The Server Labs G-Cloud 14 Service Description

G-Cloud 14 Service Definition - Google Cloud Platform

Ref: TSL/GCLOUD14/SERVICE_DESC_GOOGLE

Issue : 1.0

Date: April 2024

For: **G-Cloud-14**

Table of Contents

1 Introduction

This document provides you with a description of each of The Server Labs' Cloud Services on Google Cloud Platform

If you wish to receive further information please contact sales@theserverlabs.com

2 Service Definition

2.1 Compute Engine

Run large-scale workloads on virtual machines hosted on Google's infrastructure. Choose a VM that fits your needs and gain the performance of Google's worldwide fiber network.
For more information on the compute engine, please go to the following link:
https://cloud.google.com/compute

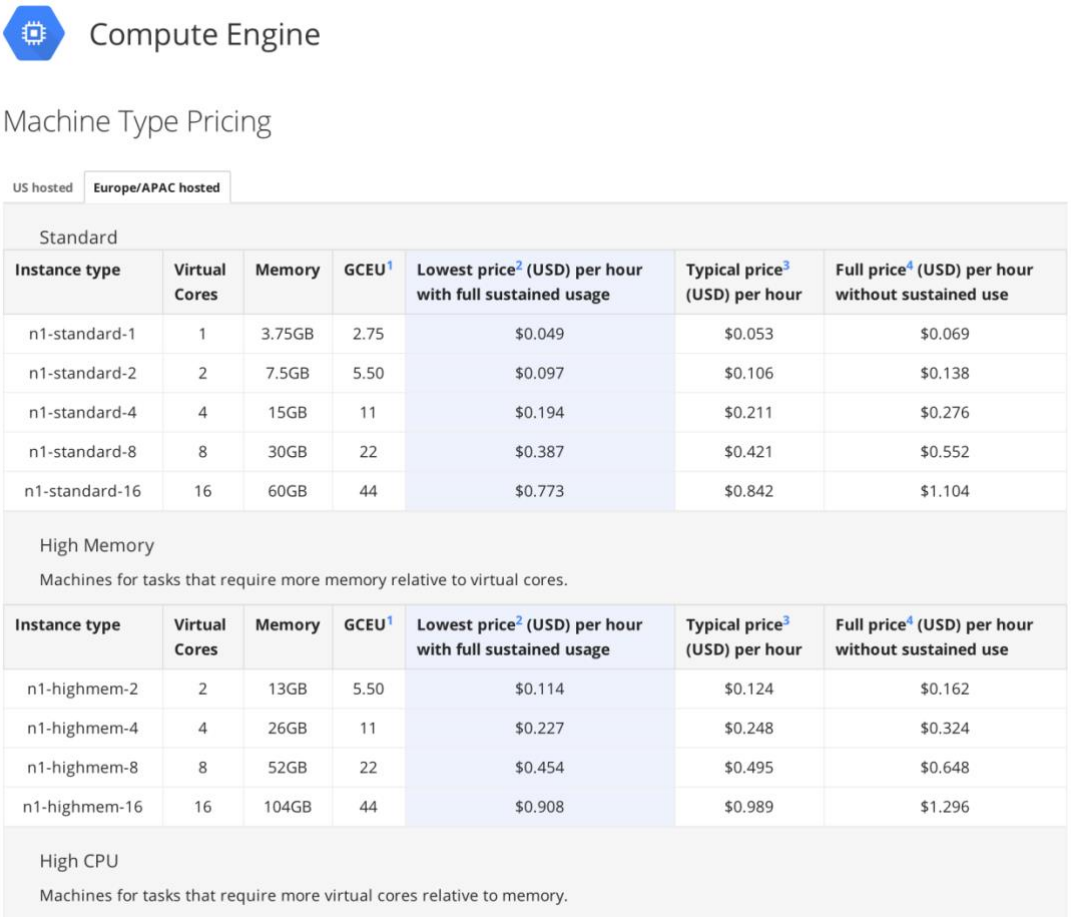

Compute Engine

Machine Type Pricing

US hosted | Europe/APAC hosted

Standard

Instance type	Virtual Cores	Memory	GCEU[1]	Lowest price[2] (USD) per hour with full sustained usage	Typical price[3] (USD) per hour	Full price[4] (USD) per hour without sustained use
n1-standard-1	1	3.75GB	2.75	$0.049	$0.053	$0.069
n1-standard-2	2	7.5GB	5.50	$0.097	$0.106	$0.138
n1-standard-4	4	15GB	11	$0.194	$0.211	$0.276
n1-standard-8	8	30GB	22	$0.387	$0.421	$0.552
n1-standard-16	16	60GB	44	$0.773	$0.842	$1.104

High Memory

Machines for tasks that require more memory relative to virtual cores.

Instance type	Virtual Cores	Memory	GCEU[1]	Lowest price[2] (USD) per hour with full sustained usage	Typical price[3] (USD) per hour	Full price[4] (USD) per hour without sustained use
n1-highmem-2	2	13GB	5.50	$0.114	$0.124	$0.162
n1-highmem-4	4	26GB	11	$0.227	$0.248	$0.324
n1-highmem-8	8	52GB	22	$0.454	$0.495	$0.648
n1-highmem-16	16	104GB	44	$0.908	$0.989	$1.296

High CPU

Machines for tasks that require more virtual cores relative to memory.

Instance type	Virtual Cores	Memory	GCEU[1]	Lowest price[2] (USD) per hour with full sustained usage	Typical price[3] (USD) per hour	Full price[4] (USD) per hour without sustained use
n1-highcpu-2	2	1.80GB	5.50	$0.061	$0.066	$0.086
n1-highcpu-4	4	3.60GB	11	$0.121	$0.132	$0.172
n1-highcpu-8	8	7.20GB	22	$0.241	$0.263	$0.344
n1-highcpu-16	16	14.40GB	44	$0.482	$0.525	$0.688

Shared Core

Machines for tasks that don't require a lot of resources but do have to remain online for long periods of time.

Instance type	Virtual Cores	Memory	GCEU[1]	Lowest price[2] (USD) per hour with full sustained usage	Typical price[3] (USD) per hour	Full price[4] (USD) per hour without sustained use
f1-micro	1	0.60GB	Shared CPU	$0.010	$0.010	$0.013
g1-small	1	1.70GB	1.38	$0.025	$0.027	$0.0347

[1] Google Compute Engine Units

[2] **Lowest price (USD) per hour with full sustained usage**: This is the price per hour with full sustained-use discounts applied, when the machine type is running for 100% of the month. To get a feel for how sustained usage pricing applies in your use cases, try the pricing calculator.

[3] **Typical price (USD) per hour**: This is the price paid per hour under average usage as calculated over all Compute Engine users, and includes sustained use discounts.

[4] **Full price (USD) per hour without sustained use**: This is the price per hour of the machine type when used less than 25% of the month; no sustained use discounts are applied because the instance was not running for at least 25% of the time.

Sustained Use Discounts

Once you use an instance for over 25% of a billing cycle, your price starts dropping. This discount is applied automatically, with no sign-up or up-front commitment required. If you use an instance for 100% of the billing cycle, you get a 30% net discount over our already low prices.

More details

Premium OS Pricing

Pricing for premium operating systems differ based on the machine type where the premium operating system image is used. For example, an f1-micro instance will be charged $0.02 per hour for a SUSE image, while an n1-standard-8 instance will be charged $0.11 per hour. All prices for premium operating systems are in addition to charges for using a machine type.

Pricing for premium operating systems are the same worldwide and do not differ based on zones or regions, as machine type prices do.

More details

Red Hat Enterprise Linux (RHEL) images	$0.06/hour for instance types with less than 8 virtual cores $0.13/hour for instance types with 8 virtual cores or more
SUSE images	$0.02/hour for f1-micro and g1-small machine types $0.11/hour for all other machine types
Windows server images	$0.02/hour for f1-micro and g1-small machine types $0.04 per core/hour for all other machine types

Network Pricing

Ingress	Free
Egress to the same Zone.	Free
Egress to a different Cloud service within the same Region.	Free
Egress to Google products (such as YouTube, Maps, Drive).	Free *
Egress to a different Zone in the same Region (per GB)	$0.01
Egress to a different Region within the US	$0.01 *
Inter-continental Egress	At Internet Egress Rate

* Promotional pricing.

Monthly Usage	Network (Egress) Worldwide Destinations (excluding China & Australia) (per GB)	Network (Egress) China Destinations (per GB) *	Network (Egress) Australia Destinations (per GB)	Network (Ingress)
0-1 TB	$0.12	$0.21	$0.19	Free
1-10 TB	$0.11	$0.18	$0.18	Free
10+ TB	$0.08	$0.15	$0.15	Free

* Pricing for China destinations is promotional for a period that ends on March 1, 2015. On this date, prices for China destinations will change to $0.23 (0-1 TB), $0.22 (1-10 TB), and $0.20 (10+ TB).

Load Balancing and Protocol Forwarding

	US	Europe
Hourly service charge	$0.025 (5 rules included) $0.010 per additional rule	$0.028 (5 rules included) $0.011 per additional rule
Per GB of data processed	$0.008	$0.009

Persistent Disk Pricing

Standard Persistent Disk Provisioned Space	$0.04 GB / month
SSD Persistent Disk Provisioned Space	$0.17 GB / month
Snapshot storage*	$0.026 GB / month
IO operations	No additional charge

* Addition of egress fees for restore will be delayed as a promotional offer until March 1, 2015.

Local SSD Pricing

Local SSD Provisioned Space	$0.218 GB / month

Image Storage

Image storage	$0.085 GB / month

IP Address Pricing

Static IP address (assigned but unused)	$0.01 / hour
Static IP address (assigned and in use)	Free
Ephemeral IP address (attached to instance)	Free

2.2 App Engine

Run your applications on a fully-managed Platform-as-a-Service (PaaS) using built-in services that make you more productive. Just download the SDK and start building immediately. For more information on SDK, please go to the following link: https://cloud.google.com/sdk

App Engine

Services
All applications are free within a usage quota that is reset daily. Paid applications can use more at the prices below.

Service	Free quota per app per day	Pricing if you exceed your free quota
Instances	28 instance hours	$0.05 / instance / hour
Cloud Datastore (NoSQL Database)	50k read/write/small 1 GB storage	$0.06 / 100k read or write ops Small operations free* $0.18 / GB / month
Outgoing Network Traffic	1 GB	$0.12 / GB
Incoming Network Traffic	1 GB	Free
Cloud Storage	5 GB	$0.026 / GB / month
Memcache	Free Usage of Shared Pool No free quota for Dedicated Pool	Free Usage of Shared Pool Dedicated Pool: $0.06 / GB / hour
Search	1000 basic operations 0.01 GB indexing documents 0.25 GB document storage 100 searches	$0.50 / 10k searches $2.00 / GB indexing documents $0.18 / GB / month Storage
Email API	100 recipients	Contact Sales
Logs API	100 MB	$0.12 per GB
Task Queue and Logs Storage	5 GB 1 GB	$0.026 / GB / month
SSL Virtual IPs	No free quota	$39 / virtual IP / month
Bundled Services		
Cron, Files API, Image Manipulation, Pagespeed, SNI SSL Certificates, Socket API, Task Queue API, URLFetch, Users API		

* Small datastore operations include calls to allocate datastore ids or keys-only queries

2.3 Cloud SQL

Store and manage data using a fully-managed, relational MySQL database. Google handles replication, patch management and database management to ensure availability and performance. For more information on Cloud SQL, please go to the following link: https://cloud.google.com/sql

Cloud SQL

Google offers two billing plans for Cloud SQL:

For developers with more traffic, we have package plans that offer a discount and help you predict your costs in advance.

For developers with lightweight applications, we offer a flexible "per use" pricing scheme. You pay only for the time you access your data. Get started with a cloud-hosted MySQL database for around $1 per day.

Packages

Tier	RAM	Included Storage	Included I/O (per day)	Charge (per day)
D0	0.125GB	0.5GB	200K	$0.36
D1	0.5GB	1GB	850K	$1.46
D2	1GB	2GB	1.7M	$2.93
D4	2GB	5GB	4M	$4.40
D8	4GB	10GB	8M	$8.78
D16	8GB	10GB	16M	$17.57
D32	16GB	10GB	32M	$35.13

Additional storage, up to 100GB, is charged at $0.24/GB per month.

External network egress is charged at $0.12/GB.

Per Use

Resource	Charge
D0 Database Instance (0.125GB RAM)	$0.025/hour
D1 Database Instance (0.5GB RAM)	$0.10/hour
D2 Database Instance (1GB RAM)	$0.19/hour
D4 Database Instance (2GB RAM)	$0.29/hour
D8 Database Instance (4GB RAM)	$0.58/hour
D16 Database Instance (8GB RAM)	$1.16/hour
D32 Database Instance (16GB RAM)	$2.31/hour
1GB Storage	$0.24/month
I/O	$0.10/million
Idle IPv4 address	$0.01/hour
External network egress	$0.12/GB

2.4 Cloud Storage

Use a durable and highly available object storage service. With global edge-caching, your users have fast access to your app's data from any location. For more information on cloud storage, please go to the following link: https://cloud.google.com/storage

Cloud Storage

Storage Pricing
per GB per month

Storage Pricing (per GB per month)	
Standard Storage	**Durable Reduced Availability Storage**
$0.026	$0.02

Network

Monthly Usage	Network (Egress) Worldwide Destinations (excluding China & Australia) (per GB)	Network (Egress) China Destinations (per GB) *	Network (Egress) Australia Destinations (per GB)	Network (Ingress)
0-1 TB	$0.12	$0.21	$0.19	Free
1-10 TB	$0.11	$0.18	$0.18	Free
10+ TB	$0.08	$0.15	$0.15	Free

* Pricing for China destinations is promotional for a period that ends on March 1, 2015. On this date, prices for China destinations will change to $0.23 (0-1 TB), $0.22 (1-10 TB), and $0.20 (10+ TB).

Data transfer to other Google Cloud services in the same region is free.

Operations

Class A operations (per 1,000 ops)	Class B operations (per 10,000 ops)	Free operations
$0.01	$0.01	Free

XML API operation classes

Operation	Class
GET Service GET Bucket (when listing objects in a bucket) PUT POST	Class A
GET Bucket (when retrieving bucket configuration) GET Object HEAD	Class B
DELETE	Free

JSON API operation classes

Operation	Class
storage.*.insert storage.*.patch storage.*.update storage.buckets.list storage.bucketAccessControls.delete storage.objectAccessControls.delete	Class A
storage.*.get storage.bucketAccessControls.list storage.objectAccessControls.list	Class B
storage.*.delete	Free

2.5 Cloud Datastore

Use a managed, NoSQL, schemaless database for storing non-relational data. Cloud Datastore automatically scales as you need it and supports transactions as well as robust, SQL-like queries. For more information on Cloud Datastore, please go to the following link: https://cloud.google.com/datastore?hl=en

Cloud Datastore

	Free quota per day	Pricing if you exceed your free quota
Stored data	1 GB total limit	$0.18 / GB / month
Write operations	50k	$0.06/100k operations
Read operations	50k	$0.06/100k operations
Small operations	Free	Free

2.6 BigQuery

Analyze Big Data in the cloud with BigQuery. Run fast, SQL-like queries against multi-terabyte datasets in seconds. Scalable and easy to use, BigQuery gives you real-time insights about your data. For more information on BigQuery, please go to the link: https://cloud.google.com/bigquery

BigQuery

BigQuery uses a columnar data structure, which means that for a given query, you are only charged for data processed in each column, not the entire table. The first 100GB of data processed per month is at no charge.

On-Demand Pricing

Resource	Pricing
Loading Data	Free
Exporting Data	Free
Storage	$0.020 per GB / month[1,4]
Interactive Queries	$5 per TB processed[2,3,4]
Batch Queries	$5 per TB processed[2,3,4]
Streaming Inserts	*Free until January 1, 2015* After January 1, 2015, $0.01 per 100,000 rows for streaming inserts.

[1] Charges rounded up to the nearest MB; minimum 1 MB data per table
[2] Charges rounded up to the nearest MB; minimum 1 MB data processed per each table referenced by a query
[3] The first 100 GB of data processed per month is at no charge
[4] Charges are based on the uncompressed data size.

Reserved Capacity Pricing
For larger, more consistent workloads, reserved capacity can save as much as 70% off On-Demand Pricing. To sign up please contact a sales representative.

2.7 Cloud Endpoints

Create RESTful services and make them accessible to iOS, Android and Javascript clients. Automatically generate client libraries to make wiring up the frontend easy. Built-in features include denial-of-service protection, OAuth 2.0 support and client key management.

Cloud Endpoints

Cloud Endpoints is available free of charge for applications running on App Engine and Compute Engine.

2.8 Translate API

Quickly and dynamically translate between thousands of available language pairs within your app, integrating with Google Translate.

Translate API

Important: Google Translate API v.2 requires billing information for all accounts before you can start using the service. See instructions below in usage fees on how to enable billing.

Google Translate API pricing is based on usage. Translation usage is calculated in millions of characters (M), where 1 M = 10^6 characters. For more information, see the Pricing FAQ.

> **Important:** Google Translate API v2 requires billing information for all accounts before you can start using the service. See instructions below on how to enable billing.

- **Usage fees:**
  - **Translation:**
    - $20 per **1 M characters** of text, where the charges are adjusted in proportion to the number of characters actually provided. For example, if you were to translate 500K characters, you would be billed $10.
  - **Language Detection:**
    - $20 per **1 M characters** of text, where the charges are adjusted in proportion to the number of characters actually provided.
- **Usage limits:**
  - Google Translate API has a default limit of 2 M chars/day. You can increase this limit up to 50 M chars/day in the Google Developers Console by following the instructions below.
  - If you need to translate more than 50 M chars/day, please contact us

To enable billing for your project, do the following:

1. Go to the Google Developers Console.
2. Select a project.
3. In the sidebar on the left, select **Billing**.
4. At the top of the page, select **Enable Billing**.
5. Fill in the form.

To view or change usage limits for your project, or to request an increase to your total limit, do the following:

1. Go to the Google Developers Console.
2. Select a project.
3. In the sidebar on the left, select **APIs & auth**, then select an API.
4. On the API's info page, select the **Quota** link near the API name.

2.9 Prediction API

Use Google's machine learning algorithms to analyze data and predict future outcomes using a familiar RESTful interface.

Prediction API

Free Quota

Usage is free for the first six months, up to the following limits per Google Developers Console project. This free quota applies even when billing is enabled, until the six-month expiration time.

- **Usage limits:**
  - **Predictions**: 100 predictions/day
  - **Hosted model predictions**: Hosted models have a usage limit of 100 predictions/day/user across all models.
  - **Training**: 5MB trained/day
  - **Streaming updates**: 100 streaming updates/day
- **Lifetime cap:** 20,000 predictions.
- **Expiration:** Free quota expires six months after activating Google Prediction for your project in the Google Developers Console.

Note that some Prediction API usage scenarios depend on Google Cloud Storage, which has different pricing terms from the Prediction API.

Paid Usage

Paid usage comes with a 99.9% availability service level agreement. Read more about the SLA here.

Monitor your usage on the Google Developers Console. Any usage beyond the free usage quota will fail if you have not enabled billing for the project. Once you exceed the free quota, the following prices and limits apply:

- **Base fee:** $10 monthly fee per Google Developers Console project.
- **Usage fees:**
  - **Prediction:**
    - **10,000** predictions/month: $0.00
    - **10,001+** predictions/month: $0.50/1,000 predictions beyond the initial 10,000
  - **Training:**
    - **$0.002/MB bulk trained** (maximum size of each dataset: 2.5GB)
    - **0-10,000 streaming updates:** $0.00
    - **10,001+ streaming updates:** $0.05/1,000 updates beyond the initial 10,000.
- **Usage limits:**
  - If you intend to make more than 40,000 predictions/day, please contact us.
  - Google Prediction has a default limit of 2,000,000 predictions/day per project for your own models.
  - Hosted models have a usage limit of 100 predictions/day/user across all demonstration models. Developers can set higher paid usage limits for models they create. Please see each model's documentation for more details.

> **Important:** These charges do *not* include any Google Cloud Storage fees that might be required to hold training data. Please see the pricing sheet for Google Cloud Storage to determine any additional costs that you might incur.

3 The Server Labs

The Server Labs (TSL) is a 100% privately founded IT Consultancy and Software Development Company with headquarters in the UK and offices in Germany and Spain and now established as a leader in Cloud Computing services. The Server Labs focuses on the design and implementation of IT architectures and advanced software engineering projects working with the most advanced technologies to provide its clients cost-effective, scalable and high performance solutions. The Server Labs has been using the Cloud since 2006 and working with its customers in the cloud since 2008 and was one of the first European partner's of Amazon Web Services.

The Server Labs has clients in many different industry areas such as space, finance and telecoms. We collaborate with our clients to obtain success, committed to innovation, enjoying what we do every day and growing with every challenge.

The Company's mission is:

- To provide ***expert services*** in the field of IT architectures and advanced software engineering
- To ***improve radically the software development process***
- **To help organisations achieve better business results through the correct use of latest technologies**
- To have ***100% satisfied clients***
- To create high quality ***innovative software solutions***, providing added value to our customers

The specific value, experience and expertise that The Server Labs can provide for e-LfH are:

1. **Technical excellence** and capability to act as lead on architectural decisions and as technology expert in software and system subjects.
2. **Architecture experience at software and system level.**
3. **Proven experience in HPC and Big Data Projects**
4. **Real cloud computing experience, at IaaS, PaaS and SaaS levels, and for both compute power as well as storage solutions in different clouds.**
5. **Quality control based on ISO9001 for all software systems developed**
6. **Technological excellence**, especially in the main technologies required for the project, including HPC, web services and security technologies.

All our architects and engineers are experts with an average of 10 years' experience in the planning, design and development of complex software systems. Our multinational team has been a pioneer in Java technologies, Object Oriented Analysis and Design and distributed architectures, and has the required hands-on experience in many state of the art technologies. In the last few years, The Server Labs has positioned itself as a leader in Cloud computing services, helping organisations move to the Cloud at all levels. For more information on current projects being undertaken by The Server Labs in Cloud computing, please see Appendix B.

Our experience working across several industries has given us a good understanding of the different requirements so we are able to provide the solution that best suits each particular business and reuse the lessons learnt in the other industry sectors when applicable. Our clients span organisations such as Banks (BNP Paribas, BBVA, Caja Madrid) , the European Space Agency (ESA), Madrid Underground (Metro de Madrid), ICCAT, Amadeus, TRAGSA, TIBCO (a leader in Messaging and Service Bus architecture systems), O2 and Telefonica, Vodafone, ORACLE, several, Sun Microsystems, TUI and Marsans travel, etc. Most recently, we have been architecting, optimising and implementing the migration into the AWS cloud of the Genomics England IT infrastructure, this work followed on from our involvement in the optimising of the HPC infrastructure for the 100k genome project.

4 Selected Customer References

BNP Paribas

Implementation of the Enterprise Development and QA platform

BNP Paribas' department for Risk Systems Development undertook a strategic change programme merging the technical architecture of a number of its key applications to streamline and unify the services they provide internally.
In order to consolidate tools and best practices we customised and implemented our own Enterprise Development Platform which covers the entire application lifecycle from development to ops with its integrated framework.

The implementation of the integrated platform enabled a substantial productivity increase based on a fully controlled, standardised development environment

How we helped

- Integration and customization /Implementation of The Server Labs development and quality control platform for BNP Paribas Risk Systems
- Consolidation of languages and tools
- Controlling and assuring the quality and standardization of development practices.
- Automation of quality and reporting processes.
- Automation of projects office.

Results

- Fully integrated and controlled development environment
- Increased project productivity and quality

Helix Nebula

European Science Cloud

The Server Labs forms part of an exclusive European partnership led by CERN, EMBL and ESA with the aim to establish a federated, sustainable and secure high-performance cloud computing platform. Supported by industrial partners it will provide stable computing capacities and services that elastically meet demand.

How we helped

- We provide our technical expertise based on the successful development of complex Cloud architectures as well as Grid processing in the Cloud both in science and industry environments
- Development of the EC2 bridge in the Helix Nebula BB

Participants

- Consortium formed by: CERN, EMBL and ESA, Atos, Capgemini, CloudSigma, Interoute, Logica, Orange Business Services, SAP, SixSq, Telefonica, Terradue, Thales, The Server Labs and T-Systems, along with the Cloud Security Alliance, the OpenNebula Project and the European Grid Infrastructure (EGI.eu).

- Flagship use cases for testing and deployment of the science Cloud:

    - Cern: HPC Processing of LHC data
    - ESA: Supersites Earthquake data
    - EMBL: Genomic processing

Results of the initiative

- Successful deployment of the science Cloud for flagship projects

ESA

Gaia project

Moving the Gaia project's data processing to the cloud provided ESA with savings of around 50% compared to using in-house hardware. At the same time it provides a level of scalability which means that the work will be notably accelerated.

ESA's ambitious Gaia project aims to create a three-dimensional map of unprecedented size and precision charting the composition, formation and evolution of over one billion stars (around 1% of our Galaxy)

We have provided state of the art architectural solutions for different areas of the project including data management, High Performance Computing, creating a development platform for all the international teams involved and project management.

How we helped

- Design of a Cloud-based solution for the data processing reducing its cost by approximately 50% whilst increasing the flexibility of the testing activities
- Collaboration in the implementation of a distributed computing framework to improve performance.
- Deployment of a development platform that allows the different teams to work in a distributed environment enabling release management, automated testing, continuous integration and quality assurance.
- Database tuning.

Results

- Substantial cost reduction: the developed solution provides an estimated reduction of 50% of Total cost of ownership (TCO)
- Higher performance: Improved performance in the distributed framework being 6x, 7x times faster (in different areas of the processing)
- Productivity boost due to our development platform.

"The Gaia AGIS Cloud experiment has been very successful for us. It indicates that bringing the data processing to The Cloud can provide us with savings of up 50% compared to using in house hardware. An additional advantage is that it gives us the ability to scale to far more processors that we could have in house which means essentially that we can finish the job sooner"
William O'Mullane, Gaia Science Operations Development Manager, European Space Agency

Eumetsat

Technology assessment for the MTG programme

When Eumetsat was planning for a near real-time data processing function for its Meteosat Third generation programme (MTG) it became apparent that this would require an unprecedented level of computing power. Its on-demand re-processing of batch data pushed compute requirements to a speed up to 30 times faster than real time with much higher data volumes than seen in the existing systems.

We have led the feasibility study and the design of prototypes to evaluate grid, cloud and supercomputing architectures to determine the most suitable option and right technologies for the MTG programme to thus enable a long-term reduction of costs.

How we helped

- Technical assessment and cost evaluation of current technologies for Eumetsat's HPC/Big Data needs
- Development of the prototypes
- Evaluation of results and recommendation

Results

- Recommendation to create an architecture based on a hybrid cloud

International Commission for the Conservation of Atlantic Tunas (ICCAT)

Electronic Bluefin Tuna Catch Documentation

The development of an Electronic Bluefin Tuna Catch Documentation system (eBCD) enables the complete tracking from catch to end-market distribution of all Atlantic bluefin tuna operations.
The new electronic system provides greater accuracy and traceability via real-time information helping to more effectively detect and prevent illegal, unreported and unregulated fishing.

Together with Tragsa we are part of the project team to design, develop and implement an Electronic Bluefin Tuna Catch Documentation system (eBCD) for the International Commission for the Conservation of Atlantic Tunas (ICCAT) replacing the previously used manual and paper-based process with a real-time information system.

The solution will enable the complete electronic operability of the catch documentation programme, tracking from catch to end-market distribution all Atlantic bluefin tuna operations.
It will maximize use and performance for end users, providing maximum security and reliability to the system.

How we helped

- Architecture, design and implementation of the eBCD
- Application quality assurance
- Cloud hosting and infrastructure
- User training and maintenance support

Results

- Management of 5000 eBCD/year
- Support for more than 100 concurrent users
- High speed connectivity between service provider and ICCAT Secretariat

Amadeus

Implementation of the TSL Continuous Integration Platform

Amadeus' Internal Information Systems department decided to implement the TSL Continuous Integration platform into their development process in order to optimise the software development cycles and improve development quality, control and automation of the delivered applications.

The Internal Information Systems department works with many developers on a variety of projects. Therefore a standardised platform was vital in order to efficiently reassign team members to activities according to their availability ensuring that neither results nor schedules are compromised. The development platform allows constant tracking of project development health and quality ensuring that simultaneous tasks and projects are running smoothly.

How we helped

- Implementation of the TSL Continuous Integration platform
- Creation of a common framework providing clear mechanisms, standards, tools and conventions enabling the effective execution of both internal as well as external outsourced projects
- Optimization of the current development lifecycle as well as overall development quality, control and automation of the delivered applications
- Seamless transition to Agile development technologies

Results

- Reduction of error prone conditions
- Constant measurability of project health

"With the support of *The Server Labs (TSL)* we have recently undertaken a development platform and processes optimisation project for the development of our internal applications in Madrid. The project outcome has been more than satisfactory.

In addition to implementing effective quality control, *TSL* has helped us to create a framework where best practices are formalised and enforced through a Continuous Integration and Release Management platform.

Because we work with many developers on a variety of projects it is crucial to have a standardised platform in place so team members can be reassigned to activities according to their availability. This way results and schedules are not compromised. We track project development health and quality constantly which is key to keeping simultaneous tasks and projects running smoothly."

Luis-Fernando Gonzalez
Web Solutions Manager, Group Internal Information Systems
Amadeus IT Group SA

CNIO

Migrating the data processing to the Cloud

The Spanish national cancer research organisation CNIO experienced that its vast computing requirements increasingly led to processing bottlenecks causing delays in sequencing projects and research programmes.

The Server Labs have worked with CNIO to help them overcome the bottleneck of data processing in its research programmes.

In an initial feasibility study, The Server Labs established how to transfer CNIO's computing requirements to the Cloud, aiming to make the processing more agile whilst substantially bringing down cost and reducing the need for in-house infrastructure.

How we helped

- Assessment to transfer CNIO's computing requirements to the Cloud enabling agile processing whilst reducing the need for in-house infrastructure
- Created a framework which incorporates all of the strengths of the cloud, in particular data durability, publishing mechanisms and audit trails to make results of experiment reproducible
- Designed and delivered data transfer

Results

- Development of a robust cloud-based platform to perform on-demand genomic processing tasks (at the same time enabling experiment results to be more easily reproduced, stored and published.)

IT Cost reduction of ca. 30% through Cloud-based processing

BBVA

1. Cloud Security

BBVA needed to develop a new secure cloud model able to keep confidential data safe whilst providing a highly productive platform based on Openstack and SDN for the development of internal applications. The new private cloud environment uses continuous delivery with automated security mechanisms.

In addition to substantially reducing the risk of security breaches based on web services the automation of the most important security operations tasks are now based on SDN solutions allowing a reduction of development time by over 20-30%.

How we helped

- Design of the project architecture and definition of the security platform and API services
- Deployment of the security platform and all operational processes
- Development of security policies
- Defining Best Practices for security

Results

- Automation of the most important security tasks based on SDN solutions
- Security policy is already applied by default decreasing develop time by over 20-30%
- Quick response time when facing changes in the integration processes
- Reducing the risk of security incidents based on web services

2. Corporate Security Application

The Server Labs (TSL) has developed and implemented an entirely new version of BBVA's corporate security application Faro. The application which has won a national award for the best application in the area of security has been unified and streamlined throughout all countries in which BBVA operates making its management and maintenance more flexible, transparent and cost efficient.

The new flexible, efficient and scalable application supports the management of all security functions of the BBVA group including for example dashboards, financial control, incident management, inventories and security personnel.

How we helped

- Development and implementation of a new and revised version of BBVA's global corporate security application Faro
- Deployment

- Designing the modular, multi country architecture
- Coordinating the multiple actors (Client, data centre, development team) for the go-live
- Deployment of TEST and development environments in the cloud

Results

- Flexible, transparent and cost efficient application management and maintenance
- Highly agile application
- Cost reduction - estimated 20% decrease of FTE in the corporate security area

Working with The Server Labs has been a very positive experience. This has also been demonstrated by the success and public recognition of our latest programme, the development of a unique technology platform for our global corporate security management.

The Server Labs has been instrumental in the design and development of its scalable and streamlined architecture enabling the worldwide rollout.

The Server Labs really stood out with its superior technical knowledge on all levels, its responsiveness and agility, and acting as our right hand throughout the entire project.

Inés Díaz Ochagavia
Director Production Corporate Security
BBVA

3. Continuous Integration Platform

With over 50 systems exchanging financial information using more than 2200 interfaces, maintenance costs and operational risk are high. BBVA wanted to significantly reduce the number of interfaces, unify the data format and thus be able to lower maintenance costs, quickly respond to changes, and centralize monitoring.

BBVA designed its new continuous integration and quality environment employing The Server Labs Enterprise Development Platform in the Cloud.

The development platform does not only control quality and improve productivity but is also instrumental enabling homogenized and standardised developments for all of BBVA's development projects, and as such providing a corporate reference architecture.

How we helped

- Supporting project organization and management
- Simplification of existing processes
- Development of integration processes, using a common framework.
- Design and development of common components, allowing later reutilization in other projects
- Design and development of auxiliary tools (maintenance, monitoring of existing processes)

Results

- Fast response to changes in the integration processes.
- Reduction of incidents in the integration software in operational environments.
- Standardization of developments.

4. Technical audit of Tugestionline

The Server Labs provided a technical audit driving the design and execution of a series of tests to boost the performance of BBVA's TuGestionLine. The web application, constantly growing in users, provides online accounting counselling for businesses and individuals.

How we helped

- Full analysis of the system architecture
- Design and execution of a series of tests to provide an accurate diagnostic about the web application's performance, quality and reliability
- Identification and quantification of possible improvement points to optimize performance, system stability and ability to deal with load peaks
- Load and capacity tests

Results

- Delivered a full performance analysis with a complete set of code fixes, reconfigurations and recommendations in to improve functionalities
- Web load time was substantially reduced tangibly improving the user experience

Webtools

Migrating e-encuesta.com to the Cloud

The growing business of Webtools' leading online survey application e-encuesta.com with its pronounced peak-times was increasingly causing server overloads limiting the number of clients that could be served simultaneously.

In order to give the application the ability to flexibly react to changing demand it was necessary to change its data centre-based system to a Cloud-based structure.
We successfully adapted the architectural base and executed the smooth migration of e-encuesta.com from a traditional data processing centre to the public Cloud using The Server Labs Rapid Cloud Transition as its central tool.

How we helped

- Designing a robust and efficient cloud architecture specific for the client's project.
- Security upgrades and enhancements to the administration of application.
- Development of a self-healing architecture, allowing that when problems are detected the architecture is able to react and reset the environment
- Provide a pre-emptive support to the client in order to make sure the application is evolving properly in time and the system is up to date.

Results

- Substantial reduction of operating cost
- Service now fully aligned to e-encuesta's needs
- Resources elastically aligned with e-encuesta's usage
- Improved performance and user experience: much higher response time and uptime
- Almost unlimited scalability supports its business in its continuous growth
- Increased reliability and stability of the environment

The Server Labs

Founded in 2004, THE SERVER LABS is specialized in the design and agile deployment of technology solutions in the area of enterprise integration, software engineering, cloud computing architectures and Big Data. Through our offices in the UK, Spain and Germany we offer our customers cutting-edge, cost-effective, scalable and high performance solutions enabling them to stay competitive with the most suitable and innovative IT technologies and solutions.

The list of our customers includes renowned organisations such as the European Space Agency, leading banks in Europe and Government bodies.
Our services span from early strategic technology consulting and architecture services to the full implementation and agile deployment solutions.

More information about The Server Labs is available on www.theserverlabs.com

Metro de Madrid

1. Enhancement of real-time driver information system

Metro de Madrid required a set of new functionalities and developments to enhance the efficiency and reliability of ARCO, the real-time information system for its train drivers.

The improved design of ARCO brings scalability to the system through a new component, allowing growth in both lines and functionalities, and increasing the system's reliability. At the same time it helps resolve capacity issues, bringing down use of CPU for the exchanged components, freeing up valuable compute power for other systems.

Delivered on-time and successfully installed on a live environment, we provided an iterative delivery solution, based on Agile methodologies, developing and releasing the new functionalities on a progressive basis.

We aimed to improve the overall software quality with each new developed component, as well as by refactoring the ones to be modified.

How we helped

- Development of an iterative delivery solution based on Agile methodologies for progressive development and release of the new functionalities
- Improvement of the overall software quality through new development components and refactoring of those to be modified

Results

- Successfully delivered a new version of ARCO, packaging items of several software layers
- Notable improvement of the overall project management initiating the use of collaborative tools like Jira

2. Code Audits

We performed a number of code audits on applications developed by Metro in order to provide visibility of the quality of the code thus enabling informed decisions about the future courses of the respective projects.

How we helped

- Comprehensive analysis of the applications in question and quality of their code
- Flagging up potential problems and areas of improvement

Results

- Metro gained a complete overview of the code quality and was able to make information-based decisions on the further project development
- In some instance the analysis enabled a strong position for renegotiation with previous providers the improvement of code quality
- Documentation of alternative solutions
- Overall reduction of code decreasing the potential for errors

Unit 4

Scalability certification of the Unit4 Health application

In order to present their Health application to medium and large hospitals Unit4 wanted to validate its performance and sustainability with a substantial load of users.
After creating a virtual cloud environment we ran tests simulating a real life load with up to 4,000 concurrent users injecting 4 new users every second.

The simulations confirmed that Unit4's Health application was performing well above average results achieved by reference hospitals.

How we helped

- Creating a virtual environment using Amazon Web Services
- Planning, set up and execution of simulations
- Application and Database fine tuning, and recommendations for optimal configuration.

Results

- Various simulations were successfully concluded confirming the application's ability to handle a large number of concurrent users surpassing the results of reference hospitals.