



Esri International User Conference | San Diego, CA
Technical Workshops | 14.07.11

Enterprise Geodatabase Administration – Tips and Tricks

Jim McAbee

Tony Wakim

Enterprise Geodatabase Tips and Tricks

Goals:

- Issues common to all ArcSDE technology consumers.
- Avoid RDBMS specific issues (e.g. Only PostgreSQL)

...our database is too slow...

...I can't add a field to a feature class...


...what is "direct connect" ?...

...how do I edit and publish maps and avoid locks?...

Assumed Knowledge:

- Use of RDBMS tools
- Use of Standard ArcMAP, ArcGIS Server operations

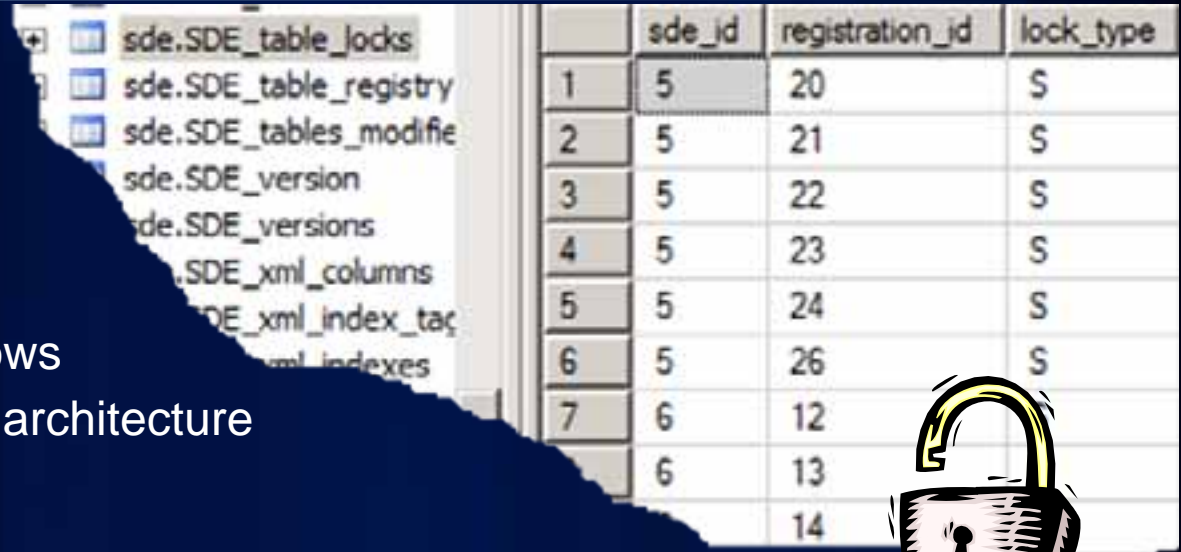
IT Stack



Schema Locks: Tips and Tricks

Schema Locking
depends upon:

- Geodatabase design best practices
- Well defined data management workflows
- Proper Geodatabase architecture



A screenshot of a database table named 'sde.SDE_table_locks'. The table has four columns: an index, 'sde_id', 'registration_id', and 'lock_type'. The data shows seven rows of locks, all with 'lock_type' 'S' (Shared). A stylized padlock icon is overlaid on the bottom right of the table.

	sde_id	registration_id	lock_type
1	5	20	S
2	5	21	S
3	5	22	S
4	5	23	S
5	5	24	S
6	5	26	S
7	6	12	
	6	13	
		14	

Lock types

- Table, Layer, State, Object
- Table – shared most common – will prevent modification of table or feature class (object) structure, permissions, etc...

Some methods of working around schema locks

- Geodatabase Architecture – multiple geodatabases
- For data synchronization – GP – Delete and Append features
- For viewing – SDE and RDBMS views
- ArcGIS Server - SchemaLockingEnabled

Geodatabase Object Management

Tips on Registration and Schema changes

Adding a field at RDBMS level

- column_registry table
- owner view w/ArcCatalog or run sdetable -o describe

database_name	table_name	owner	column_name
GISDATA	BLOCKS	GISDATA	BLOCK
GISDATA	BLOCKS	GISDATA	OBJECTID_1
GISDATA	BLOCKS	GISDATA	Res
GISDATA	BLOCKS	GISDATA	Shape
GISDATA	DISTDIAM	GISDATA	Angle
GISDATA	DISTDIAM	GISDATA	AnnotationClassID
GISDATA	DISTDIAM	GISDATA	Bold
GISDATA	DISTDIAM	GISDATA	CharacterSpacing
GISDATA	DISTDIAM	GISDATA	CharacterWidth

• “Non-Geodatabase” Table Registration

<http://webhelp.esri.com/arcgisserver/9.3.1/dotNet/index.htm#geodatabases/enhanci-205816615.htm>

- 3 data management tiers
 - ArcObjects
 - ArcSDE
 - RDBMS
- Objectid - user/sde managed unique non-null integer field

What ObjectID field provides

- If your table lacks an ObjectID field, you won't be able to
 - Select the features in the layer on the map in any way.
 - Apply [definition query](#) to display subset of records/ features.
 - Create relates.
 - Start an edit session and edit the attributes.

	OBJECTID	PROPERTY_I
938	961	1961.00000000
939	962	1962.00000000
940	963	1963.00000000
941	964	1964.00000000
942	965	1965.00000000

10.0

sde.GDB_ITEMRELATIONSHIPS
 sde.GDB_ITEMRELATIONSHIPTYPES
 sde.GDB_ITEMS
 sde.GDB_ITEMTYPES
 sde.GDB_REPLICATALOG
 sde.GDB_TABLES_LAST_MODIFIED

sde.GDB_ANNOSYMBOLS
 sde.GDB_ATTRRULES
 sde.GDB_CODEDDOMAINS
 sde.GDB_DEFAULTVALUES
 sde.GDB_DOMAINS
 sde.GDB_EDGECONNRULES
 sde.GDB_EXTENSIONDATASETS
 sde.GDB_EXTENSIONS
 sde.GDB_FEATURECLASSES
 sde.GDB_FEATUREDATASET
 sde.GDB_FIELDINFO
 sde.GDB_GEOMNETWORKS
 sde.GDB_HISTORICALMARKERS
 sde.GDB_JNCONNRULES
 sde.GDB_NETCLASSES
 sde.GDB_NETWEIGHTASOCS
 sde.GDB_NETWEIGHTS
 sde.GDB_NETWORKS
 sde.GDB_OBJECTCLASSES
 sde.GDB_RANGEDOMAINS
 sde.GDB_RASTERCATALOGS
 sde.GDB_RELCLASSES
 sde.GDB_RELEASE
 sde.GDB_RELRULES
 sde.GDB_REPLICADATASETS
 sde.GDB_REPLICATALOG
 sde.GDB_REPLICAS
 sde.GDB_REPLICASEX
 sde.GDB_SPATIALRULES
 sde.GDB_SUBTYPES

Design: Log File Recommendations

- **3 Types – Session, Shared and Stand-Alone**
- **Use default architecture for RDBMS**
 - **Shared for Oracle (IDs removed upon unselect versus disconnect)**
 - **Global temporary tables – see ESRI KB article 32161**
 - **Session for SQL Server**
- **Use session if multiple users have same login**
 - **Default for SQL Server (session/tempdb – minimizes logging)**
 - **Avoids table contention (e.g., many users connecting with a single login)**
- **Use pool if users are prohibited from creating objects**
 - **If pool unavailable, ArcSDE will attempt a user-owned table**

Design: Geodatabase Architecture

- **Single vs. Multiple Geodatabases**

- **Uses/Requirements**


- Vector and Raster, Editing and Publishing, Production and Development, Departmental
 - Performance – use specific tuning
 - Management/Administration (e.g. locking issues)
 - High Availability (HA), Disaster Recovery (DR)

- **Implementation**

- Multiple Instances (e.g. Oracle)
 - Multiple Databases or Named Instances (e.g. SQL Server)

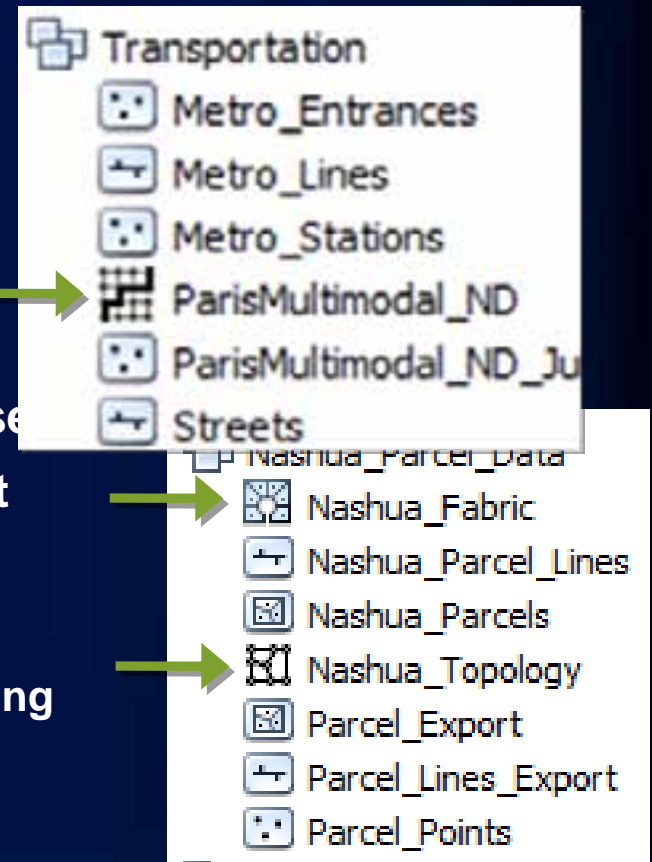
- **Data Distribution and Synchronization**

- Geodatabase Replication
 - Export/Import
 - Database Cloning or Replication



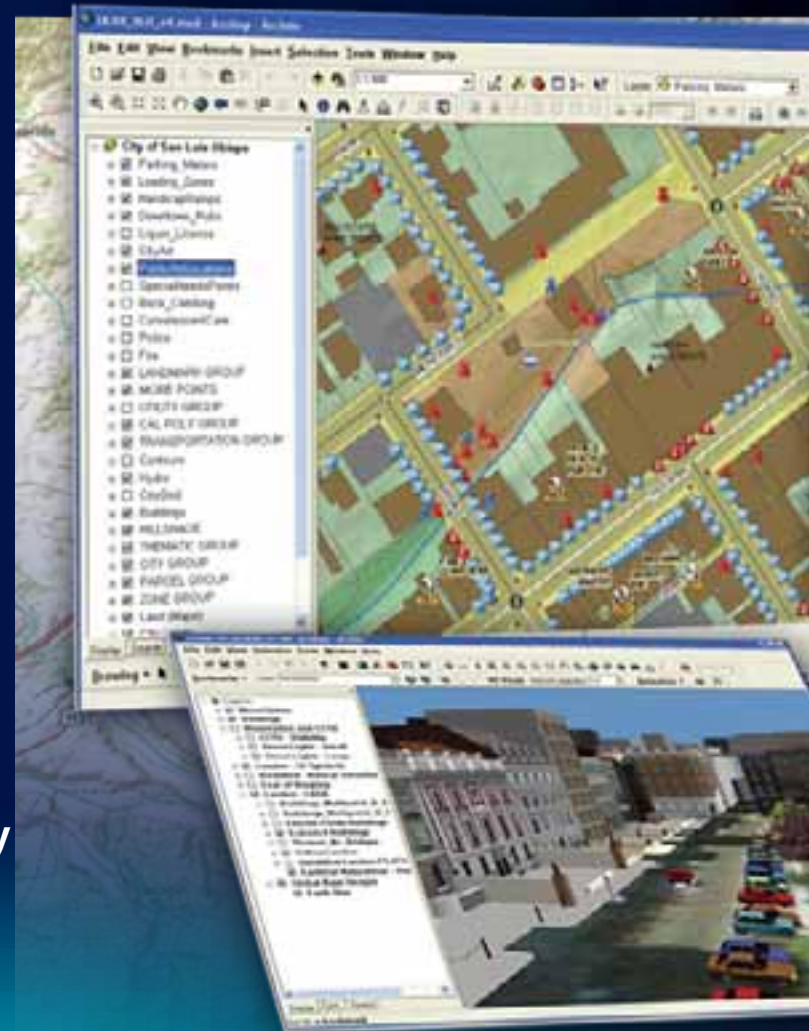
Feature Datasets

- **Designed to ensure spatial coincidence**
 - Required for many types of behavior
 - Geometric networks, topologies, etc...
- **Considerations for use**
 - All feature classes are instantiated
 - Privileges are granted/revoked for all classes
 - Registering as versioned occurs at dataset level
 - Locks can apply to all feature classes
 - Spatial Views, Separate Editing and Publishing geodatabases, etc... can sometimes help
- **Avoid using for organizational purposes**

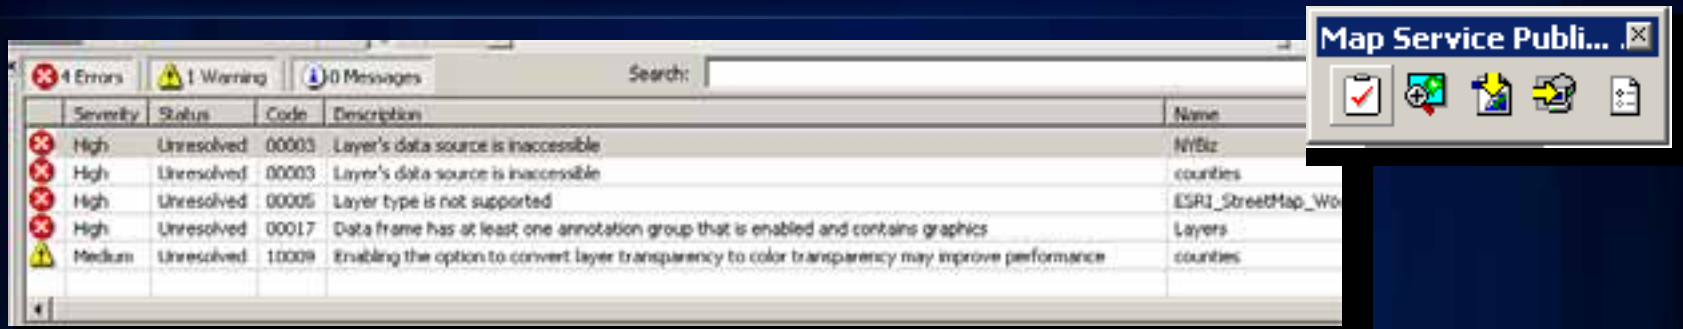


Labeling, Symbols, Sub-Queries and Renderers

- **Impact of labeling and symbols**
 - Causes a 2nd SQL query on every layer (feature, then label attribute)
 - Enable Map Cache or use Annotation
 - Try alternative symbology types
 - E.g. Graduated color vs graduated symbol
- **Impact of Sub-Queries**
 - Can reduce amount of information returned by filtering attributes
 - Can cause extra load on the database if columns not indexed
- **Impact of Renderer**
 - More complexity in rendering and symbology increase CPU load, especially in Citrix/Terminal server architectures.
 - Switch to more appropriate renderers.



Tools and Tips for Map Documents



- map document performance analysis tool at 9.3.1
- New .msd document @ 9.3.1
 - Map Service Definition file for 9.3.1 ArcGIS Server fast drawing engine
- Basemap Layer @ 10.x
 - uses a combination of in-memory and disk caches. Limit your map to less than 5 basemap layers as possible to limit memory consumption
- MxdPerfStat (arcscrips.esri.com)

Item	Layer Name	At Scale	Refresh Time (sec)	Recommendations	Features	Vertices	Labeling	Geography Phase (sec)	Graphics Phase (sec)	Cursor Phase (sec)	DB MS CPU	DB MS LIO	DB MS PIO	DBMS
2	STUDENT.parcel_1	50,000	8.22	set scale dependency; run DBMS trace;	30,628	153,140	True	4.27	3.84	5.74	2.37	69,995		esriDBMS_Oracle

ArcMap best practices for users

- **Avoid full display**
 - ArcMap magnifier and overview windows
 - Scale dependencies
 - Use spatial bookmarks
- **Set selectable layers**
- **Keep table of contents and symbology simple**
- **Use keep only matching records option with joins**
- **Use feature (10.x)/map cache (9.x)**
- **Use basemap layer @ 10 and display cache**



Caching Tips

- ArcGIS Desktop Feature/Map Cache
- Display Cache
- ArcGIS Server Map Cache

About ArcGIS caches

[Resource Center](#) > [Professional Library](#) > [Mapping a](#)

ArcGIS software refers to caches in a number of ways. This can be confusing to understand. Consider using each?

This topic helps to explain three of the common

The ArcMap feature cache

This feature cache is used to improve performance when displaying a local area network (LAN) in the ArcMap. It stores records in memory on your computer to speed up work tasks performed within



ArcGIS Resource Center

Server 10

Map caches (2D)

- What is map cache
- A quick tour of
- Map caching co
- The map cachin
- Exporting and i
- Labels in map c
- Local cache dir

Tips and best practices for map caches

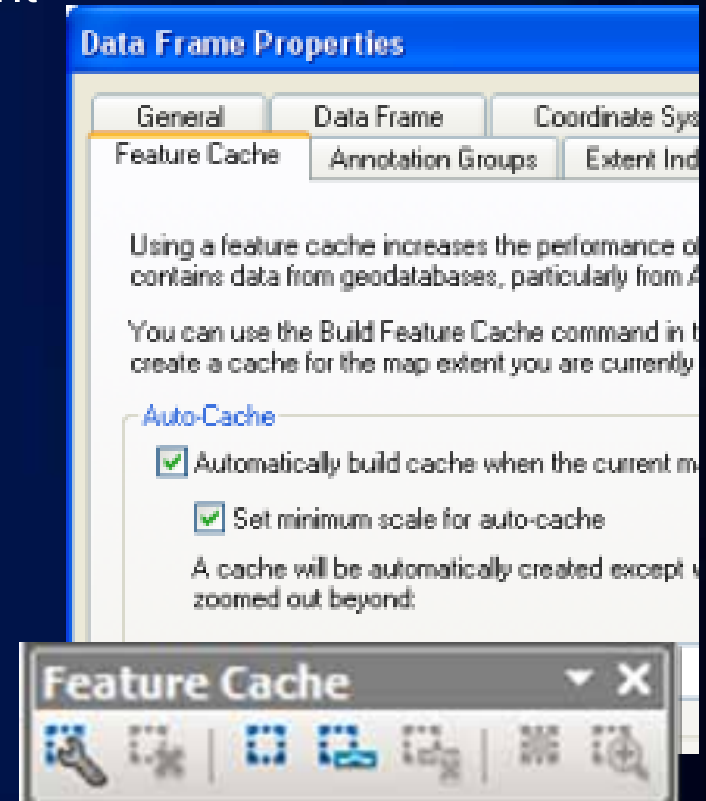
[Resource Center](#) > [Publishing services](#) > [Caching services](#) > [Map](#)

The tips below can help you achieve the best appearance

- [Preparing the map document](#)
- [Creating the cache](#)
- [Maintaining the cache](#)

ArcGIS Desktop Feature/Map Cache

- Client side caching over given spatial extent
- Can speed up queries
 - Reduces roundtrips to the database
- When to use?
 - If making many spatial queries within a common extent
 - If working with several features within a certain geographic area.
 - When editing non-simple Geodatabase features, e.g. Geometric Networks
 - When editing and snapping enabled. Each snap requires a round-trip(s) to the database unless there is a cache
- Do NOT use for non-versioned editing (short-transaction)



ArcGIS Server Map Cache

- Tiles pre-rendered at fixed scales
- Rapid display – fast map display engine (msd)
- Improved performance AND scalability for Geodatabase
 - Helps to reduce SQL activity



Versioned Editing Performance

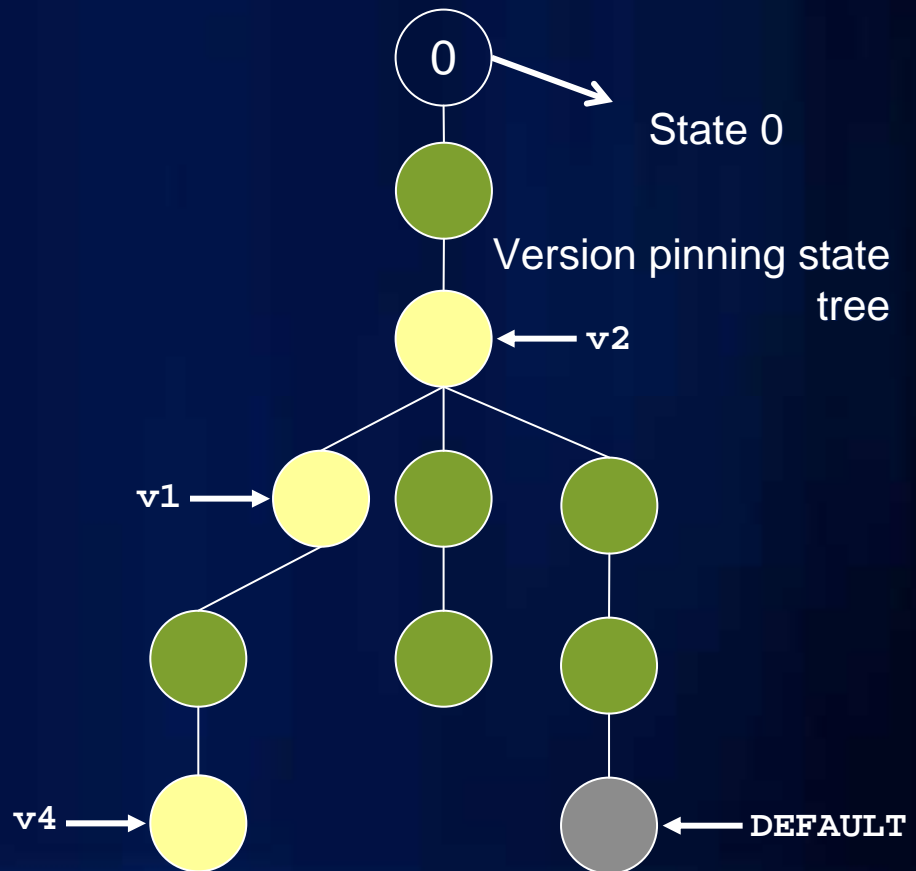
- **Performance affected by:**

- Volume of states
- Stale statistics
- Unnecessary versions*
- Proper workflow

- **Manage your version**

- Reconcile
- Post
- Compress

- Replication also based on versioning



Automation of Reconcile / Post

- **KB36809: Report the recommended reconcile order using SQL in SQL Server**
- **KB35735: Report the recommended reconcile order using SQL in Oracle**
- **Use GP Tools and script to python**

Thank You

<http://www.esri.com/sessionevals>

