Design of a testbed for R&D in network architectures

Antônio Abelém[1], Iara Machado[2], Michael Stanton[2,3] e Tereza Cristina M.B. Carvalho[4]

[1]Universidade Federal do Pará (UFPA)

[2]Rede Nacional de Ensino e Pesquisa (RNP)

[3]On secondment from Universidade Federal Fluminense (UFF)

[4]Universidade de São Paulo (USP)

abelem@ufpa.br, {iara, michael}@rnp.br, carvalho@larc.usp.br

Abstract. *This paper discusses the current state of research into large-scale network architectures to substitute the current Internet architecture. Such a "clean slate" approach to "Future Internet" architecture is being adopted in a number of important initiatives in several countries. In the case of this INCT project, such an approach is important as a way to develop Web Infrastructure for the future. The paper also assesses several alternatives for possible adoption in the INCT Future Internet Architectures subproject.*

1. Introduction - relation to Webscience

The description of the scientific programme of the INCT Web Science states that 'the "Web Infrastructure" layer will deal with the Web as a technological means to ensure scientific, technological and societal progress, dealing primarily on the question of how to scale to meet performance or reliability expectations. It will thus advance results on areas such as computer networks, integrity and dependable computing, ensuring security in data transfer and communications, and distributed and parallel execution of the hundreds of thousands of processes needed by novel applications – therefore contributing to all other layers.'

The objective of the "Future Internet Architectures" subproject is to participate in and to contribute to the growing global community which is developing research activities in "Future Internet" (FI), a term used to characterise a number of recent initiatives to develop technologies used to support the Internet, and therefore the Web. FI is characterised by a "clean-slate" or "disruptive" approach, in which backward compatibility with current IP technology is not considered essential. FI initiatives are currently receiving government and other financial support in a number of countries (EUA, UE, Japão, Coréia), and the current research initiatives are increasing in geographical scope, with the onset of research collaborations between initatives in different countries. Currently this is a field which is still little explored in Brazil, although we shall mention below some FI initiatives already underway.

This paper is organized as follows. In Section 2, we describe motivation for experimental research into Future Internet. In Section 3, we provide a review of Future Internet testbed Initiatives. In Section 4, we present some Brazilian testbed initiatives.

In Section 5 we discuss the design choice for our national testbed. Finally, Section 6 concludes the paper and lists intended future work.

2. Motivation for experimental research into FI

In the last decades networks have become part of the critical infrastructure of businesses, homes and schools. In this context, the work of networking researchers became more relevant, but their chance of making an impact is more remote. The reduction in real-world impact of any given network innovation is because the enormous installed base of equipment and protocols, and the reluctance to experiment with production traffic, which have created an exceedingly high barrier to entry for new ideas. Today, in the main countries of the world, there is almost no practical way to experiment with new network protocols in sufficiently realistic settings to gain the confidence needed for their widespread deployment. The result is that most new ideas from the networking research community go untried and untested; hence the commonly held belief that the network infrastructure has “ossified”.

Having recognized the problem, the networking community is investing in alternative solutions, like programmable networks, such as GENI [GENI 2009] a proposed nationwide research facility for experimenting with new network architectures and distributed systems. Virtualized programmable networks could lower the barrier to entry for new ideas, increasing the rate of innovation in the network infrastructure.

Virtualization of computers has long been used, and is today widely available on common platforms, such as the Intel architecture. In principle, such an approach permits the logical replication of the original computer architecture, permitting multiple instances of the this architecture, each of which can run its own operating system and applications. Here, virtualization is accomplished by the sharing of processors and I/O devices using time slicing, and by virtual memory techniques. Virtualization of networks is more recent, and is accomplished by the use of virtual routers and the multiplexing of links between them. A number of current techniques for network virtualization are discussed in [Choudhury 2010].

These programmable networks call for programmable switches and routers that using virtualization can process packets for multiple isolated experimental networks simultaneously. For example, in GENI it is envisaged that a researcher will be allocated a slice of resources across the whole network, consisting of a portion of network links, packet processing elements (e.g. routers) and end-hosts; researchers program their slices to behave as they wish. A slice could extend across the backbone, into access networks, into college campuses, industrial research labs, and include wiring closets, wireless networks, and sensor networks.

Beyond the issues mentioned above, the study of alternative architectures for the future Internet is important for others several reasons, that this is a truly effective way to resolve many of the outstanding problems which currently afflict Web use. Some of the disadvantages of continued persistence in the use of the current architecture include:

- imminent exhaustion of the currently available space of IPv4 endpoint identifiers, causing a "balkanization" of the Internet, without widely available true global connectivity;

- increased costs of IP routers, due to the non-scalable nature of the internal routing tables, and of the performance requirements to process IP packet headers at line speed on very high-speed links, thus restraining network growth;
- immense investments in palliative measures to counter such security problems as are currently caused by spam, denial of service and outright information crimes;.
- difficulty of combining access transparency and application performance for mobile users.

The adoption of an alternative architecture can alter this situation. It is highly relevant to note that the pursuit of alternative solutions has already begun in parts of the developed world.

3. Review of FI testbed Initiatives

PlanetLab [Peterson 2002] [PlanetLab 2009] is a geographically distributed overlay platform designed to support the deployment and evaluation of new network services.

As a distributed system, PlanetLab is characterized by a unique set of relationships which make the design requirements for its operating system different from traditional hosting services or timesharing systems. The PlanetLab organization and the institutions that own and host PlanetLab nodes share control of PlanetLab nodes. The Researchers have access a distributed set of machines that must be shared in a way they will find useful through a PlanetLab “account”, called a “slice”, and implemented using a technique called distributed virtualization.

VINI [Bavier 2006] is related to the PlanetLab project. It is a virtual network infrastructure that allows network researchers to evaluate their protocols and services in the wide area. VINI allows researchers to deploy and evaluate their ideas with real routing software, traffic loads, and network events. To provide researchers flexibility in designing their experiments, VINI supports simultaneous experiments with arbitrary network topologies on a shared physical infrastructure.

GENI [GENI 2009] is a US National Science Foundation initiative to build a virtual laboratory for exploring future internets at scale. After 3 years of design studies, the construction of the GENI laboratory began in 2008, and the first virtual laboratories are being deployed. Researchers may download software into GENI-compatible nodes to control how those nodes behave. Whenever feasible, nodes implement virtual machines, which allow multiple researchers to simultaneously share the infrastructure and each experiment runs within its own, isolated slice created end-to-end across the experiment’s GENI resources. GENI experiments will be an interconnected set of reserved resources on platforms in diverse locations. Researchers will remotely discover, reserve, configure, program, debug, operate, manage, and tear down distributed systems established across parts of the GENI suite.

The Future Internet Research and Experimentation (FIRE) [FIRE 2009] is a European initiative, apparently influenced by GENI, addressed at creating a multidisciplinary research environment for investigating and experimentally validating highly innovative and revolutionary ideas for new networking and service paradigms. FIRE is promoting the concept of experimentally-driven research, joining the two ends of academic-driven visionary research and industry-driven testing and experimentation.

Differently from GENI, FIRE does not plan to build a single integrated testbed infrastructure, but rather seeks to support a number of independently designed testbeds, which can later be federated to provide interoperation. The first testbed projects supported by FIRE included the projects OneLab, based on PlanetLab, PanLab, an initiative undertaken by a consortium of telecommunications providers [PanLab 2009], and Federica, a large-scale testbed based on network virtualization characteristics of a family of commercial routers, and developed and deployed by a consortium of research and education networks in Europe [FEDERICA 2009].

The New Generation Network Architecture [AKARI 2008] is a FI initiative in Japan, aimed at designing and deploying a clean-slate architecture replacement for the Internet by 2015.

Within the US, some attention has been paid to the question of providing effective programmable network elements (routers and switches) at low cost. The approach used in VINI uses COTS hardware, in the form of Intel-based PCs. Such an approach leads to a low performance network element, as the basic hardware is not at all optimised for high-performance I/O and packet switching. Nick McKeown and his team at Stanford University has concentrated on the problem of producing low-cost programmable network elements of acceptably high performance. Their first proposal is NetFPGA, a high-performance I/O extension board for PCs, providing 4 Gigabit Ethernet ports. Their most recent contribution is an architectural alteration to network element design, where high performance switching hardware is combined with a table-based implementation of the control plane, which can easily be modified by the user, in this case, an experimental network designer. The resulting architecture, known as OpenFlow [McKeown 2008], is designed as an extension to production network element design, and a growing number of switch and router manufacturers already sell OpenFlow-capable hardware, or have advanced plans to do so.

OpenFlow provides an open protocol to program the flow table in different switches and routers. A network administrator can partition traffic into production and research flows. Researchers can control their own flows - by choosing the routes their packets follow and the processing they receive. In this way, researchers can try new routing protocols, security models, addressing schemes, and even alternatives to IP. On the same network, the production traffic is isolated and processed in the same way as today.

4. Brazilian testbeds

Following their first contact in 2001, a partnership was formed between RNP and CPqD to build an optical testbed network in the states of Rio de Janeiro and São Paulo, and to carry out and promote experimental research in network and related application technologies using the testbed as a laboratory. Project GIGA, as the initiative became known, obtained about US$20M in funding from FUNTTEL (Fund for the Development of Telecommunications Technology), of which two thirds were used for the R&D activities. Using dark fibres freely lent by four telcos, the project lit up about 750 km of fibre, including the use of DWDM between the cities of Campinas, São Paulo, São José dos Campos, in São Paulo state, and the city of Rio de Janeiro, extensions to Cachoeira Paulista, in São Paulo state, and Niterói and Petrópolis, in Rio de Janeiro state, and metro networks in Campinas, São Paulo and Rio de Janeiro (see Figure 1). Including laboratories belonging to the collaborating telcos, around 25 institutions were directly served by the resulting testbed network, which used 1 and 10 Gigabit Ethernet technology to provide network transport between collaborating institutions [Scarabucci 2005].

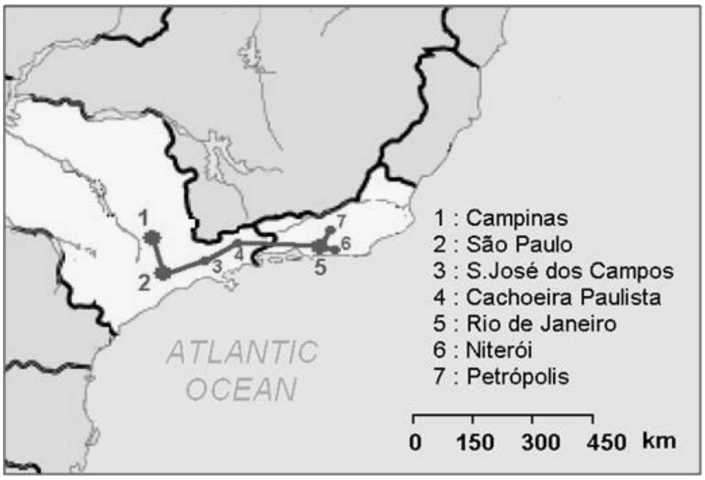

Figure 1: Location of the GIGA testbed network

GIGA received funding between 2003 and 2007, and involved a great many research institutions throughout Brazil. RNP coordinated its own research programme, with its own scientific committee of 4 well-known researchers. This committee received and evaluated 39 proposals, with the assistance of 19 ad-hoc referees from the LARC-SBC community, and approved 33 for support. Due to continuity problems with FUNTTEL funding in 2005 and part of 2006, which caused interruptions in funding to all supported projects, only 27 projects were effectively concluded after funding was restored in 2006. These included 5 in advanced network technologies, and 22 in distributed applications.

RNP's GIGA research programme effectively involved more than 500 participants from 35 Brazilian and 6 foreign research centres, and 13 companies. Academic production amounted to more than 700 documents, including 110 master's and 28 doctoral theses, 180 articles in periodicals, 240 in conferences, 2 books and 16 book chapters. Technical production included around 90 technical reports and manuals, 54 products (conceptual, functional and product prototypes) and 10 services (prototypes), of which 27 were in use at the end of 2008. 12 technologies were effectively transferred to companies, and transfer of technology has been agreed in two further cases. A detailed view of the results of this research programme is also available through the record of the final project workshop, which was held in September, 2007 [GIGA 2007].

Funding for the original GIGA testbed project ended in 2007, when both RNP and CPqD submitted proposals to FUNTTEL for continuation projects, for what would become known as GIGA Phase 2. The RNP proposal, which included adapting the testbed for carrying out research into new Internet architectures, was not considered.

Happily, the CPqD project was approved in part and, after a long delay, finally begun in early 2009. In spite of not having received specific funding for its Phase 2 proposal, RNP continues to partner CPqD in the planning and execution of its Phase 2 activities, and the testbed network continues to be used to support selected projects. The main emphases in this phase will be on optical technologies and on research into Future Internet architectures.

In 2011, RNP will launch the sixth version of its national backbone network, which will reach 24 of the 27 national capitals with bandwidth varying between 3 and 10 Gbps (see the map in Figure 2). This backbone network, complemented by the optical metro networks already partially deployed in capital cities housing points of presence (PoPs) of the backbone network, will be able to provide Gigabit Ethernet service to the campi of around 200 instituions by 2011. It is intended to make available for experimental (testbed) use a significant fraction of the available backbone bandwidth. Effectively, this will enable the expansion of at least a part of the project GIGA testbed to reach a further 21 states and the federal capital, facilitating large scale participation in such experimental research activities.

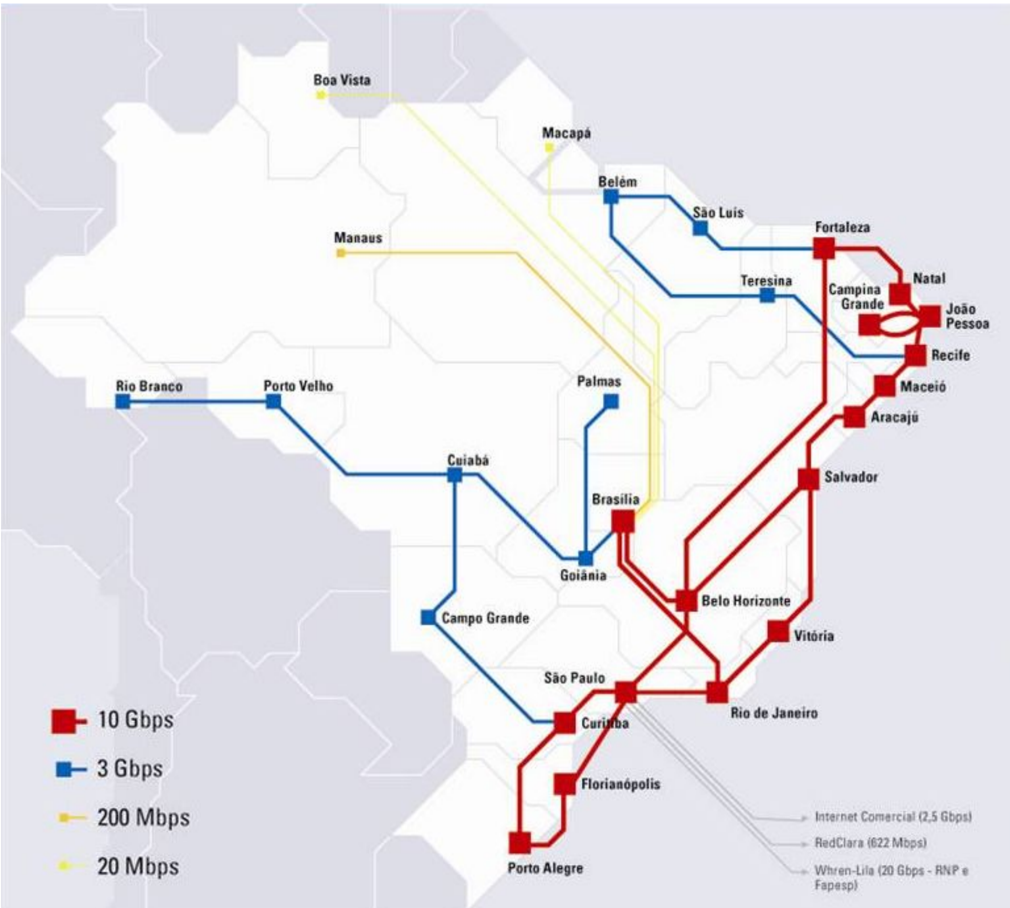

Figure 2;Proposed RNP Phase 6 backbone core in 2010

5. Design choices

The original Future Internet Architecture (FIA) subproject, was written in 2008, and proposed the adoption of the VINI platform [Bavier 2006], based on COTS hardware in the form of standard PCs. As we have noted already, such a choice necessarily leads to a low-performance network.

The delay in effectively beginning this project has allowed us to consider other alternatives, which were not available in 2008, especially the NetFPGA and OpenFlow proposals from Stanford University, which have already been widely adopted for testbeds in network design. In fact, OpenFlow is now an important component of the GENI testbed, and is being widely deployed in this environment. OpenFlow is also the technology of choice for other projects, most notably the Ofelia project in Europe, which will begin in late 2010, coordinated from the University of Essex.

At the recent Workshop on Experimental Research on Future Internet (WPEIF) organised in Gramado in May, 2010, there was also evidence of wider interest in OpenFlow within the Brazilian research community, with news of separate projects from several institutions not currently part of this INCT project. Additionally it should be noted that two Brazilian equipment manufacturers are also considering supporting OpenFlow extensions to their production hardware.

A final choice on which support environment will be used will shortly be taken in the light of these developments.

6. Conclusion

Whichever design choice ends up being taken for the FIA testbed, this will soon lead to testbed deployment within the participating institutions, using large-scale connectivity from RNP, to permit their integration, It is also intended to support the integration of the parallel deployment of a wireless testbed, to be developed and used as another activity within the FIA subproject.

7. References

[AKARI 2008] AKARI project (2008), "New Generation Network Architecture: AKARI Conceptual Design" (ver1.1), http://akari-project.nict.go.jp/eng/concept-design/AKARI_fulltext_e_preliminary.pdf , accessed on 15/12/2009.

[Bavier 2006] Bavier, A. et al. (2006), "In VINI Veritas: Realistic and controlled network experimentation", In: Proc. ACM SIGCOMM, September 2006.

[Choudhury 2010] Choudhury, N.M.M.K., e Boutaba, R. (2010), "A survey of network virtualization", In: Computer Networks, Volume 54, March 2010, p. 862-876.

[FEDERICA 2009] http://www.fp7-federica.eu/ , accessed on 15/12/2009.

[FIRE 2009] Future Internet Research and Experimentation, http://cordis.europa.eu/fp7/ict/fire/ , accessed on 15/12/2009.

[GENI 2009] Global Environment for Network Innovations, http://www.geni.net/ , accessed on 15/12/2009.

[GIGA 2007] R&D Workshop for Project GIGA/RNP, September 2007, LNCC, Petrópolis, RJ, Brazil. Conference agenda and record available at http://indico.rnp.br/conferenceDisplay.py?confId=33 , accessed on 13/12/2009.

[GLIF 2009] http://www.glif.is/ , accessed on 15/12/2009.

[McKeown 2008] McKeown, N. et al. (2008), "OpenFlow: Enabling Innovation in Campus Networks", In: ACM SIGCOMM Computer Communication Review, Volume 38, Number 2, April 2008, p. 69-74.

[PanLab 2009] http://www.panlab.net/ , accessed on 15/12/2009.

[perfSONAR 2009] http://www.perfsonar.net/ , accessed on 13/12/2009.

[Peterson 2002] Peterson, L., Anderson, T., Culler, D. and Roscoe, T. (2002) "A Blueprint for Introducing Disruptive Technology into the Internet", In: Proceedings of ACM HotNets-I Workshop, Princeton, New Jersey, USA, October 2002.

[PlanetLab 2009] http://www.planet-lab.org/ , accessed on 15/12/2009.

[Sampaio 2007] Sampaio, L. et al. (2007), "Implementing and Deploying Network Monitoring Service Oriented Architectures - Brazilian National Education and Research Network Measurement Experiments" (Invited Paper). In: Proc. 5th Latin American Network Operations and Management Symposium - LANOMS, 2007, Petrópolis, RJ, Brazil, v. 1. p. 28-37.

[RedeCipó 2009] Rede Cipó (a testbed for automatic provisioning of circuits), http://wiki.rnp.br/pages/viewpage.action?pageId=26969360 , accessed on 15/12/2009.

[Scarabucci 2005] Scarabucci, R.R., Stanton, M.A. et al. (2005), "Project GIGA – High-speed Experimental Network", In: First International Conference on Testbeds and Research Infrastructures for the DEvelopment of NeTworks and COMmunities (TRIDENTCOM'05), Trento, Itália, 02/2005, p. 242-251.