

HYDRO INTERNATIONAL INTERVIEWS STEVE HAGAN

It's All About Location, Location, Location



The rate of data growth is climbing exponentially, thanks to new sensors, new social networks and new ways to create and capture data. In response, Oracle is tackling four issues on behalf of its customers: supporting big data, simplifying IT, providing deep analytics and enabling implementations either on the cloud or on-premise. Hydro International speaks to Steve Hagan, vice president of development for server technologies at Oracle Corporation about the organisation's relationships with key geospatial software and data suppliers, its special concentration on very large databases, and its focus on addressing "spatial intelligence" and the mainstreaming of location into business intelligence and numerous related applications.

Is Oracle a total solution provider for the geospatial industry?

No, we are not a total solution provider. For more than 15 years our strategy has been to incorporate geospatial capabilities into all our information technologies – from database systems, to application servers, development tools, BI and analytics, and to our many applications products. We make it easy for database developers, Java programmers, implementers of Service Oriented Architecture systems and package and bespoke applications to incorporate geospatial data and analysis into their systems. Another key part of our strategy has been to provide a secure, scalable, high-performance platform that includes advanced geospatial capabilities to support the largest, most demanding, GIS and geospatial applications in the world.

Because we incorporate spatial analysis and visualisation into many of Oracle's vertical market applications – for example our Utilities Advanced Outage Analytics products – in the context of that application we offer what you would call a total solution. The same is true for a number of other Oracle applications in transport, business intelligence and telecommunications. But generally speaking we are a provider of technology infrastructure and components that add tremendous value to GIS systems and geospatial solutions in conjunction with many data, tool, solution and implementation partners.

How does Oracle interact with the geospatial crowdsourcing projects/community-based mapping?

This is a great development. We have partners that offer Open Street Map in Oracle Spatial and for display with our visualisation component MapViewer. It is very easy to mash up real-time and near-real time crowdsourced data with our MapViewer. We also have Big Data, Hadoop and NoSQL technologies that simplify the consumption and processing of unstructured crowdsourced content.

How does Oracle view Opensource GIS developments?

Oracle has been a big proponent of Opensource in general with Linux, MySQL, Apache, Eclipse and many other technologies. In the GIS arena, we work with GDAL, many of the OSGEO technologies, and are active on the recent activities in Opensource GIS at the Eclipse Foundation.

Oracle has dedicated solutions for large enterprise organisations. What about small and medium-sized enterprises?

Our systems do scale to address the requirements of the largest organisations, but our geospatial offerings are part of many products designed for small and medium-sized enterprises. Firstly, every Oracle database – from our free Oracle Database Express Edition to the Oracle Standard Edition and Standard Edition One – includes full vector geometry support with hundreds of spatial operators, geographic and geodetic coordinate support, coordinate transformation and many other powerful features for containment and proximity analysis. All of our Oracle WebLogic Server products used by all enterprise sizes include the MapViewer web map server and visualisation technology as a standard feature. Our Java IDE, Oracle JDeveloper, includes a Java map component – Geomap. The Oracle Complex Event Processing product supports real-time geospatial stream processing capabilities. And many of our applications products, such as Oracle

Primavera and the Oracle Utilities products, support spatial analysis and visualisation. We really believe spatial analysis and visualisation are valuable for every class of application.

How does Oracle address the ever growing amount of data in the geospatial industry?

We have specialised analytics, data types and storage types for Raster data, 3D and Lidar, networks and topological datasets. These make it possible to manage these data, perform rapid query and analysis, and deliver results through open APIs and web services. Oracle Spatial is designed for parallel architectures that enable us to parallelise spatial operations – query, load, update – on massive datasets. Oracle Database uses advanced partitioning capabilities to enable single logical datasets that can be accessed in parallel with global partitioned indexing. We support features such as Transportable Tablespaces, used by data providers like Navteq and TomTom, which allows datasets that would take hours to load into other systems to be attached, fully formatted and indexed, in minutes. Oracle supports advanced compression which saves money and storage cost that is especially effective with point, point cloud and DEM datasets.

We have recently introduced new engineered systems that process Big Data using Hadoop technology for filtering and categorising unstructured data with high performance connectivity to Oracle Databases and Data Warehouses. Working with extremely large datasets in geospatial environments is something with which we have a lot of experience and it is something in which we are continually investing and improving.

In Oracle's vision, should all geo-data become publicly available?

There is a saying that "it takes all kinds to make a world." It is clear that publicly available geo-data can be a great catalyst for development and innovation, but it would be hard to make a case for all geo-data in all situations.

Could Oracle Cloud technology help those countries with smaller budgets and less geospatial domain knowledge to map their countries in an affordable way?

Yes, cloud technology is one way to start small and scale up without an initial capital investment. But we really don't view cloud as a 'low-cost option'. You can start with Oracle Standard Edition and commodity systems with small capital investments as well. We like to think of the cloud as a flexible deployment option, which is why we design our technology and applications in a way that lets you deploy on-premise, in the cloud or in a hybrid approach.

Are there any developments the geospatial industry should be taking into account?

There are three significant developments. Stream processing is one. There is a vast amount of real-time sensor and sensor-like data that will be filtered and evaluated based on complex rules and conditions, many with a spatial component. Every person carrying a GPS-enabled phone is a human sensor telling us how congested the highway is, where they have gone when the flood waters are rising, or what the hot nightclub spot is on a Saturday night. Spatial stream processing, as supported in Oracle Complex Event Processing, has huge potential for the geospatial industry.

Another significant development is the emergence of map and spatial enabled business intelligence technology as an alternative platform to GIS for business applications. This is the context in which business and non-geospatial literate users are comfortable interacting with powerful geo-analytics. And in-memory BI systems, like Oracle Exalytics, have the ability to deliver these analytic results with sub-second response times.

Engineered systems that are tuned and optimised to perform specific tasks – database operations, BI analysis, Big Data processing – will change the cost and performance curves dramatically. We are already seeing these systems being deployed with remarkable results.

A significant number of players in hydrography (oil and gas, harbours, dredging etc.) are demanding paper-based results of marine geospatial data and are storing it in this way. How could this behaviour be changed?

On the one hand they have a point – not too many CD or mag tapes have been around as long as the Gutenberg Bible. But devices, with the capability of incorporating real-time information and augmented reality may be sufficiently compelling to change behaviour. Look at the kind of thing Nokia is doing with augmented reality on its latest smartphones, where with no context you can find detailed information about the world around you just by pointing the viewer. You can see how this might appeal to industry. Point the device and see what is underground or in the harbour, based on the geospatial and CAD data.

The industry actively wants to collect once, store once and share with many. Especially large data quantities from multi-beam bathymetry, backscatter and above all water column data have to be stored. The data quantities are massive, especially with respect to water column data. How is Oracle storing these types of data streams without having to drastically increase the amount of hardware?

All multi-beam bathymetry data can be stored as Oracle Spatial GeoRaster objects: this is in fact stored as a DEM (Digital Elevation Model) as the bathymetry data gives the depth information at each location. The backscatter data can be thought of as background image (so it can be stored as an 8 bit raster). Therefore, all of the bathymetry data and the associated data can be stored as GeoRaster objects. Since we can compress this data in the database, the storage will be optimised.

In what way is Oracle able to store geophysical data?

Again, most of the geophysical data can be stored as Raster data using Oracle Spatial GeoRaster objects.

Hydrographic offices are still distributing ENC cells (geographically limited by a boundary) to their users (navigators on board vessels) by CD or by file download, while the ENC user-community is demanding Google-like technologies. How does Oracle envision maritime information being stored and distributed to end-users in the maritime industry?

If the data is stored in GeoRaster, we currently have the technology to make this data available via web browser to end -sers using MapViewer. This will give the users the Google-like interface for easy access to the data. With advances in disconnected browsing, you will be able to get the same experience while connected through the web or after downloading to the device.

Where do you think software manufacturers should be focusing their efforts?

Reducing complexity. Complexity dramatically increases cost both in initial implementation and over the life of the project. In addition, usability – presenting spatial capabilities in the context in which the user is comfortable. If we do these two things, we will bring much greater value to our customers and free up countless resources.

<https://www.hydro-international.com/content/article/it-s-all-about-location-location-location>
