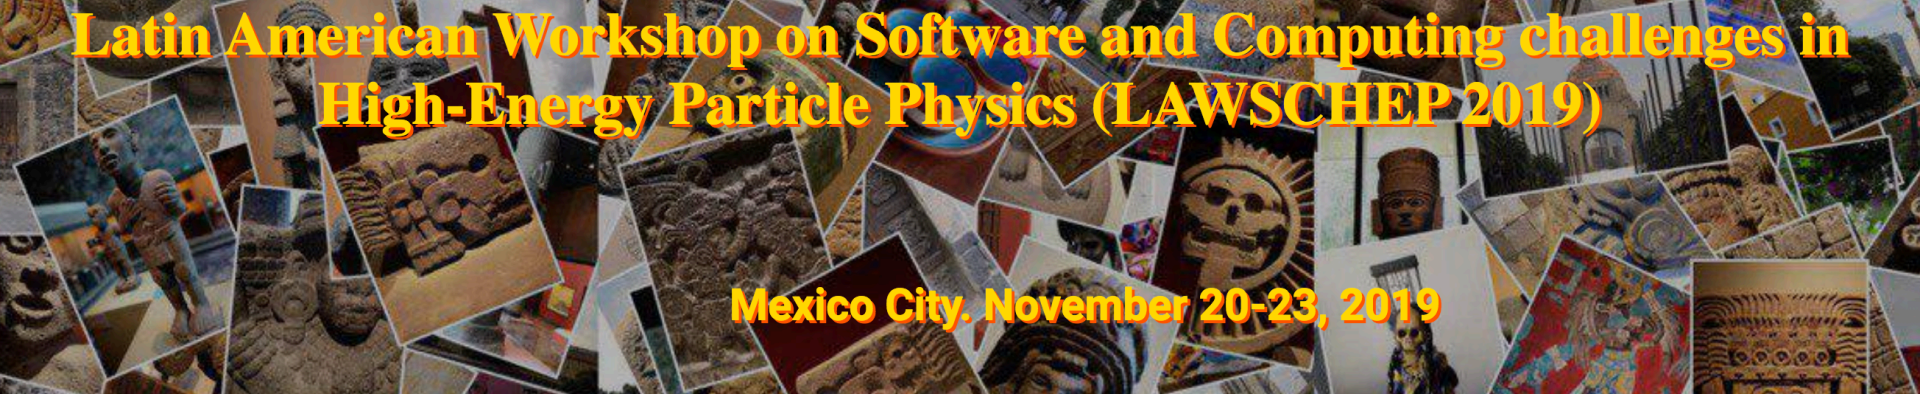

HEP Computing Infrastructures in L.A. (and Brazil in particular)

ROGÉRIO L. IOPE

São Paulo State University - Center for Scientific Computing

HEPGrid Project (Brazil)

UNIVERSIDADE DO ESTADO DO RIO DE JANEIRO (UERJ)

HEPGrid Project at UERJ
CMS Tier-2 (T2_BR_UERJ)

Research team - CMS members and users

Alberto Santoro (Principal Investigator)
Antonio Vilela Pereira
André Sznajder
Carley Martins
Clemencia Mora Herrera
Dilson de Jesus Damião
Helena Malbouisson
Hélio Nogima
Luiz Mundim
Sandro Fonseca
Vitor Oguri
Wagner Carvalho
Wanda Prado

Postdoctoral Researchers

Eliza Melo
Patrícia Rebello Teles
Sheila Amaral
Walter Aldá

Computer Systems Analysts and Engineers

Ana Beatriz Franco
Caio Costa
Douglas Milanez Marques
Eduardo Azevedo Revoredo
Renan Bernardo Valadão

HEPGrid - Hardware Infrastructure

- Five racks with 70 servers - 700 cores
  - Dell servers models:
    - PowerEdge 2950 (38)
    - PowerEdge R510 (12)
    - PowerEdge R710 (7)
    - PowerEdge R720xd (11)
    - PowerEdge R730 (2)
  - Processors:
    - Intel Xeon E5410 – 2,3GHz
    - Intel Xeon E5530 – 2,4GHz
    - Intel Xeon X5650 – 2,66GHz
- 2GB RAM per core
- 1PB (RAW) distributed storage (HDFS)
  - Hard disks: SATA 7,2 RPM and SAS 7,2 RPM
  - Sizes: 1TB, 2TB and 4 TB
- Total HEP-SPEC: 8,0 K

Cluster Topology

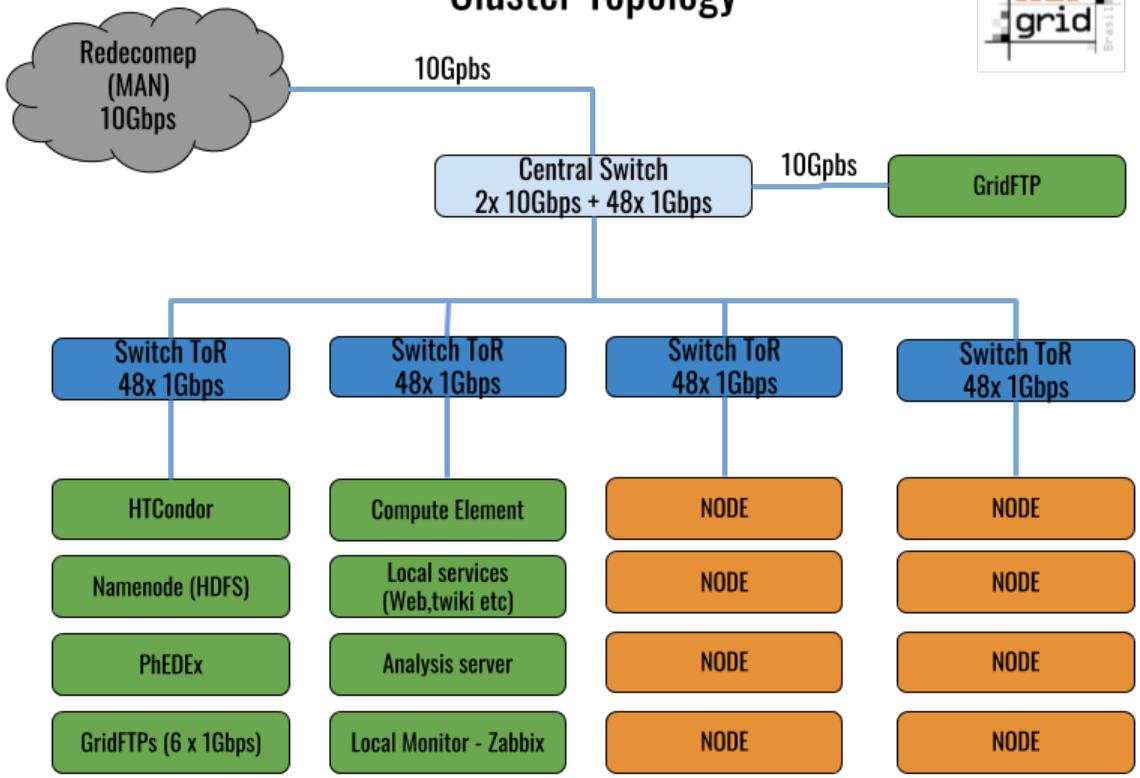

Local network - 1 Gbps

T2_BR_UERJ Cluster

- The infrastructure is fully dedicated do CMS
- Tied to U.S. Open Science Grid (OSG)
- Part of WLCG under the Latin America Federation
- Should be part of the LHCONE in 2020 (February)

CMS Group at UERJ
http://cms.uerj.br/group/

SAMPA Project (Brazil)

UNIVERSIDADE DE SÃO PAULO (USP)

SAMPA Project

USP Physics Institute

Member of the ALICE Collaboration

Main Researchers (Physicists)
Alexandre Suaide
Marcelo Munhoz

Computer Systems Analyst
Ricardo Romão

SAMPA Chip

- a new ASIC for the ALICE TPC and MCH upgrades

https://iopscience.iop.org/article/10.1088/1748-0221/11/02/C02088

SAMPA - Hardware and Network Infrastructure

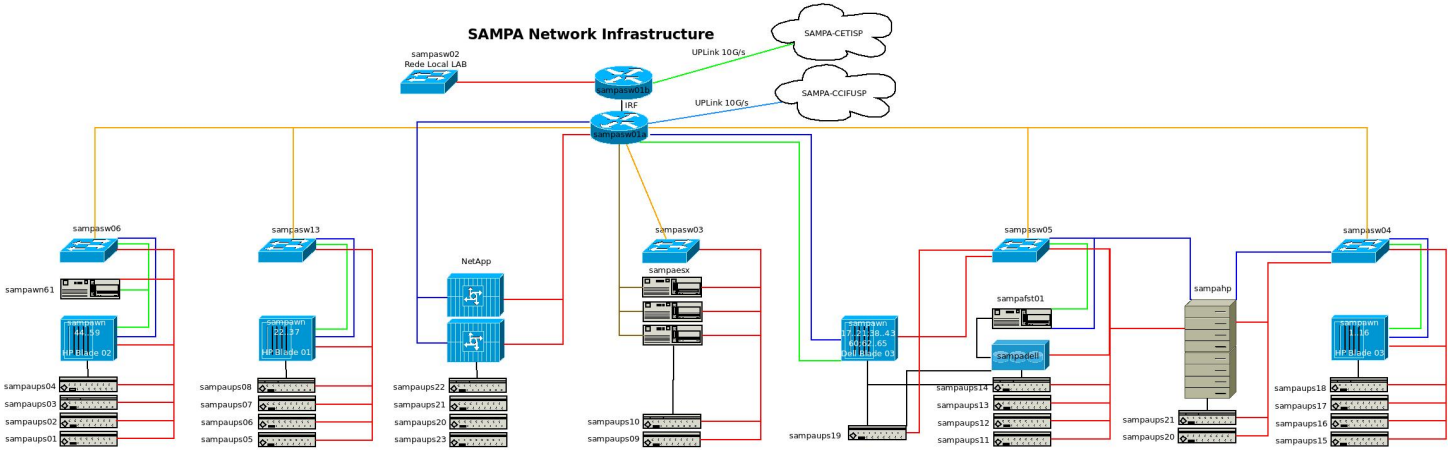

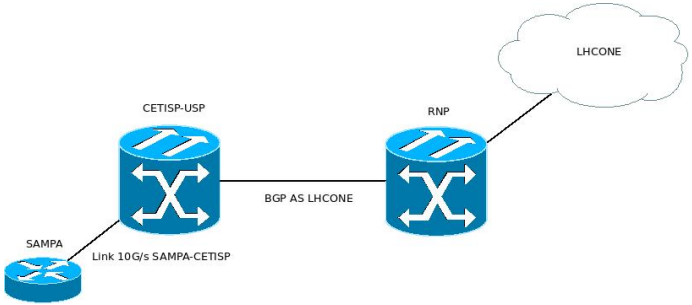

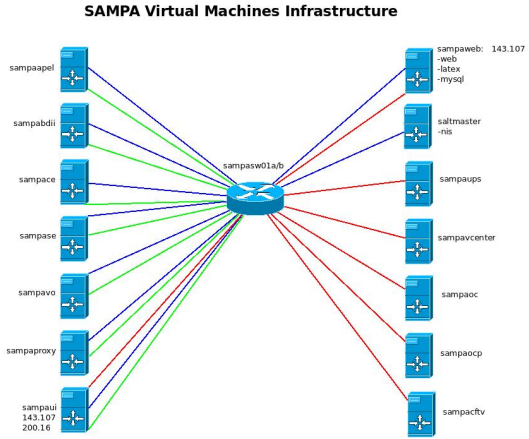

SAMPA - Total number of Jobs (by VO and Month)

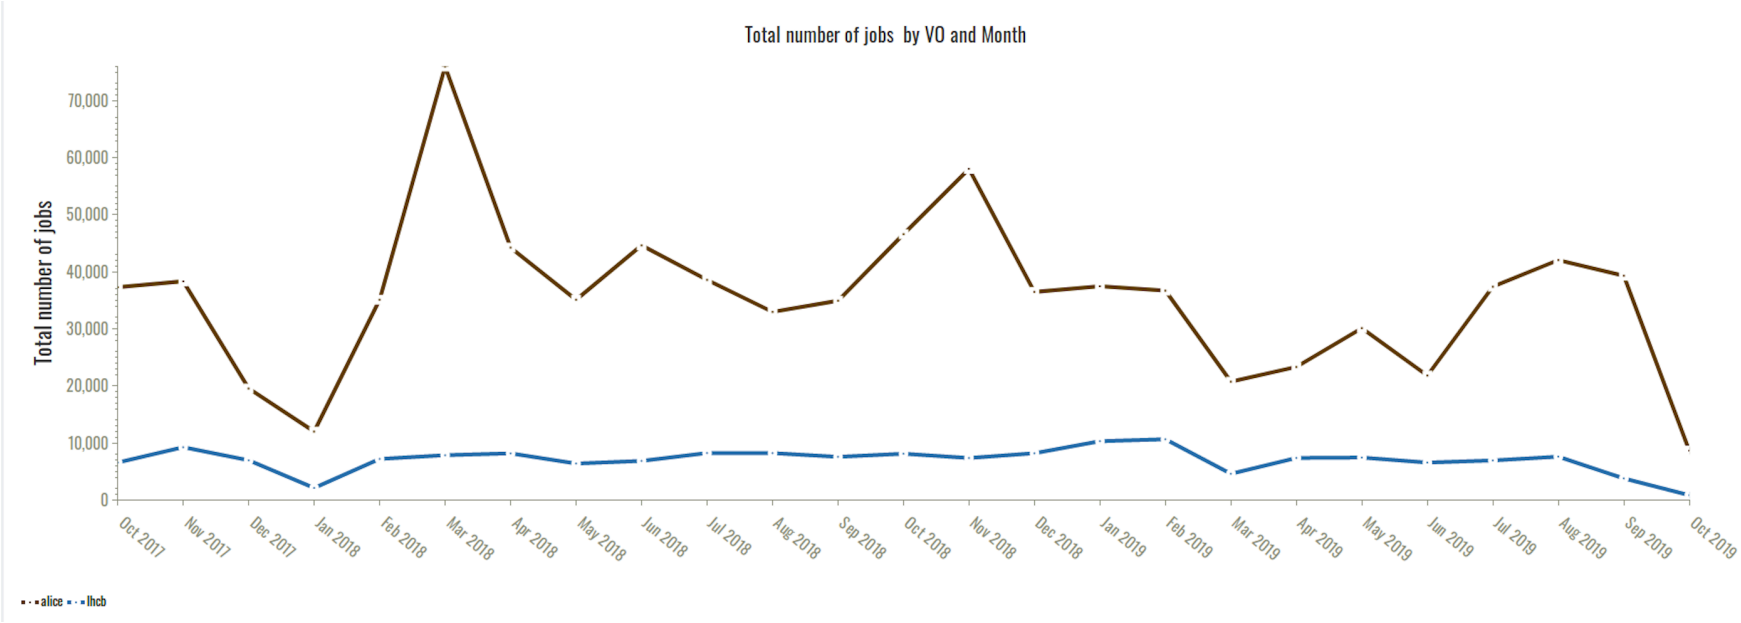

SAMPA - Availability & Reliability Report

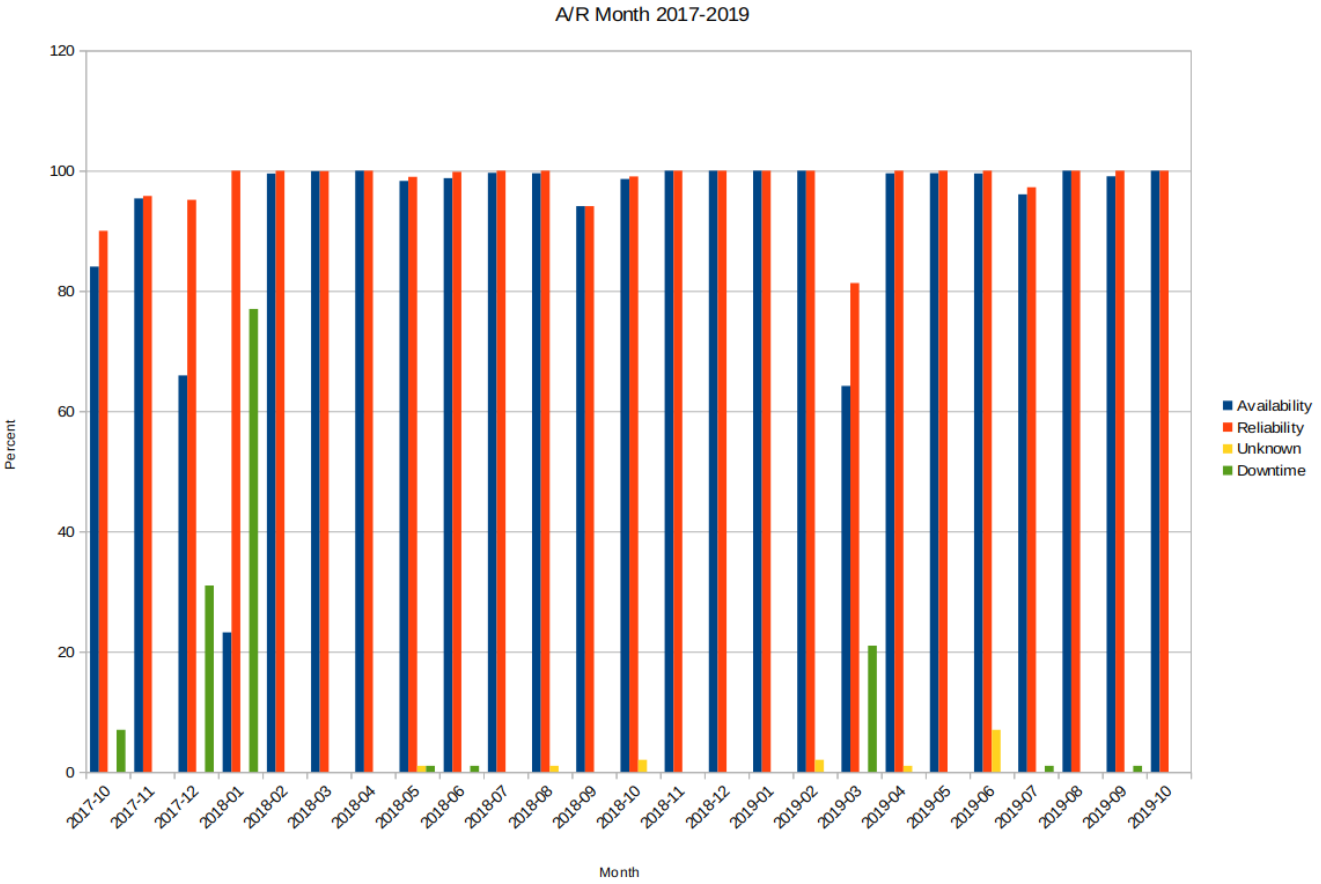

SPRACE Project (Brazil)

UNIVERSIDADE ESTADUAL PAULISTA (UNESP)

UNIVERSIDADE FEDERAL DO ABC (UFABC)

SPRACE - São Paulo Research and Analysis Center

Fundamental Research in High Energy Physics

- Physics analysis: Beyond SM and Heavy Ion Collisions
- Data processing and storage
- Scientific instrumentation

Innovation

- ❑ R&D on new technologies
  - Partnership with the private sector
  - Joint ventures with academic institutions

Outreach

- ❑ Share the knowledge with society
  - Poster, sites, game, Masterclass, etc.
  - High School teachers

SPRACE - Research and technical teams

Researchers / CMS members and users

Sérgio Novaes, Professor (Principal Investigator)
Eduardo Gregores, Professor @ UFABC
Pedro Mercadante, Professor @ UFABC
Sandra Padula, Researcher @ UNESP
Thiago Tomei, Researcher @ UNESP

Postdoctoral Researchers

César Bernardes
Luigi Calligaris

Research Fellows

Jefferson Coelho
Silvio Stanzani (PhD)

Scientific Computing staff (UNESP)

Angelo Santos (PhD)
Jadir Silva
Márcio Costa
Raphael Cóbe (PhD)
Rogério Iope (PhD)

Electrical Engineer (UNESP)

André Cascadan

Technical Support (UNESP)

Allan Szu
Ricardo Aguiar
Sidney Santos

SPRACE

CMS Tiers WLCG

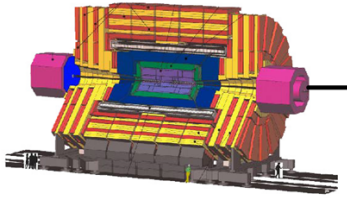

300 Tbps

Online System

500 Gbps

CERN *Switzerland*

1200 Km

3x100 Gbps

Wigner *Hungary*

10/40/100 Gbps

- Canada
- France
  - IN2P3
    - IRFU
      - IPHC
        - LLR
          - Belgium
            - China
- Germany
  - DESY
    - RWTH
      - Austria
        - Poland
- Italy
  - Bari
    - Legnaro
      - Pisa
        - Roma
- Netherlands
- Nordic
  - Greece
  - Hungary
  - India
  - Pakistan
  - Switzerland
  - Thailand
  - Turkey
  - Ukraine
- Korea
  - KISTI
- Russia NRC
- Russia JINR
  - IHEP
    - INR
      - ITEP
        - JINR
          - PNPI
  - CIEMAT
    - IFCA
      - Portugal
- Spain
- Taiwan
  - NCHC
- UK
  - Brunel
    - Bristol
      - IC
        - RAL
          - Estonia
            - Finland
- USA BNL
- USA FNAL
  - Caltech
    - Florida
      - MIT
        - Nebraska
  - Purdue
    - UCSD
      - Wisconsin
        - UERJ

SPRACE

SPRACE - Hardware Infrastructure

❑ Worker nodes (total: 128 physical servers)

- Physical Cores: 1792
- HT Cores: 2688
- HEPSpec06: 29702
- 128 cores: 3GB/core
- 768 cores: 4GB/core
- 512 cores: 2GB/core
- 1280 cores: 3.2GB/core

❑ Controller servers

- CE: HTCondor-CE job gateway and HTCondor job scheduler
- SE: dCache distributed storage system
- Shared filesystems: NFS
- 02 proxy services (Frontier Squid)
- Support services: monitoring, VM servers, DNS server, DHCP, etc.

❑ 13 dCache Storage Pool Servers

- 2.4 PiB = 2.7 PB (effective space)
- 10/40 Gbps NIC

❑ XRootD servers

- 01 Local Redirector
- 05 Data Servers

SPRACE

SPRACE - Network Infrastructure

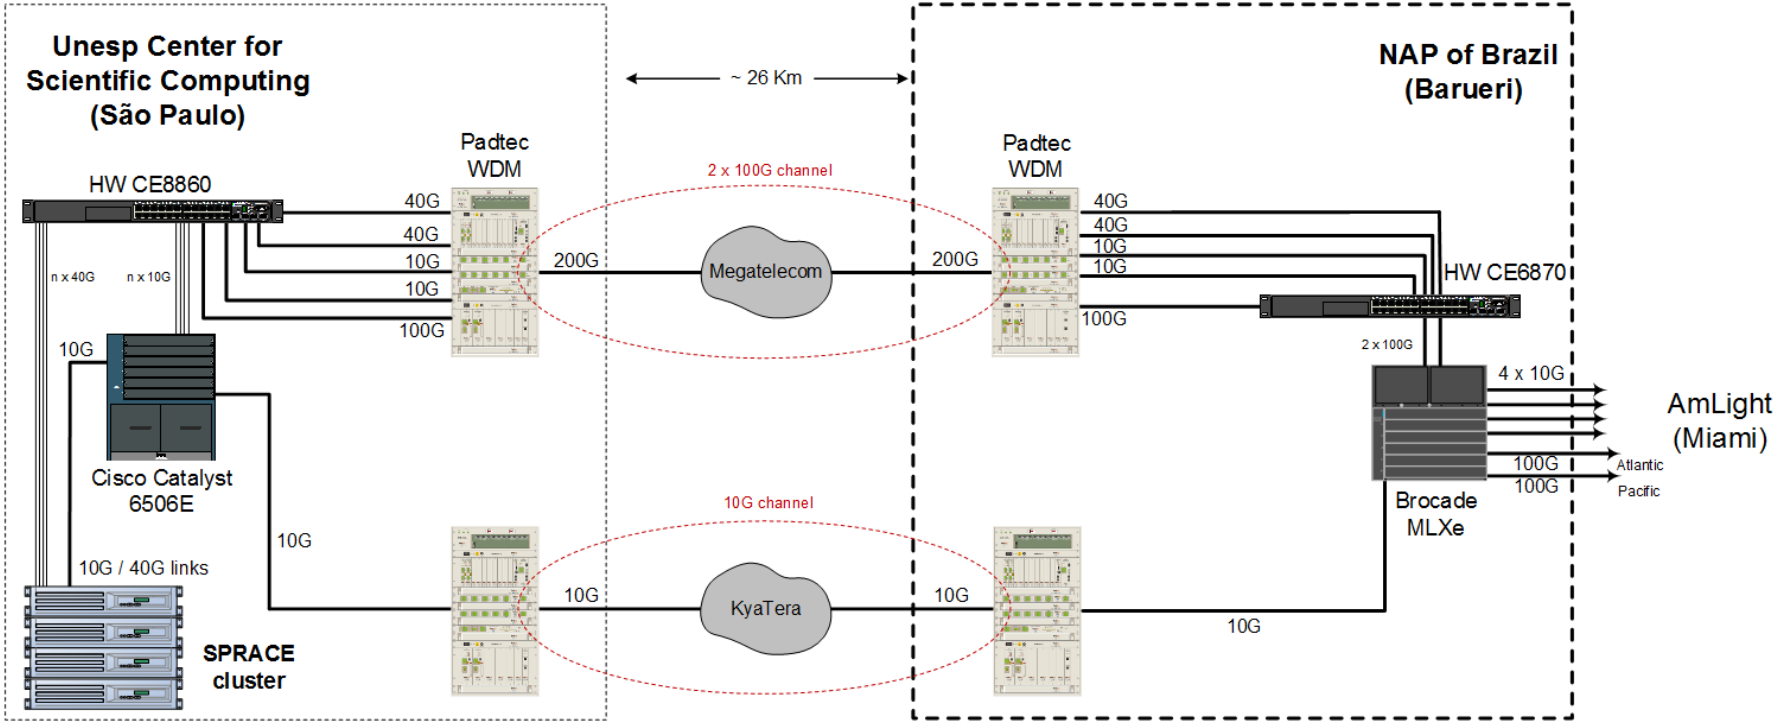

Tier-2 Availability and Reliablity Report

CMS

August 2019

Federation Summary - Sorted by Availability

Color coding: N/A <30% <60% <90% >=90%

Availability Algorithm: (CREAM-CE + ARC-CE + HTCONDOR-CE) * all SRM

Federation	Availability	Reliability
BR-SP-SPRACE	100%	100%
CERN-PROD	100%	100%
CH-CHIPP-CSCS	100%	100%
GR-Ioannina-HEP	100%	100%
T2_US_Nebraska	100%	100%
T2_US_UCSD	100%	100%
T2_US_Wisconsin	100%	100%
UA-Tier2-Federation	100%	100%
UK-London-Tier2	100%	100%
AT-HEPHY-VIENNA-UIBK	99%	99%
FR-GRIF	99%	99%
FR-IN2P3-CC-T2	99%	99%
PT-LIP-LCG-Tier2	99%	99%
T2_US_Purdue	99%	99%
UK-SouthGrid	99%	99%
FI-HIP-T2	98%	98%
HU-HGCC-T2	98%	98%
KR-KISTI-GSDC-02	98%	98%

Federation	Availability	Reliability
T2_US_Caltech	98%	98%
T2_US_Florida	98%	98%
T2_US_MIT	98%	98%
TW-CMS-T2	98%	98%
FR-IN2P3-IPHC	97%	100%
IN-INDIACMS-TIFR	97%	97%
ES-CMS-T2	96%	96%
PK-CMS-T2	96%	96%
RU-RDIG	96%	97%
IT-INFN-T2	95%	95%
DE-DESY-RWTH-CMS-T2	92%	92%
EE-NICPB	91%	91%
BE-TIER2	88%	95%
TR-Tier2-federation	84%	84%
T2-LATINAMERICA	82%	82%
CN-IHEP	81%	99%
PL-TIER2-WLCG	73%	73%

GridUnesp: SPRACE spinoff

First Campus Grid in Latin America

- Processing: 80 TFlops (3,000+ cores)
- Storage: 300 TB
- Networking: 10 Gbps

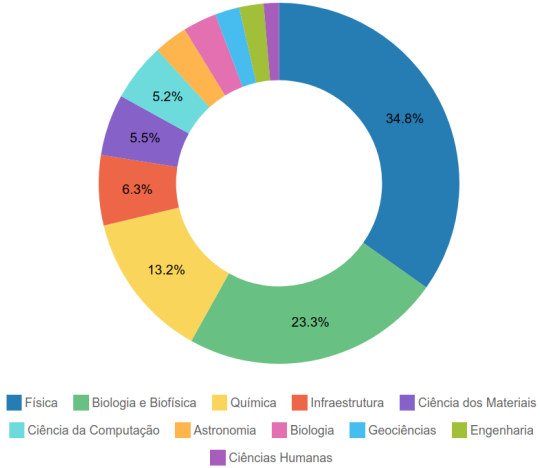

Research Fields

GridUnesp impact

- Only OSG VO outside USA
- HPC for **90+ projects** & **520+ users**
- Processes data also for
  - Higgs, LIGO, IceCube, etc.

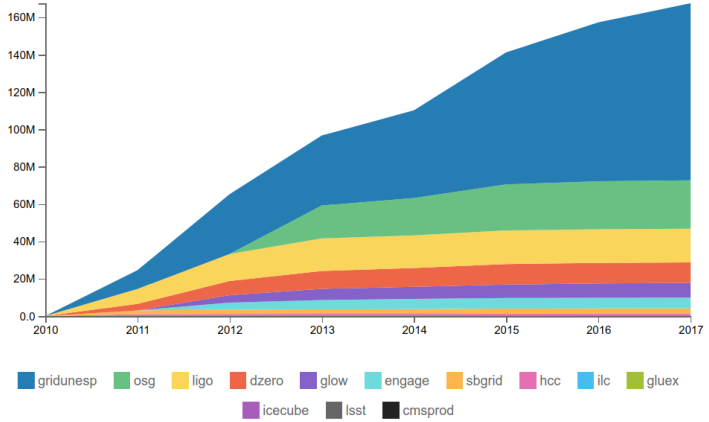

Wall Clock Hours

SPRACE - Intel R&D Projects

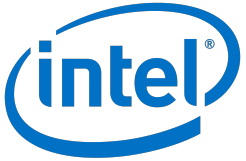

Intel Parallel Computing Center (IPCC)

- Parallelization of Geant (Geometry &Tracking)
- Broad impact
  - HEP: detector simulation
  - Radiation-hard electronics
  - Medical applications (dosimetry)
- Goals
  - Development of GeantV: massive parallelism natively
  - Test vector-coprocessor prototypes in hybrid systems
  - Analyze the performance of Geant4 X GeantV

Intel Modern Code Program (IMC)

- 1700+ students trained
- 7 International training events
- 26 tutorials at Brazilian Institutions

Center Excellence in Machine Learning

- High Energy Physics
- Proof of Concepts
  - SERPRO, Financial Institutions, etc.
  - R&D, consulting, and training in ML
- Wide range of applications
  - Image recognition for medical diagnosis, agribusiness, surveillance

SPRACE - Huawei R&D Project

Development of a Software-Defined Networking (SDN) Controller

- Kytos: an Open-source SDN Platform
- Plug and Play, responsive Web UI
- Easily scalable

Stress WAN for Data-Intensive Science

- Demonstrations at annual Supercomputing Conference
- 2016: New record of Data transmission North-South Hemispheres: **97.56 Gbps**
- 2017: Intercontinental ring: **350 Gbps**

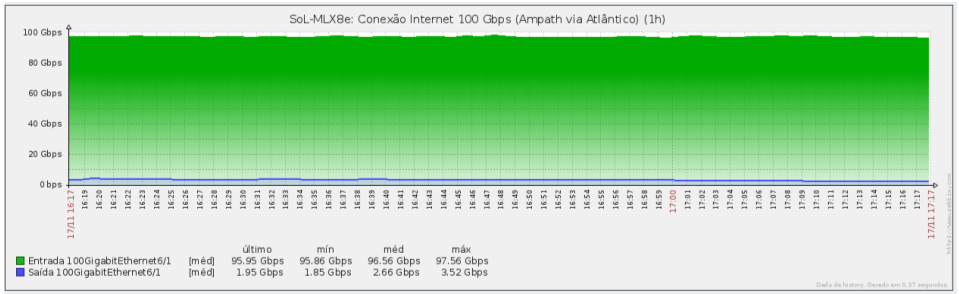

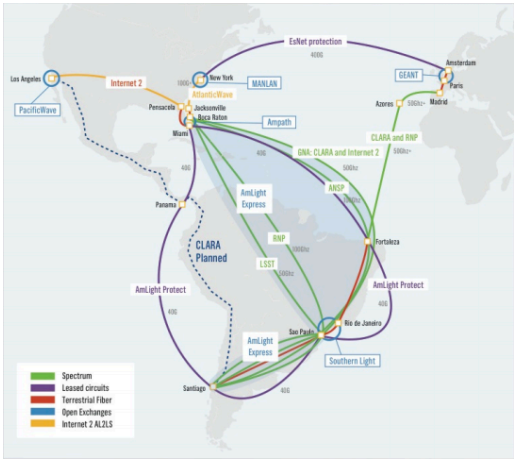

SPRACE - Machine Learning for Physics

Case study: High Luminosity LHC

- Sqrt(s) = 14 TeV
- Luminosity: 7.5x10^{34} $cm^{-2}s^{-1}$
- Pileup: 200

Extremely difficult tracking

- Timing increase by factor 50 to 120
- Efficiency comparable to regular LHC

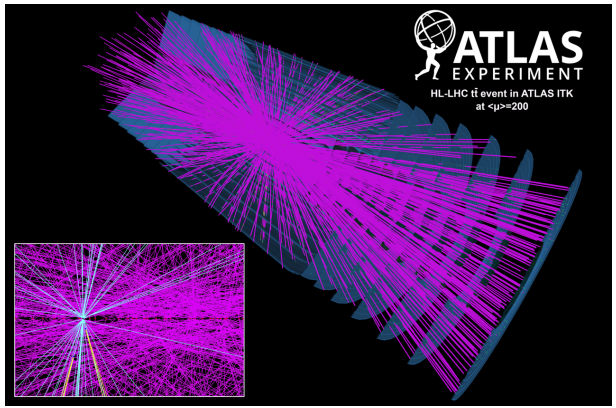

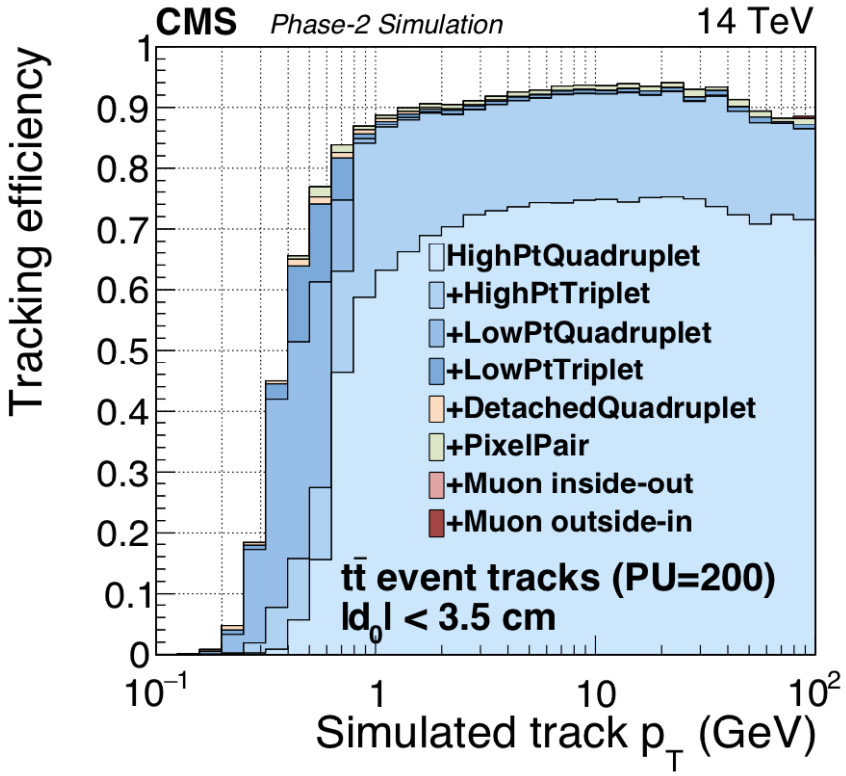

CBPF
(Brazilian Center for Research in Physics)

CBPF - LHCb Tier-2D

LHCb in Rio:

- **CBPF**
- **UFRJ Physics Institute**
- **PUC RJ**

11 Researchers

4 Technology / Computing specialists

3 Post doctors

5 PhD students

1 Master student

9 undergrad students

Total: 33

CBPF achievement: LHCb jobs in an HPC system

Santos Dumont

- **Atos-Bull hybrid HPC system installed at LNCC, in Petropolis**
- **~ 18.000 cores (including cores from GPUs and Xeon-Phis)**
- **Top 500 until 2016**

https://www.top500.org/system/178569

DIRAC

- **LHCb's Grid Workload and Data Management System**
- **Collaborative project to run DIRAC in the Brazilian Sdumont HPC system**
- **Jobs started running in SDumont in April 2018**

CBPF achievement: LHCb jobs in an HPC system

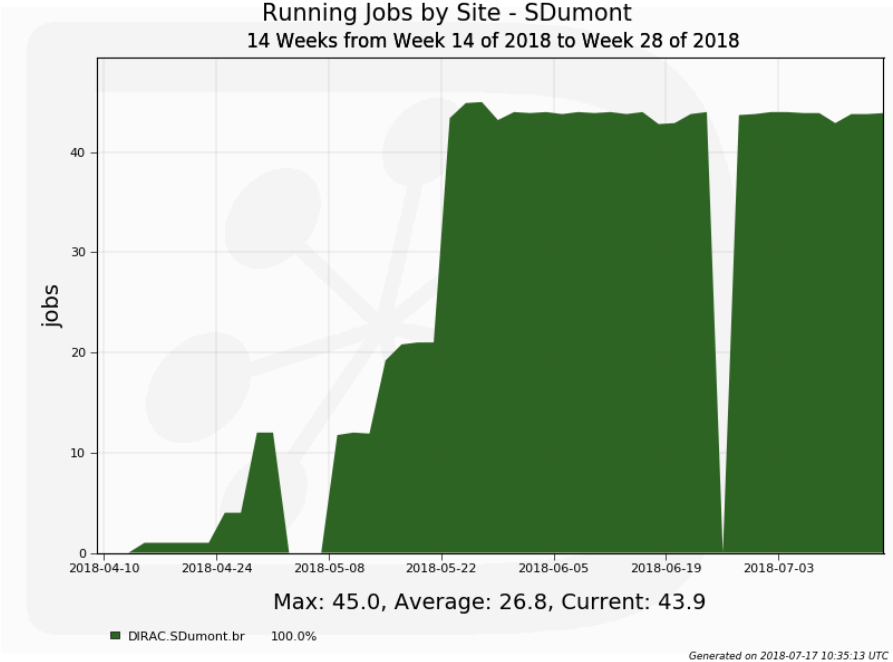

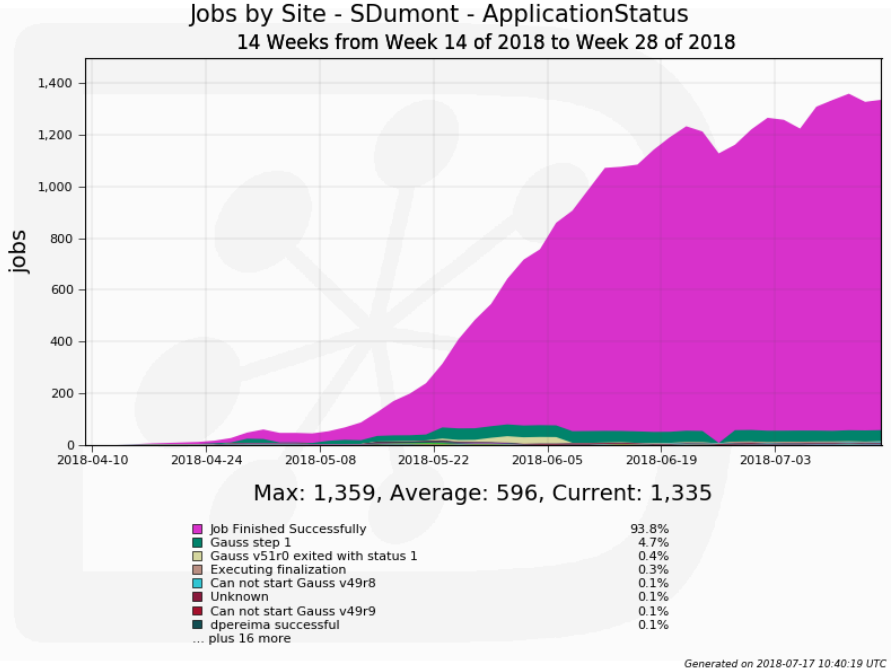

USM (Chile)
(Universidad Técnica Federico Santa Maria)

USM - Research team and projects

Research team

- Alfonso Zerwekh
- Claudio Dib
- Hayk Hakobyan
- Luis Salinas Carrasco
- William Brooks
- Yuri Ivanov

USM contributed on the development of the ClaRA framework

- ClaRA - **Cla**s12 **R**econstruction and **A**nalysis Framework
- a SOA based software framework for network distributed physics data processing applications
- Well defined, reusable, loosely coupled components
  - Multi-threaded event processing
  - Distributed event processing
  - Focus on data that is moving and transforming in the system
- implements Python, Java, and C++ in its Services layer
- includes multi-threaded and cloud/batch farm support

- References:
  - https://clasweb.jlab.org/wiki/index.php/CLAS12_Software
  - https://clasweb.jlab.org/wiki/index.php/CLARA
  - https://clas12svn.jlab.org/repos/
  - A talk about ClaRA can be found at https://indico.cern.ch/event/251143/

USM - Research work and projects (cont.)

For the ATLAS Collaboration

- USM team contributed on the development of a trigger emulation program for the ATLAS electron and photon trigger systems
  - Jorge Lopez

- R&D on a new ML project, which will perform event classification in ATLAS di-Higgs event candidates, which is a case where there is far more background than signal (referred to as a class-imbalance problem)
  - Raquel Pezoa

- Recently the USM team have implemented an analysis software to access offline and Fast Tracker trigger (FTK) full-event track and primary vertex collections
  - They used this software to determine the expected performance of the track multiplicity trigger for a pPb MC dataset and for both low and high luminosity pp collision datasets that included simulated FTK tracks
  - Carolina Robles

USM - Research work and projects (cont.)

For the ATLAS Collaboration (cont.)

- USM have also made major contributions to the Event Index system of ATLAS
  - Fedor Prokoshin
  - Ref.: https://www.researchgate.net/project/ATLAS-EventIndex

- at Fermilab, in MINERvA project
  - 3-person effort to use ML to reconstruct neutrino events, with the aim of contributing not only to **MINERvA** but also to pave the way for similar work for **DUNE** in the future
  - Anushree Ghosh (Postdoc researcher @ USM)
  - Ref.: https://arxiv.org/abs/1808.08332

PUCP
(Pontifícia Universidad Católica del Perú)

GRID computing at PUCP

Berkeley **O**pen **I**nfrastructure for **N**etwork **C**omputing is a platform for distributed high throughput computing

- Worker nodes come from **volunteer computing.**
- **BOINC** takes resources not used for the **volunteer**.
- Free software can be used.

Some **BOINC** projects:

SETI INSTITUTE

Looking for extraterrestrial life

ROSETTA

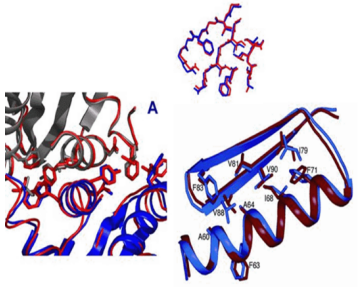

3D shape of the proteins

IBM

world community grid

Resp.: Prof. Alberto Martin Gago Medina
PUCP - Dept. Académico de Ciencias - Sección Física

GRID computing at PUCP - LEGION

- In its peak ~500 computers placed at different labs in the campus.
- Quad Core, 4GB RAM computers
- The users work with Windows

- **VMware** with **Scientific Linux (CERN)** were installed in the working nodes.
- **VMWare** takes 50% from the RAM and CPU.
- 1.6 TeraFLOP of computing power

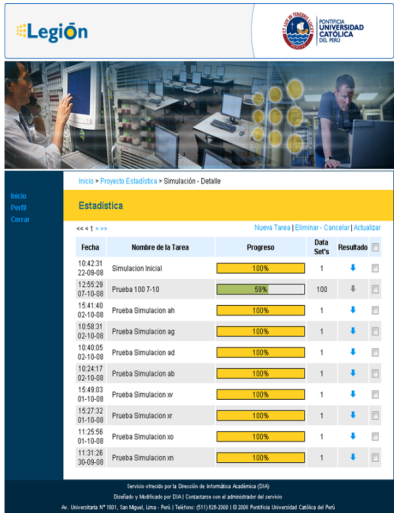

HEP-computing at PUCP

Several millions of ALICE collisions

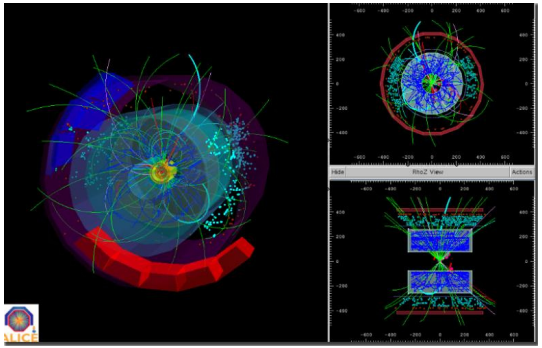

3D Modelling extragalactic neutrino fluxes

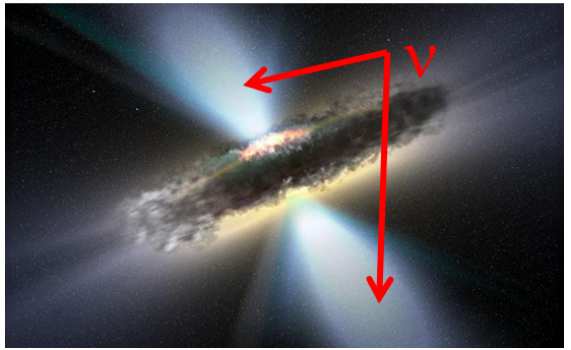

Probing Earth-crust cavities with a neutrino beam

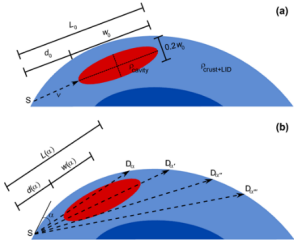

Simulations were equivalent to more than 45 years in a single CPU. In LEGION it took only 50 days.

The AmLight-Exp Project

The AmLight Express and Protect (AmLight-ExP) project is a 5-year National Science Foundation (NSF) award to FIU (OAC-1451018), and with support from the Association of Universities for Research in Astronomy (AURA) and the AmLight Consortium.

The AmLight Consortium members include FIU, ANSP, RNP, REUNA, RedCLARA, AURA, Florida LambdaRail (FLR), Internet2, Telecom Italia Sparkle, and Angola Cables.

https://www.amlight.net/?p=3935

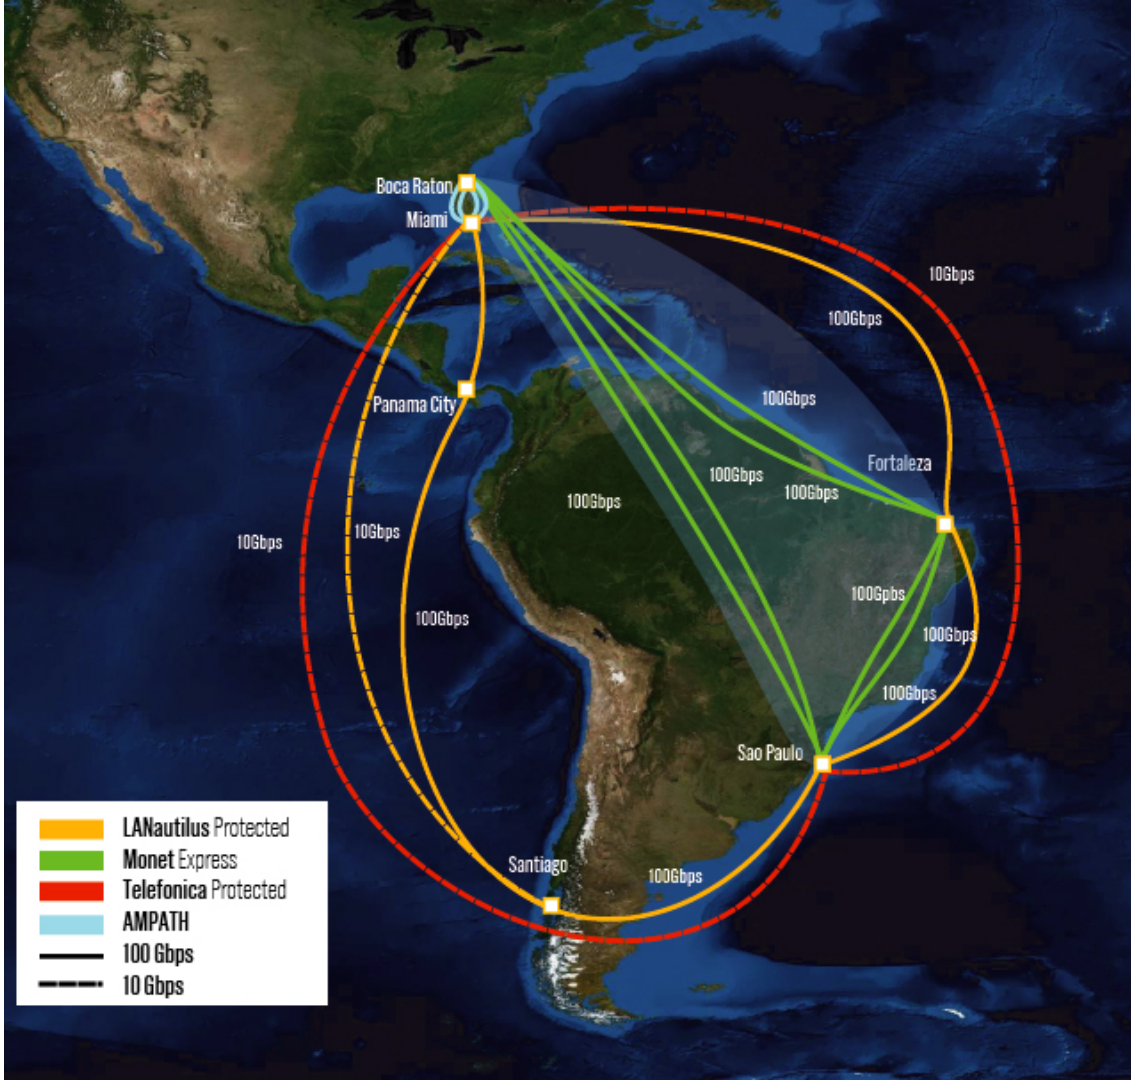

Thank you