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

eXperimental Infrastructures for the Future Internet

D5.1: Procedures and Protocols for XIFI Federation

Revision: v1.1

Work package	WP5
Task	Task T5.1
Due date	30/09/2013
Submission date	October 2013
Deliverable lead	ORANGE
Authors	Thierry Milin (ORANGE), Silvio Cretti (CREATE-NET), Federico Facca (CREATE-NET), Ivan Biasi (TN), Bernd Bochow (FRAUNHOFER), Thomas Guenther (FRAUNHOFER), Sergio Morant (ILB), Eamonn Power (WIT), Joe Tynan (WIT), Matthias Baumgart (DT), Marco Palazzotti (ENG), Alessio Martorelli (ENG)
Reviewers	Anastasius Gavras (EURESCOM), Daniele Giai Pron (TI)
Abstract	This deliverable defines a first version of operational and technical installation procedures and protocols for a new infrastructure to implement and to follow in order to become a running XIFI federation node. Also it contains definition of operational and technical requirements to joining nodes as well as a quick online test to evaluate their compliance with XIFI. In addition, this deliverable provides support procedures for joining infrastructures and end-users accessing XIFI nodes. Finally, this work defines protocols for maintenance of the XIFI tools and FI services.
Keywords	Federated platform, requirements, constraints, survey, architecture, federation models

Document Revision History

Version	Date	Description of change	List of contributor(s)
V1.0	11.09. 2013	Version ready for internal review	Thierry Milin (Orange) et al.
V1.1	22.10.2013	Published version	Thierry Milin (Orange) et al.

Disclaimer

Copyright notice

Project co-funded by the European Commission in the 7th Framework Programme (2007-2013)		
Nature of the Deliverable:	R (Report)	
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 work defines a first version of operational and technical installation procedures and protocols for a new infrastructure to implement and to follow in order to become a running XIFI federation node. Such an infrastructure can be integrated and deployed within the federation with minimum effort and minimum potential conflicts with existing operations. However, before this step the infrastructure has to be checked on its compliance with the XIFI minimum operational and technical requirements which are defined in this deliverable. This check activity is carried out using the quick online test, developed within the scope of this work.

Since the infrastructure is becoming a XIFI node, it might require additional assistance. For this kind of situation the document includes deployment support procedures for joining infrastructures. XIFI ensures support services not only for node-newcomers, but also for Future Internet (FI) developers. Therefore the work also provides procedures for end-users accessing XIFI nodes.

The XIFI federation operates based on the specific XIFI tools and FI services, which have to be maintained on a permanent basis. For this reason, this deliverable also contains the maintenance protocols for the XIFI tools and FI services.

Above all, this work delivers:

- The operational and technical requirements to new infrastructures;
- The quick online test;
- The general procedures that must be implemented;
- The physical and logical levels installation procedures;
- The new infrastructure deployment support procedures and protocols;
- The Future Internet developers support procedures;
- The protocols for the XIFI tools and FI services maintenance activities.

The target audience of this deliverable is:

- The XIFI federation office in order to evaluate whether a candidate infrastructure meets the minimum technical and operational requirements;
- The five initial nodes of the XIFI federation and the new infrastructures that will join, following the installation procedures and protocols;
- 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;
- Developers and maintenance experts of XIFI tools and FI services who will apply the procedures and use the protocols for the maintenance of the XIFI tools and FI services, hosted by the nodes.

As a result, this deliverable is a base for the XIFI nodes operation, assistance, and maintenance as well as XIFI federation extension support deliverables within the same work package, allowing the XIFI federation to monitor and to improve all the stakeholders' experience within the operational XIFI environment.

TABLE OF CONTENTS

LIST OF FIGURES

LIST OF TABLES

ABBREVIATIONS

FI	Future Internet
FI-PPP	Future Internet Public-Private Partnership Programme
FI-LAB	Future Internet Lab
API	Applications Programming Interface
GE	Generic Enabler
SE	Specific Enabler
IaaS	Infrastructure as a Service
PaaS	Platform as a Service
SaaS	Software as a Service
UC	Use Case
VLAN	Virtual Local Area Network
SMEs	Small and Medium Enterprises
SLA	Service Level Agreement
NREN	National Research and Education Network
LTE	Long Term Evolution
EIT	European Institute of Innovation and Technology
ICT	Information and Communications Technology
FIRE	Future Internet Research and Experimentation
QoS	Quality of Service
INFINITY	INfrastructures for the Future Internet commuNITY
XIPI	eXperimental Infrastructures for Public private partnership in Innovation
SFA	Slice-based Federation Architecture
FRCP	Federated Resource Control Protocol

1 INTRODUCTION

The XIFI platform is the community cloud for European FI-PPP developers enabled by advanced FI infrastructures in Europe. The FI-PPP [1] is an ambitious programme by the European Commission part of the Framework Programme 7 that aims at exploring the potential of a common platform for Future Internet technologies to establish new business ecosystems. XIFI, through this community cloud, will provide a marketplace to access:

- Web-based services offered by FI-PPP (i.e. the Generic Enablers developed by FI-WARE [2] and the Specific Enablers provided by Use Case Trials;
- Advanced Future Internet infrastructures that provide capacities and data to empower the applications developed by early adopters of FI-PPP technologies.

XIFI, as part of the FI-PPP programme, and following the principle "eat your own dog food", is based on FI-PPP technologies delivered by FI-WARE (the so called enablers). As such, not only XIFI provides FI-PPP technologies to developers that are then able to validate them through their applications (whether they are part of Large Trials or they will be part of Phase 3 SMEs and web entrepreneurs), it is in itself an adopter of FI-PPP technologies. XIFI, through FI-WARE enablers (in particular the Cloud Chapter) deploys a community cloud that federates different infrastructures across Europe (composed of a data centre and potentially additional advanced infrastructure services such as sensor networks and wireless antennas).

The XIFI project has the objective of setting up and operating a Future Internet federation for the need of expanding existing fragmented infrastructure limitations within Europe to cope with large trial deployments. The federation is formed by integrating heterogeneous test infrastructures throughout Europe and providing a sustainable marketplace for different stakeholders of the federation. To construct the federation, infrastructures are required to follow common and consistent procedures and protocols in their operations inside the federation so that a new infrastructure can join the federation with minimum effort and minimum potentials conflicts with existing operations. It is important to highlights that XIFI is based on FI-PPP technologies, and in particular on FI-WARE Generic Enablers for the Cloud Hosting and Provisioning functionalities. This poses a number of requirements on technical and operational procedures that are kept into consideration in this document. In the next sections of the document, the XIFI node (infrastructure) protocols and installation procedures are discussed within the scope of this work. Protocols and procedures may evolve over time according to lessons learnt and new needs elicited by XIFI users and new XIFI federation members. Thus the procedures contained in this document will be kept up-to-date on the wiki as soon as they evolve and will be included as Annex in the release of periodic activity report of work package 5 (i.e. within D5.3 and D5.5).

1.1 A visual map of XIFI main building blocks (you are here!)

XIFI offers a marketplace to European large-scale trial developers to access FI-PPP technologies and Future Internet infrastructures. The marketplace provides access to Generic Enablers (GEs) developed by the FI-WARE project, and potentially to the Specific Enablers (SEs) developed by Use Case and Early Trials projects, through a highly available and reliable "federation" of infrastructures. To complement and support the above-mentioned technologies (GEs and SEs) XIFI leverages upon FI infrastructures with different characteristics (e.g. in terms of location, user community, quality of service, special hardware, etc.). XIFI aims at illustrating the potential of such marketplace through different showcases that will act as demonstrators of the XIFI service offer. For example, one of the showcases will illustrate how developers can take advantage of the XIFI multi-site infrastructure to build distributed applications with high-availability set-up and controlled QoS across the different used sites. The XIFI service offer comprises two main parts: the platform, i.e. the "virtual" marketplace that allows end-users to browse through, configure and access enablers and infrastructures in preparation for their experimentations, and the operational services, i.e. the set of activities that go

around the platform to provide a comprehensive "package" to XIFI end-users. The operational services and the platform go beyond pure technical considerations: they offer a summary vision of technical and business aspects that constitute the XIFI service offering. The XIFI platform is conceived in the context of the community cloud deployment model, and offers all three established service models of cloud platforms: Software as a Service (SaaS), Platform as a Service (PaaS) and Infrastructure as a Service (IaaS) [3]. The XIFI platform is composed of different elements, such as:

- User Interfaces comprising tools for browsing, discovering, recommending, configuring, allocating and deploying resources; such interfaces provide also means to interact with operational services provided by infrastructure owners, such as developers support and SLA management;
- Federated Cloud and Service Management for the aggregation of different resources available through the federation, the shared security and identity management across the federation, the software automation for the installation of new nodes and new services on top of nodes;
- Dynamic Network Management to support connectivity configuration across federation nodes at the level of single services, supporting changing user demand;
- Resource Monitoring that supports the active and passive collection of data from physical and virtual sources, providing the capacity to gain access to meaningful information on infrastructure and service availability.

The XIFI operational services are a set of fundamental services to support the management of the XIFI platform as well as address the needs of different actors, i.e. infrastructure operators and application developers. Most of these services are defined as protocols and procedures to ensure the operational continuity of the XIFI platform. Such operational activities cover management of the nodes – in particular Level 1 and Level 2 support to developers – support for infrastructure owners in deployment and maintenance of the platform (with special focus on new infrastructures joining XIFI during its second year) and training for developers and infrastructures owners (it will be built exploiting the showcases and referring to the documentation available from the technical activities over the project).

XIFI building blocks are showed in figure below:

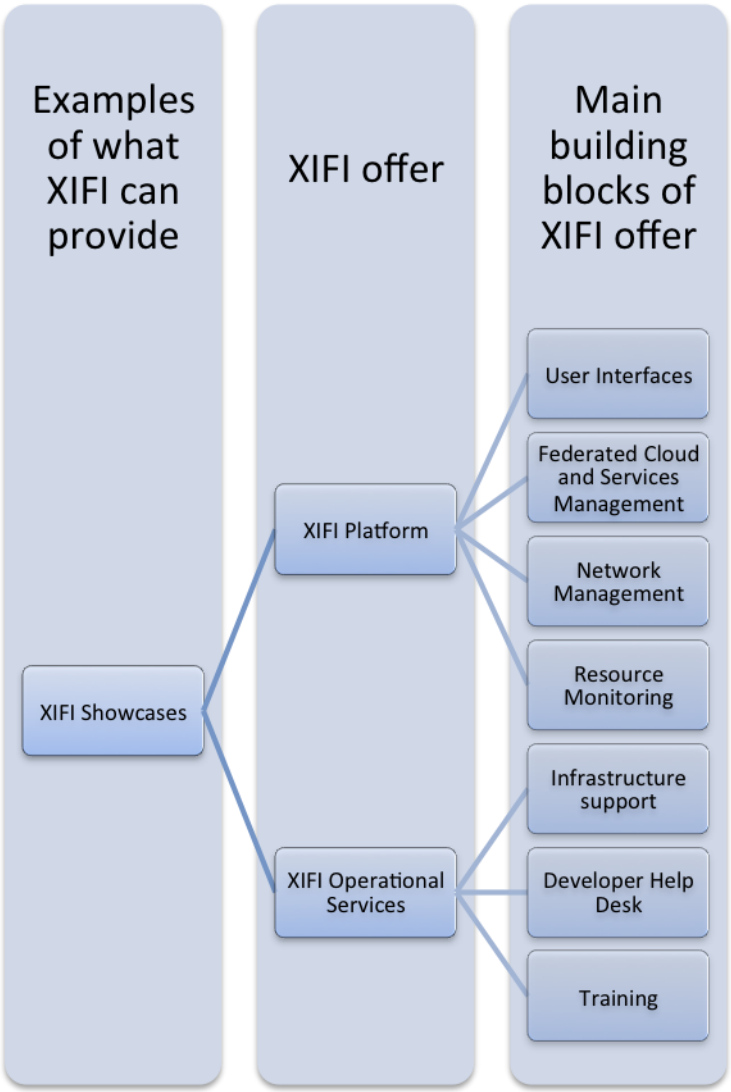

Figure 1: Visual Map of XIFI main building blocks

In this deliverable we will cover aspects related to protocols and procedures to ensure operational activities of the XIFI Platform.

1.2 Scope

The XIFI project represents a creative and innovative step towards the future, building a solid foundation for making progress in the Future Internet developments. XIFI creates a new notion of the federation by providing a single starting point as marketplace and assuring interchangeable and interoperable environment for all types of stakeholders of the federation throughout Europe.

The XIFI federation is made of heterogeneous test infrastructures called nodes. The initial node set of the federation consists of the five XIFI Core nodes, located in Berlin (Germany), Waterford (Ireland), Brittany (France), Seville (Spain) and Trento (Italy). These nodes create a fully functional federation so that other potential nodes can join and be part of it.

The work ensures definition of operational and technical requirements to be met by a joining node as well as support procedures to aid this node in integration and deployment within the federation. This makes XIFI a flexible environment for potential nodes that have reached minimum level of compliance and are ready for a certain commitment to the main principles and objectives of the federation.

Since the nature of the federation is heterogeneous, the work foresees different compliance levels of candidate nodes with XIFI. Therefore the infrastructures are provided with a quick online test as an entry point for an initial evaluation of their compliance with the XIFI federation and its requirements. To assist infrastructures on this level, the work defines the methodological support service which also includes a help desk. Besides the guidelines, documentations and a help desk, the work covers procedures of two support services more for a joining node. The assignment of one of the support services to a candidate node is carried out by XIFI, depending on the level of compliance.

The level of compliance is selected on a case by case basis; however the lowest level is presented by XIFI according to the minimum set of XIFI's requirements to be met by a joining node.

In addition to the definition of requirements, procedures and protocols for a joining infrastructure, the work aims at defining procedures that support end-users accessing XIFI services.

Last but not least, the work sets the protocols for maintenance of the XIFI tools and FI services, hosted by XIFI nodes.

Summarising, this deliverable addressed the following aspects:

- Definition of operational and technical requirements to joining nodes;
- Operational and technical procedures to implement and to follow by joining nodes;
- Quick online compliance evaluation test for joining infrastructures;
- Procedures of support services for joining infrastructures;
- Procedures to support end-users accessing XIFI nodes;
- XIFI tools and FI services maintenance protocols.

1.3 XIFI Federation Policies

In this section we list the set of policies that drives the definition of procedures for the set-up and operation of XIFI federation.

- XIFI provides Level 1 and Level 2 technical support [4] to Future Internet Developers; for network and hardware related issues related to single nodes, XIFI will cover also Level 3. Level 3 support related to Generic Enablers is delegated to FI-WARE. Level 3 related to XIFI code is delegated to XIFI developers.
- XIFI provides Level 1 and Level 2 technical support to XIFI nodes with respect to XIFI software stack installation and operation. Level 3 is provided only for software components developed by XIFI, Level 3 support related to Generic Enablers is delegated to FI-WARE.
- XIFI maintains the software stack installed on top of XIFI nodes in line with FI-LAB [5] releases (that include tested version of FI-WARE Generic Enablers) and XIFI tools releases.
- XIFI supports the inclusion of new infrastructures. A specific service for the inclusion of new infrastructures provides the needed knowledge transfer and technical support in the activity.
- XIFI may support the connection of new infrastructures related to FI-PPP Trials. Such infrastructure may be, for example, data centres connected as private capacity to the FI-PPP trials, data providers (i.e. sensor networks or other source of data), or connectivity providers required to connect FI-PPP trials services to XIFI nodes.
- XIFI and XIFI node operators are not liable for the loss of data by developers. XIFI nodes, where available, may offer backup services to prevent such events.
- XIFI deals with the definition of Usage Terms for the access to XIFI node resources. XIFI users will have to agree to such Usage Terms in order to access resources provisioned by XIFI nodes. The definition keeps into consideration the potential different needs of federation

members.

- XIFI will adopt Usage Terms compliant with FI-LAB usage terms with respect to access to Generic Enablers.

1.4 XIFI Federation, Federation Extension, and Evolution of Procedures

XIFI includes five initial nodes distributed over Europe. These nodes form the core of the XIFI federation. Through an Open Call, XIFI will enlarge the federation by encouraging additional nodes to join the federation. Additional nodes may include use case related infrastructures, EIT ICT Labs infrastructures and test beds, FIRE facilities and infrastructures, and other Future Internet infrastructures available in Europe. The Open Call and related administrative matters are managed by the Federation Office.

The set-up and federation of the five initial nodes will allow assessing procedures, protocols and tools required to establish and maintain a federation. This is considered an important step to develop, specify and document the minimum technical and operational requirements a node has to comply with, in order to federate and to host GEs. It is also an important step to gain experience and to better understand technical, operational and legal requirements of a wide-area federated heterogeneous infrastructure.

In order to federate, a new infrastructure has to satisfy a certain level of compliance with requirements documented in the subsequent sections. Since the federation process is a technical and business process at the same time, the level of compliance also reflects requirements on technical compliance and compliance with the business objectives implemented by the marketplace. Although business compliance grounds upon technical compliance, expressed through GE compliance to the extent required for implementing the envisaged business model, this work is focusing on technical compliance, in particular:

1. Compliance with procedures and workflows for node operation including node interoperability and interface compliance (addressed by this work).
2. Compliance of infrastructure monitoring capacities provided by an infrastructure respecting federation requirements (addressed through other activities in the project).

Further infrastructure compliance and interoperability issues may come up in the context of GE hosting but are considered as off-topic for this deliverable. These issues may need to be addressed in the context of other work packages, in particular in the scope of WP3.

Compliance with procedures and workflows relies on acceptance of operational requirements, on endorsement of tools required to operate and to federate, on providing physical level connectivity, and of supporting the federation's dissemination and exploitation targets. This includes the ability to:

- Host the minimal set of tools that guarantee the connection of the new infrastructure to the federation, and a minimal set of services such as cloud hosting and infrastructure monitoring as set forth in subsequent sections;
- Support the instantiation of showcases related to XIFI and FI technologies that target specific regional or local developers; and
- Support the organization of events to advertise XIFI and FI technologies to local or regional entrepreneurs, notably SMEs, and researchers.

Hence, the level of compliance here also quantifies not only the compliance of a new node with requirements related to joining the federation, but also includes compliance with the overall targets of the federation to create tangible results and business impact. This implies, for example, that a new infrastructure also needs to comply with the overall maintenance procedures of the federation to behave in a coordinated way regarding its contribution to the federation's offer to the user.

Compliance of infrastructure monitoring capacities in particular requires obtaining basic

information from each new infrastructure about its monitoring capacities and support, about resources provided, and about the QoS of these resources provided to the federation. This ensures that all federated infrastructures can be monitored regarding their availability and performance, which is the foundation of any service level agreement.

Compliance and interoperability of GEs ensures that all new infrastructures are able to host FI-WARE GEs. GE compliance in this sense consists of:

- Compliance of the infrastructure to host GEs.

Test suites will be provided to verify the level of compliance, including but not limited to interface compliance tests on the level of interface primitives and their parameters and interoperability tests on the functional and procedural level.

A set of requirements (technical and operational) and procedures (physical and logical) formalise the compliance specifications. Multiple levels of compliance will be possible, potentially based on the degree of coverage of requirements, whether judging on business compliance or on the value of the contribution to the federation.

A specification of the minimum level of compliance required to join the federation, is addressed by subsequent sections and is made part of the Open Call text, hence making them mandatory for all new infrastructures.

New infrastructures may go beyond minimum requirements by offering additional services, such as access to sensor networks, back-up facilities and support, mobile infrastructures, or other dedicated domain services. Since these additional services will enlarge the offer of the federation to future internet experimenters and developers, their offer may be considered a significant benefit enabling further business opportunities, and may require evolving procedures defined here later in the process.

Compliance with minimum requirements of an infrastructure aiming to join the federation must be verifiable during the node set-up prior to joining the federation and during federation lifetime as an integral part of the maintenance procedures. Pre-federation compliance testing is a mandatory step in order to protect the federation from potential harm through new infrastructures while online conformance testing will need to continuously verify that a certain infrastructure is not impairing the federation in whole. Hence, the latter is also included in the maintenance procedures since it may be required, in the worst case, to disconnect an infrastructure from the federation.

1.5 Infrastructures versus Nodes

It is important to clarify since now, that while a node is an infrastructure, not all infrastructures are nodes in the context of XIFI Federation. Generally speaking XIFI consider an infrastructure any ICT facility that:

1. offers capacity to host GEs to build backends for FI applications
2. offers connectivity to internet and GEANT network
3. offers services to developers: support, backup, ...
4. offers additional capacities such as: sensing environment, advanced wireless connectivity
5. offers access to user communities

To be evaluated as a node to be part of the XIFI Federation, the first 3 items are compulsory, the others are would-like. Within XIFI we will also consider other infrastructures, but not as members of the Federation. Thus not having the rights and duties that the membership entails.

In general XIFI may deal with three categories of infrastructures:

- Full member nodes: i.e. infrastructures that are part of the XIFI Consortium and as such of XIFI Federation;

- FI-PPP member nodes/infrastructures: i.e. infrastructures that are part of other FI-PPP projects, and that are connected to XIFI as private/shared resource for the FI-PPP community cloud (node) or as "data providers" (infrastructure).
- Third party nodes/infrastructures: i.e. infrastructures that are neither part of the XIFI consortium nor of any other FI-PPP project and that are connected to XIFI as private/shared resource for the FI-PPP community cloud (node) or as "data providers" (infrastructure).

1.6 Lessons Learned From Previous/Other Research Projects

1.6.1 INFINITY

Through collaboration with organisations across Europe, the INFINITY project captures and communicates information about available test infrastructures and any interoperability requirements and issues. INFINITY documents any usage-related operational constraints and identifies and fosters federation opportunities that could facilitate large scale experimentation and testing. A dynamic innovative repository XIPI [6] based on a set of community-driven Web tools was created to promote the evolving vision of available infrastructures "as a living organism". This is supported by a methodology that will promote a consistent categorisation of the infrastructure resources, thereby facilitating a mapping between Use Case requirements and infrastructure offerings. The efficient gathering of data about the available infrastructures is ensured by including key representatives of the important public and private infrastructure stakeholders directly in the consortium and/or as members of a "Concertation Board". The Web repository, the close co-operation with the FI-PPP Facilitation CSA and the specific expertise and relationships of the partners will:

- Support the Core Platform by consolidating detail about existing and emerging advanced infrastructures, and help define the required Generic Enablers for seamless integration and enable new and innovative experimentation;
- Stimulate infrastructure owners to effectively "bridge the gap" between their current capabilities and Use Case requirements, thereby encouraging investment in upgrades and standards to realise these opportunities, and leading to greater sustainability.

The end goal of INFINITY is to establish a pan-European Future Internet testing infrastructure.

Technical Approach of INFINITY:

- Gather, analyse, classify and evaluate and organise information about Future Internet infrastructure and usage profiles across Europe: INFINITY Common Description framework to be used for ICT infrastructures description, classification and profiling;
- Make that information available to other FI initiatives through a "live" web enabled repository;
- Encourage and support interaction, collaboration and experimentation between application and infrastructure owners, operators, industry, local and regional authorities and end users;
- Detect the technical and operational constraints for the experimentation with the ICT infrastructures in the Future Internet;
- Communicate and promote the activities of European Future Internet research capacity and infrastructures world-wide;
- Help prepare a route map to the future through integration of identified FI infrastructures in Phases 2 and 3.

XIPI – the widest collection of data on Future Internet related infrastructures:

- 150 infrastructures registered;

- A tool to support Future Internet trials in Europe.

INFINITY project highlights a complex and dynamic Future Internet infrastructures landscape in Europe where there is no a one model fits all. Different infrastructures (even in the same experimental arena) adopt different technologies and different operational approaches. This diversity, one side, is a stimulus to Future Internet experimentation by providing different playgrounds that can be used to develop different applications in different contexts. On the other side, it is a challenge to experimenter (or Future Internet Developers) that needs to deals in different ways to different infrastructures they would like to combine in their experimental set-up. INFINITY provided a first step to facilitate this activity by describing infrastructures using a common format, but other steps are needed on the operational side and on the technical one to simplify life to developers. XIFI will leverage on this experience to tackle some of the problems evidenced by INFINITY, for example by elaborating uniform usage terms for the infrastructure part of the federation.

1.6.2 BONFIRE

The BonFIRE Project (Building service testbeds for Future Internet Research and Experimentation) designs, builds and operates a multi-site cloud facility to support applications, services and systems research targeting the Internet of Services community within the Future Internet. BonFIRE gives researchers access to an experimental facility which enables large scale experimentation of their systems and applications, the evaluation of cross-cutting effects of converged service and network infrastructures and the assessment of socio-economic and other non-technological impact.

How does the infrastructure work? BonFIRE operates a Cloud facility based on an Infrastructure as a Service delivery model with guidelines, policies and best practices for experimentation. BonFIRE adopts a federated multi-platform approach providing interconnection and interoperation between novel service and networking testbeds. The platform will offer advanced services and tools for services research including cloud federation, virtual machine management, service modelling, service lifecycle management, service level agreements, quality of service monitoring and analytics. Where appropriate BonFIRE reuses and adapts existing tools from other FIRE projects such as Panlab [7], Federica [8] and DEISA [9]

BonFIRE supports experimentation and testing of innovative scenarios from the Internet of Services research community specifically focused on the convergence of services and networks. Three scenarios are envisaged:

- Extended cloud scenario: tests are run on cloud computing sites interconnected through public Internet. Properties within the site are controlled but the properties of the network are not.
- Cloud with a controlled experimental network scenario: nodes are now connected through an emulated virtual internet, allowing both server and network properties to be controlled.
- Extended Cloud: with complex network implications (using networks slices) and involving federation with other FIRE infrastructures.

BonFIRE provides innovative methods for describing, deploying, managing, executing, measuring and removing experiments including:

- Uniform test description and deployment descriptors for all the scenarios (including crosscutting tests);
- Cloud resource federation through the federation of clouds in different administrative domains that provide physical resources to BonFIRE;
- User-friendly user interfaces at the facility's entry point with an easy to use portal.

1.6.3 Fed4FIRE

In recent years numerous projects for building FIRE facilities have been launched, each targeting a

specific community within the Future Internet ecosystem. The goal of the Fed4FIRE project [10] is to federate these different facilities using a common federation framework. Such a federation can prove to be beneficial in several ways. First of all, it enables innovative experiments that break the boundaries of these domains. It also allows experimenters to more easily find the right resources to translate their ideas into actual experiments, to easily gain access to different nodes on different testbeds, to use the same experimenter tools across the different testbeds etc. This means that the experimenters can focus more on their research tasks than on the practical aspects of experimentation. The federation is also useful from the infrastructure providers' point of view, since they can reuse common tools developed by the federation, they can reach a larger community of possible experimenters through the federation, etc. Currently 13 testbeds are members of the Fed4FIRE federation, introducing a diverge set of Future Internet technologies. During the course of this project more testbeds will join the federation using an open call mechanism to selected testbeds for inclusion in the Fed4FIRE project.

The corresponding federation architecture that is defined by the project is characterized by a preference for distributed components. This way, the federation would not be compromised if, in the short or long term, individual testbeds would discontinue their support of the federation. The general policy is that experimenter tools should always be able to directly interact with the different testbeds, and should not be forced to pass through some central Fed4FIRE component. However, some non-critical central federation-level components are also included in the architecture for convenience purposes.

A critical aspect in such a highly distributed approach is the adoption of common interfaces in the federation, and making sure that every member of the federation is fully compliant with them. Therefore, Fed4FIRE is developing a new software tool that focuses on acceptance testing of the required interfaces for testbed federation. This test suite enables the rigorous testing and integration activities that are needed when federating a highly heterogeneous set of testbeds with the intention to realize a fully operational federation. The tool focuses on logical tests (including all steps of the experiment workflow) and adds specific interface tests and negative testing (are things breakable?) where needed.

In the context of resource discovery, reservation and provisioning, the adoption of the Slice-Based Federation Architecture (SFA) is a key element in Fed4FIRE. Therefore, the first focal point of this compliance-testing tool is the support of manual and automatic nightly testing of the entire SFA API (including the Aggregate Manager, Slice Manager and Registry APIs). This testing functionality is entirely developed in Java, allowing greater flexibility in development of both the test suite and future Java SFA client tools. For the manual testing of the SFA interface of any given testbed, both a command line and a graphical user interface are provided. The automatic (nightly) testing of testbeds is run from within a Jenkins platform, posting the test reports on a website and sending emails in case of problems. The test suite will be released as open source software, and will easily allow for extensions through a plugin system. This way, other important Fed4FIRE federation interfaces will also be added to the testing suite as the project continues. The Federated Resource Control Protocol (FRCP) is such an interface; another example is the adopted interface for identity authentication.

1.6.4 FI-WARE

1.6.4.1 The FI-WARE testbed

From the point of view of the infrastructure owner's, various actions must be taken into account when the XIFI infrastructure came to a production level. During the exploitation of the FI-ware testbed, one of the problems has been all security-related aspects. The provision of infrastructure (virtual servers/machines for the development of components and applications) to people, organizations or research groups without any experience in security is a big problem. In many cases, these servers are compromised, and many malicious software are installed. In fact, the RedIRIS security incident response team has registered many incidents where FI-ware machines are involved.

The security aspects should be fixed in the infrastructure design phase. Because cloud services are providing new infrastructure on demand, it's a bit complicated to apply filters in advance. Therefore it is necessary to implement the network tools to detect attacks, and also to investigate what has happened when a machine has been compromised.

The infrastructure's operation involves planned maintenance tasks and incidents (unexpected situations). The coordination between infrastructure owners, and other agents whom are using these capabilities is very important. The creations of the channel to broadcast information about maintenance tasks or incidents could be very useful.

1.6.4.2 The Future Internet Lab (FI-LAB)

During the development and exploitation of the production infrastructure for OIL, our experience has shown that OpenStack is very sensitive to unexpected infrastructure failures (disks, network or machines faults). It is important to work in a fault-tolerant system that allows the moving of virtual resources between hardware (in case of failure of any component). XIFI should improve the mechanism to get fault tolerant infrastructure. . Additionally, OIL was available in two geographically separated datacenters. If the data centers have no connectivity at level 2 (ethernet) is necessary to install two different cloud managers, one per data centers, which means that the datacenters resources are managed individually

1.7 Document Convention

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

1.8 Intended Audience and Reading Suggestions

The target audience of this deliverable is:

- The XIFI federation office in order to evaluate whether a candidate infrastructure meets the minimum technical and operational requirements;
- The five initial nodes of the XIFI federation and the new infrastructures that will join, following the installation procedures and protocols;
- 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;
- Developers and maintenance experts of XIFI tools and FI services who will apply the procedures and use the protocols for the maintenance of the XIFI tools and FI services, hosted by the nodes.

2 GENERAL PROCEDURES

In this section we define a set of general procedures that are required to implement the other procedures listed in the document. These procedures are related to the identification and assignment of operational roles for:

- Management of nodes;
- Developer support;
- Infrastructure support.

2.1 For the Management of Nodes

Each node should provide a reference person for the following roles:

- Node Manager: the main contact for the node and the person in charge of decision on how to apply XIFI policies and procedures in the node.
- System Administrator: the person in charge for the physical set-up of servers, the installation of server management software and its configuration.
- Network Administrator: the person in charge for the physical set-up of network (internal and external access), the installation of network management software and its configuration.
- Node Help Desk: the person in charge for the support of user requests specific to a node.

For each reference person the node should provide:

- Full name
- Email contact
- Phone contact (only for net, sys and help desk)
- Availability (only for net, sys and help desk)

This information should be kept up-to-date in the following area of the wiki:

Management of XIFI Nodes: Reference People
(http://wiki.fi-xifi.eu/Fi-ppp:Management_of_XIFI_Nodes)

2.2 For the Developer Support

The developer support is based on a shared facility that act Level 1 support - the facility is shared with infrastructure support. Different roles to run the developer support are required.

- Level 1 Help Desk: the person in charge to filter tickets incoming to the shared facility.
- Node Help Desk: the person in charge for the support of user requests specific to a node (XIFI user is a developer), he provides Level 2 support for the Node and GEs. Level 3 is demanded to System Admin or Network Administrators in case of infrastructure issue, to Software Component Support in case of GEs or XIFI federation tools.
- Software Component Support: the person in charge of providing the support for a specific GE or XIFI federation tool.

Each person listed above should provide:

- Full Name
- Email contact

- Register in the shared facility that act as Level 1 support tool (they will be assigned to an area according to their role)

This information should be kept updated directly in the shared facility.

2.3 For the Infrastructure Support

The infrastructure support is based on a shared facility that act Level 1 support - the facility is shared with developer support. Different roles to run the infrastructure support are required.

- Level 1 Help Desk: the person in charge to filter tickets incoming to the shared facility.
- Federation Manager: the person in charge of the federation office and of the process of including new nodes.
- Federation Deployment Help Desk: the person charge of providing the support for node deployment.
- Software Component Support: the person in charge of providing the support for a specific GE or XIFI federation tool.

Each person listed above should provide:

- Full Name
- Email contact
- Register in the shared facility that act as Level 1 support tool (they will be assigned to an area according to their role)

This information should be kept updated directly in the shared facility.

3 PROCEDURES FOR JOINING THE FEDERATION

XIFI will mostly adopt, as early investigated in D1.1 [11], variants of the One Shop Stop or Matchmaker federation models. In particular, the reservation of standard resources made available by the federation participants will occur through a single entry point provided by XIFI (the federator). Federation participants will be able to decide the amount of resources available to end-users and configure quotas for them, but the overall provisioning and deployment will occur through the federator. This requires that all federation participants comply with a reference software architecture and hardware architecture that will enable them to publish their resources in the single entry point provided by the federator. Nonstandard resources and ad-hoc complex services, such as set-up of LTE antennas or deployment of isolated cloud environments, will be managed by the single infrastructure owners; still the request will be originated at the federator.

From a technical point of view XIFI considers three types of infrastructures capacities:

- Computing capacities: i.e. infrastructures that provide hosting capacities for provisioning of GEs (e.g. Data Centres);
- Data capacities: i.e. infrastructures that provide data sources that can be connected to GEs (e.g. Smart Cities or Sensor Networks);
- Transport capacities: i.e. infrastructures that provide connectivity to support GEs provisioning and GEs access to/from data and users (e.g. NRENs).

A single infrastructure can combine all the three capacities listed above. An infrastructure to be a node in XIFI should combine at least computing and transport capacity. Other infrastructure maybe connected to XIFI but not considered nodes.

From an administrative point of view XIFI considers three categories of infrastructures:

- Full member nodes: i.e. infrastructures that are part of XIFI consortium;
- FI-PPP member nodes/infrastructures: i.e. infrastructures that are part of other FI-PPP projects;
- Third party nodes/infrastructures: i.e. infrastructures that are neither part of the XIFI consortium nor of any other FI-PPP project.

Full member nodes are required to provide computing capacities and transport capacities. The following paragraph describes the operational procedures to join the federation by nodes offering computing and transport capacities and enable support for data capacities.

The summary workflow diagram provides an overview of the distinct steps required for a new node to join the federation and outlines the interaction between the roles that implement the procedures. This process is structured into four distinct phases that will be mapped on the physical and logical procedures discussed in more detail by subsequent sections. In particular, the description of the four phases will be used in the following to determine mandatory and optional installation procedures as well as prerequisites and requirements including responsibilities of the roles. The process and tools to join the federation will be discussed in further detail in the scope of D2.1 [12].

3.1 Node inclusion process

The four phases of the node inclusion process are:

1. Application to join the federation;
2. Facility bootstrap;
3. Infrastructure bootstrap;
4. Federation extension.

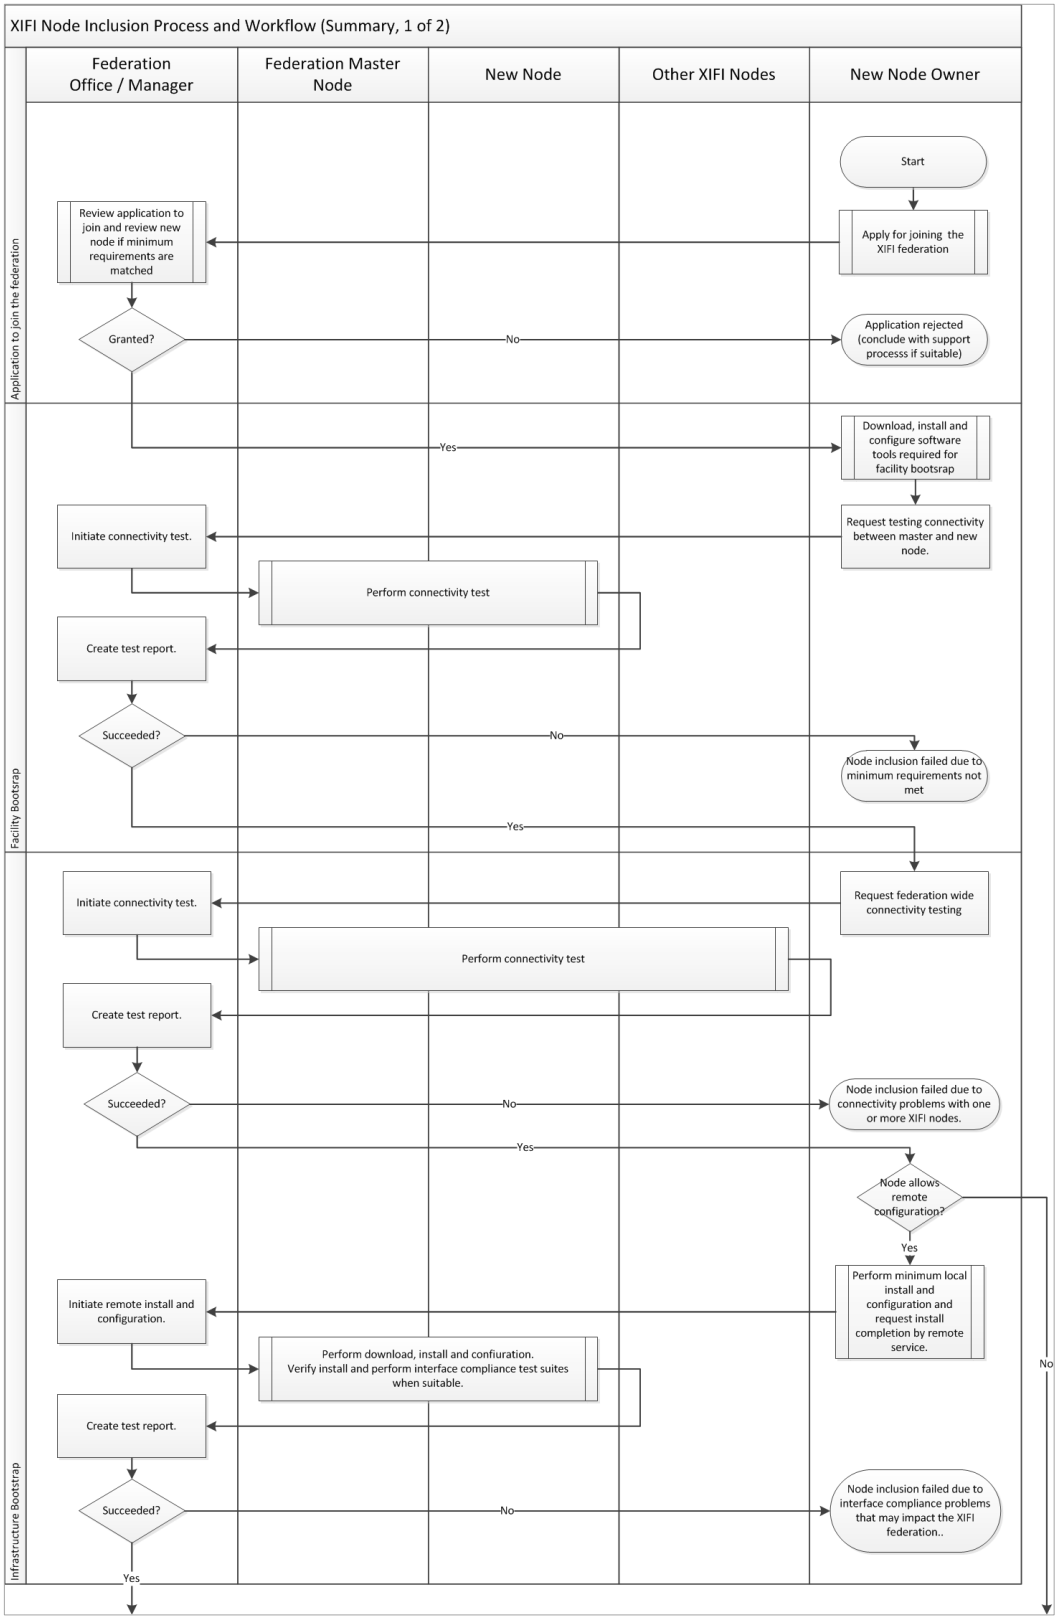

Figure 2: XIFI node inclusion process and workflow (part 1)

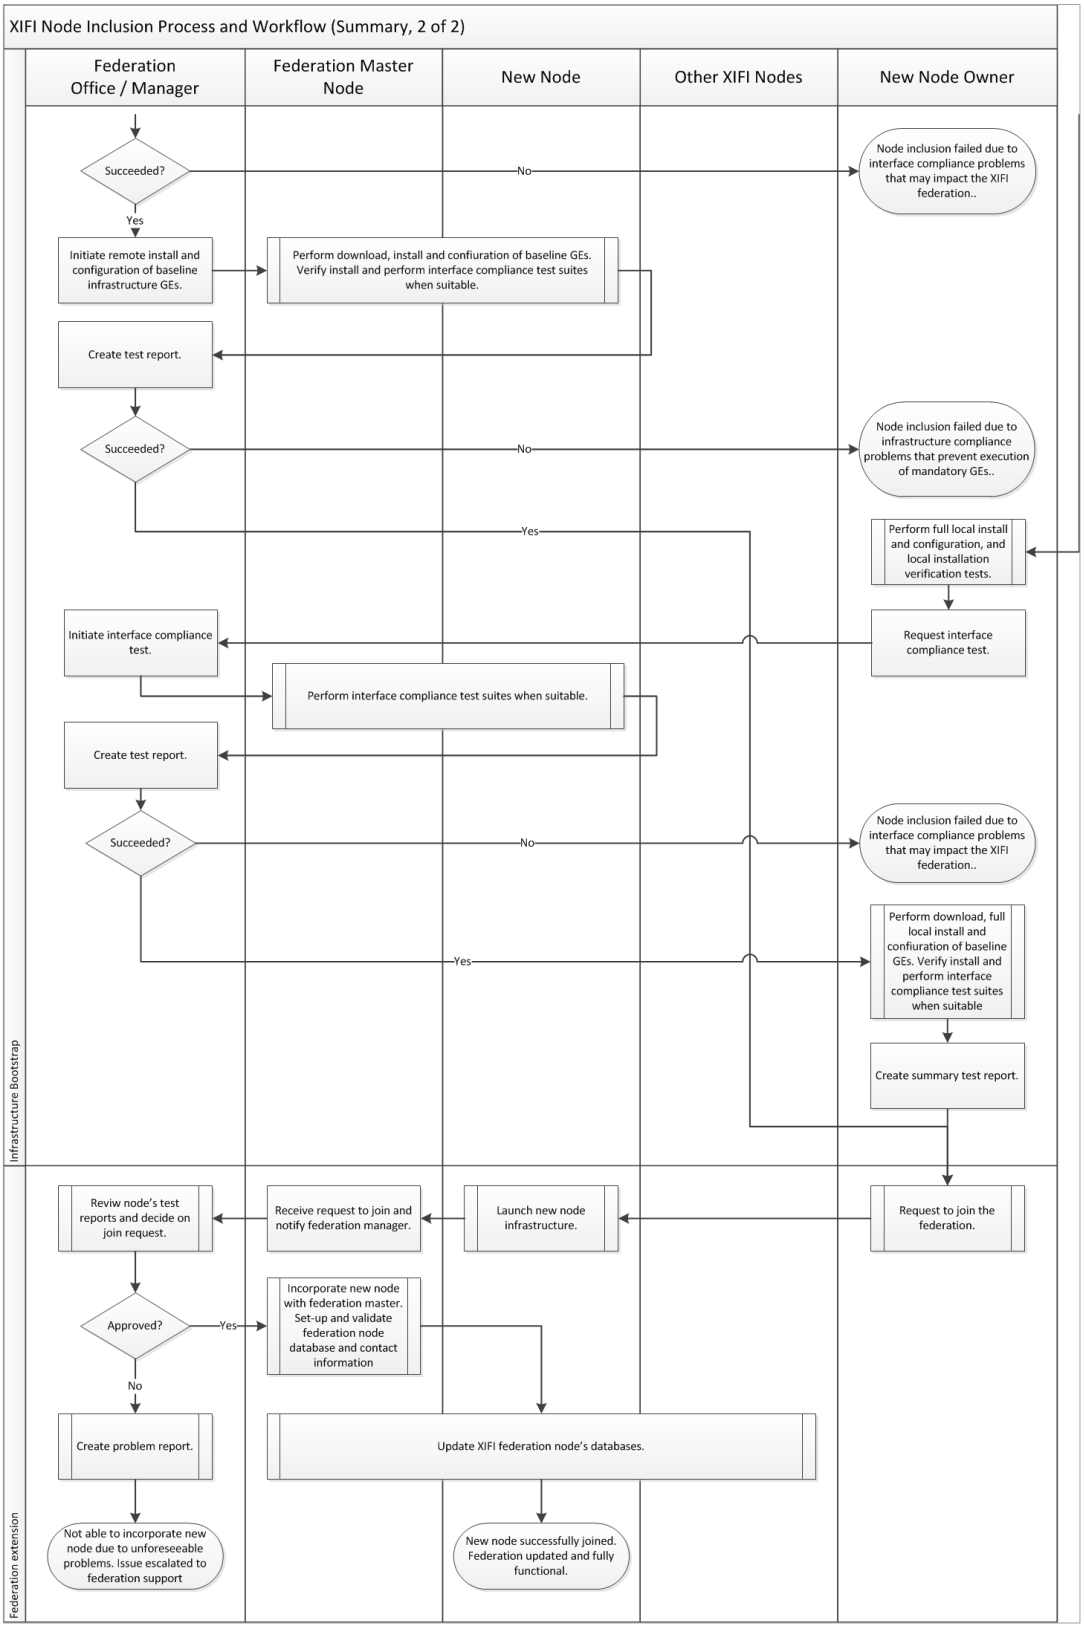

Figure 3: XIFI node inclusion process and workflow (part 2)

An **application to join the federation** is usually issued by the owner of a new node with the objective to join the federation. All new nodes need to run through this procedure. The federation office has the responsibility to review the new nodes application along with the information provided by the node

owner and shall decide if the new nodes can satisfy the minimum node requirements. The federation office takes care of this process directly in case of FI-PPP member nodes, or indirectly through the Open Call [13] process by which infrastructures become Full member nodes. In any case, the federation office supervises the process. Minimum node requirements are set forth in this document and shall be fulfilled by all XIFI nodes in order to avoid a new node to cause unwanted impact on the federation or degrading the federation's performance. It is up to the federation office either to reject an application to join, to accept a join request or to conditionally accept the application. Conditional acceptance will require support actions to guide the new node through a successful application process. The whole process is considered to take place between the new node owner and a delegate of the federation office (i.e. the federation manager) usually being human persons. Thus, this phase mostly relies on logical level procedures detailed below. An accepted application may go under a negotiation process involving the Federation Manager and other actors representing XIFI (e.g. Technical Manager and Infrastructures Manager). This process aims at defining the responsibilities among XIFI and the new node to support the process of joining the federation. For example if the new infrastructure to join the XIFI federation requires a specific driver for its monitoring solution, the two parties have to agree the responsibilities for the development of such driver.

The **facility bootstrap** aims to enable a new node to communicate with a federation master node. This process involves a number of software tools to implement the various intermediate steps. It also involves strong interaction between the infrastructures manager and the new node owner. In general, the node owner has to provide minimum functionality for her node to establish the data communication path and has to install and configure basic software components. The infrastructure manager and the technical manager verify that the new node fulfils mandatory requirements before the next phase can be approached and, if needed, may support the node owner in the configuration of his deployment architecture.

The facility bootstrap includes the following (non-exhaustive) list of sub-tasks.

- Physical architecture deployment. In this sub-task, the node owner, in compliancy with supported deployment architectures by XIFI and taking into account the consultancy of XIFI Technical Manager, defines the physical deployment architecture. Accordingly she sets-up the physical infrastructure.
- Download, installation, configuration and local functional verification of the software components installed.
- Requesting remote tests provided by the federation master to verify that minimum communication requirements between the new node and the federation master node are matched and that remotely accessible interfaces of the new node comply with the requirements and compatibility level of the XIFI federation set forth in D2.1 [12].
- If supported by the new node, installation of local monitoring facilities that can support the communication tests, and installation of software tools that can support remote assistance, installation, configuration and test procedures.

The approval process shall be supported by software tools that can create test logs and summary reports and assist the XIFI Infrastructure Manager to set-up and perform the test suites required. This facility herein is referred to as the federation management dashboard. The facility bootstrap mostly consists of physical level procedures. Few logical level procedures in this phase address the interaction between new node owner and federation manager in cases of problems arising that require supportive actions and the use of the federation help desk.

The **infrastructure bootstrap** aims to ensure that a new node is able to communicate with other nodes of the XIFI federation and that it provides all mandatory services and functionality required for participating in the XIFI federation. This phase will establish local support functionality, namely the local infrastructure monitoring capacity, cloud hosting capacities and service life-cycle management. Additionally, authorization and authentication services for accessing the infrastructure shall be provided in this phase. The infrastructure bootstrap physical procedures are dedicated to install and

configure the services and functions properly either by local node management (i.e. by node owner actions) or by remote maintenance capacities (i.e. by infrastructure manager actions) being either a temporary solution (i.e. accessible only during node bootstrap relaxing security and privacy requirements temporarily) or as a permanent service integrated with the node-local service provisioning framework (i.e. by providing suitable FI-WARE GEs).

The infrastructure bootstrap includes the following (non-exhaustive) list of sub-tasks.

- Software deployment architecture. In this sub-task, the node owner, in compliancy with supported deployment architectures by XIFI and taking into account the consultancy of XIFI Technical Manager, defines the software deployment architecture.
- Download, installation, configuration and local functional verification of supporting software tools and of a baseline set of FI-WARE GEs. This sub-task is performed accordingly to the software deployment architecture defined in the previous sub-task.
- Performing local installation verification tests and requesting federation-wide connectivity tests involving all or a sub-set of XIFI nodes in the federation upon discretion of the infrastructure manager.
- Interface compliance tests at multiple stages of the infrastructure bootstrap, in particular regarding services and functions accessible from the federation but being installed and configured by the node owner.
- Minimum functional tests verifying that all functional capacities required, to join the federation are accessible and behave in a sane way considering for example service response times and formatting of the response (performance and correctness).

Test results shall be recorded by the infrastructure manager to ensure that a later decision by the infrastructure manager upon approval of a node being allowed to join the federation is well documented. Upon acceptance of the test logs, future changes (e.g. regarding compliance or performance) of the node can be judged if they are relevant for the service level agreements between the new node and the XIFI infrastructure.

The final phase from the perspective of joining the federation is the **federation extension** phase. Upon completion of this phase the new node shall be able to operate in the federation and to offer services. The federation extension has two main facets:

1. On the physical level the new node shall be able to contribute to the federation-wide service management and resource provisioning and shall provide infrastructure services for security and infrastructure monitoring. Additionally it shall provide all necessary cloud management services to deploy and execute a user experiment or use case.
2. The final step of the federation extension phase is in publishing the new infrastructure and its offered services to the federation. Although including minor technical enablement procedures, this step consists of logical (here service level) procedures.

The naming of the phase is deduced from the fact that after completion, the XIFI federation is extended by one new infrastructure and that new capacities may be published.

3.2 Requirements

3.2.1 Technical Requirements

XIFI federation is a community cloud that offers services to FI-PPP early trials and FI-PPP developers. As such, an infrastructure joining XIFI needs to meet the technical requirements that enable the infrastructure to be connected to XIFI federation. The requirements discussed in this section derive from the analysis conducted in D1.1 [11] on the Use Case Projects' needs, on the architectures of the five core nodes, and on the FI-WARE Cloud Chapter (that constitutes the base of the software

architecture used by XIFI for providing computation capacities).

The minimal requirements listed below regard the hardware capacity needed by an infrastructure to connect to XIFI federation. Other specific requirements are being raised by Use Case projects and Early Trials within the Architectural Board. A consolidated version will be included in the XIFI Open Call [13] to select new XIFI nodes according to the Early Trial projects' needs.

The requirements can be distinguished among connectivity/network requirements, hardware requirements and software requirements. Hereunder all these minimal requirements are specified

3.2.1.1 Connectivity/Network Requirements

The connectivity capacity will be used for two aims: connect to the backbone of XIFI federated cloud to support node management operations, and provide connectivity to deployed service for end-users.

- 1 Gbps GEANT connectivity for the backbone
    - Backbone will be implemented in IPv6.
- 100 Mbps Internet connectivity for end-users
    - Service to end-users will be provided in dual stack IPv4/IPv6.
- SEC Firewall.

3.2.1.2 Hardware Requirements

The data centre capacity will be used for deploying FI-WARE platform and GEs and the XIFI services that will allow connectivity to the federation. To this aim, data centres are required to be able to host OpenStack or to have it already available. The basic required software stack is the one described in FI-WARE documentation DCRM installation [14] additional services will be provided by XIFI for network and data centre monitoring. Such services in general terms will leverage on top of existing monitoring tools provided by the data centre.

Minimal hardware capacities to be provided by a node are:

- At least 100 CPU cores (Core types: Intel VT-x or AMD AMD-v);
- RAM 2 GB x Core;
- HD 20 GB x Core.

The exact configuration of the servers depends largely on the deployment architecture as discussed in the following sections of this deliverable. The average configuration presented here refers mainly to the computational nodes.

3.2.1.3 Software Requirements

The reference deployment of XIFI nodes assumes that the following software requirements are met:

- Operating system: Ubuntu 12.04;
- Hypervisor: KVM;
- IaaS Manager: availability to install OpenStack Grizzly [15] version (all basic components plus Quantum and Swift).

It is planned to make available in the Infrastructure Toolbox a set of tools that facilitate the installation of the software components mentioned before. Nevertheless joining node may have already existing OpenStack installation on their node on top of which they can install additional software needed to join the XIFI federation.

The aforementioned requirements should be considered as the minimal set each infrastructure,

interested to joining the XIFI federation, should satisfy. It is also important to highlight that, due to the possible workload generated by the experimentation activities, an infrastructure should be able to provide further capacity with respect to the minimal described above: a suggested target is the possibility to double the minimal capacity.

Being able to provide other services (different from the computational and storage capacity detailed above), like for example sensor networks or mobile networks is an added value that can foster the infrastructure usage by the future internet developers and experimenters.

3.2.2 Operational Requirements

This section provides the minimal set of high level operational requirements an infrastructure that is part of the XIFI federation or that wants to join the XIFI federation should satisfy. These requirements can be used as a first "filter" imposing a necessary condition to XIFI federation participation. Exceptions can be allowed in particular cases: for example if an infrastructure provides unique services or capabilities that are needed by an experimenter and are not present in the XIFI federation. The operational requirements have been established taking into account the information collected through the surveys for the infrastructure owners of the five core backbone nodes and for the Use Case projects (see D1.1 [11]). This information has been elaborated in order to obtain a trade-off between the needs of the experimenters and the offer of the infrastructure owners.

Requirement	Comment
Data security and protection (for experimental data)	The experimental data should be kept private to the experimenter/developer that produced it. Only the experimenter can decide if his/her data can be published.
Availability	Since the XIFI federation is devoted to support early trials the availability requirements are lower than for production and mission critical systems. Nevertheless an availability >95% is required. This corresponds to an accumulated downtime of less than 2-3- weeks per year.
Support available on working hours	Each infrastructure should provide support from Monday to Friday during working hours.
Backup & recovery services (for experimental data)	The infrastructure should provide an on demand backup and recovery service for experimental data.

Table 1: Operational requirements

3.2.3 Federation Promotion

Members of the XIFI federation are required to contribute to activities related the promotion of the XIFI federation. As such they should be available to:

- Support the organization of events to advertise XIFI and FI technologies to local/regional entrepreneurs, SMEs, and researchers.
- Support the instantiation of XIFI showcases that target specific regional/local initiatives such as LivingLabs.

In general, federation promotion is organised as a separate work item in the project and is the responsibility of the XIFI federation office. The activities are documented in D9.1 [16] and D9.2 [17].

3.2.4 Quick Online Test

The XIFI Federation provides assistance in the form of a Quick Online Test on the requirements and criteria needed to become part of its federation. The Quick Online Test comprises five sections; namely XIFI federation introduction, legal compliance, operational compliance and technical compliance as well as contact details.

- XIFI Federation Introduction: Give an understanding on what the federation is endeavouring to achieve. It outlines a basic description of a XIFI node and its types.
- Legal compliance: Explains the legal requirements to be adhered to, if a node wishing to join the XIFI federation. Items include identity management, Memorandum of Understanding, Acceptable Usage Policy, Intellectual Property Rights and Confidentiality.
- Operational compliance: This clarifies the necessary procedures to support node functionality. It describes the need for a site local helpdesk and operational XIFI federating applications.
- Technical compliance: This describes the required resources to implement a XIFI node. It elaborates the following node aspects: Hardware, Software, Networking and any other additional IT support services.
- Contact details: For interested infrastructures joining the XIFI Federation, this page provides contact details, location and potential roles/levels they might want to participate in.

The Quick Online Test can be executed via the web portal which is currently hosted at http://qot.lab.tssg.org and is being migrated to the main XIFI web site at http://www.fi-xifi.eu. A description of the Quick Online Text can be found in Appendix A.

3.3 Deployment Architecture Reference Model

This section discusses the reference model for the physical and software deployment of a XIFI node based on OpenStack Grizzly, FI-WARE add-ons and XIFI tools. It is important to understand that there is no one model fits all, since the deployment architecture depends on, and is heavily related to resources available in an infrastructure. Dealing with existing infrastructures that connect to the federation has the impact that the deployment architecture must be adapted according to the existing hardware and to the planned upgrade of the infrastructure.

The following text is largely inspired by best practices in the deployment and operations of OpenStack based-clouds [18][19].

3.3.1 Concepts

3.3.1.1 Physical Equipment

In the deployment of a cloud-based data centre we deal with interconnected physical equipment that composes the physical architecture of the data centre. The most important equipment types for the definition of the deployment architecture are:

- Rack: Modern servers and network equipment are designed to be installed in a framework called a rack. A rack contains multiple mounting slots called bays, each designed to hold a hardware unit. Hardware may occupy more than one unit. Recent evolution for high density servers, introduces blade servers that are hosted in a blade (which allows packing several hardware component in a blade enclosure). Blade enclosures are mounted within racks.
- Server: A server is a node in the data centre (usually hosted in a rack) that offers computation and storage capacities. A server node in a cloud-based data centre may have different role according to his hardware configuration, and hence being able to host different services (that correspond to a given role). Generally speaking, server equipped with large number of CPUs

and RAM are more efficient for computational tasks, while server equipped with large amount of hard drives are more efficient for storage tasks. This discussion will become clearer in the next paragraphs that discuss node roles in an OpenStack based-cloud.

- Switch: Hardware equipment that allows the physical interconnection of different server nodes. Like a server, a switch may have different roles according to the network services it provides (e.g. management network or data network).

3.3.1.2 Node Roles

In a cloud environment, servers usually have different roles. In the following discussion we take into consideration roles usually adopted in OpenStack deployments. These roles are:

- **Controller (node)**. A controller node provides the central management for multi-node OpenStack deployments.
- **Compute (node)**. A compute node provides the computational capacity (i.e. virtual machines) to OpenStack deployments.
- **Block storage (node)**. A block storage node provides non ephemeral storage for virtual machines.
- **Object storage (node)**. An object storage node provides access to large storage solutions via Web APIs.
- **Object proxy (node)**. A proxy that distribute the objects to different storage nodes according to replica settings and region availability settings.
- **Network management (node)**. A network management node provides (dynamic) configuration on the VLANs that interconnect the VMs.

Furthermore XIFI will consider the following roles:

- **Load balancer (node)**. A node that in high-availability configurations, provides load balancing of requests among the available redundant services.
- **Monitor (node)**. A monitor node provides monitoring of resources included in a XIFI node.
- **Deployment (node)**. A deployment node provides the ability to control the deployment of a XIFI node, including a monitor node and all other nodes needed to run OpenStack and FI-WARE extensions.

It is important to underline that a node may serve different roles according to the OpenStack services it runs. Given the difference of type of service, different roles may perform better on different type of hardware. Accordingly, certain roles should not be covered by the same machine in a well-designed cloud deployment. In the next paragraphs we discuss quickly the different services. The distribution of services on actual nodes defines the role of a node and the architecture of the OpenStack deployment, according to the type of configuration of the services.

3.3.1.3 OpenStack Services

The XIFI installation of OpenStack considers the following services distributed on the nodes [20]:

- **Nova** [21]: Provides the management of computational resources. It includes three basic services: the scheduler, to define where the VM will be allocated, an API to remotely control the scheduler, the compute service that actually provides the VM on the single nodes and other support services.
- **Neutron** [22]: provides network management for OpenStack. It includes the following services: server to manage the network as service functionality for Nova, agent to apply the

configuration of the single nodes, DHCP-agent to automatically assign IPs to VMs, and other services. It requires specific plugins to configure the different network apparatus (e.g. OpenVSwitch [23]).

- **Glance** [24]: provides image management for OpenStack. It includes the following services: a registry that provides a catalogue of available VM templates and an API to control the services. Different back end are available for glance [25].
- **Keystone** [26]: provides identity and service registry functionalities.
- **Cinder** [27]: provides block storage (i.e. volumes) functionalities for OpenStack. It includes three basic services: the scheduler to define where the volume will be stored, an API to remotely control the scheduler, and the volume service that actually provides the storage;
- **Swift** [28]: provides object storage functionalities. It includes the following services: the proxy to accept API requests and to route them to storage services, the object storage that take care of the actual storage.
- **Horizon** [29]: provides a graphical user interface to support management and self-provisioning of cloud resources for the services mentioned above.

3.3.1.4 Network Services

As mentioned above, Neutron requires an actual plugin to be able to configure switches and creating VLANs in an OpenStack cluster.

- **DOVE**: The reference plugin for XIFI is a customized version of IBM's Distributed Overlay Virtual Ethernet (DOVE), provided by FI-WARE. DOVE is an SDN management solution for data centres that allows traffic shaping inside the data centre. It is based on OpenVSwitch.
- **OpenVSwitch**: as an alternative; we foresee the adoption of the standard version of OpenVSwitch.

Other network services are required to support the inter-node XIFI connectivity. These are currently under development. More details will be provided in the next version of this guide.

3.3.1.5 Other Services

XIFI deployment will require other services:

- **HAProxy**: to provide load balancing across OpenStack and FI-WARE APIs in the high-availability configuration.
- **XIFI Monitoring Management Middleware**: a middleware that is currently under development in XIFI to integrate physical and virtual infrastructure monitoring data collected from the nodes. The XIFI Monitoring Management Middleware provides adapter mechanisms for monitoring tools adopted by infrastructures (e.g. OpenNMS [30], Perfsonar [31]).

3.3.2 Physical Deployment Models

The physical architecture of a node influences the software architecture and QoS characteristics such as availability of services. Servers are usually hosted in racks and if all servers, for example, playing the role of a controller are in the same rack and power to the rack is interrupted, the cluster may not be available externally even though other services may be still running in other racks. Similar issues apply in case of switches. Therefore when possible, it is better to plan the physical architecture without a single point of failure.

3.3.2.1 Basic Physical Deployment

In a basic physical deployment, resources are not redundant and are not made resilient. In the simplest case, we will have a rack (or more racks) with a single power source that will power all servers part of the node including the switches that connect the servers. In the simplest configuration this requires:

- 1 controller node;
- 3+ compute nodes;
- 1 manager node;
- 1 switch 24 port (OpenFlow enabled).

Optionally, we can include as well:

- 3+ object storage nodes.

Such flat configuration is not recommended for production nodes, unless high-availability deployment cannot be achieved. Production nodes should refer to the high-availability deployment. Corresponding service architecture deployment is discussed separately.

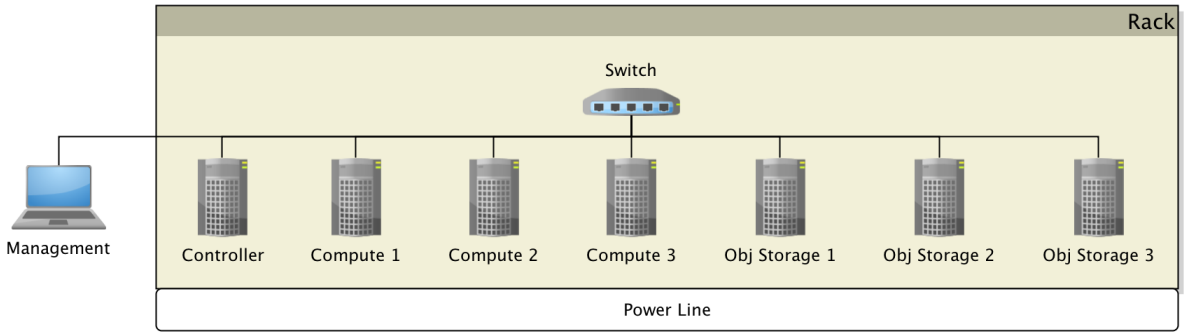

Figure 4: Basic physical deployment

Node Type	Recommended Hardware	Comments
Controller	• Processor: quad core • Memory: 8 GB RAM • Disk space: 1TB • Network: one 1 GB Network Interface Card (NIC)	More disk space is required if you do not plan to deploy an object storage or other storage backend to act as back end for the VM registry (2TB). RAID configuration is suggested to increase controller reliability. See [32] and [33].
Compute	• Processor: 16 or 32 cores • Memory: 64 GB RAM or 128 GB RAM • Disk space: 2 TB • Network: 2x1 GB Network Interface Card (NIC)	If you adopt a 16 core server, you should have 6+ servers. RAID configuration can be used but it is not recommended. The disk space, unless you have also a SAN in your data centre, will be as well used for block storage services (volumes) in shared modality. See [32] and [34].

Object Storage	• Processor: quad core • Memory: 8 GB RAM • Disk space: optimized for cost per GB (at least 4TB per node) • Network: one 1 GB Network Interface Card (NIC)	RAID configuration is highly discouraged. See [35].

Table 2: Hardware recommendations for basic physical deployments

3.3.2.2 High Availability Physical Deployment

In a high availability physical deployment, resources are redundant and they are located to be resilient. The objective of high availability deployment is to minimize:

- System downtime - the unavailability of a user-facing service beyond a specified maximum amount of time, and
- Data loss - the accidental deletion or destruction of data.

To avoid system downtime and data loss it is important to avoid the presence of single point of failure. Either in the hardware or in the software. In this section we highlight the deployment from the hardware perspective.

We assume to have two (or more) racks where the nodes are replicated with separate line power supply. This will ensure that if a power line will go down and hence turn off a rack, the second power line will be still accessible. As better alternative it is possible to consider single racks with support for 2 independent power lines. In this case all equipment in the rack should be equipped with 2 power supply units attached to the 2 power lines of the rack. This reference configuration requires:

- 2+ controller node;
- 6+ number of compute nodes;
- 3+ object storage nodes;
- switch 24 port 1GB and 10GB up-link (OpenFlow enabled);
- 1 manager node (also a laptop may do the work).

This reference configuration is the recommended one for XIFI nodes. Tweaks may be applied according to specificity of XIFI nodes.

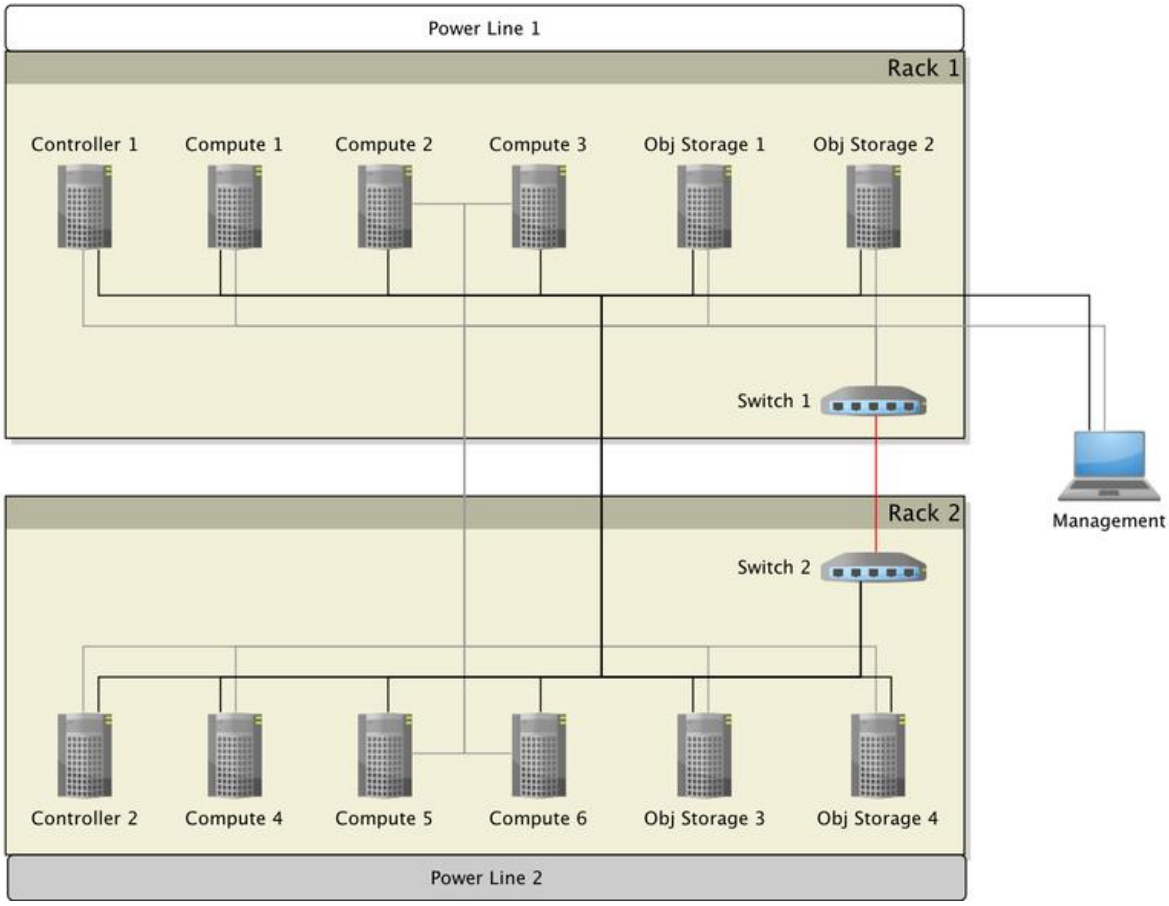

Figure 5: High availability physical deployment

Services are discussed separately in the service architecture section.

Node Type	Recommended Hardware	Comments
Controller	• Processor: 4 or 8 core • Memory: 12 GB RAM • Disk space: 1TB • Network: 2 x 1 GB Network Interface Card (NIC)	More disk space is required if you do not plan to deploy an object storage or other storage backend to act as back end for the VM registry (2TB). RAID configuration is suggested to increase controller reliability. See [32] and [33].
Compute	• Processor: 16 or 32 cores • Memory: 64 GB RAM or 128 GB RAM • Disk space: 2 TB • Network: 2x1 GB Network Interface Card (NIC)	If you adopt a 16 core server, you should have 12+ servers. RAID configuration can be used but it is not recommended. The disk space, unless you have also a SAN in your data centre, will be as well used for block storage services (volumes) in shared modality. See [32] and [34].
Object Storage	• Processor: quad core • Memory: 8 GB RAM • Disk space: optimized for cost per GB (at least 4TB per node) • Network: 2 x 1 GB Network Interface Card (NIC)	RAID configuration is highly discouraged. See Object Storage System Requirements [35]

Table 3: Hardware recommendations for high availability physical deployments

3.3.3 Services Architecture Deployment Models

In the previous section we discussed the physical deployment and listed the nodes type needed for that. But we didn't enter in any details regarding the services to be deployed on the nodes. Different services are corresponding to the different roles [19].

3.3.3.1 Basic Architecture

In the basic deployment services are not configured in high-availability. In this section we details which services are supposed to run on the different nodes discussed in the section 3.3.2.1 “Basic Physical Deployment”. It is important to underline that we foresee the computational node to cover as well the block-storage node role through the set-up of a shared filesystem (e.g. NFS). Also, we foresee the installation of XIFI specific services on the controller node, if the node offers enough capacity to run them.

The controller node will host all the services related to the management of the XIFI node. The services include:

- The nova-scheduler service, that allocates VMs on the compute nodes.
- The cinder-scheduler service, that allocates block storage on the compute nodes.
- The glance-registry service, that manages the images and VM templates. The backend for the registry maybe the controller node, or the Object Storage if included in the deployment architecture.
- The neutron-server service, that manages the VM networks.
- The swift-proxy service (optional), that manages request to the object storage nodes.
- The nova-api service, that exposes the APIs to interact with the nova-scheduler.
- The cinder-api service, that exposes the APIs to interact with the cinder-scheduler.
- The glance-api service, that exposes the APIs to interact with the glance-registry. If the object storage nodes are deployed, we recommend their usage as back-end for glance [25].
- The keystone service, that manages OpenStack services in a node.
- The horizon service, that provides a dashboard for the management of OpenStack in a node.
- The IdM GE service, that provides identity management for users.
- The SLM GE service, that provides scalability and elasticity management. it is connected the SOM GE service, hosted in the Main XIFI node.
- The XIFI-MMM service, that collects monitoring data for physical appliances.

The DCRM GEs is not listed as it is essentially a plugin to nova-scheduler and neutron.

The compute node will host all the services related to the provisioning of VMs and block storage. The services include:

- The nova-compute service, that manages VMs on the local node.
- The cinder-volume service, that manages block storage on the local node.
- The neutron-agent service, that manages VM networks on the local node.

The object storage node (optional) will host all the services related to the provisioning of object storage. The services include:

- The swift-account-server service, that handles listing of containers.
- The swift-container-server service, that handles listing of stored objects.
- The swift-object-server service, that provides actual object storage capability.

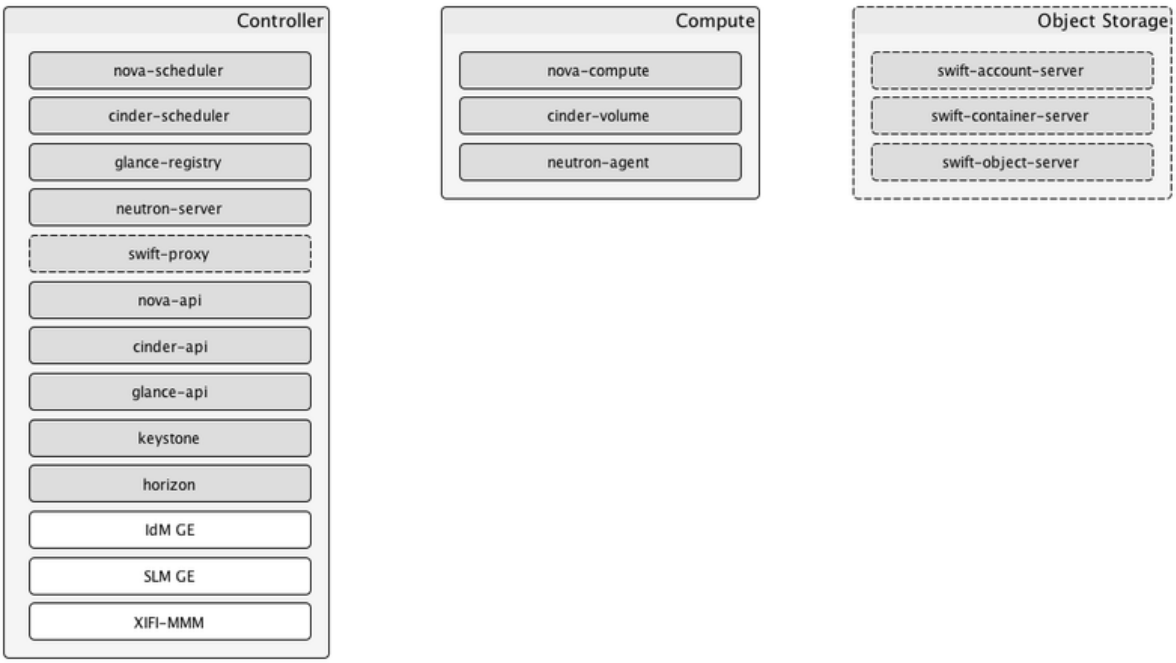

Figure 6: Service per node in the basic architecture deployment model

3.3.3.2 High Availability Architecture

In the high availability services are redundant and they are located to be resilient. In this section we details which services are supposed to run on the different nodes discussed in the section 0 "Table 2: Hardware **recommendations for basic physical deployments**

High Availability Physical Deployment". In this section we discuss the deployment from the software perspective. The deployment, except the injection of services to support high-availability of the controllers (the other are in high-availability modality by default so to say), is very similar to the basic one. In fact in OpenStack, computational, block storage and object storage nodes, are handled by the different scheduler to provide high-availability. The issue is to guarantee high-availability as well to the controller.

Generally speaking, high-availability can be provided in two modalities:

- active/passive: in this configuration, a main controller operates the resources and in case of a problem, it switches the request to a backup controller.
- active/active: in this configuration, a number of controller operates the resources at the same time, in case of a problem to a controller instance, requests are not issued anymore to that node.

In this paragraph we refer to the active/active configuration.

The controller node will host, additionally to the services mentioned in the previous section, the services needed to ensure the high-availability of the controller node:

- ha-proxy service, that provides load balancing across OpenStack and FI-WARE APIs in the high-availability configuration.
- pacemaker service[36], that provides high-availability for neutron and other services.
- galera service[37], that provides high availability for databases used by the different services.
- RabbitMQ service, present as well in the basic deployment, should be configured for high-availability policy support.

More information is available in [38].

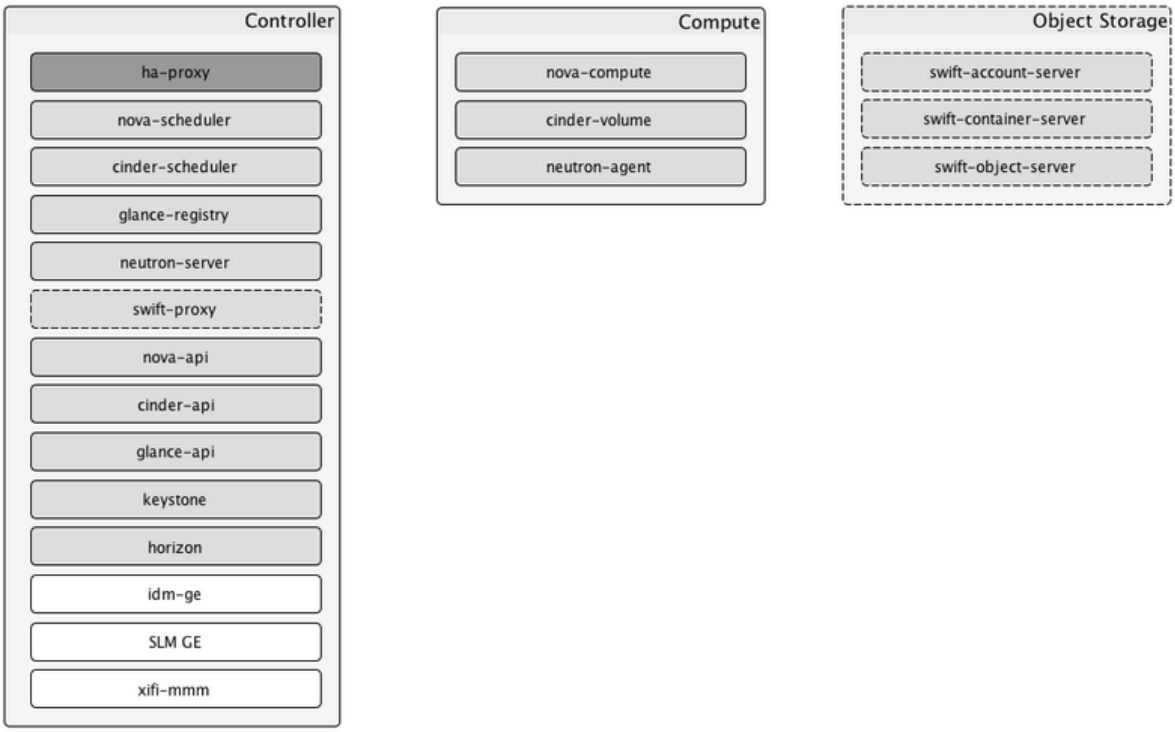

Figure 7: Service per node in the high availability architecture deployment model

3.3.3.3 Block Storage Configuration

Different modalities to run block storage services are possible in OpenStack, refer to [19] [39] for a complete discussion. In this section we refer to the default configuration selected for XIFI, that relies on shared filesystem across the compute nodes. In this configuration, each compute node has a large storage capacity that is share through a distributed filesystem that allows the disks of the different compute node to be seen as a single drive. This configuration allows for high scalability and easy live-migration and does not require for dedicated nodes to the block storage. Of course, the solution may have drawbacks such as high network i/o in case the block is accessed from a remote virtual machine instead than locally. The default shared filesystem solution foreseen in XIFI is NFS.

3.3.3.4 Network Configuration

In order to set up correctly, the network Openstack distinguishes four logical networks [40]:

- Management network
- Data network
- External network
- API network

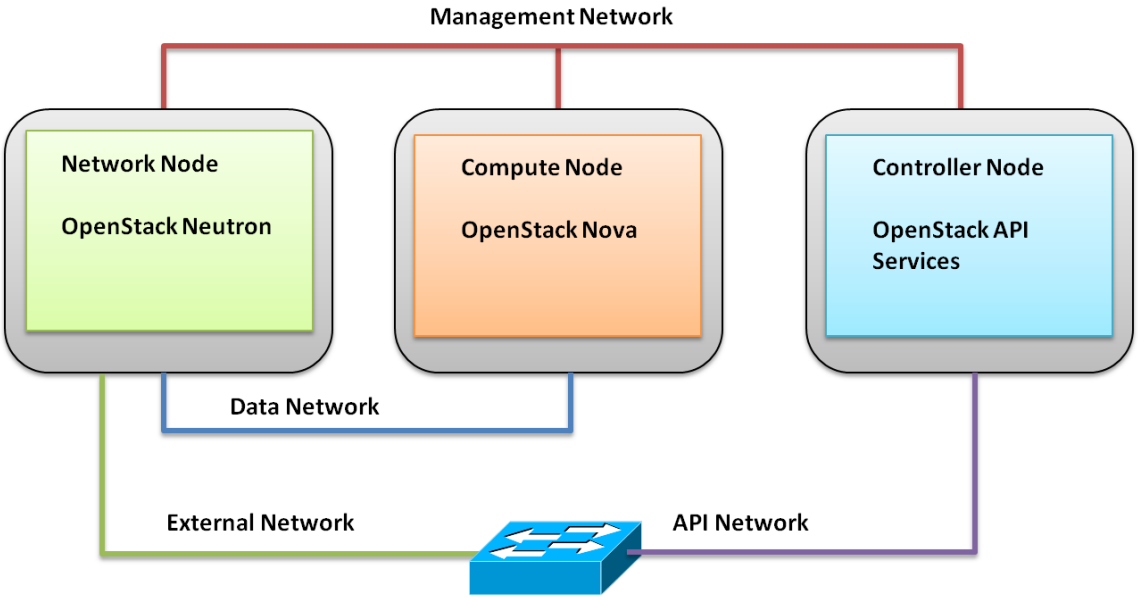

Figure 8: OpenStack Networking Architecture

Basically the first two are intended as "internal networks". The former is used for management purpose and provides connection between OpenStack components. The latter is used for VM data communication. These networks should have IP addresses reachable only within data centre. As stated in OpenStack Operations Guide[41] the recommended option is to use a management network with separate switch and separate NICs: "this segregation prevents system administration and monitoring system access from being disrupted by traffic generated by the guests themselves".

The external network and the API network are intended as "public networks" because they allow inbound connections: the IP addresses should be reachable by anyone on the Internet. In particularly the external network is used to provide both VMs inbound and outbound connection from VMs, whereas the API network exposes all OpenStack APIs to tenants. These networks can be merged into a single network.
If it is not possible to have physically separated networks, i.e. different switches and NICs for different networks, is possible to use VLANs to segregate network data. In this case the switch must be configured accordingly (the configuration steps are different for different networks). For instance, assuming nodes with at least 3 NICs, we can use the interfaces as follow:

- eth0: Management Network
- eth1: External/API Network
- eth2: Data Network

Each interface must be configured with the correct VLAN ID.

3.4 Installation Procedures

Federation Networking

In order to join the federation a node must be connected to internet and be reachable by the master. For this purpose a reliable internet connection is obviously needed. Network redundancy is strongly recommended, which means that every component should be present twice. There should be two firewalls sharing a "heartbeat" connection to ensure fail-over capabilities. Both firewalls must be connected to the router that gives internet access so that, in case of a failure, the other firewall still guarantees internet connectivity. Each firewall should be connected to a different switch. If a switch fails the hosts residing in the network will not lose internet access. It is very important that the switches are linked using a switch stack method or just by choosing a port to trunk the switches. Last but not least, each physical host should have two network interfaces in fail-over mode connected to two different switches. If a switch fails the failing network interface link will be discarded and the communication will switch to the other interface, which in turn is connected on the other switch.

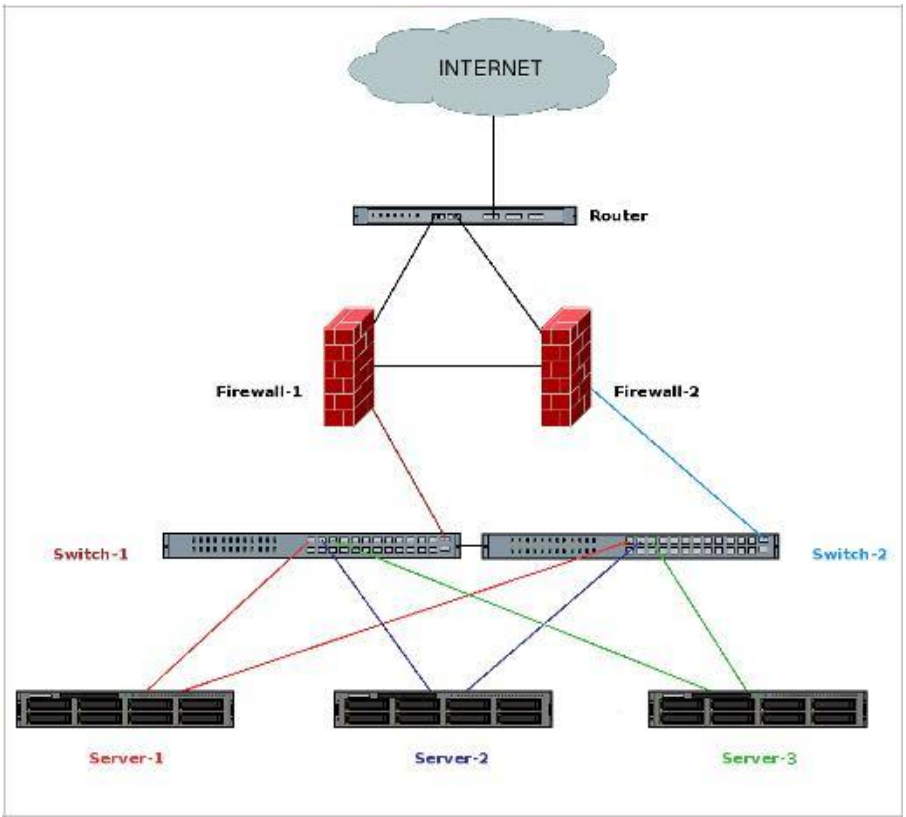

Figure 9: Federation Networking

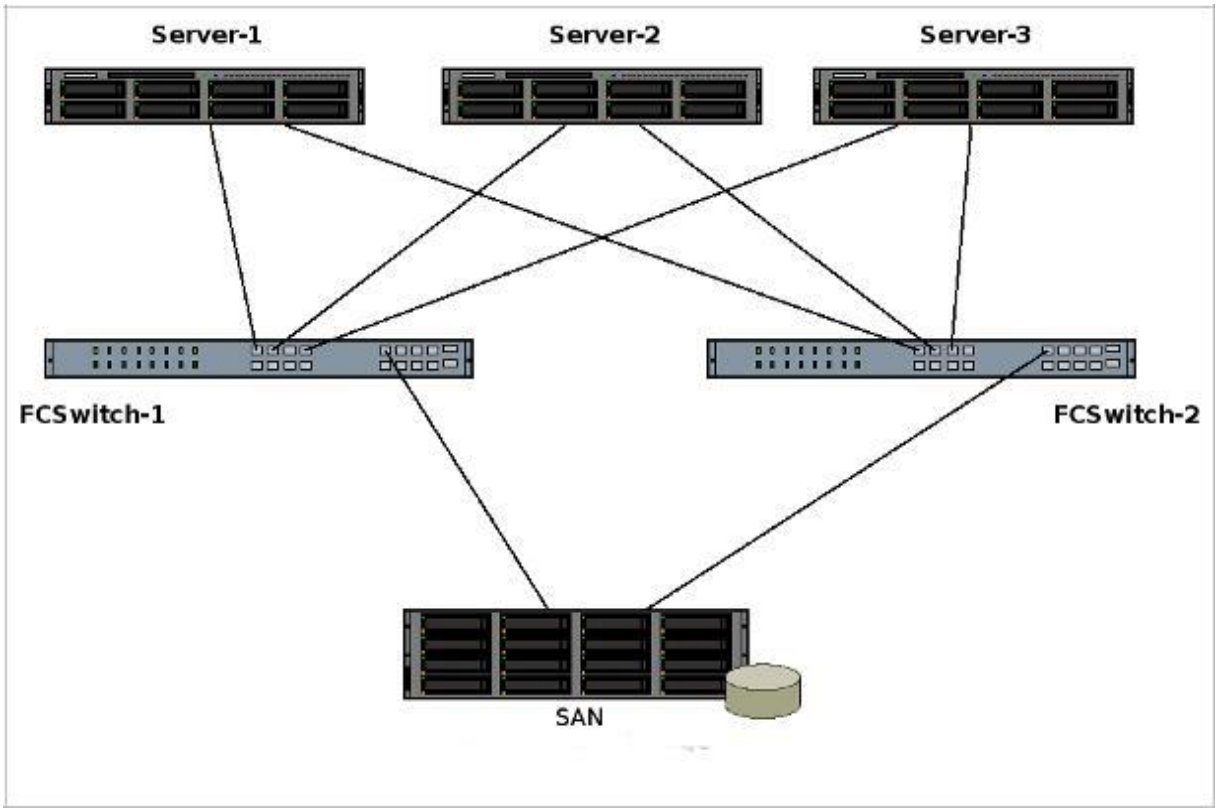

Figure 10: Node storage

3.4.1 Physical Level Procedures

The framework of supporting procedures, protocols and tools required as primitives to implement the processes for set-up, configuration, deployment and maintenance herein is denoted as the physical level procedures. In contrast to the logical level procedures discussed in subsequent sections, the physical level mainly deals with platform and service provisioning close to the node infrastructure and to the underlying communications infrastructure that connects XIFI nodes. The latter includes procedures and protocols to establish, verify and maintain reliable and secure end-to-end application layer connectivity across the XIFI federation.

This section provides an initial list of procedures described in a semi-formal way that are considered a repository of primitives that may be used in the scope of implementing the bootstrap, operation and maintenance of XIFI nodes in a federation. Clearly, this list is not complete for the time being and will definitely grow over time, progressing along with the evolution of the federation tools and the GE deployment. The procedures given in the following can be roughly categorized as:

- Testing (connectivity, performance, conformance);
- Service maintenance (software distribution, installation, verification, updating, patching);
- Configuration (nodes and services, configuration, testing and maintenance);
- Security (installation and maintenance of security functions, credential provisioning);
- Monitoring and reporting (installation and maintenance of node monitoring functions, logging, reporting tools).

3.4.1.1 Testing

Testing procedures aim to verify the network layer connectivity between XIFI nodes, to verify access

to services (i.e. federation tools and FI services), to verify conformance of the service API baseline functions, and to verify the minimum performance requirements for (some of) these. Tools required are for:

- Connection testing (client and server capable to create and accept multiple test connections);
- Service detection and identification (e.g. port scanner to find a service port and to get the service identification);
- API conformance testing (i.e. a baseline service client testing availability and correctness of responses for a tool or GE in a non-intrusive way);
- Performance testing (e.g. bandwidth, communication round-trip delay, service/tool response time for initial and consecutive requests).

Testing procedures are used by the maintenance management processes described in subsequent sections and may be callable from the user, security and first level support portals for a quick functional verification and presence detection.

Procedure	Purpose
Point-to-point connectivity test	Verify connectivity between two XIFI nodes on the transport layer using IP protocols.
Point-to-multipoint connectivity test	Verify connectivity between multiple meshed XIFI nodes on the transport layer using IP protocols-
Service detection	Scan a well-defined port range used by known XIFI tools/services for a given IP address and try to determine the service type. Then probe if the remote service is responding to a request to identify itself.
Service API verification	Select a suitable service test client from a repository of clients that matches the identified service. Then identify the API(s) provided by version and protocol supported and conduct a baseline API test that matches the API assumed. Determine if the responses and data format of response payload matches expectations. Produce an API test report stating if the tested API provides the minimum set of requirements for this specific API.
Connection performance test	Verify connection performance for a well-defined set of metrics and compare results with given minimum values. Produce a test report stating if the link under test performs better than the minimum requirements.
Service performance test	Select a suitable test client for the service under test matching the APIs supported. Verify initial and sustained response times of the service under test. Produce a test report stating if the service under test performs better than the minimum requirements.

Table 4: List of testing procedures

3.4.1.2 Service maintenance

Service maintenance procedures define how to distribute, install, verify, update and remove federation tools and FI services. They partly rely on the testing procedures described above and are used by the maintenance management processes described in subsequent sections. Tools required are for:

- Creating software distribution packages (incl. installation scripts; for a certain physical distribution media or for download);
- Verifying completeness and sanity of the installation (as far as reasonable and not covered by

platform tools already available);

- Interface testing (client software that usually is specific for a certain service; must be provided or supported by the software provider/developer along with suitable documentation and maintenance instructions);
- Tools to distribute and apply bug fixes or patches (excl. update package handling, which usually is part of a distribution package tool-chain);
- Software removal tools (as far as not already provided by the platform; pre-removal dependency checks are mandatory for this tool, so it is specific for a target that shall be removed, or some subset of similar targets to be handled jointly).

Service maintenance procedures are also used by the maintenance management processes described in subsequent sections since software maintenance both is a process to bootstrap and to maintain a XIFI node and its integration with the federation, hence utilizing the same primitives and tools.

Procedure	Purpose
Create a software distribution package	Bundling a software package for later installation on a target node including installation, test and maintenance scripts if suitable. Distinguish between uses for initial or update purposes and consider deployment platform variations if possible. Always include safe removal scripts and instructions/documentation. The distribution package in XIFI is the Infrastructure Toolbox.

Table 5: List of service maintenance procedures

3.4.1.3 Configuration

Configuration procedures aim to configure federation tools and FI services with regards to the FI services provided by a XIFI node towards the federation and the end-user. In contrast to the maintenance processes that partly overlap with this goal, the configuration addressed here is on-demand and driven by end-user and use case project demands. Hence, configuration procedures must be accessible through the end-user's portal and should be robust regarding malicious use. In particular, the configuration of 'non-conventional' services (i.e. node specific infrastructure access or specific/private services) is considered as in scope. In order to realize and verify configurations, configuration procedures defined here need to consider and make use of dependencies between tools and FI services and, in particular, must minimize the risk of failure or misconfiguration, or have to include fail-over, snapshot and back-up/restore configurations. Tools required are for:

- Node configuration and testing (as far as not already provided by the platform; incl. OSI layer 4 and 5 testing using specific clients);
- Service configuration and testing (mostly based on particular test clients and on the configuration interfaces of the particular tools; generic configuration tools are only available for the most basic services);
- Procedures regarding security, reliability and robustness verification enhancements are consider here in principle (a security dashboard might be foreseen utilizing these procedures but is not yet evolved very far; so far, these procedures cannot be detailed further and are for further study).

It is important to note that the configuration of user credentials and identity management is partly a configuration and a security issue and thus is handled by procedures from both categories.

Procedure	Purpose
Install a permanent local client/server for connectivity	Provide an endpoint for OSI layer 2...5 protocols that allow a quick test of network connectivity and of support for protocols required to detect firewall

Procedure	Purpose
tests	or routing problems on a single click.

Table 6: List of service configuration procedures

3.4.1.4 Security

Security related procedures aim to control and configure in particular access to a XIFI node through a federated service (e.g. the federation portal), and are mainly addressing user identification, authentication and authorization, implement resource access policies, and maintain the trusted and secure communication between nodes in the XIFI federation. In particular, procedures of this category will provide the means to create communicate and revoke credentials. A particular focus is set on the integration of 'non-conventional services' (e.g. regarding their particular vulnerabilities due to a deep integration with the local hardware platforms, potential legal constraints, or certain privacy concerns). The monitoring sub-system can be considered subject to security and privacy constraints, for example. Tools required are for:

- User management (as far as not already provided by the FI-WARE Identity Management (IdM) GE; regarding access control to XIFI node resources and roles of users in the scope of the federation, where a distinction between end-users and particular manager roles is needed);
- Public/private key provisioning and revocation (as far as not already provided by the FI-WARE IdM GE; in particular authorization of accessing 'non-conventional services' or exclusive use of singular resources);
- Secure connection handling and trust establishment (e.g. proper configuration of SSH/SSL and HTTPS set-up and testing - the latter with regards to nodes newly joining the XIFI federation and their compliance with the security and privacy requirements of the federation).

Procedure	Purpose
Add an end-user to the federation	Make a user and his role known to the federation. Distribute credentials to all or some XIFI nodes and get confirmation from the node(s) addressed that the new user has been accepted by the node in his proposed role.

Table 7: List of security procedures

3.4.1.5 Monitoring and Reporting

Procedures related to monitoring and reporting aim to support bootstrapping and operating a XIFI node and to log and automatically create log reports for documentation and approval of certain logical procedures as discussed in the following section. Monitoring and reporting also is relevant for verifying if a XIFI node complies with the federation requirements and if it is compliant with any service level agreements established. Distinct procedures are required for node bootstrap and node operations. The latter is putting the focus on the monitoring needed to determine and document proper node operations when in use by end-users, while the former puts the focus on supporting initial installation, configuration and sanity checking of a XIFI node and thus has testing objectives rather than monitoring objectives. Thus, different tools are required for the two objectives namely for:

- Monitoring the bootstrap process (with regards to testing compliance with FI service hosting capacity requirements such as storage and computing resources available, and communication performance requirements for inter-node communication; also monitoring availability and performance of certain platform services and their compliance with federation requirements);
- Monitoring regular operations with a focus on XIFI node performance (with an objective of quantifying a nodes utilization/load and its contribution to the federation; also regarding the

metrics determining if a node complies with its service level agreements);

- Monitoring regular operations with a focus on failure detection (regarding the detection of short-term service disruptions and permanent failures of a service enabling to decide if a XIFI node may need to be taken out of service; also regarding detection of a lack of compliance on interface and performance level after applying configuration changes or updates that may harm the federation).

Procedure	Purpose
Produce a conformance test report	Provide documentation that a certain test for conformance with requirements has been conducted. Follow a formalized layout suitable for the kind of test conducted, documenting at least the date and time, type of test conducted, identification of individual and organization conducting the test, identification and kind of test facility, test method, test accuracy limits and test results. Also produce a summary statement if the test succeeded or failed. Optionally produce a recommendation for the use of the test results provided.

Table 8: List of monitoring and reporting procedures

3.4.2 Logical Level Procedures

Once a master-node structure has been created everything is manageable from the master node: it has visibility over the federation virtualization environment. It is very important to understand that each node must be a complete self-working virtualization environment and functional checks must be done to accomplish the master needs. This way, each node has to know just its own environment, without having to know the behaviour of the master or other nodes.

To test the service availability each federation must pass a simple checklist:

- Storage visibility;
- Network compliance;
- Creation of new virtual machines;
- Check the communication with the master (ping command).

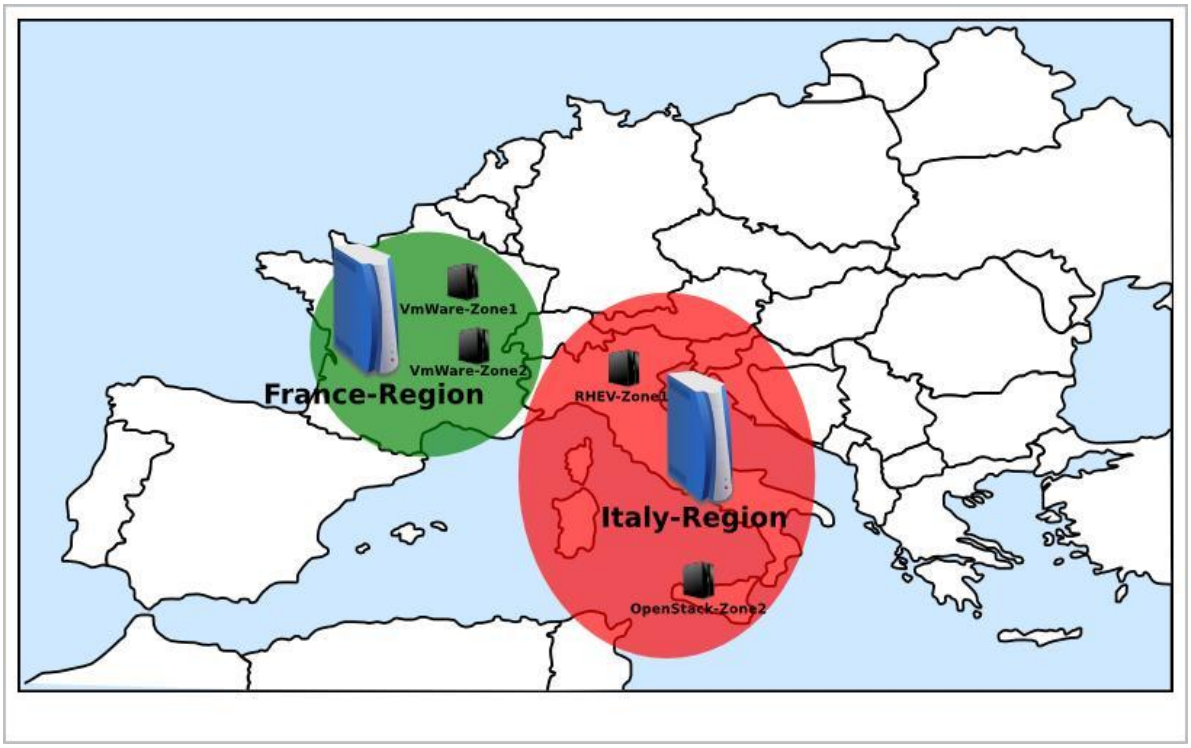

Figure 11: Federations associated with regions

Each node can belong to a specific zone which can belong to a region. Regions could represent a geographic location and a zone is the way the master separates different geographically distributed sites or environments. A zone is used to isolate traffic within a region.

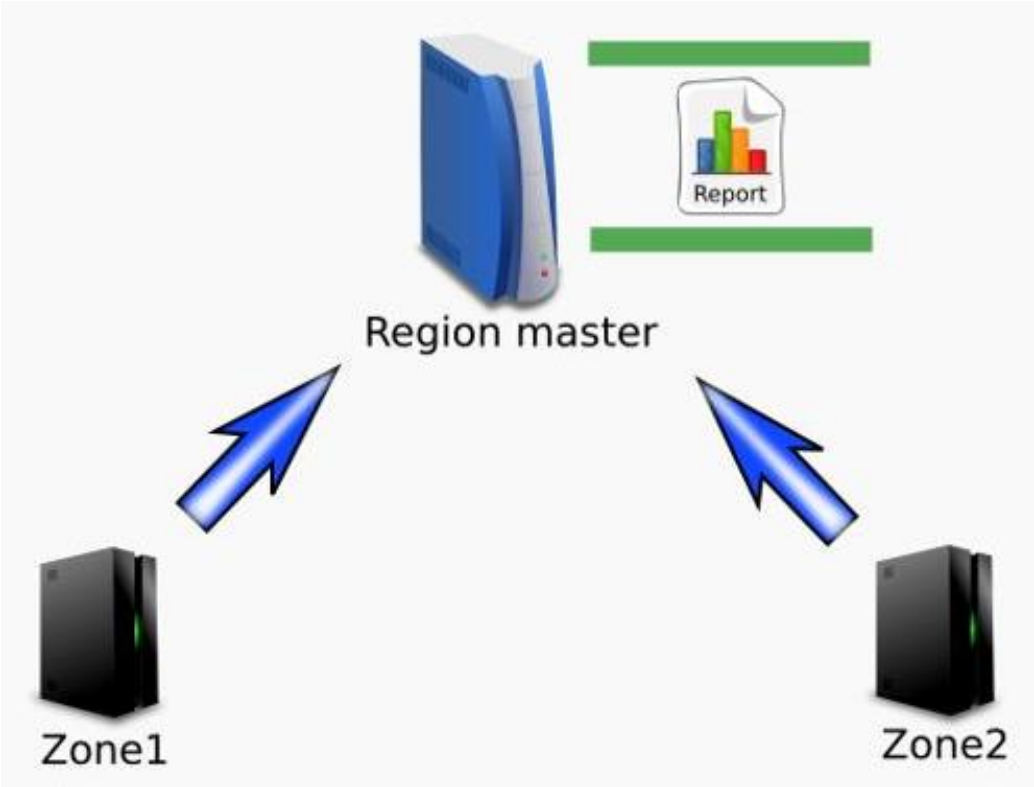

Figure 12: Report generation per region

A region is also responsible for underlying regions and zones and for summary report generation. It keeps a database containing the underlying layers information tracking the complete infrastructure and virtualization environment status.

When a new virtual machine needs to be created the user interacts with the master. Via a user interface he can choose in which region and zone the new virtual machine will reside, implicitly determining the virtualization environment. Every action can and should be performed from the master. The request is submitted to the master who dispatches the action to the right node.

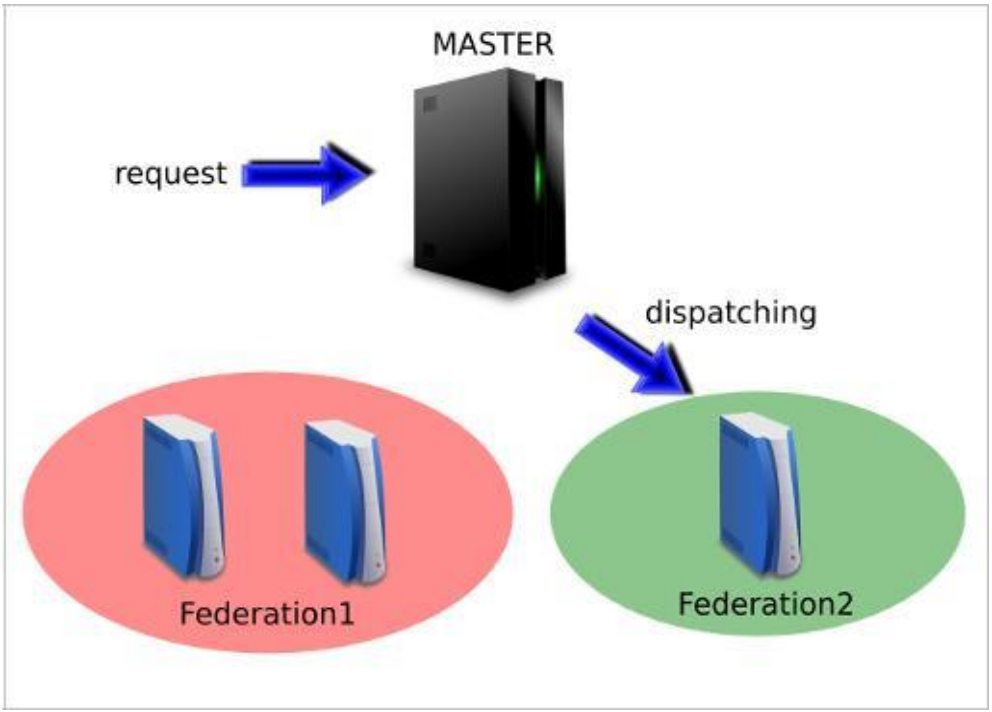

Figure 13: Dispatching of requests from region master to federation member

In XIFI we do not have a single stack for the whole federation, but the nodes must be logically connected to the master in order to build the infrastructure. In this way users can issue the creation of a resource (a virtual machine) by accessing the web portal in the master region. For this purpose we can imagine a kind of hierarchy in the infrastructure allowed by OpenStack regions. Regions have a separate API endpoints for each installation, allowing for a more discrete separation between nodes. Thus, users wishing to run instances across sites have to explicitly select a region. Currently the OpenStack dashboard (Horizon - at the time of writing) can manage a single region, for this reason each region needs its own dashboard. This means also that once a request has been issued to the master, it must be redirected to the right region API service.

In XIFI nodes all of the OpenStack components must be configured so to be integrated with the entire infrastructure and to fulfill the requirements of XIFI project. Specifically the components must be configured so to allow the use of OpenStack regions. Follow some advice about how to install the components in the nodes.

OpenStack Identity: each node should not have a dedicated Identity Manager, but should join the overall account manager, managed by the federation master (IdM). It is very important that each component in the federation node points to the correct centralized Identity Manager IdM.

OpenStack Compute: in each node should be present at least one instance of this component to allow the creation of virtual machines. Compute component is composed by the hypervisor and gives virtualization capabilities.

OpenStack Image Service: this service could be present both in the node and in the master node.

Anyway, we should consider possible latency issues due to the configuration of the node pointing to the OpenStack Image Service of the federation master. Thus, a node local Image Service is highly recommended.

OpenStack Networking: this component is responsible to maintain an aligned configuration of the infrastructure networking, reflecting the overall network connections requirements (vlans/routes/addressing).

OpenStack Dashboard: each node has a dashboard to manage and control all the Openstack components.

OpenStack Object Storage: this component should be configured on every node to provide enough space for the virtual machines. The storage configured in each node could be addressable from the master node or other nodes, but this solution should be avoided due to latency issues.

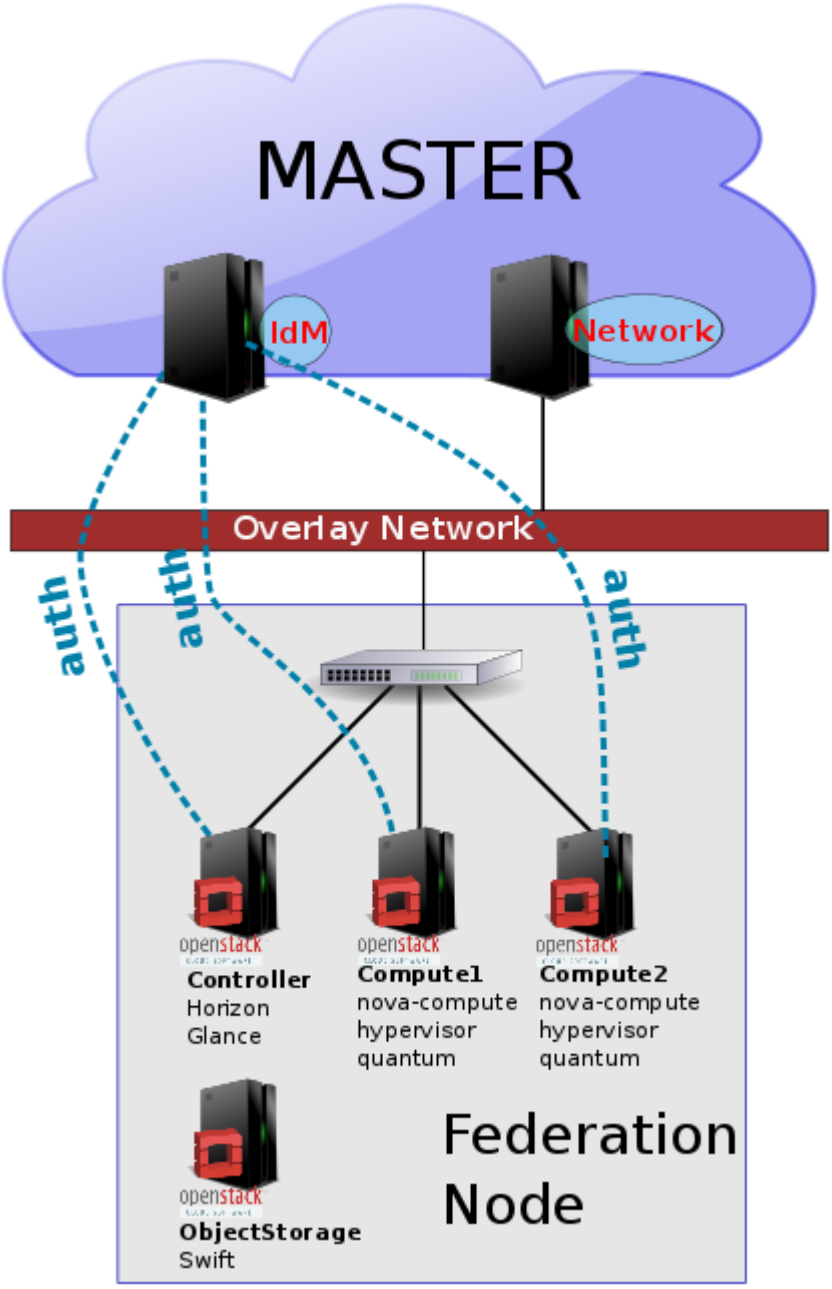

Figure 14: Components interaction

3.5 Federation Joining Support Levels

This section describes the different levels of support that will be given to an infrastructure willing to join the federation as well as the decision process of the support level to be provided. Three different levels of support are defined:

- Methodological support. The XIFI project provides guidelines and documentations on how to include a given infrastructure or service into the federation taking into consideration its potential relation with Generic Enablers, Federation services and technical and operational infrastructural constraints. A helpdesk service is also provided to support the integration of new infrastructures by their owners. This support is generally accessible only to any infrastructure/node including **Third party nodes/infrastructures**.
- Shared deployment support. The XIFI project shares the cost of the integration activities with the third parties; it coordinates the work on the XIFI federation side to ensure compatibility with Generic Enablers, Federation services and technical and operational infrastructural constraints. This support is generally accessible only to **FI-PPP member nodes/infrastructures** or **Full member nodes**. **Third party nodes/infrastructures** can access this support level only if the data sources they may provide to XIFI federation is of great value for XIFI adopters.
- Full technical/operational support. The XIFI project drives the integration activities with the partial support of the owners of new infrastructures to ensure compatibility with Generic Enablers, Federation services and technical and operational constraints. This support is generally dedicated only to **Full member nodes**.

The following subsections will describe how the XIFI federation will decide on the appropriated level of support to be delivered to new infrastructure in a negotiation process and describe for each level of support which support services will be included as well as support services that will be excluded for a given support level. The support levels built upon each other, that means that each extended level will include the services of the lower level support level (e.g. the Shared Deployment Support Level will include all services for the Methodological support level).

3.5.1 Support Level Decision Process

The Support Level Decision Process is a mechanism through which support requests are validated for further action based on the support agreement in place between the node operator and the XIFI federation.

In a traditional access control structure, a protection mechanism mediates access to protected objects (passive system entities) by subjects (active system entities) depending on a set of preconditions. In the case of the Support Level Decision Process, the base process treats the node operator (or their delegated contact) as a subject and the support resources as objects to be meditated by the process.

The rules used by this process to validate a support request are contained in the support agreement with the given node operator.

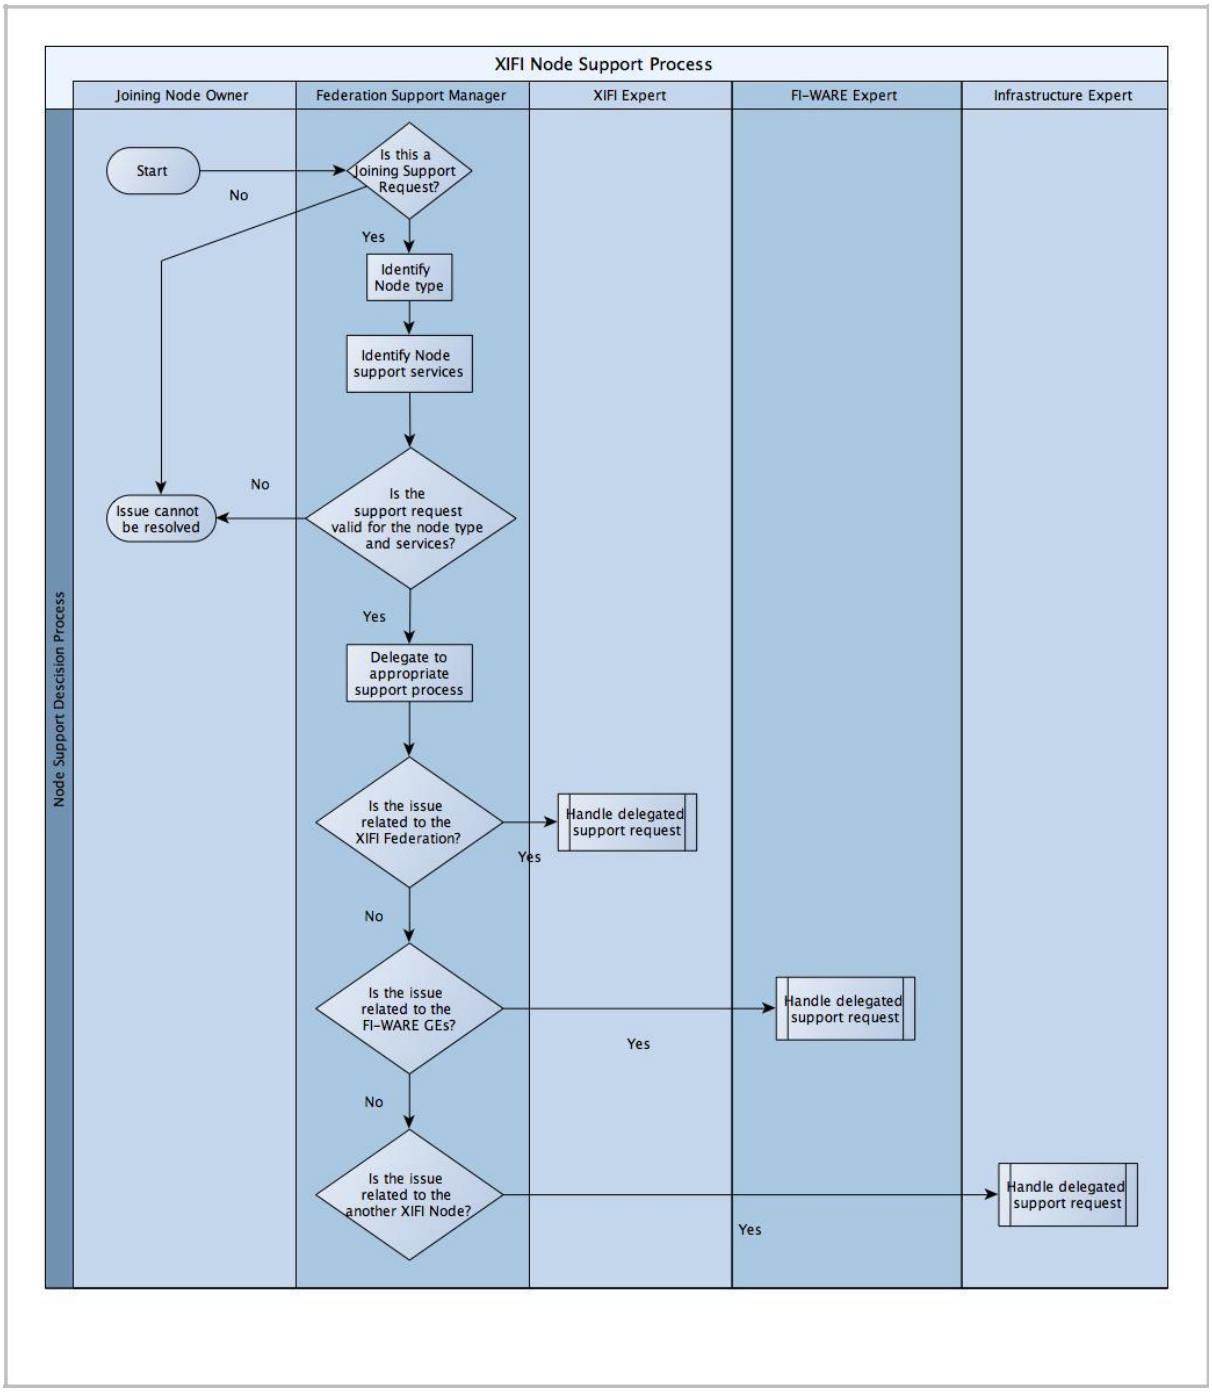

Figure 15: XIFI node support process

In the diagram, the support request is initiated by the node owner. This is then handled by the Federation Support Manager (or identified delegate). They verify the identity of the support requester and check the support agreement. If the support request exceeds the scope of the support agreement, the requestor is notified that the request cannot be resolved. If the request is within the scope of the support agreement, the process of delegation to the appropriate support process is started. From the point of view of the Support Level Decision Process, each of these processes is a predefined process with its own capability of interaction (via the XIFI support help desk) and resolution.

3.5.2 Methodological Support Service

The methodological support XIFI is offering to nodes willing to join the federation comprises multiple actions. The methodological support service of XIFI includes the following services:

- Access to documentation: all documents helping in informing and helping an infrastructure that wishes to join the XIFI federation will be provided. This includes but is not limited to the following documents:
    - 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 (included in this deliverable). This documentation describes the procedures and process an infrastructure has to take 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 FI-WARE 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 FI-WARE on these GEs (Catalogue entries, Admin- and Installation Guide, User- and Programmer-Guide and the unit testing plan.
    - Documentation on connecting different datasources to GEs (for nodes and infrastructures): This includes all information on the GEs offered by FI-WARE 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, ...
    - 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.
- Test tool (only for nodes): The test tool is used in order to check whether the technical requirements are fulfilled by a node that wishes to join. This test tool will be developed by XIFI.
- The Infrastructure Toolbox (only for nodes): a software distribution that facilitate 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 ticket;
    - Monitoring the status of the tickets;
    - Contact XIFI help desk.

For practical reasons there will be only one XIFI-helpdesk that process tickets and help-requests. 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 in most cases to an expert. Help requests on issues not related to set up as a new node or connection of an infrastructure data source will not be processed.

- Depending on frequent topics that are arisen at the helpdesk a FAQ will be provided in the wiki pages of XIFI.

The methodological Support service excludes all services not mentioned in this section. This applies especially for support services that lead to further development efforts like the

implementation of new adapter. The only exception to this are tickets that deal with bugs and problems investigated in the test tool and in the infrastructure toolbox.

3.5.3 Shared Deployment Support Service

The shared deployment support offers a more comprehensive support to potential nodes wishing to join the federation or, in some cases, to infrastructures willing to connect to a XIFI node. Besides the support already provided by the methodological support the following additional support services are offered:

- Potential nodes needing special contracts or agreements will be offered contact to the responsible persons within XIFI. This should basically already be a part of the negotiation process, but for questions that arise after the negotiation with the potential new node further support will be needed.
- Potential infrastructures with relevant data sources for the XIFI federation will be supported into connecting such data sources to a node part of the Federation, so as to offer such data source to the XIFI developers.
- Potential nodes not completely fulfilling the requirements needed to join the federation are guided through the necessary steps in order to fulfil the minimal requirements for joining as a node.
- An extended helpdesk offering also direct contact support via phone or video conferences (maybe including desktop sharing) helps to overcome problems that occur in the procedure of joining as a node. Basically this means that a XIFI expert dedicated to a certain ticket tries to support via direct contact with the person who issued the ticket in order to give him support to solve a certain issue.
- Step by step video tutorial may be offered to explain certain steps for joining as a new node.
- Necessary contact points to experts within XIFI will be established in order to solve a certain problem within the integration process (e.g. development of a certain adapter needed for joining the XIFI federation).

In order to help in connecting a new node a special adapter may be needed for certain infrastructures. The shared deployment process includes support in the way that is helps in providing information how this adapter may be built. Implementation requests besides minor modification of existing adapters are excluded.

3.5.4 Full Technical/Operational Support Service

The full technical/operational support offers the most comprehensive support from XIFI side. It offers nearly all-round carefree package to potential nodes willing to join XIFI. It includes all activities of the methodological and shared deployment service. Beyond that it offers also on-site support to the potential new node. However all infrastructure components have to be provided by the potential new node. Access to the potential new node needs to be provided to XIFI support people.

- Experts access the potential new infrastructure node directly to deploy and install necessary components at the potential nodes.
- Tools needed in order to join the federation will be deployed (or even first developed) and used by XIFI support people to establish the node connection.

4 FUTURE INTERNET DEVELOPER SUPPORT

This chapter describes the way the federation handles the support for Future Internet developers. It takes as a starting point the Use Cases described in D1.1. End user support is provided via a helpdesk tool available through the XIFI federation office portal. The objective of this section is, based on the following **Use Case diagram**, to describe:

- The support process;
- The flows and the actions to be perform at each stage of the process;
- The different roles assumed by the actors involved on the support process.

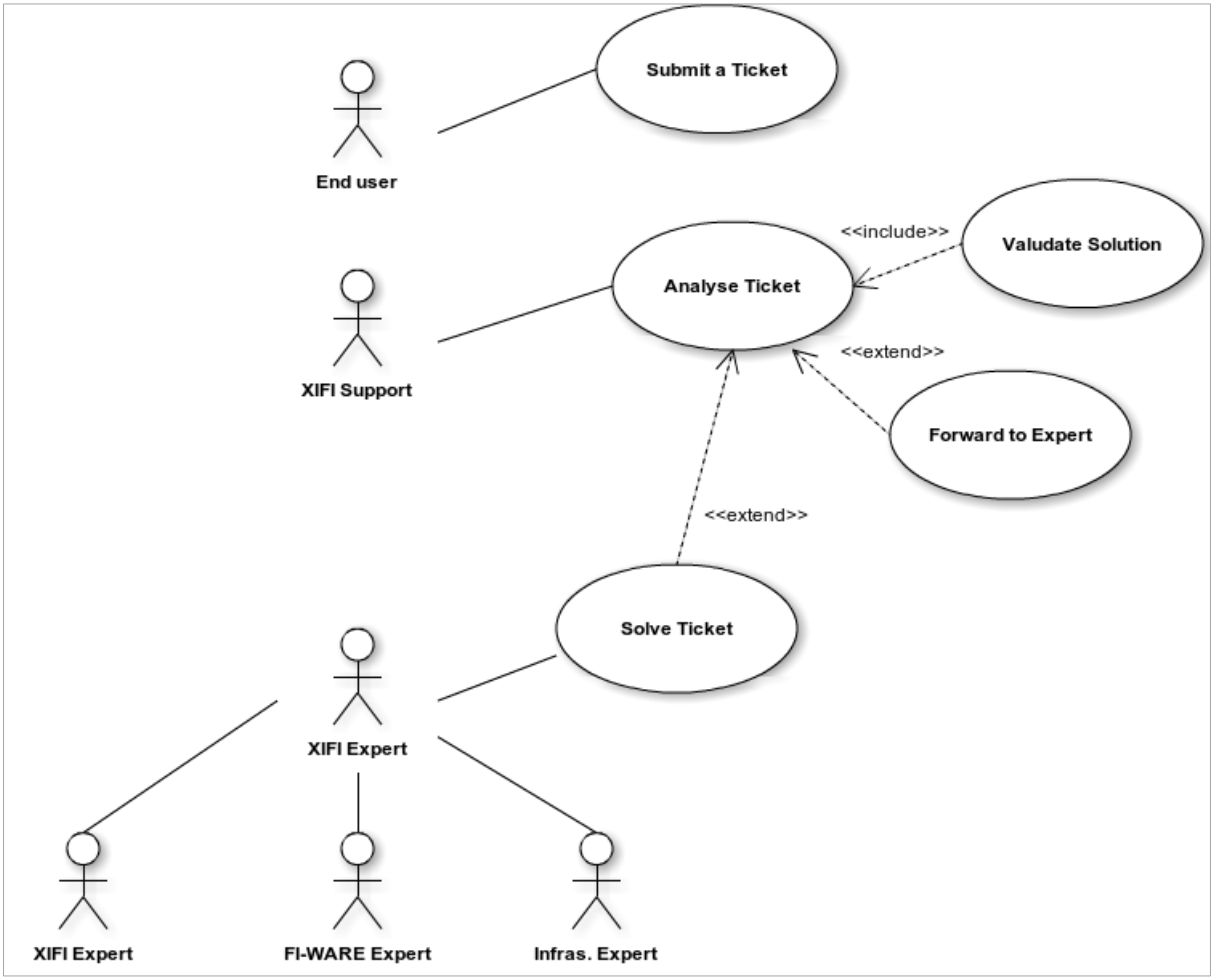

Figure 16: Support for Future Internet developers

4.1 Support Process

The support process aims to solve issues related to XIFI services and resources, address support requests and track problems until their resolution. The basic process dialog is presented below:

1. The end user identifies an issue or a request he wants to address to XIFI.
2. The end user connects to the Helpdesk through the XIFI portal.
3. The end user opens a ticket describing the problem and addresses it to the general XIFI helpdesk portal or to a specific XIFI node.
4. XIFI (first level of) support either identify the solution in the current list of FAQ, or pass it to

XIFI node help desk for second level of support.

5. XIFI support analyses, diagnoses and isolates the problem (second level of support). At this point it can either:
    - Identify a solution and propose it to the user;
    - Report the issue to an expert (third level of support).:
        - If the issue is related with XIFI services, it is forwarded to a XIFI expert;
        - If the issue is related with FI-WARE GEs, it is forwarded to a FI-WARE expert;
        - If the issue is related with an infrastructure resource, it is forwarded to the infrastructure expert;
        - Assuming the issue is reported an expert will perform an analysis and provide a solution.
6. XIFI support validates that the solution fixes the problem, possibly through an interaction with the user.
7. XIFI support provides the solution to the user.
8. User accepts the solution.
9. XIFI support closes the ticket.

Through the creation of a knowledge database, which allows knowledge sharing among XIFI support, XIFI support aims at minimising the dispatching of tickets to second level support. The main objective should be to solve most opened tickets already at first level support.

Requirements

- The end user has a contract with XIFI and has access to the XIFI helpdesk tool through the federation portal;
- The solution proposed should preserve the SLA contracted with the end user.

Deployment

The helpdesk tool provided and operated by the XIFI federation office will be based on existing popular ticketing, and bug tracing systems. At least three systems are under consideration, whereas the final deployment will be a compromise between wide acceptance, familiarity, complexity of deployment and consultation at the FI-PPP architecture board. The systems being evaluated are:

- JIRA by Altassian: used in FI-WARE for issue tracking among GE developers;
- Bugzilla: used by the Mozilla Foundation for issue tracking of Mozilla projects. The project members have experience with Bugzilla;
- Trac: a further popular wiki and issue tracking system for software development projects with which the project members have experience.

The favorite solution so far is the adoption of JIRA on the same platform as FI-WARE to have single entry point for users and PPP members. This is particular important for handover of tickets from XIFI to FI-WARE/TF.

4.2 Roles Description

Based on the Use Case diagram, the following roles have been identified:

End user

The end user is the Future Internet developer requesting support upon an issue related to the integration/operation of a Future Internet application within the XIFI federation. For the sake of clarification, the end user is not the final customer of the service or application being deployed by

the future Internet developer. The end user is the Future Internet developer himself. He is the end user of the XIFI federation and XIFI infrastructure services.

XIFI Support

It represents the group of experts who directly receive the support requests coming from Future Internet developers. The main activities are:

- Communicate with end user;
- Analyse issues and find the root causes;
- Identify the right expert to forward the ticket to;
- Validate the solutions proposed.

XIFI Expert

It is any individual appointed by a XIFI partner, who has particular skills or competences on one (or several) of the XIFI federation components. A XIFI expert is engaged only with issues related to the components developed in and offered by the XIFI project. The main activities are:

- Analyse issues related with XIFI component;
- Propose bug fixes for XIFI components.

FI-WARE Expert

It is any individual appointed by a FI-WARE partner providing support for one or more of the GEs that are included in the XIFI federation nodes. A FI-WARE expert is engaged only with issues related to FI-WARE GEs. The main activities are:

- Analyse the exchange between the application/SE and the generic enabler;
- Troubleshoot issues with the GE;
- Fix issues requiring a modification (bug fix) of the GE.

Infrastructure Expert

It is an individual appointed by an infrastructure node operator providing support for the infrastructure in which the application, the GEs, and the SEs are hosted. Each infrastructure must provide its own group of experts. The main activities are:

- Troubleshooting issues related with application hosting;
- Fix connectivity related issues.

4.3 Flow Description

Based on the interaction diagram illustrated below, the following actions have been identified:

Submit Ticket

The end user identifies an issue related to the deployment of its application within the XIFI federation. He logs into the federation helpdesk and creates a new ticket describing the issue. Once the ticket form is completed, he submits it to the XIFI support

Analyse Problem

The XIFI support team will analyse the issue based on the description provided on the ticket. If the description is not clear enough, they will contact the end user in order to gather complementary information. The analysis should lead to identify the root cause of the issue. There are to possibilities:

- The identification of the solution whenever it is possible, either by

- Consultation of the knowledge database in order to find an already identified solution
- The troubleshooting performed on the environment.

- Request expert support in case the analysis leads to issues related to the infrastructure, the enablers, or the XIFI components, and which requires specific competences to solve the issue.

Ask Expert

If the issue can't be solved by XIFI support, the ticket will be redirected to the people in charge of the component identified as root cause of the issue.

Validate Solution

Whenever possible, the solution should be tested by the XIFI support before making it available to the end user. This action requires the possibility to recreate the end user's environment for the validation.

Provide Solution

The solution will be described on the ticket with the information required to apply it, and it will be forwarded to the end-user.

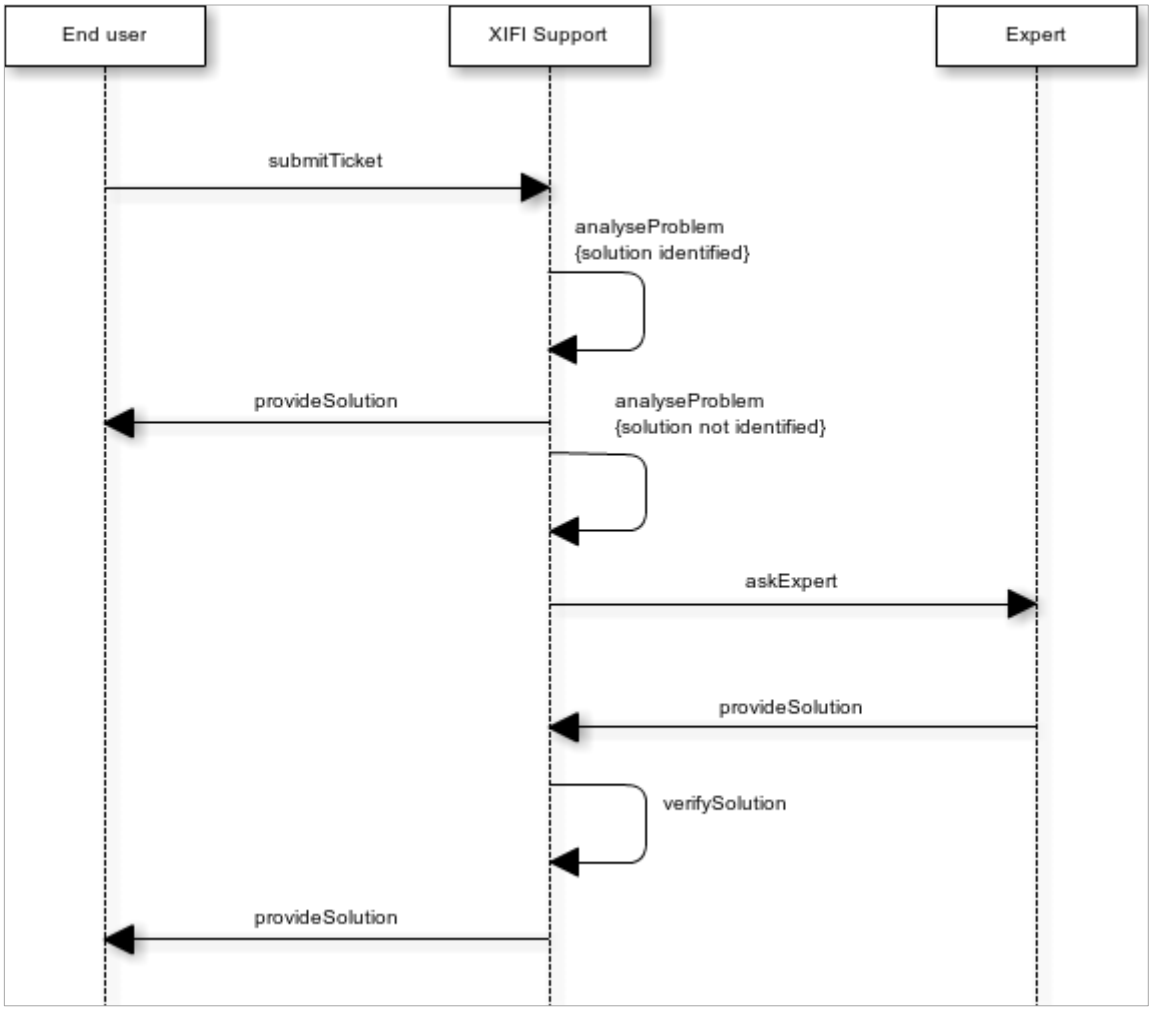

Figure 17: Developer support interaction diagram

In order to prevent spam on the deployed ticketing system future Internet developers will be required to register with the system. The XIFI support personnel will approve requests of future Internet developers and enable their account based on an e-mail to the support e-mail list.

Only developers with accounts for the ticketing system will be able to submit tickets to the ticketing system.

4.4 Identified Requirements

From the previous sections, the following requirements are identified, in order to provide support to end users as described:

Helpdesk

A helpdesk tool based on an established ticketing system must be available. This tool is provided for the FI-PPP by the **CONCORD project**. In order to be able to describe the ticket flows more precisely, it is required to have information about the implementation of the tool.

Knowledge Database

The helpdesk tool should provide the possibility to create a knowledge database.

Test Environment

In order to be able to validate certain types of issues a test environment is required to be available for the XIFI support.

5 MAINTENANCE

5.1 Maintenance Management

This section is addressing maintenance management of the XIFI federation. Please note that both terms 'maintenance management process' and 'maintenance process' are used in the following. The distinction is important since the latter is referring to a particular maintenance task usually implementing a particular sequence of actions to resolve an issue, while the former is referring to a management process, in particular dealing with the formalization of implementing the process and framework that includes a potentially large number of maintenance processes and how they are utilized and evaluated regarding their benefit in terms of an agreed set of quality of service metrics.

Maintenance management covers all topics that have to be taken into account for operating a reliable facility. Maintenance management is a cooperative process that consists of defining a maintenance strategy and of implementing this strategy. A maintenance objective must be defined when developing a suitable strategy. A maintenance strategy thus has a particular goal, defines a sequence of actions required to reach that goal, and names the stakeholders that are required to contribute. A valid maintenance strategy must consider risk (potentially based on a threat analysis) and cost factors (potentially based on a business case analysis) and it should consider metrics that allow judging the outcome when implemented to support a quality assurance process.

For the time being we define the main objective for XIFI maintenance as the optimization of the overall availability of the federation, which has two main aspects namely maximizing the up-time for all individual XIFI nodes and maximizing the redundancy across the XIFI federation by suitable fail-over strategies.

Maintenance processes have to be defined that determine which components and responsible persons must be involved. It is very important to highlight that the defined processes have to be followed by the people that are involved in the particular maintenance process. The maintenance management defines the following events that can occur:

- Scheduled maintenance

  Example: Regular Hardware or software updates; usually scheduled per node minimizing down-times collaboratively.

- Unscheduled maintenance

  Example: Replacement of defective equipment or of misbehaving software; includes restoring back-up states; unplanned or scheduled on short notice usually responding to urgent action requirements; may incur node or federation down-times.

- Incidents

  Example: Failure of major node, federation or communication infrastructure, potentially due to physical damage; usually involves significant down-times with barely predictable duration; requires preparation of incident handling processes including risk management strategies; may require subsequent unscheduled maintenance processes being initiated.

Maintenance processes and maintenance management processes involve a number of stakeholders which in the scope of the XIFI maintenance include

- Federation manager (federation maintenance supervisor).

  A federation manager is considered the first point of contact in case a maintenance process involves more than a single XIFI node. The particular process must define if the federation manager has to be informed about the particular process being initiated, has to be involved as a mediator among different nodes or with the first level support of the federation, or has to be actively involved to coordinate independent node actions to avoid federation down-times.

- Node manager (node maintenance supervisor).

  A node manager is responsible for the maintenance of a single node following both local maintenance processes and federation maintenance processes. In case a local maintenance process may affect operations of the federation the node manager also responsible for interacting with the federation manager. In case of an incident, node managers play an active role in federation management to minimize the federation-wide impact of a local incident.

- End users.

  Users may be involved in maintenance processes in various ways. Users may need to be informed about scheduled or unscheduled maintenance processes as soon as their use of the XIFI federation is affected. They may be involved to support the maintenance process by postponing activities, moving their activities across the federation to free up a particular node, or to backup and restore their results to bridge a certain foreseeable downtime of the federation, of a node or of a particular service. Users may also cause a maintenance process to initiate either through interaction with the first level support or through causing an incident.

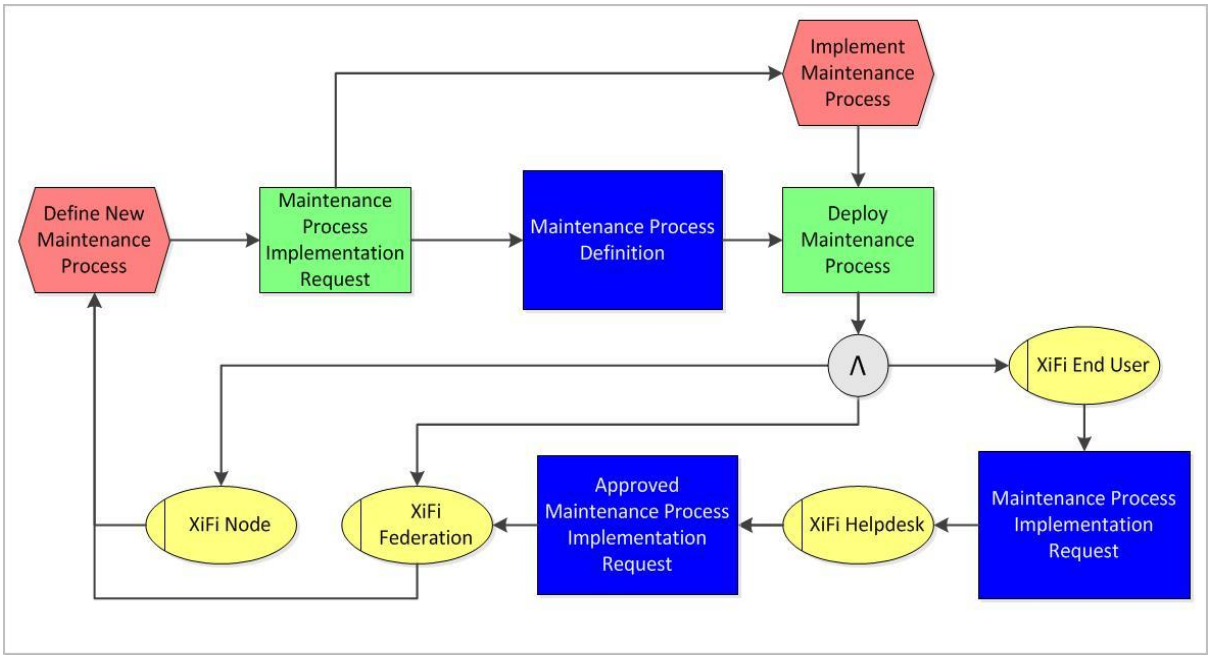

Figure 18: Maintenance management process for implementing a maintenance process

A maintenance process when defined must provide the same information as a use case description (in fact, a maintenance task it is a certain use case of the XIFI federation). In particular, the following information must be provided to describe the process:

- Pre- and post-conditions.
- Stakeholders involved and their responsibilities.
- A main sequence of actions to be taken.
- Error and Termination conditions.
- Expected outcomes and result metrics.
- Reporting requirements (for filing and effectiveness evaluation).

In contrast to a regular use case a detailed list of actions to be taken and responsibilities for each of those actions enables proper implementation. Special consideration is required in case the process is designed to handle an incident since multiple processes may apply to a single incident and their interaction impacts the effectiveness of handling an incident (i.e. minimizing down-time). This may

require a coordinating role (e.g. that of an incident response manager).

The complete set of maintenance processes when deployed implements the maintenance management provided that a framework exists that allows to apply metrics and processes to monitor and judge the effectiveness of maintenance management processes (i.e. enabling to determine the Quality of Service by a quality manager role).

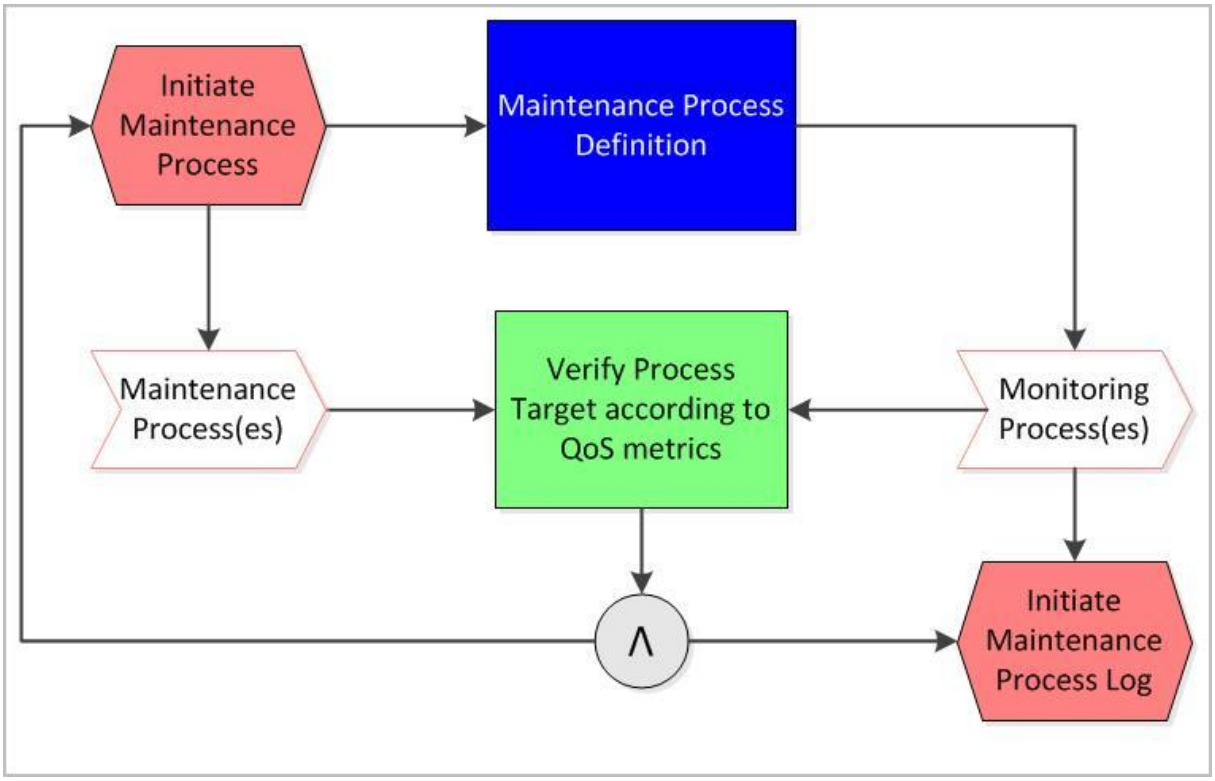

Figure 19: Maintenance management process for the maintenance process QoS evaluation

5.2 Maintenance of XIFI Node

A number of strategies will be set up covering the main areas of the XIFI federation that require maintenance management. At the time of writing a focus is set on providing process descriptions that address scheduled maintenance, leaving incident handling and unscheduled maintenance processes for later stages in the federation provision.

5.2.1 Physical Infrastructure

This includes XIFI node hardware and software maintenance processes such as regular updates and preventive measures to minimize the risk of hardware failures that may cause node down-times. Additionally, federated communications infrastructure connecting XIFI nodes must be included in suitable maintenance management processes even if not under the control of the federation manager. Although XIFI node maintenance in principle can be conducted without involving other XIFI federated resources, suitable processes need to be implemented and followed to coordinate with other federated resources and to minimize the impact of a particular node's down-time on the federation by scheduling maintenance periods into 'safe periods', for example. Dedicated processes must be defined to accommodate the following situations:

- Hardware and software upgrades for a XIFI node potentially including multiple reboots and rollbacks. The maintenance management process must consider coordinated periods and should consider allocating backups for services to bridge the down-time of a certain XIFI node by another node for longer down-times. These are mostly scheduled maintenance processes

but may be unscheduled in case a certain hardware/software problem must be solved urgently. Clearly, this does not cover federated services that may need to be ungraded in a federation-wide coordinated process (i.e. maintenance of federation tools and GEs).

- Communication system reconfiguration maintenance processes in general apply to the federation and may cause loss of connectivity (i.e. virtually node down-times) for one or more XIFI nodes. Since these may have a noticeable impact on the performance of the XIFI federation during the period of reconfiguration and shortly after, they have to be carefully planned, coordinated and conducted leaving sufficient room for node preparation and reconnection. A suitable maintenance process requires involvement of all XIFI nodes and of the federation manager since even not affected by the reconfiguration itself a particular XIFI node may need to support the process by actively reconnecting and connection testing. Communication system reconfiguration may be required for integrating a new node are new services of all nodes into the XIFI federation.
- XIFI node reconfiguration due to service changes or for resolving performance issues. These usually are unscheduled maintenance processes often going unnoticed because they don't cause noticeable down-times. They may be performed locally without involving other XIFI federated nodes but indication or support may be needed in case an updated configuration has to be verified (e.g. to re-establish permanent communication links between nodes). Such reconfigurations may require a dedicated maintenance management process only for the reason to log the event a scheduled down-time in the SLA logs and not as a XIFI node failure.
- Security software patches are unscheduled (high-priority) maintenance processes. They cannot be delayed, potentially causing a federation-wide down-time, but applying a particular patch in a row could be synchronized with other (similarly affected) XIFI nodes to minimize the joint down-time of a particular resource or service. Although the resulting maintenance management process may have its intricacies, it is of some relevance to implement a coordinated process due to the potentially high frequency of security related updates potentially on a daily schedule.

5.2.2 XIFI Federation Tools

In contrast to XIFI node maintenance as discussed above the maintenance of XIFI federation tools imposes a number of interoperability and timing requirements on the maintenance processes to ensure that the XIFI federation is not jeopardized because of (temporarily) inconsistent tools distributed across the XIFI federation. Thus, prior to defining a maintenance management process that deploys a new or updated version of a federation tool, for example, this particular tool must be evaluated regarding function, API and protocol dependencies, and regarding its backward compatibility. Additionally, the tool needs to provide rollback capability to recover from a failed update process if needed. In consequence, the maintenance processes that apply to federation tools always involve multiple stakeholders, must be collaborative and should be synchronized across the federation if involving concurrent sub-processes. When defining a suitable maintenance process, and prior to applying it to the operational XIFI federation, the utilisation of a sandbox (e.g. a particular virtual machine) or a laboratory test setup for testing the deployment should be considered. Dedicated processes must be defined to accommodate the following situations:

- Federation tool updates are either scheduled (regular updates) or unscheduled (fixes). Special care has to be taken if the deployment tool itself is subject of an update or fix. It is assumed that distinct processes are required to cover various configurations within the federation, if a single XIFI node, multiple nodes, the federation in whole, a XIFI master node, or 'unconventional services' are affected, for example. The preferred method is that a maintenance process definition has to be provided by the tool developer along with any new release of the tool.
- Federation tool initial deployment is a scheduled maintenance process that requires the

definition of one or more processes to satisfy the requirements of distinct situations such as distributing to one or more nodes via communication means (i.e. by download or subscription service) or via some off-line distribution media (i.e. by download from a server or via DVD) either introducing new functionality or replacing existing functionality provided by one or more former tools. It is recommendable to define distinct maintenance processes since pre-and post-conditions are rather different (e.g. regarding the need to inactivate/remove former tools), chaining of processes should be supported (e.g. to simplify a potential rollback to the previous deployment state), and concurrent processing may be needed (e.g. to de-register a former tool and to re-register an updated tool throughout the federation.

- Maintenance of the tool deployment tool is considered a dedicated maintenance process due to the tool's strict interoperability and rollback requirements. Any failure of the update or initial deployment of this tool would likely cause a XIFI node being disconnected from the federation and would require an incident handling process to re-integrate with the federation. Therefore, the processes defined (i.e. for initial deployment or update/fix, and for automated or manual deployment) should consider various failure conditions and should handle a failure avoiding any potential dis-functional states. In addition, it is of uttermost importance to achieve a consistent state regarding the functionality provided throughout the XIFI federation in case of a potential disruption of the maintenance process.

In general, federation tools are essential for keeping the XIFI federation in a safe operational state. Thus, some of the tools may provide redundancy to avoid a single point-of-failure while others may need to keep interoperability with alternative implementations for sustainability or scalability reasons. In consequence, related maintenance processes and maintenance management processes need to consider more than one target. That is, the process should jointly operate on a tool and its redundant/alternate companion such that maintenance objective, rollback and fail-over conditions, and quality-of-service metric should apply to both if applicable.

5.2.3 Service Host

Maintenance of the service host is a scheduled process similar to federation tool maintenance with respect to the central relevance of all processes being fail-safe as outlined above. To minimize down-times (in this case the down-time of the XIFI node affected, since FI services will fail completely due to unavailability of the service host) maintenance processes here need to involve the federation manager and other XIFI node managers that may be able to act as a fail-over for FI services hosted by the XIFI node affected. It is suggested to define a management process which handles all tools providing the service host function jointly (i.e. considering it as a single module) regardless if an update of all functions or a patch/fix for a dedicated function is performed. This will ensure consistency of the service host. It is important to note that the same argument holds for the node monitoring functions. That is, the node monitoring must be considered in conjunction with the service host functions for initial deployment, but may be considered by dedicated maintenance processes regarding updates or fixes. This exception is mainly due to the fact that the monitoring subsystem is comparably deep linking with the node platform (i.e. the physical infrastructure) but may need special consideration with regards to the sequence of processes affecting the service host and the monitoring sub-system (addressing the chicken-egg problem of the service host relying on the node monitoring and vice-verse).

As for any software maintenance process, regular updates of the FI-WARE cloud hosting GEs need to be scheduled in accordance with the general road-map issued by the FI-WARE project [42]. Since there is a strong dependency between FI-WARE cloud hosting GEs and several underlying software products such as OpenStack that have their own release cycles, any update must be approved by the maintainer of the FI-WARE cloud hosting GEs prior to scheduling update processes for the XIFI federation. It is of uttermost importance to avoid updating any software module utilized by the FI-WARE cloud hosting GEs concurrently since this may cause failure of a XIFI node to host any other GE in consequence. Hence, there will be no unscheduled incident procedures that apply in particular to

FI-WARE cloud hosting, except that FI-WARE will develop a process to handle solicited update request from the XIFI federation.

5.2.4 FI Services

Maintenance of FI services are scheduled processes (updates) or unscheduled processes (patches/fixes resolving urgent issues) having similar constraints than federation tool and service host maintenance processes regarding fail-safe operation. In particular, FI service maintenance is strongly linked to end-user activities and hence relies on involving end-users and the federation help-desk as stakeholders in maintenance management. It is suggested to involve end-users in developing requirements for related maintenance processes, in particular in defining maintenance objectives and metrics since these will be driven by experiments and use case projects mainly. In addition, maintenance process descriptions should be deployed along with the service by the service developer as discussed above for the case of deployment tools. In contrast to the above, dependencies exist not only from initial deployment time (requiring a maintenance process to consider multiple FI services jointly) but may also develop over time through their joint use by an experiment or application. Maintenance processes thus may need to consider fail-safe approaches for collections of FI services keeping track of interoperability constraints for all services of a collection. In practice this requirement may result in the need to delay an updates of a certain FI service until an update for related FI services are made available. On this level of dependency the maintenance management process needs to consider release management and risk management objectives as well as well as related business cases in conjunction with its focus maintenance objective.

6 CONCLUSIONS

The federation unites heterogeneous test infrastructures with different levels of compliance with XIFI and its requirements, thus a deployment support service is assigned to a joining node in the integration phase. So the work delivers a quick online test to help a candidate node define its level of compliance with the federation. In addition to the test, the work identifies operational and technical requirements of XIFI to a potential node as well the procedures to follow within each XIFI support level service to assist the node in integration into the federation.

Therefore having support services for joining infrastructures eases the integration process and helps to avoid technical and operation conflicts with the existing XIFI environment. It is evident, this is an essential part of the joining node process, where new nodes are asked to follow the operational and technical installation procedures defined by this work, which makes XIFI run uninterrupted.

After the integration process, stability of the federation is ensured by setting up the XIFI tools and FI services maintenance protocols developed by this deliverable, which make nodes within the XIFI federation operate correctly and allow all the stakeholders as well as just joined infrastructures to benefit from this capacity building and infrastructures project.

Furthermore, the work contains procedures to support XIFI end-users such as FI-WARE GE developers as well UC experimenters and developers as in accessing nodes of the federation.

Above all, in order to assure a smooth integration process and after effective day by day operations both the XIFI federation and the joining node must follow the procedures and protocols defined in this work for installation, support and maintenance phases.

As a result, this deliverable defines operational and technical procedures and protocols for a new infrastructure to implement and to follow in order to become a running XIFI federation node. However, it is a base for the XIFI nodes operation, assistance, and maintenance as well as XIFI federation extension support deliverables within the work package, allowing the XIFI federation to monitor and to improve all the stakeholders' experience within the operational XIFI environment.

REFERENCES

[1] *The Future Internet Public-Private Partnership Programme*, http://www.fi-ppp.eu

[2] *FI-WARE*, http://www.fi-ware.eu

[3] *The NIST Definition of Cloud Computing*. (National Institute of Science and Technology, September 2011). Retrieved 08:59, August 9, 2013, from http://csrc.nist.gov/publications/nistpubs/800-145/SP800-145.pdf, Section 2

[4] Technical support. (2013, August 7). In Wikipedia, the Free Encyclopaedia. Retrieved 08:59, August 9, 2013, from http://en.wikipedia.org/w/index.php?title=Technical_support&oldid=567491903#Tier.2FLevel_1_.28T1.2FL1.29.

[5] *Future Internet Lab*, http://lab.fi-ware.eu

[6] *eXperimental Infrastructures for Public private partnership in Innovation*, http://www.xipi.eu

[7] Pan-*European Laboratory Infrastructure Implementation*, http://www.panlab.net

[8] Federated E-infrastructure Dedicated to European Researchers Innovating in Computing network Architectures, http://www.fp7-federica.eu

[9] *Distributed European Infrastructure for Supercomputing Applications*, http://www.deisa.eu

[10] *Federation framework For Future Internet Research and Experimentation*, http://www.fed4fire.eu

[11] *D1.1 - XIFI Core Concepts, Requirements and Architecture Draft* (XIFI Consortium, 2013), https://www.fi-xifi.eu/Publications/D1.1-XIFI-Core-Concepts-Requirements-and-Architecture-Draft1.0.pdf

[12] *D2.1 - XIFI Handbook v1* (XIFI Consortium, 2013), http://www.fi-xifi.eu/Publications/D2.1-XIFI-Handbook-v1v1.0.pdf

[13] XIFI Consortium, *XIFI Open Call*, http://www.fi-xifi.eu

[14] *IaaS Data Centre Resource Management - Installation and Administration Guide* (FI-WARE Consortium), https://forge.fi-ware.eu/plugins/mediawiki/wiki/fiware/index.php/IaaS_Data_Center_Resource_Management_-_Installation_and_Administration_Guide

[15] *OpenStack Grizzly*, http://www.openstack.org/software/grizzly

[16] XIFI Consortium, *D9.1 - Dissemination and promotion plan*, available from www.fi-xifi.eu

[17] XIFI Consortium, *D9.2 - XIFI office – description and establishment*, available from www.fi-xifi.eu

[18] *Open source software for building private and public clouds*, http://www.openstack.org

[19] *OpenStack Operations Guide*, http://docs.openstack.org/trunk/openstack-ops/content/

[20] *OpenStack Conceptual Architecture*, http://docs.openstack.org/grizzly/openstack-compute/admin/content/conceptual-architecture.html

[21] *OpenStack Nova*, https://wiki.openstack.org/wiki/Nova

[22] *OpenStack Neutron*, https://wiki.openstack.org/wiki/Neutron

[23] *Open Virtual Switch*, http://openvswitch.org

[24] *OpenStack Glance*, https://wiki.openstack.org/wiki/Glance

[25] *OpenStack Image Management*, http://docs.openstack.org/grizzly/openstack-compute/admin/content/ch_image_mgmt.html

[26] *OpenStack Keystone*, https://wiki.openstack.org/wiki/Keystone

[27] *OpenStack Cinder*, https://wiki.openstack.org/wiki/Cinder

[28] *OpenStack Swift*, https://wiki.openstack.org/wiki/Swift

[29] *OpenStack Horizon*, https://wiki.openstack.org/wiki/Horizon

[30] *Open Network Management Application Platform*, http://www.opennms.org

[31] *Infrastructure for Network Performance Monitoring*, http://www.perfsonar.net

[32] *OpenStack Compute and Image System Requirements*, http://docs.openstack.org/grizzly/openstack-compute/admin/content/compute-system-requirements.html

[33] *OpenStack Cloud Controller Design*, http://docs.openstack.org/grizzly/openstack-ops/content/cloud_controller_design.html

[34] *OpenStack Instance Storage Solutions*, http://docs.openstack.org/grizzly/openstack-ops/content/compute_nodes.html#instance_storage

[35] *OpenStack System Requirements*, http://docs.openstack.org/grizzly/openstack-object-storage/admin/content/object-storage-system-requirements.html

[36] *Pacemaker*, http://clusterlabs.org/wiki/Pacemaker

[37] *Galera*, http://www.codership.com/wiki/doku.php?id=galera_wiki

[38] *OpenStack High Availability Guide*, http://docs.openstack.org/high-availability-guide/content/index.html

[39] *OpenStack Block Storage Service Administration Guide*, http://docs.openstack.org/grizzly/openstack-block-strage/admin/content/index.html

[40] *OpenStack Networking Administration Guide*, http://docs.openstack.org/grizzly/openstack-network/admin/content/connectivity.html

[41] *OpenStack Operations Guide*, http://docs.openstack.org/trunk/openstack-ops/content/network_design.html

[42] *Releases and Sprints numbering, with mapping to calendar dates FI-WARE releases*, https://forge.fi-ware.eu/plugins/mediawiki/wiki/fiware/index.php/Releases_and_Sprints_numbering,_with_mapping_to_calendar_dates FI-WARE releases

APPENDIX A QUICK ONLINE TEST

A.1 Page / Tabbed Dialog 1

XIFI Federation

XIFI will establish a sustainable marketplace for trial infrastructures and Future Internet services. XIFI will achieve this vision by integrating and federating a multiplicity of heterogeneous environments – starting from the generic and specific enablers provided by the FI-WARE core platform and the FI-PPP use cases and early trials. Through this approach XIFI will demonstrate and validate the potential and capabilities of a unified market for Future Internet facilities overcoming a number of existing limitations to the current set of Future Internet experimental infrastructures available across Europe, such as fragmentation, interoperability and scalability.

XIFI will also extend its efforts to include the results of other Future Internet services and R&D work. Initially the federation of infrastructures will consist of a core backbone five nodes located in five different European countries enabled with the Technology Foundation services (from the FI-PPP project FI-WARE) to be ready before the start of FI-PPP phase 3 (at month 12 of XIFI). This initial set will be enlarged during the second year with new use cases and collaborating local and regional infrastructures. XIFI will provide significant added value to Future Internet service and application developers. Specifically XIFI will:

- Facilitate unified access to large-scale infrastructures by providing a single entry point for users;
- Provide access to generic enablers with assured QoS and reliability that go beyond best effort;
- Offer a federation service through which the infrastructures can offer their capabilities using new and existing business models;
- Enable infrastructures to be shared across different use cases XIFI will provide training, support and assistance including integration guidelines and the promotion of best practice between large-scale trials and infrastructure nodes.

These activities aim at facilitating the uptake and continued use of the FI-PPP results.

XIFI Node

A federation of infrastructures requires node operations to be defined according to some common and consistent procedures. Such procedures should be minimally invasive and demanding for new infrastructures joining, so as to minimize conflicts with existing operational procedures of single infrastructures. In order to ensure proper operation of XIFI nodes we need to define proper procedures and protocols to establish the federation. Procedures and protocols should take into consideration that several nodes contribute to the XIFI federation and that each node in the federation hosts federation tools. Nodes are deployed in a two tier deployment Master nodes and Slave nodes.

Node Types

The XIFI Federation of infrastructures approach is to have a two tier system, with additional services requires for a Master deployment:

- Master node type: DNS server, Backup services, Hypervisor, cloud controller, XIFI Federation software, Generic Enabler repository, minimum Network connection 10Gig.
- Slave node type: Hypervisor, cloud controller, XIFI Federation software, public IP, minimum Network connection 1Gig.

A.2 Page / Tabbed Dialog 2

Legally Compliant

This Standard Service Agreement (SSA) shows the text of the contract between the XIFI Federation and new member. Upon receiving these agreements, one copy should be signed and returned the XIFI consortium.

Article 1 – Definitions

XIFI Federation: Is the deployment and management of multiple external and internal cloud computing services to match business needs Member: A legal entity that has entered into the XIFI Federation. Client: A legal entity that want to join the XIFI Federation.

Article 2 – Memorandum of Understanding

The joining client will complete and conform to XIFI Memorandum of Understanding.

Article 3 – Acceptable Usage Policy

The joining client adheres to the Acceptable Usage Policy (AUP) that sets out the measures to provide a high quality, reliable service to Platform users, protect privacy and security of users, infrastructures and networks, encourage responsible usage of the Platform resources and comply with relevant legislation.

Article 4 – Intellectual Property Rights

XIFI Members will be required to draw up their own specific agreements. Any IPR created by participating organisations using XIFI is owned by the organisation(s) that created the Intellectual Property. All background intellectual property created and owned by project partners remains sole property of those partners.

Article 5 – Confidentiality

Confidentiality rules will be incorporated into the Memorandum of Understanding. The following principles will apply:

- Each client party undertakes, subject to any requirements imposed by law, that it will treat the other party's information marked “confidential” or which from its very nature is obviously confidential with the same degree of care as it employs with regard to its own confidential information of a similar type or nature.
- Neither client or member party will intentionally disclose the other's confidential information to third parties other than those of its employees, consultants and sub-contractors who need to have such information for the purposes of the delivering services, and shall ensure that such recipients shall be bound by the same confidentiality obligations as are set out in this clause.

A.3 Page / Tabbed Dialog 3

Operationally Compliant

To ensure that the joining XIFI node can and are operationally compliant with helpdesk and XIFI federated software. Client Helpdesk support

Service Level Agreement

A service level agreement (SLA) will be written up for the each client. It outlines the XIFI Federation procedures used to process requests and problems that are reported. It will contain the following information.

- Node Helpdesk Hours of Operation: Assistance will normally be available from AM to PM, Monday through Friday or 365 days a year. Plan when help desk will be not available: public holidays, administrative closings, maintenance windows, etc.

- Node Contact Information : To report problems or request services for node administrative support the following contact information:
    - Phone: 00 353 51 2****;
    - Email: helpdesk@xifinodesite.com;
    - Web: www.support.xifinodesite.com;
- Node Support Levels, Priorities and Response Times: The joining node must conform to XIFI federation support requirements where it outlines the criteria for setting priority and associated response / competition times.

Operational XIFI Federation Software

The joining XIFI node will host and have fully operationally XIFI Data Centre Resource Management (DCRM) instance and a functional federated monitoring and network adapters installed. An operational XIFI node should have the availably of Backup and Recovery services.

A.4 Page / Tabbed Dialog 4

Technically Compliant

This is to ensure that client seeking to join the XIFI federation meet the minimum technical resources.

Item	Master Network Specifications	Slave Network Specifications
Core count	200	200
Ram	2 Gigs per core	2 Gigs per core
Disk	20 GB storage space per core	20 GB storage space per core

Table 9: Hardware Resource Pool

Item	Master Network Specifications	Slave Network Specifications
Log aggregator	CollectD	CollectD
Cloud Controller	DCRM instance	DCRM instance
DNS	DNS Server and services	
Node infrastructure Helpdesk ticketing system	Required	Required
Hypervisor	KVM preferred OS (Ubuntu 12.04)	KVM preferred OS (Ubuntu 12.04)
XIFI adapters	Required	Required

Table 10: Software Resources

Item	Master Network Specifications	Slave Network Specifications
Connectivity	Ethernet 10 Gb/s	Ethernet 1 Gb/s
Public IP pool (IPv4)	60 IP v4 address	60 IP v4 address
Public IP pool (IPv6)	IPver6 enabled	IPver6 enabled

Table 11: Network Resources

Item	Master Network Specifications	Slave Network Specifications
Shared Disk	500 GBytes	500 GBytes
Backup size	2 TBytes	2 TBytes
Firewall	Transparent instance	Transparent instance
Develop XIFI adaptors for local client node instance.	(site specific)	(site specific)

Table 12: Miscellaneous Requirements

This form will allow interested XIFI clients provide contact details and a method to inform the XIFI federation of its request. This information is passed to the XIFI federation via email and given the appropriate action.

Organisation:

- Org-name: Company A;
- Org-type: Research;
- Org-address: Address line1;
- Org-address: Address line2;
- Org-address: Address line3;
- Org-address: Address line4;
- Org-fax: ++ country code (0) ** ********;
- Org-www-country-prefix: (e.g. .ie, .co.uk, .de, .fr, etc).

Node Type:

- Node-Type: XIFI: XIFI Master/ Slave Management Details:
- Man-person: Person A;
- Man-address: Address line1;
- Man-address: Address line2;
- Man-address: Address line3;
- Man-address: Address line4;
- Man-phone: ++ country code (0) ** ********;
- Man-Email: Person.A@companyA.com;

Support Details:

- Sup-role: Tech support Group;
- Sup-address: Address line1;
- Sup-address: Address line2;
- Sup-address: Address line3;
- Sup-address: Address line4;
- Sup-phone: ++ country code (0) ** ********;
- Sup-Email: support@companyA.com.

