

## HotMiddlebox : Hot Topics in Middleboxes and Network Function Virtualization

### 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 where software-based processing runs on virtualized, 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. This trend towards virtualizing network functions is called NFV and has gained a lot of traction in the industry in the past years, to the point where we standards are being discussed and initial deployments are sprouting. It is also foreseen that such in-network commodity infrastructure will be used not only by operators, but also by third parties, and operators may become miniature cloud providers.

It is also well known that middleboxes ossify the Internet - they force all traffic to “look” like existing protocols for security and performance reasons; even app-level protocols have been hardwired into the network (e.g. protocol specific proxies for HTTP). The net effect is that extending the core protocols (e.g. TCP, even HTTP) has become increasingly difficult, and new applications must hide their traffic just to get through the network by using tunneling of various forms. In short, middleboxes have pushed the endpoints to use less efficient protocols, and this trend will continue.

This workshop focuses 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.
- revisiting the architectural implications of middleboxes and proposing feasible solutions that can be embedded into software middleboxes, before they are widely adopted.

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?
- 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, deployment, as well as Internet architecture implications of middleboxes. 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
- Security issues regarding middleboxes
- 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
- 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.
- Internet architecture implications of middleboxes.

## Important Dates

Abstract Registration: March 24th, 5pm ET

Submission: Fri March 31st, 5pm ET

Notification: April 30th, 2015

Camera ready: May 15th, 2015

Workshop date: August 21, 2015

## Steering Committee Members:

Bob Briscoe, BT, UK

Christos Kolias, Orange, USA

Sylvia Ratnasamy -- U Berkeley, USA

Vyas Sekar, CMU, USA

## Co-Chairs:

Theophilus Benson, Duke University, USA

Costin Raiciu, University Politehnica of Bucharest, Romania

## Technical Program Committee Members:

Pedro Aranda, Telefonica, Spain

Olivier Bonaventure, UCL, Belgium

Jon Crowcroft, Cambridge University, UK

Yan Cai, Google, USA

Lars Eggert, NetApp, Germany

Felipe Huici, NEC, Germany

Nate Foster, Cornell, USA

Dongsu Han, Kaist, Korea

Hani Jamjoom, IBM, USA

Changhoon Kim, Barefoot Network, USA

Li Erran Li, Bell Labs, USA

Dave Meyer, Brocade, USA

Andrew Moore, Cambridge University, UK

Laurent Mathy, U of Liege

Robin Sommer, ICSI, USA

Minlan Yu, USC, USA