

A view of Internet Traffic Shifts at ISP and IXPs during the COVID-19 Pandemic

Anja Feldmann, Oliver Gasser, Franziska Lichtblau, Enric Pujol, Ingmar Poese, Christoph Dietzel, Daniel Wagner, Matthias Wichtlhuber, Juan Tapiador, Narseo Vallina-Rodriguez, Oliver Hohlfeld, Georgios Smaragdakis

Abstract. In this position paper, we report on a measurement study on Internet traffic shifts due to the COVID-19 pandemic using data from a diverse set of vantage points (one ISP, three IXPs, a metropolitan educational network, and a mobile operator). We observe that the traffic volume increased by 15-20% almost within a week—while overall still modest, this constitutes a large increase within this short time period. However, despite this surge, we observe that the Internet infrastructure is able to handle the new volume, as most traffic shifts occur outside of traditional peak hours. When looking directly at the traffic sources, it turns out that, while hypergiants still contribute a significant fraction of traffic, we see (1) a higher percentage increase in traffic of non-hypergiants, and (2) traffic increases in applications that people use when at home, such as Web conferencing, VPN, and gaming. While many networks see increased traffic demands, in particular, those providing services to residential users, academic networks experience major overall decreases. Yet, in these networks, we can observe substantial increases when considering applications associated to remote working and lecturing.

Key points:

- Relative traffic volume changes follow demand changes—causing “moderate” increases of 15-20% during lockdown for the ISP/IXPs in our study, but decreases up to 55% at the education network. Even after the lockdown, some trends remain: 20% at one IXP but only 6% at the tier-1 ISP.
- Most traffic increases happen during non-traditional peak hours. Daily traffic patterns are moving to weekend-like patterns.
- Online entertainment demands account for hypergiant traffic surge. Yet, the need for remote working increases the relative traffic share of applications like VPN and conferencing tools by more than 200%. At the same time, the traffic share for other traffic classes decreases substantially, e.g., traffic related to education, social media, and—for some periods—CDNs.
- At the IXPs, we observe that port utilization increases. This phenomenon is mostly explained by a higher traffic demand from residential users.
- Traffic changes are diverse, increasing in some network ports while decreasing in others. One example of the latter is the educational network, where we observe a significant drop in traffic volume on workdays after the lockdown measures loosened, with a maximum decrease of up to 55%. Yet, remote working and lecturing cause a surge in incoming traffic, e.g., for email and VPN connections. The EDU traffic shift is antagonistic, yet complementary, to the observations made in other vantage points.

Introductory Note

As a result of the ongoing COVID-19 pandemic, the population had to depend on their residential Internet connectivity for work, education, social activities, and entertainment. This opens questions on *i*) how traffic characteristics changed and *ii*) if these changes challenged the Internet infrastructure or operation and ultimately if internet operation need to be altered as a result. In this position paper, we summarize a measurement study on Internet traffic shifts due to the COVID-19 pandemic that will appear at the ACM Internet Measurement Conference 2020. Our study provides a empirical and multi-provider perspective on traffic shifts by using data from a diverse set of vantage points: one major tier-1 ISP, three IXPs of which two are located in Europe and one in the US, and one metropolitan area educational network. We summarize the most relevant findings and conclude with a discussion of lessons learned relevant to this workshop.

Vantage Points

We utilize network flows collected at vantage points at the backbone and peering points of a major Tier-1 Internet Service Provider (ISP), at the core of the Internet (IXPs), and at the edge (a metropolitan university network, a mobile operator).

- The ISP is a large Central European ISP that provides service to more than 15 million fixed line subscribers and also operates a transit network (Tier-1).
- We consider three major Internet Exchange Points (IXPs) in our study. The first one has more than 900 members, is located in Central Europe (IXP-CE) and has peak traffic of more than 8 Tbps. The second one has more than 170 members, is located in Southern Europe. The third one has 250 members, is located at the US East Coast.
- REDImadrid academic network interconnecting 16 independent universities and research centers in the region of Madrid. It serves nearly 290,000 users including students, faculty, researchers, student halls, WiFi networks (including Eduroam), and administrative and support staff. From each vantage point we analyze traffic flows to reason about COVID-19 related traffic shifts.
- The mobile operator also located in Europe with more than 40 million customers.

These vantage points enable us to holistically study the effects of the COVID-19 pandemic both from the network edge (ISP-CE/EDU) and the Internet core (IXPs).

Major Observations

Traffic volume changes

We observe a significant traffic evolution in 2020 at multiple Internet vantage points (ISP and IXP). In Figure 1 we show traffic changes from January 2020 until June 2020 for five different networks:

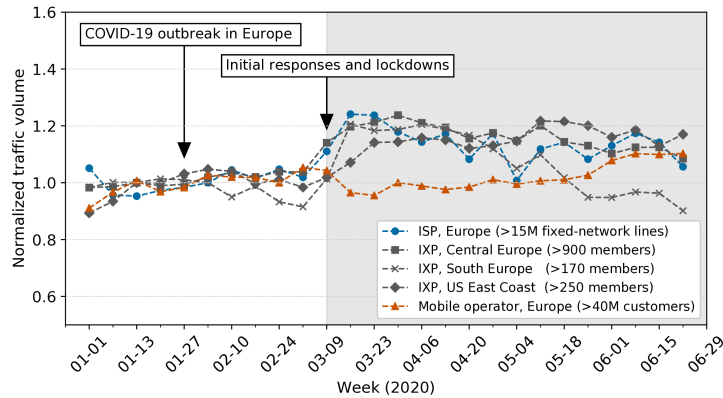


Figure 1: Traffic changes from January 2020 until June 2020 at multiple vantage points

Traffic demands for broadband connectivity, as observed at an ISP in Central Europe as well as at a major IXP in Central Europe and an IXP in Southern Europe increased slowly at the beginning of the outbreak and then more rapidly by more than 20% after the lockdowns started. The traffic increase at the IXP at the US East Coast trails the other data sources since the lockdown occurred several weeks later. While we observe this phenomenon at the ISP and IXP vantage points, one difference between them is that the relative traffic increase at the IXP seems to persist longer while traffic demand at the ISP decreases quickly towards May. This correlates with the first partial opening of the economy, including shop reopenings in this region in mid-April and further relaxations including school openings in a second wave in May. The decrease in mobile traffic can be explained by the fact, that people did not go out that frequently and would therefore use their home Wi-Fi more often instead of their phone's mobile data plan.

Some of the lockdowns were lifted or relaxed around May 2020. As people were allowed to perform some of their daily habits outside of their home again, we see a decrease of the traffic at the IXPs and the ISP; instead mobile traffic is now growing again.

Our findings align with insights offered by two reports published by Google, reports by Comcast, Nokia, TeleGeography, and two reports from Akamai.

Workday-weekend patterns

In light of the global COVID-19 pandemic a total growth of traffic is somewhat expected. More relevant for the operations of networks is how exactly usage patterns are shifting, e.g., during the day or on different days of a week. With the pandemic lockdown in March, this workday traffic pattern shifts towards a continuous weekend-like pattern.

Figure 2 shows a traffic pattern at the Central European ISP for three days: February 19, February 22, and March 25. The Internet's regular workday traffic patterns are significantly different from weekend patterns. On workdays, traffic peaks are concentrated in the evenings, see Figure 2. For instance, Wednesday, February 19 vs. Saturday, February 22, 2020: With the pandemic lockdown in March, this workday traffic pattern shifts towards a continuous weekend-like pattern, as can be seen in the daily pattern for March 25, 2020. More specifically, we call a traffic pattern a workday pattern if the traffic spikes in the evening hours and a weekend pattern if its main activity gains significant momentum from approximately 9:00 to 10:00 am.




Figure 2: Workday vs. weekend patterns before and after the lockdown

On a weekend day (orange bars) the pattern looks different, with a much steeper increase during the morning hours and a slower growth during the day, again reaching the traffic peak at around 21:00 in the evening. Since more people are staying at home during the day on a weekend compared to a working day, this behavior is affecting the traffic pattern as well.

And finally, when we investigate the traffic pattern of a working day during lockdown (gray bars), we see that it much more resembles a weekend day than a working day before the pandemic. This nicely visualizes the effect of lockdown measures on Internet traffic patterns.

We now classify every day based on its traffic pattern to being more workday-like or weekend-like.




Figure 3: Traffic on workdays during lockdown look more like weekend traffic

In Figure 3 we show the result of this classification. In the upper part of the graph we show days classified as weekend-like, in the lower part of the graph we show days exhibiting workday-like traffic patterns. If the classification is in line with the actual day (workday or weekend) the bars are colored blue, otherwise they are colored in orange.

We find that up to mid-March, most weekend days are classified as weekend-like days and most workdays as workday-like days. After mid March, however, the majority of all days are classified as weekend-like, no matter if they are workdays or weekends, and we therefore see a lot of misclassified working days (orange bars). The only exception is the holiday period at the beginning of the year. This pattern changes drastically once the confinement measures are implemented: Almost all days are classified as weekend-like.

As previously seen in Figure 2, people using the Internet from home during the day exhibits more of a weekend-like traffic pattern. Additionally, as can be seen by the increasing length of the bars starting around mid March, we see an increase in the overall traffic per day. This increase, however, is not equally distributed over the whole day but is mostly happening in off-peak hours, which can be seen in Figure 2.

Growth of specific classes of traffic

These observations raise the question of the cause for this significant traffic growth and shift in patterns, given that many people are staying at home for all purposes, e.g., working from home, remote education, performing online social activities, or consuming entertainment content. The increased demand in entertainment, e.g., video streaming or gaming, may imply an increase in hypergiant traffic.

After having analyzed changes in traffic volumes and diurnal patterns, let's now look at specific classes of traffic. Since port-based classification mixes together a lot of traffic using common ports such as TCP/80 or TCP/443, we use a combination of port-based and AS-based classifications to classify the traffic at the ISP into different groups.




Figure 4: Change of traffic patterns for specific classes of traffic

Figure 4 shows the traffic changes comparing the months of March, April, and June to our *base* week of February. We group traffic into seven different traffic classes. Traffic changes are shown for each hour of the day for all days of the week. While it would go beyond the scope of this article to discuss each and every change that is visible in the figure, we will highlight the most interesting ones.

We see a strong increase in the traffic associated to web conferencing, video, and gaming traffic in March as a result of the increasing user demand for solutions like Zoom or Microsoft Teams. Also, as people spend more hours at home, they tend to watch videos or play games, thus increasing entertainment traffic demands. Interestingly, we also see a decrease in educational traffic in these vantage points.

In April and June, web conferencing traffic is still high compared to the pre-pandemic scenario, while we see a slight decrease in CDN and social media traffic. During these months many people are still working from home, but restrictions have been lifted or relaxed, which leads to an increase in in-person social activities and a decrease in online ones. We will continue our measurements in search of the *new normal*.

We note that traffic changes are diverse and highly dependent on the vantage point. For instance, traffic shifts in the REDIMadrid academic network show an antagonistic but complementary behavior. While we observe a 55% drop in traffic volume on workdays, even after the lockdown measures loosened, remote working and lecturing cause a surge in incoming traffic for email, web, and VPN connections.

Lessons Learned

Internet operation during the pandemic: a success story. Unexpectedly, the Internet held up to this unforeseen demand with no reports of large scale outages or failures. At the beginning of the pandemic, changes in user demand for online services raised concerns for network operators, e.g., to keep networks running smoothly especially for life-critical organizations such as hospitals. In fact, the pandemic increased the demand for applications supporting remote teaching and working to guarantee social distancing as shown in our analysis across all vantage points. The Internet could handle this new load due to the flexibility and elasticity that cloud services offer, and the increasing connectivity of cloud providers. Our results confirm that most of the applications with the highest absolute and relative increases are cloud-based. Moreover, the adoption of best practices on designing, operating, and provisioning networks contributed to the smooth transition to the new normal. Due to the advances in network automation and deployment, e.g., automated configuration management and robots installing cross connects at IXPs without human involvement, it was possible to cope with the increased demand. For example, DE-CIX Dubai managed to quickly enable new ports within a week for Microsoft which was selected as the country’s remote teaching solution for high schools. In summary, our study demonstrates that over-provisioning, network management, and automation are key to provide resilient networks that can sustain drastic and unexpected shifts in demand such as those experienced during the COVID-19 pandemic.

Taming the traffic increase. In our study, we report an increase in traffic in the order of 15-20% within days after the lockdown began. This is in line with reports of ISPs and CDNs as well as IXPs. Typically, ISPs and CDNs are prepared for a traffic increase of 30% in a single year period. While these are yearly plannings, the pandemic created substantial shifts within only a few days. As a result, ISPs either needed to benefit from over-provisioned capacity—e.g., to handle unexpected traffic spikes such as attacks or flash-crowd events—or add capacity very quickly. We observed port capacity increases in the order of 1,500 Gbps (3%) across many IXP members at one observed IXP alone. Beyond our datasets, some networks publicly reported that traffic shifts due to the pandemic resulted in partial connectivity issues and required new interconnections.

When we turn our attention to traffic peaks, we notice that the increase is even smaller. Traffic engineering focuses on peak traffic increase as this requires more network resources. The effect of the pandemic fills the valleys during the working hours and has a moderate increase in the peak traffic, which can be handled by well-provisioned networks that are prepared for sudden surges of peak traffic by 30% or more, due to attacks, flash-crowds, and link failures that shift traffic to other links. One concern that network operators raised in March brought awareness to network instabilities that might occur due to traffic shifts. While on the one hand we find no evidence that the traffic shifts due to the pandemic impact network operation of our vantage points, individual links experience

drastic increases in traffic—way beyond the overall 15-20%. Such increases arise unexpectedly to some network operators and may create a need for port upgrades. On the other hand, the vantage points in this paper range from extremely large to moderate sizes with sufficient resources and a lot of experience in network provisioning and resilience. In general, smaller networks with limited resources may not be able to plan with sufficient spare capacities and fast enough reaction times to compensate for such sudden changes in demand.

Substantial shift in traffic pattern. From a network operator perspective, coping with the pandemic has required some port capacity upgrades but otherwise does not appear to impact operation. The ability of network operators to quickly add capacity when needed highlights that the Internet infrastructure works well at large, despite some challenges to access data centers imposed by the lockdown. From the perspective of the traffic mix, the pandemic, however, results in substantial changes in traffic, ranging from shifted diurnal pattern to traffic composition. This represents a remarkable shift in Internet traffic that is, based on our observations, handled surprisingly well by the Internet core at large supposedly because many operators are prepared and can react quickly to new demands. While the pandemic represents a rather extreme and exceptional case, one may argue that with the growing intertwining of the Internet and our modern society such events can occur more often. In any case, the COVID-19 pandemic highlights that user behavior can change quickly and network operators need to be prepared for sudden demand changes.

Further Details in the Paper

This position paper focuses on discussing the big picture of our study on changes that we have seen in Internet traffic due to the pandemic to discuss lessons learned relevant for discussion during the IAB workshop. We refer the reader to our IMC paper for details on the measurements and further aspects studied.

The Lockdown Effect: Implications of the COVID-19 Pandemic on Internet Traffic Anja Feldmann, Oliver Gasser, Franziska Lichtblau, Enric Pujol, Ingmar Poese, Christoph Dietzel, Daniel Wagner, Matthias Wichtlhuber, Juan Tapiador, Narseo Vallina-Rodriguez, Oliver Hohlfeld, Georgios Smaragdakis
ACM Internet Measurement Conference (IMC) 2020

The preprint of the paper is available at (<https://arxiv.org/abs/2008.10959>).

The paper sheds light on further aspects:

- How does the growth for hypergiant ASes differ compared to other ASes?
- How much did QUIC traffic increase during the lockdown?
- How does gaming traffic change at the Southern European IXP?
- Is VPN traffic increasing?
- What effect does the lockdown have on traffic at a large metropolitan academic network?

Questions to be discussed with workshop participants

- How did the situation during the lockdown affect smaller networks with less resources?
- Are there any “stories” to be told regarding mitigation strategies to tame the traffic increase
- How did the traffic levels/patterns look like after the lockdown for other networks?
- Are there better strategies out there than overprovisioning?