



HAL
open science

The Central Problem with Distributed Content

Kevin Vermeulen, Loqman Salamatian, Sang Hoon Kim, Matt Calder, Ethan Katz-Bassett

► **To cite this version:**

Kevin Vermeulen, Loqman Salamatian, Sang Hoon Kim, Matt Calder, Ethan Katz-Bassett. The Central Problem with Distributed Content: Common CDN Deployments Centralize Traffic In A Risky Way. HotNets: Twenty-Second ACM Workshop on Hot Topics in Networks, Nov 2023, Cambridge, United States. 10.1145/3626111.3628213 . hal-04273837

HAL Id: hal-04273837

<https://hal.science/hal-04273837>

Submitted on 7 Nov 2023

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

The Central Problem with Distributed Content

Common CDN Deployments Centralize Traffic In A Risky Way

Kevin Vermeulen
LAAS-CNRS

Loqman Salamatian
Columbia University

Sang Hoon Kim
Columbia University

Matt Calder
Columbia University

Ethan Katz-Bassett
Columbia University

ABSTRACT

Google, Netflix, Meta, and Akamai serve content to users from offnet servers in thousands of ISPs. These offnets benefit both services and ISPs, via better performance and reduced interdomain and WAN traffic. We argue that this widespread distribution of servers leads to a concentration of traffic and a previously unacknowledged risk, as many ISPs colocate offnets from multiple providers. This trend contributes to many Internet users likely accessing multiple popular services and fetching the majority of their Internet traffic from a single facility – perhaps even a single rack – creating shared resources and a correlated risk in cases of failures, attacks, and overload. Alternate ways to access the services often lack sufficient capacity and share resources with more services, creating the potential for cascading failures.

CCS CONCEPTS

• **Networks** → **Network reliability; Network measurement; Network structure.**

ACM Reference Format:

Kevin Vermeulen, Loqman Salamatian, Sang Hoon Kim, Matt Calder, and Ethan Katz-Bassett. 2023. The Central Problem with Distributed Content: Common CDN Deployments Centralize Traffic In A Risky Way. In *The 22nd ACM Workshop on Hot Topics in Networks (HotNets '23)*, November 28–29, 2023, Cambridge, MA, USA. ACM, New York, NY, USA, 9 pages. <https://doi.org/10.1145/3626111.3628213>

1 INTRODUCTION

Google, Netflix, Meta, and Akamai have *offnet* servers in many ISPs, serving content to the ISPs' users and customers. The offnets benefit users, services, and ISPs: better performance and reduced interdomain and WAN traffic.

Publication rights licensed to ACM. ACM acknowledges that this contribution was authored or co-authored by an employee, contractor or affiliate of a national government. As such, the Government retains a nonexclusive, royalty-free right to publish or reproduce this article, or to allow others to do so, for Government purposes only.

HotNets '23, November 28–29, 2023, Cambridge, MA, USA

© 2023 Copyright held by the owner/author(s). Publication rights licensed to ACM.

ACM ISBN 979-8-4007-0415-4/23/11.

<https://doi.org/10.1145/3626111.3628213>

We argue that this widespread distribution of servers (paradoxically?) leads to a concentration of traffic and a previously unacknowledged risk. This trend contributes to a scenario in which many Internet users likely access popular multiple services and fetch the majority of their Internet traffic from a single facility – perhaps even a single rack – creating shared resources and correlated risk in cases of failures, attacks, and overload. Alternate paths for accessing the services often lack sufficient capacity and share resources with even more services, creating the potential for cascading failures.

More and more Internet users receive an ever-increasing amount of Internet traffic from *offnet* servers in access or transit networks (collectively, ISPs), rather than from servers in content or cloud provider data centers. This access pattern results from two trends. First, content is consolidating, as a small number of services hosted by an even smaller number of providers are responsible for a growing fraction of Internet content – especially when considering the traffic volume stemming from popular video services. Second, some of the largest of these providers have offnets deployed in a growing number of user networks. Akamai used to be the only provider with a large offnet footprint. Google followed suit over a decade ago [12]. In recent years, Netflix and Meta deployed offnets in thousands of ISPs [21].

These *hypergiants* follow the same high-level approach for deploying offnet servers. If an ISP meets criteria such as traffic demand and hosting capabilities [18, 25, 43], the hypergiant may agree to deploy an offnet. The hypergiant will supply the server, which the ISP hosts in a facility and supplies with power and network connectivity. The ISP provides the hypergiant with a BGP feed of IP prefixes it is willing to serve from the offnet. The hypergiant may direct requests from those prefixes to fetch content from the offnet.

These offnets are often colocated. First, the offnets host content from popular services with similar and overlapping user bases, and so ISPs often host offnets from multiple hypergiants. Measurements from 2021 revealed that, of the 4500 ISPs that hosted an offnet from at least one of Google, Netflix, Meta, and Akamai, more than 60% hosted offnets from multiple providers, and ISPs tended to host more hypergiants over time [21]. Second, for various reasons of efficiency, an

ISP hosting multiple offnets has reasons to collocate them. Our measurements reveal that 81-95% of such ISPs collocate offnets from multiple hypergiants.

This collocation increases the chances that servers from distinct hypergiants, but serving common users, might experience simultaneous problems. Facility-wide outages will impact all hosted servers. Congestion from one hypergiant’s offnet may impact routes shared with offnets from other hypergiants. This congestion can come from sudden increases in demand, as we will present evidence suggesting that offnets are running near capacity, with little ability to absorb sudden increases such as could be caused by flash crowds, DoS attacks, and bad updates of offnet software. In cases when an event impacts multiple locations of a hypergiant’s offnets, it may impact other hypergiants at those locations.

When an offnet is unreachable or overloaded, the excess demand can be served by other offnets, via dedicated links from hypergiants, via Internet Exchange Points (IXPs), or via transit providers. We will argue that evidence exists that these alternatives often have limited spare capacity. Further, IXPs and transit providers are resources shared with other services. In total, the high rates of collocation of the offnets providing some of the largest volume services on the Internet and the frequently limited available headroom of both the offnets and alternate content delivery paths create the potential for a perfect storm of overload and cascading failure.

2 MORE AND MORE INTERNET TRAFFIC COMES FROM OFFNETS

Akamai was an early leader in Internet content delivery, and its current deployment includes 350,000 offnet servers in 134 countries and over 1000 ISPs [3]. More recently, driven in large part by the explosion of streaming video, Google, Netflix, and Meta launched and rapidly expanded their own offnet deployments, with a 2021 paper revealing that these three hypergiants each had offnets in over 2000 ISPs [21].

2.1 Offnets serve large fractions of traffic

These 4 providers account for most Internet traffic! According to Sandvine Google serves 21% of Internet traffic, Netflix serves 9%, and Meta serves 15% [48]. Akamai claims to serve 15-20% of web traffic [51]. Much of this traffic comes from offnets. Offnet servers typically serve 70-90% of Google traffic [23]. Netflix claims its offnets serve 95% of its traffic [19]. One network reports that its Google offnets deliver ≈ 20 Gbps at peak per location (80% of its Google traffic), its Netflix offnets deliver ≈ 30 Gbps ($> 90\%$ of its Netflix traffic), its Meta offnets deliver ≈ 20 Gbps (86%), and its Akamai offnets deliver ≈ 20 Gbps (75%) [47]. (Given that this network’s percentages match overall claims from Google and Netflix, we suspect its Meta and Akamai numbers are reasonable.) This

| Hypergiant | # of ISPs with offnets | |
|------------|------------------------|---------------|
| | 2021/04 | 2023/04 |
| Google | 3810 | 4697 (+23.2%) |
| Netflix | 2115 | 2906 (+37.4%) |
| Meta | 2214 | 2588 (+16.9%) |
| Akamai | 1094 | 1094 (+0.0%) |

Table 1: # of ISPs hosting offnets in 2021 [21] and 2023.

network’s offnets deliver up to ≈ 90 Gbps, compared to < 15 Gbps coming from these hypergiants over interdomain links.

2.2 Offnet deployments continue to grow

We update the methodology and results from the 2021 paper that uncovered offnet footprints [21], showing that Google, Netflix, and Meta have expanded their footprints significantly in the last 2 years. That paper demonstrated how to uncover hypergiants’ servers by looking for their TLS certificates in Internet-wide scans of port 443. If an IP address of an ISP other than a hypergiant hosts a certificate of the hypergiant, then the IP address corresponds to an offnet server of the hypergiant, hosted in the ISP. That paper pointed out ways in which the methodology was fragile: “[hypergiants can modify their certificate content by altering fields that [the methodology] currently use[s] to infer ownership and to extract fingerprints,” listing possible modifications including that a hypergiant could (1) “remove the Organization entry from the Subject Name of the EE certificate” or (2) “use unique domain names per offnet deployment”.

Google now does not include the Organization entry (1), and Meta now uses different domain names for different offnet deployments (2), and so we modified the methodology to account for these changes. For Google, instead of inspecting the Organization subfield from Subject Name field, we use the CN (Common Name) field. If a TLS certificate’s CN field matches *.googlevideo.com and passes the other checks from the 2021 methodology, we consider the server to be a Google server. Meta began using site-specific names like CN=*.fhan14-4.fna.fbcdn.net and CN=*.fbhx2-2.fna.fbcdn.net (han is Hanoi, Vietnam, and bhx is Birmingham, UK). This naming convention means that offnets have different names than onnet servers. So, whereas the 2021 methodology looked for names that exactly match onnets, we check for the pattern *.fbcdn.net.

We apply this approach to a 2023 Censys IPv4 scan. We do not use HTTP fingerprints, which have little effect [21].

Results. We found 261K offnet IP addresses for Google, Netflix, Meta, and Akamai, across 5516 ISPs. Table 1 compares our results to 2021 results. Google, Netflix, and Meta have significantly expanded. Google went from having offnets in 3810 ISPs to 4697 ISPs (123.2%). Netflix expanded from 2115

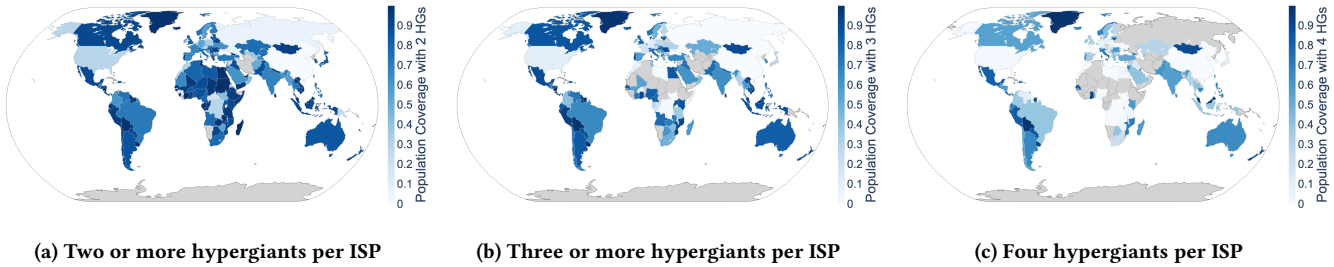


Figure 1: Per-country Internet user population in ISPs hosting offnets from multiple of Akamai, Google, Netflix, and Meta.

ISPs to 2907 ISPs (137.4%). Meta grew from 2214 ISPs to 2588 ISPs (116.9%). Akamai held at 1094 ISPs.

3 OFFNETS FROM MULTIPLE PROVIDERS ARE OFTEN COLOCATED

We argue that it is likely common for offnets for multiple hypergiants to be colocated (§3.1). Our measurements support this intuition by suggesting frequent colocation (§3.2). This colocation creates shared fate that opens the door to the potential for cascading failures (§3.3).

3.1 ISPs host multiple hypergiants and have reason to colocate the offnets

Services like YouTube, Netflix, and Instagram share overlapping user bases and hence deliver large amounts of traffic to overlapping sets of ISPs. These ISPs have incentives to host offnets to serve their clients, and so many ISPs that host offnets for one of the hypergiants are likely to have reason to host offnets for others. The results support this hypothesis. Of the 5516 ISPs that host an offnet for at least one of Google/Akamai/Meta/Netflix, 3382 host offnets for at least two, 1880 for at least three, and 505 host offnets for all four. This result indicates an increase in cohosting since 2021, when ≈ 2840 hosted at least two, ≈ 1690 hosted at least three, and ≈ 430 hosted all four [21]. This change and similar longitudinal results in the 2021 paper suggest that multi-hypergiant hosting will continue to increase over time.

Figure 1 shows world maps colored by the fraction of a country’s Internet users that are in ISPs that host offnets from two or more of Akamai, Google, Netflix, or Meta, using the APNIC ISP population dataset [27]. In many countries, the majority of Internet users are in ISPs hosting offnets from at least 2 hypergiants. In Europe and Africa, many fewer users are in ISPs that host 3 hypergiants (Figure 1b vs. Figure 1a), while other continents see more minor differences. In Figure 1c, a few countries have all or nearly all of their Internet users in ISPs that host all four hypergiants: Mexico, Bolivia, Uruguay, New Zealand, Mongolia, and Greenland. In these countries, most users can fetch content for many popular services from these local servers. These results likely

underestimate the use of offnets, which can also serve users downstream from a transit provider.

If an ISP hosts offnets from multiple hypergiants, there are good reasons to colocate them. Colocating servers that play a similar role (those provided by outside hypergiants to serve users) simplifies management and amortizes costs. Popular colocation facilities offer interconnection with multiple hypergiants, creating a situation where colocation becomes both convenient and a logical strategy for network optimization. Many smaller ISPs interconnect with other networks in only a single location and may situate offnets nearby to facilitate their cache fills. Larger ISPs will want to host offnets in locations convenient to the users they serve, minimizing the distance they need to carry the traffic to users.

3.2 Evidence of widespread colocation

Using the methodology we discuss next, we conducted measurements and preliminary analysis that revealed widespread colocation of offnet servers from multiple hypergiants. For example, in at least 46% of the ISPs hosting Netflix offnets, *all* Netflix servers are in facilities that also host servers from other hypergiants. Anecdotally, an operator at one of the hypergiants confirmed to us that not only are offnets from multiple providers often colocated in the same facility, but it is “super common” for them to be in the same rack.

We use an existing technique to cluster offnets based on the similarity of latency measurements from distributed sites, which validation showed could cluster colocated Google servers, including differentiating between multiple facilities in a city [12]. This technique uses OPTICS [5], which does not require the number or size of clusters a priori. We measured latencies to the 261K offnet IP addresses from the 163 M-Lab sites worldwide [22]. OPTICS takes a “steepness” parameter x_i , from 0 to 1, that determines how dense points must be to be considered a cluster [5]. We run the clustering twice per ISP, with two extreme values of 0.1 and 0.9, likely bounding the actual colocation. Appendix A has details.

| | x_i | Sole HG | Multiple HGs : % offnets colocated with another HG | | | |
|------------------------------------|-------|---------|--|-------------|------|-----|
| | | | 0% (0%, 50%) | [50%, 100%) | 100% | |
| % of ISPs that host the hypergiant | | | | | | |
| Google | 0.1 | 31% | 15% | 12% | 9% | 33% |
| | 0.9 | 31% | 2% | 2% | 3% | 62% |
| Akamai | 0.1 | 16% | 25% | 36% | 7% | 16% |
| | 0.9 | 16% | 7% | 4% | 15% | 58% |
| Meta | 0.1 | 6% | 23% | 27% | 12% | 32% |
| | 0.9 | 6% | 4% | 2% | 4% | 84% |
| Netflix | 0.1 | 12% | 21% | 10% | 11% | 46% |
| | 0.9 | 12% | 8% | 2% | 7% | 71% |

Table 2: A significant fraction of offnets are colocated with offnets from other hypergiants. Of ISPs that host multiple hypergiants, the columns bucket ISPs based on the % of offnets from the row’s hypergiant that are colocated with offnets from another hypergiant (buckets are {0%, (0%, 50%), [50%, 100%), 100%}). Each row sums to 100%.

Validation. To assess the accuracy of our collocation inferences, we employed the evaluation technique from the original paper by checking the consistency of location-related information contained in the hostnames of IP addresses within a cluster [12]. This validation is incomplete, as many IP addresses do not have reverse DNS entries, and many reverse DNS entries do not have obvious location information. However, this step can highlight inaccuracies in the clustering when IP addresses that appear to be hosted in different cities are colocated. To associate IP addresses with hostnames, we used Rapid7 PTR records [46]. To derive locations from the hostnames, we employed HOIHO [34] and focused on clusters with two or more IP addresses with identified locations. We manually corrected certain inaccuracies where HOIHO appeared to misinterpret hostnames (*e.g.*, it interpreted host as Hostert, LU). For $x_i = 0.1$, 60 clusters had two or more hostnames with identified locations. Within this subset, 55 clusters only included hostnames from a single city, an additional 3 included multiple locations within a single metropolitan area (*i.e.*, suburbs of London and Paris), and 2 included different cities in the same country. With the more conservative parameter setting of $x_i = 0.9$, our approach identified 34 clusters with two or more hostnames with identified locations. Within this subset, 30 clusters contained only hostnames from a single city, while 2 included multiple locations in a single metropolitan area, and 2 included multiple cities in the same country. These discrepancies within a cluster may be errors in clustering, errors in HOIHO location inferences, or stale/incorrect locations in hostnames [57].

Results. Many offnets are colocated with offnets belonging to other hypergiants. Table 2 shows, for each hypergiant, the percentage of ISPs hosting only that hypergiant, and the percentages of ISPs hosting that hypergiant and others, categorized by the fraction of the hypergiant’s offnet IP addresses from the ISP that are colocated with other hypergiants. There are two rows for each hypergiant, corresponding to $x_i = 0.1$ and $x_i = 0.9$, representing the uncertainty of our clustering

algorithm. Some ISPs only host one hypergiant: from 6% of those hosting Netflix to 31% of those hosting Google.

For ISPs that host multiple hypergiants, even though the exact numbers vary across clustering parameters and hypergiants, all cases support our claim that collocation of offnets from multiple hypergiants is common. Most ISPs colocate at least some offnets for all hypergiants – those that do not vary only from 2%–15% of ISPs for Google to 7%–25% for Akamai. A large percent of deployments are fully colocated: 33%–62% of ISPs host all their Google offnets with offnets from other hypergiants, 16%–58% of ISPs for Akamai, 32%–84% of ISPs for Meta, and 46%–71% of ISPs for Netflix.

In both clusterings, compared to the other hypergiants, a larger percentage of ISPs deploy some Akamai offnets colocated and some not (19%–45% of ISPs are neither no collocation nor full collocation). We hypothesize that this result reflects the fact that many Akamai deployments date from many years before the other hypergiants began deploying offnets and, hence, may predate current operational practices at the ISPs. In our most conservative clustering, Akamai has at least half of its offnets colocated with other hypergiants in $7% + 16% = 23%$ of ISPs. All other hypergiants and all other parameter settings suggest much more extensive collocation: from 42% of ISPs collocating at least half of Google offnets in the most conservative clustering to 88% of ISPs collocating at least half of Meta offnets in the less conservative clustering.

This level of collocation likely leads to some users fetching the majority of their Internet content from a single facility! With existing methodologies, it is impossible to know which users are served from which offnets. An earlier technique provided such results for Google in 2013 [12], but it only works if the hypergiant uses DNS to direct users to specific offnet locations for a given hostname such as `www.google.com`. Google no longer does so, and instead Google, Netflix, and Meta generally direct users to a particular offnet for cached content by embedding customized URLs into web pages returned to users (*e.g.*, `fhan14-4.fna.fbcdn.net`), while hosting their web pages on onnet and cloud locations. Akamai does use DNS to direct users to offnets, but it only accepts EDNS Client Subnet queries from allow-listed DNS resolvers and so is not compatible with the earlier technique.

Since we cannot know exactly which users are served from a facility hosting offnets, for each ISP we focus on the facility hosting the most hypergiants and estimate the fraction of traffic it serves for the (possibly subset of) users in the ISP that it serves. We use the estimates of Internet traffic share and caching efficiency from Section 2.1 to estimate that a Google offnet cache can serve $21%$ (% of total Internet traffic that is from Google) \times $80%$ (% of Google traffic that an offnet can serve) = $17%$ of the total Internet traffic for clients it serves, and a Netflix offnet can serve $9% \times 95% = 9%$ of a client’s traffic. Similarly,

a Meta offnet can serve $15\% \times 86\% = 13\%$ of a user’s traffic, and an Akamai offnet can serve $17.5\% \times 75\% = 13\%$ of a user’s traffic. A facility hosting all four hypergiants can serve $17\% + 9\% + 13\% + 13\% = 52\%$ of a user’s traffic!




Figure 2: For users in ISPs with offnets, there is usually a facility that can potentially serve much of their traffic.

Based on these estimates, Figure 2 shows the share of traffic per Internet user that potentially comes from the same facility (giving the range as the clustering parameter x_i varies). According to the APNIC dataset estimating the number of users on the Internet per ISP [27], 56% of Internet users are in ISPs with an offnet from at least one of the four hypergiants. Of these users, 71%–82% are in an ISP with a facility that hosts multiple hypergiants’ offnets collectively capable of delivering at least 25% of their traffic, and 18%–31% (representing 10%–17% of all Internet users) are in ISPs with a facility that hosts offnets from all four hypergiants. For these users, 52% of their traffic could be coming from a single facility! The actual numbers are likely higher as we do not include ISPs served from offnets in their providers.

3.3 Colocation creates correlated risk

Risks become correlated when multiple hypergiants are collocated. Shared physical infrastructure, such as power and cooling systems, introduces mutual vulnerability to outages. Traffic surges from one hypergiant might monopolize the available bandwidth, inadvertently impeding other hypergiants. Such surges could be caused by flash crowds, misconfigurations, or denial of service attacks. The interconnected ecosystem in which offnets are hosted can potentially cause cascading failures, where a single problem ripples across and then out of the entire facility, causing widespread disruption to users. These same facilities will increasingly host edge computing for critical and performance-sensitive applications. As these applications often demand high availability and low latency, disruptions from traffic overloads or infrastructure failures can have severe consequences.

In addition to these risks, the concentration of a large fraction of content within shared facilities could inadvertently streamline the control and filtering of content, especially

in countries where the government has a strong presence in the telecommunications market [13]. Instead of dealing with decentralized content sources to monitor, authorities can exert control at a handful of local choke points.

4 THE INTERNET LACKS CAPACITY TO HANDLE SPILLOVER FROM OFFNETS

4.1 Offnets run near capacity

Evidence suggests that offnets run with limited headroom to handle additional traffic. A study found that, prior to the COVID lockdown, Netflix offnets in some European ISPs delivered 63% of Netflix traffic, with the rest delivered across interdomain boundaries. When Netflix traffic spiked 58% during lockdown, the traffic from offnets only increased by 20%, whereas interdomain Netflix traffic more than doubled [32]! This result suggests that the offnets were already running near capacity, and so the excess demand had to spill over to interdomain links. Our analysis of traffic to 530 residential apartments supports this claim. During low traffic times of day, the vast majority of traffic comes from nearby servers, including Netflix and Akamai offnets hosted in the ISP. During peak periods, a higher fraction of traffic from the same services instead comes from more distant servers.

When a hypergiant runs out of capacity at an offnet location that is fully utilized or unavailable, it can deliver the overflow from other offnet locations within the same ISP or over an interdomain boundary from a server hosted in the ISP’s hierarchy of providers (offnet) or at the hypergiant itself (onnet) [24, 42]. We clustered offnet IP addresses into sites (§3) and found that 75.3%–91.2% of ISPs have only a single Netflix site, 37.8%–64.3% have only a single Meta site, 34.3%–78.4% have only a single Google site, and 34.6%–75.1% have only a single Akamai site. In these cases, the spillover would have to be served across interdomain boundaries. Even ISPs with multiple offnet sites from the same hypergiant may find spillover served from outside their networks. The other sites also likely run near capacity [24], and spillover periods may correlate in cases of geographically small ISPs with correlated peak times, spikes in traffic due to content popularity, and bad software updates that impact multiple offnets.

4.2 Insufficient dedicated peering exists

When traffic outweighs local offnet capacity, overflow traffic is served across interdomain links, either from the hypergiant or an upstream provider (which may itself host an offnet or receive the traffic over an interdomain link). Many ISPs that host a particular hypergiant’s offnets do not peer with that hypergiant (§4.2.1), many that do peer with the hypergiant do so only over shared links (§4.2.1), and the dedicated links that do exist often lack sufficient capacity (§4.2.2).

4.2.1 Many ISPs lack any dedicated peering capacity. We issued 21M traceroutes from Google Cloud in August 2023 to determine which ISPs Google peers with, using a methodology that earlier work validated against ground truth from Google [6]. (We cannot run measurements from Meta, Netflix, or Akamai). We established Virtual Machines (VMs) across all of Google’s regions, a total of 112 locations. From each VM, we issued a traceroute to a single IP address per /24 announced to the global Internet. To map IP addresses to ISPs, we followed the technique from the earlier work [6], except that we prioritize mapping with Euro-IX [17] data over PeeringDB data [44], based on from prior work [35]. This technique maps IXP addresses to the ISPs that use them. We inferred an ISP as a peer if any traceroute has a Google IP address *directly* followed by one mapped to the ISP.

Results. Of 4697 ISPs with Google offnets, 1798 (38.2%) peer with Google. For an additional 626 (13.3%), only unresponsive hops separate Google and the ISP in our traceroutes, suggesting the possibility of a peering. For the remaining 2273 ISPs with Google offnets (48.4%), our traceroutes reveal no evidence of peering, so Google traffic that is unable to be served from an offnet must come from the ISP’s provider. Among 9207 ISPs that peer with Google, 5735 (62.2%) peer via an IXP in at least one traceroute, and 3920 (42.5%) only appear to be connected through an IXP. Section 4.3 considers cases when traffic comes from a provider or via an IXP. Outside of IXPs, peering uses private network interconnects.

4.2.2 Dedicated peering that exists often lacks sufficient capacity. Although private interconnections (PNIs) provide dedicated capacity for traffic from a hypergiant, PNIs frequently lack sufficient bandwidth even under normal conditions. Hypergiants cannot unilaterally upgrade capacity as demand grows, and getting ISPs to upgrade can take months or even be impossible [49]. Google demand during peak periods exceeded capacity by an average of at least 13%, requiring rerouting of traffic [56]. Microsoft reported that workloads sometimes exceed the capacity of individual peering links [37]. Similarly, a study from Meta (at the time, Facebook) found that most Meta sites are capacity-constrained on at least some paths, and some sites are constrained on most paths [49]! In fact, 10% of Meta PNI experienced periods in which traffic demand was twice the capacity! In these periods, users either experience degraded performance due to congestion, or their traffic is rerouted via providers (§4.3). This excess demand occurred during normal operating conditions, and the situation could be much worse if offnets were overloaded or unavailable (§3.3). Given that dedicated connections are frequently overloaded under normal conditions, when 86% of Meta traffic comes from offnets (§2.1), the situation can get out of control if the offnets become unavailable, and the traffic needs to be served via interdomain links.

4.3 Spillover to IXPs and transit providers causes collateral damage

When an ISP does not have (enough) dedicated PNI capacity to a hypergiant, the overflow goes to links shared with other traffic. The hypergiant lacks insight into competing traffic so does not know how much traffic can be sent without causing congestion [49]. The shared link can be to a provider or can be a shared IXP fabric. Even without failures, neither transit providers nor IXPs have enough capacity to handle hypergiant traffic without congestion [49, 50]. In overload/failure scenarios, especially those that cause colocated offnets from multiple hypergiants to failover to the same shared routes, the collateral damage to other services can be significant.

5 RELATED WORK

Some studies map individual hypergiants [8, 9, 12, 52, 53, 55]. Section 2.2 used a technique from a paper that looked at whether offnets from multiple hypergiants are in the same ISP [21]. Our paper goes further to identify offnets colocated in the same facilities, highlighting the associated risk.

Given the Internet’s role in critical operations, there has been interest in understanding risk factors. Past efforts investigated the disruption risks of natural disasters [4, 28, 38] and shared physical conduits [16, 33, 36]. Hypergiants investigated failures in their own networks [20, 26, 40]. Our paper is in this vein and adds to the bigger picture by concentrating on the risk with shared collocation of offnets.

6 DISCUSSION

Our paper is at the intersection of two trends. First, the Internet has tended towards centralization, with a small number of hypergiants now responsible for most Internet traffic, a mix of their own popular services and services that benefit from deploying on their clouds and content delivery networks. This centralization brings the benefits of the cloud, but also risks and tradeoffs [15, 29, 31, 39, 41]. Second, some of these hypergiants have extensive offnet deployments [21], hosting servers in thousands of ISPs to decrease costs for ISPs and improve performance for users and services.

These trends combine to add a new dimension of risk, because of two common aspects of offnet deployments that our paper highlights. First, ISPs often colocate offnets from multiple hypergiants. So, while offnets decentralize content delivery from the perspective of any particular hypergiant, this collocation of some of the largest hypergiants centralizes traffic for many ISPs and users. Second, offnets allow the delivery of much more traffic to users than can otherwise be served, due to: (1) offnets that operate at close to capacity, with excess traffic spilling over to interdomain links; (2) insufficient dedicated peering, causing traffic to spill over to IXPs and transit; and (3) capacity constraints at IXPs and transit

providers. These conditions increase the risk of collateral damage and correlated failures of Internet services.

High-level approaches to mitigating risks of centralization, such as isolating resources and enforcing policies (via regulatory mandates, incentives, standards, and/or published best practices) may be common across the cloud and offnets, but differences between the settings add additional challenges. First, the cloud has much more ability to enforce isolation and provide elastic resources to handle overload and attacks while limiting collateral damage. The cloud was designed from the ground up to host multiple services via virtualization, and the same entity operates the machines, the network, and the facility. In contrast, each hypergiant provides and operates its own physical machines as offnets, and ISPs generally host them in existing ISP or third-party colocation facilities and have designed their networks primarily for providing access, not hosting high-volume third-party servers. Second, even for third-party tenants, a cloud provider has some load control, including via multi-path load balancing, VM placement, and, often, by providing the load balancing and/or service redirection service that select between multiple VMs and/or sites hosting the same service. ISPs have little control over the rate of queries or traffic served from offnets they host, and they generally have fewer path options than within a cloud data center and less capacity than within a cloud data center or WAN. Third, cloud hosting is a paid service, with transparent business agreements and SLAs, which is often not true for offnet hosting. Fourth, a handful of cloud providers host the lion’s share of cloud-based services, and each has legions of engineers and operators, whereas offnets are distributed across thousands of independent ISPs with widely varying operational sophistication and resources. Some offnet deployments can be at greater risk for outages and degradations than a hypergiant’s own facilities. The hypergiant has limited visibility and no control over the resilience of backup power, path diversity, or operational practices of the hosting ISP or of colocated hypergiants.

Despite these challenges, it is worth considering both technical and policy approaches. Technical directions could include isolation mechanisms deployed in colocation facilities, ISPs, IXPs, and transit, to protect capacity for each hypergiant and for other Internet traffic [11]; or approaches to share information and enable coordination among hypergiants as well as ISPs that host them [45]. Policy approaches could have similarities with existing compliance policies that dictate, for example, the physical security and backup power requirements for data centers hosting particular types of content [2, 7, 30]. A contemporaneous law journal article called for mandated public disclosure of content delivery infrastructure to allow risk assessment [39], focusing on the risks caused by content consolidation. Such an approach could also reveal and provide a basis for assessing the added risk

of colocation of infrastructure from multiple hypergiants. Hypergiants could update their published best practices for offnet deployments [24, 42] to address how ISPs can avoid common failure modes when hosting multiple hypergiants.

The time is ripe to consider how to mitigate these risks and ensure the resilience of Internet-based services, in conjunction with ongoing related discussions: contentious European Commission involvement in the tussle between hypergiants and access ISPs on cost sharing [14, 54] and calls to consider the role of offnets in network neutrality [1] and in the future of the public Internet [10]. We hope our paper is but the first step.

ACKNOWLEDGEMENTS

Our shepherd, Anees Shaikh, and the reviewers provided great feedback that improved our paper. We are grateful to Seung Joon Rhee for help with measurements. Marjory Blumenthal, Nick Merrill, Tejas Narechania, Aurojit Panda, Scott Shenker, and Vishal Misra provided valuable feedback on our work. This paper builds in part on our collaboration and discussions with our SIGCOMM 2021 co-authors [21]. This work was funded in part by NSF award CNS-2028550.

A CLUSTERING METHODOLOGY

For each latency value, we issued 8 pings, and took the latency of the ping that got the second smallest latency (following the existing technique). We discard 12K offnet IP addresses that did not respond to any of our measurements and 1.9K IP addresses where the latencies we measured could not possibly have come from a single destination (based on latencies from known M-Lab geolocations and the speed of light). Most remaining offnet IP addresses were very responsive, with all offnet addresses in 90% of ISPs responding to all 163 M-Lab sites. We discard ISPs that have fewer than 100 M-Lab sites with successful measurements to all offnets, to have enough data for accurate clustering.

For each of the remaining 5151 ISPs, in turn, we cluster the offnet IP addresses in the ISP using OPTICS following the approach in prior work: for each pair of IP addresses, we calculate the distance as the (normalized) Manhattan distance after excluding measurements from the 20% of M-Lab sites that have the largest latency discrepancy between the two addresses [12]. The OPTICS algorithm takes two parameters: n_{\min} , the minimum number of IP addresses that can form a cluster of the elements and x_i (§3.2). The prior publication that clustered offnet IP addresses did not specify the parameter values used in its results. We set $n_{\min} = 2$ so that clusters can be as small as two IP addresses. OPTICS will not assign an IP address to a cluster if no address is within a short distance, in which case we consider the offnet as not colocated.

REFERENCES

- [1] Muhammad Abdullah, Pavlos Nikolopoulos, and Katerina Argyraki. Caching and Neutrality. *ACM HotNets*, 2023.
- [2] Akamai. Akamai | Information Security Compliance, 2023. URL <https://www.akamai.com/legal/compliance>.
- [3] Akamai. Facts and Figures, 2023. URL <https://www.akamai.com/company/facts-figures.jsp>.
- [4] Scott Anderson, Carol Barford, and Paul Barford. Five Alarms: Assessing the Vulnerability of US Cellular Communication Infrastructure to Wildfires. *ACM IMC*, 2020. URL <https://doi.org/10.1145/3419394.3423663>.
- [5] Mihael Ankerst, Markus M. Breunig, Hans-Peter Kriegel, and Jörg Sander. OPTICS: Ordering Points to Identify the Clustering Structure. *ACM SIGMOD*, 1999. URL <https://doi.org/10.1145/304182.304187>.
- [6] Todd Arnold, Jia He, Weifan Jiang, Matt Calder, Italo Cunha, Vasileios Giotsas, and Ethan Katz-Bassett. Cloud Provider Connectivity in the Flat Internet. *ACM IMC*, 2020. URL <https://doi.org/10.1145/3419394.3423613>.
- [7] Azure. Azure Compliance Documentation, 2023. URL <https://learn.microsoft.com/en-us/azure/compliance/>.
- [8] Anurag Bhatia. Mapping Facebook’s FNA (CDN) Nodes Across the World!, 2018. URL <https://anuragbhatia.com/2018/03/networking/isp-column/mapping-facebooks-fna-cdn-nodes-across-the-world/>.
- [9] Anurag Bhatia. Facebook FNA Node Update, 2019. URL <https://anuragbhatia.com/2019/11/networking/isp-column/facebook-fna-node-update/>.
- [10] Marjory Blumenthal, Ramesh Govindan, Ethan Katz-Bassett, Arvind Krishnamurthy, James McCauley, Nick Merrill, Tejas Narechania, Aurojit Panda, and Scott Shenker. Can We Save The Public Internet? 2023.
- [11] Lloyd Brown, Ganesh Ananthanarayanan, Ethan Katz-Bassett, Arvind Krishnamurthy, Sylvia Ratnasamy, Michael Schapira, and Scott Shenker. On the Future of Congestion Control for the Public Internet. *ACM HotNets*, 2020. URL <https://doi.org/10.1145/3422604.3425939>.
- [12] Matt Calder, Xun Fan, Zi Hu, Ethan Katz-Bassett, John Heidemann, and Ramesh Govindan. Mapping the Expansion of Google’s Serving Infrastructure. *ACM IMC*, 2013. URL <https://doi.org/10.1145/2504730.2504754>.
- [13] Esteban Carisimo, Alexander Gamero-Garrido, Alex C. Snoeren, and Alberto Dainotti. Identifying ASes of State-Owned Internet Operators. In *Proceedings of the 21st ACM Internet Measurement Conference*, *ACM IMC*, 2021. URL <https://doi.org/10.1145/3487552.3487822>.
- [14] Chao Liu and Ernesto Falcon. There is Nothing Fair About the European Commission’s “Fair Share” Proposal, 2023. URL <https://www.eff.org/deeplinks/2023/06/there-nothing-fair-about-european-commissions-fair-share-proposal>.
- [15] Trinh Viet Doan, Roland van Rijswijk-Deij, Oliver Hohlfeld, and Vaibhav Bajpai. An Empirical View on Consolidation of the Web. *ACM Transaction on Internet Technology*, 2022. URL <https://doi.org/10.1145/3503158>.
- [16] Ramakrishnan Durairajan, Paul Barford, Joel Sommers, and Walter Willinger. InterTubes: A Study of the US Long-Haul Fiber-Optic Infrastructure. *ACM SIGCOMM*, 2015. URL <https://doi.org/10.1145/2785956.2787499>.
- [17] Euro-IX. IXPDB. URL <https://ixpdb.euro-ix.net/en/ix-f/ixp-database/>.
- [18] Facebook. FNA Installation and Operation Guide, 2017. URL https://kupdf.net/download/fna-installation-and-operation-guide-v804_597ff4a4dc0d6055102bb17e_pdf.
- [19] Dean Garfield. Red Light Green Light? No to Network Usage Fees, 2021. URL <https://about.netflix.com/en/news/red-light-green-light-no-to-network-usage-fees>.
- [20] Monia Ghobadi and Ratul Mahajan. Optical Layer Failures in a Large Backbone. *ACM IMC*, 2016. URL <https://doi.org/10.1145/2987443.2987483>.
- [21] Petros Gigis, Matt Calder, Lefteris Manassakis, George Nomikos, Vasileios Kotronis, Xenofontas Dimitropoulos, Ethan Katz-Bassett, and Georgios Smaragdakis. Seven Years in the Life of Hypergiants’ off-Nets. *ACM SIGCOMM*, 2021. URL <https://doi.org/10.1145/3452296.3472928>.
- [22] Phillipa Gill, Christophe Diot, Lai Yi Ohlsen, Matt Mathis, and Stephen Soltesz. M-Lab: User Initiated Internet Data for the Research Community. *ACM SIGCOMM Computer Communication Review*, 2022. URL <https://dl.acm.org/doi/10.1145/3523230.3523236>.
- [23] Google. Edge Nodes GGC, 2023. URL <https://peering.google.com/#/options/google-global-cache>.
- [24] Google. Multi-node Deployments - Interconnect Help, 2023. URL <https://support.google.com/interconnect/answer/9051938>.
- [25] Google. Preparing your Network, 2023. URL https://support.google.com/interconnect/answer/9059054?hl=en&ref_topic=7659203&sjid=2857398392929088225-NA.
- [26] Ramesh Govindan, Ina Minei, Mahesh Kallahalla, Bikash Koley, and Amin Vahdat. Evolve or Die: High-Availability Design Principles Drawn from Google’s Network Infrastructure. *ACM SIGCOMM*, 2016. URL <https://doi.org/10.1145/2934872.2934891>.
- [27] Geoff Huston. How Big is that Network, 2014. URL <https://labs.apnic.net/?p=526>.
- [28] Sangeetha Abdu Jyothi. Solar Superstorms: Planning for an Internet Apocalypse. *ACM SIGCOMM*, 2021. URL <https://doi.org/10.1145/3452296.3472916>.
- [29] Aqsa Kashaf, Vyas Sekar, and Yuvraj Agarwal. Analyzing Third Party Service Dependencies in Modern Web Services: Have We Learned from the Mirai-Dyn Incident? *ACM IMC*, 2020. URL <https://doi.org/10.1145/3419394.3423664>.
- [30] Thomas Koch, Ke Li, Calvin Ardi, Matt Calder, John Heidemann, and Ethan Katz-Bassett. Anycast in Context: A Tale of Two Systems. *ACM SIGCOMM*, 2021. URL <https://doi.org/10.1145/3452296.3472891>.
- [31] Rashna Kumar, Sana Asif, Elise Lee, and Fabian E. Bustamante. Each at Its Own Pace: Third-Party Dependency and Centralization Around the World. *ACM SIGMETRICS*, 2023. URL <https://doi.org/10.1145/3579437>.
- [32] Craig Labovitz. Pandemic Impact on Global Internet Traffic, 2020. URL <https://www.nanog.org/news-stories/nanog-tv/nanog-79-webcast/effects-covid-19-lockdowns-service-provider-networks/>.
- [33] Shucheng Liu, Zachary S. Bischof, Ishaan Madan, Peter K. Chan, and Fabián E. Bustamante. Out of Sight, Not Out of Mind: A User-View on the Criticality of the Submarine Cable Network. *ACM IMC*, 2020. URL <https://doi.org/10.1145/3419394.3423633>.
- [34] Matthew Luckie, Bradley Huffaker, Alexander Marder, Zachary Bischof, Marianne Fletcher, and K. Learning to Extract Geographic Information from Internet Router Hostnames. *ACM CoNEXT*, 2021. URL <https://doi.org/10.1145/3485983.3494869>.
- [35] A Marder, kc claffy, and A Snoeren. Inferring Cloud Interconnections: Validation, Geolocation, and Routing Behavior. *PAM*, 2021. URL https://doi.org/10.1007/978-3-030-72582-2_14.
- [36] A Marder, Z Zhang, R Padmanabhan, R Mok, B Huffaker, M Luckie, A Dainotti, kc claffy, A Snoeren, and A Schulman. Access Denied: Assessing Physical Risks to Internet Access Networks. *USENIX Security Symposium*, 2023. URL <https://www.usenix.org/system/files/usenixsecurity23-marder.pdf>.
- [37] Michael Markovitch, Sharad Agarwal, Rodrigo Fonseca, Ryan Beckett, Chuanji Zhang, Irena Atov, and Somesh Chaturmohta. TIPSYP: Predicting Where Traffic Will Ingress a WAN. *ACM SIGCOMM*, 2022. URL <https://doi.org/10.1145/3544216.3544234>.
- [38] Juno Mayer, Valerie Sahakian, Emilie Hooft, Douglas Toomey, and Ramakrishnan Durairajan. On the Resilience of Internet Infrastructures in Pacific Northwest to Earthquakes. *PAM*, 2021. URL

- https://link.springer.com/chapter/10.1007/978-3-030-72582-2_15.
- [39] Nick Merrill and Tejas N Narechania. Inside the Internet. *Duke Law Journal Online*, 2023. URL https://papers.ssrn.com/sol3/papers.cfm?abstract_id=4466419.
- [40] Justin Meza, Tianyin Xu, Kaushik Veeraraghavan, and Onur Mutlu. A Large Scale Study of Data Center Network Reliability. ACM IMC, 2018. URL <https://doi.org/10.1145/3278532.3278566>.
- [41] Giovane C. M. Moura, Sebastian Castro, Wes Hardaker, Maarten Wullink, and Cristian Hesselman. Clouding up the Internet: How Centralized is DNS Traffic Becoming? ACM IMC, 2020. URL <https://doi.org/10.1145/3419394.3423625>.
- [42] Netflix. Netflix | Open Connect Sample Architectures, 2023. URL <https://openconnect.netflix.com/en/#sample-architectures>.
- [43] Netflix. Requirements for Deploying Embedded Appliances, 2023. URL <https://openconnect.zendesk.com/hc/en-us/articles/360034538352>.
- [44] PeeringDB. PeeringDB. URL <http://www.peeringdb.com>.
- [45] Enric Pujol, Ingmar Poese, Johannes Zerwas, Georgios Smaragdakis, and Anja Feldmann. Steering Hyper-Giants' Traffic at Scale. ACM CoNEXT, 2019. URL <https://doi.org/10.1145/3359989.3365430>.
- [46] Rapid7. Project Sonar, Reverse DNS, 2023. URL <https://www.rapid7.com/research/project-sonar/>.
- [47] Reddit. Which (CDN) Caching Appliances Do You Run at Your ISP? Which Gives the Biggest Savings?, 2023. URL https://www.reddit.com/r/networking/comments/11ddyxg/which_cdn_caching_appliances_do_you_run_at_your/.
- [48] Sandvine. Global Internet Phenomena Report 2023, 2023. URL <https://www.sandvine.com/global-internet-phenomena-report-2023>.
- [49] Brandon Schlinker, Hyojeong Kim, Timothy Cui, Ethan Katz-Bassett, Harsha V. Madhyastha, Italo Cunha, James Quinn, Saif Hasan, Petr Lapukhov, and Hongyi Zeng. Engineering Egress with Edge Fabric: Steering Oceans of Content to the World. ACM SIGCOMM, 2017. URL <https://doi.org/10.1145/3098822.3098853>.
- [50] Brandon Schlinker, Italo Cunha, Yi-Ching Chiu, Srikanth Sundaresan, and Ethan Katz-Bassett. Internet Performance from Facebook's Edge. ACM IMC, 2019. URL <https://doi.org/10.1145/3355369.3355567>.
- [51] Kyle Schomp, Onkar Bhardwaj, Eymen Kurdoglu, Mashooq Muhaimen, and Ramesh K. Sitaraman. Akamai DNS: Providing Authoritative Answers to the World's Queries. ACM SIGCOMM, 2020. URL <https://doi.org/10.1145/3387514.3405881>.
- [52] Florian Streibelt, Jan Böttger, Nikolaos Chatzis, Georgios Smaragdakis, and Anja Feldmann. Exploring EDNS-Client-Subnet Adopters in Your Free Time. ACM IMC, 2013. URL <https://doi.org/10.1145/2504730.2504767>.
- [53] Ao-Jan Su, David R. Choffnes, Aleksandar Kuzmanovic, and Fabián E. Bustamante. Drafting behind Akamai (Travelocity-Based Detouring). ACM SIGCOMM, 2006. URL <https://doi.org/10.1145/1159913.1159962>.
- [54] Telefonica. What is Fair Share, 2023. URL <https://www.telefonica.com/en/communication-room/blog/what-is-fair-share/>.
- [55] Sipat Triukose, Zhihua Wen, and Michael Rabinovich. Measuring a Commercial Content Delivery Network. ACM WWW, 2011. URL <https://doi.org/10.1145/1963405.1963472>.
- [56] Kok-Kiong Yap, Murtaza Motiwala, Jeremy Rahe, Steve Padgett, Matthew Holliman, Gary Baldus, Marcus Hines, Taeun Kim, Ashok Narayanan, Ankur Jain, Victor Lin, Colin Rice, Brian Rogan, Arjun Singh, Bert Tanaka, Manish Verma, Puneet Sood, Mukarram Tariq, Matt Tierney, Dzevad Trumic, Vytautas Valancius, Calvin Ying, Mahesh Kallahalla, Bikash Koley, and Amin Vahdat. Taking the Edge off with Espresso: Scale, Reliability and Programmability for Global Internet Peering. ACM SIGCOMM, 2017. URL <https://doi.org/10.1145/3098822.3098854>.
- [57] Ming Zhang, Yaoping Ruan, Vivek Pai, and Jennifer Rexford. How DNS Misnaming Distorts Internet Topology Mapping. USENIX ATC, 2006. URL <https://dl.acm.org/doi/10.5555/1267359.1267393>.