



Enterprise Geographic Information Servers: A New Information System Architecture

An ESRI® White Paper • October 2003

Copyright © 2003 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, the ESRI globe logo, www.esri.com, and @esri.com 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.

Enterprise Geographic Information Servers: A New Information System Architecture


An ESRI White Paper

Contents	Page
Introduction.....	1
GIS Subsystems	1
Enterprise Geographic Information Servers	4
Conclusion	6
References.....	6

Enterprise Geographic Information Servers: A New Information System Architecture

Introduction Throughout the course of the last 40 years, geographic information system (GIS) technology has been implemented using a variety of contemporary computer architectures (Coleman 1999). The earliest systems were built to run on mainframe and minicomputers, then came workstations, next PCs, and now system designers are recommending Web-based technologies for distributed systems (Peng and Tsou 2003). A new architecture is currently emerging based on some new technologies that allow all GIS functions to run in a centralized server environment and be accessed from any device on a network. This white paper will describe this new enterprise geographic information server approach and its implications for building truly distributed enterprise GIS.

GIS Subsystems GIS comprises three major subsystems that are concerned with mapping, spatial analysis (geoprocessing), and data management. A useful, general purpose GIS needs capabilities in all these areas.



The Three Subsystems That Are Needed for a Well-Rounded GIS Platform

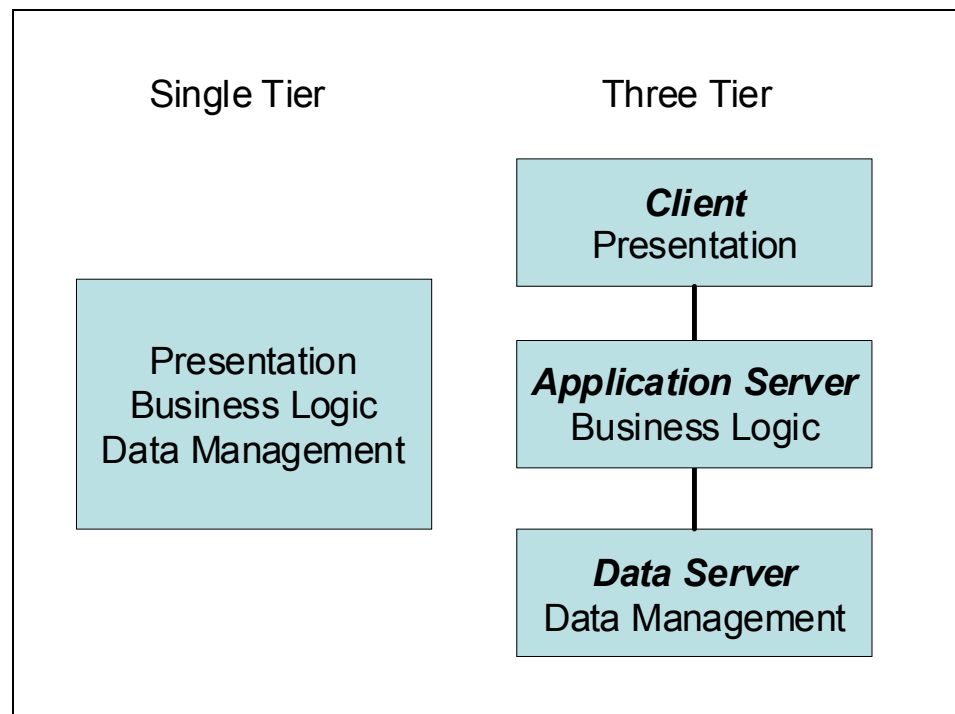
The mapping subsystem deals with the manipulation and visualization of geographic information. It is here that users will find projection and datum transformations, map to page transformations, symbology models, map displays, two-dimensional and three-dimensional visualization tools, editing components, and the framework for interacting with geographic data in the form of a map.

The spatial analysis subsystem implements a series of geographic analysis functions for the two classic geographic operations of proximity and overlay analysis as well as functions for data conversion (import computer-aided design [CAD] data, export graphics, etc.), grid analysis (watershed modeling, intervisibility analysis, etc.), and three-dimensional analysis (calculation of slope, aspect, etc.).

In the past few years it has become common practice to use commercial off-the-shelf database management system (DBMS) products, such as Microsoft® Access, IBM® DB2®, Oracle®, and Microsoft SQL Server™, to store and manage geographic information (Longley et al. 2001). DBMS products have been widely used in GIS because they allow users to create a single, centralized data repository (avoiding redundancy and duplication); they facilitate data sharing (by establishing de facto standards); they support multiuser editing of continuous geographic databases; and they allow users to employ DBMS backup and recovery tools. However, the limited support for advanced geographic data types (Shekar and Chawla 2003); the weaknesses of SQL as a geographic data access and programming language (Egenhofer 1982); and concerns about scalability for operations, such as topology management (Hoel et al. 2003), have

instigated a reexamination of the role of DBMS in GIS and have prompted the development of enterprise geographic information servers.


The three GIS subsystems can be implemented on a single computer (single-tier architecture), two computers (two-tier architecture), or in the case of large systems, they are usually spread across three separate machines (three-tier architecture). Presently, almost without exception, all the major GIS software systems have been developed to operate on desktop class personal computers running the Microsoft Windows operating system. In a single-tier architecture the presentation (user interface), business logic (geographic processing functions), and data management functions all run on the same machine. At the other extreme, in a three-tier situation the presentation tier (including the mapping and spatial analysis subsystems) runs as a client on a desktop PC, the business logic for data access runs on a middle-tier application server, and the data management tier (DBMS) runs on a server. In this type of configuration the client has sophisticated capabilities and, in advanced systems, can comprise several hundreds of megabytes of code (a so-called "thick" client implementation). This configuration has served GIS users well for the last decade and will continue to do so for advanced GIS tasks such as editing, cartographic compilation, three-dimensional modeling, and spatial analysis. However, in recent years users have become dissatisfied with this type of architecture for less demanding tasks because the cost of client machines and software is prohibitive, it is expensive to maintain and update widely distributed PCs—especially if users cannot do it themselves, and it is difficult to maintain corporate standards throughout a widely dispersed organization.



Single- and Three-Tier Information System Architectures

Enterprise Geographic Information Servers

An alternate, more centralized architecture is emerging that addresses these issues. At the heart of this new architecture approach is the idea that the application server tier can be enhanced to encompass not only the data management GIS subsystem but also the mapping and spatial analysis subsystems. Running all the GIS components on a server means that a thin client can be used to initiate processing requests and display the results of GIS tasks. Such centralized GIS implementations are now being referred to as enterprise geographic information servers. They allow both thin (browser) and thick (desktop GIS) distributed clients to access data and processing capabilities over a standard local or wide area network. The server itself has two parts. The server object manager is responsible for interfacing with the network (e.g., a Web server) and for allocating requests to server containers that perform the actual work. Both parts can run on the same CPU, be spread over multiple machines (for scalability), or run multiple CPUs on the same machine. The server containers obtain the data required to fulfill a task from a data server (a DBMS).



Enterprise Geographic Information Server Architecture

Enterprise geographic information servers have a number of characteristics that make them ideal for serving the needs of distributed GIS users.

- Low cost of maintenance. Because all data and processing capabilities are centralized in one location, costs of upgrade and maintenance are minimized (e.g., software updates do not need to be pushed out to hundreds of PCs). GIS servers can be installed alongside an organization's other servers, and the same information technology (IT) staff can be used to maintain both.

- Easy scalability. Adding additional resources can easily be achieved by installing new memory or disks in an existing server or by installing a new server, without taking the system down. Moreover, the resources required are not expensive state-of-the-art hardware components but widely available, lower cost, general purpose machines.
- Standards-based access to data and processing. A key factor that will encourage the uptake of enterprise geographic information servers is that they have open standards-based application programming interfaces (APIs) that expose all the data and processing capabilities for use by developers. This means that users of all types can access data and processing across a network from any location. The key standards that are supported include Extensible Markup Language/Simple Object Access Protocol (XML/SOAP) for messaging and data transfer over Web (http) network connections, Java™ for building local and wide area network enterprise applications, and .NET for building local and wide area network enterprise applications. Such standards offer a rich environment for communicating with an advanced application server and for building true enterprise applications.
- Support for all key GIS data types and functions. In order to address the widest range of users and applications, enterprise geographic information servers support all common GIS data types and functions. Supported data types include two-dimensional and three-dimensional vectors, rasters, annotation, attributes, addresses, CAD, survey measurements and computations, dimensions, networks, and topologies. The list of functions includes a full range of mapping, spatial analysis, and data management tools such as topographic and thematic mapping, editing, geocoding, network analysis, data management, three-dimensional analysis, raster (spatial) analysis, and geoprocessing.

Enterprise geographic information servers are well suited to simple and medium complexity GIS tasks that can be undertaken on a remote server, accessed over a network. Typical tasks include creating a topographic map centered on a user-specified location, at a user-specified scale; simple editing (e.g., heads-up, on-screen land parcel creation and boundary update, or changing the attributes of streetlights and road signage); and determining the shortest route for a vehicle between multiple home delivery stops. Complex applications that require high interactive performance and extensive local processing power are not well suited to wide area network deployment and will continue to be implemented on dedicated local computer resources such as a desktop PC.

Enterprise geographic information servers are good examples of the new class of information systems called enterprise application servers. Similar general purpose systems have been developed by major vendors such as BEA (WebLogic), IBM (WebSphere), Microsoft (Windows® Server), and Oracle (Oracle Application Server). They all rely on and enhance the capabilities of DBMS by providing enterprise application development tools, network services, Web service interfaces, and scalability frameworks. In addition to GIS, a number of other application areas have also adopted this server implementation approach. Enterprise resource planning (Oracle, PeopleSoft®, and SAP™) and customer relationship management (Microsoft, Pivotal®, Siebel®) systems are both examples of server-centric applications.

Conclusion Enterprise geographic information servers are a new type of GIS architecture that supports access to potentially thousands of distributed users from a centralized location. The industry standards-based approach to data management, application development, and data and processing access makes them ideal for providing GIS services to a wide range of distributed users. Similarly, the low cost of maintenance, high scalability, and excellent reliability means that organizations can create compelling business cases that demonstrate a good return on investment. In the next few years many new GIS user organizations will build systems based on this new server-centric architecture, and many existing user organizations will transition their system to it.

References Coleman, D.J., 1999, "Geographical Information Systems in Networked Environments," in Longley, P.A., M.F. Goodchild, D.J. Maguire, and D.W. Rhind, *Geographical Information Systems: Principles, Techniques, Applications, and Management*. New York: John Wiley & Sons, 317–29.

Egenhofer, M.J., 1992, "Why not SQL!" *International Journal of Geographical Information Systems*, Vol. 6, 71–85.

Hoel, E., S. Menon, and S. Morehouse, 2003, "Building a Robust Relational Implementation of Topology," *Proceeding of 8th International Symposium on Spatial and Temporal Databases*, 16 pp.

Longley, P.A., M.F. Goodchild, D.J. Maguire, and D.W. Rhind, 2001, *Geographic Information Systems and Science*. New York: John Wiley & Sons.

Peng, Zhong–Ren and Ming–Hsiang Tsou, 2003, *Internet GIS: Distributed Geographic Information Services for the Internet and Wireless Networks*. Hoboken, New Jersey: John Wiley & Sons.

Shekar, Shashi and Sanjay Chawla, 2003, *Spatial Databases: A Tour*. Upper Saddle River: Prentice Hall.

Acknowledgement This white paper was created using the article, "Enterprise Geographic Information Servers," written by Dr. David Maguire for the *GIS Development* magazine, Volume 7, Issue 8, August 2003.