

Esri's Spatial Database Engine

IT4GIS
Keith T. Weber, GISP
GIS Director
ISU-GIS Training and Research Center

Pocatello | Idaho Falls | Meridian | Twin Falls



Today's Topics


- Part 1:
 - What is ArcSDE?
 - Why use ArcSDE?
 - ArcSDE Data Structure
 - How is data stored within ArcSDE?
- Part 2:
 - Enterprise workflow: Versioning and Replication

Pocatello | Idaho Falls | Meridian | Twin Falls



What is ArcSDE?

- A spatial database engine that works on an RDBMS.
- Helps to serve geospatial data to clients via a network



Pocatello | Idaho Falls | Meridian | Twin Falls



Why use ArcSDE?

- **Advantages:**


- Data loss/integrity degradation through versioning
- Centralized data management
- Enterprise GIS
- Geo-spatial data is immediately usable



Pocatello | Idaho Falls | Meridian | Twin Falls

Idaho State
UNIVERSITY

Why use ArcSDE? (cont'd)



- **Disadvantages**

- Data management role
- RDBMS administration
- Capital expenditure

Pocatello | Idaho Falls | Meridian | Twin Falls

Idaho State
UNIVERSITY

To Use ArcSDE...or Not ...

- What will help make this decision?

- ROI
- TCO
- Is this the correct technology for the problem?

Pocatello | Idaho Falls | Meridian | Twin Falls


Idaho State
UNIVERSITY


ArcGIS Data Structures

Vector objects
Shapefiles
Coverages

Raster objects
Grids
Images

GDB







Pocatello | Idaho Falls | Meridian | Twin Falls

The GDB

- Can store tables (data values), vector feature classes, and raster layers






Pocatello | Idaho Falls | Meridian | Twin Falls

Layers and Layer Files

- All GIS Datasets are considered LAYERS in ArcMap.
- A LAYER FILE is a file that you save in ArcMap to retain customized settings.
 - This file refers to the LAYER (shape file, coverage, grid, or feature class)
 - It displays the data with your saved visualization settings, textual annotation, etc.



Pocatello | Idaho Falls | Meridian | Twin Falls

Workspaces

- Arc/Info
 - Info folder
 - Geodata sets (coverages, grids, TINs)
- Collection of ArcView shape files
- Geodatabases

GeoDatabases

- Personal (going away with ArcGIS Pro)
- File-based
- ArcSDE Personal
- ArcSDE Professional (or Enterprise)

Personal Geodatabases

- Uses the MS Access Jet Database engine
 - *Note: Do not open/edit these with MS Access*
- Limitations
 - 2GB (Access)
 - Only vector feature classes are actually stored inside the Access database
 - 4 users but only one editor
 - Does not support versioning
 - No longer supported with ArcGIS Pro



File-based Geodatabase

- fgDB
- Stores vector and raster layers in the file/folder structure.
- Limitations
 - Multi-user (max = 10)
 - 1 Editor (no versioning)
 - Max size is 1 TB



Pocatello | Idaho Falls | Meridian | Twin Falls

Idaho State
UNIVERSITY

ArcSDE Personal

- Uses MS SQL Server Express
- Limitations
 - 10 GB
 - Supports versioning/replication but only one editor



Pocatello | Idaho Falls | Meridian | Twin Falls

Idaho State
UNIVERSITY

ArcSDE Professional Geodatabases

- Uses DB2, Oracle, Informix, SQL Server, etc.
- No software size limits and unlimited number of users
- Can accommodate vector and raster data



Pocatello | Idaho Falls | Meridian | Twin Falls

Idaho State
UNIVERSITY

Given all these differences, there are really many similarities

Pocatello | Idaho Falls | Meridian | Twin Falls

Idaho State
UNIVERSITY

Geospatial Data Storage (Vector)

- Geo-spatial data are stored as Feature classes
- Non-spatial data are stored as stand-alone tables
- Vector data is handled by DB2's Spatial Extender. ArcSDE is a "doorman".



Pocatello | Idaho Falls | Meridian | Twin Falls

Idaho State
UNIVERSITY

Geo-spatial Data Storage (Raster)

- Two methods
 - Stand-alone raster data set
 - Mosaic
- ArcSDE is not the best solution to store raster GIS data for the Enterprise
 - Size considerations
 - Performance issues
- Raster data is handled by ArcSDE

Pocatello | Idaho Falls | Meridian | Twin Falls

Idaho State
UNIVERSITY

Internal Data Storage

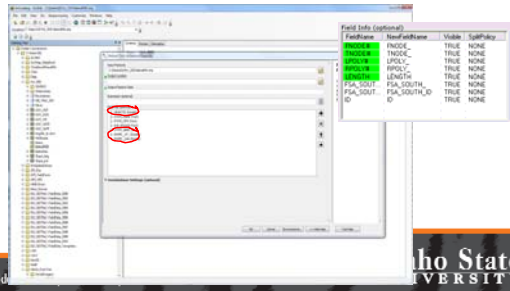
- Within the DB2 RDBMS
 - All data is stored within table spaces –referred to by Configuration Keyword.
 - A Configuration Keyword points to a set of two table spaces:
 - Attribute table space
 - Coords table space

Pocatello | Idaho Falls | Meridian | Twin Falls

Idaho State
UNIVERSITY

Loading Vector Data into a GDB

- PART 1: Stand-alone feature classes



Pocatello | Idaho Falls

Idaho State
UNIVERSITY

Leveraging Common Computing Architecture The Future...



- Works seamlessly across all devices
- Reduces need for custom applications
- Platform for integration with other business systems
- Cross organizational collaboration
- Ready to use content and services
- Content management system

Think about it...

- Object-relational databases have *native* geospatial capabilities
- ArcGIS for Server can make geospatial data available to the Enterprise
- Do we need ArcSDE?

Pocatello | Idaho Falls | Meridian | Twin Falls

Idaho State
UNIVERSITY

Questions...



Pocatello | Idaho Falls | Meridian | Twin Falls

Idaho State
UNIVERSITY

Geodatabases in an Enterprise Workflow

Keith T. Weber, GISP
GIS Director, ISU
GIS Training and Research Center

Pocatello | Idaho Falls | Meridian | Twin Falls

Idaho State
UNIVERSITY

Presentation and Discussion

UNDERSTANDING AND MANAGING WORKFLOW

Pocatello | Idaho Falls | Meridian | Twin Falls

Idaho State
UNIVERSITY

Let's Get Started

- GIS is...
 - Data-driven
 - Powerful
 - Dynamic




• Adjectives
• Adjectives

Pocatello | Idaho Falls | Meridian | Twin Falls

Idaho State
UNIVERSITY

GIS Data Life Cycle

Create Data



Pocatello | Idaho Falls | Meridian | Twin Falls

Idaho State
UNIVERSITY

The Bottleneck

- Distributing the new edition



Postallo | Idaho Falls | Meridian | Twin Falls

Idaho State
UNIVERSITY

The Solution

- Networks and the Internet



Postallo | Idaho Falls | Meridian | Twin Falls

Idaho State
UNIVERSITY

A New Problem is Born

- "MY" version

Postallo | Idaho Falls | Meridian | Twin Falls

Idaho State
UNIVERSITY

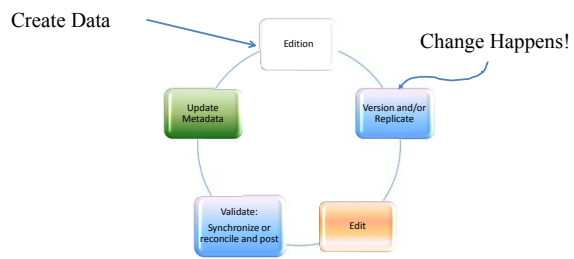
GIS Grows Up!

- RDBMS
 - Keep the benefit of network connectivity
 - Eliminate the problem of “MY” version
 - Eliminate the bottleneck
 - And, change the cycle of events

Pocatello | Idaho Falls | Meridian | Twin Falls

Idaho State
UNIVERSITY

GIS Data Life Cycle



Pocatello | Idaho Falls | Meridian | Twin Falls

Idaho State
UNIVERSITY

Backup vs. Versioning

- Backups and archiving are still critical steps for the enterprise.
- BUT, not part of the GIS Life Cycle any longer

Pocatello | Idaho Falls | Meridian | Twin Falls

Idaho State
UNIVERSITY

In the Beginning...

- Backups were made in case we really messed up
- Edits were made to the original
- Copies of the “clean” new edition were distributed

Today...

- The original [parent] is versioned [a child is born]
- Edits are made to the child, not the parent
- “Clean” edits are copied [synchronized or posted] to the parent.

Benefits Of This Approach

- Brainstorm!!!
 - Minimize downtime
 - Processes completed within the RDBMS



The Role of Backups

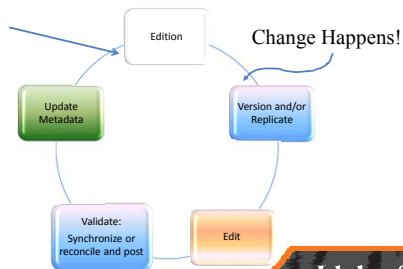
- Data retention and deletion
- Legal requirements

Pocatello | Idaho Falls | Meridian | Twin Falls

Idaho State
UNIVERSITY

GIS Data Life Cycle...Today

Create Data



Pocatello | Idaho Falls | Meridian | Twin Falls

Idaho State
UNIVERSITY

Questions/Discussion?



Pocatello | Idaho Falls | Meridian | Twin Falls

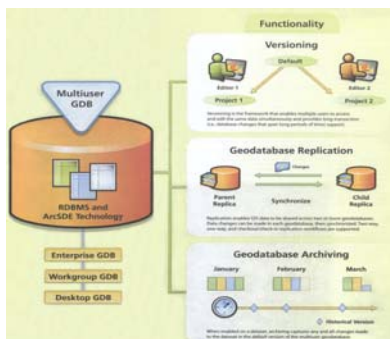
Idaho State
UNIVERSITY

Presentation and Discussion

REPLICATION AND VERSIONING

Pocatello | Idaho Falls | Meridian | Twin Falls

Idaho State
UNIVERSITY




Pocatello | Idaho Falls | Meridian | Twin Falls

Idaho State
UNIVERSITY

What is Replication?

- Duplication
- Copying
- Mirroring




Pocatello | Idaho Falls | Meridian | Twin Falls

Idaho State
UNIVERSITY

True Replication...


- Does not need ArcGIS
- Every RDBMS can be replicated natively
- However, using ArcGIS to perform the replication
 - Is easy
 - Supports GIS workflows better



Pocatello | Idaho Falls | Meridian | Twin Falls

Idaho State
UNIVERSITY

Why Replicate?



- Enable *disconnected* editing for:
 - Performance/load balancing
 - Network load reduction
 - Publishing data to subscribers

Pocatello | Idaho Falls | Meridian | Twin Falls

Idaho State
UNIVERSITY

Network Load Reduction

- The network is a primary bottleneck in the Enterprise

Pocatello | Idaho Falls | Meridian | Twin Falls

Idaho State
UNIVERSITY

How Do I Replicate?

- We will cover this with the hands-on exercise
- As an overview...
 - Version the database
 - Replicate the database
 - Edit/update
 - Synchronize changes with the parent

Pocatello | Idaho Falls | Meridian | Twin Falls

Idaho State
UNIVERSITY

So Replication is Versioning


- No... but replication uses a versioned database

Pocatello | Idaho Falls | Meridian | Twin Falls

Idaho State
UNIVERSITY

What is Versioning?

- One database
- Parent edition (tables) remains live/usable
- Child edition(s) simultaneously edited
- Roll-up is seamless




Pocatello | Idaho Falls | Meridian | Twin Falls

Idaho State
UNIVERSITY


Versioning: Principal Concepts

- Edits are stored in “Supporting Tables”
- Geographic changes (linework) are stored in Supporting Vector Tables
- Attribute changes are stored in Supporting Delta Tables.



Delta Tables

- A = Add (insert)
- D = Delete
- U = Update (delete existing then add)




A Tree is Formed

- As versions are created and changes are made, a tree grows
 - Q: What kind of tree?
 - A: A State Tree




Sort of an Upside-down Tree



Pocatello | Idaho Falls | Meridian | Twin Falls

Idaho State
UNIVERSITY

The State Tree



Pocatello | Idaho Falls | Meridian | Twin Falls

Idaho State
UNIVERSITY

Multiple Versions

- Multiple versions are allowed
 - Versions can be based upon location (north edits, south edits), projects (sub-divisions), or other logic decided upon by the GIS Manager.
- Batch reconcile and post are supported

Pocatello | Idaho Falls | Meridian | Twin Falls

Idaho State
UNIVERSITY

The Day of Reconciliation

- Arthur's Court sub-division edits have been completed
 - Time to reconcile
 - This process looks for conflicts
 - Once all conflicts have been resolved...
 - Reconciliation is complete

Pocatello | Idaho Falls | Meridian | Twin Falls

Idaho State
UNIVERSITY

Post

- To roll-up the edits back to the “trunk of the state tree” we Post

Pocatello | Idaho Falls | Meridian | Twin Falls

Idaho State
UNIVERSITY

Considerations

- Performance can degrade with active databases
 - Workflow itself can generate unnecessary versions
 - Delta tables will become large over time

Pocatello | Idaho Falls | Meridian | Twin Falls

Idaho State
UNIVERSITY

The Cure

- For many of these ArcGIS-centric performance issues is *compressing the database*
 - Moves common rows from delta tables into base tables
 - Reduces depth of the state tree by removing states no longer needed

Compression Example



Questions/Discussion?



Hands-On Exercise

- Practice both replication and versioning

Pocatello | Idaho Falls | Meridian | Twin Falls

Idaho State
UNIVERSITY

Your Assignment

- Complete the exercise handouts
 - Connecting to and using SDE on DB2
 - Practice both replication and versioning
- Read the PDFs in the SDE exercise folder
- Visit the URL link for Spatial Data Server and explore this topic

Pocatello | Idaho Falls | Meridian | Twin Falls

Idaho State
UNIVERSITY



Key Concepts

- ArcSDE is an **engine** residing between a spatially-enabled RDBMS and ArcGIS desktop.
- ArcSDE and the GDB enables GIS for the Enterprise
- ArcSDE reduces data management responsibilities.
- Understand Enterprise workflow

Pocatello | Idaho Falls | Meridian | Twin Falls

Idaho State
UNIVERSITY
