IEEE International Conference on Communications 2013: IEEE ICC'13 - Workshop on Cloud Convergence: challenges for future infrastructures and services (WCC 2013)

Experimenting Content-Centric Networks in the Future Internet Testbed Environment

Pedro Henrique V. Guimarães, Lyno Henrique G. Ferraz, João Vitor Torres,
Diogo M. F. Mattos, Andrés F. Murillo P., Martin E. Andreoni L.,
Igor D. Alvarenga, Claudia S. C. Rodrigues, Otto Carlos M. B. Duarte
Grupo de Teleinformática e Automação
Universidade Federal do Rio de Janeiro (COPPE/UFRJ)
Rio de Janeiro – RJ – Brasil
Email:{guimaraes,lyno,jvitor,menezes,afmurillo,martin,alvarenga,susie,otto}@gta.ufrj.br

***Abstract*—Future Internet Testbed with Security (FITS) is a testbed for experimenting Next-Generation Internet proposals that provides two virtualization schemes based on Xen and on OpenFlow. Experimenting new protocol proposals for the Future Internet requires a realistic condition environment for packet forwarding. FITS nodes are spread in Brazilian and European universities. In this paper, we present FITS and we use it to test a Content Centric Network (CCN), which is one of the main proposals for the Future Internet. The experiment creates a virtual network on the testbed with CCNx stack and measure the file transfer performance under real Internet traffic conditions. The results show that CCN presents an overhead of 19% when compared with the conventional TCP/IP stack. Nevertheless, CCN outperforms TCP as the number of consumers increases and CCN download time is approximately 25% smaller than TCP on the Internet.**

I. Introduction

The rapid Internet growth and access popularization has completely transformed the Internet. The original client-server communication model becomes a multimedia content distribution network and, as consequence, a new Internet is required. Future Internet architecture proposals and protocol stack experimentation demand real-scale and real-traffic testing. Network virtualization paradigm offers different virtual routers, which share a physical router, in order to simultaneously provide different network services [1]. Therefore, current TCP/IP Internet production traffic can be shared with other experimental networks. Key aspects of this network virtualization are isolation and packet forwarding performance. Isolation ensures independent virtual network operation, preventing malicious or fault virtual routers interference in the operation of other virtual networks. Privacy is also important and a virtual network can not eavesdrop another virtual network traffic.

We develop the Future Internet Testbed with Security (FITS) [2], an experimentation environment based on virtual networks that offers network isolation, secure access, and quality of service differentiation. FITS nodes are spread over Brazilian and European universities. This virtual network environment allows performance tests and evaluations of Future Internet proposals by virtualizing routers with Xen [3] and managing data flows with OpenFlow [4].

Concerning Future Internet proposals, the content distribution model considers content itself as the fundamental resource to share. Therefore, the network main service is content distribution instead of host-to-host communication. The Content-Centric Networking (CCN) [5], also known as Named Data Networking (NDN), is a network architecture that efficiently supports content distribution because the packet forwarding procedure is based on content names. This architecture splits content identification from its location, providing support to equal content name requests aggregation, copying and caching responses, balancing request-response pair flow on a hop-by-hop approach, flow multipath, signing, and ciphering content independently of its source and destination host. This paper presents Future Internet Testbed with Security (FITS) and an evaluation of Content-Centric Networks on it. Virtual routers run CCN protocol stack, based on CCNx [6] and OSPFN [7] software packages. FITS cooperating universities are interconnected to each other through the Internet.

The rest of this paper is organized as follows. FITS interuniversity testbed is discussed in Section II. The Content-Centric Network paradigm and its main components are presented in Section III. The experiments with CCNx implementation and its results are described at Section IV. Section V contains a final discussion and concludes this paper.

II. FITS Experimentation Platform

Experimentation testbeds provide researchers with isolated virtual network slices in which new proposals can be deployed. The advantage of testbeds in comparison to simulating or modeling is to perform tests in real-scale networks with real Internet traffic.

PlanetLab[1] network slicing is based on kernel virtualization, creating a user-space container for each slice, and connecting user-space containers in an overlay network. In PlanetLab, researches run the protocol proposals as applications on the user-space [8]. It is not possible, however, to run different operating systems. FEDERICA [9] is another testbed architecture that uses VLANs for Layer 2 virtualization and VMWare as machine virtualization platform. Thus, experimental environment in FEDERICA is a collection of virtual machines connected through VLAN-sliced virtual networks. OFELIA, on the other hand, focuses on OpenFlow-enabled switches for packet forwarding [10].

We propose Future Internet Testbed with Security (FITS)[2] that is based on open source software and does not require specific hardware configuration. FITS is agnostic to operating systems or applications running onto virtual machines. The

[1] http://www.planet-lab.org/.
[2] http://www.gta.ufrj.br/fits

978-1-4673-5753-1/13/$31.00 ©2013 IEEE

main strength of FITS is the flexible network virtualization, which is achieved by Xen virtual network routers or network slices using OpenFlow. FITS virtualizes Layer 2, allowing researchers to implement their own protocol stack over their own virtual Layer-2 network. Besides, FITS focuses on virtual network isolation and migration functionalities.

We used FITS to experiment the Content-Centric Network. FITS is a flexible, open, and shared platform that provides a virtual environment for experimenting innovative Future Internet network proposals.

FITS allows the creation of multiple virtual networks in parallel, based on virtualization tools Xen and OpenFlow. The testing environment is geographically distributed, with the collaboration of Brazilian and European institutions. Each institution participates with physical machines that act as nodes of the whole environment. The experimenting platform follows the pluralist approach that divides physical network in virtual networks, each containing its own protocol stack, routing rules and management. Therefore, FITS allows the creation of many isolated virtual networks, working over the same infrastructure for experimentation. The access control for virtual network management and creation uses a secure platform, based on OpenID [11] and secure microcontrollers. FITS also offers quality of service differentiation and virtual network migration features [12]

FITS implements virtual networks using virtual machines acting as virtual routers and/or flow virtualization. Each virtual network runs a different protocol stack and, therefore, physical machines host instances of these virtual networks. The FITS platform offers a web interface for management, flexible network virtualization through flow virtualization with OpenFlow [13] and machine virtualization with Xen [14]. Besides, the platform offers innovative functionalities, for example virtual network migration.

A. Virtual Networking

FITS architecture offers functionalities for virtual networking and for easing physical node insertion. FITS does not require specific hardware configuration. Commodity computers running Linux Operating System are physical nodes of the testbed.

FITS allows virtual network creation and migration through web services. It also provides an isolated environment for each virtual network. Xen machine virtualization guarantees CPU and memory isolation. Network isolation is achieved by two mechanisms. The first mechanism uses VLAN tag to isolate virtual network addressing space, whereas the second mechanism controls network resources and QoS by forwarding packets of each virtual network through dedicated queues with limited bandwidth. This isolation scheme is performed at the Xen control domain, called Domain 0, by controlling queues at virtual machine interfaces.

B. FITS Platform Architecture

The experimentation platform allows the participating institutions to deploy islands. Each island contains its own policies for experimentation. FITS nodes act as physical substrate for virtual network formation, in which islands are connected to each other through Generic Routing Encapsulation (GRE) tunnels and Virtual Private Networks (VPNs) to emulate virtual layer 2 links on the Internet.

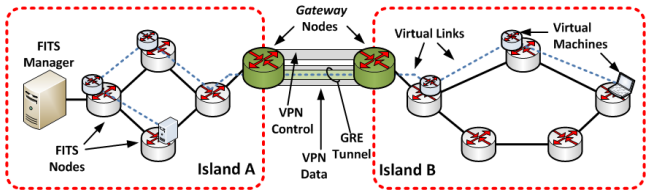

Figure 1. Connection between FITS islands. Gateway Nodes connect islands through two VPNs. GRE tunnels, inside a data-dedicated VPN, create a unique Ethernet diffusion domain and a unique Link Layer between physical nodes. Control messages are exchanged between FITS Nodes and FITS Manager. FITS Manager provides management service through a web interface.

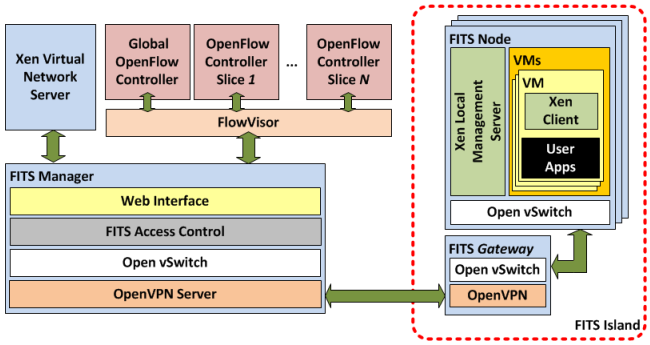

Figure 2. FITS experimentation platform. The FITS Manager controls FITS Nodes and virtual networks through secure connections between the Gateway Nodes.

There are three different types of physical nodes in FITS: FITS Manager, Connection Gateways and FITS Nodes. Figure 1 shows the network infrastructure and the role of each FITS Node. The Connection Gateway is a special node, which provides interconnection between islands by creating communication tunnels with other island Gateways. The FITS Manager coordinates the platform operations, provides user and node authentication, allocates virtual resources, creates and migrates virtual networks, and collects statistics from all FITS islands.

Figure 2 shows main FITS service components. FITS Gateway Nodes provide communication between FITS islands and Xen Virtual Network Server and Global OpenFlow Controller. The Xen Virtual Network Server creates and manages virtual networks, and connects to each Xen Local Management Server in FITS Nodes. Xen Virtual Network Server sends commands and collects usage information from Xen Local Management Servers. Xen Local Management Server connects with each Xen Client, inside the virtual machines, to retrieve information about virtual resources usage. The Global OpenFlow Controller has a global view of the OpenFlow network and connects to the FlowVisor [15]. The FlowVisor acts as an interface between OpenFlow switches and OpenFlow network and is responsible for slicing OpenFlow network into different controllers.

The FITS network operation is based on Open vSwitch [16], a software switch which implements the OpenFlow API. The Open vSwitch on FITS Nodes are connected to the OpenFlow Global Controllers that manages the platform network and virtual links in each virtual network. The OpenFlow Global Controller is the responsible for forwarding packets and for virtual network isolation.

III. Content-Centric Network

The Content-Centric Network paradigm is an alternative for the current Internet Protocol (IP), which splits location and content addressing. In this new paradigm, communication is oriented to content instead of host location. The communication is done with two types of packets: Interest and Content. Interest packets notify requests for a particular content in the network. Content packet is sent in response to an Interest packet, transmitting the desired content or part of it (chunks, as it is called in CCN). The content is sent by its original producer or by the nearest repository that contains it, for instance, a CCN router. This approach enhances mobility and multiple sources and multiple destinations communication.

A. Content-Centric Network Forwarding

The CCN model is based on Interest and Content packet parity, in which a consumer sends Interest packet whenever he desires information. Network forwards Interest to producers that reply with Content packets, traveling the reverse path of its Interest. The CCN router has three main data structures for packet forwarding: Content Store (CS), Pending Interest Table (PIT) and Forwarding Information Base (FIB). The CS repository stores temporarily the searched content, allowing local response to repeated requests. The router replaces stored contents for more relevant items using strategies like Least-Recently-Used (LRU) or Least-Frequently-Used (LFU). The FIB table is similar to IP router FIB, storing prefix-based rules, a map of the output interfaces and more specific name prefix patterns. CCN, however, allows FIB to map a prefix pattern to a list of output interfaces ordered by priority. The Pending Interest Table stores Interest packets, input and output interfaces by which the packet was forwarded but yet not replied [17].

The CCN forwarding also comprises an adaptive procedure to choose the best content forwarding path. This procedure has several mechanisms, for example: i) Interest packet time calculation between dispatch and arrival, ordering interfaces with the best metrics, ii) congestion control limiting the number of simultaneous pending Interests, iii) traffic limitation [18].

IV. Experimental Results

We implemented a CCN network as a virtual network in FITS. Using the virtualized CCN network, we measured file download time and analyzed the causes of download delays.

The CCN prototype was implemented in a virtual network over FITS (Future Internet Testbed with Security). The CCN stack was implemented using CCNx [6], version 0.7.1rc1, with OSPFN (Open Shortest Path First for Named Data Networking) [7], version 2.0. The OSPFN extends routing calculation for data names, since OSPFN disseminates data name adjacencies. The OSPFN runs over IP and calculates CCNx network prefix routes in a distributed way based on name adjacencies.

The experimental network proposed by Jacobson *et al.* is extended in this paper over the Internet using a virtual network running CCNx stack [5]. The tests allow to compare the performance of TCP and CCN when multiple consumers are downloading the same content across the Internet, at the same time from the same content producer. For the TCP, the file download uses a HTTP client, *wget*[3].

The content producer (or HTTP server in the TCP/IP network) and consumers (or clients) are connected through a single virtual router. There are 12 virtual machines acting as consumers and downloading the same 6 MB file from one content producer. The virtual network runs CCN and TCP/IP stacks. The experiment is divided into two different scenarios: first, all 14 virtual machines run inside the same physical node. This scenario was conceived to repeat similar conditions like those seen in Jacobson *et al.* [5]. The second scenario is composed by two physical nodes, which virtual machines are connected across the Internet. In this scenario, the content producer and the router are connected by a virtual link through the Internet. Therefore, bandwidth limitations between content producer and router are caused by real traffic conditions between two physical nodes. One physical node hosts the virtual machine acting as the content producer, while the second one hosts the router and the consumers. Experimental network scheme is presented in Figure 3.

The test is run by the control node that executes a measurement *script*, responsible for starting the tests. The *script* was written in Python and prepares the file to be downloaded using CCN, through CCNx implementation, by using *ccnputfile* command. As soon as the CCN experiment is over, the test is repeated using *wget*.

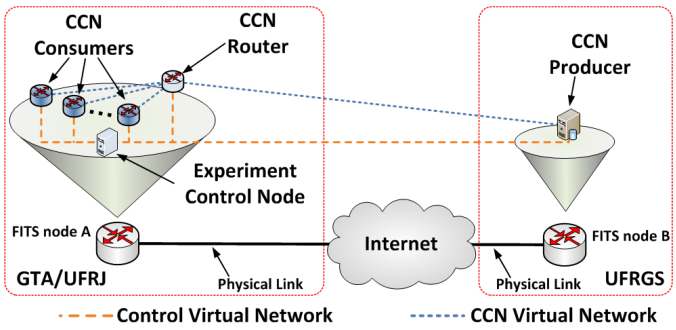

Figure 3. Virtual networks used on the experiments. Two virtual networks were created, the experimentatal network and the control network. The control network is an isolated network that sends execution commands to the experimentatal network elements. The experiments use the communication provided by experimentatal network.

In the first scenario, all the 14 virtual machines run in the same physical node and bandwidth between the virtual router and the content is limited. In the second scenario, two physical machines are chosen in FITS and they are placed on different islands and the connection limitation is due to traffic conditions on the Internet that degrades the bandwidth.

Figure 4 shows the download time of a 6 MB file as a function of the number of consumers requesting the file at the same time. Results seen in Figure 4(a) shows that when all virtual machines are in the same physical node and the content producer bandwidth is limited, CCN outperforms TCP after two consumers downloading the file at the same time. The second test scenario is executed over the Internet with consumers connected to a router, each one using a rapid connection, while the router is connected to a content producer in another island. Connection between islands are subjected

[3] http://www.gnu.org/software/wget/.

to real-traffic limitations. When the number of consumers increases, Internet traffic conditions degrades the bandwidth as on the first scenario. The TCP connection is limited by the bandwidth of the virtual link between the content producer and the router, whereas on CCN, downloading the file means downloading the file to router cache and distributing the file between consumers. This allows bandwidth usage optimization in the available connections and acceleration of content distribution process.

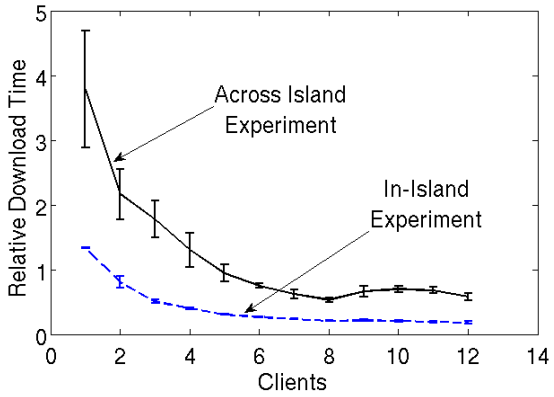

(a) All virtual machines are on a single physical node.

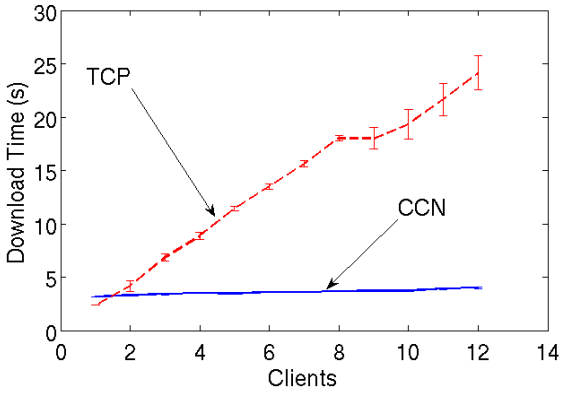

(b) Comparison between downloading a file in all virtual machines over the same physical machine (`In-Island`) and over two spread physical machines through Internet (`Across Island`).

Figure 4. Evaluation of the downloading time of 6 MB file in CCN vs. TCP scenarios.

Figure 4(a) presents results of download time when all virtual machines of the experiment are inside the same physical node. The virtual link between the content producer and the virtual router is limited to 10 Mb/s. As expected, results are close to the one seen in Jacobson *et al.* [5]. This first scenario shows the proposed platform, FITS, as a valid environment for testing the CCN proposal. As the number of consumers increases, the bandwidth limitation degrades TCP performance. This happens because, in TCP, each download shares the available bandwidth with the others. CCN, however, is more robust to bandwidth limitations. CCN uses the available bandwidth more efficiently because the router forwards only one request per content demand, but replies the content received to all interfaces that it received requests. TCP routers forwards all requests for certain IP and the HTTP server replies to each request while CCN routers forwards only one Interest for certain content, no matter how many Interests arrive after and the data producer replies with only one Content packet.

Figure 4(a) compares results seen in the first testing scenario (in-island test). Figure 4(b) compares the ratio of CCN/TCP download time `in-island` and `across-island` scenarios. The main conclusion is that with the Internet traffic conditions, CCN follow the same pattern seen in the first scenario, Figure 4(a), and CCN stack outperforms TCP as the number of consumers increases.

The second experiment consists of capturing transmitted packets between consumers and content producer in the TCP/IP and CCN. This experiment is important for the comprehension of the CCNx protocol stack implementation and the results of the first experiment. The experiment was done by capturing packets from a consumer while it consumes (or downloads) data of a content (file) in CCN or TCP/IP. For packet capture the *tcpdump*[4] was used. During the experiment, a 20 MB file was downloaded. In this way, the average transfer rate was evaluated for each scenario. The average transfer rate in the CCN scenario was 1.7 MB/s while for the TCP was about 11.8 MB/s. This difference in transfer rate is the reflex of the implementation of the two stacks. TCP is implemented in an optimized way, directly in the Operating System kernel, while CCN stack, in special the CCNx implementation is a user space application written in Java [5]. The implementation of the CCNx router, upon arrival of an Interest, passes through to the next link in the network. At a data arrival, however, the first copy of the Content package is stored in the local cache of the router and, after, is resent to the consumer that requested it. Another important fact to understand the CCNx smaller transfer rate is that CCNx 0.7.0rc1 implementation uses a default chunk size of 4 KB. This default size is defined in the *ccnputfile*, the application responsible for content distribution. The chunk size is important, because it is the minimum data transfer unit of the CCN. In the experiment we see that CCNx is based on UDP transport protocol, therefore, there is no connection between consumers and the content producer, thus the chunks are sent as UDP datagrams. Another point is that each chunk is mapped in a UDP datagram, hence, UDP datagrams have 4 KB of content. As the *Maximum Transmission Unit* (MTU) of the virtual network is 1500B, all transmitted UDP datagrams are fragmented, generating a higher delay for datagram reconstruction before sending it to the application.

Figure 5 highlights the file transfer behavior in CCN and in TCP according to time. As before mentioned, TCP presents a higher file transfer rate, ending the transfer in 13 s. CCN download, however, took 143 s. The longer time interval spent by CCN is due to smaller transfer rate of the CCNx implementation. Another point in Figure 5 is that CCN presents an initial delay of 43 s. This delay is due to downloaded file is not initially inside the nearest router cache and, thus, it is necessary that the router treats the incoming Interest of the consumers for each chunk and resends it to the content producer that has the content. After the time needed to download the file to the router, CCN file transfer rate is smaller than the TCP rate, since CCN demands the fragmentation of 4 KB datagrams, and demands data storage in the intermediary router before resending to the consumers. CCNx also counts upon a user space implementation, while TCP/IP forwarding scheme is optimized directly by the kernel. At the end, Figure 5 still

[4]http://www.tcpdump.org/.

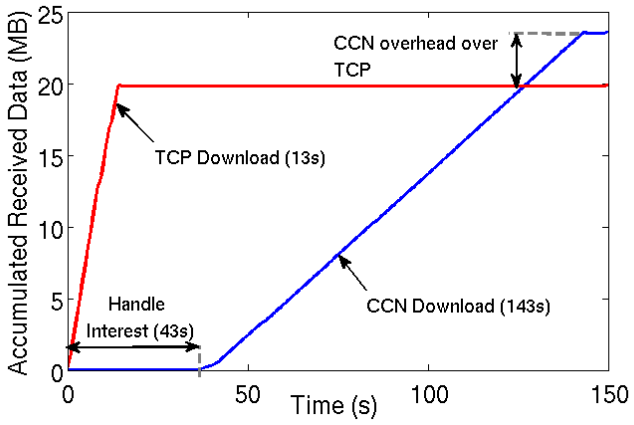

Figure 5. File download duration using a TCP application, compared with duration for downloading the same file using the CCNx stack. Original file size is 20 MB.

highlights that the CCNx implementation introduces a header overload of 19% in comparison to the same content in the HTTP application over TCP/IP stack.

V. Conclusion

This paper presented Future Internet Testbed with Security (FITS), a testbed environment for Future Internet protocols, it also presented experimental results of a real implementation of a Content-Centric Network. FITS is a collaboration of universities to test Future Internet proposals. This environment allows the creation of isolated virtual networks with secure access, Quality of Service differentiation, and virtual network migration features.

Content-Centric Network (CCN) is pointed out as one of the most viable Future Internet proposals. Nowadays, the protocol stack implementation for Content-Centric Network is the CCNx. The presented experiments compare the CCNx stack with the conventional TCP/IP protocol stack.

The results show that CCNx, when compared with TCP/IP stack, presents an overhead of 19%. Nevertheless, CCN outperforms TCP as the number of consumers increases and CCN download time is approximately 25% smaller than TCP when working with 12 consumers downloading content from another FITS island across the Internet. As future works we will study and develop new routing mechanisms for CCN and experiment new Future Internet proposals on the FITS.

Acknowledgments

The authors would like to thank FINEP, FUNTEL, CNPq, CAPES, FAPERJ, and UOL for their financial support. The authors would also like to thank the institutions that collaborate with FITS test platform. A special acknowledgment to Professors Luciano Gaspary and Marinho Barcellos and researchers Éderson Vieira and Lucas Muller of UFRGS; Prof. Igor Moraes of UFF; Prof. Marcelo Rubistein and researcher Leopoldo Mauricio of UERJ; Profs. Edmundo Madeira and Nelson Fonseca and researcher Estéban Rodriguez of UNICAMP; Prof. Artur Ziviani of LNCC; Prof. André dos Santos and researcher Edgar Tarton of UECE; Prof. Djamel Sadok and researcher Marcelo Santos of UFPE; Prof. Cesar Marcondes of UFSCar; Prof. Paulo Veríssimo and researchers Oleksandr Malichevskyy and Diego Kreutz of Universidade de Lisboa, which were essential for the installation and subsequent experiments was possible. Authors would like to thank GTA/UFRJ's team that, in the last three years, worked on virtualization and made the experimentations possible.

References

[1] N. Fernandes, M. Moreira, I. Moraes, L. Ferraz, R. Couto, H. Carvalho, M. Campista, L. Costa, and O. Duarte, "Virtual networks: isolation, performance, and trends," *Annals of Telecommunications*, vol. 66, pp. 339–355, 2011.

[2] FITS, "Future Internet Testbed with Security," 2012. [Online]. Available: http://www.gta.ufrj.br/fits/

[3] P. S. Pisa, R. S. Couto, H. E. T. Carvalho, D. J. S. Neto, N. C. Fernandes, M. E. M. Campista, L. H. M. K. Costa, O. C. M. B. Duarte, and G. Pujolle, "VNEXT: Virtual NEtwork management for Xen-based Testbeds," in *2011 International Conference on the Network of the Future (NoF'11)*, Paris, France, Nov. 2011, pp. 41–45.

[4] D. Mattos, N. Fernandes, V. da Costa, L. Cardoso, M. Campista, L. Costa, and O. Duarte, "OMNI: OpenFlow MaNagement Infrastructure," in *Network of the Future (NOF), 2011 International Conference on the*, nov. 2011, pp. 52 –56.

[5] V. Jacobson, D. Smetters, J. Thornton, M. Plass, N. Briggs, and R. Braynard, "Networking named content," in *Proceedings of the 5th international conference on Emerging networking experiments and technologies*. ACM, 2009, pp. 1–12.

[6] "CCNx Project," 2011. [Online]. Available: http://www.ccnx.org/

[7] L. Wang, A. Hoque, C. Yi, A. Alyyan, and B. Zhang, "OSPFN: An OSPF Based Routing Protocol for Named Data Networking," University of Memphis and University of Arizona, Tech. Rep., Jul. 2012.

[8] B. Chun, D. Culler, T. Roscoe, A. Bavier, L. Peterson, M. Wawrzoniak, and M. Bowman, "PlanetLab: an overlay testbed for broad-coverage services," *ACM SIGCOMM Computer Communication Review*, vol. 33, no. 3, pp. 3–12, 2003.

[9] P. Szegedi, S. Figuerola, M. Campanella, V. Maglaris, and C. Cervelló-Pastor, "With evolution for revolution: Managing FEDERICA for future internet research," *Communications Magazine, IEEE*, vol. 47, no. 7, pp. 34–39, 2009.

[10] A. Köpsel and H. Woesner, "OFELIA–pan-european test facility for openflow experimentation," *Towards a Service-Based Internet*, pp. 311–312, 2011.

[11] D. Recordon and D. Reed, "OpenID 2.0: A platform for user-centric identity management," in *ACM Workshop on Digital Identity Management (DIM)*, 2006, pp. 11–16.

[12] P. S. Pisa, N. C. Fernandes, H. E. T. Carvalho, M. D. D. Moreira, M. E. M. Campista, L. H. M. K. Costa, and O. C. M. B. Duarte, "OpenFlow and Xen-based virtual network migration," in *Communications: Wireless in Developing Countries and Networks of the Future*. Springer Boston, 2010, vol. 327, pp. 170–181.

[13] N. McKeown, T. Anderson, H. Balakrishnan, G. Parulkar, L. Peterson, J. Rexford, S., and J. Turner, "OpenFlow: Enabling Innovation in Campus Networks," *ACM SIGCOMM Computer Communication Review*, vol. 38, no. 2, pp. 69–74, Apr. 2008.

[14] N. Egi, A. Greenhalgh, M. Handley, M. Hoerdt, F. Huici, and L. Mathy, "Towards high performance virtual routers on commodity hardware," in *Proceedings of the 2008 ACM CoNEXT Conference*. ACM, 2008, p. 20.

[15] R. Sherwood, G. Gibb, K.-K. Yap, G. Appenzeller, M. Casado, N. McKeown, and G. Parulkar, "Can the production network be the testbed?" in *USENIX Conference on Operating Systems Design and Implementation (OSDI)*, 2010, pp. 1–6.

[16] B. Pfaff, J. Pettit, T. Koponen, K. Amidon, M. Casado, and S. Shenker, "Extending networking into the virtualization layer," in *ACM Workshop on Hot Topics in Networks*, Oct. 2009.

[17] L. Z. *et al.*, "Named Data Networking (NDN) Project," Tech. Rep., Oct. 2010.

[18] C. Yia, A. Afanasyevb, I. Moiseenkob, L. Wangc, B. Zhanga, and L. Zhangb, "A Case for Stateful Forwarding Plane," University of Arizona, University of California, Los Angeles and University of Memphis, Tech. Rep., 2012.