

Geodatabase Archiving: Introduction to Concepts and Capabilities

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, and welcome to the ESRI Instructional Series podcast on geodatabase archiving. I'm Owen Evans from the D.C. technology center at ESRI in Washington, D.C., and I specialize in geodata management and ArcSDE. There will be several podcasts in this series on geodatabase archiving. The first, which you are listening to now, will present an introduction to the concepts and capabilities of archiving, while subsequent podcasts will cover ways to leverage this new capability, and also get into the specifics of how users interact with archived data. This podcast is intended for ArcGIS Desktop users and database administrators. Please note, the functionality discussed here is scheduled to be included in ArcGIS 9.2. The information presented, therefore, is subject to change pending the ArcGIS 9.2 final release.

What is geodatabase archiving? Geodatabase archiving is a new capability that will be introduced at ArcGIS 9.2 that enables the geodatabase to record changes to your data. It is essentially a way of preserving your data through time, without storing additional complete copies of your datasets for each moment you require. This archived data can be viewed, queried, and analyzed by ArcMap and other applications. When you enable archiving on a feature class, changes to that feature class are automatically archived in a separate table called an archive class. There is one archive class per feature class on which archiving is enabled. Archiving can be enabled on tables, stand-alone features classes, or entire feature datasets.

How does the geodatabase implement archiving? Currently (that is at version 9.1 and earlier releases of ArcGIS), to preserve a snapshot of your data, you would need to keep a version of your geodatabase that corresponds to a specific point in time, such as the completion of a new utility project, or the end of a fiscal year. This type of version references all feature classes contained in your geodatabase, whether you are interested in preserving a historical record of them or not. The other way to save a snapshot would be to create a separate copy of your data for each moment you require. With the new archiving functionality, the geodatabase will archive changes to the feature classes you specify, so it won't waste space storing information about layers that are not important to your organization's history. A point to note here is, that only changes made and posted to the Default version of your geodatabase are archived. Changes to other named versions are not archived.

How is archiving enabled? To enable archiving, there are some conditions that must first be met. First, your data must reside in an ArcSDE 9.2 geodatabase. Since the archiving feature will be introduced at version 9.2, there is no way to enable this feature on pre-9.2 databases, even if you

are using version 9.2 of ArcGIS Desktop. I should also note that archiving works with any ArcSDE geodatabase. This includes ArcSDE Enterprise Edition, as well as the new Workgroup and Personal editions of ArcSDE that are being introduced at 9.2. Second, the data on which you plan to enable archiving must be registered as versioned. Be sure that you do not register the data as versioned with the new save edits to base option. Data must be fully versioned to participate in archiving. Once these conditions are met, you can enable archiving by right-clicking on your feature class and then selecting Enable Archiving from the Archiving pull-right menu.

What happens when you enable archiving? When you enable archiving, a table called an archive class is created for each class you are enabling. All rows in the Default version for the classes being enabled are copied to the archive class. The archive has the same schema as the original class with the addition of three new fields: two date attributes named GDB_FROM_DATE and GDB_TO_DATE, and a numeric GDB archive OID attribute that uniquely identifies each row. The GDB_FROM_DATE and GDB_TO_DATE attributes store the birth date and death date of features in the archive class. These essentially describe the lifetime of the feature as it is represented for any given moment in time. If you want to find out when archiving was enabled on a particular feature class, right-click it in ArcCatalog and go to the General tab of its properties. This dialog box contains information about if and when archiving was enabled, and also if and how it is registered as version.

What happens when you edit or post changes to a Default version that has archive-enabled data? There are a few concepts that we should discuss at this point. These are the concepts of valid time and transaction time. Valid time is the actual time an event occurred in the real world. For example, the moment when a legal transaction is recorded, such as a deed of ownership. Transaction time is the time at which that change is recorded in the database. Geodatabase archiving uses transaction time to record changes, and this is an important point to remember when working with archiving.

There are three types of changes that can happen in a geodatabase: a new feature can be created, an existing feature can be modified, or an existing feature can be deleted. These three changes are handled in different ways by the archive process. Let's look at the first case.

If a new feature is added to the Default version of your archive-enabled feature class, a new row is created in the archive with the attributes of the new feature. The GDB_FROM_DATE attribute

in that row is set to the date of the archive operation, and the GDB_TO_DATE is set to infinity, which is represented by December 31 in the year 9999. A row with a to-date of infinity corresponds to a feature that is currently active in the database. A special case of this situation occurs when you initially enable archiving on a feature class. All features in a Default version are copied to the archive, and the GDB_FROM_DATE in all rows is set to the moment when archiving was enabled. The GDB to date in all rows is set to infinity, since the copied rows correspond to active features in the Default version. Otherwise, rows are added to your archive either through direct edits of the Default version or by posting edits to Default from another version.

In the second case, when an existing feature is modified or updated, the current date is entered into the DB_TO_DATE attribute as the death date in the row that corresponds to that feature in the archive class. Then, a new row is inserted with the modified attributes for the feature. The inserted row obtains the time stamp of the archive operation as its birth date, and infinity as its death date.

When a feature is deleted, the DB_TO_DATE in the corresponding row in the archive is updated with the moment of the delete. Recall: The archiving process uses the database transaction time when recording changes to the archive classes.

There is a fourth type of change that you can make to a feature class, and that is a schema change. If you add or move an attribute of an archive-enabled feature class, the archive will account for this change. See the ArcGIS Desktop 9.2 Help for details about how this is handled.

How is archiving disabled? To disable archiving, right-click on your feature class and select Disable Archiving from the Archiving pull-right menu. This will terminate the growth of the archive, but it will not delete the archive from the geodatabase. You can still access archived data associated with this feature class, however, new changes to it will no longer be captured. If you delete a feature class that is or was archive-enabled, the archive class will not be deleted. It will be exposed as a stand-alone class and will appear in ArcCatalog with a feature class name appended with *_H*. For example, if you deleted an archive-enabled layer called *parcels*, an archive class would appear at the root level of your geodatabase called *parcels_H*. You may need to refresh your view to see the archive after deleting a feature class. Archiving can be re-enabled after it has been disabled. However, this will create a brand new archive class. An existing

archive cannot be reassociated with an existing feature class when re-enabling archiving. There is no concept of suspending the archive process for a period of time and then re-enabling it where it left off. This will produce a gap in the archive, and this is not supported.

More information on geodatabase archiving will be available in future podcasts in this series and in the ArcGIS Desktop Help once version 9.2 is released.

This concludes this session of ESRI's Instructional Series podcast. Thanks for tuning in, and stay tuned for future broadcasts.