



Esri International User Conference | San Diego, CA  
Technical Workshops | July 12 & 14, 2011

# Administration for SQL Server - Advanced

Thomas Dunn

Tony Wakim

# Scope of this Technical Workshop

- **Advanced topics related to managing Enterprise SQL Server geodatabases**
- **Prerequisites**
  - **SQL Server Administration – An Introduction, or**
  - **Basic knowledge of administering an ArcSDE for SQL Server geodatabase**
  - **Working knowledge of ArcSDE Technology basics**
  - **Working knowledge of SQL Server basics**

Fulton County Dept. of Health and Wellness/District 3, Unit 2, Office



DeKalb County Board

# Agenda

Administration

Performance



SQL Server Spatial  
Types



Fulton County Dept. of Health and Wellness/District 3, Unit 3

DeKalb County Board

# Administration



# Enhanced Data Type Support in ArcGIS 10

- **Native SQL Server XML**
  - Alternative to SDE\_XML
  - Choose with DB\_XML or SDE\_XML configuration keyword
- **varbinary(max)**
  - Replaces deprecated type *image* for large binary objects
- **varchar(max) & nvarchar(max)**
  - replaces deprecated types *text* and *ntext* for large text objects
- **Datetime2**
  - replaces datetime
  - Larger date range and larger default fractional precision


# SQL Server Support

- **SQL Server 2008 R2 is supported in 9.3.1 SP2 and 10**
  - Same level of support and functionality as with earlier versions of SQL Server 2008
- **SQL Server 2008 Express delivered with ArcGIS**
  - 32-bit and 64-bit
- **SQL Server 2005 not supported in 10.1**
- **SQL Server 2000 not supported in 10.0**

# Upgrading Geodatabases to 10

- **GDB upgrades to 10 use a geoprocessing tool**
  - **Named “Upgrade Geodatabase”**
  - **Callable from ArcGIS Desktop or a Python script**
- **Replaces upgrades previously done (pre-10) from ArcSDE post-installer or from *sdesetup -o upgrade***
- **See <http://help.arcgis.com/> topic “A quick tour of geodatabase upgrades”**

# Upgrading a Geodatabase






# Storage Type Migration GP Tool

- In-line migration of vector (feature) storage from SDEBINARY to Geometry or Geography
- No need to reload data
- Use Migrate Storage geoprocessing tool
  - Also available in *sdelayer -o migrate*
- See <http://help.arcgis.com> topic “Data migration from one storage type to another”

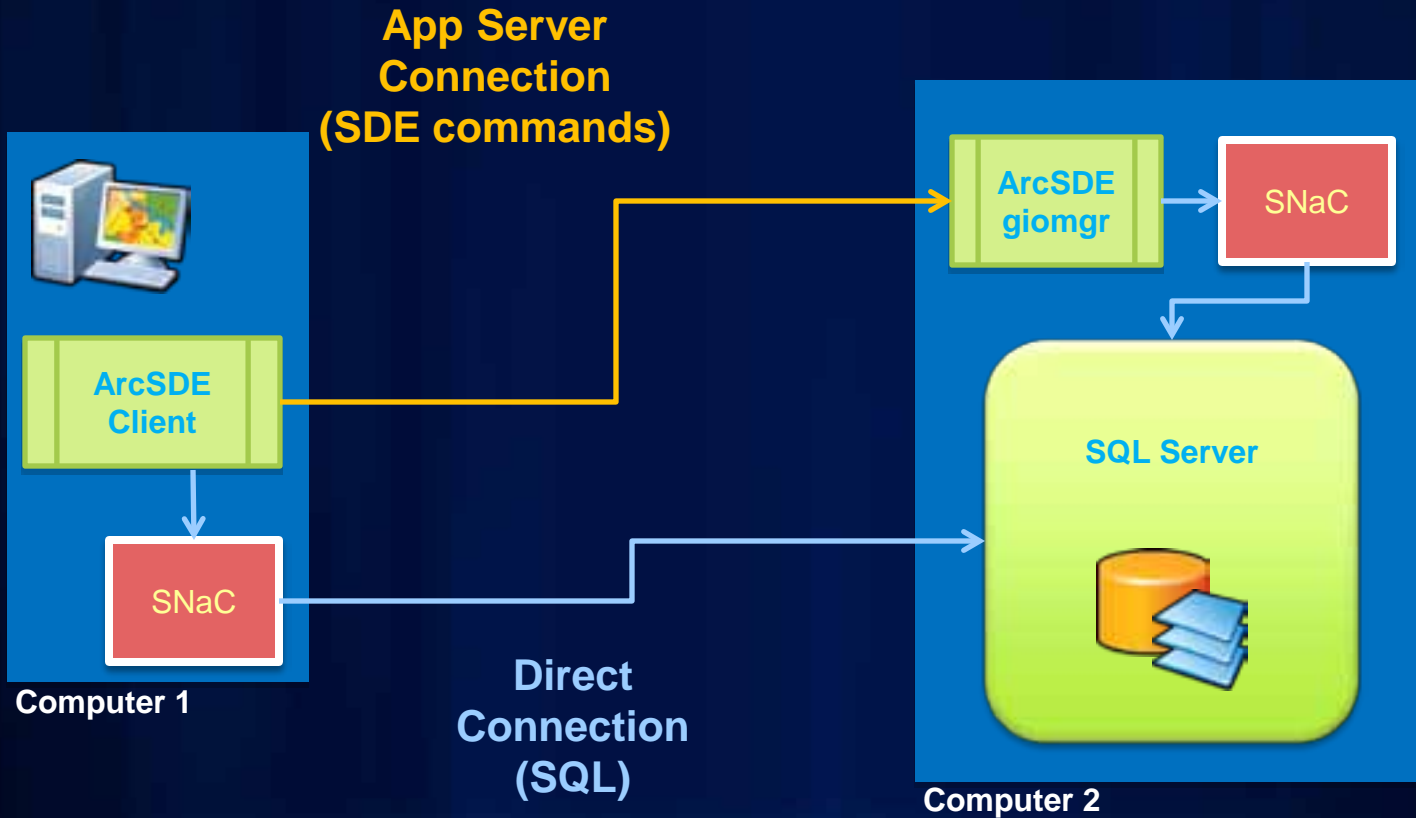
# Storage Type Migration



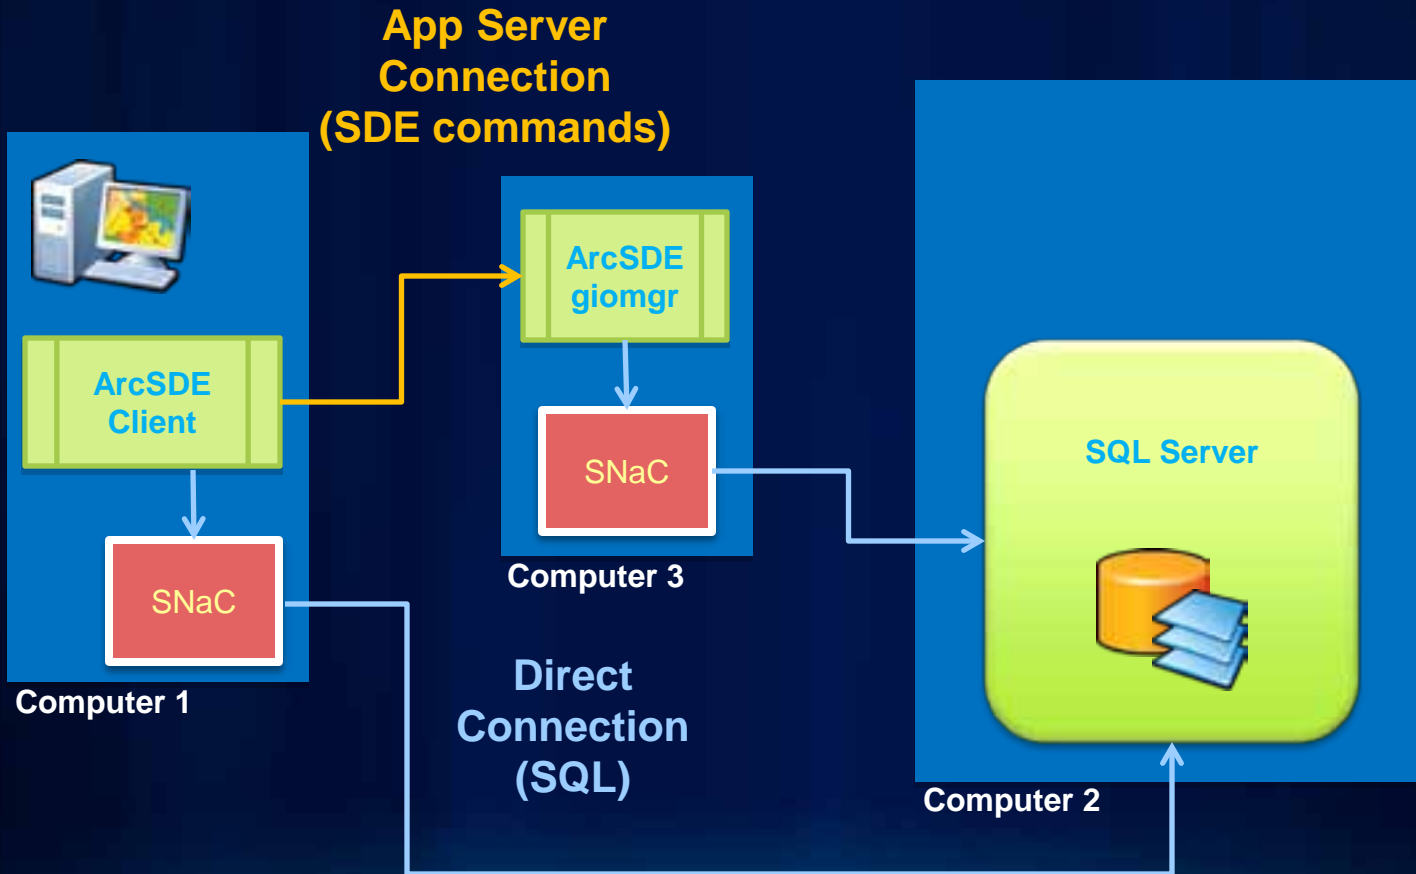
## Connections: SQL Native Client

- **Required for connections to SQL Server**
  - Installed with ArcGIS Desktop, Server and Reader
  - Also available on Engine and ArcSDE SDK media
  - Download sqlncli.msi from [SS2008 SP2 Feature Pack](#)
- **Use SQL Native Client software at least as new as the SQL Server database engine**
  - Provides support for all data types and commands
  - MDAC not supported: use SQL Native Client instead
- **See Microsoft's Data Access Technologies Roadmap**
  - History and current state of Microsoft data access technology
  - <http://msdn.microsoft.com/en-us/library/ms810810.aspx>

# SQL Native Client Location



# SQL Native Client Location



## **Native 64-bit Builds of ArcSDE**

- **64-bit ArcSDE builds for most operating systems available for ArcSDE 9.3 SP1 or later**
- **Typically a performance improvement over 32-bit ArcSDE running on 64-bit platforms**
- **Much larger memory address range with 64-bit**
  - **Limited by O/S and database, potentially up to 16.8 million TB**
  - **Windows Server 2008 Datacenter & Enterprise supports 2 TB RAM; Standard & Web editions support 32 GB**

Fulton County Dept. of Health and Wellness/District 3, Unit 3, 01/10/10



DeKalb County Board

# Performance

Keeping an eye on  
performance

Things that seem like a  
good idea, but aren't

Performance  
diagnostics



# Keeping an Eye on Performance

- **Compress your geodatabase periodically**
  - Particularly when many edits are performed
  - Removes unneeded rows from tables queried often
    - See “Enterprise Geodatabase Tips & Tricks” workshop
- **Pre-extend MDF and LDF files to avoid autogrow**
  - Autogrow is a safety net, not an automated file management tool
  - With default settings, there is potential to fragment files
    - <http://support.microsoft.com/kb/315512>
- **Back up your data**
  - Minimizing downtime is part of performance planning
  - Create a restore plan first, then a backup plan
    - <http://msdn.microsoft.com/en-us/library/ms191239.aspx>



# Keeping an Eye on Performance

- **Beware of poorly performing queries and views**
  - **Prototype, trace, examine execution plan during view development, not just afterward**
  - **Use the Activity Monitor to look for recent expensive queries in production system**
- **Statistics**
  - **Don't forget statistics are needed on the GDB system table statistics, especially SDE\_states, SDE\_state\_lineages, SDE\_mvtables\_modified, GDB\_Items**
  - **[http://msdn.microsoft.com/en-us/library/dd535534\(v=sql.100\).aspx](http://msdn.microsoft.com/en-us/library/dd535534(v=sql.100).aspx)**
  - **ArcGIS Resource Center article “Statistics in SQL Server databases”**
  - **[Technical Article 24518](#) “How can ArcSDE performance be improved?” (SQL to update statistics and rebuild indexes)**

# Statistics

Demo

**Things that Look Like  
Good Ideas,  
but Aren't**

## **Auto close (database property)**

- **Indicates that database will shut down and free resources after last user disconnects**
- **Reallocating resources takes time, delaying the next connection**
- **Recommendation: Leave this turned OFF**
- **This setting is deprecated by Microsoft**

## **XACT\_ABORT (server instance connection property)**

- **Causes automatic rollback on some errors**
- **Usually “off”, sometimes this property is set “on” by DBA for all new connections to the server**
- **Some “errors” are expected (e.g. table does not exist)**
- **Leave this turned OFF or ArcSDE editing will fail**

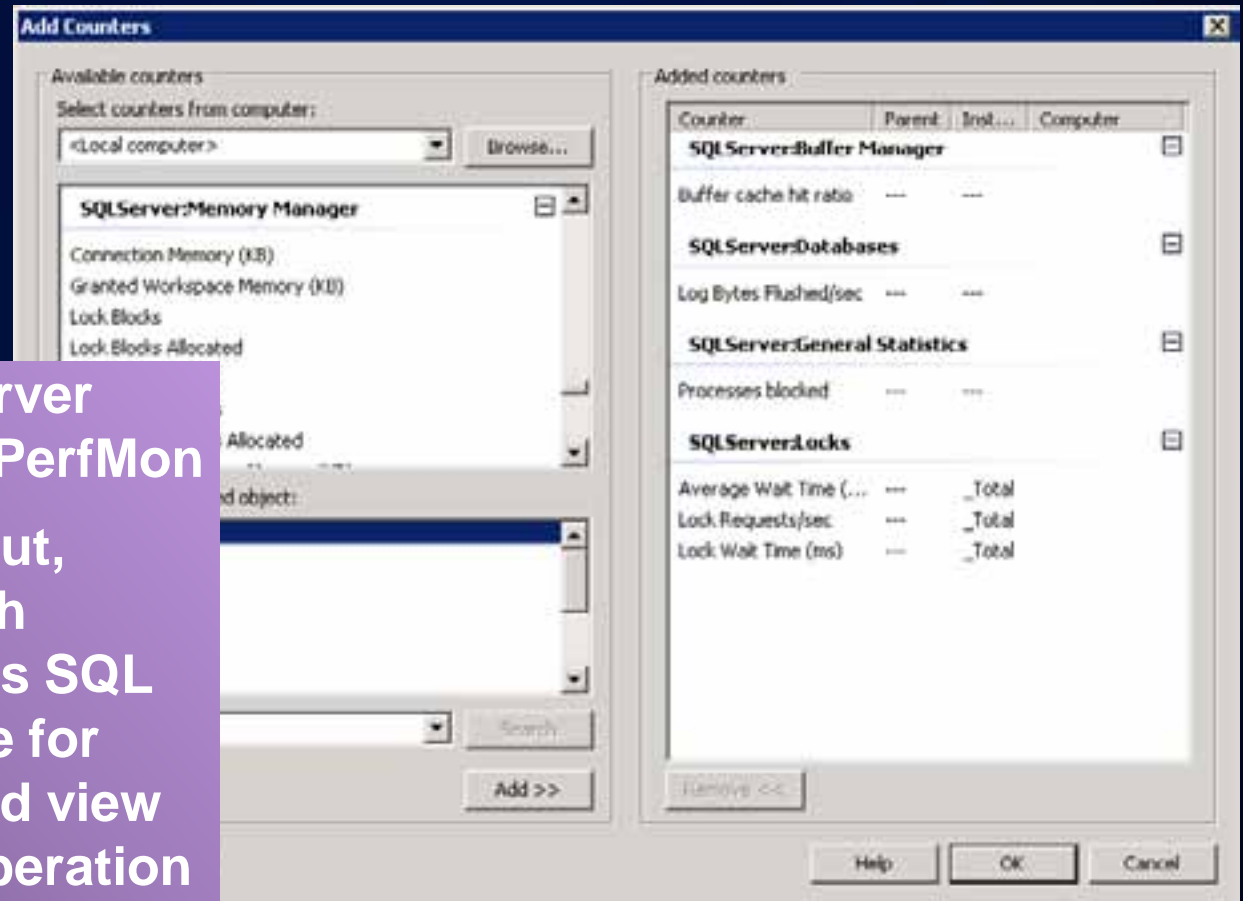
# Performance Diagnostics

# SQL Server Profiler

- **SQL Server Profiler** traces what is happening as it occurs
- Wide array of events and event properties to trace
- Build custom tracing templates containing selected event and parameter filters needed for various scenarios
- Traces can be saved for later analysis, and overlaid with Performance Monitor Counter Logs
- Able to trace remote server instances, including Express

# SQL Server Counters in Windows Performance Monitor

- Add SQL Server counters to PerfMon
- Record output, combine with simultaneous SQL Profiler trace for more detailed view of system operation







# Spatial Type Topics

Geometry and  
Geography

Creating feature  
classes

Spatial indexing

Spatial references



# Geometry and Geography Spatial Data Types

## Built-in CLR types in SQL Server 2008

- **Geometry**
  - Any coordinate system
  - Any spatial extent
  - Cartesian (flat Earth) model, calculations
- **Geography**
  - Fixed set of geographic coordinate systems
  - Global extent (maximum)
  - Ellipsoid (round Earth) model, calculations



# Comparing Geometry and Geography Types

	<b>geometry</b>	<b>geography</b>
Line interpolation	Cartesian	Great Circle
Spatial Domain	Defined by coordinate system	Sphere (Earth)
Units in coordinate system	Defined by coordinate system	Degrees
Units from calculations	Same as coordinate system units	Defined by coordinate system – meters or feet
Supported coordinate systems	Any: SRID value has no inherent meaning, but match is required for comparison	Defined in <code>sys.spatial_reference_systems</code> table

# Creating Spatial Type Feature Classes

- **Spatial type column is in Business & Adds tables**
- **SQL Server maintains the spatial index**
  - **Using parameters from the SDE\_DBTUNE table**
- **Only one value for Shape.STSrid allowed per table**
  - **Enforced with a check constraint**

## **SDE\_GEOMETRY<n> Table**

- **Side table created automatically**
  - **Named SDE\_GEOMETRY<registration\_id>**
- **Contains curves, point IDs, CAD elements: items that cannot be represented natively in Geometry or Geography type**
- **Table contains rows only for each feature that has one of these elements**
- **Joins are made to the side table only after it contains rows**

# Creating Feature Classes

Demo

# Geography Spatial Type Rules

- **Must have SRID defined in SQL Server data dictionary (SQL Server rule)**
- **Cannot violate hemisphere rule (SQL Server rule)**
  - <http://blogs.msdn.com/b/isaac/archive/2009/02/06/the-geography-hemisphere-limitation.aspx>
- **Cannot have Z or M dimensions (ArcSDE rule)**
  - ArcSDE imposes limitation on Z & M coordinates for Geography type
  - Hemisphere rule requires SQL Server to validate shape on insert
  - For performance reasons, ArcSDE uses Well-Known Binary (WKB) interface for reading & writing to shape
  - SQL Server WKB implementation uses 1999 OCG specification for WKB, which does not support Z or M

# Spatial References




- **Defines how a representation of a feature is related to the real world feature it represents**
  - **Conversely, defines how coordinates are assigned to a real world feature**
- **Necessary to understand location of and relationships between features**
- **The meaning of “SRID” meaning varies with context**



# Spatial Reference Identifier (SRID)

- **Geometry data type: Arbitrary, but required**
  - Geometry methods that operate on two objects require that their SRID's match, or will return null
- **Geography data type: fixed set of coordinate systems**
  - `sys.spatial_reference_systems`
- **ArcGIS: Complete spatial reference**
  - `sde_spatial_references`
  - `AUTH_SRID` column refers to SRID used in objects in spatial type column

# SQL Server Spatial Index



- 4-level tessellated, adjustable grid hierarchy
- Spatial Indexing overview whitepaper:
  - <http://msdn.microsoft.com/en-us/library/bb964712.aspx>

# SQL Commands to Create a Spatial Index



```
CREATE SPATIAL INDEX spidx_geometry
ON parcels (geometry_col)
USING GEOMETRY_GRID WITH (
  BOUNDING_BOX = (0, 0, 500, 200),
  GRIDS = (LOW, LOW, MEDIUM, HIGH),
  CELLS_PER_OBJECT = 64);
```

Geometry  
needs a  
bounding box

```
CREATE SPATIAL INDEX spidx_geography
ON world_rivers (geography_col)
USING GEOGRAPHY_GRID WITH (
  GRIDS = (MEDIUM, MEDIUM, MEDIUM, MEDIUM),
  CELLS_PER_OBJECT = 64);
```

# Spatial Index Options

- **GRIDS** – number of cells at each grid level
  - LOW = 16, MEDIUM = 64, HIGH = 256
  - Default = MEDIUM, MEDIUM, MEDIUM, MEDIUM
- **CELLS PER OBJECT** – maximum cells that can be recorded as intersecting each object
  - Indirectly specifies lowest grid level used based on object size
  - Default is 16



# Using SQL Server Spatial Index with ArcGIS

- **Spatial index is created automatically for all feature classes**
  - Uses DBTUNE configuration parameters  
B\_MS\_SPINDEX (business table) and A\_MS\_SPINDEX (Adds table)
- **Spatial indexes can be changed**
- **ArcGIS does not use hints for spatial indexes**
  - See Esri technical article [36617](#)

# Optimizing Spatial Index

- Relationship of extent of spatial index and feature class extent can have large effect on efficiency
- `sp_help_spatial_geometry_index` (et al) can measure efficiency of a spatial index
  - Returns properties of spatial index
    - Creation parameters, total/average row and cell counts, cell sizes
  - Calculates index efficiency based on a sample query
    - $\text{Output rows} / \text{primary filter rows} = \text{efficiency}$

# Evaluating Spatial Index Effectiveness

1. **Trace the query using SQL Server Profiler**
2. **Examine execution plan to verify that spatial index is being used**
3. **Execute traced spatial query directly in SQL Server Management Studio to determine if performance is different**
4. **Evaluate spatial index parameters to see if it can be optimized**
5. **Consider whether an attribute-constrained query may be faster**

# Profiler

Demo



# Related Workshops

- **Using SQL Profiler to Troubleshoot**
  - Tuesday 5:00pm Geodatabase Management Demo Theatre
  - Wednesday 1:00 pm Geodatabase Management Demo Theatre
- **Using SQL with your Geodatabase**
  - Thursday 10:30am Geodatabase Management Demo Theatre
- **Upgrading to ArcGIS 10.0 Geodatabases**
  - Thursday 1:30pm Room 23B
- **Troubleshooting Performance Issues with Enterprise Geodatabases**
  - Thursday 10:40am Room 24A
- **Python – Automating Geodatabase Administration**
  - Thursday 11:05am Room 24A
- **Using Spatial Data in ArcGIS with Query Layers**
  - Thursday 1:55pm Room 23B

# Related Workshops

- **Using Spatial Data in ArcGIS with Query Layers**
  - Thursday 1:55pm Room 23B
- **Implementing Database Roles in the Enterprise Geodatabase**
  - Thursday 3:15pm Room 3
- **Enterprise Geodatabase Administration – Tips and Tricks**
  - Thursday 3:40pm Room 3

**Questions?**

**Session evaluations at**  
**<http://www.esri.com/sessionevals>**



esri