Grant Agreement No.: 604590
Instrument: Large scale integrating project (IP)
Call Identifier: FP7-2012-ICT-FI

eXperimental Infrastructures for the Future Internet

D5.6: XIFI Federation Completion and Support

Revision: v1.0

Work package	WP 5
Task	Task 5.5
Due date	31/03/2015
Submission date	30/04/2015
Deliverable lead	Telecom Italia S.p.a (TI)
Authors	Daniele Giai Pron (TI), Marco Cipriani (TI), Thierry Milin (Orange), Joe Tynan (WIT), Riwal Kerherve (ILB), Ioannis Igoumenos (UTH), Rudolf Vohnout (CESNET), Jan Kundrát (CESNET), Sándor Laki (WIGNER), József Stéger (WIGNER), Joaquin Iranzo (ATOS), Cristian Cristelotti (TN), Genci Tallabaci (TN), Fernando López (TID), Wojbor Bogacki (PSNC), Jacek Kochan (PSNC), Vicent Borja Torres (iMinds), Thijs Walcarius (iMinds), Panagiotis Demestichas (UPRC), Vera-Alexandra Stavroulaki (UPRC), Angelos Rouskas (UPRC), Aristi Galani (UPRC), Demetrios Kelaidonis (UPRC), Marios Logothetis (UPRC), Konstantinos Tsagkaris (UPRC), Aimilia Bantouna (UPRC), Panagiotis Vlacheas (UPRC), Georgios Poulios (UPRC), Vassileios Foteinos (UPRC), Roland Elverljung (Acreo Swedish ICT), Stéphane Junique (Acreo Swedish ICT), Seán Murphy (ZHAW)
Reviewers	Federico Michele Facca (CNET), Laura Pucci (ENG)

Abstract	This document reports on the federation extension activities, describing support mechanisms associated, inclusion of a new infrastructures, issues, encountered, technical and operational feedbacks.
Keywords	New nodes, Deployment plan, federation extension, Validation Tests, Lessons Learned, Node Status

Document Revision History

Version	Date	Description of change	List of contributor(s)
V1.0	29.04.2015	Final and reviewed document version	Marco Cipriani (TI), Daniele Giai Pron (TI) et al.

Disclaimer

Copyright notice

Project co-funded by the European Commission in the 7th Framework Programme (2007-2013)		
Nature of the Deliverable:	P (Prototype)	
Dissemination Level		
PU	Public	✓
PP	Restricted to other programme participants (including the Commission Services)	
RE	Restricted to bodies determined by the XIFI project	
CO	Confidential to XIFI project and Commission Services	

[1] http://creativecommons.org/licenses/by-nc-nd/3.0/deed.en_US

EXECUTIVE SUMMARY

This document describes the federation extension activities run by the XIFI project to increase the capacity of the FIWARE Lab. In particular the document, describes: the support mechanisms associated to the creation of the federation; the activities of inclusion of new infrastructures and issues encountered in running the federation process; technical and operational feedbacks derived from the lessons learnt by applying the federation process. This document reports the activities performed to set up the federation of the XIFI nodes from Month Nine (M9)[6] to Month Twenty Four (M24) of XIFI project.

During this timeframe, there, has been the first Open Call in order to choose new infrastructures to join the federation and expand the FIWARE Lab.

As a result of the Open Call, additional eleven[2] nodes were selected beyond the five initial ones, increasing the FIWARE Lab capacities to more than 2500 cores from the initial 1000, and from 5 European regions to 13 European regions (some including more than one node per region, like France) to offer the needed capacity to support FIWARE Accelerators.

The Open Call selected the new partners according to:

- Technical capacity and scientific excellence
- Fulfilment of nodes requirements [1]
- Level of Commitment to FI-PPP program [2]
- Coverage of the 5 Large Trial Projects of Phase 2 of the FI-PPP Program
- Coverage of European countries not covered by the five initial nodes
- Commitment to contribute to the sustainability of XIFI federation after the project's end.

Beyond the federation into FIWARE Lab of those nodes selected through the Open Call, four infrastructures – on voluntary bases – asked to become part of the federation through the association process. Currently one of them is successfully federated and the others are in the process of completing the federation.

The federation process, as reported in the document, took longer than expected and also evidenced a number of technical challenges that were not foreseen at the start of the project. This document discusses lessons learnt and suggests improvements to the current deployment of the FIWARE Lab and procedures for the management of the FIWARE Lab. The main challenge that is under resolution, is the overbooking of resources and lack of control over resources. The resolution of this challenge is currently on-going through the enactment of resource access policies for users. It is foreseen that the application of such policies will drastically mitigate the issues related to overbooking of resources, guaranteeing a more fair and transparent usage of resources by developers. Other relevant issues are as well discussed in this document with related ongoing actions and recommendations for the upcoming period, not only to XIFI nodes, but as well as to FI-Core project.

[2] Due to financial issues of GOWEX, one of the partners providing new infrastructures to the federation, the nodes taking part to the process have been eleven rather than the twelve initially selected.

TABLE OF CONTENTS

LIST OF FIGURES

LIST OF TABLES

ABBREVIATIONS

AES	Advanced Encryption Standard
BGP	Border Gateway Protocol
BW	BandWidth
CMS	Content Management System
CNAME	Canonical Name
CoC	Carrier of Carriers
DCRM	Data Center Resource Management
DMZ	DeMilitarized Zone
DNS	Domain Name Service
DoW	Description of Work
EBM	Exploitation and Business Modeling (Working Group)
EC (EU)	European Commission (European Union)
EoMPLS	Ethernet over MPLS
ESP	Encapsulating Security Payload
FI	Future Internet
FI-PPP	Future Internet Public-Private-Partnership
FIRE	Future Internet Research and Experimentation
FP7	Framework Programme 7
FQDN	Fully qualified domain name
GE	Generic Enabler
GRE	Generic Routing Encapsulation
IaaS	Infrastructure as a Service
IDM	Identity Management
IKE	Internet Key Exchange
ILO	Integrated Lights-Out
IoS	Internet of Services
IPMI	Intelligent Platform Management Interface
IPR	Intellectual Property Rights
IPSEC	Internet Protocol Security
KVM	Kernel-based Virtual Machine
L2TP	Layer 2 Tunneling Protocol
L2-VPN	Layer 2 VPN
L3-VPN	Layer 3 VPN
LACP	Link Aggregation control protocol
MAC	Media Access Control address

MD-VPN	Multi Domain Virtual Private Network
MPLS	MultiProtocol Label Switching
NAT	Network Address Translation
NFS	Networking File system
NREN	National Research and Education Network
NRPE	Nagios Remote Plugin Executor
NTP	Network Time Protocol
OAM	Operations, Administration and Management
OOB	Out Of Band
OSI	Open Systems Interconnection
P2P	Point to Point
PaaS	Platform as a Service
PDU	Power Distribution Unit
PE	Provide Edge router
PoC	Proof of Concept
QoS	Quality of Service
RFC	Request For Comments
SaaS	Software as a Service
SDC	Software Deployment and Configuration
SFTP	Secure File Transfer Protocol
SHA	Secure Hash Algorithm
SLA	Service Level Agreement
SSH	Secure SHell
UC	Use Case
VM	Virtual Machine
VPLS	Virtual Private LAN Services
VPN	Virtual Private Network
VRF	Virtual Routing & Forwarding
WP	Work Package
XIFI	eXperimental Infrastructures for the Future Internet….

1 INTRODUCTION

1.1 Scope

The process of federation of new nodes is a core activity of XIFI in support of FIWARE and the FI-PPP programme; the introduction of new nodes must be managed with sustainable and efficient technical solutions in order to minimize the phases of definition, deployment and validation of a new node with the goal to have the resources of the new node available for the users of the FIWARE Lab (http://lab.fiware.org).
As described in [1], XIFI notion of federation aims at providing a single starting point as marketplace and assuring interchangeable and interoperable environment for all types of stakeholders of the federation throughout Europe. This document describes how a new node is introduced into the XIFI federation (federation process model), which are the technical steps to achieve the federation of a node (deployment process) and which are the validation tests used to assure that a node can be operational in the federation and in FIWARE Lab (validation process).

This document concludes with an analysis of the lessons learned during the federation activities with the goal to improve the processes for the federation activities that will continue in FI-Core or beyond.

Currently the (administrative) federation behind FIWARE Lab comprises 20 nodes and as such it can cope with large trial deployments and can serve the various needs of a broad set of FI users and experimenters. The interconnection of the infrastructures is supported by the national research networks (NRENs) and the pan-European GÉANT facility.

The nodes belonging to XIFI infrastructures federation can be found in the following table, for Overall Status meaning please refer to paragraph 4.2

Node	Country	Organization	Node category	Overall Status
Gent	Belgium	iMinds VZW	Full Member - Open Call	Operational
Prague	Czech Republic	Czech Education and Scientific NETwork (CESNET)	Full Member - Open Call	Up and Running
Lannion	France	Association Images & Réseaux (ILB)	Full Member – Initial consortium	Up and Running
Sophia Antipolis	France	Association Plate-Forme Télécom (Com4Innov)	Full Member - Open Call	Not Operational
Berlin	Germany	Deutsche Telekom / Fraunhofer FOKUS	Full Member – Initial consortium	Up and Running
PiraeusN	Greece	NeuroPublic A.E. PLIROFORIKIS & EPIKOINONION	Full Member - Open Call	Up and Running
PiraeusU	Greece	University of Piraeus Research Center (UPRC)	Full Member - Open Call	Up and Running

Node	Country	Organization	Node category	Overall Status
Volos	Greece	University of Thessaly (UTH)	Full Member - Open Call	Up and Running
Budapest	Hungary	Wigner Research Centre for Physics	Full Member - Open Call	Operational
Waterford	Ireland	Waterford Institute of Technology	Full Member – Initial consortium	Operational
Trento	Italy	Trentino Network SrL	Full Member – Initial consortium	Operational
Poznan	Poland	Poznan Supercomputing and Networking Center (PSNC)	Full Member - Open Call	Operational
Spain	Spain	Entidad Pública Empresarial Red.es/RedIRIS	Full Member – Initial consortium	Operational
Karlskrona	Sweden	Blekinge Institute of Technology (BTH)	Full Member - Open Call	Operational
Stockholm	Sweden	ACREO Swedish ICT AB	Full Member - Open Call	Up and Running
Zurich	Switzerland	Zurich University of Applied Sciences (ZHAW)	Full Member - Open Call	Operational
Crete	Greece	Technical University of Crete - Associate Node	Associate Member	Up and Running
Messina	Italy	Università di Messina - Associate Node – Not yet federated	Associate Member	Ongoing
Mexico	Mexico	Infotec - Associate Node	Associate Member	Operational
Wroclaw	Poland	Wroclaw University of Technology – Associated node – Not yet federated	Associate Member	Ongoing

Table 1: Nodes belonging to the XIFI federation

On the 6th May Gowex was contacted by email and was provided with technical XIFI documentation and references. Due to bankruptcy, this new partner was not active during the project and was not federated and is not referenced in this deliverable.

TID resources available to support this node were reallocated to support and improve in particular the sanity check tool (section 3.3)

1.2 Document conventions

The formatting of the document is compliant with the deliverable template provided by the XIFI project. No other specific convention has been applied.

1.3 Intended audience

The target audience of this deliverable is:

- The Infrastructures owners of the federation of XIFI, the federation includes the original five nodes and the twelve new Infrastructures coming from the first XIFI open call and the associated nodes
- The new infrastructures that will join the federation in the future
- Experts and technical personnel providing deployment support and end-user support activities. These activities will be fulfilled by the support entity of the XIFI federation
- Stakeholders of FI-Core project

1.4 Reading suggestions

The document is divided into seven chapters.

Chapter 1: Introduction

- Provides the information required to understand the rest of the document

Chapter 2: Federation Process

- Describe the process used inside XIFI project for supporting the federation extension activity

Chapter 3: Validation Process

- Describe the Validation Process used in XIFI in order to define a node as fully operational and to assure that the nodes is "up and running" in terms of the functionalities for the users

Chapter 4: Federation Status

- Describe the overall and the specific operational status of all the nodes and the history of the technical activities of the federation of the new nodes

Chapter 5: Lessons Learned

- Includes the FAQs used for the deployment of the new nodes as a tool of sharing knowledge and describe the lessons learned during the XIFI project based on a survey among all the nodes and the main project stakeholders

Chapter 6: Plan and Recommendations for the future

- This section include the on-going actions to mitigate and resolve issues in the FIWARE Lab and recommendations to FI-Core project for further improvements of the technology behind the Lab.

Chapter 7: Conclusions

2 FEDERATION PROCESS

The preliminary definition of the process to support the new nodes that are in the process of joining the XIFI federation was defined in the XIFI DOW [6]. It consists in assigning a support responsibility of a specific geographical partition of the European states, to a task 5.5 partner[6]. In this Chapter we describe the process followed in XIFI to support the federation. This process has been updated many times during the lifecycle of the federation activities in order to be more efficient and sustainable.

2.1 Description of the process

The methodological support (as described in [1]) is offered by XIFI to nodes willing to join the federation. The methodological support service of XIFI includes the following services:

- Access to documentation: all documents helping in informing and helping an Infrastructure Owner that wishes to join the XIFI federation will be provided. This includes - but is not limited to - the following documents (for some references, see the section 2.1.2):
    - XIFI Handbook (only for nodes): The XIFI Handbook describes in detail how to deploy a new XIFI node and how to install and configure FI services on top of it.
    - Procedures and protocols for XIFI federation: This documentation describes the procedures and the process an infrastructure has to follow in order to join the federation as well as the operational tasks that have to be fulfilled after joining the federation at a high level.
    - Documentation on needed GEs to be installed (only for nodes): This includes all information on the GEs provided by FIWARE that need to be set-up in order to run a node (e.g. DCRM, Monitoring GE). This information is mainly based on the information provided by FIWARE on these GEs (Catalogue entries, Admin-and Installation Guide, User-and Programmer-Guide and the unit testing plan).
    - Documentation on connecting different data sources to GEs: This includes all information on the GEs offered by FIWARE that are hosted in XIFI nodes and that can be used as data consumers of different data sources offered by an infrastructure, such as: Sensors, Smart Cities, Open Data, and so on.
    - Documentation on the test tool (only for nodes): The test tool is provided to a potential new node in order to test the fulfilment of the technical requirements. The information on how to use this tool will be provided to the user.
- The Infrastructure Toolbox (only for nodes): a software distribution that facilitates the installation of components needed for nodes. A helpdesk is established as a contact point for leading potential partners to the correct guidelines and help on the usage of tools. This helpdesk should provide the means for:
    - Requesting support;
    - Issuing a new JIRA ticket;
    - Monitoring the status of the tickets;
    - Contact help desk.

For practical reasons there will be only one central helpdesk that processes tickets and help-requests, the helpdesk is shared among all FIWARE activities.

The support process for potential nodes will be similar to the support process for internet developers described in the next section, besides that the ticket will be issued by a joining node and the forwarding will be done in most cases to an expert. Help requests on issues not related to the set up as a new node or to the connection of an infrastructure data source will not be processed. Depending on

frequent topics that are raised to the helpdesk a FAQ will be provided in the wiki pages of XIFI and will be moved to the FIWARE wiki before the of the project. The methodological support service excludes all services not mentioned in this section. This applies especially to support services that lead to further development efforts like the implementation of new adapters. The only exception to this are tickets that deal with bugs and problems investigated in the test tool and in the infrastructure toolbox.

In this paragraph we go deeper in the methodological support service process provided by the XIFI project for the inclusion of new nodes. We can define the deployment phase of a new node as the timeframe period that goes from the acceptance of the node in the federation until the operational phase (availability in the FIWARE Lab). The starting point of this phase is the acceptance of the new node (e.g. through the Open Call) and the exit point is when the new node is on production in the FIWARE Lab and accessible by users. As described in the chapter 3, a node is accepted in the operational phase only when the node passes a validation test. As defined in the Task 5.5 (see [6]) of WP5 of XIFI project, every new node is assigned to a geographical partition managed by a Task 5.5 partner; this partner will guide the new nodes through the process of joining the federation, in the context defined by the methodological support.

The tools used to support new node infrastructures are:

- Task 5.5 partner coaching
- Official XIFI Documentation
- ITBox (Infrastructure Toolbox)
- Helpdesk (JIRA – issue tracking software)
- FAQ
- Dedicated Meeting for new nodes.
- Training Sessions for new nodes ([7])

2.1.1 Partner coaching

The process of helping new nodes in joining the XIFI federation is guided by Task 5.5 partners, and it is based on a geographical partition, as showed below:

- Telecom Italia: South-East Europe (Italy, Austria, Hungary, Serbia, Croatia, Bosnia, Slovenia, Slovakia, Montenegro, Albania, Macedonia, Bulgaria, Romania, Greece and Turkey), specifically:
    - NeuroPublic S.A (Greece)
    - University of Piraeus (Greece)
    - University of Thessaly (Greece)
    - Wigner Research Center for Physics (Hungary)
- Waterford Institute of Technology: North Europe (Ireland, UK, Norway, Sweden and Finland), specifically:
    - Blekinge Institute of Technology (Sweden)
    - Acreo Swedish ICT AB (Sweden)
- Deutsche Telekom: Central-East Europe (Germany, Czech Republic, Denmark, Poland, Switzerland, Latvia, Estonia and Lithuania), specifically:
    - CESNET, z.s.p.o (Czech Republic)

- Poznan Supercomputing and Networking Center (Poland)
- Zürcher Hochschule für Angewandte Wissenschaften (Switzerland)

- ORANGE: Central-West Europe (France, Belgium, Netherlands and Luxemburg), specifically:
  - ASSOCIATION PLATE-FORME TELECOM (France)
  - iMinds (Belgium)

We will refer to "new node coach" as the role of a partner (task 5.5) to support the process of joining the federation of a new node. The partner assigned is the first contact that supports the new node in the process defined. Examples of support are: solving issues, analysing problems and gathering information. In terms of roles it has to be clarified that the coaching activity is not technical, since the technical support is covered by XIFI experts, as specified in [1], in case of a technical problem that has been raised during the deployment, the coach can help the new node to find a solution using the specific tools or processes.

2.1.2 Official XIFI Documentation

For the new nodes joining the federation the reference documents are:

- D1.1b XIFI core concepts, requirements and architecture [8]
- D5.1: Procedures and Protocols for XIFI federation [1]
- D5.2: XIFI Core Backbone [5]
- D2.4 (replace D2.1): XIFI Handbook v2 [9]
- D2.5 (replace D2.2): APIs and Tools for Infrastructure Federation v2 [4]
- D3.5 (replace D3.2): Infrastructures monitoring and interoperability adaptation components API v2 [11]
- D3.4 : XIFI infrastructure network adaptation mechanisms API [10]

2.1.3 ITBox

The ITBox tool [12] is the software package delivered by XIFI to help new nodes in the process of deploying the required services in their infrastructure and in joining the federation. The new nodes that are going to use the ITBox kit for deployment purpose are facilitated through the process of XIFI federation. ITBox includes all packages needed to successfully deploy a new XIFI node.

2.1.4 Helpdesk

The access to the Helpdesk is granted, also for the new infrastructure owners, through JIRA ticketing system, even in the deployment phase. As described in Figure 6, during the phase of deployment, the new infrastructures owners can use the helpdesk, for having support related to bugs.

2.1.5 FAQ

The section of the FAQ in the Task 5.5 of WP5 wiki page [15] is an important tool for the new infrastructures. FAQ is a liquid document, it means that it will be constantly maintained during the lifetime of the XIFI project. Main arguments or categories of the FAQ are:

- questions/answers related to technical issues
- How To related to technical topics

- suggestions and hints
- best practices

The actors of the FAQ are:

- XIFI Project Members
- XIFI infrastructure owner
- Task 5.5 partners

The process of feeding the FAQ is "in progress": it means that every XIFI project member can propose a new topic in the FAQ. Every new entry will be moderated by Task 5.5 members to ensure that it belongs to the categories defined above. Also, in terms of Task 5.5 activities, some partners are assigned to specific contents or arguments of the FAQ, based on their technological knowledge.

2.1.6 Dedicated meetings

In order to help the new nodes in starting their activities, some dedicated meetings are organized to:

- give new nodes the most important information needed to start activities (e.g. which is the geographical partner that will coach the new node);
- share technological issue, problems but also comments and hints between the new nodes and the stakeholders;
- share the status of the deployment plan.

The frequency of the meetings can vary from two weeks to one week (one week during the deployment phase or when there are important goal to achieve).

2.1.7 Training Sessions for new nodes

As described in [7], some dedicated training sessions were held for the new nodes (e.g. Madrid Session)

2.2 Process update: Associated Nodes

Associated Nodes are Infrastructures belonging to XIFI federation but not included in XIFI project. An associated node gives to XIFI users the same types of services of the nodes that are already federated.

2.2.1 Process to manage Associated Nodes

The process of federation in XIFI for an associated node is summarized in the following steps:

- an infrastructure owner requests the XIFI federation to the federation manager; the federation manager evaluates the request and if it is accepted the request to join, the Infrastructure Owner signs an agreement [2] for the inclusion of the node in the federation.
- after the signing, the node manager of the associated node is included in a ML (federation-nodes@fi-xifi.eu) that will be used to give basic support and information; in the federation-nodes@fi-xifi.eu mailing list are included:
    - XIFI Infrastructures Owners (Node Managers)
    - WP Leaders
    - XIFI Federation Manager, XIFI Technical Manager
    - Task 5.5 coach

- The purpose of the mailing related to the new associated nodes is to:
    - Inform all the stakeholders of the progress done in the federation process by the associated nodes
    - give to the associated node IO a basic tool for requesting help in case of problems, since the associated node is not part of the project this will be managed by partners with a best effort approach
- some information are given to the new associated node:
    - access to XIFI Wiki in read-only mode
    - information related to basic documentation
- a deployment plan is asked to the associated node and it will be tracked in WP5 task5.5
- the deployment plan is partitioned in some phases:
    - Connectivity: this is managed by the associated nodes
    - HW Procurement: this is managed by the associated nodes
    - Cloud platform Installation and Configuration: this is managed by the associated nodes with the help of XIFI documentation, task 5.5 FAQs, federations ML
    - Federation (FIWARE Lab Joining): this step of the federation is managed with the support of a XIFI expert; the XIFI experts will interact with the associated node using a defined document containing all the data needed to finalize the process. When this step has been successfully performed, the associate node is federated inside the cloud portal. Details about this process can be found in 2.3.
    - Monitoring: this step of the federation is managed with the support of a XIFI expert, the XIFI process will interact with the associate node using a defined document containing all the data needed to finalize the process. When this step has been successfully performed, the associate node is visible in infographic[13]web page.
    - Finalization of functionalities tests: a node is federated if the validation tests are successfully run on the node itself.

After these steps are finalized, the associate node belongs to XIFI federation.

2.2.1.1 Account for XIFI wiki (read-only for Associated Nodes)

In order to have an account for having read-only access to XIFI wiki, there are some steps to be performed by the user (associated node), as described below:

1. go to http://sso.fi-xifi.eu/openam/UI/Login?module=SelfReg&realm=/ and register yourself (don't use emails as usernames - wiki does not support that)
2. send an email to dev-tools-support@fi-xifi.eu ask for account activation (provide username)
3. wait until you receive a notification from the administrator that your account has been activated
4. login into the wiki using the Login and then Login(SSO) link (wiki.fi-xifi.eu)

2.3 Process of Federation Deployment

We can define the deployment phase as the timeframe period that goes from the acceptance of the node in XIFI until the production phase. The starting point of this phase is the acceptance of the new node (e.g. through the Open Call) and the exit point is when the new node is on production in XIFI.

The Phases of the Deployment of the federation process refer to the specification of a Pre-Deployment Phase and a Validation Tests Phase.

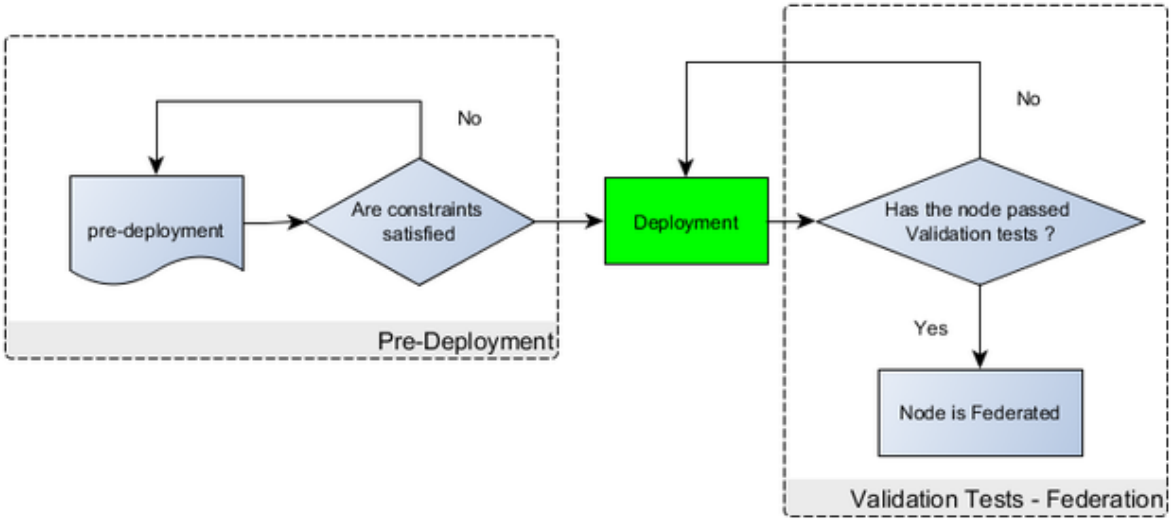

Figure 1: Federation Deployment phases

As shown in the previous figure the pre-Deployment Phase is related to the preparation of the Deployment of the node. During this phase the infrastructure owners, with the support of XIFI coach, assures that all the requirements [1]for the federation of the node are satisfied. After the deployment the node is considered fully operational in XIFI only when the set of validation tests are successfully passed (see 3.2).

2.3.1.1 Deployment process in details

The Deployment process consists in the sequence of the following phases:

- Connectivity: connectivity to Network is assured by local NREN of the node, the requirements are defined in [2].
- HW Procurement: for the requirements refers to [1].
- Cloud Infrastructure Installation and configuration: the installation and the configuration of OpenStack can be achieved with the help of ITBox or through a manual installation.
- FIWARE Lab Joining: during this phase the node is tested through some preliminary tests (see below); if the results of the tests are successful, the node is configured in the Cloud Portal, and in the centralized keystone proxy.
- Monitoring: this phase is started when the previous phase is successfully completed; during this phase the monitoring package is configured using the access specifications defined in the FIWARE Lab Joining phase, this step can be considered successfully achieved when the node is properly configured and visible in infographic.
- GEs: this phase can be done in parallel with the Monitoring step, otherwise it is usually performed after the monitoring phase to assure that the node is properly working. During this step the images, GEs and the blueprint images are loaded from a master node (Spanish node).
- Validation Tests: this is the last step and it is performed through a centralized tool, when all the tests have been successfully executed it can be considered fully operational in the federation of XIFI nodes.

In the following diagram the deployment phases are defined:

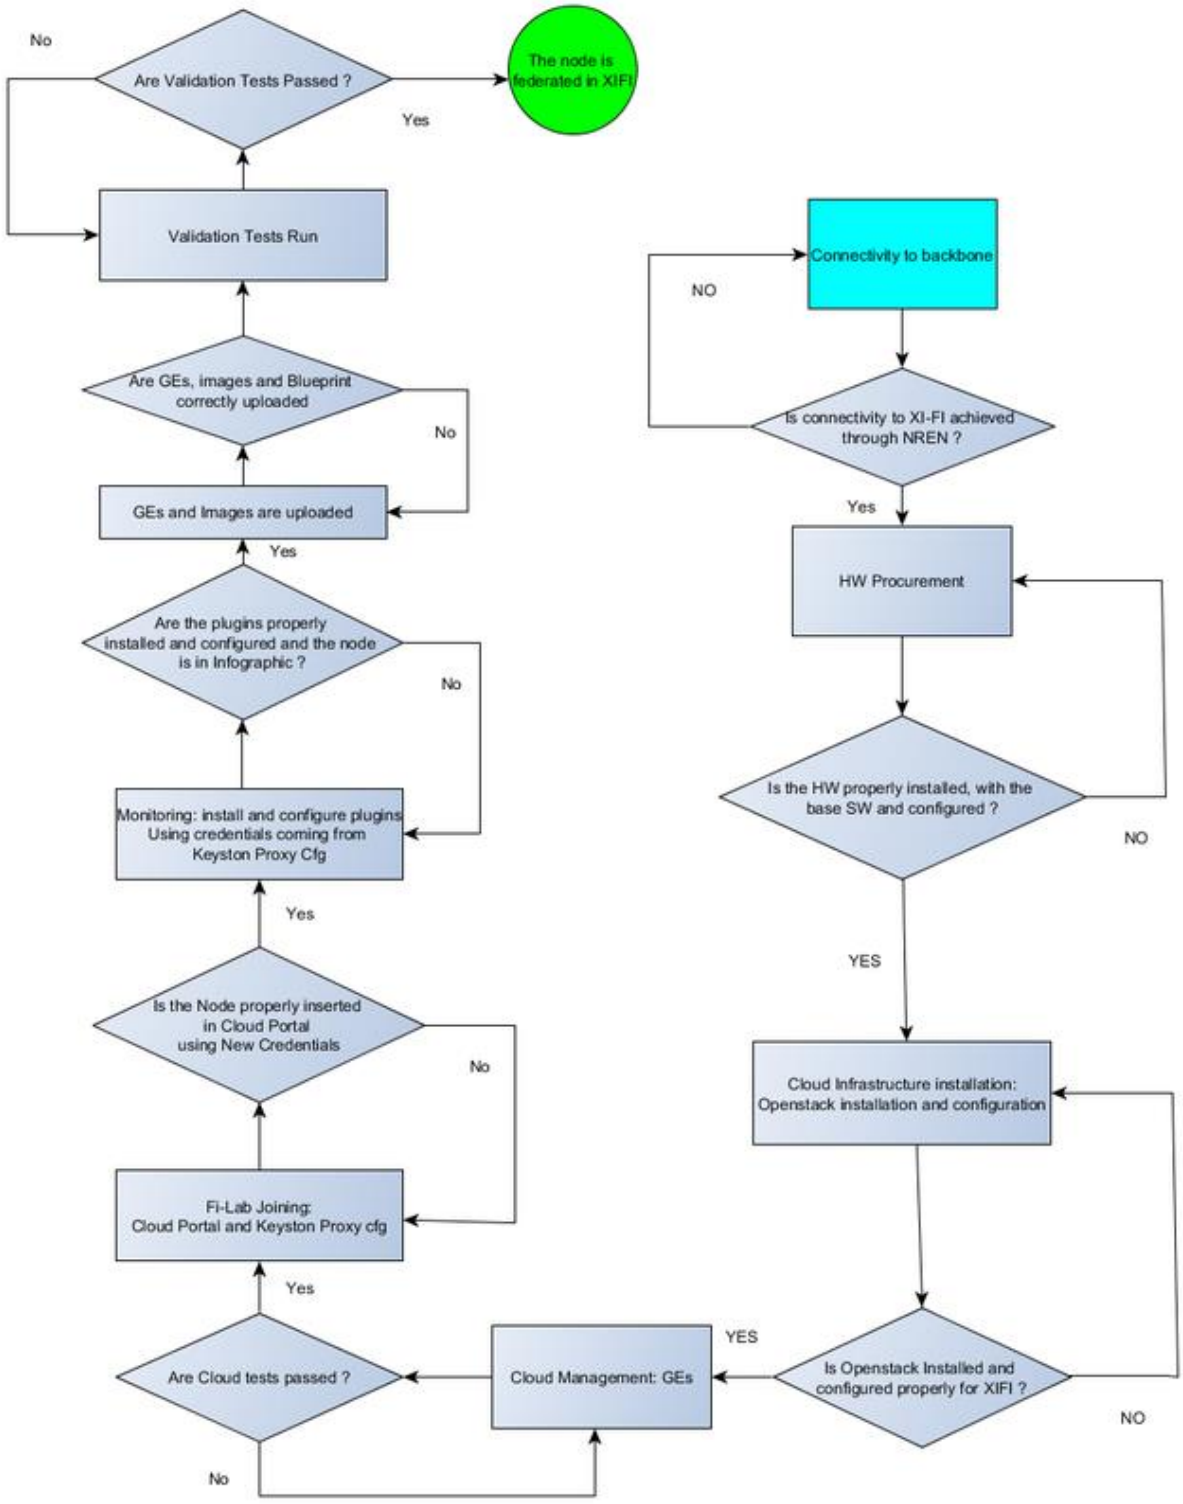

Figure 2: Federation Deployment diagram

As described in the above diagram, a set of tests (Cloud portal tests) were identified to be run before configuring a node in the cloud portal and assigning properly security access management (e.g. keystone proxy configuration).

The purpose of those tests is to verify that the node has the basic features of the cloud platform correctly working; those tests are:

- a node can deploy a VM to a user
- a node can deploy an image (deploy an image = register an image in glace)
- a node can deploy a network
- a node can attach a floating IP to a VM and it is reachable from outside

Those tests are performed by the node itself and it is part of the process of the configuration of the node in the cloud portal and in the centralized IDM (Keystone Proxy).

3 VALIDATION PROCESS

The operational status of a node is showed through the XIFI monitoring tools but it is complemented with a set of validation tests, described below, that are run by a centralized tool. Those tests will be improved in the FI-Core project, anyway a stable and robust set of tests is already used in XIFI project.

3.1 Status Management

A node is considered **Up&Running** only if all the tests defined as mandatory in the table 3.2 are successfully passed, the mandatory tests are listed in the table 3.2 and labelled as "Required for Up&Running node status" in the Mandatory Tests column. The validation tests are automatically run on all the federated nodes with a daily frequency (see 3.3).

The tests results are used daily by Infrastructures Owners (IOs) as part of the operational activities related to the node. In case of failures of mandatory tests the IOs will operate all the necessaries actions on the node in order to pass all the tests. The tests not marked as mandatory are also very important as complementary tests in order to maintain the node fully operational.

3.2 Validation Tests

This activity is taking the results offered by the development of the activities in the FI-Core project in order to show any possible problems in any of the OpenStack functionalities. These tests will be increased during the next months in order to test automatically all the needed functionalities. Currently, the activated tests are the following:

ID	Group	Type	Description	Expected results	Mandatory Tests
1	Base	Content	Check if the Region has flavors.	* The region has one or more flavors	--

ID	Group	Type	Description	Expected results	Mandatory Tests
2	Base	Content	Check if the Region has images.	* The region has one or more images	--
3	Base	Content	Check if the Region has images with 'init' in the name.	* The region has SDC-Aware images (with 'init' in its name) to be used by PaaS Manager	--
4	Base	TestSupport	Check if the Region has the BASE_IMAGE_ NAME used for testing.	* Exists a synchronized image BASE_IMAGE_NAME in the image region list (it will be used in deployment test cases)	--
5	Base	SecurityOps	Check if it is possible to create a new Security Group with rules.	* Sec. Group creation request is successful * Rules creation request is successful * Sec. Group ID is returned	--
6	Base	SecurityOps	Check if it is possible to create a new Keypair.	* KeyPair creation request is successful * Private key is returned	--
7	Base	SecurityOps	Allocate a public IP	* Public IP allocation request is successful * Pool name for the region must be configured in settings.json properties file. * Public IP ID is returned	Required for "Up&Running" node status
8	With Networks	NetworkOps	Check if it is possible to create a new Network with subnets	* Network creation request is successful * SubNetwork creation request is successful * New network is returned * New network has ACTIVE status	--
9	With Networks	NetworkOps	Check if there are external networks configured in the Region	* The region has one or more 'external' networks configured	--

ID	Group	Type	Description	Expected results	Mandatory Tests
10	With Networks	RouterOps	Check if it is possible to create a new Router without setting the Gateway	* Router creation request is successful (without Gateway) * New router is returned * New router has ACTIVE status	--
11	With Networks	RouterOps	Check if it is possible to create a new Router, with a default Gateway	* Router creation request is successful (with a default Gateway) ** If an external network for the region has been configured in settings.json, test case will use that one. ** Else, test cases will use the first external network found. * New router is returned * New router has ACTIVE status	--
12	With Networks	ServerOps	Check if it is possible to deploy a new Instance: Params: Name, FlavorID, ImageID, new NetworkID	* It can retrieve a flavor (""small"" or the last one if no ""small"" are found) * It can retrieve an image with name BASE_IMAGE_NAME * It can create a new network and sub-net * Instance creation request is successful (using above params in the request) * The instance has been deployed with ACTIVE status after 300 seconds MAX	Required for "Up&Running" node status
13	With Networks	ServerOps	Check if it is possible to deploy a new Instance: Params: Name, FlavorID, ImageID, new NetworkID, Metadatas	* It can retrieve a flavor (""small"" or the last one if no ""small"" are found) * It can retrieve an image with name BASE_IMAGE_NAME * It can create a new network and sub-net * Instance creation request is successful (using above params in the request) and custom metadatas * The instance has been deployed with ACTIVE status after 300 seconds MAX	Required for "Up&Running" node status

ID	Group	Type	Description	Expected results	Mandatory Tests
14	With Networks	ServerOps	Check if it is possible to deploy a new Instance: Params: Name, FlavorID, ImageID, new NetworkID, new Keypair	* It can retrieve a flavor (""small"" or the last one if no ""small"" are found) * It can retrieve an image with name BASE_IMAGE_NAME * It can create a new network and sub-net * It can create a new KeyPair value * Instance creation request is successful (using above params in the request) * The instance has been deployed with ACTIVE status after 300 seconds MAX	Required for "Up&Running" node status
15	With Networks	ServerOps	Check if it is possible to deploy a new Instance: Params: Name, FlavorID, ImageID, new NetworkID, new Sec. Group	* It can retrieve a flavor (""small"" or the last one if no ""small"" are found) * It can retrieve an image with name BASE_IMAGE_NAME * It can create a new network and sub-net * It can create a new Sec. Group with with one rule * Instance creation request is successful (using above params in the request) * The instance has been deployed with ACTIVE status after 300 seconds MAX	Required for "Up&Running" node status
16	With Networks	ServerOps	Check if it is possible to deploy a new Instance: Params: Name, FlavorID, ImageID, NetworkID, Sec. Group, keypair, metadata	* It can retrieve a flavor (""small"" or the last one if no ""small"" are found) * It can retrieve an image with name BASE_IMAGE_NAME * It can create a new network and sub-net * It can create a new KeyPair value * It can create a new Sec. Group with with one rule * Instance creation request is successful (using above params in the request) and custom metadatas * The instance has been	Required for "Up&Running" node status

ID	Group	Type	Description	Expected results	Mandatory Tests
				deployed with ACTIVE status after 300 seconds MAX	
17	Without Networks	ServerOps	Check if it is possible to deploy a new Instance: Name, FlavorID, ImageID, Metadatas	* It can retrieve a flavor (""small"" or the last one if no ""small"" are found) * It can retrieve an image with name BASE_IMAGE_NAME * Instance creation request is successful (using above params in the request) and custom metadatas * The instance has been deployed with ACTIVE status after 300 seconds MAX	Required for "Up&Running" node status
18	Without Networks	ServerOps	Check if it is possible to deploy a new Instance: Name, FlavorID, ImageID, new Keypair	* It can retrieve a flavor (""small"" or the last one if no ""small"" are found) * It can retrieve an image with name BASE_IMAGE_NAME * It can create a new KeyPair value * Instance creation request is successful (using above params in the request) * The instance has been deployed with ACTIVE status after 300 seconds MAX	Required for "Up&Running" node status
19	Without Networks	ServerOps	Check if it is possible to deploy a new Instance: Name, FlavorID, ImageID, new Sec. Group	* It can retrieve a flavor (""small"" or the last one if no ""small"" are found) * It can retrieve an image with name BASE_IMAGE_NAME * It can create a new Sec. Group with with one rule * Instance creation request is successful (using above params in the request) * The instance has been deployed with ACTIVE status after 300 seconds MAX	Required for "Up&Running" node status
20	Without Networks	ServerOps	Check if it is possible to deploy a new Instance: Name,	* It can retrieve a flavor (""small"" or the last one if no ""small"" are found) * It can retrieve an image with name BASE_IMAGE_NAME	Required for "Up&Running" node status

ID	Group	Type	Description	Expected results	Mandatory Tests
			FlavorID, ImageID, Sec. Group, keypair, metadata	* It can create a new KeyPair value * It can create a new Sec. Group with with one rule * Instance creation request is successful (using above params in the request) and custom metadatas * The instance has been deployed with ACTIVE status after 300 seconds MAX	

Table 2: List of Validation Tests for the XIFI nodes

There are several fields in order to identify each of the tests. The description of each of them is the following:

Group. It gives detail about the type of test that it is to realize. The possible values are:

- Base, it corresponds to all basic tests like management of flavors and images (Content Type), security groups and keypair (SecurityOps Type) and confirm that the image that it is used in the test (TestSupport Type).
- With Networks, it corresponds with the different tests related to the deployment and configuration of virtual machine using network functionalities (neutron service). You have two options in that case, the operations related to the creation of the networks and subnetworks (NetworkOps type) and the operations related to the creation and configuration of routers (RouterOps type).
- Without Networks, it corresponds to the operations to deploy a server instance without networks specification (ServerOps type).

Type. Provide details about the type of operations that we want to implement. It is related to the group field and the different values are described in the previous field.

Description. Give us a description about the operation that we want to test together with the parameters that we want to use.

Expected results. Give us a description about the possible response of the operations that the test has launched.

Mandatory Tests. Show us if one test has to be mandatory. All the mandatory tests have violet colour as background.

3.3 Automated Test Tool

We have deployed in an instance of FIWARE Lab and associated to the FI-Core project a continuous integration tool that automatically launches the execution of the validation test. This tool is Jenkins and you can find it at the URL http://fi-health.lab.fi-ware.org:8080/ .

Currently, the Jenkins home page is the following:

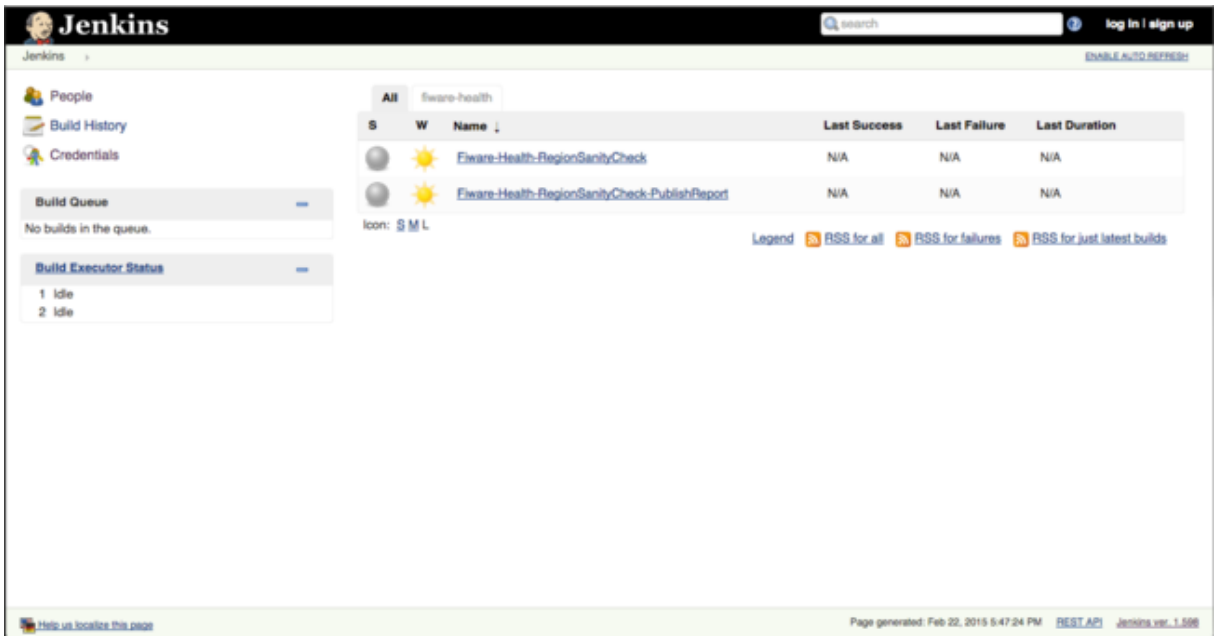

Figure 3: Home Page Jenkins

In the above figure are currently shown two jobs:

- The first one, whose name is Fiware-Health-RegionSanityCheck, executes periodically the validation tests and produces a text file with the execution results.
- The second one, whose name is Fiware-Health-RegionSanityCheck-PublishReport, publishes to a web server the content of the last execution of the validation test so that Infrastructure Owners can access to it and see any problem detected. The URL of this page is http://fi-health.lab.fiware.org/RegionSanityCheck/

The validation test report will have the following aspect:

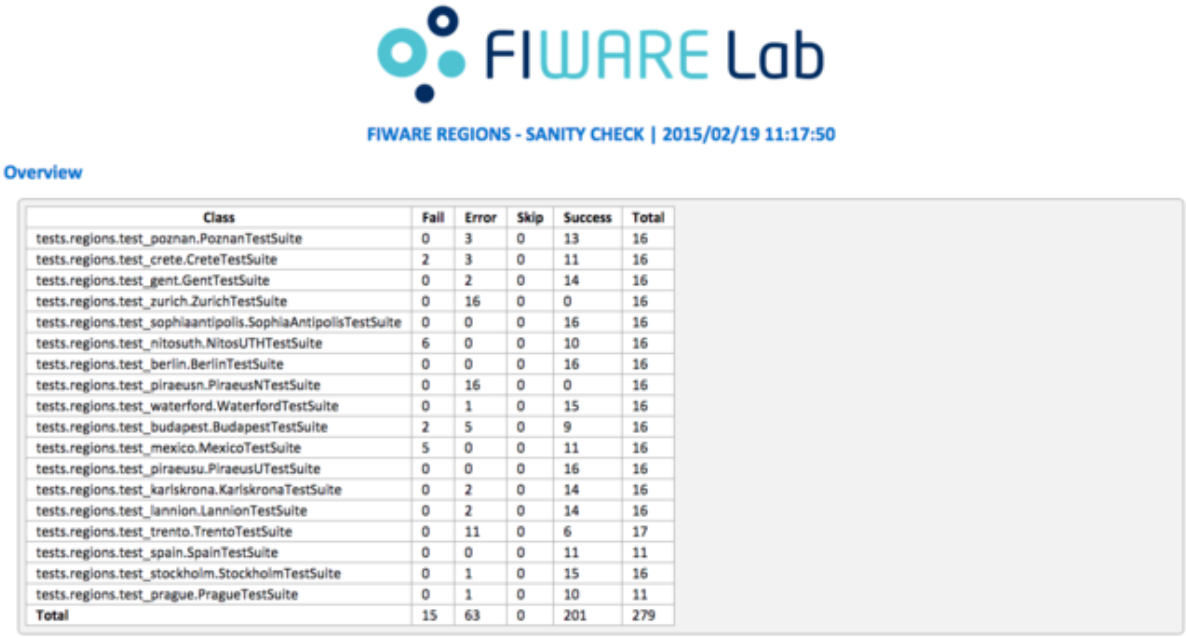

FIWARE REGIONS - SANITY CHECK | 2015/02/19 11:17:50

Overview

Class	Fail	Error	Skip	Success	Total
tests.regions.test_poznan.PoznanTestSuite	0	3	0	13	16
tests.regions.test_crete.CreteTestSuite	2	3	0	11	16
tests.regions.test_gent.GentTestSuite	0	2	0	14	16
tests.regions.test_zurich.ZurichTestSuite	0	16	0	0	16
tests.regions.test_sophiaantipolis.SophiaAntipolisTestSuite	0	0	0	16	16
tests.regions.test_nitosuth.NitosUTHTestSuite	6	0	0	10	16
tests.regions.test_berlin.BerlinTestSuite	0	0	0	16	16
tests.regions.test_piraeusn.PiraeusNTestSuite	0	16	0	0	16
tests.regions.test_waterford.WaterfordTestSuite	0	1	0	15	16
tests.regions.test_budapest.BudapestTestSuite	2	5	0	9	16
tests.regions.test_mexico.MexicoTestSuite	5	0	0	11	16
tests.regions.test_piraeusu.PiraeusUTestSuite	0	0	0	16	16
tests.regions.test_kariskrona.KariskronaTestSuite	0	2	0	14	16
tests.regions.test_lannion.LannionTestSuite	0	2	0	14	16
tests.regions.test_trento.TrentoTestSuite	0	11	0	6	17
tests.regions.test_spain.SpainTestSuite	0	0	0	11	11
tests.regions.test_stockholm.StockholmTestSuite	0	1	0	15	16
tests.regions.test_prague.PragueTestSuite	0	1	0	10	11
Total	15	63	0	201	279

Figure 4: Validation Test Report

If we consider a specific region, we could see detailed information about the different tests that were executed. The green lines correspond to the success test and the red lines correspond to the failed or error tests. We can see an example in the following image related to the Sophia Antipolis FIWARE Lab node:

tests.regions.test_sophiaantipolis.SophiaAntipolisTestSuite (0 failures, 0 errors)

- test_allocate_ip
- test_base_image_for_testing_exists
- test_create_keypair
- test_create_network_and_subnet
- test_create_router_external_network
- test_create_router_no_external_network
- test_create_security_group_and_rules
- test_deploy_instance_with_all_params
- test_deploy_instance_with_new_network
- test_deploy_instance_with_new_network_and_keypair
- test_deploy_instance_with_new_network_and_metadatas
- test_deploy_instance_with_new_network_and_sec_group
- test_flavors_not_empty
- test_images_not_empty
- test_there_are_external_networks
- test_there_are_init_images

Figure 5: Validation Test Report for one region

The difference between a failure and an error is that the first we obtain a value in the execution of the test but it is not the expected value at all. By contrast, an error means that we cannot execute the test due to an error communication or whatever type of error that should be investigated. In the following release of this component, we have planned the automatic communication with each infrastructure owners in order that they receive as soon as possible any detected problems on their FIWARE Lab node.

4 FEDERATION STATUS

The purpose of this chapter is to give an overview of the status of the nodes of XIFI federation. The federation activities of XIFI are related to the new nodes included in the project through an Open Call. In the next paragraph it is summarized the plan of the federation of the new nodes using the phases defined in the chapter 2.

4.1 Federation New Nodes Activities

In the following table is showed the federation plan finalization for the new nodes (the 11 nodes related to XIFI Open Call [14]).

Organization (Country)	Release Date Connectivity to XiFi Core Backbone	Release Date HW procurement	Release Date Cloud Infrastructure Installation	Release Date Cloud Management (GE)	Release Date Monitoring	Release Date Fi-LAB Joining
NeuroPublic (GR)	30/Sept/2014	05/Sept/2014	09/Sept/2014	09/Sept/2014	02/Oct/2014	07/Oct/2014
UPRC (GR)	04/Sept/2014	15/Sept/2014	26/Sept/2014	26/Sept/2014	26/Sept/2014	13/Oct/2014
UTH (GR)	31/Aug/2014	20/July/2014	25/Sept/2014	25/Sept/2014	29/Sept/2014	09/Oct/2014
Wigner (HU)	31/July/2014	30/Apr/2014	30/Aug/2014	15/Sept/2014	30/Sept/2014	10/Oct/2014
BTH (SE)	15/June/2014	23/Sept/2014	15/Sept/2014	15/Sept/2014	02/Oct/2014	13/Oct/2014
Acreo (SE)	05/May/2014	25/May/2014	31/Aug/2014	31/Aug/2014	02/Oct/2014	20/Oct/2014
Com4innov (FR)	08/Aug/2014	31/Aug/2014	18/Sept/2014	18/Sept/2014	27/Oct/2014	21/Oct/2014
CESNET (CZ)	02/May/2014	1/Apr/2014	21/May/2014	30/June/2014	07/Oct/2014	30/May/2014
PSNC (PL)	30/June/2014	30/June/2014	26/Sept/2014	26/Sept/2014	29/Sept/2014	13/Oct/2014
ZHAW (CH)	21/Sept/2014	07/Aug/2014	02/Oct/2014	03/Oct/2014	09/Oct/2014	16/Oct/2014
iMinds (BE)	25/Sept/2014	01/May/2014	15/Sept/2014	30/Sept/2014	09/Oct/2014	13/Oct/2014

Table 3: Deployment Plan of the New Nodes (Open Call)

Legend of the table above:

- Connectivity to XIFI Core Backbone:

MD-VPN connectivity through the local NREN (information and example at [5].

- HW procurement:

It means hardware procured and deployed with the base operating system

- Cloud Infrastructure Installation:

This is basically the OpenStack installation. It can be installed manually (more information at FAQ section) or with ITBox tool (reference at [12]).

- Cloud Management (GE):

If you will use ITBox for OpenStack installation, this step is included. Otherwise you will need install manually the DCRM GE (reference at FAQ section). After the installation of the DCRM you would download from the XIFI Master Node the images of the commons GEs and also the SDC_aware images to be used by the PaaS Manager (reference at FAQ section).

- Monitoring:

If you will use ITBox for OpenStack installation, this step is included. Otherwise you will need to install manually the plugins (reference at FAQ section).

- Fi-LAB Joining and Infographic:

This is essentially the installation and configuration of the Keystone Proxy module, the finalization of the monitoring Configuration (using the credentials coming from Keystone Proxy) and the steps needed to have the node visible on Infographic (more information at FAQ section).

4.2 Status of all Federation Nodes

The content of this paragraph was provided directly by the IOs through a survey run between 16th February and 3rd March 2015. Instead, the results of the validation tests are referred to the 13th of March. The purpose is to give an overall view of the operational status in terms of availability of the services.

Generally speaking, we define the following status of the node:

- **Up & Running**: the node has passed with no error all the mandatory tests.
- **Operational**: the node has failed some tests due to lack of resources (e.g. floating IP). The node is able to run all the services and to manage the resources already allocated, but it is not able to provide further resources to the users.
- **Not Operational**: the errors found in the tests detect technical problems that do not allow to the node of being operational.

We will use this definition for each node in the following paragraphs.

4.2.1 Status of NeuroPublic node

Node Name	OpenStack release	Overall status of the node	Validation tests results
PiraeusN	Icehouse	Up and Running	0 Errors/16 Passed

Table 4: Status of NeuroPublic node

4.2.1.1 Description of the node status

NeuroPublic node runs OpenStack Icehouse version deployed on Ubuntu 12.04.1 with the use of Mirantis Fuel deployment tool version 5.1.1. Apart from the core components of the node NeuroPublic uses a proxmox VE server which hosts the monitoring virtual machines. There is one virtual machine running Ubuntu 12.04.5 on which ngsi_adapter (version 1.1.1) and ngsi_event_broker (version 1.3.1) run, and one virtual machine running CentOS 6.5 on which the Orion context broker (version 0.15) is located.

NeuroPublic deployed a High Availability architecture that consists of:

- 1x Fuel deployment server used to deploy the whole environment.
- 3x Controller Nodes which also host the neutron services and the Ceph storage services.
- 4x Compute nodes
- 1x Proxmox VE hosting the monitoring VM's

4.2.1.2 History of the node availability

NeuroPublic node was up and running OpenStack Grizzly since 7/10/2014[3].

Late November 2014 it was decided that NeuroPublic node will upgrade to Icehouse so the node was not operational from 1/12/2014 to 18/12/2014.

From 18/12/2014 NeuroPublic was operational in a "beta" phase solving problems the migration to Icehouse has caused.

From 12/1/2015, the node was fully operational running OpenStack Icehouse

4.2.2 Status of UPRC node

Node Name	OpenStack release	Overall status of the node	Validation tests results
PiraeusU	Icehouse	Up and Running	0 Errors /16 Passed

Table 5: Status of UPRC node

4.2.2.1 Description of the node status

PiraeusU node hosts OpenStack Icehouse release, on CentOS 6.5. It is comprised by 6 different nodes that are allocated as following:

- 1 x Controller node (8 Cores, 16 GB RAM, 1 GBps Connectivity, 0,6TB storage).
- 3 x Compute and Cinder node (32 Cores, 64 GB RAM, 1 GBps Connectivity, 0,5TB storage).
- 1 x Compute and Cinder (16 Cores, 32 GB RAM, 1 GBps Connectivity, 0,5TB storage).
- 1 x Monitoring node (8 Cores, 16 GB RAM, 1 GBps Connectivity, 0,2TB storage).

For the networking the node exploits the offered capabilities by neutron network agent. It provides two different external networks that essentially constitute the floating IP pools for the node.

The first external network named as public-ext-net-1 and includes a subnet with prefix /26 (50 public IPs available for the node tenants), whereas the second network is named as 'federation-ext-net-1' and includes a subnet with /20 prefix (4096 md-vpn IPs for the node tenants). Further to that, node respects and follows the defined policies by XIFI community for the FIWARE resources usage. In particular by defining specific quotas for all tenants, limits the resources usage, such as the VM instances (3 per tenants), the Floating IPs (1 per tenant) and so on.

For the networking PiraeusU node uses 1 l3_agent that manages the two different external networks,

[3] The date format inside specific reporting activities will be from this point until the end of the document DD/MM/YYYY.

where as it provides a private network (the 'node-int-net-1') as shared network for the tenants. Each tenant can use directly the private shared network while creates its instances or alternatively can create it own private network by assigning it in turn to its own instances. Moreover, each tenant specific private network can be associated with the external networks (public-ext-net-* and/or federation-ext-net-*) by creating and configuring its own private router through the corresponding operations that are provided by cloud portal.

Finally, each tenant can create its own volumes that can be attached to its VM instances so as to extend their storage capabilities, whereas the tenant has the ability to create snapshots of their instances so as to ensure a backup of its user-specific data.

4.2.2.2 History of the node availability

PiraeusU node became available for the first time on 13/10/2014 and it was Up and Running until the first half of January 2015.

The period between 15/01/2015 to 21/01/2015 it was Down due to it was in upgrade process from OS Grizzly to OS Icehouse release.

By 21/01/2015 until today the node is Up and Running with OpenStack Icehouse release.

4.2.3 Status of UTH node

Node Name	OpenStack release	Overall status of the node	Validation tests results
Volos	Icehouse	Up and Running	0 Errors/16 Passed

Table 6: Status of UTH node

4.2.3.1 Description of the node status

Volos node comprises 7 HP blade servers, each one contributing with 32 cores and 98 Gb of RAM. In addition, 1 HP 380 hosting the controller node and the monitoring one. As far as the Storage is concerned VOLOS is using the shared storage approach. The setup is aligned with the official documentation of OpenStack and Mirantis. At the moment the node is hosting OpenStack Icehouse version and it is "live migration" enabled.

Furthermore, the node is configured following the XIFI documentation, i.e. D2.1 (XIFI Handbook) and it is aligned with the recommended policies and directions (quotas appliance, etc).

Components	Description	Comments
Servers	7 x HP blade servers 1 x HP DL380p server 1 x Storage Server	

Components	Description	Comments
Total Capacity	· CPU: 256 Cores · RAM: 784 GB · HDD capacity: 13 TB o Shared storage(CEPH): 9TB	
Per server capacity	Blade Servers · CPU : 2 processors (32 cores per server) · RAM : 98 GB · HDD (local): 2 x 300 6G SAS 10k disks · Network: o 2 x 10Gbit Ethernet NIC o 4 x 1Gbit Ethernet NIC DL380p server · CPU : 2 processors (24 cores) · RAM : 64 GB · HDD (local): 5 x 450 6G SAS 10k disks · Network: o 2 x 10Gbit Ethernet NIC o 4 x 1Gbit Ethernet NIC Storage Server: · CPU: 1 processor (8 cores) · RAM: 16 GB · HDD: 9TB · Network: 2x 10Gbit Ethernet NIC	
Switch	HP 5412r with 1Gb and 10Gb modules	OpenFlow 1.3
Firewall	Software solution	

Table 7: Description of UTH node

4.2.3.2 History of the node availability

Volos node became available on 18/10/2014 on the federation portal and was available until the end of December 2014.

From the beginning of January 2015 until the end of the same month the Volos node was not available in the portal due to upgrading of OpenStack Icehouse.

Since then the node is up and running.

4.2.4 Status of Wigner node

Node Name	OpenStack release	Overall status of the node	Validation tests results
Budapest	Grizzly	Operational	12 Passed/ 2 Errors

Table 8: Status of Wigner node

Note that the errors in the validation tests are due to the reason that there is no free IP address in the public floating IP pool. Currently, all the 150 IPs are under usage. The public IP pool will be extended after the migration. We also have to note that currently Budapest site is slightly overloaded, hosting 150 active VMs. However, after Icehouse migration, with the involvement of new computational resources previously mentioned, this overloaded situation is expected to be solved.

4.2.4.1 Description of the node status

The node is currently upgrading to OpenStack Icehouse. The upgrade to Juno is awaited for confirmation of compatibility with cloud portal and GEs. Originally 2 blade servers out of the 16 were dedicated to experimenting with Icehouse, but Budapest node managed to involve additional 16 blade servers. At the moment, 14 servers (112 CPU cores, 224 GB RAM, 60TB shared disk) are dedicated to the Grizzly installation, and 18 take part in the Icehouse deployment.

After the migration, 28 blade servers (224 cores, 548 GB RAM, 60 TB shared disk) are intended to be offered through FIWARE Lab, and 4 servers (32 cores) are kept for the preparation of a Juno environment.

The node uses two storage servers (SAN over InfiniBand) of capacities 10 TB and 50 TB for serving the root filesystems of blade servers and for user created volumes, respectively. IPs from two floating IP range can be assigned to the VMs: 1. a /23 public IPv4 range (almost 512 addresses, currently 150 are in the IP pool of Grizzly installation) from the /16 domain of Wigner RCP, 2. a /20 MDVPN range.

4.2.4.2 History of the node availability

- The node was federated at the beginning of October. Since then, it operates continuously with only minor interruptions. In the early stages, some issues were reported by end users, which were handled rapidly by the node support team.
- A two days long outage was on 26-27 January 2015, due to a scheduled maintenance that affected the energy supply of the entire cluster. The stop was announced to the user community in due time, few weeks earlier.
- Further scheduled outage is planned during March when Icehouse migration is going to be done.

4.2.5 Status of CESNET node

Node Name	OpenStack release	Overall status of the node	Validation tests results
Prague	Grizzly	Up and Running	0 Errors/11 Passed

Table 9: Status of CESNET node

4.2.5.1 Description of the node status

The node is still following original Grizzly release XIFI recommendation with IPv6 support. Migration to Juno is now under testing. Detailed HW configuration is described in the following table:

Component	Description
Total Capacity	216 CPUs, 1152 GB RAM, 13 TB disk space, n x 10G network
Servers	8x SuperMicro H8DGT (A+ 2122TG) 4x Apple Mac Pro (Z0PK)
Per server capacity	SuperMirco: 24 cores AMD Opteron 6344, 128 GB RAM, four hot-plug 10k drives, RAID10 for performance and reliability, SSD Apple: 6 Intel cores, 3.5GHz, 32 GB RAM, 256 GB PCI-e SSD, 4x Gig-Ethernet
Network	Cisco Nexus 3172PQ, SDN support, SFP+, QSFP, 10G connection to CESNET's backbone and GÉANT, n x 10G uplinks, IPv6 routing Juniper EX4300-24T, SDN support, 1000Base-T, SFP+, QSFP

Table 10: Hardware and Network components of Prague node

4.2.5.2 Description of the problems experienced

- In early 2014, a user launched a single VM which started generating 800 Mbps of a Denial-of-Service network traffic targeted against a number of foreign systems. The administrators of the identity management service didn't provide us with details about the identity of the user who owned this VM. We had no choice but to forcefully terminate it. We do not know whether the operators of the central Keystone instance blocked the user's identity.
- Number of the VMs running in Prague has grown steadily over the months. Because the XIFI project doesn't have an established process of reclaiming resources of inactive users, we had to pause the VMs and terminate them after a timeout.

4.2.5.3 Validation test "test_allocate_ip" failed

It reflects the state of utilization of CESNET resources: people are using VMs in Prague, and many of them allocated IPv4 addresses. This needs to be solved on the federation-wide basis. This is communicated well by the actual error message: NoMoreFloatingIps: Zero floating IPs available. Solution: currently, CESNET solved the lack of IPv4 addresses by assigning to all users VMs a publicly-routable IPv6 address by default.

4.2.5.4 History of the node availability

- Up and running on Grizzly since 05/2014 in limited mode.
- Fully federated during summer 2014.
- Migration to Juno expected during March 2015.

4.2.6 Status of PSNC node

Node Name	OpenStack release	Overall status of the node	Validation tests results
Poznan	Icehouse	Operational	13 Success, 3 errors

Table 11: Status of PSNC node

Failed tests are related to assigning of floating IP, which is blocked in our node. Due to lack of public IP addresses we decided to manually assign public IP addresses to tenants on demand, after sending request to us.

4.2.6.1 Description of the node status

PSNC node has 20 8-cores servers with 8-12 GB RAM with connection to SAN utilizing 5TB of external storage.

The node was deployed using Fuel 5.1.1. It is composed of Fuel node, 1 controller, 1 storage and 17 compute nodes. There are 2 external networks: public with a /27 subnet and a MDVPN XIFI federated network /24 subnet.

Currently, the PSNC node is working with the following resources:

- 120 cores
- 156 GB of RAM
- 5 TB of disks

Tests that are failed:

- test_allocate_ip,
- test_create_router_external_network,
- test_create_router_no_external_network

4.2.6.2 History of the node availability

The node PSNC was up and running from 13/10/2014.

The node PSNC had a downtime from 22/12/2014 to 9/1/2015 due to upgrade to Icehouse OpenStack.

The node PSNC was up and running from 9/1/2015.

4.2.7 Status of ZHAW node

Node Name	OpenStack release	Overall status of the node	Validation tests results
Zurich	Icehouse	Operational	2 Errors/14 Passed

Table 12: Status of ZHAW node

Failed tests are due to lack of floating IP resources.

4.2.7.1 Description of the node status

The ZHAW node is operational. It comprises 1 Fuel master deployment node, 8 nodes running OpenStack - 1 controller node and 6 compute nodes and 1 storage node.

Currently, the ZHAW node is working with the following resources:

- 208 cores
- 1.6 TB RAM
- 44 TB Disk storage

It has connectivity to the Internet via the aggregated ZHAW high-speed connection as well as high speed direct access to the MD-VPN via a GEANT POP in Geneva. The node contains a Pica8 3290 SDN/Openflow capable switch.

4.2.7.2 History of the node availability

The node was federated in early October 2014. While most of the services were working, there were some issues with the routing arising from the instability of the Grizzly routing agent - a known problem in OpenStack.

It was not possible to solve this within the project so the solution was to move to Icehouse as quickly as possible.

During this period, problems arose with the federation connectivity which required input from the Swiss NREN to address.

The node is live and stable as of late February 2015.

4.2.8 Status of Com4Innov node

Node Name	OpenStack release	Overall status of the node	Validation tests results
SophiaAntipolis	Grizzly	Not Operational	16 Errors

Table 13: Status of Comm4Innov node

4.2.8.1 Description of the node status

SophiaAntipolis node used ITBox 2.3.4.1 to deploy OpenStack and current XIFI tools.

The deployment used and respected the technical requirements described in D5.1.

The node is composed by one virtual server 2 cpus / 2 Gb of Ram for the Fuel/ITBox usage.

3 Dell poweredge R820 providing 96 cores and 192 Gb of Ram memory with 4,8 TB of storage (3 x 1,6 TB per host in RAID 5) to delivering the OpenStack/XIFI current services.

We use a 1Gb internal network link vlan-based in our datacenter. The Internet connection is provided by the datacenter Internet service provider with 1Gb link.

The node is connected to the MDVPN by a dedicated link based on Vlan provided by RENATER, which provided the Public IP subnet 193.48.247.192/26 we have allocated.

We have allocated XIFI 10.0.224.0/20 subnet provided by the federation.

The following table shows the hardware resources of SophiaAntipolis node:

Hardware resources SophiaAntipolis	
Quantity	**Description**
1	Dell PowerEdge R820 32 CPU Intel XEON E5-4620 at 2,2 GHz 64 Gb of Ram, 1,6 To RAID 5 ESXi 5.5 > Virtual Machine ItBox 1.3.4.1 2 cores, 2Gb RAM, 120Gb HDD
1	Server Controller: Dell PowerEdge R820 32 CPU Intel XEON E5-4620 at 2,2 GHz 64 Gb de Ram, 1,6 To RAID 5, 4 x1Gb Nic
1	Cinder & Compute Server: Dell PowerEdge R820 32 CPU Intel XEON E5-4620 at 2,2 GHz 64 Gb de Ram, 1,6 To RAID 5, 4 x1Gb Nic Volume Cinder: 1 To Volume Compute: 600 Gb
1	Cinder & Compute Server: Dell PowerEdge R820 32 CPU Intel XEON E5-4620 at 2,2 GHz 64 Gb de Ram, 1,6 To RAID 5, 4 x1Gb Nic Volume Compute: 1,6 To
1	Switch Pica8 Pronto 3290 48 Ports Openflow

Table 14: SophiaAntipolis Hardware resources

Statistics of the node usage:

```
root@node-1:~# nova --os-region SophiaAntipolis usage-list
Usage from 2015-01-28 to 2015-02-26:
+----------------------------------+-----------+--------------+-----------+-----------+
| Tenant ID                        | Instances | RAM MB-Hours | CPU Hours | Disk GB-Hours |
+----------------------------------+-----------+--------------+-----------+-----------+
| 00000000000000000000000000000135 | 1         | 344064.04    | 672.00    | 0.00      |
| 00000000000000000000000000002107 | 4         | 11010049.22  | 5376.00   | 107520.01 |
| 00000000000000000000000000004463 | 2         | 170515.27    | 333.04    | 0.00      |
| 00000000000000000000000000005111 | 3         | 2064384.23   | 2016.00   | 13440.00  |
| 00000000000000000000000000006470 | 1         | 319322.95    | 623.68    | 0.00      |
| 00000000000000000000000000007157 | 1         | 2752512.31   | 1344.00   | 26880.00  |
| 00000000000000000000000000007161 | 1         | 2752512.31   | 1344.00   | 26880.00  |
| 00000000000000000000000000007968 | 1         | 2752512.31   | 1344.00   | 26880.00  |
| 00000000000000000000000000008419 | 2         | 317479.18    | 620.08    | 0.00      |
| 00000000000000000000000000008430 | 5         | 13762561.53  | 6720.00   | 134400.01 |
| 00000000000000000000000000008485 | 1         | 5505024.61   | 2688.00   | 53760.01  |
| 00000000000000000000000000008822 | 4         | 288668.44    | 140.95    | 2819.03   |
| 00000000000000000000000000008916 | 1         | 11010049.22  | 5376.00   | 107520.01 |
| 00000000000000000000000000009883 | 3         | 8257536.92   | 4032.00   | 80640.01  |
| 00000000000000000000000000009884 | 3         | 8257536.92   | 4032.00   | 80640.01  |
| 00000000000000000000000000009886 | 3         | 8257536.92   | 4032.00   | 80640.01  |
| 00000000000000000000000000009889 | 3         | 8257536.92   | 4032.00   | 80640.01  |
| 00000000000000000000000000009891 | 1         | 11010049.22  | 5376.00   | 107520.01 |
| 00000000000000000000000000010166 | 1         | 1376256.15   | 672.00    | 13440.00  |
| 00000000000000000000000000010259 | 2         | 1748.05      | 1.28      | 14.23     |
| 00000000000000000000000000010260 | 1         | 3993.60      | 1.95      | 39.00     |
| 00000000000000000000000000010340 | 1         | 344064.04    | 672.00    | 0.00      |
| 00000000000000000000000000010422 | 1         | 1376256.15   | 672.00    | 13440.00  |
| 00000000000000000000000000010538 | 3         | 1090.56      | 0.53      | 10.65     |
| 00000000000000000000000000010547 | 4         | 7200.43      | 3.52      | 70.32     |
| 00000000000000000000000000010769 | 3         | 92654.93     | 45.24     | 904.83    |
| 00000000000000000000000000010815 | 1         | 1989442.55   | 971.41    | 19428.15  |
| 00000000000000000000000000011075 | 1         | 364638.44    | 178.05    | 3560.92   |
| 00000000000000000000000000011233 | 31        | 119.47       | 0.06      | 1.17      |
| 0000000000000000000000000000admin | 1         | 22668.02     | 44.27     | 0.00      |
+----------------------------------+-----------+--------------+-----------+-----------+
```

Figure 6: Usage statistics SophiaAntipolis

4.2.8.2 History of the node availability

The SophiaAntipolis node is available since November 2014.

The node had the following issues:

- Instability of the node between Controller and Compute:

  Why: Puppet service, launched by cron service, rewrote some configurations files and perturb some functionalities by wrong syntaxes.

  The problem was temporally fixed until we change the ITBox version from 1.3.4.0 to 2.3.4.1.

- Instabilities and errors coming from undetectable error during OpenStack deployment:

During OpenStack deployment no errors appeared and all deployments were finished successfully.

After many issues and problems encountered, SophiaAntipolis discovered after investigating, that some OpenStack services were missing or in error status.

All issues were manually fixed.

On the node, there is currently a problem about the volume creation/attachment, fixing ongoing.

4.2.9 Status of iMinds node

Node Name	OpenStack release	Overall status of the node	Validation tests results
Gent	Grizzly	Operational	11 Errors/ 5 Passed

Table 15: Status of iMinds node

The tests show that there was a communication problem with the keystone proxy, anyway the node, in the following week (after 13/03/2015), had failed only the tests related to missing of floating IP resources.

4.2.9.1 Description of the node status

OpenStack Grizzly version is deployed with OS Ubuntu 12.04LTS from Ubuntu Cloud Archive repository.

iMinds Gent node is composed by:

- 8 physical servers.
- 128 CPU cores.
- 384Gb Ram.
- 8Tb Storage.

Infrastructure network is connected to the MD-VPN, currently used for monitoring purposes and user VMs, and 2 public IPv4 pools (one for Federation control services, other to provide floating public IP's to users).

Also, we have dedicated 3 physical servers with OpenStack Juno version, for testing purposes with the objective to upgrade the whole node to that running version when all the Federation tests finish and all compatibility issues were solved.

Every day the node is supervised and services status is continuously checked by monitoring scripts. The allocated floating IPs, that are not used within a week, are de-allocated to free them to other users.

The main problem is that all users ask for a Public Floating IP and they are not enough for everybody. The problem is expected to be solved with the next OpenStack version Juno that has IPv6 support.

4.2.9.2 History of the node availability

The node is up and running since 14/10/2014, and it's operational since then. No downtime has been noticed.

Monitoring scripts are used to supervise the operation of the node services.

4.2.10 Status of Acreo node

Node Name	OpenStack release	Overall status of the node	Validation tests results
Stockholm	Grizzly	Up and running	0 Errors/ 16 Passed

Table 16: Status of Stockholm node

4.2.10.1 Description of the node status

The Stockholm Node runs OpenStack Grizzly. It was deployed with help of ITBox 1.3.4.0. The OpenStack environment is composed by the following components:

- A Controller running in a virtual machine
- Cinder service running on its own virtual machine
- A monitoring machine on its own in a virtual machine
- A ContextBroker on its own in a virtual machine
- Three compute nodes (1 compute node comprises: 32 cores, 128Gb RAM, 2TB), one of which is used in the current Grizzly deployment while the other two are used for the undergoing upgrade to Icehouse.

An allocation of 256 public IPv4 addresses is used by the node.

The node is used extensively by users creating one time test environments that are very often left unused after a short period. This exhausts our IPv4 pool rapidly, preventing the allocating of floating IPs and allocating IPs for the routers. However, deleting resources is a difficult exercise since we do not know which users they belong to (users are anonymous to us managing the node).

Statistics of the node usage over the month passed:

Usage from 2015-02-03 to 2015-03-04:

Tenant ID	Instances	RAM MB-Hours	CPU Hours	Disk GB-Hours
00000000000000000000000000000135	2	1720320.11	1344.00	13440.00
00000000000000000000000000003807	1	237345.36	463.57	0.00
00000000000000000000000000004979	1	11010048.71	5376.00	107520.01
00000000000000000000000000005584	4	2918384.87	1789.51	26069.75
00000000000000000000000000005768	7	4523788.47	2885.02	39670.06
00000000000000000000000000007161	2	8257536.53	4032.00	80640.01
00000000000000000000000000007827	4	2614040.10	1276.39	25527.74
00000000000000000000000000008392	1	5505024.36	2688.00	53760.00
00000000000000000000000000008715	6	1041707.57	624.30	9401.91
00000000000000000000000000008822	3	1162456.60	567.61	11352.12
00000000000000000000000000009785	1	2752512.18	1344.00	26880.00
00000000000000000000000000009985	1	344064.02	672.00	0.00
00000000000000000000000000010079	1	1376256.09	672.00	13440.00
00000000000000000000000000010384	3	61156.46	29.86	597.23
00000000000000000000000000010684	1	334054.06	652.45	0.00
00000000000000000000000000011233	65	263.96	0.13	2.58
00000000000000000000000000011283	1	51865.94	25.33	506.50
00000000000000000000000000011432	2	1114111.54	544.00	10880.00
00000000000000000000000000011660	6	1317581.94	643.35	12867.01

Figure 7: Stockholm Node Statistics

4.2.10.2 History of the node availability

The history of the node has had the following main events:

- We reached full integration including status monitoring in the Infographics Status Page on 20/10/2014.
- We have had intermittent service unavailability due to resource depletion since our node has become online. A long term solution to this problem requires better control of which users have access to the resources. The XIFI project has proposed to segment users into two categories:
    - Users that are registered, have support from the node and access to resources that are agreed upon;
    - Test users, which may access whatever resources are still available, with resource allocation lifetime of two weeks.
- 11/2/2015: Our site suffered from a general power failure that took our node down and required several hours to recover.
- 3/3/2015: We had a maintenance reboot to update service account names and passwords.

4.2.11 Status of BTH node

Node Name	OpenStack release	Overall status of the node	Validation tests results
Karlskrona	Grizzly	Operational	3 Errors/13 Passed

Table 17: Status of Karlskrona node

Failed tests are related to missing of floating IP resources.

4.2.11.1 Description of the node status

The node is currently upgrading to Icehouse. The node would like to upgrade to Juno awaiting confirmation that GEs required are available for Juno release. Migration to Icehouse currently consumes a number (50%) of CPUs taken from the computing node pool. The Karlskrona node is accessible through the portal. Node information is available through the Infographics page in the portal. The node is currently passing daily checks. Temporary failures are observed during component update periods and when the floating IP pool is exhausted which can happen frequently during busy periods given the pool limits in relation to the number of active users. The Karlskrona node has access to a complete C network, and intends to grow its floating IP.

Currently the Karlskrona node is working with the following resources:

- 64 cores
- 256Gb ram
- 2TB disk storage

The same amount of resources are allocated to the Icehouse upgrade. Integration into Network Test Automation System is under way.

4.2.11.2 History of the node availability

The node was up and running beginning mid-October 2014, and has been operational since then.

No downtime has been noticed.

4.2.12 Status Berlin node

Node Name	OpenStack release	Overall status of the node	Validation tests results
Berlin	Grizzly	Up and running	0 Errors/16 Passed

Table 18: Status of Berlin node

4.2.12.1 Description of the node status

- The node is currently upgrading to Icehouse. The node would like to upgrade to Juno awaiting confirmation that GEs required are available for Juno release. Migration to Icehouse currently consumes a number of CPUs taken from the computing node pool.
- The Berlin node is accessible through the portal. Node information is available through the Infographics page in the portal.

  The node is currently passing daily checks. Temporary failures are observed during component update periods and when the floating IP pool is exhausted which can happen frequently during busy periods given the pool limits in relation to the number of active users.

  The Berlin nodes aims to double its floating IP pool size going into effect with the migration of the OpenStack release and the new user policies, which allows to install dedicated pools for trial and community users.
- The Berlin node currently provides (nominal, i.e. including the resources currently assigned to the migration effort):
  - 96 Cores
  - 384 GB Ram
  - 3TByte redundant Disk storage
  - an out-door wireless access network.

4.2.12.2 History of the node availability

- The node was up and running beginning mid-March 2014. The Berlin node was demonstrating first showcases in November 2013 (not fully federated) and was hosting first training sessions in May 2014 (4 federated nodes).
- Longer down-times have not been observed (max. Saturday/Sunday) but access though the portal was limited rather often - usually in consequence of portal and keystone issues. Meanwhile cloud portal and Identity management can be considered quite stable and noticeable down-times usually take place only in consequence of OpenStack issues.

4.2.13 Status of Lannion node

Node Name	OpenStack release	Overall status of the node	Validation tests results
Lannion	Grizzly	Up and Running	0 Errors/16 Passed

Table 19: Status of Lannion node

4.2.13.1 Description of the node status

Currently the production Lannion node is running on Grizzly. A new node is being installed with Fuel 6.0 and is running on Juno. Once the new node will be federated and fully tested, we will migrate to it the data of the actual production node and then the hardware will be also migrated.

Actual Hardware resources assigned to the Production node are:

- 96 Cores
- 256 GB Ram
- 32 TB Disks
- /26 Subnet
- After the migration to Juno, the resources assigned to the Lannion node will be:
- 224 Cores
- 768GB Ram
- 64TB Disks
- /24 + /26 Subnet

Usage from 2015-02-02 to 2015-03-03:

Tenant ID	Instances	RAM MB-Hours	CPU Hours	Disk GB-Hours
00000000000000000000000000000081	1	1376256.70	672.00	13440.01
00000000000000000000000000000084	1	496319.05	242.34	4846.87
00000000000000000000000000001173	1	2752513.40	1344.00	26880.01
00000000000000000000000000002553	1	951327.81	464.52	9290.31
00000000000000000000000000002900	2	69.26	0.14	0.00
00000000000000000000000000002983	8	12730374.47	6720.00	120960.06
00000000000000000000000000002988	42	770992.59	374.61	7378.10
00000000000000000000000000003005	3	1290334.39	630.05	12600.92
00000000000000000000000000003199	1	1376256.70	672.00	13440.01
00000000000000000000000000003273	126	1203703.94	587.75	11754.92
00000000000000000000000000003437	5	11010053.60	5376.00	107520.05
00000000000000000000000000003851	1	2752513.40	1344.00	26880.01
00000000000000000000000000003997	3	6193155.15	4032.00	0.00
00000000000000000000000000004004	1	5505026.80	2688.00	67200.03
00000000000000000000000000004017	2	2752513.40	1344.00	26880.01
00000000000000000000000000004259	1	1708.94	3.34	0.00
00000000000000000000000000004511	3	3096577.57	2016.00	26880.01
00000000000000000000000000005416	1	344064.17	672.00	0.00
00000000000000000000000000005854	2	21403.01	41.80	0.00
00000000000000000000000000005964	1	2752513.40	1344.00	26880.01
00000000000000000000000000007281	1	1376256.70	672.00	13440.01
00000000000000000000000000007572	2	638716.38	1247.49	0.00
00000000000000000000000000007573	1	2752513.40	1344.00	26880.01
00000000000000000000000000007584	1	2752513.40	1344.00	26880.01
00000000000000000000000000007621	1	2752513.40	1344.00	26880.01
00000000000000000000000000007787	2	2752513.40	1344.00	26880.01
00000000000000000000000000007968	1	2752513.40	1344.00	26880.01
00000000000000000000000000008201	1	1376256.70	672.00	13440.01
00000000000000000000000000008576	32	25896867.53	12146.88	232891.69
00000000000000000000000000008916	1	2752513.40	1344.00	26880.01
00000000000000000000000000008950	1	2752513.40	1344.00	26880.01
00000000000000000000000000008978	1	47612.49	23.25	464.97
00000000000000000000000000009784	2	2752513.40	1344.00	26880.01
00000000000000000000000000010099	1	1376256.70	672.00	13440.01
00000000000000000000000000010328	1	8612.68	16.82	0.00
00000000000000000000000000011233	45	134.26	0.07	1.31
00000000000000000000000000011586	1	48834.26	95.38	0.00
00000000000000000000000000011595	3	33232.17	16.23	324.53

Figure 8: Lannion Node Statistics

4.2.13.2 History of the node availability

- The Lannion was publicly available since middle of March 2014
- We had faced 2 mains instability issues:
  - From time to time, we lost the connectivity to the instances due to the instability of Quantum processes in High Availability:
  - When: From mid of March to September 2014
  - Service affecting: Yes, time to time experienced some loss of connectivity to instances
  - Why: Quantum processes were badly managed in High Availability. Mirantis scripts in this release were quite new
  - Solution: Bugs where founds in Mirantis scripts that are managing the monitoring status of quantum processes handled by Pacemaker HA service. A lot of bugs were found in this scripts but mostly it failed to find the right PID. We improve the commands that were in this scripts and it fixed our instability.
- From time to time, our node got Over flooded of MySQL database and the requests to our node hang up
  - When: From June to October 2014
  - Service affecting: Yes, when it happened, the node did not respond to requests.
  - Why: From time to time, it happened that the MySQL Database was over flooded by requests and then cannot answer anymore.
  - Solution: After some search, we found that the problem was linked with the operations done on the node from the Cloud Portal, as most of the time, service were blocked when we were connected to the cloud portal. After some deeper search, we saw that some of our accounts got Cloud Portal "admin" privilege checked on the Fiware Lab Account Manager. When performing Neutron related requests as administrator causes heavy load on the database, which end up locking the access to the database and preventing other service from running properly. This behaviour is described in a bug tracked by the OpenStack community (https://review.openstack.org/#/c/47651/). It has been resolved by unchecking the "admin" privilege in the account of Fiware Lab portal.
- Crash of one of the 2 controllers:
  - When: from 09/06/2014 to 15/06/2014
  - Service affecting: Yes
  - Why: The crash of one of the disk in one of our 2 controllers, crash the server itself. Pacemaker, in charge of handling the cluster, did not appreciate well the loss of one of the 2 controllers and keep on rebooting the processes that it was handling.
  - Solution: By removing the crashed node from pacemaker configuration, it stabilised the node. The node have been fully restored using a backup image of the node.
- Node maintenance:
  - When: 27/02/2015 at morning
  - Service affecting: Yes, some disruptions to the instances
  - Why: Modification of the credentials to the keystone proxy

4.2.14 Status of Spain node

Node Name	OpenStack release	Overall status of the node	Validation tests results
Spain	Essex	Operational	6 Errors/12 Passed

Table 20: Status of Spain node

Spain Node run out of resources. There is a migration ongoing on a new node spain2.

4.2.14.1 Description of the node status

Currently, the Spain node is working with the following resources:

- 544 cores
- 2176 GB or RAM
- 17 Tb for disk instances + 10 Tb for base instances
- 1 Tb for nova volume
- 0.6 Tb for object storage (swift)
- 1280 Floating IPs

Regarding the information of each node the available statistics about the usage of such resources are the following:

Host	CPU Total	CPU Used now	CPU Used max	RAM Total	RAM Used now	RAM Used max	HDD Total	HDD Used now	HDD Used max
aytoma-host02-1	32	201	203	128903	122096	316416	909	570	4980

Host	CPU Total	CPU Used now	CPU Used max	RAM Total	RAM Used now	RAM Used max	HDD Total	HDD Used now	HDD Used max
aytoma-host02-2	32	208	208	128903	123821	326144	909	517	5000
aytoma-host02-3	32	204	207	128903	121885	321024	909	480	5100
aytoma-host02-4	32	203	206	128903	127462	290816	909	542	4610
aytoma-host02-5	32	198	202	128903	122490	302592	909	493	4850
aytoma-host02-6	32	192	198	128903	119549	302592	909	474	4620
aytoma-host02-7	32	199	199	128903	95614	273408	909	530	4450
aytoma-host02-8	32	203	212	128903	125286	325120	909	612	5280
aytoma-host02-9	32	147	149	128903	109178	256006	909	405	3690
aytoma-host02-10	32	180	180	128903	107459	255488	909	413	4100
aytoma-host02-11	32	195	196	128903	119637	301568	909	534	4690
aytoma-host02-12	32	161	162	128903	106561	245760	909	432	3900
aytoma-host02-13	32	197	199	128903	115471	276480	909	606	4400
aytoma-host01-3	32	197	200	128903	123640	288768	10239	9428	4530
aytose-host02-1	32	37	43	128904	15803	44554	912	515	740
aytose-host02-2	32	112	123	128904	59888	147480	912	575	2370

Host	CPU Total	CPU Used now	CPU Used max	RAM Total	RAM Used now	RAM Used max	HDD Total	HDD Used now	HDD Used max
aytose-host02-3	32	117	123	128904	64059	151060	912	561	2320
aytose-host02-4	32	122	122	128904	38122	159250	912	584	2430

Table 21: Available Statistics of usage hosts in Spain node

The resume of those data can be found in the following table:

CPU Total	CPU Used now	CPU Used max	RAM Total	RAM Used now	RAM Used max	HDD Total	HDD Used now	HDD Used max
576	3073	3132	2320258	1818021	4584526	25704	18271	72060

Table 22: Total resume of usage statistics in the Spain node

4.2.14.2 History of the node availability

- The node has been publicly available since the 1/9/2013 with some availability issues due to a previous version of the system and lack of resources, hardware with poor performance and network connectivity. Previously to it, we have a FIWARE Testbed since December 2012 used for the test of the development of FIWARE (currently also working).
- Downtime on 10/12/2013 due to a Santander datacenter flood without great consequences but we needed to switch off the hosts.
- The Spanish node was deployed in a small datacenter in Santander (Spain) and it was migrated one year ago (September 2014) to Seville (Spain), this represented some downtimes.
- Downtime from 17/10/2014 12:54 to 16:55 due to a fiber cut.
- Downtime from 30/12/2014 09:20 to 30/12/2014 14:44 due to a fiber cut.
- From December, 2014 to January, 2015 it was not possible to deploy new resources in Spain due to the lack of computational resources. As shown previously in the statistics, the resources are really oversubscribed -- It was not possible to assign new resources. However, all the resources previously assigned did not stop working.
- Downtime from 08/01/2015 07:00 to 08/01/2015 08:35 due to a fiber cut.
- By the end of January 2015, and looking forward to an important hosts movement from one datacenter to another which would mean a 2 days unavailability of the Spanish node, we had to move again the resources from the Seville's datacenter to the Malaga's one. This meant some minor downtimes.
- There was a downtime from 21/02/2015 22:00 to 22/02/2015 10:30 due to a hacker attack.
- There was a micro downtime from 3/03/2015 9:00 to 5/03/2015 17:00 due to management task related to poor improvement of Malaga fiber link.

4.2.15 Status of Trento node

Node Name	OpenStack release	Overall status of the node	Validation tests results
Trento	Grizzly	Operational	1 Errors/ 15 Passed

Table 23: Status of Trento node

Failed tests are due to missing floating IP resources.

4.2.15.1 Description of the node status

The Trento Node was operational on OpenStack Grizzly deployed on Ubuntu 12.04 LTS, by use of Mirantis Fuel 3.2.1 version, from early January 2014 until 9/3/2015.

Trento Node was composed by 7 Dell R210 and 6 Dell R715. Controllers, Monitoring and Object Storage were Dell R210:

The Trento Node was configured as follows:

- 3 Controller nodes (HA mode)
- 6 Compute nodes
- 3 Object Storage nodes
- 1 Fuel 3.2.1 node

and the hardware resources assigned to the node were:

- 120 cores
- 480 GB of RAM
- 27 TB of Disk

From 9/3/2015, the Trento Node is hosting OpenStack Icehouse version, deployed with Fuel 5.1.1 release. The node is deployed in HA-mode and it uses Ceph RBD for images & volumes, with replication factor equal to 3. The node is accessible through the portal and the information about the node can be consulted through the Infographics page. The node is currently composed by 3 controllers and 2 computes and the hardware resources assigned are:

- 76 cores
- 176 GB of RAM
- 9,2 TB of Disk

The hardware integration with the previously used hardware is ongoing. The final configuration will have 3 controller/Ceph-osd nodes, 3 Ceph-osd nodes and 8 compute nodes, besides the monitoring dedicated node.

The node usage for February 2015 as follows:

Usage from 2015-02-01 to 2015-03-01:

Tenant ID	Instances	RAM MB-Hours	CPU Hours	Disk GB-Hours
00000000000000000000000000000081	5	704611.56	688.10	6880.97
00000000000000000000000000000125	1	688128.00	672.00	6720.00
00000000000000000000000000000135	1	344064.00	672.00	0.00
00000000000000000000000000000348	1	239.50	0.23	2.34
00000000000000000000000000000442	3	3440640.00	3360.00	47040.00
00000000000000000000000000000495	1	302608.64	591.03	0.00
00000000000000000000000000000852	1	1376256.00	1344.00	20160.00
00000000000000000000000000001489	1	344064.00	672.00	0.00
00000000000000000000000000002682	1	344064.00	672.00	0.00
00000000000000000000000000002782	12	12730484.05	9408.11	161282.83
00000000000000000000000000002798	2	4128768.00	2016.00	53760.00
00000000000000000000000000002862	2	27525120.00	10752.00	174720.00
00000000000000000000000000002982	1	6.26	0.01	0.00
00000000000000000000000000003327	2	688128.00	1344.00	0.00
00000000000000000000000000003373	4	2752512.00	2688.00	26880.00
00000000000000000000000000003478	1	5505024.00	2688.00	107520.00
00000000000000000000000000003495	1	688128.00	672.00	6720.00
00000000000000000000000000003954	5	9889103.64	4828.66	153763.65
00000000000000000000000000004189	1	2185.96	4.27	106.74
00000000000000000000000000004213	1	688128.00	672.00	6720.00
00000000000000000000000000004219	1	688128.00	1344.00	33600.00
00000000000000000000000000004242	1	688128.00	672.00	6720.00
00000000000000000000000000004259	9	53.90	0.06	0.47
00000000000000000000000000004504	1	688128.00	1344.00	33600.00
00000000000000000000000000004520	1	8678.97	8.48	211.89
00000000000000000000000000004835	1	344064.00	672.00	0.00
00000000000000000000000000006104	1	1376256.00	672.00	13440.00
00000000000000000000000000006435	1	1376256.00	1344.00	33600.00
00000000000000000000000000006858	1	688128.00	1344.00	33600.00
00000000000000000000000000006916	1	344064.00	672.00	0.00
00000000000000000000000000007573	1	1376256.00	1344.00	33600.00
00000000000000000000000000007774	2	1376256.00	1344.00	13440.00
00000000000000000000000000007827	1	155355.31	303.43	0.00
00000000000000000000000000007968	1	688128.00	672.00	6720.00
00000000000000000000000000008201	1	688128.00	672.00	6720.00
00000000000000000000000000008576	1	62.29	0.12	0.00
00000000000000000000000000008642	1	688128.00	1344.00	33600.00
00000000000000000000000000008914	1	688128.00	672.00	6720.00
00000000000000000000000000008916	1	688128.00	1344.00	33600.00
00000000000000000000000000008950	2	1376256.00	1344.00	13440.00
00000000000000000000000000009094	2	688128.00	1344.00	0.00
00000000000000000000000000009353	1	688128.00	672.00	6720.00
00000000000000000000000000009451	2	1032192.00	1344.00	6720.00
00000000000000000000000000009534	2	1032192.00	1344.00	6720.00
00000000000000000000000000009570	1	344064.00	672.00	0.00
00000000000000000000000000009634	1	688128.00	672.00	6720.00
00000000000000000000000000009859	2	1032192.00	1344.00	6720.00
00000000000000000000000000010086	1	688128.00	672.00	6720.00
00000000000000000000000000010150	2	1376256.00	1344.00	13440.00
00000000000000000000000000010153	1	107355.73	209.68	0.00
00000000000000000000000000010274	1	18404.98	35.95	0.00
00000000000000000000000000010374	1	688128.00	672.00	6720.00
00000000000000000000000000010440	1	688128.00	672.00	6720.00
00000000000000000000000000010651	1	478732.80	467.51	4675.12
00000000000000000000000000010672	1	1296085.90	1265.71	18985.63
00000000000000000000000000011110	1	30974.86	60.50	0.00
00000000000000000000000000011233	19	89.32	0.09	0.87
00000000000000000000000000011266	4	373939.20	405.32	5238.37
00000000000000000000000000011407	1	44.23	0.09	0.00
00000000000000000000000000011586	1	28759.18	56.17	0.00
00000000000000000000000000011661	3	20910.36	20.42	204.20
dbb1d0ef27704663b2336e60cafbb8db	2	5505024.00	2688.00	53760.00

Figure 9: Statistics of usage in Trento node

Trento dedicates also 3 servers Fuji-Siemens RX200 S3 to provide a Test node. Main objective of the Test Node is the Generic Enablers testing from the partners and also for the IO update tests for different components.

About the Floating IPs issue, it is planned to add a new pool of Floating IPs. For security matters, the IO will follow a set of procedures regarding the identity verification of the user before issuing a Floating IP.

4.2.15.2 History of the node availability

Trento Node is up and running on OpenStack Grizzly since beginning of January 2014.

- 19-20/6/2014 → Additional Reconfiguration of the Node
- 2/7/2014 → Second External Network Configuration
- 23-24/7/2014 → time-out issue connecting to the Keystone-Proxy from Trento Node.

- 29/7/2014 9.30 am CEST - 12:30 CEST → Restore of Mysql Database Corrupted
- 4-5/8/2014 → Instable connection to the cloud portal (other nodes experienced the same issue)
- 23-24/2/2015 → Security Issue Maintenance
- 3-9/3/2015 → Migration to Icehouse

4.2.16 Status of Waterford node

Node Name	OpenStack release	Overall status of the node	Validation tests results
Waterford	Grizzly	Operational	3 Errors/13 Passed

Table 24: Status of Waterford node

Failed tests are due to missing floating IP resources.

4.2.16.1 Description of the node status

The node is currently running on OpenStack Grizzly. We currently have a second system running in parallel, where are testing the migration to fuel 6 and ITbox 2 with Ceph running Glance, Cinder and Swift components. This second instance is current on loan for the duration of the migration. The Waterford node operational state can be observed in two main ways: one being real time via infographic tools and the second runs daily against API hosted on Waterford XIFI controller. Currently the operational state of these checks shown via these two mediums is up. On occasions we notice we were failing on floating IP allocation. This is due to IP pool being fully allocated. During the migration we plan to add another class C to our servicing offering.

Hardware resources the Waterford node current has assigned:

- 112 compute cores, with 364 GBytes ram and 1.5 TBytes of disk.
- 16 controllers cores, with 128 GBytes ram and 1 TBytes of disk.
- 16 NFS server cores, with 64 GBytes ram and 1 TBytes of disk.
- 6 ITBox cores, with 8 GBytes ram and 1 TBytes of disk.
- 16 motoring cores, with 8 GBytes ram and 300 GBytes of disk.

For migration (on loan for migration)

- 32 compute cores, with 128 GBytes ram and 600 GBytes of disk.
- 16 controllers cores with 128 GBytes ram and 1 TBytes of disk.
- 24 Ceph cores server with 64 GBytes ram and 2 TBytes of disk.
- 6 ITBox cores, with 8 GBytes ram and 1TBytes of disk

Network resources:

- IPv4 address supplied by HEAnet, 384 IPv4 host addresses.

4.2.16.2 History of the node availability

The Waterford node went through two maintenance cycles, the first cycle need a re-install of the node the second needed reboot of Waterford OpenStack controller service.

Waterford IT maintenance cycle 1 (Node outage)

When: This happened during July 2014 and lasted for 1 and half weeks.

Service effecting: yes

Why: Because we were an early adopter of OpenStack and the ITBox installation tool was not released, we decided to install OpenStack grizzly manually. This initially worked well and was in place for about 10 months. We then, on the root Operating System (OS) of each node hosted, executed a software update "apt-get update". This action proved crucial as Open Virtual Switch software got updated and then stopped functioning. The issue was: OVS plugin agent on the compute node issued all the commands to Open Virtual Switch, and in turn Open Virtual Switch never created the interface in the compute nodes namespace. Result being, VMs hosted on the compute node could not connect to under lying network interface.

How it was resolved: After days of testing and reconfiguring, we were unable to find a working solution to the issue. We then decided to port running Virtual Machines (VM) to a new OpenStack install based on ITBox. This initially did not work as ITBox put incorrect labels on disk slices, which had the result that when the host OS tried to use swap the compute server failed instantly. We found and implemented a work around for the label swap issue, then proceeded with the install. We ported the VM to the new system and power them on. It is this install base we are currently running our platform on today.

Waterford IT maintenance cycle 2 (Admin account failure)

When: This happened during September 2014 and lasted for 3 days.

Service effecting: limited

Why: The admin tenant account we used to deploy configuration changes with stopped functioning on our local node. This caused difficulties in quantum as quantum's external networks are built using the tenantID. Slowly external networks started to deteriorate, and got to a point where floating IP's failed.

How it was resolved: Rebooted OpenStack controller. This issue happened again in December 2014 but we now have a work around in place.

4.3 Federation Nodes Journal

This paragraph is a report of the deployment activity for the new nodes. In terms of lifecycle a node life can be partitioned in three steps:

- pre-deployment: the new node has been chosen as a result of the new call but it has not already started the deployment, in this phase the Infrastructure Owner of the new node will have to manage to satisfy the constraints and the requirement to start the deployment (e.g. connectivity to the backbone, Hardware procurement,...).
- deployment: in this phase the node is installing and configuring the platform (connectivity, Hw, Sw) to federate to node in XIFI.
- support: the node has been successfully deployed and is going in the production phase.

The period of interest of this report is the pre-deployment and deployment phase. This Chapter is divided in geographical areas as defined above:

- South-East Europe: Italy, Austria, Hungary, Serbia, Croatia, Bosnia, Slovenia, Slovakia, Montenegro, Albania, Macedonia, Bulgaria, Romania, Greece and Turkey
- Central-East Europe: Germany, Czech Republic, Denmark, Poland, Switzerland, Latvia, Estonia and Lithuania
- South-West Europe: Spain and Portugal
- Central-West Europe: France, Belgium, Netherlands and Luxemburg
- North Europe: Ireland, UK, Norway, Sweden and Finland

4.3.1 Geographical Area South-East Europe

The South-East Europe area is composed by the following new nodes:

- NeuroPublic
- UPRC
- UTH
- Wigner

and it is managed by Telecom Italia as task5.5 partner.

4.3.1.1 NeuroPublic

On 2/5/2014 NeuroPublic was contacted by email and was provided with technical XIFI documentation and references. In the same mail it was asked to provide the deployment plan. On the 14/05/2014 NeuroPublic provided the plan. With respect to the XIFI core backbone connectivity, as a temporary solution, they had been talking with DANTE to establish an L3VPN tunnel over Internet with another XIFI node already part of the XIFI Federation.

On 3/7/2014, during the CC of the SE Europe Area nodes, NeuroPublic said that will use the backup connectivity solution from Geant. To discuss that, a dedicated call was organized with DANTE and all the Greek partners for 10/7/2014.

During the call held on 10/7/2014, about the connectivity status, NeuroPublic was in talks with ISP & GRNet investigating the possibility of a VPN between NeuroPublic and GRNet network. NeuroPublic informed Michael Enrico (DANTE) about the progress and said that they would give update as soon as they have more info. About hardware status, NeuroPublic is still in talks with suppliers. They think to be able to have the hardware by the end of July.

On 13/8/2014, NeuroPublic said they were in talks with their ISP and GRNET in order to implement a direct 300Mbps VPN link from NeuroPublic to GRNet and connect to XIFI. The hardware delivery for both the servers and the OpenFlow switch was expected around beginning of September.

On 27/8/2014, about the connectivity, NeuroPublic said that they were still waiting for the official proposal from their ISP for the 300Mbps VPN link from NeuroPublic to GRNET. About the HW, they had received the servers (on 25/8/2014); the OpenFlow switch was supposed be delivered by the first week of September.

On 3/9/2014, NeuroPublic said that they received the official proposal from their ISP for the VPN to GRNET. They scheduled to start Cloud installation and the other activities after having received the OpenFlow switch (expected for September 3rd).

On 10/9/2014, NeuroPublic said they signed the official proposal from their ISP in order to implement a direct 300Mbps VPN link from NeuroPublic to GRNet and connect to XIFi. They are waiting for update on circuit delivery date from ISP. Furthermore, they added that the all hardware was deployed (OpenFlow switch included). They installed OpenStack using ITBox 1.3.4.0 . About the Monitoring, they used ITBox 1.3.4.0 to install that, but it needed to be properly configured (they had pending issues and they were discussing them with Alessandro Martellone).

On 17/9/2014, NeuroPublic said that the Monitoring installation was going ahead: they had also opened a Jira ticket to be able to solve the previously reported issue about the monitoring configuration.

On 1/10/2014, NeuroPublic said that they were connected to the Geant network and now they can ping the other nodes. About Federation (keystone proxy), they sent an email to Alvaro Alonso (UPM) with their IPs.

On 7/10/2014, NeuroPublic said that they had completed the Monitoring installation (Nagios and Nrpe components were manually installed). Nevertheless, in the Infographic web page, the Nagios "buttons" were greyed out. They received some help from Attilio Broglio (Create Net) that they had to add to nagios.cfg file:

- broker_module=/usr/local/nagios/lib/ngsi_event_broker_xifi.so -r <REGIONID> -u http://<IP_NGSI_ADAPTER>:1337

NeuroPublic found out that this broker module was not built by ITBox. So, they built it by themselves and they added the configuration files to Nagios. Now, their Nagios log reads:

- [1412755454] ngsi_event_broker_xifi - Request sent to http://10.0.80.5:1337/check_nova_cert?id=Athens_Neuropublic:10.60.2.4&type=vm
- [1412755454] ngsi_event_broker_xifi - Request sent to http://10.0.80.5:1337/check_local_disk?id=Athens_Neuropublic:10.60.2.5&type=host
- [1412755484] ngsi_event_broker_xifi - Request sent to http://10.0.80.5:1337/check_swap?id=Athens_Neuropublic:10.60.2.3&type=vm

Still there was no data in their Infographic and a test sent from Attilio Broglio using "./xifi_query_cb_monitored_objects.sh -b 10.0.80.5:1026 -t host_service" gave them a 404 error. They told that they would try to resolve this issue as soon as they can. Lastly, NeuroPublic completed the Federation activity (keystone proxy): in cloud portal they can create VMs, assign floating IP, access VM's from outside.

On 14/10/2014, NeuroPublic solved the Infographic web page issue: now all services are green.

On 12/10/2014, NeuroPublic finished the migration to Icehouse OpenStack release. They completed successfully all checks: the node is up and running.

4.3.1.2 UPRC

On 2/5/2014 UPRC was contacted by email and was provided with technical XIFI documentation and references. In the same mail it was asked to provide the deployment plan. On the 12th May UPRC provided the plan, now available in the wiki. About the connectivity, they wrote that GRNET should provide the MD-VPN service on early September.

During the call held on 10/7/2014, UPRC said that they were in the process of obtaining the hardware. About connectivity, they were in communication with GRNET regarding the alternative solution.

In the meeting on 16/7/2014, UPRC said that, about the connectivity, the backup solution would be exploited for their connectivity with Geant by the end of August. UPRC and GRNET are sending information about the VLANs. About the HW, because of delay due to vacation period, HW procurement was to be achieved in late August.

On 24/7/2014, UPRC said that UPRC Noc communicated with GRNET and sent the necessary VLAN info. The procedure for the connectivity to XIFI core backbone has been started.

during the call on 13/8/2014, UPRC said they had received from GRNET the cpe switch that will be used for the connection with XIFI cloud . About the hardware, because of delay due to vacation period, HW procurement was to be achieved in early September. All other activities were scheduled for September.

On 27/8/2014, UPRC confirmed that the connectivity with GEANT was to be accomplished during the first week of September.

On 3/9/2014, UPRC confirmed that the connectivity to GEANT was going to be accomplished before the end of current week. The HW procurement to be completed within September. They said that progress was being made on Cloud Installation and Monitoring, and further significant progress was expected by the forthcoming meeting scheduled in September.

On 10/9/2014, UPRC said that, about connectivity to GEANT, the CPE (Customer Premises Equipment) switch was installed, the network circuit among GRNET and UPRC is ready and the MD-VPN functionality is available. Testing MD-VPN connection was to be completed in parallel with OpenStack installation. About Cloud Infrastructure, currently UPRC was finalizing a bare metal ITBox installation. Some CentOS and Mirantis Fuel configuration activities should be finalized. ITBox as well as some slave nodes have already been tested in a standalone server and they worked correctly. By the end of this week (12/09), barring unforeseen, the overall server infrastructure, including switches, was to be optimized and will work. In the current architecture plan there are 5 servers (1 ITBox Master Node and 4 Slave Nodes that are distributed as Compute, Controller and Storage Nodes), while this plan was to be extended after HW equipment procurement, with 3 or 4 additional servers.

On 17/9/2014, about the connectivity to GEANT, UPRC said that they had configured the IP ranges of the demanded VLANs on the router and the respective switches. They were collaborating with GRNET for testing of MD-VPN connection. The process was supposed to be completed by the end of the week. About the Hardware, they were working with company ACTIVE, to installing/testing hardware components, and integrating with software parts, in order to deliver the node. In the meantime, ITBox was installed and fully configured. Three bare metal installations of OpenStack on the first 3 nodes that are separated as Compute, Controller+Monitoring, Cinder. In addition, that week was to be added 1 additional server that belongs to the new HW, whereas 3 new servers, provided by ACTIVE company, were going to be added in the node by the end of the following week. Further to that, they were going to create Trunk for the VLANs, because their servers have 2 NICs, and during this week they will test the performance of this approach, so as to achieve the best possible network configuration. UPRC will participate in the remote session that will be performed on 19/09 about the keystone proxy installation/configuration, so as to extract valuable information that will help to a faster integration of the node with the keystone and by extension will lead to the successful federation of the node.

On 1/10/2014, UPRC confirmed that the connectivity to GEANT is ok. About the infrastructure, they has finalized the cloud installation and management achievement by using the ITBox 1.3.4 version. They have 1 Physical server hosts (a bare metal) of the ITBox, while 4 different types of networks were setup and configured for the smooth operation of the infrastructure. Currently 2 additional physical servers works in the infrastructure:

- 1 server that hosts the Controller and Monitoring components (including the GEs: ContextBroker, NGSI Adapter and ODC modules)
- 1 server that hosts a Compute and Cinder node fully available and tested by OpenStack dashboard environment.

They created some instances so as to ensure that everything works correctly, as well as to check the monitoring functionalities of the Nagios module hosted in the monitoring server. Some bugs in the installation of the OpenStack modules on the servers for the monitoring were fixed manually, so as to solve some minor problems in the monitoring operation. UPRC could share relevant instruction, based on its experience so far, for the debugging in Ubuntu 12.04 environment for the monitoring nodes, by aiming to help other nodes to accelerate their installation of the monitoring module. The monitoring module was successfully installed in a server, and Nagios components were tested and work correctly by presenting the available services per node. In addition, the NRPE Nagios modules in the Compute and Cinder node, were tested and they work without problems. Consequently, the monitoring installation, configuration and operation achievement were successfully performed. About Federation (keystone proxy), UPRC was working on the development of the external network by using the OpenStack Customization XIFI guide. There had some issues with the interconnectivity from outside that were going to be elaborated so as to be solved immediately, before next Friday 3/10.

On 8/10/2014, UPRC almost completed the Federation (keystone proxy) activity: they achieved the VMs creation and floating IPs assignment. They had to fix only an issue about the external VM access.

On 13/10/2014, UPRC finished the Federation activity.

During the period By October 2014 to the first half of January 2015 (15/01/2015) the PiraeusU node was Up and Running.

The period between 15/01/2015 to 21/01/2015 of January it was Down because it was in upgrade process from OS Grizzly to OS Icehouse release. PiraeusU technical team used the Mirantis FUEL 5.1.1 release for the installation and the setup of the OpenStack Icehouse on the node, with the following basic configuration:

- Hypervisor: KVM.
- Storage nodes: Cinder LVM over iSCSI for volumes.
- Networking: Neutron agent with OVS VLAN splinters hard trunks workaround enabled.

Before the node federation, they tested its operations, by exploiting the OpenStack CLI API, as well as the capabilities provided by OpenStack Horizon dashboard Web-based UI. After some specific reconfigurations in the FIWARE OS services endpoints the node was provided for external use over the FIWARE Cloud portal, while it was monitored through the XIFI monitoring infrastructure. For the federation of PiraeusU node, after its successful setup, there were some reconfiguration options in the existing setup so as to achieve the integration with the FIWARE infrastructure. In particular:

- It was updated the corresponding point(s) that refer to the authentication service endpoints so as to be compatible with the keystone authentication process, which is hosted by XIFI infrastructure in the public keystone in the region of Spain.
- It was performed the update of the neutron, nova and cinder configuration files in the controller node so as to agree with the credentials that correspond to the PiraeusU node software agents.

Further to that, PiraeusU node technical team, having previous experience on the setup of the monitoring infrastructure (as it is described above in the current section), performed the manual setup and configuration of the Node Monitoring System (NMS) in the infrastructure. For the monitoring our NMS includes the Nagios 3.4.1 over Ubuntu 12.04.1, while a set of additional components for the XIFI monitoring were included and configuration so as to achieve a successful federation in the monitoring part, as well. In particular:

- It was installed the latest version of the ODC on the controller node, whereas it was updated the corresponding odc.conf file with the appropriate configuration options, in order to provide access to the PiraeusU node information that, among others, are related with the available VMs, active user, available networks, active floating IPs, and so on.
- It was installed the FIWARE NGSI Adapter (v1.1.1) and the FIWARE Context Broker GE - Orion (v0.13).

For the realization of the above steps, it was important the exploitation of existing guidelines provided through the deliverable documents that have been written in the context of the XIFI project. Specifically, there were followed the guidelines that are provided in the corresponding wiki pages and FIWARE website: a) for the OpenStack Data Collector module, at http://wiki.fi-xifi.eu/Public:OpenStackDataCollector, b) for the Orion Context Broker (v0.13), at http://catalogue.fiware.org/enablers/publishsubscribe-context-broker-orion-context-broker, c) for the NGSI Adapter (v1.1.1) at https://github.com/telefonicaid/fiware-monitoring/tree/v3.5.2/ngsi_adapter, d) the deliverable document D3.2 for the Infrastructures monitoring and interoperability adaptation components toolkit and API (Revision v.1.1), whereas finally after the finalization of the above setup, it was filled the requested information for the public access on PiraeusU Orion GE instance (Public Context Broker), at http://wiki.fi-xifi.eu/Xifi:Wp5:Context_Broker_Public_IP_Address.

After the successful realization of the above actions, the PiraeusU node is Up and Running with OS Icehouse release, by 21/1/2015 until today.

4.3.1.3 UTH

On 2/5/2014 UTH was contacted by email and was provided with technical XIFI documentation and references. In the same mail it was asked to provide the deployment plan. On the 12th May UTH provided the plan. They added to have some problems with the lack of MD-VPN support from their NREN provider. There was a negotiation with GRNET: UTH was confident that the Greek NREN could provide MD-VPN services by the end of July.

During the call held on 10/7/2014, UTH said that they foresaw to finish the basic set up sooner than expected. The hardware was obtained and they were going to start the deployment on 14/7/2014. As far as the connectivity, their NREN (GRNET) and university's NOC are in touch. UTH would be able to provide more info next week.

On 16/7/2014, during the call of the new nodes, about the connectivity UTH confirmed that GRNET and UTH's NOC had started the procedure of the connectivity deployment. VLAN info had been exchanged too. The HW deployment was to be available by the end of July.

On 24/7/2014, UTH said that GRNET had generated the L2VPN connection with UTH university. After concluding the deployment of the node, UTH was to proceed to the connection with the rest of the XIFI cloud.

On 13/8/2014, UTH said they had started the process of connection with GRNET and the GEANT: the MD-VPN functionality was then available and, despite of vacation period, they would try to connect by the end of August. UTH installed the hardware except from the Openflow switch. They were expecting its arrival in September. No matter how they connected a simple switch in order to make all the configurations. The installation process they were following was the ITBox one. UTH finished the preparation of the master node, meaning the installation of the ITBox image, and they were in the process of configuring the network interfaces according to the WP deliverables.

On 27/8/2014, UTH said that the network circuit among GRNET and UTH was ready, as well as among GRNET and GEANT. They were on hold regarding the due connectivity with XIFI cloud. They were expecting GEANT to confirm the network status and provide the clarifications needed for the conclusion of the process.

On 3/9/2014, UTH said that the connectivity with GEANT, and thus, the rest of the XIFI community, was established. They were able to ping other online nodes. About the HW, they were still expecting the Openflow switch. Nevertheless they were going to use a surrogate one, the Pronto 3290. Until the end of following week they were going to finish all the installations they had to do with OpenStack and the XIFI accompanied tools.

On 10/9/2014, UTH stated that they had installed the ITBox machine and they were ready to deploy the OpenStack environment by the end of the week. The environment configuration was the following:

- 1 controller + cinder
- 1 controller + swift proxy
- 4 compute nodes
- 1 VM for monitoring
- 1 storage

About the Monitoring, they were waiting for some clarification from XIFI community, to be able to finalize the installation.

On 1/10/2014, about connectivity to GEANT, UTH said that the BGP protocol was configured as well. About the Cloud Infrastructure, they installed the ITBOX 1.3.2.1 version and they deployed an environment with the following nodes:

- 1 controller node

- 1 monitor node
- 2 compute/cinder nodes

The Monitoring configuration was completed: the ping test was confirmed by both sides. The Federation (keystone proxy) activity was in progress: they filled the appropriate online doc with their IPs.

On 9/10/2014, UTH concluded the federation process and became a fully operational node of the XIFI cloud, contributing with 196 cores, 576GB RAM, 9,6TB of storage capacity. They also provided network connectivity and switching through the usage of an Openflow modular HP 5412r switch.

On 20/1/2015, UTH completed the migration to Icehouse OpenStack release via the usage of Mirantis Fuel 5.1.1 suite.

4.3.1.4 Wigner

On 2/5/2014, Wigner was contacted by email and was provided with technical XIFI documentation and references. In the same mail it was asked to provide the deployment plan. On 12/5/2014 Wigner provided the plan. They said also that the installation of the required software stack on the SGI hardware should take more time than expected, however the installation should be completed before the estimated dates. Furthermore, Wigner added to have a plan B, practicable by the end of September, in case other hardware was needed. About the connectivity, their NREN could setup the MP-VPN in few weeks.

On 3/7/2014, during the CC of the SE Europe Area nodes, Wigner said to have some problems because, at that moment, it was not possible to use ITBox for the deployment on their servers (diskless blades + gateway node + NFS storage node). CREATE-NET got in touch with Wigner to provide a solution.

During the call on 10/7/2014, Wigner said:

- as far as concerns the connectivity, to get the node outside the firewall, they needed a new optical connection to the node. Constructions started within 2 weeks.
- about the software:
    - from scratch (following the HandBook & docs in net) not using ITBox. They had diskless nodes: rootfs via nfs.
- Controller node:
    - Installed / configured:
        - MySQL, RabbitMQ, NOVA, Glance, Keystone, Quantum, Horizon, Quantum-l3-agent, Quantum-dhcp-agent
    - Issues:
        - cannot run properly the “keystone ec2-credentials-create” command Really needed?
        - Quantum-openvswitch configuration in progress
        - Cinder configuration in progress

At that moment, Wigner was trying to fix the issues.

- Compute node: the necessary packages were installed and the configuration was in progress.

On 16/7/2014, during the new nodes meeting, Wigner presented their progresses:

- as far as concerns the connectivity to GEANT, MD-VPN end-point was already configured by the local NREN (NIIF/HUNGARNET). Virtual Routing and Forwarding configurations and the deployment of additional fiber were ongoing. The new optical cable had already been

acquired and its deployment was ongoing. It was expected that at most in two weeks Wigner could start testing the connectivity. This needed to fully separate their XIFI node from the Wigner infrastructure because of security reasons. A few day delay was possible. The initial public floating IP range assigned to XIFI was148.6.80.0/24.

- about the HW, SGI Altix ICE 8200 XE, 16 blade servers were ready and dedicated to host XIFI components.
- about Cloud Infrastructure & Management, since their blade servers were diskless, using an NFS file system, they had started the manual installation of OpenStack and other components on Ubuntu 12.04. Meanwhile, Wigner was in touch with Alessandro Martellone (CREATE-NET) in order to check if it was possible to use ITBox tool in Wigner node configuration or not. The manual installation of an OpenStack controller and a compute node were almost finished. Their experiences were collected in a manual installation guide (cookbook) that could interest other partners with diskless nodes.
- about the Monitoring, Wigner helped Jose Gonzalez to test the NAM tutorial before publishing it.

On 13/8/2014, Wigner presented the status of the node:

- as far as concerns the connectivity to GEANT, MD-VPN end-point was already configured by their NREN (NIIF/HUNGARNET), vrf configurations at their side was done. The new optical cable had already been deployed and tested. Testing MDVPN connection was ongoing in cooperation with their NREN. Wigner confirmed that the initial public floating IP range assigned to XIFI was 148.6.80.0/24.
- Cloud Infrastructure & Management: the manual installation of an OpenStack controller and a compute node was done. A setup with 1 controller and 15 compute nodes were up and running, and were under testing.
- issues: networking issues with OVS: sometimes the VMs could not be reached from outside. It was strange: they saw who had ARP requests on the outgoing interface of the virtual router, but there was no reply. While monitoring the physical interface they saw that the reply packets arrived correctly encapsulated in GRE packets. It seemed that those packets were lost between the GRE tunnel endpoint and the br-tun on the controller. They were working to solve this issue.

On 27/8/2014, Wigner said that a possible cause of the networking issues with OVS was the misconfigured network. Thus, the reinstallation of both the controllers and the compute nodes was required. To be finished and tested by 31 August.

On 3/9/2014, Wigner said they had solved the OVS issue, probably caused by a misconfigured network. They would have added the secondary external network within the day and they were going to launch DCRM installation this week.

On 17/9/2014, about the Monitoring, Wigner said that Nagios installation had almost finished: they had some issues with the manual installation because some of the plugins were missing, but it had been solved by now. Only rabbitmq plugin was missing (not configured properly), but they were going to finalize it within the following day. They got D3.5 for Internal review which describes the manual installation of most of the monitoring components.

On 1/10/2014, Wigner said that the Monitoring installation was completed and the following components were installed and tested: Nagios, ODC, NAM, DEM, NGSI Adapter and Context Broker. They used the Mirantis-Fuel Ubuntu repository to install rabbitmq and libvirt plugins. Context Broker is connected to BigData GE in Trento. The connection was working and Trento was able to reach their Context Broker. About the Federation (keystone proxy), they tested their local installation, before sending the request to Alvaro Alonso. The list of public service endpoints was to be sent soon. They were waiting a response from Fernando Lopez about the access for GE images, but, , GEs would be installed by the Spanish node after the federation.

On 10/10/2014, Wigner completed the federation process: all services were green in the infographic site, the VMs could be deployed through the cloud portal, the floating IPs could be assigned to them and the VMs be reached from outside.

In January 2015, Wigner was creating a XIFI services migration plan from Grizzly to Icehouse. The safest path was chosen. The former Grizzly services are to be maintained in a parallel fashion, while a clean new Icehouse environment was set up on different servers. At next, when the new system would be properly tested and both systems federated, migration could begin, which was expected approximately by the mid of March. In the 2nd week of January installation began in two spare blades and a new 50 TB storage was installed, which hosted the root file system. In the new setup the faster InfiniBand network was used.

On 4/2/2015, Wigner said that the migration to Icehouse could be completed by the end of February or middle of March. They had already deployed Icehouse in separate blade servers and they were testing the environment. The installation of monitoring components was ongoing. Wigner led internal negotiations and as a result 16 more blade servers could be allocated to the XIFI operation. With these hardware resources the Icehouse installation is comparable in size and safely, and it could take over the load of the Grizzly servers. Experimentation with Grizzly-Icehouse migration was ongoing as well.

Between 25-27/2/2015, the Grizzly Controller and network node operated by Wigner had been under a severe DoS attack and the node was forged. A memory resident automatically launching Trojan malware (Linux/XOR.DDoS) was found among system processes. The process was isolated and the generated DoS traffic was blocked by securing firewall rules. The node was down for approximately 2 hours in the early afternoon while the filesystem was cleaned from the malware. After rebooting the node and restarting services, running VMs are reachable. This fact was not introduced any data loss of the system users.

4.3.2 Geographical Area Central-East Europe

The Central-East Europe area is composed by the following new nodes:

- CESNET
- PSNC
- ZHAW

and it is managed by Deutsche Telekom as task5.5 partner.

4.3.2.1 CESNET

CESNET was chosen to be the first node from the open call to be connected to XIFI. A timeline was scheduled to achieve the integration of CESNET before the review meeting in May. The schedule was considered tight and risky. The initial deployment plan was contributed in the wiki and updated several times during the time of the integration for various reasons.

CESNET required to get insights in the behaviour of the IT-Box before installing. CESNET chose not to use ITBox and focus on deployment via StackForge's Puppet modules instead. The connection to the MD-VPN could be established with support from Dante till the 2nd of May.
It was considered what the minimal set-up components would be in order to establish the integration of CESNET in time. CESNET achieved the main cloud infrastructure installation on the 21/5/2014. It was immediately started to set up the monitoring system with some direct support of Create-NET. In order to achieve a timely integration it was considered to run the VM for context Broker for the monitoring not inside the CESNET node. Some firewall issues still needed to be fixed on the 22/5/2014 to get the data from the OpenStack Data collector. The installation of the DCRM that was not directly needed by the monitoring system was postponed.

The concept to connect to the central keystone was discussed and anticipated to be established on the 26/5/2014.

On 23/5/2014 it seems well achievable to have CESNET connected to the federation before the review meeting in Brussels.

On 30/5/2014 CESNET switched to using the central Keystone instance in Spain. The site used nova-network rather than Neutron in order to get high-availability and IPv6 support, features missing from Neutron in the Grizzly release. The status of CESNET node could already be reviewed in the infographic page: http://infographic.lab.fiware.org/status/

In August CESNET enhanced the internal monitoring and alert system. On 17/9/2014 CESNET upgraded 8 host servers to 128GB ram and 2×10Gbit ethernet cards per host. Also each compute node was equipped with new SSD disk.

After the Dublin meeting (24/9/2014) CESNET also finalized a federated monitoring tool chain and Prague node has appeared in status and infographic pages.

On 12/12/2014 CESNET finished configuration of internal backup solution of the infrastructure. On daily basis they were backing up the glance images and the Git repository of puppet configurations. Two times per day they were backing up a dump of OpenStack's databases (glance, cinder, nova). From the first days of 2015 CESNET started preparation and work on upgrade to OS Juno. CESNET earmarked a part of the infrastructure to prepare resources for a separated instance of new Open Stack instance.

Technical and operational feedbacks

The question arose whether it would be possible to run the OpenStack controller as a VM rather than on a dedicated physical machine. It was estimated that it might be possible but also considered that performance issues might come up later on in that approach.

It was clarified that the use of IT-Box for installation is not mandatory but strongly recommended to ease up the installation and set-up.

On February 2015 we were still waiting for approval to free resources occupied by testing VMs for completion of upgrade process.

4.3.2.2 PSNC

A contact was established and necessary documentation was provided to PSNC.

On 21/5/2014 a short telco was held to discuss the needed steps for integration. Some questions arose by PSNC that were documented in the FAQ of the wiki. The deployment plan was provided to the wiki by PSNC.

On 12/6/2014 a short F2F meeting took place during the general meeting in Berlin. Further steps were discussed.

On 16/6/2014 all servers and switch MX-80 was installed on rack.

On 30/6/2014: switch MX-80 and all servers were ready to use. XIFI Core Backbone connection was finished. 12 servers 12 GB and 8 servers 8 GB (all with 8 cores) were provided. The MD-VPN service was configured.

On the first week of August PSNC was decided to use ITBox and Fuel to install and configured XIFI node. PSNC installed and configure ITBOX 1.3.4.0, next updated to ITBOX 1.3.2.1RC. PSNC had configured Controller Host, Monitoring host, Computes and Storage Hosts. PSNC’s installed and configured Fiber Channel connection to our Storage Host to increase the capacity of hard disk space.

On 2/9/2014 the internal networks were configured: Administrative Network, Management Network, Storage Network, Federation Network. We started to install FIWARE components following the

XIFI's core backbone 5.2 and XIFI's Handbook v1.

On 26/9/2014 ITBox tool installation and OpenStack installation was finished. All nodes could be reached. The pull of Public IPs was also ready. PSNC still had some problem with monitoring software and had to configuring it manually.

On 29/9/2014 monitoring and configuration files for XIFI were ready.

On 8/10/2014: Registration Status was OK, and PSNC was ready to downloaded GE Images.

On 13/10/2014: GE images was downloaded and node were connected to FIWARE Lab.

PSNC was contributing with 120 cores, 156G RAM, 5 TB of hdd capacity. PSNC uses an openflow switch MX-80 to provide a 10Gb/s connection to XIFI backbone.

Software installed on XIFI nodes in PSNC:

- ITBox ver. 1.3.4.1,
- Nogios ver. 3.4.1,
- OpenStack ver. Grizzly.

On 22/12/2014 the PSNC node was a down due to upgrade software to Icehouse OpenStack.

On 9/1/2015 the OpenStack was updated to Icehouse version, but there were some problems with monitoring subsystem. Server administer was contacting with XIFI support to solve this problem.

On the 16/2/2015 the problem with configuration of monitoring was solved, and node was fully operational.

To provide a seamless cooperation with monitoring subsystem PSNC has installed:

- ngsi adapter 1.1.1
- ngsi event broker 1.3.1
- context broker 0.15.0
- ngsi parsers from FIWARE svn v.3.5.2-xifi-2

4.3.2.3 ZHAW

The process of deploying the Zurich XIFI node started in April 2014. The initial work involved reading the documentation to understand the architecture of the system and determining how the available resources in ZHAW could be integrated into XIFI. At this time, ZHAW also recommenced previous conversations that they had with both ZHAW Central IT and the SWITCH (the Swiss NREN) regarding high speed connectivity to XIFI.

At the time the project commenced, ZHAW got the opportunity to acquire new hardware which they could dedicate to the project. Consequently, they had to specify the hardware and liaise with suppliers to determine the best offer. In the end, ZHAW chose IBM x3550 M4 servers which are significantly resourced with 192GB RAM, dual processors and 6*600GB disks. The hardware was delivered on Aug 7 2014.

Concurrently, GEANT, who had responsibility within the project for providing the federated connectivity solution orchestrated a solution which required a long distance connection to a GEANT POP in Geneva, as SWITCH does not support MD-VPN. This was a somewhat different requirement than that which ZHAW had originally envisaged with SWITCH and consequently, the internal network configuration within ZHAW had to be modified – this involved installing a new switch which had a direct L2 optical connection to the SWITCH endpoint in ZHAW. This was configured on 21/9/2014.

Installation of OpenStack was performed using ITBox. This did take some time and involved multiple installations of the system as it was sometimes not clear what had been properly installed. More specifically, in some cases, there were difficulties while installing the monitoring components which became apparent after the installation and while ZHAW was configuring the system. It was also necessary to configure the second network interface which was involved to modify the routing agents and bridging configuration on the nodes. The base installation of (non-federated) OpenStack was performed on 2/10/2014.

Configuring the monitoring components was done next. This involved installing some of the components on the monitoring node by hand (Nagios, context broker) and then configuring them appropriately for the Zurich node. The monitoring configuration was completed on 3/10/2014.

Having finalized the monitoring configuration, the final step was to federate which involved modifying the configuration files to point to a different authentication service (keystone) hosted on the Spanish node. This was completed in early October 2014.

The original installation of OpenStack was somewhat problematic as it was based on the Grizzly release. The particular deployment decisions which ZHAW chose (deploying the MD-VPN network as the primary interface in the deployment tool) caused problems subsequently as having two independent routing agents was never stable – this is a known problem within Grizzly. This meant that the system was operational and federated, but exhibited problems for some months.

The ultimate solution to the problem was to install Icehouse: as Icehouse has much more stable routing agent implementation in which all routing is performed within a single agent, the double agent approach (and associated problems) did not manifest. Installation of Icehouse took some time: there were a couple of issues that arose in the deployment. These include:

- The Icehouse variant of Fuel has new default options which had a significant impact on the resulting configuration – it was not obvious that the new default is not to allocate public facing IP address to each node. ZHAW performed a deployment with this setting as it is more economical with IP addresses; however as the federation requires each node to be able to communicate with the authentication service directly this did not work. Some workarounds were tried including providing ssh tunnels however, this did not work seamlessly. Ultimately it was necessary to perform a reinstallation in which the new default was not selected.
- The long-distance connection to Geneva demonstrated some problems. It took a couple of weeks working with GEANT, SWITCH and ZHAW central IT to diagnose this problem: ultimately, it was a problem in the configuration of the long-distance tunnel which SWITCH needed to resolve.
- The use of a new monitoring node causes problems: an Ubuntu solution was used initially, but eventually this had to be changed to CentOS as the monitoring solution does not work well with Ubuntu.

A fully monitored, operational variant of Icehouse was deployed at the end of February 2015.

4.3.3 Geographical Area South-West Europe

The only additional node foreseen in South-West Europe was Gowex. Due to the bankruptcy of Gowex, no activity run in this area.

4.3.4 Geographical Area Central-West Europe

The Central-West Europe area is composed by the following new nodes:

- Com4innov
- iMinds

and it is managed by ORANGE as task5.5 partner.

4.3.4.1 Com4Innov

On 12/5/2014, Com4Innov was contacted by email and was asked to provide the deployment plan. One conference call was organized by Orange on 16/5/2014 to give some details.

Status on 12/6/2014 about Connectivity: The administrative tasks are done, contract has been signed with RENATER and the interconnection between the data center where our equipment are located and INRIA who is hosting Renater access node is proceeding. They were awaiting from Renater the date of connection and also information regarding the availability of MD-VPN node in Sophia. With respect to the hardware,the list of equipment was defined and the order was in progress.

On 1/9/2014, the hardware was delivered to Com4Innov datacenter.

On 3/9/2014, the OpenStack deployment was concluded, Com4innov began to install FIWARE components following the XIFI's core backbone 5.2 and XIFI's Handbook v1.

On 17/9/2014, the XIFI MD-VPN has been successfully installed. Com4innov nodes were reachable from all Federation's nodes and vice versa. Joining federation Handbook was followed but OpenStack nodes present many issues on Grizzly release after endpoint registration.

On 16/10/2014, the Com4Innov node was visible on Infographics webpage. Context Broker and Ngsi reported all information successfully. The services were partially operational: Com4Innov was experiencing many issues on Nova & Cinder services (help was asked by Com4Innov on Wp5 mailing list to have feedback and advices from other nodes). They were also waiting a range of public IPs from Renater.

On 10/11/2014, Com4Innov received a range of public IPs from Renater. The range was used for OpenStack floating IPs.

On the 18/12/2014, the Controller node has crashed. The server was reinstalled and the last backup was restored.

On 5/1/2015: Node name was changed from "C4I" to "SophiaAntipolis".

On 14/1/2015, a problem was identified on Cinder component. Volumes creations and attachments were not available. Com4Innov followed the instructions and the advices provided by Cesnet and the storage node was restored. But, although the configuration was changed, the problem was still present.

4.3.4.2 IMINDS

On 12/5/ 2014, iMinds was contacted by email and was asked to provide the deployment plan. iMinds confirmed having discussion with Belnet about the MD-VPN set-up. BELNET has decided to deploy the MD-VPN capability in their domain. They will need a few weeks to implement it.

On 22/5/2014, Orange confirmed to iMinds the MD-VPN service is deployed by GEANT and detailed information can be requested directly by email to GEANT partner.

On the 16/7/2014, iMinds began the evaluation and testing of the available tools for the deployment of OpenStack and XIFI software components into its existing Virtual Wall infrastructure.

On the 31/7/2014 it was decided that ITBox and Fuel were insufficiently compatible with the existing Virtual Wall management infrastructure. Because of this, iMinds decided to create custom Chef deployment scripts for deployment. The development of the Chef-scripts to deploy OpenStack was finished by the 4th of September.

On 17/9/2014, the OpenStack deployment was concluded, and iMinds started with the manual configuration of the XIFI software components, as described in the XIFI Handbook v1.

On the 26/9/2014, BELNET resolved the last configuration issues to enable the connectivity of the iMinds site to the XIFI MD-VPN. On the 30th of September, the deployment and configuration of the monitoring service was also finished.

Between the 2nd and 6th of October, iMinds debugged its configuration, as random kernel panics appeared on the controller node. This was resolved by upgrading the kernel and Openvswitch-versions, after which OpenStack was redeployed. On the 6th of October iMinds started the registration process with Alvaro Alonso to become fully federated. This process was completed on October, 10. iMinds requested an new, larger IPv4 address pool from Belnet on the 7th of October. This was received and configured on the 15th of October.

On 1/12/2014, TID uploaded and synchronized GE images on the iMinds XIFI-node.

On 7/1/2015, iMinds started the plan and test configurations to migration the current OpenStack Grizzly to the next version Juno, using 4 physical servers that were running under capacity.

On 13/1/2015, Rogue Spam User attacked the Federation starting to create VM. Gent node was affected with the creation of 1 VM with 2048MB RAM. The user was unable to create more because the RAM memory quota. This VM was deleted.

On 20/1/2015, Floating IP's pool exhausted. iMinds discovered a problem since the users that have an IP assigned doesn't have any vm's running. Users who allocate an unused public IP's were avoiding others use.

On 4/2/2015, iMinds was informed that the Federation is not ready to upgrade OpenStack to Juno version, first is wait to fix all compatibility issues. The upgrade is postponed.

On 12/2/2015, iMinds updated monitoring node with new component release OpenDataCollector 2.2.

On 23/2/2015, they started to apply a policy in order to maintain the floating IP pool. If an allocated IP is unused for an entire week, this will be manually deallocated and returned to the pool to be used by another user. This was a temporary solution before iMinds can upgrade to OpenStack Juno and start using IPv6.

On 2/3/2015, iMinds carried out the migration to more secure users and passwords of the OpenStack services, after receiving the new credentials and being activated by Alvaro Alonso.

4.3.5 Geographical Area North Europe

The North Europe Europe area is composed by the following new nodes:

- Acreo
- BTH

and it is managed by WIT as task5.5 partner.

4.3.5.1 Acreo

Contacted via email on 29/4/2014, asked to update the Deployment Plan for new Nodes on XIFI wiki. Contacted via email on 5/6/2014, asked to revisit the Deployment Plan to check mile stones are correct. Issue with XIFI General Agreement (GA) clarified.

Contacted via email on May 7th tested MD-VPN running (network prefix)

XifiRouter1#show ip route vrf XIFI

10.0.0.0/8 is variably subnetted, 6 subnets, 2 masks

C 10.0.0.0/20 is directly connected, GigabitEthernet0/2

L 10.0.0.1/32 is directly connected, GigabitEthernet0/2

B 10.0.16.0/20 [20/0] via 188.1.201.9, 2w6d

B *10.0.32.0/20 [20/0] via 193.51.178.40, 1w5d*

B *10.0.48.0/20 [20/0] via 193.51.178.40, 1w5d*

B ***10.0.144.0/20 [20/0] via 130.242.80.54, 4d07h***

Skype call with on issue with ITBox (discuss trunks, Vlans there ITBox roles, preinstall networking conditions for ITbox)

Skype call on the 6/6/2014 with Jonas and Roland on to discuss ITBox, OpenStack controller and compute nodes (provided basic guides on how to get VM and networking up and running from the command prompt). We notice that the networking node had the :OpenStack OVS plugin installed but no OpenVSwitch deployed via ITBox.

Quick catch up at XIFI Technical Meeting in Berlin.

On 2/9/2014 call on Skype regarding a couple of issues in OpenStack customizations and finding correlated info in the wiki.

On 4/9/2014 main conversation about network configuration issues.

On 3/10/2014 Skype call about monitoring followed by email.

On 8-9/12/2014 email conversation about volume attachment management.

4.3.5.2 BTH

On 29/4/2014 Karlskrona (BTH) was contacted by email, and asked to fill in a Deployment plan and to start tracking issues that they might experience when deploying the node.

On 3/6/2014 the node reported that they were connected to the GEANT-MDVPN, and requested the node's allocated IP range within the MDVPN.

On 3/6/2014 had a Skype chat with the Karlskrona node with regards to the interface naming within ITBox. Asked the node to update their deployment plan.

On 4/6/2014 informed the Karlskrona node that the route was advertised Irish PE router.

On 3/9/2014 contacted the Karlskrona node via Skype, as the node reported having problems with the ITBox deployment. Especially the network configuration. The problem was related to interface naming, vlan tagging and Ethernet switches not behaving as they should.

On 11/9/2014 the node informed that they started writing a how to wiki, http://wiki.fi-xifi.eu/Xifi:Wp5:how_to_deploy_a_xifi_node.

On 29/9/2014 was contacted by the node as they had some challenges with the IP allocations to the computers running OpenStack, and OpenStack allocating the same IPs to VMs.

On 30/9/2014 the node stated that they deployed an OpenStack environment using ITBOX 1.3.4.0. The environment had the following configuration:

- 1 Controller Node
- 4 Compute + Cinder Nodes
- One Linux Router acting as gateway, for both MDVPN and Internet access, based on Ubuntu 12.04 LTS.
- One ITBox 1.3.4.0 computer
- One Monitoring node, installed from a standard Ubuntu 14.04 LTS.
- Two switches, one for ITBox and the other for the OpenStack network.

On 1/10/2014 was contacted by the node asking as they wondered that information they should have

received when they federated. The node also informed that they were working hard on getting the monitoring working.

On 1/10/2014 contacted both Swedish nodes (Karlskrona and Stockholm) with instructions on how to install a cinder node on a Grizzly OpenStack deployment.

On 6/10/2014 was contacted by the node stating that they were federated, and informed them that they now had to rebuild all the components that were built with the old admin tenant id.

On 13/10/2014, forwarded an email to the node regarding VNC access from the Cloud portal. The node was informed by Attilio that the node was subscribed to the master node, and it was OK.

On 19/11/2014, removed on compute node as to have a node for Icehouse testing.

On 20/11/2014, tested to deploy Icehouse on one compute node, followed '10-minutes' instructions. Worked, but operates stand-alone.

On 24/11/2014, tested to deploy Icehouse and Juno on a few old computers, using plain Mirantis Fuel. Computer having to small hardware for a serious test.

On 16/1/2015, detected problem with VNC connection to VMs. Submitted Jira Ticket.

On 21/1/2015, the node obtained an IBM BladeCenter with 13 HS21 blades. BTH will use this to deploy Icehouse, but initially it needs to be upgrade to meet the 4G/core requirement.

On 9/2/2015, finalized the configuration to get live migration to work on the Grizzly (operational) environment.

On 11/2/2015, the node tried to update the ContextBroker to the latest version. This caused the node to disappear from the infographics page, and was not resolved until we reverted back to the supported version within XIFI.

On 17-18/2/2015, attended the final XIFI meeting in Prague.

On 24/2/2015, obtained the needed hardware for the BladeCenter. Starting to test deployment via ITBox 1.3.4.0.

On 2/3/2015, the node noticed that a new ITBox with Icehouse support was about to be released, waits until the 10th before proceeding.

On 3/3/2015, needed to remove yet one compute node as it was attached to a UPS that needed replacement. Did live migration of all tenants on to the remaining two compute nodes. Shutdown the compute node.

On 3/3/2015, updated all passwords on all relevant services.

On 4/3/2015, all auto tests were executed successfully.

4.3.6 Associate Nodes

The Associate Nodes are the following:

- Intellicloud (Crete)
- University of Messina (Italy)
- Mexico

Intellicloud and University of Messina nodes are managed by Telecom Italia. Mexico node is managed by Telefonica.

4.3.6.1 Intellicloud (Crete node)

On 19/5/2014, TSI-TUC signed an agreement to becoming an associate member of the XIFI infrastructure federation. A few mails had been exchanged before where TSI-TUC presented its plans for supporting XIFI (FIWARE LAB), the existing infrastructure, its plans for extension along with a short description of its organization and personnel. Note that TSI-TUC has the status of partner in project FI-STAR (www.fi-star.eu), with the role of supporting deployment of 8 medical use cases on an already established stable OpenStack cloud environment.

On 24/6/2014, TSI-TUC attended the XIFI training session in Madrid which offered a good opportunity to get acquainted with the XIFI technologies and solutions and to establish a communication channel with the XIFI experts that would help its colleagues in a smooth integration in the federation.

On 4/7/2014, TSI-TUC received its new equipment and started the OpenStack (Grizzly on Ubuntu 12.04) deployment using ITBox 1.3.4.0 and experimented on many topologies of its internal network, facing a lot of installation problems. The final environment configuration followed a HA scheme:

- 1 controller + monitoring
- 1 controller + cinder
- 1 compute + cinder
- 1 compute

(TSI-TUC aims to incorporate into the Crete node infrastructure the existing separate infrastructure dedicated to the FI-STAR project, adding eventually 8 more compute nodes.)

On 9/9/2014, GRNET (the local NREN) established the deployment of the MD-VPN capability in their domain.

On 26/9/2014, XIFI MD-VPN connection has been successfully installed.

On 23/10/2014, TSI-TUC started the procedure of joining the XIFI Federation, facing a lot of problems that mainly had to do with the new configuration of the OpenStack Environment and its connection with Keystone Proxy. Help was asked to the XIFI experts and other XIFI partners.

On 11/11/ 2014, the Keystone-Proxy connectivity has been established. Then a long period started with self-tests on the node of Crete, resolving a lot of problems that had to do with VM creation, HA scheme configuration, errors mainly on dhcp agents, network conflicts, Nagios, NRPE, etc.

On 9/12/2014, the Crete node was completely synchronized with VM images of the federation.

On 20/1/2015, the Crete node started the finalization of the integration on the XIFI Federation: monitoring and infographic configuration. NGSI Adapter and OpenStack Data Collector (ODC) had to be established and connected with the existing Context Broker.

On the 29/1/2015, the Crete node could be seen in the infographic page.

4.3.6.2 University of Messina

On 5/5/ 2014, after enquiry from the same University, UniMe got info from TI about formal processes for getting associate node status and joining the federation.

On 12/5/2014, UniMe received info about upcoming training sessions for FI-Ops.

On 19/5/2014, UniMe provided preliminary information about potential hardware to be devoted to the Federation.

On 24-25/6/2014, a representative from UniMe attended FI-Ops training sessions in Madrid.

On 1/8/2014, UniMe sent more info about hardware, in particular with regards to forthcoming hardware procurements, including network upgrades which would positively impact the node and its connectivity to the Federation.

On 8/9/2014, UniMe received info about upcoming training sessions for development with FIWARE.

On 3-5/10/2014, a representative from UniMe attended FI-Dev training sessions in Milan.

On 1/12/2014, UniMe was acknowledged official entry into the Federation.

On 2/12/2014, UniMe was asked to provide mailing addresses and register on the Wiki for access to internal documentation.

On 5/12/2014, UniMe followed up by informing TI that registration to the Wiki was effective, and that mailing addresses were set up according to the request.

On 15/1/2015, representatives from UniMe met with GARR PoP people in order to discuss the MD-VPN requirement and officially submit the request.

On 16/1/2015, UniMe completed the set up (basic hardware enablement and OS install) of the blade servers for the node.

On 3/2/2015, UniMe got confirmation about the last procurements, in particular about a 10G Vyatta-based router/firewall and 2x 10G Cisco switches.

On 24/2/2015, UniMe enquired about the status of the MD-VPN activation, getting back a confirmation that this procedure was experiencing delays temporarily and that ETA was unknown at that time.

On 24/2/2015, there was a follow-up call from TI to track progress with UniMe activities, in particular asking for node support info and journal updates. TI also suggested switching from MD-VPN to a dedicated site-to-site VPN with another node of the Federation (possibly doing as Trento in leveraging Lannion for that), as a stopgap solution in the meantime the MD-VPN request progresses. A contact point in Trento has been provided for further info on actions to take in order to set up such a VPN.

On 27/2/2015,, UniMe sent a mail to set up a call with the aforementioned contact point in Trento.

On 2/3/2015, UniMe sent required node support info, per the corresponding Wiki page and deliverable drafting.

4.3.6.3 Mexico

On 10/4/2014, Infotec was contacted by email and was asked to provide the deployment plan, infrastructure available and planning for the deployment of the rest of the components. Some exchanges of emails to complete the description of the equipment needed to deploy a node.

On 3/6/2014, Update of the process by Infotec by phone and schedule of the activities to develop in the Campus Party Mexico and the different advances that they have in the Mexican node. During this call we planning the different actions to deal in the Campus Party in order to inform about the collaboration between InfoTec and Telefónica regarding the configuration of the FIWARE Lab Mexican node.

On 24/6/2014, Campus Party Mexico where we have several presentation meetings in order to see what they need to use, install and configure the Mexican Node. Resolution of some network specific problems related to the configuration of neutron in case of Infotec. Analysis of the connectivity with MDVPN through the RedClara network connected to the CUDI (Corporación Universitaria para el Desarrollo de Internet).

On 29/8/2014, Infotec finalize the deployment of the mock up in Mexico, D.F. and inform about the process to buy the new infrastructure of the Mexican node. They continue in contact with CUDI in order to connect to the MDVPN through RedClara.

On 11/9/2014, Infotec inform about the process to buy new infrastructure (over 1000 cores) but they have some problems with the acquisition of the equipments. In the meantime, they provide a mock up working in Aguas Calientes where finally will be deployed the Mexican node. The connection through MD-VPN will be realized with a connection with the VPN that Red.es is configuring in order to resolve the problem with the delay in the communication with RedClara. TID prepare documentation about the different things that they have to take into consideration in the installation of the node. Decision should be taken in order to install Havana or Icehouse. TID recommends the installation of Icehouse.

On 15/9/2014, TID review the user management in case of Mexico specially of there is any legal issue regarding the management of data. Besides, we talk about the configuration of the IdM-Keystone Proxy in order to configure the access of the user to the new FIWARE Lab node. Red.es ask about the connectivity between RedClara and GEANT in order to see if there could be any potential problem in the interconnection between the two networks. In case of any problem, we continue with the plan to connect through the VPN configured by Red.es in order to connect to the FIWARE Lab. We decide to reschedule the deployment of the monitoring components after we have a stable version of the node.

On 23/9/2014. The contacts between GEANT and RedClara continue in the same line. Mexico installs Icehouse version of OpenStack. TID will check internally the functionality of the node. We planned to update the images in the Mexican node in order to align with the images that we have in the FIWARE Lab nodes. Recommendations about the definition of quotas, We suggest to move from 50 public IPs to a less amount of them. TID check the node and some comments are translated to the Mexico team, which include the impossibility to see the console of the virtual machine. Currently, the Object Storage is not installed in Mexico, they are waiting to receive new resources in order to install this OpenStack service. The language should be english in order to join to the federation.

On 29/9/2014. TID comment that the configuration of the network is correct but some works should be done in order to have the node completely functional. Mexico need to create a router with the link port in the external network (207.249.127.33). Assign to this router, Mexico team has to create a new interface to this router with a local network (in this case 192.168.10.0/24). Finally, in the deployment of the virtual machine assign a floating IP to the node. We recommend that it is not work to put virtual machines directly over the external network, it is not a good politics of use of OpenStack resources. The actual configuration of the services is:

- keystone (just temporal before joining the federation)

  http://filab.infotec.net.mx:5000/v2.0 (publicurl, internalurl)

  http://filab.infotec.net.mx:35357/v2.0 (adminurl)
- nova

  http://filab.infotec.net.mx:8774/v2/%(tenant_id)s (publicurl, internalurl, adminurl)
- glance

  http://filab.infotec.net.mx:9292 (publicurl, internalurl, adminurl)
- neutron

  http://filab.infotec.net.mx:9696 (publicurl, internalurl, adminurl)
- cinder v1

  http://filab.infotec.net.mx:8776/v1/%(tenant_id)s (publicurl, internalurl, adminurl)
- object storage

  Not installed.

On 30/9/2014. Discussion about the installation of the Monitoring components following the indications of TID. The URL should be something like http://filab.infotec.net.mx:1026/monitoring/regions/

On 26/11/2014. Due to the delay of the interconnection between CUDI and GEANT, we decide to move directly to the VPN connectivity between Spain and Mexico. We decide to wait to have the new infrastructure to install the monitoring components. Scheduled by beginning of 2015.

On 31/1/2015. Mexico inform that they have the new infrastructure available in Aguascalientes. They are working with the provider of the architecture and in the next week we receive new information about the status of the node.

On 14/2/2015. Maintenance operation is done in order to configure the new infrastructure of the node in order to offer power supply in high availability. We develop some test with the new infrastructure in order to see that all is working properly, some errors was detected and Mexico support team is working to resolve them.

5 LESSONS LEARNED& KNOWLEDGE MANAGEMENT

The purpose of this chapter is to analyse the activity of node federation with the goal to collect useful information for FI-Core and for future activities.

The first paragraph is the analysis of a survey between IOs and project stakeholders; the goal of the survey was to collect all the 'lessons learned', classify them and produce some recommendations useful for the prosecution of the federation activity.

The second paragraph reports all the work done in terms of sharing knowledge, solutions, 'how to do' documentation arisen during the federation activity; this activity has been reported in a set of Frequently Asked Questions that are shared in the wiki [15].

5.1 Lessons learned

In the second week of February 2015 we prepared and run a survey between the IOs and some stakeholders of the XIFI project, in order to collect 'lessons learned' indication related to the activity of federation of the project. The data has been elaborated and analysed and the result of this activity is reported here.

In order to give transparent information, all the 'lessons learned' are reported as they were **originally produced**. The data has been classified and analysed in order to give some recommendations for the prosecution of the federation activity in FI-Core and, generally, for the future. In the next paragraph of this chapter the lessons learned have been partitioned with a classification elaborated during the analysis of the survey.

Lessons learned concerns:

- ITBox;
- Knowledge sharing improvement;
- Federation process management;
- Architecture;
- Maintenance and operations.

5.1.1 ITBOX

This group of lessons learned are related to the federation tool ITBox.

ID	Lesson learned	Analysis
ITBox1	Directions and tutorials provided will never constitute a straight forward way to complete the task. Each system has its own specifics and the administrator should find the work around needed. An example is that the initial ITBOX version was not supporting the HP vendor due to software-hardware incompatibility being created from the HP's smart array. The newer version of the same tool does not have this issue since it was identified and resolved by the developers.	Two main versions of ITBox were released during the project to satisfy initial requirements. The next ITBox releases should consider common additional requirements. A study for each hardware/software specificity should be evaluated by the ITBox developers.

ID	Lesson learned	Analysis
ITBox2	We found it difficult to do the deployment with ITBox - the tool was too raw when we needed to do the deployment during Summer of last year .	From the early releases, the ITBox development roadmap has improved the functionalities offering the automatic deployment of the most relevant software packages. The ITBox roadmap should be aligned as much as possible with new OpenStack releases.
ITBox3	Integration processes where there is no ability to test the configuration prior to integration are far from optimal - it would be much better to have a clearly defined test suite - passing this should mean that integration will pass	XIFI considers several architectures and ITBox is not a magical box. For commercial deployment, add more technical HW and SW requirements
ITBox4	Deployment Issues: do not fully trust Fuel/ITBox -> It is not a Magic Black Box, to have a fully operational node some manual work is required, especially during the federation process.	Same as ITBox1
ITBox5	We have faced the problem that ITBox does not support diskless nodes. To bypass this problem we have decided to follow the procedure of manual installation. The problems we faced during the installation of Grizzly and Icehouse were solved, and the manual installation of monitoring components were also done for both systems. The description of the installation procedures for diskless blades are summarized in internal documents that are planned to be published in March.	Same as ITBox1
ITBox6	Installing the tools of the federation and federating our node take time. As we were the first nodes to be federated, procedures were not finalized. We switched from federating our keystones to centralizing it to a single point after seeing replication problems in the Mysql database. The tenants owned by the local keystone becomes kind of ghost and has to be redeployed.	Issue is solved thanks to the maturity of the tool, this lesson learned refers to the early installation and usage of ITBox.
ITBox7	High availability deployment is the reference architecture proposed by the federation (D5.2). In itself High availability involves more complexity. It adds new processes handling replication of database, load balancing in between service processes, and so on. This knowledge need as well time to be acquired. Adding High Availability on an already complex architecture	Same as ITBox1, with a specific reference to HA (High Availability).

ID	Lesson learned	Analysis
	just make the node more complex. In this version of fuel used to build the ITBox, High Availability on a Ubuntu environment was new, and then the Mirantis scripts added to run the High availability had several bugs. In consequence, a node freshly installed in High Availability using this version of the ITBox was unstable and cannot be put live as it was. A lot of debugging was necessary to correct the scripts in order to stabilize the node.	

Table 25: Lesson Learned ITBox

5.1.2 Knowledge Sharing improvement

This group of lessons learned are related to the Knowledge Sharing improvement process between the new nodes.

ID	Lesson learned	Analysis
KS1	Collaboration and sharing of knowledge among peers and partners is a key element to complete the federation process. This way it is possible to solve technical issues emerging during the federation and/ or the administration of a deployed system. The collaboration can be direct, via emails, or through well-defined tools such as Jira platform. Either way, sharing of knowledge and the existence of a highly motivated and experienced support community can provide the necessary means for efficient problem resolving.	Each Infrastructure Operator should be invited to push more knowledge in the tools, to encourage the use of interactive tool like IRC. The search engine included in the wiki is too raw and should be improved.

ID	Lesson learned	Analysis
KS2	The experience obtained during the first federation process provided the means so as to complete the second federation phase, i.e. upgrading to Icehouse, in very short period of time.	See KS1
KS3	Better information sharing around the specific configurations of the nodes and recommendations on preferred configurations would have been helpful (as many of the nodes ended up solving the same problems independently). For example, in the initial deployment, it was not clear whether to choose the public network as the default for Fuel or the MD-VPN (in the end, it seems that the public network was the better choice).	See KS1
KS4	The project should have mandated a much more rapid response information sharing mechanism - IRC is typically used in these contexts, or some other chat mechanism. Some of us set this up on an ad hoc basis and it was useful when we were under pressure to install things, but having everyone on it and having a constant conversation would have helped solve problems much more rapidly.	See KS1
KS5	Also there is a need of tools to manage this, issues trackers (Jira was not very useful during deployment), collaboration tools (in order to permit to the nodes to share knowledge), solutions.	The collaboration tools adopted in the project are Jira, Redmine and the wiki. Probably in some circumstances the process governing the usage of these tools was not well defined
KS6	Operators must document *all* installation steps if they want to be able to fix thinks at a later time. Preferably this should be done in a wiki that allows collaborative editing. The XIFI infrastructure is very complex, heterogeneous and highly evolving, which makes keeping detailed, up-to-date documentation a very challenging task. We appreciate the great efforts by the authors of the XIFI documents in handling this task, but we believe it is in the IOs best interests to also document their installation in order to keep track of the small but important details that deviate from or are not available in the current XIFI documentation. These details should be forwarded to the XIFI authors so current documentation can be updated and improved.	see KS1. Each Infrastructure Owner should track (and share) step by step installation.
KS7	Documentation needs to be better organized and classified: • We started by using deliverables to install our node, in particular for the monitoring – they were outdated. The wiki should have been used instead (our mistake).	see KS1

ID	Lesson learned	Analysis
	• We learned very late how the search function in the wiki works. It should search the whole wiki by default. Having a non-functional search in the wiki led to a lot of unnecessary extra work.	

Table 26: Lesson Learned Knowledge Sharing

5.1.3 Federation project management

This group of lessons learned are related to the federation project management, specifically how to improve the management of the federation activities.

ID	Lesson learned	Analysis
FPM1	There was a congestion of deployment activity during the last period, more or less all nodes started their deployment very late and a lot of nodes went through the date of 22nd of October. To avoid a problem like that it is needed to introduce a management responsibility (e.g. federation manager). It is necessary to prepare a proper deployment plan for all the nodes, track the plan and take actions in case the plan is not on schedule. The missing of a clear overall responsibility for the deployment is one of the reason of the failure of the deployment (asking for a deployment plan to the nodes must be avoided).	The inclusion during the same period (September-October) of about ten new nodes was not efficient due to lack of maturity of the tools and availability of some component owners. However sometimes happened that the IO delayed too much the start of the federation activities with respect to the plan defined. Need to define a plan to avoid critical paths to XIFI experts to manage critical issues.
FPM2	Missing or not useful documentation for a step-by-step deployment. To solve this, it would be useful to have trials: choose 1 or 2 nodes to start early deployment, track and document (and improvement documentation) and use this experience for the other nodes.	Step by step installation documentation should be evaluated and validated in the Testbed prior use by the nodes. In case of specific requirements that cannot be satisfied by the Testbed, the installation manual should be evaluated (by the IO and the component owner) on a node before use.
FPM3	Technical commitment of component owners for solving issues as far as issues and problems are found (internal reference the monitoring case).	Hot topic that appeared after the first release of the ITBox due to expert bandwidth to solve critical maturity issues. Need to define a deployment plan to avoid critical paths to XIFI experts to manage critical issues

Table 27: Lesson Learned Federation Project Management

5.1.4 Architecture

This group of lessons learned are related to the architecture, specifically how to improve the architecture of a node and the architecture of the federation.

ID	Lesson learned	Analysis
ARCH1	There's one simple rule to this: The higher the federated regions the larger the repetitive HTTP traffic it is cursed. The most common example to this is the highly repetitive query for a token. This happens both when a service identifies itself and when a service validates a user token. The response of the validation query is about 36K long but it will increase the number of them when a new FIWARE Lab node joins the federation. This has 2 consequences: • A high increase in the network traffic because every region does the same queries to the keystone. • Increase of the effective load in the centralized Keystone-Proxy. • Decrease in responsiveness of the different services for each service due to they have to wait for the responses which are coming through the internet. In order to resolve the third consequence and partially improve the second one, the keystone proxy has been configured to operate in an active-active high availability configuration. However, the effect of the second consequence hasn't been reduced. Maybe we could suggest developing local “keystone-proxy” caches in order to prevent sending so many repetitive queries with their corresponding repetitive answers to a centralized keystone-proxy. This action would not only decrease the traffic through the internet and the centralized keystone-proxy load but it would also improve the services regions performance as they should had to wait, in the most cases, for local responses and not from responses coming from the Internet.	It is suggested to implement local “keystone-proxy” caches in order to prevent sending many repetitive queries with their corresponding repetitive answers to a centralized keystone-proxy. This in fact is exactly what the architecture designed in XIFI is proposing. Unfortunately the current deployment is different with respect to the architecture proposed. In the future the deployment should be aligned as much as possible to the architecture.

ID	Lesson learned	Analysis
ARCH2	The only common glue for the federated regions is the Keystone-Proxy and the Cloud portal; it causes a problem specially related to the images that users can instantiate: They might differ. As every region has its own Image Service, the images updated from the GE Owners, which should be the referenced images, are only updated in Spain region. It means that a synchronized problem could appear between the Spain node and the rest of FIWARE Lab nodes. Besides, the image obsolescence can also happen with the basic images (CentOS, Ubuntu, Debian...). In order to solve this problem, a new service has been developed, which synchronizes the public Spanish node images with the other regions in order to have the images synchronized. However, as the number of federated nodes increases, this procedure is likely to become ineffective since it has to synchronize much data in little time from Spain node to the rest of the nodes. Maybe, it could be more effective a cascade synchronization schema of data. Maybe we could suggest developing local "keystone-proxy" caches in order to prevent sending so many repetitive queries with their corresponding repetitive answers to a centralized keystone-proxy. This action would not only decrease the traffic through the internet and the centralized keystone-proxy load but it would also improve the services regions performance as they should had to wait, in the most cases, for local responses and not from responses coming from the Internet.	See ARCH1 Moreover a tool in order to synchronize the images among the nodes has been developed and deployed.
ARCH3	DNS is a great service but when the number of HTTP queries increases with the new FIWARE Lab nodes, the number of DNS queries proportionally increases too. This fact decreases performance and increases the network traffic. An easy solution is caching the names in /etc/hosts and using "localhost" when possible in order not to query DNS servers at least once in every HTTP query. Other possible solution could be a DNS cache, which can improve the performance.	It is possible to distribute the DNS. So there is the possibility to have caching.
ARCH4	The usage of the federated nodes should be anticipated and done in consultation with the node. By being federated, we lost the control of a part of our OpenStack node. For instance, the user management part is no more in control of the infrastructure owners. Each Infrastructure owner has its own policy depending of course of its country and the laws that are applied. Multiple events already happened, like connecting our node to FIWARE Lab, where consulting the nodes was	1) It is recommended to give more control on the user information managed by the IDM 2) A local Identity Manager, independent from the IDM of the federation is expected by the nodes. All this is already considered in

ID	Lesson learned	Analysis
	omitted.	the architecture. The problem is the current deployment. See ARCH 1.
ARCH5	The delegation of Public IP means that the persons behind these IPs have their identities known and verified. This process is not a simple task as it implies some manual tasks behind. Until now, this problem is not solved by the federation and by default people with not verified identities can log into the cloud portal. As a consequence, this task should be taken by the Infrastructure owners themselves and as a default configuration, they apply restricted quotas where Public IP allocation is not allowed.	Reference User Policy implementation in XIFI. See ARCH 4.
ARCH6	Some FIDevelopersdeveloped User Acceptance Tests to evaluate the regions and complained about the availability of some services	A set of tests in order to verify periodically the status of the nodes has been developed.
ARCH7	The Identity Management (and KP) should be located in each node (and not only dependent from another in another country)	See ARCH 4

Table 28: Lesson Learned Architecture

5.1.5 Maintenance and Operation

This group of lessons learned are related to maintenance and operations activities.

ID	Lesson learned	Analysis
OPMAINT1	Documented installation steps should transformed into scripts in order to automatize and save time during operation, maintenance and re-installation.	Invite and help IOs to create and share their automatic tools and scripts for maintenance and re-installation

ID	Lesson learned	Analysis
OPMAINT2	The deployment of the monitoring component has been upgraded several times at the beginning of the federation deployment, with not a clear upgrade strategy and test yet.	1) This concern is related to the maturity of ITBox for the first deployments. 2) Unfortunately the strategy for sub-system release update was not defined from the beginning and this has affected the first nodes deployment 3) IOs should implement a pre-production dedicated platform.
OPMAINT3	Despite the FIWARE Lab Infographic provides OpenStack services, the Federation Manager do not have access to a detailed and historical availability of the API status of the regions.	The monitoring APIs are currently used by the monitoring dashboard in order to provide both current and historical data for each VM. Before the end of the project also historical data on the availability of the physical hosts will be provided.

Table 29: Lesson Learned Maintenance and Operation

5.2 Sharing best practises: FAQ

In this paragraph is reported the FAQ that was elaborated in coordination with WP2, WP3 and WP4 for the inclusion of new nodes and updated at M24 of XIFI Project.

This FAQ is the outcome of a continuous technical discussion since the beginning of the inclusion process with new nodes.

The FAQs are partitioned in the following topics:

- Documents
- Software Kits
- Connectivity and Network
- Cloud Installation
- Cloud Management (GEs)
- Monitoring
- Security/FIWARE Lab Joining
- Grizzly issues
- Icehouse issues

- Juno issues

Each of these topics is structured hereafter under the form question/answer.

5.2.1 XIFI Reference documentation

- **Where is the XIFI Documentation?**

  Official Documentation and the final version of the deliverables can be found at https://www.fi-xifi.eu/publications/deliverables.html

- **What are the main documents to read to deploy a new node?**

  Here is a short list of what documents should be read and why:

  - *D1.1: Big picture of the project. It provides useful information to be able to understand the rest of the deliverables.*
  - *D5.1: Entry point for all WP5 related tasks. Here you will find a description of all OpenStack modules, the definition of the reference architecture for deployment and lots of references to interesting documentation.*
  - *D5.2: This document describes the deployment plan for network backbone connectivity, and OpenStack on the five legacy nodes. You will also find some guide lines in order to join the federation and the perform basic service testing .*
  - *D2.1: Deployment How-To (handbook) with manual and automated (ITBox) OpenStack deployment procedures.*

  Optionally, the following deliverables can also be useful:

  - *D2.2: APIs and Tools for Infrastructure Federation v1.*
  - *D3.4: Federation Network adaptation mechanism. It describes the components that will allow to handle the tenant network orchestration based on the SDN capabilities of Openflow.*

  Before to perform the deployment, it can be useful to read the Fuel pre-installation guide at http://docs.mirantis.com/fuel/fuel-4.1/pre-install-guide.html. It describes pretty well the ways to deploy the cloud.
  Also can be useful to have a look at: http://wiki.fi-xifi.eu/Xifi:Wp5:how_to_deploy_a_xifi_node

- **Is there some videos about XIFI installation and configuration or XIFI in general ?**

  Some links on XIFI Channel of Youtube:

  - *Architecture of Federated Platform (Madrid 2014) http://youtu.be/Cqn92GvGIVc*
  - *Monitoring Webinar (Madrid 2014) http://youtu.be/xS1hcsLRkDc*
  - *Infrastructure ToolBox (ITBox)(Madrid 2014) http://youtu.be/Yng3m2uVNxw*
  - *Components That You Need (Madrid 2014) http://youtu.be/FT2is8IfaVo*
  - *Marketplace and Resource Catalogue http://youtu.be/L6KmyaXdJq8*
  - *Process for Joining the Federation http://youtu.be/gMaLbZMqm98*
  - *ITBox & Monitoring Webinar http://youtu.be/lHcJA2w07nY*

- **Where is further documentation available?**

  - *The slides presented by Alessandro Martellone during the webinar "ITBox & Monitoring" (18/09/2014) are available at: https://bscw.fi-xifi.eu/bscw/bscw.cgi/d90678/ITBOX-Telco-webinar_new_nodes_20140918.pdf . You can find more information about what was discussed during the same webinar in the minutes (section "Question & Answer") at https://docs.google.com/document/d/1-IuMZkHj__Bh-XjpfnkfikpgoaF5SDBRaHPXy6ooFrs/edit*
  - *The slides presented by Alvaro Alonso during the webinar "Keystone Proxy and Security Configuration" (19/09/2014) are available at: https://bscw.fi-xifi.eu/bscw/bscw.cgi/d90902/Keystone_Proxy_Configuration_webinar_new_nodes_20140919.pdf . You can find more information about what was discussed during the same webinar in the minutes (section "Question & Answer") at https://docs.google.com/document/d/1e6WYr8YET_HAa9FIwmEAC2g1U8P7lu9VoxhQhSflawM/edit*

5.2.2 Software to perform the node installation and configuration

- **Where are the official XIFI kits references?**

  All the Software Kits are related with XIFI components. There are alist of public and available components,where you can find details about the architecture, installation guide, developer guide, user manuals and more.

- **Where are the official XIFI kits needed to deploy a new node ?**

  A new node, which want to be federated, needs to deploy the main component:

  - *Infrastructure Toolbox component: It is the principal component and supports the automated installation of the main components of a XIFI node. The download version and some configuration parameters are provided by the portal, following the registration of the new node. Monitoring and network adapters are included (or linked) in the ITBox distribution from the adapters repository, while the same applies to GEs and related software needed to complete the XIFI node installation. ITBox is distributed as an ISO image which contains an installer for ITBox Master Server. The following components are deployed automatically by the ITBox, using the graphical user interface:*

  - *Create a new OpenStack environment.*
  - *Deploy DataCenter Resource Management (DCRM) GE (multi-node).*
  - *Enable Quantum.*
  - *Deploy Nagios and adapters.*
  - *Deploy OpenStack Data Collector and the Context Broker GE.*
  - *NGSI Adapter.*
  - *Installation of Keystone Proxy and IDM.*

  The next versions of ITBox will include more components, adapters and tests, in order to simplify the federation of the new nodes, as the Galera Arbitrator and test scenarios from continuous integration. For more details, you can see the TBox description and the course of FIWARE Ops for Infrastructure, chapter "A guide to the Infrastructure Toolbox" in the eLearning platform
  Other components that should be installed and there are not still included in the ITBox:

- *Network Passive Monitoring (NPM) Adapter: it is in charge of managing the Network Passive Monitoring (NPM) of the XIFI Infrastructure Monitoring Middleware The NPM Adapter enables the inclusion of both resource elements to be monitored and the KPIs (Key Performance Indicators) to be collected via the SNMP protocol. The details are in the NPM Adapter description.*
- *Datacenter & Enablers Monitoring (DEM) Adapter: It is responsible for performing the monitoring (along with the necessary adaptation mechanisms) of the node and the Generic Enablers in terms of computing and storage resources. It is one of the components belonging to the XIFI Infrastructure Monitoring Middleware (XIMM), as it provides monitoring and adaptation functionalities. The details are in the DEM Adapter description.*
- *Network Active Monitoring (NAM) Adapter provides a multi-domain monitoring mechanism able to handle latency and bandwidth-related tests along a set of points of interest within the federated community.*
- *Security Prove (SIEM Agent) collects security events and send them to the Service Level SIEM server.*
- *Installation of Cloud Portal that needs to be installed in the different nodes.*

- **Where are the official XIFI kits needed to deploy a master node?**

  The Infrastructure Owners, who wants to be federated, don't need to deploy the components, which are related with the federation layer and the GUI Portal. These IOs only need to deploy the components related with the new node, since the federation layer is deployed and ready to be used. You can see the above question to know the necessary components in order to install a new node.
  Nevertheless, you can find in the architecture diagram the necessary components to deploy a master node, which covers the XIFI federation platform. The following list summarizes and indicates the links of these necessary components.

  - *Federated Identity Management aims at providing a federated management of identities for members of XIFI federation. This component is also distributed on all the nodes of the federation and stores identities of authorized users.*

  - *Federation Monitoring aims at providing a common framework for storing, aggregating and publishing the monitored data collected by the different monitoring adapters provided by the XIMM module. Part of this component is also distributed on all the nodes of the federation and elaborates monitoring data leveraging on big data analysis techniques.*

  - *Cloud Portal provides a support for the users of the XIFI cloud infrastructure and platform to manage their services and resources deployed in cloud. It is implemented in a form of a Web GUI following the example of the portals that today's common cloud infrastructure managers have.*

  - *Marketplace & Resource Catalogue. XIFI Portal represents the component enclosed in the XIFI Federation Platform from which XIFI end users will gain access to the services and tools provided by the federation. Therefore, each of these services shall be accessible through a graphical representation (a Graphical User Interface - GUI) supported by this portal. Resource Catalogue & Recommendation tool is in charge of matching the infrastructure catalogue (from where end users browse the available XIFI services and their status), with the business logic required to fulfil the operations.*

- *SLA Manager provides mechanisms to support service level agreements management in the federated environment, based on WS-Agreement specification. The component allows the direct interaction among the different actors through the graphical user interface and is designed to be part of XIFI portal.*

- *Federation Manager is the central registration point for new infrastructures and their services.*

- *Infographics and Status Pages provides information on the infrastructure capacities and status of FIWARE-Lab infrastructure services.*

- *Deployment and Configuration Adapter (DCA) provides deployment of multiple GEs and XIFI components upon XIFI infrastructure, check of available resources prior to the Deployment and persistency of information related to the deployed GE instances.*

- *Interoperability Tool supports development and testing of FIWARE based applications and services; in particular the tool focuses on interoperability problems.*

- *Security Dashboard supports the management and visualization of security incident-related events and data. Beside core functionalities offered by such a dashboard, the Security Dashboard will also provide reporting capabilities.*

- *Security Monitoring is responsible for the collection and correlation of security monitoring data. Part of this component (Security Prove) is also distributed on all the nodes of the federation and collect security events and send them to the Service Level SIEM server.*

- *Monitoring Dashboard shows to users and Infrastructures Owners live and historical data about the status of the Federation environment resources in XIFI.*

5.2.3 Connectivity and Network

- **The OpenCall specified that the XIFI traffic will run over IPv6, but the deliverables talk about IPv4. Would you mind commenting?**

  There are several OpenStack components that does not support (or partially support) IPv6. In particular, L3 Agent in Grizzly does not support IPv6 forwarding (http://docs.openstack.org/grizzly/openstack-network/admin/content/ch_limitations.html). A more detailed report can be found here: http://www.nephos6.com/pdf/OpenStack-on-IPv6.pdf.

- **Following from above - OpenStack does not support IPv6 - what does this mean for IPv6 support within the project?**

  We have identified several limitations during the first year of the project. Geant MD-VPN service is not fully IPv6 compliant for the moment and we didn't have the information about IPv6 compliance of all the GEs at the right time to prepare the initial deployment on IPv6. Never less it remains a goal of the project to provide an IPv6 federation backbone with dual stack access to services IPv4/IPv6

- **What version of OpenFlow the Switch shall support. Is it V1.3 or later? What is the time frame for OpenFlow usage in this project?**

  To simplify the deployment across heterogeneous infrastructures, the choice seems is oriented to an implementation at the source which means that Openflow rules will be pushed to the compute nodes OVS and not to the physical switch. OpenFlow capabilities of the switch is not needed.

- **For those (new) nodes that do not have MD-VPN support from their NREN, what is the alternative connectivity solution?**

  There two alternative solutions depending on the nature of the limitation. They are described in D5.2 and here is a short description :

  - *If the node is connected to an NREN not providing MD-VPN service, the solution proposed by GEANT is to provide a VPN proxy that will forward traffic between GEANT PoP and the NREN PoP where the node is located. The infrastructure will be considered as doing part of the MD-VPN by itself*
  - *In case the node is not connected to a NREN or if the MD-VPN deployment delay is too long, a point to point L3-VPN connectivity with a XIFI node (probably The Spanish Node) that use the MD-VPN will be required.*

- **Who provides the other end of the VPN connection for the L3-VPN solution?**

  The best choice seems to be the Spanish Node as they have a 10Gbps dedicated Link. Any other node connected to the MD-VPN is suitable of providing this access also

- **How is the MD-VPN going to be used by the OpenStack user VM instances? D5.2 is extremely short on details.**

  MD-VPN connectivity has 2 usages:

  - *Connecting the federation components (monitoring, authentication..), either physical or virtual, across the federation nodes*
  - *Allow the multi-site tenant deployment by creating GRE tunnels between sites for which there are instances belonging to the same tenant. D3.4 covers this kind of deployment*

- **What is the relation of the MD-VPN private IPv4 network and the OpenStack's internal networks (API, block,...)?**

  Each node can deploy the components at is will. There several illustrations on D5.2 of the legacy nodes deployment. My suggestion is to dedicate one external network for the federation connectivity and use one subnet (/24 should be enough) of the private federation network to connect all the OpenStack components. This way the node will be available from the rest of the federation through the MD-VPN

- **Which IP-Ranges should be configured for the internal OpenStack network?**

  Administrative, management, data and storage networks, as described in D5.2 can be handled with the datacenter addressing plan, although using the federation private addressing plan is allowed. However, I would rather recommend to use the datacenter addressing plan as this will simplify the OAM local tasks.

- **Are there any node specific constraints concerning the network (IP-addresses, ports)? Which?**

    - *Respect federation private addressing plan assignment, it is defined in the deliverable 5.2, section 3.2 "IP Address Plan"*
    - *Segregate XIFI MD-VPN routing from datacenter routing to prevent routing issues. It is recommended to implement a dedicated VRF for XIFI MD-VPN routing*
    - *Preserve the default ports in order to simplify configuration management (not mandatory)*

- **Who will provide the configuration, credentials or keys to connect a new infrastructure to XIFI root keystone?**

    Keystone Proxy actual configuration can be found at the article http://wiki.fi-xifi.eu/KeystoneConfigFile. Installation procedure and support contact can be found at http://wiki.fi-xifi.eu/Public:Security_Proxy.

- **How should a new infrastructure be visible from xifi by vpn - they should see all of the infrastructure or only the external ip address??**

    Only the external network to which the federation services are connected is required to be accessible through the MD-VPN and thus using the federation private addressing plan. Multi-site tenant connectivity requires also to be done via the MD-VPN implementation details are not available at the time of the writing

- **Should a new infrastructure configure the network topology basing on specified IP address pool?**

    Please check D5.2 and previous questions on this FAQ

- **How to setup addition of extra networks (eg: public ip addresses)?**

    Look at D5.2 . About adding secondary external networks, there is a wiki at http://wiki.fi-xifi.eu/Xifi:Wp5:_OpenStack_Customizations.

5.2.4 Cloud Platform installation

- **For a new node with 8 physical servers, we are thinking to the following deployment architecture:**

    - **1x ITBox node**
    - **3x Controller (including monitoring, network & storage) nodes**
    - **4x Compute nodes**

    Is it a feasible architecture plan (also including HA) or there is some better design?

    You could consider also another strategy: to combine two controllers in HA with a galera arbitrator and use the free physical server as monitoring node.

- **Should every new node have the HA or it is up to the node to decide?**

  The HA deployment is recommended for production nodes. For more information take a look at http://wiki.fi-xifi.eu/Xifi:Wp5:d51#Physical_Deployment_Models

- **What do you need to do prior to node deployment?**

  It's essential to read the listed documents at [FAQ Documents section](#) to get a clear idea of the architectures you can deploy and as described in the training you should complete the MD-VPN connection and prepare your hardware on the datacenter.

- **How can I set the nodes to PXE boot?**

  The servers are usually configured to look for PXE if no OS is present on the disk. You don't need other action than setting all nodes administrative network on eth0, so Fuel can act as PXE server.

- **How do we install Cloud Environment using ITBox tool?**

  All information regarding ITBox, including the installation manual, are available at http://wiki.fi-xifi.eu/Public:InfrastructureToolbox

- **Is there some useful scenario for using ITBox ?**

  The ITBox is a tool that simplifies the deployment, setup and operation of distributed cloud infrastructures related to FIWARE Lab. It has been designed to deploy several models: multi-node, multi-node with High Availability using two controllers, multi-node with High Availability using at least three controllers. These models are designed according to the best practices recommended by OpenStack specialists. You can find more information at https://bscw.fi-xifi.eu/pub/bscw.cgi/d58595/XIFI-D2.1-XIFI_Handbook_v1.pdf in the section "4.3.2 Automated Installation with Infrastructure Toolbox".

- **Where can we find some use cases about Cloud Installation?**

  You can find a basic deployment example at https://bscw.fi-xifi.eu/pub/bscw.cgi/d58595/XIFI-D2.1-XIFI_Handbook_v1.pdf in the section "4.3 Installing OpenStack and Cloud Service".

- **Is there some configuration reference for public flavors?**

  m1.small:
  Memory: 2048MB,

  VCPUS: 1,

  Root: 10GB,

  Ephemeral: 20Gb,

  Swap: 0MB,

  RXTX Factor: 1.0

  m1.tiny:

  Memory: 512MB,

  VCPUS: 1,

  Root: 0GB,

Ephemeral: 0Gb,

Swap: 0MB,

RXTX Factor: 1.0

Quotas:

metadata_items: 128

volumes: 10

gigabytes: 50

ram: 25000

security_group_rules: 30

instances: 3

fixed_ips: 10

security_groups: 20

injected_file_content_bytes: 10240

floating_ips: 1

injected_files: 5

cores: 6

quota_metadata_items=128

quota_volumes=10

quota_gigabytes=50

quota_ram=25000

quota_security_group_rules=30

quota_instances=3

quota_security_groups=20

quota_injected_file_content_bytes=10240

quota_floating_ips=1

quota_injected_files=5

quota_cores=6

quota_fixed_ips=10

5.2.5 GEs, images and blueprints installation

- **Which are the suggested GEs to install in a new node?**

  The system administrator of an official node does not need to install any GE image. This is the work of the synchronization tool, which is run on the master region.

- **Do I need to download any other image?**

  Both GE images and base images are uploaded to your node by the synchronization tool, which is run in the master region. All images are uploaded to your node using the admin account.
  Of course, you can provide also your own extra images, but these images will not be synchronized to other regions. Only images from master region are synchronized.

- **How do we register nid property of GE images in Glance metadata? How do we register image to work properly with PaaS Manager/SDC Manager in Glance metadata?**

  This work is done in the master region. Some metadata is registered and added as attributes in the images. This metadata is synchronized with the image content to each region from the master region.

- **How do we register image to work properly with PaaS Manager/SDC Manager in Glance metadata?**

  Actually, you do not need to register image. You just provide the corresponding image to the master node in order to deploy automatically to all the nodes. In case that you want to maintain only in your FIWARE Lab node but you want to use it with PaaS Manager and SDC Manager, you need to install cloud-init.
  To install cloud-init, just execute the command:

  ```
  $ yum install cloud-init
  ```

  In case of RedHat, Fedora or CentOS linux, in case of Ubuntu, Debian you should execute:

  ```
  $ apt-get install cloud-init
  ```

  The PaaS Manager will inject automatically in the image the user-data in order to do the operations that it is needed in order to work with it.

  Besides, you need to specify that your image can be used by the PaaS Manager and SDC through the specification of the metadata parameters:

Configuration of Glance Metadata	
SDC-aware Image name	**Glance command**
MyNewImage	$ glance image-create --name MyNewImage --disk-format qcow2 --container-format ovf --size 4107534336 --min-disk 0 --min-ram 0 --is-public True --is-protected False --property sdc_aware=Yes --file <name of the file of the corresponding downloaded image>

Table 30: Configuration of Glance Metadata

- **What have we to do if we want to use Ceph? How can I change the format of the image from qcow2 to raw?**

  In case that you want to install Ceph and not Swift like your object storage system, you should translate the format of the image from the qcow2 to raw format due to it is strongly recommended to use this image format. Taking into account and assuming that you have kvm-img tool, you can use the quemu-img to do the conversion from qcow to raw format:
  `$ qemu-img convert image.qcow image.raw`

  It is also possible to upload qcow images to Ceph and then convert to raw: http://www.sebastien-han.fr/blog/2014/11/11/openstack-glance-import-images-and-convert-them-directly-in-ceph/

5.2.6 Monitoring installation

- **How to setup the deployment of monitoring?**

  The easiest method is to use a dedicated node and deploy through the ITBox. In case of French node, the monitoring is running on VMs, but, at the moment, there is not a complete configuration guide.

- **Can the monitoring VMs be deployed through OpenStack?**

  It depends what you mean by "Monitoring VMs". If it means "monitored VMs" then the answer is yes: the monitor plug-in will be automatically embedded in each VM image so as to deploy it together with the VM.

- **How can I install Monitoring manually?**

  Some components should be installed:

  1. Nagios
  2. NRPE on each physical host
  3. Context Broker (v13)
  4. NGSI Adapter (v1.1.1)
  5. OpenStack Data Collector (v2.1)

  Moreover everything should be configured properly. Details can be found in the wiki pages and in the deliverables D2.1, D2.2 and D3.2.

- **Do I need the MD-VPN connection?**

  Yes, MD-VPN connection is required. In order to test it, try to reach the 10.0.32.20 (the IP of the Trento node)

- **Which ports should be opened in the Monitoring?**

  The monitoring node should have the 1026 and 1337 ports open. These are the default port for the CB and ngsi_adapter

- **Why I don't see my region in the infographic status page?**

  The monitoring APIs, which provide information to the infographic status, drop all the regions that have not retrieved new information for 3 hours. This problem is caused by the ODC status. If the ODC is not working, no new info is collected

- **Why do I see my region with wrong data (VM,etc..)?**

  This problem is caused by the wrong credentials in the odc.conf. If the credentials are not the authorized (federated), the ODC cannot collect the proper info

- **How can I install the new version 1.1.1 of NGSI adapter that supports timestamps?**

  You can find at https://forge.fiware.org/frs/?group_id=7&release_id=529#cloud-monitoring-3-5-2-title-content the new NGSI Adapter 1.1.1 including the timestamp feature and integrating all the parsers needed to process monitoring raw data.

  - *Installation manual at https://github.com/telefonicaid/fiware-monitoring/tree/develop/ngsi_adapter assumes a package repository is already configured so packages can be installed with tools such as "apt-get". Given we have no repository at the moment, as an alternative IOs can download such file and install the package with tools like "gdebi" (also installs dependencies) or "dpkg".*
  - *As described at https://github.com/telefonicaid/fiware-monitoring/tree/develop/ngsi_adapter#usage, now a Linux service "ngsi_adapter" is setup (but no started) by package installer. Installation directory has been normalized to /opt/fiware/{component}, so it'd be /opt/fiware/ngsi_adapter in this case. Directory for log files has been also normalized, and by default would be /var/log/ngsi_adapter/ .*
  - *This package supports Ubuntu 12.04 LTS distribution. The node.js "official" version for such distro is too old: during package installation the version on node.js is checked and, if an upgrade is needed, installer prepares system files to get newest version from NodeSource (and not from Chris Lea's PPA, as before).*

- **What are the processes to be monitored? And what's the command to check the processes status?**

  The following is the configuration of 'nrpe.cfg' Nagios on a controller at the Trento node. In order to run the monitoring on Status page, you need to monitor the processes listed below. Obviously, the command could be different on other nodes depending on the installation.

  command[check_nova_scheduler]=/usr/local/nagios/libexec/check_procs -c1:1 -C nova-scheduler

  command[check_nova_api]=/usr/local/nagios/libexec/check_procs -c1: -C nova-api

  command[check_nova_novncproxy]=/usr/local/nagios/libexec/check_procs -c1:1 -C nova-novncproxy

  command[check_nova_conductor]=/usr/local/nagios/libexec/check_procs -c1:1 -C nova-conductor

  command[check_nova_cert]=/usr/local/nagios/libexec/check_procs -c1:1 -C nova-cert

  command[check_nova_consoleauth]=/usr/local/nagios/libexec/check_procs -c1:1 -C nova-consoleaut

  command[check_nova_objectstore]=/usr/local/nagios/libexec/check_procs -c1:1 -C nova-objectstor

  command[check_cinder_scheduler]=/usr/local/nagios/libexec/check_procs -c1:1 -C cinder-schedule

  command[check_cinder_api]=/usr/local/nagios/libexec/check_procs -c1:1 -C cinder-api

  command[check_quantum_openvswitch_agent]=/usr/local/nagios/libexec/check_procs -c1:1 -C python -a quantum-openvswitch-agent

  command[check_quantum_l3_agent]=/usr/local/nagios/libexec/check_procs -c1:1 -C python -a quantum-l3-agent

  command[check_quantum_dhcp_agent]=/usr/local/nagios/libexec/check_procs -c1:1 -C python -a quantum-dhcp-agent

command[check_quantum_server]=/usr/local/nagios/libexec/check_procs -c1:1 -C python -a quantum-server

command[check_quantum_metadata_agent]=/usr/local/nagios/libexec/check_procs -c1:1 -C python -a quantum-metadata-agent

command[check_glance_api]=/usr/local/nagios/libexec/check_procs -c1: -C glance-api

command[check_glance_registry]=/usr/local/nagios/libexec/check_procs -c1: -C glance-registry

- **How can finalize the Monitoring configuration after the federation (Keystone proxy cfg) ?**

  To finalizing monitoring configuration (file odc.conf) you need to have authorization infos from Keystone Proxy cfg people (Alvaro aalonsog@dit.upm.es), you need user, password and tenant-id to configure the odc.conf e.g.:

  username = IDMUser (coming from federation)

  password = IDMPwd (coming from federation)

  tenant_name = IDMTenantName (coming from federation)

  auth_url = http://130.206.82.10:4730/v2.0/

  regionName = <YourRegionID>

  regionId = <YourRegionID>

  location = <YourState (ISO 3166-1 :alpha2)>

  latitude = <00.00 use this format>

  longitude= <00.00 use this format>

  agentUrl= <IPwhereNSGI_AdapterIsInstalled>:<port(default1337)>

- **Where can I find some cfg examples of monitoring components and a simple guide for monitoring installation and configuration?**

  here: https://bscw.fi-xifi.eu/bscw/bscw.cgi/91263

- **Is there some step by step instructions to configure Metadata Services in case you have used ITBox to deploy OpenStack?**

  If you used ITBox/FUEL to deploy OpenStack, you've maybe experienced metadata-server cannot be reached by VMs.
  Basically, using Quantum with overlapping IPs, there is a quantum proxy process for each namespace (working on port 9697), the so-called ns-metadata-proxy, that redirects the traffic to nova metadata (working on port 8775) through quantum-metadata-agent.

  Here you can find an explanation with a valuable figure http://techbackground.blogspot.it/2013/06/metadata-via-quantum-router.html

  ---

  STEP#1

  You have to configure the file /etc/quantum/l3_agent.ini modifying this values:

  metadata_port = 9697 (instead of 8775)

  metadata_ip=<ip_br-mgmt> (instead of 169.254.169.254)

Where ip_br-mgmt represent the IP of the related controller (e.g., 192.168.0.3).

[For HA Nodes only]: Using HA you have to change the configuration in all the controllers, putting the corresponding IP related to the OVS bridge br-mgmt of the server.

STEP#2

Then you have to kill all the ns-metadata-proxy with the wrong port (8775), e.g.,

ps aux | grep metadata

root 28665 0.0 0.1 90860 22504 ? S :Jun23 0:00 python /usr/bin/quantum-ns-metadata-proxy :--pid_file=/var/lib/quantum/external/pids/35897224-6e8d-4e11-afff-7b3f6716d923.pid --log-file=/var/log/quantum/ns-metadata-proxy.log :--router_id= d2e9c2bb-a725-4397-9cec-35c7d0f55bbd --state_path=/var/lib/quantum --metadata_port=8775 –verbose

Restart quantum-l3-agent (NOTE: use crm CLI if your deployment is in HA)

Restart quantum-metadata-agent (NOTE: use crm CLI if your deployment is in HA)

Restart quantum-server on the controllers

STEP#3

Now check if the all ns-metadata-proxy have been restarted with the previously configured port (9697), e.g.,:

root 27034 0.0 0.1 91524 23728 ? S Jun26 0:00 python /usr/bin/quantum-ns-metadata-proxy --pid_file=/var/lib/quantum/external/pids/d2e9c2bb-a725-4397-9cec-35c7d0f55bbd.pid --log-file=/var/log/quantum/ns-metadata-proxy.log --router_id=d2e9c2bb-a725-4397-9cec-35c7d0f55bbd --state_path=/var/lib/quantum --metadata_port=9697 --verbose

---

STEP#4

The last step is to try to deploy a VM, access on it and look if the metadata is working properly, e.g.,:

curl http://169.254.169.254/openstack

If it works you have such type information of output:

curl http://169.254.169.254/openstack

2012-08-10

2013-04-04

latest

5.2.7 Federated identity management

- **What is needed to setup federated IDM, what components have to be installed?**

  In order to setup the Federated IDM you have to install an instance of KeyRock and an instance of Security Proxy. Then you have to connect both. The instructions and manuals are here:

  - *Federated IdM: http://wiki.fi-xifi.eu/Public:Federated_Identity_Management*
  - *Security_Proxy: http://wiki.fi-xifi.eu/Public:Security_Proxy*

- **What is needed to federate my node?**

  Before Federating be sure that your node:

  - *can deploy a VM.*
  - *can deploy an image (deploy and image = register a image in glance).*
  - *can deploy a network.*
  - *can attach a floating ip to a VM and it is reachable from outside.*

  Then you have to send the public endpoints of your node to the Identity Manager administrators. They will provide you the needed credentials to perform admin tasks.

- **Where can we find some examples of Keystone Proxy configuration**

  You can go to the documentation of the Security Proxy Security Proxy where you can find a description of the configuration of the Keystone Proxy.

- **After the federation how can I administer my OpenStack platform?**

  After the federation you need to use CLI as described in chapter 6 of D5.2 https://bscw.fi-xifi.eu/pub/bscw.cgi/d64414/XIFI-D5.2-XIFI_Core_Backbone.pdf

- **How to take advantage of the Security components in my applications?**

  In FIWARE Academy website you will find some courses of how to use security http://edu.fiware.org/course/category.php?id=8

5.2.8 Grizzly issues

- **Create a VM with no networks result to an AttributeError**

  If no network is configured and the tenant tries to create an instance the whole procedure crashes with the following error:

  2014-10-22 09:39:59.412 WARNING nova.network.quantumv2.api [req-5199dcf4-590a-4054-b21b-f4ec6f1562e0 admin 00000000000000000000000000admin] [instance: 115353fe-4c58-451a-93db-8a8b7ba1ebe3] No network configured!
  2014-10-22 09:40:00.491 ERROR nova.compute.manager [req-5199dcf4-590a-4054-b21b-f4ec6f1562e0 admin 00000000000000000000000000admin] [instance: 115353fe-4c58-451a-93db-8a8b7ba1ebe3] ***Instance failed to spawn***
  2014-10-22 09:40:00.491 2615 TRACE nova.compute.manager [instance: 115353fe-4c58-451a-93db-8a8b7ba1ebe3] Traceback (most recent call last):
  2014-10-22 09:40:00.491 2615 TRACE nova.compute.manager [instance: 115353fe-4c58-451a-93db-8a8b7ba1ebe3] File "/usr/lib/python2.7/dist-packages/nova/compute/manager.py", line 1118, in _spawn
  2014-10-22 09:40:00.491 2615 TRACE nova.compute.manager [instance: 115353fe-4c58-451a-93db-8a8b7ba1ebe3] self._legacy_nw_info(network_info),
  2014-10-22 09:40:00.491 2615 TRACE nova.compute.manager [instance: 115353fe-4c58-451a-93db-8a8b7ba1ebe3] File "/usr/lib/python2.7/dist-packages/nova/compute/manager.py", line 703, in _legacy_nw_info
  2014-10-22 09:40:00.491 2615 TRACE nova.compute.manager [instance: 115353fe-4c58-451a-93db-8a8b7ba1ebe3] network_info = network_info.legacy()

2014-10-22 09:40:00.491 2615 TRACE nova.compute.manager [instance: 115353fe-4c58-451a-93db-8a8b7ba1ebe3] ***AttributeError: 'list' object has no attribute 'legacy'***
2014-10-22 09:40:00.491 2615 TRACE nova.compute.manager [instance: 115353fe-4c58-451a-93db-8a8b7ba1ebe3]
2014-10-22 09:40:00.579 AUDIT nova.compute.manager [req-5199dcf4-590a-4054-b21b-f4ec6f1562e0 admin 00000000000000000000000000admin] [instance: 115353fe-4c58-451a-93db-8a8b7ba1ebe3] ***Terminating instance***
2014-10-22 09:40:02.055 2615 ERROR nova.virt.libvirt.driver [-] [instance: 115353fe-4c58-451a-93db-8a8b7ba1ebe3] During wait destroy, instance disappeared.
2014-10-22 09:40:02.059 INFO nova.virt.libvirt.driver [req-5199dcf4-590a-4054-b21b-f4ec6f1562e0 admin 00000000000000000000000000admin] [instance: 115353fe-4c58-451a-93db-8a8b7ba1ebe3] ***Deleting instance files*** */var/lib/nova/instances/115353fe-4c58-451a-93db-8a8b7ba1ebe3*

The bug description and the corresponding fix can be found here:
https://bugs.launchpad.net/nova/+bug/1222781
https://github.com/openstack/nova/commit/d367ab62720988231524accab33488ed1d6c8555
https://github.com/openstack/nova/commit/0dbcdc87d86031c9e3af29a72b830bda57b1161c

After applying the changes to every node, restart the services:

- *Controller node*

  ls /etc/init.d/ | grep nova | while read x;do service $x restart;done

- *Compute nodes*

  service nova-compute restart

After the fix is applied the creation of the instance returns with a success status
2014-10-24 14:41:21.047 WARNING nova.network.quantumv2.api [req-53d67507-e17d-4eb2-820b-a769c3339edd admin 00000000000000000000000000admin] [instance: 2829d2e8-df46-433a-a12b-24345bce8d9f] ***No network configured!***
2014-10-24 14:41:22.309 INFO nova.virt.libvirt.driver [req-53d67507-e17d-4eb2-820b-a769c3339edd admin 00000000000000000000000000admin] [instance: 2829d2e8-df46-433a-a12b-24345bce8d9f] ***Creating image***
2014-10-24 14:41:22.431 INFO nova.virt.libvirt.driver [req-53d67507-e17d-4eb2-820b-a769c3339edd admin 00000000000000000000000000admin] [instance: 2829d2e8-df46-433a-a12b-24345bce8d9f] Injecting key into image ad62f2ed-b688-4ebf-b2e4-eac89fbd286f <code>2014-10-24 14:41:24.994 35544 INFO nova.compute.manager [-] Lifecycle event 0 on VM 2829d2e8-df46-433a-a12b-24345bce8d9f
2014-10-24 14:41:25.006 35544 INFO nova.virt.libvirt.driver [-] [instance: 2829d2e8-df46-433a-a12b-24345bce8d9f] ***Instance spawned successfully****.*

- **DHCP agent fails**

```
Command: ['sudo', 'quantum-rootwrap', '/etc/quantum/rootwrap.conf', 'ip', 'netns',
'exec','qdhcp-efaaffc5-314a-40c7-bfef-05a18cee090b', 'kill', '-9', '14659']
 Exit code: 1
Stdout:
Stderr: 'Cannot open network namespace: No such file or directory\n'
 2014-10-21 16:39:04     INFO [quantum.agent.dhcp_agent] Synchronizing state
 2014-10-21 16:39:04    ERROR [quantum.agent.dhcp_agent] Unable to disable dhcp.
Traceback (most recent call last):
File "/usr/lib/python2.7/dist-packages/quantum/agent/dhcp_agent.py", line 131, in
call_drivergetattr(driver, action)()
```



```
File "/usr/lib/python2.7/dist-packages/quantum/agent/linux/dhcp.py", line 139, in
disableip_wrapper.netns.execute(cmd)
File "/usr/lib/python2.7/dist-packages/quantum/agent/linux/ip_lib.py", line 414, in
execute check_exit_code=check_exit_code)
File "/usr/lib/python2.7/dist-packages/quantum/agent/linux/utils.py", line 61, in
executeraise RuntimeError(e)
```

Solution A
This is a known and old issue regarding the dhcp agent as stated below:
https://bugs.launchpad.net/neutron/+bug/1052535
The solution can be found in here:
https://answers.launchpad.net/neutron/+question/206604

- *Reboot the controller*
- *quantum net-list*
- *quantum subnet-list*
- *ip netns list*
- *ip netns exec <id_of_qdhcp> bash - the id should much one of the networks id that has dhcp enabled and is also enlisted in the netns list*

Solution B
Possibility of a network conflict is causing the problem. The recommended approach is to recreate the namespaces and taps and remove the networks that create the problem(if the network cannot be removed due to port usage for example, check the database)
Command recommendations:

- *ovs-vsctl add-port br-int tap83384c19-b9*
- *ip netns add qdhcp-efaaffc5-314a-40c7-bfef-05a18cee090b*
- *ip netns list - verify that the namespaces are present*

5.2.9 Icehouse issues

- **How to configure the MDVPN using Icehouse? Is it the same as Grizzly?**

In OpenStack Grizzly to add a second external network (e.g., MDVPN) it is necessary to add a second layer-3 agent and bind each Layer-3 agent to its respective network. This caused some troubles especially concerning metadata agents and high availability scenarios. In OpenStack Icehouse there is a way to attach the secondary external network to one L3-agent, using on-link routes. So it is suggested to use this approach.

- **After the migration to Icehouse there are some issues on the XIFI node. I can create private-networks, but they are in "down state" and I cannot attach VMs to them.**

Check if Neutron server and agents are up (i.e., neutron agent-list) and restart them. If the problem still remains check the logs.

- **Using Ceph as Glance backend I noticed that VMs deployment is too slow.**

Ceph does not support QCOW image format. You have convert the OpenStack images in RAW format (e.g. qemu-img convert -f raw -O qcow2 centos64.dsk centos64.qcow2)and re-upload them in Glance.

- **I have some issues with the Orion context Broker after configuring the Monitoring Module.**

Change the MongoDB configuration, deactivating the authentication, i.e., use auth=false option.

5.2.10 Juno issues

- **Clean vxLANs network**

We need sometimes the configuration of some vxLANs network during the OpenStack (Juno) installation in order to test that the Neutron server is running properly. It should produce the creation of endpoints in the vxLAN tunnel and the final problem that they are mixed and disappeared. We need to clean the testing network in order to resolve this issue. If we execute the following commands:

```
# ovs-vsctl  show
 ....
   Bridge br-tun
   ......
      Port "vxlan-0a000401"
         Interface "vxlan-0a000401"
            type: vxlan
            options: {df_default="true", in_key=flow, local_ip="10.0.3.19", out_key=flow,
remote_ip="10.0.4.1"}
      Port "vxlan-0a000413"
         Interface "vxlan-0a000413"
            type: vxlan
            options: {df_default="true", in_key=flow, local_ip="10.0.3.19", out_key=flow,
remote_ip="10.0.4.19"}
      Port "vxlan-0a000301"
         Interface "vxlan-0a000301"
            type: vxlan
            options: {df_default="true", in_key=flow, local_ip="10.0.3.19", out_key=flow,
remote_ip="10.0.3.1"}
 ...
```

We could see that in that case the ports with the IPs 10.0.4.x should not be there. Nevertheless, we cannot delete manually due to OpenStack will create it afterward because those ports are fisically in the Neutron DB.

```
mysql> show tables where Tables_in_neutron like '%vxlan%';
+-----------------------------------+
| Tables_in_neutron                 |
+-----------------------------------+
| cisco_n1kv_vxlan_allocations      |
| ml2_vxlan_allocations             |
| ml2_vxlan_endpoints               |
+-----------------------------------+
3 rows in set (0.00 sec)

mysql> select * from ml2_vxlan_endpoints;
+------------+------------+
| ip_address | udp_port   |
+------------+------------+
```

```
| 10.0.3.1    |      4789|
| 10.0.3.19   |      4789 |
| 10.0.4.1    |      4789 |
| 10.0.4.19   |      4789 |
+-------------+-------------+
4 rows in set (0.00 sec)
```

We could see the different IPs associate to the ml2_vxlan_endpoints and excatly the two IPs that we do not like to see there (10.0.4.x). We delete manually from the DB there IPs

```
mysql> delete from ml2_vxlan_endpoints where ip_address like '10.0.4.%';
Query OK, 2 rows affected (0.00 sec)
```

Last but not least, we need to restart the ml2 plugin:

```
# service neutron-plugin-openvswitch-agent restart
```

And we have no more those IPs in the vxLANs.

- **Nothing work without Security Group**

We have detected in Juno, maybe it should be the same in versions after Essex, that nothing work by default if we do not use security groups.

- **SSH command halts with no output**

This problem is described in https://ask.openstack.org/en/question/30502/can-ping-vm-instance-but-cant-ssh-ssh-command-halts-with-no-output/

Firstly, we need to ensure that dnsmasq process is running.

```
# pgrep -fl dnsmasq
6994 dnsmasq
24516 dnsmasq
25029 dnsmasq
```

Secondly, we have to get details about the virtual machine in order to know what is the network that we are using.

```
$ nova list
+--------------------------------------+------------+--------+---
---------+-------------+----------------------+
| ID                                   | Name       | Status |
Task State | Power State | Networks             |
+--------------------------------------+------------+--------+---
---------+-------------+----------------------+
| 62f7e6c3-f15a-4127-9cb3-3930a1abe6bc | myserver1  | ACTIVE | -
| Running     | demo-net=192.168.1.9  |
+--------------------------------------+------------+--------+---
---------+-------------+----------------------+
```

We can see that the name of the network is demo-net, the next step is to obtain details about the different networks that we have in order to obtain the identifier of that network.

```
# neutron net-list
+--------------------------------------+----------+--------------
--------------------------------------+
| id                                   | name     | subnets
|
+--------------------------------------+----------+--------------
--------------------------------------+
| 9babfff6-7553-4947-a709-e4a419e6fb6f | ext-net  | d15dd853-
09aa-4dc9-81a8-f2166bfdc17e                 |
| e9dcf592-c282-49ca-9c75-2ace3dd0997e | demo-net | 6795f2f2-
a103-4602-ab07-c5ea0ac737e7 192.168.1.0/24 |
+--------------------------------------+----------+--------------
--------------------------------------+
```

We can identify now what is the qdhcp-network <networkUUID> namespace corresponding to the demo-net network. We can check is really this namespace exists by the execution of the following command.

```
$ ip netns | grep qdhcp
qdhcp-706dd746-e059-4476-b780-e2dcda4ecfd8
qdhcp-90689a39-75cb-450b-a43e-f647186f59c2
qdhcp-e9dcf592-c282-49ca-9c75-2ace3dd0997e
```

From the list of namespaces, we can obtain the corresponding one to the network (demo-net) of the virtual machine (myserver1). the next step, we can check the ssh connectivity from the controller machine running the following command:

```
# ip netns exec qdhcp-e9dcf592-c282-49ca-9c75-2ace3dd0997e ssh -vvv cirros@192.168.1.9
22
 OpenSSH_6.6.1, OpenSSL 1.0.1f 6 Jan 2014
 debug1: Reading configuration data /root/.ssh/config
 ...
 debug2: mac_setup: setup hmac-md5
 debug1: kex: client->server aes128-ctr hmac-md5 none
 debug2: bits set: 1032/2048
 debug1: sending SSH2_MSG_KEXDH_INIT
 debug1: expecting SSH2_MSG_KEXDH_REPLY
```

We could see that in that point the SSH halts. The problem is related to the size of the MTU packets. In case of the example (Cirros image), if we change this value by 1400, the problem is resolved. To change this value, just execute the following command.

```
# sudo ip link set eth0 mtu 1400
```

And now, the ssh connectivity works again, but we need to update the neutron configuration in order to reflect this changes too. We need to edit the file ***etc/neutron/dhcp_agent.ini***

```
[DEFAULT]
 ...
```

```
dnsmasq_config_file = /etc/neutron/dnsmasq-neutron.conf
...
```

*and in the file **/etc/neutron/dnsmasq-neutron.conf** we need to chang the following attribute:*

```
...
 dhcp-option-force=26,1400
...
```

Finally, we need to kill the dnsmasq process and restart the neutron-dhcp-agent

```
$ killall dnsmasq
$ service neutron-dhcp-agent restart
```

- **Kernel Panic in Neutron**

This problem is explained in https://ask.openstack.org/en/question/29147/ssh-to-a-vm-causes-kernel-panic-on-icehouse-neutron-host/. OVS module produce a Kernel panic in Neutron with vxLANs. The procedure to resolve it is executing the following commands.

```
# ethtool -K eth0 gro off
# ethtool -K eth0 gso off
```

6 PLAN AND RECOMMENDATIONS FOR THE FUTURE

From the experience gained within the federation process and the support activities in the lab, we can summarize the problems that affected the activity and the availability of resources in the FIWARE Lab three main issues:

- **Technical complexity.** Complexity of mastering OpenStack and immaturity of some OpenStack components caused several operational issues to the nodes.
- **Resource availability.** Due to the underestimation of the external usage of FIWARE Lab resources by unknown parties, the Lab is often short in resources.
- **Heterogeneity.** The freedom of nodes to purse different configuration within the FIWARE Lab, leads in some cases to confusion among the users.

The issues, on-going mitigation plans and actions are discussed in the following paragraphs.

6.1 Technical complexity

OpenStack is composed of several components that are not easy to master. Clearly several node operators underestimated this, causing a delays and congestions in the deployment phase. On the other side, the deployment of eleven additional nodes into the federation, and their opening up to a large number of external users, pushed to the limits the release of OpenStack that initially was supported by FIWARE (Grizzly). Indeed, Grizzly was the first release where Neutron was included, and Neutron has been one of the component creating more stability issues (especially in the high-availability configuration). Nova-Network, adopted by CESNET instead of Neutron, proved to be much more stable.

A number of actions to improve the technical understanding of OpenStack by operators were put in place within the project since its beginning. Still, while the background of some infrastructures improved to the needed level, in other cases, this is still lacking beyond.

6.1.1 On-going actions

To resolve technical issues related to the complexity and immaturity of components four main actions were enacted:

1. **Migration of nodes to more stable released of OpenStack** (i.e. Ice House or newer). With the availability of more stable (and performing) releases of OpenStack, the decision was to replace Grizzly with up-to-date stable and supported releases. The activity is ongoing since December 2014; time to upgrade depends on the capacity of the nodes to provide a parallel environment with needed resources to migrate users. Unfortunately, the maturity of OpenStack still discourages the upgrade of an existing deployment. The detailed plans for each infrastructure are included in the maintenance activities described in D5.5. By end of May 2015 all nodes are planned to run on Ice House or Juno.
2. **Uniform network configuration of OpenStack.** OpenStack, beyond providing different network management tools, enables multiple configurations. Initially within the project provided only a general reference configuration, but nodes had a certain degree of freedom. The different implementation of networks by different nodes, made more complex support operation to the nodes. Thus increasing the number of faced issues. Currently all nodes, except the ones not running Neutron (Spain and Prague) have a common network set-up.
3. **Health-checks to improve automatic assessment of OpenStack status**. The complexity of OpenStack is challenging. It is not enough to monitor if OpenStack services are up to ensure that the system is running properly. Infrastructure operators need to assess their infrastructure daily. To facilitate this task, a number of automated tests were developed, tests were run on daily basis, and a quite effective view on status of nodes was provided.

4. **Reward / penalise nodes according to performances**. To make the nodes more responsive and target improvements of their service quality, a karma mechanism was defined. The karma points relates to results of automatic checks, usage by developers of the node, number of open tickets, and other parameters. The mechanism proved to be effective in pushing nodes toward higher achievements. Unfortunately, the mechanism is not enough when the functionality of the nodes is compromised by other factors, such as: i) resource overload; ii) malfunctioning of central components (see recommendation 2).

6.1.2 Recommendations for the future

Following the already on-going actions, we recommend as well:

1. **Improving OpenStack skills**. It is recommended that nodes part of the Lab or that will be part in the future, provide solid skills in OpenStack operation. The continuous support by developers cannot replace node operators, and should be used only for actual software bugs. With the opening of the new Open Call, FI-Core should carefully selected only nodes having a successful track record in operating OpenStack infrastructures.
2. **Reduce single point of failures**. In the current deployment of the FIWARE Lab, all nodes connect to the same service catalogue and identity provider service (the so called keystone-proxy). This component proved to be a single point of failure in several situations causing disruptions in the FIWARE Lab. In the XIFI project, different solutions were proposed to solve this issue at an architecture level, but none was ever implemented so far due to privacy management constraints posed by FIWARE Lab identity managers. A new solution within the project was proposed leveraging on delegation to reduce the issue, but: 1) this only move the problem to a third component (even if through caching, access to the identity delegation provider will be much more limited); 2) this would require some components in the FIWARE Lab (namely the Cloud portal and the Blueprint manager) to change their authentication process. We recommend FI-Core not to ignore any further this issue, and proceed either supporting the developed solution, either providing an alternative solution.

6.2 Resource availability

Since its opening, FIWARE Lab has been accessible in a total self-provisioning modality without any restriction by any type of users without prioritization. This resulted in an overload of the nodes, in particular as regards public IPs. Beyond that, this caused a number of fraudulent actions to occur within the Lab without possibility to track real identities of authors of such actions.

Different solutions have been explored, including the potential adoption of IPv6, which unfortunately is not yet fully supported within OpenStack. Beyond that, node operators, without proper access to users' information and contacts, could not detect which ones to prioritise.

6.2.1 On-going actions

To resolve issues related to resource availability two main actions were enacted:

1. **Increase resources made available by nodes.** Nodes were requested to provide resources beyond the initially planned ones. i.e., most of the nodes almost doubled the number of public IPs available to end-users. But still, if such resources are not controlled, they get over allocated as well. Information on resources made available is presented in the maintenance activities reported by D5.5.
2. **Implement user access policies.** Since November 2014 a number of actions have been taken in order to provide a better control on resources. In particular, a key point has been recognized in the establishment of user access policies. The policies have been largely defined before end of 2014. Unfortunately, different issues faced during the deployment of policies slowed down

the process that was supposed to be enacted in February. The process has anyhow initiated through manual enactment and it will be completed between April and May with the deployment of a new release of the Identity Manager. We believe the completion of this action will be crucial into improving the quality of service to Phase III users.

6.2.2 Recommendation for the future

Following the already on-going actions, we recommend as well:

1. **Improve visibility to users of resources status.** While "resources over" does not mean that a system is not working properly (it only means it reached is maximum capacity), it is important that this information is transparently shared to end-users. FI-Core should improve the infographics developed within XIFI to enrich the information shared with users, includingthe number of available resources and number of available trial accounts.
2. **Improve nodes pro-activeness.** On one side XIFI (and FI-Core) should push nodes to be more pro-active in the management of resources, by releasing resources from low priority tenants and inactive ones. On the other, FI-Core should ensure that nodes can promptly access to contact information needed to provide better support to users and to assess the user right to use resources.
3. **Adopt IPv6.** As earlier mentioned, up to Juno, IPv6 is not fully supported in OpenStack. FI-Core should strive to contribute or push the completion of development of Neutron support for IPv6. As consequence nodes should migrate to IPv6 when OpenStack Neutron will introduce support for it.
4. **Change application deployment approach.** The problem of resources overload, it is not only related to the mismanagement of resources or to their scarcity. Novel PaaS tools, mostly building on containers (such as Cloudfoundry or Heroku), do not require users to have access to public IPs to deploy and run applications. Such tools adopt other solutions such as natting or reverse proxying in combination with advanced code deployment solutions not requiring ssh access to the applications (e.g. git code retrieving or file-upload based deployment). We encourage FI-Core to explore as well such solutions, which we believe not only could reduce the number of needed resources, but provide more controlled and secured usage of resources and simplify the application deployment to users.

6.3 Heterogeneity

One of the major initial goals of XIFI was to support heterogeneity. This was seen as an added value for developers. While this may be the case for some specific resources (e.g. availability of sensor networks, radio antennas), this proved not to be the case for the access to Cloud resources. By monitoring FIWARE Lab Help Desk tickets, we realized that several of the issues reported in association with the new nodes where related to different network configuration modality supported by nodes using Nova-network and nodes using Neutron (See also on-going action 2 in Section 6.2.1). Users – due to the fact that tutorials were referring to Nova-network configuration, and that Neutron is more complex to master[4] - were facing several difficulties into using Neutron-enabled nodes.

6.3.1 On-going actions

To resolve issues related to heterogeneity a main action have been enacted:

[4] Complexity of usage is directly proportional to the powerful capabilities provided by Neutron, which allows for private network owned by single tenant ensuring higher security and configurability of tenant networks.

- **Adopt Neutron as reference network solution.** In line with the on-going action 1 and 2 in Section 6.2.1, all nodes will be migrated to an OpenStack version that supports Neutron and will adopt Neutron as network management service. The migration plans for each infrastructure are included in the maintenance activities described in D5.5. Currently, Spain is already providing a new environment based on Juno that will shortly replace the one on Essex. Users will be supported as much as possible in the migration from the old environment to the new one.

6.3.2 Recommendations for the future

Following the already on-going actions, we recommend as well:

- **Improve documentation and information to users.** Being FIWARE Lab based on a federation of infrastructures, there will be always a certain degree of heterogeneity. Trying to remove all these differences will result into a less powerful tool for end-users. Instead of "flattening" the offer, the offer should be better presented and documented, possibly leveraging on the tools developed for the scope by XIFI.

7 CONCLUSIONS

The content of this deliverable is about the activities related to the extension of XIFI federation.

- In terms of arguments this document has addressed the following aspects:
- Definition of a process inside the project XIFI to manage the federations of new nodes
- Reporting of the deployment of the new nodes
- The definition of a Validation process to assure all the nodes are 'up&running', also as a check point for the new nodes becoming operational
- Knowledge Sharing (FAQs)
- Analysis on Lessons Learned during the federation activities
- Considerations on the future of the federation

As described in the document, all the XIFI nodes (five initial nodes and eleven new nodes coming from the Open Call) are operational inside the federation. Four associated nodes were added to this set: two nodes, University of Crete and Infotec (Mexico) are fully operational; the other two nodes, University of Messina and Wroclaw University of Technology, are currently in the process of federation.

These results have been obtained through the continuous improvement of tools and process for the federation.

The main improvements were:

- the sharing of know-how with all the IOs (through the wiki and FAQ), acquired by the nodes more ahead in the process of federation deployment.
- the definition of specific training sessions for the new nodes (with the collaboration of WP7) related to important deployment topics and sharing with IOs of some answers from XIFI experts related to deployment technical issues.
- the definition of a set of operational tests for the validation of the federation process. Those tests are also important for the daily operational check of every node.

Of course, the processes and the XIFI federation tools will be able to be further improved. For this reason a survey has been done, involving all the infrastructure owners and the stakeholders involved in the task 5.5, with the goal to collect a set of “lessons learned” acquired during the federation process. The results of this survey are available inside this document: they provide useful information to be used for subsequent phases of the project and for future activities related to the federation of new nodes.

The experience gained by running the federation revealed some key issues in FIWARE Lab. In this document we shortly discuss on-going actions to solve them and further recommendation (in particular to FI-Core) to further improve the usability and stability of the Lab. As discussed the main challenge that is under resolution, it is the introduction of resource access policies for users. We foresee that the application of such policies will mitigate drastically the issues related to overbooking of resources, guaranteeing a more fair and transparent usage of resources by developers.

As the last point, the additional goal of this document is to become a reference of the federation of XIFI inside FI-Core and for the future of federation.

REFERENCES

[1] D5.1 - Procedures and Protocols for XIFI federation, https://bscw.fi-xifi.eu/pub/bscw.cgi/d44719/XIFI-D5.1-Procedures_and_protocols_for_XIFI_federation.pdf

[2] D9.2 EURESCOM XIFI office – description and establishment,https://bscw.fi-xifi.eu/pub/bscw.cgi/d64939/XIFI-D9.2b-XIFI_office_description_and_establishment.pdf

[3] D1.5 Federated Platform Architecture v2,https://bscw.fi-xifi.eu/bscw/bscw.cgi/d91459/XIFI-D1%205-Federated_Platform_Architecture_v2.pdf

[4] D2.5 APIs and Tools for Infrastructure Federation v2, https://bscw.fi-xifi.eu/bscw/bscw.cgi/d91470/XIFI-D2%205-APIs_and_Tools_for_Infrastructure_Federation_v2.pdf

[5] D5.2, Report on XIFI Core Backbone Deployment, https://bscw.fi-xifi.eu/pub/bscw.cgi/d64414/XIFI-D5.2-XIFI_Core_Backbone.pdf

[6] Description Of Work, eXperimental Infrastructures for the Future Internet, FP7-2012-ICT-FI XIFI – IP Proposal

[7] D7.3, XIFI Training Strategy and Material (v2), https://bscw.fi-xifi.eu/bscw/bscw.cgi/d91153/XIFI-D7%203-XIFI_Training_Strategy_and_Material_%28v2%29.pdf

[8] D1.1, XIFI core concepts, requirements and architecture, https://bscw.fi-xifi.eu/pub/bscw.cgi/d60668/XIFI-D1.1b-XIFI_core_concepts_requirements_and_architecture_draft.pdf

[9] D2.4, XIFI Handbook v2, https://bscw.fi-xifi.eu/bscw/bscw.cgi/d91683/XIFI-D2%204-XIFI_Handbook_v2.pdf

[10] D3.4, XIFI infrastructure network adaptation mechanisms API, https://bscw.fi-xifi.eu/pub/bscw.cgi/d64447/XIFI-D3.4-XIFI_Infrastructure_Network_Adaptation_Mechanisms_API.pdf

[11] D3.5, Infrastructures monitoring and interoperability adaptation components API v2, https://bscw.fi-xifi.eu/pub/bscw.cgi/d91656/XIFI-D3%205-Infrastructures_monitoring_and_interoperability_adaptation_components_API_v2.pdf

[12] ITBox – XIFI wiki page, http://wiki.fi-xifi.eu/Public:InfrastructureToolbox

[13] Infographic web page, http://infographic.lab.fiware.org/status

[14] XIFI Open Call, https://www.fi-xifi.eu/open-call.html

[15] Private web page of XIFI Task 5.5, http://wiki.fi-xifi.eu/Xifi:Wp5:t5.5

