

Umbrella: a deployable SDN-enabled IXP Switching Fabric

Marc Bruyere¹, Gianni Antichi³, Eder L. Fernandes⁴, Remy Lapeyrade²,
Ignacio Castro⁴, Steve Uhlig⁴, Philippe Owezarski², Andrew W. Moore³

¹Information Technology Center, University of Tokyo, JP ³Computer Laboratory, University of Cambridge, UK

⁴Queen Mary University of London, UK ²CNRS, LAAS, FR

ABSTRACT

Software Defined Internet eXchange Points (SDXs) are a promising solution to the long-standing limitations and problems of interdomain routing. While proposed SDX architectures have improved the scalability of the control plane, these solutions have ignored the underlying fabric upon which they should be deployed. This work makes the case for a new fabric architecture that proposes stronger control and data plane separation.

KEYWORDS

Software Defined Networking, OpenFlow, Internet eXchange Point

1 INTRODUCTION

Internet eXchange Points (IXPs) are a central element of the Internet ecosystem: IXPs can carry huge traffic volumes and interconnect a multitude of networks of different types [1]. The fundamental service provided by IXPs is a Layer 2 neutral facility where heterogeneous networks exchange IP traffic. While IXPs are the ideal vehicle to extend the benefits promised by Software Defined Networking (SDN) to the interdomain level [4], reliability and scalability are essential aspects of an IXP that cannot be compromised by the introduction of SDN. In particular, the impact of control channel disruptions or outages can cause severe disturbances, potentially affecting hundreds of networks and huge traffic volumes [2]. Control plane failures in large scale deployed SDN networks outweigh largely data or management combined together [3].

In this paper we make the case for a stronger control and data plane separation and propose Umbrella, a novel approach to IXP fabric management that can be deployed in any IXP topology to reduce the risks of a fabric excessively dependent on the control plane. Umbrella leverages SDN programmability to tackle part of the control traffic, removing the actual mac learning mechanism in legacy IXPs networks. For this, the broadcasted Address Resolution Protocol (ARP) traffic is directly programmed within the data plane. Secondly, we use a Layer 2 encode path to minimize the resources, management cost and extend scalability. This approach greatly simplifies the management of the fabric: the only role of the controller is supervising the network, leveraging its global knowledge.

The main contributions of this paper are:

- We make the case for a stronger control and data plane separation in SDN IXPs.
- We introduce the Umbrella architecture and show how it leverages SDN programmability within the data plane.

2 THE CASE FOR STRONGER CONTROL AND DATA PLANE SEPARATION

Previous works have shown how OpenFlow (OF) [6] could be deployed at the exchange [4] using an IXP fabric with a central controller for all the peering routers at the IXP. In such architecture, the SDN controller would be co-located with the Route Server (RS) to ensure that the SDN and BGP control planes can talk to each other with a minimal delay. Despite advantages such as richer policies, one challenge remains: data plane issues may affect control plane messages, leading to a slow or unresponsive control plane, further aggravating the effect on the data plane. The critical problem resides in the centralized ARP-proxy: delays in the control channel might lead to all the connection oriented mechanisms (i.e., BGP, TCP) failing. For example, if the ARP messages of the peering router A get delayed inside the IXP fabric to reach the SDN controller: all the BGP sessions between the router A and its peers would suffer, forcing (in the worst case scenario) establishing new connections.

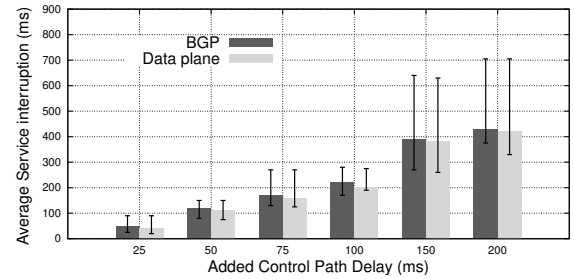


Figure 1: The dependency between control and data plane.

We show in Fig. 1 the disruption caused by ARP delays on the data plane or BGP sessions. We emulated the above SDN scenario of a central controller co-located with the RS on Mininet, a Mininet extension to build complex networks. We instantiated virtual containers acting as peering routers: one for the RS and two more working as client hosts directly connected to one peering router each. We represented the IXP as a single open vSwitch coupled with a Ryu controller acting also as an ARP-proxy, as in [4]. Fig. 1 shows that, even a small delay of a few tens of milliseconds for ARP messages may trigger much larger disruptions on the data plane or BGP. Given the large volumes of traffic IXPs' critical role, such disruptions are not acceptable [7].

3 A NEW SDN FABRIC FOR IXPS

IXPs apply strict rules [5] to limit the side effects of a layer-2 shared broadcast domain, e.g., the MAC address of the router with which

the member connects to the peering fabric must be known in advance. Only then the IXP will allocate an ethernet port on the edge switch and configure a MAC filtering Access Control List (ACL) with that MAC address. The location of all the member's routers is thus known to the IXP. For this reason, Umbrella eliminates the need of location discovery mechanisms based on broadcast packets, i.e., ARP request, IPv6 neighbor discovery, making unnecessary the active ARP-proxy daemon proposed in previous SDN-enabled IXP solutions. Umbrella makes on-the-fly translation of broadcast packets into unicast by exploiting the OF ability to rewrite the destination MAC address of a frame matching a rule.

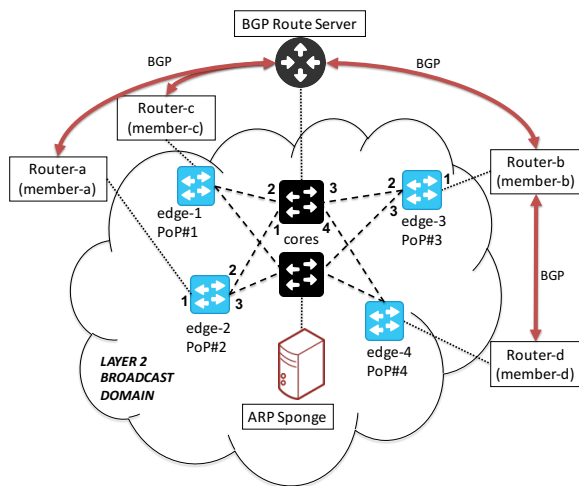


Figure 2: Typical topology of a medium to large IXP.

When the source and destination are directly connected to the same edge switch, no encoding is needed, and the edge switch directly replaces the broadcast destination address with the target MAC destination address. In an IPv6 scenario, the OF match pattern indicated in the edge switch needs to be on the IPv6 ND target field of the incoming ICMPv6 Neighbor Solicitation packet. The matching table on the edge switch should maintain an association between IPv6 addresses and their location, as in the IPv4 case.

Umbrella's forwarding mechanism allows reusing legacy switches in the core, limiting the burden (and costs) of upgrading the hardware. A core switch only needs to forward packets based on simple access filtering rules, whereas the edge switches need OF-like capabilities to rewrite the layer-2 destination field. While this approach is directly applicable to single-core IXP fabrics, it cannot handle multiple-hops fabrics. With a single hop, the core switch would expect the output port to be encoded in the most significant byte of the destination MAC address. In the multi-hop case, since a packet can traverse multiple core switches, a new encoding scheme is needed to distinguish the output ports at different core switches. This is a fairly common case in hypercube-like topologies, such as the ones adopted by LINX or MSK-IX.

4 CONCLUSION

REFERENCES