

Microsoft® **SQL Server™** 2008

Code Name: "Katmai"

An Overview of Features in the Next Release of SQL Server

Chad Boyd

chad.boyd@gmail.com

<http://blogs.msdn.com/chadboyd>

<http://www.mssqltips.com>

What's old?

- Getting more serious about deprecation policy
- Deprecation announcements since SQL 7.0
 - Formalized policy in SQL Server 2005
- “Katmai” implements the policy
 - First CTP
- Examples:
 - 60/65/70 compatibility levels
 - BACKUP LOG WITH NOLOG/TRUNCATE ONLY
 - sp_addalias, sp_addgroup, etc.
- For future deprecation, we'll help you detect usage

Microsoft® Introducing: SQL Server™ 2008

Enterprise Data Platform

- Secure, trusted platform for your data
- Optimize and accelerate your system performance

Enterprise Data Platform

- Productive policy-based management of your infrastructure



Dynamic Development Beyond Relational

- Accelerate your development with entities
- Synchronize your data from anywhere



Dynamic Development Beyond Relational

- Store and consume any type of data
- Deliver Location Intelligence within your applications



Pervasive Insight *Pervasive Insight*

- Integrate all your data in the Enterprise Data Warehouse
- Reach all your users with scalable BI platform
- Empower every user with actionable insights



Secure, Trusted Platform For Your Data



Transparent Data Encryption

External Key Management

Data Auditing

Pluggable CPU

Enhanced Database Mirroring

Extended Events Tracing

● **Protect your information**

- Encrypt your data without requiring an application re-write
- Consolidate security keys within the data center
- Integrated auditing support

● **Increase reliability of applications**

- Add system resources without affecting your users
- Leverage database mirroring to increase reliability
 - Automatic page repair
 - Improved performance with log record flush (leveraging new compression algorithms)

Transparent Data Encryption

- Encryption *with* application transparency
 - Expands the SQL Server 2005 encryption offering
- Database level scope
 - Introduces Database Encryption Key (DEK)
- Data at rest protection
 - Backups are encrypted optionally

```
CREATE DATABASE ENCRYPTION KEY  
    WITH ALGORITHM = AES_128  
    ENCRYPTION BY SERVER CERTIFICATE ent_cert
```

```
ALTER DATABASE <database_name>  
    SET ENCRYPTION {ON | OFF}
```

Auditing

- From building blocks in SQL Server 2005 to out-of-the-box solution in SQL Server 2008



- AUDIT is a first Class Server Object
- Granular audit actions on objects and/or users
- Multiple outputs (File, Windows Application Log, Security Event Log)
- High Performance – based on Extended Events
- Built in tools for consolidation, support for AS/RS

Auditing

```
CREATE AUDIT HIPAA_Audit  
TO FILE  
(  
  FILENAME=' \\PRO1\Aud\HIP_ADT.aud' ,  
  MAX_SIZE=100 MB,  
  RESERVE_DISK_SPACE  
)  
WITH (SHUTDOWN_ON_FAILURE = ON);
```

```
CREATE AUDIT SPECIFICATION  
SvrAC  
ON SERVER  
TO HIPAA_Audit  
  ADD FAILED_LOGIN_GROUP;
```

```
CREATE AUDIT SPECIFICATION  
AuditAC  
ON DATABASE  
TO HIPAA_Audit  
  ADD SELECT ON  
  table::Customers(payment);
```

Extended Events (xEvt)

- High performance lightweight tracing infrastructure
 - Easier to instrument
 - Suitable for production environments
 - Reduced troubleshooting time
- Integrated with ETW (Event Tracing for Windows)
 - Enables activity correlation at unprecedented levels

Provides the foundation for other high performance tracing subsystems in SQL Server (e.g. Audit)

Lays the foundation for key diagnostic and supportability enhancements for troubleshooting SQL Server

Optimized And Predictable Performance



Data Compression

Backup Compression

Performance Data Collection

Performance System Analysis

Filtered Indexes

Plan Freezing

Resource Governor

- **Better Performance**
 - Increase query and backup performance
 - Reduce your storage costs
- **Optimize System Performance**
 - Monitor your system by collecting and storing system events
 - Analyze and report on your system performance
- **Predictable Performance**
 - Manage concurrent workloads on single system
 - Prevent runaway queries
 - Ensure consistent query performance
 - Easy-to-use plan guides

Compression

- Data and Backup Compression
 - Different implementations and usage
- Data compression
 - Basic approach is to store data more efficiently
 - 2 Levels of compression (ROW and PAGE)
 - Smaller in-memory footprint one of many additional benefits besides on-disk savings
- Backup compression
 - In addition to data compression (can benefit from using data compression as well) – Traditional Compression
 - New sp_configure option to control default
 - ON or OFF via new BACKUP option

<http://blogs.msdn.com/chadboyd/archive/2007/07/28/katmai-sql-2008-data-compression-including-backup-compression.aspx>

Data Compression

Dateld	CarrierTracking	OfferID	PriceDisc
20070601	4911-403C-98	10	0.00
20070601	4911-403C-99	10	0.00
20070602	6431	10	0.00
20070602	6431-4D57-83	10	0.00
20070602	6431-4D57-84	10	0.00
20070602	6431-4D57-85	10	100.00
20070603	4E0A-4F89-AE	10	0.00



Dateld	CarrierTracking	OfferID	PriceDisc
20070601	4911-403C-98	10	0.00
20070601	4911-403C-99	10	0.00
20070602	6431	10	0.00
20070602	6431-4D57-83	10	0.00
20070602	6431-4D57-84	10	0.00
20070602	6431-4D57-85	10	100.00
20070603	4E0A-4F89-AE	10	0.00

Row Compression

- SQL Server 2008 row compression for all fixed-length data types
 - e.g. int, bigint, etc.

Data Compression

Dateld	CarrierTracking	OfferID	PriceDisc
20070601	4911-403C-98	10	0.00
20070601	4911-403C-99	10	0.00
20070602	6431	10	0.00
20070602	6431-4D57-83	10	0.00
20070602	6431-4D57-84	10	0.00
20070602	6431-4D57-85	10	100.00
20070603	4E0A-4F89-AE	10	0.00

Dateld	CarrierTracking	OfferID	PriceDisc
1 2007060	2 4911-403C-9	3 6431-4D57-8	
1 1	2 8	10	0.00
1 1	2 9	10	0.00
1 2	3 4	10	0.00
1 2	3 3	10	0.00
1 2	3 4	10	0.00
1 2	3 5	10	100.00
1 3	4E0A-4F89-AE	10	0.00

Prefix Compression

- A prefix list is stored in the page for common prefixes
- Individual values are replaced by
 - Token for the prefix
 - Suffix for the value

Data Compression

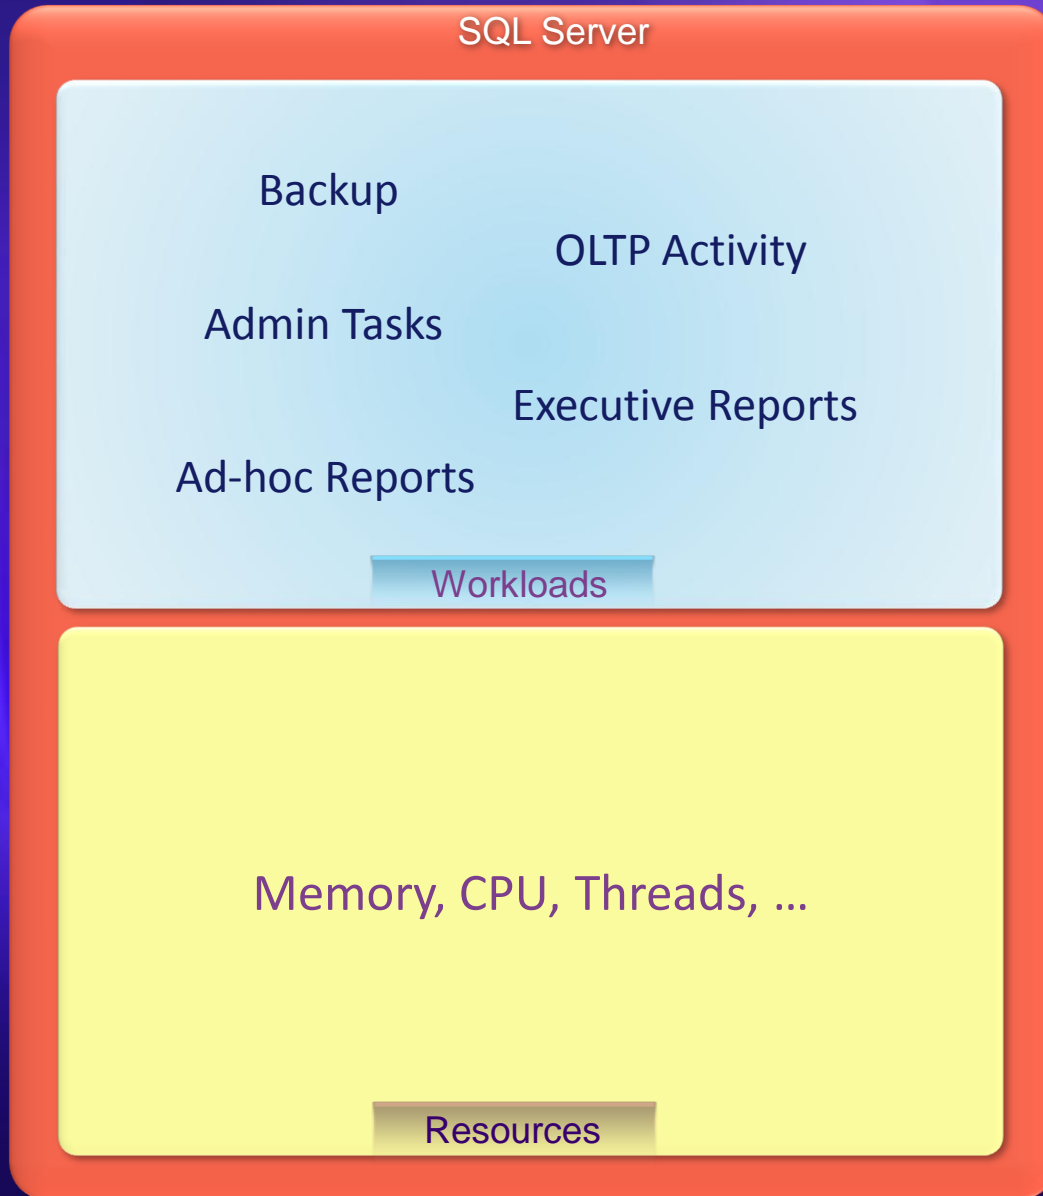
Dateld	CarrierTracking	OfferID	PriceDisc
1 2007060	2 4911-403C-9	3 6431-4D57-8	
1 1	2 8	10	0.00
1 1	2 9	10	0.00
1 2	3 4 →	10	0.00
1 2	3 3	10	0.00
1 2	3 4	10	0.00
1 2	3 5	10	100.00
1 3	4E0A-4F89-AE	10	0.00

Dateld	CarrierTracking	OfferID	PriceDisc
1 2007060	2 4911-403C-9	3 6431-4D57-8	
1 1	2 2	3 10	4 0.00
1 1	2 8	3	4
1 1	2 9	3	4
1 2	3 4 →	3	4
1 2	3 3	3	4
1 2	3 4	3	4
1 2	3 5	3	100.00
1 3	4E0A-4F89-AE	3	4

Dictionary Compression

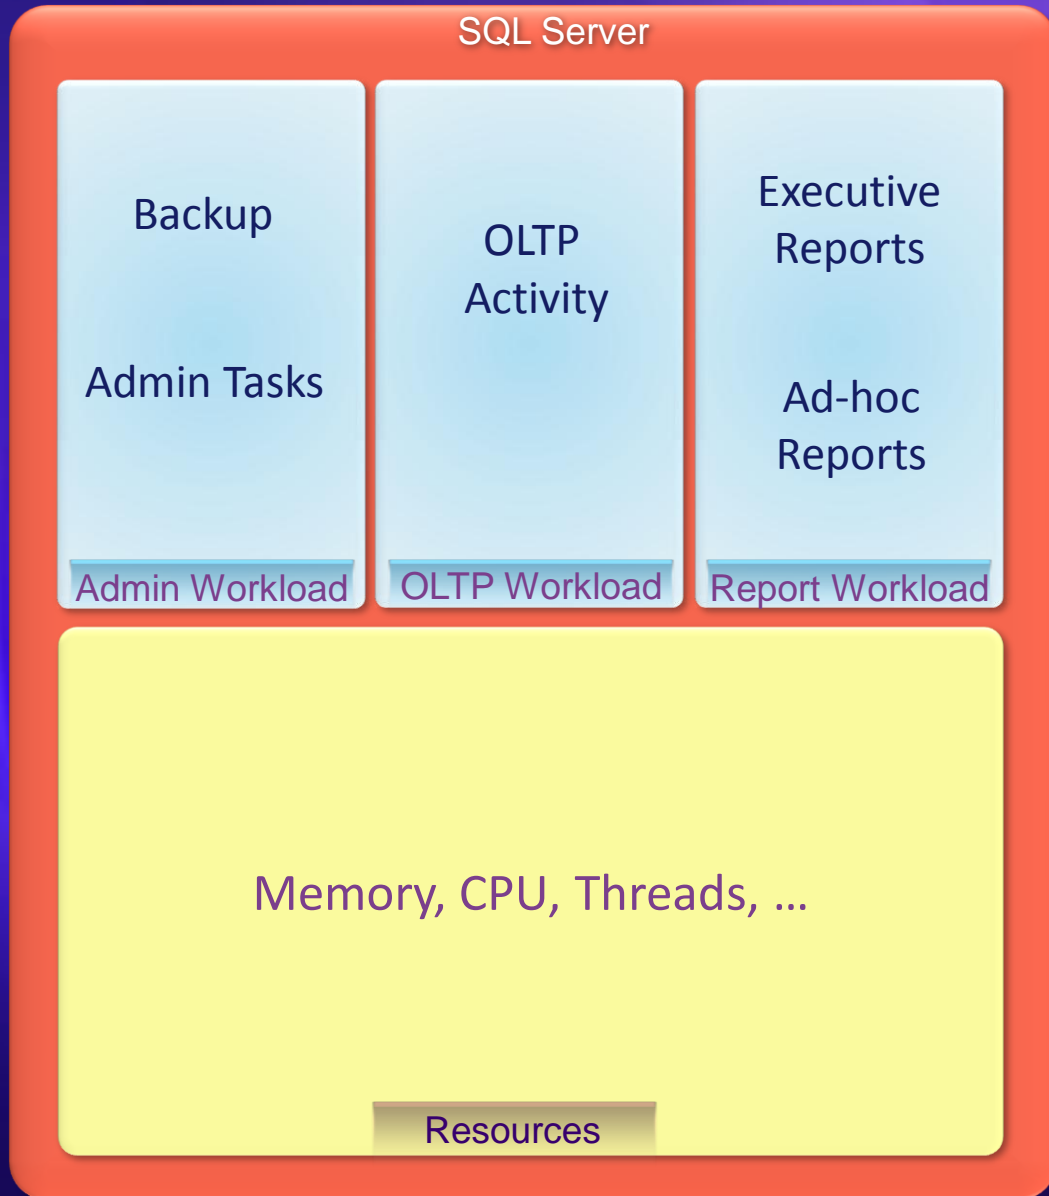
- A common value dictionary is stored in the page
- Common values are replaced by tokens

SQL 2005 Resource Management



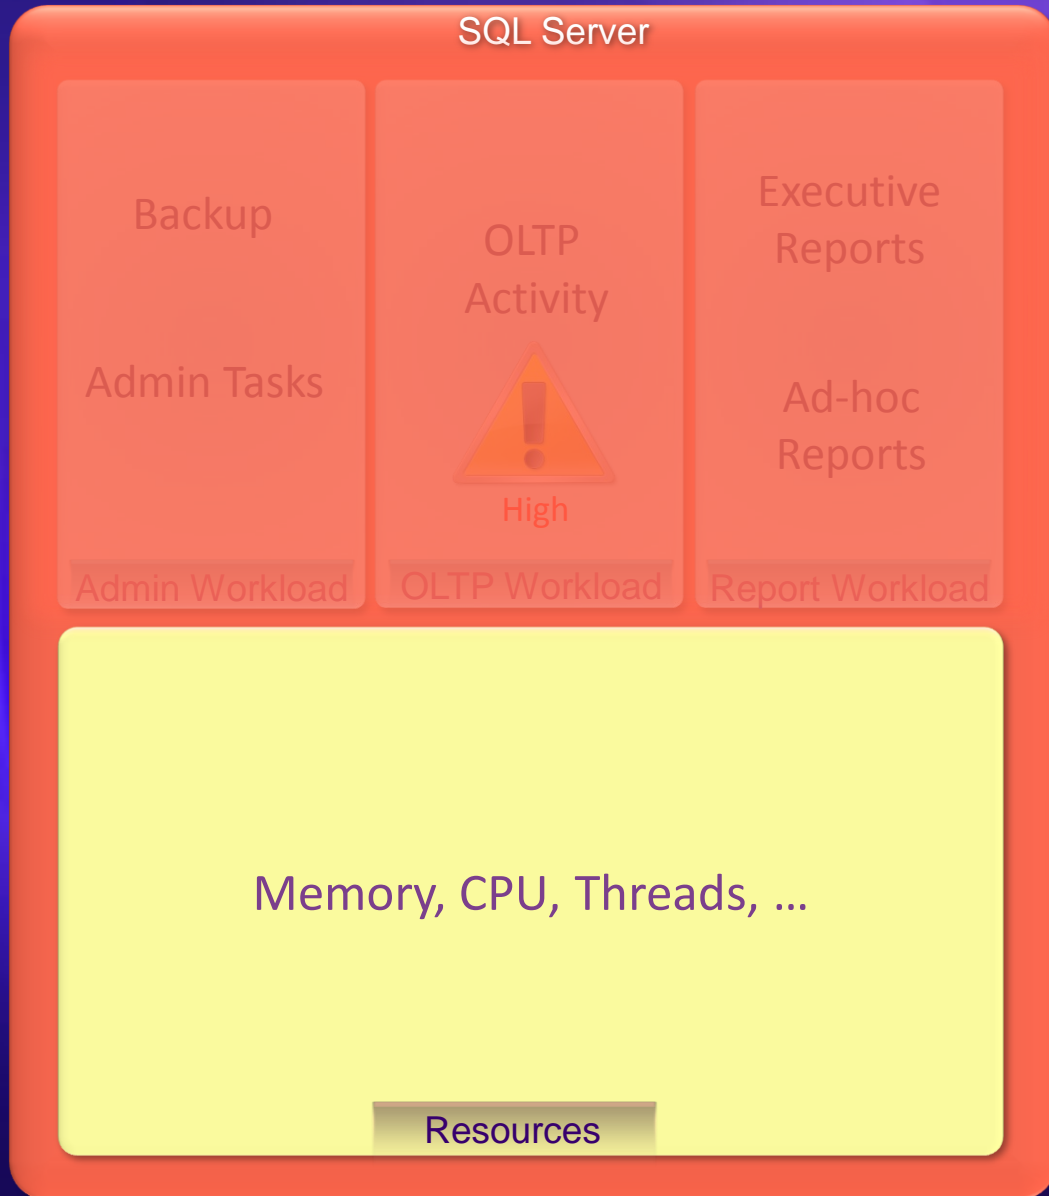
- Single resource pool
- Database engine doesn't differentiate workloads
- Best effort resource sharing

Resource Governor – Workloads



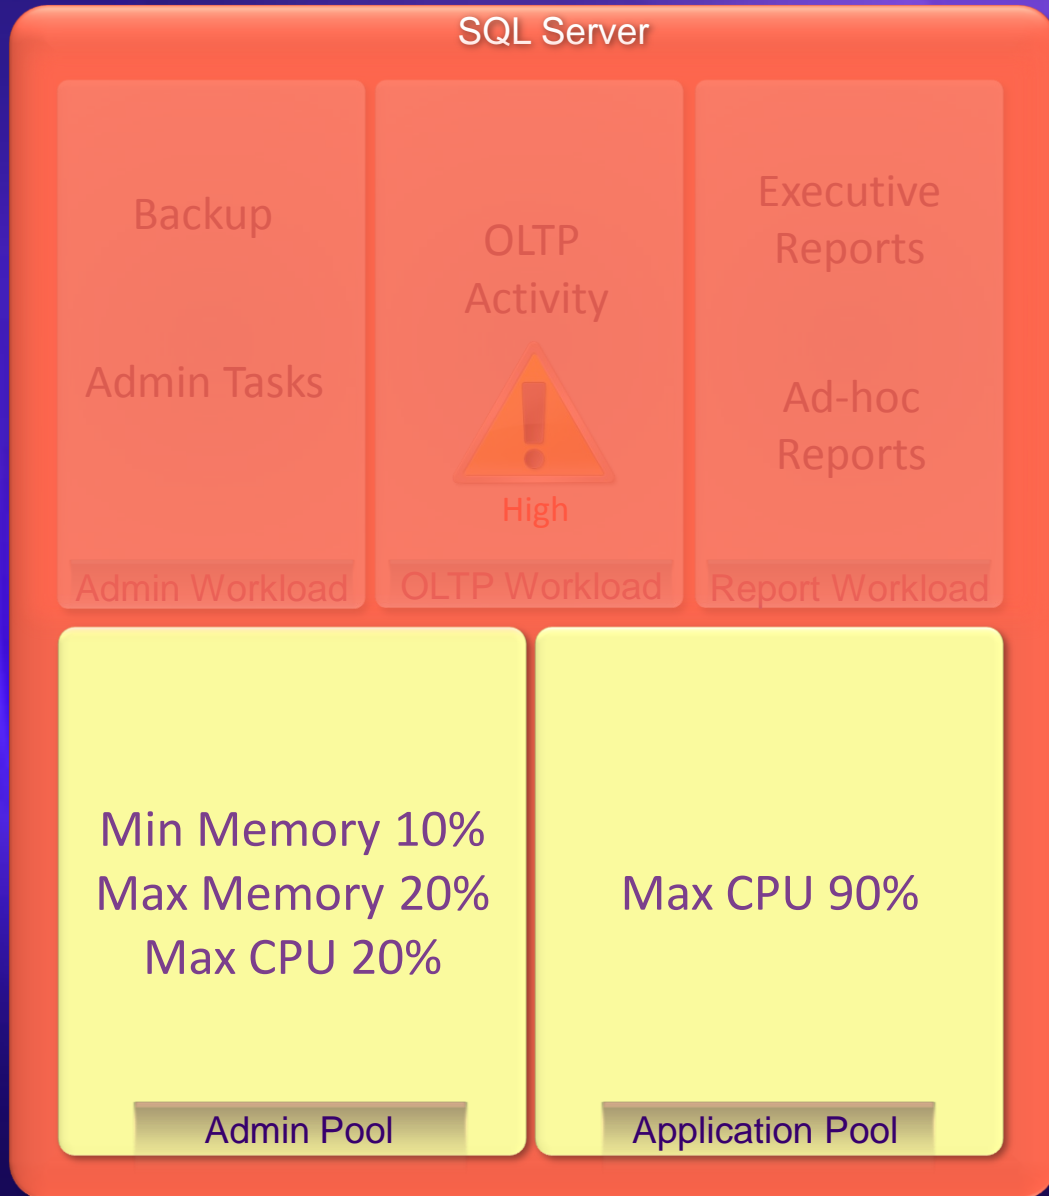
- Ability to differentiate workloads
 - e.g. app_name, login
- Per-request limits
 - Max memory %
 - Max CPU time
 - Grant timeout
 - Max Requests
- Resource monitoring

Resource Governor – Importance



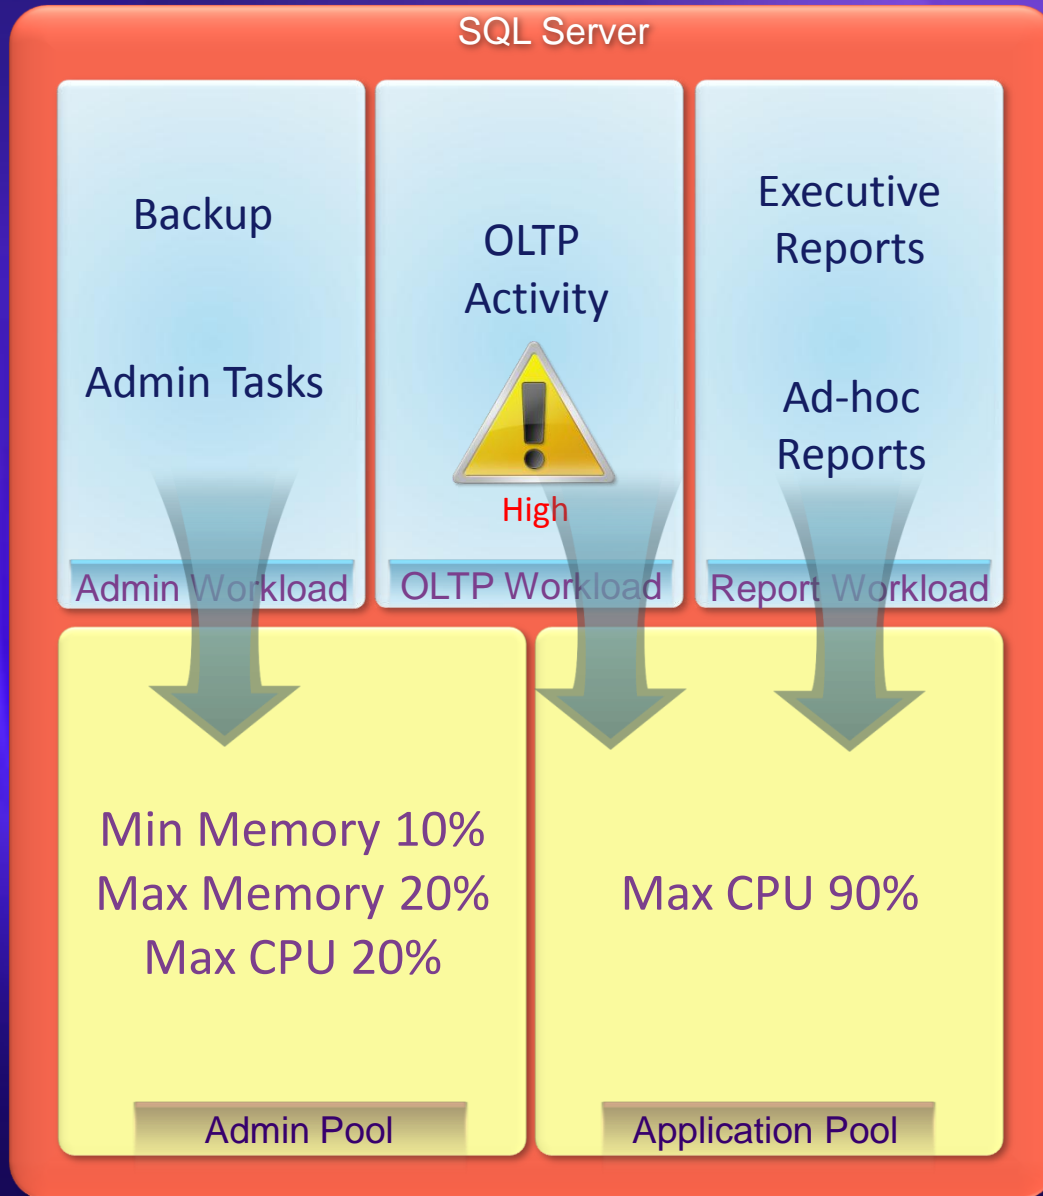
- A workload can have an importance label
 - Low
 - Medium
 - High
- Gives resource allocation preference to workloads based on importance

Resource Governor – Pools



- Resource pool: A virtual subset of physical database engine resources
- Provides controls to specify
 - Min Memory %
 - Max Memory %
 - Min CPU %
 - Max CPU %
 - Max DOP
- Resource monitoring
- Up to 20 resource pools

Resource Governor



Putting it all together

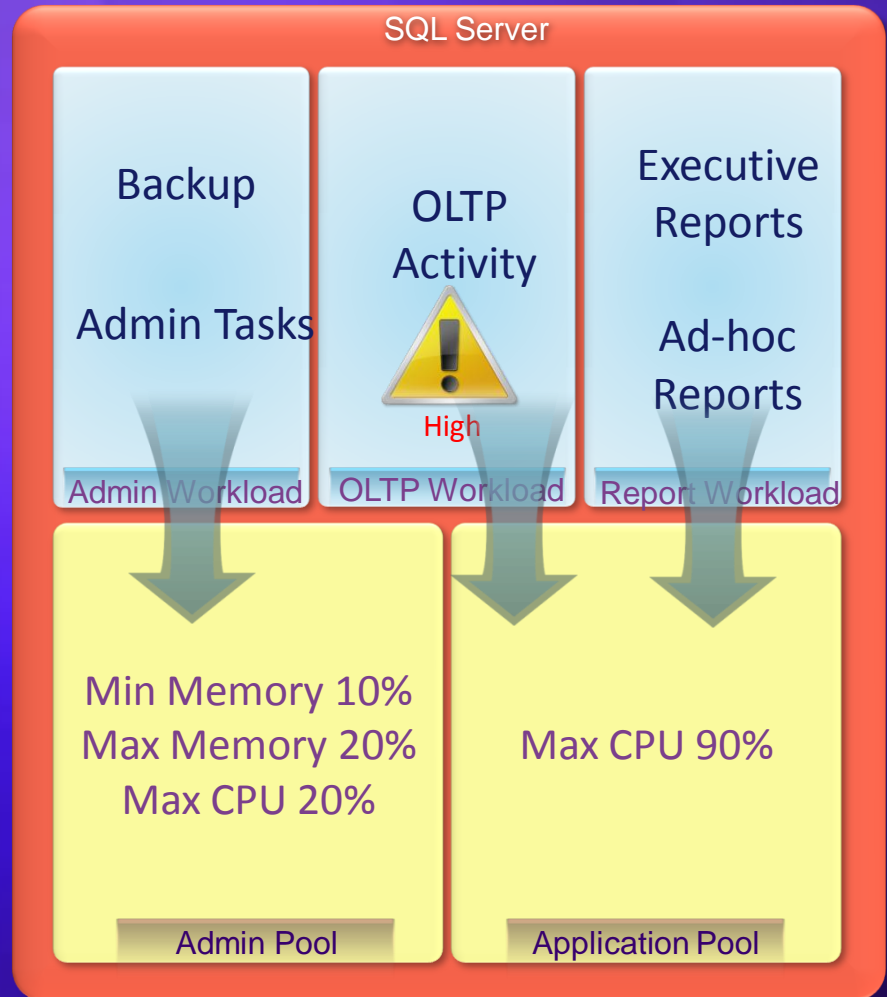
- Workloads are mapped to Resource Pools (n : 1)
- Online changes of groups/pools
- SQL Server 2005 = default group + default pool

Main Benefit

- Prevent run-away queries

demo

Resource Governor




```
use master
```

```
GO
```

```
-- Create pool for monitoring
```

```
CREATE RESOURCE POOL pSlow
```

```
WITH (MAX_CPU_PERCENT = 100)
```

```
CREATE RESOURCE POOL pFast
```

```
WITH (MAX_CPU_PERCENT = 100)
```

```
-- group definitions
```

```
CREATE WORKLOAD GROUP gSlow
```

```
USING pSlow
```

```
CREATE WORKLOAD GROUP gFast
```

```
USING pFast
```

```
GO
```

Messages

Command(s) completed successfully.

Query executed successfully.

Ready

Another computer is nearby
Click here to send files to the computer JHOLT2.

(54) master 00:00:00 0 rows
Ch 1 DMS

Query Plan Freezing

- SQL Server 2005 introduced plan guides and USE PLAN hint
- SQL Server 2008 enhances plan guides
 - Easier creation based on plan cache entries
 - DML support
- Many use cases
 - Smooth Upgrades
 - Disallow plan changes for critical plans
 - Plan fixing for ISV applications
 - Stability between test and production systems

Filtered Indexes

● Scenarios

- Indexing type specific/sparse properties
- Indexing Active/Hot data
- Indexing data in specific partitions

● Examples

- Product Catalogs
 - Product specific properties
- Data Warehouses
 - Sales data for the current month
- Inventory systems
 - Index items not marked as 'Deleted'

Filtered Indexes

pk	c1	sc 1	sc 2	sc 3	sc 4	sc 5	sc 6	sc 7	sc 8	sc 9
1	A	1								9
2	B		2		4					
3	C						6	7		
4	D	1				5				
5	E				4				8	
6	F			3						9
7	G					5		7		
8	H		2						8	

Filtered Indexes

pk	c1	sc 1	sc 2	sc 3	sc 4	sc 5	sc 6	sc 7	sc 8	sc 9
1	A	1								9
2	B		2		4					
3	C						6	7		
4	D	1				5				
5	E				4				8	
6	F			3						9
7	G					5		7		
8	H		2						8	

- **Create index i1 on T(sc1) where C1=A or C1=D;**
select sc1 from T where c1=A and sc1>5

Filtered Indexes

pk	c1	sc 1	sc 2	sc 3	sc 4	sc 5	sc 6	sc 7	sc 8	sc 9
1	A	1								9
2	B		2		4					
3	C						6	7		
4	D	1				5				
5	E				4				8	
6	F			3						9
7	G					5		7		
8	H		2						8	

- **Create index i1 on T(sc1) where C1=A or C1=D;**
select sc1 from T where c1=A and sc1>5
- **Create index i2 on T(sc7) where sc7 IS NOT NULL;**
select sc1 from T where sc7=5

Productive Policy-based Management of your infrastructure



Declarative Management Framework

Server Group Management

Streamlined Installation

Enterprise System Management

- **Spend less time on ongoing operations**

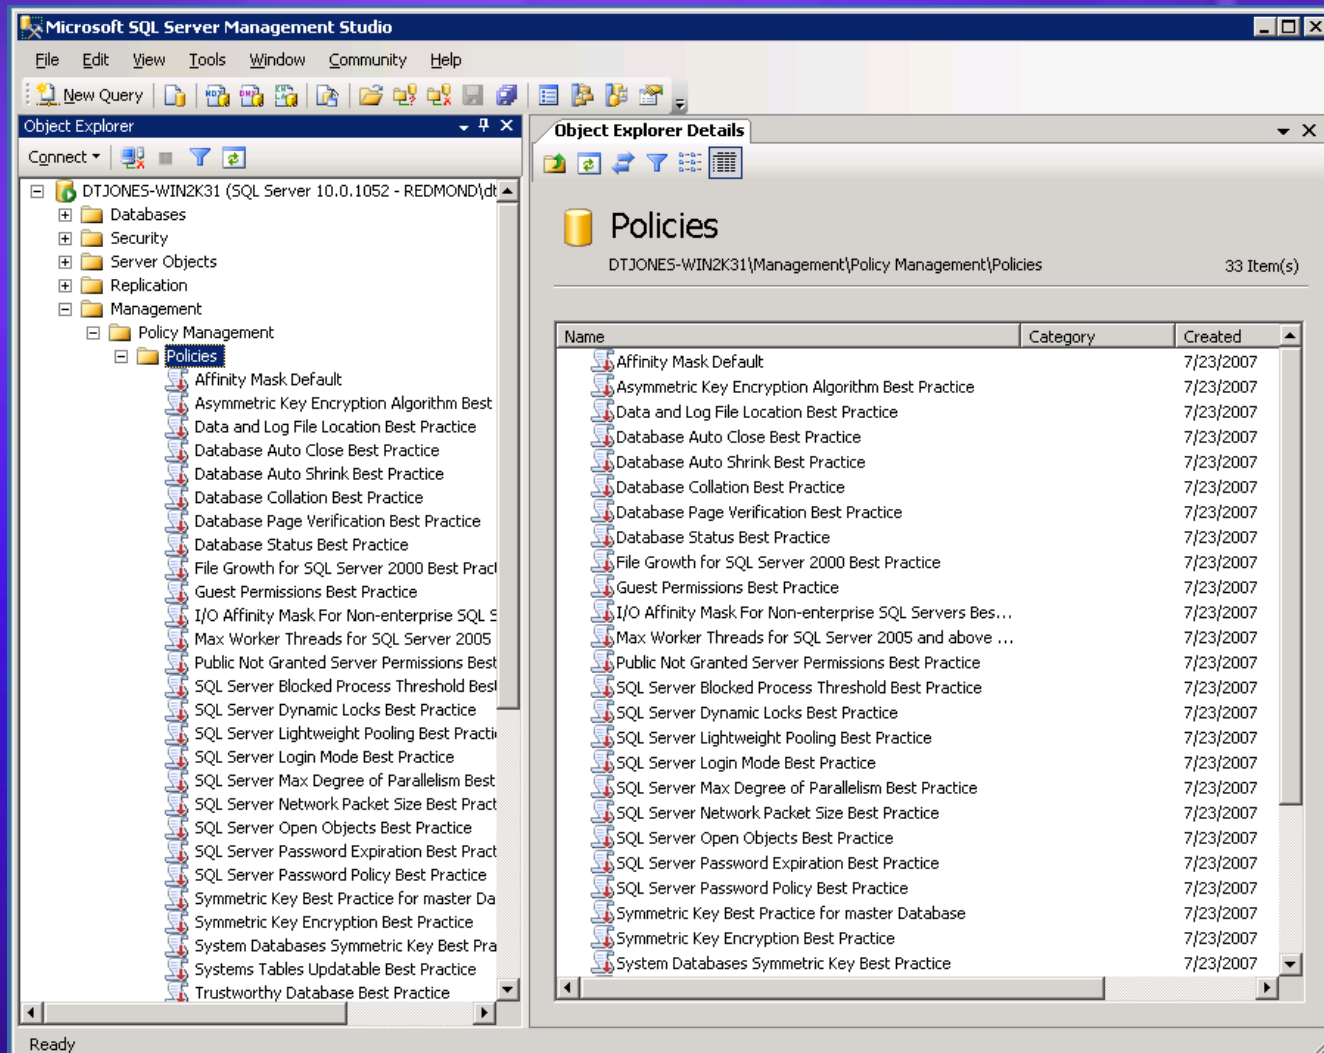
- Manage via policies instead of scripts
- Define Enterprise wide data management policies
- Automated monitoring and enforcement of policies
- Simplify your installation and configuration

- **Integrated with your enterprise system management**

- Define Policies that are compliant with System Definition Model
- Manage your data and system infrastructure with Microsoft System Center

demo

DMF



Microsoft® Introducing: SQL Server™ 2008

Enterprise Data Platform

- Secure, trusted platform for your data
- Optimized and predictable system performance
- Productive policy-based management of your infrastructure



Dynamic Development

- Accelerate your development with entities
- Synchronize your data from anywhere



Beyond Relational

- Store and consume any type of data
- Deliver Location Intelligence within your applications



Pervasive Insight

- Integrate all your data in the Enterprise Data Warehouse
- Reach all your users with scalable BI platform
- Empower every user with actionable insights



Accelerate Your Development



Entity Data Model

LINQ

Visual Entity Designer

Entity Aware Adapters

Table Valued Parameters

Merge

Broker Enhancements

- **Develop with business entities**
 - Define your business not tables
 - Model your complex business relationships
 - Retrieve entities instead of rows and columns
 - Connect to SQL Server with entity aware adapters
- **Reliable application performance**
 - Store entities locally on your clients

T-SQL “Delighters”

- Can you spot the features?

```
DECLARE @t int = 5;
```

```
INSERT dbo.myT  
VALUES ('WA', @t), ('FL', @t+1);
```

```
UPDATE dbo.myT  
SET instances+=1;
```

- Continued investment and innovation in T-SQL

Service Broker

- Diagnostics Tool
 - Check new configuration before deploying it
 - Report errors found and recommendations on how to fix them
- SQL Server Management Studio
 - Adding more integrated Broker capabilities e.g. “add” / “edit” actions (queue/service)
- Conversation Priority

Service Broker

- Conversation priority
 - Send and process messages ahead of others

```
CREATE BROKER PRIORITY customer_waiting  
FOR CONVERSATION  
SET ( CONTRACT_NAME = 'http://POS',  
      PRIORITY_LEVEL = <1-10> )
```

- Affects the order messages from different conversations are put on the wire
- Affects the order in which a conversation group is RECEIVE'd off the queue.

Table-valued Parameters (TVP)

- Parameters of type “Table Type”
- Input parameters on SPs/Functions
- New “*ReadOnly*” keyword is needed.
- TVPs are scoped within the SP/function body
- Optimized to scale and perform better for large data
- Behaves like BCP inside server

demo

Table Valued Parameters

```
use tempdb
```

```
GO
```

```
-- Set up table type to be used as parameter
```

```
-- and table to insert data into
```

```
create type t_tvptable as table (i int)
```

```
create table t (a int)
```

```
GO
```

```
-- Set up proc to accept new type
```

```
create proc usp_tvptinsert (@p t_tvptable readonly) as
```

```
insert t select * from @p
```

```
GO
```

```
-- Insert some values
```

Messages

Command(s) completed successfully.

MERGE

- New DML statement that combines multiple DML operations
 - Building block for more efficient ETL
 - SQL-2006 compliant implementation
 - Performance Improvements
 - Optimized loading from flat files -
OPENQUERY(BULK...)
 - Includes both SCD-1 and SCD-2 semantics
 - Includes DELETE semantics

MERGE

- New DML statement that combines multiple DML operations
 - Single statement, cleaner solution

Source

XXXXXX XXX XXX
XXXX XXX

XXXXXXXXXXXX
X XXX XXXX XX
XX XXXX
XXXXXX XXX XX

Source can be
any table or
query

MERGE

- New DML statement that combines multiple DML operations
 - Single statement, cleaner solution

Source

XXXXXX XXX XXX
XXXX XXX

XXXXXXXXXXXX
X XXX XXXX XX
XX XXXX
XXXXXX XXX XX

Target

XXXXXX X
XXXX

XXX XXX

Target can be
any table or
updateable
view

MERGE

- New DML statement that combines multiple DML operations
 - Single statement, cleaner solution



MERGE

- New DML statement that combines multiple DML operations
 - Single statement, cleaner solution

Source

XXXXXX XXX XXX
XXXX XXX

XXXXXXXXXXXX
X XXX XXXX XX
XX XXXX
XXXXXX XXX XX

If no match,
INSERT

Target

XXXXXX XXX XXX
XXXX XXX

XXXXXXXXXXXX
X XXX XXXX XX
XX XXXX
XXXXXX XXX XX

XXX XXX

MERGE

- New DML statement that combines multiple DML operations
 - Single statement, cleaner solution

Source

XXXXXX XXX XXX
XXXX XXX

XXXXXXXXXXXX
X XXX XXXX XX
XX XXXX
XXXXXX XXX XX

If source not
matched,
DELETE

Target

XXXXXX XXX XXX
XXXX XXX

XXXXXXXXXXXX
X XXX XXXX XX
XX XXXX
XXXXXX XXX XX

XXX XXX

MERGE

```
MERGE Stock S
  USING Trades T
  ON S.Stock = T.Stock
  WHEN MATCHED AND (Qty + Delta = 0) THEN
    DELETE -- delete stock if Qty reaches 0
  WHEN MATCHED THEN
    -- delete takes precedence on update
    UPDATE SET Qty += Delta
  WHEN NOT MATCHED THEN
    INSERT VALUES (Stock, Delta)
```


MERGE

```
MERGE Stock S
  USING Trades T
  ON S.Stock = T.Stock
  WHEN MATCHED AND (Qty + Delta = 0) THEN
    DELETE -- delete stock if Qty reaches 0
  WHEN MATCHED THEN
    -- delete takes precedence on update
    UPDATE SET Qty += Delta
  WHEN NOT MATCHED THEN
    INSERT VALUES (Stock, Delta)
  OUTPUT $action, T.Stock, inserted.Delta;
```

demo

MERGE

- 1) Simple usage
- 2) More complex usage

File Edit View Query Project Tools Window Community Help

New Query

master

Execute

Object Explorer

(local).master - _ng Merge Output.sql

Object Explorer Details

Connect

Database!

Adventure

Database

Tables

Sys

cd

cd

cd

cd

cd

db

db

db

Hu

Hu

Hu

Hu

Hu

Hu

Hu

Hu

Hu

Hu

Hu

Hu

Hu

Hu

Hu

Hu

Hu

Hu

Hu

Hu

```
--Set up the test tables and data
```

```
USE [AdventureWorks]
```

```
GO
```

```
CREATE TABLE [dbo].[SCD2_Source] (
```

```
    [BookGroupCode] [char](1) NOT NULL,
```

```
    [BookCode] [char](2) NOT NULL,
```

```
    [BookDescAbbr] [varchar](6) NULL,
```

```
    [BookDesc] [varchar](15) NULL,
```

```
    CONSTRAINT [PK_SCD2_Source] PRIMARY KEY CLUSTERED ([BookGroupCode], [BookCode])
```

Messages

Nonqualified transactions are being rolled back. Estimated rollback completion: 100%.

Query executed successfully.

(local) (10.0 CTP) REDMOND\granti (53) master 00:00:04 0 rows

Ready

Ln 1

Col 1

Ch 1

INS

Microsoft® Introducing: SQL Server™ 2008

Enterprise Data Platform

- Secure, trusted platform for your data
- Optimized and predictable system performance
- Productive policy-based management of your infrastructure



Dynamic Development

- Accelerate your development with entities
- Synchronize your data from anywhere



Beyond Relational

- Store and consume any type of data
- Deliver Location Intelligence within your applications



Pervasive Insight

- Integrate all your data in the Enterprise Data Warehouse
- Reach all your users with scalable BI platform
- Empower every user with actionable insights



Store and Consume Any Type of Data



FILESTREAM data type

Integrated Full Text Search

Sparse Columns

Large User Defined Types
(including CLR > 8k)

Date/Time Data Type

XML Enhancements

- **Seamless transition between relational and non-relational**
 - Store your documents in cost effective storage platform
 - Access documents as data
 - Encode your complex hierarchies within XML
 - Extend the database to meet your unique needs
- **Find your data**
 - Query across relational and text data
- **Build global applications**
 - Capture precise time from clients across the globe

Date & Time Enhancements

- 4 new data types
- SQL Standard compatible, align with .NET

Date

- Date only
- From 0001-01-01 to 9999-01-01 in Gregorian calendar

Time (n)

- Time only
- Optional user specifiable fractional precisions up to 100 nanoseconds

DateTimeOffset(n)

- Time-zone aware/preserved UTC datetime
- Optional user specifiable fractional precisions up to 100 nanoseconds

DateTime2 (n)

- Large date range
- Optional user specifiable fractional precisions up to 100 nanoseconds (Default)
- Time-zone NOT aware

Date and Time Data Types

```
CREATE TABLE t1 (c1 DATE, c2 TIME(3),  
    c3 DATETIME2(7) NOT NULL DEFAULT GETDATE(),  
    c4 DATETIMEOFFSET CHECK  
        (c4<CAST(GETDATE() AS DATETIMEOFFSET(0)))  
);
```

```
INSERT INTO t1 VALUES ('0001-01-01', '23:59:59',  
    '0001-12-21 23:59:59.1234567',  
    '0001-10-21 23:59:59.1234567 -07:00');
```

```
INSERT INTO t1 VALUES ('9999-12-31', '23:59:59',  
    '9999-12-31 23:59:59.1234567',  
    '1111-10-21 23:59:59.1234567 -07:00');
```

```
SELECT c4,  
    DATEPART(TZOFFSET, c4),  
    DATEPART(ISO_WEEK, c4),  
    DATEPART(MICROSECOND, c4) FROM t1;
```

demo

New Date/Time Types

CLR Integration

- User-defined types
 - Support for large instances (>8000)
- User-define aggregates
 - Support for large serialization (>8000)
 - Multi-column input
- Table-valued functions
 - Support for Order declaration
 - Enforced at runtime

```
CREATE FUNCTION myTVF()  
RETURNS TABLE (@i int, @j int, @k int)  
WITH ORDER i ASC, j DESC  
AS EXTERNAL NAME ...
```

XML Upgrades

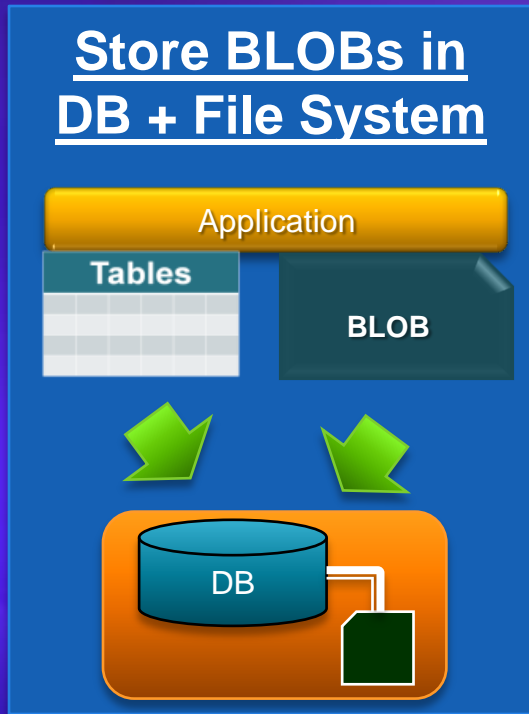
- Improved XML Schema Validation
 - Support for lax validation
 - Full support for storing & validating Office 12 Document formats
 - Full xs:dateTime support
 - Support for no timezone values
 - timezone preservation
 - Support for lists and union types
- Added support for let-clause in XQuery
- Added support insert sql:variable("@xml") into /a/b

Integrated Full Text (completely)

- Current Issues:
 - Indexes stored outside SQL Server lead to manageability challenges
 - Mixed query performance suffers from having to pull complete full-text result set
- Integrated FTS:
 - Full-Text Indexes fully integrated into SQL Server
 - Mixed queries perform and scale

```
SELECT * FROM candidates
WHERE CONTAINS(resume, '"SQL Server"')
      AND ZipCode = '98052'
```

Filestream Storage



- Storage Attribute on VARBINARY(MAX)
- Unstructured data stored directly in the file system (requires NTFS)
- Dual Programming Model
 - TSQL (Same as SQL BLOB)
 - Win32 Streaming APIs with T-SQL transactional semantics
- Data Consistency
- Integrated Manageability
 - Back Up / Restore
 - Administration
- Size limit is the file system volume size
- SQL Server Security Stack

Deliver Location Intelligence Within your applications



GEOMETRY data type

GEOGRAPHY data type

Virtual Earth Integration

HierarchyID data type
(ORDPath)

New Index Types

- **Build location-aware applications**
 - Capture location data from across your organization
 - Integrate location intelligence to your existing applications
 - Visualize your location information
- **Standard Spatial applications**
 - Integrate with industry standard spatial applications
 - Supported by 3rd parties

Spatial in SQL Server 2008

- Provide vector support
- Target geospatial
 - Spatial data which is referenced to a location on the Earth
 - Typically uses spherical coordinates or projected planar coordinates
- No restriction on data actually being geospatial
- Can be a big, complex, important problem


Location aware applications become mainstream!

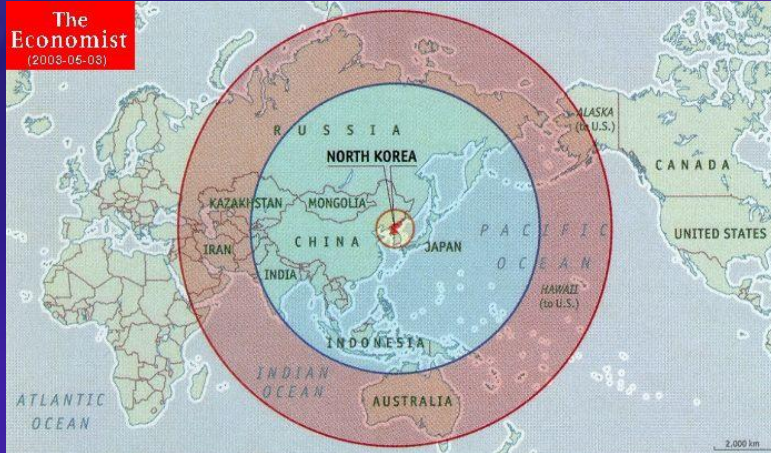
Round Earth / Flat Earth Models

Geodetic



Planar





North Korea's missile threat

Type	Maximum range	Payload	Status
Nodong	1,300 km (810 miles)	700 kg (1,550 pounds)	Currently deployed
Taepodong-1	Up to 10,000 km	Several hundred kg	Test failed 1998, not yet operational
Taepodong-2	10,000-15,000 km	Several hundred kg	Not yet tested




North Korea's missile threat

Type	Maximum range	Payload	Status
Nodong	1,300 km (810 miles)	700 kg (1,550 pounds)	Currently deployed
Taepodong-1	Up to 10,000 km	Several hundred kg	Test failed 1998, not yet operational
Taepodong-2	10,000-15,000 km	Several hundred kg	Not yet tested



Sample Query



SELECT *
FROM roads
WHERE roads.geom.Intersects(@ms)=1
Which roads intersect
Microsoft's main campus?

Two New Spatial Data Types

Geometry

- Geometry can store instances of various types
 - Points, Line strings, Polygons
 - Collections of the above
- Methods for computing
 - Spatial relationships: intersects, disjoint, etc.
 - Spatial constructions: intersection, union, etc.
 - Metric functions: distance, area
 - Planar is conceptually simpler, but more specialized



Geography

- Very similar interface to geometry
 - Some methods have different semantics
- Most commonly available user data is geodetic

Spatial Indexing supported for both data types



Hierarchy Management

- New data type: HierarchyId
- Rich built-in methods for manipulating hierarchies
 - Simplifies storage/querying of hierarchical data
- Many usages – (forum/mailing threads, content management, files/folders, etc.)
- Depth-first and Breadth-first indexes



demo

HierarchyID



Microsoft® Introducing: SQL Server™ 2008

Enterprise Data Platform

- Secure, trusted platform for your data
- Optimized and predictable system performance
- Productive policy-based management of your infrastructure



Dynamic Development

- Accelerate your development with entities
- Synchronize your data from anywhere



Beyond Relational

- Store and consume any type of data
- Deliver Location Intelligence within your applications



Pervasive Insight

- Integrate all your data in the Enterprise Data Warehouse
- Reach all your users with scalable BI platform
- Empower every user with actionable insights



Enterprise Data Warehouse



Partitioned Table Parallelism

Star Join

Data Compression

Resource Governor

Persistent Lookups

Change Data Capture


MERGE

Data Profiling

- Scale and Manage large number of users and data
 - Improve Query performance on large tables
 - Optimize Queries for data warehousing scenarios
 - Increase I/O performance with efficient and cost effective data storage
 - Manage concurrent workloads of ad-hoc queries, reporting and analysis
- Integrate growing volumes of data
 - Optimize ETL performance by identifying data in your largest tables
 - Reduce the data load volumes by capturing operational changes in data
 - Simplify the insert and update data processing
 - Profile your information to identify dirty data

Change Data Capture

- Mechanism to easily track changes on a table
 - Changes captured from the log asynchronously
 - Information on what changed at the source
 - Enabled per table, retention-based cleanup
 - Information about what changed, operation that occurred
- Table-Valued Functions (TVF) to query change data



Change Data Capture

- Mechanism to easily track changes on a table
 - Changes captured from the log asynchronously
 - Information on what changed at the source
 - Enabled per table, retention-based cleanup
 - Information about what changed, operation that occurred
- Table-Valued Functions (TVF) to query change data

```
sys.sp_cdc_enable_db_change_data_capture  
sys.sp_cdc_enable_table_change_data_capture
```

```
cdc.fn_cdc_get_all_changes_<instance>  
cdc.fn_cdc_get_net_changes_<instance>
```

demo

Change
Data
Capture



- 1) Basic Usage
- 2) In Practice (SSIS)

File Edit View Query Project Tools Window Community Help

New Query

master

Execute

Object Explorer

Connect

. (SQL Server 10.0.1056 - REDMOND\granti)

Databases

System Databases

Database Snapshots

AdventureWorks

AdventureWorksDW

EDM_Demo_402

LookupRefDB

MDW

REAL_Source_Sample_V6

REAL_Warehouse_Sample_V6

TDWI2007

Security

Server Objects

Replication

Management

SQL Server Agent

(local).master - ...ects\Demos\CDC.sql Object Explorer Details

```
-----DEMO PART 1-----  
--Enable CDC on DB  
use AdventureWorks  
go  
  
exec sys.sp_cdc_enable_db_change_data_capture  
go  
  
--Check DB is enabled for CDC  
select name, is_cdc_enabled from sys.databases  
where name = 'AdventureWorks'  
go  
  
--Enable a table for CDC  
sys.sp_cdc_enable_table_change_data_capture  
    @source_schema = 'Person'  
    , @source_name = 'Contact'  
    --, @capture_instance = 'Person_Contact'  
    , @supports_net_changes = 1  
    , @role_name = null
```

Messages

Command(s) completed successfully.

Query executed successfully.

(local) (10.0 CTP)

REDMOND\granti (53)

master

00:00:02

0 rows

Ready

Ln 2

Col 1

Ch 1

INS

Demos

Microsoft SQL Serv...

Recurring ETL - Mi...

97%

11:00 AM

Recurring ETL - Microsoft Visual Studio

File Edit View Project Build Team Debug Data Format SSIS Test Tools Window Community Help

Development

Dim_Store_CDC_Before.dtsx [Design] Dim_Store_CDC_After.dtsx [Design]

Control Flow Data Flow Event Handlers Package Explorer

Data Flow Task: DFT Process Dim Store

SRC DimStore

STAT Source

DER Coalesce

SCD

Inferred Output

New Output

SCD-2 Output

SCD-1 Output

SRC DimStore: Filtered list of BI

DER Coalesce: Replaces any NI values in the incoming source d defined unknown strings and non-nullable target fields.

Connection Managers

SQL_REAL_Configuration SQL_REAL_Warehouse SQL_REAL_Source

Solution Explorer

SSIS Packages

AS_Dim_Process.dtsx AS_Dim_Process_MT.dtsx AS_Dim_Process_PT.dtsx AS_Fact_Process.dtsx AS_Fact_Process_MT.dtsx AS_Fact_Process_PT.dtsx Dim_Buyer.dtsx Dim_Buyer_CDC.dtsx Dim_DCVendor.dtsx Dim_DistributionCenter.dtsx Dim_Item.dtsx Dim_Item_Lookup.dtsx

Properties

DFT Process Dim Store Task

Disable False

DisableEventHandlers False

FailPackageOnFailure False

FailParentOnFailure False

MaximumErrorCount 1

Forced Execution Value

Edit Breakpoints

Name

Specifies the name of the object.

Item(s) Saved

Windows Taskbar

Demos Microsoft SQL Serv... Recurring ETL - Mi... 97% 11:13 AM

Grouping Sets

- Extension to the GROUP BY clause
- Lets you define multiple groupings in the same query
- Produces a single result set that is equivalent to a UNION ALL of differently grouped rows.

Makes aggregation querying and reporting
easier and faster

demo

Grouping Sets

Reach All your Users with Scalable BI platform



Pervasive
Insight

Scalable Report Engine

Scale out Analysis

Subspace Computations

New Cube Design Tools

Best Practice Design Alerts

Scalable Backup Tools


IIS Agnostic Report Deployment

- Deliver insights throughout your organization
 - Deliver reports of any size at enterprise scale
 - Scale out through read-only Analysis Services storage
 - Enhance analytical capabilities with more complex computations and aggregations
- Deploy and manage your BI infrastructure
 - Streamline development of the analysis infrastructure with new cube design tools
 - Optimize cube design with real time best practice alerts
 - Backup cubes with enhanced scalability
 - Deploy Reporting Services without IIS dependency

Subspace Computation

Goal: MDX Query Performance

Minimize the space that calculations are performed on




- Cube space is generally “sparse”
 - Values only exist for small number of combinations of dimension keys
- SQL Server 2005 evaluates expressions on complete space

Subspace Computation

Goal: MDX Query Performance

Minimize the space that calculations are performed on



- SQL Server 2008 divides the space to separate calculated members, regular members and empty space
 - SP2 includes some subspace computation
- Analysis and evaluation of calculations for each cell
- Null / default values substituted on subspace basis

Empower Every User With Actionable Insight



New Word Rendering

Improved Excel Rendering

Data Mining Add-Ins for Excel

Report Builder Enhancements

More Flexible Report Layout

Rich-Text Support

Enhanced Data Visualization

Data Mining Engine Improvements

- Deliver information via Microsoft Office
 - Render reports to Microsoft Word
 - Enjoy improved rendering to Microsoft Excel
 - Bring data mining to new, much broader audience
- Enable users to create powerful reports
 - Build powerful ad-hoc reports
 - Create reports with any structure using Tablix
 - Add rich text regions your reports
 - Embed powerful graphical data visualizations into reports (Dundas)
- Empower users with enhanced analysis
 - Empower users with enhanced write back scenarios
 - Accelerate end user prediction capabilities through enhanced Data Mining structures

Questions?

Chad Boyd

chad.boyd@gmail.com

<http://blogs.msdn.com/chadboyd>

<http://www.mssqltips.com>

Data Compression Blog:

<http://blogs.msdn.com/chadboyd/archive/2007/07/28/katmai-sql-2008-data-compression-including-backup-compression.aspx>

Feature List Blog:

<http://blogs.msdn.com/chadboyd/archive/2007/07/26/katmai-sql-2008-the-list-of-new-features.aspx>

Upcoming Katmai Presentations in the DC area:

<http://blogs.msdn.com/chadboyd/archive/2007/07/18/katmai-sql-server-2008-presentations-in-the-dc-area-novasql-caparea-net-cmap.aspx>

Future Webcasts and Tips:

<http://www.mssqltips.com>

Sql Server 2008 CTP's:

<https://connect.microsoft.com/SQLServer/content/content.aspx?ContentID=5395>