
Public Review for
**Open Connect Everywhere: A Glimpse at
the Internet Ecosystem through the Lens
of the Netflix CDN**

Timm Bottger, Felix Cuadrado, Gareth Tyson
Ignacio Castro, Steve Uhlig

In this measurement paper the authors expose the deployment of the Netflix CDN, focusing on two large countries: USA and Brazil. In contrast to the other big video players, Netflix delivers its content without operating neither a backbone network nor data centers. To reach its customers instead, Netflix relies on deploying servers at IXPs and inside ISPs. The detailed analysis of the Netflix infrastructure deployment performed by the authors becomes an opportunity to study the role that IXPs play as a platform to enable large-scale content delivery and to reach a world-wide customer base and identifying some of the challenges involved.

The reviewers found this paper a pleasant read and agreed that it delivers relevant results uncovering information about how a *hypergiant's* network operates and highlighting how it is possible to deliver a large volume of data without a proprietary or leased infrastructure. Reviewers also observed that the methodology is not easily generalizable but appreciated the methods used and how they were validated.

Several papers in the past years have discussed about the *flattening* of the Internet. This paper contributes to shed light on this phenomenon with detailed empirical data.

Public review written by
Alberto Dainotti
CAIDA

Open Connect Everywhere: A Glimpse at the Internet Ecosystem through the Lens of the Netflix CDN

Timm Böttger
Queen Mary University of London
timm.boettger@qmul.ac.uk

Felix Cuadrado
Queen Mary University of London
felix.cuadrado@qmul.ac.uk

Gareth Tyson
Queen Mary University of London
gareth.tyson@qmul.ac.uk

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

Steve Uhlig
Queen Mary University of London
steve.uhlig@qmul.ac.uk

ABSTRACT

The importance of IXPs to interconnect different networks and exchange traffic locally has been well studied over the last few years. However, far less is known about the role IXPs play as a platform to enable large-scale content delivery and to reach a world-wide customer base. In this paper, we study the infrastructure deployment of a content hypergiant, Netflix, and show that the combined worldwide IXP substrate is the major corner stone of its Content Delivery Network. This highlights the additional role that IXPs play in the Internet ecosystem, not just in terms of interconnection, but also allowing players such as Netflix to deliver significant amounts of traffic.

CCS CONCEPTS

• **Networks** → **Network measurement; Public Internet; Physical topologies; Network servers;**

KEYWORDS

Internet eXchange Points, Hypergiants, Content Delivery Networks, Netflix

1 INTRODUCTION

Originally designed as a research network, the Internet has evolved into a massive-scale platform for multimedia delivery. This transformation has been possible thanks to many underlying technical evolutions and innovations, stretching the Internet way beyond its original design. In this paper, we focus on two such shifts that are dramatically impacting the way the Internet operates today. First, a topological flattening has been observed [15], driven partly by the expansion of Internet Exchange Points (IXPs). These IXPs commoditise the interconnection of networks [16], and significantly lower the cost of network operations. Previous studies uncovered a rich and varied network ecosystem inside an IXP, so large that it fundamentally questions our current knowledge of the AS-level topology [1]. Second, consumption of online content, especially video material, has steadily grown, sparking the deployment of content delivery infrastructures deep inside the network, e.g., ISP caches, on a global scale.

When combining the above two observations, we begin to see a greater emphasis on traffic being generated and exchanged locally, rather than following the traditional hierarchy. This process, led by so-called hypergiants [16] (e.g., Google, Facebook), has radically altered the location of network “hot spots”, reducing the importance

of the traditional tier-1 networks and re-asserting the edge as the principal playground for innovation. Although previous studies have shown that individual IXPs are important for today’s network interconnection landscape [1, 16], there yet is no thorough analysis of the role the IXP ecosystem plays to support major content delivery players.

One of these major players or hypergiants is Netflix. Since 2012, Netflix has been deploying its own content delivery infrastructure, named Open Connect. It relies on server locations near the edge, strategically located close to its user base. In contrast to other hypergiants (e.g. Google, Facebook), Netflix operates neither a backbone network¹ nor datacenters [20]. Instead Netflix pre-loads content on its servers during off-peak times to reduce the need for transit traffic².

In this paper, we have performed the first large-scale measurement study of the Open Connect infrastructure. Using a range of techniques, we have discovered servers present at locations around the world. Using location information provided in the server names, we study the regional footprints of the deployed infrastructure and expose a variety of regional Internet ecosystems. Our results not only reveal the dependence that Netflix has on these regional ecosystems, but also highlight the combined ability of the many IXPs world-wide to deliver huge amounts of traffic on a local scale. They bypass the traditional tier-1 and transit networks, thus underpinning the fact that hypergiants like Netflix contribute to the flattening of the Internet.

To summarise, in this paper we make the following contributions:

- (1) We describe the infrastructure deployment of a content hypergiant (Netflix), which delivers large amounts of traffic from over 500 locations world-wide.
- (2) We provide evidence for the vastly understated ability of the many IXPs world-wide to deliver large amounts of traffic on a global scale: The world-wide footprint of IXPs enables Netflix to operate a global content delivery system, with very limited transit traffic, and without operating a backbone or owning datacenters.

An accompanying technical report detailing aspects of this paper is available online [5].

¹<https://openconnect.netflix.com/en/peering-locations/>

²<https://openconnect.netflix.com/en/fill/>

2 METHODOLOGY

In this section we describe the methodology we use to discover servers deployed by Netflix. We briefly describe the relevant implementation details of the Open Connect infrastructure (§2.1), describe the actual collection process in detail (§2.2) and validate the obtained data (§2.3).

2.1 Open Connect Infrastructure

Netflix uses Amazon Web Services (AWS) for most of its computing tasks. Such computing tasks for example are serving the website, the main application logic and the recommendation system, but also tasks related to video pre-processing and transcoding. The actual video content however is exclusively delivered through Netflix's own CDN Open Connect [20]. It is only this delivery infrastructure that we examine in this study.

To better understand how individual video clients are assigned to content servers, we ran a measurement campaign using HTTP proxies from a multitude of vantage points. We used the browser plugin Hola for this, which gave us vantage through 753 different IPs in 94 ASes. When a client requests a video file, the main application logic directly instructs the client which content servers to use. It (typically) hands out three domain names. The client then directly requests the video content from these servers via HTTPs.

The server names are very specific. They include information on the physical cache location and a cache number. This detailed naming structure makes it unlikely that names resolve to more than one IP address. This is consistent with what Netflix publishes on the naming convention of servers³. Nevertheless, we used Planetlab to confirm that each name only resolves to a single and always the same IP address, independent of the client's location. These findings, although more detailed, are in line with what Netflix publishes on how client redirection works⁴. Examples of server names used by Netflix are shown in Figure 1.

```
ipv4_1-lagg0-c020.1.lhr001.ix.nflxvideo.net  
ipv6_1-lagg0-c002.1.lhr005.bt.isp.nflxvideo.net
```

Figure 1: Examples of Netflix server names.

We conjecture that the meaning of the individual components of a name are as described in Figure 2. We will revisit the correctness of these assumptions later in this section.

```
ipv4 / ipv6: IP protocol version.  
lagg0: Type of network card. We also found other NICs  
(i.e., cxgbe0, ixl0, mlx5en0, mce0).  
c020: Server counter for a given location.  
lhr001: IATA airport code of a location with counter.  
bt.isp / ix: Network (type) identifier; server operated in-  
side ISP British Telecom or at an IXP
```

Figure 2: Components of a Netflix server name.

³<https://openconnect.netflix.com/en/portal-naming>

⁴<http://oc.nflxvideo.net/docs/OpenConnect-Deployment-Guide.pdf>

For the remainder of this paper, we will use the IATA airport code to infer the physical location of a server and the network identifier to distinguish between ISP and IXP servers⁵. Whenever we refer to the location of a server, we will use the airport code only without the counter, i.e., for three servers deployed at *lax001*, *lax002* and *lax003*, the location will be *lax* only, and the location *lax* will have three servers deployed.

2.2 Crawling DNS

To unveil the Open Connect network, we use a DNS crawler which enumerates and tries to resolve all domain names matching the above scheme. If a domain name can be resolved to an IP address, we assume that we found a Netflix server.

Note that ignoring the structured nature of the names and simply iterating over all possible character sequences is practically infeasible and not desirable.⁶ To narrow down the search space and limit the load on the DNS servers, the crawler is fed with lists of airport codes and ISP names, so that only DNS names for valid airports codes and ISPs are constructed. We further limit the number of probed DNS names, if no IP address is retrieved for a specific location and network operator. We also rely on DNS server behaviour standardised in RFC8020 [4] to prune empty DNS subtrees with a single query.

We used the following data sources to generate the input lists of airport codes and ISP names fed to the crawler:

Wikipedia We relied on Wikipedia to compile a list of IATA airport codes. While Wikipedia also has information on ISPs, extracting this information from Wikipedia is much more cumbersome, as it is spread across many pages and summary pages often are not updated frequently. We thus used additional sources to compile a list of ISPs.

Certificate Transparency In the specific case of Netflix, we can leverage the Certificate Transparency (CT) project, to generate a list of relevant ISP names. The Google-driven project aims to increase Internet security by providing datastores of all issued SSL/TLS certificates, which are distributed amongst independent entities and cryptographically secured [17]. These datastores allow individuals to verify certificate issuance. They can be used, for example, to detect rogue certificates issued without a genuine certificate request. The peculiarity of Netflix to use subdomains for the airport code and network (type) identifier, requires their servers to use separate SSL/TLS certificates for each server location⁷. These certificates are committed as individual log entries to CT. We can use these log entries to infer ISP names and airport codes used by Netflix. In addition, Google, through the CT project, discovered a non-authorised pre-certificate for its domains issued by Symantec's Thawte CA [3]. As a consequence, Google requested Symantec to log all issued

⁵Netflix does not distinguish between public IXPs and private peering facilities, but qualifies both as IXPs via the 'ix' part of the server names. This is reasonable if both options are viewed as just a means for delivering traffic. For the remainder of this paper we will adopt this view as well.

⁶Assuming an alphabet of 26 characters plus '.', '-', '_' as special characters and a prefix length of at least 30 characters (c.f. Fig. 1), enumerating all 29³⁰ possible combinations in one year's time would require roughly 2³⁶ DNS queries per second.

⁷A wildcard SSL/TLS certificate issued for *.nflxvideo.net will not be accepted as valid for the actual server domains.

	ISP	IXP	total
Servers	4,152	4,340	8,492
Locations	569	52	578
ASNs	743	1	744
ISP names	700	-	-

Table 1: Data Set Overview.

certificates with CT. As Netflix uses Symantec certificates for all its video delivery servers, we expect the CT logs to have complete coverage on the certificates used by Netflix’s video delivery servers.

Peering DB To cope with the unlikely event that an ISP is not discoverable by using certificate logs as outlined above, we extracted all network names from PeeringDB. We used these names and all subsets of them as possible inputs for our ISP list.

Unless explicitly stated otherwise, the data used in this paper was collected on May 15 2017.

2.3 Data Validation

To complement our CT logs ground truth, we can use a map by Netflix of their Open Connect infrastructure, published in a blog entry [21] dating from March 2016. Our measurements are highly consistent with this map. A comparison of the two makes it obvious that we in general observe the same global coverage and relative weight of individual regions. However, our measurements were done more than a year after Netflix’s data release and show significant additions and developments in certain regions. Netflix’s data only reveals qualitative information, while our measurements yield quantifiable results. Furthermore, our measurements identify the ISP networks where Netflix has deployed servers.

All in all, we are confident that we observed a complete enough part of Netflix’s video delivery infrastructure, allowing us to draw conclusions for those regions of the world, in which Netflix has a significant presence. For the following sections we will thus treat our data as a ground truth on Open Connect.

3 THE LOCAL ECOSYSTEMS OF THE INTERNET

In this section, we describe the infrastructure deployment by Netflix in more detail. Our goal is to illustrate the diversity of the various local ecosystems that are part of the Internet and assess the role of IXPs in each ecosystem. We look at the largest deployments of Netflix servers in each continent, and expose different types of deployments in terms of relative importance of ISP and IXP footprint.

We start our sample of local ecosystems with the largest market of Netflix, the USA (§3.2). We follow with an emerging, though already large, market for Netflix, Brazil (§3.3).

3.1 Data Overview

An overview of the gathered data set is shown in Table 1. In total we discovered 8,492 servers, of which 4,340 (51%) are deployed within IXPs and 4,152 are deployed in ISPs. We observe servers at 569 different ISP and 52 different IXP locations, where a single location

is a single airport code (see also §2.1). Our measurements reveal servers inside 700 different ISPs. While the IPs of all IXP servers are announced by the same AS (ASN 2906, Netflix), the IPs of the ISPs servers are announced by 743 ASs (which is more than the number of ISPs we observe). This happens because some ISPs use multiple AS numbers.

Comparing the sheer number of ISP networks versus the relatively fewer IXPs where Netflix servers are deployed, we can already conclude that Netflix strategically chooses the IXPs where it is present, which are relatively few in numbers. This is in contrast to ISP deployments, where its servers are scattered across hundreds of ISPs. From this, we can expect very different granularities in Netflix IXP and ISP deployments, with fine-grained deployment in ISPs, while IXP deployments are likely to be more significant in terms of number of servers.

These different granularities also appear when looking at the geographical footprint of Open Connect. Figure 3 shows the server locations on a world map. Green dots indicate an IXP server location, blue dots indicate an ISP server location. The marker sizes are scaled by the number of servers at a location. Although Netflix offers its service globally, its servers are predominantly present in Western countries, their deployment mostly focuses on the Americas and Europe, and to a smaller extent on Australia.

The largest deployment, by far, with 4,253 servers is in the US, followed by 901 servers in Brazil and 565 servers in the Canada. The United Kingdom and Mexico complete the top five countries. The complete data set with server counts for all countries is available upon request.

3.2 USA

We begin our look at local ecosystems with the United States of America. USA is the region with most Netflix customers by far [22], and is supported by the largest server deployment of any country. Netflix has 3,246 IXP and 1,007 ISP servers deployed in the USA. Those servers are spread across 24 IXP and 205 ISP locations, reaching into 211 different ISPs.

We look first at the IXP deployment, given its numeric dominance for delivering content in the USA. Such a strong IXP deployment is perhaps surprising, given that according to the public information from PeeringDB, we find no American IXP in the top five of largest IXPs world-wide in terms of either members or capacity. Nonetheless, there is a significant number of IXPs across the country. Netflix has taken advantage of this footprint, and is present at 24 IXP locations (as identified by airport codes). The deployment covers the major metropolitan areas, picking the largest US IXPs according to PeeringDB member count. Netflix is present at nine of the ten largest IXPs in the USA, and 15 of the largest 20.

Netflix’s deployment at IXPs typically involves a significant number of servers, whereas deployment inside ISPs is more fine-grained. We encounter IXP deployments at 24 different locations, the largest consisting of 360 servers. For ISPs, the largest deployment in a single location consists of a mere 14 servers. However, ISP servers are installed at 205 locations in total. Deployment at ISPs therefore appears to complement the geographical reach of the IXP deployment, over a higher number of locations, but with relatively small deployment sizes at each location compared to IXP ones.

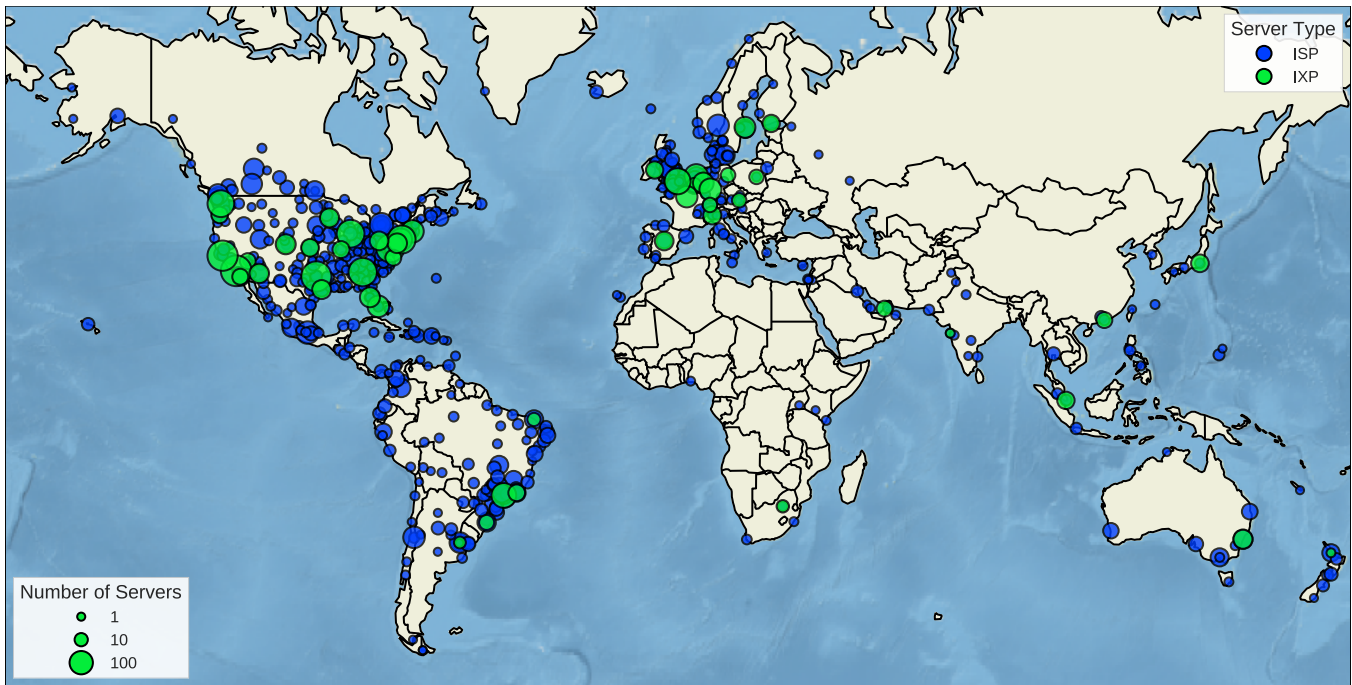


Figure 3: Netflix server deployment. Each marker denotes a location, the marker sizes are scaled by deployment size.

Note the absence⁸ of Netflix deployment inside four major ISPs (AT&T, Comcast, Time Warner Cable and Verizon⁹), as shown in Table 2. The explanation for this absence is that these ISPs publicly refused to deploy Netflix servers. Instead, they insisted on signing paid peering contracts with Netflix [19, 23]. This makes sense given the strong position of these ISPs in the US market.

ISP servers in the USA are hosted by smaller players. When contrasting ISPs with the Netflix ISP Speed Index¹⁰, we observe that those ISPs which do not deploy servers provide similar performance results as those which have Netflix servers deployed. This suggests that deploying Netflix servers inside an ISP network does not automatically imply better performance, at least according to Netflix’s own ISP Speed Index.

In summary, we observe that the USA has an IXP ecosystem mature enough, so that the available IXPs are sufficient for Netflix to rely primarily on IXPs to reach its large customer base. This comes in as a surprise, given that based on research literature little is known about the US IXP ecosystem, especially in comparison to the European one [12]. Furthermore, relying on IXP deployments, and not having deployments inside some ISPs, does not appear to have negative consequences on performance as reported by Netflix, highlighting again the usability of IXPs for large-scale content delivery.

⁸To discard the possibility of a measurement error, we included all reasonable abbreviations of these ISP names as input for the DNS crawler. However, even after this extensive search, we could not discover further servers.

⁹We discovered three Netflix servers in Verizon’s network, which do not offer a significant advantage in traffic savings for such a large network, but might be part of a trial.

¹⁰<http://ispsspeedindex.netflix.com/>

USA			
AT&T	-	Mediacom	-
Bright House	-	Optimum	-
CenturyLink	113 / 11	Spectrum	-
Charter	-	Suddenlink	68 / 31
Comcast	-	TWC	-
Cox	-	Verizon	3 / 2
Frontier	19 / 3	Windstream	31 / 11

Table 2: Netflix servers deployed inside US ISPs. ISPs are taken from Netflix’s ISP Speed Index. The left number denotes the number of servers in an ISP, the right one the number of locations those servers are deployed at. ISPs listed multiple times in the index (e.g., due to different broadband connection types), are listed only once in this table.

3.3 Brazil

Our second chosen local ecosystem is Brazil. This emerging market has the second largest Netflix server deployment despite not being an English-speaking country. Netflix offers in Brazil a substantial selection of content with at least subtitles in Portuguese for a fraction of the cost of cable TV. The deployment consists of 901 servers, 713 servers inside ISPs and 188 servers at IXPs. Unlike the USA, servers in Brazil are primarily located inside ISPs. ISP servers are deployed inside 187 ISPs, covering 58 locations across the vast Brazilian geography, but mostly along the Eastern coastal regions where most people live.

In strong contrast to the USA, IXP servers are only deployed at 3 locations on the South East Coast (São Paulo, Rio de Janeiro, and Porto Alegre) and at one location on the North East Coast

(Fortaleza). In Brazil, Netflix has a limited IXP server deployment (in both locations and number of servers), despite a reasonably large number of available IXP locations (25 locations in total according to [6], in the USA Netflix uses 24 IXP locations (§3.2)). Deploying servers in IXPs has to be more cost efficient for Netflix due to economies of scale and a simpler contractual situation with fewer parties involved compared to ISP deployment. The observed deployment suggests an IXP ecosystem with limited capacity to reach Netflix customers. This limitation might be caused by a multitude of factors, including lack of capacity on the IXP switching fabric or the inability to host additional servers at IXPs.

According to PeeringDB data, the three IXPs on the South East Coast Netflix is present at, are also the largest ones, in terms of number of members. The IXP in Fortaleza is the seventh largest in Brazil. The Brazilian IXP infrastructure is developed by IX.br, a non-profit initiative. IX.br explicitly aims to improve the Internet connectivity deficiencies of the north, west and central regions, by providing a collection of exchange points. However, we see that Netflix only uses the IXP facilities at 3 (São Paulo, Rio de Janeiro, Fortaleza) of the 5 largest metropolitan areas, all located on the East coast. The vast majority of IXPs in Brazil have a small number of peers, and more importantly lack content providers, and private companies except in the South East [6].

Brazil has a developing Internet infrastructure. External metrics such as the Netflix Speed Index figures¹⁰ show much lower bandwidth figures compared to the other top Netflix markets. Whereas IXPs by nature aim at fostering local access ecosystems, the edge Internet infrastructure must be strong enough for service providers to operate purely from these exchange points. Otherwise, deployment inside ISPs seems necessary.

3.4 Summary

In this section, we illustrated two local Internet ecosystems, as seen through Netflix's server deployment. Our choice of local ecosystems has shown how the specifics of each local ecosystem translate into very different outcomes in terms of server deployment. We observed that ecosystems where developed IXPs are available typically lead to significant IXP server deployment. However, we also observed that to reach a large customer base, which is geographically scattered, ISP deployment is often necessary to compensate for the limited footprint of the local IXPs.

4 DISCUSSION

In this section we will discuss our most important findings regarding the current state of the IXP ecosystem and its usability as a base for content delivery.

One peculiarity of the way Netflix delivers its content, is that, in contrast to the other big video players by traffic volume (YouTube and Amazon Video), it does so without operating a backbone network¹¹. To reach its customers, Netflix instead relies on deploying servers at IXPs and inside ISPs. These deployment sites form self-sufficient islands, capable of serving the local customer demand more or less independently. Netflix's pre-fetching approach to populate content on its servers is key to reduce the amount

of transit traffic, i.e., traffic between the servers holding the original content and the copies placed on the deployment sites. The backbone-less and light in transit approach of Netflix contributes to the observed phenomenon of Internet flattening. Instead of flowing through the traditional Internet hierarchy (tier-1s), Internet traffic goes through more and larger direct interconnects between networks at the edge. To deliver its traffic, Netflix chooses IXP locations, as well as ISPs that are not in the traditional core of the Internet, therefore bypassing the traditional Internet hierarchy and inevitably contributing to the observed flattening.

The case of Netflix demonstrates that large-scale traffic delivery from edge locations (esp. IXP locations) is possible. We believe that reporting this approach followed by Netflix is important, as it illustrates its feasibility, but also the challenges that come with it, in terms of being able to exploit the very different local ecosystems of the Internet. This will hopefully inspire other small and large players to follow a similar approach, at least for some parts of their content, which then may in turn exacerbate the flattening phenomenon.

Netflix not only does not operate a backbone, but it nowadays also does not operate a single datacenter either¹¹. Instead, Netflix serves its traffic from servers deployed in colocation housing locations at or in close proximity to IXPs. These locations allow Netflix to operate without its own datacenters, as those locations essentially provide all the features of a regular datacenter. One drawback of such an approach is the space restrictions in these locations that might limit their usability for large deployments. Nevertheless, for Netflix's needs focused on data storage and data transfer, not operating its own datacenters seems to work. To our knowledge, it is the first time such a worldwide deployment is exposed, based on a strategic use of IXP facilities as a datacenter replacement. From this, we learn that the benefit of IXPs is not limited to network interconnection [1], but that they also facilitate the deployment of large server bases at locations with strategically beneficial network connectivity.

5 RELATED WORK

As one of the major players in video content delivery, Netflix's role in the Internet directly illustrates the observations from Labovitz et al. [16], back in 2010. Indeed, Labovitz et al. [16] observed a new trend, whereby traffic was seen to flow directly between large content providers, datacenters, CDNs and consumer networks, away from large transit providers. Subsequent studies investigated the potential implications of more direct interconnections on the Internet [9, 13, 14, 18].

Due to the success of players such as Netflix, the rise in video traffic observed by Labovitz et al. [16] has only continued. Our study of the server deployment of Netflix at the edge of the Internet, and the corresponding traffic delivered to end-users, makes the observations of Labovitz et al. [16] even more relevant today. Despite their importance in the Internet ecosystem, only a few studies have targeted IXPs [1, 2, 8, 10–12] and their role in the Internet. The work from Augustin et al. [2] aimed at systematically mapping IXP infrastructures through large-scale active measurements, leading to the first evidence of the huge number of IXPs around the world. Ager et al. [1] studied the ecosystem and traffic of one

¹¹<https://openconnect.netflix.com/en/peering-locations/>

of the largest European IXPs, while Restrepo et al. [8] looked at two smaller European IXPs. Subsequent studies from Chatzis et al. [10–12] reinforced the critical role played by IXPs in the Internet ecosystem.

IXPs are a major component supporting the peering ecosystem of the Internet. To this day, however, the role of IXPs world-wide in supporting the delivery of large amounts of traffic close to end-users has been understated. Indeed, despite the large number of IXPs known to exist [2], the largest of them having hundreds of members and delivering daily traffic volumes in the petabyte range, their relative importance for content delivery was largely unreported. In this work, we uncovered the importance that IXPs play in enabling a player such as Netflix to deliver its traffic to its large and worldwide customer base. We observed that despite preferring to deploy servers within ISP networks, a majority of Netflix servers exploit the strategic location and ecosystems provided by IXPs all around the world. Labovitz et al. [16] indicated a significant shift in the mental map of the Internet, with traffic being increasingly delivered directly between large content providers and consumer networks, away from large transit providers. Our work adds another piece of evidence for this shift, with a direct observation of a large video delivery provider doing this by strategically exploiting the rich ecosystem that many IXPs provide.

Mapping the server deployment and expansion of a large content player has been done before. Calder et al. [7] developed techniques that enumerate IP addresses of servers of the Google infrastructure, found their geographic location, and identified the association between clients and clusters of servers. To do this accurately, they use the EDNS-client-subnet DNS extension to measure which clients a service maps to which of its serving sites. Different from our work, Calder et al. [7] focused on the accuracy of the server mapping and geolocation, necessary given the size and complexity of the Google infrastructure. In this paper, we focus on the types of locations where Netflix has chosen to deploy its server infrastructure. Overall, we are not overly concerned with the mapping of the servers itself, as Netflix runs a single service, contrary to Google. Rather, our focus is on the implications of Netflix’s server deployment strategy, with the lens it provides on the Internet ecosystem.

6 SUMMARY

In this work, we studied the global footprint of one content hyper-giant, Netflix, to gain a new perspective on the current Internet. We exposed the approach used by Netflix to deliver massive amounts of traffic from over 500 world-wide locations with neither a backbone nor datacenters. It does so by deploying its own servers at IXP locations as well as in ISP networks. By studying the deployment of its servers, we highlighted regional differences in the deployment, by sampling the diversity of local ecosystems that collectively make up the Internet. The Netflix lens provides evidence for the vastly understated ability of the many IXPs world-wide to deliver large amounts of traffic on a global scale. The world-wide footprint of IXPs is the major corner stone of Open Connect and enables Netflix to operate a global content delivery system, with very limited transit traffic, and without operating a backbone or owning datacenters.

ACKNOWLEDGMENTS

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

The authors would like to thank Robert Edmonds for his valuable comments on how to employ RFC8020 [4]. The authors further would like to thank Gea-Suan Lin for the valuable comments on how to employ the information made available by the Certificate Transparency project [17].

REFERENCES

- [1] Bernhard Ager, Nikolaos Chatzis, Anja Feldmann, Nadi Sarrar, Steve Uhlig, and Walter Willinger. 2012. Anatomy of a Large European IXP. In *Proc. of ACM SIGCOMM*.
- [2] Brice Augustin, Balachander Krishnamurthy, and Walter Willinger. 2009. IXPs: Mapped?. In *Proc. of ACM IMC*.
- [3] Google Security Blog. 2015. Sustaining Digital Certificate Security. (2015). Retrieved January 11, 2018 from <https://security.googleblog.com/2015/10/sustaining-digital-certificate-security.html>
- [4] Stéphane Bortzmeyer and Shumon Huque. 2016. NXDOMAIN: There Really Is Nothing Underneath. RFC 8020. (Nov. 2016). <https://doi.org/10.17487/RFC8020>
- [5] Timm Böttger, Félix Cuadrado, Gareth Tyson, Ignacio Castro, and Steve Uhlig. 2016. Open Connect Everywhere: A Glimpse at the Internet Ecosystem through the Lens of the Netflix CDN. (2016). arXiv:1606.05519
- [6] Samuel Henrique Bucke Brito, Mateus AS Santos, Ramon dos Reis Fontes, Danny A Lachos Perez, and Christian Esteve Rothenberg. 2016. Dissecting the Largest National Ecosystem of Public Internet eXchange Points in Brazil. In *Proc. of PAM*.
- [7] Matt Calder, Xun Fan, Zi Hu, Ethan Katz-Bassett, John Heidemann, and Ramesh Govindan. 2013. Mapping the Expansion of Google’s Serving Infrastructure. In *Proc. of ACM IMC*.
- [8] Juan Camilo Cardona Restrepo and Rade Stanojevic. 2012. IXP Traffic: A Macroscopic View. In *Proc. of ACM LANC*.
- [9] Ignacio Castro, Juan Camilo Cardona, Sergey Gorinsky, and Pierre Francois. 2014. Remote Peering: More Peering Without Internet Flattening. In *Proc. of ACM CoNEXT*.
- [10] 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 *Proc. of ACM IMC*.
- [11] Nikolaos Chatzis, Georgios Smaragdakis, Anja Feldmann, and Walter Willinger. 2013. There is More to IXPs Than Meets the Eye. *ACM CCR* 43, 5 (Nov. 2013).
- [12] Nikolaos Chatzis, Georgios Smaragdakis, Anja Feldmann, and Walter Willinger. 2015. Quo Vadis Open-IX? *ACM CCR* 45, 1 (Jan. 2015).
- [13] Amogh Dhamdhere and Constantine Dovrolis. 2010. The Internet Is Flat: Modeling the Transition from a Transit Hierarchy to a Peering Mesh. In *Proc. of ACM CoNEXT*.
- [14] Amogh Dhamdhere and Constantine Dovrolis. 2011. Twelve Years in the Evolution of the Internet Ecosystem. *IEEE/ACM ToN* 19, 5 (2011), 1420–1433.
- [15] Phillipa Gill, Martin Arlitt, Zongpeng Li, and Anirban Mahanti. 2008. The Flattening Internet Topology: Natural Evolution, Unlikely Barnacles or Contrived Collapse? In *Proc. of PAM*.
- [16] Craig Labovitz, Scott Iekel-Johnson, Danny McPherson, Jon Oberheide, and Farnam Jahanian. 2010. Internet inter-domain traffic. *Proc. of ACM SIGCOMM*.
- [17] Ben Laurie, Adam Langley, and Emilia Kasper. 2013. Certificate Transparency. RFC 6962. (June 2013). <https://doi.org/10.17487/RFC6962>
- [18] Richard TB Ma, John Lui, and Vishal Misra. 2015. Evolution of the Internet Economic Ecosystem. *IEEE/ACM ToN* 23, 1 (2015), 85–98.
- [19] Netflix. 2014. Internet Tolls And The Case For Strong Net Neutrality. (2014). Retrieved January 11, 2018 from <https://media.netflix.com/en/company-blog/internet-tolls-and-the-case-for-strong-net-neutrality/>
- [20] Netflix. 2016. Completing the Netflix Cloud Migration. (2016). Retrieved January 11, 2018 from <https://media.netflix.com/en/company-blog/completing-the-netflix-cloud-migration/>
- [21] Netflix. 2016. How Netflix Works With ISPs Around the Globe to Deliver a Great Viewing Experience. (2016). Retrieved January 11, 2018 from <https://media.netflix.com/en/company-blog/how-netflix-works-with-isps-around-the-globe-to-deliver-a-great-viewing-experience>
- [22] Digital TV Research. 2014. Number of Netflix paying streaming subscribers in 3rd quarter 2014. (2014). Retrieved January 11, 2018 from <https://www.statista.com/statistics/324050/number-netflix-paying-streaming-subscribers/>
- [23] Time. 2014. Verizon Won’t Use Netflix’s Hardware to Boost Streaming Speeds. (2014). Retrieved January 11, 2018 from <http://time.com/2866004/verizon-netflix/>