Project Acronym	**Fed4FIRE**
Project Title	**Federation for FIRE**
Instrument	**Large scale integrating project (IP)**
Call identifier	**FP7-ICT-2011-8**
Project number	**318389**
Project website	**www.fed4fire.eu**

D3.3 - Report on first cycle development

Work package	WP3
Tasks	All
Due date	31/03/2104
Submission date	DD/MM/YYYY
Deliverable lead	Timur Friedman (UPMC)
Version	1.0
Authors	Ciro Scognamiglio (UPMC) Georgios Androulidakis (NTUA) Carlos Bermudo (i2Cat) Aris Dadoukis (UTH) Donatos Stavropoulos (UTH) Wim Vandenberghe (iMinds) SuengYong Park (KOREN) Changwoo Kim (KOREN) Thierry Rakotoarivelo (NICTA)
Reviewers	Julien Lefeuvre (Inria) Peter Van Daele (iMinds)

Abstract	This document describes the work done in WP3 to support cycle 1 of Fed4FIRE.
Keywords	Infrastructure, community

Nature of the deliverable	R	Report	X
	P	Prototype	
	D	Demonstrator	
	O	Other	
Dissemination level	PU	Public	X
	PP	Restricted to other programme participants (including the Commission)	
	RE	Restricted to a group specified by the consortium (including the Commision)	
	CO	Confidential, only for members of the consortium (including the Commission)	

Disclaimer

The information, documentation and figures available in this deliverable, is written by the Fed4FIRE (Federation for FIRE) *– project consortium under EC co-financing contract FP7-ICT-318389 and does not necessarily reflect the views of the European Commission. The European Commission is not liable for any use that may be made of the information contained herein.*

Executive Summary

The Fed4FIRE architecture addresses the needs of the experimenter communities in three key areas: experiment lifecycle, measurement and monitoring, and trustworthiness. The architecture is evolving as time progresses and the requirements of the federation (including its sustainability) become better understood. The architecture is being implemented in a series of cycles. This document describes the work done to implement the features of cycle 1 in the architecture-focused test facilities PlanetLab Europe, w-iLab-t, VirtualWall, OFELIA, Nitos, Netmode, Norbit and Koren.

Ideally, all test facilities would have been able to implement all features of cycle 1 of the architecture. However the technical level of difficulty associated with implementing some features differed between test facilities, and therefore a prioritised subset of the ideal cycle 1 features became the target. Fed4FIRE is an evolutionary project and this process of prioritisation of development is ongoing, driven by the principle of offering benefits and value to experimenters. In some cases this has involved temporary work-arounds to ensure that critical functionality is available. The table below gives the status of each test facility at the end of cycle 1.

Architectural Features	PlanetLab Europe	W-iLab.t	VirtualWall
Deploy SFA Aggregate Manager	GENI AM v3 implemented	GENI AM v3 Implemented	GENI AM v3 Implemented
Extend GetVersion struct with testbed and tools info	Implemented	Implemented	Implemented
Adoption of OMF6	Implemented	Implemented	Implemented
Provide facility monitoring OML stream	Implemented	Implemented	Implemented
Provide infrastructure monitoring OML stream	Planned for cycle 2	Partially planned for cycle 2	Planned for cycle 2
Deploy OML for experiment measurement	Planned for future cycles	Planned for future cycles	Planned for future cycles
Provide an ID provider compliant with SFA and X.509 certificates to local users	Implemented	N/A	Implemented
Automatically download root certificates of Fed4FIRE partners from the certificate directory	Implemented	Implemented	Implemented

Architectural Features	OFELIA	NITOS	NETMODE
Deploy SFA Aggregate Manager	GENI AM v2 (v3 in cycle 2)	GENI AM v2 (v3 in cycle 2)	GENI AM v2 (v3 in cycle 2)
Extend GetVersion struct with testbed and tools info	Implemented	Implemented	Implemented
Adoption of OMF6	Planned for cycle 2	Implemented	Implemented
Provide facility monitoring OML stream	Implemented	Implemented	Implemented
Provide infrastructure monitoring OML stream	Planned for cycle 2	Planned for cycle 2	Planned for cycle 2
Deploy OML for experiment measurement	Planned for future cycles	Planned for future cycles	Planned for future cycles
Provide an ID provider compliant with SFA and X.509 certificates to local users	N/A	Implemented	Implemented
Automatically download root certificates of Fed4FIRE partners from the certificate directory	Implemented	Implemented	Implemented

Architectural Features	NORBIT	KOREN
Deploy SFA Aggregate Manager	Early version deployed, further implementation planned for cycle 2	GENI AM v2 (v3 in cycle 2)
Extend GetVersion struct with testbed and tools info	Planned for future cycles	Implemented
Adoption of OMF6	Implemented	Planned for Cycle 2
Provide facility monitoring OML stream	Implemented	Implemented
Provide infrastructure monitoring OML stream	Implemented	Planned for Cycle 2
Deploy OML for experiment measurement	Implemented	Planned for future cycles
Provide an ID provider compliant with SFA and X.509 certificates to local users	N/A	N/A
Automatically download root certificates of Fed4FIRE partners from the certificate directory	Planned for cycle 2	Implemented

Table 1 Status of Fed4FIRE WP3 facilities at M18

Acronyms and Abbreviations

AM	Aggregate Manager
API	Application Programming Interface
CA	Certificate Authority
EC	Experiment Controller
FLS	First Level Support
FRCP	Federated Resource Control Protocol
GENI	Global Environment for Network Innovations
NEPI	Network Experiment Programming Interface
OCCI	Open Cloud Computing Interface
OF	OpenFlow
OMF	Control and Management Framework
OML	OML Measurement Library
OMSP	OML Measurement Stream Protocol
QoS	Quality of Service
PDP	Policy Decision Point
RSpec	Resource Specification
SFA	Slice Federation Architecture
SNAA	Sensor Network Authentication and Authorization
VM	Virtual Machine
WSN	Wireless Sensor Network
XMPP	Extensible Messaging and Presence Protocol

Table of Contents

1 Introduction

This document describes the work done in workpackage 3 of the Fed4FIRE project to implement the features agreed for cycle 1 of the Fed4FIRE architecture in the test facilities PlanetLab Europe, w-iLab-t, VirtualWall, Ofelia, Nitos, Netmode, Norbit and Koren. The period of this report is up to month 18 of the project.

The architectural features are described in deliverables D2.1 [1], D5.1 [2], D6.1 [3] and D7.1 [5]. The features described on these documents represent the 'ideal' state of cycle 1 of Fed4FIRE. However it was recognised within the project that the technical level of difficulty associated with implementing some features differed between test facilities, and therefore a prioritised subset of the ideal cycle 1 features (described in milestone MS3.1 [6]) became the target. Fed4FIRE is an evolutionary project and the process of prioritisation of development is ongoing, driven by the principle of offering benefits and value to experimenters. In some cases this has involved temporary work-arounds to ensure that critical functionality is available.

The document is structured as follows:

- This introduction reviews briefly the functionality proposed in cycle 1 of the project
- Subsequent sections detail the progress in each test facility in WP3 for the three architectural areas of experiment control, experiment monitoring and trust and security.
- The progress towards cycle 2 goals is outlined for each test facility.

The Fed4FIRE architecture for cycle 1 of the project is described in deliverable D2.1 "First federation architecture". The architecture addresses the needs of the experimenter communities in three key areas: experiment lifecycle management, measurement and monitoring, and trustworthiness.

1.1 Experiment lifecycle management

Experiment lifecycle management covers the ability for experimenters to discover and provision resources, and to control running experiments across federated test facilities. This is implemented in the architecture by centrally provided features such as a portal, and by features implemented within each test facility, and is described in [2].

The decision was taken to adopt Slice Federation Architecture (SFA) as the standard for resource description and discovery, and the Federated Resource Control Protocol FRCP standards for experiment control. Consideration was given to the relative ease with which these protocols could be adopted by a facility, and hence a uniform level of acceptance was not specified in cycle 1 across all test facilities. For example many test facilities already used SFA and hence had little or no work to do to be compliant with the architecture. Others had considerable work to do to match the concepts of their existing resource management framework to SFA and implement changes, and hence were not expected to become SFA compliant in cycle 1.

Fed4FIRE has adopted Control and Management Framework (OMF) [8] for experiment control.

1.2 Measurement and monitoring

Measurement and monitoring has three purposes in Fed4FIRE:

- To collect data on the health of the underlying facilities so that they can be kept up and running. This is called **facility monitoring** in Fed4FIRE
- To collect performance statistics on the resources used by an experiment, which are of interest to the experimenter. This is called **infrastructure monitoring** in Fed4FIRE
- To collect experiment-defined data, also for use by the experimenter. This is called **experiment monitoring** in Fed4FIRE.

The approach to cycle 1 is described [3]; facility monitoring was the priority for cycle 1. The decision was taken to adopt OML as the standard for providing monitoring for all three types of information stream.

1.3 Trust and security

A well-defined trust and security framework is essential for a federated network to operate. The specifications for cycle 1 are described in [5]. In cycle 1 the focus was on a basic security framework for resource discovery and provisioning, following the decision to adopt SFA.

The focus on cycle 1 is user management and authentication. Users can obtain X.509 certificates from Fed4FIRE identity providers. The certificates can be presented to any test facility, which makes a decision as to whether it trusts the identity provider that signed the certificate. A test facility does this by reference to directory of root certificates from Fed4FIRE ID providers. The trust decisions are made at the test facility level.

2 PlanetLab Europe

2.1 Facility description

PlanetLab Europe is the European portion of the publicly available PlanetLab testbed and is a part of the OneLab experimental facility.

Established in 2002 PlanetLab is a global network of computers available as a testbed for computer networking and distributed systems research. As of December 2011, PlanetLab was composed of 1024 nodes at 530 sites worldwide. Each research project runs a "slice" that gives experimenters access to a virtual machine on each node attached to that slice. See the PLE statistics page for more details: http://onelab.eu/index.php/testbeds/onelab-testbeds/planetlab-europe/ple-statistics.html.

Accounts are available to persons affiliated with corporations and universities that host PlanetLab nodes. Those who join PlanetLab Europe have access to the entire system. They also participate in the initiatives built around PlanetLab in Europe.

PlanetLab members actively participate in developing tools for the greater good of the community, and as a result each user has a wide choice of tools to use in order to complete regular slice maintenance tasks.

There are a number of free, public services that have been deployed on PlanetLab, including CoDeeN, the Coral Content Distribution Network, and Open DHT.

PlanetLab Europe operates under the direction of Timur Friedman of UPMC Sorbonne Universités, working in collaboration with the Institut National de Recherche en Informatique et en Automatique (INRIA).

PlanetLab Europe and OneLab

PlanetLab Europe is a key testbed within the OneLab experimental facility for Future Internet technologies. OneLab is extending PlanetLab Europe into new environments, beyond the classic wired internet. OneLab is deepening PlanetLab Europe by incorporating new monitoring tools. OneLab is federating PlanetLab Europe, both with other PlanetLabs worldwide and with other types of testbeds.

The OneLab experimental facility is funded by several different research projects, including OpenLab and NOVI within the European Commission's FIRE Unit, along with national and international projects F-Lab, FIT and FIBRE.

2.2 Cycle 1: Overview

The vision for Fed4FIRE is to make the federation easy to use for experimenters. One of the ways of realizing this vision is to have uniform mechanisms for the processes involved in setting up and running experiments. Fed4FIRE cycle 1 dealt with three major process areas:

- Experiment lifecycle management: describing, discovering and provisioning resources
- Experiment measurement and monitoring: providing measurements at the facility, infrastructure and experiment levels.
- Trust and security: providing a federation-wide approach to authentication.

The following sections deal with progress in PlanetLab Europe in cycle 1 in each of these areas.

2.3 Cycle 1: Experiment workflow and lifecycle management

2.3.1 Resource Discovery, requirements, reservation and provisioning

In Fed4FIRE the mechanism chosen for implementing resource discovery, requirements, reservation and provisoning is the Slice Federation Architecture (SFA) standard (specifically GENI AM APi v3), together with ontology based resource specifications (Rspecs).

PlanetLab Europe has a native implementation of SFA that has been upgraded and enhanced during Cycle 1 development.

2.3.2 Experiment Control

In PlanetLab Europe the means of interacting with one's slivers is through ssh. In addition, a slice can optionally be created as "OMF-Friendly" in which case it is possible to control its related slivers through an OMF Experiment Controller. PLE OMF support has been upgraded to the latest version 6.0.

2.4 Cycle 1: Measurement and monitoring

2.4.1 Facility Monitoring

PlanetLab Europe uses a separate tool, but tailored for PlanetLab deployments, named MyOps, which allows implementing escalation policies. For example, if one site has a node down, first, messages are sent to the technical contact, then to the PI, and after a while the site loses some of its capabilities (fewer slices). All this workflow is entirely automated under MyOps, which also provides raw data on the status of nodes – see http://planetlab.eu/monitor/site.

Nagios is also used as an infrastructure monitoring and alerting solution. Nagios provides information on the status of the infrastructure and generic services like DNS, SSH or HTTP. Services specific to PlanetLab testbed are also monitored: SFA AM, Registry, OMF etc.

Both sources have been used to provide monitoring information to the Fed4Fire FLS service by the use of OML streams.

2.4.2 Infrastructure Monitoring

PlanetLab Europe uses TopHat to aggregate measurement sources such as TDMI and others. One difficulty encountered by the PlanetLab Europe operations team was that the potentially very useful data gathered by MyOps are not easily accessible through an API or other query-based techniques, and so MyOps does not lend itself to the creation of a gateway so as to be made available through TopHat. There clearly was a need for both: (1) aggregating data about nodes from various places (MyOps being one, CoMon slicestat being another one, and we found quite a few other sources of potentially high interest), and (2) providing a decent querying interface to these data. In a first step, we leveraged the internal "tags" mechanism right in the MyPLC DB to extend it with such external data. In a federated world, and in particular with respect to MySlice, it might make sense to design a separate tool for hosting this aggregated data.

Regarding monitoring of experimentation metrics, OML is considered to be a suitable candidate to deploy, and, thus, will be deployed in PlanetLab Europe. From a user point of view, TopHat will be used to query the OML database through a specific gateway, which has to be developed.

The metrics that are measured and therefore will also be supported in the Fed4FIRE experiment monitoring OML stream are:

- TopHat
  - Traceroute measurements between each pair of PlanetLab nodes

  - For each IP hop, we could provide more information (ASN, country, hostname, etc.).
- CoMon slicestat data (http://codeen.cs.princeton.edu/slicestat/)
  - Slice name
  - Slice context id
  - CPU consumption (%)
  - Physical memory consumption (%)
  - Physical memory consumption (in KB)
  - Virtual memory consumption (in KB)
  - Number of processes
  - Average sending bandwidth for last 1 min (in Kbps)
  - Average sending bandwidth for last 5 min (in Kbps)
  - Average sending bandwidth for last 15 min (in Kbps)
  - Average receiving bandwidth for last 1 min (in Kbps)
  - Average receiving bandwidth for last 5 min (in Kbps)
  - Average receiving bandwidth for last 15 min (in Kbps)
  - Local IP address of this node
  - Number of active processes - that is, processes using the CPU cycle at the moment

2.4.3 Experiment Monitoring

Slicestat data can be queried directly on a node in real time; additionally a service collects those informations in a central server with a time resolution of 15 minutes. The resulting data can be queried either directly (data is stored on a PostgreSQL database) or through a JSON based RESTful interface that can be used to easily retrieve statistics on usage and load averages.

Experimenters can use those services to monitor the trend (in terms of resource usage) of their running experiment.

A plugin for the Portal has been developed to easily display such information.

2.5 Cycle 1: Trust and security

2.5.1 Identity Provider

PlanetLab Europe supports X.509 certificate identity authentication and a certificate-based user identity management has been implemented and integrated.

2.5.2 Certificate Directory

PlanetLab Europe provides a self-signed root certificate to the certificate directory and also fetches the other certificates to trust them.

2.6 Progress towards cycle 2

2.6.1 Interconnectivity

PlanetLab Europe resources are already accessible through a public IPv4 address. IPv6 support is currently being added even though full support will depends on the network infrastructure hosting the resources. Resources that are part of those institutions already supporting IPv6 in their infrastructure are also accessible with an IPv6 address.

PlanetLab Europe added OpenFlow capabilities through a modified version of OpenVSwtich called sliver-ovs. Experimenters will be able to create an OpenFlow overlay network by specifying the links between PLE nodes. The possibility to interconnect PLE and other OpenFlow capable testbeds will be tested.

3 Virtual Wall and w-iLab.t

3.1 Facility description

The **Virtual Wall** is hosted at and operated by iMinds iLab.t. The hardware consists of 110x dual CPU nodes (2x 6 cores), 24 GB RAM, 1x harddisk 250GB, Intel Xeon E5645 (2.40GHz), interconnected with a Force 10 C300 switch (multiple gigabit interfaces per node). The hardware can be used as bare metal hardware (operating system running directly on the machine) or virtualized through openVZ containers or XEN virtualization. XEN Virtualization comes into two flavours: using shared nodes (non-exclusive) where VMs are running on physical hosts which are shared with other experimenters, or using physical nodes (exclusive) which are exclusively used by your experiment. You have full control on all XEN parameters as you have root access to the DOM0. Multiple operating systems are supported, e.g. Linux, FreeBSD, Windows 7, Fedora.

Network impairment (delay, packet loss, bandwidth limitation) is possible on links between nodes and is implemented with software impairment. Some of the nodes are connected to an OpenFlow switch to be able to do OpenFlow experiments in a combination of servers, software OpenFlow switches and real OpenFlow switches.

w-iLab.t is hosted at and operated by iMinds iLab.t. The wireless testbed is located in an unmanned utility room (size: 66m x 22.5m). There is almost no external radio interference. At this location, hardware is hosted at 60 spots. Every spot is equipped with:

- 1 embedded PC with 2 Wi-Fi a/b/g/n interfaces and 1 IEEE 802.15.1 (Bluetooth) interface
- a custom iMinds-Rmoni sensor node with an IEEE 802.15.4 interface
- an “environment emulator” board (enabling unique features of the testbed including the triggering of repeatable digital or analog I/O events at the sensor nodes, real-time monitoring of the power consumption, and battery capacity emulation.

There are two additional possibilities:

- A number of cognitive radio platforms (including USRPs) as well as specialized spectrum scanning engines are available at this location. This enables state-of-the art research in the field of cognitive radio and cognitive networking.
- 20 mobile robots

3.2 Cycle 1: Overview

iMinds has changed the management software of its Virtual Wall and w-iLab.t testbeds in such a way that the adaptations needed to support the Fed4FIRE architecture can be developed jointly for both testbeds. This optimization allows the more efficient use of available manpower devoted to the testbed-side developments in the context of Fed4FIRE. To be more concrete:

- The Virtual Wall used to rely on Emulab for its management, and did not have a full experiment control framework in place.

- w-iLab.t used to rely on OMF for both its management and for experiment control.
- Now both testbeds rely on Emulab for its management, and use the experiment control part of OMF for experiment control.

3.3 Cycle 1: Experiment workflow and lifecycle management

3.3.1 Resource Discovery, requirements, reservation and provisioning

Both testbeds have been exposed through SFA on a publicly reachable IP address (on IPv4 while the nodes themselves are reachable over the public internet via ssh on IPv6). Topology creation support has been added through the SFA API and SFA GetVersion call has been extended as required.

The Virtual Wall and w-iLab.t testbeds run the emulab software optionally combined with OMF for experiment control and OML for measurements. (NEPI will be supported if the integration is successful).

Emulab supports SFA GENI AMv2 and AMv3 APIs with GENI RSpecs v3.

3.3.2 Experiment Control

FRCP (OMF6) and support of the Fed4FIRE portal and jFed have been extended.

In both testbeds OMF6 experiment and resource controllers have been deployed on top of emulab (which will be used for discovery and provisioning).

3.4 Cycle 1: Measurement and monitoring

3.4.1 Facility Monitoring

Both testbeds provided local facility monitoring information to the FLS dashboard through an OML stream.

3.4.2 Infrastructure Monitoring

This has not yet been implemented in cycle 1, since this is only tackled by WP6 in cycle 2.

3.4.3 Experiment Monitoring

For both testbeds OML, for experimenter monitoring, has been deployed on top of the testbeds.

3.5 Cycle 1: Trust and security

3.5.1 Identity provider

iMinds has provided a Fed4FIRE-compliant (SFA and X.509) identity provider based on emulab, that allows to create Fed4FIRE identities for local users.

3.5.2 Certificate directory

The Virtual Wall and w-iLab.t testbed handed over their self-signed root certificate to the certificate directory and fetch the other testbeds certificates in order to authorize access to our testbeds to all Fed4FIRE users.

3.6 Progress towards cycle 2

3.6.1 Trust and Security

The main focus in cycle2 is to adopt the Policy Decision Point (PDP) of WP7 in order to make OMF6 secure in a federated context (see [7] for more details), to support the SLA framework and to support hard reservations.

3.6.2 Interconnection

iMinds has made all resources on both testbeds publically reachable through IPv6, which makes them reachable by experimenters on layer 3 over the plain Internet (because of shortage of IPv4 addresses, IPv6 is a necessity).

iMinds has initiated connectivity at layer 2 to Internet2 in the US. The goal is to have dynamic layer 2 connectivity (stitching of VLANs) to other testbeds and to the US. Within Europe, Geant and Autobahn are investigated for this, including the development of some proof-of-concept demonstrators.

3.6.3 SLA support and Infrastructure monitoring for federation services

In cycle 2 iMinds will be the first party to adopt support for SLA's. The SLA type that will be implemented by the Virtual Wall and w-iLab.t testbeds guarantees X uptime rate for Y rate of the resources during the sliver lifetime. For this it will deploy the SLA management module of WP7 at its testbeds. In order for this module to be able to function, iMinds will also support infrastructure monitoring for federation services as described in [4]. Infrastructure monitoring for experimenters will only be supported in cycle 3.

4 OFELIA

4.1 Facility description

OFELIA facility refers to **University of Bristol** (UNIVBRIS) and **i2CAT** OFELIA islands. UNIVBRIS and i2CAT OFELIA testbeds comprise OpenFlow-capable L2 switches and servers with virtual machines that act as traffic sources and sinks. UNIVBRIS also has OpenFlow-capable Reconfigurable Optical Add/Drop Multiplexers (ROADMs) but these resources will not be available for integration in cycle 1.

Access to the OFELIA testbed is free and open to any experimenter. In the testbed, experimenters can define and deploy virtual machines on any of the available servers of the interconnected islands. These virtual machines can be linked in a topology defined by the experimenters over the OpenFlow-enabled switches available. The experimenter's traffic over this topology is isolated from other experiments and users can define and set controllers, which dictated the behavior of the switches in their topology.

4.2 Cycle 1: Overview

OFELIA facility is composed of several islands (in Fed4FIRE cases, available islands are i2CAT and University of Bristol). Simplifying, each island exposes the managers of the different resources it can provide: Virtualization Aggregate Manager (VT AM) for the virtual machines resources and OptIn Aggregate Manager (OF AM) for the OpenFlow resources.

Main focus for OFELIA facility during cycle 1 was to provide the experimenter workflow and lifecycle management. To achieve this, the testbed must expose its resources through an SFA API. OFELIA facility has modified its aggregate managers to provide them with the SFA API by developing a custom wrapper for the AMs based on the SFAwrapper implementation. Additionally the AMs are in continuous development and enhancement will be added in future cycles.

OFELIA facility is also providing a basic monitoring of the infrastructure for Fed4FIRE making use of the existing monitoring tools of the OFELIA project.

4.3 Cycle 1: Experiment workflow and lifecycle management

4.3.1 Resource Discovery, requirements, reservation and provisioning

As commented in previous section, the mechanisms used in Fed4FIRE for management of resources are SFA and RSpecs. OFELIA does not provide this natively, so the Aggregate Managers for the resources where modified to include an SFA API based in GENI v2 with provision to update it to v3 in coming cycles. Both the virtualization and the OpenFlow AMs are able to expose and provision their resources through SFA GENI API v2 implemented as a wrapper based on Generic SFAwrapper tool.

4.3.2 Experiment Control

In OFELIA, once the resources are provisioned the VMs can be directly controlled by accessing them through ssh. The OF switches are controlled by providing them the address of a controller tool that modifies their flowtables according to the experimenter's needs.

It is being studied to include FRCP in the vms so they can be controlled remotely.

4.4 Cycle 1: Measurement and monitoring

4.4.1 Facility Monitoring

OFELIA testbeds use the Zenoss monitoring framework for facility and infrastructure monitoring. An OML script has been developed to retrieve the monitoring data stored in the backend of Zenoss and stream this data to the central Fed4FIRE OML collection server. Using this streamed OML data, the First Level Support website of Fed4FIRE is able to show the internal status of the OFELIA islands and when the internal status of the facility was last measured by Zenoss.

Testbed Name	Ping latency (ms)	GetVersion Status	Free Resources	Internal testbed monitoring status	Last check internal status
Ofelia (Bristol openflow)	11.83	ok	40	ok	2014-04-29 10:10:02+00
Ofelia (Bristol vtam)	11.13	ok	6	ok	2014-04-29 10:10:02+00
Ofelia (i2CAT openflow)	11.11	ok	43	ok	2014-04-29 10:10:02+00
Ofelia (i2CAT vtam)	11.22	ok	4	ok	2014-04-29 10:10:02+00

Figure 1: Display of the facility monitoring of the OFELIA islands on the First Level Support website

In addition, the SFA-enabled aggregate managers of the OFELIA islands can provide an indicative number of free resources that are available on the islands. For example, a SFA ListResources API call on one of the OpenFlow aggregate managers will return the number of free VLANS available for allocation to experimenters. This feature of the OFELIA aggregate managers is used by the First Level Support website to display the amount of free resources.

4.4.2 Infrastructure Monitoring

The OFELIA islands have extended the number of monitoring points in Zenoss monitoring framework. Currently, the following monitoring points are available for our compute resources:

1. Host Name (i.e. machine ID);
2. Net Bandwidth (available bandwidth in Mbps);
3. Net Used (% of bandwidth);
4. Net RX and errors (bytes received by NIC);
5. Net TX and erros (bytes sent through NIC);
6. Net Info (NIC identification, e.g. model, manufacture);
7. CPU Used (%);
8. CPU Info (CPU identification, e.g. model, manufacturer);
9. Memo Total Size (in KB);
10. Memory Used (% RAM);
11. Memo Swap Used (%);
12. Memo Info (Memory identification, e.g. model, manufacturer);
13. Disk Total Size (in KB);
14. Disk Used (%); and
15. Ports status (i.e. we are monitoring that our aggregate manager ports are up).

The OFELIA islands also monitor the OpenFlow packet switches of their networks with the following monitoring points being measured:

1. CPU used (%);
2. memory used (MB);
3. Net RX and errors in bytes and packet per interface;
4. Net TX and erros in bytes and packets per interface; and
5. Openflow stats via pox controller module (built in-house).

The monitoring information specified above will be provided to the central Fed4FIRE collection server and/or portal as required from WP6. Moreover, it would be possible to use this unprocessed monitoring information for further statistical analysis such as getting the peak link traffic over an extended period of time.

4.4.3 Experiment Monitoring

Experimenters using OFELIA testbeds need two types of experiment monitoring: OpenFlow network monitoring and VM monitoring. For packet OpenFlow network monitoring, it is common that experimenters use their OpenFlow controllers to monitor the OpenFlow rules that they install in their slice. This type of monitoring is useful because many OpenFlow applications, which sit on top of OpenFlow controllers, need to directly have access to this information to function properly. For example, traffic engineering applications is one category of such applications. OFELIA islands will support experimenters by providing pre-configured VM images with OpenFlow controllers which already have some network monitoring modules. For the VM monitoring, OFELIA islands are investigating the integration of OML as an option to experimenters by providing preconfigured VMs with OML already installed.

4.5 Cycle 1: Trust and security

4.5.1 Identity provider

The current identity system is based on the LDAP protocol, although AMs have been modified to accept Fed4FIRE certificates and allow Fed4FIRE users manage the OFELIA AMS. To allow local OFELIA users to gain access to the other testbeds within the Fed4FIRE federation, they should be given an appropriate identity using SFA X.509 certificates. This can be done in two ways: by providing our own identity provider (e.g. as part of the to be adopted AMSoil), or by utilizing the central Fed4FIRE identity provider to provide these credentials to the OFELIA local users. At the moment some further investigation is needed to identify the best strategy and, as OFELIA is not only composed by i2CAT and UnivBRIS islands, providing local OFELIA users the needed identity to authenticate themselves as valid Fed4FIRE users on the other testbeds of the federation is a feature that needs to be carefully studied.

4.5.2 Certificate directory

The root certificates of the other testbeds will be retrieved from the certificate directory in order to be able to authenticate valid Fed4FIRE users.

4.5.3 Rules based authorization

OCF has a Policy Engine, which uses the pyPElib library [13], which allows establishing a rule based authorization mechanism for the requests for virtual machines resources. At the moment WP7 has not yet made any decisions or recommendations regarding the preferred mechanisms for rules based authorization. Instead, WP7 is currently exploring the different technical possibilities to define and implement this aspect of the federation, intending to define a detailed strategy for development and deployment of rules based authorization in cycle 2. During this current exploration, the OFELIA testbeds will provide their experience previously gained with pyPElib.

4.6 Progress towards cycle 2

i2CAT and UnivBRIS will continue the significant efforts that they are investing to make the OFELIA islands meet all the requirements of Fed4FIRE. In line with this overall goal, OFELIA islands have identified some key objectives that they plan to achieve in cycle 2:

1. OFELIA islands are currently designing a new website which will provide extensive details on the different features of the islands and how to best use the Fed4FIRE tools, e.g. JFed and Fed4FIRE portal, with the islands. More important, this website will contain several tutorials to help experimenters get started with experimenting on the islands.
2. OFELIA islands are currently extending and testing the range of monitoring points for infrastructure and experiment monitoring possible through their Zenoss monitoring framework. The OFELIA islands are also extending the OML script for streaming the monitoring data from the Zenoss backend to the central Fed4FIRE collection server.
3. OFELIA islands are improving the existing MySlice plugins for the islands by extending the functionalities of the plugins as well as make it more intuitive for the users.
4. OFELIA islands are extending their L2 connectivity with other Fed4FIRE partners. For example, we are investigating the setup of a GEANT autobahn lightpath between BonFIRE in Edinburgh and OFELIA in Bristol.
5. OFELIA islands are developing the SFA GENI v3 API that will be available in addition to the current SFA GENI v2 API. This is being done by adapting the VT AM or replacing in OFELIA project the actual OF AM by a GENI developed one called FOAM [14] which already includes SFA GENI v3 API. In the second case, the availability of the AM with v3 API depends on when the AM is finally developed by the GENI team.
6. OFELIA islands is investigating the provision of VM images which are preconfigured with different software stacks such as OML and OpenFlow controllers with monitoring modules.
7. OFELIA islands are investigating the deployment of an experiment control framework, which will involve the deploying of a resource and/or experiment controller as an option in the experimenter's VM and an XMPP server.

5 NITOS

5.1 Facility description

The NITOS FI facility developed by the **University of Thessaly, Greece** (UTH), attracts numerous researchers from around the globe, thus rendering UTH a key player among the testbeds comprising the FIRE (Future Internet Research Initiative) initiative in Europe. The main experimental components of NITOS are:

- A wireless experimentation testbed, which consists of powerful nodes (some of them mobile), that feature multiple wireless interfaces and allow for experimentation with heterogeneous (WiFi, Bluetooth, ZigBee) wireless technologies. NITOS is about to be extended to a meso-scale testbed, by acquiring WiMAX and LTE Base Stations and by also enabling WiMAX/LTE connectivity to the wireless nodes.
- A software defined radio (SDR) testbed that consists of Universal Software Radio Peripheral (USRP) devices attached to the NITOS wireless nodes. USRPs allow the researcher to program a number of physical layer features (e.g. modulation), thereby enabling dedicated PHY layer or cross-layer research.
- A Software Defined Networking (SDN) testbed that consists of multiple OpenFlow technology enabled switches, connected to the NITOS nodes, thus enabling experimentation with switching and routing networking protocols. Experimentation using the OpenFlow technology can be combined with the wireless networking one, hence enabling the construction of more heterogeneous experimental scenarios.
- A testbed for conducting video-transmission (wired or wireless) related experimentation, which consists of high definition digital cameras, mounted on the NITOS nodes. This component can be combined with the wired (OpenFlow) and wireless testbeds mentioned above, enabling the study of video transmission over heterogeneous communication technologies.

NITOS testbed is open to the research community 24/7 and it is remotely accessible through the NITOS reservation tool. Parallel experimentation of different users is enabled, through the utilization of the NITOS scheduler software. The testbed is based on open-source software that allows the design and implementation of new algorithms, enabling new functionalities on the existing hardware. Through OMF, NITOS supports evaluation of protocols and applications under real world settings and through SFA enables federation with the other testbeds.

5.2 Cycle 1: Overview

During the first cycle of development there have been major updates in NITOS testbed. One of the most significant was the adoption of OMF6 along with the integration of NITOS Scheduler as a new entity in OMF, which is called “Broker”. Furthermore SFA AM has been updated from that of the Generic SFA Wrapper’s to an interface of “Broker”, capable of carrying all SFA commands.

5.3 Cycle 1: Experiment workflow and lifecycle management

5.3.1 Resource Discovery, requirements, reservation and provisioning

In the first cycle an instance of the generic SFA wrapper has been deployed in production state and its operation has been evaluated as robust and bug free. In addition to this wrapper NITOS exposes testbed resources through SFA by implementing an XML-RPC interface as part of Broker's communication layer, exposing this way the AM v2 API with RSpecs GENI v3 extended with reservation information. This local instance of Broker is deployed in development state and it is used for testing purposes.

For reservation purposes a Myslice plugin has been implemented, which interacts with the Broker through SFA. This plugin is implemented in a generic way, thus it can be easily adopted by all federated testbeds that provide exclusive resources or want to support in advance reservations of their resources.

5.3.2 Interact with the Future Reservation broker

As mentioned in 5.3.1 a local instance of the Broker has already been deployed in NITOS, which acts as a local reservation system. This local Broker will be contacted by the Reservation Broker, which will operate centrally at the federation level in cycle 2.

5.3.3 Experiment Control

For experimental control OMF6 is supported by NITOS through the provision of images that feature OMF6 Resource Controllers.

5.4 Cycle 1: Measurement and monitoring

5.4.1 Facility Monitoring

In order to monitor NITOS facility information are gathered and simple OML streams are used to populate a database with this information, which is then used by the FLS.

5.4.2 Infrastructure Monitoring

Infrastructure monitoring for Federation services will be implemented in the second cycle by gathering information from monitoring tools that are either already developed by NITOS or will be developed for that purpose. Infrastructure monitoring regarding the experimenter will be implemented in the third cycle of Fed4FIRE.

5.4.3 Experiment Monitoring

NITOS is using OML for collecting measurements from experiments. Experimenters can exploit the OML server that NITOS is providing for facilitating the procedure of measurement collection.

5.5 Cycle 1: Trust and security

5.5.1 Identity provider

An Identity provider with SFA X.509 certificates is supported. Currently, in NITOS we have generated a self-signed X.509 root certificate and every user of NITOS has a unique certificate signed by that root certificate. This root certificate has already been uploaded on F4F's Certificate directory.

5.5.2 Certificate directory

NITOS has implemented an automated procedure in order to upload its own certificate and fetching all the other testbeds' certificates from the F4F's Certificate directory. This enables the authentication of all valid F4F users.

5.6 Progress towards cycle 2

In cycle 2 there will be the following tasks for NITOS:

- Infrastructure Monitoring regarding Federation services using OML and monitoring tools, as described in 5.4.2.
- Broker deployment in production state with GENI AM v3 API and GENI RSpecs v3 extended with reservation information.
- Improvements in all the tools that are already deployed (MySlice plugin, OMF6, etc) with the goal of making them more robust.

6 NETMODE

6.1 Facility description

NETMODE Wireless Testbed consists of a control server and 20 wireless nodes. The testbed control server (Controller) runs a webserver with a GUI (custom GUI) and an SSH server that gives access to the wireless nodes. Inside the testbed, all nodes are connected via a switch. The testbed consists of two kinds of nodes:

18 Alix-based (Nodes 1-18)

- alix3d2 board
- 100Mbit Ethernet port
- 2 802.11 a/b/g interfaces
- 1GB flash card storage device

2 PC-based (Nodes 19-20)

- Intel Atom CPU
- 1Gbit Ethernet port
- 2 802.11 a/b/g/n interfaces
- 250 GB hard disk

Control and management functionalities of the Testbed are based on the OMF framework. NETMODE testbed is SFA-enabled by the deployment of the NITOS broker which acts as an SFA Aggregate Manager. The wireless nodes are scattered around the 3rd floor and the roof of the ECE building at the National Technical University of Athens.
For more information visit: http://www.netmode.ntua.gr/testbed/

6.2 Cycle 1: Overview

The vision for Fed4FIRE is to make the resources of the federated testbeds easily accessible for experimenters. One of the ways of realizing this vision is to have uniform mechanisms for the processes involved in setting up and running experiments. Fed4FIRE cycle 1 dealt with three major process areas:

- Experiment lifecycle management: describing, discovering and provisioning resources
- Experiment measurement and monitoring: providing measurements at the facility, infrastructure and experiment levels.
- Trust and security: providing a federation-wide approach to authentication.

The following sections describe the progress for NETMODE testbed in cycle 1 in each of these areas.

6.3 Cycle 1: Experiment workflow and lifecycle management

6.3.1 Resource Discovery, requirements, reservation and provisioning

NETMODE exposes testbed resources through SFA AM APIv2 with RSpecs v2 extended with reservation information (leases). This is achieved via the deployment of the NITOS broker which acts as an SFA enabled Aggregate Manager.

6.3.2 Experiment Control

NETMODE testbed has deployed OMF6 experiment and resource controllers. Therefore the experimenter can use OEDL (OMF Experiment Description Language) to describe and execute his experiments.

6.4 Cycle 1: Measurement and monitoring

6.4.1 Facility Monitoring

For facility monitoring, NETMODE testbed has deployed Nagios and Zabbix. A plugin for exporting facility monitoring information to a central OML repository has been developed and deployed.

6.4.2 Infrastructure Monitoring

Regarding Infrastructure monitoring, NETMODE testbed has deployed Nagios and custom scripts.

The metrics that are measured and therefore will also be supported in the Fed4FIRE infrastructure monitoring as OML streams (cycle 2) are the following:

Node Status

- Ping
- SSH access
- Wireless link quality

Other possible metrics such as CPU load, memory, etc. may be considered as well.

6.4.3 Experiment Monitoring

Experimenter measurements: OML is already deployed in baseline image for the wireless nodes to support the collection of measurements during the experiments.

6.5 Cycle 1: Trust and security

6.5.1 Identity provider

An Identity provider with support of SFA X.509 certificates has been deployed (SFA Registry) in NETMODE testbed in order to provide its local users with the capability to authenticate as valid Fed4FIRE users on the other testbeds of the federation.

6.5.2 Certificate directory

NETMODE testbed has uploaded its root certificate to the certificate directory and has fetched other testbeds' certificates to the trusted_root certificate directory of the AM.

6.6 Progress towards cycle 2

6.6.1 Interconnectivity

NETMODE resources (wireless nodes) at the moment have private IP addresses and are accessible through a gateway which as a public IPv4 address. IPv6 support is currently being added at the wireless nodes in order to be accessible from the public Internet.

7 KOREN

7.1 Facility description

KOREN (KOrea advanced REsearch Network) is a non-profit testbed network infrastructure established for facilitating research and development and international joint research cooperation. It provides quality broadband network testbed for domestic and international research activities to the industry, academia, and research institutions, enabling testing of future network technologies and supporting R&D on advanced applications. KOREN has multiple islands in Korea, spread across several cities. OpenFlow islands are located in three major Cities, Seoul, Pusan and Daejeon.

Each KOREN OpenFlow island comprises OpenFlow-capable L2 switches, servers with virtual machines and Juniper switch that are attached to WDM (Wave Division Multiplexing) switch. Juniper switch is a gateway to dynamic circuit network (also acknowledged as DCN, ION or Autobahn) that provides the on-demand WAN connectivity. One hundred VLANs are dedicated for dynamic circuit switching and it makes KOREN afford one hundred concurrent experiments over WAN.

7.2 Cycle 1: Overview

The goal of KOREN during cycle 1 is to transform the existing infrastructure federation-ready. Existing infrastructure is tightly integrated with the domestic KOREN management system, albeit it is following GENI APIs and SFA requirements. Major efforts have been focused on the re-architecting so that the infrastructure is independent of in-house management system. After the re-architecting, KOREN team made the necessary modification and installation of the software so that F4F management system better access and control KOREN infrastructure. One of the ways of making F4F management system better accesses KOREN infrastructure is to have uniform mechanisms for the processes involved in setting up and running experiments. For cycle 1 dealt with three major process areas:

- Experiment lifecycle management: describing, discovering and provisioning resources
- Experiment measurement and monitoring: providing measurements at the facility, infrastructure and experiment levels.
- Trust and security: providing a federation-wide approach to authentication.

The following sections deal with progress in KOREN in cycle in each of these areas.

7.3 Cycle 1: Experiment workflow and lifecycle management

7.3.1 Resource Discovery, requirements, reservation and provisioning

In the Fed4FIRE mechanism chosen for implementing resource discovery, requirements, reservation and provisioning is the Slice Federation Architecture (SFA) standard (specifically GENI AM API v3), together with ontology based resource specifications (Rspecs).

KOREN also has such mechanism, but it was tightly integrated with the domestic manager. KOREN team separated the dependency and made the infrastructure exposed and discoverable to any authorized management systems. After cycle 1, OpenFlow resources in KOREN are discoverable by certified any F4F management system.

7.3.2 Experiment Control

When a user experiments an OpenFlow test, one typically reserves OpenFlow switch(s) and associated VMs. As VMs would work as traffic source or traffic sync, they have to have physical connection to OpenFlow switch. FOAM, the OpenFlow AM, can make a reservation of OpenFlow switch. But it does not have any mechanism to assign VMs that are physically attached to the switch.

Previously, KOREN resolved the issue in a propriety manner. But after the Federation mechanism is in place, we may not be able to rely on the propriety mechanism anymore. KOREN team is working on solution that would enable a user to identify VMS that have direct physical connection to the reserved OpenFlow switch. With this solution, a user can make a proper reservation on VMs that are attached to the reserved OpenFlow switch. The solution would be available in cycle 2.

7.4 Cycle 1: Measurement and monitoring

7.4.1 Facility Monitoring

KOREN testbeds have been in use on the OML stream for facility monitoring about simple facility information. To enhance facility monitoring framework for experiment measurement, we has been planned to provide the facility monitoring data stored in the backend of OML server and inject the stream of data to the central Fed4FIRE OML collection server.

It's being customizing OML script for streamed specified OML data and we'll show the internal status of the KOREN islands on the First Level Support website dashboard of Fed4FIRE.

In addition, the SFA-enabled aggregate manager (AM) of the KOREN can provide an indicative number of free physical OpenFlow resources that are available on the island. For example, a ListResources API of SFA call on the OpenFlow AM will return the number of free VLANs available for allocation to experimenters. This feature of the KOREN AM is used by the First Level Support website dashboard to display the amount of free OpenFlow resources.

7.4.2 Infrastructure Monitoring

KOREN is currently upgrading its monitoring system to Zabbix. Previously, it was using PerfSonar. At this stage, Zabbix monitors CPU load and Utilization of the hosts on KOREN OpenFlow islands. This monitoring information will be provided to the central Fed4FIRE collection server as required from WP6.

As KOREN upgrade proceeds, Zabbix will also monitor OpenFlow switches and legacy switches as well as servers. Flow stats and other intrinsic OpenFlow stats are already available to OpenFlow Controller, but KOREN will also provide the stat information of the physical switch, such as packet count per port,

error rates on physical port, through Zabbix monitoring system. It is not decided if the other partners need this extened information. But if such demand arises, we provide it through OML.

7.4.3 Experiment Monitoring

To use KOREN testbeds for experiments, users typically need to monitor switch(s) and VMs. For intrinsic OpenFlow switch monitoring, users' own OpenFlow controllers monitor the flows set and other parameters. This type of monitoring is useful because many Openflow applications, which sit above core of OF controllers, need to have access to this monitoring information to function properly. Although KOREN already provides this type of monitoring, we consider we may need additional monitoring capability so that we can monitor OpenFlow switch itself and associated VMs. At the moment, we are not aware of the availability of such tools. And if necessary, we may come up with the solutions that would help the experimenters to monitor. This is one of the issues KOREN would pursue in future cycles.

7.5 Cycle 1: Trust and security

7.5.1 Identity provider

The identity system of KOREN has provided the X.509 certification identity authentication and a certificate-based user identity. To allow local users to gain access to the other testbeds within the Fed4FIRE federation, they should be given an appropriate identity using SFA X.509 certificates. For cycle 1, KOREN relies for this on the central Fed4FIRE identity provider.

7.5.2 Certificate directory

The root certificates of the other testbeds will be retrieved from the certificate directory in order to be able to authenticate valid Fed4FIRE users.

7.6 Progress towards cycle 2

KOREN will dedicate us to do best effort way to produce actual results that make to meet all the requirements of Fed4FIRE. The overall vision of KOREN is to achieve the plan that we expect in the cycle 2:

1. KOREN is being prepared to enhance the GENI AM API v3 that will be available, including the current supported GENI AM API v2. It's also considering availability of AM with AM API v3 when GENI team support GENI AM API v3 for AM.

2. KOREN islands are currently extending and testing the monitoring measurement points for facility and infrastructure monitoring in accordance with our testbed through OML stream and Zabbix. It's also extending the OML script for streaming the monitoring data from backend server to central Fed4FIRE collection server.

3. KOREN will provide the solution that helps experimenters to identify VMs that have the direct connection to the reserved OpenFlow switch.

4. KOREN is preparing to provide the L2 connectivity to the Fed4Fire partners through L2 over L3 tunneling.

5. KOREN islands will provide OpenFlow switches, virtual machines and on-demand WAN connectivity with DCN. KOREN will also deploy OMF6 on its testbeds. OMF6 with support for OpenFlow resources should facilitate the optional integration of OMF for experiment control.

8 NORBIT

8.1 Facility description

The NORBIT (NICTA-ORBIT) facility is a testbed developed by NICTA, Sydney, Australia. It consists primarily of 40 nodes deployed across 3 office floor levels at NICTA's research laboratory in Sydney. Each of these nodes have the following capability:

- VIA Esther processor 1000MHz
- 1Gb RAM
- 2 x 1GB Ethernet interfaces
- 2 x 802.11 a/b/g/n interfaces (with 2 antennas per interface)
- additionnal USB ports allow local experimenters to add capabilities (e.g. Wi-Spy dongles)
- Chassis Manager Controller card, allowing remote switching on/off and resetting of nodes

For each node, one of the Ethernet interface is used for experiment control traffic and measurement collection, while the other interface is free for the experimenter to use. This last 'experimental' interface is connected to an OpenFlow switch for each node of a given office floor. The 3 OpenFlow switches (one per floor) are connected together. Experimenters can have full use of the nodes and the OpenFlow switches for their experiments.

The NORBIT testbed provides the OMF6 framework for experimenters to design and orchestrate the execution of their experiments. In addition, it also provides 2 OML servers to collect any measurements from their experiments. One of these servers is using a SQLite3 backend, and the other a PostgreSQL backend.

Standard OS images are provided to experimenters as baseline OS to load on the nodes for their experiments (e.g. Ubuntu, Debian, Fedora). However, experimenters can also create, save, and load their custom OS images.

8.2 Cycle 1: Overview

As mentioned previously Fed4FIRE cycle 1 deals with three major process areas:

- Experiment lifecycle management: describing, discovering and provisioning resources
- Experiment measurement and monitoring: providing measurements at the facility, infrastructure and experiment levels.
- Trust and security: providing a federation-wide approach to authentication.

The following sections describe the progress in addressing these points for the NORBIT testbed.

8.3 Cycle 1: Experiment workflow and lifecycle management

8.3.1 Resource Discovery, requirements, reservation and provisioning

Fed4FIRE selected the Slice Federation Architecture (SFA) standard (GENI AM APi v3) as its de-facto mechanism to discover, reserve, and provision resources using specifications captured in Rspec documents.

The NORBIT development team has been working on an OMF-SFA service, which offers an SFA API to allow the discovery, reservation and provisioning of OMF-controlled resources. An early version of this service has been testbed at NORBIT, and the full code is available at:

https://github.com/mytestbed/omf_sfa

8.3.2 Experiment Control

As previously mentioned, NORBIT provides by default a full OMF v6 deployment, which may be used by experimenters to design and control the orchestration of their experiments. The NORBIT team has large number of experiment descriptions available as reference for experimenters. The use of OMF v6 also allows experimenters to use resources from different testbeds (e.g. NORBIT, PlanetLab Europe) within the same experiments. This has been demonstrated at different venues (e.g. Fed4FIRE first review).

8.4 Cycle 1: Measurement and monitoring

8.4.1 Facility Monitoring

NORBIT has a simple background running process to monitor the availability of each node. In addition, the server running many of the provided services (e.g. the OML server) is monitored using Nagios. These monitoring sources provide OML streams, which are further collected to allow operators to monitor the testbed's status, and to also be used in Fed4FIRE FLS.

8.4.2 Infrastructure Monitoring

NORBIT has an automated weekly experiment that monitors the status of the wireless connectivity between all nodes of the tested across the different floors. This process collects different metrics (e.g. RSSI, SNR, frame retransmission, etc…) as OML streams. These streams are forwarded to an OML server, which records them in a database. These infrastructure measurements are then available for experimenters who may want to use them to plan their own experiments. For example, it can be used to select a set of nodes across the testbed which would ensure 3-hops paths when using 802.11 a.

In addition the OMF v6 tools (both the Experiment Control and the individual Resource Controller on each node) are also instrumented using OML. Therefore, an experimenter can measure the performance of the tools that she is using to run experiments and take that into account when interpreting her results (e.g. how many control messages were sent between the EC and the RC, how long did it take for a control message to be processed, etc…)

8.4.3 Experiment Monitoring

NORBIT provides many OML-instrumented applications, which can be used by experimenters within their experiments. Examples of such OML-enabled applications are: iperf, collectd, gpsd, ping, WattsUp? Power Monitor, WPA Supplicant monitor.

In addition, NORBIT also maintain a list of OML-instrumented applications, which were contributed by third parties such as other experimenters in other universities. Examples of such contributed applications are: VLC, btclient, tinyhttpd, wget.

The complete list of OML-instrumented applications available to NORBIT experimenters is located at:

http://mytestbed.net/projects/omlapp/wiki/OML-instrumented_Applications

8.5 Cycle 1: Trust and security

8.5.1 Identity Provider

While OMF v6 does support authentication, NORBIT currently does not deploy remote identity authentication management.

8.5.2 Certificate Directory

The root certificates of the other testbeds will be retrieved from the certificate directory in order to be able to authenticate valid Fed4FIRE users.

8.6 Progress towards cycle 2

The NORBIT team has been involved in the design of the authorisation scheme based on chained assertions and the related Policy Decision Point entity developed in WP7. The outcome of that task will be deployed on the NORBIT testbed as part of cycle 2.

9 Conclusions

WP3 testbeds have made good progress implementing the feature of cycle 1 of the Fed4FIRE architecture. These developments will allow experimenters from the first open call to make use of these test facilities to carry out their experiments. The table below shows the progress to date. The user experience will be enriched as the architecture matures and further features are implemented.

Architectural Features	PlanetLab Europe	W-iLab.t	VirtualWall
Deploy SFA Aggregate Manager	GENI AM v3 implemented	GENI AM v3 Implemented	GENI AM v3 Implemented
Extend GetVersion struct with testbed and tools info	Implemented	Implemented	Implemented
Adoption of OMF6	Implemented	Implemented	Implemented
Provide facility monitoring OML stream	Implemented	Implemented	Implemented
Provide infrastructure monitoring OML stream	Planned for cycle 2	Planned for cycle 2	Planned for cycle 2
Deploy OML for experiment measurement	Planned for future cycles	Planned for future cycles	Planned for future cycles
Provide an ID provider compliant with SFA and X.509 certificates to local users	Implemented	N/A	Implemented
Automatically download root certificates of Fed4FIRE partners from the certificate directory	Implemented	Implemented	Implemented

Architectural Features	OFELIA	NITOS	NETMODE
Deploy SFA Aggregate Manager	GENI AM v2 (v3 in cycle 2)	GENI AM v2 (v3 in cycle 2)	GENI AM v2 (v3 in cycle 2)
Extend GetVersion struct with testbed and tools info	Implemented	Implemented	Implemented
Adoption of OMF6	Planned for cycle 2	Implemented	Implemented
Provide facility monitoring OML stream	Implemented	Implemented	Implemented
Provide infrastructure monitoring OML stream	Planned for cycle 2	Planned for cycle 2	Planned for cycle 2
Deploy OML for experiment measurement	Planned for future cycles	Planned for future cycles	Planned for future cycles

Provide an ID provider compliant with SFA and X.509 certificates to local users	N/A	Implemented	Implemented
Automatically download root certificates of Fed4FIRE partners from the certificate directory	Implemented	Implemented	Implemented

Architectural Features	**NORBIT**	**KOREN**
Deploy SFA Aggregate Manager	Early version deployed, further implementation planned for cycle 2	GENI AM v3
Extend GetVersion struct with testbed and tools info	Planned for future cycles	Implemented
Adoption of OMF6	Implemented	Planned for Cycle 2
Provide facility monitoring OML stream	Implemented	Implemented
Provide infrastructure monitoring OML stream	Implemented	Planned for Cycle 2
Deploy OML for experiment measurement	Implemented	Planned for future cycles
Provide an ID provider compliant with SFA and X.509 certificates to local users	N/A	N/A
Automatically download root certificates of Fed4FIRE partners from the certificate directory	Planned for cycle 2	Implemented

References

[1] Fed4FIRE Architecture Workpackage D2.1 “First federation architecture”
[2] Fed4FIRE Experiment Lifecycle Management Workpackage D5.1 “Detailed specifications for first cycle ready”
[3] Fed4FIRE Measuring and Monitoring Workpackage D6.1 “Detailed specifications for first cycle ready”
[4] Fed4FIRE Measuring and Monitoring Workpackage D6.2 “Detailed specifications regarding monitoring and measurement for second cycle”
[5] Fed4FIRE Trustworthiness Workpackage D7.1 “Detailed specifications for first cycle ready”
[6] Fed4FIRE Service and Applications WorkPackage “MS4.1 First design specifications for facilities”
[7] Fed4FIRE Trustworthiness Workpackage D7.2 “Detailed specifications for second cycle ready”
[8] Control and Management Framework http://mytestbed.net/projects/omf/
[9] Open Cloud Computing Interface http://www.occi-wg.org
[10]SFAWrap, http://sfawrap.info
[11] GENI Aggregate Manager API Version 2, http://groups.geni.net/geni/wiki/GAPI_AM_API_V2
[12] FI-Lab, The Open Innovation Lab, http://lab.fi-ware.org
[13] pyPElib python Policy Engine library, https://github.com/fp7-ofelia/pypelib
[14] FOAM, GENI OpenFlow Aggregate Manager, http://groups.geni.net/geni/wiki/OpenFlow/FOAM