



Success Story: The Server Labs

The Server Labs uses Amazon Web Services and Oracle to help Gaia build the largest, most precise 3D map of our galaxy

The Server Labs (TSL) is an European IT company focused on providing expert services in the fields of IT architectures and advanced software engineering. One of their current, high-profile clients is the European Space Agency (ESA) and they are helping the ESA build the science operations infrastructure for the Gaia project. Gaia is a global space astrometry mission with a goal of making the largest, most precise three-dimensional map of our Galaxy by surveying an unprecedented number of stars - more than one billion.

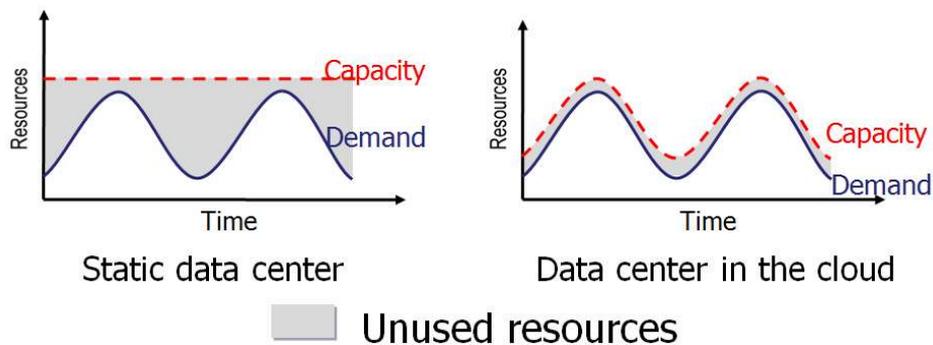
In order to help Gaia reach its goal, TSL team has built the Astrometric Global Iterative Solution (AGIS) to process all the observations produced by the satellite (1 billion stars x 80 observations x 10 readouts). This requires a tremendous amount of data processing. As an example of the magnitude of this project: if it took one millisecond to process one image, it would take 30 years of data processing time on a single processor. Thus, TSL developed their own grid/distributed computing system based on data processing trains.

Originally, TSL began testing the technology on their in-house cluster. The team estimated that with the current data set, it would cost approximately 1.5 million Euros for in-house data processing. But the amount of data was increasing each month and each year. They needed a more scalable, cost effective solution.

Paul Parsons, Founder & CTO of The Server Labs says, "With the magnitude of data and processing power required and the fact that the processing for AGIS is not continuous, it made an ideal candidate for the cloud. Every 6 months we need to process all the observations in as short a time as possible (typically two weeks) and AWS could help us do

that. After due diligence, we discovered that Amazon Web Services had the most functionality of all the private clouds, and we especially liked the self serve aspect.”

The Server Labs estimated that processing the full 1 billion stars data set with 6 years of data would cost \$463,929 on AWS. Whereas, it would cost \$972,147 to purchase the amount of servers required to analyze the same data in the same amount of time, not including bandwidth, electricity, or storage. These calculations helped TSL realize that AWS’s on-demand model would be cheaper and more efficient than buying and maintaining the hardware internally.



Parsons recalls the migration process: “It took us 20 man days to port the software to EC2, and most of that time was spent configuring Oracle. To move the architecture to AWS, we created a 64 bit Amazon Machine Image (AMI) running Oracle Database 11g Enterprise Edition using Automated Storage Management (ASM) on top of Amazon Elastic Block Store (EBS). For the grid software, we created another AMI capable of running the 3 different types of data train used in AGIS. As all the software is written in Java, this process was quite straight forward. To get everything running we had to change only 4 lines of code to solve a thread synchronization problem that only occurred in virtual machines.”

Today, The Server Labs has moved its entire architecture lab and public web site to the cloud and estimates saving 60% of their monthly infrastructure bill. For more on the Gaia project visit http://www.esa.int/esaSC/120377_index_0_m.html or visit The Server Labs at <http://www.theserverlabs.com/>.