SBRC 2010

28th Brazilian Symposium on Computer Networks and
Distributed Systems
24-28 May 2010
Gramado, RS, Brazil

1st Workshop on Experimental Research on the Future Internet (WPEIF)

Published by
Sociedade Brasileira de Computação (SBC)

Organizing Committee
Michael Stanton (RNP, UFF)
Marcos Rogerio Salvador (CPqD)
Antônio Jorge Gomes Abelém (UFPA)
Luciano Paschoal Gaspary (UFRGS)
Marinho Pilla Barcellos (UFRGS)

Organization
Instituto de Informática
Universidade Federal do Rio Grande do Sul (UFRGS)

Promoted by
Sociedade Brasileira de Computação (SBC)
Laboratório Nacional de Redes de Computadores (LARC)

Cover: Josué Klafke Sperb
Editorial Production: Flávio Roberto Santos, Roben Castagna Lunardi, Matheus Lehmann, Rafael Santos Bezerra, Luciano Paschoal Gaspary, Marinho Pilla Barcellos.

Additional Symposium Record:
Sociedade Brasileira de Computação (SBC)
Av. Bento Gonçalves, 9500 - Setor 4 - Prédio 43.412 - Sala 219
Bairro Agronomia - CEP 91.509-900 - Porto Alegre - RS
Phone: +55 51 33086835
E-mail: sbc@sbc.org.br

International Data for Cataloguing in Publication (CIP)

Workshop de Pesquisa Experimental na Internet do Futuro (1. : 2010 : Gramado, RS).

Anais / I Workshop de Pesquisa Experimental na Internet do Futuro ; organizadores Michael Stanton... et al. − Porto Alegre : SBC, c2010.
35 p.

ISSN 2177-496X

1. Redes de computadores. 2. Sistemas distribuídos. I. Stanton, Michael. II. Título.

Promotion

Sociedade Brasileira de Computação (SBC)

Promotion

Council

2009-2013
Virgílio Almeida (UFMG)
Flávio Rech Wagner (UFRGS)
Silvio Romero de Lemos Meira (UFPE)
Itana Maria de Souza Gimenes (UEM)
Jacques Wainer (UNICAMP)

2007-2011
Cláudia Maria Bauzer Medeiros (UNICAMP)
Roberto da Silva Bigonha (UFMG)
Cláudio Leonardo Lucchesi (UNICAMP)
Daltro José Nunes (UFRGS)
André Ponce de Leon F. de Carvalho (USP)

Additional Members - 2009-2011
Geraldo B. Xexeo (UFRJ)
Taisy Silva Weber (UFRGS)
Marta Lima de Queiroz Mattoso (UFRJ)
Raul Sidnei Wazlawick (UFSC)
Renata Vieira (PUCRS)

Laboratório Nacional de Redes de Computadores (LARC)

Board of Directors

Director of the Scientific Technical Council
Artur Ziviani (LNCC)

Executive Director
Célio Vinicius Neves de Albuquerque (UFF)

Vice-Director of the Scientific Technical Council
Flávia Coimbra Delicato (UFRN)

Executive Vice-Director
Luciano Paschoal Gaspary (UFRGS)

Institutional Members
CEFET-CE, CEFET-PR, IME, INPE/MCT, LNCC, PUCPR, PUC-RIO, SESU/MEC, UECE, UERJ, UFAM, UFBA, UFC, UFCG, UFES, UFF, UFMG, UFPA, UFPB, UFPE, UFPR, UFRGS, UFRJ, UFRN, UFSC, UFSCAR, UNICAMP, UNIFACS, USP.

Organization

Organizing Committee

General Chairs
Luciano Paschoal Gaspary (UFRGS)
Marinho Pilla Barcellos (UFRGS)

Technical Program Committee Chairs
Luci Pirmez (UFRJ)
Thaís Vasconcelos Batista (UFRN)

International Talks and Tutorials Chair
Lisandro Zambenedetti Granville (UFRGS)

Panels Chair
José Marcos Silva Nogueira (UFMG)

National Tutorials Chair
Carlos Alberto Kamienski (UFABC)

Workshops Chair
Antônio Jorge Gomes Abelém (UFPA)

Tools Session Chair
Nazareno Andrade (UFCG)

Steering Committee
Artur Ziviani (LNCC)
Carlos André Guimarães Ferraz (UFPE)
Célio Vinicius Neves de Albuquerque (UFF)
Francisco Vilar Brasileiro (UFCG)
Lisandro Zambenedetti Granville (UFRGS)
Luís Henrique Maciel Kosmalski Costa (UFRJ)
Marcelo Gonçalves Rubinstein (UERJ)
Nelson Luis Saldanha da Fonseca (UNICAMP)
Paulo André da Silva Gonçalves (UFPE)

Organization

Local Organizing Committee
Adler Hoff Schmidt (UFRGS)
Alan Mezzomo (UFRGS)
Alessandro Huber dos Santos (UFRGS)
Bruno Lopes Dalmazo (UFRGS)
Carlos Alberto da Silveira Junior (UFRGS)
Carlos Raniery Paula dos Santos (UFRGS)
Cristiano Bonato Both (UFRGS)
Flávio Roberto Santos (UFRGS)
Jair Santanna (UFRGS)
Jéferson Campos Nobre (UFRGS)
Juliano Wickboldt (UFRGS)
Leonardo Richter Bays (UFRGS)
Lourdes Tassinari (UFRGS)
Luís Armando Bianchin (UFRGS)
Luis Otávio Luz Soares (UFRGS)
Marcos Ennes Barreto (UFRGS)
Matheus Brenner Lehmann (UFRGS)
Pedro Arthur Pinheiro Rosa Duarte (UFRGS)
Pietro Biasuz (UFRGS)
Rafael Pereira Esteves (UFRGS)
Rafael Kunst (UFRGS)
Rafael Santos Bezerra (UFRGS)
Ricardo Luis dos Santos (UFRGS)
Roben Castagna Lunardi (UFRGS)
Rodolfo Stoffel Antunes (UFRGS)
Rodrigo Mansilha (UFRGS)
Weverton Luis da Costa Cordeiro (UFRGS)

Message from the SBRC General Chairs

Bem-vindo(a) ao XXVIII Simpósio Brasileiro de Redes de Computadores e Sistemas Distribuídos (SBRC 2010)! Esta edição do simpósio está sendo realizada de 24 a 28 de maio de 2010 na pitoresca cidade de Gramado, RS. Promovido pela Sociedade Brasileira de Computação (SBC) e pelo Laboratório Nacional de Redes de Computadores (LARC) desde 1983, o SBRC 2010 almeja não menos que honrar com uma tradição de quase 30 anos: ser reconhecido como o mais importante evento científico em redes de computadores e sistemas distribuídos do país, e um dos mais concorridos em Informática. Mais do que isso, pretende estimular intercâmbio de idéias e discussões qualificadas, aproximá-lo(a) de temas de pesquisa efervescentes e fomentar saudável aproximação entre estudantes, pesquisadores, professores e profissionais.

Para atingir os objetivos supracitados, reunimos um grupo muito especial de professores atuantes em nossa comunidade que, com o nosso apoio, executou com êxito a tarefa de construir um **Programa Técnico** de altíssima qualidade. O SBRC 2010 abrange as seguintes atividades: 20 sessões técnicas de artigos completos, cobrindo uma grande gama de problemas em redes de computadores e sistemas distribuídos; 2 sessões técnicas para apresentações de ferramentas; 5 minicursos ministrados de forma didática, por professores da área, sobre temas atuais; 3 palestras e 3 tutoriais sobre tópicos de pesquisa avançados, apresentados por especialistas nacionais e estrangeiros; e 3 painéis versando sobre assuntos de relevância no momento. Completa a programação técnica a realização de 8 *workshops* satélites em temas específicos: WRNP, WGRS, WTR, WSE, WTF, WCGA, WP2P e WPEIF. Não podemos deixar de ressaltar o **Programa Social**, organizado em torno da temática "*vinho*", simbolizando uma comunidade de pesquisa madura e que, com o passar dos anos, se aprimora e refina cada vez mais.

Além da ênfase na qualidade do programa técnico e social, o SBRC 2010 ambiciona deixar, como marca registrada, seu esforço na busca por excelência organizacional. Tal tem sido perseguido há mais de dois anos e exigido muita determinação, dedicação e esforço de uma equipe afinada de organização local, composta por estudantes, técnicos administrativos e professores. O efeito desse esforço pode ser percebido em elementos simples, mas diferenciais, tais como uniformização de datas de submissão de trabalhos, portal *sempre* atualizado com as últimas informações, comunicação sistemática com potenciais participantes e pronto atendimento a qualquer dúvida. O nosso principal objetivo com essa iniciativa foi e continua sendo oferecer uma elevada *qualidade de experiência* a você, colega participante!

Gostaríamos de agradecer aos membros do Comitê de Organização Geral e Local que, por conta de seu trabalho voluntário e incansável, ajudaram a construir um evento que julgamos de ótimo nível. Gostaríamos de agradecer, também, à SBC, pelo apoio prestado ao longo das muitas etapas da organização, e aos patrocinadores, pelo incentivo à divulgação de atividades de pesquisa conduzidas no País e pela confiança depositada neste fórum. Por fim, nossos agradecimentos ao Instituto de Informática da UFRGS, por viabilizar a realização, pela quarta vez, de um evento do porte do SBRC.

Sejam bem-vindos à Serra Gaúcha para o "SBRC do Vinho"! Desejamos que desfrutem de uma semana agradável e proveitosa!

Luciano Paschoal Gaspary
Marinho Pilla Barcellos
SBRC 2010 General Chairs

WPEIF Program Committee

Antônio Abelém, Federal University of Pará, Brazil
Antonio Marcos Alberti, National Institute of Telecommunications, Brazil
Cesar Marcondes, Federal University of São Carlos, Brazil
José Augusto Suruagy Monteiro, University of Salvador, Brazil
Lisandro Zambenedetti Granville, Federal University of Rio Grande do Sul, Brazil
Maurício Magalhães, University of Campinas, Brazil
Nelson Luis Saldanha da Fonseca, University of Campinas, Brazil
Paulo André da Silva Gonçalves, Federal University of Pernambuco, Brazil
Priscila Solis, University of Brasília, Brazil
Tânia Regina Tronco, Center for Research and Development in Telecommunications, Brazil
Tereza Cristina M. B. Carvalho, University of São Paulo, Brazil

Index

1st Workshop on Experimental Research on the Future Internet

♦

Opening Session

Future Internet initiatives

Michael Stanton

Directorate of R&D – Rede Nacional de Ensino e Pesquisa (RNP)
Rua Lauro Muller 116/1103 – Botafogo – 22290-906 – Rio de Janeiro – RJ – Brazil
michael@rnp.br

(on secondment from Computing Institute, Universidade Federal Fluminense – UFF)

Abstract. *Large-scale Future Internet testbeds are beginning to be deployed in North America, the EU, Japan and Korea. In the USA, NSF launched its GENI (Global Environment for Network Innovations) programme in 2005 and, after several years spent on design, began in 2008 to formulate and deploy an experimental facility to support R&D into new network architectures [GENI 2009]. The EU launched its FIRE (Future Internet Research and Experimentation) programme, also in 2008, based initially on a number of existing testbed projects: OneLab, Panlab, FEDERICA and Phosphorus [FIRE 2009]. Meanwhile, in Japan the AKARI project was launched to design a New Internet by 2015 [AKARI 2008].*

There are ostensibly several similarities between these different proposals, especially in the technologies adopted. In principle, all these testbeds seek to support simultaneous use by concurrent projects (architectures). To carry this out, extensive virtualisation is carried out, both of network resources, including switches, and of processing and storage devices available on the network. This latter facility was originally included as a fundamental part of PlanetLab technology [Peterson 2002], and this has now been extended into network virtualization by variants of PlanetLab, such as VINI, which enable the virtualization of a level 3 router based on a PC [Bavier 2006]

The most general model is that of GENI, which supposes the existence of a level 2 transport service linking network nodes containing programmable and virtualisable routers, as well as processing and storage elements. Among the programmable routers, apart from the VINI model, are such designs as OpenFlow and NetFPGA [McKeown 2008]. On the other hand, the FEDERICA project has adopted the use of production IP routers which support router virtualization [FEDERICA 2009].

One thing is quite clear: there is considerable interest in interoperation of these different testbeds, leading to collaboration around the globe. In Brazil, several invitations have been received to participate in testbed projects which were proposed to GENI in 2009. Therefore, in the planning of a Brazilian Future Internet experimental facility, future interoperation with foreign partners is of great importance.

It should be mentioned that a couple of Brazilian Future Internet R&D projects are already underway: Horizon and WebScience

Horizon is a project to study new Internet architectures, which is being jointly carried out by a consortium of French and Brazilian universities, together with industrial partners, and funded by their respective governments [Horizon 2010].

Web Science is a large consortium of more than 100 researchers from several leading universities, which is being funded for 3 to 5 years of research activity by CNPq, under its National Institutes of Science and Technology programme [WebScience 2010]. RNP and a group of researchers from 5 universities have included in this project the establishment of a VINI-style testbed for experimental research into Future Internet architectures.

Lastly, interest has been expressed at government level in coordinating officially funded projects in the Future Internet area between Brazil and the EU, with a first call expected to be published in 2010.

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.

[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.

[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.

[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.

Biography

Michael Stanton is Director of Research and Development at RNP. After a Ph.D. in mathematics at Cambridge University in 1971, he has taught at several universities in Brazil, since 1994 as professor of computer networking at the Universidade Federal Fluminense (UFF) in Niterói, Rio de Janeiro state. Between 1986 and 2003, he helped to kick-start research and education networking in Brazil, including the setting-up and running of both a regional network in Rio de Janeiro state (Rede-Rio) and RNP. He returned to RNP in 2001, with responsibility for R&D and RNP involvement in new networking and large-scale collaboration projects.

1st Workshop on Experimental Research on the Future Internet

♦

Technical Session 1

Tools and Platforms for Future Internet Experimental Research

Innovating in Your Network with OpenFlow: A Hands-on Tutorial

Brandon Heller[1], Masayoshi Kobayashi[2]

[1]Stanford University, Stanford, USA

[2]System Platforms Research Labs, NEC Corporation, Japan

`{brandonh,mkobaya1}@stanford.edu`

Abstract. *OpenFlow is an open interface for remotely controlling the forwarding tables in network switches, routers, and access points. Upon this low-level primitive, researchers can build networks with new high-level properties. For example, OpenFlow enables more secure default-off networks, wireless networks with smooth handoffs, scalable data center networks, host mobility, more energy-efficient networks and new wide-area networks – to name a few. This tutorial is your opportunity to gain hands-on experience with the platforms and debugging tools most useful for developing network control applications on OpenFlow. Following an introduction, each participant will turn the provided hub controller into an Ethernet switch, then a flow-based switch, and finally a firewall or router - you get to choose. Along the way, you'll learn the full suite of OpenFlow debugging tools: you'll view flow tables with dpctl, dissect packets with Wireshark, visualize with LAVI, slice with FlowVisor, and simulate a multi-switch, multi-host network with Mininet on your laptop. After the tutorial, you can apply what you've learned to physical networks based on software switches, NetFPGAs, or even hardware switches at line rate.*

The only requirement is to bring a laptop; no experience is required.

References

McKeown, N. et al. (2008) "OpenFlow: enabling innovation in campus networks" In: ACM SIGCOMM Computer Communication Review, 38(2):69–74, April 2008.

Biographies

Brandon Heller is a third-year Ph.D. student in Computer Science at Stanford University, with research interests in data center networks and energy efficiency. Brandon currently maintains the OpenFlow specification. His most recent work was ElasticTree, a system to dynamically optimize the energy consumption of data center networks, while taking care to respect performance and fault tolerance considerations (see the NSDI 2010 ElasticTree paper). Other projects include powernet.stanford.edu, a measurement infrastructure for understanding IT-related power consumption at the scale of an entire building, and Ripcord, a modular platform for building data center networks.

Masayoshi Kobayashi received his Bachelor's and Master's engineering degrees from Kyoto University, Japan in 1995 and 1997, respectively. In 1997, he joined NEC and became involved in research including high-speed routers, TCP congestion control and network measurements. Since 2007, he has been with Prof. Nick McKeown’s group at Stanford as a visiting researcher from System Platforms Research Laboratories, NEC Corporation, Japan. Masa was one of the first to be involved in the OpenFlow project, leading the mixed wired and wireless production deployment at Stanford, in addition to contributing to award-winning OpenFlow demos at SIGCOMM in 2008 and 2009.

National infrastructures for Future Internet Research and Development

Marcos Rogério Salvador[1], Michael Stanton[2]

[1]Centro de Pesquisa e Desenvolvimento em Telecommunicações (CPqD)
Rod. Campinas – Mogi-Mirim km 118,5 – 13086-902 – Campinas –SP – Brazil

[2] Directorate of R&D – Rede Nacional de Ensino e Pesquisa (RNP)
Rua Lauro Muller 116/1103 – Botafogo – 22290-906 – Rio de Janeiro – RJ – Brazil
(on secondment from Computing Institute, Universidade Federal Fluminense – UFF)

marcosrs@cpqd.com.br, michael@rnp.br

Abstract. *This paper features the GIGA testbed network [Scarabucci 2005] and the academic network Ipê [Stanton 2010], two of the most relevant networks in Brazil for the support of experimental research. It describes the characteristics and supported functionality of these networks, discusses their limitations and elaborates on how they can be evolved to support Future Internet research and development at various levels by way of large-scale experimentation.*

The GIGA testbed is a large-scale experimental network that can support experiments at any layer. The network has approximately 800km of fiber and interconnects almost 70 labs in 23 universities in 7 cities in the southeast region of Brazil. It currently consists of a static WDM layer and a 10G/1Gb Ethernet layer controlled by a GMPLS control plane, and it is connected to network Ipê and to experimental and academic international networks at transmission rates in the order of gigabits per second. In the near future the WDM layer will evolve to support dynamic reconfiguration and transmission rates of up to 100Gb/s, and the GMPLS control plane will control both layers in an integrated way. 10G OpenFlow-enabled Ethernet switches, currently under development, should gradually replace conventional Ethernet switches to support Future Internet experimental research. FlowVisor will be installed in the network to virtualize these OpenFlow Ethernet switches and, hence, allow for the harmonious coexistence of production traffic, multiple experiments with OpenFlow-based traffic and GMPLS-driven circuit traffic.

The Ipê network is the Brazilian national research and education network, operated by RNP, which provides connectivity to more than 300 institutions throughout Brazil, through PoPs (points of presence) in all 26 state capitals and the national capital. MPLS technology allows simultaneous use of this network for production IP traffic as well as L2VPNs for providing end to end level 2 circuits. The central high-speeed core of the Ipê network currently consists of 10 PoPs linked by 2.5 and 10 Gbps circuits.

By 4Q2010, a new version of this MPLS network will be deployed to reach 14 capitals at 10 Gbps and 10 capitals at 3 Gbps. By the same date, last mile

access to 26 out of 27 PoPs will be provided by wholly-owned optical metro networks, using 1 and 10 Gbps Ethernet links. This will enable at least 1 Gbps access to the more than 200 client institutions located in capital cities. A significant fraction of the capacity of the new Ipê network is intended to support level 3 network experiments, and will be used to extend the geographically limited coverage of the level 2/3 facilities of the GIGA testbed to institutions in the rest of Brazil, using L2VPNs for VLAN tunneling.

To enable international collaboration in experimental Future Internet research, the 20 Gbps of international connectivity which RNP shares with Fapesp-maintained networks in the state of São Paulo, will be used to enable the federation of the GIGA and Ipê experimental network facilities with similar resources in other countries.

References

Scarabucci, R. 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, Italy, February, 2005, p. 242-251.

Stanton, M.A. (2010), “RNP experiences and expectations in future Internet research and development”, In: Tronco, T.R. (ed.), New Network Architectures: The Path to the Future Internet, Springer, 2010 (to appear).

Biographies

Marcos Rogério Salvador obtained his Ph.D. degree in 2003 from University of Twente, in the Netherlands, and since then works for CPqD, where he is currently manager of technology evolution in the areas of optical communications, wireless communications, next-generation networks and electronic transactions. Over his years at CPqD, Marcos has coordinated and executed various high-budget R&D projects in optical networking technologies (e.g., optical packet/burst switching, optical control plane, traffic engineering, topology discovery), resulting in prizes, products transferred to companies, patents and dozens of papers published in scientific events and periodicals of relevance in the area.

Michael Stanton is Director of Research and Development at RNP. After a Ph.D. in mathematics at Cambridge University in 1971, he has taught at several universities in Brazil, since 1994 as professor of computer networking at the Universidade Federal Fluminense (UFF) in Niterói, Rio de Janeiro state. Between 1986 and 2003, he helped to kick-start research and education networking in Brazil, including the setting-up and running of both a regional network in Rio de Janeiro state (Rede-Rio) and RNP. He returned to RNP in 2001, with responsibility for R&D and RNP involvement in new networking and large-scale collaboration projects.

Future Internet Network Management

José Augusto Suruagy Monteiro

Núcleo de Pesquisas em Redes e Computação (NUPERC)

Universidade Salvador (UNIFACS)

R. Ponciano de Oliveira, 126 – Rio Vermelho – 41950-275 Salvador - Brazil

suruagy@unifacs.br

Abstract. *Even in today's technology, end to end performance may not be what users would expect from high speed backbones. In order to identify possible performance holes, it is desirable to instrument such networks, from the user's lab, through his campus, regional, national and international backbones with tools that can perform tests in order to obtain delay and throughput measurements. One such measurement infrastructure is perfSONAR [Hanemann et al. 2005], a service oriented architecture for multi-domain network monitoring. perfSONAR have been deployed in several National Research and Education Networks (NRENs), including Internet2, ESNet, Géant, and RNP, through the MonIPÊ service [Sampaio et al. 2007].*

In testbeds for future Internet experiments such as GENI, OneLab, and Federica, where each experiment runs on a given slice it is even more important to instrument the testbed in order to collect all relevant data for each experiment, including not only network related measurements but also operational data which can help operations personnel on monitoring and troubleshooting the infrastructure itself.

GENI (Global Environment for Network Innovations) is a network research and engineering infrastructure which is currently being designed. GENI key concepts are the following: programmability, virtualization, federation, and slice-based experimentation [GENI 2009]. perfSONAR is already being used in the ProtoGENI infrastructure to monitor Internet2 backbone and other parts of it.

In this presentation we review the proposed measurement infrastructures for Future Internet testbeds, including the GENI Instrumentation and Measurement System (GIMS) [Barford et al. 2009], and PlanetLab Europe's TopHat measurement service [Borgeau et al. 2010].This paper features the GIGA testbed network [Scarabucci 2005] and the academic network Ipê [Stanton 2010], two of the most relevant networks in Brazil for the support of experimental research. It describes the characteristics and supported functionality of these networks, discusses their limitations and elaborates on how they can be evolved to support Future Internet research and development at various levels by way of large-scale experimentation.

References

Barford, P., Blodgett, M., Crovella, M., and Sommers, J.. (2009) Requirements and Specifications for the Instrumentation and Measurement Systems for GENI: Version 0.2, May 2009. http://groups.geni.net/geni/attachment/wiki/MeasurementSystem/ MeasurementSysSpec.pdf.

Bourgeau, T., Augé, J., Friedman, T. (2010) TopHat: supporting experiments through measurement infrastructure federation. Accepted to TridentCom'2010, 18-20 May 2010, Berlin, Germany.

GENI (2009). System Requirements Document. Document ID: GENI-SE-SY-RW-02.0. Available at http://groups.geni.net/geni/attachment/wiki/SysReqDoc/GENI-SE-SY-RQ-02.0.pdf

Hanemann, A., Boote, J. W., Boyd, E., Durand, J., Kudarimoti, L., Lapacz, R., Swany, M., Trocha, S., and Zurawski ,J.. (2005). Perfsonar: A service oriented architecture for multi-domain network monitoring. In Proceedings of the Third International Conference on Service Oriented Computing (ICSOC 2005). ACM Sigsoft and Sigweb.

Sampaio, L., Koga, I., Costa, R., Monteiro, H., Monteiro, J. A. S., Vetter, F., Fernandes, G., Vetter, M.. Implementing and Deploying Network Monitoring Service Oriented Architectures - Brazilian National Education and Research Network Measurement Experiments (Invited Paper). In: 5th Latin American Network Operations and Management Symposium (LANOMS), 2007. Petrópolis, RJ, p. 28-37.

Biographies

José Augusto Suruagy Monteiro received his B.S. in Electrical Engineering from Universidade Federal de Pernambuco (Brazil) in 1979, his M.Sc. in Electrical Engineering, from Universidade de São Paulo (Brazil) in 1982, and his Ph.D. in Computer Science from University of California, Los Angeles, in 1990. He is currently a faculty member at Universidade Salvador (Brazil) where his research interests include Internet measurements, network resources allocation, and traffic engineering.

1st Workshop on Experimental Research on the Future Internet

♦

Technical Session 2

Future Internet Experimental Research Works 1

LISP as a solution for Internet scalability

Jerônimo Bezerra, Humberto Galiza, Luciano Porto Barreto,
Luiz Cláudio Mendonça, Claudete Alves

[1]Ponto de Presença da Rede Nacional de Pesquisa (PoP-BA/RNP)
Av. Ademar de Barros, s/n, UFBA, Ondina – 40110-140 – Salvador – Ba – Brasil

[1]Laboratório de Sistemas Distribuídos – Departamento de Ciência da Computação –
Universidade Federal da Bahia (LASID/DCC/UFBA)
Av. Ademar de Barros, s/n, UFBA, Ondina – 40110-140 – Salvador – Ba - Brasil

{jab,galiza,lportoba,mendonca,claudete}@ufba.br

Abstract. *The exponential growth of BGP routing tables and mobility issues are two major problems in the Internet today. This paper presents an overview of LISP and our initial experiences in joining the LISP+ALT network. We also present our goals on using LISP as a testbed for Internet research in mobility, scalability, management and security.*

The Internet Architecture Board (IAB) Workshop in October 2006 [MEYER, 2007] pointed out the need to deal with current and future problems related to routing scalability and addressing for the Internet. Besides the exhaustion of IPv4 addresses, a major problem is the exponential growth of BGP routing tables due to network configuration and practices such as site multihoming and traffic engeneering. Mobility is a leading issue as well, which requires low latency and needs to scale. Therefore, the way the Internet works today does not provide an adequate solution for mobile users and scalable routing.

BGP is the core routing protocol of the Internet. It manages a table of IP networks (prefixes) which defines how autonomous systems are connected in order to perform routing decisions. The lack of prefix aggregation is causing BGP tables to grow continuously. Several researchers and practitioners believe that this may dangerously affect the operation of the Internet in the near future. It is important to state that IPv6 does not solve the problem, since it is based on the same routing infrastructure as IPv4. Therefore, the problem will not simply go away with IPv6.

It is common knowledge nowadays that the source of these problems is mainly due to the fact that IP addresses are both used for location and identification. This so called "IP address overloading" lessens mobility and imposes several complications to the routing infrastructure. Hence, there is an emergent need to decouple identification and location in order to provide flexibility and mobility for Internet users and administrators.

LISP [MEY 2008] [FARINACCI,2010] is one such a solution which splits the addressing scheme into two namespaces: identifiers and locators. LISP basically inserts one level of indirection at the network stack to perform packet encapsulation and transparent routing. LISP employs a dynamic encapsulation scheme in which outside addresses refer to locators and inside addresses refer to user`s IDs. As such, LISP packets are transparently routed in the Internet and it requires a few changes in software and no specific *configuration at network hosts. Furthermore, LISP does not require huge modifications to the network infrastructure and it is interoperable with IPv4 and IPv6.*

In Latin America, the number of announced IP prefixes is twice as big as the rest of the world [YANNUZZI,2009]. At the end of 2009, we conducted a study to verify the benefits of using LISP. This work showed a potential reduction up to 89% of the BGP routing table, mainly due to prefix aggregation. Afterwards, we contacted Lisp Working Group and joined the LISP+ALT network [FULLER,2009] at the beginning of 2010. To the best of our knowledge, there are only two institutions (UFBA in Brazil and LACNIC in Uruguay) in the LISP+ALT network in Latin America. We hope that other institutions get motivated to join this network as well. Finally, our main goal is to use the LISP+ALT network as a testbed to perform studies on the analysis of LISP implementations, internet scalability, mobility and security.

References

Farinacci, D., et al. (2010), Locator/ID Separation Protocol (LISP). Internet Draft.

Fuller, V., et al. (2009), LISP Alternative Topology (LISP-ALT). Internet Draft.

Meyer, D, et al. (2007), Report from the IAB Workshop on Routing and Addressing. RFC 4984, September 2007.

Meyer, D. (2008). The Locator Identifier Separation Protocol (LISP). The Internet Protocol Journal, vol. 11, No 1, 2008.

Yannuzzi, M. et at. (2009). Managing Interdomain Traffic in Latin America : A new perspective on LISP. IEEE Communications Magazine, vol. 47, p. 40-48, July 2009.

Biographies

Jerônimo Aguiar Bezerra obtained his Computer B.S. degree in 2009 and he is currently a MSc. student in Mechatronics at the Federal University of Bahia, where he works at the Point of Presence of RNP in Bahia.

Humberto Galiza obtained his Computer B.S. degree in 2009 from Federal University of Bahia and he is currently a graduate student in Complementary Applications on Military Sciences at the Brazilian Army Administration School – EsAEx.

Luciano Porto Barreto holds a Ph.D. degree in 2003 from the University of Rennes I (INRIA/IRISA) in France. He is currently a lecturer at the Computer Science Department of Federal University of Bahia.

Luiz Cláudio Mendonça obtained his Computer B.S. degree in 1982 from Federal University of Bahia, where, currently, is the Chief Executive Officer of Internet Data Center.

Claudete Alves holds a Computer M.Sc. degree in 1988 from Federal University of Paraíba and, actually, is coordinator of Point of Presence in Bahia of RNP.

Flat Routing in Internet-like Topologies

Fábio L. Verdi[1], **Rafael Pasquini**[2], **Maurício F. Magalhães**[2]

[1]Federal University of São Carlos - 18052-780
Sorocaba - SP - Brazil

[2]State University of Campinas - 13083-970
Campinas - SP - Brazil

verdi@ufscar.br, {pasquini, mauricio}@dca.fee.unicamp.br

Abstract. *A frequent subject in forums, academia and industry is the evolution of the Internet in terms of routing. Basically, the scalability in the Default Free Zone related to (1) the growing rate of the routing tables and (2) the convergence of the routing system, are pointed by routing experts as the main concerns of the current mechanism. Several approaches have emerged, but they normally require a mapping system to translate from identifiers to locators (IP). We contribute with this discussion by introducing our XOR-based Flat Routing mechanism for Internet-like topologies. Essentially, the proposed mechanism routes directly on top of flat ASes identifiers, eliminating the need for mapping systems. In this work we present a mechanism for building the routing tables over the XOR-based scenario in conjunction with a reachability service developed using the concepts of Landmark and Bloom filters. We also present our recent results related to the usage of the flat routing mechanism in data centers based on hipercube topologies.*

1. General overview

This proposal employs an XOR-based routing mechanism which is used to build a mesh network structure, as opposed to the virtual ring organization proposed in VRR [Caesar et al. 2006a] and ROFL [Caesar et al. 2006b]. Basically, XOR-based routing mechanisms are available in the literature [Ford 2003, Maymounkov and Mazières 2001], and this work leverages their routing tables organization model and forwarding mechanism, proposing a brand new process for building the routing tables over such XOR-based routing mechanism which (1) removes the need for any kind of underlay (tunneling) network providing communication between nodes, and (2) considers the concept of physical proximity in number of hops for building the routing tables. The first difference has the benefit of totally eliminating the need for mapping systems, and the second difference provides the fundamental basis for controlling the signaling overhead (convergence of the overall routing system) due to the adopted regionalism approach.

The convergence is also pointed by routing experts as one important scale-limiting factor of current routing mechanisms. Since the current mechanisms require 100% of information (global knowledge) in the DFZ about the network condition, a change in a given part requires messaging throughout all the network in order to converge to the updated network condition. Consequently, a routing mechanism which requires only regional messaging (not global convergence), and regional information (not global knowledge) in the routing tables, becomes a promising research topic to tackle some problems so far identified.

The proposed routing mechanism is evaluated over several Internet-like (Power-law) topologies. All the evaluations were conducted using our developed emulation tool, and the thorough evaluation considers more than 620 million paths (it was computed 100% of the paths for all evaluated topologies), allowing an insightful discussion around the proposed routing mechanism.

References

Caesar, M., Castro, M., Nightingale, E. B., O, G., and Rowstron, A. (2006a). Virtual Ring Routing: Network Routing Inspired by DHTs. *In Proceedings of the ACM SIGCOMM*, pages 351 – 362.

Caesar, M., Lakshminarayanan, K., Condie, T., Stoica, I., Kannan, J., and Shenker, S. (2006b). ROFL: Routing on Flat Labels. *In Proceedings of the ACM SIGCOMM*, pages 363 – 374.

Ford, B. (2003). Scalable internet routing on topology-independent node identities. *MIT technical report conducted as part of the IRIS project (http://project-iris.net)*.

Maymounkov, P. and Mazières, D. (2001). Kademlia: A peer-to-peer information system based on the xor metric. *In 1st International Workshop on Peer-to-Peer Systems (IPTPS)*.

Fábio Luciano Verdi is currently a Professor at Federal University of São Carlos (UFSCar), Campus Sorocaba. He received his Master degree in Computer Science and Ph.D. degree in Electrical Engineering both from State University of Campinas (UNICAMP). His main interests include computer networks, mobility, routing, service oriented architectures, Future Internet, data centers and cloud computing.

OpenMesh: OpenFlow in Wireless Mesh Networks

Billy Anderson[1], Fernando Farias[1], Vagner Nascimento[1], Antônio Abelém[2]

[1]Programa de Pós Graduação em Engenharia Elétrica (PPGEE) - Universidade Federal do Pará (UFPA)
[2]Programa de Pós Graduação em Ciências da Computação (PPGCC) - Universidade Federal do Pará (UFPA)

{billy,fernnf,vagner,abelem}@ufpa.br

Abstract. *The adoption of Wireless Mesh Networks (WMNs) is a solution to provide last mile indoor and outdoor Internet access and is gaining an important attention academic and industry research groups in recent years. WMNs will support the distribution of diverse type of services, ranging from battlefield surveillance to high quality mobile audio and video applications.*

As an alternative to proprietary solutions, IEEE 802.11 equipments with an extended firmware can be used as an open source and low-cost solution in Linux environments [Tsarmpopoulos et al. 2005]. This type of solution allows the easy creation of digital and multimedia environments, enabling the distribution of new services for fixed and mobile devices and attracting new customers.

There is an open source WMN solution used in a production network in the Federal University of Pará (UFPA). The WMN that was deployed in an area that has buildings with an average height of eight meters, with a predominance of large trees, typical of the Amazon region, and also with high rate of rainfall.

The WMN located in the UFPA has six mesh kits. These kits are made from air-tight box to hold the wireless routers and omnidirectional antenna of 18.5 dBi gain. The wireless routers used in the network are Linksys WRT54GL.

Since it is a production network, is difficult to run experimental protocols on it. However with the aim to solve this problem, the OpenFlow solution [McKeown et al. 2008] was propose, to allow tests with new network architectures and protocols over a production network.

On this work we show the use of OpenFlow protocol solutions and FlowVisor [Sherwaood et al. 2009] in a OpenMesh [Pinheiro et al. 2009] network, to provide a laboratory in a production networking. The OpenFlow was deployed in ours router, that have the OpenWRT firmware already, and the FlowVisor was used to created slices, where multimedia traffic was inserted to evaluate the network metric, mainly the delay produced by the redirecting of flows over a wireless multi-hop network.

References

Tsarmpopoulos, N., Kalavros, I., Lalis, S.: A Low-Cost and Simple-to-Deploy Peer-to-Peer Wireless Network based on Open Source Linux Routers. In: First international Conference on Testbeds and Research infrastructures For the Development of Networks and Communities. IEEE Computer Society, Washington, DC, 92-97 (2005)

N. McKeown, T. Anderson, H. Balakrishnan, G. Parulkar, L. Peterson, J. Rexford, S. Shenker, and J. Turner. Openflow: enabling innovation in campus networks. SIGCOMM Comput. Commun. Rev., 38(2):69–74, 2008.

Pinheiro, B., Nascimento, V., Moreira, W., Abelém, A.: Abaré: A deployment and management framework for wireless mesh network. In: IEEE Latin-American Conference on Communications. Colombia (2009)

Rob Sherwood, Michael Chan, Adam Covington, Glen Gibb, Mario Flajslik, Nikhil Handigol, Te-Yuan Huang, Peyman Kazemian, Masayoshi Kobayashi, Jad Naous, Srinivasan Seetharaman, David Underhill, Tatsuya Yabe, Kok-Kiong Yap, Yiannis Yiakoumis, Hongyi Zeng, Guido Appenzeller, Ramesh Johari, Nick McKeown, and Guru Parulkar. Carving research slices out of your production networks with OpenFlow. In Proceedings of ACM SIGCOMM (Demo), Barcelona, Spain, August 2009.

Biographies

Billy Anderson is member of Study Group on Computer Networks and Multimedia Communications (GERCOM), graduated in Computer Science for Federal University of Para(UFPA). Currently, is a Ph.D. student of pos-graduate program in electrical engineering at UFPA. Has experience in program language, wireless mesh network and management.

Fernando Farias received a B.Sc degree in Mathematics from Federal University of Pará in 2003 and his M.Sc. degree in Electrical Engineering from the same institution in 2008. He is now a Ph.D. student at the Technology Institute, Federal University of Pará, Belém , Brazil. His main research activities are in the field of networking and include, next-generation optical networks, traffic engineering and network management. His current research is related to management issues for next-generation optical networks.

Vagner Nascimento Born in Macapa, Amapa, Brazil, on January 23, 1981.It concludes his graduate degree in Computer Science at the University Center of Pará (CESUPA) in 2005. Active student of the Graduate Program in Computer Science (PPGCC) at Federal University of Pará (UFPA). As a member of the Study Group on Computer Networks and Multimedia Communications (GERCOM) published articles in wireless mesh networks, TCP and cross-layer.

Antônio Jorge Gomes Abelém is Associate Professor on the Computer Science Faculty at the Federal University of Pará (UFPA) in Belém, capital of the Brazilian state of Pará. He holds a doctorate in Informatics from the Catholic University of Rio de Janeiro (PUC-Rio). He was the General Coordinator of the 25th Brazilian Symposium of Computer Networks and Distributed Systems (SBRC 2007), which was held in Belém. He is local coordinator of MetroBel Project, which has developed an optical metropolitan area network (MAN) project in Belém, based on community ownership of the telecommunications infrastructure, and used to interconnect the higher education and research institutions in Belém. He is a member of the IEEE Communications society and Brazilian computer society (SBC). He represented the Brazilian Computing Society (SBC) as a member of the Administrative Council of the Brazilian National Research and Education Network, RNP (2007-2009).

Remote Centralized IP Routing Protocol using Openflow

Marcelo R. Nascimento[1,2], Christian E. Rothenberg[1,2], Marcos R. Salvador[2], Mauricio F. Magalhaes[1]

[1]Faculdade de Engenharia Elétrica e de Computação – Unicamp
Cidade Universitária – Barão Geraldo – 13083-970 – Campinas – SP – Brazil

[2]Centro de Pesquisa e Desenvolvimento em Telecommunicações (CPqD)
Rod. Campinas – Mogi-Mirim km 118,5 – 13086-902 – Campinas –SP – Brazil

mauricio@dca.fee.unicamp.br, marcosrd@cpqd.com.br, esteve@cpqd.com.br, marcelon@cpqd.com.br

Abstract. *In this work, we take a necessary intermediate step when moving towards programmable (software-defined) networks, namely transparent interaction with legacy networks. While the OpenFlow [1,2] model already provides means to treat operational traffic with legacy protocols embedded in the switches themselves, we explore the feasibility of completely moving a legacy routing protocol stack (Quagga [3]) to logically centralized controllers and using the OpenFlow protocol (together with the OpenFlow switch configuration protocol) as the solely communication channel with the forwarding engines. More precisely, we propose an OpenFlow controller application acting as a proxy between Quagga instances and their physical counterparts. Basically, we intend to replicate the discovered physical topology by stitching the virtual interfaces of the VMs running Quagga.*

This piece of work gets us closer to benefiting from the OpenFlow model in terms of rapid innovation and rich network control opportunities. In a short term, the core benefits of our external Quagga integration with OpenFlow include:

1.- Cheap network gear with minimal embedded software.

2.- Avoid re-writing legacy protocols in a centralized fashion.

3.- Ensure interoperability with legacy network elements.

Along this way, we expect to contribute to the revisions of the OpenFlow table abstraction, and the requirements of the OpenFlow configuration protocol.

As a result of this exploratory work we do not only intend to contribute to the feasibility of progressively adopting OpenFlow but we look forward to devise opportunities in offering "legacy routing emulation services." This way, network operators can re-use their current practices of network management and approach network virtualization by safely offering third parties (e.g., customers) controlled access to isolated Quagga instances in order to customize the routing of their "contracted" network slices. In the long term however, network virtualization should move beyond the current network

virtualization model based on overlaying a virtual network of multiple virtual routers on top of a shared physical infrastructure and propose a more convenient, user-friendly “Platform as a Service” model for networking [4].

References

[1] The OpenFlow Switch Consortium. http://www.openflowswitch.org

[2] N. McKeown, T. Anderson, H. Balakrishnan, G. Parulkar, L. Peterson, J. Rexford, S. Shenker, and J. Turner. OpenFlow: enabling innovation in campus networks. ACM SIGCOMM Computer Communication Review, 38(2):69–74, April 2008.

[3] Quagga Software Routing Suite. http://www.quagga.net

[4] Eric Keller and Jennifer Rexford, "The 'Platform as a Service' model for networking," to appear in Proc. Internet Network Management Workshop and Workshop on Research in Enterprise Networking, April 2010.

Biography

Marcelo Ribeiro Nascimento graduated in 2008 from State University of Campinas (São Paulo, Brazil) in Electrical Engineering. Since them he works for CPqD in charge of embedded software development and research in telecommunication area. In 2009 he started a master's degree program in Computer Networks at State University of Campinas, focusing on innovation network technologies.

Towards an Agent-based NOX/OpenFlow Platform for the Internet

Alexandre Passito, Edjard Mota, Rodrigo Braga

[1]Laboratório de Computação Inteligente e Autônoma (LabCIA)
Departamento de Ciência da Computação (DCC)
Universidade Federal do Amazonas (UFAM)
Manaus – AM – Brazil

{passito,edjard,rodrigo.braga}@dcc.ufam.edu.br

Abstract. *Concerns about security, mobility, routing, quality-of-service, middle-boxes led to revisions in the basic architectural framework of the Internet, such as new schemes for addressing, new features for switching/forwarding equipment and deployment of new services using overlay networks. Some of these revisions depend on network domains to cooperating in order to achieve scalability. For example, [Liu et al. 2008] proposes a filter-based DDoS defense system which depends on an infrastructure service to be deployed for each network domain in order to block attack traffic. Somewhat solution seems hard to be applied at large scale because each network has its own administrative boundaries and internal policies, and incentives to cooperation can be hard to achieve.*

To tackle the highly distributed nature of the Internet and such increasingly complex interactions, each administrative domain should be modeled as an autonomous society of agents (or a multiagent system). A multiagent system, [Wooldridge 2009], is one that consists of a number of agents, which interact with one another on behalf of owners with different goal and motivations. In order to successfully interact, these agents will thus require the ability to cooperate, coordinate and negotiate with each other.

This ongoing research aims to enhance network operating systems [Gude et al. 2008] with agent capabilities, such as reactivity, pro-activity and social ability. This approach will enable the building of artifacts for the autonomous control of networks, allowing networks to self-govern their behavior, but only within the constraints of the goals that the system as whole seeks to achieve. To tackle large scale Internet problems, social abilities like cooperation and negotiation is being used to make agents interact with other network domains. Using such high-level and centralized abstraction of the network will reduce the complexity of building agents in a too complex and often uncertain environment. This feature expressively reduces the burden to construct a translation layer into each agent to cope with different network vendors. From the network operating system viewpoint, agents are used as an efficient manner to build autonomous network control artifacts. Applications, now characterized as agents, can be used to build self-managed networks and exploring autonomous solutions for configuration, optimization, recuperation and protection.

The implementation/experimentation of the agent-based network operating system uses the NOX/OpenFlow [Gude et al. 2008] platform and is under development by LabCIA. Our preliminary result using the platform under construction

is self-protecting networks against flooding DDoS attacks [Braga et al. 2010]. The detection method uses Self Organizing Maps, an unsupervised artificial neural network, trained with features of the network traffic. Taking advantage of the abstraction layer, it is possible to extract flow-based information from all OF switches registered by NOX. Furthermore, there is great flexibility to add/remove OF switches into/from the detection loop. The benefit of this possibility is that if there is a change in the network topology, it is possible to autonomously adapt to it adding switches more relevant to detection or removing those less important. This new method also yielded low rates of false alarms and high rates of attack detections using flow information instead of per-packet information.
For mitigating detected attacks, we plan to allow the cooperation between agents from different network domains. For example, if network A detects that an attack has been launched, it could inform B that is under attack and ask B to filter the packets in the origin. After negotiation, B uses network biddings of NOX/OpenFlow to block the attack source.
We stand on the assumption that OpenFlow will gradually be adopted by network domains. OpenFlow is not a toy protocol. It makes part of a world-wide consortium trying to innovate in networks research and its deployment is reaching large scale infra-structures like GENI and Internet2. Additionally, a new proposal of IP and Transport unification through OpenFlow increases its possibility to future large scale adoption. As soon this assumption holds, our research will be valuable to show how intelligent scalable architectures can be used to build the future Internet.

References

Braga, R., Mota, E., and Passito, A. (2010). Lightweight DDoS Flooding Attack Detection Using NOX/OpenFlow. In *Under submisstion to IEEE Conference on Local Computer Networks (LCN)*.

Gude, N., Koponen, T., Pettit, J., Pfaff, B., Casado, M., McKeown, N., and Shenker, S. (2008). NOX: Towards an Operating System for Networks. *ACM SIGCOMM Computer Communication Review*, 38(3):105–110.

Liu, X., Yang, X., and Lu, Y. (2008). To filter or to authorize: Network-layer DoS defense against multimillion-node botnets. In *ACM SIGCOMM*, pages 195–206. ACM.

Wooldridge, M. (2009). *An Introduction to MultiAgent Systems*. John Wiley & Sons, 2 edition.

Biographies

Alexandre Passito is M.Sc. (2008) and Ph.D. student in Informatics (UFAM).

Edjard Mota is Ph.D in Artificial Intelligence (The University of Edinburgh, 1998).

Rodrigo Braga is M.Sc. student in Informatics (UFAM).

Mobility: Challenges and Opportunities

Luiz Claudio Schara Magalhães, PhD

Depto de Eng. de Telecomunicações – Universidade Federal Fluminense (UFF)
Rua Passo da Pátria 156 – 24220-121 – Niterói – RJ – Brazil

schara@midiacom.uff.br

Abstract. *This presentation focuses on two main ideas: that mobile devices should be mobility-aware, and that hiperconectivity (increasing the number of paths at the edge of the Internet by creating large "cooperative" networks) is a way of creating resiliency in the increasingly vital resource which is Internet access.*

When the Internet was created, the type of mobile devices we have now were virtually unknown. In fact, all computers connected to the network were mainframes with rare exceptions. Therefore, tying together identification and routing in a single identifier made sense. When mobility was introduced, the solution chosen to support it was to hide mobility. Mobile-IP maintains a home network as the canonical address of the mobile host, and mobility is achieved by tunneling from this home network to the real position of the mobile. Although this allows traditional applications to be used without change, this makes mobile devices second class citizens of the Internet. Their routing will be necessarily more inefficient than their fixed counterparts, and the devices are not able to take advantage of their adaptability.

As Internet access becomes ubiquitous, being used in all public and private instances of life, it is imperative to guarantee that services will not fail when a link fail. Moreover, the loosely hierarchical architecture of the Internet, which was a result of interconnecting autonomous networks, becomes strained when networks lose their communal aspect. When the Internet was created by connecting research and University networks, the individual networks had users which had great affinity, and traffic was in part internal. In contrast, current commercial providers create groupings that are random, and traffic tends to be extra network, generating topologically inefficient routing depending on the peering between networks.

Mobile devices generally have multiple network interfaces. For example, the well know Nokia N95 cell phone has five radios and seven antennas, plus both an USB and an infrared port. This allows a N95 to have simultaneous infrared, wifi, Bluetooth and 3G links. Alas, normal IP routing will only allow one of these interfaces to be used at a time. An aggregate link, made up of all available physical links will greatly improve throughput and resiliency of communication. If a mobile host can change the set of links being used at any time, mobility can be achieved simply by trading a failing link with a new one that becomes available.

Hiperconnectivity helps link availability by creating a framework where cooperative networks can be created. Instead of postulating the connection to a central provider, the idea is to create networks interconnecting physical neighbors. The link technology to do that is already available, and mesh networks are one instance of that technique. Another example is the community network called REDECOMEP in Brazil, where dark fiber is launched to interconnect research institutions and Universities, as well as municipal and state institutions. To create a higher degree of interconnection at the edges allows for alternative paths in case of failure, and for better routing of local traffic. There are many research and even business opportunities for cooperative networks, and it is a great solution for providing good, high bandwidth connection at low cost for better last mile capillarity.

Hiperconnectivity creates an environment with many opportunities for network access, which can be explored by mobile hosts, especially if they are able to use well connectivity-rich environments. This provides fertile ground for new applications, which can take advantage of local, low cost connectivity as well as location information to provide new services for mobile hosts.

References

Carrano, R et al, Multihop MAC: IEEE 802.11s Wireless Mesh Networks, Encyclopedia on Ad Hoc and Ubiquituous Computing, 2009

Magalhães, L, A Transport Layer Approach to Host Mobility, PhD Dissertation, 2002

Biography

Luiz Claudio Schara Magalhães has a PhD in Computer Science from the University of Illinois at Urbana-Champaign. He has been a professor at the Telecommunications Department at Universidade Federal Fluminense since 1994, instructor at UIUC, Visiting Scholar at HP-Labs in Palo Alto and is currently working for the Brazilian government on the deployment of mesh-network capable educational laptops. His main research interests lie in the area of mobility, large scale infrastructure to support mobile nodes, hiperconnectivity (high-redundancy last mile access) and device environment awareness and cooperation.

1st Workshop on Experimental Research on the Future Internet

♦

Technical Session 3

Future Internet Experimental Research Works 2

Research infrastructures for a Cloud-driven future Internet

Christian Esteve Rothenberg and Maurício Ferreira Magalhães

Faculdade de Engenharia Elétrica e de Computação (FEEC)
Universidade Estadual de Campinas (Unicamp)
Caixa Postal 6101 – 13083-970 – Campinas – SP – Brazil

{chesteve,mauricio}@dca.fee.unicamp.br

Abstract. *The availability of vast commoditized processing and storage technologies together with the success of the Internet has enabled the realization of a computing mode known as Cloud Computing, also touted as the Fifth Utility. Although it is still an evolving paradigm, there is little disagreement in recognizing that the Cloud is here to stay. Consequently, the inter-networking requirements of the Cloud should be regarded as a major driving force shaping the future Internet, i.e., the set of protocols governing the global-scale interaction of computing processes.*

In a recent talk [1], Internet evangelist Vint Cerf has noted that "the Cloud represents a new layer in the Internet architecture and, like the many layers that have been invented before; it is an open opportunity to add functionality to an increasingly global network." Indeed, it can be argued that the emergence of the Cloud is another example of the history repeating itself. We can draw an analogue between Cloud Computing and the time when simple terminals devices (i.e., thin clients) accessed mainframes running applications and hosting users´ data. Also noteworthy, there are three Cloud initiatives that have an analogue in the Internet's past, including (i) the rising importance of academia, (ii) the increasing interest in interoperability among cloud vendors, and (iii) the carrier (infrastructure owner) interest in new service opportunities.

While the transformations of the Cloud in the service area are more-or-less more obvious, the implications of the core networking technologies are less predictable. It has been reported that Cloud computing has contributed to a flatter Internet topology [2] as a consequence of content providers deploying their own wide-area networks, closer to users and bypassing Tier-1 ISPs on many paths. In addition to peering, other profound impacts may include mobile offload for large data volumes from 3G/4G networks or further developments of application content infrastructures [3].

When distributing the Cloud data center footprint closer to the users, network providers are in a good position to leverage their infrastructure and knowledge (e.g., topology, usage patterns) while embracing the Cloud business model. A noteworthy technology enabler to become smart pipes is current industry trend towards open programmable network platforms (e.g., OpenFlow, Cisco IOS, Juniper Junos), aka network operating systems providing true network support for challenging applications, enabling rapid introduction of network innovations, and custom networks for different segments.

While the architecture for Inter-Cloud standards is still very much in its early stages [4] and the market implications are unclear, we may however expect that a global market of cloud resources (e.g., computational, storage, networks) will eventually

emerge. We may speculate that the geo-distribution of data centers may drive the development of a cloud-oriented connectivity market beyond traditional multi-homing.

Driven by the same forces that are leading to an architectural transformation of the data centers, we should expect innovative inter-networking solutions, re-examining cost/control trade-offs and revisiting the incentives for a broader adoption of mature protocols (e.g., VPNs, IPv6, IP Multicast, Secure BGP, DNS Security Extensions), novel routing mechanism (e.g., LISP, energy-aware protocols, or even XML routing) and radical information-centric inter-networking architectures (e.g., CCN, PSIRP).

Future Internet projects need to move from technical reports to running code and large scale prototypes that illuminate the benefits of beyond-IP proposals while devising an incentive-proven migration path is devised. This pragmatic step in future Internet research becomes clear in initiatives like GENI or the EU PPP and the EU new flagship large-scale integrating project called SAIL. As a key step, experimental research infrastructures should be fed by the requirements and opportunities of the Inter-Cloud, exposing the research community to enterprise-level requirements, yielding realistic traces of cloud workloads and all in all foster networking research around cloud computing developments by federating heterogeneous datacenters. By embracing open source developments (e.g., Eucalyptus) and vendor-independent open APIs (e.g., OpenFlow), existing infrastructures can be leveraged setting the path of the next generation of network devices. Remarkable approaches towards those Cloud-driven research infrastructures include OpenCirrus [5] and the GENICloud project, which among other goals aim at providing a federation interface for compute clusters running Eucalyptus open-source software infrastructure.

References

[1] Cloud Computing and the Internet, Available online (last access on 03/08/09): http://googleresearch.blogspot.com/2009/04/cloud-computing-and-internet.html

[2] Gill et al., “The Flattening Internet Topology: Natural Evolution, Unsightly Barnacles or Contrived Collapse?” PAM 08

[3] Bill St. Arnaud, A personal perspective on the evolving Internet and Research and Education Networks, Draft, Latest Version February 15, 2010. [Available online]

[4] Bernstein el al. "Blueprint for the Intercloud - Protocols and Formats for Cloud Computing Interoperability," In ICIW 2009.

[5] Avetisyan et al., "Open Cirrus: A Global Cloud Computing Testbed," IEEE Computer, vol 43, no 4, April 2010, pp 42-50.

Biography

Christian Esteve Rothenberg is a research scientist at CPqD, Brazil, and currently works towards his PhD on compact forwarding methods in data-centric networks at University of Campinas (Unicamp), Brazil. He holds a Telecommunication Engineering degree from the Technical University of Madrid (UPM), Spain, and a German Diplom in Electrical Engineering and Information Technology from the Darmstadt University of Technology.

Publish/subscribe architecture with transparent mobility support

Gustavo Baptista[1], Markus Endler[1], Vagner Sacramento[2]

[1] Departamento de Informática – Pontifícia Universidade Católica do Rio de Janeiro (PUC-Rio). Caixa Postal 38.097 – 22.451-900 – Rio de Janeiro – RJ – Brazil

[2] Instituto de Informática – Universidade Federal de Goiás (UFG)
Caixa Postal 131 – 74.001-970 – Goiânia – GO – Brazil

{gbaptista,endler}@inf.puc-rio.br, vagner@inf.ufg.br

Abstract. *In this work, we present the development of a publish/subscribe architecture that supports mobility and disconnection of event producers and consumers effectively and in a flexible manner, as well as firewall and NAT traversal, enabling the system to be used on the entire Internet, and not only within network domains. Unlike other work about distributed pub/sub systems e.g. [Mühl, Ulbrich et al. 2004], which interprets the mobility of a device as being the re-association with a new event server (or broker) within the pub/sub overlay network, in our work we consider the device mobility as a predicted or unpredicted change of its IP address. This second form of mobility happens when the event producer/consumer switches between network domains or enters a network protected by firewalls/ NATs, while possibly still being associated with the same event broker. Hence, our notion of mobility support is orthogonal to the common notion of mobility support in distributed pub/sub architectures, and our support can be directly incorporated into distributed pub/sub systems. Thus, the main contribution of our work is this approach to mobility on pub/sub systems with the combination of different concepts and technologies (e.g. application layer mobility management, reliable protocols and disconnection detection) to achieve this goal. A central issue is to find an effective and flexible solution for mobility management [Henderson and Works 2003; Eddy 2004]. We present the development of a SIP Mobility Support Layer, an implementation of a mobility management mechanism at the application layer using SIP (Session Initiation Protocol) [Rosenberg, Schulzrinne et al. 2002] based on the work presented by [Wedlund and Schulzrinne 1999]. This layer has been implemented as an API that can be used by any service or application that requires transparent mobility and connectivity management. For addressing the problem of temporary disconnections of event producers and consumers, we introduced a customizable disconnection detection and handling mechanism, using TCP and an implementation of Reliable UDP [Stevens and Narten 1990], and also subscriptions storage and notifications queuing mechanisms. The referred pub/sub system has been implemented in Java, and the client APIs were ported to the Android Operating System.*

References

Eddy, W. (2004). "At what layer does mobility belong?" Communications Magazine, IEEE 42(10): 155-159.

Henderson, T. and B. Works (2003). "Host mobility for IP networks: a comparison." Network, IEEE 17(6): 18-26.

Mühl, G., A. Ulbrich, et al. (2004). "Disseminating Information to Mobile Clients Using Publish-Subscribe." IEEE INTERNET COMPUTING: 46-53.

Rosenberg, J., H. Schulzrinne, et al. (2002). "RFC3261: SIP: Session Initiation Protocol." RFC Editor United States.

Stevens, W. and T. Narten (1990). "Unix network programming." ACM SIGCOMM Computer Communication Review 20(2): 8-9.

Wedlund, E. and H. Schulzrinne (1999). Mobility support using SIP, ACM New York, NY, USA.

Biographies

Gustavo Baptista obtained the M.Sc. in Computer Science from Pontifícia Universidade Católica do Rio de Janeiro (PUC-Rio). Currently he is part of a research group at LAC (Laboratory for Advanced Collaboration), also at PUC-Rio, where he is a PhD candidate. His main research interests include Mobile, Ubiquitous and Cloud Computing. He has participated in projects related do middleware for distributed and context-aware mobile applications such as the MoCA (Mobile Collaboration Architecture) project. He has also participated in projects with Oi Telecom, Accenture and Tecgraf.

Markus Endler obtained the Dr. rer. nat. in Computer Science from the Technical University of Berlin, and the Professor Livre-docente title from the University of São Paulo. From 1989 to 1993 he worked as a researcher at the GMD Research Institute Karlsruhe (Germany), and from 1994 to 2000 as an Assistant Professor at the Institute of Mathematics and Statistics of the University of São Paulo. In 2001 he joined the Department of Informatics of the Pontifícia Universidade Católica in Rio de Janeiro (PUC-Rio), where he is currently Associate Professor. His main research interests include Mobile, Ubiquitous and Cloud Computing. He has supervised 8 Doctoral and 17 Master theses, and has published 90+ scientific papers.

Vagner Sacramento obtained his Doctoral degree from the Department of Informatics at the Pontifícia Universidade Católica in Rio de Janeiro (PUC-Rio) in 2006, and his M.Sc. in Computer Science from the Federal University of Rio Grande do Norte (UFRN) in 2002. He is currently adjunct professor at the Federal University of Goiás (UFG). His research interests include middleware for distributed systems and mobile computing, context-awareness, location-based services, RFID applications and sensor networks.

PrimoGENI – Developing GENI Aggregates for Real-Time Large-Scale Network Simulation

Jason Liu[1], Julio Ibarra[2], Heidi Alvarez[2]
Florida International University, Miami, Florida, USA.
[1]School of Computer Science and Information Sciences
[2]Center for Internet Augmented Research and Assessment
`liux@cs.fiu.edu, {Julio, heidi}@fiu.edu`

Abstract. *The Global Environment for Network Innovations (GENI) is an NSF-funded program to develop a nationwide research facility for developing programmable networking environments and experimenting with new network architectures and distributed systems. PrimoGENI is a project at FIU that will augment the GENI suite of interoperable infrastructure to enable large-scale experiments consisting of simulated, emulated and physical network entities. Support for large-scale simulation is currently missing in GENI. Simulation can significantly improve the flexibility and scalability of network experiments beyond what can be achieved on physical and emulation testbeds [Liu 2008a]. For example, simulation allows easier incorporation of abstract mathematical models (e.g., the epidemic worm propagation model and fluid network traffic model) and, through the use of parallel and distributed simulation, can accommodate extremely large and detailed network models. Real-time simulation refers to the technique of supporting the execution of large-scale detailed network models in real time [Liu 2008b]. We will discuss the PrimoGENI project plans to integrate simulation capabilities into the GENI "ecosystem" through real-time simulation methodologies.*

GENI is building a meso-scale prototype infrastructure that is deploying OpenFlow switches at Internet2 and National LambdaRail backbone Points of Presence (POPs), and also at least 8 U.S. university campuses. The objective of the meso-scale build out is to "get real experiments up and running (Elliot, 2010)". A GENI experiment is an interconnected set of reserved resources on platforms in diverse locations, such that each experiment is instantiated on shared infrastructure, yet runs within its own isolated slice, created end-to-end across an interoperable suite of federated infrastructure (GENI-SY-RQ, 2008). We will discuss FIU's plans to participate in the GENI meso-scale prototype deployment, which would extend the OpenFlow infrastructure to Miami. We will also discuss our plans to create an OpenFlow testbed between the AMPATH Internet Exchange Point, in Miami, and the SouthernLight Exchange Point, in São Paulo, to support simulation/emulation experiments between the U.S. and Brazil.

References

Elliot, C. (2010) "GENI Spiral 2: Exploring Networks of the Future", GENI Engineering Conference (GEC7), http://groups.geni.net/geni/wiki/Gec7Agenda

GENI. Global Environment for Network Innovations. Web site http://geni.net

GENI-SY-RQ. (2008) “GENI: System Requirements Document”, Document ID: GENI-SE-SY-RQ-01.7, The GENI Project Office BBN Technologies 10 Moulton Street Cambridge, MA 02138 USA

Liu, J. (2008a) “A primer for real-time simulation of large-scale networks”. In Proceedings of the 41st Annual Simulation Symposium (ANSS’08), pages 307–323.

Liu, J. (2008b) “Immersive real-time large-scale network simulation: a research summary”, NSFNGS Workshop, in conjunction with the 22nd IEEE International Parallel and Distributed Processing Symposium (IPDPS'08), Miami, FL, April 14-18.

Biographies

Dr. Jason Liu (PI) is an Assistant Professor at the School of Computing and Information Sciences of Florida International University (FIU). His current research focuses on fundamental technologies supporting real-time immersive large-scale network simulations. His research group designed and implemented PRIME, which is a high-performance network simulation and emulation framework. He received his B.A. degree from Beijing University of Technology in China in 1993, an M.S. degree from College of William and Mary in 2000, and a Ph.D. degree from Dartmouth College in 2003.

Mr. Julio Ibarra (Co-PI) is Assistant VP for Technology Augmented Research at FIU, where he is responsible for strategic planning and development of advanced research networking services, including the development and management of the AMPATH Internet Exchange Point for Research and Education networks. He is the PI of the Américas Lightpaths (AmLight) Project - an NSF International Research Network Connections (IRNC) grant. He has B.S. and M.S. degrees in computer science from Florida International University, and Ph.D. degree from Twente University (in progress).

Dr. Heidi Alvarez (Co-PI) is the Director of the Center for Internet Augmented Research and Assessment (CIARA) at FIU. She is responsible for supporting the center’s directive to offer Internet, Internet2 and next-generation high-performance networking services to research and education institutions. Dr. Alvarez is the PI for the CyberBridges NSF CI-TEAM demonstration and implementation projects. She received her B.S. degree from the Univesity of Miami, in Education, an M.A. degree from FIU, in English, and a Ph.D. from the Rotterdam School of Management (RSM) at Erasmus University.

Author Index