IPv6 Addresses as Content Names in Information-Centric Networking

Suman Srinivasan, Henning Schulzrinne
Columbia University
{sumans,hgs}@cs.columbia.edu

Abstract

Content is quickly beginning to emerge as the core of Internet and networking applications today. Among the most important research issues with content is the problem of addressing and naming content, since a robust and naming-centric networking strategy will enable the building of next-generation Internet architectures that can easily scale content demands correctly. We propose a counter-intuitive approach to solving the naming problem, by using IPv6 addresses as content names. We explain our proposal and architecture for using IPv6 addresses for content names, and argue that using IPv6 addresses for naming content will allow us to solve the problems of routing and directory services associated with naming.

1. Introduction

Content is quickly becoming the core feature of the Internet. However, the foundations of the Internet and the various protocols that run on top of it were mostly built several decades ago and are host-based. Several research projects attempt to address this. Naming schemes, such as i3 [1], attempt to solve the content problem by looking at the aspect of naming. Content-centric networking, such as CCNx [2] and XIA [3], aim to replace the IP-based Internet stack with one based on content and content names.

We propose solving the content issue through a counter-intuitive proposal: using IPv6 addresses for content names. Using IPv6 addresses for content names solves the content networking problem, and at the same time, IPv6 provides an extensive architecture for handling issues related to routing, security, etc.. In other words, we propose solving the content networking problem by mapping content names to a resource that addresses network problems comprehensively: IPv6 addresses.

2. Why IPv6 for Content Names?

At first, the idea of using IPv6 addresses as content names might seem counter-intuitive. After all, aren't IPv6 addresses meant to represent host names? We argue that while IPv6 addresses are meant to represent host names, there is nothing in the networking mechanism that limits them to being used as host names. Through a simple lookup mechanism (proposed in the architecture), we can achieve this.

In addition to the basic implementation, the use of IPv6 addresses as content names also inherently brings with it the many advantages of IPv6 addresses to content. The amount of work devoted to IPv6 means that we get these features for "free" in any ICN/CCN implementation based on IPv6 content names! While the space for this abstract does not allow us to expand on these features of IPv6 as they pertain to content naming, we hope to expand on it for a full paper.

3. Mechanism/Architecture

While we are not able to arbitrarily assign IPv6 addresses for our content, we can use the prefixes of IPv6 addresses already assigned to content publishers as the prefix for our new set of IPv6 content names. We segment the IPv6 address so as to be able to differentiate between the publisher and the content name.

We implement the mapping of content names to IPv6 addresses as follows: we parse the content name into publisher name and content name. A lookup of the existing IPv6 address space for the content publisher is done, and that is used as a prefix for the IPv6 address. For the second 64-bits, an MD5 hash of the content name is done, and the first 64-bits of the hash are converted into the last "quads" of the IPv6 address. Once a content name request is made and published, it is registered with a central registry so that a reverse lookup of the IPv6 address can be easily performed. More details about the architecture are in the poster slide.

4. Conclusion

In this poster, we described our proposal for using IPv6 addresses for content naming. We also presented our initial implementation. We believe that using IPv6 addresses for naming content will allow for a robust naming scheme for the future Internet, while addressing the fundamental issues of caching and routing.

5. References

[1] J. Kannan, A. K. K. Lakshminarayanan, I. Stoica, K. Wehrle, "Supporting Legacy Applications Over i3", IEEE Computer 2004
[2] V. Jacobson, D. Smetters, J. Thornton, M. Plass, N. Briggs, R. L. Braynard, "Networking named content", Proc. of CoNEXT '09, Rome, Italy
[3] eXpressive Internet Architecture (XIA), http://www.cs.cmu.edu/~xia/, retreived December 17, 2010

IPv6 Addresses as Content Names in Information-Centric Networking

Suman Srinivasan, Henning Schulzrinne
Columbia University

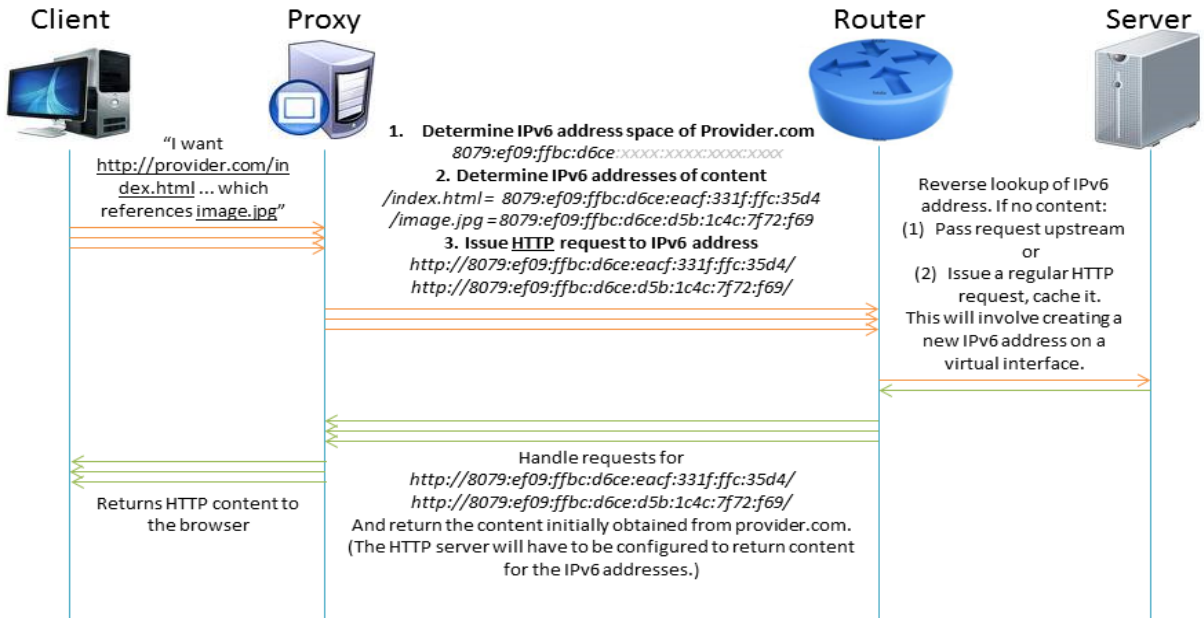

The architecture diagram of our IPv6 content addressing system. In our system, the regular browser makes a HTTP request through a proxy, which translates HTTP requests to an IPv6 content addressing system. The request is sent out over the network, until a router on path that has the content responds to the request. The proxy then translates the retrieved content back into a HTTP response to the user's browser.

IPv6 Features That Are Useful for Content Networking

- IPSec: security
- Multicasting: streaming video
- No fragmentation
- Better mobile support
- Jumbograms
- IPv6-over-IPv4 bridging mechanisms: use of our IPv6 content naming proposal in archaic or today's networks

Sample Name to IPv6 Mappings

- **8079:1b37:2650:3af8:***1d78:a723:dee0:2522*
  http://nytimes.com/content/video/file.mp4
- **8079:1b37:2650:3af8:***eacf:331f:ffc:35d4*
  http://nytimes.com/index.html

Currently implemented

- Content address registry as a web service (built in PHP) connected to a MySQL database.
- Requests to set/get content names and their corresponding IPv6 address mapping are done through simple put and get requests. We do plan to make this more scalable and hierarchical in the future.

Current/Future Work

- Starting work on the full implementation of the IPv6 content naming architecture.
- Use netfilter, particularly libnetfilter_queue and its Python language bindings, to handle and serve IPv6 content naming and addresses.