Check for updates

Can We Save The Public Internet?

Marjory Blumenthal[1], Ramesh Govindan[2], Ethan Katz-Bassett[3], Arvind Krishnamurthy[4], James McCauley[5], Nick Merrill[6], Tejas Narechania[6], Aurojit Panda[7], Scott Shenker[1,6]
[1] ICSI, [2] USC, [3] Columbia, [4] UWashington, [5] Mount Holyoke, [6] UC Berkeley, [7] NYU (Authors in alphabetical order)
Contact:shenker@icsi.berkeley.edu

This article is an editorial note submitted to CCR. It has NOT been peer reviewed.
The authors take full responsibility for this article's technical content. Comments can be posted through CCR Online.

PREFACE

In the pursuit of making technical progress, our community sometimes loses sight of the broader picture. One example of this phenomenon is the current trajectory of the Internet. The past decade brought great improvements in both the deployed infrastructure – with its greater speeds, extended reach, and additional functionality – and the conceptual foundations, with our community now having a far greater understanding of verification, control-plane designs, programmable hardware, and other fundamental issues. However, significantly less attention has been paid to changes in the overall structure of the Internet ecosystem, which increasingly revolves around extensive private networks associated with a large content and/or cloud provider. While these developments have certainly resulted in tangible short-term benefits, we believe that in the long-term this trend threatens the future of what we call the "public Internet."

To raise awareness and ignite a debate about these concerns, in 2021 some of us wrote a CCR editorial entitled "Revitalizing the Public Internet by Making It Extensible." In terms of our goals, the editorial was a colossal failure; it caused more confusion than concern, and was more ignored than igniting. One of the painful lessons of writing is that if, after a careful editing process, your prose still confuses people, then the problem almost certainly lies in your own lack of conceptual clarity. *Nostra culpa.*

We spent the past year seeking that clarity, and have identified a central issue underlying the Internet's problematic trends: the tension between *interconnection* and *enhancement*. This captures the essence of the problem in a way that eluded us in 2021. Rather than write about this tension in a manner only accessible to highly technical readers, we produced a document that is intended for a much wider audience, which seems appropriate given the critical role the Internet plays in so many aspects of our lives. This document is entitled "Can We Save The Public Internet?" and follows this preface; the standalone version can be found at [5].

To reach a general audience, we wanted to articulate our insight as clearly and concisely as possible, explaining the basic problem from start to finish in roughly two pages of text (not counting the addendum)! As a result, we did not cite previous work (most notably the early warnings from Geoff Huston [9]), nor document our findings (papers such as [2, 3, 6–8, 10–13, 15–19] quantify some of these trends), nor deal with additional factors in the modern Internet, such as IXPs (as described in [1]).

While this lack of completeness is the unfortunate cost of brevity, we hope the concision helps our community view this document not as a recitation of fully-documented facts, but as a challenge to answer two fundamental questions:

(1) *Do we believe that current trends are causing the Internet to drift away from being a neutral platform for application innovation?* There are two components to answering this question. One is to provide quantitative measures of the underlying infrastructure that reveal the extent to which the current Internet departs from the classical Internet's use of transit for connectivity. The other is to explore the economic impact of the current structure. For instance, the carrier-dominated Internet allowed the hyperscalers to arise, and they are now the most powerful incumbents; however, there are reasons to doubt that the current hyperscaler-dominated Internet would support the emergence of the next generation of successors. Any preliminary evidence for or against this conclusion would be most valuable.

(2) *If the answer to the first question is yes, what should be done to reverse this trend?* This question also has two components. The first is exploring a set of appropriate technical responses. We have put forward the notion of the InterEdge, but we hope the community develops a variety of possible approaches and engages in an active debate about their relative merits. The second is considering what role our community should play in making any of these possible solutions a reality. The Internet's creation was a victory of the research community over the incumbent telecommunications ecosystem. Given the extensive ties between today's research community and the current hyperscalers, do we have the courage and independence to repeat this success?

Abstract

The goal of this short document is to explain why recent developments in the Internet's infrastructure are problematic. As context, we note that the Internet was originally designed to provide a simple universal service – global end-to-end packet delivery – on which a wide variety of end-user applications could be built. The early Internet supported this packet-delivery service via an interconnected collection of commercial Internet Service Providers (ISPs) that we will refer to collectively as the "public Internet." The Internet has fulfilled its packet-delivery mission far beyond all expectations and is now the dominant global communications infrastructure. By providing a level playing field on which new applications could be deployed, the Internet has enabled a degree of innovation that no one could have foreseen. To improve performance for some common applications, "enhancements" such as caching (as in content-delivery networks) have been gradually added to the Internet. The resulting performance improvements are so significant that such enhancements are now effectively necessary to meet current content delivery demands. Despite these tangible benefits, this document argues that the way these enhancements are currently deployed seriously undermines the sustainability of the public Internet and could lead to an Internet infrastructure that reaches fewer people and is largely concentrated among only a few large-scale providers. We wrote this document because we fear that these developments are now decidedly tipping the Internet's playing field towards those who can deploy these enhancements at massive scale, which in turn will limit the degree to which the future Internet can support unfettered innovation. This document begins by explaining our concerns but goes on to articulate how this unfortunate fate can be avoided. To provide more depth for those who seek it, we provide a separate addendum with further detail.

CCS Concepts

• **Networks** → **Network architectures**;

Keywords

Internet architecture, Public Internet, Internet enhancements

Why Should We Worry About Internet Enhancements?

The public Internet is a set of densely interconnected ISPs, and its basic end-to-end packet-delivery service relies on transit ISPs to carry traffic between the ISP where traffic originates and the ISP where traffic terminates. This interconnection enables the Internet to achieve global service without requiring global infrastructures by allowing new ISPs with limited geographic expanse to join the public Internet by merely connecting with a transit ISP (and paying any associated costs). While tremendously successful during the Internet's initial decades, this delivery-via-transit-ISPs paradigm is no longer the dominant delivery model. Instead, today, most Internet traffic goes directly between a last-mile network (as provided by an access ISP) and a content-delivery infrastructure (as deployed by Google, Netflix, Fastly, Cloudflare, and others) that, in addition to packet delivery, also provides enhanced functions that significantly improve performance. These enhancements include caching (as provided by content-delivery networks, hereafter CDNs) and load balancing (which spreads requests across datacenters). These content-delivery enhancements are not standardized, nor are they part of normal Internet interconnection arrangements. As a result, while the basic packet-delivery functions of these infrastructures can interconnect (i.e., can leverage each other to achieve end-to-end packet delivery), their enhancements do not interconnect in the same way; for example, Cloudflare and Fastly cannot leverage each other's content-delivery infrastructures to provide enhancements for their customers. Thus, a collection of geographically-scoped content-delivery infrastructures cannot interconnect to provide global coverage. This tension between interconnection and enhancement, which are mutually exclusive in today's Internet, is central to the rest of this document; the former without the latter (which is what the public Internet provides) has worse performance, but the latter without the former (which is the dominant paradigm today) reduces competition and coverage, and seriously undermines the public Internet. We now explain these concerns in more detail.

In contrast to the public Internet, where the interconnection of ISPs can create global coverage out of local infrastructures, the lack of interconnection between content-delivery infrastructures means that global coverage for enhancements requires global content-delivery infrastructures. For example, Google, Azure, Amazon, Netflix, Meta, Cloudflare, and several others have deployed, in parallel, their own vast content-delivery infrastructures. This need to build out large-scale infrastructures creates a high barrier to entry, limiting the number of entities capable of entering this market and increasing the possibility of anti-competitive behavior.

The lack of interconnection also creates a strong incentive for these parallel content-delivery infrastructures to focus their deployments in the same high-revenue regions, thereby limiting coverage and deepening the digital divide (e.g., their coverage is far better in North America and Europe than in other parts of the world). Of course, the divide between having access to enhanced content delivery and only having access to the public Internet is far less severe than the divide between having Internet access and not having any access at all, but the gap between enhanced and unenhanced may increase over time as more enhancements are deployed and they become essential for an increasing number of applications.

Even more troubling for the long-term is that, because most traffic today does not involve a transit ISP but instead goes directly between an access ISP and a content-delivery infrastructure, the role of (and therefore the relative investment in) the public Internet is rapidly diminishing. The emphasis on these enhancing content-delivery infrastructures is putting the public Internet at risk, all for the sake of providing better service for a subset of users and a subset of the applications they use. The public Internet is far too important to let it wither in silence, particularly when, as we argue below, this fate is avoidable.

Can We Have Both Interconnection and Enhancement?

This previously unaddressed question is crucial because combining interconnection with enhancement would create an "enhanced public Internet" that couples the broad coverage and vigorous competition of the public Internet with the improved performance of today's content-delivery infrastructures. To describe how we might create an enhanced public Internet, we note that the current enhancements are implemented in software running on clusters of general-purpose servers at the network edge. Interconnection requires standardization of these enhancements, and we propose achieving standardization through open-source implementations of each enhancement. Interconnection also requires that enhancement providers agree about how they cooperatively ensure that all traffic has the relevant enhancements applied at the appropriate network edges. For convenience, we will refer to this set of standardized open-source implementations and the rules for cooperation as the InterEdge, to reflect that it would bring Internet-style interconnection to the variety of enhancements now being applied at the network edge.

These InterEdge clusters (which, to provide uniformity, must support all standardized InterEdge enhancements, and may also offer additional experimental or proprietary services) could be deployed by a variety of commercial entities – including access ISPs, cloud providers, traditional CDNs, Internet exchange points (IXPs), and various emerging edge providers – so the InterEdge's level of coverage and competition would be more similar to the interconnected public Internet than to the current parallel deployments of content-delivery infrastructures. Moreover, because new enhancements only involve the deployment of software at the network edge, the set of standardized InterEdge enhancements can grow over time, giving it an evolvability more similar to the content-delivery infrastructures than to the unchanging packet-delivery service of the public Internet, where the packet-forwarding functionality is largely implemented in specialized hardware.

Importantly, the InterEdge would not replace existing content-delivery infrastructures but would only provide an alternative option. While, initially, the InterEdge might not be as high-performance as the finely-tuned content-delivery infrastructures deployed today, the InterEdge would have an important structural advantage over them. Today, all enhancements must be implemented in terms of existing protocols, because application vendors are justifiably unwilling to bind themselves to any particular provider of proprietary enhancements (unless, of course, the enhancement provider is also the application vendor, as with Google with Android and Chrome). In contrast, embracing a new InterEdge enhancement, even one that does not use existing protocols, does not tie an application vendor to any particular enhancement provider, as such open-source enhancements could be deployed by any InterEdge provider. Thus, the InterEdge would be an interconnected yet evolvable platform, making the enhanced public Internet far more prepared to meet the needs of future applications than the public Internet or today's content-delivery infrastructures. The question then is not whether the InterEdge would be desirable, but whether it is technically and economically viable.

Is the InterEdge Technically and Economically Viable?

There are many technical challenges in building the InterEdge, including scalability, resilience, security, and performance. However, these challenges are already being actively addressed and, because they are common problems faced in many large-scale distributed systems, we have no doubt that they can be adequately met.

The economic issues of the InterEdge are more subtle, as commercial viability requires that some entities in the ecosystem are willing to become first-adopters of the InterEdge. The dominant players in the current cloud and content-delivery ecosystem are the highly-integrated and hyperscaled cloud/content providers who have their own large content-delivery infrastructures, such as Microsoft and Google. This is in contrast to two other segments of the ecosystem: cloud and content providers who are not big enough to deploy their own large-scale content-delivery infrastructures (e.g., IBM) and content-delivery infrastructures that do not have their own fully-featured cloud computing or first-party content offerings (e.g., Cloudflare). The early adoption of the InterEdge by these two market segments would be mutually beneficial; the smaller cloud and content providers would have access to advanced content-delivery infrastructures that do not lock them in, and the InterEdge providers would have a growing set of paying customers. The resulting symbiosis could establish the InterEdge's economic viability.

It is unlikely that the integrated hyperscalers would join the InterEdge in its early years, but as the InterEdge achieves significant market share, they might begin to support the standardized InterEdge enhancements in their own infrastructure for two reasons. First, if application vendors start leveraging the InterEdge enhancements – which could happen if the InterEdge offered enhancements that were critical for some emerging applications, such as data aggregation for federated learning – then the hyperscalers would have an incentive to also support those enhancements in their own content-delivery infrastructure. Second, by adopting the InterEdge's set of enhancements, they could then more flexibly augment their own infrastructure by leveraging the InterEdge providers in locations that their own infrastructure does not sufficiently cover.

This document provides a short and shallow introduction to the deep and complicated issues in today's Internet. However, we believe that the bottom line is quite simple: today's trends pose a threat to the public Internet and the level playing field it has provided for application innovation. We believe that this unfortunate fate can be avoided if we can combine the interconnection of the public Internet with today's enhancements. The result would create an enhanced and evolvable but still level playing field for innovation. This is a goal well within our reach, but it will take support among a broad coalition of participants – from industry leaders to policymakers to open-source development communities – to make it happen. If you are interested in learning more, please contact Scott Shenker at shenker@icsi.berkeley.edu.

Addendum to "Can We Save The Public Internet?"

The previous document briefly summarizes why we must and how we can save the public Internet, but in pursuit of brevity, it necessarily glosses over many points that deserve more careful discussion. This addendum addresses a small subset of those points.

We first clarify our terminology. We use the term ISP broadly to refer to any provider of Internet service, whether access or transit. More crucially, the term "public" in the phrase "public Internet" is used to indicate that the Internet was originally designed to provide general-purpose connectivity between any two connected users. The nature of the ISPs carrying the traffic is typically commercial, but they provide – to the general public via an open market – a generic common carriage service. This is in stark contrast to the content-delivery infrastructures that only carry their own traffic or the traffic of those that use their cloud infrastructure. To clarify this distinction, we now provide a taxonomy of content-delivery infrastructures.

Each content-delivery infrastructure has a set of edge-based clusters implementing the enhancements, and some have their own private network providing communication between those clusters (the alternative being to use the public Internet to provide the intercluster communication). This leads to several ways content providers can leverage current content-delivery infrastructures:

- Content provider with its own clusters and private network (e.g., Google, Azure, Meta)
- Content provider with its own clusters but uses the public Internet to provide communication between those clusters (e.g., Netflix)
- Content provider using a third-party content-delivery infrastructure that has its own clusters and private network (e.g., customers of Google, Azure, and Amazon)
- Content provider using a third-party content-delivery infrastructure that has its own clusters but uses the public Internet to provide communication between those clusters (e.g., customers of Fastly)

The history of how enhancements came to be deployed illuminates some important issues. Standalone CDNs (i.e., those having a set of clusters but using the public Internet for intercluster communication) arose soon after the Internet went commercial to meet the needs of content providers (e.g., Akamai was founded in 1998). These early standalone CDNs did not threaten the public Internet because they utilized the public Internet to carry their traffic between their clusters, and thus were consistent with the delivery-via-transit-ISPs paradigm. The threat only came later, when the hyperscalers started building out their own content-delivery infrastructures paired with a private network. These private networks are now carrying an increasing share of the Internet's traffic, thereby reducing the relative demand for (and investment in) transit capacity.

One might ask why enhancements, particularly the CDNs that emerged so early in the Internet's history, did not become a normal part of the services ISPs offer to their customers. CDNs lower user-experienced latency, but this improvement is typically not enough to convince users to pay additional access fees. However, the improvements tend to cause users to stay online longer, resulting in more revenue for the content provider. Thus, it was the content providers, not content consumers, who were eager to pay for these enhancements. However, no single ISP could offer content providers global coverage and a single point of payment, so CDNs such as Akamai arose to fill this market gap soon after the Internet went commercial.

There are probably many reasons why the early ISPs did not succeed in creating an interconnected set of CDNs, much like what we are proposing for the InterEdge, but here we focus on the two main challenges such an approach would face today. The first is keeping pace with the feature set offered by existing content-delivery infrastructures, and the second is agreeing on how to provide a single point of payment. The first challenge is met by the use of open-source software rather than formal written standards, which allows the feature-set of InterEdge enhancements to evolve more easily over time. While the resulting pace of change would be slower than that of existing content-delivery infrastructures, the differential in evolution speed would likely be outweighed by the larger set of InterEdge providers.

The second challenge, a single point-of-payment, is significantly more difficult. We first note that the InterEdge would not change current ISP interconnection agreements; we expect a clean separation between the payments for enhancements and the payments between ISPs and by users to ISPs for packet delivery, which will remain the current "pay for access" approach. So the question becomes: how could the InterEdge provide a single point-of-payment for enhancements? Answering this question requires some additional context.

The current Internet is built around access charging and recursive transfers: you pay your ISP to access the Internet (where access covers both sending and receiving traffic), and then your ISP pays its provider, and so on, via a recursive set of pairwise payments. In contrast, the payments for enhanced services are primarily borne by content providers, who pay for the content-delivery services either by building their own content-delivery infrastructure or by paying a third-party content-delivery infrastructure. Content providers typically only pay a single content-delivery infrastructure to provide the geographic coverage they need (some contract with a few CDNs, but for convenience we will focus on single points-of-payment). The challenge for the InterEdge is that a single content provider might be reaching a set of users who have many different InterEdge providers. The content provider cannot be expected to set up separate contracts with all of the InterEdge providers that their customers might be using. To rectify this, we propose that InterEdge providers be required, as part of the cooperation agreement, to publish a set of "posted prices" (which can vary from cluster to cluster and be based on the volume of service rendered). We then assume that a set of brokers will arise that make arrangements with the providers of InterEdge clusters in all regions. The brokers would ask content providers where they want enhancement coverage and their expected resource requirements (e.g., cache space), and the

broker would provide a price for that service and charge accordingly, providing a single point-of-payment. The brokers would select which InterEdge providers to use in a given region, based on whatever criteria they choose, so there is no need for an explicit agreement among InterEdge providers about who serves which customers. Involving brokers in the InterEdge is a decidedly complicating and perhaps problematic step, but we think it is a necessary one. Moreover, the use of brokers has been discussed elsewhere [14], and some companies have created their own internal brokers ("meta-CDNs" [4]) by contracting with multiple CDNs. Because the duties of the brokers are relatively lightweight – akin to travel sites – we expect that there will be many competing brokers, lowering concerns about their taking too large a share of revenue.

References

[1] B. Ager, N. Chatzis, A. Feldmann, N. Sarrar, S. Uhlig, and W. Willinger. Anatomy of a Large European IXP. In *SIGCOMM*, 2012.

[2] T. Arnold, E. Gürmeriçliler, G. Essig, A. Gupta, M. Calder, V. Giotsas, and E. Katz-Bassett. (How Much) Does a Private WAN Improve Cloud Performance? In *INFOCOM*, 2020.

[3] T. Arnold, J. He, W. Jiang, M. Calder, I. Cunha, V. Giotsas, and E. Katz-Bassett. Cloud Provider Connectivity in the Flat Internet. In *IMC*, 2020.

[4] J. Blendin, F. Bendfeldt, I. Poese, B. Koldehofe, and O. Hohlfeld. Dissecting Apple's Meta-CDN during an iOS Update. In *IMC*, IMC '18, 2018. doi: 10.1145/3278532.3278567.

[5] M. Blumenthal, R. Govindan, E. Katz-Bassett, A. Krishnamurthy, J. McCauley, N. Merrill, T. Narechania, A. Panda, and S. Shenker. Can We Save The Public Internet? https://drive.google.com/file/d/1-JQU17yzqbNbdEO0bE8-zewmmczIFjP/view?usp=share_link.

[6] T. Flach, N. Dukkipati, A. Terzis, B. Raghavan, N. Cardwell, Y. Cheng, A. Jain, S. Hao, E. Katz-Bassett, and R. Govindan. Reducing Web Latency: The Virtue of Gentle Aggression. In *SIGCOMM*, 2013.

[7] A. Ganjam, J. Jiang, X. Liu, V. Sekar, F. Siddiqi, I. Stoica, J. Zhan, and H. Zhang. C3: Internet-Scale Control Plane for Video Quality Optimization. In *NSDI*, 2015.

[8] P. Gigis, M. Calder, L. Manassakis, G. Nomikos, V. Kotronis, X. Dimitropoulos, E. Katz-Bassett, and G. Smaragdakis. Seven Years in the Life of Hypergiants' off-Nets. In *SIGCOMM*, 2021.

[9] G. Huston. The death of transit? Web, Oct 2016. https://blog.apnic.net/2016/10/28/the-death-of-transit/.

[10] U. Krishnaswamy, R. Singh, P. Mattes, P.-A. C. Bissonnette, N. Bjørner, Z. Nasrin, S. Kothari, P. Reddy, J. Abeln, S. Kandula, H. Raj, L. Irun-Briz, J. Gaudette, and E. Lan. OneWAN is better than two: Unifying a split WAN architecture. In *NSDI*, 2023.

[11] C. Labovitz. Effects of COVID 19 Lockdowns on Service Provider Networks. https://www.nanog.org/news-stories/nanog-tv/nanog-79-webcast/effects-covid-19-lockdowns-service-provider-networks/.

[12] C. Labovitz, S. Iekel-Johnson, D. McPherson, J. Oberheide, and F. Jahanian. Internet Inter-Domain Traffic. In *SIGCOMM*, 2010.

[13] N. Merrill and T. N. Narechania. Inside the Internet. *Duke Law Journal Online*, 73:35–76, 2023.

[14] M. K. Mukerjee, I. N. Bozkurt, D. Ray, B. M. Maggs, S. Seshan, and H. Zhang. Redesigning CDN-Broker Interactions for Improved Content Delivery. In *CoNEXT*, 2017.

[15] A. Pathak, Y. A. Wang, C. Huang, A. Greenberg, Y. C. Hu, R. Kern, J. Li, and K. W. Ross. Measuring and Evaluating TCP Splitting for Cloud Services. In *PAM*, 2010.

[16] B. Schlinker, H. Kim, T. Cui, E. Katz-Bassett, H. V. Madhyastha, I. Cunha, J. Quinn, S. Hasan, P. Lapukhov, and H. Zeng. Engineering Egress with Edge Fabric: Steering Oceans of Content to the World. In *SIGCOMM*, 2017.

[17] K. Vermeulen, L. Salamatian, S. H. Kim, M. Calder, and E. Katz-Bassett. The Central Problem with Distributed Content: Common CDN Deployments Centralize Traffic In A Risky Way. In *HotNets*, New York, NY, USA, 2023.

[18] F. Wohlfart, N. Chatzis, C. Dabanoglu, G. Carle, and W. Willinger. Leveraging Interconnections for Performance: The Serving Infrastructure of a Large CDN. In *SIGCOMM*, 2018.

[19] K.-K. Yap, M. Motiwala, J. Rahe, S. Padgett, M. Holliman, G. Baldus, M. Hines, T. Kim, A. Narayanan, A. Jain, V. Lin, C. Rice, B. Rogan, A. Singh, B. Tanaka, M. Verma, P. Sood, M. Tariq, M. Tierney, D. Trumic, V. Valancius, C. Ying, M. Kallahalla, B. Koley, and A. Vahdat. Taking the Edge off with Espresso: Scale, Reliability and Programmability for Global Internet Peering. In *SIGCOMM*, 2017.