

# Inter-Domain Networking Innovation on Steroids: Empowering IXPs with SDN Capabilities

Marco Chiesa, Christoph Dietzel, Gianni Antichi, Marc Bruyère, Ignacio Castro, Mitch Gusat, Thomas King, Andrew W. Moore, Thanh Dang Nguyen, Philippe Owezarski, Steve Uhlig, and Marco Canini

The authors give a comprehensive overview of use cases for SDN at IXPs, which leverage the superior vantage point of an IXP to introduce advanced features like load balancing and DDoS mitigation. They discuss the benefits of SDN solutions by analyzing real-world data from one of the largest IXPs. They also leverage insights into IXP operations to shape benefits not only for members but also for operators.

## ABSTRACT

While innovation in inter-domain routing has remained stagnant for over a decade, Internet exchange points (IXPs) are consolidating their role as economically advantageous interconnection points for reducing path latencies and exchanging ever increasing amounts of traffic. As such, IXPs appear as a natural place to foster network innovation and assess the benefits of SDN, a recent technological trend that has already boosted innovation within data center networks. In this article, we give a comprehensive overview of use cases for SDN at IXPs, which leverage the superior vantage point of an IXP to introduce advanced features like load balancing and DDoS mitigation. We discuss the benefits of SDN solutions by analyzing real-world data from one of the largest IXPs. We also leverage insights into IXP operations to shape benefits not only for members but also for operators.

## INTRODUCTION

The growth of demand for high-performance online services is posing tremendous challenges on independent networks (i.e., autonomous systems [ASs]), which carry the traffic of these services and form the Internet. Pressing requirements, such as lower latencies and higher bandwidth, have pushed ASs to move from the strict hierarchical, transit-based interconnection model of the early commercial Internet toward a denser and flatter structure. Internet exchange points (IXPs) are playing a leading role during this transition phase by providing a simple layer 2 broadcast domain to which members connect and through which they exchange IP traffic, possibly with any other member (Fig. 1, left). After establishing peering relationships, IXP members exchange routing information by means of the Border Gateway Protocol (BGP), the de facto standard inter-domain routing protocol. IXPs have become central to the Internet peering ecosystem by attracting increasing numbers of members and, consequently, traffic. Nowadays, about 80 percent of the address space is reach-

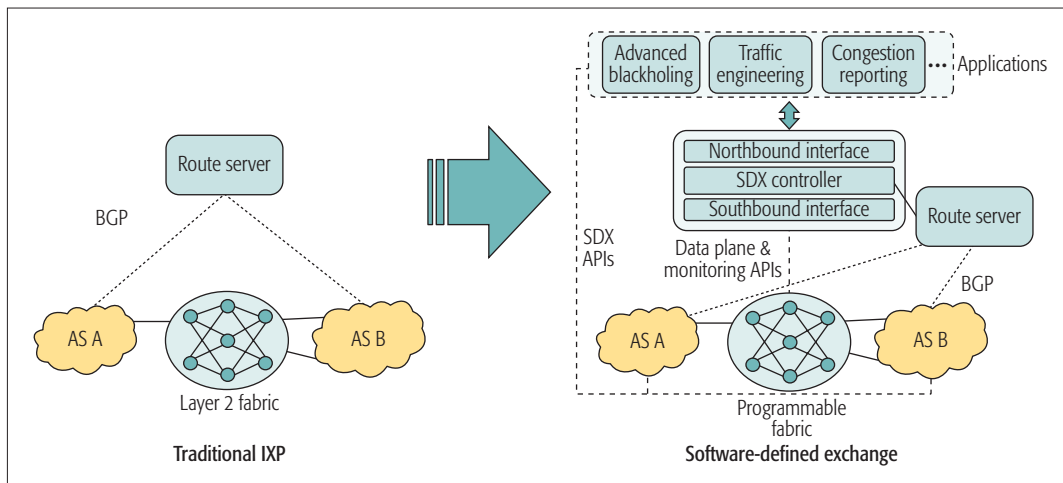
able through more than 350 existing IXPs [1]. The largest IXPs interconnect hundreds of ASs and carry traffic volumes comparable to those of Tier 1 transit providers [2].

Over a decade of work has gone into proposing modifications to the BGP routing control plane to improve its security, and make it easier to manage and troubleshoot [3]. As these proposals require substantial global changes in BGP, unfortunately there has been no significant adoption. As such, inter-domain routing still suffers from well-known shortcomings of BGP such as its coarse-grained control of traffic based on just destination IP prefixes, and indirect control of how remote networks forward traffic.

As others have argued [4], we also deem IXPs an ideal place to spur innovation in the Internet ecosystem. First, IXPs are convergence points for a large number of ASs. Any improvement that can be deployed at a large IXP has direct impact on hundreds of members. Second, IXPs' physical networks are fairly limited in size. This eases any physical migration toward novel network architectures and infrastructure. Third, IXPs are interested in offering services that go beyond simple layer 2 connectivity. To simplify peering, many IXPs already operate route servers, which allow IXP members to peer with many other ASs via a single BGP session to a route server. Finally, IXPs have strong economic incentives to embrace innovation for reducing operating costs in the face of continuous traffic growth in their networks.

Software-defined networking (SDN) recently emerged as a new paradigm that enables network programmability to facilitate management and enhance flexibility [5]. SDN supports logically centralized control of a collection of programmable switches that can act on traffic at a fine level of granularity. In contrast, BGP, which has been the standard for exchanging inter-domain reachability information for decades, only provides limited and indirect control knobs to network operators. BGP operates at a coarse-grained level based on IP destination prefixes. Despite SDN's potential, its success so far has

*Marco Chiesa, and Thanh Dang Nguyen are with Université Catholique de Louvain; Christoph Dietzel is with DE-CIX and TU Berlin; Gianni Antichi and Andrew W. Moore are with the University of Cambridge; Marc Bruyère and Philippe Owezarski are with LAAS-CNRS; Ignacio Castro and Steve Uhlig are with Queen Mary University of London; Mitch Gusat is with IBM Research; Thomas King is with DE-CIX. Marco Canini was with Université Catholique de Louvain, and is now with King Abdullah University of Science and Technology (KAUST).*



**Figure 1.** The evolution from a traditional IXP fabric to a richer SDN-enabled IXP: a software-defined exchange (SDX). The SDX exposes APIs to IXP members through which they can consume advanced services. For compatibility, IXP members exchange routes using BGP via a route server. Applications establish high-level policies and behaviors for the underlying IXP fabric. The SDX controller mediates access to the programmable IXP fabric and exposes monitoring information.

Traffic-engineering refers to the steering of traffic flows by network operators to improve performance. Achieving these goals depends on the accuracy of the network state estimation, the algorithms used to compute the optimal routing paths, and the specific features of the routing scheme adopted.

been restricted to environments such as intra-domain routing and data centers [6].

In this article, we advocate that IXPs offer an exciting opportunity for inter-domain networking to benefit from the advantages of SDN deployment in neutral, dense environments, thus giving rise to the concept of software-defined exchanges (SDXs). By looking at real data from one of the largest IXPs, we demonstrate the need of enhanced network capabilities. Further, we show the potential benefits of SDXs by presenting use cases. We advance existing work [4, 7–9] by proposing use cases that benefit not only the IXP members but also the IXP management itself. Also, we show evidence of the need of a breakthrough in network management at the IXP. We believe that SDN programmability and its finer-grained control capabilities at large IXPs will lead to novel peering arrangements, greater responsiveness, and easier network management.

The work presented in this article is part of an ongoing effort within the ENDEAVOUR project,<sup>1</sup> a research and innovation action funded by the European Union’s Horizon 2020 program, which began in January 2015. By focusing on SDN deployment at IXPs, the project aims to create a flexible SDN ecosystem, which can support a service marketplace composed of cloud/data centers, networked applications, and the underlying interconnection fabric.

Figure 1 depicts the architectural evolution 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 with 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 well defined rules and programs the underlying data plane through the southbound interface (e.g., OpenFlow).

The rest of the article is organized as follows.

We provide an overview of those customized applications, including fine-grained traffic engineering, mitigation of denial of service (DoS) attacks, broadcast message reduction, layer 2 label switching, and other emerging use cases. We discuss some of the challenges and the road ahead for SDXs, and finally, we conclude.

## USE-CASES

This section presents some of the benefits we believe an SDX approach brings to the current generation of IXPs from both operators’ and members’ perspectives.

### IXP MEMBER’S BENEFITS

**Traffic Engineering:** Traffic engineering (TE) refers to the steering of traffic flows by network operators to improve performance. Achieving these goals depends on the accuracy of the network state estimation, the algorithms used to compute the optimal routing paths, and the specific features of the routing scheme adopted (e.g., per-destination routing, hash-based load balancing).

At an IXP, ASs exchanging large traffic volumes connect with multiple ports (e.g., multiple 100GE). These ASs employ TE to load balance traffic through these ports and avoid congestion while attaining high port bandwidth utilization. The outcome of the TE operations is the result of the interplay between the inbound TE policies of the traffic receiver and the outbound TE policies of the originator. While the inbound policies specify what type of traffic can be received through each port alongside information about the route used to forward that traffic, outbound policies involve a mechanism to compute where to send the originated traffic.

Well-known limitations of BGP hinder what TE goals network operators can achieve. Operators must resort to indirect BGP configuration mechanisms (e.g., AS path prepending, communities, and selective announcements). It is then hoped that configuration changes have the desired effect, for example, incoming traffic is evenly split among the IXP member ports.

<sup>1</sup> <https://www.h2020-endeavour.eu/>




Figure 2. Inbound traffic load balancing in a large IXP.

Certain types of control are simply not possible. This problem is particularly acute at IXPs where the wide range of independent and inconsistent peering policies might clash. The lack of network programmability poses an additional problem that can lead to human mistakes and reduces AS responsiveness to network events.

Despite the aforementioned pitfalls of current solutions, TE is widely performed by IXP members. Figure 2 shows how members with multiple ports (roughly 15 percent of IXP members) load balance their inbound traffic across them. Each IXP member is depicted by a vertical bar, and the partitioning within the same bar represents the relative traffic volume per port. Observe that 27 out of 103 IXP members receive more than 99 percent of their traffic through a single port, probably using the second one as a backup port. The remaining members perform inbound TE more intensively. We advocate that TE operations should be simplified by means of more intuitive primitives and easier network programmability.

We believe that SDN substantially enhances TE operations. First, it allows network operators to express their TE routing policies with a specification language that is more fine-grained than the legacy BGP-based one. Operators can effectively use this greater expressiveness to effectively steer the flows of traffic according to layer 2–4 attributes. Second, to perform TE, IXP members have to estimate their incoming traffic volumes. IXPs can leverage SDN monitoring capabilities to provide an interface for accessing a global coherent view of the state of the network, which also includes incoming traffic volume statistics. In addition, TE can directly be outsourced to the IXP, which applies the policies specified by the IXP members. Automating TE engineering leads to less human error, whereas outsourcing the process frees IXP member resources. Third, SDN programmable networks allow operators to choose TE algorithms based on their performance goals and scalability limits. For instance, SDN-based load balancing can be performed in either a static or dynamic manner. The former leverages weighted hash-based per-flow load balancing mechanisms, which are available in the OpenFlow 1.3 standard. It allows an IXP network to spread the incoming traffic on behalf of the receiver member without any knowledge of the traffic volume. The latter enables each IXP member to further optimize TE whenever the former approach does not achieve the desired optimization goals, as is the case whenever a few large flows of traffic are hashed to the same port. In that case, fine-grained monitoring and routing capabilities can be used to detect and reroute those flows of traffic that cause traffic imbalances.

<sup>2</sup> <http://blog.cloudflare.com/the-ddos-that-almost-broke-the-internet>

<sup>3</sup> [www.de-cix.net/products-services/de-cix-frankfurt/blackholing/](http://www.de-cix.net/products-services/de-cix-frankfurt/blackholing/)

**Advanced Blackholing:** Distributed DoS (DDoS) attacks are a serious threat to the Internet ecosystem.<sup>2</sup> To cope with this threat, some IXPs offer blackholing services to their members.<sup>3</sup> Blackholing in general is a technique that allows an AS to ask its neighbors to drop packets destined toward a certain IP prefix. DDoS attacks, which typically affect a limited number of IP destination addresses, cannot be mitigated by withdrawing the BGP announcement of the entire IP prefix that contains the target of the attack, since this would also affect large portions of legitimate traffic. With blackholing, an AS first has to detect that it is receiving malicious traffic targeting a specific small contiguous set of IP prefixes. To stop the incoming flow, it sends a special BGP announcement for the affected IP prefix to its upstream neighbor that originates this traffic. Upon receiving such a message, each packet destined to that IP prefix is deflected toward a dummy medium access control (MAC) address by means of Address Resolution Protocol (ARP) mechanisms. All packets with this destination MAC are discarded via layer 2 filtering rules implemented via access control lists at the IXP switching fabric.

Current blackholing implementations at IXPs are too coarse-grained [10]. Although AS operators are capable of filtering on layer 4 ports within their own ASs, fine-grained blackholing for multiple AS border routers is severely limited by ACL hardware constraints (e.g., static nature and limited number of rules). Also, it does not protect the AS peering link at the IXP, that is, it can be overwhelmed and congested. Furthermore, blackholing cannot distinguish between legitimate and malicious traffic: all packets destined for the blackholed IP prefix are dropped, thus disconnecting its upstream networks. To make things more complex, blackholing lacks programmability and must be manually triggered after the DDoS attack is detected. Since the AS operators cannot observe the traffic volumes (e.g., attack termination) and patterns (e.g., destination port mix), identifying the beginning and termination of an attack is cumbersome.

We envision SDN-enabled blackholing to overcome those limitations. SDN, and OpenFlow in particular, allows operators to specify fine-grained drop policies and eases the blackholing process, hence minimizing the risk of misconfiguration. Using OpenFlow, ASs can detect a DDoS attack by monitoring the traffic properties and define fine-grained drop rules to discard the unwanted packets. The IXP can then provide an interface, such as an application programming interface (API), so that members can express their very own precise drop rules and have them automatically implemented when an attack is detected. The IXP can also provide insights on the blackholed traffic by monitoring the corresponding flows through the OpenFlow flow counters. While some traffic might still be unintentionally blackholed, OpenFlow rate-limiting capabilities can alleviate the problem. By limiting the traffic (according to specific header fields) toward the attacked members to a non-critical volume, the legitimate traffic can still stand a chance of reaching its destination.

**Controlling Broadcast:** IXPs interconnect members' routers through a shared layer 2 broadcast domain. As the scale of IXPs keeps increasing (the biggest IXPs nowadays count more than 600 members), traditional address resolution mechanisms that rely on broadcast solutions, such as ARP or neighbor discovery (ND), pose a challenge for faultless operation and stability. Broadcast packets are needlessly processed by all the routers connected at the IXP, consuming an excessive amount of each router's CPU capacity. Thus, broadcast mitigation and filtering at the IXP's edge becomes crucial. Figure 3 shows an increase of 10 percent in the volume of broadcast traffic (i.e., location discovery) at a large IXP in only 15 months. This growth continues to happen in spite of IXPs' strict rules for the configuration of members' routers<sup>4</sup> and techniques such as ARP sponge.<sup>5</sup> In contrast to a traditional switch, an SDN switch by default can drop all packets not matching any of the installed forwarding rules. IXP operators can thus program the IXP fabric to solve the problems exposed above. We observe that the topology of an IXP fabric is fairly stable over time, with the addition or removal of members happening at a sufficiently low frequency. Because the location of all members' routers is stable and well known to the IXP operators, broadcast traffic can be eliminated by just programming the switching SDN-enabled fabric to exclusively transport packets to the requested destination by ARP and ND packets.

**Layer 2 Label Switching:** Large IXPs, especially those with multiple points of presence, need complex and robust infrastructures to satisfy members' requirements, such as robustness and scalability. These IXPs use a transparent layer 3 infrastructure for internally addressing these requirements, while providing an emulated layer 2 fabric to their members.<sup>6</sup> Unfortunately, exploiting path diversity within an IXP fabric remains a problem. Traditional solutions such as multiprotocol label switching (MPLS) and virtual private LAN service (VPLS) help but come at the cost of increased complexity. Despite multiple recent attempts, the two main standards bodies, IEEE 802.1aq<sup>7</sup> and TRILL,<sup>8</sup> have not yet converged to a uniquely and universally accepted solution. As this divergence effectively precludes interoperability across vendors, a solution is required. We believe that SDN can help to simplify this complexity while preserving the benefits of layer 3 protocols, that is, leveraging the IXP architecture and reducing the protocol and management overhead. Fine-grained label switching capabilities such as those provided by OpenFlow can be used to route the traffic efficiently over multiple paths across the switching fabric. Ingress packets can be matched with installed flow rules and annotated with a label, either with an MPLS label or encoded in any arbitrary header field [7], which shall be removed at the IXP's egress interface.

### EMERGING NOVEL SERVICES

**Port-Based Congestion Reporting:** Managing congestion is an essential factor for an IXP and its members. However, in spite of experiencing

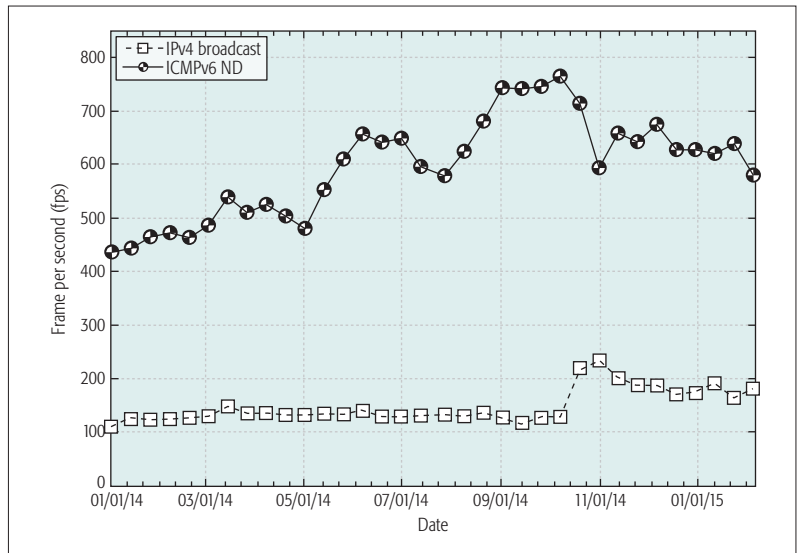


Figure 3. ARP requests and ICMPv6 ND frames per second at a large IXP.




Figure 4. CDF of port utilization in percent showing that 43 ports exceed 100 percent utilization.

significant and persistent congestion at multiple peering links, ASs and IXPs have no primary means of controlling congestion. That is, as the traffic sources and destinations are beyond its domain, an IXP cannot rely on the traditional congestion notification mechanisms such as explicit congestion notification (Internet Engineering Task Force [IETF] RFC 3168).

Nowadays, congestion at the level of a member's port is a significant problem. To illustrate, Fig. 4 shows our measurements from one of the largest IXPs during a seven-day period, where we monitor using a five-minute interval of the traffic volume flowing from the IXP into each IXP member's port. We found that 43 (~6 percent) of the 760 IXP member's ports suffered at least once from congestion (i.e., more than 100 percent utilization). Once a member's egress

<sup>4</sup> <https://www.euro-ix.net/networks/configuration-samples/>

<sup>5</sup> <https://ams-ix.net/technical/specifications-descriptions/controlling-arp-traffic-on-ams-ix-platform>

<sup>6</sup> <https://ams-ix.net/technical/ams-ix-infrastructure/the-ams-ix-mpls-vpls-infrastructure>

<sup>7</sup> <http://www.ieee802.org/1/pages/802.1aq.html>

<sup>8</sup> <http://tools.ietf.org/html/rfc6325>

Building upon this increasing diversity, IXPs could become a marketplace for transport. Benefiting from the co-location of ASes, the existing infrastructure and transparency standards, the IXP could help supply to meet the demand of transport over the Internet in all its different flavors.

port reaches 100 percent utilization, excess traffic starts filling buffers and is dropped once buffers are full. By accounting for all traffic that traverses the IXP, our measurement captures the severity of this congestion. For instance, a 200-Mbit and a 10-Gbit port experienced traffic rates at 473 and 214 percent of their capacity, respectively. Moreover, we found that utilization at 114 ports exceeds 85 percent.

Empowering IXPs with SDN, and in particular with OpenFlow capabilities, would ease both the per port utilization and packet loss monitoring process for each member. Furthermore, SDN enables highly dynamic solutions wherein the ASs can be informed whenever a specific congestion level is reached. For instance, since version 1.5.1, OpenFlow includes push-based counters monitoring triggered by predefined thresholds. IXP members could act upon congestion reports and apply different routing policies based on network conditions.

**IXP as Transport Marketplace:** IXPs have dynamized peering interconnections by enabling ASs to automatically peer through the route server with those ASs willing to peer with any other network. Furthermore, going beyond their original goal of being the central place for peering relationships, IXPs have expanded the transport opportunities. IXPs are frequently interconnected by remote peering providers, which offer layer 2 transport between IXPs [1]. This interconnection between IXPs is frequently directly promoted by the IXPs themselves. There is even anecdotal evidence of transit providers offering their transport services at IXPs.

Building on this increasing diversity, IXPs could become a marketplace for transport. Benefiting from the co-location of ASs, and the existing infrastructure and transparency standards, the IXP could help supply to meet the demand of transport over the Internet in all its different flavors.

Despite these developments, current interconnection practices are still characterized by a cumbersome and slow contract negotiation process. We believe that SDN can play a key role in supporting IXPs to meet those challenges. As with other SDN applications, APIs at the IXP can be leveraged to facilitate the choice of how and whether to interconnect. For instance, an AS might declare its peering conditions through the API, and all networks meeting them can automatically peer.

**Service Chaining:** Increasingly, middleboxes (e.g., firewalls, VPN gateways, proxies, intrusion detection systems, WAN optimizers or scrubbers) play a critical role in the performance and security of many networks. Operators commonly want to chain together multiple middleboxes to create a processing path of desired network functions, also called a *network service* [11]. Service chaining is emerging as an appealing solution to enable operators to dynamically deploy network services without having to make hardwire physical middleboxes. However, using traditional protocols like BGP or MPLS, service chains are difficult to deploy and change. Operators must carefully configure protocols to “hijack” wanted traffic and steer it through the sequence of middleboxes. Moreover, such mechanisms steer traffic based only on IP destination addresses, limiting the overall service granularity. In competitive markets,

with rapid innovation at the application layer, this limits operators’ ability to address emerging use cases and business models. Instead, an SDN approach is a well suited alternative to enable service chaining given its logically centralized management and configurable forwarding rules at fine granularity. Thus, inter-domain routing convergence points like IXPs appear to be promising locations to deploy service chains.

## DISCUSSION AND OUTLOOK

In the previous section, we presented a selection of appealing use cases that highlight the potential for innovation with SDXs. We now turn our discussion on broad research challenges toward bringing SDN at IXPs within grasp.

### SCALABILITY

SDXs will interconnect hundreds of networks and expose rich service APIs. This creates scalability challenges that will need to be addressed. For instance, given the size of global Internet routing, which counts over 500,000 IP prefixes, naively supporting arbitrary SDN applications may result in solutions that cannot even cope with the scale of medium-sized IXPs. Moreover, the growing interest in remote peering [1] will pose even more challenges in designing scalable solutions for large-sized IXPs. In fact, despite the SDN fine-grained control capabilities over traffic, hardware forwarding table sizes are a constrained resource. These considerations demand techniques such as those explored in the iSDX design [7] to efficiently compute and encode flexible routing policies and to minimize the number of data plane updates due to changes in inter-domain reachability information.

### RELIABILITY

A recent incident at AMS-IX in May 2015,<sup>9</sup> wherein a forwarding loop affected the IXP for a brief time, demonstrated how configuration mistakes or failures can dramatically impact the network. In this regard, the rise of new applications enabled by SDN will pose new challenges in ensuring a stable and reliable network operation. According to a recent report [12], many experts from the networking and formal methods communities believe that despite their importance, tools for programming and reasoning about networks are still in a state of infancy. We believe this also applies to SDXes specifically.

### SECURITY

Increased Internet security is a highly sought after goal and a desired use case by operators, particularly for detecting and preventing DDoS. Beyond DDoS mitigation, SDXs might enable new architectures that can help address broad classes of network attacks by design. For example, these architectures can provide an opportunity to reconsider the line of research on network capabilities, where embedding costs into traffic could act as a deterrent for attackers [13].

### PRIVACY

SDXs will provide a multitude of services beyond a traditional layer 2 interconnection including network functions (e.g., caches, optimizers, packet scrubbers) that will be deployed at SDXs. This raises questions regarding the privacy of

<sup>9</sup> <https://ams-ix.net/newsitems/195>

processing traffic at exchanges and the neutrality of SDXs [14]. Who controls the network functions? Who specifies what traffic traverses which network functions? How will it be regulated and what auditing will be performed?

### BUSINESS CONFIDENTIALITY

Members will interact with SDXs' control software through APIs to consume services. However, the correct or efficient consumption of these services might require exchange of information that is traditionally considered proprietary due to its business-critical nature, such as peering and routing policies. Therefore, new solutions (e.g., [15]) are needed to align conflicting objectives such as service consumption while avoiding leakage of confidential information. This is very challenging as possible approaches such as secure multi-party computations might be prohibitively compute intensive in practical settings. Beyond the technical solutions, we believe that to a large extent the problem lies in the vague legal framework for IXP-data disclosure rather than on whether the IXP is a reliable partner for neutral information sharing. A clearer legal framework will definitely help foster innovation at IXPs.

### TRANSITIONING TO SDN

Despite the advantages that SDN offers for innovating wide-area traffic delivery, SDXs will require a clear migration path from existing systems to deploy SDN software and hardware. We observe that a first step has been made to show the viability of deploying SDN hardware side by side with existing production equipment at public IXPs [16, 17], and also best practices exist for migration in enterprise data center, campus [18], and carrier networks.<sup>10</sup> However, actual experience with production deployments remain limited, and research that further spurs early adoption in operational networks will be crucial toward reaching a wide-spread deployment.

### SUMMARY

We explore the benefits of casting SDN concepts into Internet exchange points for fostering innovation in the Internet peering ecosystem. Based on multiple discussions with operators from the IXP community and an analysis of data traces from one of the largest IXP in the world, we illustrate several use cases that stand to benefit from advanced SDN capabilities at IXPs. Finally, we discuss open problems toward reaching the goal of widespread SDN deployment at IXPs.

### ACKNOWLEDGMENTS

We thank Josh Bailey, Yatish Kumar, David Meyer, and Jennifer Rexford for providing advisory guidance to the project. This research is (in part) supported by the European Union's Horizon 2020 research and innovation programme under the ENDEAVOUR project (grant agreement 644960).

### REFERENCES

- [1] I. Castro *et al.*, "Remote Peering: More Peering without Internet Flattening," *CoNEXT*, 2014.
- [2] B. Ager *et al.*, "Anatomy of a Large European IXP," *SIGCOMM*, 2012.
- [3] K. R. B. Butler *et al.*, "A Survey of BGP Security Issues and Solutions," *Proc. IEEE*, vol. 98, no. 1, 2010; <http://dx.doi.org/10.1109/JPROC.2009.2034031>, pp. 100–22.

- [4] N. Feamster *et al.*, "SDX: A Software Defined Internet Exchange," *Open Networking Summit*, 2013.
- [5] H. Kim and N. Feamster, "Improving Network Management with Software Defined Networking," *IEEE Commun. Mag.*, vol. 51, no. 2, Feb. 2013, pp. 114–19.
- [6] D. Kreutz *et al.*, "Software-Defined Networking: A Comprehensive Survey," *Proc. IEEE*, vol. 103, no. 1, 2015, pp. 14–76.
- [7] A. Gupta *et al.*, "An Industrial-Scale Software Defined Internet Exchange Point," *NSDI*, 2016.
- [8] A. Gupta *et al.*, "SDX: A Software Defined Internet Exchange," *SIGCOMM*, 2014.
- [9] S. Whyte, "Project CARDIGAN An SDN Controlled Exchange Fabric," 2012, <https://goo.gl/eTTgXG>.
- [10] C. Dietzel, A. Feldmann, and T. King, "Blackholing at IXPs: On the Effectiveness of DDoS Mitigation in the Wild," *PAM*, 2016.
- [11] S. Palkar *et al.*, "E2: A Framework for NFV Applications," *SOSP*, 2015, pp. 121–36.
- [12] N. Björner *et al.*, "Formal Foundations for Networking (Dagstuhl Seminar 15071)," *Dagstuhl Reports*, vol. 5, no. 2, 2015, pp. 44–63.
- [13] C. A. Shue *et al.*, "On Building Inexpensive Network Capabilities," *SIGCOMM Comp. Commun. Rev.*, vol. 42, no. 2, Apr. 2012; <http://doi.acm.org/10.1145/2185376.2185386>, pp. 72–79.
- [14] J. Sherry *et al.*, "BlindBox: Deep Packet Inspection over Encrypted Traffic," *SIGCOMM*, 2015, pp. 213–26.
- [15] M. Chiesa *et al.*, "Towards Securing Internet eXchange Points Against Curious onlookers," *Applied Networking Research Wksp.*, July 2016.
- [16] J. Stringer *et al.*, "Cardigan: Deploying a Distributed Routing Fabric," *HotSDN*, 2013.
- [17] LightReading, "Pica8 Powers French TOUJX SDN-Driven Internet Exchange," June 2015; <http://ubm.io/1VcOSLE>.
- [18] M. Canini *et al.*, "Software-Defined Networks: Incremental Deployment with Panopticon," *Computer*, vol. 47, no. 11, Nov. 2014, pp. 56–60.

### BIOGRAPHIES

MARCO CHIESA [M] ([marco.chiesa@uclouvain.be](mailto:marco.chiesa@uclouvain.be)) is a postdoctoral researcher at Université Catholique de Louvain, Belgium. His research interests include routing optimization and security. He received a Ph.D. in computer science and engineering from Roma Tre University.

CHRISTOPH DIETZEL ([christoph.dietzel@de-cix.net](mailto:christoph.dietzel@de-cix.net)) has been a member of the DE-CIX R&D team since June 2014, responsible for several research efforts and also involved in numerous projects funded by the public sector. Moreover, he has been a Ph.D. student in the INET group, advised by Anja Feldmann at Technische Universität Berlin, since the end of 2014. His ongoing research interests focus on Internet measurements and security, routing, and traffic classification. He is also highly interested in IXP-related aspects of the Internet ecosystem.

GIANNI ANTICHI ([gianni.antichi@cl.cam.ac.uk](mailto:gianni.antichi@cl.cam.ac.uk)) is a senior researcher at the Network and Operative System group in the Computer Laboratory at the University of Cambridge, United Kingdom. His research interests include software-defined networking, network measurements, and hardware accelerated networking systems. He received a Ph.D. in computer science and engineering from the University of Pisa in 2011.

MARC BRUYÈRE ([marc.bruyere@laas.fr](mailto:marc.bruyere@laas.fr)) is a Ph.D. student at the LAAS CNRS, Toulouse, France. He designed and deployed the first European OpenFlow IXP fabric for the ToulX. He started his career in 1996 working for Club-Internet.fr, and for Cisco, Vivendi Universal, Credit Suisse First Boston, Airbus/Dimension Data, Force10 Networks, and Dell. He is a Cisco Certified Internetwork Expert. He has been involved in the NetFPGA project for a few years, and his Ph.D. thesis is about measurements in an IXP OpenFlow/SDN environment.

IGNACIO CASTRO ([i.castro@qmul.ac.uk](mailto:i.castro@qmul.ac.uk)) received a Ph.D. degree from the Internet Interdisciplinary Institute (IN3) while researching at the Institute IMDEA Networks in 2015. He is currently a postdoctoral researcher at Queen Mary University of London. His research interests focus on the economics of Internet interconnections.

MITCH GUSAT [M] ([mig@zurich.ibm.com](mailto:mig@zurich.ibm.com)) is a researcher at the IBM Zurich Research Laboratory. His current focus is on data center performance, and distributed deep learning and feedback. He has contributed to the design and standardization of CEE, IBA, and RapidIO, while also advising Master's and Ph.D. students from several European universities. He is member of ACM and holds a few dozen patents related to SDN, transports, HPC architectures, switching, and scheduling.

THOMAS KING ([thomas.king@de-cix.net](mailto:thomas.king@de-cix.net)) received an M.Sc. degree (Diplom) in computer science and business administration from the University of Mannheim, Germany, in 2004, and a Ph.D. degree at the chair of Computer Networks from the University of Mannheim in 2008. He was head of the Research & Development department at DE-CIX until the end of 2015. In 2016, he has been promoted to the newly created position of CIO of DE-CIX.

ANDREW W. MOORE [M] ([andrew.moore@cl.cam.ac.uk](mailto:andrew.moore@cl.cam.ac.uk)) is a reader in systems at the University of Cambridge Computer Laboratory, England, where he

Beyond the technical solutions, we believe that to a large extent the problem lies in the vague legal framework for IXP-data disclosure rather than on whether the IXP is a reliable partner for neutral information sharing. A clearer legal framework will definitely help foster innovation at IXPs.

<sup>10</sup> <https://www.opennetworking.org/images/stories/downloads/sdn-resources/use-cases/Migration-WG-Use-Cases.pdf>

---

is part of the Systems Research Group working on issues of network and computer architecture. His research interests include enabling open network research and education using the NetFPGA platform; other research pursuits include low-power energy-aware networking, and novel network and systems data center architectures.

THANH DANG NGUYEN (thanh.nguyen@uclouvain.be) is a postdoctoral researcher at Université Catholique de Louvain, Belgium. His research interests include software-defined networking and large-scale distributed computing. He obtained his Ph.D. in information science from Japan Advanced Institute of Science and Technology (JAIST) in 2015. He is a member of ACM.

PHILIPPE OWEZARSKI (owe@laas.fr) is director of research at CNRS (the French center for scientific research), working at the Laboratory for Analysis and Architecture of Systems, Toulouse, France. He received a Ph.D. in computer science in 1996 from Paul Sabatier University, Toulouse III, and his habilita-

tion for advising research in 2006. His main interests deal with next generation Internet, more specifically taking advantage of IP networks monitoring and machine learning for enforcing quality of service and security.

STEVE UHLIG (steve.uhlig@qmul.ac.uk) received a Ph.D. degree in applied sciences from the University of Louvain (2004). He is currently a professor of networks at Queen Mary University of London. His research interests are focused on the large-scale behavior of the Internet, Internet measurements, software-defined networking, and content delivery. He is a member of ACM.

MARCO CANINI [M] (marco@kaust.edu.sa) is an assistant professor in assistant professor in computer science at King Abdullah University of Science and Technology (KAUST). His research interests include software-defined networking and large-scale and distributed cloud computing. He received a Ph.D. in computer science and engineering from the University of Genoa. He is a member of ACM and USENIX.