
What's New In SQL Server 2000?

Brian Knight, MCSE

wrox

Topics to be Covered

- Release goals
- Server performance and scalability
- Multiple instances
- Tool enhancements
- New Data Transformation Services (DTS) features
- Relational engine
- Security

WFO

Release Goals

- Increase scalability and reliability
- e-Commerce and Internet support
- Easiest to build and manage DBMS
- Keep dominance in usability market
 - Improve Data Transformation Services (DTS)
 - Make programming easier
- Continue to dominate cost/performance market

WFO

Performance

- DBCC now uses parallel processors
 - Uses row locks (SS7 took table locks)
 - Can use object names (not ID's)
- Improved scans
 - Better read-ahead utilizing all processors
- Merry-go-round scans
 - When user 1 begins a table scan and user 2 begins the same scan, user 2 "piggybacks" on the same scan
 - User 2 finishes what he missed when he completes the first scan

Indexes Improvements

- Index creations
 - Use the TempDB to better performance
 - Uses parallel processors
- New DBCC command (**DBCC INDEXDEFRAG**)
 - Defrags internal and external fragmentation
 - Reorganizes leaf nodes based on logical order
 - Quickly uses page locks
 - Much faster than a rebuild of the index and online operation
 - Use **DBCC SHOWCONTIG** to see if needed

Scaling-Up

- Dynamic asynchronous I/O
 - Increases automatically with demand
 - Can no longer configure this option
- Support for 64 GB of RAM and 32 processors
 - With Windows 2000 DataCenter
 - Enterprise Edition of SQL Server only
- SAN Support
 - Great for high bandwidth applications

Indexed Views

- Includes ability to put indexes on computed columns
- Content of view is automatically persisted and updated
- Great for month-end data warehouse crunch time
 - Can include join and aggregate info
 - Great for data warehouse loading
- **SCHEMABINDING** must be turned on
- Performance hit and DB size grows as view is "materialized"
- Can't reference other views
- Must use 2-part names for tables referenced

Distributed Partitioned Views (DPV)

- Allows many servers to appear as one (federated)
- Down payment in SS7 with linked servers
- Best scale-out feature in this release - improving with Yukon
- Must have a uniform primary key
 - State, Country, Year, for example
- Uses check constraints
- Constraint must use part of the primary key
- Built on OLE-DB
 - Query optimizer has been improved to handle DPVs
 - Key T-SQL syntax is **UNION ALL**

Multiple Instances

- Allows you to run multiple copies of SQL Server on one server
- Performance Monitor and Event Viewer instance aware
- Ideal for development environment
- Ideal for application hosting companies who need individual billing based on CPU
- Always one default instance plus up to 16 additional instances
- If you install SS2K on a SS7 server, the SS7 is the default
- Only use when needed – each instance takes the same amount of resources

Tool Enhancements

- Tools are Active Directory aware
- Improved Index Tuning Wizard
- Much more usable English Query and Full Text Search
- Much better Query Analyzer
 - Object browser and search
 - Query templates that can be customized
 - Improved tracing and Showplans
 - Stored Procedure debugger
 - Drag-and-drop queries
 - Multiple result sets in a grid
 - Bookmark your code
 - Shortcut queries

- Create a sample template
- Use the template
- Perform a drag and drop query from browser

wpro

Data Transformation Services

- Much of the programming needed in SS7 no longer needed with the Dynamic Properties Task
- New tasks, MSMQ, FTP, Dynamic Properties, Copy Database
- Old tasks improved: multiphase data pump
- Pass packages parameters
- Input and output parameters now available
- Can now save packages as VB files
 - Great way to learn the object model
- See "Fundamentals of DTS" presentation for details

User Defined Functions

- Now you can create your own functions and call them from T-SQL
- Future releases expect to see VBScripting abilities
- Can act as a "parameterized" view
 - In-line functions return table record sets
- Scalar-valued
 - Return a single value based on the input value
 - Select **functionname(column1)**

Cascading DRI

- Automatically cascade updates and deletes from Primary Key to Foreign Key tables

```
CREATE TABLE region (  
    Region_Name Varchar(50) NOT NULL  
    PRIMARY KEY)
```

```
CREATE TABLE employee (  
    EmployeeID int,  
    Territory Varchar(50) NOT NULL  
    REFERENCES region  
    ON UPDATE CASCADE  
    ON DELETE NO ACTION)
```

Instead Of and After Triggers

- **INSTEAD OF** triggers fire in the place of the normal triggering action
- **AFTER** triggers execute after the row is inserted
- Can be placed on a view or table
 - Allow views to be updated
- Can reject part of a load but not fail the entire load

WFO

Instead of Code Example

```
CREATE VIEW AllOrders AS
    Select OrderID, Amount, CustomerID, Region
    FROM OrderEast
UNION ALL
    Select OrderID, Amount, CustomerID, Region
    FROM OrderWest

CREATE TRIGGER IO_Trig_INS_AllOrders ON AllOrders
INSTEAD OF INSET AS
BEGIN
    INSERT INTO OrderEast
        Select OrderID, Amount, CustomerID, Region
        FROM inserted WHERE Region = 'East'
    INSERT INTO OrderWest
        Select OrderID, Amount, CustomerID, Region
        FROM inserted WHERE Region = 'West'
END
```

Collation Changes

- Previously, you couldn't restore a database with a different collation onto your SQL Server
 - You would have to rebuild the Master database
 - In SS2K, this is no longer a problem
- Can now set collation at the column or database level
- Great for multi-lingual applications
- Code example:

```
CREATE DATABASE Wrox COLLATE  
Latin1_General_CS_AI
```

```
CREATE TABLE Orders (OrderID int,  
                      OrderRegion Varchar(10) COLLATE  
French_CS_AS)
```

New Data Types

- **sql_variant** allows you to store any data type except text, ntext, image, timestamp
 - Great for storing metadata
- **bigint** supports 8 byte integers
 - Use when you need numbers larger than **int** can support
 - Most functions still return **int** values unless noted
 - **COUNT_BIG** is a new function that uses **bigint**
- **table** holds row sets temporarily until needed
 - Similar to a temp table but faster

Table Data Type Example

```
DECLARE @VariableName TABLE
    (columna int PRIMARY KEY, columnb varchar(20))

INSERT INTO @VariableName VALUES (1, '123 Main St')
INSERT INTO @VariableName VALUES (2, '55 Riverside Av')

SELECT * FROM @VariableName

GO
```

XML Integration

- Access "self-defining" data easily on the web
- Support for
 - Structured
 - Semi-structured
 - Unstructured
- UpdateGrams will be a web release after RTM
- Can perform direct queries or use templates
- See the "Using XML and SQL Server 2000 Integration" presentation for details

Log Shipping

- Incrementally creates a backup of transaction log and applies it to another server
- Allows you to create
 - A warm standby server or servers
 - A read-only reporting database
- Technique existed since early releases but now is automated
- Available for Enterprise Edition only
- Use Maintenance Plan Wizard to setup and manage with Log Shipping Monitor

Security

- Kerberos aware
- Better out of the box security
- Better auditing
 - Done with SQL Profiler
 - Support for log file rollover
- C2-level security
 - Must be enabled with **sp_configure**
 - Uses profiler to audit
- Role based security
 - New BCP role
- SSL communication encryption

SQL Server Resources

➤ Web

- <http://www.microsoft.com/sql>
- <http://www.swynk.com>
- <http://msdn.microsoft.com/sql>

➤ Newsgroups

- microsoft.public.sqlserver.server
- microsoft.public.sqlserver.tools

Questions?

Brian Knight, MCSE
bknight@sqlservercentral.com

WROX