Contents Page

Presented at AURISA 98 - The 26th Annual Conference of AURISA Perth, Western Australia, 23-27 November 1998

The Importance of Metadata Engines in Spatial Data Infrastructures

***Andrew Phillips*[1],**
***Ian Williamson*[2] and**
***Chukwudozie Ezigbalike*[3]**

[1] *Graduate Student*
Department of Geomatics
The University of Melbourne
Parkville, Victoria, 3052
Phone: 03 93444509
Email: honey@sunrise.sli.unimelb.edu

[2] *Professor of Surveying and Land Information*
Department of Geomatics
The University of Melbourne
Parkville, Victoria, 3052
Phone: 03 93444431
Email: i.williamson@eng.unimelb.edu.au

[3] *Lecturer*
Department of Geomatics
The University of Melbourne
Parkville, Victoria, 3052
Email: d.ezigbalike@eng.unimelb.edu.au

ABSTRACT

With the rapid increase in the development of spatial data collection tools, such as GPS and remote sensing technologies, the amount of spatial data being collected and stored on computer networks is becoming vast. Knowing what spatial data is already available for an area is of great benefit to many spatial data users as duplication of spatial data collection and processing is a time consuming and costly exercise. Descriptive data about the data sets are maintained as metadata to provide such information to spatial data users.

Current Spatial Data Infrastructure (SDI) technology allows search engines to search metadata systems within a network to discover what spatial data is in existence. As SDIs evolve, it is expected that spatial data sets will be accessed "live" from their distributed locations, rather than being downloaded before being used. With this development, it is anticipated that metadata systems will evolve into metadata engines and will again be background tools that are used for querying spatial data sets that are distributed across a network.

This paper will describe the terms metadata and metadata engines in relation to SDIs. The paper will also discuss the trend towards distributed processing of spatial data and the evolving SDIs. Details will be given about the processes and problems of developing a prototype, using a public-domain package, ISITE, as a starting point, that will allow for simple distributed processing of spatial data to occur.

KEYWORDS: Spatial data infrastructures, metadata, and metadata engines.

INTRODUCTION

The rapid advancement in spatial data capture technologies such as the Global Positioning System (GPS), satellite imaging and total stations have all made the capture of digital spatial data a relatively quick and easy process. This has meant that in the last decade or so the amount of digital spatial data in existence has increased significantly. There is now a vast amount of spatial data, stored by numerous organisations at various locations across the globe. Much of these data are not being used as effectively as they should. Referring to the LANDSAT images, the Vice President of the USA said:

> *In spite of the great need for the information, the vast majority of those images have never fired a single neuron in a single human brain. Instead, they are stored in electronic silos of data. (Gore 1998)*

Recently there has been a greater focus on how to use the spatial data that are collected and stored in the expansive "electronic silos" to their full potential. Spatial Data Infrastructures (SDIs) are a key component in allowing for the better utilisation of spatial data throughout the spatial data industry. Current SDI technology allows search engines to search metadata systems within a network to discover what spatial data is in existence. It is expected that as SDIs evolve spatial data sets will be accessed "live" from their distributed locations, rather than being downloaded before being used. With this development, it is anticipated that metadata systems will evolve into metadata engines and will again be background tools that are used for querying spatial data sets that are distributed across a network.

This paper firstly describes the terms SDI, metadata and metadata engines along with the importance of metadata engines in relation to SDIs. The paper then discusses the trend towards distributed processing of spatial data through metadata, metadata engines and the evolving SDIs. Finally details are given about the development of a prototype that allows for simple distributed processing of spatial data to occur.

SPATIAL DATA INFRASTRUCTURES

Software packages that are capable of integrating spatial and non-spatial data to yield the spatial information that is used in decision making are commonly referred to as Geographic Information Systems (GIS). They are computer-based equipment, procedures and techniques for manipulating spatial or map data. They are used mostly on a project basis to perform particular analysis. GIS, when used in such a way, acquires its data by digitising or scanning the relevant maps and/or by using field techniques that accumulate the data in digital form. At this level, GIS is used as a tool.

In most cases data that are collected for a particular project are useful for other unrelated projects. With the recent "commoditisation" of data and information this fact has become even more pertinent. The costs involved with data collection are taken into account in project planning, along with attempting to maximise the use of the data from a project. Most people in the GIS industry would realise that some data required for particular decisions are transient and may no longer be available to collect when required. Decisions concerning agricultural practices are a good example. Often these decisions will require data on the environment spanning over several years. These data must be collected when they are available, even if the need for them are not present at the time of collection, as they cannot be collected at the time that they are needed. Therefore there is a need for this type of data to be collected and placed in databases and made accessible for others to use. Databases of this type become shared resources, which are maintained continuously. A database, which has been created, maintained and exploited using a GIS- the tool- is itself often referred to as a GIS. Thus at this level, GIS is used as a resource.

Collaboration and cooperation of several disciplines and a proper strategic plan are usually required for the maintenance of these GIS resources. Coordinating authorities are needed to coordinate the state or national level government organisations, and sometimes private corporations that have been assigned custodianship responsibilities and use privileges for subsets of the data. General community users are then able to expect the data to be available, and with network technology, to be accessible transparently. At this level, the GIS have now acquired the status of an infrastructure- the spatial data infrastructure (SDI).

SDIs are just like other forms of better-known infrastructure, such as roads, powerlines and railway. The whole concept of infrastructure, SDIs included, is that they allow participating members of the community to use them. They are available and taken for granted, even if we have to pay for the right to use them, for example through vehicle registration, railway tickets, etc. The general users essentially do not care how they work or who makes them work, as long as they work.

SDIs comprise the fundamental datasets (spatial data resource) as well as the interrelationships between these datasets, the management of them, and the means of access to, and distribution of, those data (Scanlon 1998). The FGDC(1996) defined a SDI as an "umbrella of policies, standards, and procedures under which organisations and technologies interact to foster more efficient use, management, and production of geospatial data." It further explained that it "consists of organisations and individuals who generate or use geospatial data, of the technologies that facilitate use and transfer of geospatial data, and of the actual data." It should at no stage be assumed that SDIs are all about networks and technology (FGDC 1996; Masser 1998). A SDI will not function, no matter how good the networking and technology, if communication channels, standards and procedures, partnerships and data have not been developed.

METADATA

Metadata is commonly defined as "data about data" (ANZLIC 1996; Kildow 1996; ANZLIC 1997). There are two different forms of metadata. The first, and oldest, form of metadata occurs within GIS, CAD packages and databases where it is "[an] underlying set of rules which tells a software program how to handle data" (Wilson 1998). Database management systems have for a long time used metadata to describe the internal layout of the data schemes within them (Codd 1990; Korth and Silberschatz 1991). This description of the internal schemes is then used by the database management system to construct the results of user queries. By using the metadata the system knows where to find the results for the query.

The second form of metadata is a recent development. Metadata have become products in their own rights, especially in the spatial data management field, where they are used to describe the characteristics of datasets. Characteristics like the custodian, description of the data, geographic extents of the data, currency of the data, storage format, data quality, contact information to inquire about the dataset are all described. In this context metadata is extremely important for spatial data as it allows a potential user of a dataset to determine whether the dataset is useful to them or not. Metadata systems can be established that allow users to search the metadata records for the datasets located on a network. From the results they are able to determine if there are any datasets that may be of interest to them, how to gain access to them, any constraints on using them, etc. Such an application is often referred to as a spatial data directory or in some cases a clearinghouse.

Metadata of this type are extremely important as they facilitate the more efficient use of spatial data. This is achieved by allowing potential users of spatial data to search for datasets that may suit their needs. They can look at the metadata record for a dataset and see if it meets the criteria for use that they have set. The record will also tell the searcher the access rights/constraints of the dataset. All this is very important as it is usually cheaper to purchase that spatial dataset from another party that has produced it for another purpose than it is to reproduce the dataset oneself. The last thing organisations want to do is to duplicate work that has already been completed by another party.

METADATA ENGINES

A metadata engine is an application that is used by database management systems (DBMS) to extract and display the results to a user's query. They work by parsing the query and then consulting the data dictionary for the database, which contains metadata that outlines the internal structure of the database. By comparing the parsed query with the metadata in the data directory the user's query is able to be resolved (Korth and Silberschatz 1991). The metadata engine works completely in the background with no direct interaction with the user of the database. The user of the database does not even have to know that the engine exists. All the user is concerned with is writing the query and then getting the right results returned. Obviously the type of metadata that are being referred to here are the first type mentioned previously.

Metadata engines should not be confused with metadata systems. A metadata system is very similar to a search engine. They allow a user to search metadata records, which have been produced to describe the characteristics of a dataset, and determine whether they wish to gain access to the dataset. Data directories and clearinghouses both use metadata systems to allow users to search them. They both contain databases that hold the individual metadata records for each dataset that is available. These databases are searched by keyword, geographical location, date, etc. and return the individual metadata records to the user for them to view. By viewing these metadata records the user is able to determine whether the dataset is of use to them, whether it meets their accuracy requirements, any access constraints, who to contact to gain access to the dataset, etc. In the case of a clearinghouse there is also the capability to download the dataset online. However there is no capability to query the dataset online, whereas a metadata engine has this capability. An example of a metadata system is that of GI connections in Victoria, located at http://www.giconnections.vic.gov.au .

At the present time there appears to be no true metadata engines in existence that allow the distributed processing of spatial data over the WWW. Distributed processing is the term used when a distributed database is set up that allows the querying of autonomous databases that are located over a network. To the user of a system that allowed distributed processing of spatial data it would appear as if they were just using one integrated database. It should be transparent to them that the data that returns after they submit a query is actually returning from possibly two or more autonomous databases.

THE IMPORTANCE OF METADATA ENGINES IN SDIs

The importance of metadata engines is probably best emphasised when one looks at how decision makers use spatial data to make their decisions. Currently decisions requiring spatial data are generally made with the assistance of a stand-alone GIS (Glover 1997). This GIS does of course have a limited amount of data on a limited number of topics stored. Thus the decision that one makes while using this GIS is obviously biased by that data. The best decisions are made when as much information as possible is taken into consideration. To get as much information as possible into the decision making process the decision maker could keep adding more and more data to their closed GIS, or they could undertake a distributed approach where they essential have an open GIS. This approach takes advantage of other relevant datasets that are available over the network and consults them to gain the results that the user wants (Ordnance Survey of Great Britian 1996). This approach is the better one as

1. The physical storage in the users own GIS is minimised;
2. The data that they are using in the remote databases is more likely to be up to date as it is likely to be being drawn from the data custodians own database, or at least a mirror database; and
3. It is likely to be much more economical to access small pieces of the required datasets remotely when it is necessary, rather than purchasing the whole dataset and the updates.

When implemented the approach of distributed processing allows the creation of what could be called a "virtual database" (Glover 1997). The users of a "virtual database" would be able to use it as if they were accessing a single database that is located on their own machine. They do not need to know that the "virtual database" that they are using is in fact made up of any number of distinct databases located at any number of different locations. They databases could also be under the control of differing custodians from both the private and public sectors as well as being stored in differing proprietary database systems. All this is hidden from the user.

A good example of where a "virtual database" could be implemented is in any state or national government throughout the world. Each of the departments that exist within the government are likely to be custodians for one or more spatial datasets and are hence responsible for collecting, maintaining and distributing the data from those datasets. This is part of their core business. It is also part of their core business to use that data to make decisions. It is likely that other datasets maintained by other government departments are used in conjunction with their own datasets to make these decisions. It is not part of their core business to duplicate the collection of these datasets and whole dataset transfers between the departments is costly.

If a "virtual database" were established between the government departments it would be possible for a user to locate a parcel boundary on their own GIS application and then, via a network, reference required themes, such as base mapping, surface geology, environmental constraints, flood hazards, land ownership, planning zones and transportation routes. These themes may be owned and maintained by any number of other government departments and agencies, however the user is unaware of this and has the capability to analyse them spatially without the need for file transfer and data duplication (Glover 1997). With the advances in communication technologies meaning faster networks this style of application becomes a reality.

One of the key concepts that is required in order to make distributed processing work is that of metadata engines. It will be the metadata engine that will undertake the metadata and query management that will enable the user to query several remote autonomous databases simultaneously. Metadata engines will act in the same fashion in a distributed processing system for spatial data as they do in stand-alone GIS/databases. Just as they use metadata in stand-alone GIS/databases to tell them where the relative components are within the database to satisfy the users query, they will be able to determine the internal architecture of distributed system though the use of metadata. If the metadata engine works as intended its existence should be transparent to the user and they should be able to use the system as if they were only accessing a single database.

Another example of where distributed processing is likely to be used in the future was recently outlined by Al Gore(1998) when discussing his vision for a "Digital Earth". He said:

> "Imagine, for example, a young child going to a Digital Earth exhibit at a local museum. After donning a head-mounted display, she sees the Earth as it appears from space. Using a data glove, she zooms in, using higher and higher resolutions, to see continents, the regions, countries, cities, and finally individual houses, trees, and other natural and man made objects. Having found an area of the planet she is interested in exploring, she takes the equivalent of a "magic carpet ride" through a 3-D visualisation of the terrain. Of course, terrain is only one of the many kinds of data with which she can interact. Using the systems' voice recognition capabilities, she is able to request information on land cover, distribution of plant and animal species, real-time weather, roads, political boundaries, and population. She can also visualise the environmental information that she and other students all over the world have collected as part of the GLOBE project. This information can be seamlessly fused with the digital map or terrain data. She can get more information on many of the objects she sees by using her data glove to click on a hyperlink. To prepare for her family's vacation to Yellowstone National Park, for example, she plans the perfect hike to see the geysers, bison, and bighorn sheep that she has just read about. In fact, she can follow the trail visually from start to finish before she

> ever leaves the museum in her home town."

While this scenario may seemed far fetched to some, most of the technologies and capabilities that would be required to build a Digital Earth are either being developed, or indeed already here. Essentially the ideas outlined by Gore are just an extension of the principles of distributed processing of spatial data. Once distributed processing systems for spatial data are in circulation, which is not that far away, extensions to the concept like virtual reality interfaces will make the Digital Earth idea of Gore possible. This of course assumes that the network and computer speeds to complete the required data processing are sufficient. The rapidly increasing speed of development of both of these technologies suggests this should not be a problem.

One might say "Why does the Digital Earth have to be developed using a distributed processing approach? Why can't all the data be integrated into one enormous data warehouse?" True all the data could be integrated into one data warehouse, computer speeds and storage are progressing rapidly enough to make this possible. The real problem with the data warehouse approach is that of data custodianships. The datasets that will be involved in the Digital Earth project will be owned by many different organisations in government and the private sector from all around the world. These organisations will wish to remain in control of their datasets. If all the datasets are merged into one data warehouse the individual organisations lose control of their datasets. It is fair to assume that most countries involved in such a project would be very keen to maintain control over the datasets that are able to be accessed. By undertaking a distributed approach, each of the custodians retain control of their datasets as they are accessed at their site via a network.

DISTRIBUTED PROCESSING PROTOTYPE

The Concept

As part of the work of a research group working on cadastral issues at the University of Melbourne, a prototype has been developed to help demonstrate some of the benefits of distributed processing of spatial data. One of the key processes that has to occur for distributed processing to become a reality is for metadata, for spatial data, to move away from being an end product in its own right to being a background tool. As a background tool it would be used in a very similar way as in a stand-alone GIS/database which uses it for describing the internal make up of the GIS/databases. In this case the metadata would describe the make up of the "virtual database". The metadata would outline where each database was located, its format, access constraints and so on, all of which would be needed to access the remote databases over a network. A metadata engine could use these metadata to parse user queries so that they can be sent to the individual databases. The individual databases would resolve the queries and then return the results to the metadata engine for recomposition for the user to view. The prototype that was developed modified the source code of the software Isite to accomplish this. Figure 1 shows the architecture of this style of system.

Figure 1

Figure 1: Distributed Processing Prototype.

Isite

Isite is a software package that uses the Z39.50 ANSI/NISO, a client/server based protocol for information retrieval (ANSI 1992), standard communication tools to access databases that are stored on a distributed network. The software is public domain software, produced by the Center for Networked Information Discovery and Retrieval (CNIDR), and is freely available for anyone to download and set up a site that could be accessed by remote users interested in their data (Gamiel and Warnock 1994). The clearinghouse that has been established in the USA by the FGDC uses the Isite software as does the Prototype Distributed Database Directory recently established in Australia by ERIN (Hatton 1997; FGDC 1996). The architecture of the Isite system is shown in Figure 2.

Figure 2

Figure 2: Isite Information System Architecture (Gamiel and Warnock 1994).

In the Figure 2 it can be seen that there are two main methods for searching databases that are located on a remote server from your local system. The first method is simply to spawn a Z39.50 client from a local Unix machine with Isite installed on it. This client will make contact with the remote server that is currently running a zserver session and the search of the relevant database can be made through the Search API.

The second method, and more convenient, is to search the remote databases through a WWW browser interface. In this method an initialisation HTML form is used to establish a Http to Z39.50 stateful gateway with the remote server, which is running a zserver session. This once again allows the databases to be searched through the Search API. "Stateful" means that a Z39.50 session is initialised only once and is interactively used by a stateless WWW browser. There is an inactivity timer on the gateway that automatically closes the session after a pre-defined amount of time with no input from the user (Gamiel and Warnock 1994). Figure 3 shows the architecture of the stateful gateway.

In Figure 3 the combination of the http server, zgate and zcon applications represents the gateway running on a single machine. The process that occurs in the event of the user of the WWW browser submitting a query is as follows:

- The WWW browser connects to the http server and posts a HTML form containing

information related to a new or existing Z39.50 session.
- The zgate CGI application then parses that form and either starts a new zcon process or connects to an existing process.
- The user’s request is then passed from zgate to the appropriate zcon, which in turn communicates with the remote Z39.50 server.
- The remote Z39.50 server then passes the results from the query back firstly to the zcon process, and then in turn back to zgate, the http server and finally the WWW browser where the original query was made.
- At the WWW browser the results are then displayed for the user.
- The zgate CGI process then exits, however the associated zcon process stays alive, holding the Z39.50 connection open until a predefined period of inactivity is exceeded, upon which time the zcon process exits (Gamiel and Warnock 1994) .

The applications that are possible when using Isite require accessing text search facilities. For this reason CNIDR developed a Search API (SAPI) which generalises access to arbitrary database systems via a common API (Gamiel and Warnock 1994). Any application that links with this API inherits the functionality of any database system that may reside "behind" the API. For the system administrator this requires the maintenance of one or more text files that describe the databases that are currently available, their location on the file system, etc (Gamiel and Warnock 1994). For the users of the system they should be unaware of the details behind the search. The aforementioned databases are in fact databases that contain the metadata records for all the datasets that exist for an organisation.

Figure 3

Figure 3: Http to Z39.50 Stateful Gateway Architecture (Gamiel and Warnock 1994).

The Data Model

The prototype that has been developed uses a data model that allows the metadata for each of the datasets to be stored at either the central server where the metadata engine is located or on the data custodian’s own server. This approach is taken for two main reasons:

1. The data custodians have control over their metadata when it they are is located on their own server. This makes it easier to keep the metadata to date.
2. It is likely that in the early stages of development some data custodians may not have servers that are capable of participating in the system. Hence they will need to have their metadata located on another server. This is also likely to mean that their data will not be accessible on-line, however the metadata record that is located on the central server will give details on how the user can gain access to it through other means.

Each dataset that is accessible in the system has two metadata "records". These are a metadata record that describes the characteristics of the dataset, the second type described earlier, and an entry in a file that essentially contains the details on how to access the dataset. This entry in the file is the equivalent of the metadata that exists within a stand-alone database and describes the internal architecture of the database. In the case of this prototype the internal architecture metadata represents the architecture of the network, telling the metadata engine where to find, and how to access each of the distributed datasets. It is this metadata that is the key to querying several datasets at once on the network. The metadata that describes the characteristics of the datasets is simply used as it is in data directories or clearinghouses to find datasets that are of interest to the user.

Isite is used in the prototype as a metadata engine by modifying the source code to allow extra metadata to be returned to the engine from the individual databases that hold each metadata record. The original Isite source code returned a hyperlink to the individual metadata records that satisfied the users query. The modified code not only returns this hyperlink, but also extra metadata from the metadata file. This extra metadata is placed in a temporary file on the central server for later access by a simple spatial data display tool. Figure 4 shows the architecture of the prototype that has been developed.

Figure 4

Figure 4: Architecture of the distributed processing prototypes metadata engine.

The engine has been linked to a simple spatial data display tool developed by Polley(1998). The link is made via the temporary file that was filled with the access metadata returned by the query. The spatial data viewer has a menu item that allows the users to select what spatial data they wish to display. This menu item obtains the possible datasets from the temporary file. The display tool can only display spatial data, it has no facility to query the datasets that it displays, other than to change the origin that the user is focused on and to zoom in and out, pan and fit the data.

The prototype is not a true metadata engine as defined earlier. This prototype does have user interaction, unlike the traditional metadata engine found in stand-alone databases. User interaction is required to allow the user to select which datasets they wish to use from the list of datasets that is returned after the user searches for a particular dataset or group of datasets. This enables the user to find new datasets that may have recently come into existence on the network, and allows the user to choose which datasets they wish to use, as not all datasets that are returned via the search may be of relevance.

The sequence of events that is undertaken by the user and the prototype in order to view one or more spatial datasets using the data viewer is as follows:

1. User establishes a http to z39.50 stateful gateway by loading a web page that is located at the central server;
2. User enters the search criteria for the datasets that they wish to find and then hits the submit button on the page;
3. Zserver then sends the search criteria to each of the remote metadata databases, whilst searching its own metadata database using Isearch. Any metadata records that satisfy the search criteria have a hyperlink for them returned the Zserver, and their corresponding entry in the metadata file written to the temporary file;
4. The remote metadata databases are searched by the remote systems copy of Isearch. When a metadata record is found that satisfies the search criteria a hyperlink to the record, along with the metadata corresponding to that metadata record in the metadata file, is encoded into one string which is returned to the central servers Zserver where it is decoded and the metadata that was in the metadata file is written to the temporary file;
5. The hyperlinks to the individual metadata records that satisfied the users search criteria are shown on a web page, along with a button that starts the spatial data viewer. Each of the individual metadata records can be viewed by simply clicking on the hyperlink with the mouse. The data viewer is started by the user pressing the button with the mouse;
6. The user displays the spatial datasets they wish by using the "Available Server Data" button. This button displays a form that allows the user to choose the datasets that they wish to display. Once the user has selected the datasets that they wish to display they simply hit the "Load Data" button;
7. The spatial data viewer then uses the metadata that has been written to the temporary file earlier to access the remote/local databases systems that each of the datasets is located in. The results are returned to the spatial data viewers by the remote/local databases systems where they are displayed.

There are several limitations of the prototype which include:

- The spatial data viewer has essentially no GIS capabilities. It is simply a data viewer that cannot query the spatial data as it has no interface to do so.
- Only one user can use the system at anytime. At this stage the prototype has the name of the temporary file hard coded into it. This should not be a great problem to fix as a random file name can be generated from the process identification number created by the zserver application.
- Firewalls are not catered for in this prototype. If an organisation has a firewall between the outside network and their own internal network the metadata and spatial data will not

be able to be accessed. The easiest way to solve this problem would be to have a copy of the metadata and spatial data located on a server that is not protected by the firewall.

CONCLUSION

Due to the rapid advancements in spatial data collection technologies there has started to be more focus on the best ways to use spatial data to its upmost, rather than the best ways to collect spatial data that may be stored in "electronic silos" and never get used. The development of spatial data infrastructures to their full potential is an important aspect in achieving maximum usage of spatial data as they allow for data to be more easily shared amongst different users.

Metadata engines will be an important component in the not to distant future as they are the key component that allows for the distributed processing of spatial data. Distributed processing of spatial data is one of the fundamental concepts that will have to be mastered in order to achieve such projects as The Digital Earth that has been foreseen by the likes of Al Gore, the vice president of the United States.

The prototype that has been developed at the University of Melbourne is a simple example of how a system allows the distributed processing of spatial data to occur through the use of a metadata engine. The system was developed using the Isite software made freely available by the Center for Networked Information Discovery and Retrieval in the United States and demonstrates basic distributed processing properties.

ACKNOWLEDGEMENTS

The authors wish to gratefully acknowledge the support of the Department of Infrastructure of Victoria and Land Victoria. However, the views expressed in the paper are those of the authors and do not necessarily reflect the views of the Department of Infrastructure or Land Victoria.

REFERENCES

ANSI (1992) "Application Service Definition and Protocol Specification for Open Systems Interconnection" Accessed 20/7/1998, http://www.cni.org/pub/NISO/docs/Z39.50-1992/.

ANZLIC (1996). "ANZLIC Guidelines: Core Metadata Elements Version 1" Report: ANZLIC Working Group on Metadata, July 1996,

ANZLIC (1997) "Core Metadata Elements for Land and Geographic Directories in Australia and New Zealand" The Australian New Zealand Land Information Council, Accessed 5/6/1997, http://www.auslig.gov.au/pipc/anzlic/metaelem.html.

Codd, E. F. (1990). "The Relational Model for Database Management", Addison - Wesley, 538, 2 2.

FGDC (1996) "The National Spatial Data Infrastructure" Federal Geographic Data Committee (US), Accessed 3/3/97, http://www.fgdc.gov/NSDI/Nsdi.html.

Gamiel, K. and A. Warnock (1994). "The Isite Information System, Version 1.05", The Clearinghouse for Networked Information Discovery and Retrieval

Glover, J. (1997). "OPEN GISTM - Removing the barriers to a spatial data infrastructure." Proceedings AGI 96, Birmingham 24-26 Sept, Association for Geographic Information, London pp2.21-2.22.5.

Gore, A. (1998). "The Digital Earth." The Australian Surveyor Vol 43, No. 2 : 89-91 .

Hatton, B. (1997) "Prototype Distributed Data Directory" ERIN, Accessed 4/11/97, http://www.erin.gov.au/database/metadata/isite/index.html

Kildow, M. (1996) "The value of Metadata (A NSDI report)" US Fisheries and Wildlife Services, Accessed 4/9/97, http://www.r1.fws.gov/metadata/meta.html.

Korth, H. F. and A. Silberschatz (1991). "Database Systems Concepts". Singapore, McGraw - Hill, 694, 2 2.

Masser, I. (1998). "Governments and Geographic Information". London, Taylor and Francis, 121, 1 1.

Ordnance Survey of Great Britian. (1996). "National Geospatial Database". Proceedings of the NGD Seminar, July 1996.

Polley, I. (1998). *Personal Communication* "Cadastral Data and the WWW" 22/7/98.

Scanlon, J. (1998). "The Australian Spatial Data Infrastructure and Business". Adding a Spatial Dimension to Business, Mapping Sciences Institute Australia, Fremantle Western Australia, Promaco Conventions Pty Ltd.

Wilson, E. (1998) "Manuals go click" The Age, Melbourne, 30/6/98.

Author/s:
PHILLIPS, ANDREW; WILLIAMSON, IAN; Ezigbalike, Chukwudozie

Title:
The importance of metadata engines in spatial data infrastructures

Date:
1998

Citation:
Phillips, A., Williamson, I. and Ezigbalike, C. (1998). The importance of metadata engines in spatial data infrastructures. In Proceedings, AURISA '98, Perth, Western Australia.

Publication Status:
Published

Persistent Link:
http://hdl.handle.net/11343/33916

File Description:
The importance of metadata engines in spatial data infrastructures