Boosting Advanced Computational Applications and Resources in Latin America through Collaboration and Sharing

Esteban Mocskos
Universidad de Buenos Aires, CONICET

Carlos J. Barrios H.
Universidad Industrial de Santander

Harold Castro
Universidad de los Andes

Dennis Cazar Ramírez
Universidad San Francisco de Quito

Sergio Nesmachnow
Universidad de la República

Rafael Mayo-García
CIEMAT, Harvard University

Although computing research and facilities in Latin America have been developing steadily, a remarkable gap nevertheless remains in the availability of resources and specialized human resources compared to other regions. RICAP (Red Iberoamericana de Computación de Altas Prestaciones, or Ibero-American Network for High-Performance Computing) aims to fill this gap by means of a strategic and advanced computational infrastructure that includes both high-performance and high-throughput computing platforms. This network will also develop software tools to facilitate this infrastructure's accessibility and computational efficiency to encourage use by the Latin American computing community as well as focus on strengthening and training human resources in both parallel programming techniques and large-scale computing platform operations.

Copublished by the IEEE CS and the AIP
1521-9615/18/$33 ©2018 IEEE

Red Iberoamerica de Computación de Altas Prestaciones (Ibero-American Network for High-Performance Computing)— known as RICAP—is a new consortium that was created to provide Latin America with advanced ICT services.[1] In particular, a strategic cloud-based infrastructure based on advanced architecture for scientific computing has been developed that includes both high-performance computing (HPC) and high-throughput computing (HTC).

RICAP—which is funded by the Ibero-American Programme on Science and Technology for Development, known as CYTED—began on 1 January 2017 and will run until at least 31 December 2020. At the kick-off meeting, RICAP included

- 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); and
- an experimental Latin American consortium in the field of physics (LAGO).

Three additional institutions joined the consortium in 2017: Uniandes (Colombia), CEDIA (Ecuador), and the National Supercomputing Center/Universidad de la República (Uruguay).

Figure 1. Countries with a RICAP presence: Argentina, Brazil, Colombia, Costa Rica, Ecuador, Spain, Mexico, and Uruguay. The international LAGO consortium joins several Latin American countries (see http://lagoproject.org/ for further details).

This network aims to boost the development of different tools to ease access and improve computational efficiency of HPC infrastructure. Another key objective is to encourage infrastructure use through dissemination and outreach actions that are expected to attract users from various 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 extend services throughout Latin America and the Caribbean. RICAP is strongly promoting the integration of new nodes (partners) in the Ibero-American countries that were not initially attached to it, as demonstrated by its integration of Uniandes, CEDIA, National Supercomputing Center, and Universidad de la República.

The project seeks to make HPC resources more readily available to the community, thereby providing a real alternative to proprietary services located outside the region. Specific objectives include the following:

- effective interconnection of high-performance open services from the clusters provided by RICAP (both supercomputing and cloud access);
- implementation and subsequent promotion of solutions for the access and operation of this software-based network;
- 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;
- promotion and transfer of knowledge of RICAP solutions and services by way of courses and seminars for administrators and end users on the latest technologies in HPC and HTC; and
- collaboration with other national and regional initiatives (RedCLARA, H2020, and others to be identified).

Achieving these objectives will allow 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—to carry out new activities. These groups will also be able to collaborate with Spanish research centers such as CIEMAT or BSC-CNS as well as 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 for advanced academic networks; and ICT solutions that support research.

ANTECEDENTS

RICAP is a recent collaborative network based on the efforts of other previous and present initiatives. As the first seeds, the work by RedCLARA[2] and the HTC projects articulated around it (for example, the EELA series,[3] CHAIN series,[4] and MAGIC[5]) should be mentioned. The 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 have promoted computer and computational sciences in Central and South America.

Some years ago, an FP7 (7th Framework Programme for Research and Technological Development) collaboration project called Red Iberoamericana de Supercomputación (RISC)[6] was established as a network to support the coordination of supercomputing research between the EU and Latin America. It aimed to deepen strategic R&D cooperation between the two regions in the field of HPC by building a multinational and multistakeholder community that involved significant representation of relevant researchers, policy makers, and users). RISC worked to identify common needs, research issues, and opportunities in the transition to multicore 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 (Advanced Computing Services for Latin America and the Caribbean), or SCALAC. Similar to RICAP, SCALAC is a consortium of centers in several Latin American countries dedicated to the development of HPC and scientific computing in Latin America. These centers have been organized, with the support of RedCLARA, to jointly and collaboratively provide advanced computing services in the region. SCALAC members are centers that have advanced computing resources, 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 on 1 March 2013 at the Universidad Industrial de Santander in Bucaramanga, Colombia.

Jointly with these projects, there is an effort in the region to promote HPC through international conferences that unite researchers to present their activities and latest developments. Such networking activities are the seeds of future collaborations. In the HPC field, two major conferences were organized in Latin America up to 2014: CLCAR and HPCLATAM. In 2014, these two conferences merged into a single major event that gathers the entire HPC/distributed computing research community: the Latin America High Performance Computing Conference (CARLA). The most recent CARLA was held in Buenos Aires in September 2017,[7] with more than 100 participants from more than 25 countries representing America, Europe, Asia, and Oceania. It is also worth mentioning the recent organization of a flagship conference on cluster, grid, and cloud services—the 16th Annual IEEE/ACM International Symposium on Cluster, Cloud, and Grid Computing—held in 2016 in Cartagena, Colombia.

Another important topic in any scientific discipline is training, which will be discussed later in this article.

PLANS TO PROVIDE REAL ADDED VALUE TO THE IBERO-AMERICAN COMMUNITY

This section provides information about the infrastructure being made available for free by RICAP and the methodology for accessing it.

Accessing RICAP Resources

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 federated. However, in RICAP, this federation will be kept as simple as possible to facilitate easy access and administration.

Access to the strategic infrastructure of RICAP will be carried out in two ways. First, 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 also be granted to the final user (from any Ibero-American country, not only those that have a partner within RICAP). Second, RICAP will enable a cloud infrastructure especially suitable for HTC. It will be accessible via either command line or a user-friendly web interface. The availability of the cloud infrastructure will be continuous and uninterrupted.

This scenario requires open-source solutions that improve computational efficiency in heterogeneous and dynamic environments. Therefore, RICAP will develop and provide tools that maximize this computational efficiency beyond what has already been 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 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, and message-passing libraries. They will be transparent to the end user and exploitable by administrators. An exploratory work can be found in "A Simple Model to Exploit Reliable Algorithms in Cloud Federations."[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. Container management in virtualized HPC environments will also be studied.

Work in HPC environments and 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.

Infrastructure

Table 1 describes the strategic HPC infrastructure of RICAP, which is connected internally by Infiniband and externally with fiber optics by the corresponding academic networks associated with RedCLARA.

Table 1. RICAP's HPC infrastructure.

Academic network	Country	Description of cluster
BSC	Spain	General-purpose cluster with 165,488 Intel Platinum cores in 3,456 nodes, with more than 394 Tbytes of main memory and 25 Pbytes of storage
CIEMAT	Spain	One cluster with 680 Intel Gold processors and 456 cores Xeon Phi, one cluster of ~100,000 Nvidia cores, two cloud nodes with ~950 cores CPU, and more than 1 Pbyte of storage
CEDIA	Ecuador	12 computing nodes with 322 Intel Xeon cores, 1 Tbyte RAM, 6 Tbytes of storage, and 5,760 Nvidia cores
CINVESTAV	México	SGI ICE-XA (CPU) and SGI ICE-X (GPU) with 8,900 cores and peak performance of 429 Tflops. Storage Type Chandelier Seagate ClusterStor 9000 of 1 Pbytes
CSC-CONICET	Argentina	AMD Opteron cluster of 4,096 cores and 16,384 Nvidia cores, 8,192 Gbytes of RAM, and storage space of 72 Tbytes
CUDI	México	CUDI cloud computing service
UBA	Argentina	One cluster with 30 nodes with two AMD 6320, 64 Gbytes of RAM, 24 Tesla K20c, 2 heavy nodes with extra memory (512 Gbytes of RAM each), and 7 nodes with two Intel E5-2630 v4. Another infrastructure consists of 4 nodes with two AMD Opteron 6320 configured to support cloud computing.
UCR	Costa Rica	Multiple clusters with capacity of ~80 cores CPU, ~25,000 Nvidia cores, and ~1,450 Xeon Phi cores
UFRGS	Brazil	256-node cluster with 19,968 CUDA cores
UIS	Colombia	Cluster of 24 nodes (2.4 GHz and 16 Gbytes RAM) for training activities, and cluster with 128 Nvidia Fermi Tesla (104 Gbytes of RAM and 4 Intel Haswell processors per node)
Uniandes	Colombia	One cluster with 1,808 cores with HT (8 Tbytes RAM) jointly with 160 Tbytes storage
National Supercomputing Center	Uruguay	Cluster-UY infrastructure; 29 nodes with 576 cores and 1.28 Tbytes of RAM, and 128 Xeon Phi cores with 16 Gbytes of RAM

This freely accessible infrastructure already in production will be federated thanks to RICAP and will ensure 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 access to this valuable resource.

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 participate in calls for accessing HPC resources or directly use the network cloud. All resource providers 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 other institutions form a balanced consortium between consolidated and emerging groups. Also, they have a wide baggage of participation in FP7 and H22020 projects in the e-Infrastructures sector.

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 that can reproduce any software environment and make it possible to deploy such environments on desktop computers, even while there are end users (students) running their applications. In this way, institutions with no HPC facilities might work with the same software stack used by RICAP providers and be able to teach or test applications and, in some cases, get the results they are looking for.

Work Plan

Below is a brief list of the activities and milestones that have already been or are expected to be achieved during the 2017–2020 period and are of interest to the different user groups in Ibero-American universities and other scientific and technological entities.

2017:

- Federation and production of the strategic cloud infrastructure
- Design of the first call for access to HPC infrastructure and implementation of the web form, associated dissemination and outreach of the call, award of 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 first HPC call
- Analysis of the results obtained by the new solutions
- Completion of the second HPC call
- Implementation of new solutions combining 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 combining 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

Note that the network will be managed through an executive committee formed by a representative of each RICAP group. The committee will meet remotely once a month as well as face-to-face once a year. This executive committee will appoint experts to evaluate applications for access to HPC resources.

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

EXPECTED RESULTS

During the initial four years of RICAP, it is expected to achieve multiple milestones and results that will improve various scientific-technological and social aspects.

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, granting access to large computing facilities to any researcher or group with a need for computational services as long as it presents a project of quality and interest and is technologically feasible within the capabilities offered.

For this reason, it can be assured that RICAP will have at least as much benefit and impact in Latin America. And in social terms, its impact 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 do not; therefore, RICAP offers these researchers the opportunity to enjoy this computing capacity for free.

Moreover, RICAP is open to all scientific, technological, and social fields and will be able to serve—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.

Some users have already been preidentified and will be able to use the RICAP capabilities from the outset. Based on the experience 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.

Also of note is the participation of LAGO,[14] which has a working group that specializes in simulation that ensures use of the HTC network for studies of cosmic rays and their effect on the health of air crews and 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 chance to approach the solution of ambitious and complex problems that could not otherwise be attempted by many Latin American researchers.

Training

Training is a cornerstone of RICAP. The network counts among its members the BSC, one of the PRACE Tier-0 partners. Within the activities of this great European consortium, a wide variety of courses is organized by its teams, which can be extended to Latin America with funds from the PRACE consortium itself. The one- or two-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, and ISUM.[16]

For these conferences and others, a strong collaboration with RedCLARA has been established. Thus, TICAL has effectively integrated community streaming into its Thematic Encounters program, so people not in attendance can still follow the courses locally.

Courses to be taught as part of RICAP include the following:

- 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),
- Intel MIC and GPU programming (PRACE),
- methodologies for efficient execution of tasks in HPC and HTC environments,
- performance analysis and tools (PRACE),
- an HPC-CSC intensive school, and
- an SC-CAMP supercomputing and distributed systems camping school (SCCAMP Consortium).

As the list shows, 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 taught within RICAP will be promoted by the academic entities of the network within its degree and master's courses to encourage attendance by not only computer science students but also students in other departments (science, engineering, statistics, and so on) to whom knowledge of the network is applicable. LAGO also has its own human resources training program in which the teaching activities of RICAP will be promoted.

To complement training and dissemination activities, several events are regularly organized in Argentina, Uruguay, Brazil, 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. With respect to networking activities highlighting HPC in Latin America, WHPCEuroLatam is periodically allocated as part of ISC-HPC,[19] and an Americas HPC Collaboration BoF was recently held as part of Supercomputing 2017.[20]

All didactic materials generated (presentations, exercises, videos) will be posted on the network webpage and made permanently available for free on the RICAP website (http://www.red-ricap.org).

CONCLUSIONS

By making a huge amount of computational resources freely available, RICAP will generate a multitude of new activities in different scientific and social areas of Latin American society. This is particularly significant today, with computing so hugely integrated into the generation of scientific, social, and engineering knowledge in both the public and private sectors. It will also contribute toward improved social equality because access to the infrastructure will be direct (cloud) or competitive concurrency (HPC) based, meaning that scientific interest will prevail and end users without local access to HPC and HTC environments will be able to carry out their work.

The originality of the network lies in the fact that regional access is not a reality in Latin America as a whole, but only at a national level in some countries. Therefore, RICAP represents a tremendous advance for scientific communities across all areas in the region.

ACKNOWLEDGMENTS

This work was 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*, 2018; http://www.red-ricap.org.
2. *The RedCLARA Association*; http://www.redclara.net/.
3. *EELA: E-Infrastructure Shared between Europe and Latin America*; https://www.eu-eela.eu/first-phase.php.
4. *CHAIN Series of Projects*; http://www.chain-project.eu.
5. *MAGIC Project*; http://magic-project.eu/.
6. *RISC Project*; http://www.risc-project.eu/.
7. *CARLA 2017: Latin America High Performance Computing Conference*; http://carla2017.ccarla.org/.
8. *PRACE: Partnership for Advanced Computing in Europe*; http://www.prace-ri.eu/.
9. "Cloud Compute," *FedCloud Initiative*; https://www.egi.eu/infrastructure/cloud/.
10. M. Rodríguez-Pascual, I.M. Llorente, and R. Mayo-García, "Montera: A Framework for Efficient Execution of Monte Carlo Codes on Grid Infrastructures," *Computing and Informatics*, vol. 32, 2013, pp. 113–144.
11. A.J. Rubio-Montero et al., "GWpilot: Enabling Multi-Level Scheduling in Distributed Infrastructures with Gridway and Pilot Jobs," *Future Generation Computer Systems*, vol. 45, 2015, pp. 25–52.
12. A.J. Rubio-Montero, E. Huedo, and R. Mayo-García, "Scheduling Multiple Virtual Environments in Cloud Federations for Distributed Calculations," *Future Generation Computer Systems*, vol. 74, 2017, pp. 90–103.
13. A.J. Rubio-Montero, M. Rodríguez-Pascual, and R. Mayo-García, "A Simple Model to Exploit Reliable Algorithms in Cloud Federations," *Soft Computing*, vol. 21, 2017, pp. 4543–4555.
14. H. Asorey et al., "The Latin American Giant Observatory: A Successful Collaboration in Latin America based on Cosmic Rays and Computer Science Domains," *Proc. 16th IEEE/ACM Int'l Symp. Cluster, Cloud, and Grid Computing*, 2016, pp. 707–711.
15. *TICAL Conference*; http://tical.redclara.net.
16. *ISUM 2018: 9th Int'l Supercomputing Conf. in Mexico*; http://www.isum.mx/.
17. *ECAR 2017: High Performance Computing Camp*; http://ecar2017.hpclatam.org/agenda.
18. *SC-Camp 2017: 8th Int'l Supercomputing Camp*; http://www.sc-camp.org/2017/.
19. *ISC-HPC 2018 Conference*; http://isc-hpc.com.
20. *The Americas HPC Birds-of-Feather*; http://www.sc3.uis.edu.co/bofamericas-hpc-collaboration/.

ABOUT THE AUTHORS

Esteban Mocskos is a professor of computer science at the Universidad de Buenos Aires, and researcher at CONICET (Centro de Simulación Computación de Aplicaciones Tecnológicas, or Computational Simulation Center for Technological Applications). His research interests include distributed systems, parallel programming and architectures, modeling, and simulation. Mocskos received a PhD in computer science from the Universidad de Buenos Aires. Contact him at emocskos@dc.uba.ar.

Carlos J. Barrios H. is director of the High Performance and Scientific Computing Center and an assistant professor at the Universidad Industrial de Santander. His research interests include HPC and hybrid/high-performance architectures for science and ultrascale systems

involving design, performance evaluation and implementation mechanisms. Barrios received a PhD in computer science from Nice–Côte d'Azur University. Contact him at cbarrios@uis.edu.co.

Harold Castro is an associate professor of ICT at the Universidad de los Andes. His research interests include cloud computing, high-performance computing (HPC), and big data. Castro received a PhD in computer science from the Institut National Polytechnique de Gernoble. Contact him at hcastro@uniandes.edu.co.

Dennis Cazar Ramírez is a professor of physics at the Universidad San Francisco de Quito. His research interests focus on the development of methods and instrumentation for astroparticle detection. Cazar received a PhD in electronics and information technology from the Universitá degli Studi di Pavia. Contact him at dcazar@usfq.edu.ec.

Sergio Nesmachnow is a professor and researcher at the Universidad de la República, PEDECIBA, and ANII. His research interests include HPC, scientific computing, metaheuristics and smart cities. He received a PhD in computer science from the Universidad de la República. Contact him at sergion@fing.edu.uy.

Rafael Mayo-García is a scientific researcher at the Centro de Investigaciones Energéticas Medioambientales y Tecnológicas (CIEMAT) and a Fellow at Harvard University. His research interests include HPC, scientific computing, and computational physics. He received a PhD in physics from the Universidad Complutense de Madrid. Contact him at rafael.mayo@ciemat.es.