Internet Security Liberated via Software Defined Exchanges

Kuang-Ching Wang
Clemson University
Clemson, South Carolina
kwang@clemson.edu

Richard R. Brooks
Clemson University
Clemson, South Carolina
rrb@clemson.edu

Geddings Barrineau
Clemson University
Clemson, South Carolina
cbarrin@clemson.edu

Jonathan Oakley
Clemson University
Clemson, South Carolina
joakley@clemson.edu

Lu Yu
Clemson University
Clemson, South Carolina
lyu@clemson.edu

Qing Wang
Clemson University
Clemson, South Carolina
qw@clemson.edu

ABSTRACT

With software defined networking and network function virtualization technologies, networks can be programmed to have customized processing and paths for different traffic at manageable costs and for massive numbers of applications. Now, picture a future Internet where each entity - a person, an organization, or an autonomous system - has the ability to choose how traffic in their respective network sessions is routed and processed between itself and its counterparts. The network is, essentially, liberated from today's homogeneous IP-based routing and limited connection options. To realize such a network paradigm, we propose a software defined exchange architecture that can provide the needed network programmability, session-level customization, and scale. We present a case study for traffic-analysis-resistant communication among individuals, campuses, or web services, where IP addresses no longer need to have a one-to-one correspondence with service providers.

CCS CONCEPTS

• **Networks → Network design principles**; **Programmable networks**; **Public Internet**;

KEYWORDS

SDN, SDX, Cloud, Censorship Circumvention, GENI, PEERING, Internet Architecture, Wide Area Network

ACM Reference Format:
Kuang-Ching Wang, Richard R. Brooks, Geddings Barrineau, Jonathan Oakley, Lu Yu, and Qing Wang. 2018. Internet Security Liberated via Software Defined Exchanges. In *SDN-NFV Sec'18: 2018 ACM International Workshop on Security in Software Defined Networks & Network Function Virtualization, March 19–21, 2018, Tempe, AZ, USA*. ACM, New York, NY, USA, 4 pages. https://doi.org/10.1145/3180465.3180475

SDN-NFV Sec'18, March 19–21, 2018, Tempe, AZ, USA

ACM ISBN 978-1-4503-5635-0/18/03…$15.00
https://doi.org/10.1145/3180465.3180475

1 INTRODUCTION

The advent of software defined networking (SDN) and network function virtualization (NFV) technologies signals the readiness to create programmable network solutions at scale. Put in the context of the Internet, this means opportunities to realize customized network services at unprecedented scales and in flexible scopes. With NFV, the Internet traffic can be transformed where needed with a virtualized network function (VNF) instance. With SDN, such VNFs can be programmatically controlled by software controllers via web APIs, serving diverse stakeholders' needs at web scale.

When it comes to network security, contemporary solutions strongly reflect two primary characteristics of the Internet architecture today - *end-to-end packet delivery* [7] and *IP routing based on autonomous systems* (ASes). With the end-to-end principle, traffic processing specific to applications and users takes place within boundaries of individual ASes - hereafter referred to as the *Internet edge*. Beyond AS boundaries, traffic enters the *Internet core* and gets forwarded based on *peering agreements* amongst ASes and Internet Exchange Point (IXP) providers. Today, Internet security solutions largely focus on the Internet edge, through encryption, access control, authentication, and intrusion detection and prevention. Recently, network security headlines have reflected vulnerability in the Internet core, through such attacks as border gateway protocol (BGP) hijacking [4, 17] and nation-state censorship [2]. These attacks target IP routing and IP addressed based detection and blocking of traffic in the core. Currently, BGP routing depends heavily on manual configuration, and BGP hijacking frequently masquerades as fat fingering. One big vulnerability is the near global visibility of traffic flows. Managing the scope of where the source and destination is visible makes DDoS and MITM attacks more difficult.

To begin addressing such risks in the Internet core, efforts are underway to further secure the BGP protocol and domain name systems (DNS). Nonetheless, the complexities in global AS peering relationships and the fact that all packets have their IP addresses in plain text leaves some types of attacks extremely difficult to combat. In this paper, an alternative approach is proposed to deter such security threats rooted in the correlation of IP addresses and individual persons, organizations, or applications. While the concept can be applied in different scopes, this paper focuses on its implementation as a distributed software defined exchange (SDX). Bearing similarities with an IXP but with distinctive differences in design and capabilities, the proposed SDX approach utilizes distributed data centers that, beyond traditional IXP peering services, allows

different types of customers(say, individuals v.s. ASes) to instantiate individually customizable VNFs, much like leasing VMs on a public cloud, at the SDX's data centers for application-specific needs. In this paper, a use case is presented for achieving traffic analysis resistant networking (TARN). With VNFs running software SDN switches, the SDX allows customers to communicate with other individuals or content providers using specially designed IP address hopping schemes so that the communication becomes incredibly hard for third parties to detect, eavesdrop, or block using firewalls.

The proposed SDX approach has significant implications to the future landscape of the Internet. It suggests a fundamental change in Internet users' choice of how they communicate over the Internet. This also comes with a change in the relationships among Internet users, AS operators, IXP/SDX, and content providers. In the context of network security, opportunities arise in the research, development, and business offerings of security solutions beyond today's solutions. 2 The remainder of the paper is organized as follows. Section 2 discusses backgrounds and related work. Section 3 presents the proposed SDX architecture and implications. Section 4 presents the TARN case study. Section 5 concludes the paper.

2 BACKGROUNDS AND RELATED WORK

Today, the Internet is made up of nearly sixty thousand ASes interconnected by hundreds of IXPs worldwide [3]. As D. Clark et al. discussed in [7], while interconnectivity in the Internet is rich, complex, and driven by a wide spectrum of contracting relationships, the focus has been on the *value flow* across ASes. Contracts were negotiated based on perceived values of whether or not to allow traffic through their ASes. While the practice has contributed to the growth of the Internet, it also limited the Internet's ability to support certain applications.

A recent, fine-grain examination of Internet traffic at one of the largest European IXPs has revealed an interesting view of the Internet beyond the traditional, traffic-agnostic AS-level view [5]. Specifically, due to the adoption of content distribution network (CDN) strategies by major content providers, they have been hosting their servers in third party ASes (e.g., cloud providers or customers) to accelerate content distribution, resulting in a global footprint of their administered servers. For example, as seen at the IXP studied, Akamai peers with and hosts servers in some 400 other ASes. In a security context, this highlights how such providers may desire a cross-AS security service footprint under its own customization and control, which is only possible via overlay networks (e.g., VPNs) and hard to achieve "down-to-the-wire" today.

An early work on the SDX concept was by Feamster, et al. [9], exploring the use of SDN at an IXP to enable finer grain, application specific peering beyond what BGP is capable of today. Since then, various forms of SDX have been proposed. In [13], SDX interconnects multiple network domains at flexible layers (not limited to layer 3 or BGP) via signaling among federated network controllers. At the 2013 SDN Program Review workshop [6], an expanded view of SDX was discussed as a software defined infrastructure, providing both SDN interconnect and inline compute services that can be custom programmed via software APIs.

One of the most discussed security threats of the day is traffic analysis. Even with encrypted traffic, there are sufficient clues to detect activities of specific applications, devices, or content access. In [1], it was shown how encrypted sensor devices can leak significant information of the owners' private activities by simply observing network flow traffic patterns. One of the modes of attacks depend on the source and/or destination IP addresses of observed flows. In Section 4, the use case "TARN" tackles this exact vulnerability by eliminating an adversary's ability to use IP addresses to identify target flows. The work in [11] also attempted to mutate IP addresses for end hosts to avoid unwanted traffic analysis. The solution, however, uses random addresses within the same subnet of a source node, while the proposed TARN applies randomization across wide area networks.

3 SOFTWARE DEFINED EXCHANGE

In this paper, SDX is defined as a system that provides inter-AS Internet connectivity. Figure 1 illustrates the basic components that make up the infrastructure for a SDX provider, who operates multiple SDX data centers at geographically distributed locations. Each SDX data center (DC) is like an IXP facility today. Considering practical economics, we expect SDXs to connect with other IXPs and transit ASes to attain global Internet reach.

SDX differs from an IXP in that, beyond the switching fabric, it hosts a scalable NFV cloud. With the NFV cloud, the SDX can support not only legacy BGP-based peering among ASes but also *customized peering methods* and *consumer network services*. Both require programmable, high performance packet processing for large numbers of distinct customers at public cloud scale.

Customized peering methods: In [9], application-specific peering was suggested as a SDX use case. In this paper, a new *secure and dynamic prefix* (SDP) peering method is studied to support the TARN use case. With SDP, the SDX provider dynamically allocates a set of multiple prefixes to a customer (an AS or a person) at any given time. The allocation is secure, i.e., not disclosed to anyone, and it can be used by the VNFs instantiated by that customer. Pushing the envelope even further, the capability can potentially completely relieve the need for an AS to hold and manage its own IP prefix and let the SDX provide it. More discussion of its implications is in Section 4.

Consumer network services: IXPs today support peering requests from ASes. In contrast, the proposed SDX adopts a public cloud model so as to support massive numbers of customers. Beyond ASes, any consumer can "purchase" one or more VNF-based network service instances and apply them to their Internet connections. Figure 2 illustrates the service models for a SDX in contrast with that of IXPs today.

Conceptually, *SDX = Cloud + NFV + SDN*. By directly serving consumers in volumes as a public cloud, and using VNFs to enable flexible offering of useful services, Internet routing, especially in the core, becomes significantly more heterogeneous than today. This makes the Internet more responsive to special needs of new types of applications while remains scalable. SDX can offer other resources, such as IP prefixes, as well for use by applications, e.g., TARN.

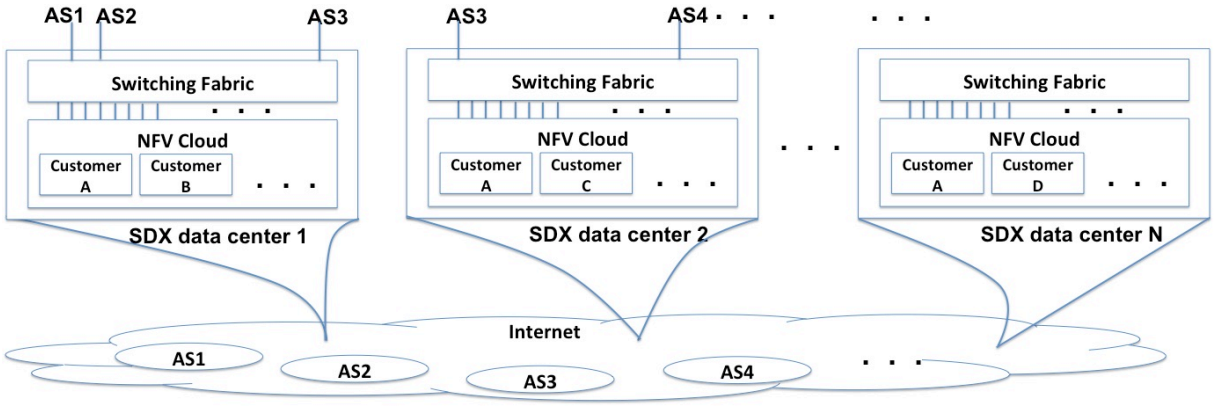

Figure 1: Distributed SDX infrastructure.

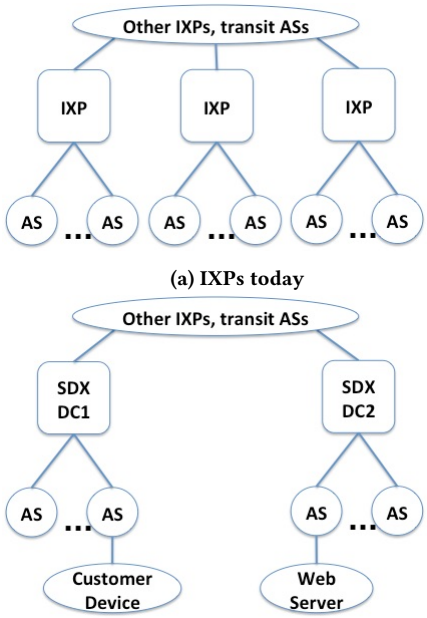

Figure 2: Comparing service models.

4 TARN: A CASE STUDY

The TARN study's objective is to combat threats of traffic analysis that may result in active Internet communication sessions being detected, blocked, or the parties at either end being tracked and persecuted. Such threats are prominent in the Internet today in certain parts of the world. Most solutions today are based on proxied tunnels, and such solutions are prone to their proxies being detected. Since traffic analysis predominantly uses IP addresses of communication sessions as basis, tremendous benefits can be gained by making it difficult to correlate IP addresses to potential targets - such as websites or organizations. This requires tackling one of the Internet's major assumption from the very beginning, that it is composed of ASes, each of which ownes one or more statically assigned IP prefixes. In the case of censorship, it has been common for adversaries to block IP addresses of specific websites or all addresses belonging to specific IP prefixes. Today, there is no real way of disguising a destination's IP prefix other than using a proxy such as Tor [15], I2P [10], Psiphon [14], or Lantern [12], trusting the proxies have not been detected.

TARN was first introduced in [16]. Three deployment strategies (end host, campus gateway, and SDX) were proposed with different trust assumptions and BGP routing update needs. With SDX, TARN can be invoked as a service for individuals or entire ASes. SDX-based TARN provides traffic analysis resistant communication between any two SDX DCs. Customers at both ends connect to a nearby SDX DC via a secure connection (VPN or dedicated circuit). With TARN, destination addresses for communication sessions will appear random. To support the service, the SDX provider will need to own a large pool of IP prefixes and dynamically assign them for use by different customers. With IPv6, there is abundant available prefixes to support the need. Such a prefix ownership and use model, however, is unforeseen. The SDX provider is expected to have a very large, potentially global, service footprint, with SDX DCs distributed across the area. For ease of management and refreshed protection, the SDX owned prefixes can be partitioned for use by different SDX DCs and such allocation can be regularly updated.

4.1 Assumptions

The SDX-based TARN makes a number of assumptions:

- The SDX supports a range of services, including TARN.
- Companies contract with the SDX to offer TARN-enabled service, e.g., access to its website.
- Individuals contract with the SDX for the service.
- Companies may require individual registeration for added levels of protection.
- ASes have a layer 2 connection to the SDX. Individuals have a VPN connection to the SDX.
- When VPN is not allowed or risky for use in certain areas, an alternative "end-host based TARN" was described in [16].

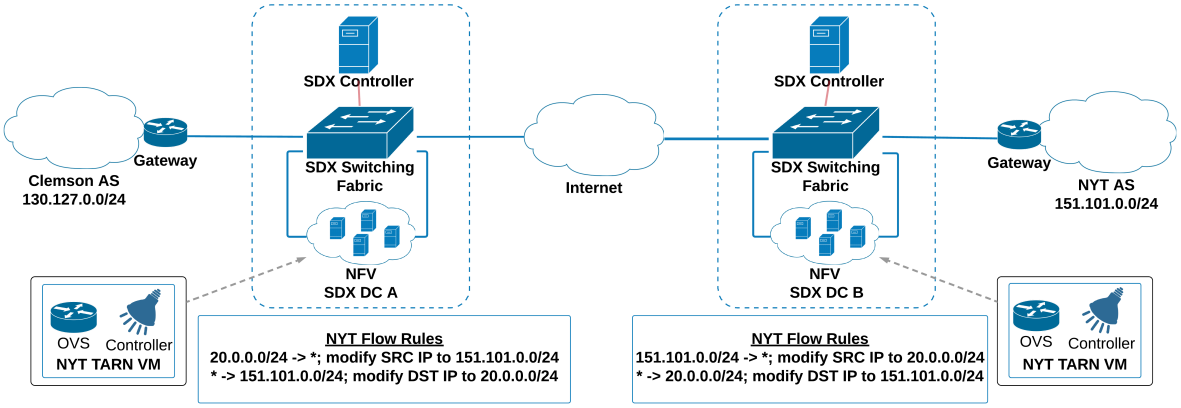

Figure 3: A SDX TARN implementation.

4.2 Implementation

Figure 3 illustrates an example based on one possible implementation for the SDX-based TARN. In the example, a fictitious company, say New York Times (NYT) contracts with the SDX to enable TARN-enabled access to its website for customers with such needs. Consider such a customer in AS A signing up the service, followed by accessing NYT by way of SDX DC A. Note that each SDX DC has a SDN-based fabric to steer customer traffic to and from service VNFs in the NFV cloud. Inside the NFV cloud, service-specific VNFs are instantiated as VMs. We assume a TARN VNF runs an Open Vswitch (OVS) controlled by a SDN controller. Our prototype builds on the FloodLight open source controller [8]. In this example, the SDX instantiates one NYT TARN VNF at each DC. Traffic entering the VNF with NYT's original prefix 151.101.0.0/24 as destination, which gets rewritten into a short-lived, "external" prefix 20.0.0.0/24 for traversing the "unsafe" Internet between DC A and B. External prefixes are drawn from a pool of SDX owned prefixes, which get reallocated to different SDX DCs over time and their BGP routing updated accordingly. At DC B, the NYT TARN VNF will restore the destination address with its original prefix before forwarding to NYT. OpenFlow rewrite flow rules are inserted accordingly in OVSes in the NYT TARN VNFs.

5 SUMMARY AND FUTURE WORK

In this paper, a SDX architecture is proposed to enable new forms of network services at cloud scale based on SDN and NFV. For security, it allows new services to alter traffic in the Internet core. TARN uses the SDX to control the scope of IP visibility in the core. The authors have been prototyping TARN for experimental validation on the US National Science Foundation sponsored GENI and PEERING testbeds. The authors are working with international partners to validate TARN at multi-nation scale.

ACKNOWLEDGMENTS

This material is based upon work sponsored by the National Science Foundation under Grant No. 1643020. Any opinions, findings, and conclusions or recommendations expressed in this material are those of the authors and do not necessarily reflect the views of the National Science Foundation.

REFERENCES

[1] Noah Apthorpe, Dillon Reisman, Srikanth Sundaresan, Arvind Narayanan, and Nick Feamster. 2017. Spying on the smart home: Privacy attacks and defenses on encrypted iot traffic. *arXiv preprint arXiv:1708.05044* (2017).
[2] Simurgh Aryan, Homa Aryan, and J Alex Halderman. 2013. Internet Censorship in Iran: A First Look.. In *FOCI*.
[3] Tony Bates, Philip Smith, and Geoff Huston. 2018. CIDR Report. (2018). http://www.cidr-report.org/as2.0/.
[4] Russell Brandom. 2015. A network error routed traffic for the UK's nuclear weapons agency through Russian telecom. (March 2015). https://www.theverge.com/2015/3/13/8208413/uk-nuclear-weapons-russia-traffic-redirect.
[5] Nikolaos Chatzis, Georgios Smaragdakis, Jan Böttger, Thomas Krenc, and Anja Feldmann. 2013. On the benefits of using a large IXP as an Internet vantage point. In *Proceedings of the 2013 conference on Internet measurement conference*. ACM, 333–346.
[6] Vince Dattoria, Inder Monga, Bryan Lyles, Kevin Thompson, and Grant Miller. 2013. SDN Program Review Final Report. (December 2013). https://www.orau.gov/sdnpr2013/.
[7] Peyman Faratin, David D Clark, Steven Bauer, William Lehr, Patrick W Gilmore, and Arthur Berger. 2008. The growing complexity of Internet interconnection. (2008).
[8] Floodlight 2018. Project Floodlight. (2018). Retrieved January 26, 2018 from http://www.projectfloodlight.org/
[9] Arpit Gupta, Laurent Vanbever, Muhammad Shahbaz, Sean P Donovan, Brandon Schlinker, Nick Feamster, Jennifer Rexford, Scott Shenker, Russ Clark, and Ethan Katz-Bassett. 2015. SDX: A software defined internet exchange. *ACM SIGCOMM Computer Communication Review* 44, 4 (2015), 551–562.
[10] I2P 2018. The Invisible Internet Project. (2018). https://psiphon.ca/en/download.html.
[11] Panos Kampanakis, Harry Perros, and Tsegereda Beyene. 2014. SDN-based solutions for moving target defense network protection. In *World of Wireless, Mobile and Multimedia Networks (WoWMoM), 2014 IEEE 15th International Symposium on a*. IEEE, 1–6.
[12] Lantern 2018. Lantern. (2018). Retrieved January 26, 2018 from https://getlantern.org/en_US/
[13] Joe Mambretti, Jim Chen, and Fei Yeh. 2014. Software-defined network exchanges (SDXs): Architecture, services, capabilities, and foundation technologies. In *Teletraffic Congress (ITC), 2014 26th International*. IEEE, 1–6.
[14] Psiphon 2018. Psiphon. (2018). https://psiphon.ca/en/download.html.
[15] Tor 2018. Tor. (2018). https://www.torproject.org/projects/torbrowser.html.en.
[16] Lu Yu, Qing Wang, Geddings Barrineau, Jon Oakley, Richard R Brooks, and Kuang-Ching Wang. 2017. TARN: A SDN-based Traffic Analysis Resistant Network Architecture. *arXiv preprint arXiv:1709.00782* (2017).
[17] Kim Zeeter. 2017. Someone's Been Siphoning Data Through a Huge Security Hole in the Internet. (April 2017). https://www.wired.com/2013/12/bgp-hijacking-belarus-iceland/.