

PInCom project: SaaS Big Data Platform for and Communication Channels

Juan Manuel Lombardo, Miguel Ángel López, Felipe Mirón, Susana Velasco, Juan Pablo Sevilla and Juan Mellado

Fundación I+D del Software Libre, Granada, Spain

Abstract — The problem of optimization will be addressed in this article, based on the premise that the successful implementation of Big Data solutions requires as a determining factor not only effective -it is assumed- but the efficiency of the responsiveness of management information get the best value offered by the digital and technological environment for gaining knowledge. In adopting Big Data strategies should be identified storage technologies and appropriate extraction to enable professionals and companies from different sectors to realize the full potential of the data. A success story is the solution PInCom: Intelligent-Communications Platform that aims customer loyalty by sending multimedia communications across heterogeneous transmission channels.

Keywords — Optimization, NoSQL, Relational DBMSs, Big Data Analytics.

I. OVERVIEW

IN the Fundación I+D del Software Libre (Fidesol), as a research center, we do analysis and experimentation with technological trends to improve the achievement of project goals as well as provide new elements of innovation research organizations and SMEs. We have a research line focused on Big Data and optimization analysis of large data sets to provide our beneficiaries generating value and improving business processes.

Research in “Advanced Analytics of Big Data” [1] is essential to help organizations consolidate technological advantage obtained from the discovery of patterns, trends and useful information to help develop reliable diagnostic and prognostic in many areas.

Companies that take advantage of the potential of data reach a better position to plan, strengthens their capabilities for knowledge discovery and the prediction of behavior in uncertainty scenarios. Knowledge discovery in databases continues extending to almost every field where large amount of data are stored and processed (databases, system logs, activity logs, etc.) [2].

Potential of Big Data was shown in a McKinsey [3] market study. For example, retail trade, using the full capability offered by the Big Data could increase their profits by 60%.

Also, European governments could save more than € 100 thousand millions in operating efficiency improvements.

According to the EMC Chief Marketing Officer, Jonathan Martin, inside study carried out about digital universe in 2014, conversion of the organizations defined by software companies is supported: "IT must press the reset button and find new ways to optimize storage and take advantage of these masses of data" [4].

We Fidesol create Big Data efficient and quality solutions that respond to the most demanding needs and satisfy the expectations of our customers.

We are introducing PInCom project, an intelligent platform for

communication, that shows by using more appropriated advanced strategies and Big Data optimization techniques, that improving of results is possible.

This article is structured as described: Section II shows a summary of PInCom platform, objectives and added value to ICTs area. Under Section III can be found system specifications and an analysis of open source technologies applied for Big Data treatment in this database-focused SaaS platform. Section IV describes system architecture and cloud deployment as well as its benefits. Section V reports the results of the developed pilot to verify and ensure the compliance with the objectives fixed. Final Section VI talks about conclusions extracted from the adoption of this Big Data use.

II. PINCOM PROJECT

The team has also developed the “PInCom” project, aligned with a line of research with greater continuity for Fidesol, which is Cloud Computing [5], technological area of interest in our company for its universal and ubiquitous character.

The Cloud is an essential component of any application and, according to analysts, the Cloud services market will grow exponentially in the next few years, along with the Internet of Things industry [6].

Therefore, it is considered that may be of interest to integrate aspects of cloud computing with the treatment of large volumes of data and the potential and versatility of messaging through various formats and platforms, especially on mobile devices, covering another Fidesol major lines, such as R & D applied to mobile technology.

The importance of PInCom project - Intelligent Communication Platform - for Fidesol is based on the quantitative importance of critical mass or target population of the project, ie, the high number of users who have access to those technologies and media that allow access to existing communications networks to receive information of interest. The main objective is to offer a cloud service for ubiquitous marketing and loyalty by sending messages using different protocols and communication channels.

PInCom contribution to ICTs is expressed by:

- The system is scalable and highly available since its inception. This is because the system will be available on a cloud computing platform that allows to increase the computing power of the system as needed.
- It supports sending to several heterogeneous communication services, both sending the same information to all media or, by spreading and diversifying depending on the communication channel and the message you want to provide.
- It is an open platform that easily facilitates the incorporation of new services.

PInCom defines the building and deployment of a high performance Big Data SaaS platform efficiently able to process a large volume of communications. Its cloud architecture provides:

- High availability: The system ensures operational continuity, which means that the workload of the system does not affect the number of connections allowed.
- Fault tolerance: The system must be able to continue normal operation after a software/hardware failure.
- Scalability: The hardware features of the system are improved as required, while growing the number of connections or database computing/extension.
- PInCom can dispatch a large number of applications, with low latency, so that it is usable simultaneously by a large number of customers.

The server must meet the objectives of providing a new fully functional platform, able to handle a lot of traffic and send messages by using intelligent routing through the use of standardized protocols, acting as a gateway bursting and managing information security when they connect.

III. TECHNICAL DESCRIPTION OF THE SYSTEM

The System is a modular application that offers different cloud services:

- User interaction: The PInCom platform provides a web interface that allows complete management of the system after user authentication. On the other hand, allows through a mobile web and native mobile applications, any user with an Internet connection would be able to send communications through the gateway.
- Interconnection with other systems: The system receives requests to send message through standard IT protocols, and transfers these requests to other service providers using their interfaces.
- Intelligent System for optical route detection: PInCom uses heuristics and connections to different nodes in communication networks to select the optimal route for each message, preventing congestion and fault tolerance connection with each supplier.




Fig. 1. PInCom's components

PInCom is offered as a modular system that will offer different services deployed in a cloud offering:

- High availability.
- Fault tolerance.
- Scalability.

This will increase the computing requirements as needed.

But to provide this service is not enough to have a conventional clustered database manager considering the huge amount of data to be processed. Due to its specific characteristics (3 Vs) [7], Big Data requires technologies to collect, store and manage very large amount of various data, at high speed in order to allow information optimal processing and get the required results.

A. Analysis of Open Source NoSQL distributed databases

Relational databases have been the choice of many developers and system administrators, opposite to adoption of noSQL-BD, at an early stage. Relational Database Management Systems (RDBMS) are designed for sets of small but frequent transactions of reads and writes, compared with NoSQL databases which are optimized for transactions requiring high load and speed of delivery. The loss of flexibility at run time compared to conventional systems SQL, is offset by significant gains in scalability and performance when it comes to certain data models [8].

The NoSQL BD contains complex information structures in a single record, different information each time, providing polymorphism of data, accessibility and performance, to have the information linked in the same block of data. The ability to process a large amount of data is the ability to run queries type Map-Reduce [9], which can be launched on all nodes at the same time and collect the results to return.

In addition, such systems operate in memory and the data dump is done periodically. This allows very fast write operations. Instead the data security is sacrificed, existing the risk of loss and inconsistency that is treated decreasing the time of offset disk dump and validating the writing operations by more than one node. Have to highlight from a list of 150 NoSQL databases: Cassandra [10], with storage capacity for real-time processing and Terracotta [11] that provides organizations through cloud services increased scalability, availability and high real-time performance in their activities and processes.

B. Cassandra and Terracotta systems evaluation results

Researching results (Cassandra and Terracotta) proposed for the new PInCom architecture -whose purpose is to provide the high availability system at all levels and scalability-, are set out below.

Cassandra evaluation

Using a NoSQL database in PInCom and more specifically, Cassandra, arises for two reasons:

- Performance: The performance of the system in the processing of repetitive tasks on the database (query and modify users and communications) may be increased by providing the data cluster of new nodes.
- Stability: The use of a system of distributed databases with automatic data replication and fault tolerant allow PInCom remain stable even when one of the storage nodes collapses.

The main features that make Cassandra an interesting alternative to MySQL according to where more reads and writes are performed are:

- Linear scalability: Performance increases proportional to the number of nodes.
- High Availability: Cassandra's architecture is fully distributed. There are no single points of failure or bottlenecks. All cluster nodes are identical.
- Fault tolerance: To ensure this, the data is automatically replicated. If a cluster node goes down, service is not stopped and housed data inside the node is still available in the cluster.

After testing using Fidesol infrastructure, it was found that: Node-level performance in Cassandra node insertions concerning the volume of data is slightly worse than that in MySQL. This problem can be

solved by distributing data between cluster nodes. This can always be sized so that the nodes do not have an excessive amount of data.




Fig. 2. Cassandra and MySQL queries performance concerning data volume

Read and update operations performance at node level in Cassandra concerning data volume, grows with a lowest rhythm than in MySQL.




Fig. 3. Cassandra and MySQL queries performance concerning data volume

Cluster-level in Cassandra grows linear, duplicating that in MySQL when using four nodes. This is main advantage for Cassandra, due to cluster performance can be adapted to the load of the system by adding new nodes.




Fig. 4. Cluster-level Cassandra performance concerning number of nodes

This performance is divided by the replication factor (number of replicas of each data stored in the cluster). In tests the value designated for the replication factor has been one.

The stability of a Cassandra node is not higher than a MySQL server when the load limit is maintained. In both cases, the server collapses and, after a given time, the service becomes unavailable. However, since Cassandra offers high service availability and data, this problem at the node level that would lose criticality in a system without high availability.

Terracotta evaluation

Terracotta is proposed as a solution for sharing data between nodes in order to:

- Parallel processing of client requests.
- Establish a mechanism of coordination between processing nodes for distribution of connections to routes, ensuring that for every path, the corresponding connection is established on a single node and the distribution of connections between nodes is consistent, based on some defined criteria (number of connections per node, sum of shipping rates of each node path, etc).

Fidesol testing is focused on the “BigMemory” product in its free version. This version does not have all the functionality offered by the paid version, being limited in cache size allowed, allowed number of customers, etc. However, the basic skills are the same, so it is presented as a viable alternative to test Terracotta as a solution to the sharing of data between processing nodes in PInCom.

Validity is verified, although there is a downside: BigMemory doesn’t has the ability to transfer objects partially, so that sharing of data structures of arbitrary size (which will grow accordingly to the increased load on the system) will be a limitation for horizontal scalability.

The following graph shows the evolution of the time required for access to an object in the BigMemory cache depending on the size. BigMemory is found that does not have the partial load capacity of objects, so that the trend is linear. This is a limitation on system scalability, as the growing load of incoming requests to the system will be an increasing load of communication between processing nodes and Terracotta servers.




Fig. 5. Time evolution graph to access and object located in BigMemory cache concerning its size.

Therefore, Terracotta BigMemory is not a valid solution. One possible alternative is the use of this software at a lower level and there is not a commercial product for this purpose. Terracotta DSO (Distributed Shared Objects) is the basis of Terracotta products and supports the partial load of certain data structures [12].

This partial data loading allows horizontal-scaling of PInCom processing cluster without overloading the network traffic. However, test results have not been positive: Terracotta DSO can be integrated with Glassfish V2, but not higher, the latter being essential for the deployment of PInCom.

Moreover, Terracotta DSO can be integrated at JVM level.

Fidesol testing have not been positive about it, being apparently not feasible such integration in JDK7. In JDK6 it has not been achieved verification successfully.

In addition DSO is no longer available within the packages offered. Because of this, the use of alternative solution for data sharing arises, such as:

- Hazelcast: Apache, supported by benchmarks and success stories.
- Cassandra: it’s also considered the possibility of using NoSQL database to house the shared data structure. This would mean a

loss of capacity in terms of timing and consistency that Terracotta and Hazelcast [13] are offering, and must be implemented at the application level in PInCom.

IV. PINCOM ARCHITECTURE

Below can be found PInCom architecture and the benefits of its implementation:

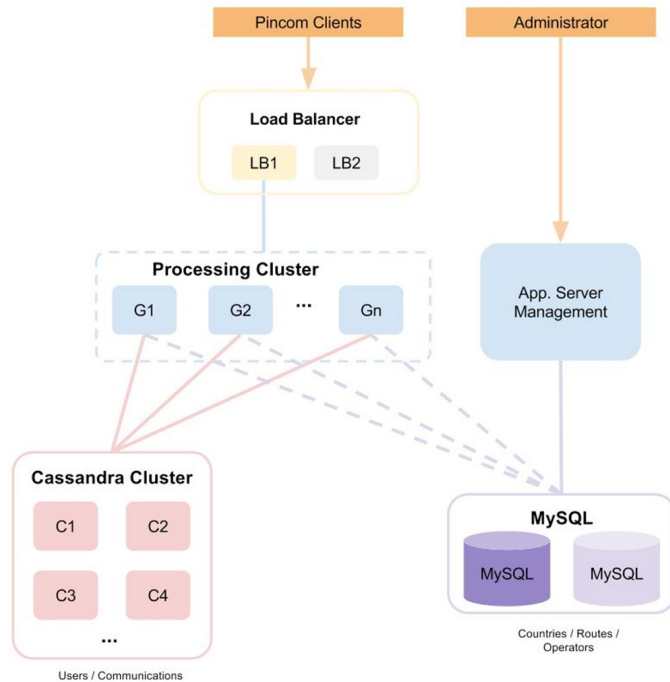


Fig. 6. Proposed architecture

Load Balancer

This machine will be performing balancing tasks toward processing nodes. Balancing will be made at TCP level. Due to task simplicity, it is not necessary a high-performance machine, being enough with the use of a basic machine. However this machine is considered as a single point of failure, so the component goes important for the system high availability.

Processing cluster

All processing nodes (known as G1 and G2) implement parallel processing requests from clients. Moreover, each connection to a route can only be established in a processing node, so there will be a mechanism for the distribution of routes between nodes. To make this work in a coordinated way, nodes use a series of data structures shared by Cassandra:

- Incoming client-requests queue: This structure contains requests that have been accepted by the system but have not yet been processed.
- Routed queue requests: There is a structure of this type for each system path. Each element inserted into one of these queues contains information of a request made by a client and the list of routing tables that will be used for sending this message.
- Routing table: Processing nodes will use this structure together with a coordination protocol for consistently sharing routes.

Administration server

This machine houses the system management application, where the administrator user can manage the static system information in. Due to its function, it is not required its scalability and so a single machine is

dedicated.

MySQL server

MySQL settings will not be different from the proposal so far: two nodes with master-master replication and high availability through heartbeat. This database is used for the following purposes:

- Storage for application management (user management, countries, operators, routes, etc). As the management application has no requirements for scalability, database on which it is support neither does it have.
- Preload of static information for processing nodes. Once this information has been preloaded, processing nodes will only access to Cassandra Cluster for reads and writes.
- Update of static information in processing nodes: When the administration application makes a relevant modification, it will request reload the static information to the processing nodes.

Cassandra cluster

The proposed configuration includes at least 4 nodes and replication factor 2. This will provide high availability and an acceptably high performance.

For the processing cluster to take advantage of data cluster performance, the processing nodes create sets of connections to random nodes of the data cluster and balances the load between these connections. The size of the connection set does not match the data cluster size, as this would compromise the scalability of the system.

It also acts as a synchronization tool between the processing nodes, storing shared structured. The synchronization between the processing nodes and preservation of data consistency is full implemented in these nodes.

V. PILOT PHASE

In order to evaluate the PInCom platform functionalities, a pilot phase has been performed to verify the capacity, scalability, performance and efficiency of the system, critical factors to success.

- Verification of the system load for different sizes of processing cluster and data cluster, preserving the scaling ratio between these clusters.
- Evaluation of the system performance with different volumes of requests for different cluster sizes.

A. Project pilot definition

Pilot was built in an iterative way, going from 2 to 12 nodes for processing cluster and Cassandra Cluster, two nodes for MySQL replication and one node for administration service. So we can verify the scalability of the platform and performance that provides this scalability in already mentioned cluster.

The hardware supporting the nodes is as follows:

For processing cluster, servers with 32 Gb of RAM, 8-core, redundant power supplies and RAID 5 have been deployed. In this case RAID 5 has been chosen, for its high performance in reading operations and its high fault tolerance, this ensures that no loss of service if a disk fails.

For Cassandra cluster, servers with 32 Gb of RAM, 8-core, redundant power supplies and RAID 10 have been deployed. We have chosen RAID 10 due to the huge amounts of data that Cassandra writes, as it combines high performance in writing data of RAID 0 and fault tolerance of RAID 0.

For MySQL configuration we have chosen servers with 32 GB of RAM, 8-core, redundant power supplies and RAID 1. We have selected a RAID 1 for it, because storing more static data, high performance is not required and results in cost saving.

Have deployed a server with 32 GB of RAM, 8-core, and redundant power supply and RAID 1 for the management server. Same applies that in MySQL server. It is not required high performance since this server can only be accessed by the system administrator

B. Project pilot development

The following chart shows the maximum load supported by the system, depending on the processing cluster size and data cluster. The number of nodes of both cluster is shown on the horizontal axis -the tests have been performed scaling both cluster at the same time-. In the vertical axis the maximum rate of requests per second that the system can reach is shown.




Fig. 7. Maximum load vs. cluster size

The following chart shows the performance of the system to different rates of requests received for different cluster sizes.




Fig. 8. Throughput vs input load

C. Project pilot results

As shown in the charts, it has achieved the aim of desired scalability and performance since the performance increases almost linearly with the cluster size.

To increase the volume of communications to manage while keeping response level and high-speed of computing capacity, the potential of processing and data clusters should be increased.

In addition, high availability and fault tolerance is managed from computing nodes through copies so that it will be able to respond to a fault from database. At the database level with replication is accomplished

the same function. A fault in a node will be detectable at runtime and will not compromise the system dramatically.

VI. CONCLUSION

We have presented the assessment of the information management tools to setting up PInCom, an efficient, available, fault-tolerant and scalable system. Result shows that PInCom is able to deliver a large volume of communications with low latency, so that it can be used by a large number of customers simultaneously. The deployment of multiple computing nodes on an application server allows each node to process different requests in parallel improving overall system speed.

PInCom is fully scalable and system capacity can easily be extended by increasing the number of nodes in the processing cluster, without implying any changes in the code.

Furthermore in system architecture design, has been carried out implementation of non-relational database to manage large volumes of information: The relational databases are not able to handle the size, the speed of delivery nor the diversity of Big Data Formats.

NoSQL databases are the response of scalability and performance that Big Data needs. NoSQL are non-relational, distributed, Open Source and scalable horizontally. This scalability is achieved by increasing the number of nodes also gaining availability. However, not all NoSQL systems have the capabilities required for our communication project.

Before implementing any Big Data technology must study their characteristics and assess whether their qualities correspond with the requirements of the system to be developed.

Also, in the proposed architecture for PInCom it is shown the coexistence between different technologies to obtain the desired configuration. The use of relational database manager is incorporated for the treatment of specific information. Clustered relational database system can solve the problem of inconsistency of own NoSQL system data [14].

NoSQL systems allow management of a large volume of data, but don't offer guarantee of reliability that information requires. Its major disadvantage is working with distributed and sometimes replicated data since it gives priority to the availability and immediacy to supply data. It also raises the possibility of creating different versions and inconsistencies [15].

NoSQL systems hardly reaches the consistency of relational databases due to the presence of integrity constraints, concurrency, concurrent updates, etc. Could be possible not has the latest update of the data, which can be somewhat chaotic.

Specialists remarks that NoSQL is an alternative but it does not answer all the data storage issue. It is necessary to keep complex relationships between structured data, and not all NoSQL systems can support it [16].

In the management and analysis of structured and unstructured data based Big Data technologies, NoSQL systems represent a revolution in how they are stored and the data processing is performed.

As Big Data solutions are growing, use cases where relational systems can be attached to the paradigm of this model for optimization and data analysis are discovered.

The key is to choose the most appropriate solutions and integrate them conveniently for the treatment of information very quickly and efficiently in order to get competitive edge shape.

REFERENCES

[1] L. Columbus, "84% Of Enterprises See Big Data Analytics Changing Their Industries' Competitive Landscapes In The Next Year", 2014. <http://>

www.forbes.com/sites/louiscolombus/2014/10/19/84-of-enterprises-see-big-data-analytics-changing-their-industries-competitive-landscapes-in-the-next-year/#1cb8b3863250

- [2] J. Campo-Ávila, R. Conejo, F. Triguero, R. Morales-Bueno, "Mining Web-based Educational Systems to Predict Student Learning Achievements", *International Journal of Interactive Multimedia and Artificial Intelligence*, Vol III, no. 2, España, 2015, pp. 49.
- [3] MckInsey Global Institute, "Big Data: The next frontier for innovation, competition, and productivity", 2011. http://www.mckinsey.com/insights/business_technology/big_data_the_next_frontier_for_innovation
- [4] EMC2, Martin J., "Aprovechamiento del poder del universo digital", 2014, <http://spain.emc.com/leadership/digital-universe/index.htm>
- [5] P.M Timothy, The NIST Definition of Cloud Computing, NIST, 2011. <http://nvlpubs.nist.gov/nistpubs/Legacy/SP/nistspecialpublication800-145.pdf>
- [6] J. Iglesias, G. Amat, "New Challenges on Crossplatform Digital", *International Journal of Interactive Multimedia and Artificial Intelligence*, Vol III, no. 2, España, 2015, pp 28-29.
- [7] IBM, "The four V's of Big Data", 2014. http://www.ibmbigdatahub.com/sites/default/files/infographic_file/4-Vs-of-big-data.jpg
- [8] G. Magaña, "El big data y la convergencia con los insights sociales", 2015, www.bib.uia.mx/tesis/pdf/016079/016079.pdf
- [9] IBM, "What is Map Reduce. About MapReduce", 2012. <https://www-01.ibm.com/software/data/infosphere/hadoop/mapreduce/>
- [10] The Apache Software Foundation, "Cassandra", 2010. <http://cassandra.apache.org/>
- [11] Software AG, Terracotta Enterprise Suite, 2011. <http://www.terracotta.org/products/enterprise-suite>
- [12] O. Letizi, "Terracotta DSO Data Structures Guide", 2011. <https://confluence.terracotta.org/display/docs/Home>
- [13] Hazelcast, "The Leading Open Source In Memory Data Grid: Distributed computing, simplified", 2012. <http://hazelcast.org/>
- [14] M.Levi, S. Barjaktarovic, "Proceedings of the XIII International Symposium SymOrg 2012: Innovative management and Business Performance", 2012, pp. 979-980.
- [15] S. Mei, "Why You Should Never Use MongoDB", 2013. <http://www.sarahmei.com/blog/2013/11/11/why-you-should-never-use-mongodb/>
- [16] T. Patel, T. Eltaieb, "Relational Database vs NoSQL", *Journal of Multidisciplinary Engineering Science and Technology (JMEST)*, Vol. 2 Issue 4, 2015, pp. 691-695.



Juan Manuel Lombardo, PhD in Computer Science from the Pontifical University of Salamanca, was graduated in Economics and Business Administration in the University of Granada, Spain, Diploma of Advanced Studies (DEA) in Economics from UNED, Research Sufficiency in Business Science from the Complutense University of Madrid and Diploma of Advanced Studies (DEA) in Sociology from the Pontifical University of Salamanca. He is CEO at Fidesol and Professor at Andalusia Business School. Dr. Lombardo is the author of numerous articles and research papers published in journals and books of national and international conferences. Visiting Professor at the Private Technical University of Loja (UTPL Ecuador), The National University of the Northeast (Argentina), University Francisco José de Caldas (Colombia), Catholic University of Colombia, Catholic University of Ibarra (Ecuador), University of Lisbon (Portugal) and National Engineering University (Peru). Member of the Knowledge Management committee of AEC (Spanish Association for Quality) and the Institute CICTES (Ibero-American Centre on Science, Technology and Society) and trustees of the Fundación Descubre for scientific communication.




Miguel Ángel López, has a degree in Engineering in Technical Engineering in Computer Systems from the University of Almería, graduates in Computer Engineering and Master in Softcomputing and Intelligent Systems from the University of Granada. Currently he is CTO. at Fidesol where performs different roles on the projects. His research focuses on distributed systems, management, integration and analysis of data, robotics, fuzzy logic systems, and the development of virtual reality environments for different purposes.



Felipe Mirón is a member of Technology Transference team, and coordinator of international projects at Fidesol. Born in Granada November 1976, he received his degree in Computer Engineering in 2000. Prior to his coalescence to Fidesol, he worked for Public Administration at Región de Murcia, leading a team to change record processing on Public Social Services. He also joined development team as an Analyst at Onda Cero Radio staff in Madrid, and worked as DBA and developer at Credit Suisse in La Florida (Madrid). He has collaborated with Fidesol team in writing several articles published in "Cuaderno de Avances TIC" 2013.



Susana Velasco, has a Technical Engineer in Computer from the University of Granada. In the past, she worked in manufacturing, financial and service sector enterprises as software engineer and analyst programmer. She has also been R&D engineer and QA engineer at Telefónica I+D and Intecna Soluciones. Currently she is a researcher and QA manager of Fidesol. Her research interests include quality assurance, software quality management systems, Ambient Intelligence (AmI) systems and devices, and new generation of ICT technologies.



Juan Pablo Sevilla, is a student of Computer Engineering in the University of Granada. He currently works as a researcher and software architect for Fidesol in Granada. Previously, he worked as a researcher for the University of Granada. His main fields of work are Java technologies, distributed computing and web applications.



Juan Mellado, has a Technician in Computer Systems Administration and Senior Technician in Automatic Control Systems, with CCNA Certified and LPIC-1 Certified. He has worked since 2001 as a system administrator in different companies like Caton Sistemas Alternativos and Intecna Soluciones. Since 2011 works as a systems administrator at Fidesol.