An Enterprise Geodatabase: Montgomery County, Maryland
ESRI 2004 User Conference – Paper #1674

1.0 ABSTRACT

Montgomery County, Maryland has a long history of supporting GIS applications with a SDE/ArcSDE based geodatabase. By migrating into the latest ArcSDE release (ArcSDE8.3), the County geodatabase has been advanced a big step and rose from a small scale, localized database into an enterprise level database. The geodatabase powered by the Oracle enterprise RDBMS and coupled with enhanced ArcSDE functionality is serving the GIS users, IT specialists and engineers with a vast amount of vector data, raster data, attribute data, and CAD drawings. In addition to hosting and serving the GIS specific data collection, the geodatabase is also connected to multiple County databases, which enables data sharing and integration from heterogeneous data sources throughout the County. The details of system and architecture design, data organization, and the range of IT applications the geodatabase is supporting will be discussed.

2.0 BACKGROUND

One of the most important functions of the GIS team in the Department of Technology Services (DTS), Montgomery County is to provide database service to all County GIS users. The County GIS data collection encompasses a wide range of spatial and non-spatial data including vector data, raster/imagery data and attribute tables such as Census 2000 publications. The GIS data collection has been growing ever since 1992, the beginning of the County GIS program. It is now a terabyte inventory of data. Data originally required a single UNIX workstation for storage has expanded over the years into a collection that occupied multiple UNIX and Windows workstations and mini-servers.

In 1996, the GIS team began the construction of a centralized geodatabase using ESRI's SDE technology. SDE 3.0 was installed and configured on a Sun E4000 server equipped with two 167MHz processors. The SDE has been subsequently upgraded into ArcSDE8.x in the following years, and the database management system was Oracle v7.3.2. The deployment of the first County geodatabase was successful; it hosted the property parcel data layer and all of the planimetric data layers – roads, building footprints, hydro features, transportation features, etc... Early ArcView users enjoyed the operation of simple connection and one-stop access of most of the spatial data layers.

Due to the fact that most of the GIS work was performed in a conventional Arc/INFO workstation environment before year 2000, and the fact that County GIS analysts mainly used spatial data in the form of Arc/INFO coverages, there was not a lot of effort devoted to enhancing the geodatabase. With the release of ArcGIS desktop from ESRI in recent years, more and more GIS analysts and end users switched from Arc/INFO workstations and ArcView 3.x into the use of ArcGIS, which led to the increased interest and demand for enhancing and upgrading the functionality of the early geodatabase.

The year 2003 marked the turning point for the County geodatabase. By gaining the ownership of a more powerful UNIX server (Sun E450 with two 400MHz processors) with increased storage capacity, GIS team was ready to develop an enterprise-level geodatabase for the County that would host the entire primary GIS data collections. Time and labor were invested to make the new geodatabase a one-stop data service for all GIS analysts, web application developers and other end users. Taking advantage of increased functionality of ArcSDE 8.3, the new geodatabase has been redesigned with a more logical structure. Data layers are organized into categories by natural relations. The new geodatabase is also able to serve a total of 45 gigabytes of color ortho-images as a single raster dataset. By building 13 layers of image pyramids, the entire county image can be displayed on users' screens almost instantaneously. The geodatabase was formally put into production at the end of 2003, and it is now the data server for all GIS users and internet/intranet applications.

3.0 SYSTEM ARCHITECTURE, DATABASE DESIGN AND CREATION

3.1 Architecture Layout

In 2003, the Montgomery County GIS data collection was consolidated into three major data servers. The latest addition (GIS-DB3) is the host of the County geodatabase. User workstations are connected to the data servers through the local Ethernet network at 100Mbps, which rides (at 1 giga-bit per second) on the County fiber-optic network backbone, as shown in Figure 3-1. The County's GIS program and use of technology have evolved over the years following ESRI's strategic development pattern. This is reflected in the County's GIS computing environment layout, which is shown in Figure 3-2.

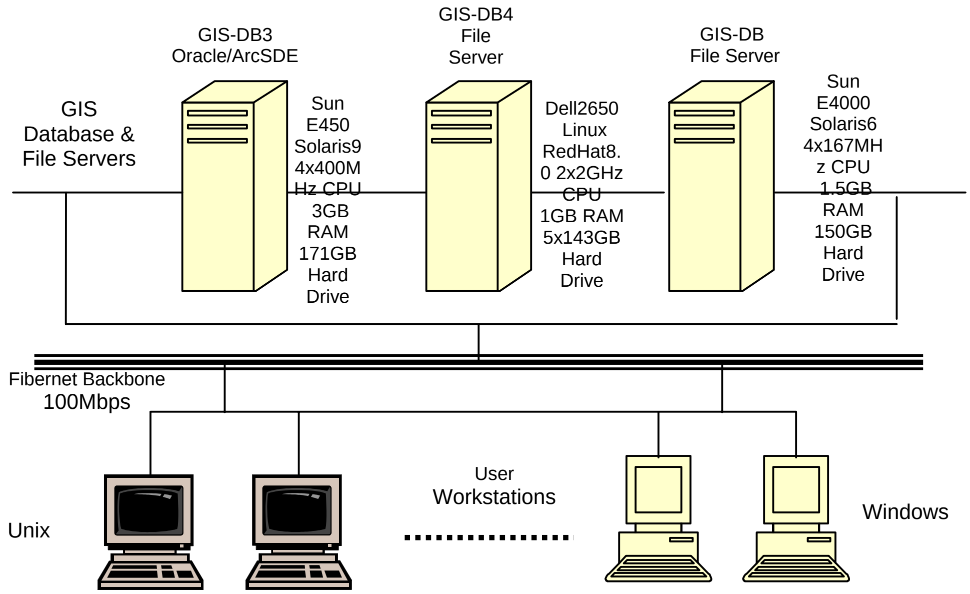

Figure 3-1. GIS Data Server System Architecture

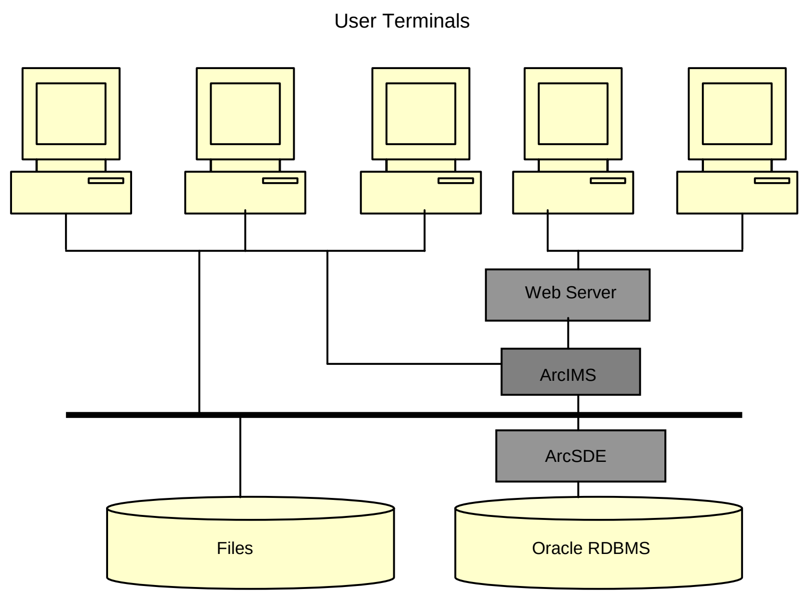

Figure 3-2. GIS Computing Environment

3.2 Database Software Installation, Configuration and Tuning

Although ArcSDE 8.3 is compatible with both Oracle v8.1.7 and Oracle v9.0.4, the decision was made to install Oracle v8.1.7 because of its proven stability. Following the successful installation of the Oracle database management system, ArcSDE8.3 was installed. A considerable amount of time was spent configuring and tuning both the Oracle instance and the ArcSDE service because this is the key to the success of a well-performed database. Following both Oracle and ESRI configuration guidelines and utilizing the server's three gigabyte memory space, the newly created geodatabase was successfully configured and tuned and hence optimized the database performance.

3.3 Database Design and Data Organization

Montgomery County's GIS data collection comprises different data types, including vector, raster, CAD drawings and attribute tables. Good database design and data organization are essential to facilitating user orientation and simplifying accessibility of the data. Taking advantage of the new geo-dataset concepts in ArcGIS and ArcSDE, vector data layers are organized into different datasets, each containing naturally associated data layers. For example, the DISTRICTS dataset contains all service/administrative area polygons, such as election districts, school service areas, fire boxes, police districts etc…; the LOCATIONS dataset contains all location points, such as points (or places) of interest, hospital locations, school locations etc... Figure 3-3 captures a snapshot of the content of the datasets NATURE&REC (nature and recreation); it features bikeways, parks, scenic routes and trails.

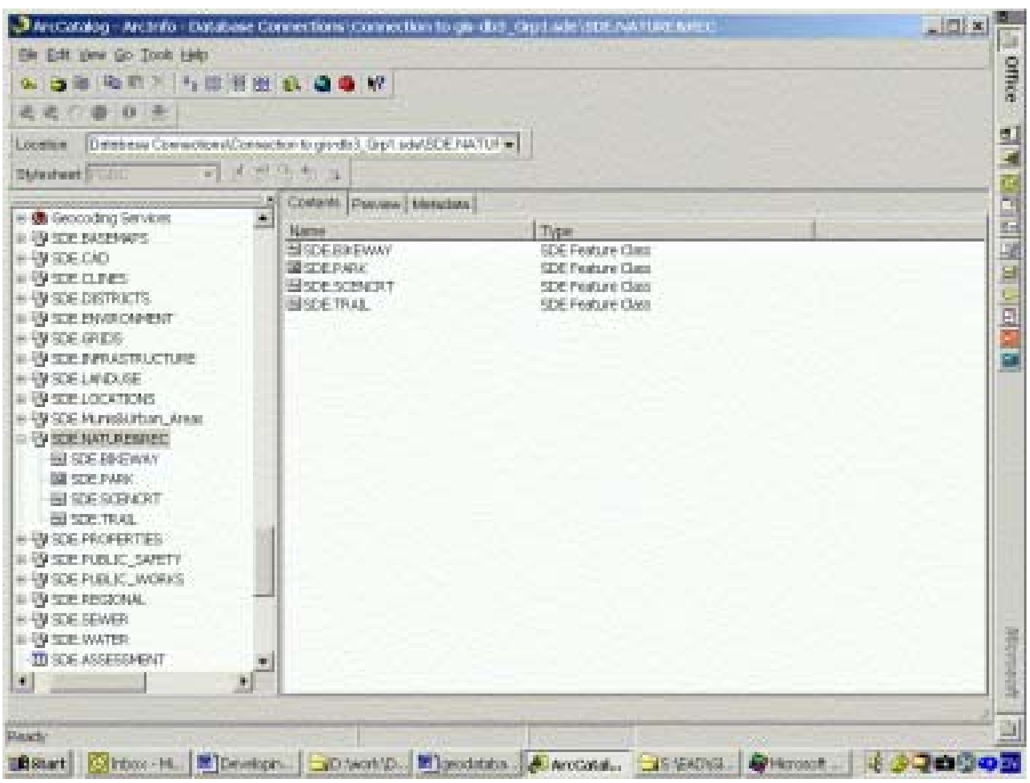

Figure 3-3 Content of Dataset Nature &Rec

CAD drawings are stored in a separate category; each planimetric data layer is stored as a feature class in the CAD dataset. The 2002 color ortho-image of the County is stored in a single raster dataset ORTHO02.

3.4 Database Completion and Deployment

Once the database was designed, the geodata was loaded into the centralized geodatabase from different data servers and disk drives. Vector data was loaded using the ArcGIS catalog manager ArcCatalog; raster data was loaded using ArcSDE command SDERASTER. Uploading the entire county raster images of 2672 TIFF files into a single ArcSDE raster dataset represented quite a challenge. The prerequisite for merging multiple image files into a single raster data map in ArcSDE is that each image file has to have the exact pixel registration as its neighboring image files, which means that pixels from multiple image files line up exactly, and the corners of those image files meet at the exact x, y coordinates. A batch job was written to clean up the TIFF files to ensure that all adjacent image files had the same coordinates. Another batch job was executed to merge all TIFF files into a single mosaic while the TIFF files were loaded into ArcSDE one by one. Eight layers of image pyramids (each has reduced resolution for the purpose of fast rendering of images) were built for the raster dataset. Finally, the dataset was analyzed using Oracle optimization tool. Test results revealed that the speed for raster image retrieval and rendering greatly outperformed the previous method of serving images through the use of individual TIFF files. Staff investment in this effort proved to be very worthwhile.

Since both Oracle and ArcSDE communicate to their client software through TCP/IP protocol, deployment of the geodatabase to all users and departments is much easier than the old way of mounting file systems. Most end users were able to enter the database server information into their desktops by following a set of simple instructions. Few end users required phone services or office visits. The deployment was quick and smooth, yielding a big savings on the GIS system administrator’s time.

4.0 APPLICATION SUPPORTED

4.1 Back-End of Web Based Applications

Currently, 17 feature datasets have been created in the ArcSDE geodatabase and 112 feature classes are served to County users. The new geodatabase is now the host of all GIS intranet/internet applications. The array of applications backed by the database includes County GIS web site, Snow Removal application, Permitting Center GIS component, and recently implemented Water Quality Protection Charge (WQPC) system. The GIS web site supplies all primary GIS data layers stored the geodatabase to County workers through intranet, and supplies about two dozens of GIS data layers to public through internet. Snow removal application is an internet based application that provides snow removing status to citizens during snow storms, and at the same time allows citizens to report road conditions and file complains. Both Permitting Center and WQPC are intranet applications that are used by County workers to provide customer services to citizens.

4.2 Support Heads-Up Digitizing

With the fast delivery of ortho-image of aerial photos, the ArcSDE geodatabase is becoming an image server for many to perform heads-up digitizing work. Users have discovered the flexibility of zoom in, zoom out, and pan either the entire County aerial photos or specific areas in a matter of seconds, and decided it best facilitates the heads up digitizing of ground features. GIS users are utilizing this feature to create new or maintain the existing data layers, such as storm water facilities, building roof lines, parking lots and etc.

4.3 Support Server-Size Geocoding

The Department of Technology Services (DTS) at the County has been maintaining a good quality street centerline file (CLINE) sine early 1980's. A couple of server-side geocoding services of different geocoding styles have been created against the CLINE file, and stored in ArcSDE geodatabase. A collaborative effort from multiple GIS analysts was also made in creating a master alias file that incorporates all possible spellings of street names and street types. This master alias file is also stored together with the server-side geocoding services. The result is a centralized geocoding server for all GIS analysts and end-users.

4.4 Support Wide-Area Data Access and Integration to Other Databases

With the infrastructure of wide-area network (WAN) already in place at the County, enabling database access through WAN becomes simple. By opening the Oracle and ArcSDE TCP ports to privileged user locations, the public schools, cities and some other local government agencies have gained access to Montgomery County DTS geodatabase. It is greatly enhancing the means of data sharing and data exchange between neighboring government agencies.

Since the County geodatabase is powered by Oracle RDBMS, enabling data access between other Oracle and SQL Server databases becomes straight forward by utilizing open database connection (ODBC). One good example is the Solid Waste Customer Support application. Solid Waste maintains their recycle and refuse route geodata in geodatabase, and they also maintain another customer support database on a separate Oracle server. The customer support database stores information such as recycle/refuse pick-up status, service calls, contract company information and so on. By establishing an ODBC connection between ArcSDE geodatabase and Solid Waste customer support database, GIS users are able to link the dynamic information from the customer support database to the recycle and refuse route service areas. This allows Solid Waste GIS workers to build statistics, identify potential problems, and anticipate the needs to redesign the route.

5.0 CONCLUSION

The geodatabase is gaining popularity not only to the GIS community but also to IT specialists, engineers, desktop operators and web applications. Because the data integration capability between the centralized geodatabase and other departmental databases, more and more departments are ready to build new IT applications with embedded or integrated GIS components. For example, the Department of Public Works is conducting a feasibility study of Cityworks, an asset management software from Azceta Systems, for the purpose of managing and handling department assets, work orders and other public service businesses. Cityworks is a typical application that requires the data support from both geodatabase and other business databases.

The Montgomery County geodatabase has truly grown from a small scale, localized geodatabase, that served primarily GIS users, into an enterprise level geodatabase. It has established a solid foundation for an enterprise GIS. It is now an integral part of the IT infrastructure at the County.

Lian Chen, Ph.D.
Sr. IT Specialist
Montgomery County, Maryland
Lian.chen@montgomerycountymd.gov