

INTERNET PERFORMANCE REPORT

COVID-19 Impact Edition

2020 EDITION

The Internet Performance Report: COVID-19 Impact Edition is a measurement-based study of Internet health, examining the various network infrastructures critical to modern content delivery, namely those belonging to ISPs, as well as cloud, CDN and DNS providers.

CONTENTS

3	EXECUTIVE SUMMARY
5	METHODOLOGY
5	The ThousandEyes Platform
6	Outage Detection
6	Performance Data Collection
7	Important Disclaimers
8	FINDINGS AND TAKEAWAYS
11	Internet Service Providers
24	Cloud Providers
33	CDN and DNS Providers
39	CONCLUSION AND RECOMMENDATIONS



EXECUTIVE SUMMARY

The *Internet Performance Report: COVID-19 Impact Edition* is a measurement-based study of Internet health, examining the various network infrastructures critical to modern content delivery, namely those belonging to Internet Service Providers (ISPs), as well as cloud, Content Delivery Network (CDN) and Domain Name System (DNS) providers.

By tracking network disruptions and key performance indicators for the various components of the Internet delivery chain between January and July of 2020, this report uncovers important insights into the resilience and behavior of the global Internet under “normal” conditions, as well as under conditions without precedent in its history.

Now, more than ever, the Internet has become a critical lifeline for billions of people around the world impacted by COVID-19. A vast portion of the workforce is working remotely and the Internet has effectively become a digital tether to the world, connecting every one of us to critical services and each other. Despite countries around the world slowly beginning to reopen their economies, many enterprises will continue to support remote workers for the foreseeable future—which means the Internet is increasingly the “x-factor” for employee productivity and business continuity.

While the Internet has come to be regarded as a key facilitator of enterprises’ digital transformation initiatives, the Internet is still a “black box” for many organizations, which puts IT investments into public cloud, SaaS, and SD-WAN at risk, and it makes managing a remote workforce more challenging.

With this report, we hope to provide visibility into the operational and performance signatures that underlie the Internet, both on a global level and regionally, so that IT leaders and practitioners can better understand this “network of networks.”

Key findings from the report include:

- **A precipitous rise in network disruptions beginning in March 2020**
- **Not all outages are created equal—and more disruptions do not indicate systemic performance degradation**
- **Macro traffic shifts impact Internet-related infrastructures**
- **Cloud provider networks demonstrate greater resilience than ISP networks**
- **Regional availability and performance differs across ISP and cloud networks**

Enterprise IT teams can draw conclusions from the findings in this study to make informed architecture and digital supply chain decisions.

INTERNET PRIMER

When using the Internet, there are a number of fundamental services and protocols that are required to work together in order to ensure digital experience is as optimal and efficient for all users as possible. This section defines and describes those services and associated protocols, their operation and role in the service delivery chain.

Any conversation around the scale and scope of the Internet warrants a discussion of the various services and protocols that are integral to its functionality. Widely referred to as a “network of networks,” the Internet relies on certain foundational protocols, such as Border Gateway Protocol (BGP) and DNS, in order for networks to communicate and users to quickly locate and transit to their intended destinations.

In addition, services like CDNs and public cloud providers are being increasingly relied on for geographically distributed services, aimed at bringing content closer to end users—with the added benefit of reducing loads on backbone networks. Without these services (or some home-grown alternative), business SaaS applications and digital entertainment such as gaming and media streaming would not have been viable. Together, these protocols and services work to provide a fast and resilient Internet to every user and, as such, they are very much integral to its definition. For these reasons, we’ve expanded coverage of the “Internet” in this report to include cloud networks, as well as CDN and DNS services.

For an in-depth explanation of the protocols and services discussed above, the problems they address, and the underlying vulnerabilities they introduce, see the following companion eBook: [Internet Fundamentals: Underlying Network Infrastructures Explained](#)

Below is an example of how ISPs, DNS, CDNs, and cloud providers would, together, facilitate a user interaction with a web service.

As customer digital experience is a primary objective of many enterprises today, let’s focus on an example where customers from around the world are accessing a website over the Internet. Every digital user goes through four steps in their journey to your brand. Each step is supported by different Internet infrastructure and providers. The first and most critical step is to look up your Internet address (IP address) via the Domain Name System (DNS), as it allows every subsequent step to take place. Second, once a URL is mapped to an IP address, the user journey proceeds across various Internet service providers (ISPs), using the border gateway protocol (BGP), to where the app or website is hosted. Third, users connect to a content delivery network (CDN) provider like

Akamai for locally cached content. Fourth, either users are redirected to a hosted application to fetch content, or the content is retrieved by the CDN (either hosted on-prem or within a cloud provider) and served directly to the user.

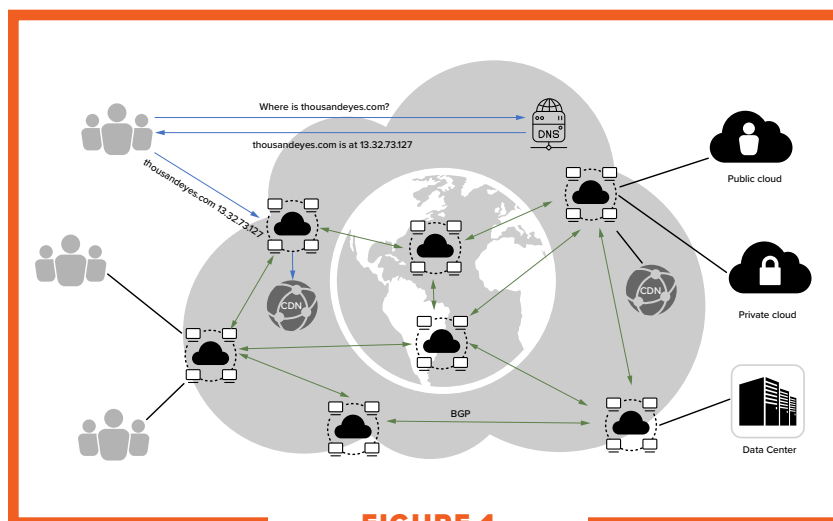


FIGURE 1

Example showing customer digital experience over the Internet.

REPORT METHODOLOGY

The findings presented in this report are based on measurements collected between January and July 2020. These measures were taken using the ThousandEyes platform, which yielded both performance metrics, as well as large-scale detection of network outages within publicly routed networks.

THE THOUSANDEYES PLATFORM

The ThousandEyes platform uses active monitoring techniques to gather network metrics such as loss, latency and jitter along with in-depth path metrics with detailed layer 3 hops. The platform also leverages a collective, abstracted data set derived from all network tests to detect outage events in various service provider networks. Both aspects of the ThousandEyes platform are used to collect the data that is surfaced in this report.



FIGURE 2

ThousandEyes active monitoring infrastructure powers the Internet Performance Report.

OUTAGE DETECTION

ThousandEyes leverages aggregated network telemetry data from ThousandEyes global sensor network to detect network outage events taking place across ISP, public cloud and edge service networks such as CDN, DNS and SECaaS. Network outages are detected based on billions of measurements to hundreds of apps and services crossing thousands of publicly routed networks each day. Outages are defined as terminal events, with 100% packet loss in the same ASN during a given period of time. A proprietary algorithm is used to reduce triggers based on transitory performance degradation, rather than meaningful service disruption. Therefore, thresholds are in place in terms of impact on ASN infrastructure, as well as ThousandEyes sensors and destination services. Multiple probing techniques are used, each utilizing different packet streams and protocols, which prevents false positives in the data set. This outage detection mechanism is highly specific, enabling isolation of traffic termination incidents down to the interface IPs and geo-location of the infrastructure involved.

PERFORMANCE DATA COLLECTION

Many of the findings presented in this report are based on data gathered from continuously monitoring bi-directional network performance such as latency, packet loss and jitter between fixed vantage points connected to ISP networks within three global regions—North America, Europe, and Asia-Pacific—as well as cloud Inter-region performance for Azure, AWS, and GCP within those same geographic regions.

Many of the insights, conclusions and recommendations provided in this report were derived from analysis of these metrics. While the principles of data collection such as metrics gathered remained consistent within the study, multiple test methodologies were deployed.

INTRA-REGION ISP MEASUREMENTS	US BROADBAND ISP MEASUREMENTS	CLOUD INTER-REGION MEASUREMENTS
Network performance metrics gathered from regionally distributed vantage points connecting to one another via various in-region ISP networks, including North America, Europe, and APAC.	Network performance metrics for two broadband service providers across six US cities.	Inter-region performance within Azure, AWS, and GCP within three regions—North America, Europe, and APAC.

IMPORTANT DISCLAIMERS

Active performance measurements in this report are collected by ThousandEyes vantage point agents that are connected to Tier 2, Tier 3 and broadband ISPs, as well as cloud backbone networks in the case of agents in cloud provider hosting regions. Measurements taken from other locations or ISP connections in similar locations may yield different results. This highlights the importance and complexity of geo-location and network peerings as factors in network performance on a global basis.

The results presented in this report have been gathered during the timeframe mentioned in the methodology section. They are intended to illustrate deviations from baseline and performance changes over time, not to comprehensively define performance expectations in a given region. Measurements taken from similar or different locations before or after the documented time frames could yield different results, which is why it is critical to continuously monitor and measure for changes.

Outages described in this report are based on ThousandEyes globally distributed vantage points. Given the massive, distributed nature of the Internet, no organization can claim total visibility into all network disruptions; however, the outage figures provided are illustrative of availability trends over time.

ThousandEyes vantage points are used by hundreds of the world's largest enterprises, financial institutions, cloud and SaaS providers to actively monitor and provide real-time business and operational insights. ThousandEyes visibility data is trusted to automate service path remediation for large-scale Internet-dependent services. Vantage point agents and monitoring methodologies are continuously optimized for accuracy.

Enterprises looking to establish their specific performance baselines and operational metrics should utilize the data in this report as a guide and collect performance measurements from their own data center, office and VPC locations.



FINDINGS AND TAKEAWAYS

While outages occur in the course of normal network operations, global network disruptions increased dramatically in March 2020, coinciding with pandemic-related social distancing orders taking effect across nearly every region. This increase in disruptions, however, should not be taken as an indication that Internet infrastructures did not “hold up” under the strain of rapidly shifted traffic patterns. Evidence presented below, based on large-scale network measurements collected between January and July 2020, suggest that recent increases in disruptions are not due to network duress, but instead the result of operators adjusting their networks to accommodate changes in traffic pattern and load. Further, some of these disruptions may not have meaningfully impacted Internet users, as in particular regions and types of providers, a large proportion of them take place outside of traditional business hours.

Considering the sudden, unprecedented change in traffic flows brought on by the pandemic, the Internet has held up surprisingly well overall, with no systemic performance degradation. A temporary elevation in disruptions may indicate a normal adaptation by network operators in response to macro traffic shifts, and enterprises should consider planning for greater Internet volatility when these shifts occur.



FINDING

EVIDENCE



Global Internet disruptions saw an unprecedented, dramatic rise in March, coinciding with pandemic-related “shelter-in-place” orders, and remained elevated through the first half of 2020 compared to pre-pandemic levels.

- 63% increase in global disruptions in March over January
- 45% higher in June compared to January
- APAC ISPs - 99% more outages in March vs January
- North America ISPs - 65% more outages in March vs January



Increased disruptions across provider networks did not appear due to network duress caused by congestion, but rather the result of increased traffic engineering activity.

- Performance indicators, such as traffic delay, loss, and jitter generally remained within tolerable ranges between March and June, showing no evidence of systemic duress.
- Network disruptions post-February were longer and involved more infrastructure, indicating control plane related disruptions.



Cloud provider networks are more stable than ISP networks; however, when outages do occur in their networks, they are more likely to impact users.

- Despite an unusual rise in outages across every region (up to 150% in North America), their overall numbers were relatively low.
- Between January and July, cloud providers experienced ~400 outages versus more than ~4500 in ISP networks (excluding China).
- A large portion of outages, particularly in North America, occur during peak business hours, potentially increasing their impact on users.



ISPs experienced more outages than cloud provider networks; however, their business impact varied by region, possibly reflecting local workplace and/or organization practices.

- North American ISP outages are highest outside of business hours, while outages in EMEA and APAC are more distributed from a time of day standpoint.
- Fewer disruptions take place on weekends than weekdays across all regions and types of providers—possibly indicating fewer staff resources available to make network changes and/or lower overall traffic levels during this period.



The networks of key services critical to service reachability and optimal delivery of traffic—CDN and DNS providers—held up well under increased and changing traffic conditions.

- CDN and managed DNS providers experienced few outages in the first half of 2020
- Public DNS resolver performance changed between February and March, possibly due to traffic pattern changes brought on by work from home activity

GLOBAL NETWORK DISRUPTIONS PRE AND POST-PANDEMIC

Looking at disruptions in key traffic delivery networks between January and July of 2020 provides useful insights into the health and behavior of the global Internet under “normal” conditions, as well as under conditions without precedent in its history.

Since outages are often based on unpredictable events their numbers can vary dramatically from one day to the next, but significant month-over-month increases are unusual and usually temporary. For example, in October 2019, several large network operators experienced a series of disruptions that pushed outages to 19% higher than the previous month; however, the following month, disruptions declined to typical levels which range between 700-800 per month.

January and February outage levels were similar to previous periods; however, March brought a precipitous rise in global network disruptions that coincided with large-scale traffic shifts brought on by lock-down and shelter-in-place orders.




FIGURE 3

Global network outages increased 63% in March over January levels and remained elevated through the first half of 2020.

Globally, network outages were 63% higher in March compared to January, and although they declined slightly in subsequent months, by July they had not yet returned to pre-March global levels. But not every provider experienced similar levels of disruption.

Although both CDN and DNS providers did see a few unusual spikes in outages, overall they maintained good availability between January and July 2020. ISP and cloud providers, however, experienced significant increases in network outages during that same period.

INTERNET SERVICE PROVIDERS

Note: The term “Internet service provider” is used here to refer to a broad category of network transport providers, including telecom, transit, and broadband providers.




FIGURE 4

Network disruptions rose sharply between January and March 2020, for both ISPs and cloud providers

Internet service provider outage levels vary by region; however, every region saw an increase in network disruptions post-February. The most significant increase occurred in two regions, North America and Asia-Pacific (excluding China). Both of these regions experienced a substantial increase in outages in March, with outages nearly 100% higher in Asia-Pacific (excluding China) and 65% higher in North America compared to January. After peaking in March, outages gradually declined in both regions over the next few months, eventually plateauing in North America, and returning to pre-March levels in Asia-Pacific.

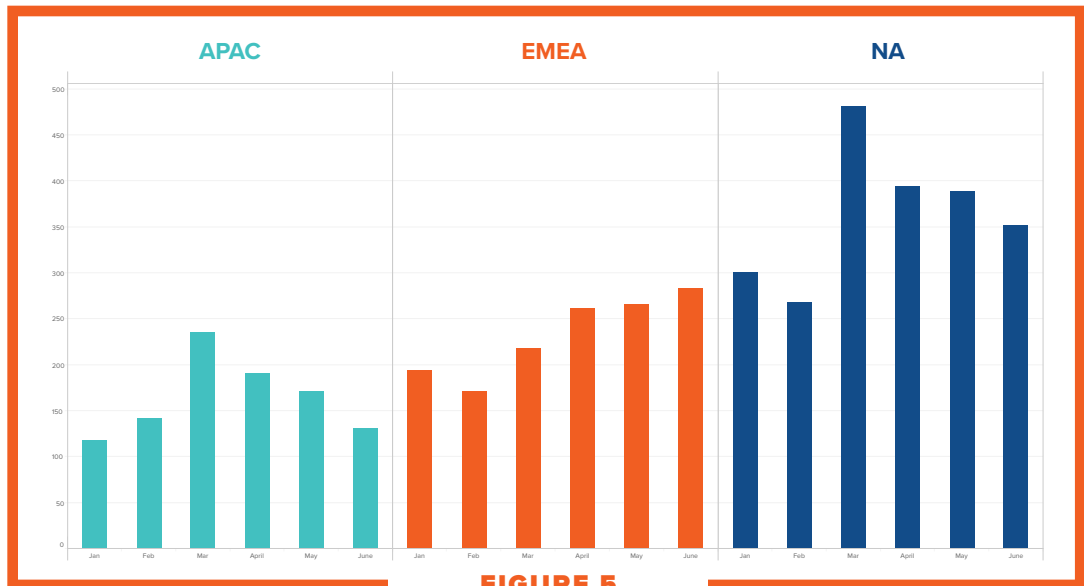


FIGURE 5

ISP outages varied by region

ISP outages in EMEA, however, followed a very different pattern, with less dramatic but steady month-over-month increases, and June marking a new regional record for the number of outages. Variations in national responses to the pandemic, in terms of ordering and lifting lockdowns, may have staggered disruptions, leading to less concentrated increases in that region.

Looking at disruptions by week, it's common to see large fluctuations, both pre and post-March. For example, in North America, the second and third weeks of January had nearly 100 outages each, but dropped to under 50 in the fourth week. These peaks and valleys continued through June, but March was notable in that the first, third and fourth weeks each broke previously set records for weekly outage levels, first breaking 100 and continuing to increase or stay above that level until early April.

The highest weekly levels in March coincided with or followed state-wide shelter-in-place orders. Two of the most populous states in the US, California and New York, issued lock down orders the week of March 16th, transitioning a large portion of the workforce to the home. Other states and counties issued similar orders that week and in the following week.



FIGURE 6

North American outages fluctuate by week (Monday-Sunday), but the three-week moving average shows March incline and subsequent plateauing.

Given the sudden, large-scale change in traffic patterns and volume (as more activities were conducted digitally), a corresponding increase in network issues may not be surprising to some; however, the low number of outages in early April and May suggest that outage levels don't directly correlate with traffic levels, as according to ISP, Internet exchange, and content delivery providers, traffic volumes were still increasing during this period.

Weekly network outage patterns in Asia-Pacific share some similarities with North America, with sharp increases and decreases throughout the examined period, but higher than normal concentrations in March and parts of April. Both regions also had a spike in disruptions in January, but the increase was more pronounced in Asia-Pacific (relative to its typical levels), and occurred only during the week of January 12th, whereas in North America, the increase started in the previous week.



FIGURE 7

APAC outages fluctuate by week (Monday-Sunday), but the three-week moving average shows March incline and post-April decline

The January increase in disruptions aligns with some parts of Asia-Pacific returning to the workforce after end-of-year holidays, which in North America begins earlier for many workers, possibly explaining a rise there beginning the week of January 6th. EMEA experienced a similar spike in outages in January, though it occurred even later in the month than Asia-Pacific, during the week of January 19th.




FIGURE 8

EMEA outages fluctuate by week (Monday-Sunday), but the three-week moving average shows a steady increase throughout the first half of 2020.

Unlike North America and Asia-Pacific, March was unremarkable in EMEA from an outage standpoint, though its final week did have higher than typical levels that increased slightly every other week over the next six weeks. Their numbers finally seemed to stabilize in the latter part of May; however, early June brought another sharp increase that set a new weekly record for the region.

Network disruptions were fairly low across all of the regions covered in this study in the final week of June. For Asia-Pacific, this follows an overall downward trend; however, in North America and EMEA, it could either indicate increasing normalization or simply a temporary ebb.

OUTAGE IMPACT ON INFRASTRUCTURE AVAILABILITY VARIED

While the number of outages within a region or provider is useful to understand overall stability, the impact of outages on service availability can vary widely depending on their scope (i.e. the amount of infrastructure impacted). A small number of outages could conceivably have more of an impact on network availability than a large number of outages if they were severe enough. When regional outages are viewed from the standpoint of the number of interfaces within provider networks unavailable at any given time (as a percentage of monitored interfaces), several notable findings emerge.

All three regions share similar outage impacts, where increases and decreases in infrastructure availability either coincide with or are slightly staggered in relation to one another. While many large ISPs operate across multiple regions and may experience an outage incident that impacts different geographical parts of its network, those incidents are not common and wouldn't account for the consistently similar availability trends. Given when the most significant peaks are occurring, they are likely related to social influences common across those regions, such as returning to the workplace after a holiday period (as suggested earlier) or working from home due to pandemic-related mandates.




FIGURE 9

Percent of ISP network infrastructure unavailable, relative to monitored infrastructure.

Normalization based on infrastructure coverage and impact also contextualizes outage numbers. For example, even though the number of outages in EMEA the week of January 20th (78) was almost double that of January 6th (39), the impact on infrastructure was less severe than the earlier week. Asia-Pacific is also shown to have suffered more availability damage than North America and EMEA in the first half of 2020. Despite higher numbers of disruptions observed in North America, ISP networks in that region are not “less available” than those of Asia-Pacific and EMEA.

NETWORK PERFORMANCE INDICATES NO SYSTEMIC DURESS

Given that increases in network disruptions (and lower infrastructure availability) coincided with large-scale traffic shifts throughout the first half of 2020—whether post end of year holidays or post shelter-in-place/reopening orders—is it conceivable that traffic loads constrained network availability across provider networks, leading to increased disruptions?

A massive traffic spike, such as one generated by a large-scale DDoS attack or BGP route leak, can overwhelm networks, leading to extreme congestion within provider networks that effectively prevents large portions of traffic from reaching a service. However, in considering the root cause of the increased level of disruptions, it's useful to understand the broader performance context, as the occurrence of outages is not, on its own, evidence of network duress.

Network duress due to traffic surges would first manifest as degraded performance. Elevated latency, jitter (latency variance) and packet loss can be symptoms of congestion, as router queuing may delay traffic, create variability in interpacket gaps, and lead to traffic drops, which may then trigger packet retransmissions, creating further delay and, potentially, more congestion and loss.

However, none of the routes examined between major transit points within each region showed signs of systemic degradation throughout the first half of 2020. Network jitter remained within tolerable thresholds across all regions, averaging less than 1ms and under 14ms in 95% of the collected data set.

		Jan	Feb	Mar	April	May	June
NA	Packet Loss (avg)	1%	1%	1%	3%	5%	3%
	Packet Loss (95P)	100%	100%	100%	100%	100%	100%
	Jitter (avg)	0.35	0.33	0.35	0.61	0.61	0.63
	Jitter (95P)	11	12	10	13	14	14
EMEA	Packet Loss (avg)	0%	0%	1%	1%	0%	0%
	Packet Loss (95P)	100%	100%	34%	100%	44%	100%
	Jitter (avg)	0.18	0.22	0.23	0.23	0.26	0.21
	Jitter (95P)	6	7	6	6	6	5
APAC	Packet Loss (avg)	2%	2%	3%	2%	1%	0%
	Packet Loss (95P)	100%	100%	100%	100%	100%	100%
	Jitter (avg)	0.25	0.30	0.34	0.32	0.31	0.31
	Jitter (95P)	2	4	8	6	3	4

FIGURE 10

Network performance metrics stayed within acceptable thresholds between January and July 2020.

Packet loss was lowest in EMEA, remaining at or under 1% each month. It also stayed fairly low in Asia-Pacific, despite an increase to 3% in March. In North America, however, there was a rise in packet loss beginning in April and peaking in May at 5%, though this level of loss may not necessarily be disruptive in terms of user experience.

Perceived performance, particularly for sensitive applications, such as VoIP, depends on the inter-relationship of latency, jitter and packet loss. For example, elevated packet loss may have a lower impact on performance if jitter is low and vice versa. As either latency, packet loss, or jitter increase, overall degradation tolerance lowers, so optimal thresholds for each of these will vary depending on network conditions, as well as application sensitivity.

Some amount of network delay is normal and expected, as the rate at which light travels through fiber is bound by the laws of physics and is based, at a minimum, on the distance between two points. Unnecessary delay, however, can be introduced by sub-optimal routes, network congestion and packet loss, so it is a potential indicator of both network health, as well as network state changes (e.g. route changes).

Average latency each month relative to January (expressed as a percentage increase or decrease) was significantly different between regions. In Asia-Pacific, network latency improved over January in each subsequent month, with June latency (aggregated across monitored in-region routes (see Methodology section for details) 14% faster than January. In EMEA, latency increased slightly February through April, but returned to January levels in May. In North America, latency increased slightly February through April, but returned to January levels in May.




FIGURE 11
Average percent change in latency over January

The largest change in latency was in North America, starting in April and continuing through June, when it was 18% higher than January.

Looking more closely at where the increased loss and latency was taking place in North America between April and July, the sources were isolated to a specific location within two ISPs—a broadband provider in San Jose, California and a transit provider between Dallas and Mexico. Higher levels of loss or some congestion within these providers may have contributed to higher latency in some cases, though route changes were also found to be a factor. For example, in the case of the transit provider, increased latency coincided with a change in route between various cities within the U.S. and Mexico City.

These two providers/locations were outliers from a performance degradation standpoint. When their performance is excluded from the data set, overall ISP performance in North America is similar to pre-pandemic levels.

DISRUPTIONS DUE TO INCREASED TRAFFIC ENGINEERING

Since network congestion does not appear to be responsible for increased network disruptions, it's useful to examine the outages themselves, as the characteristics of an outage can provide clues to its underlying cause. For example, an outage within a specific autonomous system (AS) affecting a lot of infrastructure and taking place across multiple locations simultaneously, is likely caused by a misconfiguration or control plane issue, rather than a router failure or other localized issue, such as power outage or damaged cable.

Modern network operators architect their networks to avoid single points of failure and are generally resilient from a layer 2 standpoint, so outages resulting from infrastructure failure are not common. Exceptions do occur, however, as when a router failure at a DE-CIX exchange point caused a significant service disruption in 2018, or when two fibers cut later that year led to a widespread Comcast outage. Though in the latter example, the impact of the fiber cuts on the network's control plane is what ultimately led to the massive scale of the incident. Additionally, the notion that physical infrastructure would globally become more fragile post-February is unrealistic, so infrastructure issues are not likely wholly to blame for the increases.

Eliminating both widespread network congestion, as well as infrastructure-related issues as the primary cause of higher than normal levels of outages in each region leaves internal network state changes or control plane issues (e.g. BGP route leaks) as potential sources. Both of these are attributable to network engineering activity, which reportedly increased as a result of traffic shifts brought on by shelter in place announcements.




FIGURE 12

Average number of interfaces per outage

¹Based on [public statements by ISPs and IXPs](#), as well as privately communicated anecdotes by providers and their enterprise customers.

The size and duration of outages—both pre and post pandemic—particularly in North America and EMEA, strongly indicates network state changes (intentional or not) as a significant contributor to their numbers. In EMEA, the average number of interfaces per outage both pre and post February was between six and eight. The average number of interfaces per outage incident in North America increased from approximately six in January to nearly eight in March, coinciding with the largest rise in disruptions in that region.

Network outage size did not vary much in Asia-Pacific between January and July, staying between five and six throughout; however, as with other regions, these ranges point to control plane related disruptions.

Network outage duration followed a slightly different regional pattern, where outages in North America and Asia-Pacific outages were, on average, longer in April, while in EMEA outages were longest in February and May.




FIGURE 13

ISP outage duration by region.

Long-lasting outages that involve more than a handful of network interfaces may be the result of planned downtime or the unintended result of network changes.

ISP networks are complex, diverse environments, so when changes or adjustments are made to accommodate peers and customers, they can have unintended, far-reaching consequences. Although downtime may be anticipated for some changes and planned as part of a maintenance window or implemented outside of peak usage hours, some changes may need to be made during typical usage hours.

Network operators assess service requirements based on known or predicted traffic patterns. In the case of the “known,” this is typically based on the physical location of the user population, pattern of usage, along with the type of services accessed. The “predicted” is typically based on foreseen events, such as seasonal holidays, major sporting competitions, streaming game releases, or shopping rituals (e.g. Cyber Monday in the U.S.). Given the unprecedented scenario network operators faced in the wake of the COVID-19 pandemic, any assertion of known and predicted traffic usage was lost. This meant that any provisioning or changes required were reactive, potentially increasing the likelihood of them disrupting service availability.

USER IMPACT VARIES BY REGION

The timing of an outage is one of the most important factors in determining its impact on users. The occurrence of an outage and its scope and duration is less meaningful if the outage occurs during a period when users are unlikely to be impacted, for example very early morning hours or on weekends when Internet usage is typically lower.

As outages increased at various points throughout the first half of 2020, the day of week and time of day patterns seen in each region prior to March persisted, suggesting that activities related to shelter in place traffic shifts did not significantly impact local operational practices.

Across all regions, fewer network disruptions take place on weekends than weekdays. As outages increased, more outages occurred on these days; however, they constituted less than 25% of all outages across nearly every region between January and July.




FIGURE 14

Weekend (Saturday-Sunday) outages by time of day

This pattern was most prominent in North America, where during months with higher levels of disruptions, only 15-17% occurred during weekends. An even distribution based on the number of days in a week would be just under 30%. While the difference is not extreme, it is notable that weekends do not appear to be periods during which more maintenance or downtime is planned.




FIGURE 15

Less than 25% of outages in North America occur on weekends

EMEA is similar to North America in this respect, consistently remaining under 25% each month.




FIGURE 16

Less than 25% of outages in EMEA occur on weekends

Asia-Pacific was the only exception to this pattern, with May and June hitting 29% and 27% respectively.



FIGURE 17

More outages take place on weekends in APAC than other regions

During weekdays, regional preferences again held constant irrespective of traffic volume. Looking at weekday outages from the standpoint of time of day occurrence and highlighting those falling within traditional business hours of 9AM to 6PM, localized by region, important differences emerge.

In North America, ISP network disruptions appear to be influenced by traditional workforce cycles. Notably, outage density is much higher outside of traditional business hours, especially concentrated between midnight and 4AM, when Internet usage is lower, and maintenance or change windows are less likely to be disruptive to users.




FIGURE 18

North American outages by time of day, Monday-Friday (9AM-6PM ET highlighted)

This “off hours” preference is evident in January and February but becomes particularly clear in March as the outage data set increases.

EMEA ISPs have higher levels of outages in early morning hours—again, indicative of planned change windows—however, aside from those hours, more outages take place during business hours than, for example, early evening hours. So outages in that region may be slightly more disruptive from a business continuity standpoint than those in North America.




FIGURE 19

EMEA outages by time of day, Monday-Friday (9AM-6PM GMT+2 highlighted)

Outages in Asia-Pacific were more indiscriminate from a time of day standpoint. Fewer outages took place within traditional non-peak usage (early morning) hours, with higher concentrations taking place during business and early evening hours. Given time of day occurrence, outages in Asia-Pacific, though on average smaller and shorter in duration than other regions, appear to be more disruptive to local users. However, these disruptions would be brief and likely impact fewer users.





FIGURE 20

APAC outages by time of day, Monday-Friday (9AM-6PM JST highlighted)


Looking at the number of outages occurring within traditional business hours (Monday to Friday, 9AM to 6PM local time) in relation to outages outside of that window (i.e. weekends and non-business weekday hours), far fewer impact typical business operations than the total number observed for each region.

In North America, only approximately 25% of outages occur within business hours, and in March, when outages were at their peak in the region, less than 21% took place during business hours.




North American outages based on local business hours impact

In EMEA, approximately 30% of outages occur within typical business hours, with the rest taking place either outside of these hours on weekdays or over the weekend.



EMEA outages based on local business hours impact

While the Asia-Pacific region had more monthly variation in user impact between January and July 2020, only approximately 29% impacted business users overall, largely due to a higher volume of outages taking place on weekends versus EMEA. Given that EMEA outages involved more infrastructure and were longer lasting, when they did take place during business hours, their effect would have been more profound than those in Asia-Pacific.



Asia-Pacific outages based on local business hours impact

Both Asia-Pacific and EMEA saw a proportional increase in business impacting outages as overall outages increased in their respective regions, while that was not the case in North America, suggesting that operators in North America were able to exercise more control over the timing of network changes in the wake of shelter-in-place announcements to limit their disruption to business users.

TAKEAWAYS

- While ISPs globally experienced an unusual increase in network outages post-February, their severity, duration and user impact varied by region, with North American operators less likely to disrupt services during traditional weekday business hours.
- Understanding local (and specific) operator practices is key to ensuring enterprises can plan and effectively communicate with stakeholders and vendors to ensure minimal business disruption even when unforeseen traffic shifts require ISPs to make more frequent network state changes.

CLOUD PROVIDERS

Like ISPs, cloud providers experienced unusual increases in network outages post-February; however, these were more notable given the relatively few disruptions typically observed.




FIGURE 24

Cloud provider network outages January-June 2020

Both pre and post pandemic, cloud provider networks were more stable than ISPs, contributing only 10% of all outages observed compared to 80% for ISPs—even though they are equally represented in the data set.




FIGURE 25

Cloud versus ISP network outages January-June 2020

Cloud provider networks serve as the backbone for the delivery of many ubiquitous digital services, making them effectively an extension of the Internet. In particular, Microsoft and Google have highly distributed network edges, and advertise routes to services globally regardless of where the service is hosted. Which effectively onboards users into their network closer to their location. While some services are served from their edge, there is a considerable amount of intra-cloud traffic that would have likely increased around the same time that many ISPs and Internet exchange providers began experiencing traffic surges. In particular, work from home activity increased demand for cloud services on multiple fronts. Many of the most popular video conferencing applications are hosted within public cloud providers, all of which saw a surge in demand beginning in March. Around the same time, some enterprises took advantage of cloud infrastructure to quickly scale up service capacity to maintain business continuity. For example, deploying VPN concentrators on cloud infrastructure to handle the demands of a workforce now entirely outside of the traditional enterprise perimeter.

Increased service demand did appear to impact cloud providers, but nowhere near ISP levels, which experienced network disruptions at 10x that of cloud providers. All things being equal from the standpoint of infrastructure coverage (quantity of network interfaces monitored), why would cloud providers, operating massive, complex global networks, demonstrate greater resilience than ISPs?

There are several factors that likely account for the extreme differences between these network operators. Cloud provider networks were built out more recently than most ISP networks, so they likely do not have the technical debt of longstanding operators. Many cloud providers also operate bespoke, software-defined networks purpose fit for their needs, allowing them to innovate and address issues faster, with little dependence on an external network vendor.

In contrast, ISP networks principally use commercially available infrastructure running vertically integrated monolithic operating systems, with release cycles and bug fixes under the control of a third party. In addition to service complexity that can build over time, infrastructure complexity using combinations of fixed, mobile and hybrid networks, results in an enormous level of technical and operational weight, ultimately hindering the ability to address issues and conditions quickly.

Cloud provider networks are advantaged not only from a resiliency standpoint but also in terms of network performance, showing few signs of degradation throughout the period.

CLOUD PROVIDERS MAINTAIN NETWORK PERFORMANCE

Cloud provider networks performed consistently throughout the first half of 2020, with virtually no elevation in packet loss or jitter between regions within North America, Europe, and Asia-Pacific for any of the three major cloud providers.

Inter-region performance for AWS, Azure, and GCP within the same geographic region (e.g. AWS us-west-1 to us-east-1) showed less variation than that observed from a network outage standpoint, as each region was similarly high-performing.

		NA						EMEA						APAC					
		Jan	Feb	Mar	Apr	May	Jun	Jan	Feb	Mar	Apr	May	Jun	Jan	Feb	Mar	Apr	May	Jun
aws	Packet Loss (avg)	0%	0%	0%	0%	0%	0%	0%	0%	0%	0%	0%	0%	0%	0%	0%	0%	0%	0%
	Packet Loss (95P)	0%	0%	0%	0%	0%	0%	0%	0%	0%	0%	0%	0%	0%	0%	0%	0%	0%	0%
	Jitter (avg)	0.0ms	0.0ms	0.0ms	0.1ms	0.0ms	0.1ms	0.2ms	0.2ms	0.2ms	0.2ms	0.2ms	0.2ms	0.0ms	0.0ms	0.0ms	0.0ms	0.0ms	0.0ms
	Jitter (95P)	0.2ms	0.2ms	0.3ms	0.3ms	0.2ms	0.4ms	0.6ms	0.6ms	0.6ms	0.6ms	0.6ms	0.6ms	0.2ms	0.2ms	0.2ms	0.2ms	0.2ms	0.2ms
Microsoft Azure	Packet Loss (avg)	0%	0%	0%	0%	0%	0%	0%	0%	0%	0%	0%	0%	0%	0%	0%	0%	0%	0%
	Packet Loss (95P)	0%	0%	0%	0%	0%	0%	0%	0%	0%	0%	0%	0%	0%	0%	0%	0%	0%	0%
	Jitter (avg)	0.2ms	0.2ms	0.2ms	0.2ms	0.2ms	0.2ms	0.2ms	0.2ms	0.2ms	0.2ms	0.2ms	0.2ms	0.2ms	0.2ms	0.2ms	0.2ms	0.2ms	0.2ms
GCP	Packet Loss (avg)	1%	0%	0%	0%	0%	0%	0%	0%	0%	0%	0%	0%	0%	0%	0%	0%	0%	0%
	Packet Loss (95P)	0%	0%	0%	0%	0%	0%	0%	0%	0%	0%	0%	0%	0%	0%	0%	0%	0%	0%
	Jitter (avg)	0.1ms	0.1ms	0.1ms	0.1ms	0.1ms	0.1ms	0.1ms	0.1ms	0.1ms	0.1ms	0.1ms	0.1ms	0.1ms	0.1ms	0.1ms	0.1ms	0.1ms	0.1ms
	Jitter (95P)	0.5ms	0.5ms	0.5ms	0.5ms	0.5ms	0.5ms	0.5ms	0.4ms	0.4ms	0.5ms	0.5ms	0.5ms	0.5ms	0.5ms	0.5ms	0.4ms	0.5ms	0.5ms

FIGURE 26

Packet loss and jitter — average and 95th percentile for AWS, Azure and GCP

Latency variations were minimal, with average latency staying consistent or slightly improving in most regions throughout the first half of 2020. Most of these variations amounted to small changes in performance, often of less than a few milliseconds. Figure 27 below uses January as a baseline to look at latency variation over time for each provider within their networks in each geographic region. A negative percentage indicates an improvement in latency (lower than baseline) and a percentage increase represents a latency increase over January.




FIGURE 27

Network latency variation — average monthly percentage change over January for AWS, Azure, and GCP

NOT ALL CLOUD REGIONS ARE EQUALLY STABLE

Despite a smaller, more homogeneous set of operators, cloud provider network resiliency varied considerably depending on region, with more outages in EMEA and Asia-Pacific than North America. The distribution of these outages throughout the first half of 2020 fluctuated independently of one another and don't appear to be as heavily influenced by traffic and usage increases reported post lockdowns in most regions.




FIGURE 28

Cloud provider network outages varied by region

Despite fewer outages overall, the duration of outages compared to ISPs was either on-par or longer, indicating either maintenance activity or less agile recovery from network changes that unintentionally impacted availability.




FIGURE 29

Duration — CSP network outage duration by region

Given their longer duration relative to other providers, outages taking place in cloud networks could be more disruptive to users depending on when they take place.

OUTAGE IMPACT ON USERS

Understanding the true impact of an outage on users isn't necessarily determined by if an outage occurred, but by when an outage occurred. The context of an outage, such as one that occurred on a weekday or on a weekend, or during business hours or the middle of the night, can mean the difference between an outage that makes headlines and one that goes unnoticed.

As noted earlier in this report, many ISP outages—particularly in North America—take place outside of business hours. In contrast (and despite their lower numbers), cloud provider network outages tend to take place more frequently during business hours compared to other periods.

Across all regions, very few outages occur on weekends. While this pattern is evident with ISPs as well, it is proportionally more extreme amongst cloud providers.




FIGURE 30

Regional outages by time of day, weekends


Traffic levels within ISP networks are generally lower on weekends. With cloud providers, this pattern may be even more pronounced, both from a services standpoint, as well as network backbone usage. Given that cloud networks are ultimately in service of core, revenue-generating offerings, such as compute, platform, or application suites, it may be that fewer reactive measures are required to ensure the reachability of those services during the slower weekend period. This pattern highlights the on-demand nature of cloud provider environments in which outage-causing activities like provisioning or optimizing service delivery are more likely to be made when users are most active.

In North America, this appears to be the case, as 98% of network outages occurred during weekdays, with no related lockdown pattern reflected. Outages were just as likely to occur during weekdays as their numbers increased (e.g. April and June), as they were during periods of less outage activity.



North American network outages, weekday versus weekend

The ratio of weekday to weekend outages in EMEA was similar to North America, with 85% of outages taking place during weekdays, and only 15% occurring on the weekends.



EMEA network outages, weekday versus weekend




FIGURE 33

APAC Cloud provider network outages Weekday compared to Weekend


This pattern was also replicated in Asia-Pacific, with most incidents occurring during weekdays.

The pattern in relation to the time of day outages occurred did not alter significantly throughout the months of January through July, which implies that normal operations were in effect, even during the post-lockdown period. This pattern differs from ISPs, underscoring the architectural and operational differences between these providers.

Network outages were just as likely to take place during traditional business hours (9AM-6PM, weekdays) within each region, than at other times.

FIGURE 34

Cloud network outages occur indiscriminately throughout weekdays



In North America, 44% of outages were likely to impact business users, as they fell within a 9AM-6PM ET time window.




FIGURE 35

Business versus non-business impacting network outages in North America

In EMEA, however, only 35% of all outages took place during business hours (9AM-6PM GMT+2) on weekdays versus other weekday times and weekends.




FIGURE 36

Business versus non-business impacting network outages in EMEA

Cloud network outages in Asia-Pacific were the least disruptive from the standpoint of local business hours (9AM-6PM JST), as just 24% occurred during that period.

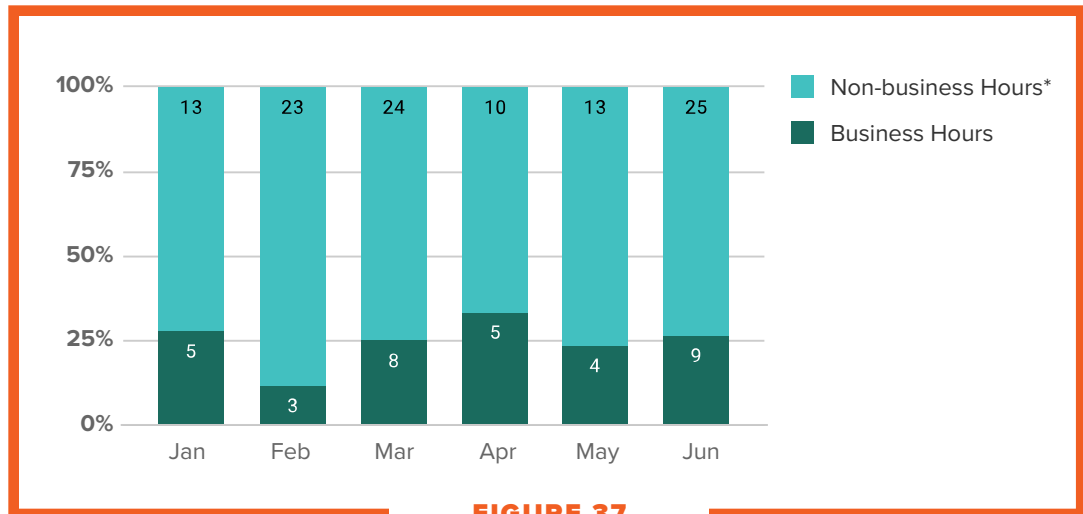


FIGURE 37

Business versus non-business impacting network outages in APAC

Despite looking at these outages from a regional perspective, it is important to remember that some cloud providers may take a centralized approach in terms of network control, reflected by timing that is dictated by a North American time zone. For example, an outage that impacts APAC during non-business hours is likely to have been initiated by a change during business hours in North America.

Cloud providers, particularly the ones with the most market share, have sophisticated network operations teams. The fact that a large percentage of outages take place during business hours may indicate a dynamic provisioning to meet demand and/or a Continuous Integration Continuous Deployment (CI/CD) approach, where incremental changes are continuously made to scale and optimize their networks to meet demand—demand that may largely be driven by a provider’s biggest market, which for most, is North America.

TAKEAWAYS

- Cloud provider networks demonstrated greater overall stability both pre and post-pandemic compared to ISP networks; however, they appear to operate under an agile, on-demand approach to maintenance and provisioning that may result in outage activity potentially impacting users during traditional business hours, particularly in North America.
- Despite less discriminate occurrence, cloud network outages in North America are far fewer than ISPs in that region, as well as cloud networks in other regions.

CDN AND DNS PROVIDERS

CDN and DNS services are critical to the optimal delivery of nearly every modern application, whether that application is business or consumer-oriented. The DNS enables users to find services over the Internet, and managed DNS providers serve two critical functions—hosting authoritative records (typically, from highly distributed infrastructure), as well as directing users to the optimal service site based on their location and the performance of the application the users are trying to reach. CDNs are also essential to the functioning of the modern Internet, as they enable media content and bandwidth-intensive applications to be distributed to users quickly while also reducing the load on Internet backbone infrastructures.

Both of these providers are, by virtue of the services they offer, highly distributed and scalable, which may explain why they experience few disruptions across their edge network infrastructure.

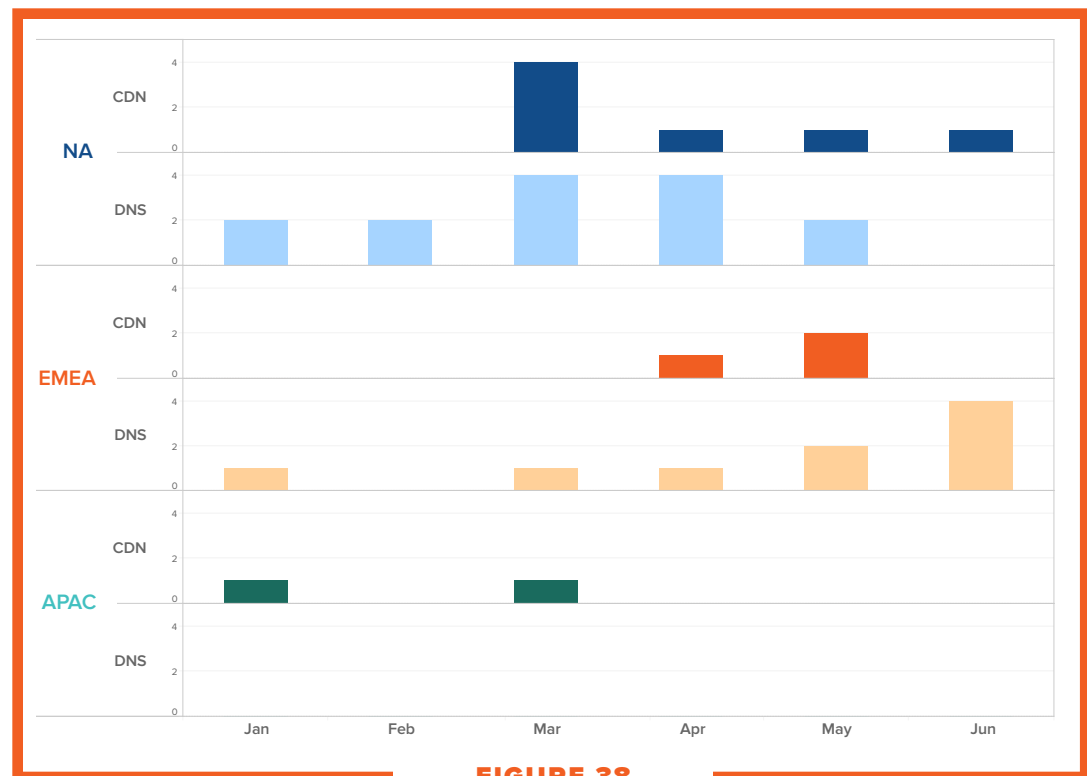


FIGURE 38

CDN and Managed DNS provider outages were uncommon across every region

Between January and July, just a few outage incidents were seen across each region. This stability was reflected not just in low numbers of outages, but also the patterns observed when outages occurred. In almost all cases, outage patterns were uniform and indicative of either maintenance or automation that has gone awry as opposed to infrastructure failure or network congestion. The most number of outages observed within CDN provider networks in North America was four. While all four occurred in March, they did not appear to be related to network congestion or infrastructure failure, as they were clustered close together within a single provider's infrastructure, and their size and repetition suggested automation cycling or some maintenance event.

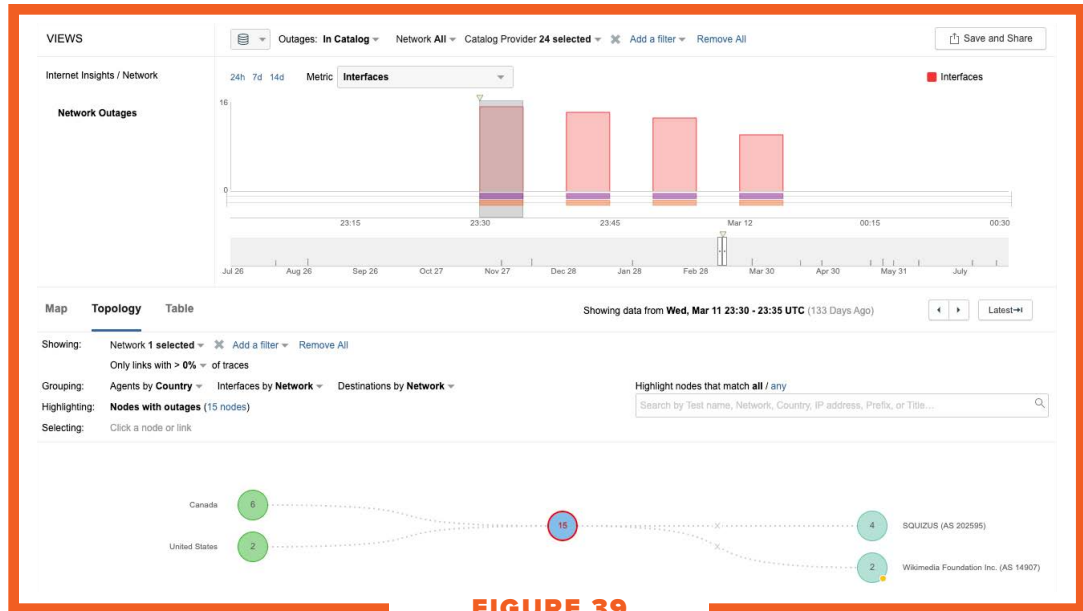


FIGURE 39
Four staggered, relatively small disruptions in one CDN's North American network

PUBLIC DNS RESOLVERS

In addition to CDN and managed DNS providers, public DNS resolvers and their performance play an important role in user experience, as several of them have very large numbers of users around the globe. Subsequently, they provide a good indicator of the quality of user experience.

As the Internet developed, two distinct environments emerged. The enterprise network, which is commonly encountered at the physical workplace, and the consumer network, which is typically deployed across residential infrastructure. Each has different characteristics and usage patterns, particularly from a time and day of week standpoint, as well as the types of services accessed. Many enterprises run their own DNS resolvers to service internal users. Others may leverage public DNS resolvers—particularly if local IT resources are less available.

Several public DNS resolvers, run by Google, Cloudflare, and Cisco, are used by a mix of business and consumer users around the globe. Typically, access methods and location determine how and what resolvers are used; for example, a user physically located in an office would be subjected to an organization's policies, whereas a remote user accessing corporate resources would most likely leverage a VPN and will potentially have different policies and configurations that subsequently determine how they resolve and connect to applications

When looking at the performance of these key services in February and March, we found that resolution times for DNS resolvers across all regions were consistently within acceptable usage tolerances and showed only occasional deviations, which could be attributed to latency on the path to the resolver, and not a result of DNS resolver performance. However, interesting performance patterns emerged for these resolvers depending on the provider and region.

Typically, DNS resolution times for services influenced by business policies tend to be higher (meaning slower to return a response) during weekdays (possibly indicative of greater usage of the service during that period), while services more heavily used by consumers reflect higher resolution times during traditional, non-business hours (including the weekend). As noted above, however, higher resolution time

is relative to either weekdays or weekends and does not indicate that any of the resolvers performed outside acceptable levels during the observed period. For DNS resolution to not significantly impact a users experience, you would typically expect to see the complete resolution process to be completed between 20ms to 120ms. The average resolution time observed across all providers, globally, during the period was 54ms, well under the optimal threshold. All resolver response times were measured based on a query for the A record of example.com. Given the average resolution times were within the optimal threshold, 95th percentile was used to highlight the different performance patterns.

DNS PERFORMANCE PATTERNS INFLUENCED BY USERS

In comparing February and March resolution times for OpenDNS for users in Asia-Pacific, we can see performance patterns change between these months. In February, during weekdays, response time is higher, particularly during the day. During weekends, however, response time is significantly lower. The following month, in March, weekday response time is closer to that of February weekends, and consistently lower than those typically seen for weekdays in February, which may indicate that either fewer users in that region were actively working from home or else enterprises using OpenDNS as a resolver were perhaps not fully enforcing policies around its usage at that time.




FIGURE 40
OpenDNS APAC response time patterns February-March

In contrast, OpenDNS' performance for EMEA users shows the opposite pattern, with resolution time lowest on weekdays in February. However, as in Asia-Pacific, March weekday performance was similar to February weekends, perhaps indicating that workers at home were defaulting to DNS services used at home pre-lockdown, or else were using VPN software for work-related activities. Most VPN software allows for client-side split tunneling, where enterprise traffic uses a VPN to connect to the enterprise network, and the rest of the traffic goes directly to the Internet. DNS is typically also split based on domains. When workers transitioned from the office setup to the home environment, more non-sensitive traffic would likely, initially (as enterprises worked to scale VPN capacity), have gone directly over the Internet and used a public resolver, accounting for usage higher than typically seen during weekdays prior to the pandemic.




FIGURE 41

OpenDNS, EMEA response time patterns February-March

Cloudflare’s DNS resolver performance in Asia-Pacific has a similar pattern to that of OpenDNS EMEA, with February weekend response times higher versus weekday, which may point to higher consumer-type usage. In March, weekday performance mirrored that of weekend performance in February, yet again suggesting that as the workforce migrated to the home environment, they may have initially consumed core services, such as DNS, in a similar way to that of their weekend (presumably, personal) usage.





FIGURE 42


Cloudflare, APAC response time patterns February-March

This consumer-type pattern was similar for Cloudflare users in North America, where pre-lockdown weekend response times were slower than those on weekdays, and March weekday performance was, again, similar to February weekend performance.



Cloudflare, North America response time patterns February-March

Google DNS resolver performance seemed to be less influenced by lockdown conditions and increased work-from-home activity. In particular in EMEA, where its resolver service was the only one measured showing weekday performance consistent between February and March.



Google, EMEA response time patterns February-March

TAKEAWAYS

- Overall, CDN and Managed DNS providers experienced few network disruptions in the first half of 2020, despite traffic and usage levels reportedly increasing beginning in March.
- The performance of public DNS resolvers also held up well, remaining high-performing throughout the observed period. However, nearly all experienced a shift in usage patterns, where weekday resolution times in March closely matched that of weekend resolution times in February. This may be due to workers at home in March defaulting to the DNS service typically used in their home environment prior to lockdown, or else were using VPN split tunneling.



CONCLUSION AND RECOMMENDATIONS

Despite an increase in network disruptions post-pandemic, the state of the Internet is healthy. The networks of services critical for modern application delivery, such as CDN providers, continued to be highly available, mitigating the load on Internet backbone infrastructure. ISPs implemented network changes to meet service needs, while in many instances minimizing the disruptive impact of these changes on businesses. Overall, Internet-related infrastructures have held up well, suggesting overall healthy capacity, scalability, and operator agility needed to adjust to unforeseen demands.

FINDING	RECOMMENDATION
Cloud provider networks are more stable than ISP networks, experiencing fewer network disruptions; however, when disruptions do occur, they are likely to impact users.	Consider your risk tolerance when choosing an application delivery architecture, taking into account local ISP performance and cloud network preferences and service tiers. For example, in a volatile ISP market, cloud networks could be relied on to mitigate the risk of business disruption due to outage incidents; however, this will vary by cloud provider, as some providers favor their backbones by default, while others employ hot-potato routing, except where a premium routing service is used.
Macro traffic shifts require compensatory actions by network operators that, depending on factors such as operator sophistication and technical debt, may lead to an increase in network disruptions.	Based on your service portfolio, plan for different traffic loads and/or distribution around workforce practices and significant events. Communicate these requirements when selecting an operator, and consider basing selection criteria on service assurance as opposed to simply connectivity.
Not all outages are equal.	Factor in outage characteristics and potential user impact when evaluating providers based on their resilience.
Provider operational practices vary by region—likely tied to local workplace norms—which may impact businesses differently. For example, ISP operators in North America tend to make network changes outside of business hours, but that practice is not consistent across all regions and providers.	Understand how operators work locally, including when changes are more likely to be made. Identifying operational practices and working closely with providers may help avoid business interruptions.
Though the impact was not systemic, some providers and locations experienced more network degradation than others post-pandemic.	Leverage active visibility to identify performance bottlenecks or suboptimal routing impacting your customers, employees, or applications, and use this data to effectively communicate and resolve issues with external network operators.



ThousandEyes 

201 Mission Street, Suite 1700
San Francisco, CA 94105
(415) 231-5674

www.thousandeyes.com

About ThousandEyes

ThousandEyes delivers Network Intelligence—immediate visibility into experience for every user and application delivered over any network, so companies can deliver superior digital experiences, modernize their enterprise WAN, and successfully migrate to the cloud.

© 1992–2020 Cisco Systems, Inc. All rights reserved. All product and company names are trademarks™ or registered® trademarks of their respective holders. Use of them does not imply any affiliation with or endorsement by them.