



Interoperability in Enterprise GIS

An ESRI® White Paper • April 2005

Copyright © 2005 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 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.

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, ArcGIS, 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.

Interoperability in Enterprise GIS

An ESRI White Paper

| Contents | Page |
|--|------|
| Introduction..... | 1 |
| Background and Issues | 2 |
| Strategies for Achieving Interoperability..... | 5 |
| Loosely Coupled, Message-Based Services | 5 |
| Common Data Models | 6 |
| Supporting Legacy Applications..... | 8 |
| Conclusion | 8 |
| References and Further Reading..... | 9 |

Interoperability in Enterprise GIS

Introduction

Interoperability can be defined as the ability to exchange and use information across a large, heterogeneous network made up of several local information system nodes. The purpose of this white paper is to summarize ESRI's strategy for supporting an interoperable enterprise geographic information system (GIS) suitable for large mapping organizations, given current technology, practices, and trends in the software industry.

From previous generations of geospatial "stovepipes," which were monolithic applications that handled all aspects of the user's requirements, commercial GIS has evolved into component-based frameworks operating within federated information systems integrated with commercial relational database management systems (RDBMS). Multiple tiers of applications are now developed by various organizations using different languages, data store platforms, and user interaction environments. These application tiers generally perform various subsets of tasks more efficiently and with greater maintainability than before. And just as software technology has emerged to support federated systems, it has also created the need to share data across these multiple independent application systems as well as with legacy systems still in use.

Interoperability is the key for enabling disparate users utilizing different commercial or custom GIS software to access and make productive and efficient use of shared geospatial data and other information. However, interoperability covers multiple distinct areas of capability when it comes to geospatial data. Some degree of interoperability is enabled when the GIS industry as a whole agrees on standard data types, formats, and Web services such as the Open Geospatial Consortium, Inc. (OGC), Simple Features Specifications; Geography Markup Language (GML); Web Map Services (WMS); Web Feature Services (WFS); and metadata catalog services. ESRI® ArcGIS® software directly supports these formats and services with as much ease and transparency to the user as possible.

Efficiency and high performance in enterprise GIS are difficult goals to achieve, but it becomes especially difficult when designing a system for interoperability. This has mostly to do with the overhead of making function calls across process boundaries to take advantage of multiple vendors' software capabilities. The most efficient software application is one that can run in a single process, because then all function calls can be made in the same address space, and network latency is not an issue. By dividing an application into components that run in separate processes, overall performance can drop by several orders of magnitude, even if all the processes are on the same computer. Remote calls over a network have even more to do: "Data may have to [be] marshaled, security may need to be checked, packets may need to be routed through switches. If the two processes are running on machines on opposite sides of the globe, the speed of light may be a factor." (Fowler 2003, p. 388.) Given the inevitable high overhead, it is

essential to keep interaction between interoperable software components to the minimum, which means using coarse-grained interfaces instead of fine-grained interfaces, thereby passing larger sets of data with fewer Web service exchanges.

Another level of interoperability can be achieved when a specific GIS user community defines and agrees on a common data model that captures the essential geospatial features and other concepts used in its application domain such as for land parcels, addressing, hydrography, or homeland security. This requires more direct and active involvement by the user to manage his/her application schemas and to take steps to capture and share the data in such a way that others in the same community of interest can make the most productive use of it. ESRI cannot dictate these data models but has been very active in facilitating and supporting many different user communities to develop and share their own domain specific data models. More information on these is available at <http://support.esri.com/datamodels>. The remainder of this paper presents ESRI's technical strategy for supporting interoperability from within its applications.

Background and Issues

Large organizations have considerable investment in their enterprise information systems, which increasingly involve geospatial data. This is especially true of large mapping agencies. Because of the large number of consumers and multiple software applications, some of these organizations have long struggled to develop organization-wide standards for data formats, data models, application program interfaces (API), and data exchange mechanisms. Not only is the pace of technology development difficult to keep up with, but the competing needs between differing user communities result in much of the standards development efforts falling short from the beginning. In the past, defense and other federal agencies in the United States developed custom systems often called government off-the-shelf (GOTS) software but have increasingly turned toward commercial off-the-shelf (COTS) and standards-based COTS (SCOTS) solutions. The focus of this paper is on ESRI's approach to SCOTS for supporting such organizations. Here are some of the issues we have to consider.

- Federal agencies today need to support numerous different resource sponsors, some of whom require that the enterprise information systems support their particular legacy user applications. These might be government- or academia-developed applications, each with its own data store, processing, and presentation components that do not interoperate uniformly with information from other systems. However, each of these meets certain analytical needs of a group of consumers, and so much time and funding has been invested in these systems that it is very difficult for their managers to justify switching to newer technology and approaches.
- Geospatial data is increasingly being integrated with more general business data in corporate and government information systems. Geospatial data requires considerable storage, processing, and data throughput. This coupled with the processing and data throughput load of SQL RDBMS stretches most systems to their capability limits.
- Enterprise data tends to be centralized within an organization, but some enterprise data remains in departmental- and agency-level databases (islands of isolation) because of any number of factors such as replication for more efficient access, desire

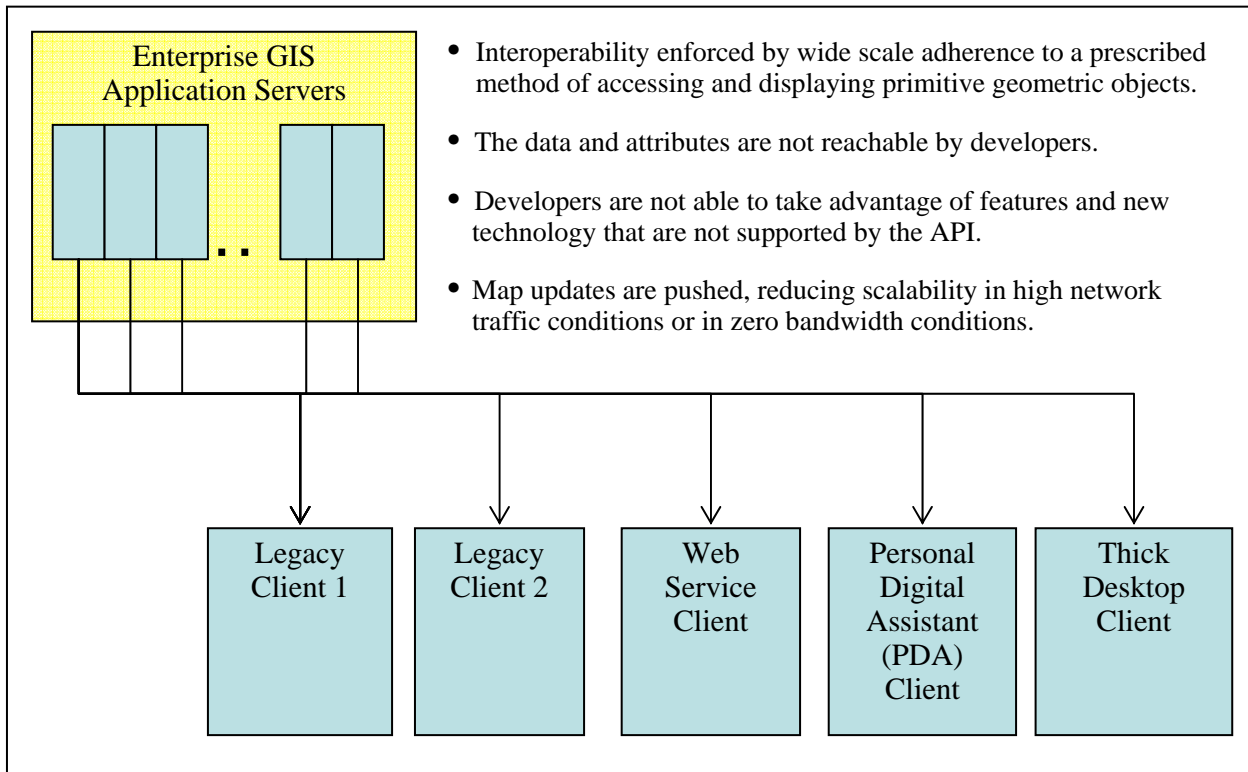
for local autonomy, differences in desired attribution of the data, and so on. Efforts to fully centralize enterprise data are essentially futile and can even lead to reduced system efficiency.

- In many enterprise systems, there is a need to perform updates on parts of a centralized database while disconnected from the host or network.
- Geospatial application vendors, as in other application domains, compete on the basis of their capabilities and expertise in specific problem areas. No single vendor could or should be expected to have the best solution for all applications. Therefore, it is expected that vendors will find ways to share information of general use because it benefits vendors and customers alike. In today's atmosphere of increased security awareness and emergency preparedness, this has become a mandate, not an option.
- Many federal agencies and large corporations are so decentralized that different divisions have developed wholly different ways of managing and using their enterprise data. It is increasingly important to support interoperable access to this enterprise information. The context of national defense, homeland security, and other emergency response adds urgency to achieving interoperability.

In the past, many government-developed systems focused on standardizing their fine-grained APIs for data management, processing, and display capabilities. A driving force behind this focus was the notion that by prescribing an API, independent developers could and would write interchangeable components for subsystems of the application. In practice, however, software development required considerable investment of time and learning for new programmers to undertake and complete. Furthermore, it has proven difficult to "mesh gears" between mature GIS software and these prescriptive APIs; it is often expedient or even necessary to write new applications from scratch, another considerable investment. But even then, these applications would suffer from the lack of flexibility to best apply the vendor's skill set to solving the real problems in data store access, geoprocessing, and map display. This would further hinder the ability and likelihood that these government applications could evolve at a fast enough pace to keep up with improvements in technology and its customers' demands.

This brings up what is perhaps the most onerous issue for supporting legacy GIS—that the central information system is expected to push the actual geometric primitives to its various consumers' client applications, adhering to the low-level geometry API defined for each type of client (see Figure 1 below). This poses a considerable burden on the central information servers to keep abreast of the distinct and inevitable evolution of each client application. It also hinders much of the potential analytical use of the geospatial data being served, because it is mandated to serve only symbolized geometry, ostensibly to support thin and ultrathin clients. This is a serious shortcoming to desktop and other clients who can take advantage of more aspects of the geospatial data than just the geometry. Much of the geospatial information in an enterprise GIS has rich and meaningful attribution that can be essential to business-related analysis. Another drawback of this approach is that data consumers have no alternative way of accessing the data when they are disconnected from the network.

Figure 1
Interoperability Through API



In response to the numerous issues and requirements for information exchange, a myriad of standards development efforts has been underway for more than 10 years.

- The Digital Geographic Information Working Group has been concerned with geographic information exchange (DIGEST, for Digital Geographic Information Exchange Standard) specific to the defense and intelligence communities in the North Atlantic Treaty Organization (NATO), while the International Hydrographic Organization (IHO) has been concerned with similar standards for maritime navigation (S-57). ArcGIS software products and those of other GIS vendors support these formats.
- A collaboration of federal agencies and academia developed still another such standard for civilian use (Spatial Data Transfer Standard [SDTS]). While this data exchange format is in very limited use in the United States, it is supported in current ESRI products.
- The International Organization for Standardization (ISO) and OGC are concerned with interoperability standards for the geospatial industry (ISO 191xx, OGC Simple Features, GML, WMS, WFS, catalog services, and so on). Many of these standards

are in widespread use among GIS software vendors and are supported in the current ArcGIS software products.

- The World Wide Web Consortium, Organization for the Advancement of Structured Information Standards, Internet Engineering Task Force, and other software industry standards bodies have developed an extensive set of technology standards related to the Extensible Markup Language (XML) and Simple Object Access Protocol (SOAP). These are in widespread use across the software industry and are in a continual state of rapid evolution. A useful summary of software standards organizations and programs may be found at <http://www.service-architecture.com/Web-services/articles/organizations.html>. ESRI keeps abreast of these standards developments and incorporates those that become pertinent to its own software development projects as they mature.

Clearly, these are a lot of activities and developments with which to keep up. The software industry standards bodies have a proven track record for achieving productive collaboration and agile development among their members for the general benefit of all, regardless of competitive pressures that may exist in the marketplace. This collaboration and agility among software and geospatial industry vendors should be compelling reasons for defense and other federal agencies to take advantage of this development; this is the means by which our national defense can best maintain its superiority in information and intelligence management.

Strategies for Achieving Interoperability

Loosely Coupled, Message-Based Services

Interoperability was defined previously as the ability to exchange and use information across a large, heterogeneous network made up of several local information system nodes. The broadest understanding within the software industry for the best way of achieving interoperability is through coarse-grained, loosely coupled protocols that can operate among a federation of heterogeneous information systems using message-based services. The focus on interoperability should be on protocols for message-based GIS services (e.g., Web service protocols such as WMS, WFS, and XML/SOAP). This is how ArcGIS has been architected and will continue to follow the evolution of these standards. This approach addresses the issues presented above in the following ways:

- The use of coarse-grained, stateless, message-based protocols reduces the coupling between system components, resulting in less communication traffic taking place between application service providers and consumers. This generally improves overall system performance.
- Coarse-grained protocols have received extensive attention in the software industry standards arena, resulting in greatly reduced barriers for new developers to learn the systems, apply common software industry development tools, and begin to add value in various ways. Technology such as ArcGIS Server further aids the developer by managing the creation of Web service providers as needed.

- A message-based approach allows the distributed architecture to reduce dependency on direct connection between server and client. In situations with limited or unpredictable network connections, this is essential for a robust system design.
- Message-based approaches are inherently more scalable as network traffic increases. With a fine-grained API approach, scalability is typically constrained because of the need for synchronous client/server connection. But with loosely coupled, coarse-grained messaging, a wide range of scalability can be supported without changing the application architecture simply by adding hardware and layers of application servers.

To make the distinction between fine-grained and coarse-grained APIs clearer, consider a simple example of exchanging street address data between a server and client. A fine-grained interface would have separate accessing methods for each component of an address such as name, street, city, state, country, and postal code. However, communicating at this level between separate applications for each individual address record would be painfully slow in comparison with the use of a coarse-grained interface. In the coarse-grained approach, each application may have a detailed domain model of addresses but could send addresses between applications in terms of a tabular structure, such as a record set, encoded in an XML document. Each record set might contain a complete list of addresses including all associated attribution. There would be orders of magnitude less interprocess and network traffic generated using an XML record set in a single message request, in place of a series of fine-grained API calls.

Common Data Models

Another aspect of interoperability has to do with the data models being shared among the various collaborating applications. This consideration has driven the development of a host of data format standards such as the Vector Product Format (VPF) family, National Imagery Transmission Format (NITF), IHO's S-57, SDTS, and numerous others. Each of these serves a number of requirements, but these are very complex formats requiring large investments by software companies to support and maintain. Furthermore, they are better suited to batch transmission of whole data sets rather than for incremental updates within interactive sessions. None of these are candidates for connections between enterprise GIS systems and their consumers.

There are essentially two kinds of data models that need to be considered: a low-level data model that defines a basic unit of data exchange and a higher level data model that captures the essential concepts of geospatial features, business objects, and their relationships in the user's application domain. For example, VPF defines a low-level organization structure for representing geometric primitives, related attributes, cross-references, metadata, and so on. In addition, each of the VPF product specifications, such as Vector Map Level 0 (VMap 0) or the Digital Nautical Chart (DNC), defined thematic sets of feature classes, each feature's specific attribution, and scales at which the features would be captured in a given coverage. The core VPF specification could be thought of as the physical data model or data format, and each of the VPF products, such as VMap 0 or the DNC, could be thought of as a logical data model. The logical organization of coverages and features in a given VPF product could be implemented any number of ways; the core VPF physical model was just one such way. SDTS follows a similar approach in that it defines a core physical model, and each of its profiles defines

an extension of that model. However, SDTS plus a given profile still only defines a physical model such as for point, vector, or raster data. There are no community-wide logical models based on SDTS.

Within the geospatial software industry, a number of commercial data formats have emerged as ad hoc standards in the sense that they are very widely used. The ESRI shapefile specification is one of these as well as Intergraph DGN, AutoCad DXF, and MapInfo MID or MIF formats. While vendor specific, these are nevertheless freely published formats. Their relative simplicity in comparison with VPF, NITF, or SDTS is their greatest advantage, as this makes it easy for third party developers to read and write to these formats. A prominent third party developer, Safe Software, has even staked out a market niche by developing the Feature Manipulation Engine, a sophisticated tool for translating between the hundreds of standard and commercial data formats.

However, it is understandable that defense and other federal agencies may not want to favor any one vendor over the others by standardizing on a vendor specific format. For this reason, the OGC has pursued development of GML shortly after XML appeared on the scene in the late 1990s. GML version 3 appears to support most, if not all, typical enterprise GIS requirements for data exchange and can be used for both incremental updates and whole data sets. Because of its roots in XML, GML is versatile and can be adapted to a wide range of application conditions and environments. Further work is needed, however, to address GML's complexity and large file sizes. At a minimum, a simple profile of GML is needed to support true interoperability. The current work in progress to this end is called GML Level 0 Profile for WFS. This will most closely harmonize with the previous OGC specification for the Simple Features geometry model and greatly improve the interoperability of this specification for many applications.

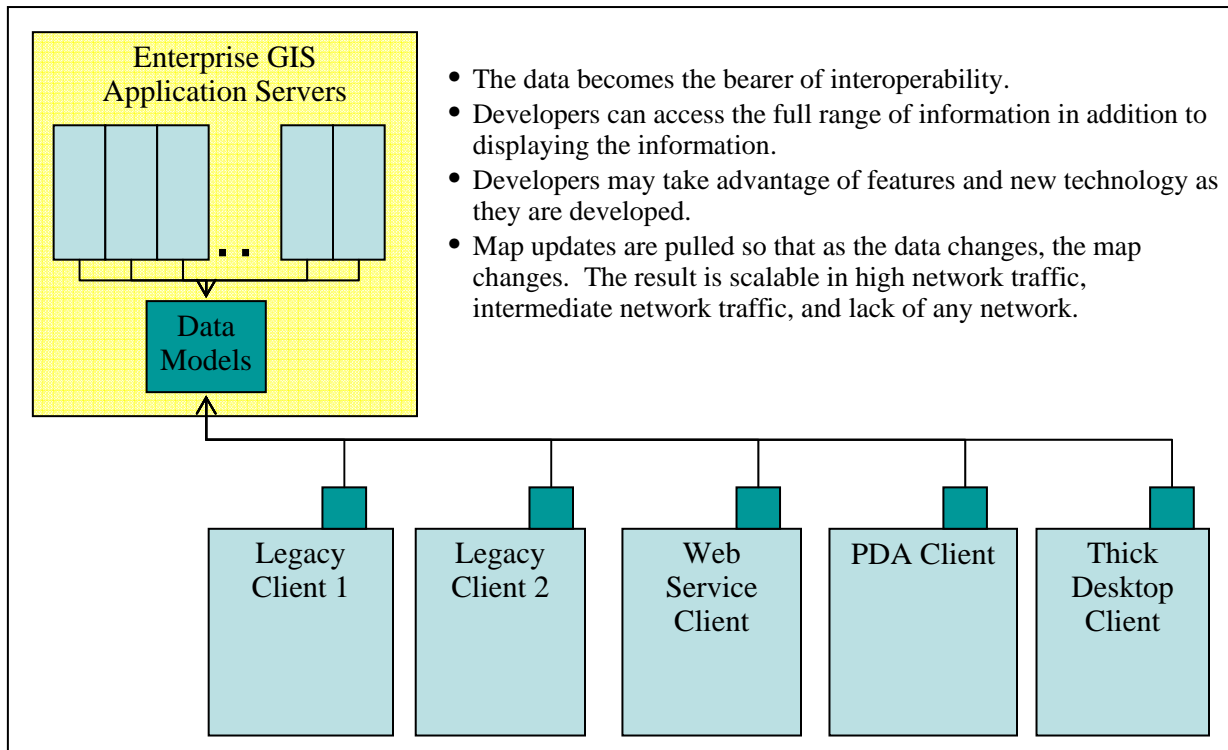
The real key to designing a physical data format for greatest interoperability is to keep the data format as simple as possible. This means keeping the definition of features and their geometry to the bare minimum, for example, not including topological primitives within the data set. This immensely reduces file size and complexity and makes the data accessible to the widest range of sophistication across GIS vendors and other types of applications. The reason the ESRI shapefile specification is so widely used is not just because of all the installations of ESRI software but because it is such a simple format that it is easy to create and edit.

A higher level of conceptual and/or logical data model standardization may be useful to achieve, as this would reduce the need to translate data from one application logical data model to another. However, many federal agencies have found it difficult to achieve consensus on higher level conceptual and logical data models because of the time-consuming and challenging task of reconciling concepts and terms across their constituent communities. This need not and should not impede an organization's progress in starting at the physical data model level. The most important task is simply to share whatever data has been captured between interoperable components, and this can be achieved with existing technology and practices, as just described.

Supporting Legacy Applications

For interoperability of geospatial information with legacy applications, we recommend the following approach. First, pick a standard, simple physical data format such as GML 3 Level 0. Then, instead of pushing updates to the clients, make the clients ask for (pull) the updates from the enterprise GIS when needed, using a coarse-grained, loosely coupled API. This way, the enterprise GIS can make its complete, rich set of feature data and attribution available to those applications that can use them. Other applications that need a prescribed format of geometry can be "wrapped" with interface agents that pull the data needed and express the data in the format needed. Figure 2 below illustrates this approach. The key is to design the API to be *permissive* rather than *prescriptive*.

**Figure 2
Interoperability Through Data Models**



Another benefit of this approach is improved support for disconnected applications. When a client is going to be disconnected for a period of time, a subset of the essential database could first be copied to the client's hard disk. The application wrapper for each client can handle getting the data needed from the local data store. Depending on application requirements, changes to the data made by the client could then be merged back to the main database upon reconnection.

Conclusion

Interoperability is an essential design goal for any enterprise GIS system, but this will not be achieved through standardization of fine-grained APIs. The experience of GIS vendors and data publishers within the OGC, not to mention the larger software industry, has confirmed this. The first family of OGC technology specifications, collectively

referred to as Simple Features, did not deliver directly in terms of interoperability because the Simple Features APIs for Component Object Model (COM) and the Common Object Request Broker Architecture (CORBA) were not widely implemented (only one vendor has implemented all these, which does not constitute interoperability). However, the specification of feature geometry and coordinate systems that emerged from the Simple Features specifications has become central to the new generation of Web-based specifications.

The bearer of interoperability must be the data model. The data model must be rich enough and adaptive enough to contain all that is needed by an enterprise GIS software's anticipated clients. It must also be simple enough structurally that it is easily and widely supported by multiple vendors' software. Coarse-grained APIs working with this data model will provide the best performance and most efficient communication and processing of geospatial information among various application clients.

The approach laid out in this paper achieves the following important goals:

- Enterprise GIS systems need interoperability solutions that are scalable across a range of bandwidths from zero to very high connectivity.
- Some GIS users need to be able to work with legacy systems for the foreseeable future, but this should not prevent others from taking advantage of current and future capabilities of information technology as they appear.
- Any API defined for an enterprise GIS should be *permissive*, rather than *prescriptive*, in support of differing data requirements for its various client applications. GIS consumers should not be limited to getting only what they can use in their legacy applications today but should be able to pull additional information from the central data store as their applications evolve to handle it (such as commercial GIS can today).

References and Further Reading

- Fowler, Martin, *Patterns of Enterprise Application Architecture*, Addison-Wesley, 2003.
- Schulte, Roy, *The Future of Application Integration*, Gartner Report K1, APN13, November 2004.
- Erl, Thomas, *Service-Oriented Architecture: A Field Guide to Integrating XML and Web Services*, Prentice Hall PTR, 2004.
- Spatial Data Standards and GIS Interoperability*, ESRI White Paper, January 2003.
- ArcGIS: Engineered for Interoperability*, ESRI White Paper, August 2004.