FUTURE INTERNET TESTBEDS
EXPERIMENTATION BETWEEN
BRAZIL AND EUROPE

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***

D2.5. Report on the deployment of the first 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 16
Actual submission date: 28/05/2013
Start date of project: June 1st 2011 Duration: 34 months
Version: v.1

Project co-funded by the European Commission in the 7th 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)	✓

FP7 Grant Agreement No.	288356
CNPq Grant Agreement No.:	590022/2011-3
Project Name	Future Internet testbeds/experimentation between BRazil and Europe – EU
Document Name	FIBRE-D2.5
Document Title	D2.5: Report on the deployment of the first version of the control and monitoring framework for the FIBRE-BR facilities
Workpackage	WP2
Authors	Cesar Marcondes (UFSCar) Ettore Zugliani (UFSCar) Natalia Castro Fernandes (UFF) Kleber Vieira Cardoso (UFG) Raphael Augusto Dourado (UFPE) José Augusto Suruagy Monteiro (UFPE) Adriano Spínola (UNIFACS) Joberto Martins (UNIFACS) Ricardo de Freitas Gesuatto (UFSCar)
Editor	Cesar Marcondes (UFSCar)
Reviewers	
Delivery Date	
Version	V1.0

Abstract

The first version of the FIBRE CMF was reached at M16 (month) of the project as it described here. The main result of this first version FIBRE-BR CMF was the adaptation and combination of 3 CMFs (OCF, OMF and ProtoGeni), instead of a single one, in order to match the variety of resources (netFPGAs, ORBIT nodes, OF switches) and all the requirements described in D2.1, and previously discussed in deliverables D2.2 and D2.3.

This document starts by describing a holistic view of the FIBRE CMF combination, the addressing scheme, the identification through LDAP, and we further exemplify by giving a particular configuration of an island - the UFSCar island - to show how hardware and software (available at the moment of this report) were organized and how it is being replicated in all other islands. We describe some preliminary tests and we also show the status of deployment of this new CMF software in all islands.

In following sections, we discuss the features added to the CMFs, in order to cope with, the differences of the Brazilian testbed environment compared to other environments where the original CMFs (in isolation) were deployed in Europe and USA. In particular, we describe the modifications developed on the OFELIA control framework for the FIBRE-BR scenario. For example, one of the developments accomplished, was to make OCF installation and library dependencies more “independent” of the underlying Linux distribution. For example, while the latest OCF framework makes substantial use of a specific *Debian* distro, on the other hand, netFPGA bitstreams have been using a more stable environment based on CentOS. Thus, it justifies the need to have OCF ported and compatible with a CentOs environment.

Furthermore, due to the nature of the Brazilian FIBRE-BR islands – have only 1 big virtualization server and 3 rackmount PCs based on netFPGA – the FIBRE-BR environment need to make use, at the same time, of the netFPGAs as OpenFlow switches and as “small” virtualized servers based on Xen (OFELIA Xen Agent – OXA – alongside with netFPGA) to perform OpenFlow experiments correctly. Therefore, we had to modify the environment to use, simultaneously, Xen linux bridges and an openvswitch running side by side in a CentOS environment.

Other developments of the FIBRE-BR CMF also include modifications performed on the OMF control framework to cope with a centralized scheduling mechanism, such that it is possible to federate OMF resources from different islands. The idea was to use the early implementation of the NITOS Scheduler [Anadiotis 2010] and encapsulate it in a virtual machine to provide this scheduler that best suits the Brazilian testbed needs, especially in terms of wireless resources.

Another important result, in terms of integration of the federated resources, was that the WP2 team decided to create a phase 0 integration plan with all islands based on VPN tunnels interconnecting openVPN servers running inside VMs, and thus providing interconnection of the control and experiment Xen bridges per island, in a national overlay network.

Finally, following the discussion of these main changes in isolated CMFs to create the combined FIBRE-BR CMF, we also describe the first version of the monitoring solution. We

start by discussing the use of ZenOSS [ZenOSS 2013] as a monitoring tool within the islands, for example, by reading information from the underlying XEN environment. Afterwards, we discuss the use of 2 “spare” servers that would work as PerfSonar bandwidth test and latency test managers. These servers were added to the scenario, as a bonus, for the testbed, it was not in the original proposal. By using these servers, the Brazilian islands can verify the amount of available bandwidth and latency on every experiment on-demand. In addition, the latency server is connected to a GPS-device to make precise delay measurements.

We finish the document concluding the version 1 of the FIBRE-BR software by showing how the deployment status currently holds.

TABLE OF CONTENTS

List of Figures

No table of figures entries found.

List of Tables

1 Acronyms

AM	Aggregate Manager
API	Application Programming Interface
CP	Control Plane
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
p.	Page
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

2 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.

[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. Tassiulas, 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. 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.

3 Holistic view of the FIBRE-BR Island

The purpose of this section is to give some overview idea of how the individual islands are organized in terms of hardware and software, and the software part is the focus on this document. We describe how the islands are built up from two logical networks (control and experiment). The management network, the one used for the test bed manager to setup the resources and configurations also re-uses the control network based on the premises of the secure communication of the web services through self-signed certificates.

Afterwards, we describe the interconnection and the addressing scheme from the island point-of-view to be used internally for services, for VMs, and OF switches. We describe the identification mechanism of the FIBRE-BR CMF. Finally, in the last section we give an example of a running island - UFSCar testbed - and tests performed to show the running elements.

3.1 Recap of a typical Brazilian architecture per island

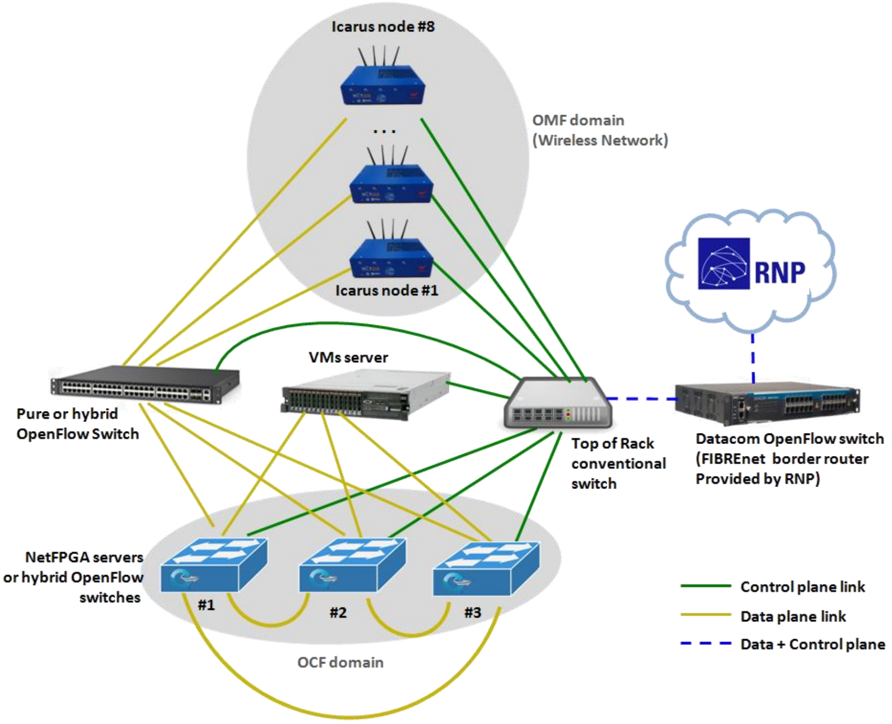

Figure 1 - FIBRE-BR typical island

A typical Brazilian island has the following hardware, interconnection and planes, according to Figure 1 – FIBRE-BR typical island. As the rest of the equipments are still not available, for example the pure OpenFlow switch and the ToR switch, and as they are added, we will perform additions to the software, for the missing specific components in version 2 of the FIBRE CMF. In this figure 1, we can verify only two logical networks:

Control network (green lines): The control network is composed by a “set” of separated linux bridges at the servers (VMs Server and NetFPGA Servers) called “control bridges”. And these bridges are connected by the control switch that allows the communication of the OCF (expedient, opt-in, vt-manager) and OMF (am, rc, ec) specific services, performed securely by cryptographic certificates. It also enables the access to the physical machines of the testbed, by allowing ssh to be performed on the control network (secured by passwords). Finally, and most importantly, the control network allows users to have access to their experiments (their slices). The control network has a NAT access from the main server in order to allow, experimenters (and their VMs) and the testbed itself to use the internet to update packages easily.

Experiment or Data Plane network: (yellow lines): this is the network that interconnects the OpenFlow switches, Virtual Machines (VMs) and inter-islands interconnections, available for the experimenter. It is important to remind that the resources (flowspaces and VMs) are contained within a slice controlled by the FIBRE-BR control framework.

At the time of this writing, the hardware available to setup Figure 1, was the VM Server, and 3 netFPGAs servers, where 1 of the netFPGAs is being used as a second VM server and 2 netFPGA are being used as OpenFlow switches (Figure 2). Later on, we will describe the software developed and currently in tests that will allow the netFPGAs to perform both as servers and OpenFlow switches at the same time. In terms of the OMF infrastructure, it uses only the server to instantiate VMs and manages the OMF resources.

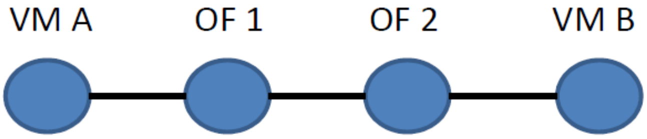

Figure 2 - Linear Topology using Available FIBRE-BR Resources

3.2 Interconnection Details (both OCF and OMF)

As decided by the WP2 team, the islands are being interconnected, in a phase called PHASE 0, by the use of VPN tunnels (running openVPN VMs). Thus, either the control network, or the experiment network, for both OCF and OMF software, would ensemble a single bus, from the experimenter point of view. This way, all resources will be available in all islands, requiring coordination in order to avoid conflicts, especially in terms of addressing. Thus, following, we describe the proposed address scheme being used.

3.3 FIBRE-BR addressing scheme

The WP2 team agreed, temporarily, on an addressing scheme proposal that suits both naming and low-level addressing for the island resources. It could be redefined in the future based on experiences with the European and American Future Internet testbeds.

3.3.1 FIBRE-BR Naming (DNS)

We start by describing the naming scheme. Since it could be quite demanding to remember all the services running in a particular island, the idea of the naming scheme is to store in a DNS database (from *.fibre.org.br domain) every service that needs to be accessed easily by a human.

For example, let's suppose an administrator wants to check the monitoring service of the UNIFACS Island. The idea is to convey the service name along with the institution name to reach the specific service.

service.institution.fibre.org.br

i.e. monitoring.unifacs.fibre.org.br

The listing of all services is currently being created by the WP2 team and it will be available in wiki of the project.

3.3.2 Holistic Private IP Addresses Allocation

Based on the fact of the PHASE 0 of the project being based on VPN tunnels and a common shared bus for the control and experiment network, it is mandatory for the project to have an addressing scheme that can distinguish OF switches in the FIBREnet, OF switches in the islands, linux-bridges, VMs, server IP addresses and so on. It is important to remind that this is a provisory address scheme, since it could have compatibility issue with other address schemes in the future.

We start by describing that all components OCF and OMF, in all islands, will use a private IP address space 10.x.x.x/8. There is a substantial amount of addresses in using this subnet for any practical purposes.

According to Figure 1, the Datacom OpenFlow switch (FIBREnet border router Provided by RNP) will have the following address scheme, based on the PoP phone area code, easy to remember.

FIBREnet Address Scheme

10.1(Phone Area Code).0.1

Datacom Switches WAN address scheme

PoP	Endereço
PoP-RJ *	10.121.0.1/32
PoP-SP *	10.111.0.1/32
PoP-ES	10.127.0.1/32
PoP-BA *	10.171.0.1/32
PoP-PE *	10.181.0.1/32
PoP-PA *	10.191.0.1/32
PoP-GO *	10.162.0.1/32
PoP-DF *	10.161.0.1/32
PoP-RS	10.151.0.1/32
PoP-MG	10.131.0.1/32

* - initial RNP Point of Presence Datacom Switches, others are for future use.

Following, the islands interconnection plan needs to differentiate the servers, the VMs, the bridges, of every single component within an island. In order to facilitate this differentiation, since VPN requires some specific code for every member participating, we leverage the VPN code in the internal IP addressing scheme of the islands.

Islands Address Scheme

10. VPN Code . Service/Equipment/Control Framework . Host

VPN Code per institution

Institution	✓ Code
✓ UFRJ	✓ 1
✓ UFPA	✓ 3
✓ UNIFACS	✓ 4
✓ UFPE	✓ 5
✓ USP	✓ 6
✓ UFF	✓ 7
✓ UFSCar	✓ 8
✓ RNP	✓ 9
✓ UFG	✓ 10

✓ CPqD	✓ 12
✓ New Institutions	✓ N + 1

The third octet of the island private address scheme can be used to define the framework being used, for example OCF, OMF or ProtoGENI can run concurrently using different address spaces.

Control Frameworks Codes (for specific components)

OCF Code	10
Generic Format	10.X.**10**.0-255/24
Example (UFRJ Scenario):	10.**1**.**10**.0/24

OMF Code	11
Generic Format	10.X.**11**.0-255/24
Example (CPqD Scenario):	10.**12**.**11**.0/24

ProtoGENI Code	12
Generic Format	10.X.**12**.0-255/24
Example (UFPA Scenario):	10.**3**.**12**.0/24

By setting up the third octet as ZERO, the final octet can define a series of equipments and services running within an island.

Equipment and Services Address Scheme

10**.** VPNCode **.** 0**.** Service/Equipment

Examples of Equipment and Services Address Scheme

Top of Rack Switch (ToR)	**1**
Generic Format	10.X.0.**1**/32
Example (UFPA ToR)	10.**3**.0.**1**/32
Example (CPqD ToR)	10.**12**.0.**1**/32
Notes	All other physical equipment would be allocated in the range: 3~49.

VM Server Address	**2**

Generic Format	10.X.0.2/32
Example (UFPA VM Server)	10.3.0.2/32

LDAP Server Address	50
Generic Format	10.X.0.50/32
Example (UFRJ LDAP Server)	10.1.0.50/32
Example (CPqD LDAP Server)	10.12.0.50/32

ZenOSS Server Address	60
Generic Format	10.X.0.60/32
Example (UFRJ ZenOSS Server)	10.1.0.60/32
Example (UFPA ZenOSS Server)	10.3.0.60/32

Finally, the VMs MAC address range provided by the every island and allocated by the OCF or OMF control frameworks is performed by using a generic format of 2 fixed octets based on VPN code + 4 first octets directly from the upper IP address scheme.

MAC Layer Address Scheme for the VMs

2 fixed bytes + 4 first bytes of the IP private address

In summary, we have proposed an address scheme to work with the needs of all the CMFs in the same address space. This private address scheme network will be reached by a VPN single-on point of entrance, and therefore the FIBRE-BR testbed will need few valid IP addresses (from each institution) to reach the following important services (LDAP, DNS, Portal / VPN) and monitoring services (PerfSonar).

3.4 Identity management in FIBRE

The first version of the control and management framework of FIBRE provides federated authentication through LDAP (Lightweight Directory Access Protocol). Hence, both OMF (ORBIT Management Framework) and OCF (OFELIA Control Framework) are configured to authenticate users using LDAP.

We use the same LDAP schema in all entities, to simplify the interoperability. The schema attributes were selected based on the schema of Eduroam, a project to federate identities through different universities and research centres spread over the world. The idea was to use a complete description of the researchers' identities and also to allow, in the future, using Eduroam databases to authenticate researchers in FIBRE-BR.

We added a new attribute to this schema in order to perform access control in the federated testbed. Current versions of OCF and OMF allow researchers to create an account on the testbed portal. The access to the resources, however, depends on the analysis of a research project sent by e-mail. Hence, we added a Boolean attribute userEnable to manually control the access of new FIBRE-BR users.

The FIBRE-BR LDAP tree is rooted in the NOC. All the islands have a local database and a reference to NOC to perform the authentication of users of other islands. The NOC also maintains copies of the databases of all FIBRE-BR islands to reduce authentication times and to increase reliability, in case of temporary failure in an island.

3.5 UFSCar Testbed Example

In order to provide a running example of the first version of the FIBRE-BR CMF in use, we briefly describe the UFSCar running island and tests performed to show the running elements. We focus on the OCF elements but OMF elements are also running concurrently providing experimental setup from the VM server (IBM Server).

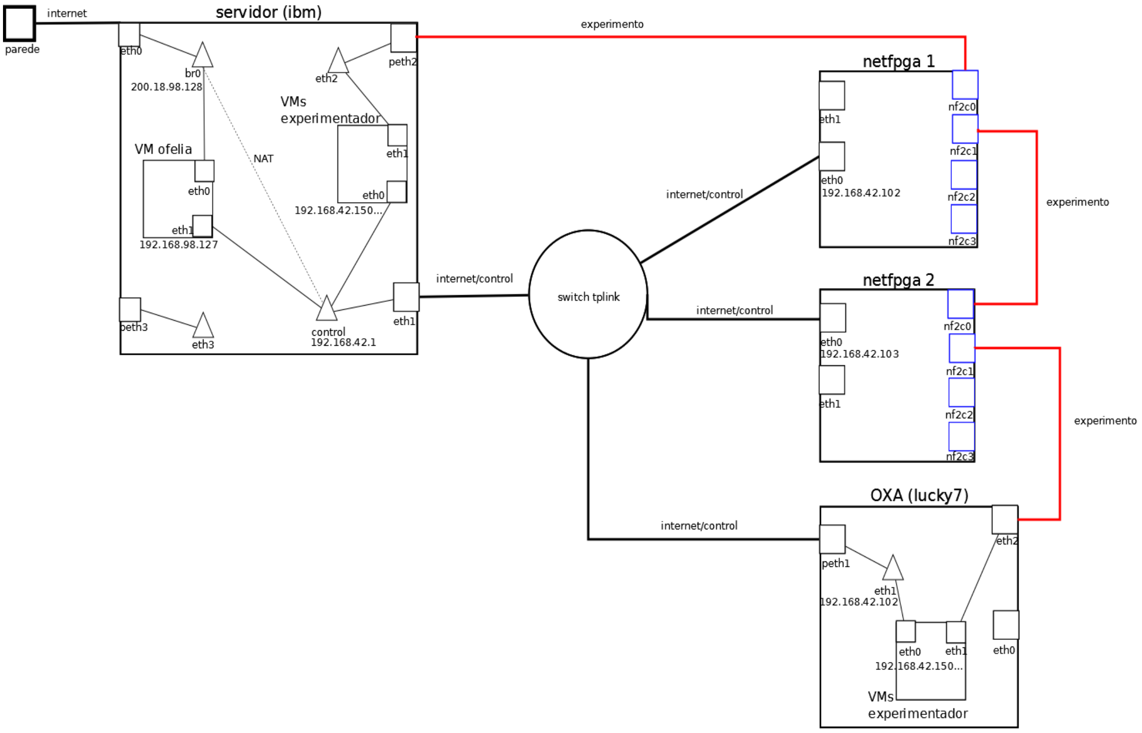

Figure 3 - UFSCar OCF-related Running Testbed

In Figure 3, we can identify 4 machines of the UFSCar testbed (servidor (ibm) + netfpga1 + netfpga 2 + lucky7) representing the linear topology of Figure 2. The OCF elements are interconnected by bridges represented by small triangles. The control network is being plugged on a "spare" TPLink switch and the experimental network is the linear interconnection between netFPGAs and the servers. The netFPGAs of the center are acting only as OpenFlow switches using the netFPGA bitstream for hardware-based flow tables and OVS software. The "experimentador VM" is the one used by the experimenter and thus is it on-demand allocated. In this setup, we didn't allocate the proposed address scheme (previous section) yet, this is the next step.

We point out that, there is an "OCF portal VM" inside the IBM server. This is very peculiar of the FIBRE-BR setup. This OCF portal contains all the expedient software, the opt-in, the vt_manager. In DOM0 of the IBM server, we've installed the Flowvisor. Furthermore, we provide internet connectivity for the VMs and for the servers using a NAT service in the IBM server (br0 bridge with NAT).

3.5.1 Simple testing procedures

Following are some testing procedures that can be used to verify all the configurations are in conformance with this version 1 of the FIBRE-BR CMF.

1) Login at the VM Server (IBM server) as root (or sudoer) using ssh -X
2) Verify xen daemon is up and running
    a. ps –aux | grep xen
    b. xm list
3) Verify the bridges configuration (control bridge, experiment bridge)
    a. vim /etc/xen/ofelia.cfg
    b. vif = [MAC1 – br0, MAC2 – control, MAC3 – eth2]
4) Verify all bridges are up and running, NAT is enabled, DNS is enables, network interfaces
    a. vim /etc/rc.local (force bridges to stay up and configure NAT)
    b. vim /etc/resolv.conf
    c. vim /etc/network/interfaces
5) Connect to the console of OCF portal VM to verify configurations
    a. xm console ocf-server
    b. vim /etc/network/interfaces
    c. if you have created a VM throught the OCF system before, try the same

6) Verify the OFELIA XEN AGENT (OXA) running

a. ps -ef | grep oxa (it uses libvirt)

b. Check the configurations

i. vim /etc/init.d/OfeliaLauncher.sh (remember to point to python2.6)

7) Verify disk space and files

a. Check the /data directory (where OXA is using for the VMs)

b. The files are supposed to be located at /data/oxa/cache/vms/UUID#1/UUDI#2/

i. Example. UUID#1 = c2577cb7-5da0-4ed7-b108-9a449003d11f

c. Check the offline configuration of a specific VM

i. vim ibm4-backup2.conf

8) Verify if Flowvisor is configured and running

a. vim /etc/flowvisor (UFSCar FlowVisor is installed in DOM0)

b. ps –ef | grep flowvisor (check if it is running)

9) Verify the rest of the equipment (verify if netFPGAs are up and running)

a. ssh root@192.168.42.103

b. ps -ef | grep of (in particular, pay attention to–d parameter of ofdatapath (it must be coherent with the OFELIA DPID terminology, and also check if it is pointing to the correct FlowVisor IP address and port through the control network IP range - 192.168.42.1:6633

c. Remember that nf2c0 = port 1 in FlowVisor

10) Access the OCF Portal VM

a. ssh root@200.18.98.127

b. Check if all services are up and running

i. service apache2 status

ii. service mysql status

c. Access the OCF interface as administrator (root) – don't forget https

11) Verify that root can see all the project, pending requests, manage all the users

12) Verify that the aggregates are up and running (opt-in and vt_manager)

13) Access both the specific web portals of the aggregate managers

14) Create an account for a user

15) Verify that the user can ask permission to create and project and access the aggregate managers

16) Create a slice for that user

17) Create virtual machines in both servers

18) Create an OpenFlow slice based on VLAN ID

19) Start the VMs of the user based on a pre-configured UFSCar VM that comes with POX installed

20) Access the VMs through the OCF Portal (in the future, the VPN portal)

21) Run the following commands on the VMs

    a. ssh root@your_VM_IP (password: ofelia)

    b. configure VLAN interfaces on both VMs

        i. modprobe 8021q

        ii. ifconfig eth1 up

        iii. vconfig add eth1 222 (222 is the VLAN ID of this user)

        iv. ifconfig eth1.222 18.18.18.18/24 up

        v. ping 18.18.18.19

    c. This will not work, unless, the other VM have similar configuration

    d. The FlowVisor is pointing to this VM and then by running

        i. “./pox.py log.pretty_log forwarding.learnig_switch”

22) Ping Works !!! DONE !!!

4 OCF Modifications tailored to FIBRE-BR

4.1 Introduction / Motivation

The FIBRE Brazilian islands have some particularities; one of them is the fact that they are mostly heterogeneous. This is a problem not only for users, since they need to understand how every island work, but it is also difficult to maintain them functional as a whole. One of our project decisions was to use NetFPGAs on the islands; thus requiring the use of the Linux distribution CentOS.

To fully use the available resources on the islands and have a more interesting topology, our idea was to also use the NetFPGA servers as OXA (Ofelia Xen Agent) servers. The problem is that OXA also needs one specific distribution, in this case Debian based distributions. As a solution, we provided a working infrastructure for OXA to run in different environments, in a way that the islands can progressively become more homogenous in terms of software. Using the NetFPGA and other servers as OXA servers, we encourage the research and development of different tools for the Ofelia Control Framework.

4.2 Porting OXA to CentOS

The Ofelia Control Framework software has been developed with a Debian Linux environment in mind. It is known to work on Debian Squeeze (the officially supported platform) and Wheezy; and according to comments on GitHub it should also work on Ubuntu with no modification.

However, the supported environment for the NetFPGA is a slightly modified CentOS 5.x installation, a distribution based on Red Hat Linux sources. Although these distributions still provide a GNU/Linux compatible userland, it's common to identify diverting aspects. Different development views and implementations of a feature often result in incompatible package managers, conflicting naming conventions, and disparate configuration files and initialization hierarchy.

In order to install and run Ofelia XEN Agent (**OXA**) successfully, a number modifications are needed in both the destination environment and in the software. The first step was to analyze the components related to visualization in Ofelia. By correlating each packaged dependency to equivalent packages in CentOS it's possible to identify an additional problem: the concept of “stability” may vary somewhat according to the distribution's policy. Key packages, such as Python, Libvirt and the XEN hypervisor, are outdated in CentOS 5.x when compared to Debian Squeeze. Most notably, Python 2.4 lacks many syntactic features widely used in OCF; an unofficial port of Python 2.6 can be found on trusted repositories, but other packages that depend on Python are not available.

As a solution for dependencies, the following trusted repositories were added to YUM, the package manager available on CentOS:

- Extra Packages for Enterprise Linux (EPEL)

```
wget http://dl.fedoraproject.org/pub/epel/5/i386/epel-release-5-4.noarch.rpm
rpm -Uvh epel-release-5*.rpm
```

- IUS Community Project (IUS)

```
wget dl.iuscommunity.org/pub/ius/stable/Redhat/5/i386/ius-release-1.0-10.ius.el5.noarch.rpm
rpm -Uvh ius-release-*.rpm
```

These repositories only provides the basic Python 2.6 packages. Additional dependencies were manually compiled and packaged, resulting in the deployment of a custom repository for FIBRE related packages. Temporarily available at http://www2.comp.ufscar.br/~ricardofg/centos5/ , the repository contains additional Python packages needed for OCF, as well a newer Libvirt version and an updated XEN package.

After all dependencies were packaged and installed successfully, the installation scripts in OXA were modified to provide detection of distributions, selection of the correct package manager and configuration tools, and installation of the adequate dependencies.

The first prototype version implemented these features in a high level manner. This version modifies the installation scripts (under `vt_manager/src/python/agent/tools/versions/default/install`) and adds an "**osdetect**" script based on Linux Standard Base (LSB) helpers to detect the environment and adjust the installation accordingly. A LSB-compatible installation script was also proposed; however this solution was found to be infeasible, as not all distributions strictly adhere to the standards proposed on LSB.

Further discussion with the I2Cat team on email and GitHub suggested re-factoring **ofver** itself in order to better accommodate multiple distributions in a modular way, without affecting code complexity as support grows over time. This is currently a work in progress.

Instantiating VMs with OCF still proved unsuccessful, with failed attempts to create a new VM and VMs not starting. A set of shell commands hardcoded directly were found in OXA's XEN provisioning infrastructure (for example **FileHdManager.py**). Creating a symbolic link for **nice** and **ionice** at **/usr/bin** solved the first problem in a simplistic manner.

As noted above, the next version of the software should use the modified **ofver** infrastructure as a clean way to provide any platform-specific versions of files.

5 OMF and NITOS Scheduler modifications tailored to FIBRE-BR

In [Report D2.2], we enumerated the main issues related to OMF: 1) resource allocation, and 2) authentication and authorization. In order to deal with these issues, we have integrated OMF with NITOS Scheduler [Anadiotis 2010]. NITOS Scheduler depends on Joomla CMS to provide the authentication and authorization capabilities. The integration is not trivial and it has demanded intensive interaction with NITOS Scheduler's developers from NitLab. Besides, the FIBRE project has as one of the main goals the federation. OMF 6 is being developed with federation-ready, but the Brazilian deployment was made based on the stable OMF 5.4. The OMF 5.4 allows a basic federation approach based on peering of XMPP servers using Server 2 Server protocol [S2S XMPP]. However, there is no isolation between experimenters under this federation approach.

In the first version of the OMF and the NITOS Scheduler that are being deployed in Brazilian islands, we have addressed the issues described. In the following, we present the details about the NS_OMF-BR, i.e. the Brazilian version of the CMF composed of NITOS Scheduler and OMF.

Figure 4 presents the general overview of the NS_OMF-BR. Beside NITOS Scheduler and OMF, it is important to highlight the LDAP component that is a key step towards the federation.

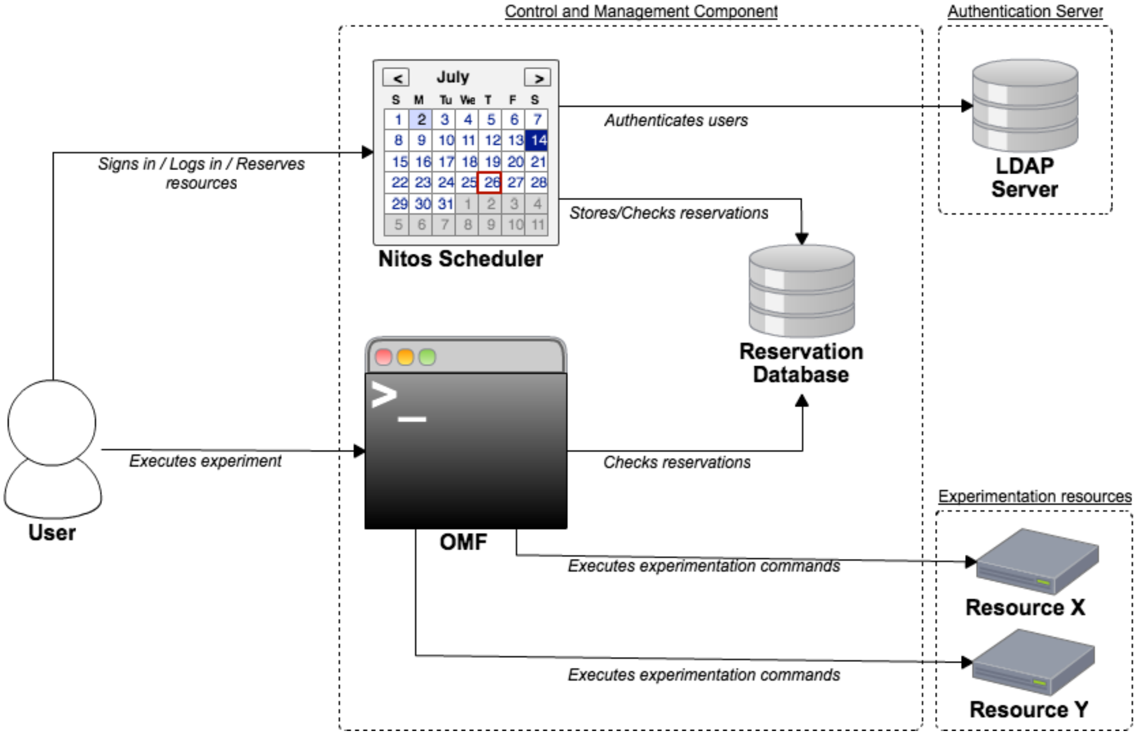
Control and Management Component
Authentication Server
July
Authenticates users
LDAP Server
Signs in / Logs in / Reserves resources
Stores/Checks reservations
Nitos Scheduler
Reservation Database
Executes experiment
Checks reservations
User
OMF
Experimentation resources
Executes experimentation commands
Resource X
Executes experimentation commands
Resource Y

Figure 4: NS_OMF-BR overview.

In the Brazilian version of NITOS Scheduler, we used the core PHP code but we do not employed Joomla CMS as the front-end and user module. We developed a customized Web front-end and user module that is directly integrated with a LDAP database. Since FIBRE-BR intends to use an integrated LDAP directory service to the whole federation, the Brazilian NITOS Scheduler is ready to be used in a distributed manner.

The NS_OMF-BR employs a unified inventory database for NITOS Scheduler and OMF. Besides, the database of every island can be accessed in a centralized manner in order to have a complete view of the resources in the federation and NS_OMF-BR also allows allocating the resources in this centralized manner. Figure 2 illustrates this integration and also the shared authentication and authorization viewed that is offered by the integrated LDAP directory service.

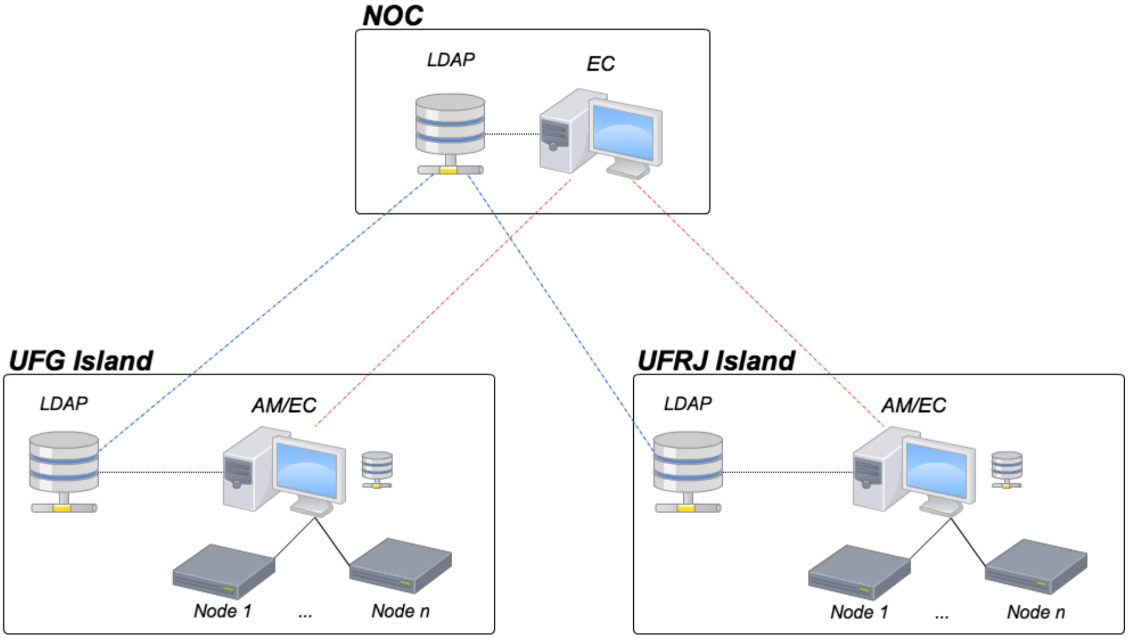

Figure 5: Integration between NS_OMF-BR islands.

The process of user account creation in the EC [Report D2.2] is now simplified by employing the PAM module called `mkhomedir`. This module creates automatically the user directory in the operating system at first access if the user exists in the LDAP database.

The vanilla EC does not control which nodes are available for an experimenter according to his/her allocation in the NITOS Scheduler. We modified the EC code to get the allocation information filled by NITOS Scheduler. Nowadays, the NITOS Scheduler original developers, from UTH, also offer a modified EC, but that was not the case when we forked the code. So we followed a slight different approach that includes also the verification of the "`omf exec`" command. Actually, the experimentation script is parsed in order to allow running the experiment only in the previous allocated nodes. The original branch of the EC verifies only "`omf load`" and "`omf tell`" commands.

The federation based on peering of XMPP servers using S2S protocol demands for a slice in the AM [Report D2.2] containing all federated resources. As a consequence, the experiments' IDs of all federated experiments are all related to the federated slice. Our modified EC employs the user ID, and not the slice, as the experiment ID. This makes easier to associate users and their experiments.

6 Network Infrastructure Monitoring

The FIBRE-BR proposed 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.

6.1 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 have decided to deploy 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 6 shows the current deployment status of these servers among the FIBRE-BR islands. The function of the additional server named “ZenOSS” depicted in Figure 6 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 are in the process of deploying the perfSONAR Monitoring and Debugging Dashboard (MaDDash) on the NOC facility. This tool provides a global view of the network performance and will serve as the entry point for accessing the measurements collected by the pSPT servers. Figures 7 and 8 show a series of screenshots of our prototype installation, collecting measurements from FIBRE islands.

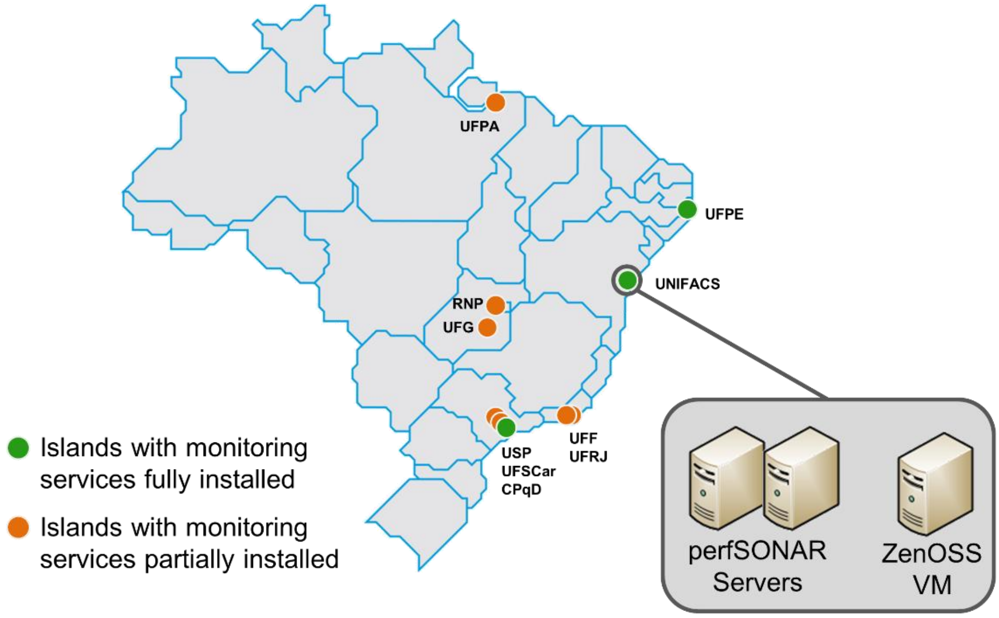
UFPA
UFPE
UNIFACS
RNP
UFG
USP
UFSCar
CPqD
UFF
UFRJ
Islands with monitoring services fully installed
Islands with monitoring services partially installed
perfSONAR Servers
ZenOSS VM

Figure 6. Monitoring infrastructure deployment

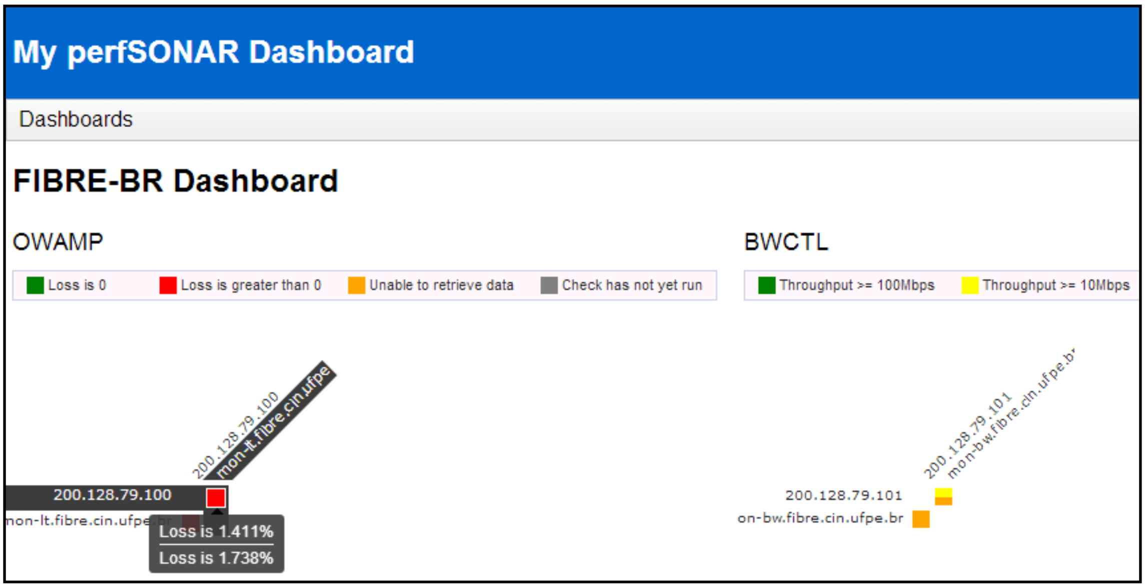

My perfSONAR Dashboard

Dashboards

FIBRE-BR Dashboard

OWAMP

🟩 Loss is 0 🟥 Loss is greater than 0 🟧 Unable to retrieve data ▨ Check has not yet run

BWCTL

🟩 Throughput >= 100Mbps 🟨 Throughput >= 10Mbps

200.128.79.100 mon-lt.fibre.cin.ufpe

200.128.79.100

non-lt.fibre.cin.ufpe

Loss is 1.411%

Loss is 1.738%

200.128.79.101 mon-bw.fibre.cin.ufpe.br

200.128.79.101

on-bw.fibre.cin.ufpe.br

Figure 7. MaDDash homepage: Dashboard with measurements for each node pair

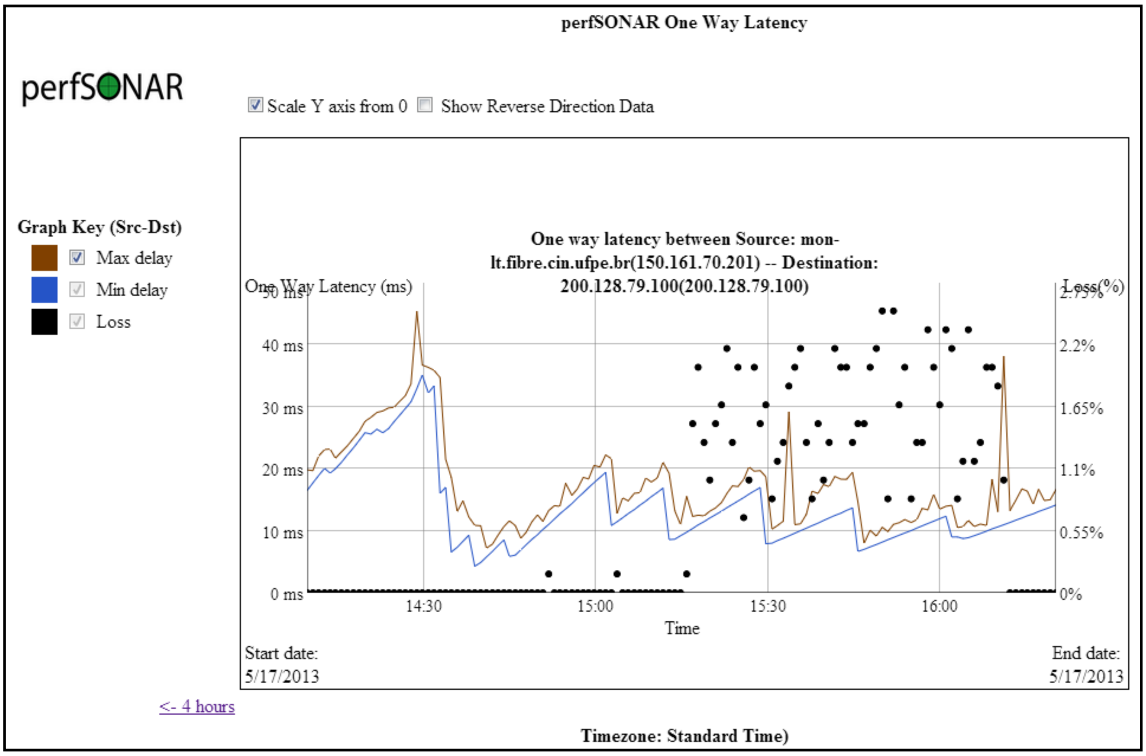

perfSONAR One Way Latency

☑ Scale Y axis from 0 □ Show Reverse Direction Data

Graph Key (Src-Dst)

- ☑ Max delay
- ☑ Min delay
- ☑ Loss

Figure 8. One-Way Latency Graphic Results

Since latency measurements require a highly synchronized clock, each island received a GPS kit (model ublox LEA6T) to be coupled to the node performing latency measurements. The installation and configuration of these kits are in progress and in the meantime, the NTP servers available at each PoP are being used for synchronization.

6.2 ZenOSS Deployment

To fulfill the infrastructure monitoring requirements of each island we are using Zenoss Core v4.2 [Zenoss 2013]. This open source management tool provides a wide collection of advanced resources to collect, display and analyze measurement data collected via SNMP, SSH and WMI (Windows Management Instrumentation). Zenoss is being deployed in each island alongside the pSPT servers. 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 9.

To better cope with the characteristics of the hardware available at each island, 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 10).

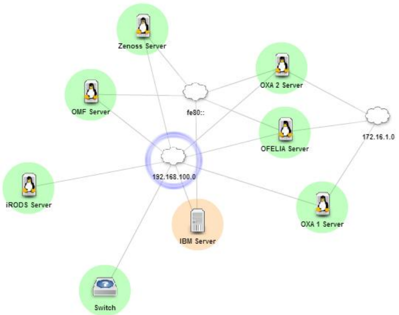

Figure 9: ZENOSS view of an infrastructure's topology.

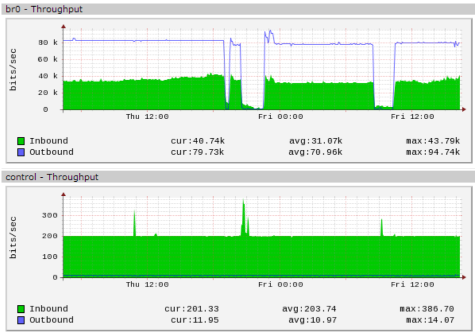

Figure 10: Detailed graphic of NIC's throughput

In this first phase, the NOC will be able to access the data of each island through a webpage coupled to the Zenoss Dashboard (shown in Figure 11), which contains links to each island's Zenoss server. With this solution, the NOC will be capable to monitor all islands without jeopardizing the network with SNMP messages. While this is a temporary setup, we are in the process of developing a more refined solution to deal with the centralized access to the distributed Zenoss servers.

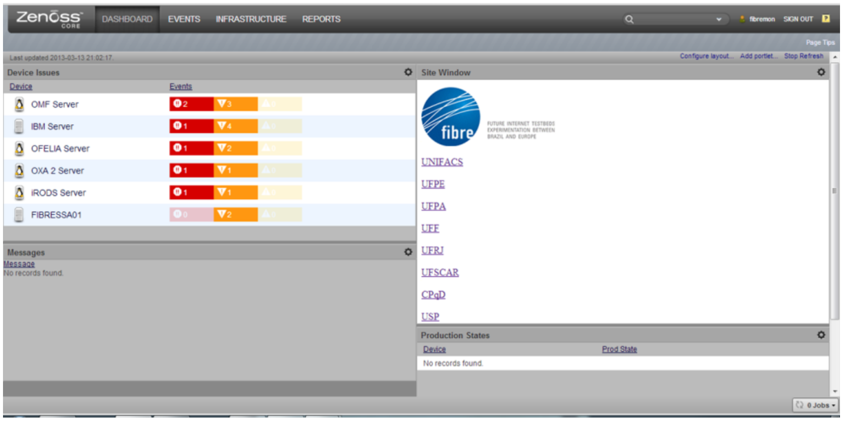

Last updated 2013-03-13 21:02:17.

Device Issues

Device	Events		
OMF Server	2	3	0
IBM Server	1	4	0
OFELIA Server	1	2	0
OXA 2 Server	1	1	0
iRODS Server	1	1	0
FIBRESSA01	0	2	0

Messages

Message

No records found.

Site Window

UNIFACS

UFPE

UFPA

UFF

UFRJ

UFSCAR

CPqD

USP

Production States

Device	Prod State
No records found.	

Figure 11: Main Dashboard of Zenoss with a customized FIBRE's WebSite.

7 FIBRE-BR CMF Deployment Status

The FIBRE-BR CMF is being deployed in all Brazilian sites, with all the components. Some of the components are still missing the hardware, for example the ORBIT nodes, other are in the testing phase, for example the LDAP and VPN services. Following is the complete status table.

	UFPA	USP	UFF	UFSCar	UFRJ	UFG	UNIFACS	UFPE	RNP	CPqD
OCF	**Yes**	**Yes**	**Yes**	**Yes**	**Yes**	**Yes**	**Yes**	**Yes**	**Yes**	**Yes**
OMF	**Yes**	**Yes**	**Yes**	**Yes**	**Yes**	**Yes**	**Yes**	**Yes**	**No**	**No**
netFPGA	**Yes**	**Yes**	**Yes**	**Yes**	**Yes**	**Yes**	**Yes**	**Yes**	**No**	**Yes**
ZenOSS	**Yes**	**Yes**	**Yes**	**Yes**	**Yes**	**Yes**	**Yes**	**Yes**	**Yes**	**Yes**
PerfSonar	**Yes**	**Yes**	**Yes**	**Yes**	**Yes**	**Yes**	**Yes**	**Yes**	**Yes**	**Yes**
ORBIT nodes	**No**	**No**	**No**	**No**	**No**	**Yes**	**No**	**No**	**No**	**No**
VPN	**No**	**No**	**No**	**No**	**Yes**	**Yes**	**No**	**No**	**No**	**No**
LDAP	**No**	**No**	**Yes**	**No**	**No**	**No**	**No**	**No**	**No**	**No**

UFF was performing LDAP experiments, and UFRJ and UFG have already federated their OMF resources by the use of VPN. Finally, the RNP site is the NOC and thus it is not receiving ORBIT nodes, neither installing OMF nor netFPGAs.

8 Conclusions and Future Work

In summary, this deliverable reviewed the effort in building the first version of the FIBRE-BR CMF. This CMF is a composition and adaptation of 3 other CMFs: OMF, OCF and ProtoGeni. Moreover, the Brazilian side resources are atypical and need special attention, for example, turn the netFPGA rack mount PCs into capable small VM servers is important to give the project more resources and flexibility for the users.

We are going to continue debugging and improving this software base. In fact, some of the adaptations being developed in this WP2, and described here In OCF and OMF contexts, have their particular FORKs in git projects. This git process allows the developers of the main OCF tools, to reuse some of these adaptation schemes, for other purposes.

As a future plan, we intend to integrate the OCF and OMF under a common LDAP hierarchy infrastructure. Change the addressing scheme of all islands, and provide the VPN connectivity between islands and explore QinQ functionalities in the FIBREnet switches to expose the SDN core to the experimenter.

"*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