FUTURE INTERNET TESTBEDS
EXPERIMENTATION BETWEEN
BRAZIL AND EUROPE

Grant Agreement No.: 288356 (FP7) and 590022/2011-3 (CNPq)

# FIBRE

## Future Internet testbeds/experimentation between BRazil and Europe

Instrument: ***Collaborative Project***
Thematic Priority: ***[ICT-2011.10.1 EU-Brazil] Research and Development cooperation, topic c) Future Internet – experimental facilities***

## D 4.1 Report on the Federation Requirement Analysis

Author: WP4
Revised by: WP3
Due date of the Deliverable: Month 12
Actual submission date: 30/09/2012
Start date of project: June 1st 2011 Duration: 34 months
version: v.1.0

| **Project co-funded by the European Commission in the 7th Framework Programme (2007-2013)** | | |
|---|---|---|
| **Dissemination Level** | | |
| **PU** | Public | ✓ |
| **PP** | Restricted to other programme participants (including the Commission Services) | |
| **RE** | Restricted to a group specified by the consortium (including the Commission Services) | |
| **CO** | Confidential, only for members of the consortium (including the Commission Services) | |

## Abstract

The purpose of this deliverable is to define the requirements for the federation of Brazilian (FIBRE-BR) and European (FIBRE-EU) Experimental Facilities. These requirements are based on different aspects of the architecture, security, resource allocation and policies.

# TABLE OF CONTENTS

# 1 Acronyms

| | |
|---|---|
| API | Application Programming Interface |
| CDS | Content Data Server |
| CN-DS | Communication Networks and/or Distributed System |
| CNPq | Brazil's Council for Scientific and Technological Development |
| CPqD | Telecommunications Research and Development Centre |
| EU | European Union |
| FI | Future Internet |
| FIBRE | Future Internet testbeds / experimentation between Brazil and Europe |
| FP7 | Seventh Framework Programme |
| GENI | Global Environment for Network Innovations |
| ICT | Information and Communication Technologies |
| IP | Internet Protocol |
| I&M | Instrumentation and Measurements |
| LDAP | Lightweight Directory Access Protocol |
| MS | Milestone |
| NITOS | Network Implementation Testbed using Open Source platforms |
| NOC | Network Operations Center |
| NTP | Network Time Protocol |
| ORBIT | Open-Access Research Testbed for Next-Generation Wireless Networks |
| OF | OpenFlow |
| OFELIA | OpenFlow in Europe: Linking Infrastructure and Applications |
| OGF | Open Grid Forum |
| OMF | cOntrol, Management and Measurement Framework |
| OML | ORBIT Measurement Library |
| ORCA | Open Resource Control Architecture |
| OS | Operation System |
| RFC | Request for Comments |
| RNP | National Research and Education Network |
| UFF | Federal Fluminense University |
| WP | Work Package |
| WP2 | Building and operating the Brazilian facility |

# 2 Scope

This document is the result of the task T4.1 of WP4 of the FIBRE project.

The key words "MUST", "MUST NOT", "REQUIRED", "SHALL", "SHALL NOT", "SHOULD", "SHOULD NOT", "RECOMMENDED", "MAY", and "OPTIONAL" in this document are to be interpreted as described in RFC 2119 [Bradner 1997].

# 3 Reference Documents

[Altmann, 2008] Altmann, J., Courcoubetis, C., Stamoulis, G. D., Dramitinos, D., Rayna, T., Risch M., and Bannink. C., GridEcon: A Market Place for Computing Resources. In 5th International Workshop, GECON 2008, 2008.

[Antoniadis, 2004] Antoniadis, P., Courcoubetis, C., and Mason, R., Comparing Economic Incentives in Peer- to-Peer Networks. Special Issue on Network Economics, Computer Networks, Elsevier, 45(1):133-146, 2004.

[Antoniadis, 2010] Antoniadis, P., Fdida S., Misra V. - Federation of virtualized infrastructures: sharing the value of diversity, Conext 2010, Philadelphia, December 2010.

[Bradner 1997] S. Bradner, "Key words for use in RFCs to Indicate Requirement Levels", RFC 2119, March 1997.

[Cerf, 2009] Cerf V., Davie B., Greenberg A., Landau S., Sincoskie D. - FIND Observer Panel Report. April 9, 2009.

[Clark, 2002] Clark D.D, Wroclawski J., Sollins K.R., and Braden R. Tussle in cyberspace: defining tomorrow's internet. Proc. ACM Sigcomm 2002, pp. 347-356.

[Crowcroft, 2008] Crowcroft J. Toward a Network Architecture that does Everything, Communications of the ACM, vol. 51, no. 1, pp. 74-77, January 2008.

[Economides, 2006] Economides, N., The Economics of the Internet Backbone. In S. Majumdar, I. Vogel-sang, and M. Cave, editors, Handbook of Telecommunications Economics. North Holland: Amsterdam, 2006.

[Feamster, 2006] Feamster N., Gao L., and Rexford J., "How to lease the Internet in your spare time," ACM SIGCOMM Computer Communications Review, p. 61-64, January 2007. A longer version appears as Georgia Tech Technical Report GT-CSS-06-10, August 2006.

[Feigenbaum, 2002] Feigenbaum, J., Papadimitriou, C., Sami R., and Shenker, S., A BGP-based Mechanism for Lowest-Cost Routing. In ACM Symposium on Principles of Distriubted Computing, (PODC), 2002.

[Feldman, 2005] Feldman, M. and Chuang, J., Overcoming Free-Riding Behaviour in Peer-to-Peer Systems. ACM Sigcom Exchanges, 5(4):11-20, 2005.

[Fisher, 2007] Fisher D., US National Science Foundation and the Future Internet Design. ACM SIGCOMM Computer Communication Review, Volume 37, Number 3, July 2007.

[Gavras, 2007] Gavras A., Karila A., Fdida S. - Future Internet Research and Experimentation: The FIRE Initiative. ACM SIGCOMM Computer Communication, Review, Volume 37, Number 3, July 2007.

[Handley, 2006] Handley M., Why The Internet Only Just Work - BT Technology Journal, No 3 : Vol. Vol 24, July 2006.

[He, 2006] He, L. and Walrand, J., Pricing and Revenue Sharing Strategies for Internet Service Providers, IEEE Journal on Selected Areas in Communications, May 2006.

[Jacobson, 2009] Jacobson V., Smetters D., Thornton J., Plass M., Briggs N., Braynard R., - Networking Named Content, Conext 2009, Rome, December 2009.

[Karlsson, 2006] Karlsson G., Lenders V. and May M. Delay-tolerant broadcasting [Conference] Applications, Technologies, Architectures, and Protocols for Computer Communication: Proceedings of the 2006 SIGCOMM workshop on Challenged networks. - 2006. - Vol. 11. - pp. 197-204

[Keller, 2008] Keller, A., Hossmann, T., May, M., Bouabene, G., Jelger, C., and Tschudin, C., A System Architecture for Evolving Protocol Stacks, In Proceedings of 17th International Conference on Computer Communications and Networks, 2008 (ICCCN '08).

[Kelly, 1998] Kelly, F., Maulloo, A., and Tan, D., “Rate control for communication networks: Shadow price proportional fairness and stability,” *J. Oper. Res. Soc.*, vol. 49, pp. 237–252, 1998.

[Odlyzko , 2009] Odlyzko A. Network neutrality, search neutrality, and the never-ending conflict between efficiency and fairness in markets, Review of Network Economics, vol. 8, no. 1, March 2009, pp. 40-60.

[Shakkottai , 2006] Shakkottai, S. and Srikant, R., Economics of Network Pricing with Multiple ISPs. IEEE/ACM Transactions on Networking (TON), 14(6):1233-1245, 2006.

[Spyropoulos , 2007] Spyropoulos T., Fdida S., Kirkpatrick S., “Future Internet: fundamentals and measurement”, ACM SIGCOMM Computer Communication Review, Volume 37 Issue 2, March 2007.

# 4 Introduction

In general, a particular communication is a path between the endpoints that passes through intermediate technologies; occasionally with alternate paths available. One of the primary sources of evolution is the competition between paths for traffic. Metrics to be developed here include per hop and per path values for carried traffic; larger values create targets for competition. Another hop metric is the “work” or “difficulty” of the hop. We introduce Polymorphism, Virtualization and Federation as observations and concepts of the recent development of the Internet, we relate them to an expanded view of network architecture and finally, to the context of network testbeds.

Polymorphism refers to the different technologies and/or administrative domains that characterize the alternative paths or parts of the same path. It encompasses the various forms of diversity that could make the graph overwhelmingly complex. In order of decreasing importance:

- Multiple ownership, and hence policies,
- Multiple technical alternatives, for example wireless technologies,
- Simple growth in size.

Virtualization enables the parallel deployment of different types of networks on top of the same infrastructure transforming the traditional debate between the « clean-state » and « evolutionary » approaches for the future Internet architecture to a question of interconnection or possibly different architectures. In other words, the future Internet will be inevitably polymorphic and the key functionality that will ensure global connectivity and competition is federation.

Federation occurs when two or more organizations each devote some of their resources to a shared purpose, for example: interdomain routing, email delivery, DNS service, or mutual completion of telephony services. Federation is meaningful at all layers of the network architecture: from the lowest infrastructure layer (e.g., scientific grids and experimental facilities) to the highest application layer (e.g., openID). Each federation technique has two components: a set of **standards**, which provides connectivity, and a set of **policies**, which states what levels of connectivity are allowed (e.g., the amount of resources shared between the federated parties). This means that federation is both a technical and an economical problem.

There are different types of complexity that might arise in a polymorphic network without federation. For example, users need to keep up with different devices and/or interfaces in order to access different sources of information or communicate with a different set of people. They also face various types of lock-in situations and dependence on service providers. In the worst case certain parts of the network could be completely isolated from each other reducing the potential value of the network at large.

Intuitively, polymorphism increases the complexity and federation reduces it. However, federation itself is a costly process that increases the complexity for providers. And interestingly, virtualization becomes today an enabler for the creation of parallel, isolated networks built on top of the same infrastructure.

However, critical parts of the network (users, geographically dependent infrastructure, e.g., wireless access networks or sensors, specialized technology, etc.) will in general belong to different organizations. This is at least one of the main objectives for regulators and governments who need to encourage competition at local and global levels. This means that federation is necessary at one level or another and its main principles will shape significantly the evolution of the future Internet.

So what are the next steps in refining these ideas? The first is to look at contemporary federation systems and extract common design principles and best practices; the second is to use these to move toward a predictive model that includes non−technical factors, such as the desire of enterprises to create “sustainable competitive advantage” or “proprietary systems”, sometimes called monopolies or near monopolies.

This process will provide a transition methodology, which is one of the most critical issues for a successful evolution. In the following, we clarify the “Federation” concept that extends the peering solution currently in operation to provide the means for a global sharing of resources available in different domains and at different layers of the future Internet architecture. We analyze its different instances and explore the main principles that should guide federation at different layers of the network architecture.

The case of federation of network testbeds has been discussed since long and preliminary studies have been conducted (Antoniadis, 2010). We will illustrate how federation can be enabled in such environments that provide an appropriate playground to explore such principles.

# 5 Federation Rationale and Principles

We expose in this section the rationale, state of the art and basic principles for federation, supported by virtualization and required by polymorphism. Then, we provide a few examples to illustrate how the basic concepts were used in various contexts.

## 5.1 State of the art and challenges

There has been a tremendous number of papers to provide an analysis of the major flaws of the current Internet as well as its main grand challenges for the years to come (Handley, 2006), (Crowcroft, 2008), (Spyropoulos, 2007), (Gavras, 2007), (Clark, 2002), (Fisher, 2007), (Cerf, 2009). The main outcome is that opposite forces shape the current internet with requirements hard to satisfy all at the same time with the current architecture (trust, security, mobility, data-centric, …).

One main observation is that diversity will increase further. The system is going to evolve from a network of computers to Networked Systems that will take various forms. Virtual worlds are emerging and links with the real world will be established with sensors, tags and various “things”. Those emerging environments will be populated with a variety of devices and networks with an increased heterogeneity. Most of those will be mobile and dynamic both in time and space. Functionalities will continue to expand. Future networks will have to scale in size but also in functionalities to cover. In addition, they will have to support a broad spectrum of service semantics, ranging from best effort to high-quality, from open to highly trustable and secured. In order to run, they will use different set of resources provided by independent domains, requiring resource accountability. A consequence is that we might want to extend the waist of the reference model, introducing the concept of federation instead of a homogeneous abstraction (Keller et al, 2008). It is recognized that the service provided by the Internet is moving from connectivity to content and should be designed as an enabler for service creation and competition. An obvious consequence is that the cost incurred could be an increased management complexity. Besides technical evolutions, a user-centric approach to system design should be considered and other factors than technology will be instrumental, such as economics, social behaviours or regulations.

From an economic point of view one needs to understand under which assumptions federation is beneficial for the involved parties and the network as a whole and what types of federation agreements could help the system to reach the desirable equilibrium. To this end, there is a growing literature on Internet economics studying independently resource allocation issues in the Internet (Kelly 1998), ISP interconnection agreements (Feigenbaum, 2002), (Economides, 2006), (Shakkottai, 2006), (Antoniadis, 2010), (He, 2006), Grid economics (Altmann, 2008), incentive mechanisms for p2p systems (Feldman, 2005), (Antoniadis, 2004) and more. Federation economics will have to address the heterogeneity and polymorphism of the network domains involved and their complex interactions. How to compare and value multi-dimensional resources, to what extent the future Internet economy should be regulated or be designed as a free market in order to achieve a globally efficient allocation and provision of resources are only some of the questions that need to be answered.

## 5.2 Polymorphism and virtualization

Looking at the future, one should question if a single protocol will fit all needs in a world where diversity is fundamental. In order to illustrate the problem, we provide three parallel evolutions at different aspects/layers of the Internet architecture that demonstrate its polymorphism and will discuss how they could be integrated.

### 5.2.1 The Network is Wireless (and somehow optical)

At the physical layer, it is obvious to observe the fast development of wireless technologies and networks. Originally, they were designed to support mobility in basic services, like in GSM, and then they were also considered to ease network deployment and management (Wifi Hot-spot, Mesh networking, Ad hoc networks, Hybrid networks…). Wireless covers a broad range of systems, going from tiny devices with energy constraints to cellular-based systems covering a metropolitan area. The frequency resource is, to some extent, limited by regulation but can be better-utilized using solutions based on cognitive radio. New components have benefited from technological developments towards small device integration, providing opportunities to sense our environment or tag objects around us. Again, it is not realistic to hope that a single technology will support all needs as the basic requirements are so different, ranging from high-capacity networks to low energy consumption ones. Specific services are often attached to the Wireless world such as handover management or geolocalization.

The interconnection of all these diverse network technologies will need to cope with a variety of technical and socio-economic issues, for example, different performance characteristics, different routing schemes, and different ownership models. We will elaborate more about these issues in the following.

### 5.2.2 The Network is a Database

Although, part of the Internet success is based on its packet architecture, customers at large do not care about packets! Traffic engineering is now moving from flows to services and this is made feasible with the availability of DPI (Deep Packet Inspection) in communication devices as well as packet classification and content-based switching or routing. Similarly, the current service model of the Internet was designed to address a packet to a given receiver with multicast extensions being provided a decade ago. However, the current preferred model is a "shared" one, "Data to Many", supported by the development of the Web2.0 and the large diversity of end systems involved. The same observation holds for Enterprise networks with the emergence of SOA (Service Oriented Applications), the utilization of ESBs (Enterprise Service Buses) and FESBs (Federated ESBs). Moving content involves storage and transmission. The right trade-off between these resources has always been an issue depending on their cost and availability, the heterogeneity and dynamicity of content and the communication service model. Various solutions have been developed over years such as Client-Server, Push, Podcasting, CDNs, PubSub. The Publish and Subscribe paradigm seems to have a strong advantage over other models and is particularly well adapted to various environments such as DTN (Fall, 2003) (Karlsson, 2006) or Enterprise networking (banking applications) and is a promising one for the evolution of the Internet if some interesting challenges can be overcome such as scaling this service model to the global internet. Alike, the newly introduced CCN model (Jacobson 2009) is of increasing interest. Under an evolutionary approach for the future Internet the question that arises is how to interconnect content-centric networks with traditional Internet domains. Which types of federation agreements will be meaningful in this case? It is easy to see that the current peering and transit agreements are not incentive compatible for cases when storage instead of bandwidth is the main resource that needs to be offered in a federation.

### 5.2.3 The Network is the people

From the perspective of applications, it was envisioned for years that connectivity would be on by default. Then, various constraints and emerging environments emphasized situations where infrastructure is lacking or damaged and connectivity is intermittent or too expensive. It was then questioned if one can take benefit from encounters, providing transient or ephemeral connectivity by inter-contact opportunities to serve communication services. Mobile environments also trigger specific design concerns as it was observed that mobility is in general very context sensitive. Wireless is now a reality and social networks too, meaning that there exist correlations between social groups and network connectivity that one can exploit to provide services. The DTN concept has emerged (Karlsson, 2006) with various targets, both in situations where infrastructure is abundant or non-existent. The specific nature of these environments triggers interesting research and operational questions related to the service that can be delivered and its own performance.

In addition, the issue of ownership comes into play. When networks or services are built through the collaboration of independent organizations or even individual users how to allocate resources internally and most importantly how to connect them with external systems become a very challenging questions that the concept of federation for the future Internet should address.

### 5.2.4 Virtualization: The Network is one and many

The above examples illustrate some different shapes and requirements that the network could take in the future. Each of these evolutions is addressing a given environment where the objects and constraints are quite particular. The objective to design an architecture with an Hourglass model will come at cost to accommodate the diversity of its numerous components.

In the situation where we have to face very heterogeneous environments, how can we envision to enable them to co-exist? It will be useful if, by definition, there exists a solution to run concurrently multiple networks over the same physical substrate. This is supported, to some extent, by network virtualization (Feamster, 2007). At the end, the network at large should be seen as a global shared resource, which is virtualized and made accessible at scale. Virtualization is therefore a strategic component to accommodate different instances of the network into a single framework.

Indeed, virtualized networked systems are emerging. Today, there is already a rationale for going to virtualized servers in Enterprise networks and to progressively connect virtualized resources. There is also a trend for these resources and components to be less persistent. This is the case with mobile and nomadic terminals and networks but also because of the objective to lower the entry cost for connectivity, allowing on-demand-networking.

Virtualization enables the possibility to run and operate multiple heterogeneous networks on the same physical substrate. It should ease the process of network creation supporting opportunities to continuously deploy and operate customized solutions when available and economically viable. Many challenges are faced to reach this objective in the future. On the one hand, network equipments should evolve in order to support virtualization at high-speed and at scale in number of slices. In addition, some resources are hard and costly to virtualize, such as wireless or sensor devices. On the other hand, the management tools for infrastructure virtualization have also to be designed and deployed, including traffic management, security and interconnection. Virtual resources will have to be properly represented in order to provide an efficient utilization even in the situation of very heterogeneous resources.

Note that above Virtualization is considered as “1 to n” relation (slicing); however it can be “n to 1” relation (aggregation).

Nevertheless, virtualization is certainly an important architecture component and an enabler for the future networks that has the potential to support our vision for the coexistence of multiple concurrent networks, a Polymorphic Internet, and also to provide new business opportunities.

To take full advantage of the benefits of virtualization, an opposite force is required to ensure global connectivity and competition: The glue to manage and secure the Polymorphic Internet will be provided by the **Federation principle**. Federation could be horizontal or vertical involving different levels of cooperation between independent organizations as we analyze in the following.

## 5.3 Federation and complexity

The possible coexistence of multiple concurrent networks creates a new level of complexity that until now the hourglass model hided pretty well. The reason is that one needs to devise ways to interconnect these different networks. There are numerous examples in the past and today that demonstrate that the lack of federation has been responsible for increased complexity for users, although more profits for the biggest players. Social networks like facebook and twitter are perhaps the most interesting case of polymorphism at the application layer for which the lack of federation reduces the maximum value that users could get out of these systems.

Of course, the requirement of federation becomes even more challenging when isolation might be desirable for security or profit reasons. In this case, virtualization creates an opposite force that facilitates the parallel deployment of networks at different layers (e.g., IPv4 and IPv6 or content centric networks or different social networks like facebook and twitter).

The central question that we wish to address is related to the ability to interoperate virtualized infrastructures supporting heterogeneous protocols and services. The goal is to achieve increased coverage at different layers and/or enable resource sharing between independent domains. Federation is more than interconnection. It covers API, policies, governance, trust and economics. Of course, interoperability should be achieved at different levels, such as naming, service discovery or resource management.

Federation will govern the interoperability of independent networks, each managed by a given authority. Alike the domain concept in the current internet, a similar environment will be defined as it is unlikely that a single entity will deploy alone the concurrent networking environments mentioned above. A domain is considered as an independent set of resources providing services managed by a trusted administrative authority. Therefore, a domain has a value by itself but will also often benefit for being associated with other domains in order to achieve scale or heterogeneity. Users of a domain have access to its resources but would, in certain circumstances, benefit from accessing services offered by other federated domains. The governance of the global shared resources provided by the federation of domains is therefore distributed. It requires local policies to control local resource access but also external policies to grant access to external users. Different federation architectures can be considered, ranging from bilateral agreements to more scalable peering models. Key issues are to enforce a federation model that supports incentives mechanisms for sharing and rewarding policies that favour access to resources and services.

ISP interconnection is a typical example of a federation implemented through a set of bilateral or multilateral agreements. In this case federation is necessary for any domain that wishes to have access to the whole Internet population. Indeed, the global Internet has grown progressively through the interconnection agreements of up to 20000 autonomous systems. These agreements characterize the business relationships between the authorities governing each autonomous system and define the rules for forwarding traffic between them. Federation is more than just peering. The Internet allows anyone to use resources made available to access a variety of services. Alike the SOA (Service Oriented Architecture) or the Cloud approach, these resources are not limited to network bandwidth but include storage and computing. This diversity and heterogeneity requires a more precise accountability of resources, and the ability to compute the value of every Autonomous System according to the service that it delivers. So, a federation of ASes in the polymorphic Internet will also share common features with GRIDS and Clouds, often seen as computational markets run by resources brokers. The idea is that instead of having to agree on specific resource provision/allocation rules between facilities, one could imagine the creation of a global market. Organizations gain credit from providing their resources, which can be then used by their users to buy resources from other organizations. Which is the most suitable approach to guide federation policies? Rule-based federation policies or free markets?

Note that federation is in principle more than standardization and interoperability. Two domains will federate not only by using the same interface to talk with their customers but also by sharing unused resources. In the case of Grids the objective is to take advantage of the bursty demand and take advantage of unused resources. In the case of ISPs, the objective of peering is to reduce the transit costs and improve the quality of service.

A common definition for Federation in a political context is informative: A federation is a union comprising a number of partially self-governing regions united by a central ("federal") government under a common set of objectives. Federation will then comprise 1) credentials between the different regions or domains, which is the term used often in the networking world to describe an independent organization, and 2) policies to enforce their governance Federation is then enforced by the recognition that the federated domains share common objectives, otherwise being meaningless.

In the following we present some examples of federation at different layers and different types of polymorphism (technological and/or administrative). We then make a first step towards an abstraction of federation that could help us understand its generic properties and dependencies.

### 5.3.1 Standardization/Addressing

The fundamental requirement of federation is a unified characterization/addressing of resources and some sort of standardized interface for communication and/or service deployment.

At the extreme case, federation schemes need not involve data transfer. Today's 13 digit optical barcodes of today evolved from the 12 digit US barcode system, along with the desire by the non−US world to be an independent numbering authority. We might theorize that providing mechanism for merging with another numbering system without requiring authorization from the other system would be a desirable characteristic. The RFID numbering system in use by EPC can be compatible with EPC's optical bar code system, but the aircraft industry RFID system is numbered independently from the EPC's, in part to avoid paying for EPC numbers. German RFID emphasizes the need to coordinate number spaces, but also to be able to operate without an available root, or even if the root is hostile. The lesson here is that federation should allow easy concatenation (without central control) but also provide local autonomy, even if secession is never really anticipated.

### 5.3.2 Resource sharing: IP – CCN federation

At the network layer in addition to agreeing to a common addressing and communication API, independent providers need to also share resources. In the case of IP the main resource shared is bandwidth. This additional requirement for federation has lead to a wide variety of federation/interconnection agreements between ISPs today ranging from the multilateral peering agreement of Tier-1 ISPs, the more complicated (paid or not) peering agreements between smaller ISPs, transit agreements, etc.

How such agreements will evolve when the different networks run different routing protocols? For example how an interconnection architecture between IP and CCN networks would look like? Besides the technical aspects, a major challenge in this scenario comes from the fact that in the case of IP interconnection the main resource shared is bandwidth while in a CCN-based network the main resource shared is storage.

### 5.3.3 Infrastructure sharing: Cloud federation

Virtualization enables a much more generic infrastructure sharing including storage, processing power, etc. For example, a potential cloud federation could lead to a global infrastructure pool used on demand for a variety of services at different layers of the network architecture. Differences in model, ownership, and just a plain rapidly evolving technology create an obvious desire for standards and interoperability by users or regulators, but little benefit to proprietary cloud vendors. How to balance users needs and creation of economical value? This question is addressed in the following section 5.4.

## 5.4 A generic economic model of federation

Modelling the utility and costs of the different actors and their interactions will play a key role in studying the design and implementation of appropriate policies that will guide future federation agreements. The outcome of this analysis will also help us to determine the necessary elements that should be included in the polymorphic convergence layer to allow the various tussles to be resolved according to the corresponding socio-economic environment.

The main agents of this economy are:

- Customers (either end-users or institutions promoting research activities) who pay for different high-level services independently of how they are composed by the different contributions of the federated domains,
- Independent domains who pay a certain cost for their contribution to the aggregated resources and acquire a certain percentage of the total profit either directly or through internal payments based on the federation policies

The value of a federation agreement between two or more independent domains can have two different dimensions: 1) diversity and 2) capacity. Diversity refers to the value acquired from the fact that the interconnection of two domains increases the available choice of their users for services and communication. An extreme form of federation in terms of diversity is the notion of interoperability, which has been brought up for example in recent discussions concerning Cloud federation[1]. In this scenario two domains need to agree on certain interoperability standards in order for their customers to be able to access a wider set of services at the same time. No resource exchange takes place. On the other extreme, in the case of shared Grids for example, independent organizations federate in order to increase the

[1] [http://cloud-standards.org/]

available capacity of their applications to exploit the burstiness of demand. In this scenario, it is only the total amount of raw resources that counts and not where these resources reside.

ISP interconnection lies somewhere in the middle since federation's main objective is to achieve global connectivity with low cost (coverage dimension) but still federated domains might need to consume bandwidth to provide good service to each other (capacity dimension).

It is interesting that these two different types of value shared between domains have different semantics from an economic point of view. The value from diversity is non-rivalrous under low demand and has public good characteristics while capacity is a rivalrous good, subject to congestion effects and different levels of quality of service. Considering diversity as the main value generator can have significant implications on the way profits should be shared among federated domains supporting a single service. In this case, the assumption on the type of demand is critical (Antoniadis et al. 2010). This means that defining the resource unit to be exchanged and its valuation in a multiplicity of federations between highly diverse domains participating in different layers of the internet network architecture is a challenging task.

Then a federation policy should be agreed and enforced between the parties involved that will define the required contributions and the corresponding profits or access rights of each party. There are many alternatives: free markets with different payment schemes, reciprocity mechanisms, and public good approaches. Clearly, the final equilibrium of the system depends heavily on the form of the policies, and their implementation. Different policies should be evaluated based on their effectiveness in achieving the following objectives:

- to provide incentives for investments on infrastructure and resource sharing
- to protect local users and resources from congestion and other forms of abuse
- to increase the total value of the system and the corresponding profits
- to allocate resources and share profits according to a high-level objective.

There are various key modelling decisions that need to be made to this end. First, one needs to characterize the different types of resources provided, decide on their relative value and identify the corresponding costs. Then whether one needs to optimize for fairness or efficiency (Odlyzko, 2009) can significantly affect the theoretically optimal policies and the complexity for their implementation and enforcement. Finally, the actual demand, political aspects, the ability to set a global constitution instead of a series of bilateral or multilateral agreements, and various technical constraints (e.g., accounting) will play also significant role on how the system will evolve under different mechanisms.

Clearly, we are not in a position to provide final answers on how federation will evolve in the future Internet. The strategies of the big players, the way the net neutrality conflict will be resolved, the advances of the wireless technology are only some of the developments that will further affect the form of federation agreements and policies. Economic modelling can help us understand certain trade-offs and dominant system characteristics and build protocols and abstractions that will allow the system to adapt to the socio-economic forces.

Therefore, the federation of testbeds offer an appropriate playground to explore further these concepts in a practical environment that is under deployment in various contexts such as FIRE and GENI.

# 6 Federation Building Blocks and Requirements, Policy and Technology

The concept of federation is central to address the complexity of a polymorphic Internet and federation standards and policies should be defined at different layers of the network architecture in order to ensure interoperability and maximum connectivity of users and services, to encourage strong competition of infrastructure, network, and service providers, and maximize speed and reliability of communications with an efficient usage of the underlying infrastructure.

Similar to what was presented in the WP2 requirements for the FIBRE-BR system in terms of federation; we present below the federation requirements to the FIBRE system as a whole, considering the main building blocks of a successful federation at different layers.

- Architecture: resource provisioning, inter-communication, management and measurements.

**[FR01]** The FIBRE architecture MUST enable federation, which is the management and inter-operation of multiple independently administered resources, referred to here as islands, which are owned by multiple distinct organisations both participating in the FIBRE project and external to it, and possibly using different wired and wireless technologies.

**[FR02]** The FIBRE system SHALL provide a framework for the communication among federated organizations, which enables the inter-operation of the Brazilian and European islands, as well as with other external testbeds.

**[FR03]** The FIBRE system SHALL allow administrators to manage the federation. This involves activities such as including new islands, installing and updating system federation software, and monitoring the federation in terms of performance, functionality, and security.

**[FR04]** The FIBRE system SHALL provide a framework for gathering measurements and results from experiments run across multiple federated resources.

- Authentication: Different organizations need to be able to authenticate and trust each other's users

**[FR05]** The FIBRE system SHALL provide the means to enable organizations to carry out secure authentication of users from federated organizations, possibly providing different levels of clearance and access according to organization policies and the needs of the users.

- Addressing: Resources need to characterized and identified in a unified way

**[FR06]** The FIBRE system SHALL allow federated organizations to describe resources that they contribute to the federation in a common format.

**[FR07]** The adopted format SHOULD be flexible enough to permit the description of resources from different technologies.

- Security and resource allocation: Restrictions and access privileges need to be enforced for different users belonging to different organizations

**[FR08]** The FIBRE system SHALL provide mechanisms to allocate resources between experiments being run concurrently.

**[FR09]** The FIBRE system SHALL provide mechanisms to enforce resource sharing policies between experiments being run concurrently.

- API: Standardized communication and/or service deployment

**[FR10]** The FIBRE system SHALL offer access to the federated resources through open interfaces.

**[FR11]** The FIBRE system SHALL include a portal that allows experimenters to acquire, schedule, and release federated resources.

- Policies: resource allocation rules based on contribution

**[FR12]** The FIBRE system SHALL allow federated organizations to declare usage policies for substrate facilities under their control, and to provide mechanisms for enforcing those policies.

# 7 Conclusion

The evolution of the Internet is of utmost importance to our economy and our society just because it is playing a central and crucial role as the main enabler of our digital era. It is widely recognized that, although the Internet is a major success and one of the largest system build by humans, it was designed with assumptions and a target environment that faces more difficulties to cope with the future objectives:  stronger security, more privacy, better mobility support, enhanced reliability, better manageability and robust service guarantees are only examples of areas where innovation is needed.

We foresee various concurrent networks being deployed and customized to provide their specific service. An increased diversity and functionality of the networks and their components questions the architecture to support their interoperability and continuous deployment. We claim that the future Internet will therefore be polymorphic to enable several networking environments, each with their own features and strengths, to be deployed and coexist on a permanent basis. Network architecture should have its own laws and principles. We expect that virtualization and federation are the pillars of such architecture. We provide an analysis of these concepts when applied to existing systems and discuss how they should be refined in the future.

In addition, we consider the testbeds framework as a playground to explore the practicality of federation. The FIBRE project offers the opportunity to bridge two different testbed communities to explore adequate federation architecture and its relevant components.

"*This work makes use of results produced by the FIBRE project, co-funded by the Brazilian Council* for Scientific and Technological Development (CNPq) and by the *European Commission within its Seventh Framework Programme."*

END OF DOCUMENT