

SDN-enabled Traffic Engineering and Advanced Blackholing at IXPs

Christoph Dietzel
TU Berlin / DE-CIX
christoph@inet.tu-berlin.de

Gianni Antichi
University of Cambridge
gianni.antichi@cl.cam.ac.uk

Ignacio Castro
Queen Mary, University London
i.castro@qmul.ac.uk

Eder L Fernandes
Queen Mary, University London
e.leao@qmul.ac.uk

Marco Chiesa
Université catholique de Louvain
marco.chiesa@uclouvain.be

Daniel Kopp
DE-CIX
daniel.kopp@de-cix.net

ABSTRACT

While the clean slate approach proposed by Software Defined Networking (SDN) promises radical changes in the stagnant state of network management, SDN innovation has not gone beyond the intra-domain level. For the inter-domain ecosystem to benefit from the advantages of SDN, Internet Exchange Points (IXPs) are the ideal place: a central interconnection hub through which a large share of the Internet can be affected. In this demo, we showcase the ENDEAVOUR platform: a new software defined exchange approach readily deployable in commercial IXPs. We demonstrate here our implementations of traffic engineering and Distributed Denial of Service mitigation, as well as how member networks cash in on the advanced SDN-features of ENDEAVOUR, typically not available in legacy networks.

CCS Concepts

•**Networks** → **Network design principles; Routing protocols; Programmable networks; Network management;**

Keywords

Software Defined Networking, Internet eXchange Points

1. INTRODUCTION

By facilitating Autonomous System (AS) interconnection, Internet eXchange Points (IXPs) are a critical element in the Internet ecosystem [1, 2]. IXPs provide a simple layer-2 switching fabric to which members connect and exchange traffic with other collocated ASes. Once a peering arrangement is established, IXP members rely on the Border Gateway Protocol (BGP), the de-facto standard inter-domain routing protocol, to exchange routing information.

Because of this central role, IXPs are the ideal place to innovate the inter-domain ecosystem. The limited size of IXPs, in terms of deployed equipment and geographical distribution, eases the physical migration towards novel network architectures and infrastructure. Furthermore, it is typ-

ically in the interest of the IXPs to offer services that go beyond simple layer-2 connectivity. For instance, many IXPs already operate route servers, which simplify peering by allowing IXP members to peer with other ASes via a single BGP session to a route server. Facing a continuous traffic growth in their networks, IXPs have a strong economic incentive to embrace innovation to reduce operating costs.

This demo showcases the ongoing effort within the ENDEAVOUR project By deploying Software Defined Networking (SDN) at IXPs, ENDEAVOUR aims to create a flexible SDN ecosystem that supports a service marketplace composed of data centers, networked applications, and the underlying interconnection fabric. The ENDEAVOUR architecture advances existing work on SDXs [5] by providing an SDX architecture deployable in single and multi-hop topologies, as well as by proposing use cases that benefit not only the IXP members, but also IXP management itself. Figure 1 depicts the rich architecture that we envision as IXPs transition to SDN. At the data plane level, programmable SDN switches forward traffic according to fine-grained forwarding and filtering rules generated by the control plane (i.e., the SDX controller) in accordance to declarative high-level goals established by applications. In this scenario, a number of different applications can be instantiated to accomplish various use cases (e.g., Traffic Engineering, Advanced Blackholing). The SDX controller then translates these goals into rules and programs the data plane through the southbound interface, i.e., OpenFlow. In this demo, we advance the high-level ideas of ENDEAVOUR [3] with the demonstration of our implementation for SDN-enabled Traffic Engineering and Advanced Blackholing.

Traffic Engineering (TE) allows network operators to express their routing policies. At IXPs, member ASes that exchange large traffic volumes, typically demand connections through multiple ports (e.g., multiple 100GE). These ASes ideally employ TE to load-balance traffic through these ports and minimize congestion while attaining high port utilization. The outcome depends on an interplay between the in-

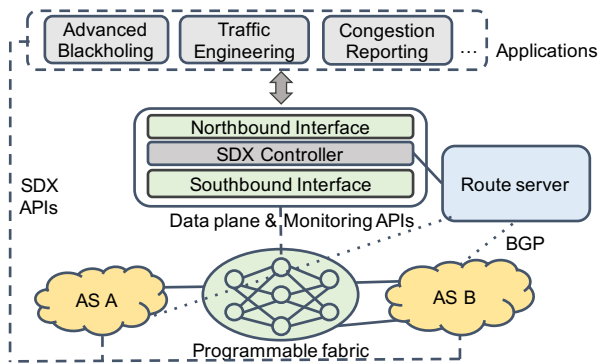


Figure 1: Rich architecture of an SDN-enabled IXP.

bound TE policies of the traffic receiver and the outbound TE policies of the originator. As of today, well-known limitations of BGP restrict what network operators can achieve. Operators must resort to indirect BGP configuration mechanisms, e.g., AS path prepending, communities, and/or selective announcements. This problem is particularly acute at IXPs where the wide range of independent and inconsistent peering policies might clash. Resolving such conflicts is a manual, and thus error-prone, process which may lead to further inconsistencies.

Blackholing allows networks to counter the frequent Distributed Denial of Service (DDoS) attacks that ASes face. Legacy blackholing allows an AS to request neighbors to drop packets towards a certain IP prefix. Unfortunately, the approach is rather deficient: it has not the adequate granularity (exclusively destination-based) and fails to distinguish between legitimate and malicious traffic: all packets destined for the blackholed IP prefix are dropped, thus disconnecting its upstream networks [4].

2. DEMONSTRATION DESCRIPTION

While past SDX examples, such as [5], consider IXP topologies with a single hop, ENDEAVOUR extends previous efforts to work in multi-hop topologies, where members are connected to edge switches and the traffic is aggregated with multiple core switches. This is a fairly common configuration, as the DE-CIX and AMS-IX examples show. The implementation distributes the member’s policies across the respective edge switches. While outbound policies are present only in the respective participant edges, inbound policies are installed in every edge switch because the decision process occurs in the first hop. When the edge resolves the next hop, a labeling mechanism encodes the path of the packet to follow in the fabric. We showcase the ENDEAVOUR capabilities with a virtual environment created on Mininet.

2.1 SDN-enabled Traffic Engineering

SDN improves Traffic Engineering operations at IXPs by going beyond the exclusively IP destination-based approach of BGP. As we show in this demo, SDN allows networks to steer traffic flows according to layer 2-4 attributes. The demonstration of the SDN-enabled *Traffic Engineering* features four IXP members connected to different edge switches in an edge-core-edge topology. Members have outbound and inbound policies configured. We show the fine-grained sys-

tem capabilities over a multi-hop architecture, by enabling an outbound policy to direct traffic on port 80 to a designated member. By leveraging the ENDEAVOUR monitoring capabilities, graphical statistics are displayed to show that traffic flows accordingly to the established policies.

2.2 Advanced Blackholing

SDN helps to overcome the mentioned limitations in blackholing by allowing operators to specify fine-grained drop policies and automating the blackholing process, hence reducing response time and overhead. The IXP can provide an interface, e.g., an API, for the members to express their precise drop rules and have them automatically implemented as an attack is detected. The IXP can also provide insights in the blackholed traffic by monitoring it through the OpenFlow counters. While some traffic might still be unintentionally blackholed, OpenFlow rate-limiting capabilities can alleviate the problem by limiting the traffic (based on specific header fields) towards the attacked members.

This demo demonstrates *Advanced Blackholing* in a topology setup comprising three IXP members connected to the IXP’s fabric. Initially, we show a set of traffic flows. Over time, one member activates a set of different blackholing policies, affecting traffic destined to the destination port 53. The policies are translated into flow rules which are then installed within the IXP fabric, as a result; matching traffic flows are discarded. In contrast to commodity blackholing, the blackholed traffic is still visible. Finally, the installed rules are removed and the traffic returns to its initial pattern.

Videos are available at https://youtu.be/6_adVvSJ7FA and https://youtu.be/enm_17jbEvs.

Acknowledgments: This research is (in part) supported by European Union’s Horizon 2020 research and innovation programme under the ENDEAVOUR project (grant agreement 644960).

3. REFERENCES

- [1] B. Ager, N. Chatzis, A. Feldmann, N. Sarrar, S. Uhlig, and W. Willinger. Anatomy of a Large European IXP. In *SIGCOMM*. ACM, 2012.
- [2] I. Castro, J. C. Cardona, S. Gorinsky, and P. Francois. Remote Peering: More Peering without Internet Flattening. In *CoNEXT*. ACM, 2014.
- [3] M. Chiesa, C. Dietzel, G. Antichi, et al. Inter-domain Networking Innovation on Steroids: Empowering IXPs with SDN Capabilities. *IEEE Communications Magazine*, 54(10):102–108, 2016.
- [4] C. Dietzel, A. Feldmann, and T. King. Blackholing at IXPs: On the Effectiveness of DDoS Mitigation in the Wild. In *PAM*, 2016.
- [5] A. Gupta, R. MacDavid, R. Birkner, M. Canini, N. Feamster, J. Rexford, and L. Vanbever. An Industrial-Scale Software Defined Internet Exchange Point. In *NSDI*. USENIX, 2016.