Poster: Private Internet: A Global End-to-End Service Model

Souvik Das
Department of Computer Science
The University of Texas at Dallas, USA
souvikdas95@utdallas.edu

Kamil Sarac
Department of Computer Science
The University of Texas at Dallas, USA
ksarac@utdallas.edu

***Abstract*—The public Internet is a network of autonomously owned and operated networks. Outdated peering policies and lack of end-to-end performance guarantees are causing its ossification which have led large cloud and content providers to build their own global private backbone infrastructures. As much as these private backbones help eliminate public transit for content hosted across their networks, content hosted elsewhere is still carried over the public Internet. In this poster, we propose a model where these private backbone operators collaborate with the access-networks of content providers and consumers to implement end-to-end network services with better performance characteristics than the public Internet. We call the resulting end-to-end service domain as a ”Private Internet”.**

I. Introduction

The Internet today is a network of Autonomous Systems (AS) that interconnect using bilateral agreements and policies. Operators at the edges deploy access-networks that connect the providers and consumers of content and services to the Internet. Access-networks run a variety of protocols to connect the end-users, however the ASes at large operate solely on Border Gateway Protocol (BGP). BGP dictates several types of relationships between pairs of ASes with access-networks paying to regional providers and regional providers paying to backbone providers in the Internet today.

Today's Internet is mainly used for high-value content delivery. The majority of content carried over the public Internet *transits* between backbones in the Internet core. As much as the content generates financial benefits for the content providers, the current peering policies in the public Internet define a financial model for the ISPs that does not provide a mechanism to share revenue generated by content delivery [1]. Moreover, with very little scope and control of traffic beyond the ”next hop”, it is difficult to guarantee optimality of the path for a content of interest with any degree of preference. Since there is no authority to hold responsible for the end-to-end flow of traffic or value, the Internet is falling short of meeting requirements, i.e. getting ossified [2].

Major players, such as Google, Amazon and Facebook, are building their own private backbones to overcome the limitations of public *transit*. These backbones peer with numerous access-networks spread across globally for the purpose of improving the delivery of content hosted in their backbones. In this poster, we propose a scenario where these major players also provide similar services for content hosted elsewhere. We propose a collaboration between a major player (e.g. Google) and the access-networks to support end-to-end provisioning and transit for content hosted in these access-networks (Section II). We also propose a high-level design of an architecture that can leverage this concept to enable global end-to-end path control (Section III). Finally, we conclude by briefly discussing our ongoing work (Section IV).

978-1-6654-4131-5/21/$31.00 ©2021 IEEE

II. Vision

Today large content providers with global infrastructures (CPGI), e.g. Google [3], provision content hosted in their networks as well as supply the required wide-area transit for it. On the other hand, there are various content-providers (CP) as well as content-consumers (CC) that are located at the access-networks (AN) that use public Internet as their transit provider (TP). We envision a scenario where a CPGI establishes peering with many ANs to virtually extend its infrastructure to become an end-to-end TP. In this model, an AN can consider a CPGI as an infrastructure client in its network and give CPGI the ability to control its traffic as it flows in the AN. This results in an end-to-end global network (that we term as ”Private Internet” or PI) controlled by a single entity, i.e. CPGI. Thanks to network slicing technology, we can now envision multiple PIs to operate simultaneously over the same AN.

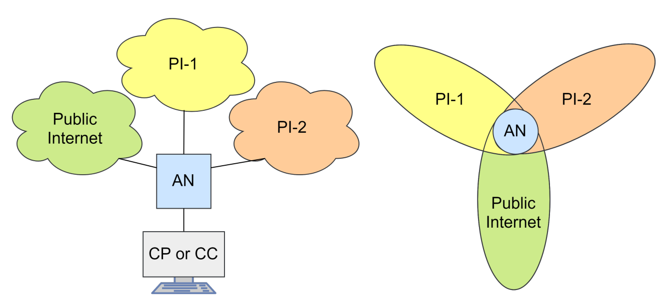

Fig. 1. Physical view (left) and logical view (right) of the proposed model.

In this model, CPGIs become Private Internet Operators (PIOs), ANs become part of a PI infrastructure, and PIO becomes the network service provider of CPs as well as CCs connected to those ANs. Given that the PIO operates over the end-to-end PI, it can introduce various service guarantees to end users (both CPs and CCs) and can be held accountable for these guarantees. The potential benefit of the proposed system includes: (i) for CPs and CCs, they have alternative ways of communication with potential end-to-end guarantees, (ii) for ANs, they have PIOs as new infrastructure costumers, and (iii) for CPGI, becoming PIOs to provide transit service introduces new business opportunities. Figure 1 presents a conceptual depiction of the proposed model.

In a related work [4], the authors proposed the concept of ”Zero-Hop Networking” which aims to introduce CPGI control into ANs of CCs only. In contrast, we introduce the ability to control the entire end-to-end network by PIOs (or CPGIs). Moreover, our model promotes a new business domain for participating CPGIs with the incentive to support a wider range of end-to-end services.

III. Proposed Architecture

In today's Internet, ANs use peerings with other networks to introduce alternative paths to meet the growing traffic demand by CPs and CCs. However, once the traffic goes across the peering points, no one can be held responsible for the performance of the rest of the path. In our proposal, PIOs control the end-to-end path between the endpoints. The endpoints obtain service-level agreements (SLAs) from the PIOs and PIOs determine the optimal path between the endpoints across the PI based on these SLAs. PIOs implement these service guarantees in cooperation with the ANs.

When an endpoint, such as CP wants to reach a remote endpoint, such as CC over a PI network, its request is captured by the PIO of the PI. The PIO then establishes an end-to-end path that fulfills the SLAs of both the endpoints. This path consists of three parts: sender, receiver and backbone. The sender part comprises of the path from the endpoint in the source AN to the selected peering location between the source AN and PI's backbone; the receiver part comprises of the path from the selected peering location between PI's backbone and the destination AN to the endpoint in the destination AN; and the backbone part comprises of the path that connects the two peering locations that the PI's backbone has with the source and destination ANs across the PIO infrastructure. These paths are obtained as a result of optimization over multiple available alternative paths across multiple peering locations in each of these parts. We assume that the ANs operate their networks using SDN and the PIO runs an SDN application on top of it to communicate with the ANs i.e. PIO is an infrastructure client of the ANs. This enables the PIO to obtain information about the alternative paths in the sender and receiver parts from the corresponding ANs. We also assume that the PIO knows the status of the alternative paths that connects the two parts within its backbone. Once the PIO processes these information in its backend SDN application and determines the end-to-end path, it can communicate the sender and receiver paths to the SDN controllers at the ANs to install them into their networks.

Figure 2 shows an example of the architecture with two access-networks, AN-1 and AN-2 that operate their networks using SDN and have peerings with public Internet as well as a

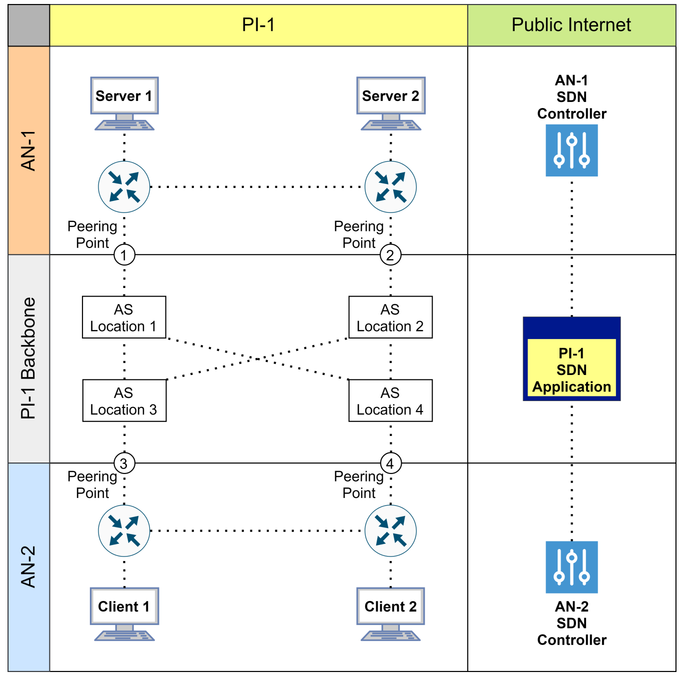

Fig. 2. Example of the architecture

”Private Internet”, PI-1. AN-1 and AN-2 join PI-1 at peering points $\{1, 2\}$ and $\{3, 4\}$ respectively. AN-1 hosts two CPs server-1 and server-2; AN-2 hosts two CCs client-1 and client-2. PIO-1 runs a server functioning as a client SDN application for the controllers in the ANs. When server-1 sends a packet to client-1, the SDN controller at AN-1 captures it from the SDN switches in AN-1 and requests PIO-1 to determine a path for it. PIO-1 evaluates this request against the SLAs of server-1 and client-1 and deploys a path between them across AN-1, AN-2 and its backbone. The role of public Internet in this example is to enable the SDN application in PIO-1 and the SDN controllers in both the ANs to communicate control messages over the Northbound interfaces.

IV. Ongoing Work

Currently, we are working on several components of the proposed architecture such as how a PIO ensures that its traffic is isolated from other PIOs at the ANs; how the backbone part of the end-to-end path is deployed; how endpoints can determine which PI to utilize from among multiple PIs available at the ANs; and how to implement routing and control in such a multi-network environment in a scalable way.

References

[1] Y. Harchol *et al.*, “A public option for the core,” in *Proceedings of ACM SIGCOMM*, August 2020.

[2] M. Ammar, “Ex uno pluria: The service-infrastructure cycle, ossification, and the fragmentation of the internet,” *SIGCOMM Comput. Commun. Rev.*, p. 56–63, 2018.

[3] K.-K. Yap *et al.*, “Taking the edge off with espresso: Scale, reliability and programmability for global internet peering,” in *Proceedings of ACM SIGCOMM*, August 2017.

[4] M. Ammar, E. Zegura, and Y. Zhao, “A vision for zero-hop networking (zen),” in *Proceedings of IEEE ICDCS*, June 2017.