Computer Networks 63 (2014) 33–47

Contents lists available at ScienceDirect

Computer Networks

journal homepage: www.elsevier.com/locate/comnet

Evaluation of tiered routing protocol in floating cloud tiered internet architecture

Yoshihiro Nozaki [a,*], Parth Bakshi [a], Hasan Tuncer [a], Nirmala Shenoy [b]

[a] Golisano College of Computing and Information Sciences, Rochester Institute of Technology, Rochester, NY 14623, USA
[b] Department of Networking, Security, and Systems Administration, Rochester Institute of Technology, Rochester, NY 14623, USA

ARTICLE INFO

Article history:
Received 2 July 2012
Received in revised form 10 September 2013
Accepted 7 November 2013
Available online 22 November 2013

Keywords:
Internet architecture
Routing
Testbed

ABSTRACT

Clean slate future Internet initiatives have been ongoing for a few years. An important consideration in the eventual deployment of solutions for such Internet architectures is the testing and validation of the design and its scalability in realistic network environments. Large scale emulation and experimentation testbeds sponsored and funded by major research organizations worldwide provide a suitable platform for the purpose. In this article, we present the implementation details of a new network and routing protocol that entirely replaces IP and its routing protocols from the protocol stack to provide efficient routing and forwarding of packets in a clean slate Floating Cloud Tiered (FCT) Internet architecture. The FCT architecture leverages the tier structure existing among ISPs, and has a new addressing and routing schema based on tiers. In this article, the implementation and evaluation details of the network protocol with these two features, namely the tiered addressing and tier-based routing using the Global Environmental for Network Innovations (GENI) testbed are presented. The performance of the protocol is also compared with Open Shortest Path First (OSPF) implemented over the GENI testbed for identical network topologies.

1. Introduction

The Internet has evolved massively and impressively since its inception several decades ago. The TCP/IP protocol suite designed for the Internet has sustained tremendous growth despite the fact that it was designed for technologies of yesteryear. The sustenance is also attributed to the evolutionary research directed to overcome the vulnerabilities and limitations of the protocol suite as they surfaced. However, this trend of resolving vulnerabilities has resulted in patchwork point-solutions [1], and an Internet that is highly complex in operation, difficult to manage, more vulnerable to emerging threats, and highly brittle and fragile in the face of new developments and requirements.

The alarming trends in the current Internet evolution eventually led to several clean slate Future Internet initiatives around the world such as the Future Internet Design (FIND) [2] and Future Internet Architecture Project [3] by the National Science Foundation (NSF) in the United States, the Seventh Framework Program (FP7) [4] by the European Union, AKARI project [5] by Japan, and the *12th Five-Year Plan* projects [6] by the Ministry of Science and Technology (MOST) in China. These programs support research efforts that target challenges such as routing, scalability, mobility, security, and reliability among others – towards an ideal future Internet architecture. An important consideration in the design of Internet architectures are testing and validation of the design and scalability using realistic network scenarios in a near realistic experimental setup. The Autonomous Systems (ASs) and Internet Service Providers (ISPs) that construct the current Internet would not be willing to expose their networks to the risk of such experimentation, nor would they be willing to reveal information of their

* Corresponding author. Tel.: +1 (585) 475 4887.
E-mail addresses: yxn4279@rit.edu (Y. Nozaki), pab8754@rit.edu (P. Bakshi), hxt3948@rit.edu (H. Tuncer), nxsvks@rit.edu (N. Shenoy).

1389-1286/$ - see front matter © 2013 Elsevier B.V. All rights reserved.
http://dx.doi.org/10.1016/j.bjp.2013.11.010

internal network topologies and implementations. Research communities have hence implemented open virtual large-scale testbeds using virtualization technologies. Large scale emulation and experimentation testbeds for this purpose are another effort sponsored and funded by major research organizations in the world; these include Global Environment for Network Innovations (GENI) [7] by NSF in the United States, the Future Internet Research and Experimentation (FIRE) project [8] a part of FP7 in the European Union, the Japan Gigabit Network 2 Plus (JGN2plus) [9] and the China Next Generation Internet (CNGI) [10] testbeds in Asia. The research project described in this article was funded under the NSF FIND program and tested and evaluated on Emulab, which is part of GENI. The major goals of this project were (i) to address the two related problems of routing scalability and addressing, through an approach which would leverage the existing structures in the current Internet architecture, (ii) to propose a solution that is acceptable to the ISP community that supports the Internet and lastly (iii) to provide a transition platform and mechanism which is essential to the ultimate successful deployment of the proposed design. Given the goals, it was decided to explicitly use the tiered relationships adopted by ISPs in their business models. The tiered relationships are based on a tiered structure that has the benefits of both hierarchical and distributed architectures. Each tier can have clear and well defined functionalities to incorporate and improve manageability and controllability, while still availing services from an upper tier and providing service to entities in a lower tier. Fault isolation can be improved as the structure allows for locating possible failure points due to the inherent structural properties of tiers [11].

Huston [12] observed the growth of BGP (Border Gateway Protocol) routing table sizes for years and showed that one of the main contributions for the increasing the table size is due to an increasing number of multi-homed ASes. Because many of the ASes are moving from a single-homed connection to multihoming and peering, the BGP table size has rapidly increased as the result of an increasingly dense interconnected AS mesh at the edge of the Internet. Furthermore, logical links achieved by Multiprotocol Label Switching (MPLS) technology have introduced meshed topologies within an ISP. The Level 3's router topology presented in [13] is highly meshed because of this. Flat and highly meshed network structures provide high redundancy, which comes at the cost of reduced efficiency as more and more complex routers are necessary to discover and maintain routes as the network grows in size, a fact that is apparently alarming when one notices the processing and operational conditions of core routers today [14]. Routing loops, looping packets, and high network convergence times are also the costs attributed to meshed network structures. Scalability is difficult to achieve under these conditions. Meshed networks are also hard to upgrade, troubleshoot, and optimize unless they are designed using a simple and hierarchical model [11].

In comparison, the tiered network structure described in this project adopts a tier-based routing and forwarding and a suitably designed tiered addressing scheme that reduces route maintenance by several magnitudes, as is evident from the results discussed in Section 6. The tiered structure does not cancel the benefits of the underlying meshed connectivity, as they still continue to exist, operationally, optimally combining hierarchy and meshing. With the modularity introduced through the concept of network clouds and nesting, the structure affords a high level of scalability [15]. The tier concept is common among ISPs, but it can also be noted within an ISP network which comprises several Point of Presences (POPs), inside which three tiers can be identified; tier 1 comprises backbone routers, tier 2 comprises distribution routers, and tier 3 comprises access routers. In the proposed design, each set of routers is identified as a network cloud and then associated to a tier. The design takes into consideration the eventual transition and deployment through MPLS, where differentiation of tiers and tier-based forwarding can be achieved through labels and label stacking respectively.

In this article, a description of the proposed future Internet architecture, its main features, namely the tiered addressing scheme and routing, and their attributes are presented. A Linux based router is implemented to support the tiered addressing and routing in a way which operates just above layer 2, and bypasses the IP layer and the IP routing protocols. The routers are deployed on the GENI testbed to run the routing protocol on suitably designed topologies. Scripts were written to setup the test case scenarios, to generate traffic, and to collect data to validate and evaluate the performance of the proposed architecture on the GENI testbed. In addition to the detailed testbed implementation, the performance improvements achieved by the routing protocol under the proposed architecture is compared with Open Shortest Path First (OSPF).

This article is organized as follows. In Section 2, the related future Internet architectures are presented. Section 3 covers the main concepts of the proposed future Internet architecture, the comparison of the new tiered addressing scheme to IP addressing, and the comparison of the new tiered routing scheme to IP based routing. The detailed operations of the proposed tiered routing protocol in an intra-domain environment and OSPF routing protocol are discussed in Section 4. Implementation and evaluation methods used for the testbed are presented in Section 5. In Section 6, we analyze the performance of the proposed tiered routing protocol compared to OSPF. Further discussions of the capabilities of TRP are provided in Section 7. Concluding remarks are presented in Section 8, followed by Acknowledgements and References.

2. Related work

A significant amount of research has been conducted towards resolving different issues prevalent in the current Internet. Those closely related to our work in terms of the targeted goals and articles that demonstrate a variety of approaches towards bringing solution to different aspects of the Internet's scalability problem are included in this section. They are presented under two subsections: