A Next Generation Internet State Management Framework

Xin Huang and Tilman Wolf
Department of Electrical and Computer Engineering
University of Massachusetts at Amherst
{xhuang,wolf}@ecs.umass.edu

ABSTRACT

In the next generation Internet architecture, more functionality in the form of "services" and packet processing features will be placed inside the network rather than on end-systems. This trend will increase the amount of network state information maintained by the systems inside the network. We present a framework for structuring and managing network state information that can be used in the next generation Internet architecture.

1. INTRODUCTION

The current Internet architecture can not be easily extended to accommodate new requirements from diverse applications, heterogenous network paradigms (e.g., sensor network, mobile network, etc.), and security and manageability demands. The networking community is in the process of considering a complete redesign of the Internet architecture [1] and a major trend among network architecture designs is to place more functionality and network state information inside the network [2, 3]. Instead of constraining routers to simply forward packets and leaving all processing tasks to the end-systems, programmable processing components (e.g., network processors) on routers are used to implement complex packet processing services.

The trend towards computation on routers is complemented by an increasing number of stateful applications, such as IP traceback, stateful firewalls, worm detection, and packet inspection, to meet security and manageability requirements of the next generation Internet [4]. We expect that these trends will lead to a network architecture where network systems maintain a large amount of state information and need to exchange this information with other network systems (i.e., routers and end-systems). To avoid complexities associated with numerous point solutions for each type of state information exchange, we believe that it is important to design a solution to unify network state information management.

One of the main challenges in this research is to find a balanced solution that is sufficiently general, but is efficient and feasible to implement. In this paper, we propose a framework for structuring and managing network state information that can be applied to a number of different network systems. Specifically, the contributions of this work are: (1) A structuring of network state information into a unified three-dimensioned state space that serves as the basis for state information management; (2) A unified addressing mechanism to manage network state information across different systems.

2. STATE MANAGEMENT FRAMEWORK

The proposed framework includes two major parts: a definition of the state information space and the processes for state information management.

2.1 Network State Information Space

In our work, the term "network state information" refers to the state information that is maintained in the memory of network systems and typically used for packet processing. Some illustrative examples of state information inside routers are routing tables, forwarding information bases, packet counters, and flow classification information. Such state information is used in the data path as well as in the control path.

We organize the state information into a three-dimensional state space, as illustrated in Figure 1. Each dimension provides a different point of reference for state in the network. The three dimensions of our framework are:

- **Space Dimension**: Packets traverse different routers to reach their destination. These routers and end-systems maintain their own state and thus are considered as separate nodes along the space dimension.

*CoNEXT'*07, December 10-13, 2007, New York, NY, U.S.A.

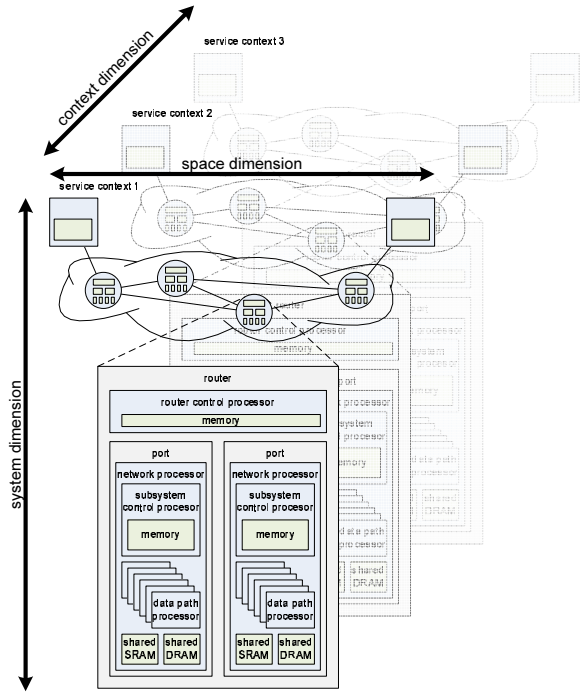

Figure 1: State Information Space

- **System Dimension**: A router (or end-system) may consist of different levels of components that maintain their own memory and network state (e.g., data path packet processing system, control processor, etc.). The number of levels and the type of system components vary with the system architecture of each node.
- **Context Dimension**: State can be associated with different services that are implemented on routers. Traffic belonging to different services may use different processing functions and thus requires access to different state information.

2.2 State Information Management

Our proposed state information management system can be roughly compared to the concept of virtual memory [5] on a workstation computer. While virtual memory hides details of the physical location of data, our state information management framework hides the complexities of accessing and exchanging state information spread over the network between network entities.

Our proposed state information management framework includes the following three major components:

- **State Information References**: Our framework provides access to network state through a common interface. The whole state information space can be referenced through a virtual address. This virtual address consists of the three dimensions mentioned above: space, system, and context.
- **Address Mapping**: In order to access to the state information, the virtual space address needs to be translated into a physical memory address on the network system where the data is stored. This functionality can be achieved by using a mapping process similar to how it is done for segments in virtual memory. Isolation and access control are two important issues that need to be considered in this context to ensure that safe and secure processing.
- **Access Mechanism**: When a physical memory address on a network system has been identified, the actual memory access needs to be performed. Different access mechanism (e.g., shared memory, DMA, message passing, or the combination of serval) are necessary for local and non-local memory accesses. The framework hides the complexity of these mechanisms from the user.

3. CONCLUSION AND FUTURE WORKS

We have presented a framework for structuring and managing network state information. We believe that finding a unified design to manage all state information in network is important. Such a system can hide the complexities of accessing and exchanging state information across heterogeneous system implementations. Our proposed framework for referencing and accessing state information may simplify the design and implementation of network services and other packet processing tasks in the next-generation Internet. This architecture is a first step in this direction and a basis for continued research and discussion.

4. REFERENCES

[1] A. Feldmann. Internet clean-slate design: what and why. In *SIGCOMM Computer Communication Review* 37,3 (July 2007), 59-64.

[2] T. Wolf. Service-centric end-to-end abstractions in next-generation networks. In*Proc. of Fifteenth IEEE International Conference on Computer Communications and Networks(ICCCN)*, page 79-86, Arlington, VA, Oct 2006.

[3] R. Yates, D. Raychaudhuri, S. Paul, and J. Kurose. Postcards from the edge: a cache-and-forward architecture for the future internet. http://www.nets-find.net/Postcards.php.

[4] J. Verdu, M. Nemirovsky, J. Garcia, and M. Valero. Architecture impact of stateful networking applications. In *Proc. of ACM/IEEE Symposium on Architectures for Networking and Communication Systems(ANCS)*, page 11-18, Princeton, NJ, Oct 2005.

[5] A. Silberschatz, J.L. Peterson, and P.B. Galvin. Operating System Concepts. third edition, Addison-Wesley Publishing Company, 1991.