

# Managing Large Systems with DB2 UDB

Chris Eaton  
IBM Toronto Lab  
1150 Eglinton Ave. E  
North York, Ont.  
(416) 448-2587

ceaton@ca.ibm.com

## 1. ABSTRACT

**In this talk, we will describe the usability challenges facing large distributed corporations. As well, we will discuss what IBM's DB2 Universal Database is doing to address these complex issues.**

### 1.1 Keywords

IBM DB2 UDB, usability

## 2. INTRODUCTION

The challenges that are effecting large distributed customers today with data warehousing go well beyond the performance and scalability issues of the database engine. Database administrators are faced not only with large complex databases but also with the issues of managing vast amounts of data, large numbers of users and large numbers of physical as well as logical objects on multiple platforms.

Clearly the challenges facing database administration tool vendors today is to mask the complexity of today's large systems, to automate administration tasks, and to remove the need to continually monitor the system. In this talk I will describe the direction that IBM is taking in each of these areas.

## 3. ANALYSIS AND AUTOMATION

Historically, database administrators have been heavily involved in object management. That is, maintaining tables, tablespaces, indexes, disk space, memory usage, backups, etc. However, this is not where administrators

can be of most benefit. By having tools create, schedule, perform and monitor these types of tasks, it leaves the administrators with time to perform analysis where their experience can be of most benefit. In order for database administration tools to perform these functions they must deal with several aspects of large systems. The following sections will describe how we can 1) mask the complexity of large systems, 2) monitor the system, and 3) provide the administrator with the tools and data necessary to review current as well as historical events that are effecting the system.

### 3.1 Complexity

As open systems grow in size, they also grow in complexity in terms of the number of objects to manage. Even if the database schema does not become complex, the sheer volume of data necessitates the need to manage not only larger numbers of disks but larger clusters as well. It is important for the database administration tools to mask this complexity by making object management location independent. That is, an administrator should not be required to know on what physical machine an object resides in order to perform an action on that object. DB2 UDB handles this by displaying the database objects as being node independent. An administrator does not need to know on what nodes a table exists to be able to alter it, drop it, load data into it, etc.

Complexity also encompasses the large number of tunable parameters that exist. These parameters are important to be able to customize a system to a customers particular environment but they also introduce complexity in the permutations and combinations that can be changed. To mask this complexity DB2 UDB provides a Performance Wizard that will ask specific questions and then make recommendations based on the administrators environment. However, wizards are not the ultimate solution especially for those administrators that we would consider to be "power users". For these types of users, we provide a notebook where all of the parameters can be reviewed/updated. But to assist users in this area, parameters are placed into logical groupings with context sensitive help easily accessible.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. To copy otherwise, to republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee.  
SIGMOD '98 Seattle, WA, USA  
© 1998 ACM 0-89791-995-5/98/006...\$5.00

### **3.2 Monitoring**

In order to provide administrators with more time to perform analysis, it is important to make monitoring as non intrusive as possible. In order to do this, DB2 UDB will provide "canned" monitor packages. What this means is that if you are concerned with deadlocks, you would start the deadlock canned monitor. This monitor will turn on the necessary switches and begin to gather the data elements necessary for the system to watch for deadlocks. When such an event occurs, the appropriate data elements are gathered, a notification is sent to the administrator and the analysis can take place at a future time. This frees the administrator from having to monitor the system manually, since he/she will be notified if an event occurs for which they wish to react to. Dozens of canned monitors are provided with the product but the administrator has the ability to create their own and distribute them throughout their enterprise.

### **3.3 Using historical and live data**

Capturing historical data gives the administrator/analyst the information necessary to find problems in the past as well as compare current information with historical. Another important aspect of providing historical as well as live monitor information is to allow the administrator to analyze the information to look for trends and to be able to

react to them. DB2 will allow you to store historical data for any monitored object.

### **3.4 Capacity planning**

Trend analysis also aids in the area of capacity planning. In order to maintain a large warehouse, administrators need the information that will allow them to plan for future growth. By providing DBAs with a historical analysis of their space usage, resource usage, etc. they are more prepared for future requirements and can act instead of having to react.

DB2 UDB also provides the information necessary for DBAs to perform impact analysis. That is, it give the administrator information related to particular actions. As an example, the administrator may want to know what views reference a specific table before dropping it.

## **4. CONCLUSIONS**

Given the ever increasing size of enterprise warehouses and datamarts, DBA tools vendors must be constantly improving their products to be able to mask the complexity of large systems, to automatically notify or take actions when specific events occur in the system, and to present data to the DBA in such a way as to allow them to perform analysis and take actions rather than forcing them to continuously manage objects.