Supporting Multi-Vendor Applications

Using an Open Spatial Database

Since 2002 GITA in the US has conducted an annual survey of the North American utility industry. One of the questions that is asked is whether the organization uses more than one GIS. As you can see, most of the organizations reported using more than one GIS and sharing data between them. The survey also asked what tools were used to share spatial data, and Safe Software's Feature Manipulation Engine (FME) was named most often as the third-party application for data sharing. The reality for many organizations is that sharing data between different vendors' applications has involved redundant data: multiple copies in the different formats supported by the different GIS vendors.

By Geoff Zeiss

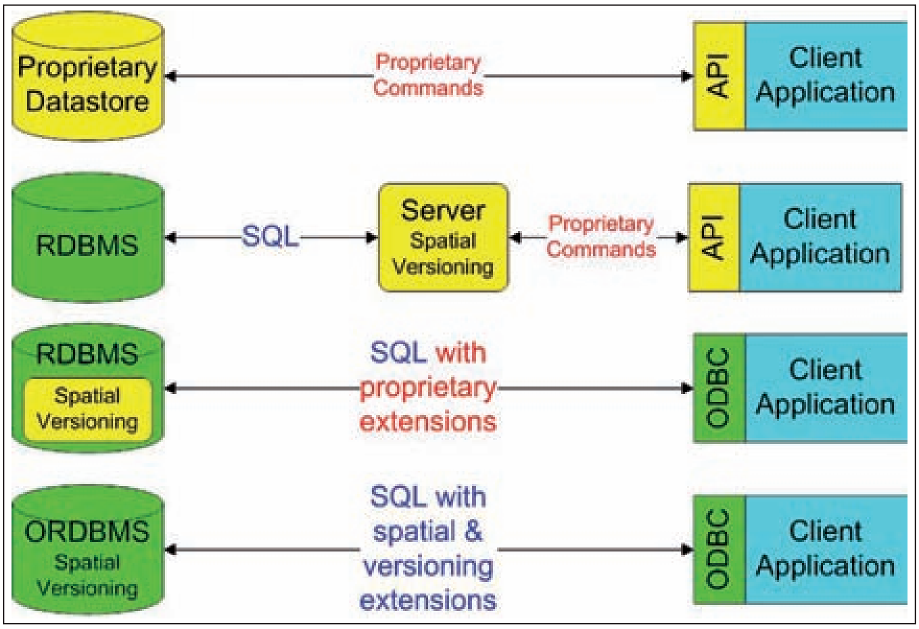

Geospatial Architectures.

Single Point of Truth

Many people have been looking at open spatial databases as a way of replacing multiple files with a single point of truth. The promise is becoming a reality. At GITA last year, two US municipalities reported how they had implemented multi-vendor interoperability based on an open spatially-enabled relational database management system (RDBMS).

Throughout the world, utilities and telecommunications firms manage infrastructure in basically the same way and are facing similar challenges. If you look at the information flow in these organizations, the most obvious thing that strikes you is the problem of silos or islands of information. The second thing is that the information flow in these organizations is for the most part based on paper. For example, the Engineering group uses CAD, the Records (sometimes called Network Documentation) group uses GIS, and the flow of information between these two groups is paper. The result is redundant processes and backlogs. In addition, the aging of the work force exacerbates what is already a critical problem because there is no effective mechanism for transferring the knowledge in the heads of experienced workers to the facilities database where it can be accessed by younger, less experienced workers. Three components are critical to addressing this problem: a single point of truth implemented as a centralized, spatially-enabled RDBMS; a way to exchange electronic as-builts between Engineering and Records; and what I call field force enfranchisement.

Liberating

Widespread support of spatially-enabled relational database management systems by GIS vendors has enabled customers to begin to adopt geospatially-enabled RDBMSs as a shared enterprise datastore. The business benefits of choosing an open geospatially-enabled RDBMS as the enterprise single point of truth is that applications from different vendors can share spatial information in a common repository. This approach to spatial data management liberates customers from vendor lock-in, enabling them to buy best-of-breed applications.

Managing spatial data has evolved over the years from proprietary files, through proprietary schemas for storing spatial data in an RDBMS, to the current state where you can store just about any spatial data, including topology, in a modern object-relational database management system (ORDBMS). The advantage of an ORDBMS is that spatial data is accessible through an open query language, SQL, and an open interface standard such as ODBC, JDBC, and OLEDB.

All-relational

In the area of network infrastructure management, beginning in the early 90s solution vendors such as GeoVision pioneered the use of relational database management systems for storing geospatial data. Such systems were deployed by early adopters, typically large utility and telecommunications firms and municipal governments, to manage their network infrastructure. These solutions were marketed as all-relational to distinguish them from traditional GIS applications which used relational technology for feature properties but invariably stored geospatial data in proprietary files external to the RDBMS. All-relational systems were remarkably successful and are still deployed at major utility and telecommunications firms and municipalities around the world. However, one of the disadvantages of these systems is that the data model or schema used to store geospatial data was specific to the solution vendor and required reverse engineering to enable the sharing of geospatial data with other vendors' products.

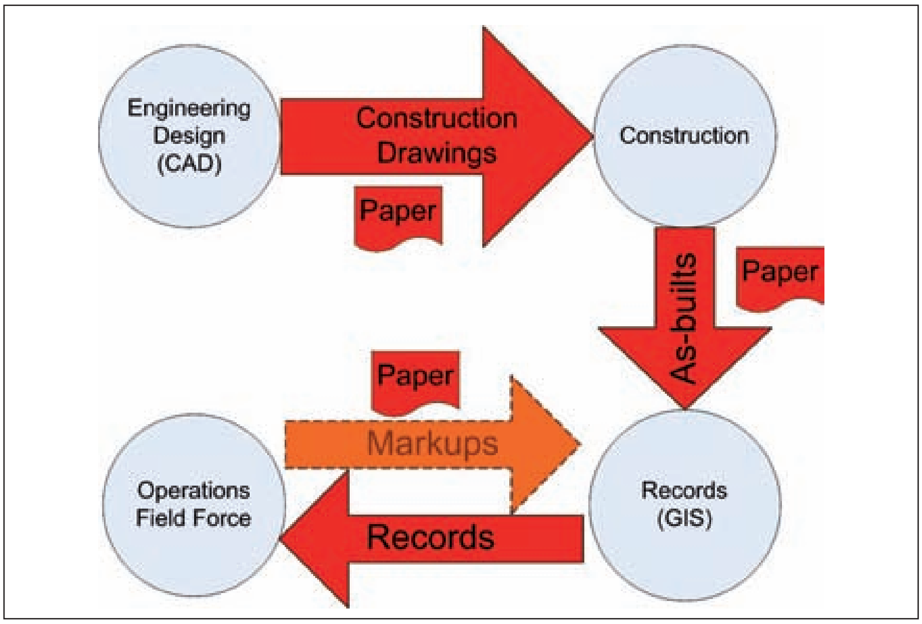

Infrastructure Management Lifecycle Paper.

Object-relational

Michael Stonebraker and pioneering RDBMS vendors such as Illustra introduced what was referred to at the time as object-relational database management systems. ORDBMS differentiated themselves from traditional relational database management systems in their ability to store complicated data structures in each cell of a table. These data types included geospatial data, time series, and other structures that were difficult to store in a simple two-dimensional table. This advance in RDBMS technology meant that now the storage of geospatial data and the SQL for manipulating geospatial data were defined by the RDBMS vendor, not by the geospatial application vendor. The important implication was that two GIS vendors supporting the same RDBMS could share data and avoid the data redundancy associated with file import/export.

Developing a Standard

A major breakthrough in sharing spatial data was the development of the Open Geospatial Consortium's (OGC) standard Simple Feature Specification (SFS) for SQL which is incorporated by most RDBMS systems, including both closed and open source. Although the SFS only supports simple features, points, lines, and closed polygons, it has become widely adopted and specific implementations by commercial RDBMS vendors are now supported by most geospatial vendors.

However, there are some wrinkles. Sharing basic geospatial data types covered by the SFS standards such as points, lines, and closed polygons is supported by most geospatial vendors. But things get more complicated when you want to share text, symbolization, layer definitions, topology, and long transactions. The OGC is making progress in addressing these issues. For example, an extension to the SFS to support text was approved recently.

No Paper-based Infrastructure

At GITA 2006 two municipalities, Tacoma and San Jose, gave presentations that described operational systems that shared spatial data between applications from multiple vendors, based on a spatially-enabled RDBMS. As I remember the vendors involved included ESRI, Autodesk, Intergraph, MapInfo, and possibly others. The most important implications of what these municipalities have done are 1) organizations are recognizing the business value of having a single point of truth for spatial data and 2) it is technically feasible to implement a central spatially-enabled RDBMS. To put it simply, technology is no longer the excuse for maintaining a slow and expensive paper-based infrastructure management system.

In both municipalities, the major challenge involved managing text, metadata, and stylization, and each municipality addressed this in somewhat different ways. At Tacoma RDBMS tools such as triggers were used to maintain the metadata required for each vendor's application. At San Jose a third party text management tool was used to share text. The important conclusion is that by either using built-in RDBMS tools or third-party add-ons, it is feasible to address the problems associated with managing metadata, stylization, and text.

Final Challenge

Spatially-enabled RDBMS technology and support by geospatial application vendors has advanced to the point where spatial data can be shared in a secure, highly available environment between applications from different geospatial vendors. This provides a tremendous business benefit to utility and telecommunications firms as well as municipal government organizations in managing their network infrastructure because these organizations can now build solutions by integrating best-of breed applications and solutions from multiple vendors. At the present time this permits telecommunications and utility firms to address the paper-based flow of information between Engineering, Records, and Operations, thereby improving customer responsiveness and reducing costs. There are even signs that the final challenge, improving the flow of information between construction and Engineering, may also be beginning to be addressed.

Geoff Zeiss (geoff.zeiss@autodesk.com) is Director of Technology at Autodesk Inc.

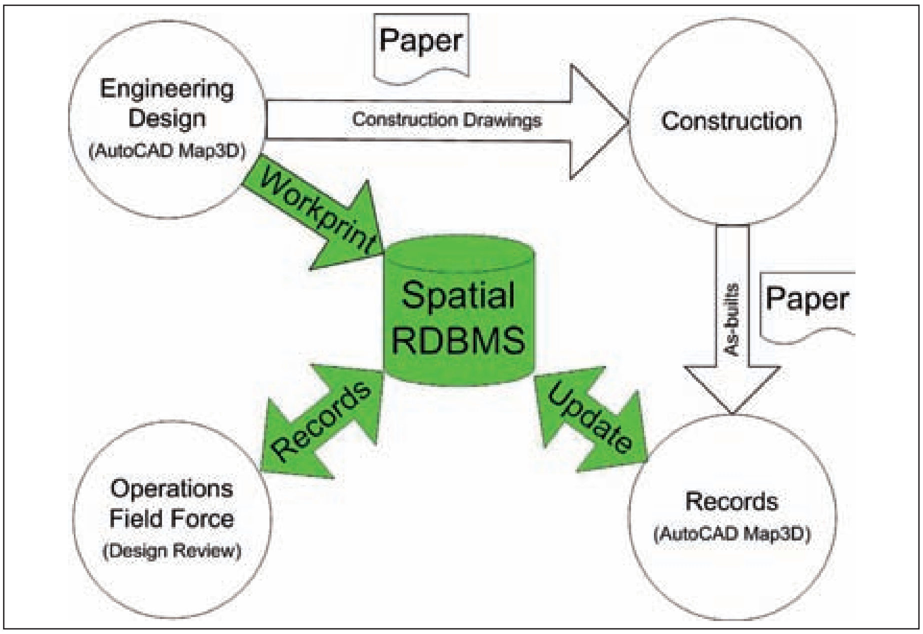

Infrastructure Management Lifecycle New.