



Environmental Systems Research Institute, Inc., 380 New York St., Redlands, CA 92373-8100 USA • TEL 909-793-2853 • FAX 909-307-3014

ESRI Systems Integration Technical Brief

ArcSDE Geodatabase Replication Overview

Brian Price
Senior Systems Consultant
Robert Kircher
Senior Technical Designer
ESRI, Redlands, CA

September 20, 2004

Rev. 2

Copyright © 2004 ESRI
All rights reserved.
Printed in the United States of America.

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 Manager, ESRI, 380 New York Street, Redlands, CA 92373-8100, USA.

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

U.S. GOVERNMENT RESTRICTED/LIMITED RIGHTS

Any software, documentation, and/or data delivered hereunder is subject to the terms of the License Agreement. In no event shall the U.S. Government acquire greater than RESTRICTED/LIMITED RIGHTS. At a minimum, use, duplication, or disclosure by the U.S. Government is subject to restrictions as set forth in FAR §52.227-14 Alternates I, II, and III (JUN 1987); FAR §52.227-19 (JUN 1987) and/or FAR §12.211/12.212 (Commercial Technical Data/Computer Software); and DFARS §252.227-7015 (NOV 1995) (Technical Data) and/or DFARS §227.7202 (Computer Software), as applicable. Contractor/Manufacturer is ESRI, 380 New York Street, Redlands, CA 92373-8100, USA.

@esri.com, 3D Analyst, ADF, AML, ARC/INFO, ArcAtlas, ArcCAD, ArcCatalog, ArcCOGO, ArcData, ArcDoc, ArcEdit, ArcEditor, ArcEurope, ArcExplorer, ArcExpress, ArcFM, ArcGIS, ArcGrid, ArcIMS, ArcInfo Librarian, ArcInfo, ArcInfo—Professional GIS, ArcInfo—The World's GIS, 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 People, FormEdit, Geographic Design System, Geography Matters, Geography Network, GIS by ESRI, GIS Day, GIS for Everyone, GISData Server, *InsiteMAP*, 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, Rent-a-Tech, *RouteMAP*, SDE, SML, Spatial Database Engine, StreetEditor, StreetMap, TABLES, the ARC/INFO logo, the ArcAtlas logo, the ArcCAD logo, the ArcCAD WorkBench logo, the ArcCOGO logo, the ArcData logo, the ArcData Online logo, the ArcEdit logo, the ArcEurope 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 ArcUSA 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 ArcWorld logo, the Atlas GIS logo, the Avenue logo, the *BusinessMAP* logo, the Data Automation Kit logo, the Digital Chart of the World logo, the ESRI 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.

Table of Contents

Background	4
Scope	4
Why Replicate Data?	5
Performance/Load Balancing.....	5
Network Load Reduction	5
System/Data Security.....	6
Spatial Data Partitioning	6
Disconnected Computing.....	6
Enterprise-wide use of GIS Data	6
Server Consolidation/Reorganization	7
High-Availability Architectures.....	7
Disaster Recovery	7
Common Replication Strategies.....	7
Export/Import (or database dump/restore).....	8
Data Storage-level Replication	8
Snapshot Replication	10
Transactional Replication (Updateable Snapshots)	11
Filtered/Partial Replication	11
Disconnected Editing.....	12
Distributed Geodatabase Replication.....	14
Multi-Master Replication.....	15
Replication Misconceptions.....	16
ArcGIS, GeoDatabase, and ArcSDE Replication Challenges.....	17
Conclusion	19
Bibliography.....	20

Background

Traditionally moving and duplicating spatial data around an enterprise GIS has been a daunting endeavor. Whether it was done manually or through some special automation, keeping data synchronized and timely without overly burdening the system challenged even the most ambitious GIS administrator. With the increased usage of geodatabases such as those based on ESRI's ArcSDE and with replication services found in DBMS', many ArcGIS enterprises are exploring the notion of employing DBMS replication more than ever to share and duplicate spatial data.

Moreover, GIS and IT enterprises are being pressed to grow their databases, grow their infrastructure, support remote users, provide more spatial data services, and achieve adequate levels of performance and service while doing so. Having multiple data servers and maintaining them via replication seems a logical method to satisfy these demands.

Although building an effective replication infrastructure is certainly possible, designing, creating, and maintaining one can be a very complex endeavor. In addition, replicating data in an ArcSDE geodatabase environment requires special accommodations.

This paper explores the realities of replicating data in an ArcGIS enterprise using an ArcSDE geodatabase, and an industry standard DBMS. It is targeted for system implementation architects to help them understand what is possible and what is not (at a fairly high level) in designing ESRI GIS systems. The paper includes discussions about:

- ***Why** replicate data?*
- *What are the common replication **strategies**.*
- *What are common **misconceptions** about replication?*
- *What are the typical **challenges and risks** in building a replicating geodatabase?*

The ArcSDE Configuration and Tuning Guides describe how customers replicate data in a specific DBMS, so this discussion explores the “when”, “what’s” and “why’s” when considering a replication scheme for an organization.

Scope

For this discussion, “replication” is the process of duplicating spatial data stored in an industry standard DBMS, like Oracle or SQL Server, managed by ESRI ArcSDE. We will look at vendor software database replication, storage-level replication such as disk mirroring, and vendor software geodatabase replication. Clearly other options such as custom programming are available for organizations to move, copy and synchronize data in their organization. A discussion of these is beyond the scope of this document, however. We also will not address the details of functionality or implementation techniques as these are very product-dependant, and can change quickly. Consult the database vendor, third-party vendor (if appropriate) and ESRI Support website for additional help in these areas.

Why Replicate Data?

The simple answer is to create more than one copy of a geodatabase. So why do we need multiple copies? Databases have been copied/replicated over the years for many different reasons and purposes. The following discussion outlines some of the more prevalent uses of replicated data with an emphasis on their pertinence to an ESRI GIS setting. These include:

- Performance/Load Balancing
- Network Load Reduction
- System/Data Security
- Spatial Data Partitioning
- Disconnected Computing
- Enterprise Use of GIS Data
- Server Consolidation/Reorganization
- High-availability Architectures
- Disaster Recovery

Performance/Load Balancing

Replication can be used to spread the geodatabase processing load across several servers. This is most often considered when organization's users are dispersed geographically. ArcSDE servers local to those users can provide faster, more reliable access to the data.

Moreover, some enterprises organize their users around specific functions, like OLTP vs OLAP (DSS, data marts, data warehouses), and create separate, but duplicate database instances in which to work from. For instance, because the GIS administrator noticed that the map products group was really bogging down the centralized database, he decided to partition them from the editors. So, database1 became a high transactional system, doing a lot of editing and spatial data maintenance, (Online Transactional Processing - OLTP) and be dedicated to server1 to maintain key levels of performance. Meanwhile, all the new or updated data could be replicated to database2/server2 for producing reports or map products (Online Analytical Processing - OLAP).

Network Load Reduction

When GIS users are geographically dispersed, replication can be used to distribute data to ArcSDE servers that are close to the users of that data. This configuration can reduce the network load required to support communications between interactive users and the ArcSDE database, and for printing and plotting maps which requires moving large amount of data from the workstation (or applications server) to the printer. It can also provide higher availability to those remote users who otherwise would be dependant on WAN connections to access data.

Of course there is no free lunch - the bandwidth required for the replication process must be considered. If replication can occur during off-peak hours or utilize non-network resources (i.e. sneaker-net), then a net reduction in network traffic can be obtained for the enterprise during peak hours.

System/Data Security

Organizations often publish GIS data and services to the web or to external organizations using ArcIMS or ArcGIS Server. Maintaining data security in this environment requires careful network design and server placement. Placing a production server outside the secure network, for instance, so that web users can access it is clearly not advisable. So several strategies have evolved over time to provide the data publishing services while minimizing exposure to production databases. This usually involves placing the application servers (ArcIMS, AGS) for external users in a network demilitarized zone (DMZ) and placing the data source for those servers either inside the secure network, or on a separate read-only database inside the DMZ. When an additional data server is placed in the DMZ the data can be maintained via replication from the master database by opening up the ArcSDE port for Application Server connections or the DBMS client port for Direct Connections. Snapshot replication is a common method for this architecture.

Spatial Data Partitioning

In larger organizations there are often several different functions operating somewhat independently but which must coordinate with each other and with a central office. Ideally a GIS could support this kind of organization by providing methods for these independent offices to work on their geographic or thematic areas of interest on local databases while maintaining the ability to consolidate these separate databases into a single corporate data store as required. This allows organizations to create and enforce data standards while permitting independent data maintenance on separate data stores. Two-way replication is required to support this.

Disconnected Computing

Users may wish to view and edit GIS data utilizing laptops, PDAs and other mobile devices which are not well connected to the organizational network. They need to access their data locally without requiring a full-time connection to the central ArcSDE database. A disconnected computing solution allows them to create an extract of their data from the master database, view and edit that data, and then synchronize their changes with a master database periodically.

Enterprise use of GIS Data

Increasingly, organizations are deploying a variety of applications, many of which may not be GIS-centric, that require access to GIS data in specific formats or views which are different than the master database schema. Examples of this might be ERP software such as SAP or Peoplesoft which have spatial enabled components. Deployment templates enable organizations to create multiple customized snapshot environments relatively quickly. Users can utilize variables to customize each snapshot environment for its individual needs. For example, one could create a sales force automation template containing variables for each sales regions and salespersons.

Server Consolidation/Reorganization

Database replication has been used to assist in the process of consolidating databases, moving databases from one instance or server to another and upgrading versions of DBMS. Replication services offer a way to migrate data across the infrastructure. More commonly database export/import or database dump/restore functions are used for this function in the GIS environment as it is less complex and more reliable.

High-Availability Architectures

Highly-available systems recover quickly (seconds to very few minutes) from a server or other platform failure. Replication can be used to improve the availability of GIS applications by providing them with alternative data access options in the case of a data server or database failure. If the database server crashes, for instance, users can continue to access other servers where the data has been previously replicated.

A more common method for achieving high-availability is utilization of database clustering with automatic failover capabilities. With this solution several ArcSDE servers with access to a common data store (where the database tables are located) are clustered and configured for failover. In this case the data is not replicated, but if the primary database server fails, a secondary server can continue to provide database services.

High-availability and disaster recovery are separate topics, although certainly related. Disaster recovery implies a complete data center or other larger-scale outage. Replication for disaster recovery is covered in the next section.

For additional information on ArcSDE High Availability you can download/view a technical paper on the ESRI Systems Integration website at <http://www.esri.com/systemsint/kbase/techbriefs.html>.

Disaster Recovery

Disaster recovery usually requires the recovery of data and servers as a result of a catastrophe to a data center. A disaster recovery site requires standby application and data servers, and the access to the most current data available. Although DBMS' offer specific backup and recovery services to backup and restore a database offsite, some enterprises have used database replication to store a snapshot of the database on a remote server. This works well for read-only applications in particular. Storage-level replication can also be used to provide a near real-time copy of data to data stores in standby data centers.

Common Replication Strategies

To support the business and technical requirements described above, a variety of techniques can be employed. In general the solutions come from four different sources:

- ESRI ArcGIS/ArcSDE functionality
- Database Management Systems
- Third-party database and storage replication solutions
- Customer-built replication implementations

In this section we will outline the various replication strategies that are available starting with the least complex, most common with the least risk, and then proceeding to the more complex solutions. Some of the issues to consider when contemplating a replication solution include:

- Complexity of implementation and ongoing management.
- Cost of implementation and ongoing management.
- Robustness of solution (reliability, availability)
- Impact on server performance due to replication processes
- Network Utilization (bandwidth availability and impact)
- Geodatabase Integrity Issues
- System/Data Security Issues
- Data “Freshness” Requirements (synchronous vs asynchronous, frequency of replication)

We will touch on the most important issues for each strategy as we move through these topics. At the same time we will point out the most likely potential uses for each, and typical sources for acquiring the technology.

Export/Import (or database dump/restore)

This utilizes standard facilities which come with the DBMS software and/or ArcSDE. These can be executed manually or automated to some degree. They can be used for backing up databases, creating read-only versions of databases, and re-building databases. This technique uses DBMS (or ArcSDE) specific commands and has few geodatabase specific ramifications. Therefore this is very reliable and a low risk method of providing geodatabase copies and is supported by all DBMS supported by ArcSDE.

Since the import/export function does not require network services (export copies can go directly to disk or tape), it will have little or no effect on network utilization. It does require resources on the publishing and subscribing servers, but this processing can be scheduled to occur during off-peak hours. Because it is straightforward and low-risk (even with geodatabases) it is heavily used throughout the GIS community.

Export/import techniques do not provide synchronous/real-time data replication as they copy the database at a point-in-time and there is no method for incrementally updating the copies. This method of copying data is very secure as it usually does not require moving information across a Wide Area Network (WAN).

Potential Uses:

- Server Consolidation/Reorganization
- Disaster Recovery

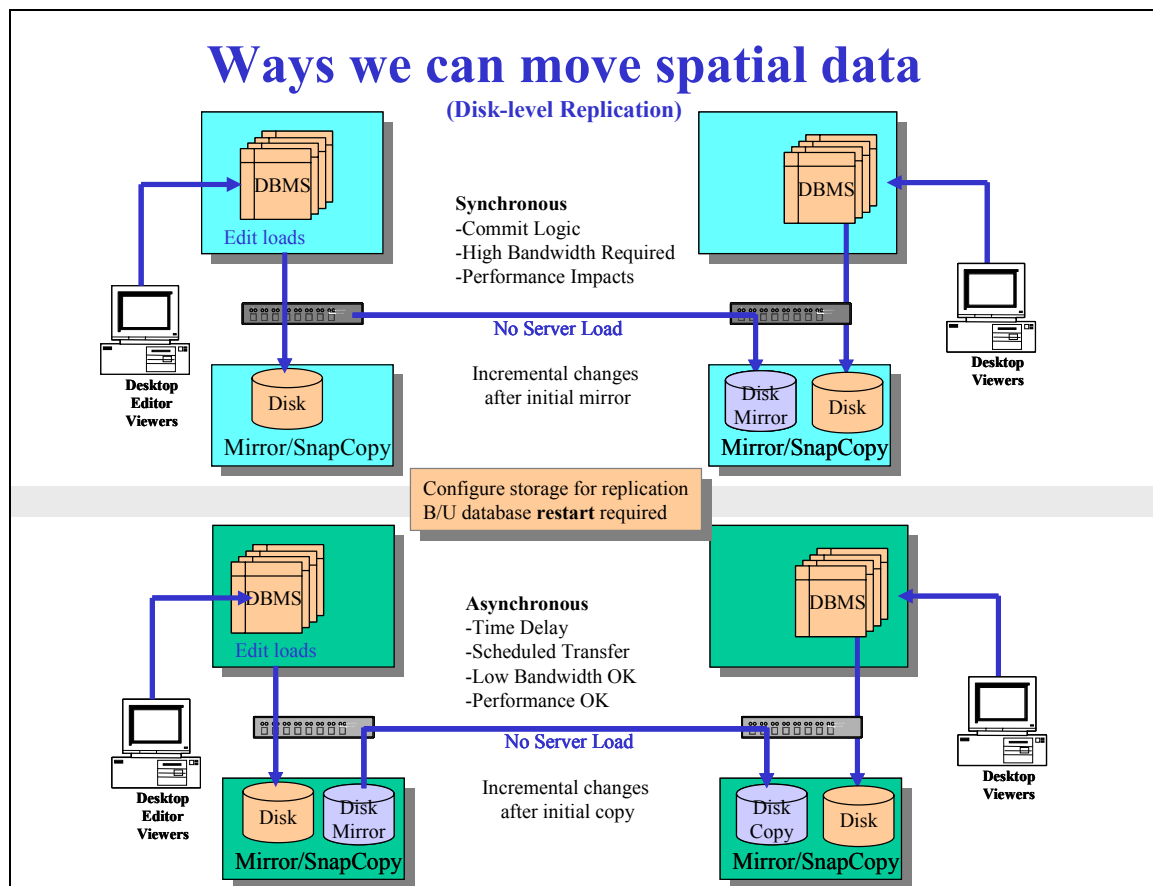
Data Storage-level Replication

This form of replication does not rely on ArcSDE or database functionality and is implemented through data storage subsystem features. In essence it provides near real-time hardware level replication of the database which has a number of advantages in the geodatabase environment. Using hardware-level replication reduces the DB server

overhead of replication by offloading it to the storage subsystem. Neither ArcSDE nor the underlying database software is even aware that this data replication/copying is occurring. This has implications for any database attempting to access the disk storage being maintained from the primary storage.

If there is a geodatabase instance connected to the secondary storage it will be unaware of any changes made to the underlying data since it did not create those changes. The DBMS in-memory data structures do not represent the actual state of the database, so database reads and writes will likely fail. So in general, the secondary data storage is not accessed unless the primary data storage system fails. The geodatabase instance accessing the secondary storage would need to be re-started before it could properly access the updated database. (This is required when the database management software is not managing the updating of database tables – the in-memory data structures do not reflect the physical storage). This synchronous form of replication can be used effectively for maintaining near real-time data to a standby disaster recovery site, for instance.

One of the main advantages of this technique is that it avoids the complexities of trying to use database or other software to maintain geodatabase integrity for the replicas. Replication is at the disk block level, so the replica(s) can be fully functioning versioned databases.



Storage-level replication over a WAN or LAN certainly will impact the network as every change to every block on the primary storage subsystem will need to be propagated to the secondary system. As with most replication solutions, an analysis of required network bandwidth would be required as part of the planning process.

Data storage vendors will be the best source of information on the performance and functionality of these systems.

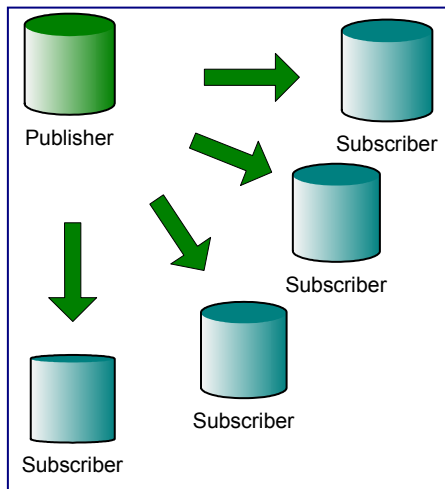
Potential Uses:

- High-availability Architectures
- Disaster Recovery

Snapshot Replication

This can be used for creating read-only versions, or backup copies of databases. This is used in organizations to reduce network traffic for remote user GIS software usage and printing/plotting by providing users with a local read-only ArcSDE instance. It is also used to maintain data for Internet/Intranet publishing servers such as ArcIMS and ArcGIS Server. It uses functionality provided with the third-party DBMS or other replication products and the implementation details vary with the individual product. It is in use in the GIS community, primarily with Oracle and MS SQL Server based solutions.

Snapshot replication creates an image of the database at a point in time. In the Oracle environment, replicating spatial objects may require some data conversion. ArcGIS contains tools that can be used for this conversion.



Snapshot replication requires that the database tables be copied from one location to another. Clearly this could have a major impact to the network depending on timing, frequency of replication, size of the database, number of updates, and the available network bandwidth. A geodatabase can vary in size from 1 GB to 20 TB or more so a determination of the feasibility of this method would normally include a network suitability analysis. Creating a snapshot will also place a load on the master database but this

processing could also occur during off-peak hours to reduce impact to the interactive users. Since data is being moved across a network there could be security implications as well. Note that this is for creating read-only databases and editing is not supported on these replicas.

Snapshot replication could be used for disaster recovery, especially for read-only databases. For editing databases the situation becomes more complex, but should work in a non-versioned, non-disconnected editing environment. The replica would require sequence-ids to be updated before it would be editable. Child replicas of the master database would generally not be able to post their changes to the new parent database since the sequence-ids are not synchronized on the replica. Customers should thoroughly

test their replication disaster recovery configuration to insure that it meets the needs of the organization.

Potential Uses:

- Performance/Load Balancing
- Network Load Reduction
- System/Data Security
- Enterprise Use of GIS Data
- Disaster Recovery

Transactional Replication (Updateable Snapshots)

Some products provide the capabilities for periodically updating snapshots/database dumps instead of re-creating entirely from scratch each time. The transaction update process can take place periodically or can provide near real-time data synchronization. The choice will depend on the requirements for real-time data and network bandwidth considerations.

In general this method reduces the amount of network traffic required to support replication as only changed data is moved across the network. To maintain data integrity periodically a complete snapshot refresh is required. Not all snapshot replication products support this in an ESRI geodatabase. OmniReplicator from Lakeview Technology is one example of a solution that is in use with ESRI ArcSDE as is Transactional replication with MS SQL Server. The Distributed Geodatabase functionality in ArcGIS 9.1 will also provide this capability for editable databases and will more readily support the complexities of a versioned geodatabase environment, especially for editing.

Updateable snapshot technology could also be used for creating disaster recovery instances in some scenarios. See the discussion in Snapshot Replication.

Potential Uses:

- Performance/Load Balancing
- Network Load Reduction
- System/Data Security
- Enterprise Use of GIS Data
- Disaster Recovery

Filtered/Partial Replication

This form of replication is used to publish read-only subsets of databases, usually re-organized for a specific business purpose and with database schemas modified to support the data views required. This typically requires third-party database software functionality. It is most commonly used for data warehousing and data mart applications, for supporting remote printing, or with decision support applications with specific data requirements. In the GIS world, a common implementation of this is with Oracle databases utilizing the “Materialized View” replication function. Other products such as Lakeview Technology’s OmniReplicator provide similar functionality as well. The

Disconnected Editing function in ArcGIS 8.3 and higher, and the future ArcGIS 9.1 Distributed Database functions will provide this functionality for editable databases.

Adding spatial objects to these materialized views can require special user created scripts (or specific vendor software support), and some data conversion may be required in an Oracle environment to support the Oracle Spatial data format.

This form of replication may use less network bandwidth as only subsets of the publishing database need to be moved across the network. As with snapshot replication, with some solutions the replicated databases can be incrementally updated – they do not require complete replacement and near real-time synchronization is possible. Creating a materialized view replica will place a load on the master and replica databases, but this processing could also occur off-peak hours to reduce impact to the peak-time users of the server and the network.

Potential Uses:

- Performance/Load Balancing
- Network Load Reduction
- System/Data Security
- Enterprise Use of GIS Data
- Disaster Recovery

Disconnected Editing

There are several options available for this purpose, including third-party industry solution products such as Go Sync! Tadpole, and Miner and Miner's Mobile Designer for utilities. Most of these currently use ESRI personal geodatabases created from the master database, and provide methods for applying the PGDB field changes to the ArcSDE master geodatabase. The implementations vary, but several support ArcPad based editing and other mobile devices

The ESRI ArcGIS Disconnected Editing function is available in ArcGIS 8.3 and above. This allows users to check-out a version of the master database to their desktop, laptop, PDA or server, disconnect from the master database and from the network, perform edits on their local version, reconnect to the master database network, and then post the changes to the master database through a reconcile and post process. This is an effective solution for disconnected computing applications such as field editing of GIS data using mobile devices such as laptops or PDAs, or remote or regional office GIS editing. This works best when the user has periodic access to the LAN where the data server is located. Otherwise the database must be created over the WAN which, depending on the size of the DB and available network bandwidth, may not be feasible. The master update process (reconcile/post) requires only the changed data to transit the network and this minimizes the impact to the network.

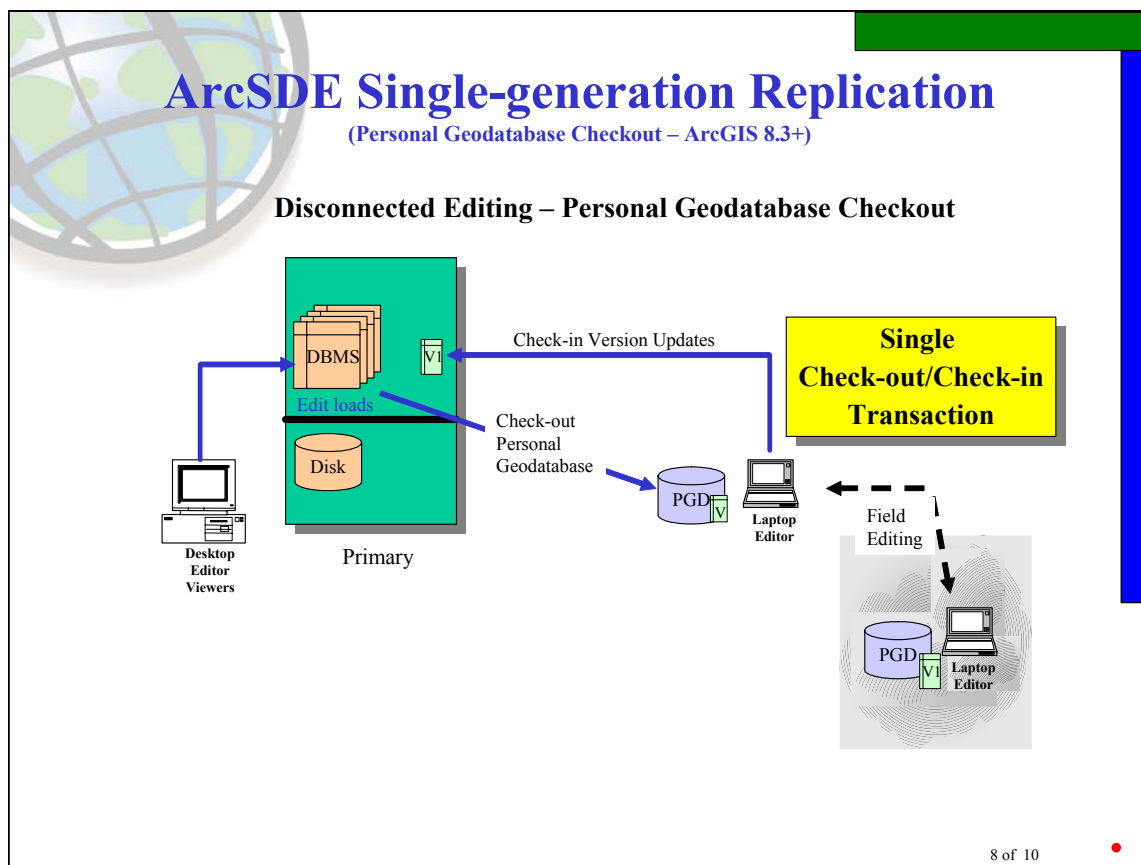
The reconcile/post process provides a resolution method for conflicts that can occur when multiple people have edited/changed the same feature, attribute or other database object concurrently. Conflicts are identified and then users can decide which changes should be

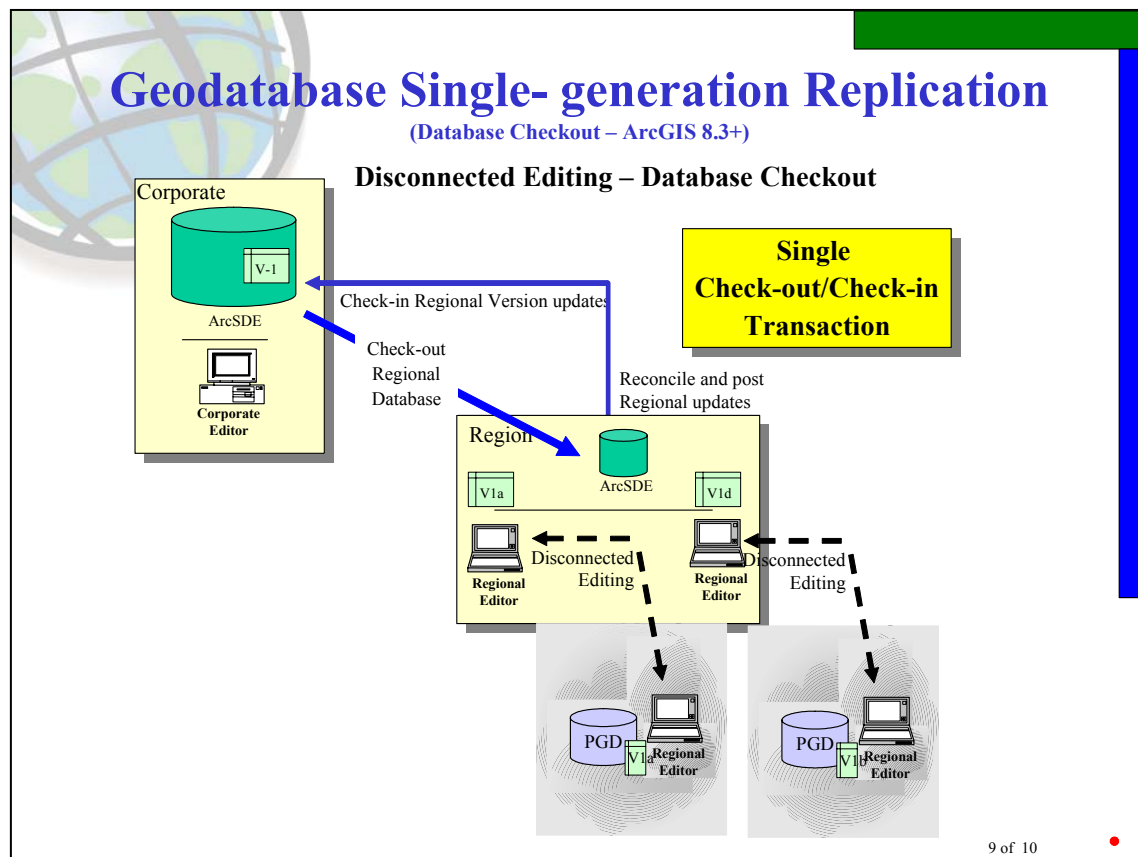
posted to the master database. This overcomes one of the main problems with editing on multiple databases and then incorporating those changes into a master database.

The diagrams below show the two ‘flavors’ of this technology; Personal Geodatabase, and full ArcSDE database. Additional functionality will be provided in ArcGIS 9.1 and is discussed in the ‘Distributed Geodatabase’ topic below.

Potential Uses:

- Performance/Load Balancing
- Network Load Reduction
- Spatial Data Partitioning
- Disconnected Computing



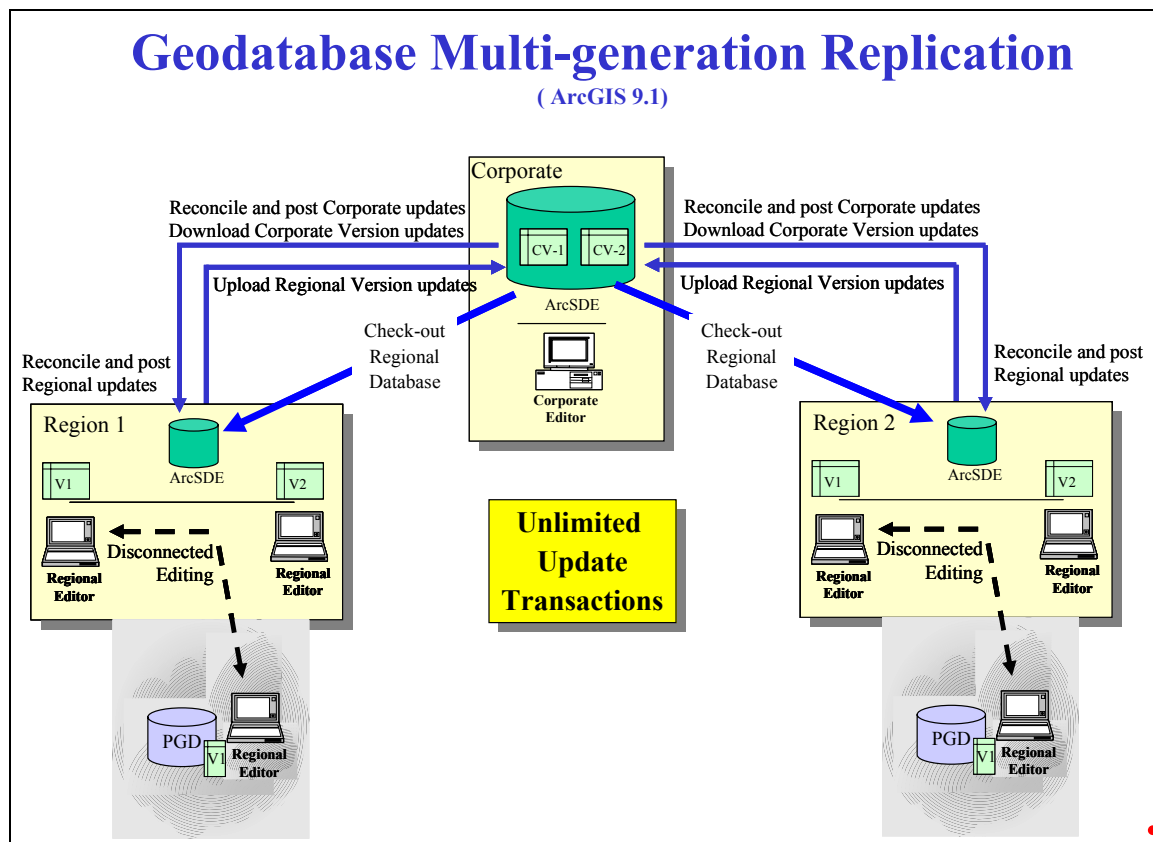


Distributed Geodatabase Replication

This is a future ESRI ArcSDE technology which will be available in ArcSDE 9.1. It will extend the Disconnected Editing functionality to allow multiple version check-out and check-ins and to incrementally synchronize multiple editing databases. Child replicas can either be personal geodatabases or enterprise ArcSDE databases, and child geodatabases can be parents to other replicas. The incremental synchronization from parent to child geodatabases should reduce network bandwidth requirements for maintaining multiple databases in remote locations. Since network bandwidth is always a concern when operating over a WAN, the reduced requirements for this solution make it a more feasible alternative for customers.

The reconcile and post process will be similar to that available with disconnected editing, with improvements to help automate the process and to support edit transaction reconcile and post capabilities. The diagram below shows an example of this technology for an organization with a central office and several regional offices; each with editable databases.

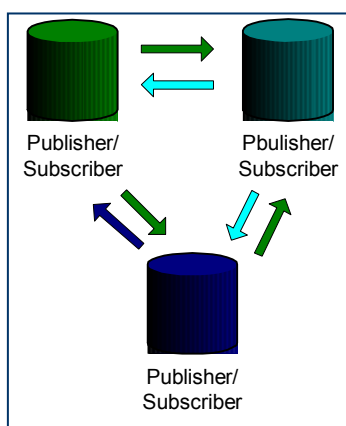
A properly designed distributed geodatabase configuration could also provide a disaster recovery mechanism, assuming that the entire master geodatabase had been replicated to a remote site and maintained periodically.



Potential Uses:

- Performance/Load Balancing
- Network Load Reduction
- Spatial Data Partitioning
- Disconnected Computing
- Disaster Recovery

Multi-Master Replication



This is also known as peer-to-peer replication, merge replication, or broadcast replication. It is the most complex form of replication and is difficult to implement in an ESRI geodatabase setting. The intent for this architecture is to allow concurrent editing on multiple database instances and the synchronizations of those instances through replication processes.

This technique does not generally work in an ArcSDE versioned geodatabase environment with native database and third-party replication products. STATE_IDS are an integral part of the ArcSDE versioning mechanism, and their consistency is absolutely required to ensure a healthy GeoDatabase. As a result, ArcSDE versioning does not support the multi-master/merge replication architecture because the various publishers would not synchronize STATE_IDS accurately. Additional details on

this and other challenges to multi-master replication are discussed in the ArcGIS, ArcSDE Geodatabase Replication challenges section below.

ESRI's disconnected editing function available with ArcGIS 8.3 (and higher), and the future ArcGIS 9.1 distributed geodatabase functions are designed to overcome these obstacles.

Potential Uses:

- Performance/Load Balancing
- Network Load Reduction
- Disconnected Computing
- Disaster Recovery

Replication Misconceptions

Replication is a loaded word, often oversimplified and misunderstood. Below are a list of the common misconceptions about creating a replication service, and replicating data across a GIS enterprise.

- **“It will lighten the network load.”**

Not necessarily true. In fact, there are many scenarios where more traffic will be added to the network. Another option is to keep the centralized database and move the remote user's GIS applications to the same datacenter, hosting the software on application servers such as Windows Terminal Server or Citrix. Thin-clients can access their applications and data over the WAN, thus minimizing network utilization.

- **“I need to replicate everything.”**

Although replication promises ease of use on paper, the notion of replicating an entire database easily must be carefully considered. Analyze your requirements and carefully select, if possible, only those objects required for the task. Without this selection process you may find your DBMS server and network to be overworked trying to synchronize the replicas.

- **“Local data will outperform data centrally.”**

Replication adds overhead to the database server and network which can impact those users accessing the central database. A central database architecture may be a better choice in many cases. As mentioned earlier, a WTS/Citrix solution can be a better method for supporting remote users.

- **“I need data real-time.”**

The allure of replicating data as soon as it changes is soon overshadowed by the realities of actually processing and propagating data. Configuring real or near real time replication is expensive, difficult to configure and maintain, and requires high performing elaborate hardware and infrastructure. This is beyond the budget and viability of many GIS enterprises.

- **“To support remote users I need to replicate their data to them”**

Remote users can be supported in several different ways. One option is to provide them with a local DBMS hosting spatial data and maintain it through replication from a master database (or another replica). As mentioned previously a WTS/Citrix solution can be more effective.

ArcGIS, GeoDatabase, and ArcSDE Replication Challenges

As mentioned earlier, most of the database management systems offer replication services which can work out-of-the-box for many applications. The ArcSDE geodatabase environment, however, adds a layer of intelligence and complexity to the standard database which challenges these replication services. We have mentioned some of these previously, but we list some of them here for users to consider when implementing a replication solution. This is not meant to be a complete list, but it should provide some idea of what organizations face when implementing replication in an ESRI geodatabase setting. Note that the ESRI Disconnected Editing and future 9.1 Distributed Geodatabase functions either deal with or avoid these issues. These challenges pertain to general DBMS, third-party or custom replication solutions when applied to ESRI GIS.

1. Volume of data changes and SQL per spatial edit

Editing an enterprise GeoDatabase in ArcGIS is done through a mechanism called “versioning” in ArcSDE. A single edit in ArcSDE can produce changes to multiple DBMS tables which might need to be replicated to subscribers. Do not assume that a single feature class change will produce a single DBMS change. It could be that one change produces many. This in turn could impact the performance of the replicas depending on the profile of the edits.

2. Configuring GeoDatabase objects to replicate

A GeoDatabase object, like a feature class, often requires several related DBMS objects to operate properly (like feature tables, business table, spatial index tables, triggers, indexes, etc). Configuring and maintaining replication subscriber schemas correctly will challenge the DBA.

3. Synchronizing versions and state_ids

ArcSDE versioning depends heavily on keeping track of and synchronizing the “state” of edits in STATE_IDS in the feature class “delta” tables and ArcSDE system tables. In order for an edit to correctly complete not only does the data need to be committed to the business and delta tables, but the STATE_Ids need to be committed to the system tables as well. A replication scheme in a versioned environment will be more complicated. Consider using the ‘default’ version for replication.

4. Storing x,y values as binary types

ArcSDE typically stores the actual geometric x,y values in a binary array in the DBMS (this is a configurable parameter in some ArcSDE support DBMS). In any case, some DBMS replication services will not replicate particular binary types (e.g. Oracle LONG RAW). However, replication services that support binary types are provided in 3rd party offerings or storage level replication.

5. Cannot currently do multi master/merge updates

STATE_IDs are an integral part of the ArcSDE versioning mechanism, and their consistency is absolutely required to ensure a healthy GeoDatabase. As a result, when using ArcSDE versioning the multi-master/merge replication architecture is not supported because the various publishers do not have a mechanism for synchronizing STATE_IDs accurately. ArcSDE 9.1 will provide facilities for a multi-master database editing and synchronization through the distributed geodatabase functionality.

Multi-master database configurations also require some method of reconciling conflicts when several users in different databases modify the same feature or other database objects. This can be a complex undertaking. As mentioned previously, the disconnected editing in ArcGIS 8.3 (and higher) has methods for resolving these conflicts, and the ArcGIS 9.1 distributed geodatabase architecture will build on that functionality.

6. SDE_logfile issues

Some client side operations require the use of ArcSDE logfiles (which are actually DBMS tables, SDE_logfile_data, SDE_logfiles) like large selections or tracing a geometric network. The logfiles are temporary locations for information about GeoDatabase objects for a specific operation that is likely not important to the subscribers. So, users will likely not need to replicate ArcSDE logfiles, and thus potentially reduce a lot of replication activity.

7. TABLE_LOCKS issues

ArcSDE uses its system table TABLE_LOCKS to during every GeoDatabase edit. The volume of activity on that table can become quite heavy. Also, the entries in that table are specific to an edit on that instance. Thus, users will not likely include this object in replicas.

8. Replicating versioning conflicts

Often subscribers will only be interested in the data comprising the DEFAULT version of a versioned database. However, the replication solution will replicate all versions and version information, including potential conflicts in the child versions. Users may want to consider replicating a versioned database after all versions have been replicated and compressed to minimize the number of conflicts and data.

9. Getting all the DBMS objects that participate in a specific edit/change

A versioned GeoDatabase requires several objects to operate correctly. The ArcSDE configuration guides list the objects participating in a given versioned edit. Careful consideration of the objects to configure for replication will be required. Also, some ArcSDE system tables do not have unique indexes which make them very difficult to synchronize.

10. ACID updates to ensure an operable database

Because an enterprise GeoDatabase edit is comprised of many edits to various DBMS objects (business table, delta tables, system tables), users will need to make sure that a single edit passes the ACID test on subscribers so they don't become out of sync or corrupted.

11. Do you want to see all versions or just DEFAULT?

Native DBMS replication services are not aware of ArcSDE's versioning mechanism, and, as a result, cannot publish just the DEFAULT version or any specific versions.

12. Hot ArcSDE system tables

There are several ArcSDE system tables at the heart over every edit, and thus become a contention point and impact performance. These objects (see configuration documentation) will need to be carefully organized in the physical database configuration to minimize the contention.

13. Compressing can cause a lot of replication activity

The compression operation (to optimize a versioned database) trims the states in a GeoDatabase and tries to consolidate edits in a given object. As a result, there may be high replication activity after a compression is done.

Conclusion

The following is a summary of key recommendations when considering geodatabase replication.

- Carefully consider why you need to replicate early in your planning and whether it is really necessary.
- Know that geodatabase replication in an ESRI GIS environment is possible, especially for view-only applications.
- In general the more simple solutions will be easier to implement and manage over time, will be less expensive and cause the least amount of headaches.
- Consider storage/disk-level mirroring or replication as it avoids many of the problems that one can incur in geodatabase software-level replication.
- ESRI solutions such as disconnected editing and the future distributed geodatabase functionality in 9.1 avoid many replication headaches (especially in editing environments) and will work well if their functionality meets the requirements for the target environment.
- When the options listed above are not feasible for the target environment, consider third-party products which are known players in the GIS marketplace and which have reference ESRI GIS customers to backup their marketing materials.
- Thoroughly test the replication solution prior to committing to it. Many schemes which sound good on paper may, in reality, not provide the functionality required.
- Utilize vendor DBMS documentation, ESRI ArcSDE documentation and the ESRI Support web to support planning and implementation efforts. In particular, check the ESRI Knowledge Base and User Forums, searching on keywords such as 'replication' and 'ArcSDE'.

As in any IT or GIS project consider the costs, benefits and risks in the analysis of the overall feasibility of replication in your environment.

Bibliography

Microsoft SQL Server Replication -

www.microsoft.com/technet/prodtechnol/sql/2000/books/c09ppcsq.mspx.

Oracle Replication – www.oracle.com

Lakeview OmniReplicator - www.omnireplicator.com

Tadpole Cartesia GoSync! - www.tadpolecartesia.com/html/products/gosync/

Miner & Miner Mobile Designer -

<http://www.miner.com/products/index.shtml#Des>.

ESRI Knowledge Base web site -

<http://support.esri.com/index.cfm?fa=knowledgeBase.gateway>.

ESRI User Forum web site - <http://support.esri.com/index.cfm?fa=forums.gateway>.

ESRI ArcSDE Configuration and Tuning Guides – There is one of these for each supported database and they contain information on how to set up replication services.

ESRI Support Project Center – The Project Center centralizes available GIS information and provides a framework for users to plan, design, implement and maintain a successful GIS. Here you will find information ranging from high-level strategies and planning to the latest information on applications, system requirements, programming, testing, and tuning techniques, as well as resources that will help you to create successful GIS solutions.

<http://support.esri.com/index.cfm?fa=projectCenter.gateway>.