



CDB Software, Inc.

Notes from IDUG 2006

July 18, 2006

Introduction

The annual IDUG conference was held this year in the convention center in Tampa, Florida. IBM continued to educate the market on DB2 V8, while finally whetting people's appetite about the newly announced V9. CDB has compiled some information that our attendees learned at all the various presentations that took place over the week. CDB has compiled some of the topics that were most important at IDUG. These include:

1. "DBA Beware! DB2's Increasing Limits,"
2. "XML takes Center Stage," and
3. "WARNING! Read this only if Authorized."

DBA BEWARE! DB2's Increasing Limits

DB2 V8 has increased the limits for many items:

1. 2MB SQL statements
2. 225 way joins
3. 128 byte table names
4. 32K character literals

The amount of data that can be stored has been increased with support for 128TB table spaces. This can be implemented with 4096 partitions and on top of this each partition can have multiple DPSIs. As users take advantage of these new limits, housekeeping will become much more of a headache with so many objects to copy, reorganize, etc. In DB2 V9 a partitioned table space is to be introduced which will add partitions automatically as it runs out of space. In the Keynote speech it was stated that one customer already had a table space over 23 TB in size!

Partly to allow for these large objects, DB2 V8 has 64 bit addressing which means that up to 16 Exabytes (you may never have even heard that

word before!) can be addressed. It also allows the use of large buffer pools up to a limit of 1TB. However a warning note was sounded that for performance, the buffer pools need to be backed by real storage of the same size. Unfortunately real storage is currently limited to 128GB due to a limitation in z/OS, so don't over commit your real storage with very large buffer pools; remember some real storage will also be needed by software apart from DB2 – what a waste!

DB2 V8 in conjunction with z/OS 1.5 increased the number of open datasets to 65,000. This will obviously help with the proliferation of a partitioned table space with high numbers of partitions and DPSIs. Opening and closing these datasets can take a significant amount of time when DB2 starts or stops. To address part of this problem, DB2 V9 will have an option to automatically open datasets for specified objects when initializing, thus helping to make DB2 available faster.

XML TAKES CENTER STAGE

In the Keynote speech it was stated that 40% of data is already in XML format compared to only 20% in DB2; thus it was not surprising to find that a significant part of DB2 V9 is extending the support for XML. Prior to V9 XML was supported through DB2 extenders, but V9 will integrate features into the engine: an XML data type, native storage of XML documents, XML schemas in the catalog, SQL support of XPath language, and much more.

As part of the support for XML, DB2 will introduce text search for XML (and LOBs, CHAR and VARCHAR). DB2 will utilize a text index server running on another platform which could be zLinux. TCP/IP communications will be established with OmniFind servers who will be passed the object and the search requirements.



OmniFind will return an index to the required text and this will persist in DB2 as a table. It seems an unusual implementation for a DB2 function, but time will tell!

WARNING! READ THIS ONLY IF AUTHORIZED

Security was another major theme for DB2 development due to the growth in legislation demanding security and auditing. This is cast in two flavors: security and encryption. The main thrust for security is row level security where each row contains a seclabel which determines who can see the row. Any utility externalizing the row, or changing it, needs to have a security level which allows it to see the row (e.g. UNLOAD). Those utilities which don't (e.g. COPY, RECOVER), do not have their row level security checked for each row as it is processed. One anomaly here is REORG; a normal REORG does not require row level security clearance, but if you want to DISCARD rows then it will.

Encryption has been around for some time, but IBM is now moving in a new direction. Instead of performing encryption using processor cycles in the mainframe, IBM intends to enhance their storage subsystems (DS8000) to perform the encryption, as they have powerful processors (P5) which have spare cycles that can be used. This will make encryption much more practical, and liable to be much more widely used.

About CDB Software

CDB Software, Inc. is a leader in data management solutions for DB2 z/OS. CDB focuses its business on DB2 for z/OS to provide unique and innovative solutions that enable companies to expand their DB2 system to meet business needs while controlling the overall cost of the mainframe. Founded in 1985, CDB is a privately held corporation based in Houston, Texas with offices worldwide.

For more information visit:

www.cdbsoftware.com



CDB Software, Inc.
11200 Richmond Ave.
Houston, TX 77082