Securing the network against malicious programmable switches

Supervisors: Prof. Ankit Singla and Prof. Laurent Vanbever

Our computing ecosystem relies heavily on highly available and performant networks. This critical role has meant that networked systems have long been a high-value target for malicious actors, with high-profile attacks reported with alarming frequency and increasing severity. While the most common attacks target end systems or services like DNS, increasingly, networking devices themselves are coming under attack. The motivation is transparent: compromising network routers and switches provides visibility across large swathes of individual end devices. For many attackers, being able to compromise network devices is thus a higher priority [1].

We explore a new potential threat: that attackers try to cause visible and obvious deterioration in networked applications, while making the diagnosis of these problems challenging.

Such a threat can be posed by an attacker with control over programmable switches inside a target network. The attacker can cleverly choose *which* packets to manipulate such that only a small number of packets are affected, but the applications suffer a lot. Typical network monitoring systems that estimate packet losses, *e.g.*, by Microsoft [2] and Facebook [3], will be largely ineffective against such attacks.

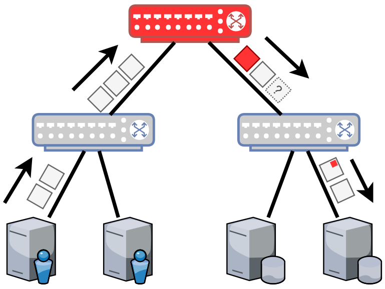

Figure 1: *The malicious switch can drop, delay or modify passing packets, and generate new ones.*

Our initial results show that a surprisingly small number of packets attacked can translate to extremely poor application-layer performance. We must thus design a monitoring system which mitigates such attacks. To achieve this goal, the proposed work will explore the following questions:

- How are attacks from malicious programmable switches different from traditional attacks? Which creative ways are there to exploit the programmable behavior? What is the maximum damage that can be done while remaining undetected by today's monitoring systems?
- At what granularity should a monitoring system operate to detect these attacks? What degree of malicious behavior is detectable given the noise inherently present within operational networks?
- Is it essential for the network monitoring system to be application-aware? If so, which application information should be made available to it?

This project involves design, analysis, and implementation components, but can be tailored to emphasize any of these, depending on student skills and ability. It will involve creative thinking from both the attacker perspective, as well as from the countermeasure perspective.

[1] NSA laughs at PCs, prefers hacking routers and switches. Kim Zetter (2013).
`https://www.wired.com/2013/09/nsa-router-hacking/`
[2] 007: Democratically Finding the Cause of Packet Drops. Arzani et al. (2018).
[3] Passive Realtime Datacenter Fault Detection and Localization. Roy et al. (2017).

Contact: Simon Kassing, Ankit Singla, or Laurent Vanbever: {kassings, asingla, lvanbever}@ethz.ch