

Location Intelligence: Meeting IT Expectations

The benefits of deploying location intelligence across the enterprise can only be optimized when making an informed choice when considering a vendor's ability to meet the rigorous standards of today's IT department.

“We designed Envinsa to seamlessly integrate into a standard IT environment”

—GEORGE MOON, CTO, MAPINFO CORP

Spatial information, commonly known as “location”, relates to involving, or having the nature of space. Spatial is not constrained to a geographic location however most common business uses of spatial information deal with how spatial information is tied to a location on the earth.

Miriam-Webster defines Intelligence as “the ability to learn or understand, or the ability to apply knowledge to manipulate one's environment.” Combining these terms alludes to how you achieve an understanding of the spatial aspect of information and apply it to achieve significant competitive advantage.

Location Intelligence enables a business to measure, compare and analyze its data from business operations, in conjunction with external data such as transportation networks, regulatory jurisdictions, market characteristics or its own customers. Location Intelligence solutions comprise:

- The analytical capabilities to quantify, measure, compare, analyze and predict spatial data patterns
- Technology that is easy to use, scalable, deployed to where your users are and integrated into your business applications and systems

- Reference data, including geographic data (cities, states, streets) and the attributes of those data (demographics, consumer buying patterns, areas of high risk from earthquakes, etc.)
- Analytical competency and domain knowledge of the business problems at hand to enable the development of applications uniquely suited to specific operational and analytic processes.

Location Intelligence uses a spatial framework to answer complex business problems:

- Where are my best customers located?
- How would this outage impact revenue?
- Where can we open profitable new stores?
- How can we offer greater value to our citizens?

Location Intelligence: Meeting IT Expectations

Enterprise IT Challenges

The IT staff in an organization is responsible for ensuring the computing infrastructure is sound and working to the benefit of its company. Its customers are the employees of the company. An IT organization leverages the domain expertise of these employees whose role it is to fully understand the benefits and capabilities of a location-based system.

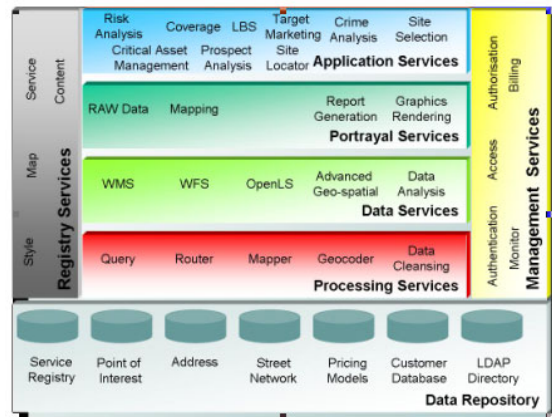
This allows the IT department to focus on ensuring that any software under its responsibility can be administered, can scale as demand dictates and can meet the expectations of its customers. All of this should be accomplished while fitting into an existing infrastructure and without harming existing systems and networks. MapInfo Corporation understands the needs of an enterprise's IT organization and provides proven technology that enables an organization to obtain maximum benefit from location capabilities.

The demands of IT will depend on the needs of its customers. MapInfo Corporation provides location-based software that fits smoothly into the existing network for the enterprise, can access corporately held database systems, can scale to thousands of users when necessary and solves mission critical problems for the organization. It has been built using an outwardly facing web services architecture that behaves as a first class IT citizen. This reduces the pain felt by any organization that purchases a set of technologies that force them to learn unfamiliar skills and does not fit into its existing infrastructure. For location capabilities to benefit an entire organization without being destined to be a point or small department solution, it must take into consideration IT needs.

MapInfo's location platform is Envinsa™. Envinsa provides what is expected in a secure, robust, highly scalable service-oriented architecture and addresses the most critical questions from an IT organization providing a sense of familiarity and confidence from a support capability. The Envinsa platform and services are deployed on J2EE compliant application servers and the monitoring software included for IT's use interacts with these services.

MapInfo Corporation is also a platinum Microsoft developer. The Envinsa location services provides a .NET SDK as well as our MapXtreme® .NET portrayal software. As well, the Envinsa web services definition language

(WSDL) can be used to create client stubs for various languages. MapXtreme.NET has been used to provide a SOA for portrayal and data access services as well. MapXtreme.NET was designed to be a first class .NET implementation and not just a wrapper of legacy code.



Conceptual Architecture of Envinsa

Common IT Considerations

Monitoring Resources and Software Health


In an enterprise system where many systems and processes are working, the IT organization will need to know how to check whether the geospatial software is functioning as expected. MapInfo provides this capability built into its Envinsa location platform with the Enterprise Manager. Both the Simple Network Management Protocol (SNMP) and the Java Management Extension (JMX) are supported.

Being compatible with SNMP allows monitoring with familiar system management software products such as HP OpenView Network Node Manager (<http://openview.hp.com>). SNMP management consists of a network manager and an SNMP agent. A network manager is the software running on a workstation through which the network administrator monitors and controls the different hardware and software systems that comprise a network. The agent is a piece of software running in the network equipment that implements the SNMP protocol. SNMP defines exactly how a network manager communicates with an SNMP agent.

JMX technology allows you to configure, manage and monitor Java applications along with non-Java hardware resources using remote or local management tools in a way that is parallel to SNMP. For more information about Java management and the JMX standard, refer to SUN's web site at <http://java.sun.com/products/JavaManagement>. The Java Management Extensions Instrumentation and Agent Specification document outlines the JMX standard. JMX has been adopted as a standard part of the J2EE platform and our support in Envinsa provides what is becoming expected from IT.

Setting Alert and Response Levels

Not only do you want to monitor the resources and health of the location system, you need to be alerted if there are problems. System administrators usually have several systems they are monitoring. Without some type of notification, they may not know they need to respond to a problem until a customer calls. In a mission-critical operation, this may be too late. Envinsa provides a mechanism to set various alerts depending on the need of the system administrator:




There is also control at a lower level application layer for setting alerts. Supported with the application layer and built into the Enterprise Manager is the ability to remotely send alerts not just to the monitor but to also send emails or SMS messages depending on the severity of an alert. Monitoring as mentioned can be used with any SNMP or JMX monitoring software or for convenience in the Enterprise Manager.

Remote Monitoring and System Stop and Start

IT may not have direct access to the physical hardware in a large organization where the software is running but they still need to be able to be alerted and respond to any difficulties when they are detected. IT personnel on call have a huge advantage if they can get a remote alert and then successfully respond to the difficulty from home or their wireless PDA if they are in the field.

MapInfo with Envinsa provides alerting as mentioned already above. Using a peer-to-peer mechanism in the Enterprise Manager the status of a service can be checked with green showing a service available and red not available.



You can also shut down, start up or restart any service on the network from the Enterprise Manager regardless of location of the service. This gives it the flexibility necessary for wide deployment and management. This is a must for wide-scale, mission-critical deployment within any large corporation.

Logging Control

Logging is necessary to look at past history and help track down difficulties. Another important consideration is that the log does not grow unconstrained and fill a disk with the log information causing other problems. Logging is supported in Envinsa either by using the supplied interfaces or with the Enterprise Manager to provide information used for managing the software in the enterprise.

Location Intelligence: Meeting IT Expectations

There are four main log files used by Envinsa:

- Application server log file.
- Trace log file.
- Transaction log file
- Application log file

The application log file contains system and software error messages. The trace log file has statistics about an individual location-based service. The information gathered includes such items as counter, rank, and frequency of hits. Transaction log files record requests made to any of the Envinsa Web Services on a per service bases for maximum granularity of monitoring. This log file can be used for billing or monitoring usage.

The log files can be managed in a number of ways. The number of log entries to include in a log file can be set. Log values can be displayed in a console window. There is support for directing the output to a web server log. Logging information from different sources can be controlled by severity.

Meeting 5 g's Needs

The more mission-critical the software is to an enterprise's operation, the more important constant uptime will be. The Envinsa SOA was designed from the beginning to be robust. It was built to be a first class citizen of the J2EE environment and to fit into any compliant J2EE container. A decision was made to support the W3C simple object access protocol (SOAP) and discovery through the universal description discovery integration (UDDI) standard. Envinsa was built and tested to meet maximum availability requirements.

The Envinsa platform with its services have had to undergo the scrutiny and testing of a multiple billion dollar company that acted as the prime for a worldwide contract. Envinsa not only passed, but exceeded expectations of the prime. A comment made was that there are many companies that claim an individual service that runs fast under certain circumstances however what is more important is how well many services work together as a unit. Envinsa was recognized for its overall throughput and stability in providing a set of services for complex applications. MapInfo owns the development of all of the underlying services, software and data to ensure the services do work well together.

Allowing the services to work in any J2EE compliant container means the choices for an organization on how they have or want to deploy are not dictated by the vendor. Today's J2EE environments have proven themselves in extremely large operations to support cluster deployments, redundancy and other capabilities necessary for long uptimes. J2EE runs in most operating systems used by various corporations and is often run under UNIX or Linux when long uptime is necessary.

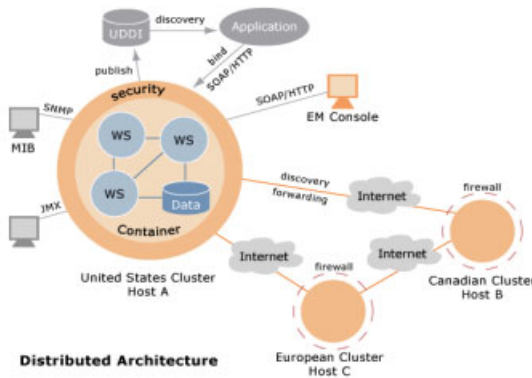
Peer-To-Peer Environment

Envinsa supports a peer- to-peer mechanism that provides even more options for deployment in the enterprise. The Enterprise Manager's peer-2-peer architecture allows for a direct discovery process. Those components registered with an Enterprise Manager peer anywhere on the local or wide area network and even running in different J2EE containers are discovered by any other Enterprise Manager peer. This allows services to forward requests, and lets proxy services be set when a service is offline allowing redirection of unavailable services giving companies even more uptime guarantees.

As a result of the Enterprise Manager's discovery capabilities and service forwarding, an installation can be distributed across a network or the Internet. This gives an enterprise a great deal of flexibility in how it implements a global solution.

The following diagram illustrates an installation that includes geographic distribution of services. In this example the web services (WS) are hosted for three different geographic regions, the U.S, Canada, and Europe, by three different offices (Host A, Host B, and Host C). The different clusters can discover services and forward requests across the Internet. In this example the likely situation would be to provide worldwide service with requests for location information using a supported ISO country code to forward a request to the appropriate country cluster.

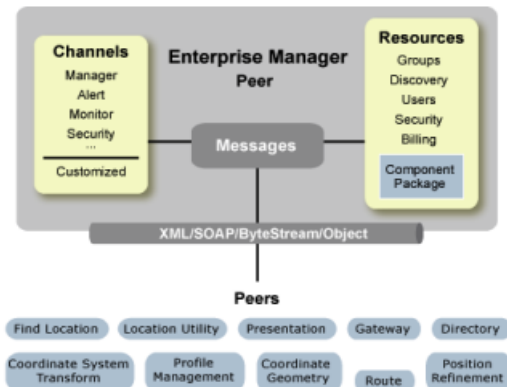
The distribution above can be used among a corporation's existing computing facilities within a country for added security and availability. The Enterprise Manager supports a network of nodes called peers as above. Peers operate independently inside the Enterprise Manager network and have similar functions. There are no specialized directory servers or a centralized resource repository. Therefore,



peers use the network to locate and interact with other peers. The work is performed and the information is shared through interaction between Communication between peers can be direct (one peer to another peer) or propagated throughout the group of peers to maintain high availability.

The peer-to-peer environment allows resources to be distributed across a network, instead of in a centralized repository—thereby sharing resources locally or at various physically separated sites. Let’s look at a simplified example. Instead of one person storing an entire encyclopedia of books at their desk, the encyclopedia can be distributed across several peoples’ desks. When a person (a node) in the group (peer group) wants a different volume of the encyclopedia they can “discover” who has the book by asking the other members of the group.

Channels are used to communicate between peers. Channels are specialized for tasks such as discovery of other peers and resources, alert messages and monitoring thus looking like:



Data Management

A final important piece for high availability is minimizing any downtime necessary for data maintenance. Location Intelligence systems have substantial data requirements including frequent refreshes. There are also many different sources of spatial data. In fact more than 80 percent of existing corporate databases contain spatial information. Having an address is having location data. Thus as a corporation understands the value of using location data in its operations, adding new data sources will be necessary.

The Content Manager provides a convenient application where an organization’s spatial data and metadata can be loaded, updated and defined. It provides an easy way to visually organize the source data and make it available to users. The focus of the Content Manager is on the information provided in the data rather than the structure, storage or databases tables of the data source itself thus avoiding the need to learn new techniques for data management.

The Content Manager provides the mechanism to allow you to set up connections for your data sources such as Oracle, Informix, SQL Server and others that are supported by the Envinsa Location Platform, or as dynamic data sources such as a dynamic Point of Interest (POI) data that are fed live over the Internet. Using the Data Access Component of Envinsa any currently unsupported source can be added.

After the data sources are identified, the content manager focuses on the information or content in the data source, rather than the data source itself. This provides a view of the data sources that can be useful to a variety of users depending on their roles in an organization. For example, you can create one content specifically for an application that requires only three attributes of the ten available in a particular data source. The same data source can then be used to create content for a different user who requires all ten data attributes. Contents permit the data sources to be used as efficiently as possible by all end users while at the same time honoring any of the database management systems previously.

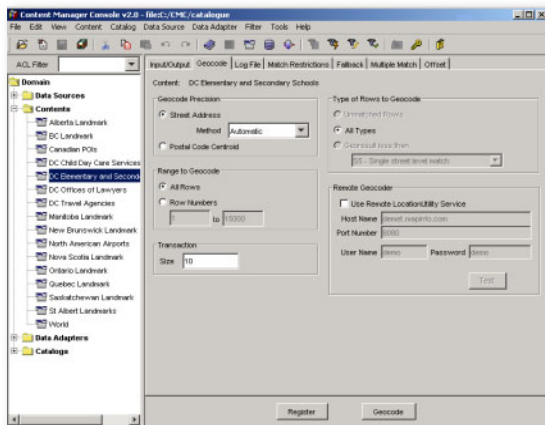
Location Intelligence: Meeting IT Expectations

6

Data Organization

References to the content defined for various roles can be organized into a hierarchical directory structure of Catalogs and Categories. This enables users to quickly and easily find the information they need, and also allows a content to be referenced several times in different places in the Catalog. This directory of catalogs and contents can be created in the Catalogs section of the domain.

By setting up your data using the Content Manager, Envinsa web services have access to the data. All the web services have access to the same Content Manager domains. This means that if one service, such as the Web Feature Service, makes updates to the data, all the other services will see this change. New data sources are added without any down time. The complexity of spatial data is hidden as well eliminating the need for GIS specialists in the IT organization. For example to add customers from a company database as data to be used in location-enabled operations is as simple as referencing the address column in the data source and indicating the data is to be spatially used and letting the built in geocoding ability determine the location of the customers to be displayed or used in location-based analysis:



There is of course built in capabilities for the more advanced user to refine the results if necessary. Again all possible without bringing the system down or needed undue GIS expertise.

Enterprise Security

In a secure system you need to ensure that a request or response is:

- Private and confidential. Accomplished by encrypting the message.
- From a trusted source. Accomplished by authenticating the message using digital signatures.
- In tact and has not been tampered with. Using certificates and digital signatures can also insure the integrity of the message.

Envinsa web services support security in the form of both encryption and authentication to meet all of the above and there are several levels of security available:

- *No security*—any valid request is processed, regardless of whether or not the requestor has an account on the system.
- *Authentication*—a user name and password must be included in the request. The requester must have a user account registered for the web services they can access. Envinsa supports LDAP for authentication.
- *Signing*—a digital signature is included in the request. Only those requests containing a valid registered certificate are processed.
- *Encryption*—requests must be encrypted.

Security can be set up per web server or namespace, or both. This gives the necessary flexibility to apply various levels of security to specific installations on your network. For example, it is possible to set up security only on those services that are exposed on the Internet and not use security for web services that are only used internally.

Interoperability

Whether or not you will be combining Web Services with services from other vendors, Envinsa Web Services follow industry standards, so that they will interoperate with other standard services. For example, you may run the Envinsa Gateway Service within a DMZ, because it can communicate with a third-party Gateway Mobile Location Center (GMLC). You will need to know the processing requirements of these other services and plan accordingly; you will also require a firewall.

DBMS Compatibilities

IT does spend a considerable amount of its time ensuring the DBMS is finely tuned. IT is sensitive to any access that may result in adding more enterprise licenses or requiring additional network demand or hardware.

MapInfo Corporation does not force a DBMS administrator to learn new skills. MapInfo provides the means to connect to many different data sources including the major DBMS offerings without subverting the reasons for having a DBMS. MapInfo software relies on the roles, privileges and other security mechanisms inherent in the corporate DBMS and does not impose a separate, questionable mechanism. MapInfo also allows end user customers or third-party providers to build their own data interfaces into a DBMS. Thus it is possible to have multiple DBMS or other data sources being accessed in a single application without adding complex middleware software. A single query can also span multiple data sources.

Scalability

This may be in terms of the number of users and/or in terms of data volumes. In the case of a location-based offering for a telecommunications company, it may be necessary to support millions of transactions.

Envinsa provides, as explained above, the ability to seamlessly integrate into the J2EE capabilities that allow scaling. Envinsa also provides additional means of balancing outside of a set of clusters as mentioned thus allowing service requests to be balanced according to load either within a cluster or across clusters or both.

Scalability can be measured in a number of ways. This can be in terms of the number of concurrent users, the volume of data that needs to be managed or both. A 'scalable' system may scale for volume of data but not the number of concurrent users. Scaling for both concurrency and volume significantly increases the difficulty when designing software. Add in the complexity within the geospatial software; simple point buffer or buffer of complex shapes and some type of spatial interaction modeling 'best choice', and you can have orders of magnitude differences in computational and data volume performance.

Deployment and Integration

Envinsa provides considerable latitude for deployment and depending on the necessary use case. The use case or cases may illustrate a quick increase in the number of users over time, or an increase in volume of data as the system grows and will help in the understanding of the nature of the data; the mix of static or dynamic. There are too many examples of the 'demo' working well but implementation never meeting the user or purchaser/builder expectations. MapInfo Corporation guarantees that not just the demo works, but we deliver the specified solution.

MapInfo software accesses corporate databases and maintains security integrity in a number of ways. It can be done very loosely whereby there are essentially duplicate data. You get a snapshot from the corporate DBMS. This may work fine when the data are static and there is no need for transactional update of the data. Once there is more dynamic data or where transactions are needed, then the connections become more critical.

A slightly tighter integration is a loosely federated approach. This is where there can be separation of spatial data and DBMS stored attribute data. This minimizes the impact on IT; unfortunately this loose approach means it is difficult to guarantee spatial integrity. The connection between the spatial data and attribute data in the DBMS is 'fragile'. That may be acceptable provided any update of the attribute data is within a proper DBMS transaction (through something like jdbc or obdc) however if periodically the spatial connection for the DBMS has to be re-built you do not have a robust location intelligence solution.

Middleware Options

There are middleware solutions that try to provide some of the advantages of a robust DBMS. That is, they try and maintain some form of spatial integrity. Few "geospatial" providers have IT acceptable spatially extended DBMS middleware. Often the middleware requires the query builder to work in the query language of the middleware.

Even though there are sql query standards, all major DBMS vendors extend the standards so you may not be able to do all queries that you want in order to take advantage of IT's efforts. Also some "geospatial" vendors

provide the much easier to implement sql 89 syntax. All major vendors support at least ISO SQL 92 and more likely ISO SQL 99 (see <http://www.iso.org/iso/en/StandardsQueryFormHandler.StandardsQueryFormHandler?scope=CATALOGUE&sortOrder=ISO&committee=ALL&isoDocType=ALL&title=true&keyword=SQL> for ISO and http://www.dba-oracle.com/art_sql_iso_99.htm for an example of how major vendors make enhancements). So if the spatial middleware requires a different syntax than IT uses, that may be a problem. Compounding this are legitimate concerns about security mechanisms, multi-phase commit support and numerous other issues that IT already addressed with their DBMS vendor and likely are not too keen to re-address.

If middleware is to be considered, it may best be in light of a federated DBMS approach.

Part of the ISO effort has been to define multi-media extensions for SQL including spatial. See the ISO reference above and in particular "Information technology—Database languages—SQL Multimedia and Application Packages —Part 3: Spatial—ISO/IEC 13249-3:1999".

This certainly is the most 'IT friendly' approach if there are standards-based spatial extensions for its existing DBMS. Of course there will still be questions about increase in network traffic and additional demands on computational capacity. A query like:

```
SELECT b.customer_name, b.street, b.city, b.postal_code
FROM building AS b, flood_regions AS f
WHERE b.building_footprint.ST_Within(f.flood_geometry)
= 1 and
f.name = "Albert"
```

is a simple example that likely is not too intensive in computation or volume requirement (also <http://www.btw2003.de/proceedings/paper/68.pdf> and <http://www.acm.org/sigmod/record/issues/0112/standards.pdf> if you are interested in more examples). IT will need to know a bit more about queries though to be able to understand what could be a runaway though request (which they already periodically see and manage). For example:

```
SELECT ST_Intersection(p.spatial_geometry,
r.spatial_geometry)
FROM parcel AS p, river AS r
WHERE r.spatial_geometry.ST_Intersect(p.spatial_
geometry) = 1
```

for North America may take awhile and overwhelm your DBMS resources. Even though IT may not have to know about GIS, they will have to know enough to anticipate certain conditions. Thus even though an ISO SQL approach is the most IT friendly, it will require more education of the IT group than a very loose federation where the solution is perhaps only peripherally 'enterprise'.

In Summary

Since the broader use of location throughout an organization is becoming mainstream, education regarding spatial capabilities is likely worthwhile for the long term benefit of a project. Especially given the location-based extensions for spatial interaction modeling that are starting to appear as extensions to DBMS and which will further enable 'GIS' to be mainstream business. Although some of the concepts and terminology introduced may be new, it is critical that the foundation of the technology itself is based on a solid IT platform.



Schlosser Geographic Systems, inc. (SGSI)

One Union Square, 600 University St
Seattle, WA 98101 USA
1.206.224.0800

www.sgsi.com

MapInfo Corporate Headquarters

One Global View
Troy, NY 12180-8399