

# Ordnance Survey Adopts Grid Computing for Virtually Unlimited Scalability



Ordnance Survey  
Southampton, England  
www.ordnancesurvey.co.uk

## Industry:

Public Sector

## Annual Revenue:

US\$200 million

## Employees:

1,500

## Oracle Products & Services:

Oracle Database  
Oracle Real Application Clusters  
Oracle Spatial  
Oracle Enterprise Manager  
Grid Control  
Oracle Consulting

## Key Benefits

- Created a highly reliable and flexible IT foundation for the organization's Mercury 2 geospatial data mapping system
- Deployed a Linux-based system to provide extra power at lower cost
- Gained the ability to process 1.6 billion rows of spatial data within a maximum 10-hour deadline
- Simplified database management and streamlined storage management

*“The virtually unlimited scalability of Oracle’s grid computing solution enabled us to build a system that could easily process 1.6 billion rows of data. We can now extrapolate production performance with confidence.” – Marc Von Weber, Programme Manager, Mercury 2 Project, Ordnance Survey*

Whether it is hikers walking coast-to-coast, tourists trying to navigate their way around London, or a construction company surveying a job site, the undisputed source for definitive geographic information on Great Britain is Ordnance Survey. For over 200 years, the organization has been providing accurate, reliable, and detailed geographic information for government, businesses, and anyone with a penchant for maps and geography.

Mapping information is vital to government and businesses. Independent experts estimate that its mapping data underpins 10% of all economic activity in the United Kingdom, estimated at \$200 billion a year.

Since 2001, Ordnance Survey has been engaged in improving the currency and accuracy of its authoritative and widely used digital mapping of Great Britain, OS MasterMap®. Because the map is a huge and ongoing commitment to the nation, the organization has sought to improve the reliability and reduce the total cost of ownership of the systems used to publish the data.

To achieve this goal, Ordnance Survey developed Mercury 2, a custom-written map data processing application using Java and PL/SQL (Oracle’s Procedural Language extension to SQL) to process and deliver digital map data to its customers.

## The Mercury 2 Project

To support Mercury 2, Ordnance Survey needed a stable IT infrastructure that could not only handle massive amounts of data, but could also quickly and inexpensively scale to deliver significantly more storage and power when needed.

“We wanted to create an affordable and flexible architecture, so that if we needed to increase the throughput, we could add a few

nodes at a time without having to buy an expensive server” said Marc Von Weber, programme manager for Ordnance Survey’s Mercury 2 project.

In January 2005, Ordnance Survey turned to Oracle to install an enterprise grid computing solution based on Oracle Database 10g and Oracle Real Application Clusters 10g. For ultimate flexibility and manageability, the organization enhanced the system with two tools: Oracle Automatic Storage Management and Oracle Enterprise Manager10g Grid Control.

### **200 Years of Accumulated Geometrics**

The roots of Great Britain’s Ordnance Survey go back to 1791 when the government realized that in planning adequate defenses to repel the threat of invasion from Napoleon, the South coast of England needed to be accurately mapped. So the government instructed the Board of Ordnance- the defence ministry of the day- to speed the necessary work. The first Ordnance Survey map—of the county of Kent—was published in 1801. Two hundred years later—in 2001—the organization began publication of its flagship digital product, OS MasterMap.

Considered the most comprehensive large-scale mapping database in the world, OS MasterMap uses a 660-gigabyte database that records every fixed feature of Great Britain larger than one meter in one continuous digital map.

Keeping the database current—Ordnance Survey claims the map is never more than six months out of date—is a huge job for the organization. Some 478 million topographical identifiers (TOIDs ®) have been assigned to the map. Every 24 hours, it has the potential to process 1.6 billion geometric data points, which it stores on Oracle Spatial, a spatial database management platform that provides support for all geospatial data type and models.

### **Defining the Infrastructure**

In 2004, the organization began looking at ways to create the Ordnance Survey IT infrastructure of the future. In considering its new architecture, Ordnance Survey had three non-negotiable criteria: power, flexibility, and storage.

After extensive study, the organization concluded that it could meet all three criteria by combining open source Linux with Oracle’s grid computing architecture based on Oracle Real

Application Clusters. Oracle's grid architecture makes it possible to coordinate large numbers of low-cost servers and corresponding storage so they operate like one large computer. Oracle Real Application Clusters makes it possible to run one or more databases in such an environment.

By utilizing clustering technology in their grid architecture, Ordnance Survey can add more nodes and ramp up its system at any time—without having to replace the system via an expensive hardware upgrade. “What we were really interested in was the total cost of ownership over a long period of time,” said Tim Warr, principal consultant for Ordnance Survey. “Oracle offered that opportunity with Oracle Real Application Clusters.”

“With our new grid architecture, we can now more easily project our IT costs moving forward,” he continued. “We can never predict how many customers or orders we are going to have, so we needed the flexibility to accommodate almost any volume of requests.”

To assure maximum flexibility, Ordnance Survey configured the system on HP ProLiant servers, grouped into one 10-node cluster with four CPUs per node. It connected the system via an interconnect using an InfiniBand architecture. “We used InfiniBand because it was capable of meeting the speeds that we were processing,” Von Weber said. “And all of that is connected to an IBM RS6000 fiber-channel array and some 3.5 terabytes of data that we spread across various partitions.”

### **Meeting Service Level Agreements (SLAs)**

The high availability delivered by the Oracle grid solution has enabled Ordnance Survey to keep its vast storehouse of data ready for use. This, in turn, has improved customer service by ensuring that users can access the data whenever and wherever they want. Ordnance Survey considers this essential for the service. “It's very, very important that we meet our SLAs,” Von Weber said.

“We have SLAs with our key customers, and for some there are financial implications if the service falls below the required level,” he continued. “The new system is exceeding our expectations. We're now well within all of our SLAs.”

**“With a full Oracle infrastructure software stack, we can perform, we can scale, and we can manage.”**

Marc Von Weber  
Program Manager,  
Mercury 2 Project  
Ordnance Survey

### **Maximizing the System**

Ordnance Survey is also making extensive use of Oracle Automatic Storage Management (ASM), a feature of Oracle Database 10g that provides a simple, consistent storage management interface across all server and storage platforms. The key feature of ASM is that it automatically spreads data evenly across all available storage resources to optimize performance and utilization. This even distribution of database files makes manual performance tuning obsolete—a boon to the organization’s database administrators (DBAs).

Furthermore, Oracle Automatic Storage Management lets the organization alter the storage configuration without having to take the database offline—ensuring system availability. The solution also automatically rebalances files across the disk group after the disk configuration has been altered.

“We’ve leveraged many elements of the Oracle Database 10g environment,” Von Weber said. “We’ve used partitioning, and spatial indexes, interlinking the partitioning with the spatial indexes because they have to be very closely coupled for our data extract.”

### **Capacity for Tomorrow**

Capacity is Ordnance Survey’s ongoing concern, making scalability paramount. For one thing, spatial data processing consumes great quantities of processing power. In addition, the organization’s massive volume of data puts a huge demand on all its input/output (I/O) programs and devices.

“Scalability was a key driver behind our implementation of Oracle Real Application Clusters,” Von Weber said. “At the moment, we still have a lot of customers using our older products. We have a fairly aggressive customer migration program underway, and what we didn’t want to have to do is buy architecture that meets the final demand and have it sit there for two years until our customers make the switch.”

The scalability of Oracle’s grid technology ensures the required capacity and flexibility to meet Ordnance Survey’s unique demands. To cover its bases, Ordnance Survey has 10 nodes—each one with 20 gigabytes of random access memory (RAM) per node and the capacity to upgrade to 64 gigabytes. In addition, all

The Cisco Infiniband switch ports can be bundled to increase bandwidth. At any given time, the system automatically allocates processing resources to cover demand. In the event of a failure, Ordnance Survey can reallocate processing resources to ensure that high-priority services continue to meet required service levels.

“We’ve got plenty of capacity,” he said. “We’re only going to hit a problem if, say, eight of our 10 nodes fail. Then the system will be struggling to meet demand. But if eight of them fail, we’re not doing our jobs as system administrators.”

### **No Sacrifice in Speed**

In the past, speed had been an issue for Ordnance Survey. At peak times queues could form, raising the possibility of backlogs and logjams. “Our Oracle solution delivered performance and speed,” Warr said. “It processes orders very quickly. We never have a queue anymore. We might have one or two complex orders coming in a one-hour period, but two hours later they’ve been cleared completely from the system.”

### **Toward Peak Performance**

Ordnance Survey also added Enterprise Manager 10g Grid Control to its stack. Grid Control provides a single tool that now monitors and manages many of the organization’s Oracle applications.

“It seemed wrong not to take advantage of Oracle Grid Control.” Von Weber said, “What’s interesting is that we’re also now using it very effectively for new projects that come on board.”

### **A Natural Transition**

“Oracle was already in our technology strategy as our preferred database solution, and Oracle Spatial is our preferred way of storing and managing spatial data. So, it fitted in very well with our existing technical strategy,” Warr said.

Oracle Spatial also has built-in indexing and validation features that bring added value to the organization. “The thing to remember with OS MasterMap on Mercury 2 is that it’s a very big database,” Von Weber added. “There are 1.6 billion special features. With Oracle Spatial, we were able to index it very effectively. When we load that data, we validate it straight away using the built-in Oracle Spatial validation functions.”

### Implementation Process

Ordnance Survey began implementing its new grid architecture in 2005 and went live 18 months later. An in-house team from Ordnance Survey Information Systems along with Oracle Consulting managed the deployment. “We completed most of the development process in November 2005 and then tested the system until July 2006,” Warr said.

The project involved writing thousands of lines of Java and PL/SQL, as well as extensive use of Oracle Spatial, none of which bogged down the process. The project team designed the approach to ensure that the database was carefully partitioned and the code written to exploit the parallelism of the grid to its fullest.

“We spent a lot of time testing with data,” Warr said. “There is good reason for that. The quality of our data is essentially our brand and our reputation, so, it’s very important to us that all customer orders are error free.”

### Why Oracle?

Oracle Real Application Clusters and grid technology were a natural choice for Ordnance Survey, which had been using Oracle Spatial for years before it began considering new options for OS MasterMap.

Ordnance Survey selected Oracle specifically to take advantage of the Oracle grid and Oracle Real Application Cluster environment, which was essential to creating the Linux-based system the organization wanted to adopt for maximum scalability and lower costs. Ordnance Survey also had a history of using Oracle products, and saw no reason to switch. “We’ve been using Oracle as a tool within Ordnance Survey since 1988,” Warr said.

*Ordnance Survey has been providing accurate, reliable, detailed geographic information of the Great Britain for more than 200 years.*