

Author: Bill Ennis

Introduction

SQL Server 2000 is in the process of being released. This paper will introduce some of the enhancements that have been made – you can't use them if you don't know they're there! In addition I will mention some things that I gained by attending a recent seminar introducing SQL Server 2000 that may even be in version 7.0!

Resources

Let's start with a link to the presentation that I attended. You can find it at:

<http://www.microsoft.com/midwest/developer>

In addition a book was identified as a great XML resource. It is from Microsoft Press and named Step By Step XML.

Prerequisites and versions

Before you get started it makes sense to define the versions and platforms they are supported on. The information listed here is not exhaustive but should aid you in getting started. The versions of SQL Server 2000 that will be available include:

Version	Platform Support	SQL Server 2000 for Windows CE	Windows CE OS
----------------	-------------------------	--------------------------------	---------------

Desktop Engine (was MSDE)

All Windows NT 4 and Windows 2000, Windows Me and Windows 98

SQL Server Personal Edition (was Desktop)

2 CPU's, 1 max on 98	All Windows NT 4 and Windows 2000, Windows Me and Windows 98
	Standard Edition

Up to 4 CPU's, up to 2 GB RAM	Windows NT4 Server, Windows2000 Server, Advanced
-------------------------------	--

and Data Center
with limited license)
Windows NT4 and Windows 2000

Developer Edition (full Enterprise
All

Enterprise Edition

Unique Features include Partitioned Views, Log Shipping, Large Memory and SMP support, Multi-Language Pack, Advanced OLAP, Enhanced Failover Clustering Windows NT4 Server, Windows2000 Server, Advanced and Data Center

New Features

Scalability and Performance

Microsoft will provide the option to **scale up (shared everything architecture) or scale out (shared nothing architecture)** with the release of SQL Server 2000. They have reported that ON THE SAME hardware running NT version 4 and SQL Server 7.0 they were able to achieve a 37% increase in performance when upgraded to Windows 2000 and SQL Server 2000.

Setting up a clustered SQL Server is reported to be much easier in 2000. Keep in mind that you must have the appropriate hardware when setting up a clustered solution. In fact, you should notify the hardware vendor of your intention of using clustering when ordering the hardware so that the hardware vendor may suggest the appropriate setup for your needs.

SQL 2000 supports **online reorganization** of tables and indexes.

Multiple Instances of SQL Server can reside on the same machine. In addition SQL Server 2000 instances can reside on the same machine that is running a SQL Server 7.0 instance.

In certain hardware configuration SQL Server 2000 can address up to 64 GB of RAM.

Differential backup speed has been improved.

Support for “serverless” backup has been added to enable high-end disk storage solutions such as EMC to break a mirror and perform a full database backup in the time it takes to break the mirror (noted as less than a minute).

Indexes on computed columns have been added to the product. For example, if you frequently add or subtract two columns in a table you can now index on the result of this computation.

Index creation can now be performed **in parallel**.

Indexed views (or materialized views) are now available. This allows the realization or pre-stored result of a query. This is great in a heavy read environment but is not suggested in an OLTP environment.

Partitioned Views have been added to the product. A partitioned view is SQL Server's answer to Oracle's Partitioning and Informix's fragmentation. It basically allows you to distribute table data according to an algorithm. The view is basically a UNION ALL across underlying physical tables. In this scenario you would need to have a stored procedure or some application code that would interpret which physical table that data needs to be stored in when inserts occur. Updates also need to be aware. One suggestion pointed out was to use instead of triggers on the view and redirect to a stored procedure that puts the data in the correct place. All in all this one area where SQL 2000 is behind Oracle. Partitioning in Oracle does all of this for you. However, partitioning is an add one feature of Oracle.

Linked Servers have been improved. Linked Servers allow SQL Server instances to talk to each other. One often-used feature of Linked servers is inserting data from a database in one server by selecting the data out of another.

Backup and Recovery Enhancements

Application Point in Time Recovery has been added. This will allow an application to mark the transaction logs with an event that signals an important point in processing has been achieved or that the system is in a consistent state that may make sense to recover to at some point in the future. Before this only point in time recovery was possible which left you guessing what time to recover to.

Log shipping is now supported. The log shipping feature transports the transaction logs at periodic intervals from a “hot” server to a “warm-standby” server. This can be used in addition to replication and/or clustering for disaster recovery. Log Shipping is available in a service pack for 7.0 but was undocumented and had no GUI support. In 200 there is a wizard setup for log shipping.

Replication can now support ad/drop columns. This was a big headache previously since you needed to turn off replication when you applied any schema changes.

DTS supports save as Visual Basic so that you can maintain DTS jobs after creation.

A copy database wizard has been added. This can be used to copy or move a database from one SQL Server instance to another. It is also the suggested way to upgrade a 7.0 database to a 2000 database since a straight backup and recover will not create all database objects in 2000 format. The actual structure of database files is the same in 2000 which is why it is possible to restore a database that was backed up in 7.0 to a 2000 instance. Copy database also creates logins, stored procedures, and jobs associated with a database.

Recovery Models have been revamped. There are now 3 recovery models (simple, bulk-logged, and full). They replace the need to specify the truncate log on checkpoint and select-into bulk copy options of a database. The main difference is allowing select-into bulk copy in all online logging models without failure. In most situations you will probably want to run in full mode (or at least start out that way). Full Recovery Mode is also the default.

Application Hooks

Cascading updates and deletes are supported. This allows you to perform an update or

delete on a parent and have the changes ripple through the data model. A word of caution: With this feature it is possible to delete an entire database under the right circumstances. Be careful!

Bigint , variant, and table types have been added. Bigint is an 8 bytes number. Variants are like Visual Basic variants. And tables are a row/column structure that may be returned.

Documentation has been improved with 2000.

Query Analyzer now has a Transact SQL debugger. Previously this was only available from Visual Studio. Query Analyzer also has an object browser. You can define hot keys in the query analyzer; such as alt-I to perform a sp_lock procedure invocation.

Instead of triggers have been added. These allow you to provide an alternate action when DML is invoked on a table or view. I.e. when an insert is invoked on a view you can actually catch the request and redirect it to a stored procedure.

Data Access security can now be defined at the database, table, or column level. In addition C2 security support has been enhanced significantly. This is needed for US Government standards. This involves audit logging and translation of the audit logs.

User defined functions have been added. These are similar to stored procedures except you can use them within an SQL select statement. They can be very powerful. User defined functions can return a value or a table (new data type). Functions can be used in any expression. You can pass a table type to a function. Here's an example:

SELECT * from large_ship_orders(15)-> In this case we are passing 15 dollars as an argument that assists in defining a large ship order. The large_ship_order function encapsulates the user from the details. Behind the scenes you may actually construct a ship order from several tables if need be.

XML support has improved. The SQL select construct has added a FOR XML clause that effectively returns an XML document. There are also some utilities being developed to perform bulk loads with XML. I did not see any support for exporting DTD's or XML schemas that correspond to query results that return XML (this would be extremely valuable for information transfer). In addition, data is returned in a format that comprises of an entity per row with column values defined as entity attributes. This seemed quite limited. There may be a way to override this but it was not pointed out at the Summit.

IIS hooks have been added for supporting database access via http. You can now perform a query through IIS through a URL that points to a SQL Server database using IIS for authentication.

OLAP – they talked about OLAP a bit. It looks as though it has been improved quite a bit. I have not used these features so I will not comment since I don't have much opinion. They did have a couple of vendors with some interesting products (**knowsys and maxim software**). These products allow complex data analysis and reporting.

Miscellaneous

I queried the presenter a couple of times regarding some of the nuances of SQL Server either experienced by myself or others. Here were the questions and responses.

Question	Response
	We have had some problems copying large data strings from the clipboard into varchar(8000) database columns.

You need to upgrade your OLE DB or ODBC Driver. Try upgrading to MDAC 2.5 or 2.6.

Are there any plans to implement row level triggers?

No, due to implementation details this will be very difficult to get into SQL Server. I.e. not coming any time soon.

Conclusion

Overall, I was impressed with the feature set of SQL Server 2000. Microsoft has made significant progress with the 2000 release and is catching up with Oracle in many high-end areas (if they would only supply a Unix port!). I actually inquired on the possibility of a Unix port; don't hold your breath!