

Quarter 3, 2006 Vol. 11, Issue 3

<http://www.dbmag.intelligententerprise.com/story/showArticle.jhtml?articleID=191502571>

## User Views

Read: [Quality Time](#) By David Beulke

Read: [Success Stories](#) By Stuart Litel

### Quality Time

**Follow through on best performance practices.**

By David Beulke

Today, database tables commonly contain billions of rows. At the same time, many businesses are transforming their processes for 24x7 availability. As a result, design, scalability, and recoverability are more important than ever.

Meeting database performance goals requires that you make the most of your procedures, existing database infrastructure, and new DB2 features. A mature corporate data-management group knows all the processes and procedures a successful database project requires.

But what happens when business needs require a schedule that doesn't allow time for all of these steps? You implement on the business's schedule, then figure out how to make the database design better with minimal coding changes. You can dramatically improve performance by using database table views, adding, combining, and splitting tables. Each of these approaches can improve performance, availability, and recoverability with a minimum of application impacts. Even when you have to bypass a few steps up front, don't let up on addressing the data issues after implementation. There's never time to do it right, it seems, but there's always time to do it over.

Instead of duplicating databases, you can extend them to satisfy that new 24x7 or Internet processing requirement. DB2 can handle very large transaction loads. You know your systems, data, and business processes inside and out. Use that knowledge to incorporate new data tables or columns for the new processing. Many times, new projects want their own copy of data that already exists. Avoid making unnecessary copies of critical data; the last thing a business needs is to duplicate its Sarbanes-Oxley compliance requirements or increase the exposure risks for sensitive data. These arguments are prime motivators for upper management, so use them to properly focus the discussion. Adding data columns or new tables can extend the database — a better choice than duplicating

the environment, which only gobbles up critical CPU capacity, disk requirements, and processing-window considerations.

DB2 version 9 for Linux, Unix, and Windows and DB2 version 8 for z/OS offer features that make it easier to incorporate these new systems and tables into existing databases. DB2 V8 for z/OS offers the ability to define 4,096 partitions, rotate them, and archive or cut old data partitions out. These capabilities allow the database designer to separate old data from new; isolating active data, improving buffer pool efficiencies, and using DB2 resources to address performance appropriately. Clustering and partitioning are now separate in DB2 for z/OS. As a result, designers have more options for grouping data and eliminating sorts.

DB2 9, previously known by the code name "Viper," offers new range-partitioning and row-compression features for improving database performance. You're probably familiar with the flexibility provided by UNION ALL views; range partitioning offers some of the same functionality. Range partitioning separates active data from older (or less-frequently used) data and includes an option for parallel application design.

Row compression in DB2 9 is similar to what's available in DB2 for z/OS. It offers potentially enormous I/O and CPU savings, because it enables five to 10 times the number of rows per DB2 page I/O. Although DB2 row compression may not be appropriate for high INSERT or UPDATE applications (because of the compression overhead), it makes a tremendous difference for the majority of business applications. I've used compression on many systems where it has saved CPU and I/O for all aspects of processing. Combining these new features within new systems or extending them into existing systems offers DB2 shops excellent options for meeting aggressive business schedules.

The International DB2 Users Group (IDUG) North America conference in Tampa, Fla., this spring was a great opportunity to learn about DB2 9 features. The upcoming IDUG conference in Vienna, Austria, in October promises more great presentations. Find out more at [www.idug.org](http://www.idug.org).

---

*[David Beulke](#), a past president of [IDUG](#), is a DB2 consultant, author, and lecturer who specializes in database performance, data warehouses, and Internet applications.*

---

## Success Stories

# Informix users tout its strengths.

## By Stuart Litel

**I recently had dinner with the CTO of a dot-com company that uses Informix. The CTO told me a story about how the Informix side of the company is profitable, does three times the number of transactions, and has enjoyed 98.7 percent uptime. When a visitor hits the company's Web site, the screen is painted in less than a second using Informix Dynamic Server (IDS) 10 on 14 Intel-based Linux servers.**

**My dinner companion's company was recently purchased by a business that uses 345 servers from a company whose headquarters is in the Pacific Northwest. Those 345 servers run the same company's database software, which processes only one-third the number of transactions that my friend's Informix system handles. To handle the number of transactions Informix handles on 14 servers, the company would need about**

**1,000 servers running that other database software.**

**On a recent trip, I met with the executive staff of a well-known company. This company uses IDS to run its main business, and the "O"ther database to handle its Web business. A member of the management team told me that they have fewer problems with the "O"ther database vs. IDS. I asked a few questions about the problems, and it came out that IDS is doing 10 times the volume of the "O"ther database.**

**That team would have learned something at the International Informix Users Group (IIUG) and International DB2 Users Group conference in Tampa, Fla., in May. The second joint conference was even more successful than the first, with more sessions, more attendees, and more activities. Among the highlights was a preview of some of the features we'll be**

**seeing in the next IDS release (scheduled for early 2007). These features include:**

- . Greater spatial and geodetic support**
- . Enhanced unstructured data support**
- . Greater HDR data protection**
- . Performance enhancements**
- . Enhanced system availability**
- . Reduced maintenance**
- . Configurable footprint**
- . Enhanced backup/restore**
- . More interoperability with IBM software.**

**Congratulations go out to IBM's John F. Miller for winning the Best Overall Informix Speaker Session and Bloomberg's Art Kagel for winning Best User Speaker.**

**The Annual IIUG Directors Award, given by the IIUG Board of Directors for best overall contribution to the Informix community, went to John F. Miller and**

## **James Edmiston.**

**If you missed this year's conference, plan to attend the event in Vienna, Austria, in October or the joint conference in San Jose, Calif., next May. You'll find the Informix presentations from the recent conference in the Software Download section of the [IIUG Member site](#).**

---

***[Stuart Litel](#) is CTO of [Kazer Inc.](#) and president of the [International Informix Users Group](#).***

[Return to Article](#)