The Tiger Meets the Termite Mound: Toward Robust Inter-Domain SDN

John Wroclawski – jtw@isi.edu

To date, SDN development has focused primarily on creating rich function and management clarity, in environments fundamentally characterized by centralized logical control. From early roots in the research community to today's commercially deployed OpenFlow use cases, the principal focus has been technical implementation of extended functionality and separation of control and data path, in the context of executing policies and functions determined within a single organization.

In stark contrast, the successful creation and sustenance of large-scale networks such as the Internet is *at its heart* a distributed coordination problem. That is, successful coordination across the competing interests and objectives of multiple organizations, operators, providers, and users is the fundamental challenge: more important than pure function, because it is the entire basis of that function. It should be noted that the Internet is not unique in this respect. The same basic problem arises in large decentralized enterprises, multi-agency government networks, cross-organization ad-hoc networks set up in times of emergency, and nearly all other decentralized settings.

Consequently, perhaps the most basic long-term research challenge facing "Interdomain SDN" is the creation of architecture and structure that supports successful, sustainable distributed coordination in this new environment.

To explore this problem in a bit more detail, we consider some lessons from the Internet, asking why it succeeded where competing network architectures did not. While any number of technical and non-technical factors played a role, three high-level ideas can be singled out:

- A narrow narrow waist: in contrast to other then-competing architectures, the IP protocol designers intentionally specified a single, extremely simple "narrow waist" service model and implementing protocol to serve as the universal spanning layer in the stack. By providing this simple, crisply defined fixed point in the architecture, the IP protocol enabled trivial interconnection and traffic flow across domains, as well as 30 years of evolution below, in communications technologies, and above, in areas such as routing, services and applications.
- Information hiding: as a consequence of the simple service model and IP's multi-level routing approach, the architecture strongly minimizes the information that must be shared between domains, allowing the vast majority of information about a domain's capabilities, design and operation to be kept internal for engineering or competitive reasons.
- Economic simplicity: again as a consequence of the simple, narrow-waist service model and other technical factors, and in stark contrast to some then-alternative architectures, the Internet was able to thrive for many years with a very simple economic model for traffic peering, inter-domain interconnect, and end user service functions. While this model is now evolving towards somewhat more complexity, the ability to provide economic simplicity and low barrier to entry for a wide range of parties was crucial to IP's early commercial success.

A principal challenge facing Interdomain SDN, then, is that of applying these lessons to an emerging new environment where high function, individualized configuration, and dynamic programmability are perceived as central capabilities. The crucial observation is that, at first glance, these SDN goals and characteristics are in more or less direct conflict with the architectural concepts and lessons learned above.

Consequently, we are forced to consider at least three distinct paths to resolving this conflict. First, we may attempt to *balance* the conflicting forces inherent in this tension to achieve the best long-term result, which ideally requires gaining sufficient understanding of the strength of each force in the competition; second, we may attempt to *develop new, higher-level abstractions* that can provide the same architectural benefits to SDN that IP's concrete simplicities provided to the Internet; and third, we may attempt to *identify entirely new approaches to distributed coordination* that produce the desired results in the context of SDN's rich and dynamic functionality. Each of these threads offers a promising and exciting line of future research, with many possibilities for distinct and focused concrete efforts.

Bio: John Wroclawski is Director of the Internet and Networked Systems Division of the University of Southern California's Information Sciences Institute, with responsibility for the strategic direction of this 40-person research organization. His personal research interests include the architecture, technology and protocols of large, decentralized communication systems such as the Internet, and architectural aspects of cyber-physical systems. At ISI, he also serves as chief scientist of the DETER Cybersecurity testbed, a major DHS and NSF funded project aimed at transforming the effectiveness, rigor and scientific merit of experimental cybersecurity research through the development of new methodologies and infrastructures for the field. In the broader community, he served from its inception on the planning group for NSF's GENI project, co-chaired from 2006-2008 the working group charged with overall technical architecture of the GENI facility, and continues as an interested observer of the program.