

FIT@BR - a Future Internet Testbed in Brazil

Antonio Abelem¹, Michael Stanton^{2,3,*}, Iara Machado², Marcos Salvador⁴, Luiz Magalhaes³, Natalia Fernandes³, Sand Correa⁵, Kleber Cardoso⁵, Cesar Marcondes⁶, Joberto Martins⁷, Jose Monteiro⁸, Tereza Carvalho⁹ and José Rezende¹⁰

¹ UFPA – Federal University of Pará, Belém, PA, Brazil

² RNP – National Education and Research Network, Rio de Janeiro, RJ, Brazil

³ UFF – Federal Fluminense University, Niterói, RJ, Brazil

⁴ CPqD – Telecom Research and Development Center, Campinas, SP, Brazil

⁵ UFG – Federal University of Goiás, Goiânia, GO, Brazil

⁶ UFSCar – Federal University of São Carlos, São Carlos, SP, Brazil

⁷ UNIFACS – Salvador University, Salvador, BA, Brazil

⁸ UFPE – Federal University of Pernambuco, Recife, PE, Brazil

⁹ USP – University of São Paulo, SP, Brazil

¹⁰ UFRJ – Federal University of Rio de Janeiro, Rio de Janeiro, RJ, Brazil

E-Mails: abelem@ufpa.br; michael@rnp.br; iara@rnp.br; marcosrs@cpqd.com.br; schara@midia.com.uff.br; natalia@midia.com.uff.br; sand@inf.ufg.br; kleber@inf.ufg.br; marcondes@dc.ufscar.br; joberto@unifacs.br; suruagy@cin.ufpe.br; carvalho@larc.usp.br; rezende@land.ufrj.br

* Author to whom correspondence should be addressed; Fax: +55-21-2279-3731

Abstract: A major objective of the Brazil-EU FIBRE project is the deployment in Brazil of FIT@BR, a wide-area network testbed to support user experimentation in the design and validation of new network architectures and applications. In such a testbed, a high degree of automated resource sharing between experimenters is required, and the testbed itself must be instrumented so that precise measurements and accounting of both user and facility resources may be carried out. In this article, we describe the design and implementation of the Control and Monitoring Framework (CMF) for the FIT@BR testbed, which is based on three CMFs developed in existing testbed projects. In order to take best advantage of different testbed functionalities at different sites, FIT@BR is being created as a federated testbed, which will facilitate future interoperability with international initiatives.

Keywords: Future Internet testbed, CMF – control and monitoring framework, monitoring, federation of testbeds.

1. Introduction

FIBRE (Future Internet testbeds/experimentation between Brazil and Europe)¹ is one of five projects that were approved in response to the 2010 Brazil-EU Coordinated Call in ICT, jointly funded by CNPq (the Brazilian Council for Scientific and Technological Development) and by the 7th Framework Programme (FP7) of the European Commission. The main objective of FIBRE is to create a common space between the EU and Brazil for Future Internet (FI) experimental research into network infrastructure and distributed applications. Prior to FIBRE, such facilities already were operated, or were being built, by partners in this project from both sides of the Atlantic Ocean. FIBRE was designed so that such a space would enable and encourage closer and more extensive BR-EU cooperation in FI research and experimentation, as well as strengthening the participation of both communities in the increasingly important global collaborations in this important area of network research and development.

The EU-side partners in FIBRE (i2CAT, NWX, NICTA, UnivBRIST, UPMC and UTH) are also participants in the EC's FIRE (Future Internet Research and Experimentation)² testbed projects CHANGE³, OpenLab⁴ and OFELIA⁵.

An important characteristic of OFELIA is its leveraging of the OpenFlow (OF) approach from Stanford [1]. The participation of OpenLab partners allows extensions of the OFELIA approach to new testbed environments and use cases not included in OFELIA, especially in the fields of wireless communications. In this latter area, considerable expertise in designing, building and evaluating large-scale testbed systems is brought to the project through the participation of National ICT Australia (NICTA), which has been a major contributor to the development of the OMF control framework [2].

Brazilian partners have brought to the project experience in participating in different FI testbed projects. One of these is the Project GIGA, managed by CPqD and RNP, which built and still operates a wide-area optical network testbed beginning in 2003 [3], which was afterwards redirected to FI experimentation in 2009. A second relevant project is the INCT (National Institutes for Science and Technology, a CNPq programme) Web Science project⁶, which effectively began in 2010, and whose Future Internet Architectures subproject currently involves RNP and 4 university partners (UFF, UFPA, UFPE, and USP), with expertise in optical and wireless networks, simulation and emulation studies, and network monitoring. The remaining

¹ <http://www.fibre-ict.eu>

² <http://www.ict-fire.eu>

³ <http://www.change-project.eu>

⁴ <http://www.ict-openlab.eu>

⁵ <http://www.fp7-ofelia.eu>

⁶ <http://webscience.org.br>

university partners in the present proposal (UFG, UNIFACS, UFRJ and UFSCar) bring to the collaboration expertise in wireless communications and network software development.

This article concentrates on the development of the Brazilian side of the FIBRE proposal, also known as FIBRE-BR, particularly the construction of the FIT@BR testbed facility and its interaction with other testbeds, both within the FIBRE-EU consortium and outside it.

2. Physical infrastructure of the FIT@BR testbed

The testbed, already partially complete, includes sites (aka islands) at each of the ten partner institutions in this project. Each site has a common nucleus of OpenFlow-capable switches, some based on NetFPGA and others on production-quality switches, together with their controller(s), as well as a cluster of compute and storage servers, appropriately virtualised, and (usually) a cluster of virtualised wireless nodes. Each site will propose its own possible extensions, integrating site-specific resources to FIBRE, such as wireless access testbeds (WiFi, WiMax, 3G/4G), OF-enabled equipment, optical networks or even more complex testbeds with heterogeneous resources and their own control framework (e.g.: the Emulab⁷ cluster at USP). Figure 1 illustrates a FIT@BR site, its common facilities and external connectivity.

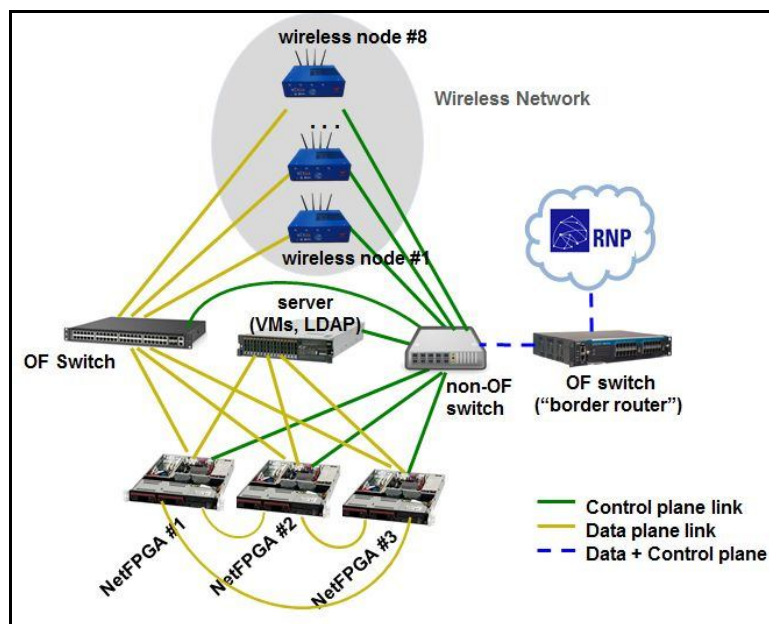


Figure 1. Overview of a FIT@BR island.

These sites have been interconnected using private (level-2) links over wide-area and metropolitan networks for research and education use. These include RNP’s national backbone

⁷ <http://www.emulab.net>

network and the GIGA testbed network jointly maintained by RNP and CPqD. RNP-owned metropolitan networks will be used for access where necessary, and RNP international connections will permit federation with overseas testbeds, such as OFELIA (at i2CAT and UnivBRIST) and NITOS (at UTH) [4].

3. Control Monitoring and Measurement Frameworks (CMFs)

In FIT@BR it was decided to use three CMFs instead of one: OFELIA Control Framework (OCF)⁸, ProtoGENI⁹, and OMF¹⁰ [2]. The use of more than one CMF represents a gain for the project because it enables FIT@BR to be seen as a federation from the beginning, permitting the simultaneous orchestration of three complementary classes of resource: OpenFlow resources, emulated resources and wireless resources.

The OFELIA Control Framework (OCF) was originally created in the context of the OFELIA testbed project but today it is supported by a larger community including FIBRE and the GEANT network. OCF is synchronized with the GENI¹¹ initiative in the US and follows the Slice-based Facility Architecture (SFA) orientation [5]. The first OCF implementation was based on Expedient and the Opt-in manager [6], both deployed in Stanford University, but it has evolved to a completely independent suite, which also introduces the VT-Manager to handle virtualizing servers. From the point of view of experimenters (or network researchers) the available underlying network substrate is fully controllable using explicit and dynamic configurations based on OpenFlow abstractions like FlowSpace.

OMF [2] is a CMF with the focus on controlling and managing network devices. It was developed based on XMPP (eXtensible Messaging and Presence Protocol) in the Ruby language. The OMF suite also provides OML (OMF Monitoring Library) [7], which allows instrumentation of applications for collecting measurements. Currently, OMF depends on a suite of software components in order to provide services such as dynamic IP addressing, remote booting, remote OS imaging, data archiving, etc. The current version of OMF supports a large number of different wired and wireless resources, such as PC-like devices, Android phones, routers and sensors, among others.

ProtoGENI is a CMF from the University of Utah. It is based on an enhanced version of the management software for the Emulab testbed. ProtoGENI was created to provide the integration between Emulab and other testbeds in order to build the Cluster C facility of GENI. It is SFA-oriented for testbed federation.

⁸ <http://fp7-ofelia.github.com/ocf>

⁹ <http://www.protogeni.net>

¹⁰ <http://mytestbed.net>

¹¹ <http://www.geni.net>

One of the FIBRE goals is to design a framework where all the CMFs in use can complement each other, as well as federate different instances of the same CMF.

4. Federation

In FIBRE the federation approach enables two or more organisations to devote some of their resources to a shared purpose. Each federation technique has two components: a set of standards, which provides connectivity, and a set of policies, which states what levels of connectivity are allowed (e.g., the amount of resources shared between the federated parties). From the technical point of view, FIBRE will adopt a Sliced-based Federation Architecture (SFA) [5]. However, each CMF installed in the FIBRE islands must implement web services compliant with SFA 2.0. Among the provided functionalities, SFA enables federation members to track and exchange services among islands. The idea is to provide a distributed management facility through a common and safe API that builds an extra layer between the CMFs and the users. In fact, features like member authentication and resource discovery and accountability are among the most important characteristics of the SFA. Unfortunately, standardization of the use policies still needs to be developed.

MySlice¹² is a software layer that enables the creation of a federation abstraction. Such an interface is based on a web client allowing users to interact with the great volume of results generated by each island testbed.

The FIBRE federated architecture will allow users to access the testbed through an integrated interface for either experimental or control planes, which provides a common access to the underlying different CMFs. In addition, enhancements are being developed for collection and visualisation of results to help researchers to deal with the experiment as a whole, i.e. providing an integrated view of the experiment.

5. Instrumentation and monitoring

A major FIT@BR goal is to provide Instrumentation and Measurement (I&M) facilities in order to help experimenters, network administrators and researchers to define experiments and collect infrastructure and/or experiment specific data.

The basic requirement for the FIT@BR I&M Architecture is the capability to configure, monitor, collect, and display both infrastructure and experiment specific data for distinct federated or individual CMF aggregates. Besides that, the architecture also includes or considers the set of requirements adopted by the FIT@BR experimental facility which includes substrate measurements, experiment measurements, privacy of measured data, link measurements, measurement data transmission and storage, component locations, time services, and federation.

¹² <http://myslice.info>

The building blocks of the FIT@BR I&M architecture leverage from current CMFs' measurement capabilities. They create an integration facility (MDIP – Measurement Data Integration Point) that conveys the measured data to other FIT@BR I&M services and users through a common schema and protocol, according to its policy and experimental network requirements.

The Orchestration and Configuration Services act on behalf of the users allowing them to configure, to define measurement points, and to orchestrate the measurement data collection facilities according to each individual CMF.

The I&M Portal is another component of the I&M architecture. Its main functionality is to provide a user friendly interface to control and access the measured data. As required by any data sharing facility, the experiment data will be available to users according to a defined policy.

The Measurement Data Integration Point (MDIP) conforms the collected data from the available CMFs to FIT@BR I&M standard format, representation and distribution (including visualization). This service includes all measurement data processing related aspects such as, message format, message transport protocol and/or service, access privileges, and common data storage or on-the-fly data distribution.

For each CMF in FIT@BR, MDIP requires a specific I&M implementation in order to deal with tools and architecture involved, data access and format issues. Initially, three cases are considered: OMF, ProtoGENI and OFELIA.

As an illustration of the architectural approach, in the OMF scenario, MDIP interfaces with the OML server and sends the requested data to the I&M Persistent Data Repository. In this case, a SQL-to-SQL adjustment is required and data format issues are conformed to NM-WG standard. As a second illustration case, the integration with ProtoGENI CMF benefits from the LAMP (Leveraging and Abstracting Measurements with perfSONAR)¹³ architecture, which is based on perfSONAR [8]. In this specific case, MDIP data manipulation is much simpler due to the fact that perfSONAR format and schema have been adopted by FIT@BR I&M architecture.

FIT@BR I&M Architecture is intended to be an evolutionary I&M architecture in the sense that, firstly, it evolves from integrating single aggregates belonging to common CMFs and, secondly, it integrates federated aggregates from multiple CMFs.

6. Operation

The operation of FIT@BR testbed is managed and monitored by the NOC (Network Operation Centre), which is responsible for controlling and monitoring the network assets of the testbed and monitoring the services provided, to support its users.

The NOC provides support at 3 levels:

¹³ <http://groups.geni.net/geni/wiki/LAMP>

- Level 1 – acts as a single point of contact (SPOC) for the end user. This level of support should deal with well-known problems, while escalating to Level 2 incidents related to a specific island, and/or to Level 3 incidents that do not have a known resolution and are not related to any particular island.

- Level 2 – triggered by Level 1, when it is not able to solve an incident. Level 2 is offered in a distributed fashion. If a particular island is unable to solve the incident, it should be escalated to Level 3 for resolution.

- Level 3 – consists of a geographically distributed development team responsible for fixing software and hardware problems.

FIT@BR testbed user authentication is carried out using a LDAP [9] directory in each island synchronised with a LDAP directory at the NOC. LDAP allows authentication with all CMFs (OMF, OCF, and ProtoGENI). The road map for future authentication is to migrate to a Federated model using Shibboleth¹⁴ through the Brazilian Academic Network Federation, CAFe¹⁵ (“Comunidade Acadêmica FEderada - Federated Academic Community”), and integrate this solution to MySlice, an integrated access portal for all CMFs.

8. Perspectives and conclusion

The FIBRE project effectively began in October, 2011, and will terminate in March, 2014. The FIT@BR testbed is planned to be operational in the second semestre of 2013, at least to support the research activities of the development team, and to be open for a wider user community in 2014. Interest has been generated in Brazil and the testbed is currently being extended to include islands at six additional universities in 2013.

International collaboration with the FIBRE-EU consortium includes federation of the testbed resources before the end of the FIBRE project, and the demonstration of the use cases which include participants from both continents. In addition, informal collaboration is already underway with some projects of the GENI community, especially involving the ProtoGENI site at USP, and we expect an increase in this activity in the near future.

The main motivations of the Brazilian participants in the FIBRE project were to gain experience in building a large-scale FI testbed in Brazil, and to raise the level of interest in research into FI architectures. The first of these has been accomplished, and it is felt that the second will be attained through effective international collaboration in experimental research. Encouraging signs of further interest in such international collaboration were given at a meeting of representatives of testbed projects from the US, Canada, EU, Japan, Australia and Brazil (FIT@BR), at the 17th GENI Engineering Conference (GEC17) in July, 2013, where it was decided to attempt to demonstrate their effective federation at GEC18 in October, 2013.

¹⁴ <http://shibboleth.net>

¹⁵ <https://www.rnp.br/en/services/cafe.html>

Acknowledgements

The FIBRE project is funded under Grant Agreement numbers 288356 (EC/FP7) and 590022/2011-3 (CNPq). The Web Science project is funded under Grant Agreement number 557128/2009-9 (CNPq).

References

1. McKeown, N. et al., OpenFlow: Enabling Innovation in Campus Networks, *ACM SIGCOMM Computer Communication Review*, 2008, 2, 69-74.
2. T. Rakotoarivelo, M. Ott, G. Jourjon, I. Seskar, OMF: A Control and Management Framework for Networking Testbeds, *ACM SIGOPS Operating Systems Review*, 2010, 4.
3. Scarabucci, R.R., et al. (2005), Project GIGA – High-speed Experimental Network. In *First International Conference on Testbeds and Research Infrastructures for the DEvelopment of NeTworks and COMMunities (TRIDENTCOM'05)*; Trento, Italy, 2005; pp. 242-251.
4. Welcome to NITOS Testbed Scheduler. Available online at: <http://nitlab.inf.uth.gr/NITlab/index.php/scheduler>, accessed on 14/6/2013.
5. L. Peterson, R. Ricci, A. Falk, J. Chase, Slice-Based Federation Architecture, version 2.0; July 2010; Available online at: <http://groups.geni.net/geni/attachment/wiki/SliceFedArch/SFA2.0.pdf>, accessed on 14/6/2013.
6. Expedient: A Pluggable Centralized GENI Control Framework. Available online at: <http://yuba.stanford.edu/~jnaous/expedient>, accessed on 14/6/2013.
7. Jolyon White, Guillaum Jourjon, Thierry Rakotoarivelo, and Max Ott. Measurement architectures for network experiments with disconnected mobile nodes. In *TridentCom 2010*, May 2010.
8. Hanemann, A., Boote, J. W., Boyd, E. L., Durand, J., Kudarimoti, L., Lapacz, R., Swany, D. M., Zurawski, J., Trocha, S. PerfSONAR: A Service Oriented Architecture for Multi-Domain Network Monitoring. In *Proceedings of the Third International Conference on Service Oriented Computing*; Springer Verlag; LNCS 3826, pp. 241-254, 2005.
9. Lightweight Directory Access Protocol (LDAP): Technical Specification Road Map, RFC4115, 2006.

© 2013 by the authors; licensee Asia Pacific Advanced Network. This article is an open-access article distributed under the terms and conditions of the Creative Commons Attribution license (<http://creativecommons.org/licenses/by/3.0/>).