HotMiddlebox : Hot Topics in Middleboxes and Network Function Virtualization

http://conferences.sigcomm.org/co-next/2013/workshops/HotMiddlebox/index.html

Introduction

Modern networks increasingly rely on advanced network processing functions for a wide spectrum of crucial functions ranging from security (firewalls, IDSes, traffic scrubbers), traffic shaping (rate limiters, load balancers), dealing with address space exhaustion (NATs) or improving the performance of network applications (traffic accelerators, caches, proxies), to name a few. Such "network appliances" or "middleboxes" are a critical piece of the network infrastructure and represent, to a first-order approximation, the de-facto approach for network evolution in response to changing performance, security, and policy compliance requirements.

However, most of this functionality is implemented in costly, hard-to-modify dedicated hardware, making the network difficult to evolve or adapt to changing traffic requirements. Recent work seeks to address this issue by shifting network processing from a world of dedicated hardware to one built on software-based processing running on (sometimes virtualized and shared) platforms built on commodity hardware servers, switches, and storage. This vision of "software-based" network services enables new in-network functions to be rapidly instantiated, on-demand, and at places in the network where it is most needed, without having to modify the underlying hardware. The scope of this workshop focuses both on the design of the data plane to support advanced services as well as the control plane functions necessary to manage these advanced data plane functions. In some sense, this vision is complementary to ongoing efforts in the SDN community, where the focus has largely been on the control plane and assuming a commodity data plane.

While our workshop builds on the recent promise of realizing high-performance network processing on commodity hardware, many questions remain open:

- What are the best virtualization technologies for implementing high-performance network functions?
- What are the challenges when trying to push them to rates of 10Gb and beyond?
- How do we provide the best possible isolation, both in terms of software isolation but also performance?
- How do we ensure that middlebox modules from different entities running on the same platform are assigned to the available hardware in an optimal way?
- How can we provide quick instantiation of processing (in the order of milliseconds or less)?
- What control plane abstractions are necessary to manage such advanced and stateful services?

The HotMiddlebox workshop will serve as an avenue to showcase and discuss ongoing work

from both academic and industry efforts in this space and to identify key challenges and potential solutions, with the ultimate goal of providing a roadmap for practical deployment in operational networks.

Scope of the workshop

We encourage the submission of work-in-progress papers in the area of middlebox design, implementation, measurement, management, and deployment. We look for submissions of previously unpublished work on topics including, but not limited to, the following:

- Performance optimizations of network stacks on virtualized systems
- Verification of unknown code running on shared middlebox platforms
- Extensible software stacks for rapid implementation of new middlebox functions
- Mechanisms for migration of stateful middleboxes
- Resource allocation mechanisms for shared/virtualized middlebox platforms
- Integrating new software middleboxes into legacy networks
- Backend storage/memory architectures for middleboxes
- Management abstractions and policy language frameworks for middleboxes
- Experiences in deploying software-based middleboxes in operational networks
- Connections to Software-Defined Networking
- Deployment and use of middleboxes in the cloud
- Measurements of middleboxes in enterprise, ISP, and data center networks.
- Novel security, performance, and monitoring applications atop middleboxes
- Challenges for policy verification in the context of middlebox services.

Submission Guidelines

Each submission must be a single PDF file no longer than six (6) pages in length (in two-column, 10-point format) including references, following the LaTeX style file. Papers should be submitted electronically via the submission site [to be announced]. Papers must include the author name and affiliation for single-blind peer reviewing by the program committee.

Accepted papers will be published in the ACM Digital Library. Publication at Hot Middlebox is not intended to preclude later publication of an extended version of the paper. Authors of accepted papers are expected to present their papers at the workshop

Important Dates

Submission: Fri August 30th, 5pm ET
Notification: September 30th
Camera ready: October 24th
Workshop date: December 9, 2013

Co-Chairs
Felipe Huici, NEC Europe
Vyas Sekar, Stony Brook University

Technical Program Committee:

Aditya Akella, UWisconsin
Katerina Argyraki, EPFL
Jon Crowcroft, Cambridge
Lars Eggert, NetApp
Norbert Egi, Huawei
Mark Handley, University College London
Giuseppe Lettieri, University of Pisa
Diego Lopez, Telefonica
Anil Madhavapeddy, Cambridge
KyoungSoo Park, KAIST
Peter Pietzuch, Imperial
Sylvia Ratnasamy, Berkeley
Jennifer Rexford, Princeton
Luigi Rizzo, University of Pisa
Anees Shaikh, IBM
Robin Sommer, ICSI
Nick Feamster, GTech
Julian Chesterfield, OnApp/XenSource
Teemu Koponen, Nicira / VMware
Christian Macciocco, Intel Labs
Bob Briscoe, British Telecom