



Large Scale Outage Visibility on the Control Plane

Leonard Becker
TU Berlin

Oliver Hohlfeld
Brandenburg University of
Technology

Georgios Smaragdakis
TU Delft

ABSTRACT

With the increasing cloud usage for access to fast and well-connected computational power, cloud outages have also become a growing risk for businesses and individuals alike. We derive a method to analyze publicly available BGP data to measure the visibility of cloud providers' outages on the Internet control plane. We then utilize this method to analyze an outage of Cloudflare, a large DNS and content provider. Cloudflare's outage study shows that visible traces can be found in BGP, enabling data-driven outage studies.

CCS CONCEPTS

• **Networks** → **Network measurement.**

ACM Reference Format:

Leonard Becker, Oliver Hohlfeld, and Georgios Smaragdakis. 2021. Large Scale Outage Visibility on the Control Plane. In *CoNEXT Student Workshop 2021 (CoNEXT-SW '21)*, December 7, 2021, Virtual Event, Germany. ACM, New York, NY, USA, 2 pages. <https://doi.org/10.1145/3488658.3493788>

1 INTRODUCTION

The Internet was built as a robust *infrastructure* that can tolerate infrastructure failures. Yet, the concentration of Internet *services* to few but large clusters (e.g., data centers) can challenge its resilience to failures. In this regard, attacks and infrastructure failures of such clusters that concentrate many services can lead to noticeable outages. While outages of various forms have been studied before [1, 3, 5], studying the *visibility* of outages in service provider networks itself has received only little attention [9]; an aspect we address.

In this paper, we analyze the visibility of such network-specific large-scale outages on the Internet's control plane. That is, we analyze an outage in the Cloudflare network on July 17, 2020 [2] from the perspective of RIPE RIS BGP collectors to study how and where such an outage is visible, thereby contributing to the limited work in this space. Our work further shows a BGP-based approach to study outages. The goal of our work is to make a first step towards data-driven outage monitoring systems for real-time outage detection.

2 THE CLOUDFLARE OUTAGE IN 2020

Cloudflare offers content delivery (CDN), Internet security (e.g., DDoS mitigation), and domain name (DNS) services. On July 17, 2020 a configuration error in their backbone network resulted in a major outage that made hosted services unreachable [2, 7]. This

outage resulted in a 50% traffic drop for a duration of 27 minutes. The goal of our work is to use this particular case to study the visibility of such kinds of outages on the Internet's control plane with the goal of enabling data-driven outage detecting systems.

We next summarize the incident timeline according to Cloudflare's reports [2, 7]. At **20:25 UTC** the link between the datacenters in Newark (EWR) and Chicago (ORD) started having issues due to unknown reasons. This led to congestion of the link between Ashburn (IAD) and Atlanta (ATL). To resolve the congestion, the configuration of ATL was changed at **21:12 UTC**. However, instead of removing traffic from ATL, the faulty configuration lead to all routes being leaked into the backbone, which made ATL attract more traffic and congest even faster. At **21:39 UTC** ATL was removed from the backbone completely which immediately fixed the issue. At **21:47 UTC** all customer facing services were restored as the edge network operated normally again. Only some internal services for logs and metrics were still impacted by congestion, but these issues were resolved at **22:10 UTC**.

3 METHODOLOGY

The goal of our study is to analyze the *visibility* of a major infrastructure outage on the Internet's control plane (i.e., BGP). Our analysis is thus based on archived BGP data publicly available from RIPE RIS route collectors that provide an extensive amount of BGP data for many years. *Collectors* are special routers of which the only purpose is to receive BGP update messages. They are located in several countries, and are most commonly peering with many service providers directly inside large Internet Exchange Points.

To keep the noise of regular operation low, we focus on three RIPE RIS collectors in proximity to the location of the outage: New York (**rrc11**), Palo Alto (**rrc14**) and Miami (**rrc16**). We focus our analysis on two main attributes available in BGP update messages—an approach to enable the study of further outages. Namely, the number of IP prefix updates and the BGP communities contained in these messages. During normal operation, a certain amount of prefix announcements and withdrawals are to be expected. Outages on the other hand are usually rooted in hardware failures or misconfigurations, both of which can be resolved by rerouting and steering traffic. Due to the time sensitivity of critical outages it would come to no surprise to see a sudden surge in prefix announcements or withdrawals being published by a network experiencing such an outage. Secondly, community tags can be used to direct and steer traffic by passing certain messages to routers. As such, sudden changes in the number of announcements of certain communities can be used as an quantitative indicator for active traffic steering, without knowing the semantics of a certain community value. We analyze Cloudflare's outage by visualizing the changes of these attributes over time in suitable plots. At the same time, we attempt to correlate sudden changes in the plots with important timestamps extracted from outage reports published by Cloudflare.



This work is licensed under a Creative Commons Attribution International 4.0 License.

CoNEXT-SW '21, December 7, 2021, Virtual Event, Germany

© 2021 Copyright held by the owner/author(s).

ACM ISBN 978-1-4503-9133-7/21/12.

<https://doi.org/10.1145/3488658.3493788>




Figure 1: BGP updates by Cloudflare recorded on three collectors during the outage of 17.07.2020.

4 VISIBILITY OF THE OUTAGE IN BGP

Increase in BGP update messages. We show the number of BGP update messages received by the three collectors from AS 13335 (Cloudflare) in Figure 1. The solid blue line shows the number of BGP messages received during the outage. The vertical grey bars indicate different events during the outage described in Section 2. First, we notice an increase in BGP message at the time of the misconfiguration at 21:12 UTC. This event is clearly visible at all vantage points, though with differing intensity. To set a baseline, we show the number of BGP updates received at the same time one week earlier during normal operation (dashed orange line). It shows that the numbers outside the outage roughly align with those from a week ago and that the numbers during the outage are significantly higher. Thus, unusual volumes of BGP updates indicate network configuration changes (due to the outage).

Change semantics reflected in BGP communities. While deviations in the volume of announcements are one change indicator, they do not reveal the semantic of a change. Routing policies (e.g., biasing path, peer selection, or performing traffic engineering in general) are typically realized by tagging BGP announcements with communities. If we look at the communities distributed by Cloudflare during the outage, we can attempt to infer some information about the anatomy of the outage, e.g., affected colocations.

Figure 2 shows a histogram of community values that were included in BGP updates by Cloudflare between 20:00 UTC and 23:59 UTC. We consider only community values that were announced more than 50 times across all three collectors. Clearly, some community values are announced more often than others. The number of announcements visible at **rrc16** is about an order of magnitude lower than in other collectors. To validate our observations we contacted Cloudflare. They confirmed that they rely on BGP communities for internal traffic management through the backbone, as well as route visibility through internal route collection. They also shared the mapping of some communities. Indeed, some of the observed BGP communities, i.e., **13335:10027** ("ATL01"), **13335:19000** ("NORTH-AMERICA") and **13335:20520** ("SITE-LOCAL-ROUTE"), were announced due to the faulty configuration, as they do exactly match the communities being added to the leaked routes.




Figure 2: Histogram of community values announced by Cloudflare between 20:00–23:59 UTC on 17.07.2020.

5 DISCUSSION AND NEXT STEPS

Our analysis shows that the number of prefix updates is a simple but noisy activity change detector. A surge in control plane activity combined with BGP communities changes is an excellent indicator to infer changes in the network when the semantics of these communities are known. With these insights, we were able to infer that a route leak had occurred not only into the network's backbone but also to the Internet. This was not mentioned in the incident report but later confirmed by Cloudflare. Building a data-driven outage detector is a non-trivial task given the noisy nature of control-plane messages. Yet, we posit that combining different control plane signals can pave the way for building such a detector as the next step in our work. We also plan to investigate other outages, e.g., the Google outage [6] and Facebook outage [4], as well as possible disruptions of popular applications due to such outages [8].

ACKNOWLEDGMENTS

This work was partially funded by the ERC Starting Grant Resolu-tioNet (679158) and BMBF BIFOLD 01IS18025A and 01IS18037A.

REFERENCES

- [1] K. Benson, A. Dainotti, K. Claffy, and E. Aben. 2013. Gaining insight into AS-level outages through analysis of Internet Background Radiation. In *INFOCOM Workshops*.
- [2] Cloudflare. 2020. Cloudflare Network and Resolver Issues. <https://www.cloudflarestatus.com/incidents/b888fyhbygb8>.
- [3] A. Dainotti, C. Squarcella, E. Aben, K. C. Claffy, M. Chiesa, M. Russo, and A. Pescapé. 2011. Analysis of Country-Wide Internet Outages Caused by Censorship. In *IMC*.
- [4] Facebook. 2021. More details about the October 4 outage. <https://engineering.fb.com/2021/10/05/networking-traffic/outage-details/>.
- [5] V. Giotsas, C. Dietzel, G. Smaragdakis, A. Feldmann, A. Berger, and E. Aben. 2017. Detecting Peering Infrastructure Outages in the Wild. In *SIGCOMM*.
- [6] Google. 2019. Google Cloud Networking Incident #19009. <https://status.cloud.google.com/incident/cloud-networking/19009>.
- [7] J. Graham-Cumming. 2020. Cloudflare outage on July 17, 2020. <https://blog.cloudflare.com/cloudflare-outage-on-july-17-2020/>.
- [8] A. Kashaf, V. Sekar, and Y. Agarwal. 2020. Analyzing Third Party Service Dependencies in Modern Web Services: Have We Learned from the Mirai-Dyn Incident?. In *ACM IMC*.
- [9] T. Wan and P. C. van Oorschot. 2006. Analysis of BGP prefix origins during Google's May 2005 outage. In *IEEE IPDPS*.