Boosting the use and development of advanced computational applications and resources in Latin America through collaboration and sharing

Esteban Mocskos[1,2], Carlos J. Barrios[3], Harold Castro[4], Dennis Cazar[5], Sergio Nesmachnow[6], and Rafael Mayo-García[7*]
(On behalf of the RICAP Network)

[1] Universidad de Buenos Aires, Facultad de Ciencias Exactas y Naturales, Departamento de Computación, (C1428EGA), Argentina
[2] Centro de Simulación Computacional p/Aplic. Tecnológicas/CSC-CONICET, Godoy Cruz 2390, Buenos Aires (C1425FQD), Argentina.
[3] School of System Engineering, Universidad Industrial de Santander, Bucaramanga, Colombia
[4] Computing and Systems Engineering Department, Universidad de los Andes, Bogota, Colombia
[5] Dpto. de Física, Universidad San Francisco de Quito, Quito, Ecuador
[6] Facultad de Ingeniería, Universidad de la República, Montevideo, Uruguay
[7] CIEMAT, Avda. Complutense, 40, 28040 Madrid, Spain
[*]Corresponding Author

Abstract

Computation consolidates as a highly dynamic discipline impacting in a broad range of scientific domains. In Latin America, computing research and facilities have been steadily developing during the last decades. In spite of this, still there is a remarkable gap in terms of availability of resources and specialized human resources when comparing with other regions or countries. The integration of HPC communities within Latin America and connection with other regions could boost the development of this field. RICAP (Red Iberoamericana de Computación de Altas Prestaciones, i.e. Ibero-American Network for High Performance Computing, http://www.red-ricap.org) aims at filling this gap by means of a strategic and advanced computational infrastructure that includes both High Performance Computing (HPC) and High Throughput Computing (HTC) platforms. This network will also develop different software tools to facilitate access and computational efficiency to the infrastructure, which is being encouraged to be freely used by the Latin American community.

Keywords

Ibero-American Network; HPC; HTC; Supercomputing Infrastructure; Cloud computing.

1. Introduction

RICAP (Red Iberoamerica de Computación de Altas Prestaciones, i.e. Ibero-American Network for High Performance Computing) [1] is a new consortium created to provide the Latin American region with an advanced ICT service. In particular, a strategic infrastructure based on advanced architecture for scientific computing is available, including both High Performance Computing (HPC) and High Throughput Computing (HTC, based on cloud infrastructure)

Fig1. Countries with RICAP presence (Argentina, Brazil, Colombia, Costa Rica, Ecuador, España, Mexico, and Uruguay). The international LAGO consortium joins several Latin AMerican countries (see http://lagoproject.org/ for further detail)

RICAP is funded by the CYTED Programme and was started on January 1, 2017. It will run at least till December 31, 2020. At the kick-off meeting the members of RICAP were:

- Eight Latin American computing centers: CSC-CONICET (Argentina); UFRGS (Brazil); SC3-UIS (Colombia); UCR (Costa Rica); CIEMAT (Spain), BSC-CNS (Spain); and, CUDI and CINVESTAV (Mexico);
- One of the world's largest supercomputing vendors (FUJITSU)
- An experimental Latin American consortium in the field of physics (LAGO).

Throughout 2017, three new institutions have joined the consortium:

- Uniandes (Colombia), CEDIA (Ecuador), and National Supercomputing Center/Universidad de la República (Uruguay)

This network aims to boost the development of different tools to ease access and computational efficiency HPC infrastructure. Another key objective is encouraging its use through different dissemination and outreach actions expected to attract users from different universities and other scientific and industrial areas. RICAP is supported by RedCLARA, which develops and operates the only Latin American advanced Internet network, and by SCALAC, an advanced computing consortium to reach services in Latin America and Caribbean region. RICAP is strongly promoting the integration of new nodes (partners) in the Ibero-American countries that were not initially attached to it as it is demonstrated by the integration of Uniandes, CEDIA and National Supercomputing Center, Universidad de la República.

One the main project objective is improving the availability of HPC resources to the community looking to provide a real alternative to proprietary services located outside the region. Specific objectives include:

- The effective interconnection of high performance open services from the clusters provided by RICAP (both supercomputing and cloud access);
- The implementation and subsequent promotion of solutions for the access and operation of this software-based network;
- The design and development of open source tools that improve the computational efficiency of the infrastructure (especially in an environment such as the cloud) in an unattended and dynamic way;
- The promotion and transfer of knowledge of the RICAP solutions and services through the provision of courses and seminars for administrators and end users with the latest technologies in the field of HPC and HTC;
- The collaboration with other national and regional initiatives (RedCLARA, H2020 and others to be detected).

With the achievement of these objectives it will be possible to carry out new activities by different Latin American groups for whom it was previously extremely difficult to perform large scale research on either data analysis or simulation because of the lack of sufficient computing power. These groups will be able also to collaborate with Spanish research centres such as CIEMAT or BSC-CNS and with leading Latin American groups belonging to RICAP. In this sense, RICAP brings value to the scientific and technological community in several of its fundamental dimensions: Infrastructure and software development; Value added services of advanced academic networks; and ICT solutions to support research.

2. Antecedents

RICAP is a recent collaboration network that is based in efforts made by other previous and present initiatives. As a first seed, it should be mentioned the work carried out by RedCLARA [2] and the HTC projects that have been articulated around it (EELA series [3], CHAIN series [4], MAGIC [5], etc.). This Panamerican association of National Research and Education Networks has provided not only connectivity to the Academia, but direct support to research communities, specific calls to access to infrastructure and many other activities that promoted computer and computational sciences in Central and South America.

Some years ago, a FP7 (i.e. 7 the Framework Programme for Research and Technological Development) collaboration project called RISC, “Red Iberoamericana de Supercomputación” [6] was set up as a network for supporting the coordination of Supercomputing research between Europe and Latin America. It aimed at deepening strategic R&D cooperation between Europe (EU) and Latin America (LA) in the field of High Performance Computing (HPC) by building a multinational and multi-stakeholder community that involved a significant representation of the relevant HPC R&D EU and LA actors (researchers, policy makers, users). RISC worked on identifying common needs, research issues, and opportunities for cooperative R&D on HPC between EU and LA in the transition to multi-core architectures across the computing spectrum and relevant programming paradigms, algorithms and modeling approaches, thus setting the basis for the formulation of a global strategy for future research.

Another very important initiative that is still on the way is the Servicio de Cómputo Avanzado para Latinoamérica y el Caribe (SCALAC). Similarly to RICAP, SCALAC is a consortium of centers in several Latin American countries dedicated to the development of High Performance Computing and scientific computing in Latin America. These centers have been organized, with

the support of RedCLARA, to provide jointly and collaboratively advanced computing services in the region. SCALAC members are centers that have advanced computing resources, which belong to universities and scientific research institutions, and are supported by the advanced research and education networks of their countries. The formal launch of the SCALAC community was the 1st. March 2013 at the Industrial University of Santander (UIS), in the city of Bucaramanga, Colombia.

Jointly with projects, there is an effort in the region to promote the HPC activities with international conferences that join international researchers to present their activities and latest developments. Not only this, networking activities being held are the seed of future collaborations. In the HPC field, two major conferences were organized in Latin America up to 2014: CLCAR (most countries from Mesoamérica) and HPCLATAM (most countries from South América). These two conferences merged in 2014 in a single major event that gathers the whole HPC/distributed computing research community: the current Latin America High Performance Computing Conference (CARLA). The 2017 edition of the conference was held in Buenos Aires in September 2017 [7], having more than 100 participants from more than 25 countries from America, Europe, Asia, and Oceania. The Latin American HPC community also supports other initiatives; it is also worth mentioning the recent organization of a flagship conference on cluster, grid and cloud services (16th Annual IEEE/ACM International Symposium on Cluster, Cloud, and Grid Computing, CCGRID), in 2016 in Cartagena, Colombia.

Another important topic in any scientific discipline is training. Some information regarding this topic can be found below in subsection 4.2.

3. Plans to provide a real added value to the Ibero American community

In this section, information about the infrastructure that is freely available by RICAP and the methodology followed for accessing it will be detailed.

3.1 Accessing RICAP resources

Due to the nature of this network, several aspects should be defined: different methodologies related to access to the strategic infrastructure provided by RICAP; development of new solutions to improve the infrastructures exploitation; and, actions related to transfer of knowledge, and dissemination and outreach.

With respect to the first, the proposed methodology is similar to that followed by large computing infrastructures such as PRACE [8] in HPC or FedCloud [9] in HTC, in which different nodes hosting computing clusters and computing capacities are somehow federated. However, in RICAP this federation will be as simple as possible in order to facilitate easy access and administration.

Access to the strategic infrastructure of RICAP will be carried out in two ways. On the one hand, there will be online open calls for proposals for the use of supercomputers from which CPU and/or accelerator hours (GPU and Xeon Phi) will be granted by a committee designated by RICAP. In this call, storage and transfer capacities of data will be granted as well to the final user (from any Ibero-American country, not only those that have a partner within RICAP).

On the other hand, RICAP will enable a cloud infrastructure especially suitable for HTC. Its access will be via either command line or a user-friendly web interface. The availability of the cloud infrastructure will be continuous and uninterrupted over time.

This scenario forces to have open source solutions that improve computational efficiency in heterogeneous and dynamic environments. Therefore, RICAP will develop and provide tools that maximize this computational efficiency from those that have already been previously successfully put into production [10-12].

In relation to the HPC environment, fault-tolerant developments that will equally improve the use of supercomputers will also be used and have also been conveniently tested by the RICAP groups. These solutions are based on managers and tools mostly used in the Top500 such as resource managers, checkpoint and load balancing libraries, message passing libraries, etc. They will be transparent to the end user and are focused on their exploitation by the administrators. An exploratory work can be found in [13].

All of these developments will converge into tools that combine dynamic checkpointing and fault tolerance capabilities with cloud job resource managers over CPU and accelerators. Also, container management in virtualized HPC environments will be studied.

Both works on HPC environments and those related to the cloud will produce user statistics and usage modes, as well as a measurement of computation times, tools used, and success stories. This information will be useful not only in the design and development of future administrative applications, but also in the improvement of existing ones in each RICAP center.

3.2 Infrastructure

The strategic HPC infrastructure of RICAP, which is connected internally by Infiniband and outside with fiber optics by the corresponding academic networks associated with RedCLARA, is composed of:

- BSC (Spain): A general purpose cluster composed of 165,488 Intel Platinium cores in 3,456 nodes, with more than 394TB of main memory and 25PB of storage.
- CIEMAT (Spain): 1 cluster with 680 Intel Gold processors and 456 cores Xeon Phi, 1 cluster of ~100,000 Nvidia cores, two cloud nodes with ~950 cores CPU and more than 1 PB of storage.
- CEDIA (Ecuador): 12 computing nodes with 322 Intel Xeon cores, 1TB RAM & 6TB of storage. Also, 5760 Nvidia cores.
- CINVESTAV (México): SGI ICE-XA (CPU) and SGI ICE-X (GPU) with 8,900 cores and a peak performance of 429 Tflops. Storage Type Chandelier Seagate ClusterStor 9000 of 1PB
- CSC-CONICET (Argentina): An AMD Opteron cluster of 4,096 cores and 16,384 Nvidia cores. It has 8,192 GB of RAM and a storage space of 72TB.
- CUDI (México): CUDI cloud computing service
- UCR (Costa Rica): Multiple clusters with capacity of ~80 cores CPU, ~25,000 Nvidia cores, and ~1,450 Xeon Phi cores
- UFRGS (Brasil): A 256-node cluster with 19,968 CUDA cores
- UIS (Colombia): A cluster of 24 nodes (2.4GHz and 16GB RAM) for training activities and a cluster with 128 NVIDIA FERMI Tesla (104GB in RAM and 4 Intel Haswell Processors by node)
- Uniandes (Colombia): One cluster with 1,808 cores with HT (8 TB RAM) jointly with 160 TB of storage.
- National Supercomputing Center (Uruguay): Cluster-UY infrastructure; 29 nodes with 576 cores and 1.28 TB of RAM, and 128 Xeon Phi cores with 16GB of RAM.

This freely accessible infrastructure already in production will be federated thanks to RICAP and ensures the achievement of the network's objectives. Up until now, access to major supercomputers was limited to only a few countries and groups in the region, but RICAP is expected to change this situation and leverage the access to this valuable resources.

It should be borne in mind that RICAP is made up of companies and resource providers (the groups listed above), but also direct suppliers of use cases (BSC, CIEMAT, CINVESTAV, CSC-CONICET and its associated institutions, LAGO and UCR) that will carry out the initial tests of the strategic infrastructure and that will later attend the calls for accessing HPC resources or directly use the cloud of the Network. All resource providers also have extensive experience in the tasks related to RICAP both on the exploitation and federation of computer infrastructures and on R&D work and, together with the rest of institutions, form a balanced consortium between consolidated and emerging groups. Also, they have a wide baggage of participation in 7th Framework Programme and H22020 projects in the sector of the e-Infrastructures.

The aforementioned infrastructure is complemented by the availability of a desktop cloud that can be used as training or development platform. This platform allows the creation of virtual clusters reproducing any software environment and making possible to deploy such environments on desktop computers, even while there are end-users (i.e. students) running their applications. In this way, institutions with no HPC facilities may work with the same software stack used by RICAP providers and they can teach or test applications and, in some cases, get the results they are looking for.

3.3 Work plan

Below is a brief list of the activities and milestones that are expected to be achieved over the period 2017-2020 and which are expected to be of interest to the different user groups that are working in Ibero-American universities and other scientific and technological entities mainly (but not exclusively).

2017
- Federation and production of the strategic cloud infrastructure.
- Design of the 1st call for access to HPC infrastructure and implementation of the web form. Associated dissemination and outreach of the call. Awarded computing resources and use of the infrastructure by the selected groups.
- Integration into the infrastructure of the new efficiency solutions provided by RICAP
- General RICAP dissemination and outreach activities

2018
- Update and operation of the HPC and cloud infrastructures
- Analysis and debugging of possible errors derived from the 1st HPC call.
- Analysis of the results obtained by the new solutions
- Completion of the second HPC call.
- Implementation of new solutions that combine the use of tools in the cloud with accelerators
- General RICAP dissemination and outreach activities

2019
- Update and operation of the HPC and cloud infrastructures
- Analysis of the results obtained by the new solutions
- Completion of the third and fourth HPC calls.

- Implementation of new solutions that combine the use of tools in the cloud with fault tolerance techniques
- General RICAP dissemination and outreach activities

2020
- Update and operation of the HPC and cloud infrastructures
- Analysis of the results obtained by the new solutions
- Completion of the fifth and sixth HPC calls.
- Implementation of new solutions devoted to computational efficiency with containers
- General RICAP dissemination and outreach activities

It should be noted that the network will be managed through an Executive Committee that will meet remotely once a month and that will be formed by a representative of each group of RICAP. In addition, an annual face-to-face meeting will be held. This Executive Committee will appoint experts to evaluate applications for access to HPC resources.

The first RICAP call for requesting computing resources was launched on June 2017.

4. Expected results

During the initial four years of existence of RICAP, it is expected to be able to reach different milestones and results that improve different scientific-technological and social aspects.

4.1 Benefits and expected impact for Latin America

Strategic infrastructures such as PRACE and EGI (FedCloud) in Europe or the Spanish Supercomputing Network in Spain have had enormous success and a notable impact on the scientific community in their regions of influence, making it possible to access large computing facilities to any researcher or group with a need for computational services as long as it presents a project of quality, interest, and technologically feasible within the capabilities offered.

For this reason, it can be assured that the benefits and impact that RICAP is expected to have in Latin America should not be less. On the contrary, in social terms it will be even greater because it will make available to researchers and groups a computing capacity that in their countries can be practically impossible to materialize. Some of the countries that participate in RICAP have world-class supercomputing facilities, but some others in Latin America don't, so thanks to the possibility offered by RICAP they researchers will be able to enjoy this computing capacity for free.

Moreover, RICAP is open to all scientific, technological, and social fields and it will be able to attend on equal terms users of any Ibero-American country. Therefore, it offers a clear commitment to social equity. Also, as indicated above, a mechanism will be articulated so that more Latin American institutions can be integrated into the activities of the Thematic Network.

However, some users have already been pre-identified. They will be able to use the RICAP capabilities from the outset. Based on the experience of some of its founding partners, such as BSC, CIEMAT, CINVESTAV, CSC-CONICET (and its associated institutions such as CNEA or UBA), and LAGO, tests will be carried out in areas such as Wind Energy, Medical Physics, Nuclear and radiological Safety, Plasmas Physics, Geophysics, Air Pollution, Radar Network Design and Wireless Communication, etc.

Also, it is noteworthy the participation in the Network of the Latin American consortium LAGO [14], which has a Working Group specialized in simulation that ensures the use of the HTC network for studies of cosmic rays and its effect on the health of crews and air travelers. LAGO will also apply to access the HPC resources under competitive concurrence.

Thanks to the capacities and solutions provided by RICAP, all these communities will have the possibility to approach the solution of more ambitious and complex problems, which could not be raised by some Latin American researchers without this Network.

4.2 Training

As aforementioned, training is cornerstone in RICAP. The Network counts among its members with the BSC, one of the PRACE Tier-0 partners. Within the activities of this great European consortium, there is a very wide variety of courses organized by its teams, which can be extended to Latin America with funds from the PRACE consortium itself. The 1-day or 2-day theoretical-practical seminars that emerge from this collaboration with PRACE will serve as training for staff developing their work in the HPC / HTC field, a fact that will increase their employability.

In addition to these PRACE courses, there will also be seminars within the academic program directly focused on the use of the strategic network provided by RICAP. These tutorials will be organized in conjunction with events and conferences held in Latin America such as TICAL [15], CARLA, ISUM [16], etc.

For these conferences and others, a strong collaboration with RedCLARA has been set. Thus, there is an effective integration into its program of Thematic Encounters of the TICAL community streaming, so the impact of the training RICAP activities will be even greater and thus people who are not attending in person those courses can follow them locally.

Examples of courses that will be taught within the activities of RICAP are:
- Access and use of computational resources provided by RICAP
- Understanding High Performance Computing Cluster Management
- Hands-on Introduction to HPC (PRACE)
- Message-Passing Programming with MPI (PRACE)
- The Intel MIC and GPU programming (PRACE)
- Methodologies for the efficient execution of tasks in HPC and HTC environments
- Performance analysis and tools (PRACE)
- HPC-CSC Intensive School
- SC-CAMP Supercomputing and Distributed Systems Camping School (SCCAMP Consortium)

As can be seen from the list above, this training plan is designed to encompass and be profitable for different academic levels.

Another important point to note is that the courses and seminars that are taught within RICAP will be promoted by the academic entities of the Network within its Degree and Master courses so that the students can attend them, not only in the Faculties of Informatics, but also in others where the knowledge of the Network (Science, Engineering, Statistics, etc.) is applicable. LAGO also has its own human resources training program in which the teaching activities of RICAP will be promoted.

In order to complement training and dissemination activities, several events are regularly organised in Argentina, Uruguay, Brasil, Costa, Rica, Mexico, Colombia, and other Latin American countries. The High Performance Computing School (ECAR) [17] and the International SuperComputing Camp [18] are two examples about it. With respect to networking activities highlighting HPC in Latin America, it is worth mentioning WHPCEuroLatam, which is periodically allocated as part of ISC-HPC [19], or the recent Americas HPC Collaboration BoF held as part of Supercomputing'17 [20].

All the didactic material generated (presentations, exercises, videos in which the course is recorded) will be posted on the web page of the Network and will be available permanently for free.

5. Conclusions

Thanks to RICAP, it will be possible to generate a multitude of new activities in different scientific and social areas of Latin American society, since a huge amount of computational resources will be made available for free. This is so because today, computing is hugely integrated in the generation of knowledge in scientific, social, and engineering environments, in the public as well as the private sectors. In addition, access to the infrastructure will be direct (cloud) or competitive concurrency (HPC) based, in which scientific interest will prevail, so that end users without local access to HPC and HTC environments will from now on can carry out their work thanks to RICAP, which will improve social equality.

The originality of the new Network is that this regional access is not a reality in Latin America as a whole, but only at a national level in some countries, which is a tremendous advance for the scientific community in any area in the region.

Acknowledgements

This work has been partially funded by the Thematic Network 517RT0529 RICAP (Ibero-American Network for High Performance Computing) of the Ibero-American Program of Science and Technology for Development (CYTED).

References

[1] The RICAP Network, available at http://www.red-ricap.org
[2] The RedCLARA association, available at http://www.redclara.net/
[3] The EELA series of projects, available at https://www.gisela-grid.eu/
[4] The CHAIN series of projects, available at http://www.chain-project.eu/
[5] The MAGIC project, available at http://magic-project.eu/
[6] The RISC project, available at http://www.risc-project.eu/
[7] The CARLA conference, available at http://carla2017.ccarla.org/
[8] The PRACE project, available at http://www.prace-ri.eu/
[9] The FedCloud initiative, available at https://www.egi.eu/infrastructure/cloud/
[10] M. Rodríguez-Pascual, I.M. Llorente, R. Mayo-García. "Montera: a framework for efficient execution of Monte Carlo codes on Grid infrastructures". Computing and Informatics 32, 113-144 (2013)
[11] A.J. Rubio-Montero, E. Huedo, F. Castejón, R. Mayo-García. "GWpilot: Enabling multi-level scheduling in distributed infrastructures with GridWay and pilot jobs". Future Generation Computer Systems 45, 25-52 (2015).

[12] A.J. Rubio-Montero, E. Huedo, R. Mayo-García. “Scheduling multiple virtual environments in cloud federations for distributed calculations”. Future Generation Computer Systems 74, 90-103 (2017)
[13] A.J. Rubio-Montero, M. Rodríguez-Pascual, R. Mayo-García. “A simple model to exploit reliable algorithms in cloud federations”. Soft Computing 21, 4543-4555 (2017).
[14] H. Asorey et al. “The Latin American Giant Observatory: a successful collaboration in Latin America based on Cosmic Rays and computer science domains” 16th IEEE/ACM International Symposium on Cluster, Cloud, and Grid Computing, pp. 707–711 (2016).
[15] The TICAL conference, available at http://tical.redclara.net/
[16] The ISUM conference, available at http://www.isum.mx/
[17] The ECAR school, available at http://ecar2017.hpclatam.org/agenda/
[18] The SC-Camp school, available at http://www.sc-camp.org/2017/
[19] The ISC-HPC conference, available at http://isc-hpc.com/
[20] The Americas HPC Birds-of-Feather, available at http://www.sc3.uis.edu.co/bofamericas-hpc-collaboration/