix across a broad variety of applications. Today, the organization operates several release versions on many hardware and software platforms, but Villegas is firmly set on a path to consolidate and simplify. An innovation in the most recent release of Informix will help DHL leverage its infrastructure investment while migrating and updating: Flexible Grid.

Flexible Grid for flexible consolidation

Introduced with Informix release 11.7, Flexible Grid allows organizations to create a grid—with two servers or thousands of them. The grid can contain a mix of hardware, operating systems (AIX, Linux, Solaris, and Windows), and versions of Informix. Informix offers centralized, simultaneous administration for all servers in the grid. It balances workloads across the grid. It practically eliminates planned downtime and reduces the number of DBAs needed. With one command from a single location, administrators can create new nodes and propagate the required Data Definition Language (DDL) across the grid to all the necessary locations. Each node is automatically configured for the hardware it is running on, including the number of cores. Backup and restore to and from the cloud are supported. For the many sites that DHL must operate—most without IT staff—these features will be invaluable.

> "The Flexible Grid capability will let us continue to use old equipment and add new equipment to scale out with different types of hardware and operating systems."
>
> —**Felipe A. Villegas Moran,** *CIO, DHL Mexico*

The administrative benefits are only part of the story. Optimizations in Informix 11.7 have dramatically improved query and transaction performance, and have leveraged physical changes to the POWER7 platform to make transparent improvement possible. But equipment has a useful life and a planned replacement schedule that is designed to make the best use of funds. "The Flexible Grid capability will let us continue to use old equipment and add new equipment to scale out with different types of hardware and operating systems. Pressure for systems convergence will drive platform convergence, and Informix Flexible Grid will enable that to happen much more easily—on our schedule," Villegas says.

Flexible Grid will permit the continuous use of the applications even as platforms and the operating systems on them change. It will not be necessary to upgrade or refresh all of them at the same time. Storage provisioning is automatic—meaning that intervention at that level often will not be required, creating dramatic administrative time savings.

Additional innovation

The Informix community is accustomed to the investment IBM has continued to make in the platform it relies on. From the free downloadable Developer Edition to the top-of-the-line Ultimate Edition, IBM has added numerous features to Informix in recent years. Enhanced high availability, encryption, and text search were added in 2007 and 2008. Data warehouse features, Cognos integration, and in-memory support (via solidDB) followed in 2009. And in 2010, with release 11.7, grid computing and new features for application embedding were on the menu.

Recent innovations leveraging the extensibility of Informix, such as the Spatial DataBlade for location-based data and the real-time feed capabilities enabled by time series data accompanied by the TimeSeries Real-Time Loader DataBlade, will likely add great value—especially for DHL's mission. IBM has laid out a road map for the next several years that includes support for Hadoop, multi-temperature data, and the extension of grid to non-database sources.

The infrastructure that Flexible Grid provides will make it possible to seamlessly exploit these features as DHL continues to build out its network. Doing it on your own terms, at your own pace, and within your financial constraints sounds like a win for DHL.

RESOURCES

IBM Informix 11.7: ibm.com/Informix

Informix Flexible Grid: ibm.com/software/data/informix/flexgrid

Server and Storage Focus

Hardware Reliability Features That Maximize Uptime

One hundred percent reliability is nearly impossible to attain, but that doesn't stop you from trying to get there. In this installment of Server and Storage Focus, we welcome back IBM experts Bob Zuber, worldwide marketing manager for System x Enterprise Servers, and Bina Hallman, director of entry and midrange disk storage, to walk us through features built into IBM systems that help you keep the business running no matter what.

How Can Reliability Be Built In?

Q: Customers are looking for systems that are reliable and designed to work day in and day out with minimal downtime. Let's face it: if the infrastructure isn't running, no one in the IT department is sleeping. How are you helping them get in a good eight hours?

Zuber: As x86 systems have become more powerful, they are being used for mission-critical applications and businesses are demanding high availability. As a result, many x86 enterprise servers now offer advanced reliability features such as hot-swap components and monitoring of system health. However, IBM® eX5-based systems go beyond one-off features and take an overall solution approach to availability that incorporates server hardware, systems management, and software.

For example, IBM System x® and IBM BladeCenter® servers combine the features of Predictive Failure Analysis (PFA), IBM Systems Director VMControl™, and VMware® vSphere™ vMotion™ to deliver a high-availability solution for virtualized environments. PFA gives key components in the IBM xSeries® range of servers the ability to monitor their own health and generate an alert up to 48 hours before failure actually occurs. This combination allows the data center manager to pre-program VMControl to automatically respond to specified PFAs by triggering vMotion to move virtualized machines from a compromised system to a stable host. This means that virtualized workloads stay up and running, even when the IT department isn't looking.

Another example of one IBM server-level reliability feature is Memory ProteXion™. This exclusive capability is integrated into our eX5-based servers and enables the system to stay up and running, even with as many as two failed DIMMs. Memory ProteXion works with off-the-shelf DIMMs, and the server just routes data around the failed chips automatically. As more mission-critical applications continue to run on these x86 enterprise servers, the demand for memory will increase and this type of superior reliability is just one way that IBM ensures clients will have access to their critical data.

Hallman: Reliability and uptime are very important aspects of any data center solution purchase but are often overlooked or not examined closely enough by the people making purchasing decisions. I think this occurs because—unlike performance, price, and capacity parameters—the information that enables a customer to evaluate the quality, reliability, and the expectation of uptime of a system can be difficult to obtain and to compare between systems from different vendors. Most people who purchase cars, washing machines, and many other consumer goods will seek out a wide array of reviews and opinions online and in magazines before making a purchasing decision. Even greater rigor should be applied when making a data center solution purchasing decision.

There are also ways to evaluate the potential reliability of a system without extensive testing. Start by looking at the fundamental design of a system. Has the supplier done its homework? Storage systems should have many built-in redundancy features to provide the highest level of uptime possible. Parts, components, and drives will almost certainly fail at some point, but proper design can minimize the likelihood that failures will cause downtime or data loss.

The IBM System Storage™ DS5000 series has many redundant design features, from redundant controllers and power supplies to redundant cable paths leading to the various hard drive expansion units. All drives with the DS5000 systems are dual-pathed for redundant communication to the storage controllers. The DS5000

series also has many RAID options, so you can choose the degree of protection—balanced against cost of implementation due to factors such as disk utilization—provided to reduce the risk of data loss due to disk failure. In addition to hardware-level redundancy, the DS5000 has a number of proactive failure alert monitors that notify the customer of potential pending component failures. These proactive failure alerts enable customers to plan component replacement prior to failure.

The DS5000 systems also have optional software features designed to further reduce the risk of data loss. For example, IBM FlashCopy®, VolumeCopy, and Enhanced Remote Mirroring features enable IBM DS5000 series users to set different levels of information protection for parts or all of the data stored on a system and also configure a remote storage site to support the user base in case an entire storage system is taken out of service by some unplanned event such as a natural disaster. This multi-level protection in the hardware and software design leads to very high confidence levels regarding data protection with the IBM DS5000 storage system.

Another consideration that you should take into account is the track record of both the supplier and the system it is delivering. IBM has a very good reputation for quality and standing behind its solution offerings. The IBM DS5000, for example, is the seventh generation of a storage architecture that has more than 110,000 systems deployed worldwide. More than 700 petabytes of capacity has been deployed on this architecture.[1] The DS5000 storage systems have a "five nines" or 99.999 percent level of reliability.[2] Now that should help our customers sleep a little better!

Physical Footprint and System Efficiency

Q: Data center space is at a premium, and organizations have become much more sensitive to system power and cooling costs. We have all seen significant improvements in the density of processors and hard drives. How does that translate into the density and efficiency of servers and storage?

IBM BladeCenter HX5

IBM System x3690 X5

IBM System x3850 X5

Hallman: Good question! We have designed a number of features into our offerings to deliver dense and efficient storage systems. The DS5000 systems have been designed to allow users to mix all supported drive types within the same EXP5000 expansion unit. For example, this means a customer looking to deploy a tiered IBM DB2® storage system can put solid-state drives (SSDs), Fibre Channel drives, and large-capacity SATA drives all in the same DS5000 expansion unit. This design ensures the highest possible density and utilization of the expansion units. An article titled "Unleashing the Value of Solid-State Drives for DB2 Workloads," by Sunil Kamath of IBM, in *IBM Data Management* magazine Issue 1, 2010, gave an excellent overview on how to take advantage of the high-performance SSDs within an IBM DB2 environment. A simple feature such as drive-intermix eliminates the need to buy a second expansion unit for different drive types, meaning that you only have to power one expansion unit instead of two—that alone can give you serious power savings. It makes you wonder why more manufacturers of storage systems don't have this built-in efficiency.

Another exciting, leading-edge storage offering from IBM is the EXP5060 expansion unit for the DS5000 storage system. The EXP5060 supports up to 60 high-capacity SATA drives in a 4U enclosure. Using the largest currently available SATA drives, a customer can get 120 terabytes of raw storage in 4U of rack space! Not only is it extremely dense storage, but it is also about 10 percent more power efficient than a comparable configuration of equal capacity. The IBM DS5000 storage system and associated expansion units deliver outstanding levels of density and efficiency to our customers.

Zuber: Today's new generation of eX5-based servers deliver increased processing power and memory capacity, enabling you to do more with less and consolidate to a smaller number of servers. IBM eX5-based systems are capable of achieving performance never before seen in the x86 space. Our four-socket x3850 X5 just achieved a TPC-C benchmark score of more than 2.3 million transactions per minute.[3] Compare this to the predecessor, the IBM x3950 M2, which achieved 1.2 million transactions per minute—but needed eight processors to do so.[4] The new x3850 X5 has nearly doubled transaction-processing performance with half the processors, which translates to a huge savings in terms of space and energy.

[1] Deployment figures as reported by IBM.

[2] Reliability figures as reported by internal IBM quality team.

[3] TPC-C Result Highlights: IBM System x3850 X5. www.tpc.org/tpcc/results/tpcc_result_detail.asp?id=110111601

[4] TPC-C Result Highlights: IBM System x3950 M2. www.tpc.org/tpcc/results/tpcc_result_detail.asp?id=108081902

Learn more

IBM Server and Storage Systems: www.ibmdatamanagement.com

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

NEW WORLD, NEW WORK:

Data Governance and the DBA

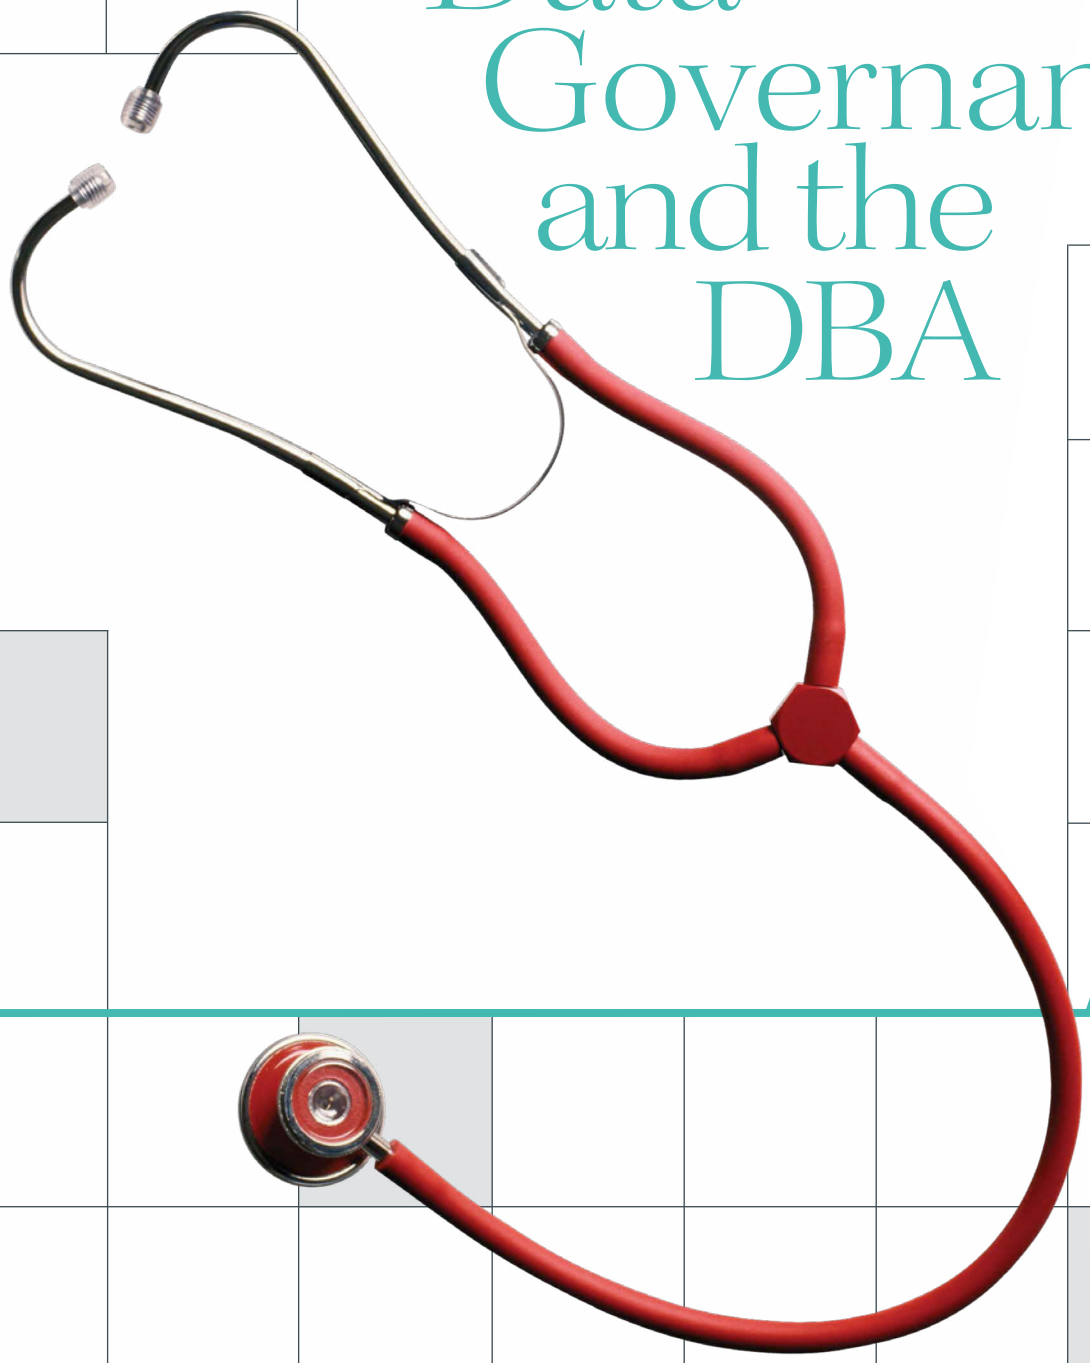

Healthcare and insurance industries offer a glimpse of the DBA's future role

By Ives Brant

Back in 1986, healthcare giant BlueCross BlueShield of Tennessee (BCBST) did something unorthodox. At the time, most large companies organizationally separated their DBAs who did physical database design from the data administrators who created logical design. "We invited DBAs and data analysts to the same table, to decide together on the logical and physical design of our corporate database environment," says Frank Brooks, director of data management and information delivery at BCBST. "We got into data governance without realizing it, just by defining commonsense practices related to our corporate data."

That decision signaled a change in how BCBST—which today has more than 3 million member-customers and 16,000 company customers—would govern data and do business in the future. "Data governance is integral to better healthcare delivery and reducing costs," says Brooks. "We currently use a large repository of data from our operational systems to help us stratify, or classify, members. Then we can steer the member into a support group for their asthma or cardiac condition. Reliable classification lets us guide them to better health practices, and helps BCBST manage costs."

The meeting back in 1986 also signaled a small but subtle change in the role that DBAs play at the company. Today, DBAs at BCBST still focus on changing data structures and adding new ones, but they are aware that SAS-70 rules and the Model Audit Rule give new importance to their work. The concepts behind data governance are changing business operations for organizations in healthcare and beyond. And as it turns out, the rise of data governance is also changing the roles and responsibilities of DBAs in important ways.

> "We invited DBAs and data analysts to the same table, to decide together on the logical and physical design of our corporate database environment."
>
> — Frank Brooks
> *Director of Data Management and Information Delivery*
> *BlueCross BlueShield of Tennessee*

The beginnings of data governance at BCBST

Just about everyone depends to some degree on the healthcare and insurance industries, and few are as essential to our well-being. They hold private, vital data on nearly all of us, millions of businesses, and thousands of healthcare providers. That data must be captured accurately, governed, protected, and held confidential. By law, it has to be retrievable upon request.

Healthcare and insurance have been leaders in data governance, though not exactly by choice. They are among the most regulated industries, because they handle so much sensitive and critical data. The best-known regulation of data: the Health Insurance Portability and Accountability Act (HIPAA), which requires healthcare providers to keep patient data confidential, yet accessible.

BCBST dove into data governance before the term existed. In 1994, a BCBST white paper led to the creation of a simple data warehouse architecture. Structured source data collected by operational systems—on patients, medical claims, and providers—was fed through extract, transform, load (ETL) processes into the data warehouse. From there, the data was delivered to business intelligence (BI) tools for analysis, and through more ETL filters to various datamarts.

By 1996, BCBST clearly understood that its data was a valuable resource and worked hard on ways to measure and improve its overall quality. One of the earliest metrics that the organization developed is still in use today: a metrics scorecard that tracks data quality over time. For example, BCBST tracks for what percentage of the company's 3 million members the database holds a Social Security number. That unique ID is the best way to identify each member. The scorecard keeps evolving, and users give their input on which data quality metrics they need most.

From there, says Brooks, "Our data warehouse architecture drove our org chart." In 1996, BCBST aligned IT staff with database architecture, segmenting three groups from IT: database administration to store and manage data, data integration for ETL processes, and BI/performance management to support analytics.

The organization's need for ongoing data quality, and to extract more value from data, drove how it approached its IT projects. BCBST established a process flow for data governance initiatives, from defining the problem and getting sponsorship, to evaluating current maturity levels, all the way to metrics and results (see Figure 1). Each project includes maintaining a metadata repository, creating a data repository, and designating a data steward. DBAs work on data structures as usual and "apply good data governance principles as they do their jobs," according to Brooks. Some projects focus on specific governance issues, and DBAs may be tasked to find creative solutions, such as customized security for the data pertaining to a particular customer.

The DBAs' understanding of what a piece of data means in the different departments and subsidiaries, as well as an awareness of semantic discrepancies, is valuable knowledge. Data that is captured correctly and that follows a uniform structure can help a liability insurer to reduce risks. For example, vehicles and other assets at a single site might belong to different corporations, but analyzing risk by location would alert the insurer to excessive exposure to a disaster striking that single site. A business glossary (that is, what does "location" mean?), master data management, and data validation can all play a role in recording the location of insured assets perfectly.

Project step	Tool applied
Define problem, obtain sponsorship, evaluate maturity	Data governance maturity questionnaire
Build road map, sample constructs, build data dictionary	InfoSphere Business Glossary
Understand data	InfoSphere Information Analyzer
Create data repository	InfoSphere Metadata Workbench
Define metrics	Cognos data governance scorecard
Appoint data stewards, manage data quality, implement master data management	InfoSphere QualityStage
Measure results	Cognos scorecard

Figure 1: *How BCBST carries out data governance projects: actions and tools*

Datamarts for specific groups of users

BCBST built a data governance application to balance and control updates to the data warehouse. The IT group runs about 1,000 processes monthly with this application, as data moves from BCBST's operational systems and external sources to the data warehouse, then to datamarts and BI analysis, building data cubes. As the data moves, it is checked intensively. Once in the data warehouse, the financial data is reconciled back to the general ledger. Then, to keep the financial datamart in sync with the general ledger, once data arrives there it is reconciled back to the data warehouse, down to the penny. "Our tools monitor everything that we do," notes Brooks.

BCBST's data-driven perspective has brought cost savings. "If you have the right information, and it's timely and stored right, then you can build anything on top—any logical process," says Brooks. "We developed a financial datamart, then showed Finance how to logically order and use their data elements for financial analysis." With its closed loop of common ETL processes and a robust BI infrastructure in place, BCBST enabled financial performance management by adding just one column to a single table. "It saved $5 million versus the approach recommended by outside consultants, and we delivered it faster," Brooks says. A coup for the home team, and an illustration of why data governance has a silver lining for DBAs. Uniformity and clarity, and alignment with the direction of the business, make it easier to complete new initiatives.

BCBST has found it effective to organize its people to address overall enterprise architecture organization in ways that reflect the teams working on its data governance (see Figure 2). Both draw upon executive, middle management, and practitioner staff. The relationship between IT and the business has evolved as well. Executives on a data governance committee provide oversight, setting policy and standards. They look for new ways to leverage BCBST data. Then, an information management committee decides how to align IT with the business priorities, and it balances demands from constituents. Finally, an execution group that includes IT and business users gets it done; this group includes the core integration, administration, and BI/analytics teams described earlier. In this latter group, DBAs contribute frequently on security to protect data confidentiality, as well as on backup and recovery considerations. "The DBAs serve an important role as our last line of defense in protecting our data assets," says Brooks.

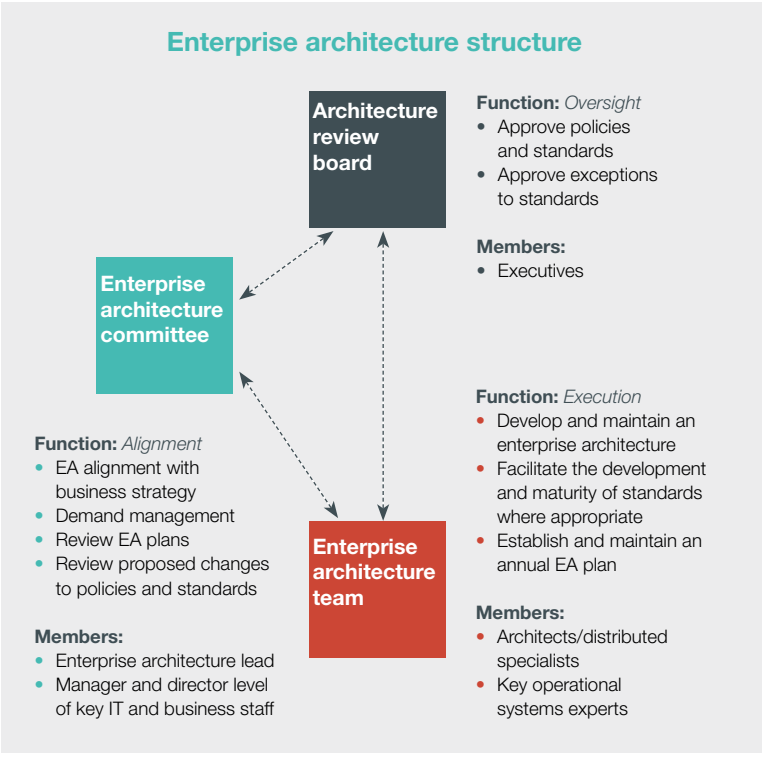

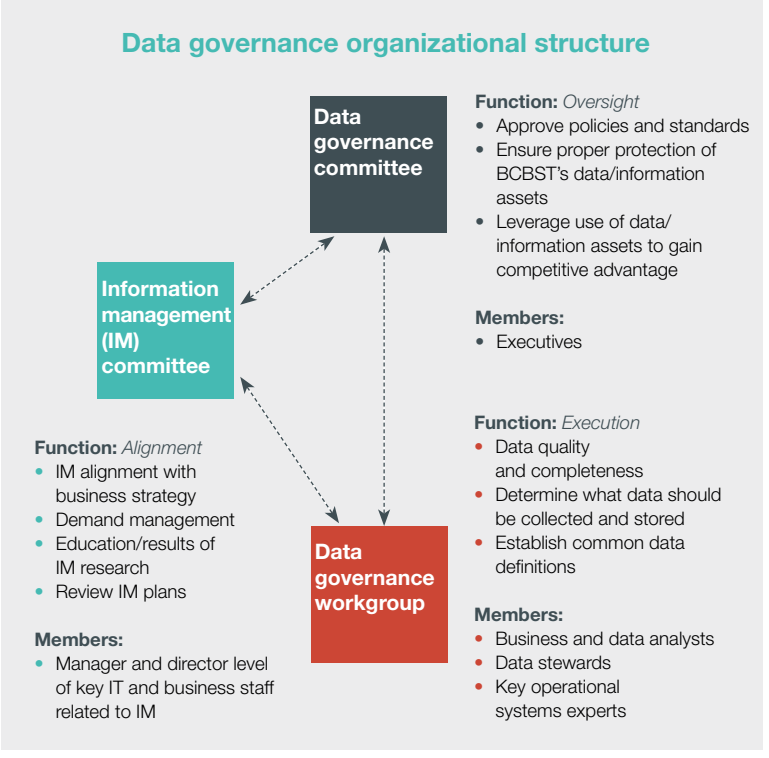

Figure 2: *BCBST aligns its people in similar ways to address data governance architecture and enterprise architecture effectively.*

Is the data ready to go out the front door?

Six years ago, BCBST started giving analytic access—in the form of IBM Cognos cubes—to its largest customers on the health statistics of their enrolled employees. "These customers are self-insured, so they are at risk. To improve the healthcare of their employees but limit costs at the same time, we enabled them to analyze their data 24x7 and store their reports," Brooks says.

> "We got into data governance without realizing it, just by defining common-sense practices related to our corporate data."
>
> — Frank Brooks
> *Director of Data Management and Information Delivery*
> *BlueCross BlueShield of Tennessee*

Security was a challenge that involved the DBAs. "We had to architect this and be careful how we loaded the cubes, so each customer can see only the information on its employees. We protect the data cubes, firewall them, and let each customer analyze their own," Brooks says.

"Insurers are becoming very involved in the healthcare of their members, on a close to real-time basis," Brooks continues. "The master data will give us a 360-degree view on everything about a member, account, or provider. The right data architecture lets us show all of it in near real time." That still requires a consistent, accurate combination of all data pertaining to a customer. Semantic discrepancies must be ironed out, data errors avoided, and duplicates eliminated.

DBAs at BCBST contribute to this evolution in four governance-related aspects. They provide effective logical and physical models, appropriate indexes, and ongoing tuning to ensure efficient loading and access. Second, they provide views and consulting to address ease of use and usefulness of the data. Next, they work with the data integration team and production application staff to ensure the overall integrity and quality of the data, so information derived from the data can be trusted. Finally, they are involved in applying records management policies related to archiving and purging data.

What's ahead: Text analytics and in-database predictive models

"In the future, we will add text analytics to gain new insights from our unstructured data, such as comments fields and certain fields embedded in electronic records. The results of text analysis combined with existing structured data will enable more powerful, accurate predictive models," says Brooks. BCBST is designing a new data mining and predictive analytics environment. "We already develop predictive models with SAS, but a more structured architecture will let us add in-database, near-real-time, predictive analytics."

Brooks sketched out one possible way to combine real data with what-if scenarios: "BCBST's patient-facing staff can gather enough data for us to offer the member a forecast—on the phone call—of his life span. 'Sir, you're 50 and based on your profile, you have an additional 12 years of expected life.' Then the kicker: 'Sir, would you like to know what changes in weight loss, exercise, and smoking will add 10 years to your expected life?'"

An architecture that reflects the imperatives of data governance

Concludes Frank Brooks, "To proactively engage with its customer base, a business needs data governance as a foundation of its IT structure." The overall database and data warehouse architecture—not just policies, roles, and tasks—should reflect the imperatives of data governance. With that consistency, IT can confidently move to a data hub environment like BCBST, and can have the ability to serve up analytics to major customers and individual history and predictions to members, give near-real-time performance management, and capture and leverage both structured and unstructured information. And as the story at BCBST shows, DBAs bring knowledge critical to the practical implementation of data governance initiatives.

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

Will your z/OS batch jobs complete or capsize?

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

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

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

Restart made simple™

Download our White Paper:

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

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

Relational Architects International

Hands-on First Look:

Informix 11.7

New options for data security and performance tuning really deliver. Plus, it's flat-out faster.

By Lester Knutsen and Art Kagel

Unless you've been hiding under a rock for the last few months, you know that IBM is pretty excited about the latest release of Informix. The company has good reason to be. During the beta program, Lester Knutsen and Art Kagel worked together to evaluate the newest release of Informix 11.7 server (formerly code-named Panther). They can tell you first-hand that there's some exciting stuff under the hood. In this article, they'll share the results of their performance testing, and they'll describe the new features and how to use them.

Informix 11.7 delivers significant speed boost

We tested Informix 11.7 using the benchmarks that we developed for the Fastest Informix DBA contests. For each test, we started with Informix 11.5, optimized it to get the best possible performance, and measured how quickly the job completed. Then, we installed 11.7 on the same hardware and measured the job again.

Out of the box, version 11.7 was faster on almost every trial. The first benchmark was a batch-billing job that generated bills for more than 100,000 customers. Informix 11.7 ran this job up to 17 percent faster than the best time posted by the previous version.

The next benchmark we tested was from the Fastest Informix DBA Contest III in 2010. It was an online transaction processing (OLTP) benchmark that used the open source BenchmarkSQL Java program to generate 100 sessions doing inserts, updates, and deletes against a database. The goal is to generate as many transactions as possible in a fixed amount of time; this JDBC benchmark closely resembles the TPC-C standard for OLTP. Again, using the same configuration as the prior version, 11.7 generated 12 percent more transactions.

Given that we didn't use any of the new features available in 11.7, these results are pretty amazing. Informix 11.7 also has a number of exciting new features for data warehouse design and development that we are still testing, using several very complex queries. Some of the new join features of 11.7 have helped us improve the performance of one query by 22 percent.

Implementing new Informix 11.7 features

When it set out to develop the 11.7 release of Informix Dynamic Server, the IBM development team pulled out all of the stops to extend and expand each area where Informix is already very strong. New features enhance OLTP data security, OLTP performance, data warehouse performance, and administration. This part of the article will address the features we think are most significant, what the benefits are, and how to implement them.

OLTP data security

What do we mean by OLTP data security? The worst-case scenario in an OLTP environment is loss of transactional data. For example, if you're a bank and a user completes a transaction, you cannot afford to lose it. Your company is on the hook for that money, whether it's a $20 ATM withdrawal or a multimillion-dollar electronic funds transfer. To take another scenario, if your primary server crashes, what do you do? If you tell the users who are mid-transaction, "Sorry. System problem. Please try again later," you will lose a customer or maybe several hundred customers.

Informix 11.7 has a new feature that IBM calls "transaction survival." We call it "uninterruptible transactions." If your primary MACH11 server goes down, any transactions begun by a user connected to any secondary server (HDR, SDS, or RSS) will be picked up by the surviving secondary servers. One of the surviving servers is promoted to be the new primary server, and the transaction continues unharmed and uninterrupted. Using SDS secondary servers and the Connection Manager, failover can be completed in mere seconds; this capability will support many organizations' service level agreements (SLAs), and most users will not even notice the delay.

The best part is that you don't have to do much at all to implement this amazing feature. No application changes, onmode flag, or SQL application programming interface

Also in 11.7:

Native time series data

Informix 11.7 also has native time series data support, which allows relational data to be combined with time series data in the same query. This native support helps improve query performance and reduce storage requirements for time series data. The customizable data loader handles data in almost any format and integrates with real-time streaming data products such as IBM InfoSphere Streams. Plus, the time series toolkit enables you to develop custom analytics to run in your database.

(API) functions are needed. A single new ONCONFIG parameter (FAILOVER_TX_TIMEOUT) sets the time that secondary servers will wait for a new primary server to make contact before rolling back open transactions after a primary server failure. Your existing applications don't have to be modified. This is all automatic and effort-free. It is not yet clear whether older applications will need to be recompiled with the latest Informix SDK version (3.70xC1 or later), but most sites will do this as part of an upgrade process anyway. After that, you just need to make sure that all of your critical applications are connected to a writable secondary server rather than connected directly to the primary server.

OLTP performance

Beyond the underlying database engine improvements, Informix 11.7 has several features you can use to further improve performance. The following sections describe three of these features in an OLTP context; however, each of these features will improve the performance of decision support and data warehouse systems as well.

What's new: Eliminate index on foreign keys. Informix 11.7 eliminates the requirement for a foreign key constraint to have an index on the foreign key. Indexes on the foreign keys referencing tables with few rows have low selectivity, and the utility of requiring the index has always been questionable. The engine does not need these indexes to enforce the constraints unless cascading deletes are enabled, and you probably don't need them for searching because most systems have composite keys containing the foreign key columns that are used for searching and filtering. It is rare in an OLTP query for the lookup tables to be selected as the primary query tables requiring an index on the dependent table to support the join. Removing this requirement will reduce the overhead of inserting and modifying data in tables that have many code columns, without causing any impact on query performance.

What to do. To create a new constraint without a supporting index, follow the constraint definition with the INDEX DISABLED clause:

```
ALTER TABLE mytable ADD CONSTRAINT FOREIGN
KEY (fkcol) REFERENCES fktable(keycol) CONSTRAINT
mytable_fk1 INDEX DISABLED;
```

If you have an existing constraint and would like to take advantage of this feature, you must drop the constraint and re-create it. There is no option to disable index

Heterogeneous grids: Everyone's invited

IBM is making quite a bit of noise about the heterogeneous grid capabilities of Informix, and rightfully so. If you have lots of old hardware around, don't throw it out. Don't sell it to a junker or to a used equipment reseller; Informix gives you the ability to use different hardware, operating systems, or even different versions of Informix from cluster to cluster. Configure your old hardware as an IDS grid node and add it to your company server farm's net computing power at near zero cost (since the boxes are already paid for and amortized). You cannot mix hardware, operating system, or Informix versions between MACH11 servers backing each other up, but you can have separate MACH11 clusters running on different platforms that are linked into a single grid using Enterprise Replication.

Heterogeneous grid capabilities are a part of the Informix Flexible Grid feature. They also automate some feats that previously required both planning and manual configuration. For example, the new grid management features and the ifxclone utility make it simple to do zero-downtime upgrades. The engine can temporarily (and automatically) convert your HDR pair into an ER pair, upgrade the Informix version on one server, clone the upgraded server, and reestablish HDR.

support for an existing constraint. When the constraint creation is complete, the index supporting the constraint is disabled and its space is dropped.

What's new: Forest of Trees indexes. Forest of Trees indexes combine many of the advantages of a hash index with those of a traditional Informix B+tree index. B+tree indexes can become rather deep with relatively few unique elements on a level in indexes on keys where the first column(s) have low selectivity but are still important to the correct processing of OLTP queries (think geographic codes such as country code, region code, state code, and combinations of these). With a Forest of Trees index, the database architect can specify one or more of the leading columns to be used to create a hash key. Then for each hash value, a separate B+tree is created that contains only the remaining key parts following the hash columns. The result is much flatter B+tree indexes. One of the drawbacks of using a Forest of Trees index is that you can't perform range scans on the hash columns (for that you will also need a pure B+tree index). But, you can perform range scans on the remaining columns in the index key. The most significant benefit will be indexes that have fewer levels and therefore faster performance.

What to do. Create a Forest of Trees index just as you would create any other index, except that after the column list, specify the HASH ON clause:

```
CREATE INDEX mytable_fot ON mytable( col_1, col_2, col_3 )
IN indx_dbs
HASH ON ( col_1, col2 ) WITH 200 BUCKETS;
```

Note that the storage clause must precede the HASH ON clause, which is not clear from the syntax diagrams in the manual.

What's new: Multiple index scans. Multiple index scans in the optimizer are another big win for OLTP. Often you have a join between two tables on one or more columns, with filters on one or more different columns in those tables. Both sets of columns may be indexed, but by different indexes. In earlier releases of Informix, the optimizer had to select the "best" single index for a query from all the indexes available based on index statistics and data distributions of the key columns. As a result, the engine needed to read many rows of data and apply the final filters using the actual row data, even though indexes of the filter columns were available. The optimizer in Informix Extended Parallel Server (XPS) could use more than one index on each table in a query, and now Informix 11.7 can as well. This capability can enable DBAs and database architects to reduce the number of long key compound indexes, which can also improve insert, delete, and update performance.

This feature could be a huge win when running queries like the one previously described against databases that are delivered as part of a third-party application—you can't control the schema, and the composite indexes that could make such queries reasonably efficient may not exist.

What to do. The best part of this feature is that there is nothing you need to do. Just make sure your statistics and data distributions are properly maintained. Once you have seen multiple index scans improve your query performance, you may want to look into replacing some compound key indexes with simpler ones and reducing the total number of indexes, which may improve delete, update, and insert processing times as well.

Your turn

Informix 11.7 expands the capabilities of our favorite database in a number of exciting ways. We're still investigating the new features, and even in this article, we've only been able to scratch the surface of what's possible. Check out Lester's Informix DBA column in this issue for more tidbits, and let us know what you think of the new version, and what you're most excited about. ✱

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

Art Kagel *(art@advancedatatools.com) is principal database consultant with Advanced DataTools Corporation. Kagel manages the Informix consulting and support team, providing remote and onsite DBA support to clients around the world. He is a current member of the IIUG board of directors and an IBM Data Champion.*

Migrating From Sybase ASE This Year?

If your resolution to keep database expenses within a modest IT budget is already colliding with the fact that there's too much month left over at the end of your money, you're not alone.

But there's good news for creative managers who are looking for new ways to make their 2011 numbers. "Standardizing DBMS products on just one or two can significantly lower database cost for licensing, administration, staffing and database administration tooling," according to an independent Forrester Research Inc. report. "Enterprises that standardize on one or two DBMS products typically save 25 percent or more," the report continues.[1]

"A query that used to take six hours now takes 15 minutes...[we] originally went to compression to save disk space and we've seen a performance improvement there, too."

Tom Holdener,
Senior Technical Specialist,
BJC HealthCare

To see a complete interview with Tom Holdener, please visit http://www.youtube.com/watch?v=1x5b4_sW8fY. For more information about IBM DB2 SQL Skin, visit:http://www.ibm.com/software/data/db2/linux-unix-windows/editions-features-sql-sybase.html.

Even better news: if you've been avoiding such standardization because of the risk and expense, new technology recently introduced makes database migrations easier than ever before, by reducing or eliminating the need to change the applications that run on legacy databases.

Some firms are already finding that this new technology—IBM DB2 SQL Skin for applications compatible with Sybase ASE—can help take the pain out of database migration by reducing or eliminating the need to rewrite applications developed for Sybase ASE. And in a world where every DBA needs additional training for software updates on each database the company owns, it can be expensive to continue to support multiple databases.

DB2 SQL Skin helps to migrate data from Sybase ASE to IBM DB2 using automated techniques. Then, queries originally written to run against a Sybase ASE database are instead run against the DB2 database. Applications continue to work as if they were still communicating with Sybase ASE, when they're really interfacing with DB2 SQL Skin and DB2. There's no disruption for end users, and no need for them even to know that migration has happened.

Developed jointly by IBM and ANTs Software, the product eliminates the need to license Sybase ASE and administer applications running on it. And because code doesn't have to undergo a major manual rewrite, or testing and validation, migration time can often be cut dramatically, reducing risk and costs.

For example, as part of its evaluation process, St. Louis–based BJC HealthCare recently migrated one of the company's two Sybase applications to run on DB2, both to avoid having to invest more in Sybase development and to free up its DBAs, who used Sybase so infrequently that administering it was turning into a huge time sink.

"If they had to apply a patch or they have an issue or anything like that, they had to do quite a bit of research to figure how to do it because they don't use it all the time," says Tom Holdener, senior technical specialist at BJC.

Holdener says the migration has had other important benefits, such as performance improvements. "A query that used to take six hours now takes 15 minutes," he says, adding that the company "originally went to compression to save disk space and we've seen a performance improvement there, too."

Furthermore, says Holdener, migration time was surprisingly short. "The new IBM DB2 SQL Skin feature built with ANTs technology allowed us to migrate a proprietary application from Sybase ASE to DB2 with virtually no changes, drastically reducing the time to a matter of weeks instead of months," he says.

These short migration times and the benefits of reduced costs are making automation-assisted migration a trend to watch throughout 2011. As Philip Howard of Bloor Research puts it, "Data migration is a big challenge for enterprises large and small, especially when they inherit applications and databases through acquisitions. I regard the new IBM and ANTs Sybase ASE to DB2 technology as a significant step forward and I expect to see more of this sort of capability introduced in the future."

[1] Yuhanna, Noel. "Take Advantage of New Ways to Save Money on Database Costs," Forrester Research Inc., February 23, 2010, p. 4

THE DATABASE REVOLUTION

By Carl Olofson

Today they're indispensable. Sixty years ago they didn't even exist as an idea. How did we get here?

Today, almost everyone who uses information technology simply takes the vast sea of underlying data for granted. Business users may be aware that a database is involved, but have little idea of the complex architecture needed to keep their data organized, related, linked, current, consistent, and available.

Even database administrators and other professionals operate within a universe of well-established concepts and reliable functionality. They depend on characteristics such as atomicity, consistency, isolation, and durability (ACID) properties as well as a range of data management technologies and methodologies that did not exist a half century ago. This article explores how these things came to be.

Early days: Start with a file system

In 1958, computers had already evolved from huge tube-driven hulks into smaller, lighter, transistorized machines. They could perform most simple management functions themselves, could load and execute application programs without rebooting, and most importantly for business, they were (relatively) affordable, practical, and manageable.

However, the data revolution was still in its infancy. Computer programs wrote data to storage devices formatted according to their own internal definitions. These data collections, eventually called "files," were normally used so that a program could generate output that would be read by the next program in a batch sequence. The last file in the sequence would be stored on tape or cards for the next time the sequence, eventually called a "job," was executed. Because a file was written and read according to definitions within the program code, there was no external data definition, making data sharing difficult and ad hoc reporting impossible. Although business users wanted to use this data for reporting or with other newer programs, there was no effective technique for collecting and reusing data in an organized manner.

As it turns out, 1958 was an important year for data. That was the year IBM produced a file system that could store and retrieve data based on data format definitions held in the file system rather than in the program. It was called the Formatted File System (FFS), and it was developed for the IBM 704 computer. Just two years earlier, IBM had also developed the first random access disk storage device, the Random Access Memory for Accounting and Control (RAMAC) 350, so there was a place to store this stuff. FFS was a critical first step toward automated database management; it spawned a series of software systems that could be regarded as the first database management systems (DBMSs).

During the next four years, the arrival of powerful, transistorized computers from several vendors as well as third-generation languages—including COBOL, ALGOL, and PL/I—resulted in a flurry of application programming, more and more application data, and a critical need to organize that data. A team of IBM developers responded with the Generalized Information Retrieval and Listing System (GIRLS) for the IBM 7090 in 1962. GIRLS improved on the FFS facilities, enabling users to collect data and easily code reports on a recurring basis. This activity could be thought of as a forerunner to data warehousing. The collection of data was referred to as a "data base" because it consisted of not one but many different record types.

The declaration of (data) independence

At the time that GIRLS was developed, a tremendous increase in computer functionality, manageability, and affordability was taking place. Computers were getting smaller and more powerful due to the development of integrated circuits (ICs) and large scale integration (LSI). Meanwhile, data could be shared even more widely using the American Standard Code for Information Interchange (ASCII) and IBM's Extended Binary Coded Decimal Interchange Code (EBCDIC).

However, data kept in the early FFS-based DBMSs of the day formed indexed sets of records having the same layout, and the records in one set could not be formally associated with the records in another set. In 1964, General Electric consultant Charles W. Bachman developed a way to enable the sharing of complex data under schematic control that would ensure that all the data forms necessary for the constituent applications would be preserved. He used a network orientation to build his design, and the result was a DBMS called the Integrated Data Store (IDS) that was designed to run under GE's GECOS operating system.

Bachman formally defined each bit of data typologically as an "element," such as "customer name" or "order number." Then, record types would be defined using these elements. Records would be stored in sets—not based on their type, but based on their relationship to the record that "owned" the set. So, for instance, a "customer" record might own the set of all the "order" records representing orders that that customer had placed. A "product" record might own the set of all the "order" records representing orders for that product. To find out what products a customer had ordered, it was possible to traverse the set of all orders placed by the customer and, for each order, find the product record that owned the set of ordered products in which that order record

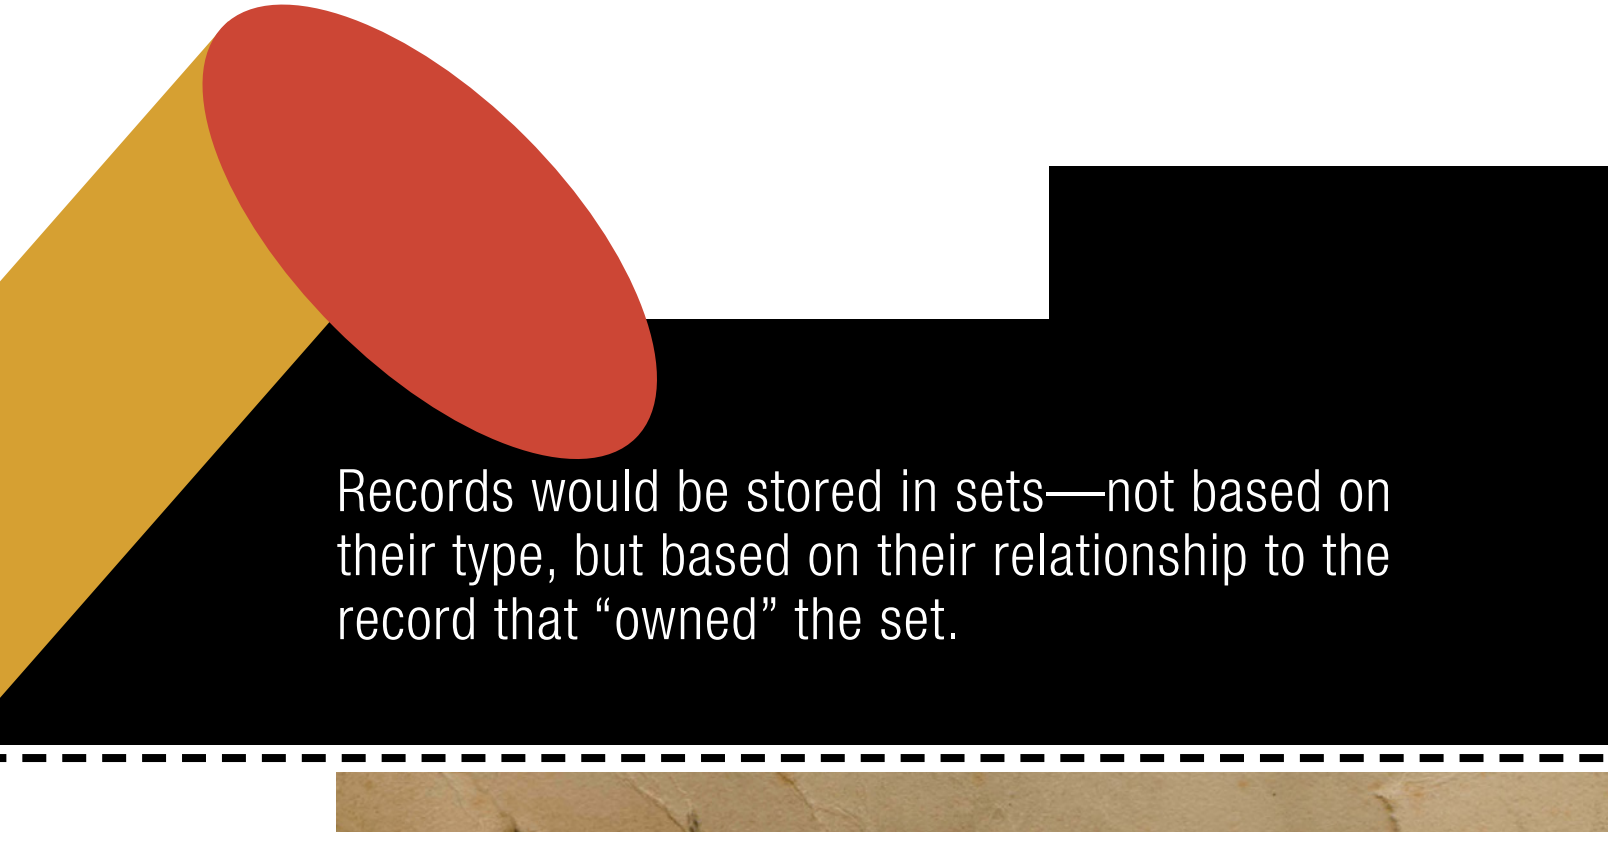

is kept. Records that had no owner were called "entrypoints" and were found based on their key value.

Although most DBMSs are not network oriented, they share these common features:

- Data can be defined at an elementary level, and it is kept in typologically consistent groups that can be related to other groups.
- Data groups can be found randomly, and associated data can be retrieved based on the way the data is defined to the database in a structure called a "schema."

In this way, the data definitions and the rules governing how they are stored and related to each other are independent of the programs that use the data. These concepts evolved from the FFS work, but they first found organized expression in IDS.

The following year, IBM produced its own stand-alone data storage and query system, called the Generalized Information Store (GIS). Over the next several years, the company executed a major project with Rockwell International to build upon GIS and produce a high-volume, large-scale DBMS for the National Aeronautics and Space Administration (NASA). Unlike Bachman's approach, which organized record types under other record types as "sets" with "owners" and "members" to form a data network, IBM and Rockwell International used a hierarchical organization of data to deliver very rapid response times in executing complex data transactions. The result was the Information Management System (IMS), which was delivered to NASA in 1969. IBM subsequently productized it for the System/360 mainframe.

The next 10 years represented the heyday of the mainframe DBMS, most of it centered on the IBM System/360 and its successor, the System/370. Network DBMSs IDMS and TOTAL, inverted-list DBMSs Adabas and Model 204, and indexed table DBMS DATACOM provided lively competition during this period. There were major DBMS offerings on other mainframes also: Honeywell, which had acquired GE's computer business, continued to offer IDS, and Burroughs had a similar network DBMS called DMS. In the end, however, it was the IBM mainframe that became dominant.

From navigational to relational

The DBMSs of the 1970s represented powerful advances in the ability of companies to collect, reuse, and report on data, and to manage large and complex, yet well-coordinated application systems. However, these DBMSs were complex and required people with considerable technical knowledge to understand them, much less design and manage them. They were also not useful for providing query support, because users had to know how to navigate the physical structure of the database to find the data.

But another quiet revolution was taking place—one that began with a research paper describing a simple way to collect, manage, and share data based on mathematical set theory. The paper was called "A Relational Model of Data for Large Shared Data Banks," and its author was an IBM engineer named E. F. Codd.

Codd had been bothered by the widespread practice in the database world of treating data and its definitions haphazardly, without a systematic approach that could lead to scalable sharing of the data across diverse systems. The problem, from his point of view, was a lack of mathematical rigor in defining the data model (some database models had no rigor whatsoever). His solution was to consider ways in which data could be defined and organized according to principles derived from mathematical set theory, allowing the data to be managed using elements of predicate logic. His work, which was laced with obscure mathematical terms, at first was barely understood in the broader data management community. Once some bright light translated "tuples," "attributes," and "relations" into "rows," "columns," and "tables," the relational concept spread like wildfire.

Codd's paper touched off a flurry of research into how his approach might be implemented practically. The first fruit of this effort was a purpose-built relationally-based data handling system built for the British Geological Survey in 1973. Called G-EXEC, it was developed on an IBM System/360 by a team led by Keith Jeffery from the British Geological Survey and Elizabeth Gill from the Atlas Computer Laboratory, along with Stephen Henley and John Cubitt from the British Geological Survey. Four years later, IBM's Jim Gray led a team of engineers to produce a prototype relational DBMS (RDBMS) called System R. (Gray went

To read an expanded version of this article, go to www.ibm.com/developerworks/data/library/dmmag/DMMag_2011_Issue1/FeatureHistory

The success of SQL/DS and its mainframe successor, DB2, helped ensure the adoption of SQL as the standard for RDBMS data access.

on to develop RDBMS technology at Tandem Computers and Digital Equipment Corporation, and he finished his impressive career at Microsoft.) By that time, Honeywell had produced an RDBMS for its Multics OS called the Multics Relational Data Store (MRDS), and Michael Stonebraker launched a project at the University of California at Berkeley to produce an RDBMS, code-named Ingres. During the following decade, these projects, as well as others from a variety of companies and targeting a variety of platforms, created the now-familiar RDBMS landscape.

The start of SQL

As part of the System R project, an IBM engineer named Donald Chamberlin led a team that developed an interactive query language called SEQUEL (a sort of acronym for Structured English Query Language) in 1974, which Chamberlin and colleague Robert Boyce rewrote in 1977. Its name was changed for legal reasons to SQL (Structured Query Language) even though it was built out to become a full data manipulation and definition language.

IBM promoted SQL as a standard and developed from System R an RDBMS called SQL/DS. The success of SQL/DS and its mainframe successor, DB2, helped ensure the adoption of SQL as the standard for RDBMS data access. The origins of DB2 go back to the OS/2 operating system: IBM took the RDBMS technology developed for OS/2 and created a powerful cross-platform RDBMS, which it called DB2 Universal Database (UDB) at first (in 1997). Today this RDBMS is generally known as DB2 for Linux, UNIX, and Windows (LUW).

Other DBMS technology developments

Some of the DBMS technology developments between 1970 and 2000 were not relational. Two of the engineers who had contributed to the development of GIRLS, Dick Pick and Don Nelson of TRW, developed a computer system with integrated database management capabilities based on GIRLS that came to be known as the Pick System. It had the additional advantage of supporting multivalued fields. It was sold as an integrated environment on a variety of minicomputers in the mid- to late 1970s.

In the late 1980s, object-oriented (OO) programming techniques and languages such as SmallTalk and C++ emerged, and soon there were DBMSs that could store and retrieve object attribute data seamlessly as extensions of the programming environment. These OO RDBMSs had sophisticated structures that supported nesting, recursion, and all the required OO characteristics, and were brought to market by such firms as Versant, Objectivity, Object Design, GemStone, and POET Software.

Some of these developments have influenced the course of the relational DBMS, resulting in extensions to the relational model such as multivalued data support and explicit support for unstructured and semi-structured data in the database. Stonebraker was looking to achieve this capability by applying OO principles in working on an object-relational DBMS (ORDBMS), which resulted in the emergence of Illustra in 1994. Informix acquired Illustra the following year, blending ORDBMS technology into a product called Informix Universal Server and inspiring Oracle and IBM to offer ORDBMS capabilities of their own.

The turn of the 21st century saw considerable consolidation. VMark (makers of UniVerse) acquired Unidata to become Ardent Software, which was also acquired by Informix, along with Red Brick, a data warehouse DBMS vendor. Informix DBMS technologies were ultimately acquired by IBM in 2001. Oracle acquired in-memory DBMS vendor TimesTen in 2005 and open-source small-footprint DBMS vendor Sleepycat in 2006.

In addition to consolidation, the DBMS industry saw considerable technical advancement by all the major vendors and some new players. These technical developments included advanced clustering support and, more recently, columnar and cell-based in-memory DBMS.

The revolution continues

Technologies underlying some products have fueled the design and development of what may be the next big leap in DBMS technology. The key-value pairs, list-oriented structures, and other technologies that offer flexibility and scalability for workloads that do not require perfect consistency are being rediscovered and offered, often in the context of the cloud, as so-called NoSQL DBMSs. In addition, object-oriented DBMS technology is being used to deliver graph databases, which are key to managing social networks and addressing other relationship-oriented Big Data problems.

In 2000, some leading analysts who watched the database industry said that DBMS was becoming a commodity—that most significant innovation is in the past. How wrong they were. The database revolution has barely begun. ✱

Carl Olofson *is a research vice president at IDC who studies the database software markets.*

A Benchmark for the Real World

IBM DB2 pureScale Software and Servers Based on Intel Xeon Processors Deliver Impressive Results in Scalability Tests

In a recent email, a DBA wrote that benchmarks are like unmatched socks; every time he turns around, he's looking at new ones. "They must be replicating when I'm not looking," he joked.

So if you've ever wearied of trying to relate lab-based performance tests to your environment, perhaps you'll be heartened by some important benchmarking news from people who understand just how critical databases are to today's enterprises.

Recently, IBM and Intel reported that they have achieved dramatic scalability and performance results on a real-world IBM® FileNet® P8–based content management application with IBM DB2® pureScale™ software running on Intel® Xeon® processor 7500 series–based servers. In such a representative environment, a cluster of eight Intel Xeon processor 7500 series–based systems processed over 5,400 documents per second—or over 19 million documents per hour—and never jeopardized the sub-second response times that users have come to expect.

Designed to mimic the database and content management needs of businesses ranging from insurance companies to construction, this benchmark tests more than just how fast, scalable, and reliable DB2 pureScale cluster technology can make existing databases and content management applications. Rather, the results show just how much DB2 pureScale clustering can do for performance and reliability for servers leveraging Intel's extremely fast Xeon processor 7500 series, and suggest how these new technologies could benefit a wide range of typical enterprises. And of course, it would be hard to find an organization that couldn't enjoy the benefits of high availability, performance, and reliability enabled by DB2 pureScale software running on Intel Xeon processor 7500 series–based servers.

Business Benefits: From Speed to Availability

Virtually no other type of application can claim the same high level of impact upon business and government as database software. Databases have become central to virtually everything that organizations do, from document processing and manufacturing to marketing and compliance. But this reliance means that any delay or bottleneck in database processing can literally slow the business down.

The quest for faster, more reliable, more scalable database technology is therefore relentless, as organizations search for ever-better systems upon which to base their enterprises. In addition, the deployment of cloud infrastructures that deliver Data-as-a-Service (DaaS) has put a premium on hardware and software that support rapid expansion or contraction.

These factors are driving intense interest in DB2 pureScale software, which uses remote direct memory access (RDMA) and other advanced techniques to cluster DB2 database servers for efficient distributed processing, with near-limitless scalability, high availability, and complete application transparency. (For more on how DB2 pureScale works, see "What is DB2 pureScale?" in *IBM Data Management* magazine, Issue 1, 2010, p. 35, http://bit.ly/auOvNt.)

DB2 pureScale Performance Study: Test Environment Configuration

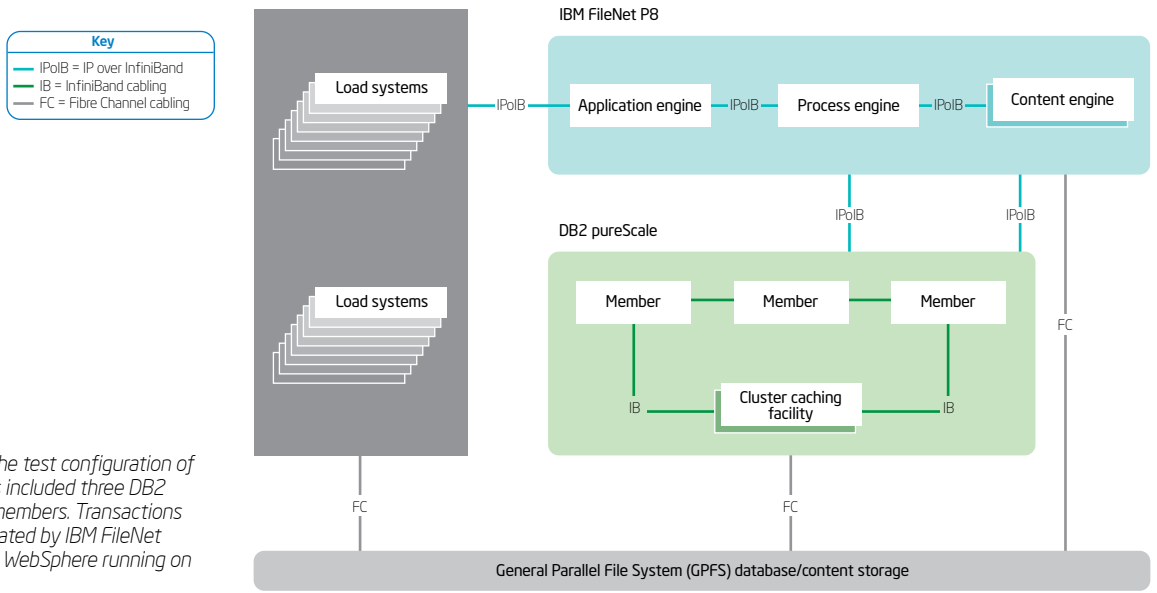

Figure 1: The test configuration of 30 systems included three DB2 pureScale members. Transactions were generated by IBM FileNet P8 and IBM WebSphere running on SLES 10.3.

In addition, DB2 pureScale software is designed to continue operations in the event of outages—planned or unplanned—and so can help improve the reliability of mission-critical applications built upon it. One of its most important capabilities is the power to rapidly expand or contract resources while continuing cloud operations, even in diverse, 24x7 environments. And that's important, because many DB2 applications run for years with little or no downtime, fueling legends of programs that haven't seen downtime since MTV was a youngster.

Benchmark Test Details

The benchmark test was designed to mimic a business with heavy processing and database activity—in this case, an insurance agency. Simulating real-world conditions, a typical content management application workload ran on more than 30 systems, generating a mixture of content management and records management transactions (such as document ingests, content searches, and content retrievals) at OLTP rates. Five systems ran IBM FileNet P8 5.0 and IBM WebSphere® 7.0 on SUSE Linux* Enterprise Server 10.3 (SLES 10.3). Three other systems acting as DB2 pureScale 9.8 member systems also ran SLES 10.3. See Figure 1 and Figure 2 for configuration details.

Each server contained four Intel Xeon processor X7560 CPUs, each with eight cores operating at 2.27 GHz. IBM selected systems with these processors in part because they are becoming a popular choice for consolidation, due to the fact that they deliver an average of 3x higher performance across a wide range of industry benchmarks, compared to the previous generation of processors.[1] This makes servers based on the Intel Xeon processor 7500 series an extremely fast, flexible choice for most enterprises.

Already, many organizations are exploiting these servers' multi-core nature to consolidate and upgrade their servers, and it's easy to see why. In other tests, the Intel Xeon processor 7500 series delivered up to 20x better performance per server than single-core processors.[2] Who wouldn't want to replace 20 aging servers with a single fast one? When one considers operating costs—which can rival purchase costs over

DB2 pureScale Performance Study: Test Environment Systems

Load systems

34 load-generation systems used to produce simulated search/retrieve/ingest operations

Software:

SuSE Linux Enterprise Server 10.3 (SLES 10.3)
Open Fabrics Enterprise Distribution 1.4.2 (OFED 1.4.2)
OpenSSL 9.8a
Remote Shell 0.17 (RSH 0.17)
Redundant Disk Array Controller 9.3 (RDAC 9.3)
General Parallel File System 3.3 (GPFS 3.3)
Java* Virtual Machine 1.6 (JVM 1.6)

Hardware:

QLogic* 4 Gb FC
Mellanox* IB

SAN hardware information

IBM System Storage™ DS5300

FileNet systems

IBM FileNet P8 5.0

One application engine, one process engine, and one to five content engines to support load

Software:

SLES 10.3
OFED 1.4.2
OpenSSL 9.8a
RSH 0.17
RDAC 9.3
GPFS 3.3
IBM WebSphere 7.0
IBM LDAP 6.2
JCC 3.59
JVM 1.6
IBM HTTP Server

Hardware:

Four Intel® Xeon® X7560 processors (2.27 GHz)
32 GB memory
QLogic 4 Gb FC
Mellanox IB

DB2 pureScale systems

DB2 pureScale 9.8

One to three members and one cluster caching facility to support increased user load

Software:

SLES 10.3
OFED 1.4.2
OpenSSL 9.8a
RSH 0.17
RDAC 9.3
GPFS 3.3

Hardware:

Four Intel Xeon X7560 processors (2.27 GHz)
32 GB memory
QLogic 4 Gb FC
Mellanox IB

Figure 2: The test environment included load-generation systems, IBM FileNet, IBM DB2 pureScale, IBM storage hardware, and servers with Intel Xeon X7560 processors.

DB2 pureScale Performance Study: Performance Results

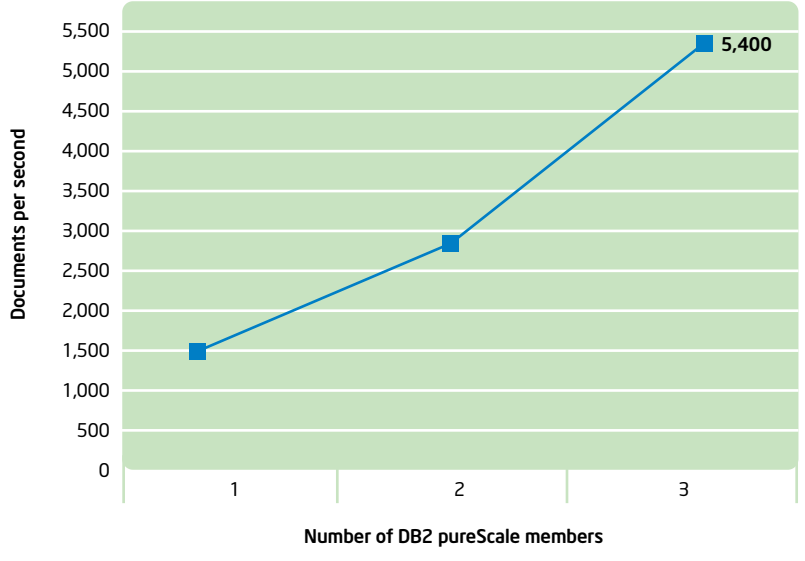

Figure 3: Benchmark results show linear scalability as the number of DB2 pureScale member systems increased.

the life of a server—such consolidation can be a huge win. Previous studies have shown that such 20-to-1 consolidations can save up to 90 percent in operating costs,[3] with such investments essentially paying for themselves in as little as a year.[4]

Impressive Scalability Results

During testing, the systems achieved an impressively high level of scalability: the cluster of eight systems processed over 5,400 documents per second, or over 19 million documents per hour at sub-second response times (see Figure 3). Not only did the systems scale linearly as the number of DB2 pureScale members was increased from one to three, but the third member was not saturated during testing, indicating that the system had the capacity to handle more throughput.

Beyond the Benchmark

So what can be learned from this type of enterprise benchmark? The most important "takeaway" is that this combination of DB2 pureScale software and servers based on the Intel Xeon processor series offers excellent scalability while maintaining high performance levels—and does so in processing- and application-intensive environments common to a wide variety of companies.

Furthermore, the scalability of the system means that with DB2 pureScale software, businesses can add state-of-the-art clustering technology to grow their database capacities to meet demand, without jeopardizing service levels. In a true business situation involving the kind of workload increases involved in the test, for example, any service-level agreements in place would have been met, and users would have continued to experience "immediate" responses, since all response times were less than one second. And because databases are so central to everything that businesses do, this means better, more efficient processing throughout the enterprise.

Now that's a benchmark result worth replicating in any organization.

Learn More

More about IBM DB2 pureScale: www.ibm.com/db2/purescale

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

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

---

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

[2] Claim: "Up to 20x performance per server" Disclaimer: Intel performance comparison using SPECjbb*2005 business operations per second between five-year-old single-core Intel® Xeon® processor 3.33 GHz–based servers and one new Intel Xeon processor X7560–based server. Performance tests and ratings are measured using specific computer systems and/or components and reflect the approximate performance of Intel products as measured by those tests. Any difference in system hardware or software design or configuration may affect actual performance. Buyers should consult other sources of information to evaluate the performance of systems or components they are considering purchasing. For more information, visit www.intel.com/performance/server.

[3] Claim: "Up to 90 percent lower operating costs" Disclaimer: Intel comparison replacing 20 five-year-old single-core Intel® Xeon® processor 3.33 GHz–based servers with one new Intel Xeon processor X7560–based server. Costs have been estimated based on internal Intel analysis and are provided for informational purposes only.

[4] Claim: "Estimated as little as one-year payback on investment" Disclaimer: Intel comparison replacing 20 five-year-old single-core Intel® Xeon® processor 3.33 GHz–based servers with one new Intel Xeon processor X7560–based server. Return on investment has been estimated based on internal Intel analysis and is provided for informational purposes only.

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

DB2 Indexes and Query Performance: Part 2

No new indexes? No problem.

In the first installment of this two-part column, I focused on the pros and cons of reducing query times by adding indexes to a DB2 database. We discovered that there is a point where aggregate index costs—more-expensive INSERT and DELETE operations, for example—outweigh a new index's query speed-up benefit. But if you don't want to increase the number of indexes defined on the table, are you at a performance tuning dead end? Of course not. You have plenty of options for accelerating queries in a DB2 environment while holding the line on indexes.

Better indexes versus more indexes

Start with an index effectiveness upgrade. Can you replace useless indexes with indexes that enhance query performance? This could be a challenge with earlier versions of DB2 if a lot of dynamic SQL statements were executed, because the SYSCAT.PACKAGEDEP catalog view (DB2 for Linux, UNIX, and Windows [LUW]) and the SYSIBM.SYSPACKDEP catalog table (DB2 for z/OS) show index dependencies only for static SQL statements.

However, DB2 9.7 for LUW and DB2 9 for z/OS introduced a handy catalog column called LASTUSED. In the SYSCAT.INDEXES catalog view (LUW) and the SYSIBM.SYSINDEXSPACESTATS catalog table (z/OS), LASTUSED shows the date on which an index was last used either to speed up a DML statement (static or dynamic) or to enforce a referential integrity constraint. If the LASTUSED column contains its default value ('0001-01-01' on LUW or NULL on the mainframe), the associated index is a good candidate for being dropped. Of course you'll want to do some additional checking, just to make sure that the index truly isn't needed.

Another way to boost index effectiveness is to add columns to an existing index, so as to increase the number of columns that can be matched with query predicates or to get index-only access for more queries. But be careful with this tactic; you do not want to replicate a table in an index by including all of the table's columns in the index key.

Consider adding a column to an existing index if the additional column is relatively short and if its inclusion in the index key would provide a major performance boost for one or more high-cost queries. With DB2 for LUW, you'd drop the existing index and re-create it with the additional column; on a DB2 for z/OS system, you can add a column to an index via an ALTER INDEX ADD COLUMN statement. DB2 10 for z/OS introduces a nifty new way to add a column to an existing index: the ADD INCLUDE option of ALTER INDEX. Check out the Resources sidebar to learn more.

No-new-index tuning options

Rewriting long-running queries is another way to boost performance without adding indexes. One technique often used is rewriting to eliminate predicates linked by OR, which are not indexable (aside from the possibility of DB2 using a multi-index access mechanism, called index ORing, that involves the union of row ID lists from two different indexes).

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

IDUG
The Worldwide DB2 User Community

IDUG Mentor Program

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

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

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

What's on the Horizon?

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

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

IDUG Conferences

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

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

MENTORING | CERTIFICATION | DB2 EXPERTS | NETWORKING | EDUCATION

Be a Part of the IDUG Community

Strength in numbers

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

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

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

Using a UNION or an IN-list are two common ways to eliminate an OR in a predicate (see Figure 1).

Using UNION to replace an OR can be especially effective if the query involves a join and the two ORed predicates reference columns in two different tables. In such a case, the OR prevents the two predicates from being applied until *after* the respective tables have been joined. Note also that in the IN-list example, both of the predicates in the rewritten query are stage 1 and indexable, leading to better index key column matching and potentially much-reduced query elapsed time.

Also in the query rewrite vein, modifying a SELECT to get rid of a range-type predicate can increase performance by improving index key column matching. As we saw in the last issue, the matching of predicates to index key columns stops after a range predicate—such as >, <, or BETWEEN—is encountered. In some cases, you can use an IN-list to replace a range predicate. For example, if COL1 is a character column of length one, a predicate like COL1 BETWEEN '2' AND '4' could be rewritten as COL1 IN ('2', '3', '4').

Eliminate OR with UNION

```
SELECT * FROM TABLE_A A, TABLE_B B WHERE A.COL3 = B.COL3 AND A.COL1 = 2 OR B.COL2 = 'CAT'
```

Becomes

```
SELECT * FROM TABLE_A A, TABLE_B B WHERE A.COL3 = B.COL3 AND A.COL1 = 2
UNION
SELECT * FROM TABLE_A A, TABLE_B B WHERE A.COL3 = B.COL3 AND B.COL2 = 'CAT'
```

Eliminate OR with an IN-list

```
SELECT * FROM TABLE_A
WHERE (COL1 = 2 AND COL2 = 'CAT') OR (COL1 = 2 AND COL2 = 'DOG')
```

Becomes

```
SELECT * FROM TABLE_A
WHERE COL1 = 2 AND COL2 IN ('CAT', 'DOG')
```

Figure 1: *Rewriting to eliminate predicates linked by* OR

Here's one more query rewrite idea: if possible, remove column expressions in predicates to make the predicates indexable. Let's say that :V1 and :V2 are host variables, with :V2 containing a date value. Now, consider this predicate:

```
(START_DATE - :V1 DAYS) < :V2
```

The arithmetic expression involving the START_DATE column makes the predicate non-indexable. If you rewrite the predicate to get the arithmetic operation on the right side of the operator ("operator" being =, >, <, etc.), it becomes indexable and you could see a big

reduction in query run time. The rewritten example predicate would look like this:

```
START_DATE < (DATE(:V2) + :V1 DAYS)
```

No query mods needed

There are also no-new-index-needed tuning actions that don't require query modification. Do you have a high-cost query with a predicate that references a column that has many duplicate values and for which the distribution of those values is very non-uniform (for example, a column of a 1-million-row table that has 1,000 distinct values, one of which appears in 900,000 of the 1 million rows)?

Let DB2 know about that uneven distribution of duplicate values by gathering value-frequency distribution and/or histogram statistics for the column via the RUNSTATS utility. Absent these statistics, DB2 assumes a uniform spread of duplicate column values across a table's rows, so in the aforementioned example it would be assumed that each of the 1,000 distinct column values appears in 1,000 of the table's 1 million rows. For DB2 for LUW, column value distribution statistics are collected via the WITH DISTRIBUTION clause of RUNSTATS. For DB2 for z/OS, value-frequency stats are generated using the FREQVAL option of RUNSTATS, and histogram statistics—available on the mainframe platform starting with DB2 9—are obtained via the HISTOGRAM option (histogram statistics have been available on DB2 for LUW systems for quite some time).

Once the catalog has been populated with value-frequency and/or histogram statistics, rebind the program (or just submit the query again, if it's dynamic SQL), and you may see a big-time performance improvement. By the way, you can get recommendations on enriching catalog statistics to potentially improve the performance of a query by running the query through the Statistics Advisor component of IBM's Data Studio tool.

The right tool for the job

When the task at hand is query performance tuning, of course you should keep the new-index option in mind. Just remember that it's a high-value card that should be played at the right time. And don't forget the other techniques available to you. Indexes aren't the only tool in the query-tuning shed, pardner.

RESOURCES

Matching predicates and index key columns ("DB2 Indexes and Query Performance: Part 1," *IBM Data Management*, Issue 4, 2010): ibm.com/developerworks/data/library/dmmag/DMMag_2010_Issue4/DataArchitect/index.html

Add columns to existing indexes with the DB2 10 ALTER INDEX statement in the DB2 10 for z/OS SQL reference: http://publib.boulder.ibm.com/infocenter/dzichelp/v2r2/topic/com.ibm.db2z10.doc.sqlref/db2z_sql_alterindex.htm

How Well Do You Know the Rules?

Understanding the rules that govern referential integrity constraints

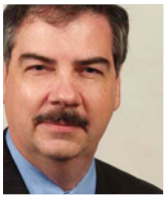

Roger E. Sanders (roger_e_sanders@yahoo.com), a consultant corporate systems engineer at EMC Corporation, is the author of 20 books on DB2 for Linux, UNIX, and Windows and a recipient of the 2010 IBM Information Champion award. His latest book is titled From Idea to Print: How to Write a Technical Article or Book and Get It Published.

If you've worked with relational databases for any length of time, you're probably aware that referential integrity (RI) constraints are used to define relationships between two or more base tables. But are you aware that a set of rules governs the behavior of RI constraints? More importantly, do you know how these rules work to safeguard data integrity when DML operations are performed against tables that are related? In this column, I'll identify the rules available and I'll illustrate how these rules control behavior when insert, update, and delete operations are performed against tables that are linked by one or more RI constraints.

Why the RI constraint rules exist

The primary reason for using an RI constraint is to ensure that data integrity is maintained whenever one table references another. When one table (the child) is linked to another (the parent) via a foreign key (one or more columns in the child table that reference columns in the parent table), the resulting RI constraint guarantees that for every row in the child table that has a value in a column that's part of a foreign key, a corresponding row will exist in the parent table. So what happens when a SQL operation tries to manipulate data in a way that would void this guarantee? Before answering this question, let's look at how data integrity could be compromised if the checks and balances provided by an RI constraint were not in place:

- An insert operation could add a record to a child table for which there is no matching record in the associated parent table.
- An update operation could change a record in a child table such that it no longer has a matching record in the associated parent table.
- An update operation could change a record in a parent table, leaving records in an associated child table that no longer have a matching record in the parent table.
- A delete operation could remove a record from a parent table, leaving records in an associated child table that no longer have a matching record in the parent table.

DB2 can either prevent such operations from being performed, or it can attempt to perform these actions in a way that will safeguard data integrity. That's where the rules that govern RI constraints come into play. Each RI constraint has three rules—an Insert rule, an Update rule, and a Delete rule—and the way in which DB2 responds to operations that threaten data integrity is controlled by the way two of these rules are defined.

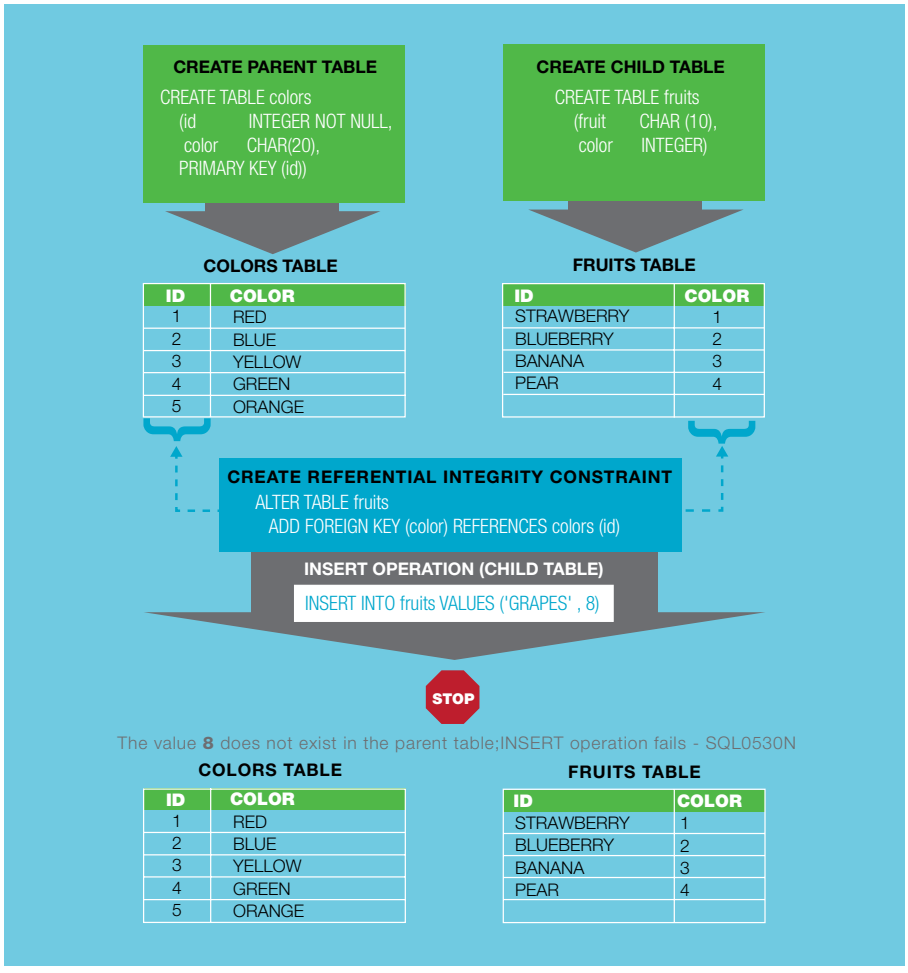

Figure 1: *How the Insert rule is enforced*

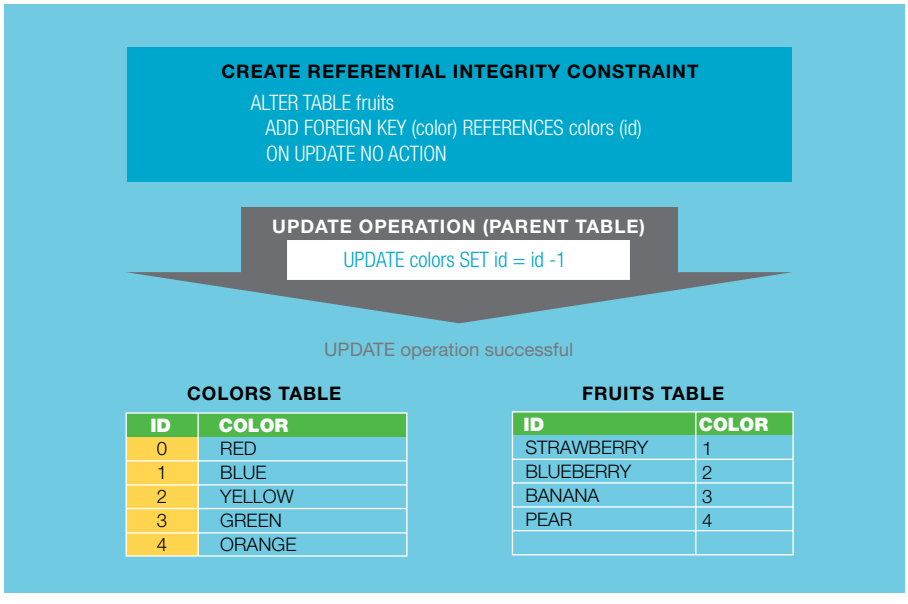

Figure 2: *Example where the* ON UPDATE NO ACTION *definition allows data in a parent table to be modified*

The Insert rule

The Insert rule guarantees that a value will never be inserted into the foreign key of a child table unless a matching value already exists in the corresponding key column(s) of the associated parent table. Any attempt to insert records into a child table that violate this rule will result in an error. In contrast, no checking is performed when records are added to the key columns of a parent table. Figure 1 illustrates how the Insert rule is enforced.

The Update rule

The Update rule controls how update operations performed against tables linked by an RI constraint are processed. Two types of behaviors are possible, depending upon how this rule is defined. Update rules can be defined as follows:

- ON UPDATE RESTRICT—When an update operation is performed on the parent table, each record in the child table will have the *same* value for its foreign key that it had before the update operation was performed. This rule is enforced before all other constraints, including other RI constraints.
- ON UPDATE NO ACTION—When an update operation is performed on either table (parent or child), each record in the child table will have a value for its foreign key that has a matching value in the key column(s) of the associated parent table. However, the value may not be the same as it was before the update operation occurred. This rule is enforced after all other constraints, including other RI constraints, are applied.

Figure 2 illustrates how the ON UPDATE NO ACTION definition will allow an update operation to be performed that the ON UPDATE RESTRICT definition would prevent. (Note that the COLORS and FRUITS tables referenced in this illustration were created as shown in Figure 1.)

Like the Insert rule, the Update rule is implicitly created as part of an RI constraint. If an Update rule is not explicitly defined, the ON UPDATE NO ACTION definition is used by default.

The Delete rule

The Delete rule controls how delete operations performed against a parent table in an RI relationship are processed. Four types of behaviors are possible, depending upon how this rule is defined. Delete rules can be defined as follows:

- ON DELETE CASCADE—When a record is deleted from the parent table, all records in the child table with matching foreign key values are also deleted.
- ON DELETE SET NULL—When a record is deleted from the parent table, all records in the child table with matching foreign key values are set to NULL (provided the columns that make up the foreign key are nullable). Other values for the dependent row are not affected.
- ON DELETE RESTRICT—When a delete operation is performed on the parent table, each row in the child table will have the *same* value for its foreign key that it had before the delete operation was performed. This rule is enforced before all other constraints, including other RI constraints that modify data such as ON DELETE CASCADE and ON DELETE SET NULL.
- ON DELETE NO ACTION—When a delete operation is performed on the parent table, each row in the child table will have the *same* value for its foreign key that it had before the delete operation was performed. This rule is enforced after all other constraints, including RI constraints that modify data such as ON DELETE CASCADE and ON DELETE SET NULL.

Figure 3 illustrates how the Delete rule is enforced when the ON DELETE CASCADE definition is used; Figure 4 illustrates how the

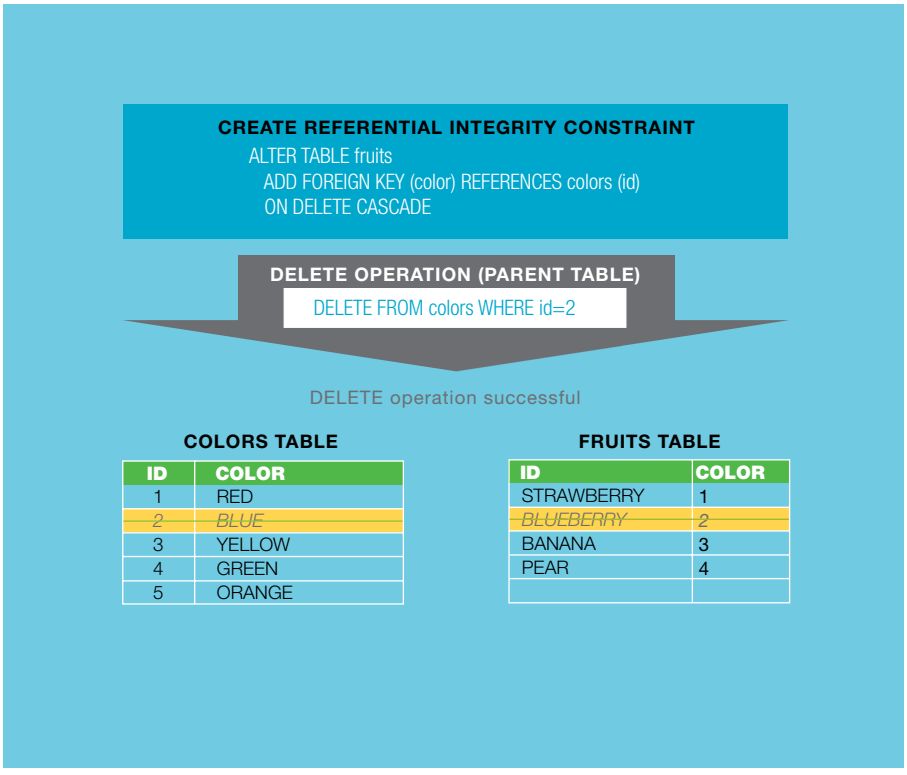

Figure 3: *Example of the* ON DELETE CASCADE *definition at work*

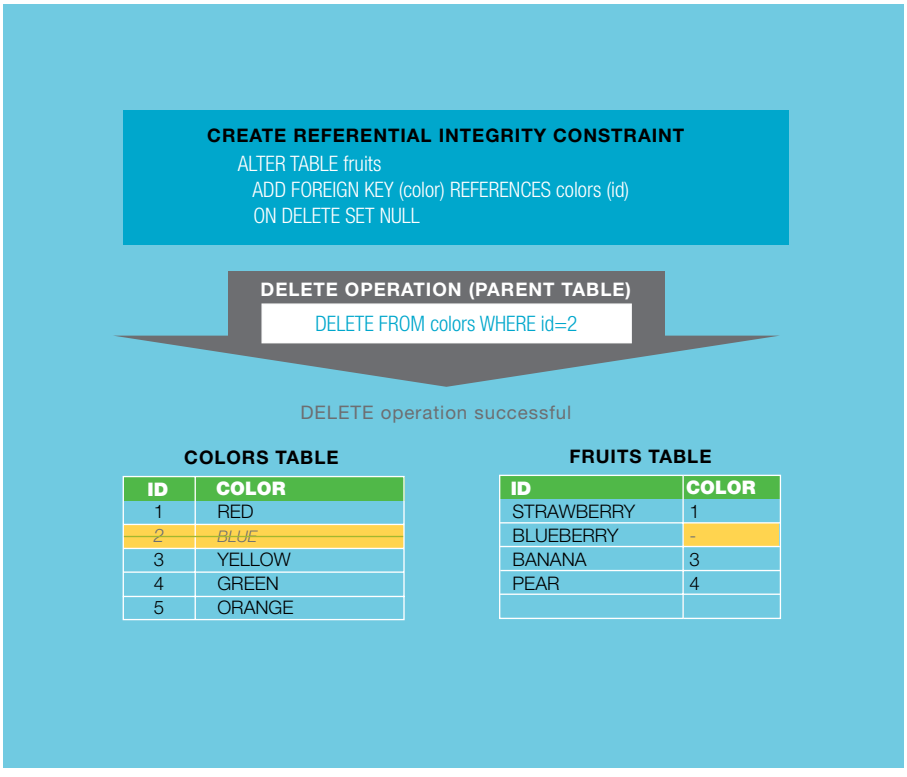

Figure 4: *Example of the* ON DELETE SET NULL *definition at work*

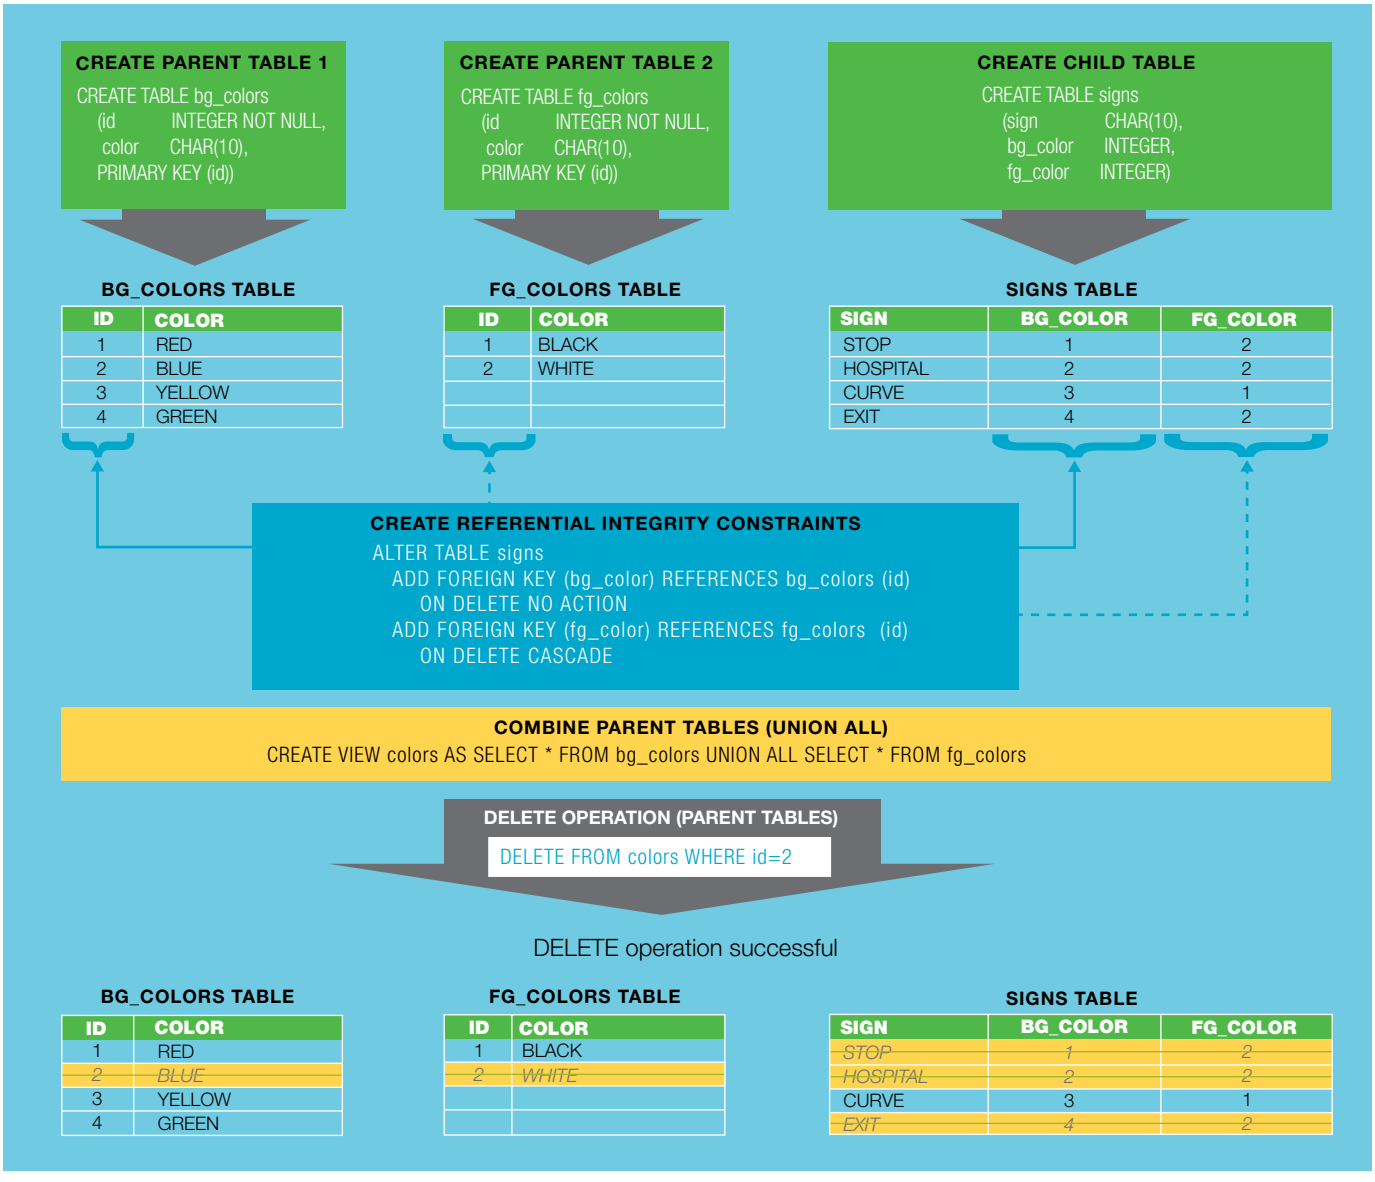

Figure 5: *Example where the* ON DELETE NO ACTION *definition allows data in parent tables to be deleted*

Delete rule is enforced when the ON DELETE SET NULL definition is used. (Note that the COLORS and FRUITS tables referenced in these illustrations were created as shown in Figure 1.)

If a table has only one RI constraint, the behavior of ON DELETE RESTRICT and ON DELETE NO ACTION is essentially the same. Where the two definitions differ can be seen when an attempt is made to delete rows from multiple parent tables that are referenced by a single child. Figure 5 shows an example where the ON DELETE NO ACTION definition will allow a delete operation to be performed that the ON DELETE RESTRICT definition would prevent. In this example, the RI constraint with the ON DELETE CASCADE definition is processed before the RI constraint with the ON DELETE NO ACTION definition, allowing records to be deleted from all of the related tables.

As with the previous rules, the Delete rule is implicitly created as part of an RI constraint. If a Delete rule is not explicitly defined, the ON DELETE NO ACTION definition is used by default.

Make the rules work for you

Within most businesses, data often must adhere to a certain set of rules and restrictions. (For example, "Every employee must be assigned to a department.") By using RI constraints, you can easily place some of the logic needed to enforce such business rules directly in a database rather than in applications that interact with the database. But when you define RI constraints, keep in mind that there are rules in place to safeguard data integrity when DML operations are performed against related tables.

New ORDER BY Information

Avoiding sorts with data-partitioned secondary indexes

In the latest releases of DB2 for z/OS, our body of knowledge about ORDER BY has to be revamped. Not only are there many new things to learn, but there are also many changes to old things that we once knew about ORDER BY. I'm going to devote the next few columns to exploring five of these thought-provoking, even worrisome, issues:

- The impact of a data-partitioned secondary index (DPSI) structure on ORDER BY sort avoidance
- The impact of using the "random" index option on ORDER BY sort avoidance
- The issue of having only a GROUP BY COL1, COL2 versus having both a GROUP BY COL1, COL2 and a follow-up (unnecessary?) ORDER BY COL1, COL2 in your SQL statement
- The potential of using DYNAMIC SCROLL cursors as a new way of forcing DB2 to avoid a sort for an ORDER BY (or any other sort syntax where an index can be used to avoid the sort)
- The new (limited) ability to use the UPDATE WHERE CURRENT OF cursor even though there is an ORDER BY in the CURSOR declaration—that is, the cursor is an "unambiguously read-only cursor"

Some background on sort avoidance

Long ago, in one of my Programmers Only columns ("Looking for a little ORDER?"), I explained that DB2 does not automatically build a result set of all qualified rows at OPEN CURSOR. If you have a long-running OPEN CURSOR, it is not because DB2 is building a result set. It is most likely because DB2 is doing a data sort at OPEN (usually to satisfy a common ORDER BY). I do not mind when folks say that DB2, when sorting, builds a SORTOUT file at OPEN CURSOR. Why? Because it's clear that the work is the result of having to sort. I do take issue with folks saying that DB2 builds a result set at OPEN because calling it a "result set" implies that it is *always* built for *every* OPEN and does not explain that the result set is actually a SORTOUT file whose creation and necessity are driven by both syntax and access path. It is not a given, nor is it a requirement.

Remember, sort syntax such as ORDER BY does not necessarily have to cause a data sort. With the appropriate access path (an index is available and by using it the right way, DB2 can return qualified rows back to your program in the desired order), the ORDER requirement can be met without sorting. In fact, the use of an index for sort avoidance is a key way to improve the performance of SQL that qualifies far more rows than will actually ever be FETCHed. If DB2 does not need to sort, then all of the work for finding a qualified row (or rows with multi-row FETCH) will be done at each FETCH. You stop FETCHing, and DB2 will stop looking for qualified rows. You will only pay the price of finding the row(s) you actually need.

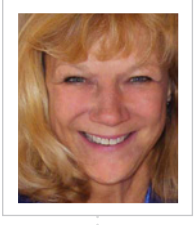

Bonnie Baker
(bkbaker@bonniebaker.com) specializes in teaching onsite classes for corporations, agencies, and DB2 user groups. She is an IBM DB2 Gold Consultant, an IBM Information Champion, a five-time winner of the International DB2 Users Group (IDUG) Best Speaker award, and a member of the IDUG Speakers'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

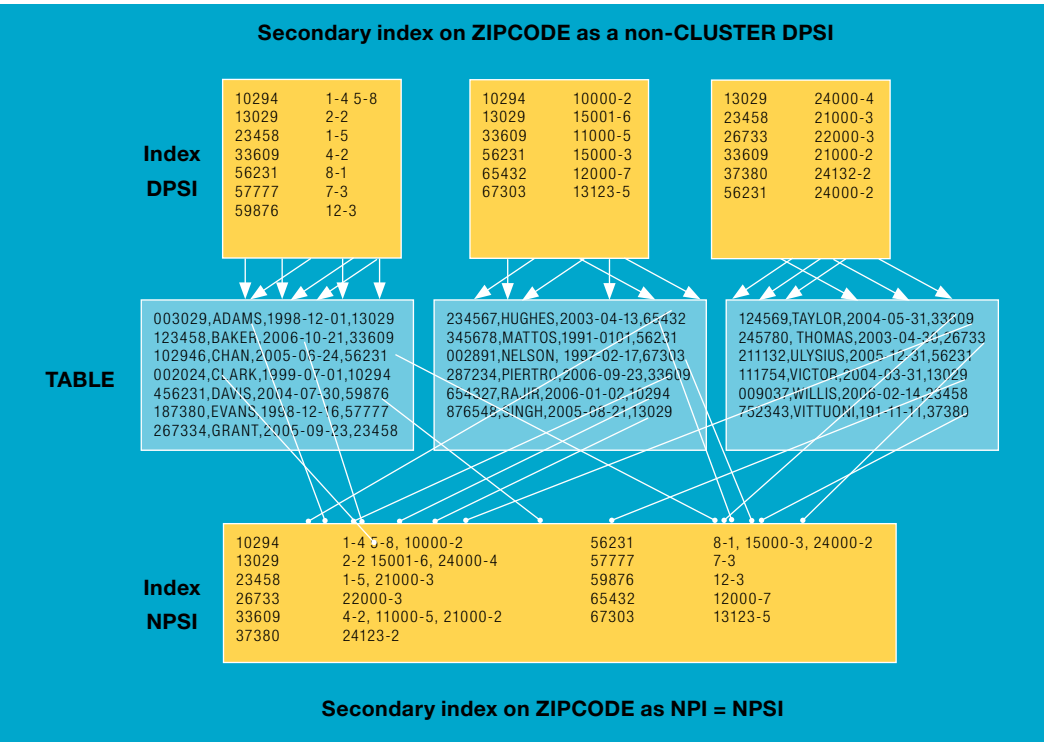

Figure 1: *Example of DPSI and NPSI indexes on the column ZIPCODE*

Now, what do I mean by using the *appropriate* index the *right* way? An appropriate index is one that will allow DB2 to retrieve the rows that you want in the order in which you want them without sorting. For example, if you have a four-column index on LASTNAME, FIRSTNAME, MIDINIT, EMPID (all ascending), and you code the following SQL statement:

```
SELECT ... FROM BIGTABLE
WHERE LASTNAME = :HVLASTNAME
```

followed by any of the following ORDER BY clauses:

```
ORDER BY FIRSTNAME, MIDINIT, EMPID
ORDER BY FIRSTNAME, MIDINIT
ORDER BY FIRSTNAME
ORDER BY FIRSTNAME DESC, MIDINIT DESC,
  EMPID DESC
ORDER BY FIRSTNAME DESC, MIDINIT DESC
ORDER BY FIRSTNAME DESC
```

then DB2 could use the four-column index the "traditional" way (index, table, index, table, etc.) to return the rows in the desired order (the same order as or the reverse order of the index). If DB2 were to use the index the nontraditional, LIST PREFETCH way, the rows would not come back in the order of the index, but rather in the order of the table. That is, the rows would come back in table CLUSTER order if the table were freshly REORGed (with no intervening maintenance or "log apply" impact on its perfect order).

When DB2 does not have an index in the appropriate order for sort avoidance or when DB2 uses any index the LIST PREFETCH way, then an ORDER BY will most certainly cause a data sort.

So, now on to the five new issues mentioned earlier.

Issue 1: The impact of DPSI structure on sorts

By creating a DPSI, you are essentially taking what would have been a single ordered list (once a non-partitioning index [NPI], now renamed non-partitioned secondary index [NPSI]) with row identifier (RID) pointers to all partitions of a table space and turning it into multiple ordered lists, each with pointers to only one partition. Since our subject index is not in the same order as our table (it is not our CLUSTER index), the order of an NPSI would be random to the *entire* table while the order of each index partition of a DPSI would be random only to its own *partition* and (we like this) has the probability of being a narrower (fewer leaf pages) and shorter (fewer levels) index. These are especially good features when you are accessing rows from a single partition. We are not so pleased with the DPSI if we must probe multiple index trees to find rows across multiple partitions. If we make this index our CLUSTER index, as an NPSI, the index would still be random to the table (partitioned on LASTNAME) as a whole, but as a DPSI, each index partition would be aligned (not random) to its table partition. Again, this is a great feature when you are accessing rows from a single partition. This is, of course, the ideal scenario. The less-than-ideal scenario of using a DPSI to access multiple partitions can be very painful from a SQL performance perspective. More on this later.

Figure 1 shows an example of two indexes: one a DPSI and the other an NPSI on the column ZIPCODE. The table is in the middle of the picture. The column data on the table are EMPID, LASTNAME, HIREDATE, and ZIPCODE. The table is currently CLUSTERed on LASTNAME.

Because the NPSI is a single ordered list, it is easy to see how DB2 could just toggle back and forth, index to table, index to table, and return rows in ZIPCODE order without sorting. The reads would be random (ouch), but for some transactions a few random reads is a small price to pay in exchange for eliminating a large data sort. In other words, many rows qualify but the program only FETCHes 10 rows; for example:

```
SELECT ... FROM BIGTABLE (for example,
  60 million rows)
WHERE ZIPCODE BETWEEN 10000 AND 70000
ORDER BY ZIPCODE
FETCH FIRST 10 ROWS ONLY
```

The FETCH FIRST clause discourages DB2 from using our index the LIST PREFETCH way (in which case, because of the intervening RID sort, DB2 would need to sort to satisfy the ORDER BY) and encourages DB2 to use the index the traditional way to avoid a large data sort.

However, look at the alternative DPSI; we have three ordered lists (one per partition). Obviously, we have no problem with sort avoidance if we code:

```
SELECT … FROM BIGTABLE (for example,
  60 million rows)
WHERE ZIPCODE BETWEEN 10000 AND 70000
  AND LASTNAME BETWEEN 'A       ' AND
  'G99999'
ORDER BY ZIPCODE
FETCH FIRST 10 ROWS ONLY
```

Even though LASTNAME is *not* part of our ZIPCODE index, DB2 knows that the rows we want can only be in Partition 1. Therefore, it can use Partition 1 of our index the traditional way to avoid a sort on ZIPCODE. But what if we code the following SQL?

```
SELECT … FROM BIGTABLE (for example,
  60 million rows)
WHERE ZIPCODE BETWEEN 10000 AND 70000
  AND LASTNAME BETWEEN 'A       ' AND
  'T9999'
ORDER BY ZIPCODE
FETCH FIRST 10 ROWS ONLY
```

DB2 knows that our desired rows cross multiple partition boundaries, and it knows that it will need to either do a table space scan across multiple partitions or use more than one index partition, each in its own order, to find each row. The FETCH FIRST clause encourages sort avoidance. So what does DB2 do?

DB2's solution

With smoke and mirrors and a bit of match/merge logic, DB2 can use multiple partitions of a DPSI for sort avoidance. Obviously, this is not as trivial as using a single index's ordered list, but it can be done. The DB2 Optimizer knows that in our case, three partitions (but in others maybe 64 or even more than 4,000 index partitions) must be read in a match/merge fashion in order to return the first 10 rows in ZIPCODE order, but DB2 can factor that possibility into its cost-based decision (see APAR PM25934 (open right now), which is refining DPSI costing). Without the FETCH FIRST (or similarly, but not exactly, without an OPTIMIZE FOR clause), then DB2 would almost certainly prefer some other access path.

The important point here is that you can choose to define indexes as DPSIs without losing that index's aptitude for sort avoidance, including the sort avoidance required for DYNAMIC SCROLL cursors. You just need to realize that we have been warned about the negatives of DPSIs in cross-partition processing, including the multiple index probes needed for sort avoidance. So take heed.

Negatives of DPSIs

Just in case you think that I am very pro-DPSI, I feel that it is important to stress some of their negatives:

- Potentially, far more probes and getpages.
- The loss of index lookaside. An NPSI would take advantage of lookaside on the index, for example, with a nested loop join in which the outer rows drive sequential access to the inner. For an NPSI, this results in index lookaside and sequential detection. For a DPSI where qualified rows are in more than one partition, every probe skips between partitions. The partition skipping breaks index lookaside and can also disrupt sequential detection.

Wrap-up

Knowing that we do not give up sort avoidance for ORDER BYs with DPSIs is good news. The next few columns will continue with the other four issues regarding ORDER BY (or the lack thereof) on our SQL.

Go Cruising with Informix in 2011

IIUG Informix Conference

Be Part of the Largest Informix Gathering in the World

- Over 80 Technical Sessions
- Free Hands on Labs
- Half and Full Day Tutorials
- Free IBM Certification Testing
- Perfect for DBAs and Developers
- Learn About New Informix 11.7 Features

Get $100 off registration
by using code **AFJR3101**

visit iiug.org/conf for more information

Overland Park Marriott, Kansas

May 15-18th, 2011

International Informix Users Group

www.iiug.org

More Musings on Informix 11.7

The best database keeps getting better…

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

Art Kagel
(art@advancedatatools.com) is principal database consultant with Advanced DataTools Corporation. Kagel manages the Informix consulting and support team, providing remote and onsite DBA support to clients around the world. He is a current member of the IIUG board of directors and an IBM Data Champion.

First of all I would like to welcome Art Kagel, who is co-writing with me for this issue. Art and I have worked together for many years and in January 2010, he joined Advanced DataTools to focus on Informix database development, support, and training. We worked together, evaluating and testing Informix 11.7 server (formerly code-named Panther) during the beta program. In this column, we wanted to share some more thoughts that we didn't have room for in the main article.

Optimizing star schema

Informix has always been a great datamart and data warehouse server. Four enhancements in 11.7 make Informix a real contender for larger warehouses and data vaults. Some of these enhancements are code merged from the old Informix Extended Parallel Server (XPS) engine, and some are new features. The new features are:

- Elimination of indexes on foreign keys
- Forest of Trees indexes
- Multiple index scans
- Star schema optimizer support

The elimination of indexes on foreign keys helps when you have a fact table with billions of rows and each added index would slow down inserts and loads. Forest of Trees indexes combine the advantages of a hash index with those of a traditional B+tree index when the first column of the index is not very selective. Multiple index scans in the optimizer are a real benefit for data warehousing, as they can reduce the number of long key compound indexes, which can also improve insert, delete, and update performance. These index scans also help online transaction processing (OLTP) databases and are described in much more detail in the main article.

However, we want to spend a bit more time on star schema optimizer support, which is an offshoot of XPS. In 11.7, when you have many dimension tables, you can create a single compound index on the fact table containing all of the dimension table keys. The optimizer will recognize this index and use the filters on the dimension tables to generate a list of valid combinations of the dimension tables' keys into a temp table. The optimizer will then join that temp table to the fact table using the compound index to quickly locate the rows that satisfy the criteria on the dimension tables. We have tested this one quite a bit and the results are impressive.

Finally, here's one we haven't tried yet: the documentation indicates that the star schema optimizer can also combine with multi-index scan technology if an appropriate compound index is not available. In this case, the optimizer uses another technology from XPS known as push-down hash joins to perform the query while examining fewer rows than would be possible using older complex query processing techniques. In this case, the first column of any index, simple or compound, is a candidate to satisfy the star join. Informix 11.7 contains an exciting new release of this feature for those of us who have been implementing Informix as a data warehouse server.

No more extent worries

There are two problems with tables that have lots of extents, or at least there used to be. Prior to 11.7, if your tables had close to 200 extents (the exact number varies based on page size, the number of "special columns," and the number of attached indexes in the table), you had to be concerned that you might fill the table's extent list and the table would no longer accept inserts or the expansion of variable length rows. That problem is gone. Informix 11.7 extends the table's extent capability to support more than 32,000 extents, making the extent list virtually unlimited due to the effects of extent doubling.

The second problem, of course, is that a table or an index with many highly active extents will suffer from performance degradation due to increased disk head movement and increased contention for head positioning with other objects in the database. Before the 11.50xC4 release, you had only a few options to reorganize a table with many active extents (and fewer for highly fragmented indexes), and all of them required that the table be taken offline for a time. You could unload the data, and then drop and re-create the table with a larger extent size; you could create a CLUSTER index on the table; or you could perform an ALTER FRAGMENT ... INIT on the table or index.

When the compression feature was added in 11.50.xC4, it brought along with it the REPACK and SHRINK SQL application programming interface (API) commands. This provided some improvement on the table reorg problem, because you could use these options to perform a partial reorganization in the background without taking the table offline. However, the REPACK option does not necessarily reduce the number of extents. It simply moves all of the data to the earliest extents possible in the existing extent list, and the SHRINK option will free up any extents that are unused after the REPACK (including any unused portion of the last partially used extent). Since most tables' largest extents are the last ones added and the smallest ones are the first added, the compression feature doesn't help much.

Informix 11.7 introduces the new DEFRAGMENT SQL API option, which moves data to new extents, creating contiguous extents that can be coalesced into a single larger extent wherever possible. This can all be done while the table or index is online. Since the engine performs this operation a page at a time as many very small transactions, very little contention is created. There is no possibility of a long transaction rollback, and best of all, since the pages are moved whole and the order of pages in the table doesn't change, each row's address does not change in the table—so no index updates are required. This operation is very fast, even compared to a REPACK.

At last, a true low-overhead online table reorg capability. Also, you can stop the operation at any time and restart it again later. Here's all you need to do to reorg your table, index, or table/index fragment into the minimum number of extents possible:

```
EXECUTE FUNCTION TASK( 'DEFRAGMENT', 'database:owner.tablename' );
```

or

```
EXECUTE FUNCTION TASK( 'DEFRAGMENT PARTITION', <partnum> );
```

Storage provisioning

We want to wrap up by touching briefly on the new storage provisioning features. We could go on for pages, but the important thing is that you can now configure your Informix instance so that you never again need to worry about a dbspace running out of storage space. You can configure it so that either existing chunks are extended, or so that a new chunk will be added to a dbspace when it fills.

You define a pool of storage, which can be RAW, COOKED, or filesystem files or even a directory name (the engine will create and allocate new files on its own within those directories). Once you have defined the storage pool, you can either use the pool when you manually add dbspaces or chunks or extend existing chunks, or you can define rules that the engine will use to automatically extend chunks or expand dbspaces with new chunks when the free space in the dbspace drops below the thresholds you define.

As John Miller III, STSM, embedability architect from IBM, says, "Informix will do it all while you are watching your favorite TV show!"

IBM invites the readers of *IBM Data Management* magazine to visit the Business Partner sites listed below.

Smarter is... Out of Chaos, Order

Lisa K. Stapleton *is a former programmer and currently a senior editor for* IBM Data Management *magazine.*

In the beginning, there was the data, and the data was good. But the data was alone, with neither hierarchy nor relation nor object. "Let there be a team," IBM said, and there was, and it was good. And IBM said, "Go forth, and produceth a system that will shepherd rocket parts, and truck parts, and every other thing that rolleth on the land, or zoometh through the air."

Well, maybe it didn't *exactly* happen that way, but IBM has been involved in some of the biggest database developments of the past 50 years. It all began in the era of punch cards and disk platters.

In the 1960s, the space program needed a way to track millions of parts. IBM's answer was IMS—a joint project with North American Rockwell—and a hierarchical structure that was perfect for representing the huge bills of material required.

"We had a model and a general idea of what we wanted to do," says Carl Chamberlin, one of the early team managers. "But I had no vision of where this was all going to go," he says. "We *never* expected to see multi-terabyte databases, or the phenomenal success of IMS, or that some IMS code from the 1970s would still be running unchanged."

What was it like to be present at the birth of modern databases? Back then, data wasn't considered an asset that could be mined—or even reused very much. "A file was created for a particular program and was rarely used for another program," says Pete Sadler, who first started teaching IMS in 1970 and who developed some features later incorporated into IMS.

Interactivity was also pretty rare. "IMS used to have an interactive query facility; it sort of came and went. There weren't that many people who thought they'd like to do dynamic queries," says Sadler.

And IBM saw that IMS was good. But it lacked a relational partner. Thus IBM created System R, System R begat DB2 and RDBMS, and IBM saw that all was good.

Why was System R chosen as the model for IBM's first relational database? "System R was fast, at a time when relational database advocates felt very vulnerable to the claim that SQL was never going to be fast enough," says Michael Blasgen, who did some of the early work on the System R project.

And new technology and new databases were fruitful and multiplied. Informix and solidDB took form and spread out upon the land, eventually to enter the realm of IBM.

Jonathan Leffler, a longtime Informix developer, remembers how different the technology world was then, and says that relational databases are now in an ideal position to exploit recent hardware improvements. "Now you have eight-core chips, and relational databases can really take advantage of that," says Leffler.

An important recent development has been databases that are optimized for near-real-time and embedded environments, where microsecond response times and very high availability are required. Ari Valtanen, one of the founders of Solid Information Technology—acquired by IBM in 2008—says that the availability of huge amounts of RAM will further transform databases.

And the data multiplied on the land, in the sea, and in the air. And the users of the end variety saw that it was good. ✱

dbi

Smarter Tools for Smarter People

In the time it takes to browse through this magazine, you could **download, install,** and **see remarkable breakthrough** performance **results** from our software

We're the world's fastest route to DB2 LUW productivity and performance. In fact, you can download, install, and actually see results in minutes instead of days, weeks, or months.

Only DBI offers all these reasons to be your smart choice for performance monitoring and tuning tools:

DBI	IBM OPM*	
☑	☐	Patented SQL Workload Cost Analysis
☑	☐	Tuning Results in 5 Mouse Clicks or Less
☑	☐	Automatically Tracks Database Changes
☑	☐	Trend Charts Plot Change Events on Graphs
☑	☐	Identifies SQL driving I/O to a Table
☑	☐	Summarizes Programs, Users, and their SQL Costs
☑	☐	Shows All Users that ran a particular SQL

Contact DBI to experience our results-oriented Proof of Concept process and discover why **Smarter People choose Smarter Tools** for DB2 performance and cost optimization. Act now and receive a free white paper by Scott Hayes, IBM DB2 GOLD Consultant, which discusses critical performance measurements and provides complimentary SQL commands.

www.DBISoftware.com/smarter

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

POWERFUL SCALABLE RELIABLE & EFFICIENT

IBM System x Servers and DS5000 Storage Systems for your database applications

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

They need:

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

IBM System x3690 X5

IBM System x3850 X5

IBM System Storage DS5300

IBM System Storage EXP5060

IBM System x Servers and DS5000 Storage Systems *Deliver*

IBM

For more information visit
www.ibmdatamanagement.com