FUTURE INTERNET TESTBEDS
EXPERIMENTATION BETWEEN
BRAZIL AND EUROPE

Grant Agreement No.: 288356

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***

D2.2 Report on the design of the control monitoring framework

Author: **Cesar Augusto C. Marcondes**
Revised by: Leonardo Bergesio, Dimitris Giatsios, Sebastià Sallent
Due date of the Deliverable: Month 10
Actual submission date: Jun 20 2012
Start date of project: June 1st 2011 Duration: 34 months
version: v.1.0

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
Project Name	Future Internet testbeds/experimentation between BRazil and Europe – EU
Document Name	Report on the design of the control monitoring framework
Document Title	Report on the design of the control monitoring framework
Workpackage	WP2
Authors	Name (Organization)
Antônio Jorge G. Abelém	UFPA
Cesar Augusto C. Marcondes	UFScar
Iara Machado	RNP
JobertoS. B. Martins	UNIFACS
José Augusto Suruagy Monteiro	UFPE
José Ferreira de Rezende	UFRJ
Kleber Vieira Cardoso	UFG
Marco AntonioTorrez Roja	USP
Michael Antony Stanton	RNP
Tereza C. M. B. Carvalho	USP
Editor	Cesar Augusto C. Marcondes, Iara Machado
Reviewers	Leonardo Bergesio, Sebastià Sallent (i2CAT), Dimitris Giatsios (UTH)
Delivery Date	20/06/2012
Version	V1.0

Abstract

The purpose of this deliverable is to provide information about the defined CMFs (OFELIA, OMF and ProtoGENI) that are planned for the Brazilian Experimental Facility, also referred to here as the FIBRE-BR system. This information includes, a review of the CMFs chosen, an overall assessment of their compliance to the requirements acquired in T2.1, and an initial proposal plan for the deployment of the control and monitoring of the FIBRE-BR.

TABLE OF CONTENTS

1 Acronyms 5
2 Scope 7
3 Reference Documents 8
4 Introduction 9
5 CMF survey 10
5.1 OFELIA 10
5.2 OMF 14
5.3 ProtoGENI 17
6 Compliance to FIBRE-BR requirements 20
6.1 Methodology Adopted 20
6.2 Compliance and Gaps Identification 20
6.2.1 Req UR - User Requirement 27
6.2.2 Req UC – Use Case Requirements 28
6.2.3 Req AR – Architecture Requirements 29
6.2.4 Req ES - Experiment Support 31
6.2.5 Req IM – Instrumentation and Measurements Requirements 32
6.2.6 Req IF – Infrastructure Requirements 33
6.2.7 Req SE - Security Requirements 33
6.2.8 Req OP - Operations Requirements 34
7 CMF Definition and FIBRE-BR Design 36
7.1 CMF Definition 36
7.2 FIBRE-BR Monitoring Integration Approach 36
7.2.1 FIBRE-BR Instrumentation & Measurement (I&M) Architecture 37
7.2.2 MDIP Implementation Aspects 38
8 FIBRE-BR Initial Proposal Plan 40
8.1 Phase 0 40
8.2 Phase 1 40
8.3 Phase 2 40
8.4 Phase 3 40

1 Acronyms

AM	Aggregate Manager
API	Application Programming Interface
BT	Bluetooth
CDS	Content Data Server
CMF	Control and Monitoring Framework
CN-DS	Communication Networks and/or Distributed System
CNPq	Brazil's Council for Scientific and Technological Development
CPqD	Telecommunications Research and Development Centre
EC	Experiment Controller
ED	Experiment Description
EMULAB	Network Emulation Testbed
EU	European Union
FI	Future Internet
FIBRE	Future Internet testbeds / experimentation between Brazil and Europe
FOAM	FlowVisorOpenFlow Aggregate Manager
FP7	Seventh Framework Programme
FPGA	Field Programmable Gate Array
Gbps	Gigabits per second
GENI	Global Environment for Network Innovations
GMPLS	Generalised MultiProtocol Label Switching
HTTP	HyperText Transfer Protocol
IaaS	Infrastructure as a Service
ICT	Information and Communication Technologies
IP	Internet Protocol
I&M	Instrumentation and Measurement
INSTOOLS	Instrumentation Tools for a GENI Prototype
LAMP	Leveraging and Abstracting Measurements with perfSONAR
LDAP	Lightweight Directory Access Protocol
MA	Measurement Archive
MCS	Measurement Collection Server
MDIP	Measurement Data Integration Point
ML	Measurement Library
MS	Milestone
MTBF	Mean Time between Failures.
MTTR	Mean Time to Repair.
NaaS	Network as a Service
NITOS	Network Implementation Testbed using Open Source platforms

NM-WG	Network Measurement Working Group
NOC	Network Operations Center
NTP	Network Time Protocol
OMNI	OpenFlow Management Infrastructure
OCF	OFELIA Control Framework
OEDL	OMF Experiment Description Language
OF	OpenFlow
OFELIA	OpenFlow in Europe: Linking Infrastructure and Applications
OMF	cOntrol, Management and measurement Framework
OMNI	OpenFlow Management Infrastructure
OML	OMF Monitoring Library
ORBIT	Open-Access Research Testbed for Next-Generation Wireless Networks
RC	Resource Controller
RRD	Round Robin Database
SNMP	Simple Network Management Protocol
SOAP	Simple Object Access Protocol
SQL	Structured Query Language
XML	eXtensibleMarkup Language
XMPP	eXtensible Messaging and Presence Protocol

2 Scope

This document is the result of the task T2.3 of WP2 of the FIBRE project. The design and implementation of the FIBRE-BR control and monitoring framework (CMF) will be driven by this document.

3 Reference Documents

FIBRE's Description of Work, available under request to WP1.

FIBRE FP7 Project Public Deliverable 2.1

[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: 08-may-2012

[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", International ICST Conference on Testbeds and Research Infrastructures for the Development of Networks and Communities (TridentCom), 2010.

[Naous, 2011] JadNaous.Path-Policy Compliant Networking and Platform for Heterogeneous IaaS Management.PhD Thesis. Stanford. March 2011. [White, 2010] White, J. (2010). A Tutorial Introduction to OML. GEC9. Available at http://omf.mytestbed.net/projects/oml.

[Duerig et al. 2012]Duerig, J., Ricci, R., Stoller, L., Strum, M., Wong, G., Carpenter, C., Fei, Z., Griffioen, J., Nasir, H., Reed, J., and Wu, X. (2012). Getting started withGENI: a user tutorial. SIGCOMM Computer Communications Review, 42(1):72–77.

[Hanemann 2005]Hanemann, A., Boote, J. W., Boyd, E. L., Durand, J., Kudarimoti, L., Lapacz, R., Swany, D. M., Zurawski, J., Trocha, S., (2005) "PerfSONAR: A Service Oriented Architecture for Multi-Domain Network Monitoring", In "Proceedings of the Third International Conference on Service Oriented Computing", Springer Verlag, LNCS 3826, pp. 241-254, ACM SIGSOFT and SIGWEB, Amsterdam, The Netherlands.

[Swany and Boyd 2011] Swany, M. and Boyd, E. (2011). Leveraging and abstracting measurements with perfSONAR (LAMP). Available at: http://groups.geni.net/geni/wiki/LAMP. Last access: 02/21/2012.

4 Introduction

The Brazilian testbed, which is being designed and built in WP2, include 9 (nine) islands, one per each of the nine Brazilian partners in this project. Each island will be located at the corresponding partner's venue. Each individual island has OpenFlow enabled equipments, Orbit wireless nodes and a virtualized computer server. At USP, the FIBRE-BR island includes an EMULAB cluster.

The CMFs evaluated by Task 2.3 were OFELIA Control Framework (OCF), OMF, and ProtoGENI.

OFC because we are interested in using the same FIBRE-EU CMF to facilitate the federation and also, based on FIBRE-BR requirements, to investigate enhancements to the OFELIA CMF in which Brazilian researchers could collaborate.

The same rationale applies to the case of the OMF CMF.

ProtoGENi is the CMF for EMULAB currently existing in, at least, one of the Brazilian islands and we intend to integrate it in the final FIBRE-BR facility.

5 CMF survey

5.1 OFELIA

The OFELIA testbed is an initiative of the European Union 7th Framework Programme (FP7) that provides a unique large scale experimental facility devoted to the experimentation of OpenFlow related projects. This project has two major contributions to the international community, the public infrastructure itself (based on server virtualization and commercial OpenFlow switches –NEC and other vendors and switches) and, in addition, it provides an open-source version of the OFELIA Control Framework (OCF) synchronized with other initiatives in USA, following the evolution of Expedient and Opt-in manager [Naous, 2011], and the introduction of the VT-Manager, all publically available (Codebasin[1]) and intended to be of easy setup.

From the point of view of experimenters (or network researchers) the underlying network substrate available in OFELIA is fully controllable using explicit and dynamic configurations based on OpenFlow abstractions like FlowSpace. Once the FlowSpace is set up, the researcher can proceed with the allocation of a controller, either remote or in a local virtual machine to test his new idea. The resources are booked according to the experiment needs and accountability is done through a GENI-like slice abstraction.

As of this writing, the OFELIA facility has 8 islands with varied resources, from optical equipments to wireless devices, and these islands are to be interconnected using dedicated virtual tunnels among the islands. In each island there is a varied set of resources that will allow multi-layer and multi-technology experimentation of future networks. Another important point, is that OFELIA allows deep programmability of network resources granting researchers the use of the infrastructure as a Network-as-a-Service (NaaS) type of cloud service.

From the point of view of the OFELIA architecture, it was based on the control software, Expedient and Opt-In [Naous, 2011]. Moreover, it diverged in many ways from the original version. Expedient consists of several subsystems: an object relational mapping (ORM) database, a base platform subsystem, and at least two types of plug-ins, the so-called connectors and user clients. These plugins cooperate to provide the user interface for creating and managing slices across multiple island providers. The architecture depicted originally in Naous work can be represented by Figure 1. In OFELIA, regarding the Expedient sub modules, they were substantially improved, in particular the WebUI, that is the most used and important part for the experimenters and the database.

[1] CodeBasin OFELIA CF website - http://codebasin.net/redmine/projects/ocf/wiki

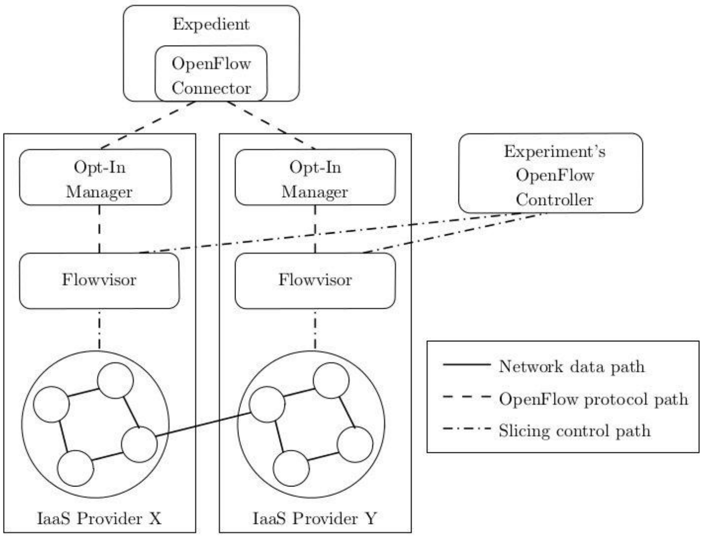

Figure 1: Expedient architecture.

Source: [Naous, 2011]

In order to slice the OpenFlow resources to support multiple experiments and users, the Expedient system, from which OFELIA inherits some of the most important features, make use of 4 layers: at the bottom (identified in the figure 1 as IaaS Provider resources are the OpenFlow switches, then on top of them, the Flowvisor, the Opt-In Manager, and finally the OpenFlow Expedient connector. The Flowvisor is the most necessary part to isolate and support multiple experiments. At an abstract level, Flowvisor is a transparent proxy between many islands with OpenFlow switches and OpenFlow controllers used in a certain slice. Basically, the Flowvisor(s) present in each island (or provider) monitors OpenFlow protocol messages from and to the controllers, ensuring that each slice defined by the FlowSpace (a set of header values defined to isolate experiments) operates on traffic within that space.

The OFELIA CF, in addition to the expedient and opt-in pair, make use of an additional resource of a virtualization service, in order to provide, virtual machines to the users as end-hosts or in order to, allocate local OpenFlow controllers to test the experiments. Figure 2 (extracted from Public Deliverable 5.1 of OFELIA FP7 Project) presents the main components of the OFELIA framework and the internal configuration of each OFELIA island.

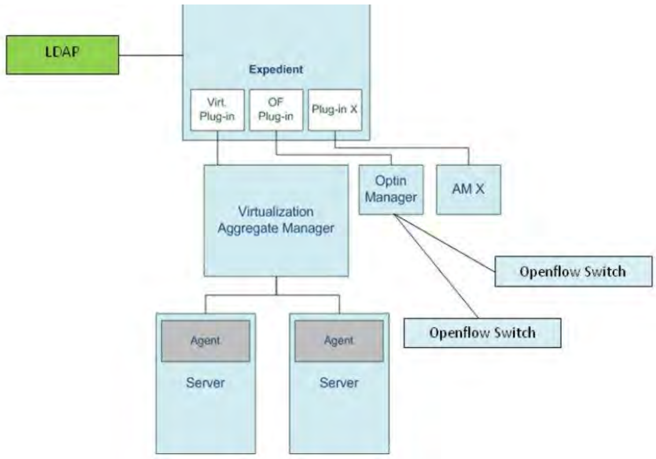

Figure 2 OFELIA: framework main components.

Source: OFELIA FP7 Project Public Deliverable 5.1

As it can be verified, the OFELIA CF system uses LDAP as the main software authenticator, and provides the Virtualization Aggregate Manager based on Xen machines (called OXA – OFELIA Xen Agent) which can allocate on-demand virtual machines.

Inside each OFELIA island, there is a complex set of configurations that allows researchers to experiment remotely through VPN connections, and the slicing is done using Flowvisor (FV in the figure 3) that virtualizes the OpenFlow switches (X in Figure 3) to many controllers (NOX red and green).

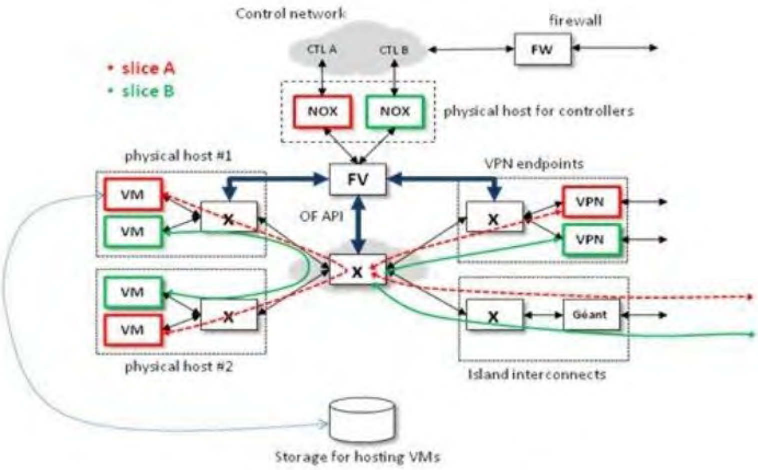

Figure 3: Basic OFELIA island.

Source OFELIA FP7 Project Public Deliverable 5.1

During the initial months of the FIBRE project, the UFScar team set a challenge to replicate, in a small scale, all the settings of the i2CAT OFELIA island. During that period, had a very supportive help from i2CAT, a total of 100 emails from January to March. They managed to install a local experimental testbed in a small set of 4 machines, one server, two XEN agents and one netFPGA-based Openflow switch. During that time, it was created a full mini-course on OFELIA, from how to install it, how to manage it, how to use it and which was presented at SBRC (Brazilian Computer Networks and Distributed Systems Symposium 2012). Additional contributions from the UFSCar team included:

- Creation a new template virtual machine file and hack the "ofver" setup to point to a different TEMPLATE URL, given that the OFELIA template VM was strictly tied to the LDAP / VPN infrastructure of i2CAT.
- Discovery of a set of minor configuration issues in the installation scripts, such as the kernel VM that was hardcoded at the boot, assuming amd64, thus it was necessary to change to x86 in the local files (not "ofver").
- Discovery that 2 VMs at the same physical agent can ping each other without the openflow controller due to the Linux bridge used by Xen permissiveness.
- Trial to install a set of Linksys WRT54G switches with the firmware replaced by openwrt + pantou, however without success till now. The Openwrt system in order to do port isolation requires the equipment to rewrite VLANs with specific numbers (1 to 5) instead of using the VLAN numbers used by virtual machines and managed by Flowvisor. This impedes the Linksys switches to be used as openflow switches integrated with the OFELIA CF.
- Trial to interconnect several netFPGA-based OpenFlow switches in a linear topology but they did not appear in the topology discovery applet. If only one netFPGA is used, there is no problem. The problem seems to be related to a bug in the netFPGA that do not communicate correctly with the Flowvisor Topology Discover feature. This is going to be further analyzed since all the Brazilian islands are going to usenet FPGA-based OpenFlow switches.
- Finding of some issues on setting the controller over the VPN (a remote controller), especially if it requires the restart of the slice from the Flowvisor point of view.
- Discovery of many operational issues during the months of testing, including things like DHCP errors, VMs didn't start due to LDAP, firewall problems, VPN setup not activated in one OFELIA site (a security breach), but all these were fixed in a short time once notified.

Based on our experience with OFELIA, we strongly recommend it as the main tool for the OpenFlow-based infrastructure part of the FIBRE-BR system. The software is robust, flexible, and easily manageable and already in production at the European partners.

The Evolution of the OFELIA Control Framework

The OFELIA project is in development and its CMF is undergoing constant modifications to suit the new demands of Future Internet experimentation. In its current phase, the OFELIA CMF consists of Expedient software, OPT-IN and VT AM. During the next phases of the OFELIA project (Phase 1 and Phase 2) they will be migrating to the software FOAM and OMNI, with a more flexible architecture and inclusion of new features.

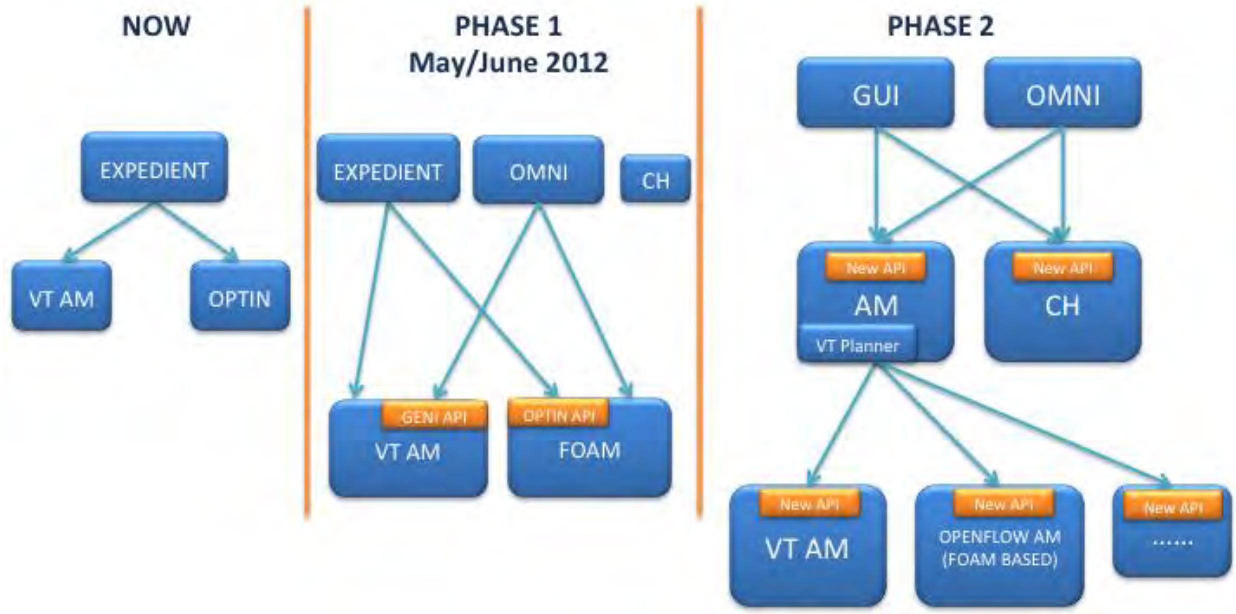

Figure 4: Planned Evolution of OFELIA CMF.

Source: Slides presented by Leonardo Bergesio during Architecture Meeting in 5th march 2012 by videoconference

FOAM is an aggregate manager for OpenFlow resources in the GENI infrastructure. It uses GENI v3 RSpecs with OpenFlow extensions. FOAM will replace the OpenFlow AM Opt-in Manager . Currently, it is only an Aggregate Manager, not yet an end-user opt-in. Foamctl is the primary interface for administrative control of FOAM. This is the interface by which administrator can approve and reject slivers, as well as configure essential FOAM options.

OMNI is a GENI experimenter tool that communicates with GENI Aggregate Managers via the GENI AM API. In summary, OMNI is a command line client for the GENI Aggregate Manager API. The OMNI client can also communicate with control frameworks in order to create slices, delete slices, and enumerate available GENI Aggregate Managers.

5.2 OMF

OMF [Rakotoarivelo 2010] is a framework developed in Ruby language and based on XMPP (eXtensible Messaging and Presence Protocol) with the focus on controlling and managing network devices. The OMF suite also provides OML (OMF Monitoring Library), which allows instrumentation of applications for collecting measurements. Actually, OMF depends on a suite of software components in order to provide services such as dynamic IP addressing, remote booting, remote OS imaging, data archiving, etc.

OMF was originally developed for the ORBIT wireless testbed at Winlab, Rutgers University. This former software only supported the specific ORBIT hardware resources (i.e. custom built static nodes with 802.11wireless devices), and provided a limited experiment description language. In contrast, the current OMF [OMF 2012] supports a large number of different wired and wireless resources, such as PC-like devices, Android phones, routers, sensors, among others.

The OMF architecture consists of 3 logical planes: control, measurement, and management. This architecture is not completely clear in version OMF 5.4, but it is the approach defined for the forthcoming version 6. The control plane includes the OMF tools that a researcher uses to describe his/her experiments, and the OMF entities responsible for orchestrating them. The measurement plane includes the OMF tools to instrument an experiment, and the corresponding OMF entities to collect and store measured data. The management plane includes the OMF functions and entities to provision and configure the resources, which are provided by the testbed facilities and used by the experiments.

Figure 5 presents a system view of an OMF deployment. The main OMF elements are detailed below.

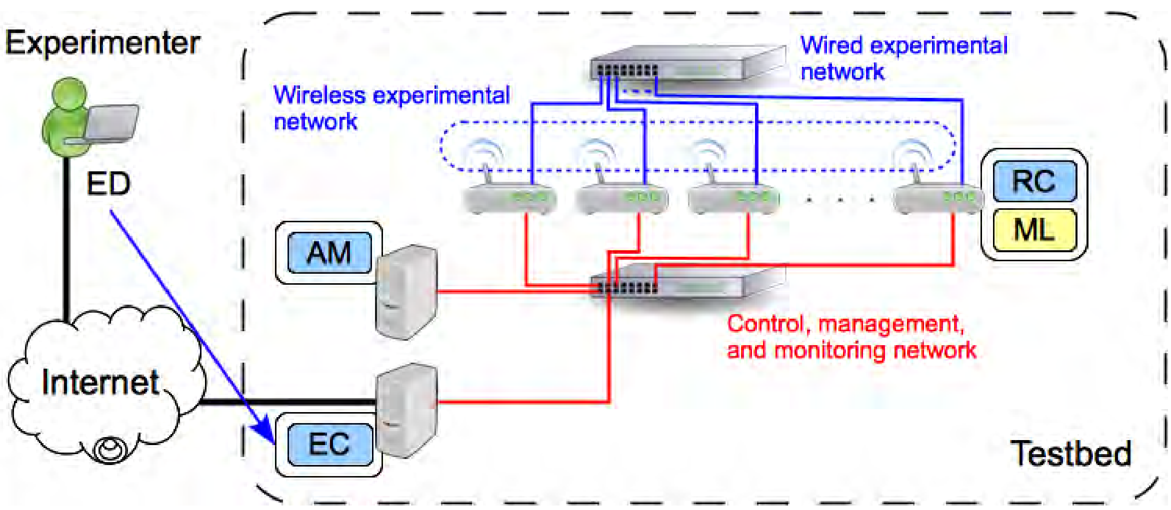

Figure 5: OMF system overview.

A domain specific language (the OMF Experiment Description Language – OEDL) allows an experimenter to write an Experiment Description (ED), which details the resource requirements, their initial configuration, and a state machine describing the time/event-triggered actions required to realize the experiment.

The experimenter submits the ED to an Experiment Controller (EC), which is the control entity orchestrating the experiment on behalf of the experimenter. The EC issues requests on the management plane to configure the resources as specified in the ED. Once the experiment prerequisites are met, the EC sends directives to the Resource Controller (RC) associated with each resource. Actually, the configuration of the resources is also done by the RC. Following the ED's state machine, the EC sends directives to the RCs, which are the control entities responsible for executing actions on the resources.

The measurement plane consists of two parts, namely the Measurement Collection Server (MCS) and the Measurement Library (ML). OML allows recording of any relevant metrics and instrumented applications. The resulting streams are forwarded to the MCS either in real-time or batch mode to minimize interference with the experiment itself. The MCS stores them in a SQL database created for each experiment instance. The user can directly run SQL queries against it, or retrieve a data dump from it. Standing queries can feed events back into the EC to support steerable experiments.

The batch mode is monitoring part the disconnected operation mode, which also involves specific configuration in OMF elements. The capability of performing experiments with partial disconnection of mobile devices is an important feature which allows tests with smart phones, cars, DTN devices, etc.

Another important feature is the capability to visualize the measurements at real-time on the experimenter's Web browser. This is useful for the beginners, but it can also help the experts to easily verify if the expected results are being generated at very begin of an experiment.

OMF combines a set of management services into an Aggregate Manager (AM). Thus, the AM is a collection of services, which can be deployed across multiple servers for performance and redundancy reasons. AM accepts request from the ECs or the testbed operator, and sends corresponding commands to the RC running on each resource. MCS may be treated as one of the services of the AM.

Experience with OMF

OMF provides an easy tutorial about its installation, operation, and use in a testbed. There are some missing details in the tutorial, but they can be easily discovered by searching in configuration files, source code or examples.

Versions 5.3 and 5.4 have been evaluated, and this last one is running stable at UFG. Also they has run many tests with 802.11, Ethernet and virtualized devices. In general, the tests involved data transfers employing an instrumented Iperf. We have run UDP and TCP tests, and these tests have worked without problems. The results were the expected ones and the collected measurements were correctly persisted in the database. We have also successfully tested the functionalities of turning on/off the nodes, checking nodes status, and OS imaging.

The main drawbacks found till now for OMF CMF are some missing functionalities, mainly:

- resource allocation; and
- authentication and authorization.

We have tried to integrate OMF with NITOS Scheduler [Anadiotis 2010] in order to get these functionalities. 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, which have kindly answered a lot of questions.

In general, OMF is a lightweight service, but its simplest installation is based on Openfire XMPP server, which is written in Java and has a high demand for CPU and memory resources. Besides, the installation tutorial employs Dnsmasq as the software for dynamic IP configuration, name resolving, and remote boot. For a small prototype testbed, this software is fine. However, a large testbed intended to be in production eventually will need a more mature and stable software, e.g. ISC DHCP and BIND.

5.3 ProtoGENI

ProtoGENI is a control and monitoring framework software from the University of Utah. It is based on an enhanced version of the Emulab testbed management software. The Emulab testbed is used to perform experimental network and research about distributed systems.

ProtoGENI was created to provide the integration between Emulab and other testbeds in order to build the Cluster C facility of GENI. This control framework was design based on the SFA (Slice-based Facility Architecture) for federation among the testbeds.

The SFA components implemented by ProtoGENI are:

- The Clearing House behaves as a central registry, coordinating the federation;
- The Aggregate Manager and Slice Authority are responsible for the operation and resource control;
- The RSpec is a data interchange format to describe resources properties consistently between testbeds.

Figure 6 presents an overview of the ProtoGENI federation architecture, presenting ProtoGENI main components and their interactions.

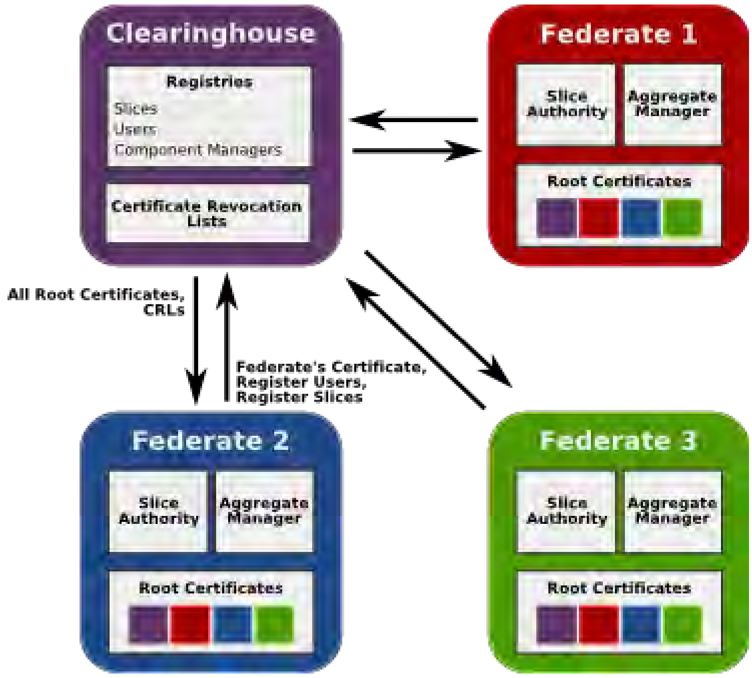

Figure 6: ProtoGENI federation architecture.

Source: ProtoGENI official site http://www.protogeni.net

ProtoGENI provides the interface to allocate and manage resources from Emulab testbed and backbone nodes:

- The Emulab resources are composed of local nodes and virtual network links;
- The backbone nodes consist of PCs, NetFPGA cards and an Ethernet switch.

For the utilization of ProtoGENI, the experimenter must define the desired resources in a RSpec file or using the Flack interface (a Flash-based GUI). The RSpec file can also contain the network topology of the experiment. After the submission, the requested resources will be allocated automatically and the virtual links between the nodes will be created as it was defined in the experiment topology. The experimenter can access the provided resources by SSH (for any node), as well as configure and run the experiment. The instrumentation and monitoring of the experiments are achieved by the ProtoGENI tools allowing passive and active measurements. The INSTOOLS toolkit can be used mainly for passive measurements while LAMP (based on perfSonar) can be used for active measurements and some passive measurements.

Experience with ProtoGENI

An Emulab testbed was created at USP and is being used since 2008. In order to integrate this testbed with the FIBRE-BR infrastructure the ProtoGENI CMF is required.

The ProtoGENI CMF installation files for Emulab are available within the Emulab source code. To enable ProtoGENI features in Emulab, the main server's configuration was changed, by turning on the ProtoGENI flag, and the server was reinstalled.

Any new ProtoGENI installation federates by default with the ProtoGENI from Utah. It is possible to create your own federation, independent from Utah, but the Clearinghouse also needs to be installed and new certificates must be created. The USP testbed was configured to be the first island ofan independent federation.

The infrastructure was verified by executing some experiments. These experiments were developed to test the resource description format, the topology creation interface, the resource allocation, and the management services.

6 Compliance to FIBRE-BR requirements

6.1 Methodology Adopted

In order to verify the compliance of the FIBRE-BR requirements with the CMFs evaluated, each specific requirement was analyzed for each CMF and one spreadsheet was elaborated in order to identify requirement gaps for each CMF.

For each requirement, the methodology consists in labelling the CMF support the functionality as follows:

- Yes
- No
- Partially
- Do not apply
- Planned

The requirements are grouped by "type of requirement". Colours are used to represent conformance or not and other information in relation to the analysed CMF as follows:

"**Yes"** (green), "**No**" (red), "**Partially**" (yellow), "**Do not apply**" (grey), "**Planned**"(blue);

This color code is similar to the regular traffic lights. In addition, we color the " Do not apply" as Grey, in the sense that for the evaluator of the CMF this shouldn't matter, and we color the "Planned" as Blue, corresponding to blueprints for the requirements. We start by presenting the colortable, presenting the results for every requirement number in a certain type of requirement (the rows) – the numbers are the same from D2.1, and the evaluations done by each partner (the columns).

The I&M (Instrumentation and Monitoring) evaluation was done only in relation to certain requirements specifically related to I&M.

Comments to explain conformance, gaps or additional comments are also eventually included.

6.2 Compliance and Gaps Identification

This section presents a summarization and brief discussion about the requirements evaluation process in order to propose the solution space for the FIBRE D2.1 requirements. In the regards of the identification of these gaps, we used an earlier version of D2.1 that has a slightly larger number of requirements to fulfil. Therefore, in order to make easier to understand, we will re-describe all of them below.

[Req UR01] The FIBRE-BR system MUST support researchers with different levels of experience, ranging from graduate students to experienced research engineers.

[Req UR02] A FIBRE-BR experimenter MUST provide their credentials to be authenticated and authorised by the system, in order to have remote access to the experimentation facilities.

[Req UR03] A FIBRE-BR experimenter MUST be able to discover available resources in the experimentation facilities, and to reserve those needed for his or her experiment, in such a manner as not to interfere unnecessarily with other experiments being carried out concurrently.

[Req UR04] A FIBRE-BR experimenter MUST be able to configure the allocated resources.

[Req UR05] During the period reserved for an experiment, a FIBRE-BR experimenter MUST be able to start, pause, stop or reset an experiment at any given moment.

[Req UR06] An experimenter MUST have full control over the experiment running.

[Req UR07] An experimenter or the facility operator MUST be able to terminate an experiment that, due to mal-function, is affecting negatively the performance or behaviour of other experiments or even of the facility as a whole.

[Req UR08] The system SHOULD enable an experimenter to change the configuration of the allocated resources, moving back to a previous snapshot of the "experiment deployment".

[Req UR09] Taking into consideration the resource limits defined in step "resources allocation", the experimenter MUST also be able to grow or shrink the elements involved in the experiment.

[Req UR10] A FIBRE-BR experimenter MUST be able to archive the experiment setup and its results, in order to permit post-mortem analysis.

[Req UR11] Experiment archives SHOULD be made available for other experimenters to validate, re-run, and make comparisons.

[Req UR12] An experiment archive MUST include a human-readable description of the experiment, as well as setup information, experiment results and any special considerations for repeating the experiment.

[Req UR13] An experiment archive SHOULD also include the experiment description in a formal language.

[Req UR14] A FIBRE-BR experimenter MUST have access to FIBRE-BR usage guidelines, and to a communication channel to report problems.

[Req UC01] The FIBRE-BR system MUST support the following components: content sever, OpenFlow switch, WiMax base station, Orbit nodes, OpenFlow-compatible nodes, GPS nodes.

[Req UC02] The FIBRE-BR system SHOULD support the export to the experimenter of the physical properties of RF channels, such as signal-to-noise relation or interference conditions near nodes, when available.

[Req UC03] The FIBRE-BR testbed MUST provide the means for researchers to investigate all issues involved vertical handoffs.

[Req UC04] To allow for investigation into vertical handoffs, nodes involved (laptops or smartphones) MUST possess two or more interface types, such as WiFi, WiMax, GSM, etc.

[Req UC05] Some predefined report templates, based on the experiment type, SHOULD be provided. The FIBRE-BR system MAY also provide the means to generate customised reports based on the selection of some of the metrics made available in the experiment logs.

[Req UC06] The FIBRE-BR system SHOULD allow user profile data to be exported and/or shared, so that the parameters already selected can be imported by any other user of the environment, subject to the owner's permission.

[Req UC07] The FIBRE-BR system MUST support the following components: OpenFlow switches, dynamically-reconfigurable WDM equipment with OpenFlow support, content delivery server to store HD and 4K quality media, content delivery application to control delivery of streams to subscribers, 4K projection facility to display 4K-quality media streaming that is compatible with the FIBRE UK island́s projection facility, HDTV to display HD-quality media streams

[Req UC08] The FIBRE-BR system MUST provide enough capacity to store 4K experiment data, i.e. 3 or 4 4K-videos (~ 5TB).

[Req UC09] The participating FIBRE-BR island in this use case MUST provide energy, air conditioning, cabling, racks, floor space and physical security to host servers, projectors and communications equipment.

[Req UC10] In order to participate in this Use Case, FIBRE islands in EU and in BR MUST be interconnected by high performance network links. More specifically, content delivery servers MUST be connected, and it MUST be possible to connect any user with one or another content delivery server at a time.

[Req UC11] In order to run experiments defined in Use Case 3, the participating FIBRE-BR island MUST deploy a GMPLS Stack as a NOX application, in order to be able dynamically to reconfigure the switching elements to redirect users from one server to another, according to criteria to be defined on demand.

[Req AR01] The FIBRE-BR architecture MUST enable federation, which is the management and inter-operation of multiple independently administered resources, referred to here as islands, which are owned by multiple distinct organisations both participating in the FIBRE project and external to it, and possibly using different wired and wireless technologies.

[Req AR02] The FIBRE-BR system SHALL provide a framework for the communication among federated organizations, which enables the inter-operation of the Brazilian and European islands, as well as with other external testbeds.

[Req AR03] The FIBRE-BR system SHALL provide the means to enable organizations to carry out secure authentication of users from federated organizations, possibly providing different levels of clearance and access according to organization policies and the needs of the users.

[Req AR04] The FIBRE-BR system SHALL allow administrators to manage the federation. This involves activities such as including new islands, installing and updating system federation software, and monitoring the federation in terms of performance, functionality, and security.

[Req AR05] The FIBRE-BR system SHALL allow federated organizations to describe resources that they contribute to the federation in a common format.

[Req AR06] The adopted format SHOULD be flexible enough to permit the description of resources from different technologies.

[Req AR07] The FIBRE-BR system SHALL provide mechanisms to allocate resources between experiments being run concurrently.

[Req AR08] The FIBRE-BR system SHALL provide mechanisms to enforce resource sharing policies between experiments being run concurrently.

[Req AR09] The FIBRE-BR system SHALL allow federated organizations to declare usage policies for substrate facilities under their control, and to provide mechanisms for enforcing those policies.

[Req AR10] The FIBRE-BR system SHALL include a portal that allows experimenters to acquire, schedule, and release federated resources.

[Req AR11] The FIBRE-BR system SHALL offer access to the federated resources through open interfaces.

[Req AR12] The FIBRE-BR system SHALL provide a framework for gathering measurements and results from experiments run across multiple federated resources.

[Req AR13] The FIBRE-BR system SHALL federate different experimental facilities and technologies and, additionally, be extensible enough to support future technologies still to be developed.

[Req AR14] The FIBRE-BR system MUST provide customisable mechanisms for routing, signalling, resource allocation and the acceleration of information processing.

[Req AR15] The FIBRE-BR system MUST provide tools for configuring, managing, monitoring and debugging end-to-end slices.

[Req AR16] The FIBRE-BR system MUST include a means for providing isolation between slices (except when the slices share resources), in order to ensure that independent experiments do not interfere with one another.

[Req AR17] These isolation mechanisms MAY vary for different types of experiments.

[Req AR18] These isolation mechanisms MUST be sufficiently robust to make

reproducible experiments possible.

[Req AR19] A user MUST be able to assess the level of isolation of his slice, and the degree of interference caused by other slices.

[Req AR20] The FIBRE-BR system MUST provide mechanisms to control resources, such as adding resources to and/or subtracting them from an existing slice, in order to grow or shrink a long-lived experiment.

[Req AR21] Thus, the researcher MUST NOT be forced to tear down and re-create a slice, in order to make changes to its configuration.

[Req AR22] The FIBRE-BR system MUST support controlled interconnection of slices to each other and to external networks, such as the current Internet.

[Req AR23] The FIBRE-BR system MUST provide integrated monitoring and measurement capabilities across all provided facilities.

[Req AR24] The FIBRE-BR system MUST support wired, optical and wireless environments and interconnections between them.

[Req AR25] The FIBRE-BR system MUST provide integrated resource management (reservation, announcement and discovery) for all provided facilities.

[Req AR26] The FIBRE-BR system MUST provide support for multiple wireless technologies (e.g. Ad-Hoc, Wi-Fi, WiMax, Bluetooth, etc.).

[Req AR27] The FIBRE-BR system MUST support horizontal and/or vertical handover in wireless facilities.

[Req AR28] The FIBRE-BR system MUST provide explicitly defined and extensible system interfaces to facilitate the incorporation of additional technologies.

[Req ES01] The FIBRE-BR system MUST provide the basic tools to distribute and/or install experimental code and data over every island.

[Req ES02] The FIBRE-BR system MUST provide a basic interface for the experimenter to be able to check the status of his experiment.

[Req ES03] The FIBRE-BR system MUST provide a basic lightweight mechanism to monitor experiment misbehaviour and trigger instant alarms in case of problems.

[Req ES04] The FIBRE-BR system MUST provide a basic method for terminating an experiment, either by using timers or a stop button, and the system MUST provide clean-up routines to release facility resources allocated to an experiment.

[Req ES05] The FIBRE-BR system MUST make an experiment accountable for any misuse of resources.

[Req ES06] The FIBRE-BR system SHOULD provide support for experimental realism. This implies adding support to automatic tools and features that make the environment more realistic and repeatable.

[Req ES07] The FIBRE-BR system SHOULD support repeatable behaviour for any selected experiment.

[Req ES08] The FIBRE-BR system MUST provide a basic mechanism to support intentional degradation of some components.

[Req ES09] The FIBRE-BR system MUST provide support for a virtualized management plane, separate from the data plane, which could be used by system operators for low-level monitoring of experiment resources in any phase.

[Req ES10] The FIBRE-BR system MUST provide a broad range of realistic background traffic to be added to the experiment at any time.

[Req IM01] The FIBRE-BR experimental facility MUST include an Instrumentation and Measurement (I&M) subsystem in order to provide data gathering, analysis, and archival capabilities to network operators, experimenters and island administrators in a federated environment.

[Req IM02] The FIBRE-BR I&M subsystem MUST provide performance and operational measurements, including power usage, of the physical equipment, which provides support to the FIBRE-BR system itself.

[Req IM03] The FIBRE-BR I&M subsystem MUST provide a monitoring interface for experiments.

[Req IM04] The FIBRE-BR I&M subsystem SHALL implement a mechanism to protect its measured data as defined by the FIBRE-BR access policy.

[Req IM05] Besides the equipment that is part of the substrate, optical, wired and wireless links SHOULD be monitored.

[Req IM06] The FIBRE-BR I&M subsystem MUST allow measured data to be collected and stored.

[Req IM07] The physical location of the FIBRE-BR components SHOULD be informed to the users.

[Req IM08] Time synchronization is essential for any measurement infrastructure. Ideally, accuracy in the order of microseconds SHOULD be provided. However, due to cost concerns, in some components accuracy in the order of milliseconds can be tolerated.

[Req IM09] The FIBRE-BR I&M subsystem SHALL support integration among wired and wireless islands, and SHOULD support integration with external I&M infrastructures.

[Req IF01] The FIBRE-BR facility MUST include a collection of building block components permitting the creation of virtual networks.

[Req IF02] The FIBRE-BR infrastructure MUST be expansible, permitting new components to be added after facility construction.

[Req IF03] The infrastructure MUST include the computational resources necessary to build wide-area services and applications, as well as initial implementation of new network elements.

[Req IF04] The infrastructure SHOULD include core network data processing functions for high-speed, high-volume traffic flows.

[Req IF05] The infrastructure MUST include data forwarding functionality at the boundary between access networks and a high-speed backbone.

[Req IF06] The infrastructure MUST include proxies and other forwarding functionality within a wireless network.

[Req IF07] The infrastructure MUST include mobile devices intended to run applications that give end-users access to experimental services available on the combined wired/wireless substrate.

[Req IF08] The infrastructure MUST include lightpath interconnection between FIBRE-BR core nodes, forming a nationwide backbone network supporting rates of at least 1 Gbps.

[Req IF09] The infrastructure MUST include access circuits to connect FIBRE-BR Edge Sites to the FIBRE-BR Core, supporting rates of at least 1 Gbps.

[Req IF10] The infrastructure MUST include at least one 802.11-based mesh wireless subnet to provide real-world experimental support for ad-hoc and mesh network research based on an emerging generation of short-range radios.

[Req IF11] The infrastructure SHOULD include at least one Wide-Area 3G/WiMax- based Wireless Subnet intended to provide open-access 3G/WiMax radios for wide area coverage, along with short-range 802.11 class radios for hotspot and hybrid service models.

[Req IF12] The infrastructure SHOULD include at least one cognitive radio subnet intended to support experimental development and validation of emerging spectrum allocation, access, and negotiation models.

[Req IF13] The infrastructure SHOULD include at least one Application-Specific Sensor Subnet capable of supporting research on both underlying protocols and specific applications of sensor networks.

[Req IF14] The infrastructure MUST include at least one Emulation Grid that allow researchers to introduce and utilize controlled traffic and network conditions within an experimental framework.

[Req IF15] The institution MUST support enough space for at least one FIBRE-BR rack with the components described above. These racks MUST be energized and air conditioned according to project requirements.

[Req IF16] The institution MUST provide a dedicated access link to the experimental facilities in order not to create disturbances in the production network.

[Req SE01] All participating experimental facilities MUST adopt and maintain the security best practices for networking, operating systems, and applications provided by FIBRE Project.

[Req SE02] The FIBRE-BR system SHALL only assign an identity to a user that was approved in a defined registration process.

[Req SE03] All components, institutions, researchers, and slices MUST be assigned a unique identity.

[Req SE04] The FIBRE-BR system MUST require user compliance with the defined authentication and authorisation policies before the resource allocation.

[Req SE05] The FIBRE-BR system MUST permit system activity to be audited and traced back to the responsible entity.

[Req SE06] The FIBRE-BR system MUST provide the least privilege to each component of the system in order to grant exactly the privileges it needs to perform its tasks and no more.

[Req SE07] The FIBRE-BR system MUST provide access control mechanisms to ensure that resources are only granted to those users who are entitled to them.

[Req SE08] The FIBRE-BR system MUST provide mechanisms to establish secure communications between system entities.

[Req SE09] The FIBRE-BR system MUST ensure integrity and availability of the experimenters' data during the lifespan of the experience.

[Req SE10] The FIBRE-BR system MUST isolate slices from each other, so that a sliver cannot access objects (e.g., files, ports, processes) or eavesdrop network traffic from another slice, without permission.

[Req SE11] The FIBRE-BR system MUST isolate slices from the experimental facility infrastructure, so that a slice cannot compromise facility operation.

[Req SE12] The FIBRE-BR system MUST provide mechanisms to establish trust relationships between system entities.

[Req SE13] The FIBRE-BR system MUST support authorisation and access control of users from other experimental facilities.

[Req SE14] The FIBRE-BR system MUST provide mechanisms to guarantee that only authorised traffic is allowed among experimental facilities.

[Req SE15] The FIBRE-BR system MUST provide control mechanisms for inbound/outbound traffic, in order to prevent attacks and/or unauthorized access from/to the Internet.

[Req OP01] The FIBRE-BR system MUST be supported by a minimum support infrastructure.

[Req OP02] The FIBRE-BR system MUST at least offers 3 levels of support.

[Req OP03] The FIBRE-BR system MUST use a ticketing system tool as a communication channel to record issues users meet in using the system.

[Req OP04] The FIBRE-BR ticketing system SHOULD be accessible to users and to different levels of support, and MUST allow a ticket to be assigned to different teams and/or individuals.

[Req OP05] The FIBRE-BR system MUST be able to send notifications by e-mail and maintain schedules and other timestamps, such as opening and closing times of events. One desirable feature would be the ability to share maintenance calendars of the different islands.

[Req OP06] The FIBRE-BR system MUST be intuitive to use and user friendly.

[Req OP07] The FIBRE-BR system SHOULD allow the recording of important events, the definition of workflows and the customisation of forms.

Relevant information that SHOULD be recorded in the ticketing system will include: island information, user information, events.

[Req OP09] Every island SHOULD operate a highly operational environment.

[Req OP10] All components of the infrastructure SHOULD be monitored to ensure visibility of the performance of each island.

[Req OP11] During the evolution of the project, targets SHOULD be defined based on the current operational environment.

[Req OP12] A monthly report of the components of each island SHOULD be generated and made available through the project Wiki.

[Req OP13] A centralized monitoring infrastructure SHOULD maintain monitoring information about all elements of each island and generate a monthly report, to be made available through the project Wiki.

[Req OP14] This monitoring infrastructure should monitor the following metrics for each component: availability, downtime, number of failures, MTTR - Mean Time to Repair, MTBF - Mean Time between Failures, downtime for maintenance, availability (ignoring downtime during scheduled maintenance).

[Req OP15] The FIBRE-BR system MUST provide the means for researchers to identify, define and choose specific flows and monitored nodes in a categorised way, offering filtering by node local, node model, traffic type, etc.

[Req OP16] The FIBRE-BR system MUST provide the means to classify traffic registers (logs) as active or passive traffic. Passive traffic is a genuine use of the facility with no specific testing purpose, while active traffic is data generated for a specific experiment.

[Req OP17] The FIBRE-BR system MUST provide the means to discover whether other experiments are in progress at a given time.

[Req OP18] The FIBRE-BR system MUST provide a single screen with all scheduled experiments, and MAY provide the means for researchers to share their logs and results online.

[Req OP19] The FIBRE-BR system MUST provide the means for researchers to schedule their experiments by the selection of a minimum set of parameters.

[Req OP20] The FIBRE-BR system MUST generate logs in open source format.

The following sections describe the gaps identified and the solution space for them.

6.2.1 Req UR - User Requirement

These are the "user requirements", which are sets of features that the CMF should support in order to help the user in the experiment accomplishment.

Requirement	OFELIA	FOAM	OMF	ProtoGENI	I&M OFELIA	I&M OMF	I&M ProtoGENI
Req UR01							
Req UR02							
Req UR03							
Req UR04							
Req UR05							
Req UR06							
Req UR07							
Req UR08							
Req UR09							
Req UR10							
Req UR11							
Req UR12							
Req UR13							
Req UR14							

It follows a brief discussion about user requirements not supported by specific CMFs.

UR8 - This requirement is about the need of experiment snapshots (pause and resume for later use). Currently, from the point of view of OFELIA, the VT_Manager sends commands to a Xen daemon. Being so, it is intrinsically supportable but this is not in the interface neither it is automatized. Since the experimenter could do it manually, we review the requirement lowering the need from MUST to SHOULD and we will incorporate changes in the VT_Manager or either other CMF tools that evolve.

UR9 - In OFELIA, growing or shrinking resources can be done at any time, for example, one can redefine the flowspace or network resources. Regarding VMs, it is always possible to add/delete VMs, although, it could be not possible to change the VM configuration (HD size, RAM, once created). However, restarting the slice is not the same as restarting the experiment.

UR10 - In OFELIA 0.1, there exist the possibility of storing the experiment data through a remote access by the experimenter. However, since the virtual machine (end-host) is based on a template, at every new experiment it generates a VM zeroed from the template. Thus, at every new restart tools and results could be lost. This functionality is planned to be incorporated in next releases of VT_Manager where each user could have his own VM and storage block. Other possibility is to document that the transfer of results should be done by the experimenter as the part of the experiment cycle.

UR12 - OFELIA supports this requirement partially (high level description of projects, slices and flowspaces). As such, a more detailed documentation could be done manually, since OFELIA system differently than OMF and ProtoGENI has no formal configuration of topology and order of events as it is done in ns-2 format (protoGENI) and OEDL format (OMF).

6.2.2 Req UC – Use Case Requirements

These are "use-case specific requirements". Since the project has a goal to present showcases and demos, there is a need to support some specificity in certain islands. Since FIBRE-BR is adopting a multi-CMF structure, many requirements that are supported by one CMF and not supported by another CMF are in compliance due to this approach.

Requirement	OFELIA	FOAM	OMF	ProtoGENI	I&M OFELIA	I&M OMF	I&M ProtoGENI
Req UC01							
Req UC02							
Req UC03							_
Req UC04							
Req UC05							
Req UC06							
Req UC07							
Req UC08							
Req UC09							
Req UC10							
Req UC11							

UC1 - Since ProtoGENI has the current support of only few types of resources, like PCs, emulation and virtual machines, it may not support alone the diversity of heterogeneous resources of FIBRE. However, in the case of FIBRE-BR, as mentioned, we will use more than one CMF at the same time. In other words, FIBRE-BR will work with OFELIA, OMF and

ProtoGENI and whatever resource that ProtoGENI do not support, it will be supported by the rest.

UC2, UC3, UC4 - Similar reasoning as UC1, we take for granted that OMF will support many wireless resources, included the ones planned in FIBRE itself, like LTE and Wimax.

UC5 and UC6 - These are monitoring related requirements and are discussed in the I&M requirements.

UC7 - This is a specific requirement of certain islands. CMF should support minimally an interface so that the experimenter can access the resource by regular means available at the CMF. For example, the reconfigurable WDM-openflow switch could be seen just a regular openflow switch from the OFELIA CMF point of view, therefore transparently to the user. Other example is about the 4K video streaming servers, since they are too demanding, the idea is to allocate them in a physically separated machine and access it remotely from a VM from OFELIA and control it remotely using OMF.

UC11 - Same reasoning as UC7.

6.2.3 Req AR – Architecture Requirements

These requirements raise the need of federation support from the CMFs and the support of typical abstractions used to accomplish the federation such as slices for accountability and clearinghouse for authentication. Therefore, several of these requirements are seen just as if the CMFs provide the mechanisms to do it, not the way the federation, messages and semantics, is going to be done. All the federation detailed discussion and specifications are going to be considered in WP4.

Requirement	OFELIA	FOAM	OMF	ProtoGENI	I&M OFELIA	I&M OMF	I&M ProtoGENI
Req AR01	blue	blue	blue	green			
Req AR02	blue	blue	blue	green			
Req AR03	grey	yellow	blue	green			
Req AR04	blue	blue	blue	yellow	blue	grey	yellow
Req AR05	blue	blue	blue	green			
Req AR06	green	green	blue	yellow			
Req AR07	green	green	blue	red			
Req AR08	blue	red	blue	red			
Req AR09	blue	blue	blue	red			
Req AR10	blue	blue	blue	green			
Req AR11	blue	red	green	green			
Req AR12	red	red	yellow	yellow	blue	yellow	yellow
Req AR13	blue	green	blue	yellow	blue	grey	yellow
Req AR14	green	green	yellow	red			
Req AR15	green	yellow	blue	yellow			
Req AR16	green	green	blue	yellow			

Req AR17							
Req AR18							
Req AR19							
Req AR20							
Req AR21							
Req AR22							
Req AR23							
Req AR24							
Req AR25							
Req AR26							
Req AR							
Req AR28							

AR7, AR8, AR9 - Since the beginning, Expedient and GENI (and OFELIA) we designed to support users from outside (the Internet) to participate in some experiment within the experimental infrastructure. Therefore, this makes possible also for one experiment participate, or be the input, for other experiment using a web interface (opt-in users). This is something planned to be tested in the OFELIA, and later this functionality could be incorporated.

AR11 - After discussion with the evaluators, we got a consensus that in AR11, OFELIA has the mechanism support of SFA in the aggregate manager level, and SFA is an open interface. Therefore, it is fully supported.

AR12 - It is related to measurements and addressed in Section 7.2 (FIBRE-BR Monitoring Integration Approach).

AR14 - This requirement is about out-of-band network management, such that allocation of resources, routing and signalling can be done without affecting the rest of the requirements. Although this is supported by default in ProtoGENI, it was verified that depending how the islands are structured with separated control and management networks, this can be supported by all CMFs.

AR17 and AR18 - In the specific case of OFELIA, since there is the isolation mechanism of OFELIA - the FlowSpace, this could support the guarantees for data plan isolation from the point of view of security. Packets and events from one experiment are not seen by other. This is physical isolation on layer 2 mostly and other layers if needed, however the FlowSpace do not isolate the bandwidth of the experiments. For this network part, the idea is to set VLAN and / or virtual circuits per experiment that isolate the bandwidth. Such work is going to be addressed in the NOC operation plan of the FIBRE-BR testbed. On the other hand, in OMF whenever allocating the antennas or access points, usually, there is one use per antenna so the isolating is robust. Finally, from the point of view of virtual machines, it is not possible to guarantee this robust isolation, unless the scheduler of the CPU is real time and have precise guarantees of CPU and memory usage, thus this can be considered partially supported.

AR19 - Same reasoning as AR12.

AR20 - Same reasoning as UR09. Growing and shrinking resource usage in the federation in real time could lead to problems, thus, we consider that it is not a best practice for now.

AR21 - This requirement is too demanding and very difficult to realize without a major retrofit of all the CMFs. In the case of ProtoGENI, this is not possible to be done, since the allocation of resources and topological changes come from a ns-2 file, therefore, as an example, it is really challenging to change ProtoGENI to have this support. After discussion, we see the partial accomplishment of this requirement by changing the attitude of the experimenter, which is, the experimenter should stop the experiment, change it, and run it again, not change it in real time. Due to all these discussion, we change this requirement to SHOULD.

AR22 - Same reasoning as UC7.

AR23 - Same reasoning as AR12.

AR24 - As discussed previously, by using a multi-CMF approach within the FIBRE-BR system, we also leverage the support of heterogeneity in the FIBRE federated system, therefore it might be the case that certain types of resources are going to be allocated by certain CMFs.

AR26 - Same reasoning as AR24.

AR27 - Same reasoning as AR24, in the particular case of wireless devices, this could be done by OMF.

6.2.4 Req ES - Experiment Support

The "experiment support" requirements support the experiment running inside the infrastructure and share a common space of solutions depending on how the NOC and operations are going to be planned and deployed. Aspects such as, verification if the testbed is active, verification if the testbed is being used in a malicious way, among other, are issues.

Requirement	OFELIA	FOAM	OMF	ProtoGENI	I&M OFELIA	I&M OMF	I&M ProtoGENI
Req ES01							
Req ES02							
Req ES03							
Req ES04							
Req ES05							
Req ES06							
Req ES07							
Req ES08							
Req ES09							
Req ES10							

ES3 - This requirement is going to be fully supported according to a plan for the FIBRE-BR NOC. Possible suggestions of tools to be used are NAGIOS and CACTI, quite useful monitoring tools to detect the stability and operational status of the system.

ES5 - Same reasoning as ES3 and also by the policies that are going to be proposed in the federation discussion in WP4.

ES8 - From the point of view of the experimenter in a wireless network, this type of degradation could be accomplished by changing the antenna gain in the time of transmission. Other types of degradation could also be emulated such as in ProtoGENI, it could be loss probability or varied delay encodedan ns-2 file and deployed using the emulator. In this way, we can verify that from the point of view of the multitude of CMFs this is partially supported. In relation to specific equipments that add more degradation such as jamming generators, there is no current plan to acquire this equipment at this stage, but if it happens by other sources in the future, they could be added using conventional interfaces based on OMF.

ES9 - Same reasoning as ES3, in other words, it could be necessary to isolate the management plan using a VPN approach, and the whole system could be monitored in low level by the NOC FIBRE-BR.

ES10 - Same reasoning as UC8, from the point of view of the VT_Manager there has been planned the use of more than one type of VM template support such that these VMs could be available to experimenters with everything installed, for example, realistic traffic generators.

6.2.5 Req IM – Instrumentation and Measurements Requirements

These are requirements related to the instrumentation and measurement in the FIBRE-BR testbed and in particular, in the slices themselves and in certain cases, at the level of experiments, and the influence of the testbed in the experiments' results.

Requirement	OFELIA	FOAM	OMF	ProtoGENI	I&M OFELIA	I&M OMF	I&M ProtoGENI
Req IM01	[red]	[grey]	[blue]	[green]	[blue]	[grey]	[green]
Req IM02	[red]	[grey]	[yellow]	[yellow]	[blue]	[grey]	[green]
Req IM03	[red]	[grey]	[green]	[green]	[blue]	[yellow]	[green]
Req IM04	[grey]	[grey]	[green]	[grey]	[blue]	[grey]	[green]
Req IM05	[red]	[grey]	[yellow]	[grey]	[blue]	[yellow]	[green]
Req IM06	[grey]	[grey]	[green]	[green]	[blue]	[green]	[green]
Req IM07	[green]	[grey]	[green]	[green]	[green]	[yellow]	[green]
Req IM08	[grey]	[green]	[grey]	[grey]	[green]	[grey]	[grey]
Req IM09	[blue]	[grey]	[blue]	[yellow]	[yellow]	[red]	[blue]

IM1, IM2, IM3, IM5, IM9 - this has been a concern when evaluating OFELIA 0.1, but there are plans to add more monitoring features in newer version of the OFELIA CMF. We also address these requirements from the point of view of FIBRE-BR by proposing in Section 7.2 - FIBRE-BR

Monitoring Integration Approach. In terms of end-hosts, the monitoring has several tools ready such as the OML library and the INSTOOLS and PerfSonar on ProtoGENI.

6.2.6 Req IF – Infrastructure Requirements

These are requirements related to the infrastructure itself.

Requirement	OFELIA	FOAM	OMF	ProtoGENI	I&M OFELIA	I&M OMF	I&M ProtoGENI
Req IF01							
Req IF02							
Req IF03							
Req IF04							
Req IF05							
Req IF06							
Req IF07							
Req IF08							
Req IF09							
Req IF10							
Req IF11							
Req IF12							
Req IF13							
Req IF14							
Req IF15							
Req IF16							

IF6, IF11 - This requirement is fulfilled by the use of multiple CMF in the FIBRE-BR system, in our discussing, the view is that OMF will accommodate the last mile access to wireless networks resources, while the ProtoGENI will give access to a last mile emulated local network and finally OFELIA would provide last mile virtual machines to interconnect to OpenFlow resources also controlled by OFELIA. In summary, this requirement is covered in all aspects by the multiple CMF approach.

6.2.7 Req SE - Security Requirements

These types of requirements are closely related to the security and accountability of the system. Therefore, from the point of view of the CMFs this is usually mapped in the Clearing House abstraction. Thus, most of these requirements are related to the user rights and access control to the system.

Requirement	OFELIA	FOAM	OMF	ProtoGENI	I&M OFELIA	I&M OMF	I&M ProtoGENI
Req SE01							
Req SE02							
Req SE03							

Req SE04							
Req SE05							
Req SE06							
Req SE07							
Req SE08							
Req SE09							
Req SE10							
Req SE11							
Req SE12							
Req SE13							
Req SE14							
Req SE15							

SE6 - It is important that the resources be properly managed from the security point of view, thus, there is an obligation to have a minimal set of manageable resources per experiment or slice. In the case of OFELIA, whenever a project and further slices are created there is a mandatory need to put at least one VT_Manager and one Opt_in associated with it. Therefore, after discussion about this requirement, we consider it supported.

SE7 - In the OFELIA CMF there is the possibility of tight control of resource allocation since the requests of FlowSpaces are sent by email to the administrator and he can grant on a one-by-one basis manually. There is a plan to make this decision making of granting resources based on matching authorization and site policy in an automatic mode. Therefore, this requirement is partially supported.

SE9 - Same reasoning as UR10.

SE12, SE13, SE14 - These are related to federation management, from the point of view of establishing trust and policies per institution, verifying what can be done and what cannot, among federations FIBRE-BR and FIBRE-EU, thus this requirement will be better defined and the solution space proposed in WP4.

SE15 - This requirement focus on preventing the use of the infrastructure as a Distributed Denial of Service tool to attack points in the Internet. We discussed and verified that since the federated Clearing House will provide a unique access to an experimenter and this is done using secure connections like VPN, the system is secured in this VPN access. Otherwise, in the case of attack, a shutdown of the experiment will be done by the NOC and the experiment traffic will be blocked from the infrastructure to outside.

6.2.8 Req OP - Operations Requirements

The majority of evaluators of CMFs marked these as not applied, since from the point of view of CMFs this not something expected from them, there are requirements related to a NOC and the regular operation of the infrastructure in a regular basis. Moreover, we are planning to discuss and propose a NOC FIBRE-BR Plan at next stages of the FIBRE-BR project.

Requirement	OFELIA	FOAM	OMF	ProtoGENI	I&M OFELIA	I&M OMF	I&M ProtoGENI
Req OP01							
Req OP02							
Req OP03							
Req OP04							
Req OP05							
Req OP06							
Req OP07							
Req OP08							
Req OP09							
Req OP10							
Req OP11							
Req OP12							
Req OP13							
Req OP14							
Req OP15							
Req OP16							
Req OP17							
Req OP18							
Req OP19							
Req OP20							

OP6 up to OP18 - These are requirements related to the NOC Plan to be discussed in WP2 and presented in the next stage of FIBRE-BR project. It will cover all parts as web portal of the FIBRE, wiki, calendar, components availability verification, and everything else related to operations.

7 CMF Definition and FIBRE-BR Design

7.1 CMF Definition

As a result of a comprehensive study of the three CMFs described above, i.e. OFELIA, OMF and ProtoGENI. And despite the Description of Work (DoW) in which we have stated previously to take ONE Control Framework, we have decided to adopt all of the three in the FIBRE-BR facility. By considering the thorough evaluation, the multiple resources we have at hand, and the goal of federation to achieve the intercommunication among CMFs, we think that using multiple CMF on each island can allow them to operate independently as a single instance and the federation process can start in the construction of the FIBRE-BR instead of a later step.

In fact, these three CMFs will be deployed in the nine different islands according to the following table:

Tablea1 – FIBRE-BR CMFs to be deployed by island

	OFELIA	OMF	ProtoGENI
RNP	X		
UFF	X	X	
UFRJ		X	
CPqD	X	X	
UFSCAR	X	X	
USP	X	X	X
UFG	X	X	
UFPA	X	X	
UNIFACS	X	X	

As shown in the table, the ProtoGENI will only be deployed in the USP island since this institution is the sole to have an Emulab platform. The RNP will use OFELIA CMF to control the backbone resources required to interconnect the Brazilian islands. The UFRJ will deploy OMF only because this institution will be focused in wireless resources only and in addition in the BR/EU federation.

Another general comment is that, the inclusion of OCF will allow a deeper testing of the software by Brazilian partners as well as identify possibilities to improve it. Additionally, based on the Use-Cases and Federation Working Packages, there will be a need also to contribute to the extension of OCF capabilities and this can also be applied to the other control frameworks.

The deployment plan of these CMFs will be described in Section 7.3.

7.2 FIBRE-BR Monitoring Integration Approach

One of the FIBRE-BR goals is to provide Instrumentation and Measurement facilities in order to help experimenters, network administrators and researchers to define experiments and collect infrastructure and/or experiment specific data.

As already mentioned, the FIBRE-BR testbed will use three different control and monitoring frameworks in its nine islands: OFELIA, OMF and ProtoGENI.

Monitoring in OMF is achieved by the Orbit Measurement Library (OML) which adopts a client/server architecture. OML is composed by the OML-Service, the OML-Client library, and the OML-Collection-Server. The client library runs on the application code. Measurement data sent from clients may be temporarily stored locally on the OML-Service but are destined to the OML-Collection Server that stores them in a backend SQL database [White, 2010].

ProtoGENI uses conventional network monitoring tools to measure the testbed's infrastructure as a whole (nodes, routers, links, etc). ProtoGENI uses Instrumentation Tools (INSTOOLS) to gather data from traffic generated from the experiments [Duerig et al. 2012], which is similar to OML, and Leveraging and Abstracting Measurements with perfSONAR (LAMP) [Swany and Boyd 2011].

Our target is to provide, possibly, with a maximum reuse of the available CMFs, monitoring services over a new integrated and federated network structure. Some of the various aspects involved include the virtualized equipment, networks and monitored data; the collected data control access; and, finally, the multiple CMFs I&M data integration.

The followingsection introduces the FIBRE-BR I&M Architecture. It integrates the diverse I&M services, tools and facilities from the multiple CMFs, allowing FIBRE-BR users, possibly transparently to each specific CMF, to benefit from the corresponding infrastructure and experiment specific measurement data.

7.2.1 FIBRE-BR Instrumentation & Measurement (I&M) Architecture

The basic requirement for the FIBRE-BR Instrumentation and Measurement (I&M) Architecture (FIBRE-BR I&M) is the capability to configure, monitor, collect, and display both infrastructure and experiment specific data for distinct federated or individual CMF aggregates. Besides that, the architecture also includes or considers the set of requirements adopted by the FIBRE-BR experimental facility which includes substrate measurements, experiment measurements, privacy of measured data, link measurements, measurement data transmission and storage, component locations, time services and federation. It is important to notice that FIBRE I&M, fundamentally, provides some sort of interface, API or access to the existing CMFs.

The building blocks of the FIBRE-BR I&M architecture, shown on

Figure 6, leverage from current CMFs measurement capabilities. They create an integration facility that conveys the measured data to other FIBRE-BR I&M services and users through a common schema and protocol, according to its policy and experimental network requirements.

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

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

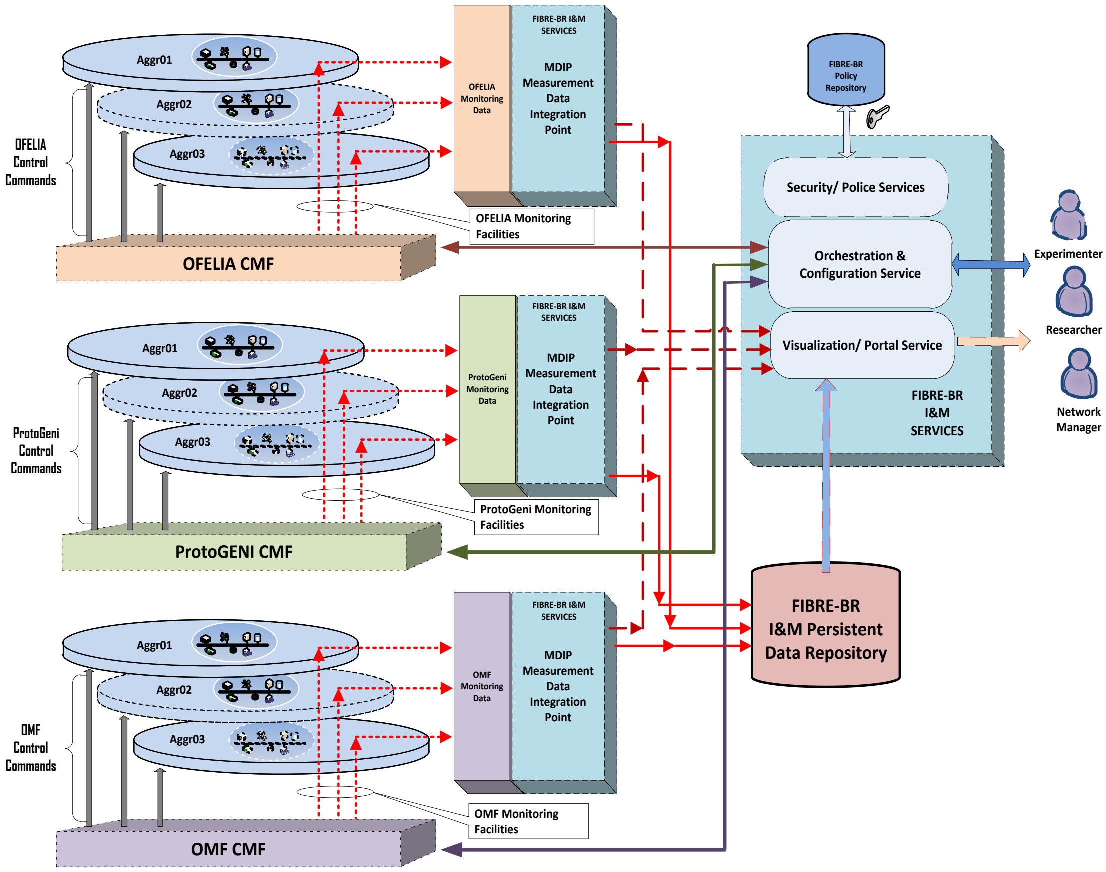

Figure 7: FIBRE-BR I&M Architecture Building Blocks

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

Finally, in terms of persistent data, the architecture has a storage strategy that allows users to retrieve data from their own or from others previous experiments, according to their access privileges. The persistent storage option is an experimenter decision that must comply with FIBRE-BR retention policy.

7.2.2 MDIP Implementation Aspects

MDIP interfaces with each CMF I&M tool that, in turn, has its own data format and storage mechanisms. MDIP deals mainly with data format and storage issues by, for instance: (1) using SOAP WebServices (XML-RPC) to make data available from one I&M tool to the portal; (2) possibly converting the collected data to a common data format among the different I&M tools; (3) using Measurement Archive (MA) services to export data stored, for example in RRD or SQL databases.

MDIP is strongly based on perfSONAR Measurement Archive service which is able to expose collected data using SNMP, SQL Database and measured data stored in Round Robin Databases (RRD) format [Hanemann 2005]. MDIP provides also datanormalization and/or conversion to Network Measurements – Working Group (NM-WG) standard, which is XML based, in orderto export the collected data to the FIBRE-BR I&M Visualization/Portal services and to store them in the I&M Persistent Data Repository.

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

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

Final Considerations

FIBRE-BR I&M Architecture is designed to configure and collect data from different CMFs. It is intended to be an instrumentation and measurement evolutionary architecture in the sense that, firstly, it evolves from integrating single aggregates belonging to common CMFs and, secondly, integrating federated aggregates from multiple CMFs. The proposed architecture aims at solving most of the associated challenges in the I&M scenario of Future Internet applications.

8 FIBRE-BR Initial Proposal Plan

In order to overcome all the issues and doubts about types, requirements and differences among CFs (Control Frameworks), it was considered a joint deployment process divided into four (4) phases described in the following Subsections.

8.1 Phase 0

This phase consists of installation and operation of CMFs at isolated islands. The partners should install and configure the CMFs software and dependencies in a FIBRE server at each island according with defined in Table a1.

Finally, the partners should run simple experiments at each island and use the mailing list to report problems, results or doubts.

8.2 Phase 1

In this phase the Island will be interconnected over the RNP backbone using static topologies, providing a path between islands configured in a static way. With the islands connected, experiments will be run using different islands to test the interconnections among them.

8.3 Phase 2

While the phase 1 is running, the requirements for the dynamic topology will be listed and discussed. After, the Island will be interconnected on a dynamic topology. The solution used in Fibre in this phase could be developed during the project or based in an existing solution.

8.4 Phase 3

The phase 3 will deploy the SFA-based federation within FIBRE-BR and with FIBRE-EU.All the island will be federated and it will be possible to run experiments using resources from Brazilian or together with European partners.

The details about these Phases will be provided by the Fibre Deployment Report in Task 2.2.

"*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