

Geodatabase Archiving: Working with Archived Data

Transcript

Copyright © 2006 ESRI
All rights reserved.

The information contained in this document is the exclusive property of ESRI. This work is protected under United States copyright law and other international copyright treaties and conventions. No part of this work may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying and recording, or by any information storage or retrieval system, except as expressly permitted in writing by ESRI. All requests should be sent to Attention: Contracts and Legal Services Manager, ESRI, 380 New York Street, Redlands, CA 92373-8100, USA.

The information contained in this document is subject to change without notice.

@esri.com, 3D Analyst, ADF, AML, ARC/INFO, ArcAtlas, ArcCAD, ArcCatalog, ArcCOGO, ArcData, ArcDoc, ArcEdit, ArcEditor, ArcEurope, ArcExplorer, ArcExpress, ArcFM, ArcGIS, ArcGlobe, ArcGrid, ArcIMS, ArcInfo Librarian, ArcInfo, ArcInfo-Professional GIS, ArcInfo-The World's GIS, ArcLocation, ArcLogistics, ArcMap, ArcNetwork, ArcNews, ArcObjects, ArcOpen, ArcPad, ArcPlot, ArcPress, ArcQuest, ArcReader, ArcScan, ArcScene, ArcSchool, ArcSDE, ArcSdl, ArcStorm, ArcSurvey, ArcTIN, ArcToolbox, ArcTools, ArcUSA, ArcUser, ArcView, ArcVoyager, ArcWatch, ArcWeb, ArcWorld, Atlas GIS, AtlasWare, Avenue, BusinessMAP, Database Integrator, DBI Kit, ESRI, ESRI-Team GIS, ESRI-The GIS Company, ESRI-The GIS People, FormEdit, Geographic Design System, Geography Matters, Geography Network, GIS by ESRI, GIS Day, GIS for Everyone, GISData Server, InsiteMAP, JTX, MapBeans, MapCafé, MapObjects, ModelBuilder, MOLE, NetEngine, PC ARC/INFO, PC ARCPLOT, PC ARCSHELL, PC DATA CONVERSION, PC STARTER KIT, PC TABLES, PC ARCEDIT, PC NETWORK, PC OVERLAY, PLTS, Rent-a-Tech, RouteMAP, SDE, SML, Spatial Database Engine, StreetEditor, StreetMap, TABLES, the ARC/INFO logo, the ArcCAD logo, the ArcCAD WorkBench logo, the ArcCOGO logo, the ArcData logo, the ArcData Online logo, the ArcEdit logo, the ArcExplorer logo, the ArcExpress logo, the ArcFM logo, the ArcFM Viewer logo, the ArcGIS logo, the ArcGrid logo, the ArcIMS logo, the ArcInfo logo, the ArcLogistics Route logo, the ArcNetwork logo, the ArcPad logo, the ArcPlot logo, the ArcPress for ArcView logo, the ArcPress logo, the ArcScan logo, the ArcScene logo, the ArcSDE CAD Client logo, the ArcSDE logo, the ArcStorm logo, the ArcTIN logo, the ArcTools logo, the ArcView 3D Analyst logo, the ArcView Business Analyst logo, the ArcView Data Publisher logo, the ArcView GIS logo, the ArcView Image Analysis logo, the ArcView Internet Map Server logo, the ArcView logo, the ArcView Network Analyst logo, the ArcView Spatial Analyst logo, the ArcView StreetMap 2000 logo, the ArcView StreetMap logo, the ArcView Tracking Analyst logo, the Atlas GIS logo, the Avenue logo, the BusinessMAP logo, the Data Automation Kit logo, the ESRI ArcAtlas Data logo, the ESRI ArcEurope Data logo, the ESRI ArcScene Data logo, the ESRI ArcUSA Data logo, the ESRI ArcWorld Data logo, the ESRI Digital Chart of the World Data logo, the ESRI globe logo, the ESRI Press logo, the Geography Network logo, the MapCafé logo, the MapObjects Internet Map Server logo, the MapObjects logo, the MOLE logo, the NetEngine logo, the PC ARC/INFO logo, the Production Line Tool Set logo, the RouteMAP IMS logo, the RouteMAP logo, the SDE logo, The World's Leading Desktop GIS, Water Writes, www.esri.com, www.geographynetwork.com, www.gisday.com, and Your Personal Geographic Information System are trademarks, registered trademarks, or service marks of ESRI in the United States, the European Community, or certain other jurisdictions.

Other companies and products mentioned herein are trademarks or registered trademarks of their respective trademark owners.

Hello, thanks for downloading this edition of ESRI's Instructional Series podcast. I'm Owen Evans from the D.C. Technology Center at ESRI in Vienna, Virginia and I specialize in geodata management and SDE technology.

Geodatabase archiving is a new capability that is being introduced at ArcGIS 9.2 that enables the geodatabase to record changes to your data. What you're listening to now is the second podcast on geodatabase archiving. It will cover historical versions and how to use them and introduce several tools that ArcGIS Desktop provides for working with historical data. This material is intended for ArcGIS Desktop users, database administrators, and ArcObjects developers. For more information on the basics of this new capability, check out my first podcast, *Geodatabase Archiving: Introduction to Concepts and Capabilities*.

At ArcGIS 9.1 there was a single type of version that could exist in a geodatabase, the transactional version. A transactional version is a named pointer to a state of the geodatabase and is used to provide editing access in a versioned environment. Transactional versions allow users to connect to different sets of edits that have been made to your geodatabase.

Historical versions are new at 9.2 and exist in geodatabases that contain archived data. Historical versions point to moments in time and allow you to view the geometry and attributes of archived records as they existed at that moment. Connecting to historical versions allows read-only access to the geodatabase to view the changes to archive-enabled classes. Once you have connected to a historical version, you define the specific moment you wish to view. You can connect to any moment in time. And, if data has been archived for that time, the display will reflect the data at that moment.

Historical versions are assembled from queries on the archive class and are completely independent from transactional versions. They do not reference the usual versioning tables such as the Adds and Deletes tables. You can connect to historical versions of your database just like you connect to transactional versions. This is done in ArcMap by right-clicking on a workspace in the Source tab of the table of contents and selecting Change Version. In the Change Version dialog you now have a choice as to what type of version you wish to view. Selecting Historical from the Version Type drop-down menu changes the dialog to allow the selection of a historical version. The server where the geodatabase exists and the ArcGIS Desktop client machine can be in different time zones, and it is important to remember that all archive operations are logged using the time on the geodatabase server.

There is a button on the Change Version dialog that allows you to obtain the current time on the database server. This button can be used to connect to the current time through the historical version or simply to compare the database time to your local time. It is possible to save connections to historical versions in ArcCatalog for personal, workgroup, and enterprise ArcSDE geodatabases. To do this for an enterprise geodatabase, go to the Database Connections node in ArcCatalog and double-click Add Spatial Database Connection. In the Connection Properties dialog, fill in the connection information for

your geodatabase and in the Connection Details section, click the Change button. Then select a historical version for the new database connection. This process is slightly different for personal and workgroup ArcSDE geodatabases. Consult the ArcGIS Desktop Help topic called *An overview of database server connections* for more information.

So what are the different tools that we can use to work with historical data? The first one is the Historical Marker Manager, which can be opened from the new Geodatabase History toolbar. Note that this tool is only enabled when you have a workspace that contains archived data selected in the Source tab of the table of contents. The Historical Marker Manager allows you to create, modify, and delete named aliases that correspond to moments in time that have meaning to you and your colleagues.

For example, before I start making edits that correspond to the recent construction of a new housing development called Hidden Lakes, I can create a historical marker that corresponds to the present moment called Hidden Lakes Preconstruction. You can name your historical markers anything you like except Default, which is a reserved name. The Default historical marker is automatically generated, cannot be deleted, and references the current state of the archive classes. This marker is equivalent to the Default transactional version.

Historical markers are accessible from all dialogs where you are prompted to choose a historical version. The marker names and the moments they reference are stored in the geodatabase in a table called `gdb_historical` markers that is owned by either SDE or DBL. All database users have privileges to add, modify, and delete historical markers and can see all markers that have been created.

I mentioned that the archive process uses the database transaction time to record changes. Historical markers provide a way to associate database transaction times with text descriptions of valid time events. Let's suppose that the construction of our fictitious development was completed on July 14th, 2006. The corresponding edits to our geodatabase might not have been posted until August 24th, 2006. However, if I create a descriptive historical marker called Hidden Lakes Construction Complete for the moment when these edits were posted, users of this geodatabase do not need to be concerned with the difference between transaction time and valid time. They can simply connect to the historical version that corresponds to the marker I created.

The History Viewer is another tool that helps users work with archived data. This tool is a dockable window in ArcMap that can be made visible from the Geodatabase History toolbar. If my map document does not have any connections to historical versions, the History Viewer will be disabled. To enable it, simply change the version of one of the workspaces in the active data frame to a historical version. When activated, the History Viewer displays a list of all the layers and tables in my table of contents that have historical data, so you know what map layers the historical versions reference. If there are archive data layers for more than one geodatabase in the active data frame, the drop-down at the top of the History Viewer will allow you to select between the historical

versions for these geodatabases. ArcMap can only display one historical version at a time.

When you enter a date and time in the History Viewer, the map view shows the archive data for the selected moment for any map layers that have associated archive classes. So if I select January 1st, 2006 at 12:37 p.m., the map view will change to show how the features in my archive layers looked at that moment. Remember, any historical markers that have been created are available for all users to choose from in the History Viewer. So even after I post the edits for the Hidden Lakes Development to my Default transactional version, I could select my Hidden Lakes Preconstruction marker to see what the data looked like before this new development existed.

Another tool for working with historical data is the Add Historical Archive tool. This tool becomes enabled if a layer with an associated archive is selected in the table of contents. Clicking on this tool will add to the table of contents the archive class for the selected layer. Any number of archive layers can be added to a map document and more than one archive layer can be added for the same archive class. These historical archive layers behave just like any other layer in a map document and there are several ways to utilize them.

Since the historical information is stored as attributes in the archive class, you can symbolize the archive layer by one of the date attributes such as the GDB To date. For example, if a red to green color ramp was used to symbolize sewer lines, the older pipes would be orange or red, while the more recently installed pipes would be yellow and green. For a polygon feature class like a parcel layer, where the features are often joined or subdivided, this might not be useful since many features in the archive layer could overlap. In this case, you could add several archive layers for your parcel feature class to the table of contents and set up a different definition query for each one. Each query would reference a different date or range of dates. These layers could be made transparent so that all the features that existed could be seen at the same time. You can also use the Identify tool to identify on the archive class at a particular point. This will show all the changes that have occurred to the data at that location.

The availability of historical data offers the opportunity to create several types of historical queries. These queries can be written in the Select By Attributes dialog or the Definition Query field in the archive layer properties. One type of query is a moment-in-time query. An example of this would show the state of the parcel data on May 13th, 2005 at 12:15 p.m. It would display all the features that existed in the Default version at that moment in time. The syntax for this query would be: GDB From date is less than or equal to May 13th, 2005 at 12:15 p.m. AND GDB To date is greater than May 13th, 2005 at 12:15 pm. If you write a moment-in-time query referencing a point before archiving was enabled, your query will return no data.

Another type of query you can write is a range query. An example of this type of query would be to select all water quality sampling results for June 2005 that are greater than 500 milligrams per liter. The syntax for this query would be: GDB From date is greater

than or equal to June 1, 2005 at 12 a.m. AND GDB To date is less than or equal to June 30, 2005 at 11:59 and 59 seconds p.m. AND test result is greater than 500. Dates must be in single quotes and in proper date format for the query syntax to be valid. Querying the archive class for a particular feature without specifying a date will return all the archived changes to that feature. For example, if you query on the parcel data for APN546, the result set would be every edit that has been archived for that particular parcel, answering the question, how has this parcel changed over time?

Also, you could write a query for a specific attribute without specifying a date and this would return all features that satisfy the attribute query since archiving was enabled. Keeping with our water quality example, you could query a lake feature class for archived features with values greater than 750 milligrams per liter. This would tell you which lakes exceeded this value at any time since archiving was enabled. Queries on the archive table are very efficient since the query is running against a single indexed table and does not involve table charts.

Finally, a word for database administrators about keywords for history tables. You can create keywords that define the storage parameters for history tables by appending underscore archive (`_archive`) to an existing keyword; for example, `defaults_archive`. When a history table is created, the keyword for the archiving table is noted and if there is a corresponding keyword for the history table, the parameters specified for that keyword are used to create the history tables.

For more information on geodatabase archiving, consult the ArcGIS Desktop 9.2 Help which is currently live on the Web at webhelp.esri.com. Also look for the three-day instructor-led course called [Managing Editing Workflows in a Multiuser Geodatabase](#) offered by ESRI Educational Services.

Thanks for tuning in to this edition of ESRI's Instructional Series podcast and stay tuned for future broadcasts.