



FUTURE INTERNET TESTBEDS  
EXPERIMENTATION BETWEEN  
BRAZIL AND EUROPE



**BRAZIL**  
PAÍS RICO E PAÍS SEM POBREZA

Grant Agreement No.: 288356 (FP7)  
CNPq Grant Agreement No.: 590022/2011-3

## FIBRE-EU

Future Internet testbeds/experimentation between BRazil and Europe –  
EU

Instrument: *Collaborative Project*

Thematic Priority: *[ICT-2011.10.1 EU-Brazil] Research and Development cooperation,  
topic c) Future Internet – experimental facilities*

### D2.6. Report on the deployment of the second version of the control and monitoring framework for the FIBRE-BR facilities

Author: WP2 – Cesar Marcondes (UFSCar)

Revised by:


Due date of the Deliverable: Month 34

Actual submission date: 31/03/2014


Start date of project: June 1<sup>st</sup> 2011 Duration: 34 months

Version: v.1

Project co-funded by the European Commission in the 7 <sup>th</sup> Framework Programme (2007-2013)		
Dissemination Level		
PU	Public	✓
PP	Restricted to other programme participants (including the Commission Services)	
RE	Restricted to a group specified by the consortium (including the Commission Services)	
CO	Confidential, only for members of the consortium (including the Commission Services)	


	D2.6 Report on the deployment of the second version of the control and monitoring framework for the FIBRE-BR facilities	Doc FIBRE-D2.6
		Date 31/03/2014

<b>FP7 Grant Agreement No.</b>	288356
<b>CNPq Grant Agreement No.:</b>	590022/2011-3
<b>Project Name</b>	Future Internet testbeds/experimentation between BRazil and Europe – EU
<b>Document Name</b>	FIBRE-D2.6
<b>Document Title</b>	D2.5: Report on the deployment of the second version of the control and monitoring framework for the FIBRE-BR facilities
<b>Workpackage</b>	WP2
<b>Authors</b>	Cesar Marcondes (UFSCar) Ricardo de Freitas Gesuatto (UFSCar) Kleber Vieira Cardoso (UFG) Sand Luz Corrêa (UFG) Natalia Castro Fernandes (UFF) Edelberto Franco Silva (UFF) Débora C. Muchaluat-Saade (UFF) Noemi Rodrigues (PUC) José Augusto Suruagy (UFPE) Raphael Dourado (UFPE) Joberto S. B. Martins (UNIFACS) Adriano Spínola (UNIFACS)
<b>Editor</b>	Cesar Marcondes (UFSCar)
<b>Reviewers</b>	Carlos Bermudo (i2Cat) Sebastià Sallent (i2Cat/UPC)
<b>Delivery Date</b>	March 31 <sup>th</sup> 2014
<b>Version</b>	V1.0

	D2.6 Report on the deployment of the second version of the control and monitoring framework for the FIBRE-BR facilities	Doc FIBRE-D2.6
		Date 31/03/2014


## Abstract

The second version of the FIBRE CMF was reached at M34 (month) of the project as it is described here. There have been several major changes since the first version. As it can be verified in [Report 2.5], we described a more island centric view of the testbed, the major components and configuration and changes made to cope with Brazilian island peculiarities. While in this document, Deliverable 2.6 – Second version of CMF, there are several steps that turn FIBRE-BR testbed more ready to deploy for a large audience of experimenters. The contributions on the second version of CMF are organized in the following subsections: in section 2, we describe the topology that interconnect all the BR islands, this was a major advancement, by using the FIBREnet overlay backbone on top of RNP backbone, and OpenFlow controllers to configure the paths, therefore, it enabled the individual islands to interoperate at full performance, without relying on VPN tunnels. In Section 3, we describe a major contribution on developing a truly multi-CMF authentication environment, this software allows CMFs (like OMF and OCF) to fully integrate and also combines the registry flexibility and readiness of LDAP servers (used in the earlier phases of this work), along with state of the art slice authentication standard based on GENI Slice Federation (SFA). Therefore, allowing the creation of a single-logon web portal to experimenters explained in Section 4. In terms of monitoring features, Section 5, focus on two dimensions: infrastructure and experiment monitoring. Each one of these dimensions have their particular challenge and these are addressed by several tools, like PerfSONAR performance servers and ZenOSS, customized to make available easily cross islands measurements. Finally, in Section 6, we presented some further modifications done on the CMFs that form the FIBRE-BR testbed, OCF and OMF. The modifications were done on some very specific issues related to BR LDAP authentication, robustness to errors and testbed instabilities, and improved security, all these created for OCF. And finally, improved federation by completely redesigning NITOS OMF Scheduler, in a new scheduler called LABORA. We finish the document, in Section 7, concluding the version 2 of the FIBRE-BR software by showing how the deployment status currently holds. It can be noticed that the address scheme changed from [Report 2.5] due to turn the address space compatible to Europe for federation purposes, such that each continent has a 10.0.0.0/16 block. Hence, it can be concluded that the FIBRE BR testbed is ready to start receiving experimenters.

	D2.6 Report on the deployment of the second version of the control and monitoring framework for the FIBRE-BR facilities	Doc FIBRE-D2.6
		Date 31/03/2014

## TABLE OF CONTENTS

1	Overview of first version of FIBRE-BR CMF and accomplishments of the second version .....	9
2	FIBREnet Architecture and Deployment Tests .....	10
3	Multi-CMF Authentication using LDAP and SFA .....	14
3.1	Credential translation.....	15
3.2	Credential Translation and Access Control Modules.....	15
3.3	Credential Delegation Module .....	17
3.4	Implementation and Results .....	18
4	Single Logon Web Portal Customization .....	20
5	Monitoring Infrastructure and Experiments .....	22
5.1	perfSONAR Performance Toolkit Deployment .....	22
5.2	ZenOSS Deployment.....	24
5.3	Experiment Monitoring Support .....	25
6	Modifications of OCF and OMF Software for Second Version .....	29
6.1	OCF Modifications for Second Version.....	29
6.1.1	LDAP Integration on OCF VM Template .....	29
6.1.2	Adding Robustness to OXA by Auto-Recovery .....	31
6.1.3	Adding Security and Network Connectivity to OXA VMs .....	32
6.2	OMF Modifications for Second Version .....	34
7	Islands Deployment Status .....	37
8	Conclusions and Future Work .....	41


	D2.6 Report on the deployment of the second version of the control and monitoring framework for the FIBRE-BR facilities	Doc FIBRE-D2.6
		Date 31/03/2014

## List of Figures

Figure 1 - Data and Control Plane Topology in 1e Phase.....	11
Figure 2 - FIBRE-BR overlay network topology in 1p phase.....	12
Figure 3 - High-level view of the federation integration proposal.....	15
Figure 4 - Communication between the elements of proposed architecture. ....	17
Figure 5 - Complete authentication scheme using MySlice.....	20
Figure 6 - MySlice login screen with the CAFé authentication option. ....	21
Figure 7 - Testbed resources shown by MySlice.....	21
Figure 9 - MaDDash homepage: a dashboard with measurements for each node pair .....	23
Figure 8 - Monitoring infrastructure status .....	23
Figure 10 - ZENOSS view of an infrastructure's topology .....	24
Figure 11 - Detailed Graphic of NIC's Throughput .....	24
Figure 12 - Zenoss Monitored Infrastructure .....	25
Figure 13 - The Orchestration Service as an intermediary between GUIs and Slices .....	26
Figure 14 – The Orchestration Service Complete Architecture .....	28
Figure 15 - OCF Login Authentication using FIBRE BR LDAP .....	30
Figure 16 - SSH Authentication using jean.menossi@ufscar.br registered in LDAP .....	30
Figure 17 - Internal Bridge br0 access to the Internet.....	32
Figure 18 – LABORA Scheduler architecture.....	34
Figure 19 – Web interface of the LABORA Scheduler .....	35
Figure 20 – Data organization in the LABORA Scheduler.....	36


## List of Tables

Table 1 - The Orchestration Service REST API (simplified).....	26
Table 2 - Current status of the services of FIBRE-BR islands.....	37

	D2.6 Report on the deployment of the second version of the control and monitoring framework for the FIBRE-BR facilities	Doc FIBRE-D2.6 Date 31/03/2014
--	---	-----------------------------------

## Acronyms

AM	Aggregate Manager
API	Application Programming Interface
CP	Control Plane
CMF	Control and Monitoring Framework
DPID	DataPath Identifier
FIBRE	Future Internet testbeds / experimentation between Brazil and Europe
FPGA	Field-Programmable Gate Array
FV	FlowVisor
GUI	Graphical User Interface
HTTP	HyperText Transfer Protocol
IM	Island Manager
LLDP	Link Layer Discovery Protocol
MS	Milestone
NE	Network Element
NMI	Network Management Interface
OCF	OFELIA Control Framework
OF	OpenFlow
OFELIA	OpenFlow in Europe: Linking Infrastructure and Applications
OMF	cOntrol, Management and Measurement Framework
OML	ORBIT Measurement Library
OSS	Open Source Software
QoS	Quality of Service
SDN	Software Defined Networking
SFA	Slice-based Federation Architecture
SNMP	Simple Network Management Protocol
TCP	Transmission Control Protocol
UDP	User Datagram Protocol
UNI	User Network Interface
VLAN	Virtual Local Area Network
VM	Virtual Machine
VT	Virtualization Technology
WP1	Project Management
WP2	Building and operating the Brazilian facility
WP3	Building and operating the European facility
WP4	Federation of facilities
WP5	Development of technology pilots and showcases
WP6	Dissemination and collaboration

	D2.6 Report on the deployment of the second version of the control and monitoring framework for the FIBRE-BR facilities	Doc FIBRE-D2.6
		Date 31/03/2014

## Reference Documents

[Report D2.2] Report on the design of the control monitoring framework, WP2 Team, 2012.

[Report D2.3] Report on implementation and testing of the FIBRE-BR facilities, WP2 Team, 2012.

[Report D2.4] FIBRE-BR operational plan, WP2 Team, 2012.

[Report D2.5] D2.5 Report on the deployment of the first version of the control and monitoring framework for the FIBRE-BR facilities, WP2 Team, 2013.

[Rakotoarivelo 2010] T. Rakotoarivelo, M. Ott, G. Jourjon, Iv. Seskar, "OMF: a control and management framework for networking testbeds", ACM SIGOPS Operating Systems Review archive Volume 43, Issue 4, 2010.

[OMF 2012] OMF official site. <http://mytestbed.net>. Last visited: 15-may-2013.

[Anadiotis 2010] A. C. Anadiotis, A. Apostolaras, D. Syrivelis, T. Korakis, L. Tassioulas, L. Rodriguez, I. Seskar, M. Ott, "Towards Maximizing Wireless Testbed Utilization using Spectrum Slicing", Internatinonal ICST Conference on Testbeds and Research Infrastructures for the Development of Networks and Communities (TridentCom), 2010.

[S2S XMPP] Openfire Server 2 Server connection.

[http://mytestbed.net/projects/omf54/wiki/Openfire\\_s2s](http://mytestbed.net/projects/omf54/wiki/Openfire_s2s). Last visited: 15-may-2013

[Internet2 2013] pSPT Official Website. <http://psps.perfsonar.net/toolkit/>. Last visited: 21-may-2013.

[ZenOSS 2013] ZenOSS Community version official website. <http://community.zenoss.org>. Last visited: 21-may-2013.


[SFA] L. Peterson, R. Ricci, A. Falk, J. Chase, "Slice-Based Federation Architecture", 2010.

[SFA WRAP] SFA Wrap official website. <http://sfawrap.info>. Last visited: 31-mar-2014.

[ITU-T 2009] ITU-T, "NGN identity management framework," 2009, recommendation Y.2720.

[Jensen 2012] J. Jensen, "Federated identity management challenges," in Availability, Reliability and Security (ARES), 2012 Seventh International Conference on, 2012, pp. 230–235.

[Dhungana 2013] R. Dhungana, A. Mohammad, A. Sharma, and I. Schoen, "Identity management framework for cloud networking infrastructure," in Innovations in Information Technology (IIT), 2013 9th International Conference on, 2013, pp. 13–17.

	D2.6 Report on the deployment of the second version of the control and monitoring framework for the FIBRE-BR facilities	Doc     FIBRE-D2.6  Date    31/03/2014
--	---	--

[Leskinen 2012] J. Leskinen, “Evaluation criteria for future identity management,” in Trust, Security and Privacy in Computing and Communications (Trust-Com), 2012 IEEE 11th International Conference on, 2012, pp. 801–806.

[Niinimäki 2004] M. Niinimäki, J. White, W. de Cerff, J. Hahkala, T. Niemi, and M. Pitkanen, “Using virtual organizations membership system with edg’s grid security and database access,” in Database and Expert Systems Applications, 2004. Proceedings. 15th International Workshop on, 2004, pp. 517–522.

[Lee 2005] Y.-J. Lee, “A dynamic virtual organization solution for web-services based grid middleware,” in Database and Expert Systems Applications, 2005. Proceedings. Sixteenth International Workshop on, 2005, pp. 40–44.


[Scavo 2005] S. Scavo, T. e Cantor, “Shibboleth architecture,” Tech. Rep., Jan. 2005. [Online]. Available: <http://shibboleth.internet2.edu/docs/draft-mace-shibboleth-tech-overview-latest.pdf>

[OASIS 2005] OASIS, Security Assertion Markup Language (SAML) v2.0, Std., 2005.

[Karp 2009] A. H. Karp, H. Haury, and M. H. Davis, “From abac to zbac: The evolution of access control models,” Hewlett-Packard laboratories, Tech. Rep., 2009. [Online]. Available: <http://www.hpl.hp.com/techreports/2009/HPL-2009-30.pdf>

[ProtoGeni] ProtoGeni ClearingHouse, <http://www.protonet.net/wiki/ClearingHouse> Desc, March, 2013.



	D2.6 Report on the deployment of the second version of the control and monitoring framework for the FIBRE-BR facilities	Doc FIBRE-D2.6
		Date 31/03/2014

## 1 Overview of first version of FIBRE-BR CMF and accomplishments of the second version

The purpose of this section is to give some review of deliverable 2.5, the so called First version of the FIBRE-BR CMF, and furthermore highlight the accomplishments executed in the second version of the CMF on FIBRE BR. The first version of the FIBRE CMF was based on a combination of two CMFs, OCF (OFELIA) and OMF, and we took a particular path by integrating initially these CMFs through a simple and unified LDAP authentication and VPN tunnels. Due to the hardware constraints in the Brazilian side, a complex bridge configuration was made, separating in the virtualization environment, the control and data (or experiment) planes. And, we tackle some issues related to heterogeneous equipments like netFPGA servers, OpenFlow switches, and also a centralized federation of OMF resources, VPN stitching, and a first version of basic monitoring using specific servers and GPS.


Since the first version, we identify several issues for a complete and scalable testbed. The issues were related to open questions in the Future Internet space, like:

- 1) How to interconnect islands using an SDN-like Layer-2 island-to-island approach?
- 2) How to perform experiments using a standard federated authentication while remaining compatible with legacy directory databases?
- 3) How to provide a single-logon experience to the user to access all resources?
- 4) How to keep measurements of the infrastructure along with experiments measurements in a clean way?
- 5) How to improve availability, security and improve federation using the two CMFs that FIBRE CMF relies upon?

Some of these questions motivated the second version of the FIBRE CMF, as we will present in the rest of this deliverable. The text is organized in such way that each question is answered fully or partially on each major modification and test performed. All the explanations are self-contained and present a concise overview of all the effort put in this task. A complete view of the evolution of the FIBRE CMF is the wiki documentation<sup>1</sup>, more than 130 wiki pages (in a process of being translated to English).

In the following, we start the deliverable giving some overview idea of the overlay backbone called FIBREnet that allows individual islands to get interconnected using Layer 2 connectivity.

<sup>1</sup> <https://wiki.rnp.br/display/fibre/Home>

	D2.6 Report on the deployment of the second version of the control and monitoring framework for the FIBRE-BR facilities	Doc FIBRE-D2.6 Date 31/03/2014
--	---	-----------------------------------

## 2 FIBREnet Architecture and Deployment Tests

The FIBRENet is a VPLS overlay network created upon the RNP's backbone (Ipê network), in which the network plane is divided in two parts: an experiment plane and a control plane. The control plane in FIBRENet is prepared to offer control through the central FIBRE-BR NOC. Thus the experimenter can manage and monitor control information among FIBRE-BR islands. In other words, this channel allows communication between control frameworks of all islands and to collect monitoring information. For this an overlay network was built on the backbone of the RNP.


Virtual Private LAN Service (VPLS) is a way to provide an Ethernet multipoint-to-multipoint communication over IP / MPLS networks. It allows geographically dispersed sites to share the same Ethernet broadcast domain by connecting them through pseudo-wires. The technologies that can be used as Ethernet over MPLS pseudo-wires are, L2TPv3 or even GRE. There are two IETF RFCs (RFC 4761 and RFC 4762) describing VPLS establishment.

The control plane is an overlay mesh network and is used to traffic the data corresponding to the communication of the control frameworks, like OCF and OMF, and the gathering of monitoring information from the components used in the testbed.

The network overlay data plane is used to route traffic between the participants of the experiments using dedicated circuits islands between PoPs. The use of each circuit is determined by identifiers - VLAN IDs - to be defined by administrators.

The network overlay control plane will be used to control data sharing between islands such as monitoring information, OpenFlow control messages between device and controller or remote access for management equipment FIBRE-Net. Furthermore, the control plane network includes the NOC FIBRE-BR, which will operate and control of FIBRE-Net backbone, used by the experimenters during the construction of their inter-island topologies.

The FIBRENet have two phases 1p e 1e. The construction of the backbone FIBRE-Net will be held in two phases: **1e** and **1p**. On the 1e phase is developed on way experimental for executing and learning tests using a reduced number of islands. The second, 1p phase, it is FIBRENet network in production, with all project participants islands.

	D2.6 Report on the deployment of the second version of the control and monitoring framework for the FIBRE-BR facilities	Doc FIBRE-D2.6
		Date 31/03/2014

## 2.1 Topology 1e Phase

On 1e stage, the FIBRE-Net has been developed on experimental way with the participation of some islands strategically chosen for facilitating communication between the islands and PoPs. Among the participating institutions were chosen:

- 1- UFPA (Federal University of Pará ) and PoP-PA (Point-of-Presence of Pará)
- 2- UFG (Federal University of Goiás) and PoP-GO (Point-of-Presence of Goiás)
- 3- USP (University of São Paulo) and PoP-SP (Point-of-Presence São Paulo)

The Phase 1 has a ring topology based on three nodes (island) and one aggregate node, which is located in the Point-of-Presence of Distrito Federal (PoP-DF). Moreover, this phase two overlay networks are created for transport of both traffics the control and data plane. The goal is able to ensure logic or traffic isolation. The data and control plano topology is illustrated by Figure 1(a) and (b).

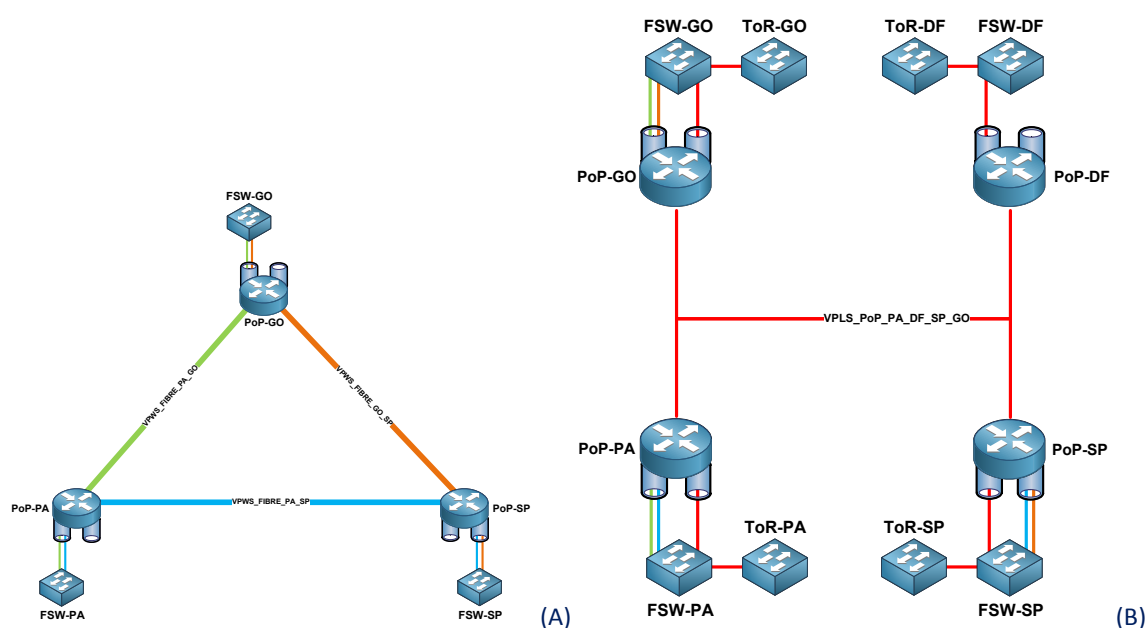


Figure 1 - Data and Control Plane Topology in 1e Phase

## 2.2 Topology 1p Phase

The topology of the 1p phase of the FIBRENet overlay network is depicted in Figure 2.




Figure 2 - FIBRE-BR overlay network topology in 1p phase

The data plane is used to switching traffic among participants of the experiments using an overlay point-to-point network. In this plane, the Flowvisor is the component responsible to segregate the experiments from each other, through the creation of circuits, where each circuit is determined by identifiers, known as VLAN IDs.

## 2.3 Deployment Test

The FIBRENet Tests aim to enhance the environment in the best possible way for making experiments and the performance evaluate of network, and create a troubleshoot base.

### 2.3.1 Types of Tests

#### 2.3.1.1 Experiments

##### a) Test 1

Name: Verifying the Establishment of Untagged Experiments.

Objective: Checks if the experiments creation on islands, without to use vlan tag, is functioning correctly and if the information is being configured on the switches.

##### b) Test 2


Name: Verifying the Establishment of tagged Experiments.

Objective: Checks if the experiments creation on island using vlan tag is functioning correctly and if the information is being configured on the switches.

##### c) Test 3

Name: Verifying the Establishment of tagged Experiments among island.

Objective: Checks if the experiments creation among islands using vlan tag is functioning correctly and if the information is being configured on the switches for each island.

	D2.6 Report on the deployment of the second version of the control and monitoring framework for the FIBRE-BR facilities	Doc FIBRE-D2.6
		Date 31/03/2014

### 2.3.1.2 *Performace*

#### a) Test 1

Name: Checking the numbers of flows are able to establish on FIBRENet OpenFlow switches.

Objective: Check the maximum flows that are able to establish in each switch of FIBRENet.

#### b) Test 2

Name: Checking the numbers of rules are able to establish on FIBRENet OpenFlow switches.

Objective: Check the maximum rules that are able to establish on forwarding table of each FIBRENet switch.

#### c) Test 3

Name: Checking the numbers of experiments are able to establish on FIBRENet environment.

Objective: Check the maximum experiments are able to establish on each FIBRENet switch.

### 2.3.1.3 *Connectivity*

#### a) Test 1


Name: Connectivity checking in VPWS channel between the islands of participants FIBRENet.

Objective: Verify connectivity using the Experiments VLAN among islands and also to measure the maximum rate allowed or achieved on this channel. Using packets flows generated by iperf.

#### b) Test 2

Name: Connectivity checking in VPLS channel between the islands of participants FIBRENet.

Objective: Verify connectivity using the Control VLAN among islands and also to measure the maximum rate allowed or achieved on this channel. Using packets flows generated by iperf

	D2.6 Report on the deployment of the second version of the control and monitoring framework for the FIBRE-BR facilities	Doc FIBRE-D2.6
		Date 31/03/2014


### 3 Multi-CMF Authentication using LDAP and SFA

Identity Management (IdM) is the set of processes and technologies used for guaranteeing the identity of an entity. IdM ensures the quality of identity information such as identifiers, credentials, and attributes and uses it for authentication, authorization, and accounting processes. Authentication procedures focus on confirming the identity of an entity that is, checking that an entity is who it claims to be. Authorization mechanisms define the access rights to resources associated to an identity. Authorization procedures describe these access rights to ensure that they are obeyed. Finally, accounting refers to track network resource consumption by users for capacity planning and billing [ITU-T 2009, Jensen 2012].

In recent years, the use of academic authentication and authorization (A&A) federations to control access to resources became popular [Dhungana 2013, Leskinen 2012]. In Brazil, for instance, academic researchers access scientific publication repositories using identities of the academic federation CAFe. Hence, there is no need to duplicate user information in local databases. On the other hand, there are applications and services whose access is restricted to certain members of specific institutions, such as the participants of an inter-institutional project. Those groups are called virtual organizations [Niinimäki 2004, Lee 2005]. FIBRE is an example of a virtual organization. In order to access distributed resources, partner institutions establish agreements. In this scenario, IdM appears as a strong requirement for establishing the trust environment among participants. Also, sharing tools and resources with mutual trust entities requires a federation. In FIBRE, we use SFA [SFA] to federate resources among different islands. In SFA-based testbeds, users supply their credentials to get access authorization to a set of resources located in different institutions, such as a set of computers and a minimal specified bandwidth.

Although SFA is the most important initiative to create a federation of FI testbeds, it presents open issues related to A&A. Briefly speaking, this occurs because its proposal is focused on interconnecting resources through a resource federation. The A&A ends up in background, composed only of a simple authentication mechanism based on X.509 certificates and static profiles. Hence, we proposed a new authentication and authorization method to federate SFA-based testbeds. Our proposal integrates A&A federations based on Shibboleth [Scavo 2005] or LDAP and SFA. Shibboleth implements the SAML standard [OASIS 2005] and also supports the Attribute-Based Access Control (ABAC) concept [Karp 2009]. Using ABAC, it is possible to implement more granular and dynamic access policies. Moreover, Shibboleth is used by the Brazilian academic federation, CAFe, and also by eduGAIN. Our proposal translates the SAML credentials generated by the academic federation in the SFA X.509 certificates after verifying whether the user has the required attributes of the virtual organization.

Our proposal was developed on MySlice portal, in order to integrate different authentication methods and SFA. MySlice is an FI testbed resource manager tool that provides an intuitive user interface and implements SFA to establish communication with resources.

	D2.6 Report on the deployment of the second version of the control and monitoring framework for the FIBRE-BR facilities	Doc FIBRE-D2.6
		Date 31/03/2014

We implemented the proposed A&A architecture using a real experimentation laboratory called GIDLab. GIDLab provides a mirror of the CAFE federation and of the FIBRE LDAP database, which serve as experimental environment for new applications that use the federation. This laboratory also offers virtual machines, in which we installed MySlice platform and some virtual testbeds. This implementation allowed us to validate the proposal and to evaluate the security features, comparing our scheme to other proposals.

### 3.1 Credential translation

Figure 3 illustrates a high-level view of our proposal. The key idea is to develop a mechanism to integrate the identity federation using Shibboleth or LDAP with the resource federation provided by SFA. Therefore, we introduced a mechanism to translate the credentials and check the attributes for access control.

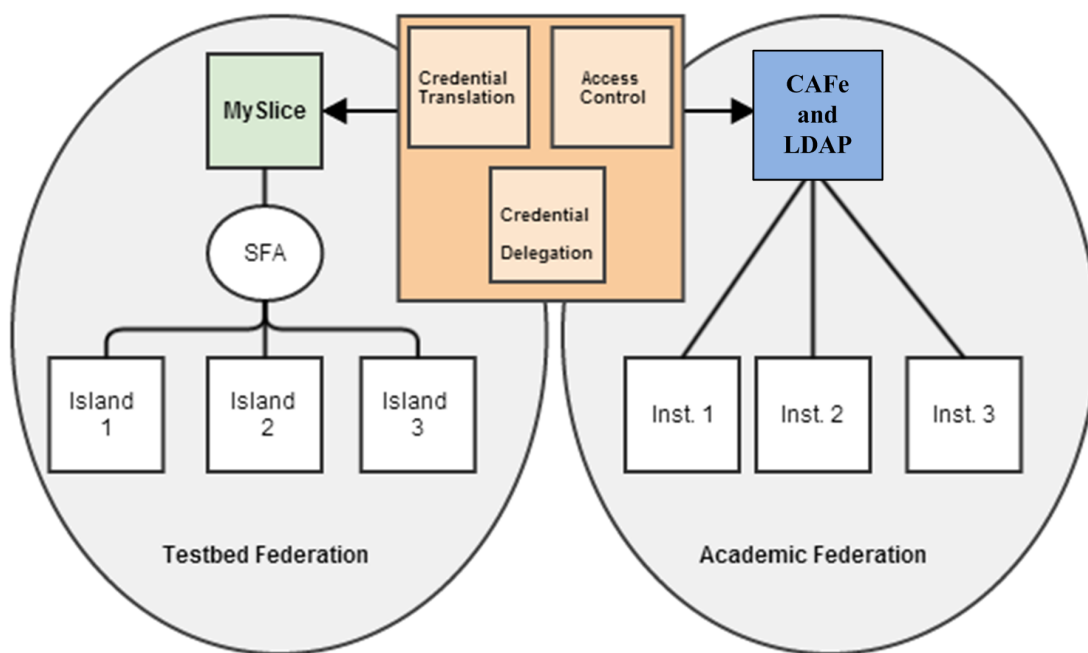



Figure 3 - High-level view of the federation integration proposal.

### 3.2 Credential Translation and Access Control Modules

In the context of this work, IdPs are represented by universities and research participants of the FI testbed federation. SPs are represented by the resource management tools, which require that users present their credentials to get access authorization to use testbed resources. In the implementation of our proposal, the SP which is responsible for allocating resources, authenticate and authorize users, and is represented by MySlice management web



	D2.6 Report on the deployment of the second version of the control and monitoring framework for the FIBRE-BR facilities	Doc FIBRE-D2.6
		Date 31/03/2014

tool (used by user to access the SFA testbed federation). A user may, for example, request the allocation of a number of computers and of a minimal bandwidth to connect them. MySlice must check whether the user has the required access rights to use those resources.

We propose a trust federation involving café/LDAP and the Brazilian testbeds bound together using a MySlice portal. We have implemented a credential translator that takes the authentication and authorization assertives emitted by CAfe or LDAP and generates the credentials expected by MySlice. Our work:

- was proposed to create a specific module for communication between federations, allowing validation tests of the solution (Credential Translation Module -- CTM);
- implementation is based on the new version of MySlice (Django);
- was integrated into various types based on the user identification and authentication model;
- used a module to delegate credentials (Credential Delegation Module -- CDM) and increase the user safety;
- used the ABAC concept in a real environment experimentation.

Figure 4 shows the architecture of the solution we propose. Basically, four entities are involved: the user interface (MySlice) with its user database (SQLite); CAFe Federation with the Discovery Service (DS) and the user's home IdP with its LDAP database; the communication tool with the testbed federation, implemented by Manifold, with its user database (SQLite); and testbed itself, represented in the picture by dummy PlanetLab, and its user database called SFA Registry (based on PostgreSQL).



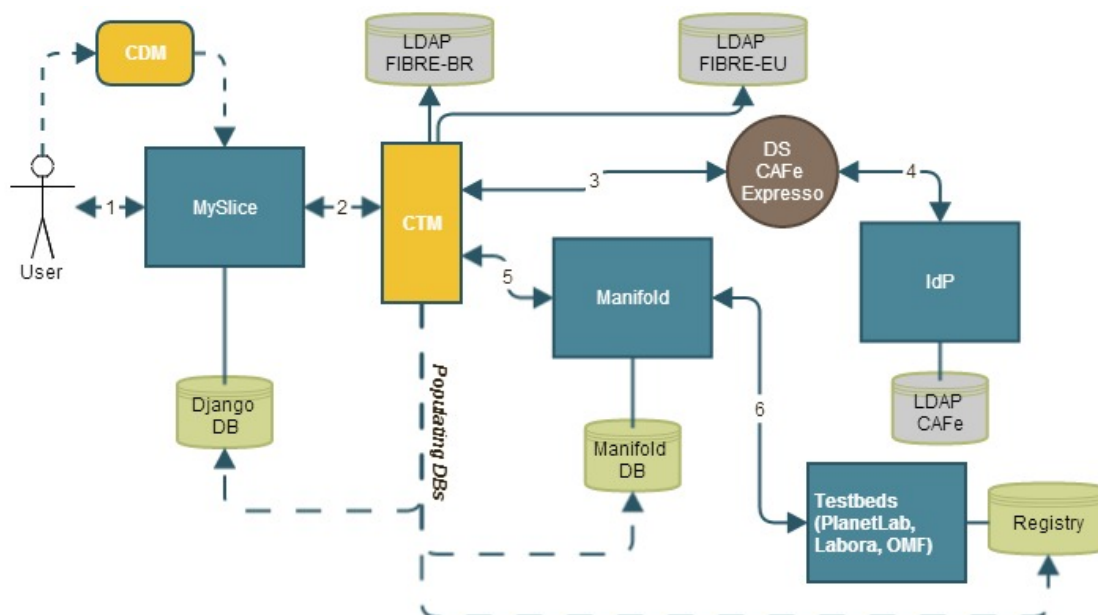



Figure 4 - Communication between the elements of proposed architecture.

Because this architecture is based on four different databases, two from MySlice -Django DB and Manifold DB-, one from SFA –the SFA Registry-, and the IdP LDAP database, a special attention is provided to synchronization. Our module synchronizes the databases from MySlice and SFA whenever a new user is added from CAFE or from FIBRE LDAP.

Figure 4 also shows the flow of events. First the user selects the CAFE authentication (1). Then, after passing the Credential Translation Module (CTM) (2), he is redirected to the discovery service of the A&A federation (3) to select his institution. After selecting his home institution, the user needs to authenticate in his IdP (4). If the authentication is successful and the access control module grants the permissions, the user can access the portal services and manage testbed resources (5 and 6).

### 3.3 Credential Delegation Module

In MySlice, there is an additional step in which the user must send, in a specific area of the portal, its private and public keys to generate SFA certificates, which represents a security threat. The main reason for this is that an SFA certificate is required for each resource that a user wants to use. Thus, when a user sends or generates (storing) his key pair in MySlice, this allows the portal to act on behalf of the user in the search and resources allocation process in the testbeds. Even though the private key should be known only by the user, passing the key to the portal avoids the need that the user locally generates self-certificates. This same process, although critical since the private key is stored in the user database, has been

	D2.6 Report on the deployment of the second version of the control and monitoring framework for the FIBRE-BR facilities	Doc FIBRE-D2.6
		Date 31/03/2014

adopted by other FI associations, such as GENI [ProtoGeni]. In order to eliminate this vulnerability, the user would need to generate credentials of each resource using an external tool (run in BASH) to delegate their credentials to MySlice. This solution, nevertheless, is not easy for the user.


To overcome this issue, we propose the Credential Delegation Module (CDM). This module runs in the user machine in order to generate the credentials and send them to the portal in a transparent way to the user. This transparency is important because the credential delegation is performed every time a new resource is created or when the credentials of an existing resource expire. Indeed, usually resource credentials have a short lifetime, necessary only to carry out the desired experiments, which means that the user would need to generate certificates many times.

Using the proposed module, the user loses flexibility, because he needs to install a program in his machine. On the other hand, the security is increased without making it difficult to use the federated testbeds.

### 3.4 Implementation and Results

The study was conducted using a real experimental laboratory (GIDLab), which supports the creation of virtual machines (VM) and where mirrors of the CAFé federation, FIBRE LDAP-BR, OFELIA LDAP, and of the MySlice platforms are installed. Thus, it was possible to install all three architecture components as shown in Figure 4 and create a PlanetLab dummy testbed above the SFA federation.

After the authentication, CTM calls the access control module, which is based on attribute verification. In the current implementation, only an ABAC proof of concept was developed, using the CAFé attribute affiliation. We check this attribute to see if the user is a student, otherwise he will not have access to the FI testbed. The access control of our architecture can be viewed in the high-level description of system verification in Algorithm 1.

	D2.6 Report on the deployment of the second version of the control and monitoring framework for the FIBRE-BR facilities	Doc FIBRE-D2.6
		Date 31/03/2014

---

**Algorithm 1: Steps of A&A proposal.**

---

**Input:** User attributes.  
**Output:** A&A – access allowed or denied.

```

1 if affiliation  $\neq$  student then
2   | access denied;
3   | if User exist in MySlice then
4   |   | perform A&A for this user;
5   | else
6   |   | register user in MySlice/Manifold and SFA
   |   | Registry;
7   | end
8 end

```

---

Algorithm 1 - Steps of A&A proposal.

According to the user profile, three different actions are provided by CTM:

1. CAFé user already exists in MySlice database and is automatically logged in the system.
2. CAFé user does not have permission to access MySlice as an ABAC decision.
3. CAFé user does not exist in MySlice but he is allowed to access the testbed. Therefore, his account is automatically created in MySlice, Manifold and SFA databases.

Considering item 3, before adding a new user in the databases, it is necessary to associate him to an HRN (Human Readable Name), and allow him to generate his key pair through an administration interface account. This is currently supported by MySlice when CTM is used. But if the user does not want to generate his key pair, he can send his public key and delegate all SFA credentials to MySlice, using the CDM (Credential Delegation Module). CDM allows the user to store his private key only in his machine. In current implementation, the certificates generated using CDM must be loaded in the web portal by the user. A next step will be to integrate the automatic transmission of the credential to the web service.

We know the efforts in the context of SFA and GENI for developing tools to perform the authentication and access control through the use of Shibboleth. However, so far we have not found this intention as an available tool. We therefore believe that this work has a great potential for integration into the FIBRE project and can be extended to all other research initiatives in FI, since we are respecting the premises of open-source.

## 4 Single Logon Web Portal Customization

From the point of view of MySlice, it is possible to authenticate a FIBRE user through three different methods, CAFe academic federation, FIBRE-BR LDAP, and FIBRE-EU LDAP (based on OFELIA LDAP database), as show in Figure 5.

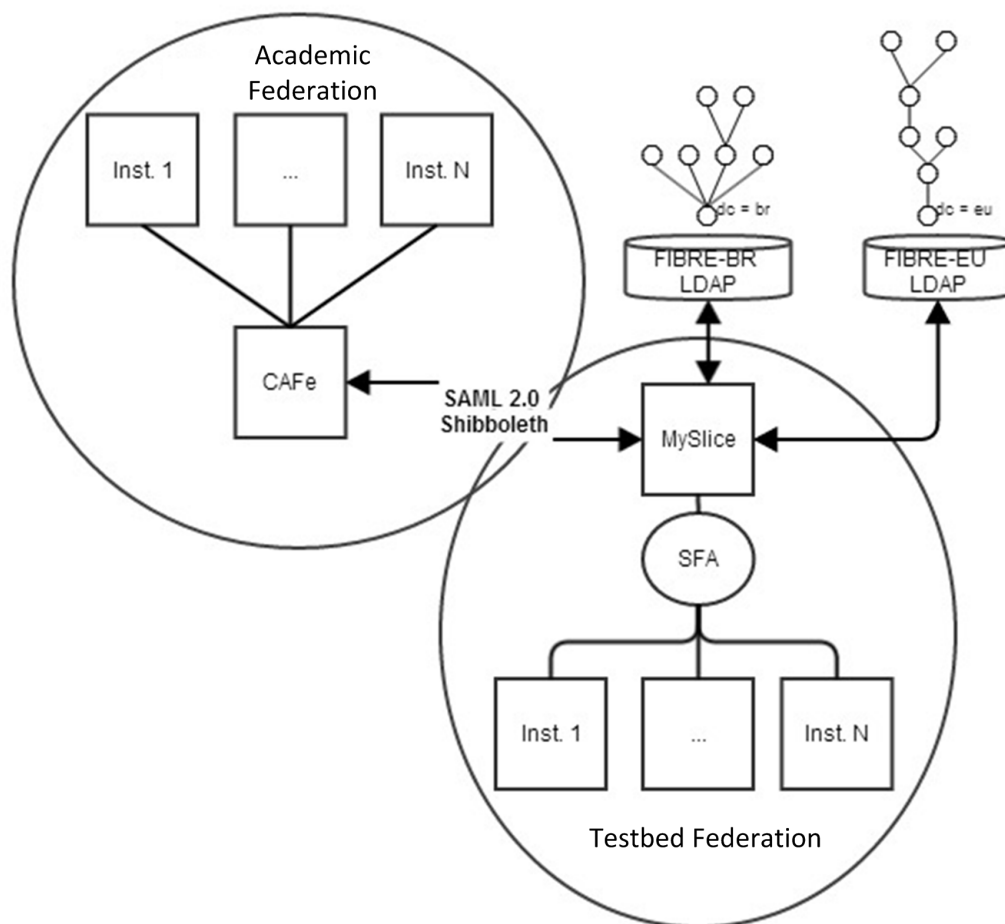

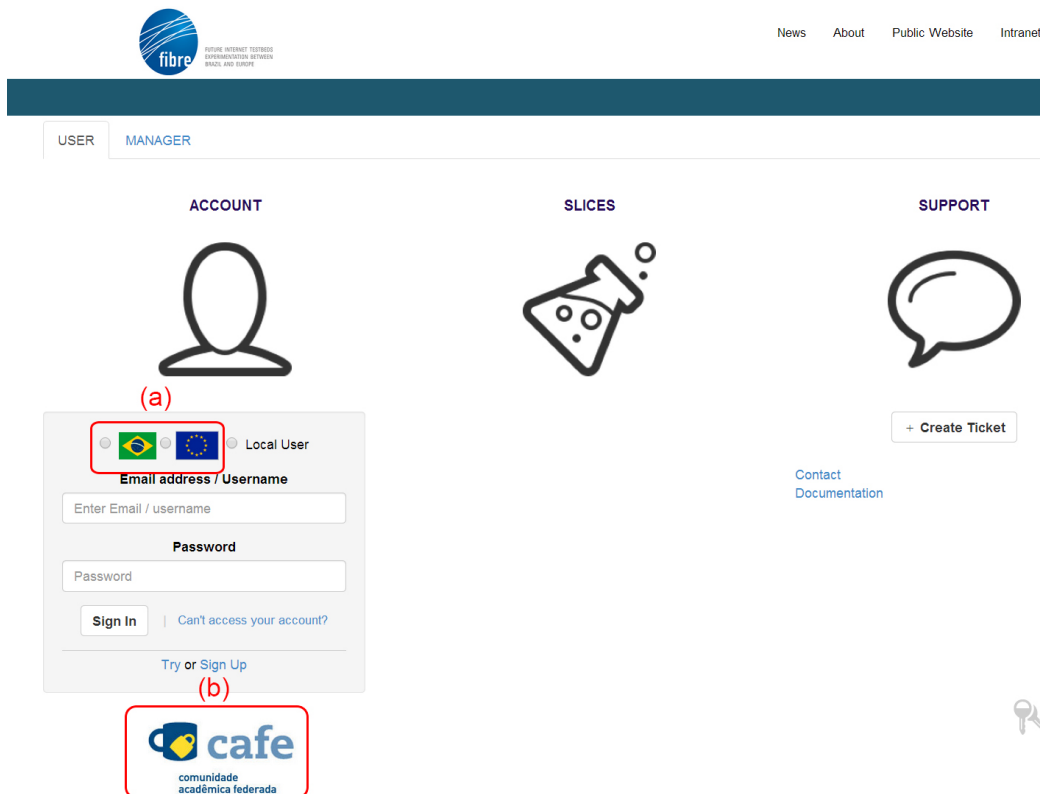


Figure 5 - Complete authentication scheme using MySlice.

The CAFe and LDAP authentication options can be seen in the front page of FIBRE MySlice Portal, as show in Figure 6. When a user selects CAFe as the authentication mechanism, highlighted as (b) in the Figure, he is automatically redirected to CTM, responsible for redirecting, collecting, and processing the identity federated authentication from CAFe. The Credential Delegation Module is used after this process to generate all the required credentials without providing the user's private key to the portal. If the user is registered in the FIBRE LDAP database, he must select either in Brazil or Europe LDAP and then provide his login and password.

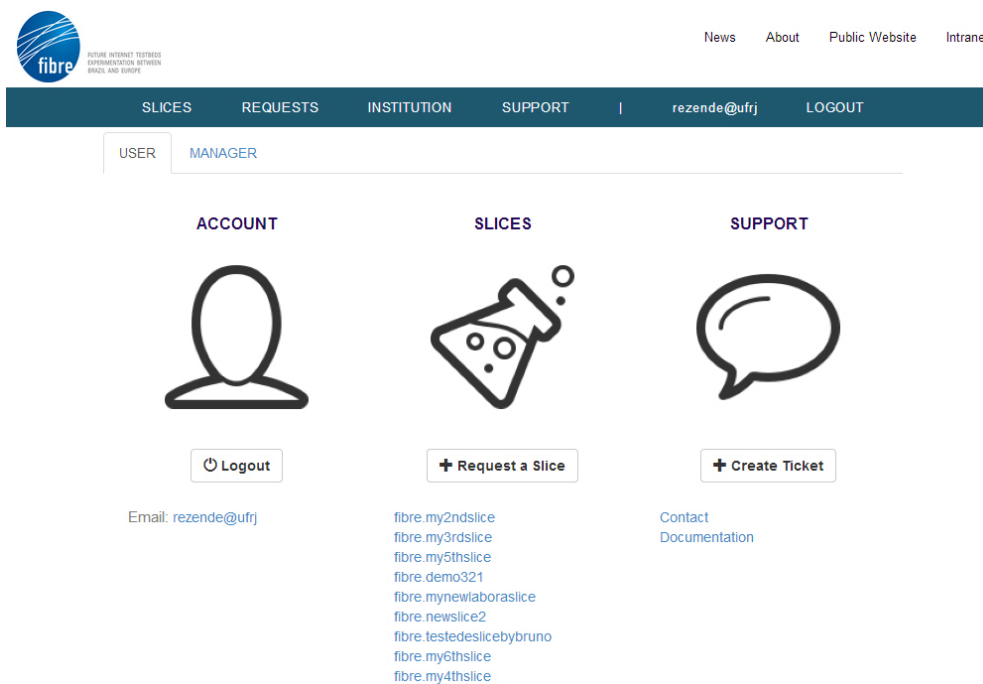
	D2.6 Report on the deployment of the second version of the control and monitoring framework for the FIBRE-BR facilities	Doc FIBRE-D2.6
		Date 31/03/2014



The image shows the MySlice login interface. At the top, there's a header with the 'fibre' logo and navigation links: News, About, Public Website, Intranet. Below this is a dark blue bar with 'USER' and 'MANAGER' tabs. The main content area has three columns: ACCOUNT, SLICES, and SUPPORT. The ACCOUNT column features a user icon, a login form with fields for 'Email address / Username' and 'Password', and a 'Sign In' button. A red box labeled '(a)' highlights the 'Local User' option and the 'Email address / Username' field. Below the login form, a red box labeled '(b)' highlights the 'CAFé' logo, which is the 'comunidade acadêmica federada' logo. The SLICES column has a flask icon. The SUPPORT column has a speech bubble icon and a '+ Create Ticket' button. At the bottom right, there's a 'Contact Documentation' link and a key icon.


Figure 6 - MySlice login screen with the CAFé authentication option.

After these steps, the user is authenticated and is able to access the testbed resources through MySlice, as shown in Figure 7.



The image shows the MySlice dashboard after authentication. The header is the same as in Figure 6. Below the dark blue bar, there's a navigation bar with links: SLICES, REQUESTS, INSTITUTION, SUPPORT, and a user profile section showing 'rezende@ufrj' and a 'LOGOUT' button. The main content area has three columns: ACCOUNT, SLICES, and SUPPORT. The ACCOUNT column shows a user icon, a 'Logout' button, and the email 'rezende@ufrj'. The SLICES column shows a flask icon, a '+ Request a Slice' button, and a list of testbed resources: fibre.my2ndslic, fibre.my3rdslic, fibre.my5thslic, fibre.demo321, fibre.mynewlaboraslice, fibre.newslic2, fibre.testedeslicbybruno, fibre.my6thslic, and fibre.my4thslic. The SUPPORT column shows a speech bubble icon, a '+ Create Ticket' button, and a 'Contact Documentation' link.

Figure 7 - Testbed resources shown by MySlice.

	D2.6 Report on the deployment of the second version of the control and monitoring framework for the FIBRE-BR facilities	Doc FIBRE-D2.6 Date 31/03/2014
---	---	-----------------------------------

## 5 Monitoring Infrastructure and Experiments

The FIBRE-BR monitoring architecture is composed of two dimensions: network/infrastructure monitoring and experiment monitoring. In this report, we are going to focus on the deployment status of the network/infrastructure monitoring tools, which include the perfSONAR Performance Toolkit servers and ZenOSS, and the current experiment monitoring support.

### 5.1 perfSONAR Performance Toolkit Deployment

In the context of the FIBRE-BR testbed, the knowledge of the delay, available bandwidth and active routes among each island will be extremely important to the experimenter and to the Network Operation Center (NOC). With this information, the experimenter will be able to choose the resources that better reproduce the scenario expected for his/her experiment, and also latter correlate the results of his/her experiment with the state of the network.

In order to fulfill this requirement, we deployed instances of the perfSONAR Performance Toolkit on each island of our testbed. The perfSONAR Performance Toolkit (pSPT) [Internet2 2013] is a customized CentOS image with a comprehensive set of pre-installed and configured measurement and monitoring tools. Once installed and set a few network parameters, the pSPT is ready to perform, store, and display measurements like One-way delay (OWD), available bandwidth, and traceroute, just to name a few.

The adopted deployment strategy was the installation of two pSPT servers on each island, one for latency and loss-related measurements and the other for bandwidth-related measurements. The use of two separate servers is necessary in order to avoid interference between the bandwidth and latency measurements. Figure 8 shows the current status of these servers among the FIBRE-BR islands, most of them fully operational and minor islands have faced some hardware's issues. The function of the additional server named "ZenOSS" depicted in Figure 8 will be explained in the next subsection.

Each of these two pSPT servers is performing delay, throughput and traceroute tests with each one of its counterparts on the other islands. Therefore, we have "full mesh" measurement data for the aforementioned metrics, which provides a complete "weathermap" of the testbed's network.

In order to graphically present this data to the users, we deployed the perfSONAR Monitoring and Debugging Dashboard (MaDDash) on the NOC facility. This tool provides a global view of the network performance and serves as the entry point for accessing the measurements collected by the pSPT servers. Figure 9 shows a screenshot of our current installation, collecting measurements from the FIBRE-BR islands.

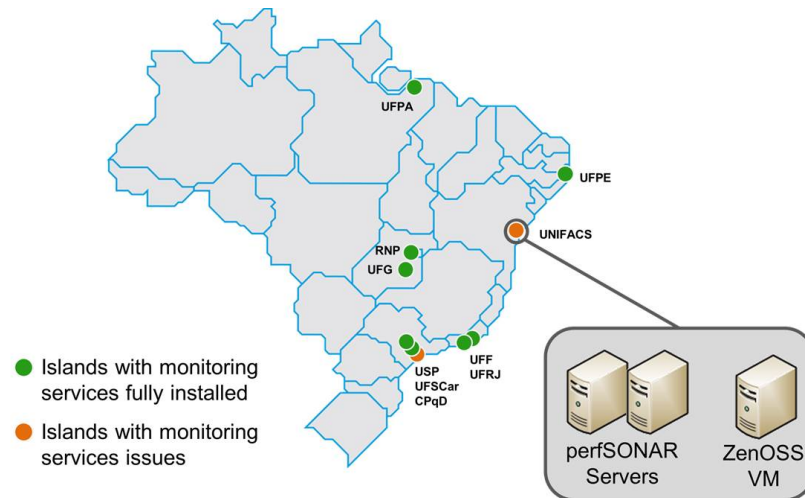


Figure 8 - Monitoring infrastructure status

## FIBRE-BR Dashboard

### OWAMP




Figure 9 - MaDDash homepage: a dashboard with measurements for each node pair



## 5.2 ZenOSS Deployment

The basic monitoring infrastructure of each island uses Zenoss Core v4.2 [Zenoss 2013]. This open source management tool provides a wide collection of advanced resources to collect, display and analyzes measurement data collected via SNMP, SSH and WMI (Windows Management Instrumentation). Zenoss is deployed in each island alongside the pSPT servers and a central one is hosted in the Network Operation Center (NOC). However, while the latter are being installed in “physical” hardware, the former is being installed as a VM on the IBM server.

A view of the island infrastructure’s topology through Zenoss can be seen on Figure 10.

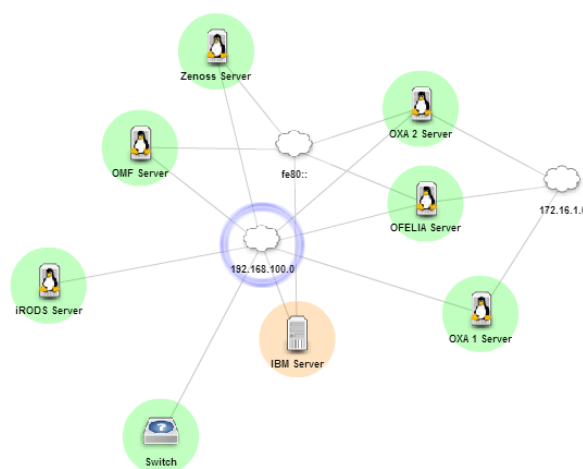


Figure 10 - ZENOSS view of an infrastructure’s topology

In effect, we are using a set of Zenoss extensions (ZenPacks). One of them is the “Xen Virtual Hosts Monitor”, which delivers monitoring information regarding the Xen hypervisor, used at the main server (IBM Server); and the other one is the “Interface Graphs” extension, which delivers more detailed graphs of the NICs (Figure 11).




Figure 11 - Detailed Graphic of NIC's Throughput





In actual implementation, a refined solution to deal with the centralized access to the distributed Zenoss servers is taken. Although each island can monitor its own resources, NOC is responsible for aggregating all monitored resources from all islands and through events, alerts mechanisms forward issues to the island managers to solve the possible errors. Figure 12 depicts a list of monitored resources from NOC, grouped by islands.


Device	IP Address	Device Class	Production State	Events
ibm.cpod.fibre.org.br	10.12.0.30	/Ping	Production	1
ibm.noc.fibre.org.br	10.0.0.30	/ServerLinux	Production	
ibm.rnp.fibre.org.br	10.9.0.30	/Ping	Production	1
ibm.ufpa.fibre.org.br	10.10.0.30	/Ping	Production	
ibm.ufpe.fibre.org.br	10.3.0.30	/Ping	Production	
ibm.ufsc.fibre.org.br	10.5.0.30	/Ping	Production	
ibm.ufes.fibre.org.br	10.1.0.30	/Ping	Production	
ibm.ufscar.fibre.org.br	10.8.0.30	/Ping	Production	1
ibm.unifacs.fibre.org.br	10.4.0.30	/ServerLinux	Production	1
ibm.usp.fibre.org.br	10.6.0.30	/ServerLinux	Production	1
ldap.noc.fibre.org.br	10.0.0.50	/ServerLinux	Production	1
ldap.uff.fibre.org.br	10.7.0.50	/Ping	Production	
ldap.usp.fibre.org.br	10.6.0.50	/ServerLinux	Production	1
mlf.noc.fibre.org.br	10.0.0.3	/ServerLinux	Production	2
mon.noc.fibre.org.br	10.0.0.80	/ServerLinux	Production	1
mon.usp.fibre.org.br	10.6.0.80	/ServerLinux	Production	1
netfpga1.usp.fibre.org.br	10.6.0.10	/Ping	Production	
netfpga2.usp.fibre.org.br	10.6.0.11	/ServerLinux	Production	2
netfpga3.usp.fibre.org.br	10.6.0.12	/ServerLinux	Production	2
ocf.noc.fibre.org.br	10.0.0.100	/ServerLinux	Production	1
ocf.usp.fibre.org.br	10.6.0.100	/ServerLinux	Production	1
omf.noc.fibre.org.br	10.0.11.200	/ServerLinux	Production	1
perfsonar.noc.fibre.org.br	10.0.0.60	/ServerLinux	Production	1
syslog.noc.fibre.org.br	10.0.0.81	/ServerLinux	Production	1
vpn.noc.fibre.org.br	10.0.0.70	/ServerLinux	Production	1
vpn.usp.fibre.org.br	10.6.0.70	/ServerLinux	Production	1

Figure 12 - Zenoss Monitored Infrastructure

## 5.3 Experiment Monitoring Support

For experiment monitoring support, in this second version of our CMF we implemented the use case 4 (Inter CMF) presented in Deliverable 4.4, where a common measurement tool is deployed in selected nodes at a given (possibly federated) slice.

Use case 4 involves the use of two components. The first one is the Orchestration Service, which will (1) install and configure a common tool on the endpoints of a measurement – a process we call instrumentation – and (2) provide an API where any kind of GUI can request the setup of a measurement between the measurement points. The second component is a set of gateways deployed across the slice to register the measurement results in the Measurement Information (MI) Service and make them available through a perfSONAR-like

	D2.6 Report on the deployment of the second version of the control and monitoring framework for the FIBRE-BR facilities	Doc FIBRE-D2.6
		Date 31/03/2014

API. We will first present our Orchestration Service implementation, then the gateway implementation, and finish showing a complete workflow of this use case.

The Orchestration Service acts as an intermediary between the GUIs and the slices, as shown in Figure 13. It saves the GUIs from the hassle of configuring and managing (possibly) inter-CMF measurements, which tends to be a complex task, worth of a specialized service. The Orchestration Service abstracts all this complexity behind a REST API.





Figure 13 - The Orchestration Service as an intermediary between GUIs and Slices

Table 1 presents a simplified description of the methods offered by the Orchestration Service REST API and their functions. The first thing a GUI needs to do is register the slice in the Orchestration Service using the `addSlice()` method. After that, the GUI can ask the Orchestration Service to prepare a node as a Measurement Point using the `addMNode()` method. This preparation (instrumentation) consists of installing a set of tools and services responsible for performing measurements and enabling remote configuration.

Table 1 - The Orchestration Service REST API (simplified).

Method	Parameters	Description
<code>addSlice()</code>	sliceName userName userEmail	Registers a slice in the Orchestration Service.
<code>addMNode()</code>	sliceID nodeName nodeIP	Asks the Orchestration Service to prepare a node (instrumentize) so it can perform measurements and be remotely configurable.
<code>addMeasurement()</code>	sliceID type sourceIP destinationIP parameters	Asks the Orchestration Service to set up a measurement between two nodes ( <i>sourceIP</i> and <i>destinationIP</i> ). The <i>type</i> parameter identifies which kind of measurement should be performed (currently we support one-way delay and available bandwidth measurements, but <i>traceroute</i> and two-way delay can also be easily implemented) and the <i>parameters</i> parameter (a key-value list) specifies details about the measurement such as periodicity,

	D2.6 Report on the deployment of the second version of the control and monitoring framework for the FIBRE-BR facilities	Doc FIBRE-D2.6 Date 31/03/2014
--	---	-----------------------------------

		the protocol to be used (UDP or TCP), etc.
--	--	--

However, FIBRE is a multi-CMF environment, and each CMF usually provides nodes with different operating systems, or different versions of the same operation system. This way, there is no one size fits all approach possible when comes to instrumenting nodes. This led us to implement the instrumentation function in the Orchestration Service following a plugin-based architecture – one plugin for each environment (operating system, or a specific version of an operating system).

In the current implementation, every plugin must install two specific tools during the instrumentation process: the perfSONAR-PS BUOY service and the perfSONAR Mesh Configuration service. The perfSONAR-PS BUOY service (henceforth pS-BUOY) acts as the common measurement tool that will make tests between nodes from different CMFs possible.

The second tool, the perfSONAR Mesh Configuration service, acts as a remote configuration agent. The Mesh Configuration service usually works as follows: the network operator writes a configuration file describing all tests he wants to perform in the network and publishes this file in some web server. Alongside each Measurement Point (MP) there is a Mesh Configuration agent who periodically fetches this configuration file, identify which tests this MP must perform, and configure the measurement tools – such as pS-BUOY – accordingly.

Since each slice on a testbed is an independent overlay network, we used the Mesh Configuration service to manage all inter-CMF measurements inside a slice. Each slice has its own configuration file and all Mesh agents inside the slice follows the definitions of this file. This way, the Orchestration Service is able to remotely configure the pS-BUOY instances running inside any slice.

Now, let us get back to usage workflow of the Orchestration Service API. After instrumenting at least two nodes using the addMNode() method, a GUI can call the addMeasurement() method to configure a measurement in the slice. The Orchestration Service manages the measurements on each slice by creating and editing Mesh configuration files in JSON and saving them on a web server. Each modification on this slice configuration is saved on this file, and is automatically detected by the Mesh agents. When a Mesh agent identifies that the node it is responsible for should perform a measurement to some other node, it automatically configures pS-BUOY to do so, and the measurement is thus started.





Figure 14 – The Orchestration Service Complete Architecture

Figure 14 shows a more detailed view of the Orchestration Service architecture. The service's two modules behind the REST API are: the Configuration Editor, responsible for managing the JSON configuration files and publishing them in a web server (currently we use an Apache server, and the Instrumentation Manager, which coordinates the instrumentation plugins, choosing the right one depending of the target environment. We implemented these two modules and the REST API using the Python programming language and the Flask web framework.

	D2.6 Report on the deployment of the second version of the control and monitoring framework for the FIBRE-BR facilities	Doc FIBRE-D2.6 Date 31/03/2014
--	---	-----------------------------------

## 6 Modifications of OCF and OMF Software for Second Version

In this section, we describe the main modifications made on OCF and OMF software, in order to support, the second version of FIBRE CMF. The FIBRE Second Version CMF is the combination of these software pieces along with new customization for the Brazilian specific requirements (LDAP user authentication, lack of a stable, in terms of energy and connectivity, infrastructure, insecure network, etc).


### 6.1 OCF Modifications for Second Version

We start by describing the modifications done specifically in the OCF CMF. The OCF is the SDN based framework as well as the virtualization framework, and some of the issues were related to the difficulty in provide authentication per project in the OCF using the proposed specific Brazilian LDAP scheme, that is different from the OFELIA LDAP scheme. In order to accomplish this task, the UFSCar team developed an alternative method of authentication per VM using the VM Template as basis and preparing the internal configuration to authenticate the SSH access through LDAP. The second contribution in terms of software components for OCF was the Auto-Recovery procedure added to the FIBRE OXA installation. This customization feature was added due to a typical problem in Brazilian infrastructure at universities and labs, usually there is a lack of energy UPS. Thus, in places like UFRJ and UFSCar in 2013 that went through institutionally over major power grid reformulations, they experience several cases of FIBRE islands machine being shut off suddenly. Then, the OXA software based on Xen kernel would produce phantom VM machines and some unbootable ones. Thus, the UFSCar team added the recovery procedure already deployed on most FIBRE islands. A final contribution on the OCF CMF was done in terms of adding extra layer of security to the OFELIA CMF, using Linux Firewalls and NATs to provide Internet access to VMs.

#### 6.1.1 LDAP Integration on OCF VM Template

One of needs of the project in point of view of researcher was the personal login in each virtual machine created by OXA. In this way, the researcher can log using SSH in each virtual machine created with your own login, which was previously created on OCF portal. Thus, it excludes the need of pattern password for all slices, for this procedure, the LDAP local or remote service should be functional.

The original OCF template (based on the OFELIA official template) can be used to create virtual machines but doesn't treat the point of personal use for login. Therefore, a new template was made including the necessary packets for direct authentication and only for the LDAP server. Now, different experiments can only be access by the researches authorized to it, like a virtual machine of another experiment. All changes were made directly on default image on OXA, in tar.gz format, changes like the installation of basic packets to LDAP such as:

	D2.6 Report on the deployment of the second version of the control and monitoring framework for the FIBRE-BR facilities	Doc    FIBRE-D2.6  Date    31/03/2014
---	---	---

- libpam-ldap
- nscd
- libnss-ldap

The respective configuration files of LDAP point to DNS local address for login request. Therefore, the OCF portal participate with same configuration for access, narrow the login only for user who already registered on LDAP server. The following screenshot (Figure 15 and 16) presents this feature. As a matter of fact, the OFELIA OCF have the functionality of separate the researches by project, but this functionality was not explored yet, waiting for final unified scheme of LDAP between EU and BR.

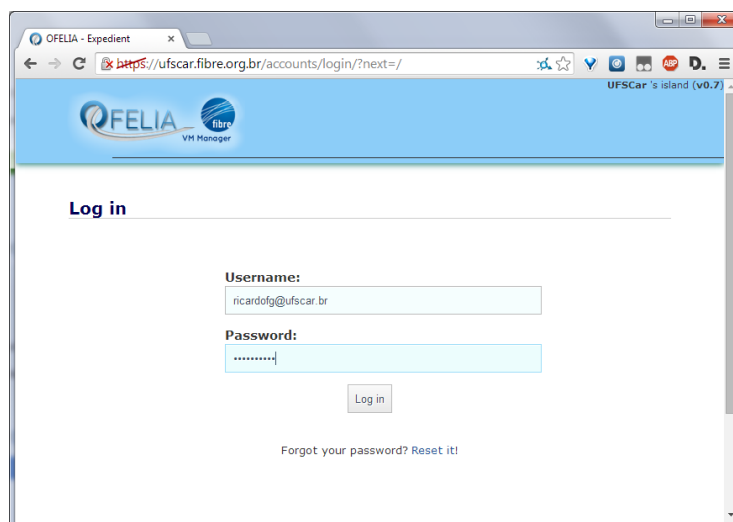


Figure 15 - OCF Login Authentication using FIBRE BR LDAP

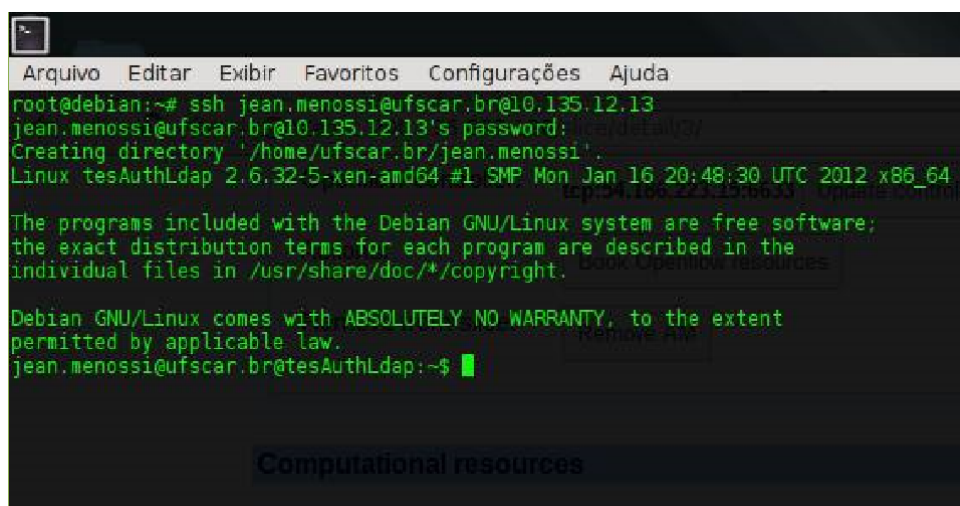



Figure 16 - SSH Authentication using jean.menossi@ufscar.br registered in LDAP

	D2.6 Report on the deployment of the second version of the control and monitoring framework for the FIBRE-BR facilities	Doc FIBRE-D2.6
		Date 31/03/2014

### 6.1.2 Adding Robustness to OXA by Auto-Recovery

The auto-recovery for essential island services is important in case of failures or an emergency stop. This procedure was developed by UFSCar team and it is fully documented in the FIBRE BR Wiki page<sup>2</sup>. It is considered optional for the islands to deploy; although it is highly recommended due to Brazilian islands do not have UPS installed. These services should be started when the operating system boots and keep running for the island's lifetime; whereas experimenter resources do not depend on such warranties (especially in 'red button' stop).

In order to ensure the main operating system (Dom0) always have at least 1GB RAM available, it is recommended not to use memory ballooning for the Dom0 in XEN's configuration files. This measure prevents administrative access issues in case of over-provisioning of virtual machine resources.

A possible solution for the service VMs to start automatically at boot time is to create an init script to manage virtualization. While flexible, this solution requires a high degree of scripting to achieve basic functionality. Obtaining state information correctly, to support management decisions, is also an issue with this approach.

Our solution includes the use of XEN stack tools to manage these virtualization resources. To this end, it is possible to use the "xendomains" daemon. Further customization of this daemon's settings is done by editing the /etc/defaults/xendomains file. Domains managed by XEN itself are referred as "managed domains" in this approach. Managed domains require the use of different primitives for management with the XM stack, such as using "xm new" instead of the more known "xm create" command.

Therefore, using managed domains enables resource management to be done with a higher lever interface. For ease of use, it is possible to deploy and use the Libvirt interface as a management tool. The Libvirt stands out from other solutions due to the large number of hypervisors supported in a homogeneous API, and the array of features and tools already implemented on top of this library. Once installed, almost none configuration is needed before use. Virtual machines can then be set to autostart using virsh, or any other Libvirt-enabled application.

After researching possible solutions in order of abstractions and features, it was opted to use either xendomains or preferably Libvirt to manage auto-recovery for essential island services. According to analysis of UFSCar internal logs, in the period of 4 months after the auto-recovery had been deployed, we observe 4 failures that were recovered using the proposed auto-start. If the system was not in place, the VMs would not been up all this time, and manual recovery would be necessary.

<sup>2</sup> <https://wiki.rnp.br/pages/viewpage.action?pageId=81428782>



### 6.1.3 Adding Security and Network Connectivity to OXA VMs

It was developed a study case, using the OCF OXA machine at the UFSCar island that provides an augmented security for the island perimeter. This is done by configuring in the main server that is physically connected to the Top of the Rack switch, the traffic that is coming from and to the Internet reaches 3 different internal bridges. Those bridges are:

- br0 (Internet)
- control (Internal Control Network)
- exp (Experimentation Network)

According to the figure 17, all the “service” virtual machines, the ones considered to support the production of the island, like the DNS server, FlowVisor, etc, are virtualized inside the IBM server (Servidor IBM) and thus they participate in the control and br0 network segments. Usually, the machines created by the experimenter can participate only on the “exp” and “control” network segments. Thus, our approach is to let the “control” bridge give temporary access to the Internet, such that the experimenters can download packages and software for their experiment. This introduction is important to understand the Firewall configurations, we develop the script based on the Linux iptables.





Figure 17 - Internal Bridge br0 access to the Internet



	D2.6 Report on the deployment of the second version of the control and monitoring framework for the FIBRE-BR facilities	Doc FIBRE-D2.6
		Date 31/03/2014

The Firewall script is an internal development, an early version can be found here<sup>3</sup>. We used some initial variables to define the services and thus reconfigure the firewall based on each island. These are:

OCF Portal <https://EXTERNALIP>

VT\_Manager <https://EXTERNALIP:8443>

Optin <https://EXTERNALIP:8445>

Nitos <https://EXTERNALIP:18443>

Perfsonar 1 <http://EXTERNALIP:3180>

Perfsonar 2 <http://EXTERNALIP:3280>

According to the firewall configurations, we determine the following chain of events:

- 1) Default policies for the default chains
  - a. forward drop
  - b. input drop
  - c. output accept
- 2) Basic Access
  - a. IBM server accessing local loopback
  - b. SSH access through port 22
  - c. Allowing Experimenters VMs to access the FlowVisor on port 8080
  - d. Allowing test IBM ping reachability from FIBREnet
  - e. Allowing internal communication OXA IBM and control network and Internet
  - f. Allowing flowvisor access to control network
  - g. Allowing zenoss access to control network
- 3) Pre-Routing and NAT
  - a. Treating Open Connections to be accepted
  - b. Allowing Forwarding to OCF Modules (all of them) to be accessed
  - c. Allowing ZenOSS to be accessed from Internet
  - d. Allowing VPN to be accessed from Internet
  - e. Allowing NITOS to be accessed from Internet
  - f. Allowing LDAP to be accessed from Internet
  - g. Allowing PerfSonar to be accessed from Internet
  - h. Doing the Reverse Proxy and NAT
  - i. NAT from Internet to IBM Server to default services

4) PostRouting the FIBREnet traffic to Internet using NAT

5) Enabling VM OCFs access to the Internet

```
for interface in `brctl show | grep vif | sed -e "s/\t//gi"`; do
$ipt -I FORWARD -m physdev --physdev-in $interface --physdev-is-bridged -j ACCEPT
$ipt -I FORWARD -m physdev --physdev-out $interface --physdev-is-bridged -j ACCEPT
done
```

---

3

[https://s3.amazonaws.com/prod\\_object\\_assets/assets/9588544869955/Firewall.pdf?AWSAccessKeyId=AKIAI7NUHQYARXR2GGCQ&Expires=1396278367&Signature=TulaLSxmhYl6cXLHcPM6jZrarAE%3D](https://s3.amazonaws.com/prod_object_assets/assets/9588544869955/Firewall.pdf?AWSAccessKeyId=AKIAI7NUHQYARXR2GGCQ&Expires=1396278367&Signature=TulaLSxmhYl6cXLHcPM6jZrarAE%3D)

## 6.2 OMF Modifications for Second Version

In [Report D2.5], we presented the details about the NS\_OMF-BR, i.e. the Brazilian version of the CMF composed of NITOS Scheduler and OMF. As described in [Report D2.2], both software systems are important because OMF alone does not provide resource discovery and allocation, and user authentication and authorization, which are covered by NITOS Scheduler. However, nor the original NITOS Scheduler neither OMF 5.4, deployed in Brazilian islands, are ready to federate on standard way, by using SFA [SFA]. In the first Brazilian version of the CMF, we have employed a basic federation approach based on peering of XMPP servers using Server 2 Server protocol.

As the second Brazilian version of the CMF, we have completed redesigned the system for resource discovery and allocation, and user authentication and authorization, which we named LABORA Scheduler or LS. By design, LS is SFA-ready and driver has been developed for proper federation. Besides, OMF 5.4 has been also changed in order to deal with federation. We still kept the basic federation based on peering of XMPP servers, but this implementation is used in an administrative way, i.e. the only testbed operators need to worry about and only during the initial deployment. In the following we describe the LABORA Scheduler SFA Compliant.

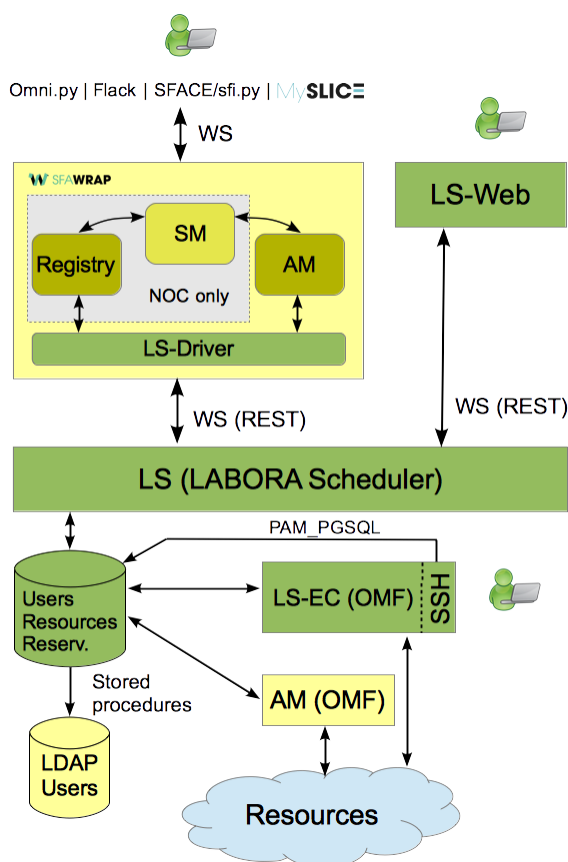


Figure 18 – LABORA Scheduler architecture


	D2.6 Report on the deployment of the second version of the control and monitoring framework for the FIBRE-BR facilities	Doc FIBRE-D2.6
		Date 31/03/2014

Figure 18 presents the software architecture of the LABORA Scheduler. The core module, named LS (LABORA Scheduler), offers all functions needed for resource discovery and allocation, and user authentication and authorization. As example, some important functions are: `get_resources()`, `get_slices()`, `add_slice()`, `add_user_to_slice()`, `add_resource_to_slice()`, `authenticate_user()`, `get_users()`, etc. It is worth noting that the SFA concepts, such as resource and slice, are already present in the core of LS. All functions are accessible by a RESTful API, making easy to plug it into different front-ends.

A Web-based front-end, named LS-Web, has been developed to offer basically the same interface developed in the first Brazilian version of the CMF. Actually, there are some minor improvements related to the look-and-feel, as illustrated in Figure 19.



Figure 19 – Web interface of the LABORA Scheduler

The first Brazilian version of the CMF operated with a MySQL database with some replicated information due to the need to access information from the scheduler and from OMF. By inference, originally, the redundancies were not treated. However, in the second Brazilian version of the CMF, we have made a redesign of the data scheme, as described in Figure 20. Now, both LS and OMF access a unique PostgreSQL database and can consistently follow any relationship. Again, it is worth noting that the data scheme has also been designed with federation in mind. The choice of the PostgreSQL was based on the mature implementation of the store procedures, which have been used to keep updated the user information in the LDAP

database. Naturally, OMF EC (Experiment Controller) and AM (Aggregate Manager) has been modified to connect to the PostgreSQL database and to properly handle the new data scheme. This modification is especially important for the OMF EC because it employs the database to enforce the reservation control. The login/SSH access is now based on PAM\_PGSQL, a PAM module that authenticates/authorizes users that exist in PostgreSQL database.

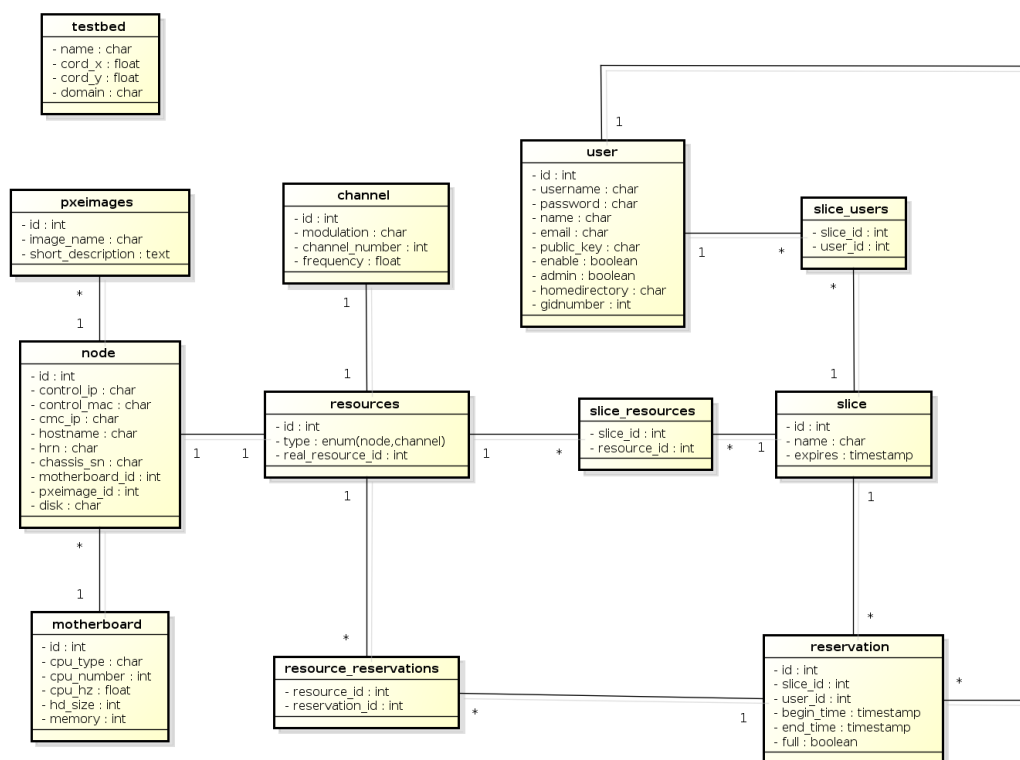



Figure 20 – Data organization in the LABORA Scheduler

Finally, LABORA Scheduler became full SFA compliant by the development of the LS-Driver and its integration into the SFA Wrap. By deploying the SFA Wrap [SFA WRAP] and configuration of the LS-Driver, an island can have its resource available in a federated manner. Thus, users from different islands can discovery, reserve and release resources, since these users have the credentials for proper authentication and authorization.


	D2.6 Report on the deployment of the second version of the control and monitoring framework for the FIBRE-BR facilities	Doc FIBRE-D2.6
		Date 31/03/2014

## 7 Islands Deployment Status


We have evolved substantially in terms of the FIBRE CMF infrastructure. The majority of the FIBRE-BR islands have all (or almost all) the components and services deployed. Table 2 describes the status of the deployment of services in FIBRE-BR islands and their private addresses according to the new standard for Brazilian island (see deliverable D2.7).

Table 2 - Current status of the services of FIBRE-BR islands

Island	Service	Status	Address
<b>CPqD</b>	Portal	Operational	portal.cpqd.fibre.org.br
	OCF	Operational (v0.7)	10.138.0.100
	OMF	Node control not working	10.138.11.200
	PerfSONAR	Operational	10.138.0.60 / 61
	LDAP	Operational	10.138.0.50
	DNS	Operational	10.138.0.80
	ZenOSS	Operational	10.138.0.80
	VPN	Operational	10.138.0.70
<b>UFRJ</b>	Portal	Operational	portal.ufrj.fibre.org.br
	OCF	Operational (v0.7)	10.129.0.100
	OMF	Operational	10.129.11.200
	PerfSONAR	Operational	10.129.0.60 / 61
	LDAP	Updating configuration	10.129.0.50
	DNS	Operational	10.129.0.80
	ZenOSS	Operational	10.129.0.80
	VPN	Operational	10.129.0.70
<b>RNP</b>	Portal	Operational	portal.rnp.fibre.org.br
	OCF	Operational (v0.7)	200.130.15.176
	OMF	Does not apply	-
	PerfSONAR	Operational	10.136.0.60 / 61
	LDAP	Operational	10.136.0.50
	DNS	Operational	10.136.0.80
	ZenOSS	Operational	10.136.0.80
	VPN	Does not apply	-

	D2.6 Report on the deployment of the second version of the control and monitoring framework for the FIBRE-BR facilities	Doc FIBRE-D2.6
		Date 31/03/2014

<b>UFPE</b>	Portal	Operational	portal.ufpe.fibre.org.br
	OCF	Operational	150.161.70.201 / 10.132.0.100
	OMF	Operational	150.161.70.203 / 10.132.11.200
	PerfSONAR	Operational	10.132.0.60 / 10.132.0.61
	LDAP	Updating configuration	10.132.0.50
	DNS	Operational	10.132.0.80
	ZenOSS	Operational	10.132.0.80
<b>UFPA</b>	Portal	Operational	portal.ufpa.fibre.org.br
	OCF	Operational (v0.7)	200.129.132.80 / 10.130.0.100
	OMF	Operational	200.129.132.77 / 10.130.11.200
	PerfSONAR	Operational	10.130.0.60 / 61
	LDAP	Operational	200.129.132.79 / 10.130.0.50
	DNS	Operational	10.130.0.80
	ZenOSS	Operational	200.129.132.81 / 10.130.0.80
	VPN	Operational	10.130.0.70
<b>USP</b>	Portal	Configuring	portal.usp.fibre.org.br
	OCF	Operational(v0.7)	10.133.0.100
	OMF	Configuring	10.133.11.200
	PerfSONAR	Partially Operational (hardware problem)	143.107.111.81
	LDAP	Operational	10.133.0.50
	DNS	Operational	10.133.0.80
	ZenOSS	Operational	10.133.0.80
	VPN	Does not apply	
<b>UFG</b>	Portal	Operational	portal.ufg.fibre.org.br
	OCF	Operational(v0.7)	200.137.197.232
	OMF	Operational	200.137.197.213
	PerfSONAR	Operational	10.137.0.60 / 10.137.0.61
	LDAP	Updating configuration	10.137.0.50


	D2.6 Report on the deployment of the second version of the control and monitoring framework for the FIBRE-BR facilities	Doc FIBRE-D2.6
		Date 31/03/2014

<b>UFSCAR</b>	DNS	Operational	10.137.0.80
	ZenOSS	Operational	10.137.0.80
	VPN	Does not apply	-
	Portal	Operational	portal.ufscar.fibre.org.br
	OCF	Operational (v0.7)	200.136.237.130
	OMF	Operational	10.135.0.90
	PerfSONAR	Operational	10.135.0.60 / 10.135.0.61
	LDAP	Operational	10.135.0.50
	DNS	Operational	10.135.0.110
<b>UNIFACS</b>	ZenOSS	Operational	10.135.0.80
	VPN	Operational	200.136.237.131/10.135.0.70
	Portal	Operational	portal.unifacs.fibre.org.br
	OCF	Hardware problem	10.131.0.100
	OMF	Hardware problem	200.128.79.104
	PerfSONAR	Hardware problem	10.131.0.60-64
	LDAP	Hardware problem	10.131.0.50
	DNS	Hardware problem	10.131.0.80
	ZenOSS	Hardware problem	10.131.0.80
<b>UFF</b>	VPN	Hardware problem	10.131.0.70
	Portal	Operational	portal.uff.fibre.org.br
	OCF	Operational (v0.7)	200.20.10.88
	OMF	Operational	200.20.10.188
	PerfSONAR	Operational	10.134.0.62
	LDAP	Operational	10.134.0.50
	DNS	Operational	10.134.0.80
	ZenOSS	Operational	10.134.0.80
	VPN	Operational	10.134.0.70

As we can verify in Table 2, all OCF instances have been updated to the most recent version (0.7). Few islands reported some hardware failures in its components (such as the UNIFACS virtualization server), and their replacements are already being provided.


Some islands also reported some difficulties while updating their LDAP schema to a new UID standard defined recently due to some incompatibilities with the definition for EU islands. We are currently migrating all LDAP user base to this new format to enable federated authentication between the testbeds.

Additionally, the so-called Phase 1 is being deployed, by using the FIBREnet Backbone and removing the VPN tunnels according to the network connectivity of the members to the

	D2.6 Report on the deployment of the second version of the control and monitoring framework for the FIBRE-BR facilities	Doc FIBRE-D2.6
		Date 31/03/2014


FIBREnet backbone. In cases where an island is ready for Phase 1, the VPN service status is listed as “does not apply”.



	D2.6 Report on the deployment of the second version of the control and monitoring framework for the FIBRE-BR facilities	Doc FIBRE-D2.6
		Date 31/03/2014

## 8 Conclusions and Future Work

In summary, this deliverable reviewed the effort in building the second version of the FIBRE-BR CMF. We overcome major issues like improving CMF and authentication to better support federation, creating an overlay backbone with full line rate performance on top of the RNP national backbone. As we continue to finish up these modifications and deploying the software homogeneously in all islands, we expect that the testbed will be ready to receive experimenters all over the world, and be open to the public in October 2014. Despite of that, we are already receiving some beta experimenters especially from Brazilian networking courses, like one currently being taught at UFSCar, and other testers are master and doctorate students doing research using the federated testbed for their work. These early experimenters will help improve further the software and guide us the organization of the helpdesk of open issues in the future.

	D2.6 Report on the deployment of the second version of the control and monitoring framework for the FIBRE-BR facilities	Doc FIBRE-D2.6
		Date 31/03/2014

"This work makes use of results produced by the FIBRE project, co-funded by the Brazilian Council for Scientific and Technological Development (CNPq) and by the European Commission within its Seventh Framework Programme."

END OF DOCUMENT